

Analysis of the second flight of ANtarctic Impulsive Transient
Antenna with a focus on filtering techniques

DISSERTATION

Presented in Partial Fulfillment of the Requirements for
the Degree Doctor of Philosophy in the
Graduate School of The Ohio State University

By

Brian T. Dailey,

Graduate Program in Physics

The Ohio State University

2017

Dissertation Committee:

Prof. Amy L. Connolly, Adviser

Prof. John F. Beacom

Prof. Klaus Honscheid

Prof. Andrew Heckler

© Copyright by
Brian T. Dailey
2017

Abstract

The observed cutoff in the cosmic ray spectrum leads to a highly motivated expectation of an ultra-high energy (UHE) neutrino flux, coming from interactions between the cosmic rays and cosmic microwave background photons. Although no UHE neutrinos have yet been detected; better background separation and removal will help accelerate the search.

Past flights of the ANtarctic Impulsive Transient Antenna (ANITA) experiment have set the strongest limits on the UHE neutrino flux above 10^{19} eV. Due to the advanced sensitivity of future flights to both signal and anthropogenic backgrounds, the techniques used in the past analyses may not be sufficient to remove backgrounds.

Here, we discuss processes developed for this analysis. First, we discuss newly techniques to filter event waveforms in both the amplitude and phase spectra. These new techniques were applied to the ANITA-2 experiment data set. We discuss a new technique developed that uses equal area bins of ice on the Antarctic continent. Further, we define a set of analysis cuts, how the analysis cuts were optimized for maximum sensitivity for UHE neutrinos, how the number of background and neutrino events were estimated. For our search, we used the maximal Kotera *et. al.* 2010 flux model [1] and optimized based on this model. After optimization, we found zero events from the 10% sample passing all cuts.

These techniques will prove useful for future flights of ANITA as the sensitivity of the instrument increases. The optimization procedure can also provide a starting point for future analysis. The filtering technique shown here decreased misreconstruction in pointing of events. The HealPix method, while requiring further refinement, shows promise by retaining valuable areas of ice that may have been removed from previous analyses.

To my wife, Bridget, and my kids, David and Samantha. Without you, I would have
finished years ago.

ACKNOWLEDGMENTS

I would like to thank my advisor Amy Connolly for helping and supporting me during my Ph.D. Her support and guidance helped me through many rough patches in both research and personal life. It was an honor and a pleasure to work in her group and on the ANtarctic Impulsive Transient Antenna (ANITA) experiment.

I would also like to thank members of my group: Carl Pfindner, Patrick Allison, Jordan Hanson, Sam Stafford, Oindree Banerjee, Brian Clark, and Kaeli Hughes. Each supported me in their own way whether helping with technical problems, sounding out ideas, helping create plots, and reading my dissertation. Thank you for everything.

I thank my committee members: John Beacom, James Beatty, Klaus Honscheid and Andrew Heckler for filling in at the last moment.

I appreciate all the feedback from the members of the ANITA Collaboration for the advice and feedback on my research.

Thank you to Jeffrey and Jessica Welsh. Whenever I needed a break, you were always there to help get my mind off things and help me relax.

Finally I would like to thank my family. Without the love and support of my wife and children I would not have been able to finish this Ph.D. To my parents, siblings and in-laws: thank you for being there when times turned bad and we needed an extra hand.

VITA

2005High School Diploma, Mt. Gilead High School, Mt. Gilead, OH
2009B.S. Physics, Kent State University, Kent, OH
2011M.S. Physics, The Ohio State University, Columbus, OH

PUBLICATIONS

Characteristics of Four Upward-pointing Cosmic-ray-like Events Observed with ANITA,
P.W. Gorham *et. al.*, Phys. Rev. Lett., 117 (2016) [2]

FIELDS OF STUDY

Major Field: Physics

Studies in Experimental Astroparticle Physics: Ultra-High Energy Neutrinos:
Professor Amy Connolly

Contents

	Page
Abstract	ii
Dedication	iv
Acknowledgments	v
Vita	vi
List of Tables	xii
List of Figures	xiii

Chapters:

1. Introduction	1
1.1 Neutrino Astronomy	1
1.1.1 Solar Neutrinos	2
1.1.2 Supernova 1987A	2

1.1.3	IceCube	3
1.2	UHE Astronomy	5
1.2.1	Cosmic Rays	5
1.2.2	Neutrinos	7
1.3	Production of UHE Neutrinos	8
1.3.1	Direct Production	8
1.3.2	GZK Neutrinos	9
1.4	Detection of UHE particles	10
1.4.1	Optical Cherenkov	10
1.4.2	Askaryan Effect	11
1.5	Geo-Magnetic Radiation from Cosmic Rays	14
2.	The ANtarctic Impulsive Transient Antenna	15
2.1	ANITA-1	16
2.2	ANITA-2	20
2.2.1	Gondola	20
2.2.2	Power System	22
2.2.3	Position and Orientation	23
2.2.4	Radio Frequency Signal Chain	23
2.2.5	Front End Signal Processing	28
2.2.6	Triggering	30
2.2.7	Level 3 Trigger	33
2.2.8	Phi Masking	33
2.2.9	Trigger Efficiency	34

2.3	Flight Software and Data Storage	35
2.4	Flight Performance	37
2.4.1	Ground Calibration Pulser	37
2.4.2	Performance of ANITA-2	39
3.	Techniques for Filtering of Background Contamination	42
3.1	Continuous Wave Contamination	42
3.2	Adjusting the Power Spectrum for Notch Filters	43
3.2.1	No Fill	49
3.3	Filling in Notched Bins	51
3.3.1	Rayleigh	51
3.3.2	Interpolated Filter	54
3.3.3	Wiener (Optimal) Filter	56
3.3.4	Comparison of techniques	60
3.4	Phase Corrections	67
3.4.1	Understanding Phase	67
3.4.2	Shifting Methods	76
3.4.3	Interpolated Phase	81
3.4.4	Geometric Method	88
3.4.5	Phase Filtering Results	93
3.5	SineSubtraction	98
3.6	Filtering Results	99
3.6.1	Simulated Events	104

4.	Re-Analysis of ANITA-2 DATA	110
4.1	Blinding	110
4.2	Analysis Strategy	111
4.3	Analysis Techniques	112
4.3.1	Filtering	112
4.3.2	Radio Interferometry	112
4.3.3	HealPix	119
4.4	Training Samples	124
4.4.1	Taylor Dome sample	124
4.4.2	Thermal Sample	124
4.4.3	10% sample	125
4.4.4	Simulated sample	125
4.5	Quality Events	131
4.5.1	Non-RF events	132
4.5.2	Quality Cuts	133
4.5.3	Background Cuts	135
5.	Optimization Procedure and Results	143
5.1	Optimization Procedure	143
5.2	Final Cuts	151
6.	Conclusions	159

Appendices:

A. Calculation of Error Ellipse	161
B. Plots for HealPix Bins	166
Bibliography	182

List of Tables

Table	Page
3.1 Number of over and underflows for θ and ϕ reconstruction.	62
3.2 Number of over and underflows for θ and ϕ reconstruction based on phase filtering algorithms	93
4.1 Number of events cut by each cut if cuts were applied sequentially and if that cut were applied last.	133
4.2 Number of events cut by each cut if cuts were applied sequentially and if that cut were applied last.	134
5.1 HealPix Bins included in Analysis	156
5.2 Number of events cut by final cuts	157

List of Figures

Figure	Page
1.1 A depiction of IceCube	4
1.2 Reconstruction direction of IceCube's extra-solar neutrinos	4
1.3 Cosmic ray spectrum	6
1.4 Cartoon of GZK Horizon	8
1.5 Depiction of patterns from neutrino interactions	11
1.6 Observation of Askaryan effect in ice	13
1.7 Waveform from SLAC test	14
2.1 Cartoon of ANITA	16

2.2	ANITA-1 Flux Limit	17
2.3	ANITA-1 Cosmic Ray Waveforms	18
2.4	ANITA-1 Cosmic Ray polarization vs Earth's Magnetic Field Polarization	19
2.5	Picture of ANITA-2	21
2.6	ANITA-2 RF Signal Chain	24
2.7	ANITA Antenna Transmission Coefficients	25
2.8	ANITA Antenna Beam Pattern	26
2.9	ANITA Impulse Response	27
2.10	ANITA-2 Impulse Response excluding Antenna	27
2.11	Gain and Noise Temperature from RFCMs	29
2.12	Banding of ANITA Trigger	31

2.13	L3 Trigger Efficiency	34
2.14	L3 Trigger Efficiency with different rings of antennas turned off	35
2.15	ANITA-2 Flight Software	36
2.16	ANITA-2 Flight Path	38
2.17	ANITA-2 Dead Time Fraction	40
3.1	Impulse Response for ANITA-2	45
3.2	Difference in thermal baseline placement due to iterative technique	48
3.3	Taylor Dome event with no filtering	50
3.4	Coherently Summed waveform for Taylor Dome Event with no filtering	51
3.5	Single Antenna waveform of Taylor Dome event with no-fill method	52
3.6	Coherently Summed waveform for Taylor Dome Event with no-fill method	52
3.7	Taylor Dome event with Rayleigh method	55

3.8	Antenna Waveform for Taylor Dome Event with Rayleigh Method . . .	56
3.9	Taylor Dome event with Interpolated method	57
3.10	Coherently Summed waveform for Taylor Dome Event with Interpolation Method	58
3.11	Taylor Dome event with Wiener method	61
3.12	Coherently Summed waveform for Taylor Dome Event with Wiener filter	62
3.13	Efficiency of amplitude filters	63
3.14	Taylor Dome Phi Reconstruction	64
3.15	Taylor Dome Theta Reconstruction	65
3.16	SNR distribution after amplitude filtering for Taylor Dome sample . . .	66
3.17	Boxcar Windowed Sine Wave of 400MHz frequency	69
3.18	Amplitude and Phase Spectra of ANITA-2 Impulse Response	71
3.19	Boxcar Windowed Sine Wave plus Impulse	72

3.20	Time Domain, Amplitude and Phase Spectra of a waveform	73
3.21	Boxcar Windowed Sine Wave plus thermal noise	74
3.22	Boxcar Windowed function of two sine waves added together	75
3.23	Simple Mean Shift Regions	78
3.24	Taylor Dome event with simple mean shift method	79
3.25	Coherently Summed waveform for Taylor Dome Event with Simple Mean Shift Method	80
3.26	Complex Mean Shift Regions	82
3.27	Taylor Dome event with complex mean shift method	83
3.28	Coherently Summed waveform for Taylor Dome Event with Complex Mean Shift Method	84
3.29	Time Shift of Impulse Response Waveform	85
3.30	Time Shift of Simulated Waveform	86

3.31 Results of Interpolated Phase	87
3.32 Coherently Summed waveform for Taylor Dome Event with interpolated phase method	88
3.33 Results of Thermal Event with Interpolated Phase	89
3.34 Geometric Method Phasor Diagram	90
3.35 Geometric Method geometry diagram	91
3.36 Taylor Dome event with Geom Method	94
3.37 Coherently Summed waveform for Taylor Dome Event with geometric method	95
3.38 Taylor Dome Phi Reconstruction after phase filtering	95
3.39 Taylor Dome Theta Reconstruction after phase filtering	96
3.40 SNR distribution after phase adjustments for Taylor Dome sample	97
3.41 An example of fitting for sine subtraction method	100

3.42 Taylor Dome event with sine subtraction method	101
3.43 Coherently Summed waveform for Taylor Dome Event with sine subtraction method	102
3.44 Taylor Dome Reconstruction in Phi after Filtering	103
3.45 Taylor Dome Reconstruction in Theta after Filtering	103
3.46 RMS error in Theta and Phi angle	104
3.47 Frequency Distribution of 1st example CW event	105
3.48 Frequency Distribution of 2nd example CW event	106
3.49 Frequency Distribution of 3rd example CW event	106
3.50 Reconstruction of Simulated Events	107
3.51 Reconstruction of Simulated Events vs SNR	107
3.52 RMS of Pointing Resolution vs SNR	108

3.53	Difference in reconstruction of Taylor Dome sample by two filtering methods	109
4.1	A constructed waveform use to demonstrate cross-correlation.	115
4.2	Cross Correlation of constructed waveforms	116
4.3	New Cross Correlation of constructed waveforms	117
4.4	Interferometric Map of example Taylor Dome event	118
4.5	Refined Interferometric Map of example Taylor Dome event	119
4.6	Interferometric Map of example Taylor Dome event	120
4.7	Coherently summed waveform of Taylor Dome event	120
4.8	10 % sample placed into HealPix Bins.	122
4.9	An example of error ellipse on HealPix bins.	123
4.10	Antenna chosen in noise sample with largest voltage peak	128
4.11	Antenna chosen in no-noise sample with largest voltage peak	129

4.12	Trigger Efficiency Distributions	130
4.13	Efficiency curves of the simulated set and Taylor Dome set	131
4.14	Ratio of Peaks Distribution	136
4.15	Peak Value of Cross Correlation Distribution	137
4.16	Peak Hilbert Distribution	138
4.17	Polarization Fraction Distribution	140
4.18	Elevation Angle Distribution	141
5.1	Rotated Cross Correlation for Taylor Dome	144
5.2	2D scatter plot for Rotated Cross Correlation Cut	146
5.3	The differential plot of a HealPix Bin used for optimization	147
5.4	The likelihood distribution of a HealPix Bin	148
5.5	P-Value distribution for slope	149

5.6	$\frac{\text{Signal}}{S_{\text{up}}}$ as a function of y-intercept	152
5.7	Rotated Cross Correlation distribution with optimized cut value	153
5.8	Location and numbering of HealPix bins used in this analysis	154
5.9	Comparison of efficiency of analysis cuts between the old and new analyses	158
B.1	Optimization plots for Bin 3012	167
B.2	Optimization plots for Bin 3030	168
B.3	Optimization plots for Bin 3031	169
B.4	Optimization plots for Bin 3032	170
B.5	Optimization plots for Bin 3033	171
B.6	Optimization plots for Bin 3045	172
B.7	Optimization plots for Bin 3048	173

B.8 Optimization plots for Bin 3051	174
B.9 Optimization plots for Bin 3053	175
B.10 Optimization plots for Bin 3057	176
B.11 Optimization plots for Bin 3061	177
B.12 Optimization plots for Bin 3062	178
B.13 Optimization plots for Bin 3063	179
B.14 Optimization plots for Bin 3066	180
B.15 Optimization plots for Bin 3069	181

Chapter 1

Introduction

The study of UHE neutrinos is ultimately a study of the most distant and powerful sources in the universe. During the past century, scientists have used particles to probe and study the universe. The Earth is constantly struck by charged particles called cosmic rays (CR), and the study of those particles has been extensive and has given many new insights into sources of highly energetic particles in the universe. With the study of cosmic neutrinos, we can push the boundaries of our understanding further still, finding new sources and broadening our understanding of the most powerful sources of particles in the universe.

UHE neutrino astronomy (above 10^{18} eV) is rather new in particle astrophysics and promises to open up a new frontier in the highest energy regimes where other particles, such as cosmic rays or gamma rays, are limited by their attenuation as they travel cosmic distances. These particles may even provide details on sources, such as type and location of UHE Cosmic Rays (UHECRs), a mystery which has not been solved.

1.1 Neutrino Astronomy

Neutrino astronomy began from a quest to understand how our sun works. It led to an understanding of the properties of the neutrino, then to using the neutrino to

probe for sources in the universe. Neutrino astronomy has a very limited number of known sources. The first known source is the sun (solar neutrinos). The first non-solar astronomical neutrinos were observed from Supernova 1987A. A constant source of neutrinos are atmospheric neutrinos. These neutrinos are created when charged particles interact with the Earth's atmosphere. Since 2013, IceCube has discovered tens of events from outside our solar system, believed to be extra-galactic, although no known sources have been identified. [3–5].

1.1.1 Solar Neutrinos

Ray Davis won a Nobel Prize for his Homestake Mines experiment measuring the solar neutrino flux [6]. This experiment was designed to find and measure the solar neutrino flux by measuring ν_e collisions. Davis only saw 27% of the expected flux. This result paved the way for other experiments [7, 8] to find the deficit, the "solar neutrino problem".

The deficit of solar neutrinos was paramount to our understanding of neutrino oscillations. Our current understanding is that the neutrino flavor eigenstates (ν_e, ν_μ, ν_τ) are superpositions of the neutrino mass eigenstates (ν_1, ν_2, ν_3). Neutrinos have mass, so a neutrino can oscillate between flavors as it travel through space. The mixing of the different flavor states has been studied by both interactions in Earth's atmosphere and by neutrinos created in nuclear reactors and particle accelerators.

1.1.2 Supernova 1987A

In 1987, neutrino astronomy had a stroke of luck. Various experiments were studying solar and atmospheric neutrinos when a supernova, Supernova 1987A, occurred. Supernova 1987A occurred in the Large Magellanic Cloud and was observed on February 23, 1987. Kamiokande [9] and IMB [10] both saw a burst of neutrinos

corresponding to this supernova, 20 total events over 13 seconds. Kamiokande and IMB were both water Cherenkov detectors, using photomultiplier tubes to detect the visible Cherenkov radiation as described in Section 1.4.1. Using these 20 neutrinos, the two experiments were able to place an estimate on the total energy output of the supernova, $(6 \pm 2) \times 10^{46}$ Joules, an upper limit of the mass of the electron neutrino, and limits of the neutrino lifetime [11,12]. Most importantly, analysis of these neutrinos proved consistent with models of the supernova process. It proved that neutrinos are an important particle in astronomy and could be used to probe sources in ways that other particles could not. Since this supernova, no other extra-solar events were detected until IceCube began detecting neutrinos.

1.1.3 IceCube

The IceCube experiment has begun detecting extra-solar neutrinos since 2013. IceCube is an optical Cherenkov (see section 1.4.1) experiment buried in the Antarctic ice near the south pole, with 1 km^3 of ice instrumented. (Figure 1.1).

Using digital optical modules (DOMs) placed onto 86 strings buried 1450-2450 m in the ice, IceCube relies on detecting optical Cherenkov light from neutrino interactions in the ice. They have seen an excess above the expected atmospheric neutrinos flux inside the detector, leading to a discovery of 28 extra-galactic neutrino events with a maximum energy of $\sim 10 \text{ PeV}$ [14, 15]. Reconstruction locations for those events can be seen in Figure 1.2.

This discovery of extra-solar neutrinos is a milestone in neutrino astronomy. It has shown that neutrino astronomy is on the cusp of great progress. These few events have given a new life to the field with new models arising and searches for apparent sources being performed. It is an exciting time for neutrino astronomy as multiple searches for possible sources [3–5] have occurred. Furthermore, studies into the particle aspect

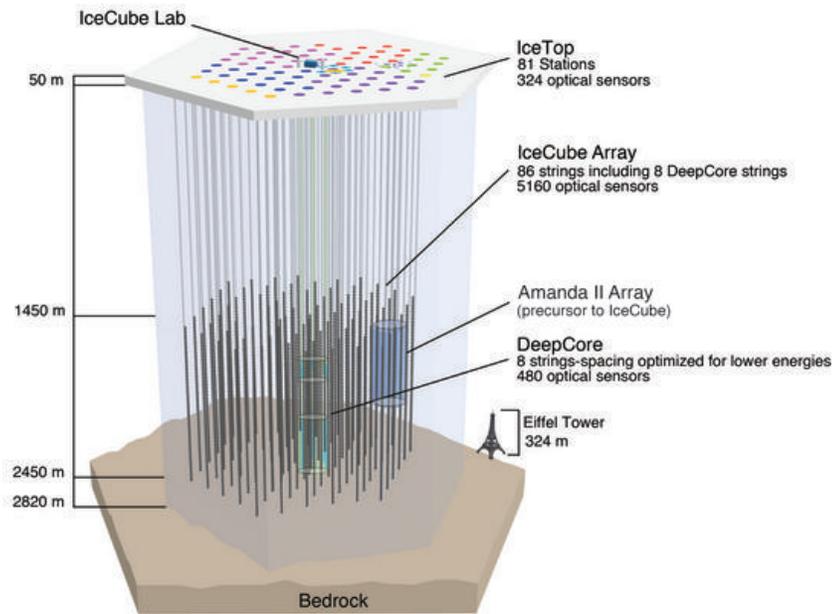


Figure 1.1: A depiction of IceCube. From [13]

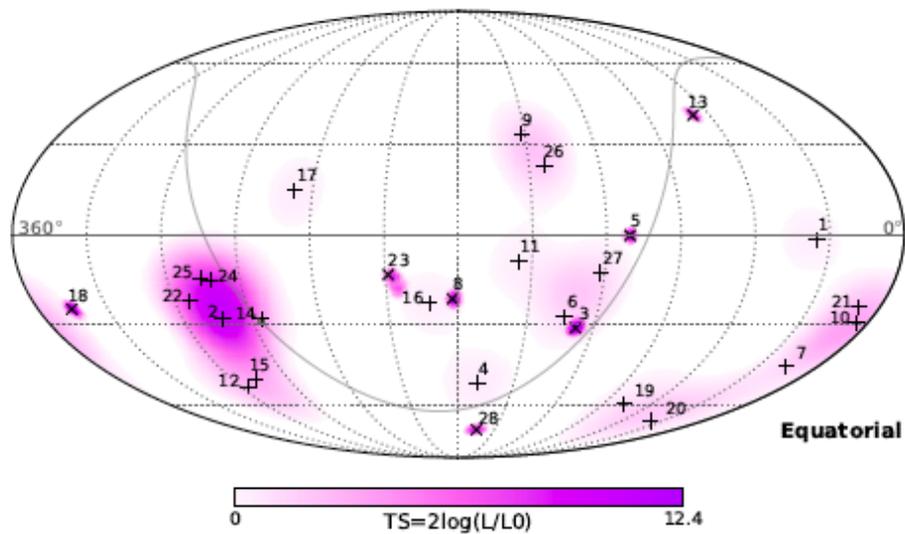


Figure 1.2: A look at reconstruction direction on the 26 extra-solar neutrinos discovered by IceCube. There was no clustering significance from these events. From [15].

of neutrinos can be investigated as well. The Standard Model predicts a cross section for the neutrino as [16]:

$$\sigma = 2.501 \times 10^{-39} \left(\frac{E_\nu}{10^9} \right)^{0.3076} \text{ cm}^2 \quad (1.1)$$

Past studies of the cross section were done at various energies [17–20]. With these new interactions, testing of this cross section at higher energies is now possible.

1.2 UHE Astronomy

UHE astronomy attempts to detect particles produced by sources in the universe: photons, cosmic rays, and neutrinos. Many experiments [21–24] were built to measure photons and cosmic rays interactions. These experiments told us much about the universe, but are limited, especially at ultra-high energies. The highest energy gamma rays measured by Fermi Lat are in the TeV range [25] and the cosmic ray flux begins to die away above 10^{19} eV [22, 23], leaving neutrinos to probe this energy space.

1.2.1 Cosmic Rays

Cosmic rays are charged particles traveling through space. At lower energies, the composition of these particles is well known. In the UHE regime, these particles are believed to be mostly proton or iron nuclei, but the exact contents are still unknown [27–30]. Due to their charged nature, the arrival directions of these particles are isotropic. The cosmic rays are bent in the magnetic fields and cannot be used to point to sources except possibly at the highest energies.

Measurement of these cosmic rays has produced the flux seen in Figure 1.3. The cause of the shifts in spectral power, called the knee ($\sim 10^7$ GeV) and ankle ($\sim 10^{10}$ GeV) respectively, are unknown. The knee is thought to be occur due to the

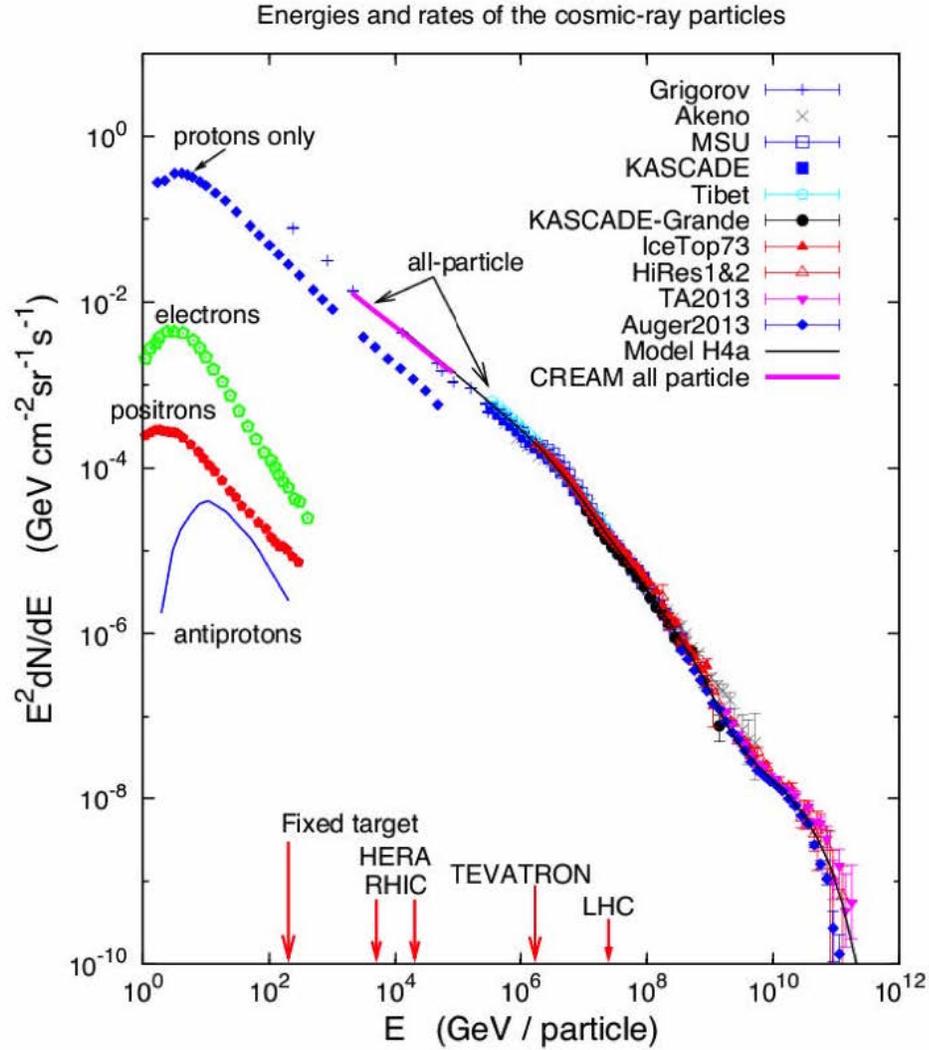


Figure 1.3: Cosmic Ray spectrum as measured by various experiments. There is a shift if the spectral power at $\sim 10^7$ GeV called the knee, and another shift at $\sim 10^{10}$ GeV called the ankle. The arrows along the bottom axis show the equivalent energy of collisions from colliders. From [26]

limited strength of the magnetic fields in our galaxy. At those energies, the cosmic rays are capable of leaving the galaxy, resulting in a lower number of particles hitting the Earth. A second theory is that the particle accelerators in our galaxy begin to fall off at this energy [31]. The ankle can be caused by different effects, but one theory is the cosmic ray spectrum changing from galactic spectrum to extra-galactic. The absence of cosmic ray particles at the highest energies is quite a mystery. The leading theory calls for an interaction of these cosmic rays with the cosmic microwave background, leading to a reduction in the energy of the cosmic ray. This drop in flux could also be caused by a cutoff in the injection spectrum, i.e. no sources being able to accelerate cosmic rays to these energies. This question is one of the first neutrino astrophysics hopes to answer. Due to the lack of particles at these ultra-high energies, cosmic rays are not the ideal particles to investigate this energy region.

1.2.2 Neutrinos

Neutrinos are a good cosmic messenger at these ultra-high energies. Unlike cosmic rays, neutrinos are uncharged so they are not bent by galactic magnetic fields. Additionally, neutrinos only interact weakly, which allows them to point back to their source. Cosmic rays at above 10^{19} eV interact with the cosmic microwave background (CMB) via the delta resonance, leaving a relatively short (< 50 Mpc) observation distance for cosmic ray astronomy (Figure 1.4). Cosmic rays can survive the trip to the Earth if the energy is below this threshold, but cannot be used to find sources due to magnetic fields bending the flight path. Weak, uncharged neutrinos are good candidates for an unattenuated look at the most distant and powerful sources in our universe.

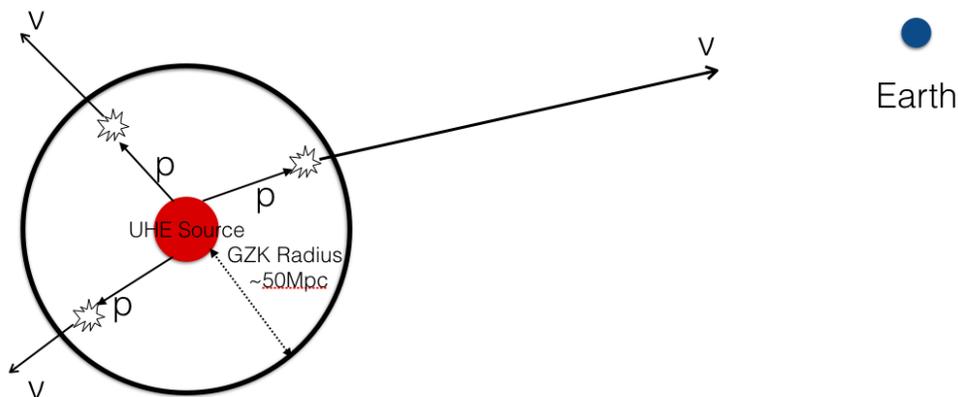


Figure 1.4: A cartoon of the GZK process, not to scale. Description of this process can be found in Section 1.3.2.

1.3 Production of UHE Neutrinos

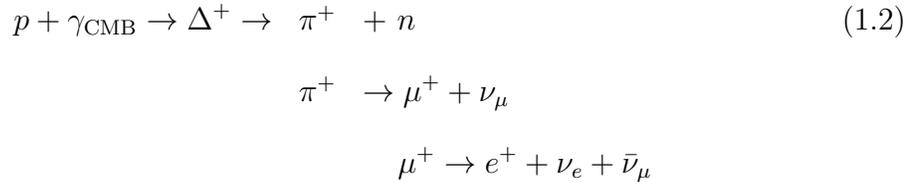
There are various theories that explain potential sources of UHE neutrinos. Neutrinos may be created directly at a source (direct production) that is also producing cosmic rays. There are also theories that have neutrinos being created outside the original sources of cosmic rays through a number of interactions.

1.3.1 Direct Production

Various experiments have measured cosmic rays of energy above 10^{19} eV. Possible sources for these UHE particles include active galactic nuclei (AGN) and gamma ray bursts (GRBs). These objects are theorized to accelerate charged objects via Fermi shock mechanism. The collisions that occur during the shock wave could also produce UHE neutrinos as well as the cosmic rays seen [32].

1.3.2 GZK Neutrinos

The shift in cosmic ray spectrum at the ankle is believed to be a shift to a extragalactic sources. For particles of that high energy, cosmic rays can interact with a photon of the cosmic microwave background (CMB). This interaction typically occurs within 50 Mega-parsecs of the source, going through a delta resonance and producing a pion and a proton or neutron (Figure 1.4). The charged pion will decay, producing UHE neutrinos. This process is called the Greisen-Zatsepin-Kuzman (GZK) process ([33,34]). This process promises neutrinos as long as there are protons being produced at sources above $10^{19.5}$ eV.



While these GZK neutrinos are guaranteed as long as the parent particles are being produced, the flux of the neutrinos is unknown. Many models [35–37] have been created to predict the flux while also molding to known observations: the IceCube flux and the cut-off in cosmic rays observed by Pierre Auger Observatory (Auger) and The High Resolution Fly’s Eye Experiment (HiRes). There are many parameters that can be fine tuned for these models, such as the source evolution and injection spectrum. Since these neutrinos arise from cosmic rays, the assumed composition of UHECRs plays a very important part in the expected flux. There is debate on whether UHECRs are mostly protons or a heavier nuclei such as iron, and this difference impacts the flux of neutrinos due to the delta resonance in the chain of reactions.

Detecting GZK neutrinos can provide new information for astroparticle astronomy. Since the range of interaction on these cosmic rays is relatively small (< 50 Mpc), these resultant neutrinos will point back to their source, giving clues about the particle accelerators in the universe. Another bit of information we would learn is that these sources do not have an intrinsic cut-off at the same energies that the GZK process would turn on. It could be possible that sources creating UHECRs have a cut-off and do not produce cosmic rays above certain energies. By finding GZK neutrinos, it would show that the loss of those UHECRs is due to the GZK process turning on, rather than the particles not being produced.

1.4 Detection of UHE particles

One of the best ways to find high energy particles is through use of optical Cherenkov light. Many experiments were designed around this technique including Kamiokande, IMB, Super Kamiokande and IceCube. These experiments were designed to find particles at energies lower than 10^{18} eV, where the neutrino flux is larger. UHE neutrinos are expected to impact the Earth at a rate of $\sim 1 / \text{km}^2 / \text{century}$. One way compensate for this low flux is to have a large volume for the detector. This presents logistic problems when dealing with optical light due to attenuation lengths in ice and water being on the order of meters.

1.4.1 Optical Cherenkov

Cherenkov radiation occurs when a charged particle travels through a dielectric medium faster than light can travel through that medium. This radiation is broad-band and increases with frequency. Photomultiplier tubes (PMTs) inside the medium can detect this radiation.

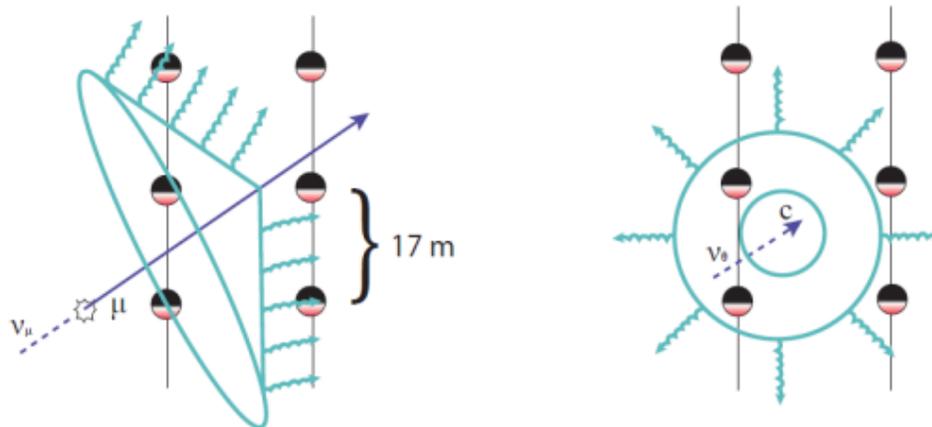


Figure 1.5: A cartoon showing the radiation from different neutrino interactions. Left, muon neutrinos will produce a muon that will leave a track in the detector. Right, electron and tau neutrinos will produce the corresponding lepton which will interact in a short length, creating a spherical pattern in the PMT trigger pattern. Figure from [38]

The charged particles needed for this technique are created when a neutrino interacts in the medium. The neutrino interaction creates a shower of charged particles that propagate through the medium. These particles each produce Cherenkov radiation seen by the photomultiplier tubes. Flavor differentiation of the neutrinos occurs by the pattern of arrival times of the Cherenkov radiation in the PMTs. Electron and tau flavored neutrinos produce a spherical pattern due to the short lifetime of the leptons in the medium, while muon neutrino interactions have long tracks through the detector. Examples of these timing patterns can be seen in Figure 1.5, illustrated by the IceCube detector [38].

1.4.2 Askaryan Effect

A different method of detection of charged particles has been developed that relies on particle showers, rather than individual particles, propagating through a dielectric

medium, such as ice. When a UHE neutrino interacts in ice, it will produce a shower of particles. This shower will develop a negative charge asymmetry ($\sim 20\%$) due to its propagation through matter, with electrons are drawn into the shower while positrons annihilate. The Cherenkov radiation from this excess charge is coherent if the wavelength of the radiation is large compared to the Molière radius of the shower. In ice, the Molière radius is ~ 10 cm, which places the coherent emission in the radio regime. Gurgen Askaryan predicted this strong coherent radio emission from showers in a dielectric [39, 40].

This use of coherent radio emission brings many benefits. With optical Cherenkov light, the power of the radiation scales linearly with the number of particles in the shower. When the emission becomes coherent, the power scales as the number of particles in the shower squared ($P \propto N^2$). For UHE particles, the strongest emission is in the radio regime due to the coherence. Another bonus for radio is the attenuation length of the electromagnetic signal. In ice, the attenuation length of optical light is on the order of 100m [41] while the attenuation length of radio in ice is 1 km [42]. This difference in attenuation lengths allows for a much large volume of ice to be instrumented since the density of detectors will go down.

This coherent emission was first observed at a SLAC beam test in 2001 using a sand target [43] and confirmed using a salt target [44]. In 2006, the ANITA instrument was used to detect the radio emission from an ice target [45]. Proof of the $P \propto N^2$ can be seen in Figure 1.6. From this test, it was seen that the emission is characterized by a broadband spectrum, linear polarization, and coherence. A typical waveform from the 2006 SLAC beam test can be seen in Figure 1.7, with a typical rise time of ~ 100 ps.

Many experiments are now using this technique to explore the UHE regime. Experiments such as RICE, GLUE, ANITA, ARA, and ARIANNA [46–50] are using

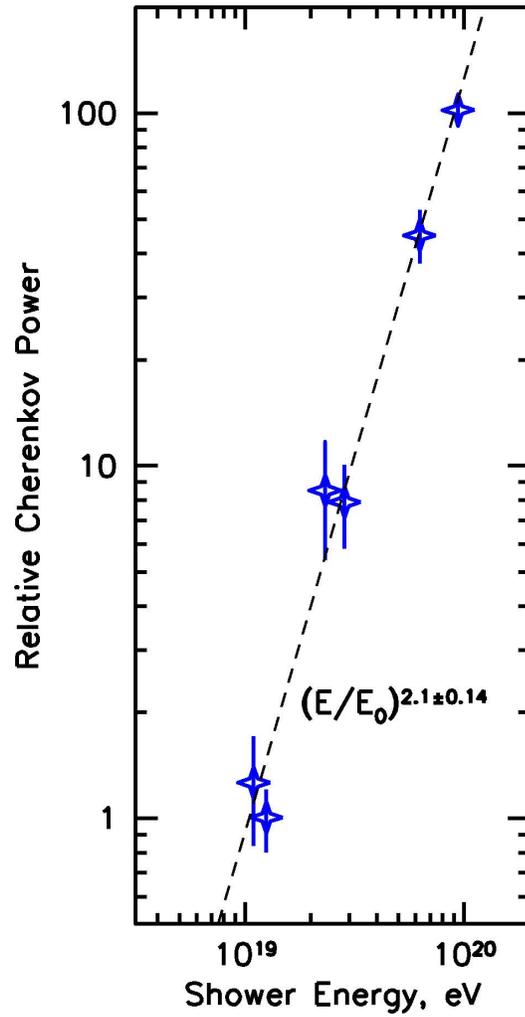


Figure 1.6: Observation of Askaryan effect in ice. Power in the effect is proportional to the number of particles squared. From [45]

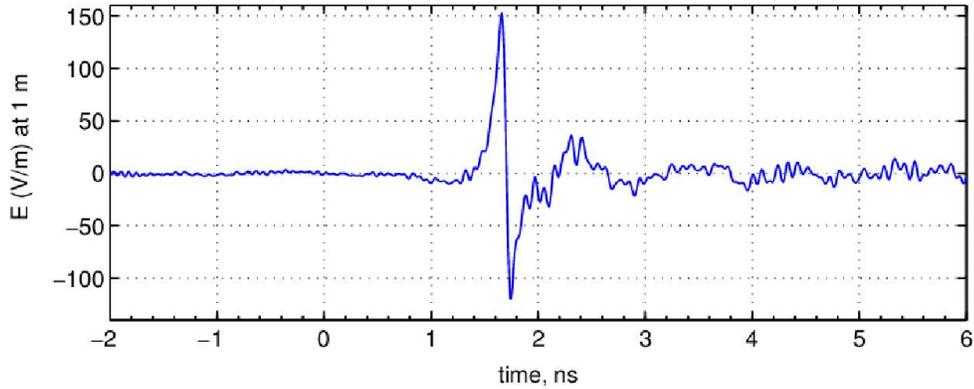


Figure 1.7: A typical waveform from the SLAC test. From [45]

the radio Cherenkov technique in differing mediums in an attempt to find this UHE neutrino flux.

1.5 Geo-Magnetic Radiation from Cosmic Rays

When cosmic rays enter Earth’s atmosphere, they produce particle showers. The electrons and positrons in those showers will be bent by Earth’s magnetic field, and the particles will spiral around the field lines. This acceleration will cause forward-beamed synchrotron radiation, which is coherent on scales larger than the shower. The shower size is on the order of a meter, which places coherent radiation at or below a few hundred MHz. Similar to the coherence from Askaryan Radiation, the power in the coherent emission also scales as a square of the energy. The cosmic ray signal is characterized by lower frequency content than neutrino emission, linear polarization and coherence. Geo-synchrotron emission models can be found in references [51, 52]. Cosmic Rays have been detected from this emission [53] but ANITA was the first experiment to detect cosmic rays in the UHE regime using this technique.

Chapter 2

The ANtarctic Impulsive Transient Antenna

The ANtarctica Impulsive Transient Antenna (ANITA) was designed to be the most sensitive experiment for detection of UHE neutrinos ($E_\nu > 10^{19} eV$). ANITA is a balloon-borne array of radio antennas, which observes the Antarctic ice and looks for impulsive radio signals expected from a UHE neutrino interaction in the ice.

Figure 2.1 depicts an event in the ice. A neutrino that interacts in the ice will produce a Cherenkov cone that is coherent in the radio regime. Part of this coherent signal will pass through the ice-air surface and can be observed by ANITA. The characteristic feature of this signal would be its short (nanoseconds) duration, it would be a broadband, up to ~ 1 GHz, impulsive signal. Due to the geometry of the problem, the only neutrinos ANITA is expected to see are slightly up-going (Earth-skimming) events, up with respect to the orientation of the payload. At this energy, the Earth is opaque to neutrinos. Neutrinos that are mostly up-going will be absorbed by the Earth. Neutrino events which are down-going often do not have their signal escape the ice and propagate to the antennas due to the orientation of the Cherenkov cone.

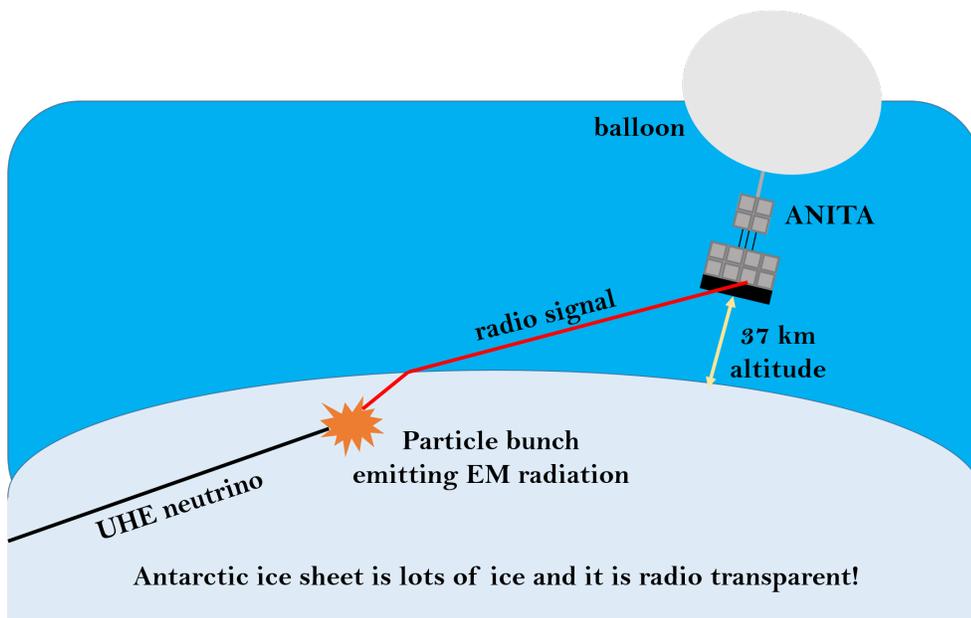


Figure 2.1: A cartoon depicting event geometry for ANITA. Picture from O. Banerjee.

2.1 ANITA-1

The first flight of ANITA was launched in December 2006 and flew for 35 days. The live-time of the experiment was 17.3 days. Analysis showed no neutrino events during this time [48]. This allowed a limit to be placed on the UHE neutrino flux seen in Figure 2.2.

ANITA-1 observed 16 cosmic rays events during the flight [54]. These events show a high degree of similarity in both shape and frequency content as seen in Figure 2.3. In Figure 2.4, we can see the correlation between the direction of the geomagnetic field at the shower location and the polarization of the event. This shows the path of the charges are being bent by the Earth's magnetic field and radiating. The red line shows the expected measured polarization, which is in good agreement with the points.

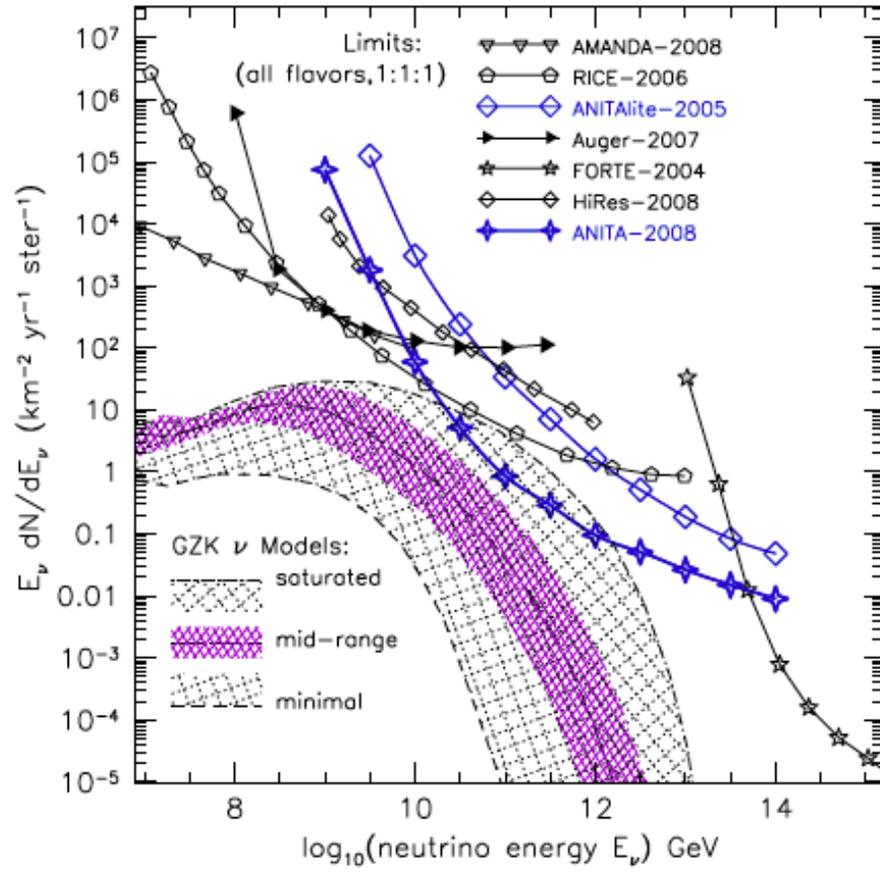


Figure 2.2: A look at the flux limits from the ANITA-1 flight. This flight set the lower flux limits above 10^{18} eV at the time. From [48]

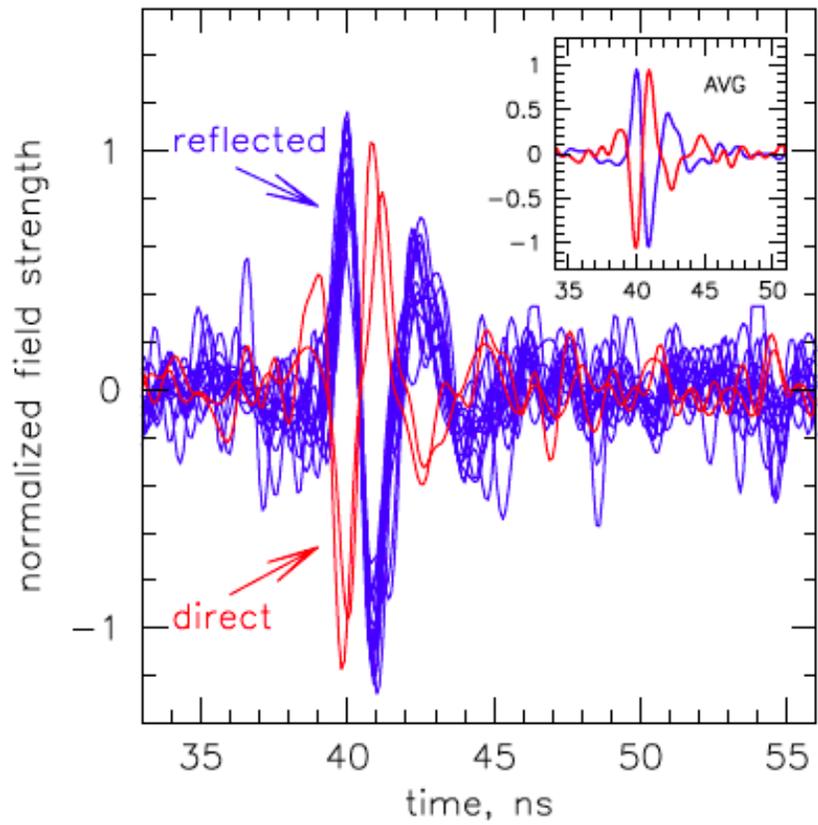


Figure 2.3: A look at the waveforms of the 16 cosmic ray events seen by ANITA-1. 14 of the events were reflected off the ice (blue) while 2 events were direct (red). The inset shows the average waveforms. From [54]

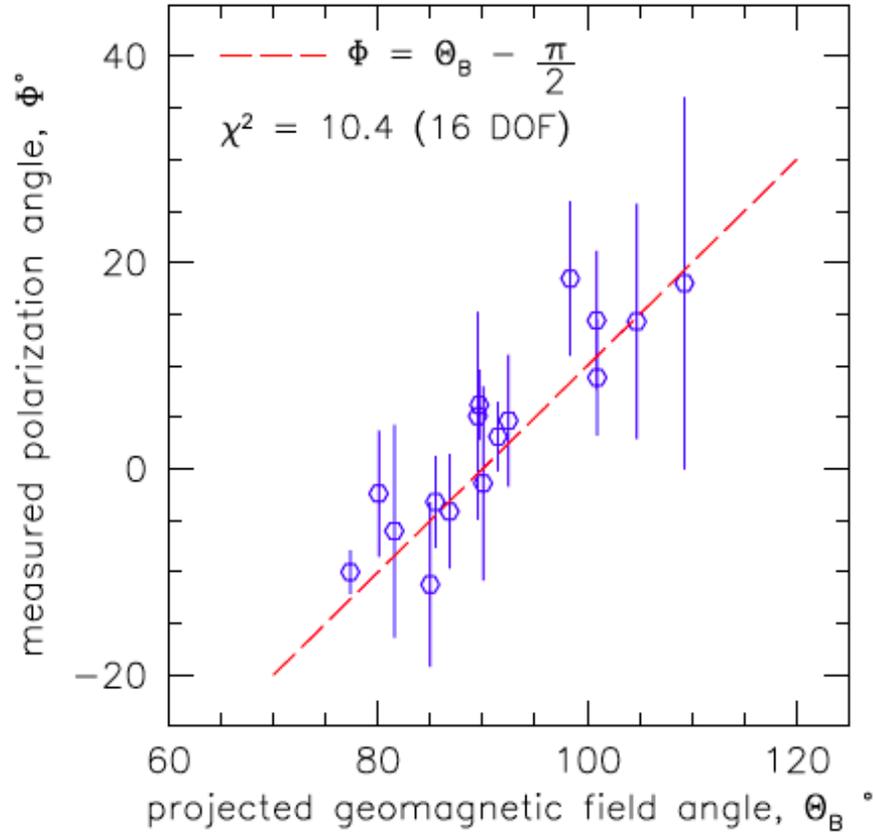


Figure 2.4: The angle of Earth’s magnetic field at the location of the shower, Θ_B , predicts the polarization of the event plotted against the polarization of the event Φ . The red dashed line is the expected polarization angle, and is in good agreement with the data points. From [54].

