

THE UNIVERSITY OF CHICAGO

RADIO DETECTION OF ULTRA-HIGH ENERGY NEUTRINOS

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY

ANDREW LUDWIG

CHICAGO, ILLINOIS

JUNE 2019

Copyright © 2019 by Andrew Ludwig
All Rights Reserved

For my grandfather, Francis Ludwig, who inspired me to be a scientist.

TABLE OF CONTENTS

LIST OF FIGURES	viii
LIST OF TABLESxxviii
ACKNOWLEDGMENTS	xxx
ABSTRACT	xxxii
1 INTRODUCTION	1
1.1 Neutrino Physics and Astronomy	1
1.1.1 Solar Neutrinos	1
1.1.2 Neutrino Oscillations	2
1.1.3 Neutrino Flux Measurements	3
1.1.4 Neutrino Source Observations	3
1.2 UHE Neutrino Astronomy	4
1.2.1 Multi-messenger Observations	5
1.2.2 Fundamental Physics Possibilities	6
1.3 Cosmic Ray Astronomy	7
1.3.1 Cosmic Ray Spectrum	7
1.3.2 The GZK Effect	12
1.3.3 Production of UHE Cosmic Rays and Neutrinos	15
1.4 Radio Detection of Neutrinos and Cosmic Rays	15
1.4.1 Askaryan Radiation	15
1.4.2 Extensive Air Showers	19
1.4.3 Geomagnetic Radiation	20
1.5 Introduction to the ANtarctic Impulsive Transient Antenna (ANITA)	21
1.5.1 Previous Results	24
1.5.2 ANITA-IV	24
1.5.3 Complementary Experiments	25
2 ANITA INSTRUMENT AND FLIGHT	26
2.1 Antennas	26
2.2 Signal Processing	29
2.2.1 AMPAs	30
2.2.2 IRFCM and TUFFs	31
2.3 Trigger Path	33
2.3.1 Hybrids	33
2.3.2 L0 Trigger	33
2.3.3 L1 Trigger	34
2.3.4 L2 Trigger	35
2.3.5 Global Trigger	35
2.3.6 Minimum-bias Triggers	35

2.4	Digitizer Path	36
2.5	Orientation	36
2.5.1	GPS	37
2.5.2	Sun Sensors	37
2.5.3	Magnetometer	38
2.6	Flight Computer	38
2.6.1	Data Storage	38
2.6.2	Telemetry	39
2.6.3	Prioritizer	39
2.6.4	Ground Based Monitoring	40
2.7	Balloon and Flight	41
3	CALIBRATION	44
3.1	LABRADOR Timing Calibration	45
3.1.1	Sample-to-sample Timing	45
3.1.2	Write Pointer Wrap Around Timing	46
3.1.3	Temperature Correction	46
3.1.4	Channel to Channel Delays	46
3.2	ADC to Voltage	48
3.3	System Impulse Response	50
3.3.1	Signal Chain Responses	50
3.3.2	Antenna Responses	51
3.3.3	TUFF Responses	53
3.4	Calibration Pulses	56
3.4.1	LDB	57
3.4.2	WAIS	57
3.4.3	HiCal	60
3.5	Phase Centers	60
3.5.1	Photogrammetry	61
3.5.2	Optimization Using Calibration Pulses	61
3.5.3	Resolutions	64
3.6	Trigger Efficiency	65
3.7	Relative Hpol and Vpol Delays	68
3.8	Pointing Offset	68
4	ANALYSIS	72
4.1	Software Overview	73
4.1.1	libRootFftwWrapper	73
4.1.2	anitaTreeMaker	73
4.1.3	eventReaderRoot	73
4.1.4	AnitaAnalysisFramework	74
4.1.5	UCorrelator	74
4.1.6	anitaAnalysisTools	75
4.1.7	anitaMagicDisplay	75

4.1.8	<code>anitaEventCorrelator</code>	75
4.1.9	<code>icemc</code>	77
4.2	Data	77
4.2.1	Thermal Samples	77
4.2.2	Pulsar Events	78
4.2.3	Simulated Neutrinos	78
4.3	Blinding	78
4.3.1	Askaryan Neutrino Blinding	79
4.3.2	Extensive Air Shower Blinding	80
4.4	Digital Filtering	80
4.4.1	Sine Subtraction Filtering	80
4.4.2	Channel BH13 Transfer Function	82
4.5	Event Reconstruction	83
4.5.1	Interferometry	84
4.5.2	Tracing to the Source	86
4.6	Coherently Summed Waveform	87
4.7	De-dispersion	87
4.7.1	Picking the Correct Impulse Response	88
4.8	Polarimetry	88
4.9	Impulsivity	90
4.10	Quality Cuts	90
4.11	Thermal Cuts	96
4.12	Clustering Cuts	99
4.12.1	Angular Log-likelihood	101
4.12.2	Setting the Clustering Cut	103
4.13	Cosmic Ray Specific Analysis Techniques	104
4.13.1	Template Analysis	104
4.13.2	Geomagnetically Induced Polarization Angle	105
4.13.3	Linear Polarization Fraction	106
5	RESULTS	107
5.1	Background Estimates	107
5.1.1	General Anthropogenic Background Estimate	108
5.1.2	EAS Channel: Cosmic Rays and Possible Inverted EAS Events	109
5.1.3	Askaryan Neutrinos	114
5.1.4	Alternative Background Estimate	115
5.1.5	Isolated Background Singlets	116
5.2	EAS Channel (Hpol) Unblinding	116
5.2.1	Further Analysis of Candidates	153
5.3	Askaryan Neutrino Channel (Vpol) Unblinding	160
5.3.1	Posterior Analysis of Candidates	160
5.3.2	Diffuse Neutrino Flux Limit	163
5.4	Summary of Analysis Results	167

5.5	Future Outlook	169
6	THERMAL NOISE CORRELATION STUDIES	171
6.1	Introduction	171
6.2	Thermal Noise Correlation Measurements	171
6.3	Thermal Noise Correlation Simulations	174
6.4	Analysis	175
6.5	Results and Comparison	176
	REFERENCES	181

LIST OF FIGURES

1.1	The astrophysical neutrino flux as measured by IceCube. Plot taken from [21].	3
1.2	Much of the universe is unobservable with photons or protons. The hatched blue region is all of energy-distance phase space that is not observable with photons, the hatched red region is all of energy-distance phase space that is not observable with protons. Various astrophysical object distances are also shown. Plot courtesy of Dr. Peter Gorham.	5
1.3	Auger’s measurement of composition of the CR spectrum using the average of the X_{max} distribution as a function of energy. Each dashed colored line represents a pure spectrum of a single element, red is H, grey is He, green is Ne, and blue is Fe. Figure is taken from [38]. TA’s measurement of composition of the CR spectrum using the average of the X_{max} distribution as a function of energy. The red line also represents various models a pure proton composition here, while the pink line is a pure nitrogen composition and the blue line represents a pure iron spectrum. As you can see, the data from TA indicates a spectrum much closer to pure proton than for Auger, however, the discrepancy between the two measurements are within systematic uncertainties [37]. Figure is from [36].	8
1.4	In the top panel is the CR flux as a function of energy-per-nucleus as measured by various experiments. The flux is multiplied by $E^{2.6}$ to more easily see the changes in spectral slope. Labeled with arrows are the knee, 2nd knee and ankle of the spectrum. In the bottom panel is the same flux zoomed in about the highest energy ankle region. These UHE fluxes are only measured by Auger and TA, and their discrepancy is within systematic uncertainties up to $10^{19.4}$ eV, at which point the discrepancy is not fully explained [37]. Figures are from [39].	9

1.5	An updated Hillas diagram taken from [45]. Various sources are shown in $B - R$ phase-space. The phase-space above the blue (red) line is the allowed region for acceleration of protons (iron nuclei) above 10^{21} eV (10^{20} eV).	11
1.6	The fraction of CRs that survive from a distance $> D$ is plotted vs D for H, He, CNO, and Fe at $E > 60$ EeV, and H at $E > 40$ EeV and $E > 100$ EeV. A solid gray line showing the 50% survival point for various nuclei compositions is also plotted. Only protons or iron can survive for distances greater than 50 Mpc at $E > 60$ EeV. Figure taken from [45].	13
1.7	Two models for neutrino flux are shown for different CR composition assumptions. Black is an all proton flux and dashed blue is a mixed composition (between pure proton and pure iron fluxes). The expected neutrino flux scales downwards with increasingly heavy UHECR flux compositions. Figure is taken and slightly modified from [49].	14
1.8	The current leading detections and limits of neutrinos at high energies from IceCube are shown [58, 59, 60], along with various models. Models displayed are for low-luminosity GRBs [53] and AGNs [51], along with an extrapolation of the IceCube flux. All allowed GZK fluxes from UHECR constraints according to [61] are shown in shaded grey. Figure is taken and slightly modified from [62]. . . .	16
1.9	Figure is from [65], the paper detailing the first observation of the Askaryan effect originating in ice. The top panel shows the angular dependence of the recorded electric field for the full 200-800 MHz band of the experiment. The bottom panel is the same as the top panel, but split into smaller frequency bands. High frequency power falls off more quickly at angles away from the Cherenkov cone angle than low frequency power because it loses coherence.	18
1.10	A recorded Askaryan pulse, produced in ice, and recorded by the ANITA-I electronics. Figure from [65].	19

1.11	The radio signal flux density as a function of viewing angle for various frequencies (denoted by color). This distribution is from a simulated $10^{18.4}$ eV CR with a Cherenkov angle at $\psi = 0.7$. Similar to radiation from the Askaryan effect, high frequencies fall off quickly when viewed away from the Cherenkov cone angle. Figure is taken from [75].	21
1.12	Sources ANITA is sensitive to, minus thermal noise. Picture taken from [80].	23
2.1	Response as a function of angle from boresight in the E plane for antennas for various frequencies. Gain is referenced to the maximum gain of the antenna. Data is from Antenna Research Associates, Inc., plot is from [80].	27
2.2	The ANITA-IV payload pre-flight. An additional drop-down ring of PV panels was added to the bottom ring of antennas before flight.	28
2.3	There is a 1.96 cm separation between the horizontally polarized and vertically polarized antenna feeds for ANITA horn antennas.	29
2.4	The ANITA-IV system. See text for information on each component.	30
2.5	Internals of the AMPA. Picture from [98].	31
2.6	The gain in dB vs frequency in Hz (left) and noise temperature in Kelvin vs frequency in Hz (right) for all 98 of the AMPAs. Each AMPA is in a different color, with the thick black line in the middle corresponding to the mean value. Plots taken from [99].	32
2.7	Trigger path assembly. The blue rectangular pieces are the 90° hybrid couplers, below them are the RF-tight boxes housing the SHORTs. A little more than half of the hybrids have been installed, in the end this will be symmetrical. Also pictured are power distribution systems on the right, along with cabling. Picture was taken during assembly in Hawaii and is used courtesy of John Russell.	34

2.8	Payload heading in degrees as calculated by ADU5-A (blue) and sun sensor (red) for one run of flight data, taken from [101]. They track very closely, indicating that the sun sensors would be a viable backup for heading information if the GPS antennas failed.	37
2.9	The AWARE program for monitoring ANITA during flight. Buttons on the left change what information is displayed. The current display shows all 96 waveforms for an example RF-triggered event.	40
2.10	The ANITA-IV payload pictured with inflated balloon during launch. The balloon will inflate to a much greater size as it rises and the pressure drops. During launch, the ANITA payload is held aloft by a truck-mounted crane. Launch must take place during a day with perfect conditions (still winds, clear skies), because the balloon is delicate and incredibly expensive. ANITA-IV had one launch attempt scrubbed due to weather before a successful flight.	41
2.11	ANITA-IV after terminating the flight. Picture taken by and used courtesy of Christian Miki. Crumpled antennas can be seen on the bottom ring of the instrument. This photo was taken during recovery of the data, the remote location of which necessitated use of the pictured plane.	42
2.12	The ANITA-IV flight path. The first pass around the continent is shown in red, the second pass is shown in blue, and the third pass, where we got stuck in a holding pattern and had to terminate the flight, is shown in green. The Antarctic background map is colored by ice thickness.	43
3.1	This colormap shows the mean calculated time delays compared to the reference channel for each channel, broken up by LAB chip. The color axis is the time delay in ns. The X-axis is the SURF and the Y-axis is the channel and chip on that SURF. The reference channel is SURF 6, channel 3. Delays vary by SURF because each SURF is cabled differently.	47

3.2	This colormap shows the correction from ADC counts to mV for the input calibration signal, broken up by LAB chip. The X-axis is the SURF and the Y-axis is the channel and chip on that SURF. Calibration constants vary between 2 and 3.5 mV per ADC count between channels.	49
3.3	An example signal chain impulse response is shown. The Y-axis is unitless. This is the bottom Vpol channel 7 (BV07), but is representative of every channel besides BH13 and TH03.	51
3.4	The average antenna impulse response is shown. Compared with the signal chains, the antennas had much more similar impulse responses.	52
3.5	Shown are the full impulse responses (signal chain convolved with antennas) for all horizontally-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring. TH03 and BH13 are shown in red because they had significantly different impulse responses from the average response. The impulse response differences are caused by different filters along the signal chain.	54
3.6	Shown are the full impulse responses (signal chain convolved with antennas) for all vertically-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring. All of the vertical channels had very similar impulse responses, although one of the channels (BV14) died during flight.	54
3.7	Shown are the power spectra of the full impulse responses (signal chain convolved with antennas) for all horizontally-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring. TH03 and BH13 are in red because they had different impulse responses from the average response.	55

3.8	Shown are the power spectra of the full impulse responses (signal chain convolved with antennas) for all vertically-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring.	55
3.9	Magnitude responses for all seven TUFF configurations that were employed during the flight are shown. Magnitude is in dB and frequency is in Hz. Notch 1 (nominal frequency 260) was almost always on throughout the flight. Notch 2 (nominal frequency 375) was on most of the flight, and had to be tuned. Notch 3 (nominal frequency 460) only had to be turned on near large Antarctic bases such as McMurdo and South Pole Station.	57
3.10	The ANITA-IV flight path is shown along with the locations of the pulsing stations. The pulsing station at LDB is the magenta star and the pulsing station at WAIS is the cyan star. ANITA flew almost directly over WAIS, allowing us to collect a large amount of high-quality calibration data.	58
3.11	A picture of the WAIS pulsing station with all components labeled. The LDB pulser setup was much the same. Picture taken from an ANITA internal note [106].	59
3.12	An example of a vertically polarized WAIS calibration pulse, with the phase response of the instrument removed. It is extremely impulsive, and very linearly polarized.	60
3.13	An example of a HiCal calibration pulse, with the phase response of the instrument removed. It is much less impulsive than WAIS.	61
3.14	3D model of ANITA-IV, as produced by PhotoModeler Scanner 2015 software. Image taken from [109].	62

3.15	The pointing resolution using the photogrammetry numbers as the phase centers. There is a mean offset of 0.61° and a resolution of 0.87° in ϕ and a mean offset of 0.04° and a resolution of 0.24° in θ	66
3.16	The pointing resolution using the phase centers determined from fitting using the WAIS pulses. There is a mean offset of 0.01° and a resolution of 0.42° in ϕ and a mean offset of 0.01° and a resolution of 0.14° in θ	66
3.17	The top panel shows θ resolution as a function of SNR as calculated from WAIS pulses using the calibrated phase centers and the bottom panel is the same thing in ϕ . The color axis for both is simply the number of WAIS pulses. SNR is calculated using the coherently summed waveform, so it is very roughly $\sqrt{15}$ times larger than single antenna SNR.	67
3.18	The top panel shows the trigger efficiency according to the pre-flight scans. The 50% point (indicated by gray lines) is at approximately 4.2 SNR. Vertical error bars are statistical, while horizontal error bars are negligible because SNR was precisely set. Data is from an ANITA internal note [110]. The bottom panel shows the trigger efficiency according to WAIS data. The 50% point (indicated by gray lines) is at approximately 4.1 SNR. Vertical error bars are statistical, while horizontal error bars are indicative of the binning of the data. Note the different X-axis limits. The efficiency as derived from in-flight pulses appears to be slightly better than from data taken pre-flight, but likely due to slightly different definitions of SNR. SNR as defined in the top plot is taken on an oscilloscope directly before the SURF, whereas in the WAIS data it is calculated from what is recorded by the SURF.	69
3.19	The offset in pointing in ϕ as a function of incoming pulse polarization angle. Pointing is done using the “interesting” peak, as defined in Chapter 4.5.1. . . .	70

3.20	This plot shows pointing offsets in ϕ for when using Hpol and Vpol to reconstruct the 45° WAIS pulses. The y-axis is the difference between ϕ reconstructed using Hpol channels and the actual ϕ wais. The x-axis is the difference between ϕ reconstructed using Vpol channels and actual ϕ wais. Hpol and Vpol point at approximately the same place, but the 45° and -45° pulses have opposite $d\phi$ offsets.	71
3.21	This plot shows pointing offsets in ϕ for the 45° WAIS pulses as a function of incoming θ angle to WAIS. Pointing is done using the “interesting” peak, as defined in Chapter 4.5.1. The y-axis is the difference between $\phi_{reconstructed}$ and ϕ_{wais} and the x-axis is θ_{wais} . 45° and -45° pulses have opposite $d\phi$ offsets and this offset only appears at $\theta > 14^\circ$	71
4.1	A WAIS calibration pulse as displayed in <code>anitaMagicDisplay</code> . The red and green waveforms are the vertically and horizontally polarized channels that triggered, and the black and blue waveforms are the vertically and horizontally polarized channels that did not trigger. The buttons in the top left correspond to different ways to view each event, and the buttons in the top right allow the user to navigate through events. In the center top there is useful information about each event displayed.	76
4.2	When all three notches are on, the sine subtraction algorithm can fail by falsely identifying an area between notches as a peak and subtracting it. This is shown with event 36785931, a cosmic ray. The blue waveform is unfiltered and the red waveform is after sine subtraction filtering. The signal has lost a significant amount of low-frequency power.	82

4.3	This is the interferometric map for a WAIS calibration pulse event. The z-axis (color axis) is the average correlation value for the waveforms going into the map, and the x and y-axes are ϕ and θ , respectively. For this map the peak is very clear, because it is a high SNR pulser event. The three black crosses on the map are the peaks found by the peak finding algorithm, where the biggest cross is placed on the biggest peak. There is a red star right at the center of the largest black cross indicating the true direction of the WAIS pulser.	85
4.4	A single sample digitizer glitch event example.	92
4.5	A step function event example. Step function events have badly distorted pointing and coherently-summed waveforms.	92
4.6	A payload blast event as viewed in <code>anitaMagicDisplay</code> . Notice the large, low frequency dominated waveforms present in the bottom and middle rings. This one was identified by a cut on the ratio of maximum voltages in the bottom to top ring.	94
4.7	The coherently-summed waveform and frequency spectrum of that waveform for the payload blast event shown in Fig. 4.6. Notice the huge amount of power between the low end of the band and the first notch, especially in the vertical polarization.	95
4.8	The black line that maximizes separation between signal and background in the space of power window gradient and difference between Hilbert envelope peak values of the coherently-summed and coherently-summed dedispersed waveforms, as calculated by ROOT's TMVA package. Pink points are simulated neutrinos, while the blue through yellow color axis is made up of events that reconstruct above the payload. Moving this line upwards would increase sample purity, while moving it downwards would increase efficiency.	97

4.9	Cumulative Fisher score distributions for testing sample of background (top, red) and signal (bottom, blue) events. Vertical black lines are drawn at expected thermal background leakage of 0.5, 0.05 and 0.01 events.	99
4.10	Fisher scores for various types of events. Thermal events in red, WAIS pulses in cyan, and simulated neutrinos in blue. The final Fisher score cut was set at -1 (dashed line) to allow for an expected background leakage of 0.05 events. . . .	100
4.11	Clustering's asymmetrical nature is shown. The top panel shows event A and event B 's error ellipses when viewed from above, with the X denoting where the event pointed. In the middle panel we see how where events A and B were viewed from. The bottom panel clearly shows that while B is less than 2σ from A when viewed from A 's frame, event A is over 3σ from B when viewed from B 's frame. Diagram taken from Ben Rotter's thesis [80].	102
4.12	Simulated Askaryan neutrino efficiency is on the Y-axis and $-2\log(\mathcal{L})$ is on the X-axis. Efficiency is flat at around 84% between $-2\log(\mathcal{L})$ values of 4 and 10. . .	103
4.13	The EAS template created using ten ANITA-IV events. This is the dedispersed version, zoomed in about the peak.	105
5.1	Vertically polarized scope trace taken using a Seavey antenna pointed at a snowmobile. Image courtesy of David Saltzberg and Konstantin Belov [127].	108
5.2	An illustrative general example of the ABCD method of background estimation. Region A is signal-like. Regions B, C, and D are background. V_1 and V_2 are uncorrelated for background events. The prescription for estimating the background using this method is detailed in the text.	110
5.3	The distribution of correlation values for the coherently summed WAIS pulser events. The procedure I followed to set a cut using this distribution is detailed in the text.	112

5.4	The distribution of differences from the expected geomagnetic polarization angles for the ANITA-III EAS candidate events. This distribution and the gaussian fit to it was used to determine the geomagnetic polarization angle difference cut on ANITA-IV EAS events. Figure taken from [80].	112
5.5	The distribution of linear polarization fraction values for the WAIS pulses. Using this distribution as a guide, I set the cut on EAS linear polarization at 0.625.	113
5.6	The power spectrum and waveforms for the dedispersed and coherently summed background singlet event are shown. Vertically polarized signal is red and horizontally polarized signal is blue. This is clearly not an EAS or Askaryan neutrino candidate, but it appears very isolated.	117
5.7	Event 50549772, which was moved from the signal region to the sideband region after a bug was found in the geomagnetic angle calculation code. The coherently summed waveform is in the top panel, the dedispersed waveform is in the middle panel and the power spectrum is in the bottom panel. The dedispersed waveform has a strange second peak after the first peak, which is different from the normal tripolar structure of an EAS, also the power spectrum peaks in the mid-band, which is unexpected for an actual EAS event.	120
5.8	The approximate location of EAS candidates on the continent are plotted (cyan X) along with the payload viewing location (magenta triangle).	121

5.9	The difference between the expected geomagnetically induced polarization angle and the measured polarization angle for EAS candidates. Cyan events are reflected cosmic rays, black events are direct cosmic rays, and other events are colored individually as denoted in the legend. The two events that were moved out of the signal region, 17904564 and 50549772 are the farthest from the central black line, in red and pink, respectively. Vertical error bars are based off of the ANITA-III cosmic ray events, error on expected polarization angles are estimated at 2° , from uncertainties in geomagnetic angle calculations.	122
5.10	All 28 of the EAS candidate's coherently summed waveforms. Black is the horizontally polarized channel, blue is the vertically polarized channel.	123
5.11	All 28 of the EAS candidate's dedispersed coherently summed waveforms. Black is the horizontally polarized channel, blue is the vertically polarized channel. The window is ± 15 ns about the peak intensity.	124
5.12	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is one of two CRs detected with all three notches on, and has a distinct frequency spectrum that peaks a bit higher than most.	125
5.13	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is one of two direct events.	126
5.14	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	127

5.15	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	128
5.16	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	129
5.17	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is one of the unusual polarity events.	130
5.18	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	131
5.19	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	132
5.20	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	133

5.21	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is the second of two CRs detected with all three notches enabled, and shows a similar spectrum to the other one.	134
5.22	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	135
5.23	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	136
5.24	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	137
5.25	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	138
5.26	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. The second of two direct cosmic rays.	139

5.27	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	140
5.28	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	141
5.29	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	142
5.30	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	143
5.31	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	144
5.32	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	145

5.33	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. The second of two unusual polarity events.	146
5.34	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. A low SNR cosmic ray that almost looks unipolar in shape.	147
5.35	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	148
5.36	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. A very off-cone cosmic ray. Notice how dominant the low frequencies are in the power spectrum, and how different the coherently summed template appears from the others.	149
5.37	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	150
5.38	Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.	151

5.39 Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. 152

5.40 Candidate waveforms are pictured in blue, with template waveforms as black dotted lines. Waveforms are dedispersed and windowed ± 5 ns about the zero-crossing nearest the peak of Stokes I. Red dotted lines are the reference point the waveform and template are aligned to. The template is aligned to the reflected CR polarity. 154

5.41 The windowed, dedispersed correlation for each EAS candidate. Windowing is done ± 5 ns from the zero-crossing nearest the peak value of Stokes I for the waveform. A positive cross-correlation value at the dotted red line indicates a normal reflected CR polarity by the windowed correlation metric, a negative value indicates polarity consistent with a direct CR. Dotted black lines are at \pm maximum absolute cross-correlation coefficient values, this can be used to estimate the PSL value. The astute reader will note that event 36785931 has a PSL value that opposite what the fiducial indicates for polarity. We concluded that this is due to distortion of the waveform caused by having all three notches on for this event, and categorize it as a normal, reflected CR, as two out of three metrics indicate a reflected CR polarity. 155

5.42 Results of the fourier phase method for each EAS candidate. The plots are color-coded so that red (positive phase) is reflected polarity and blue (negative phase) is direct polarity. The fourier phase method is described in the text, and the reference point of the waveform is shown in Fig. 5.40. 156

5.43	The PSL is plotted vs the correlation value vs a template. On the left in blue are the EAS candidates with polarity consistent with a direct CR, on the right in red are the EAS candidates with polarity consistent with a reflected CR. The three on the right with low PSL values are the two events with all three notches on, 4098827 and 36785931, and 88992443, which appears to be very off-cone. Event 36785931 has a PSL below 1 as well.	157
5.44	ANITA's flight path is shown overlaid with simulated neutrinos passing or failing clustering cuts with $-2\log(\mathcal{L}) = 8$. Simulated events that fail clustering as shown in magenta and passing events are shown in cyan. The bulk of efficiency loss comes from events near South Pole Station.	161
5.45	Event 36019849 (the lone neutrino candidate in the Askaryan channel), dedispersed and coherently summed, with the vertically polarized waveform in red and horizontally polarized waveform in blue. It is clearly a very impulsive, vertically polarized signal. The bottom panel shows the power spectrum of the top panel in dB.	162
5.46	Event 17904564 (the isolated singlet that made it into the Hpol signal region by mistake), dedispersed and coherently summed, with the vertically polarized waveform in red and horizontally polarized waveform in blue. It is a very impulsive event, with nearly equal power in Hpol and Vpol. An interesting event to be sure, but there is not much more we can say about it. The bottom panel shows the power spectrum of the top panel in dB.	164
5.47	Event 36019849 with its 5-sigma error ellipse projected onto the continent. The nearest known location of human activity, a fixed-wing landing site on Deverall Island, is also shown.	165

5.48	The effective area as a function of energy used in calculating the ANITA-IV limit. This is calculated using the geometric average of two independent Monte Carlo simulations.	167
5.49	The all-flavor diffuse UHE neutrino flux limit for ANITA, as well as Auger and IceCube [133, 60], along with predicted fluxes from Kotera and Ahlers models [116, 132]. The combined limit is the most stringent above energies greater than $10^{19.5}$ eV.	168
6.1	Simulated trigger efficiencies for 3, 8 and 16-antenna phased arrays summed into a single beam. Efficiency is drawn for rates of 1000 Hz, 100 Hz, and 10 Hz, denoted by the black-solid, dashed and gray-solid lines, respectively. Figure taken from [136].	172
6.2	Thermal noise correlation measurements setup in the PSL anechoic chamber. Antennas are placed in each other's nulls, and antenna output is passed through amplification and filtering along the way to the Tektronix scope, where it is recorded. Grey, shaded triangles represent the RF attenuating foam that lines the anechoic chamber. Figure taken from [136].	173
6.3	A picture of the constructed GNO dipole is pictured on the left. The middle panel shows the measured S_{11} of all constructed antennas (in color) and the HFSS simulated S_{11} in dashed black. On the right panel is the measured radiation pattern of a single antenna (solid color) and the simulated radiation pattern (dashed color). Figure taken from [136].	173

6.4	Histograms of the maximum cross-correlation coefficient within a causal window (left) and all cross-correlation coefficients within a causal window (right). Black is simulated uncorrelated noise, blue is uncorrelated noise taken with terminated amplifiers, and red is the data taken using the ARA Vpol bicone. All three of these are consistent with being from the same distribution, showing that noise between closely spaced antennas is uncorrelated. Figure taken from [136]. . . .	176
6.5	The mean of the maximum correlation value for 500 trials of simulated data (blue) and data recorded with the ARA bicones (green). As the distance between the antenna feeds decreases, the thermal noise eventually becomes correlated, but the required distance is unachievable. Data was taken down to the smallest physically realizable distance. The simulated correlation coefficient maxes out around 0.8 because roughly 20% of thermal noise in the system is uncorrelated noise introduced by the amplifiers. Figure taken from [136].	177
6.6	Diagram representing the setup during the tests with a transmitting antenna. Transmitting antenna and receiving antennas were all the same antenna type. Signal chain is the same as in Fig. 6.2. Figure taken from [136].	178
6.7	Voltage vs time traces of each channel on the left and the combined beam on the right. Channel 1 is solid red and channel 2 is dashed blue on the left and the ideal beam as combined in software is solid green, while the beam as combined in hardware is shown in dashed orange. Channel 1 and channel 2 are 2.6 and 2.7 SNR, respectively, while the hardware and software combined beams are 3.1 and 3.3 SNR, respectively. Figure taken from [136].	179

LIST OF TABLES

3.1	TUFF configurations and the fraction of the flight they were active. The configurations with very low fractions of flight time were failed experiments. For instance, we turned off the 2nd notch twice, and trigger rates raised so drastically that we immediately turned it back on.	56
4.1	Table of quality cuts and their effect on the complete sample, the WAIS pulses, and the simulated neutrino sample. V Ratios stands for the suite of voltage ratios cuts, P Ratios stands for the suite of power ratio cuts, and Imp. vs P stands for the impulsivity vs. power cuts, all of which were implemented to remove payload blasts. For more details about each cut, see text. Total quality cut efficiency on WAIS pulses is $\sim 97\%$ and total efficiency on simulated neutrinos is $\sim 98.2\%$	96
4.2	Table of variables that go into the Fisher discriminant, and values that come out, ranked by discriminatory power. LPF stands for linear polarization fraction, Wind. is short for windowed, Imp. is short for impulsivity, and PWG is power window gradient. The Fisher score for each event is obtained by multiplying each variable's value by the given coefficient and adding an overall offset (-49.178 for my analysis). Discriminatory power represents the amount of separation you can achieve between signal and background distributions of your training sample by using only that variable. Minimum, maximum, and mean values of each variable for all events in the training sample are also included, to give a sense of how the coefficient might actually act on each variable.	98
5.1	The values for each EAS candidate event for CR-like cut variables. Event 50549772 is reported at the bottom separately.	119

5.2	The values for each EAS candidate event for polarity metrics, with their respective chances to have their polarity misidentified in parentheses. A 0 for misidentification chance means that there were no misidentified polarities in 10000 independent trials for that candidate event. The coherent correlation, PSL and Fourier phase are positive for normal reflected CR polarity and negative for direct CR polarity. Only one of the events has disagreement between the three polarity metrics, event 36785931. Event 50549772 is shown separated by a horizontal line at the bottom.	159
5.3	Probabilities for the two mystery event candidates to be background under several different sets of assumptions. Scenarios and assumptions are more fully described in the text.	160
5.4	Summary of the effect of cuts in my analysis. The quoted efficiency is the cumulative efficiency on Monte Carlo generated neutrinos for all cuts when performed in sequence. Quality, thermal and clustering cuts are detailed in Chapter 4. The Askaryan neutrino analysis is left with one event in the signal region at the end of all of these cuts.	161
5.5	Likelihood ratio calculations for the Askaryan neutrino candidate. The likelihood ratio is derived from simulated neutrinos and clustered events. Observables are chosen because they should be very strong in true neutrino candidates. The total product is the ratio of the likelihoods of the event to be signal to background. Event 36019849 appears to be about 6 times more likely to be background than signal according to this choice of variables.	166

ACKNOWLEDGMENTS

First and foremost I would like to thank my boss, mentor and advisor Abigail Vieregg for her guidance and support throughout the entire Ph.D. process.

Next, I want to thank everyone on the ANITA collaboration for allowing me to work on such an interesting project while providing useful advice along the way. I want to especially thank Cosmin Deaconu for both creating most of the framework I used in my analysis and teaching me everything I know about ROOT and vim along the way. Thanks also to Eric Oberla for always being available to help. Special thanks to Linda Cremonesi for all of her work on `icemc`, which made my analysis possible. I very much appreciate Steven Prohira's willingness to answer whatever physics questions I've had for the past couple years. Thanks to my co-analyst of ANITA-IV, Peng Cao, and by extension, his advisors John Clem and Dave Seckel, for being there to talk through ideas about the analysis and help discover and fix bugs. I want to thank Ben Strutt and Ryan Nichol for helping create almost all of ANITA's software. I specifically want to thank Ben Rotter for three things: a beautiful thesis that I borrowed a few figures from, a well thought-out cosmic ray analysis that inspired my approach, and a lot of code for deriving impulse responses. Nan Wang merits a special thanks for always being willing to do any job and do it well. David Saltzberg, Peter Gorham, Steph Wissel and Dave Besson have all given me valuable feedback and direction on ANITA phone calls throughout my analysis, thank you all. A very heartfelt thanks goes out to Christian Miki and Brian Hill for somehow making a summer spent in Palestine, Texas a fun experience. Thanks to the OSU crew of Amy Connolly, Patrick Allison, Oindree Banerjee, Jacob Gordon and Keith McBride for their work on many aspects of the project, but especially `icemc` and the TUFFs. Thanks to John "JRuss" Russell and Luke Batten for all of their work on many aspects of ANITA-IV.

Thanks also to my friends and officemates in the UChicago Physics program. Thanks to Ryan McGeehan, Carlos Blanco and Kaeli Hughes for sharing an office with me and

providing valuable conversation and insights. An extra special thanks goes out to Akash Dixit and Zoheyr Doctor for living with me for two years and generally being very good roommates.

I also would like to thank my friends and family away from the physics world, but especially my parents and my sister, who have supported me always.

Finally, thanks to Denise, my beloved.

ABSTRACT

Ultra-high energy (UHE) neutrinos are a promising messenger particle, and their discovery would be an important step towards understanding the universe at energies above 10^{18} eV. The ANtarctic Impulsive Transient Antenna (ANITA) is a NASA long duration balloon payload that aims to measure the UHE neutrino flux by detecting Askaryan radiation in the radio band from neutrinos interacting in the Antarctic ice. Previous ANITA flights have also demonstrated an ability to detect geomagnetically induced radio emission from UHE cosmic rays, both from emission beamed directly at the payload, and from emission that has been reflected off the Antarctic ice.

In this thesis I present information about the fourth flight of ANITA, which took place during December of 2016. Contained within are details of the instrument, the calibration procedures, the analysis, and finally, the results. In the neutrino channel of my analysis I find 1 candidate event on an expected background of $0.64^{+0.69}_{-0.45}$ events. While this is not a discovery, it does allow us to set the world leading limit at energies between $10^{19.5}$ eV and 10^{21} eV. In the cosmic ray and extensive air shower channel of my analysis, I find 28 events, 26 of which are likely candidates for normal cosmic ray induced extensive air showers. The remaining two events reconstruct to locations below the horizon, but lack the 180° phase inversion seen in reflected cosmic ray events.

CHAPTER 1

INTRODUCTION

1.1 Neutrino Physics and Astronomy

Wolfgang Pauli originally proposed a light, neutral particle in 1930 in order to explain the missing energy and momentum in the β -decay process [1]. Enrico Fermi took this idea and ran with it, christening the particle neutrino, or “little neutral one” in Italian, and making it the basis of his theory of β -decay in 1932 [2, 3]. In 1953, the existence of the neutrino was confirmed experimentally by the Cowan-Reines experiment [4], by directly measuring the rate of inverse β -decay near an active nuclear reactor.

1.1.1 Solar Neutrinos

Once neutrinos were confirmed to exist, the natural next step was using them for astronomy applications. The first such experiment was Ray Davis’s Homestake mine experiment, which began in the late 1960s to measure the neutrino flux coming from the Sun [5]. The experiment measured between 27.5% and 33.3% of the ν_e flux predicted by various models of nuclear interactions taking place in the Sun [6, 7, 8]. Other experiments observed similar deficits in the ν_e flux, including Kamiokande and Super-Kamiokande [9, 10, 11], Sage [12], and GALLEX [13]. Super Kamiokande showed evidence for neutrino oscillations, which would explain the flux deficit, at the 90% confidence level in 1998 [14], but the solar neutrino problem was not conclusively solved until the Sudbury Neutrino Observatory (SNO) came along in the early 2000s. SNO was able to measure both the total neutrino flux and the electron neutrino flux from the Sun, and measured a flux that agreed with model prediction in aggregate, but with a different flavor composition [15].

