

UNIVERSITY OF CALIFORNIA,
IRVINE

The Performance and Initial Results of the ARIANNA Prototype

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Jordan Christian Hanson

Dissertation Committee:
Professor Steven Barwick, Chair
Professor Gaurang Yodh
Professor Stuart Kleinfelder

2013

DEDICATION

I would like to make a special dedication to the man who inspired in me a passion for nature and a love of science. Timothy O'Halloran was a science teacher and football coach at Holland Hall middle school, in Tulsa, OK where I grew up. Throughout the 4th-8th grades, he inspired and taught my classmates and I, both during the summer and the regular school year. He was the epitome of an educator of young people: knowledgeable, but also able to instill wonder. The Tim S. O'Halloran Young Alumni award for teaching excellence was established in 2009. May he rest in peace.

"The heart of the discerning acquires knowledge, for the ears of the wise seek it out." -
Proverbs 18:15

"Who is like the wise man? Who knows the explanation of things? Wisdom brightens a man's face and changes its hard appearance." - *Ecclesiastes 8:1*

Omnia cum omnibus junguntur

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	ix
ACKNOWLEDGMENTS	xi
CURRICULUM VITAE	xii
ABSTRACT OF THE DISSERTATION	xiv
1 Introduction	1
1.1 The Cosmic Ray Origin Problem	1
1.2 What is an ultra-high energy neutrino, and why are they interesting?	9
1.3 The Diffuse Flux of GZK Neutrinos, and the Askaryan Effect	13
1.3.1 Diffuse Signals	13
1.3.2 The Askaryan Effect	23
1.3.3 Experimental Confirmations of the Askaryan Effect	27
1.4 Potential New Physics with ARIANNA	30
1.4.1 Flavor physics of cosmogenic neutrinos	30
1.4.2 An example of Lorentz Invariance Violation characterization with ARIANNA	33
2 Properties of the Ice in Moore’s Bay	36
2.1 Depth of the RIS in Moore’s Bay, and Experimental Geometries	38
2.2 Absorption and Reflection Properties of the Ross Ice Shelf in Moore’s Bay	46
2.2.1 Dielectric Properties of Ice, Seawater, and the Debye Model	46
2.2.2 Experimental Technique and Antenna Equations	54
2.2.3 Results	62
2.3 Polarization Measurement	74
2.4 Surface Waves	77
2.4.1 Experimental Setup	77
2.4.2 Results	81
2.4.3 Analysis	85

3	The Icicle1 Prototype Detector	93
3.1	The Prototype Station - Software, control, and communications	94
3.1.1	Station Electronics	94
3.1.2	The IceSoft control system	100
3.1.3	Configuring the Run States	102
3.2	The Icicle1 Prototype - Power systems and Signal Antennas	107
3.2.1	Solar and Wind Hardware, and the Basic Charging Circuit	107
3.2.2	Several properties of the Log-Periodic Dipole Array Antennas	112
3.2.3	Pulsed LPDA experiments in Aldrich Park	128
3.2.4	Heartbeat Events from Different Periods	142
4	Icicle 1 Prototype Data Analysis	145
4.1	Housekeeping data, 2009-12	147
4.2	The Icicle1 Prototype event classes and the properties of thermal background fluctuations	159
4.2.1	IceSoft event classes - forced triggers, thermal triggers, and heartbeat triggers	159
4.2.2	Mathematical properties of thermal noise, informing the Monte Carlo simulations	163
4.2.3	Forced trigger comparisons to Monte Carlo	173
4.2.4	Analysis of heartbeat triggers	182
4.3	Time Differences Analysis	187
4.3.1	2009-10 Season: 12.2 days of live-time.	188
4.3.2	2010-11 Season: 20.58 days of live-time.	192
4.3.3	2011-12 season: 62.53 days of live-time	196
4.4	Description of Waveform variables, comparisons to Monte Carlo simulation .	203
4.5	Searching for Low-Planarity Events satisfying all other Cuts	233
4.6	Demonstrations of the Icicle1 Prototype Trigger	249
5	Neutrino Sensitivity of Prototype Station, and Flux Limits	256
5.1	Effective Volume, Flux Limits	257
5.2	Sensitivity to UHE neutrinos	270
6	Conclusion	276
	Bibliography	280
	Appendices	286
A	Selected Antenna derivations	286
B	Information Regarding the shelfmc Monte Carlo software	291
B.1	Compiling and running <i>shelfmc</i>	292
B.2	The steering file - input.txt	292
B.3	The declarations file - declarations.hh	294

LIST OF FIGURES

	Page
1.1 UHECR spectrum	3
1.2 UHECR spatial correlation	8
1.3 UHECR arrival directions	8
1.4 p- γ total cross-section	10
1.5 The GZK cutoff	11
1.6 $\nu - \gamma$ total cross-section	12
1.7 UHECR chemical composition	15
1.8 UHE neutrino flux models, 1	20
1.9 UHE neutrino flux models, 2	20
1.10 More recent UHE ν flux predictions	22
1.11 Askaryan radiation geometry	24
1.12 Properties of Askaryan pulses	26
1.13 Measurements of the Askaryan effect, 1	28
1.14 Measurements of the Askaryan effect, 2	29
1.15 LIV, 1	35
1.16 LIV, 2	35
2.1 Reflection studies geometry	37
2.2 Reflection studies equipment	42
2.3 Example reflections, 1	44
2.4 Example reflections, 2	45
2.5 Ice permittivity vs. frequency	48
2.6 Attenuation length vs. temperature	53
2.7 Calibration pulse between Seavey and LPDA	55
2.8 LPDA VSWR	56
2.9 LPDA VSWR, zoomed	57
2.10 Calibration, direct, and angled waveforms (2010)	63
2.11 Calibration, direct, and angled waveforms (2011)	64
2.12 Raw reflection FFTs (squared)	65
2.13 Reflection coefficient and attenuation length results (frequency averaged) . .	66
2.14 Snow effects, 1	68
2.15 Snow effects, 2	69
2.16 Auxillary figure for 2011 reflection studies	70
2.17 Attenuation length versus frequency, 1	71

2.18	Attenuation length versus frequency, 2	72
2.19	The polarization fraction, F , is shown as a function of frequency for transmission through air (dashed line) and for the angled bounce configuration at the ARIANNA site (solid line) in 2010.	76
2.20	The polarization fraction, F , is shown as a function of frequency for transmission through air (dashed line) and for the direct bounce configuration at the ARIANNA site (solid line) in 2011.	76
2.21	Surface wave experimental setup	78
2.22	Surface wave results, 1	83
2.23	Surface wave results, 2	84
2.24	Shadowing effect schematic	86
3.1	Satellite photograph of Moore’s Bay	95
3.2	Icicle1 prototype layout	96
3.3	Software and data management in Icicle 1	101
3.4	Icicle1 and Icicle 2 solar frame	108
3.5	Prototype charging circuit	109
3.6	Prototype photographs (2010)	111
3.7	Log-periodic dipole array, 1	113
3.8	Log-periodic dipole array, 2	113
3.9	Log-periodic dipole array, 3	115
3.10	Log-periodic dipole array, 5	117
3.11	Gain versus detector efficiency	118
3.12	Pockel cell driver output	121
3.13	Aldrich park tests, 1	122
3.14	Aldrich park tests, 2	123
3.15	Effective height confirmation	124
3.16	Effective height vs. time (estimate)	125
3.17	(Left): <i>Bore sight</i> configuration. (Right) Configuration 1.	128
3.18	(Left): Configuration 1 (45°). (Right) Configuration 1 (90°).	129
3.19	(Left): Configuration 2. (Right) Configuration 2 (45°).	129
3.20	Configuration 2 (90°).	129
3.21	Aldrich park studies geometry, 1	130
3.22	Aldrich park studies geometry, 2	131
3.23	Power vs. distance, LPDA	132
3.24	Peak voltage vs. distance, LPDA	132
3.25	Typical pulse from Aldrich park studies	134
3.26	Gain vs. frequency, 1	135
3.27	Power spectral density, LPDA, 1	135
3.28	Gain vs. frequency, 2	137
3.29	Power spectral density, LPDA, 2	137
3.30	Gain vs. frequency, 3	139
3.31	Power spectral density, LPDA, 3	139
3.32	LPDA configurations, con’d	140
3.33	LPDA positions for the Icicle 1 prototype	141

3.34	North channel heartbeats, different times	143
3.35	East channel heartbeats, different times	143
3.36	South channel heartbeats, different times	144
4.1	Icicle1 housekeeping data set, first season	148
4.2	Daily battery periodicity	149
4.3	Madgetech temperature loggers, first year	151
4.4	Global housekeeping plot, three seasons of Icicle1 operations	153
4.5	Aerogen-6 current vs. wind speed	154
4.6	Histogram of wind speeds, Moore's Bay	155
4.7	Icicle1 and Icicle2 housekeeping data correlation	155
4.8	Wind/live-time simulation	157
4.9	Live-time from solar power (estimate)	158
4.10	ARIANNA amplifier gain vs. frequency	169
4.11	Power spectral density of forced triggers by channel	174
4.12	Planarity distribution for Monte Carlo forced triggers	175
4.13	Planarity distribution for forced triggers	176
4.14	Action of the v_{pp} algorithm	178
4.15	Amplitude distribution of forced triggers, white noise Monte Carlo	179
4.16	Amplitude distribution of forced triggers, realistic Monte Carlo	179
4.17	Forced trigger example	180
4.18	Microsecond times vs. time	181
4.19	Δt vs. time, first season	183
4.20	Heartbeat example waveforms	186
4.21	Δt of thermal triggers, first season	189
4.22	Δt vs. time, first season	190
4.23	Further Δt graphs	191
4.24	Δt vs. time, second season	193
4.25	Δt vs. time, second season, con'd	194
4.26	Δt distribution of thermal triggers, second season	195
4.27	Raw Δt vs. time, third season	198
4.28	Δt vs. time, period 1 of third season	199
4.29	Δt distributions for third season	200
4.30	Thermal trigger example, second season	204
4.31	2009-10 heartbeat example	209
4.32	2010-11 heartbeat example	210
4.33	T_{pp} delay from Icicle1 trigger	211
4.34	T_{pp} vs. v_{pp} , 2009-10	213
4.35	T_{pp} vs. v_{pp} , 2010-11	213
4.36	T_{pp} vs. v_{pp} , 2011-12	214
4.37	MC thermal events w/ trigger delay	215
4.38	Amplitude distributions for thermal triggers, first season	218
4.39	Amplitude distributions for thermal triggers, second season	218
4.40	Amplitude distributions for thermal triggers, third season	219
4.41	Slope of amplitude dist., first season, vs. time	220

4.42	Slope of amplitude dist., second season, vs. time	221
4.43	Slope of amplitude dist., third season, vs. time	222
4.44	Slope of amplitude dist., MC, vs. time	223
4.45	First season amplitude, comparison to MC	227
4.46	First season West amplitude, comparison to MC	228
4.47	Second season amplitude, comparison to MC	229
4.48	Second season West amplitude, comparison to MC	230
4.49	Third season amplitude, comparison to MC	231
4.50	Third season West amplitude	232
4.51	MC threshold demonstration	234
4.52	Final results, first season	236
4.53	Final results, second season	239
4.54	Random non-thermal events cut	240
4.55	Heartbeat event that is shifted in time	242
4.56	Hkd execution times	242
4.57	Third season noise events	244
4.58	Third season noise events	245
4.59	Final results, third season	246
4.60	Amplitude distribution, North (2011-12)	247
4.61	Amplitude distribution, East (2011-12)	248
4.62	Amplitude distribution, South (2011-12)	248
4.63	Demonstrations of trigger	251
4.64	Demonstrations of trigger, 2	253
4.65	Demonstrations of trigger, 3	253
4.66	Demonstrations of trigger, 4	254
5.1	Short-time Fourier transform of 2011 reflection	265
5.2	Short-time Fourier transform of 2011 reflection	265
5.3	Effective aperture of a 31×31 ARIANNA array, HRA, and single station . .	268
5.4	Final flux limits from ARIANNA prototype Icicle1	272

LIST OF TABLES

	Page
2.1 Ross ice shelf depth data	40
2.2 Shelf depth timing	41
2.3 Attenuation length fit parameters	73
2.4 Surface wave setup timing delays	79
3.1 Power consumption	110
3.2 Power consumption	111
3.3 Log-periodic dipole array, 4	116
4.1 Start/stop dates	152
4.2 Calculations of trigger rates	172
4.3 Trigger rates and event amplitudes	173
4.4 Heartbeat time offsets	185
4.5 Live-time summary	187
4.6 Δt cuts, first season	192
4.7 Second season event totals, first and second period	192
4.8 Δt cuts, second season	196
4.9 The five periods of the third season	197
4.10 Event rates from Δt distributions, five periods of third season	201
4.11 Δt cuts for third season	202
4.12 Δt cuts for third season, con'd	202
4.13 Definition of causality variable	205
4.14 Causality and MC	206
4.15 Causality and MC, 2	208
4.16 Trigger delay and MC, 2009-10	216
4.17 Trigger delay and MC, 2010-11	216
4.18 Trigger delay and MC, 2011-12	216
4.19 Slopes of amplitude distributions	217
4.20 MC and amplitude cut	223
4.21 DAC threshold settings, three seasons	224
4.22 Amplitude cuts, second season	231
4.23 Amplitudes, third season	231
4.24 Final results, first season	236
4.25 Final results, second season	239
4.26 Random non-thermal event, cuts	240

4.27	Third season noise events	244
4.28	Event cleaning summary, third season	245
4.29	Final results, third season	246
4.30	Demonstrations of trigger	252
5.1	MC settings to produce $V_{eff}\Omega$	269
5.2	Predicted neutrino counts, 3.0 year live-time	274

ACKNOWLEDGMENTS

First and foremost I would like to thank my parents, Drs. Christian and Carol Hanson, and my sister Sarah, for instilling in me a love of learning, the courage to pursue my goals, and for their financial support of my education.

I would like to thank my advisor, Dr. Steven Barwick, for his extensive time and expertise, his knowledge and all that he has taught me about our field, and for guiding me through my doctoral education. I would also like to thank Dr. Barwick for the opportunity to journey to Antarctica, which was a wonderful experience. I would also like to thank Dr. Stuart Kleinfelder, for his time and dedication as well.

A special thanks is in order to Drs. Gaurang Yodh and Dr. Kleinfelder for participating in my dissertation defense committee.

To all of my friends and colleagues at the Department of Physics and Astronomy at the University of California, Irvine, I would like to express my heartfelt gratitude for their knowledge and support.

Specifically, I would like to thank my research group: Dr. Kamlesh Dookayka, Joulien Tatar, Dr. Corey Reed, Thorin Duffin, Mashid Roumi, Liang Zou, Zongnan Fan, and Dr. Eric Berg. These individuals accomplished a great deal of work in a short time, and I consider myself lucky to have been a part of such a team.

The Antarctic Ross Ice-shelf ANtenna Network Array (ARIANNA) project was supported by the U.S. NSF Office of Polar Programs and U.S. NSF Physics Division, award numbers 0839133 and 0970175, and by the Department of Energy under contract DE-AC-76SF-00098. I would also like to thank the dedicated and hardworking members of Raytheon Polar Services for their logistical, technical and medical support during the Antarctic expeditions.

CURRICULUM VITAE

Jordan Christian Hanson

EDUCATION

Doctor of Philosophy in Physics University of California, Irvine	2013 <i>Irvine, California</i>
Master of Science in Physics University of California, Irvine	2008 <i>Irvine, California</i>
Bachelor of Science in Physics (Intensive track) Yale University	2007 <i>New Haven, Connecticut</i>

RESEARCH EXPERIENCE

Graduate Research Assistant University of California, Irvine	2007–2013 <i>Irvine, California</i>
Research Assistant Yale University	Summer 2006 <i>New Haven, Connecticut</i>
Research Assistant Los Alamos National Laboratory	Summer 2005 <i>Los Alamos, New Mexico</i>

TEACHING EXPERIENCE

Teaching Assistant University of California, Irvine	2007–2008 <i>Irvine, California</i>
Online Teaching Coordinator University of California, Irvine	Summer 2009 <i>Irvine, California</i>

REFEREED JOURNAL PUBLICATIONS

- Design and Performance of the Autonomous Data Acquisition System for the ARIANNA High Energy Neutrino Detector** 2013
IEEE Symposium on Radiation Measurements and Applications
- Multi-GHz Waveform Sampling and Digitization with Real-Time Pattern-Matching Trigger Generation** 2013
IEEE Symposium on Radiation Measurements and Applications
- A prototype station for ARIANNA: A detector for cosmic neutrinos** 2010
Nuclear Instruments and Methods in Physics Research A
- Detection and Imaging of He₂ Molecules in Superfluid Helium** 2008
Physics Review Letters

CONFERENCE PROCEEDINGS

- Ross Ice Shelf Thickness, Radio-frequency Attenuation and Reflectivity: Implications for the ARIANNA UHE Neutrino Detector** 2011
Proceedings of the 32nd International Cosmic Ray Conference, Beijing, China

INVITED TALKS

- Developing the Next Generation of UHE Neutrino Detectors in Antarctica
Seminar in High Energy Physics at the University of Kansas, Lawrence, KS

PROFESSIONAL MEMBERSHIPS

- American Physical Society, Member

ABSTRACT OF THE DISSERTATION

The Performance and Initial Results of the ARIANNA Prototype

By

Jordan Christian Hanson

Doctor of Philosophy in Physics

University of California, Irvine, 2013

Professor Steven Barwick, Chair

The Antarctic Ross Ice-shelf Antenna Neutrino Array (ARIANNA) experiment exploits serendipitous properties of the Ross Ice Shelf in Western Antarctica. The ice shelf forms the fiducial volume of an ultra-high energy (UHE) neutrino detector capable of observing cosmogenic neutrinos with energies in excess of 10^{17} eV. The clarity of the shelf ice and the reflectivity of the ocean-ice interface enhance the detection of radio-frequency (RF) electromagnetic pulses created by neutrino interactions via the Askaryan effect. An array of autonomous electronics stations outfitted with radio antennas listen for these pulses.

A prototype station was designed in 2009, using sustainable power and RF trigger and digitization electronics. It was deployed in Moore's Bay in December 2009 during the Austral summer, and additional data was collected in two subsequent seasons after the system rebooted automatically during Austral spring. This data located and helped to remove local anthropogenic noise. A total of 90.4 days of live-time was achieved, with thermal noise as the single background. Additionally, data characterizing the environment of Moore's Bay was collected and used in the development of future power systems and RF electronics.

The depth and dielectric properties of the ice beneath the detectors were calculated using data taken during the expeditions. The linear fit to the frequency-dependent, temperature-averaged attenuation length of radio waves is $\langle L \rangle = (500 \pm 30 - (0.18 \pm 0.05)\nu[\text{MHz}])$ m,

and the reflection coefficient at the oceanic interface is $0.70 \leq \sqrt{R} \leq 1.0$ at 68% C.L. The depth of the shelf is 576 ± 2 m. The possibility of surface effects capable of boosting detector sensitivity was also revealed in these studies.

Finally, the RF data collected by the first ARIANNA station was analyzed and compared to Monte Carlo simulations modeling the detector response. These studies have produced the first upper limit on the differential UHE neutrino flux from ARIANNA, and the integral flux limit assuming a E^{-2} spectrum is $E^2\phi < 10^{-5}$ GeV cm⁻² s⁻¹ sr⁻¹. These limits demonstrate that the completed version of ARIANNA will have the capability to observe UHE cosmogenic neutrinos.

Chapter 1

Introduction

1.1 The Cosmic Ray Origin Problem

The primary objectives of physics and astronomy were somewhat different before the beginning of the twentieth century. Astronomers recorded the motions, positions, and brightness of stars and planets, in order to learn what lies beyond our own local solar environment. Physicists studied matter, attempting to analyze its fundamental constituents, beginning with basic substances and progressing to the chemical elements, the atom, and eventually the smallest sub-atomic particles. Although this distinction between these two branches of science still holds meaning today, the discovery of radiation penetrating the atmosphere in the early 1900's has required scientists to undertake cross-disciplinary research in order to explain it fully.

Italian physicist Domenico Pacini, in 1911 [4], was able to establish that mysterious particles that seemed to ionize air spontaneously could not be attributed solely to isotopes such as radium or thorium from the Earth's crust. This simple experiment compared the decrease in the rate of spontaneous discharge of an electroscope just above and 3 meters below the

surface of the ocean. Having subtracted known radioactive backgrounds, Pacini concluded that the absorption coefficient of water was attenuating the radiation which had to be coming from above. In the other extreme, German physicist Victor Hess measured the *increase* in the rate of spontaneous electroscope discharge between 1-5 km above the Earth's surface, in a balloon [55]. This resulted in Hess sharing the Nobel prize in 1936, with Carl Anderson, who discovered the positron. It was Robert Millikan who eventually introduced the term *cosmic rays* to describe the ionizing radiation originating from deep space. From these simple experiments emerged a scientific field known as *astroparticle physics*, focused on the detection and study of cosmic rays, in the form of hadrons, leptons, and photons. We now have a detailed knowledge of the flux with respect to energy of hadronic cosmic rays, which extends over a remarkable sixteen orders of magnitude (figure 1.1). While the overall trend is a power law, there are at least two changes in the spectral index, called the *knee* and the *ankle*. Below 100 GeV, the masses of the nuclei vary widely and the simpler approach is to plot the flux versus the kinetic energy per nucleon. Above 100 GeV, given that the proton mass is 0.938 GeV, these units are almost identical [80].

The spectrum of cosmic rays below 10^{20} eV follows a power law of the form $E/E_0 = \exp(-\gamma E)$. The *spectral index* γ changes at least twice, at 10^6 GeV and 10^{10} GeV. A basic argument for a strictly galactic population of accelerators at or below 10^6 GeV (the knee) is that the magnetic gyro-radius of the accelerated particles is much smaller than the typical radius of our galaxy. If we take the component of the particle momentum which is orthogonal to the magnetic field, $P \approx 10^7$ GeV, a magnetic flux strength of $B \approx 1$ nT (typical for the Milky Way), and assume the particle is a proton such that its charge $Z = +1$, we obtain a gyro-radius that is four orders of magnitude smaller than the galactic radius (≈ 35 kpc):

$$\rho_{gyro}[kpc] = \frac{3.3P[GeV]}{B[T]|Z|} \frac{1}{3.1 \times 10^{19} m \text{ kpc}^{-1}} \approx 10^{-3} kpc \quad (1.1)$$




Figure 1.1: The flux with respect to the natural logarithm of the energy per nucleus of cosmic rays. [80]

While one parsec is much smaller than the radius of the galaxy, galactic cosmic rays do begin to diffuse out of the confinement of the galaxy once the gyro-radius reach a few parsecs [50]. This is known as the leaky box model, which must account for radioactive decay and secondary cosmic ray creation in addition to diffusion out of the magnetic bottle of the galaxy. Another basic argument for the galactic nature of sub-knee cosmic rays is that the energetics of objects like supernova remnants allows us to explain (via first and second order Fermi acceleration) how particles of such energies could have been accelerated.

Suppose a supernova produces a gaseous shock front with a front and back side, traveling at velocities v_{front} and v_{back} , and that a charged particle traveling anti-parallel to the shock front at velocity v collides with the back side and begins to reflect between the front and back. When it is traveling anti-parallel to the gas, it gains an energy $\Delta E = (1/2)m(v + v_{back}) - (1/2)mv^2$. When it is traveling parallel to the gas, it loses an energy $\Delta E = (1/2)m(v - v_{front}) - (1/2)mv^2$. On average (assuming the initial velocity $v \gg v_{front}$ and v_{back}), the energy gain is $\Delta E \approx mv\Delta v$, where $\Delta v = v_{back} - v_{front}$. We get to assume that the leading edge of the shock front decelerates earlier than the trailing edge because the supernova is not taking place in a total vacuum, but is surrounded by inter-stellar matter such as dust and gas. Thus, the leading edge will interact with this matter and decelerate. As long as there is an appreciable Δv , this process of first order Fermi acceleration can lead to fractional energy gains that can accelerate particles to at least 100 TeV.

$$\frac{\Delta E}{E} \approx 2 \frac{\Delta v}{v} \tag{1.2}$$

Under a fully relativistic treatment, the factor of 2 on the right hand side of the equation becomes (4/3), but the energy gain is still linear in Δv . The derivation of the energy gain of second-order Fermi acceleration is exact whether one uses relativistic arguments or simple

kinematic ones. In this scheme, a population of gas clouds with intrinsic magnetic fields has an average velocity \bar{u} , and there are energy gains and losses for a particle of velocity v with parallel and anti-parallel collisions:

$$\Delta E_{gain} - \Delta E_{loss} = (1/2)m(2uv + u^2) - (1/2)m(-2uv + u^2) = 2\left(\frac{u}{v}\right)^2 \quad (1.3)$$

Given that the cloud velocities are non-relativistic, the energy gains per pair of collisions will be small under these circumstances, and the velocity of the particle needs to be relativistic already if this mechanism is to accelerate particles to energies close to the knee. Accelerated nuclei also have to live long enough not to create secondary cosmic rays or lose significant energy by interacting with non-magnetic material. One possible scenario is that first-order Fermi acceleration precedes second-order, thus providing the required injection spectrum.

In order to push the primary cosmic ray energy above the knee, special mechanisms are required beyond simple Fermi acceleration. What follows is a quick example of the energetics of a binary star system with a compact object such as a neutron star. Suppose a neutron star pulls gas away from its companion in the binary, and the gas gains kinetic energy as it falls into the gravitational potential. We have (for in-falling protons with mass m_p)

$$\Delta E = -G \int_{\infty}^{R_{pulsar}} \frac{m_p M_{pulsar}}{r^2} dr = G \frac{m_p M_{pulsar}}{R_{pulsar}} \approx 70 MeV \quad (1.4)$$

The above calculation assumes a pulsar mass of 10^{30} kg, and a radius of 20 km. Because the energy per particle reaches relativistic a relativistic regime, we may assume it has a speed comparable to the speed of light. Using $(1/2)m_p v^2$ for the kinetic energy yields about half

the speed of light. Neutron stars can have magnetic fields as large as 10^6T , and the distance between ordinary star and neutron star could be $\approx 10^5$ m. Assuming the (variable) magnetic field strength is perpendicular to the trajectory of the in-falling protons, we have

$$E = \int \mathbf{F} \cdot d\mathbf{s} = q \int \mathbf{v} \cdot \mathbf{B} ds = ecB\Delta s \approx 10^{19}eV \quad (1.5)$$

So it is at least plausible that simple kinematics and electromagnetism could accelerate cosmic rays to the highest observed energies in the vicinity of a special astrophysical object. Active Galactic Nuclei (AGN), are an even more dramatic example of enormous energy generation, and a potential candidate for acceleration. From equation 1.1, assuming an approximately massless proton such that the momentum and energy are equal, we have a magnetic gyro-radius of about 1 kpc, which would quickly diffuse out of the host galaxy and become an extra-galactic cosmic ray. If the ultra-high energy cosmic ray (UHECR) travels through the Milky Way a distance of one kpc before arriving at the Earth, the magnetic deflection would shift the trajectory of the UHECR away from the original trajectory by $\approx 15^\circ$. Thus arises the difficulty in pinning down the origin of UHECRs: even if the cascade of particles they create in the Earth's atmosphere is easily detected (giving quantities like the energy, mass and charge), the uncertainties in the arrival direction limit spatial correlations with known astrophysical objects like neutron binaries and AGN [6].

Figure 1.2 demonstrates how good this type of correlation has become with the most recent UHECR data. Plotted on the y-axis is the correlation coefficient describing how many events are within a 3.1° radius centered on an AGN from the Veron-Cetty and 12th Veron AGN catalogues, including 1,2 and 3- σ statistical errors. The x-axis is the number of detected UHECRs above 6×10^{19} eV. The AGN chosen from the catalogues have to be within 75 Mpc of the Earth. If AGNs really are the source of the cosmic rays above a certain energy, and

the magnetic deflection was not too strong (e.g. the cosmic rays propagate a much shorter distance with respect to their gyro-radii through the Milky Way), the spatial correlation coefficient would approach unity. However, the data is ambiguous, since the limiting value of this correlation coefficient is $0.38_{-0.06}^{+0.07}$.

In figure 1.3, the black dots represent 69 UHECRs with energy greater than 55 EeV detected at the Pierre Auger observatory up to 2009, in galactic coordinates. The blue circles represent the locations of the AGNs that are close enough to allow UHECR propagation to the Earth without an interaction with the cosmic microwave background (see below). The relative shading of the AGN indicators represents relative exposure, with the dark blue circles having the most exposure.

Accelerators of cosmic rays could also create high energy gamma rays, either by inverse Compton scattering, or neutral pion decay (equation 1.7). Gamma rays provide an uncharged messenger which would preserve directional information in the presence of magnetic fields. In fact, there is currently a class of experiments [19] [48] [21] [20] detecting gamma rays from within the Milky Way at energies less than 1 TeV. Processes like Compton scattering and $p\text{-}\gamma$ interactions limit gamma-ray astronomy at red-shifts of $z > 1$ to energies less than ≈ 100 GeV, making the astrophysics cosmological in distance, but lower in energy. At energies increasingly greater than a TeV, $\gamma\gamma$ interactions between gamma rays and intergalactic infrared photons, and eventually the cosmic microwave background (CMB), begin to take place. This limits instead the red-shift to $z \ll 1$, making the astrophysics high-energy, but restricts distances to the local galactic environment. It is the strength of the electromagnetic interaction which prevents gamma rays from carrying information across cosmological distances. Neutrinos, which are electrically neutral and weakly-interacting, stream freely to the Earth from sources with $z \gg 1$ at high energy. Thus the hope is that UHE neutrinos will facilitate the study of high-energy astrophysical accelerators throughout the universe, and help us to solve the cosmic ray origin problem.




Figure 1.2: This data is from a recent analysis [6] done by the Pierre Auger Collaboration, a ground-based UHECR observatory in South America.




Figure 1.3: This data represents the arrival directions of UHECRs, in galactic coordinates, from [6].

1.2 What is an ultra-high energy neutrino, and why are they interesting?

Being electrically neutral and only weakly interacting, neutrinos preserve their trajectories and can propagate cosmological distances. Greisen, Zatsepin, and Kuzmin realized in 1966 that, given the existence of UHECRs, there must be a process whereby their flux is attenuated by interactions with the CMB [49] [86]. Three years later, Berezhinsky and Zatsepin pointed out [10] that this process leads to the creation of UHE neutrinos. The attenuation of the flux of UHECRs is known as the GZK process:

$$p^+ + \gamma_{CMB} \rightarrow \Delta^+ \quad 1232 \text{ MeV}/c^2 \quad (1.6)$$

$$\Delta^+ \rightarrow p^+ + \pi^0 \rightarrow p^+ + 2\gamma \quad (1.7)$$

$$\Delta^+ \rightarrow n + \pi^+ \rightarrow n + \mu^+ + \bar{\nu}_\mu \rightarrow n + e^+ + \bar{\nu}_\mu + \nu_\mu + \nu_e \quad (1.8)$$

Assuming we can boost to the rest frame of the proton, equation 1.8 is equivalent to the Δ^+ resonance at a photon energy of 300 MeV for a fixed-target collider (figure 1.4). On the y-axis we have the total p- γ cross section, in microbarns, for photons of energy ϵ' incident on fixed target protons. The 340 MeV resonance is the Δ^+ resonance. There are several resonances slightly above the main peak, until multi-particle production begins and the trend becomes smooth. In the astrophysical case, the proton must find a relatively higher energy CMB photon (10^{-3} eV in order to explain why the GZK cutoff is less than a few times 10^{19} eV. This figure comes from reference [80].




Figure 1.4: The total cross section of the $p\text{-}\gamma$ interaction shows a resonance around 300 MeV. Boosting to the rest frame of a UHECR proton makes a CMB photon appear energetic enough to produce a Δ^+ . [80]

This resonance decays in one of two ways. In equation 1.7, the Δ^+ baryon decays into a neutral pion and a lower energy cosmic ray, with 2/3 probability. Or, the baryon could decay into a neutron and charged pion, as in equation 1.8, with 1/3 probability. Direct pion production via virtual meson exchange occurs less frequently in these $p\text{-}\gamma$ interactions, but this process favors charged pion creation, and thus brings the charged to neutral pion ratio up to $\approx 1:1$ [22]. In the case of equation 1.7, the branching ratio for neutral pions to gamma rays is greater than 99%. If the reaction in equation 1.8 proceeds, we have three neutrinos, each with about 5% of the cosmic ray energy. The neutrino energies are dictated by the amount of inelasticity of the baryon decay, and how evenly distributed the energy is between the four leptons. The above processes are depicted for cosmic ray protons, however these processes also take place for heavier nuclei, after photo-disintegration takes place. The initial flavor ratio is 1:2:0, but the neutrinos will undergo sufficient flavor oscillations to be detected as $\approx 1:1:1$ given the cosmological distances involved (see section 1.4.1).

The GZK process saps energy from cosmic rays above the ankle energy, thereby cutting off the end of the spectrum. This is easiest to observe when the flux is multiplied by the cube




Figure 1.5: The flux of the very highest energy cosmic rays above 10^{17} eV demonstrates a cutoff in the spectrum. The flux times the energy cubed is plotted on the y-axis [83].

of the energy, in order to reveal the deviation from the power law above 5×10^{19} eV (figure 1.5). The interaction length is such that after 100 Mpc, most of the UHECRs disappear, restricting UHECR observation to the local universe. The data in this figure come from the HiRes, Telescope Array, Auger and AGASA collaborations [83]. While the data from AGASA is thought to be incorrect, the other three experiments all show an attenuation effect above 10^{20} eV. In this case, the Telescope Array data comes only from the hardware known as the surface detector (which excludes the fluorescence detectors). The broken power law fit to the Telescope Array data finds the ankle at $(4.9 \pm 0.3) \times 10^{18}$ eV, and the GZK suppression at $(4.8 \pm 0.1) \times 10^{19}$ eV.

Conversely, neutrinos from GZK interactions have an enormous interaction length compared to that of GZK cosmic rays. As an example, consider figure 1.6, which depicts the neutrino-photon cross section at low and high energies. At ≈ 1 TeV, the maximum total cross section is 10^{-2} picobarns, or about 10^{-26} cm². Let us assume that a 10^{19} eV neutrino interacts




Figure 1.6: Left: the total neutrino-photon cross section at lower energies (in pb) versus center of mass energy [18]. Right: the same cross section as the energy increases beyond the W-boson resonance.

predominantly with CMB photons, which have a number density of $\approx 400 \text{ cm}^{-3}$, and have an average energy of 10^{-3} eV . Then $\sqrt{s} \approx 2\sqrt{10^{-3} \cdot 10^{19}} = 0.2 \text{ GeV}$. In figure ??, we note that the cross section falls as the sixth power of energy if the center of mass energy is below the mass of the W-boson, and at 0.2 GeV the total cross section is less than 10^{-15} pb . Using this number for the total cross section (and the number density of CMB photons), we obtain an interaction length, $L = 1/n\sigma$ which is larger than the known universe. Fortunately for UHE neutrino astronomy, the CMB photon-UHE neutrino cross section predicted by the plotted equation implies virtually no neutrino attenuation. The neutrino-photon cross-section calculations come from [18].

1.3 The Diffuse Flux of GZK Neutrinos, and the Askaryan Effect

Although no UHE neutrinos have been observed, there are many astrophysical models and theories which predict potentially observable neutrino fluxes with respect to energy. What follows is a description of these fluxes, how they depend on energy, and what factors affect their normalization.

1.3.1 Diffuse Signals

A diffuse signal of neutrinos is such that the data displays no angular anisotropy, and we cannot pin down the source of the neutrinos as any particular astrophysical object. GZK neutrinos are an example of a diffuse signal in part because there is no large anisotropy in the UHECR data (figure 1.2); if we could identify point sources of cosmic rays capable of producing GZK neutrinos, then GZK neutrinos found to arrive from the same direction could be considered point sources as well. Of course this doesn't happen: the cosmic ray point source would have to be close enough to the Earth for direct detection of the protons/heavy nuclei, but far enough from the Earth to allow the GZK process to take place. It is more likely that GZK neutrinos arrive from across the universe, and have some small anisotropy, indicating that they do originate from point sources but that this effect is averaged out over time and space. This amount of smearing will depend on details such as the density of cosmic ray accelerators as a function of redshift (see below). Figure 1.8 (below) depicts the energy spectrum of GZK neutrinos (multiplied by the square of energy to reveal deviations from an E^{-2} power law), along with several models accounting for various theoretical assumptions.

There are six factors that affect the normalization of the flux of GZK neutrinos which we will discuss in turn. First, the chemical composition of the UHECRs shifts the neutrino

flux by at least two orders of magnitude. The blue shaded region in figure 1.8 assumes that the UHECR composition is strictly protonic above the ankle in the cosmic ray energy spectrum. The decaying delta baryon in equation 1.8 originates from a cosmic ray proton with a center of mass energy (along with the CMB photon) above the 1232 MeV threshold. Suppose instead that the cosmic ray begins as a nucleus of iron, stripped of all electrons. To undergo the GZK process, the Fe nucleus must first photo-disintegrate, producing the nucleon to create the delta resonance. The nucleon must have a fraction of the total energy of the nucleus, E_{CR} , and in the case of iron $E_n \approx E_{CR}/56$. That means the cosmic ray nucleus must have fifty-six times the energy of an equivalent proton to produce a GZK nucleon. Requiring that the cosmic ray nucleus be 56 times as energetic lowers the flux, and ultimately the neutrino flux. This argument is shown by the pink shaded region in figure 1.8, where the flux (times energy squared) has been reduced by a factor of 100.

The measurement of the composition of the cosmic rays is not yet settled (figure 1.7). The fluxes at the relevant energies are too small to measure directly by satellite or balloon-borne detectors, so indirect methods are employed by utilizing the earth atmosphere as a large calorimeter. The charge of the primary particle is related to the depth in the atmosphere of maximum particle density of the developing cascade, known as the shower maximum (X_{max}), but the interpretation of this data is complicated by an inadequate understanding of the interaction mechanisms at such high energies, and systematic differences between the detectors. The data are consistent with pure protons, compositions comparable to galactic mixtures in the GeV/nucleon-TeV/nucleon range, and even pure iron is not excluded at energies near the GZK cut-off [61].

The y-axis in figure 1.7 is the average value of X_{max} depth in g/cm². The x-axis is the energy of the cosmic ray primary, and we can see that while the global analysis favors proton models at higher energies, there is a spread in the data, which seems to be trending back toward iron models above 10¹⁹ eV. The highest-energy data point published by Auger is

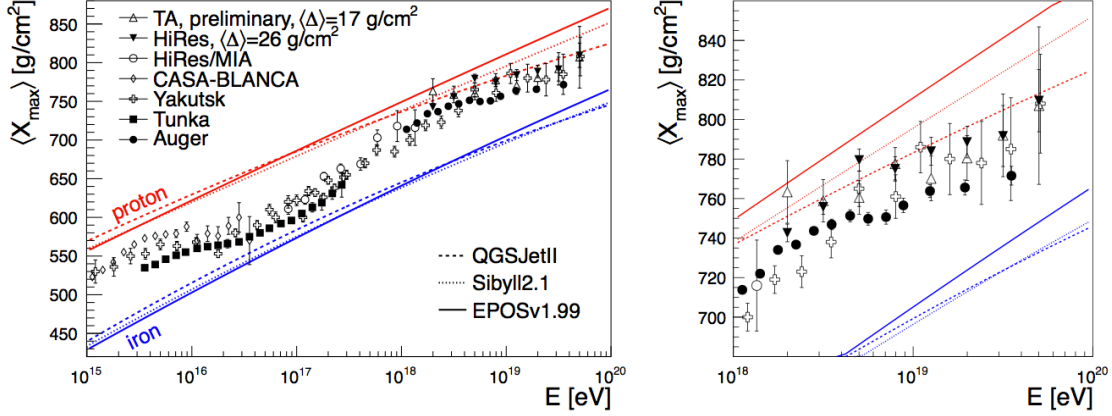


Figure 1.7: A global analysis of cosmic ray composition data, including ground-based air shower experiments and air-fluorescence experiments. This figure comes from Kampert and Unger (2012) [61].

directly between iron and proton models. The authors in this reference are careful not to rule out heavy or light composition above 10^{18} eV, citing experimental uncertainties and the range of predictions by Monte Carlo analysis. The best statistics come from Auger, however the central values of $\langle X_{max} \rangle$ from the various experiments are slightly different.

The second factor that affects the prediction of the UHE neutrino flux is the shape of the UHECR energy spectrum. A simple way to model the power law form with a cut-off is to use a Heaviside step function [68]:

$$\Phi(E, z) = fG(z)E^{-\alpha}\Theta(E_{max} - E) \quad (1.9)$$

In equation 1.9, G is some function of the redshift z , discussed below, and f is a simple normalization. One method to approach the solid blue line in figure 1.8 is to allow E_{max} to increase above 10^{21} eV [68], arguing that we only observe the truncated version of the UHECR spectrum after the GZK effect has taken place. Another choice for the injection

spectrum, however, is a power law with an exponential cut-off:

$$\Phi(E, z) = fG(z)E^{-\alpha} \exp\left(-\frac{E}{E_{max}}\right) \quad (1.10)$$

In equation 1.10, E_{max} is now the energy where the GZK effect begins to turn on, allowing for slight variations in the threshold. The width of the delta resonance around 340 MeV in figure 1.4, for example, could cause the cut-off energy to vary by a factor of ≈ 3 . The observed cosmic ray data does in fact follow a broken power law above the ankle, with the cutoff energy beginning around 4.8×10^{19} eV. Since the data display a power law with an exponential cut-off, it is reasonable to assume that the cosmic ray sources follow this pattern as well. If the parameter α is tuned lower, then the number of UHE neutrinos and the average UHE neutrino energy both increase, because the number of higher energy cosmic rays increases. The best estimate for the value of α is ≈ 2.3 . Similarly, the normalization must be chosen to match flux measurements, and the neutrino flux increases with increasing normalization. Finally, the the dip in the UHECR spectrum below the GZK cutoff is due to direct lepton pair production: $p^+\gamma \rightarrow p^+e^+e^-$. This interaction only plays a role well below the Δ^+ resonance. Near the GZK energy the $p^+\gamma \rightarrow \Delta^+$ interaction dominates because it proceeds with only one QED vertex at tree-level, whereas $p^+\gamma \rightarrow p^+e^+e^-$ requires two QED vertices which makes the matrix element less probable.

The third factor that affects the GZK neutrino flux is the set of assumptions made about cosmology. Measurements like the angular power spectrum of the CMB temperature from the COBE and WMAP satellites have shown that we live in an expanding, accelerating universe, represented by the *concordance model* with $\Omega_\Lambda = 0.7$ and $\Omega_M = 0.3$. The flux of cosmic rays weakly depends on these cosmological factors [70], because the source distribution

with respect to redshift is proportional to the quantity

$$\frac{dt}{dz} = H_0^{-1}(1+z)^{-1} (\Omega_M(1+z)^3 + \Omega_\Lambda + (1 - \Omega_M - \Omega_\Lambda)(1+z)^2)^{-1/2} \quad (1.11)$$

Taking $\Omega_M = 1$ and $\Omega_\Lambda = 0$ gives the Einstein-de Sitter universe, but using the concordance parameters yields a dt/dz larger by a factor of $\sqrt{3}$ in the limit $z > 1$. Overall the enhancement to the UHE neutrino flux is about 60% at all energies.

The fourth factor that affects the neutrino flux is the definition of the function $G(z)$ in equations 1.9 and 1.10. This function controls the density of cosmic ray sources per unit redshift and co-moving volume, e.g.

$$G(z) = H(z) \frac{dt}{dz} \Phi_0(E) \quad (1.12)$$

Here all the action is placed inside the function $H(z)$. The function $H(z)$ must have the property $H(0) = 1$, because the density of astrophysical accelerators is $\rho(z) = H(z)\rho(0)$. Because the source of UHECRs and therefore UHE neutrinos is unknown, we must speculate as to the exact form of $H(z)$. There are several logical choices, and what follows is a sketch of each of them. First, we could assume the trivial case where the evolution of the Universe is unrelated to the physics of cosmic ray production, and set $H(z) = 1$ for all z . This leads to conservative estimates for the neutrino flux. The simplest thing to do beyond the trivial case is to put in the constraint for zero redshift and assume some power law index, n controlling the increase of cosmic ray sources with increasing redshift: $(1+z)^n$. Most calculations allow this function track the same function for some other class of objects (figure ??), such as

star formation [1], AGN [70] [84], or gamma-ray bursts (GRBs) [53]. Each model assumes a source injection spectrum with index 2 and cutoff energy 10^{21} eV. The highest rates come from GRB evolution, followed by AGN and star-formation. In order of increasing neutrino flux, we have:

$$\begin{aligned}
H_{SFR}(z) &= (1+z)^3, \quad z < 1.3 \\
H_{SFR}(z) &= (1+1.3)^3, \quad z \geq 1.3, \quad z < 6 \\
H_{SFR}(z) &= 0, \quad z \geq 6
\end{aligned} \tag{1.13}$$

$$\begin{aligned}
H_{AGN}(z) &= (1+z)^3, \quad z < 1.9 \\
H_{AGN}(z) &= (1+1.9)^3, \quad z \geq 1.9, \quad z < 2.7 \\
H_{AGN}(z) &= (1+1.9)^3 \exp(-(z-2.7)/2.7)
\end{aligned} \tag{1.14}$$

$$\begin{aligned}
H_{GRB}(z) &= (1+z)^{4.8}, \quad z < 1 \\
H_{GRB}(z) &= (1+z)^{1.4}, \quad z \geq 1, \quad z \leq 4.5 \\
H_{GRB}(z) &= (1+z)^{-5.6}, \quad z > 4.5
\end{aligned} \tag{1.15}$$

The fifth factor which introduces variations in the GZK neutrino spectrum is any uncertainty or variability in the dynamics of the cosmic-ray/CMB photon interaction. This includes the energy dependence of the $p\gamma$ cross section (figure 1.4) and the inelasticity of the interaction itself. If the energy cut-off of the cosmic ray injection spectrum increases, then the GZK process can sample lower energy CMB photons, placing the interaction in the delta resonance at 1232 MeV. However, the average CMB photon energy increases linearly with redshift, so

cosmic ray sources of the same energy at higher redshifts see the higher energy tail of the $p\gamma$ cross section where multi-particle production begins. These processes have cross-sections which are lower by a factor of 5. The inelasticity of the $p\gamma$ interaction controls how much energy is transferred to the pion and divided amongst the leptons in equations 1.7 and 1.8. The value usually assumed is that 20% of the cosmic ray energy moves to the charged pion, but this depends on the maximum energy of the cosmic ray. Some references ([68], [70]) vary either the proton injection index or the maximum energy to get more UHECRs above 10^{20} eV. The inelasticity increases [80] to ≈ 0.5 at these energies, making the fraction of the total energy obtained by each neutrino larger than 10%. The peak in the blue flux around 10^{11} GeV in figure 1.8 is an example of such a neutrino flux.

The final factor that affects the GZK neutrino flux is oscillatory nature of the neutrino weak eigenstates. Unlike the charged leptons, the neutrino mass eigenstates are much different from the weak eigenstates, leading to flavor oscillations in vacuum. Assuming there is no sterile neutrino (and this remains an open puzzle: [77]), the overall neutrino flux is unaffected by oscillation physics. However, the flux in each flavor is affected, in that the standard GZK ratio of 1:2:0 averages over many oscillation lengths to a ratio at Earth of $\approx 1:1:1$. How precisely the events are distributed between the three flavors, given uncertainties of the measured oscillation parameters from terrestrial experiments is covered in section 1.4.1. If the detection mechanism for GZK neutrinos is the *Askaryan effect*, then the detector is sensitive to all three flavors, and the measured flux does not have to be corrected. However, detecting neutrinos in this way leaves us with the challenge of interpreting radio-frequency pulses as tau, muon, or electron type neutrinos. The physics of the Askaryan effect is the focus of the next section.

On the y-axis in figure 1.8, we have the differential flux with respect to energy (GeV), multiplied by the square of the energy. The models represented by the blue shaded region come from [68], and assumes only protons as the primary cosmic ray particle. One choice that

leads to the large flux above 10^{10} GeV is the high cut-off energy of the injection spectrum. The upper blue line represents the most liberal model, tuning in addition the maximum redshift of cosmic ray sources, and cosmological source evolution parameters to their most favorable values. This procedure moves the flux above the Waxman-Bahcall [17] bound, around 10^{-7} . The black lines are derived from reference [70], which contains calculations for no source evolution function $H(z)$, standard assumptions about $H(z)$, and an $H(z)$ following GRB evolution (strong evolution).

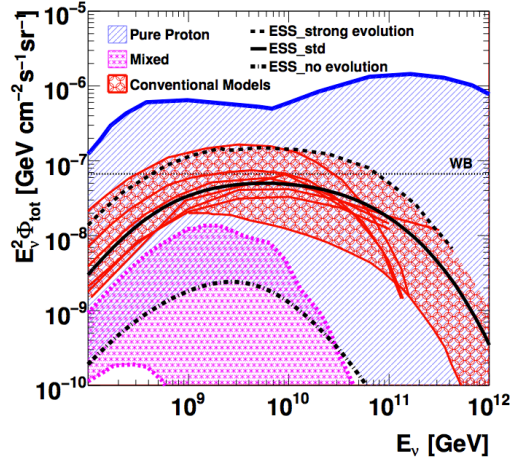


Figure 1.8: UHE GZK- ν (all-flavor) flux predictions ($\times E[\text{GeV}]^2$) are show here, coming from the following references: pure proton [68], conventional models: [70], [53], mixed: [70], [24], [53], Waxman-Bahcall: [17]. Other constraints include [22], [23], and a good summary can be found in [61].




Figure 1.9: (Left): UHE GZK neutrino flux predictions assuming different $H(z)$ functions. (Right): GRB evolution, varying the index of the injection spectrum produces observable UHE neutrino fluxes. [53].

Figure 1.8 contains a variety of models of the UHE neutrino flux from the last decade of research. However, experiments such as IceCube and ANITA have helped to eliminate some of the more extreme and optimistic models (for example the upper proton curve from [68]). Figure 1.10 comes from [61], and contains a more recent summary of calculations. The absolute fluxes are slightly lower due to a better understanding of both the GZK cutoff shape and the overall UHECR flux. One open question is whether the most optimistic pure-proton models in figure 1.10 violate the recent Fermi-LAT diffuse gamma ray limit [22] on the neutrino flux; the answer to this question lies in how well constrained the extra-galactic background light and extra-galactic magnetic fields are [14] [2]. The GZK pure-proton models vary by source evolution function: the top of the red area in this figure corresponds to FR-II type galactic redshift evolution, while the bottom is the standard evolution of star formation. The GZK pure iron models also vary by source evolution, and the flux is lower by an order of magnitude.

The grey area is more complex, as it includes different galactic to extra-galactic source transition models. The top of the grey area represents transition models where the dip in the spectrum between the ankle and the GZK cutoff is caused by the Bethe-Heitler process (e^+e^- pair-production by cosmic rays incident on CMB photons). The lower side of the grey area represents models where the transition occurs right at the ankle. In both scenarios, the cosmic ray maximum energy and energy spectral hardness must be tuned to match current data. Notice that the ESS (standard, fig. 4 from [70]) model from figure 1.8 (corresponding to a $H(z)$ function $\propto (1+z)^3$) represents about the same flux as the lower curve of the red band the more recent calculation in figure 1.10. The only difference is that the $H(z)$ function in the former flattens at $z = 1.9$ rather than $z = 1.3$, and cuts off exponentially. The ESS (fig. 4) calculation refers to AGN-like evolution. The ESS (strong) model assumes $H(z) \propto (1+z)^4$, and is closer to the upper red curve in the more recent figure, which represents FR-II type evolution. Because the FR-II type evolution already produces fluxes that are close to the Fermi-LAT diffuse gamma-ray bound (depending on what assumptions

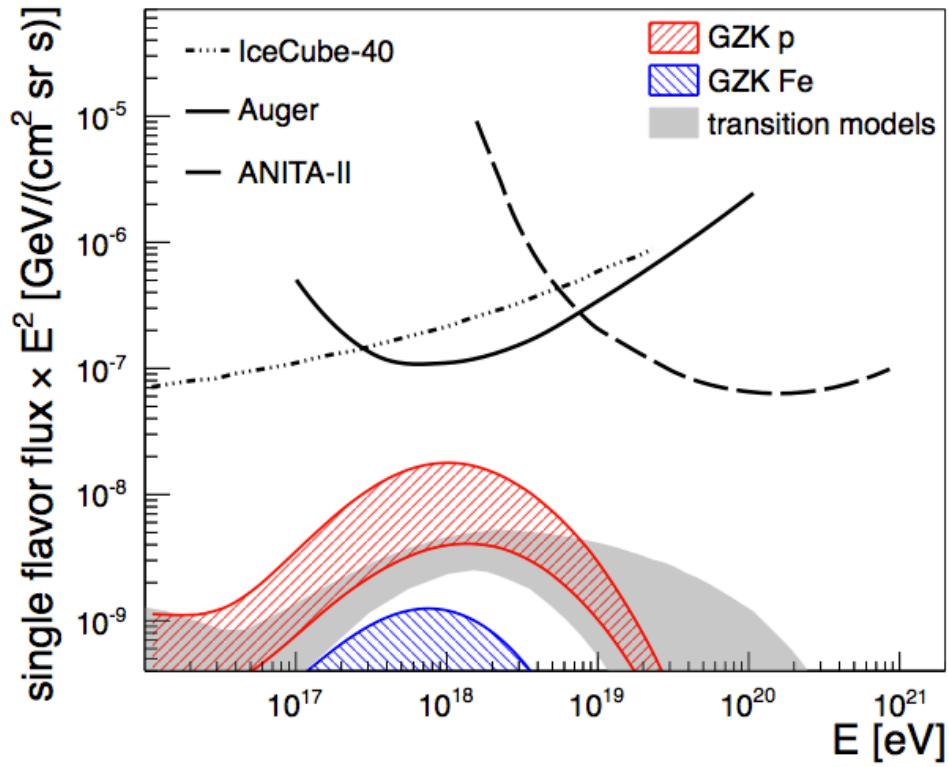


Figure 1.10: More recent predictions of the (single-flavor) UHE ν flux. Experimental upper limits are given by the black lines, including IceCube [43], ANITA-II [37], Auger [51]. Models of GZK neutrino production for proton and iron primaries, and various transition models come from [61].

are made), the ESS (strong) curve is somewhat optimistic.

1.3.2 The Askaryan Effect

Gurgen Askaryan realized in 1962 [5] that high-energy neutrino interactions produce linearly polarized radio-frequency (RF) pulses due to a decelerating negative charge excess in the hadronic and electromagnetic showers (figure 1.11). Thus, a class of experiments [30] [26] [36] [31] has developed which focuses on the detection of these Askaryan pulses originating in favorable dielectric media such as ice. This effect is similar to Cherenkov radiation in ice, where energy radiated by a charged particle is emitted at a characteristic Cherenkov angle ([58], chapter 13) when the velocity $\beta = |\mathbf{v}|/c$ is greater than $1/n$.

$$\cos(\theta_c) = \frac{1}{\beta\sqrt{\epsilon(\omega)}} \quad (1.16)$$

The square root reduces to the constant index of refraction in the limit of small dispersive effects. In an electromagnetic shower that produces Askaryan radiation, the negative charge excess can be as large as 25% of the total number of positrons and electrons, forming a decelerating total charge that radiates power proportional to the square of the charge excess. Figure 1.11 includes a schematic of the polarized radiation from the charge excess developing along the z-axis. Since the shower energy is proportional to the excess charge, the radiated power is proportional to the square of the shower energy. These scaling relationships cause (conveniently) the electric field in the far field of the radiation to be proportional to the shower energy, and thus the voltage across antenna terminals as well. Along the z-axis, the shower develops a negative charge excess $Q(z')$. The Poynting vector of the radio wave in this case is \mathbf{u} , and the wave is linearly polarized along $\mathbf{p} = -\hat{\mathbf{u}} \times (\hat{\mathbf{u}} \times \hat{\mathbf{z}})$. The shower maximum occurs after a few meters (see figure 1.12), and the shower can be stretched if the LPM effect takes place.




Figure 1.11: A schematic of the Askaryan radiation resulting from a hadronic or electromagnetic shower induced by a UHE neutrino. The graphic comes from reference [60].

The electric field in the time domain has been calculated analytically and modeled modeled for different viewing angles [59] [60]. Here the viewing angle is the angle away from the Cherenkov angle the observer views the Askaryan radiation. The highest frequencies (and narrowest pulse widths), which peak around several GHz, are only observable at or within 1-2 degrees of the Cherenkov angle (figure 1.12). As the viewing angle increases, the observed electric field in the time domain is dominated by frequencies of order 100 MHz, changing the pulse properties in the time-domain. The charge excess develops over several meters, before reaching a maximum and dissipating after 10-20 m. This dissipation distance depends on the density of the dielectric medium, and how strongly the Landau-Pomeranchuk-Migdal (LPM) effect plays a role [62]. The LPM effect causes the bremsstrahlung interaction and pair production to become suppressed, because the momentum transfer is inversely proportional to the square of the energy. For increasing energy (decreasing momentum transfer), the Heisenberg uncertainty principle dictates that the interaction must occur over many lattice

constants in the ice, resulting in deconstructive interference.

To obtain a first approximation to the radiated electric field, the authors of [59] assume that the charged particles move in straight lines, and allow for the Fraunhofer approximation: $|\mathbf{x} - \mathbf{x}_0 - \mathbf{v}t'| \approx R - \mathbf{v} \cdot \hat{\mathbf{u}}t'$. In this notation, \mathbf{x} is the observation coordinate, \mathbf{x}_0 is the position of the charged particle along its straight track, and $\mathbf{v} \cdot \hat{\mathbf{u}}t'$ is the component of the particle velocity along the Cherenkov angle (figure 1.11). The prediction from this simple model for a single negatively charged particle is that there should be two delta functions of opposite polarity that approach each other in time as the observation angle approaches the Cherenkov angle. Conversely, the larger the viewing angle, the farther apart in time the two pulses arrive at the observation point. If the viewing angle is positive (the observer sees the shower from an angle greater than the Cherenkov angle), the bi-polar pulse is delayed. The earlier arrival time in the opposite case appears to violate causality, but the fact that the charged particle is moving faster than light in the medium accounts for this.

Finally, when realistic current distributions describing the development of the charge excess are added to the model, the electric field model becomes more realistic: instead of two delta functions, we have a bipolar pulse of a given width. The width is controlled by the radial spread of charge away from the shower axis, with wider pulses coming from wider charge distributions. These main features of the Askaryan effect are depicted in figure 1.12. The pulse width (FWHM) at top left is less than 1 ns because the viewing angle is zero. As the view angle increases, the pulse widens. For electromagnetic showers, the charge excess, and therefore the electric field pulse, is stretched due to the LMP effect. The density of water is 1 g/cc (and ice is ≈ 0.92 g/cc), so these showers develop over 5-10 m, but longer if the LPM effect takes place.