2.2 ANITA-2

2.2.1 Gondola

The Columbia Scientific Balloon Facility (CSBF) has strict requirements for the size and weight of a payload. In addition to those requirements, the payload needs to be easily disassembled for shipping and retrieval after the flight. This led to the gondola being crafted out of hollow aluminum tubes connected by joints. There were 40 highly directional radio antennas arranged in four layers on the ANITA-2 payload. The top two layers of antennas both had eight antennas each, spaced at 45° apart. The two layers have an offset between them of 22.5° for better coverage. These two rings are considered one full layer for triggering. The middle layer of antennas consists of 16 antennas spaced 22.5° apart. The bottom layer of antennas, nadir ring, contains eight antennas, spaced 45° apart. This layer of antennas is stowed inside the gondola during the launch and are deployed after liftoff. The lowest layer is aligned with the second layer of antennas. Each antenna is given a downward cant angle of 10° below the horizontal to maximize the amount of ice seen by each antenna.

Contained on the deck on top of the middle layer of antennas is the CSBF Science Instrument Box (SIP). The SIP is powered and operated independently of ANITA. It is used for flight control (e.g. ballast and flight termination) as well as all telemetry. It also provides a separate navigational system than that used in ANITA.

The telemetry provided by the SIP performed two useful functions during the flight. Firstly, it provided a data link to the payload. Where possible, the payload used line-of-sight (LoS) transmission and used Tracking and Data Relay Satellite System (TDRSS) at other times. This data link allowed monitoring of systems to ensure the payload functioning correctly (temperature, battery life, etc). This telemetry was also able to transmit a small sample of events to ground shortly after they were



Figure 2.5: A picture of the ANITA-2 payload after launch. The lowest layer of antennas were folded underneath the payload during launch and deployed during ascent. From [54].

recorded. The rate of telemetry was much lower than the data acquisition rate, so a prioritization program was created to ensure high priority events were sent first along with a small background sample. This backup was to ensure that if the payload was irretrievable or the data was corrupted, we would still have a small sample of events that contained possible neutrino candidates. The second use of the SIP telemetry was the ability to send basic commands to the payload. These commands were used to set trigger thresholds and to turn on/off power to certain systems.

The ANITA instrument box and the battery box were also contained on the deck of the payload. Many of the components on the payload were painted white to reflect sunlight and help to regulate temperature as the payload flies at high altitude, where heat loss through conduction and convection is minimal. Components that produce a large amount of heat were connected to a radiator plate on the instrument box which was covered with Teflon tape to help to radiate the heat.

2.2.2 Power System

The power for ANITA was provided by an omni-directional photo-voltaic (PV) array that was located between the middle ring and the nadir ring on the payload. The outputs from the array led to a charge controller. The charge controller distributed a steady 24V to the payload, using DC-DC converters that were implemented to provide a set of voltages (12V, -12V, 3.3V, 5V) used in various systems. The charge controller was also connected to a battery farm of 12V lead-acid batteries. The charge controller was capable of pulling the necessary power from the batteries to provide the needed 24V. When the PV array was providing enough power alone, the charge controller was able to charge the battery farm.

2.2.3 Position and Orientation

ANITA relies heavily on knowing the location and orientation of the payload to a high precision. It is necessary to have this info to differentiate between anthropogenic noise and possible signal events. ANITA-2 used three different GPS systems during its flight to ensure no loss of data.

Two of the GPS systems were square antenna arrays. Each array consisted of four Magellan ADU5 GPS antennas and were situated on the top of the gondola. These arrays provided heading information that was accurate to $< 0.2^\circ$ and pitch and roll values accurate to $< 0.5^\circ$. The third GPS system was a Thales G12 antenna, also placed on top of the gondola. This GPS system also updated absolute time on the flight computer through its Network Time Protocol (NTP) server.

There were a host of other sensors and arrays that were used as backup for these GPS systems. The SIP provided by CSBF contained a GPS system. Four sun-sensor instruments, as well as a magnetometer and accelerometer were located on the deck of the gondola. These precautions, specifically the sun sensors, proved useful as ANITA-2 experienced a loss in GPS data for $\sim 5\%$ of the flight.

2.2.4 Radio Frequency Signal Chain

Radio frequency (RF) signals observed by ANITA-2 were filtered, amplified and digitized from each of the eighty channels before being written. A schematic of this signal chain is shown in Figure 2.6, from the antennas through the signal path, the trigger logic, and readout. The signal chain will be described in detail in the following sections. ANITA-2 only triggered using VPol information, which accounts for the slight differences between hardware in polarizations.

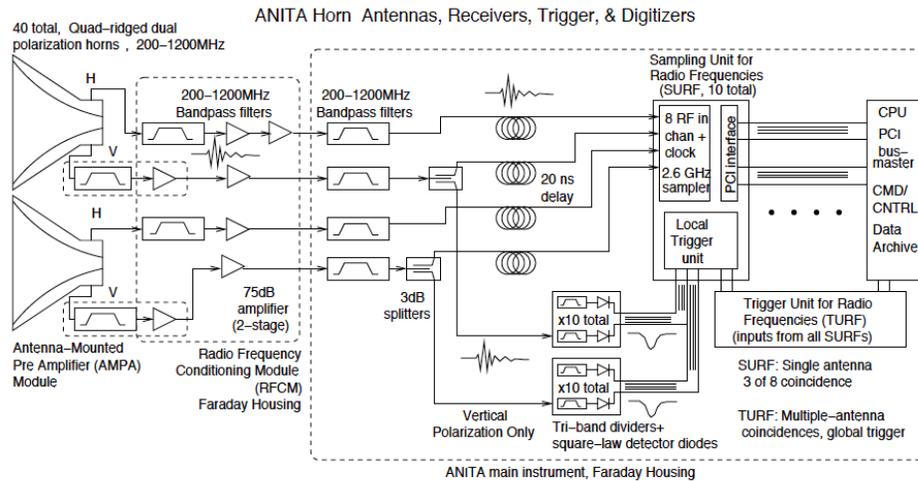


Figure 2.6: The signal chain for ANITA-2. Adapted from ANITA-1. From [55].

2.2.4.1 Antennas

ANITA-2 used custom built quad-ridged horn antennas built by Seavey Engineering. These antennas were designed to have a 3 dB bandwidth between 200 and 1200 MHz, shown in Figure 2.7. Each antenna had a face of 0.8 m on a side and had two feeds which allowed for two linear polarizations. ANITA-2 recorded data in the vertical polarization (VPol) and horizontal polarization (HPol), resulting in 80 RF channels.

These antennas were highly directional, with an on-axis gain of 10 dBi. The 3 dB point for these antennas is around 30° . The beam pattern also depends on frequency, with a wide beam at low frequencies and a more narrow beam at higher frequencies. The antennas were spaced at 22.5° to ensure optimal overlap between beam patterns while still having full azimuthal coverage. Antennas were given a downward cant angle of 10° to place the horizon and most observable ice area inside the 3 dB beam

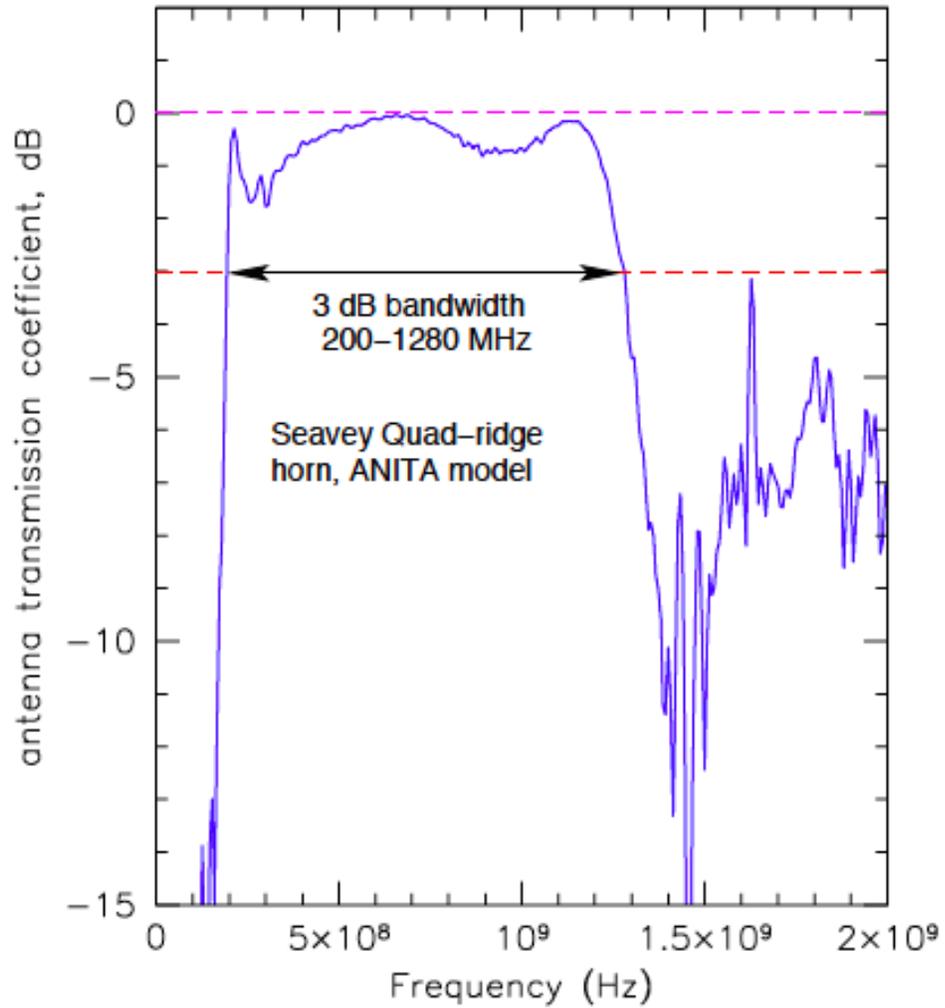


Figure 2.7: Seavey transmission coefficient as a function of frequency. This shows the frequency range in which the antennas are most sensitive. From [56].

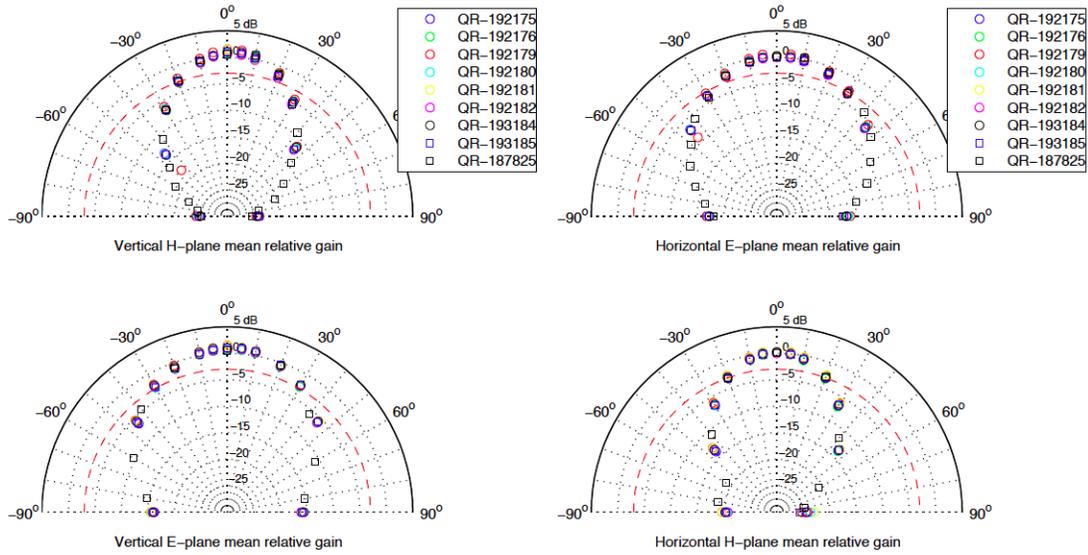


Figure 2.8: The beam pattern in dB of 9 different antennas for ANITA-2. Left plots are for VPol for both E and H planes. Right plots are similar but for HPol. From [56].

width. The beam pattern can be seen in both vertical and horizontal polarizations in Figure 2.8.

The antenna impulse response must be well known to help distinguish noise from impulsive, neutrino-like events. This response is important for our understanding of the triggering of the system as well as analysis. A flat group delay across the frequency band for the antennas is essential as this is needed for a fast impulse response. The impulse response is shown in Figure 2.9. While the ringing occurs for tens of nanoseconds, most of the power is contained in the first few nanoseconds of the pulse. An impulse response was created by adjusting the ANITA-1 impulse response, seen in Figure 2.10.

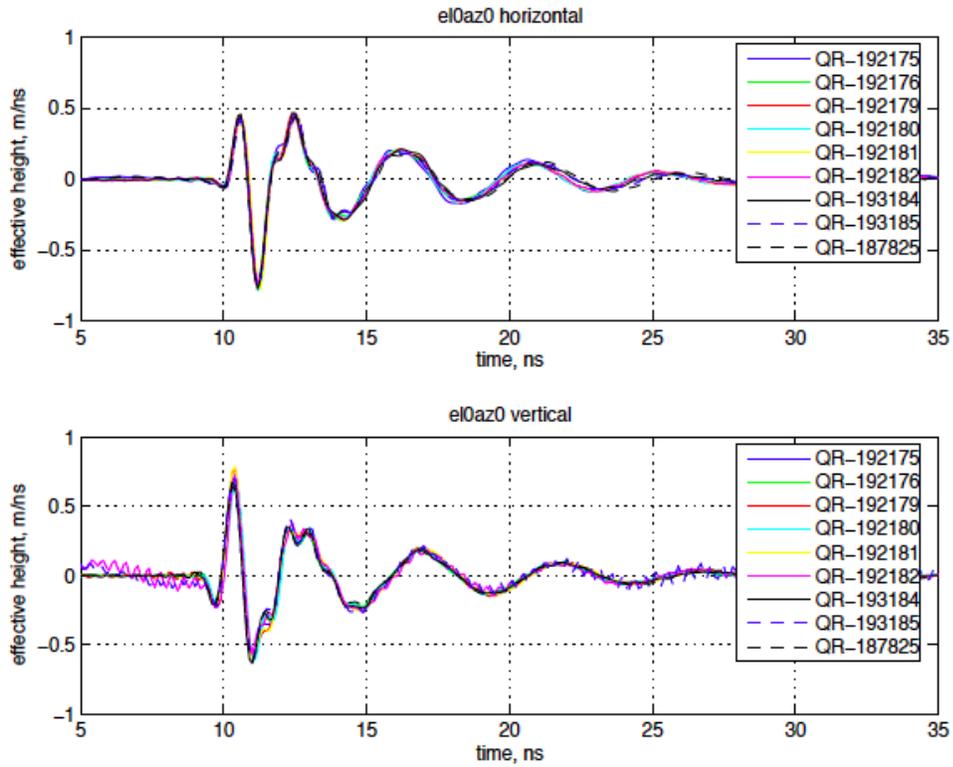


Figure 2.9: The impulse response of nine different antennas. The top plot is for HPol and bottom is for VPol. From [56].

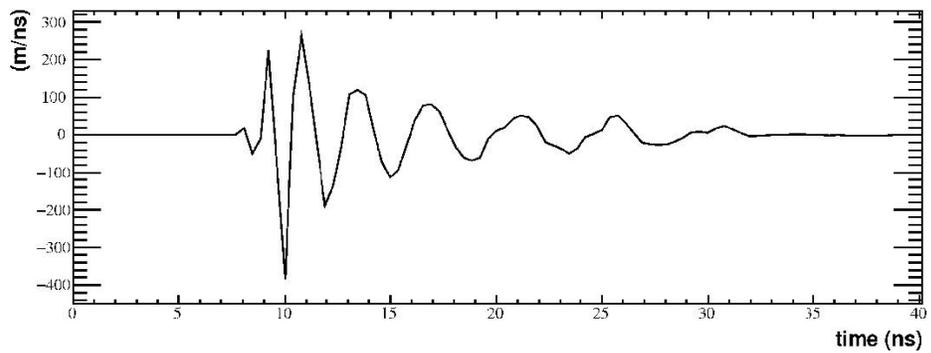


Figure 2.10: The impulse of ANITA-2, full RF chain excluding the antenna. From [55].

2.2.5 Front End Signal Processing

Each of the 80 channels present in ANITA-2 were amplified and filtered near the antenna output. Having the amplifier near the antenna output reduces cable noise and reduces transfer losses in the cable. The amplification was done by low-noise amplifiers (LNAs).

There were slight differences between VPol channels and HPol channels. For VPol channels, there is an antenna-mounted preamplifier module (AMPA) which was directly connect to the antenna output. The AMPA houses a low-noise bandpass filter, followed by an LNA, which has a power limiter on the input. Downstream from the AMPA was the radio frequency conditioning module (RFCM). This module provided the 12V needed by the AMPA by sending the power over the signal cable, which the AMPA decoupled from the signal in a Bias T. The RFCM also provided second stage amplification (with a 3 dB attenuator between the AMPA and RFCM to dampen oscillations). For HPol channels, both amplifier stages occurred within the RFCM. The HPol channels had a slightly higher noise figure in amplifiers and filters for the first stage of amplification and filtering than the VPol channels.

Due to the differences in amplifiers and filters, the VPol channels have lower noise and lower gain than the HPol channels (~ 40 K, ~ 2 dB). The noise temperature and gain can be seen in Figure 2.11. This choice to have VPol channels differing from HPol came from time and money constraints. Improvements in the hardware were focused on VPol channels due the expectation that the neutrino signal would be vertically polarized.

The signal is then transmitted to the ANITA instrument box where various actions occur. First, the signal undergoes a second stage band pass filter (200-1200 MHz). The VPol signal is then sent through a 3 dB splitter, where one path goes to the data acquisitions unit (DAQ) and the other is held in a delay cable. If a trigger is seen, the

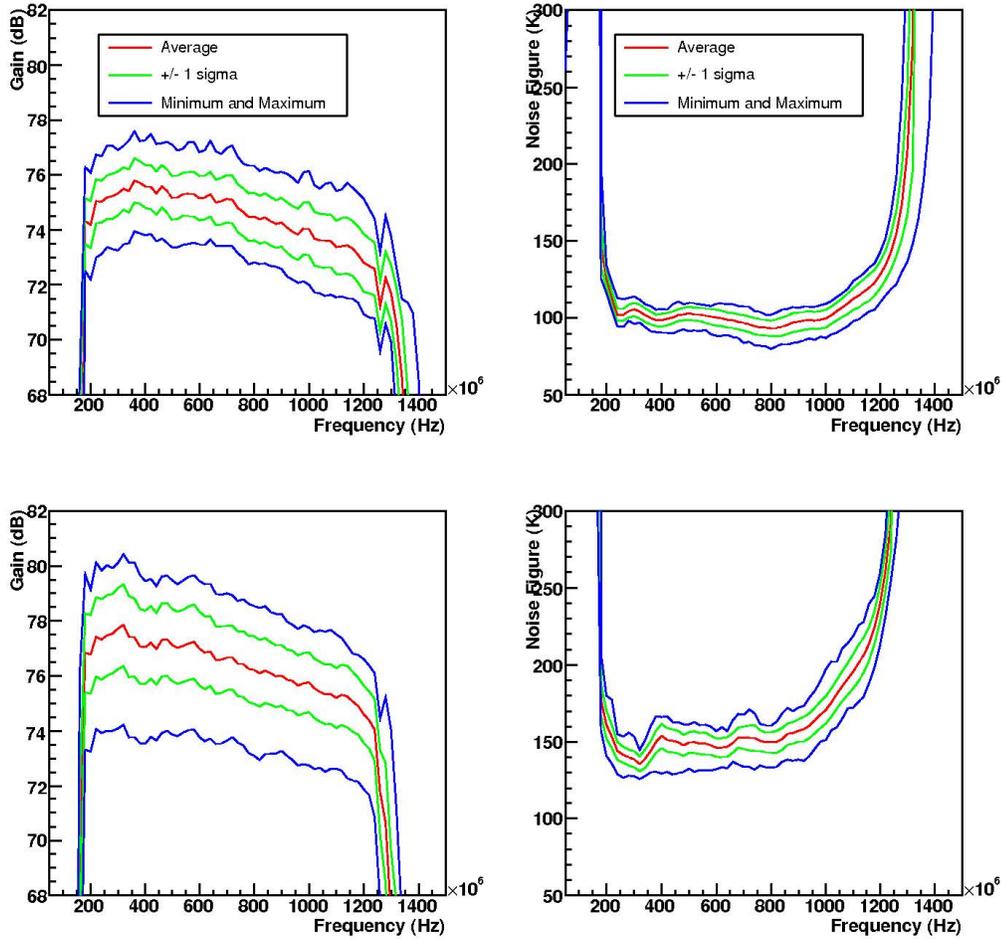


Figure 2.11: The average, 1 standard deviation over all channels, maximum and minimum gain and noise temperature from RFCMs for VPol and HPol channels. Top plots are VPol, bottom are HPol. From [55].

signal from the delay cable is processed and recorded. Since ANITA-2 only triggered on VPol signals, there is no splitter for the HPol signal.

Inside the DAQ, the signal path was fed into Sampling Unit for Radio Frequency (SURF) boards. Each SURF board has eight input RF channels (four VPol and four HPol). A total of ten SURF boards were used in the data-taking of the eighty RF channels. To reduce cross-talk in cables, no two adjacent antennas were fed into the same SURF. On the SURF, the signal from each channel was split and fed into four parallel Large Analog Bandwidth Recorder and Digitizer with Ordered Readout (LABRADOR) chips. The LABRADOR chip had a ninth channel in which a common 125 MHz clock was fed.

The LABRADOR chips contained a 260 element Switched Capacitor Array (SCA). The SCA continuously samples the signal at a rate of 2.6 GSamples/sec. When a trigger is issued, the sampling in one LABRADOR is frozen and read out into a digitizer. Each channel had four LABRADOR chips, which is a buffer depth of four, allowing re-triggering from the other three chips. If all four buffers are full, ANITA cannot re-trigger until one of the buffers becomes clear which takes 30 ms.

2.2.6 Triggering

The triggering system of ANITA-2 was an analog trigger designed to ride thermal noise levels. The trigger system used contained multiple levels and placed requirements on both the broadband signal and directional information. ANITA-2 only used VPol signals for triggering as radio emissions from a neutrino interaction were expected to be predominately vertically polarized if observed by ANITA.

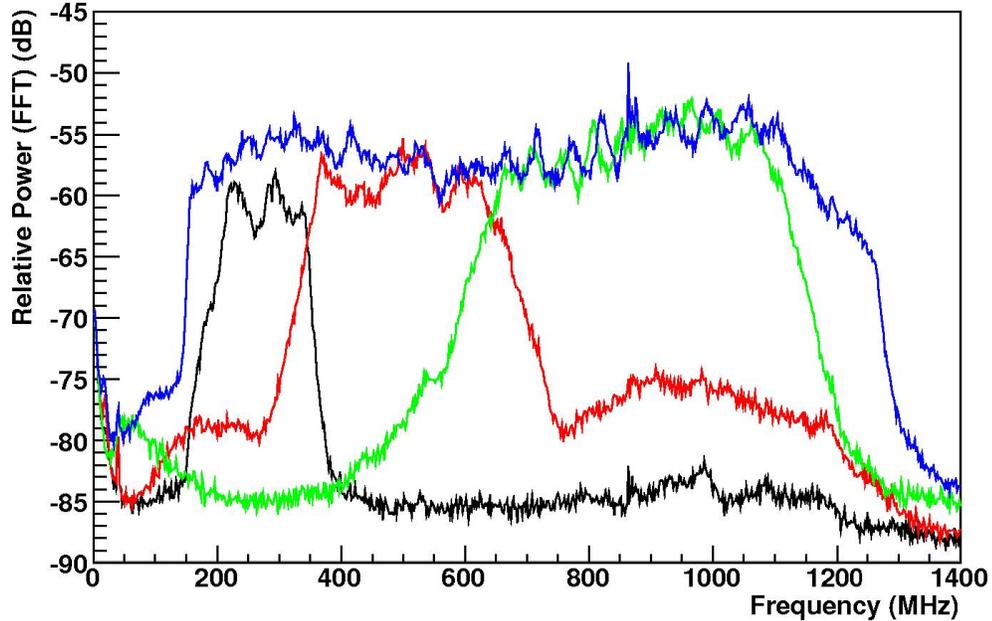


Figure 2.12: Banding of the trigger. Each color represents a different band. From [55].

2.2.6.1 SURF High Occupancy RF Trigger (SHORT) Boards

The VPol signal channel is split into two parts and one branch, the trigger branch, is fed into the SURF High Occupancy RF Trigger (SHORT) Boards. Here, the trigger signal was broken down into four bands, a full band and three sub-bands. This splitting was done to help combat triggering on man-made noise which tends to be continuous-wave (CW) which is narrow banded. Figure 2.12 shows the trigger banding on a representative channel. The 3dB points were 200-350 MHz for the low band, 350-700MHz or the mid band, 600-1150 MHz for the high band, and 150-1240 MHz for the full band. After each channel is split into four frequency bands, the signals enter a tunnel diode which acts as a square-law detector and outputs a unipolar pulse that tracked the power passing through the diode.

2.2.6.2 Level 1 Trigger

The output from the diode is filtered and sent to a field-programmable gate array (FPGA) on the one of the ten SURF boards. Inside the FPGA, signals passed through a discriminator with a programmable threshold. This programmable threshold allowed each band to ride thermal noise levels, monitoring their own rates. The thresholds were set by software which routinely compared user-defined target trigger rates to the actual trigger rates in the band and adjusted threshold accordingly. Typical rates in the top and middle rings of antennas were 14 MHz, 8 MHz, 8 MHz, 1 MHz for the low, mid, high, and full bands respectively. For the nadir ring, the single-band rates were 200 kHz, 200 kHz, 200 kHz, 40 kHz for the low, mid, high, and full bands.

A level one trigger (L1) can be thought of as an antenna trigger. It requires that 2 of out 3 sub-bands and the full-band exceed the threshold within a 10ns window. This requirement means the triggering signal must have some broadband aspect. L1 triggers occur in the top and middle rings at ~ 200 kHz per antenna during flight. L1 triggers in the nadir ring were at 800Hz due to a less stringent requirement for L2 and L3 triggers (see [2.2.6.3](#) and [2.2.7](#)).

2.2.6.3 Level 2 Trigger

L2 triggers were based on the ring geometry of the ANITA payload. The antennas for ANITA-2 were divided into 3 rings: a top, middle, and nadir ring. The top ring consisted of 16 antennas separated into 2 layers, the middle layer had 16 antennas in a single layer, while the nadir ring had only 8 drop-down antennas. When the SURF detects a L1 trigger, it would report this immediately to the Trigger Unit for Radio Frequency (TURF). This TURF handles the logic for both L2 and L3 triggers and

is located on the backplane of the cPCI crate. The TURF would determine whether the conditions were met for a L2 trigger.

For the top and middle rings, two out of three neighboring antennas had to have a L1 trigger within 20 ns of each other to pass the L2 requirement. An L2 trigger occurs at ~ 3 kHz per pair of antennas in these two rings. Since the nadir ring was more sparsely populated, there was no coincidence for a L2 trigger. Instead, a channel in this ring that passed L1 also passed the L2 trigger. This is why the L1 trigger rate in the nadir rings was much smaller than the other two rings. The total rate of L2 triggers for the entire payload was 50kHz.

2.2.7 Level 3 Trigger

The L3 logic was based on coincidence of the L2 trigger. When there was an L2 trigger in the same phi-sector on two of the three rings, the global L3 trigger was passed. These global triggers occurred at a rate of ~ 10 Hz although higher rates were observed during the flight.

2.2.8 Phi Masking

During the ANITA-1 flight, there were significant periods of flight time when the payload was experiencing large amount of CW signals due to human bases. These CW signals would overwhelm our trigger and we were not able to trigger near thermal noise, reducing our livetime. To help combat this, ANITA-2 was given software to mask out phi sectors where the trigger rate was too high. This phi masking would monitor the L3 trigger rate and if a phi-sector's trigger rate remained high (above ~ 20 Hz) for too long (~ 20 seconds) the phi sector would be masked. The phi mask would turn off global triggers from that phi sector but the phi mask for each phi-sector was independent, allowing other phi-sectors to still trigger. The phi mask

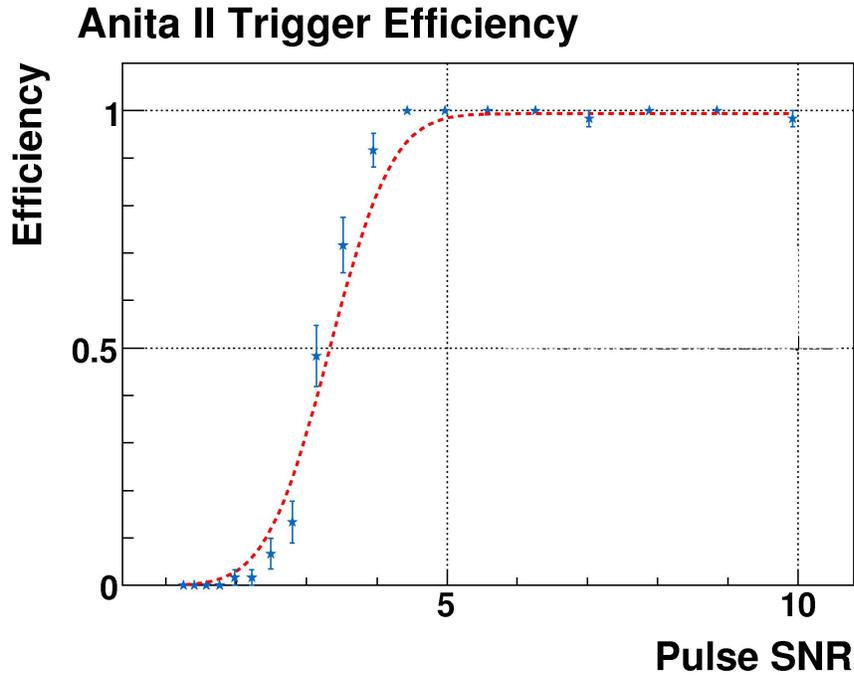


Figure 2.13: The trigger efficiency of L3 (global) triggers. Measured before flight as a function of signal to noise of input impulse. From [55].

would turn off when the trigger rate in that phi-sector reduced to a more acceptable rate.

2.2.9 Trigger Efficiency

Prior to flight, the efficiency of the trigger was tested by injecting impulsive signals with varying strength into the signal chain with the fraction of impulses causing a trigger recorded. Figure 2.13 shows the trigger efficiency as a function of SNR with all antenna channels contributing to the trigger, as would be the case during flight. In this configuration, the trigger is 50 % efficient at an SNR of 3.5 and approaches 100 % efficiency at an SNR of 5.

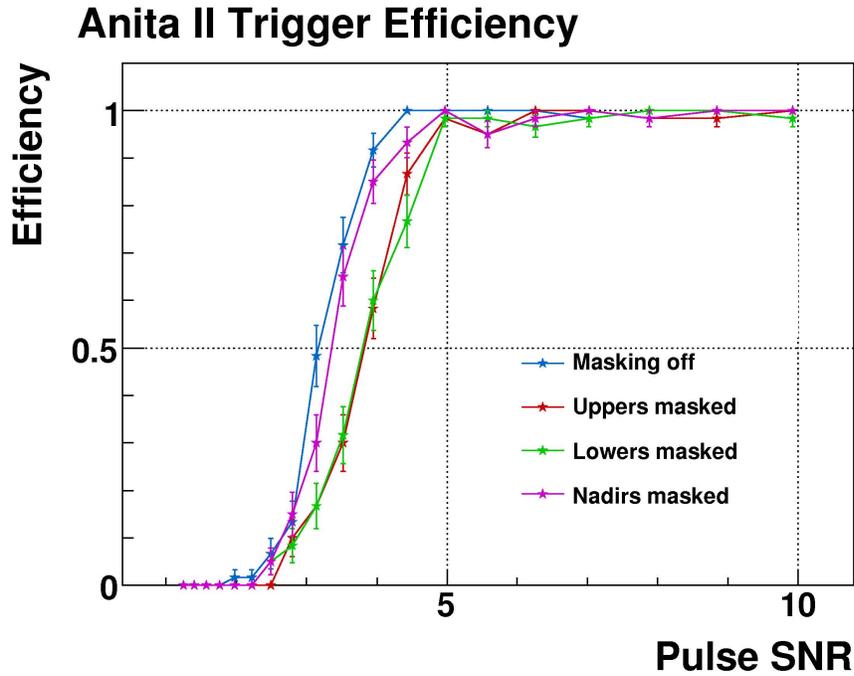


Figure 2.14: The trigger efficiency of L3 (global) triggers. Efficiency was measured as different rings were turned off. From [55].

The effect of removing a ring of antennas from contributing to a trigger was also investigated with results shown in Figure 2.14. The nadir ring contributes less to the trigger efficiency due to the lower number of antennas and the less strict logic for an L2 nadir trigger. This trigger is still an improvement over the ANITA-1 trigger, which had a 50% trigger efficiency at an SNR of 6 [57].

2.3 Flight Software and Data Storage

The ANITA-2 flight computer was a cPCI single-board computer, housed in the main instrument box. The flight software was split into a few autonomous programs, following the hierarchy shown in Figure 2.15.

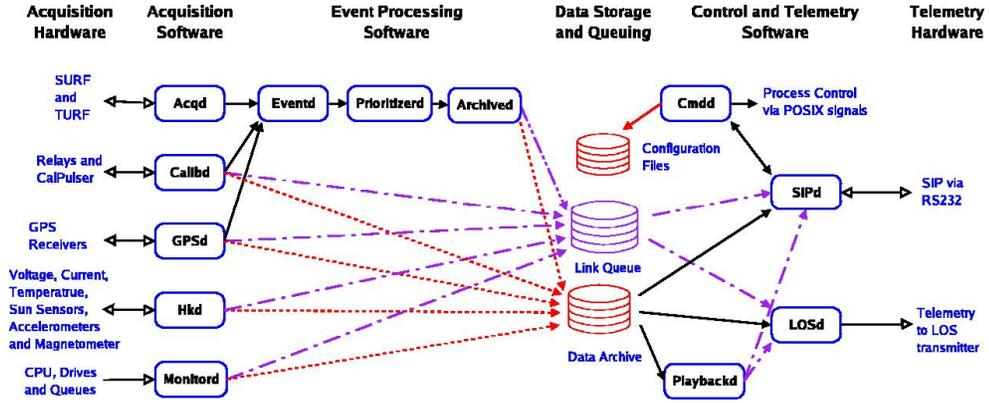


Figure 2.15: A schematic for flight software. From [56].

The RF data received by ANITA-2 was processed by the acquisition daemon (Acqcd). The Acqcd was tasked with acquiring the waveform and housekeeping data from the SURFS and triggering info from the TURFS as well as dynamically updating thresholds and phi-masking. After being processed by Acqcd, the event data was written using the event daemon (Eventd), the prioritizer daemon (Prioritizerd), and the archive daemon (Archived). Prioritizerd also assigned a relative significance on a 1-9 integer scale, with 1 being the most neutrino like. Events with a sufficiently high priority were then queued and relayed to the ground via Line of Sight (LOS) or SIP. This was a safe guard against the event that if the payload was unrecoverable, we would have the most valuable data.

The calibration daemon (calibd) controlled the relays to the on-board calibration pulser. The GPS daemon (GPSd) controlled the readout from the three suites of on-board GPS units. The housekeeping daemon (Hkd) read housekeeping data such as voltages, currents, temperatures and sun-sensor data. The monitor daemon (Monitord) monitored the free space on the drives and the length of queues for writing data.

ANITA-2 had three forms of storage on board. This redundancy was implemented to ensure that even if two forms of storage was corrupted, there would still be usable data. There were two arrays of eight 1TB MTRON solid state drives along with spinning hard drive inside the "neobrick", a shock-mounted pressure vessel.

2.4 Flight Performance

ANITA-2 launched on December 21st, 2008 and was in flight for a total of 31 days with a total of 28.5 days of live-time. Figure 2.16 shows the flight path ANITA-2 took around the continent, including 2 passes over the deep quiet ice of East Antarctica. The flight was ended on January 21st, 2009 where the payload was cut from the balloon and crash landed near Siple Dome. ANITA-2 recorded ~ 27 M events, with ~ 21 M RF triggered events.

2.4.1 Ground Calibration Pulser

ANITA-2 used ground calibration pulsing systems were used during the flight to send pulses to the payload. These pulses were used to monitor the payload during flight, as well as calculating efficiencies and help inform analysis techniques. Two calibration pulsers were used during the flight, one located at Taylor Dome (77.88° S, 158.46° E) and the other at Williams Field.

The site at Taylor Dome consisted of a discone antenna buried into a 90m borehole in the ice. The antenna would transmit impulsive signals at a rate of 1Hz through the duration of the ANITA-2 flight. The impulse system was powered by PV cells and was set-up and tested by a team from UCLA. The system was autonomous through the flight and were seen on all three passes of Taylor Dome.

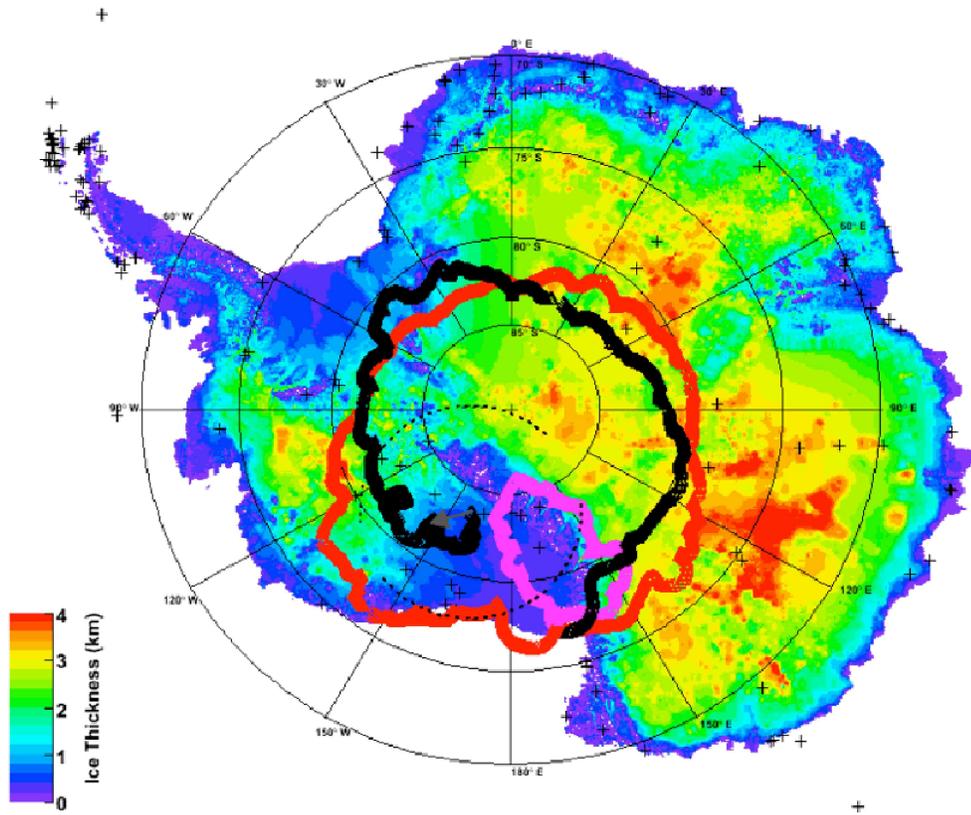


Figure 2.16: The flight path of ANITA-2. The payload looped around the south pole before performing a bigger loop of the continent. The different colored lines show different loops around the continent. From [55].

The Williams Field pulsing system also discone antenna buried in the ice, but only to a depth of 25m. This pulsing systems also had a Seavey antenna which allowed for different polarizations to be broadcast:vertical, horizontal, and 45°. While the Taylor Dome system required a thermal trigger to be seen by ANITA-2, the Williams Field setup was synchronized to ANITA-2's GPS controlled PPS triggering systems. This allowed the Williams Field pulses to occur during a PPS trigger, ensuring the pulse was recorded.

2.4.2 Performance of ANITA-2

ANITA-2 had a much more successful flight than ANITA-1, which suffered from large amounts of dead-time due to rebooting of the flight computer and high trigger rate due to human bases. Fractional dead-time and event rate over the duration of the flight can be seen in Figure 2.17. The total livetime of the flight was 28.5 days.

There were a few issues that occurred during the flight that reduced sensitivity. Shortly after launch, channel 2V failed, with the channel's LNA seemingly lacking power. Although power was restored intermittently, channel 2V was not used in any analysis of the flight data. Worries over the charge controller led to turning off the nadir antennas (by turning of the RFCMs) for part of the flight, $\sim 5\%$. During this time, the remaining 32 antennas were able to perform normally. GPS data was lost for $\sim 5\%$, but the GPS information was extracted from sun-sensor and CSBF GPS data, although that data is less accurate. Lastly, the MTRON solid state drives experienced a problem with writing to the disk, causing disagreement with header information. This occurred for approximately 30,000 events which were removed from the analysis.

ANITA-2 experienced a full recovery, with every critical piece of hardware recovered and sent off the ice. All three sets of data were undamaged during landing. Only

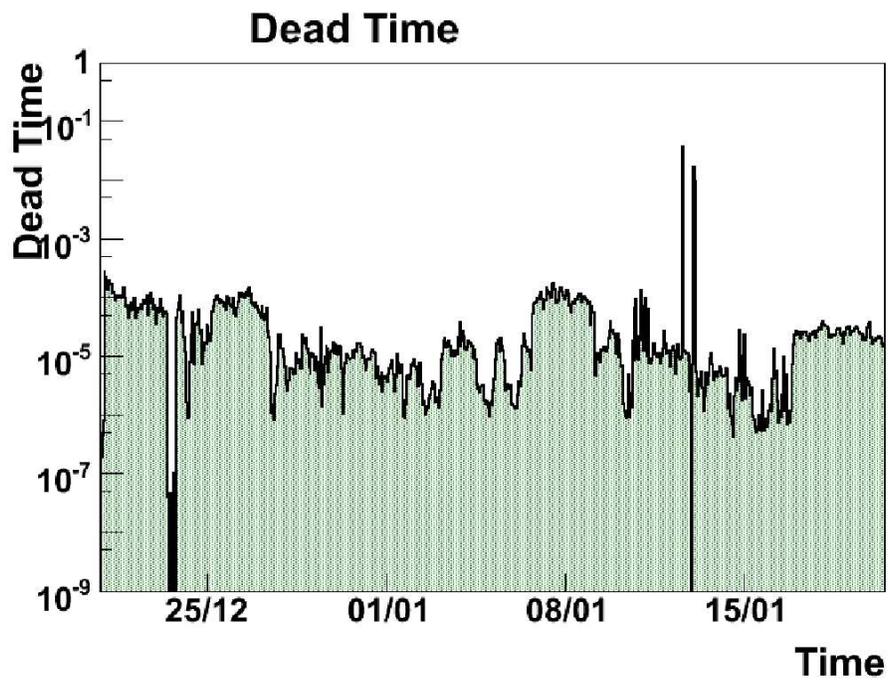


Figure 2.17: The deadtime fraction for ANITA-2 flight. X-axis is label as day/month from 2008-2009. From [55].

a few antennas and gondola parts were not able to be sent off the ice by the end of Austral Summer.

Chapter 3

Techniques for Filtering of Background Contamination

In this chapter, we will discuss an investigation of techniques for filtering waveforms in the frequency domain. Many of the events recorded during the flight have background contamination from anthropogenic sources. This background noise needs to be removed to find possible pulses in the data. Before going through the techniques, we will discuss the nature of the contamination and how filtering these waveforms will benefit us. We will then discuss notch filters and various techniques used to implement those filters, then we will discuss our approach, which is to notch filter while appropriately adjusting phases in the notch. We will also discuss a different technique for background removal called Sine Subtraction. Finally, we will discuss the results of our investigation.

3.1 Continuous Wave Contamination

As discussed before, we are looking for a radio frequency impulse that exists for a short period of time (on the order of nanoseconds). Unfortunately, radio frequency signals are commonly produced where humans are present. A significant background for radio Cherenkov experiments is therefore anthropogenic noise. These radio waves produced by humans are called continuous wave (CW) due to the constant amplitude

and frequency of the radio wave. These CW waves tend to have higher voltages at the payload than expected thermal noise. As our experiment is designed to trigger near thermal noise, these strong sources can mask potential impulsive events. Human activity can also produce impulsive signals but those backgrounds must be dealt with in a different way due to their similarity to the sought after neutrino impulse.

Filtering techniques were developed to help reduce the impact of CW noise on our neutrino search. These techniques were based on a notch filter. A notch filter works by removing power due to the CW background and allowing a glimpse at any underlying pulse present in the waveform. While not the only technique that could be used to accomplish this purpose, it was the technique that was easily modified by changing how the notch was filled and how we treated the phase inside the notch.

Notch filtering is based upon the fact that each CW signal uses a limited number of frequencies (carrier frequencies) for transmitting their signal. By reducing the power at the CW carrier frequencies, we can reduce the CW signal from the original waveform and see what power is left in the other frequencies. The investigation performed was to modify the notch filtering technique in various ways and test their impact on reconstructing signals and rejecting background events. The technique was modified by changing how the notch was filled in. The goal was to reduce the impact of the CW in the waveform in power. The primary metric we used for testing was reconstruction of known calibration pulser events.

3.2 Adjusting the Power Spectrum for Notch Filters

The first step in imposing a notch filter is identifying what frequency bins have the signature we wish to remove, which in our case is CW. To do this, we must first create a baseline spectrum for each antenna that can be used for comparison to identify CW. After these baselines are created for each antenna, we group antennas

together in different groupings and create baseline spectra for these groupings. We then compare the magnitude of an event’s amplitude spectrum for each unique group to the baseline for that group and any frequency bins that have a magnitude higher than 4 dB above the baseline spectral magnitude are flagged as part of the notch.

The 4 dB threshold was chosen based on the variation of our impulse response function (Figure 3.1). The impulse response’s maximum deviation away from the mean was just under 4 dB, so any impulse we expect to see may also have that same range. As the purpose of this threshold is to remove CW power, only events with power in excess of 4 dB above average will be flagged.

To create the baseline, we must find a sample of events that would not contain possible neutrino signals for which we are searching or CW. From the data set, we were able to identify events that were downward-going (i.e., not coming from ice). Since these events are not originating from the ice, there should be no signal from the Askaryan effect or CW from bases. This sample would give us an understanding of each antenna individually, in terms of expected thermal noise at that antenna. We called this sample the ‘thermal’ sample, as it is used to help create a baseline representing the expectation due to thermal noise, called a thermal baseline.

For each antenna, we wish to create a thermal baseline for comparison. To accomplish this, we took the waveform measured from an antenna for each event in the sample and performed a Fourier transform. We then added the amplitude spectra and divided by the number of events used. After performing that average, we then smoothed the spectrum through use of a box average. This smoothing uses seven bins, but for bins near the start and end of the waveform we reduce the number of bins used in the average down to three or five bins due to end of the waveform. This final smoothed spectrum serves as a baseline against which we compare signals received in that antenna.

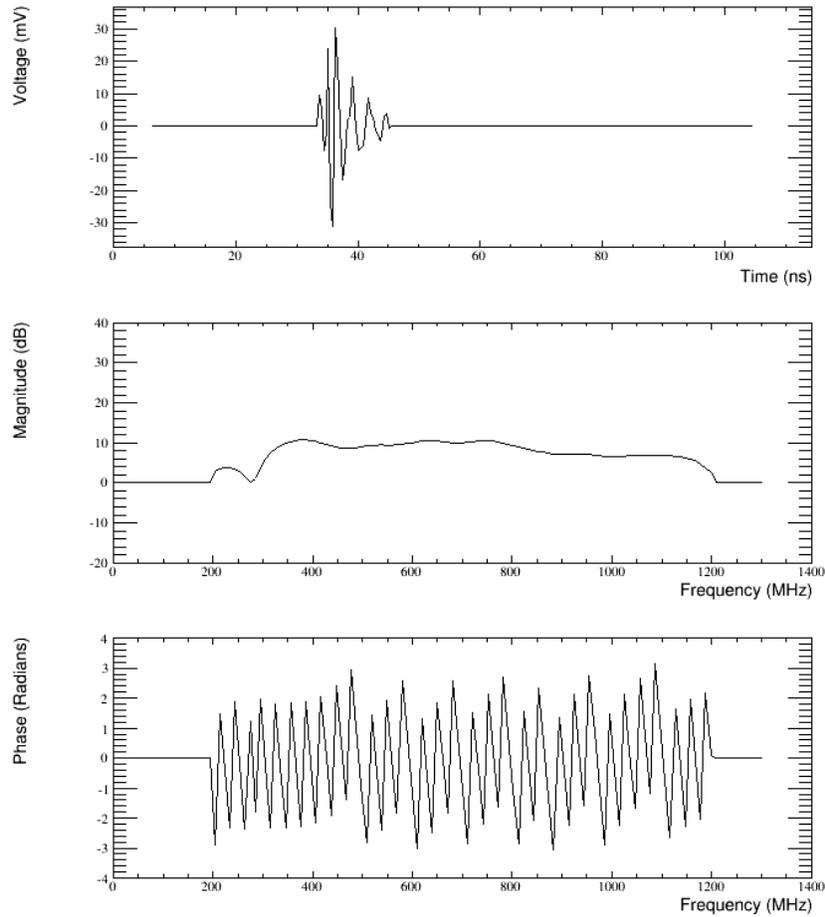


Figure 3.1: The impulse response for ANITA-2. Top Plot: Time domain. Middle plot: Amplitude spectrum. Bottom plot: Phase spectrum. The amplitude spectrum varies by less than 4 dB from its average. The phase spectrum is linear, but due to wrapping appears to jump very often.

In an individual event, there are forty antennas, and any may need CW removed. This can be done on an antenna by antenna basis, but there are concerns about thermal noise. Thermal noise in a single antenna can vary widely from frequency bin to frequency bin. A strong fluctuation would cause a frequency bin to be flagged for notching. One way to combat this is to add many antennas together. Thermal noise would add incoherently while CW noise would add coherently. Therefore, thermal noise voltages would grow as \sqrt{N} while CW voltages would grow as N , where N is the number of antennas used.