1.1.2 Neutrino Oscillations

SNO's result was the first experimental confirmation of neutrino oscillations. Neutrino flavor (ν_e, ν_μ, ν_τ) and mass (ν_1, ν_2, ν_3) eigenstates are not the same thing. In fact, flavor eigenstates are superpositions of mass eigenstates, which can be represented using the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [16, 17], U :

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.1)$$

where U is commonly parameterized as:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{-i\delta_{CP}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \quad (1.2)$$

where δ_{CP} is the charge parity violating phase, and s_{ij} and c_{ij} represent $\sin \theta_{ij}$ and $\cos \theta_{ij}$, respectively, where θ_{ij} is the mixing angle between the i th and j th neutrino mass eigenstate. If U were the identity matrix, neutrino flavor and mass eigenstates would be equivalent.

The probability of oscillating from one flavor to another turns out to be dependent on the difference between the squares of the masses, δm^2 , the mixing angles between the mass eigenstates, θ , the neutrino energy, E , and the traveled distance, L . A suite of experiments have been designed at different L and E scales to probe various mixing angles and mass differences both with terrestrial sources like in T2K [18] and Daya Bay [19], and astronomical sources like in Kamiokande and Super-Kamiokande [11].




Figure 1.1: The astrophysical neutrino flux as measured by IceCube. Plot taken from [21].

1.1.3 Neutrino Flux Measurements

Experiments like Super-Kamiokande and SNO measure the flux of neutrinos at low energies. The next natural step was to push those measurements to higher energies. The first experiment to successfully do so was IceCube, which grew out of an experiment called AMANDA. In 2013, IceCube was able to record the first ever neutrino events in the PeV energy range, the highest energy neutrinos ever detected at the time [20]. Since that initial discovery, IceCube has mapped out the diffuse high energy flux of neutrinos out to a few PeV [21]. IceCube’s measured neutrino flux at energies between 10^{12} and 10^{16} eV is shown in Fig. 1.1. Out past the edge of the IceCube flux, the shape of the spectrum is still unknown.

1.1.4 Neutrino Source Observations

Apart from the Sun, there have been two other observed sources of neutrinos, the first of which, Supernova 1987A, is named after the year in which it was observed. It produced a detectable burst of neutrinos shortly after the star powering it collapsed towards becoming

a neutron star. Neutrino emission is a natural part of the neutron star formation process:

$$e^- + p \rightarrow n + \nu_e \tag{1.3}$$

After the initial burst, there is an additional leakage of neutrinos diffusing out past the confines of the newly formed dense star and escaping. SN1987A produced many ν_e events, 20 of which were recorded by Kamiokande and the IMB water Cherenkov detector [22, 23].

More recently, IceCube recorded a tantalizing coincidence between neutrinos and the blazar TXS 0506+056 [24]. A blazar is a luminous active galactic nucleus (AGN) with its relativistic jet beamed straight towards Earth. In September 2017 IceCube detected a 290 TeV neutrino from the TXS 0506+056 blazar, which was flaring in gamma-rays (emitting higher than normal gamma-ray emission) at the time [25]. After discovering this coincidence, IceCube performed a search through archived data for evidence of other neutrinos from the same source and discovered a flare of neutrinos in 2014-2015 [24]. At this time in 2014-2015, the blazar was not flaring in gamma-rays [24], so there is some mystery about how multi-messenger flares fit together. This is the highest energy source of neutrinos observed to date.

1.2 UHE Neutrino Astronomy

Neutrinos are ideal messenger particles for astrophysical sources because they are weakly-interacting, neutral particles. This allows them to travel for very long distances with low probabilities of interacting with anything, with their trajectories unaltered by the magnetic fields they propagate through. Neutrino astronomy opens areas of the universe that are inaccessible to conventional photon or cosmic ray (CR) astronomy. Sources that are too distant or opaque for CRs and photons can be viewed in the light of neutrinos. A schematic representation of the portion of the universe unobservable to photons and cosmic rays in




Figure 1.2: Much of the universe is unobservable with photons or protons. The hatched blue region is all of energy-distance phase space that is not observable with photons, the hatched red region is all of energy-distance phase space that is not observable with protons. Various astrophysical object distances are also shown. Plot courtesy of Dr. Peter Gorham.

Energy-Distance phase space is shown in Fig. 1.2.

1.2.1 Multi-messenger Observations

Further observations of neutrino sources, especially in the ultra-high energy (UHE) regime, would open up the field of UHE multi-messenger studies. Neutrinos observed directly from sources are indicators of hadronic processes that will also produce cosmic rays. UHECRs are deflected by magnetic fields, and therefore not usable in multi-messenger astronomy, however, due to the Griesen-Zatsepin-Kuzmin (GZK) effect (see Chap. 1.3.2 for more), neutrinos

and gamma rays can act as stand-ins for UHECRs. Coincident neutrino and gamma ray observations of a source would be a very powerful indicator of its identity as a UHECR accelerator [26]. Certain gravitational wave sources, such as neutron star mergers and tidal disruption events are expected to produce accelerators powerful enough to produce UHECRs and neutrinos [27]. Multi-messenger observations are the path towards having complete knowledge of all of the processes in an astrophysical source. Multiple messenger particle are also necessary to view the entirety of a source, for instance, the innermost core of a source may be opaque to gamma rays but transparent to neutrinos.

1.2.2 Fundamental Physics Possibilities

UHE neutrinos probe an energy and distance regime that is otherwise unreachable on Earth. This allows unique windows into particle physics that will complement and expand on what is being done in current and future terrestrial collider experiments. There are many fundamental physics questions that UHE neutrinos can probe, I will highlight the two that I find most interesting here, but for a more in-depth discussion, see [28].

One question that UHE neutrinos could help answer is whether neutrino cross sections behave as expected in the UHE regime, and whether there are any beyond the standard model (BSM) modifications to cross sections. Currently the highest energy measured cross sections for neutrinos are from IceCube [29] and match standard model expectations within uncertainties. If certain BSM theories like, for instance, large extra dimensions were accurate, neutrino cross sections would greatly differ from standard model cross sections at higher energies [30].

Another question answerable by UHE neutrinos is whether flavor mixing happens as expected in the UHE regime. If flavors mix as expected by the standard model, the composition of UHE neutrinos detected on Earth should be close to maximally mixed [31]. Certain BSM effects like neutrino decay or Lorentz violation would allow for very different neutrino flavor

compositions.

1.3 Cosmic Ray Astronomy

Cosmic rays (CR) were discovered in the early 1900s. Physicist and Jesuit priest Theodor Wulf believed that radiation measured on the Earth came purely from terrestrial sources, and should wane as a function of height above the Earth. In order to test this, he measured ionizing radiation as he ascended the Eiffel tower and found it dropped much less swiftly than theory would predict [32]. In 1912, Victor Hess followed up on this experiment by systematically measuring the level of ionizing radiation at various altitudes by way of a hot air balloon [33]. Hess saw that ionizing radiation actually began increasing above a certain altitude, indicating an additional extraterrestrial origin for radiation.

1.3.1 Cosmic Ray Spectrum

In the century following the discovery of CRs, many experiments were designed to map out the composition, sources, and energy spectrum of cosmic rays. At the highest energies, recent measurements from Auger probing the composition of CRs using maximum shower depth, X_{max} , suggest a composition mixed between protons and heavier nuclei [34], trending heavier above $\sim 10^{18.5}$ eV. At ultra-high energies, X_{max} can be observed directly via fluorescence detectors, and measurements of it can be used to determine the mass of the primary particle, A , because the average depth of the shower maximum, $\langle X_{max} \rangle$ is expected to follow this relation [35]:

$$\langle X_{max} \rangle = \alpha (\ln E - \langle \ln A \rangle) + \beta \quad (1.4)$$

where the coefficients α and β depend on the nature of hadronic interactions, and E is the energy of the incident CR. However, while Auger's measurement suggests a heavier composition, Telescope Array's (TA) measurement suggests a lighter mixed composition [36],

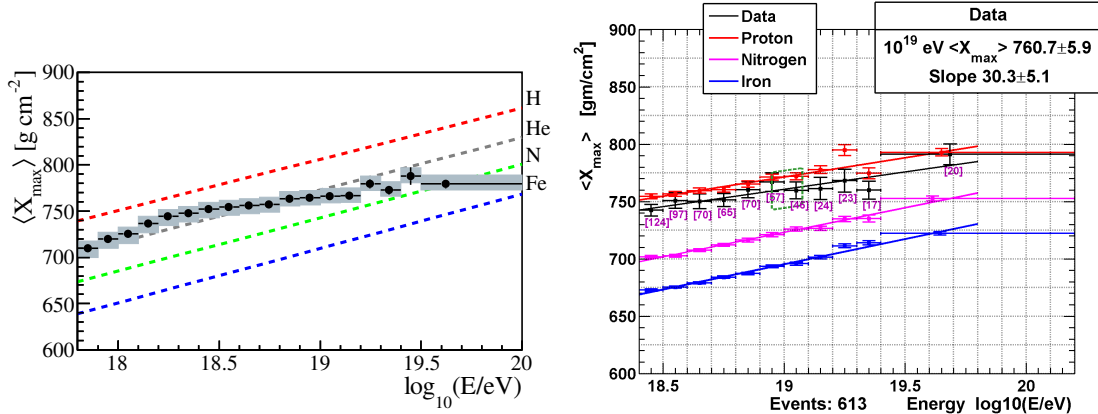


Figure 1.3: Auger’s measurement of composition of the CR spectrum using the average of the X_{max} distribution as a function of energy. Each dashed colored line represents a pure spectrum of a single element, red is H, grey is He, green is Ne, and blue is Fe. Figure is taken from [38]. TA’s measurement of composition of the CR spectrum using the average of the X_{max} distribution as a function of energy. The red line also represents various models a pure proton composition here, while the pink line is a pure nitrogen composition and the blue line represents a pure iron spectrum. As you can see, the data from TA indicates a spectrum much closer to pure proton than for Auger, however, the discrepancy between the two measurements are within systematic uncertainties [37]. Figure is from [36].

although the discrepancy is within systematic uncertainties [37]. Fig. 1.3 shows the predicted composition of CRs as a function of CR energy, with Auger’s measurement in the left panel and TA’s in the right panel. The spectrum follows a broken power law, with three changes of index at the “knee” ($\sim 3 \times 10^{15}$), “second knee” or “iron knee” ($\sim 10^{17}$) and “ankle” ($\sim 3 \times 10^{18}$) of the spectrum, as shown in Fig. 1.4.

Accelerating CRs to these energies is a challenge, and is modeled by two distinct mechanisms, first and second order Fermi shock acceleration. First order Fermi shock acceleration, or diffusive shock acceleration, requires a shock front propagating through interstellar medium [40]. Charged particles are accelerated by scattering off the shock front, diffusing across it and being scattered again, and so on, gaining energy each time.

Second order Fermi shock acceleration occurs in regions of space with large magnetic inhomogeneities, called magnetic clouds [41]. Charged particles reflect off these magnetic clouds in elastic collisions, because the cloud is very massive compared to the particle,

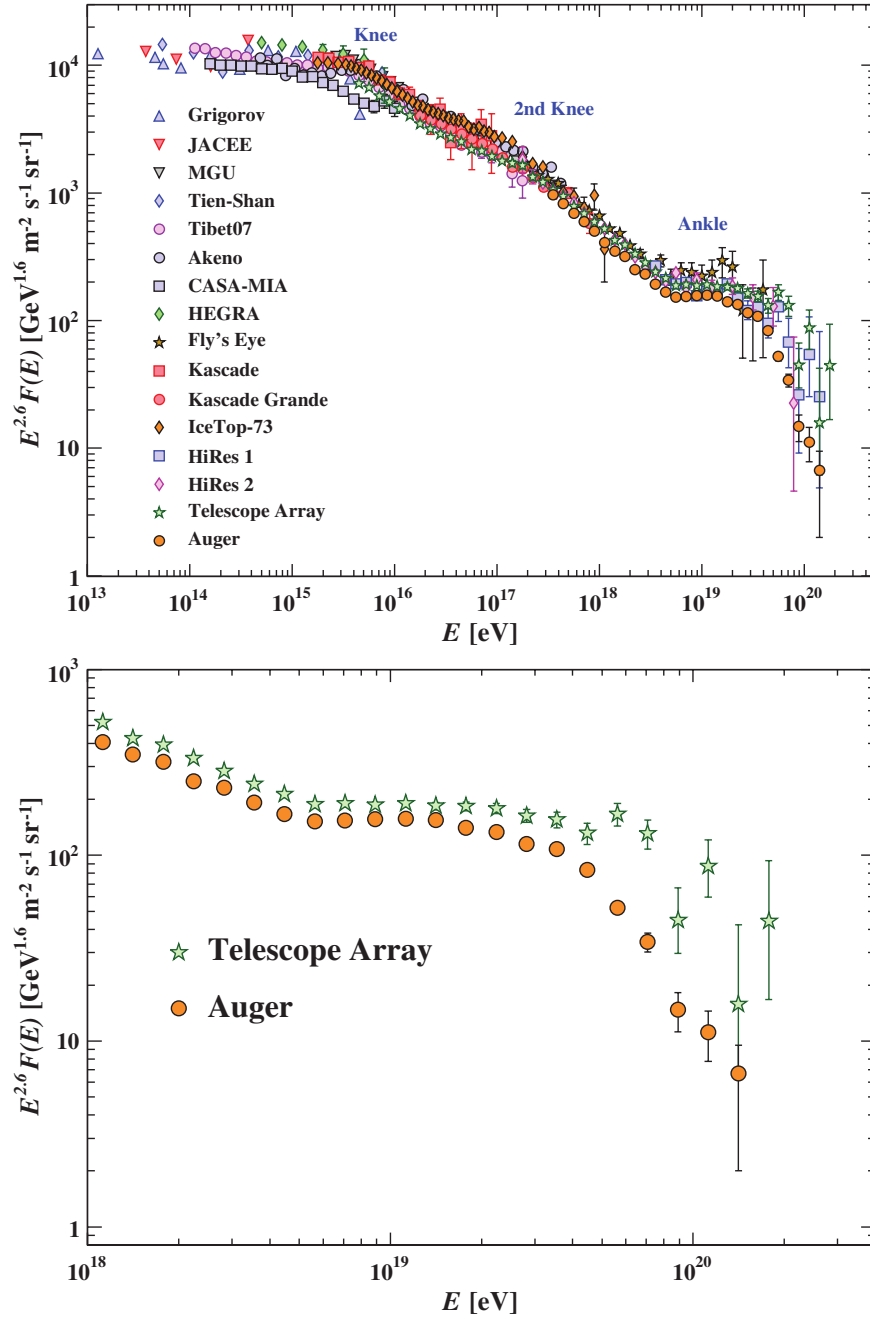


Figure 1.4: In the top panel is the CR flux as a function of energy-per-nucleus as measured by various experiments. The flux is multiplied by $E^{2.6}$ to more easily see the changes in spectral slope. Labeled with arrows are the knee, 2nd knee and ankle of the spectrum. In the bottom panel is the same flux zoomed in about the highest energy ankle region. These UHE fluxes are only measured by Auger and TA, and their discrepancy is within systematic uncertainties up to $10^{19.4}$ eV, at which point the discrepancy is not fully explained [37]. Figures are from [39].

gaining energy by reflecting off an approaching cloud and losing energy by reflecting off of a receding one. Collision with an approaching cloud is more probable than with a receding one, so a charged particle will gain energy over time. First order shock acceleration in supernova remnants can account for cosmic ray acceleration up to the knee and possibly the ankle of the spectrum [42].

There are two interpretations for why the spectral index of the CR spectrum steepens at the knee. First is that beyond a certain energy, CRs can escape the galactic magnetic fields that would otherwise restrict them to our galaxy. If galactic CRs can escape out of the confines of the Milky Way, we have a lower chance of detecting them on Earth. This happens if the radius of curvature of a charged particle in a magnetic field, or Larmor radius, is larger than the bounds of the Milky Way. The Larmor radius, r_l , is defined as:

$$r_l = \frac{p_{\perp}}{|q|B} \quad (1.5)$$

where p_{\perp} is the momentum of the particle perpendicular to the magnetic field, B , and q is the charge of the particle. This explanation of the knee assumes that Galactic sources strong enough to accelerate CRs to these energies exist, which is still an open question.

The second possible explanation for the change in index at the knee of the spectrum is that shock fronts from Galactic sources cannot accelerate cosmic rays above a certain energy. Particles will escape the shock front without being diffused across it once they have reached a certain energy. This maximum energy, E_{max} is given by the Hillas criterion [43]:

$$E_{max} = \frac{|q|BR}{c} \quad (1.6)$$

where R is the size of the accelerator, and this equation is obtained by enforcing that the Larmor radius not exceed the size of the accelerator. Potential sources of UHECRs obeying the Hillas criterion can be represented on a Hillas diagram, where candidate sources are




Figure 1.5: An updated Hillas diagram taken from [45]. Various sources are shown in $B - R$ phase-space. The phase-space above the blue (red) line is the allowed region for acceleration of protons (iron nuclei) above 10^{21} eV (10^{20} eV).

plotted as a function of B and R , as in Fig. 1.5.

The second knee occurs naturally if one assumes that the spectral composition up to the knee is proton dominated. There should be a corresponding knee for the iron portion of the spectrum at $\sim 10^{17}$ in either possible explanation for the first knee. In the first scenario, the Larmor radius, given by Eq. 1.5, is decreased by a factor of 26, or the atomic number of iron, which means the CR's momentum must increase by the same factor before leaking out of the Milky Way. In the second scenario, the E_{max} given by Eq. 1.6 is increased by a factor of 26, again, the atomic number of iron. This second knee is sometimes called the “iron knee” for this reason.

The spectral index changes again at the ankle, this time becoming flatter. The common interpretation of this is that the spectrum is bolstered by the presence of a flux of extragalactic cosmic rays [44].

1.3.2 The GZK Effect

Cosmic rays at energies past the ankle, $> 10^{18}$ eV are categorized as ultra-high energy. Past this ankle, at $\sim 10^{19.5}$ eV, the spectrum changes again. The measured CR flux drops off very sharply. This happens because UHECRs at this energy interact with the 2.7 K CMB photons with a center-of-mass energy at a Δ^+ resonance, producing a nucleon and a pion in the process:

$$p + \gamma_{CMB} \rightarrow \Delta^+ \rightarrow n(p) + \pi^+(\pi^0) \quad (1.7)$$

This process was predicted independently by Griesen [46], and Zatsepin and Kuzmin [47], and named the GZK process in their honor. The GZK process makes it extremely unlikely for CRs with energy per nuclei above this threshold to propagate more than a few tens of Mpc from their sources. In fact, only protons and iron nuclei may reach the earth from distances $\gtrsim 50$ Mpc, because light composite nuclei will photo-disintegrate with CMB photons [45]. The fraction of CRs for different nuclei arriving at energies > 60 EeV is shown in Fig. 1.6, clearly showing only protons and iron nuclei are viable CR candidates at distances $\gtrsim 50$ Mpc and energies > 60 EeV.

While the GZK effect makes CR astronomy difficult, it opens up the field of neutrino astronomy. Beresinsky and Zatsepin noted that the decay chain from the initial GZK effect will produce three neutrinos if one of the GZK products is a charged pion [48]:

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (1.8)$$

followed by:

$$\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu \quad (1.9)$$

Each of these GZK neutrinos carries away a few percent of incident proton's energy, producing three neutrinos with energies $\sim 10^{18}$ eV. A GZK dominated neutrino spectrum would




Figure 1.6: The fraction of CRs that survive from a distance $> D$ is plotted vs D for H, He, CNO, and Fe at $E > 60$ EeV, and H at $E > 40$ EeV and $E > 100$ EeV. A solid gray line showing the 50% survival point for various nuclei compositions is also plotted. Only protons or iron can survive for distances greater than 50 Mpc at $E > 60$ EeV. Figure taken from [45].

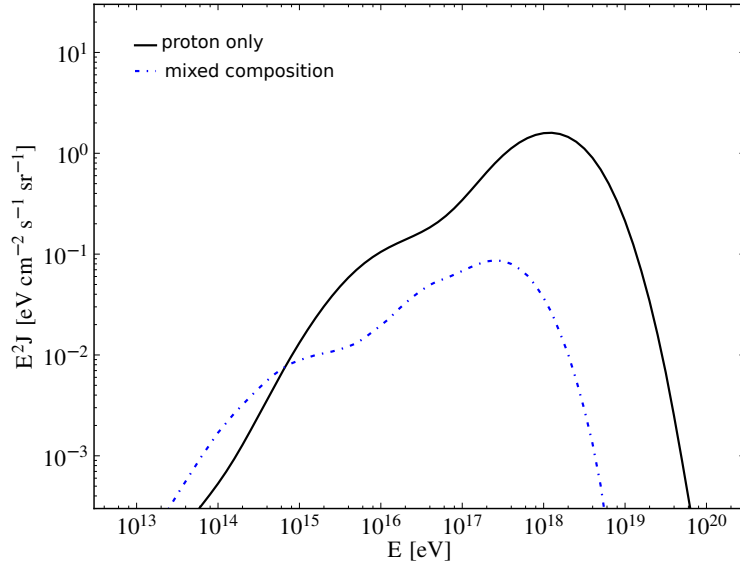


Figure 1.7: Two models for neutrino flux are shown for different CR composition assumptions. Black is an all proton flux and dashed blue is a mixed composition (between pure proton and pure iron fluxes). The expected neutrino flux scales downwards with increasingly heavy UHECR flux compositions. Figure is taken and slightly modified from [49].

therefore peak near 10^{18} eV. There is a secondary peak in the spectrum from neutron decays near 10^{16} eV. Both of these peaks are scaled about an order of magnitude in energy for more distant sources because the CMB was hotter at earlier times, which lowers the threshold for Δ^+ resonance, and neutrino energies are reduced by redshift. Specific GZK neutrino spectrum peak positions are also dependent on cosmological parameters and source composition. The flux of neutrinos from this process is highly dependent on UHECR element composition.

Heavier elements have a higher GZK process threshold because the average energy per nucleon must be above the Δ^+ resonance threshold, and average energy per nucleon is reduced by the number of nucleons. Photodisintegration is the main process by which heavy UHECRs lose energy:

$$A + \gamma_{CMB} \rightarrow (A - nN) + nN \quad (1.10)$$

where the atomic weight of the primary UHECR, A , is reduced by n nucleons, N . This causes the GZK neutrino flux to decrease by orders of magnitude for heavier compositions of

UHECRs. Fig. 1.7 shows this dependence on UHECR composition by plotting an all proton flux and a mixed composition flux, with the same source evolution assumptions.

1.3.3 Production of UHE Cosmic Rays and Neutrinos

Using the GZK process and the mechanisms of shock acceleration, theorists have built many “bottom-up” models of UHECR and neutrino sources. These models are bottom-up in the sense that all of them start with lower energy charged particles that are accelerated to higher energies by rotating induction or Fermi shock acceleration. Models include AGN [50, 51], pulsars [52], low-luminosity gamma-ray bursts (GRB) [53], tidal disruption events (TDE) [54], binary neutron star mergers [55], and more. Some of these models are shown, along with the neutrino fluxes and limits as measured by IceCube in Fig. 1.8. More exotic “top-down” models, in which cosmic rays are created at UHE levels, have mostly been ruled out by experiment [56], although some, based on super-heavy dark matter decays, still exist [57].

1.4 Radio Detection of Neutrinos and Cosmic Rays

The flux of UHECRs and UHE neutrinos is unfeasibly low for direct observation, so secondary detection techniques are a necessity. UHE neutrinos have the additional challenge of interacting only weakly with matter. The best way around these constraints is instrumenting and monitoring a huge detector volume.

1.4.1 Askaryan Radiation

Neutrinos interact in two distinct ways, through neutral current (NC) and charged current (CC) interactions. Both of these interactions are mediated by the exchange of a boson, W^\pm in the CC case, and Z^0 in the NC case. It is not possible to determine the flavor of the incident neutrino in NC interaction and the resulting shower is hadronic caused by nucleon

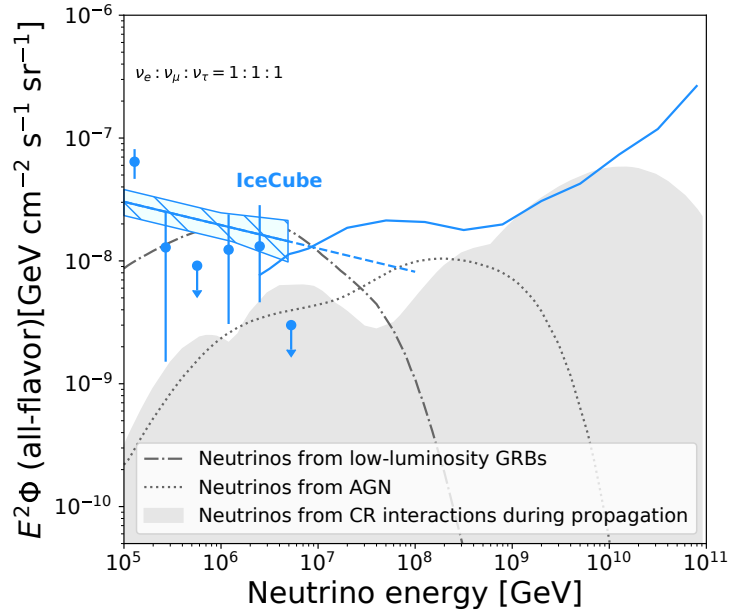


Figure 1.8: The current leading detections and limits of neutrinos at high energies from IceCube are shown [58, 59, 60], along with various models. Models displayed are for low-luminosity GRBs [53] and AGNs [51], along with an extrapolation of the IceCube flux. All allowed GZK fluxes from UHECR constraints according to [61] are shown in shaded grey. Figure is taken and slightly modified from [62].

scattering, which turns into an EM shower. In a CC interaction, the flavor of the incident neutrino and produced lepton are the same, and the resulting shower is mostly EM in nature.

In the 1960s Gurgen Askaryan recognized that an electromagnetic shower produced in a dielectric would give rise to coherent Cherenkov radiation [63, 64]. The Cherenkov radiation energy E for a single particle is given by the Frank-Tamm formula:

$$\frac{dE}{dx d\omega} = \frac{q^2}{4\pi} \mu(\omega) \omega \left(1 - \frac{1}{\beta^2 n^2(\omega)} \right) \quad (1.11)$$

where x is length, ω is frequency, q is the charge of the particle, $\mu(\omega)$ is the frequency dependent dielectric magnetic permeability, β is fractional speed of light, $\frac{u}{c}$, and $n(\omega)$ is the frequency dependent refractive index of the medium. As the electromagnetic shower develops, a $\sim 20\%$ negative charge excess forms because electrons are Compton scattered into the shower, while positrons annihilate. Like Cherenkov radiation, Askaryan emission is radially polarized and in a cone, with the cone angle, θ_c given by:

$$\cos(\theta_c) = \frac{1}{n\beta} \quad (1.12)$$

where n and β are the same as in Eq. 1.11. The shower is characterized in the transverse direction by its Molière radius, which is a material dependent quantity. At wavelengths larger than the Molière radius, the Cherenkov radiation from all particles interferes constructively. One way to think about this is that wavelengths this large see the shower as a single particle. A dense dielectric material is advantageous because the shower will be confined to a smaller disk with a smaller radius, making smaller wavelengths coherent. As the viewing angle for Askaryan radiation becomes more “off-cone,” coherence is lost at high frequencies, shown in Fig. 1.9. The coherently emitted power scales with the square of the incident particle energy, as implied by Eq. 1.11. Power also scales with number of particles in the shower times the charge excess fraction squared.

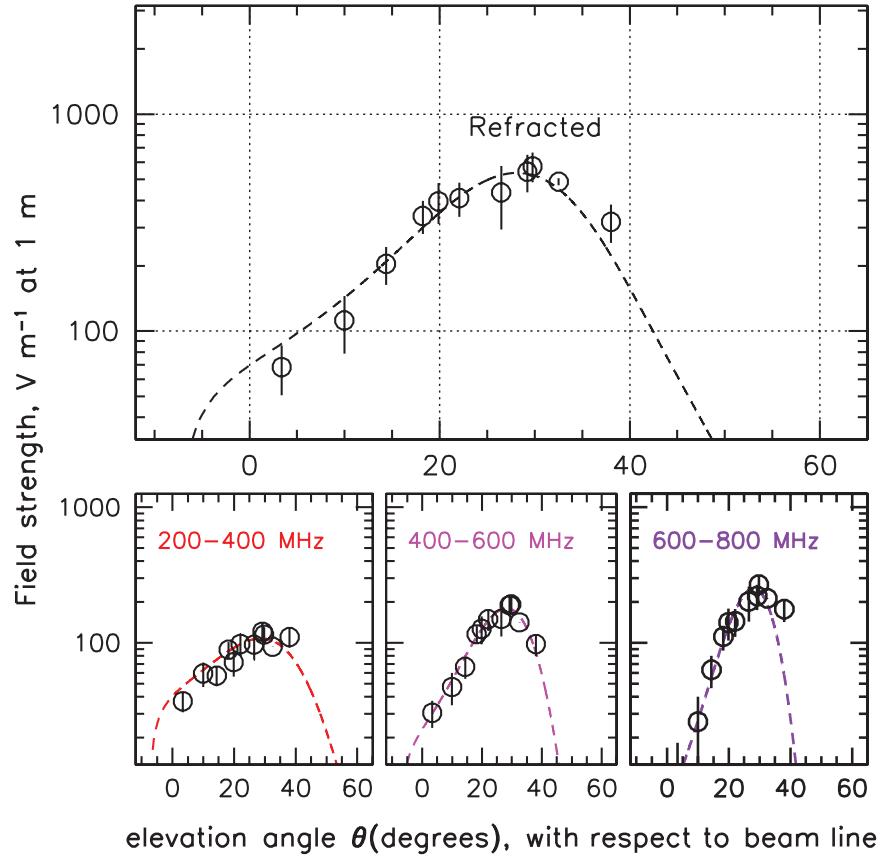


Figure 1.9: Figure is from [65], the paper detailing the first observation of the Askaryan effect originating in ice. The top panel shows the angular dependence of the recorded electric field for the full 200-800 MHz band of the experiment. The bottom panel is the same as the top panel, but split into smaller frequency bands. High frequency power falls off more quickly at angles away from the Cherenkov cone angle than low frequency power because it loses coherence.




Figure 1.10: A recorded Askaryan pulse, produced in ice, and recorded by the ANITA-I electronics. Figure from [65].

Askaryan radiation was experimentally confirmed and measured in 2001 at SLAC using silica sand as the dielectric material [66]. Later the effect was also measured in salt [67], and then finally in ice in 2007, using ANITA horn antennas as the detector [65]. In ice, Askaryan radiation is coherent out to a few GHz. Fig. 1.10 shows an Askaryan pulse from the ice measurements.

1.4.2 Extensive Air Showers

Pierre Auger discovered that cosmic rays that enter the Earth’s atmosphere with energy $> \sim 10^{15}$ eV will cause a particle cascade, known as an extensive air shower (EAS) [68]. Radiation from an EAS was predicted to be Askaryan in nature. Radio emission was observed in 1964 [69] and confirmed by other experiments after [70, 71], but measured radiation was inconsistent with Askaryan dominated emission. Further theoretical work postulated that EAS radiation should actually be dominated by geomagnetic radiation [72], and the Auger experiment has measured that only $\sim 14\%$ of cosmic ray shower emission is Askaryan on average [73].

1.4.3 Geomagnetic Radiation

Geomagnetic radiation occurs when the electrons and positrons produced in an EAS are split by traveling through the Earth's magnetic field. Each particle will experience the Lorentz force, where \bar{v} is the velocity of a particle, \bar{B} is the Earth's magnetic field:

$$\bar{F} = q(\bar{v} \times \bar{B}) \quad (1.13)$$

Electrons and positrons have opposite sign, so the Lorentz force pushes them apart. As the charges travel, they spiral around the magnetic field lines, emitting something akin to synchrotron radiation. This effect has been measured experimentally in beam tests [74].

Radiation created this way has a few important experimental signatures. Similarly to Askaryan radiation, the particle shower size sets the scale of coherent radiation. In an EAS, the shower is larger than the Askaryan particle disk, so emitted radiation is coherent out only to 100s of MHz, and the power drops off quickly beyond that. Radiation is very forward beamed, and due to the atmospheric refractive index, emitted in a cone-like pattern about the shower axis. Even though the cone-like beaming comes from different effects, as with Askaryan radiation, the measured radiation becomes more low frequency dominated as the viewing angle becomes more off-cone because high frequencies become incoherent, as shown in Fig. 1.11. Lastly, a radiation from an EAS will have a very distinct polarization content. The shower will be linearly polarized in a direction perpendicular to both the shower axis and the magnetic field. Near the south pole, this implies mostly horizontally polarized radio emission from cosmic ray showers.




Figure 1.11: The radio signal flux density as a function of viewing angle for various frequencies (denoted by color). This distribution is from a simulated $10^{18.4}$ eV CR with a Cherenkov angle at $\psi = 0.7$. Similar to radiation from the Askaryan effect, high frequencies fall off quickly when viewed away from the Cherenkov cone angle. Figure is taken from [75].

1.5 Introduction to the ANtarctic Impulsive Transient Antenna (ANITA)

ANITA is a NASA Long Duration Balloon (LDB) payload that is sensitive to UHE neutrinos and radiation from EAS. It is designed not to directly measure these particles, but the radio emission from their interactions with matter. Neutrinos and EAS events have distinct signatures to ANITA. As I explained in 1.4.3, EAS events will have polarization angles that are both predictable, and near horizontal. ANITA can see EAS events directly or after the emission reflects off of the ice. There are two types of direct events, a direct CR type event, and the decay of an upgoing τ lepton.

If a ν_τ undergoes a charged current interaction near the Earth’s surface, a τ lepton can escape from the Earth and shower in the air, a detection channel proposed here [76]. This looks almost exactly like a normal, cosmic ray induced EAS event, except that it appears to be “up-going,” or coming from the Earth rather than the sky. Past ANITA flights have seen two events [77, 78] that appear compatible with this signal channel, however recent simulations disfavor the standard model ν_τ hypothesis [79].

Due to the much larger solid angle, reflection off of the ice is a much more likely detection mechanism. Direct and reflected EAS events are distinguishable by a relative phase inversion (called a “polarity flip”) that happens upon reflection from the ice. Polarity is always in the same direction because EAS events are very narrowly forward-beamed, so we never see a backlobe of the emission.

Askaryan neutrino events will mostly be seen as vertically polarized because the only allowed neutrino geometries are grazing, so for Askaryan radiation emitted within the ice below ANITA, everything except for angles near the “top” of the emission cone will be totally internally reflected by the ice surface. Additionally, the vertically polarized component of the emission transmits better than the horizontally polarized component. The cone is radially polarized, so the radiation that ANITA records is approximately vertically polarized. Polarity for Askaryan neutrino events are also well defined, since we should always see the top part of the cone. In the case of more complicated scenarios, such as reflections off of water below ice shelves, or strange geometries where we see the bottom of the Cherenkov cone, polarity can be opposite the default expectation.

In addition to the science events viewable by ANITA, there are two main backgrounds, thermal and anthropogenic noise. Thermal noise comes from the combination of ice and sky in ANITA’s view, and makes up the vast majority of recorded events due to our low threshold trigger scheme. Anthropogenic noise makes up the majority of the non-thermal recorded events and comes in two distinct types, continuous waveform (CW) and impulsive noise. Narrowband CW anthropogenic noise is mainly from satellites and is dealt with in mostly in hardware through the use of notch filters, while impulsive anthropogenic noise can be very similar in appearance to signal events and must be dealt with in analysis. A cartoon representation of ANITA and all of its detection channels (including anthropogenic noise) is shown in Fig. 1.12.

ANITA flies in Antarctica primarily because ice has a long attenuation length in the




Figure 1.12: Sources ANITA is sensitive to, minus thermal noise. Picture taken from [80].

radio frequencies, but Antarctica also provides low background due to relatively few people, . The radio attenuation length has been measured at multiple places in Antarctica including the South Pole [81], Ross Ice Shelf [82], and Taylor Dome [83].

1.5.1 Previous Results

ANITA has published results from its three previous flights on a variety of topics. ANITA-I saw no candidate neutrino events [84], setting the world leading limit on the diffuse UHE neutrino flux between $\sim 10^{19.5}$ and $\sim 10^{21}$ eV, and recorded 16 CR events [85]. Much later, an anomalous CR-like event with polarity opposite expected was discovered [77]. ANITA-II improved on the ANITA-I diffuse neutrino limit [86], but saw very few CRs due to a change in the trigger scheme. In addition to the diffuse neutrino search performed on the ANITA-II data, we also produced a GRB point source search [87]. ANITA-III set an improved combined neutrino limit [88] and saw 28 CRs, along with another unusual polarity CR [78].

