AUTONOMOUS ULTRA-LOW POWER

ELF/VLF RECEIVER SYSTEMS

A DISSERTATION

SUBMITTED TO THE DEPARTMENT OF ELECTRICAL

ENGINEERING

AND THE COMMITTEE ON GRADUATE STUDIES

OF STANFORD UNIVERSITY

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

Max Edward Klein

November 2009

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

_____

(Umran Inan)   Principal Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

_____

(Ivan Linscott)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

_____

(Dwight Nishimura)

Approved for the University Committee on Graduate Studies.

.

This dissertation is dedicated to my parents Dana *and* John;

to my sister, Meredith;

and to my fiancé, Amanda.

Abstract

ELF/VLF radio signals, from approximately 300Hz to 30KHz, are commonly used for submarine communications, ionospheric remote sensing, geophysical prospecting, and studies of the near-Earth space environment. Naturally occurring ELF/VLF emissions caused by nearly every lightning strike can be detected for thousands of miles and provide an abundance of wave-particle interaction possibilities in the radiation belts. Any applied studies of these events, however, require electromagnetic waves in the ELF/VLF band to be amplified, captured, analyzed, and stored. Specific studies often require the placement of receivers in extremely remote locations such as Antarctica or the middle of an ocean, and are almost always located far away from power sources to decrease noise. These scientific requirements drive the creation of new ELF/VLF receiver systems. Two new receiver systems were designed for use in terrestrial, typically Antarctic, locations. Both new systems utilize 10-100 times less power than the lowest-power comparable existing ELF/VLF receivers and were designed for remote unmanned operation in extreme environments without external power.

The so called Penguin system, comprised of a hybrid microcontroller and FPGA architecture, removes the overhead of a general purpose CPU to provide the most streamlined processing for data acquisition possible while still maintaining a relatively traditional sampling architecture. The drastically reduced architecture of the Penguin system, compared to traditional receiver systems, on average consumes less

power than a typical LED indicator lamp while capturing high fidelity ELF/VLF magnetic field snapshots every fifteen minutes. The low power and thermal requirements of the Penguin system enables low-cost remote studies of medium to large timescale phenomena such as Chorus and Auroral Hiss without the need for local power. The system has been deployed and operated at the United States Amundsen Scott South Pole Station in Antarctica.

A second ELF/VLF receiver architecture, the VLF Advanced Technology platform, or "VAT", removes the typical CPU from the sampling and recording loop further reducing power requirements and physical system size while gaining the ability to record continuously. This radically new architecture enables future scientific studies of the fine structure in time and frequency of long-term events, such as the onset mechanisms of natural Chorus emissions, due to the extremely low power requirements. The system architecture is also greatly applicable to any continuous time recording system, including but not limited to acoustic and electromagnetic arrays for subsurface imaging systems, ionospheric remote sensing, and optical sensors.

Acknowledgements

The saying often goes something like "if you want to see the world, join the military." I guess whoever coined the phrase wasn't aware of the Stanford VLF group! Over the last five years I've been afforded the rare opportunity to visit and work in some of the most remote stretches of the planet, from Midway Atoll to midway down the South Pacific Ocean, culminating with the most extreme: the South Pole.

First and foremost, I would like to thank Prof. Umran Inan for bringing me into the VLF group. Even with a large group of students, he still made time to meet with everyone one-on-one every week. His tireless dedication to his students and VLF science is readily observable by all those around him and is truly appreciated. I would also like to thank Dr. Ivan Linscott for his advising, both technical and personal, throughout the Ph.D. program. Prof. Dwight Nishimura has my sincerest gratitude not only for being on my Ph.D. committee, but also for his mentoring of my teaching, especially with advice on dealing with tough classroom problems. The chair for my Defense, Prof. Per Enge, not only was kind enough to support my work and lead the discussion, but also taught one of the most interesting classes I took during my graduate education career. Behind the scenes, Dr. Maria Spasojevic and Dr. Dave Lauben provided invaluable insight and information on Antarctic ELF/VLF scientific studies, and general instrumentation and systems, respectively.

Over the years of outlandish ELF/VLF projects I have had the privilege to work with some of the finest engineers and students. At the forefront, Jeffrey Chang

provided consistent and excellent support, including system designs, schematics, board layouts, and debugging. Many of the projects would not have been possible without his dedication. I had the honor of learning an immense amount from Dr. Evans Paschal, our ultimate source for all things analog, and thank him for his extreme detail in explanations. Sarah Harriman, whose LNA is the front end of Penguin, provided significant help in moving Penguin from a bench-top recorder to a fully integrated receiver system. Charles Wang and Benjamin Mossawir were some of the best lab-mates one could ask for; I am grateful for their advice during the wee hours of the night. Bradford Bonney greatly aided in the construction of two buoys, along with Vikram Malhotra, Tim Su, Koji Gardiner, Daryl Ng, and Aaron Bronzan. Mark Gilbert, Yuhao Ding, and Wen Qi Chin provided excellent coding support during and after the 2007 REU program. Patrick Blaes did an amazing job for anyone, let alone an undergraduate, on Verilog coding of the modules used in the VAT CPLD, and continues to be an asset to our research group.

Last but certainly not least, I am eternally grateful to my family for their continued support throughout my academic career. For over 20 years my parents and sister have been there to push me forward, catch me when I fall, and support me in more ways than imaginable. My lovely fiancé Amanda has been amazing throughout the process, helping prod me to completion while continuously being so understanding of crazy schedules and esoteric terms. Thank you all.

<div align="right">Max E. Klein</div>

Stanford, CA
September 2009

Contents

List of Tables

List of Figures

Chapter 1

Introduction

1.1 ELF/VLF Science

Extremely Low Frequency (ELF) and Very Low Frequency (VLF) radio emissions, both naturally occurring as well as man made, provide a rich phenomenological spectrum for investigative remote sensing of geophysical phenomena and the near-Earth space environment. These low-frequency signals, ranging from 300 Hz to 30 kHz, can be ducted along the Earth's magnetic field lines, causing ground-based transmission to travel through regions of the ionosphere, interact with particles in that area, and return as modified signals to the ground [13]. While staying on Earth, the ionosphere and Earth itself form a waveguide, allowing ELF/VLF transmissions to be ducted along the surface. While traveling through this Earth-ionosphere waveguide, pernerbations in the ionosphere or Earth's surface can result in changes in signal strength and phase, allowing sensing of lightning and other events over extremely long distances with a signal ELF/VLF receiver [62].

From the most polar locations on the planet, explorations of the far reaches of the Earth's ionosphere and magnetosphere are possible due to the shape of the Earth's magnetic field lines [49]. Studies are abundant in Antarctica, leading to a

better understanding of the Aurora, charged particle populations in the ionosphere, and the connectivity between the sun and Earth [26]. Experimental data is now available demonstrating how manmade transmissions can be used to intentionally alter the trapped particle population of the inner radiation belts [30]. All of these studies, however, rely on a common need for accurate, high-resolution ELF/VLF data to examine the natural and manmade environment, detect new events and phenomena, and provide ground-truth for theoretical models. The receiver systems capable of continuously receiving this data across the globe and recording them with high fidelity are the topic of this research.

1.2   Sensing in Remote Locations

ELF/VLF reception and recording location choices are driven by the scientific requirements of each individual study. As described above, these requirements often push the signal locations to high L-shells, or high geomagnetic latitude. Studies performed by Stanford University often take place in the far reaches of Alaska and across the Antarctic Plateau. Collecting data in these remote locations requires reliable systems that not only meet minimum signal sensitivity and signal integrity requirements, but are also capable of autonomous operation without human intervention or maintenance for extended periods of time.

A paradox exists in collecting ELF/VLF data. The instruments and receivers necessary all require power of some level to operate. However, the sites with the most pristine ELF/VLF environments are far away from power lines and generators, which can generate significant levels of ELF/VLF interference [43, 9]. The AWESOME receivers deployed in Alaska minimize this problem through both careful selection of the receiver sites as well as separation of the antennas and LNA from the remainder of the system, and local power source, by a significant standoff (Section 2.3.1). Some

of the most pristine, and scientifically interesting, ELF/VLF receiver sites occur in Antarctica [22]. In these ultra-remote locations, neither power lines nor generators are present, requiring any instrumentation to contain its own power source. While solar may produce significant amounts of power during the summer months, the nearly 6 months of total darkness do not make this a viable option for continuous power. Minimizing the required system power is paramount for the success of such installations.

Many sites, such as those in Alaska, may be close to human populations who can occasionally swap and mail hard drives or other storage media for mass data transfer back to the science community. Other locations, especially polar sites, are without significant local bandwidth, internet or otherwise, and may not even have a human population present for the majority of the year. AGO field sites, described in Section 2.3.2, are totally unattended except for occasional visits dedicated to system maintenance, occurring once every few years. Clearly, these locations represent some of the most challenging environments to collect scientific data, and are largely unexplored due the the current unavailability of low cost, low mass, low power reception systems. New instruments are essential for massive data gathering operations in these regions.

1.3 Unique Receiver Requirements

Reception and storage of ELF/VLF signals fall into a unique regime in terms of signal bandwidth and required fidelity. A broadband ELF/VLF receiver is required to sample at 100 kHz with 16-bit resolution, as discussed in detail in Chapter 2. This sampling requirement fits uncomfortably between existing Supervisory Control And Data Acquisition, or SCADA, networks and traditional radio receivers. SCADA networks, typically used for industrial monitoring and control applications, have

significantly slower data rates, typically around several kilobits per second [38, 8]. Due to this significantly slower data rate, simple power saving measures, such as turning on and off the sensor and/or processor between samples, is more practical than with ELF/VLF reception. Conversely, modern software defined radios typically have a much higher bandwidth, and thus demand extremely fast analog to digital converters and high powered commodity processors [10].

The unique bandwidth of ELF/VLF reception systems leaves a gap in traditional power optimization mechanisms and architectures, especially when coupled with the need for low electromagnetic emissions within the 300 Hz to 30 kHz band. Baseline receiver systems, detailed in the next chapter, are able to provide good quality data, but are impractical for widespread deployment in extreme and remote environments due to power limitations.

Two new receiver systems, including a revolutionary new sampling architecture, are presented in the following chapters. These two systems each reduce the power requirements of existing ELF/VLF receiver systems by over one order of magnitude, while maintaing or exceeding parity in data quality. The first system, Penguin, utilizes a traditional receiver and sampling architecture with a highly customized processor, optimized for ELF/VLF sampling and recording tasks. Penguin is operational on the Antarctic Plateau near the South Pole as a totally self-contained unit weighing less than 50 pounds total, with a majority of the weight due to mechanical structures; the batteries themselves weigh under 20 pounds.

The second receiver system, the VLF Advanced Technology receiver, or VAT, removes the traditional general purpose CPU from the sampling loop, dramatically reducing power consumption while still enabling continuous, real-time sampling of multiple ELF/VLF broadband channels. VAT reduces the continuous average power consumption by fivefold from the Penguin system, creating the smallest and most power efficient ELF/VLF receiver system known at this time. This revolutionary

sampling system architecture provides better than order of magnitude power improvement, and is generally applicable to many mixed-signal sampling systems of similar bandwidth.

Both Penguin and VAT enable new scientific studies that were not previously possible with existing instrumentation. The dramatic reduction in power consumption and physical footprint allow these systems to be deployed in some of the most remote locations on the planet. Extensive lab testing and field demonstrations indicate that the systems as designed are capable of operating unattended in the coldest naturally occurring climates on Earth while maintaing proper operation. With these new capabilities, large distributed interferometric arrays are now possible on the Antarctic Plateau or anywhere else on Earth with low cost hardware and greatly reduced deployment complexity and expense.

Chapter 2

VLF Receiver Background

In order to have an in-depth discussion of VLF receivers at the system level, it is critical that a common vocabulary be established. The required background includes generic radio and electronic terms as well as a primer on traditional direct conversion VLF receiver topology. Once the terminology is established, a review and discussion of existing systems will provide the requisite background to identify prior shortcomings and underscore the contributions of the system described in this work.

2.1 Terminology

2.1.1 ELF/VLF

ELF/VLF literally stands for Extremely Low Frequency / Very Low Frequency, and refers to two specific bands in the RF spectrum. In the telecommunications community, ELF is defined as frequencies between 3 Hz and 30 Hz, while VLF is defined as 3 - 30 kHz, however in the scientific community ELF is the region between 300 Hz and 3 kHz, while VLF is 3 kHz to 30 kHz [60]. Since this paper is

directed at the scientific and academic community, it will use the later convention, and may also refer to an ELF/VLF receiver as simply a "VLF" receiver for brevity. In ITU / telecommunication terms, the discussed receivers are technically ULF/VLF receivers. All systems discussed in this paper thus cover the 300 Hz to 30 kHz band unless explicitly stated otherwise.

2.1.2   Broadband

A VLF receiver is said to be "broadband" if it records the entire ELF/VLF band, nominally 300Hz to 30KHz as defined by the scientific community. This contrasts to a narrowband receiver, such as components of the AGO system described in Section 2.3.2, which sample only a small frequency range which is totally contained within this band. As a practical matter, the broadband ELF/VLF receivers discussed in this paper sample at 100 ksps unless otherwise noted. This oversampling, in relation to the Nyquist rate, allows for finite filter roll-off and provides greater alias rejection [35]. Note that while the temporal, and thus frequency, resolution is specified by the "broadband" qualifier, the amplitude resolution and bandwidth are not. All systems discussed contain a 16-bit digitizer, unless otherwise stated. This does not, however, imply that collected data necessarily has 16 ENOB (see Section 2.1.4).

2.1.3   Zero IF / Direct Conversion

All discussed broadband ELF/VLF receiver systems are direct conversion, meaning that the raw RF signal is amplified, filtered, and then directly digitized and/or recorded [1]. In contrast to a heterodyne system, where the RF signal is mixed with a local oscillator to an Intermediate Frequency (IF), a direct conversion receiver does not require any RF oscillators to operate, resulting in significant power and cost savings [48]. "Zero IF" is generally synonymous with direct conversion, as it

alludes to an intermediate frequency of zero, or no frequency shifting of the original signal. Due to the relatively low carrier frequencies of VLF, modern analog-to-digital converters are more than capable of sampling sufficiently fast to capture the entire ELF/VLF band without downconversion.

2.1.4   ENOB

The Effective Number of Bits, or ENOB, is a measure of the effective resolution of an ADC, and is tested by injecting a single sine wave and examining the output. The ENOB is directly related to the decibel SINAD, or Signal to Noise And Distortion, of a sinusoidal test signal and is expressed in units of bits. Note that the system-wide ENOB is dependent upon the amplitude of the stimulus signal. Thus, a poor RF front-end connected to an extremely high performing ADC may result in a system-wide ENOB lower than that of a poorer performing ADC but with a better front-end. From the equation, it is clear that anything harming the SNDR or SINAD of the system will also adversely impact the ENOB as revealed in Equation 2.1.

$$ENOB = \frac{SINAD_{dB} - 1.76}{6.02} \tag{2.1}$$

2.1.5   Firmware and Software

In the discussions of digital system architectures, both firmware and software will be covered. To avoid confusion, these two terms must be explicitly defined for our local context. Firmware signifies the arrangement or interconnection of electrical logic. Both Field Programable Gate Arrays and Complex Programable Logic Devices contain firmware, which describes how their respective logic fabrics are configured. Software is defined to be codes running on a given hardware processing unit. Software encompasses the programming that is executed inside of a microcontroller,

on a personal computer, as well as that executed within a CPU that is defined in an FPGA. The implementation of the CPU itself on an FPGA, however, is considered to be firmware.

2.2   Receiver Topology

A typical broadband VLF receiver, shown below in Figure 2.1, consists of an antenna, Low Noise Amplifier (LNA), Anti-Aliasing Filter (AAF), a time reference, and a recording or communications device. In a digital VLF receiver, which will exclusively be discussed in this paper, the recording or communications devices consists of an analog-to-digital converter (ADC) plus a processing device and storage or communications channel. Historically, the LNA is placed as electrically close as possible to the antenna to reduce the amount of cable losses before the received signal is amplified. The remainder of the system, from the AAF through the storage or communication system, is placed further away, greatly reducing the received system noise from these components through the antenna.
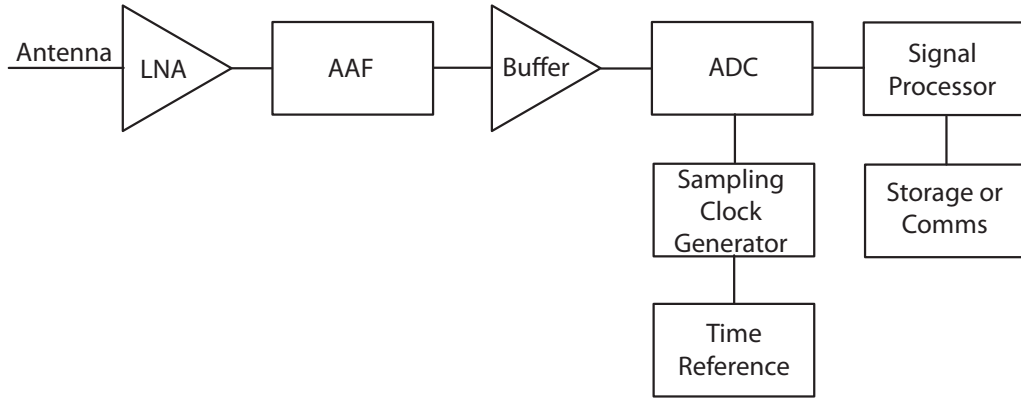

Figure 2.1: Block diagram of a typical zero-IF VLF receiver

Narrowband VLF receivers have several different architectures depending on the

desired frequency resolution and power constraints. One type of receiver, often re-
ferred to as "channelized," has a bank of bandpass filters just after the buffer and
before the ADC. Each bandpass filter only passes the frequency range desired for a
particular channel, with both the center frequency and bandwidth set by scientific
requirements. The output of each bandpass filter can then be fed to an integra-
tor to allow for extremely slow sampling, often with a sampling rate of around 1
Hz. This low sampling rate allows for ultra-low power computational systems and
minimal storage requirements. The analog bandpass filters also can consume signif-
icant amounts of power due to the amplification required to maintain the desired
signal strength through an array of many parallel bandpass filters. One example of
such a system is the Automated Geophysical Observatory, explored in more detail
in Section 2.3.2. Unfortunately the data collected from this type of narrowband
architecture typically has a frequency resolution of several kHz and temporal res-
olution of around 1 second. While adequate for many scientific studies of natural
phenomena, this low resolution data is unable to capture fine structure in time or
frequency, and thus is unable to record or reproduce the full detail of the received
events.

Another narrowband architecture with significantly greater resolution is targeted
at receiving manmade ELF/VLF transmissions. This type of narrowband receiver
locks onto a manmade transmitter, demodulates and removes the data from the
carrier wave, and then exclusively examines the carrier itself. Currently, the de-
modulation of the signal can be performed readily by commodity PC hardware,
as is the case with the AWESOME receiver discussed in Section 2.3.1 (note, how-
ever, that this receiver is also capable of full broadband ELF/VLF reception as
well). This receiver class is useful for scientific studies where the behavior of natural
phenomena is observed through amplitude and phase changes of VLF transmission
paths across the world [41]. For these studies to be accurate, especially if multiple

sites are used for interferometry such as with the Holographic Array for Ionospheric Lightning Research, the transmission sources must have a highly accurate and stable carrier wave, as is true with many naval transmitters, and the receivers must have a similarly accurate and stable time and frequency base, as is the case with GPS-synchronized receivers [32]. With these conditions met, the low data rate narrowband data can be integrated over minutes or hours to detect extremely small changes in the wave propagation path.

With any of the receiver topologies discussed so far, there is still a choice as to when a system should record and for how long. A continuously sampling full broadband system with all three possible magnetic axes being received would consume a considerable amount of power and recording space. For example, assuming 16-bits per channel per sample with a 100 kHz sampling rate, nearly 52 gigabytes of data would be collected each day from each receiver. With modern storage technologies, this vast quantity of data can be stored provided the storage medium is able to be exchanged every few weeks. For remote sites, like those in Antarctica, the extreme reaches of Alaska, or autonomous platforms in the ocean, this is not practical, and mechanisms of data reduction are required.

The easiest method to both reduce the quantity of data collected as well as possibly the energy consumed by the receiver is to simply cycle the recording on and off. Such a synoptic recording schedule is used in the AGO system for its broadband snapshots while the lower data rate narrowband data is continuously recorded. Another method is to only have the receivers operating when the researchers are expecting events to occur. Scheduling of the recording periods is done both with the South Pacific Buoys (Section 2.3.3) as well as the AWESOME systems. Lastly, the huge amount of broadband data can be compressed for instance by either decimating it to remove higher-frequency components that are not of interest during a particular experiment, or to save a processed snapshot of the data for evaluation.

The buoys both decimate the data to an effective sample rate of 10 kHz, due to the low-frequency nature of the experiments, and are able to store both raw data as well as a JPEG image of a spectrogram of any given section of the data, greatly reducing the data transfer required to analyze signals from the extremely remote platform.

These and other data reduction techniques are particularly prevalent on remote, power-constrained, autonomous ELF/VLF receiver platforms. As discussed throughout this paper, while storage capacity has been increasing at great speeds, maintaining the lowest power possible while operating in some of the harshest climates on Earth is still a significant challenge. We will briefly look at the current "gold standard" of VLF data quality, the AWESOME receiver, as well as two previously deployed autonomous receiver systems. With the benefits and limitations of these systems understood, we then proceed to describing two different receiver platforms, the subject of this dissertation, that greatly increase data collection capabilities in extreme environments.

2.3 Existing VLF Receivers

2.3.1 AWESOME

The AWESOME receiver, or the "Atmospheric Weather Educational System for Observation and Modeling of Electromagnetics," is a ground-based ELF/VLF receiver system. The AWESOME system is extremely sensitive, capable of receiving signals as low as several tens of femtoTesla per root Hertz, with true broadband coverage of 300 Hz to nearly 50 kHz [29]. The receiver typically consists of two large crossed loop magnetic field antennas with a co-located preamplifier or LNA. The antennas and LNA are impedance matched and sized to obtain the maximum sensitivity, limited primarily by atmospheric or local noise. The antennas and LNA are connected

by a significant length of Belden shielded multi-conductor cable to an indoor line receiver supporting data acquisition. This cable allows the antenna to be located a significant distance away from the rest of the system, and preferably away from power lines, generators, and any other sources of ELF/VLF interference. The line receiver filters the signals, digitizes them with a COTS sampling system into a PC, and runs custom software to control the sampling intervals, perform any processing such as narrowband extraction, and record the data to disk or to a remote site. A block diagram of the AWESOME receiver system is shown in Figure 2.2.
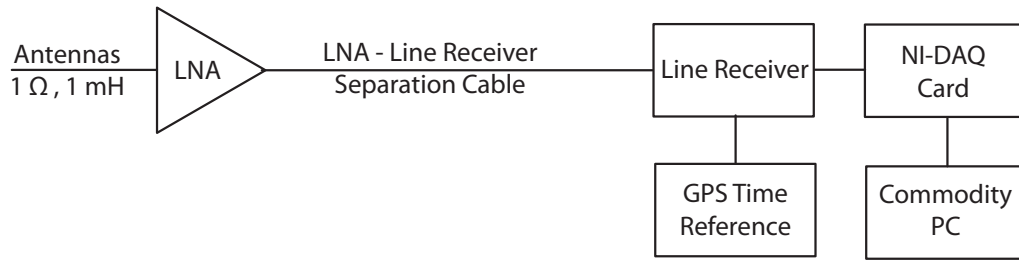
Figure 2.2: Block diagram of AWESOME receiver

The system is relatively inexpensive to produce, and benefits from a GPS-synchronized sampling clock, allowing for widespread interferometric use [52]. While the system produces superb data quality and is the baseline of quality for many new systems, it is limited in deployment locations. The system requires nominally 30 Watts not including the PC, or conservatively estimated at 60 Watts total when combined with a current laptop. Most deployments, however, use a regular desktop PC plus external hard drives for data transfer, equating to over 200 Watts of power. This high power requirement limits AWESOME receivers to locations with generator or

AC line power. Unfortunately, these locations also will typically have a higher background ELF/VLF noise level than those without local power. Extremely remote locations, including some islands and most of the Antarctic Plateau, are unsuitable for AWESOME deployments.

The AWESOME system contains a traditional direct-conversion broadband receiver topology as detailed in Section 2.2. For its particular implementation, the ADC is a National Instruments Data Acquisition, or NI-DAQ, Card contained in a PC, sampling at 100 kHz with 16-bit resolution, timed off of a GPS synchronized clocking circuit. Likewise, all signal processing and recording takes place in the same COTS computer. The AAF is an active circuit, allowing for extremely sharp filter roll-off and minimal signal loss, but at an obvious cost in required power. The LNA and antennas are matched to a nominal 1 ohm, 1 milliHenry impedance, which is a standard used in all other system discussed in this document.