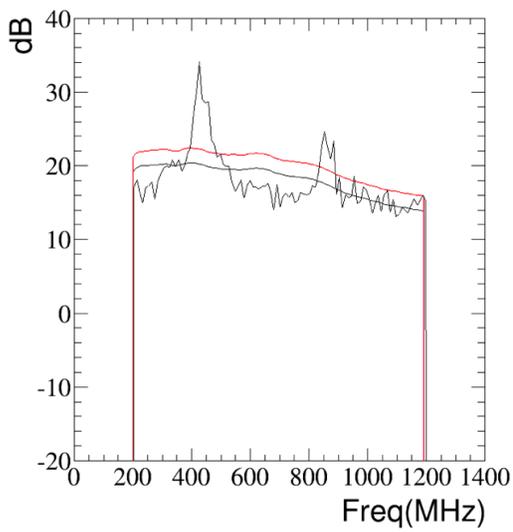
In the past analysis, analysis 1.0 [55], all 40 antennas were added together for filtering. We chose to use smaller groups because CW signal does not hit all 40 antennas equally. The CW will be strong on one side, or just a few antennas, and by adding all 40 antennas together we can weaken the CW peak due to some antennas not receiving that CW signal. Choosing which antennas to add together is a simple process. Nine antennas were chosen due to their representing approximately three phi sectors. It is not exactly three phi sectors due to the nadir ring of antennas on ANITA-2. The nadir ring of antennas did not have an antenna for every phi sector, instead having one antenna for every other phi sector. We can create a group of antennas by asking which nine antennas are closest to a given azimuthal direction. By looping through all possible azimuthal directions, we can see how many unique groupings of antennas there are, unique meaning no two groups share all nine antennas. In this way, we found that there were 22 unique groupings of antennas.

The method we used to reduce thermal fluctuations was to add the amplitude spectra of the nine antenna waveforms together. The amplitude spectra of the waveforms from the individual antennas were averaged together using the root-mean-square (RMS) in each frequency bin to produce an average spectrum. The thermal

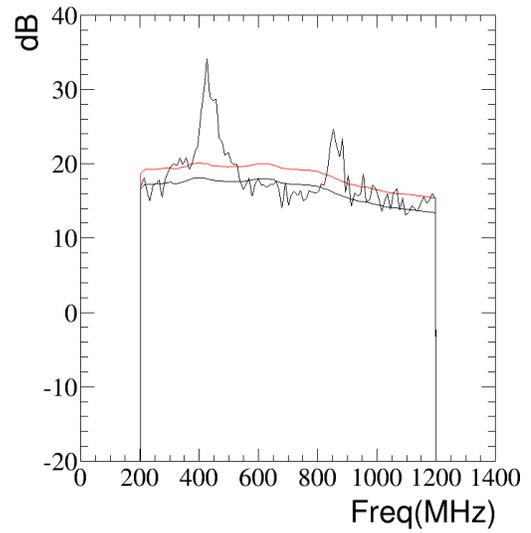
baselines for these antennas, as described above, were also averaged together in the same way for comparison to the event spectra.

Before we can use the thermal baseline for comparison to decide which frequency bins to flag for notching, the thermal baseline needs to have some corrections applied. These corrections are made for two reasons. The first reason is that the average power in a waveform seen at the payload is not consistent between events. This could be due to higher thermal noise or more background noise when that event was recorded. For a fair comparison, we want to be sure that the thermal baseline we created has the same average power as the average power of thermal noise for that waveform.

Finding the average power of the thermal noise of the waveform requires a recursive procedure that also involves flagging the spectrum for notches. This is due to not knowing in advance which bins contain extra power due to CW contamination. Thus, this procedure is done at the same time as flagging frequency bins for notching in a recursive manner. We first average the power of all frequency bins of the waveform and scale the thermal baseline so that it has that same average value. We then flag any frequency bins above the threshold (4 dB) for all antennas in the group. Once a frequency bin has been flagged, we extend away from that peak until the amplitude becomes 3dB above the thermal baseline, flagging all frequency bins along the way. We then recompute the average power of the waveform, excluding any bins we have flagged, which lowers the thermal baseline. We then use that new baseline to reassess which frequency bins are flagged as being above threshold. This recursive procedure, averaging and flagging, is done until the average of the signal moves less than 0.01 dB from the previous iteration. Typically, this is achieved in less than 5 iterations of the procedure. This recursive procedure allows us to place the thermal baseline at the right level for each event and to flag frequency bins contaminated by CW. The difference in baseline placement can be seen in Figure 3.2.



(a) The placement of the thermal baseline for this event before iterative process. Black line is thermal baseline, red is base +2dB.



(b) The placement of the thermal baseline for this event after iterative process. Black line is thermal baseline, red is base +2dB.

Figure 3.2: The left plot shows the placement of the thermal baseline before accounting for peaks, while the right plot shows placement after we account for peaks.

We perform this flagging procedure for all unique groupings of antennas, keeping track of which frequencies are being flagged for each antenna individually. After performing the procedure on all the groups, each antenna will have a list of frequencies to be filtered from its individual waveform. We will perform a notch filter on each individual antenna, that is, we set the amplitude spectrum to zero in the bins that have been flagged as CW contaminated but leave the non-flagged bins untouched.

We investigate the effects of further treatment of the bins where the amplitude spectrum has been removed, detailed in the following sections. In Figure 3.3, we have a Taylor Dome event that is used for comparison plots for each method and Figure 3.4 shows the coherently summed waveform for the same event.

3.2.1 No Fill

The easiest method to understand, which was used in analysis 1.0, is a no-fill method. As the name suggests, we will notch the places in each individual antenna where CW was flagged by setting the amplitude in those bins to 0.

This method has two advantages. We do not have to any extra power put in by the filter that the signal would not have. By having those bins set to zero, we also do not have to worry about the impact of the CW-contaminated phases.

The disadvantages for this method are simple. By not filling in the notch, we are reducing the power in the waveform. Since we are targetting specific frequencies, we may see residual CW-like features in the time domain waveform induced by the notch due to the properties of Fourier analysis, i.e. ringing in the time domain.

We can see how this method affects the example event we selected in Figure 3.5.

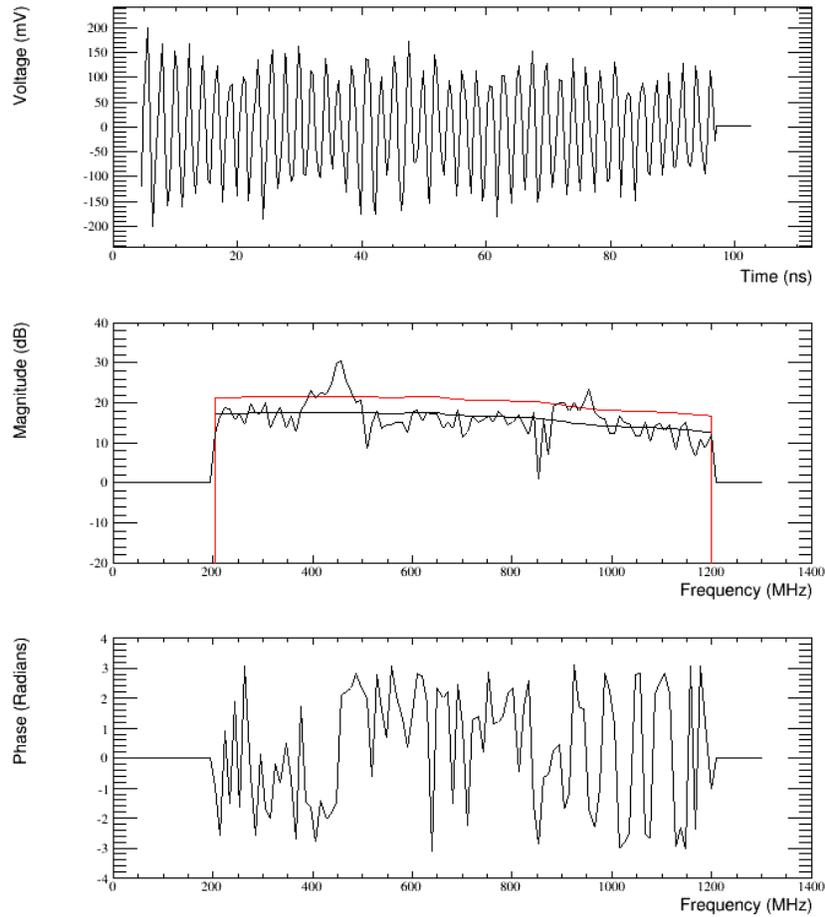


Figure 3.3: Waveform for antenna 12 from event 687109, a Taylor Dome event. It has two CW peaks at 436 and 924 MHz that must be filtered. Top: Time Domain waveform. Middle: Amplitude Spectrum of this waveform. The solid black line is the thermal baseline, the red line is thermal baseline + 4dB. Bottom: Phase Spectrum

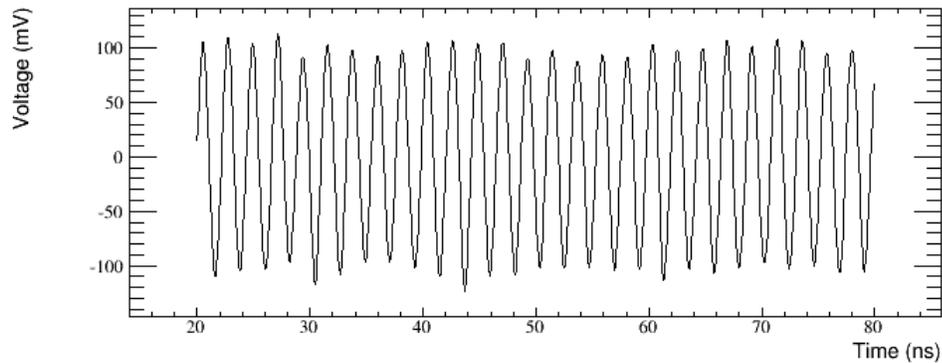


Figure 3.4: Coherently summed waveform for event 687109. Notice the sine-wave like structure that is indicative of CW.

3.3 Filling in Notched Bins

Rather than leaving the notch frequencies with a zero spectral amplitude in some frequency bins, we wish to approximate what the spectrum would have looked like if there were no CW present. How to approximate what would have been present is a matter of discussion, described in sections [3.3.1](#), [3.3.2](#), and [3.3.3](#)

3.3.1 Rayleigh

While we know thermal noise will be present and help set a baseline average, we do not know whether there is a signal present in an event. As a conservative method, we can fill in the notched frequencies with just thermal noise. While this may degrade any neutrino signal that could possibly be present, we will not be falsely adding signal to events in which there is none.

Thermal noise spectral amplitudes follows a Rayleigh Distribution (where f is the probability density function, σ is a scale parameter related to the mean, x is spectral

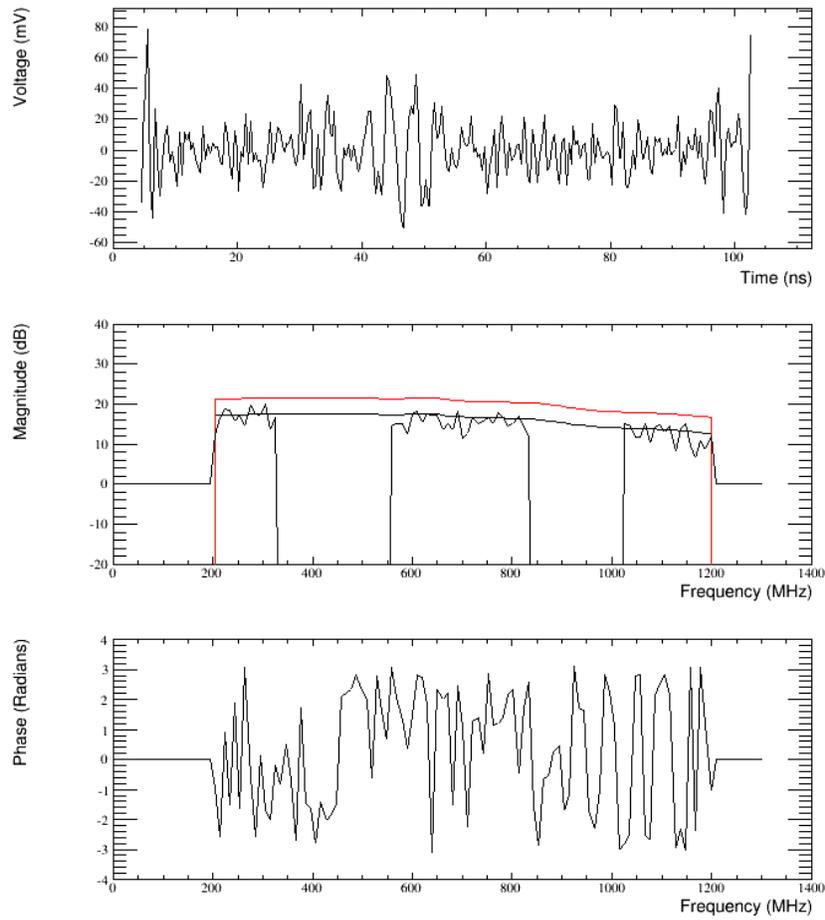


Figure 3.5: Same antenna for event 687109. CW peaks (seen in Figure 3.3 have been removed by notches. Top: Time Domain. Middle: Amplitude Spectrum, Black line is thermal baseline, Red line is thermal +4dB. Bottom: Phase Spectrum

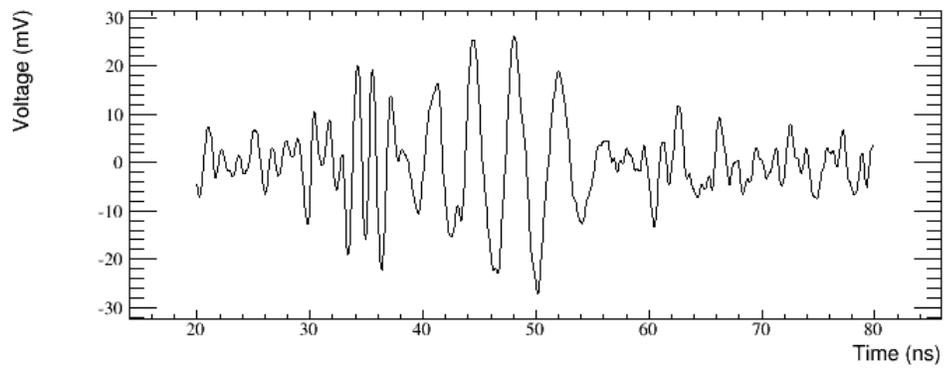


Figure 3.6: Coherently summed waveform for event 687109.

amplitude):

$$f(x, \sigma) = \frac{x}{\sigma^2} e^{-\frac{x^2}{2\sigma^2}} \quad x \geq 0 \quad (3.1)$$

We can use this distribution to produce a spectral amplitude in each frequency bin. To do this, we will need to find σ for use in each bin. This is quite simple, as for a Rayleigh distribution σ is easily related to the mean of the distribution:

$$\text{mean} = \sigma \sqrt{\frac{\pi}{2}} \quad (3.2)$$

The thermal baseline we have created for each individual antenna can be used as the mean of the Rayleigh distribution for each bin, allowing us to find σ for each frequency bin.

To select the spectral amplitude, we will use the Inverse Transform method. For this method we need to use the cumulative distribution function, which is the integral of the probability distribution function.

$$F(x) = 1 - e^{-\frac{x^2}{2\sigma^2}} \quad x \geq 0 \quad (3.3)$$

$$x = \sqrt{-2\sigma^2 \ln(1 - F(x))} \quad 0 \leq F(x) \leq 1 \quad (3.4)$$

Using Equation 3.4, we can find a new amplitude for each frequency bin flagged by the notch filter. We merely need to draw a random number from zero to one and set it equal to $F(x)$ and use the value for σ we calculated from equation 3.2.

The advantage of this technique is that we are approximating what the bin could have looked like without the CW contamination, using a Rayleigh distribution. This should give a reasonable waveform that we could have seen.

There are also some disadvantages. We are filling the bins in with pure thermal noise. We cannot approximate signal in those bins for fear of injecting a signal that is not there. Secondly, these magnitudes are created using randomly generated numbers which can be hard to manage, as a large fluctuation in a single bin can cause an event to pass or fail analysis cuts. Finally, this filter is acausal, so the power from the notched bins was spread through the waveform, including earlier times which is physically impossible.

An example of this filter can be seen in Figure 3.7 and the coherently summed waveform seen in Figure 3.8.

3.3.2 Interpolated Filter

While the Rayleigh method of filling in the notched bins is simple, there is yet a simpler way of filling in the notched bins. We could simply interpolate the spectral amplitude from one side of the notch to the other. Since we expect thermal noise to be roughly equal in magnitude for all frequency bins (discounting antenna gain), an interpolation will give each bin a correct level for the power for the bin without any worry of the impact of fluctuations.

There are various types of interpolation, but for simplicity, we decided to use a linear interpolation. Knowing the amplitude on either side of the notch, we can derive a linear equation using a slope and intercept based on those two points. We then use this equation to fill in the notch with the calculated values.

The advantage of this technique include not having to worry about random numbers or fluctuations. The event will have the same magnitude in the specified bin every time we run our analysis code.

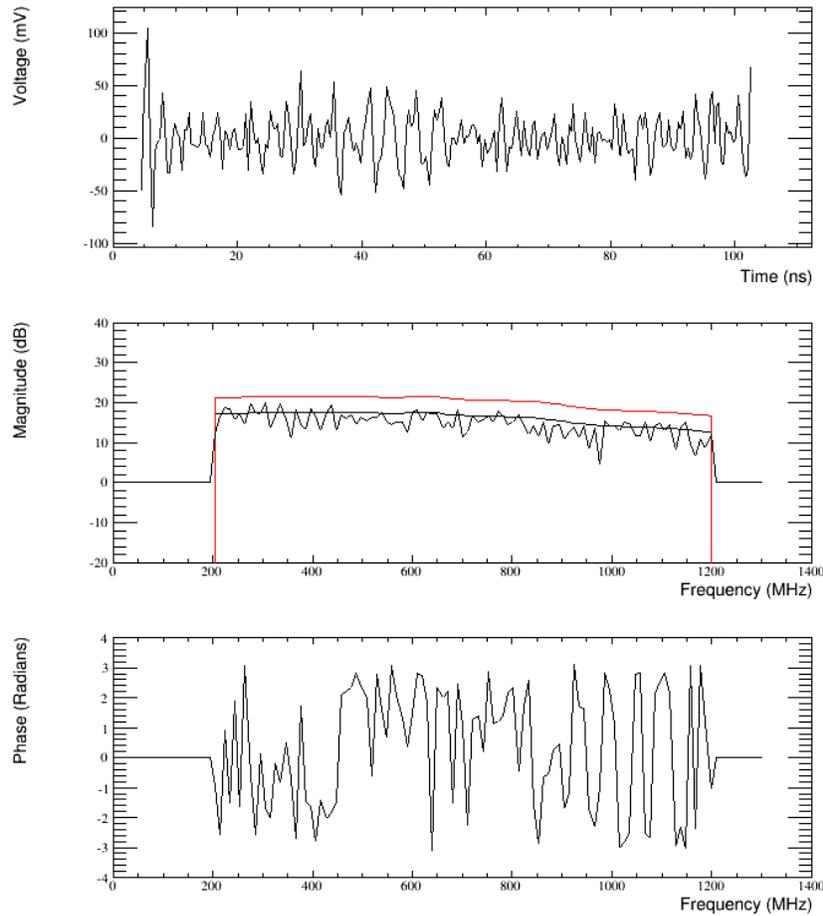


Figure 3.7: Single antenna waveform for event 687109. CW peaks (seen in Figure 3.3) have been removed. The notches have been filled in using the Rayleigh method. It is now hard to find the location where the CW was. The problem with this method is the noise that is used to fill in the notch is randomly generated, making it difficult to fill the notch the same way twice. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red line is thermal baseline +4dB. Bottom: Phase Spectrum

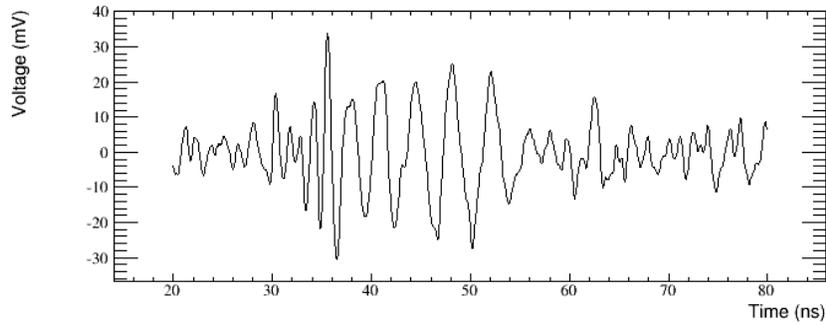


Figure 3.8: Coherently summed waveform for event 687109 after Rayleigh method. An impulse like structure can now be seen.

To our disadvantage, this technique is not based on 'reality', as a straight line would never be observed in a frequency spectrum. This may produce unintentional effects on the waveform. As with the Rayleigh filter, this interpolated filter is acausal.

An example of this filter can be seen in Figure 3.9 with the coherently summed waveform seen in Figure 3.10.

3.3.3 Wiener (Optimal) Filter

A common technique that is used to remove a corrupted signal and produce the original signal is called the optimal filter. Originally derived by Norbert Wiener, it makes use of a filter that can produce a new signal as close as possible to the original. The implementation we are using is described in [58]. Suppose we have a measured signal:

$$c(t) = s(t) + n(t) \quad (3.5)$$

$$C(f) = S(f) + N(f) \quad (3.6)$$

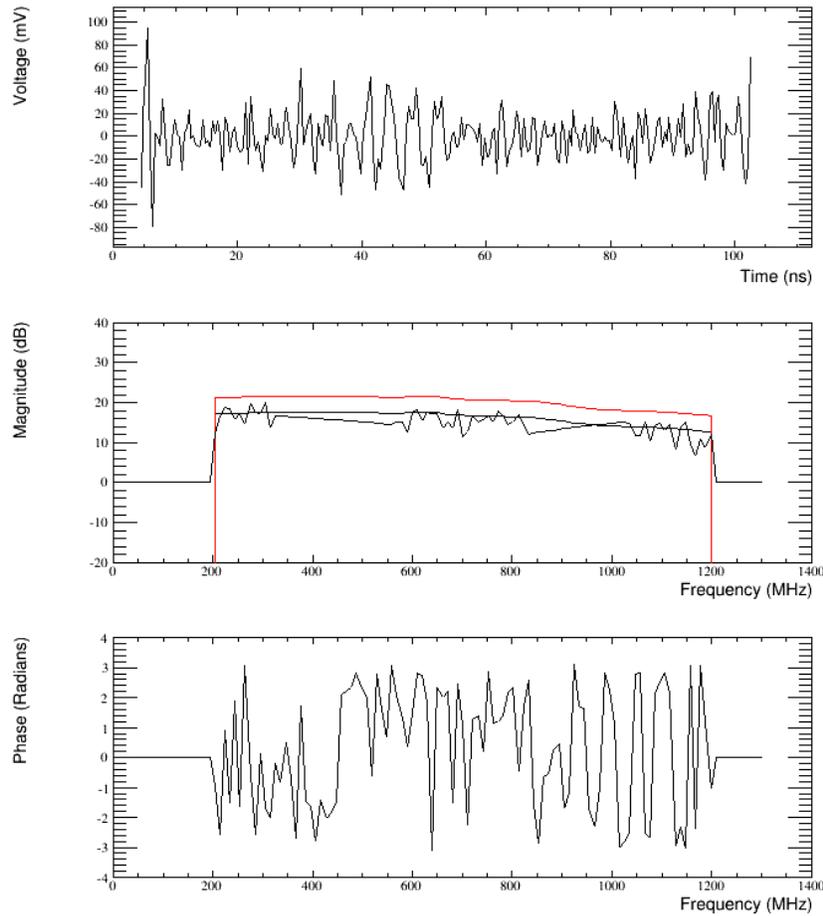


Figure 3.9: Single antenna waveform for event 687109. CW peaks (seen in Figure 3.3) have been removed. The notches have been filled in using the interpolated method. It is easy to see where notches occurred in the amplitude spectrum, but the average power should be consistent with thermal noise. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red is thermal baseline + 4dB. Bottom: Phase Spectrum

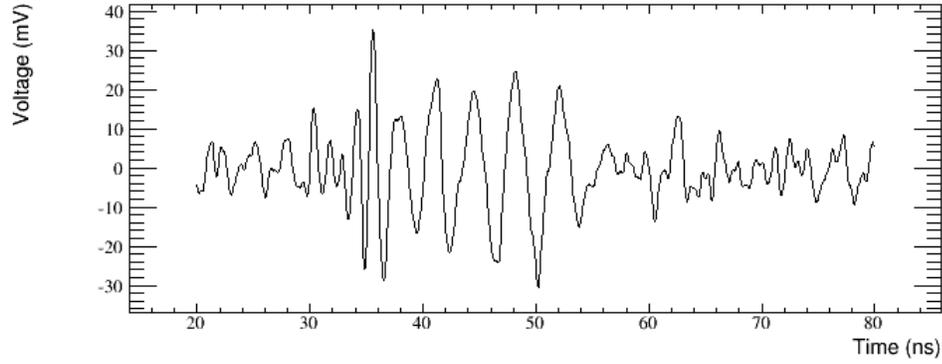


Figure 3.10: Coherently summed waveform for event 687109 after Interpolation Method filtering.

where $c(t)$ is the measured (corrupted) signal, $s(t)$ is the ‘true’ signal (without noise), $n(t)$ is the noise component, and $C(f)$, $S(f)$ and $N(f)$ are the Fourier transforms, respectively.

We wish to produce a filter $\Phi(f)$ such that when applied, we produce $\tilde{S}(f)$

$$\tilde{S}(f) = C(f)\Phi(f) \quad (3.7)$$

which is close to $S(f)$. We minimize by using least squares method, where

$$I = \int_{-\infty}^{\infty} |\tilde{S}(f) - S(f)|^2 df \text{ is minimized.} \quad (3.8)$$

Plugging equations 3.6 and 3.7 into 3.8, we can see

$$\begin{aligned} I &= \int_{-\infty}^{\infty} \left| [S(f) + N(f)]\Phi(f) - S(f) \right|^2 df \\ &= \int_{-\infty}^{\infty} \left| S(f)[\Phi(f) - 1] + N(f)\Phi(f) \right|^2 df \end{aligned} \quad (3.9)$$

If $S(f)$ and $N(f)$ are uncorrelated, which in our case they are, then the product of the cross terms is zero when integrated over frequency, which leaves us with

$$= \int_{-\infty}^{\infty} |S(f)|^2 |1 - \Phi(f)|^2 + |N(f)|^2 |\Phi(f)|^2 df \quad (3.10)$$

We must minimize the integrand with respect to $\Phi(f)$ for every value of f . Differentiating the integrand with respect to $\Phi(f)$, and solving for $\Phi(f)$ we get

$$\Phi(f) = \frac{|S(f)|^2}{|S(f)|^2 + |N(f)|^2} \quad (3.11)$$

This equation gives us a filter, that when applied to the measured waveform, gives an approximate waveform that is close to the signal that would have been measured if no noise is present! We only need to have an estimate of signal, $S(f)$, and of the noise, $N(f)$. The latter is where we encounter a problem. Knowing the exact nature of the noise is not trivial as we only measure a corrupted signal. But since the signal we expect and the noise are uncorrelated, we can use

$$|S(f)|^2 + |N(f)|^2 \approx |C(f)|^2 \quad (3.12)$$

To use this in filtering is quite easy. We need only an idea of what the signal looks like in the notched region and what we measured. The signal in the region will be approximated by the thermal baseline, as used in the past filters. The value we use to fill in the bin is calculated as

$$\tilde{S}(f) = C(f)\Phi(f) = C(f) \frac{|S(f)|^2}{|C(f)|^2} \quad (3.13)$$

An example of this filter can be seen in Figure 3.11 and the coherently summed waveform can be seen in Figure 3.12.

The advantages of this technique are the well documented usage and the rigorous mathematical proof. This filter can also be made causal with some care, which was not done in this work. The disadvantage of the technique is in using an approximate $S(f)$, which may or may not be good enough for this purpose. When the CW is much higher than the expected signal, the filtered amplitude spectrum is reminiscent of the no-fill method. Additionally, if the filter is applied to a region where the measured signal is less than the expected signal, this filter boosts those frequency bins, causing CW-like peaks in the amplitude spectra. This feature is seen in Figure 3.11 around 850 MHz in the amplitude spectrum.

3.3.4 Comparison of techniques

To compare these different types of filters we used a sample of events, calibration pulsar events from Taylor Dome. These events are useful because we know they should have an impulse-like waveform, and the known location of the source helps us see how the filters affect reconstruction. Also, many Taylor Dome events have CW contamination due to interceding known human bases that produce radio noise.

We judge filter performance on a few criteria. The first criterion is the efficiency of the past analysis cuts with the new filters applied. We want a filter than can remove CW but keep as many signal-like events as possible. The second criterion is how well the events point back to the known location of the pulsar source. This can be judged in payload coordinates, azimuthal and polar coordinates θ and ϕ relative to axes on ANITA. The last criterion is how the SNR of the events is distributed. An impulse with low noise will have a high SNR, so we are hoping to recover some of that using

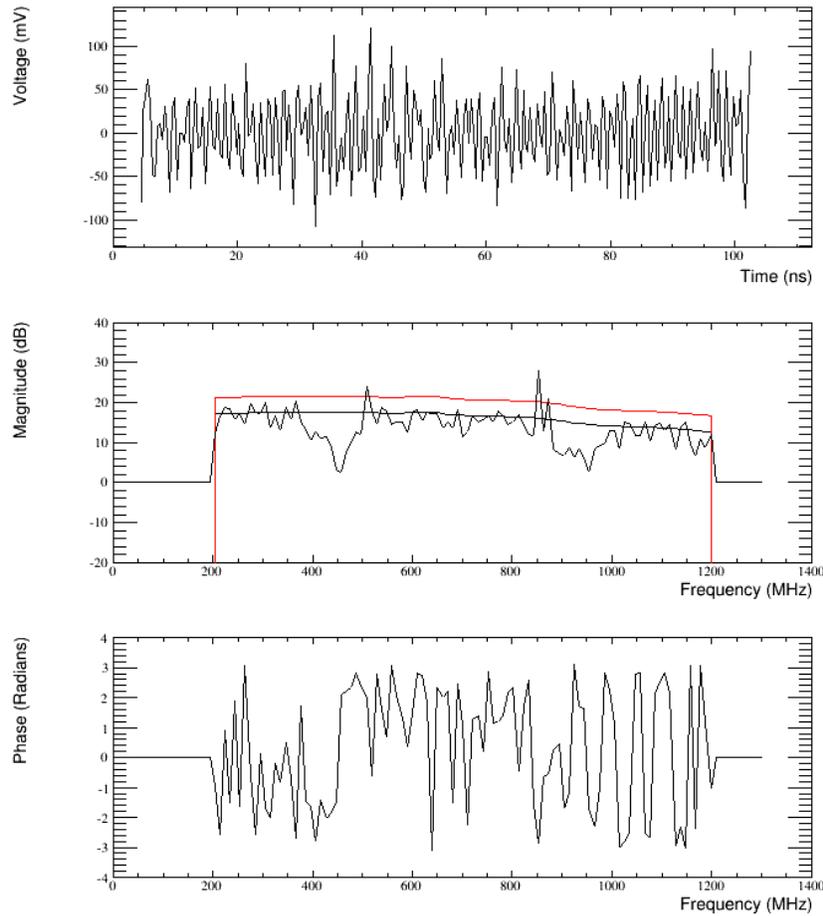


Figure 3.11: Antenna 12 waveform for event 687109. CW peaks (seen in Figure 3.3) have been removed. The CW has been removed by use of the Wiener filter. The location of the CW is easily seen as the Wiener filter reduces the signal in those regions drastically. Additionally, this filter can produce CW-like peaks after filtering as seen at ~ 850 MHz. Top: Time Domain. Middle: Amplitude Spectra. Black line is thermal baseline, red line is thermal baseline + 4dB. Bottom: Phase Spectrum.

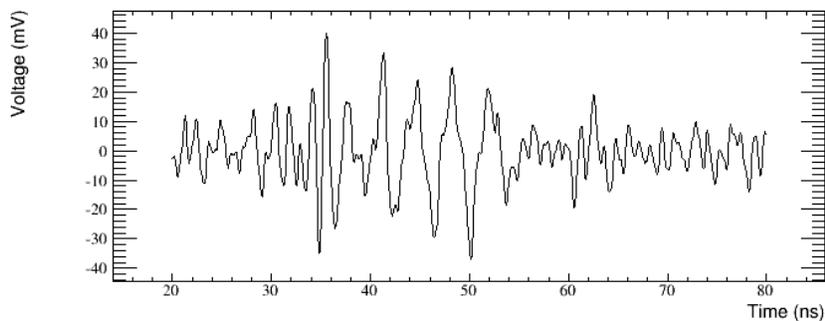


Figure 3.12: Coherently summed waveform for event 687109 after Wiener filter.

Amplitude Filter	Underflows (θ/ϕ)	Overflows (θ/ϕ)
No-Fill	456 / 154	2 / 61
Rayleigh	562 / 178	3 / 60
Wiener	953 / 369	39 / 51
Interpolated	671 / 193,	3 / 55

Table 3.1: Number of over and underflows for θ and ϕ reconstruction.

these filters. While only an optimized analysis with all cuts can determine the best filter, we can use these metrics to pick out a couple of the best candidate filters.

For comparison purposes, we will have a non-filtered sample plotted against the same sample run through the discussed samples.

From 3.13, we can see that the interpolated and Rayleigh method both have better efficiency at passing analysis cuts at various SNR than the wiener and no-fill filters. In Figure 3.15 and Figure 3.14, we can see that the Rayleigh and Interpolated Methods have the best reconstruction in both θ and ϕ coordinates, with Table 3.1 showing the under/over flows for the distributions, and are roughly equal. The distribution in SNR (Figure 3.16) confirms that interpolated and Rayleigh are nearly equivalent.

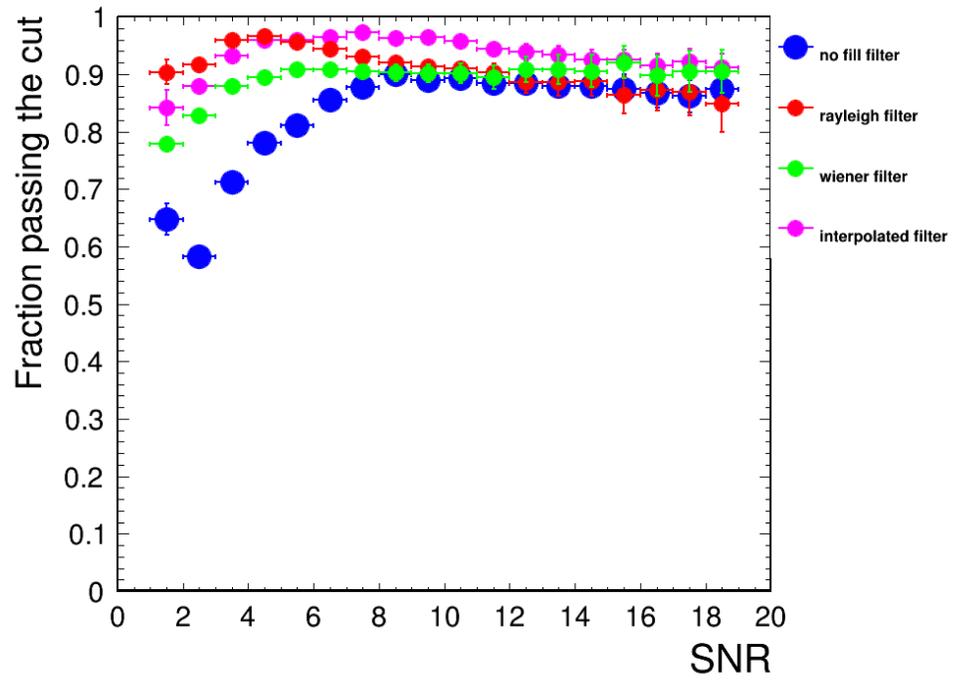


Figure 3.13: The efficiency of passing events for different amplitude filters. Both Rayleigh and interpolated filters seem to be the best at lower SNR, with interpolated remaining highly efficient at high SNR.

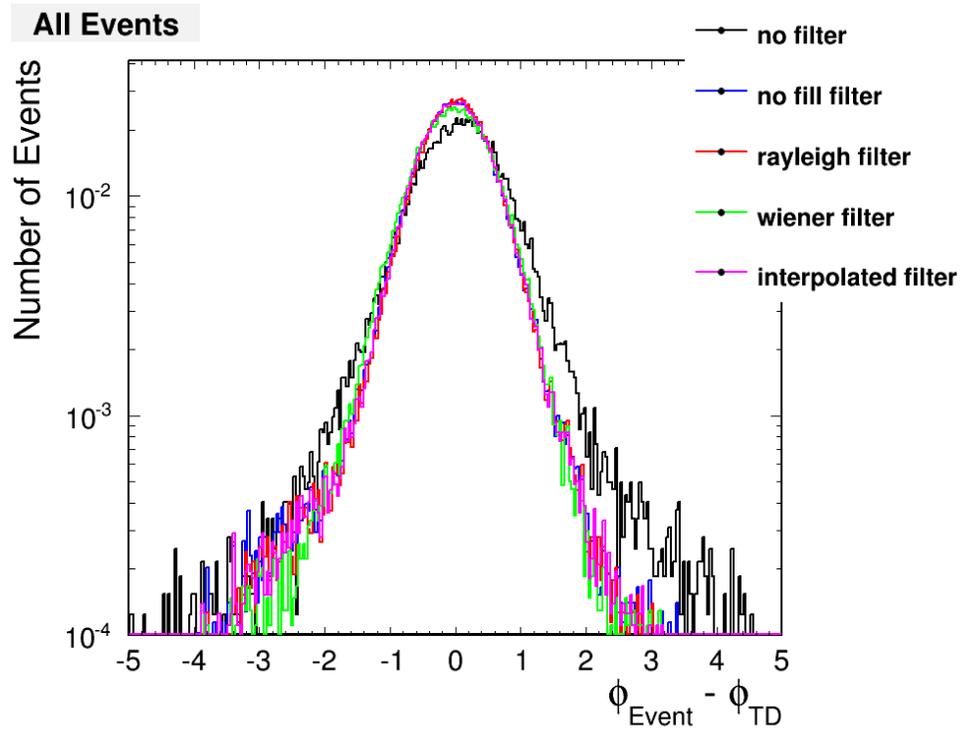


Figure 3.14: Error in ϕ reconstruction of Taylor Dome events for amplitude filters. All methods are an improvement over no filtering.

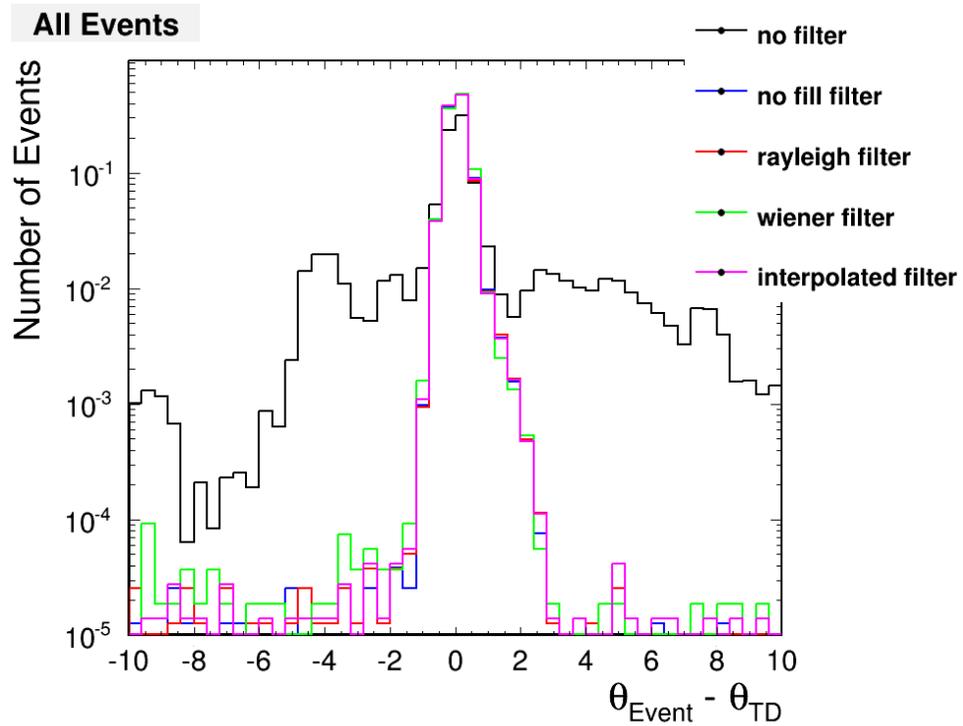


Figure 3.15: Error θ reconstruction of Taylor Dome events for amplitude filters. All methods are an improvement over no filtering, but the Wiener filter has more events on the edge of the distribution than the other methods.

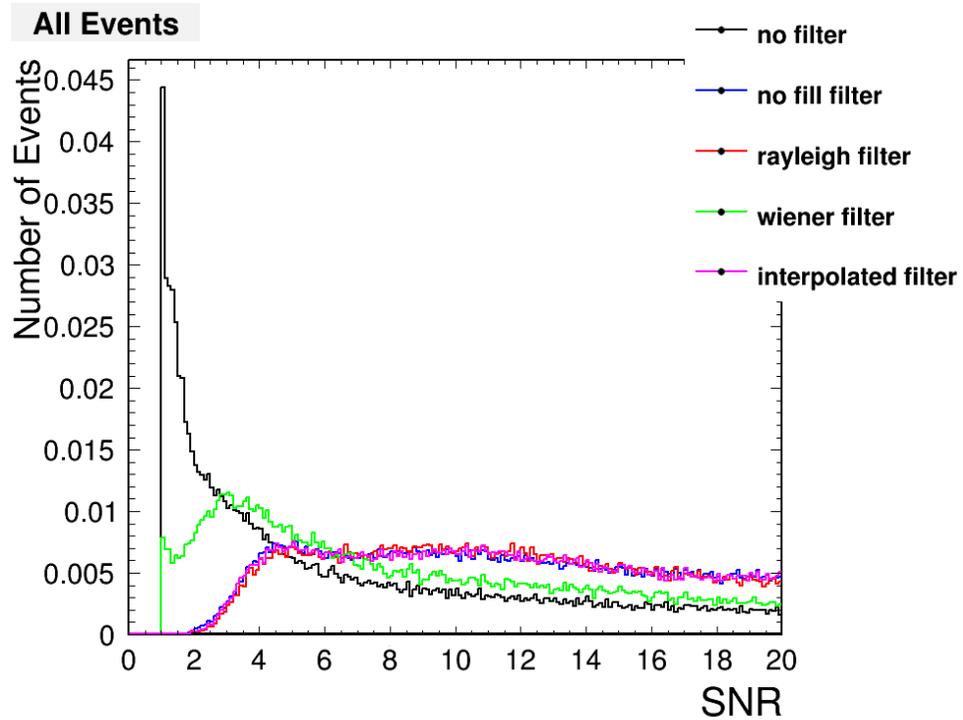


Figure 3.16: SNR distribution for Taylor Dome sample after amplitude filtering. The black line shows how the SNR of non-filtered events are located at much lower SNR than filtered events. The no-fill, Rayleigh, and interpolated methods all seem to have the same distribution of SNRs after filtering. The Wiener filter is not producing a strong impulse-like structure like the other methods, reducing the SNR of events compared to the other methods.

From this information, we decided that the Interpolation Method as described in Section 3.3.2 was the best filter to use for setting the amplitude spectrum. While the Rayleigh Method is just as good from reconstruction standpoint, we were worried about the indeterminate nature of using random numbers to select spectral amplitudes. The Interpolated Method gives the same results without that cause of concern.

3.4 Phase Corrections

In section 3.2, we discussed possible notch filters that can be used to filter the amplitude of the frequency spectrum. However, none of those filters changed the phase of the waveform. The phase spectrum is an important part of any waveform, so applying a correction to the phase could help with reconstruction and SNR. We devised, implemented, and tested different methods for adjusting the phase spectrum of waveforms.

3.4.1 Understanding Phase

To better understand the phase spectrum of a waveform, we first started to investigate a pure sine wave. From Fourier analysis, we know that an infinite, single frequency sine wave will have exactly one frequency with a non-zero complex magnitude and phase. With real data, we know this is not true due to the finite time window in which data is taken. This window, the window function, cuts short the waveform causing multiple frequencies to have non-zero complex magnitudes and phases. The data waveform of a recorded sine wave, $d(t)$, can be thought of as function of signal waveform (a sine function), $s(t)$ and the window function, $w(t)$.

$$d(t) = s(t)w(t) \tag{3.14}$$

The most common window function is a rectangular box function

$$w(t) = \Theta(t - t_i) - \Theta(t - t_f) \quad (3.15)$$

where t_i is the earliest time of the window, t_f is the latest time of the window, and

$$\Theta(t - t_i) = \begin{cases} 0 & t < t_i \\ 1 & t > t_i \end{cases} \quad (3.16)$$

known as the Heaviside step function.

While there are other windowing functions that can be used, past analyses have used the rectangular boxcar function so that is what was investigated. The phase spectrum of the waveform will be the phase of the convolution between the signal and window functions. We cannot escape the windowing function being present, so we need to understand how it affects the phase.

Our investigation started off simply, we wrote down an equation for finite sine wave

$$x(t) = \sin(2\pi f_0 t + \phi_0) \left(\Theta(t_0) - \Theta(t_f) \right) \quad (3.17)$$

where f_0 is the frequency of sine wave, ϕ_0 is the initial phase, t_0 is the starting time of the boxcar window, and t_f is the final time of the boxcar window.

With this analytic function, we can do the Fourier transform by hand:

$$X(f) = -i \left(\frac{T}{2} \right) \left(e^{-i\phi} e^{2\pi i(f_0 - f)t_{12}} \text{sinc} \left[\pi T(f_0 - f) \right] - e^{-i\phi} e^{2\pi i(f_0 + f)t_{12}} \text{sinc} \left[\pi T(f_0 + f) \right] \right) \quad (3.18)$$

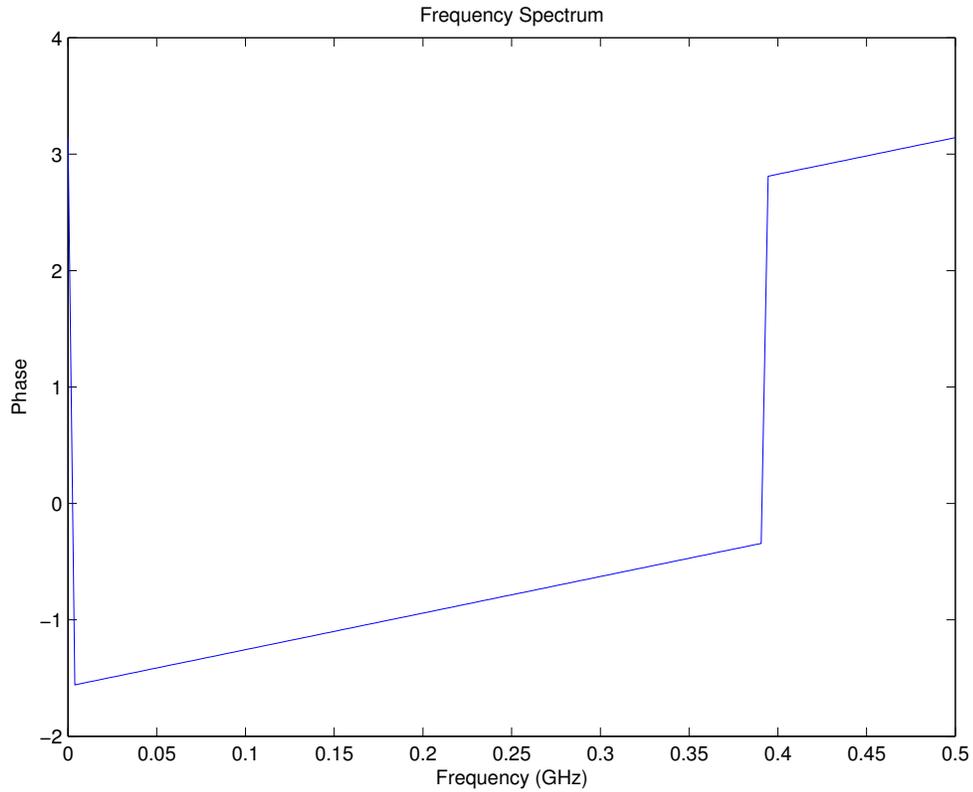


Figure 3.17: Phase Spectrum of a windowed sine function at 400MHz. We can see a linear trend for most frequency bins, but at the frequency of the sine wave (.4 GHz for this plot) there is a large jump in the phase.

where $t_{12} = \frac{t_f - t_0}{2}$ is the middle time of box car window function and $T = t_f - t_0$ is the length of the boxcar window.

The phase spectrum of this equation can be seen in Figure 3.17. We can clearly see a feature at the frequency of the sine wave. For a single sine wave, the magnitude of this jump is always π .

This large jump is characteristic of a boxcar windowed sine wave, so it could be an indicator of CW noise. But to see if this jump in the phase spectrum can be used as a discriminator, we will need to check how the phase are affected by various circumstances. We want to know what the phase of the impulse looks like (Figure

3.18), how the phase of an impulse added to a pure sine (CW) looks (Figure 3.19), how thermal noise may affect the phase (Figure 3.21), and how the phase is affected by having more than one sine wave present with different frequency dependencies (Figure 3.22).

From Figure 3.18, we can see that the phase spectrum of the impulse response is linear. When coupled with a pure sine wave of significant amplitude, we would still have a large jump at the frequency of the sine wave as seen in Figure 3.19. We noticed that the CW contamination is not limited to just the one frequency of the sine wave. Due to the windowing function, the CW power is spread out to multiple frequency bins, giving a higher amplitude in all the bins than what the impulse response would have. In those frequency bins, the CW phase dominates the impulse phase, giving a different shape for those frequency bins than for bins far away from the frequency of the sine wave. This tells us that we cannot just correct for the jump, but must correct for the surrounding frequency bins as well. This can be demonstrated by plotting the components together, as seen in Figure 3.20.

There are more complications to the phase spectrum. In Figure 3.21, we added thermal noise to a windowed sine wave. We can see that the jump is much harder to spot in the phase spectrum due to the thermal noise fluctuations. That does not mean there is not a jump because of CW; it is just harder to find.