ANITA's main science results have been the diffuse neutrino flux limit, with the best constraints between $10^{19.5}$ and 10^{21} eV. The two anomalous events detected by ANITA have generated a lot of interest. One possible physics explanation for them is a tau lepton, created by a tau neutrino, decaying and causing an air shower outside of the Earth. However, this is a disfavored explanation [79], as both of these events come from very steep angles that are inconsistent with standard model cross sections [89]. A recent explanation put forth is the interpretation of these signals as transition radiation created when the transverse geomagnetic currents cross the air ice boundary [90].

1.5.2 ANITA-IV

The most recent flight of ANITA took place in December of 2016 for 28 days. It is the fourth flight under the ANITA name. ANITA-I flew in 2006-2007 [84], ANITA-II flew in 2007-2008 [86], and ANITA-III flew in 2014-2015 [88]. ANITA-IV (the main subject of this thesis)

retained many of the key features of the previous three ANITA payloads, with significant upgrades. The most substantial upgrades from ANITA-III are the addition of three tunable notch filters on each channel, used for rejecting continuous waveform noise [91], and the implementation of a trigger that requires a signal to be mainly linearly polarized. These upgrades will be described, along with the rest of the instrument in Chapter 2. Chapter 3 will deal with the calibration of the instrument. Next, Chapter 4 details the analysis of the ANITA-IV flight data, and finally, in Chapter 5 I discuss everything found in the analysis.

1.5.3 Complementary Experiments

ANITA is no longer the only experiment looking for neutrinos in the radio band. Most similar to ANITA are the ground-based radio detection experiments the Askaryan Radio Array (ARA) [92] and the Antarctic Ross Ice-shelf ANtenna Neutrino Array (ARIANNA) [93]. Both of these experiments consist of radio frequency (RF) antennas embedded in Antarctic ice, with ARA at the South Pole and ARIANNA on the Ross Ice Shelf. They are closer to the shower position than ANITA is, which gives them increased sensitivity over ANITA at lower energies, but their decreased detector volume makes them less sensitive than ANITA above $\sim 10^{19}$ eV.

The Giant Radio Array for Neutrino Detection (GRAND) [94], the Beamforming Elevated Array for COsmic Neutrinos (BEACON) [95] and the Taiwan Astroparticle Radiowave Observatory for Geo-synchrotron Emission (TAROGÉ) [96] also look for neutrinos in radio, however they are concerned only with ν_τ where the others listed can see all flavors.

The Probe of Extreme Multi-Messenger Astrophysics (POEMMA) [97], will also search for UHE cosmic rays and tau neutrinos, but in the optical band. It consists of two satellites in formation with one another, each with fluorescence telescopes designed to observe air Cherenkov signals.

CHAPTER 2

ANITA INSTRUMENT AND FLIGHT

In the following chapter I will describe the ANITA-IV hardware in-depth, starting from the antennas and following the signal chain. At a certain point along the signal chain there is a split between the trigger and digitizer path, both of which I will examine in detail. After detailing both the trigger and digitizer paths, I will discuss some of the peripherals, not directly related to recording signals, but important just the same.

2.1 Antennas

ANITA has used similar antennas since its inception. They are dual-polarized (horizontal and vertical), high-gain, broadband horn antennas from Antenna Research Associates, Inc. The band is ~ 180 MHz to ~ 1200 MHz. ANITA-IV consists of 48 antennas, in rings of 16, arranged in a cylindrical pattern, the arrangement is shown in Fig. 2.2. Antennas are arranged in columns, so that antennas from the top, middle, and bottom rings in each column are aligned with one another. These columns cover 22.5° each and are known as phi sectors. The beam of the antenna has a $\sim 60^\circ$ FWHM diameter on average throughout the band, which provides ~ 10 dBi gain on average. Antenna beam as a function of angle is shown in Fig. 2.1.

This antenna arrangement allows between 9 and 15 antennas to view each event (3-5 phi sectors), which allows each event to have many baselines for interferometric pointing, as well as an overall \sqrt{N} increase in SNR for the coherently summed waveforms, where N is the number of antennas used to create the coherently summed waveform. All antennas are pointed downwards 10° from the horizontal in order point their maximum response slightly below the horizon, which is the most probable location for the signals we are interested in. At 40 km above the ground (our approximate height at float), the Earth's horizon is $\sim 6^\circ$

ANITA3/4 Seavey Factory Antenna Data




Figure 2.1: Response as a function of angle from boresight in the E plane for antennas for various frequencies. Gain is referenced to the maximum gain of the antenna. Data is from Antenna Research Associates, Inc., plot is from [80].




Figure 2.2: The ANITA-IV payload pre-flight. An additional drop-down ring of PV panels was added to the bottom ring of antennas before flight.

below the horizontal.

ANITA’s antennas have dual, orthogonal polarization in order to map out the complete Stokes parameters of any recorded signal. This is important for distinguishing science events from thermal or anthropogenic noise, because we expect a high degree of linear polarization, at characteristic angles. Thermal noise should be completely unpolarized, and anthropogenic noise could be polarized in any manner. Having dual polarization requires dual feeds, and due to design constraints, the feeds must be offset by some amount, in our case 1.96 cm. This 1.96 cm offset contributes to a difference in phase center positions between the two polarizations, as well as a relative timing offset that must be corrected for. The feed offset is shown in Fig. 2.3.

The antenna size is dictated by the minimum frequency and the desired directivity. We




Figure 2.3: There is a 1.96 cm separation between the horizontally polarized and vertically polarized antenna feeds for ANITA horn antennas.

chose the minimum frequency to be 180 MHz due to constraints on the payload size, as well as to mitigate CW interference in the VHF band. We also chose a highly directional antenna with low dispersion, the quad ridge horn. The increased directivity improves signal response without increasing received noise power. Antenna temperatures are roughly ~ 130 K from the combination of ice and sky in view of the antennas. Using highly directional antennas does not come without cost; they impose the need for more antennas to cover the desired field of view. With more antennas comes more digitization channels, and an increased complexity overall.

2.2 Signal Processing

Due to the dual polarization of each antenna, we have two channels for each antenna, for a total of 96 channels. A diagram of the RF signal processing system is shown in Fig. 2.4. The horizontally polarized (Hpol) and vertically polarized (Vpol) channels are treated identically.

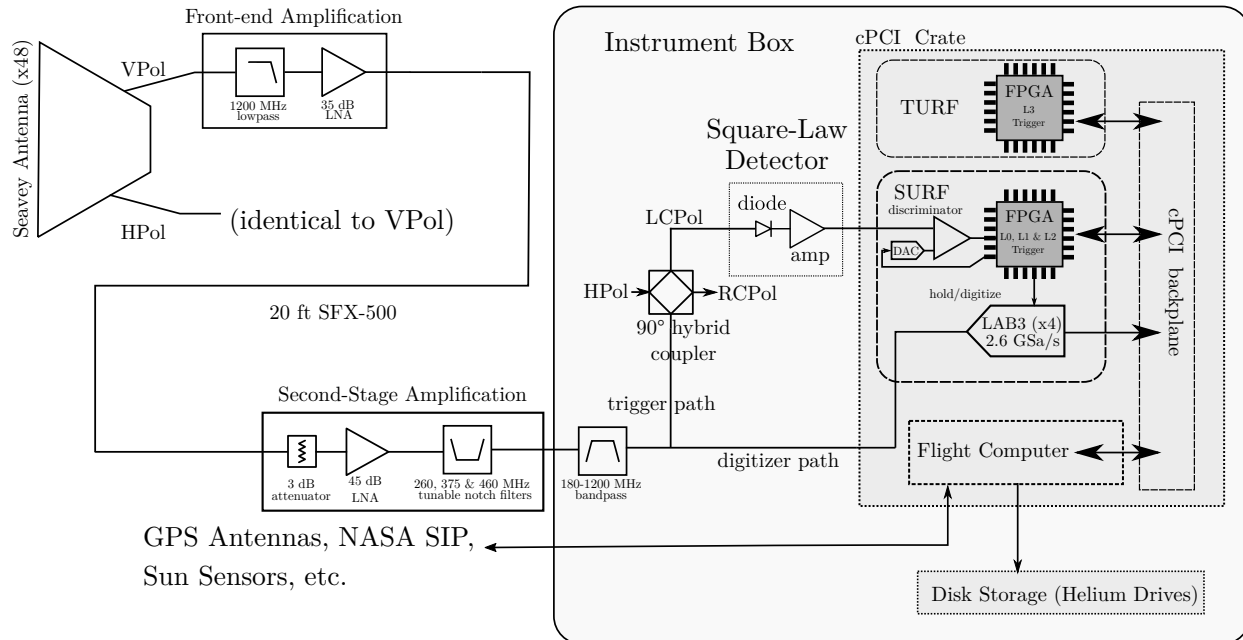


Figure 2.4: The ANITA-IV system. See text for information on each component.

Signals received on the antennas are low-pass filtered and amplified in our front end amplifiers, then piped to the second stage amplifiers where they are notch and then band-pass filtered. At this point there is a fork in the road where the received signal is split into the trigger and digitizer paths. If an event passes the triggering requirements, it is recorded onto disk storage at the end of the digitizer path. I will expand upon the entire signal processing system in the following sections.

2.2.1 AMPAs

Front end amplification is done by the custom built Antenna Mounted Pre-Amplifier (AMPA). The AMPA enclosure contains a 1200 MHz low-pass filter and a 35 dB Low Noise Amplifier (LNA). Only a low-pass filter is used because the antenna acts as a high-pass filter, and adding an additional high-pass would raise our noise temperature. The low-pass filter is at 1200 MHz because our digitizers sample at 2.6 Gs/s, but our sampling speed is not constant, so our effective Nyquist sampling frequency is ~ 1.2 GHz. First stage filtering is done in




Figure 2.5: Internals of the AMPA. Picture from [98].

the AMPA, immediately after the antenna, in order to prevent our pre-amplifier saturating due to out-of-band noise. Internals of the AMPA are pictured in Fig. 2.5. The gain and noise temperature vs frequency, as derived from measurements taken during integration in Palestine, Texas, are shown in Fig. 2.6.

2.2.2 IRFCM and TUFFs

After the AMPAs, the signal is run through about 20 ft of SFX-500 cable to the Internal Radio Frequency Conditioning Modules (IRFCM). There are four custom built IRFCMs, each of which is comprised of 45 dB LNAs and a set of three tunable notch filters per channel. This set of notch filters is on a board called the Tunable Universal Filter Frontend (TUFF). Notch filters were added for this flight to counteract CW satellite interference that significantly impacted the ANITA-III flight livetime. This satellite interference was dealt with in ANITA-III by restricting portions of the payload from triggering (“masking”).




Figure 2.6: The gain in dB vs frequency in Hz (left) and noise temperature in Kelvin vs frequency in Hz (right) for all 98 of the AMPAs. Each AMPA is in a different color, with the thick black line in the middle corresponding to the mean value. Plots taken from [99].

These satellites are geosynchronous and always in view of ANITA, so the north-facing half of the payload was almost always masked throughout ANITA-III’s flight.

ANITA-IV’s notch filters were installed with their default center frequencies at 260 MHz, 375 MHz, and 460 MHz. These frequencies were chosen to combat specific origins of CW interference, 260 MHz and 375 MHz for Department of Defense Satellites, and 460 MHz for the communications systems of Antarctic bases. Notches can be switched on and off, as well as tuned to different central frequencies during the flight, depending on the payload’s current noise environment. Notches 1 and 2 (260 MHz and 375 MHz) were on for the majority of the flight, and were tuned no more than 10 MHz in either direction from their center frequencies. Notch 3 (460 MHz) was only turned on when in view of large Antarctic bases. The addition of the TUFFs significantly decreased the fraction of the payload that had to be masked (from always $> 50\%$ of the payload to never $> 30\%$), while also keeping our deadtime low (6.7% average over the course of the flight).

After passing the the IRFCMs, the signal is band-pass filtered between 180 MHz and 1200 MHz by Lark Engineering filters, to filter out any out-of-band noise added by the amplifiers, before being split into trigger and digitizer paths.

2.3 Trigger Path

ANITA-IV was set to trigger at about 50 Hz, where each triggered event is about a 100 ns snapshot. If you were to stack all of our triggered events together, we only record about 0.0005% of the flight. This means we have to trigger very selectively on segments of time that are likely to contain interesting events. In the following sections I will detail how we have set up the trigger in order to have a low threshold without incurring too much deadtime. All of the triggering is done within the Triggering Unit for Radio Frequency (TURF).

2.3.1 Hybrids

The first step along the trigger path is the 90° hybrid coupler. Each pair of channels has one hybrid coupler assigned to it. It takes the horizontally polarized and vertically polarized channels from a single antenna and mixes them into left-circular polarization (LCP) and right-circular polarization (RCP) signals. If a signal coming into the hybrid is linearly polarized along any axis, the hybrid output will be approximately equal in LCP and RCP. Our expected science triggers are approximately linearly polarized, while thermal noise is largely unpolarized, and CW from satellites is elliptically polarized, so we can use this coincidence to restrict our triggering to only the most likely signals.

2.3.2 L0 Trigger

In order to restrict triggering to somewhat impulsive signals we use a square-law power detector. We use a tunnel diode, amplify its output, transform it from a single-ended to differential pair, and feed that into an FPGA comparator circuit. The combination of diode, amplifier and transformer make up the SURF High-Occupancy RF Trigger unit (SHORT). Hybrids and SHORTs can be seen, along with power distribution systems and some cabling, in Fig. 2.7. This is done for both outputs of all hybrids simultaneously (96 channels), and if




Figure 2.7: Trigger path assembly. The blue rectangular pieces are the 90° hybrid couplers, below them are the RF-tight boxes housing the SHORTs. A little more than half of the hybrids have been installed, in the end this will be symmetrical. Also pictured are power distribution systems on the right, along with cabling. Picture was taken during assembly in Hawaii and is used courtesy of John Russell.

any of these outputs exceed the threshold required by the comparator circuit, a zeroth-level (L0) trigger is issued. Trigger thresholds are automatically adjusted in real time to keep the trigger rate at approximately 50 Hz. Zeroth-level triggers are entirely dependent on the power in the signal. Thresholds for the zeroth-level trigger were adjusted to keep it at the target rate, which was between 5 MHz and 6 MHz.

2.3.3 L1 Trigger

If a zeroth-level trigger is issued, a check for the first-level (L1) trigger is initiated. A first-level trigger is only issued if both the LCP and RCP outputs of the hybrid that issued the L0 trigger exceed the required threshold within 4 ns of one another. This requirement enforces linear polarization along any axis.

2.3.4 L2 Trigger

A second-level (L2) trigger is checked for if a first-level trigger is issued. It is imposed at the level of the phi sector and is satisfied by the coincidence of two or more channels. A second-level trigger begins by delaying the signal from the middle and bottom rings by 4 ns, in order to bias against triggering on signals where the top and middle or top and bottom issue L1 triggers at the same time. The size of each coincidence window depends on which ring issued the first L1 trigger. They are set to 12 ns for the bottom to top ring, 8 ns for the middle to top ring, and 4 ns for the bottom to middle ring. Windows are set this way to preferentially trigger on signals from below the horizon.

2.3.5 Global Trigger

A third-level, L3, or global trigger is issued if two L2 triggers are issued in adjacent phi sectors within 10 ns of one another. If a global trigger is issued, the digitized signal is read out and recorded, unless the all four levels of the digitizer buffer is full. We tune the thresholds so that the global trigger rate is approximately 50 Hz to prevent incurring too much deadtime.

2.3.6 Minimum-bias Triggers

In order to characterize our noise environment, we take a set of triggers forced by the software at regular intervals over the course of the flight, called minimum-bias triggers. This set of events is very different than the rest because it is not “biased” in the sense that it has not passed our triggering criteria. All events taken this way are marked as a separate subset of triggers by our software. ANITA records these triggers at approximately 3 Hz throughout the entire flight.

2.4 Digitizer Path

The signals propagated through the instrument must be digitized and saved to disk before they can later be analyzed. This is done using an array of custom analog to digital converter (ADC) application specific integrated circuits (ASIC). ANITA-IV continues to use the Large Analog Bandwidth Recorded And Digital Ordered Readout (LABRADOR) or LAB3 chips [100], that have been used on all ANITA flights so far. The LAB3 chip has 12 bits (one of which is thrown out to conserve space) and samples at 2.6 GS/s. Each LAB3 chip is made up of nine channels, eight RF analog input channels and one clock channel. Each RF channel has a 260 long switched-capacitor array (SCA) ADC, resulting in 100 ns window size. Timing between samples is uneven, which leads to a frequency response that needs to be calibrated for each chip individually. It is important to note that the LAB3 chip has a diminished bandwidth from the quoted 180-1200 MHz ANITA bandwidth. It has a 3 dB point at around 900 MHz, 300 MHz lower than the desired frequency cutoff.

We bundle four LAB3 chips to each Sampling Unit for RF (SURF) board. The SURF board allows four LAB3 chips to observe and measure the same 33.3 MHz clock signal on their ninth (clock) channels. The relative phase difference between the measured clock signal between the SURFs allows us to correct for any trigger timing jitter.

2.5 Orientation

The ANITA-IV payload is attached to a free-floating, free-rotating balloon, whose direction, speed and orientation are all at the mercy of the winds. It's very important for our analysis that we know have accurate attitude information at all times. In order to characterize this properly we have multiple GPS systems, as well as sun sensors and magnetometers as back-up, redundant systems.




Figure 2.8: Payload heading in degrees as calculated by ADU5-A (blue) and sun sensor (red) for one run of flight data, taken from [101]. They track very closely, indicating that the sun sensors would be a viable backup for heading information if the GPS antennas failed.

2.5.1 GPS

Our GPS information comes from three independent GPS antennas. There are two sets of four ADU5 GPS antennas and one G12 antenna, for a total of nine antennas, all read out at 1 Hz. The ADU5 arrays independently report the longitude, latitude, altitude, heading, pitch and roll of the payload. The G12 antenna reports only position and velocity information. These combined GPS systems report the position and orientation of ANITA to a high accuracy ($< 0.5^\circ$).

2.5.2 Sun Sensors

If the GPS antennas were to fail, we have two backup systems that could be used to determine orientation of location. One of these is our system of sun sensors. The Sun is always in our view, because it is above the horizon in the Austral summer. With accurate timing information, which we have from several redundant systems, we can use our knowledge of the Sun’s position in the sky to determine heading information. Only our payload heading can be determined using the sun sensors, because the Sun is a single source. We did not have to use the sun sensors because our GPS antennas remained on-line the entire flight, but we did validate this method using flight data, as demonstrated in Fig. 2.8.

2.5.3 Magnetometer

Another backup to our GPS system was a magnetometer we had attached to the payload. It can be used to measure our position by measuring the local magnetic field. In order to do this properly, we would need to use an accurate model of the Earth's magnetic field and assume an elevation for our payload. We also did not use the magnetometer for this purpose because the our GPS antennas did not go out.

2.6 Flight Computer

The ANITA flight computer coordinates the whole operation. It gathers data from all sources, writes that data to disk, and sends some portion of that data to those on the ground monitoring, while also acting on any commands sent to it from the ground. It is a single board computer running Fedora Linux, housed alongside the TURF and SURF boards. The flight computer works via a series of daemons, with most of the work done by the acquisition daemon (`Acqd`). `Acqd` gather data from the SURF and TURF boards and uses that data to automatically change the triggering thresholds and phi masking. Examples of other daemons include those in charge of monitoring the GPS (`GPSd`), the hard disks (`Monitord`) and the prioritizer (`Prioritizerd`). For more information about the flight software consult Ben Strutt's thesis [102].

2.6.1 Data Storage

Data storage is an important concern for ANITA because we cannot transmit all of our data in real time. After our flight is concluded, a team of people has to fly out to the crash site and collect the hard drives. That means our hard drives must be large enough to ensure we don't run out of space (because we can't add any during the flight), be robust to temperature variations, be able to survive the end of a flight, and have redundancy,

should something happen to part of our system. Our storage system consists of two different methods. First we have two 8 TB helium-filled spinning disk drives. Both of these were written to simultaneously with the same information. Our second data storage method was an array of six 1 TB Solid State Drives (SSD) designed by National Taiwan University (NTU). All of our data survived the flight and was successfully recovered.

2.6.2 Telemetry

During flight it is important we have at least some downlink ability to receive data for diagnostic purposes, and uplink ability for sending commands (such as TUFF tuning). We have five different ways of accomplishing data transfer during flight. The first of these is the Line-Of-Sight (LOS) transmitter. As the name implies, this is only useful when we can directly see the payload. It transmits data the most quickly of these telemetry systems. The second is the Iridium Openport link, a satellite transmitter. Third and fourth are the NASA Tracking and Data Relay Satellite System (TDRSS), with both fast and slow listed as different systems. Last, both in this list and in terms of data transmission speed is the Iridium low rate. The combination of these systems allowed our ground based team to monitor flight operations and send commands throughout the flight.

2.6.3 Prioritizer

There are a few conceivable scenarios in which we recover no data. Either all of our drives could fail, or the payload could land somewhere unrecoverable, such as a deep crevasse or in the ocean. To mitigate the damage of this happening, we have a Graphical Processing Unit (GPU) based “prioritizer” categorize each incoming event based on its priority. The priority value of an event is an integer ranging from 1 to 9, where 1 is the highest priority and 9 is the lowest. The range for normal events is 1-6, while 7, 8 and 9 have special meanings. A value of 7 is assigned with the queue is overly long, 8 indicates an event with strong CW

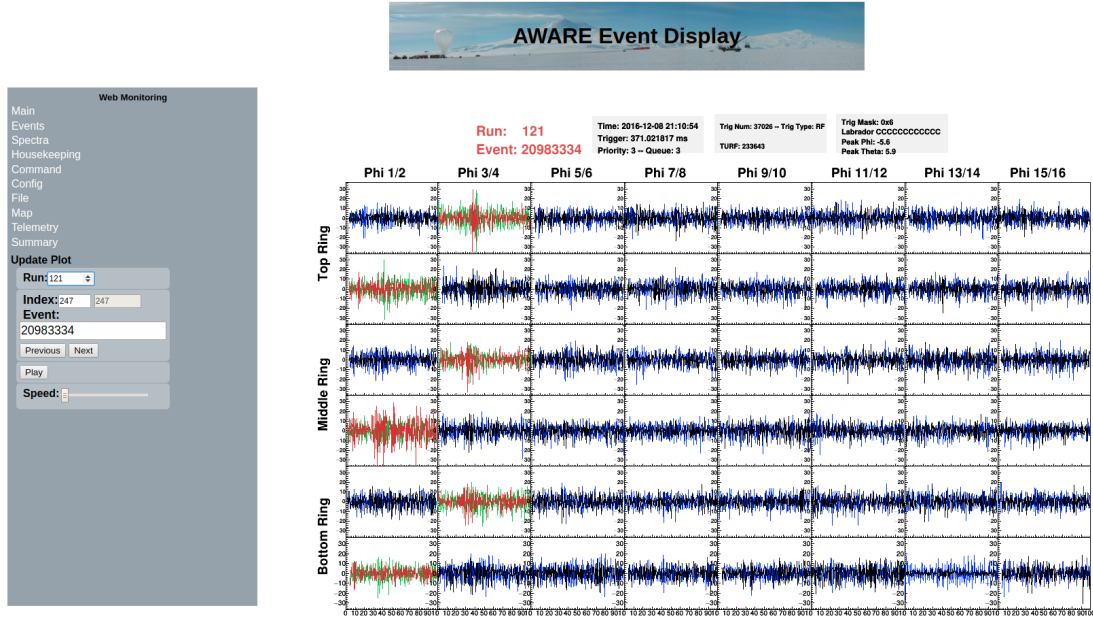


Figure 2.9: The AWARE program for monitoring ANITA during flight. Buttons on the left change what information is displayed. The current display shows all 96 waveforms for an example RF-triggered event.

contamination, and 9 is reserved for SURF saturation. Priority values are assigned based on interferometric image peaks and Hilbert envelope peaks for each event. For further details of prioritizer design and operation see [102]. Because the instrument was recovered, prioritizer values were not used in analysis.

2.6.4 Ground Based Monitoring

ANITA-IV had two event display programs for the purpose of monitoring the instrument while it was in flight. AWARE is the older of the two programs, written by Ryan Nichol at UCL. It is a website that reads in JSON files. An example plot from the program is shown in Fig. 2.9. The newer of the two programs is called WebANITA, written by Peng Cao at University of Delaware. WebANITA is written in the Python Flask framework, and uses SQLAlchemy to query databases and display graphs. For more information about the development and use of WebANITA see [103].



Figure 2.10: The ANITA-IV payload pictured with inflated balloon during launch. The balloon will inflate to a much greater size as it rises and the pressure drops. During launch, the ANITA payload is held aloft by a truck-mounted crane. Launch must take place during a day with perfect conditions (still winds, clear skies), because the balloon is delicate and incredibly expensive. ANITA-IV had one launch attempt scrubbed due to weather before a successful flight.

2.7 Balloon and Flight

The ANITA payload is attached to a NASA long-duration balloon that must be handled very carefully before flight. It holds 34 million cubic liters of Helium which brings the instrument to roughly 40 km above the Antarctic continent at float altitudes. The balloon is rated to survive for up to 60 days. We were in no danger of running up onto this limit, as our flight lasted around 28 days. The balloon expands as it ascends, until it equilibrates and reaches float altitudes. Once at float, the balloon and instrument are completely out of our control and go wherever the winds take them. ANITA is pictured with an the balloon in Fig. 2.10.

The Antarctic Austral summer brings with it circumpolar winds, which means an approximately circular orbit around the South Pole for our instrument. The necessity of recovering




Figure 2.11: ANITA-IV after terminating the flight. Picture taken by and used courtesy of Christian Miki. Crumpled antennas can be seen on the bottom ring of the instrument. This photo was taken during recovery of the data, the remote location of which necessitated use of the pictured plane.

the data means we must terminate the flight if ANITA breaks from its circular trajectory and head northwards, towards the sea. When a flight is terminated, the payload detaches from the balloon and plummets towards the ground. A parachute ejects and slows the payload down, but the bottom two rings of antennas still act as a crumple zone and are fairly destroyed. ANITA-IV was terminated near the South Pole, in a relatively easily accessible place, and the data was recovered successfully shortly after. A photo of ANITA post-crash is shown in Fig. 2.11. The flight path of the payload is plotted in Fig. 2.12.

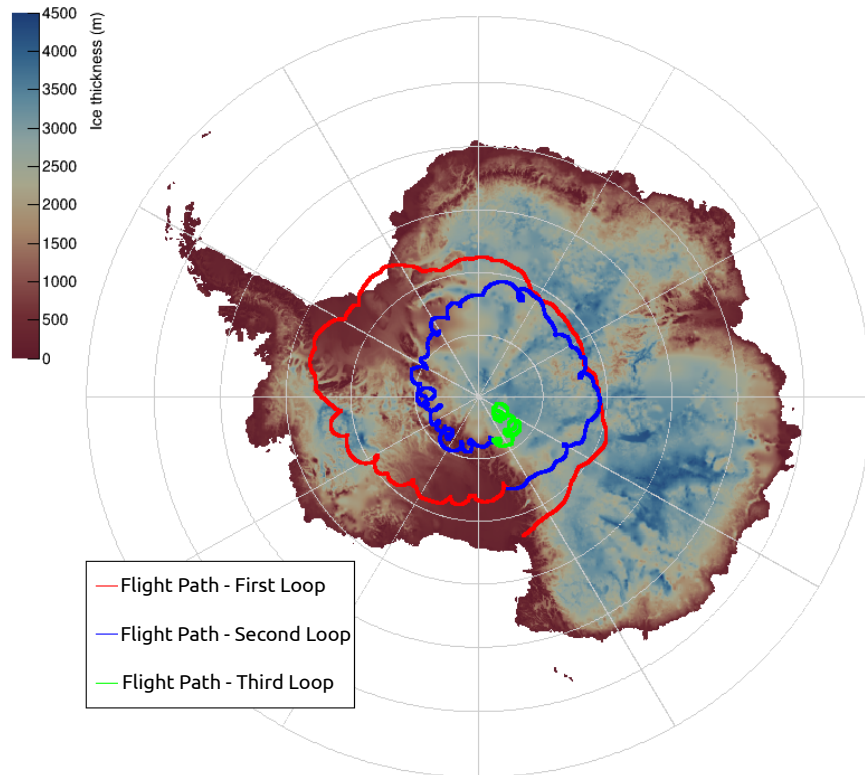


Figure 2.12: The ANITA-IV flight path. The first pass around the continent is shown in red, the second pass is shown in blue, and the third pass, where we got stuck in a holding pattern and had to terminate the flight, is shown in green. The Antarctic background map is colored by ice thickness.

CHAPTER 3

CALIBRATION

The calibration process is essential for translating the recorded voltages back to the incident electric fields that produced them. Nearly every component must be calibrated or measured before it is integrated with the rest of the instrument. As a result, the calibration process is highly involved and time-consuming. There are many considerations that need to be taken into account to ensure calibration is done correctly. The sampling time base must be measured and its non-uniformity corrected for post-digitization. Relative delays between channels are corrected for, in order to have accurate phase information. The reported count values for the LAB3 ADC must be translated to real units (voltage). Highly precise locations of the antenna phase centers need to be determined in order to have accurate pointing and signal reconstruction. Magnitude and phase responses are meticulously determined for each channel in order to deconvolve the system response out of the measured signals. The trigger efficiency needs to be quantified in order to report science results accurately.

We use multiple data sets in order to do all of these calibrations. There are data sets recorded before ANITA flew, in Texas, Hawaii, Antarctica, and other places, but there are also pieces of calibration we could only do once the entire instrument was assembled and on-line, something that only happened during the actual flight. For these purposes we had calibration pulser stations at WAIS (West Antarctic Ice Sheet) and the LDB (Long Duration Ballooning) Facility. We also had two small balloon payloads, known as HiCal-2A and HiCal-2B, follow ANITA and pulse. Using these various datasets we were able to calibrate the entire instrument.

3.1 LABRADOR Timing Calibration

Before any real science can be done, we must calibrate the LABRADOR digitizer chips. LAB3 chips use a ring oscillator to control timing of when to successively connect and disconnect an input line to cells in a Switched Capacitor Array (SCA) with 9 rows and 260 columns, leaving a stored charge on each capacitor related to the input voltage. Each LAB3 chip has two phases at which it does this successive sampling, Ripple Carry Out (RCO) 0 and 1. When sampling reaches a certain capacitor, it loops back to the beginning, which takes a finite time. The RCO functions as a write pointer, and the stored charge it deposits in capacitors it samples is left until a trigger is issued, upon which each capacitor is digitized by an ADC and read out by the flight computer.

The LAB3 chip has many timing related quantities that need to be calibrated, including sample-to-sample timing, wrap around timing for the write pointer, temperature corrections, and calibration of the overall relative delays due to cabling from the antennas.

The LAB chips installed in ANITA-IV were recovered from the ANITA-III instrument and re-used, so all the calibration of the aforementioned quantities, save for the relative cable delays from the antennas, was already complete. I will discuss them here in a cursory manner, but for more information see two theses on ANITA-III [80, 102].

3.1.1 Sample-to-sample Timing

The digitizers have an uneven sample-to-sample time differences, δt . These were determined by digitizing many sine waves of random phase. Given enough recorded events, the probability of a sine wave crossing zero between two samples is directly proportional to the time between them. The mean δt for each sample is $\frac{1}{2.6}$ ns, and the uncertainty on each sample δt is approximately 3.3 ps [102].

3.1.2 Write Pointer Wrap Around Timing

Once the write pointer passes the 255th sample it loops back to the 0th sample, while the remaining four capacitors are used to fill the wrap around time. The timing is dependent on the RCO phase, so there are two separate calibration constants for each chip. The wrap around time was also measured using sine wave data, where the difference in phase of a fitted sine wave between sample 255 and sample 0 gives the wrap around time. There is a significant difference in wraparound time depending on the RCO phase, on the order of 0.45 ns, with the mean wrap around time taking somewhere between 1 and 1.8 ns [102].

3.1.3 Temperature Correction

Sampling frequencies change with changing temperature. This is a real worry for ANITA, because we are in a balloon that will warm and cool depending on payload rotation, time of day, and latitude. The temperature was corrected for by using one single 17 hour run of sine wave data, where the instrument box temperature was carefully monitored. Temperature variations for the RCO wrap around time were on $\mathcal{O}(10\%)$, while sample to sample dt variations were much smaller, $\mathcal{O}(0.1\%)$ [102].

3.1.4 Channel to Channel Delays

While in Antarctica we calibrated the relative delays between channels. This is the only digitizer calibration that needed to be done separately for ANITA-IV. The difference in channel to channel delays here is caused mainly by different cable lengths. For these measurements we used a 16-way splitter hooked up to a pulse generator to plumb the same pulse into 16 channels at a time. We also had a 2-way splitter before the 16-way splitter so we could keep channel BV1 (Bottom Vertical polarization 1) connected at all times to act as a reference channel. The cables from the splitters to the channel inputs at the AMPAs were all various lengths, but the group delay and S-parameter information for each one was carefully

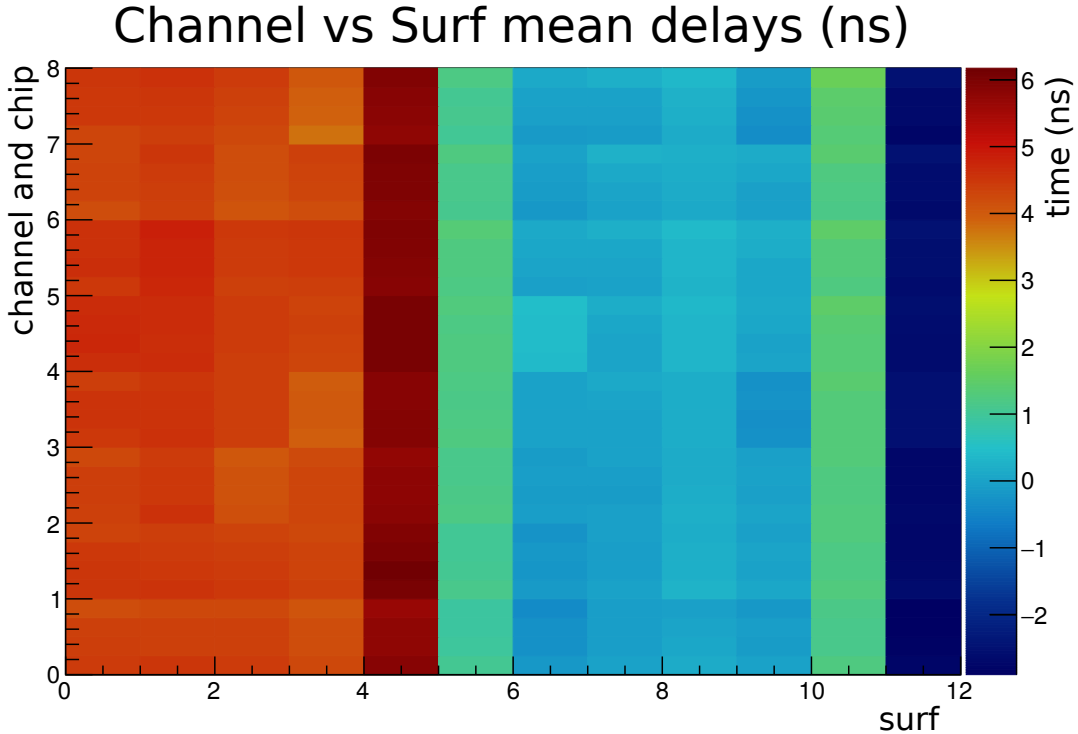


Figure 3.1: This colormap shows the mean calculated time delays compared to the reference channel for each channel, broken up by LAB chip. The color axis is the time delay in ns. The X-axis is the SURF and the Y-axis is the channel and chip on that SURF. The reference channel is SURF 6, channel 3. Delays vary by SURF because each SURF is cabled differently.

measured and recorded and disseminated via ANITA internal notes [104, 105]. We rotated cables in and out of this setup six times (so as to record all 96 channels). This setup was also used to calibrate the impulse responses and ADC count to voltage corrections.

The channel to channel delays were calculated using the pulses from the pulse generator. For each event, pulses from each channel were correlated against the reference channel. The time delay required for maximum correlation is the relative delay between channels. We only look at positive correlation, as all pulses should have the same polarity. This was done for many events, and the mean of all the correlations is shown in Fig. 3.1. The calculated time delays are applied in software to the digitized waveforms before doing anything else.

One channel (BH13) showed a significant spread in the maximum correlation time with the reference channel. This was due to its impulse response being significantly different

from all other channels. In order to correct for this channel I had to restrict the range of correlation manually.