2.3.2 Automated Geophysical Observatory (AGO)

The Automated Geophysical Observatory is a decades-long experiment to study phenomena at high geomagnetic latitude. Of particular interest is the current AGO VLF receiver system, deployed as a subsystem of the AGO electronics in the mid 1990s [54]. This VLF receiver provides synoptic "broadband" data with initially 2.5 kHz bandwidth, and more recently 10 kHz bandwidth, as well as several configurable channelized feeds. Both the broadband and narrowband data are filtered in analog, allowing for minimal computational requirements. The narrowband channels are of varying bandwidths between 500 Hz and 16 kHz, which are integrated and sampled at up to 2 Hz. A representative AGO system block diagram is shown in Figure 2.3. The extremely low sampling rate reduces the storage requirements of the system, which often must operate for a year or more unattended, and only have low data rate Iridium connections for communications. Aliasing on these channels does not

occur since the 500 - 16,000 Hz bandwidth is integrated prior to sampling, however the temporal resolution of events is limited to 0.5 seconds and the frequency resolution is equivalent to a particular channel bandwidth. High temporal and frequency resolution data are available from the AGO system during the synoptic recording periods, but are limited to a maximum frequency of 10 kHz and only currently record for one minute out of every fifteen minutes.

The VLF portions of the AGO system consumes approximately 30 Watts, allowing for the system to operate on the Antarctic Plateau with onboard power, but requiring extremely costly maintenance trips each year to provide more fuel for the power system. The system is also housed in a temperature-controlled hut, and is not currently rated for outdoor use without additional protection. The AGO system has provided a plethora of scientifically-useful data from Antarctica for a number of years. The great expense of installing and maintaining each AGO site, coupled with the lack of full broadband VLF data, however, leaves significant room for improvement. New systems that can provide high temporal and frequency resolution while being as robust, low-power, and physically small as possible would provide both better data quality as well as lower operational costs in such an extreme location.

2.3.3 South Pacific Buoys

The South Pacific Buoys were a research platform designed by Stanford University and deployed in early March 2007. The buoys were related to the High Frequency Active Auroral Research Program, or HAARP, a large HF heater located outside of Chistochina, AK. The buoys themselves were designed to provide reception of HAARP-produced ELF/VLF signals at its geomagnetic conjugate point. The particular research focused on the one-hop phenomena and natural methods of VLF amplification and dispersion [23]. These signals were expected to occur well within
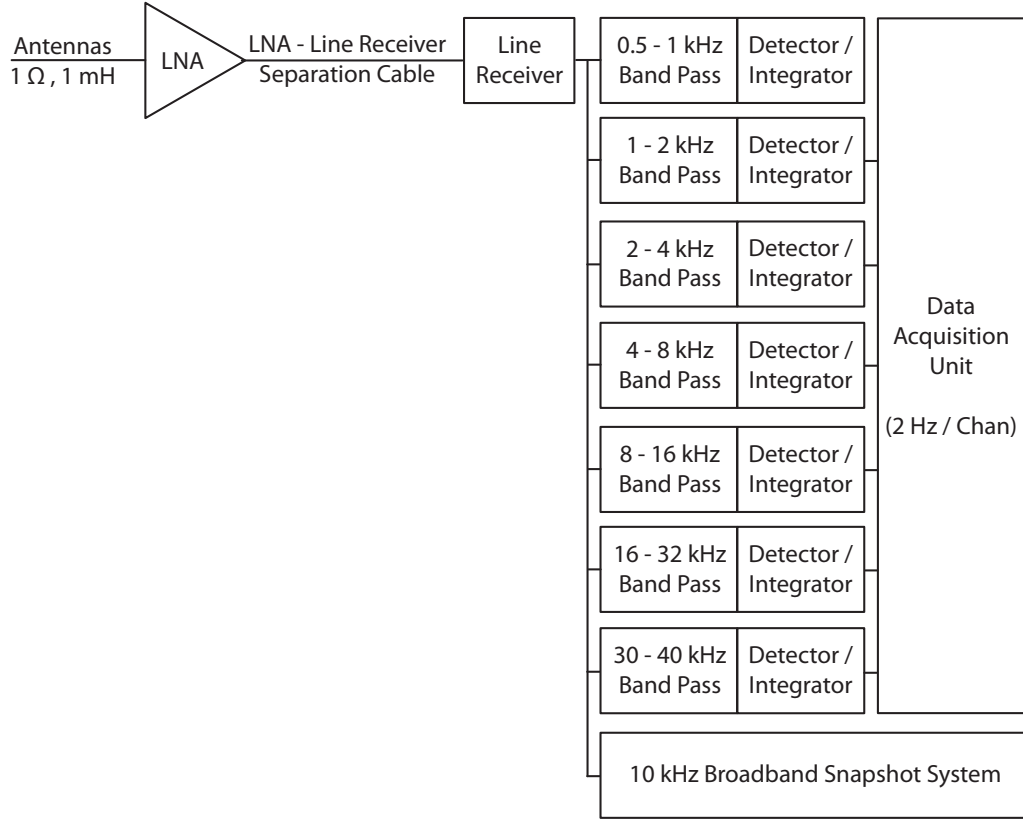
Figure 2.3: Block diagram of a sample AGO receiver

the ELF band, however the system was designed to include full ELF/VLF broadband and narrowband capabilities for added research while at the conjugate point, as well as for deployment of identical electronics systems at other locations, such as Tern Island which shipped prior to buoy deployment.

The buoys were limited in antenna size by the radome structure, however still contained two six foot square vertical antennas plus one 5.5 foot circular horizontal antenna, providing measurements of all three axes of magnetic field. The LNA was similar to that of the AWESOME receiver, and an active AAF was also used. Due to physical restrictions of fitting onto the buoy, the main system electronics were virtually co-located with the antenna, being directly below the center of the antenna by only a few feet. To minimize system noise, the main electronics were housed in

mumetal, a magnetic shielding nickel-iron alloy, and significant effort was placed on power system noise shielding and filtering. The custom electronics digitized the VLF radio signals at 16-bit resolution with 100 kHz sample rate, also synchronized with Universal Time by GPS allowing for interferometric studies similar to AWESOME. The processing system consisted of a Texas Instruments TMS320 digital signal processor plus a stack of up to four CompactFlash cards. The processor was capable of continuous time recording at full broadband sampling, as well as near-real-time narrowband demodulation. A block diagram of the buoy electronics is shown in Figure 2.4.
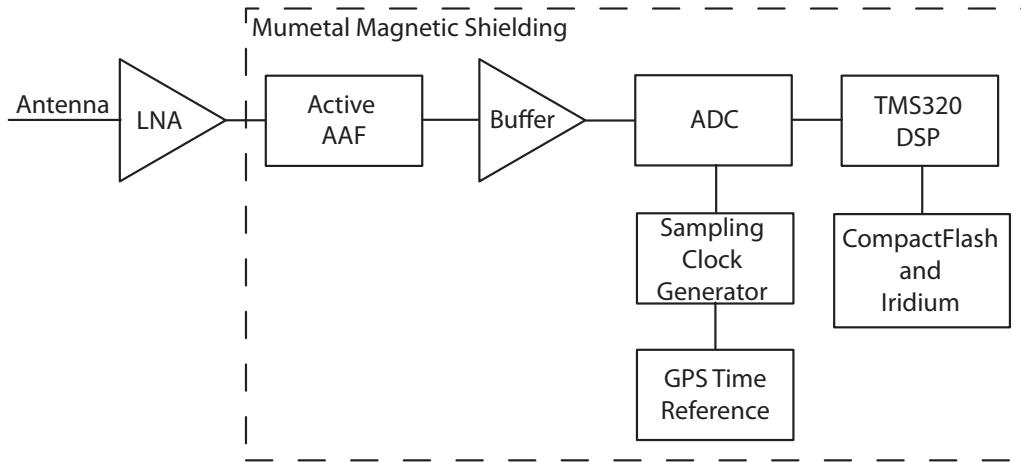
Figure 2.4: Block diagram of the South Pacific Buoy

Although designed to exclusively retrieve its schedule and transmit collected data via Iridium, the buoy electronics can be operated without satellite communications for pre-scheduled recordings and processing, maintaining the data exclusively on the CompactFlash cards. In this mode, the buoys consume approximately 15 Watts for a single system. Although deployed as double-redundant on each buoy, the system

was designed to run only a single system at a time, failing over in the event of a hardware failure. The buoys were able to recharge themselves via solar panels, and were rated for a one-year on-station operational life. These systems were able to function autonomously and in an extreme environment, however without sunlight, as is the case in Antarctica for half of the year, the systems would be unable to recharge. The buoys each shipped with 2,000 pounds of sealed lead acid batteries to ensure operation throughout the low, but still present, sunlight winter months. The large battery bank would not operate at temperatures significantly below freezing, and are extremely heavy making them impractical for Antarctic deployment.

2.4   Existing Receiver Limitations

An ELF/VLF receiver system can be characterized by many different performance metrics, including frequency response, minimum detectable signal, sampling rate, sampling resolution, time synchronization, size, weight, average power and peak power. With a firm grasp on signal integrity and sensitivity properties, as demonstrated on the AWESOME and buoy receiver systems, attention can turn to time and frequency resolution, size, weight, and power.

For remote observations, it is critical that the total system volume and mass, including the power source, is as small as possible to reduce deployment costs. Large systems, such as AGO, are only feasible with significant funding and even then are limited in the number of deployment sites. The buoy system is simply too large, when the required batteries for antarctic operation are added, for deployment to most remote sites. The AWESOME system has a reasonable mass, and is able to be hand-carried by two people on an aircraft if necessary, however a battery bank able to supply several hundred watts continuously for a year is extremely large.

Power requirements clearly can drive the total system mass required for remote

deployments. Systems with average power consumption as low as 10 Watts still requires 87,600 Watt hours per year, or a 7,300 Amp hour battery at 12 Volts. More than an order of magnitude drop in power requirements is necessary with current battery technology to maintain a man-portable power supply for a year of operation. Both novel systems presented in the remaining chapters of this paper exceed this order of magnitude requirement while maintaining comparable data quality.

Data quality itself is another parameter where room for optimization exists on embedded platforms. The AGO system trades reduced computation requirements for extremely poor time and frequency resolution of its data; fine structure features of less than one second are not observable, and frequency resolution is limited to hundreds of Hertz or worse. The exception to this is during the short synoptic broadband window, where a 10 kHz of high fidelity data is available. Collecting the full 30 kHz broadband ELF/VLF data continuously, as is done with AWESOME and optionally the buoys, captures all time and frequency information of the studied signals, leaving no information to imagination.
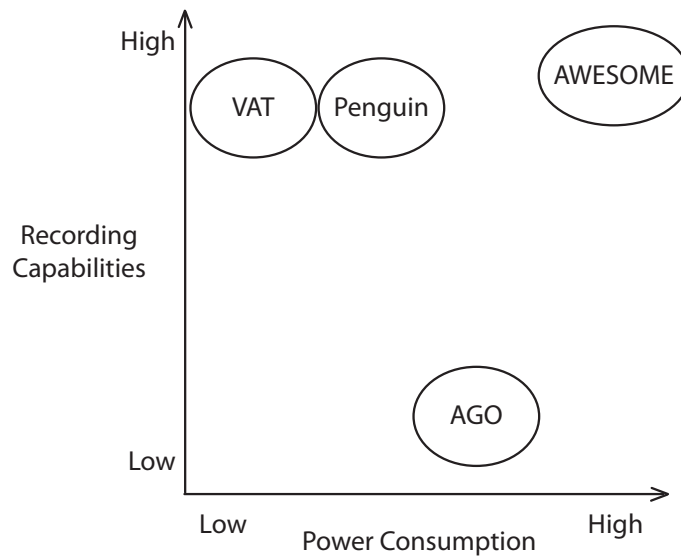
Figure 2.5: Comparison of VLF Receiver Systems

2.5   Contributions of New Systems

This work addresses a longstanding problem in the data acquisition world: high performance and long duration sampling in harsh environments with minimal power consumption. Existing systems discussed have advanced the state of the art to allow VLF recordings in extremely remote locations and harsh environments, but still require at least an order of magnitude more power than desired. To develop a system with better power performance by such a large margin, while still maintaining ruggedness, the problem was approached by an architectural solution.

The proposed architectures focus on the sampling control, sampling read back, data storage, and scheduling systems. These systems are complex, and each involved significant amounts of hardware logic. While the system behavior is deterministic, it is sufficiently complex that standard timing and power simulations may not prove the actual, real-world performance of the system as a whole, especially in the extreme environments targeted. Due to the complexity of the solutions presented, and especially the interdependence of the system components on various pieces of disparate ICs, the architectures have been implemented and tested in laboratory and operational conditions, proving, and sometimes exceeding, the predicted improvements from the design stage. The underlying new architecture design represents a significant leap in the state of the art for sampling systems in the target frequency regime, while the straightforward implementation proves not only the performance, but utility of new systems based on the described architecture.

The following chapters describe in detail both a highly refined traditional sampling architecture used to capture synoptic (or optionally continuous) but full-bandwidth broadband data as well as an entirely new sampling architecture for ultra-low power continuous sampling of ELF/VLF data. The Penguin system replaces a traditional microprocessor for the data acquisition system with a highly

customized processor within an FPGA, along with a microcontroller for ultra-low power sleep modes of the system. The VAT system, subsequently described, is a radically new architecture, removing the general purpose CPU from the sampling loop, replacing it with a CPLD containing a highly-tuned sampling and recording controller.

Both new systems provide greater than one order of magnitude power savings, even while operated in continuous sampling modes, extremely small physical footprints, and demonstrated reliable operation in extreme environments. These new capabilities permit experiments in extreme locations that were previously impossible or financially and logistically impractical. The novel ultra-low power architecture developed for VAT is also widely applicable to any continuous sampling system, including acoustic, sonar, and optical devices.

Chapter 3

Penguin Micro-VLF Receiver

3.1 Background

The Penguin Micro-VLF Receiver is an ELF/VLF receiver system designed for exploration of scientific phenomena which are most readily studied in the polar regions. These phenomena take place in the outer regions of the Earth's magnetosphere, and the resulting VLF waves are often ducted along the magnetic field line back to the Earth's surface. Since the location of the phenomena are in the outer regions of the magnetosphere, or at high "L-Shell", the waves are most easily received at high geomagnetic latitude, requiring the placement of such receiver systems in the extreme environment of the polar regions [49]. The specific phenomena to be studied, such as Auroral Hiss and Chorus, require broadband reception of VLF signals over long, multi-hour, periods. Due to the remoteness of the Antarctic Plateau where the system is deployed, it is imperative that the Penguin system is capable of autonomous operations for at least one year per deployment, while withstanding the extreme environment and lack of external power. Additionally, since deployment operations are extremely expensive with highly restricted logistics, the system must be as lightweight and low volume as possible.

Existing VLF receivers, as outlined in Section 2.3, consume tens of watts of power for broadband VLF reception and physically are rather massive, especially once a one year supply of batteries is added. These systems often include high power CPUs which are always running at full speed, even if idle, while the VLF data is recorded, such as with AWESOME. Some utilize a more power-efficient DSP, such as the buoy or AGO system, however even these more efficient DSPs are power hungry (several watts of power at idle), and during a majority of the sampling process the DSP is also idle. Having high clock rate digital devices running while idle during the time between sampling cycles is clearly not the most efficient use of power, although it is the typical implementation method of current sampling systems.

To combat these problems, the Penguin system utilizes a Field Programable Gate Array (FPGA), which is effectively a sea of programmable logic, to implement a customized CPU that minimally meets the sampling system requirement to conserve as much power as possible. By stripping out any extra peripherals and core CPU items which are not necessary, such as MMUs or advanced computational instructions, there is less logic to be clocked or to leak power and thus less power consumed overall. Also, by reducing the CPU clock speed to the minimum frequency for the device to keep up with the required sampling, the dynamic power is further reduced. These architectural changes remove the Penguin system from being simply the next incarnation of a standard architecture, implemented on the newest silicon, to a significantly more refined processing flow, providing performance enhancements over traditional architectures regardless of the underlying hardware. The overall Penguin architecture represents a highly-optimized digital sampling system for VLF data acquisition that exceeds the power performance of any other known comparable VLF system by more than an order of magnitude. The architectural design process hints at the anticipated performance improvements, while the implementation and deployment of this new architecture prove the end performance, operational success,

and utility of the new system design.

3.2 Derivation of Requirements

The first, and arguably most important, step of scientific instrument design is to define the requirements of the device as completely as possible. We will take a ground-up approach and start with the scientific constrains and requirements, followed by derived requirements from the scientific objectives, and lastly will consider logistical and operational concerns.

For the study of high L-Shell phenomena such as Auroral Hiss and natural Chorus emissions, we can directly determine the VLF bandwidth requirements and start to determine the minimal system dynamic range. Studies have indicated that Chorus is typically band limited to under 2.5 kHz [11]. Hiss, on the other hand, is present over a much larger band which includes the entirety of the 30 kHz VLF band [34, 39]. Thus, due to the larger frequency constraints of Hiss, Penguin must sample the entire 30 kHz ELF/VLF band. For Nyquist sampling, the sample rate or F_s must be at least 60 kHz [35]. However, this assumes that the anti-aliasing filter has perfect frequency response with infinite cutoff for out of band signals, which is not at all realizable. On a practical note, to achieve 50dB of stopband rejection an additional 10%, or 6 kHz, of sample rate is sufficient for an active-element anti-aliasing filter. These active filters, however, can consume significant amounts of power (hundreds of milliwatts), and due to the ultra-low power requirement for Penguin are less than ideal. To combat this waste of power, Penguin uses a passive element anti-aliasing filter which consumes approximately 20 mW of power. This passive filter, however, requires nearly 20 kHz of roll-off to reach comparable stopband performance. Thus, the required sample rate is much closer to 100 kHz. This trade-off selection is valid only if the required increase in power due to the higher sampling frequency is less

than the additional power that an active filter would required, and is proven to be true later in this chapter. Additionally, since it is often useful to determine the physical origination location of a signal, two orthogonal receiver channels are necessary, which follows the classical antenna type for an ELF/VLF ground-based receiver.

The dynamic range requirements of the Penguin system are not as clear-cut as the frequency requirements. Dynamic range encompasses information about the background noise level of the environment and receiver system where it is installed as well as the signal strengths of both the weakest and strongest signals present that are desired to be studied and thus must be recorded in an undistorted fashion. Existing ELF/VLF receiver systems on the Antarctic Plateau to study the same phenomena, such as the AGO systems, have a 12-bit ADC at the front end without any automatic gain control, thus limiting the dynamic range to at most 74 dB. Data from these systems, however, typically exhibit closer to 50 dB of dynamic range which has been sufficient for all scientific studies thus far. As such, the Penguin system is designed to have a dynamic range of at least 50 dB end to end, with a desire to increase this range as much as practical.

Along with frequency and amplitude requirements on the data come time. Time is important for both how much data must be collected (i.e. data volume) as well as how accurately the time stamp of the data must be. First it is necessary to determine how much data must be collected to still be useful. Chorus and Hiss are long term phenomena which develop over hours, however it is not necessarily practical to record all of this data due to power and storage limitations. Minimizing the volume of data recorded eases meeting the longevity requirements of the system. In consultation with the data end users, it is determined that a 2-second snapshot once every fifteen minutes would be ample to study the ambient ELF/VLF environment. Combined with the larger bandwidth of Penguin and possibly higher dynamic range, Penguin

will provide better broadband data than is currently available from the large and expensive AGO systems. This snapshot mode also greatly reduces data compared to continuous recording.

With a determination of the volume of data complete, the absolute timing accuracy of the system must be specified. For the studies being conducted with Penguin, it is ideal for at least some of the data from each snapshot collected by each of multiple Penguin units across the Antarctic Continent to overlap. More formally, the frequency departure of the timing system for each Penguin receiver must be small enough to ensure that no system drifts more than one second per time synchronization period. With the readily available nature of simple crystal oscillators with normalized frequency departures of less than 10^{-11}, time synchronization can be as infrequent as once per day to easily maintain this requirement [2]. Since time synchronization is a potentially costly operation in terms of power, limiting this synchronization frequency is another method of reducing overall power consumption.

Now that bandwidth, dynamic range, and duration have been established, the amount of storage space can be calculated. To maximize the dynamic range and ensure that the sampling and digital systems are not the limiting factor, a 16-bit analog-to-digital converter (ADC) is selected for Penguin. This results in one 16-bit word (or 2 bytes) per sample per channel. Taking this data from two independent channels with $F_s = 100$ kHz results in 400,000 bytes of data per second. With a recording schedule of two seconds every 15 minutes, 76.8 MBytes are sampled per day or just over 28 GBytes per year. In addition to this 28 GBytes of ELF/VLF data, several hundred MBytes should be available for system logs and timekeeping information, resulting in a storage requirement of approximately 30 GBytes per year which is feasible with currently available COTS flash storage mediums.

Power and physical volume should be minimized as low as practically possible. Existing receiver systems require hundreds if not thousands of pounds of batteries

or fuel to operate per year, with the AGO systems requiring an entire small building to operate. Penguin must be transportable for extremely remote deployment, with weight limits typically in the hundreds of pounds for cargo. Less power and less weight are obviously ideal, and should be the top priority and consideration in design tradeoffs after the basic scientific requirements previously outlined are met.

Lastly, the Penguin system must be able to operate in one of the harshest environments on the planet. Temperatures in Antarctica are extraordinarily cold ($\leq -50°$C), with 6 months of continual darkness each year. The snow itself can be used as an insulator, maintaining a relatively constant -55° C throughout the year even as temperatures plunge below -70° C in the winter [58]. Components selected for use must be able to either operate in these extreme temperatures, or be sufficiently insulated from the outside environment, possibly with the addition of heaters, to function properly. While many integrated circuits are rated to this temperature by binning of production chips, it is important to note at that many battery technologies freeze and cease to function at significantly warmer temperatures than are expected with Penguin.

3.3 System Architecture

3.3.1 Background

While Penguin is designed to take as much as possible from legacy ELF/VLF receiver systems, the key to Penguin's success is to modernize them to the most recently available technologies, and reduce the system power while maximally increasing the system performance. By utilizing legacy system architectures, risk to the Penguin program is minimized as compared with taking a radical approach using an unproven and/or novel architecture. Additionally, the lessons learned over

years of systems design can be incorporated to end with a polished finished product. Lastly, due to advances in semiconductor, battery, and circuit technologies, extremely careful design of the new Penguin system can greatly reduce power consumption as compared to any other known system, while maintaing as good, if not better quality of data. Penguin as a system serves to illustrate the best possible performance from these legacy high-level system designs. The VAT system, a radical new architecture described in chapter 4, uses Penguin as the baseline best possible traditional system to demonstrate the extreme performance possible by changing the fundamental architecture of ELF/VLF receiver systems.

The traditional architecture adopted by Penguin is shown in Figure 2.1, where the "Signal Processor" block is the main system CPU. Penguin, however, replaces this commodity CPU with an FPGA and customized CPU to critically meet the system performance requirement without having additional wasted capabilities that are unused but still consume power. For the innovative design trade studies and decisions necessary to accomplish order of magnitude better power consumption over existing systems an in-depth look of each sampling subsystem is required.

3.3.2   Analog Front End

The analog front end used in the Penguin system is the results of a separate in-house development effort of a magnetic field low noise amplifier application specific integrated circuit. The second fabricated version of this chip, designed by Stanford VLF Graduate Student Sarah Harriman, is used in the deployed Penguin system. The specific LNA ASICs used in the system provide greater than 60 dB of spurious free dynamic range (SFDR) across the entire ELF/VLF band. The implementation details of the LNA itself are outside the scope of this document as they are not the author's work. The use of the LNA in the system, however, is critical for determining the minimum requirements for the remainder of the system.

As indicated in Section 3.2, a passive anti-aliasing filter is required to reduce the system power consumption. A seventh order Cauer elliptical filter is selected to provide approximately 80 dB of stopband rejection for 100 kHz sample rates, with the maximum unaliased frequency of 30 kHz [56]. The filter is based around a chain of capacitors and inductors to form the actual filter, plus an Analog Devices AD8606 instrumentation-grade operational amplifier to counteract the filter losses as well as add system gain to match the front end with the analog-to-digital converter.

The implemented filter utilized custom wound inductors with an extremely high Q to provide the sharpest roll-off possible. The custom inductors contain a pot-core with an internal gap to reduce the temperature dependence of the impedance. In-house testing of commercial, non-gapped cores resulted in 10 dB of attenuation at the maximum unaliased frequency and 30 dB less stopband rejection when cooled to -55° C, as compared to the identical circuit with custom gapped pot-core inductors.

A separate AAF is required for each of the two VLF channels. Both AAFs are designed to fit on a single card which directly plugs into the digital printed circuit board, which contains the sampling, processing, and data storage systems. By having the AAF as a removable module, different filter topologies, bandwidths, roll-offs, and gains may be inserted into the Penguin system without significant modifications for easy adaptability to future scientific needs or differing LNAs. This modular approach has been extremely beneficial in the AGO and AWESOME systems.