In real data events, there can be multiple sources of CW occurring at different carrier frequencies. How the phase of multiple sine waves of differing frequencies interact has not been considered yet. You can see the results in Figure 3.22. We took two sines wave with frequencies 700MHZ and 900MHZ and added them together. We have seen that a single sine wave has a jump of π in the phase. When two or more sine waves are present, that is no longer true. This is supported analytically as well, as the jump will depend on the frequency of both sine waves and the starting phase of

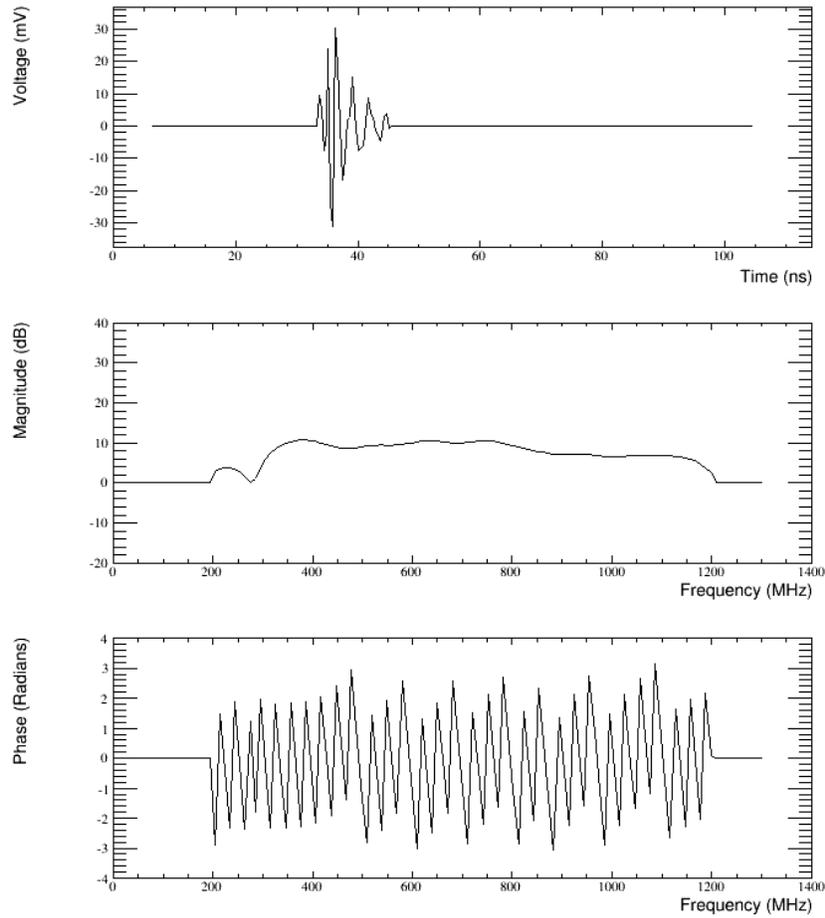


Figure 3.18: Amplitude and phase spectra of ANITA-2 impulse response. With no noise, the impulse response phase spectrum looks very jagged. This is due to wrapping (the phase exists only between $-\pi$ and π so the values that go over or under 'wrap' around to the other value. If unwrapped, the phase of the impulse response would be linear.

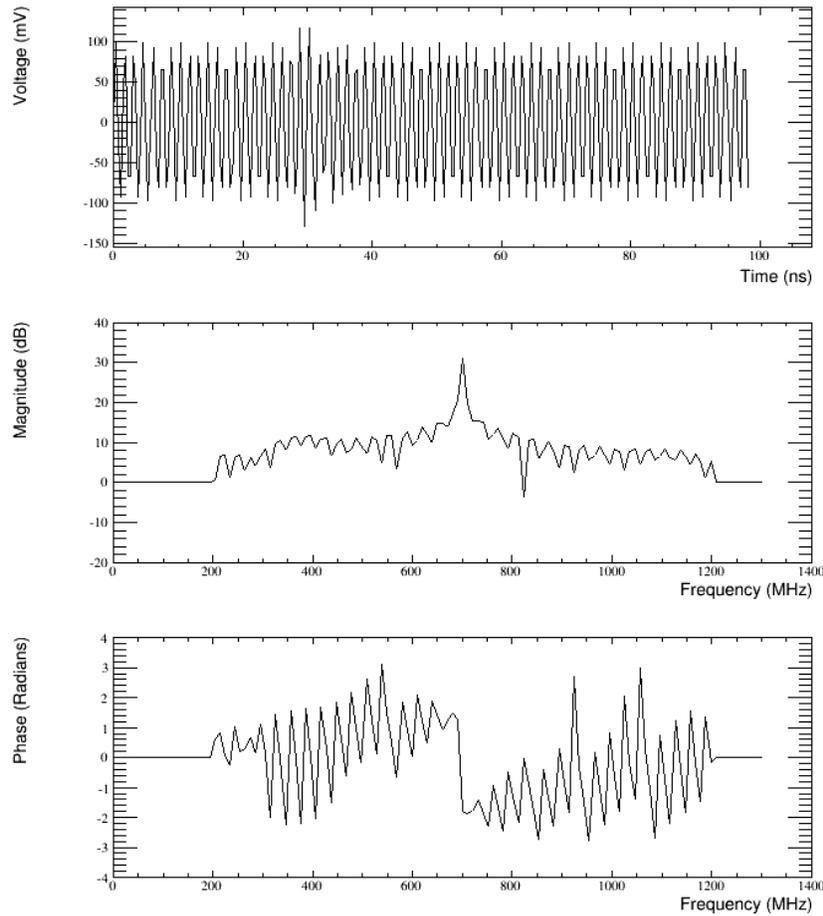


Figure 3.19: Boxcar windowed function of sine wave plus impulse. Top plot shows the time domain of the waveform. Middle plot shows the amplitude spectrum. Notice the strong peak at 700 MHz where the sine wave frequency is located. The bottom plot shows the phase spectrum of the waveform. There is a jump in the phase at 700 MHz. The impulse phase response is clearly seen in the frequency bins higher than 800MHz, but is damped in the frequency regime where CW is located.

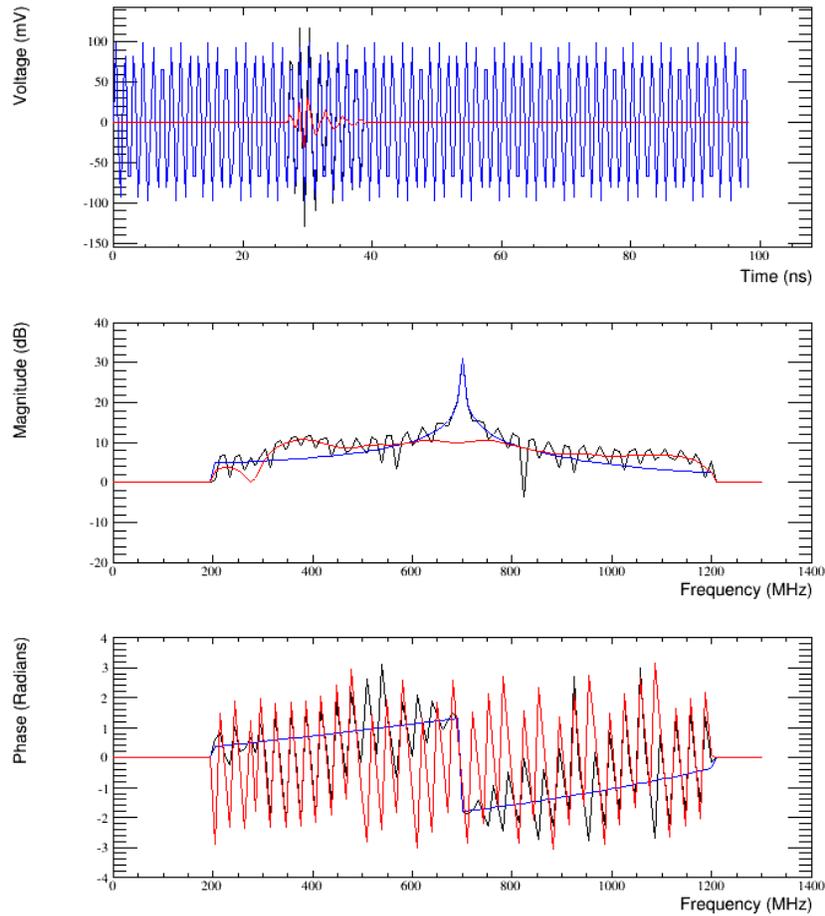


Figure 3.20: Components of Waveform. Top: Time domain, Middle: Amplitude spectrum, Bottom: Phase spectrum. By looking at the components of the individual waveform, we can see how the CW contaminates more than one bin. Blue line is sine wave with frequency of 700 MHz, red line is impulse, black is sine + impulse. Notice how the phase mimics the blue line as you get closer to the frequency of the sine wave.

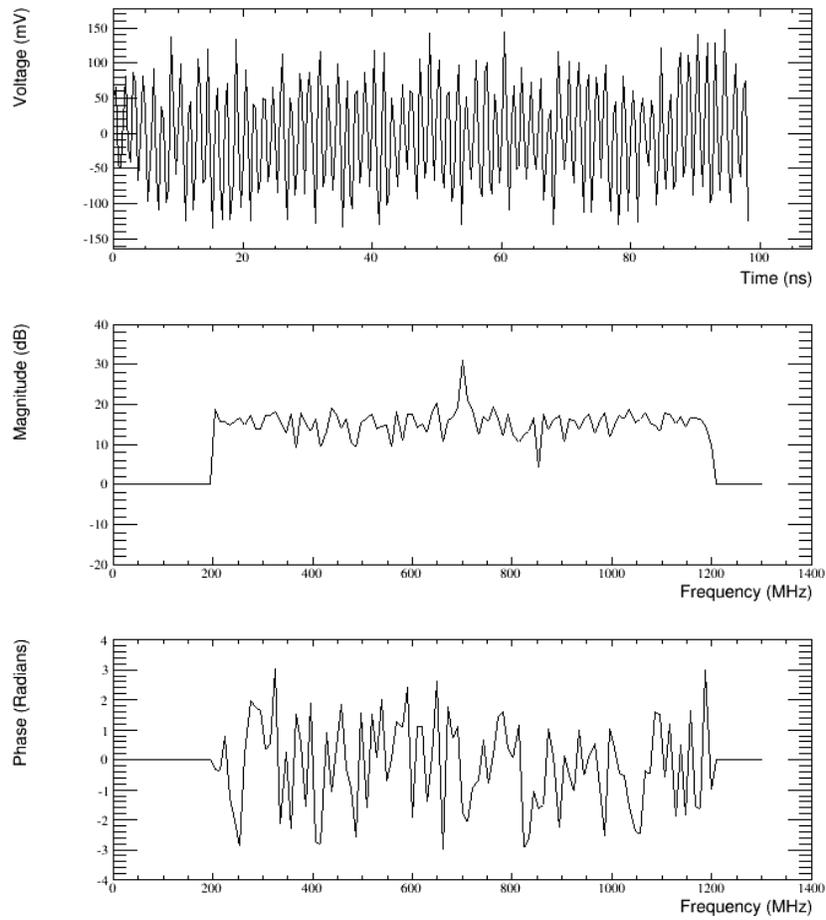


Figure 3.21: Boxcar windowed function of sine wave plus noise.

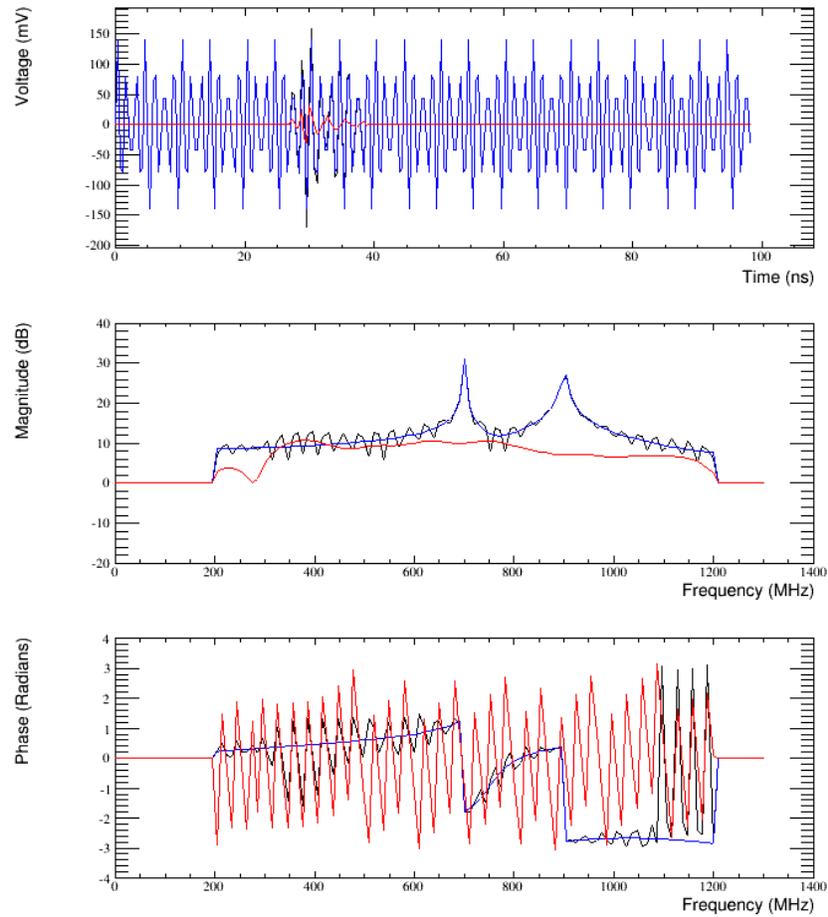


Figure 3.22: Boxcar windowed function of 2 sine waves added together with an impulse. Sine waves have amplitude, frequency of 100 mV, 700 MHz and 50 mV, 900 MHz.

both sine waves. Since we cannot know the original phase of the CW, we developed techniques to attempt to correct the jump we know is present due to the CW.

In the following subsections, we will discuss the techniques we developed and tested for adjusting the phase of a waveform. We will talk about a simple mean shift method, a complex mean shift method, interpolating phase method, and finally a geometric method. This adjusting of the phase occurs after the notch filter has been performed on the amplitude spectrum.

The simple mean shift and complex mean shift method attempt to shift the phase spectrum of the waveform such that the average of the phase is zero. The interpolated method attempts to strengthen any impulse that may be in the waveform by interpolating the phase in a certain set of circumstances. The geometric method is an attempt to extract the original phase of the signal and thermal noise from the phase information of the data. Figure 3.9 provides a Taylor Dome event that will be used for comparison for the various methods.

3.4.2 Shifting Methods

The first two techniques investigated are attempts to take advantage of three useful pieces of information: The CW causes a non-zero jump at its carrier frequency, the average phase of the impulse and average phase of thermal noise is zero, and the average phase of CW will not be zero. The impulse phase is increasing rapidly and frequently wrapping from π to $-\pi$, so the average value is approximately zero. The phase of thermal noise is random, so its average will also be approximately zero. Using this, we can attempt to take the average of the phase for a given waveform inside the notched region (the same notched region calculated for the notch filter). We can shift the phase in those frequency bins so it has a mean of zero.

3.4.2.1 Simple Mean Shift

From figures 3.17 and 3.22, we can see that for any CW contamination of the data there is a discrete jump in the phase. We wish to remove this jump in an attempt to help identify any impulse that may be in the waveform, as well as weaken any remaining CW. From the notch filter, we have already identified the frequency bins in which CW rests, the notch region. The jump in the phase will occur at the frequency of the CW, or the cut frequency. A cut frequency is a frequency found during the notch filtering process, where the amplitude spectrum peaks were flagged for filtering.

This jump gives very clear regions on which to work: the frequency bins below the cut frequency and the frequency bins above the cut frequency. These ranges extend from a cut frequency to the next cut frequency or the end of the waveform, whichever occurs first. The simplest concept we wanted to try to correct for this jump was a shift of the mean of the phase spectrum. We want the mean of phase to be zero, due to the mean phase of thermal noise and mean phase of the impulse response being zero. We take the average of the phase in each range, then shift the phase of each frequency bin by that average, such that the mean for the range is zero. See Figure 3.23 for a clear example.

This technique was created as a simple test to see if a shift can help with pointing resolution and SNR distribution. We do not expect this to be the best technique since we are treating the notched region (where CW exists) the same as outside of that region.

Figure 3.36 demonstrates the technique on the sample event, with the coherently summed waveform seen in Figure 3.25.

The advantage of this technique is its simplicity. It is easy to implement due to the low number of regions and having only to keep track of the cut frequencies.

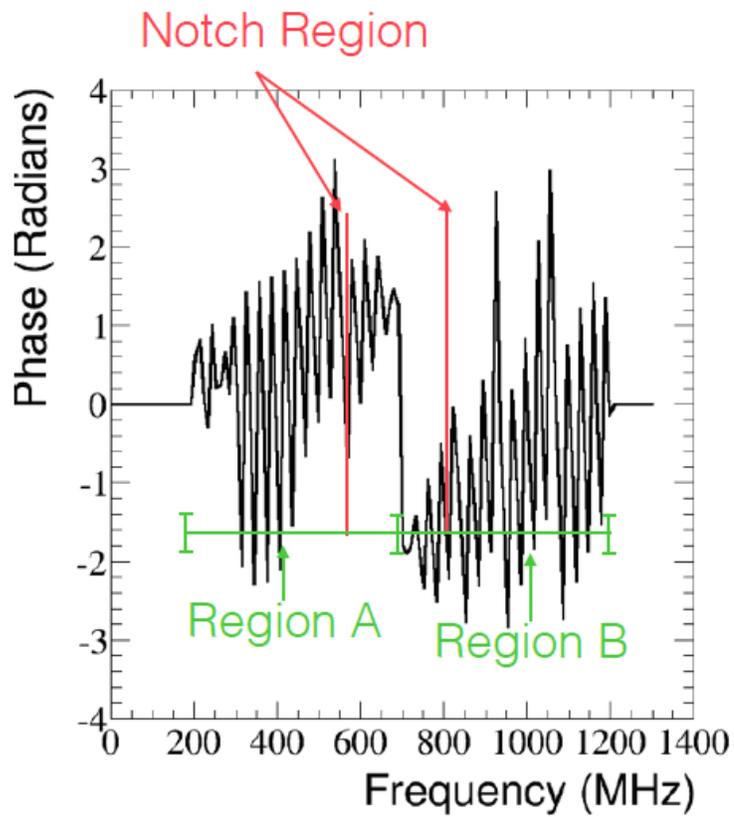


Figure 3.23: Simple Mean Shift Regions. A simple look at how we split up a region containing CW to calculate the mean.

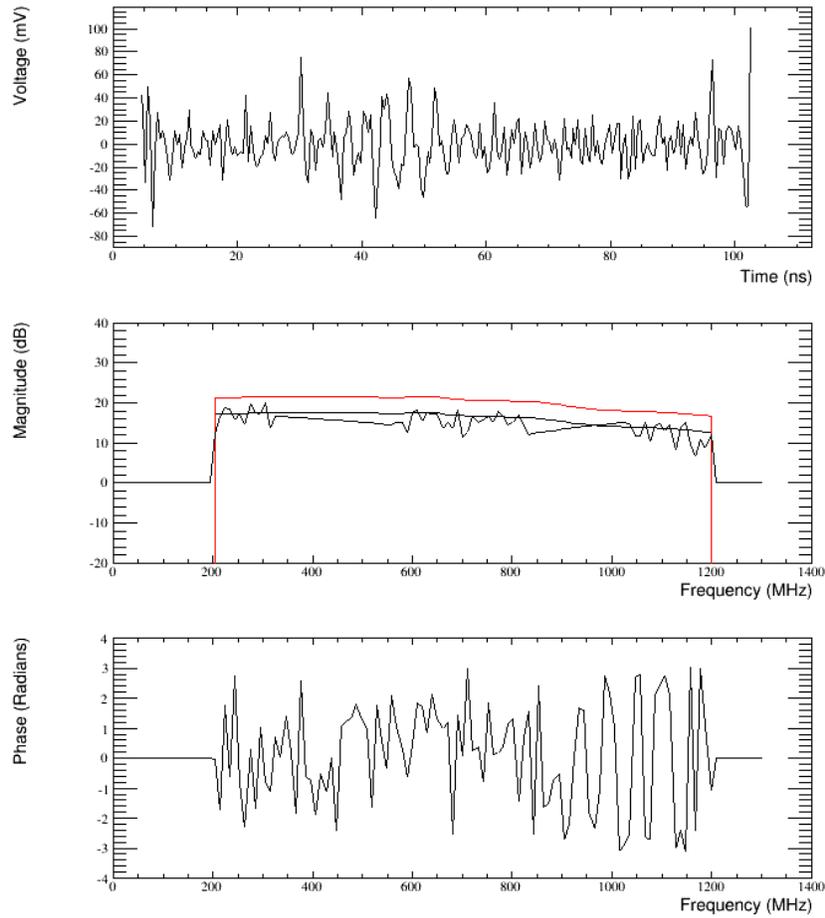


Figure 3.24: Single antenna waveform for event 687109. The amplitude spectrum has been filtered using the interpolation method. The phase has been adjusted by using the Simple Mean Shift Method. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red is thermal baseline + 4dB. Bottom: Phase Spectrum

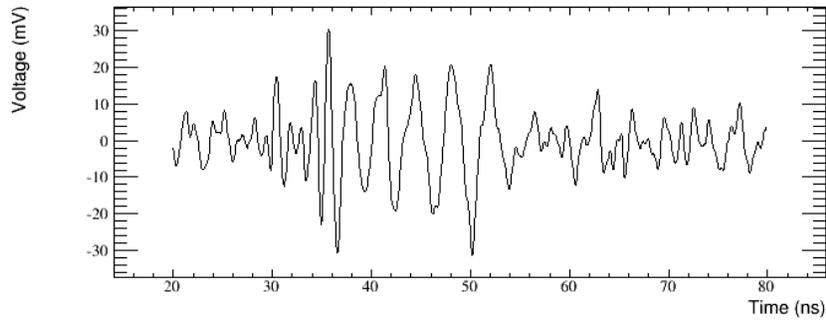


Figure 3.25: Coherently summed waveform for event 687109 after interpolation method filtering on the amplitude spectrum and simple mean shift method on the phase.

A disadvantage for this technique is the fact that this shift does not always erase the jump completely. This technique also treats the notched region where we know CW is present the same as outside the notch region where CW may not be present. Lastly, we have seen that when the CW magnitude is much larger than the impulse magnitude, the phase behaves more like the CW phase rather than the impulse phase. This technique does not correct for that.

3.4.2.2 Complex Mean Shift

Having tested the Simple Mean Shift (3.4.2.1), we wanted to attempt to correct some of the problems associated with that technique. Simple mean shift treated all frequency bins equally, whether CW is present in that bin or not. Since the presence of CW causes the phases to behave differently, those bins will have a different average than bins without CW. The way we wish to correct this problem is to have more regions in which we take an average.

From the notch filtering technique, we have regions in the frequency domain in which CW is present. Inside these notches, the CW can be a major factor in the phase,

as seen in Figure 3.20. The CW contamination does extend outside this region, but its impact on the phase is smaller. As in section 3.4.2.1, we can split up the phase spectrum into regions by using the cut frequency. Unlike with the simple mean shift, we will also use the notch region boundaries as split points for regions as well. See Figure 3.26 for an example of the regions. Just like in Simple Mean Shift, we will take the average of the phase in each distinct region and adjust the frequency bins in that region by the average. This will create a zero mean for each region.

We can test this technique on the Taylor Dome sample and see how it affects the pointing resolution and SNR distribution of those events. We expect this to perform better than the simple mean shift since we are treating regions where the phase behaves differently in a different manner. A single antenna example can be seen in Figure 3.27 and the coherently summed waveform is shown in Figure 3.28.

This technique was created to correct some problems from the simple mean shift technique. By creating more regions based on where CW is present in the frequency spectrum, we are able to treat bins with large CW contamination (inside the notch region) differently than bins that do not have such a large contamination. This should help reduce the contamination in the phase of the frequency spectrum. This technique does not address the problem of limited range in the phase spectrum when the CW contamination is large. It shifts the mean, but the reduced range is still present.

3.4.3 Interpolated Phase

In a different approach, we developed a technique that would improve any impulse present while also reducing CW contamination. The idea stemmed from the fact that the phase of the impulse response was linear and the knowledge that time shifting the time domain waveform changes the phase spectrum in the frequency domain.

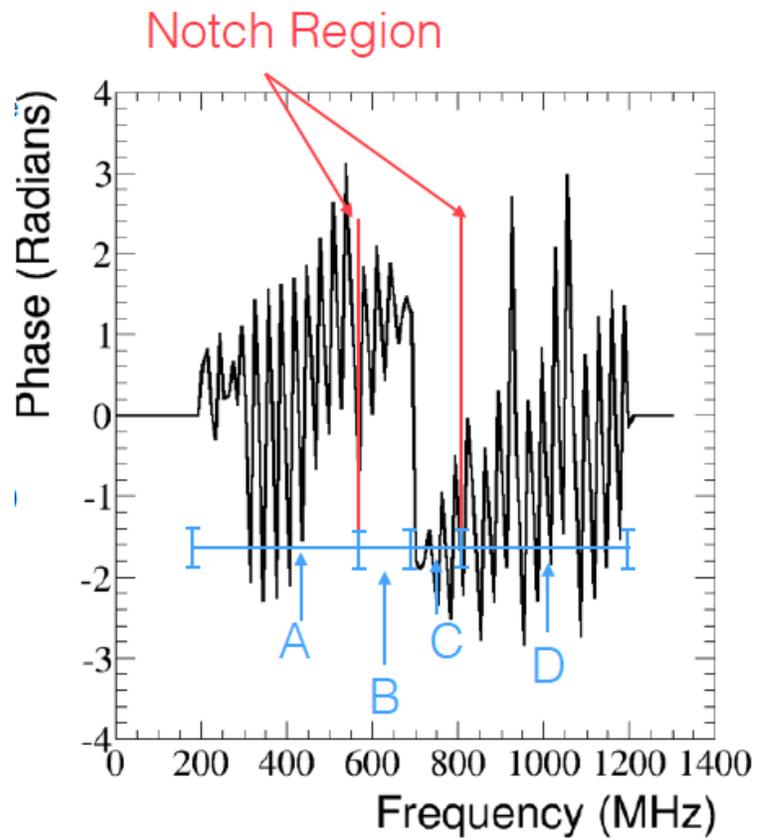


Figure 3.26: Complex Mean Shift Regions. A look at how we split up the phase spectrum into region containing CW to calculate the mean for the complex mean shift technique. Region A exists from the start of the spectrum to the lower boundary of the notch region. Region B extends from that boundary to the cut frequency. Region C extends from the cut frequency to the end of the notch region. Region D extends from the end of the notch region to either the lower boundary of the next notch region or to the end of the waveform, whichever occurs first. If there is more than one notch region, we would keep creating regions as mentioned before until we reach the end of the waveform.

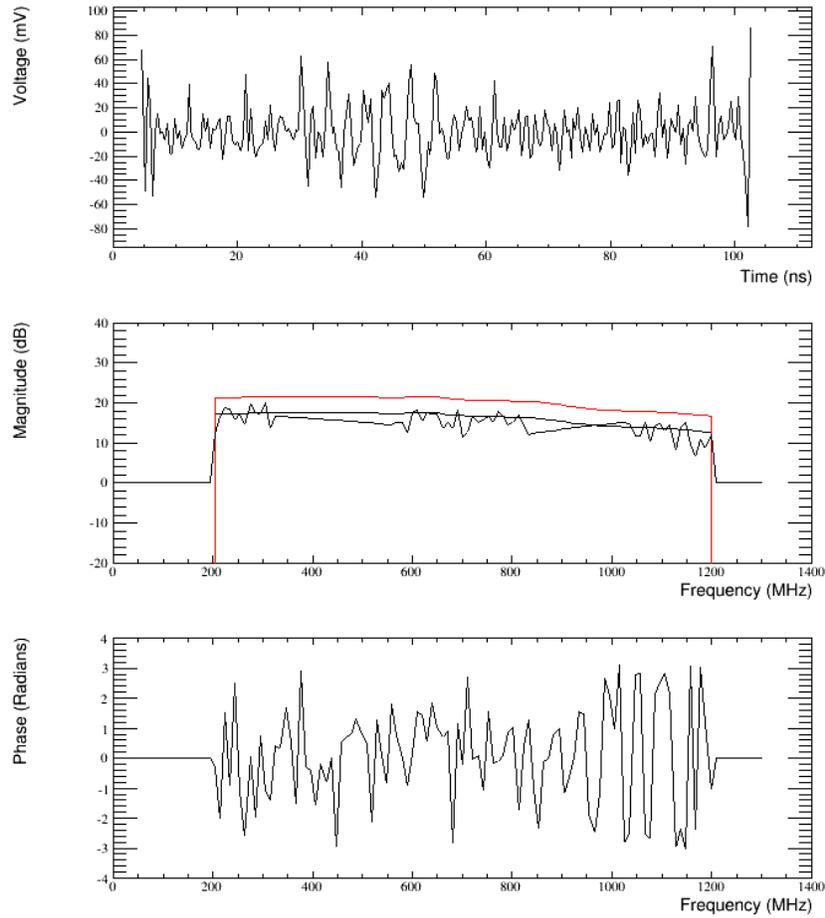


Figure 3.27: Single antenna waveform for event 687109. The amplitude spectrum has been filtered using the interpolation method. The phase has been adjusted by using the Complex Mean Shift Method. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red is thermal baseline + 4dB. Bottom: Phase Spectrum

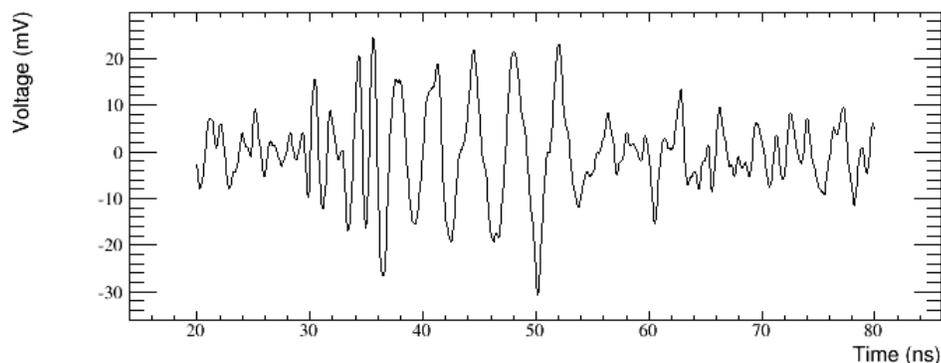


Figure 3.28: Coherently summed waveform for event 687109 after interpolation method filtering on the amplitude spectrum and complex mean shift method on the phase.

We know that if we shift a delta function impulse in the time domain such that the pulse was centered at zero time, the phase of that waveform in the frequency domain would be zero. Any change from zero would be the result of contamination from either thermal noise or from CW. If we interpolate the phase across the bins that hold the contamination, we would strengthen the pulse (because the phases in those frequency bins are now pulse like) while simultaneously removing CW's impact on the phase. We could then shift the time domain back to its original time. In practice, we cannot get a flat phase spectrum in the frequency domain for a real impulse but we can get something flatter by time shifting the waveform (Figure 3.29).

The process is as follows: after performing the notch filter, we have a time domain waveform that may contain an impulse. We find the largest voltage peak, V_{peak} in the middle two-thirds of that waveform and the time associated with that peak, t_{peak} . We chose the middle two-thirds to exclude any ringing that occurs from the notch filter. We then shift the time domain waveform by t_{peak} so the peak voltage is now in

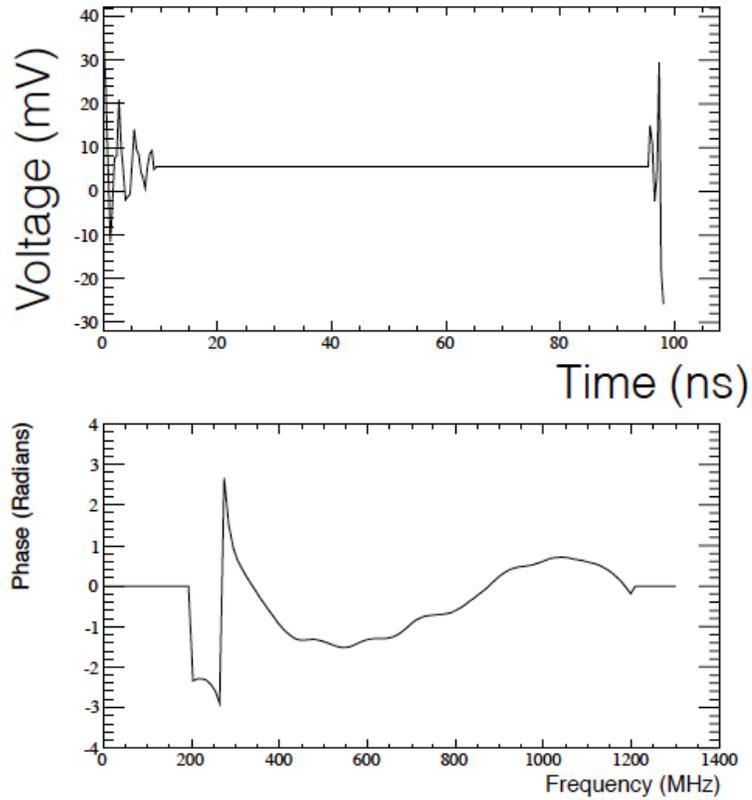


Figure 3.29: Time Shift of Impulse Response Waveform. By time shifting the impulse response waveform such that the largest voltage occurs in the first time bin, and wrapping negative times around and appending them to the end of the waveform, we can see that we have a phase that is much flatter than the non-shifted waveform Figure 3.18.

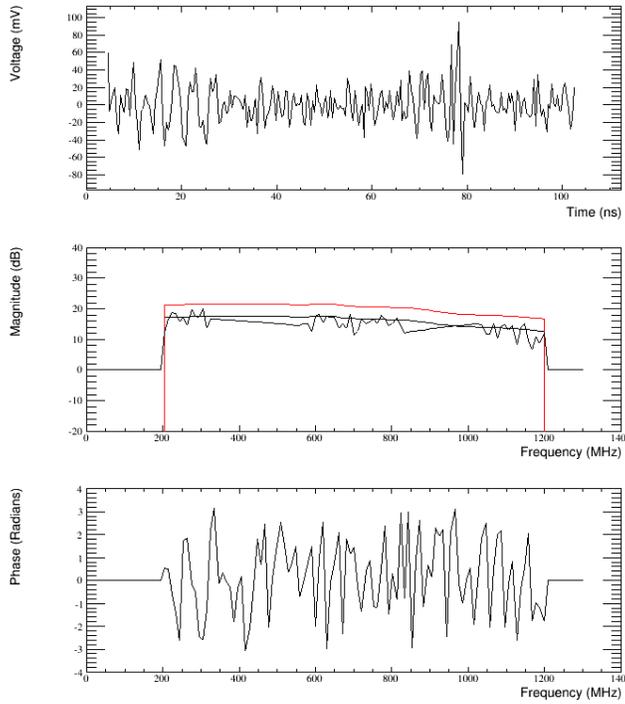


Figure 3.30: Time Shift of a Taylor Dome event. This event has been filtered with the Interpolated notch filter, then time shifted such that the largest non-ringing voltage is in the first bin. The wrapping in the time domain occurs by appending the 'negative' time values to the end of the waveform instead, causing a spike due to the ringing of the amplitude filter.

the first time bin. The waveform that went to negative time is wrapped around and appended to the end of the time domain (Figure 3.30).

After this time shift and wrapping are performed, we take the Fourier transform of the wave. Note that the ringing after the time shift looks like an impulse. In Figure 3.30, looking at the phase spectrum, we can see that we have an impulse-like structure from the ringing that is strongest inside the notch region. By doing a flat interpolation across the notch region, we remove the CW phase and replace it something more like the phase of the true impulse. Finally, we take the Fourier transform back to the time domain, then re-shift the time domain by t_{peak} . From

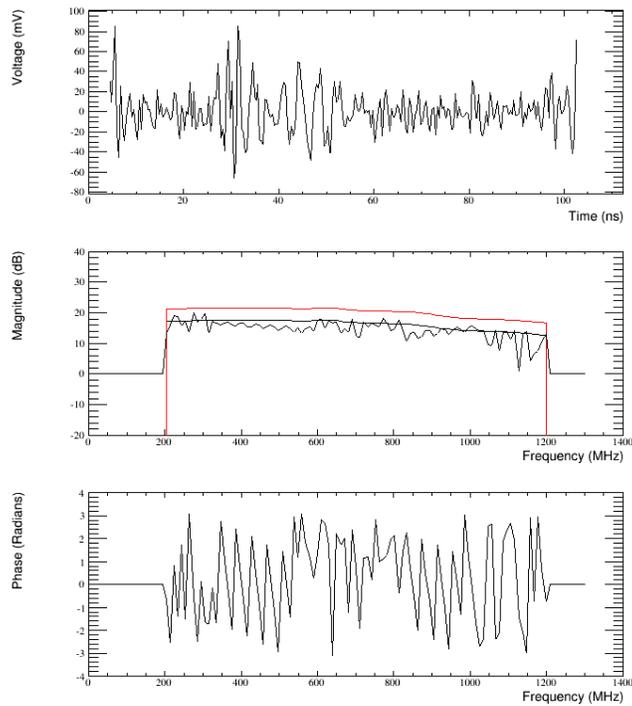


Figure 3.31: Results of Interpolated Phase. We can see how the phase spectrum has changed by performing this technique. In the time domain (top plot), we see that the impulse in the middle of the waveform now has a higher voltage, giving a better SNR. The ringing on the edges of the waveform have been reduced as well. The phase spectrum is more impulse like in the notched regions. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red is thermal baseline + 4dB. Bottom: Phase Spectrum

Figure 3.31, it seems the technique works as expected. We strengthened the impulse while also reducing the ringing that was present. The coherently summed waveform can be seen in Figure 3.32.

As a test of this technique, we will try this same process on a simulated waveform that has no impulse present, a pure sine wave plus thermal noise. With no impulse, the technique should attempt to strengthen a non-existent pulse and nothing should occur. We would change the phase of thermal noise and CW, but since there is no

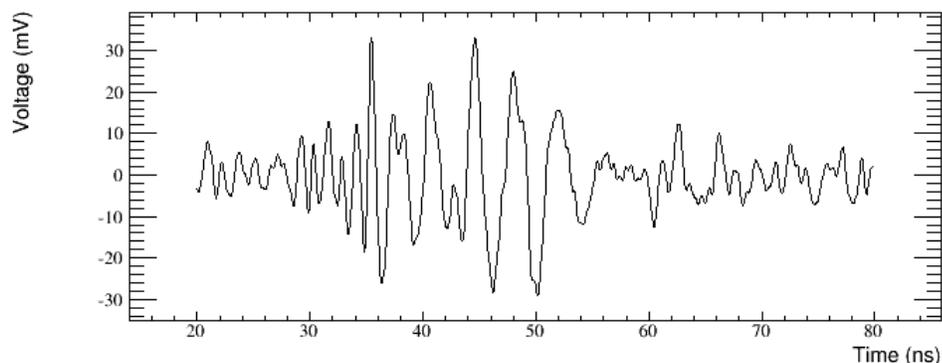


Figure 3.32: Coherently summed waveform for event 687109 after interpolation method filtering on the amplitude spectrum and interpolated phase method on the phase.

signal we would not care and this event would be caught by analysis cuts. We can see the results in Figure 3.33.

We can see a problem has occurred. We have created an impulse in the time domain by using this technique. In hindsight, this is quite obvious. We were putting an impulse-like structure into the phase spectrum, which when Fourier transformed back, should obviously give a impulse like structure in the time domain. This technique cannot be used due to that reason.

3.4.4 Geometric Method

The last technique we explored for filtering in the phase domain is called the geometric method. This method was designed in hopes of fixing the problem of limited variation in the phase from where CW dominates the frequency spectrum by using geometry of phasors. By looking at the addition of thermal noise, impulse, and CW as the phasors they are, we may be able to create a filter to help reduce the CW.

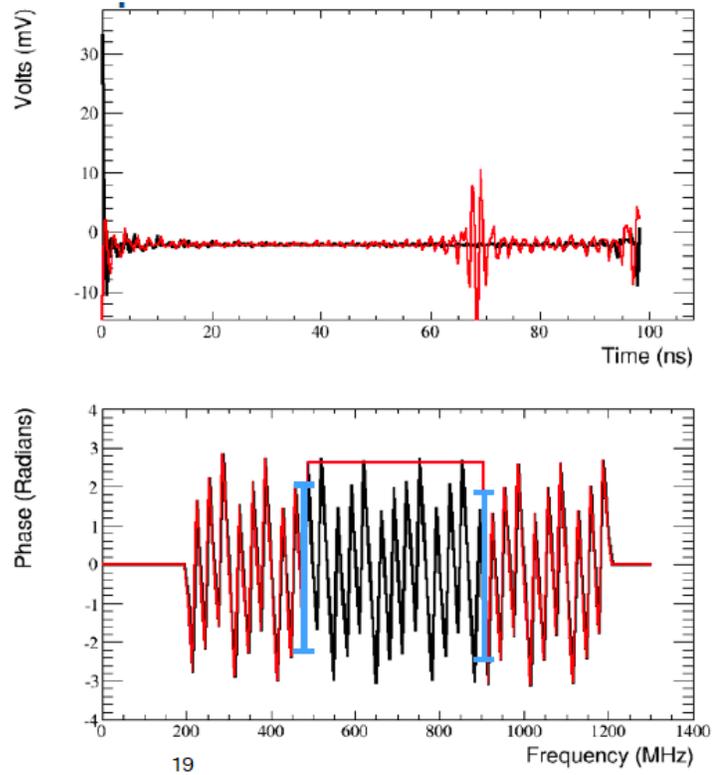


Figure 3.33: Results of Interpolated Phase Filter on a waveform with CW and thermal noise with no signal. In the top plot, we can see how the black line (pre-phase interpolation) has ringing, but no other large features. The red line (post phase interpolation) now seems to have an impulse. The bottom plot shows that the phase was changed between the two blue lines. The black line was the original phase of the filtered waveform. The red line is after phase interpolation. This technique creates an impulse!

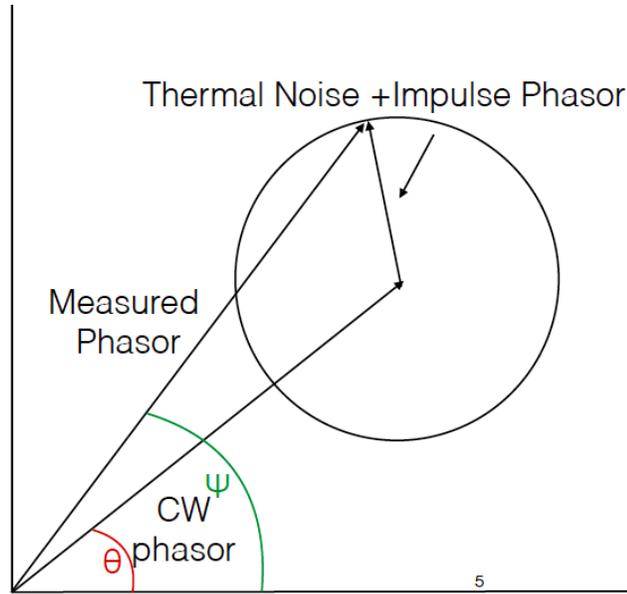


Figure 3.34: Geometric look at Phasors. By treating the sources as phasors, we can see how the measured phasor compares to the possible sources.

In a frequency bin, the CW noise will have a phase θ . The signal and thermal noise could have any phase as they are independent of each other and of the CW. Combining this idea with the magnitudes (where we expect the magnitude of the CW to be larger than signal plus thermal noise for events that caused a notch), we can see that in phasor space, (Figure 3.34), we have a ring of possible measured phasors. All possible phasors measured should lie in this ring for a set signal and thermal noise strength.

Our goal is to use our knowledge of geometry to be able to write an equation that relates the angle of the measured phasor to the angle of the signal plus thermal noise phasor and the angle of the CW noise. We hope that this equation does not rely on the magnitudes of the different phasors. We are able to relate our quantities together through geometry, as seen in Figure 3.35.

Reading from Figure 3.35, we can see that

Solving for γ ,

$$\gamma = \theta + \arccos \left[\frac{\sin(2(\psi - \theta))}{2 \sin(\delta)} \left[1 \pm \sqrt{1 - \sec^2(\psi - \theta) \cos^2(\delta)} \right] \right] \quad (3.22)$$

To use this equation, we must have estimates of θ and δ . From Figure 3.20, by looking at the the frequencies where the CW magnitude is large we can see the the measured phase varies around the CW phase. By taking the average of the phase in those bins, we can get an estimate of the CW phase, θ . From Figure 3.35, δ is the measure of the largest phase difference from the CW phase. We can get an estimate for δ by taking the largest deviation from that average phase for the frequency bins inside the notch region, which is a constant over that region.

In Equation 3.22, there are a few problems we must address before implementation of this technique. The first problem is that since we are taking an average to find θ , we must somehow address what to do for the frequency bin in which the step occurs. The second problem is that Equation 3.22 gives two possible answers. Choosing the correct solution is not trivial. Lastly, the size of δ must be less than 90° . If this limit is not observed, then equation 3.22 has no solution due to the geometry in Figure 3.35. This limit causes problems near the edge of the notch region where the CW magnitude is comparable to the signal and thermal magnitude (frequency bins where the phase starts to vary more).

To fix these problems we adopted a few limits for this technique. The first is that we will apply this technique to the frequency bins inside the notch region except for the cut frequency where the step in phase should occur and one bin to either side of that frequency bin. This is to ensure that the average phase we are calculating is correct, as attempting to average over the large step gives incorrect results. To correct the other two problems, we decided to allow delta to change from bin to bin. We set $\delta = \psi - \theta$, where ψ is the phase in that frequency bin. This solution was

Phase Filter	UnderFlows (θ/ϕ)	Overflows (θ/ϕ)
No Phase Change	671 / 193,	3 / 55
Simple Mean Shift	1113 / 426	35 / 83
Complex Mean Shift	1129 / 441	48 / 82
Geom Method	647 / 194,	2 / 73

Table 3.2: Number of over and underflows for θ and ϕ reconstruction based on phase filtering algorithms

tested and seemed to give the best results. In physical terms, this is assuming the signal plus thermal noise phase is the largest variation away from the CW phase in that frequency bin. When choosing this value for δ , the solution in Equation 3.22 collapses down to 1 solution which always exists.

After testing, the geometric method reduces the impact of the CW on the phase spectrum. As the current goal is to improve the existing filter, this method may be the best yet. It changes the phase of the waveform to a close approximation of signal and thermal noise phase.

3.4.5 Phase Filtering Results

As with the notch filters, we can compare these techniques using our Taylor Dome sample. We can look at the pointing resolution for reconstructing these events in both theta and phi coordinates, while also looking at the distribution of SNR for the events. We will use the Interpolated filter method as a baseline for comparison. We will compare our phase techniques to see if we obtain some improvement over just the notch filter alone, as seen in Figures 3.38 and 3.39 with under/overflows presented in Table 3.2. Figure 3.40 shows the SNR distribution of these filters.

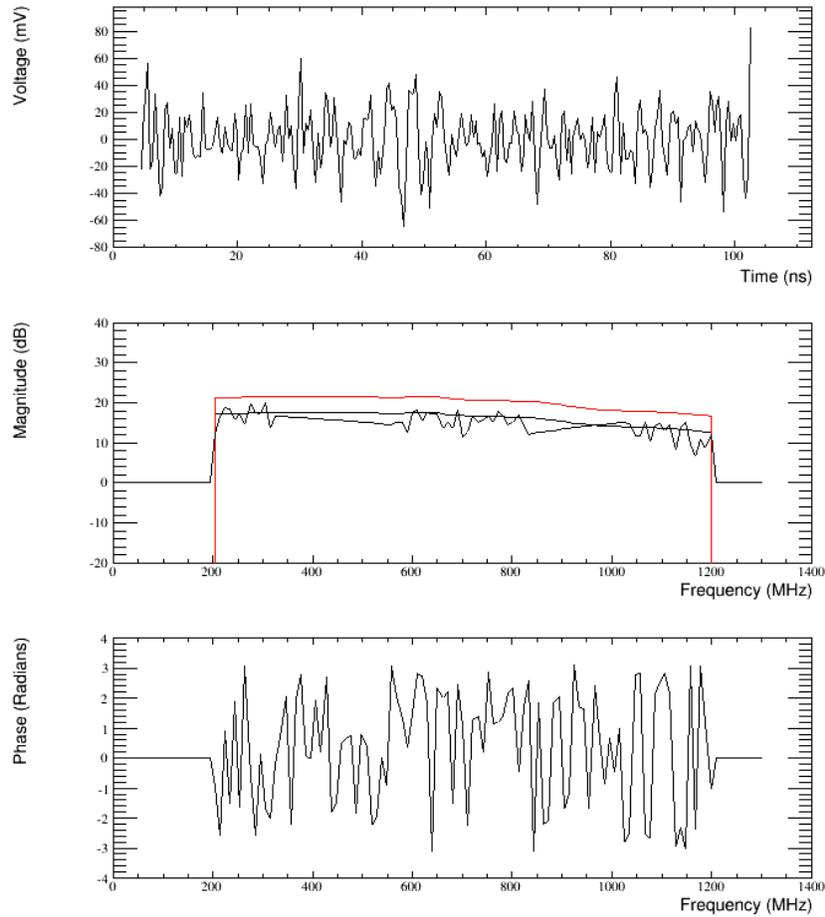


Figure 3.36: Single antenna waveform for event 687109. The amplitude spectrum has been filtered using the interpolation method. The phase has been adjusted by using the geometric Method. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red is thermal baseline + 4dB. Bottom: Phase Spectrum

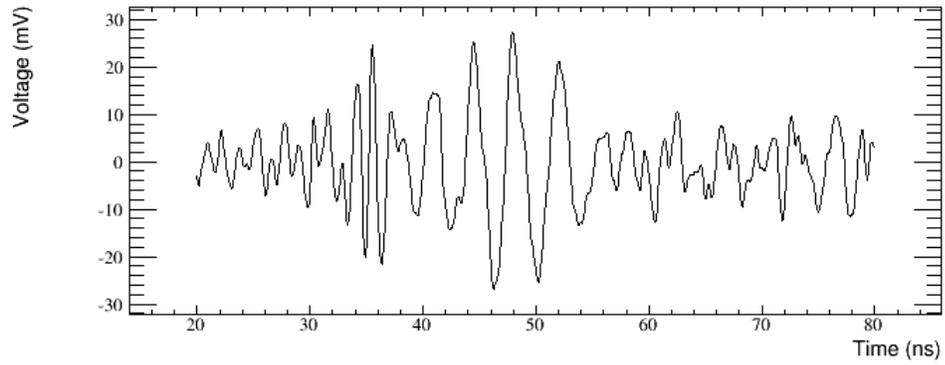


Figure 3.37: Coherently summed waveform for event 687109 after interpolation method filtering on the amplitude spectrum and geometric method on the phase.

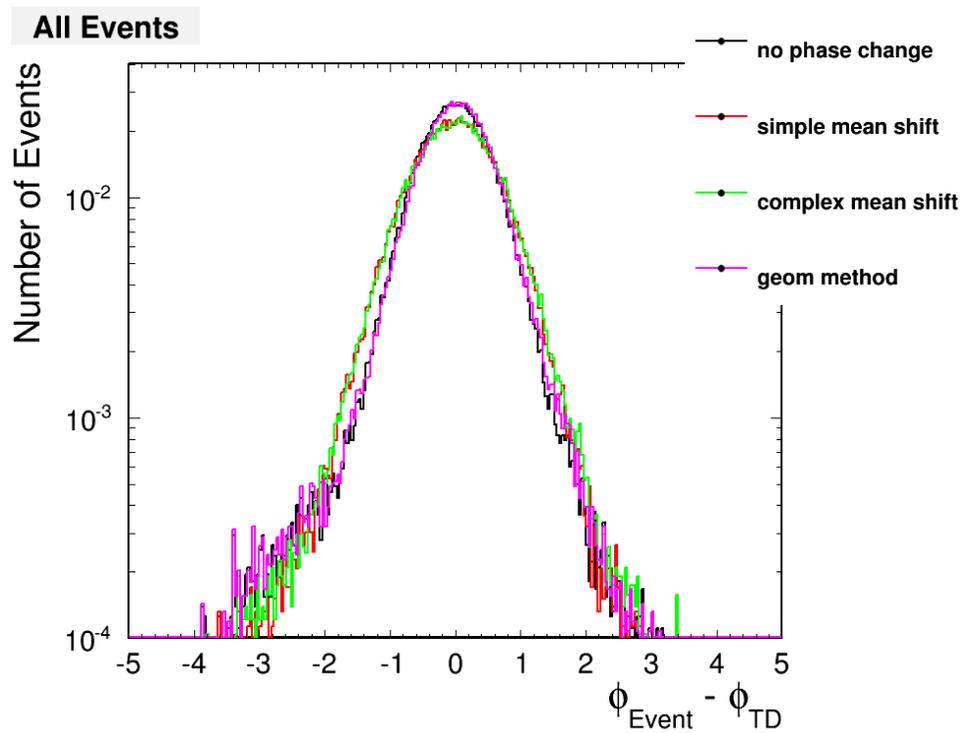


Figure 3.38: Error in ϕ reconstruction of Taylor Dome events for phase filters.

