

A BINNED SEARCH FOR ULTRA HIGH
ENERGY NEUTRINO WITH DATA FROM
THE THRID FLIGHT OF THE
ANTARCTIC IMPULSIVE TRANSIANT
ANTENNA

DISSERTATION

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

By

Jacob William Hunter Gordon, B.S., M.S.

Graduate Program in Physics

The Ohio State University

2018

Dissertation Committee:

Amy Connolly, Advisor

John Beacom

Klaus Honscheid

Richard Hughes

DRAFT JANUARY 2, 2018

© Copyright by

Jacob William Hunter Gordon
2018

ABSTRACT

Ultra-High Energy (UHE) Neutrinos are an important and under explored component of multimesseger astronomy. The Antarctic Impulsive Transient Antenna (ANITA) is a balloon born discovery experiment, searching for the first neutrinos in an energy region above Icecube’s energy band. ANITA flies over Antarctica, one of the most radio quite areas on the planet. Anthropogenic noise however, is still an enormous problem for ANITA. We discuss here the addition of the Tunable Universal Frontend Filter (TUFF) as a method of filtering continuous wave (CW) noise for ANITA-IV, as well as the reintroduction of the Hybrid, a signal polarization converter, added to offer a means of avoiding false triggers on Satellite CW noise. We also discuss the analysis of ANITA-III data, using a binned analysis method which, splits the ice of Antarctica into smaller regions, with the end goal of discovering UHE neutrino candidates or setting a limit on neutrino production models.

To Shannon, my fiance. Now that I am done we can finally plan that wedding!

ACKNOWLEDGMENTS

I would like to thank my adviser, Jim Beatty, who took me in off the cold dark streets of adviserless graduate school and guided me through my graduate school journey.

I would like also like to thank my substititue adviser, Amy Connolly. She has been invaluable as I work on the analysis side of ANITA, and has been kind enough to dedicate large amounts of her time to assisting me.

I thank Patrick Allison, who was wellspring of hardware wisdom within our labratory.

I thank senior graduate students James Stapleton, Ryan Hupe and Sam Stafford who have now moved onto bigger and better things, for their invaluable guadance and mentorship.

I thank Oindree Banerjee for working with me and helping me so much in my thesis research, and Brian Clark for being a much needed sounding board who helped me both solve problems and de-stress.

I thank my Committee, John Beacom, Klaus Honscheid for their sage advice, and Richard Hughes for agreeing to fill in on my committee with so little notice.

Finally I thank my loving fiance Shannon O'Connor, my sister Katie Gordon, and my parents Doug and Brenda Gordon for their constant support and unyeilding belief in me.

VITA

1988	Born, Muskegon, MI
2007	High School Diploma, Reeths Puffer High School, Muskegon, MI
2011	B.S. Physics and Astronomy, Michigan State University, East Lansing, MI
2013	M.S. Physics, The Ohio State University, Columbus, OH

Publications

The ANITA Collaboration. Dynamic tunable notch filters for the antarctic impulsive transient antenna (ANITA), The ANITA Collaboration, arXiv eprint (2017) [1]

Fields of Study

Major Field: Physics

Studies in:

Experimental Astroparticle Physics: Ultra-high Energy Cosmic Rays
Professor Jim Beatty

Experimental Astroparticle Physics: Ultra-high Energy Neutrinos
Professors Jim Beatty and Amy Connolly

Table of Contents

	Page
Abstract	ii
Dedication	iii
Acknowledgments	iv
Vita	v
List of Figures	ix
List of Tables	xii

Chapters

1 Neutrino Astronomy	1
1.1 Introduction	1
1.2 Brief History of Neutrino Astronomy	2
1.2.1 Solar Neutrinos	2
1.2.2 Supernova 1987A Neutrinos	3
1.2.3 IceCube Neutrinos	4
1.3 Why Neutrinos	4
1.3.1 Neutrinos as Observables	5
1.3.2 The GZK process and BZ neutrinos	5
1.4 Neutrinos in dense media	6
1.5 Introduction to ANITA	7
2 ANITA	11
2.1 ANITA-I	11
2.1.1 Detector	11
2.1.2 Instrument Box	13
2.1.3 Triggering	15
2.1.4 The ANITA-I Flight	16
2.1.5 ANITA-I Results	16
2.2 Anita-II	18
2.2.1 ANITA-II Hardware changes	18
2.2.2 ANITA-II Trigger Changes	20
2.2.3 The ANITA-II Flight	21
2.2.4 ANITA-II Results	22
2.3 Anita-III	23
2.3.1 ANITA-III RF Frontend Hardware Changes	23

2.3.2	ANITA-III Trigger Changes	25
2.3.3	The ANITA-III Flight	25
2.3.4	ANITA-III Results	26
2.4	Anita-IV	27
2.4.1	ANITA-IV RF Frontend Hardware Changes	27
2.4.2	ANITA-IV Trigger Changes	28
2.4.3	The ANITA-IV Flight	29
2.4.4	ANITA-IV Results	30
3	The TUFF, Tunable Universal Frontend Filter	31
3.1	TUFF Design	31
3.1.1	Notch Filter Design	32
3.1.2	Amplifications and Bias Tee	34
3.1.3	Microcontroller	35
3.1.4	Designed Integration into the existing ANITA Experiment	35
3.2	TUFF Construction	36
3.3	TUFF Debugging and Testing	37
3.3.1	Test Mode	39
3.3.2	Thermal Stress Testing	39
3.3.3	Tuff Integration	39
3.3.4	Vacuum chamber Testing	41
3.4	TUFF Programming	42
3.5	TUFF Master, TUFF Master Booster and TUFF Loom	43
3.5.1	TUFF Master Programming	43
3.6	TUFF performance during ANITA-4 Flight	44
3.6.1	Notch 1 During Flight	44
3.6.2	Notch 2 During Flight	45
3.6.3	Notch 3 During Flight	47
3.6.4	Results from the TUFF being included in ANITA-IV	47
4	The Hybrid	53
4.1	How the Hybrid Works	53
4.2	Simulation of the Hybrid	55
4.3	Characterizing the Hybrid	56
4.3.1	Experimental setup	56
4.3.2	Measurements	58
4.4	Transforming data into a means of Data driven Simulation	62
4.5	Ideal simulation verses Data driven simulation	64
5	The ANITA-III Binned Analysis	68
5.1	Past Analyses	68
5.1.1	Clustering Analysis	68
5.1.2	The ANITA-II Binned Analysis	69
5.2	The ANITA-III Binned Analysis	69
5.2.1	Healpix	70
5.2.2	Blinding Strategy and Data Sets	70
5.2.3	The Calculation of Analysis Parameters	71

5.2.4	Quality Cuts	75
5.2.5	Stage 1 Analysis Cuts	78
5.2.6	Stage 2 Analysis Cuts	79
5.2.7	Final Analysis Cuts	81
5.2.8	Background estimate	86
5.2.9	Background Estimate Systematic Uncertainty	89
5.2.10	Optimization	98
6	Discussion and Conclusions	107
6.1	Results with the 10% data set	107
6.1.1	Events Passing in our 10% data set	107
6.2	Analysis Efficiency	108
6.3	Results from our 90% sideband data set	109
6.4	Results from our 90% H-pol data set	110
6.5	Discussion	111
6.6	Future Analysis Improvements	112
6.6.1	Improved Payload blast cut	113
6.6.2	Definition of SNR	113
6.6.3	Size of the Training Data Set	113
6.6.4	Rigidity	114
6.6.5	Optimizing for Sensitivity	114
6.6.6	Gerrymandering our Healpix bins	114
Bibliography		115
Appendices		
A	Fitting with Root's Weighted log likelihood Option	119
A.1	What 'LW' means to Root and when to use it	119
A.2	What Root does when you pass Fit the 'LW' options	119
B	Fitting a 2d Gaussian to a standard error ellipse	122
C	Step by step instructions on programming the TUFF Master	124
C.1	Step 1: Preparing the TUFF Master to be programmed	124
C.2	Step 2: Preparing your computer	124
C.3	Step 3: Setting up Energia	125
C.4	Step 4: Open Energia's Serial Monitor	125
C.5	Step 5: Programming or uploading the TUFF Master	125
C.6	Step 6: Assigning the iRFCM Addresss	126
D	Plots for Passing Healpix Bins in H and V-pol Channels	127

List of Figures

Figure	Page
1.1 SN1987A Events	3
1.2 Cosmic Ray Spectrum	9
1.3 ANITA-I and ANITA-II Flux Limit	10
2.1 ANITA-I Gondola	12
2.2 Antenna Responce	13
2.3 ANITA-I Block Diagram	14
2.4 ANITA-I Flight Path	17
2.5 ANITA-I Flux Limit	18
2.6 ANITA-II Photo	19
2.7 ANITA-II Flight Path	22
2.8 ANITA-III Gondola	24
2.9 ANITA-III Flight Path	26
2.10 ANITA-III Landing	27
2.11 ANITA-IV Flight Path	30
3.1 TUFF Photo	32
3.2 TUFF Notch Circuit Diagram	33
3.3 Single Channel TUFF Photo	34
3.4 iRFCM Photo	36
3.5 TUFF Gain	38
3.6 TUFF Sandwich Photo	40
3.7 TUFF Master Photo	44
3.8 TUFF Serial Loom	45
3.9 Flight Notch Status	45
3.10 ANITA-III Observed CW Peak	46
3.11 ANITA-IV Flight DAC Threshold	49
3.12 ANITA-III and ANITA-IV phi Masking	50
3.13 ANITA-III and ANITA-IV Digitization Livetime	51
3.14 ANITA-III and ANITA-IV Instrument Livetime	52
4.1 Hybrid Photo	54

4.2	Network Analyser Photo	57
4.3	Hybrid Experiment Diagram	58
4.4	Hybrid Gain	59
4.5	Hybrid Phase Delay	60
4.6	Hybrid Phase Difference	61
4.7	Hybrid Pulse Data	62
4.8	Hybrid Simulated (Hilbert) vs Measurements	65
4.9	Hybrid Simulated (Data Driven) vs Measurements	66
4.10	Hybrid Simulaion vs Measurements	67
5.1	Healpix Map	71
5.2	Geometric Filter	73
5.3	Antenna Correlation Circle	74
5.4	Correlation Map Example	75
5.5	SNR Noise and Signal Sampling	76
5.6	Sattelite Stripes	82
5.7	Healpix bin Status	85
5.8	Exponential Fit Example	87
5.9	P-value Pseudo Experiments Histogram	90
5.10	P-value Histogram	91
5.11	Example Exponential and Powerlaw Fit	95
5.12	Example Background Experiment Distribution	98
5.13	Example Likelihood and Smeared Likelihoos Functions	103
5.14	Example plot of s and s_{up}	104
5.15	Example plot of s/s_{up}	105
6.1	ANITA-II Analysis Efficiency	108
6.2	ANITA-III Binned Analysis Efficiency	109
6.3	10% and 90% LD distributions for Healpix bin 3009	112
6.4	10% and 90% LD distributions for Healpix bin 3031	113
B.1	Error Ellipse Parameters	123
D.1	Plots for Healpix bin 3014 in V-pol	128
D.2	Plots for Healpix bin 3037 in V-pol	129
D.3	Plots for Healpix bin 3013 in V-pol	130
D.4	Plots for Healpix bin 3015 in V-pol	131
D.5	Plots for Healpix bin 3012 in V-pol	132
D.6	Plots for Healpix bin 2969 in V-pol	133
D.7	Plots for Healpix bin 3016 in V-pol	134
D.8	Plots for Healpix bin 2998 in V-pol	135
D.9	Plots for Healpix bin 2967 in V-pol	136
D.10	Plots for Healpix bin 3028 in V-pol	137
D.11	Plots for Healpix bin 2970 in V-pol	138
D.12	Plots for Healpix bin 3032 in V-pol	139
D.13	Plots for Healpix bin 2936 in V-pol	140
D.14	Plots for Healpix bin 3029 in V-pol	141

D.15 Plots for Healpix bin 3025 in V-pol	142
D.16 Plots for Healpix bin 3003 in V-pol	143
D.17 Plots for Healpix bin 3024 in V-pol	144
D.18 Plots for Healpix bin 3026 in V-pol	145
D.19 Plots for Healpix bin 3004 in V-pol	146
D.20 Plots for Healpix bin 3046 in V-pol	147
D.21 Plots for Healpix bin 2972 in V-pol	148
D.22 Plots for Healpix bin 3007 in V-pol	149
D.23 Plots for Healpix bin 3013 in H-pol	150
D.24 Plots for Healpix bin 3014 in H-pol	151
D.25 Plots for Healpix bin 3015 in H-pol	152
D.26 Plots for Healpix bin 2936 in H-pol	153
D.27 Plots for Healpix bin 2991 in H-pol	154
D.28 Plots for Healpix bin 2990 in H-pol	155
D.29 Plots for Healpix bin 2988 in H-pol	156
D.30 Plots for Healpix bin 3004 in H-pol	157
D.31 Plots for Healpix bin 3031 in H-pol	158
D.32 Plots for Healpix bin 3030 in H-pol	159
D.33 Plots for Healpix bin 2901 in H-pol	160
D.34 Plots for Healpix bin 3010 in H-pol	161

List of Tables

Table	Page
3.1 TUFF Address Table	43
5.1 Quality Cuts Table	78
5.2 Stage 1 Analysis Cuts Table	80
5.3 Stage 2 Analysis Cuts	83
5.4 Final Analysis Cuts	86
5.5 Spillover Example Table	94
5.6 V-pol Passing Healpix bins Table	99
5.7 H-pol Passing Healpix bins Table	100
5.8 V-pol Healpix Optimization Table	105
5.9 H-pol Healpix Optimization Table	106
6.1 Passing Events Table, 10% dataset	108
6.2 Events passing LD Cut in sideband Healpix Bins	110
6.3 Events passing LD Cut in 90% H-pol sample	111

Chapter 1

NEUTRINO ASTRONOMY

1.1 Introduction

People have always looked up. There used to be nothing good on TV, just the night sky to stare at. All of the information that those early astronomers received about the universe beyond our planet was coming to them through optical light. Astronomy continued this way through the invention of the telescope and the discovery of other planets.

In the last hundred years Astronomers have started using light outside of the visible spectrum. Radio, X-Ray and Gamma Ray observations have yielded new information about everything from our own sun to distant galaxies. Wide-spectrum observations have allowed detailed spectroscopy of gas in the interstellar medium. Ultra-violet light has allowed us to make direct observations of gas ejected from stars as they evolve [2]. Observations of the Cosmic Microwave Background have given astronomers a glimpse into the very early universe. Countless astronomical objects not observable in the optical spectrum have been discovered and been critical in the development of the understanding of our universe, but most of what is coming at the earth from space is not light.

In the last fifty years scientists and experiments have begun analyzing non-electromagnetic charged particles arriving at earth from space, also known as cosmic rays. It is from these particles that Cosmic Ray Astronomy started to grow. These cosmic rays offer insights into both particle physics and astronomy. Manmade particle accelerators cannot accelerate particles even to even 1/100000th the energies of the typical particles observed by The Pierre Auger Observatory and the Telescope Array observe [3] [4]. These

particles offer a check on theoretical models of particle interactions at energies we have not been able to test in laboratory settings and raise many interesting questions of their own. Where are they coming from? What is accelerating them and how? What they are made of?

In just the last few years, new observables such as gravitational waves and high energy neutrinos have expanded Astronomers sources of information even further. In the last decade both IceCube and LIGO have announced the first detections of extra galactic high energy neutrinos and gravitational waves respectively [5] [6]. The Antarctic Impulsive Transient Antenna (ANITA) is an experiment searching for a population of ultrahigh energy neutrinos more energetic than the ones seen by IceCube that have thus far never been observed.

1.2 Brief History of Neutrino Astronomy

Neutrino Astronomy is still a young field. The Neutrino was not even discovered until about 60 years ago, though it had been theorized since the 1930s. Neutrinos interact only very weakly with other matter, making them very difficult to observe. In fact up until a few years ago only two sources of Astronomical Neutrinos had ever been seen, our own Sun, and Supernova 1987A. Now the IceCube experiment has seen a few handfuls of High energy neutrinos as well.

1.2.1 Solar Neutrinos

The very first neutrino detector, the Homestake Solar Neutrino experiment, was a mine experiment sensitive only to electron neutrinos of energies above ~ 0.8 MeV. It was designed to measure the flux of neutrinos expected from our sun's nuclear processes [7]. It began construction in 1965 and relied on interactions of neutrinos with ^{37}Cl dissolved in water. The interaction caused the production of ^{37}Ar and an electron which was then observed [7]. The Kamiokande-II experiment, also an underground water based neutrino detector, confirmed the Homestake experiment's findings [8]. Even though these experiments were in agreement,

they were in disagreement with theoretical predictions. They were only measuring about one third of the expected solar neutrino flux. These early measurements and their disagreement with theory spurred interest and investment into solar neutrino experiments, and eventually led to the understanding of neutrino oscillations that we have now. These experiments were critical Nobel Prize winning stepping stones in the development of neutrino physics.

1.2.2 Supernova 1987A Neutrinos

The second source of astronomical neutrinos came on February 23rd 1987. When a blue super giant star in the nearby dwarf galaxy the Large Magellanic Cloud collapsed under its own gravity. A burst of neutrinos from SN 1987A were detected by three separate experiments well before the light from the event reached earth. 25 neutrinos arrived and were detected over roughly 13 seconds, as shown in Figure 1.1. About three hours later telescopes observed the supernova's photon emission [9].

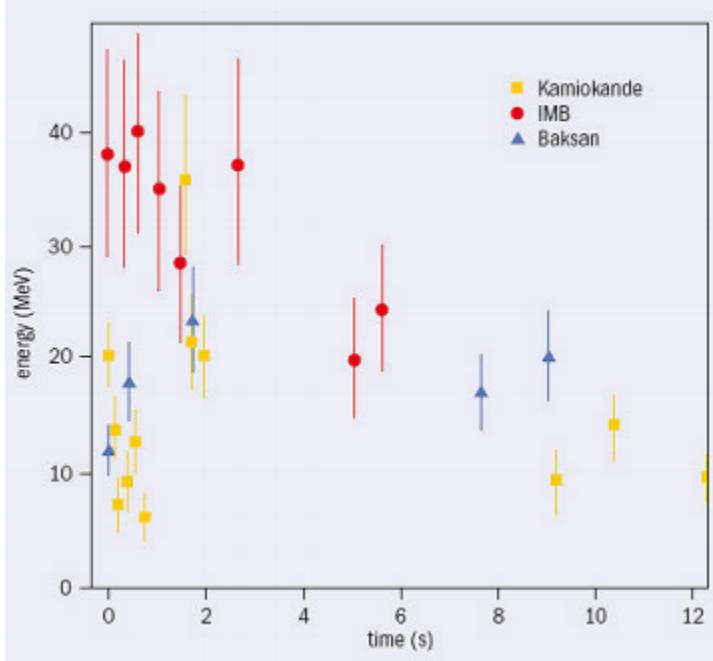


Figure 1.1: Plot showing the 25 SN1987A neutrinos detected within 13 seconds by all three groups [9].

It had long been speculated that the run-away nuclear fusion that occurs during a core collapse supernova would create an enormous flux of neutrinos. Models even relied on that neutrino exodus as an outward force to help eject the outer shell of the star [10]. Up until SN 1987A, the supernova neutrino flux had never been observed. This event was a huge experimental confirmation of scientist's understanding of at least the basics of core-collapse supernova. There hasn't been another supernova close enough to earth to observe a substantial neutrino flux from since 1987, but were one to happen the number of neutrinos observed by the numerous larger, more sensitive and currently active neutrino experiments would dwarf SN 1987A's 25 observed neutrinos.

1.2.3 IceCube Neutrinos

The IceCube experiment is an in-ice experiment made up of a three dimensional lattice of optical detectors. Each optical detector is designed to detect the optical Cherenkov radiation produced by the daughter particles of high energy neutrinos that interacted inside Antarctica's ice. In 2013 they announced that they had detected a pair of high energy neutrinos. Shortly after they revealed they had seen 28 neutrino candidate events in only two years of observing. These neutrinos had energies up to the PeV range, making them not only the highest energy neutrinos ever detected by a wide margin, but also possibly first ever neutrinos detected from a source not within the Milky Way, or its satellite galaxies [11]. Since then they have seen many more neutrinos, effectively making IceCube the first ever Neutrino Observatory looking outside our solar system. IceCube marks an important step forward into a future of intentional extragalactic neutrino astronomy.

1.3 Why Neutrinos

Neutrinos only interact very weakly with ordinary matter. This makes them extremely hard to detect, because they can pass through a lot of matter before interacting. Photons and cosmic rays do not have this issue. Given how difficult neutrinos are to detect effectively, it is valid to ask why it is worthwhile to look for them.

1.3.1 Neutrinos as Observables

Observations with photons have a critical problem. They suffer attenuation from dust in the interstellar medium and our own atmosphere. We cannot see what is happening on the inside of a star using photons because any photon created inside a star will interact before leaving it. Beyond that, dust in space can alter and stop photons before they reach the earth. Then once they arrive here, earth based telescopes suffer from interference due to the earth's atmosphere.

Cosmic rays can suffer less from attenuation at very high energies than photons. However, because cosmic rays are charged particles their trajectories through space are bent by magnetic fields. This makes any cosmic rays that come from further away than the local area of our galaxy near impossible to point back to an exact source with, especially without detailed knowledge of the galactic magnetic field, which is currently unknown.

Neutrinos don't suffer from either of these problems. They are neutral, thus ignoring magnetic fields. This means they will always point back to where they were created. They interact only very weakly with matter thus suffering next to no attenuation from intergalactic dust. They can even escape dense mediums that no other particles could, like for example the center of a dense star during a supernova.