Penguin does allow, however, for a significant separation between the low noise amplifier and the remainder of the electronics, most importantly the digital system. This separation is most common with AWESOME system deployments, where the AAF and receiver PC are located indoors near power lines, while the antenna, LNA, and line driving circuitry are located hundreds or thousands of meters away from such interference sources. Penguin, due to its extremely low power consumption during operations and careful system design, has minimal radiated emissions and

thus does not require large separations for interference-free data. The separation capability, however, is maintained in Penguin for these legacy reasons.

3.3.3 Sampling System

The sampling system of Penguin directly interfaces between the main processing unit and the analog signals provided by the front end AAF described above. In the traditional architecture used in Penguin, the analog-to-digital converter is the only component that is exclusively part of the sampling system; the main CPU is implied as controlling the ADC as required to perform sampling operations. The principal requirements on the ADC itself include input signal type and voltage range, supported sampling rates, sampling resolution, performance (SFDR/SINAD/ENOB), supply voltages, digital interface type, and power consumption.

As the most basic requirement, the ADC must be able to properly interface electrically with the front end. Penguin utilizes a differential signaling type out of the LNA to minimize induced signals on the wires [44]. Single-ended signaling provides no common mode rejection, and is ill-suited for long runs of sensitive signals, thus making differential signaling especially important for the potentially long link between the LNA and AAF/sampling system. The AAF itself is intrinsically single-ended, but immediately converts back to differential to again maintain good noise immunity between the AAF and the actual ADC inputs. Thus, for maximum signal integrity, a differential input ADC must be used.

The performance of the ADC, namely the Spurious Free Dynamic Range (SFDR), Signal plus Noise And Distortion (SINAD), and Effective Number of Bits (ENOB) must be selected to at least meet, if not exceed, the similar specifications of the entire front end section plus the anticipated signals. The LNA, as tested in a laboratory setting, provides not less than 60 dB of SFDR, while the AAF provides at least 80 dB of stopband rejection. Thus, the ADC must have at least 80 dB of SFDR to

avoid being a limiting factor in system performance. This results in a requirement of at least 13 ENOB for the currently used LNA and AAF.

The signal to be sampled is not only filtered but also amplified within the AAF section to better match the voltage rails of the ADC. The voltage rails of the AAF are set to 3.3 V to accept the maximal 2.5 V signal from the LNA with minimal distortion. Amplifying the signal to levels too close to the voltage rails of the operational amplifier can result in signal distortion; the devices are most linear for small signal levels. Thus, we should expect the signals leaving the AAF to also be limited to approximately 2.5 V maximum. On the ADC side, however, maximizing the rail span of the signal is important for utilizing the full dynamic range of the converter. This is the typical tug-of-war between the analog designer's desire to keep the signal levels low to maximize linearity of the components in use versus the digital/ADC designer's desire to maximize the span of the signals compared to the ADC rails to maintain the highest dynamic range. As a compromise in Penguin, a 3.3 V supply is used in the AAF gain stages while keeping the maximal signal amplitude below 2.5 V. The ADC, however, is selected to have a 2.5 V rail span to match the analog signal, while having a 3.3 V digital I/O level to mesh with the selected processor without the need for level converters. 3.3 V is selected for the main processor to also match the voltage requirements of the selected nonvolatile storage medium discussed in Section 3.3.6.

Several digital interfaces are available to connect the ADC to the main processor, with parallel versus serial being the first branching point in the decision. Parallel will provide a separate signal path for each bit of the sampled value, plus control lines. While this allows the data to be read out from the ADC in a single CPU clock cycle, it does require a significant number of I/O lines and signals to be routed on the board. Serial interfaces, on the other hand, allow for typically 2-4 I/O lines to be used for the same data, however additional complexity is necessary inside of the

CPU to process the bit stream. Many CPUs include serial interfaces as dedicated peripherals, so this additional processing requirement is minimal. Thus, a serial interface will maintain minimal complexity in the physical system with no other impact on system performance with this architecture.

As discussed in Section 3.2, the ADC samples at 100 kHz, resulting in 200 kBytes per second or 1.6 MBits per second of data output assuming a 16-bit ADC. At face value one might interpret this to indicate that the digital interface must be able to transfer 1.6 Mbps, however this is incorrect. Most ADCs require that the readback of a sample be completed prior to the commencement of the next sample and hold. The actual analog-to-digital conversion is also not an instant process, even on a Flash ADC, and thus less than the entire sample period is available for readback, with often only about half of the period free. This timing limitation varies greatly by each specific ADC, but must be accounted for when designing the digital interface. Assuming that 50% of the sampling period is available for readback, the Penguin digital interface must operate at least 3.2 Mbps. Serial Peripheral Interface (SPI) is the most common serial interface used by ADCs with multi-megabit per second bandwidths, and thus is the ideal candidate for use in Penguin.

Combining all of the requirements of this section along with a survey of currently available COTS ADCs, the Analog Devices AD7687 is selected [20]. This ADC provides 96 dB of dynamic range, and over 93 dB of SINAD at the desired frequency and voltage operating conditions present in Penguin. The differential input allows for a fully differential analog signal path depending on the specific AAF used in the application, while the SPI digital interface allows for a minimum pin-count interface between the host processor and ADC. Lastly, the AD7687 is selected for its extremely low power consumption compared to other available devices, which is paramount for long-term autonomous operation.

3.3.4 Time Keeping and Synchronization

As driven by the scientific data requirements outlined in Section 3.2, Penguin must internally know the time with better than one second accuracy at any given instant. This absolute timing requirement translates to a normalized frequency departure, $Y(t)$, of one part in 86,400 or approximately 10^{-5} [3]. This $Y(t)$ requirement is extremely easy to obtain with modern crystal oscillators, which typically provide $Y(t)$ from 10^{-8} to 10^{-9}. Given $Y(t) = 10^{-8}$, one can expect approximately an 860 μs departure per day, which is significantly longer than a single sampling period at 100 kHz. While Penguin is not required to have subsample absolute timing resolution, it is important to understand that such resolution is possible with a crystal oscillator time base given that it is resynchronized at least every 16 minutes, given the assumed $Y(t) = 10^{-8}$.

To minimize power consumption while still maintaing rather accurate time information, Penguin is designed to resynchronize its internal oscillator with a time standard once per day, yielding an expected normalized frequency departure of 860 μs. This design allows for data from multiple Penguin sites to be compared with general ease, however still without interferometry capabilities which are discussed later in this section. The time standard most available to Penguin in such remote locations is GPS, which provides 10 ns level time accuracy globally utilizing relatively small and low power receivers [36]. The selected GPS receiver is an i-Lotus (formerly Motorola) M12M timing receiver [28]. This receiver provides a 1 pulse per second (PPS) output that is synchronized to UTC within 20 ns. In addition, this timing output can be delayed by an arbitrary, programable amount to compensate for cable and output propagation delays.

The main timekeeping system consists of a high-precision quartz crystal oscillator combined with a discrete logic counter and two discrete logic dividers. The main crystal, clocked at 2 MHz, is divided down to the sampling clock frequency of 100

kHz. This 100 kHz signal is then continuously counted in a discrete, low power logic, 40-bit counter. This counter allows for more than 100 days of timekeeping; however, note that a 32-bit counter, which is the next smallest available counter constructed from COTS 8-bit parts, would be unable to keep track of time for even a single day. In addition to the 100 kHz counter, the same signal is also divided down to create a 1 PPS signal, used to determine the instant that the system should begin sampling on the integer second boundary of wall clock time. Lastly, the exact UTC date and time when the counters and dividers were last re-synchronized and reset is maintained in non-volatile memory in the system control microcontroller. In total, the system thus tracks the last time the internal clock was reset as well as an internal stopwatch of the estimated time elapsed since then. Combining these two pieces of information allows the system to know the current UTC time to the accuracy of the internal crystal oscillator. The system also knows how much time has elapsed since the last resynchronization and thus can schedule time resynchronizations as necessary. The principle hardware constituants of the timekeeping and resynchronization system are shown in Figure 3.1.
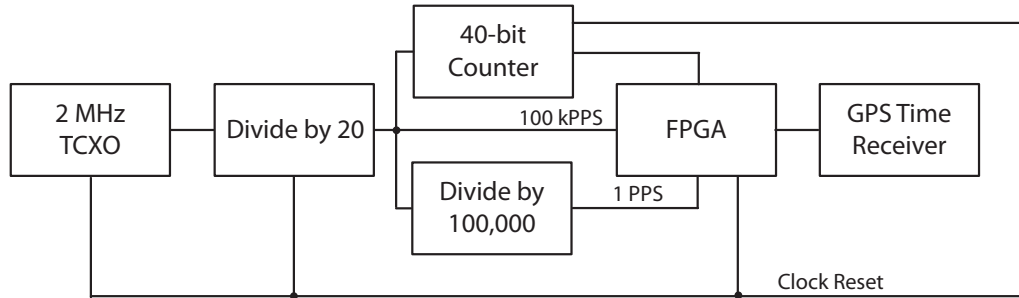

Figure 3.1: Penguin Time Keeping System

The resynchronization process occurs due to either the system powering on from a previously powerless state or when the time since last update is greater than 24 hours. In the first case, during bootstrap the system control microcontroller clears the last synchronization time from its internal memory, forcing a time update when

the main system processor comes online. In the later case, the main system processor directly determines from a read of the timekeeping counter that at least 24 hours has elapsed. In either case, the actual synchronization process is the same. First, the main processor must be running in addition to the system control microcontroller (see Section 3.3.7 for details). Next, the main system processor instructs the system control processor to power on the GPS receiver. This functionality is not directly handled by the main system processor to avoid accidentally powering on the GPS system due to glitches on the main system processor's I/O lines during power-up. If the GPS power is directly controlled by the main system processor, as in earlier revisions of the Penguin system, the internal 1 PPS signal would be corrupted on a random basis.

With the GPS powered up, the main system processor initializes the GPS receiver for proper settings, such as elevation blanking, UTC time mode, and the requisite 1 PPS delay time. The GPS is then polled once per second until both GPS signal lock is achieved as well as the number of leap seconds is updated. The leap seconds count is required to calculate true UTC from the GPS time, and is only received once per GPS navigation message, which may take up to 12.5 minutes to receive after lock on the first GPS satellite is acquired [17]. Once locked, the 1 PPS counter is automatically resynchronized to the edge of the 1 PPS signal from the GPS. The 100 kHz counter is read back on one of these 1 PPS edges, as well as the exact true UTC time for the GPS receiver. These two pieces of information are recorded to a log file and are used to back out the drift of the Penguin oscillator for post-processing of the data after the system is recovered from the field. On the next 1 PPS edge, the 100 kHz counter is cleared and the 100 kHz generator is reset, and the UTC time of the reset is stored in the system control microcontroller's internal memory. This process assures that both the 100 kHz signal as well as the internal 1 PPS signal are within one 2 MHz period (500 nS) of true UTC.

The 500 ns accuracy limit is the direct result of using a 2 MHz internal crystal. By replacing the crystal with a VCO or other device with adjustable phase, this 500 nS limit can be removed, however at a significant power penalty. This limit also is quickly overshadowed by the frequency departure of the crystal itself and thus is not a prime candidate for optimizing system time performance. The system should, however, be calibrated at the anticipated operating temperature to remove any UTC offsets. Since the clear signal to the counters and dividers has a non-zero propagation time, especially when the delay of the GPS 1 PPS edge detector is added, the system clock has an intrinsic offset from actual UTC. This offset, however, is directly due to the components involved and board layout, and thus can be calibrated for when each board is manufactured. The calibration process consists of running the system and allowing it to resynchronize with GPS without any propagation delay factor programmed into the GPS receiver. The output 1 PPS edge from the Penguin system clock should then be observed on an oscilloscope with a scan rate of at least 100 nS per division, if not faster. On another channel of the oscilloscope, a reference 1 PPS signal, possibly directly from the onboard GPS receiver, should be simultaneously viewed. The offset between these two signals is then measured and programmed into the Penguin system as the desired GPS cable delay factor. This process can be iterated if necessary to minimize the timing difference.

If a higher level of timing precision is necessary, the GPS resynchronization period can be decreased, or for maximal precision the GPS receiver can be used to constantly discipline a 100 kHz clock source. There are many methods possible for constant GPS disciplining. The method used in both the AWESOME receivers as well as the buoys utilized a high frequency Voltage Controlled Oscillator (VCO) running at 40 MHz, which was directly disciplined and divided down to the required 100 kHz clock. By fine tuning the frequency of this HF oscillator, and allowing for the divide down factor to be digitally controllable, the accuracy of the output clock

is limited by the accuracy of the GPS 1 PPS signal, which is typically better than 20 ns. As implemented on the buoy, this system requires on the order of 100 mW of power continuously to operate, largely due to the power requirements of current GPS receivers, which is typically very acceptable when interferometric measurements are required of the system. Penguin does not implement continual clock disciplining since it is not necessary to meet the current Penguin scientific objectives.

3.3.5   Scheduler and Power Control

To minimize power consumption, the Penguin system is designed to continuously maintain timekeeping information as described in the previous section, while duty cycling the main system processor to only be awake for the required timekeeping maintenance and well as scientific data sampling, processing, and storage operations. The wake and sleep control of the main system processor is under the purview of the system control microcontroller. This controller is responsible for controlling power to the GPS receiver, remembering the UTC time and date of the last time synchronization, and waking/sleeping the main system processor as necessary to carry out its duties.

Power control of the GPS receiver requires only a single General Purpose Input/Output (GPIO) pin, and minimal RAM is necessary to store the date and time of last synchronization, thus nearly any microcontroller suffices for these needs. The system scheduler is implemented in C and requires a single GPIO line for power control of the main system processor plus some type of serial communication method with the main processor. Clearly, the RAM, flash, and IO requirements of the system control microcontroller are not major drivers for the selection of the particular device. Power, however, is of paramount concern since this microcontroller will be operating continuously throughout the deployment of a Penguin system.

There are three major factors to consider when looking for an ultra-low power

microcontroller: the required clock speed, the running power consumption at this clock rate, and any sleep or low-power modes and their associated power consumptions. The required clock rate of the Penguin system control microcontroller is extremely low, as it is only required to check for scheduled activity once per second. The always running 100 kHz onboard clock provides an alluring clock source for the microcontroller to avoid operating another clock continuously. Simulations of the scheduling software indicates that a 100 kHz clock provides more than sufficient speed to meet timing requirements. Thus, microcontrollers of interest should operate from an external 100 kHz clock source. The actual run-mode power consumption and power saving features of the candidate processors is then compared to select the best candidate for the job.

A search and review of currently available microcontrollers in light of the above requirement, coupled with our past experience and code base for Atmel products, yields the ATtiny series of microcontrollers as the ideal candidate. Specifically, the ATtiny84 is selected due to its spacious 8 kB of flash and 2 kB of RAM, along with an internal SPI controller for communications with the main system processor and sufficient GPIO lines for its required power control operations [15]. While the microcontroller's power draw contains both a static component plus a frequency dependent component, when measured in bench tests replicating the required system functionality, it was shown to consume approximately 30 mW when operating at 100 kHz without any other power conservation measures in place. By activating ultra-low power sleep modes on the processor, allowing it to operate in normal run mode for only extremely brief periods as required to complete its tasking (10 ms) on each tick of the 1 PPS signal, the system control microcontroller power draw is reduced to sub-milliwatt levels.

3.3.6   Processing and Storage Systems

The main system processor is the unit with the most concentrated responsibility for proper system operation. This processor is required to check the time and resynchronize it as necessary, wait for the exact moment to begin sampling, and subsequently control both of the ADCs throughout the required sampling period, process the raw data if desired, and store the data on an associated long-term nonvolatile storage medium. With the time synchronization process details in Section 3.3.4, attention should be given to the sampling process.

The first step in sampling is to be sure that the proper instant has been reach to being collecting data. To make this determination, the processor needs direct access to the system 1 PPS line. Additionally, the processor needs access to the time of last synchronization, via SPI from the system control microcontroller, as well as the number of 100 kHz ticks since that synchronization, available through 3 GPIO lines to the clock counter shift registers. During the sampling process, the processor must send the sample clock to both ADCs, wait for a conversion complete signal, and the retrieve the data via a SPI bus. Providing a separate SPI interface to each ADC allows for simultaneous readback of the data from both ADCs, which can be advantageous since a single SPI read consists of many processor clock cycles. A typical SPI bus requires a clock (SCK), Master Out Slave In data (MOSI), Master In Slave Out (MISO), and a Slave Select (SS or CS) [31]. Since the ADCs each have their own dedicated SPI controller, as long as the processor can disable the SCK when not reading back data the SS lines are not required. This results in only 8 I/O lines necessary from the processor for both ADCs, with two each of: sampling clock, SCK, MISO, and MOSI. If pin count on the main processor becomes an issue, the sampling clock can be reduced to a single output shared among the ADCs, however this removes the ability to independently skew the ADC conversion signals inside of the processor to accommodate for differing circuit board propagation delays. Since

SPI operates as a shift register in each node, it is also possible to have only a single SPI controller for both ADCs and chain them as necessary to reduce pin count as well.

With timekeeping and sampling electrical needs established for the main processor, the method of long-term digital nonvolatile storage of the end data must be determined. The science requirements detailed in Section 3.2 indicate that approximately 30 GB of total data, including system logs, will be generated each year. At the time of initial design, multiple options COTS options were available including: 16 GB CompactFlash (CF) cards, 8 GB Secure Digital High Capacity (SDHC) cards, 16 and 32 GB Solid State Disks (SSDs), as well as 320 GB traditional magnetic storage 2.5" hard drives. The last option is least desirable due to its extremely high power draw compared to all other options as well as reliance on mechanical parts which have not been tested to the extremely low temperature requirements for Penguin. SSDs tend to have greater performance than a CF or SD card, however their increase in performance is offset by a similar increase in power for the additional RAM buffers and processing that is present internal to the SSD. The slight speed advantages (30 MBps for SSD versus 15 MBps for CF according to various product packaging) is not likely to be encountered by Penguin due to its relatively low data throughput of less than 1 MB every two seconds. The increase in required power, from around 100 mW for CF compared with greater than 1 W for Penguin, it much more worrisome, and thus SSD is removed from consideration.

The last two options, CF and SDHC, are in close competition. Both have comparable power requirements, while the CF cards generally available have higher capacities than SDHC due to their physically larger size. SDHC, however, is accessed by a serial interface requiring less than 8 wires for full speed operation, while CF cards when operating in True IDE mode require 27 I/O lines. Given the concern for total system quality over complexity of PCB layout, CF cards are selected due

to their superior storage capacities and larger variety of available cards.

Selection of a particular flash card, or even flash storage technology as a whole, for use on the Antarctic Plateau requires vigilant testing to ensure survivability of the device. Manufacturers generally rate their cards to 0° C, while the ambient buried temperature is anticipated to be -55° C as per Section 3.2. Studies are available of a known Low Temperature Data Retention (LTDR) issue caused by program/erase cycling of flash memories at temperatures under 200° C, however these experiments are typically done near room temperature and are not necessarily indicative of performance at extremely low temperatures [33]. In-house testing of candidate medias is clearly necessary to ensure system operations.

A test fixture of early Penguin prototype hardware with an onboard FPGA and two CompactFlash slots is used as the electrical basis for the test. The FPGA, independently tested to extremely low temperatures, is programmed to continuously write and read data from a single CF card, and output the results of its testing in real time through a standard RS-232 port. The actual write tests are designed to maximize coverage of the flash cells. Because CF is based on NAND flash memories, entire 512-byte blocks must be written or erased [12]. Additionally, internal wear leveling control may be present inside of the CF device, causing different physical memory cells to be used each time the same logical block is written or erased [14]. Thus, to fully test the operation of the memory at temperature the entire card must be written.

Several write patterns are selected, namely all zeros, all ones, alternating zeros and ones, alternating ones and zeros, and lastly a hash of the blocks logical address. The first four tests write the same pattern to every byte on the card, and then read back the entire card to verify the pattern before erasing the card and moving on to the next test. The last test allows for blocks to be written and read on a random address basis, and uncovers any issues in the Flash Translation Layer where

a logical memory address may be incorrectly linked to a physical memory cell. All
five tests are run on the self-contained hardware in a thermal chamber at -55° C
while an external PC logged the tests results from the RS-232 port. One additional
read of the last test cycle is then executed at room temperature, followed by one
cycle of full erase-write-read tests to ensure continued flash survival. Testing takes
approximately 72 hours to complete, and multiple cards were tested. No failures
or bit errors were detected. An additional test of a single pattern across the entire
card was performed at -70° C and -105° C, again without any observed failures. As
a result of these amazing test results, the Sandisk 16 GB Extreme III flash card is
selected for deployment use in Penguin.

With confidence in the selected flash storage medium, a runtime RAM is the
only missing component to finalize processor pin count requirements. The Pen-
guin system as outlined by the immediate science needs does not require significant
onboard data processing capabilities. For future deployed use, however, it is desir-
able for onboard processing, such as real-time FFTs, to be possible with minimal
if any hardware modifications. Such signal processing requires a significant amount
of RAM which should be included in the Penguin design. Synchronous Dynamic
RAM, SDRAM, is the most common current technology of RAM for modern micro-
processors. Single Data Rate (SDR) is more than sufficient for the needs of Penguin
and offers significant power savings over DDR memories. For fastest processing, the
selected SDRAM should be at least 16-bits wide, since the science data being col-
lected and processed is also 16 bits. Utilizing a more narrow SDRAM would greatly
increase system latency, while a wider SDRAM may be able to provide additional
performance improvements if 32-bit floating point math is used in signal process-
ing. Careful selection of COTS SDRAM modules available leads to the selection of
the Micron Technologies MT48LCxxxM16 series of modules. These modules are all
16-bits wide with pin-compatible options of up to 512 Mbits per chip.

Assuming the worst case, electrically speaking, of a 512 Mbit SDRAM in Penguin, along with all of the other I/O needs discussed in this section, the main system processor must have at least 118 I/O lines available as per Table 3.3.6. If a traditional microprocessor were used, this pin count would include the external memory and address bus, as required to interface with the SDRAM. For Penguin, one of the primary goals is to reduce the power footprint compared to a system utilizing a hard microprocessor. To have the flexibility to remove unnecessary peripherals and critically clock what remains at the lowest frequency, a Field Programable Gate Array is necessary. The FPGA allows for a full processor of arbitrary design to be instantiated in it without any changes to the physical hardware or circuit board. Due to heritage and an existing code base, an Altera Cyclone II series of FPGA is selected. The exact model must not only support the 118 I/O lines required but also have a sufficiently high number of internal Logic Elements (LEs) to support the instantiation of a full processor plus and required signal processing.

Device	I/O Pins	Quantity
100 kHz Clock Counter	4	1
AD7687 ADC	4	2
CompactFlash	27	2
External Clocks	1	3
GPS UART + GPS 1PPS	3	1
MT48LCxxxM16 SDRAM	39	1
RS-232 Debug Port	4	1
System Control Microcontroller (SPI)	3	1
Total	**118**	

Table 3.1: Pin counts required for Penguin main system processor

Simultaions of the processor architecture, described in Section 3.5.1, indicate that 5k LEs are necessary for instantiation of the desired processor. Additionally, a hardware FFT engine tightly coupled with this processor as well as a Direct

Memory Access (DMA) controller can be added for a cost of 10k additional LEs. For a maximally flexible hardware platform, the Penguin FPGA housing the main system processor needs at least 15k LEs. The closest available FPGA in the Cyclone II line contains 20k LEs and is thus an ideal candidate. This same FPGA is available both in a 256-pin Fineline Ball Grid Array package as well as a 240-pin Quad Flat Pack. The later is physically large leading to dramatically increased lead inductance, making signal integrity of high slew rate signals, such as those associated with the SDRAM, much more difficult. Thus, the 256-pin FBGA is the preferred, resulting in the use of a EP2C20F256N6 chip in Penguin.

3.3.7   Digital Hardware Architecture

With all of the individual pieces selected for the Penguin system, the entire digital system architecture can be described. The system has a flat, linear layout as detailed in Figure 3.2. A precision quartz crystal and associated oscillator circuit generate a 2 MHz cock. This clock is fed into a divider to generate a 100 kHz clock for sampling. This 100 kHz clock is fed into an additional divider to generate a 1 PPS signal, to the system control microcontroller for its own clock, the main system processor FPGA for sampling use, as well as into a counter to keep track of the time since the last GPS synchronization. The 1 PPS is used both by the system control microcontroller for scheduling as well as the main system processor for determining the sampling start instant. The system control microcontroller interfaces with the voltage regulators to wake and seep the main system processor and GPS receiver, and includes a separate SPI interface directly to the main system processor to receive schedule updates and timekeeping information.