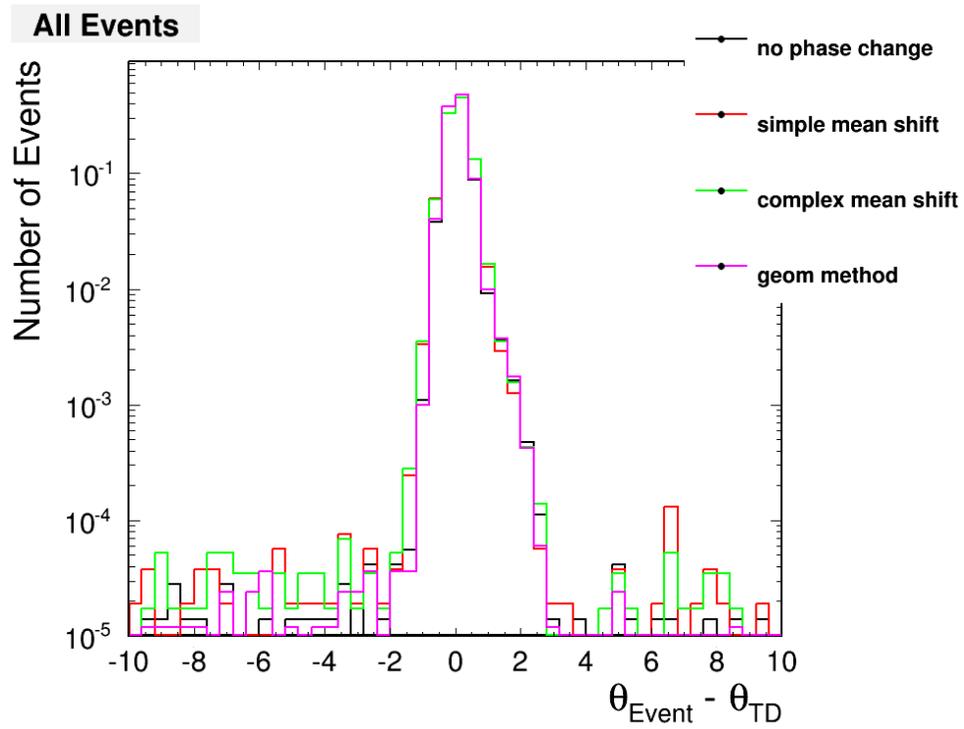


Figure 3.39: Error in θ reconstruction of Taylor Dome events for phase filters.

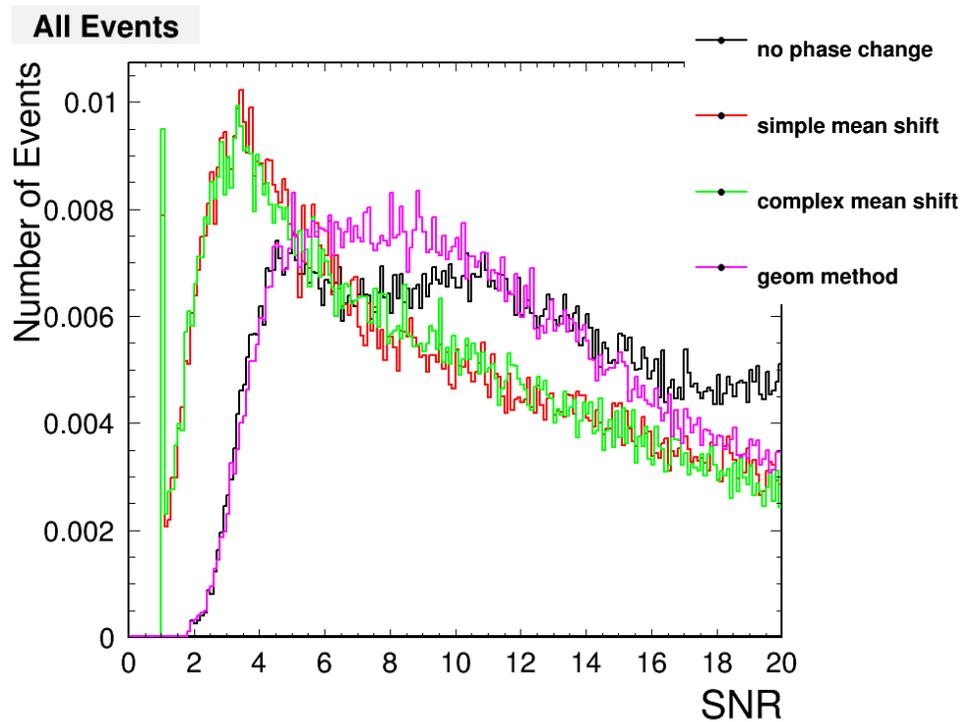


Figure 3.40: SNR distribution for Taylor Dome sample after phase filtering. Simple mean shift and complex mean shift have moved events to lower SNR compared to no phase change. The geometric method has a small shift to lower SNRs.

From these plots, we can see that the geometric method is the best method for phase filtering. It produces the best pointing resolution and has the best SNR distribution of events.

3.5 SineSubtraction

While investigating these different techniques for both magnitude and phase filtering, another technique was developed that does not rely on notch filtering. This technique instead relies on Fourier series, separating time domain waveforms into an infinite combination of sine waves with differing frequencies. Since CW power is contained in a few frequencies, we may be able to fit the sine wave and subtract out that fit. This technique was coined Sine Subtraction.

We can decompose any time domain signal into a combination of sine waves, with each sine wave having its own amplitude and phase. Equation 3.23 shows how this decomposition works.

$$s(t) = \frac{a_0}{2} + \sum_{n=1}^{\infty} a_n \sin(2\pi nft + \phi_n) \quad (3.23)$$

The method we use for this technique is very simple. We take the Fourier transform of the time domain waveform and find the frequency bin that contains the most power. Taking that frequency, we will fit a sine wave with that frequency, allowing the fit function to choose best amplitude and phase. After the fit, we subtract the fit from the time domain waveform, creating a new waveform. We repeat this procedure on the new waveform, finding the frequency with the largest power and fitting, until the change in power from the last iteration to the current iteration is less than 10%.

When fitting the sine wave to the time domain waveform, we must place limits on the parameters we use. For the frequency, we know we want a frequency inside the bin but there is still room for the frequency to change. The frequency can vary up to half a bin width ($\frac{\Delta f}{2}$). We allow this varying to get the best fit as the Fourier transform is limited to certain frequencies due to the width of our time sample, but the sine fit does not have that limit. The second limit we impose is that the phase can only exist between $[0, 2\pi]$. A fit of the original waveform and the result of the subtraction in the time domain can be seen in Figure 3.41.

This technique has some advantages and disadvantages. It is a very simple technique to understand and easy to implement. We do not have to worry about filling in any notches since we are never taking away 100% of the power in a bin. Since we are fitting the phase component, this technique also helps change the phase of the frequency domain. The only downside to this technique is the fitter. The fitter takes time and memory to work, which can cause problems if not handled properly. Since this code was developed by another person, it was difficult to get it working. The fitter used was very sensitive to the parameters and the computer environment used and did not give good fits to the time domain, skewing the results.

3.6 Filtering Results

We now need to see which form of filter is the best for use in the neutrino search. From Sections 3.3.4 and 3.4.5, we have concluded that the interpolated filter in the magnitude of the frequency domain, paired with the geometric method in the phase of the frequency domain is the best notch filter we could create. We can compare the Sine Subtraction technique against this to get the best filter possible.

We will use the Taylor Dome sample for use in the comparison. For comparison purposes, we will be comparing the old filter method, the interpolated filter, the

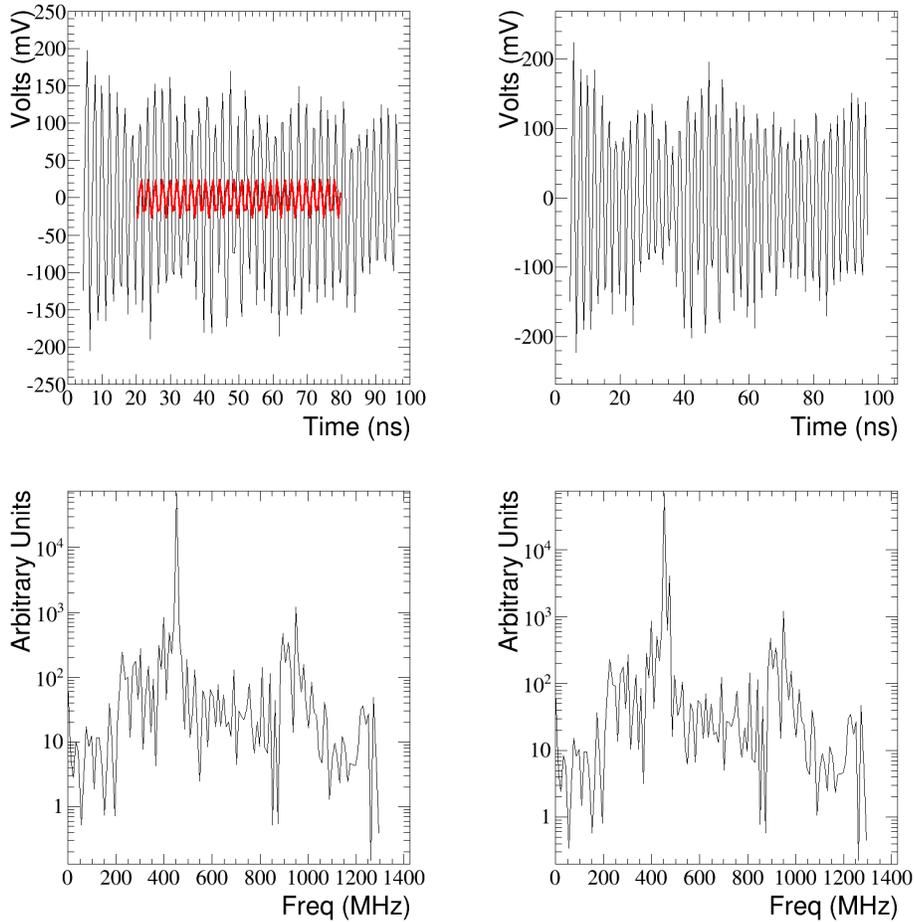


Figure 3.41: A look at sine subtraction fitting a single antenna waveform. The top left plot shows the time domain of the waveform, with the red line being the fit with the chosen magnitude and phase. The bottom plot shows the amplitude spectrum for the original waveform. The top right plot shows the time domain after one sine subtraction. The bottom right plot shows the amplitude spectrum for that time domain waveform. The fit is not a very good, a problem this method kept having. The fit should remove more power, reducing the peak in the amplitude spectrum. More investigation could provide a fix and better results in the future.

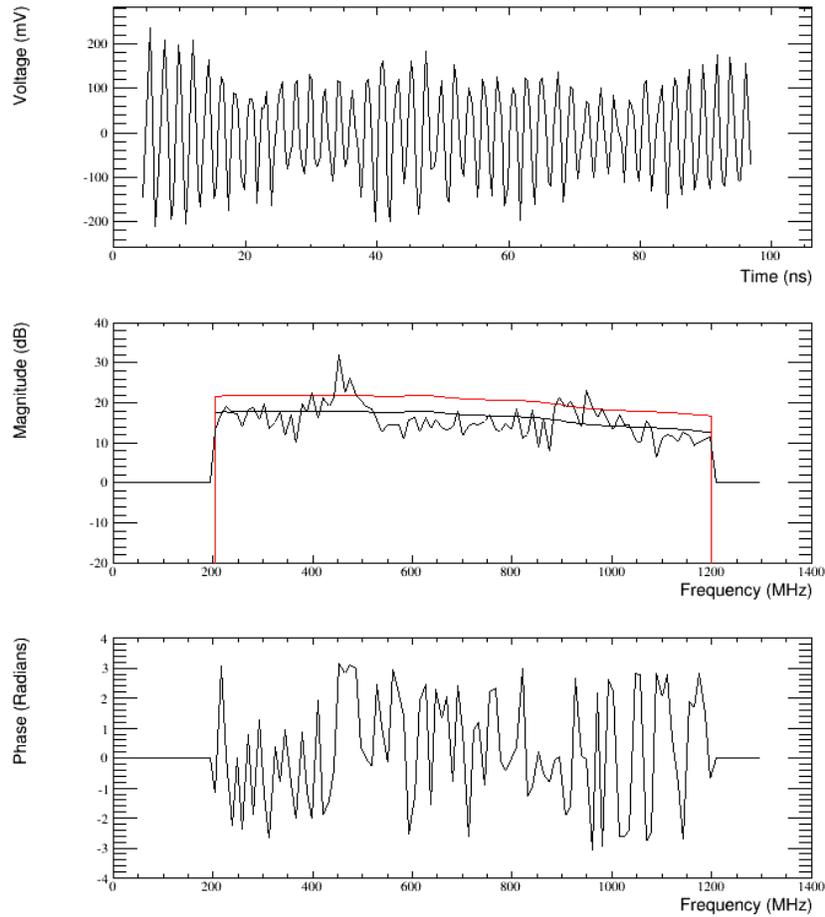


Figure 3.42: Single antenna waveform for event 687109. The waveform has been filtered using the sine subtraction method. Top: Time Domain. Middle: Amplitude Spectrum. Black line is thermal baseline, red is thermal baseline + 4dB. Bottom: Phase Spectrum

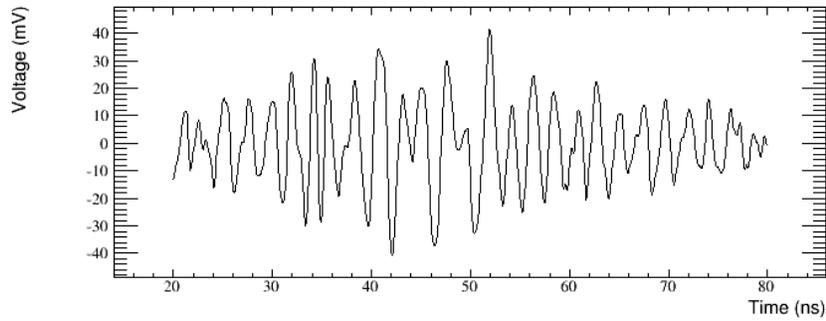


Figure 3.43: Coherently summed waveform for event 687109 after sine subtraction method has been applied.

interpolated + geometric method, and finally sine subtraction. We will look at the pointing resolution in theta and phi, then look at the pointing resolution as a function of SNR, followed by the RMS of the pointing resolution.

From Figure 3.44, we can see that all of the new filters perform better than the old method. These filters have a better resolution due to a smaller tail in both theta and phi resolution. Looking at the resolution as a function of SNR (Figure 3.45), we can see that the old method has a large error in theta and phi at low SNR. In theta, all 3 new filters perform nearly identically. But in phi reconstruction, sine subtraction seems to have more variation than the other 2 filters. At high SNR, the geometric method has the best phi reconstruction, although it is very slight.

The calibration pulser is near large human activity and the cut to select these events is not perfect. There could easily be events in the Taylor Dome sample that do not have a pulse in the waveform as expected. Since we cannot know if an event has a pulse in it a priori of filtering, we need a new test to see how these filters react to only pulse events.

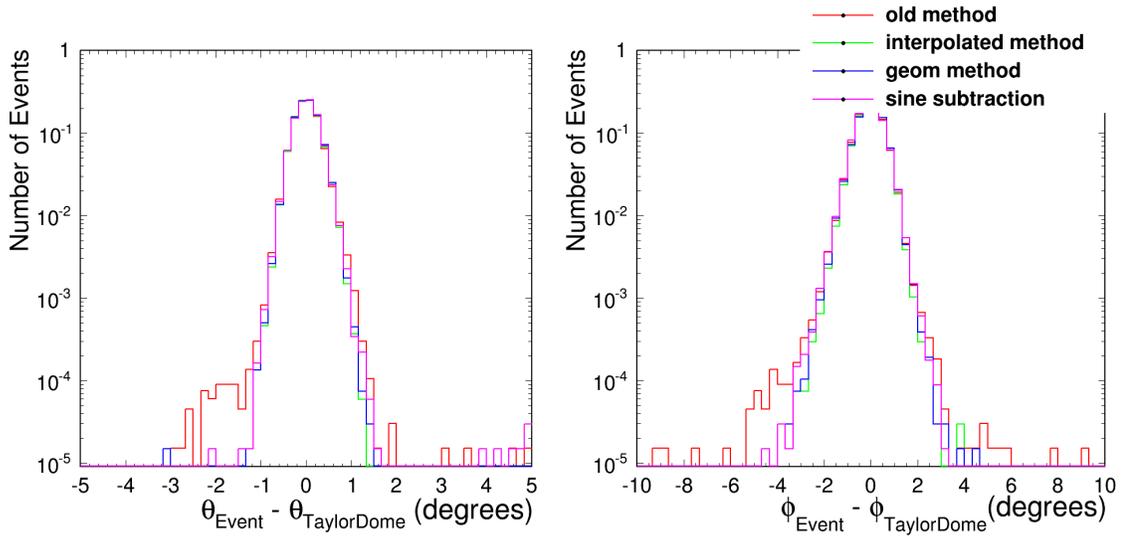


Figure 3.44: Reconstruction in the theta and phi angle from ANITA.

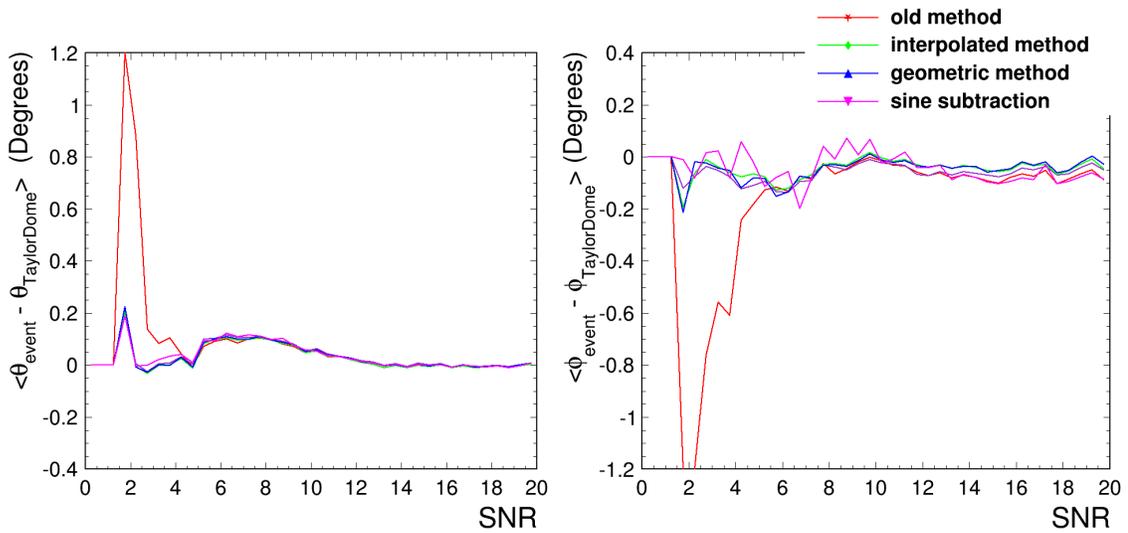


Figure 3.45: Reconstruction in the theta and phi angle as a function of SNR from ANITA.

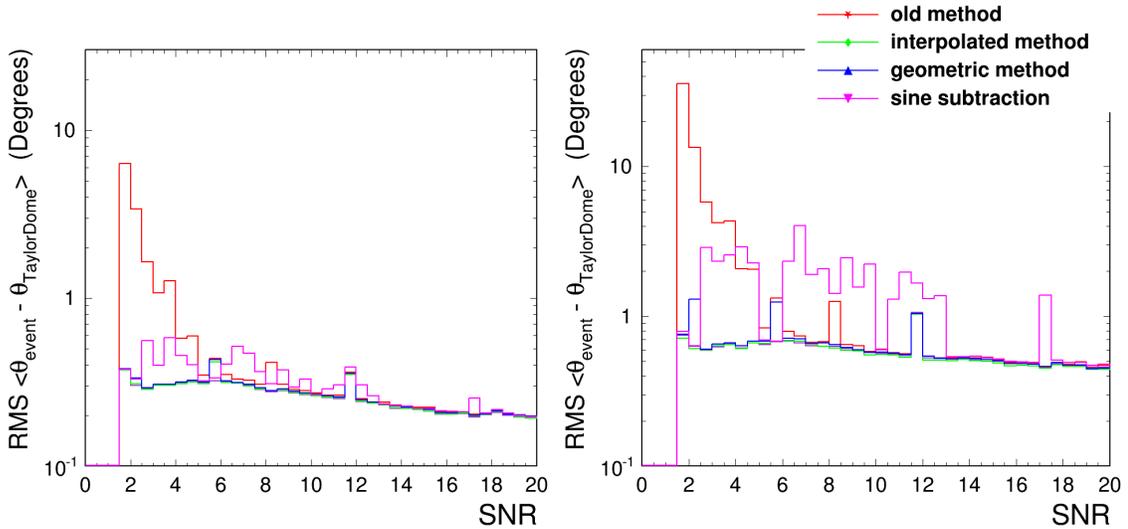


Figure 3.46: RMS of the pointing resolution in the theta and phi angle as a function of SNR from ANITA.

3.6.1 Simulated Events

In many experiments, there is a Monte-Carlo simulation to help calculate various quantities needed for that experiment. Our Monte-Carlo is called "icemc" and we use it to simulate a flight of ANITA and calculate the effective area of our experiment. We do this by modeling neutrinos and propagating them through the Earth until they interact in the ice. Then, we have the Askaryan radiation propagate to a simulated version of the ANITA detector which records the signal as the real ANITA detector would. This simulation allows one to create a sample of simulated neutrinos where we know everything possible about the event.

To use these events to test for the power of the filtering techniques, we had to make a few changes to the waveforms. Firstly, we turned off thermal noise in the simulation. This creates a neutrino signal that is purely Askaryan radiation. We did this because we will be adding CW signal from data onto the simulated waveform

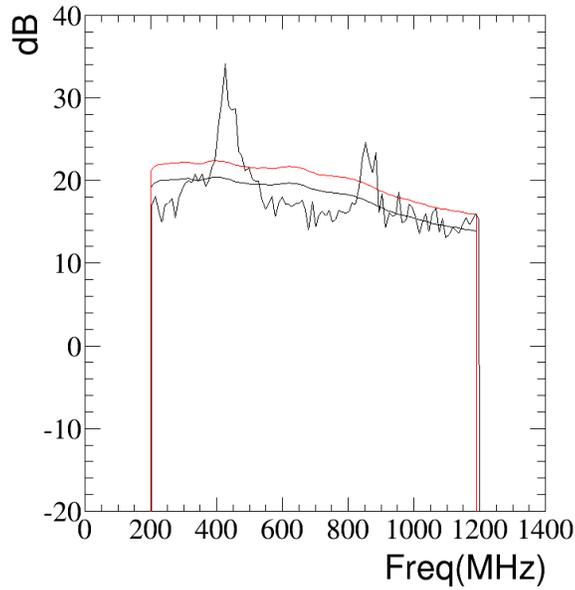


Figure 3.47: A look at the frequency distribution of one of the representative CW events used for adding to the simulated neutrino.

and did not want two forms of noise present. Secondly, we chose 3 representative CW events with varying levels of power in the CW signal for use in testing shown in figures 3.47, 3.48,3.49. For every simulated neutrino, we chose one of the three CW events at random and added the two waveforms together to form a simulated neutrino with CW. These events could now be run through the filters to test how the filters affect reconstruction.

This is a useful test because we know for each event in this simulated plus CW sample has a pulse, we also know the exact location of where the waveform exits the ice. This allows us to create similar plots to what we did for the Taylor Dome sample using Latitude and Longitude rather than theta and phi.

From Figures 3.50, 3.51 and 3.52, we can see no real difference between the filters when a pulse is present. Due to the better response on real data and the effort

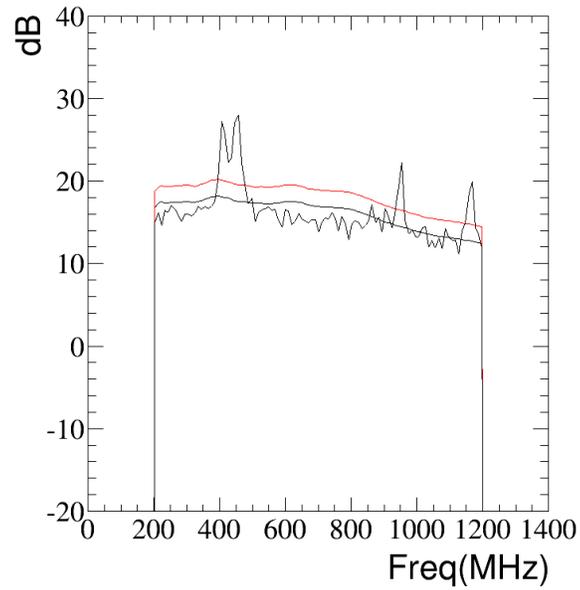


Figure 3.48: A look at the frequency distribution of one of the representative CW events used for adding to the simulated neutrino.

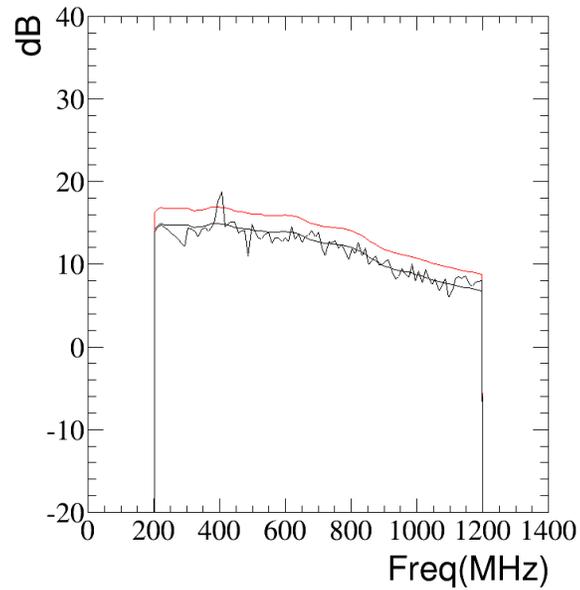


Figure 3.49: A look at the frequency distribution of one of the representative CW events used for adding to the simulated neutrino.

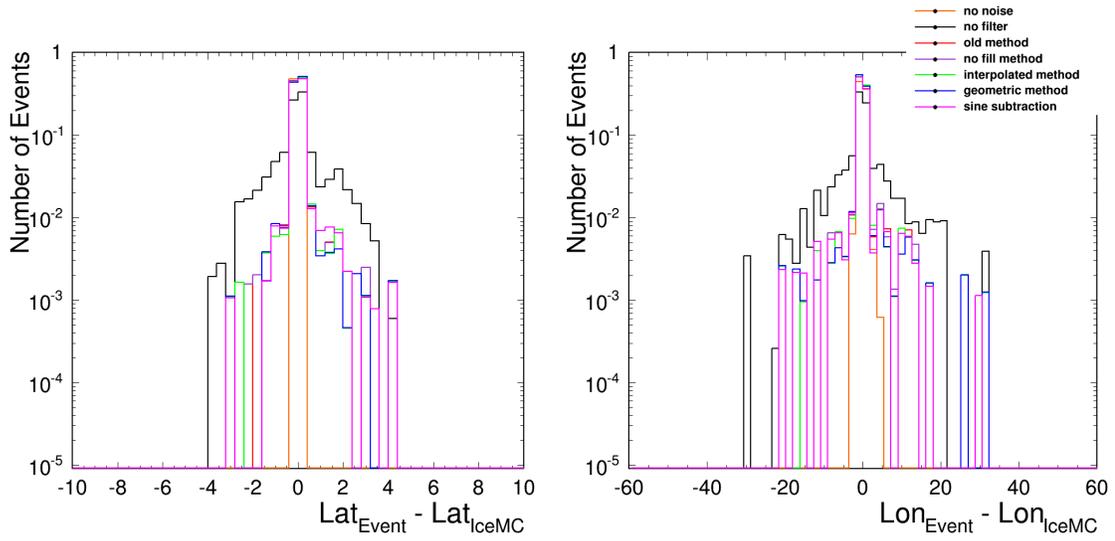


Figure 3.50: Reconstruction of simulated events to surface location.

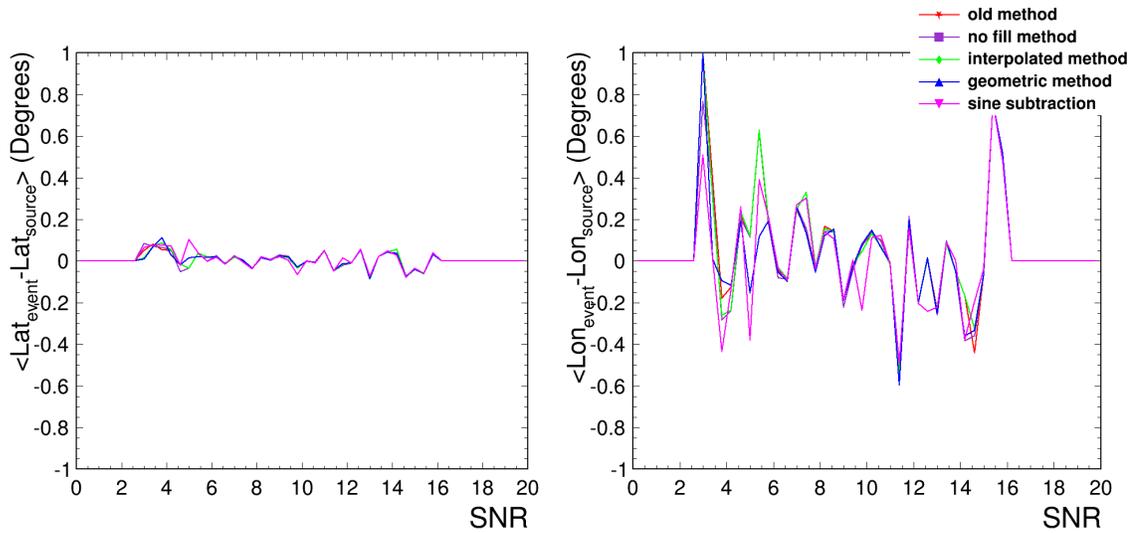


Figure 3.51: Reconstruction of simulated events to surface location as a function of SNR.

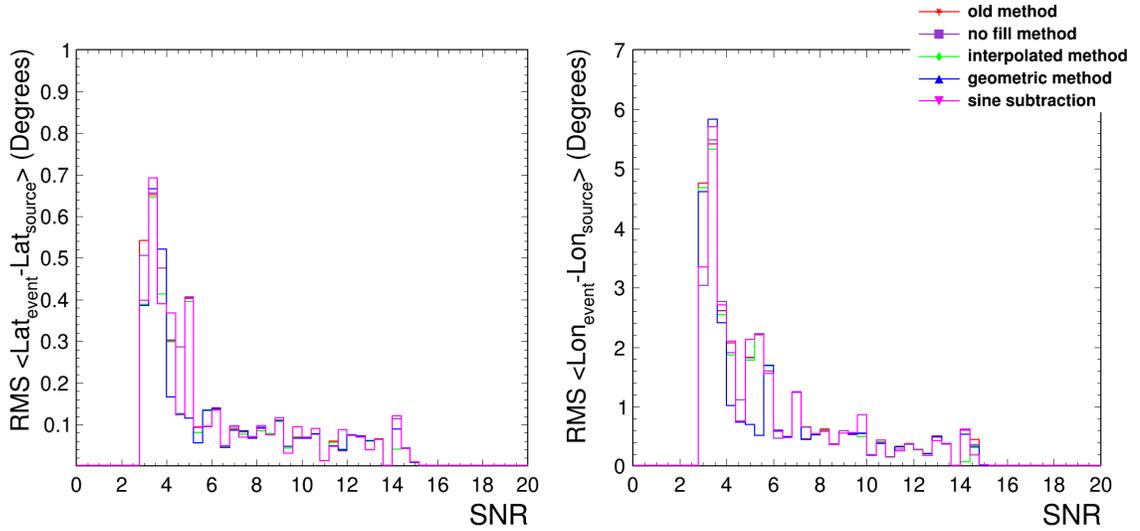
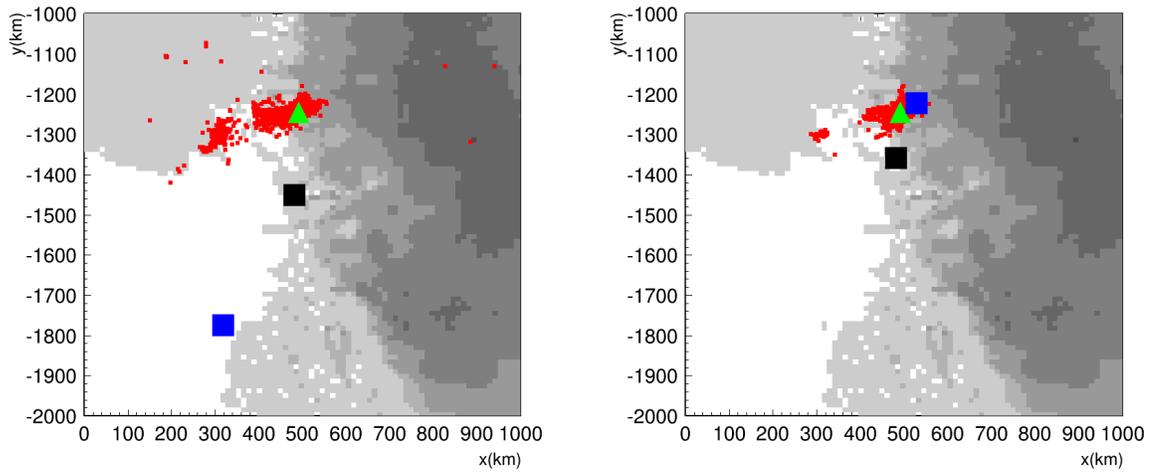


Figure 3.52: RMS of the pointing resolution as a function of SNR.

required to fix the sine subtraction technique, we chose to use the interpolated filter and geometric method as the filter for our analysis.

To help visualize the improvements in reconstruction due to the new method of filtering, we used the Taylor Dome sample and placed individual markers at source locations on a map of Antarctica. We produced this plot for both the old method of filtering and with the new method of interpolated amplitude filter and geometric method to compare. Results can be seen in Figure 3.53.



(a) Reconstruction of Taylor Dome events using old method of filtering.

(b) Reconstruction of Taylor Dome events using new method of filtering: interpolated amplitude and geometric method.

Figure 3.53: Difference in reconstruction of Taylor Dome events by old method (right) and new interpolated+geometric method (left). The green triangle shows the location of Taylor Dome. The blue and black squares show two events with a large discrepancy between the two method of filtering. The red squares are Taylor Dome events. The new filtering method has a much more contained sample than the previous methods.

Chapter 4

Re-Analysis of ANITA-2 DATA

ANITA-2 recorded ~ 21.4 M RF triggered events during its flight. Simulation and theory leads us to believe that ANITA-2 would see only a few neutrinos at most, depending on the flux model used. From this, it is clear that nearly all of the recorded events are background (non-neutrino or non-cosmic-ray) events. To have a strong statistical significance for any UHE event found, an analysis of ANITA-2 data must reliably remove these background signals to (lower levels).

4.1 Blinding

When performing a discovery type of analysis, one must be worried about biasing the result. Without some sort of blinding on the sample set, there is a risk that the analyzer will continue with the analysis until they obtain a result they believe or want to be correct. Blinding the result also reduces statistical bias, where the analyzer may change various values to include/exclude events from the signal bin.

The blinding strategy for this analysis is twofold. After producing a sample of quality events (Section 4.5), we split the sample into two parts. One part contains 10% of the events, the other 90%. The various cuts and analysis techniques will be developed and optimized using the 10% sample. Once the cuts are set to optimize for the best limit, the same cuts are then applied to the 90% sample.

The 10% sample was created in the following manner. To avoid any bias, we did not want to choose every 10th event for this subsample. We also did not want to randomly choose from all the events in the subsample, as we wanted to sample all times and the entire flight path of ANITA-2. To combat these problems, we separated the total sample into groups of 10 events sequentially, and then chose one random event from each group. This allows for even sampling from the flight, as each group of 10 occur near the same time, but the randomness allows for variable spacing between events. An analysis performed by the ARA collaboration used this technique [59].

A second method of blinding was introduced into the data itself. A select number of calibration pulser waveforms were inserted throughout the data set. The exact number and details of these inserted events were unknown to us while performing the analysis. After performing the analysis, we will remove the remaining calibration events from the signal box. A second benefit of this blinding method is a rough signal efficiency using the inserted events.

4.2 Analysis Strategy

ANITA-2 has two main background sources: thermal and anthropogenic. These two sources have very different characteristics and will have different cuts that attempt to reduce each. To remove thermally triggered events, we can require events to have multiple antennas possessing directional coherence of the signal. This simple requirement produces a sample of directional events. Anthropogenic events are harder to remove and a novel way to reduce their background will be discussed in Section 4.3.3.

4.3 Analysis Techniques

There are various techniques we developed that form the base for this analysis. These are filtering, interferometric reconstruction, and HealPix program.

4.3.1 Filtering

A majority of events triggered in ANITA-2 are contaminated by some form of continuous wave (CW) signal. CW noise may trigger the payload by itself, but often there is CW contamination in a different phi sector than the impulse signal that causes a trigger. Since CW noise reconstructs well, we must remove as much of that noise as possible. A more extensive look at filtering techniques and their importance can be found in Chapter 3.

To do this, we use an adaptive filter based on the filter used in analysis 1.0. In groups of 9, amplitude spectra of antennas were added together, and any peaks above 4 dB were notched using a variable width notch. The notch was then filled in by interpolating across the notch (see Section 3.3.2). In addition to the notch, the phase spectra of the antennas were adjusted using the Geometric method (Section 3.4.4). The functions used to identify CW peaks, notch the amplitude spectra, and apply the Geometric method to the phase are, respectively, `MyCorrelator::adaptiveFilterPartialPayload`, `MyCorrelator::applyAdaptiveFilter_singleAnt`, and `MyCorrelator::GeomMethod`.

4.3.2 Radio Interferometry

To remove thermal noise events, we use an interferometric image. This interferometric image is produced using all active antennas in a given polarization. This image

provides information on signal strength and reconstruction direction as a function of payload coordinates.

Cross-correlation is a powerful tool for finding the reconstruction direction of an event. Using a cross-correlation, we can evaluate how well two waveforms match for a given time delay between the two waveforms. The time delay between any two antennas is a function of incoming direction and the geometry of the payload. Therefore, taking the time delay that corresponds to the maximum cross-correlation for two antennas, we have a set of directions from which the waveform could have come. By introducing more pairs of antennas, these sets of directions will begin to overlap, leaving one spot with high correlation for many pairs of antennas.

Calculating the value of the cross-correlation is a multi-part process. First, we calculate the cross-correlation values between two antennas for various time delays, creating a cross-correlation function. The equation for cross-correlation value, $C(m)$ is

$$C(m) = \frac{\sum_{n=0}^N \psi_1(n)\psi_2(m+n)}{\sqrt{\sum_{n=0}^N \psi_1^2(n)}\sqrt{\sum_{n=0}^N \psi_2^2(n)}} \quad (4.1)$$

where ψ_1 and ψ_2 are the time domain waveforms of the two antennas, m is the offset between the two waveforms, n is the bin number of the waveform, and N is the number of *overlapping* bins between the two waveforms for the given delay.

The denominator of this function is new to ANITA analyses. Previously, the cross-correlation was normalized by the RMS of the individual waveforms, using the entire waveform, independent of m . This presents problems as the denominator is constant, even if there are only a few bins overlapping. Figures 4.1 and 4.2 demonstrate the problem. That normalization highly favors correlation values with the most

bins overlapping. This new normalization relies on the voltages in the overlapping bins only, giving the cross-correlation calculated at each time delay between the two waveforms equal weight, evident in Figure 4.3.

After constructing the cross-correlation function, we calculate the delay for each bin in the interferometric image, where $\theta = 0^\circ$ is pointed straight up, ($65^\circ < \theta < 150^\circ$) and azimuth ($0^\circ \leq \phi \leq 360^\circ$). Using that time delay, we pull the cross-correlation value out of the cross-correlation function. we then place the cross-correlation value into the interferometric image bin.

This process is done for each pair of antennas used. we allow each antenna to pair with each antenna within two phi-sectors and allow each pair of antennas to contribute to the interferometric image within 75° of the center of both antenna beam patterns. The nadir ring is included in these pairings, unless the nadir RFCMs were off. After all antenna pairings are performed, each bin in the interferometric map is normalized by the number of antenna pairs contributing to that bin.

During the analysis process, we create two interferometric maps in each polarization for each event. The first map is a coarse map, binning the sky into .03 by 2° ($\cos(\theta), \phi$) bins (using `MyCorrelator::DoCorrelationMap`). After finding the peak in the correlation map, we create a new map centered on the peak location, spanning 22.5° in θ and 45° in ϕ , with a .005 by 0.3° resolution (using `MyCorrelator::DoRefinedMap`). Using Akima interpolation, we interpolate across the fine binning to get the final reconstruction direction (using `MyCorrelator::DoInterpolationPeakFinding`). we create maps for Vertical and Horizontal polarizations.

Figure 4.4 shows a VPol interferometric map for a calibration pulser event, with Figure 4.5 showing the refined map, and Figure 4.6 shows a VPol interferometric map for thermal triggered event.

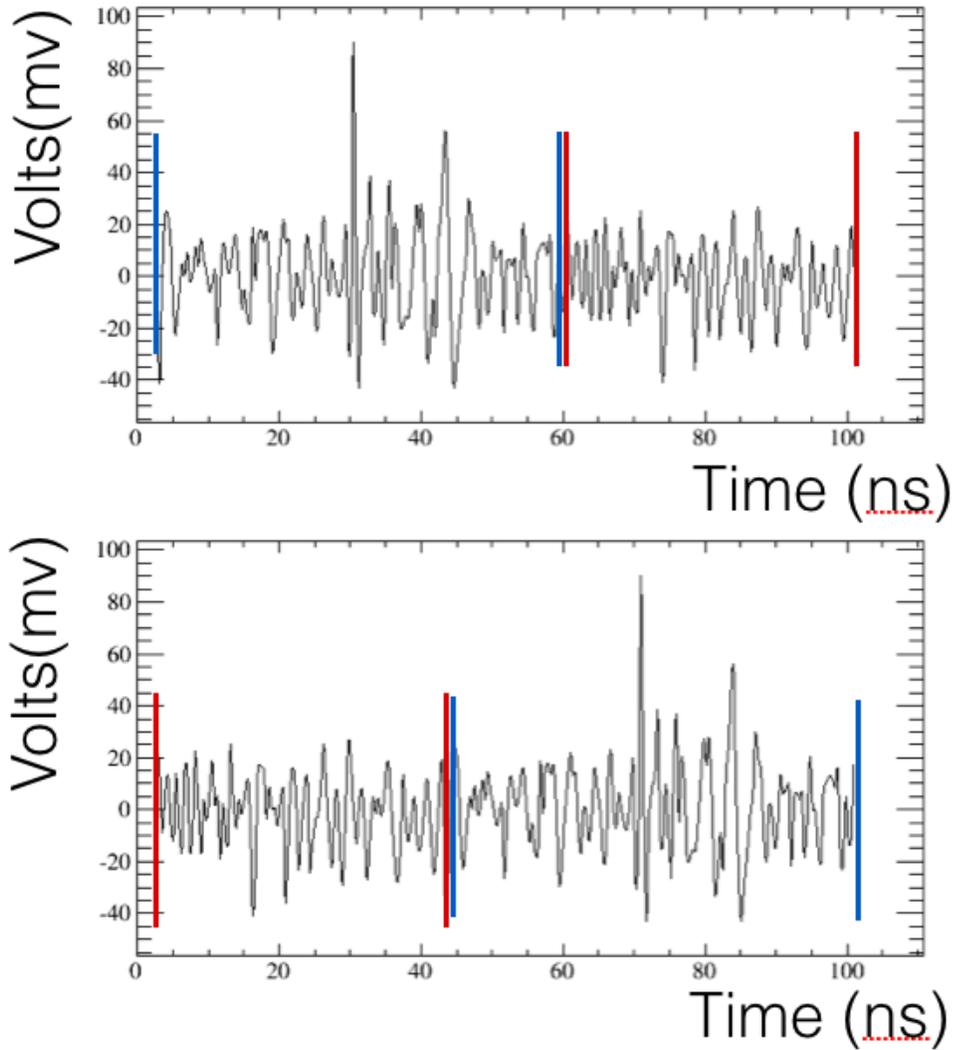


Figure 4.1: To demonstrate the normalization of the cross-correlation, we rearranged a waveform and correlated the new shifted waveform against the original. The top plot shows the original waveform, while the bottom shows the shifted. The parts of the waveform between the red lines are identical in both waveforms, and the parts of the waveform between the blue lines are identical. When the corresponding pieces overlap during the correlation, there should be a strong peak in correlation value.

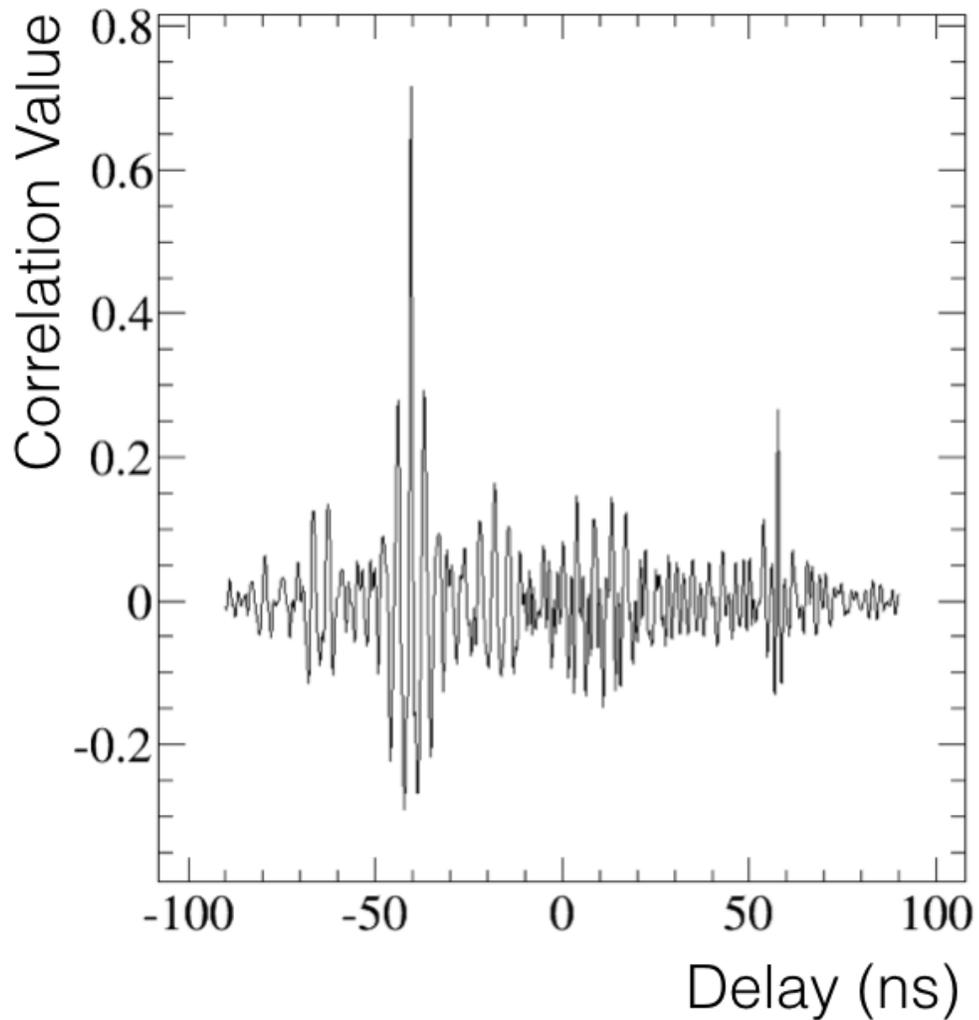


Figure 4.2: By cross correlating the two waveforms seen in Figure 4.1, we can see there are two distinct peaks. These peaks occur when the blue zones overlap, and the second when the red zones overlap. The correlation value for the second peak is much smaller than the first, even through both zones correlate perfectly.

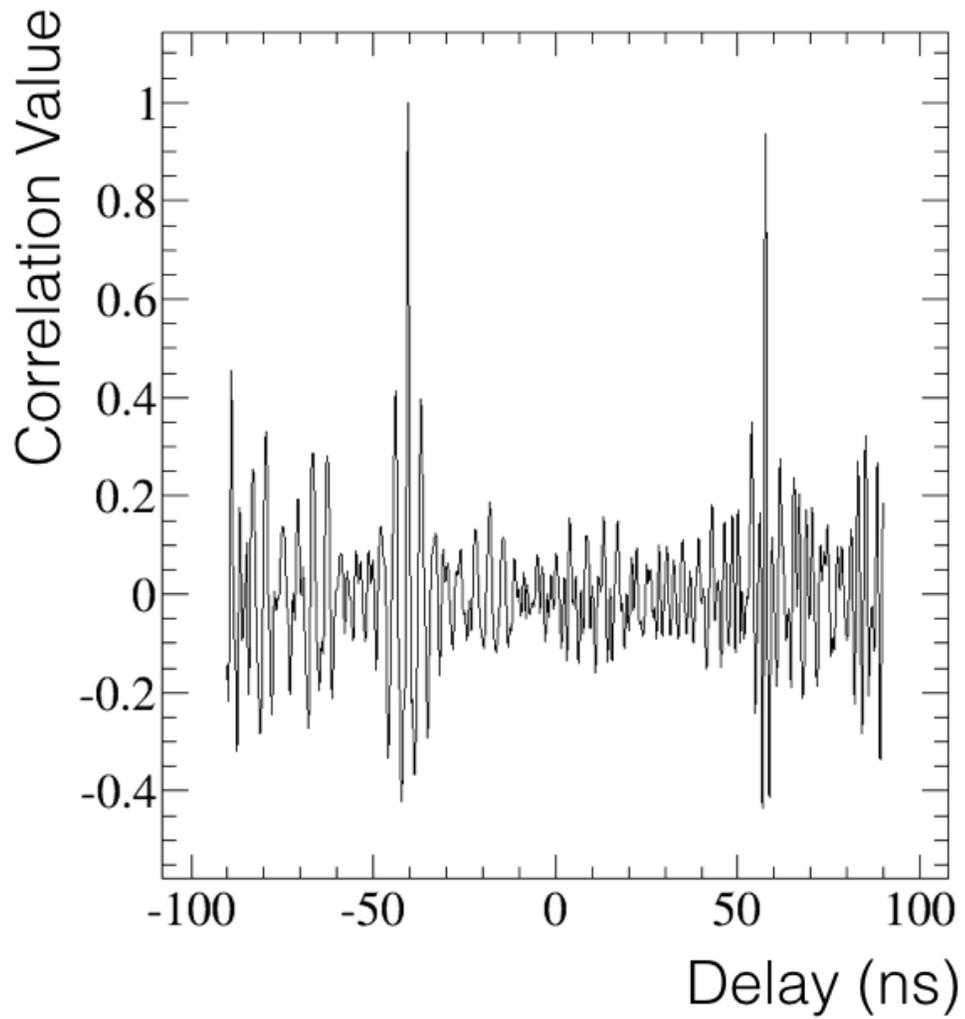


Figure 4.3: By changing the normalization of the cross-correlation to depend on only overlapping bins, the two peaks seen in Figure 4.2 are now equal in size, and further, they are both equal to one since the zones match perfectly.