3.2 ADC to Voltage

The LAB3 chips use SCAs to store sequential snapshots of the incoming signal. These snapshots are not stored as voltages directly. They are instead digitized and stored as “counts” on the ADC (analog to digital converter). An ADC count is determined by comparing the voltage stored on each capacitor to a ramp signal that stops a Gray code counter at the time of comparator activation. The digitized signal is a measurement of the time required for the ramp signal to reach parity with the voltage stored in the SCA bin. There are variations in the conversion factor from ADC count to voltage from chip to chip. ADC to voltage calibrations from between ANITA-III and ANITA-IV should have been the same because we used the same LAB3 chips, but because the different impulse response of channel BH13 was never noticed or corrected for in ANITA-III, I decided to re-do the calibration.

We used the same setup detailed in the channel to channel delays calibration section to calibrate the ADC to voltage conversion. The pulse was measured before the SURF input for the reference channel and then again for each channel from the LABs using the ANITA software suite. We compared the peak to peak voltage of the two measurements, and normalized all signals recorded on the LAB to the reference channel. Fig 3.2 shows there was significant variation in between each ADC. We determined the variation was in the ADC because all amplifiers in the system were already measured and determined to have approximately the same gain. This is just a first order voltage calibration, which is why we compare the pulse at the SURF input to the recorded pulse. The system impulse response is necessary to compare the pulse at the system input to the recorded pulse.

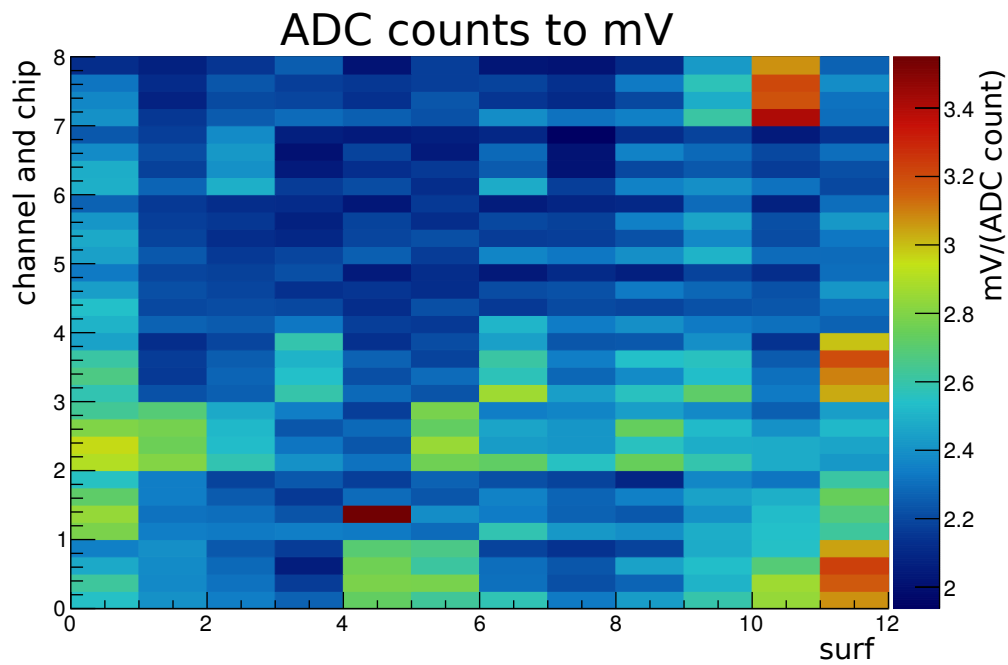


Figure 3.2: This colormap shows the correction from ADC counts to mV for the input calibration signal, broken up by LAB chip. The X-axis is the SURF and the Y-axis is the channel and chip on that SURF. Calibration constants vary between 2 and 3.5 mV per ADC count between channels.

3.3 System Impulse Response

The difference between the incident electric field and the output of the digitizers is the system impulse response. An impulse response is defined as what your system would record if it were given a delta function impulse input. The system response comes mainly from two components, the signal chain and the antenna. The signal chain can be further broken down into individual components, including AMPAs, SURFs, TUFFs and cables. The full system impulse response is a convolution of the antenna, the signal chain, and whichever TUFF notches are active, and can be represented as:

$$h_{full}(t) = \sqrt{\frac{50\Omega}{377\Omega}} (h_{ant}(t) * h_{sig}(t) * h_{TUFF}(t)) \quad (3.1)$$

where $h(t)$ represents an impulse response, $\sqrt{\frac{50\Omega}{377\Omega}}$ is the coupling of free space to the 50 Ω impedance of the instrument, $*$ is the convolution operation, and the subscripts *ant*, *sig* and *TUFF* represent the antenna, signal chain and TUFF notches, respectively. h_{full} has units of meters, from the antenna effective height. The incident electric field, E_0 is then turned into a measured voltage, V_{meas} by the equation:

$$V_{meas}(t) = h_{full}(t) * E_0(t) \quad (3.2)$$

In our system, the dominant effect on signal shape is due to the group delay induced by the the Lark Engineering filters and the TUFF notches.

3.3.1 Signal Chain Responses

The signal chain impulse response was determined using the same dataset used to determine the ADC to voltage calibration as well as the relative channel to channel delays. The same broadband impulse was injected directly into the AMPA for each channel through a series




Figure 3.3: An example signal chain impulse response is shown. The Y-axis is unitless. This is the bottom Vpol channel 7 (BV07), but is representative of every channel besides BH13 and TH03.

of cables (all of which had to be measured), bypassing the antenna. This same signal was also measured at the source via an oscilloscope. The oscilloscope measured input signal is necessary as a reference to compare to, to determine the full complex transfer function. After the extra cables that are not actually part of the instrument are deconvolved out, we are left with an input signal measured at the oscilloscope and an output signal measured at the SURF. Deconvolving out the input signal from the output signal gives the signal chain impulse response. An example signal chain impulse response is shown in Fig. 3.3.

3.3.2 Antenna Responses

The impulse response of the antennas is a difficult to obtain accurately because both the gain and phase responses vary as a function of angle. There are few data sets that were useful for making the antenna impulse response. The manufacturer, Antenna Research Associates,




Figure 3.4: The average antenna impulse response is shown. Compared with the signal chains, the antennas had much more similar impulse responses.

made measurements of gain as a function of angle for some of the antennas. There are several sets of data taken from ground based configurations in Hawaii and Palestine. We also could potentially use the data taken with calibration pulsers in flight.

An antenna is characterized by its effective height, h_{eff} , which relates the incident electric field to the transmitted voltage. Effective height is found by measuring antenna S_{11} , which is the amount of input power transmitted out of the antenna as a function of frequency, and gain pattern, which is where that power is transmitted as a function of angle. The full complex antenna response for the on-boresight angle was computed by Ben Rotter during ANITA-III, see [80] for more detail. The average antenna impulse response is shown in Fig. 3.4.

There are two things lacking from our impulse response measurements that could be improved upon in future flights. The bigger problem of the two is our off-axis measurements. Gain response should be maximum on-boresight and drop-off quickly after 22.5 degrees off-

boresight. Phase response also varies with angle. Both gain and phase response patterns will change shape with frequency. We don't have sufficient data to accurately correct for off-axis effects at all frequencies, so I ended up using no correction at all. The second problem is leakage from one polarization to another. From measurements done in Palestine, this is a small effect, but still causes uncertainties in polarimetry that could be improved upon.

Our antenna responses convolved with the signal chain responses makes up the base case of the full impulse responses. Impulse responses change as different TUFF configurations are convolved in, but the base case for each channel is always the same. All of the Hpol and Vpol responses are plotted on top of one another in Fig. 3.5 and Fig. 3.6, respectively. Hpol and Vpol response power spectra are plotted on top of one another in Fig. 3.7 and Fig. 3.8, respectively. The Vpol impulses are all remarkably similar. The Hpol impulses have two channels that are different enough from the average to be noticeable. Channels TH03 and BH13 differ significantly from the rest due to different filter responses, and are plotted in red. These differences are accounted for in analysis.

3.3.3 TUFF Responses

The TUFFs had notches nominally at 260 MHz, 375 MHz and 460 MHz. They were tuned to slightly different frequencies throughout the flight as well as turned on and off. Tuning them or turning them on or off changes the impulse response. The total impulse response for the entire instrument is the antenna convolved with the signal chain convolved with the TUFF response. Specific impulse responses are chosen in software based on the time of the flight and the known TUFF configuration at that point.

Due to a lack of both time and foresight, we did not record the impulse responses of every TUFF and every configuration. Instead, we modeled the TUFF impulse responses. A total of seven different TUFF responses were used at any point throughout the flight. Fig. 3.9 shows all of the responses that we used. The most commonly used notch configuration was

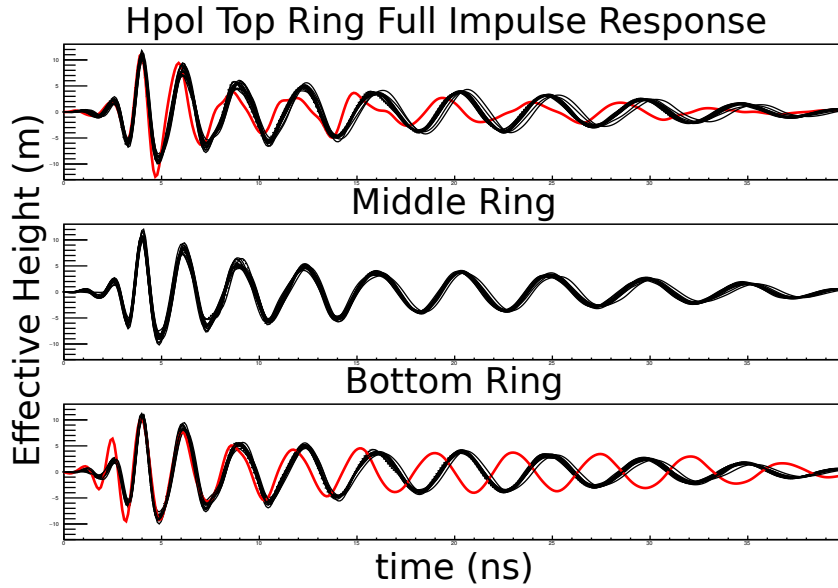


Figure 3.5: Shown are the full impulse responses (signal chain convolved with antennas) for all horizontally-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring. TH03 and BH13 are shown in red because they had significantly different impulse responses from the average response. The impulse response differences are caused by different filters along the signal chain.




Figure 3.6: Shown are the full impulse responses (signal chain convolved with antennas) for all vertically-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring. All of the vertical channels had very similar impulse responses, although one of the channels (BV14) died during flight.

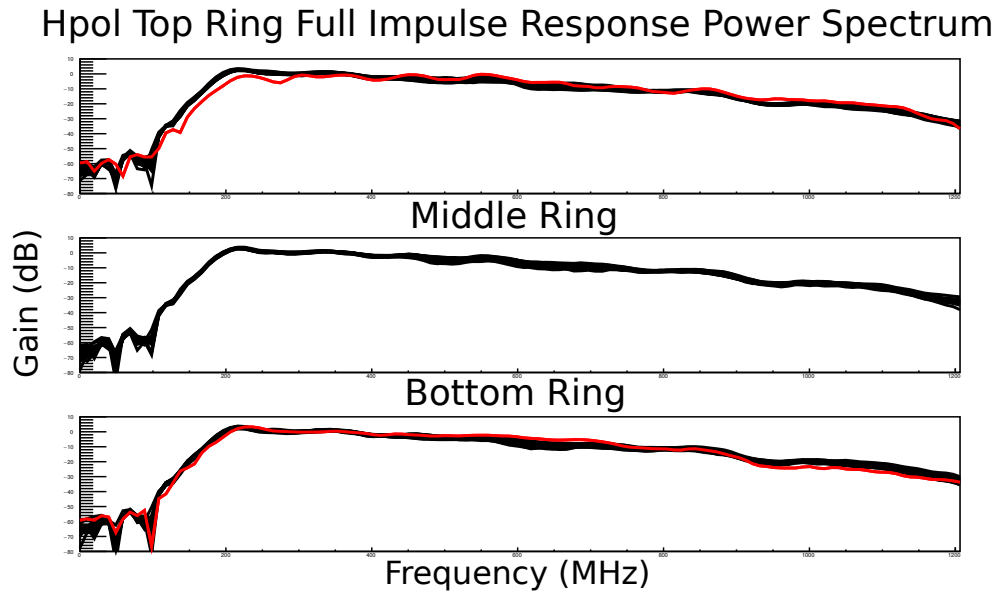


Figure 3.7: Shown are the power spectra of the full impulse responses (signal chain convolved with antennas) for all horizontally-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring. TH03 and BH13 are in red because they had different impulse responses from the average response.

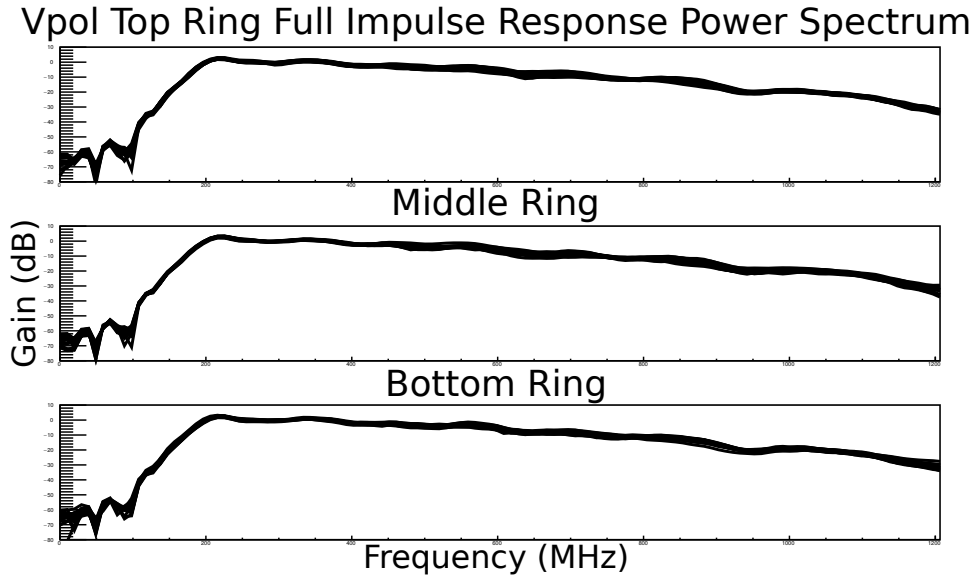


Figure 3.8: Shown are the power spectra of the full impulse responses (signal chain convolved with antennas) for all vertically-polarized channels, with all notches off. The top panel is the top ring of antennas, the middle panel is the middle ring and the bottom panel is the bottom ring.

Notch 1 (MHz)	Notch 2 (MHz)	Notch 3 (MHz)	Fraction of Flight
260	OFF	OFF	4.1×10^{-4}
250	375	OFF	1.3×10^{-4}
260	365	OFF	1.8×10^{-1}
260	375	OFF	5.3×10^{-1}
260	385	OFF	1.2×10^{-1}
260	OFF	460	4.8×10^{-4}
260	375	460	1.7×10^{-1}

Table 3.1: TUFF configurations and the fraction of the flight they were active. The configurations with very low fractions of flight time were failed experiments. For instance, we turned off the 2nd notch twice, and trigger rates raised so drastically that we immediately turned it back on.

notches at 260 MHz and 375 MHz for 53% of the flight. Table 3.1 shows the fraction of the flight each TUFF configuration was enabled for.

If TUFFs or a TUFF-like system is used in the next flight of ANITA, it is essential we take impulse response data for each TUFF and configuration separately. The responses of channels BH13 and TH03 demonstrate that filter response can vary dramatically and unexpectedly from what is expected.

3.4 Calibration Pulses

ANITA-IV had an extensive set of calibration pulses. There were three completely separate datasets of calibration pulses, all potentially useful for different things. Two of the pulser stations were ground based and there were two balloon-borne calibration pulsers that followed the payload. The locations of these two stations, along with the ANITA-IV flight path are shown in Fig. 3.10. All pulsers transmitted high power, broadband pulses meant to mimic science signals. Pulses were used for many different purposes, from phase center identification to ice studies.




Figure 3.9: Magnitude responses for all seven TUFF configurations that were employed during the flight are shown. Magnitude is in dB and frequency is in Hz. Notch 1 (nominal frequency 260) was almost always on throughout the flight. Notch 2 (nominal frequency 375) was on most of the flight, and had to be tuned. Notch 3 (nominal frequency 460) only had to be turned on near large Antarctic bases such as McMurdo and South Pole Station.

3.4.1 LDB

Pulsing from LDB was done immediately after launch. We used an antenna very similar to the Antenna Research Associates antennas on the payload to transmit pulses. Due to the extremely noisy environment surrounding McMurdo, we had to pulse at a high power and high rate to ensure that we were triggering on these pulses. This hurt our ability to identify pulses because they are not synced to the GPS second. It's also hard to disentangle these pulses from other anthropogenic sources in the area in analysis because both the pulses and the anthropogenic noise are coming from the same direction. Overall, these pulses were not as useful in calibration as the WAIS or HiCal pulses.

3.4.2 WAIS

Two ANITA team members were sent to WAIS divide to pulse shortly after launch [106]. They ran a pulsing station together using an Antenna Research Associates horn as the




Figure 3.10: The ANITA-IV flight path is shown along with the locations of the pulsing stations. The pulsing station at LDB is the magenta star and the pulsing station at WAIS is the cyan star. ANITA flew almost directly over WAIS, allowing us to collect a large amount of high-quality calibration data.




Figure 3.11: A picture of the WAIS pulsing station with all components labeled. The LDB pulser setup was much the same. Picture taken from an ANITA internal note [106].

transmitting antenna. The Hpol and Vpol feeds of the pulsing antenna were each connected to a different pulser, with slightly different frequency spectra. Both the Hpol and Vpol channels of the antenna pulsed at 1 Hz, with Hpol and Vpol separated by $10 \mu\text{s}$. The pulsing team adjusted the antenna to better point at the payload every 45 minutes to 1 hour. A picture of the WAIS pulsing station is shown in Fig. 3.11, and an example WAIS pulse is shown in Fig. 3.12.

Because of a favorable flight pattern, WAIS was in view of the instrument for multiple days, which allowed us to collect a lot of useful calibration data. We were able to probe a wide range of θ receiving angle, all of ϕ receiving angle, and multiple polarization angles for transmitted signals. In all, over 200 thousand pulses were identified, and this was our most useful data set for instrument calibration.




Figure 3.12: An example of a vertically polarized WAIS calibration pulse, with the phase response of the instrument removed. It is extremely impulsive, and very linearly polarized.

3.4.3 HiCal

HiCal is a balloon-borne calibration source consisting of a piezoelectric spark generator driven by a small DC motor attached to a bicone antenna. ANITA-IV had two HiCal modules trail it as calibration sources, 2A and 2B. HiCal periodically reported timestamps and GPS locations via Iridium link, which were used to tag pulses and remove them from the signal sample in analysis. The most useful feature of HiCal is the ability to see both direct and reflected pulses. A direct HiCal pulse is shown in Fig. 3.13. This allows a probe of Antarctic ice Fresnel coefficients at a variety of angles and frequencies, as HiCal’s pulse is broadband. For more information about HiCal see [107, 108].

3.5 Phase Centers

In order to have the pointing resolution we require, we need very precise phase center positions. Photogrammetry gives us a good first guess, but we can refine the phase center positions using WAIS pulses. Phase center errors are by far the dominant error on timing, as inter-SURF timing resolution is on the order of 10 ps.




Figure 3.13: An example of a HiCal calibration pulse, with the phase response of the instrument removed. It is much less impulsive than WAIS.

3.5.1 Photogrammetry

Initial measurements of the positions of the antenna locations were done by analyzing a series of standardized photographs of the payload taken at many angles around the payload. These photographs were also used to determine placement of the instrument box and GPS antennas. Once the antenna positions are determined from photos, the phase centers are assumed to be ~ 20 cm inwards from the face of the antenna. The Vpol phase centers are assumed to be slightly closer to the face of the antenna than the Hpol phase centers, due to the relative positions of the Hpol and Vpol feeds (1.96 cm apart). A 3D model of the payload constructed using the photogrammetry photos is pictured in Fig. 3.14.

3.5.2 Optimization Using Calibration Pulses

In practice, the estimates of phase center position from photogrammetry need corrections on order of centimeters per antenna. To determine more precise phase center positions we use the roughly 200 thousand calibration pulses from WAIS, as they were taken with precise timing, from a well known location, at all angles of payload rotation.




Figure 3.14: 3D model of ANITA-IV, as produced by PhotoModeler Scanner 2015 software. Image taken from [109].

The first step of this optimization is selecting pulses. This is relatively easily achieved due to pulses being sent at a precise time, twice a second, once exactly on the GPS second, and once 10 μs after. Each pulse is sent through a different feed (Hpol or Vpol), so there is one Hpol and one Vpol pulse per second. Pulses are initially selected by taking all RF triggers (not minimum-bias triggers) that fall within 1200 ns of the expected trigger time. The expected trigger time is calculated by propagating an RF pulse from the pulser to the payload. We chose a 1200 ns requirement by plotting $t_{triggered} - t_{calculated}$ for all suspected WAIS pulses, fitting a Gaussian to the resulting distribution, and including all pulses 2σ away from the peak. Once we had more precise phase centers, we validated this method of selection by imposing a pointing requirement, and found it removed less than 1% of events tagged as WAIS.

With our selection of WAIS events in hand, we look at $\Delta t_{expected}$ and $\Delta t_{measured}$ for each antenna pair in the three phi sectors most closely facing the WAIS pulser. $\Delta t_{expected}$ is the the time difference in received signals expected from payload geometry for an antenna pair. $\Delta t_{measured}$ is the time of maximum cross-correlation of the received signal for an antenna pair. The correlation is done after upsampling 40x by frequency padding. Our main variable in the following calculations is $\Delta T = \Delta t_{measured} - \Delta t_{expected}$.

The next step is to pick an objective function to minimize. I tried a few different approaches but found that what worked best was minimizing the mean ΔT for nearest neighbor antenna pairs. This approach leads to an objective function with 192 free parameters (48 each of $\Delta CableDelay$, Δr , Δz , $\Delta \phi$):

$$f(\Delta CableDelay, \Delta r, \Delta z, \Delta \phi) = \sum_j^{AntennaPairs} < \Delta T_j^2 > \quad (3.3)$$

The Hpol and Vpol phase centers are minimized separately using the same objective function.

In order to find a reasonable fit, some things needed to be excluded. Runs 129-132, when

the pulser was tilted to produce pulses polarized at $\pm 45^\circ$, were all excluded. I also excluded any single correlation with $\Delta T > 1$ ns (about 5% of Hpol events and 9% of Vpol events). Only the three phi sectors of antennas facing WAIS and their correlations with the closest antennas in those 3 phi sectors were considered.

Empirically I found that this minimization needs to be done in parts (if all variables are allowed to float, the procedure does not converge due to correlations between variables). The ordering that I found worked best was:

- (1) Minimize the objective function while allowing only the cable delays to vary.
- (2) Use the cable delays found in step (1) to minimize the objective function while varying only the Δr .
- (3) Use the cable delays and Δr to find $\Delta\phi$.
- (4) Use variables found in previous steps to find Δz .
- (5) Use $\Delta\phi$, Δz and the cable delays to fit for Δr once more.

The results of the fit are stored in a calibration file that is used for validation.

The final phase center positions found in this way are roughly ~ 36 cm inwards from the face for Hpol and ~ 21.5 cm inwards from the face for Vpol. The discrepancy is unintuitive, but consistent with both the ANITA-III findings and XFDTD simulations of the horn antennas.

3.5.3 Resolutions

The ultimate goal of finding more accurate phase center positions is enhanced pointing resolution. In order to evaluate our pointing resolution, we plot the difference between the known WAIS position and the reconstructed WAIS position in ϕ and θ . Our final phase center positions remove pointing offsets in ϕ . The removed offsets were 0.8° in ϕ for Hpol and 0.4° in ϕ for Vpol. Very small offsets were also removed in θ , less than 0.1° . Pointing resolution for the uncalibrated (photogrammetry positions) phase centers σ was 0.75° in ϕ and 0.22° in θ for Hpol and 0.65° in ϕ and 0.23° in θ for Vpol. Uncalibrated pointing

resolutions for Hpol and Vpol combined are shown in Fig. 3.15. The final calibrated pointing resolution σ was 0.31° in ϕ and 0.13° in θ for Hpol and 0.37° in ϕ and 0.14° in θ for Vpol. Calibrated pointing resolutions for Hpol and Vpol combined are shown in Fig. 3.16. Pointing resolution varies strongly as a function of SNR, and this is shown in both θ and ϕ is Fig. 3.17.

In future experiments it would be preferable to have multiple calibration pulser locations. Using only WAIS for both calibration and validation is circular. We know that phase center position varies significantly with frequency content, so having different pulser sources would allow us to have different phase centers for different frequency makeups. This would cut down on the pointing uncertainty we have for things that are dissimilar to WAIS. If multiple stations at different locations proves impossible, then having multiple pulsers or a pulser with tunable frequency content would be servicable. We also should take data with all TUFF configurations that may be used throughout the flight. In ANITA-IV, all WAIS calibration data was taken with a single TUFF configuration. Different TUFF configurations necessarily imply different frequency content of recorded signals, so the phase centers positions will vary.

3.6 Trigger Efficiency

Trigger efficiency is very important to overall instrument sensitivity. It varies with SNR and elevation angle of the signal. We attempted to quantify both of these effects using two different setups, a pre-flight scan performed in the NASA LDB facility, and using the previously mentioned WAIS calibration pulsers during flight.

Our pre-flight efficiency testing was done by injecting impulses into like-polarized antennas in neighboring phi sectors. We varied the SNR of the injected impulses to map out the response curve of the firmware. The results of these tests can be seen in Fig. 3.18.

The overall trigger efficiency of the payload was determined using the hundreds of thousands of WAIS pulses we recorded. In order to calculate efficiency, we divide the number of recorded signals by the number of expected signals given a 1 Hz pulsing rate, and bin




Figure 3.15: The pointing resolution using the photogrammetry numbers as the phase centers. There is a mean offset of 0.61° and a resolution of 0.87° in ϕ and a mean offset of 0.04° and a resolution of 0.24° in θ .




Figure 3.16: The pointing resolution using the phase centers determined from fitting using the WAIS pulses. There is a mean offset of 0.01° and a resolution of 0.42° in ϕ and a mean offset of 0.01° and a resolution of 0.14° in θ .

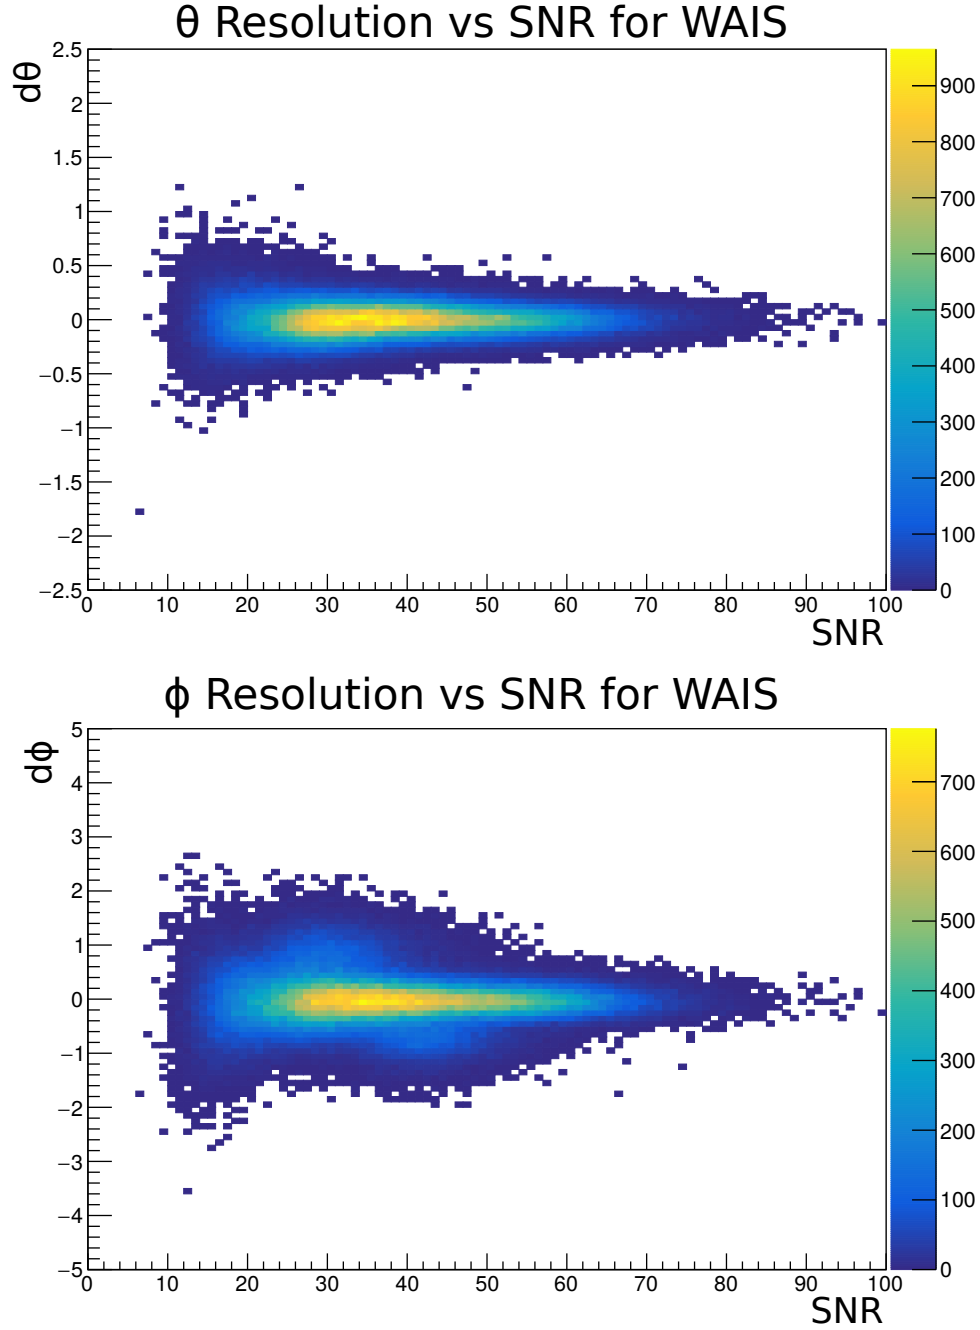


Figure 3.17: The top panel shows θ resolution as a function of SNR as calculated from WAIS pulses using the calibrated phase centers and the bottom panel is the same thing in ϕ . The color axis for both is simply the number of WAIS pulses. SNR is calculated using the coherently summed waveform, so it is very roughly $\sqrt{15}$ times larger than single antenna SNR.

that according to SNR. The trigger efficiency according to WAIS data is shown in Fig. 3.18. There is little difference in trigger efficiency in the perfect conditions of the LDB lab tests versus the reality of using WAIS pulses. Our trigger efficiency 50% point is near ~ 4.1 SNR in either case, where SNR is defined using the single most on-boresight antenna.

3.7 Relative Hpol and Vpol Delays

The relative delay time between the Hpol and Vpol channels is especially important with our trigger scheme. If there were a large relative delay between channels, any linearly polarized signal not directly along the Hpol or Vpol axes would appear to be non-linearly polarized, and therefore would have reduced trigger efficiency. Due to the offset in feed position for Hpol and Vpol, there is a 65 ps difference in timing that is induced by the antenna. During instrument integration and testing in Palestine, we very carefully corrected out the difference induced by antenna feeds using slightly different length cables for the Hpol and Vpol channels of each antenna. Using the WAIS pulser data taken from when the pulsing antenna was tilted to 45° we were able to confirm that the timing offset was well compensated for and correct the last few picosecond differences present in each channel.

3.8 Pointing Offset

One outstanding unsolved mystery was a pointing offset in ϕ that occurred for the 45° polarized WAIS pulses. Depending on whether the pulse came in at $\pm 45^\circ$, there is an offset in ϕ pointing of roughly $\pm 1^\circ$. The offset as a function of pulse polarization angle is shown in Fig. 3.19. Since these pulses are polarized at $\pm 45^\circ$, we record significant power in Hpol and Vpol, and can reconstruct pulses using either polarization. Fig. 3.20 shows that whether pointing using Hpol or Vpol, the result is the same. Fig. 3.21 shows that the offset in ϕ pointing only appears at θ angles $> 14^\circ$ to WAIS, and appears to continuously get worse




Figure 3.18: The top panel shows the trigger efficiency according to the pre-flight scans. The 50% point (indicated by gray lines) is at approximately 4.2 SNR. Vertical error bars are statistical, while horizontal error bars are negligible because SNR was precisely set. Data is from an ANITA internal note [110]. The bottom panel shows the trigger efficiency according to WAIS data. The 50% point (indicated by gray lines) is at approximately 4.1 SNR. Vertical error bars are statistical, while horizontal error bars are indicative of the binning of the data. Note the different X-axis limits. The efficiency as derived from in-flight pulses appears to be slightly better than from data taken pre-flight, but likely due to slightly different definitions of SNR. SNR as defined in the top plot is taken on an oscilloscope directly before the SURF, whereas in the WAIS data it is calculated from what is recorded by the SURF.




Figure 3.19: The offset in pointing in ϕ as a function of incoming pulse polarization angle. Pointing is done using the “interesting” peak, as defined in Chapter 4.5.1.

with increasing θ . We never figured out how to fix this offset, but found that it didn’t matter for clustering because it only affected events with large θ values, which are clustered mostly by the distance metric detailed in Chapter 4.12.




Figure 3.20: This plot shows pointing offsets in ϕ for when using Hpol and Vpol to reconstruct the 45° WAIS pulses. The y-axis is the difference between ϕ reconstructed using Hpol channels and the actual ϕ wais. The x-axis is the difference between ϕ reconstructed using Vpol channels and actual ϕ wais. Hpol and Vpol point at approximately the same place, but the 45° and -45° pulses have opposite $d\phi$ offsets.




Figure 3.21: This plot shows pointing offsets in ϕ for the 45° WAIS pulses as a function of incoming θ angle to WAIS. Pointing is done using the “interesting” peak, as defined in Chapter 4.5.1. The y-axis is the difference between $\phi_{reconstructed}$ and ϕ_{wais} and the x-axis is θ_{wais} . 45° and -45° pulses have opposite $d\phi$ offsets and this offset only appears at $\theta > 14^\circ$.

CHAPTER 4

ANALYSIS

The goal of this ANITA analysis was to achieve the highest possible sensitivity while remaining as unbiased as possible. To that effect, both analyses were performed blind, with details of the blinding described later in this chapter. Our science events are expected to be impulsive, highly linearly polarized (vertically polarized for Askaryan neutrinos and horizontally polarized for extensive air showers), and isolated from other events. The strategy of this analysis is to remove the two major sources of background, thermal and anthropogenic noise, until the highest sensitivity is achieved, with high efficiency and low background in the signal region. My analysis methods are a combination of the best of past ANITA analyses and my own ideas in an attempt to perform the best analysis possible. Results of the Askaryan neutrino search are published in [111].

We have an additional handle in the CR search that allows us to achieve lower backgrounds. This choice was made due to the *mystery events* found in ANITA-I and III analysis [77, 78]. These two events are tantalizing but unexpected, so previous analyses attempted an after the fact background estimate and likelihood analysis, whereas we would like to do the entire analysis while remaining blind.

In the following chapter I will give a complete description of both the Askaryan neutrino and EAS analyses, including the choices I made and the reasoning behind them. I begin the chapter with an overview of the software, all of which is publicly available at <https://github.com/anitaNeutrino>. From there I will describe the available data, as well as the blinding strategies employed. Then I describe the steps to turn raw data into useful data: filtering, interferometry, constructing a coherently-summed waveform, dedispersion and extracting relevant quantities from these constructions. Next I will talk about removing events we don't want, starting with non-quality events, moving on to thermal noise, and ultimately removing anthropogenic events. Finally I will talk about making a background

estimate and unblinding.

4.1 Software Overview

The ANITA analysis software has been developed over many years, by many members of the collaboration. The following section describes all of the packages that were important to my analysis. All of these tools were written in C++ using the ROOT Data Analysis Framework [112].

4.1.1 libRootFftwWrapper

The Fastest Fourier Transform in the West (FFTW) is a collection of routines for quickly computing discrete fourier transforms using C [113]. It was not designed to be used with ROOT objects, so `libRootFftwWrapper` was designed to provide a wrapper for it. Ryan Nichol wrote most of the initial software in the `libRootFftwWrapper` package, but it has grown to be more than just a wrapper for FFTW since its inception. It now includes a variety of ANITA related routines, especially related to sampling, filtering and polarimetry.

4.1.2 anitaTreeMaker

Before anything can be done with the data, it needs to be converted to a usable format. Raw ANITA data is written in a compressed C++ binary format that has very fast write speeds and requires relatively little storage space. The `anitaTreeMaker` repository is filled with scripts to convert all of the raw ANITA data into ROOT objects.