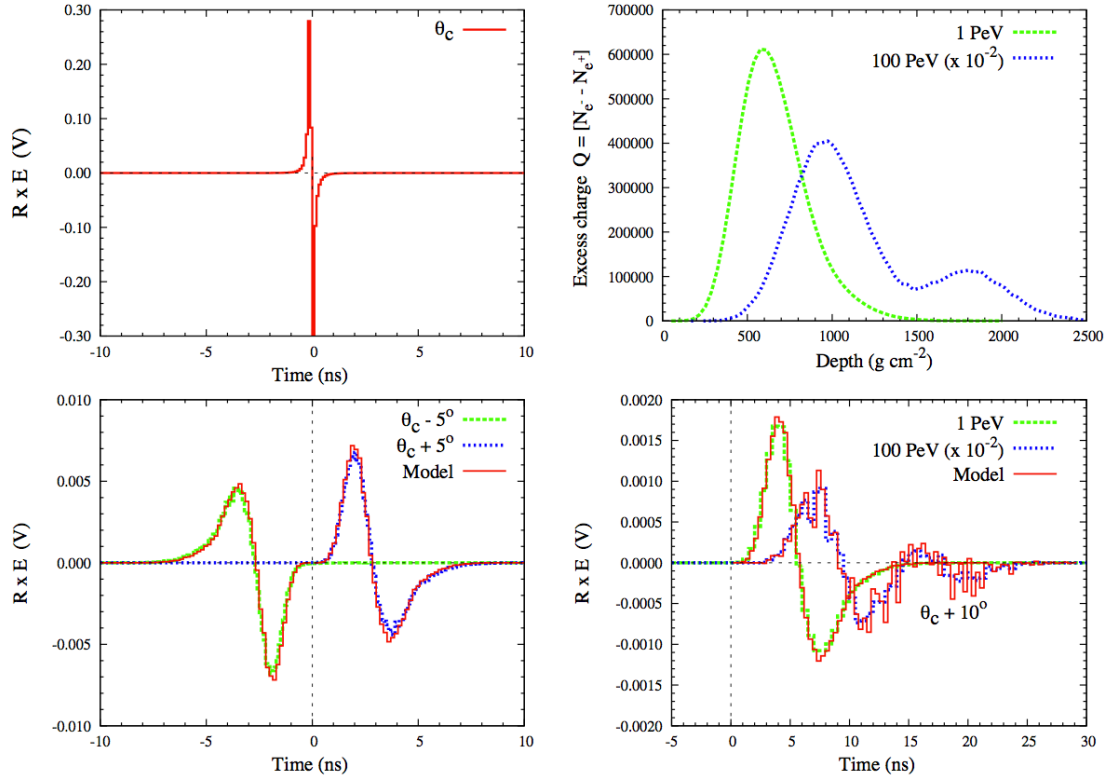


Figure 1.12: (Top left): a bipolar pulse observed 1 m away from the maximum charge excess of a shower with 1 PeV of energy [59]. (Bottom left): the same shower as the top panel, but observed at a viewing angle of $\pm 5^\circ$. (Top right): The development of the charge excess $Q(z')$ and (bottom right) the radiated electric field for a 1 PeV and 100 PeV shower.

1.3.3 Experimental Confirmations of the Askaryan Effect

There have been several experimental confirmations of the Askaryan effect in laboratory settings. Two are described here, both having taken place at the Stanford Linear Accelerator facility. The first [38] [34] utilized silica sand as the target dielectric material. Silica sand has similar RF properties as ice. Specifically, the refractive index of pure, cold ice is 1.78, and that of silica sand is 1.67. The Cherenkov angle in ice is $\approx 55^\circ$, and that of sand is $\approx 51^\circ$, at microwave frequencies, although silica is about 2.9 times as dense. The experimental setup of this initial confirmation of the Askaryan effect used GeV bremsstrahlung photons from 28.5 GeV electrons to create electromagnetic showers in the silica target (figure 1.13). Initiating the showers with photons eliminates the transition radiation caused by charged particles entering a volume with a new index of refraction. The peak electric field strength radiated from the silica varied with the lateral charge distribution, in the same fashion as $Q(z')$ in figure 1.11. The Kamata-Nishimura-Greisen (KNG) approximation predicts the charge excess in figure 1.12, and the electric field strength observed as a function the distance along the shower axis (the z -axis from figure 1.11) varies in the same fashion. In this experiment, the excess charge only propagates several hundred centimeters due to the density of silica. In addition to the effects shown in figure 1.13, the pulses are linearly polarized, and the arrival times indicate the charge excess moves at the speed of light rather than c/n . The electric field strength as a function of frequency is given by equation 1.17, where E_s is the shower energy, $f_0 = 2500$ MHz is the decoherence frequency for silica, R is the distance from the shower, and N is the overall normalization.

$$R|\mathbf{E}| = N \left(\frac{E_s}{TeV} \right) \frac{f}{f_0} \left(\frac{1}{1 + 0.4(f/f_0)^{1.44}} \right) \quad (1.17)$$

The ANITA detector served as the receiver for the next SLAC test [82], and the dielectric




Figure 1.13: (Left): A measurement of the charge profile $Q(z')$ radiating electric fields at the Cherenkov angle [38]. The inset is an example of a de-convolved radiated pulse. The amplitude of these pulses (right) matches expectations; the observed peak electric field should vary linearly with the total shower energy due to the coherence of the radiation, and the index of the observed power law is 0.95 ± 0.05 . The electric field as a function of frequency is consistent with equation 1.17.

was ice maintained at a temperature between -20 to -5 C, for minimal RF absorption (see chapter 2). In this experiment, however, 28.5 GeV leptons were not converted to photons before creating the shower, and about 15% of the negative charge excess was due to the incident electrons. Despite this systematic, and the potential for transition radiation, the Cherenkov radiation was isolated and characterized in the same manner as before. Figure 1.14 displays the data in the same form as figure 1.13, except that radiated power versus shower energy, rather than radiated electric field, is shown in figure 1.14. The radiated power follows the expected quadratic scaling with shower energy. Further work [34] was done to calculate specifically the electric field pulse in the time domain, which will be discussed below in section 3.2.2. The same model is used for the electric field strength versus frequency, and the ANITA horn data match it when accounting for the 40% systematic error. Discone and bicone antennas were included for timing and phase information.




Figure 1.14: (Left): This data is similar to the data plotted in figure 1.13 (right), except the dielectric is ice rather than silica [34]. (Right): the radiated Cherenkov power scales as the shower energy squared, again due to the coherence of the radiation from the total negative charge in the shower. The power law exponent is 2.1 ± 0.14 .

1.4 Potential New Physics with ARIANNA

In addition to solving the cosmic ray origin problem, a sample of GZK neutrinos collected in Antarctica would boost our understanding of UHE neutrino physics. This section provides two examples of physical measurements that could be done with a neutrino sample of ARIANNA. The first section outlines calculations done by Xing and Zhou (2006) [87]. The second summarizes work done by Stecker and Scully (2010) [76], demonstrating how neutrinos observed by ARIANNA could quantify an example of Lorentz-invariance violation (LIV).

1.4.1 Flavor physics of cosmogenic neutrinos

We have already discussed how the flavor ratio of GZK neutrinos begins as 1:2:0 (e, μ, τ), and how this must average to $\approx 1:1:1$ over many oscillation periods, given the 10-100 Mpc that could separate the nearest UHECR sources from Earth. A typical oscillation length calculation assumes just two flavors of neutrino, electron and one additional flavor (which can be thought of as a superposition of muon and tau flavors), with the mixing matrix [7]

$$U = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{1.18}$$

This matrix rotates any two-dimensional vector by an angle θ . In this case, the components of the initial vector are mass eigenstates, and the result is a vector with flavor eigenstates. After solving the time-dependent Schrodinger equation, and letting R be the distance travelled by an approximately massless neutrino in time t , the probability that an electron neutrino

oscillates and is detected as some other flavor is

$$|\langle \nu_x | \nu_e \rangle_t|^2 = \sin^2 2\theta \sin^2 \left(\frac{\pi R}{\lambda} \right) \quad (1.19)$$

In the above equation, λ is the oscillation length:

$$\lambda = \frac{4\pi E \hbar}{\Delta m^2 c^3} = 2.48 (E [MeV]) / (\Delta m^2 [eV^2]) m \quad (1.20)$$

Therefore the flavor most likely to be detected from an event where the neutrino began as electron flavor depends on how many factors of $\lambda < 10$ pc the neutrino travels ($E_\nu \approx 10^{20}$ eV, $\Delta m^2 \approx 10^{-3}$ eV²). For the nearest conceivable UHECR sources, these neutrinos would oscillate approximately 1 million times, and we can assume that the detected flavor distribution will reflect an average. Thus the complex, oscillatory terms that result from the solution of the time-dependent Schrodinger equation drop out, and we are left with just the numbers from the full PMNS matrix elements (with the usual three mixing angles θ_{ij} and CP-violating phase δ):

$$\phi_\beta^T = \sum_\alpha \sum_i |V_{\alpha i}|^2 |V_{\beta i}|^2 \phi_\alpha \quad (1.21)$$

In equation 1.21, α and i refer to any of the three flavors and β refers to the flavor detected at the neutrino telescope. In fact, if $|V_{\mu i}| = |V_{\tau i}|$ exactly, then the flavor pattern observed at the telescope is 1:1:1. From terrestrial oscillation experiments [87], we know that this is not

an exact symmetry, but instead this allows for

$$\phi_e^T : \phi_\mu^T : \phi_\tau^T = (1 - 2\Delta) : (1 + \Delta + \bar{\Delta}) : (1 + \Delta + \bar{\Delta}) \quad (1.22)$$

$$\Delta = \frac{1}{2} \sin^2 2\theta_{12} \sin \epsilon - \frac{1}{4} \sin 4\theta_{12} \sin \theta_{13} \cos \delta \quad (1.23)$$

$$\bar{\Delta} = (4 - \sin^2 2\theta_{12}) \sin^2 \epsilon + \quad (1.24)$$

$$(\sin^2 2\theta_{12} \sin^2 \theta_{13}) \cos^2 \delta +$$

$$\sin 4\theta_{12} \sin \epsilon \sin \theta_{13} \cos \delta$$

$$\epsilon = \theta_{23} - \pi/4 \quad (1.25)$$

Because θ_{23} is close to $\pi/4$ (39°), ϵ is a small parameter. Currently, the best fit for θ_{13} is also small, about 9° , and there is some indication from global analyses that $\delta \approx \pi$ [28]. Combining uncertainties in a global analysis, one finds that the parameters Δ and $\bar{\Delta}$ are only constrained to about $\pm 10\%$. Assuming that ARIANNA can identify the flavor of at least one type of neutrino, whether this comes from the LPM effect (electron type), or tau double-bang or tau regeneration events (tau type neutrino), it is possible to measure Δ , since the ratio of one flavor of neutrino (muon or tau) to the other two can be expressed as $R_x = 1/2 + 3/4\Delta$. If x represents electron type, then the ratio is $R_e = 1/2 - 3/4\Delta$.

Suppose the statistical error on the number of detected neutrinos in each flavor follows Poisson statistics: $\sigma_x = \sqrt{N_x}$. Here, $R_\tau = R_\tau / (R_e + R_\mu)$, and we have also $\Delta = 4/3(R_\tau - 1/2)$.

Letting $N = N_e + N_\mu$, we have

$$\frac{\sigma_\Delta}{\Delta} = \frac{(N_\tau + N(N_\tau/N)^2)^{1/2}}{N(R_\tau - 1/2)} \quad (1.26)$$

These type of analysis, however, would be extremely difficult to perform. If ARIANNA were to observe 500 neutrinos (approximately 10 Antarctic seasons' worth of data - see chapter 5), the fractional error of the Δ parameter would still be around 0.9, if the observed flavor distribution was 0.32:0.32:0.36. The fractional error of Δ is inversely proportional to the difference between the null hypothesis of R_τ (1/2) and the observed R_τ ; thus the statistics get better only if there is a large deviation from standard oscillation physics at ultra-high energies. Suppose the observed flavor breakdown is a more extreme example: 0.25:0.25:0.5, with 500 neutrinos. The fractional error is then just 18%, and gets increasingly better as the effect becomes stronger. Despite the reference above indicating δ is approximately π , however, the CP-violating phase is the least constrained parameter at the moment. This actually allows one to tune the Δ parameter to the level of ± 0.1 when tuning δ from 0 to 2π .

1.4.2 An example of Lorentz Invariance Violation characterization with ARIANNA

Scully and Stecker give examples [76] [75] of how introducing small LIV effects can alter the GZK cutoff of UHECRs, and the UHE neutrinos. One example is to introduce a different maximum achievable velocity (MAV) between the proton and the pions. In the relativistic limit that both particles are massless, each would have a velocity of $\approx c$. With a difference in

MAV $\delta_{\pi p}$, the GZK process is only kinematically allowed the following equation is satisfied.

$$\delta_{\pi p} \leq 3.23 \times 10^{-24} (\omega/\omega_0)^2 \quad (1.27)$$

In equation ??, ω is the energy of the CMB photon, ω_0 is the average energy of the CMB photon satisfying $\omega_0 = kT_{CMB} = 2.35 \times 10^{-4}$ eV, and T_{CMB} is the average CMB temperature, 2.725 ± 0.02 K. Consider a UHECR of moderate energy, 10^{17} eV, which would require some boost to observe it in the rest frame. The interacting CMB photon, which actually has an energy of less than one meV, appears to be a 300 MeV photon because of this boost. Now if we shift upwards the UHECR energy, the boost increases and the CMB photon must have a lower energy to appear the same in the proton rest frame. However, this places it below the limit of equation ??, and thus the interaction is forbidden. Another way to understand this effect is to consider that the LIV effect restricts the kinematic phase space for this interaction, such that the restriction gets stronger with increasing energy, and the inelasticity of the interaction approaches zero. With the GZK process restricted, the UHECR spectrum recovers and we observe the consequences in the UHE neutrino spectrum. Figures 1.15 and 1.16 detail the shifts in cosmic ray and neutrino spectra. Both figures come from [76]. Not only is the overall neutrino flux diminished by increasing the amount of LIV, but the average energy shifts lower as well.




Figure 1.15: The y-axis shows the energy squared times the differential cosmic ray flux, and the x-axis shows the energy of the cosmic ray. Recent data from the Pierre Auger Observatory is shown with statistical uncertainties. As the strength of LIV increases, the cosmic ray spectrum recovers more quickly. [76]




Figure 1.16: The UHE neutrino flux, multiplied by one factor of the energy, versus energy, for increasingly strong tunings of the LIV effect in question. The upper dashed curve is the simulated sensitivity of the full version of ARIANNA with a 6-month live-time, and the lower curve is the same with a five year live-time. [76]

Chapter 2

Properties of the Ice in Moore's Bay

In order to understand a potential detection of an Askaryan pulse from Antarctic ice, we must understand the dielectric properties of the ice. By transmitting short-duration (≈ 1 ns) RF pulses down through the ice and studying the reflections from the ocean-ice interface, we can probe the attenuation length of the bulk ice and the reflection coefficient of the interface over the bandwidth of interest to ARIANNA [0.1-1 GHz]. As the pulses propagate downwards through the ice shelf, they gradually lose energy due to absorption and are therefore attenuated in amplitude. This effect is in addition to the loss in energy density from the wave spreading in three dimensions (a $1/r$ effect). The wave experiences specular reflection at the ice-ocean interface, and propagates back to the surface. In general, this technique is known as radio-echo sounding, and has been used to study the properties of glaciers and ice-shelves in various locations on the Ross Ice Shelf (RIS) and the high plateau [66] [67] [27].

Results from an exploratory expedition in 2006 [40] indicate large attenuation lengths assuming a perfect reflection coefficient (100% reflected energy). During that expedition, the transmitting antenna (a Seavey radio horn used by the ANITA experiment) propagated

the RF pulse directly down, and an identical antenna served as the receiver. This single direct-bounce geometry (figure 2.1) renders indistinguishable the losses due to the reflection coefficient and absorption in the bulk ice. In order to measure these separately, a second geometry was included in addition to the direct-bounce in December 2010 [54]. By introducing a ≈ 1 km baseline between the transmitting and receiving antennas (angled-bounce geometry), we can distinguish between reflection losses and absorption losses because each geometry has a different path length, but the same reflection properties. In December 2011, a repetition of the 2010 tests was completed with the the goal of studying higher frequencies by using a shorter baseline, along with surface wave effects that could enlarge the effective volume of the experiment if confirmed [46]. In the sections that follow we will discuss measurements of the ice shelf depth and density, index of refraction, attenuation length, reflection coefficient, polarization preservation, and surface propagation properties.




Figure 2.1: (Left): The direct bounce has a path length of twice the ice shelf depth (578 m), and during the 2006 expedition this was the only geometry tested. (Right): On the 2010 and 2011 expeditions, the angled bounce geometry was added, with a 1 km baseline (2010) and a 0.5 km baseline (2011). Surface propagation effects (section 2.3) were studied by orienting the transmitter and receiver directly facing each other in the angled bounce geometry, allowing potential signals to travel directly through the firn.

2.1 Depth of the RIS in Moore’s Bay, and Experimental Geometries

The upper 60-70 m of the ice shelf is a mixture of snow and ice known as firn, which has a density at the upper layer of ≈ 0.4 g/cc (from direct measurements taken in 2009). This density result is in agreement with the result of 0.36 g/cc from [16]. The density and thus the index of refraction has an exponential form, according to the Schytt model [40] [74]:

$$\begin{aligned} n(z) &= 1.78 \equiv n && (\geq 67m) \\ n(z) &= n_0 + p \exp(-z/q) && (< 67m) \end{aligned} \tag{2.1}$$

In the model above, $n_0 = 1.86$, $p = -0.55$, and $q = 35.4$ m, assuming an upper firn density of ≈ 0.4 g/cc. In this notation, increasing the depth requires a negative value of z . Although we do not have experimental errors on the parameters in this model, a different model accounting for the firn as a constant correction was used in 2009 [31], and we obtain a depth which agrees within statistical and systematic errors (see table 2.1). In the exponential model, the exponential variation in the specific gravity (and thus refractive index) was calibrated with measurements taken at Williams Field near McMurdo station. As the depth increases, the index of refraction grows to the bulk ice value of $n = 1.78$. One possible flaw in this model is that the depth at Williams Field is around 200 m rather than the 580 m measured in Moore’s Bay, leaving open the possibility that the extent of the firn is shallower than Moore’s Bay. Given the time delay between RF pulse generation and reception, this model

can be integrated to produce the depth of the ice, $d/2$:

$$c\Delta t = \int_0^d n(z)dz \quad (2.2)$$

Using D_f for the depth of the firn, integrating leaves us with an ice depth of

$$d_{ice} = d/2 = \frac{c\Delta t}{2n} - \frac{D_f(n_0 - n)}{n} + \frac{qp}{n} (\exp(-D_f/q) - 1) \quad (2.3)$$

The second two terms are corrections associated with the firn. Assuming errors in reflection time, index of refraction, and firn depth, the error in d_{ice} is

$$\sigma_{d_{ice}} = \sqrt{\left(\frac{\sigma_t c}{2n}\right)^2 + \left(\frac{\sigma_n c \Delta t}{2n^2}\right)^2 + k\sigma_{D_f}^2} \quad (2.4)$$

The index of refraction cross terms involving σ_n and firn properties have been neglected; these terms represent small corrections compared to the second term. In the third term, $k = 2 \times 10^{-3}$, derived from the other constants. Even with an error of 10 m in the depth of the firn, the third term beneath the square root in equation 2.4 is an order of magnitude below the other terms. We are therefore justified in dropping it, to simplify and obtain

$$\sigma_d = \frac{c}{2n} \sqrt{\left(\frac{\sigma_n}{n}\right)^2 \Delta t^2 + \sigma_t^2} \quad (2.5)$$

Year	Time Delay (ns)	Calculated Depth (m)
2006	6783 ± 10	577.5 ± 10
2009	6745 ± 15	572 ± 6
2010	6772 ± 14	576 ± 3.5
2011	6816 ± 12	580 ± 3.5

Table 2.1: A summary of the RIS shelf depth data since 2006. In 2006, the data [40] was taken 1 km from the location of the Icicle1 prototype, and subsequent data is from the protostation location. A 5% error in the propagation time through all coaxial cables is assumed, and error calculations are done to the nearest nanosecond. The extra precision from the 2009 data relative to 2006 comes from assuming a more precise index of refraction (see [31]). In this work, error analysis was performed differently. However, we obtain similar errors to 2010 and 2011 if we use the 2009 numbers with equation 2.1.

This model yields an average shelf-depth over the four expeditions (2006 and 2009-11) of $d_{ice} = 576 \pm 2$ m (error in the mean, to the nearest meter). Statistical error contributions come from timing, reflected pulse shape, and the index n . The real part of the dielectric constant of ice ϵ'_i is stable between $10^6 - 10^{13}$ Hz at 3.2 ± 0.02 [52]. However, there is not good agreement on the uncertainty on this number, so here we simply assume an uncertainty of ≈ 0.01 . Using $n \approx \sqrt{\epsilon'_i}$, we have $\sigma_n \approx 0.01$ if we round to the nearest hundredth. Systematic errors include the assumption that the direct bounce path length is perfectly straight, and uncertainties in the coaxial cable lengths, both of which affect the uncertainty in the reflection time. Timing uncertainties are summarized in table 2.2, and lastly, the constant terms and the normalization of the exponential in equation 2.1 are treated as exact unless stated otherwise.

The measurement of the physical time delay caused by the ice shelf depth must be separated from the systematic delay introduced by the experimental setup. Figure 2.2 depicts the various components used in the 2010 setup which combine to produce a systematic delay. The setup in 2009 was similar, except that the pulser was an order of magnitude lower in maximum amplitude capability, making it suitable only for depth measurements and not absorption. Because the reflected pulse was only several mV above noise backgrounds, the

Year	Meas. Delay (ns)	Phys. Delay	Stat. error	Sys. error	Pulse	Total
2009	-	6745	-	-	-	15
2010	7060	6772	5.0	8.0	10	14
2011	6964	6816	4.0	5.0	10	12

Table 2.2: A summary of timing uncertainties, subtracting systematic delays. The error estimate in the reflection time for the 2010 data has been revised down from that quoted in [54]. In that work, an error calculation similar the 2009 [31] analysis was performed, which assumed the pulse width contributes to statistical errors in the reflection time. For the 2010 and 2011 data, we assume a 10 ns uncertainty based on not the width of the whole reflected pulse, but FWHM width of the strongest peaks.

authors of (NIM) had to include the entire pulse width into the timing uncertainties. For the 2010 and 2011 data, we instead take the full-width at half-maximum (FWHM) around the peak voltages in the reflected pulse. Alternatively, one could instead use the time it takes the high-frequency components (which are recorded first due to the LDPA structure) to rise above thermal backgrounds. In either case, 10 ns is a good conservative number, and when folded into the timing uncertainties we obtain smaller but comparable errors to 2009.

The Pockel Cell Driver (PCD) is capable of delivering a 1 kV (2.5 kV in 2011) negative polarity, 1-ns wide pulse to the transmitter (model HYPS PCD from Grant Applied Physics). In 2011, the HYPS PCD pulser amplitude was upgraded from 1 kV to 2.5 kV, keeping the pulse width constant. The PCD and 1 GHz scope were triggered by a TTL pulse from a BNC 555 Delay Generator, with a programmable delay. The delays used in the 2010 and 2011 seasons are quoted in table 2.2 (measured delay). The low-noise amplifier (LNA) was a Miteq 62.4 dB, 10-1000 MHz unit. The low-pass and high-pass filters (mini-circuits NLP-1200+, NHP-50+) were used for all tests to protect the data from external noise. An example of low-band noise came from the Honda portable generator used to power the electronics, and in 2010 we used the NHP-50+ to block this noise. The NHP-50+ high-pass filter has a quoted range of 41-800 MHz, a stop-band attenuation below 41 MHz of >40 dB, and a 3 dB point of 37 MHz. In the passband, the NHP-50+ introduces only a fraction of a dB in losses, and above 800 MHz the manufacturer has no data. We verified that introducing filters did

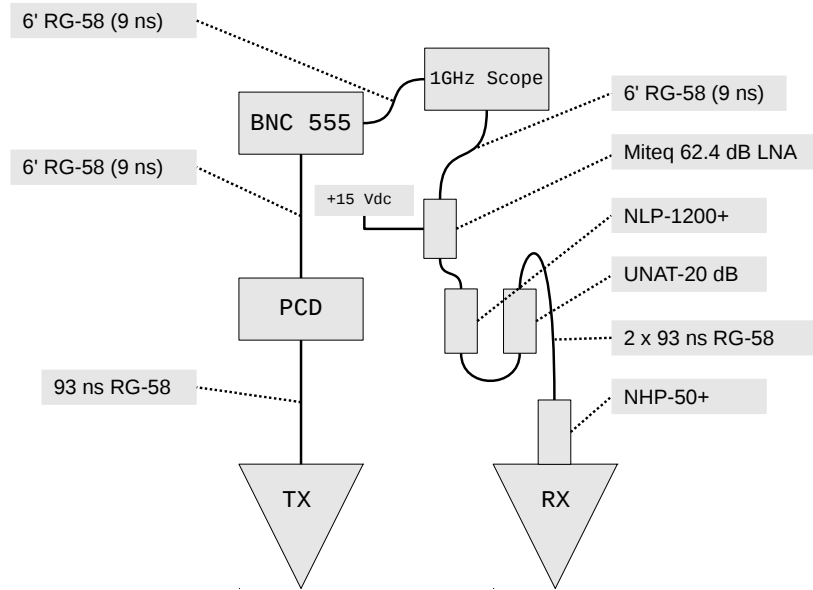


Figure 2.2: Experimental setup for reflection studies in 2010. The setup was improved in 2011, as described in the text. The pulse was generated by a Grant Applied Physics HYPS Pockel Cell Driver (PCD), and transmitted through an LPDA in 2010, and a Seavey radio horn from the ANITA experiment in 2011. When using the Seavey radio horn, we chose the polarization channel such that the transmitter was co-polarized with the receiver.

not alter the signal data, the main components being between 80 and 900 MHz. In 2011, we kept this filter in the system for consistency even though we switched away from the Honda generator to an DC-AC inverter which powered the electronics from a car battery. We were unable to identify any noise coming from the inverter at the same magnitude as the Honda generator, however we remained skeptical and chose to keep the filter in place.

The NLP-1200+ has a quoted pass band of DC-1000 MHz, where the 3 dB point is at 1200 MHz. Like the high-pass, this filter has a stop band attenuation of > 40 dB. The purpose of this filter was to block noise from wireless ethernet transmitters. For the majority of the 2010 expedition we had no internet access, so this filter was superfluous, but again it has no effect on the data. In 2011 we had much cleaner internet access, so we kept this filter in place as well. The 20 dB attenuator (UNAT-20) was used occasionally to achieve manageable signal amplitudes through the LNA. In 2010, all coaxial cables were RG-58 type (expected time

delays are shown). The transmitting (TX) and receiving (RX) antennas were separated by 18.7 m in 2010 and 23.0 m in 2011. These are identical to the station antennas (Creative Design Corp. CLP5130-2) which are log-periodic dipole arrays. In 2011, a 10+20 m length of LMR-600 coax replaced the RG-58 before the LNA on the RX side, and a 10 m length of LMR-600 replaced the RG-58 on the TX side. Finally, the TX antenna was replaced with a Seavey radio horn, identical to those used in the ANITA experiment [30]. While it has a higher low-frequency cutoff than the LPDA (200 MHz in air), the Seavey transmits more high frequency power (figure 2.4). The latter effect is so strong it is visually evident in the data, and the direct bounce pulses exhibit similar envelopes in time. The LPDA transmitter data is larger in amplitude only because this antenna has a lower low-frequency cutoff; because of the shape of the LPDA, the lowest frequencies cause the highest amplitudes in the waveform.

Examples of averaged, reflected pulses for the direct and angled bounce configurations are shown below in figure 2.3 below. In both cases, the Miteq 62.4 dB amplifier was used. The linearity of this amplifier extends to about ± 1 V, and a 3 dB attenuator was used to make the direct bounce data fit within this range, but no attenuation was required for the 543 m baseline angled bounce data. An amplitude difference of -3 dB is expected between the two signals. If we assume a 500 m attenuation length, the expected difference in amplitude is within 5% of -3 dB if we assume an attenuation length of 500 m (-1 dB from geometric losses and -2 dB from absorption). Thus applying a 3 dB attenuator to the direct bounce data makes the two pulses equal in amplitude. In 2011, we noticed that the amplitudes were changing by 10% due to several nanoseconds of jitter from the BNC 555 delay generator while it was tuned to delays $> 6 \mu\text{s}$, so we used the scope delay instead. Because this issue was not observed in 2010, we concluded that the BNC 555 delay generator was no longer matching its performance specifications.

The error reduction from the final depth measurement in 2011 was achieved by using 4 ns RG-58 and 10-20 m LMR-600 coaxial cables to send timing signals through the experimental




Figure 2.3: Averaged, reflected waveforms for the direct (upper) and angled bounce (lower). The direct bounce data has an attenuation of 3 dB, but the angled bounce has no attenuation.

apparatus , rather than using the 93 ns RG-58 cables of 2010 (see figure 2.2 and table 2.2). The LMR-600 cables also improved the high-frequency transmission of the system, aiding in the measurement of the attenuation length at frequencies above 400 MHz (see next section). Table 2.2 depicts the delay totals from 2011 for the direct bounce. To be conservative, the combination of cable length and wave-speed in the coaxial cables is assumed to be accurate only to 5% in all cases. In table 2.2, the total error in reflection time from 2009 comes from (NIM), which totals errors from cable lengths, geometry, and arrival time. The timing of the radio echo soundings is important to establish precisely, as the errors in the depth propagate further into the measurement of the attenuation length and reflection coefficient measurements.

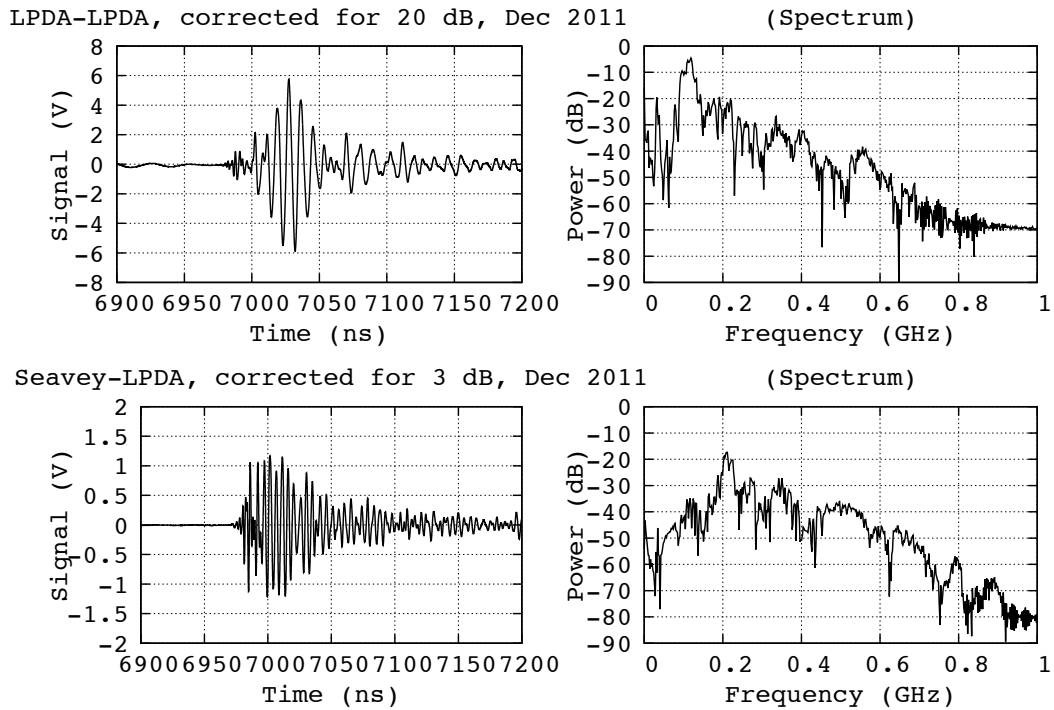


Figure 2.4: These pulses are from the direct bounce configuration, where the transmitter is alternatively the LPDA antenna (top) and the Seavey horn (bottom). Both waveforms have been corrected for attenuators used to collect this data. The y-axis units are relative, but can be compared between LPDA and Seavey.

2.2 Absorption and Reflection Properties of the Ross Ice Shelf in Moore's Bay

In this section, we derive from measurements performed across several seasons the dielectric properties of the ice which will form the fiducial volume of the ARIANNA detector. We use radio-frequency techniques common to glaciology to deduce the answers, but first we lay down the theory of electromagnetic wave propagation in lossy media. Specifically, we review a simplified model for the *attenuation length* of pure, mono-crystalline ice, and we also review the *reflection coefficient* experienced by the waves at the ocean-ice interface beneath the shelf. Following the theory, we have the experimental setup and finally the results in the last section. We conclude with a special section on surface waves and the potential benefits of surface waves to UHE neutrino detection.

2.2.1 Dielectric Properties of Ice, Seawater, and the Debye Model

Radio-glaciology is essentially the application of solutions of Maxwell's equations to glacial ice. We shall see in the following calculations how a simple model can come close to explaining the dielectric behavior of ice shelves in Antarctica. Note, however, that we will eventually have to account for unknown quantities like the concentration of impurities in the ice, and non-uniformity of crystal structure. Because it is a dielectric material, ice has a complex relative permittivity:

$$\epsilon = \epsilon' - i\epsilon'' = \epsilon'(1 - i \tan \delta) \tag{2.6}$$

where the *loss tangent* $\tan \delta = \epsilon''/\epsilon' = \sigma/\omega\epsilon_0\epsilon'$. The parameter σ is the conductivity,

which approaches infinity for conductors, making the loss tangent approach infinity. This is another way of stating that electromagnetic waves cannot penetrate conductors; the electric field inside a conductor is zero because conductors are infinitely polarized compared to dielectrics. The Ross Ice Shelf system is essentially a volume of dielectric ice above a good conductor, the seawater. In this section, ϵ_1 refers to the bulk ice and ϵ_2 refers to the ocean and will tend to be either infinite or very large compared to the ice. For example, the reflection coefficient for the electric field at the ocean-ice interface is

$$\sqrt{R} = \frac{\sqrt{\epsilon_1} - \sqrt{\epsilon_2}}{\sqrt{\epsilon_1} + \sqrt{\epsilon_2}} \quad (2.7)$$

We use \sqrt{R} for the (complex) electric field coefficient throughout the chapter because R is usually reserved for the power reflection coefficient in the literature. Taking ϵ_2 to infinity gives $\sqrt{R} = -1$, or $|\sqrt{R}| = 1$. The minus sign simply indicates a 180° phase shift, and we're usually just interested in the magnitude of the reflected electric field. The reflection coefficient of the ocean-ice interface therefore depends on how infinite the properties of conductive seawater are compared to the bulk ice. Broadly speaking, the conductivity and dielectric constants of ice are related to how energetically favorable it is to polarize the ice crystal. C.S. Neal [66] performs a calculation of the reflection coefficient based on estimates of the dielectric constants of seawater and ice. He defines the reflection coefficient (for intensity) via the *intrinsic impedance* $\eta = \sqrt{\mu_0/\epsilon_0\epsilon}$:

$$R = \left| \frac{\eta_2 - \eta_1}{\eta_2 + \eta_1} \right|^2 \quad (2.8)$$

However, neither ice or seawater is magnetic, and C.S. Neal quotes Smith and Evans [44]

for the relative dielectric constants: $3.2(1-0.007i)$ for ice, and $77(1-11.3i)$ for seawater [0° C, 60 MHz]. The corresponding electric field reflection coefficient works out to be 0.91, representing the expectation for an idea flat interface. If the measured reflection coefficient is much lower, then the interface must contain some roughness that lowers the amount of coherently reflected energy.




Figure 2.5: The real part ϵ' of the permittivity versus the frequency of the external polarizing electric field. The frequencies of interest to ARIANNA are strictly in the flat region where the frequency is so far from $1/\tau$ at 0° (8.3 kHz) that ϵ' is dominated by the high-frequency constant value $\epsilon'_\infty = 3.2 \pm 0.02$. [52].

The polarization of ice crystal originates from orientational defects associated with proton (H^+) transitions. The energy is provided by the external time-varying electric field, and rate of transitions follow a Boltzmann distribution, with probability $\propto \exp(-E_a/k_B T)$. E_a is the activation energy (typically between 0.1 and 1 eV), T is temperature in Kelvin, and k_B is Boltzmann's constant. We define a relaxation time, τ , which is inversely proportional to the transition rate. The temperature dependence in τ plays a role in the formulation of the attenuation length (below), however the frequencies interesting to neutrino physics ($10^8 - 10^9$ Hz) are far from any $1/\tau$ values corresponding to quantum mechanical rotational or vibrational modes of the ice molecules (figure 2.5). Given the relaxation time, the high

frequency and static (DC) limiting values ϵ'_∞ and ϵ'_s , with $\Delta\epsilon = \epsilon'_s - \epsilon'_\infty$, the Debye model gives

$$\epsilon/\epsilon_0 = \epsilon'_\infty + \frac{\Delta\epsilon}{1 + i\omega\tau} \quad (2.9)$$

The quantity $\Delta\epsilon$ is known as the *dispersion strength*. This equation implies a loss tangent of

$$\tan \delta = \Delta\epsilon\omega\tau/(\epsilon'_s + \epsilon'_\infty(\omega\tau)^2) \approx \frac{\Delta\epsilon}{\epsilon'_\infty}(\omega\tau)^{-1} \quad (2.10)$$

The approximation holds because $(\omega\tau)^2 \gg 1$ in the frequency range of ARIANNA (equivalently, $\omega \gg 1/\tau$). Notice that in this version of the Debye model, with a single relaxation time under the given approximation, the quantity $f \tan \delta$ is a *constant*. This is only true to first order, as there could be other frequency dependent effects in the ice. However, the expectation is, with $f \tan \delta$ roughly constant, the attenuating properties of Antarctic ice will be flat with respect to frequency. Next, we can obtain the loss in dB/m due to the complex dielectric constant [73] (and index of refraction), and we will see specifically how the temperature affects the results.

Consider an electromagnetic plane wave with wavenumber k propagating through a dielectric with frequency ω :

$$\mathbf{E} = \mathbf{E}_0 \exp\{i(kx - \omega t)\} \quad (2.11)$$

The complex index of refraction n is defined in the same way as ϵ : $n = n' + in''$. In the ice, the electromagnetic wavelength elongates: $k \rightarrow nk$, and as a result we have

$$\mathbf{E} = \mathbf{E}_0 \exp\{i(n'kx - \omega t) - n''kx\} \quad (2.12)$$

which implies that $1/n''k$ is the distance the wave travels before the amplitude drops by a factor of e . The index of refraction is related to the dielectric constant in the following way:

$$n = \sqrt{\epsilon} = \sqrt{\epsilon' - i\epsilon''} \approx \sqrt{\epsilon'}(1 - i/2 \tan \delta) \quad (2.13)$$

in the limit that the loss tangent is much smaller than unity. Again, this is a good approximation because we are far from any frequencies close to $1/\tau$. Thus, $n'' = -(\sqrt{\epsilon'}/2) \tan \delta$ and we can define the loss (in factors of e) per meter in amplitude as the *attenuation constant*:

$$\alpha = n''k = -(k\sqrt{\epsilon'}/2) \tan \delta = -(\pi\nu/c)\sqrt{\epsilon'} \tan \delta \quad (2.14)$$

Multiplying by a factor of 2 produces the power loss (in factors of e) per meter, rather than amplitude. Affixing the pre-factor $10\log(e) = 4.34$ changes the units from factors of e per m to dB/m:

$$\alpha = -(2 \cdot 4.34\pi\nu/c)\sqrt{\epsilon'} \tan \delta = -8.68(\pi\nu/c)\sqrt{\epsilon'} \tan \delta \quad (dB/m) \quad (2.15)$$

Substituting equation for the loss tangent, we are left with

$$\alpha = -8.68 \left(\frac{1}{2c} \right) \frac{\Delta\epsilon}{\sqrt{\epsilon'_{\infty}}} \tau^{-1} \quad (dB/m) \quad (2.16)$$

Recall that τ is inversely proportional to a molecular transition rate which follows the Boltzmann distribution: $\tau = A \exp(E_a/k_B T)$, and A is just a normalization with units of seconds. We have $\tau = 2 \times 10^{-5}$ seconds at 0° C, with $E_a = 0.58$ eV [52]. Solving for A , we get $A = 4.1 \times 10^{-16}$ sec. Substituting this functional form into equation 2.16, we arrive at the result for the dielectric loss as a function temperature:

$$\begin{aligned} \alpha &= -8.36 \left(\frac{1}{2cA} \right) \frac{\Delta\epsilon}{n} \exp(-E_a/k_B T) \quad (dB/m) \\ &\approx -18 \times 10^8 \exp(-E_a/k_B T) \quad (dB/m) \end{aligned} \quad (2.17)$$

Equation 2.17 demonstrates a qualitative rule about the dielectric losses in the ice shelf: radio-frequency propagation is exponentially better in colder ice. Assuming no other frequency dependencies, the α works out to ≈ -15 dB/km at 260 K. We have substituted the index of refraction for RF frequencies $n = \sqrt{\epsilon'_{\infty}}$. The quantity $\Delta\epsilon/\sqrt{\epsilon'_{\infty}} = \Delta\epsilon/n$ has a value of ≈ 53 , but depends on temperature because ϵ'_s increases with decreasing temperature. The attenuation constant is convenient for calculating losses in dB over a given distance. Conversely, the *attenuation length*

$$L_{\alpha} = 1/\ln 10^{\alpha/20} \quad (2.18)$$

is equal to the length in meters over which the electric field amplitude drops by one factor of e . It is this quantity that is derived from the data below, with the expectation that it will only vary weakly with frequency. Using the answer obtained above for α (in [31] the average ice temperature is assumed to be $\approx -10^\circ$ C), we see that the attenuation length should be about 600 m for pure mono-crystalline ice at that temperature.

We shall see that this naive estimate is close to the actual results below, however we must note that the temperature dependence is not actually as strong as the basic Debye model for pure mono-crystalline ice. The actual temperature profile of ice at the South pole we have from the AMANDA and IceCube experimental efforts [25]. The authors of [36] assume a model for the temperature dependence of the attenuation constant from Matsuoka et al. [41], obtaining a result for the full depth-dependent electric field attenuation length. Using the temperature versus depth data from [25] the attenuation length versus temperature information can be extracted from the South pole data.

Figure 2.6 shows the resulting data fitted with equation 2.18, with equation 2.17 substituted in for the attenuation factor. We find that the activation energy E_a is closer to 0.24 eV, rather than 0.58 eV from the pure mono-crystalline model. Model parameters that differ from the simple Debye model are likely indicating factors such as ice impurity and varying crystal orientation. The all-depth average for the ice underneath the Askaryan Radio Array (820_{-65}^{120} m) is shown in blue [36], and the same measurement at the same frequency is shown for ARIANNA (440_{-30}^{30} m) in red [54] [46]. The average ice temperature (over depth) at the South Pole is near -34.7 C, which brings up the all-depth attenuation length average. Note that the average is far from the long attenuation lengths near the surface ice (at colder temperatures) because signals do not solely experience the coldest temperatures. Additionally, the ARIANNA number (from further analysis below) is from a smoothed fit to the data; the single frequency 300 MHz measurement is actually 530_{+30}^{-30} m. This value is representative of the maximum attenuation lengths experienced by ARIANNA, especially at lower frequencies

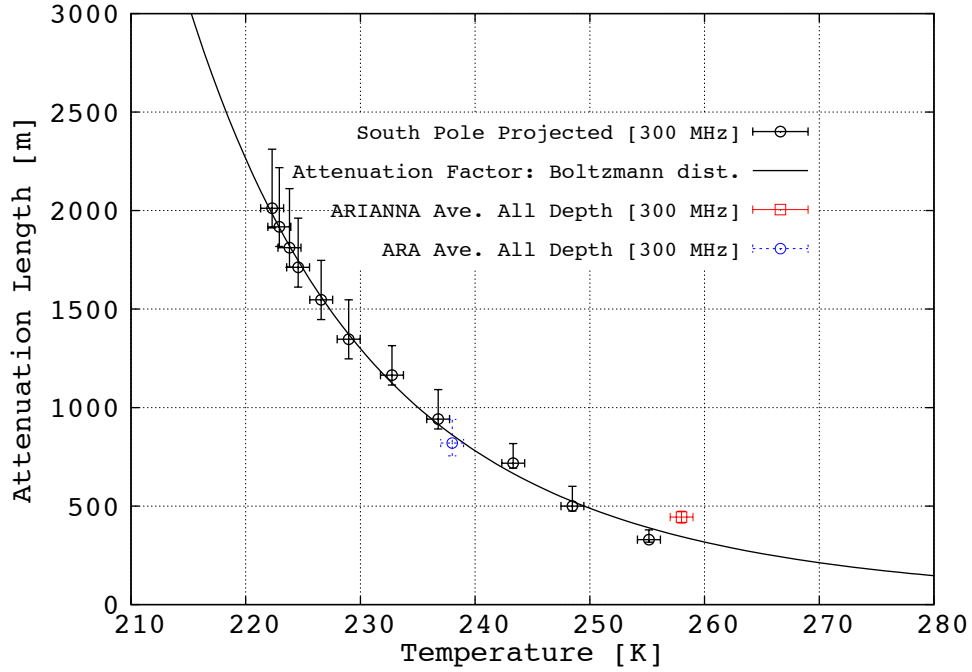


Figure 2.6: The ARIANNA 300 MHz all-depth attenuation length average is placed at -15 C. The Askaryan Radio Array result is also taken at 300 MHz. The black data is the projected South Pole results using a simple Debye model, with adjusted fit parameters (for example, $E_a = 0.24$ eV) and temperature data from [25]. The black data are not a pure measurement.

inaccessible to ARA [100-200] MHz.

2.2.2 Experimental Technique and Antenna Equations

The experimental setup for the attenuation length and reflection coefficient measurements is the same as for the ice depth measurement. The angled-bounce baseline between transmitter and receiver was determined from a handheld GPS tracker to be 977 ± 7 m, making the angled bounce path length 1517 ± 8 m. In the 2011 season, the baseline was adjusted to 543 ± 7 m, for an angled path length of 1272 ± 7 m. A longer baseline allows us to study ice absorption at low frequencies, where the difference between the direct and angled path lengths is comparable to one factor of the attenuation length around 100 MHz (≈ 500 m). The same difference for the case of the shorter baseline is less than one factor of the attenuation length (300-400 m) for frequencies greater than 500 MHz, which allows us to probe the high frequency absorption properties of the ice shelf.

To obtain a calibration signal for the unattenuated, non-reflected power, the receiver and transmitter were rotated to face each other 1 m above the surface, separated by 19 ± 0.02 m and 23.0 ± 0.02 m in 2010 and 2011 respectively. These distances are in the far field of the LPDA/Seavey antennas, where the far-field criterion is the usual $d \gg 2D^2/\lambda \approx 1.5$ m. The amplifier was removed because it isn't necessary at such short distances through air, however, attenuators are still required to keep the received signals in the linear range of the scope. The gain of the Miteq amplifier (62.5 dB) is essentially flat, so removing it does not change the response of the system. An example of an averaged, Seavey to LPDA pulse is shown in figure 2.7 below. The analysis of the December 2010 data focuses on the 90-180 MHz bandwidth, where the measured power was well above thermal noise for the three antenna configurations. An RF noise source verified that cable losses were negligible in this range. The high frequency strategy implemented in 2011 allowed us to probe simultaneously the attenuation length and reflection coefficient in the bandwidth 325-750 MHz, frequencies where there were both measurable losses from attenuation and enough power above thermal noise backgrounds to make a good measurement.




Figure 2.7: This is an averaged signal taken during the 2011 Moore’s Bay expedition, where the transmitter is the Seavey radio horn and the receiver is an LPDA. There is 3 dB of attenuation and the distance is 23.0 m through the air, with no amplifier.

In December 2010, the direct-bounce and angled-bounce geometries had identical transmitting and receiving antennas, oriented downwards in both cases. The signal was generated by a 1 kilovolt, 1 ns wide negative polarity pulse from a HYPSPockel-cell driver in 2010, and in 2011 a new unit provided 2.5 kV. A 17 element, 6 dBi gain log-periodic dipole array (LPDA) antenna was chosen (Creative Design Corp. CLP5130-2) for its relatively uniform frequency characteristics and low VSWR in air and snow. This is the same antenna model used in the prototype detector *Icicle1*. Although the lowest frequency quoted for this antenna is 105 MHz, it becomes sensitive down to 80 MHz in the snow because the index of refraction of the upper firn at Moore’s bay is $n \approx 1.3$, which lowers the wave speed ($k \rightarrow nk$ in equation 2.12). Formally, the $VSWR = (1 + \rho)/(1 - \rho)$, where ρ is the reflection coefficient for a signal traveling through the transmission line toward the TX antenna. The VSWR quantifies the amount of energy that is radiated, and which frequencies are not radiated by the antenna/dielectric system. If the $VSWR \gg 1$, this indicates that electromagnetic energy is reflected back to the transmitter, rather than radiated, at that frequency. In fact,

using this simple picture of the antenna radiation, we can derive the fractional change in the index of refraction using $105 \text{ MHz} / 80 \text{ MHz} = 1.3$. Assuming $n_{air} = 1$ exactly, we have obtained $n = 1.3$ for the firm surface, which is the answer we obtain if we take $z \rightarrow 0$ in the model for $n(z)$ above.

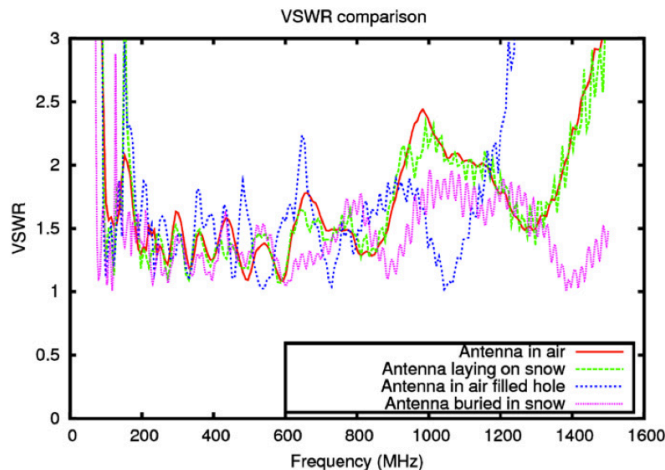


Figure 2.8: The voltage standing wave ratio (VSWR) for the ARIANNA LPDA in various positions on the surface near the prototype location [31]

The VSWR is close to one for all relevant frequencies in this work (figures 2.8 and 2.9). The returned signal from the receiver was filtered with both a NHP-50+ high-pass filter and a NLP-1200+ low-pass filter, and amplified by a 1 GHz Miteq AM-1660 low noise amplifier (with a noise figure $< 1.5 \text{ dB}$ and gain 62.5 dB) before being recorded on a 1 GHz bandwidth oscilloscope. To obtain manageable amplitudes, we attenuated the signal with Mini-circuits coaxial attenuators (as needed). The oscilloscope and HYPS Pockel cell driver were both triggered via a BNC 555 delay generator in 2010, such that the physical delay in table 2.2 could be measured. In 2011, an upgraded scope had enough delay to allow us to zero out the the delay from the BNC 555, decreasing the jitter. In the angled-bounce geometries both years, $977 \pm 7 \text{ m}$ (2010) and $543 \pm 7 \text{ m}$ (2011) RG8X coaxial cables carried the signal from the BNC 555 to trigger the oscilloscope.

Once the signal leaves the TX antenna, geometric spreading and absorption begin to sap

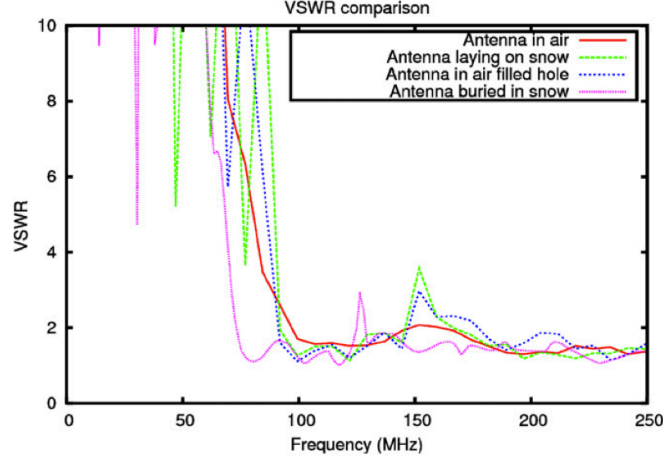


Figure 2.9: The same data as figure 2.8, but zoomed in on the low-frequency modes, revealing the $k \rightarrow nk$ effect, where the lower bound to the LPDA bandwidth shifts upon being buried in the snow [31].

energy density from the wave. The Friis equation calculates the power received P_r in a lossless medium at a given wavelength. For two identical antennas, separated by a distance d with intrinsic gain G_a , it becomes

$$P_r = \left(\frac{G_a \lambda}{4\pi d} \right)^2 = P_0 / d^2 \quad (2.19)$$

In the basic Friis formula, the only loss in power is due to the geometric spreading of the wave in three dimensions. The intrinsic G_a is related to the beam width (Appendix A), and λ is the wavelength. The variation in LPDA intrinsic gain is small for the frequency interval of this study. P_0 was determined from the short distance configuration with the antennas rotated to point toward each other ($d = 19.0 \pm 0.02$ m in 2010 and 23.0 ± 0.02 m in 2011). To account for absorption losses and possible losses upon reflection, the Friis equation is

modified to

$$P_r = \frac{P_0}{d^2} (RG^2) \exp\left(-2\frac{d}{\langle L_\alpha \rangle}\right) \quad (2.20)$$

The factor R is the reflection coefficient, defined for power, making the reflection coefficient for the electric field \sqrt{R} . The factor of 2 is required in the exponential if L_α is to be the attenuation length of the electric field and not the power. The brackets around the attenuation length indicate that the experiment probes the depth-averaged attenuation length, with the understanding that the ice temperature varies from -35° at the surface (during the winter months) to the freezing point of seawater (-2° C) at the bottom. The actual attenuation length experienced by the neutrino signals will vary with depth and therefore temperature according to equation 2.18. The factor G^2 accounts for the difference in relative antenna gain if the received signal is slightly off-axis. That is, for a single antenna, $G = G_a(\theta)/G_a(0)$, so $G \in [0, 1]$. We must assume that the transmitted angle is equal to the receiving angle in this case, but this is not an unreasonable assumption given the experimental geometry. G is close to 1 for the direct bounce tests, but G must be quantified for the angled bounce tests. In 2010, at an angle of 49.9° in the H-plane of the LPDA, we found $G = 0.80$. Similar numbers for 2011 are discussed below.

Equation 2.20 can be linearized such that the slope is $-1/\langle L_\alpha \rangle$:

$$\ln \sqrt{\frac{d^2 P_r}{RG^2}} = \ln \sqrt{P_0} - d/\langle L_\alpha \rangle \quad (2.21)$$

There are **three** cases we examine: the calibration (in air) data, and the direct and angled

bounces. For the case of the calibration pulse, the data is neither reflected or attenuated, so we have

$$\ln \sqrt{\frac{d^2 P_r}{1}} = \ln \sqrt{P_0} \quad (2.22)$$

Thus the first term in equation 2.21 can be solved for, and the remaining two unknowns (\sqrt{R} and $\langle L_\alpha \rangle$) can be solved for using the results from the direct and angled bounces (two equations and two unknowns). In practice, we simply plot the left hand side of equation 2.21 versus d for the **three** cases (each with its own path length) and perform a linear fit. Once we chose a hypothetical value for \sqrt{R} , we fit the trend line and assess the goodness of fit. By scanning over a range of \sqrt{R} values, we find the best possible fit. The goodness of fit is quantified by the χ^2 value with one degree of freedom:

$$\chi^2 = \sum_{i=1}^3 \left(\frac{X_i - X_{fit}}{\sigma_i} \right)^2 \quad (2.23)$$

There is only one degree of freedom because although the constant in the linear fit is constrained by the calibration pulse, and we have three data points, we're solving for two parameters (\sqrt{R} and $\langle L_\alpha \rangle$). The values of χ^2 with one degree of freedom are 1.07 and 3.84 for the 68% and 95% confidence limits, respectively. Contour plots in the next section will show the results for these confidence limits. The errors in the denominator of the χ^2 equation are obtained through error propagation of the expression on the left hand side of equation

2.21:

$$\sigma_i = \sqrt{(\sigma_{d_i}/d_i)^2 + (\sigma_{P_i}/P_i)^2}/4 \quad (2.24)$$

The statistical errors in the power come from the standard deviation of the power in the frequency bin, divided by the square root of the number of measurements in the bin. The raw power spectra are plotted below in the next section, where we find that it is easiest to average over a range of frequencies, because we are trying to reveal small effects (i.e. the difference between losses from absorption and reflection). The errors in the path lengths are known from section 2.1, where we study the depth of the ice shelf. The calibration pulse distances were 19 ± 0.02 m and 23 ± 0.02 m for the 2010 and 2011 setups, respectively.

In the 2011 season, we attempted to align the antennas during the angled-bounce test such that the received signal was on the bore-sight of the LPDA. However, this proved difficult given the gradient in the index of refraction of the upper firn and the lack of resolution of antenna angle with respect to the true horizontal. Ultimately, we estimate that the G factor corresponds to an off-axis angle of $\approx 10^\circ$. In 2011, we reverted to the Seavey radio horn transmitter in order to broadcast more power at higher frequencies. The Seavey radio horn has an average gain of 10 dBi and thus a narrower beam angle than the LPDA, such that $G = 0.82$ (at 10° off-axis) similar to the 2010 season. Ultimately this is still an improvement, since we were able to achieve a similar G factor with a TX antenna with a higher intrinsic gain at high frequencies ($G_a = 6 - 7$ dBi for the LPDA). The added value of using the Seavey transmitter in addition to the higher gain is the increase in radiated power at higher frequencies. The measurements were also performed when the ambient temperature was close to -10° , rather than near or above 0° C as in 2010. This eliminated a ≈ 20 dB loss at frequencies above ≈ 200 MHz due to liquid water content in the snow, present because of

unseasonably warm temperatures.

2.2.3 Results

Figures 2.10 and 2.11 presents typical averaged waveforms acquired with this technique during the December 2010 and December 2011 expeditions. We use amplifiers and attenuators where necessary to record voltage amplitudes that are within the linear range of our instrumentation. The linear range of the Miteq amplifier is ≈ 1 V, and the oscilloscopes used in these experiments remain linear to ± 10 V. Though some of this data has been shown already in section 2.1, it is worth displaying it again here to remind the reader that the powers received and transmitted in equation 2.20 come from reflected time-dependent waveforms. Because the 2010 setup used a 977 m baseline, the angled bounce data experienced additional factors of the attenuation length and a longer path length. However, the 2011 data used a 543 m baseline, meaning the attenuation length and path length effects were more muted.

The raw averaged power spectra from the time-dependent waveforms are shown in figure 2.12. These figures are constructed from averaging the modulus-squared of the FFT of the time-dependent waveforms. The errors are the standard deviations of power in each bin, divided by the square root of the number of measurements in each bin. There are several important features to discuss. The first is that the spectra obtained in 2010, including the calibration pulse in snow show heavy attenuation above 200 MHz. Because this effect is observed for the calibration pulse (in snow) as well as the direct and angled bounce cases, we conclude that the environment caused it. One simple explanation is that because the ambient temperature was near 0° C, there was significant liquid water content in the top layer of snow that caused high-frequency attenuation. In the 2011 data, the ambient temperature was closer to -10° C due to an earlier November deployment, and the effect has disappeared (see figures 2.12, and 2.14 below).

The 2011 data demonstrates the correct trend in the band [325-750] MHz, where the angled




Figure 2.10: These waveforms are typical of the 2010 setup and are averaged over successive Pockel Cell pulses. The separation between transmitter and receiver for the calibration pulse (left) was 19 m. The direct bounce (middle) was amplified with the Miteq LNA, and required 100 averages. The angled bounce (right) was also amplified and required 1000, due to the smaller amplitudes involved. Here we have corrected the amplitudes for the attenuation, but not the amplification.

bounce is less powerful than the direct bounce. At high frequencies, noise plays too large a role and we cannot distinguish the two geometries. Our expectation at low frequencies is that the two spectra not exhibit much difference. The difference in path length is only 10%, and the expected attenuation lengths are near 500 m at frequencies below 325 MHz. There are, however, several data points from the angled bounce spectrum that are larger by $\approx 1\sigma$. While this effect is not fully understood, it could be that we introduced it when we attempted to align the receiver and transmitter such that the receiver observed the reflection at bore-sight.