1.3.2 The GZK process and BZ neutrinos

Another reason that neutrinos are worth looking for is that the observation or lack of observation of them offers tests on astronomical models. One clear place we know to look is for BZ neutrinos, or neutrinos created as a byproduct of the GZK process first predicted by Berezinsky and Zatsepin [12]. The GZK process, proposed by Greisen[13], Zatsepin and Kuzmin[14], happens when protons have such high energies that they interact in a resonance with cosmic microwave background photons. The process begins to become important at energies above $5*10^{19}$ eV (for protons). The interaction is believed to create a delta particle which then decays. Some of the daughter particles from that decay are neutrinos [15]. The original proton has such a huge amount of energy the neutrinos created as byproducts are

also quite energetic, with neutrinos created from the domain decay channel containing close to 5% of the energy of the original proton . A similar interaction can happen between other ultra-high energy nuclei and the CMB (cosmic microwave background) as well.

$$\begin{aligned}
p^+ + \gamma &\rightarrow \Delta \\
\Delta &\rightarrow p^+ + \pi^0 \\
\pi^0 &\rightarrow 2\gamma \\
\Delta &\rightarrow n + \pi^+ \\
\pi^+ &\rightarrow \mu^+ + \nu_\mu \\
\mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu
\end{aligned} \tag{1.1}$$

Cosmic ray experiments have measured the flux of protons and other nuclei up to energies beyond the GZK process and observe a steep drop off in cosmic ray flux around the energy the GZK process is predicted at, as shown in Figure 1.2 at energies above $5*10^{19}$ eV. This is strong (but circumstantial) evidence that out in space the GZK process is happening, which means many ultrahigh energy (UHE) neutrinos should be being created as well.

The GZK process is one specific example of a potential source of UHE neutrinos but other UHE interactions should also be creating neutrinos. Any time a process involves the creation of leptons, matching neutrinos will also be created. Cosmic rays are observed out to energies up to 10^{20} eV and beyond. The models for the incredibly energetic sources of these cosmic rays also predict fluxes of neutrinos that become potential observables for UHE neutrino experiments.

1.4 Neutrinos in dense mediums

Based on the predicted flux of ultrahigh energy neutrinos (from BZ models), there are not very many passing through earth in the first place. Only one is expected per square kilometer per century [25]. In addition, every individual neutrino only has a low chance of interacting within a given area. If we want to observe an UHE neutrino then we will either

need to observe a very large area, or observe for a very long time. Scientists (and their funders) are impatient though. This is where Askaryan radiation can help.

When a neutrino interacts inside a medium it will create an electromagnetic shower. As the electromagnetic shower moves through the medium, electrons from the medium get swept up into the shower. In addition, positrons in the electromagnetic shower are able to annihilate with matter in the medium. These two affects combine to cause the shower to have a negative charge excess. The charge excess will Cherenkov radiate as it moves through the medium. If the size of the charge excess is much smaller than the wavelength of the Cherenkov radiation that is being observed, then the radiation will coherently add giving a substantial signal a boost in power at certain frequencies [26].

For many materials the conditions for coherence are met by frequencies in the Radio energy band. This makes materials like ice and dry rock, which are mostly transparent to radio emission, excellent candidates for observing Askaryan radiation from a neutrino interaction [26]. Radio signals can travel many kilometers in ice. This can help experiments such as ANITA and The Askaryan Radio Array (ARA) counteract the low flux and low cross section of UHE neutrinos, by increasing their effective observing area. For ARA with the help of sparse instrumentation of the ice, and with ANITA by observing a large area of the ice from high in the sky.

1.5 Introduction to ANITA

ANITA, The Antarctic Impulsive Transient Antenna, is a balloon based UHE neutrino discovery experiment. It is launched in Antarctica using a NASA Long Duration Balloon. Under the Balloon is the experiment, which consists of three rings of eight or sixteen quad ridge radio antennas, and an experiment box where data processing and storage occurs. The antennas point radially outward from the center of the gondola, and ten degrees down from horizon towards the ice, searching for impulsive radio signals which escape from the surface of the ice. Surface skimming neutrinos which interact inside the ice will create electromagnetic showers which in turn create an impulsive Askaryan radio wave. ANITA

floats at about 40000m above the ice surface and is thus able to survey approximately 1.5 million square kilometers [27].

ANITA's quad ridge horns are able to observe both horizontal and vertical polarized signals between 200-1200MHz. Upward going Neutrino signals are expected to be vertically polarized, while downward going cosmic ray signals are expected to be horizontally polarized. The horizontal and vertical polarized signals can also be combined into circularly polarized left and right polarizations. Signals from satellites, their reflection off the ice, and perhaps other sources of anthropogenic noise are expected to be circularly polarized. Using the relative strength of the different polarization signals is a powerful trigger discriminate for ANTIA.

ANTIA has now flown four times. The ANITA-I and II flights did not detect any UHE neutrinos, and set limits on the possible flux of UHE neutrinos, shown in Figure 1.3.

ANITA-III and IV analysis is ongoing. The different ANITA flights and their results will be discussed in Chapter 2 in much more detail. Specific hardware components of ANITA's signal path will be discussed in Chapters 3 and 4. The TUFF in chapter 3 and the hybrid in chapter 4. An alternate method of data analysis, the Binned Analysis and its application to ANITA-II and ANITA-III data will be discussed in Chapter 5.

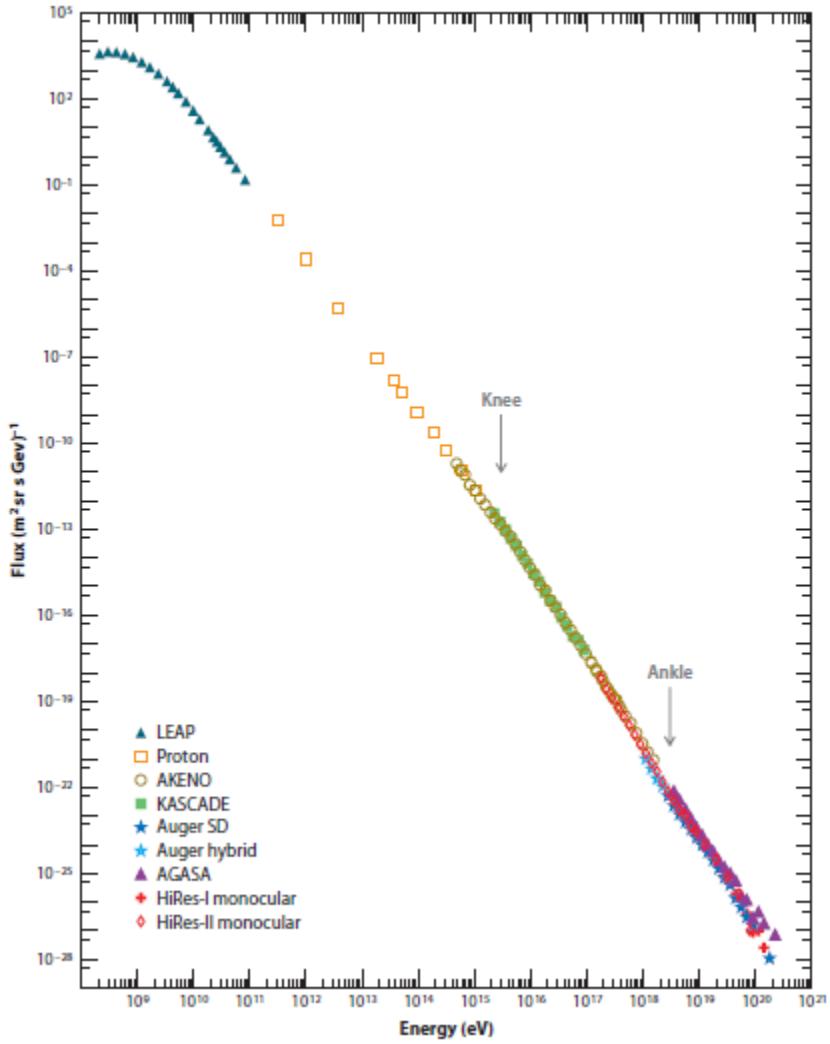


Figure 1.2: Plot showing the energy spectrum of high energy cosmic rays as seen by several experiments. The observed breaks in the power law spectrum are labeled with their common names as the ‘knee’ and ‘ankle’. The steepening of the spectrum at energies above 5×10^{19} eV could be due to the GZK process [15]. Data from LEAP [16], Proton [17], AKENO [18], KASCADE [19], Auger surface detector [20], Auger hybrid [21], AGASA [22], HiRes-I monocular [23], and HiRes-II monocular [23]. The LEAP proton-only data was scaled to the all-particle spectrum as described in [24]. Figure from [15].

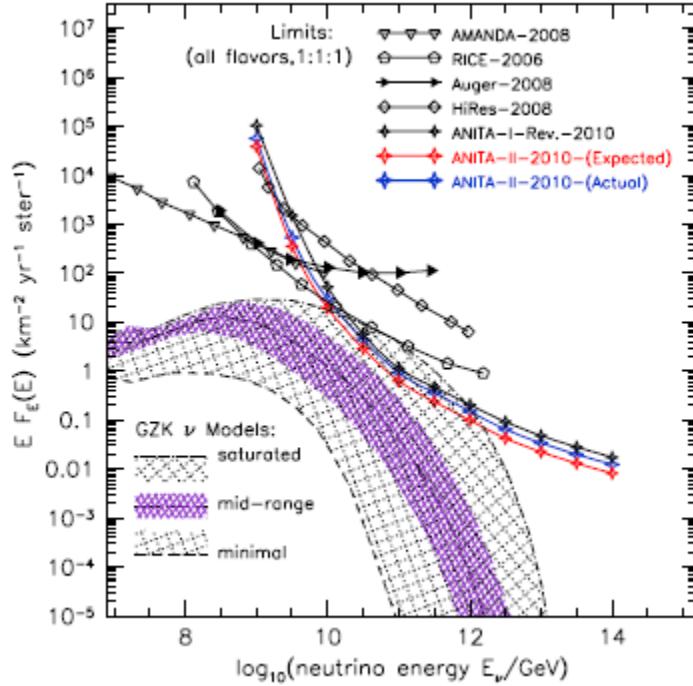


Figure 1.3: In purple are the expected flux of BZ neutrinos based on a variety of models. The red curve shows the expected ANITA-II limit, the blue curve shows the actual ANITA-II limit. Other limits are from AMANDA [28], RICE [29], Auger [20], HiRes [23], and ANITA-I [30]. BZ neutrino models are based on [25] [31] [32] [33] [34] [35], [36]. Figure from [37].

Chapter 2

ANITA

As of 2017 ANITA has had four flights. Each flight has undergone several changes in an attempt to improve various aspects of the experiment. In this chapter each of the four flights are discussed. For ANITA-I a general outline of the experiment is given, followed by a discussion of the flight and ANITA-I's major results. For ANITA-II through IV the updates over the previous experiments to the radio frequency (RF) front end and the trigger systems are summarized, followed by a discussion of flight and the major results.

2.1 ANITA-I

The first ANITA long duration balloon flight flew in the southern hemisphere's summer season of 2006-2007. Figure 2.1 shows ANITA-I's gondola in Antarctica pre launch.

2.1.1 Detector

For ANITA's first flight there were two rings of sixteen quad ridge antenna. The Top ring consists of two rings of 8 antenna. Each antenna has its pointing rotated 45 degrees in azimuth from its nearest neighbors and the two rings have their pointing offset by 22.5 degrees in azimuth from each other. This approximates the layout of the bottom ring's 16 antenna, each with their pointing rotated 22.5 degrees from its neighbors in azimuth. All thirty two antenna are tilted down from horizontal by ten degrees to optimize their sensitivity to the upper ward traveling impulsive radio signals from surface skimming neutrino interactions in the ice. The layout of the antennas can be seen in Figure 2.1.



Figure 2.1: This photo shows the ANITA-I experiment ready for launch in Antarctica [27].

The antenna themselves are custom antenna produced for the ANITA experiment originally by Seavey Engineering, Inc. The full-width-at-half-maximum of the beam is at roughly 45 degrees. This means that with 22.5 degree azimuthal offset in pointing between adjacent antennas, there is significant overlap between antenna beams. The frequency response of the antennas is shown in Figure 2.2. The antennas are sensitive to signals from 200MHz to slightly above 1200MHz.

All 32 antennas output both vertical and horizontal polarized signals to the ANITA Instrument Box. Before the signals get to the instrument box they pass through a 200-1200MHz bandpass filter and 75dB of amplification. The 75dB of amplification occurs in two stages. The amplifiers and filters combined are referred to as the Radio Frequency Conditioning Modulal (RFCM) and are stored inside of a faraday housing.

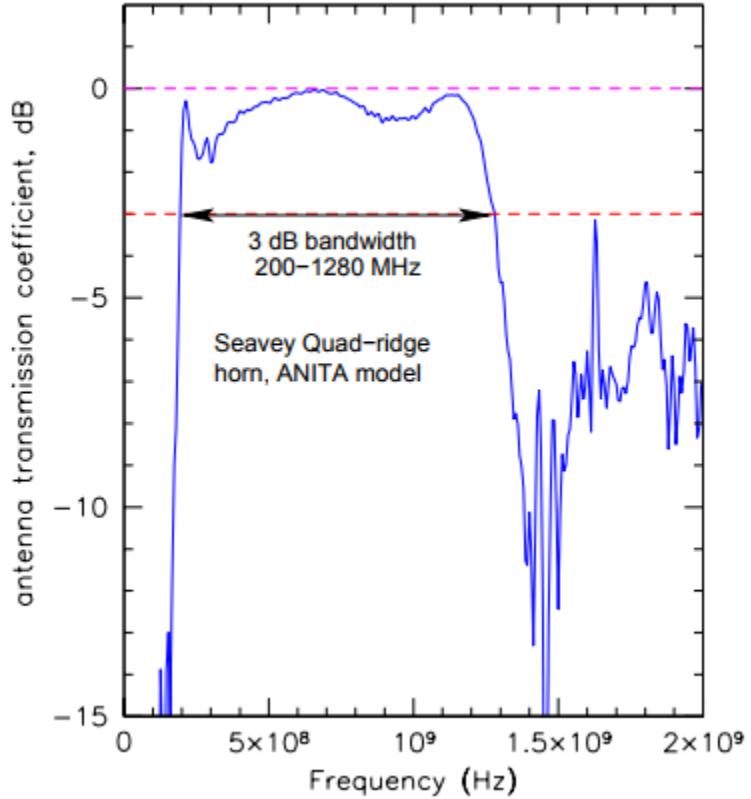


Figure 2.2: Typical transmission coefficient for a signal going into one of the quad-ridge horn antennas used by the ANITA experiment [27].

2.1.2 Instrument Box

ANITA’s instrument box sits near the center of the gondola where it rests just above the bottom ring of 16 antenna. It contains ANITA’s triggering logic computer systems and data storage inside a large faraday housing. It can be seen as the right most and shiner of the two boxes in Figure 2.1. A Block diagram of ANITA-I’s primary RF signal chain is shown in Figure 2.3. It shows the filtering and amplification a signal undergoes in the RFCM before entering into ANITA’s instrument box, as well as how the signals are processed inside the instrument box.

After traveling through a second set of 200-1200MHz bandpass filters the incoming signals are split. One set of signals, the upper set in Figure 2.3, continues on to the Sampling Unit for Radio Frequencies (SURF) boards after being delayed by 20ns. The

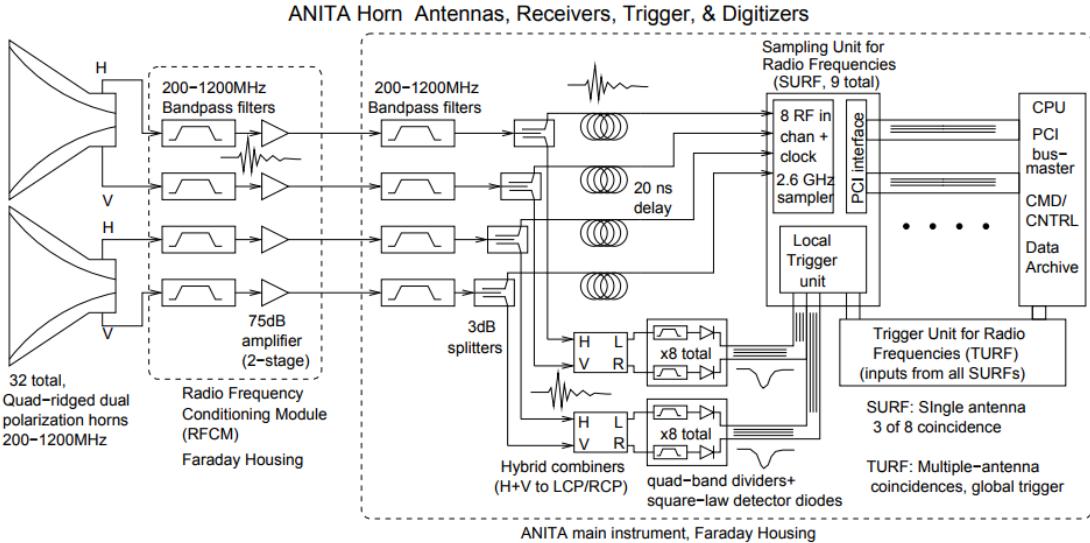


Figure 2.3: Block diagram showing ANITA-I's antenna, triggers and digitizer [27].

SURF boards, as well as the Triggering Unit for Radio Frequencies (TURF) boards, will be discussed in more detail in Section 2.1.3.

The other set of signals, the lower set in Figure 2.3 travel through the hybrids, which convert the incoming horizontally polarized (H-pol) and vertically polarized (V-pol) signals into outgoing left circularly polarized (LCP) and right circularly polarized (RCP) signals. The Hybrids will be described in detail in Chapter 4.

After the signals pass through the hybrids they are split into 4 separate bands. The bands are roughly 200-330MHz, 355-515MHz, 515-785MHz and 785-1200MHz. Impulsive Askaryan signals from neutrinos are expected to be broad band, where as many sources of anthropogenic noise are strong in only one of these bands and thermal noise in one band is independent of the thermal noise in the other bands. Analyzing the bands seperately and searching for coincidence helps reject both anthropogenic and thermal noise.

In the case of a trigger, the upper path in Figure 2.3 is digitized and saved to data archives within the instrument box. High priority digitized data is sent back down to the ground via slow, bandwidth limited, satellite networks. This was needed so that during flight it can be determined whether or not ANITA was operating as intended. If it is not

working as it is supposed to, commands can be sent back up to ANITA’s onboard CPU through those same satellite networks. When the balloon is within line of sight of the launch site, McMuerdo Base, data can be transmitted over a much higher bandwidth line of sight transmitter.

While ANITA is in flight the experiment is solar powered and thus must operate on a tight power budget. For the ANITA-I flight that was about 400W [27]. In addition because it is a balloon experiment there were tight size and weight constraints that had to be taken into account during the design of the ANITA instrument box.

2.1.3 Triggering

ANITA is a power, space and weight limited experiment. This means it cannot be recording all of the data the antennas see. It needs some way to prioritize the important neutrino like data. That is the function he of the SURF and the TURF, ANITA-I’s data triggering boards. The conditions required for ANITA-I’s different trigger levels are outlined below.

L1 Trigger

The L1 Trigger is the single antenna trigger. An L1 Trigger is issued when any three of the eight signal bands (four for LCP and four for RCP) for a given antenna go above a given threshold within 10ns of each other. This single band threshold is constantly being adjusted within the SURF board so that signals exceed the threshold at a rate of 2.6-2.8 MHz under thermal noise conditions. This is done to maximize the trigger rate while also minimizing the amount of deadtime ANITA-I experiences.

L2 Trigger

The L2 Trigger looks for coincidence between neighboring antennas. When the SURF determines that a L1 Trigger has happened it tells the TURF. The TURF saves that information and checks to see if any two adjacent antennas in either the upper or lower ring of antennas have both had an L1 Trigger in the same 20ns window. If there has been, then a L2 Trigger is issued. These happen at about a rate of 2.5kHz per antenna pair for thermal noise.

L3 Trigger

The L3 Trigger looks for coincidence between the two rings of antennas. A L3 Trigger is issued when a L2 Trigger happens within both the upper and lower ring within a 30ns window. In addition, one of the antennas from the top ring, and one of the antennas from the bottom must share the same azimuthal pointing, or phi sector. When an L3 Trigger occurs the signal is digitized. For ANITA-I they happened at roughly a rate of four or five L3 Triggers per second for thermal noise [27].

2.1.4 The ANITA-I Flight

During the Antarctic summer, a Polar vortex sets up above the Antarctic continent. This polar vortex is normally very circular and occurs in the upper stratosphere where ANITA flies. ANITA relies on this polar vortex to carry it around the Antarctic continent and keep it over the ice, from which it seeks to observe neutrino signals. An image of ANITA-I's flight path as it was carried by the polar vortex is shown in Figure 2.4. ANITA-I completed 3 full rotations and one partial rotation around Antarctica during flight over its 35 days of air time. Much of that time however was unfortunately spent near large sources of anthropogenic noise, or over thinner areas of ice. Additionally on board flight computers experienced stability issues. The effective livetime of ANITA-I was 17.25 days [37].

2.1.5 ANITA-I Results

No neutrino candidate events were detected during ANITA-I's flight. Limits were set on the allowed flux of UHE neutrinos based on the lack of detection. At the time these limits were the most competitive from any experiment. They are shown in Figure 2.5 along with limits from several other experiments from pre 2008 era for historical comparison [38]. Note that the ANITA-I limit was later revised to what is shown in Figure 1.3 in Chapter 1.