The main system processor FPGA is directly connected to the ADCs for both sampling control as well as data read back. Additionally, the FPGA has its own clock for the internal processor, as well as an external SDRAM and two CF card slots.
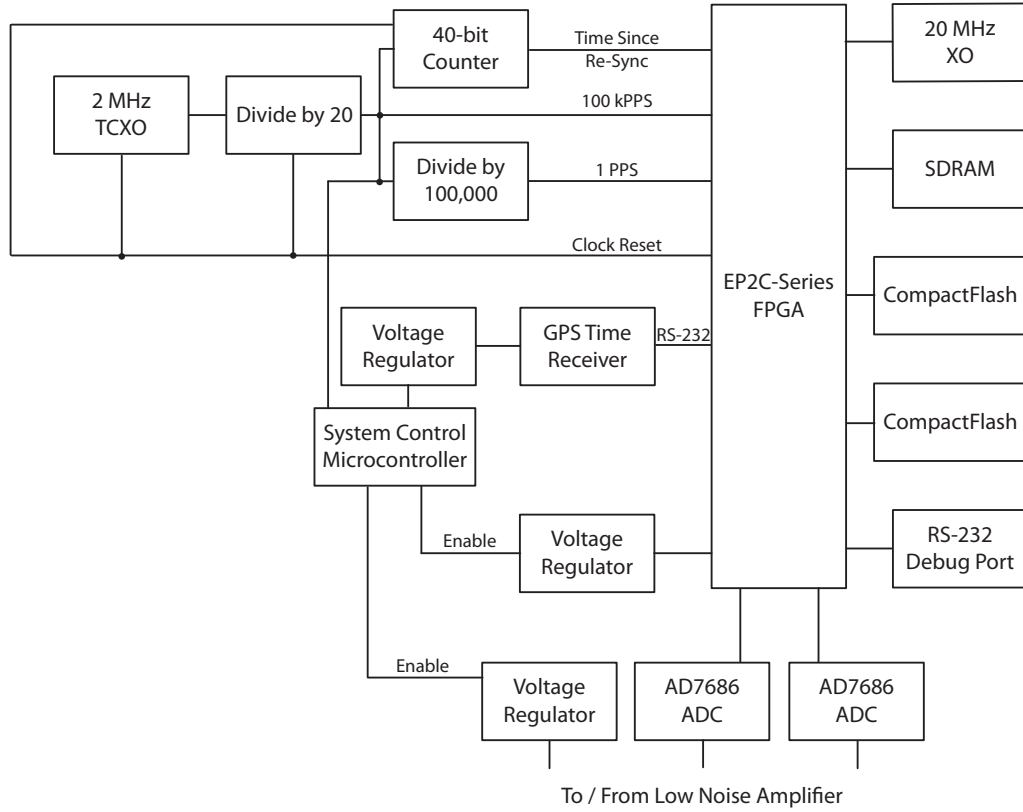
Figure 3.2: Penguin Digital Architecture

The system processor contains two RS-232 connections to external peripherals: one for control signals to and timing information from the GPS receiver and a second external RS-232 serial port for debugging as well as system status information.

3.4 Power Budget

As a result of the limited times when the GPS receiver or system processor need to be operational, the voltage regulators are split into three distinct power domains. These separate power interconnections allow for the main system processor and GPS receiver to each be independently turned on or off, consuming zero power while off as compared to some standby current common when a CPU is on but in a sleep mode.

The result are three power domains: domain 0 containing the 2 MHz clock and associated dividers, counters, and system control microcontroller, domain 1 containing the ADCs, main system processor, and associated SDRAM and CF cards, and domain 2 containing the GPS receiver. Power domain 0 is always operational so that the system can keep count of absolute UTC time and wake the remainder of the system as necessary to capture the desired data or maintain the clock. Power domain 1 is active for the entirety of any sampling, processing, and recording operations, and is also required during GPS time synchronization. Power domain 3 is only active during the daily time sync. Table 3.4 below lists the major components from the hardware architecture, their associated power domains, and a per-component power estimate derived from datasheets and historical systems. The power estimates are made assuming a linear regulator with an input supply voltage of 3.6 V.

Device	Power Domain	Power (mW)	Quantity
100 kHz Clock Divider & Counter	0	3	1
AD7687 ADC	1	3	2
Analog Front End (LNA+AAF)	1	50	1
CompactFlash	1	60	2
External Clock, 2 MHz	0	5	1
External Clock, 20 MHz FPGA	1	5	1
FPGA, Main System Processor	1	400	1
GPS Receiver	2	100	1
MT48LCxxxM16 SDRAM	1	300	1
System Control Microcontroller	0	2	1

Table 3.2: Power domains and estimated draw of Penguin peripherals

The raw data contained in Table 3.4 is sufficient to estimate the average power consumption of the Penguin system, which is critical for both power system and thermal design. The average power of a system by definition is the total power used by a system in each of its operating modes divided by the duty cycle percentage of

each modes. More formally, the average power P_{tot} is related to each operational mode or power domain i by Equation 3.1 where D_i is the duty cycle for the given component and P_i is the average power during the on-time of the same component.

$$P_{tot} = \sum D_i \times P_i \tag{3.1}$$

Components with multiple modes of operation, for example the system control microcontroller with a nominally running power plus a nominally sleeping power, can be split into distinct components for use in Equation 3.1. Following the system control microcontroller as a trivial example, we know that the ATtiny84 consumes 100 μW of power during its sleep period, which is 900 ms per second, and 10 mW during its wake period, which is 100 ms per second. Plugging this information into Equation 3.1 yields an average power of 1.09 mW for the system control microcontroller. This value with a duty cycle of 1 can then be used in a complete system power budget.

In addition to the information in Table 3.4, the duty cycle of each power domain must be known to form a total system power budget. The duty cycle for power domain 0 is known to be 1 since it is defined as always running. Power domain 1 is only active during data collection, processing, and storage, which takes place every 15 minutes. The active time for this operation must be estimated. To collect the two seconds of data, the system must wake up prior to the start of sampling, wait for the exact instant to sample, spend two seconds sampling and potentially a few more seconds for any post-processing and recording. As an extremely conservative estimate, it is estimated that domain 0 will be operational for 30 seconds every 15 minutes, or a duty cycle of 0.0333. Additionally, domain 0 must be active for time resynchronization. The average case for time synchronization is 7 minutes for a GPS cold start, satellite acquisition, and UTC offset reception. Since this occurs once

per day, the additional partial duty cycle is 7 minutes per 1440 minutes or 0.00486. Combining the two utilization factors of power domain 0 yields a total duty cycle of 0.0382. The duty cycle for power domain 2 is simply the duty cycle of a GPS synchronization, previously shown to be 0.00486.

Combining these duty cycle estimates with the estimated per-component power draw from Table 3.4 and Equation 3.1 results in an average power estimate for the Penguin system of 44 mW. Additionally, a peak-power estimate is necessary for proper sizing of the power supplies. This worst case number is achieved when all peripherals are on and operating at their maximal current draw. The estimated Penguin peak power is 1 W or 280 mA at 3.6 V.

3.4.1   Battery Selection

To operate autonomously on the Antarctic Plateau, Penguin must contain an internal battery supply. During the 6 months of total darkness in the winter, the plateau experiences relatively little wind and thus both solar and wind power systems are not feasible. With an operating temperature of -55° C, most common battery technologies provide minimal, if any, capacity [45]. Both primary and secondary cell batteries are explored, although due to the lack of wind and sunlight secondary cells would only be recharged once per year during the annual field maintenance trip.

Traditional Alkaline batteries typically freeze and cease to function around -20° C, while specially formulated chemistries are available that do not cease to operate until -40° C [4]. Lead Acid batteries also are based on a sulfuric acid electrolyte which freezes close to 0° C as the battery discharges, rendering them useless in such a cold environment [27]. These batteries are also extremely heavy and thus less than ideal for Antarctic deployment. Nickle Metal Hydride (NiMH) cells are only well-preforming to 0° C, with advanced versions available that operate down to -30° C [21]. Penguin, however, must operate at -55° C, and therefore none of these typical

battery candidates could survive the Antarctic Plateau.

Wet cathode batteries tend to have significantly better performance at extremely cold temperatures, however they are generally more hazardous to ship than any other technology discussed here [59]. Lithium-thionyl chloride batteries are one example of such design, are commonly used in military systems, and are readily available from various vendors. Saft provides several lines of lithium-thionyl chloride that are rated for operation at -60° C, well meeting Penguin system requirements [51].

Out of the available COTS batteries, the Saft LSH20 is selected for the base cell in the Penguin battery module[50]. These cells supply up to 20 A-hr with a nominal open circuit voltage of 3.66 V. Saft, the manufacturer, rates the cells for use to -60° C and provides nominal discharge curves down to -40° C. Utilizing this information, at -55° C we anticipate 8 A-hr per cell. With an average power estimate of 44 mW for the penguin system running at 3.6 V, 385 W-hr or 107 A-hr at 3.6 volts are required per year. With the Saft LSH20 derated for the low temperature, this equates to less than 14 cells in parallel.

It is critical to note that during discharge of all current battery technologies, the supplied voltage drops a significant amount over the discharge cycle. Additionally, the open circuit voltage of a lithium thionyl-chloride cell is inversely proportional to the ambient temperature [50]. In deployed operation, the cell voltage is closer to 3.3 V at mid discharge and low temperature. To avoid possible electromagnetic interference from the power regulation circuitry which could corrupt the VLF science data, linear regulators are used in Penguin rather than more power efficient switching regulators (see Section 5.2.1 for more information). One disadvantage of linear voltage regulators for battery powered devices is that the input voltage, ie the supply voltage from the battery, must be at least some given amount higher than the output voltage of the regulator [55]. This voltage difference, or dropout, can be as low as several hundred millivolts or as high as the gate oxides and thermal dissipation of

the regulator will allow.

The practical implication of the non-negative dropout voltage for Penguin is that the battery module must supply at least 3.3 V plus the dropout (nominally an additional 300 mV) in order to power the system logic which runs at 1.8, 2.5, and 3.3 V. The Saft LSH20 will initially meet this requirement with ease, however with little use at the extreme temperatures expected the voltage will quickly drop below this critical threshold according to published discharge curves [50]. As a result, the Penguin battery module consists of series-connected pairs of LSH20 cells all together in parallel. This design offers 7.2 V initial operating voltage, and will remain above the 3.6 V dropout throughout an extraordinarily high percentage of energy discharge of the cell. The downside is that twice as many cells are required as may initially be expected. The resulting Penguin battery module consists of 15 pairs of LSH20 cells, providing nearly two years of energy storage. This overspecification of the battery capacity is partially required to allow for pre-deployment testing of the full system, including battery module, for qualification of the actual deployed components and workmanship. Additional spare capacity, powering the Penguin system through one complete austral winter plus both the austral summer preceding and following, allows for increased flexibility of the deployment and recovery scheduling during the summer seasons when transportation to the Antarctic Plateau is possible.

Lithium thionyl chloride cells require special care in use to prevent venting or explosion. High discharge rates and reverse biasing of a cell can result in catastrophic destruction of the battery module, possibly resulting in a secondary fire [57]. To avoid these potentially disastrous situations, two specific protection measures are implemented in the battery pack. To prevent reverse biasing a more depleted cell from a less depleted cell, all parallel connected elements are each independently wired through a diode into the main battery bus. The voltage drop of the diode must be considered in voltage headroom calculations for the regulators, requiring

an additional 700 mV. Secondly, to prevent short circuits or other unexpected high current draws from causing thermal runaway and system destruction, Positive Temperature Coefficient (PTC) resistors are placed in series with each pair of cells. PTC resistors cause minimal loss of voltage and power during normal operations while providing a high resistance path when low impedance loads are presented to the battery pack [37].

The Penguin battery module, including plastic housing, weighs approximately 10 pounds, consumes less than 0.75 ft^3, and provides two years of operating time to the Penguin system as designed. At room temperature, the energy density is greater than 4 kW-hr per ft ^{3}. Rated for -60° C operation, the described battery pack supplies Penguin with 18 months of maintenance free power once deployed, in addition of up to six months of pre-deployment testing of the field hardware.

3.5 Firmware and Software Architecture

With the physical Penguin hardware fully described, there is one critical piece of system firmware and two additional pieces of software that are required for the system to operate as intended. First and foremost, the currently blank FPGA requires a firmware implementation of a microprocessor used as the main system processor. This "soft core" or firmware implemented processor subsequently requires operational software. Lastly, the system control microcontroller requires software for scheduling and timekeeping.

3.5.1 System Firmware

The firmware of the main system processor is one of the key places that excess components and wasted power are removed. As discussed in Section 3.3.1, one of the

radical improvements made in Penguin is the elimination of a standard, fully featured commodity central processing unit and replacing it with a more streamlined, lower powered (both in power consumption and computational ability) processor implemented in an FPGA. By defining the processor in firmware, custom peripherals, such as FFT processors, can be implemented and tightly coupled to the processing unit. Additionally, these enhancements can be made at any time without hardware modifications.

To determine the amount of power necessary in the CPU, one must analyze the functions that the processor is responsible for performing and associated processor requirements. Section 3.3.6 stated that the main processor is responsible for time resynchronization, triggering and controlling the sampling process on a hard real-time deadline, processing the recorded data, and storing it to the now-selected CompactFlash cards.

Timekeeping is based on externally clocked counters which are read by shift registers. While the maximum data length is greater than 32 bits, the system only needs to compare this clock value with a pre-defined time expiration value. This comparison can occur at the bit level, or more practically as each word is clocked in through the shift registers. Timekeeping thus requires a loop counter and comparator, both of which are trivial computational needs. The re-synchronization process requires reading and storing a-bit values for the year, month, day, hour, minute, and second, as well as actuating a clock reset line with extremely high precision on the edge of a GPS clock pulse. The 8-bit data storage drives the requirement for at least an 8-bit processor internally. The precision clocking, however, might be mistakenly viewed as driving a requirement for extremely high CPU clock rates. Without external hardware, a traditional CPU would need to be clocked sufficiently fast such that its time to respond to a GPS-based interrupt would be less than the maximum time error permitted in the system. Additionally, careful software design

would be required to ensure that the clock jitter is as small as possible. With the FPGA surrounding the CPU, however, external logic is available to aid in the clock reset function. The CPU is required to meet a very loose timing requirement of "arming" this section of external hardware in less than one clock period before the desired trigger, in our case 1 second. The external hardware, a flip-flop and AND gate, then latches on the clock edge to pass the GPS clock signal to the rest line of the timekeeping circuit with a well defined and extremely small about of delay. Thus, the CPU is only required to respond to interrupts within 1 second.

Control of the analog-to-digital converters has similar requirements for extremely high timing precision and extremely low jitter on the trigger lines. With similar external hardware, the CPU is relieved of the need to meet tight real-time deadlines to initiate ADC sampling by simply gating the externally generated sampling clock as necessary. With a sample rate of 100 kHz, both ADCs must complete their sampling operation and their data must be read back every 10 us. With the ADC requiring up to 3.2 us between the start of conversion and the data being ready, the CPU only has 6.8 us to read both ADCs of data [20]. Assuming that the ADCs are read sequentially, as is the case in Penguin, all 32 bits of data must be read in this period over the SPI interface, resulting in a required SPI clock rate of at least 32 bits / 6.8 us or 4.706 MHz. Many current SPI controllers require a main CPU clock running at least twice the internal SPI speed, if not four times as fast, resulting in a CPU clock requirement of at least 10 MHz. This clock requirement does not provide for any headroom or other processing time requirements on the data.

Penguin, as designed for the 2009 deployment at the South Pole, does not require any additional signal processing on the raw VLF data. For longer term studies, how-ever, several forms of signal processing may be desired. Most immediately, recording only a FFT of the data would likely dramatically decrease the storage requirements

and increase the available duty cycle for recording. Scientific requirements established the end users of the Penguin data indicate that a 2048-point, 16-bit FFT with 25% overlap would be sufficient to conduct the required geophysical research. Such an FFT engine, including Hanning and Hamming window options, was created and tested on the Penguin system. The engine consisted of a Verilog-defined FFT function, a Verilog-defined windowing look-up table for the Hanning or Hamming filters, as well as interface logic to the processor's memory bus. The main CPU was re-compiled to include a Direct Memory Access (DMA) scatter-gather controller to allow streaming of the data from the CPU's memory to the FFT engine and back, without the continual oversight of the processor. Lab testing demonstrated that with a 50 MHz clock the FFT engine was able to operate in real-time alongside the processor's sampling operations. A further, but small, increase in processing speed is required to also simultaneously stream the processed data to non-volatile storage. The overall power implications of the increased clock frequency and added logic required for the FFT engine requires further application-specific study.

Storage of the recoded data, processed or not, requires the CPU to read and interpret the file system, locate a blank location to create the file, and transfer the actual data to the storage medium. CompactFlash is natively accessed through a 16-bit interface, and thus the CPU should have at least a 16-bit wide data path for maximum efficiency [5]. The file system used in Penguin is REVFS, a slotted file system used on both the Buoy systems (Section 2.3.3) as well as several hand-held instruments. This file system is optimized for speed and low memory footprint. Looking at the processing requirements, accessing a file location requires addition of 16-bit numbers as well as multiplication of a 8 and 16-bit number, with the result guaranteed to be confined to a 32-bit unsigned integer. Creating and updating file information also requires calculation of a checksum, necessitating bit shifts and XOR calculations. Aside from these operations, no complex math is required. Clock rates

for the storage operations are unbounded at the low end, and are capped due to our card's measured access times at approximately 30 MHz.

As a last constraint on the processor speed, the peripherals must be examined. In particular, the Synchronous Dynamic RAM used to provide massive volatile storage for buffering and future processing of the recorded data must be refreshed with regularity. The processor is designed in firmware to include a SDRAM controller to offload the refresh responsibilities from the CPU, however both the controller and the CPU should be synchronously clocked to avoid unnecessary complexities of clock boundary crossings in the main system memory path. As a practical matter, the SDRAM controller should be clocked no slower than 40 MHz to achieve reasonable performance.

Surveying all of the above computational requirements, the CPU implementations can be easily defined. A minimum data width of 16-bits is required to handle the raw sampled data as well as the CompactFlash operations. 32-bits is not strictly required, but will speed file system calculations. Aside from the most basic math, a 16-bit by 16-bit multiply is necessary, as well as bit shifting and XOR. No division or other advanced operators are required. For memory access, an independent SDRAM controller is required, and at least one SPI controller is required to interface with the ADCs. General Purpose IO is sufficient for triggering and clock maintenance operations, as well as a single Universal Synchronous/Asynchronous Receiver Transmitter (USART) to interface with the GPS. Lastly, CPU provisions are desired to allow the addition of a DMA controller in future campaigns to support real-time FFT and other signal processing operations.

For ease of real-world implementation, the Altera NIOS-II 32-bit CPU is selected for use in Penguin. The NIOS-II is a software defined processor which is tightly integrated with the Quartus II and System On Programable Chip (SOPC)

builder for firmware implementation, as well as coupled with an Eclipse-based Integrated Development Environment for software development. The processor has a 32-bit wide data bus, available peripheral modules for many types of communications including SPI and USART, as well as selectable core modules allowing its footprint to be minimized as much as possible. For Penguin, the NIOS-II is synthesized with hardware multiplication, no instruction or data cache required (however optional for increased computational performance), a SDRAM controller, three SPI controllers, two CompactFlash controllers, several GPIO / Parallel IO lines, and two USARTS (one for debug and one for GPS). Together this hardware occupies less than 25% of the available FPGA resources, allowing space for signal processing firmware. Additional firmware space is occupied with the requisite gating circuitry between the CPU GPIO lines and the high-performance external sampling clocks. This additional logic is extremely small and consists of only flip-flops and NAND gates.

3.5.2 System Software

With the firmware well defined, the main processing software is developed to control all of the peripherals and ensure proper data capture and transport. On power-up, the software first mounts the CompactFlash cards to allow for log messages to be recorded as necessary. If a drive is not able to be mounted, whether due to a corrupted file system or lack of physical card in a particular slot, the system simply skips that disk and continues with the remaining good drives. In the event that no disks are mountable, the system will reset itself and sleep for 15 minutes. This failsafe mechanism prevents the system from rapidly depleting its battery in the event of a transient failure, however allows for the system to resume operation should the fault be recovered.

With the file system mounted and logging, the system must next check the

time information. First, the main processor sends a SPI request for the date and time of last GPS synchronization to the system control microcontroller. If the microcontroller replies back with all zeros or with a date before 2007 then a cold start has occurred and the clock is not currently synchronized; the system then immediately executes a GPS time synchronization. If the date received from the microcontroller is valid, the system next reads in the value of the time counter from the shift registers. This time counter increments each second, with zero being at the last GPS synchronization; simply adding this counter value to the time store in the microcontroller yields the current date and time. If the counter value is larger than a pre-defined value, 86,400 seconds or 24 hours in the 2009 deployed system, the system sets a flag to synchronize after the current sampling run and before the system shuts down. By synchronizing this stale time after the current sampling interval, the system avoids a potentially long GPS acquisition phase from causing the system to miss a data collection interval.

The time synchronization phase is a two step process: measure the current clock drift if the system is not waking from a cold start, and setting the clock to the current time. Both of these steps require a known true time standard, for which Penguin uses the Global Positioning System as discussed in Section 3.3.4. To synchronize, the main system processor first requests that the system control microcontroller activate the GPS. This functionality is relegated to the microcontroller to prevent glitching of the FPGA lines during power-up from causing the GPS to temporarily power on, resulting in accidental modification of the current time information. Once the GPS is powered up, the main processor sends initialization strings to the receiver, including enabling of UTC (versus GPS) time, and setting the elevation mask as necessary for the deployment location. The GPS receiver then returns diagnostic information once per second indicating the number of GPS satellites locked in the correlators as well as the type of position fix, if any. Once the position is fixed

and at least four satellites are locked, the system then queries the GPS receiver for the current UTC to GPS time offset. This offset is caused by the insertion of leap seconds on UTC which are not accounted for in GPS time. The offset message is repeated once per GPS message, and thus may take up to 12 minutes to receive worst case. Once the offset is properly set, the timing information from the GPS is known to be good and the synchronization task may continue.

If the system is not synchronizing from a cold start, then time error information of the internal clock is measured and logged. To accomplish this, the processor first updates the current system time from the microcontroller's record of the time of last synchronization plus the tick count in the external timekeeping registers, which are read via shift registers. It is important to note from the electrical schematics that the system 1 PPS line is automatically re-synchronized with the GPS 1 PPS when available. Thus, by latching the timekeeping information into the shift registers on the edge of a 1 PPS trigger, the read value is the Penguin-estimated current time to the nearest clock tick, or 10 us as implemented in the deployed Penguin system, while the exact time of the latching clock edge is an integer second boundary of GPS time which is read via the USART connection. This clever auto-synchronization of the 1 PPS generator while maintaing constant internal clocking allows for arbitrarily accurate time error recording without the need for extremely fast processors, limited only by the clock rate of the internal timekeeping counters. This time error is logged to a timekeeping log file on all CompactFlash cards, and may be used to post-process the recorded data to adjust for timing drifts.

With the timing error information recorded, if available, the internal clock is synchronized with GPS. This task is accomplished by the main processor arming the reset circuit for the internal timekeeping counters. With this circuit armed, the next GPS 1 PPS edge will reset the clock. The processor then waits for this first 1 PPS edge to occur, at which point the clock reset circuit is disarmed to prevent

re-resetting on the next second. The GPS is then queried to receive the exact date and time, which is correlated with the previous 1 PPS edge that cleared Penguin's clock. This time value is then transmitted to the system control microcontroller via SPI and recorded in the log file, and the GPS receiver is turned off. Note that in the Penugin system the internal clock cannot be disciplined or tuned. This design trade-off results in approximately 100 PPM time stability, or slightly better when buried in the snow which effectively serves as an ovenized crystal oscillator, while using significantly less power than a Voltage Controller Oscillator or other tunable frequency source. The overall drift of the timekeeping system is bounded by the frequency of the time synchronizations, assuming no significant thermal change is present, and is dictated by scientific requirements.

With a cold-start clock synchronized, or the flag set to run the above synchronization process after the next sampling run, the system continues through its startup sequence. The current time, as read from the microcontroller and shift registers, is checked against the pre-programmed recording schedule. This schedule is implemented as an offset from the top of the hour plus a sampling period, and is thus designed for schedules that have a mode at one hour. Penguin as deployed is specified to record at :05, :20, :35, and :50 of each hour, which corresponds to other VLF receivers across the Antarctic Plateau. In the scheduling software, this is defined as an offset of 5 minutes and a period of 15 minutes. This simple schedule format meets the requirements for data collection and obviates the need for a more complex scheduling algorithm, resulting in less computational overhead and less code volume for potential implementation bugs.

When the current time is compared against the schdule, if there are more than 30 seconds before the next scheduled recording run the system beings the shutdown procedure. Otherwise, with less than 30 seconds left before the sampling time, the system watches the current time counters at the second boundary until exactly

one second is left to begin recording. At this point, the system arms the sampling system which allows the externally generated sampling clock to pass through the FPGA and into the ADCs. At this point, the ADCs are powered up and sampling, but the sampled data is discarded allowing for stabilization of the ADCs. The next 1 PPS signal received by the main system processor indicates that the recording period has started and that the ADC values must be recorded. At this point, an Interrupt Service Routine (ISR) is enabled. This ISR is activated after the maximum conversion time of the ADCs after each sampling clock pulse. The ISR is then responsible for initiating the SPI transfers of the data from each ADC and placing the received values into a RAM buffer. Each trigger of the ISR also increments a counter, which is compared against the desired recording time. Once the total number of ISR actuations, which equals the total number of recorded samples, equals the target number of samples, the ISR is uninstalled and the sampling system is disarmed, causing the ADCs to stop sampling.