The firm and bulk ice have different indexes of refraction (1.3 and 1.78, respectively), and the baseline was 543 m. Thus, we made an attempt to angle the receiver and transmitter




Figure 2.11: These waveforms are typical of the 2011 setup. The separation between transmitter and receiver for the calibration pulse (left) was 23 m. The direct bounce (middle) and angled bounce (right) required 100 averages each. Here we have corrected the amplitudes for the attenuation, but not the amplification.

toward each other (approximately 35° from vertical) so the refracted signal would arrive at receiver bore-sight. However, to angle the antennas, we had to dig deep trenches rather than small slots in the snow, as for the direct bounce. In 2010, we made no attempt to angle the receiver and transmitter, and therefore used slots as well. The trenches may have acted as resonators that caused constructive interference at low frequencies. The frequencies are 200-300 MHz, corresponding to ≈ 1 m wavelengths in air, the same width and depth as the trenches. In any case, figure 2.12 demonstrates how the bandwidth of [325-750] MHz demonstrates the expected attenuation for the experimental geometry used, regardless of systematics like constructive interference.

Figure 2.13 show the low and high frequency results for fitting simultaneously the attenuation length and reflection coefficients to equation 2.21. The statistical errors in the received power were obtained from the rms fluctuation over the entire frequency band. The 1 and




Figure 2.12: These data are the reflection data from the 2010 (left) and 2011 (right) expeditions. The 2010 calibration pulse data (black) was taken at 19 m, followed by the direct (red) and angled (blue) bounces (baseline: 977m). The 2011 calibration pulse data was taken at 23 m, followed by the direct (red) and angled (blue) bounces (baseline: 543 m).

$2\text{-}\sigma$ contours in these plots were obtained from the chi-squared fit. The contours match the dashed curves, which are derived by varying the reflection coefficient \sqrt{R} assumed in [40] for the quoted attenuation lengths and frequencies. Specifically, the 2010 attenuation length [80-180 MHz] is determined to be $480 \leq \langle L_\alpha \rangle \leq 510$ m at 68% C.L., in agreement with previous values from 2006, and the field reflection coefficient $0.72 \leq \sqrt{R} \leq 0.88$ at 68% C.L., close to the theoretical expectation of 0.91 from the previous section. The dashed curve represents the data from [40] at 100 MHz.

In 2011, we obtained $385 \leq \langle L_\alpha \rangle \leq 440$ m at 68 % C.L., and $0.70 \leq \sqrt{R} \leq 1.0$ at 68% C.L. in the band [325-750 MHz]. Our measurements indicate that the electric field reflection coefficient at the ARIANNA site is compatible with a flat surface, although values as small as 0.7 are also permitted by the statistics at the 68% C.L. The results lie in between the results from table 1 of [40] at 400 and 800 MHz, and matches an interpolation to 530 MHz (the middle of the bandwidth). The reflection coefficient is not expected to vary significantly with frequency for specular reflection. If the permittivity of seawater is assumed to be very large for 100-1000 MHz, then the field reflection coefficient approaches unity with no frequency

dependence [16]. In this reference, the authors claim that up to 800 MHz, the dielectric constant of water is large enough to assume 100% reflection.

Finally, because the duration of the both the transmitted and received pulses are less than 100 ns ($\approx 1\%$ of the reflection time) multi-path interference can be safely ignored. Comparing the time-dependent reflected waveforms to the calibration pulses, it is clear that the effect driving the spreading of the 1 ns, 2.5 kV PCD pulse is the antennas, and not the ice or the ice-ocean interface. If there were several scattering surfaces at the bottom of the shelf near the ocean, separated by $O(10)$ m, then we would observe time-differences of $O(100)$ ns in the arriving signals given the index of the bulk ice near the ocean. However, this is not the case, and thus we still infer a flat, specularly-reflecting surface from the data.




Figure 2.13: (Left): the lower frequency results for the electric field reflection coefficient and attenuation length from 2010. (Right): the higher frequency result from 2011. Both reflection coefficient answers encompass the expectation for specular reflection [66] [67].

There are several bandwidth considerations to be discussed. The bandwidth of the 2010 data is restricted to frequencies above 80 MHz simply because the TX and RX antennas in that case were LPDAs, in snow. This lower limit is explained by the VSWR measurements in figures 2.8 and 2.9 above. The upper limit of 180 MHz comes from the strange effect of the snow on high frequency RF propagation. The *average* temperature for the 2010 expedition (December 8th through 27th, 2010) was near 0° C. We hypothesize that there

was a significant mixture of liquid water in the upper firn snow, and the evidence for this is shown in figure ?? below. The TX and RX antennas were positioned first in air (18.7 m apart), and then in trenches cut in the snow just larger than the dimensions of the antennas, which were co-polarized with each other. The rest of the experimental setup was left untouched.

We observe a ≈ -20 dB shift in the power spectrum of the transmitted pulses as they propagate through the snow and through the receiver system, for frequencies above 180 MHz. In each measurement, the antennas in each case were situated $\approx \lambda/2$ meters above/beneath the snow, referring to the largest wavelength of the LPDA. We also see an effect between 400 and 500 MHz (wavelengths of 50-75 cm in air) in 2011, but this is likely due to the snow cavity rather than the snow itself, since it's absent between 600 MHz and thermal noise at 900 MHz. This effect is irrelevant to the 2011 absorption and reflection data since the Seavey horn was used as the TX antenna. The hypothesized liquid water effect observed in 2010 is gone from the data in 2011 (figure 2.15). The average temperature was lower by $\approx 10^\circ$ C on the 2011 expedition (mid-November 2011 to mid-December 2011), because we began the expedition earlier in the year. Since there is no snow effect in either polarization in 2011, the bandwidth is not restricted due to these effects.

During the 2011 expedition, we attempted to observe more high-frequency information in the radio echoes by shortening the path length. The trade-off is that the difference between the angled bounce path length and the direct bounce path length is only ≈ 130 m. This means that the attenuation length has to be *small enough* to observe differences between angled and direct signal paths. Because the attenuation length is larger at lower frequencies, this moves the lower frequency bound up to 325 MHz (see figure 2.16), where the angled bounce absorption is statistically different from the direct path. Above 750 MHz, the angled and direct paths are too close to thermal backgrounds to observe differences. Figure 2.16 represents the data used to make the attenuation length vs. reflection coefficient measurement with

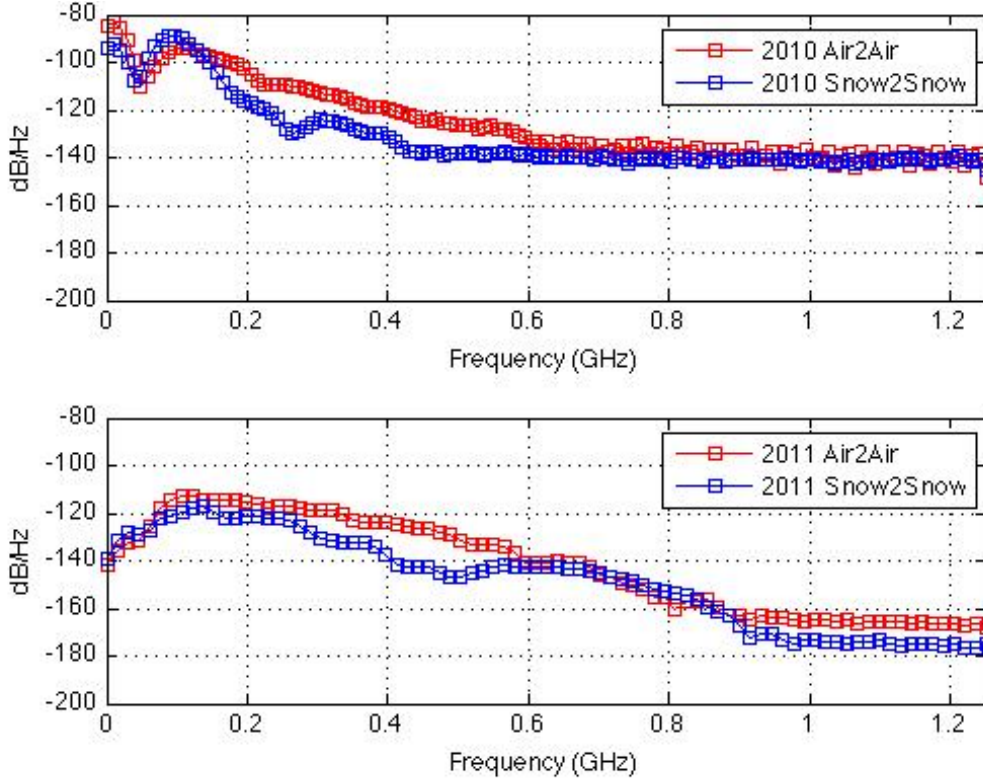


Figure 2.14: The transmitted pulse power spectra through the LPDA-LPDA system, in air and snow in Moore’s Bay (2010 and 2011). The 2010 data shows a 20 dB effect at 180 MHz and continuing until the signals drop below thermal noise.

the reduced baseline of 2011. The x-axis is the frequency in GHz, and the y-axis is the difference in dB between the direct and angled bounce power spectra. The region where we observe a difference that is inconsistent with 0 dB is the one chosen for the reflection coefficient/absorption analysis.

The analysis associated with figures 2.10-2.16 are a demonstration of how to measure the reflective and absorptive properties of the RIS simultaneously. That technique is analytically more difficult and harder to interpret than simply assuming a reflection coefficient and calculating $\langle L_\alpha \rangle$ versus frequency. Since the reflection coefficient results are consistent at low and high frequencies, it’s not unimaginable that the scale of surface roughness of the ice shelf/ocean interface are small enough to assume an ideal flat surface. If we make this assumption, then we can derive the attenuation length independently:

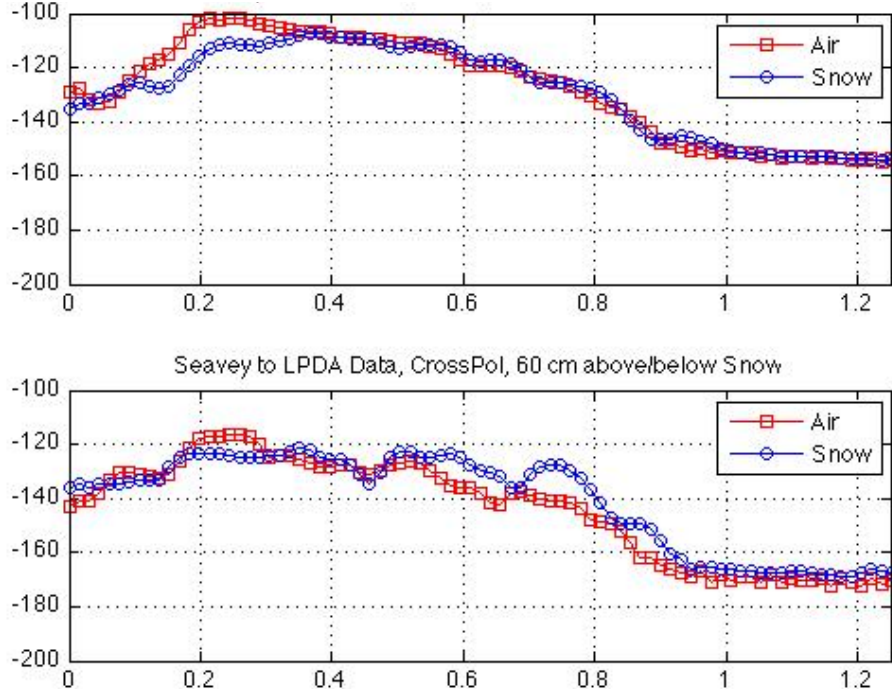


Figure 2.15: These 2011 power spectra are similar to the figure above, with the same units, however the TX antenna in this case is the Seavey radio horn. The upper graph depicts the co-polarized state, and the lower graph the cross-polarized state.

$$V_{air} \propto 1/d_{air} \quad (2.25)$$

$$V_{ice} \propto 1/d_{ice} \exp(-d_{ice}/\langle L_{\alpha} \rangle) \quad (2.26)$$

$$\langle L_{\alpha} \rangle = \frac{d_{ice}}{\ln((V_{air}d_{air})/(V_{ice}d_{ice}))} \quad (2.27)$$

In these equations, V is just the received amplitude at a given frequency, proportional to the square root of the power spectrum. In this section of the analysis, Welch's method [85] was chosen to calculate the power spectra for several reasons. Welch's method, which is similar to the periodogram method, provides a calculable improvement in the fractional




Figure 2.16: The x-axis is the frequency of the reflected pulses on the short (543 m) baseline, and the y-axis is the ratio (expressed as a difference in dB) of angled and direct bounce power. The red dashed line represents an estimate in Δ (dB) for an attenuation length of 500 m.

error in the power spectrum relative to simply squaring the FFT and averaging. This serves to suppress noise and non-physical spikes in the power spectrum that lead to errors in the attenuation length. The trade-off is in frequency resolution, which is minimal. The second advantage of Welch’s method is that the results are smoothed according to a controllable overlap between data points, serving to make the power spectrum more continuous. If large, physical trends are present (with respect to frequency) this smoothing can be set to zero. In this section, d_{ice} is the total path length and not the shelf depth, and the attenuation length is again weighted by the ice temperature at a given depth as in equation 2.18 for the attenuation constant.

Figure 2.17 (left) assumes $\sqrt{R} = 1.0$, as in [40], and figure 2.17 (right) assumes $\sqrt{R} = 0.85$ for purposes of comparison, and the 2006 data is left unscaled in both graphs. A first-order polynomial describing the 2011 data is $L = (550 \pm 40) - (0.27 \pm 0.09) \times f[\text{MHz}]$, assuming 100% reflection, and both the slope and the intercept vary by less than one standard deviation

in varying \sqrt{R} from 100% to 85%: $L = (590 \pm 50) - (0.3 \pm 0.1) \times f$ [MHz]. The errors on the attenuation length are $\pm 1\sigma$ and include statistical fluctuations in the power over the frequency bins (75 MHz wide), and systematic errors due to uncertainties in the ice-shelf depth. The errors from the 2006 data are computed in the exact same way, except the bin width is 25 MHz, giving a larger statistical error.

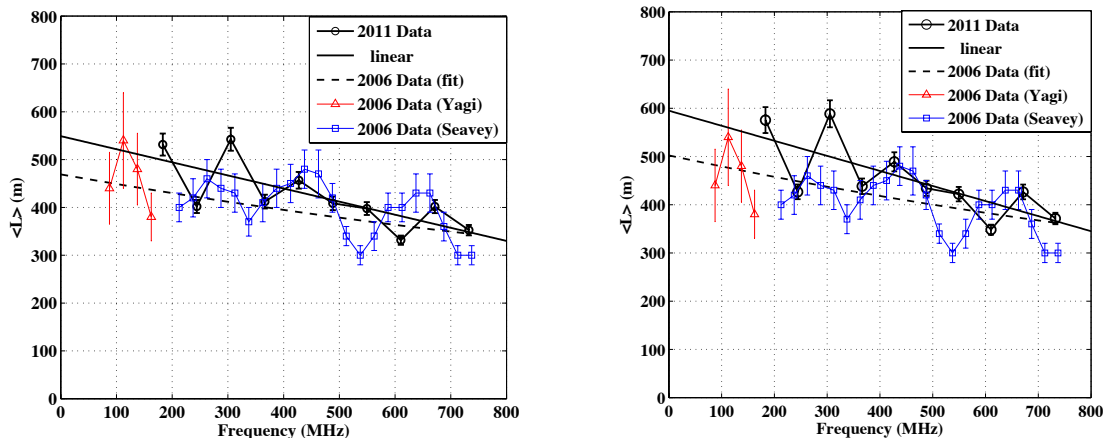


Figure 2.17: The attenuation length versus frequency, derived from equation 2.27, assuming constant reflection coefficients (left: 1.0, right: 0.85). The 2011 data is derived from the direct-bounce data (Seavey radio horn transmitter, LPDA receiver). The solid black lines are linear fits to the 2011 data. The data from [40] ($\sqrt{R} = 1.0$) is shown in both graphs, including a quadratic fit (dashed line).

Figure 2.18 uses the 2011 results with 30 MHz bins to obtain answers very close to the original study in 2006 (recall that this was 1 km away from the proto-station site, where the 2010 and 2011 data was collected). The attenuation length data is hard to trust past 750 MHz since we do not have a good measurement of the reflection coefficient above this frequency, but the data does not diverge at these frequencies either, and the linear fits should be useful for purposes of extrapolation. The fit to the 2006 data is technically quadratic, but the quadratic term is $(4.87 \pm 2.34) \times 10^{-5}$ m/MHz², making the fit nearly linear. All results are summarized in table 2.3. In figure 2.18, the finer binning admits two points which appear to be outliers at 300 and 330 MHz; however, the linear fit to the 2011 data converges with better agreement in slope with the 2006 results (see table 2.3).




Figure 2.18: Here we have the same attenuation length data as the figures above, but with a 100% reflection coefficient and 30 MHz bins instead of 75 MHz.

The values obtained for the attenuation length are comparable to, and in some cases larger than the shelf-depth, which is important for the prospects of UHE neutrino detection with ARIANNA. For example, if a down-going neutrino interacted in the bulk ice and reflected off of the bottom, hitting the detector on the way back after traveling 1000 m, these results indicate that the losses due to attenuation and reflection would only amount to about -17 dB in amplitude for an attenuation length of 500 m, which is close to the -15 dB derived above from an average temperature of -10° C. This level of absorption is easily managed considering that the ARIANNA amplifiers have gains ≈ 70 dB across the 0.1-1 GHz bandwidth.

Year	Bin Width [MHz]	\sqrt{R}	Slope [m/MHz]	Intercept [m]
2006	25	1.0	-0.205 ± 0.036	469 ± 13
2011	30	1.0	-0.18 ± 0.05	500 ± 30
2006	25	0.85	-0.23 ± 0.08	500 ± 16
2011	75	1.0	-0.27 ± 0.09	550 ± 40
2011	75	0.85	-0.30 ± 0.10	590 ± 50

Table 2.3: A summary of the fits to the attenuation length data assuming constant reflection coefficients. The 2006 data was fit quadratically, but the quadratic coefficient was $(+4.87 \pm 2.34) \times 10^{-5}$ m/MHz².

2.3 Polarization Measurement

In this section, we briefly cover the measurement of the cross-polarization fraction of the reflected bounces. The LPDA antenna is linearly polarized, and the Seavey transmitter from prior sections is also linearly polarized, though it comes with two separate feed points for horizontal and vertical polarizations. By cross-polarizing the receiver and transmitter in these reflection studies, we were able to quantify how much power leaks from the transmitted polarization into the orthogonal polarization. Reflection coefficient measurements suggest that the ocean-ice interface is a smooth, specularly reflecting surface. A smooth surface is ideal for neutrino detection, because it helps to determine the initial direction of the neutrino [15]. The radio pulse from the neutrino interaction is perfectly linearly polarized (figure 1.11). This polarization lies in the plane defined by the neutrino direction vector and Askaryan pulse propagation vector. A specularly reflecting surface preserves the polarization of the reflected signal, and therefore directional information.

We investigate the effect the ice and reflection surface beneath it has on polarization preservation by comparing the co-polarized power, $P_{||}$, to the cross-polarized power, P_{\perp} , for the angled bounce configuration. The fraction of cross-polarized power to total power is

$$F = \frac{P_{\perp}}{P_{\perp} + P_{||}} \quad (2.28)$$

Figures 2.19 and 2.20 compare the quantity F (solid curves) to a studies performed in air (dashed curves). Due to imperfections in the LPDA receivers, some power will leak into the cross-polarized configuration, representing a lower limit to F . This is estimated by air measurements with the transmitter and receiver oriented to point toward each other, and separated by 10 m to avoid near-field effects. The sudden rise to a value of 0.5 (the value

of F for unpolarized noise) at frequencies below 80 MHz is due to the LPDA response. In air, the response of the LPDA decreases dramatically below 105 MHz, whereas in snow this decrease takes place at 80 MHz (see figures 2.8-2.9 for more details).

The 2010 experimental setup was identical to the angled bounce test, LPDA transmitter and receiver, except that the LPDAs were cross-polarized. The 977 m baseline and liquid water effects described above limited the useful bandwidth to [80-180] MHz. In 2011, we replaced the transmitter with the Seavey radio horn, which made it more convenient to measure F in the direct bounce configuration. Thus, in figure 2.20, the experimental setup is identical to the 2011 direct bounce configuration. Rather than use the two feed-points on the Seavey Radio horn, the same feed point was used and the horn was physically rotated 90° .

The good agreement between the ice reflections and the air studies suggests that little power is transferred from the co-polarized direction to the cross polarized direction in the frequency range of [80-450] MHz. Regarding the higher frequencies, it is possible that not enough power is transmitted over the 1 km path length through the ice to observe a cross-polarized signal above thermal noise. The data from 2011 is being smoothed (using Welch's method [85]) more coarsely than in 2010 (the effective frequency resolution is ≈ 40 MHz), to reveal higher frequency modes of F that seem to deviate from the thermal noise value 0.5. This procedure smoothes over the effect of the finite receiver bandwidth at 80 MHz.

Thermal noise is un-polarized, having equal power in all polarizations. Thus equation 2.28 gives $(1/2)$. However some modes (e.g. near 500 and 625 MHz) are returning from the ocean with values less than 0.5. This may suggest that the reflective surface would preserve the polarization at those frequencies if we could transmit enough power to observe a clear signal. It is also interesting to note that, in an identical setup from 2006 [40], 450 MHz was also found to be the frequency where F deviated from the minimum value. However, these results further confirm the idea that ARIANNA takes advantages of the flat reflective surface beneath the ice comprising the fiducial volume.




Figure 2.19: The polarization fraction, F , is shown as a function of frequency for transmission through air (dashed line) and for the angled bounce configuration at the ARIANNA site (solid line) in 2010.




Figure 2.20: The polarization fraction, F , is shown as a function of frequency for transmission through air (dashed line) and for the direct bounce configuration at the ARIANNA site (solid line) in 2011.

2.4 Surface Waves

Recall that the index of refraction is an increasing function with increasing depth below the surface (equation 2.1). At any given depth within the firm, RF signals traveling horizontally should bend downward towards the bulk ice below. This effect is known as shadowing, and can impact the UHE neutrino detection rates [15]. In addition, there are cases of electromagnetic surface waves which can propagate for multiple kilometers given that their geometric attenuation is proportional to $1/\sqrt{\rho}$, rather than $1/\rho$, (where ρ is the distance between detector and UHE neutrino interaction) and the fact that their attenuation lengths pick up an enhancement of $\approx 2\sqrt{2}$ (Ralston). An example of a surface wave between air and metal is a Zenneck wave [71]. These properties lead to several advantages for UHE neutrino detection. In addition to enlarging the effective volume of a single station due to the enhancement of the attenuation length, the spreading of an electromagnetic wave in two dimensions rather than three means that the electric field amplitude (in the far-field approximation) will be larger by a factor of $\approx (\omega\rho/c)^{1/2}$. Finally, the geometric scaling is interesting because the volume of a surface detector grows as $2\pi\rho d\rho$, meaning the number of detectable events grows as $\sqrt{\rho}$. In summary, although there could be a shadowing effect which lowers the overall event rate, there are several effects that also boost the detection of UHE neutrinos in the presence of surface waves. An investigation was done on the 2011 expedition to look for the existence of shadowing and surface waves.

2.4.1 Experimental Setup

Recall from figure 2.1 the geometry of the angled and direct bounce tests. For the surface propagation tests, the TX and RX from the angled bounce were oriented toward one another, and the amplifier was removed. Figure 2.21 below describes the electronics and antennas, and their location and orientation. The data was taken on December 6th, 2011, and the

temperature was roughly -7°C , ensuring that there was little water content in the surface snow. The Seavey radio horn (TX) and the ARIANNA LPDA (RX) were each buried 1.5 m in the snow, co-polarized with polarizations normal to the surface. The BNC 555 delay generator controlled the trigger of the oscilloscope and the PCD. We noticed early on that there is a problem with jitter between channels A and B of the BNC 555, for delays of order microseconds only. The jitter is negligible (much smaller than 1 ns) for small delays, so we set the delay between channel A and B to zero, and measured delay using the Agilent 1 GHz scope. An RG-58 cable labeled 4 ns took the trigger pulse from the BNC 555 channel B to the HYPS pulser, which produced the usual 2.5 kV, 1ns wide, negative polarity pulse. The PCD was triggered at 5 Hz, about half of the maximum trigger rate.

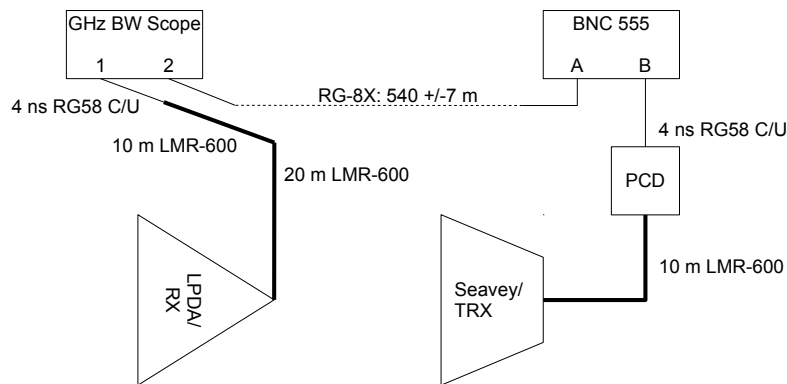


Figure 2.21: The experimental setup of the surface wave test.

Simultaneously, channel A of the BNC 555 sent a trigger pulse to the scope down the 543 meter long, custom RG8X coaxial cable. We measured this propagation time to be 2134 ns. This result is in agreement with the cables specified wave speed of 84% the speed of light. The physical length of the baseline was measured with a GPS tracker, where the mark points were taken at each antenna location. Finally the receiver caught the pulse and sent it to the scope through two LMR-600 cables 10 and 20 m in length. Those cables then attached to a 4 ns RG-58 cable which fed into the scope channel 1. The wave speed for a LMR-600 type cable is 87% of the speed of light. Test in the lab indicated that there is also a 10 ns

Component	Delay (ns)
RG-58	4
10 m LMR-600	38
20 m LMR-600	76
RG-8X	2134
PCD	10

Table 2.4: Delays associated with various components in the surface wave system.

delay introduced by the HYPS PCD itself, that is, the pulse comes out 10 ns after the unit is triggered. Table 2.4 summarizes the timing delays and the delay calculation. The antennas were buried 1.5 m beneath the snow surface, and oriented in the co-polarized state. The total systematic delay turned out to be 1964 ns. Any delay in addition to this number is due to the physical propagation time of the pulse through the firn surface.

We define the quantity Δ to represent the time after the scope is triggered that the transmitted signal arrives at the oscilloscope. From this measured quantity, we can derive the physical propagation time of the signal through the upper firn: $\Delta(ns) = (4 + 10 + 38 + \Delta t_{prop} + 76 + 38 + 4) - t_{RG8X} = \Delta t_{prop} - \Delta t_{sys}$. The quantity Δt_{sys} is 1964 ns. We measure $\Delta = 360$ ns. Let the time $t = 0$ ns correspond to when the BNC 555 triggers. Given Δ and Δt_{sys} , we can calculate the index of refraction. We have

$$\Delta = \Delta t_{prop} - \Delta t_{sys} = n\Delta x/c - \Delta t_{sys} \rightarrow n = c(\Delta + \Delta t_{sys})/\Delta x \quad (2.29)$$

The error on the index of refraction is

$$\sigma_n = (c/\Delta x)((\Delta + \Delta t_{sys})^2(\sigma_x/x)^2 + \sigma_\Delta^2 + \sigma_T^2)^{1/2} \quad (2.30)$$

From the data (shown below), the time Δ depends on where we identify the beginning of the prompt pulse, and here we admit a 10 ns uncertainty since the prompt pulse itself is 100 ns long with approximately 10 ns of rise time. We simply take the error in Δt_{sys} to be 20 ns total ($\approx 1\%$ fractional error), to account for cable delay uncertainties and any lingering jitter in the scope or PCD. The GPS tracker mark points were accurate to 5 m, so the difference between the mark points (543 m) has an error of $5\sqrt{2}$ m, or 7 m. Folding in the uncertainties, we have

$$n_{snow} = 1.29 \pm 0.02 \tag{2.31}$$

In the reference [40] above, if we set the depth equal to zero in the model for index versus depth (equation 2.1), we get 1.31, which is within one standard deviation of the answer derived here. Of course that model contains several empirically determined parameters, such as the firn depth. The depth of the firn, for example, can be observed in data from overflights of the RIS performed by the Center for Remote Sensing of Ice Sheets [13]. Introducing a 1% error in any of the empirical parameters in equation 2.1 can accommodate the difference between theory and measurement in the index of refraction. Less precise measurements of n_{snow} come from two places: snow density and VSWR measurements. Recall that in 2009 the measured snow density was between 0.32 and 0.4 g/cc (NIM), consistent with 0.36 g/cc from Dowdeswell [16]. A modified version of Looyenga's equation [52] tells us that $(n - 1)$ scales as the ratio of densities for mixed states of matter, in this case, ice and snow. (Technically, Looyenga's equation involves cube-roots of the permittivities, but switching to the index of refraction introduces only minor errors). Using 0.92 g/cc for the ice density, we obtain $n_{snow} = 1.3$ for a snow density of 0.36 g/cc. From VSWR measurements with the LPDA, we have the shift in the lower bound on the LPDA bandwidth: 105 MHz to 80 MHz. The index

should be given by the ratio: $105/80 = 1.3$. Regardless of slight deviations from $n = 1.3$, this data is in direct conflict with theoretical predictions involving surface waves that travel at the speed of light in air (c.f [52] chapter 6.3). If that were true, we would have measured $n = 1.0$, which we reject by fifteen standard deviations according to equation 2.31.

2.4.2 Results

Figure 2.22 shows the received signal (red and blue), a signal from a similar setup with a 23 m baseline (black) for comparison, and their respective power spectra in dB. The 543 m baseline data has been pre-scaled by a factor $(543/23)$ before the calculation of the power spectrum, to account for potential $1/r$ losses. This allows us to compare losses due only to the snow and firn, assuming three-dimensional wave propagation. The long-baseline data appears to have two temporal components, a prompt and delayed pulse. If we restrict the 543 m baseline data to just the prompt pulse, the power content above 400 MHz does not change, and it appears to suffer no losses at all other than three-dimensional spreading. This is surprising, since the firn is expected to have a well defined loss tangent just like the bulk ice beneath it. The hypothesis is that the wave instead radiated in two-dimensions, obeying $1/\sqrt{\rho}$ scaling with a well-defined attenuation length. This idea will be discussed in detail below.

The bottom portion shows the time-dependent waveform for the 543 m baseline (red: prompt, blue: delayed components), averaged over 100 shots through the firn (approximately 1.5 m deep) and normalized by $(543 \text{ m} / 23 \text{ m})$ to account for potential $1/r$ losses. The middle (left) portion shows the same data (red), zoomed in on the prompt component. The middle (right) portion shows the 23 m baseline time-domain pulse (black), which has no delayed component. The top portion shows the power spectra of all cases, where the red spectrum is derived from the prompt component of the 543 m baseline data. The blue spectrum is

derived from the delayed component, and the 23 m baseline data is represented in black.

We define the delayed component of the 543 m baseline pulse as after the first 100 ns of the received pulse, and it appears to have much less power per unit frequency above 400 MHz than the prompt component. To aid in the understanding of when certain frequency components arrive in the RX antenna, we also compute the spectrogram of the long baseline data, including the prompt and delayed components (figure 2.23). Our hypothesis is that the delayed component arises from scattering off of impurities in the firm, and this possibility will also be addressed below. Figure 2.23 demonstrates the chirp in the lower frequency delayed component, in that the lowest frequencies arrive later. Notice that the minimum frequency allowed by the transmitter, 200 MHz, does not arrive until 600 ns after the prompt component, whereas 300 MHz delayed power is present just after the prompt component.




Figure 2.22: Surface propagation data taken in Moore’s Bay, December 6th 2011. The black signal is the calibration pulse, the red is the prompt component of the surface pulse, and the blue is the delayed component.




Figure 2.23: This is a spectrogram of the 543 m baseline data from the previous figure. The x-axis is the frequency in GHz, the y-axis is a relative time scale in nanoseconds, and the color scale is linear, representing power. The dark red colors are the most powerful and the light blue colors are random noise.

2.4.3 Analysis

The data recorded in this particular experiment raises many questions. Our initial expectations were that the pulses would dive below the RX antenna due to the graded index of refraction. An illustration of the volume of ice beneath an ARIANNA station location which would allow Askaryan radiation to reach the detector is shown in figure 2.24. Events with neutrino vertices at points A and C should not send RF pulses to the ARIANNA station because the changing index of refraction bends the ray away. Events like B will reflect from the ocean and not be bent enough to escape detection. Notice that the effect gets worse for shallower events. This diagram, however, cannot be entirely correct, since it predicts that horizontal rays originating in the firn will always miss a surface detector. An alternative to this paradox is that there exist solutions to Maxwell's equations that both allow for bending and the creation of surface waves that can be detected. We argue here that the surface studies performed in 2011 should have mimicked event type C, in figure 2.24 below, in which the radio wave is bent downwards and away from the receiving antennas. The firn depth here is about 10 m deeper than assumed in previous sections, however, a shallower firn makes the shadowing effect less severe.

The prompt component has several important features. First, the prompt component is markedly higher in frequency (400-700 MHz) than the delayed component, and when scaled for $1/\rho$ geometric attenuation in the time-domain, it matches the power spectrum of the 23 m in-air calibration signal we use to measure P_0 in the absorption calculations above, meaning no noticeable frequency-dependent dispersion is taking place above 400 MHz. Second, assuming $1/\rho$ losses, the prompt component seems to experience no attenuation, since it has equal amplitude to the calibration pulse (after scaling by $543/23$). On the other hand, surface wave scaling ($\sqrt{543/23}$) may be assumed, which results in $\langle L_\alpha \rangle / (2\sqrt{2}) \approx 650 \pm 400$ m when averaged between 400 and 750 MHz. The large scatter in the data suggests noise is playing too large a role, and that this test should be repeated with a shorter baseline,




Figure 2.24: This figure is from the doctoral dissertation of Kamlesh Dookayka [15]. The admitted zone is cylindrically symmetric about the station location, forming a dome where direct events and reflected events like event B may trigger the station.

or a longer one with an amplifier. However, it is technically consistent with the surface wave attenuation length enhancement, since $\langle L_\alpha \rangle / (2\sqrt{2})$ is consistent with our results for the attenuation length versus frequency in the previous section.

The existence of the delayed component is hard to explain. The timing alone is difficult to understand, in that a signal which nominally should take $\Delta + \Delta t_{sys} = \Delta t_{prop} = 2324$ ns to propagate is spread out over an additional 1000 ns. Suppose the delayed components arrive later because they are lower frequency (longer wavelength) modes actually bent down by the gradient in $n(z)$, and forward-scattered by layering in the firn. Knowing the form of the index of refraction versus firn depth, we can calculate the depth of the hypothetical scattering layer. If a signal is traveling parallel to the z -axis (normal to the firn surface), the

following equation describes how long and how fast it travels a distance s :

$$c\Delta t = \int_0^s n(z)dz = 2 \int_0^{s/2} n(z)dz \quad (2.32)$$

Let B equal the baseline of 543 m between the transmitting antenna and the receiving antenna, and z_0 be the depth (for now) at which the signal scatters back off of some impurity or sudden change in index between TX and RX antennas. An example of this would be an intermediate layer of solid ice, like we observed while performing maintenance on Icicle1. We assume that the signal travels in approximately a straight line, at very shallow angle down to the impurity, scatters and heads towards the receiver. Obviously with the gradient in the index, the path will not be straight, but we approximate here by assuming the angle is shallow (see below). The path length traveled is

$$s = 2\sqrt{z_0^2 + B^2/4} \quad (2.33)$$

From the form of $n(z)$ (see prior sections), we have

$$c\Delta t = 2 \int_0^{z_0} \frac{2zn(z)dz}{(z^2 + B^2/4)^{1/2}} \quad (2.34)$$

because

$$ds = 2zdz/(z^2 + B^2/4)^{1/2} \quad (2.35)$$

Because of the form of $n(z)$, equation 2.34 breaks into two integrals, (which will be called I_1 and I_2 below), and the first is straightforward:

$$I_1 = 4n_0 \int_0^{z_0} \frac{zdz}{(z^2 + B^2/4)^{1/2}} = 2n_0B(1 - \sqrt{1 + 4(z_0/B)^2}) \approx 4n_0B(z_0/B)^2 \quad (2.36)$$

We perform a trigonometric substitution in the first step and use the binomial approximation in the second step since the scattering depth divided by the baseline is a small enough quantity. This is equivalent to making a small angle approximation on θ . For the second integral, we start with

$$I_2 = 4p \int_0^{z_0} \frac{zdz}{(z^2 + B^2/4)^{1/2}} \exp(-z/q) \quad (2.37)$$

$$I_2 = 2Bp \int_0^{z_0} \tan(\theta) \sec(\theta) \exp(-1/2(B/q) \tan \theta) d\theta \quad (2.38)$$

after making the same trigonometric substitution made in the first integral. Recall that the argument of the exponential was originally $-z/q$, which, for modest scattering depths, is smaller than unity because $q = 35.4$ m. This allows us to expand the exponential to first

order, with the understanding that we should not trust this approximation beyond scattering depths that are much larger than 35 m.

$$I_2 \approx 2Bp \int_0^{z_0} \tan(\theta) \sec(\theta) (1 - 1/2(B/q) \tan \theta) d\theta \quad (2.39)$$

This breaks things into two remaining integrals, the first of which is similar to I_1 , where we make the binomial approximation at the end. We have

$$I_2 \approx 4Bp(z_0/B)^2 - B^2(p/q)I_3 \quad (2.40)$$

$$I_3 = \int_0^{z_0} \tan^2 \theta \sec \theta d\theta \quad (2.41)$$

We can use Mathematica or tables of integrals to solve I_3 , remembering to rotate back to z coordinates at the end so the integral limits make sense. The final expression simplifies if we make the small angles approximation. Combining all the results, we have

$$c\Delta t = (4(n_0 + p)/B)z_0^2 - (Bp/q)z_0 \quad (2.42)$$

This equation is quadratic in z_0 , so we can solve (using $n_0 + p = n_{firm}$):

$$z_0 = B \frac{(Bp/q) + \sqrt{(Bp/q)^2 - 16n(c\Delta t/B)}}{8n} \quad (2.43)$$

This result gives us an approximate answer for the scattering depth as a function of the baseline distance, firm properties, and signal propagation time. Because the first term in the numerator on the right hand side is strictly negative ($p = -0.55$), we must chose the positive root of the solution to the quadratic equation. The time Δt refers to the total physical propagation time, and c is the vacuum speed of light. The problem with this calculation is that because Δt is ≈ 2800 ns, the scattering depth is of order 70 m, which is already a little large to make the approximation of equation 2.39. The expectation for the depth of the firm (according to our model) is around 70 m. Thus, we can conclude that while these numbers are correct to within 30%, a single scattering at the maximum firm depth is not the full picture.

There are several alternative explanations. First, the parameter q , which controls how quickly the density increases with increasing depth, could be smaller (fast increase in density). This is equivalent to keeping the boundary conditions constant and making the firm shallower. A second possibility is that we really do have to perform the curved path calculation with full ray-tracing, because the straight path (with varying speed) approximation is not close enough. A final possibility is that there is another effect altogether, perhaps multiple scatterings between parallel layers of ice in the firm. The latter seems the most attractive, since the difference in path lengths for scattering from the firm bottom and direct travel is small:

$$s = 2\sqrt{z_0^2 + B^2/4} \approx B(1 + 2(z_0/B)^2) \rightarrow \Delta s = 2B(z_0/B)^2 \quad (2.44)$$

Even if the entire path length difference was spent at the slowest possible speed, where $n = 1.78$, we would see a maximum delay between prompt and delayed component of $n\Delta s/c = 1.78 \cdot 18 \text{ m}/0.3 \text{ m ns}^{-1} \approx 100 \text{ ns}$.

Suppose there are just two layers of reflector, separated by a depth Δx which is small enough to consider the index of refraction constant. If the signal enters the region between the two reflection layers at an angle θ with respect to the vertical axis, then incremental path length between reflections is

$$s_i = \Delta x / \cos \theta \tag{2.45}$$

such that the total path length is $\sum^N s_i$, and for constant θ the total path length is Ns_i .

The horizontal distance traveled between reflections is

$$\Delta y = \Delta x \tan \theta \tag{2.46}$$

The number of reflections before traveling across a horizontal baseline B is

$$N = B/\Delta y = B/\Delta x \tan \theta \tag{2.47}$$

The total path length is therefore

$$s = B / \cos \theta \tan \theta \tag{2.48}$$

For $\theta = \pi/2$, we retain the horizontal baseline B , for no scattering at all. For $\theta = 0$, we have an infinite number of scatterings and the path length is infinite. For $n = 1.5$, and $\theta = 60^\circ$ (thirty degrees below horizontal), we can easily achieve propagation times of more than 3000 ns, which would be required to explain the tail end of the delayed component.

Chapter 3

The Icicle1 Prototype Detector

In December 2009, collaborators Dr. Spencer Klein and Thorsten Stezelberger (Lawrence Berkeley National Laboratory) journeyed to Moore’s Bay in Western Antarctica, to the site of a prior exploratory expedition by Dr. Steve Barwick (UC Irvine) and Dr. David Saltzberg (UCLA). They deployed the first experimental hardware for the ARIANNA project: a four-antenna, autonomously running detector. In this chapter, we will discuss the design, functionality, and data collected with this first deployed station. We will begin with describing the control software and communications, station electronics, and data management. Secondly, we will discuss the power systems and antenna hardware, with a focus on the signal antennas designed to detect Askaryan pulses.

3.1 The Prototype Station - Software, control, and communications

3.1.1 Station Electronics

The full ARIANNA array will be sited in Moore’s Bay on the Ross Ice Shelf, which is about 110 km south of McMurdo station, the main US base in Antarctica. The stations will be installed over a 30×30 km² region, with a 1 km spacing. The 1 km spacing places each station just outside the effective volume of adjacent stations [15]. Moore’s Bay is isolated from RF interference from McMurdo by a 1000 m tall ridge known as Minna Bluff. Except for a few aircraft overflights a day during the Austral summer, we expect very small levels of impulsive radio noise at the ARIANNA site. A satellite photo of the field region is shown in figure 3.1. The coordinates of the first deployed station are $77^\circ 44.523'$ S, $165^\circ 02.414'$ E.

This site was chosen for the ARIANNA array for several reasons. In addition to the favorable ice properties and ice depth (see chapter 2), the site provides the dual benefit of isolation from anthropogenic RF interference and access to wireless and satellite communications. While the geologic formations Minna Bluff and Mt. Discovery form a radio barrier around Moore’s bay, we have installed wireless communications station atop Mt. Discovery which routes wireless internet across the ice shelf from McMurdo to ARIANNA. Helicopter flights transport researchers and equipment to Moore’s bay via the pass between Minna Bluff and Mt. Discovery. Although ARIANNA stations can run autonomously, year-round wireless communications remain a goal, for data transfer and control of the *run state* of the detectors.

Figure 3.2 shows a layout of the Icicle 1 prototype system. A RF sealed box containing the CPU stack, DC-DC converters, heartbeat pulser, digitization and trigger board, low-noise amplifiers (LNA), a satellite modem, ethernet modem, and GPS circuits was buried under the snow surface. A 12 V, 100 A-hr lead-acid AGM battery (MK 8G27-DEKA) provided




Figure 3.1: This is a satellite photo of the region surrounding Ross Island, home of McMurdo station - the home of the United States Antarctic Program. Ross Island is the white island in the top center of the photograph. Minna bluff is below (South) of Ross Island, with Black and White Islands in between. Mt. Discovery is at the left-most edge of Minna Bluff.

storage for charge gathered from solar panels and a wind generator (see below). AGM stands for absorbed glass mat, in which the electrolyte is stored in woven glass fibers rather than liquid, reducing the need for ventilation and increasing charge storage surface area. The CPU stack had a PC104 form factor (Parvus CPU1421 PC104+), with three boards: a central CPU/8 GB flash-drive system that contained the operating system and ARIANNA custom station software, a four-port USB board (Diamond TMM-USB4P), and a Diamond Systems Corp. DMM-AT ADC/digital I/O board. The CPU itself is a 133 MHz AMD ELAN, running Slackware 12.2 Linux. The 8 GB solid state disk is divided equally into a system and a data partition. The interface to the data acquisition and trigger circuitry is done via the USB ports. In addition to the USB interface, the CPU uses several serial ports

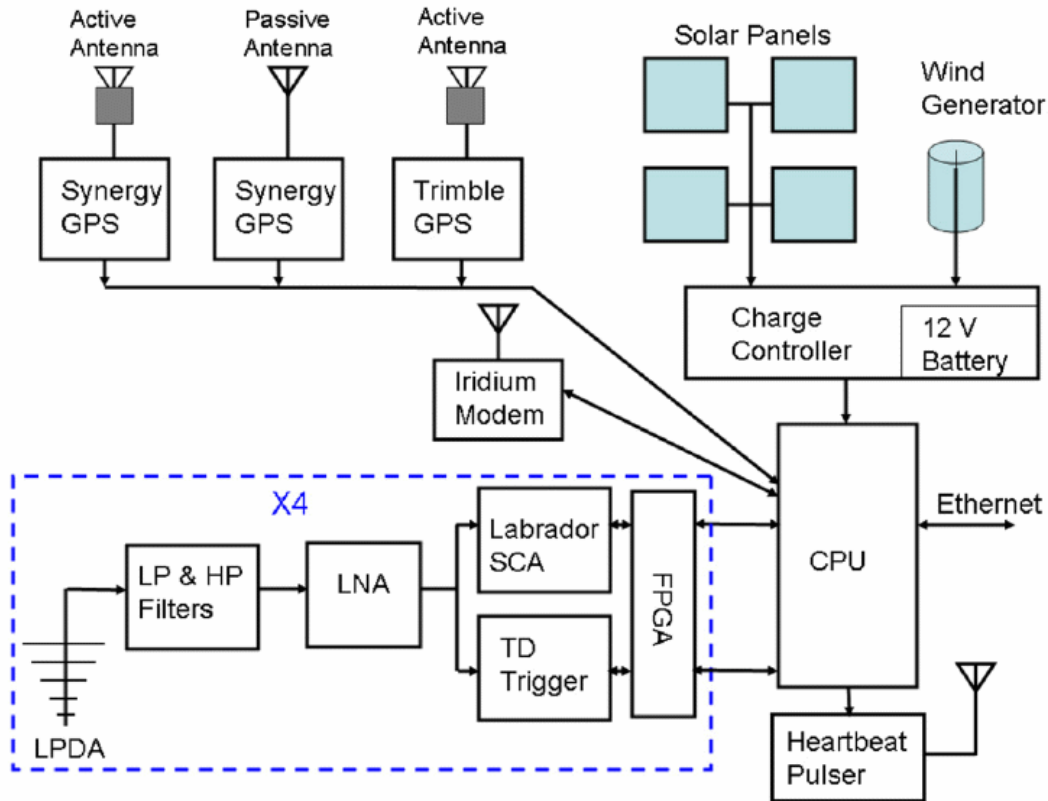


Figure 3.2: A layout of the prototype system deployed in December 2009 at $77^{\circ} 44.523' S$, $165^{\circ} 02.414' E$, Moore’s Bay, Antarctica. [31]

to talk to the Iridium data modem, a GPS receiver, and external laptops. The ARIANNA custom software, known as iceSoft, has software daemons that activate optoelectronic relays via the DMM-AT board. The relays provided 3.3 V, 5 V, and 12 V power to the various other components within the RF sealed box. The DMM-AT board also has a 12-bit, 16-channel analogue-digital converter (ADC) with the dynamic range set to +10 V. The ADC was used to make housekeeping measurements: AGM battery voltage, wind-speed, and temperature measurements were all performed successfully with this ADC (see section 4.1).

Four log-periodic dipole antennas (Create Corp. CLP5130-2 LPDA) are buried 2 m down in the snow, and connected to the electronics box via LMR-600 coaxial cable. Each antenna is filtered by a Mini-Circuits coaxial filter pair (NHP-150+ and NLP-1200+), and feeds a

custom, four-stage 70-dB LNA (100-1000 MHz). The signal is then passed to a switched capacitor array (SCA) analog-to-digital converter (ADC) and programmable trigger circuit: the CAGES board. The physics of the antennas is described in section 3.2.2. The CAGES SCAs have 256 samples and are digitized at 2.85 GHz. These are read out whenever the trigger circuit fires, requiring signals in at least two of the four antennas within a 10 ns gate. Data in the SCAs is then passed through the data acquisition daemon to be compressed and stored on the station hard drive for further analysis. IceSoft and the station daemons are described in further detail below.

Throughout the year, the station communicates via an Iridium satellite modem. During the summers, it also had a wired ethernet connection, which is connected to a wireless AFAR tower 10 m from the station. The AFAR tower communicates with the repeater installed on Mt. Discovery (about 30 km away), which in turn is linked to McMurdo Station. Work is currently underway to keep this internet connection active all year long. To minimize the possibility of radio interference on the ethernet connection, a pair of media converters switch the incoming twisted pair ethernet signal to coaxial, and vice versa inside the box, before attaching to the CPU stack. Future stations have employed special bulkhead connectors that allow RF sealed ethernet to pass through the station bulkhead without the need for power hungry coaxial converters. In February 2010, the wireless link was removed, so the station communicated exclusively via Iridium. It is configured to call home once every 15 min during the summer and 3 h during the winter, giving housekeeping data and a few sample waveforms. In subsequent seasons, the Mt. Discovery and Moore's Bay wireless equipment was left in place during the winter and proved to be useful once the sun rose again.

The data acquisition system is a modified version of the system used on the ANITA balloon flight [30]. Data are recorded with a LABRADOR [29] ASIC, which is a switched capacitor array waveform sampler. Each channel of the trigger uses a tunnel-diode detector and field-programmable gate array-based discriminator. The whole setup is controlled by a Xilinx

Spartan 3 (FPGA). The LABRADOR SCA is a single chip with 8 channels (plus a 9th reference timing channel), each containing 256 capacitors. The chip has an analog bandwidth of approximately 1 GHz, and, in ARIANNA, samples at 2.85 GS/s (2.85 GHz). Each antenna is connected to one channel. A 5th channel is connected to a 40MHz clock; this is used to calibrate the sampling rate. The other four channels are unused. The chip has 2340 Wilkinson ADCs, so it can digitize all of the stored samples in parallel, to 12 bits, in under 50 ms.

The digitization path and trigger path split just after the amplification stage. The trigger path divides the input signal into two frequency bands: 130 to 460 MHz and 650 to 990 MHz. The low-band frequencies are defined by a Mini-Circuits LFCN-320 low-pass filter and a discrete LC circuit, while the high-band is defined by a Mini-Circuits HFCN-650 high-pass filter and a LFCN-800 low-pass filter. The filters have a fairly gradual roll-off; the low-band had a 3 dB point of 460MHz, so the intermediate gap was not so important. The two bands were used to allow for improved background rejection of low or high frequency noise. Each band feeds a tunnel diode based trigger, which acts as a square-law detector. After amplification the tunnel diode feeds an FPGA-based discriminator with a programmable threshold.

The circuit threshold is electrically adjustable. Each of the 8 trigger bits (two frequency bands for four channels), are connected to the FPGA, which forms a logical trigger. For all prototype data runs, we used a trigger that subjects the outputs of the two frequency bands from a single antenna to a logical OR function, and required at least two of the antennas to detect a signal within ≈ 10 ns. This amount of time is called the gate; in fact there is another gate that checks for coincidence every 25 ns continuously, where each channel raises a continuous trigger signal for as many clock cycles as the signal is above threshold. Long (25 ns) signal pulses trigger this gate mode rather than the shorter window. During the first year of operations, the prototype station detected significant noise at 300 and 600 MHz,

most likely the subharmonics of the 2.4 GHz AFAR wireless signal used to communicate with McMurdo station. Thus, the thresholds for the lower frequency tunnel diodes were set quite high; most of the triggers were formed using the higher frequency bands. When the system triggers, the FPGA initiates a LABRADOR digitization cycle and reads out the chip. Data are compressed and can be transmitted North over Iridium satellites or wireless internet. The former only has 9600 baud of bandwidth, making it useful for looking at housekeeping measurements, and a handful of triggered events, whereas the wireless internet is the only viable mode for transmitting large amounts of data.

Only one GPS receiver is required for keeping the system clock accurate, however in 2009 several GPS receivers were deployed for testing purposes. Ultimately, the Trimble Resolution T starter kit proved to be the most reliable and accurate while consuming the least amount of power. Although GPS clocks have the ability to maintain a system clock precision of ≈ 15 ns, the precision of the unix-based clock on the PC104 CPU is only ≈ 1 ms. Over the course of the season, however, the CPU clock can also drift by up to 100 seconds. Thus the GPS timer is needed to keep the unix clock in sync with universal coordinated time (UTC).

The heartbeat pulser is an Avtech 10V, 1 ns width unit, used for generating fast calibration pulses to simulate a horizontally traveling Askaryan pulse. The signal is transmitted through LMR-600 coax and a fifth, identical LPDA antenna transmits the pulse toward the other four signal antennas from a separate location (see figure 3.33). The heartbeat was activated for 40 seconds, and then deactivated, at the start of each run in 2009 (Ariannad was set to create 6-hour runs - this was subsequently changed to 12 hours). In the 2010-11 and 2011-12 seasons, the heartbeat period was adjusted to 100 seconds to obtain more heartbeat calibration events. The Avtech pulser providing the heartbeat pulses needs a TTL pulse to trigger it. The Trimble GPS receiver produces a TTL pulse once per second, precise to 15 ns. By connecting the GPS TTL output to the Avtech, and powering the Avtech, the heartbeat pulses were fired at the signal antennas once per second. Thus, each run from

2009-10 contains 40 calibration pulses, and the next two seasons contained 100 calibration pulses per run.

3.1.2 The IceSoft control system

IceSoft is the name of the collection of control programs that govern the prototype (and eventually the Icicle2 station). The main package was written by Dr. Ryan Nichol (University College, London). For detectors manufactured after Icicle2, the software was updated to a version called SnowShovel. The original IceSoft version can be downloaded here: <http://www.hep.ucl.ac.uk/~rjn/arianna/prototype/>, however updates have been made since 2009 and can be accessed via <http://arianna.ps.uci.edu>. For more details contact either jchanson@uci.edu or jtatar@uci.edu. To compile the original version of IceSoft, a USB library, `libusb`, is required (open source availability). To run IceSoft successfully with the DMM-AT module, the proprietary library `dscud5` is also required. This library was upgraded to `dscud-6.02` in 2011 to operate the Helios CPU/ADC in Icicle2. There are seven programs associated with the original version of IceSoft:

- Ariannad - The control program which switches between low and high power configurations, and controls other daemons with POSIX signals.
- Acqd - Setting trigger conditions and reading out data from the CAGES digitiser board, over the USB 2.0 link
- GPSd - Reads time and position information from the Trimble GPS unit over the serial interface.
- GPSa - Reads time and position information from the alternative active GPS receiver (2009-10 only).

- GPSp - Reads time and position information from the alternative passive GPS receiver (2009-10 only).
- Hkd - Controls DMM-AT, optoelectronic relays, and acquires housekeeping data (temperature, wind, voltage, current).
- emailScript.sh - Simple script for connecting through the Iridium modem and emailing updates to the arianna.icebox@gmail.com account.

Formally, these programs are known as *daemons*, because they run autonomously in the background. Specifically, Ariannad executes and terminates the other daemons based on its configuration file, and the other six daemons operate according to their own associated configuration files. Thus, a user simply has to edit times and flags in the configuration files to achieve the desired *run state*. Ariannad automatically cycles the detector through several run states (described below), for acquiring data, communications, and low-power consumption.

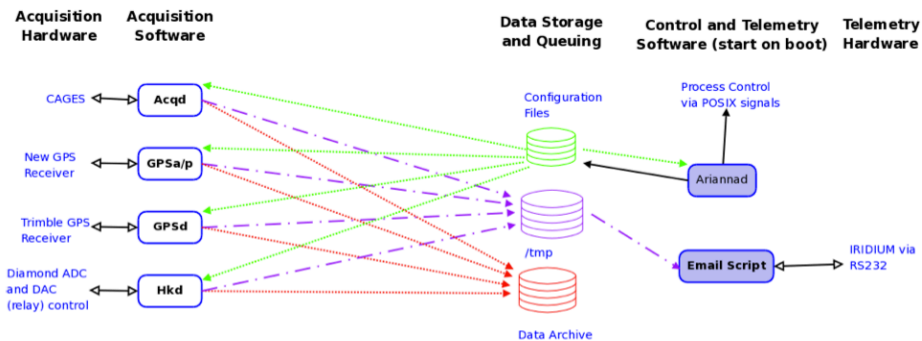


Figure 3.3: An overview of the software processes running on the first ARIANNA station.

There are two main detector modes, high-power and low-power. In the high-power mode, all of the hardware is enabled and all of the software programs are operational. In the low-power mode, only the housekeeping system is enabled by default, and the CAGES board does not operate. Ariannad controls the global run state, and it can switch to default

settings (upon a reboot) by copying the default config files to the active config files. The active config files are located in `/etc/arianna/config` on the hard disk and the defaults in `/etc/arianna/config/defaults`. The basic run state properties are:

- State 0 - Everything enabled, new runs start every 12 hours, Iridium modem always active, heartbeat pulser on for first 100 seconds of each run.
- State 1 - Only housekeeping systems enabled, Iridium modem active for 5 minutes every three hours, heartbeat pulser deactivated, switch to State 2 after 12 hours.
- State 2 - CAGES and amplifiers enabled, heartbeat pulser active for first 100 seconds of the run, switch to State 3 after 5 minutes.
- State 3 - CAGES and amplifiers deactivated, GPS units enabled, after 5 minutes switch to State 1.

Together, run states 1-3 comprise low-power mode. The main power consumers are only on for several minutes, and no two main systems are on at the same time. This cycle allowed the housekeeping system to collect and transmit data via Iridium once the sun began to set, for as long as the station's battery still had charge.

3.1.3 Configuring the Run States

For each software daemon and run state, there is an associated configuration file in `/etc/arianna/config`, with the defaults stored in `/etc/arianna/config/defaults`. For example the default file for the GPS receiver daemon in run state 0 is at `/etc/arianna/config/defaults/GPSd.config.0`. The main control daemon, Ariannad, contains several tunable parameters regarding the run other run states. Flags are entered as integers; an entry of 0 means a daemon will not be acti-

vated, an entry of 1 means a daemon will be activated. The units of the run state length are seconds.

- current run state - which run state the daemon executes first at system boot (if hard coded the run state will never switch)
- run state 0 parameters - which daemons are active in run state 0, and the time until the next run state
- run state 1 parameters - which daemons are active in run state 1, and the time until the next run state
- run state 2 parameters - which daemons are active in run state 2, and the time until the next run state
- run state 3 parameters - which daemons are active in run state 3, and the time until the next run state

Acqd stands for acquisition daemon, and it is responsible for setting trigger conditions and data taking. It requires these parameters:

- Maximum number of events - when the number of events in a run reaches this integer, data acquisition stops and must be restarted by Ariannad (typically 1000)
- Software trigger period - number of seconds between automatic triggers (typically 67)
- Enable threshold scan - used for testing purposes, allows the system to automatically measure event rate vs. threshold
- Threshold scan step size - the number of DAC counts between subsequent thresholds in a scan

- Threshold scan points per step - the number of rate measurements at each step in a scan
- ... (skipping unused parameters)
- Enable global threshold - set the same DAC value (12 bit, 0-4096) for all thresholds (rarely used)
- Global threshold value - the DAC value in counts of the global threshold
- Threshold Values - A comma separated list of 8 DAC values corresponding to the thresholds of the high and low frequency band triggers of the four channels
- ...

When the maximum number of events are collected by the acquisition program, the program terminates and begins again. This results in several files (depending on the event rate) of up to the maximum number of compressed data events corresponding to a single run. The format of the data structures can be located in the `ariannaStructures.h` file as part of the `iceSoft` package header files. Threshold scans were performed in 2009 to determine the appropriate thresholds to achieve desired event rates (discussed below). Thresholds were adjusted once known sources of man-made noise were removed from the environment of the prototype.