ANITA-I did detect several cosmic ray candidates. Cosmic rays are expected to be predominantly H-pol signals due to the charged particles interactions with the earth's magnetic field. Neutrino signals are expected to be predominantly V-pol signals because of the nature

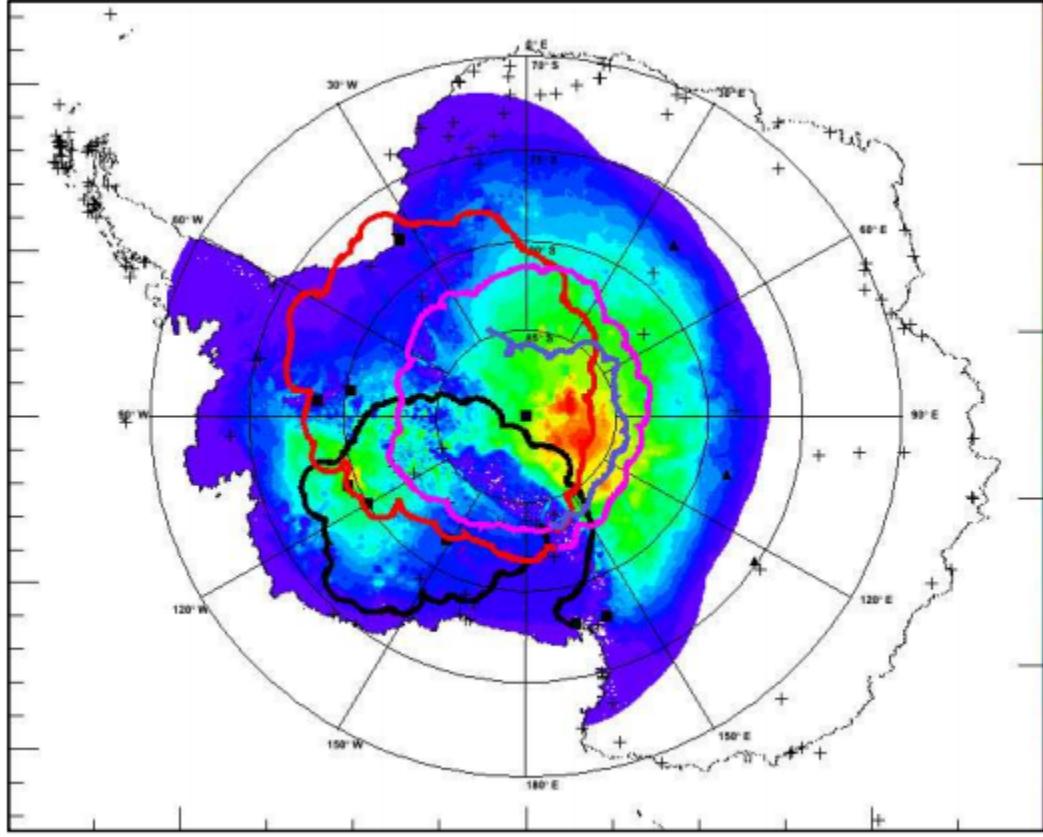


Figure 2.4: The path ANITA-I was carried over Antarctica by the polar vortex. The line color represents different loops. Chronologically it goes black, red, blue, purple. The color map on the continent represents ice depth. Red is greater than 4km. Yellow is 3-4km. Green is 2-3km. Light blue is 1-2km, and blue is less than 1km [27].

of Askaryan radiation. Two populations of cosmic ray candidates were observed. The first, and larger population are believed to be radio signals created by downward moving cosmic ray interactions in the atmosphere that are reflected off of the ice and back up to ANITA's antennas. The second population believed to be from radio signals created by earth skimming cosmic rays. They point back to below horizontal, but above the horizon. This second population shows its polarity reversed from the first population. The first population's polarity would have been reversed when it was reflected off the ice. This population of earth skimming cosmic rays had never been observed before ANITA-I [41].

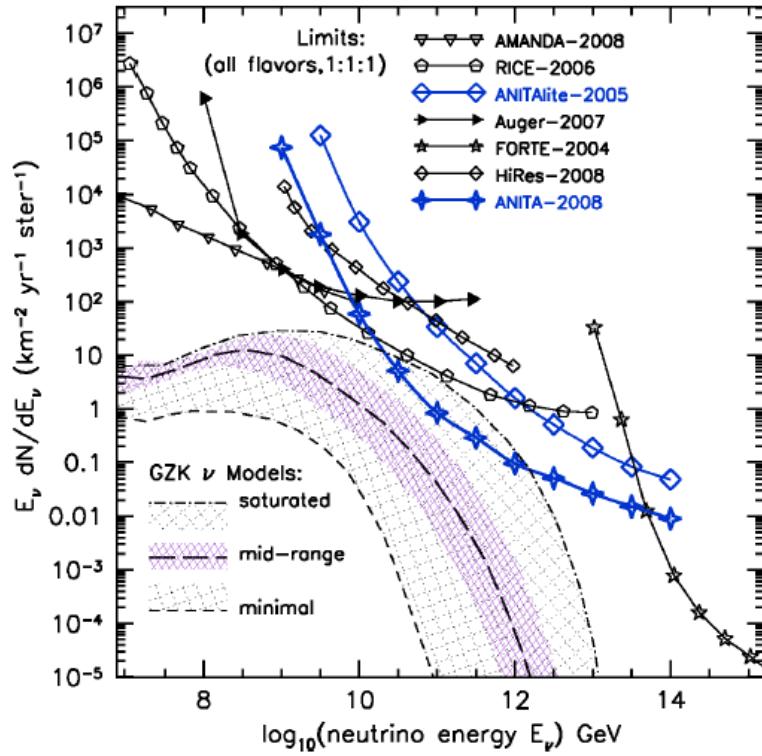


Figure 2.5: The historical limit on the flux of UHE neutrinos set by the ANITA-I experiment. Limits from experiments AMANDA [28], RICE [29], Auger [20], HiRes [23], FORTE [39] and ANITA-lite [40] are also shown. GZK neutrino model range is based off of a variety of models [31] [25] [32] [33] [34] [35] [36]. Figure from [38].

2.2 Anita-II

The flight of ANITA-II happened in the 2008-2009 Antarctic summer season. It flew for 31 days giving it 28.5 days of livetime, roughly 50% more than ANITA-I. ANITA underwent several improvements between its first and second flights which will briefly outlined in the following sections.

2.2.1 ANITA-II Hardware changes

ANITA-II added eight additional quad ridge horn antennas. Due to special constraints in the launch procedure these additional horns needed to be folded underneath the experiment and deployed after launch. Figure 2.6 shows these eight horns deployed below the bottom



Figure 2.6: Image showing the ANITA-II experiment rising. The insert is ANITA-II at a height of 35km as viewed through a telescope [37].

ring of solar panels. These eight new antenna and the other thirty two antenna are all of the same type used in the ANITA-I flight. The new horns helped to increase ANITA-II's over all aperture and pointing resolution.

Several improvements were made to the front end amplification of the signal for ANITA-II. Improvements were focused on reducing the noise figure for the vertical polarized channels. The first stage of amplification was moved to directly after the antennas so that cable noise is not amplified. The second stage of amplification, which was already called the low noise amplifier (LNA), was replaced with an even lower noise amplifier. We still just call it the LNA, for some reason ELNA didn't catch on. The third and final change to the radio frequency (RF) front end was swapping out the low and high band pass filters in the signal path for a lower noise figure band pass filter with the same frequency behavior. These changes were done only for the V-Pol because it is in the V-Pol that observations of neutrinos are expected. The first stage amplifier as well as the band pass filter and a bias-T are

all stored inside the same RF proof box. The unit as a whole is called the antenna-mounted pre-amplifier (AMPA). The second stage amplifier and another bias-T were stored int the RFCM.

Writing speed to the hard drives was also increased. This improvement allowed, on average, ANITA-II to sustain double the global trigger rate ANITA-I was capable of.

2.2.2 ANITA-II Trigger Changes

It is this same prediction for neutrino signals to appear predominatly V-pol that motivated the changes in ANITA-II's triggering logic. ANITA-II triggered only on impulsive V-pol signals, where as ANITA-I triggered on strong impulsive V and H-pol signals. Horizontally polarized data is still saved for successful V-pol triggers. This means that the Hybrids that converted the H and V-pol signals into LCP and RCP signals has been removed.

ANITA-II L1 Trigger

While for ANITA-I the L1 Trigger occurred when three of the eight signal bands (four bands for LCP and four for RCP) went over a given threshold, there are now only have four signal bands. ANITA-II did not trigger on any of the H-pol information. Additionally the four bands themselves have changed. ANITA-II had a low band, middle band, high band, and full band. The three dB points of these four bands were 200-350MHz, 330-600MHz, 630-1100MHz, and 150-1240MHz. The condition for an L1 trigger was two of the three frequency bands and the full band exceeding their individual thresholds.

ANITA-II L2 and L3 Triggers

For the two full rings of antenna the L2 and L3 triggers remained the same. However the triggering logic for the 8 new drop down antennas needed to be modified. There were fewer antennas in the drop down ring. While for the top and middle rings an L2 trigger occurs when two of three adjacent antenna have a L1 trigger within 20ns. For the bottom ring an L2 trigger happens when two adjacent antenna have a L1 trigger within 20ns. The

trigger thresholds for the bottom ring were tuned down to keep the bottom rings over all contribution to the L3 trigger proportional to the number of antenna in the ring.

Trigger Phi Masking

Additionally, phi masking was added for ANITA-II. ANITA-I suffered from being oversaturated by anthropogenic noise when it was near large groups of humans, such as McMurdo Station, the U.S. base ANITA is launched from. This led to large amounts of deadtime when only half of the antennas, the half pointing at the base, were causing the trigger saturation. To counter act this in ANITA-II a phi masking system was implemented. Phi sectors are made up of antenna with the same azimuthal pointing. When a phi sector's L3 Trigger goes above a certain threshold (typically 20Hz) for more than some amount of time (typically 30 seconds) then that phi sector's triggers are ignored. Data is still recorded and saved if other phi sectors cause a trigger, the phi sector being masked is just allowed to cause a global trigger. When the L3 Trigger rate goes below a given threshold (typically 10Hz) for some period of time (again, typically 30 seconds) then triggering is re-enabled. Each phi sector can be masked individually to allow directions facing away from trigger saturating sources to remain sensitive to neutrinos [38].

2.2.3 The ANITA-II Flight

An image of ANITA-II's flight path is shown in Figure 2.7. The first loop shown in purple looped tightly in an atypical fashion. Luckily the second loop, in red, and the third loop, in black, took much wider arcs that overlooked deeper ice. Overall ANITA-II was able to see deeper ice than ANITA-I.

In addition to traveling a better overall flight path than ANITA-I traveled, ANITA-II was able to achieve much higher live time. ANITA-II did not experience the flight computer stability issues that plagued the first flight. Additionally, phi masking allowed ANITA-II to continue searching for neutrinos with a subset of its antennas even when near large sources of anthropogenic noise. ANITA-II flew for 31 days and achieved a live time of 28.5 days, nearly doubling ANITA-I's live time.

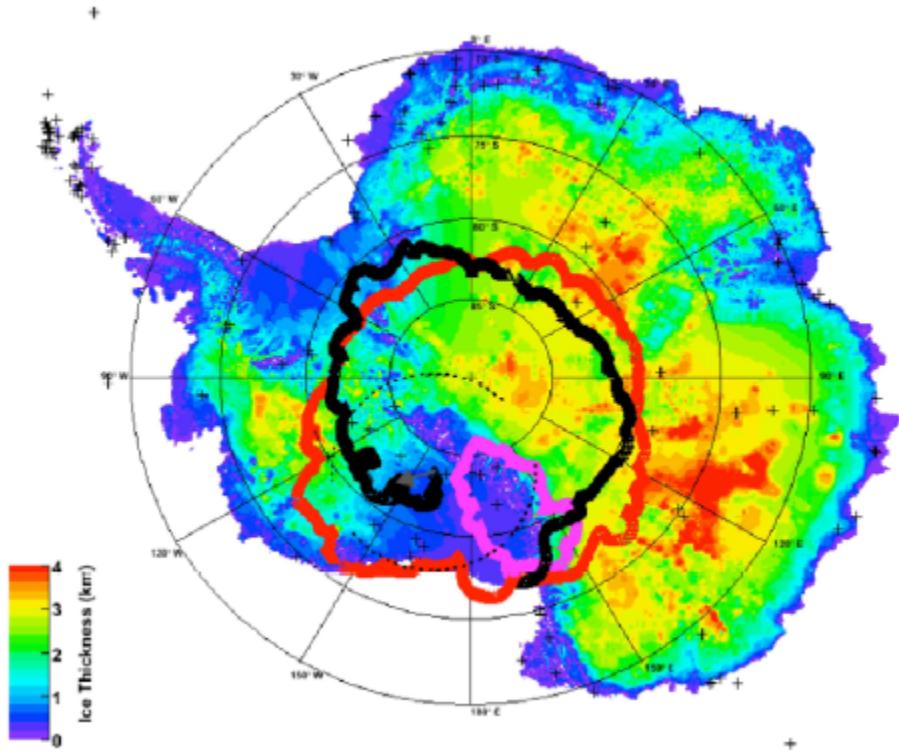


Figure 2.7: The path ANITA-II as carried over Antarctica by the polar vortex. The line color represents different loops. Chronologically it goes purple, red, black. The color map on the continent represents ice depth. Figure by R. Nichol [38].

2.2.4 ANITA-II Results

Due to ANITA-II's improvements it had about 4 times the sensitivity ANITA-I did. The limit that would be derived from this improvement in sensitivity is shown in Figure 1.3 as the red line labeled ANITA-II-2010-(Expected). However, analysis of ANITA-II's data revealed two neutrino candidates with an expected background of $0.97 +/ - 0.42$ events. Though the two events are consistent with simulations of neutrinos, the collaboration lacked significant enough statistics to identify these two events as a unique non-anthropogenic population. ANITA-II saw more events passing than it's expected background, so the limit it could set was worsened by about a factor of two. It is shown as the blue line labeled ANITA-II-2010-(ACTUAL) in Figure 1.3.

Additionally ANITA-II’s analysis revealed three horizontally polarized UHECR candidates. ANITA-I saw 16 of these, but the reduction in signal is to be expected. ANITA-II was not triggering on horizontal signals, where as ANITA-I was.

2.3 Anita-III

Anita-III flew from December to January in the Antarctic summer season of 2014-2015. It flew for 23 days and achieved a live time of 18.7 days [42]. It saw several new improvements over the ANITA-II flight. The gondola is shown in Figure 2.8 before launch. The flight path ANITA-III took unfortunately led to an earlier termination of the flight and less air time than earlier ANITA flights were able to achieve. Figure 2.9 shows the flight path ANITA-III took. After reviewing the fact that the primary science result from the first two ANITA flights were the UHE cosmic rays observed by the first ANITA flight, it was decided that ANITA needed to regain the H-pol sensitivity that it lost in its second flight. Many upgrades to ANITA-III revolve around this.

2.3.1 ANITA-III RF Frontend Hardware Changes

ANITA-III saw one major and obvious improvement easily visible in Figure 2.8. Eight antennas have been added to the bottom ring, bring all three rings of antennas up to the standard 16 antennas. This was achieved by fixing the bottom rings 16 antennas in place, and instead allowing the bottom array of solar panels to descend into their location after launch. In ANITA-II it was the bottom ring of eight antennas that lowered into place after launch and the solar panels were fixed. ANITA-I completely lacked the third bottom ring of antennas. This brings the total number of antennas in the ANITA-III experiment up to 48. In addition to the eight new antennas, the antennas themselves were upgraded to new Seavay quad-ridge horn antennas with an improved frequency response in the 200-300 MHz range [44].

In addition to the improvements to the antennas and the antenna layout, the Upgrades to the pre amplification stage that were given to the V-pol signal paths for ANITA-II were



Figure 2.8: ANITA-III gondola with solar panels raised over the new complete bottom ring of 16 antennas. Image by S. Wissel [43].

added to the H-pol and new antenna's signal paths as well. This meant that all channels now had AMPAs. The amplifiers for the bottom ring of antennas differ from the other three rings. They are referred to as drop-down antenna-mounted pre-amplifiers (DDAMPAs). The DDAMPAs contain two separate LNAs instead of the one that the AMPAs contain. The combination of the two give slightly more amplification than the AMPA's LNA and are powered by 5V instead of 12V.

The RFCM was replaced with the internal Radio Frequency Control Modual (iRFCM) which was now just inside of the Instrument Box. As with the RFCM, the iRFCM contained a bias T and the second stage of amplification. The amplifiers used in ANITA-II were replaced with the AMPLITE, which was designed at The Ohio State University. Each iRFCM handled 24 signal channels, so ANITA-III had four iRFCMs. Inside each iRFCM

was also a power distribution board responsible for powering all 24 of the AMPLITEs and bias Ts within a single iRFCM.

2.3.2 ANITA-III Trigger Changes

In order to allow for triggering on horizontally polarized cosmic ray events the trigger logic for ANITA-III was overhauled. Instead of splitting incoming signals into multiple bands like in past experiments, only the full band is triggered on, but in both H-pol and V-pol. This slightly weakened ANITA-III's neutrino sensitivity but allows for triggering on UHE cosmic rays. Because of these changes, ANITA-III only has a L1 and L3 trigger.

L1 Trigger

An L1 Trigger is issued if two of the three antennas within the same phi sector are above their threshold within 10ns of each other. Like for past ANITA's L1 Triggers, the thresholds for each full band signal path are set depending on how often the band is triggering. ANITA-III's L1 Trigger incorporates previous experiments L3 Trigger because there are not different signal bands to require coincidence between. The L1 Trigger rate for ANITA-III was about 15 kHz.

L3 Trigger

If two adjacent phi sectors both have an L1 trigger within the same 30ns window then an L3 trigger is issued. ANITA-III's L3 Trigger is similar to past experiment's L2 Triggers. The L3 trigger rate for ANITA-III was about 50 Hz.

2.3.3 The ANITA-III Flight

The ANITA-III flight did not go quite as well as previous flights. Its first loop around Antarctica started well, traveling over deep ice. However ANITA-III's ever widening spiral outward soon took it towards open water. Due to this, the ANITA-III flight needed to be ended after only 23 days, the shortest ANITA flight yet. ANITA-III's flight path can be seen in Figure 2.9. In addition to the short flight, ANITA-III had a rough landing,

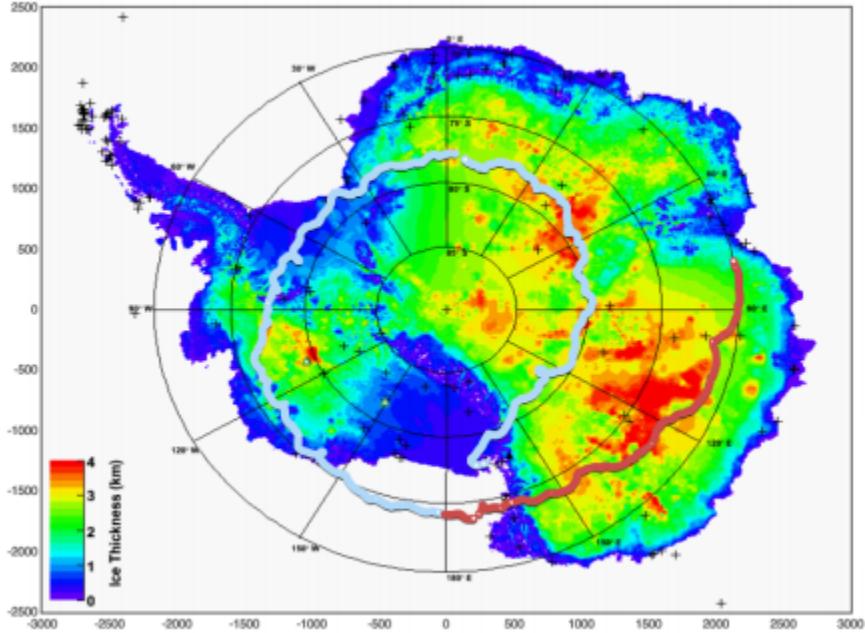


Figure 2.9: Figure shows ANITA-III’s flight path over a map of Antarctica. The first loop of the flight path is shown in gray, and the second loop is shown in red. The color on Antarctica represents the ice thickness in kilometers [44].

especially when compared to the previous two flights. Some of the damage to the upper ring of antennas can be seen in Figure 2.10 which shows ANITA-III’s final resting place. Unfortunately due to weather and other external factors a full recovery of ANITA-III was never achieved. The data storage disks and instrument box were recovered, but much of the rest of the gondola still rests, partially buried, in the snow.

2.3.4 ANITA-III Results

Analysis of ANITA-III’s data is still underway. One thing that was clear after initial analysis of the data though is that in the several years since ANITA-II the levels of electromagnetic interference became a major problem. Much of the additional interference is continuous wave (CW) circularly polarized signals at specific frequencies. It is believed that the source is geostationary satellites. This observation motivated the changes ANITA underwent as it evolved for its fourth flight.

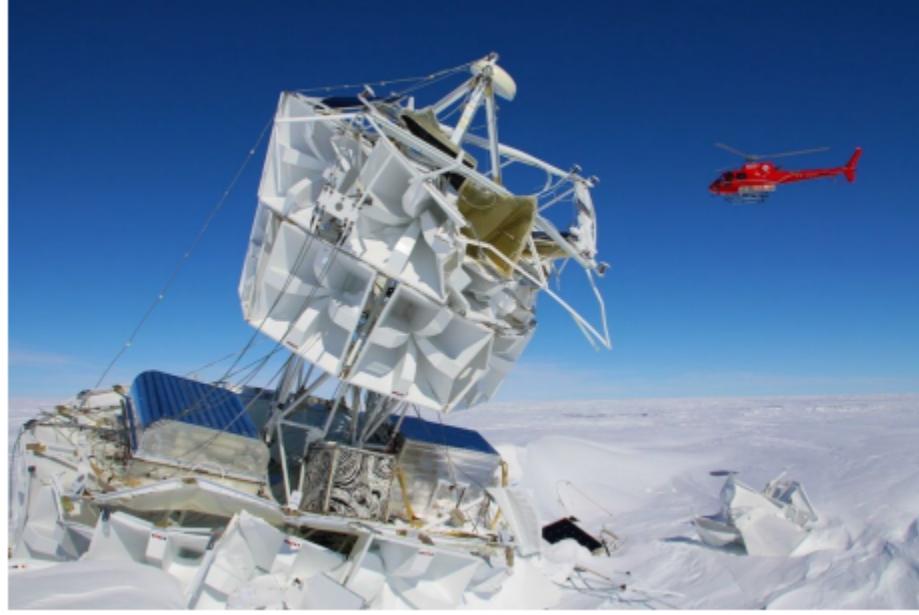


Figure 2.10: Photo of ANITA-III at its landing site. Photo credit to C. Miki and the Australian Antarctic Program [44].

2.4 Anita-IV

The ANITA-IV experiment flew in the 2016 Antarctic summer season. It was launched on December 2nd and flew until December 30th, giving it 28 days of air time. More than ANITA-III, but less than both ANITA-I and ANITA-II. It was in neutrino searching mode for 94% of its air time, giving it a live time of 27.3 days. This makes it the ANTIA flight with the second most live time, only one day less than ANITA-II, and with substantially more live time than either ANITA-I or ANITA-III. Anita-II was not triggering on H-pol signals, so ANITA-IV achieves the largest live time for UHECRs by a wide margin.

2.4.1 ANITA-IV RF Frontend Hardware Changes

Unlike previous ANITA flights, ANITA-IV did not gain any additional antennas. It used the same antenna layout as ANITA-III. It did however have several changes to its RF front end. The bandpass filters that were filtering the H and V-pol antenna signals in previous flights were replaced with lowpass filters. Additionally the custom NTU LNAs that were

being used in the 32 DDAMPAs in ANITA-III were used for all 96 channels in ANITA-IV. The combination of these two modifications reduced the noise temperature of ANITA-IV’s RF front significantly, leading to a 20% improvement in ANITA-IV’s energy threshold when compared to ANITA-III [42].