4.1.3 eventReaderRoot

The `eventReaderRoot` package was originally designed for ANITA-II and has been updated to include ANITA-III and IV. Its main purpose is to read in the ROOT-ified versions of the

raw ANITA data and turn that into calibrated data. It also includes constant, flight specific information, such as the number of channels and how they are mapped, the locations of pulsers, and information about broken channels.

4.1.4 `AnitaAnalysisFramework`

`AnitaAnalysisFramework` is a repository created to provide a standardized framework for analysis across the collaboration. It includes tools to better use digitized waveforms, including tools for easily resampling and filtering. The framework also includes data that is useful for all analyzers, such as cosmic ray templates and impulse responses for each channel. Additionally, the framework houses the `AnitaEventSummary` class. An `AnitaEventSummary` can be made for each event, which is filled with all of the important reduced quantities for that event. These summaries are created to keep track of all quantities an analyzer might want, and to reduce computation to one large ordeal, to avoid repeating costly computations. In practice it has also allowed multiple analyses to operate using a single set of summaries, generated once.

4.1.5 `UCorrelator`

The main purpose of `UCorrelator` is to fill a given `AnitaEventSummary` with the relevant reduced quantities. The `Analyzer` class does this, but getting there takes a lot of work. To process one event, the `Analyzer` must load in the calibrated data, filter the waveforms, construct an interferometric map, find the peaks of that map, construct coherently-summed and coherently-summed dedispersed waveforms for those peaks, and extract the desired reduced quantities from these objects. In order to do all of that, the `Analyzer` makes use of many other objects also in the `UCorrelator` repository. The `UCorrelator` repository also contains tools for clustering, which are not relevant for this analysis, but were used in [88].

4.1.6 `anitaAnalysisTools`

I used code found in the `anitaAnalysisTools` repository purely for the clustering step of my analysis, however, this set of tools is capable of doing everything `UCorrelator` is. The tools within were developed at the same time and for the same purpose as `UCorrelator`, for a parallel ANITA-III analysis [88].

4.1.7 `anitaMagicDisplay`

The `anitaMagicDisplay` package is used to display events. It has different settings allowing the user to look at the waveforms from each antenna, the frequency spectrum of those waveforms, or the outputs of `UCorrelator` or `anitaAnalysisTools`. Also included are quantities of interest for each event such as polarization parameters and payload orientation information. Another recent addition is the ability to change filtering strategies on the fly. I added the option to read in a “playlist” made up of event numbers so the user can restrict their viewing to only those events of interest. An example of a WAIS pulse as displayed in `anitaMagicDisplay` is shown in Fig. 4.1.

4.1.8 `anitaEventCorrelator`

The `anitaEventCorrelator` repository is no longer used for event correlation. It has become a repository mostly filled with miscellaneous functions that don’t have a home anywhere else. Much of what is housed here has to do with geography and geometry. This repository includes Antarctic base lists, objects to determine the position of HiCal, and methods for mapping and plotting events onto the Antarctic continent, among other things.

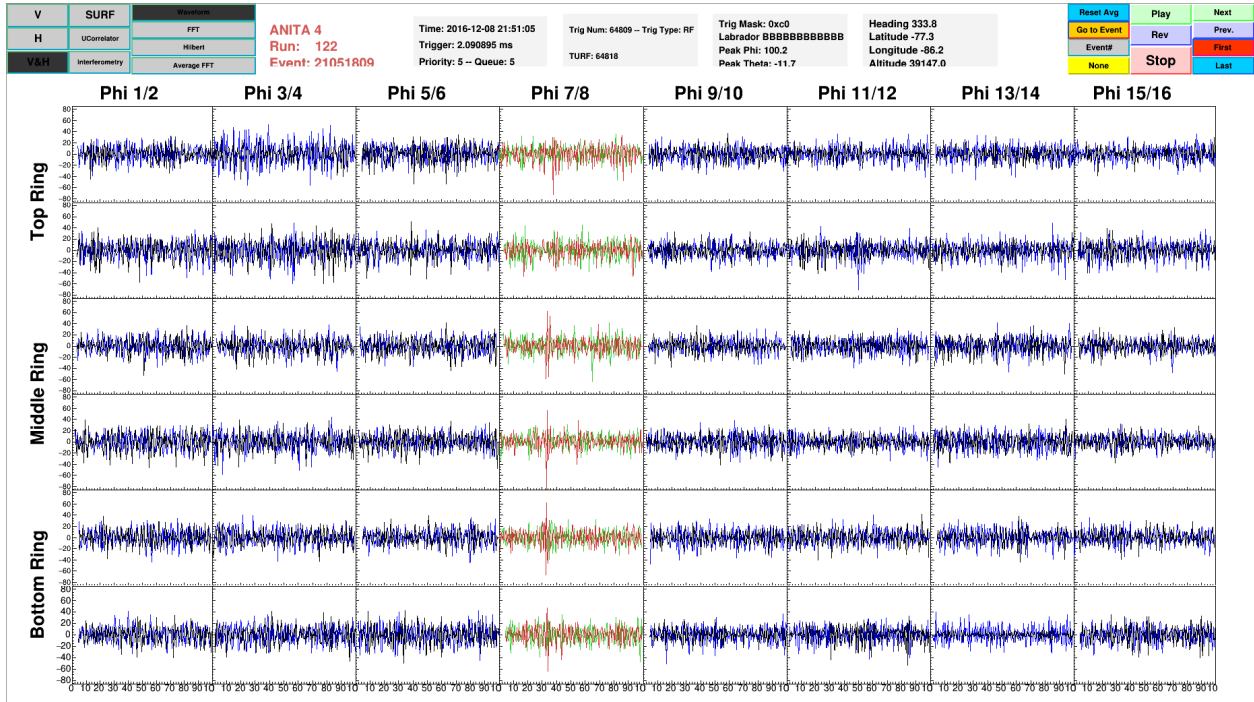


Figure 4.1: A WAIS calibration pulse as displayed in `anitaMagicDisplay`. The red and green waveforms are the vertically and horizontally polarized channels that triggered, and the black and blue waveforms are the vertically and horizontally polarized channels that did not trigger. The buttons in the top left correspond to different ways to view each event, and the buttons in the top right allow the user to navigate through events. In the center top there is useful information about each event displayed.

4.1.9 icemc

`icemc` is a Monte Carlo simulation tool used to simulate how neutrinos interact in the ice and how ANITA might view them. I used `icemc` simulated output neutrinos as my signal training set to tune my analysis cuts. It was also the tool used to calculate the ANITA-IV limit. For more information on the development, methodology, and usage of `icemc`, see [114].

4.2 Data

Recorded events from the ANITA flight can be broken down into smaller categories, each of which are useful for different purposes. Thermal noise makes up the vast majority of recorded triggered events ($\sim 99\%$). We also have two large samples from different pulsers, WAIS and HiCal, which are used mainly for calibration, as detailed in Chapter 3. There are also a large number of simulated neutrinos, produced by `icemc`, which are necessary for estimating analysis efficiency.

4.2.1 Thermal Samples

About $\sim 99\%$ of the recorded events are thermal noise, which must be cut from the analysis with very high rejection efficiency to ensure a low background. In order to effectively remove all of these thermal noise events, we need a very large sample to train our analysis on. Our thermal noise sample consists of all triggered events that reconstruct above the horizontal. We expect no impulsive events from above the horizontal because of the dispersive nature of the ionosphere. The horizontal is chosen rather than the horizon because we still expect to see a few cosmic rays from above the horizon but below the horizontal.

4.2.2 Pulsar Events

Pulsar events are an important source of impulsive events from a known location. We have three sets of pulsar events from WAIS, HiCal, and LDB. The WAIS events were used extensively in calibration, as detailed in Chapter 3. WAIS pulsar events were also used as an impulsive neutrino-like source to check analysis efficiency. HiCal was used mainly as a check on the robustness on polarity measurements and pointing. It also provided a large dataset for Antarctic surface reflectivity measurements, see [115] for more on that. I never used the LDB pulses, but there is a large dataset there that could potentially be useful.

4.2.3 Simulated Neutrinos

In order to properly train an analysis, we need signal and background training sets. The data for the signal portion of our training set is produced by the `icemc` simulation. Multiple sets of simulated neutrinos were produced over the course of the analysis. The sample that was used to make the final analysis cuts was a sample of 2 million simulated neutrinos, added to actual minimum-bias waveforms from the flight, produced following the maximum mixed-composition Kotera model energy spectrum [116]. This training sample is what is used to calculate analysis efficiency. Each simulated neutrino is produced with a weight based on a phase-space factor and earth absorption. All analysis efficiencies are calculated accounting for these weights.

4.3 Blinding

Past ANITA analyses have typically been blind, and we as a collaboration chose to continue that trend with the ANITA-IV analysis. Blind analyses are important to reduce the effects of individual's (in this case, my own) bias in determining the outcome of an experiment. Here are two ways to imagine bias strongly affecting the ANITA-IV analysis without a proper

blind analysis plan –

Suppose the analyst were an ambitious and optimistic young graduate student, with high hopes to discover the highest energy neutrino ever recorded. This analyst fully believes in both their own abilities and the experiment, and is 100% certain that ANITA saw the neutrinos it is searching for. Said analyst could tune their cuts until one (or a few) neutrino-like events remained in the signal region, then tune their background estimate to artificially inflate this evidence to discovery levels.

In the second example, suppose your analysis is being spearheaded by a burnt-out 8+ year graduate student whose only desire is to graduate and get a cushy job outside of academia. The foremost desire of this analyst is to publish a new limit as quickly and quietly as possible. They have no desire to defend a controversial result. It isn't hard to imagine this analyst tuning their cuts so that zero events end up in the signal region and calling it a day.

Bias may not always show up so blatantly in science, but it will always be present without some check on it. For a review on bias and blind analysis methods in nuclear and particle physics see [117], for an interesting article, mostly about bias in medical sciences, that was formative in my understanding of scientific bias see [118].

4.3.1 Askaryan Neutrino Blinding

The method of blinding for the Askaryan neutrino signal region is the “hidden signal box” technique, where the parameter space in which the signal resides is hidden to the analyst until all cuts are finalized and a background estimate is constructed. In the past other techniques have been used such as making cuts using only 10% of the data, or “salting” the dataset with simulated neutrinos. I ultimately decided on the hidden signal box method rather than the other two because it seemed the most fool-proof. The 10% method introduces problems with cutting thermal noise, as you only have a portion of the distribution to train on, which means you must extrapolate your cuts to more stringent levels when expanding to the whole

analysis. Salting is a perfectly reasonable blinding method, but I determined the hidden signal box to be sufficient.

4.3.2 Extensive Air Shower Blinding

Blinding in the EAS signal region was done differently. We have already seen cosmic rays and assumed that tens of them would be present in the signal region for ANITA-IV, so the full hidden signal box method was not necessary. Instead, we decided to blind the *polarity* of the events in the EAS signal region. Polarity is the only discriminator we have between direct cosmic ray events, reflected cosmic ray events, and a potential τ neutrino channel. Polarity blinding is in effect a hidden signal box around the EAS τ neutrino signal region. The way that polarity blinding is implemented is by randomly applying an overall multiplicative factor of 1 or -1 to the voltage values of the entire signal.

4.4 Digital Filtering

Despite being greatly improved from previous flights due to the addition of the TUFFs [91], our analog filtering is far from perfect. If CW interference is present and strong enough, it will affect pointing by “pulling” the peak of the interferometric map towards the CW source. In order to remove CW interference not already removed by the TUFFs, we need to have a secondary filtering step done in software. Because the CW interference background changes dramatically at different points on the continent, we need something adaptive, that can deal with different noise environments as they arise.

4.4.1 Sine Subtraction Filtering

The sine subtraction filtering algorithm was designed to be adaptive and minimize dispersion of the incident signal [119]. It works by first looking at the frequency spectrum of the event

and searching for any peaks in the magnitude. The algorithm attempts to fit a sine wave with a frequency bounded near the frequency of any peaks it finds to the signal, with variable phase and amplitude. If there is a sine wave that fits well, the software subtracts that sine wave off and compares the total power of the signal before and after subtraction, and if the total power subtracted off is above a configurable threshold, that sine wave is subtracted off and the algorithm continues iterating. If the subtracted power is below the given threshold, iteration continues until the (also configurable) failure limit is met. This acts to remove single frequency CW without adding dispersion or acausality (because filtering is done in the time domain).

In order to speed up computational efficiency, spectrograms for the entire flight are created ahead of time, which creates guidelines for which frequencies are noisy at which times during the flight. These guidelines are made by looking at a variable known as “peakiness,” which is the ratio of power in a frequency band for the spectrogram of a specific time in the flight to the power in that frequency band from the spectrogram taken with terminated amplifiers (equivalent to thermal noise). A configurable exponent, α , modifies the amount of power that needs to be subtracted to keep a solution, P , by the peakiness value, \mathcal{P} , to that exponent, like so:

$$P_{new} = \mathcal{P}^\alpha P \tag{4.1}$$

The filtering strategy that I settled on for this analysis was a sine subtraction algorithm that required $> 10\%$ power removed to keep a solution, required 3 failed iterations in a row to quit, and squared the peakiness value. In previous analyses this filtering strategy was determined by finding the filtering strategy with the best pointing resolution for calibration pulses, but because we had very little CW interference for our WAIS pulses thanks to the TUFFs, the pointing resolution didn’t change strongly as a function of filtering strategy.

My motivation for choosing this filtering strategy was that it worked well for ANITA-III and its adaptive nature allowed it to work on the ANITA-IV calibration pulses without




Figure 4.2: When all three notches are on, the sine subtraction algorithm can fail by falsely identifying an area between notches as a peak and subtracting it. This is shown with event 36785931, a cosmic ray. The blue waveform is unfiltered and the red waveform is after sine subtraction filtering. The signal has lost a significant amount of low-frequency power.

causing any pointing problems. Ex post facto analysis has suggested that this type of filtering is in some way “fooled” when all three notches are on. The two small regions between notches can appear to be narrowband CW interference. This failure mode is shown with a cosmic ray in Fig. 4.2. In a future experiment, I would make sure to take calibration data with all potential notch settings, to allow for better testing of the interaction between software filtering and hardware notch filtering.

4.4.2 Channel BH13 Transfer Function

Although not exactly filtering, the channel with a very different impulse response from the rest, BH13, must be corrected for before continuing on. To correct the impulse response of channel BH13, I first made an average impulse response of all horizontally polarized channels. Then I divided that average impulse response by the impulse response of BH13 to create

a transfer function. Before doing any further analysis, channel BH13's recorded waveform is multiplied by this transfer function to make sure steps of the analysis reliant on having similar waveforms, such as interferometry and coherent summing, aren't spuriously harmed by BH13's dissimilar impulse response.

4.5 Event Reconstruction

After filtering out CW interference we can begin the process of pointing events back to their sources. It is important for ANITA to be able to reconstruct event directions with good resolution, because the isolation of events is the most powerful discriminator between man-made and physics events. The first step of this process is to interpolate the waveforms so that they are both evenly sampled and upsampled to a finer time resolution.

As I mentioned before, our sampling rate is nominally 2.6 GS/s, with some variability in the timing separation between points. Waveforms are interpolated to be evenly sampled by iteratively fitting a cubic polynomial to a range of sampled points, and using that polynomial to determine the values at the proper evenly spaced time values, in a process called Akima spline interpolation [120]. Akima interpolation is used to ensure that every sample-to-sample $\delta t = \frac{1}{2.6}$ ns.

Evenly sampled waveforms can now be used to construct the interferometric map, but are further upsampled before constructing coherently-summed waveforms. Upsampling is accomplished by zero-padding in the Fourier domain, meaning zeros are appended to the end of the complex Fourier transform of the signal. Zero-padding in this way is desirable because it adds no dispersion (doesn't change to frequency content of the signal). I chose to zero-pad the waveforms by a factor of four before creating coherently-summed waveforms, effectively increasing the sampling rate to 10.4 GS/s.

4.5.1 Interferometry

The basic idea of interferometry is that you can combine multiple measurements of the same signal into a higher SNR or higher resolution measurement. In this analysis, I used interferometric methods that were developed for the ANITA-I analysis, detailed here [121]. Because we have similar antenna responses throughout our array of antennas, we can pairwise cross-correlate the recorded signals, where the recorded signals are delayed by a time corresponding to a given direction.

Assuming we are in the far-field, where the distance of the source D and the wavelength of the source λ obey the relation $D \gg \lambda$, we can assume an incident plane wave, and the geometric time delay $\Delta\tau_{ij}$ between antennas is given by:

$$c\Delta\tau_{ij} = \hat{\mathbf{r}} \cdot \mathbf{R}_i - \mathbf{R}_j \quad (4.2)$$

where c is the speed of light, $\hat{\mathbf{r}}$ is the incident plane wave direction, and \mathbf{R}_i and \mathbf{R}_j are the position vectors of two antennas. With $\Delta\tau_{ij}$ in hand we can calculate the cross-correlation values between antenna voltages for a given incoming direction, defined as:

$$v_i \star v_j(\hat{\mathbf{r}}) = \int_0^T dt v_i(t) v_j(t - \Delta\tau_{ij}) \quad (4.3)$$

where \star is the cross-correlation operator, and the bounds of the integral are the lengths of the waveforms. In practice, we use the dimensionless correlation coefficient, C_{ij} , defined as:

$$C_{ij} = \frac{v_i \star v_j}{\sigma_{v_i} \sigma_{v_j}} \quad (4.4)$$

which is bounded between -1 and 1. The map is filled with the summed values of C_{ij} for




Figure 4.3: This is the interferometric map for a WAIS calibration pulse event. The z-axis (color axis) is the average correlation value for the waveforms going into the map, and the x and y-axes are ϕ and θ , respectively. For this map the peak is very clear, because it is a high SNR pulsar event. The three black crosses on the map are the peaks found by the peak finding algorithm, where the biggest cross is placed on the biggest peak. There is a red star right at the center of the largest black cross indicating the true direction of the WAIS pulsar.

each incoming direction:

$$M(\hat{\mathbf{r}}) = \sum_{i=1}^{N_a} \sum_{j<1} \frac{C_{ij}(\hat{\mathbf{r}})}{N_b} \quad (4.5)$$

where M is the value of the map as a function of incoming direction, N_a is the number of antennas used, and N_b is the number of baselines, which is related to N_a by $N_b = \binom{N_a}{2}$. This analysis used 15 antennas per map direction, corresponding to 105 baselines. I used 15 antennas because the directionality of the ANITA horn antennas ensures that no more than five adjacent phi sectors can see any one event. An example interferometric map for a WAIS pulse event is shown in Fig. 4.3.

After the initial interferometric map is constructed, it needs to be searched for peaks. First, a configurable number (three for each polarization for my analysis) of rough peaks are found by looping through the map and finding any maxima, while also making sure maxima are separated by a configurable distance. Once these rough peaks are found, we search the map for a finer, zoomed in peak. We do this by creating a finely binned image

centered around the peak, and fitting a bivariate parabola to it. The direction of the peak according to these fits are our signal source hypotheses. I wanted to reduce the six peaks (3 horizontal, 3 vertical) of each interferometric map to the single most “interesting” signal source hypothesis for ease of analysis. To do this I selected the source hypothesis with the highest value of interferometric map peak multiplied by dedispersed impulsivity measure (see Chapter 4.9 for impulsivity measure description). I chose this quantity because I wanted the most impulsive and coherent peak of the map, and map peak and dedispersed impulsivity could both get “fooled” in different ways. Map peak could fail to find the most interesting peak because it could choose a strong source of coherent CW over a weak impulsive source. De-dispersed impulsivity sometimes fails by picking the sidelobes of very bright impulsive signals over the main lobe. My analysis only considered one source hypothesis from each event, chosen in this way.

4.5.2 Tracing to the Source

Now that we have the likely direction of the source, we can trace back along that vector to the source position. If the source is above the horizon that is enough information, but if the source is on the continent, we must determine exactly where. In order to do this well we must have accurate topographical information about the Antarctic continent. The topographical maps we use were made by the Radarsat Antarctic Mapping Project Digital Elevation Model Version 2 (RAMPDEM2) [122]. Topographical information provided by RAMPDEM2 is important because Antarctica is far from flat. Ray-tracing for my analysis was done without correcting for refraction, which I determined to be unimportant, unless the event appeared isolated.

4.6 Coherently Summed Waveform

After the interferometric map is constructed and searched for peaks, we can use those peak directions to create our coherently-summed waveforms. These are produced using waveforms upsampled to ~ 10 ps between points. We delay each waveform by the offset required for the source direction hypothesis, and sum them all together. The coherently-summed waveform is also produced using 15 antennas, from the 5 phi sectors in view of the signal, for the same reasoning as the interferometric map. Theoretically the SNR of these waveforms should be improved by a factor of \sqrt{N} where N is the number of antennas, because signal will add coherently and noise will add incoherently.

4.7 De-dispersion

Our signal is shaped by the antennas, cables and filters it passes through on the way to being recorded. In order to reconstruct the original signal shape we must remove the effects of the impulse response as best we can. Ideally we would like to deconvolve the entire system response, removing both phase and amplitude distortions caused by the system, but because there are frequency bands where we cannot record signal, such as the notches and outside of our band, we cannot do a full deconvolution. Instead we remove only the dispersion, or phase response of the impulse response. The method we use to do this is called all-pass deconvolution, or dedispersion. In dedispersion we only divide out the phase distortion of the impulse response, leaving the amplitude untouched, which works well enough because our amplitude response is fairly flat across the 180-600 MHz band, apart from the notches. The resulting signal has the same frequency content as the recorded signal, but with minimal group delay between frequencies, which greatly increases the power near the peak of impulsive signals.

4.7.1 Picking the Correct Impulse Response

With the addition of the TUFFs, the impulse response of ANITA went from a single quantity throughout the entire flight in ANITA-III to a function of time in ANITA-IV. There were six different configurations used over the course of the flight. I created an index containing times each configuration was active. Every time a dedispersion is done, the software checks the index of times and loads the correct impulse response.

4.8 Polarimetry

Polarization angle and content is an important discriminator for our science events. We expect a high linear polarization fraction for EAS and Askaryan neutrino events, and a well-defined polarization angle for EAS events. We transform the orthogonal dual polarization waveforms recorded by our antennas into Stokes parameters, which are components of a four-vector describing the polarization of any given signal. Stokes parameters are I , which is total intensity, Q and U , which are linear polarizations in bases rotated 45° from one another, and V , which is circular polarization.

The Stokes parameters of a signal are calculated from the complex voltage, \tilde{v} of a signal: The equations for calculating the instantaneous Stokes parameters are below:

$$\tilde{v}_i = v_j + i\mathcal{H}(v_j) \quad (4.6)$$

$$I_j = |\tilde{v}_{j,H}|^2 + |\tilde{v}_{j,V}|^2 \quad (4.7)$$

$$Q_j = |\tilde{v}_{j,H}|^2 - |\tilde{v}_{j,V}|^2 \quad (4.8)$$

$$U_j = \text{Re}(\tilde{v}_{j,H} \cdot \tilde{v}_{j,V}^*) \quad (4.9)$$

$$V_j = \text{Im}(\tilde{v}_{j,H} \cdot \tilde{v}_{j,V}^*) \quad (4.10)$$

where \mathcal{H} is the Hilbert transform operator, the subscript j is the current sample number, the subscripts H and V refer to the horizontally or vertically polarized channel of each antenna, and the superscript $*$ is the complex conjugation operator. There are two ways we calculate the Stokes parameters, the average and the windowed average.

In order to get the average Stokes parameters, we sum over the n samples in the waveforms and divide the sum by n . The windowed average Stokes parameters are calculated by finding the maximum instantaneous Stokes I, summing outwards from that point until the instantaneous Stokes I drops below a configurable threshold (I chose 80% of maximum instantaneous Stokes I), and taking the average of the sum in that window. Windowing the Stokes parameters is useful to distinguish signal events from noise because ANITA has a trigger window much longer than the expected length of events, so portions will be filled with noise, biasing our average Stokes parameters low for non-thermal events. However, windowed Stokes parameters will be biased high for noise events, because there is often a peak in Stokes I, so using a combination of these two measures is most effective.

Using either version of the Stokes parameters we can do some further calculation to find more useful quantities. The fraction of signal that is linearly polarized, L , can be calculated by:

$$L = \frac{\sqrt{(Q^2 + U^2)}}{I} \quad (4.11)$$

The fraction of signal that is polarized, P , in any manner can be calculated by:

$$P = \frac{\sqrt{(Q^2 + U^2 + V^2)}}{I} \quad (4.12)$$

Both of these quantities are effective for separating out thermal noise, as it should be largely unpolarized. Stokes parameters can also be used to calculate the angle of linear polarization, ϕ , where:

$$\tan(2\phi) = \frac{U}{Q} \quad (4.13)$$

Polarization angle is useful for increasing purity in the EAS search, as well as reconstructing direction in the Askaryan neutrino search.

4.9 Impulsivity

We are looking for impulsive signals, so we must construct a way to define how impulsive a signal is. For this we introduce two definitions of impulsiveness.

First is I , the “impulsivity” of a signal, $I = 2A - 1$, where A is the average of the CDF constructed from integrating the fractional power of a waveform outwards from the peak of its Hilbert envelope. This quantity is constrained to be < 1 , where 1 would correspond to a perfect δ function.

The second measure of impulsiveness is the power window gradient, G , which is the average difference between the smallest time windows containing 10%, 20%, 30%, 40%, and 50% of the power. Both of these quantities are recorded for both the coherently-summed and dedispersed waveforms. They were both developed for, and first used in, the ANITA-III analysis [88].

4.10 Quality Cuts

Before we can begin separating thermal events, we must take out non-quality events. These include things like digitizer glitches, events in which the amplifiers are saturated, and *payload blasts*, which are high power, low frequency dominated signals that seem to emanate from somewhere on the payload. The majority of the following quality cuts are meant to deal with payload blasts. They are difficult to remove without specialized cuts, as they can appear very impulsive and even linearly polarized, and they will not necessarily cluster together. I cannot overstate how problematic payload blasts are for our analysis.

The exact values for all of my cuts are available in `AnitaAnalysisFramework`, but I will

go over them in a cursory manner here. There are three basic categories of quality cut: is an event real, is an event a glitch, and is an event a payload blast. A real event has an interferometric map value > 0 , is an RF trigger (not a minimum-bias trigger), and points to a θ between 60 and -50 (where negative is above the payload).

There are two types of glitch events, single sample digitizer glitches, and step function glitches. Single sample digitizer glitches are single sample spikes in voltage on the order of hundreds of mV. I search for them by looking for a voltage asymmetry in waveforms larger than 500 mV. In this step I also cut out saturated amplifier events, by searching for waveforms with voltages greater than 1000 mV. Single sample glitches are restricted to a few LAB and SURF combinations, which can be found in `UCorrelator::Analyzer::fillFlags`, so I only search those combinations for glitches. An example of a single sample digitizer glitch is in Fig. 4.4. The second glitch type is the step function glitch, in which portions of the waveform are shifted upwards or downwards by some total offset. Only a single channel on the first LAB chip exhibited this issue (9MV). An example of a step function glitch is in Fig. 4.5.

Payload blasts come in many forms, so the cuts required to remove them all are necessarily more complicated. An example of a fairly typical looking blast as viewed in `anitaMagicDisplay` is shown in Fig. 4.6, and the same event, coherently-summed with the frequency spectrum on display is shown in Fig. 4.7. I attacked the payload blasts in three general ways. The first is similar to methods used in previous ANITA analyses [84, 86, 88]. Most of the power in a payload blast shows up in the middle or bottom rings, so in previous analyses, events were marked as blasts if the maximum ratio of peak to peak voltages between bottom and top or middle and top antennas in a single phi sector in one polarization was greater than some value (2.8 in my case). Because blasts usually show up in both polarizations, I extended this ratio cut to mark events as blasts if the maximum ratio for the horizontally polarized channels added to the maximum ratio for the vertically polarized channels is greater than 4.3.




Figure 4.4: A single sample digitizer glitch event example.

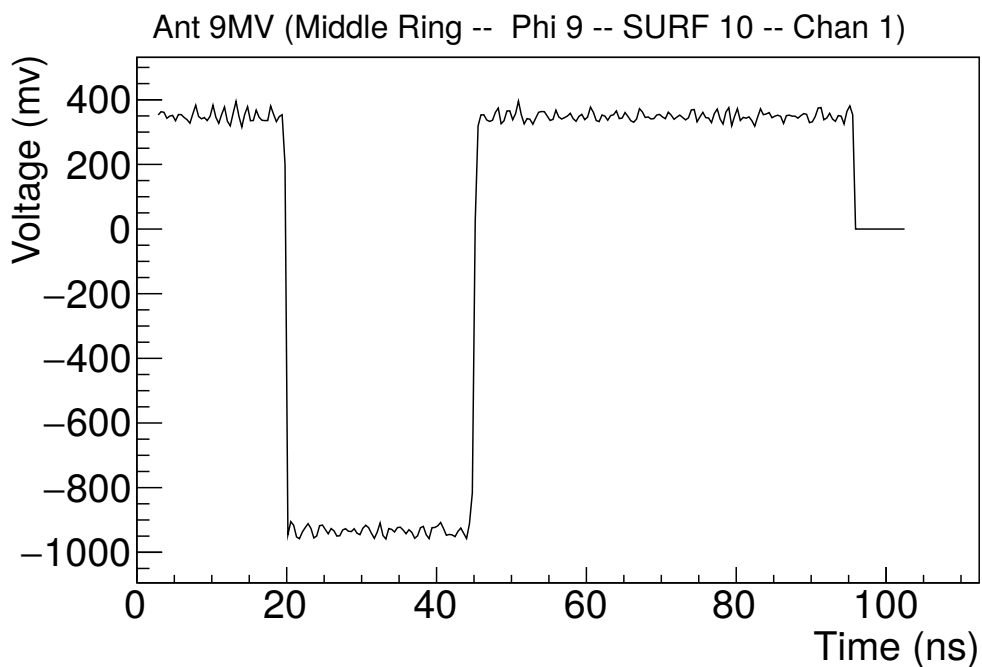


Figure 4.5: A step function event example. Step function events have badly distorted pointing and coherently-summed waveforms.

There is another rare type of blast that shows up very strongly the top rings, but not in the middle or bottom rings. These events can be removed using the same ratio previously mentioned and cutting all events below a certain value, chosen at 0.98 for my analysis. One might expect that this cutoff should be much farther from 1, but because it is the maximum ratio of voltages in the bottom or middle ring to the top ring there are 32 combinations to choose from for each polarization. If we assume that signals in the bottom, middle and top should on average be very close to the same size in terms of peak to peak voltage so that each ratio has a 50% chance of being above or below 1, then the chance for this maximum ratio value to be below 1 is only 0.5^{32} in each polarization. This is reduced even farther because we expect all real signals to come from below the payload and therefore be stronger in the bottom and middle rings. Using these voltage ratio cuts as inspiration I also developed cuts based on the maximum ratio of power in the middle or bottom rings and top rings, broken into a low power band (< 260 MHz) and total power.

The second set of cuts for dealing with payload blasts was based on their coherence. Payload blasts seem to emanate from somewhere in the near field, so they should not add coherently for a plane-wave hypothesis. I have a set of cuts based on the difference between the average value of the peaks of the Hilbert envelopes of all the waveforms going into a coherently-summed waveform and the peak value of the Hilbert envelope of the coherently-summed waveform. For a signal that adds coherently, these should be of similar magnitudes, but for a signal that adds incoherently, the average of the peak Hilbert envelope values should be much greater.

I developed my third set of cuts by plotting two variables against one another using my training data sets and seeing if there were phase space regions where payload blasts resided that simulated neutrinos and WAIS pulses didn't. There were many failed iterations before I found sets of variables that worked. The only two cuts of this type that were used in my final analysis were both measures of impulsivity vs fraction of power below 260 MHz in the

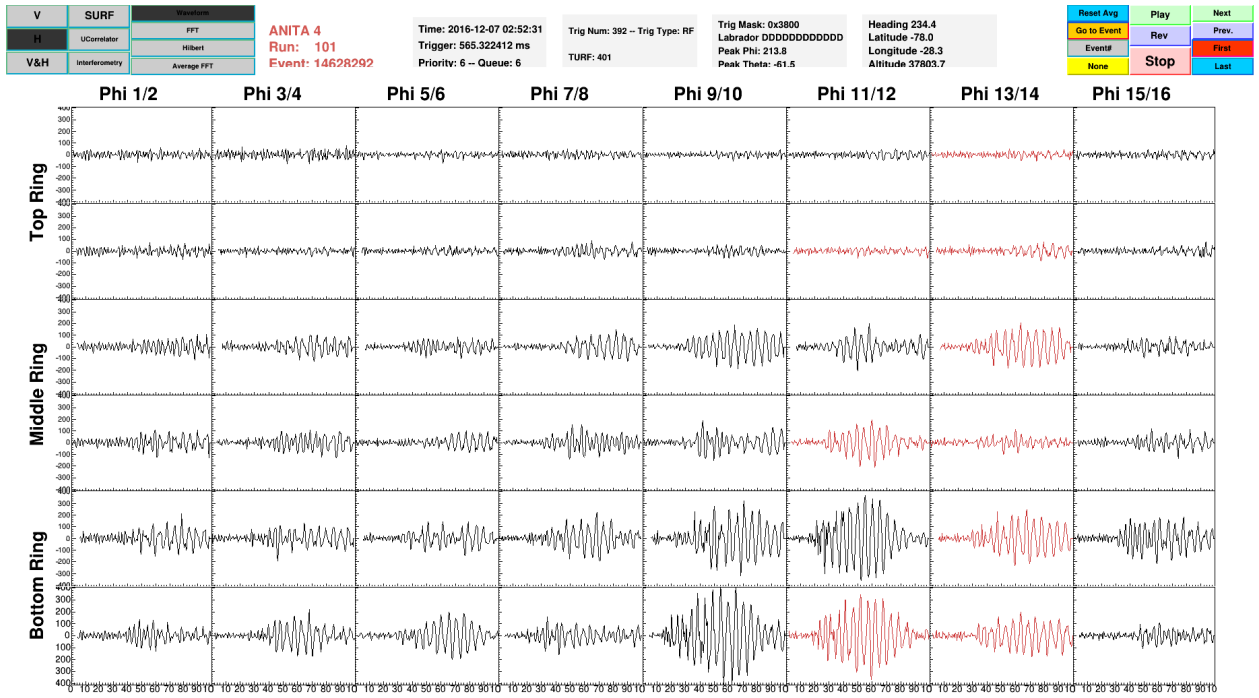


Figure 4.6: A payload blast event as viewed in `anitaMagicDisplay`. Notice the large, low frequency dominated waveforms present in the bottom and middle rings. This one was identified by a cut on the ratio of maximum voltages in the bottom to top ring.




Figure 4.7: The coherently-summed waveform and frequency spectrum of that waveform for the payload blast event shown in Fig. 4.6. Notice the huge amount of power between the low end of the band and the first notch, especially in the vertical polarization.

Quality Cuts

Cut	N_{events}^a	N_{events}^b	N_{WAIS}^a	N_{WAIS}^b	Frac. ν^a	Frac. ν^b
None	96307012	96307012	215100	215100	1.	1.
Real	89822992	89822992	215100	215100	0.983	0.983
Glitch	89767340	96248394	214992	214992	0.983	1.
Step Fn.	89755720	96277071	214425	21440	0.983	1.
V Ratios	88524896	94983427	214413	215085	0.983	1.
P Ratios	87938147	94484490	212995	213429	0.983	1.
Coherence	72124609	77891470	208650	209765	0.982	0.995
Imp. vs P	71975141	95018063	208650	215100	0.982	1.

^a Events remaining if cuts are performed in sequence

^b Events if only this cut is performed

Table 4.1: Table of quality cuts and their effect on the complete sample, the WAIS pulses, and the simulated neutrino sample. V Ratios stands for the suite of voltage ratios cuts, P Ratios stands for the suite of power ratio cuts, and Imp. vs P stands for the impulsivity vs. power cuts, all of which were implemented to remove payload blasts. For more details about each cut, see text. Total quality cut efficiency on WAIS pulses is $\sim 97\%$ and total efficiency on simulated neutrinos is $\sim 98.2\%$.

event. The effect of each of these cuts is shown in Table 4.1.

4.11 Thermal Cuts

Almost all of ANITA-IV’s dataset is made up of thermal noise events, so we need to create an analysis with extreme thermal noise rejection. In order to have the necessary rejection, we need a huge, pure sample of thermal noise. The set of triggered events that pass quality cuts and reconstruct above the payload provides the necessary sample. Next we must build a classifier to separate likely thermal noise from likely signal events. This is accomplished using a Fisher discriminant [123] in ROOT’s TMVA package [124], with simulated neutrinos as the training sample for signal and events that reconstruct above the payload as the training sample for thermal noise.