Hkd stands for housekeeping daemon, and it is responsible for two main functions: recording environmental measurements such as battery voltage and windspeed, and sending the activate/deactivate signals to the optoelectronic relays that power the other systems in the station (such as data acquisition). When Hkd receives a deactivate bit from Ariannad, Hkd is the program that shuts off power to the corresponding hardware. Hkd is governed by the following flags:

- Housekeeping period - the number of seconds between Diamond Systems DMM-AT board ADC measurements
- CAGES board state - One if CAGES receives power, zero otherwise
- Iridium modem state - One if the Iridium modem receives power, zero otherwise
- LNA power state - One if the amplifiers receive power, zero otherwise
- Trimble GPS state - One if the main GPS receiver is powered, zero otherwise
- Synergy GPS state - One if the Synergy GPS receiver is powered, zero otherwise
- Ethernet converter state - One if the ethernet to coax converters are powered, zero otherwise
- Anemometer sensor state - One if the wind sensor receives power, zero otherwise
- Heartbeat pulser state - One if the calibration pulser receives power, zero otherwise
- Heartbeat on period - the number of seconds to leave the calibration pulser powered
- Heartbeat off period - the number of seconds to leave the calibration pulser deactivated
- Iridium on period - the number of seconds to leave the satellite modem powered
- Iridium off period - the number of second to leave the satellite modem deactivated

During the summer months, the housekeeping period was one measurement every 10 seconds of the various voltages, currents, wind speed and temperatures (see next section). During the winter months when the station was running in low power mode, this period was one measurement every 15 minutes and the data was transmitted exclusively via iridium satellite. The Iridium period was usually 300 seconds. This proved to be just long enough to complete a call from the modem to the satellite network and transfer one set of housekeeping measurements, a ZDA packet, and the system uptime. Because the Iridium modem consumes

a comparable amount of power to the entire acquisition system (around 5 W), we tried to maximize the off time of the modem. In the upcoming version of the 7-station ARIANNA array, the Iridium dial-up modem has been replaced with a short-burst data (SBD) module capable of sending text messages while consuming at most one fifth of the modem's power. The SBD module also comes with a sleep mode that consumes a negligible amount of power. Similarly, since the heartbeat pulser is restricted to a rate of 1 Hz, we use the heartbeat on/off periods to control the number of heartbeat events per run so as not to clog the data with excess calibration pulses.

GPSd has the following parameters (GPSa and GPSp will not be covered since they were only used for a brief period):

- GPS clock skew - the number of seconds the unix clock is allowed to drift from the GPS clock before this daemon corrects it
- GPS clock update enable - whether or not the GPS daemon will update the unix clock
- NMEA interval - Interval (in seconds) between recording the NMEA data strings (see text).
- Enable recommend minimum course data (usually set to zero)
- Enable fix information data (usually set to one)
- Enable latitude and longitude data (redundant - set to zero)
- Enable vector track and speed over ground (usually set to zero)
- Enable detailed satellite information (usually set to one)
- Enable satellite information (usually set to zero)
- Enable date and time information (ZDA packets, usually set to one)

The skew between the unix and GPS clock was never allowed to be more than 100 seconds, and GPS correction was always enabled for the duration of the 2009-12 data runs. The acronym NMEA stands for National Marine Electronics Association, which sets an industry standard for GPS data strings - this parameter controls how often the system reads GPS information. The ZDA data structure contained the unix seconds and unix microseconds alongside the GPS seconds and microseconds, which allowed the system to correct clock skew. The other GPSd data structures (GSV, GGA, GPSSAT) contain data like latitude and longitude, and the location and number of satellites.

3.2 The Icicle1 Prototype - Power systems and Signal Antennas

3.2.1 Solar and Wind Hardware, and the Basic Charging Circuit

The station housekeeping system and CPU stack receive power directly from a 12 V, Pb-acid AGM (absorbed glass mat) battery. The voltage lines going into the RF-sealed box are filtered with Compac 100-1007 filter pins, which filter any noise picked up on these lines above 10 MHz. At 80 MHz, the LPDA signal antennas become sensitive (their VSWR drops below 2), and the filter pins have an insertion loss of ≈ 50 dB at this frequency. The Pb-acid battery nominal voltage is between 11.8 and 14 volts. At 11.8 V, the battery no longer provides sufficient current to drive the internal DC-DC converter inside the station which provides the 5 V necessary to operate the CPU stack. The battery is considered fully charged at 14 V. While the RF-sealed electronics box is buried flush to the snow surface, the station also had a solar panel and instrumentation frame that supported four 30W GE solar panels (16.2 V closed circuit) and the iridium and GPS antennas, and the anemometer (Vector Instruments Ltd. model A100L2). Figure 3.4 below is a 3D rendering of the solar




Figure 3.4: The solar frame depicted here was used for Icicle2. The Icicle1 solar frame lacked only the telescoping elements at the bottom. The aluminum pipes slide through the Kee-Lite connectors, allowing easy setup and takedown of the solar frame. Instrumentation antennas and the anemometer were fixed to the top corners above the solar panels, at a height of 7 ft. above the snow surface.

frame, with telescoping modifications introduced in 2011-12 for Icicle2.

The solar frame was constructed of standard schedule-40 aluminum pipe (1.12 inch inner bore, 1.66 inch outer diameter). Kee-lite pipe components (Simplified Building Concepts) connected the pipes together and fixed the solar panels to the horizontal cross-bars. The vertical pipes were each 7 ft. long, and the cross-bars were 3 ft. long (in the Icicle2 redesign in figure 3.4, the solar panels are oriented vertically which allowed us to shorten the cross-bars to 2 ft.). In 2009, a Forgen 1000LT wind generator was mounted in the center of the square solar array on a central pipe with supporting cross-bars. Four supporting nylon guy-wires secured the rig by tying the upper corners to bamboo rods buried 2 ft. under the snow and oriented orthogonally to the guy lines. In 2010-11 and 2011-12, the station employed an

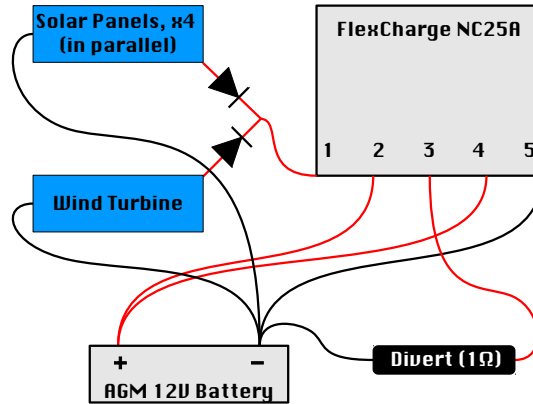


Figure 3.5: The NC25A charge controller regulated charge from the solar and wind sources onto the battery. Ports 4 and 5 together measure the battery voltage, and, if the battery is below an adjustable set point, ports 1 and 2 are connected. Otherwise, port 1 is switched to port 3.

Aerogen-6 wind turbine on a separate, 14 ft. tower located 60 ft. to the South. In this case, the blocking diode was located at the base of the tower, and 14 ft. of AWG 12 teflon coated wire carried current to the tower base. From there, 100 ft. of SEOOW type, 3-conductor, 12-AWG copper stranded cable carried the current to the solar frame where it was spliced in parallel with the solar panels.

The basic charging circuit is described in figure 3.5. The charge controller unit (FlexCharge NC25A-12) was located inside a Compac RF sealed aluminum box with the same 100-1007 filter pins filtering inputs and outputs. Filtering was necessary because of the potential for electromagnetic pulses during charge switching. When ports 4 and 5 detect that the battery is full, the charge controller switches power output from port 2 to port 3, through a 1Ω , 100 W ceramic divert load resistor instead of the battery. The RF-sealed box containing the NC25A, the divert load resistor, and the AGM battery were all contained within a plastic crate and buried $\approx 2\text{m}$ from the electronics box, with 12 AWG teflon coated power lines running between them. This circuit powered all of the station electronics described in the previous section, and the power consumption of each device is listed in table 3.1.

Device	Voltage (V)	Power consumption (W)
CAGES/USB	5.0	5.22
4 × LNA (70 dB)	5.0	1.24
Synergy GPS	5.0	0.96
Iridium sat. modem	5.0	≈ 2
Trimble GPS	12.0	1.26
Avtech pulser	12.0	1.74
2 × Ethernet conv.	12.0	7.69
CPU stack	5.0	7.02

Table 3.1: The power consumption and operating voltages for the Icycle1 proto-station electronics. (LNA: low-noise amplifier). The Iridium satellite modem causes the base current of the station to jump by $\approx 0.5A$ during a call, making the consumption 6W during this time.

Certain devices were removed throughout the three seasons as they proved to be no longer useful or relevant to the research. Once it was determined that an active antenna would not be needed for GPS time synchronization, GPS devices were removed (disconnected). This determination was based on the similar uncertainty in ZDA timing packets (≈ 20 ns) received over both the active and passive GPS systems. During the deployment of the second ARIANNA prototype (Icycle2), we learned of RF sealed feed throughs for ethernet cable. Installing this part in Icycle1 allowed the removal of the ethernet converters which had been actively switching external ethernet signals to coaxial signals before passing through the RF sealed electronics box. Table 3.2 shows the year certain devices were removed, and the reduction in power consumption.




Figure 3.6: (Left) The Icicle1 prototype in December 2010. The solar frame protrudes above 4 ft. of snow accumulation that occurred sometime during the Austral winter of 2010. The red and green flags mark the positions of the signal antennas. (Right) After uncovering the electronics box, was located flush with the original snow level. The four coaxial cables visible in at the lower right connected the signal antennas to the electronics box. (Photos by the author (left) and Raytheon Polar Services personnel (right)).

Device	Power consumption (W)	Removed?	Date
CAGES/USB	5.22	N	N/A
4 × LNA (70 dB)	1.24	N	N/A
Synergy GPS	0.96	Y	Dec 2010
Iridium sat. modem	≈ 2	N	N/A
Trimble GPS	1.26	N	N/A
Avtech pulser	1.74	N	N/A
2 × Ethernet conv.	7.69	Y	Dec 2011
CPU stack	7.02	N	N/A

Table 3.2: The power consumption and operating voltages for the Icicle1 proto-station electronics. (LNA: low-noise amplifier). The Iridium satellite modem causes the base current of the station to jump by ≈ 0.5A during a call, making the consumption 6W during this time.

3.2.2 Several properties of the Log-Periodic Dipole Array Antennas

Here we summarize the what we know about the properties of the signal antennas deployed to capture the Askaryan pulses. This includes our initial attempts to study the *effective height operator* of the LPDA, which converts incident electric fields into voltages on the transmission line feeding into ARIANNA digitizers. Understanding this operation precisely will feed into the reconstruction of the electric field created by neutrino interactions. Other numbers like the antenna positions, voltage standing wave ratio, gain, and beam-width will be documented below, and compared to a simulation designed to calculate the LPDA antenna properties based on the antenna geometry. In [15], an optimization has been performed to maximize observed neutrinos based on the antenna properties, and this optimization will be compared to actual antenna numbers.

The current ARIANNA LPDA antenna is a 17-element logarithmic $\lambda/2$ dipole array, with an alternating phase pattern. This means that dipole elements across from one another are not directly connected to each other, but attach in an alternating fashion to the upper and lower halves of the boom. Figure 3.7 shows the equivalent electrical configuration of the elements, and figure 3.8 shows a complete schematic of the specific model used for the prototype station. In what follows, the front of the antenna where the smallest dipole element is located is associated with the spherical coordinates $\theta = 90^\circ$ (polar) and $\phi = 0^\circ$ (azimuthal).

The geometry of the elements of a log-periodic dipole array obey the following two equations:




Figure 3.7: (Left) The feed point of ARIANNA LPDA is a standard N-type RF connector. The central boom (feeder line) of the antenna is comprised of an upper and lower half. Each half is electrically connected to an alternating set of dipoles. (Right) The ratio of lengths of adjacent dipoles and the ratio of the distance between adjacent dipoles are equal and constant [65].

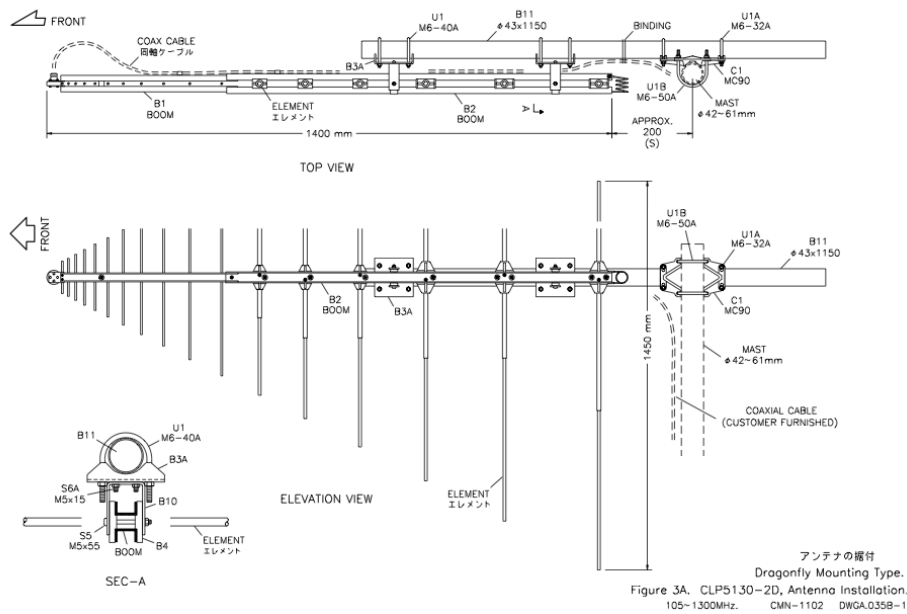


Figure 3.8: Schematic of the CLP5130-2, provided by the manufacturer. The logarithmic structure of the element geometry causes many of the antenna properties to be periodic in the logarithm of the frequency.

$$\tau = \frac{D_{n+1}}{D_n} = \frac{L_{n+1}}{L_n} = \text{constant} \quad (3.1)$$

$$\sigma = \frac{1 - \tau}{4 \tan \alpha} = \frac{D_n}{2L_n} \quad (3.2)$$

The quantities D_n and D_{n+1} are shown in figure 3.7 (right) as the distance between elements n and $n+1$, and the distance between element $n+1$ and the next adjacent element. L_n is the combined vertical length of the dipole halves of element n , and L_{n+1} is the combined length for element $n+1$. The angle 2α is the opening angle of the array. The parameter σ is also a constant, and it turns out that specifying either τ or σ automatically specifies the other for a maximum LPDA gain in dBi. The CLP5130-2 has the following parameter values: $\tau = 0.83$ and $\sigma = 0.081$ which combine to give $\alpha = 27^\circ$, the measured opening angle.

A simulation of the VSWR (no units) and the gain (in dBi) of a 17-element LPDA with the geometry of the CLP5130-2 is shown in figure 3.9. Note how the periodic trend in the VSWR extends over the entire frequency range. Ideally, we'd like to have a broadband antenna with the VSWR equal to unity, and independent of frequency so we could claim equal sensitivity at every frequency in the experimental bandwidth. Building such an antenna is impractical, and the logarithmic nature of the LPDA produces a low VSWR that is periodic in the logarithm of frequency. The VSWR fluctuates up to 2.0 with a period of ≈ 75 MHz. Indeed, measurements from 2009 suggest that the VSWR fluctuations exist, even when the LPDA is buried in the snow, but they only reach 1.5 (see section 2.2.2).

There are usually two angles associated with the narrowness of the beam width of an antenna. The *E-plane* and the *H-plane* of a linearly polarized antenna contain the electric and magnetic field vectors, respectively. For dipole arrays, the electric field lies in the same plane as the




Figure 3.9: A simulated frequency-swept analysis using code from [8]. The forward gain, in decibels over an ideal isotropic source (dBi), and the VSWR, are for a LPDA with 17-elements and the σ and τ parameters of the CLP5130-2.

dipoles. Let us assume that the spherical coordinate θ remains equal to 90° in the E-plane, and that the spherical coordinate ϕ (azimuthal) remains constant in the H-plane. We define the E-plane half-power beam width ϕ_{HPBW} as twice the angle at which the radiated power per unit solid angle $dP/d\Omega$ drops by half. Similarly, the H-plane half-power beam width is θ_{HPBW} . The *beam area* is defined in equation 3.3, where P_n is the normalized power pattern (equal to the Poynting vector divided by the maximum value of the Poynting vector).

$$\Omega_A = \iint_{4\pi} P_n(\theta, \phi) d\Omega \approx (\theta_{HPBW})(\phi_{HPBW}) \quad (3.3)$$

It turns out (see Appendix A) that the gain G of the antenna is equal to $k4\pi/\Omega_A$ where k is an efficiency factor close to unity. With the above approximation, we have $G \approx$

$4\pi/(\theta_{HPBW})(\phi_{HPBW})$. From the same simulation code that produces figure 3.9, we can also learn the beam widths versus frequency (table 3.3). The numbers in table 3.3 are in agreement with the $dP/d\Omega$ curves specified by the manufacturer in figure 3.10. The equation for the gain can be expressed as

$$G \approx 41000[sq.deg]/(\theta_{HPBW}[deg])(\phi_{HPBW}[deg]) \quad (3.4)$$

Using the numbers in table 3.3, we obtain $G \approx 4.5$ or $10 \log(4.5) = 6.5$ dBi, which is the average, frequency independent gain shown in figure 3.9. The maximum gain of the LPDA is in the direction of the spine, which is orthogonal to the dipole elements (assuming co-polarization). Another feature of the antenna we learn from graphs such as figure 3.10 is that the ratio of forward gain to backward gain, known as the FB ratio, is about 15 dB. This is important for eliminating down-going anthropogenic signals and potential cosmic-ray induced radio showers [45], since we want the antenna to be sensitive only to the ice beneath it.

frequency [MHz]	ϕ_{HPBW}	θ_{HPBW}	gain (dBi)
< 300	74°	144°	5.9
300-600	70°	128°	6.6
600-900	74°	134°	6.2
900-1200	78°	132°	6.0

Table 3.3: The results of the LPDA simulation with the FORTRAN code from Balanis. The beam widths in the E and H planes are almost constant over a wide range of frequencies, and this is reflected in the stability of the gain.

The reason we must understand the antenna gain is that it affects the neutrino sensitivity, for two reasons. The sensitivity decreases for smaller gains because the power pattern is too broad; too high a fraction of the LPDA beam pattern looks horizontally, where events are expected to experience the shadowing effect, and too little of the beam pattern looks down




Figure 3.10: The manufacturer of the CLP5130, Create Corp. provides the E-plane (dashed) and H-plane (solid) power patterns here. The radial units are in dBi and the graphs have been normalized such that the maximum gain is equal to 0 dBi. The drop-off in power with respect to azimuthal angle in the E-plane has been verified in the lab.

toward the region expected to contain upcoming or reflected Askaryan pulses. Conversely, the sensitivity decreases for large gains because the viewing volume of the antenna is too small; the gain in sensitivity for a narrow viewing region cannot compete with losses due to the reduction of solid angle. In [15] the results of simulations show that the optimized antenna gain assuming the approximate relationship above between gain and beam angles is between 7 and 10 dBi. One caveat is that given the potential observation of surface waves in December 2011, it might actually be more optimal for the antennas to in fact have a larger amount of $dP/d\Omega$ looking sideways for horizontally traveling Askaryan pulses (section 2.4.3), and thus a smaller gain.

The **effective height** of an antenna carries a simple interpretation for a simple dipole at a single frequency, and becomes slightly more complex as we treat all frequencies with more complex antenna geometries [78]. Consider a single frequency plane wave incident on a $\lambda/2$




Figure 3.11: By varying the gain of the ARIANNA antennas (with a trigger condition of 2 of 4 antennas hit at 5 times the r.m.s. of noise) and counting the fraction of detected events, the optimal gain for the ARIANNA antennas is revealed. This data is digitized and normalized from figure 4.1 of [15].

dipole of length l . The amplitude of the voltage transmitted to the load resistance by the antenna and the amplitude of the incident electric field are related by the effective height h : $V = \mathbf{h} \cdot \mathbf{E} = hE \cos \theta$, where θ is the angle between the electric field polarization and the dipole. Assuming a co-polarized state, we have $h = V/E$. For a $\lambda/2$ dipole, the current distribution is sinusoidal: $I(x) = I_0 \sin(\pi x/l)$, and we can think of the effective height as the averaged, normalized current distribution over the length of the antenna:

$$h = \frac{1}{I_0} \int_0^l I(x) dx = \frac{2}{\pi} l = \frac{\lambda}{\pi} \quad (3.5)$$

This calculation is instructive for several reasons. First, because changing the length or shape of the antenna changes the current distribution, the effective height changes non-trivially. More importantly, there is an overall scaling with wavelength, meaning the effective height

is inversely proportional to frequency:

$$h(f) = 2 \times \sqrt{\frac{G c^2 Z}{4\pi f^2 Z_0}} \quad (3.6)$$

For the derivation of this expression, see Appendix A. Here, G is the unit-less gain, c is the speed of light, f is the frequency, Z is the impedance at the antenna terminals and Z_0 is the impedance of free space. Let us assume that the dipole is in air, has 50Ω resistance at the terminals, and a standard gain of 1.64. A common trick in antenna analysis is to use 120π for the free space impedance, introducing an error of less than 0.1 percent. We obtain $h = 0.83 \frac{\lambda}{\pi}$, which is close to the ideal answer above but modified for the impedance mismatch between antenna terminals and free space.

In the summer of 2010 we performed measurements to check the above equation with the ARIANNA LPDA. This measurement followed the calculations in [34], where the signal input was a step function. In an experiment where a known signal V_{src} is sent through the antenna under test (AUT), transmitted through air a known distance r and received by the same model antenna (V_{rec}), the effective height versus the frequency ν may be derived in the Fourier domain:

$$h_N(\nu) = \sqrt{\frac{rcV_{rec}(\nu)}{i\nu V_{src}(\nu)}} \quad (3.7)$$

This expression from [34] has been normalized for simplicity:

$$h_N(\nu)/h(\nu) = 2\sqrt{Z_0 Z}/(Z_a + Z) \quad (3.8)$$

$V_{rec}(\nu)$ and $V_{src}(\nu)$ are the complex Fourier transforms of the input and received signals. Terms like the 4π and the gain have cancelled due to the division of source and receiver signals, given that the antennas are alike. We may assume that the antenna impedance Z_a is matched to the 50Ω terminal impedance Z to simplify the equations. Because the VSWR is low for the entire bandwidth, this isn't an unreasonable assumption. In fact, the same simulation used to produce figure 3.10 demonstrates that the magnitude of the antenna impedance Z_a is $50 \pm 14 \Omega$, where the main contribution to the standard deviation is the periodic variations with the logarithm of frequency (see below). Thus, the overall expression for the effective height becomes (Appendix A)

$$h(\nu) = \sqrt{\frac{ZrcV_{rec}(\nu)}{Z_0 i\nu V_{src}(\nu)}} \quad (3.9)$$

The experimental setup is displayed in figure 3.13. Two LPDA antennas were placed atop wooden poles 2 m above the cement floor in the Physical Sciences High Bay. This is not ideal since it allows for delayed reflections. The tips of the antennas were separated by 6.9 ± 0.05 m. We set the line impedance Z equal to 50Ω , the speed of light to be 0.3 m/ns (making Z_0 equal to 120π in SI units), and we assume the manufacturer's measurement of the gain, 7 dBi. Instead of a step function, we used the same Pockel Cell Driver (PCD) used in the 2010 ice properties experiments (chapter 2), attenuated by 20 dB. This eliminated the need for an amplifier, which would have been required by our function generators capable of

producing square or step pulses. Eliminating the amplifier removes a transfer function from the denominator under the square root of equations 3.9, reducing noise and complications.




Figure 3.12: Another example of an attenuated (60 dB) Pockel Cell Driver pulse. In this instance, the raw pulse is only 1 kV in negative amplitude, and the pulse width has been reduced to ≈ 1 ns with the aid of a shunt comprised of a T-junction and a short (both N-type RF connections). The trade-off is a small ripple after the pulse.

Thus, the raw signal is essentially a negative polarity, 1.4 kV amplitude δ -function. The signal was carried to the transmitting antenna by Heliax coaxial cable and from the receiver by LMR-600 coaxial cable to a 1GHz oscilloscope (Tektronix TD4S, set to 2.5 GS/s). We can conclude that the transfer function of the cables is negligible, because a system comprised of just the cables, signal generator, and scope did not alter the frequency content of the negative polarity 1 ns-wide PCD pulse. This was also true of smaller amplitude signals such as a 2.5 V step function with 2.5 ns rise-time, generated by a 20 MHz Agilent waveform generator (33220A Agilent), which have similar frequency content. The received pulses were averaged and are plotted in figure 3.14. The peak to peak voltage is within the linear range of the scope, and with the time-base set to 2.5 GS/s, Nyquist's sampling theorem guarantees that we capture frequencies up to 1.25 GHz. To calculate the form of the effective height in the time domain from the input δ function and the received pulse in figure 3.14, the complex




Figure 3.13: The LPDA antennas were separated by 6.9 m, 2 m above the ground. The LPDAs were horizontally co-polarized.

phase must be unwrapped appropriately by the complex square root in equation 3.9.

If we neglect the phase of the complex effective height and take just the magnitude, it should match the theoretical expectation in equation 3.6, up to inefficiencies like impedance mismatching. Figure 3.15 displays the experimental data compared to equation 3.6. The data covers frequencies between [75-975] MHz with 50 MHz bins, and has been shifted upward by 0.1 m, which minimizes the χ^2 per degrees of freedom. The χ^2/dof is defined in the usual way, with h_i being the data at a given frequency, h_ν being the prediction of equation 3.6, and σ_i being the standard deviation for h_i in bin i .

$$\chi^2/N = \frac{1}{N} \sum_{i=0}^N \frac{(h_i - h_\nu)^2}{\sigma_i^2} \quad (3.10)$$

The 0.1m correction could be necessary for several reasons. First, the actual antenna impedance might not be perfectly matched to the oscilloscope or PCD, which has an output impedance of 50 Ω . From equations 3.6 and 3.9, we see that the effective height in the




Figure 3.14: The average pulse received from the experimental setup shown in figure 3.13. Normally an amplifier would limit the linear vertical range to ± 1 V, but here we are only limited by that of the scope, which is ± 10 V.

frequency domain goes as \sqrt{Z} . For example, the average of the magnitude of the impedance near 100 MHz would only have to be 57Ω rather than 50 to induce this correction. Secondly, the antennas might not be separated by enough distance to escape near-field effects. In the near-field, the electric field does not strictly obey r^{-1} dependence, which is assumed in equation 3.6 where we solve for the effective height. While the separation was limited to 6.9 m, the Fraunhofer distance $2D^2/\lambda$ works out to be 3 m at 200 MHz. Because $6.9 \text{ m} > 3 \text{ m}$ we are satisfying one far-field condition, but not satisfying the condition that $r \gg \lambda_{max}$. Third, there could be destructive interference from reflections off the cement floor of the experimental area. Given the geometry of the setup, these should be delayed by only ≈ 4 ns, and the width of the pulse is 75 ns. In any case, compared to the effective height at low frequencies, this is a small correction and the chi-squared per degree of freedom minimizes at 1.67 with 18 degrees of freedom.

It is important to check that the CLP5130-2 converts incident electric fields into voltages in the expected way in the Fourier domain, because this is the starting point for the *shelfmc* simulations that predict how significant an Askaryan pulse is given an incoming angle. Sev-




Figure 3.15: Equation 3.6 matches the magnitude of the effective height measurements in the Fourier domain. The reduced chi-squared comparing the data to the curve is $\chi^2/dof = 1.67$, with $N_{dof} = 18$.

eral aspects of this conversion, however, are still not known very well. Eventually, we will want to know the time-profile of the incident electric field given the observed voltages in the data. This conversion can be performed by inverse Fourier transforming the data above, taking care to unwrap properly the phases given the complex square root. However, this process is equivalent to deconvolution because we must divide by the source signal in the Fourier domain before executing the inverse transform. Deconvolution usually magnifies noise sources. In the present data we have reflections and external noise from the UC Irvine radio station (kUCI - 88.9 MHz) (when using amplifiers).

A rough attempt at producing the time-domain version of the effective height, with the correct units of m/ns, is shown in figure 3.16. The overall amplitude is by no means fully understood, and this is meant to be a first estimate of the time-domain waveform that must be convolved with an incident electric field to obtain the result in V vs. ns. Because of the noise present, a running-average filter of 1.5 ns has been passed over this data. While this

smooths the noise away from data, the running average filter also acts as a crude low-pass filter. In this case the filter is a sinc function in the Fourier domain with the first null at $1/1.5 \text{ ns} \approx 670 \text{ MHz}$. Eliminating some frequencies in favor of others changes the shape and amplitude of the effective height waveform. Currently, all that is necessary for the *shelfmc* simulations is figure 3.15. The next-generation ARIANNA triggering analogue transient waveform digitizer (ATWD) triggers in the time-domain, meaning the convolution process will have to be fully understood.




Figure 3.16: This is an estimate of the signal that must be convolved in the time domain with an incident electric field to produce the voltage waveform observed in ARIANNA data.

One open question is how the effective height operator changes shape given a different incident angle. We already know that the gain varies smoothly with θ and ϕ coordinates. So does the effective height merely undergo an overall scaling proportional to the shift in gain? Equation 3.6 suggests that the effective height goes as the square root of the antenna gain. For the LPDA the gain is independent of frequency, so the gain should just be an overall scaling factor in transform operations. There is also the question of polarization. Does the effective

height retain an overall scaling as the LPDA and electric field vary smoothly from co- to cross-polarized orientations? We can quantify these ideas in the following two equations, where bold-face indicates a vector:

$$V_{rec}(t) = \frac{1}{2\pi r c} H_{sys} \circ h_{eff}(t, \theta, \phi) * h_{eff}(t, 0, 0) \circ \partial_t V_{src}(t) \quad \text{where} \quad (3.11)$$

$$(\mathbf{f} * \mathbf{g})(t) = (f_\theta \circ g_\theta)(t) + (f_\phi \circ g_\phi)(t) \quad \text{and} \quad (3.12)$$

$$(f \circ g)(t) = \int_{-\infty}^{\infty} f(\tau)g(t - \tau)d\tau \quad (3.13)$$

The third equation is just the definition of convolution with a wide range of applications in physics. From this equation, we can see why the units on the y-axis in figure 3.16 must be m/ns if the function f is the effective height and g is the incident electric field. The *convolution theorem* states that, in the Fourier domain, convolution is just a simple multiplication of f and g . The star-operation combines the vector dot-product with convolution (the second equation), meant to represent the effect of antenna polarization as well as convolution with effective height. Equation 3.11 implies that the effective height operator could be different in each polarization coordinate. The first equation combines these ideas into a hypothetical expression for the voltage observed at the terminals of one antenna given the input signal into an identical antenna some distance r away. The receiving antenna is oriented at angles θ and ϕ with respect to the spine of the first antenna. The function H_{sys} is a transfer function incorporating effects due to cables, amplifiers, and filters.

The transmitting antenna acts on the input pulse with the operator $2h_{eff} \circ \partial_t$ [9]. While the data above confirms the simple formula in the Fourier domain (figure 3.15), it is not a given that equations 3.11 must hold. In [45], it is however assumed that the effective height

multiplies and convolves like a vector in the Fourier domain. Figure 3.15 is essentially a derivation of $|h_{eff}(\nu)|$ with the two angles set to zero, and 3.16 is the corresponding answer in the time-domain. We will begin to perform experiments to derive the form of h_{eff} for various angles and polarizations in 2013 with the help of the newly constructed anechoic chamber at the University of Kansas, designed to help understand the experimental equipment of Antarctic researchers [13].

3.2.3 Pulsed LPDA experiments in Aldrich Park

In the summer of 2010, we decided to experiment with the signals received on LPDA antennas. In the absence of an anechoic chamber, which blocks most reflections to better than 100 dB [13], we chose to attempt these tests in Aldrich park, far from any major conductors which could cause reflections. There were several ideas we wanted to confirm about the LPDA. First, we wanted to confirm the $1/r^2$ behavior of the received power with respect to distance. Second, we wanted to record the various pulse shapes corresponding to different LPDA orientations (figures 3.7 through 3.16 above). Finally, we wanted to ensure that the modeled gain with respect to orientation matches the data independently of frequency. If the spectra for two orientations match at all frequencies, they confirm the idea that the antenna response merely undergoes an overall scaling proportional to the shift in gain.




Figure 3.17: (Left): *Bore sight* configuration. (Right) Configuration 1.




Figure 3.18: (Left): Configuration 1 (45°). (Right) Configuration 1 (90°).




Figure 3.19: (Left): Configuration 2. (Right) Configuration 2 (45°).




Figure 3.20: Configuration 2 (90°).

Observing the geometric $1/r^2$ effect is an important first step before beginning other studies such as gain with respect to angular orientation. Two LPDA antennas were placed at distances varying from 5.5 m to 10 m apart, bores facing the sky, and tine elements oriented parallel (configuration 1 in figure 3.17). This orientation was chosen to minimize reflections from underground conductors, and for comparison to the size of heartbeat pulses in Moore’s Bay. In 2009-10 the heartbeat antenna was oriented downward like the signal antennas. In subsequent seasons, we would direct it at the signal antennas. Figures 3.21-3.22 display the details of the orientation.

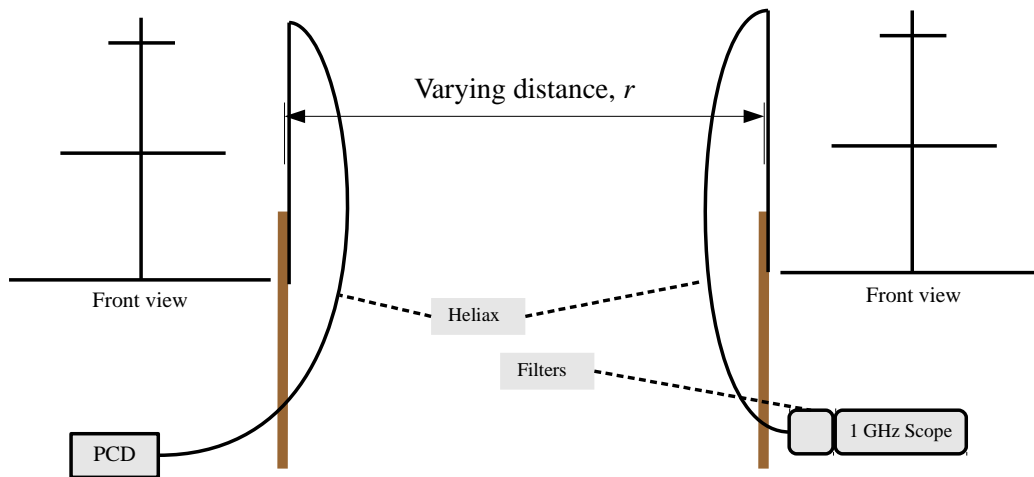


Figure 3.21: Configuration 1. Both the receiver and transmitter spines (or bores) were pointed skyward, and in the *parallel case* the tines were parallel. In the 45° and orthogonal cases, the receiver was rotated about the spine by 45° and 90° . The wooden mounts were 2 m above the ground.

The distance r was incremented by 50 cm from 5.5 m to 10 m. At each distance, three LPDA sub-orientations were checked. We refer to the case depicted in configuration 1 as the *parallel case*. The other two orientations had the receiver tines at 45° with respect to those of the transmitter, and 90° (*orthogonal case*) with respect to the transmitter. The Pockel Cell Driver was placed under an oil drum to prevent noise from the internal electronics from




Figure 3.22: Configuration 2. The receiver antenna bore was pointed upwards, with tines parallel to the transmitter antenna in the *parallel case*. In the 45° and orthogonal cases, the receiver was rotated about the spine by 45° and 90° . The transmitter tines were parallel to the ground, and each LPDA remained 2m above the ground.

interfering. Heliax coaxial cable carried the pulse out through the bottom to the transmitter. While being outside in Aldrich park placed the antennas far from most large conductors, it also exposed them to noise as well, such as kUCI (88.9 MHz), some cellular bands (≈ 900 MHz), and public safety radio communications (800 – 900 MHz). As a result the data below is filtered with both a high and low pass filter on the receiver end (NHP-300+ and NLP-600+), which allowed us to eliminate backgrounds. The corner frequency (3 dB point) of the high-pass NHP-300+ is 245 MHz, and that of the NLP-600+ low pass is 640 MHz. Figures 3.23 and 3.24 contain the results.




Figure 3.23: Plotted on the x-axis are the inverse-squared distances scaled by R_0 , with $R_0 = 5.5$ m. On the y-axis we have the relative power, with P_0 corresponding to 5.5 m. For $P \propto r^{-2}$, the expected slope is 1.0, and the linear fit is in agreement. These data are for the parallel case.




Figure 3.24: Plotted on the x-axis are the inverse distances scaled by R_0 , with $R_0 = 5.5$ m. On the y-axis we have the maximum relative amplitude, with V_0 corresponding to 5.5 m. For $V \propto r^{-1}$, the expected slope is 1.0, and the fitted slope is 1.2 ± 0.1 . These data are for the parallel case.

One way to express Parseval's theorem is that the total power of a signal in the Fourier domain is equal to the total power in the time domain. Since we are interested only in the total power at each distance for the moment, we do not have to calculate the spectra and can simply write the total power in terms of the sum of the squares of the recorded time series voltages v_i . This form leaves the units in $W\Omega$, where the resistance is 50Ω . But the power in figure 3.23 is normalized such that it does not matter what the resistance is; it is merely an overall scale factor that drops out of the analysis.

$$P = \sum_{i=0}^N v_i^2 \tag{3.14}$$

$$\sigma_P = 2\sigma_v\sqrt{P} \tag{3.15}$$

Given these two equations, the error in the relative power becomes

$$\sigma_{P/P_0} = \frac{2\sigma_v}{\sqrt{P_0}} (P/P_0 + (P/P_0)^2) \tag{3.16}$$

Although the fit in figure 3.23 is consistent with $1/r^2$ scaling with a slope of 1.0 ± 0.1 , points at the closest distances ($(r_0/r)^2 \approx 1$) seem to deviate from the trend, and we can speculate that this is due to near-field effects. The pulse amplitude should be inversely proportional to the distance: $V \propto r^{-1}$, and the fitted slope in figure 3.24 is 1.2 ± 0.1 . The points closest to $(R_0/R) = 1$ seem to deviate as with the power. Examples of the types of pulses received are shown in figure 3.25.




Figure 3.25: When the separation in the parallel configuration doubles, the peak to peak voltage drops by a factor of 2 and the power a factor of 4, as expected. This data was taken in Aldrich Park in April 2010.

It is trivial to confirm the radiation pattern of the E and H planes specified by the manufacturer (as in figure 3.10), provided one varies the angle of the receiver only in the E or H plane, for angles in $[0^\circ, 90^\circ]$. Data taken in configuration 2 (figure 3.22), however, can be used to examine how the LPDA gain changes in both planes simultaneously. The transmitting LPDA in this case was oriented such that the maximum gain is aimed at the receiver. The receiver was oriented in the same sub-configurations as in the previous experiment: parallel, 45° and 90° . Using our LPDA simulator, we can predict the gain at the measured angles and compare the power spectra at these angles to the nominal power spectrum corresponding to $\theta = 90^\circ$ and $\phi = 0^\circ$ (the alignment of figure 3.13 and 3.17 (left), denoted *bore sight* alignment). Configuration 2 is also close to the heartbeat orientation of 2010-11 and 2011-12 seasons.




Figure 3.26: The average normalized gain versus frequency for 90° in the receiver H-plane, from the Log-Perd FORTRAN code is ≈ -6 dBi.




Figure 3.27: The red/black data is the power spectral density (PSD) of the aligned case ($\theta = 90^\circ$, $\phi = 0^\circ$), scaled down by the average gain of -6 dBi at 90° in the H-plane to match the data observed at this angle (blue/grey).

The parallel case in configuration 2 is just an example of receiving at 90° in the H-plane. The simulation predicts an average gain of -6.0 ± 0.5 dBi at this angle (relative to $\theta = 90^\circ$, $\phi = 0^\circ$), and the specific gains at each frequency across the bandwidth of interest is plotted in figure 3.26. These numbers do display some periodicity with the logarithm of frequency, as expected. The lowest point is at 125 MHz, chosen to stay away from the lower bound of 105 MHz (in air) where properties like the VSWR begin to diverge. In figure 3.27, the nominal

spectrum corresponding to $\theta = 90^\circ$ and $\phi = 0^\circ$ is scaled down by the average gain 3.26 and compared to the spectrum corresponding to the parallel case. The error bands plotted are the 95% confidence levels, and the spectra match within errors.

Though the optimal answer for the spectral scaling is -8.0 dBi in the parallel case, the difference between -8.0 dBi and the answer from simulation (-6.0 ± 0.5) could likely correspond to a receiver angle that was slightly larger than 90° in the H-plane. This systematic error is not difficult to imagine considering the receiver and transmitter were located outside on a gradual slope with the receiver slightly above the transmitter. Thus, orienting the receiver directly upwards made the H-plane angle a little larger than expected, causing the drop in gain to be larger than expected by a few dBi. The power spectral density (PSD) in figure 3.27 are calculated using Welch's method [85], with a 50% overlapping rectangular window function that leads to ≈ 35 MHz frequency resolution, and excellent resolution on the y-axis (dB/Hz). Finally, the percent difference in the peak to peak voltages in the time domain of these two signals is 2.0%, assuming the -8.0 dBi scaling figure. Thus, using one scaling number, we can explain both the spectrum and the time-domain amplitude of the configuration 2, parallel case data.




Figure 3.28: The average normalized gain versus frequency for 90° in the receiver H-plane, 45° in the E-plane (azimuthal coordinate) from the Log-Perd FORTRAN code is ≈ -10.3 dBi.




Figure 3.29: Red/black data: PSD of the aligned case ($\theta = \phi = 0^\circ$), scaled by the average gain -10.3 dBi to match the (blue/grey) data at this angle ($\phi = 45^\circ$). The dip in simulated gain (fig. 3.28) near 300 MHz is reflected in the matching of the PSDs.

The next case observed in configuration 2 was the 45° case, corresponding to $\theta = 0^\circ$, $\phi = 45^\circ$. The average gain predicted by the simulation is -10.3 ± 0.8 dBi (relative to maximum). The specific gains versus frequency and the PSD of scaled bore sight data and 45 degree-case data are plotted below in figures 3.28 and 3.29. Using the predicted gain from the simulation, the PSDs agree and the percent difference in time-domain peak to peak voltages is 8%.

The third case observed in this study was the 90° case, corresponding to $\theta = 0^\circ$, $\phi = 90^\circ$. The simulation causes the average gain at $\phi = 90 \pm 1^\circ$ to diverge to negative infinity, however the gain prediction is symmetric within 10° of this point. Thus we can examine what happens when we apply the simulated gain at 80° to the bore sight data, which turns out to be -22.0 ± 1.0 dBi. The specific gains versus frequency and the PSDs of scaled bore sight data and 90 degree-case data are plotted below in figures 3.30 and 3.31. Using the predicted gain from the simulation, the PSDs agree and the percent difference in time-domain peak to peak voltages is 20%. The percent difference in peak to peak voltages is reduced to 7% if we ignore what appears to be a singular fluctuation in the bore sight negative amplitude. The scaled bore sight time domain data is plotted alongside the data for the three configurations in figure 3.32.




Figure 3.30: The average normalized gain versus frequency for 90° in the receiver H-plane, 90° in the E-plane (azimuthal coordinate) from the Log-Perd FORTRAN code is ≈ -22.0 dBi.




Figure 3.31: The red/black data is the power spectral density (PSD) of the aligned case ($\theta = \phi = 0^\circ$), scaled down by the average gain -22.0 dBi to match the (blue/grey) data observed at this angle ($\phi = 90^\circ$).

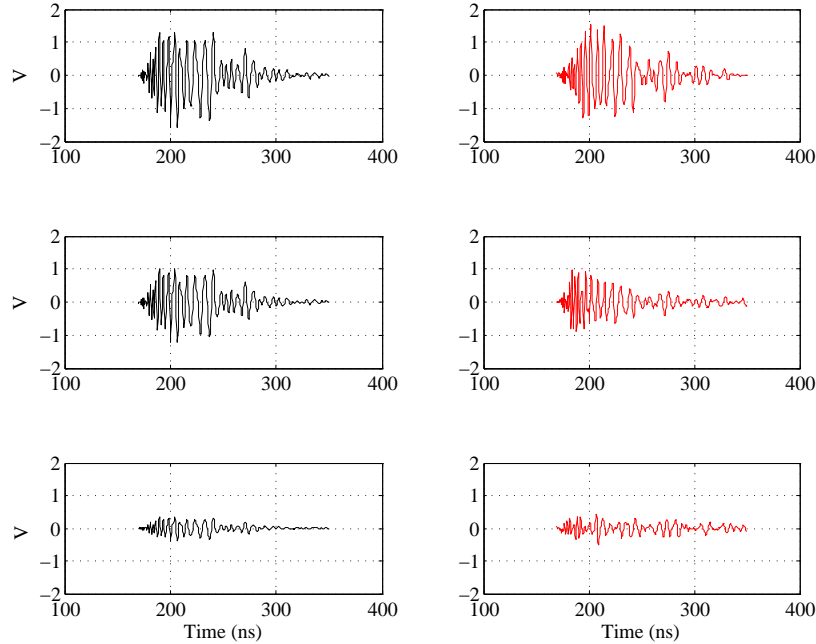


Figure 3.32: The scaled bore sight data (black) compared to the three sub-configurations (red). Top: parallel configuration. Middle: 45 degrees. Bottom: 90 degrees.

Finally, the antenna locations for the various proto-station seasons are displayed in figure 3.33. The origin of this graph is the corner of the square solar tower (figure 3.4) that held the anemometer. The central red points (North, East, South, West) mark the spine positions of the LPDAs, and the red points to either side mark the end-points of the longest tines, which are nearest to the surface. The same is true for the 2009 heartbeat position, when the heartbeat antenna was oriented in configuration 1. From Dec 2010 onward, the heartbeat was farther away, and oriented in configuration 2 with the tines at the snow surface.

The baseline distances between channel antennas were planned to be 6m, however limited length of LMR-600 cable prevented the 2009-10 setup from being exactly 6m by 6m. The figure 6m was chosen because the trigger gate of the proto-station requires that two channels trigger within ≈ 20 ns of each other (more details in the next chapter). This means that a plane wave of sufficient amplitude arriving from $\geq 40^\circ$ below the horizontal would satisfy this condition. A vast majority of signals are expected to behave this way because of the




Figure 3.33: Antenna positions, all seasons. The origin was chosen to be one corner of the solar tower (arbitrarily). The four signal LPDAs were deployed in December 2009 and were not moved. The central antenna points represent the spine location, and the edge points indicate the orientation and length of the top LPDA tine. This applies to the 2009-10 heartbeat antenna as well. In 2010-12, the heartbeat was placed farther away, and it lay flat on the snow surface with bore-sight facing the array.

shadowing effect, however, the gate will require adjustment in the presence of surface waves.

In 2009 the heartbeat antenna was in configuration 1, with different azimuthal angles for each channel. In 2010-11, we re-oriented the heartbeat antenna to be farther away, and in configuration 2. This action was taken for several reasons First, the heartbeat pulses were missing from the West channel because it was previously the last channel to be hit by the heartbeat, and the 60 ns trigger delay in T_{pp} meant that the SCA digitizer for the West channel was missing the waveform as it passed through the detector. By reorienting the heartbeat antenna, we were able to capture heartbeat pulses in the West channel as well as

the others. Using configuration 2 also shortened the width of the pulses, making them more ideal for cross-correlation studies to determine timing resolution.

3.2.4 Heartbeat Events from Different Periods

Understanding the structure of the time-dependent waveforms recorded by the detector LPDAs is essential. It is also important to verify that the response of the LPDAs and subsequent pieces of the data acquisition chain does not depend on time. The heartbeat class of events recorded by the prototype confirm that the same waveform structure is recorded from the same experimental setup examined in different periods of station operation. Shown in figures 3.34 - 3.36 are heartbeat waveforms from the austral summer of 2010, both before and after the wireless tower removal, and heartbeat waveforms from after the station restarted in October 2010.

The time structure and voltage structure of the events have not changed over the course of the year, as the antennas were buried under 4 ft. of snow. Only the beginning of the West channel pulse is digitized due to the location of the heartbeat transmitter, and a 60 ns delay associated with the trigger. In December 2010, the heartbeat transmitter was moved to a location that corrected this issue. Heartbeat data from the new location are also constant in time, and an example event (with the visible West channel pulse) is shown in chapter 4, figure 4.32.




Figure 3.34: These examples of heartbeat waveforms recorded by the proto-station North channel. (Bottom): event # 659176, run 100, from January 2010. (Middle): event # 844229, run 160, from February 2010 (after wi-fi removal). (Top): event # 17745, run 100 (after winter restart), October 2010.




Figure 3.35: These examples of heartbeat waveforms recorded by the proto-station East channel. (Bottom): event # 659176, run 100, from January 2010. (Middle): event # 844229, run 160, from February 2010 (after wi-fi removal). (Top): event # 17745, run 100 (after winter restart), October 2010.




Figure 3.36: These examples of heartbeat waveforms recorded by the proto-station South channel. (Bottom): event # 659176, run 100, from January 2010. (Middle): event # 844229, run 160, from February 2010 (after wi-fi removal). (Top): event # 17745, run 100 (after winter restart), October 2010.

Chapter 4

Icicle 1 Prototype Data Analysis

The primary purpose of the ARIANNA prototype station was to study radio frequency backgrounds of the ARIANNA site (Moore's Bay) and to begin assessment of relevant technologies for remote autonomous operation. A large range of threshold conditions, on the CAGES board, and detector operation modes were explored over the course of 3 years. Specifically, high and low-power run states were explored to study the trade-off between power consumption and live-time. These changes resulted in useful insights on station operation, control, data archiving and power generation, but it also created highly variable trigger rates. Fortunately, the operation of the prototype station was stable enough for a limited period of time each year for a dedicated search for neutrino signatures in the data samples. This chapter is devoted to the discussion of total detector operational live-time and the subset of data in stable periods that were used to search for characteristic neutrino signatures.

Due to the frequency bands of the trigger system, it is possible to collect thermally triggered events with relatively more power at high frequencies (660-990 MHz) as expected from neutrino signals, or at low frequencies (130-460 MHz), which form the dominant contribution to

the time dependent waveform. Consequently, high and low band thresholds that emphasize the high frequency band result in thermally triggered events with relatively little power in the lower frequencies. Thermal triggers from the first and third seasons are a result of such conditions. It is expected that the frequency content of each trigger band is approximately independent, so highly unusual fluctuations in power in the emphasized high frequency band would not necessarily show unusual fluctuations in variables relying on time domain quantities. We study this effect and use it to optimize the criteria to remove background events generated by thermal fluctuations. This analysis also contains Monte Carlo comparison to the expected mathematics of non-uniform bandwidth thermal noise, primarily from the work of S.O. Rice [72].

The Icicle1 prototype also collected housekeeping data in the form of voltage measurements on the ADC of the DMM-AT board in the CPU stack. The purpose of this data was to characterize the local detector environment in the form of wind speed, temperature, power consumption, and temperature. From the available housekeeping data, useful design properties, such as live-time, can be revealed. Several examples are the minimum solar panel rating (given a desired detector live-time), the maximum achievable live-time with and without wind power, and the lowest temperature our electronics must withstand.

This chapter is organized as follows: sections ?? and 4.1 cover the housekeeping data, section 4.2 defines the different event classes and discusses the mathematical properties of thermal noise, and sections 4.3 through 4.6 present the data from the 2009-10, 2010-11 and 2011-12 seasons of data, along with a characterization of the trigger action.

4.1 Housekeeping data, 2009-12

The housekeeping data reveals interesting properties of the local environment of Moore's Bay. These data will be used in subsequent seasons to design robust detectors capable of withstanding the harsh climate and providing useful data on a year-round basis. To this end, measurements were made over three seasons of the yearly temperature profile, maximum achievable live-time (as a fraction of one year), and the battery buffering necessary for operation on wind power during the winter.

The purpose of the housekeeping system is two-fold: the housekeeping daemon `hkd` (section 3.1.2) controls the list of devices receiving power via the electronic relays, and it records the data sampled by the ADC. The ADC is the 16-channel, 12-bit, 10 V dynamic range device on the DMM-AT board in the CPU stack of the station. The beginning collection of measurements in the 2009-10 season was comprised of the battery voltage, wind-speed at 7 ft. above the snow surface, the external temperature, and the temperature inside the RF sealed electronics box. Figure 4.1 shows our initial prototype data.

Figure 4.1 contains all data from the first deployment of `Icicle1` until December 2010 when ARIANNA expedition members returned to Moore's Bay to make repairs and modifications. The x-axis is time, expressed in days (Universal Coordinated Time - UTC), with the origin set to January 1st, 2010, at midnight. Around the 90th day of the year 2010, the sunlight and wind power were no longer sufficient to keep the battery charge above the shutdown voltage. At a battery voltage of ≈ 10 V, the internal DC-DC converters can no longer provide a stable 5 V voltage for the CPU to remain active, and this causes a reboot, or a shutdown.

The data gap continues from day 93 to day 284. Day 284 is October 11, 2010, when the sun has begun to rise over the Antarctic horizon enough to power the station all day. The unix clock in the CPU stack drifts from UTC times obtained from satellites. As discussed in chapter 3, the GPS daemon `GPSd` corrects the unix clock once it drifts by more than

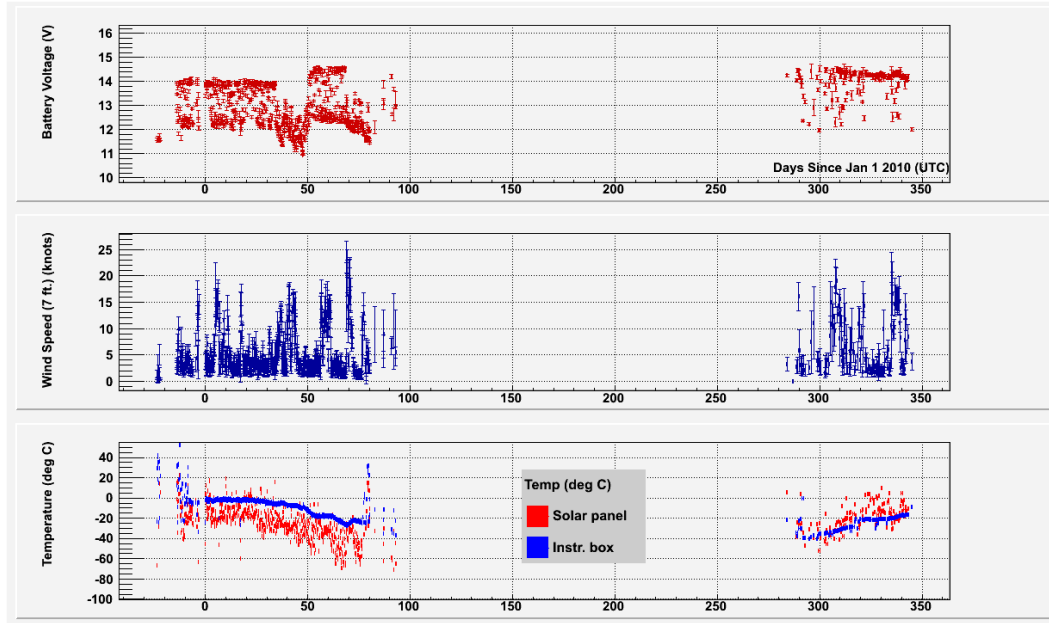


Figure 4.1: Housekeeping data set for the first season of the prototype deployment. The x-axis is time units (UTC days since January 1, 2010). The system AGM battery voltage is the upper graph, the wind speed in knots is shown in the middle, and the temperature from the back of one solar panel and inside the instrumentation box are shown at the bottom.

100 seconds from the time seen by the GPS system, which is also set to UTC. After being de-powered for several months, the CPU backup battery (responsible for holding the BIOS settings and clock time between reboots) was drained. Consequently, it generated a time from *before* the station was deployed in 2009, which in turn confused the GPS updating system. Thus, the unix times for the reboot period (days 284 to 345) are not useful. During this period, GPS times associated with the data runs were used in figure 4.1. The time dependence of the battery voltage (measured prior to the Vicor DC/DC converter system) is quite complex, and primarily associated with the charger controller. A few of the features are described next.

When the AGM battery is at full charge and connected to solar panels, the nominal voltage for a full charge is 14.4 V. The solar panels are at higher voltage than the battery (16.2 V closed circuit into a 12 V system). At a set voltage on the charge controller (14.4 V),




Figure 4.2: The battery voltage for days 4-12 in 2010 demonstrate a variation with a period of 1 day. The fitting function is $V(t) = A \sin(2\pi ft + \phi) + C$. The amplitude is 0.93 ± 0.02 V, and the frequency is 0.998 ± 0.003 days⁻¹. The phase is 0.7 ± 0.1 radians, and the constant offset is 13.14 ± 0.02 V.

the solar panel line is diverted across a resistor, and dangerous voltages are avoided. This switching occurs at a rate of ≥ 1 Hz in the lab when the battery is fully charged. There is a daily variation in the battery voltage between 12-14V. This variation is caused by the daily variation in solar intensity. Figure 4.2 illustrates the daily variation for a week in January where the conditions were especially stable. Although the effect is not necessarily simple harmonic oscillation, it is described well by a sine function that indicates simple solar variation. The nominal voltage under load is given by the downward swings near 12.2 V, and the upward swings are near the limiting voltage of the charge controller. The period of the fit is one day, and the phase indicates the high voltages are near mid-day in local time.

Just prior to 50 of 2010, the HK data indicated a unusual drop in instantaneous (not average) battery voltage which eventually lead to a reboot of the CPU. Before this happened, waveform data acquisition was deactivated, reducing the power consumption from 30 W to

10 W. The increase in windspeed and the decrease in solar power are typical of large weather fronts in the austral autumn. Once the wind dissipated and the sun reappeared, the battery voltage increased to 14.5 V (maximum), and rebooted and the station. The shift in voltage is caused by the lower power consumption, confirmed in the laboratory freezer at -30° C. After sunset, the Forgen 1000LT was needed to generate enough wattage to sustain housekeeping data collection, but this unit needs 30 knot winds to produce 1 A of current. The wind speed at the height of the generator never reached this level, limiting operations.

Finally, two thermo-electric transducers measured the temperature of the back of a solar panel and within the instrumentation box. Naively one expects the buried instrumentation box sensor to correspond to the running average of the thermistor attached to the solar panel, however inside the box the effect of the electronics on the temperature is non-negligible. The CPU stack alone consumes 7 W and throws off heat which raises the ambient temperature. Because most of the electronics were deactivated after day 50, there was a significant dip in the temperature inside the instrumentation box. The rise in temperature measurements just before shutdown is non-physical, most likely due to changing voltage levels with a low battery. The temperature sensors, which convert a temperature to a current, rely on a steady voltage input to convert the current to a voltage. Just after the restart, the instrumentation box was still at the ambient temperature. Around day 320, there was a small jump in the internal temperature, probably caused by the peripheral devices attempting to turn on after reading the default Hkd settings rather than low power settings.

In addition to the on-board temperature sensors, two Madgetech temperature data-loggers separately recorded the same temperatures for the purposes of calibration and recording an entire year in the event of a shutdown. The results are shown in figure 4.3. The data loggers reveal a two pieces of information that the on-board sensors missed. The internal temperature does reflect the running average of the external temperature in the absence of internally generated heat, and the minimum internal temperature reached under such conditions is -35°

C. Thus, all future electronics designs for ARIANNA must be able to withstand temperatures of -40°C . Ambient temperature fluctuations are expected to diminish in subsequent seasons as the stations are further insulated by layers of snow.




Figure 4.3: The solar panel data logger (red) and the internal data logger (black) recorded temperatures throughout the year at the first ARIANNA site.