After the AMPAs the signals for ANITA-IV traveled through new Tunable Universal Frontend Filter (TUFF) boards. The TUFFs were designed to be tunable notch filters, capable of removing the CW noise that plagued ANITA-III. Each board had three separate tunable notches, but during flight the notches were left at constant values of 260, 380 and 460MHz. In addition to filtering, the TUFFs replaced the AMPLITEs second stage amplification and like the AMPLITEs are inside of the iRFCMs. The TUFFs were designed and constructed at The Ohio State University in the Beatty Connolly Lab. Their design, testing, integration and performance will be discussed in much more detail in Chapter 3.

2.4.2 ANITA-IV Trigger Changes

The Hybrids that were in ANITA-I’s signal path make a return in ANITA-IV. Once again it converts the incoming linearly polarized H-pol and V-pol signals into circularly polarized LCP and RCP signals. A detailed characterization of the Hybrids can be found in Chapter 4. The change means that once again, ANITA’s trigger logic had to be updated.

L1 Trigger

For ANITA-IV an L1 trigger occurs when both the LCP and RCP signals exceed their individual thresholds within the same 4ns window. The addition of the Hybrids was done to add the coincidence requirement of LCP and RCP signals into ANITA-IV’s trigger. Signals from both UHE neutrinos and UHECRs are expected to be linearly polarized. A linearly polarized signal should have roughly equal parts LCP and RCP, whereas the CW noise that plagued ANITA-III was strong in either LCP or RCP, but not both. Additionally because thermal noise in LCP and RCP are uncorrelated, the coincidence requirement should reduce thermal noise triggers as well. Just like in past ANITA experiments the thresholds for the individual triggers that feed into an L1 trigger are adjusted in real time to keep a consistent

L1 trigger rate. This real time adjustment did not need to be used as much as in previous experiments however, likely due to the addition of the TUFFs.

L2 Trigger

Once an L1 Trigger occurs, if another antenna in the same phi sector also has a L1 trigger within a given time window than a L2 Trigger is issued. The time window varies depending on which ring of antennas sees the first L1 Trigger. This was done to give preference to the upward coming signals ANITA is searching for. The time windows for the bottom, middle, and top rings of antenna are 16ns, 12ns, and 4ns respectively.

L3 Trigger

A L3 Trigger is issued if two adjacent phi sectors both have an L2 Trigger within the same 10ns time window. Once an L3 trigger occurs the incoming data is recorded. The L3 Trigger is the same as it was in the previous ANITA experiment but with a much shorter time window allowed for the two adjacent L2 Triggers to occur within.

2.4.3 The ANITA-IV Flight

ANITA-IV’s flight path started off with a wide spiral, however as ANITA-IV continued to curl around the Antarctic continent, its path began to tightly circle back on its own course before continuing. As the flight continued, this effect became more and more pronounced. Additionally each successive loop was smaller than the last, as ANITA-IV spiraled inward towards the South Pole. ANITA-IV’s flight path is shown in Figure 2.11. Though near the end of ANITA-IV’s flight it was still functioning well, it had become stuck tightly spiraling over the same location, so it was decided that it should be dropped. On the upside, ANITA-4’s path mostly kept it away from McMurdo station, the largest source of noise in Antarctica.

The Antarctic winter season has kept us from fully recovering ANITA-IV as of now (November 2017). The data disks have been fully recovered, as well as some of the smaller more important items, such as the AMPAs. The rest of the instrument box however sat in

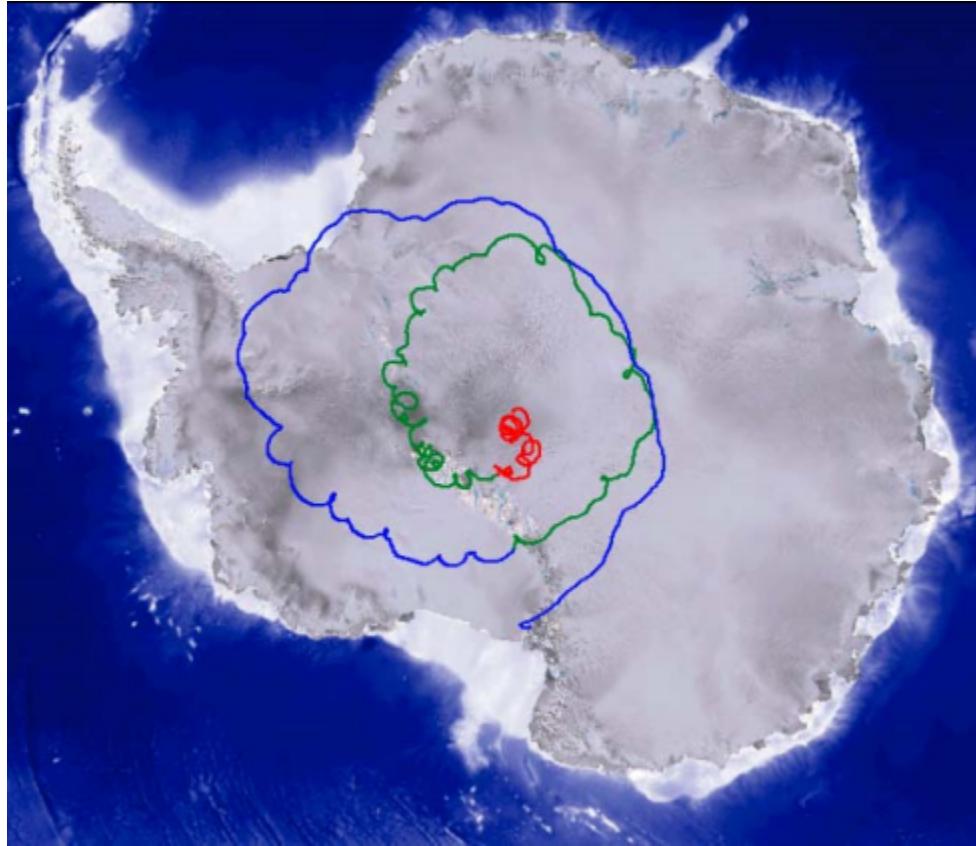


Figure 2.11: This Photo shows ANITA-IV’s flight path projected onto satellite imagery of Antarctica. This first loop is shown in blue, the second in green, and the third in red [42].

Antarctica slowly being buried by snow all winter. Full recovery is planned for November or December 2017.

2.4.4 ANITA-IV Results

Analysis of the ANITA-IV results is still underway. Some investigation into the effects of the addition of the TUFFs on the ANITA-IV flight are discussed in Chapter 3.

Chapter 3

THE TUFF, TUNABLE UNIVERSAL FRONTEND FILTER

After the ANITA-III flight, it was clear something needed to be done about the new and overwhelming CW noise from geostationary military satellites. The CW noise was observed at specific frequencies so naturally the idea of a notch filter to remove power from those specific frequencies arose. However, ANITA-III only observed this noise for a short duration during flight. There was concern within the collaboration that the noise might be more prevalent, or at different frequencies by the time the next ANITA flight took place. Due to this, the Tunable Universal Frontend Filter (TUFF) was designed. Not only would it contain three separate notch filters, but those notches could be turned on and off and moved to different frequencies. Additionally, the boards absorbed some of the functionality of other components in the ANITA-IV signal chain, now providing amplification and power for the AMPAs. Figure 3.1 shows an image of a TUFF. Each individual board handles 6 RF channels (of the 96 in ANITA-IV, one H-pol and one V-pol per antenna). The TUFFs were designed, built and tested at The Ohio State University. This Chapter outlines that process.

3.1 TUFF Design

The TUFFs were deigned to fulfill two objectives. The first was to allow tunable notch filtering so ANITA-IV could deal with satellite CW noise. The second was to replace the components that preformed the second stage of amplification, the AMPLITE, and the bias

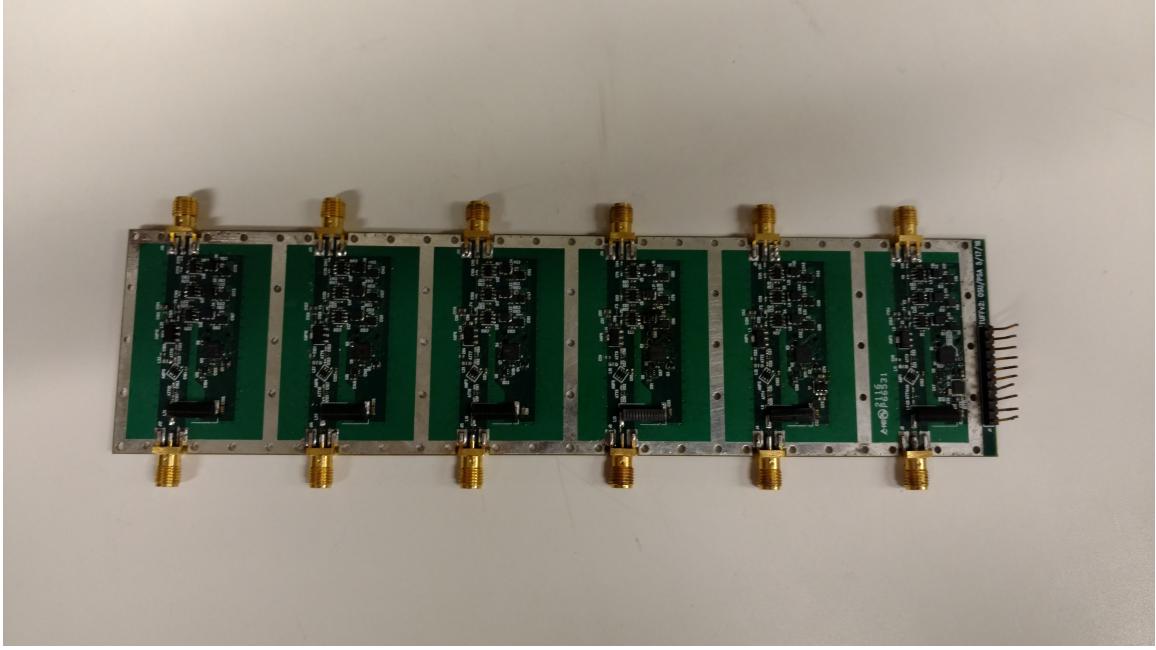


Figure 3.1: Image shows final design of the TUFF board used in ANITA-IV flight. Channel 0 is the channel furthest to the right, while channel 5 is furthest to the left.

tee that powered the AMPAs in previous experiments. The TUFF needed to do this while being light weight, small, and drawing as little power as possible due to the restrictions of ANITA being a balloon experiment.

3.1.1 Notch Filter Design

The TUFF's notches operate as LC filters. The LC circuits contain a variable capacitor to allow the notch location in frequency to be tuned to a different value.

When the inductor and capacitors in an LC circuit are driven at a specific frequency, f_0 , the circuit opens up. Other frequencies are thus filtered out. f_0 is related to the impedance of the inductor, L , and the capacitance of the capacitor, C , as follows.

$$f_0 = \frac{1}{2\pi * \sqrt{C * L}} \quad (3.1)$$

As shown in Figure 3.2, on the other side of the TUFF's LC circuits is ground. In the TUFFs, signals at the f_0 resonant frequency of one of the three notches follow the path of

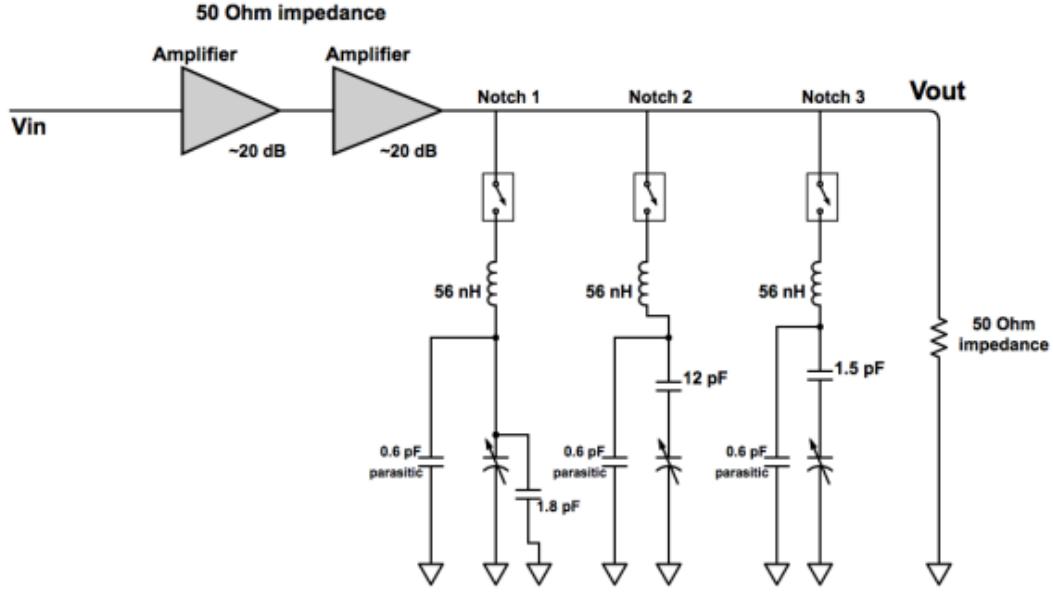


Figure 3.2: Circuit diagram containing the primary components of the TUFF notch filters [1].

least resistance to ground, and are thus removed from the RF signal path.

The TUFFs use a 56 nH inductor in their LC filters. The variable capacitors for the first notch is in parallel with a 2.4 pF capacitor. For the second notch, it is in series with a 5 pF capacitor. The third is in series with a 2.4 pF capacitor. The variable capacitor is a PE64906 variable capacitor from Peregrine Semiconductor. It can be set to 32 discrete steps in capacitance between 0.9 and 4.6 pF. Adding in these additional capacitors in series and parallel allows us to get more range in the location of the frequency notch, and additional differentiation between the notches.

In practice, the circuits on the TUFF are not ideal LC filters. The dual-pole, single-throw RF switch, shown in Figure 3.2 in each of the notches have parasitic ‘ON’ resistance. Other components in the filter also add DC resistance. The resulting resistance of the notch filters is $\sim 6\text{-}7\Omega$. Equation 3.1 holds for the resonant frequency of a LCR circuit. However, while an ideal LC circuit would create an infinitely deep notch, the actual notch is limited by the resistance in the circuit. Each notch is ~ 13 dB deep. The default positions for the

three notches are 260, 380 and 460MHz. Figure 3.3 shows the three notches outlined in their locations on a single channel of an assembled TUFF board, as well as the locations of other key components.

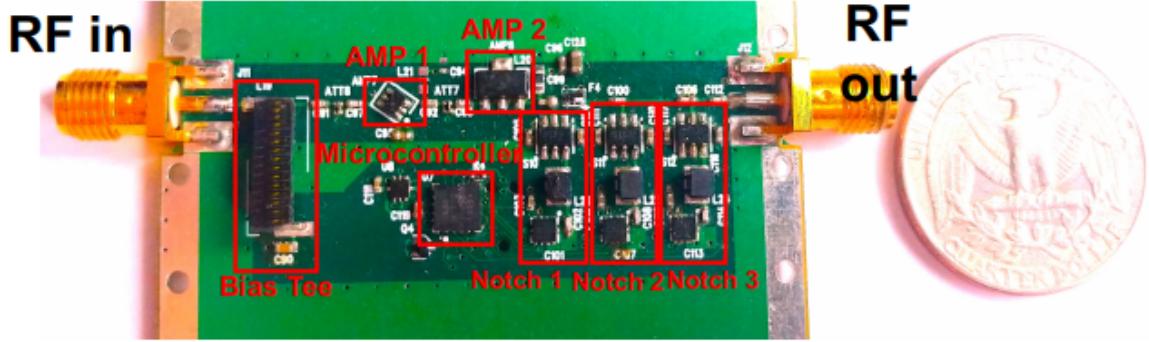


Figure 3.3: Photo of one channel of a TUFF next to a USA quarter coin for scale. Key components of the TUFF are outlined in red [1].

3.1.2 Amplifications and Bias Tee

The Amplification on the TUFFs is performed by two separate amplifiers in series. The two amplifiers combined provide ~ 45 dB of amplification. They are outlined in Figure 3.3 as AMP 1 and AMP 2. AMP 1 is a BGA2851 by NXP Semiconductors. AMP 2 is an ADL5545 by Analog Devices. There is one dB of attenuation between the two amplifiers and between AMP 1 and the Bias Tee to prevent any unpredictable interactions between the components. The two amplifiers together consume ~ 335 mW of power, which is the vast majority of a single channels power consumption. The bias tee on each channel is responsible for supplying DC power to the AMPAs which reside up the signal path on the other end of a coaxial cable. Each bias tee is constructed out of a 4310LC inductor by Coilcraft and a 0.1 F capacitor which are in series. The Inductor supplies the AMPA with the DC power. The capacitor prevents the power from entering the TUFF's signal path while allowing the RF signals coming from the AMPAs to continue through to the TUFFs amplifiers.

3.1.3 Microcontroller

Having variable capacitors would be pointless without a method to program them. This is why each TUFF board has an onboard microcontroller for each channel. They are ultra-low power MSP430G2102 microcontrollers by Texas Instruments. The microcontroller controls all three notches for a single channel on a TUFF. It can turn the notches on and off using the switches on each individual notch. It can also set the value of the variable capacitor in each notch. This was even possible during the ANITA-IV flight if notches needed to be tuned to deal with new or changing sources of CW noise. The Microcontroller has very low power consumption and 5 low power modes. In active mode, the microcontroller only consumes ~ 0.5 mW. In low power modes it can consume as little as ~ 0.2 μ W. The power consumption from the microcontrollers on each channel is negligible compared to the amplifiers.

3.1.4 Designed Integration into the existing ANITA Experiment

The TUFFs effectively replaced existing ANITA components, and were designed to fill the space left behind by those components. In past ANITA experiments the second stage amplification and bias tee were inside the Internal Radio Frequency Control Modules (iRFCMs). Figure 3.4 shows one of the iRFCMs with its lid off and TUFF boards inside. Each iRFCM handles 24 RF channels. There are 4 iRFCMs in the ANITA-IV experiment. The TUFFs were designed to sit inside the iRFCM in two stacks of two, end to end. RF shielding containers were designed for the TUFFs to help them nicely stack on top of each other. These units are referred to as TUFF sandwiches. Two TUFF sandwiches can be seen in Figure 3.4 in place inside of the iRFCM connected by short SMA cables to the iRFCM's 24 RF inputs and 24 RF outputs. Additionally, inside each iRFCM is a TUFF master and a TUFF master booster. These PCBs are the mediators between the TUFFs, the TUFF's power supply, and ANITA's CPU.

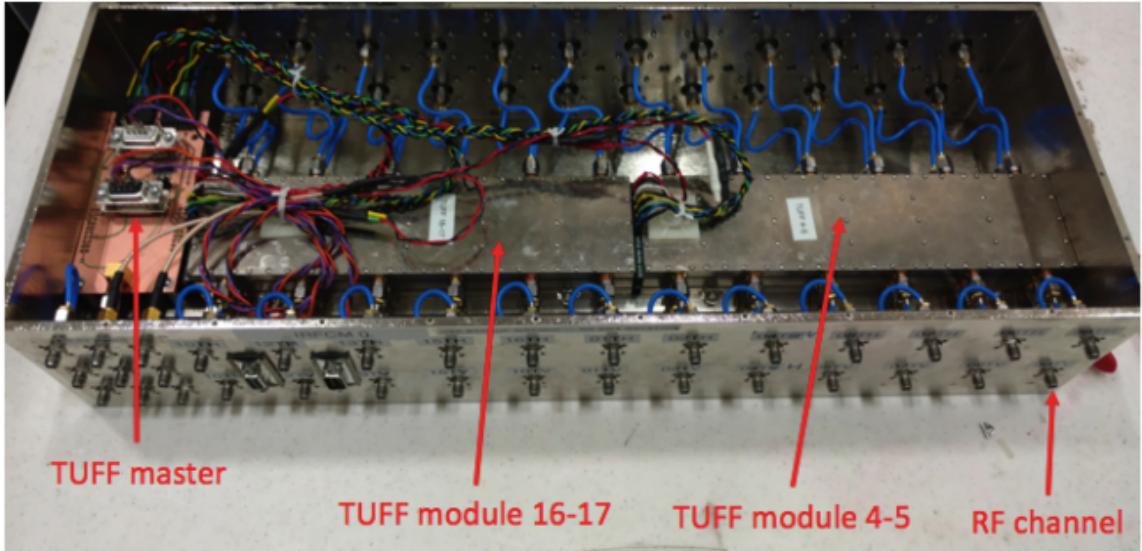


Figure 3.4: Photo showing one iRFCM with its lid off. Inside two TUFF sandwiches are connected to their RF inputs and outputs. All 4 of the TUFFs inside the 2 TUFF Sandwiches are also connected to the TUFF Master through the TUFF Master Booster for power and communication with ANITA’s CPU [1].

3.2 TUFF Construction

The TUFF printed circuit boards, PCBs, were spun by Advanced Circuits. Twenty four boards were spun but one was found to have missing vias and holes making the PCB useless. The addition of components was done by hand predominantly with the use of a rework station and soldering irons. The first round of construction saw 18 boards built and tested. Which was enough to fill all four iRFCMs and two spares. These went down to Palestine, Texas for integration into the payload. Both spares ended up being used and we returned to the Ohio State University lab with two broken TUFFs.

A second round of construction and repair was done to send to the University of Hawaii as a precaution for a final round of testing before the payload was shipped by boat to the McMurdo Base in Antarctica. During testing, three of the 96 RF channels showed atypical notch behavior. One TUFF board was repaired and one new TUFF board was constructed. A third and final round of construction was done before heading down to the McMurdo base in Antarctica. All three remaining bare TUFF boards were assembled and any previously

broken TUFFs were repaired. Three complete and tested TUFF sandwiches went south to be swapped out in case of any further problems. One half sandwich was also sent to be used as a TUFF channel bypass in an unavoidable and unfortunate circumstance. No TUFF board required repair or replacement in Antarctica.

3.3 TUFF Debugging and Testing

No testing was done until after assembly was completed for the first round of construction. In retrospect this turned out to be a mistake as it made any problems that existed much more difficult to diagnose. In future rounds of construction testing was performed incrementally after sections of the board were constructed.