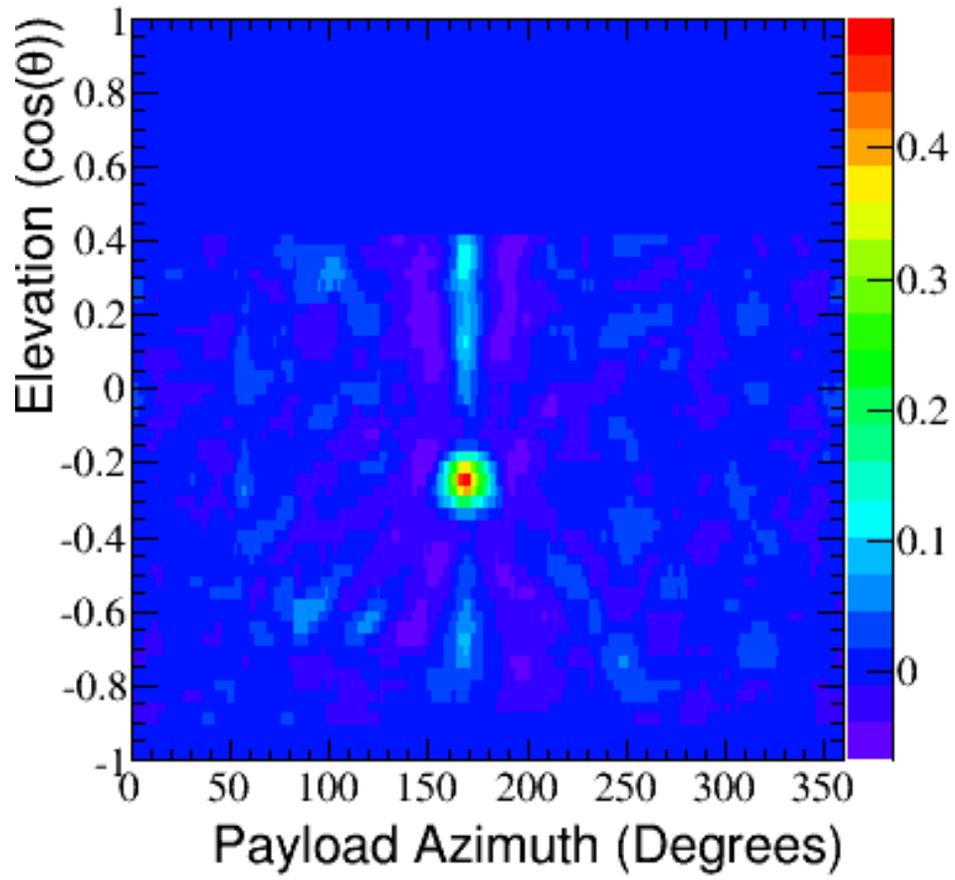


Figure 4.4: Event 1067605, a Taylor Dome event.

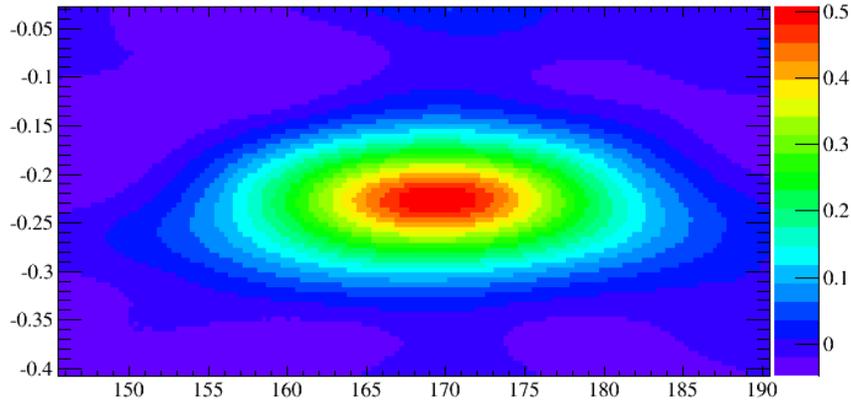


Figure 4.5: Event 1067605, a Taylor Dome event. The plot is more focused, and with finer binning than the first interferometric map.

After finding the direction of reconstruction from the peak of the correlation map, we create a coherently summed waveform from the 9 antennas closest to the reconstruction direction. The coherently summed waveform (`MyCorrelator::makeCoherentlySummedWaveform`) delays the antennas with the correct time delays using the reconstruction direction and payload geometry. The coherently summed waveform boosts the SNR of the waveform and values from this summed waveform are used for removing background events through analysis cuts, defined in Section 4.5.3. Figure 4.7 shows a coherently summed waveform for the same calibration pulser event as the one shown in Figures 4.4 and 4.5.

4.3.3 HealPix

In analysis 1.0, the previous ANITA-2 analysis, one of the cuts used to reduce background was a clustering cut. Due to the low expected flux, we do not expect neutrino events to come from the same location on the ice. An algorithm was created to determine the likelihood that either two events came from the same source, or an

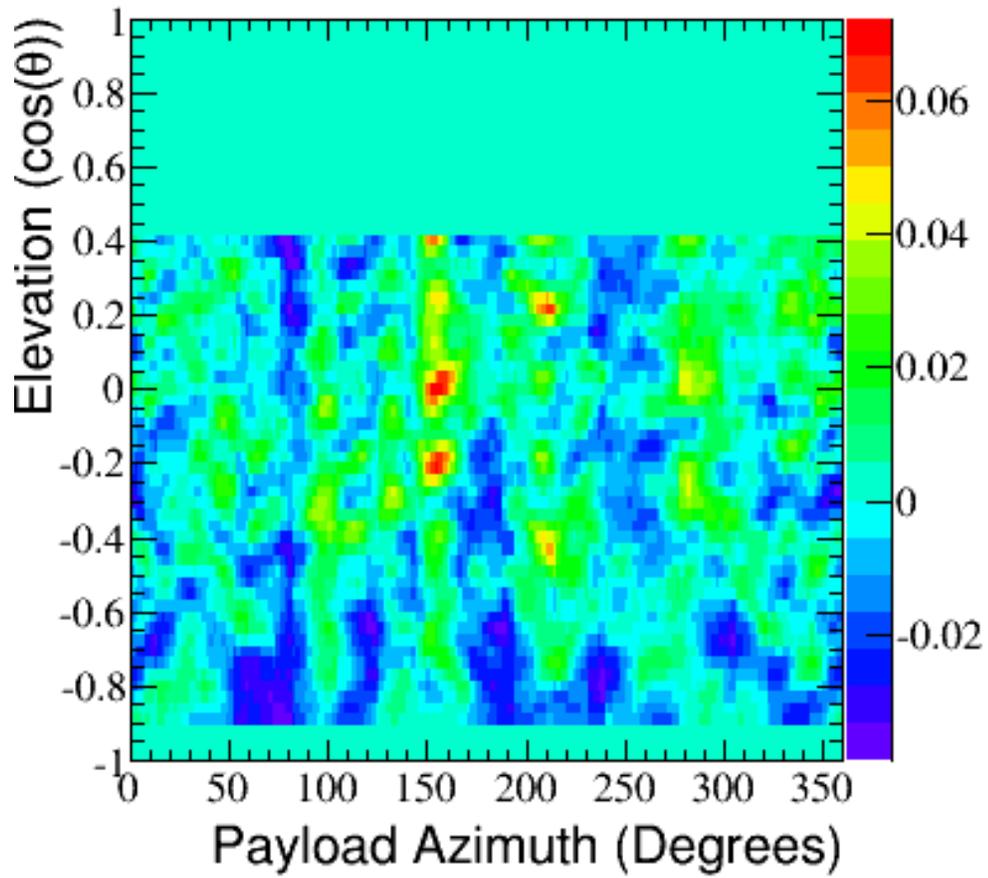


Figure 4.6: Event 544001, an event flagged as thermal.

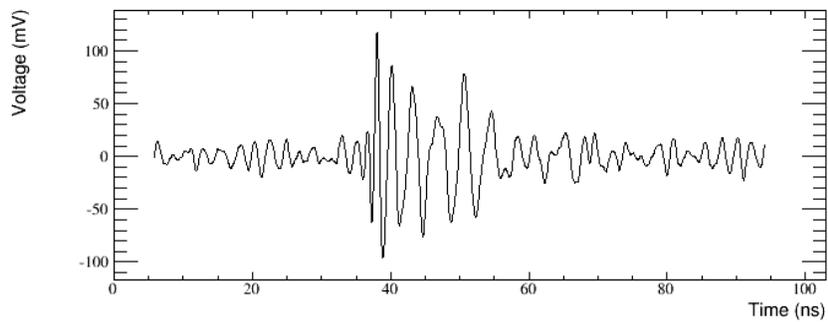


Figure 4.7: The coherently summed waveform of event 1067605.

event came from a known background location. The problem with this cut is that the likelihood algorithm reduces the amount of usable ice. As the hardware gets more advanced, we are able to reduce the trigger threshold and raise the trigger rates. The more events that reconstruct to the ice, the higher the chance that two events are within the likelihood cut for clustering. This reduces the available ice that has no events for which to find a neutrino.

To combat this, we have created a new approach to replace clustering. Instead of clustering events together, we have decided to bin up the ice in Antarctica using a program called HealPix [60]. Created by The Jet Propulsion Laboratory (JPL), this program bins space into equal sized areas. Normally used for astronomy, we appropriated the program to create equal sized bins on the surface of the earth. When events are traced from the payload to their source location on the ice, we also tag which HealPix bin the event originated from. Each event on the continent now resides in a bin and we can perform a neutrino search in each bin, with unique cuts, and combine the results together. Figure 4.8 shows the 10% sample binned into HealPix bins.

One problem with placing the events into a HealPix bin is the errors in reconstruction. We do not know the reconstruction location exactly, as both θ and ϕ have errors. We decided to use the errors when placing the events into HealPix bin by using fractional weights. We computed the area of the error ellipse (see Appendix A) for each event that is present in each HealPix bin. Each bin is then given a weight equal to the fractional amount of its error ellipse in the bin. An event's error ellipse crosses a HealPix boundary for $\sim 12\%$ of events in the 10% sample. An example of the error ellipse placed over HealPix bins can be seen in Figure 4.9.

This approach allows us to potentially use all the ice on Antarctica. HealPix bins with more background events will have different cut values and may contain a

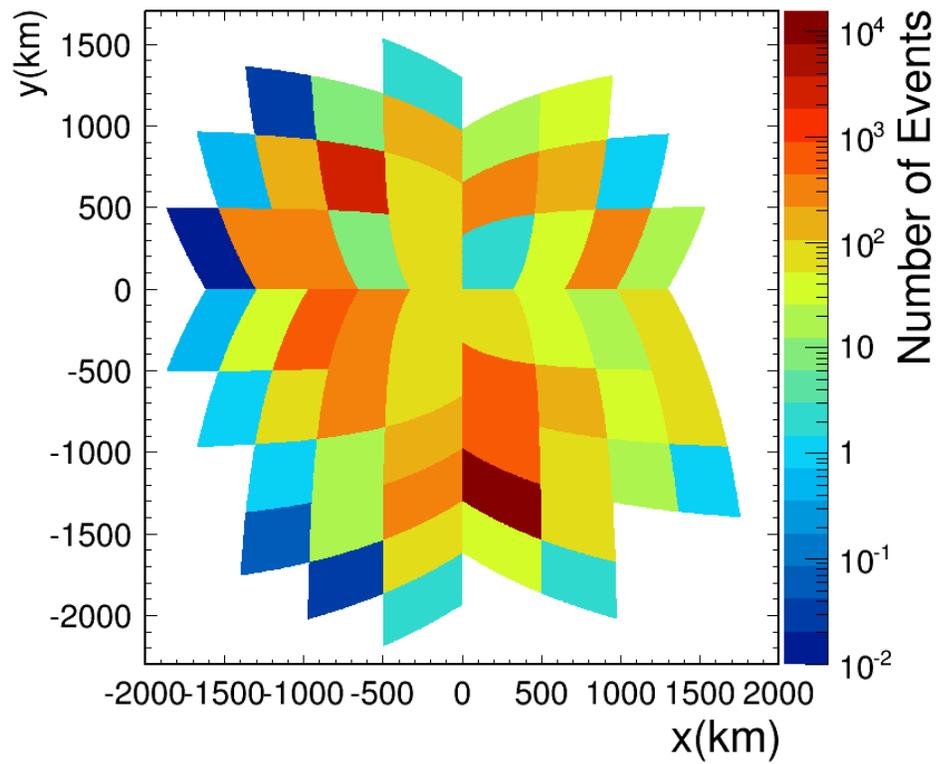


Figure 4.8: The 10% sample set placed into HealPix Bins with fractional weights where necessary.

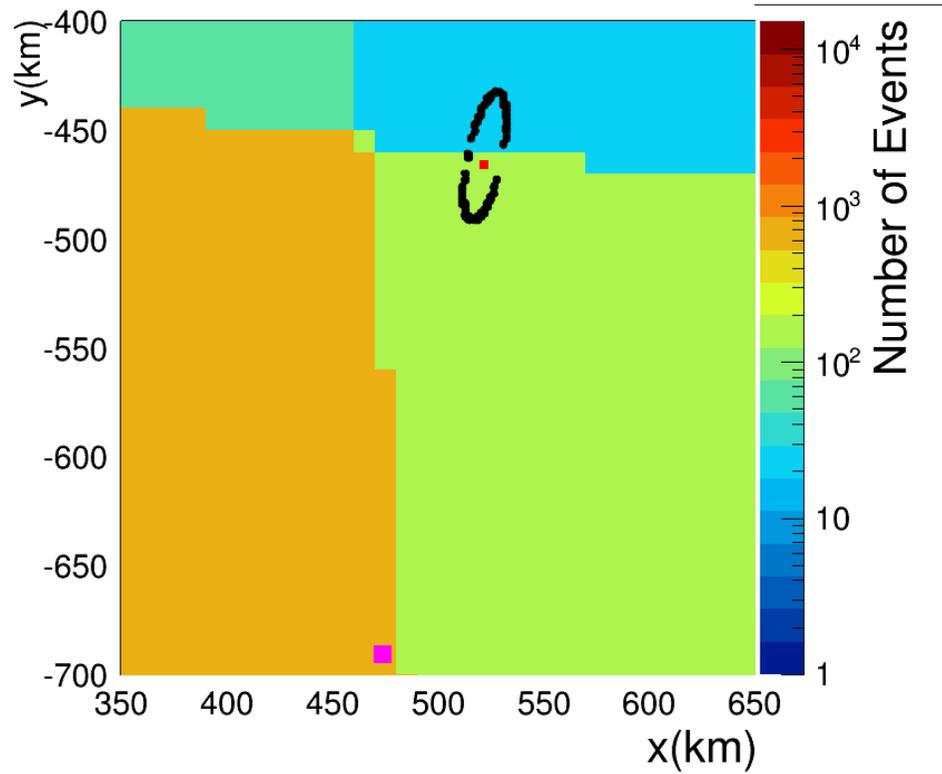


Figure 4.9: An error ellipse shown on top of HealPix bins. The pink square represents the location of ANITA for the event, the red square shows the reconstruction location of the event, with the black dots showing the location of the error ellipse. This event has fractional weights in two bins.

different expected number of background events compared to HealPix bins on quieter ice. Events that pass all cuts will be used to obtain a flux limit.

4.4 Training Samples

To create and optimize our cuts, we needed samples on which to train the analysis cuts. These samples would not be used in the final search. Each sample provides a unique take on the events seen by ANITA.

4.4.1 Taylor Dome sample

As mentioned in Section 2.4.1, there was a calibration pulser at the Taylor Dome station. The events for this sample were pulled from the full data set by a timing cut based on the trigger of the pulser and distance to the payload. These events were flagged for future analyses. This sample helps with reconstruction, such as verifying our algorithms are working, as well as providing information about how CW affects signal.

4.4.2 Thermal Sample

To help understand background, a thermal sample was created. This thermal sample was created from upward pointing (from the payload) events (originating from the sky), from Runs 30-37, 44-52, 72-152, and 168-235. These events were chosen as thermal because we do not expect neutrino signal from the sky. These runs were selected from a time when payload noise was minimal and the payload was out of sight of McMurdo. This sample provides information on pure thermal noise backgrounds. These events were used to create the thermal baselines used to compare to thresholds for continuous wave removal.

4.4.3 10% sample

As described in Section 4.1, we created a 10% sample for understanding background and designing cuts. This sample was designed to evenly sample the flight. This sample is assumed to contain only background events due to the expected low number of signal events. It includes both thermal events and events contaminated with CW. Most importantly, since these events are not upward pointing (like the thermal sample) and not limited to one location (like the Taylor Dome sample), these events can be used to represent backgrounds in HealPix bins for optimizing cuts.

4.4.4 Simulated sample

The ANITA collaboration has a Monte-Carlo simulation for all of its flights, named `icemc`. This simulation is used for many purposes, including finding the sensitivity of the payload. Only recently did the software get upgraded to output the simulated neutrinos into a real data format. This output allows the simulated events to run through the same code as the data from the flight. This allows us to create a simulated neutrino sample that can be used in comparison to the 10% sample, as this sample contains only signal (neutrino) events. This will be used along with the 10% sample for optimizing cuts in the HealPix bins.

Before `icemc` could be used in this manner, we introduced a variety of features to ensure the simulated data was the same as the real data. Firstly, the data output from `icemc` was configured to match the real data format. Secondly, the payload orientation had to be the same in both `icemc` and the analysis code. This involved rotating the payload and its antennas in `icemc` to match the physical location provided by the flight path and heading. Lastly, the phase centers of the antennas had to be adjusted slightly to match the positions used in the analysis code.

Icemc was run with various configurations. The revision used for this analysis was revision-647. Most options were default ANITA-2 options. The energy was set to Kotera max model, which produces a flux of neutrinos following that spectrum. It is currently restricted to energies above 10^{18} eV due to missing parameterization below that energy.

Icemc simulates a neutrino interaction and the signal path through the instrument. Icemc uses the flight path from ANITA-2 when finding events. It randomly chooses a location along the flight path, then it forces a neutrino interaction inside the horizon of the balloon. Each event is given a weight based on the probability of that event interacting at the set location. The signal is then propagated from the interaction location to the payload. The signal chain is simulated and the signal modified accordingly.

One question that needed to be investigated was the trigger efficiency in the simulation. This is difficult to do due to the nature of the simulation. The simulation is able to produce events near the trigger threshold before adding thermal noise to the waveform. After noise has been added, the signal may be hidden underneath the noise. Without a clear signal, the SNR is difficult to define. To determine the trigger efficiency and how it compared to the actual payload, we were required to keep track of where the signal occurred in the waveform. In this, we produced two identical signal samples, one with noise and one without. By having a noiseless sample, we can determine what bins to use when calculating SNR. SNR in this scenario is defined as peak-to-peak voltage over RMS of the first fifth of the waveform in the antenna with the largest voltage peak.

We produced 3 efficiencies to compare against the measured trigger efficiency. The efficiencies are the fraction of events passing the trigger as a function of SNR. The first curve is the most pessimistic. It takes the SNR from the noisy waveform

from the antenna with the largest peak voltage, as was done in analysis 1.0. This is pessimistic because this antenna does not necessarily have signal present in the waveform, i.e. the signal did not occur in that antenna and the waveform is pure noise. The second curve is the most optimistic. It uses the peak-to-peak of the strongest signal-only waveform of the event, over the RMS of the noisy waveform. This may not be the same antenna as was used in the first distribution. This reduces the numerator of SNR since thermal noise no longer contributes to the peak-to-peak voltage, pushing the trigger efficiency curve to lower SNRs. The third curve uses the strongest signal-only waveform to find which time bin the signal peak arrives in. Then it uses those same time bins in the same antenna with the noisy waveform to calculate the peak-to-peak voltage. This is the most realistic due to the fact that the measurements made in lab would know where the signal is located and use that bin to find the peak-to-peak voltage. Figures 4.10 and 4.11 demonstrate how the SNR is calculated in an example waveform. Figure 4.12 shows the trigger efficiencies as a function of SNR. The measured trigger efficiency lies between our most optimistic and pessimistic models. Our third curve, the most realistic, is very close to the measure efficiency, showing that icemc is performing as intended.

Lastly, we attempted to evaluate the efficiency of simulated events at passing the analysis cuts from the previous analysis. This was as simple as creating a set of simulated events and analyzing them with the analysis code. For comparison, we also created the efficiency of the Taylor Dome sample passing the same cuts. The curves can be seen in Figure 4.13. Each colored line represents the efficiency after that cut as been applied in sequence. There is a very clear difference between the simulated sample efficiency and the Taylor Dome sample efficiency. Further advancements are needed in icemc to make this comparison better, such as adding CW-like noise to

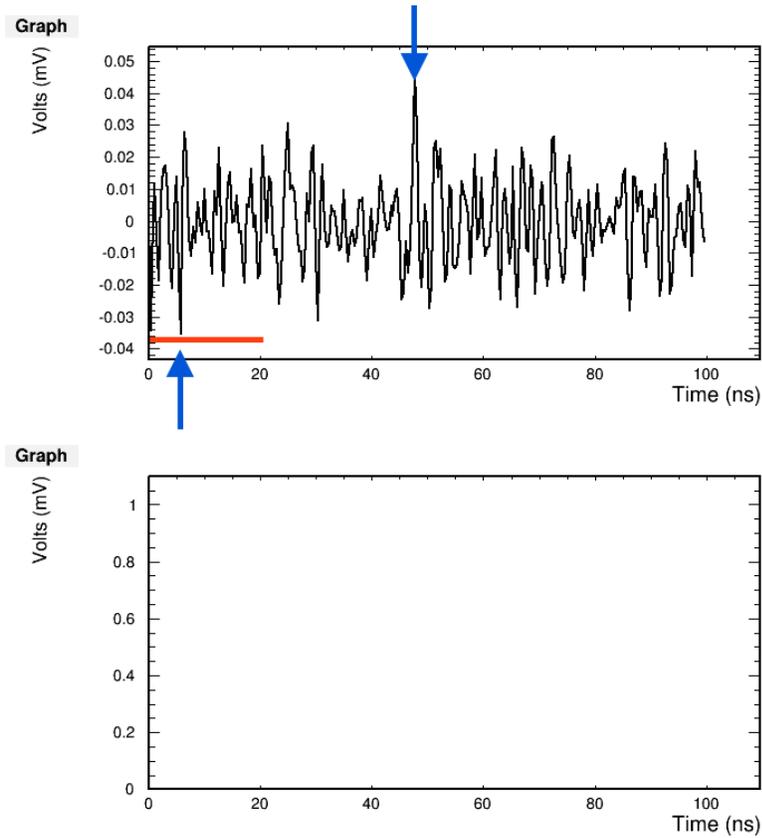


Figure 4.10: The antenna with the strongest voltage peak in the noise sample. Top plot: Noisy waveform. Bottom: Signal only waveform for this antenna. This antenna had the strongest voltage peak, but there was no signal in this antenna, as evidenced by the empty bottom plot. The red line denotes the region used to calculate the RMS for the 1st distribution. The blue arrows denote the voltages that would be used to calculate the peak-to-peak of the 1st distribution.

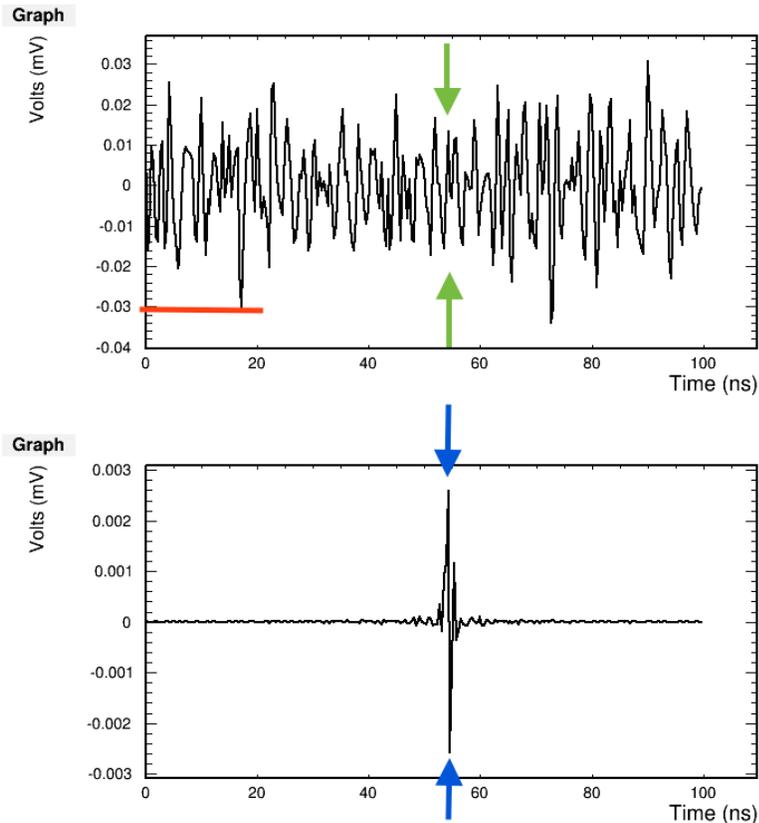


Figure 4.11: The antenna with the strongest voltage peak in the no noise sample. Top plot: Noisy waveform for this antenna. Bottom: Signal only waveform for this antenna. The red line denotes the region used to calculate the RMS for the 2nd and 3rd distributions. The blue arrows designate the voltages used to calculate the peak-to-peak for the 2nd distribution. The green arrows (which occur at the same location in time as the blue arrows) designate the bins used to calculate the peak-to-peak for the 3rd distribution.

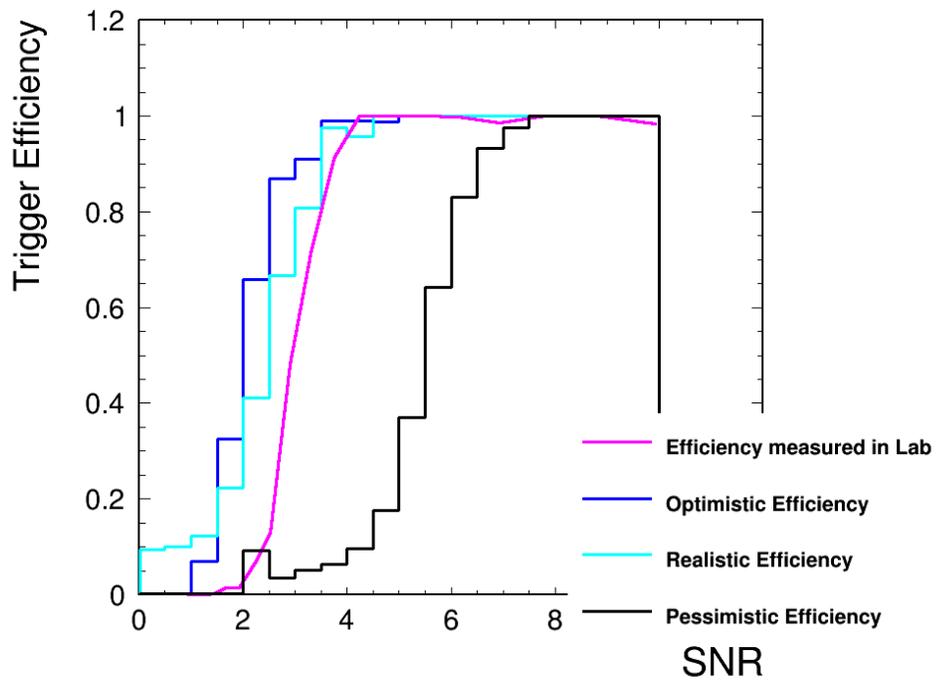
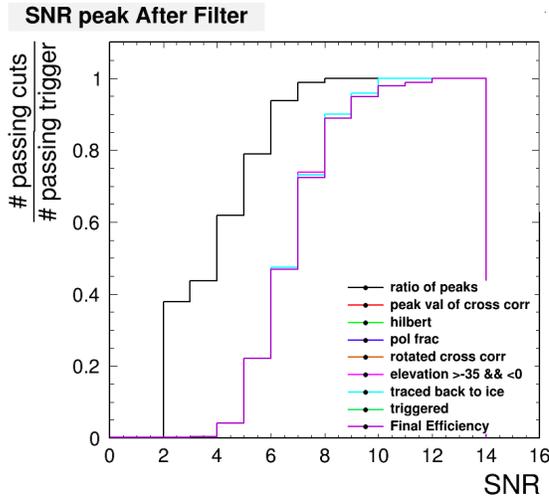
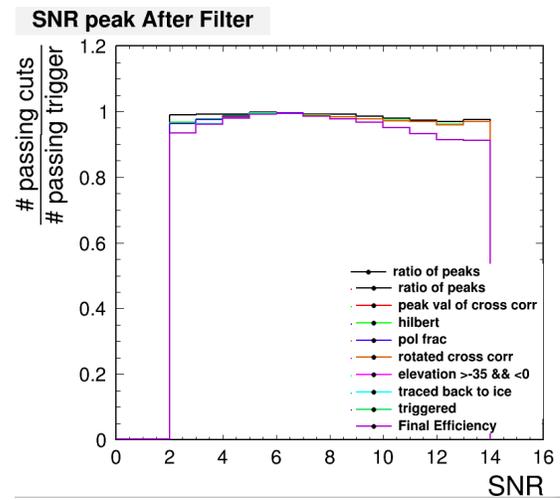


Figure 4.12: The trigger efficiency distributions as a function of SNR. The SNR/1st (black line) and optimistic/2nd (blue line) distributions provide the range of possibilities, while the realistic/3rd distribution (cyan) distribution provide the best guess to match lab measurements (magenta). While not exact, it is very close and shows that icemc is triggering correctly.



(a) Efficiency curves of applying analysis cuts to the simulated sample. Each colored line represents the efficiency after the cut has been applied, in sequence. The final line shows the final efficiency of all analysis cuts.



(b) Efficiency curves of applying analysis cuts to the simulated sample. Each colored line represents the efficiency after the cut has been applied, in sequence. The final line shows the final efficiency of all analysis cuts.

Figure 4.13: The efficiency curves of the simulated sample (left) and Taylor Dome sample (right). The goal of icemc was to have the efficiency of the simulated set nearly match the efficiency of the Taylor Dome sample. It is clear that there are issues to be worked on in the simulation, notably introducing CW noise into waveforms.

simulated events. Since CW adds power and structure to the waveform, CW contamination may help with passing of various analysis cuts. That is why upgrading icemc to mimic real signal is an important task that is still ongoing.

4.5 Quality Events

To create the 10% sample, non-quality event need to be removed. There were two sets of cuts implemented to achieve this quality sample. All events from Runs 12-262 were included in the initial sample, with 26.7M events total. Runs 1-11 were preflight data. Table 4.1 shows how many events were removed by each cut sequentially and

how many were removed by each cut if that cut were last. These cuts reduced the total sample size to 21.3M events.

4.5.1 Non-RF events

The first set of cuts were to remove obvious non-candidate events. Calibration events, forced trigger events, and events where the RFCMs were off were removed by this set of cuts. Each cut is described in detail below:

1. **Non-RF Triggers**: Requires the trigType in the header file be only RF. Removes all software triggers and PPS triggers.
2. **Williams Field Events**(MyCorrelator::isMcMBoreholeOrSeavyFromList): Removes all events that occur within a $2 \mu s$ window of the expected arrival time of a Williams Field calibration pulser event.
3. **Taylor Dome Events**(MyCorrelator::isTaylor): Removes all events within a 600 ns window of the expected arrival time of a Taylor Dome pulser event.
4. **Taylor Dome Reflection Events**(MyCorrelator::isTaylorReflection): There was a reflection between the antenna and pulser of the Taylor Dome pulser system. These events are weaker than the primary signal and arrive at the payload 750 ns after the expected time of arrival for the primary signal. These events are cut using a 900 ns window.
5. **On-Board Calibration Pulser**(MyCorrelator::isCalPulser): This cut removes any event that had a triggerTimeNS < 300 ns while the calibration pulser relay was on, noted by calibStatus variable in the header. Many of these events were also removed by non-RF trigger cut.

	Sequential Cuts	If Last Cut
Total Number of Triggered Events	26,655,876	
Non-RF (Forced) Triggers	5,142,453	3,614,077
Williams Field Events	91,827	81,821
Taylor Dome Events	115,795	115,794
Taylor Dome Reflection Events	354	354
On Board Calibration Pulser Events	3	3
Main RFCM Off Events	30	1
Tiny Peak-to-Peak Events	1418	1418
Number of Events in Analysis Sample	21,349,863	

Table 4.1: Number of events cut by each cut if cuts were applied sequentially and if that cut were applied last.

6. **Main RFCM's On**(MyCorrelator::isMainRFCMOn): This cut removes events where any of the nadir RFCMS were not on, as indicated in the calibStatus variable in the header.
7. **Tiny Peak-to-Peak**(MyCorrelator::isBigEnoughPeaktoPeak): This cut removes any event where the peak to peak of a waveform in any non-nadir antennas was less than 20 mV. This indicates that the RFCMs were not on for that event.

4.5.2 Quality Cuts

To further reduce unwanted events, a second set of cuts is applied. These cuts are meant to remove events that had problems with their data or hardware at that time, such as sync slips, payload noise and other such problems. The reduction to neutrino efficiency is insignificant. After applying these cuts, the total sample size went from 21.3 M events to 21.2 M events, a loss of less than 1% of the data. Table 4.2 shows the number of events cut by each cut sequentially and the number cut if this were the last cut applied. The cuts are describe in detail below.

	Sequential Cuts	If Last Cut
Total Number of Triggered Events	21,349,863	
Sync Slips	18,186	18,182
Bad GPS Events	16,055	16,055
SURF Saturated Events	11,570	73
High DC Offset Events	407	382
Short Traced Events	5	5
Payload Blast Events	45,867	45,867
Number of Quality Event	21,257,773	

Table 4.2: Number of events cut by each cut if cuts were applied sequentially and if that cut were applied last.

1. **Sync Slips**(MyCorrelator::isSyncSlip) A “Sync Slip” is indicated by two variable in the header: surfSlipFlag and errorFlag. These variables indicate a difference in event information between the SURFs, between the TURF and SURF, or between header and event waveforms.
2. **Bad GPS Events**(gpsBadFlag assigned in MyCorrelator::startEachFlag): During the flight, there were times when the ADU5 failed. While some of this missing information ($\sim 5\%$) was recovered using sun sensor data or the SIP GPS, the rest of data is unusable. If the closest GPS packet was further 30 seconds away, or if the nearest GPS packet was corrupt (non-physical values), the event was removed.
3. **SURF Saturated Events**(MyCorrelator::isChannelSaturated): The SURF digitization would become saturated if the peak voltage of the waveform was greater than 1.5V. If three or more channels in HPol or VPol were saturated, the event was removed. If less than three channels were saturated, the event was still used, but saturated channels were excluded from the analysis.

4. **High DC Offset Events**(MyCorrelator::isDCOffsetLarge): If an event had a mean voltage >100 mV, this indicates a problem with digitization and the event is removed. This usually occurs when the instrument is triggering at a high rate.
5. **Short Traced Events** (MyCorrelator::isShortTrace): If a waveform has less than 240 points, the event is removed.
6. **PayloadBlastEvents** (MyCorrelator::isPayloadBlast): This cut removes events in which 15 channels or 9 phi-sectors triggered with a peak voltage >400 mV. This points to payload noise that triggers many antennas as signal is expected to come from a single direction.

4.5.3 Background Cuts

To achieve the best limit, we must reduce the number of background events in the 10% sample while keeping as many events from the simulation sample. A set of cuts were designed in previous analysis [55] that were taken a starting point for this analysis. If an event passed these cuts, it would be placed into the correct HealPix bin based on reconstructed source location. The values for the cuts were decided by the distributions of the 10% sample and the simulated sample. An event is considered placed if it meets the following requirements:

1. **Ratio of Peaks:** The interferometric map used for reconstruction direction must have one peak to which to reconstruct to. The strength of the 2nd strongest peak to the 1st strongest peak of the interferometric map must be <0.85 . Figure 4.14 shows the distribution of this ratio for the 10% sample and the simulated sample. This cut was designed to remove misreconstructions.

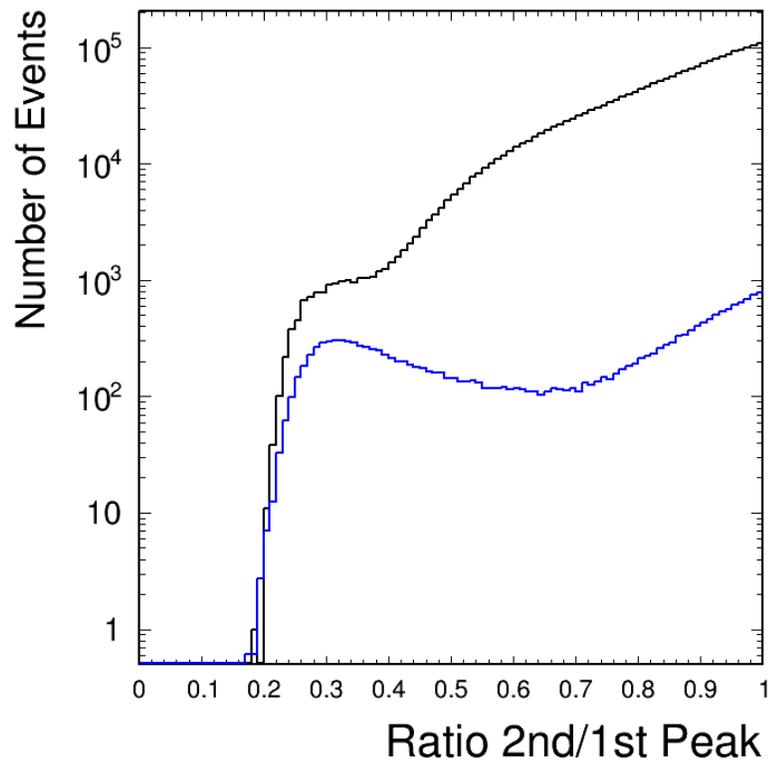


Figure 4.14: Distribution of 10% sample (black) and simulated sample (blue) for Ratio of Peaks variable.

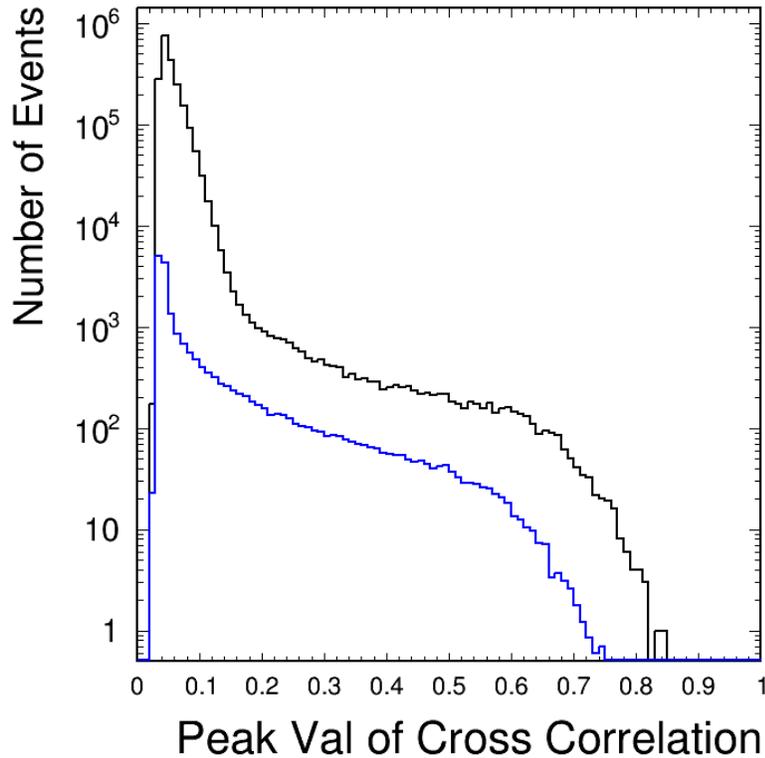


Figure 4.15: Distribution of 10% sample (black) and simulated sample (blue) for peak val variable.

2. **Peak Value:** The strength of the peak in the interferometric map must be >0.075 . This cut is designed to remove events that do not correlate strongly, such as thermal noise and onboard payload noise. Figure 4.15 shows the distribution of this parameter for the 10% sample and the simulated sample.
3. **Peak of Coherent Hilbert Envelope:** A Hilbert transform and envelope was done (`FFTtools::getHilbertEnvelope`) on the coherently summed waveform. The peak of the envelope must be > 15 mV. The distribution of envelope peaks for the 10% sample and simulated sample are shown in Figure 4.16

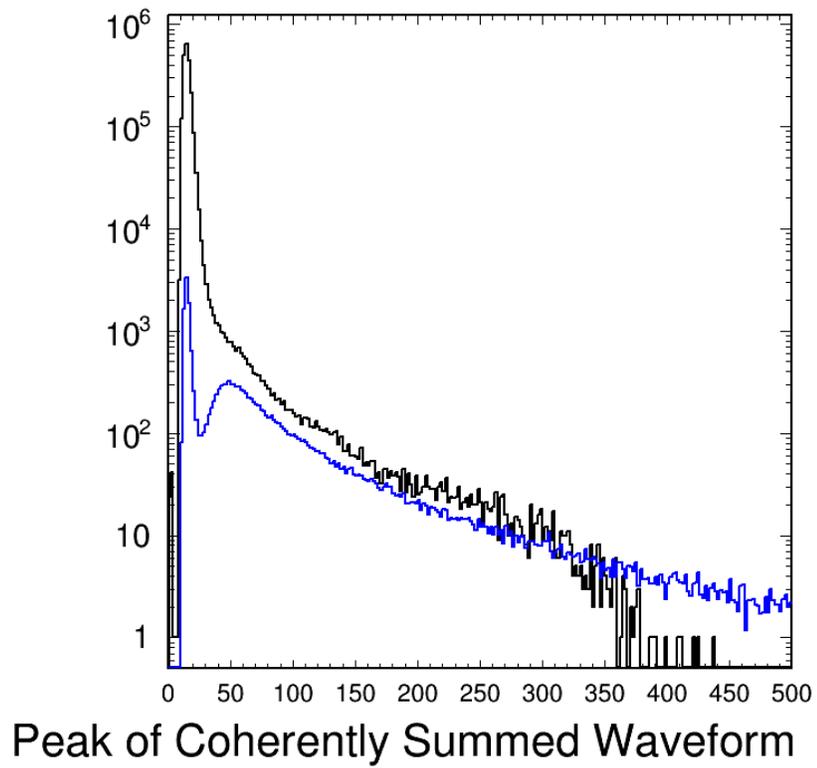


Figure 4.16: Distribution of 10% sample (black) and simulated sample (blue) for peak of Hilbert envelope.

4. **Polarization Cut:** Neutrino events are expected to be linearly polarized. By using Stokes parameters, the fraction of linearly polarization can be determined.

$$P_{\text{frac}} = \frac{\sqrt{Q^2 + U^2}}{I} \quad (4.2)$$

where

$$Q = \langle V_h^2 \rangle - \langle V_v^2 \rangle \quad (4.3)$$

$$U = 2\langle V_h V_v \rangle \quad (4.4)$$

$$I = \langle V_h^2 \rangle + \langle V_v^2 \rangle \quad (4.5)$$

V_{rmh} is the time-dependent voltage in HPol, V_{rmv} is the time-dependent voltage in VPol, and $\langle \rangle$ is the time average. This parameter plotted for the 10% sample and the simulated sample can be seen in Figure 4.17.

5. **Elevation Angle Cut:** Signals should be in the main viewing area of the antennas. $-35^\circ < \theta < -6^\circ$ This cut also help remove misreconstructions which tend to misreconstruct in the θ angle. It also removes events that reconstruct above the horizon as we expect neutrino signals to come from the ice. Figure 4.18 shows the distribution for simulated and 10% samples.
6. **Trace to Continent:** Events are expect to reconstruct to the continent. If the event does not originate from Antarctica, it is removed.
7. **Hardware Trigger Cut:** The event must have at least a L2 trigger in the direction of reconstruction (± 1 phi-sector).

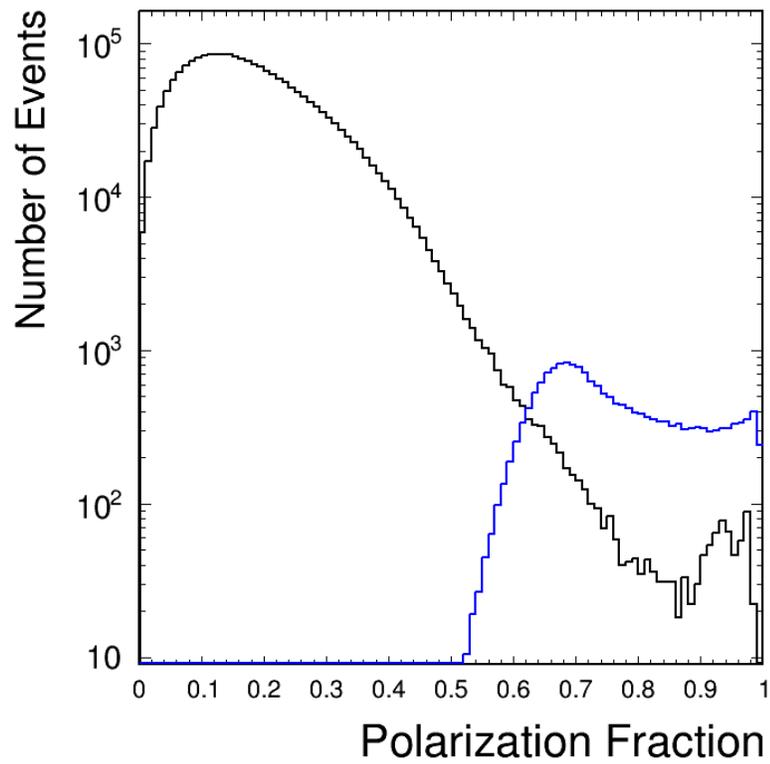


Figure 4.17: Distribution of 10% sample (black) and simulated sample (blue) for polarization fraction variable.

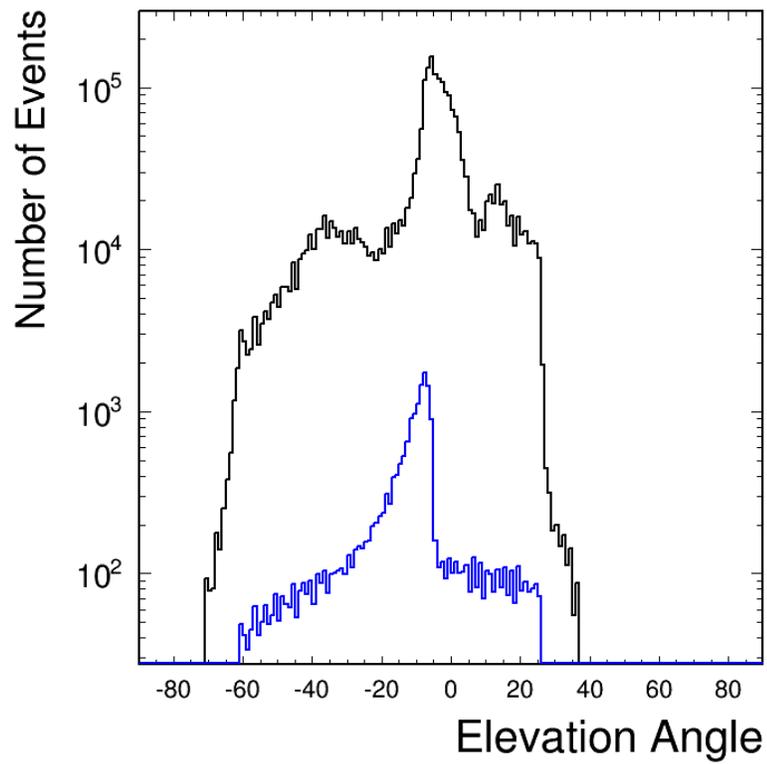


Figure 4.18: Distribution of 10% sample (black) and simulated sample (blue) for peak of elevation angle variable.

8. **Bad Noise Cut:** The SNR of the event is defined as the largest power in 1/5th of the waveform over the noise of the first 1/5th of the waveform. If the signal region lies too close to the start of the waveform, the noise sample size is reduced to ensure the signal sample is not inside the noise sample. If the noise sample is reduced below 10 ns of time, it starts to cause an artificial increase of SNR, skewing the optimization process. Events where the signal occurs early in the waveform tend to be thermal.

After having applied the background cuts to the 10% sample and the simulated sample, we placed the remaining events into the correct HealPix bins for optimization. A cut table can be seen in Table 5.2, where all cuts have been applied.

Chapter 5

Optimization Procedure and Results

We have discussed the techniques and tools needed to perform this analysis. we present the optimization procedure used for this analysis and list the results of the optimization the form a list of cuts that can be applied to the 90% sample.

5.1 Optimization Procedure

We chose to optimize the Rotated Cross Correlation Cut. we also chose to maximize our sensitivity to the maximal Kotera Model [1]. This rotated cut is a cut on the relationship between two variables, the SNR of the coherently summed waveform and the peak value of the cross correlation. An example of this distribution can be seen in Figure 5.1. The rotated cross correlation cut will be optimized on a HealPix bin-by-bin basis. The rotated cross correlation cut is performed by choosing a slope and a y-intercept for a line in this 2D space and removing any events that fall below this line. The slope and y-intercept needed to be determined.

First, we focused on finding a good slope value for the cut line. We want to find a slope of the slanted line that works for many HealPix bins. To do this, we first created differential plots for each HealPix bin. The differential plot is created by choosing a slope for the cut line, then slowly increasing the y-intercept value by steps

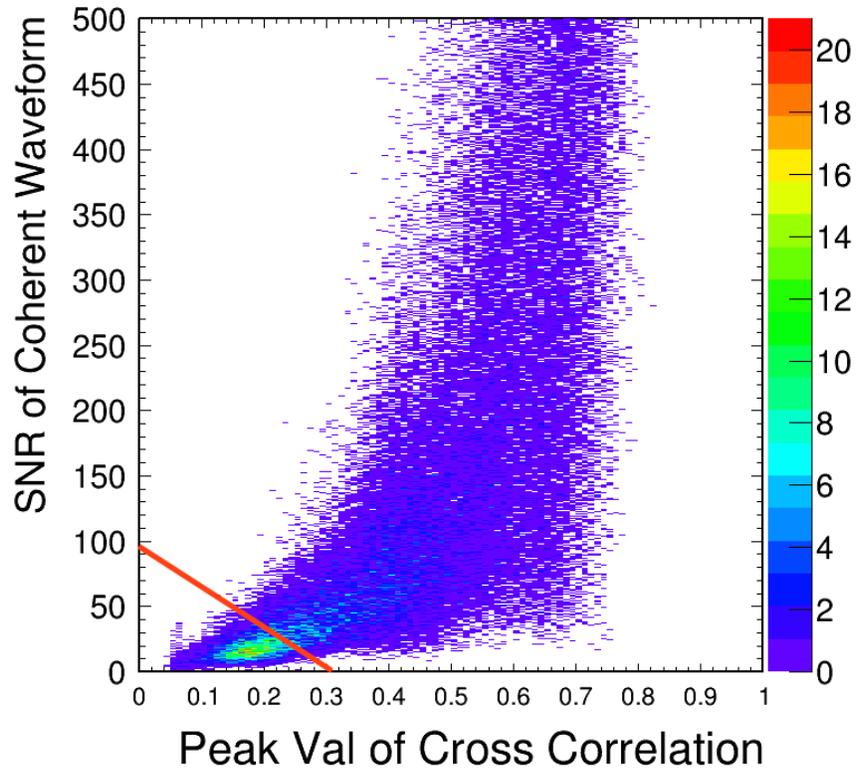


Figure 5.1: Distribution of the Taylor Dome sample for the Rotated Cross Correlation Cut. The red line represents an example of a possible cut for this distribution. Anything to the left of the red line would be removed, while anything to the right passes. Both the slope and the y-intercept value of this line need to be determined by optimizing on 10 % and simulated samples.

of 1. An example of the 2D space and slanted line can be seen in Figure 5.2. As the y-intercept value increases, we recorded the number of events cut that were removed by that step. Figure 5.3 shows an example of this differential plot. An exponential line was fitted to the falling spectrum of the differential plot and the log-likelihood value ($-2 \text{ Log}(L)$), L_0 was found from the fitting program (option 'WL' from [61]). The fit is performed from the first bin after the peak in the differential plot to the last filled bin.