With the proper number of samples in a volatile SDRAM buffer, the next step is to post process the data. In the deployed Penguin system, no post-processing is necessary and thus this step is skipped. If FFT processing is desired, however, this is the appropriate time to do so. Assuming a reference FFT implementation as discussed in Section 3.5.1, the first step is to configure the scatter-gather DMA controller. This configuration instructs the controller as to which memory addresses are the beginning and end of each channel's data set, as well as the memory address for both the input and output of the FFT engine. Once configured, the DMA is activated and parallel processing of other tasks may commence, or the main processor may wait for the DMA controller to indicate that the processing is complete. It is important to note that this tightly coupled FFT peripheral is able to process the data without CPU intervention once the DMA controller is configured, allowing the CPU to also go into a low powered sleep mode if desired. Once the post-processing

is complete, flags indicating the start and stop address of each data set should be updated before continuing to the storage step.

With the raw data buffer and with the appropriate, if any, post-processing complete, the system is ready to record the data to non-volatile storage. With multiple CompactFlash cards installed in the system, the decision was made to interleave the recording periods on different cards. By having the :05 and :35 periods on card 0 while the :20 and :50 periods are on card 1, for example, the loss of a CompactFlash card would still result in an entire year survey of VLF activity, however with half of the density as desired. The alternative option of recording onto the first card until full, and the moving to the next card, would result in a half year of data and a half year of nothing. Scientifically, this latter option is less desirable. Thus, the system compares the recording time with the above constants to determine which drive the data should be stored onto. Once determined, the system checks the mount table to see which descriptor points to that physical drive. If one is not found, such as may occur with a physically damaged or missing card, the data is discarded and the system begins shutdown procedures. Otherwise, the system creates a new file on the selected drive for each channel of data, and iterates through each data set recording every sample to the proper file. Once recorded, the system closes the files, flushing the write cache and synchronizing the file journal, and begins shutdown procedures.

The data is buffered to the SDRAM during the sampling phase and not directly streamed to disk for several reasons. First, by having the data in a volatile buffer it may be more easily transported to any post-processing peripherals or engines as desired without increasing the wear on the non-volatile flash media. Additionally, the CompactFlash drives consume significantly more current when writing data than when idle, resulting in possible fluctuations on the system power rails. These fluctuations could appear as shifts in the sampled ADC values, causing the power rail noise from the CompactFlash cards to be directly imposed on the VLF data.

The severity or detectability of this phenomena is one topic of possible future work. In Penguin, the design decision was made to buffer the data in the SDRAM to avoid any possibility of data contamination.

The first step of the shutdown procedure is to check the time synchronization flag. This flag is set during the initial timekeeping check at power on if and only if the time since last synchronization exceeds the desired resynchronization interval, or 1 day for the deployed Penguin system. Again, if valid but stale time is available at power on, it is desirable to record the data on the schedule first and then synchronize the time to avoid missing a recording period. If this flag is set, the same synchronization procedure as perviously described is called. Otherwise, the system is ready to sleep. Before powering off, the system first reads the current time from the microcontroller plus external counters, and then calculates the number of minutes and second that remain before the next sampling period. Ten seconds are subtracted from this number to allow time for the system to power on and execute the boot-up commands. This value is then recorded in the log file on all CompactFlash cards as the desired sleep interval. All filesystem buffers are then flushed, and the cards are unmounted. Lastly, the processor sends a sleep command to the micrcontroller, indicating the desired sleep interval. The microcontroller immediately powers off the FPGA, sets and alarm for the desired time, and itself enters a low powered sleep mode.

3.5.3   Microcontroller Software

The system control microcontroller is tasked with several basic but critical functions: waking and sleeping the main system as appropriate, powering on and off the GPS receiver, and maintaining information about the last GPS time synchronization. Since the microcontroller is always running, optimizing its power consumption has a dramatic effect on the total system power performance. Minimizing the required clocking rate, utilizing hardware-based low power sleeps modes, and optimizing the
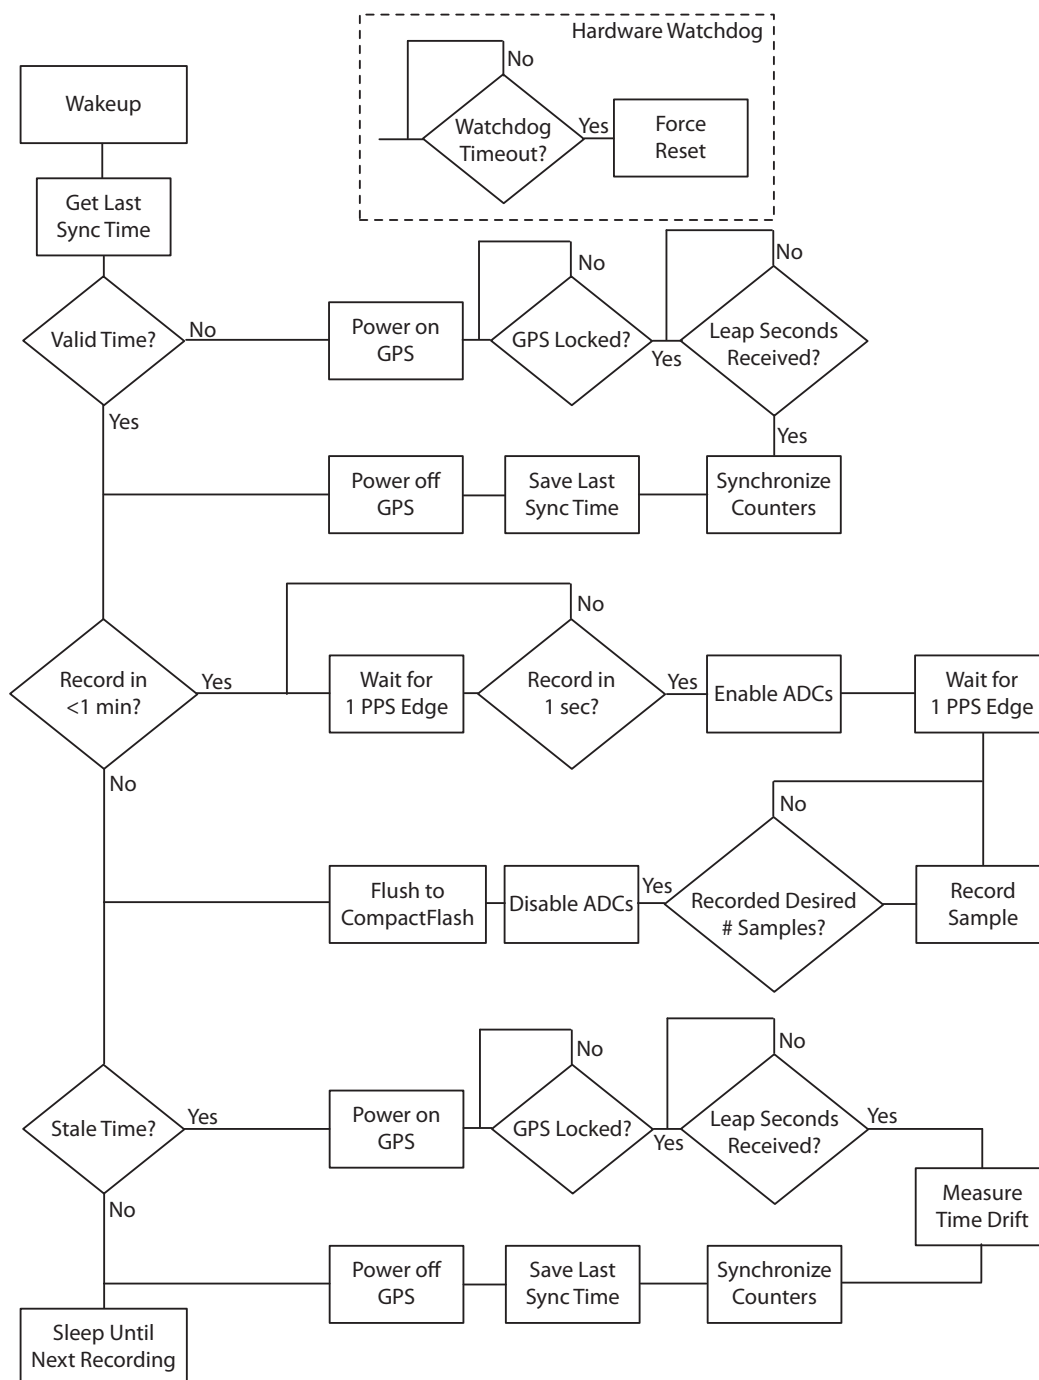
Figure 3.3: Penguin Software Flowchart

software all contribute to this goal. As discussed in Section 3.3.7, the system control microcontroller is clocked directly from the 100 kHz sampling clock source, rather than the 2 MHz frequency standard or its own high frequency internal RC oscillator. With such a slow internal clock rate, by modern standards, the active current of the microcontroller is reduced to only several hundred microamps [15]. Software design is able to reduce this by more than an additional 90%.

On initial power-up, the microcontroller allocates and zeros registers of the date and time of the last GPS synchronization and allocates space for an alarm clock timer. Next, the microcontroller updates its own watchdog timer. The watchdog for the microcontroller ensures that upon a software failure or soft-error, the Penguin system is rebooted to attempt to recover from the error. This watchdog is enabled by a fuse, set at the time of system build, and is not maskable by software. The watchdog will cause a hard system reset if it is not updated every 4 seconds by a keep-alive function. The software is designed to touch this keep-alive indicator nominally once per second, depending on the current location in the state machine.

Continuing the power-up sequence, the microcontroller then ensures that the GPS receiver is powered off, clears its SPI communications channel with the main system processor, and then powers up the main system for time synchronization. Every time the main system processor is powered on, a software-defined kill timer on the microcontroller is activated. This kill timer is implemented as a counter on the microcontroller that is incremented once per second, and cleared upon power on of the main system processor only. The kill timer is only active while the main system processor is active. If the kill timer counter exceeds 30 minutes, the main system processor is hard powered off for 30 seconds, and then re-activated with the kill timer reset to zero. The main system software is designed to never be active for more than 14 consecutive minutes unless GPS acquisition is delayed. The kill timer ensures that main system software failures, soft-errors, or other malfunctions

will not cause the system to remain on in a frozen state, rapidly depleting available power.

While the main system processor is powered on, the microcontroller may enter a sleep state to conserve power, however the SPI receiver must always remain on and active to accept commands from the main system processor. During this time, the microcontroller may receive requests to get or set the last time the unit was GPS synchronized, to enable or disable the GPS power, or to shut down the system for a defined period of time. All such requests are submitted over SPI to the microcontroller in a message packet consisting of a start byte, payload, and stop byte. The start and stop bytes are used to recover packet alignment should the microcontroller miss a byte, and also allow either side in the transaction to ignore spurious signals from power on glitches of the FPGA. All requests originate from the main system processor, and are acknowledged by the microcontroller for robustness.

Typically on power-up of the main system processor, as outlined in Figure 3.3, the microcontroller will receive a start byte followed by the request for the time of last GPS sync. Upon receipt of this request, the microcontroller replies back with the contents of the last sync registers, containing the year since 2000, month, day, hour, minute, and second bytes, terminated by the stop byte. If these values have not been previously set by the main system processor, they default to all zeros.

If the main system needs to perform a GPS (re)sync, the microcontroller will next receive the command to power on the GPS receiver. This command enables the GPS voltage regulator, and then replies back with an acknowledge to indicate that the command has successfully completed. During the synchronization, the microcontroller will then be given new time information of the current GPS synchronization. This new information overwrites the contents of the time registers, after which an acknowledge packet is again sent. As the last phase of the synchronization process, the main system will request that the GPS power is disabled. The

microcontroller shuts down the associated voltage regulator, and again replies with an acknowledge back to the main system processor.

Once the main system processor has completed its tasking, a request is sent over the SPI bus to power off the main system until the next scheduled action. The power-down request packet consists of the start byte, power-down request byte, number of minutes to power-down for, number of seconds to power-down, and finally the stop byte. This is the only command that is not acknowledged by the micrcontroller. Upon receipt of the command, the received values are checked against pre-deployment set maximums, set as 15 minutes for the South Pole deployment. If the requested time values exceed these maximum values, they are truncated to the maximal permissible extent. The microcontroller then sets an internal alarm clock timer to equal the total number of seconds of the sleep request. All peripherals, including the SPI controller, are disabled and powered off except for the digital input interrupt components.

While the majority of the Penguin system spends a significant time sleeping, the microcontroller itself must remain on to determine when to wake the rest of the system to perform data collection activities. The microcontroller is aided by monitoring the Penguin-generated 1 pulse per second line via an interrupt-enabled input pin.Due to the selected AVR architecture, the microcontroller is able to enter a low power sleep state with most of its peripherals powered down while still being able to wake on each 1 PPS rising (or falling) edge. At each 1 PPS tick, the microcontroller wakes from the low-powered sleep state, increments the alarm clock counter, and checks to see if the alarm clock counter equals the alarm clock set point. If not, the system returns to the low-powered sleep mode until the next 1 PPS signal, yielding significantly less than a 10% duty cycle during sleep. If the alarm time has been met, the system powers up the main system by enabling its voltage regulator, and the process repeats.

3.6   System Performance

With the Penguin system fully defined and constructed, both laboratory and real-world data regarding the system performance is gathered. Of particular interest is verification of scheduling and sampling clock, data integrity, and power performance. The low-level details of the RF front end are the topic of a separate research effort outside the scope of this document.

Data integrity from Penguin was verified against the AWESOME VLF receiver, the current gold-standard in use by Stanford University. This comparison was made at the Lockheed Martin Santa Cruz Test Facility, a nearby test site with minimal VLF background noise plus a calibrated AWESOME receiver site. Several hours of data was collected simultaneously by both Penguin and AWESOME, with their antennas separated by several hundred feet to minimize coupling while maintaining a relatively equivalent signal and noise environment.

The data collected shows extremely good timing accuracy between the two systems, with Penguin drifting slightly over the course of the test, but by no more than 100 PPM. This drift is expected due to the AWESOME system providing continuous GPS discipline to the internal sampling oscillator, while the Penguin system only corrects the oscillator once per day to minimize power consumption. When calibrated, the Penguin system showed better sensitivity at the lower end of the VLF band than AWESOME, however the passive AAF in Penguin causes significant degradation in received signal strength above 30 kHz. Additionally, the second channel of Penguin was attenuated approximately 20 dB down from the first channel. This reduced signal strength was localized to the LNA and determined to be due to manufacturing variations in the LNA ASIC. As shown in Figure 3.4, the overall data quality of Penguin is excellent and generally comparable to the AWESOME system.
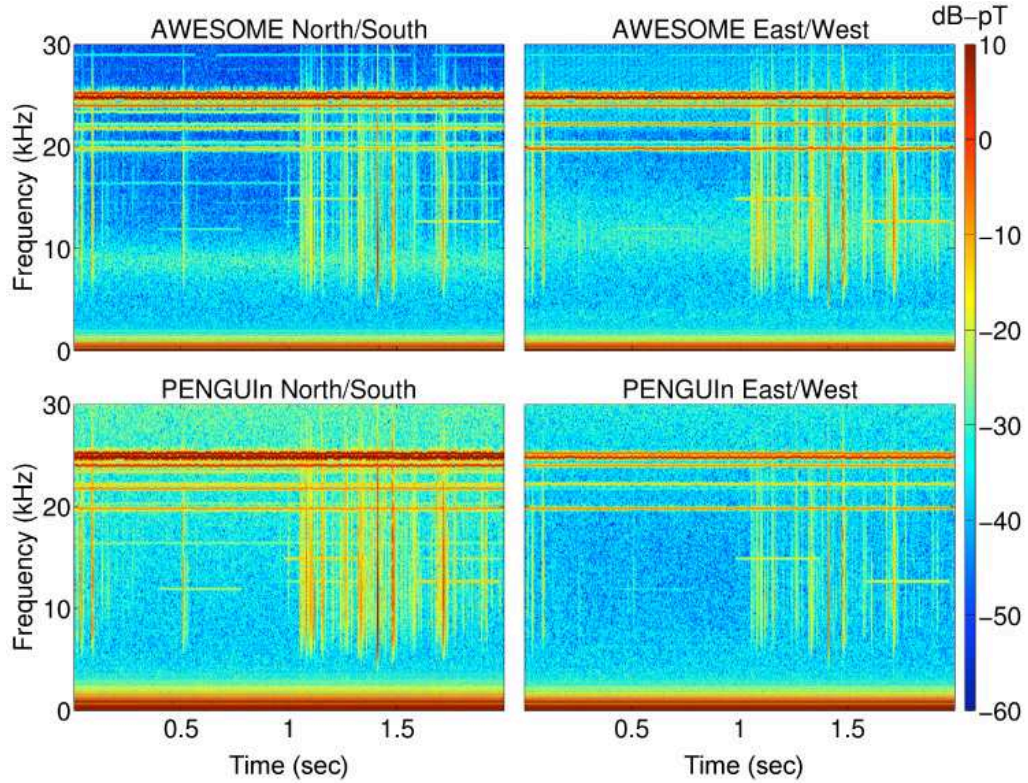
Figure 3.4: Penguin and AWESOME Simultaneous Data

The Penguin system was extensively tested both in a thermal-vacuum chamber as well as a nitrogen-flushed thermal chamber. The former was used initially due to its availability within our lab and ability to reach temperatures below -100° C with its liquid nitrogen cooling system. Individual components were verified during the early design stages in this chamber, with all selected components proving to be operational to -100° C throughout short duration testing of approximately 10 hours each.

For longer-term thermal testing, a separate nitrogen-flushed chamber was acquired due to the lower operating costs of its two-stage compressor chiller. The Penguin system as a whole was tested in this environment for approximately two months continuously to validate the long-term survivability of the system. Due to

compressor limitations, the long-term tests were conducted at the anticipated -55°
C snow temperature. No system failures of software glitches were observed at any
point during the final two-month test.
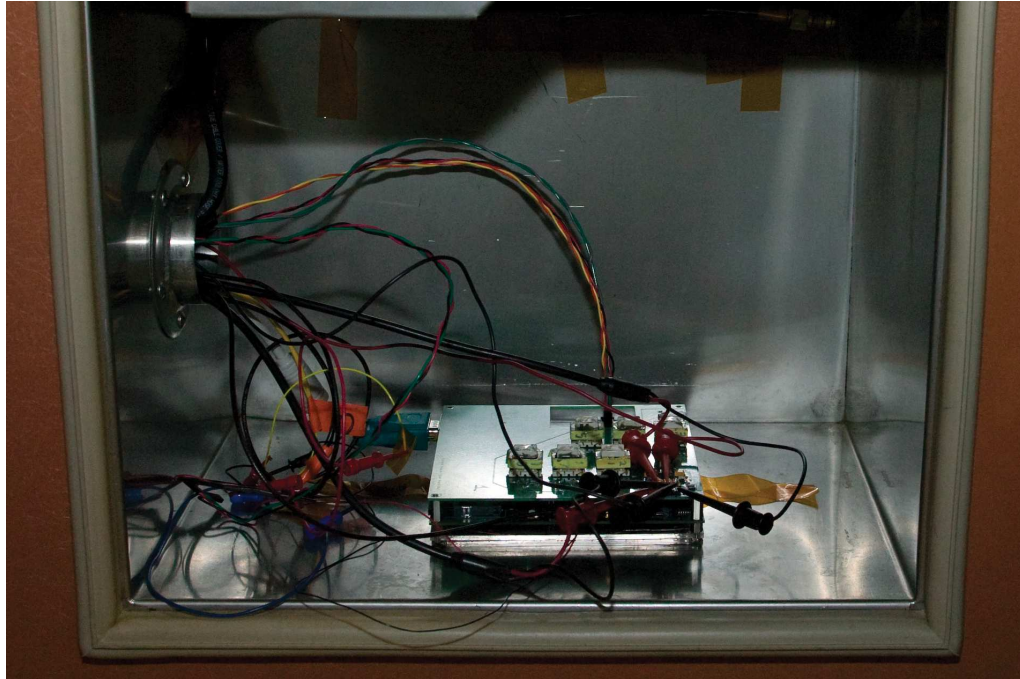
Figure 3.5: Penguin System in Nitrogen-Flushed Thermal Chamber

Initial testing of various sub-assemblies, such as the FPGA and associated mem-
ories, showed a significant temperature dependance of the consumed power. The
baseline room-temperature power draw would often more than double as the tem-
perature in the test chamber dropped. Furthermore, the draw remained high even
as the chamber temperature was returned to room temperature. Upon removal of
the test boards, signs of corrosion and water damage were present. All future tests
were made either under vacuum or in a dry-Nitrogen environment to avoid conden-
sation issues. In these dry conditions, no detectable changes to operating current
were observed on individual components as well as the entire integrated system.

In the laboratory, measurements of the subsystem power requirements were made

at temperature. The LNA plus related support hardware, located at the antenna, was measured to draw 6.4 mA with a 3.3 V supply. The AAF, including all buffering and internal voltage regulation, was measured to draw 6.6 mA with a 3.3 V supply. Both of these modules, plus the Penguin system as a whole, include internal linear voltage regulation. Because of the nature of linear regulators, the current consumption remains virtually fixed regardless of the input voltage, as long as the input voltage is within operating limits and does not significantly alter the thermal properties of the regulator. With a nominally 7.2 V battery rail on Penguin, these components were measured to consume the same current, and thus when combined require 93.6 mW of power while operating. Note, however, that these subsystems are in Power Domain 1, and as per Section 3.4 these subsystems are only powered on during sampling operations; thus, the average power consumed is significantly lower.

The digital side of the Penguin system was also independently characterized. Due to the bursty nature of the current draw and significant dependence of the draw on the processor load, measuring the power consumption with finalized software operating as a system is most accurate. Averaging over multiple runs, the entire integrated Penguin system, including LNA and AAF, was measured to consume and average of 59 mW as configured for deployment at the South Pole. The peak current observed was 180 mA bursts during write operations to the CompactFlash cards, with a nominal current of 160 mA while the processor was operating.

Significant enhancements to the software were made after the South Pole deployment. The new software allows the system to sleep longer, waking only 2-3 seconds prior to the scheduled sampling start. Additionally, an improved file seek operation was added, and the USART peripheral, utilized during debugging and to verify system operation, was removed in software, firmware, and hardware. By

combining these updates with a 3.6 V battery module, the Penguin system was measured to consume an average of only 18 mW to collect two-second recordings every 15 minutes, making Penguin the lowest power autonomous VLF synoptic recorder to date.

3.7   Deployment Data

The Penguin system was deployed in February, 2009 near the Amundsen-Scott South Pole Station. The deployed unit consisted of one insulated housing containing the Low Noise Amplifier, a second insulated housing for the Anti-Aliasing Filter, main system electronics, and battery module, as well as a 15 foot tall, 30 foot wide crossed-loop magnetic field antenna. A 250 foot spool of shielded cable separates the LNA from the main Penguin system, although only 30 feet of the cable is necessary as deployed. The additional cable remained on the spool to avoid changing the calibration and test validity of the system. Lastly, an active GPS antenna is connected to the Penguin system by an 8 foot coaxial cable.

The Penguin VLF antenna is located at 89.99353S, 98.03986E, at approximately 9314' elevation. The LNA is buried approximately 4 feet below the center of the antenna to provide better thermal stability of the electronics. The main electronics box is located 30 feet away from the antenna center, with a buried cable between the LNA and main box midway between the two antenna loops. The main Penguin system is located 6 feet below the surface, again for thermal stability, while the GPS antenna is several feet above the surface to provide adequate satellite visibility.

Approximately ten days of data were collected from the system during the initial deployment and operational verification. During this time, both the LNA and main system electronics were at the surface. The air temperature ranged from -40° C to -50° C during this time period. The data was collected at several logistically-feasible

Figure 3.6: Penguin Deployed at the South Pole

points in time to ensure proper system operation, survey the noise environment, and possibly collect scientifically-useful short-term data.

The initial 24 hours of captured data showed a slowly rising noise floor that eventually swamped out all local signals. This decrease in detected signal SNR was directly caused by the bias voltage of the LNA not tracking the physical LNA temperature; the time span of the change was due to the roughly 24 hours necessary for the initially room temperature inner insulated chamber to reach thermal equilibrium with the extremely cold outside air. Once thermal equilibrium was reached, the static bias voltage was adjusted to match the laboratory-tested values for the given ambient temperature. Once properly set, the anticipated noise floor and local signal SNR was restored. This bias voltage was subsequently re-adjusted to match the buried snow temperature prior to the final deployment. Example data with improper LNA biasing is shown in Section 5.3.