The first stage of debugging was simply to check that none of the 10 TUFF input pins were shorted to each other. The pins are, in order, Bias T V_{in} , Ground, 5 V V_{in} , Ground, Reset, Ground, Clock, Ground, Data. Several locations on the board were also checked to assure they were not shorted to the 3.3 V. Next, boards were connected to a power supply set to 5 V in order to power them through the 5 V pin to check that they were drawing approximately the correct amount of current and that that current draw was stable. The correct amount of current for a completed board to be drawing was experimentally determined to be ~ 380 mA. Once the typical current draw was experimentally identified the current limit of the power supply was set to 0.5 A to try to avoid allowing any component on the board to exceed their current limits. In the case when a board drew an irregular amount of current it was powered off immediately. Lower or higher current draws indicated incorrect behavior of one of the components that drew a substantial portion of that power, usually one of the amplifiers or either of the voltage controllers. Unstable current draws indicated that the board may have intermittent connections on one component or one of the components was in an unstable state. In either case, we attempted to fix these problems before dealing with anything else, by closely examining the board under a microscope to identify suspicious solder joints then by reflowing and fixing said joints. Common sources of problems were the 4.8 V regulator and the 3.3 V regulators along with the second amplifier

in the signal chain. If any of these were missing a connection they would not behave correctly and would draw too little or too much current. The second amplifier in particular had a very large ground plane in thermal contact to its ground pins which often lead to those pins escaping construction still unsoldered.

After we had confirmed that the board was drawing approximately the correct amount of current, we proceeded to assuring the notches and amplifiers were behaving correctly. To check the Amplifiers each channel on each TUFF was connected to a Network analyzer and powered to view the gain across the channel. A TUFF with working amplifiers shows a flat gain of ~ 45 dB from 100-1200MHz. Figure 3.5 shows the typical gain across a channel with all notches off and all notches on at default values.

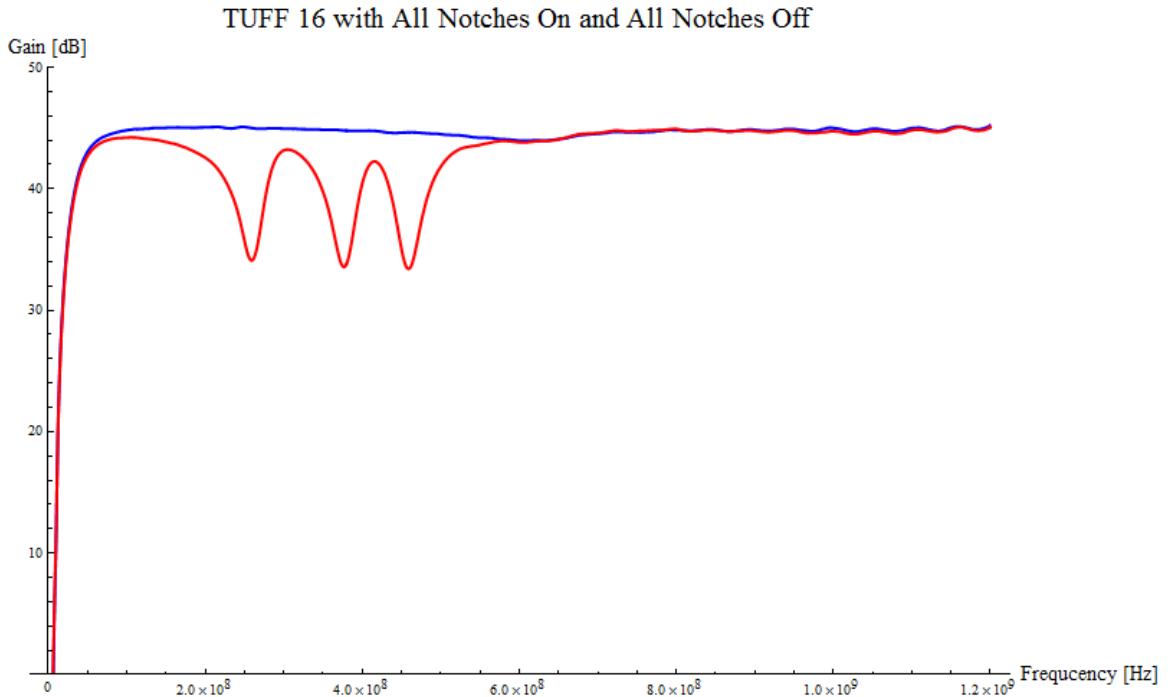


Figure 3.5: Image shows gain across channel 1 of TUFF 16. The blue line shows the TUFF with all notches off. The red line shows the TUFF with all notches on. Data taken by Oindree Banerjee.

In order to test the notches, the microcontroller on each channel needed to be pro-

grammed. The method of programming the individual channels is outlined in detail in Section 3.4. Once programmed, the functionality of the notches could be tested with a built in ‘Test Mode’.

3.3.1 Test Mode

Test mode for a TUFF could be activated by shorting CLOCK to DATA then shorting the RESET input to ground while the TUFF’s 5 V input is powered. Once activated, the TUFF will proceed through a series of set states. First, it will turn the first notch on with the variable capacitor set to its minimum (which means the notch will be at its maximum frequency position). Next, it will move the first notch to its minimum frequency position. It will then turn the first notch off. It then does the same thing with the second and third notches. This allows us to test that all three notches’ switches and variable capacitors, as well as each channel microcontroller are working correctly. If something does not work correctly, test mode narrow the problem down to only a few components which we could then fix.

3.3.2 Thermal Stress Testing

After all of the components on the TUFFs were working, the solder joints were stress tested by thermal cycling them over night. TUFFs were heated to 65° C for 15 minutes, and held there for 10 minute. They were then cooled to -35° C for 30 minutes, and held at that temperature for 15 minutes. Finally, they were heated back to room temperature for 10 minutes, and held there for another 10 minutes. This was repeatedly cycled overnight, to make sure the TUFFs and their solder joints could survive any temperature gradients they would experience during flight. After thermal testing the TUFFs again underwent checks of their current draw and notch functionality.

3.3.3 Tuff Integration

TUFFs were placed inside of their sandwiches in pairs of two after thermal testing. Once inside of their RF insulating containers the TUFFs were tested again. First, the current

draw was checked. Then, all 12 channels of each sandwich underwent test mode testing again. Nothing should have changed just by placing them inside of their cases, but, in a few cases, metal shavings from the screw holes of the sandwiches got inside of a case and caused a short. During initial testing of the cases it was also discovered that RF insulation was required inside of each channels cavity to prevent reflections inside the cavity from causing non flat gain across frequency. Figure 3.6 shows an image of a single TUFF sandwich.

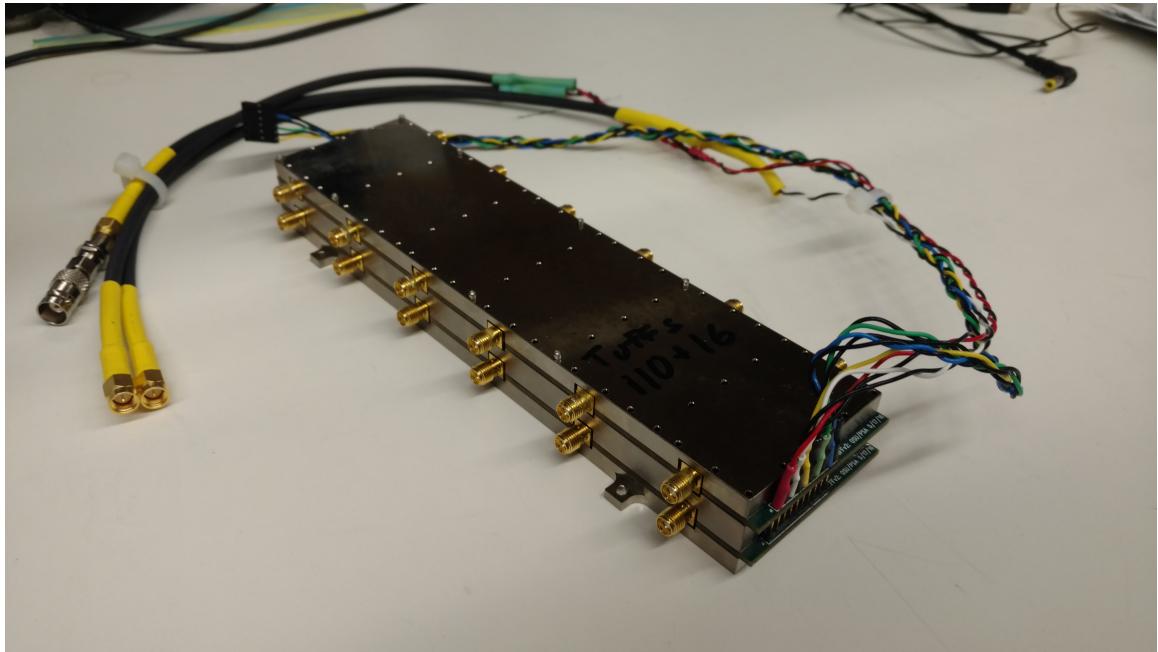


Figure 3.6: Photo shows a single TUFF sandwich containing two TUFF PCB boards. The sandwich is complete with it's power and data cables.

Once the TUFFs were found to be fully operational inside of their sandwiches their input pins were cabled up. Only the bottom TUFF in a sandwich has pins. Those pins go through the pin holes of the top TUFF in the sandwich. The bottom TUFFs in sandwiches are even numbered, while the top TUFFs are odd. Once a sandwich is ready to be wired the pins from the bottom TUFF are soldered to the pin holes from the top TUFF. The order of TUFFs pins is again Bias T V_{in} , Ground, 5 V V_{in} , Ground, Reset, Ground, Clock, Ground, Data and Ground. The wires connected directly to the TUFF and the insulating

heat shrink around the pins were color coded. Black is ground. Red is Bias T V_{in} . White is 5 V V_{in} . Yellow is reset. Green is clock, and blue is data. Each of the colored wired has its own ground, which it was twirled with. The bias T V_{in} and 5 V V_{in} wires were spliced with SMA cables which were then screwed into the iRFCM case for power. The other three were braided together and run into a 6 wire housing connector which then clipped into the TUFF Master Booster.

With the TUFF sandwiches wired, they need to be integrated into their iRFCMs. Inside of each iRFCM sits two TUFF sandwiches, as shown in Figure 3.4. They are fastened to a metal support that runs the length of the iRFCM. Each of the two TUFF sandwiches' 24 RF inputs and outputs were connected to the iRFCM's RF channel inputs and outputs by 4 inch flex SMA cables. Each connection was threadlocked using Loctite to assure no connections loosened during flight due to vibrations. On one end of the iRFCM the TUFF Master and TUFF Master Booster rested, also fastened to the metal support. Both of the TUFF sandwiches have their reset, clock and data connectors plused into the TUFF Master Booster. The TUFF sandwich's power cables screw into the iRFCM case to get power from ANITA's power supply outside of the iRFCM. The TUFF Master Booster also gets its power through a SMA cable connecting to the iRFCM case. The iRFCMs talk to each other and ANITA's CPU through DB9 Serial connections. Each TUFF Master Booster has an input and output serial connection which connects to the iRFCM case. Once outside of the iRFCM case, all of the serial cables connect to a serial loom board, which passes signals from ANITA's CPU to each of the iRFCMs. The searial loom board is shown in Figure 3.8.

3.3.4 Vacuum chamber Testing

Vacuum chamber testing was the final test the TUFFs underwent before leaving for Antarctica. At this stage the TUFFs were inside of their iRFCMs with the TUFF Masters and TUFF Master Booster boards. This was done to assute that in the vacuum like conditions the ANITA experiment operates in, the TUFFs and TUFF masters would not undergo run-away heating. The primary worry was that the microcontrollers would heat up, and without air to help them radiate this extra heat away, they would just continue heating up until they

burned themselves out. To test this the iRFCMs were placed into University of Hawaii’s vacuum chamber overnight. While inside a near vacuum was maintained. In order to simulate a stressed state for the microcontrollers, commands to turn off and on the notches and move the notches were issued every microsecond. Temperature sensors were placed on several key locations, including the TUFF Master Booster’s microcontroller. Ideally, we would have allowed the tests to continue until the temperature reached an equilibrium or came close to leaving the device’s operating range. However due to time constraints, we could not wait that long. When testing was stopped however, it was clear that the rate of temperature leveling off below max operating temperatures and we did not need to worry about runaway heating for the TUFFs or the TUFF master. After the Vacuum testing, it was determined that the TUFFs and TUFF Masters, inside of the iRFCMs, were ready to be integrated into the instrument crate to be shipped to Antarctica.

3.4 TUFF Programming

Each microcontroller on each channel of each TUFF board needed to be programmed after it was assembled. This was done using the software Code Composer Studio by Texas Instruments and a TI flash programmer. On the back of each channel of each TUFF board was a place to make electrical contact with the TUFF. The firmware for the TUFFs was written at the Ohio State University by Patrick Allison. Each microcontroller needed to be told which channel of the TUFF it belongs to. The identification also varied depending on whether the TUFF was going to be the bottom (even) or top (odd) TUFF in its TUFF sandwich. The bottom TUFF in a sandwich had its channels labeled 0-5, with channel 0 being at the bottom of the TUFF. The top TUFF in a sandwich had its channels labeled from 6-11, with channel 6 being at the bottom of the TUFF. Table 3.1 shows the address assigned to each channel within a TUFF. The address must be changed in the firmware for programming each individual channel.

Bottom TUFF		Top TUFF	
Channel	Address	Channel	Address
00	0x01	06	0x41
01	0x02	07	0x42
02	0x04	08	0x44
03	0x08	09	0x48
04	0x10	10	0x50
05	0x20	11	0x60

Table 3.1: Table shows addresses for all 12 TUFFs in a signle TUFF Sandwich

3.5 TUFF Master, TUFF Master Booster and TUFF Loom

The TUFF Master is a commercial Tiva C series Launch Pad board by Texas Instruments and it issues an iRFCM's TUFFs commands. The TUFF Master Booster is a board designed and constructed at the Ohio State University. It is responsible for Connecting the TUFF Master with the TUFFs, incoming and outgoing serial connections and incoming power. Figure 3.7 shows the TUFF Master and the TUFF Master Booster side by side. The Master is on the right while the Master Booster is on the left. The pins of the Master plug into the back of the Master Booster.

The TUFF Serial Loom board is a simple board designed to fan out commands from ANITA's CPU to all of the iRFCMs. It was designed and constructed at the Ohio State University. Figure 3.8 shows an image of the TUFF Serial Loom board. It also transfers any responses from the four iRFCM's TUFF Masters back to the CPU.

3.5.1 TUFF Master Programming

The TUFF Master was programmed through a micro USB cable using the software Energia. The firmware was written by Patrick Allison from the Ohio State University and Cosmin Deaconu from the University of Chicago. Detailed instructions on how to program the TUFF Master can be found in Appendix C.

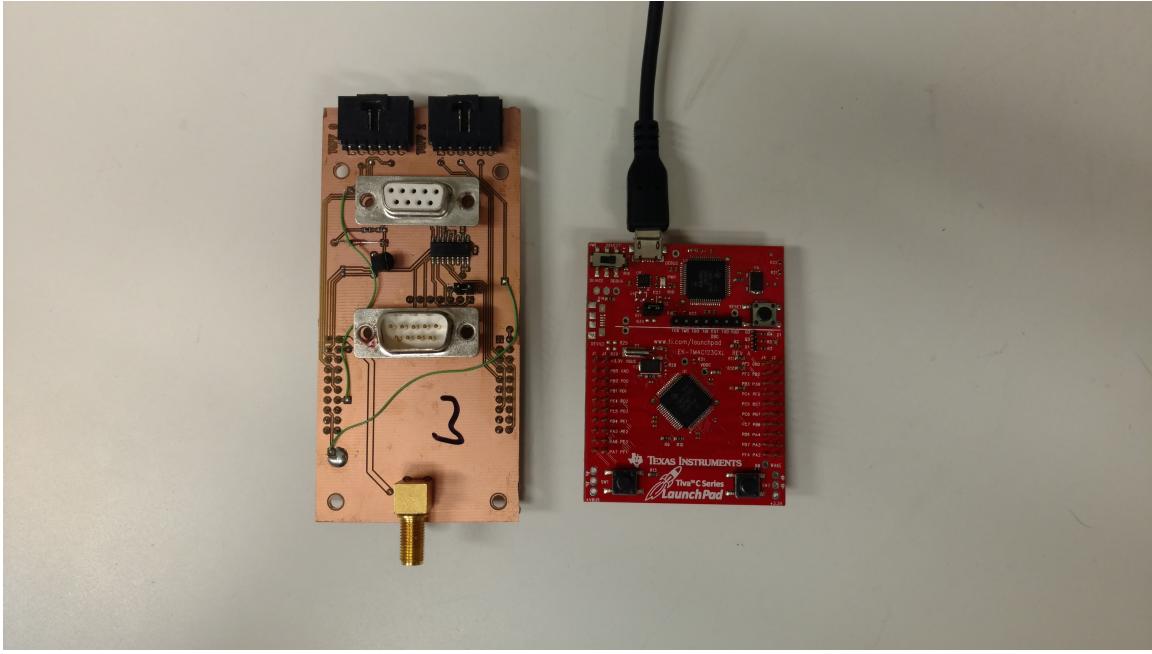


Figure 3.7: Photo shows the TUFF Master and TUFF Master Booster. The TUFF Master is on the right, the TUFF Master Booster is on the left.

3.6 TUFF performance during ANITA-4 Flight

The TUFF's notch filters were crucial to the ANITA-IV flight. Figure 3.9 shows when the three programmable notches were active during ANITA-IV's flight. Solid colors represent the notches being enabled. Hatched black lines represent the notch being inactive. On multiple occasions during the ANITA-IV flight the TUFFs notches were observed to greatly assist in blocking out CW noise and thus reduce the trigger rate.

3.6.1 Notch 1 During Flight

During ANITA-IV's flight, the 260MHz notch (notch 1) was on nearly continuously. ANITA-III saw a strong CW signal at 260MHz throughout the flight due to military satellites. The CW peaks that ANITA-III saw during flight can be seen in Figure 3.10. The same CW signal was consistently present during the ANITA-IV flight as well, which is why Notch 1 was kept on for the vast majority of the flight. On December 14th 2016 Notch 1 was retuned to be centered at 250MHz, then later that same day returned to 260MHz. This was done to

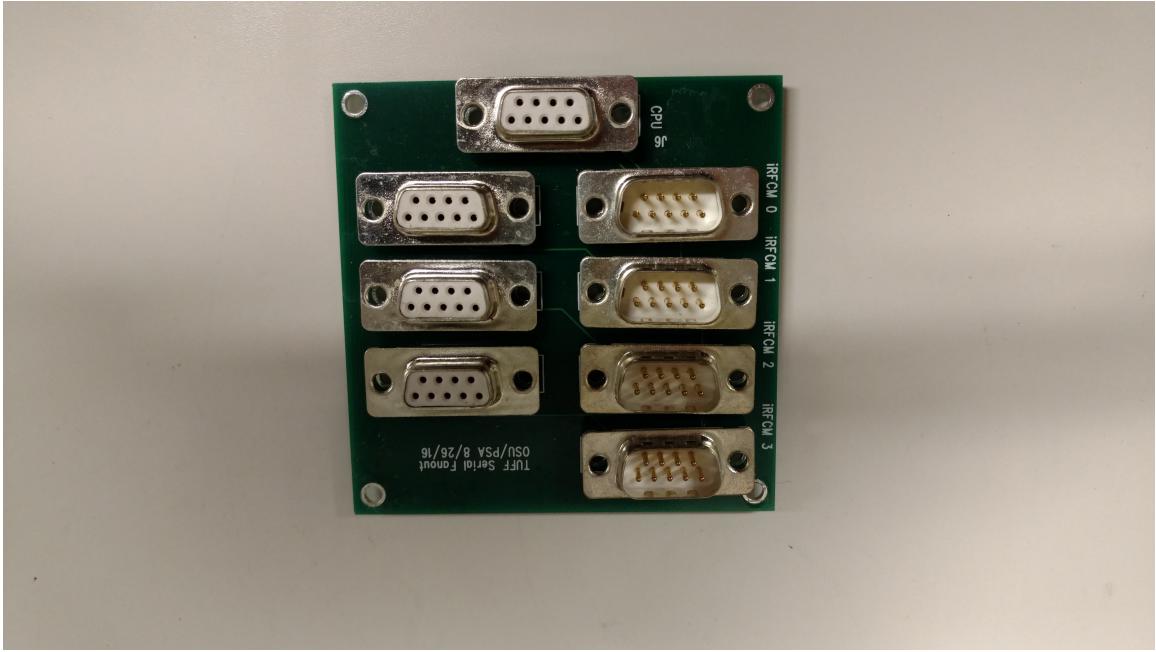


Figure 3.8: Photo shows the TUFF Serial Loom board

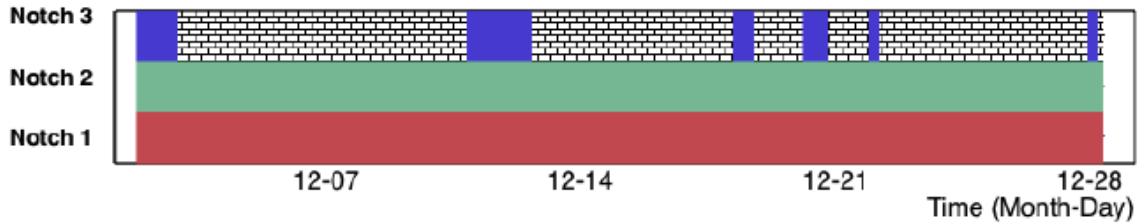


Figure 3.9: Plot shows when the three programmable notches were active during the ANITA-IV flight. Solid colors mean the notch was active. Hatched black line mean the notch was off. Notch 1 is red. Notch 2 is green. Notch 3 is blue [1].

counter act observed CW noise at 250MHz. For all of the rest of the flight however Notch 1 remained on and tuned to 260MHz.