A Fisher discriminant is a multivariate linear discriminant that finds the maximum sep-




Figure 4.8: The black line that maximizes separation between signal and background in the space of power window gradient and difference between Hilbert envelope peak values of the coherently-summed and coherently-summed dedispersed waveforms, as calculated by ROOT’s TMVA package. Pink points are simulated neutrinos, while the blue through yellow color axis is made up of events that reconstruct above the payload. Moving this line upwards would increase sample purity, while moving it downwards would increase efficiency.

aration between two classes, where separation is defined as the ratio of variance between classes to variance within classes. The Fisher discriminant can be represented as a hyperplane, which can be moved along its normal vector to increase purity or efficiency. A simple two-dimensional example, using the difference between Hilbert peak values in the coherently-summed and coherently-summed dedispersed waveforms, and the power window gradient, is shown in Fig. 4.8. Moving the black line drawn on this plot would increase purity if moved upwards, and increase efficiency if moved downwards.

Variables that go into the Fisher discriminant for my analysis are three measures of impulsiveness, four measures of linear polarization fraction, the difference between the impulsivity of the dedispersed waveform and the coherently-summed waveform, and the difference between the Hilbert envelope peak value of the dedispersed waveform and the coherently-summed waveform, weighted by total power in each waveform. The most impactful variable

Variable	Discr. Power	Coefficient	Minimum	Maximum	Mean
De-dispersed Imp.	0.943	51.560	0.162	0.966	0.447
Coherent LPF	0.887	17.237	$4 \cdot 10^{-5}$	0.998	0.432
De-dispersed LPF	0.878	-3.855	$4 \cdot 10^{-5}$	0.998	0.191
Coherent Imp.	0.827	8.443	0.010	0.874	0.433
Hilbert Peak Difference	0.756	8.285	-0.141	0.165	-0.003
Imp. Difference	0.711	4.219	-0.278	0.468	0.010
De-dispersed PWG	0.606	-0.012	1.923	229.71	77.883
De-dispersed Wind. LPF	0.395	1.024	$5 \cdot 10^{-5}$	1.	0.608
Coherent Wind. LPF	0.368	0.302	$7 \cdot 10^{-5}$	1.	0.619

Table 4.2: Table of variables that go into the Fisher discriminant, and values that come out, ranked by discriminatory power. LPF stands for linear polarization fraction, Wind. is short for windowed, Imp. is short for impulsivity, and PWG is power window gradient. The Fisher score for each event is obtained by multiplying each variable’s value by the given coefficient and adding an overall offset (-49.178 for my analysis). Discriminatory power represents the amount of separation you can achieve between signal and background distributions of your training sample by using only that variable. Minimum, maximum, and mean values of each variable for all events in the training sample are also included, to give a sense of how the coefficient might actually act on each variable.

for separating thermal noise from signal is the dedispersed impulsivity measure, in fact, Peng Cao exclusively used the dedispersed impulsivity measure to separate signal and thermal noise in his ANITA-IV analysis [103]. Variables fed into the Fisher discriminant algorithm and the values associated with them are shown Table 4.2.

Once the Fisher scores are calculated for the signal and background training samples, we can use these to set our cut. Because we have more events in the remaining (below horizontal) sample than we do in the background training (above horizontal) sample, we have to extrapolate to achieve an expected background smaller than one event. I extrapolated by fitting an exponential to the tail of the Fisher score distribution for thermal events, illustrated in Fig. 4.9. This allowed me to relate a Fisher score cut value to an expected number of thermal background events that would pass my thermal cuts. I settled on an expected thermal background of 0.05 events leaking into my clustering sample, which is already small, but is further reduced by clustering cuts, as well as Askaryan neutrino or EAS




Figure 4.9: Cumulative Fisher score distributions for testing sample of background (top, red) and signal (bottom, blue) events. Vertical black lines are drawn at expected thermal background leakage of 0.5, 0.05 and 0.01 events.

specific cuts, to the point of being negligible in the end. The resulting distribution of Fisher scores for various event types is shown in Fig. 4.10.

4.12 Clustering Cuts

Once we have removed non-quality and thermal events from our sample, we should be left only with anthropogenic and science events. The next step of this analysis is to attempt to remove all anthropogenic events through a process known as clustering. Clustering involves pointing events back to their sources on the continent, grouping together all events that




Figure 4.10: Fisher scores for various types of events. Thermal events in red, WAIS pulses in cyan, and simulated neutrinos in blue. The final Fisher score cut was set at -1 (dashed line) to allow for an expected background leakage of 0.05 events.

have the same likely origin, and removing all but the most isolated events from the final sample. The motivation for clustering is that anthropogenic events could potentially look like anything, so we don't have strong signal shape discriminators like we do for thermal events, but any place on the continent with human activity should produce many detectable signals. Physics events can come from anywhere, including locations of human activity, so it is important to tightly and efficiently localize all of our events. We estimate our efficiency losses due to clustering later. There are two metrics that determine whether an event has the same likely origin, a simple distance cut, where any two events with origins within a certain distance are considered clustered (set at 40 km between estimated source locations for this analysis), and the angular log-likelihood score, developed by Andres Romero-Wolf for the ANITA-I analysis [125] and refined by Ben Strutt [88].

4.12.1 Angular Log-likelihood

The angular likelihood metric \mathcal{L} was developed to try to properly account for the pointing resolution of ANITA as a function of SNR, and the constantly changing payload location. Suppose we measure two events, A and B , that we have projected onto the continent. We want to know the angular separation of these two events, which we can find by projecting them as if they had been seen from the same observation point. We want this quantity to be symmetric and not depend on the order in which we cluster the events, so we project A onto B and B onto A , and add them. A diagram showing the asymmetry of overlap depending on viewing location is shown in Fig. 4.11.

The angular separation of A and B as viewed from the payload's location when A was recorded are denoted $\Delta\phi_{AB}$ and $\Delta\theta_{AB}$, where ϕ is azimuth angle and θ is elevation angle. Viewing the events from the payload's location when B was recorded, the angular separation of A and B are denoted $\Delta\phi_{BA}$ and $\Delta\theta_{BA}$. There is also a pointing uncertainty for each angle and event that is a function of SNR, which we have derived using WAIS pulsers, denoted $\sigma_{\phi,A}$ and $\sigma_{\theta,A}$ for events viewed from payload location A . Using these variables we can write the log-likelihood overlap between two reconstructed event locations as:

$$-2\log(\mathcal{L}_{AB}) = \frac{\Delta\phi_{AB}^2}{\sigma_{\phi,A}} + \frac{\Delta\phi_{BA}^2}{\sigma_{\phi,B}} + \frac{\Delta\theta_{AB}^2}{\sigma_{\theta,A}} + \frac{\Delta\theta_{BA}^2}{\sigma_{\theta,B}} \quad (4.14)$$

An event is considered to be clustered with another if $-2\log(\mathcal{L}_{AB})$ is below a configurable threshold. My analysis used a clustering technique developed by Ben Strutt for the ANITA-III analysis [88], that calculates log-likelihood from the mutual best-fit point of event A and B . Compared to previous analyses, this technique has a very high efficiency on simulated events.

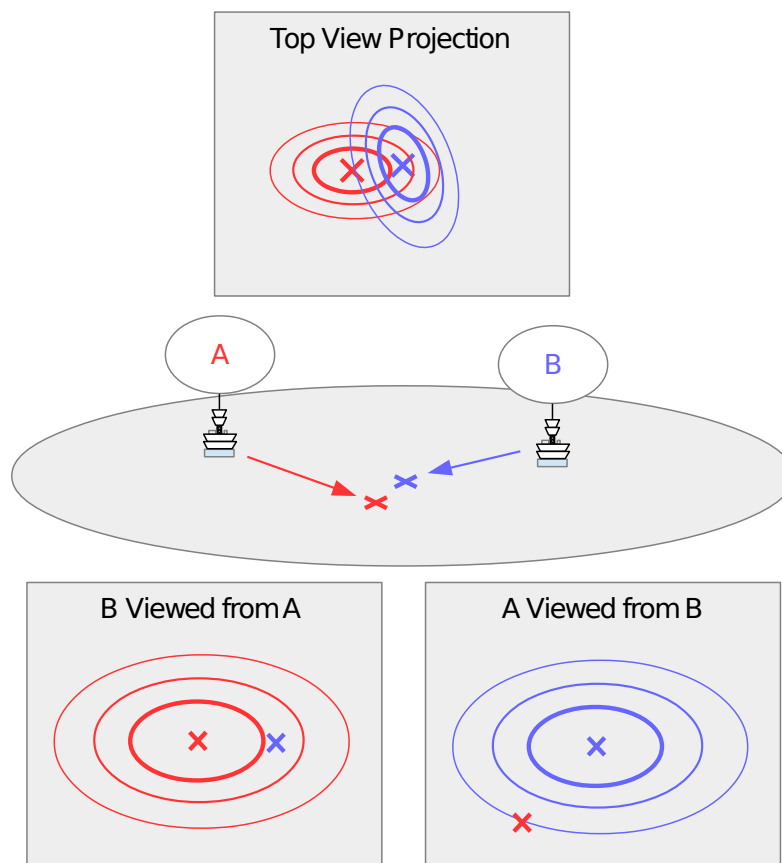


Figure 4.11: Clustering's asymmetrical nature is shown. The top panel shows event A and event B 's error ellipses when viewed from above, with the X denoting where the event pointed. In the middle panel we see how where events A and B were viewed from. The bottom panel clearly shows that while B is less than 2σ from A when viewed from A 's frame, event A is over 3σ from B when viewed from B 's frame. Diagram taken from Ben Rotter's thesis [80].




Figure 4.12: Simulated Askaryan neutrino efficiency is on the Y-axis and $-2\log(\mathcal{L})$ is on the X-axis. Efficiency is flat at around 84% between $-2\log(\mathcal{L})$ values of 4 and 10.

4.12.2 Setting the Clustering Cut

I set my clustering cuts for the Askaryan neutrino search region in order to optimize sensitivity on the Kotera SFR1 neutrino flux model [116]. Clustering cuts for the CR search were chosen in order to strike a balance between efficiency and background estimate, because we do not have a simulation with which to optimize CR search sensitivity. I chose the distance portion of the clustering cut based on WAIS pulses. The distance metric is most useful for clustering together events reconstructing from a large θ because they have a small log-likelihood footprint. A distance value of 40 km to consider events clustered effectively grouped all WAIS pulses together. Simulated Askaryan neutrino efficiency did not vary strongly as a function of $-2\log(\mathcal{L})$ between values of 4 and 10, as shown in Fig. 4.12. The background estimate was minimum at slightly different angular log-likelihood values for the Askaryan neutrino ($-2\log(\mathcal{L}) = 8$) and EAS ($-2\log(\mathcal{L}) = 6$) signal region. Estimating the background is detailed in Chapter 5.

4.13 Cosmic Ray Specific Analysis Techniques

ANITA has detected and characterized cosmic ray showers at these energies before, so there are a few additional cuts we can apply before unblinding our signal region. Since we know the expected signal shape of an EAS from previous flights, we can apply a template cut to our EAS candidates before unblinding. We also know that the polarization angle of an EAS is strongly determined by the Earth's magnetic field, allowing us to put a constraint on polarization angle. Additionally, we expect extremely strong linear polarization, so we can put a more stringent cut on linear polarization fraction.

4.13.1 Template Analysis

A previous analysis by Ben Rotter [80], used simulated cosmic ray air showers from the ZHAires simulation [126]. I took these same templates and convolved them with the ANITA-IV impulse response to produce a set of templates for my analysis. There are nine templates total, produced at 0.04° steps from the peak coherence angle, or most on-cone angle. I used the middle template to select the 10 isolated events that correlated most strongly to that template, and then combined those events to create a new template. This new template, created entirely of ANITA-IV EAS events that correlated strongly with simulated templates, is shown in Fig. 4.13.

The reasoning for a new template was that, at the time I expected to find $\mathcal{O}(100)$ EAS events, and I thought skimming the 10 best and creating a template from real data would correlate better with recorded EAS events than a simulated template would. With the new template in hand, I correlated both the coherently-summed and dedispersed template waveform against all waveforms for events that passed the thermal cuts and went into the clustering sample. Eventually I used these variables to make cuts and a background estimate, detailed in Chapter 5.




Figure 4.13: The EAS template created using ten ANITA-IV events. This is the dedispersed version, zoomed in about the peak.

4.13.2 Geomagnetically Induced Polarization Angle

An EAS will create an equal number of positively and negatively charged particles. These particles will travel through the atmosphere while being separated into positive and negatively charged groups by the Lorentz force, $\vec{F} = q(\vec{E} + \vec{v} \times \vec{B})$. Separating these charges creates radiation with linear polarization that is orthogonal to the magnetic field and shower axis.

Using the Earth's magnetic field, \vec{B} , the shower maximum, and shower direction, we can calculate the expected polarization angle for an EAS. To get the shower maximum we must trace an event from the payload to the continent, reflect it off the ice and trace that back to our shower maximum, then we must use a model of the Earth's magnetic field to get \vec{B} at that point. Expected polarization angle for up-going τ neutrino candidates is calculated similarly, but without reflection off of the ice. This usually modifies the expected angle by only a small amount, but can be a large effect.

4.13.3 Linear Polarization Fraction

Radiation from an EAS will be highly linearly polarized. The ANITA-III analyses found all detected EAS events to have a linear polarization fraction above 60% [80]. We can use this information to put constraints on the expected linear polarization fraction of any EAS we might see. All three of the aforementioned EAS specific cuts will reduce expected thermal and anthropogenic backgrounds.

CHAPTER 5

RESULTS

Using the methods laid out in the previous sections, I set my final cuts, created a background estimate, and unblinded the signal region. Results of each of these steps are presented quantitatively in the following chapter.

5.1 Background Estimates

Before unblinding our signal region we need to have an estimate of the number of background (non-signal) events that may have leaked into our signal region. ANITA's Askaryan search has two sources of background, thermal and anthropogenic events. The two backgrounds are estimated independently and combined. The EAS channel τ neutrino search has these two sources of background, as well as two additional sources of background, misidentification of event polarity and misidentification of an event from above the horizon as coming from below the horizon.

In order to make a background estimate while remaining blind, we need to have a sideband, which is a region of phase space where we expect zero signal, but we expect background events from which we can build a model to estimate background in the signal region. The sideband region is different depending on whether we are estimating thermal or anthropogenic backgrounds. It also is different between the Askaryan neutrino and EAS signal channels.

I detailed the thermal background estimate in Chapter 4.11. After thermal cuts, I expected 0.05 events leaking into the clustering sample, which is further reduced by clustering cuts, and then Askaryan neutrino and EAS cuts. In the end, the background contribution from thermal events is negligible compared to anthropogenic backgrounds (10^{-7} expected events).




Figure 5.1: Vertically polarized scope trace taken using a Seavey antenna pointed at a snowmobile. Image courtesy of David Saltzberg and Konstantin Belov [127].

5.1.1 General Anthropogenic Background Estimate

It's been clear since the inception of ANITA that anthropogenic events can appear indistinguishable from signal-like events. In 2008, David Saltzberg and Konstantin Belov recorded impulsive waveforms using an antenna pointed at a snowmobile that could, at least by eye, pass for a neutrino candidate [127]. A snowmobile waveform is pictured in Fig. 5.1. These types of events are typically produced many at a time with spatially co-located origins on the continent. However, there is a chance we record an event that is both anthropogenic and isolated, that would appear in the Askaryan neutrino or EAS signal region. An anthropogenic event can appear isolated because it actually is isolated, the classic ANITA example is a cross-country skier with a lighter, or because our imperfect trigger efficiency only recorded one event from a burst of a few. No matter the cause, isolated anthropogenic events that appear signal like are an irreducible background, so we need to have an estimate of how many we might have before opening the signal box.

The general prescription we use for estimating anthropogenic backgrounds is a standard

technique in particle physics for background estimation known as the ABCD method, similar to the On-Off problem in astronomy [128]. We plot all events as a function of two uncorrelated variables, V_1 and V_2 , and divide that phase space into four regions, A , B , C , and D . Region A is the area we expect to find all of our signal and remains hidden. Regions B , C , and D are as follows:

- Region B – signal-like in V_1 and not signal-like in V_2
- Region C – not signal-like in V_1 and signal-like in V_2
- Region D – not signal-like in either variable.

An illustrative plot of this is shown in Fig. 5.2. Since V_1 and V_2 are uncorrelated, $\frac{N_A}{N_C} = \frac{N_B}{N_D}$, where N_A is the number of events in region A , etc., for anthropogenic events. Solving this equation for N_A gives you the expected anthropogenic leakage into the signal region. We can calculate the uncertainty on this number by treating N_B , N_C , and N_D as Poisson-distributed variables, where the mean of their distribution is the detected number of events in that region. Then we repeatedly draw from these distributions to build up a distribution for N_A and quote the central 68% CL interval.

Previous ANITA analyses have used how isolated an event is, and whether it is or is not from a known location of human activity as V_1 and V_2 [84, 86, 88]. Which events are included and how to categorize them have changed from analysis to analysis. The strategy I’ve chosen for ANITA-IV is a little different from previous analyses.

5.1.2 EAS Channel: Cosmic Rays and Possible Inverted EAS Events

For the EAS signal channel ABCD background estimate, the two uncorrelated variables we use are isolation and “EAS-ness”. The EAS analysis is restricted to horizontally polarized events, which are events in which the most interesting peak was chosen from the horizontally




Figure 5.2: An illustrative general example of the ABCD method of background estimation. Region A is signal-like. Regions B, C, and D are background. V_1 and V_2 are uncorrelated for background events. The prescription for estimating the background using this method is detailed in the text.

polarized interferometric map, where picking the most interesting peak is done as detailed in Chapter 4.5.1. An event’s isolation is defined as isolated (singlet) or not isolated (in a cluster), as detailed in 4.12. The EAS-ness of an event is also a binary variable, consisting of whether an event passed cosmic ray template cuts, a geomagnetically induced polarization angle cut, and a linear polarization fraction cut. If an event passes all of these cuts it is considered EAS-like, if it fails any of them, then it is not EAS-like. The threshold to consider an event EAS-like was: > 0.55 for coherently summed template correlation, > 0.7 for dedispersed template correlation, $< 12.5^\circ$ away from the expected geomagnetic polarization angle, and a value of linear polarization fraction > 0.625 . Because we do not have a cosmic ray simulation, these cutoffs were chosen using WAIS calibration pulses and EAS events from previous ANITA flights as stand-ins for EAS events.

Correlation Cut Values

In order to pick a template cut value I used a sample of WAIS events. First I made a WAIS template using 100 high SNR WAIS events, then I correlated that against other WAIS events

to see the distribution of correlation values. I set the cut for coherently summed events first. Fig. 5.3 shows that a cut at 0.66 keeps 99.7% of WAIS events. With this 0.66 value in hand, I wanted a measure of how to scale this by how dissimilar cosmic ray waveforms could be.

WAIS pulses were transmitted on boresight of the transmitting antenna and with the same pulser, so there should be no variation of transmitted signal, but the cosmic ray signal varies a lot depending on the viewing angle from the center of the cone. To measure this dissimilarity, I correlated all of the ZHAires templates [126] against the center template (the 5th of 9 templates, where the 1st template is exactly on-cone and the last template is 0.32° off-cone). The central template was least similar to the most on-cone and least on-cone templates, with a value of $\sim 82\%$. To set my cut, I multiplied the 0.82 value I found from this template variation by the 0.66 value found from the WAIS correlation distribution to find a cut at ~ 0.55 for the coherently summed waveforms.

I followed a similar procedure to find the correlation coefficient cutoff for the dedispersed coherently summed waveforms. The cutoff for this correlation value is much higher, because we window the template about the peak of Stokes I before correlating. With windowing we cut out the noisy samples at the beginning and end of the waveform that lower correlation values, so a similar signal looks more similar.

Geomagnetic Polarization Angle Cut Values

Ben Rotter found 20 EAS events in his ANITA-III analysis [80]. His distribution of difference from the expected geomagnetic polarization angle showed a standard deviation of 4.27° when fit with a gaussian, as shown in Fig. 5.4. I chose to set my cut at $\sim 3\sigma$ of the expected linear polarization angle, or within 12.5 degrees of the expected angle. In future analyses, this number should vary with each event, by calculating an expected error on the polarization angle.




Figure 5.3: The distribution of correlation values for the coherently summed WAIS pulser events. The procedure I followed to set a cut using this distribution is detailed in the text.




Figure 5.4: The distribution of differences from the expected geomagnetic polarization angles for the ANITA-III EAS candidate events. This distribution and the gaussian fit to it was used to determine the geomagnetic polarization angle difference cut on ANITA-IV EAS events. Figure taken from [80].




Figure 5.5: The distribution of linear polarization fraction values for the WAIS pulses. Using this distribution as a guide, I set the cut on EAS linear polarization at 0.625.

Linear Polarization Fraction Cut Value

The ANITA-IV EAS sample should generally contain a higher linear polarization fraction than the ANITA-III sample did, because we have much more precise timing alignment in the Hpol and Vpol channels for each antenna. Because of this, I used WAIS pulses rather than the ANITA-III sample to set my cuts. WAIS pulses were meant to mimic physics signals, and should have comparable levels of linear polarization. As detailed in Chapter 4, I found that a combination of linear polarization metrics was more effective at finding truly linearly polarized signals than any single one, so I took the average of four metrics (windowed and unwindowed for both the coherently summed and coherently summed dedispersed waveforms). Fig. 5.5 show the linear polarization fraction distribution for WAIS pulses. I chose to cut at 0.625, because it kept 99.7% of WAIS pulses.

CR Cut Results

My analysis only considered events in small clusters to be part of the distribution that could leak into the signal region, the reasoning and cutoff for this is gone through in detail in Chapter 5.1.4. Upon unblinding my signal region, I had 30 events in the box that passed EAS-like cuts. Shortly after, I found a bug in the calculation of the expected geomagnetic polarization angle that would have resulted in 2 of those 30 events moving to the sideband region. After fixing this bug, I had three isolated, horizontally polarized events that did not pass EAS-like cuts (more on those later), giving a background estimate of $0.35^{+0.23}_{-0.15}$ events, evenly divided across polarity, at $-2\log(\mathcal{L}) = 6$. It's possible that I would have been able to further tune my cuts to get a lower background estimate without this bug, but because I had already unblinded the EAS signal region, any further tuning of cuts would have been highly biased, and I chose to go forward with the cuts chosen pre-unblinding. In addition to the statistical errors derived as described in Chapter 5.1.1, these errors also fold in a systematic error calculated by varying the cluster size threshold to 50 or 1000, which turns out to be sub-dominant to the statistical error.

5.1.3 Askaryan Neutrinos

Similar to the EAS channel, for the Askaryan neutrino channel I only considered events whose interesting peak was chosen from the vertically polarized interferometric map. The uncorrelated variables used for the Askaryan neutrino channel background estimate were isolation and total polarization fraction. Total polarization was chosen instead of linear polarization because Askaryan radiation may produce a non-negligible amount of circular polarization, but the total should be highly polarized. I chose the cut for required total polarization fraction for an event to be considered neutrino-like at > 0.92 by using the distribution of total polarization for simulated neutrinos. I maximized the expected sensitivity of the analysis to pick a final value for $-2\log(\mathcal{L})$ of 8. With one isolated event below the

total polarization fraction threshold, I arrived at a background estimate of $0.64_{-0.45}^{+0.69}$. Such a large background estimate means I would need to detect five events to be able to claim detection at 3σ .

5.1.4 Alternative Background Estimate

In addition to the standard ABCD method of background estimation, I developed a second method to get a handle on our anthropogenic background using the WAIS pulses. My first goal with this was to determine how large a cluster needed to be, in terms of number of events, in order to have zero expected events leaking out. The clustering algorithm works by putting all events that cluster together into a pseudo-base. To determine if an event belongs in a pseudo-base, it is compared with all events until it finds one that it clusters with, or, if there are no events it clusters with, it forms its own new pseudo-base. More events in a pseudo-base means more chances to cluster, or a larger footprint on the continent, which means that events that should be associated with a large population pseudo-base are less likely to falsely be considered isolated than events that should be associated with a pseudo-base of smaller population.

I found the cutoff for event leakage by taking a random sampling of N WAIS pulses and attempting to cluster them together. If any events leaked after 10^4 trials, I tried doing this with $N + 1$ WAIS events. The population required to no longer expect any singlet leakage was 125 for the log-likelihood range of interest. This cutoff is important for the ABCD background estimation as well, because events from clusters with a larger population than this should not contribute to our expected background and therefore should not be included in the ABCD calculation.

Using this cluster size threshold I calculated the approximate expected background at each log-likelihood cut value. I summed the average number of singlets that would leak out from each cluster using all clusters below the population threshold, with an error calculated

from the standard deviation of the number of leaking singlets. I found $-2\log(\mathcal{L}) = 6$ to produce 1.08 ± 0.86 total expected singlets, $-2\log(\mathcal{L}) = 8$ to produce 0.58 ± 0.51 total expected singlets, and $-2\log(\mathcal{L}) = 10$ to produce 0.54 ± 0.49 total expected singlets. For each polarization this would produce \sim half the expected singlet total, which would then be further reduced by Askaryan neutrino and EAS specific cuts. This estimate is less believable than the other estimate because it lacks a way of accounting for “true” anthropogenic singlets. It is mostly useful for setting a population threshold and having a second check on the expected background.

5.1.5 Isolated Background Singlets

I found three isolated background singlets that did not land in the signal region of either the EAS or Askaryan neutrino search channels. One of these was nearly equal parts Hpol and Vpol, so I allowed it to act as a background singlet in both the EAS and Askaryan search ABCD background estimates. The waveform and power spectrum of this event are shown in Fig. 5.6. This event has no other events within about 45 km. The nearest event is the Askaryan neutrino candidate event. There are no known places of human activity near by. It appears to be a “true” anthropogenic singlet event.

The other two events were predominantly Hpol and acted only as sideband singlets for the EAS search. They are discussed further in the following section.

5.2 EAS Channel (Hpol) Unblinding

When I opened the EAS signal region, I found 30 candidates. After fixing a bug with the calculation of expected geomagnetic polarization angle, two of those candidates moved to the sideband region and were folded into the background estimate. The first event, 17904564, is an interesting event, and is discussed in more detail in Chapter 5.3.1. The other event, 50549772 also warrants more investigation and is discussed along with the other candidates,




Figure 5.6: The power spectrum and waveforms for the dedispersed and coherently summed background singlet event are shown. Vertically polarized signal is red and horizontally polarized signal is blue. This is clearly not an EAS or Askaryan neutrino candidate, but it appears very isolated.

even though it is included in the sideband region. It is pictured in Fig. 5.7. Of these 28 candidates, 2 come unambiguously from above the horizon. Candidate CR-like variable values are listed in Table 5.1. The efficiency of the EAS search is unknown because we don't have an EAS simulation. The EAS locations on the continent are shown in Fig. 5.8. Each candidate's difference from the expected geomagnetic polarization angle is shown in Fig. 5.9. Waveforms are shown side by side for coherently summed waveforms in Fig. 5.10 and for dedispersed waveforms in Fig. 5.11. Individual waveforms for each candidate are shown, both coherently summed and dedispersed, along with power spectra, in Figs. 5.12 through 5.39.

Event	Coh. Corr. ^a	Dedisp. Corr. ^b	Geomagnetic Diff. ^c	LPF ^d	θ Angle ^e
4098827	0.76	0.76	10.29	0.80	6.17
9734523	0.84	0.93	6.40	0.89	5.68
12131787	0.84	0.98	0.62	0.91	17.22
15738420	0.85	0.97	5.24	0.91	7.56
16821419	0.74	0.96	7.27	0.91	11.27
19848917	0.60	0.79	11.42	0.64	6.73
20936205	0.83	0.98	0.73	0.93	13.59
25580797	0.76	0.97	0.81	0.94	22.31
25855454	0.80	0.98	2.56	0.85	13.16
36785931	0.71	0.83	4.40	0.93	36.44
39236841	0.82	0.96	0.87	0.91	8.39
40172984	0.83	0.97	5.62	0.96	21.77
45684620	0.81	0.97	1.17	0.98	17.33
47396999	0.84	0.99	1.00	0.94	15.27
51293223	0.79	0.95	1.15	0.87	5.46
54063721	0.71	0.95	7.77	0.89	9.88
64472798	0.75	0.96	2.28	0.90	26.11
64859493	0.68	0.90	6.28	0.93	19.97
64861754	0.73	0.97	2.23	0.87	21.88
66313236	0.71	0.97	2.41	0.83	16.23
66509677	0.77	0.96	0.58	0.98	8.88
72164985	0.85	0.94	7.42	0.93	6.16
74197411	0.58	0.71	6.88	0.85	14.77
83074427	0.71	0.94	1.20	0.97	7.35
88992443	0.63	0.74	5.21	0.84	14.49
91525988	0.77	0.93	5.25	0.80	18.70
93744271	0.77	0.95	3.60	0.95	29.97
95576190	0.78	0.96	0.52	0.80	13.65
50549772	0.56	0.79	25.11	0.63	6.43

^a Coherently summed waveform correlation with the EAS template

^b Dedispersed waveform correlation with the EAS template

^c Difference between measured and expected polarization angle

^d Linear polarization fraction

^e Angle from Horizontal

Table 5.1: The values for each EAS candidate event for CR-like cut variables. Event 50549772 is reported at the bottom separately.




Figure 5.7: Event 50549772, which was moved from the signal region to the sideband region after a bug was found in the geomagnetic angle calculation code. The coherently summed waveform is in the top panel, the dedispersed waveform is in the middle panel and the power spectrum is in the bottom panel. The dedispersed waveform has a strange second peak after the first peak, which is different from the normal tripolar structure of an EAS, also the power spectrum peaks in the mid-band, which is unexpected for an actual EAS event.

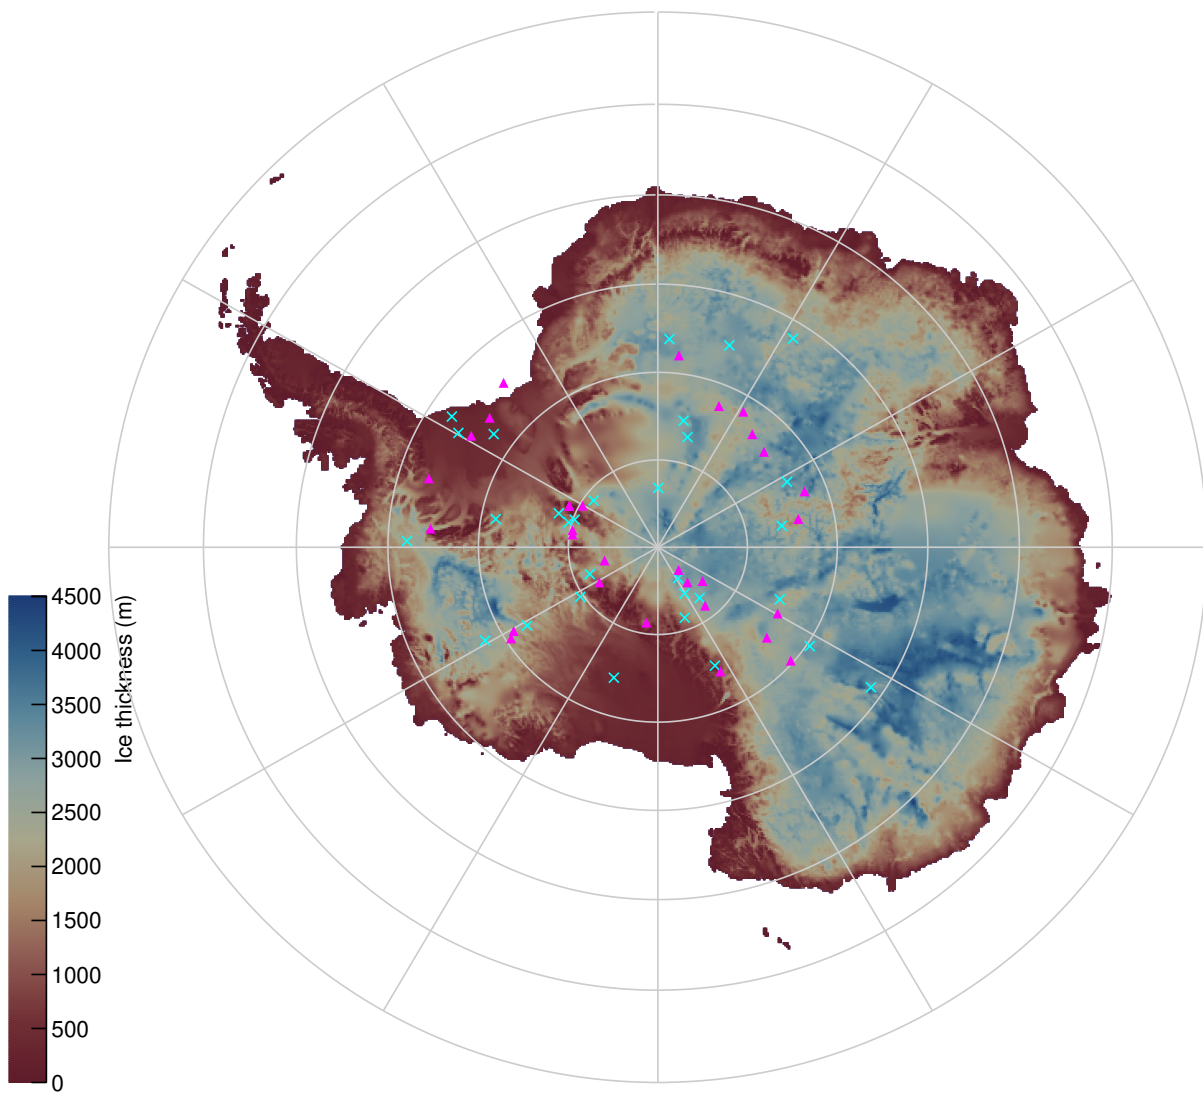


Figure 5.8: The approximate location of EAS candidates on the continent are plotted (cyan X) along with the payload viewing location (magenta triangle).




Figure 5.9: The difference between the expected geomagnetically induced polarization angle and the measured polarization angle for EAS candidates. Cyan events are reflected cosmic rays, black events are direct cosmic rays, and other events are colored individually as denoted in the legend. The two events that were moved out of the signal region, 17904564 and 50549772 are the farthest from the central black line, in red and pink, respectively. Vertical error bars are based off of the ANITA-III cosmic ray events, error on expected polarization angles are estimated at 2° , from uncertainties in geomagnetic angle calculations.

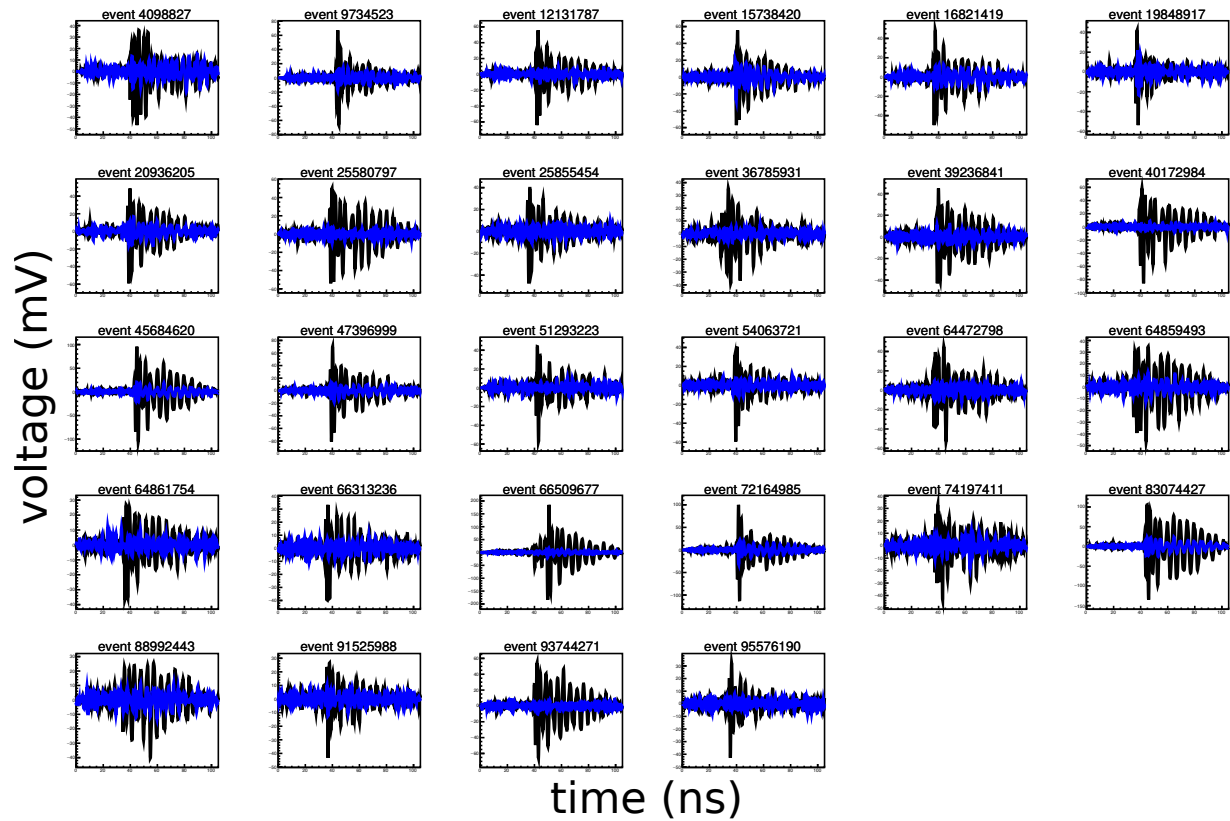


Figure 5.10: All 28 of the EAS candidate's coherently summed waveforms. Black is the horizontally polarized channel, blue is the vertically polarized channel.