On February 3, 2009 beginning at approximately 2130 hrs UTC, the Penguin system recorded a period of natural Auroral Hiss [39]. This event, lasting nearly three hours, is the first in-situ scientific data recorded by Penguin, and was also simultaneously detected at various other Antarctic ELF/VLF receiver locations, demonstrating Penguin's capacity to gather scientifically useful data.
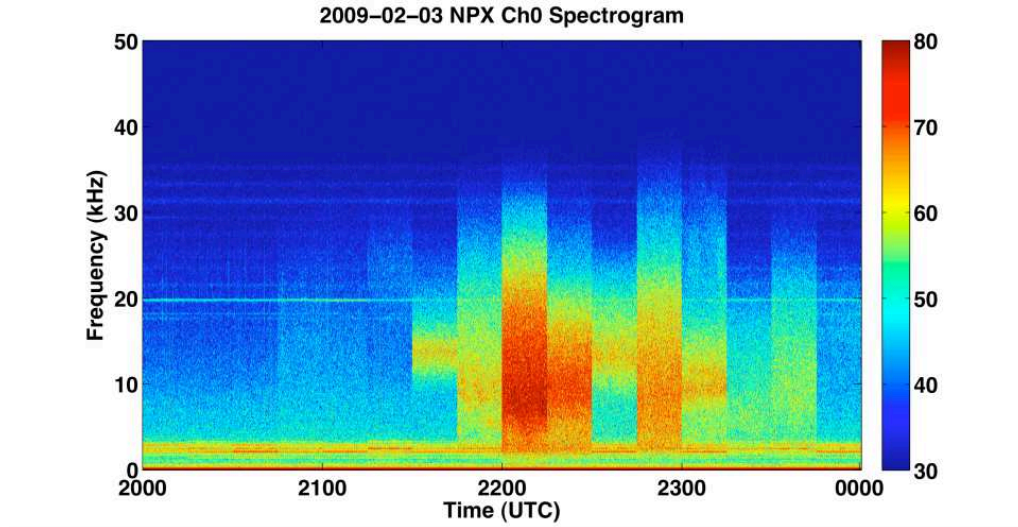
Figure 3.7: Penguin Observation of Auroral Hiss

3.8 Conclusion

The Penguin system described in this chapter pushed the developmental envelope of VLF receiver systems and is the first micro-power VLF scientific instrument deployed to the Antarctic. While maintaining a classical sampling and direct conversion architecture, the advanced digital systems designed and implemented greatly reduced the power requirements while still maintaining excellent timing resolution and received data quality. Removing the generally overpowered main processor and replacing it with the most efficient, application-specific CPU greatly reduced the on power. Additionally, the ultra-low power requirements of the system while in sleep

mode, and novel timekeeping system contributed significantly to the low average power of the system as a whole. With an ASIC VLF LNA, the requirement for high voltage rails was removed, simplifying the power system to 3.3 V and 1.8 V only.

The presented architecture is flexible enough to be easily programmed for different synoptic sampling periods and duty cycles. The system can also be upgraded to continuous GPS time disciplining for a power penalty of a few hundred milliwatts of on-power, allowing for true interferrometric measurements where necessary, above the current Penguin System assumption of linear clock drift and requirement of post-processing the timing data.

The end result of the extraordinarily low average system power requirements, high data fidelity, and low system mass and volume of Penguin enable scientific activities that were previously not possible or simply too costly to perform. With a field tested system proven at the South Pole, Penguin allows for autonomous, extremely remote, VLF synoptic data collection in one of the most harsh and extreme environments on the planet. The low system mass makes the unit easy for transportation in small aircraft to remote areas of the Antarctic Plateau, or anywhere else where transportation opportunities are limited.

Chapter 4

VLF Advanced Technology Receiver

4.1 Background

The Penguin VLF receiver has been proven to withstand extreme environmental conditions while maintaining its ability to collect high-quality broadband VLF data. Penguin is able to record synoptic broadband data at extremely low average power draw, and even if modified to sample and record continuously, is able to do so with approximately a 15x reduction in power over the next closest existing system, the South Pacific Buoys. Penguin, however, is still limited by the constraints of a traditional sampling system with a general purpose CPU in the always-on while recording sampling loop. For synoptic sampling, this architecture may allow for sufficiently low system power, however further improvements are possible with the introduction of a radically new sampling architecture.

A novel architecture introduced herein which removes the power-hungry CPU from the sampling reduces the power requirements of a continuous time recording

system by more than a factor of 4 compared to the optimized traditional architecture of Penguin when operated at 100 kHz sampling rate, and with linear gains as the sampling rate is decreased. The VLF Advanced Technology Receiver (VAT) is the result of combining available new programable logic technologies with a revolutionary new system architecture. This new system architecture directly addresses the longstanding problem of continuous sampling for extended durations in harsh environments. The architecture was simulated in parts to estimate the performance gains anticipated, however the total improvement in power and performance could only be determined by building a real system and measuring the performance. The inability to fully simulate the architecture stems from the change in performance expected from each unique commodity IC when driven by the custom logic, and the complexity of the interactions between the numerous external devices making it prohibitively difficult to fully model in current hardware description languages.

4.2 Derivation of Requirements

As with every scientific instrument, the initial design step is to analyze the scientific requirements of the data supplied. VAT was designed to study lower frequency signals such as chorus and HAARP emissions, typically contained within the lower 4 kHz of the ELF/VLF band.

From the above signals of interest, the scientific requirements dictate that the system have a usable bandwidth of at least 4 kHz and a dynamic range of at least 40 dB. Additionally, it is expected that signals within the next few octaves above the desired bandwidth which may alias down will have a relative power of approximately 0 dBc. Combining the maximum interfering signal power with the expected signal strength indicates that the AAF must have at least 80 dB of stopband rejection. Additionally, the 4 kHz maximum frequency requirement forces the maximum

unaliased signal to also be at 4 kHz.

The existing LNA has been thoroughly examined and determined to have not more than 80 dB SFDR at any point within the VLF band. Future LNAs used on the system may provide close to 90 dB of SFDR, but it is not expected that this barrier will be broken in the foreseeable future, and there is little scientific need for a larger dynamic range due to the background atmospheric noise at most receiving sites.

All known previous autonomous receivers used for scientific research of ELF/VLF phenomena at high L-shells (near the geomagnetic poles) have only collected a brief snapshot of the desired frequency band, similar to Penguin, or continuous channelized data, similar to AGO (Section 2.3.2). The synoptic mode cuts down on data storage requirements and potentially power, depending on the system architecture. To further enhance studies, it would be desirable to have continuous data collection over the entire bandwidth of interest. The continuing advances in solid-state storage technology make this feasible for moderately long campaigns at the date of publication when sampling two channels at 16bits/channel and 10 kHz sample rate. The total required data rate is thus given by Equation 4.1.

$$\frac{10,000 \text{ samples}}{\text{second} \times \text{channel}} \times 2 \text{ channels} \times \frac{2 \text{ bytes}}{\text{sample}} = 40,000 \frac{\text{bytes}}{\text{second}} \qquad (4.1)$$

At 40 KBps data rate, 128 GB of flash storage would allow for approximately 37 days of continuous operations. Operational time can also be easily extended by adding additional flash cards to the system, or duplicating the system for each channel instead of having one system recording both channels. Increasing the sampling rate to 100 kHz allowing full broadband recording would proportionally affect the system runtime. With the nominal 128 GB of flash storage, a two channel system sampling

at 100 kHz would run for nearly 3 days and 16 hours before needing new flash cards. Given the current COTS availability of 128 GB flash cards, operating at a 40 KBps data rate would easily allow for one year of autonomous operation with only 10 cards in parallel. Adding multiple cards to the system can be done with nearly zero increase in power by simply connecting all of the cards in parallel and gating the power and signal lines to each card independently with a properly sized MOSFET. Additionally, consumer flash memory devices continue to increase in capacity, and thus a 10x increase in storage density may be available in the near future for a minimal change in power.

4.3 System Architecture

4.3.1 Background

One of the main objectives of the VAT architecture is to minimize average power consumption. All known previous systems rely on the main system processor, which typically operates the file system or communications and controls signal processing systems, to be running during data collection. For example, the Buoy 1.5 system contains a Texas Instruments DSP which is responsible for all system operations except for sleep timing. This DSP nominally consumes approximately 5 Watts of power regardless of whether the system is sampling VLF data, waiting for GPS lock, or performing computationally expensive signal processing routines.

In the Penguin system, the same FPGA and FPGA configuration is used during sampling, signal processing, file system access, and timing re-synchronization, which nominally consumes 600mW. Penguin achieves its incredibly low average power consumption by waking up only for short bursts to quickly sample and record, while otherwise consuming very little power for timekeeping and sleep scheduling. As the

recording duty cycle of Penguin approaches 100%, the average power consumption will also approach 600mW, as can be derived from the data in Section 3.4.

The key understandings contributing to the VAT architecture are that 1) digital systems can wake up and sleep very quickly, typically within a few milliseconds, with minimal start-up and shutdown overhead when properly designed, 2) sampling VLF data, even at full broadband, is a relatively slow operation compared to current processor and memory bandwidths, and 3) there is no requirement that sampling, signal processing, storage, and scheduling be performed on the same hardware or within the same power domain. By combining these three critical pieces of information together it becomes obvious that the classical low-power scheduling processor and single higher-powered sampling/processing/storage processor system design is suboptimal.

To minimize power, a unique and minimally powerful device should be considered for each separable system function. The analysis surrounding the consideration of a unique component for any given system function should also examine whether or not two or more functions should be merged into a single unit for pipeline processing. For example, a decimation filter and narrowband demodulator might be merged into a single unit if the system will only be doing narrowband demodulation if and only if the system is also doing decimation. Another common example of pipelined functionality would be a windowing function and associated Fourier transform or FFT. For the first incarnation of VAT presented here, the sampling system, signal processing, and file storage were examined.

4.3.2   Sampling System

The sampling system provides the greatest potential for improvement from a traditional system. For example, the Penguin Central Processing Unit operates at 50MHz, while the ADCs only operate at 100 kHz. Let's examine an ideal system

where the CPU speed is matched by an integer multiplier to the sampling clock. In such a system only 4 traditional machine instructions are necessary per sample: read channel 0 data, store channel 0 data, read channel 1 data, and store channel 1 data. For a sample rate of 100 kHz, a single-cycle CPU should only need to be clocked at 400 kHz under these ideal conditions. For Penguin, this means the CPU is running 125 times faster than necessary, and this extra speed is just burning power. Similarly, the Buoy CPU was running 3-4 times faster than Penguin resulting in up to 500 times overclocking, and the CPU of a PC used in the AWESOME system is running nearly 10,000 times faster than the necessary rate.

In addition to the faster than necessary processing clock rates, the act of retrieving a sample from each ADC and buffering it into some type of memory is a rather trivial task, computationally speaking. In addition to the load and store operations outlined in the above ideal system, depending on the type of memory a counter may also be required to keep track of where each sample should go. With a FIFO buffer, for example, no counter is needed, but with a commodity SRAM it is critical to increment the memory address for each subsequent sample to avoid overwriting previous data.

There are several important criteria for selecting the proper analog-to-digital converter. Most important for data integrity are the analog parameters, primarily sample rate, sample resolution in bits, Effective Number of Bits (ENOB), Spur Free Dynamic Range (SFDR), Signal-to-(Noise plus Disortion) or SINAD, differential versus single-ended, and rail span or operating voltage. After the minimum data integrity requirements are met, the system designer must the select an appropriate digital interface, and examine power consumption given the expected operating conditions.

As established in the derivation of requirements, the system should have at minimum a 4 kHz usable bandwidth after AAF roll-off and any detectable aliasing.

In an ideal world, the sampling rate, F_s, could be as low as 8 kHz to meet this requirement [35], but to allow for reasonable AAF roll-off we will sample at 10 kHz minimum. Thus, the ADC must be able to support a sampling rate $F_s = 10$ kHz, with it being desirable to also support full broadband sampling should the application so desire. As defined in Section 2.1.2, broadband operates with $F_s = 100$ kHz.

The sampling resolution, ENOB, and SINAD requirements on the ADC and directly derived from the analog front end of the system, namely the Low Noise Amplifier and Anti-Aliasing Filter. The ADC should ideally be selected so that the performance of the front end is the limiting factor for signal reception and data quality. As discussed in the requirements derivation section, we are quite comfortable with a maximum front end SFDR and SINAD of 90 dB. In reality, many of the front ends used for high-performance VLF receivers have closer to 80 dB SFDR when fully characterized [25, 42, 46]. Assuming 90 dB SINAD is sufficient from the front end ADC, using Equation 2.1 approximately 14.6 bits are necessary. Commercially available ADCs typically have 12, 14, 16 or more bits of resolution, thus a 16-bit ADC is minimally required. When selecting the ADC, it is important to make sure that the 16 bit or greater ADC also meets the 90 dB SINAD requirement (Section 2.1.4).

Consideration must be given to the analog input specifications for the ADC in addition to the performance of the actual analog-to-digital conversion. Input signals can be delivered either single ended, where there is one signal line referenced to some ground, or differential, where there are two signal inputs that are opposite of each other. Differential analog signaling provides increased signal integrity over single ended signaling due to its intrinsic common-mode rejection [44]. For signals traveling over an extremely short distance in a relatively benign noise environment, one would not expect much signal interference, and thus a single ended ADC may

be sufficient. However, when the analog signal is traveling over a long distance and/or when strong interferers may be present, the use of a differential ADC is indicated. Because the systems being discussed are designed to be connected to an arbitrary LNA, differential signaling was selected in order to offer the best possible noise rejection for a wide range of operating conditions.

Once a signaling mode has been selected, the actual voltage range of the input signal should be matched to the ADC. There is significant flexibility of the output signal range from the LNA and AAF combination. The LNA itself typically has bias setting resistors that can be adjusted for small-scale variations in the gain; however the LNA typically sets the noise figure for the system and thus its gain is rarely reduced [47]. The AAF plays a less important role in setting system noise figure and is thus most readily modifiable to match the ADC. In the Penguin system, for example, approximately 10 dB of gain was present in the AAF to best match the low level LNA output to the 2.5v rail of the ADC. A higher rail on the ADC will allow for more voltage noise in raw volts before the signal integrity is compromised to the same extent as would occur with a lower voltage ADC, however the actual input dynamic range and SINAD is typically the same. Minimizing the ADC rails will allow for lower voltage system operation, thus requiring fewer batteries or the ability to avoid boost switching regulators which may increase overall system noise due to their high frequency high current switching. Experimental and real-world test results, including the Penguin system, have shown that a 2.5V ADC analog rail is sufficient to have no noticeable affect on scientific data quality while still being low enough voltage compared with the rest of the digital system to not require additional battery cells or switch mode regulators.

With the analog domain requirements established, all that remains is to select a digital interface. Peripherals can interconnect with the CPU or other ADC controlling device though serial or parallel means. The most common serial interconnects

include the Inter-Integrated Circuit (I2C) protocol or Serial Peripheral Interface. I2C officially only supports a clock rate of 400 kHz [53], and we have previously established in this section that the ADC should be able to sample at 100 kHz with a sample resolution of at least 16 bit, requiring 1.6Mbps of bandwidth per ADC. Thus, I2C is not sufficiently fast to transfer our VLF data. SPI, on the other hand, can run at tens of megahertz and is more than sufficient for our needs [31]. Both of the serial methods described can utilize only one physical wire for the data from the ADC plus one wire for clock information. SPI typically also uses one wire for chip select, although with clever manipulation of the clock signal this can be avoided for some system designs.

Parallel ADC interfaces require many more data lines between the controller and ADC, typically one for each bit of data plus at least a read and chip select. The advantage of parallel over serial interfaces is that the controller can acquire an entire 16-bit sample from the ADC in one clock cycle, while with a serial connected ADC a specialized I2C or SPI controller must run at greater than 16x the desired data copy rate. Parallel data interfaces do require more attention to signal integrity due to the number of lines available to be interfered with, however this is an already addressed issue especially at such slow clock rates [24]. For the VAT demonstration platform, a parallel interfaced ADC is desired to minimize the number of clock domains required inside of the controller and allow for a single 16-bit wide data path to be used throughout the system architecture.

The above guidelines create a subset of the available ADCs on the market which can meet our stringent analog performance requirements and are also most easily integrated into the system architecture being discussed. As long as the above requirements are met, a system should be able to be designed that will perform up to the specifications derived in Section 4.2. Within these limitations, power consumption, component availability, and form factor can all be examined. Since the

underlying principle of the VAT architecture is to minimize power consumption, selecting the lowest power ADC possible while still meeting specification is the top priority. A switched capacitor successive approximation analog-to-digital converter provided the lowest operating power currently available on the market that also met all of the performance requirements specified. This ADC architecture also has the benefit of its power consumption directly scaling by the sampling frequency, thus a higher performance part than necessary can be used in the system at a lower clock rate yielding extremely low power consumption.

4.3.3 Processing and Storage Systems

For initial scientific uses of the system, it is minimally required that the bandwidth of VLF data being studied is directly recorded to some non-volatile medium for retention over the deployment lifetime and subsequent analysis after the campaign. Some of the other existing systems, such as the South Pacific Buoys and AWESOME VLF receiver, can also track and demodulate MSK-modulated narrowband VLF signals as well as record spectrograms and FFTs of raw broadband VLF data. These capabilities are not strictly required for the scientific goals outlined in Section 4.2, however designing an initial VAT system to be extensible enough to later support these features would be desired. All of these signal processing desires would be executed in the digital domain, and thus can be accomplished by arithmetic operations on a standard CPU's Arithmetic and Logic Unit (ALU), by specialized hardware in a Digital Signal Processor CPU (DSP), or by custom logic in a Field Programable Gate Array (FPGA).

A general purpose CPU, such as is used on the AWESOME receiver, is the least efficient in operations per watt than a customized processor designed for that one specific task would be [18, 40, 61]. Commodity CPUs are monetarily inexpensive to purchase, and typically have a very broad base of existing software available for

their use; however since they are designed to process any general type of information, optimizations for MSK demodulation or FFTs are not likely present in hardware. The South Pacific Buoys utilized a Texas Instruments DSP, specifically a TMS320 variant, which does have dedicated hardware facilities for some types of convolutions and Fourier transforms. Unfortunately, the input and output data formating, such as FFT length, did not match up with the scientific needs of the program and thus the DSP was used for its general purpose processor features without reaping the benefits of the lower cost operation-per-watt available from a DSP.

Because the mathematical analysis typically performed by researchers on VLF data is rather esoteric, having dedicated processing hardware would be ideal from a power and performance perspective. Unfortunately the cost in time and money necessary to create hardware to accomplish every desired analysis is prohibitive of this method. A programable logic fabric, such as an FPGA or CPLD, is a good compromise during system design because it allows for custom hardware functionality to be added to the system at a later time through relatively simple Verilog or VHDL code modifications with requiring new physical hardware. Additionally, the processing Verilog or VHDL can be designed by another engineer or scientist while the rest of the system is being built and tested, allowing for parallel workflows without time delaying interdependencies. Unfortunately, FPGAs are typically not as power efficient as an identical ASIC design, however the previously mentioned ASIC costs are too great to go any other way [40].

Selecting between a CPLD and FPGA for a given task is rather straight forward. CPLDs generally take significantly less power to run than an FPGA, although their gate delays tend to be higher. Additionally, the gate density and availability of memory cells and registers on CPLDs is typically much lower than that of FPGAs. Thus, FPGAs are generally better suited for complex processing tasks while CPLDs can handle more simple state machines, logic, and less memory-intensive arithmetic.

CPLDs do have the added advantage over FPGAs in power-on configuration. FP-GAs typically use SRAM-like cells to hold the logic fabric interconnect information, and thus must be reprogrammed by an external or embedded device at each power on. CPLDs are typically flash or fused based, and thus their configuration is non-volatile. The added component cost and board space of an FPGA configuration device is usually not prohibitive, but it is a requirement to be aware of.

In the continual effort to reduce power consumption, it is clear from above that it would be desired to use CPLDs everywhere possible. The relatively simplistic operation of controlling the ADCs as previously discussed in Section 4.3.1 is excellent functionality for custom logic in a ultra-low power CPLD. However, since the complex processing operations desired in future applications of the system will likely require significant amounts of memory, an FPGA is more likely to be able to handle the mathematic workload and is thus better suited for signal processing. These apparently conflicting requirements for sampling control and data processing which are typically contained in the same silicon are the precise targets for multiple power domains as discussed in Section 4.3.1. Thus, the system architecture should consist of at least one CPLD to control sampling operations and the ADCs, and at least one FPGA under separate power control for signal processing.

The Complex Programable Logic Device in the system design directly control the sampling operations and thus is connected directly to the analog-to-digital converters. The CPLD is also responsible for buffering this sampled data, and is thus also directly attached to some type of RAM buffer. Lastly, the CPLD will communicate the buffered data to the controlling, processing, and storage FPGA which also must be electrically adjacent.

Once processed, the data must be stored on a non-volatile medium for later retrieval. Commodity flash storage devices are the preferred method due to their extremely robust characteristics, low power compared to mechanical devices, and

extremely high storage densities. For the cold environments where the VAT and similar systems are deployed, several CompactFlash cards have been qualified by in-house testing for operation below -100C, which is colder than any known deployment location including the South Pole during winter [58].

Determination on whether the flash storage system should have its own dedicated controller for power efficiency requires an understanding of the data flows into, inside of, and out of the subsystem. Data that is to be stored in non-volatile memory will be sourced directly from the signal processing subsystem. If no signal processing is desired for the particular scientific application of the VAT receiver then the data to the storage subsystem will be sourced directly from the ADC controller(s) and related buffer(s). Inside of the storage subsystem, all processing required to understand the format of the filesystem used on the media is handled, including how to locate available file space and where logically and electrically to write the data. This processing typically will require a fair amount of memory depending on the filesystem selected, but typically does not required complex arithmetic operations. The storage subsystem outputs whatever data is to be stored plus the requisite control signals (read and written) directly to the storage medium, in this case a CompactFlash card. CompactFlash cards can typically operate at tens of megabytes per second and thus are a candidate for short burst writes of the comparatively low bandwidth VLF data, processed or otherwise [5].

The data stream coming from the processing section has been shown to be at least wireline speed when $F_s = 100$ kHz (Section 3.5.1). Since the VAT receiver will typically be operating at $F_s = 10$ kHz, the data available from the processing system will be at roughly 10x the wireline speed of the sampling system, and can also be burst transfered if desired.

Due to the relatively simple but potentially memory intensive computation required for file system operations, the storage subsystem is most appropriate for

either a general purpose CPU or possibly an FPGA. The actual overhead for such processing is typically minimal in relationship to the bandwidth that the data itself is recorded at, and can be greatly reduced by an application specific file system.

The Field Programable Gate Array thus is responsible for retrieving the data from the sampling subsystem contained in the CPLD and its associated components, processing it, and then storing the data in a nonvolatile memory. The FPGA must then be directly attached to the CPLD for memory read back, and will likely need its own external RAM for processing operations. Lastly, the nonvolatile storage, discussed below, must also be electrically adjacent.

4.3.4   Digital Hardware Architecture

The VAT system architecture is implemented as a demonstration platform that can be adapted to fit a number of end uses. The platform components are selected based upon the design and selection criteria discussed at length earlier in this chapter.

The front end of the digital system is the analog-to-digital converter. The ADC selection criteria are discussed in detail in Section 4.3.2. Utilizing these criteria, the Analog Devices AD7623 is selected for the VAT demonstration platform [19]. This ADC meets all performance requirements, is readily available at the time of writing from several suppliers, and is available in an easy to debug TQFP package. The power consumption of this unit is linear with respect to sampling rate, and when down-clocked to our 10-100 kHz sample rate is expected to consume only several mW of power.

In addition to the ADCs, some type of memory, although not necessarily non-volatile, is necessary to buffer the sampled data between wake periods of the main system. The three most common types of memory are Static Random Access Memory (SRAM), Dynamic Random Access Memory (DRAM) and flash. Both SRAM and DRAM are volatile memories, while flash is nonvolatile. Flash memories have

a limited amount of write-erase cycles, while SRAM and DRAM are practically un-limited. SRAMs typically are available with slower access times than DRAM but faster than flash. DRAM requires additional circuitry to refresh the memory on a periodic basis, and while it provides the fastest access time of the available memo-ries, it typically also consumes the most amount of power. Due to the high power draw of DRAM and the limited longevity and still relatively high power draw of flash, SRAM is the selected memory type for the VAT demonstration platform.

Since the VAT architecture works by continuously buffering samples into this SRAM buffer and only waking the higher powered FPGA as necessary to flush this buffer, it is most important that this memory be extremely low power. Additionally, since the FPGA takes a non-zero amount of time to wake up before it can begin processing and storing the previously buffered data, it is advantageous to buffer as much data as possible between the wake period of the FPGA. Cypress Semicon-ductor, at the time of writing, has one of the lowest power SRAMs available on the market with densities up to 4 Mbit readily available from manufacturers. The extremely low power, as low as 3mW in our own in-house testing, and high density make the CY62177DV30 an ideal candidate for the development platform design.

Now that the outlying peripheral devices are selected, the requirements for the heart of the system, a Complex Programable Logic Device and a Field Programable Gate Array, can be generated. Combining the information from the above discus-sion yields Table 4.3.4 with the necessary peripheral devices, the number of input and/or output lines that each peripheral requires, the target or core device that the peripheral is connected to, and the quantity of each peripheral that is included in the system.