To find the p-value, we created many pseudo-experiments based on the fitted line. The number of events in the fitted region in each pseudo-experiments varied based on Poisson statistics. In each pseudo-experiment, a new differential plot was created and a line was then fit to the region. For every HealPix bin, 10,000 pseudo-experiments were created. The log-likelihood from each pseudo-experiment was recorded to create a log-likelihood distribution, $F(L)$ (Figure 5.4). To find the p-value, an integral was performed on this distribution from L_0 to infinity and normalized to the the total integral of the likelihood distribution. This p-value was recorded for every HealPix bin.

$$\text{p - value} = \frac{\int_{L_0}^{\infty} F(L)dL}{\int_0^{\infty} F(L)dL} \quad (5.1)$$

These p-values were used as an indicator of whether the slope for the rotated cross correlation cut was a good choice.

To choose the slope of the rotated cross correlation cut, we wanted a flat distribution in the p-value distribution. Any outliers (HealPix bins with a near 0 p-value) had to be understood and dealt with. After testing many slopes, we finally chose a slope of -38 1/mV . The p-value distribution for this choice of slope can be seen in Figure 5.5.

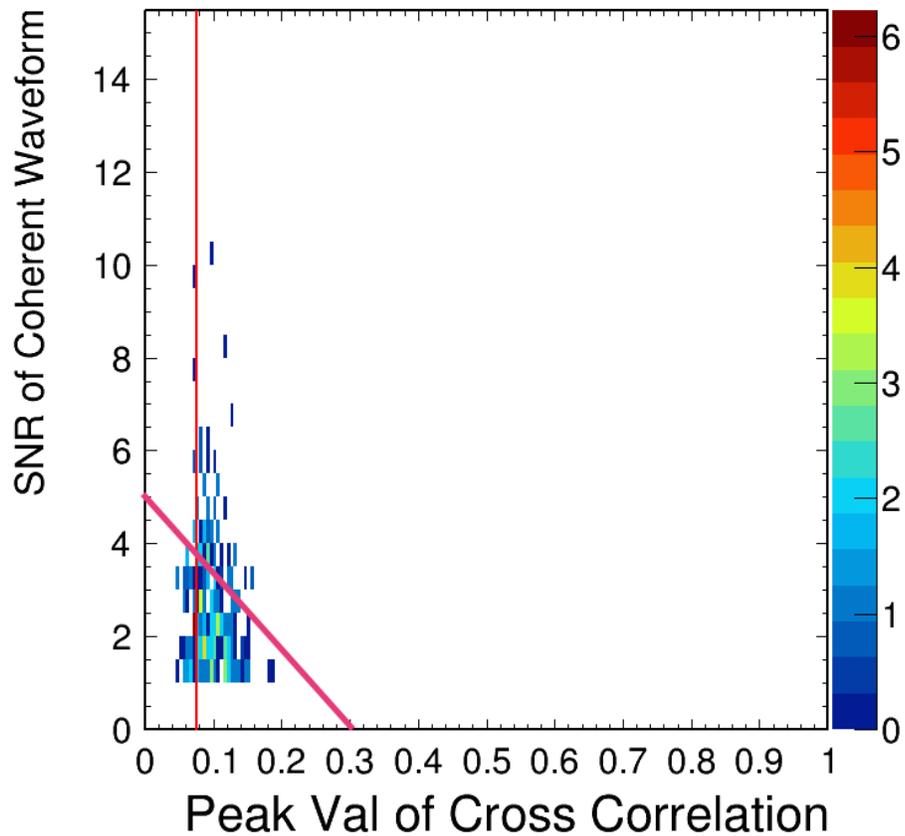


Figure 5.2: An example of the rotated cross correlation scatter plot. This is for HealPix bin 3012. The vertical line represents a previous analysis cut that would remove any events to the left of the line. The slanted line is an example of rotated cross correlation cut. The slope ($1/\text{mV}$) and y-intercept of this line must be optimized.

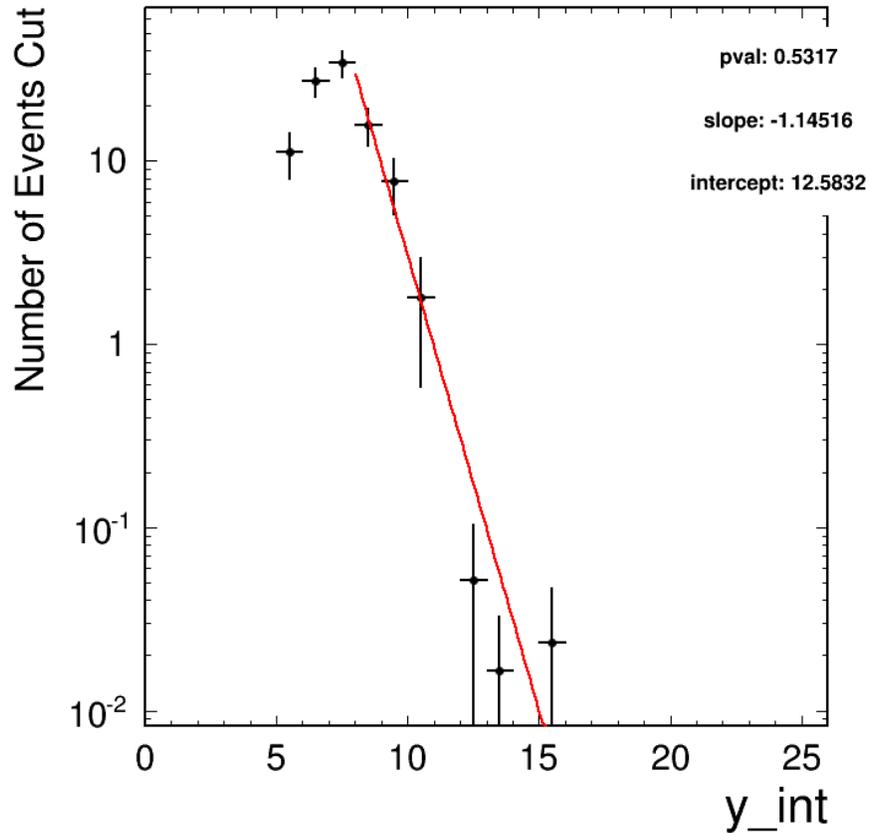


Figure 5.3: The differential plot for HealPix bin 3012. The x-axis is the y-intercept value of the slanted line show in Figure 5.2, with the y-axis showing how many weighted events were removed by moving the slanted line to that y-intercept value. This differential plot was created using a slope of -38 1/mV for the slanted line on the 2D scatter plot. The red line on this plot represents the exponential fit to the falling slope, with parameters seen in the statistics box. The fit parameters follow the equation $f(x) = e^{\text{intercept} + \text{slope} * x}$

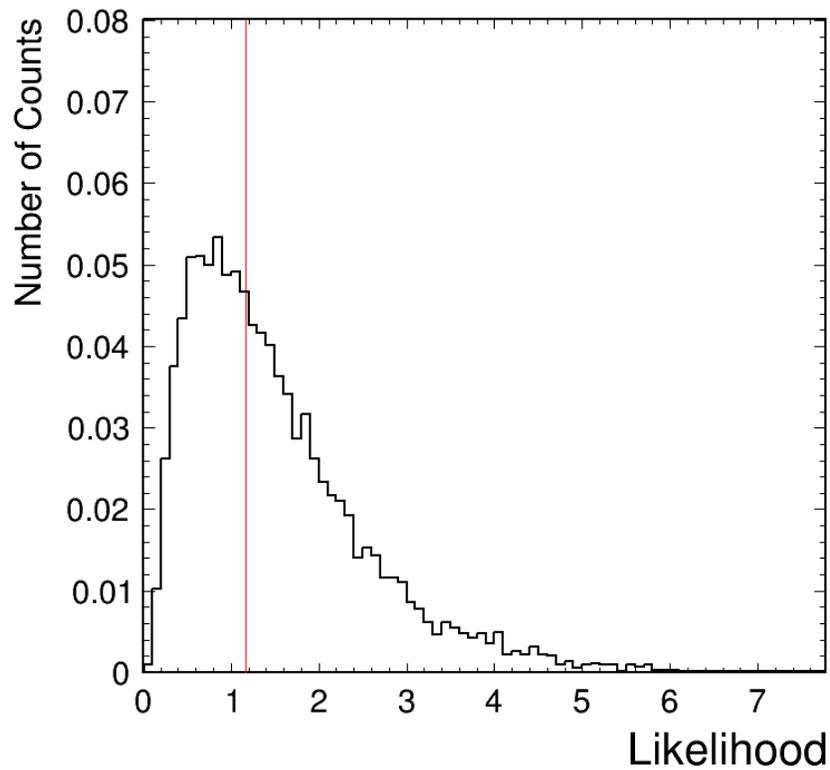


Figure 5.4: The likelihood value for HealPix bin 3012. This likelihood is based on using a slope of -38 when creating the differential plot. The p-value is the integral of the distribution from the red line to infinity.

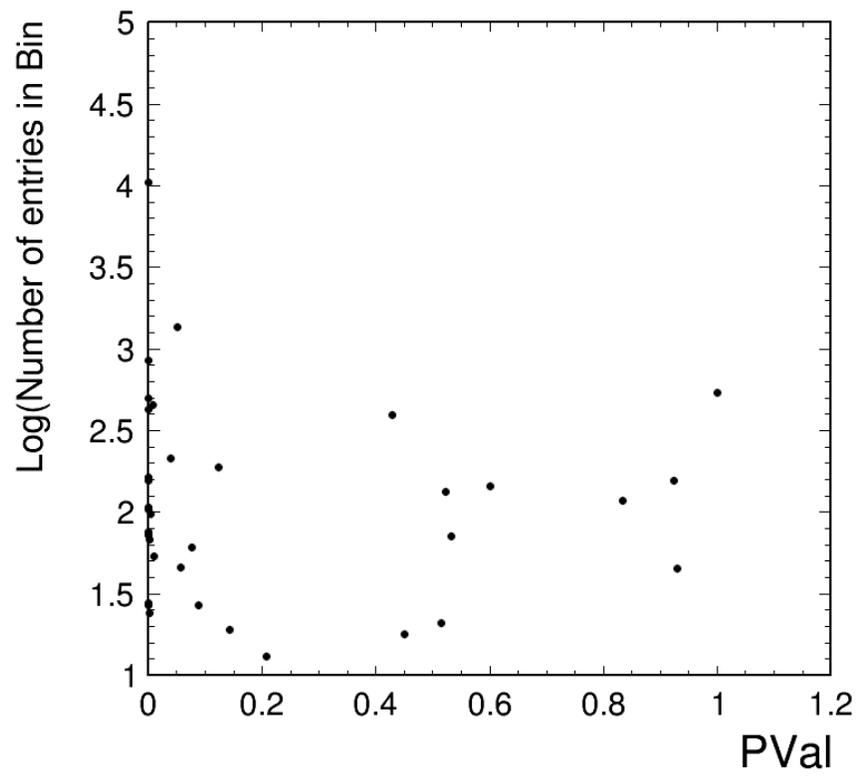


Figure 5.5: The p-value distribution of the chosen slope, -38 1/mV .

Figure 5.5 allows us to decide which HealPix bins to include in our search. Any HealPix bin with a p-value < 0.05 was removed from the analysis. This was due to large background contamination in those bins that would skew any results. With a fit so low, the differential plot does not have an exponential decay on the falling edge. Secondly, we required the fit to contain 4 bins in the differential plot to have accurate results. If any HealPix bin did not have the required number of fitted bins, it was also removed from the analysis. The final set of bins used in the analysis is seen in Table 5.1.

After choosing a slope, each HealPix bin that was included in the analysis was optimized by finding the y-intercept value for that bin. For a Poisson distribution the likelihood to detect events is

$$L(n|s) = \frac{(s + b)^n}{n!} e^{-(s+b)} \quad (5.2)$$

where $L(n|s)$ is the likelihood to detect n events with s mean signal events expected and b is the number of background events. We expect zero signal in the 10% sample due to the low flux of neutrinos, so we set $b = n$ in Equation 5.2.

To find this optimized cut value, we maximized the variable $\frac{\text{Signal}}{S_{\text{up}}}$. Signal is defined as the number of signal events (simulated neutrinos) surviving the rotated cross correlation cut with the given y-intercept value. S_{up} is defined as the maximum number of signal events allowed such that $(S_{\text{up}} + b)$ can fluctuate down to b events 10% of the time.

$$1 - \alpha = \frac{\int_0^{S_{\text{up}}} \frac{(s+b)^b}{b!} e^{-(s+b)} ds}{\int_0^{\infty} \frac{(s+b)^b}{b!} e^{-(s+b)} ds} \quad (5.3)$$

where α is 0.10 (90% confidence level). The number of background events, b , was determined by integrating the fit line on the differential plot from the cut value to infinity and multiplying by 9. This factor was to account for the larger number of

event that would be present in the 90% sample and the corresponding shape of the Poisson distribution. The variable $\frac{\text{Signal}}{S_{\text{up}}}$ was plotted as a function of y-intercept value, seen in Figure 5.6. The maximum from this distribution was chosen as the cut value for that HealPix bin. An example of the rotated cross correlation distribution with an optimized cut can be seen in Figure 5.7.

Refer to Appendix B for plots pertaining to each HealPix Bin. The bins placement in relation to one another can be seen in Figure 5.8.

5.2 Final Cuts

After finalizing the cut parameters, we can see the power of each cut. In Table 5.2, the number of events removed by that cut in sequence and if it was the sample defining cut are listed for the 10% sample.

An additional cut was placed unto the sample after the optimization. This cut is to strengthen the evidence for any possible candidate in the 90% sample. We used fractional areas based on the error ellipse to perform the optimization. When applying the cuts, we use the source location to identify the HealPix bin the event originates from. If an event passes all cuts, then the fractional area is obtained and tested against the corresponding HealPix bins. The event is removed from the analysis if more than 50% of the error ellipse is removed by this last step. This cut removes events where the majority of the error ellipse lies in a bin that it would be cut by, but a fraction of the error ellipse survives in another bin due to a different cut value.

After all cuts were applied, we found zero events to pass from the 10% sample.

A full comparison against the previous analysis is not possible without performing the analysis on the 90% sample. However, we can look at the efficiency of the cuts against the simulated sample created. While not the final efficiency, this will give us insight into how well the new techniques are performing compared to the old

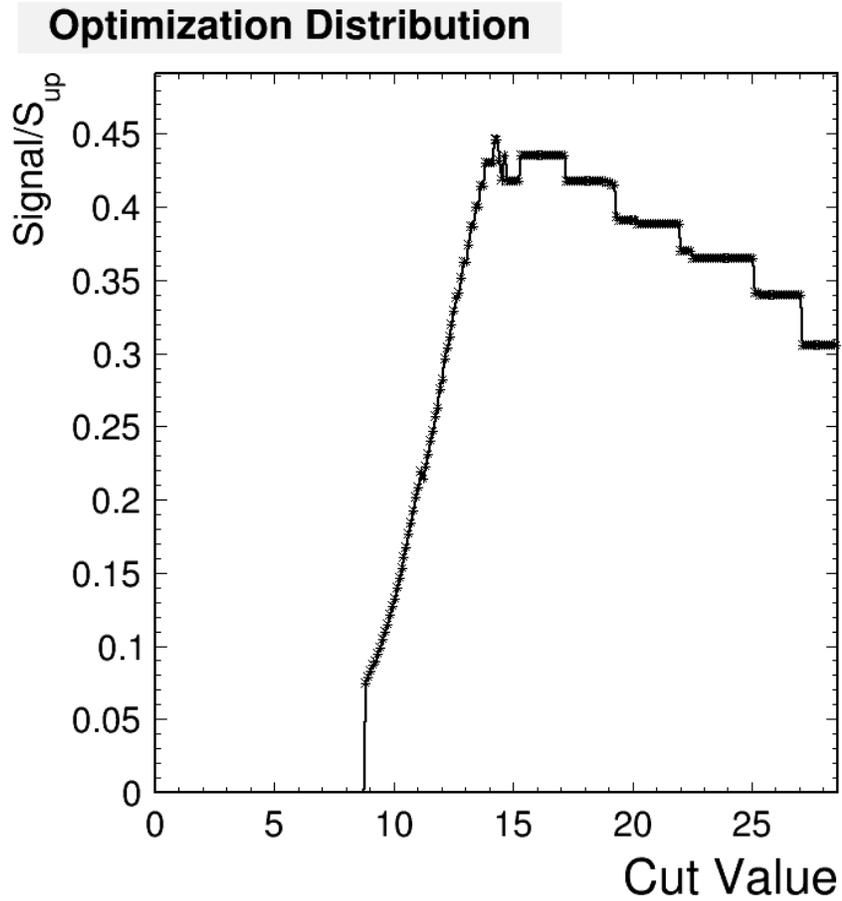


Figure 5.6: $\frac{\text{Signal}}{S_{\text{up}}}$ as a function of y-intercept for HealPix bin 3012. The y-intercept value is 14.25. The flat regions that occur in this optimization show that the size of the simulated sample was not large enough, as there were not enough simulated events in this HealPix bin to have a smooth drop-off.

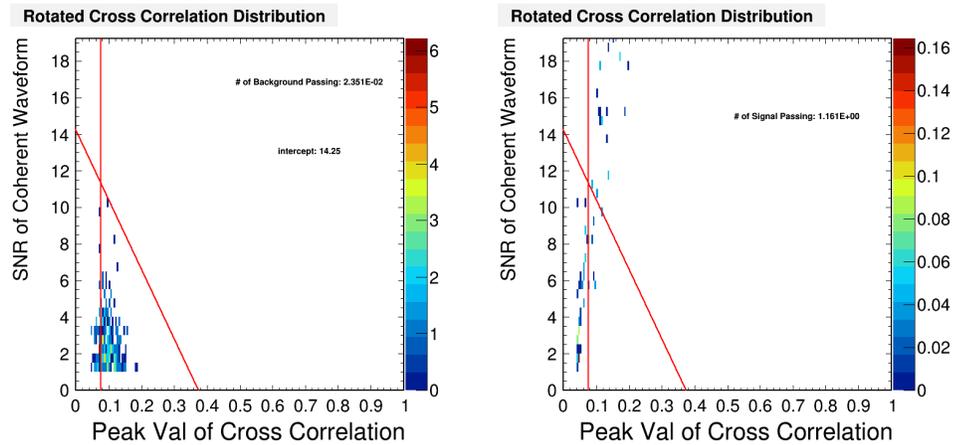


Figure 5.7: Left: The rotated cross correlation distribution for HealPix bin 3012. The slope is -38 1/mV , with a y-intercept of 14.25. There is a second cut on the peak value of the cross correlation. Any events to the left of the red vertical line or below the red slanted line are cut. Right: The rotated cross correlation distribution for simulated events for HealPix bin 3012. The optimized cut line placed onto the figure as well.

techniques. Figure 5.9 shows the efficiency of the previous analysis cuts plotted against each HealPix bin efficiency. The plot was constructed this way in an attempt to minimize that fact that clustering was not yet performed, an integral part of the previous analysis. Only after analyzing the 90% sample can we determine the amount of ice removed in the old analysis compared to the new techniques.

Figure 5.9 demonstrates a few key concepts. Several HealPix bin (3032, 3045, 3053, 3057) are as efficient as the old analysis method. Other bins are much less efficient, possibly due to the low number of events present in those HealPix bins or large backgrounds in those bins. This plot shows that while the HealPix method works, it may not be the optimal solution at this time. This method was developed for the case where the number of events on the ice becomes large and the clustering method begins to remove large quantities of ice from the analysis. That limit may

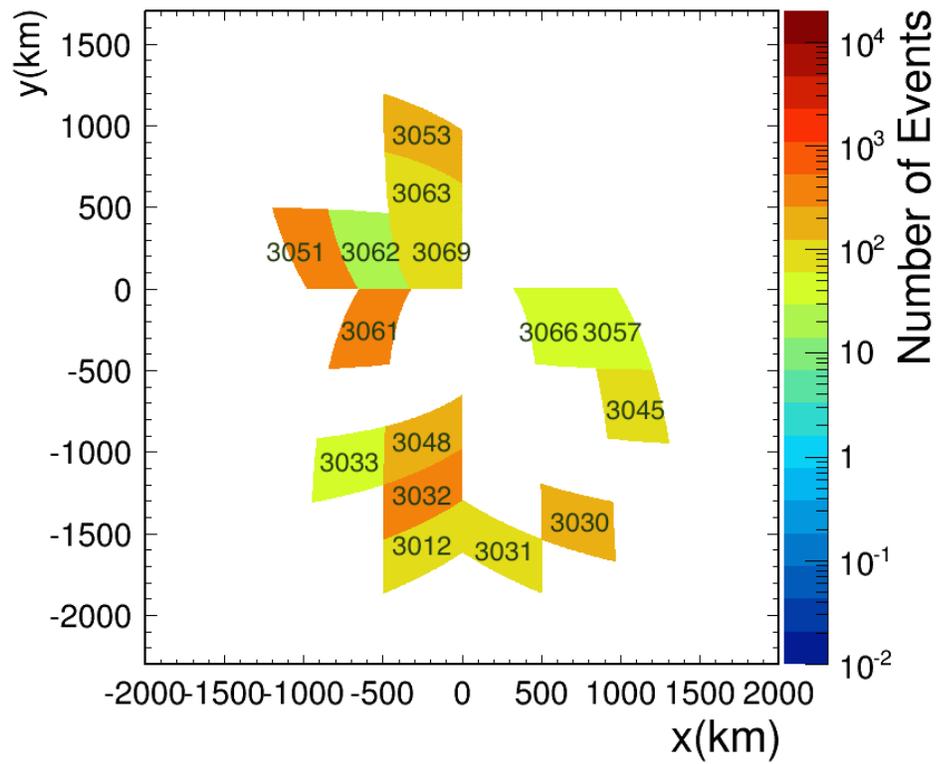


Figure 5.8: The HealPix bins used in this analysis, numbered for easy reference.

not have been reached with the current data set. However, with various tweaks this method may perform better. Developing a technique to combine HealPix bins with few events together such that those bins may be used in the analysis will allow us to keep valuable ice. Changing the size of the HealPix bin can allow for finer tuning of which bins to remove from the analysis, possibly resulting in better efficiencies.

HealPix Bin	y-intercept value	Expected Background	Sensitivity
3012	14.25	0.188	1.161
3030	17.45	0.131	141.933
3031	22.55	0.131	3.101
3032	20.45	0.057	92.240
3033	23.65	0.131	124.602
3045	15.15	0.058	191.469
3048	18.05	0.060	125.190
3051	141.95	0.412	12.581
3053	15.25	0.056	89.652
3057	22.35	0.134	152.722
3061	16.85	0.058	38.615
3062	17.05	0.208	52.820
3063	18.85	0.132	65.662
3066	13.95	0.057	24.200
3069	13.55	0.194	4.802

Table 5.1: A list of the included HealPix bins for this analysis. Y-intercept values are listed for the optimized cuts, and the expected background for the 90% sample for each bin is also listed. The last column demonstrates the number of weighted simulation events passing the cuts in that bin, giving an idea of which bins are the most sensitive to neutrinos.

	Sequential Cuts	If Last Cut
Number of Events in 10% sample		2,102,094
Number Remaining after removing HealPix Bins		194,799
Ratio of Peaks Cut	140,751	0
Peak Value of Cross Correlation	44,961	2
Peak of Hilbert Envelope	923	0
Polarization Fraction	4,501	4
Rotated Cross Correlation Cut	3,494	2,507
Elevation Cut	6	0
Traced Cut	0	0
Triggered Cut	71	19
Nadir Noise	0	0
Bad Noise	90	90
< 50% area in passing Bin	2	2
Number of Events Passing Cuts		0

Table 5.2: Number of events cut by each analysis cut if cuts were applied sequentially and if that cut were applied as the last cut with the rest of the order unchanged. The large amount of zeros in the right most column indicates redundancy, but the cuts are listed due to removing events before the optimization process.

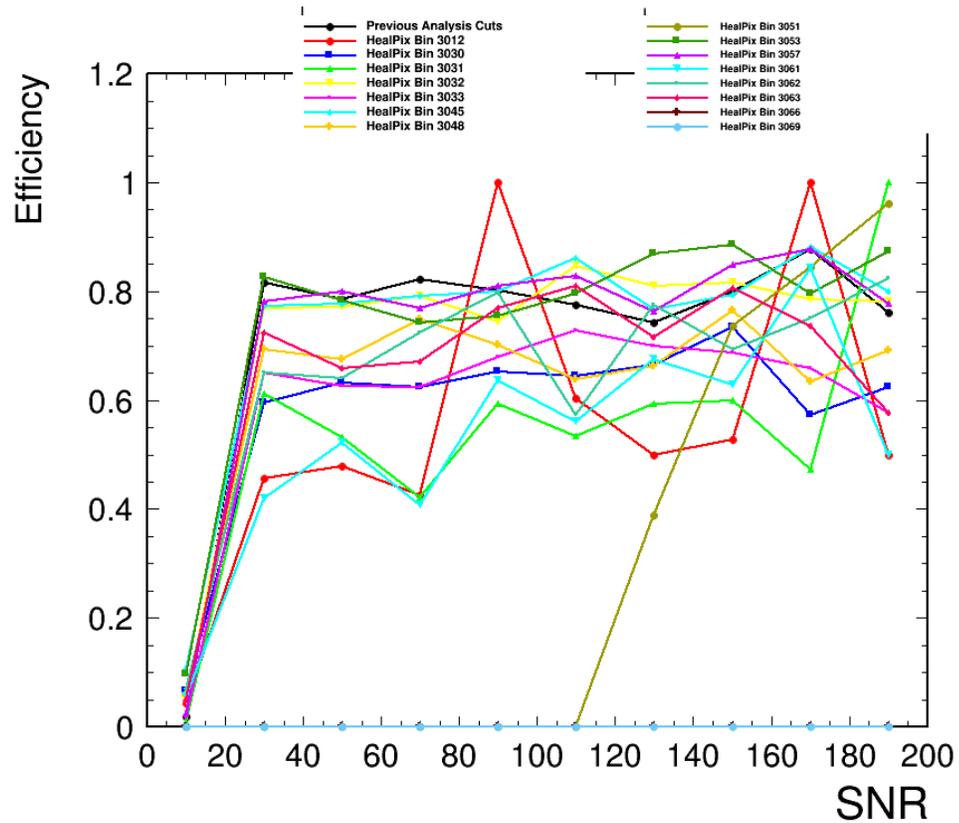


Figure 5.9: The efficiency of the analysis cuts on the simulated sample (Kotera model). The previous analysis is the black line, while the HealPix bins used in this analysis are the various colored lines. A few of the HealPix bins (3032, 3045, 3053, 3057) are about as efficient as the old analysis method, but other HealPix bins are worse. Once the 90% sample has been run, a final efficiency which includes the clustering cut of the previous analysis, can be performed.

Chapter 6

Conclusions

In this thesis, we described new techniques that can be used for future analyses. These techniques were developed and optimized using the ANITA-2 data set. With that data and a simulated set of neutrinos, cuts were optimized to prove these techniques are viable.

In Chapter 3, we discussed techniques for filtering of waveforms. We presented and tested various techniques for filtering the amplitude spectrum of events. We also presented and tested new techniques to adjust the phase spectrum of events. We tested these techniques with a sample of calibration pulsers from Taylor Dome to determine which filters give the best reconstruction.

We discussed the techniques and tool used in this analysis in Chapter 4. We demonstrated a new normalization for cross-correlation formula. Also, we talked about the upgrades to the Monte Carlo simulation, `icemc`. These upgrades allowed us to use simulated events in the same format as real data. We discussed a new technique of binning the ice using a program called `HealPix` that allows us to optimize the search for different areas of ice. This is necessary as future flight of ANITA will improve sensitivity at the trigger level, and the old method will have problems with finding candidate events due to the nature of the algorithm. Combined with this new

binning, we also weighted the events by the fraction of area that lies in each HealPix Bin. This is necessary as event reconstruction has errors.

We optimized the analysis with respect to maximal Kotera Model by optimizing the rotated cross correlation cut. We discussed the parameters of the optimization and why values were chosen. We provided an estimate of the background for the 90% sample. After optimization and apply all cuts, we found zero events passed from the 10% sample.

Having received the blessing to open the box, the next step would be to apply these optimized cuts to the 90% sample. After opening the box, the events inside would need to be analyzed. If no candidate events are observed, a flux limit could be produced.

The techniques developed here can be further expanded upon. The sine subtraction technique holds potential, but various bugs must be worked out before it proves as useful as the geometric method. The HealPix binning can be easily modified by adjusting the size of the bins. Smaller binning can provide more control of which HealPix bins to remove from the analysis due to excessive noise, but must be balanced by the fact that the fitting model needs some number of events to function correctly. Using icemc to produce real data format events is still new and must be brought along carefully to ensure no problems with analysis code. Further features can be incorporated to match data, such as CW.

Future flights of ANITA will provide more events and better sensitivity at the hardware level. These filtering techniques will help remove CW background more efficiently while providing better reconstruction. The new algorithm for binning, with further refinement, can keep more ice available for the neutrino search, resulting in a discovery.

Appendix A

Calculation of Error Ellipse

The error on the reconstruction to the ice on an event turns out to be a non-trivial problem. Following the procedure shown in [62], we can attempt the solution. The error on the reconstruction forms an elliptical cone, a cone that is defined by 2 angles.

$$\frac{x^2}{\tan^2(\theta)} + \frac{y^2}{\tan^2(\phi)} = z^2 \quad (\text{A.1})$$

This elliptical cone intersects a sphere (Earth) and that intersection defines the shape of the error on the event. Choosing the cone apex to be placed at the origin, with the z-axis along the line of sight of the cone, the cone has equation seen in equation A.1. This sphere is offset by a distance of the radius of the Earth (R) and the height of the experiment(H), where the cone makes an angle η with respect to the line connecting the origin to the center of the sphere.

$$[x + (R + H)\sin\eta]^2 + y^2 + [z - (R + H)\cos\eta]^2 = R^2 \quad (\text{A.2})$$

Substituting y^2 from eq A.1 into eq. A.2 yields

$$\left[1 - \frac{\tan^2\phi}{\tan^2\theta}\right]x^2 + 2(R + H)\sin\eta x + \sec^2\phi z^2 - 2(R + H)\cos\eta z + (R + H)^2 = R^2 \quad (\text{A.3})$$

To use eq. , we need one point on the intersection. By setting $x = 0$, we can get two values for z , where the cone enters the earth along the yz plane and where the cone exits the earth along that plane. This choice gives an equation for z :

$$z = (R + H) \cos\eta \cos^2\phi \left[1 \pm \sqrt{1 - \sec^2\phi \frac{(R + H)^2 - R^2}{(R + H)^2 \cos^2\eta}} \right] \quad (\text{A.4})$$

Choosing the value of z where the cone enters the Earth (minus solution), we can use eq [A.1](#) and the choice of $x=0$, to get two values of y .

$$y = \pm \tan\phi z = \pm \tan\phi (R + H) \cos\eta \cos^2\phi \left[1 \pm \sqrt{1 - \sec^2\phi \frac{(R + H)^2 - R^2}{(R + H)^2 \cos^2\eta}} \right] \quad (\text{A.5})$$

We now have 2 points of the intersection where the cone enters the sphere $(0, y_{\min}, z)$ and $(0, y_{\max}, z)$. We can perform an iterative process where we vary the y -coordinate from y_{\min} to y_{\max} and get the x and z coordinates. Since these y -coordinate values were derived at the zero x coordinate, we need an equation that relates x and y . Solving Eq. [A.1](#) for z and plugging into Eq. [A.2](#) and solving we end up with a quartic equation.

$$\begin{aligned}
a_4 &= \csc^4\theta \\
a_3 &= 4(R + H) \csc^2\theta \sin\eta \\
a_2 &= 2y^2 \csc^2\theta \csc^2\phi + 2 \csc^2\theta [(R + H)^2 - R^2] + 4(R + H)^2 \sin^2\eta \\
a_1 &= 4y^2 \csc^2\phi (R + H) \sin\eta + 4(R + H) \sin\eta [(R + H)^2 - R^2] \\
a_0 &= y^4 \csc^2\phi + 2y^2 [(R + H)^2 - R^2] \csc^2\phi - \frac{4y^2 (R + H)^2 \cos^2\eta}{\tan^2\phi} + [(R + H)^2 - R^2]^2 \\
0 &= a_4 x^4 + a_3 x^3 + a_2 x^2 + a_1 x + a_0
\end{aligned} \tag{A.6}$$

While it is possible to solve a quartic equation analytically, it is much easier to use a numerical method. By looping over the y -values from y_{\min} to y_{\max} , we are able to get the two x -values that correspond to the correct points, and using Eq. A.1 we can derive the z -coordinate of those points.

Now, we have the intersection points in the coordinate system where the cone apex lies at the origin. This information needs to be transformed to a coordinate system with the center of the sphere at the origin. Following along now with [62], we first need to translating the origin to the sphere center and reversing the z -direction.

$$\begin{aligned}
x' &= x + (R + H) \sin\eta \\
y' &= y \\
z' &= -z + (R + H) \sin\eta
\end{aligned} \tag{A.7}$$

We will then rotate the coordinates about the y axis such that the new z-axis lies along the line from the center of the Earth to the apex of the cone.

$$\begin{aligned}
 x'' &= x' \cos \eta - z' \sin \eta \\
 y'' &= y' \\
 z'' &= x' \sin \eta + z' \cos \eta
 \end{aligned}
 \tag{A.8}$$

To get the latitude and longitude of these intersection points involves another set of rotations. The angle from north the cone apex makes (α) and the latitude ϕ_s of the balloon position must be known. Rotating about the z'' axis by an angle of α gives

$$\begin{aligned}
 u &= y'' \sin \alpha + x'' \cos \alpha \\
 v &= y'' \cos \alpha - x'' \sin \alpha \\
 w &= z''
 \end{aligned}
 \tag{A.9}$$

followed by a rotation about the v-axis by an angle of $90^\circ - \phi_s$

$$\begin{aligned}
 u' &= u \cos \phi_s - w \sin \phi_s \\
 v' &= v \\
 w' &= u \sin \phi_s + w \cos \phi_s
 \end{aligned}
 \tag{A.10}$$

This places the w' -axis is along the earth's axis and the u' - v' plane lies along the equator. Getting the latitude ϕ and longitude λ follow as

$$\begin{aligned}\sin\phi &= w'/R \\ \sin\lambda &= \frac{v'}{\sqrt{u'^2 + v'^2}}\end{aligned}\tag{A.11}$$

where the sign of λ is determined by the signs of u' and v' . The true longitude is then calculated by adding the longitude of the cone apex (the experiment) to the calculated value.

Appendix B

Plots for HealPix Bins

In this appendix, we show plots for each HealPix bin that was kept during the optimization process. These plots include the differential plot of that bin with a fit line, the optimization parameter distribution, the Signal and S_{up} distributions, and finally the rotated cross correlation cut with optimized cut values.

In some differential plots, there are outliers to the right of the main distribution. These events are contamination from another HealPix bin. Most of the area of the error ellipse for these events occurs in a different bin, but due to the fractional weighting process, it shows up in this bin. Misreconstructions into the wrong bin can cause an event to be mis-labeled as a candidate event due to less strict cuts. This cross-talk between bins during optimization helps reduce the possibility.

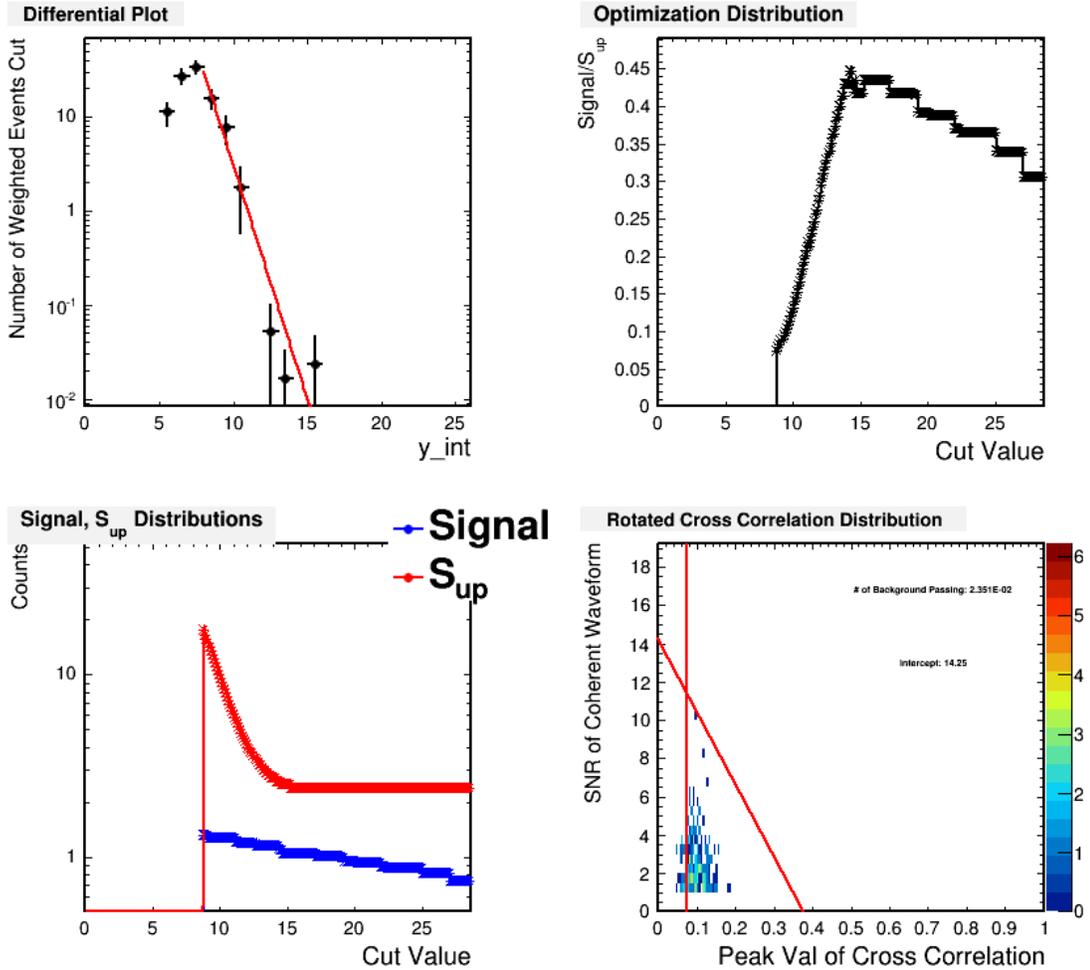


Figure B.1: The optimization plots for HealPix bin 3012. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

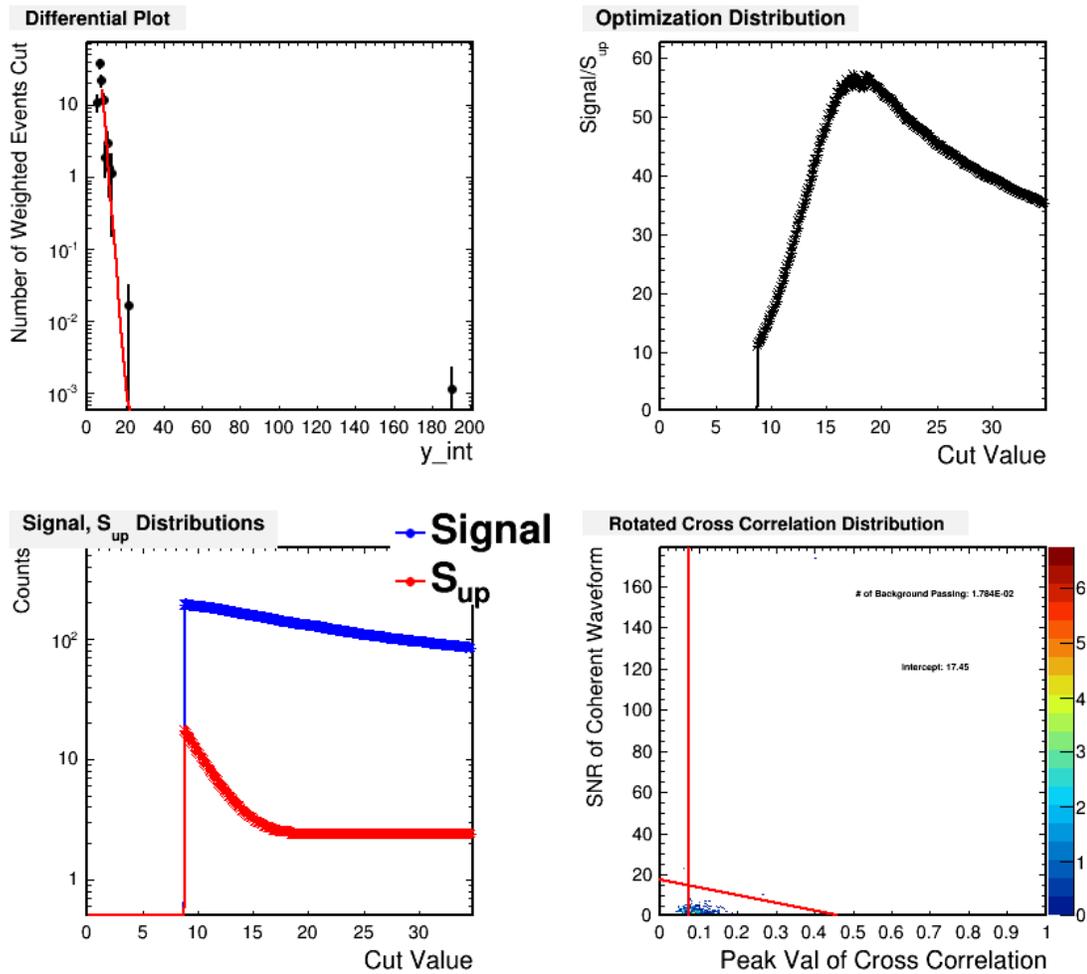


Figure B.2: The optimization plots for HealPix bin 3030. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

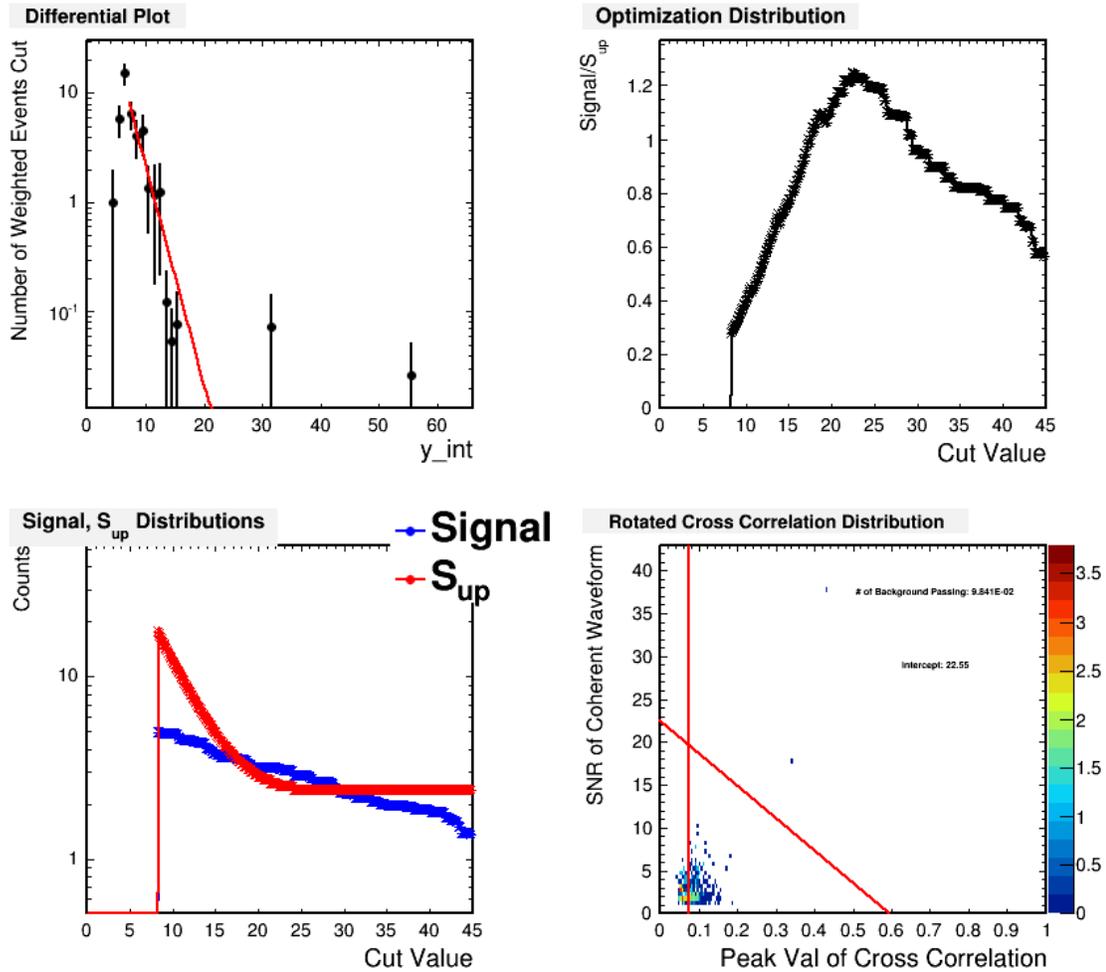


Figure B.3: The optimization plots for HealPix bin 3031. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

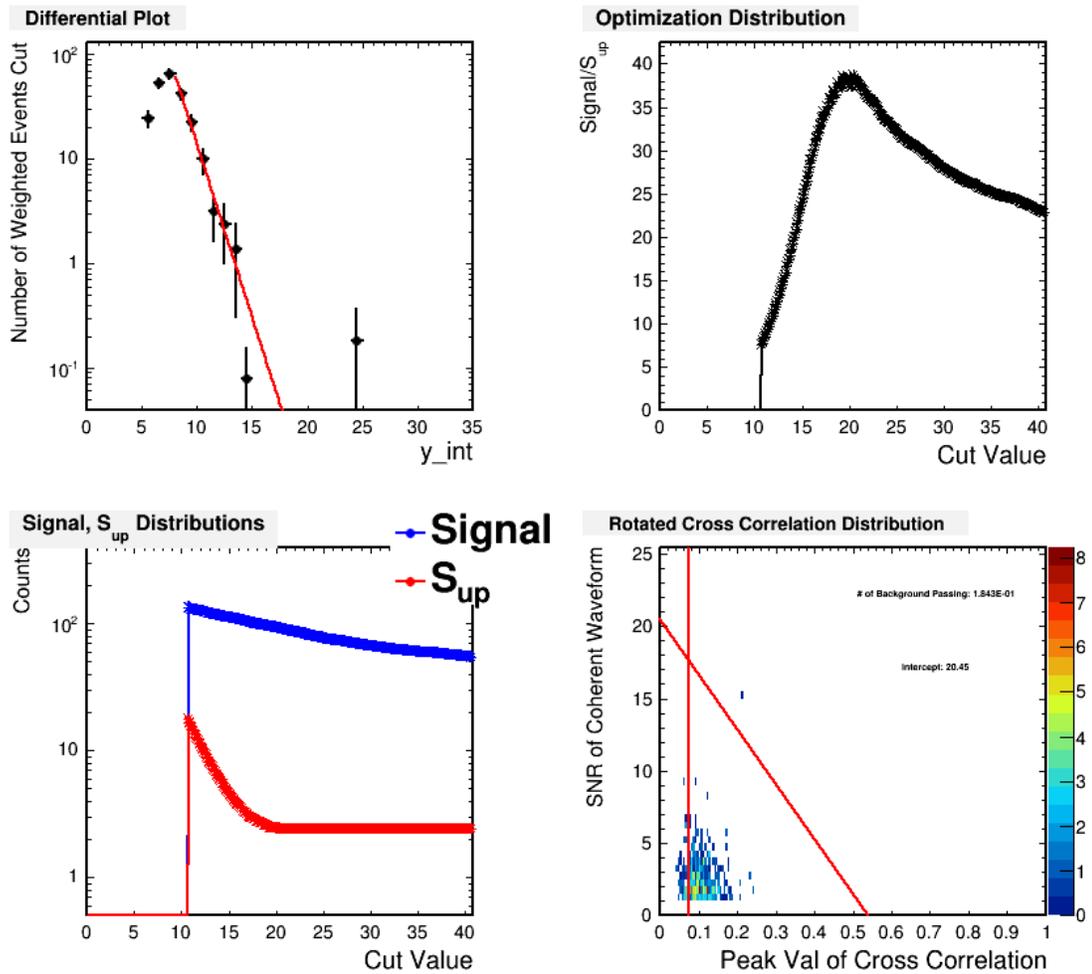


Figure B.4: The optimization plots for HealPix bin 3032. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