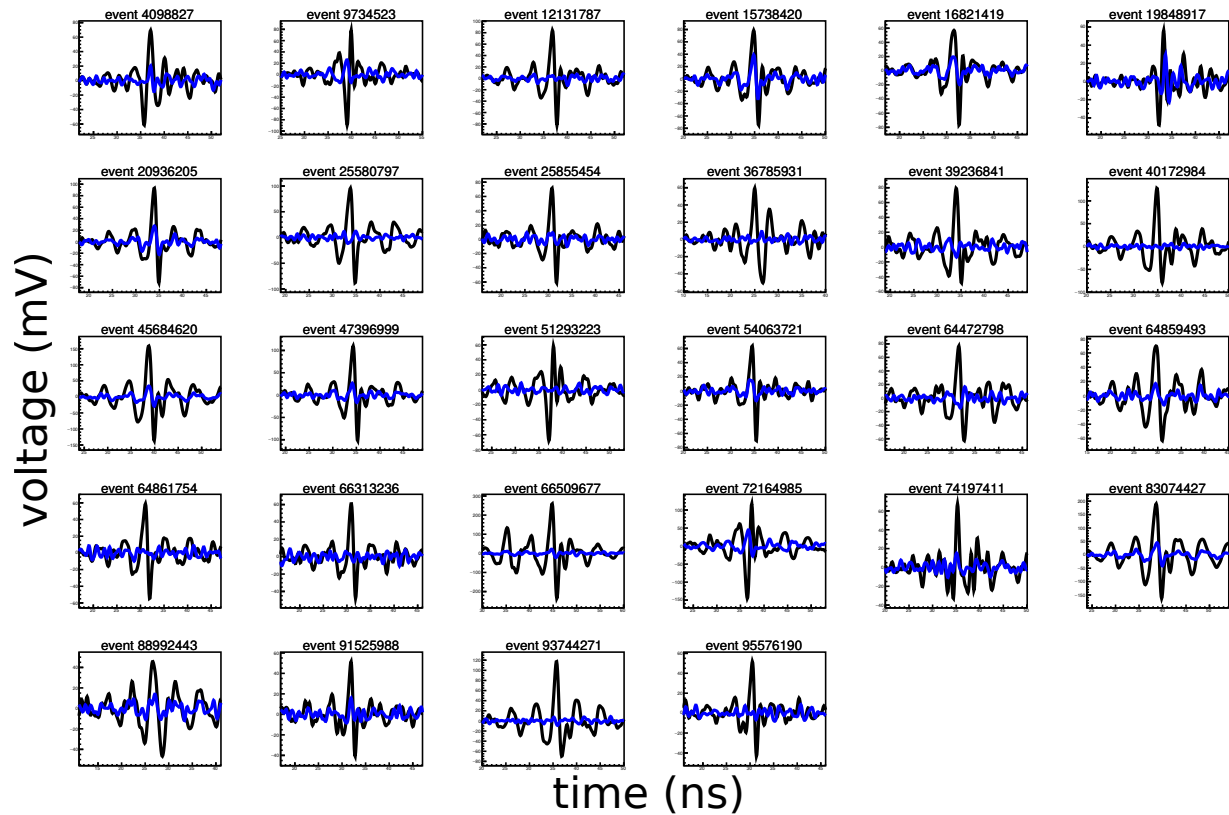


Figure 5.11: All 28 of the EAS candidate's dedispersed coherently summed waveforms. Black is the horizontally polarized channel, blue is the vertically polarized channel. The window is ± 15 ns about the peak intensity.




Figure 5.12: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is one of two CRs detected with all three notches on, and has a distinct frequency spectrum that peaks a bit higher than most.




Figure 5.13: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is one of two direct events.




Figure 5.14: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.15: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.16: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.17: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is one of the unusual polarity events.




Figure 5.18: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.19: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.20: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.21: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. This is the second of two CRs detected with all three notches enabled, and shows a similar spectrum to the other one.




Figure 5.22: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.23: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.24: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.25: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.26: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. The second of two direct cosmic rays.




Figure 5.27: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.28: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.29: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.30: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.31: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.32: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.33: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. The second of two unusual polarity events.




Figure 5.34: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. A low SNR cosmic ray that almost looks unipolar in shape.




Figure 5.35: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.36: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity. A very off-cone cosmic ray. Notice how dominant the low frequencies are in the power spectrum, and how different the coherently summed template appears from the others.




Figure 5.37: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.38: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.




Figure 5.39: Black is horizontal polarization, blue is vertical polarization. Top panel is coherently summed waveform, middle panel is dedispersed waveform, bottom panel is power spectrum. Dedispersed waveform is zoomed in ± 15 ns about the peak intensity.

5.2.1 Further Analysis of Candidates

With the signal box open, we can estimate polarity for each candidate. We have three ways to estimate polarity at the time of this writing:

1. Basic correlation of a signal with a template, where the largest magnitude value is chosen to be the polarity, this can be with the coherently summed waveform or the dedispersed waveform, with or without a fiducial reference, and can be windowed or unwindowed.
2. The peak-to-sidelobe ratio (PSL), which is the ratio of the magnitudes of the largest positive and negative peaks of the correlation.
3. The Fourier phase method, where a fiducial reference is defined (nearest zero-crossing to the peak of Stokes I), the waveform is rolled so that it starts and ends at the fiducial reference, and the phase of that new signal is inspected between 200 and 600 MHz. The quantity of interest is whether that phase is positive or negative.

After much discussion and refinement, we chose to use the PSL of the unwindowed, coherently summed waveform, the PSL of the windowed, dedispersed waveform, and the Fourier phase method to define polarity of our candidates. We found these three measures to give equivalent polarity for all of the EAS candidates except one. Windowed, dedispersed candidates are shown aligned to the EAS template, along with the zero-crossing reference point in Fig. 5.40. Candidates windowed and dedispersed cross-correlations are shown in Fig. 5.41, and results of the Fourier phase method for all candidates are shown in Fig. 5.42. PSL are correlation values for all candidates on one plot are shown in 5.43.

Two candidate events were found to have polarity consistent with that of a direct EAS, while coming from below the horizon. Both events, 19848917 and 72164985, are much closer to the horizon than previous ANITA “mystery events” [77, 78]. Event 19848917 is about 0.7° from the horizon, while 72164985 is a mere 0.2° from the horizon. Of the two events,

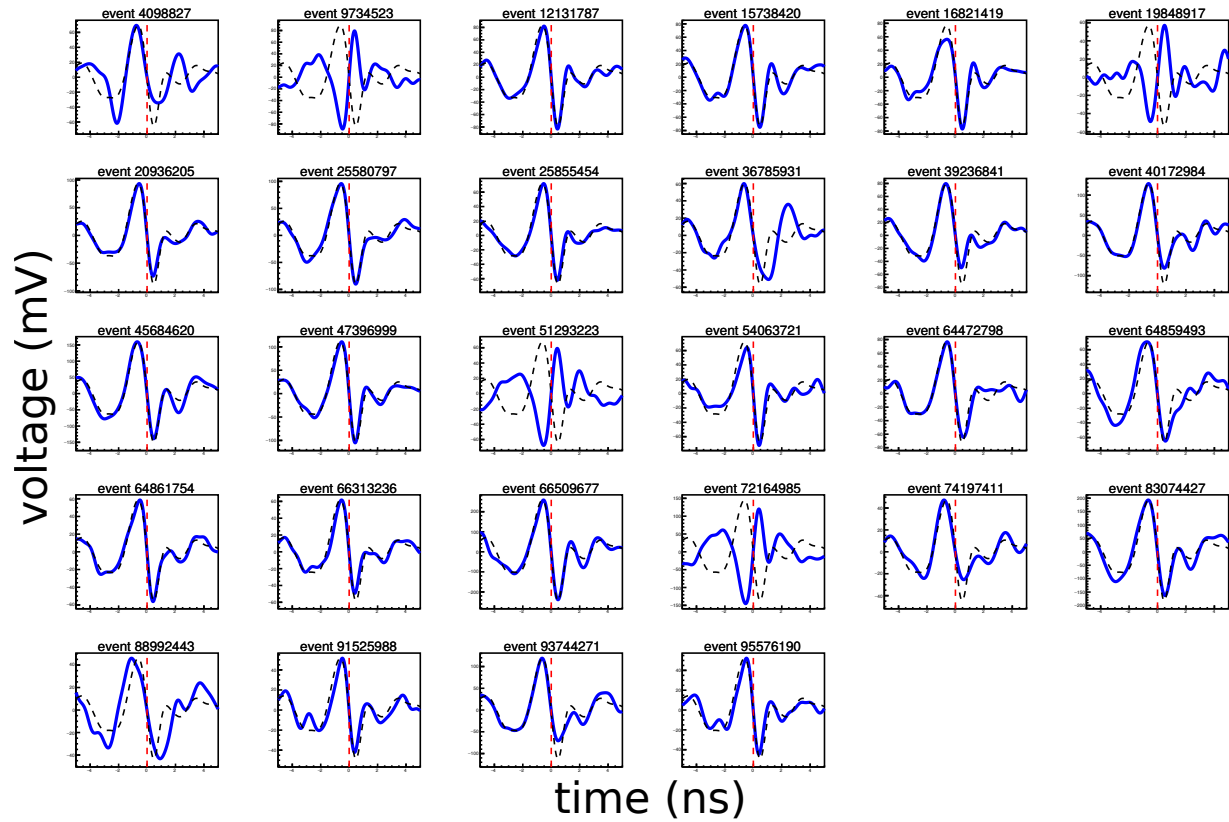


Figure 5.40: Candidate waveforms are pictured in blue, with template waveforms as black dotted lines. Waveforms are dedispersed and windowed ± 5 ns about the zero-crossing nearest the peak of Stokes I. Red dotted lines are the reference point the waveform and template are aligned to. The template is aligned to the reflected CR polarity.

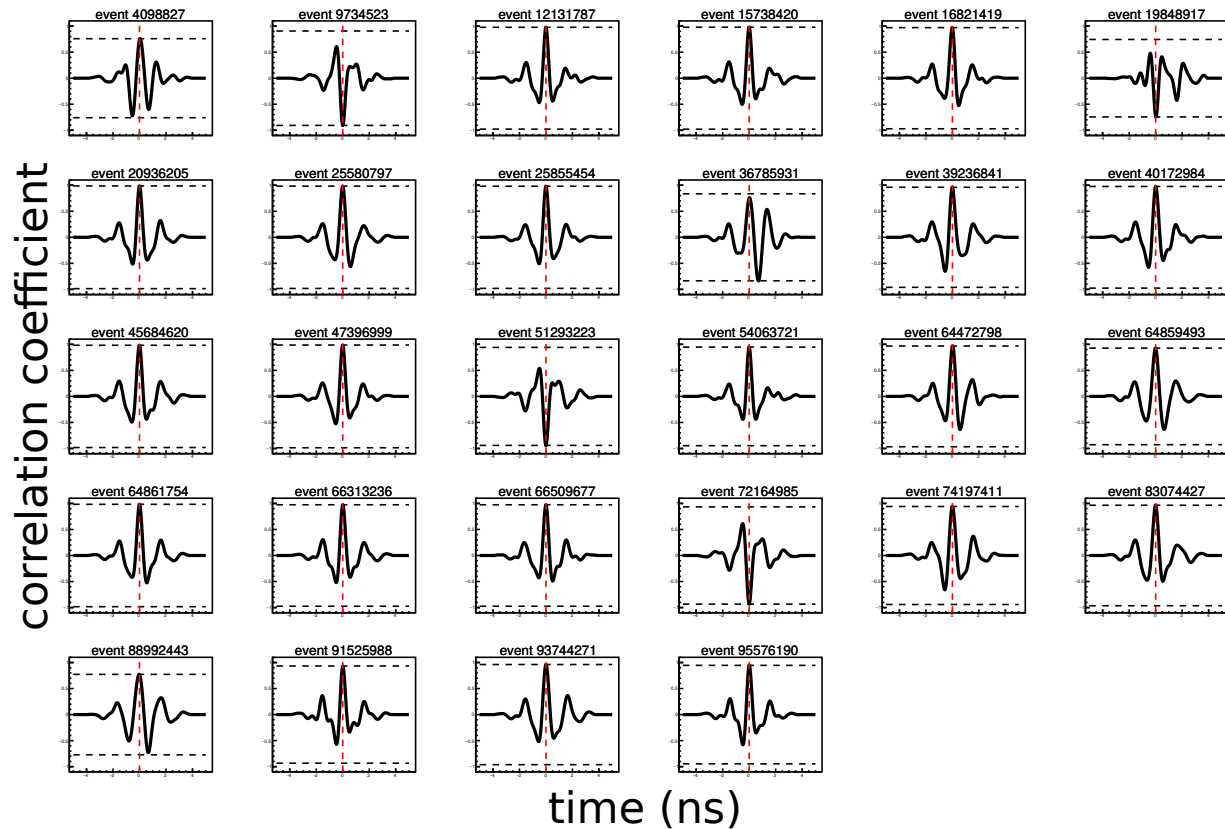


Figure 5.41: The windowed, dedispersed correlation for each EAS candidate. Windowing is done ± 5 ns from the zero-crossing nearest the peak value of Stokes I for the waveform. A positive cross-correlation value at the dotted red line indicates a normal reflected CR polarity by the windowed correlation metric, a negative value indicates polarity consistent with a direct CR. Dotted black lines are at \pm maximum absolute cross-correlation coefficient values, this can be used to estimate the PSL value. The astute reader will note that event 36785931 has a PSL value that opposite what the fiducial indicates for polarity. We concluded that this is due to distortion of the waveform caused by having all three notches on for this event, and categorize it as a normal, reflected CR, as two out of three metrics indicate a reflected CR polarity.

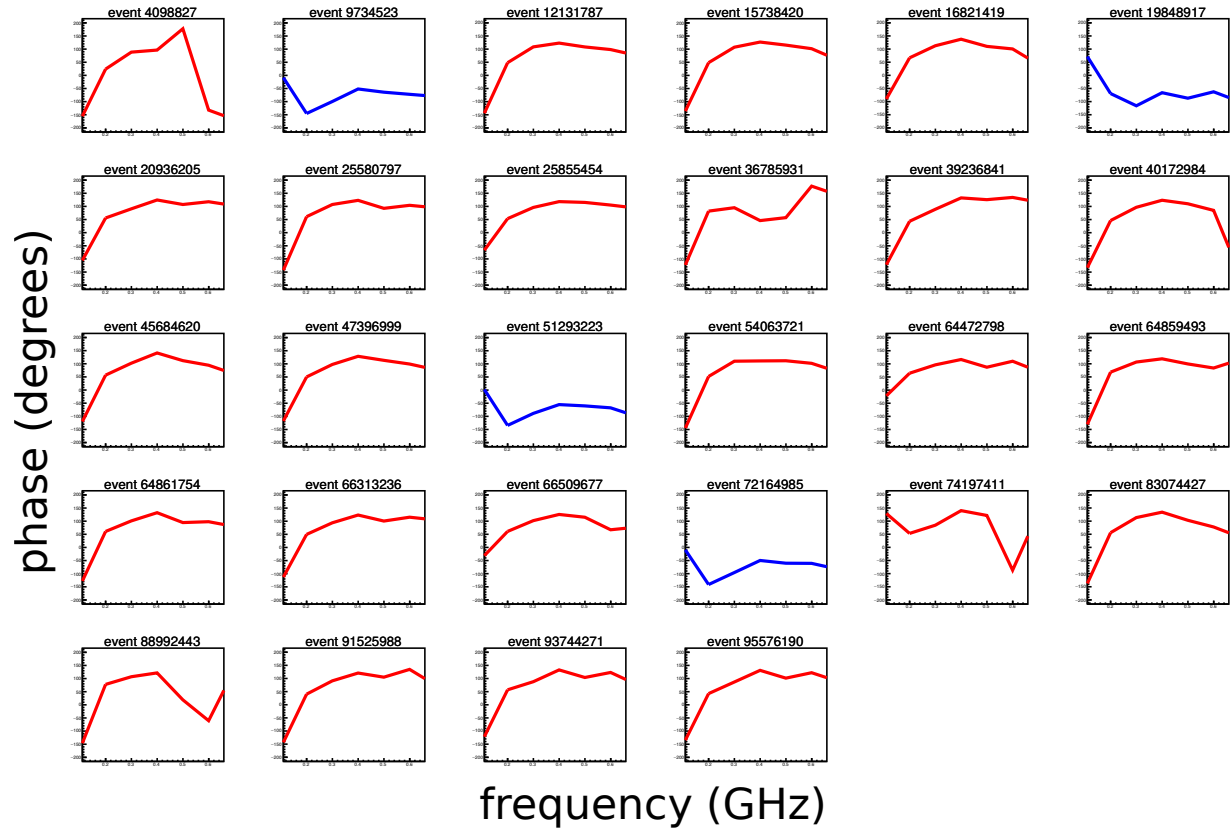


Figure 5.42: Results of the fourier phase method for each EAS candidate. The plots are color-coded so that red (positive phase) is reflected polarity and blue (negative phase) is direct polarity. The fourier phase method is described in the text, and the reference point of the waveform is shown in Fig. 5.40.




Figure 5.43: The PSL is plotted vs the correlation value vs a template. On the left in blue are the EAS candidates with polarity consistent with a direct CR, on the right in red are the EAS candidates with polarity consistent with a reflected CR. The three on the right with low PSL values are the two events with all three notches on, 4098827 and 36785931, and 88992443, which appears to be very off-cone. Event 36785931 has a PSL below 1 as well.

only 72164985 is well above the cutoff to be considered an EAS candidate. From Table 5.1 it's clear that event 72164985 is far from threshold in all four cut variables, while event 19848917 is near threshold in two of the four cuts, linear polarization fraction and difference from expected geomagnetic polarization angle. Event 50549772 also appears to have polarity consistent with a direct event, while pointing to the continent, but is left out of the following analysis because it failed the geomagnetic polarization angle cuts.

There are three ways an event can be miscategorized as a mystery event: an event could actually be background, an event could have actually come from above the horizon, or an event could have its polarity miscategorized. To determine how often we would expect these events to be anthropogenic background, we use a toy Monte Carlo simulation to draw from the distribution for our anthropogenic background, then draw from a Poisson distribution with the mean we drew. This toy simulation produces two or more CR-like events with this polarity 3.4% of the time. We conservatively bound the probability of events actually coming

from above the horizon based on distributions we measured using events seen coming from the South Pole. The third possibility, polarity miscategorization, is also estimated by toy Monte Carlo.

The chance of polarity flipping is different for each candidate. It is a function of SNR, as well as the frequency content of the waveforms, which is governed largely by the viewing angle relative to the radiation emission cone. I estimate the chance of flipping on an event by event basis for all of the candidates by finding the three most similar simulated cosmic ray templates by correlation coefficient, taking those templates, scaling them to the event SNR, adding them to nearby minimum-bias triggers, checking the polarity, and then repeating for as many times as desired. The measured polarity values and estimated chance to misidentify polarity for each candidate are shown in Table 5.2.

All three of these effects contribute to background in this channel. To quantify the background from all of these effects, we combined the Poisson background distributions, the pointing error distributions from the South Pole, and the polarity misidentification chance to create 10^5 independent trials of events with the same parameters as our measured events. We used these trials to estimate how often we would expect to see something equally or more anomalous. The estimated significance of the mystery events is shown in Table 5.3. This table includes two variations, a horizon buffer and disallowing event 72164985 to be background. The horizon buffer of 0.1° is due to the estimated Fresnel zone size for an event with this frequency content. It's unclear whether Fresnel zones effect impulsive events in the same way they effect CW, so it may or may not be a realistic scenario. The other variation, disallowing event 72164985 to be considered an anthropogenic background, is included because event 72164985 is very unlikely to be an anthropogenic background event. Only 1 event out of 366,495 impulsive, horizontally polarized anthropogenic events that pass my thermal cuts are as or more CR-like than event 72164985. For event 19848917, that same sample of events contains ~ 16000 events as or more CR-like.

Event	PSL (Mis-ID)	Dedispersed PSL (Mis-ID)	Fourier Phase (Mis-ID)
4098827	+1.25 (0)	+1.05 (0.003)	+ (0.003)
9734523	-1.51 (0)	-1.50 (0)	- (0)
12131787	+1.31 (0)	+2.09 (0)	+ (0)
15738420	+1.29 (0)	+1.95 (0)	+ (0)
16821419	+1.07 (0.0001)	+1.85 (0.004)	+ (0.004)
19848917	-1.25 (0.0003)	-1.55 (0.055)	- (0.051)
20936205	+1.24 (0)	+1.92 (0)	+ (0)
25580797	+1.12 (0)	+1.75 (0)	+ (0)
25855454	+1.32 (0)	+1.93 (0)	+ (0)
36785931	+1.22 (0)	-1.09 (0)	+ (0)
39236841	+1.41 (0)	+1.47 (0)	+ (0)
40172984	+1.32 (0)	+1.70 (0)	+ (0)
45684620	+1.18 (0)	+1.99 (0)	+ (0)
47396999	+1.34 (0)	+1.88 (0)	+ (0)
51293223	-1.43 (0)	-1.75 (0)	- (0)
54063721	+1.09 (0)	+2.15 (0.025)	+ (0.03)
64472798	+1.13 (0)	+1.52 (0)	+ (0)
64859493	+1.06 (0.044)	+1.46 (0)	+ (0.0001)
64861754	+1.08 (0.0001)	+1.87 (0.002)	+ (0.0005)
66313236	+1.16 (0.0001)	+1.87 (0)	+ (0)
66509677	+1.08 (0)	+1.99 (0)	+ (0)
72164985	-1.81 (0)	-1.52 (0)	- (0)
74197411	+1.23 (0)	+1.43 (0)	+ (0)
83074427	+1.01 (0)	+1.95 (0)	+ (0)
88992443	+1.23 (0.061)	+1.06 (0.093)	+ (0.085)
91525988	+1.38 (0)	+1.64 (0)	+ (0)
93744271	+1.16 (0)	+1.86 (0)	+ (0)
95576190	+1.45 (0.0003)	+1.63 (0.017)	+ (0.031)
50549772	-1.21 (0.001)	-1.46 (0.069)	- (0.079)

Table 5.2: The values for each EAS candidate event for polarity metrics, with their respective chances to have their polarity misidentified in parentheses. A 0 for misidentification chance means that there were no misidentified polarities in 10000 independent trials for that candidate event. The coherent correlation, PSL and Fourier phase are positive for normal reflected CR polarity and negative for direct CR polarity. Only one of the events has disagreement between the three polarity metrics, event 36785931. Event 50549772 is shown separated by a horizontal line at the bottom.

Anthropogenic Background Scenario	Horizon Buffer	Probability (sigma)
Both Possible	0°	5.5% (1.6 σ)
Both Possible	0.1°	4.7% (1.7 σ)
72164985 EAS	0°	3.8% (1.8 σ)
72164985 EAS	0.1°	2.3% (2.0 σ)

Table 5.3: Probabilities for the two mystery event candidates to be background under several different sets of assumptions. Scenarios and assumptions are more fully described in the text.

The probability of these three sources of background fluctuating to 2 or higher is between 2.3% and 5.5%, depending on assumptions.

5.3 Askaryan Neutrino Channel (Vpol) Unblinding

The total analysis efficiency for the Askaryan neutrino search was $82 \pm 2\%$, with the largest efficiency losses coming from clustering, and error coming from comparison to results with calibration pulses. The number of events remaining in the analysis as well as cut efficiencies on simulated data are shown in Table 5.4, while Fig. 5.44 gives a sense of which locations on the continent are effectively lost to imperfect clustering. Event 36019849 was the lone surviving event, its waveform and power spectrum are shown in Fig. 5.45. This event is very isolated in both distance and log-likelihood. The nearest events to event 36019849 are event 35278567 (45 km away), and event 79447264 (654 log-likelihood units away), neither of which correlate well with the candidate. The nearest known location of human activity is a fixed-wing landing site that is about 88 km distant from the event (we did not detect any events from this site). There are no other known locations of human activity within 100 km.

5.3.1 Posterior Analysis of Candidates

Event 36019849 is consistent with the expected properties of a neutrino in terms of signal shape, polarization, and exit angle. However, it originates from a place on the Ross Ice Shelf where the ice is only about 300 m deep. Fig 5.47 shows the location of the event. Only $\sim 1\%$

Cut Name	Vpol Events Remaining	Simulated Neutrino Efficiency
None	52242901	1.0
Quality	37408254	0.981
Thermal	575067	0.978
Clustering	1	0.819

Table 5.4: Summary of the effect of cuts in my analysis. The quoted efficiency is the cumulative efficiency on Monte Carlo generated neutrinos for all cuts when performed in sequence. Quality, thermal and clustering cuts are detailed in Chapter 4. The Askaryan neutrino analysis is left with one event in the signal region at the end of all of these cuts.

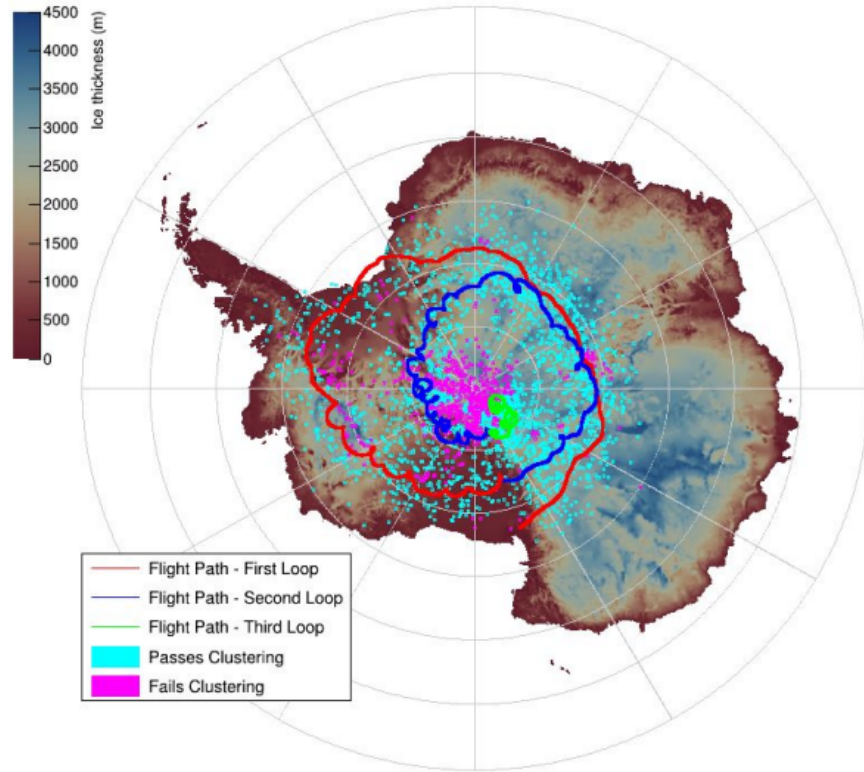


Figure 5.44: ANITA's flight path is shown overlaid with simulated neutrinos passing or failing clustering cuts with $-2\log(\mathcal{L}) = 8$. Simulated events that fail clustering as shown in magenta and passing events are shown in cyan. The bulk of efficiency loss comes from events near South Pole Station.




Figure 5.45: Event 36019849 (the lone neutrino candidate in the Askaryan channel), dedispersed and coherently summed, with the vertically polarized waveform in red and horizontally polarized waveform in blue. It is clearly a very impulsive, vertically polarized signal. The bottom panel shows the power spectrum of the top panel in dB.

of simulated neutrinos come from places on the continent with ice the shallow or shallower. One more wrinkle is that the polarity of this event is opposite that of simulated neutrinos. At this location, it is possible the event reflected off of the water before being detected, which would resolve the tension in polarity.

Because of the unusual nature of event 17904564, I decided to investigate its viability as a neutrino candidate. Its waveform is shown in Fig. 5.46. It also comes from an ice shelf, and is also flipped polarity from what is expected, so it is very similar to event 36019849. The nearest events are 17051644 (44 km) and 58980595 (186 log-likelihood units). Event 17051644 is from a cluster of 11 events, only one of which has a correlation coefficient with the candidate of > 0.55 .

In order to try to quantify how “neutrino-like” event 36019849 is, I performed a posterior likelihood analysis. Using simulated neutrino events for the signal distribution and events that passed thermal cuts but failed clustering cuts as the background distribution, we compared whether the candidate event was more signal-like or more background-like across five variables. The variables we chose were impulsivity, linear polarization fraction, power window gradient, interferometric map peak, and peak value of the Hilbert envelope. A table of the values in each variable and the total likelihood is shown in Table 5.5. I did not include event 17904564 in this table because, while event 17904564 appears very signal-like in these variables, according to `icemc` this event is extremely unlikely as a neutrino candidate, both because of the shallow ice it comes from (consistent with $\sim 1\%$ of simulated neutrinos) and its polarization angle, which is very far from vertical (consistent with $\sim 0.1\%$ of simulated neutrinos).

5.3.2 Diffuse Neutrino Flux Limit

This analysis observed one neutrino candidate in the Askaryan channel on a background of $0.64^{+0.69}_{-0.45}$. Event 36019849 was included in the Askaryan neutrino analysis, event 17904564




Figure 5.46: Event 17904564 (the isolated singlet that made it into the Hpol signal region by mistake), dedispersed and coherently summed, with the vertically polarized waveform in red and horizontally polarized waveform in blue. It is a very impulsive event, with nearly equal power in Hpol and Vpol. An interesting event to be sure, but there is not much more we can say about it. The bottom panel shows the power spectrum of the top panel in dB.




Figure 5.47: Event 36019849 with its 5-sigma error ellipse projected onto the continent. The nearest known location of human activity, a fixed-wing landing site on Deverall Island, is also shown.

Observable	Event 36019849 Likelihood Ratio
Impulsivity	0.03
Power Window Gradient	2.60
Linear Polarization	8.40
Map Peak	0.22
Hilbert Peak	1.10
Total Product	0.16

Table 5.5: Likelihood ratio calculations for the Askaryan neutrino candidate. The likelihood ratio is derived from simulated neutrinos and clustered events. Observables are chosen because they should be very strong in true neutrino candidates. The total product is the ratio of the likelihoods of the event to be signal to background. Event 36019849 appears to be about 6 times more likely to be background than signal according to this choice of variables.

was not included because it triggered in Hpol. We calculate a diffuse neutrino flux limit for ANITA using a total livetime of 24.52 days, the geometric mean of the acceptance as computed by `icemc` and an independently developed simulation from University of Hawaii (UHMC), the analysis efficiency as a function of neutrino energy, and the 90% upper limit Feldman-Cousins factor calculated for the number of events detected and expected background [129]. There is a large difference between the acceptance as calculated by `icemc` and UHMC, as large as a factor of ~ 3 at 10^{18} eV and $\sim 20\%$ at 10^{20} eV. Fig. 5.48 shows the aperture or effective area used in calculation of the published limit as a function of energy. Systematic uncertainties are included using the POissonian Limit Estimator Plus Plus(polepp) program [130, 131]. The calculated limit is shown in Fig. 5.49, along with two flux models [116, 132], the ANITA-III limit, the combined limit for ANITA I-IV, and the limits from IceCube and Auger [133, 60]. The ANITA I-IV total combined limit is made using the total number of events seen, the sum of all of the expected backgrounds, and the sum of the effective volumes for each flight weighted by their analysis efficiencies as a function of energy.

We would only expect 0.33 events from a Kotera maximum all-proton model, and 0.06 events from a Kotera maximum mixed-composition model [116]. We also set a 90% CL




Figure 5.48: The effective area as a function of energy used in calculating the ANITA-IV limit. This is calculated using the geometric average of two independent Monte Carlo simulations.

integral flux limit on a pure E_ν^{-2} spectrum between 10^{18} eV and 10^{21} eV of $E_\nu^2 \Phi_\nu \leq 2.2 \times 10^{-7} \text{ GeV cm}^{-2} \text{ s}^{-1} \text{ sr}^{-1}$ [111].

5.4 Summary of Analysis Results

My analysis of the ANITA-IV data detected 28 UHE EAS candidate events in ~ 24.5 days of livetime. 24 of these were reflected cosmic ray events, 2 were direct cosmic ray events, and 2 were unusual inverted polarity EAS signals whose origins are unknown. These two unusual inverted polarity EAS events are very different from previous ANITA “mystery events” in that they come from near the horizon rather than from steep elevation angles. Their total probability of being background is somewhere between 2.3% and 5.5%, depending on assumptions.

In the Askaryan neutrino signal region of my analysis, I found one neutrino candidate on a background of $0.64_{-0.45}^{+0.69}$. Because 1 event on a background that large is not a discovery, all we could do with ANITA-IV was set a limit on the UHE neutrino flux, further improving




Figure 5.49: The all-flavor diffuse UHE neutrino flux limit for ANITA, as well as Auger and IceCube [133, 60], along with predicted fluxes from Kotera and Ahlers models [116, 132]. The combined limit is the most stringent above energies greater than $10^{19.5}$ eV.

constraints on the flux above about $10^{19.5}$ eV.

5.5 Future Outlook

There are many things that a future ANITA-like experiment could improve upon. The first of many proposed upgrades is to the name. Chief among the proposed names is the Payload for Ultra-high Energy Observations (PUEO).

On the calibration front, we could better constrain our off-axis antenna response, cross-polarization leakage and polarization angle resolution. All of these are important for pointing and reconstruction, as well analysis. Data for each of these could be taken in Hawaii or during integration in Palestine. We also should have taken full impulse response data for each TUFF and configuration individually, but we did not. If we fly an experiment with TUFFs or TUFF-like hardware, that should be remedied. An additional calibration problem that is harder to solve is the need for different phase centers for signals of differing frequency content. To solve this we could take a programmable pulser out to WAIS and pulse using several sets of signals at different frequencies. The last obvious calibration problem I found during my analysis was the lack of calibration data taken at different TUFF configurations. All of the recorded WAIS pulser data was taken with the same TUFF configuration. There is no way of knowing if pointing resolution or phase centers change in an unpredictable way with different notch configurations.

On the analysis side of things, the number one thing that ANITA needed for this analysis was a cosmic ray simulation. Without a simulation, we are using an analysis optimized for Askaryan neutrinos, and hoping we don't cut out EAS events. It's very possible that my suite of quality cuts was sub-optimal for EAS events, because many of them hinged on detecting the low frequency dominated spectrum of a payload blast. In order to know this and tweak cuts appropriately, I would have needed simulated EAS events to test these cuts on. The same thing is true of background estimate cuts. I set these without any information about

the distributions of these quantities for actual EAS events, except what I could glean using the precious few from previous ANITA flights. A simulation seems especially pressing now that we've found a handful of mysterious inverted CR-like events.

The biggest problem in analysis that I don't have a solution for is the irreducible background from anthropogenic events. In order to improve the analysis, the number of expected background events would need to be reduced significantly. We can probably do this better on the EAS side of things with a proper simulation to map out the largest possible sideband region, but on the Askaryan neutrino side of things we seem to be somewhat stuck.

There are significant upgrades on the horizon for the hardware of future ANITA. The improvement I am most excited about is an FPGA-based phased array trigger, which has already been tested and shown to work well on ARA [134]. In addition to that there will be more antennas, most likely at with a higher minimum frequency, to save space. These antennas will be read out on higher resolution SURFv5 digitizers, using the new LAB4D chips [135]. All of these upgrades will increase livetime, trigger efficiency and sensitivity for the next generation balloon-borne radio neutrino detector, whatever its name may be.

CHAPTER 6

THERMAL NOISE CORRELATION STUDIES

6.1 Introduction

Another promising avenue for detecting UHE neutrinos is the in-ice experimental technique. Similar to ANITA, experiments in this category look for Askaryan emission from neutrinos interacting in the ice. Different from ANITA, the antennas are placed in boreholes drilled in the ice or near the surface (as in ARIANNA). The volume instrumented is lower, but so are the thresholds and backgrounds, so the energy range is different than ANITA. During the early stages of my time at UChicago, I worked towards developing an in-ice phased array system for the detection of UHE neutrinos [134, 136].