At the heart of the design, a Complex Programable Logic Device is necessary for ultra-low power sampling control and data buffering. The CPLD must be large

Device	I/O Pins	Target	Quantity
AD7623 ADC	18	CPLD	2
Compact Flash	27	FPGA	2
CPLD-FPGA Interconnect	18	Both	1
Crystal Oscillator	1	CPLD	1
Crystal Oscillator	1	FPGA	1
CY62177 SRAM	43	CPLD	2
CY62177 SRAM	43	FPGA	1
GPS	3	FPGA	1
Power Control	1	CPLD	1

Table 4.1: Pin counts of VAT platform peripherals

enough, i.e. have enough Logic Elements, so that the entire compiled sampling controller described in Section 4.5 can be loaded. Additionally, it is of critical importance that this CPLD be as low power as possible, since it is one of very few components that will always be powered up during system operation. Product offerings from Altera are investigated due to our internal knowledge and skills using their devices. The Max IIz series of CPLDs offers the lowest power dissipation of any currently available device, and the 570 series component is the largest device with the highest number of Logic Elements in the family. Since the VAT demonstration platform hardware is intended to demonstrate the VAT architecture and serve as a development platform, the largest available CPLD is desirable, especially considering the extremely minimal power penalty as compared to the next lower sized device in the same family. Externally, the selected CPLD must have a sufficiently high pin count to interface with the desired nonvolatile storage, FPGA, and any other required external peripherals. Adding up the pin counts from Table 4.3.4 indicates that the CPLD must have at least 142 user input/output lines available. This number can possibly be decreased if desired by careful multiplexing of the attached peripherals, however for the development device and to ease debugging it is preferred that each peripheral have its own dedicated IO line, resulting in a 256-pin Micro

Ball Grid Array (MBGA) package being necessary. Combining all of these criteria together, the EPM570ZM256C6N device is selected.

The other core logic device, a Field Programable Gate Array, is responsible for high-performance aspects of system operation such as signal processing and file system activities. The power consumption of this device should be minimal, however since it is only powered up when signal processing or file system operations are necessary the more critical performance metric is operations per watt, not watts of "on" power. Again exploring the Altera product line due to our familiarity with implementation on their devices, we selected the Cyclone III series of devices due to their relatively low cost (when compared to the Stratix series) but still exceptional performance per operation. The sizing of the device must again be sufficient for the firmware required for signal processing and file system operations, which is described in Section 4.5. Initial development of the firmware prior to hardware build indicated that a "25" series device would have ample internal capacity. Externally, the selected FPGA must have a sufficiently high pin count to interface with the desired nonvolatile storage, CPLD, and any other required external peripherals. Adding up the pin counts from Table 4.3.4 indicates that the FPGA must have at least 119 user input/output lines available. This number can be decreased if desired by careful multiplexing of the attached peripherals, however for the development device and to ease debugging it is preferred that each peripheral have its own dedicated IO line, resulting in a 256-pin Fineline Ball Grid Array (FBGA) package being necessary. Combining all of these criteria together, the EP3C25F256I7N device is selected.

Lastly, all of the digital and analog devices require regulated supply voltages. To enable the system to operate with absolutely minimal radiated emissions which could adversely affect the collected scientific data, the VAT demonstration platform is designed to use linear voltage regulators. These regulators reduce the input battery voltage to the desired voltage by dissipating the excess power as heat. To maximize

efficiency of the regulation system, the battery voltage must be as close to the minimum input voltage of the regulators as possible. The highest voltage devices in the system operate at 3.3 V, while the Saft LSH20 batteries, identical to those used in Penguin and discussed in Section 3.4.1, have a single cell voltage of 3.66 V. If single-cell operation is desired with this battery selection, the linear regulators must operate with at most a 300 mV dropout, which is possible with modern regulators. Ultra-low dropout linear regulators are selected for the VAT demonstration platform to meet the single-cell operation desire as well as ensure that the system noise performance is not compromised by switching regulators. Future fielded systems would likely benefit from switching power supplies, especially buck or buck/boost supplies as described in Section 5.2.1.

Combining the selected CPLD and FPGA along with the requisite peripherals outlined in Table 4.3.4 yields the majority of information necessary for construction of the physical VAT demonstration platform. The remaining details not explained in this document are considered obvious requirements and are described in depth in each component's datasheet. The resulting electrical system block diagram is shown in Figure 4.1.

4.4   Power Budget

Power budgeting as described for the Penguin system apply equally well to the VAT demonstration platform; however the unique architecture of VAT causes many of the components to be shifted into different power domains. The VAT demonstration platform has only two power domains: domain 0 containing the sampling system, and domain 1 containing the processing and recording components. Note that GPS reception and synchronization of the internal clock is not implemented in the VAT demonstration platform, however the requisite I/O lines are available
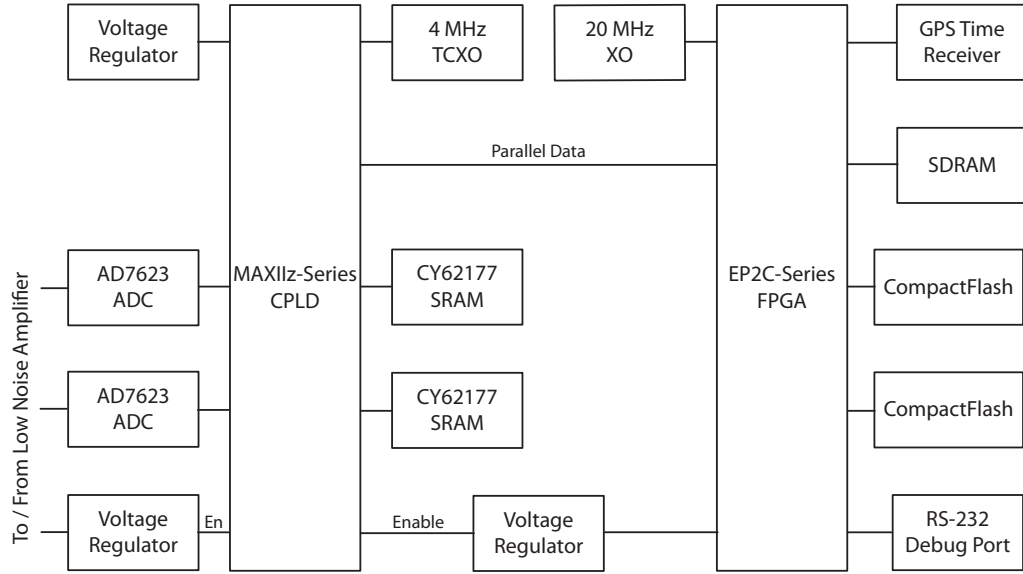
Figure 4.1: VAT Digital Architecture

on a pin header to facilitate future development. System power comparisons to the Penguin system, however, are still valid and easy to determine since the GPS time synchronization circuitry is independently powered and thus can be independently measured on Penguin and added to VAT power estimates. Table 4.4 contains the memberships of each power domain as related to the individual major components selected in Table 4.3.4. For power budgeting purposes, a power consumption estimate is made for each major peripheral. Nominal operating voltages and currents for each device are gathered from data sheets, and power estimates for programable logic devices, whether it be a FPGA, CPLD, or other type, are made based upon previous designs or simulations inside of the development tools.

Combining the major elements from Table 4.4 into a single unified average power consumption estimate again requires the duty cycle information to be considered. Recall from Section 3.4 that the power consumption is $P_{tot} = \sum D_i \times P_i$ where D_i is the duty cycle for the given component and P_i is the average power during

Device	Power Domain	Power (mW)	Quantity
AD7623 ADC	0	5.5	2
CPLD: EPM570ZM256C6N	0	10	1
Compact Flash	1	60	2
Crystal Oscillator, CPLD	0	5	1
Crystal Oscillator, FPGA	1	5	1
CY62177 SRAM, CPLD	0	5.0	2
CY62177 SRAM, FPGA	1	5.0	1
FPGA: EP3C25F256I7N	1	600	1

Table 4.2: Power domains and estimated draw of VAT platform peripherals

the on-time of the same component. Thus, the duty cycle must be determined. The VAT demonstration platform is designed to be a continuously sampling system which means that the sampling controller is always running, thus the duty cycle of power domain 0 is 100% or $D_0 = 1$. Power domain 1 is turned on only for signal processing and data storage. The exact duty cycle for this power domain is highly dependent on a number of factors such as: the signal processing performed, file system utilized, speed of compact flash, and internal clock rate of the FPGA. Assuming that the system performance is similar to that of Penguin, it is estimated that flushing the two megasamples of buffered data collected in a single buffer will take approximately two seconds. At a sampling clock rate of 10 kHz, this equates to 2 seconds every 100 seconds, or a 2% duty cycle ($D_1 = 0.02$). At a 100 kHz sample rate, the duty cycle becomes 2 seconds every 10 seconds or $D_1 = 0.20$. Generalized, the duty cycle of power domain 1 is the ratio between 2 seconds and the number of seconds that is required to collect 1,000,000 samples from each channel. Since the sample rate is known to be F_s Hz, then it is easily derived that $D_1 = \frac{2 \times F_s}{1000000}$. Using the estimated power consumptions and duty cycles discussed, the estimated power budget for the VAT demonstration platform digital electronics is described by Equation 4.2. Plugging the appropriate values into this equation yields power estimates for common sampling frequencies of $P_{F_s=10kHz} = 50.6$ mW and

$P_{F_s=100kHz} = 182$ mW. It is important to note that Equation 4.2 is only valid for sampling rates F_s which take longer than 2 seconds to collect one million points from each ADC. For sampling rates outside of this range, i.e. $F_s > 500$kHz, the currently implemented VAT architecture cannot keep up with the data rates. Increasing the FPGA clock rate and/or the sampling buffer SRAM sizes are two options to extend the maximum operating frequency of a given VAT implementation.

$$P(F_s) = \sum D_i \times P_i = 36 \text{ mW} + \frac{2 \times F_s \times 730 \text{ mW}}{1000000 \text{ samples}} \tag{4.2}$$

4.4.1   Battery Selection

The power source for the VAT demonstration must be able to operate at extremely cold environments since the system is likely to be deployed in the polar regions. Utilizing previous research described in Section 3.4.1, Saft LSH20 Lithium Thionyl Chloride batteries are again selected for their high performance in extreme temperatures as well as their extremely high power per mass density. Each D-sized cell generates approximately 3.66 V when fully charged, allowing for a 3.3 V rail to be possible from a single cell if ultra-low dropout linear regulators are used. As the batteries drain, however, the output voltage can drop significantly. The Saft provided data sheet considers a cell to be depleted when the voltage drops below 2 V. For some operational conditions, utilizing a single cell is desirable to reduce weight and physical bulk, and thus the VAT demonstration platform is designed to accommodate this mode. To fully utilize the energy in each cell, however, a future extended run-time system would need to either utilize two cells in series to provide a rail greater than 3.48 V when the batteries are nearly depleted (which is possible as well with the VAT demonstration platform, but additional power is burned off as heat), or use a buck-boost switching regulator to step up the single cell battery voltage as necessary. The use of switching regulators in VLF receivers is an area

that is lacking academically rigorous study, and is discussed as a potential item of future work in Section 5.2.1.

4.5   Firmware and Software Architecture

While extensive thought and care is required for component selection to maximize the efficiency of a VAT architecture based system as detailed in Section 4.3.4, the underlying architecture itself that is present on the programable logic fabrics is where the true innovation lies. Typical sampling systems all have a processor operating in the sampling loop. The processor often directly controls the ADCs readback as is the case in Penguin. More complex systems, such as the South Pacific Buoys, will have a small CPLD acting as glue logic to connect the ADCs to a First In, First Out (FIFO) buffer. In these FIFO buffered systems, however, the CPU is still always running during the sampling process. The FIFO buffers themselves give the CPU some relief from constant monitoring of the ADCs, allowing other tasks to be performed for brief periods of time while an interrupt or DMA controller periodically reads back the contents of the FIFOs. These buffers are usually quite small, with Buoy utilizing a 512-sample deep memory for each channel, which must be flushed no slower than once every 5.12 ms. Startup times for CPUs are typically much longer than this, with Penguin requiring approximately 60 ms to start, so duty cycling the CPU on these FIFO buffered systems is not possible. Some systems do sleep the CPU between buffer read-backs, but are typically sampling at much slower sampling rates [8].

The key innovation in the VAT architecture is the ability to control the ADCs and buffer their data for extended periods of time at extremely low power without requiring a CPU to be constantly running. Ultra-low power SRAMs, such as the Cyress CY62177 used in the VAT demonstration platform, can buffer millions of

samples while consuming only a few milliwatts of power. Assuming that the CPU can be replaced with lower powered control logic, a system buffering to this ultra-low power class of SRAMs can run for minutes between CPU operations as shown in the power budgeting in Section 4.4. The hardware architecture described thus far has an exceedingly low power consumption, but maintaining the functionality of the system lies in the design of the firmware and software. The fundamental architectural advances take place in the design of the CPLD firmware, while the FPGA firmware and associated soft core processor's software can utilize the streamlined design from Penguin.

The CPLD requirements include control of the analog-to-digital converters, readback of the ADC data, buffering of this data into the two external SRAMs, power control of the signal processing and storage FPGA, as well as clock management. To accomplish these tasks, several high-level modules are required in the firmware. Since the CPLD is fed a clock signal, at least two clock divider modules are required: one for generating the sampling clock, and one for clocking the remainder of the CPLD modules at the minimum rate possible for power savings. A module is also required to control and read data from the ADCs. This data must the be handled by a module responsible for managing access to the SRAMs for buffering the data. Since there are two SRAMs, one of which will be buffering data while the other is feeding previously recorded data to the FPGA, a multiplexer or MUX is required for each SRAM. Selecting which SRAM is used for which purpose is the responsibility of another module, dubbed the SRAM Arbiter. On the other side of the data flow, a firmware module is required to read the previously recorded data from an SRAM and pass it back to the FPGA. Lastly, a module is necessary to control the waking and sleeping of the FPGA power circuitry. These firmware modules are all internal to the CPLD, and are connected as shown in Figure 4.2.

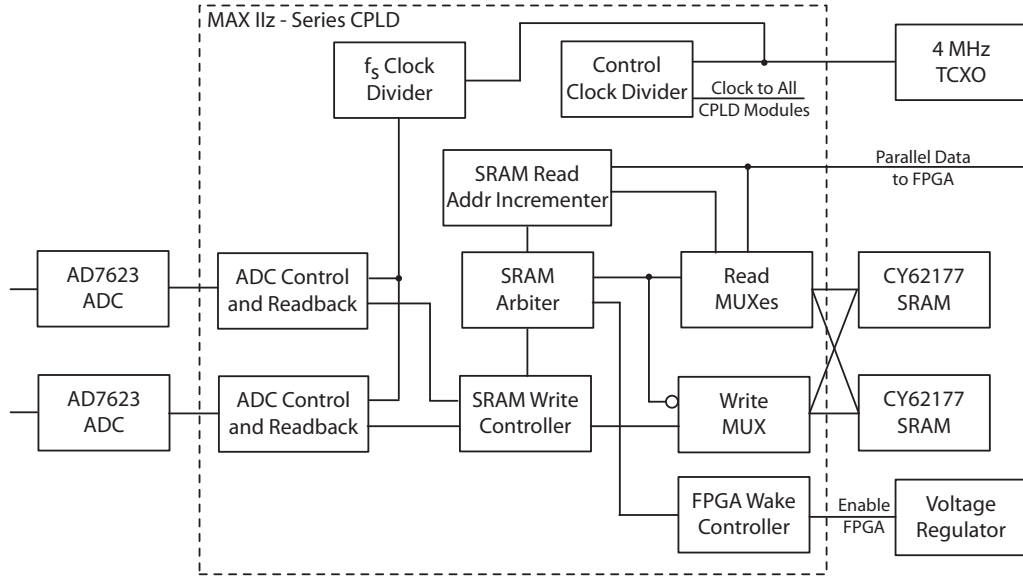On the VAT demonstration platform, the main system clock for the CPLD is

Figure 4.2: VAT CPLD Architecture

4 MHz. Internally, the CPLD must then divide this clock down to F_s to drive the sampling clock. In addition, the remaining CPLD modules must also have a clock source to perform their functions, which typically need to be at 5x the sampling clock. The FPGA readback module, however, should be clocked as fast as possible when active, and then should not be clocked at all while the FPGA is powered down. Both of these clock dividers are easily realizable in the CPLD firmware with simple counters and comparators.

Once the appropriate clock signals are available, the ADC controller module is constructed. This module is responsible for passing the required ADC control signals, namely conversion start, output enable, and readback complete, to the ADCs. The ADC controller has control and data signals directly attached to external I/O pins on the CPLD, as well as internal connections for signaling between other modules. To begin a conversion process, the ADC controller waits for a sampling clock pulse to be received from the sampling clock divider, which is immediately passed through to the ADC convert lines. Once the conversion is initiated, the

ADC controller must delay for either the maximum ADC conversion period or until a conversion complete signal is received from the ADCs. The later is most efficient and is implemented on the VAT demonstration platform, however it does require one additional signal from each ADC back to the CPLD. Once the ADC conversion is complete, the module is responsible for asserting the output enable or chip select line to each ADC, thereby allowing the ADC to present its data to the CPLD data lines. After the requisite read setup time, which is specified by the ADC manufacturer, the ADC controller then latches the 16-bit sampled data from each ADC into an internal register. With the sampled data safely received, the ADC controller de-asserts the chip select or output enable line to each ADC while also passing the first of the two 16-bit words onto the SRAM write controller module. One clock tick later, the second word can be passed to the SRAM write controller, at which point the ADC controller module sleeps until the next sampling clock pulse is detected.

The SRAM write controller module is responsible for taking in the two words of data collected on each sampling pulse from the ADC controller, serializing the data into two separate write operations, keeping track of the current write address location within the SRAM, and then actually writing the data into the SRAM via the SRAM write MUX. As each word of data to be written to the SRAM is received, it is placed in an internal register for buffering. When the buffer is not empty, the SRAM write controller retrieves the oldest buffered sample and places it on the SRAM data lines along with the currently stored SRAM address. Subsequently the controller asserts the chip select and write enable lines after the minimum setup time has elapsed to initiate the SRAM write within the SRAM itself. After a delay of the minimum write hold time, which is also determined by the SRAM, the write enable and chip select lines are de-asserted, and the internal address counter is incremented. If available, the next word of buffered data is placed on the data lines and the process is repeated. Under nominal operating conditions, two words are burst transfered to

the SRAM at a time between electrically long periods of inactivity.

Both the data and control signals from the SRAM write controller must pass through a multiplexer prior to being received at the SRAM. A second set of MUXes is required for the FPGA readback operations. While these MUXes are synthesized as standard MUXes, it is important to remember that the FPGA readback data path must have both input and output ports, while the SRAM write controller data path only requires output ports to the SRAM. For readback operation, control and address signals must propogate to the SRAM, while the data itself must propagate back, resulting in the need for a MUX both inbound and outbound. As previously described, however, SRAM writes only require data flow from the CPLD into the SRAM, including all control, address, and data signals. All of these MUXes directly interface with external I/O pins of the CPLD, which are then routed to the SRAMs.

While the primary output lines from the SRAM write controller are passed to the SRAM write MUX, one critical signal is also passed to the SRAM arbiter: an address overflow indicator. This indicator is pulsed when the address counter internal to the SRAM write controller has reached the end of the SRAM address space, when the next buffer write will be written to address zero. At this point, the SRAM arbiter is activated and switches the MUX selections after the pending write is complete. This switch causes the data path from the SRAM write controller to go to the other, currently empty, SRAM, while the FPGA readback SRAM controller is connected to the now full SRAM buffer. This timely switch ensures that samples can be continuously collected without clobbering previously stored data, assuming that the FPGA reads back the full SRAMs before the next SRAM arbiter activation. The second critical task for the SRAM arbiter is to alert the FPGA power control module that the SRAM is full and ready for readback once the MUX directions have been switched.

The FPGA power control module wakes up the FPGA each time an SRAM

buffer is full. The module has one input from the SRAM arbiter to determine when to wake up the FPGA, one input from the FPGA via a CPLD I/O pin to indicate when the FPGA operations are complete and is ready for shutdown, as well as one output line to directly enable and disable the voltage regulators for the FPGA. The implementation of this module is similar to that of a latch but also incorporates a delay prior to accepting a clear signal from the FPGA. The delay inhibits the clear line of the latch which prevents power-on glitching of the FPGA I/O lines from sending a false shutdown signal prior to even fully bootstrapping. These glitches can be observed from the FPGA both immediately upon activation of the associated voltage regulators as well as during the firmware bootstrap process.

Once the FPGA is awake and booted, the primary function of the CPLD is to pump the data from the full SRAM in the FPGA as requested. During the readback process, the sampling sections of the CPLD discussed thus far continue to operate without interruption, utilizing the other, empty SRAM buffer for the new samples. During the FPGA boot process, the FPGA readback SRAM controller resets an internal address counter to zero to ensure that the readback process starts with the oldest sample. The module then sends the address as well as the chip select and output enable signals to the SRAM to initiate the first word read. After the maximum read setup time the SRAM data lines contain the first word of data, which is passed through the FPGA readback SRAM controller and presented to the FPGA on the FPGA-CPLD I/O lines. The module then idles until the FPGA toggles the CS line, indicating that it is ready for the next word of data. The module increments the internal address counter and then repeats the read process as described. By having the CPLD maintain the current readback address pointer, no address lines are necessary between the CPLD and FPGA greatly reducing I/O count for both devices. This readback method is possible because the SRAM reads will occur in a known bounded amount of time, which can be configured into the

FPGA readback firmware and/or software. Additionally, the SRAM buffer sizes are fixed and known at compile time, so the FPGA can simply read the exact number of required samples each time without having to do additional bounds checking.

For maximum FPGA transfer speed, minimum FPGA on time, and thus minimal power consumption, the FPGA readback process should be clocked at the maximum frequency supported by both the FPGA transfer as well as the SRAM modules themselves. With a 55 ns read cycle time, the SRAMs can support a maximum readback rate of over 18 megawords per second, or 36 MBytes/sec. FPGA readback speed, however, is much more variable and is dependent upon any signal processing, the speed of the nonvolatile storage used, and its own internal clock rate.

Even with the CPLD firmware architectural advancement contributed by this research, firmware and software are still required on the FPGA for signal processing and data readback.

4.6 System Performance

Power budgets and estimations, such as those described in Section 4.4, are extremely valuable in resource planning and setting criteria for specific implementations, but are based on estimations and modeling, not real-world hardware. Taking power measurements of physical hardware is the only way to know exactly how much power a system actually draws under real-world operating conditions. The VAT demonstration platform is tested in a laboratory environment to determine real-world power consumption, initially at room temperature. As noted in Section 3.6, the similar components used on the Penguin system have nearly identical power performance at extremely low temperatures as they do at room temperature. For a field deployed VAT system, however, it is still imperative that thermal testing and thermal cycling is performed to best guarantee operations of the exact physical

parts and lots used in each system.

Since the VAT demonstration platform is used to demonstrate a highly-optimized digital system architecture for VLF reception, low power optimizations for the analog front end are outside the scope of this study. To convert the VAT demonstration platform into a field-ready receiver, a Low Noise Amplifier (LNA) and Anti-Aliasing Filter (AAF) specific to the desired sampling frequency must be selected. Due to the low power and low EMI of VAT, the traditional separation between the digital electronics and the antenna is no longer necessary, and thus the line driver circuitry can be removed for additional power savings. For reference, the LNA used on Penguin consumes 4.8 mA at 3.3 V if the line-driver operational amplifier is removed. The Penguin AAF, which is set to roll off at 30 kHz, consumes 4.6 mA at 3.3 V with its line driver operational amplifier removed. With a 3.6 V battery, the analog front end required for 30 kHz RF bandwidth operation consumes approximately 33.8 mW of power in addition to the VAT demonstration platform. As noted in Section 3.3.4, adding GPS synchronization consumes approximately 17 mW·hr per synchronization, or 100 mW for continuous oscillator disciplining.

The VAT demonstration platform power consumption is measured by independent measurements of the system during power domain 0 as well as power domain 0+1 operations (i.e. during sampling and buffering, as well as sampling and buffering and signal processing and data storage). Measurements of power domain 0 were made with a noise input to the analog-to-digital converters and without any of the analog front end elements. Table 4.6 lists the power consumption for a variety of possible sampling rates, F_s. As anticipated from the manufacturer data on the analog-to-digital converters as well as the typical operational performance of digital circuits, the power scales in an extremely linear fashion compared to sampling clock rate. A linear approximation of the power domain 0 test data results in Equation 4.3, where $I_{domain0}$ is in mA and F_s is in kHz.