During the second and third austral seasons (2010-11 and 2011-12), similar data was collected, and a higher power wind turbine was deployed in 2010. The turbine was a five-blade, furling, 100 W Aerogen-6 unit. The temperature sensors were removed, and the anemometer was transferred to Icicle2 (1 km from the proto-station). With several seasons' worth of housekeeping data, important station characteristics, such as live-time versus solar and wind capabilities, can be assessed. This data also informs our sense of the low level of power consumption needed to sustain operations all year round. Figure 4.4 shows the global time series for the battery voltage, 5V DC-DC converter line, and the wind speed, and table 4.1 summarizes the beginning and end of each data run.

The data in the first two time blocks is the same as in figure 4.1; the station began operations in 2009 and continued to October 2010. The winter shutdown is expected because of the

UTC Day since Jan 1, 2010	Start/Stop	High-power mode activated
-14.0	Start (first season)	(-14.0 to 45.75)
93.0	Stop (winter shutdown)	–
284.0	Start (winter reboot)	(354.35-411.5)
411.5	Stop (winter shutdown)	–
713.0	Start (manual reboot)	(713.0-792.0)
792.0	Stop (winter shutdown)	–

Table 4.1: These are the start/stop dates for Icicle1 operations. The first column indicates the time, the second indicates the action, and the third indicates when waveform data acquisition took place (high-power mode). See text for details.

known inefficiency of Pb-acid batteries in extreme temperatures. In December 2010, Dr. Steve Barwick and the author turned off the station to make adjustments upon arriving in Moore’s Bay. Thus, the ensuing data gap around day 350 is non-physical. Although there are nominally-digitized thermal triggers and heartbeats from the first winter reboot period, the UNIX clock was reset. This makes the data acquired between the winter reboot and the 2010 expedition more difficult to analyze because the time of each event is meaningless (these events are not included in the waveform analysis). The station continued to operate normally, observing similar wind speed structure to the prior year. The wind rarely tops 30 knots at 7 ft. above the snow surface (the anemometer was eventually 4 ft. above the surface due to snow accumulation throughout the year). Unlike the prior year, the station was operated in high-power (30 W) mode until solar power faded to non-sustainable levels on February 27th.

This strategy was chosen to acquire a maximum of thermal triggers, rather than optimize for live-time in low power (10 W) mode. The snow began to accumulate to the point that the solar panels were buried some time during the winter. As a result, the station did not recover from the 2011 winter because the buried solar panels produced no power once the sun returned. The author, along with Joulien Tatar and Dr. Eric Berg ventured again to Moore’s Bay in 2011 and were able to acquire another 90 days of data. Figure 4.4 demonstrates that

we can achieve 45% live-time with solar power only, while consuming 30 W. The live-time is estimated by considering days 0 to 360 and subtracting 200 days for the data gap. One can take any year-long period and obtain a similar answer, as long as it is taken into account that days 600-700 would have data if not for the burying .

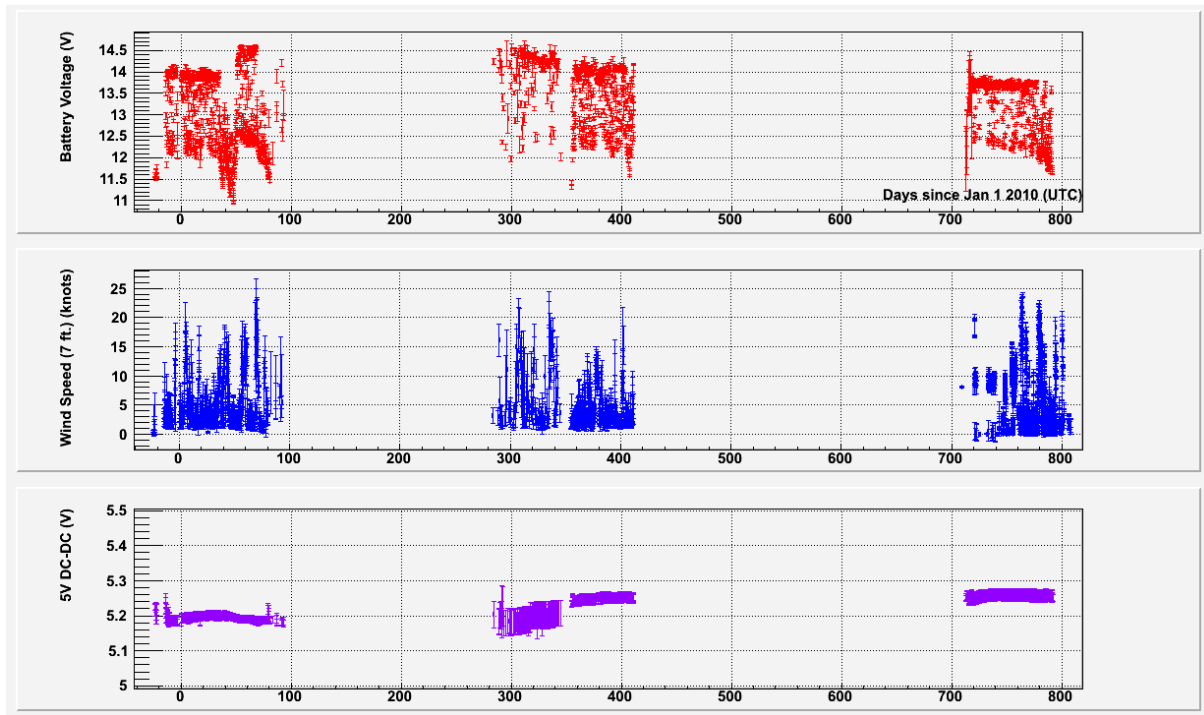


Figure 4.4: The 2009-12 housekeeping data. (Top) the global battery voltage throughout the three seasons. (Middle): the global wind speed atop the solar frame. (Bottom): the 5V DC line, responsible for powering the CPU stack.

The maximum wind speed was 43 knots, occurring on October 10, 2010, as the station rebooted for the first time after the austral winter was fading to austral spring. Figure 4.6 is a normalized histogram of the averaged (binned) wind data from figure 4.4 (middle). The average wind speed is just under 5 knots (figure 4.5), the cut-in speed of the Aerogen-6 unit. The cut-in speed is the wind speed at which the internal circuitry of the unit connects the generator to the charging lines. This way, the back-resistance from the generator will not stop the blades from turning in low winds. A 5 knot cut-in speed is typical of most wind generators of this size. While the wind speed distribution falls exponentially, and the spike




Figure 4.5: This data has been digitized from the manufacturer’s manual of the Aerogen 6 wind turbine. The cut-in speed is 5 knots and the slope is 0.75 A/knot.

at 9 knots is caused by two especially powerful storms between days 720 and 740. A simple relationship between current and windspeed for the Aerogen unit is $I(w) = 3/4(w - 5)$, where w is the wind speed in knots (the output is non-linear above 25 knots). Figure 4.6 indicates that 10 knots of wind would be required to provide 2.5 A (to operate the station in full power mode), and this windspeed is one factor of the r.m.s above the mean.

The housekeeping data from Icycle2 will be published much more extensively in [81]. The ARIANNA stations experience other variations in battery voltage besides the effect caused by the daily solar intensity variation. The Icycle1 battery voltage (red, upper) alongside the Icycle2 battery voltage (blue, lower) data is shown in figure 4.7. This data comes from the third austral season (2011-12) during which three Pb-acid batteries were deployed in Icycle2, and two Pb-acid batteries for Icycle1. The cycles are correlated in time, which is given on the




Figure 4.6: The wind data from figure 4.4 has been plotted here as a histogram.

x-axis in days since Jan 1, 2012 (5 days per division). The increased storage capacity was meant to provide a buffer between wind storms and cloud cover. The times when the voltage decreases below 12 V usually correspond to strong wind and low solar intensity, indicating a weather system moving across the detector. Because Icicle2 was deployed with 30W solar panels, rather than the 50W units of Icicle1, the battery voltage is below maximum more often despite having one additional battery.




Figure 4.7: The red data (top) is the battery voltage for the prototype station at the beginning of the 2011-12 season, and the blue data (second from top) is the Icicle2 battery voltage.

Based upon the global housekeeping data set, estimates of the future station requirements can be made. First the total livetime of a station that is powered exclusively from wind is

calculated using wind speed data from 4.4, assuming it is representative of wind conditions over the entire year. This is conservative, since average wind speeds are likely to be larger during the winter months. The simulation, developed with the help of Daniel Kennedy at the University of Kansas [12] assumes power generated by a Bergey 200W wind turbine (similar in blade geometry and power rating to Air 40 by Southwest Windpower, which was deployed in 2012). The battery storage capacity in 4.8 refers to actual storage capacity at the ARIANNA site, which is a fraction of the rated battery storage. For Pd-acid batteries, the storage capacity at -30° C (the mean average temp of the firn ice near the surface) is $\approx 30\%$ of room temperature rating. For Li ion batteries, the storage capacity is nearly 80% at -30° C. There are several positive outcomes indicated by figure 4.8, in light of new ARIANNA technology. The latest data acquisition technology is expected to require just 7.5 W of power, not 15W as assumed here. Also, solar power is sufficient to generate power for about half the year, so wind is only necessary during the winter months. Therefore, 20 Amp hr of battery storage provides an additional 50% of operational live-time during the winter months. In 2012, stations were deployed with Lithium batteries with 80 Amp hr at -30° C. Thus, if the wind storms of figure 4.4 remain frequent enough, future ARIANNA stations should be able to operate year-round.

Solar panels perform well in the Antarctic environment due to the high reflectivity of the snow. A solar power model was constructed [11] to determine how much live-time could be achieved with only solar power. Figure 4.9 provides an estimate of total live time for an ARIANNA station as a function of power production by the solar panels. The power system in the model consists of 4 solar panels, oriented to face due north, south, east and west. Each solar panel produces the specified power for normal incidence, and reduced by $\cos(\theta)$ for different astronomically accurate sun angles, where θ represents the deviation angle between the sun and normal direction of the solar panel. The solar panels have power rating specified for normal incidence for noon sun at a location of 45 degrees latitude. We have incorporated a model of solar irradiance for clear skies (typical of the Antarctic weather) as




Figure 4.8: Estimated live-time (as a fraction of 1 year) from wind generator as a function of the capacity of the battery in units of (Amp hr). In this study, the ARIANNA station is assumed to consume 15W, and that the observed wind data is typical all year round [12].

a function of zenith angle, but the model overproduces sunlight for sun angles within 5° of the horizon.

The data plotted in figure 4.9 should be interpreted as an upper limit to the fractional live-time. The four curves correspond to different loads: 7.5W (upper), 10W, 22W, and 30W (lower). We also take into account the significant additional power generated by solar panels on Antarctic snow due to significant reflection and scattering by the snow surface. A solar panel rated for 30W delivers 30% more in Antarctica for sun angles of $\approx 30^\circ$ above the horizon. Part of the live time is obtained also by running off batteries once the sun has set. The simulation uses a battery capacity of 150 A hr, again leading to an overestimate of the operational live time. The simulation shows that live-time based on solar power and large battery buffering can achieve more than 50% of the year, but optimistically it is no larger than 60% of the year. The improvements are almost as good with a much smaller battery;

the solar panels can produce enough current total current to power the station in summer. In general, a large battery is useful for buffering variations in cloud cover, and the short nights during the transitions between 24-hours of sun and darkness.




Figure 4.9: The live time as a fraction of a 1-year period versus the power rating of the four solar panels. The four curves are for varying levels of station power consumption: 7.5W (top), 10W, 22W, and 30W (bottom) [11].

Notice that the result for four 30W solar panels is close to the 45% result obtained earlier, for a 30 W station. The result from earlier, however, was obtained with data taken during low power mode, around 10 W for some of the time. The agreement between the model and the data is somewhat justifiable since the model assumes three times the storage capacity. If there was 150 A-hr of storage during the first season the station would have operated in high power mode until early April. In 2012, the new stations will consume just 7.5W, gaining an expected month in live-time compared to previous deployments.

4.2 The Icicle1 Prototype event classes and the properties of thermal background fluctuations

There are several types of waveforms collected with the prototype station: forced triggers, thermal triggers and heartbeats. Each event class serves a purpose. Forced triggers characterize properties of the thermal noise background and therefore quantities like the system temperature and the power spectral density. Studying thermal triggers helps to develop cuts that remove background events posing as neutrino signals. Heartbeat triggers are created by a pulser built into the stations to stand in for large signal events and to help understand the timing of the system.

4.2.1 IceSoft event classes - forced triggers, thermal triggers, and heartbeat triggers

During normal waveform data collection, the Icicle1 prototype and the other ARIANNA stations collect three classes of events. To establish and characterize the thermal noise collected and amplified by the LPDA and LNA system, *forced triggers* are recorded at a regular time interval, the *forced trigger period*. There are several useful calculations with forced trigger data. Forced triggers sample the underlying gaussian distribution of thermal noise in the system. A forced trigger does not satisfy the trigger settings; the station simply reads out the data in the digitizer as if it had caused a trigger. Although they are not very powerful events, forced triggers stand out from other triggers in that they occur regularly (according to the setting in the Acqd configuration file). The Icicle1 and Icicle2 do not tag forced triggers, so the regular spacing in time is crucial for identifying and removing them from populations intended for neutrino searches. Stations deployed in 2012 and later tag forced triggers.

A *thermal trigger* is a randomly occurring thermal fluctuation that satisfies the trigger settings. Any thermal trigger could actually be a neutrino or cosmic ray, because they must satisfy the trigger. They are randomly spaced in time, with a rate that depends on the threshold. Typically, if the waveforms in two different channels are above the channel threshold within 8.3 ns of one another, this satisfies the trigger settings. Known as the gate, the number 8.3 ns is two clock cycles of a $6\times$ speed 40 MHz digital clock. The channel thresholds are simply a number of ADC counts that are proportional to the event rate (a higher number of counts yields a higher event rate). The physical threshold is in the Fourier domain; the integrated power through a tunnel-diode over a 8.3 ns window must cross above a certain level. Even though the trigger is in the frequency domain, there still exists a relationship to the time domain.

Given that $x(t)$ is some square integrable signal in time, with a Fourier transform $X(f)$ in unitary normalized form, Parseval's theorem in physics and engineering is often written as

$$\int_{-\infty}^{\infty} |x(t)|^2 dt = \int_{-\infty}^{\infty} |X(f)|^2 df \quad (4.1)$$

The total power in the time domain is proportional to the square of the maximum amplitude of a signal, and the total power in the time domain is equal to the total power in the Fourier domain (equation 4.1). Therefore, the square of the amplitude of a signal is proportional to the total power in the Fourier domain. The integral of the quantity $|X(f)|^2$ (over a finite time) must be over threshold to trigger `Icicle1`. Therefore, the probability that a pulse will trigger `Icicle1` is proportional to its amplitude squared. For example, consider a gaussian pulse of the form

$$x(t) = \exp(-at^2) \tag{4.2}$$

$$X(f) = \sqrt{\pi/a} \exp(-\pi^2 f^2/a) \tag{4.3}$$

One can easily check that this transform pair satisfies Parseval's theorem by using the identity $\int \exp(-u^2) du = \sqrt{\pi}$. The amplitude of $x(t)$ is one, and if we were to scale it, the amplitude of $X(f)$ would scale with it, because the Fourier transform is a linear operator. Therefore, the integral of $|X(f)|^2$ (which is what the trigger of the prototype checks) is directly proportional to the square of the amplitude of $x(t)$. Now, the proto-station trigger cannot integrate forever; it integrates over roughly 8.3 ns. Thus, the assumption that the integral of $|X(f)|^2$ is proportional to the square of the amplitude of $x(t)$ must hold on finite (short) time scales.

Suppose that the gaussian pulse becomes increasingly narrow in the time domain, that is, $a \rightarrow \infty$. Then we need not worry; the integral of $|x(t)|^2$ covers the whole pulse, Parseval's theorem holds, and we may associate the square of the signal amplitude with the integral of the power. If $a \rightarrow 0$, then the pulse is no longer short (compared to the integration window) and the integral of $|x(t)|^2$ doesn't contain the whole pulse. However, if we identify $a = 1/2\sigma_t^2$ in the time domain, and $\pi^2/a = 1/2\sigma_f^2$ in the frequency domain, we find that the width of the signal in the time domain and the width of the spectrum in the frequency domain obey an *uncertainty principle*:

$$\sigma_t \sigma_f = \frac{1}{2\pi} \tag{4.4}$$

If the signal width becomes much longer than a trigger integration time, then the width of the distribution in the frequency domain narrows, approaching the simple case where all the power is concentrated on the DC component. In that case, we have exactly $P = V^2/R$ and the integral of the constant power over a short time scale is P times the time window, which is proportional to V^2 . Thus, for signals that are either much shorter or much longer than the trigger integration window, the integral of the power over this window is proportional to the signal amplitude squared. It turns out that [58] equation 4.4 represents a fundamental identity from the Fourier analysis of waves - the multiplication of the widths in the frequency and time domains is constant. Thus, for a variety of pulses the power in the Fourier domain is proportional to the amplitude in the time domain. Thus thermal triggers with amplitudes that are large compared to the rms noise are interesting. They are the farthest above threshold, making them better candidates for neutrino signals. A large amplitude also makes possible timing reconstructions and other cuts that make them distinguishable from a potential Askaryan pulse.

The final class of event collected by the Ice1 prototype is the *heartbeat trigger*. The heartbeat pulser is an Avtech AVP-AV-1S-P-UCIA pulser with adjustable width and height. For the data runs we set the width to be 1 ns, and the height to be between 5 and 10 V, depending on our needs at the time. The location of the heartbeat antenna is given in figure 3.33. The heartbeat pulse propagates from the heartbeat pulser, through LMR-600 coaxial cable, and is radiated from the LPDA antenna. The pulse travels back across the detector, causing a trigger with pulse amplitudes significantly larger than the average thermal trigger (see below). These pulses serve two purposes. As with the forced triggers, they are spaced evenly in time, except the period is 1 second. The housekeeping daemon activates the pulser for a short time at the beginning of a data run.

Thus at the beginning of all data runs, we observe the heartbeat triggers and know that the detector's trigger is operating normally. The second purpose of the heartbeat triggers

is to constrain the timing resolution of the detector. The heartbeat antenna was oriented in configuration 1 with four different angles for the four channel LPDAs during the first season, making timing reconstruction difficult. Though this modifies the power spectrum in a simple way (section 3.2.3), the pulse shapes are slightly different. In 2010-11 and 2011-12, the heartbeat pulser was re-oriented in configuration 2 with respect to the channel LPDAs, and the results improved enough to make an assessment of the timing resolution.

4.2.2 Mathematical properties of thermal noise, informing the Monte Carlo simulations

Any antenna and amplifier system will record thermal noise according to the system temperature. The system temperature is really a sum of several sources of Johnson noise, which we will discuss in turn. Johnson noise is merely the open circuit thermal noise across a resistance due to the temperature of the resistance. The system temperature of an antenna, transmission line and receiver/amplifier system is given by equation 4.5, assuming a lossless transmission line [63]. For the ARIANNA bandwidth, the coaxial transmission lines are LMR-600, with a 0.6 dB loss at 1 GHz for 6 m of cable.

$$T_{sys} = T_A + T_{AP} \left(\frac{1}{\epsilon_1} - 1 \right) + T_R \quad (4.5)$$

All temperature units are in Kelvin. T_{AP} is the physical temperature of the antenna and ϵ_1 is called the thermal efficiency parameter. T_A is the *antenna temperature*, thought of as the Johnson noise across the radiation resistance of the antenna. If P is the total power being received, with no direct signal other than noise, then $P = kT_A B$, where k is Boltzmann's constant, and B is the bandwidth. If the antenna is directed at a very hot region of space, the

noise fluctuations of the voltage at the antenna terminal will be larger than if the antenna is directed at a cool region of space. One way to get a handle on the antenna temperature is to direct the antenna at the region of space where the signal originates, record the temperature, and then move the antenna off the source, recording the temperature a second time (a common technique in radio astronomy). Let T_{on} be the first temperature, T_{off} be the second temperature, Ω_A be the beam area, and Ω_s be the angular extent of the source temperature. The temperature of the source of radiation is then [63].

$$T_s = \frac{\Omega_A}{\Omega_s}(T_{on} - T_{off}) \quad (4.6)$$

In the case of the ARIANNA system, we know that the noise source is the ice and snow around the antenna, and that it is isotropic:

$$T_{ice} = \frac{4\pi}{\Omega_A}(T_{on} - T_{off}) = (T_{on} - T_{off})D \quad (4.7)$$

Here we have used the familiar expression for the directivity, $D = 4\pi/\Omega_A$. The ratio of solid angles has been flipped between the two equations. In the first case, the source solid angle is subsumed by the beam area, whereas in the second case, the thermal bath solid angle subsumes the beam area. Solving for T_{on} , we have

$$T_{on} = T_{ice}/D + T_{off} \quad (4.8)$$

T_{off} is just the ambient temperature in Moore's bay, the environmental temperature the antenna would see if the ice shelf were not there. However, figure 4.3 tells us that the external ambient temperature and T_{ice} are equal on average, so we can simplify and obtain $T_{on} = T_A = T_{ice}(1 + 1/D)$. We have $D = 5.0$ from figure 3.9, and an average temperature of $\approx -17^\circ\text{C}$ from figure 4.3. We conclude that $T_A = 307\text{K}$.

The factor T_R in equation 4.5 is the *noise temperature* of the receiver, in this case the low-noise amplifier, which has a gain of approximately 70 dB. The noise introduced just by the amplifier can be thought of as the Johnson noise added to the source resistance (to which the amplifier is matched) by the amplifier - that is, the source resistance appears hotter due to the amplification. This noise is typically quantified by a *noise figure* ([57], chapter 7):

$$NF = 10 \log \left(\frac{kTBR + V_{rms}^2}{kTBR} \right) = 10 \log \left(1 + \frac{V_{rms}^2}{kTBR} \right) \quad (4.9)$$

In equation 4.9, k is Boltzmann's constant, T is the physical temperature of the amplifier/source resistance circuit, B is the bandwidth, R is the source resistance (taken to be 50Ω unless otherwise specified), and V_{rms} is the rms noise voltage at the input of the amplifier. The observed rms noise for a 50Ω terminated ARIANNA LNA is $\approx 30 \text{ mV}$, with a gain of 70 dB and a bandwidth of 0.1 to 1 GHz. Accounting for the gain, equation 4.9 gives a noise figure of 1.5 dB. The conversion between noise figure and noise temperature is

$$T_R = T_{AP}(10^{NF/10} - 1) \quad (4.10)$$

The physical temperature of the antenna, amplifier, and transmission line is assumed to be

uniform and constant. For ARIANNA, the average physical temperature can be between -35°C and 0°C . Taking an average temperature of -17°C yields a temperature in Kelvin of 256 K, and assuming the 1.5 dB noise figure produces a T_R value of 106 K. This is comparable to the answer for the receiving system of ANITA (ref. ANITA design): 140 K. For now, we assume that the thermal efficiency parameter (ϵ_1) is close to unity because the antennas are made of aluminum. Adding the noise temperature and the antenna temperature, we obtain $T_{sys} = 413$ K.

Knowing the system temperature, we can calculate the rms level of power fluctuations in the thermal noise observed by the system. From [30], the rms level of power fluctuations is

$$\Delta S = \frac{kT_{sys}}{A_{eff}\sqrt{\Delta t\Delta\nu}} [Wm^{-2}Hz^{-1}] \quad (4.11)$$

Equation 4.4 above demonstrates that $\Delta t\Delta\nu$ is a small, constant parameter. Additionally, ARIANNA does not have infinite bandwidth, so $\sqrt{\Delta t\Delta\nu} \approx 1$. All that's left is to assess the effective aperture of the LPDA antenna, A_{eff} . The LPDA has a frequency-independent gain (directivity), and therefore a frequency-independent beam area. The effective aperture is related to the beam area: $A_{eff}\Omega_A = \lambda^2$. The average effective aperture can be derived as follows. Let the bandwidth be $B = \nu_2 - \nu_1$ and the beam area be $\Omega_A = 4\pi/D$.

$$\langle A_{eff} \rangle = \frac{c^2 D}{4\pi B} \int_{\nu_1}^{\nu_2} \nu^{-2} d\nu \quad (4.12)$$

Using 0.1 GHz and 1 GHz for the limits of the integral, we obtain $\langle A_{eff} \rangle = c^2 D / (4\pi \nu_1 \nu_2) \approx 0.36 \text{ m}^{-2}$. Inserting this and T_{sys} into equation 4.11, we have $\Delta S = 1.6 \times 10^6 \text{ Jy}$. This

number can be compared to the estimated ΔS of ANITA, $\approx 2 \times 10^6$ Jy. The improvement comes from the increase in aperture for little difference in bandwidth.

This thermal noise is small compared to the typical neutrino pulse. Consider the coherent radiation emitted from a shower comprised of electrons and positrons, with an energy of 10^{17} eV. With a Bjorken inelasticity of ≈ 0.2 , and a negative charge excess of 20% (ref obs of Ask effect in ice), $N_e = 4 \times 10^6$. The total energy radiated by a single particle in this shower is [30]

$$W = \left(\frac{\pi h \alpha}{c} \right) L \left(1 - 1/(n\beta)^2 \right) (\nu_{max}^2 - \nu_{min}^2) \quad (4.13)$$

The front factor is comprised of Planck's constant, the fine structure constant, and the speed of light in vacuum. The factor $L \approx 6$ m is the track length of the shower. Here the index of refraction $n = 1.8$, the boost $\beta = 1$, and the bandwidth we take as 0.1 to 1 GHz. The total energy is proportional to the square of the charges if they radiate coherently: $W_{tot} = N_e^2 W$. Ultimately, this estimation yields 1.8×10^{-12} J of energy close to the shower. Let the Cherenkov angle be $\theta_c = 56^\circ$. The width of the Cherenkov angle can be estimated by $\Delta\theta_c = (c/\nu_{ave} L) \sin \theta_c \approx 4^\circ$. The shower emits the total energy into a solid angle of $\Omega_c = 2\pi \Delta\theta_c \sin \theta_c \approx 0.36$ sr. Thus, we have $W_{tot}/\Omega_c \approx 5 \times 10^{-12}$ J/sr.

Suppose this shower travels 1 km before hitting an ARIANNA station. The LPDA effective aperture is ≈ 0.4 m², so the Askaryan radiation would encompass 4×10^{-7} sr, and have 2×10^{-18} J of energy (accounting for the RF attenuation in a moment). The pulse is bandwidth-limited, about 2 ns wide, with a bandwidth of 500 MHz (such that $\sigma_t \sigma_f \approx 1$). So we conclude that the intensity is 4×10^{-18} W m⁻² Hz⁻¹. Suppose the pulse experiences

an average attenuation length of 400 m over the 1 km radio pulse path, recalling that the power attenuation length is one-half the field attenuation length. That means five factors of e, or a factor of 0.007. Scaling down the intensity by this factor we get 8×10^6 Jy. Remember that the thermal noise level derived about is equivalent to 1.6×10^6 Jy. Thus, the *signal to noise ratio* is $8/1.6 \approx 5$. While this is a coarse estimation, it demonstrates that the thermal noise caused by the system temperature of 406 K is small enough compared to the power of an Askaryan pulse from a UHE neutrino.

Thus far the noise has been assumed to be gaussian white noise, with equal power across the bandwidth. However the spectral response of the ARIANNA amplifier, for example, is not uniform (figure 4.10). To understand the thermal backgrounds of ARIANNA, which ultimately lead to thermal triggers, the properties of the non-uniform bandwidth must be taken into account. The analytical calculations in this section come primarily from S.O. Rice, and the technique for generating corresponding Monte Carlo waveforms is from [35]. The symbol ϕ_k will represent the normalized average PSD of the forced triggers. Given enough forced trigger events, the forced trigger data is an accurate measurement of the noise seen by the trigger, because it accounts for all frequency-dependent effects throughout the detection chain. The procedure has four steps, with N representing the number of samples in an event:

- Generate and store $2N$ spectral weights that are the square roots of the magnitudes of the desired PSD ϕ_k scaled by the number of samples: $H_k = (2N\phi_k)^{1/2}$
- Generate two sets of $2N$ independent, zero-mean, unit-variance Gaussian random variates X_k and Z_k , add them in quadrature, and multiply them by the spectral weights H_k to form $S_k = H_k(X_k + iZ_k)$.
- Apply the complex IFFT to S_k , and obtain the first N points, which we define as sequence s_n .

- Generate two events with N correlated zero-mean Gaussian random variates by separating s_n into its real and imaginary sequences: $Y_n = \Re s_n$, $W_n = \Im s_n$.




Figure 4.10: The gain of the amplifier design deployed in the prototype station is plotted on the y-axis, with 20 dB subtracted. The x-axis are 200 MHz per division, from 0 to 2 GHz.

In the actual Monte Carlo, two events are concatenated to form a full 256-sample *Icicle1* event. It is shown in the above reference that Y_n and W_n are Gaussian distributed, with zero mean. It is also shown that the autocorrelation functions of Y_n and W_n are the desired PSD, ϕ_k . The Wiener-Khinchin theorem states that the autocorrelation function for any wide-sense stationary random process is the PSD. In the next section, properties of noise events created with this procedure are compared to real data taken with the detector.

Before quantitatively comparing Monte Carlo and noise data, two important features of thermal fluctuations are worth noting, derived by S.O. Rice. These properties will not only be helpful here for estimation of event rates, but useful for simulation of the new ARIANNA

time-domain trigger, an analogue transient waveform digitizer (ATWD), relevant to Icicle3, Icicle4 and Icicle6, deployed in 2012. Consider a randomly distributed noise signal $I(t)$, with an autocorrelation function ψ_τ

$$\psi_\tau = \int_0^\infty I(t)I(t + \tau)dt \quad (4.14)$$

The probability that $I(t)$ lies between I and $I + dI$ is distributed like

$$dP = \frac{dI}{\sqrt{2\pi\psi_0}} \exp(-I^2/2\psi_0) \quad (4.15)$$

Equivalently, we identify ψ_0 with the rms deviation from zero. The first important result is that the number of zero crossings per second is calculable:

$$R = \frac{1}{\pi} \left(-\frac{\psi_0''}{\psi_0'} \right)^{1/2} \quad (4.16)$$

The primes indicate differentiation with respect to the lag, τ . This result simplifies if the power per unit frequency is *uniform*, and is limited to frequencies between f_a and f_b :

$$R = 2 \left(\frac{1}{3} \frac{f_b^3 - f_a^3}{f_b - f_a} \right)^{1/2} \quad (4.17)$$

Recall that a pulse maximum stands in for the integral of the power over a short time. The second result from S.O. Rice is that the rate of maxima per second above the value $I(t) = I_1$ is proportional to the rate of zero crossings, and in the bandwidth-limited case the expression simplifies to

$$R = \exp(-I_1^2/2\psi_0) \left(\frac{1}{3} \frac{f_b^3 - f_a^3}{f_b - f_a} \right)^{1/2} \quad (4.18)$$

Equation 4.18 allows us to generate expected single-channel event rates if we model the single channel rate as a thermal fluctuation crossing an effective threshold I_1 . Given the rate at which a single channel experiences a signal above threshold, the global thermal trigger rate can be predicted. The majority logic trigger of the prototype station requires two out of four channels to be triggered within 8.3 ns of each other (the gate $\tau = 8.3$ ns). The formula for the rate of thermal triggers firing on at least $k = 2$ channels out of N channels is [36]

$$R_{tot} = \sum_{k=2}^N \binom{N}{k} R^k \tau^{k-1} = \sum_{k=2}^N \binom{N}{k} \exp(-kI_1^2/2\psi_0) \left(\frac{1}{3} \frac{f_b^3 - f_a^3}{f_b - f_a} \right)^{k/2} \tau^{k-1} \quad (4.19)$$

For the case $N = 4$ this simplifies to ($\alpha \equiv R\tau$)

$$R_{tot} = R(6\alpha + 4\alpha^2 + \alpha^3) \approx 6R\alpha \quad (4.20)$$

Table 4.2 demonstrates predictions of the total rate given the threshold, $z = I_1/\sqrt{\psi_0}$. A threshold of $z = 5$ produces a thermal trigger rate of one event every several seconds. A

typical ARIANNA event with a primary energy of 10^{17} eV will have a signal to noise ratio of ≈ 5 . That result implies that if the threshold is to be comparable to the expected signal to noise ratio, the thermal rate leads to millions of thermal triggers per season to analyze. The next generation trigger for ARIANNA, which will be capable of identifying patterns in the time domain, should allow us to reject patternless thermal fluctuations while keeping the threshold the same (or even lower - see [81]).

Threshold, z	R	R_{tot}
3	6.76 MHz	2.74 MHz
4	0.204 MHz	2.5 kHz
5	2.27 kHz	0.31 Hz
6	9.26 Hz	5.1 μ Hz

Table 4.2: The global thermal trigger rate varies widely with z . The numbers above are for a majority logic of 2 out of 4 channels. The answer for a 0.1 Hz event rate is $z = 5.11$.

The observed thermal trigger rate tends to agree with the predictions of table 4.2. Specifically, in 2010-11 we employed both frequency bands of the prototype trigger, with roughly equal thresholds (table 4.21). From figure 4.26 below, the event rate for these thermal triggers was 0.0604 ± 0.0003 Hz. The average amplitudes $A = v/pp/2v_{rms}$ are given in table 4.3. Note that when the 2010-11 thresholds act on all frequencies uniformly, the station records amplitudes at threshold of ≈ 5 , and the trigger rate is close to 0.1 Hz. The smaller observed amplitude in the West channel indicates atypical operation, and is discussed in detail below. However, the fact that the amplitude doesn't seem to fluctuate as far above background as the other channels can be used to explain the discrepancy between the predicted rate of 0.1 Hz and the measured rate of 0.06 Hz. Removing one channel entirely from the majority logic, in this case, should drop the rate by a factor of 2. In the 2009-10 season, the low-bandwidth portion of the trigger was not used at all, and in the 2011-12 season, the thresholds (table 4.21) were much lower in the high band than the low band. The response of the LPDAs and amplifiers produce events with smaller values of $A = v_{pp}/2v_{rms}$ in this case, because the hardware emphasizes lower frequencies.

Channel	$A = v_{pp}/2v_{rms}$	RMS
North	4.8	0.9
East	5.5	1
South	5	1
West	3	1

Table 4.3: The average and rms of amplitudes $A = v_{pp}/2v_{rms}$ from 2010-11 thermal triggers. Periodic events such as heartbeats and forced triggers, as well as other self-triggered events have been removed from this data.

4.2.3 Forced trigger comparisons to Monte Carlo

To understand the noise observed in the system, comparisons to Monte Carlo simulations have been made. One important effect is that the non-uniform PSD ϕ_k of the noise causes the cross-correlations between channels to contain periodicities close to the dominant frequencies in the spectra. One statistical variable that is useful for neutrino signals is the sum of the lags from the cross-correlations of adjacent channels. The lag between two signals is defined as the value of τ that maximizes the cross-correlation:

$$\psi_{\tau,max} = \int_{-\infty}^{\infty} f(\tau_{max})g(t + \tau_{max})dt \quad (4.21)$$

The cross correlation of two channels recording the same wide-sense stationary random process yields a randomly-distributed lag between the highest and lowest possible times [64]. However, introducing the power spectra of forced triggers (figure 4.11) as the function ϕ_k in section 4.2.2 above introduces periodicities. The sum of the lags from the cross-correlations of adjacent channels (North-East, East-South, etc.) we define as the *planarity*:

$$P = \tau_{NE} + \tau_{ES} + \tau_{SW} + \tau_{WN} \quad (4.22)$$




Figure 4.11: The normalized PSD of the four Icicle1 prototype channels, forced triggers only. These spectra are not uniform, and we can use this fact to explain odd behavior in the planarity.

A true plane wave, with no signal corruption due to noise, satisfies the identity $P = 0$. To prove this, consider difference in arrival times between two channels. Let the vector \mathbf{B}_{ij} describe the baseline between two adjacent channels i and j , and let \mathbf{v} be the velocity vector of an incoming wave. The difference in arrival times for these two channels is

$$\tau_{ij} = \mathbf{B}_{ij} \cdot \hat{\mathbf{v}} / |\mathbf{v}| \quad (4.23)$$

The planarity P is therefore

$$\sum_{\text{adjacent}} \tau_{ij} \propto \sum_{\text{adjacent}} \mathbf{B}_{ij} \cdot \hat{v} = \hat{v} \cdot \sum_{\text{adjacent}} \mathbf{B}_{ij} = 0 \quad (4.24)$$

The last statement is true for any detector where the signal antennas form a polygon in a plane. Thus, the expectation when comparing thermal triggers to a true neutrino event is that a powerful neutrino event has $P = 0$ always, and thermal/forced trigger planarity values are drawn from a normal distribution, since they are the sum of four (flatly-distributed) random variables. This normal distribution can be simulated, and is fit by a gaussian function in figure 4.12, in the absence of a non-uniform PSD. The real forced trigger data has an interference pattern near $P = 0$ ns, ie it is not a simply normal distribution. However, when the forced trigger power spectrum is introduced into the simulation as described above, the same interference pattern arrises near 0 ns planarity. In spite of the interference pattern, the Gaussian fits to all the distributions lie on top of one another.




Figure 4.12: (Red): The normalized planarity distribution for 10^5 Gaussian white noise events with a perfectly flat PSD between DC and the Nyquist critical frequency. (Blue): Identical conditions, except that the event PSDs equal the forced trigger PSDs from real data.




Figure 4.13: The normalized planarity distribution from 43,000 forced triggers taken in the first half of the 2010-11 season. A similar interference pattern near 0 ns appears in the Monte Carlo above.

Another data feature that the Monte Carlo must model is the amplitude of thermal noise. We define the channel amplitude as one-half of the peak to peak voltage in the event, divided by the rms noise in the event:

$$A = \frac{1}{2} \frac{v_{pp}}{v_{rms}} \tag{4.25}$$

The motivation for this particular definition is that it is less risky than merely computing the ratio of the maximum voltage in the event and the rms noise. Fluctuations and signals that excite the LPDA antennas and are amplified by the LNAs cause ringing in the system, making them bi-modal instead of just fluctuating positively or negatively. Single direction fluctuations are uncharacteristic of neutrino signals and should be ignored by the analysis. Thus, the v_{pp} algorithm finds the fluctuation with the largest bi-modal transition around a zero-crossing (figure 4.14). The fluctuation can be either positive or negative in slope, and we can use the same function to get the time of the fluctuation. The definition of A in equation 4.25 will be used extensively in later sections.

The root-mean-square (rms) voltage of an event is calculated from the first 100 samples of the event. This requirement ensures that the rms value is not biased by pulses that form the trigger (the pulses have $T_{pp} > 60$ ns). Equation 4.26 is used to compute the rms voltage of an event, where i refers to the sample, and $\langle v \rangle$ is the mean sample value. Even though digitizer "pedestals" were measured before deployment and are subtracted during the waveform calibration process, it is not given than $\langle v \rangle = 0$ mV, so we must compute it




Figure 4.14: This is an example of how the v_{pp} algorithm locates the largest negative/positive or positive/negative transition around a zero crossing. Notice that the algorithm finds the correct answers even in the presence of a DC offset.

and remove one degree of freedom.

$$v_{rms} = \frac{1}{N-1} \sum_{i=0}^{N=100} (v_i - \langle v \rangle)^2 \quad (4.26)$$

Equation 4.25 does not refer specifically to the Gaussian significance above all backgrounds at the final stage of an analysis, but really the signal to noise ratio of a particular channel in a particular event. Equation 4.25 should just be treated like a unit-less amplitude of a potentially good event. Because the trigger in Icicle1 takes ≈ 65 ns to form in the FPGA logic (see below), the first 2/3 of each channel in each event is typically background noise. Therefore v_{rms} in equation 4.25 is taken strictly from the first 1/3 of a channel in an event to guarantee that no part of the signal leaks in and biases v_{amp} lower.

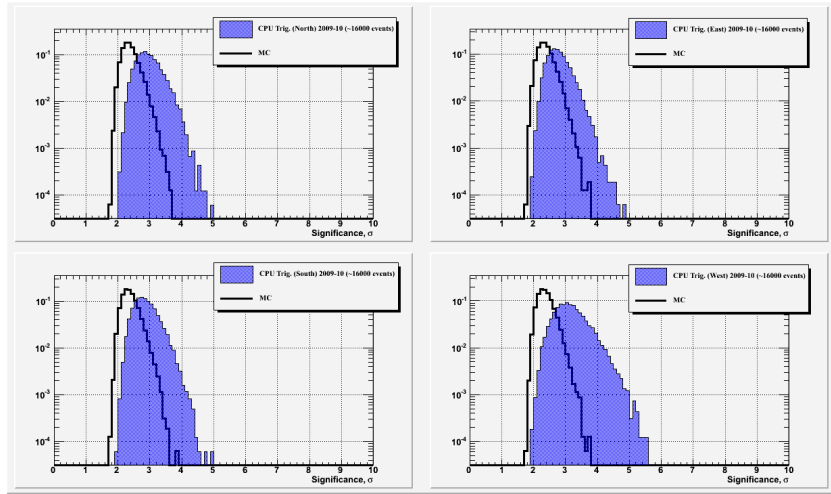


Figure 4.15: CPU triggers (forced triggers) have a characteristic amplitude distribution, peaked at $v_{amp} = \sigma \approx 3$. Monte Carlo forced trigger events, produced assuming they are made of Gaussian white noise, have a distribution peaked closer to 2.

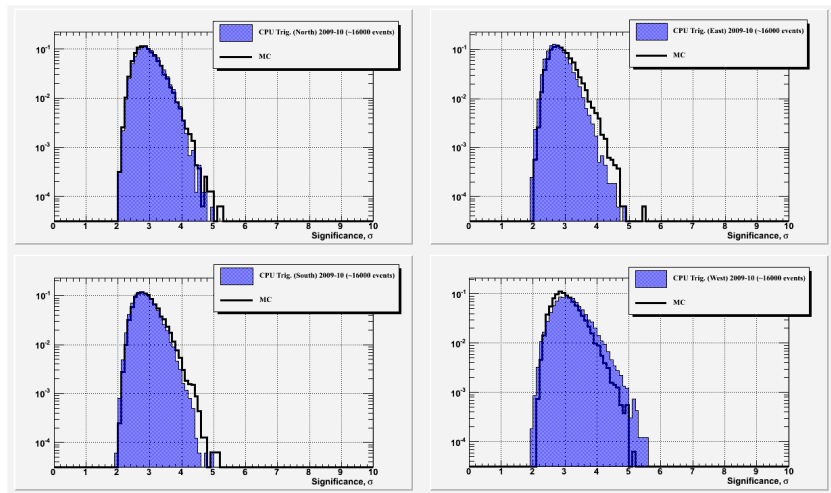


Figure 4.16: After accounting for the non-uniformity of the forced trigger PSDs, the distributions for Monte Carlo events agree with the data.

As with the planarity variable, the amplitude variable of forced triggers can be simulated through Monte Carlo techniques, as long as the non-uniform nature of the noise is taken into account. Figures 4.15 and 4.16 demonstrate that without accounting for the non-uniform PSD of forced triggers, the amplitude (significance, σ) distributions of Monte Carlo simulations do not agree with the data. Apparently favoring frequencies between 200-500 MHz (the dominant frequencies in figure 4.11) shifts upwards the amplitude distributions. The events in these graphs come from the period in the 2009-10 season after the wireless transmitting system was removed, freeing the local environment from man-made interference.

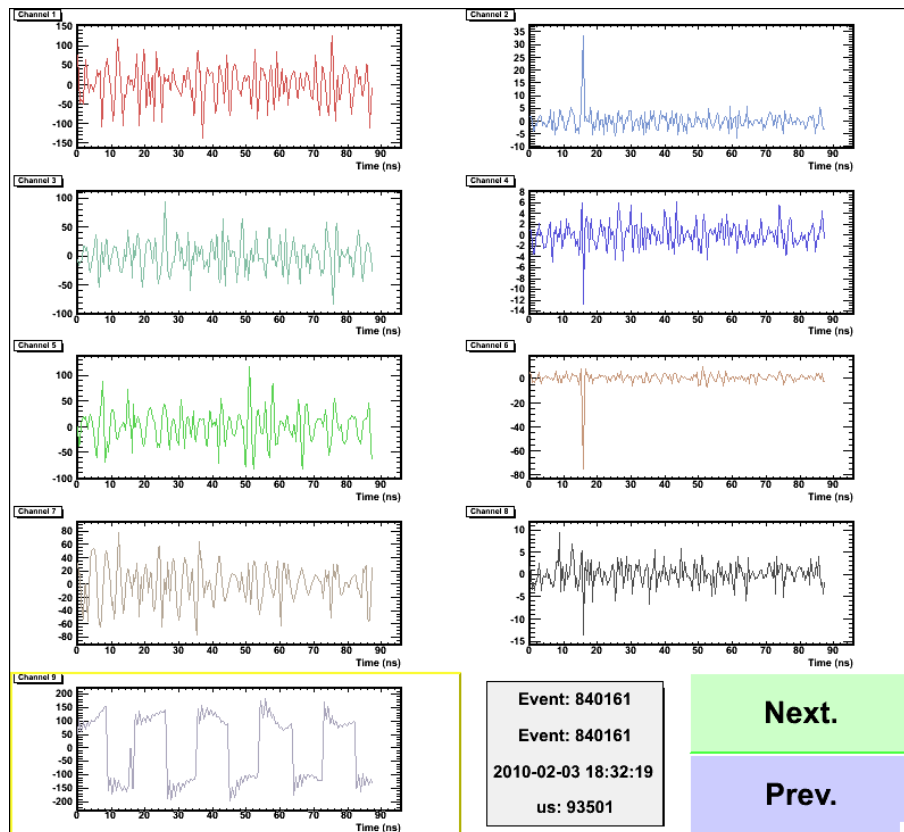


Figure 4.17: An example of a forced trigger. The North, East, South, and West channels are displayed from top to bottom on the left.

The distributions in figure 4.16 peak around 3, and we can understand this by drawing $N_{samples} = 248$ times from a standard normal distribution. (In the ARIANNA digitizer, up to 8 samples of the original 256 are removed to account for the wrap-around time in the




Figure 4.18: In the 2010-11 season, we began thermal trigger acquisition on January 18th. The forced trigger microsecond times exhibit an interesting periodicity/drift after this time, but all microsecond times are still less than 120,000 μsec .

circular buffer). After 248 trials, a normal distribution produces an event with $\sigma = v_{amp} \approx 3$ with a probability of $\approx \text{erfc}(3/\sqrt{2}) = 0.66$. It turns out that, with the slight point-to-point correlation introduced by favoring certain frequencies in the PSD, this probability increases. An example of a forced trigger is shown through the original eventPlotter GUI in figure 4.17. Note the lack of any significant fluctuations or correlated signals. The events are merely a sample of the thermal background over which thermal triggers rise above threshold. There are also four unconnected channels - a design feature we inherit from the ANITA experiment, which had eight channels per board.

A final property of forced triggers is related to the time-stamp of the event. Because iceSoft uses the unix clock to determine when a forced trigger event should be recorded, and because iceSoft is required to create a forced trigger period that is an integer number of seconds long, the forced triggers have time-stamps within ≈ 0.12 seconds of a whole second (figure 4.18, right). For example, if a forced trigger occurs at 16hr 42min 3sec, and the forced trigger period is 30 seconds, the next forced trigger occurs at 16hr 42m 33.075 \pm 0.075 sec. Although the absolute time on the unix clock in the CPU drifts and must be corrected occasionally by GPSd, the precision does not. Figure 4.18 demonstrates that, despite small issues of jitter

in the presence of thermal triggers, forced triggers can be identified by a low microsecond time. The reason the forced trigger microsecond times are not smaller is that Acqd takes a finite amount of time to execute the forced trigger command.

4.2.4 Analysis of heartbeat triggers

In addition to creating forced triggers, the proto-station also created the heartbeat triggers with the Avtech pulser. Heartbeat triggers are separated by one second in time, and there are either 40 or 100 events at the beginning of a run, depending on the season. Heartbeats from the first season were useful for station monitoring, and during subsequent seasons 100 events per run were taken to obtain good statistics for heartbeat calculations. During the first season, the heartbeat antenna was oriented in configuration 1, (spines aligned), and in the second and third seasons we adjusted it to configuration 2 (spines orthogonal). This orientation experiences less attenuation, and sharper, shorter pulses. The heartbeat transmitter was moved farther away to manage the detected pulse amplitudes.

Figure 4.19 below demonstrates how heartbeat events are identified in the data. The x-axis is the UTC time of events in days since January 1, 2010. The y-axis is the time difference Δt between consecutive events. The data comes from late in the 2009-10 season, and on day 33.5 we removed the wireless internet system that caused heavy background noise. This noise dominates all data before day 33.5, and has many events near $\Delta t = 6$ sec. During that season, the length of a data run was six hours, so there are 40 heartbeat events every six hours, with a Δt of 1 second. The other significant bunches of events with $\Delta t < 1$ second are associated with Iridium transmissions. Whenever a data run is successfully triggering on heartbeat pulses, we know that the trigger and data acquisition system are operating normally. There were **no data runs** throughout the three years of operations that failed to record heartbeats. In 2010-11 and 2011-12, we occasionally missed several heartbeats in a

heartbeat string of 100 triggers because the heartbeat antenna was farther away. This occurs $\approx 7\%$ of the time, and in such cases there are Δt (consecutive events) spikes at 2, 3, ... ns.




Figure 4.19: The time difference between consecutive events versus UTC time in days since January 1 2010. The heartbeat events stand out in both the loud noisy period and the quiet thermal background period. The heartbeats have $\Delta t = 1$ second, and occur every six hours (four runs per day).

Heartbeat events are useful in other ways. The positions of the heartbeat transmitter and detector LPDAs are known, making the timing offsets of each data channel measurable. The planarity variable should give $P = 0$ (ns) for a noiseless plane wave, but if the channel offsets do not sum to zero, then we really have $P = \epsilon$ (ns), where ϵ is some number of nanoseconds much smaller than the length of an event, but different enough from zero to confuse a neutrino and a thermal trigger. The measured and predicted time delays between channels, along with the channel offsets, are related by the following two equations:

$$\Delta t_{exp}^{ij} - \Delta t_{meas}^{ij} = \Delta t_{off}^{ij} \quad (4.27)$$

$$\Delta t_{off}^{ij} = t_{i,off} - t_{j,off} \quad (4.28)$$

The subscripts i and j represent the channels. The channel offset is $t_{i,off} = t_{i,exp} - t_{i,meas}$, where t_i is defined to be the time *after* some reference channel is hit by the calibration pulse. To calculate the $t_{i,off}$ explicitly, we must choose one of the four channels as reference (table 4.4). Because the timing information of each event will be combined into the planarity, it is necessary to know the value of ϵ . We have

$$\sum_{adjacent} \Delta t_{exp}^{ij} - \Delta t_{meas}^{ij} = \sum_{adjacent} \Delta t_{exp}^{ij} - \sum_{adjacent} \Delta t_{meas}^{ij} = \epsilon \quad (4.29)$$

Table 4.4 contains the measurements from the 2010-11 heartbeat events that we place in each sum on the right hand side. The 2010-11 heartbeats were chosen over those from 2009-10 because the West channel waveform is almost missing from the 2009-10 data, due to the global T_{pp} delay and the relatively close heartbeat antenna location that season (recall figure 3.33). The measured time differences were taken from the time of the peak to peak fluctuation given by the v_{pp} algorithm described above. The analysis includes 7013 events, located in the data by using the fact that the time stamps are one second apart, and that the events occur in the first ≈ 100 seconds in the data runs. Evaluating the sums on the right hand side of equation 4.29, it turns out that each sum is separately equal to zero within statistical errors. This implies that

$$\epsilon = \sum_{adjacent} \Delta t_{off}^{ij} = 0 \pm 0.3 [ns] \quad (4.30)$$

The error on ϵ comes from adding the errors in the second column of table 4.4 in quadrature. Although there are technical errors associated with determining the antenna positions, which propagate to the numbers in the third column of table 4.4, these errors are negligible

Channels	Δt_{ij} Measured (ns)	Δt_{ij} (ns) Predicted	$t_{i,off} - t_{j,off}$
NE	5.7 ± 0.2	7.15	1.5 ± 0.2
ES	6.7 ± 0.2	8.80	2.1 ± 0.2
SW	11.2 ± 0.1	15.17	4.0 ± 0.1
WN	-23.6 ± 0.1	-31.11	-7.5 ± 0.1
NS	12.5 ± 0.2	15.93	3.4 ± 0.2
WE	-17.9 ± 0.1	-23.96	-6.1 ± 0.1

Table 4.4: The measured and predicted Δt offsets from the 2010-11 heartbeats, along with the differences (forth column). Letting $t_{N,off} \equiv 0$ ns, we get $t_{E,off} = -1.5 \pm 0.2$ ns, $t_{S,off} = 0.7 \pm 0.3$ ns, $t_{W,off} = -3.3 \pm 0.35$ ns.

compared to the errors in the first column (we can measure antenna spine positions to a fraction of an inch). The errors in the second column are the standard error in the mean of the distributions of Δt_{meas}^{ij} , equal to σ/\sqrt{N} , where N is the number of events. What matters in the end is that we expect a plane wave to reconstruct with a planarity consistent with 0 ns.

An example of a (single, noisy) heartbeat event from 2011 is shown in figure 4.20. Being the closest channel physically to the heartbeat antenna, the West channel contains the earliest pulse. Note that the noise in the West channel is gone. This is true for all heartbeats in the 2010-11 season. This could be a symptom of something changing on that channel of the CAGES board, however the instrumentation box of Icicle1 will have to be recovered to investigate further. Assuming that the West channel is along for the ride, the other channels require ≈ 65 ns to form a trigger and readout the data.

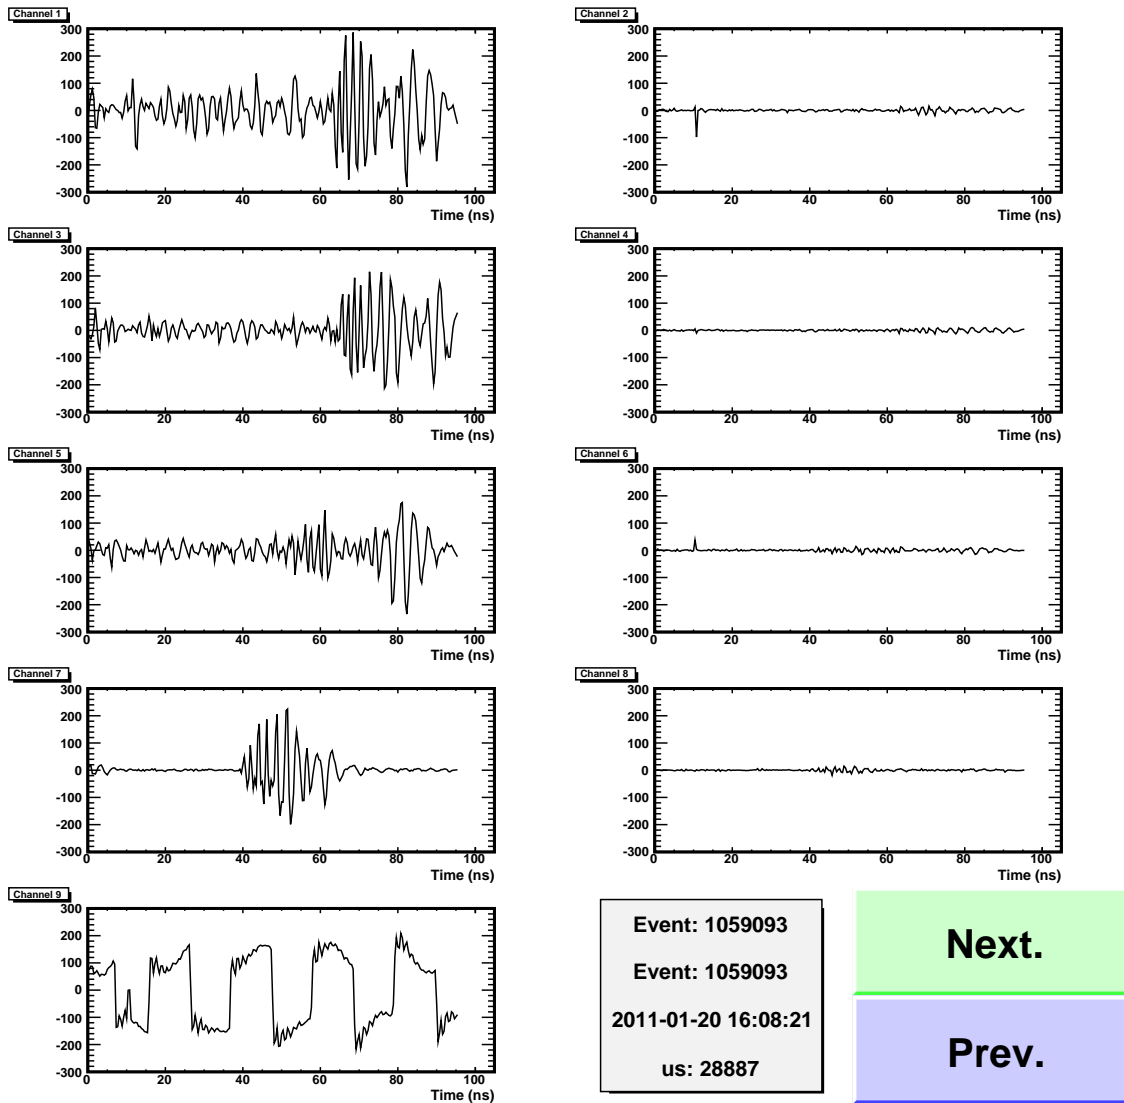


Figure 4.20: An example of a heartbeat event. The four channels on the left side are the N, E, S, and W channels from top to bottom. The non-connected channels (right side) have the same scale as the signal antennas to show the small scale of the cross-talk.

4.3 Time Differences Analysis

This section includes the analysis of the time-stamp data of thermal trigger data taken throughout the three seasons of data runs. Before beginning to analyze the waveforms of the triggered events, thermal triggers must be separated from the other event classes. Section 4.3.1 contains a quantitative discussion of these ideas, and a short analysis of the 2009-10 thermal triggers during the runs free of the interference caused by the wireless ethernet tower. Sections 4.3.2 and 4.3.3 below report similar results as section 4.3.1, but with data from the second and third seasons.

Year/Season	Start day	Stop day	Start run	Stop run	live-time (days)
2009-10	33.55	45.75	152	202	Total: 12.2
2010-11	19.17	39.75	1127	1200	Total: 20.58
2011-12 (1)	-15.0	12.35	57	123	
2011-12 (2)	22.84	23.86	153	155	
2011-12 (3)	23.95	25.66	156	161	
2011-12 (4)	26.10	31.82	162	179	Total:
2011-12 (5)	32.03	58.96	180	258	62.53

Table 4.5: A summary of the periods containing the thermal events analyzed throughout the rest of the chapter. The days quoted are measured from 00:00 hours, January 1st of 2010, 2011 and 2012, respectively, for each season.

The electronics of the w-fi tower, which belonged to and were installed by Raytheon contractors, provided excellent communications but contaminated thermal trigger data with non-random interference. The daily periodicity of this noise (figure 4.19 in the previous section) correlates with the daily battery voltage swings, and is associated with the wi-fi tower only (the noise ceased precisely when the tower was removed). Thus, we concluded that the power/charging system of the wi-fi tower was not RF shielded that season. After providing shielding for this system in subsequent seasons, we no longer observed this dramatic effect. The live-time from the first season runs from the wi-fi tower removal to station shutdown. The second season began with thresholds too high to collect thermal triggers, but low enough to collect heartbeat triggers. After the thresholds were adjusted on day 19.17 (UTC days

since January 1st, 2011, 00:00 hours), the station began to trigger on thermal noise. Data collection continued until station shutdown. The data from the final year was broken into five non-contiguous periods to avoid periods of abnormally high non-thermal triggering.

4.3.1 2009-10 Season: 12.2 days of live-time.