The phased array technique is based on coherently summing incoming waveforms from multiple antennas at the trigger level. In ANITA we do interferometry in analysis, after we've already triggered. Doing interferometry at the trigger level effectively lowers the threshold of the instrument by allowing us to trigger on lower SNR signals. SNR should increase by \sqrt{N} where N is the number of antennas going into the coherent sum, because signal should add coherently and noise should add incoherently. Fig. 6.1 shows the improvement in antenna level SNR achievable with a phased array trigger. This assumes that noise is uncorrelated, even between closely spaced antennas. I worked mainly on proving the viability of this technique by measuring and simulating thermal noise between closely spaced antennas, in order to determine whether our assumptions held.

6.2 Thermal Noise Correlation Measurements

In order to test our assumption that thermal noise should be uncorrelated between antennas, we performed measurements of noise correlation in an anechoic chamber at the Physical




Figure 6.1: Simulated trigger efficiencies for 3, 8 and 16-antenna phased arrays summed into a single beam. Efficiency is drawn for rates of 1000 Hz, 100 Hz, and 10 Hz, denoted by the black-solid, dashed and gray-solid lines, respectively. Figure taken from [136].

Sciences Lab (PSL) in Wisconsin. Fig. 6.2 shows a diagram of the setup inside the anechoic chamber. Antennas were placed in each other’s nulls to imitate the in-ice setup. We took measurements at various antennas spacings to see if thermal noise correlations varied with distance, matching our assumptions.

We took measurements with five different types of antennas, listed with their bandwidths in parentheses, Telewave 275Ds (230-330 MHz), Telewave 400Ds (360-450 MHz), ARA Vpol bicones (150-850 MHz), ARA Hpol slot antennas (200-850 MHz), and broadband dipole antennas developed by Stephanie Wissel (called GNO antennas here). A GNO dipole antenna is shown along with its response in Fig. 6.3. Antenna frequency ranges were chosen by constraints of an in-ice borehole, as well as the peak power of an Askaryan signal. We also took measurements using terminated amplifiers in the anechoic chamber to have an example of completely uncorrelated thermal noise to compare against.

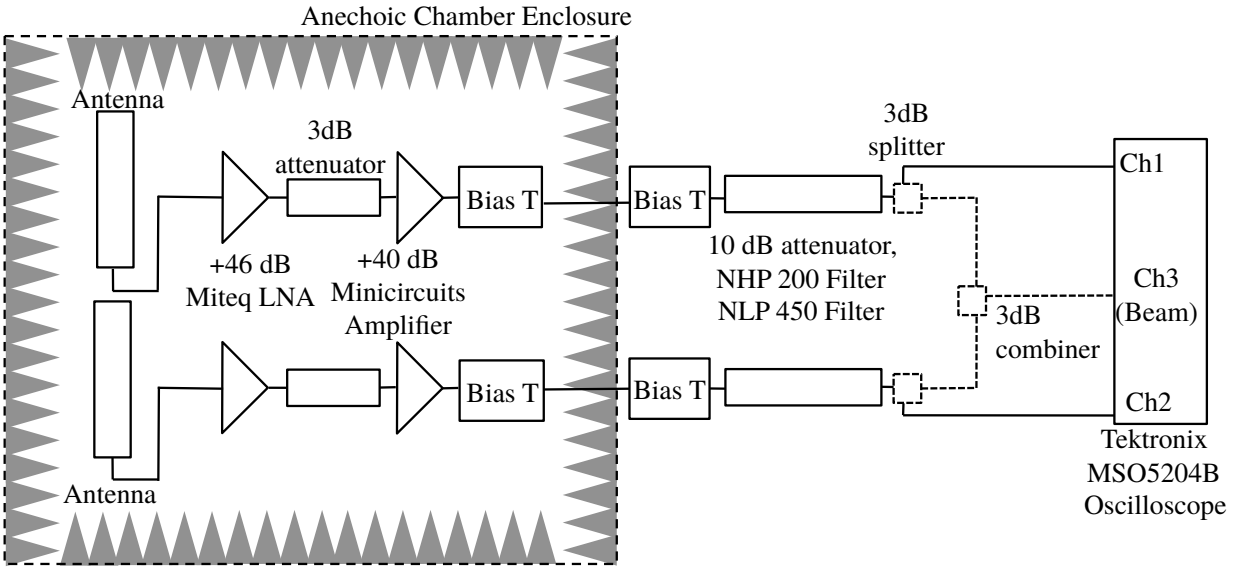


Figure 6.2: Thermal noise correlation measurements setup in the PSL anechoic chamber. Antennas are placed in each other's nulls, and antenna output is passed through amplification and filtering along the way to the Tektronix scope, where it is recorded. Grey, shaded triangles represent the RF attenuating foam that lines the anechoic chamber. Figure taken from [136].




Figure 6.3: A picture of the constructed GNO dipole is pictured on the left. The middle panel shows the measured S_{11} of all constructed antennas (in color) and the HFSS simulated S_{11} in dashed black. On the right panel is the measured radiation pattern of a single antenna (solid color) and the simulated radiation pattern (dashed color). Figure taken from [136].

6.3 Thermal Noise Correlation Simulations

In our thermal noise simulation, we treated noise sources as harmonic oscillators with an amplitude drawn from a Rayleigh distribution corresponding to 300K temperature, and a phase drawn from a uniform distribution from 0 to 2π , following the recipe from Goodman [137]. We create 10^5 oscillators, each with a single frequency drawn from a uniformly distributed band of 0 to 2 GHz. The oscillators randomly populate the walls of the chamber, also drawn from a uniform distribution. Noise from each oscillator travels to the antennas as a plane wave. There are two antennas, each at different places in the chamber, and each receives the plane wave at slightly different times.

When the plane wave radiation reaches the antenna, we apply an approximate antenna beam pattern (the beam pattern of a perfect short dipole), as well as the frequency response of the antenna, either from the antenna spec sheets or recorded data. Then we add 75 K amplifier noise in Fourier space by generating uncorrelated noise with amplitude drawn from a 75 K Rayleigh distribution and phase from a uniform 0 to 2π distribution. This is added to the rest of the noise in Fourier space and filtered through simulated Minicircuits 200 MHz high-pass and 600 MHz low-pass or 200 MHz high-pass and 450 MHz low-pass filters, depending on which test we are simulating. Filtering is done in Fourier space as well, by multiplying the signal after amplification by the frequency response as given on the datasheet for each filter provided by the Minicircuits website. This produces a waveform that, if Fourier transformed into the time domain, provides a very good simulation of the data we recorded on the scope in the anechoic chamber. Using this procedure we generate many noise waveforms that can be analyzed and compared to data.

The frequency response of the antenna depends on the type of antenna used in the simulation. For the Telewave, Inc. antennas, no datasheet measurements of antenna frequency responses could be located so we use a boxcar window on the quoted frequency band (230-330 MHz or 360-450 MHz). We do the same boxcar window for the ARA Hpol and GNO

antennas, setting all out of band frequencies to zero. However, for the ARA Vpol antennas, we have a measured frequency response that is applied along with the beam pattern.

The walls of the chamber are also adjustable, both in shape and size. The simulation default is a box shaped chamber, meant to represent the geometry of the actual anechoic chamber at the PSL in Wisconsin.

6.4 Analysis

The statistic we're most interested in is called the mean peak correlation coefficient. In order to obtain that we take two waveforms (one for each antenna, per event) and find the maximum Pearson correlation coefficient between them. We calculate the Pearson correlation coefficient by doing a correlation between two time domain waveforms and applying a normalization factor. The correlation coefficient is defined as:

$$C(x(t), y(t), \tau) = \frac{1}{N\sigma_x\sigma_y} \sum_{n=0}^N (x(t_i) - \bar{x})(y(t_i + \tau) - \bar{y}) \quad (6.1)$$

Only the values within ± 6.0 ns shift of each other are considered for the maximum correlation coefficient, which is the window of causal time for the maximum feed-to-feed spacing. We do this for all events in a certain configuration of antennas set at some distance, with 500 waveforms for each antenna spacing. Then we take the mean of all of these maximum correlation coefficient values. Each set of spacings and antennas gets this same treatment. This gives us a metric to compare simulation with data.

We also looked at the general distribution of correlation coefficients, as a validation check on the simulation. To check this, we basically ran through the same analysis of Pearson correlation coefficients, but added them all to a histogram instead of taking only the maximum value. Again, this was only done for values within the causal window of the largest feed-to-feed distance. Simulated and recorded events went through the exact same




Figure 6.4: Histograms of the maximum cross-correlation coefficient within a causal window (left) and all cross-correlation coefficients within a causal window (right). Black is simulated uncorrelated noise, blue is uncorrelated noise taken with terminated amplifiers, and red is the data taken using the ARA Vpol bicone. All three of these are consistent with being from the same distribution, showing that noise between closely spaced antennas is uncorrelated. Figure taken from [136].

analysis pipeline and were treated the same after being generated.

6.5 Results and Comparison

When we look at the distribution of correlation coefficients and maximum correlation coefficients of simulation vs data, we can see they are drawn from approximately the same distribution. When we add noise drawn from an entirely uncorrelated data set, we see that matches as well, as in Fig. 6.4. We also saw that there was no mean peak noise correlation above noise seen in the simulation until antennas were closer together than physically allowable. At very close spacings, the correlation coefficient approaches 0.8, which comes about because the fraction of completely uncorrelated noise from the 75 K amplifiers is $\frac{75}{375}$. This is shown in Fig. 6.5.

The results are consistent across all antenna types. There is no significant correlation




Figure 6.5: The mean of the maximum correlation value for 500 trials of simulated data (blue) and data recorded with the ARA bicones (green). As the distance between the antenna feeds decreases, the thermal noise eventually becomes correlated, but the required distance is unachievable. Data was taken down to the smallest physically realizable distance. The simulated correlation coefficient maxes out around 0.8 because roughly 20% of thermal noise in the system is uncorrelated noise introduced by the amplifiers. Figure taken from [136].




Figure 6.6: Diagram representing the setup during the tests with a transmitting antenna. Transmitting antenna and receiving antennas were all the same antenna type. Signal chain is the same as in Fig. 6.2. Figure taken from [136].

from thermal noise observed until antennas are closer than physically possible. This holds for different types of filters too. We do see a small deficit in noise correlations for the data when compared to simulation. We attribute this to a variety of factors not included in the simulation, such as the phase response of the antennas as a function of angle, that should only serve to further decorrelate thermal noise.

To ensure we actually receive the \sqrt{N} boost in SNR when looking at signal, we also included a transmitting antennas in some of the tests, as shown in Fig. 6.6. We sent broadband impulsive signals from the transmitting antenna to the two receiving antennas (all three antennas are the same type). The results of this measurement are shown in Fig. 6.7, where the left panel shows the signal received on each antenna, and the right panel shows the combined signal as formed in hardware and software. They are similar to the level of 15% in terms of peak-to-peak voltage. We see the SNR increase from ~ 2.6 in each antenna to ~ 3.1 when combined in hardware, which isn't quite the \sqrt{N} improvement we would expect. We attribute this discrepancy to additional noise and losses along the signal chain.




Figure 6.7: Voltage vs time traces of each channel on the left and the combined beam on the right. Channel 1 is solid red and channel 2 is dashed blue on the left and the ideal beam as combined in software is solid green, while the beam as combined in hardware is shown in dashed orange. Channel 1 and channel 2 are 2.6 and 2.7 SNR, respectively, while the hardware and software combined beams are 3.1 and 3.3 SNR, respectively. Figure taken from [136].

An in-ice phased array trigger for ARA was deployed for ARA to great success [134]. The prototype phased array trigger showed a factor of ~ 1.8 improvement in SNR over the standard ARA trigger, which corresponds to the same factor of improvement in trigger-level effective volume for neutrinos with energy between 10 and 100 PeV. We plan to continue to develop the phased array technique in the successor to ANITA-IV, as well as the proposed RNO project.

REFERENCES

- [1] W. Pauli. Neutrino letter, 1930. https://cds.cern.ch/record/83282/files/meitner_0393.pdf.
- [2] E. Fermi. Versuch einer Theorie der β -Strahlen. I. *Zeitschrift fur Physik*, 88:161–177, March 1934.
- [3] F. L. Wilson. Fermi’s Theory of Beta Decay. *American Journal of Physics*, 36:1150–1160, December 1968.
- [4] C. L. Cowan, Jr., F. Reines, F. B. Harrison, H. W. Kruse, and A. D. McGuire. Detection of the Free Neutrino: A Confirmation. *Science*, 124:103–104, July 1956.
- [5] B. T. Cleveland, T. Daily, R. Davis, Jr., J. R. Distel, K. Lande, C. K. Lee, P. S. Wildenhain, and J. Ullman. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophys. J.*, 496:505–526, 1998.
- [6] J. N. Bahcall and M. H. Pinsonneault. Solar models with helium and heavy element diffusion. *Rev. Mod. Phys.*, 67:781–808, 1995.
- [7] S. Turck-Chieze and I. Lopes. Toward a unified classical model of the sun: On the sensitivity of neutrinos and helioseismology to the microscopic physics. *Astrophys. J.*, 408:347–367, 1993.
- [8] I. J. Sackmann, A. I. Boothroyd, and W. A. Fowler. Our Sun. I - The Standard Model: Successes and Failures. *Astrophys. J.*, 360:727, 1990.
- [9] K. S. Hirata et al. Results from one thousand days of real-time, directional solar-neutrino data. *Phys. Rev. Lett.*, 65:1297–1300, September 1990.
- [10] Y. Fukuda et al. Solar Neutrino Data Covering Solar Cycle 22. *Phys. Rev. Lett.*, 77:1683–1686, August 1996.

- [11] Y. Fukuda et al. Measurements of the Solar Neutrino Flux from Super-Kamiokande's First 300 Days. *Phys. Rev. Lett.*, 81:1158–1162, August 1998.
- [12] V. N. Gavrin. Measurement of the solar neutrino capture rate in SAGE and the value of the pp-neutrino flux at the earth. *Nucl. Phys. Proc. Suppl.*, 138:87–90, 2005. [,87(2005)].
- [13] W. Hampel et al. GALLEX solar neutrino observations: Results for GALLEX IV. *Phys. Lett.*, B447:127–133, 1999.
- [14] Y. Fukuda et al. Evidence for Oscillation of Atmospheric Neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, August 1998.
- [15] S. N. Ahmed et al. Measurement of the Total Active ^8B Solar Neutrino Flux at the Sudbury Neutrino Observatory with Enhanced Neutral Current Sensitivity. *Phys. Rev. Lett.*, 92:181301, May 2004.
- [16] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Sov. Phys. JETP*, 7:172–173, 1958. [Zh. Eksp. Teor. Fiz.34,247(1957)].
- [17] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the Unified Model of Elementary Particles. *Progress of Theoretical Physics*, 28(5):870–880, November 1962.
- [18] K. Abe et al. The T2K Experiment. *Nucl. Instrum. Meth.*, A659:106–135, 2011.
- [19] X. Guo et al. A Precision measurement of the neutrino mixing angle θ_{13} using reactor antineutrinos at Daya-Bay. 2007.
- [20] M. G. Aartsen et al. First observation of PeV-energy neutrinos with IceCube. *Phys. Rev. Lett.*, 111:021103, 2013.

- [21] M. G. Aartsen et al. Observation and Characterization of a Cosmic Muon Neutrino Flux from the Northern Hemisphere using six years of IceCube data. *Astrophys. J.*, 833:3, 2016.
- [22] K. Hirata et al. Observation of a neutrino burst from the supernova SN1987A. *Phys. Rev. Lett.*, 58:1490–1493, April 1987.
- [23] R. M. Bionta et al. Observation of a neutrino burst in coincidence with supernova 1987A in the Large Magellanic Cloud. *Phys. Rev. Lett.*, 58:1494–1496, April 1987.
- [24] M. G. Aartsen et al. Neutrino emission from the direction of the blazar TXS 0506+056 prior to the IceCube-170922A alert. *Science*, 361(6398):147–151, 2018.
- [25] M. G. Aartsen et al. Multimessenger observations of a flaring blazar coincident with high-energy neutrino IceCube-170922A. *Science*, 361(6398):eaat1378, 2018.
- [26] M. Branchesi. Multi-messenger astronomy: gravitational waves, neutrinos, photons, and cosmic rays. *Journal of Physics: Conference Series*, 718:022004, May 2016.
- [27] S. Ando et al. Colloquium: Multimessenger astronomy with gravitational waves and high-energy neutrinos. *Rev. Mod. Phys.*, 85:1401–1420, October 2013.
- [28] M. Ackermann et al. Fundamental Physics with High-Energy Cosmic Neutrinos. 2019.
- [29] Y. Xu. TeV-PeV neutrino-nucleon cross section measurement with 5 years of IceCube data. *PoS*, DIS2018:019, 2018.
- [30] J. Alvarez-Muniz, F. Halzen, T. Han, and D. Hooper. Phenomenology of high-energy neutrinos in low scale quantum gravity models. *Phys. Rev. Lett.*, 88:021301, 2002.
- [31] M. Bustamante, J. F. Beacom, and W. Winter. Theoretically palatable flavor combinations of astrophysical neutrinos. *Phys. Rev. Lett.*, 115(16):161302, 2015.

- [32] T. Wulf. Beobachtungen über Strahlung hoher Durchdringungsfähigkeit auf dem Eifelturm. *Phys. Z.*, 11:811–813, 1909.
- [33] V. F. Hess. Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten. *Phys. Z.*, 13:1084–1091, 1912.
- [34] J. Bellido. Depth of maximum of air-shower profiles at the Pierre Auger Observatory: Measurements above $10^{17.2}$ eV and Composition Implications. *PoS, ICRC2017:506*, 2018. [,40(2017)].
- [35] J. Abraham et al. Measurement of the Depth of Maximum of Extensive Air Showers above 10^{18} eV. *Phys. Rev. Lett.*, 104(9):091101, March 2010.
- [36] W. Hanlon, D. Ikeda, T. Stroman, J. P. Lundquist, Y. Zhezher, and Telescope Array Collaboration. Telescope Array Composition Summary. *International Cosmic Ray Conference*, 301:536, January 2017.
- [37] *Pierre Auger Observatory and Telescope Array: Joint Contributions to the 35th International Cosmic Ray Conference (ICRC 2017)*, 2018.
- [38] A. Aab et al. Combined fit of spectrum and composition data as measured by the Pierre Auger Observatory. *JCAP*, 1704(04):038, 2017. [Erratum: *JCAP*1803,no.03,E02(2018)].
- [39] M. Tanabashi et al. Review of particle physics. *Phys. Rev. D*, 98:030001, August 2018.
- [40] L O’C. Drury. An introduction to the theory of diffusive shock acceleration of energetic particles in tenuous plasmas. *Reports on Progress in Physics*, 46(8):973–1027, August 1983.
- [41] E. Fermi. On the Origin of the Cosmic Radiation. *Phys. Rev.*, 75:1169–1174, April 1949.

- [42] L O’C. Drury. Origin of cosmic rays. *Astroparticle Physics*, 39-40:52 – 60, 2012. Cosmic Rays Topical Issue.
- [43] A. M. Hillas. The Origin of Ultra-High-Energy Cosmic Rays. *Annual Review of Astronomy and Astrophysics*, 22(1):425–444, 1984.
- [44] J. Heinze, D. Boncioli, M. Bustamante, and W. Winter. Cosmogenic neutrinos challenge the cosmic-ray proton dip model. *The Astrophysical Journal*, 825(2):122, July 2016.
- [45] K. Kotera and A. V. Olinto. The Astrophysics of Ultrahigh-Energy Cosmic Rays. *Annual Review of Astronomy and Astrophysics*, 49(1):119–153, 2011.
- [46] K. Greisen. End to the Cosmic-Ray Spectrum? *Phys. Rev. Lett.*, 16:748, 1966.
- [47] G. T. Zatsepin and V. A. Kuzmin. Upper limit of the spectrum of cosmic rays. *JETP Lett.*, 4:78–80, 1966.
- [48] V. S. Beresinsky and G. T. Zatsepin. Cosmic rays at ultra high energies (neutrino?). *Phys. Lett. B*, 28:423–424, 1969.
- [49] A. M. Taylor, M. Ahlers, and D. Hooper. Indications of Negative Evolution for the Sources of the Highest Energy Cosmic Rays. *Phys. Rev.*, D92(6):063011, 2015.
- [50] K. Murase. Active Galactic Nuclei as High-Energy Neutrino Sources. *Neutrino Astronomy - Current status, future prospects*, Eds. T. Gaisser & A. Karle (World Scientific), pages 15–31, 2017.
- [51] K. Murase, Y. Inoue, and C. D. Dermer. Diffuse Neutrino Intensity from the Inner Jets of Active Galactic Nuclei: Impacts of External Photon Fields and the Blazar Sequence. *Phys. Rev.*, D90(2):023007, 2014.

- [52] K. Fang, K. Kotera, K. Murase, and A. V. Olinto. Testing the Newborn Pulsar Origin of Ultrahigh Energy Cosmic Rays with EeV Neutrinos. *Phys. Rev.*, D90(10):103005, 2014.
- [53] D. Boncioli, D. Biehl, and W. Winter. On the common origin of cosmic rays across the ankle and diffuse neutrinos at the highest energies from low-luminosity Gamma-Ray Bursts. *ArXiv e-prints*, page arXiv:1808.07481, August 2018.
- [54] D. Biehl, D. Boncioli, C. Lunardini, and W. Winter. Tidally disrupted stars as a possible origin of both cosmic rays and neutrinos at the highest energies. *Sci. Rep.*, 8(1):10828, 2018.
- [55] K. Fang and B. D. Metzger. High-Energy Neutrinos from Millisecond Magnetars formed from the Merger of Binary Neutron Stars. *Astrophys. J.*, 849(2):153, 2017.
- [56] M. Kachelriess. The rise and fall of top-down models as main UHECR sources. In *Proceedings, 20th Rencontres de Blois on Challenges in Particle Astrophysics: Blois, France, May 18-23, 2008*, pages 215–224, 2008.
- [57] R. Aloisio, S. Matarrese, and A. V. Olinto. Super Heavy Dark Matter in light of BICEP2, Planck and Ultra High Energy Cosmic Rays Observations. *JCAP*, 1508(08):024, 2015.
- [58] C. Haack and C. Wiebusch. A measurement of the diffuse astrophysical muon neutrino flux using eight years of IceCube data. *PoS, ICRC2017:1005*, 2018.
- [59] C. Kopper. Observation of Astrophysical Neutrinos in Six Years of IceCube Data. *PoS, ICRC2017:981*, 2018.
- [60] M. G. Aartsen et al. Differential limit on the extremely-high-energy cosmic neutrino flux in the presence of astrophysical background from nine years of IceCube data. *Phys. Rev.*, D98(6):062003, 2018.

- [61] A. van Vliet, R. Alves Batista, and J. R. Hörandel. Determining the fraction of cosmic-ray protons at ultra-high energies with cosmogenic neutrinos. 2019.
- [62] M. Ackermann et al. Astrophysics Uniquely Enabled by Observations of High-Energy Cosmic Neutrinos. 2019.
- [63] G. Askaryan. Excess negative charge of an Electron-photon shower and its coherent radio emission. *Sov. Phys. JETP*, 14 (2):441–443, 1962.
- [64] G. A. Askar’yan. Coherent Radio Emission from Cosmic Showers in Air and in Dense Media. *Soviet Journal of Experimental and Theoretical Physics*, 21:658, September 1965.
- [65] P. W. Gorham et al. Observations of the Askaryan effect in ice. *Phys. Rev. Lett.*, 99(17):171101, 2007.
- [66] D. Saltzberg, P. W. Gorham, D. Walz, C. Field, R. Iverson, A. Odian, G. Resch, P. Schoessow, and D. Williams. Observation of the Askaryan Effect: Coherent Microwave Cherenkov Emission from Charge Asymmetry in High-Energy Particle Cascades. *Physical Review Letters*, 86:2802–2805, March 2001.
- [67] P. W. Gorham, D. Saltzberg, R. C. Field, E. Guillian, R. Milincic, D. Walz, and D. Williams. Accelerator measurements of the Askaryan effect in rock salt: A Roadmap toward teraton underground neutrino detectors. *Phys. Rev.*, D72:023002, 2005.
- [68] P. Auger, P. Ehrenfest, R. Maze, J. Daudin, and R. A. Fréon. Extensive Cosmic-Ray Showers. *Rev. Mod. Phys.*, 11:288–291, July 1939.
- [69] J. V. Jelley, J. H. Fruin, N. A. Porter, T. C. Weekes, F. G. Smith, and R. A. Porter. Radio Pulses from Extensive Cosmic-Ray Air Showers. *Nature*, 205:327–328, January 1965.

- [70] N. A. Porter, C. D. Long, B. McBreen, D. J. B. Murnaghan, and T. C. Weekes. Detection of radio pulses from extensive air showers, with a wide-band receiving system. *Physics Letters*, 19:415–417, November 1965.
- [71] I. A. Borzhkovskii, V. D. Volovik, V. I. Kobizskoy, and E. S. Shmatko. Measurement of the Polarization of Coherent Radio Emission of Extensive Air Showers (EAS). *Soviet Journal of Experimental and Theoretical Physics Letters*, 3:118, February 1966.
- [72] F. D. Kahn, I. Lerche, and A. C. B. Lovell. Radiation from cosmic ray air showers. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 289(1417):206–213, 1966.
- [73] A. Aab et al. Probing the radio emission from air showers with polarization measurements. *Phys. Rev.*, D89(5):052002, 2014.
- [74] K. Belov et al. Accelerator measurements of magnetically-induced radio emission from particle cascades with applications to cosmic-ray air showers. *Phys. Rev. Lett.*, 116(14):141103, 2016.
- [75] J. Alvarez-Muiz, W. R. Carvalho, D. Garca-Fernndez, H. Schoorlemmer, and E. Zas. Simulations of reflected radio signals from cosmic ray induced air showers. *Astropart. Phys.*, 66:31–38, 2015.
- [76] J. L. Feng, P. Fisher, F. Wilczek, and T. M. Yu. Observability of Earth-Skimming Ultrahigh Energy Neutrinos. *Phys. Rev. Lett.*, 88:161102, April 2002.
- [77] P. W. Gorham et al. Characteristics of Four Upward-Pointing Cosmic-Ray-like Events Observed with ANITA. *Phys. Rev. Lett.*, 117:071101, 2016.
- [78] P. W. Gorham et al. Observation of an Unusual Upward-going Cosmic-ray-like Event in the Third Flight of ANITA. *Phys. Rev. Lett.*, 121(16):161102, 2018.

- [79] A. Romero-Wolf et al. A comprehensive analysis of anomalous ANITA events disfavors a diffuse tau-neutrino flux origin. 2018.
- [80] B. Rotter. PhD thesis, University of Hawai'i, 2017.
- [81] S. W. Barwick, D. Z. Besson, P. W. Gorham, and D. Saltzberg. South Polar in situ radio-frequency ice attenuation. *J. Glaciol.*, 51:231–238, 2005.
- [82] T. Barrella, S. Barwick, and D. Saltzberg. Ross Ice Shelf in situ radio-frequency ice attenuation. *J. Glaciol.*, 57:61–66, 2011.
- [83] D. Z. Besson et al. In situ radioglaciological measurements near Taylor Dome, Antarctica and implications for UHE neutrino astronomy. *Astropart. Phys.*, 29:130–157, 2008.
- [84] P. W. Gorham et al. New Limits on the ultra-high energy cosmic neutrino flux from the ANITA experiment. *Phys. Rev. Lett.*, 103:051103, 2009.
- [85] S. Hoover et al. Observation of ultra-high-energy cosmic rays with the ANITA balloon-borne radio interferometer. *Phys. Rev. Lett.*, 105:151101, 2010.
- [86] P. W. Gorham et al. Observational constraints on the ultra-high energy cosmic neutrino flux from the second flight of the ANITA experiment. *Phys. Rev. D*, 82:022004, 2010. [Erratum: *Phys. Rev. D* 85, 049901 (2012)].
- [87] A. G. Vieregg et al. The First Limits on the Ultra-high Energy Neutrino Fluence from Gamma-Ray Bursts. *The Astrophysical Journal*, 736(1):50, July 2011.
- [88] P. W. Gorham et al. Constraints on the diffuse high-energy neutrino flux from the third flight of ANITA. *Phys. Rev. D*, 98:022001, 2018.
- [89] S. R. Klein and A. Connolly. Neutrino Absorption in the Earth, Neutrino Cross-Sections, and New Physics. In *Proceedings, 2013 Community Summer Study on the*

Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013, 2013.

- [90] K. D. de Vries and S. Prohira. Coherent transition radiation from the geomagnetically-induced current in cosmic-ray air showers: Implications for the anomalous events observed by ANITA. 2019.
- [91] P. Allison et al. Dynamic tunable notch filters for the Antarctic Impulsive Transient Antenna (ANITA). *Nucl. Instrum. Methods A*, 894:47–56, June 2018.
- [92] P. Allison et al. Design and Initial Performance of the Askaryan Radio Array Prototype EeV Neutrino Detector at the South Pole. *Astropart. Phys.*, 35:457–477, 2012.
- [93] L. Gerhardt, S. Klein, T. Stezelberger, S.W. Barwick, K. Dookayka, J. Hanson, and R. Nichol. A prototype station for ARIANNA: a detector for cosmic neutrinos. *Nucl. Instrum. Meth.*, A624:85–91, 2010.
- [94] O. Martineau-Huynh et al. The Giant Radio Array for Neutrino Detection. *EPJ Web Conf.*, 116:03005, 2016.
- [95] S. Wissel et al. A New Concept for High-Elevation Radio Detection of Tau Neutrinos. *accepted EPJ Web Conf.*, 2018. Presented at the Acoustic and Radio EeV Neutrino Detection Activities Conference 2018 (ARENA 2018), Catania, Italy, June 12–15, 2018 <https://indico.cern.ch/event/667036/contributions/3005761/>.
- [96] J. Nam. Taiwan Astroparticle Radiowave Observatory for Geo-synchrotron Emissions (TAROG). *PoS, ICRC2015:663*, 2016.
- [97] Angela V. Olinto et al. POEMMA: Probe Of Extreme Multi-Messenger Astrophysics. *PoS, ICRC2017:542*, 2018. [35,542(2017)].

- [98] O. Banerjee, J. Russell, and P. Gorham. AMPA update: testing gain with network analyzer, 2016. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/674.
- [99] L. Batten, L. Cremonesi, C. Hast, J. Russell, and N. Wang. AMPAs retested at LDB for gain and noise figure, 2016. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/681.
- [100] G. S. Varner et al. The large analog bandwidth recorder and digitizer with ordered readout (LABRADOR) ASIC. *Nucl. Instrum. Methods A*, 583:447–460, 2007.
- [101] J. Russell. Attitude determination using magnetometer or sun sensors, 2016. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/684.
- [102] B. Strutt. PhD thesis, University College London, 2016.
- [103] P. Cao. PhD thesis, University of Delaware, 2018.
- [104] A. Romero-Wolf. S parameter magnitude and group delay for 16-way splitter+cables, for MCM calib. data, 2016. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/687.
- [105] P. Allison and L. Batten. S parameter magnitude and group delay data for heliax cables (+ other setups) for McM calib. data, 2016. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/688.
- [106] B. Strutt, N. Wang, D. Saltzberg, and S. Wissel. ANITA-4 WAIS Divide pulsing log, 2017. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/699.
- [107] P. W. Gorham et al. Antarctic Surface Reflectivity Measurements from the ANITA-3 and HiCal-1 Experiments. *J. Astron. Inst.*, 6:1740002–213, 2017.

- [108] S. Prohira et al. HiCal 2: An instrument designed for calibration of the ANITA experiment and for Antarctic surface reflectivity measurements. *Nucl. Instrum. Methods A*, 918:60 – 66, 2019.
- [109] K. Belov. ANITA-IV geometry, 2017. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/702.
- [110] C. Deaconu, L. Cremonesi, P. Allison, and E. Oberla. Efficiency scans, 2016. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/691.
- [111] P. W. Gorham et al. Constraints on the ultra-high energy cosmic neutrino flux from the fourth flight of ANITA. 2019.
- [112] R. Brun and F. Rademakers. ROOT – An object oriented data analysis framework. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 389(1):81 – 86, 1997. New Computing Techniques in Physics Research V.
- [113] M. Frigo. A Fast Fourier Transform Compiler. In *Proceedings of the 1999 ACM SIGPLAN Conference on Programming Language Design and Implementation*, Atlanta, Georgia, 1999.
- [114] L. Cremonesi et al. icemc: A Monte Carlo Simulation for Cosmogenic Neutrinos interacting in the Antarctic ice as viewed by the Antarctic Impulsive Transient Antenna (ANITA). *arXiv e-prints*, page arXiv:1903.11043, March 2019.
- [115] S. Prohira et al. Antarctic surface reflectivity calculations and measurements from the ANITA-4 and HiCal-2 experiments. *Phys. Rev. D*, 98:042004, August 2018.
- [116] K. Kotera, D. Allard, and A. V. Olinto. Cosmogenic neutrinos: parameter space and detectability from PeV to ZeV. *J. Cosmol. Astropart. Phys.*, 10:13, 2010.

- [117] J. R. Klein and A. Roodman. Blind analysis in nuclear and particle physics. *Annual Review of Nuclear and Particle Science*, 55(1):141–163, 2005.
- [118] J. P. A. Ioannidis. Why most published research findings are false. *PLOS Medicine*, 2(8), August 2005.
- [119] P. Gorham. Filter CW via Sine Wave Fitting, 2015. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/621.
- [120] H. Akima. A New Method of Interpolation and Smooth Curve Fitting Based on Local Procedures. *J. ACM*, 17(4):589–602, October 1970.
- [121] A. Romero-Wolf, S. Hoover, A. Viereg, P. Gorham, and the ANITA Collaboration. An Interferometric Analysis Method for Radio Impulses from Ultra-high Energy Particle Showers. *arXiv e-prints*, page arXiv:1304.5663, April 2013.
- [122] H. Liu, K. C. Jezek, B. Li, and Z. Zhao. Radarsat Antarctic Mapping Project Digital Elevation Model Version 2 [1 km DEM]. *Radarsat Antarctic Mapping Project digital elevation model version 2, Boulder, Colorado USA: National Snow and Ice Data Center. Digital media.*, January 2001.
- [123] R. A. Fisher. The use of multiple measurements in taxonomic problems. *Ann. Eugen.*, 7(2):179–188, 1936.
- [124] A. Hoecker et al. TMVA - Toolkit for Multivariate Data Analysis. *PoS, ACAT2007:040*, 2007.
- [125] A. Romero-Wolf. PhD thesis, University of Hawai'i, 2010.
- [126] Jaime Alvarez-Muñiz, Washington R. Carvalho, and Enrique Zas. Monte Carlo simulations of radio pulses in atmospheric showers using ZHAireS. *Astroparticle Physics*, 35(6):325–341, January 2012.

- [127] K. Belov and D. Saltzberg. Skidoo and irridium phone impulsive waveforms, 2008. ANITA Internal Note https://elog.phys.hawaii.edu/elog/anita_notes/444.
- [128] T. P. Li and Y. Q. Ma. Analysis methods for results in gamma-ray astronomy. *Astrophys. J.*, 272:317–324, 1983.
- [129] G. J. Feldman and R. D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev.*, D57:3873–3889, 1998.
- [130] J. Conrad, O. Botner, A. Hallgren, and C. Pérez de los Heros. Including systematic uncertainties in confidence interval construction for Poisson statistics. *Phys. Rev. D*, 67:012002, January 2003.
- [131] J. Conrad. A program for confidence interval calculations for a Poisson process with background including systematic uncertainties: POLE 1.0. *Computer Physics Communications*, 158:117–123, April 2004.
- [132] M. Ahlers and F. Halzen. Minimal cosmogenic neutrinos. *Phys. Rev. D*, 86(8):083010, 2012.
- [133] E. Zas. Searches for neutrino fluxes in the EeV regime with the Pierre Auger Observatory. *PoS*, ICRC2017:972, 2018.
- [134] R. Bard et al. Design and Performance of an Interferometric Trigger Array for Radio Detection of High-Energy Neutrinos. 2018.
- [135] J. M. Roberts, E. Oberla, P. Allison, G. S. Varner, S. Spack, B. Fox, and B. Rotter. LAB4D: A Low Power, Multi-GSa/s, Transient Digitizer with Sampling Timebase Trimming Capabilities. *Nucl. Instrum. Meth.*, A925:92–100, 2019.
- [136] J. Avva et al. Development toward a ground-based interferometric phased array for

radio detection of high energy neutrinos. *Nuclear Instruments and Methods in Physics Research A*, 869:46–55, October 2017.

[137] J. W. Goodman. *Statistical Optics*. 1985.