3.6.2 Notch 2 During Flight

Like Notch 1, Notch 2 was on almost continuously during the ANITA-IV Flight, as shown in Figure 3.9. However it was retuned many more times than Notch 1. During the ANITA-III flight the observed CW peak at 380 MHz, shown in Figure 3.10, was intermittent. During

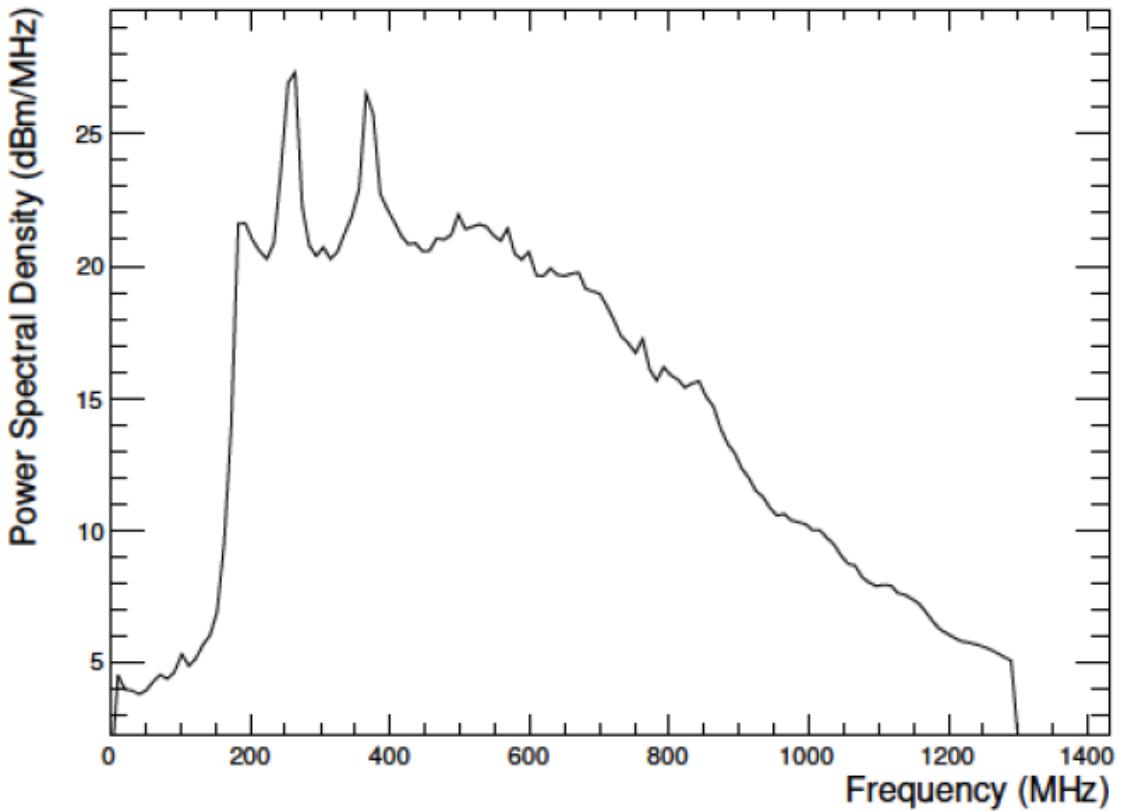


Figure 3.10: Plot shows the CW peaks at ~ 260 MHz and ~ 380 MHz observed by ANITA-III [1].

the ANITA-IV flight however it was consistently present. This change was likely due to the addition of more Military satellites into geostationary orbit between the ANITA-III and ANITA-IV's flights. Notch 2 was turned off only a few times during flight. The first time was on December 2nd 2016. The second time was on December 19th 2016. It was inactive for 16 minutes on December 2nd. Excess CW noise was observed in nearly all phi sectors around 370 MHz and the Notch 2 was reactivated. It was inactive for 10 minutes on December 19th. A dramatic rise in the trigger rate was observed and Notch 2 was turned back on. Excess CW noise was also observed in several phi sectors during the 10 minutes Notch 2 was inactive.

Notch 2's frequency location was also moved several times during the ANITA-IV Flight between December 6th and December 8th 2016. This was done to attempt to dynamically

remove CW noise in the 360 to 390 MHz range. The attempts succeeded in lowering ANITA’s trigger rate, and the notch was retuned to 380MHz for the remainder of the flight.

3.6.3 Notch 3 During Flight

Notch 3 was only used when ANITA was near large Antarctic bases. It was tuned to notch out 460 MHz CW noise for the entire flight. Notch 3’s status during flight can be seen in Figure 3.9. On December 2nd, Notch 3 was deactivated for a few moments, but quickly turned back on because ANITA-4 was still near the McMurdo base.

3.6.4 Results from the TUFF being included in ANITA-IV

During the ANITA-III flight, the two methods of preventing dead time due to high trigger rates were adjusting the individual signal band’s trigger thresholds in the L1 trigger (see Sections 2.1.3 and 2.3.2) and phi masking (see Section 2.2.2). Lowering a signal band’s threshold meant that more power was required to pass the L1 Trigger for any individual channel, making triggers less frequent. Thus, ANITA loses sensitivity to lower energy signals then it’s thresholds are lowered. Similarly, when phi sectors needed to be masked, no events coming from that phi sector can cause a trigger. So like with adjusting thresholds, Turning on phi masking causes a loss in sensitivity. Figure 3.11 shows the SURF DAC thresholds in arbitrary units plotted against time for ANITA-III and ANITA-IV. Figure 3.12 shows the fraction of the sky being masked by phi masking plotted against time for ANITA-III and ANITA-IV. Due to the added frequency notching from the TUFFs, the L1 signal band thresholds were consistently more stable and higher during the ANITA-IV flight when compared against the ANITA-III flight. Additionally, less phi masking was required during the ANITA-IV flight.

The reduction in phi masking and L1 band threshold adjustment both contribute to an increase in livetime. ANITA experiments have two different types of livetime. The first is digitization livetime. This takes into account the deadtime due to digitization of the signals. Both phi masking and adjusting the L1 individual band thresholds are designed to reduce the digitization deadtime by reducing the frequency of triggers. The second type is

instrument livetime. The instrument livetime is the digitization livetime multiplied by the fraction of unmasked phi sectors, after accounting for masked channels. Instrument livetime accounts for the amount of azimuthal ice the experiment is sensitive to at a given time.

Figure 3.13 shows the digitization livetime for ANITA-III and ANITA-IV plotted against time. Figure 3.14 shows the instrument livetime for ANITA-III and ANITA-IV plotted against time. ANITA-IV has higher digitization livetime than ANITA-III, with 92.3% digitization livetime compared to ANITA-III’s 73.7% digitization livetime. The TUFF’s notch filtering kept trigger rates due to CW noise low enough to mitigate deadtime due to digitization. The improvement in instrument livetime is even more dramatic. ANITA-IV had an instrument livetime of 91.3%, while ANITA-III had an instrument livetime of 31.6%. That is nearly a factor of 3 improvement in instrument livetime. This incredible improvement was achieved primarily due to the TUFF’s drastically reducing ANITA-IV’s reliance on phi masking. Investigations into the effects of the TUFFs on neutrino sensitivity are still underway.

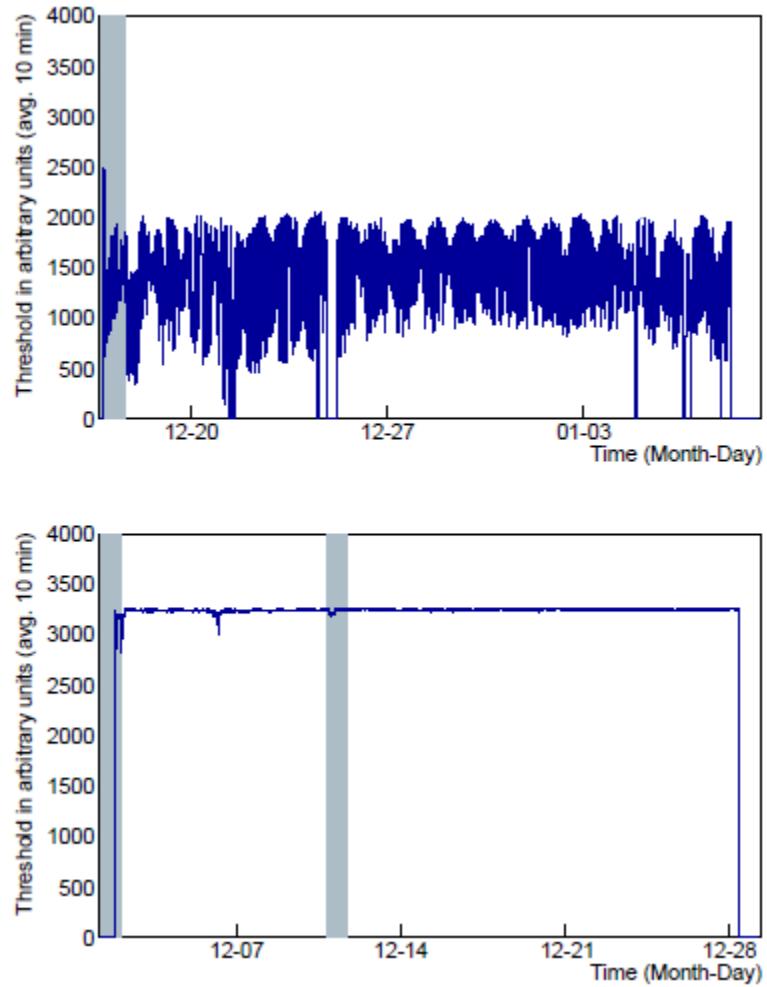


Figure 3.11: SURF DAC thresholds for a single channel in arbitrary units plotted against time for ANITA-III (top) and ANITA-IV (bottom). A lower threshold corresponds to a higher, and thus stricter, requirement on the power of the signal. The gray shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility [1].

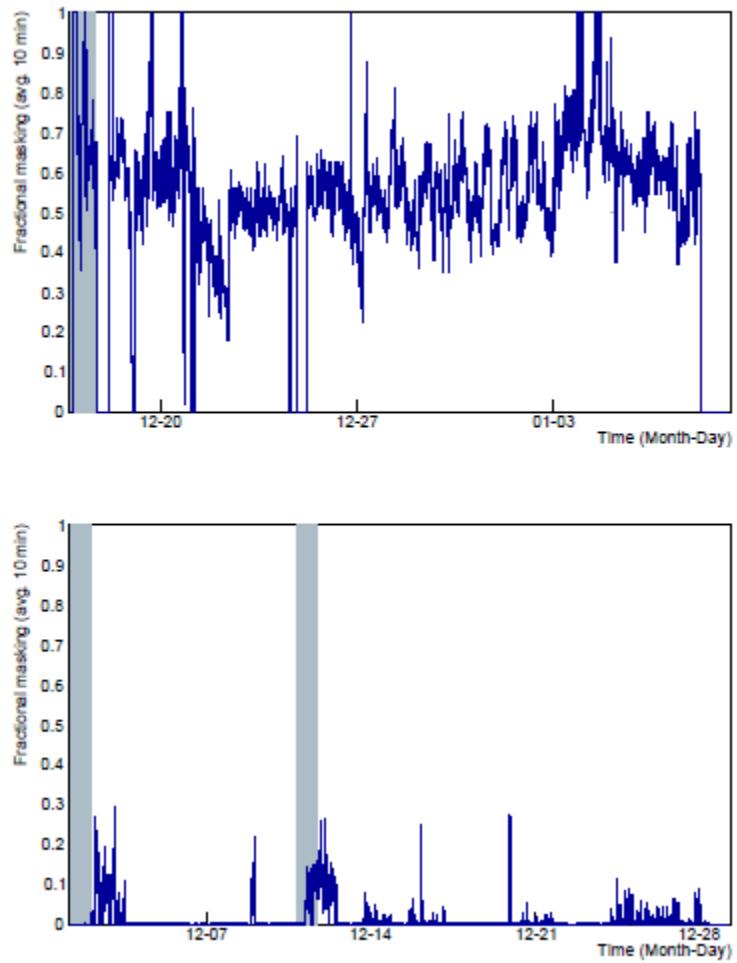


Figure 3.12: The fraction of the sky being removed by phi masking during flight plotted against time for ANITA-III (top) and ANITA-IV (bottom). The gray shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility [1].

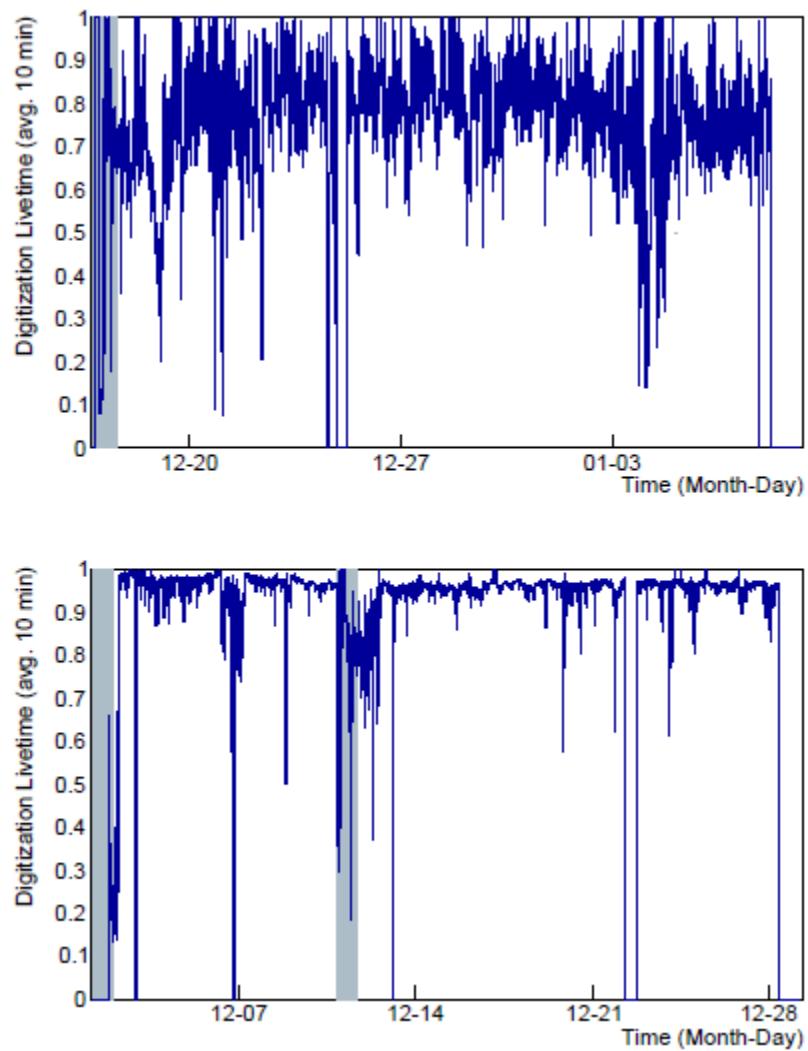


Figure 3.13: The digitization livetime for ANITA-III (top) and ANITA-IV (bottom) plotted against time. The gray shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility [1].

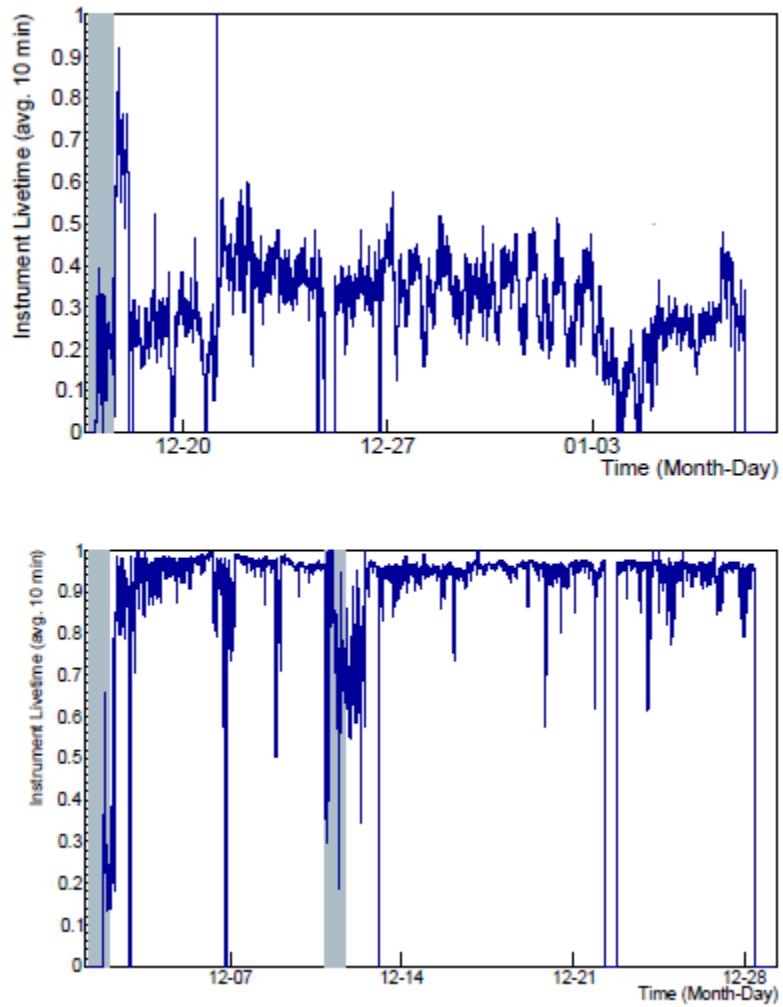


Figure 3.14: The instrument livetime for ANITA-III (top) and ANITA-IV (bottom) plotted against time. The gray shaded regions indicate when the ANITA payload was in line of sight of the NASA LDB Facility [1].

Chapter 4

THE HYBRID

As mentioned in Sections 2.1.2 and 2.4.2 the hybrids are a hardware component in the ANITA-I and ANITA-IV trigger paths intended to convert the incoming H-pol and V-pol signals into outgoing LCP and RCP signals. Figure 4.1 shows what the physical hybrid looks like. The hybrids are Anaren Model 10023-3 military grade 3dB 90 degree hybrid couplers encased in an aluminum shell. Incoming H-Pol signals go into the SMA input labeled IN. Incoming V-Pol signals go into the SMA input labeled ISOL. Outgoing RCP signals leave from the SMA output labeled 0 directly above the SMA input labeled IN in Figure 4.1. Outgoing LCP signals leave from the final SMA output labeled -90.

4.1 How the Hybrid Works

The way the hybrids achieve the basis change is by delaying the phase of the incoming signals by different amounts and then adding them together. Jones vectors in ket notation define RCP and LCP respectively as [45]

$$|R\rangle = \frac{1}{\sqrt{2}}(|H\rangle + (-i)|V\rangle) \quad (4.1)$$

$$|L\rangle = \frac{1}{\sqrt{2}}(|H\rangle + i|V\rangle) \quad (4.2)$$

In equation 4.2, multiplying by i is equivalent to phase delaying the vertically polarized signal by 90 degrees. This is something that can be implemented in hardware. However in

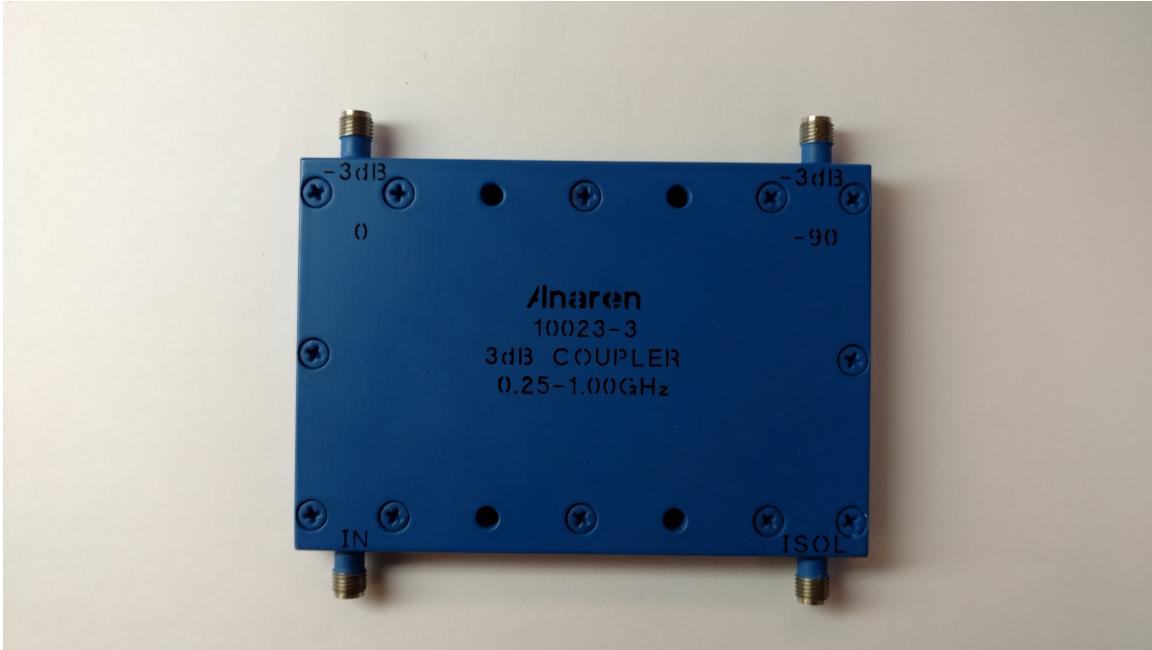


Figure 4.1: Photo of a signal antennas hybrid.

equation 4.1, multiplying by $-i$ is equivalent to moving the phase forward by 90 degrees. That would be like moving something forward in time. It is not possible. So the hybrid must not be doing exactly what equations 4.1 and 4.2 show. The important thing to preserve when converting from H-pol and V-pol to LCP and RCP, or vise versa, is the relative difference between the component signals. So as long as when converting to RCP the V-Pol signal is phase delayed by 90 more degrees than the H-pol signal, and the opposite for the LCP conversion, the transformations will work. This must be how the hybrid achieves its goal. In addition, the hybrids are totally symmetric. If you swap the signals going into the two different inputs, then the two outputs will flip as well. Equations 4.1 and 4.2 however are not symmetric. Taking this information into account equations 4.1 and 4.2 can be modified into:

$$|R\rangle' = \frac{1}{\sqrt{2}}(|V\rangle + i|H\rangle) \quad (4.3)$$

$$|L\rangle = \frac{1}{\sqrt{2}}(|H\rangle + i|V\rangle) \quad (4.4)$$

Where the equation for LCP has remained the same, and the equation for RCP has effectively been rotated by 90 degrees in phase space by multiplying the old equation by i . This modification matches the symmetry of the hybrids. These new equations still have rotations forward in phase, however. The key is that the difference of 90 degrees in phase is maintained between the two signals. The way this works in the hybrids is that both the incoming signals are delayed, but one of the incoming signals is delayed by 90 degrees more. So for example you could image that one incoming signal is being delayed by 360 degrees while the other is being delayed by 270 degrees.

4.2 Simulation of the Hybrid

Part of the ANITA analysis process is examining how a simulated version of our experiment would perform at detecting simulated neutrinos. In order to do this, the ANITA collaboration needs to simulate the entire signal path, including the hybrid. Using equations 4.3 and 4.4 makes this relatively easy.