Consider the distribution of the difference in time between consecutive, randomly occurring events. The standard definition of the Poisson distribution describes the number k of events occurring per unit time, with average rate λ :

$$P(k; \lambda t) = \frac{(\lambda t)^k \exp\{-\lambda t\}}{k!} \quad (4.31)$$

The probability that no events occur in a time interval t is distributed like $P(0; \lambda t)$. Thus, if an amount of time t passes before the next event, the probability is

$$P(0; \lambda t) = \exp\{-\lambda t\} \quad (4.32)$$

Equation 4.32 implies that if we measure the time t between events and take the natural logarithm of the distribution, the slope is $-\lambda$, the average event rate. The Icicle1 data begins with a distribution of $t = \Delta t$ that has at least two spikes above the exponential background: the heartbeats and the forced triggers at $\Delta t = 1$ sec and $\Delta t = 67.0$ sec, respectively. The 2009-10 data contained copious amounts of repetitive, anthropogenic noise from the wireless transmitter that provided internet. However, after this tower was removed at day 33.55 of 2010, the event rate became limited by the thermal backgrounds. Figure 4.21 contains the

final Δt results for the first season. The data comes from the period [33.55, 45.75] in 2010, for a run-time of 12.2 days.




Figure 4.21: The Δt distribution of just the thermal triggers (no heartbeats or forced triggers) for the data after day 33.55 of 2010. The data exhibits an exponential form, with similar mean and rms. The inverse of the fitted slope is 133 ± 2 sec/event.

To obtain figure 4.21, the UTC time stamps of the trigger data in the period [33.55,45.75] were examined. Figure 4.22 shows the Δt versus time in UTC days for this period. The straight line of events at $\Delta t = 67$ sec represents the forced triggers, and bursts at 1 second representing the heartbeats. There are occasionally events that have $\Delta t < 0.8$ seconds, corresponding to Iridium transmissions. These three types of events are tagged and cut to produce figure 4.21. Figure 4.23 shows the raw Δt distribution before forced trigger and heartbeat removal.

The table of remaining events after each Δt cut is shown below. The expected number of heartbeats comes from the number of runs times the number of heartbeats per run. The setting for this season was 40 events per run (one run was set to six hours). There is no fixed number of forced triggers. Instead we take the number of runs times the length of a run in seconds and divide by the constant forced trigger period (67 seconds). There will inevitably be thermal triggers that satisfy the cut criteria. The inefficiency introduced by the stage 1




Figure 4.22: The raw data from trigger time stamps after the removal of a noise contamination before day 33.55 in 2010. The forced triggers are represented by the line at $\Delta t = 67$ sec, and the heartbeats occur in bursts near $\Delta t = 1$ sec.

cut is small. Consider the integral of the exponential distribution (as it is defined by ROOT software) between two times t_1 and t_2 .

$$\epsilon = \int_{t_1}^{t_2} dt \exp \{A - \alpha t\} \quad (4.33)$$

First, to normalize the distribution such that the integral between zero and infinity is one, we must have a normalization constant of $N = \alpha/e^A$. After multiplying the integral by N , and do some rearranging, we obtain ($\Delta t = t_2 - t_1$).

$$\epsilon = \int_{t_1}^{t_2} dt N \exp \{A - \alpha t\} = \exp \{-\alpha t_1\} (1 - \exp \{-\alpha \Delta t\}) \quad (4.34)$$




Figure 4.23: (Left): These raw data are the same as figure 4.22, but only the Δt distribution. (Right): Zooming in to see the heartbeat and Iridium events. The excess at 2 seconds corresponds to heartbeats preceded by a missed heartbeat event.

Every Δt cut in these analyses satisfies $\alpha\Delta t < 1$, meaning we can make the approximation

$$\epsilon = \int_{t_1}^{t_2} dt N \exp\{A - \alpha t\} \approx \alpha\Delta t \exp\{-\alpha t_1\} \quad (4.35)$$

The limiting cases of this equation make sense: if the cut involves some time t_1 that is much larger than the average time between events (the inverse of the rate, α) exponentially few random events are removed. If the time t_1 is not large, however, it is the *width* of the cut that matters. Narrow cuts are efficient, and wide cuts are not. Also note that the inefficiency increases for higher rates in this regime: if there are more events at low Δt , the inefficiency will be higher. For example, cutting every event with Δt between 0.0 and 0.8 sec, with a rate of 0.0075 Hz, yields an inefficiency $\epsilon = 0.6\%$. The non-random events in this region are associated with hourly Iridium transmissions, and show up in subsequent seasons with lower thresholds, where the hourly periodicity is more apparent. The inefficiency here is actually over-estimated; we rarely have a Δt value less than 0.2 seconds because of the dead time of the detector.

The heartbeat cut is defined as any event satisfying $\Delta t = 1.0 \pm 0.02$ sec. The σ of the

gaussian fit to the heartbeat peak is 0.005 seconds, making the 0.02 width of the cut $\approx 4\sigma$ wide. The choice for such a wide cut is that we cannot have heartbeats escaping into the waveform analysis, where they pose as large amplitude plane wave neutrino candidates. The forced trigger cut is any event satisfying $\Delta t = 67.0 \pm 0.2$ sec. Similarly, the width of the forced trigger cut is $\approx 3\sigma$, because the sigma of a gaussian fit to the forced trigger peak is 0.06 seconds. It is not risky to be less stringent in removing forced triggers; they are less likely to pose as fake neutrinos at the next stage. The inefficiencies of the heartbeat and forced trigger cuts are 0.3% and 0.2%, respectively. These cuts combine for an overall inefficiency of 1.1%.

Stage	Events Remaining	Change	Expected Change	Cut
Raw	26939	0	-	None
Stage 1	24942	1997	1997	$\Delta t < 0.8$ sec
Stage 2	22785	2157	1920	$\Delta t = 1.0 \pm 0.02$ sec
Stage 3	7525	15260	15465	$\Delta t = 67.0 \pm 0.2$ sec

Table 4.6: The expected change due to the cuts in the last column can be calculated for the forced triggers and heartbeats based on how many the detector generates. The short burst events with $\Delta t < 0.8$ simply have to be counted; there is no prediction for how many there should be.

4.3.2 2010-11 Season: 20.58 days of live-time.

A similar time-difference analysis was performed on the 2010-11 data, with the same goal of obtaining a clean thermal trigger population. The 2010-2011 season presented a new

Period	Events	Begin	End	Runs	Hr/run	HBeats (exp)	Forced (exp)
Early	68415	354.35	384.17	126	6	12.6k	40.6k
Late	233045	384.17	404.75	74	6,12	6k	27k

Table 4.7: The early period corresponds to higher thresholds, and the total event count is mostly forced and heartbeat triggers. The thresholds were adjusted for the late period, giving a count dominated by thermal triggers. The late period is comprised of runs 1127 to 1200.




Figure 4.24: The x-axis is the UTC time in days since Jan 1 2010, and the y-axis is the consecutive-event time difference. Before day 384.2 after Jan 1 2010, the threshold was set too low to obtain a significant thermal trigger sample.

complication. Because we were aware of the noise problem from the wireless tower during the prior season, we ensured that the charge controller/battery system was RF-shielded. The charge controller and battery system, containing an electronic switch for the divert function, is thought to have caused the repetitive noise. The thresholds were set high enough to keep the station from always triggering on these events, thereby filling up the hard disk. Once the wireless tower was shielded in 2010-11, these thresholds proved to be too high for thermal triggers, but low enough for heartbeat triggers. The thresholds were lowered on day 384.2 (after Jan 1 2010), such that we obtained a thermal rate of approximately 0.1 Hz. The length of a run was switched to 12 hours, and a higher rate of Iridium transmissions (once per hour) on day 396.0. This reconfiguration is what is responsible for the gap in figure 4.24 around day 396.0. Table 4.7 summarizes these actions, and figure 4.25 displays the raw time difference data.

After day 404.75, the wireless transmitter was again removed, to be saved from the winter. During the 2011-12 expedition, we transferred North the rest of the 2010-11 data and learned

that the event rate hit the maximum possible rate on day 407.0. While these new events could be a sign of an Iridium malfunction caused by low voltages, it isn't presently understood. These events could have been caused by a source far from the detector (we observed external noise in Moore's bay in December 2011). This noise, however, could also have been caused by DC-AC power inverters, and has not been observed again. A third possibility could be that sun experienced a solar radio burst. This burst occurred on February 13th, 2011, between 1700 and 1900 hours UTC (ref ARA design). It was observed by the Green Bank Solar Radio Burst Spectrometer (GBSRBS), and prototype stations of the Askaryan Radio Array (ARA). Although the solar burst technically occurred in the middle of the period when we experienced the anomaly, it would be more convincing if the two anomalies began at the same time. The waveform analysis on the thermal trigger population for this season begins on day 384.17, and ends on day 404.75.




Figure 4.25: (Left): Focusing from figure 4.24 on the final two weeks of the 2010-11 data taking. The heartbeats are still present every 12 hours. After day 396, the short-burst events appear because we switched on the Iridium modem once per hour, to ensure good communications and data transfer prior to wireless comms deactivation. (Right): On day 407, we observe a dramatic rate increase - hitting the maximum trigger rate of 5 Hz of the prototype station.

For the late period, in order to begin studying the thermal triggers, Iridium noise events must be eliminated. They are assumed to have been caused by the satellite modem because the bursts are separated by one hour, and the typical Δt is less than 0.3 sec - representing a large number of fast triggers. They do not occur throughout the season, however. There is a group lasting for ≈ 3 days from day 385-387 inclusive, with 27,483 events tagged for

$\Delta t < 0.8$ sec. The number 0.8 sec is a precaution to tag events that leak up to $\Delta t \approx 0.8$ sec in rare cases. Any event in this three day period with $\Delta t < 0.8$ sec therefore does not pass. With an average $\Delta t = 0.3$ seconds, and 24 bursts per day, for an average call length of 2.1 minutes, the number of expected events is 30k. The number 2.1 minutes is obtained from prior year statistics obtained from Iridium corporation.

Another train of Iridium pulses lasts 8.0 days starting from from day 396.5, just after the station reconfiguration. The expected number of events is 80k, for a total of 110k from the two satellite modem trains. The total number of events tagged is 107,478, or 98% of the expected number of tagged events. If the average call length is off by 2%, then the prediction can agree with the number of events tagged, and this is entirely possible. The cuts described thus far are known together as stage 1. Naively applying equation 4.35 yields an inefficiency of 4.8%. However, we are not applying this cut everywhere, but applying it to just those days containing the pulse trains (53% of the data). Thus the real inefficiency is 2.5%.




Figure 4.26: The Δt distribution of just the thermal triggers (no heartbeats or forced triggers) for the 2010-11 data after day 384.17. The data exhibits an exponential form, with similar mean and rms. The inverse of the fitted slope is 16.55 ± 0.07 sec/event.

Stage 2 defines a cut that removes a sparsely distributed population of 2043 events with $\Delta t < 0.25$ sec. Although a much tighter cut could be defined to grab these events, it is not necessary since the minimum possible Δt value is 0.2 sec, making the width of this cut 0.05 sec, and the subsequent inefficiency small (0.3%). Stage 3 attempts to isolate the heartbeat triggers using the same cut from the prior season, 1.0 ± 0.02 sec. From the number of runs and heartbeat pulses per run, the expectation is 6k events, and we tag 6184 events. The prototype station ran for 10 days in the late period at 4 runs/day, and 10 days in the late period at 2 runs/day, using 100 events/run (the data shows 60 heartbeat excesses at $\Delta t = 1.0$ sec). Thus, the heartbeat sample should be 97% pure. From similar arguments of run time, the number of expected cpu triggers is 26.6k, and 27198 events were tagged using the cut 67.0 ± 0.2 sec. The forced trigger sample is thus 98% pure (stage 4). Table 4.8 summarizes. Given the event rate (figure 4.26), and fitted exponential function, the overall efficiency estimate is 97.0%, coming from inefficiencies of 2.5%, 0.3%, 0.2%, and 0.04% for stages 1-4 respectively.

Stage	Events remaining	Change	Expected change	Cut
Raw	233,045	N/A	N/A	None
Stage 1	125,567	107,478	110,000	$\Delta t < 0.8$ sec (two regions)
Stage 2	123,524	2,043	N/A	$\Delta t < 0.25$ sec (all days)
Stage 3	117,340	6,184	6,000	$\Delta t = 1.0 \pm 0.02$ sec
Stage 4	90,143	27,197	26,600	$\Delta t = 67.0 \pm 0.2$ sec

Table 4.8: The Δt cut breakdown for the 2010-11 data (late period). Expected reductions are observed with the forced trigger and heartbeat cuts, and inefficiencies are low due to the high event rate of thermal triggers.

4.3.3 2011-12 season: 62.53 days of live-time

In the 2011-12 season the author, along with Dr. Eric Berg and Joulien Tatar, returned to Moore's Bay to install Icicle2 and to restart the Icicle1 prototype. The expedition began in late November, and Icicle1 was given an additional AGM battery, and the ethernet to

coaxial converters were removed in favor of a RF-shielded ethernet bulkhead connector. This dropped the maximum power consumption to 19W (closer to 17W while the Iridium modem is off). Dropping the power consumption of the prototype by half allowed it to could operate in dimmer sunlight nearer to sunset, which extended the live-time. Only the Trimble GPS was used after the first season proved we didn't need active antennas on the GPS receiver. The power savings from that decision apply to both the 2010-11 and 2011-12 seasons. The result is that Icicle1 obtained 62.53 days of live-time with normal thermal trigger rates.

Period	Days since Jan 1 2012 (UTC)
1	-15.0 - 12.35
2	22.84 - 23.86
3	23.95 - 25.66
4	26.10 - 31.82
5	32.03 - 58.96

Table 4.9: The total live-time of the five periods in 2011-12 is 62.53 days. The periods in between begin with sudden transitions to anomalously high event rates.

The data is broken into five periods, with anomalously high rates in between these five periods. The most likely source of the noise in between the five normal periods is the AFAR system between Icicle2 and Mt. Discovery. Icicle2 is located 1 km away from Icicle1, and communicated directly with the repeater on Mt. Discovery (rather than by forming a local network with the protostation). There were several changes to the AFAR hardware mid-season, including the substitution of the Icicle2 unit with an older version. These changes were prompted by communications outages, however it is possible that these outages were actually caused by fog in between the stations and Mt. Discovery (this effect was initially observed at the end of the 2010-11 season as researchers returned to McMurdo base). Due to the high event rate, the noisy periods constitute a large quantity of data which has not yet been analyzed. Table 4.9 delineates the five periods of thermal trigger data, in UTC days since January 1, 2012, and figures 4.27 and 4.29 show the Δt distribution of the five

periods.




Figure 4.27: The Δt vs. time plot for the 2011-12 data. Days 13 to 22 represent an abnormally high event rate. Data in these periods are not included in the present analysis.

Aside from sudden transitions in the event rate, it is clear that the thermal trigger rate is slowly changing over time. Figure 4.28 shows the Δt vs. time plot of just the thermal triggers from period 1. Before and after day 5, there is roughly a factor of ≈ 2 change in the event rate. From equation 4.20, we see that the global event rate, given a 2 of 4 majority logic, is proportional to the single channel event rate squared. Thus, if the single channel event rate drops by a factor of $\sqrt{2}$, we have the observed change in global trigger rate. This amounts to a downward shift in the quantity $I_1/\sqrt{\psi_0}$ (the threshold over the rms noise level) of $\sqrt{2 \ln \sqrt{2}} \approx 0.8$. For example, if we have a threshold $I_1/\sqrt{\psi_0} \approx 5.1$ for a trigger rate of 0.1 Hz, and the rms noise level $\sqrt{\psi_0}$ increased by 20%, we would have an effective threshold of 4.3, and thus a downward shift of 0.8. It seems unlikely that the rms noise would increase by 20%, given that the ambient temperature is decreasing during all the periods. A more likely explanation is that, after three years of deployment, some voltage on the CAGES board is not being held constant.




Figure 4.28: The Δt vs. time plot for period 1 of the 2011-12 data, showing a gradual shift in the event rate. After day 5, the event rate has slowly dropped from 13.5 seconds per event to 5.5 seconds per event.

To calculate the thermal trigger event rates for the five periods, forced triggers, heartbeats, and man-made periodic excesses are removed in the usual way before fitting the slope of the resulting exponential distributions. Figure 4.29 shows the results of the fits for the five periods, and table 4.10 contains the average time between events (inverse of the average event rate) for each period. Based on the fits in table 4.10, and figure 4.29, one could draw the conclusion that the noise in between periods 1 and 2 changed the properties of the detector, producing the subsequent four periods with high rates. The prototype station was not optimized to read out data at the maximum/high rates for extended periods, however, variables such as the rms and peak to peak voltages remain more or less stable throughout the entire season (see next section).

As with the prior seasons, we must detail how forced triggers, heartbeats, and self-triggered noise sources are eliminated, and quantify any inefficiencies. Table 4.11 shows the events remaining, change in events, and expected change in events for the different periods and cut stages in the final year Δt analysis. Because so much information is being displayed, an




Figure 4.29: The Δt distributions (in seconds) for periods 1-5 of the 2011-12 season. The lower four graphs are well described by a single exponential fit between 1 and 20 seconds, whereas the upper graph deviates from a single fit, caused by the slowly changing event rate during that period.

extra column has been included that contains the formula for the expected change in events due to the forced trigger and heartbeat trigger cuts.

There is a challenge presented by the higher event rate of the 2011-12 data. Knowing equation 4.35 describes the inefficiency of a Δt cut, we see that a higher event rate drives up inefficiencies unless we compensate by narrowing the width of the cut. Whenever possible, this has been done with the four stages of cuts applied to the five periods of the data, while tagging at least the expected number of heartbeats and forced triggers. Period 4, stage 2 is the only exception to this rule, and 400 forced trigger events have no chance of posing as UHE neutrinos. With the narrower widths of the 2011-12 cuts, equation 4.35 is an even

Period	Average time between events (sec)
1	6.83 ± 0.02
2	2.53 ± 0.02
3	2.32 ± 0.01
4	2.07 ± 0.01
5	1.91 ± 0.002

Table 4.10: The period 1 number is essentially a weighted average of the slowly varying event rate during this time. The event rates in the other four periods are relatively stable.

better approximation, and we can estimate the inefficiencies.

The main results are as follows. The overall analysis efficiency from the cuts described in table 4.11 for period 1 is 98.4%, 96.1% for period 2, 96.1% for period 3, 95.3% for period 4, and 95.1% for period 5. Table 4.12 summarizes how these answers are obtained. The forced trigger cuts make little difference because they live so far from the mean of the exponential distributions. For the stage 3 cuts tasked with removing noise from Iridium transmissions, we have made use of the fact that the dead-time of the detector is 0.2 seconds, making the width of the cut < 0.3 second, 0.1 second. The average inefficiency of the five periods, based on the fractional live-time of each period, is 3.5%, making the total efficiency 96.5%.

Period	Stage (sec)	Remain.	Expect.	Obs. change	Formula
1	0 (Raw)	401333	N/A	N/A	N/A
1	1 (1.0 ± 0.015)	390673	5400	10660	$100 \cdot N_{runs}$
1	2 (67.0 ± 0.15)	354600	35000	36073	$N_{sec}/67.0$
1	3 (< 0.3)	333942	N/A	20658	N/A
2	0 (Raw)	27164	N/A	N/A	N/A
2	1 (1.0 ± 0.01)	26615	200	549	$100 \cdot N_{runs}$
2	2 (67.0 ± 0.05)	24159	1300	2456	$N_{sec}/67.0$
2	3 (< 0.3)	23234	N/A	925	N/A
3	0 (Raw)	68748	N/A	N/A	N/A
3	1 (1.0 ± 0.01)	67204	500	1544	$100 \cdot N_{runs}$
3	2 (67.0 ± 0.02)	63929	2100	3275	$N_{sec}/67.0$
3	3 (< 0.3)	57000	N/A	6929	N/A
4	0 (Raw)	259630	N/A	N/A	N/A
4	1 (1.0 ± 0.01)	252908	1200	6722	$100 \cdot N_{runs}$
4	2 (67.0 ± 0.01)	245901	7400	7007	$N_{sec}/67.0$
4	3 (< 0.3)	213939	N/A	31962	N/A
5	0 (Raw)	1385233	N/A	N/A	N/A
5	1 (1.0 ± 0.005)	1359900	5400	25333	$100 \cdot N_{runs}$
5	2 (67.0 ± 0.015)	1303141	35000	56759	$N_{sec}/67.0$
5	3 (< 0.3)	1093482	N/A	209659	N/A

Table 4.11: The Δt cuts for all periods of the 2011-12 data. Because the event rate steadily increases throughout the season, Δt cuts have to be narrowed to maintain efficiency. The heartbeat tagging always removes more events than expected - a safeguard against stray heartbeats posing as UHE neutrinos signals.

Period	Stage (sec)	ϵ	Period	Stage (sec)	ϵ
1	1 (1.0 ± 0.015)	0.19%	4	1 (1.0 ± 0.01)	0.3%
1	2 (67.0 ± 0.15)	$1.2 \cdot 10^{-4}\%$	4	2 (67.0 ± 0.01)	0.0%
1	3 (< 0.3)	1.4%	4	3 (< 0.3)	4.4%
2	1 (1.0 ± 0.01)	0.27%	5	1 (1.0 ± 0.005)	0.16%
2	2 (67.0 ± 0.05)	0.0%	5	2 (67.0 ± 0.015)	0.0%
2	3 (< 0.3)	3.7%	5	3 (< 0.3)	4.7%
3	1 (1.0 ± 0.01)	0.28%	-	-	-
3	2 (67.0 ± 0.02)	0.0%	-	-	-
3	3 (< 0.3)	3.6%	-	-	-

Table 4.12: Individual inefficiencies for the Δt cuts for the 2011-12 data. The forced trigger inefficiencies become negligible because the time between events is much smaller than the forced trigger period. Conversely, eliminating short burst events (stage 3) gets worse as the event rate increases with period.

4.4 Description of Waveform variables, comparisons to Monte Carlo simulation

Given that the non-anthropogenic portion of the data from the three seasons exhibit Poissonian time differences, we can conclude that they are generated by thermal noise backgrounds randomly triggering the station. An even more convincing argument is to show that the voltage properties of these events match expectations based on Monte Carlo simulations of thermal noise, and that the trigger conditions are being obeyed. The majority logic of Icicle1 was kept at 2 of 4 channels triggering within 10 ns for the duration of deployment. Section 4.4 discusses the various thermal trigger waveform properties analyzed, and compares them to Monte Carlo simulations. Section 4.5 demonstrates the application of waveform cuts that eliminate thermal triggers in favor of large amplitude plane waves. From the three seasons of data, there are no thermal triggers posing as fake UHE neutrino signals. An estimation of the final analysis efficiency is also computed in this section. After this result, we discuss in section 4.6 is a special section dedicated to demonstrating that the trigger produces expected effects on the thermal trigger waveforms. Figure 4.30 shows an example of a thermal trigger.

We have already defined one waveform-based variable: the *planarity*. The planarity defines the signal region where potential neutrino candidates live (equation 4.22). All events with $|P| < 300$ ps are highly interesting, because they satisfy equation 4.24. The calculation of the lags τ_{ij} between adjacent channels is performed in the Fourier domain and follows from the correlation theorem [69]. Let g and h be two real functions of time, and let G and H be the Fourier transforms of each. The correlation of g and h is defined as

$$Corr(g, h)_\tau = \int_{-\infty}^{\infty} g(t + \tau)h(t)dt \quad (4.36)$$

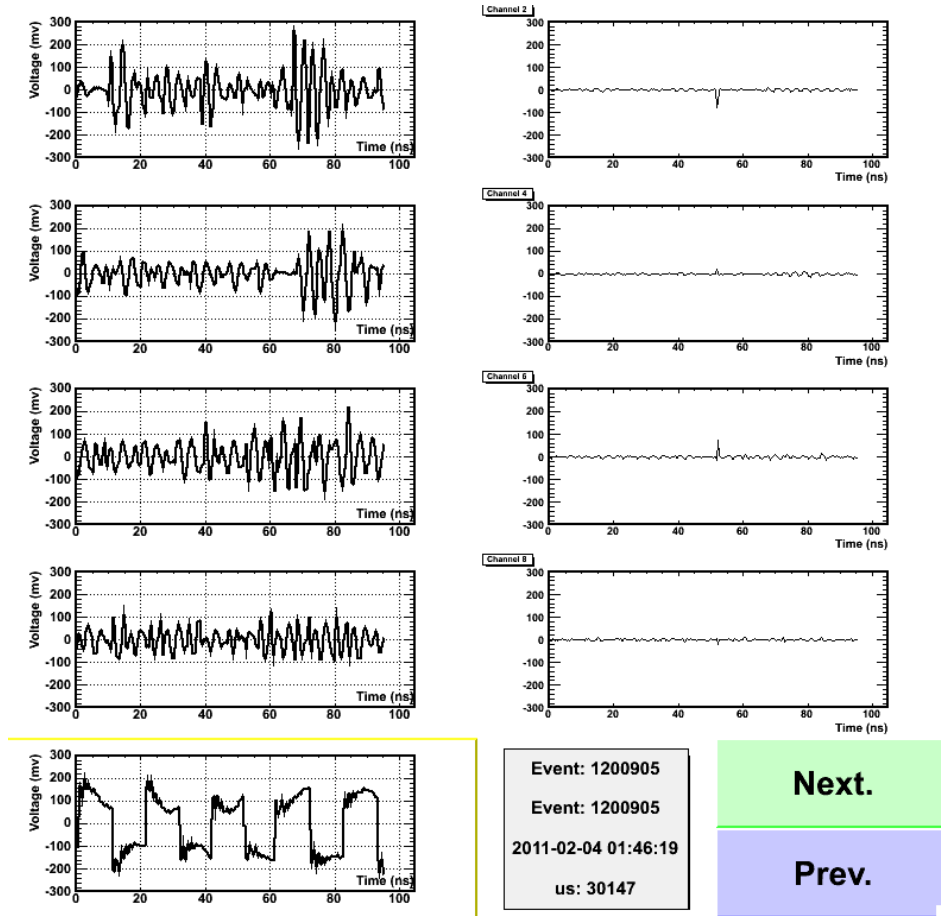


Figure 4.30: This is an example of a thermal trigger from the 2010-11 season, with the West channel still active. The North and East channel have clear fluctuations above backgrounds (top left, second to top, left), and the South and West channels show amplifier noise (second to bottom, left, and bottom left).

When this function of τ is at a maximum, the variable τ is the best choice for the lag between the two signals. The correlation theorem says that the correlation function forms a Fourier transform pair with the multiplication of G and H , provided we take the complex conjugate of H :

$$\text{Corr}(g, h)_\tau \Leftrightarrow G(f)H^*(f) \tag{4.37}$$

In fact, the choice of which function is conjugated determines the interpretation of the best lag. Suppose we take the conjugate of H (and obtain $corr(g, h)$ via an inverse transform). If the best lag is negative, this implies g is earlier than h , and vice versa. Conjugating G in the Fourier domain reverses the interpretation of the sign of the best lag.

Besides the planarity, other timing cuts are used to distinguish thermal backgrounds from signal. The **causality cut** is defined by the geometry of the detector and the wave speed in the upper firn of Moore’s Bay. The maximum possible time lag between signals in two channels arising from a common plane wave is $\pm nB/c$, where B is the baseline distance between the signal antennas. The causality cut requires that all lags between adjacent channels satisfy $|\tau_{ij}| < nB/c$. Knowing the index of refraction $n = 1.3$, and the positions of the antennas, we arrive at the criteria in table 4.13 for the causality cut. These criteria are subjected to a logical AND to satisfy the causality cut. The channel pair S,N was not originally part of the planarity variable, but after the first season the W channel began to show abnormal behavior. One can still compute the planarity variable with three channels provided τ_{SN} is calculated. A plane wave must obey causality for any pair of channels, so we include this extra channel pair in the cut. The E,W pair could have been included as well, but lag computation requires FFT algorithms which are expensive to evaluate for millions of events. Given that the S,N pair might have become necessary, it was included in the analysis.

Channels	$ \tau_{max} (ns)$
N,E	28.7
E,S	10.7
S,W	29.8
W,N	31.6
S,N	36.2

Table 4.13: The above criteria define the maximum lag between pairs of channels given the speed of light in the upper firn. The overall causality cut requires that ALL of these criteria be satisfied for a passing event.

Section 4.2.3 compares our thermal trigger Monte Carlo to forced triggers, and accounts

for the non-uniform nature of the the Icicle1 prototype bandwidth. The result is that we can explain an anomaly in the planarity variable by invoking the effect of the non-uniform bandwidth. Here, we extend this type of analysis to the other variables used to distinguish thermal triggers from neutrino signals. Table 4.14 compares the survival rates of Monte Carlo and real thermal triggers for different combinations of channels, and globally. The Monte Carlo thermal trigger population has 10,000 events, employs the non-uniform PSD measurements from forced triggers, and each generated event is subjected to a 2 of 4 majority logic with a 8.3 ns gate.

The calculated expectation for the survival probability of a given pair of channels assumes the $|\tau_{ij}|$ are flatly distributed on $[0, \frac{1}{2}t_{max}]$. The time t_{max} is just the length of the data record ($248 \text{ samples} \times 0.3508 \text{ ns/sample}$), 87.0 ns. With the exception of the channel combination involving the fluctuating signal, all channel pairs should have lags drawn randomly from a flat distribution, because either a signal is being correlated with noise, or noise with noise. This is why the pairwise survival probabilities match closely the expectations in the third column of table 4.14.

Channels	Surv. Rate (data/MC)	Expect. (null)	Cut (ns)
NE	67.7% / 68.9%	66.0%	$ \tau_{NE} < 28.7$
ES	37.7% / 28.4%	24.6%	$ \tau_{ES} < 10.7$
SW	71.9% / 68.5%	68.5%	$ \tau_{SW} < 29.8$
WN	75.7% / 74.4%	72.6%	$ \tau_{WN} < 31.6$
SN	85.3% / 84.3%	83.2%	$ \tau_{WN} < 36.2$
Global	13.1% / 9.0%	13.0%	AND all

Table 4.14: A comparison of the causality cut between 2009-10 thermal triggers and MC. The expected survival probabilities are based on a flat distribution of lags (non-participation in the trigger). The expectation for the global trigger is derived from a weighted average (see text).

For the global survival probability, we must multiply the pairwise survival probabilities of the pairs not constituting the trigger. The expectation for the survival rate of the global cut is a weighted average, where the weights are the probability that a given channel pair

creates the trigger, and the quantity averaged is the multiplication of the other pairs' survival probabilities. To put this in mathematical notation, let p_i be the survival probabilities of each channel pair, and let w_i be the weight, or the probability that pair j constitutes the trigger. We have

$$P_{exp} = \sum_i w_i \left(\prod_{j \neq i} p_j \right) \quad (4.38)$$

To determine from the data how often a given channel participates, we calculate the fraction of events with T_{pp} greater than or equal 60.0 ns. The West channel threshold, however, was set as high as possible (a relic of the wireless noise we experienced prior to the quiet period). For the causality calculations here, we ignore this effect, and we will come back to it in the discussion of the waveform amplitudes. The weights w_i are the multiplication of the two individual channel probabilities. The p_i from the data are in the second column of table 4.15. The w_i are listed below. As a check, note that they sum to $0.981 \approx 1$, as they should. The thermal trigger Monte Carlo global survival probability is only low because it takes an inconvenient amount of time to set the Monte Carlo thresholds high enough; we observe that lowering the Monte Carlo thresholds lowers the global causality cut survival rate (see below). Subsequent seasons survive the global causality cut with similar probability (in 2011-12, the rate is 10.7%).

Another cut included in our waveform analysis accounts for the **trigger delay** of the CAGES board. From now on, this cut value will be denoted T_{pp} (time of the peak to peak fluctuation). We use time of the peak to peak fluctuation in a data channel to measure T_{pp} . This is the same algorithm we use to measure the peak to peak voltage, except that the algorithm outputs the time rather than the voltage (see figure 4.14). Heartbeat events indicate that the trigger forms and stops data acquisition inside the SCA of the LABRADOR chip 65

Channels	w_i
NE	0.0758
ES	0.0782
SW	0.255
WN	0.247
NS	0.0786
EW	0.246

Table 4.15: The weights w_i used in the assessment of the causality cut accuracy above. Note that the West channel participates 3/4 of the time. Although we can't know for certain, this is probably an effect of setting the highest possible threshold in the West channel for this season.

ns after the trigger conditions are satisfied. If this delay was zero, pulses from the Avtech would arrive at the beginning of the data record. Figure 4.31 shows an example of a 2009-10 heartbeat event with the first heartbeat transmitter location, and figure 4.32 shows a heartbeat trigger typical of 2010-11 and 2011-12 with the second transmitter location (figure 3.33). Conservatively, the lowest number T_{pp} can be (for a participating channel) is the arrival time in the North and East channels, or about 60 ns.

In the first season, the heartbeat transmitter was close to the array and equidistant from the North and East channels. The typical T_{pp} is large enough, however, to hide the majority of the West channel pulse. In subsequent seasons, the West channel began to demonstrate faulty behavior. However, it is still capable of digitizing heartbeats, even though it is not in the trigger. Thus the West channel's timing properties (T_{pp} and cross-correlation techniques) are still useful.

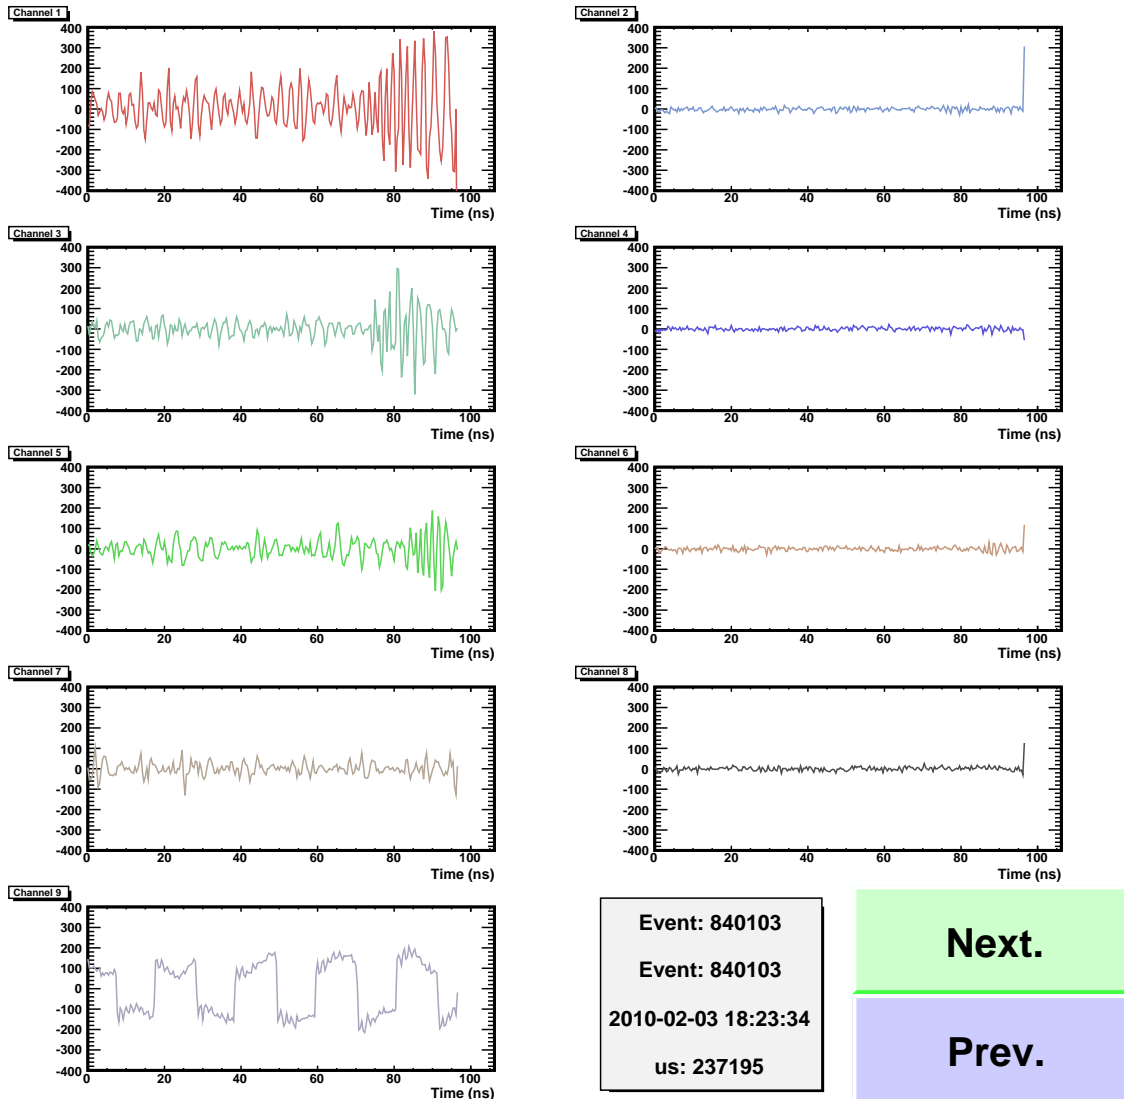


Figure 4.31: This is an example of a heartbeat event from the 2009-10 season. These events are identified by $\Delta t = 1.0 \pm 0.015$ sec, and the relative timing of the pulses in the channels. The West channel antenna buried the farthest from the heartbeat LPDA antenna. The trigger delay conspired to hide the Western pulse until we moved the heartbeat transmitter in December 2010.

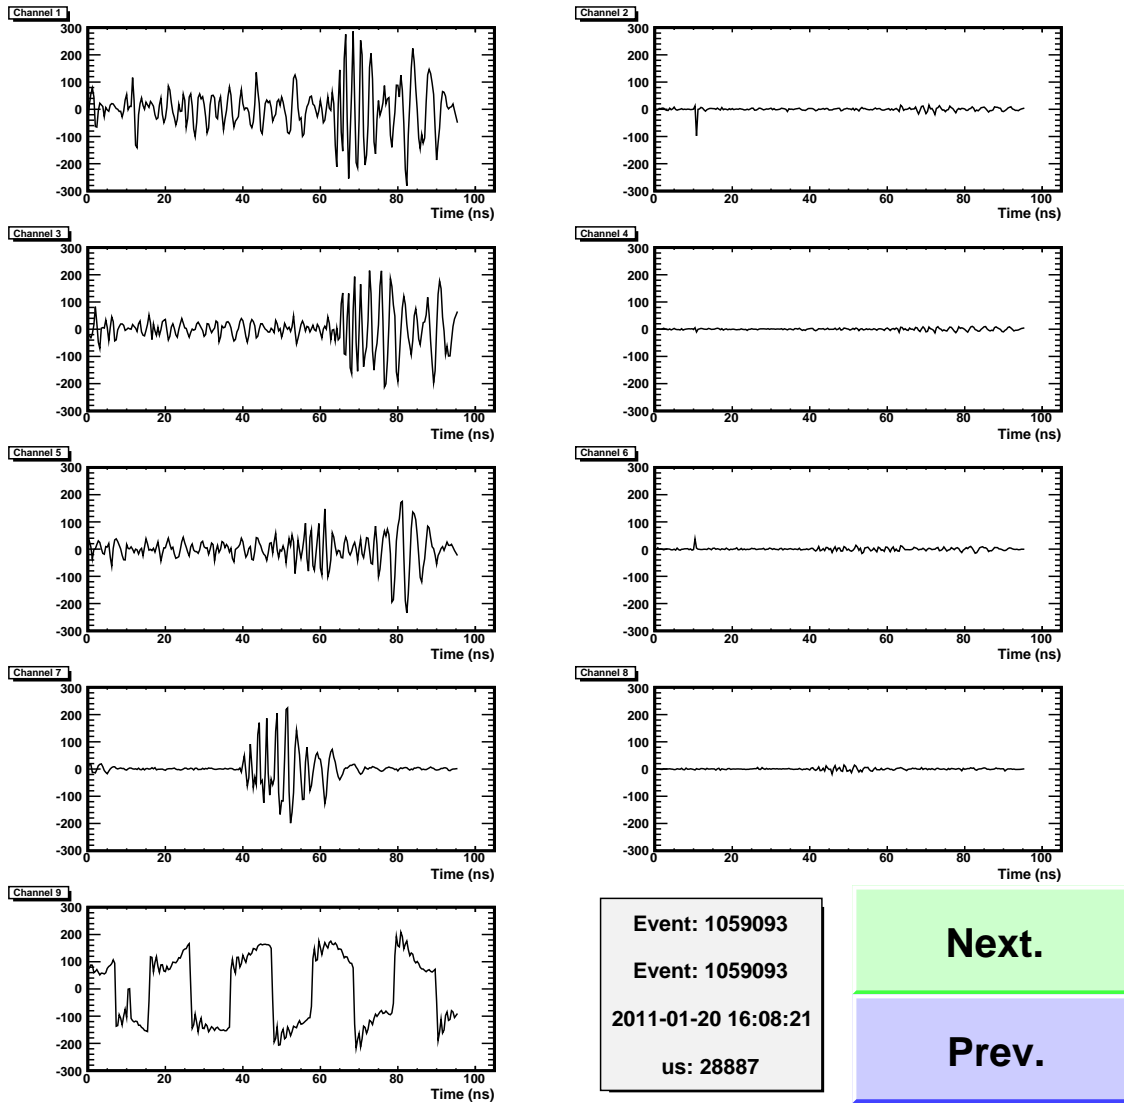


Figure 4.32: This is an example of a heartbeat event from the 2010-11 season. These events are identified by $\Delta t = 1.0 \pm 0.015$ sec, and the relative timing of the pulses in the channels. In this season and in 2011-12, the West channel pulse is now visible because the heartbeat transmitter was moved to a more favorable location.




Figure 4.33: The time T_{pp} (ns) for the North channel of heartbeat events is 65 ± 3 ns. This is the time required by the prototype station to form the trigger in hardware and begin data acquisition.

Figure 4.33 represents a measurement of the North channel heartbeat T_{pp} for the latter two seasons, with the heartbeat transmitter in the more favorable location. The peak to peak algorithm finds that the North channel (the earliest channel that participates in the trigger) waveform is not recorded until 65 ± 3 ns in the data record. In this figure, 95% of the events are between 60 and 70 ns. Eventually we will reject thermal triggers that have $T_{pp} < 60$ ns, because large fluctuations in those channels are guaranteed to be random, since they do not trigger the station. The sample of heartbeat events in figure 4.33 contains 4k events. It is important to note that placing the cut (in later sections) at $T_{pp} < 60$ ns is conservative and safe; placing it at 65 ns or higher runs the risk of excluding signal events.

Eventually we will require each channel to have $T_{pp} > 60$ ns. This cut eliminates many thermal triggers, most of which have two out of four channels satisfying $T_{pp} > 60$ ns. However, this cut cannot eliminate signal events, because the geometry of the detector prevents it.

As the signal plane wave propagates through the detector array, the largest amplitudes in the signal antennas form the trigger. Because the electric field strength does not change significantly with position (it is a plane wave), each antenna is hit by the same electric field strength. The hit angle for each antenna, however, is different. If the signal arrives at 0° in the E-plane of a signal antenna, the signal is strong. If it arrives at the E-plane null (90°), the signal is weak. Because the LPDA antennas were deployed (roughly) in a square, there is no incoming signal geometry allowing the weaker signals to arrive *before* the stronger ones. Thus, requiring $T_{pp} > 60$ ns for all four channels removes no signal, only background. While the presence of noise complicates this picture, it is accurate for large signal amplitudes.

Table 4.21 lists the high and low band threshold settings. For the first season of data, only the high frequency (650-990 MHz) band was activated for triggering. As a result, all of the thermal triggers are small in amplitude and have randomly distributed T_{pp} times. If a signal is small, the chances of the T_{pp} algorithm missing it in favor of a larger, random fluctuation is high. However, just because the low band *threshold* was non active in 2009-10 doesn't mean the signal amplitude wouldn't be large enough to avoid the problem of noise. This argument is also true for the amplitude ($A = v_{pp}/(2v_{rms})$, below) distributions. Throughout this chapter, we rely on the broadband nature of Askaryan pulses to satisfy either or both trigger bands.

Another way to probe the trigger delay in T_{pp} is to look at T_{pp} in the four channels, for thermal trigger events. In all seasons, we can observe that when v_{pp} is large, so is T_{pp} , usually above ≈ 60 ns. Figures 4.34-4.36 contain two-dimensional histograms demonstrating this effect for the 2009-10 through 2011-12 data. The rms noise changes from season to season in each channel, for several reasons. However, in each season, there are event populations of late T_{pp} values, with large v_{pp} values, suggesting the trigger is acting on the thermal noise.

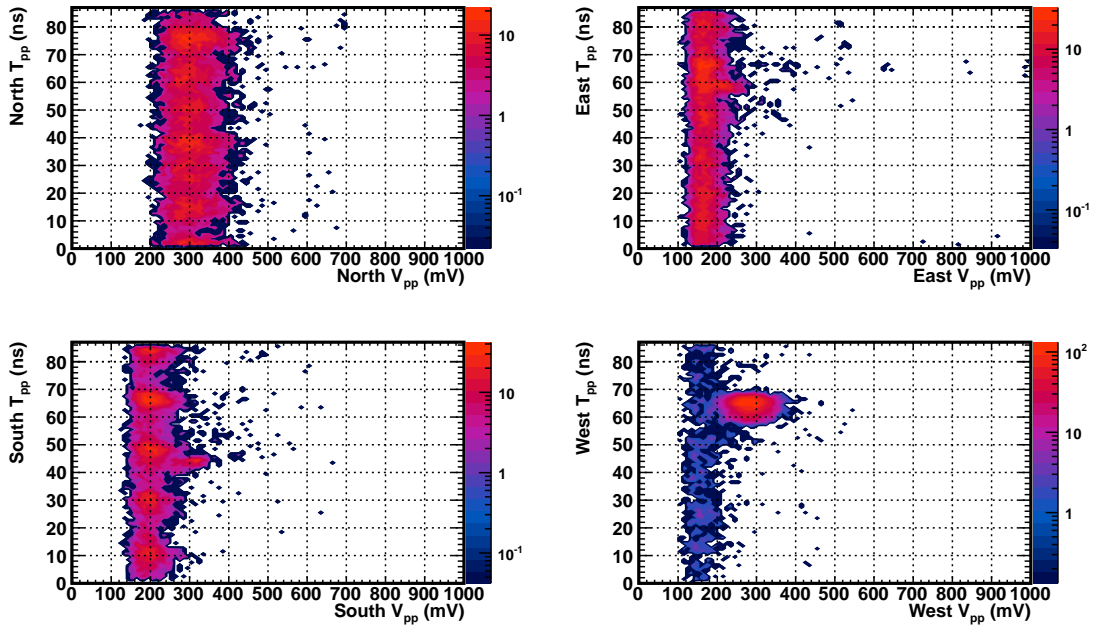


Figure 4.34: T_{pp} (ns) vs. v_{pp} (mV), 2009-10 thermal triggers, after the event-cleaning stage of the analysis.

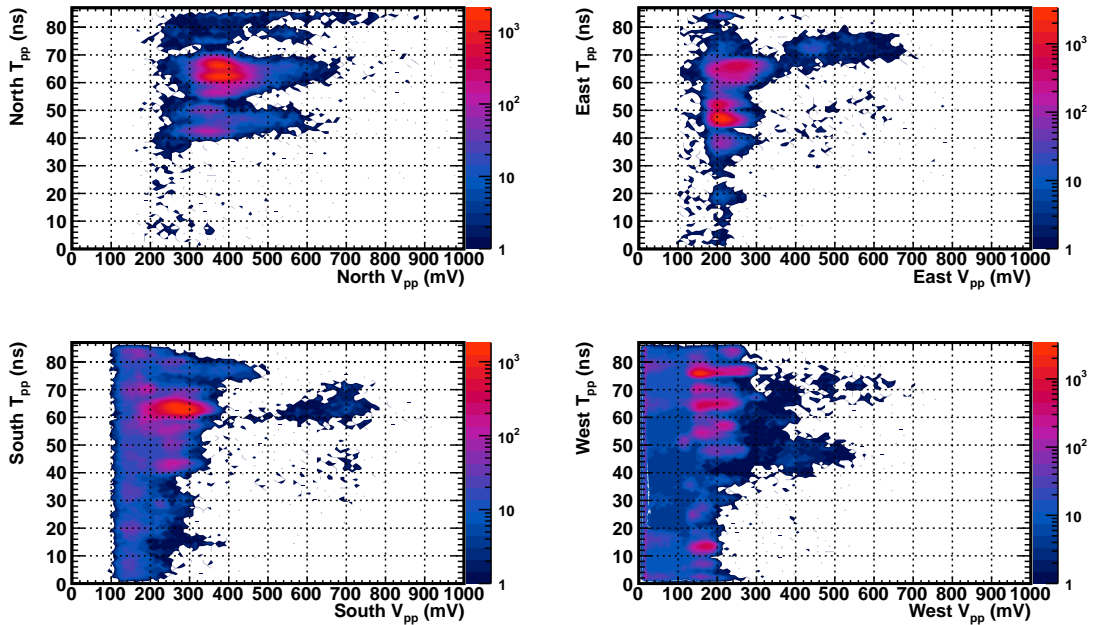


Figure 4.35: T_{pp} (ns) vs. v_{pp} (mV), 2010-11 thermal triggers, after the event-cleaning stage of the analysis.

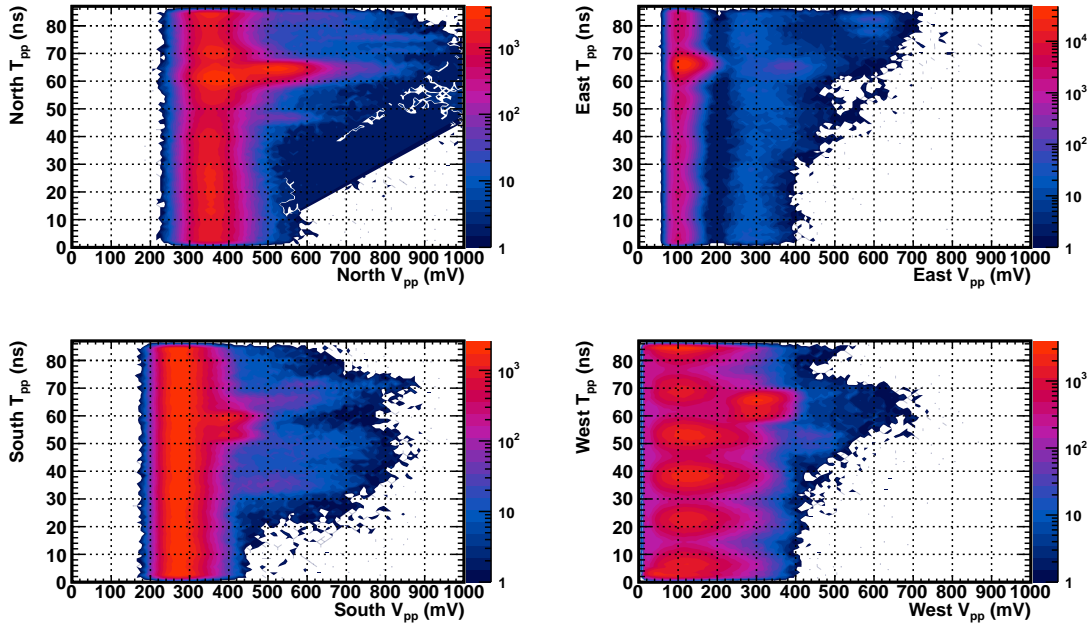


Figure 4.36: T_{pp} (ns) vs. v_{pp} (mV), 2010-11 thermal triggers, after the event-cleaning stage of the analysis.

The West channel in the 2009-10 events typically has a large pulse late in the data record. There is a clear excess in both T_{pp} after 60 ns (y-axis), and in v_{pp} itself (x-axis). For that season, there are excesses in T_{pp} after 60 ns in other channels, but they are not as separated in voltage like the West channel. Again, this is most likely caused by the abnormal West threshold setting that season. For the 2010-11 season, we made adjustments to the thresholds after a period of no thermal triggers. This action resulted in acquiring thermal triggers with the properties in figure 4.35, where both the v_{pp} voltage (x-axis) of thermal triggers stands out after the trigger delay in T_{pp} on the y-axis.

In 2010-11, the event rate was roughly ten times the event rate of the prior year, and thus ten times the events to populate figure 4.36. Figure 4.36, corresponding to 2011-12 thermal triggers, has twenty times the events as the prior year due to the increased live-time, and a similar event rate. The effect of the trigger delay in T_{pp} is evident in all four channels. Early in the 2011-12 season, the read-out characteristics of the East channel (upper right) changed,

lowering the overall voltage scale by a factor of 3 on the x-axis. However, the evidence of the trigger action is still visible (an excess after 60 ns on the y-axis).

Recall that the causality cut required all channel combinations to combine to satisfy a global cut. In effect, this implies that we require all four channels to have a decent sized signal as a result of the external RF pulse. This implies that, given the trigger delay in T_{pp} , we must require all channels to have the T_{pp} greater than or equal to the minimum delay of 60 ns. This is true, for example, of all heartbeat events, which are horizontal plane waves passing across the detector.




Figure 4.37: A Monte Carlo simulation of the effect of the trigger delay on T_{pp} and v_{pp} . The majority logic is set to 2 of 4, the gate is 10 ns wide, and the rms noise is ≈ 40 mV in each channel. Excesses for late T_{pp} values are observed, as in the thermal trigger data.

Figure 4.37 displays a Monte Carlo thermal trigger run, demonstrating the effect of adding a constant trigger delay (with jitter) to the data after a large enough fluctuation occurs in at least two channels. The Monte Carlo thresholds were set to $A > 3.0$ in each channel, for expediency, and the number of triggers taken was 10,000. Once an event with at least two

Channels	Surv. Rate (data/MC)	MC Threshold	Cut (ns)
N	34.7% / 50.4%	3.75	$T_{pp} > 60.0$
E	46.5% / 72.6%	3.75	$T_{pp} > 60.0$
S	34.8% / 45.8%	3.75	$T_{pp} > 60.0$
W	82.4% / 83.2%	3.0	$T_{pp} > 60.0$
Global	5.16% / 9.24%	--	AND all

Table 4.16: A comparison of the trigger delay cut between 2009-10 thermal triggers and MC. The cuts are based on the T_{pp} location in the thermal triggers, staying at least 3σ away from the mean, to account for jitter effects. The Monte Carlo thresholds are set to tune the channel participation rates in order to mimic the T_{pp} rates.

fluctuations within the gate of 10 ns is discovered, the event is saved. The data is rotated such that the earliest T_{pp} time is located at a trigger delay of 65 ± 3 ns (the mean value of figure 4.33). Each channel gets a distinct jitter value, and the four jitter values are drawn from normal distributions with $\sigma = 3.0$ ns. The trigger excesses, which have larger v_{pp} voltages and late T_{pp} values, are easily observed in the results. The thermal background is not gaussian white noise; it follows the PSD of the forced triggers.

Channels	Surv. Rate (data/MC)	MC Threshold	Cut (ns)
N	86.4% / 74.9%	2.5	$T_{pp} > 60.0$
E	31.7% / 31.3%	4.25	$T_{pp} > 60.0$
S	77.7% / 61.3%	2.75	$T_{pp} > 60.0$
W	66.9% / 59.4%	2.75	$T_{pp} > 60.0$
Global	7.5% / 6.6%	--	AND all

Table 4.17: A similar analysis as in table 4.16, but for the 2010-11 thermal triggers.

Channels	Surv. Rate (data/MC)	MC Threshold	Cut (ns)
N	44.8% / 56.4%	3.25	$T_{pp} > 60.0$
E	72.3% / 85.9%	2.5	$T_{pp} > 60.0$
S	46.5% / 60.4%	3.25	$T_{pp} > 50.0$
W	31.7% / 39.2%	3.5	$T_{pp} > 60.0$
Global	4.4% / 9.1%	--	AND all

Table 4.18: A similar analysis as in table 4.16 but for the 2011-12 thermal triggers.

As we saw in chapter 1, the amplitude of the RF pulse that interacts with the signal antennas is proportional to the neutrino energy. A weak signal corresponds to a lower energy neutrino, and a stronger signal indicates a higher energy. The probability that a thermal trigger poses as a neutrino increases if we consider increasingly lower maximum amplitudes. Thus, we must require all channels to have some **minimum amplitude** $A = v_{pp}/(2v_{rms})$ to eliminate those stronger thermal triggers that could pose as fake neutrinos. As we shall see, simulations and heartbeat triggers indicate that the timing resolution is good enough to reconstruct the plane wave properties of the signal, if the amplitude is large.

The unit-less amplitude $A = v_{pp}/(2v_{rms})$ for the four channels in the first season are shown in figure 4.38, and subsequent seasons in 4.39 and 4.40. Note that they appear to limit to an exponential function in all cases, albeit with differing slopes. The exponential fits are good even when the mean and rms of the distributions themselves change from season to season, suggesting that this is a physical effect rather than some effect associated with the detector.

Table 4.19 contains these fit parameters.

Season	North slope	East slope	South slope	West slope	Ave. (seas.)
2009-10	1.91 ± 0.03	2.14 ± 0.04	1.73 ± 0.03	0.95 ± 0.02	1.7 ± 0.3
2010-11	1.33 ± 0.01	0.86 ± 0.01	0.88 ± 0.01	0.85 ± 0.01	1.0 ± 0.1
2011-12	1.979 ± 0.003	2.17 ± 0.01	2.105 ± 0.004	2.87 ± 0.01	2.3 ± 0.2
Ave. (chan.)	1.7 ± 0.2	1.7 ± 0.4	1.6 ± 0.4	1.6 ± 0.7	

Table 4.19: The fitted slopes of the (falling) exponential tails of the A distributions. These number characterize how likely a large pulse is in the various channels of the thermal trigger data. A larger number indicates a softer amplitude spectrum, and a smaller number indicates a harder one.

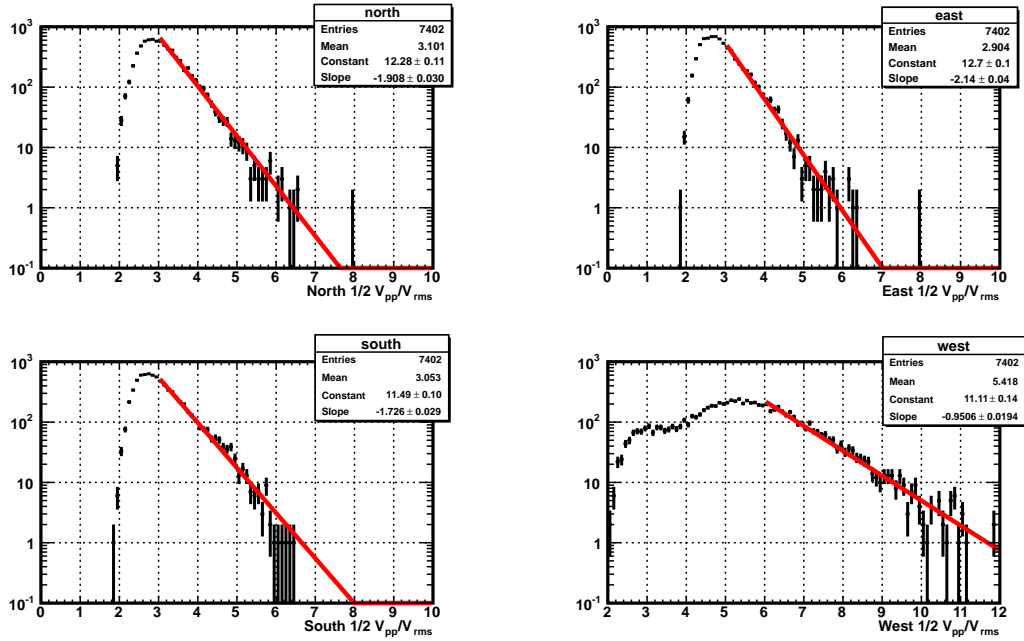


Figure 4.38: The A distribution for the 2009-10 thermal triggers.