Sampling Rate F_s	Current Consumption @ 3.6 V Battery
10 kHz	9.93 mA
30 kHz	10.96 mA
50 kHz	11.17 mA
80 kHz	11.49 mA
100 kHz	11.94 mA
140 kHz	13.00 mA
150 kHz	13.28 mA

Table 4.3: Measured VAT power domain 0 consumption

$$I_{domain0}(F_s) = 0.022 \times F_s + 9.9 \tag{4.3}$$

Measurements of power domain 1, which includes any digital signal processing as well as the data storage into the compact flash, are highly spurious. For the measurements, no additional signal processing was performed to provide a direct comparison with the Penguin system. During data transfers, the entire system power consumption (including both power domains 0 and 1), was measured to be approximately 750 mW at $F_s = 100$ kHz. This indicates that power domain 1 independently utilizes approximately 750 mW as configured for the test. Combining the information from this power domain 1 test along with the linear approximation given in 4.3, Equation 4.2 is updated to match the real world measured values in Equation 4.4. Again, the power $P(F_s)$ is measured in mW and the sampling frequency F_s is measured in kHz.

$$P(F_s) = \sum D_i \times P_i$$

$$= I_{domain0}(F_s) \times 3.6 \text{ V } \times D_0 + P_{domain1} \times D_{domain1}(F_s)$$

$$= (0.022 \times F_s + 9.9) \times 3.6 \times 1 + 750 \times \frac{2 \times F_s}{1000}$$

$$P(F_s) = 1.579 \times F_s + 35.64 \tag{4.4}$$

4.7 Conclusion

The VLF Advanced Technology architecture for broadband V LF data collection detailed in this chapter provides a radically different sampling system architecture from traditional designs. By eliminating the need for a full Central Processing Unit or Digital Signal Processor to be running throughout the sampling process, the necessary power is radically reduced. A specific sampling control system is described which replaces the traditional CPU or DSP role in sampling control, and a real-world, tangible realization of the architecture is presented in the form of the VAT demonstration platform. Measured power consumption during proper system operations is presented proving the ultra-low power nature of this novel architecture. This novel architecture fills a gap in ultra-low power sampling for sample rates that are too high to simply sleep the entire system between each sampled data points, but at the same time are slow enough to not continuously saturate the processing power of the signal processing systems.

The resulting power consumption of the VAT demonstration platform is approximately 5 times more efficient than the Penguin architecture when clocked at 100 kHz, and approximately 100 times more power efficient than the comparable subset of the Buoy 1.5 electronics for equivalent 100 kHz sampling. When compared on a practical level, the VAT platform requires approximately 24 Saft LSH20 battery cells weighing less than 5 pounds for an entire year of operation. The Buoy, however, required approximately one ton of sealed lead acid (SLA) batteries for one year of operation without recharging. The AWESOME VLF receiver when combined with a laptop requires a single 80 lb SLA battery to be exchanged once per day for field operations. Clearly, the power advantages of a VAT-based system not only permit
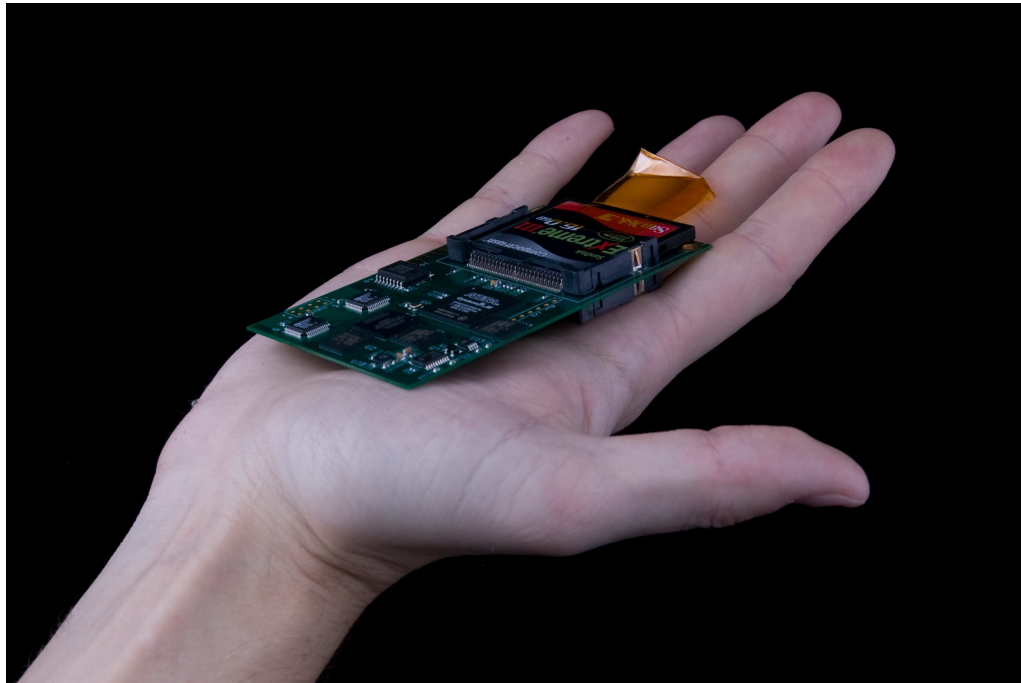
Figure 4.3: VAT ELF/VLF Receiver

extended polar operations which are currently not possible, but also greatly reduce the cost both in dollars and manpower for remote field operations of any type.

Chapter 5

Future Work

Both Penguin and VAT provide an enormous reduction in average and peak power consumption over the prior existing receiver systems. While these systems enable scientific studies not previously possible and greatly reduce the deployment costs for new ELF/VLF receiver sites, there is always room for improvement. Several studies, outlined below, have the potential to greatly advance the state-of-the-art in ELF/VLF reception, and possibly improve the Penguin and VAT designs by increasing their sensitivity, reducing system noise, and maximizing the use of available battery capacity.

5.1   Enhanced Flash Storage Media

Improvements in flash storage density and power performance will directly impact the overall performance of the VAT system described. VAT is currently capable of sampling for over a year continuously with a very modest power source. A VAT system is limited, however, by the availability of mass storage directly connected. While this problem may be worked around by arraying flash storage modules in the system, the complexity of the interconnect circuitry makes this a less than ideal

solution, and the additional components required for power control to minimize the power penalty of multiple connected cards further increases the complexity. It is likely that storage media densities will continue to grow due to consumer demand, and new low power media and technologies should be investigated for integration into the VAT system as they emerge on the consumer market.

5.2 VLF Interference Testing

5.2.1 Voltage Regulators

All of the VLF systems surveyed in Chapter 2 as well as the new Penguin and VAT receivers utilize linear voltage regulators. Linear voltage regulators work by resistively dropping the supply voltage across a transistor between the battery supply and load [55]. The transistor gate is constantly adjusted to ensure that the output voltage presented to the load is stable at a pre-defined level. Since this method relies on resistive loss to reduce incoming voltage, the battery supply must have a higher voltage than the target regulated system voltage. Additionally, the current consumption of the load at the regulated voltage is the same current drawn from the higher voltage battery bank. This excess voltage, dropped across the voltage regulator, burns off excess power according to ohms law, with the current defined by the system current draw.

Switching regulators utilize a transistor, inductor, diode, and capacitor to rapidly switch the input power on and off to an inductor, rectify the resulting current flow, and filter it to provide a constant DC voltage [7]. This architecture allows for the battery voltage to be higher than the required load voltage, as is the case with a Buck converter, or lower than the load voltage, as with a Boost converter. Hybrid Buck-Boost converters are available and allow the target load voltage to be higher

or lower than the battery voltage, allowing for more of the available battery capacity to be converted to the regulated voltage, even as the battery voltage significantly decays. Using a switch-mode voltage regulator would greatly enhance the useful battery life of any ELF/VLF receiver system, and would also provide an increase in overall system power efficiency.

The pitfall with switching regulators and ELF/VLF is with their rapid switching of a relatively high current. Switching power supplies often operate at frequencies of tens of kHz up through several MHz, and the current being switched is proportional to the system load. Clearly, as the receiver systems reduce their power demands the amount of switched current will also decrease, however due to the extreme sensitivity of the receivers to magnetic fields, these regulators have previously been generally discounted for system designs.

A rigorous study of the conducted as well as radiated emissions of various switching power supplies would be extremely useful to the ELF/VLF community. Particular areas of interest include high switching frequency, potentially maintaining the interference out of the ELF/VLF bands, synchronous switching of different regulators to provide multiple supply voltages without creating potentially in-band beat frequencies, and novel inductor designs, board layouts, and shielding to minimize radiated emissions.

5.2.2 Flash Media Noise

To greatly reduce the power requirement for autonomous ELF/VLF recordings, the Buoy systems, discussed in Section 2.3.3, as well as Penguin and VAT, utilized solid state flash memory over magnetic tape or disk. While all three of the aforementioned systems used CompactFlash due to its generally larger capacity than other competing formats, no known studies currently exist that evaluate and compare the unintentional electromagnetic emissions from various solid state memory

technologies. CompactFlash modules, NOR-based flash memories, typically have a higher peak power requirement than Secure Digital (SD) or MultiMedia Card (MMC), both NAND-based flash memories, however they also have a significantly higher read and write bandwidth, resulting in a roughly net zero power per byte performance [5, 6, 16].

A study of both real-world power consumption for available flash memory technologies, as well as their radiated and conducted emissions in the ELF/VLF band would allow for selecting the best suited COTS storage medium for future systems. Algorithms for file access and storage may also reduce the emissions during streaming writes to the storage media. In addition to flash storage, further investigation into SRAM and/or DRAM memories used by the system processor and sampling systems could further improve system noise performance by utilizing different clock rates and access patterns.

It is important to note that the studies required for both switching power supply as well as flash media noise involve complex electromagnetic interactions with literally millions of nodes internal to multiple ICs, and are thus not suitable for current EM field solvers. To properly answer these questions, hardware must be designed, built, and evaluated, and specific PCB layout and shielding practices will need to be established.

5.3 Automatic Bias Voltage Adjustment

An integral part of the Low Noise Amplifiers used in Penguin and VAT is a stable bias voltage source for the LNA ASIC. This voltage is temperature dependent, and is critical for the LNA to maintain maximum gain and a reasonable noise floor. In the LNA for Penguin and VAT, this voltage source is provided by a separate on-board adjustable linear regulator. As detailed in Section 3.7, this regulator is required

to be fine-tuned to the appropriate level once deployed. If this voltage source is not corrected for the current system temperature, the received signal is generally corrupted. Figure 5.3 shows a 24-hour period of Penguin data during which time the LNA bias voltage was set as appropriate for 25° C operating temperature. This period begins with the LNA nominally at 20° C, and cools down throughout the period to a final temperature of -55° C. Clear time-dependent, and thus temperature-dependent, degradation of the system noise floor is present in the data. After cool-down, the bias regulator was adjusted to match the nominal system temperature and the overall system performance immediately returned to expected levels.
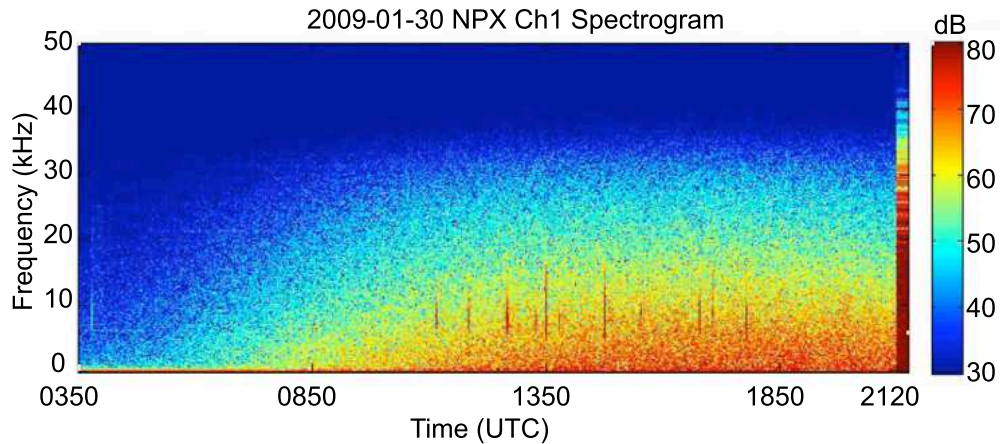

Figure 5.1: Penguin Data During Cool-down

While the Penguin system in particular is assumed to have a stable temperature environment over short, day-long periods of time, variations of the snow temperature over the course of a year can be significant at shallow depths [58]. This variation in temperature is expected to directly impact the Penguin noise floor, as the bias voltage is manually set during deployment and is not dynamically adjusted during operation. Development of a low-power circuit to automatically tune the LNA bias voltage could potentially remove this variation in data quality over the life of the

deployed system. A study of the impact for dynamically adjusting the bais voltage on the initial and final calibration data is necessary to ensure that such a circuit would not adversely affect the calibration accuracy of the data.

5.4 Integration of ASIC Signal Processors

The integration of Harriman's LNA to the Penguin and VAT systems enabled a significant amount of power that was previously consumed by the front-end to be eliminated. Tightly tuning the commodity processor architecture, as was done with Penguin, again reduced the system power footprint, and ultimately replacing the general purpose processor with a highly integrated CPLD for sampling control dramatically reduced the power even further. Following this trend, investigations into ultra-low power signal processing systems could bring demodulated narrow-band data, compressed broadband data, or new data products thus far not envisioned all well within reach of the compact autonomous systems discussed thus far. Similarly, advances in battery technology and flash storage could also lead to further overall system performance enhancements.

Bibliography

[1] AA Abidi. Direct-conversion radio transceivers for digital communications. In *1995 IEEE International Solid-State Circuits Conference, 1995. Digest of Technical Papers. 42nd ISSCC*, pages 186–187, 1995.

[2] D.W. Allan. Time and frequency (time-domain) characterization, estimation, and prediction of precision clocks and oscillators. *Ultrasonics, Ferroelectrics and Frequency Control, IEEE Transactions on*, 34(6):647–654, Nov 1987.

[3] D.W. Allan, N. Ashby, C.C. Hodge, and Hewlett-Packard Company. *The science of timekeeping*. Hewlett-Packard, 1997.

[4] K. Arisawa and T. Tamagawa. Alkaline storage battery that can be used with stability in a wide temperature range, February 2001. US Patent App. 09/794,407.

[5] CompactFlash Association. CF+ and CompactFlash Specification Revision 4.1.

[6] SD Card Association. SD Specifications Part 1 Physical Layer Simplified Specification, Version 2.00, September 2006.

[7] K.H. Billings. *Switchmode power supply handbook*. McGraw-Hill Professional, 1999.

[8] A. Bruce and R. Lee. A framework for the specification of scada data links. *Power Systems, IEEE Transactions on*, 9(1):560–564, Feb 1994.

[9] K. Bullough, ARL Tatnall, and M. Denby. Man-made ELF/VLF emissions and the radiation belts. *Nature*, 1976.

[10] E. Buracchini. The software radio concept. *Communications Magazine, IEEE*, 38(9):138–143, Sep 2000.

[11] W.J. Burtis and R.A. Helliwell. Magnetospheric chorus: Occurrence patterns and normalized frequency. *Planetary and Space Science*, 24(11):1007 – 1010, IN1–IN4, 1011–1024, 1976.

[12] P. Cappelletti. *Flash memories.* Kluwer Academic Publishers, 1999.

[13] DL Carpenter and TR Miller. Ducted magnetospheric propagation of signals from the Siple, Antarctica, VLF transmitter. *Journal of Geophysical Research*, 81(A16):2692–2700, 1976.

[14] T. Chung, D. Park, S. Park, D. Lee, S. Lee, and H. Song. System software for flash memory: A survey. *Lecture notes in computer science*, 4096:394, 2006.

[15] Atmel Corporation. ATtiny24/44/84 Preliminary Complete Datasheet Revision F (Lit. No. 8006FAVR02/07). *San Jose, CA: Atmel Corporation*, February 2007.

[16] SanDisk Corporation. SanDisk MultiMediaCard and Reduced-Size MultiMediaCard Product Manual (Doc. No. 80-36-00320), April 2005.

[17] P.H. Dana. Global Positioning System (GPS) time dissemination for real-time applications. *Real-Time Systems*, 12(1):9–40, 1997.

[18] André DeHon and John Wawrzynek. Reconfigurable computing: what, why, and implications for design automation. In *DAC '99: Proceedings of the 36th ACM/IEEE conference on Design automation*, pages 610–615, New York, NY, USA, 1999. ACM.

[19] Analog Devices. AD7623 16-Bit, 1.33 MSPS PulSAR ADC Datasheet Revision 0. *Norwood, MA: Analog Devices Inc*, July 2005.

[20] Analog Devices. AD7687 16-Bit, 1.5 LSB INL, 250 kSPS PulSAR Differential ADC in MSOP/QFN Datasheet Revision 0. *Norwood, MA: Analog Devices Inc*, April 2005.

[21] M.A. Fetcenko, S.R. Ovshinsky, B. Reichman, K. Young, C. Fierro, J. Koch, A. Zallen, W. Mays, and T. Ouchi. Recent advances in nimh battery technology. *Journal of Power Sources*, 165(2):544 – 551, 2007. IBA - HBC 2006 - Selected papers from the INTERNATIONAL BATTERY ASSOCIATION & HAWAII BATTERY CONFERENCE 2006 Waikoloa, Hawaii, USA 9-12 January 2006.

[22] AC Fraser-Smith and RA Helliwell. Global measurements of low-frequency radio noise, 1995.

[23] M. Gol-kowski, US Inan, AR Gibby, and MB Cohen. Magnetospheric amplification and emission triggering by ELF/VLF waves injected by the 3.6 MW HAARP ionospheric heater. *Journal of Geophysical Research*, 113(A10), 2008.

[24] L. Green. Understanding the importance of signal integrity. *Circuits and Devices Magazine, IEEE*, 15(6):7–10, Nov 1999.

[25] K. Hashimoto, I. Nagano, M. Yamamoto, T. Okada, I. Kimura, H. Matsumoto, and H. Oki. EXOS-D (AKEBONO) very low frequency plasma wave

instruments (VLF). *IEEE Transactions on Geoscience and Remote Sensing*, 35(2):278–286, 1997.

[26] RA Helliwell. VLF wave stimulation experiments in the magnetosphere from Siple Station, Antarctica. *Reviews of Geophysics*, 26(3), 1988.

[27] R.M. Herritty and J. Midolo. Nickel cadmium batteries for photovoltaic applications. In *Battery Conference on Applications and Advances, 1998., The Thirteenth Annual*, pages 255–257, Jan 1998.

[28] i Lotus Corporation. M12M User's Guide: Revision 1.0, May 2006.

[29] US Inan, M. Cohen, P. Scherrer, and D. Scherrer. VLF Remote-Sensing of the Lower Ionosphere with AWESOME Receivers: Solar Flares, Lightning-induced Electron Precipitation, Sudden Ionospheric Disturbances, Sprites, Gravity Waves and Gamma-ray Flares. In *2nd UN/NASA Workshop on International Heliophysical Year and Basic Space Science. Proceedings of the conference held 27 November-1 December, 2006 at Indian Institute of Astrophysics, Bangalore, Abstract Book., p. 66*, 2006.

[30] US Inan, M. Golkowski, MK Casey, RC Moore, W. Peter, P. Kulkarni, P. Kossey, E. Kennedy, S. Meth, and P. Smit. Subionospheric VLF observations of transmitter-induced precipitation of inner radiation belt electrons. *Geophysical Research Letters*, 34(2):2106, 2007.

[31] Motorola Inc. MC68HC11A8 Technical Data Book (Lit. No. MC68HC11A8/D). *Phoenix, AZ: Motorola Inc*, 1991.

[32] MP Johnson, US Inan, and DS Lauben. Subionospheric VLF signatures of oblique(nonducted) whistler-induced precipitation. *Geophysical Research Letters*, 26(23):3569–3572, 1999.

[33] P.J. Kuhn, A. Hoefler, T. Harp, B. Hornung, R. Paulsen, D. Burnett, and J.M. Higman. A reliability methodology for low temperature data retention in floating gate non-volatile memories. In *Reliability Physics Symposium, 2001. Proceedings. 39th Annual. 2001 IEEE International*, pages 266–270, 2001.

[34] T. Laaspere, WC Johnson, and LC Semprebon. Observations of auroral hiss, LHR noise, and other phenomena in the frequency range 20 Hz–540 kHz on OGO 6. *J. Geophys. Res*, 76(19):4477–4493, 1971.

[35] H.J. Landau. Sampling, data transmission, and the nyquist rate. *Proceedings of the IEEE*, 55(10):1701–1706, Oct. 1967.

[36] W. Lewandowski and C. Thomas. Gps time transfer. *Proceedings of the IEEE*, 79(7):991–1000, Jul 1991.

[37] W. Loser and C. Mattheck. Theory of the thermal switching behaviour of a PTC-resistor device. *Physica Status Solidi (a)*, 18(1), 1973.

[38] Kwok-Hong Mak and B.L. Holland. Migrating electrical power network scada systems to tcp/ip and ethernet networking. *Power Engineering Journal*, 16(6):305–311, Dec 2002.

[39] Kazuo Makita. Vlf-lf hiss emissions associated with aurora. *Memoirs of National Institute of Polar Research. Series A, Aeronomy*, 16:1–126, 19791200.

[40] D.R. Martinez, R.A. Bond, and M.M. Vai. 7.4 Embedded Signal Processing. *The Electrical Engineering Handbook*, 2006.

[41] W.M. McRae and N.R. Thomson. VLF phase and amplitude: daytime ionospheric parameters. *Journal of Atmospheric and Solar-Terrestrial Physics*, 62(7):609–618, 2000.

[42] B. Mossawir, I.R. Linscott, U.S. Inan, J.L. Roeder, J.V. Osborn, S.C. Witczak, E.E. King, and S.D. Lalumondiere. A TID and SEE Radiation-Hardened, Wideband, Low-Noise Amplifier. *IEEE Transactions on Nuclear Science*, 53(6):3439, 2006.

[43] RG Olsen. Power-transmission electromagnetics. *IEEE Antennas and propagation magazine*, 36(6):7–16, 1994.

[44] H.W. Ott and H.W. Ott. *Noise reduction techniques in electronic systems.* Wiley New York, 1988.

[45] R.A. Powers. Batteries for low power electronics. *Proceedings of the IEEE*, 83(4):687–693, Apr 1995.

[46] C. Price, M. Asfur, W. Lyons, and T. Nelson. An improved ELF/VLF method for globally geolocating sprite-producing lightning. *Geophysical Research Letters*, 29(3):1–1, 2002.

[47] T. Rappaport. *Wireless communications: principles and practice.* Prentice Hall PTR Upper Saddle River, NJ, USA, 2001.

[48] B. Razavi. Design considerations for direct-conversion receivers. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 44(6):428–435, 1997.

[49] AS Rodger, AJ Smith, and MJ Rycroft. Antarctic Studies of the Coupled Ionosphere-Magnetosphere System [and Discussion]. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, pages 271–287, 1989.

[50] Saft. Primary Lithium Battery LSH 20 (Doc. No. 31015-2-1006). *Bagnolet France: Saft Speciality Battery Group*, 2006.

[51] Saft. Primary Lithium Batteries Selector Guide (Doc. No. 31048-2-0209). *Bagnolet France: Saft Speciality Battery Group*, 2008.

[52] D. Scherrer, M. Cohen, T. Hoeksema, U. Inan, R. Mitchell, and P. Scherrer. Distributing space weather monitoring instruments and educational materials worldwide for IHY 2007: The AWESOME and SID project. *Advances in Space Research*, 42(11):1777–1785, 2008.

[53] Philips Semiconductor. The I2C-Bus Specification: Version 2.1, January 2000.

[54] DC Shafer, AD Brown, WJ Trabucco, and US Inan. A programmable and low-power ELF/VLF receiver for automatic geophysical observatories. *ANTARCTIC JOURNAL OF THE UNITED STATES*, 29:361–362, 1994.

[55] C. Simpson. Linear and switching voltage regulator fundamentals. *National Semiconductor Corporation, Santa Clara, CA, USA*, 1995.

[56] JK Skwirzynski. *Design theory and data for electrical filters.* van Nostrand, 1965.

[57] S. Subbarao, G. Halpert, and I. Stein. Safety considerations of lithium-thionyl chloride cells. *Unknown*, 1986.

[58] K.E. Trenberth and J.G. Olson. Temperature trends at the South Pole and McMurdo Sound. *Journal of Climate*, 2(10):1196–1206, 1989.

[59] C.D.S. Tuck. *Modern battery technology.* Ellis Horwood New York, 1991.

[60] International Telecommunication Union. ITU-R Recommendation V.431-7: Nomenclature of the frequency and wavelength bands used in telecommunications. *Geneva*, May 2000.

[61] T. von Sydow, B. Neumann, H. Blume, and TG Noll. Quantitative analysis of embedded FPGA-architectures for arithmetic. In *Proceedings of the 2006 IEEE international conference on Application specific Systems, Architectures and Processors (ASAP 2006), Steamboat Springs, Colorado*, volume 11, page 13, 2006.

[62] TG Wolf and US Inan. Path-dependent properties of subionospheric VLF amplitude and phase perturbations associated with lightning. *Journal of Geophysical Research*, 95(A12), 1990.