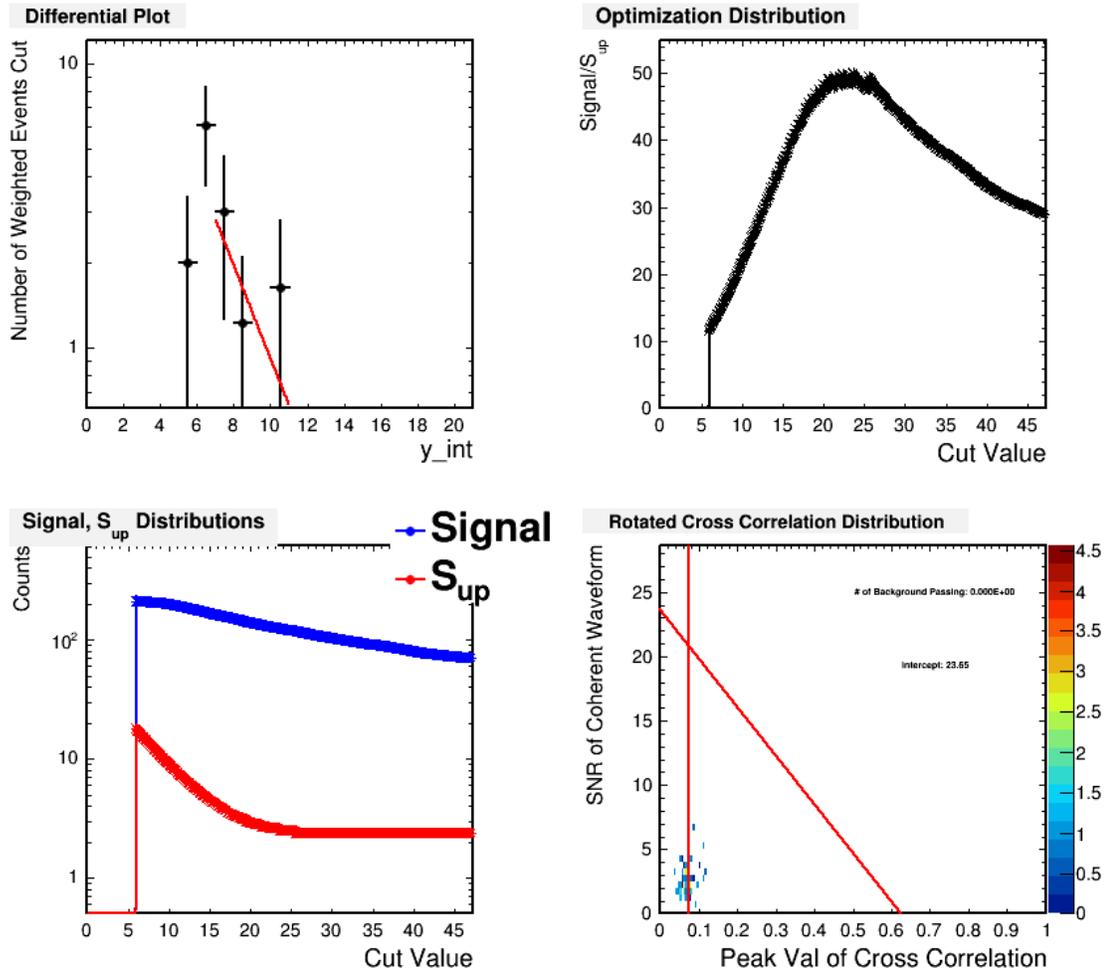


Figure B.5: The optimization plots for HealPix bin 3033. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

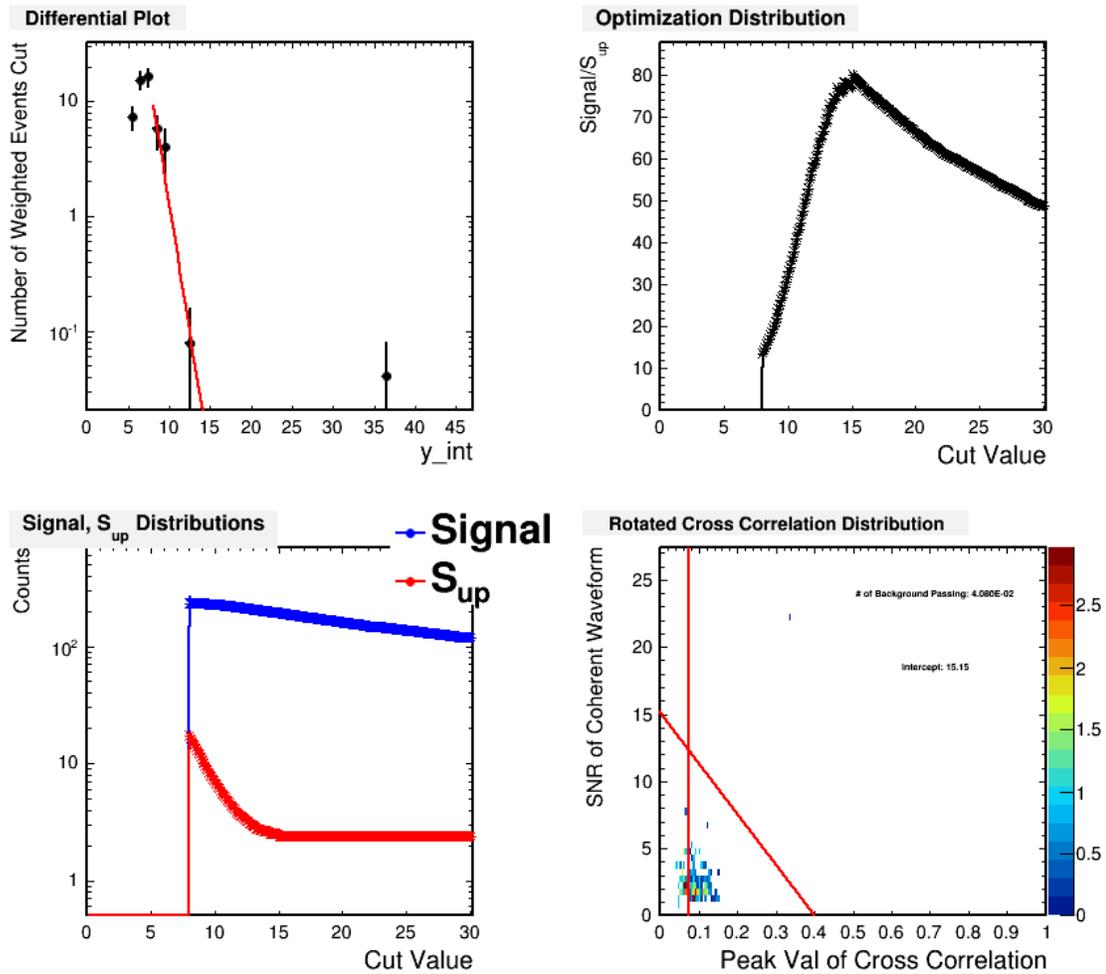


Figure B.6: The optimization plots for HealPix bin 3045. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

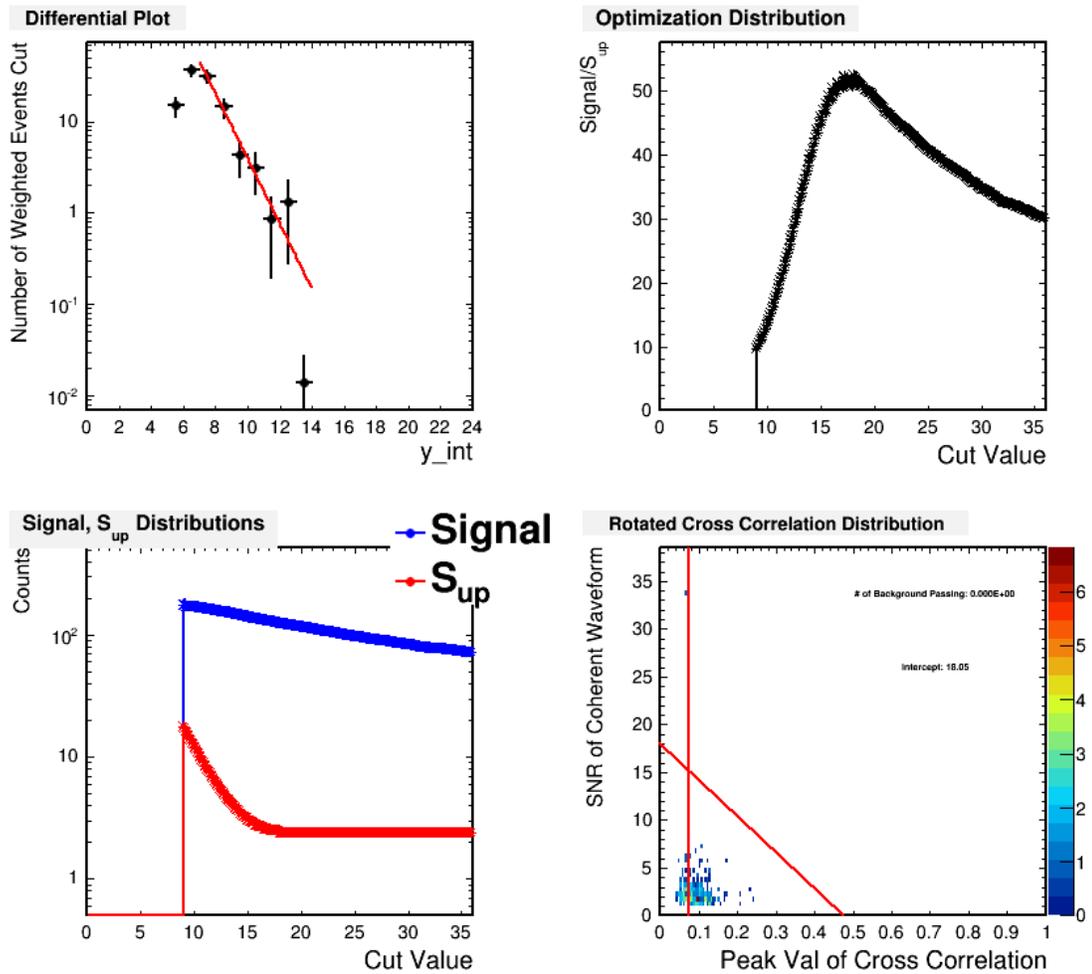


Figure B.7: The optimization plots for HealPix bin 3048. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

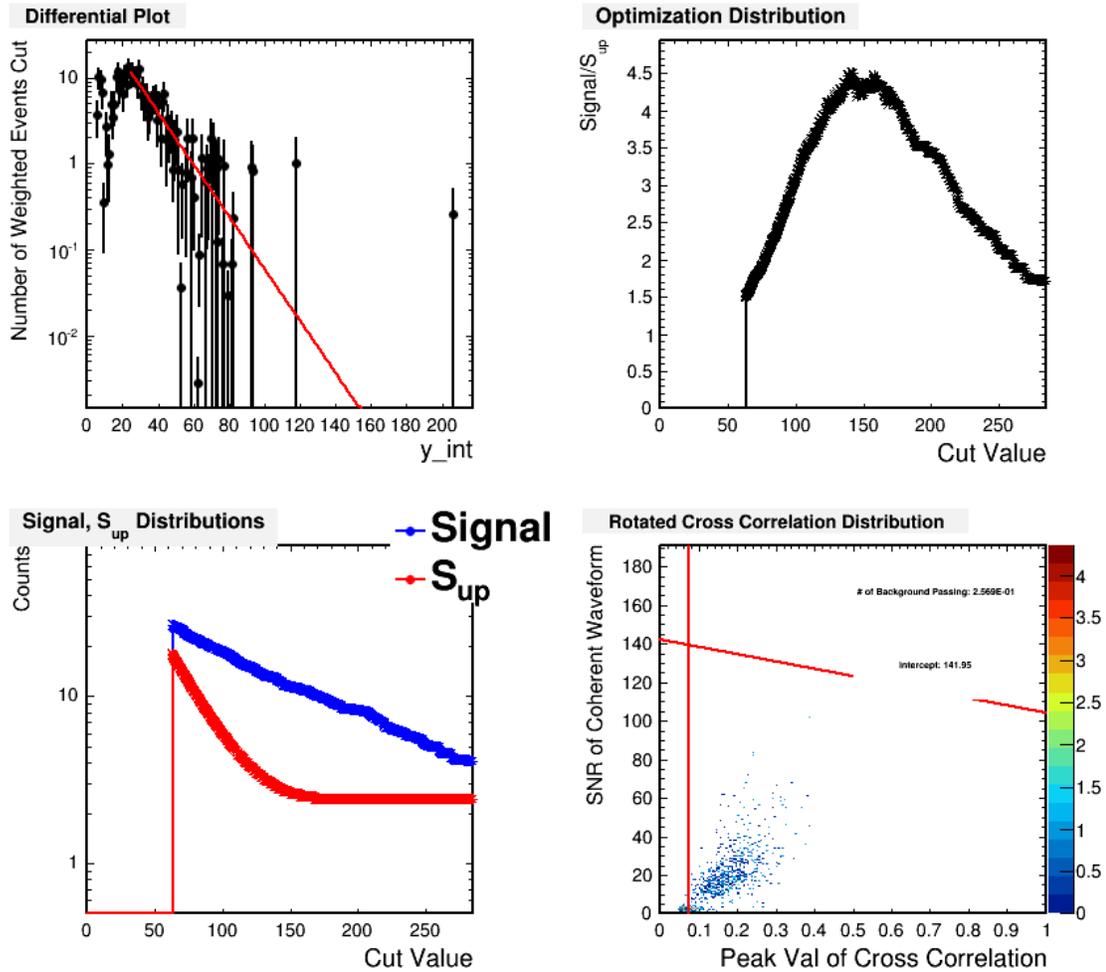


Figure B.8: The optimization plots for HealPix bin 3051. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

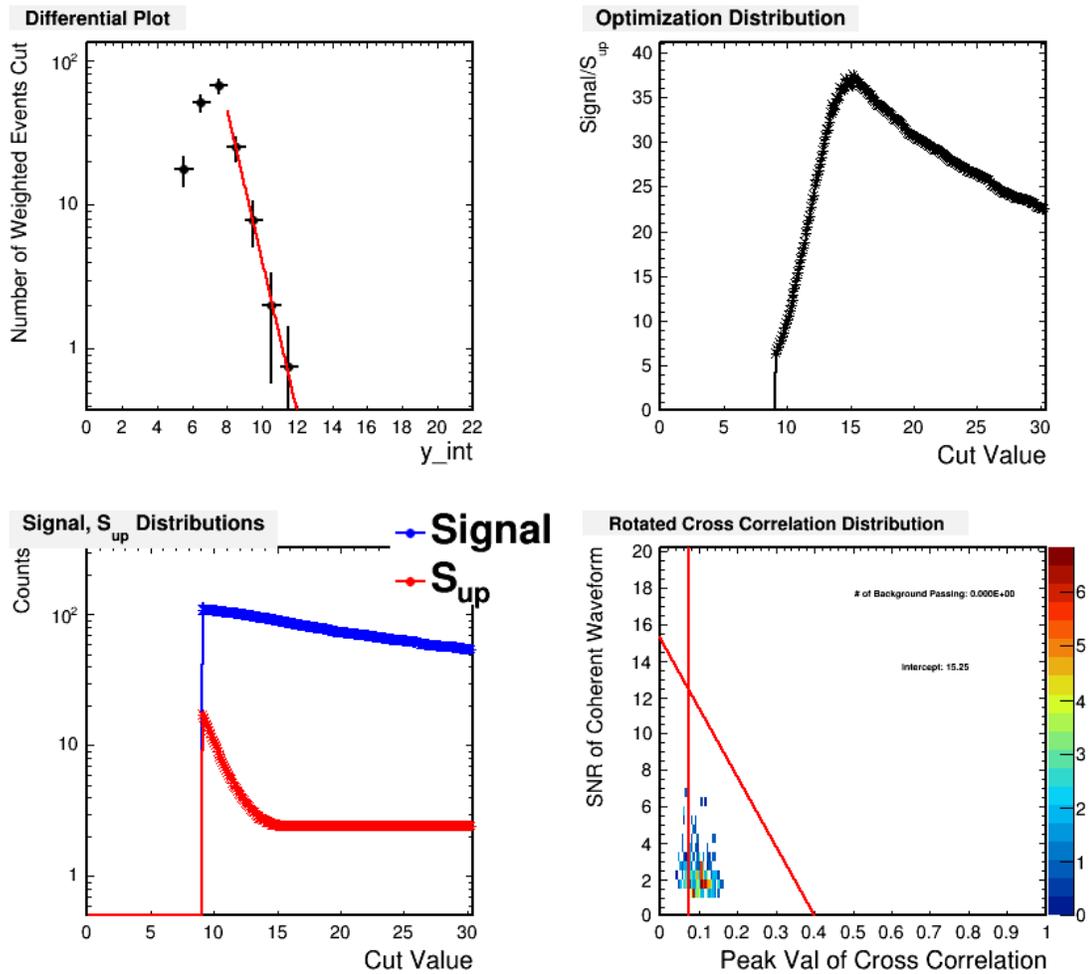


Figure B.9: The optimization plots for HealPix bin 3053. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

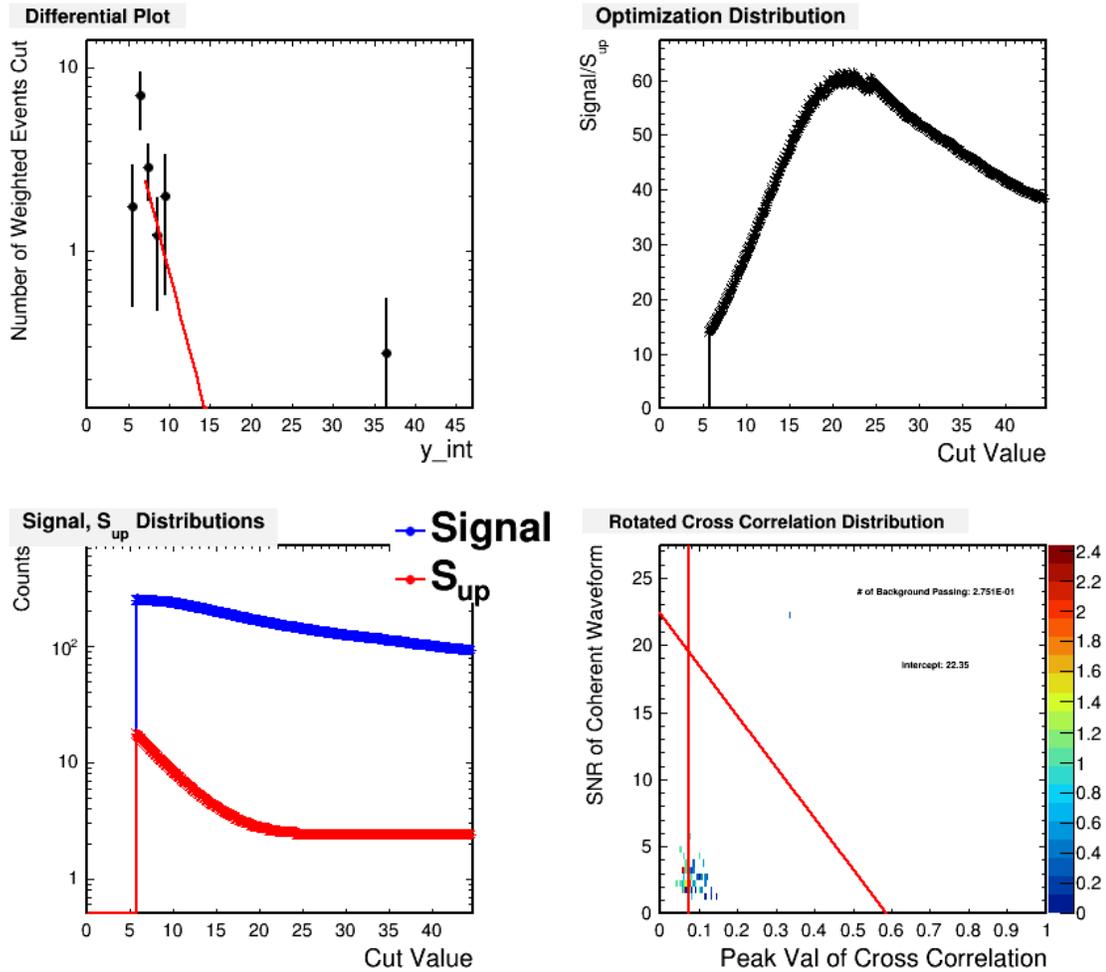


Figure B.10: The optimization plots for HealPix bin 3057. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

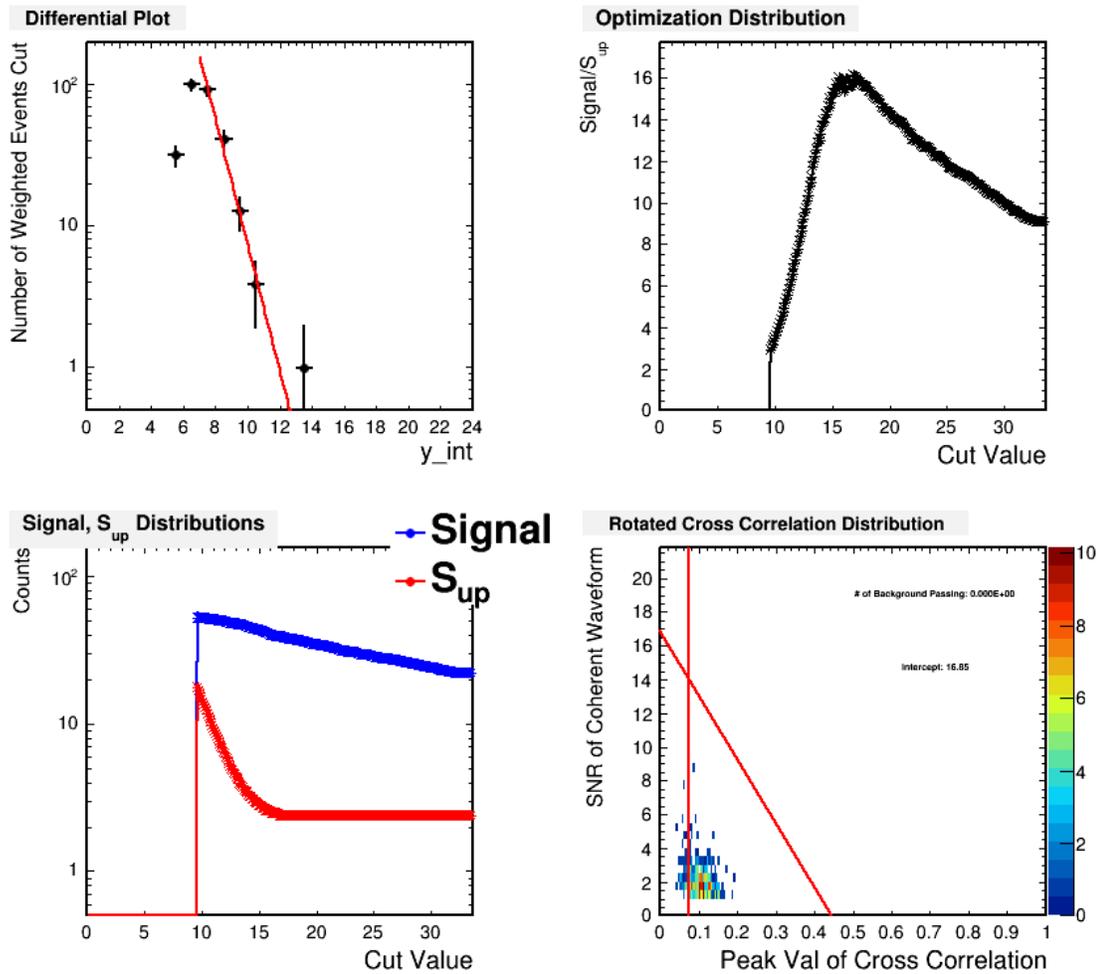


Figure B.11: The optimization plots for HealPix bin 3061. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

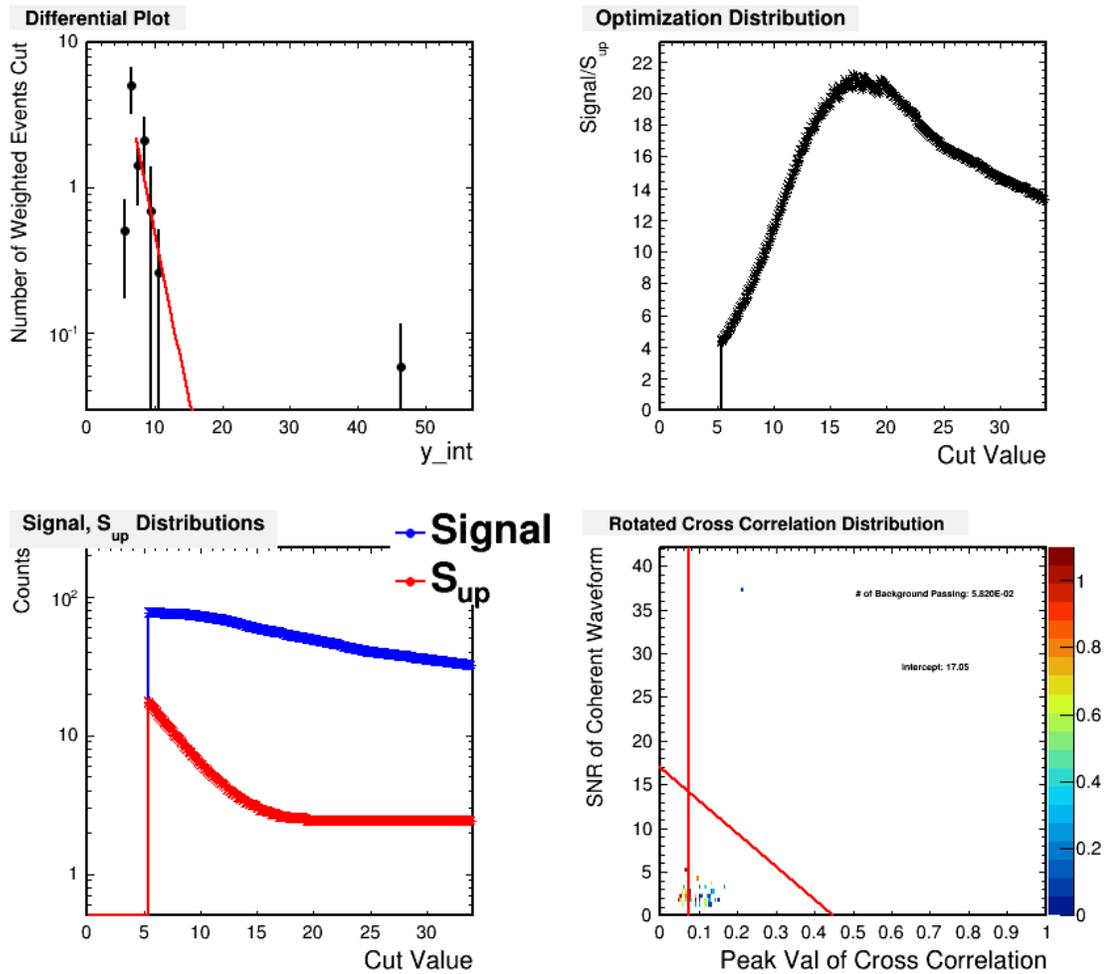


Figure B.12: The optimization plots for HealPix bin 3062. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

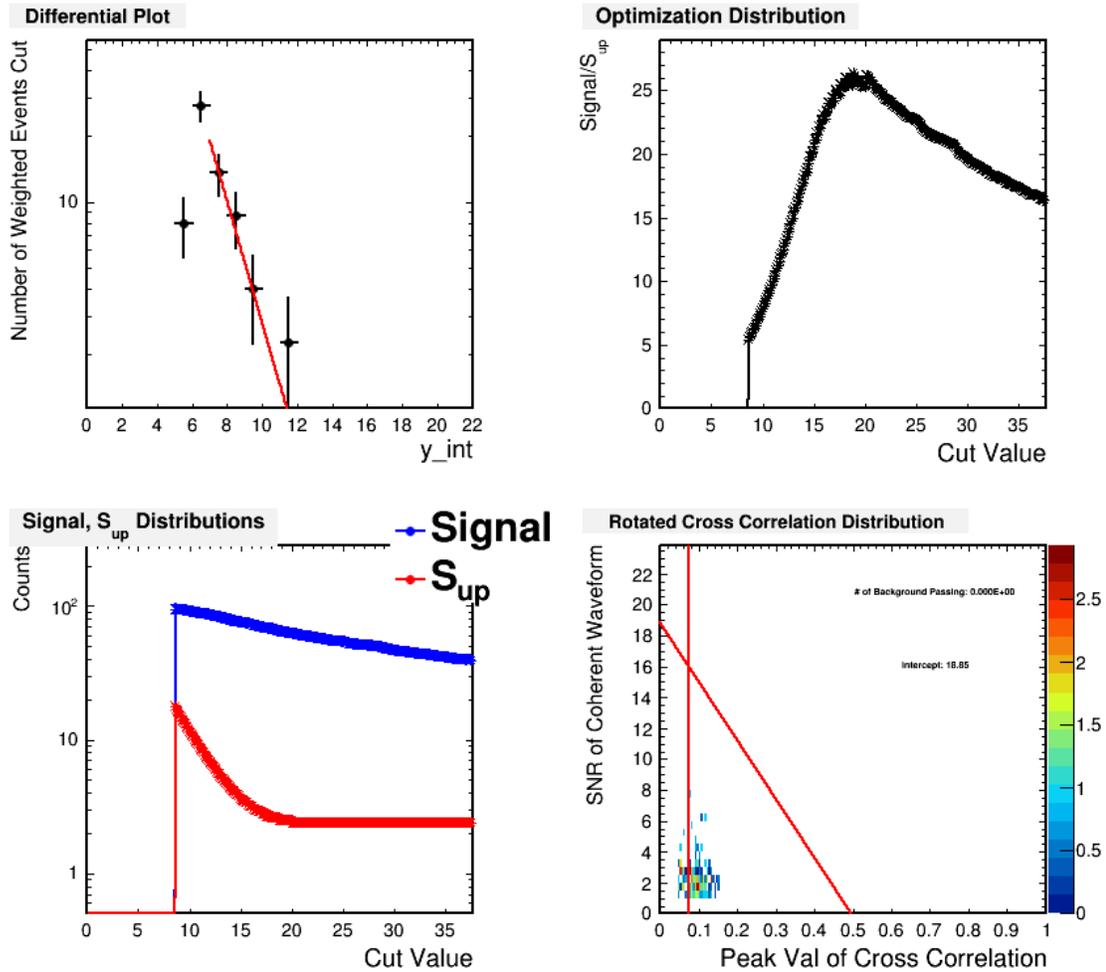


Figure B.13: The optimization plots for HealPix bin 3063. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

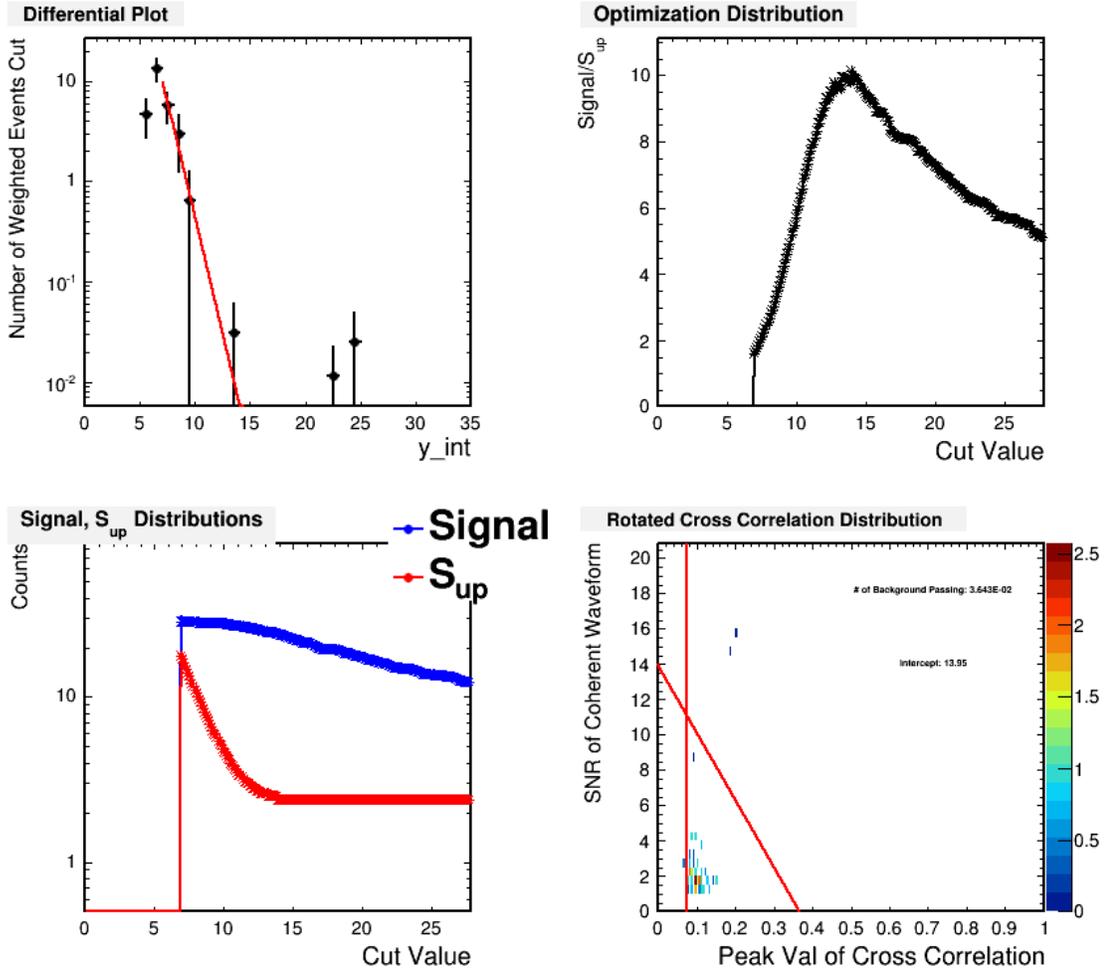


Figure B.14: The optimization plots for HealPix bin 3066. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

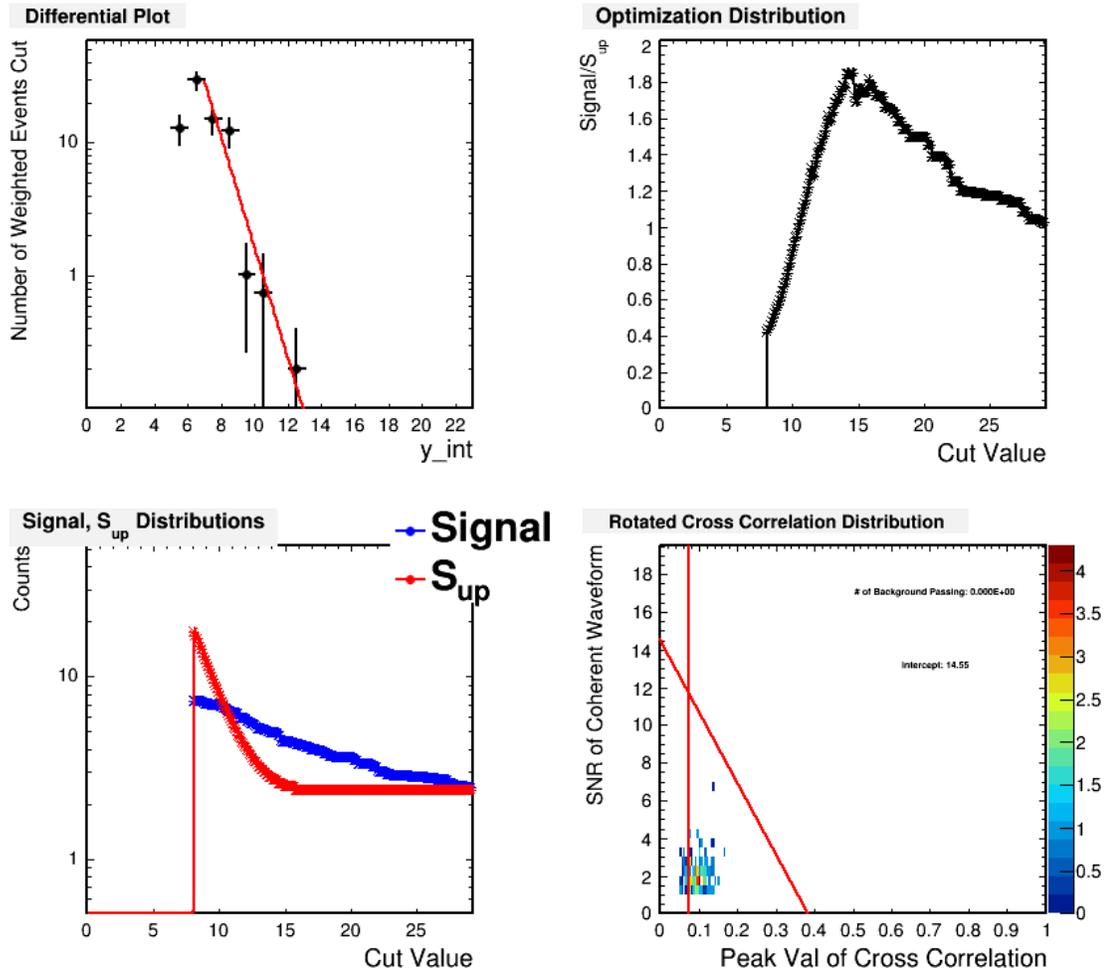


Figure B.15: The optimization plots for HealPix bin 3069. Top Left plot: Differential plot with red fit line. Top Right plot: Optimization parameter distribution. Bottom Left plot: Signal and S_{up} distributions. Bottom Right plot: Rotated cross correlation distribution with optimized slanted cut.

Bibliography

- [1] K. Kotera, D. Allard and A. Olinto *JCAP* **1010** (2010) 013 [[1009.1382](#)]. Cited on pages [ii](#) and [143](#).
- [2] P. W. Gorham *et. al.*, *Characteristics of Four Upward-pointing Cosmic-ray-like Events Observed with ANITA*, *Phys. Rev. Lett.* **117** (2016), no. 7 071101 [[1603.05218](#)]. Cited on page [vi](#).
- [3] **IceCube** Collaboration, M. G. Aartsen *et. al.*, *All-sky search for time-integrated neutrino emission from astrophysical sources with 7 years of IceCube data*, [1609.04981](#). Cited on pages [2](#) and [3](#).
- [4] **IceCube** Collaboration, M. G. Aartsen *et. al.*, *Constraints on ultra-high-energy cosmic ray sources from a search for neutrinos above 10 PeV with IceCube*, [1607.05886](#). Cited on pages [2](#) and [3](#).
- [5] P. A. The IceCube and T. A. collaborations, *Search for correlations between the arrival directions of icecube neutrino events and ultrahigh-energy cosmic rays detected by the pierre auger observatory and the telescope array*, *Journal of Cosmology and Astroparticle Physics* **2016** (2016), no. 01 037. Cited on pages [2](#) and [3](#).
- [6] 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** (1998) 505–526. Cited on page [2](#).
- [7] **Kamiokande-II** Collaboration, K. S. Hirata *et. al.*, *Results from one thousand days of real time directional solar neutrino data*, *Phys. Rev. Lett.* **65** (1990) 1297–1300. Cited on page [2](#).
- [8] **Super-Kamiokande** Collaboration, S. Fukuda *et. al.*, *Solar B-8 and hep neutrino measurements from 1258 days of Super-Kamiokande data*, *Phys. Rev. Lett.* **86** (2001) 5651–5655 [[hep-ex/0103032](#)]. Cited on page [2](#).
- [9] T. Kajita, M. Koshiba and A. Suzuki, *On the origin of the Kamiokande experiment and neutrino astrophysics*, *Eur. Phys. J.* **H37** (2012) 33–73. Cited on page [2](#).

- [10] R. Becker-Szendy *et. al.*, *IMB-3: A Large water Cherenkov detector for nucleon decay and neutrino interactions*, *Nucl. Instrum. Meth.* **A324** (1993) 363–382. Cited on page 2.
- [11] **KAMIOKANDE-II** Collaboration, K. Hirata *et. al.* *Phys.Rev.Lett.* **58** (1987) 1490–1493. Cited on page 3.
- [12] **IMB** Collaboration, R. Bionta, G. Blewitt, C. Bratton, D. Casper, A. Ciocio *et. al.* *Phys.Rev.Lett.* **58** (1987) 1494. Cited on page 3.
- [13] **IceCube** Collaboration, T. K. Gaisser, *IceCube at the Threshold*, [1507.07871](#). Cited on page 4.
- [14] **IceCube** Collaboration, M. G. Aartsen *et. al.*, *First observation of PeV-energy neutrinos with IceCube*, *Phys. Rev. Lett.* **111** (2013) 021103 [[1304.5356](#)]. Cited on page 3.
- [15] **IceCube** Collaboration, M. G. Aartsen *et. al.*, *Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector*, *Science* **342** (2013) 1242856 [[1311.5238](#)]. Cited on pages 3 and 4.
- [16] R. Gandhi, C. Quigg, M. H. Reno and I. Sarcevic, *Ultra-high-energy neutrino interactions*, *Astropart. Phys.* **5** (1996) 81–110 [[hep-ph/9512364](#)]. Cited on page 5.
- [17] A. Kusenko and T. J. Weiler, *Neutrino cross-sections at high-energies and the future observations of ultra-high-energy cosmic rays*, *Phys. Rev. Lett.* **88** (2002) 161101 [[hep-ph/0106071](#)]. Cited on page 5.
- [18] A. Connolly, *Measuring the neutrino nucleon cross section with SalSA*, *Int. J. Mod. Phys.* **A21S1** (2006) 163–167. [[163\(2006\)](#)]. Cited on page 5.
- [19] E. Borriello, A. Cuoco, G. Mangano, G. Miele, S. Pastor, O. Pisanti and P. D. Serpico, *Disentangling neutrino-nucleon cross section and high energy neutrino flux with a km³ neutrino telescope*, *Phys. Rev.* **D77** (2008) 045019 [[0711.0152](#)]. Cited on page 5.
- [20] L. A. Anchordoqui, H. Goldberg, D. Gora, T. Paul, M. Roth, S. Sarkar and L. L. Winders, *Using cosmic neutrinos to search for non-perturbative physics at the Pierre Auger Observatory*, *Phys. Rev.* **D82** (2010) 043001 [[1004.3190](#)]. Cited on page 5.
- [21] **HiRes Collaboration** Collaboration, D. Bird *et. al.*, *The Cosmic ray energy spectrum observed by the Fly’s Eye*, *Astrophys.J.* **424** (1994) 491–502. Cited on page 5.

- [22] **HiRes Collaboration** Collaboration, R. Abbasi *et. al.*, *First observation of the Greisen-Zatsepin-Kuzmin suppression*, *Phys.Rev.Lett.* **100** (2008) 101101 [[astro-ph/0703099](#)]. Cited on page 5.
- [23] **Pierre Auger Collaboration** Collaboration, J. Abraham *et. al.*, *Observation of the suppression of the flux of cosmic rays above 4×10^{19} eV*, *Phys.Rev.Lett.* **101** (2008) 061101 [[0806.4302](#)]. Cited on page 5.
- [24] **Fermi-LAT** Collaboration, F. Acero *et. al.*, *Fermi Large Area Telescope Third Source Catalog*, *Astrophys. J. Suppl.* **218** (2015), no. 2 23 [[1501.02003](#)]. Cited on page 5.
- [25] **Fermi-LAT** Collaboration, A. A. Abdo *et. al.*, *Fermi Large Area Telescope Bright Gamma-ray Source List*, *Astrophys. J. Suppl.* **183** (2009) 46–66 [[0902.1340](#)]. Cited on page 5.
- [26] P. Blasi, *Recent results in cosmic ray physics and their interpretation*, *Braz. J. Phys.* **44** (2014) 426–440 [[1312.1590](#)]. [1291(2013)]. Cited on page 6.
- [27] T. K. Gaisser *et. al.*, *Cosmic-ray composition around 10^{18} ev*, *Phys. Rev. D* **47** (Mar, 1993) 1919–1932. Cited on page 5.
- [28] **The High Resolution Fly’s Eye Collaboration** Collaboration, R. U. Abbasi *et. al.*, *Indications of proton-dominated cosmic-ray composition above 1.6 eev*, *Phys. Rev. Lett.* **104** (Apr, 2010) 161101. Cited on page 5.
- [29] J.-P. Meyer, L. O. Drury and D. C. Ellison, *Galactic cosmic rays from supernova remnants. i. a cosmic-ray composition controlled by volatility and mass-to-charge ratio*, *The Astrophysical Journal* **487** (1997), no. 1 182. Cited on page 5.
- [30] T. Abu-Zayyad *et. al.*, *Evidence for changing of cosmic ray composition between 10^{17} and 10^{18} ev from multicomponent measurements*, *Phys. Rev. Lett.* **84** (May, 2000) 4276–4279. Cited on page 5.
- [31] **Particle Data Group** Collaboration, J. Beringer *et. al.*, *Review of particle physics**, *Phys. Rev. D* **86** (Jul, 2012) 010001. Cited on page 7.
- [32] R. Blandford, P. Simeon and Y. Yuan, *Cosmic Ray Origins: An Introduction*, *Nucl. Phys. Proc. Suppl.* **256-257** (2014) 9–22 [[1409.2589](#)]. Cited on page 8.
- [33] K. Greisen, *End to the cosmic ray spectrum?*, *Phys.Rev.Lett.* **16** (1966) 748–750. Cited on page 9.
- [34] G. Zatsepin and V. Kuzmin, *Upper limit of the spectrum of cosmic rays*, *JETP Lett.* **4** (1966) 78–80. Cited on page 9.

- [35] R. Engel, D. Seckel and T. Stanev, *Neutrinos from propagation of ultrahigh-energy protons*, *Phys. Rev.* **D64** (2001) 093010 [[astro-ph/0101216](#)]. Cited on page 9.
- [36] R. J. Protheroe and P. A. Johnson, *Propagation of ultrahigh-energy protons over cosmological distances and implications for topological defect models*, *Astropart. Phys.* **4** (1996) 253 [[astro-ph/9506119](#)]. Cited on page 9.
- [37] S. Yoshida and M. Teshima, *Energy spectrum of ultrahigh-energy cosmic rays with extragalactic origin*, *Prog. Theor. Phys.* **89** (1993) 833–845. Cited on page 9.
- [38] F. Halzen and S. R. Klein, *IceCube: An Instrument for Neutrino Astronomy*, *Rev. Sci. Instrum.* **81** (2010) 081101 [[1007.1247](#)]. Cited on page 11.
- [39] G. A. Askaryan *JETP* **14** (1962) 441. Cited on page 12.
- [40] G. A. Askaryan *JETP* **21** (1965) 658. Cited on page 12.
- [41] S. W. Barwick, E. C. Berg, D. Besson, T. Duffin, J. C. Hanson, S. R. Klein, S. A. Kleinfelder, C. Reed, M. Roumi, T. Stezelberger, J. Tatar, J. Walker and L. Zou, *Radar Absorption, Basal Reflection, Thickness, and Polarization Measurements from the Ross Ice Shelf*, *ArXiv e-prints* (Oct., 2014) [[1410.7134](#)]. Cited on page 12.
- [42] M. Ackermann *et. al.*, *Optical properties of deep glacial ice at the South Pole*, *J. Geophys. Res. Atmos.* **111** (2006), no. D13 D13203. Cited on page 12.
- [43] D. Saltzberg, P. 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*, *Phys. Rev. Lett.* **86** (2001) 2802–2805 [[hep-ex/0011001](#)]. Cited on page 12.
- [44] 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** (2005) 023002 [[astro-ph/0412128](#)]. Cited on page 12.
- [45] **ANITA** Collaboration, P. W. Gorham *et. al.*, *Observations of the Askaryan effect in ice*, *Phys. Rev. Lett.* **99** (2007) 171101 [[hep-ex/0611008](#)]. Cited on pages 12, 13, and 14.
- [46] I. Kravchenko, C. Cooley, S. Hussain, D. Seckel, P. Wahrlich *et. al.*, *Rice limits on the diffuse ultrahigh energy neutrino flux*, *Phys.Rev.* **D73** (2006) 082002 [[astro-ph/0601148](#)]. Cited on page 12.

- [47] P. W. Gorham, K. M. Liewer and C. J. Naudet, *Initial results from a search for lunar radio emission from interactions of $\dot{\gamma} = 10^{19}$ eV neutrinos and cosmic rays*, in *Proceedings, 26th International Cosmic Ray Conference, August 17-25, 1999, Salt Lake City: Invited, Rapporteur, and Highlight Papers*, pp. 479–482, 1999. [astro-ph/9906504](#). Cited on page 12.
- [48] ANITA Collaboration, P. W. Gorham *et. al.*, *New Limits on the Ultra-high Energy Cosmic Neutrino Flux from the ANITA Experiment*, *Phys. Rev. Lett.* **103** (2009) 051103 [[0812.2715](#)]. Cited on pages 12, 16, and 17.
- [49] P. Allison, J. Auffenberg, R. Bard, J. Beatty, D. Besson *et. al.*, *Design and Initial Performance of the Askaryan Radio Array Prototype EeV Neutrino Detector at the South Pole*, *Astropart.Phys.* **35** (2012) 457–477 [[1105.2854](#)]. Cited on page 12.
- [50] L. Gerhardt, S. Klein, T. Stezelberger, S. Barwick, K. Dookayka *et. al.*, *A prototype station for ARIANNA: a detector for cosmic neutrinos*, *Nucl.Instrum.Meth.* **A624** (2010) 85–91 [[1005.5193](#)]. * Temporary entry *. Cited on page 12.
- [51] H. R. ALLAN, R. W. CLAY and J. K. JONES, *Radio pulses from extensive air showers*, *Nature* **227** (09, 1970) 1116–1118. Cited on page 14.
- [52] T. Huege and H. Falcke, *Radio emission from cosmic ray air showers: Coherent geosynchrotron radiation*, *Astron. Astrophys.* **412** (2003) 19–34 [[astro-ph/0309622](#)]. Cited on page 14.
- [53] 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** (01, 1965) 327–328. Cited on page 14.
- [54] S. Hoover, J. Nam, P. W. Gorham, E. Grashorn, P. Allison, S. W. Barwick, J. J. Beatty, K. Belov, D. Z. Besson, W. R. Binns, C. Chen, P. Chen, J. M. Clem, A. Connolly, P. F. Dowkontt, M. A. DuVernois, R. C. Field, D. Goldstein, A. G. Vieregg, C. Hast, M. H. Israel, A. Javaid, J. Kowalski, J. G. Learned, K. M. Liewer, J. T. Link, E. Luszczyk, S. Matsuno, B. C. Mercurio, C. Miki, P. Miočinović, C. J. Naudet, J. Ng, R. J. Nichol, K. Palladino, K. Reil, A. Romero-Wolf, M. Rosen, L. Ruckman, D. Saltzberg, D. Seckel, G. S. Varner, D. Walz and F. Wu, *Observation of ultrahigh-energy cosmic rays with the anita balloon-borne radio interferometer*, *Phys. Rev. Lett.* **105** (Oct, 2010) 151101. Cited on pages 16, 18, 19, and 21.
- [55] A. Vieregg, *The Search for Astrophysical Ultra-High Energy Neutrinos Using Radio Detection Techniques*. PhD thesis, University of California, Los Angeles, 2010. Cited on pages 24, 27, 29, 31, 34, 35, 38, 40, 46, and 135.

- [56] **ANITA** Collaboration, P. W. Gorham *et. al.*, *The Antarctic Impulsive Transient Antenna Ultra-high Energy Neutrino Detector Design, Performance, and Sensitivity for 2006-2007 Balloon Flight*, *Astropart. Phys.* **32** (2009) 10–41 [[0812.1920](#)]. Cited on pages [25](#), [26](#), [27](#), and [36](#).
- [57] P. Gorham, *Anita collaboration note*, 2008. Cited on page [35](#).
- [58] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 3 ed., 2007. Cited on page [56](#).
- [59] **ARA** Collaboration, P. Allison *et. al.*, *First Constraints on the Ultra-High Energy Neutrino Flux from a Prototype Station of the Askaryan Radio Array*, *Astropart. Phys.* **70** (2015) 62–80 [[1404.5285](#)]. Cited on page [111](#).
- [60] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke and M. Bartelman, *HEALPix - A Framework for high resolution discretization, and fast analysis of data distributed on the sphere*, *Astrophys. J.* **622** (2005) 759–771 [[astro-ph/0409513](#)]. Cited on page [121](#).
- [61] R. Brun *et. al.*, *Root Data Analysis Framework User’s Guide*. Cern. Cited on page [145](#).
- [62] I. Ruff, *The intersection of a cone and a sphere: A contribution to the geometry of satellite viewing*, *Journal of Applied Meteorology* **10** (1971), no. 3 607–609. Cited on pages [161](#) and [163](#).