The hybrid acts on a signals frequency component, thus in order to simulate the hybrids effect on a time domain impulsive signal we need to get frequency domain information for the simulated time domain data.

The most straight forward method to do this is to perform a Fourier Transform on the real, time domain data (V-pol or H-pol). Multiply the frequency domain real and imaginary data you get from your Fourier Transform by i . Then convert the frequency data back into the time domain with an inverse Fourier Transform, and add it to its linearly polarized counterpart (H-pol or V-pol). This works, but Fourier Transforms are fairly computationally intensive.

The more computationally efficient way, and the way that it is currently being performed in our ANITA analysis and simulation software, is to use a Hilbert transform. When you take a Hilbert transform of real time domain data you get back a new set of real time

domain data that has been delayed in phase by 90 degrees. This is exactly what we want. It also reduces the number of transforms necessary from two to one, making it about 50% more computationally efficient. In code what this looks like is represented by the following equations.

$$\begin{aligned} R &= \frac{1}{\sqrt{2}}(h(H) + V) \\ L &= \frac{1}{\sqrt{2}}(H + h(V)) \end{aligned} \tag{4.5}$$

Where $h(x)$ is the Hilbert transform of an arbitrary real time domain set of data x . These equations are a clear parallel of Equations 4.3 and 4.4 where now instead of multiplying by i we are taking the Hilbert transform of the data. How this method of simulating the hybrids compares to actual data will be discussed in Section 4.5

4.3 Characterizing the Hybrid

Though the hybrids were used in ANITA-I, a detailed study of their frequency response and pulse transmission was not done until after it was decided that they would be integrated into ANITA-IV to help our trigger deal with anthropogenic noise from geostationary satellites.

4.3.1 Experimental setup

In order to take detailed measurements of the hybrids frequency domain gain and phase, a hybrid was attached to a Agilent Technologies E5062A 300kHz - 3 GHz ENA Series Network Analyzer. At any given time only two of the inputs and outputs can be measured. The remaining two ports on the hybrid were 50Ω terminated. The two connected to the network analyzer used equal length double shielded coaxial SMA cables. Measurements of the two cables connected to each other by a male to male adapter were also taken. These measurements were used to calibrate out loss in power and phase due to the cables and anything internal to the Network analyzer to assure that all that anything we saw in our data was due to the hybrid. For all combinations of hybrid ports, the Log Magnitude gain,

phase, and group delay were measured for all scattering parameters, S11, S21, S12, and S22. A typical experimental set up as well as a photo of the hybrid connected to the network analyzer can be seen in Figure 4.2

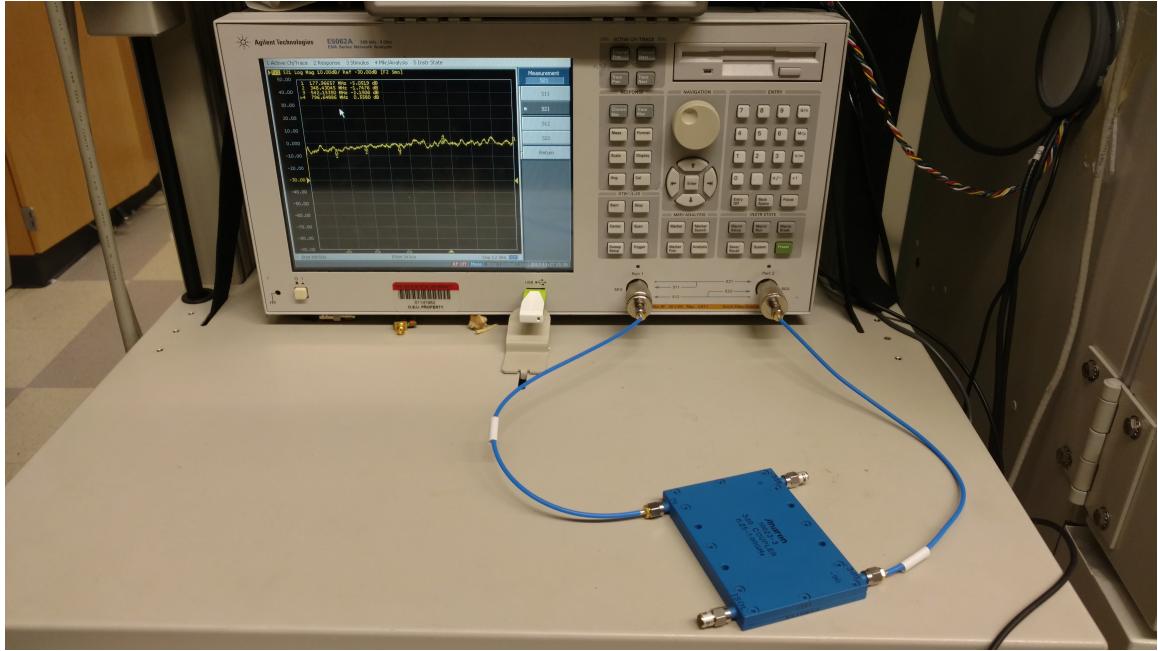


Figure 4.2: Photo shows the hybrid connected to the network analyzer used to take frequency domain measurements of the hybrid.

In addition to measurements of the frequency domain gain and phase, measurements of the hybrids effects on an askarian neutrino signal like pulse were taken using a Tektronix TDS7254B Digital Phosphor Oscilloscope. Figure 4.3 shows a diagram representing the experimental setup used to achieve this.

A pulse generator was used for creating the initial Askarian like pulse. It was then filtered through the same type of bandpass filter used in the ANITA signal path, a Lark ENG 9AB Filter. The signal was then split with a 4 way Minicircuits splitter. Each of the four paths were then passed through a high pass and low pass filter. Unfortunately our lab only had access to one Lark 5AP Lowpass filter. As a substitute the other three paths were run through a Minicircuits low pass filter. The Minicircuits filter is known to have a

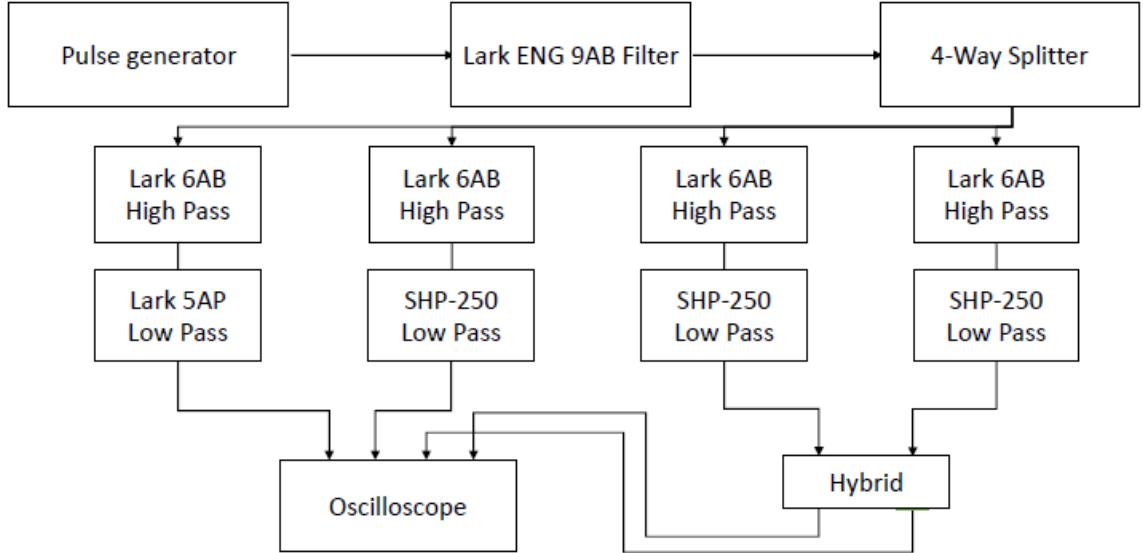


Figure 4.3: Diagram showing the experimental setup used to collect data on the hybrids effect on an askarian like pulse.

different frequency response than the Lark 5AP filter used in the ANITA flight. However because we are doing this to compare the hybrids effect on the pulse, the difference does not matter. All four paths are run through a Lark 6AB Highpass filter. One of the paths that used a Minicircuits lowpass filter and the path that used the Lark low pass filter then continue on directly to the oscilloscope, through two SMA cables in a row each. The other two paths go into the inputs of the hybrid. The outputs of the hybrid then go into the oscilloscope. All wire lengths and types are matched for each path to account for delays.

This experimental set up is limited due to the fact that the same pulse is going into both of the inputs of the hybrid. The main purpose however was to give a comparison point to simulation, and in that, it is successful. Data was also taken with one of the input of the hybrid and one of the outputs of the 4 way splitter being terminated in order to give a second configuration to compare simulation against.

4.3.2 Measurements

Network analyzer data showed excellent agreement with expected hybrid behavior. A 3db drop in signal is expected through the hybrid. Log magnitude frequency domain gain data

is shown in Figure 4.4. Inputs and outputs of the hybrid are labeled ‘A’ ‘B’ ‘C’ and ‘D’ for convenience. The bottom left port labeled IN in Figure 4.1 is ‘A’. Labeling proceeds alphabetically around the hybrid counterclockwise from there. On average the hybrid shows about -3.25dB gain instead of the expected -3db. The hybrid has near perfect symmetry. Data from A→C and from B→D, the long paths that cross the hybrid, have near identical gain, this their two lines in Figure 4.4 matching nearly perfectly. Data from A→D and from B→C, the short paths that do not crisscross the hybrid, also have near identical gain. The gain from the long and short paths are anticorrelated with each other. As used in ANITA, any output from the hybrid will be combining long and short path data together into an output, so this anticorrelation should result in a flatter total gain in frequency.

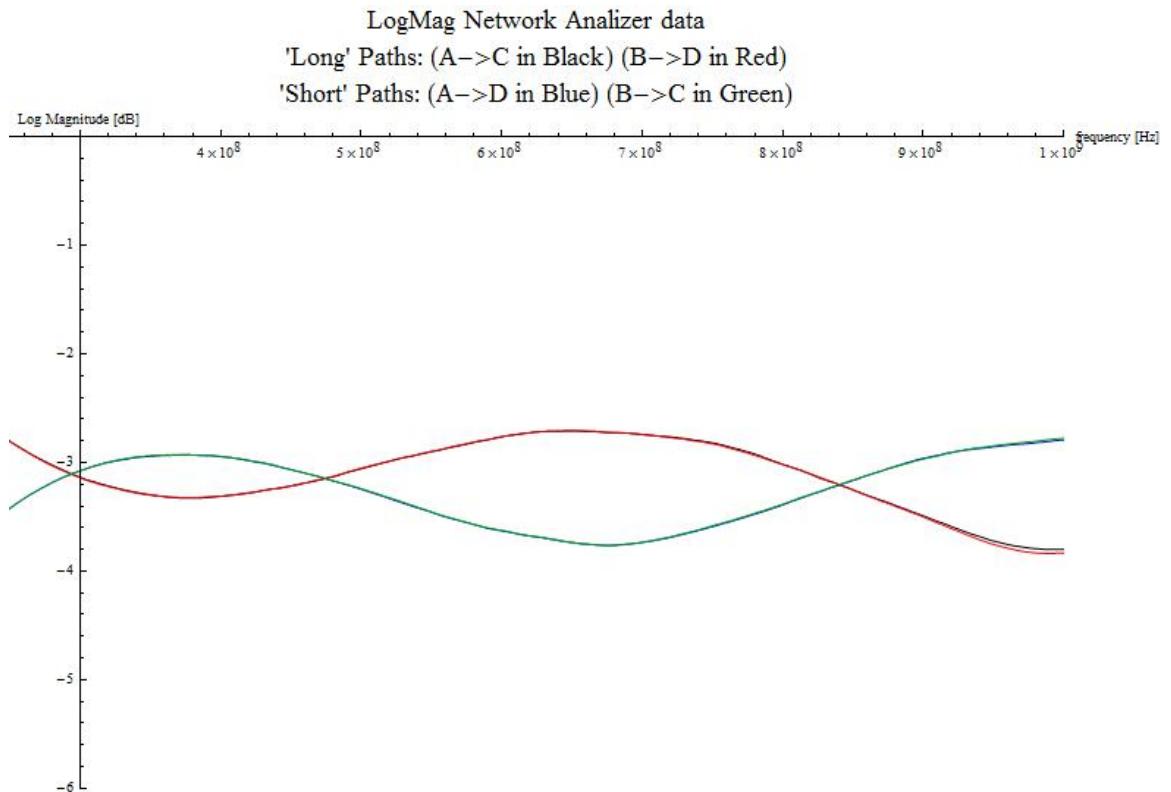


Figure 4.4: Figure shows gain of the hybrid for the 4 different input/output combinations possible. Average gain is ~ -3.25 dB.

Phase delay data also matches ideal hybrid behavior very well. Figure 4.5 shows the unwrapped phase of the long and short paths. As discussed in Section 4.1 we expect one of the paths to be delayed 90 degrees when compared to the other. That is what we observe. This is shown even more clearly in Figure 4.6, which shows the difference in phase between the two different paths that combine to create the hybrid outputs. The difference in phase, for both outputs, holds incredibly steady at 90 degrees up to about 1000MHz where it starts to fall off some. All of the network analyzer data is clear evidence that the hybrids are behaving as expected.

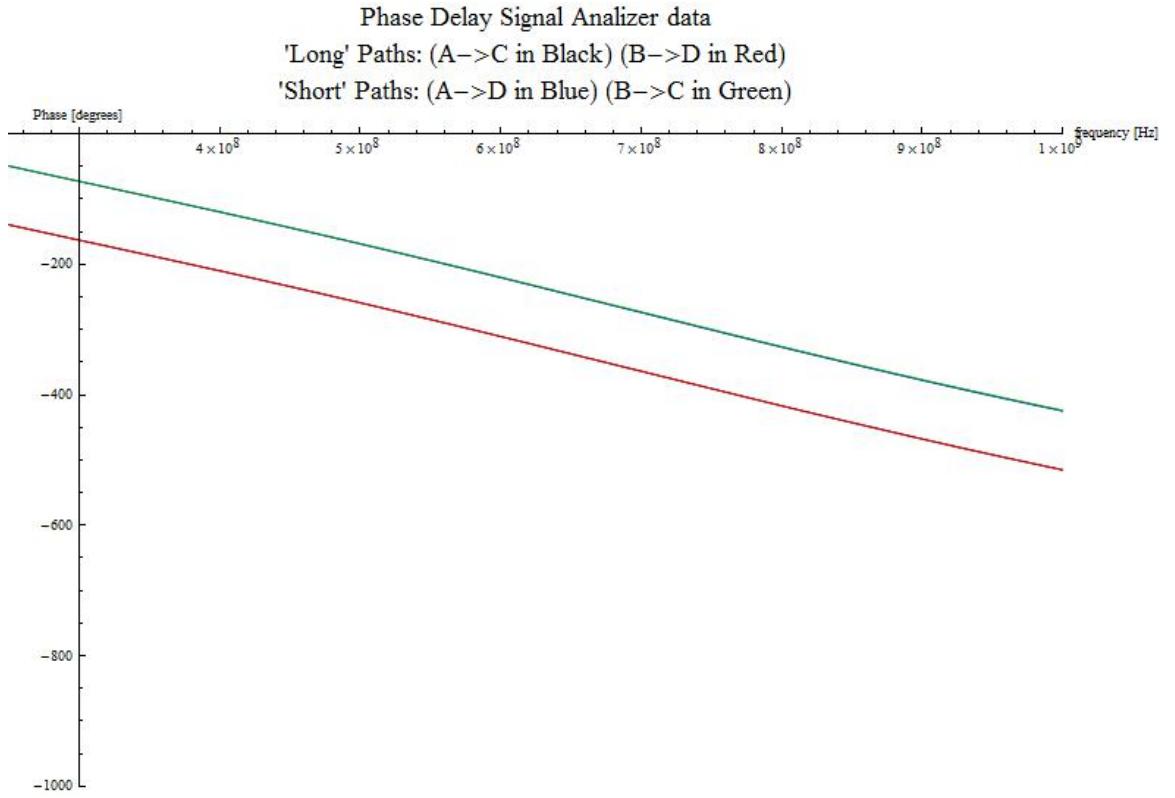


Figure 4.5: Plot shows the phase of the 4 relevant paths for a signal traveling through the hybrid. A→C is the black line. A→D is the blue line. B→C is the green line. B→D is the red line. The black and red lines lay directly on top of each other. They are the short paths through the hybrid. The blue and green lines also lay directly on top of each other. They are the long paths through the hybrid.

Examples of the hybrids effects on a pulse, as collected by the methods outlined in

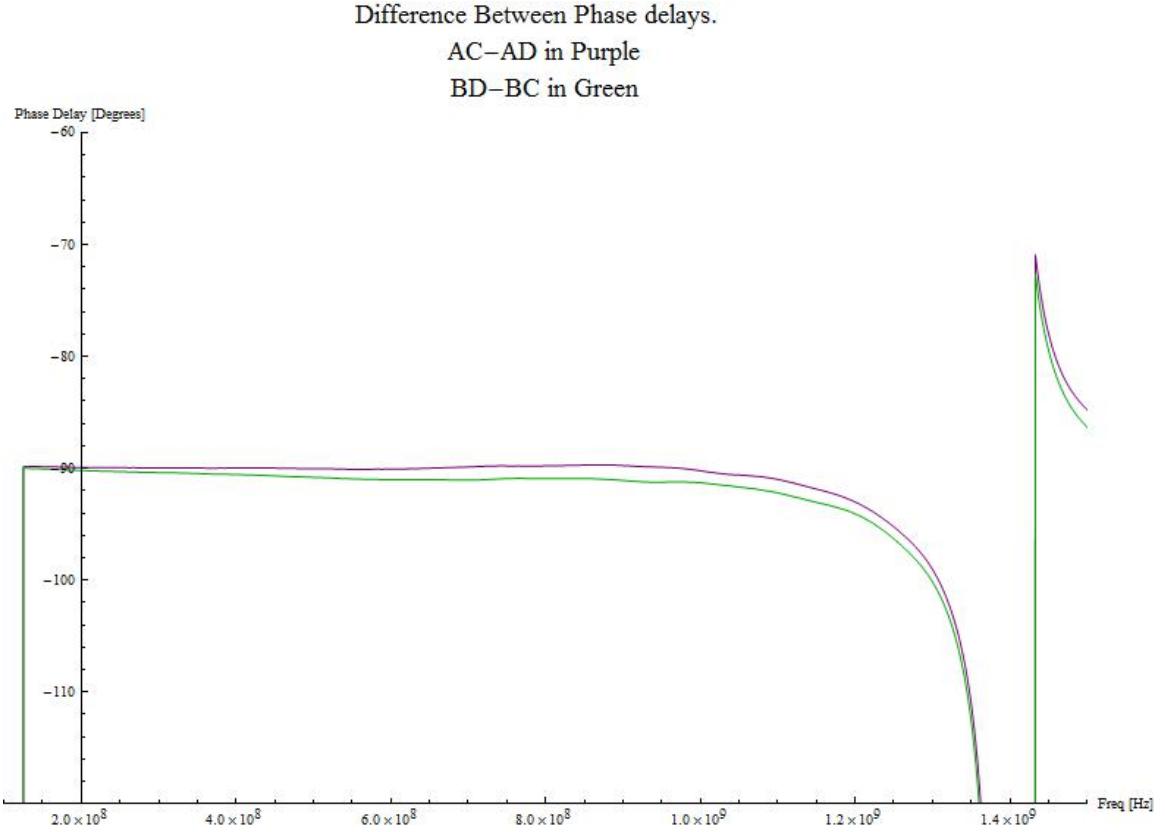


Figure 4.6: Plot shows the difference between the phase of the long paths through the hybrid and the short paths through the hybrid. The purple line is the difference of the phase of A→C and the phase of A→D. The green line is the difference of the phase of B→D and B→C. Both lines are approximately -90 degrees, which is what is expected of the hybrid.

Subsection 4.3.1 are shown in Figure 4.7. The top left shows the pulse entering the hybrid, after passing through filters. The top right and bottom left show the pulses that exit the hybrid. It is expected that these two pulses should look the same. The same signal entering both ports 'A' and 'B' would be equivalent to ANITA seeing exactly equal H-pol and V-pol signals. If both the H and V-pol signals are equal, then the LCP and RCP signals should be equal to each other as well. The bottom right shows the 3 pulses overlaid on top of each other. There is a delay in the peak of both outgoing pulses when compared to the ingoing pulse, due to propagation through the hybrid. Both outgoing pulses are in line with each other. This pulse data will be used in Section 4.5 to compare hybrid simulations against.

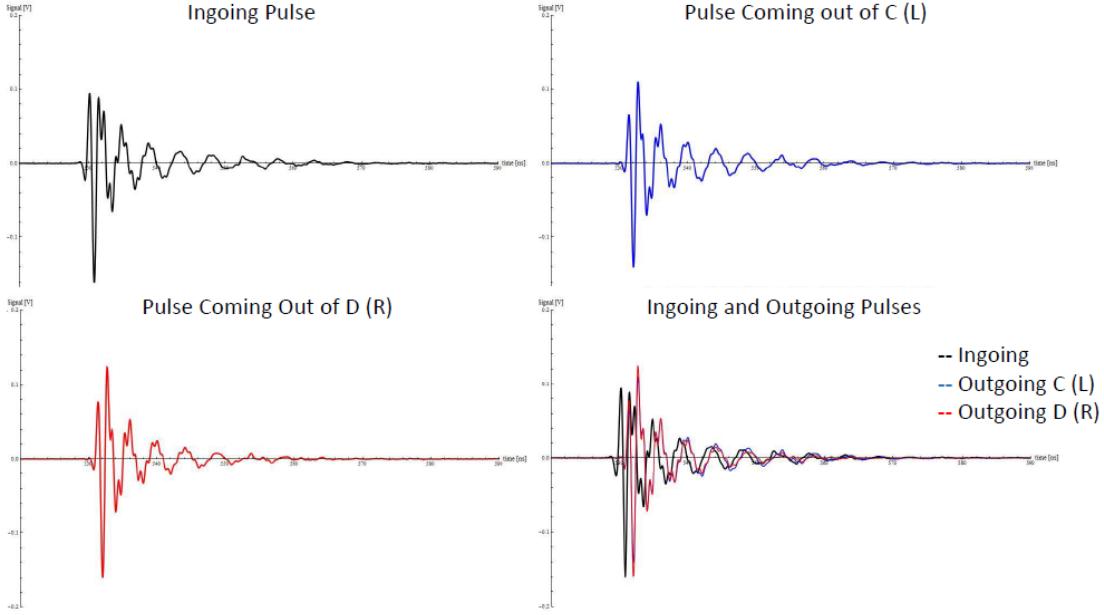


Figure 4.7: The top left shows an askarian like pulse, as seen by the hybrid, directly before entering the hybrid. The top right and bottom left plots show the pulses that exit the hybrids created when both inputs are the pulse shown in the plot on the top left. The bottom right shows the 3 pulses overlaid on top of each other.