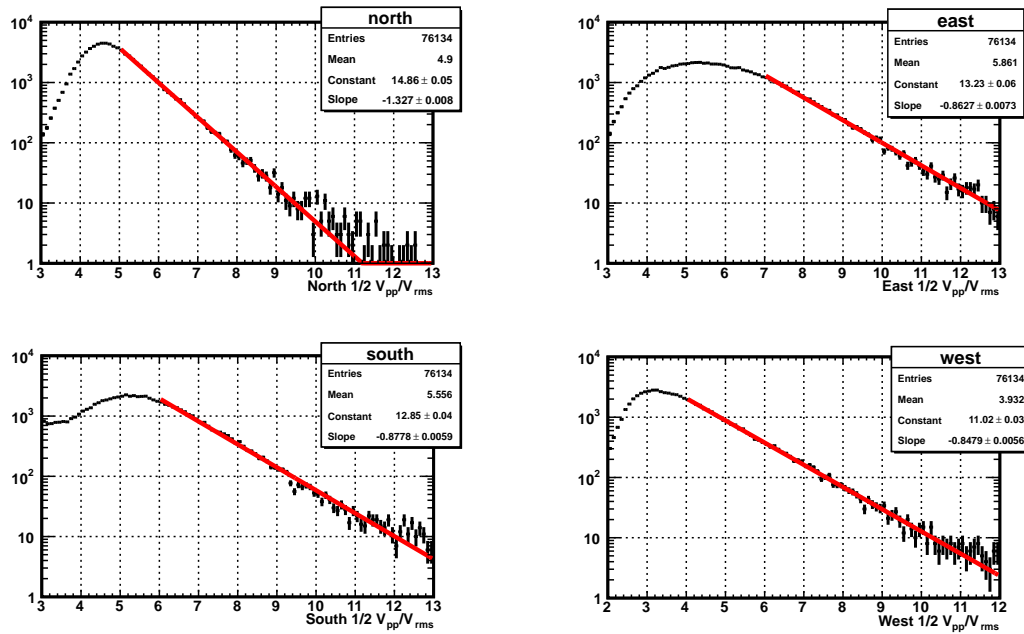


Figure 4.39: The A distribution for the 2010-11 thermal triggers.

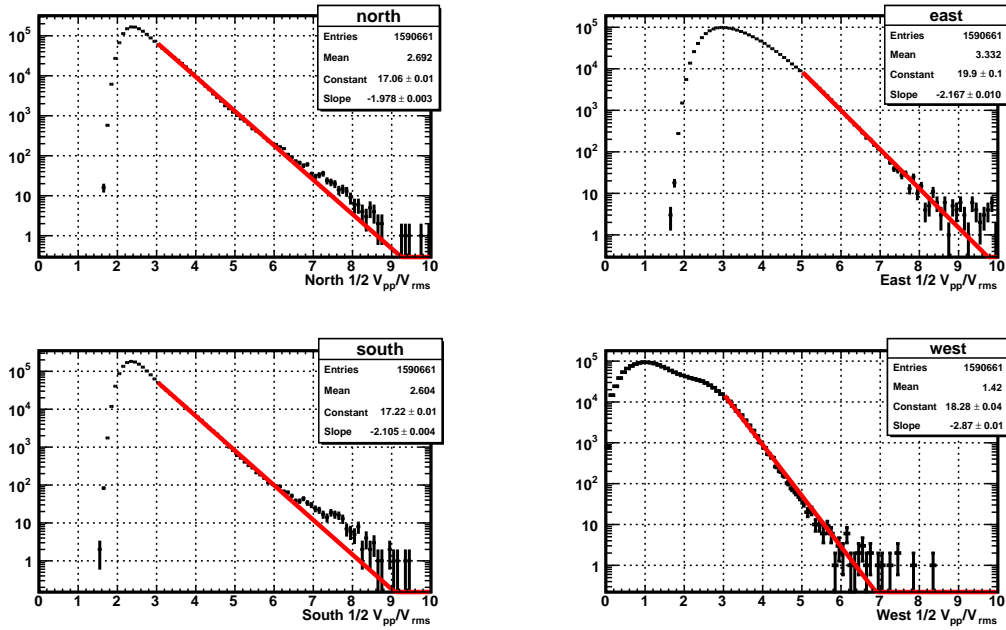


Figure 4.40: The A distribution for the 2011-12 thermal triggers. The West v_{pp} is being compared to the South v_{rms} values because the West channel rms decreased dramatically that season.

There is an interesting pattern in the slopes of table 4.19. The slope parameter represents the average significance of the thermal triggers; a smaller negative slope indicates higher significance, and a larger negative slope indicates lower significance. In the first season, as a result of noise from the wi-fi tower, the West channel threshold was set to the maximum in both frequency bands. As a result, that spectrum is much harder relative to the other channels. The global event rate in the first season (see below) is lower than in subsequent seasons because only the high band trigger was used, with relatively high thresholds in the other three channels. Despite the uneven channel participation, the results are stable throughout the run. The second season, in 2010-11, saw more democratic participation by all channels once all thresholds were lowered to begin acquiring thermal triggers, including those that triggered the low frequency band. In that case, we see that each channel has a hard spectrum, and all are stable over time. The final season, 2011-12, began just as 2010-11 ended, with hard spectra in all channels. However, the averages are lower because we see

a drift in the threshold (figure 4.43). As the season progresses, we began to trigger on less significant events and the event rate increases by a factor of about 5.




Figure 4.41: The fitted slope of the amplitude (A) distribution versus time, for the 2009-10 season (by channel). The fifth graph at the bottom is the inverse of the event rate, and the time axis is in days since January 1, 2010, UTC.

Figure 4.41 quantifies these ideas for the first season of triggers. The slope of the amplitude distribution plotted on the y-axis is stable around the averages in table 4.19 throughout the run, with slopes near -2 for the N,E and S channels, and near -1 for the W channel. Each point in figure 4.41 represents 1000 events, and the slopes for fit on the ranges $[3.5,12]$, $[4,12]$, $[3,12]$, and $[2.5,12]$ for the N,E,S and W channels, respectively. For the 2010-11 triggers (figure 4.42), the ranges had to be changed due to the increased mean of each distribution: $[5,12]$, $[6.5,12]$, $[6,12]$, and $[4,12]$. Each point in that case represents 2500 thermal events (a higher event rate allows for better statistics per point). Again, we see stability in the slope of the thermal noise pulses. The outlying points on day 28 correspond to events thought to be created by the Aerogen 6 wind generator. These events are large as 700 mV in the data and therefore make the amplitude distribution *harder*.




Figure 4.42: The same analysis as figure 4.41 applied to the 2010-11 triggers.

In the final season, a time-dependence emerges in the data. As the inverse of the event rate decreases from about 10 sec/event to 2 sec/event (figure 4.43), the event rate increases, and the slopes from the four channels get increasingly negative. These data can be explained by a drifting threshold voltage. In the final season, no threshold adjustments were made, as in the prior season. As the threshold drifts slowly down, we obtain simultaneously less significant thermal pulses and more events. The fitted ranges are [3,12] for non-West channels and [6,12] for the West channel. Because the rms noise of the West channel disappeared in this season, but we still observed high amplitude pulses being read out from this channel, we allow the South channel rms to stand in for the missing West channel rms values in this final plot. Each point in figure 4.43 represents 10,000 triggers.

Finally, figure 4.44 demonstrates that there should not be any significant time variation in the amplitude variable (as characterized by its fitted slope), based on Monte Carlo analysis. In this Monte Carlo analysis, the thresholds were held constant (at 3σ for expediency - see event rate). The noise spectrum was weighted in the Fourier domain, as usual, to account for the non-uniform bandwidth, and scaled to match the rms values of the end of the 2011-12




Figure 4.43: The same analysis as figure 4.41 applied to the 2011-12 triggers.

season. The slopes remain stable in time, around -2.

Thus the time dependence in the amplitude $A = v_{pp}/2v_{rms}$ is as stable as the trigger threshold, and represents a good variable to discriminate between neutrinos and thermal triggers. It is important to check that the waveform Monte Carlo can predict the correct number of events removed by a minimum amplitude cut, if indeed the large amplitude events are merely thermal fluctuations. This is probably not true of the West channel, with the abnormally large threshold setting (table 4.21). Table 4.20 summarizes how many events are removed channel by channel and globally for the 2009-10 data and MC. The Monte Carlo thresholds are set at different levels to control relative channel participation. The event rate is limited by computation speed, in this case, so while the event rate and slope of the amplitude distributions remain proportional, the absolute event rates should not be compared to prototype station event rates. The noise spectrum of the MC events was weighted in the Fourier domain, and scaled to match the rms values of the 2009-10 season. Similar calculations are done for the 2010-11 and 2011-12 seasons.




Figure 4.44: The same analysis as figure 4.43 applied to Monte Carlo data. Unlike the 2011-12 data, the MC simulation demonstrates no time-dependence in the slope of the amplitude, similar to the first two seasons of thermal data.

The relationship between the programmed threshold in ADC counts for both trigger bands of each channel and the equivalent threshold power is imperfect. The majority logic trigger rate provides some insight on the effective thresholds of each channel. Consider purely Gaussian fluctuations, sampled at $f_s = 2.85$ GSa/s. For pure white noise, each sample s_i is drawn

Channel	Remaining Events (Data/MC)	Cut
N	118/118 \pm 9	$A_N \geq 5.0$
E	127/125 \pm 3	$A_E \geq 5.0$
S	170/130 \pm 10	$A_S \geq 5.0$
W	4606/4520 \pm 10	$A_W \geq 5.0$

Table 4.20: Monte Carlo trials show that the 2009-10 thermal triggers survive the amplitude cuts at about the same rate as simulations. The MC thresholds are tuned to produce the amplitude distributions of the 2009-10 thermal triggers. For the N,E,S, and W channels, the MC thresholds k_i are 3.25, 3.5, 3.5, and 5.75 respectively. The West channel most likely does not participate in the trigger despite the large pulses (see below).

Channel	Year	Low/High Band	Threshold DAC value
N	2009-10	Low	0
N	2009-10	High	990
E	2009-10	Low	0
E	2009-10	High	1240
S	2009-10	Low	0
S	2009-10	High	1265
W	2009-10	Low	0
W	2009-10	High	0
N	2010-11	Low	500
N	2010-11	High	600
E	2010-11	Low	400
E	2010-11	High	500
S	2010-11	Low	1000
S	2010-11	High	400
W	2010-11	Low	0
W	2010-11	High	0
N	2011-12	Low	250
N	2011-12	High	1230
E	2011-12	Low	250
E	2011-12	High	1940
S	2011-12	Low	440
S	2011-12	High	1600
W	2011-12	Low	0
W	2011-12	High	0

Table 4.21: These DAC values can range from [0-4095], and they control the high and low frequency band threshold settings for the three seasons and four channels of data. Because of the noise experienced during the first season by the West channel from the wi-fi tower, we left the West channel threshold setting maximized (set to 0 counts).

from the distribution

$$s_i = N \exp(-z^2/2) \tag{4.39}$$

In equation 4.39, $z = v_{amp}/\sigma$ is the signal to noise ratio, σ is the rms of the thermal noise, and N is an arbitrary normalization with units of mV. Without a loss of generality, we can

pick $N = 1$ mV, since we are only interested in signal to noise ratios. Considering just one sample s_i from this distribution, the probability p_i that the corresponding signal to noise ratio z_i is greater than some threshold z'_i ($z_i > z'_i$) is

$$p_{z'} = \int_{z'}^{\infty} \exp(-z^2/2) dz \quad (4.40)$$

Letting $t = z/\sqrt{2}$, we have

$$p_{z'} = \sqrt{2} \int_{z'/\sqrt{2}}^{\infty} \exp(-t^2) dt = \sqrt{\pi/2} \operatorname{erfc}(z'/\sqrt{2}) \quad (4.41)$$

Here, erfc is the complementary error function, describing the integrated probability of a single sample fluctuating above or equal to z . Using the left hand side of equation 4.19, for the case of 2 of 3 majority logic, we have

$$R_{tot} = 3 \left(\frac{\pi}{2} \right) (\operatorname{erfc}(z'/\sqrt{2}))^2 f_s^2 \tau \quad (4.42)$$

Using $\tau = 10$ ns, $f_s = 2.85$ GS/s, and z values between 5.2 and 5.3, we can explain the event rates in figure 4.41 (bottom) for 2009-10 thermal triggers without invoking the West channel data, or non-uniform thermal noise. But this leaves us with factors of 2 uncertainty for several percent changes in z , which is hardly ideal.

Using the full version of equation 4.19 in section 4.2.2, for the special case of 2 of 3 majority

logic gives us a better approximation, because it takes into account that the bandwidth f_b might actually be less than half the sampling rate f_s . The low-noise amplifiers can be modeled with a uniform bandwidth of $[f_a-f_b]$ GHz, as defined by the 3 dB points (figure 4.10). Setting f_a to 0 GHz and f_b to 0.8 GHz in equation 4.19 and assuming 2 of 3 channels yields

$$R_{tot} \approx 3 \exp(-2I_1^2/2\psi_0) \left(\frac{1}{3} \frac{f_b^3 - f_a^3}{f_b - f_a} \right) \tau \approx 0.007Hz \quad (4.43)$$

We find a similar answer as before ($1/.007 \approx 140$ sec/event), using a threshold $z = I_1/\sqrt{\psi_0} \approx 5.25$, only now we've accounted for the finite bandwidth. The advantage of this calculation is that factors of two in the event rates can be explained with a better understanding of the actual bandwidth, rather than a change in threshold. For example, if the real high end of the bandwidth is 500 MHz, then the average time between events would be ≈ 300 sec.

As it turns out, the reason that Gaussian white noise models become insufficient for comparison to the thermal trigger data is that they do not produce slopes that match figures 4.38-4.40, and therefore they under-predict the event rate (over-predict the threshold). Figure 4.16 in section 4.2.3 demonstrates that when we assume the non-uniform bandwidth structure in the Monte Carlo, we can replicate exactly the forced trigger amplitude spectrum. Here, in figures 4.45 through 4.50, we repeat this calculation with thermal triggers of the three seasons, and demonstrate the deficit in high-amplitude events when we assume only the flat band-limited spectrum.

Figure 4.45 compares data from 2009-10 to Monte Carlo trials. The Monte Carlo trial has the same number of events, and, in the case of the non-uniform spectral weighting, the same spectrum as the forced triggers from that year. Although the spectral shapes do

not change significantly from year to year, we use just the spectrum from the season in question for comparison in all of these plots. The MC thresholds in this simulation control the mean of the distributions. The N, E, and S channel MC thresholds are 2.97, 2.50, and 2.77, respectively. Apparently using these settings produces the correct slopes to match the data exactly. The uniform and non-uniform spectral MC have otherwise identical settings. Because the West channel has such improbable amplitude, we've used the MC threshold of 3.0 and scaled the resulting distribution by a factor of 2.

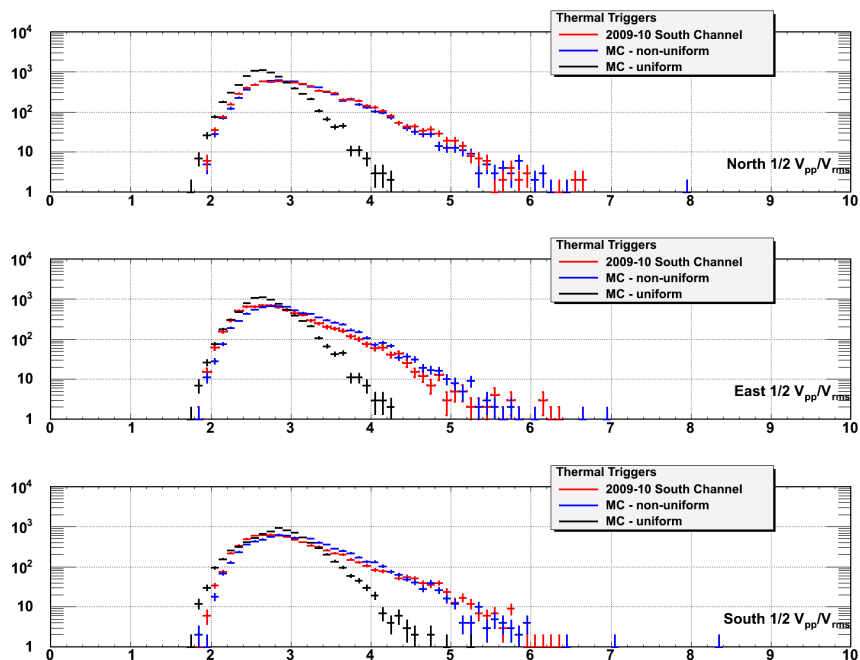


Figure 4.45: The 2009-10 non-West channel amplitude distributions, compared to MC generated according to the procedure outlined in section 4.2.2, and uniform spectral weighting (white noise). The slopes match only match when we weight the MC according to the spectral shape of the data.

Figure 4.47 compares data from 2010-11 to Monte Carlo trials. Because the threshold settings this year included the low-frequency bands, the events are more significant in the data. This makes the Monte Carlo simulations more challenging because we have to either set very large average MC thresholds and run a prohibitively long computation, or set lower thresholds and linearly scale the ensuing histograms. Here we have chosen the latter option. Each




Figure 4.46: The 2009-10 West channel amplitude distributions for data, non-uniform spectral MC, and uniform spectral MC.

non-uniform spectral MC trial uses uniform MC thresholds of 2.5. Then, each amplitude distribution (figures 4.47-4.48) is linearly scaled to match the data. The same scaling factor (≈ 1.8 , from the shift in bandwidth between white noise and real detector noise) in each case is applied to the uniform white noise version of the MC as well. Once again Gaussian white noise is insufficient to explain the amplitudes.

The West channel is again special, because the rms noise appears to be mostly zero during the 2010-11 season. However, heartbeat triggers are still read out normally, with little noise, allowing us to conclude that the channel readout is still functional. This actually isn't that hard to imagine; the trigger and digitization paths are split in the prototype station (see chapter 3). Thus, the readout (digitization) path could be fully functional, but not the trigger path. In figure 4.48 we allow the South channel rms to stand in for the missing West channel rms noise. In the MC trial, the MC threshold was set prohibitively high (7.0), just as the West channel threshold was set to maximum (0 DAC counts) throughout the season. Table 4.22 summarizes how many events are removed by the amplitude cut, how many are

predicted by the non-uniform spectral MC, and how many are predicted by the uniform spectral MC.

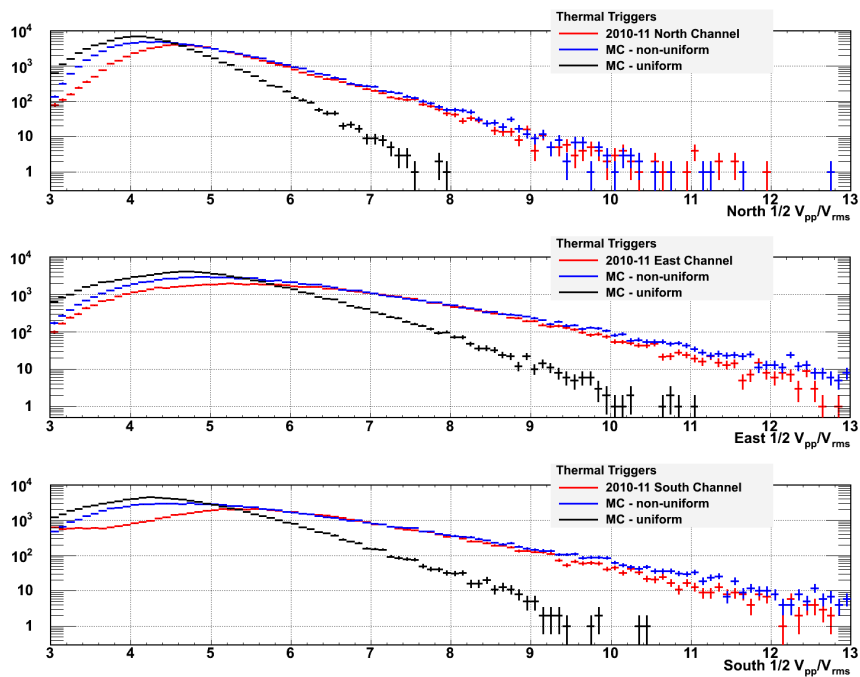


Figure 4.47: The 2010-11 non-West channel amplitude distributions, compared to MC generated according to the procedure outlined in section 4.2.2, and uniform spectral weighting (white noise). After applying a linear scaling, the non-uniform spectral MC explains the thermal distribution (see text).

Finally, we address the 2011-12 season thermal triggers. The thresholds were once again changed, this time in an attempt to trigger mostly on the upper portion of the bandwidth. However, as shown above, the threshold has begun to drift progressively lower at this stage, and the event rate steadily increased. Because of the increased event rate, and the emphasis on the upper band triggering, the means of the amplitude distributions are lowered relative to those of 2011-12, more akin to those of 2009-10. The main difference between the 2009-10 triggers and the 2011-12 triggers is that the latter population still had the low band trigger activated, whereas the former had the low band thresholds set to maximum. The West channel must again be treated separately, because the rms distribution is abnormal.

The 2011-12 distribution contains a much higher fraction of normal rms values than 2010-11,




Figure 4.48: The 2010-11 West channel amplitude distributions for data, non-uniform spectral MC, and uniform spectral MC. The South channel rms noise is used in the denominator because the West channel has lost the rms noise.

where the distribution lived almost exclusively between 2-4 mV. Placing a cut above 5.0 on the standard West channel amplitude (figure 4.50), allows 53% of the events survive. Rather than try to correct the faulty rms values, this channel is removed from the amplitude portion of the analysis. However, the West channel's timing properties are still used (the West v_{pp} values are large enough to cross-correlate). Looking back at figure 4.40, it would be too risky to let the rms values of the South channel stand in, as we did for the prior season, because that produces a distribution where fewer than 0.05% of the events survive a cut above 5.0. However, we can still rely on the timing properties of this channel, as evidenced by the West channel of heartbeat events (figure 4.20). Although the rms noise might be missing, this channel is still capable of digitizing a large RF pulse and it is therefore still capable of providing the time structure for cross-correlation, and the T_{pp} value. Finally, table 4.23 summarizes how many events are removed by the amplitude cut, how many are predicted by the non-uniform spectral MC, and how many are predicted by the uniform spectral MC.

Channel	Data	Non-uniform MC	Uniform MC
N	42.4%	33.5%	10.7%
E	71.9%	65.2%	44.3%
S	71.9%	54.3%	28.9%
W	44.1%	35.6%	8.4%

Table 4.22: The fraction of remaining events for the 2010-11 thermal triggers after a $A \geq 5.0$ has been applied, along with estimates based on two MC techniques. While the Gaussian white noise technique is only correct to within factors of 2-3, to get to the 10%-level we must invoke the non-uniform nature of the spectrum.

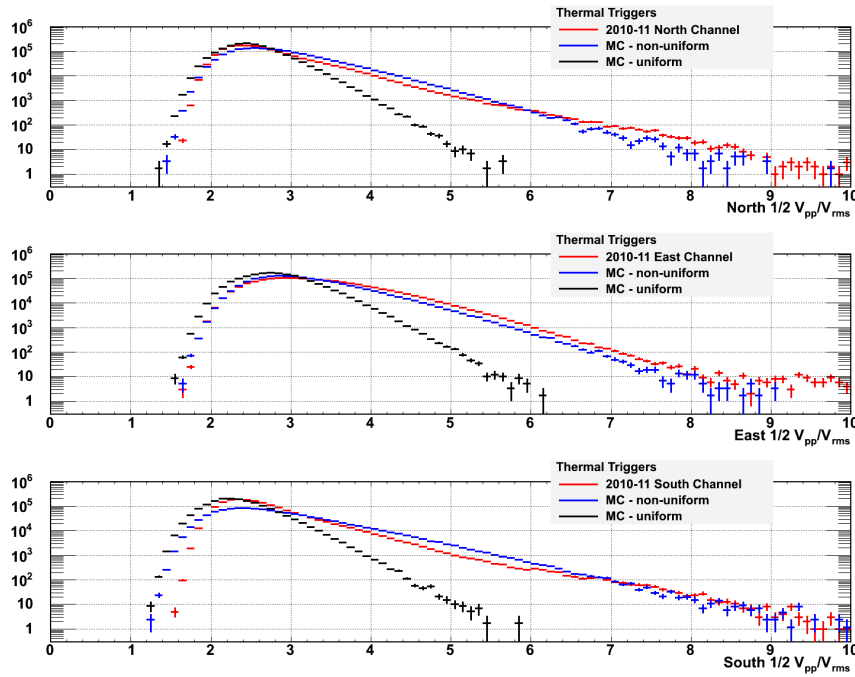


Figure 4.49: The 2011-12 non-West channel amplitude distributions, compared to MC generated according to the procedure outlined in section 4.2.2, and uniform spectral weighting (white noise).

Channel	Data	Non-uniform MC	Uniform MC
N	0.72%	0.8%	$(3.0 \times 10^{-3})\%$
E	3.0%	1.4%	$(3.0 \times 10^{-2})\%$
S	0.53%	1.3%	$(3.0 \times 10^{-3})\%$

Table 4.23: The fraction of remaining events for the 2011-12 thermal triggers after a $A \geq 5.0$ has been applied, along with estimates based on two MC techniques.





Figure 4.50: (Left): The West channel amplitude distribution for 2011-12 thermal triggers, which experiences a modified rms distribution in the denominator. (Right): The top graph is the v_{pp} voltage (mV) of the West channel, and the lower graph is the rms voltage (mV) of the West channel, for 2011-12 thermal triggers.

4.5 Searching for Low-Planarity Events satisfying all other Cuts

With an understanding of how the cuts act on realistic thermal backgrounds, thermal trigger can be eliminated in favor of signal candidates. If a thermal trigger passes all of the cuts described above, and in addition has a planarity value consistent with a plane wave, we might identify it as a neutrino instead. What follows is a summary of the tables and answers listed in the previous section applied to the data from the three seasons. An analysis efficiency is computed for each year, including inefficiencies associated with the Δt filtering and the cut variables. No thermal trigger survives every cut and still has the timing characteristics of a plane wave.

The 2009-10 thermal triggers had a lower event rate because the low-band thresholds were maximized, and the high band thresholds were large as well. The only portion of that season free of wi-fi noise was the final 12.2 days after the wi-fi removal. Thus we are only left with several thousand events and our cuts can be less stringent, and therefore more efficient. At the end of the Δt analysis we are left with 7525 events, however, 80 of these turn out to be heartbeat triggers that escape the Δt analysis, for one of two reasons. Some heartbeat triggers have no subsequent heartbeat event to form a $\Delta t = 1.0$ sec, and this occurs once per run (there are 49 runs). Sometimes the system does not record a heartbeat when it should; this is rare, but it results in a Δt value equal to a small integer multiple of 1.0 sec rather than 1.0 sec. These 80 events were eliminated on the basis that they satisfied one of these two conditions, and had large v_{pp} values (N: ≥ 500 mV, E: ≥ 400 mV, and S: ≥ 300 mV).

To estimate the inefficiency of amplitude cuts, the software package *shelfmc* [15] is used. The details of how the signal amplitudes are generated in each channel are detailed therein so they will not be repeated here. The rms noise level is generated from the usual formula for RF frequencies, $v_{rms} = \sqrt{kBT_{sys}R}$, where k is Boltzmann's constant, B is the bandwidth

(900 MHz from 100 MHz to 1000 MHz), T_{sys} is the system temperature, and R is the load resistance (taken to be 50Ω). After applying the average gain of 66 dB to the answer, we have about 30 mV_{rms} . Forced trigger rms values confirm this approximation in the various seasons. The simulation steering file requires the user to input a global threshold and gate logic; here we use 2 of 4 channels and a base threshold of 5σ . That is, the maximum channel amplitude is at least $5\sqrt{kBT_{sys}R}$ after accounting for the gain. Figure 4.51 demonstrates how the simulation output files obey this criterion. Each event in this simulation is assigned a weight according to the probability it interacts in the Earth before it reaches the ice. After creating a weighted TTree in the ROOT software to account for this interaction probability, 58.3% pass the amplitude cut in table 4.22.




Figure 4.51: The amplitude with respect to the rms thermal noise ($\sqrt{kBT_{sys}R}$) outputted by the simulation *shelfmc* after requiring a threshold of $5\times$ rms noise levels. After accounting for the events' weight, 58.3% satisfy the amplitude cut in table 4.22.

Figure 4.52 shows the planarity variable (section 4.2.3) of the 2009-10 triggers, with each

cut applied to the data in turn. The central green band represents our (naive) expectation for the planarity of a true signal event with perfect timing. Width of this band assumes a planarity precision 400 ps, or 200 ps in each of the four channels (the errors of the four lags adding in quadrature). This 200 ps assumption is probably conservative, since the error in the base planarity bin (section 4.2.4) is 300 ps (150 ps per channel). The mean of the signal region is 0 ns; the analysis of section 4.2.4 suggest the timing offsets cancel within errors. The width of the signal region is 1 ns, or 2.5 times the precision of 400 ps. This is mean to preserve 99% of the signal events with the correct timing. This same criterion is applied to each season's data. There are two other reasons why limiting the signal region to $|P| < 1.0$ ns is necessary. This analysis is restricted to the largest-amplitude events, which should have very little timing uncertainties. Secondly, defining a fixed signal region (with a fixed signal efficiency) also fixes the probability that a thermal trigger lands inside of it.




Figure 4.52: The final results for the first season’s data analysis. The cuts have been defined in section 4.4, and the remaining events are in table 4.24. The event-cleaning level represents the Δt analysis and removal of self-triggered events.

Season	Cut	Value	Events Remain.	Cut efficiency
2009-10	Event Cleaning	Δt analysis	7445	98.9%
2009-10	Causality	$ \tau_{ij} < nx_{ij}/c$	938	$\geq 99\%$
2009-10	T_{pp}	≥ 60 ns all chan.	55	$\geq 99\%$
2009-10	A	≥ 4 all chan.	10	59.2%
2009-10	Plane wave	$ P \leq 1.0$ ns	0	$\geq 99\%$
2009-10	–	–	–	58.5%

Table 4.24: These are the remaining events for each cut level, 2009-10 thermal triggers. The cut $|P| \leq 1.0$ ns encompasses 2.5 standard deviations about the mean. The composite efficiency is given at the bottom.

One question we must ask is whether or not the planarity variable reconstructs to 0 ns given the presence of noise and the amplitude requirement. With a unit-less amplitude of 4.0 in each channel, the answer is that greater than 99.7%. To obtain this result, we can superimpose, for example, the reflection pulse obtained from the 2011-12 ice analysis (chapter 2), scaling the amplitude and thermal rms values such that we obtain the desired unit-less amplitude on average. The LPDA receiver makes this pulse ≈ 100 ns wide, and the receiver system allows frequencies near 100-200 MHz to dominate the signal oscillations. With 10-20 oscillations in the signal, the cross-correlation has no trouble reconstructing properly. Using a much shorter and faster pulse template (for example, a 10 ns wide Gaussian wave packet generated by the MATLAB command `gauspuls`) degrades the integrity of the planarity reconstruction. However, the 2011-12 direct-bounce results are our best model for a real neutrino signal, because they begin as a 1 ns wide signal and are widened naturally by the ice and artificially by the receiving system. The high efficiency of the cross-correlation technique is also why we neglect inefficiencies due to the timing based cuts (`shelfmc` does not even account for noise degradation of timing). Thus, the overall analysis efficiency for this season is $.989 \times 0.592 = 0.585 = 58.5\%$.

Figure 4.53 shows the planarity variable of the 2010-11 triggers, with each cut applied to the data in turn, and table 4.25 presents the events remaining upon each cut. The central band representing the signal region is the same as before. The waveform analysis begins with 90,148 events from the end of the Δt analysis (section 4.3.2) that removed things like Iridium transmissions, heartbeats, and forced triggers. At this stage the analysis efficiency factor is 97%. Of the 90k events, 15,156 turn out to be self-triggered, but randomly occurring in time. They have consistent lags and T_{pp} times, characterized most strongly by the fact that the East channel signal arrives first, and earlier than the usual 60 ns mark. This indicates that the signals are coming from the electronics box, given that the East antenna is buried closest to the box. The North and South channel signals arrive almost simultaneously, and they are equidistant from the box (figure 4.54).

Given the location of the East channel LPDA (figure 3.33), there are very few other directions amenable to a signal arriving that early in this channel. These events are tagged with the cuts listed in table 4.26, at an efficiency of 96.7% with respect to signal. The self-triggered population from the electronics box is present in this season most likely because the low-band thresholds were activated, and the high-band thresholds were lowered. A lower global threshold reveals previously hidden noise sources. The rest of the standard waveform cuts are shown below. The total efficiency with respect to signal is $0.97 \times 0.967 \times 0.523 = .490 = 49.0\%$.




Figure 4.53: The final results for the second season's data analysis. The cuts have been defined in section 4.4, and the remaining events are in table 4.24. The event-cleaning level represents the Δt analysis and removal of self-triggered events.

Season	Cut	Value	Events Remain.	Cut efficiency
2010-11	Event Cleaning, 1	Δt analysis	90148	97.0%
2010-11	Event Cleaning, 2	Self-triggered	74992	96.7%
2010-11	Causality	$ \tau_{ij} < nx_{ij}/c$	15063	$\geq 99\%$
2010-11	T_{pp}	≥ 60 ns all chan.	2047	$\geq 99\%$
2010-11	A	≥ 5 all chan.	26	52.3%
2010-11	Plane wave	$ P \leq 1.0$ ns	0	$\geq 99\%$
2010-11	—	—	—	49.0%

Table 4.25: These are the remaining events for each cut level, 2010-11 thermal triggers. These data represent 20.58 days of live time thresholds were adjusted mid-season to begin acquiring thermal triggers. The composite efficiency is given at the bottom.

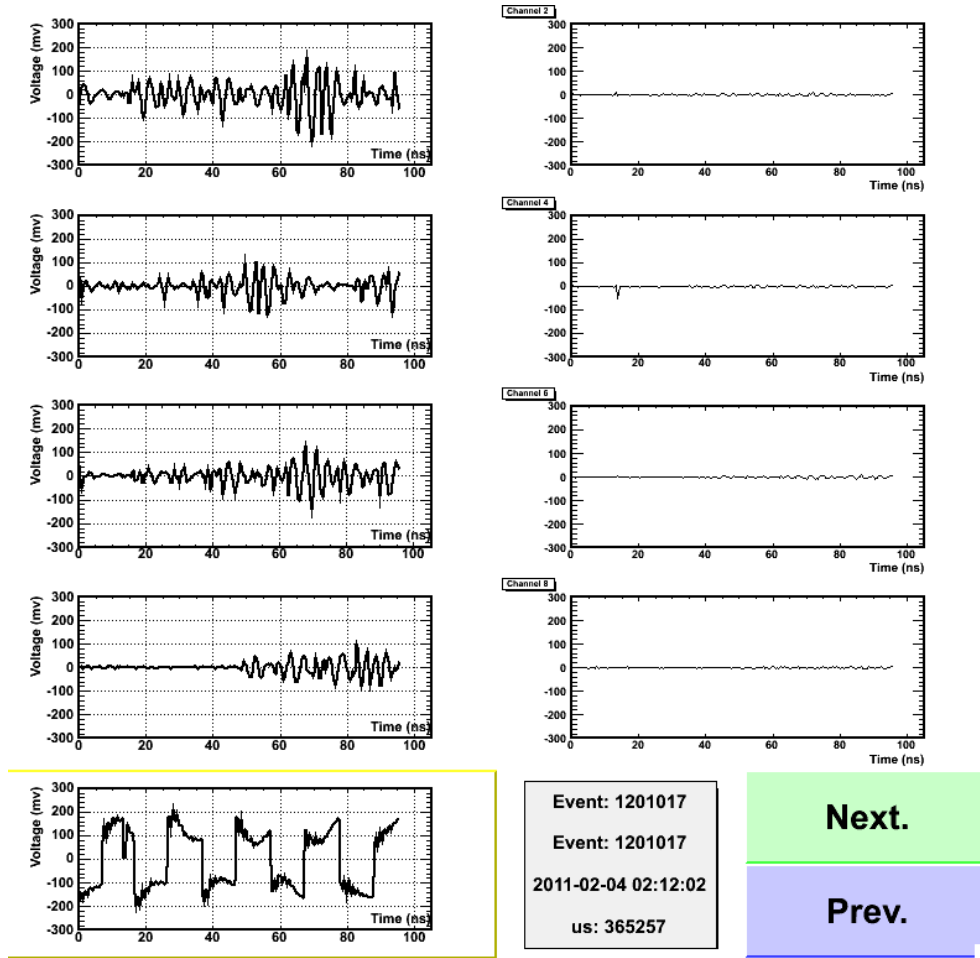


Figure 4.54: There are 15k events that have the same properties as this thermal trigger in the 2010-11 data. Although they occur randomly in time, passing the Δt analysis, we remove them on the basis that they have identical properties from waveform to waveform (table 4.26).

Cut	Efficiency Estimate
$\tau_{SN} < 6.0$ ns	$1.0 - (6.0/36.2) = 83\%$
$\tau_{NE} < 18.0$ ns	$1.0 - (18.0/28.7) = 37\%$
$\tau_{ES} < 9.0$ ns	$1.0 - (9.0/10.7) = 16\%$
East $T_{pp} < 55$ ns	$1.0 - (55/87.0) = 63\%$
AND all	$100\% - 3.3\% = 96.7\%$

Table 4.26: These cuts remove the self-triggered events discussed above with a total signal inefficiency of 3.3%. Signal lags are constrained to be causal, and this provides the denominators in the second column..

Moving on to the final season of thermal triggers, we must be mindful that the quality of the data is starting to change. The threshold is starting to drift lower, and the West channel is only useful for timing calculations. As such, the West channel is simply removed from the amplitude cut, and the waveforms from that channel are used only for timing purposes. This channel still reads out large pulses for heartbeat triggers and noise pulses, but without a clear understanding of the change in gain and the disappearance and reappearance of thermal noise in this channel's data, it is too risky to use the amplitudes. Because every heartbeat contains a West channel pulse, however, we can conclude that when there is a pulse to be read out the West channel can be relied upon to catch it. Removing the West channel from the amplitude cut makes it less restrictive with respect to Monte Carlo signal, and therefore more efficient. To be conservative, however, we assume an average efficiency for the three seasons of 50% in the next section.

Other issues crop up because of the factor of three increase in live-time. One issue is that the unix clock aboard the CPU stack of Icycle1 begins to drift over time, relying upon the GPS daemon to correct it. This drift can grow to 100 sec (according to iceSoft) before the GPSd program resets the clock. If this occurs in the beginning of a run during the heartbeat train, then the Δt of one heartbeat event will no longer be close to 1 second. During the prior two seasons the likelihood of this problem was small enough to avoid observation. After running the waveform analysis on the 2011-12, however, there are heartbeat triggers that pass all of the waveform cuts. Removing any event from the first 205 seconds of a run is sufficient to rid the analysis of heartbeats that escaped the delta T analysis. Technically, this makes the $\Delta t = 1.0 \pm 0.015$ sec redundant, but even under the worst conditions (highest event rate, largest sigma of the delta T cut) the inefficiency introduced is less than 1%.

For all runs of the 2011-12 season, the heartbeat pulser was active. The time it had power was always 100 seconds. Thus, one would expect that the cut should be placed at 101 seconds, not 205 seconds. However, there are heartbeats that survive into the final planarity plot

even after such a cut - event 2769875 is a heartbeat that occurred 204 seconds after the start of run 186, for example. The GPS daemon corrected the unix clock, adding time between heartbeats (figure 4.55). With the maximum skew set to 100 seconds by the software, one might expect a cut of 200 seconds and not 205, however delays introduced by Hkd account for the extra 5 seconds (figure 4.56).




Figure 4.55: The red marker indicates the heartbeat in question, which escaped the Δt analysis because of clock skew. The blue dots are both regular triggers and heartbeats (ie no filtering). The x-axis is the UTC time in seconds, divided by 10^6 . The y-axis is the event number.




Figure 4.56: The time of the execution of Hkd after the first event in a run, plus the heartbeat period of 100 seconds. If Hkd execution was instantaneous, x would equal 100 seconds, because Hkd is restarted just after Acqd is. However, sometimes it takes a little longer (the mean is 105 seconds and the rms is 5 seconds).

The Δt analysis for this season ends with 1.717 million events, beginning with run 57 and ending with run 258. Although technically thermal triggers begin with run 57, this run is almost five days long, indicating non-standard operations (standard operation for this season is two 12 hour runs followed by a short communications run). There are many suspicious noise events occurring in this run, and the event rate fluctuates. Additionally, researchers had to make several post-expedition trips back to the site during this time (December 15-20, 2011), allowing for the possibility of man-made noise. Thus, run 57 is removed from the analysis. This step does not introduce any signal inefficiencies, but reduces slightly the live-time. Similarly, there is a large burst of large amplitude events (figure 4.57) exclusively during run 180. Housekeeping data indicates that there are high winds during this time, however, the wind turbine was disabled in the prior season and it is unlikely to be making RF impulses as large as those of the 2010-11 season. The wind turbine of Icicle2 was fully functional at this time, but it was also 1 km distant from Icicle1 channel antennas. The burst events occur at the same time as a small data gap in the Icicle2 housekeeping data, indicating real-time AFAR and/or Iridium communications to make adjustments to detector settings. Noisy communications are an alternative explanation to pulses from wind turbines.

A final event-cleaning procedure is necessary for the 2011-12 triggers, because there is another randomly occurring self-triggered population throughout the runs. In this case, the events are distinguished by large pulses in the North and West channel that arrive first in those channels. Table 4.27 outlines the cut to remove this population of 11,095 events, and figure 4.58 shows an example of one. Because the West and North channel pulses arrive first, and simultaneously, the most likely source is the wireless tower. Adjustments to the firmware of the AFAR units were made during the final season of this analysis, as well as the wi-fi tower power system. It is possible that some RF shielding component was not reinstalled, allowing some noise pulses to escape that system and propagate to the detector. Unlike the prior season, where a noise source was revealed due to a lower threshold, these events are 3-4 times larger in amplitude and therefore cannot be associated with just the lowering of the




Figure 4.57: The large amplitude events in run 180 are on day 32.5 of the year 2012 (UTC), mainly in the North and East channels. They occur during a prolonged wind storm (Icicle2 wind data), and there is a data gap in the HK data (bottom) during the events as well.

thresholds. Table 4.28 summarizes the event cleaning, and table 4.29 shows the remaining events after the standard waveform cuts are applied, after the event-cleaning. Figure 4.59 shows the planarity of the final, clean population of 2011-12 thermal triggers.

Cut	Efficiency Estimate
North $v_{pp} > \text{all}$	$1.0 - (1/4) = 75.0\%$
East $v_{pp} > \text{South, West}$	$1.0 - (1/2) = 50.0\%$
North $T_{pp} < \text{all}$	$1.0 - (1/4) = 75.0\%$
East $T_{pp} < \text{South, West}$	$(1/2) = 50.0\%$
AND all	$100\% - 1.6\% = 98.4\%$

Table 4.27: These cuts remove the self-triggered events discussed above with a total inefficiency of 1.6%. The North channel is required to be the largest, and earliest pulse. The East channel is required to be the second-largest, and the second-earliest. Assuming the independence of these criteria, the total inefficiency is small.

Cut	Events Remaining
Post- Δt	1717295
No run 57, no run 180	1668728
Wait 205 seconds after run start	1656562
Self-triggered	1645466

Table 4.28: A summary of the event cleaning actions performed on the 2011-12 thermal triggers after the Δt analysis. These actions reduce the amount of data by 4.2%.

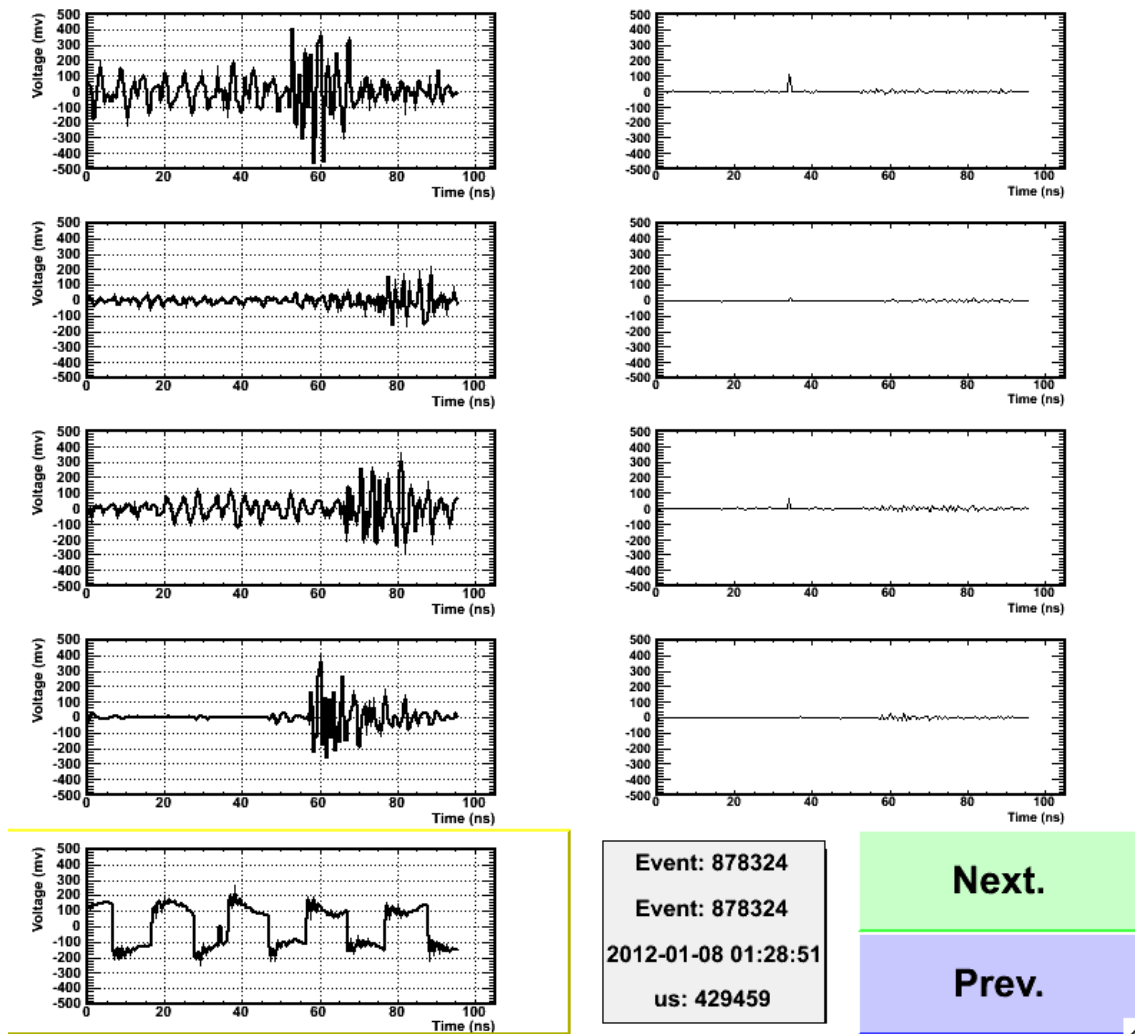


Figure 4.58: Events such as these are common in run 107, but they occur randomly in time (i.e. they are not repetitive in Δt). The cut outlined in the table above selects them with an inefficiency of 1.6% with respect to signal.




Figure 4.59: The final results for the third season’s data analysis. The cuts have been defined in section 4.4, and the remaining events are in table 4.24. The event-cleaning level represents the Δt analysis and removal of self-triggered events.

Season	Cut	Value	Events Remain.	Cut efficiency
2011-12	Event Cleaning, 1	Δt analysis	1717295	96.5%
2011-12	Event Cleaning, 2	Self-triggered	1645466	96.0%
2011-12	Causality	$ \tau_{ij} < nx_{ij}/c$	174043	$\geq 99\%$
2011-12	T_{pp}	≥ 60 ns all chan.	8077	$\geq 99\%$
2011-12	A	≥ 5 (excl. West)	15	64.2%
2011-12	Plane wave	$ P \leq 1.0$ ns	0	$\geq 99\%$
2011-12	–	–	–	59.5%

Table 4.29: These are the remaining events for each cut level, 2011-12 thermal triggers. These data represent 57.13 days of live time (see text). The Δt efficiency estimate comes from a weighted average of the five Δt period efficiencies. The cut on the amplitudes A is more efficient than prior years because it doesn’t cut on the West channel. The composite efficiency is given at the bottom.

The self-triggered and event cleaning stage inefficiency is the sum of a weighted average of the Δt inefficiencies (3.6%), plus 0.5% for the removal of the first 205 seconds of each run, plus 1.6% to remove the 11k self-triggered events. The amplitude cut carries the same as that of the 2010-11 season because the cuts are identical. The final answer for the efficiency of the 2011-12 analysis is $0.965 \times 0.960 \times 0.642 =$, or 59.5%.

As shown in table 4.21, the high-band thresholds for the 2011-12 season in Icicle1 were much lower than in the low-band. This lowers the amplitude, A , of the average thermal trigger due to the response of the data acquisition chain. For this reason alone, the cut on A in the the three non-West channels removes fractionally more thermal events than in the prior season. While the fractional reduction in the remaining number of thermal events is large between the T_{pp} cut and the combined amplitude cut, the fractional reductions involved with the individual channel amplitude cuts are similar to the other waveform cuts. Figures 4.60 through 4.62 plot the amplitude data for each individual cut.




Figure 4.60: The x-axis is the North channel amplitude distribution, plotted sequentially for each waveform cut. The numbers of events remaining after each amplitude are: 8077 (no cuts), 271 (South cut), 106 (South and East cuts), and 15 (all cuts).




Figure 4.61: The x-axis is the East channel amplitude distribution, plotted sequentially for each waveform cut. The numbers of events remaining after each amplitude are: 8077 (no cuts), 271 (South cut), 44 (South and North cuts), and 15 (all cuts).




Figure 4.62: The x-axis is the South channel amplitude distribution, plotted sequentially for each waveform cut. The numbers of events remaining after each amplitude are: 8077 (no cuts), 386 (East cut), 43 (North and East cuts), and 15 (all cuts).

Finally, it is important to quantify the probability that a thermal trigger landed in the signal region. This calculation is essentially an integral of a normal distribution, since the signal region reflects a small portion of the gaussian distributed planarity distribution. The planarity distribution width is 25 ns on average, after the T_{pp} cut has been applied. Making the calculation after the T_{pp} cut accounts for biases introduced by requiring the pulses to all arrive after the the 60 ns mark. (The cut squeezes the pulses closer in the region of the data record between 60 and 90 ns). The fraction of events that randomly fall within the signal region $|P| < \epsilon$ on the planarity normal distribution is

$$F = \frac{1}{\sqrt{2\pi\sigma^2}} \int_{-\epsilon}^{\epsilon} \exp(-x^2/2\sigma^2) dx = \frac{2}{\sqrt{2\pi\sigma^2}} \int_0^{\epsilon} \exp(-x^2/2\sigma^2) dx \quad (4.44)$$

Let the variable $t = x/\sqrt{2}\sigma$. This means that

$$F = \frac{2}{\sqrt{\pi}} \int_0^{\epsilon/\sqrt{2}\sigma} \exp(-t^2) dt = \text{erf}(\epsilon/\sqrt{2}\sigma) \quad (4.45)$$

Using $\epsilon = 1$ ns, and $\sigma = 25$ ns, we obtain 3% for the fraction F of thermal triggers that appear in the signal region. Defining the signal region as $|P| < 1$ ns therefore fixes the probability that the thermal trigger could pose as a neutrino, in addition to guaranteeing good timing reconstruction, as discussed above. There are a total of 51 events in the three seasons that survive the waveform cuts. Thus we expect 3% of them to be in the signal region, or 1.5 ± 1.2 events. The result of 0 events in the signal region is 1.25 standard deviations from the expectation.

4.6 Demonstrations of the Icicle1 Prototype Trigger

One question a skeptic might ask is how we know that the station is not randomly triggering itself throughout the duration of the deployment, rather than triggering on actual fluctuations and signals. One clear demonstration of the effect of the trigger/threshold system is to demonstrate that the events pulled out of the thermal noise are different from forced trigger data, in which we sample the background regardless of the presence of a signal. The most ideal thresholds for this calculation were used in 2010-11, when the the low and high band thresholds had more or less equal footing; in 2011-12 we focused on the lower amplitude, higher frequency triggers. In 2009-10, we did not even use the low band portion of the trigger. The amplifier/LPDA system dictates that the lower frequencies have higher amplitudes: the LPDA effective aperture is larger, and the amplifiers have slightly higher gain.

Comparing the amplitude of the thermal triggers under the conditions of 2010-11 reveals a stark contrast between forced and thermal triggers, as well as heartbeats. Figure 4.63 contains the results. This season most clearly demonstrates the effect of the trigger because the threshold was effectively higher - both trigger bandwidths were active and had roughly equal settings. Forced triggers represent small amplitudes obtained from randomly sampling the rms noise, while thermal triggers must trigger the station. Heartbeats represent a real RF pulse, nearby, and large in amplitude. The histograms are normalized to unity, and the West channel v_{pp} values are being compared to the South channel rms values in this case, because the rms noise is absent for 90% of the West channel forced and thermal triggers. This approach works well for heartbeat triggers; the West heartbeat distribution appears similar to the other channel distributions. However, the West channel thermal triggers do not have the amplitude of the other channels; this makes sense in that the West channel thresholds are maximized. The mean and rms values of the distributions in figure 4.63 are quoted in table 4.30.

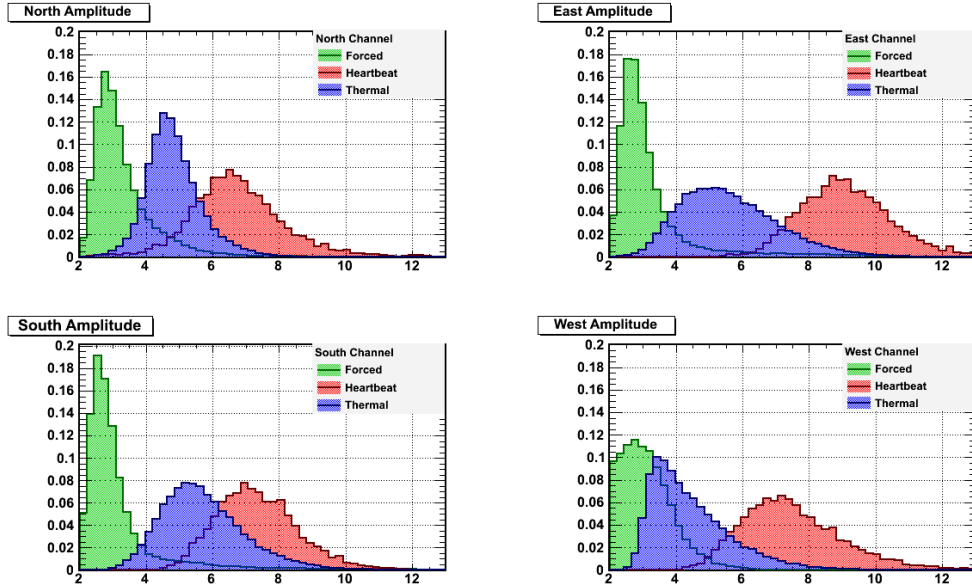


Figure 4.63: The amplitude variable, A , plotted for the three major event classes for the 2010-11 season.

Another effect that demonstrates the action of the trigger is the frequency content of heartbeat events compared to MC and thermal triggers. The dual-bandwidth nature of the trigger means that events can be caused by pulses concentrated at high frequencies (650 to 990 MHz) or low frequencies (130 to 460 MHz) with respect to the total bandwidth available. The high end of the low band, 460 MHz, is actually the 3 dB point of the combination of the filters that define the low band. As we shall see, we are interested in differences smaller than 3dB, so let's consider a signal low-frequency if the majority of the power is below 400 MHz. It is interesting to consider how much more power, for example, is below 400 MHz in an event than in above it. For a neutrino signal, we expect to observe considerable power at high frequencies. For thermal triggers, this is not necessarily true, because of the emphasis placed on low frequencies by the amplifiers and LPDAs. Figure 4.64 clarifies these ideas by plotting the average power (in dB) above 400 MHz minus the average power (in dB) below 400 MHz.

The Monte Carlo results confirm expectations for signals where the difference between the hit angle and the Cherenkov angle is less than several degrees (ie we are observing the event

Channel	Event Class	Mean Ampl.	RMS Ampl.
North	Forced	3.34	1.07
East	Forced	3.30	1.44
South	Forced	3.23	1.45
West	Forced	3.13	1.04
North	Thermal	4.85	0.87
East	Thermal	5.63	1.50
South	Thermal	5.74	1.29
West	Thermal	4.45	1.21
North	Heart	6.77	1.43
East	Heart	9.08	1.32
South	Heart	7.36	1.24
West	Heart	7.50	1.51

Table 4.30: The statistics of the distributions in figure 4.63.

on-cone). The heartbeat triggers mimic a real neutrino signal, containing high-frequency oscillations that originate from the 100 ps rise-time of the Avtech pulser. However, thermal triggers from 2010-11 exhibit a deficit in high frequency power; it is the low frequencies that are more powerful. It is essential to be clear about the decibel units: for thermal and heartbeat triggers the power in dB is expressed as $10 \log(P_\nu/P_{\nu,forced})$, so the dB represents power above forced trigger background noise. For the Monte Carlo we take the power in dB above background as well. Thus, figure 4.64 confirms that *when* there is a high frequency signal, the prototype station catches it, but when we have appropriate thresholds set, low-frequency thermal noise triggers the station.

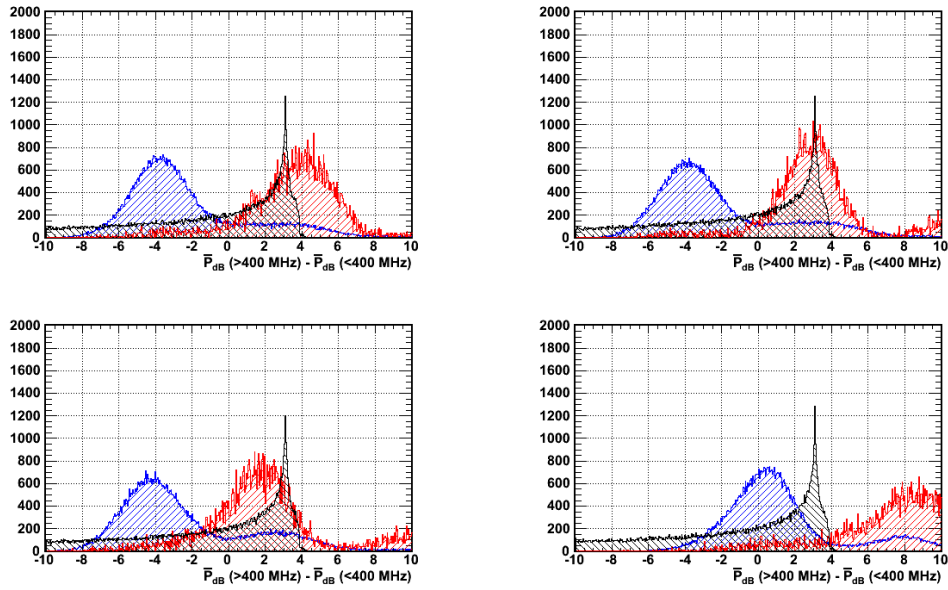


Figure 4.64: The power of thermal triggers (blue), heartbeats (red), and Monte Carlo pulses (black) is concentrated in different bands. The x-axis is the average power below 400 MHz (in dB) subtracted from the average power above 400 MHz (in dB). (North: top left, East: top right, South: bottom left, West: bottom right).




Figure 4.65: In each of the four MC channels, the signal power is concentrated at high frequencies for events that hit the detector at angles close to the Cherenkov angle. As the Askaryan pulse begins to graze the detector rather than hit it directly, we start to lose high frequency power.

There is a physical reason why the Monte Carlo answers are above 0 dB on these plots: when the Askaryan pulse hits the detector at an angle which is close to the Cherenkov cone, we retain the high frequency components. Recall from chapter 1 that the power at high frequencies drops quickly as we observe the Askaryan pulse at increasing angles from the Cherenkov cone. It is also built into the simulation and confirmed in other studies [15]. Figure 4.65 contains the results of the same simulation as the black distributions in figure 4.64. This Monte Carlo run had a threshold of 5 sigma and a majority logic of 2 of 4 antennas hit per station, and a hexagonal array of seven detectors.




Figure 4.66: The time between consecutive Δt versus time of the 2010-11 season during thermal trigger acquisition. The heartbeat events show up as a repetitive signal at $\Delta t = 1.0$ sec.

A final argument for the proper functioning of the trigger is that it is stable. Because we fire a heartbeat pulse once per second at the beginning of each data run, we would discover missing heartbeats in runs if the trigger somehow stopped functioning. Figure 4.66 shows the time between consecutive events versus time, for the 2010-11 season during thermal trigger acquisition. The heartbeat events show up at $\Delta t = 1.0$ second, once per run. Mid-season

we switched to 12 hour runs rather than 6 hour runs, and this is reflected in the figure. The same results are obtained for the other two seasons as well, proving that the trigger remained stable throughout deployment. Again, for these examinations of the trigger, the 2010-11 thermal triggers provide the best demonstration because the thresholds were even between high and low bands, and remained the most stable.

Chapter 5

Neutrino Sensitivity of Prototype Station, and Flux Limits

Given the results of the previous chapter, we should be able to set limits on the neutrino flux achievable with a varying number of stations. This calculation makes use of the simulated effective volume of a varying number of stations [15]. In addition, the Monte Carlo code accounts for other parameters of the deployment, including the majority logic, number of LPDA antennas, thresholds, and live-time. Section ?? begins with a review of the flux prediction based on no events observed, and concludes several curves that quantify the effective aperture of the prototype and arrays comprised of stations similar to the proto-station. This discussion includes modifications of the Monte Carlo code to model the precise antenna locations, and antenna polarization properties. Using this effective aperture, the final sensitivity curve reached by the analysis of the thermal trigger waveforms is presented in section ??, and it is compared to optimistic models of UHE neutrino flux. Additionally, we make predictions of the number of observed neutrinos given an array of stations like the ARIANNA prototype, assuming modest improvements in analysis efficiency and thresholds.

5.1 Effective Volume, Flux Limits

The number of observed particle interactions in a detector like ARIANNA depends on the volume observed. However, it is not necessarily the actual, or *fiducial* volume that enters this calculation. Not every neutrino that interacts within the fiducial volume is observed. Thus, if we use the fiducial volume in our predictions of the neutrino flux, we would artificially improve the sensitivity. The *effective volume* (\times solid angle) is defined as follows [15]:

$$V_{eff}\Omega = V_{fid} \frac{\rho_{ice}}{\rho_{H_2O}} 4\pi \frac{1}{n} \sum_i^m w_i \quad (5.1)$$

In the above equation V_{fid} is the fiducial volume, and the ratio of densities makes this a water equivalent quantity. The integers n and m are respectively the number of thrown neutrinos in the Monte Carlo, and the number of detected neutrinos. The thrown neutrinos are generated isotropically. The weights w_i quantify the probability that the neutrinos were not merely absorbed in the Earth. The weight is derived from the amount of matter the UHE neutrino must traverse at the given trajectory before reaching the fiducial volume.

$$w = \prod_{i=0}^n \exp\left(-\frac{x_i}{L_i}\right) = \prod_{i=0}^n \exp\left(-\frac{x_i \rho_i}{l_i}\right) \quad (5.2)$$

Here the Earth is modeled as a series of layers i of matter with varying densities ρ_i and interaction lengths L_i (m), or l_i (kgm^{-2}). The number of layers the neutrino traverses is n , and each has a thickness x_i . The UHE neutrinos observable to ARIANNA will be attenuated if the trajectory originates from 30% beneath the horizon. The average density for the Earth's crust is assumed to be 3500 kg m^{-3} , and the ice density is 917 kg m^{-3} .

The effective volume, multiplied by the steradians visible to the detector, is sometimes denoted the aperture of the detector. Another useful quantity is the *exposure*, defined as

$$A(E) = (V_{eff}(E)\Omega) \frac{\rho_{H_2O} \sigma(E)}{m_{amu}} t_{live} \quad (5.3)$$

Here ρ is the density of water, σ is the energy-dependent deep-inelastic scattering (DIS) cross section, m_{amu} is the atomic mass unit, and t_{live} is the live-time of the detector. Given equation 5.3, it can be shown that the total number of interacting particles in the effective volume is $\int A(E)\phi(E)dE$. Consider a basic formula for the number of particles interacting through a length L with a flux ϕ (particles per unit area, energy, and time) in the effective volume during a time t :

$$dN = \phi \left(\frac{V_{eff}\Omega t}{L} \right) dE \quad (5.4)$$

We identify the length L with the interaction length of a particle, or the distance over which the number of particles drops by a factor of e :

$$L = \frac{m_{amu}}{\rho_{H_2O} \sigma(E)} \quad (5.5)$$

Also, we recall the identity

$$dE = E \ln 10 d \log E \quad (5.6)$$

Substituting, we have

$$dN = \phi(E) V_{eff}(E) \Omega \left(\frac{\rho_{H_2O} \sigma(E)}{m_{amu}} \right) t_{live} E \ln 10 d \log E \quad (5.7)$$

To calculate the number of expected particles observed, we would put in the flux versus energy and integrate. We can identify within this equation the exposure, and upon integration, we have the total number of interactions observed (if the volume is the *effective volume*):

$$N = \int A(E) \phi(E) dE \quad (5.8)$$

We are free to discretize this integral, since we will want to be able to sum over energy bins and look at individual fluxes in each bin. One choice we can make is to use half-decade energy bins, such that $d \log E = 1/2$.

$$N_{exp} = \left(\frac{\rho_{H_2O} t \ln 10 \Omega}{2 m_{amu}} \right) \sum_i \phi_i V_{eff,i} E_i \sigma_i \quad (5.9)$$

To examine the number of neutrinos above a certain energy bin, we note that the flux is most likely a quickly decreasing power law, with a spectral index close to -2, such that the flux at the energy bin in question is close to the total flux above that energy. The Feldman-Cousins [47] expected upper-limit for a 90% confidence-level given no observed neutrino events in a given flavor, and at a given energy, is $N \leq N_{max} = 2.3$. The number 2.3 comes from the choice of confidence level, $1 - \alpha$. If $1 - \alpha = 0.9$, then $\alpha = 0.1$ and N_{max} is given by the formula $N_{max} = -\log \alpha \approx 2.3$. We have

$$N_{max} \geq \left(\frac{\rho_{H_2O} t_{live} \ln 10}{2m_{amu}} \right) E_i \phi_i V_{eff,i} \Omega \sigma_i \quad (5.10)$$

Finally, we solve for the flux (using the interaction length), and scale the effective volume by a factor of three to account for the three possible neutrino flavors observed:

$$\phi(E_i) = \phi_i \leq \frac{2N_{max} L(E_i)}{3 \ln 10 t_{live} E_i V_{eff,i} \Omega \epsilon} \quad (5.11)$$

The analysis efficiency factor ϵ is included in the denominator, which is the fraction of signal captured by the cuts. The effective volume has been simulated using the *shelfmc* code (Appendix B), and the results are shown in figure 5.3. The effective volume is almost linear with the number of stations, given a 1 km station separation (the slope is 0.94). The results of ARA-37, the final proposed version of the Askaryan Radio Array, come from [36]. Although the full version of ARIANNA (31 \times 31 stations) seems an improvement over the final ARA-37 results, it is unlikely to matter because the bulk of the UHE neutrinos (as we saw in chapter 1) from the GZK effect will have energies between 10^{18} and 10^{19} eV, where the two experiments have similar results.

Before using the baseline effective volume estimates of the Monte Carlo to calculate the sensitivity, we must discuss modifications made to it to more accurately model the detector. This includes both natural effects, such as the properties of the ice shelf previously not understood, and the geometry of the real detector in Moore’s Bay.

Shadowing Effect. First, It is important to note that the effective volume plotted in figure 5.3 includes the shadowing effect. The details of how the shadowing effect is modeled in the Monte Carlo are summarized in figure 2.24. There are four types of events; events that do not reflect but proceed to the detector from the bulk ice in the inclusion zone, reflected events that travel through the bulk ice to the detector in the inclusion zone, and direct and reflected events that originate in the excluded zone.

The exclusion zone is so named because the depth-dependent refractive index pulls the Askaryan pulse away from the detector if the neutrino interacts in the exclusion zone. An increase in the number of missed events decreases the effective volume (see equation 5.1). Simulation runs with and without the shadowing effect activated indicate there is a factor of two difference in the effective volume depending on the existence of the effect. In chapter 2 we discuss data taken in 2011 that prove RF pulses can travel up to ≈ 543 m without *any* shadowing at the surface. Thus, the effective volume used below for the flux prediction is a conservative estimate, which will increase if it turns out that any shadowing in Moore’s bay is smaller than modeled.

Antenna positions. Next, we have to consider the fact that in previous studies the antenna locations were treated as an ideal polygon, with a baseline of 3.75 m. While this is close to the actual deployed array of LPDA antennas, it can be made more accurate given the precision we have for the antenna locations (figure 3.33). Given the experimental geometry, the average neutrino interacts and produces a pulse that must travel usually over 1000 m before reaching the detector. The real antenna positions deviate from the ideal square assumed in the simulation by only a few meters, and thus the expectation is that the electric

field strength does not change significantly enough to affect the detection of the event.

Antenna polarizations. A stronger effect resulting from the real antenna positions is the antenna polarization. Let the x-y plane represent the surface of Moore’s Bay, and the z-axis describe the depth. The Askaryan pulse is linearly polarized; by placing linearly polarized antennas in a polygon, we allow for the possibility of polarization measurement. Ideally, the polarization of the signal antennas would alternate between x-axis and y-axis alignment. This maximizes the chance of matching the Askaryan polarization from any direction. However, the real antenna positions do not have the ideal polarization alignment, with alternating polarizations. The worst case scenario is that we see a factor of 2 reduction in the effective volume. If every antenna has the same polarization, there are two quadrants in the x-y plane from which an incoming Askaryan pulse would encounter *only* antenna null zones.

The null zone of the LPDA is located 90 degrees from bore-sight in the E-plane. The Monte Carlo models the radiation pattern of the LPDA antenna by assuming the decrease in gain is described by a two-dimensional Gaussian function in θ and ϕ (in the E and H planes). The σ parameters of this Gaussian function are θ_{HPBW} and ϕ_{HPBW} , derived in prior chapters. The Monte Carlo reveals a reduction in the effective volume when the worst possible antenna alignment is used, however it is not the expected 50% reduction but closer to 20% at 10^{17} eV, and 10% at 10^{19} eV. The discrepancy is explained by the Gaussian model; while this is a good approximation near the bore-sight of the antenna (within the HPBW of the antenna), this model does not account for the 40 dB loss in gain in the antenna null zone. Because the antenna polarizations are closer to the ideal case than the worst possible case, we simply live with the 10% level uncertainty in the effective volume calculations.

Filtering. Another source of error in the effective aperture is the presence of filters in the prototype system. The high and low pass filters are located just before the low noise amplifiers. The high pass filter was a model NHP-150+ (Mini-Circuits), with a 3dB point of 120 MHz. The low pass filter was an NLP-1200+, with a 3 dB point of 1200 MHz. The

amplifier pass-band (as defined by 40 dB drops in gain) is approximately [50-1000] MHz. The primary systematic is the effect of the low frequency cutoff, because the magnitude of the effective height (section 3.2.2) of the LPDA is inversely proportional to the frequency. Thus, an increase in the low frequency cutoff would decrease the maximum effective height and therefore the total response of the system. However, the Monte Carlo is already restricted to [100-1000] MHz, so changing the low frequency cutoff from 100 to 120 MHz is only a 10% effect. This 10% effect, which is reproducible in the $V_{eff}\Omega$ simulations, can be understood by integrating the effective height between [120-1000] MHz and [100-1000] MHz, taking the ratio of the results.

Majority logic gate. The antenna positions as deployed dictate the maximum travel time between channels for a plane wave signal. The majority logic gate is the maximum separation in time between single channel triggers that can satisfy the majority logic number (2 out of 4 channels). The issue with the gate and the physical separation between antenna channels is that the gate is shorter than the propagation time between certain antenna pairs. In fact, there are two digital gates signals per channel. One signal is raised to digital HIGH for 8.3 ns (two clock cycles of a 40 MHz clock at $6\times$ speed) if an incoming signal satisfies the channel threshold. Another signal checks to see if the threshold is satisfied each clock cycle (25 ns), and if so, it raises a gate signal for one clock cycle. The 8.3 ns gate and the 25 ns gate are always subjected to a logical OR operation. The output of this operation is the single channel trigger signal. If two channels send gate pulses that overlap, then the system latches the digitized data and records an event.

Table 4.13 shows the maximum propagation times between channel pairs used in the causality variable. All of these times are larger than the 8.3 ns signal. If the neutrino pulse in the detector was as short as it is at production ($\approx 1ns$), then Askayran pulses with propagation vectors close to the horizontal would never satisfy the majority logic. However, there are two effects that mitigate this problem. First, the analysis requires all channels to have a

large $A = v_{pp}/2v_{rms}$ value. Thus, all channels should satisfy the threshold, and raise the gate signal (8.3 ns). The East and South antennas are the closest and therefore govern the majority logic action. Monte Carlo studies show that the time difference in signal arrival in these two channels is normally distributed with a mean 0 ± 9 ns. The error in $\sigma = 9$ ns is 0.13 ns, and this answer is obtained from a weighted analysis using the weighting procedure described above. The amount of signal that satisfies the gate automatically is therefore $1 - \text{erfc}(8.3\text{ns}/9.0\text{ns}/\sqrt{2}) = 64.2\%$.

This result is certainly an underestimate, because the detection chain stretches signals like Askaryan pulses, such that the majority logic is actually easier to satisfy. Figure 5.2 demonstrates how an averaged reflection pulse from the 2011-12 ice studies contains ample power for times comparable to the largest time separation between data channels. Figure 5.2 is a spectrogram created from the modulus squared of a short-time Fourier transform of the 2011 direct bounce reflection data. The integration time for each bin along the frequency axis is 9.0 ns. The frequency resolution is 50 MHz, and the overlap between integration windows is 50%. The 2011-12 direct bounce reflection from the Seavey radio horn transmitter and the LPDA receiver had a peak to peak voltage of ≈ 2 V. This signal experienced the average path length of an Askaryan pulse from a neutrino interaction [15], including attenuation effects. Despite geometric and absorptive losses, the signal power remains high for times longer than the Icell1 gate signals, and comparable to the baselines between antenna pairs. This is true even for the high-frequency trigger band. Thus, the first data channel trigger remains HIGH long enough for the signal to propagate to the second data channel.

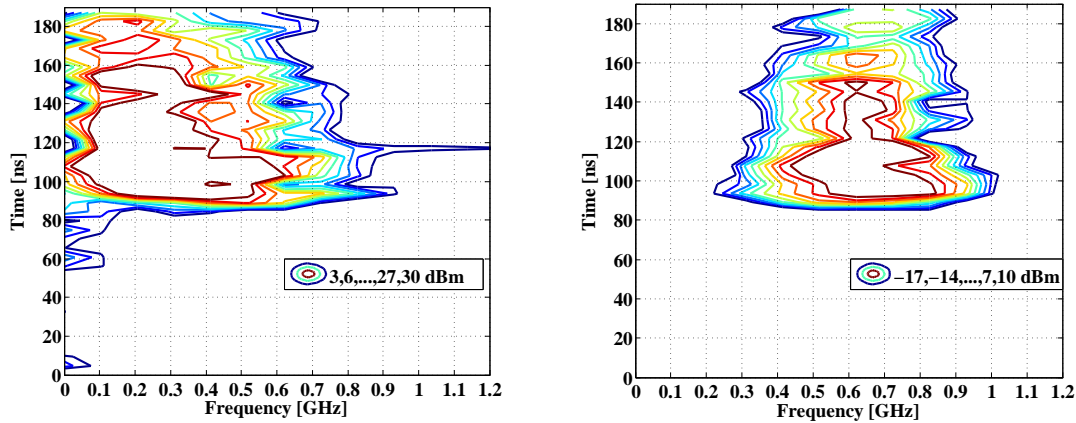


Figure 5.1: (Left) A spectrogram created from the 2011 direct bounce. The contours range from 3 dBm to 30 dBm. The lower portion of the bandwidth dominates the power. (Right) The same data as figure 5.2, but filtered with a high-pass filter (type-I Chebyshev, 660 MHz 3 dB point) to reveal the structure at high frequencies. The number 660 MHz is chosen to reflect the high-frequency trigger band. The contours range from -17 dBm to 10 dBm.




Figure 5.2: (Left) The same analysis technique as figure 5.2 applied to the Seavey to LPDA calibration signal obtained in 2011 (through 23m of air). (Right) Calibration signal, with the same high-pass filter.

The standard settings of `shelfmc` were used to produce the effective volume versus energy curve in figure 5.3. The standard parameter values are listed so the effective volume curves are reproducible. The first parameter is the neutrino energy. There are two choices in the Monte Carlo code. One can either throw neutrinos of a single energy only, or employ a known GZK spectrum. In either case the number of generated neutrinos remains fixed. To plot the effective volume vs. energy, one must generate the same population of neutrinos at each given energy. Second, the spacing in meters between stations is set to 1000 m. This setting is irrelevant for a single station, but important for tables below predicting the number of neutrinos visible to arrays of ARIANNA stations. The number of signal antennas per station, and the number of antennas to form a trigger are also tunable. Here, to accurately model the Ice1 prototype, we chose 4 signal antennas and 2 to form a trigger. The waveform analysis further requires four channels to have a large pulse, in order to reject background and reconstruct the timing, but technically introduces inefficiency.

The next parameters required by the MC describe the firn: the depth and index of refraction are set typically to 60 m and 1.3, respectively. The index of refraction we have measured and reported in chapter 2. Our understanding of the depth is less precise, between 50 and 70 m before the transition to pure ice. For example, our measurement of the depth of the Ross Ice Shelf at Moore's Bay assumes a depth-dependence function of the index of refraction based on measurements from Williams Field [40]. The transition to pure ice in this model occurs at 67 m. The Monte Carlo, however, uses 60 m. The error analysis of the ice shelf depth reveals that the firn depth can have an error of 10 m and not affect the depth measurement result. Thus, any number between 60 and 70 m should produce similar results.

Attenuation losses are calculated in the Monte Carlo based on the depth-dependence of the ice temperature [25], adapted from results at the South Pole. The average attenuation length is used to simplify the calculation of the signal power that arrives at the antenna. Armed with the average attenuation length, the path-integral over the temperature (depth)

dependance is then performed. The average attenuation with respect to frequency should be weighted by the effective aperture of the antenna:

$$\bar{\lambda} = \frac{\int_{\nu_1}^{\nu_2} A_e(\nu)\lambda(\nu)d\nu}{\int_{\nu_1}^{\nu_2} A_e(\nu)d\nu} \quad (5.12)$$

Recall that the effective aperture of the antenna at a single frequency ν is given by $A_e\Omega_A = (c/\nu)^2$, where Ω_A is the beam area. Thus, $A_e \propto 1/\nu^2$, and we have

$$\bar{\lambda} = \frac{\nu_1\nu_2}{\Delta\nu} \int_{\nu_1}^{\nu_2} \nu^{-2}\lambda(\nu)d\nu \quad (5.13)$$

Assuming $\lambda(\nu)$ follows the linear fits to the attenuation length data from chapter 2, the average attenuation length turns out to be ≈ 450 m. The next parameter is the threshold parameter is set to 5 in the MC file, and it is the usual number of standard deviations about the thermal noise level (\sqrt{kBTR}). We assume an ice-ocean reflection coefficient of 1.0, this is justified by our 2011 measurements. Prior studies have shown that lowering it to 0.5 does not have a strong impact on the effective volume results. As discussed above, the shadowing effect is also turned on, using the firn density depth-dependence. Table 5.1 summarizes the MC settings.




Figure 5.3: The effective volume times steradians (aperture) for three versions of ARIANNA: a single IceCube1 prototype station, a hexagonal grid of seven stations (Hexagonal Radio Array), and a 31 by 31 station square grid. These results are averaged over neutrino flavor, and the station separation is fixed at 1000 m. Also shown are the points for the ARA-37 simulation [36].

Effect/Setting	Value	Comments
N_ν	200000	Limits the statistics
Exponent (Energy)	17.0-21.0	10^γ eV
ATGap	1000 m	station separation
ST_Type	4	station type
N_Ant_perST	4	No. of channels
N_Ant_Trigger	2	No. of channels to trigger
Firn	1	use the firn
Firn index	1.30	from ice measurements
Firn depth	60	± 10 m
NRows	3	Hexagonal Array setting
NColumns	5	Hexagonal Array setting
Spectrum	0	Do not use if single energy
NSigma	5	Threshold
Reflect_Rate	1.0	Reflection coefficient
GZK	1	The energy distribution if Spectrum
Shadowing	1	Turn on shadowing
Depth Dependent	1	Depth-dependent firn

Table 5.1: These are the parameters taken to produce the final estimate of the effective volume.

5.2 Sensitivity to UHE neutrinos

Using equation 5.11 and the results for the detector aperture, we can derive the experimental sensitivity to various models of UHE neutrino fluxes. Taking a 90% confidence level and half-decade energy bins, we obtain the upper red curve (ARIANNA prototype) in figure 5.4. We follow the prescription from Feldman and Cousins in the previous section, where we place 2.3 events at each half-decade energy bin. Thus, this curve is a differential flux limit at the 90% confidence level. The results of several other experiments are shown, and references are given, however the corresponding effective volumes are not plotted separately. The more ambitious ARIANNA differential flux limit is calculated in the same fashion. For this limit, however, 3.0 years of live-time, 100% analysis efficiency, and a 31×31 station are assumed. This curve is thus an estimation of how sensitive the experiment can become, but it is also easily understood by scaling the upper result by number of stations and live-time. (The effective volume is linear in the number of stations, given a 1000 m station separation [15]). It is this sensitivity which is used to produce table 5.2 at the 90% confidence level.

Table 5.2 shows the predicted number of counts after performing the discretized integral in equation 5.9. The model rejection factor [56], $2.3/N_\nu$, is also shown next to the corresponding model prediction. While there are a wide variety of models affecting the GZK neutrino flux, the majority produce a non-negligible number of signal events. We can reject a model for the flux if the model rejection factor (MRF) is less than unity. We use the following function to relate the cross section to the energy: $\sigma(E) = 7.84 \times 10^{-40} E[\text{GeV}]^{0.363} \text{ m}^2$. This power law fit matches the total deep-inelastic scattering cross-section from [3], where the total cross-section means the sum of the contributions from the charged current and neutral current interactions.

All but a few of the listed models could be detected or rejected with 3.0 years of live-time. These models fall into several categories, each limiting the UHE neutrino flux in one or

more ways. These mechanisms have already been discussed in the first chapter, however it is worthwhile to recount them here. If the UHECR sources follow the evolution with respect to redshift of other astrophysical objects, then an enhancement is predicted in the neutrino flux. Without direct evidence of the source of UHECRs, the evolution functions can track a variety of phenomena: star formation rates (SFR), quasi-stellar objects (QSO), or even gamma-ray bursts (GRB). There are few kinds of astrophysical objects that do not evolve with redshift, however [33] gives at least one example. Other considerations with respect to redshift involve the curvature of the universe, that is, what value we assume for the cosmological parameter Ω_Λ . Stemming again from the uncertainty of the UHECR sources, the maximum initial energy E_{max} of the protons/nuclei at the source is more or less an unbounded number, although most references do not take values more than 2-3 orders of magnitude above the highest observed energies at Earth. E_{max} is generally assumed to scale with the charge Z of the primary particle: $E_{max}^Z = Z \cdot E_{max}$ [61] [79]. Also recall that the energy upon UHECR interaction with the CMB is redshifted by a factor $(1 + z)$, where z is the redshift. The same is true of the neutrino by-products.

Finally, the chemical composition of the UHECR primary particles can be measured, although there is no single consensus for the fraction of Fe nuclei from current experiments such as Auger, HiRes, and Telescope Array. The UHECR energy spectrum from HiRes is shown in the first chapter, where the GZK cutoff is described by a broken power law assuming pure protonic composition. Figure 1.7 demonstrates how difficult it is to draw conclusions about quantities like the iron fraction above 10^{18} eV, given experimental uncertainties. One of the main motivations behind the first UHE neutrino detection is to establish a flux that constrains the maximum amount of iron present in cosmic rays. Several more extreme examples of models with *only* iron primaries are from Ave et al. [24] and Kotera et al. [33]. ARIANNA is sensitive to both of these models with 3.0 years of live-time, and if a significant flux was observed over that time, we could rule out pure iron composition, given a standard E_{max} . Olinto et al. [23] shows that using both an E_{max} of 10^{20} eV (around the highest




Figure 5.4: The final flux limit result for this analysis is given by the upper red curve (ARIANNA prototype) at the 90% C.L. (see text for details). The 3-year estimated sensitivity is given by the red dashed curve below. Other experimental limits, including IceCube [43], ANITA-II [37], Auger [51], and RICE [32], are shown as well. Models of GZK neutrino production for proton and iron primaries, and various transition models are adapted from [61].

UHECR energy observed at Earth), and iron-dominated composition makes GZK neutrino observation almost impossible: not even with 10 years of live-time could we observe the corresponding neutrino flux. Models with pure iron composition are not considered as likely as mixed composition models because photo-disintegration of such nuclei inevitably produce protons at high energies. Thus, one must tune down E_{max}^Z to eliminate these protons, which further attenuates the proton flux at lower energies near the GZK cutoff.

Recent limits on the UHE neutrino flux have been obtained by considering the maximum energy density of diffuse gamma-rays observed by the Large Area Telescope (LAT) on the Fermi satellite [22]. In this work, the authors assume a protonic energy spectrum for the UHE cosmic rays, and make various assumptions regarding the extra-galactic background light (EBL) and inter-galactic magnetic field strengths. Arguing that the dip in the UHE cosmic ray energy spectrum before the GZK cutoff could be caused by lepton pair production (the Bethe-Heitler process), they vary the energy at which this effect begins to become important. Specifically, they vary the energy where the transition from galactic to extra-galactic cosmic rays begins, matching to the data from Hi-Res I and Hi-Res II experiments. The gamma-rays from neutral pion decays from GZK interactions, and their subsequent electromagnetic cascades on the EBL and IGM are constrained to fit the data from Fermi-LAT. For large cross-over energies (10^{19} eV) the upper red line of the GZK proton region in figure 5.4 is not in strong conflict with gamma-ray data, as it is within the 99% C.L. for the goodness-of-fit (GOF) from [22]. However the upper red curve (corresponding to FR II-type galactic source evolution) is only marginally compatible with the 99% C.L. GOF for a lower crossover energy of $10^{17.5}$ eV. Further, FR II-type galaxies exhibit strong ($\propto (1+z)^4$) evolution out to redshifts of 2. Such an exponent is also constrained by considering the maximum energy density of the diffuse gamma rays.

There are several reasons why the scientific result (the upper red curve) in figure 5.4 is located above the limits from other experiments. First, the live-time of the upper result in

Model and Reference	Model Class	Predicted N_ν	MRF ($2.3/N_\nu$)
ESS Fig. 4 ($\nu_e + \nu_\mu$) [70]	No source evo.	30.8	0.0746
Kotera (2010) Fig. 1 [33]	SFR1, Pure Proton	37.1	0.0621
ESS Fig. 9 [70]	Strong evo.	104.9	0.0219
Kalashv Fig. 2 [68]	High E_{max} , $z \leq 2$	96.1	0.0239
Barger Fig. 2 [42]	Strong evo.	114.9	0.0200
Yuksel, Kistler (2007) [53]	SFR evo.	45.4	0.0506
Yuksel, Kistler (2007) [53]	QSO evo.	55.5	0.0414
Yuksel, Kistler (2007) [53]	GRB evo.	156.1	0.0147
Ave et al. (2005) [24]	Pure Fe comp.	11.3	0.204
Todor Stanev [79]	Fe, CMB+IRB	2.40	0.956
Kotera Fig. 7 upper [33]	Mixed comp.	21.7	0.106
Kotera Fig. 7 lower [33]	Pure Fe	7.50	0.307
Fermi-LAT [22]	$E_{cross} = 10^{17.5}$ eV	15.5	0.148
Fermi-LAT [22]	$E_{cross} = 10^{18.0}$ eV	21.1	0.109
Fermi-LAT [22]	$E_{cross} = 10^{18.5}$ eV	32.9	0.0699
Fermi-LAT [22]	$E_{cross} = 10^{19.0}$ eV	42.8	0.0537
WB (1999) [17]	No source evo.	22.4	0.103
WB (1999) [17]	QSO evo.	67.1	0.0343
Olinto review (2011) [23]	Fe, $E_{max} = 100$ EeV	0.14	16
Olinto review (2011)	Mixed, $E_{max} = 10$ EeV	0.068	33
Olinto review (2011)	Proton, $E_{max} = 3$ ZeV	101.3	0.0227
Olinto review (2011)	Various protonic, SFR	37.1	0.0621

Table 5.2: The predicted number of events from various models (integrated over all neutrino flavors) over 3.0 years of live-time with a 31×31 ARIANNA array, along with the corresponding MRF ($2.3/N_\nu$). Given enough live-time and a full array of ARIANNA stations, the instrument will be sensitive enough to detect or reject almost all of the models shown, assuming no backgrounds (MRF < 1).

figure 5.4 is 90.41 days, including the removal of the 4.9 day run 57 of 2011-12. In the future, reductions in consumed power will drive the live-time to at least 50% of the year, with more live-time coming potentially from wind power. Additionally, we will rarely have the need to halt data taking to adjust thresholds, remove man-made noise sources (such as the wi-fi tower), or fix glitches in communications. Techniques learned from operating the prototype station allow us to achieve most of the live-time made available to us by the sun and wind. Secondly, the analysis efficiency of the waveform analysis is 50%, unfortunately, because of the strict amplitude cut. While this cut gets rid of a large fraction of persistent thermal backgrounds, we would like to keep more signal. On the other hand, the planarity variable requires a good signal in four channels to reconstruct the timing well enough to establish the existence of a plane wave. Future ARIANNA stations will employ a trigger which acts in the time domain, allowing for a signal threshold that is effectively lower than the ARIANNA prototype, but produces fewer overall thermal triggers [46]. The purpose of employing this trigger is to boost efficiency. In the absence of the shadowing effect, we could also gain a factor of 2 in sensitivity, but it is more likely that the increased effective volume due to weak shadowing will be between 1 and 2. Given these improvements, and the deployment of an increasing number of stations each year, the experiment will ultimately reach the required sensitivity for GZK neutrino detection.

Chapter 6

Conclusion

The prototype for the ARIANNA station, utilizing data acquisition architecture first developed for the ANITA balloon-borne high energy neutrino detector, was initially designed to explore the long-term ambient RF (especially impulsive) conditions at a potential site relatively close to McMurdo Station. The prototype station known as Icicle1, operating from December 10th, 2009 to March 1st, 2012 fulfilled this goal and much more. A measurement of the energy spectrum, direction, and flavor composition of cosmogenic neutrinos would provide significant constraints on the nature and distribution of cosmic ray accelerators in the universe. In addition, the angular distribution near the horizon provides a mechanism to measure the cross-section at energies not yet available at terrestrial accelerators.

The cross-section and energy spectrum provide unique clues in the search for physics beyond the standard model. Through simulations, detector design, power system design, Antarctic expeditions, and data analysis, we have produced in this work an upper limit on the flux of UHE neutrinos. By scaling this limit by the number of eventual remote detectors comprising the full ARIANNA array, and by scaling additionally by modest increases in live-time, this dissertation has provided evidence that ARIANNA concept, based on data collection by

remotely operated, autonomous stations located on the Ross Ice Shelf, is credible and viable.

Several milestones were reported, including autonomous data transfer and archiving through high speed wireless network links during the summer months and Iridium satellite communications when sunlight was insufficient to provide power to the wireless network. The prototype station shutdown safely during the Antarctic winter and restarted automatically when there was sufficient sunlight in Austral spring, thus demonstrating at least 50% annual live-time. Wind power can only add to the annual live-time, and the current versions of the remote detectors are assessing this possibility. The prototype station has also demonstrated robust operation across three years of deployment, standing up to the harsh Antarctic environment.

The properties of the ice that forms the fiducial volume of the experiment have been shown to be serendipitously favorable for such an experiment. The shelf depth of 576 ± 2 m is comparable to attenuation lengths we have from the linear fit to the temperature-averaged radio frequency attenuation length: $\langle L \rangle = (500 \pm 30 - (0.18 \pm 0.05)\nu[\text{MHz}])$ meters in the range [200-750] MHz. These results are consistent with the lower frequency result from 2010: 490 ± 20 meters for [80-180] MHz. It is important to note that because the shelf depth and attenuation lengths are comparable, RF neutrino signals will experience minimal absorptive losses. The reflection coefficient was found to be $0.72 \leq \sqrt{R} \leq 0.88$ at 68% C.L., for [80-180] MHz in studies performed during the field expedition in summer of 2010, and $0.70 \leq \sqrt{R} \leq 1.0$ at 68% C.L. for [325-750] MHz from a similar expedition in 2011, using a shorter baseline to reduce the path length of the bounce studies. This short baseline also revealed hints of surface effects that could boost the effective aperture of the final detector by a factor of ≈ 2 .

Using the shorter baseline in 2011, we rotated the the transmitting and receiving antenna to directly face each other and inserted them in the snow to investigate RF signal propagation for horizontal trajectories near the surface. These studies revealed hints of non-simulated

surface physics that has the potential to boost the effective aperture by as much as a factor of 2. Our naive expectation for a continuously increasing gradient of the index of refraction as a function of depth in the firn ice suggest that little RF power should propagate horizontally, but rather, the beam should arc downward. However, power was observed in our experiment.

Finally, anthropogenic RF backgrounds were identified and removed (initially, by analysis cuts and later by corrective measures to reduce the RF emission from the wireless tower power systems). We have no evidence that non-thermal, non-local RF emission was produced during this study, and this important result that suggests the ARIANNA site is suitable for expansion. The obvious and unique signatures of locally produced background events were used to remove them from the data set. The analysis of the thermal fluctuations that generated events revealed no neutrino signal candidates, leading to a flux limit of $E^2\phi < 10^{-5}$ GeV cm⁻² s⁻¹ sr⁻¹, assuming a E^{-2} differential energy spectrum.

The sensitivity to UHE neutrino flux models can be scaled by the number of future stations and live-time to achieve sensitivities which are sufficient to detect between 10-100 neutrinos per year, probing a wide range of parameters in models of cosmogenic neutrino production. These models include models that explain current UHECR data both as pure protons and protons mixed with heavier charged nuclei, unless the density of sources of the highest energy cosmic rays exhibits no redshift dependence. The only models not probed by 3.0 years of operation require pure iron or other heavy nuclei-dominated compositions, and cutoff energies comparable to the just the highest energy cosmic ray observed on Earth, with no redshift dependence. Therefore, non-observation of neutrino events by ARIANNA will significantly constrain the models of cosmic ray acceleration.

The design, deployment, calibration and operation of the prototype station has provided us with valuable knowledge and experience that will help improve the design and architecture of future station arrays. In 2011, the National Science Foundation agreed to support the development of a seven-station pilot program known as the Hexagonal Radio Array. The

first three of seven stations were deployed in December, 2012 using technology specifically optimized for ARIANNA. The new technology incorporates advance time-domain triggering and pattern recognition, a factor of 3 reduction in power consumption (boosting live-time), and advanced networking and control software. Many of these design changes are based on the experiences garnered from the ARIANNA prototype station [81]. Going forward, the ARIANNA experiment has demonstrated the possibility of using the Ross Ice Shelf to detect and study some of the fastest, rarest, and most energetic particles in the cosmos.

Bibliography

- [1] A. J. Bunker, E. R. Stanway, R. S. Ellis, and R. G. McMahon. The star formation rate of the universe at $z=6$ from the hubble ultra-deep field. *arXiv:astro-ph/0403223*, 2004.
- [2] Dennis Allard. Extragalactic propagation of ultrahigh energy cosmic-rays. *Astroparticle physics*, 39-40, 2012.
- [3] Amanda Cooper-Sarkar et al. The high energy neutrino cross-section in the Standard Model and its uncertainty. *Journal of High Energy Physics*, 42, 2011.
- [4] De Angelis. Atmospheric ionization and cosmic rays: studies and measurements before 1912. In *100 Years Cosmic Ray Physics - Anniversary of the V.F. Hess Discovery*, 2012.
- [5] Gurgen Askaryan. Coherent radio emission from cosmic showers in air and in dense media. *Soviet Journal of Experimental and Theoretical Physics*, 21:658, 1965.
- [6] The Auger Collaboration. Update on the correlation of the highest energy cosmic rays with nearby extragalactic matter. *Astroparticle Physics*, 34:314–326, 2010.
- [7] John Bahcall. *Neutrino Astrophysics*. Cambridge University Press, 1989.
- [8] C. Balanis. *Antenna Theory: Analysis and Design (3rd ed.)*. John Wiley and sons., 2005.
- [9] Carl Baum. General properties of antennas. *IEEE Transactions on Electromagnetic Compatibility*, 23(1):18–24, 2002.
- [10] V. Berezhinsky and G. Zatsepin. Cosmic rays at ultra high energies (neutrino?). *Physics Letters B*, 28:423–424, 1969.
- [11] Eric Berg and Jordan Hanson. Private communication.
- [12] Dave Besson and Daniel Kennedy. Private communication.
- [13] CReSIS. (Center for Remote Sensing of Ice Sheets). <https://www.cresis.ku.edu/news/cresis-in-the-news>.
- [14] Guillaume Decerprit and Dennis Allard. Constraints on the origin of ultra-high-energy cosmic rays from cosmogenic neutrinos and photons. *Astronomy and Astrophysics*, 17673, 2011.

- [15] Kamlesh Dookayka. *Characterizing the Search for Ultra-High Energy Neutrinos with the ARIANNA Detector*. PhD thesis, University of California at Irvine, 2011.
- [16] J. Dowdeswell and S. Evans. Investigations of the form and flow of ice sheets and glaciers using radio-echo sounding. *Reports on Progress in Physics*, 67:1821–1861, 2004.
- [17] Eli Waxman and John Bahcall. High-energy neutrinos from astrophysical sources: An upper bound. *Physical Review Letters D*, 59:023002, 1999.
- [18] Abbasabadi et al. High energy photon-neutrino elastic scattering. *Physical Review D*, 63(9):093001, 2001.
- [19] Abdo et al. A measurement of the spatial distribution of diffuse tev gamma-ray emission from the galactic plane with milagro. *Astrophysical Journal*, 688:1078–1083, 2008.
- [20] Abdo et al. Fermi large area telescope constraints on the gamma ray opacity of the universe. *Astrophysical Journal*, 723:1082–1096, 2010.
- [21] Aharonian et al. A new population of very high energy gamma-ray sources in the milky way. *Science*, 307(5717):1938–1942, 2005.
- [22] Ahlers et al. GZK neutrinos after the FERMI-LAT diffuse photon flux measurement. *Astroparticle Physics*, 34:106–115, 2010.
- [23] Angela Olinto et al. Ultrahigh energy cosmic rays and neutrinos. *Nuclear Physics B*, 217(1):231–236, 2011.
- [24] Ave et al. Cosmogenic neutrinos from ultra-high energy nuclei. *Astroparticle Physics*, 23(1):19–29, 2005.
- [25] Buford Price et al. Temperature Profile for Glacial Ice at the South Pole: Implications for life in a nearby subglacial lake. *Proceedings of the National Academy of Sciences*, 99(12), 2002.
- [26] Buitink et al. Constraints on the flux of ultra-high energy neutrinos from wsrt observations. *Astronomy and Astrophysics*, 14104, 2010.
- [27] D. Besson et al. *In situ* radioglaciological measurements near Taylor Dome, Antarctica and implications for ultra-high energy (UHE) neutrino astronomy. *Astroparticle Physics*, 29(2):130–157, 2008.
- [28] G. Fogli et al. Global analysis of neutrino masses, mixings and phases: entering the era of leptonic CP-violation searches. *Physics Review D*, 86(1):013012, 2012.
- [29] Gary Varner et al. The large analog bandwidth recorder and digitizer with ordered readout (LABRADOR) ASIC. *Nuclear Instrumentation and Methods A*, 583(2-3):447–460, 2007.

- [30] Gorham et al. The Antarctic Impulsive Transient Antenna Ultra-high Energy Neutrino Detector (ANITA): Design, performance, and sensitivity for 2006-2007 balloon flight. *Astroparticle Physics*, 32:10–41, 2009.
- [31] Hanson et al. A prototype station for ARIANNA: A detector for cosmic neutrinos. *Nuclear Instruments and Methods in Physics Research A*, 624(1):85–91, 2010.
- [32] I. Kravchenko et al. Updated results from the rice experiment and future prospects for ultra-high energy neutrino detection at the south pole. *Physical Review D*, 85(6), 2012.
- [33] K. Kotera et al. Cosmogenic neutrinos: parameter space and detectability from PeV to ZeV. *Journal of Cosmology and Astroparticle Physics*, 2010.
- [34] Miocinovic et al. Time-domain measurement of broadband coherent cherenkov radiation. *Physical Review Letters D*, 74:043002, 2006.
- [35] Norman Beaulieu et al. An FFT Method for Generating Bandlimited Gaussian Noise Variates. *European Transactions on Telecommunications*, 10(5):545–550, 1999.
- [36] P. Allison et al. Design and initial performance of the Askaryan Radio Array prototype eev neutrino detector at the south pole. *Astroparticle Physics*, 35:457–477, 2012.
- [37] P. Gorham et al. Observational constraints on the ultrahigh energy cosmic neutrino flux from the second flight of the anita experiment. *Physical Review D*, 82:022004, 2010.
- [38] Saltzberg et al. Observation of the Askaryan Effect: Coherent Microwave Cherenkov Emission from Charge Asymmetry in High Energy Particle Cascades. *Physical Review Letters*, 86:2802–2805, 2001.
- [39] Shlivinski et al. Antenna characterization in the time domain. *IEEE Transactions on Antennas and Propagation*, 45(7):1140–1149, 1997.
- [40] T. Barrella et al. Ross ice shelf in situ radio-frequency ice attenuation. *Journal of Glaciology*, 57:61–66, 2011.
- [41] T. Matsuoka et al. Effect of temperature on dielectric properties of ice in the range 539 GHz,. *Journal of Applied Physics*, 80:5884, 1996.
- [42] V. Barger et al. Ultra high energy neutrino nucleon cross section from cosmic ray experiments and neutrino telescopes. *Physics Letters B*, 642(4):333–341, 2006.
- [43] R. Abbasi et al. [IceCube Coll.]. Measurement of the atmospheric neutrino energy spectrum from 100 gev to 400 tev with icecube. *Physical Review D*, 83:012001, 2011.
- [44] S. Smith M. Ewen S. Evans. Radio echo sounding: absorption and scattering by water inclusion and ice lenses. *Journal of Glaciology*, 11(61):133–146, 1972.
- [45] Stefan Fliescher. *Antenna Devices and Measurement of Radio Emission from Cosmic Ray induced Air Showers at the Pierre Auger Observatory*. PhD thesis, Aachen University, 2011.

- [46] Stuart Kleinfelder for the ARIANNA Collaboration. Design and performance of the autonomous data acquisition system for the arianna high energy neutrino detector. *Proceedings of the Symposium on Radiation Measurements and Applications (SORMA) (peer-reviewed)*, 14B(6), 2012.
- [47] G. Feldman and R. Cousins. An FFT Method for Generating Bandlimited Gaussian Noise Variates. *European Transactions on Telecommunications*, 10(5):545–550, 1999.
- [48] Jordan Goodman and James Braun. The HAWC observatory. In *Proceedings of the 32nd International Cosmic Ray Conference*, 2011.
- [49] K. Greisen. End to the cosmic ray spectrum? *Physical Review Letters*, 16:748–750, 1966.
- [50] Claus Grupen. *Astroparticle Physics*. Springer, 2005.
- [51] Yann Guardincerri. The pierre auger observatory and ultra-high energy neutrinos: upper limits to the diffuse and point source fluxes. In *Proceedings of the 2011 International Cosmic Ray Conference, Beijing*, 2011.
- [52] Bogorodsky Bentley Gudmandsen. *Radioglaciology*. Reidel Publishing Co. (The Netherlands), 1985.
- [53] H. Yuksel and M.D. Kistler. Enhanced cosmological GRB rates and implications for cosmogenic neutrinos. *Physical Review D*, 75(8):083004, 2007.
- [54] Jordan Hanson. Ross ice shelf thickness, radio-frequency attenuation and reflectivity: Implications for the arianna uhe neutrino detector. In *Proceedings of the 32nd International Cosmic Ray Conference*, 2011.
- [55] Victor Hess. Uber beobachtungen der durchdringenden strahlung bei sieben freiballonfahrten. *Physikalische Zeitschrift*, 13:1084–1091, 1912.
- [56] Gary Hill and Katherine Rawlins. Unbiased cut selection for optimal upper limits in neutrino detectors: the model rejection potential technique. *Astroparticle Physics*, 19, 2003.
- [57] Horowitz and Hill. *The Art of Electronics*. Cambridge University Press, 1989.
- [58] John David Jackson. *Classical Electrodynamics 3rd ed.* John Wiley and sons., 1999.
- [59] Jaime Alvarez-Muniz, Andres Romero-Wolf, Enrique Zas. Cerenkov radio pulses from electromagnetic showers in the time-domain. *Physical Review D*, 81:123009, 2010.
- [60] Jaime Alvarez-Muniz, Andres Romero-Wolf, Enrique Zas. Practical and accurate calculations of askaryan radiation. *Physical Review D*, 84:103003, 2011.
- [61] Karl-Heinz Kampert and Michael Unger. Measurements of the cosmic ray composition with air shower experiments. *Astroparticle Physics*, 35(10):660–678, 2012.

- [62] Spencer Klein. The LPM effect: Comparing SLAC E-146 data with experiment. In *Proceedings of the Fourth Paris Workshop on Quantum Chromodynamics*, 1998.
- [63] John Kraus and Ronald Marhefka. *Antennas 3rd ed.* McGraw-Hill, 2003.
- [64] Henrik Madsen. *Time Series Analysis*. Taylor and Francis Group LLC, 2008.
- [65] Martin Laboka. Arctic Peak. <http://www.arcticpeak.com/antennapages/LPDA.htm>.
- [66] C.S. Neal. The dynamics of the ross ice shelf revealed by radio echo sounding. *Journal of Glaciology*, 24(90):295–307, 1979.
- [67] C.S. Neal. Radio echo determination of basal roughness. *Annals of Glaciology*, 3, 1982.
- [68] O. E. Kalashev, V. A. Kuzmin, D. V. Semikoz, and G. Sigl. Ultrahigh-energy neutrino fluxes and their constraints. *Physical Review D*, 66(6):063004, 2002.
- [69] Press, Teukolsky, Vetterling, and Flannery. *Numerical Recipes in C 2nd ed.* Cambridge University Press, 1992.
- [70] Ralph Engel, David Seckel, and Todor Stanev. Neutrinos from propagation of ultra-high energy protons. *Physical Review D*, 64:093010, 2001.
- [71] J. Ralston. Radio surf in polar ice: A new method of ultrahigh energy neutrino detection. *Physics Review D*, 71:011503(R), 2005.
- [72] S.O. Rice. Mathematical analysis of random noise. *Bell System Technical Journal*, 23:282–332, 1944.
- [73] David Saltzberg. private communication.
- [74] V. Schytt. (Norwegian-British-Swedish Antarctic Expedition). *Scientific results Norsk Polarinstitutt 194952*, pages 113–151, 1958.
- [75] Sean Scully and Floyd Stecker. Searching for new physics with ultrahigh energy cosmic rays. *New Journal of Physics*, 11:085003., 2009.
- [76] Sean Scully and Floyd Stecker. Testing Lorentz invariance with neutrinos from ultrahigh energy cosmic ray interactions. *Astroparticle Physics*, 34:575–580, 2011.
- [77] Sergei Gninenko. Resolution of puzzles from the LSND, KARMEN, and MiniBooNE experiments. *Physical Review D*, 83:015015, 2011.
- [78] W. Sorgel and W. Weisbeck. Influence of the antennas on the ultra-wide band transmission. *EURASIP Journal on the Applied Signal Processing*, 3:296–305, 2005.
- [79] Todor Stanev. Ultrahigh energy cosmic rays and neutrinos. *Nuclear Instrumentation and Methods A*, 588:215–220, 2008.
- [80] Todor Stanev. *High Energy Cosmic Rays - 2nd ed.* Springer, 2010.

- [81] Joulieu Tatar. *in preparation*. PhD thesis, University of California at Irvine, 2013.
- [82] The ANITA Collaboration et al. Observations of the askaryan effect in ice. *Physical Review Letters*, 99:171101, 2007.
- [83] The Telescope Array Collaboration. The cosmic ray energy spectrum observed with the surface detector of the telescope array experiment. *arXiv1205.5067v1*, 2012.
- [84] Eli Waxman. Cosmological origin for cosmic rays above 10^{19} eV. *Astrophysical Journal*, 452, 1995.
- [85] P. Welch. The use of fast fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms. *IEEE Transactions on Audio and Electroacoustics*, AU-15(2):70–73, 1967.
- [86] G. Zatsepin and V. Zuz'min. Upper limit of the spectrum of cosmic rays. *Soviet Journal of Experimental and Theoretical Physics Letters*, 4:78, 1966.
- [87] Zhi-zhong Xing. A further study of μ - τ symmetry breaking at neutrino telescopes after the Daya Bay and RENO measurements of θ_{13} . *Physics Letters B*, 716(1):220–224, 2012.

Appendices

A Selected Antenna derivations

There are several excellent descriptions of electrodynamic theory applied to varieties of antenna shapes: [63] [8], and [45] is a good example of Here we provide derivations of certain general properties of antennas used in the text. There are two main derivations here. The first obtains the expression for the effective height used in the Monte Carlo simulations. This operator dictates the maximum voltage fed into the data acquisition system by the log-periodic dipole array (LPDA) signal antennas. Next, we shall derive the experimental equation used to measure the effective height property of the LPDA antenna in section 3.2.2

As a first step, we must indicate how the gain of an antenna is related fundamentally to the beam area, or beam solid angle:

$$G = k \frac{4\pi}{\Omega_A} \tag{A.1}$$

Here, k is a unit-less efficiency factor that converts the *directivity* to the gain: $G = kD$. We almost always take $k \approx 1$. The directivity D is defined simply as the maximum power radiated at a given direction in the far-field of the antenna divided by the average power

radiated over the unit sphere. Using $P(\theta, \phi)$ for the radiation pattern, we can express this as

$$D = \frac{P(\theta, \phi)_{max}}{P(\theta, \phi)_{ave}} = \frac{P(\theta, \phi)_{max}}{(1/4\pi) \iint_{\Omega} P(\theta, \phi) d\Omega} \quad (\text{A.2})$$

We can rearrange the right hand side to obtain the normalized power pattern in the denominator:

$$D = \frac{4\pi}{\iint_{\Omega} P(\theta, \phi)/P(\theta, \phi)_{max} d\Omega} = \frac{4\pi}{\iint_{\Omega} P_n(\theta, \phi) d\Omega} \quad (\text{A.3})$$

The beam solid angle Ω_A is the total solid angle into which the antenna radiates: $\Omega_A = \iint_{\Omega} P_n d\Omega$. Thus the denominator above can be replaced by Ω_A and we have

$$D = \frac{4\pi}{\Omega_A} \quad (\text{A.4})$$

The gain G of an antenna with antenna impedance Z is related to the *effective height* of an antenna through the following equation, at a frequency f :

$$h(f) = 2\sqrt{\frac{Gc^2 Z}{4\pi f^2 Z_0}} \quad (\text{A.5})$$

We can derive this equation using the average intensity (or Poynting vector) of an antenna at single frequency and angle:

$$S = \frac{1}{2} \frac{|E|^2}{Z_0} \quad (\text{A.6})$$

Here, $|E|$ is the magnitude of the electric field at a the given angle and frequency, and $Z_0 \approx 120\pi$ is the impedance in Ohms of free space. The beam solid angle is related to an actual area through the wavelength: $A_e \Omega_A = \lambda^2$. The quantity A_e is defined as the antenna aperture. The power delivered to the antenna terminal by the intensity S is

$$P_{term} = S A_e \quad (\text{A.7})$$

Combining the expressions relating gain to beam area, beam area to antenna aperture, and S in terms of the electric field, we have

$$P_{term} = \frac{1}{2} \frac{|E|^2}{Z_0} \frac{\lambda^2}{\Omega_A} = \frac{1}{8\pi} \frac{|E|^2}{Z_0} G \lambda^2 \quad (\text{A.8})$$

The simplest definition of the effective height is that it is the length that multiplies the incident electric field, converting it to a voltage: $V_{term,f} = E(f)h(f)$. Below we will substitute E with V_{term}/h . The terminal power is also related to the open circuit antenna voltage and the antenna impedance (radiation resistance) in the following way, for the ideal condition of

conjugate matching [8] [45]:

$$P_{term} = |V_{term,OC}|^2 / 8Re(Z_A) \quad (\text{A.9})$$

Conjugate matching assumes that the antenna has no losses associated with impedance mismatches, because the antenna load is matched to the radiation resistance. As long as $V_{term,OC} = Eh$, we can combine equation A.9 with A.8 and solve for $h(f)$:

$$h(f) = 2\sqrt{\frac{Gc^2 Re(Z_A)}{4\pi f^2 Z_0}} \quad (\text{A.10})$$

The statement $V_{term,OC} = Eh$ is true for the case of co-polarized radiation if the antenna is linearly polarized. More generally, if θ is the angle between incident linearly polarized radiation and the polarization vector of the E-plane of the antenna, then we have $V_{term,OC} = \vec{E} \cdot \vec{h} = |E||h| \cos(\theta)$. It is equation A.10 that we employ in the Monte Carlo simulations to predict the voltage received and passed through the data acquisition systems by the LPDAs.

Next, we will explain how the magnitude of the effective height can be measured in the Fourier domain. The following derivation is a simplified version of a calculation done in [34]. The measurement of the effective height in section 3.2.2 employs this derivation. Suppose we have two copies of an antenna co-polarized and separated by a distance r . The voltage recorded on the line out of the receiving antenna V_{rec} from the voltage V_{src} radiated through

the transmitting antenna is

$$V_{rec}(t) = \frac{1}{2\pi rc} h_N(t) \circ h_N(t) \circ \frac{dV_{src}}{dt} \quad (\text{A.11})$$

Here, we assume that the antenna impedance Z_a is matched to the line impedance Z , and have taken $h_N(t) = \sqrt{Z_0/Z}h(t)$ to simplify equations. The two convolutions (the operator is \circ) of the signal by the effective height come from the transmitter and receiver, respectively. The time derivative operation comes from just the transmitter, as we have $h_{tx}(t) = 2\partial_t h_{rx}(t)$ [39] [9]. The convolution theorem states that convolution in the time domain is multiplication in the Fourier domain. Taking the Fourier transform of both sides, we have

$$\mathcal{F}\{V_{rec}\} = \frac{1}{2\pi rc} h_N^2(\nu) \mathcal{F}\{\partial_t V_{src}\} \quad (\text{A.12})$$

We have taken advantage of the associative property of multiplication in the Fourier domain on the right hand side. Using the definition of the Fourier transform with standard frequency (rather than angular frequency), we have $\mathcal{F}\{\partial_t g(t)\}_\nu = 2\pi i\nu \mathcal{F}\{g(t)\}_\nu$. Letting $V(\nu)$ represent the voltage signals in the frequency domain, and switching back to the un-normalized version of h , equation A.12 simplifies to

$$V_{rec}(\nu) = \frac{i\nu}{rc} h(\nu)^2 \frac{Z_0}{Z} V_{src}(\nu) \quad (\text{A.13})$$

Solving for h , we have

$$h(\nu) = \sqrt{\frac{ZrcV_{rec}(\nu)}{i\nu Z_0 V_{src}(\nu)}} \quad (\text{A.14})$$

This is the equation we use to calculate the effective height of the LPDA in section 3.2.2. This expression must be modified if the antenna is immersed in a medium with a uniform index of refraction n . However, in section 2.2.2 we have shown that the voltage standing wave ratio (VSWR) is very similar between air and buried in firn snow. In either case, the VSWR is typically less than 2, an indication that the antenna impedance is not being affected by the surrounding environment and that almost all power sent to the antenna is radiated. Ultimately the theoretical expression for the effective height is fed into the Monte Carlo simulations predicting the maximum voltage in the data acquisition system due to an Askaryan pulse.

B Information Regarding the *shelfmc* Monte Carlo software

This is a very brief description of how to modify and operate the *shelfmc* Monte Carlo software package. The Monte Carlo simulation takes advantage of ROOT libraries and structures, as part of the ROOT software package, ubiquitous in particle physics. ROOT can be downloaded and installed from <http://root.cern.ch/drupal/>. To install ROOT, one simply has to follow these instructions: <http://root.cern.ch/drupal/content/installing-root-source>. Installation scripts in the Makefile format are provided by those who maintain the ROOT package.

B.1 Compiling and running *shelfmc*

Similar to the ROOT packages and libraries, *shelfmc* is compiled by executing Makefile scripts. Entering "make clean" at the command prompt removes older compiled code, after which "make" compiles the *shelfmc* executable. The file "declaration.hh" is a header file controlling certain physical parameters in the software. For example, physical constants such as the speed of light, and instrumentation parameters like bandwidth are all contained in this file.

B.2 The steering file - input.txt

The steering file, "input.txt," controls several key station properties, ice properties, and neutrino properties. This file is especially useful for performing studies where one or more parameters is varied with respect to energy. This file has been reproduced below, with some example parameter values.

Inputs for ARIANNA simulation, do not change order

- 200000 NNU - The number of neutrinos thrown in the Monte Carlo
- 17.5 EXPONENT - The energy (exponent) for all neutrinos. This is overridden if SPECTRUM is 1.
- 1000 ATGap - The spacing between stations for a grid array of stations, in meters.
- 4 ST TYPE - Do not change.
- 4 N Ant perST - Antennas per station.
- 2 N Ant Trigger - Number of antennas required to be above threshold.

- 30 Z for ST TYPE 2 - Outdated parameter.
- 1 FIRN - If 1, the effects of the firn are taken into account.
- 1.30 NFIRN - The index of refraction of the firn.
- 60 FIRNDEPTH - Depth of the firn in meters.
- 3 NROWS 12 initially, set to 3 for HEXAGONAL
- 5 NCOLS 12 initially, set to 5 for HEXAGONAL
- 0 SCATTER - Do not change.
- 1 SCATTER WIDTH - Do not change.
- 0 SPECTRUM - Use 1 here if a neutrino energy spectrum is desired, rather than a single energy
- 0 DIPOLE - Outdated parameter.
- 0 CONST ATTENLENGTH - Allows precise control of the ice attenuation length for RF frequencies.
- 500 ATTEN UP - Average attenuation length
- 250 ATTEN DOWN - Average attenuation length (one-way)
- 5 NSIGMA - The signal to noise ratio required in N Ant Trigger stations to trigger an event.
- 1. ATTEN FACTOR - Change in the attenuation length.
- 1.0 REFLECT RATE - The reflection coefficient at the ice/ocean interface.
- 1 GZK - The type of neutrino spectrum

- 0 FANFLUX - The type of neutrino spectrum
- 1 WIDESPECTRUM - The type of neutrino spectrum
- 1 SHADOWING - Accounts for the shadowing effect
- 1 DEPTH DEPENDENT N - Depth dependent firm (density, index)

B.3 The declarations file - declarations.hh

The header file `declarations.hh` contains other important physical parameters of the stations, such as bandwidth and filtering, and the radius of the polygon of antennas that comprise a single station's array. The station bandwidth can be edited using the parameters `FREQ_LOW` and `FREQ_HIGH`, where the units are in MHz. To model the effect of the filters in `Icicle1`, `FREQ_LOW` was raised from 100 MHz to 120 MHz to reflect the 3 dB point of the filters. The parameter `ST4_R` governs the antenna radius. This parameter can be varied from 1-100 m with little effect on the effective aperture of the detector. Finally, because `declarations.hh` is a header file, adjusting these parameters requires re-compilation of the `shelfmc` software package.