4.4 Transforming data into a means of Data driven Simulation

Using the Network analyzer data discussed in the previous section should contain all of the information necessary to accurately transform signals entering the hybrid into signals exiting the hybrid. The first step is to take the phase frequency domain data, and the gain frequency domain data for each frequency and combine them into a single complex number. This can be done with equation 4.6.

$$S(f_i) = 10^{\frac{Gain(f_i)}{20}} \left(\cos \frac{\pi}{180} PhaseDelay(f_i) + i \sin \frac{\pi}{180} PhaseDelay(f_i) \right) \quad (4.6)$$

This needs to be done for each of the four different hybrid paths, so that sets of data for $S_{A \rightarrow C}(f_i)$, $S_{A \rightarrow D}(f_i)$, $S_{B \rightarrow C}(f_i)$, and $S_{B \rightarrow D}(f_i)$ are generated. These are effectively transfer functions for each of the possible paths though the hybrid.

The next step is to Fourier transform the time domain data you want to simulate passing through the hybrid. ANITA uses the hybrid to convert incoming H and V-pol signals into

outgoing LCP and RCP signals, so let us refer to our incoming data as $V(t_i)$ and $H(t_i)$. For this example the $H(t_i)$ data will be entering the hybrid at port A and the $V(t_i)$ data will be entering the hybrid at port B. The Fourier transformed data is then $\hat{V}(f_i)$ and $\hat{H}(f_i)$.

At this stage care must be taken to assure that the $\hat{V}(f_i)$ and $\hat{H}(f_i)$ data and the $S_{A \rightarrow C}(f_i)$, $S_{A \rightarrow D}(f_i)$, $S_{B \rightarrow C}(f_i)$, and $S_{B \rightarrow D}(f_i)$ data have the same frequencies. In all likelihood they do not. If they do not then the $S(f_i)$ data must be interpolated so it will match the $\hat{V}(f_i)$ and $\hat{H}(f_i)$ data's frequencies.

Now multiply the Fourier transformed data by the transfer functions for each of the possible paths for each of the frequency data positions. This will result in the components of the outgoing signal that originate from each of the inputs. The two paths that combine to form an output signal can not be added together in the frequency domain however. The data must be inverse Fourier transformed back into the time domain, and then added together.

$$\begin{aligned} L(t_i) &= \hat{\mathcal{F}}(\hat{H}(f_i) * S_{A \rightarrow C}(f_i)) + \hat{\mathcal{F}}(\hat{V}(f_i) * S_{B \rightarrow C}(f_i)) \\ R(t_i) &= \hat{\mathcal{F}}(\hat{H}(f_i) * S_{A \rightarrow D}(f_i)) + \hat{\mathcal{F}}(\hat{V}(f_i) * S_{B \rightarrow D}(f_i)) \end{aligned} \quad (4.7)$$

Where the Fourier and inverse Fourier transforms are represented by:

$$X(\hat{f}_i) = \mathcal{F}(X(t_i))X(t_i) = \hat{\mathcal{F}}(X(\hat{f}_i)) \quad (4.8)$$

If the ports that the V-pol and H-pol signals are entering from are switched then so too are the ports that the RCP and LCP are exiting from. In our example with H-pol singals entering through port A, and V-pol signals entering through port B, LCP signals will be exiting from port C, and RCP signals will be exiting from port D. This can be seen by comparing Equations 4.7, 4.3 and 4.4 with Figure 4.5. The long paths through the hybrid are delayed 90 degrees more than the short paths. This means that multiplying by the $S(f_i)$ for a long path is like multiplying those paths by i . Thus reproducing the effects expected from Equations 4.4 and 4.3.

4.5 Ideal simulation verses Data driven simulation

Both of the methods of simulating the hybrid outlined in the previous sections reproduce the observed effects the hybrid has on Pulse data incredibly well. Figure 4.8 shows the pulse data measured to be coming out of the hybrid compared to the ideal simulation method outlined in Section 4.2. The Ideal method slightly over estimates the peaks, but does much better than necessary for ANITA data. Real and simulated ANITA data has so much noise this amount of precision is unnessesary. The simulation peaks in Figure 4.8 were lined up artificially with the measurements by adding in a roughly 2.4 ns delay to the simulated data. Note however that the pulses coming out of C and D, though not shown here, are aligned with eachother.

Figure 4.9 shows a comparison between the data driven method of hybrid simulation and actual measured pulse data. The data driven method also matches measurements to well within ANITA's margins of error, given the amount of noise in ANITA data. It does seem to match slightly worse on some of the mid amplitude peaks than the ideal method. Figure 4.10 shows both simulation methods as well as the measured pulse data.

One might assume that the Data driven method would be superior because it takes into account any non ideal behavior of the hybrids. In practice however, while both methods match the hybrid's outputs well enough, the Ideal simulation method does a slightly better job. One advantage the Data Driven simulation method has is that it correctly reproduces the time delay caused by the hybrid in the signal, whereas the Ideal simulation method does not. The Ideal method however is more than two times as computationally efficient, and is already implemented into both AnitaTools and IceMC. The Data Driven method requires one additional transform, the storage of four arrays of complex numbers, and additional interpolation.

The hybrids preform nearly ideally within their operating range. Their behavior is now well understood. If simulation of their effects on data is needed (for example, ANITA-II and ANITA-III did not have the hybrids on board to look at LCP and RCP signals) then I would recommend using the methods built into AnitaTools and IceMC or implementing

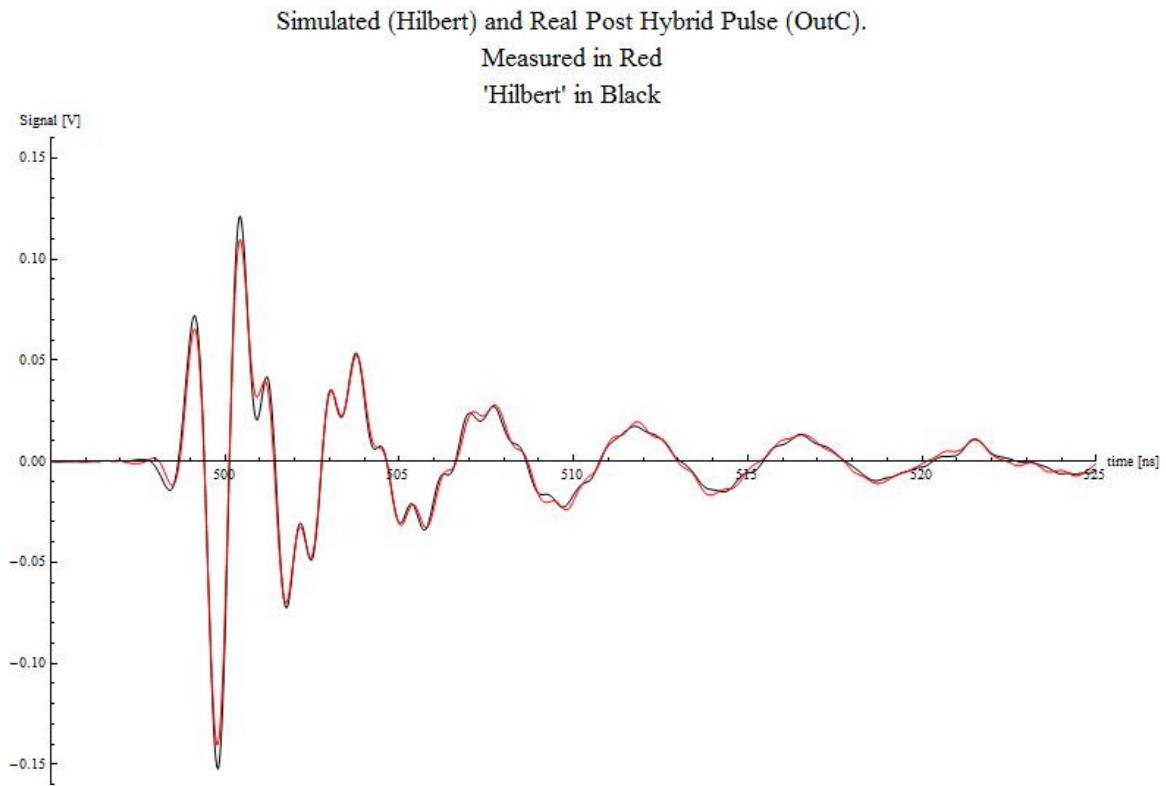


Figure 4.8: Figure shows the simulated Hilbert method (in black) compared to data of the input pulse processed by the hybrid (in red) for a pulse coming out of C. Peaks were aligned by hand by adding a few ns offset to the black line.

the Ideal Hilbert method outlined in Section 4.2.

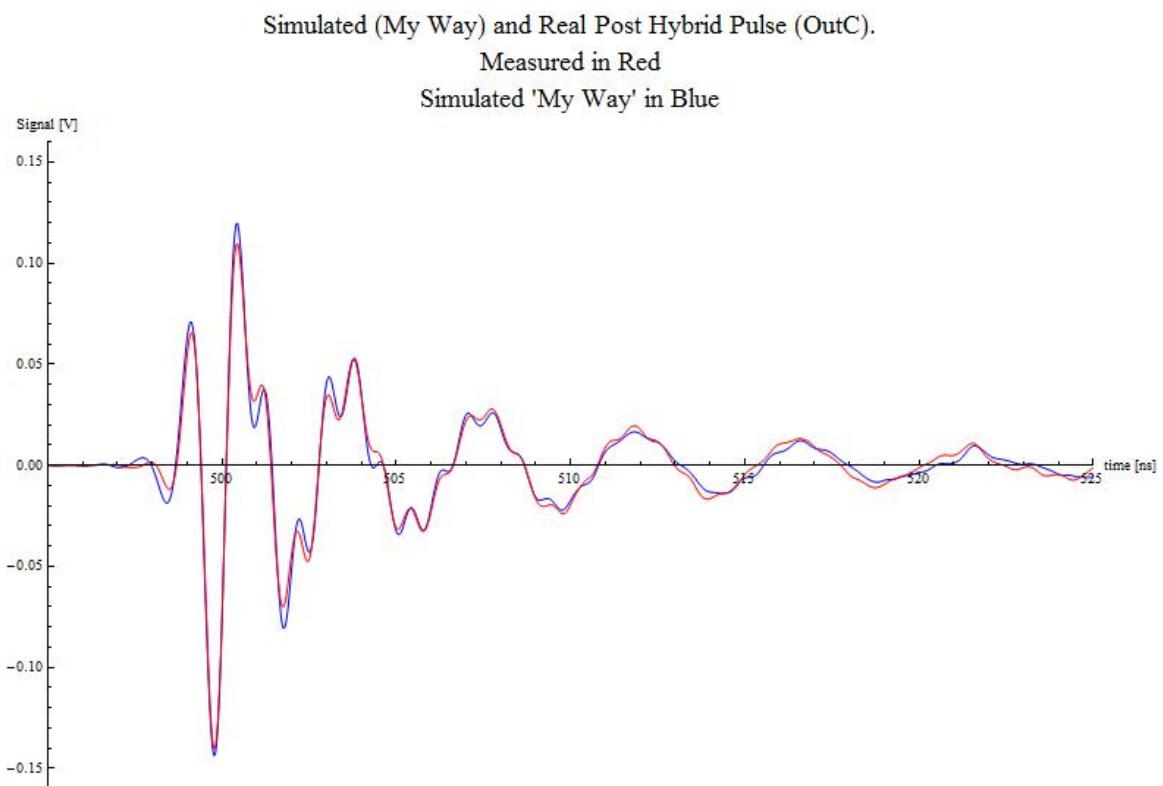


Figure 4.9: Figure shows the simulated datadriven method (in blue) compared to data of the input pulse processed by the hybrid (in red) for a pulse coming out of C.

Simulated and Real Post Hybrid Pulse (OutC).

Measured in Red

'Hilbert' in Black

'My Way' in Blue

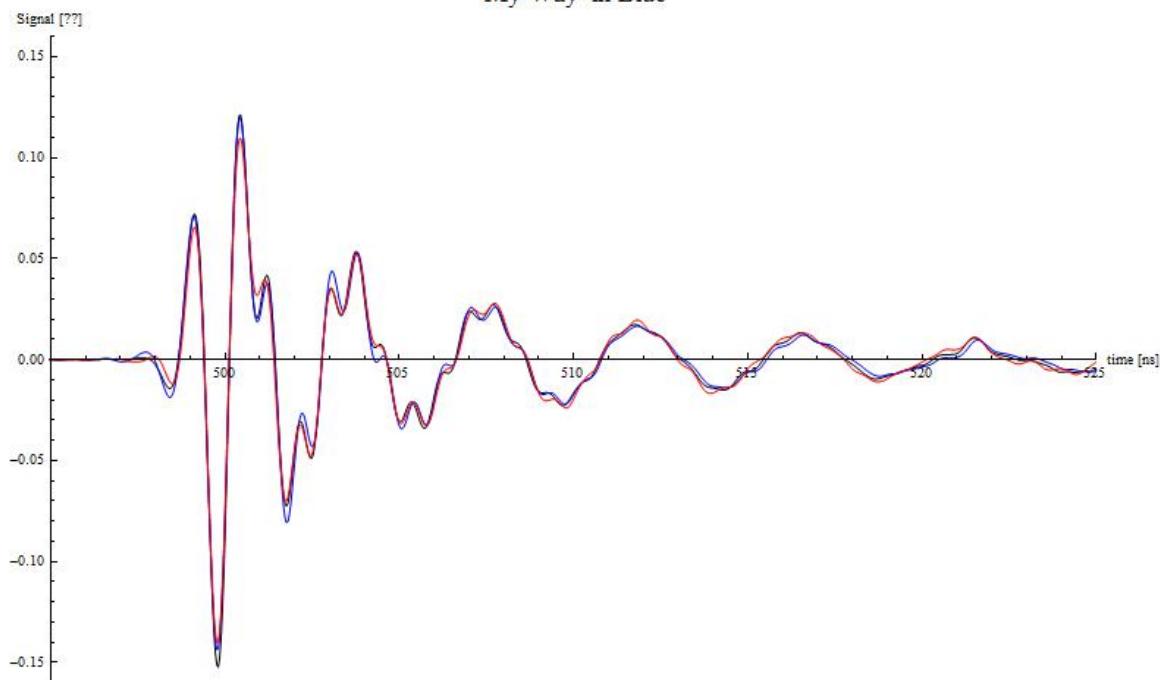


Figure 4.10: Figure shows the simulated Hilbert method (in black) and the simulated datadriven method (in blue) compared to data of the input pulse processed by the hybrid (in red) for a pulse coming out of C.

Chapter 5

THE ANITA-III BINNED ANALYSIS

This chapter is a detailed description of the ANITA-III binned analysis. Analysis of the full data set is still ongoing. Initial results from the 10% data set are discussed, as well as results from analysis of the low sensitivity sideband Healpix bins, and the full 90% H-pol data.

5.1 Past Analyses

This chapter is about a neutrino search using a binned analysis of the ANITA-III data set. However, before this analysis there have been several ANITA flights and different analysis methods. Before we talk about the ANITA-III binned analysis neutrino search, this section will briefly discuss the past neutrino searches done on ANITA data sets that led to this analysis.

5.1.1 Clustering Analysis

The primary analysis preformed on the ANITA-I and ANITA-II data was a clustering analysis. The clustering analysis relies on the assumption that we do not expect many neutrino signals to be coming from the same area of the ice, while we do expect many anthropogenic signals from the same area of the ice. When multiple events are detected from the same area of the ice, it is assumed to be a base or hotspot for human activity, and those events are cut. It is Abby Vieregg's ANITA-II clustering analysis that led to the ANITA Collaboration's published neutrino limits (seen in figure 1.3) [38].

5.1.2 The ANITA-II Binned Analysis

Following Abby Vieregg's ANITA-II clustering analysis, Brian Dailey worked to develop an alternative to the clustering method. Instead of removing any events observed near other passing events, this new binned method relies on nearby events to estimate the background ANITA observes in that area of the ice. Healpix was used to split Antarctica into roughly 50 equal area bins, each of which had its own background estimate calculated. Using those backgrounds, a linear discriminant cut which combined the event's SNR and correlation peak was optimized for each bin to give the strongest limit on neutrino models possible. After setting the analysis cuts on the 10% data sample, and analyzing the 90% data sample, approximately 20 V-pol events were observed on an expected background of 2.6 [46]. A neutrino limit was never set, however, the events that Brian saw passing in his 90% data set have informed future binned analyses on how to improve moving forward.

5.2 The ANITA-III Binned Analysis

The ANITA-III binned analysis was started by Sam Stafford in late 2016/early 2017. It updated Brian Daily's ANITA-II Binned Analysis for the ANITA-III flight by adding additional cuts based on circular polarization information and sun reflections off the ice. It however saw several events passing even in the 10% training data set, and so it was determined that further investigation was needed before analysis could proceed to the full data set [47]. That is where my analysis comes in. I have added several systematic uncertainties to the calculation of our background estimate, as well as an additional cut on events believed to have been influenced by satellite CW noise, with the help of Oindree Banerjee. In addition, the binned analysis now has an additional layer to its optimization of the linear discriminant cut intended to incorporate the results found in all of the Healpix bins into the optimization. The orientation of the Healpix bins has, in the past, always just been the Healpix default. An optimization of Healpix orientation was added, to check for orientations that better fit the data ANITA-III collected during flight. A detailed description of the entire analysis follows.

5.2.1 Healpix

Just like in the ANITA-II binned analysis, this analysis uses Healpix to cut Antarctica into smaller sections. Healpix works by cutting a sphere into 12 base pixels of equal area. Then to get finer, smaller pixels each of those 12 base pixels is cut into 4 new smaller equal area pixels. This is repeated, yielding ever smaller sections of the sphere. The order of a given healpix map is the number of times the pixels have been cut into four smaller pixels [48]. An order 0 healpix map has 12 pixels. An order 1 healpix map has $12 * 4 = 48$ pixels. An order 3 Healpix map has $12 * 4^3 = 768$ pixels. Figure 5.1 gives a visual representation of the process of splitting larger pixels into smaller ones, as well as offering an example of the shape and organization of the pixels. For our analysis we are using an order 4 Healpix map with 3072 pixels. We commonly refer to our pixels as Healpix bins. All Healpix pixels are approximately equal area and organized into equal latitude rings.

5.2.2 Blinding Strategy and Data Sets

To avoid biasing ourselves while setting up and optimizing our analysis cuts, we look at the data in stages. The first data set we use is the 10% data set. This data set consists of one of every ten events spread throughout the duration of the flight. All of our analysis cuts are set and/or optimized on this 10% data set.

The next data set, and our first step of unbinding, is to look at the 90% data set for the Healpix bins which are cut from our analysis for having low sensitivity to neutrinos. These are bins that should have a characteristic background sample for the 90% data set. Using these few Healpix bins the analysis cuts are tweaked and tuned to better reject background as necessary.

After this, we look at the 90% H-pol data set. We do not expect neutrinos in this data set, while we do expect cosmic rays. A preliminary ANITA-III UHECR search found well over a dozen. In order to further refine our background cuts using this data set we will attempt to identify cosmic ray events using methods refined by other analyses.

The final data set is the full 90% V-pol data set. At this point our analysis cuts are set

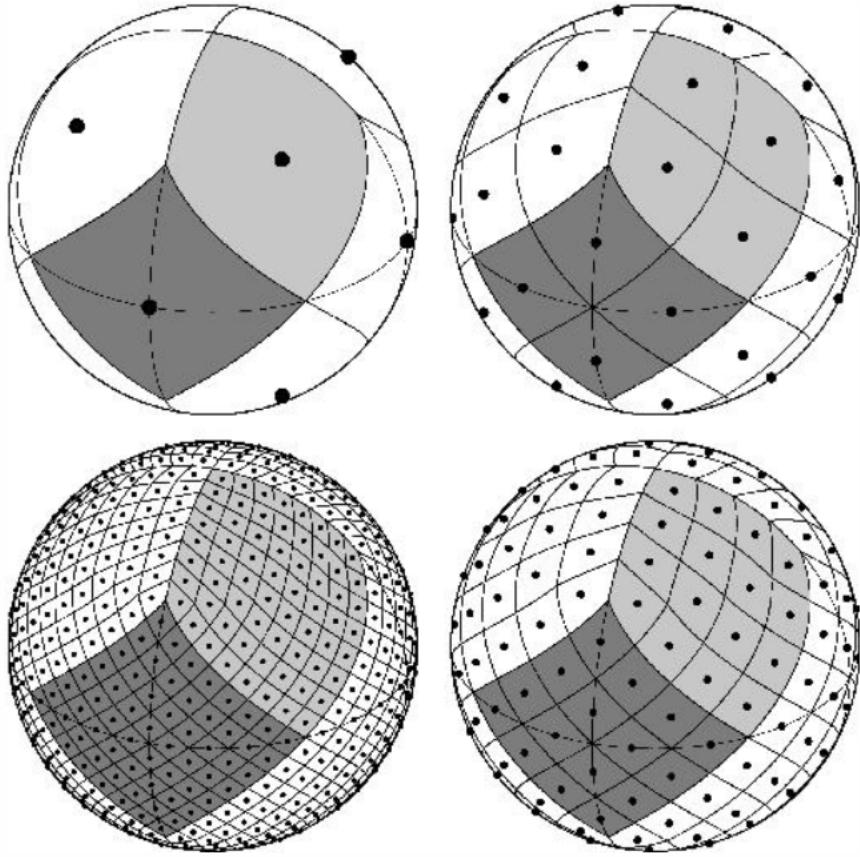


Figure 5.1: Top Left shows an order 0 Healpix map with 12 pixels. Top Right shows an order 1 Healpix map. Bottom right shows an order 2 Healpix map, while the bottom left shows an order 3 Healpix map. Dots represent the pixel centers. All pixels are of equal area, and bin centers lay on equal latitude rings [48].

and events that pass, pass. Assuming a number of events pass our cuts, which are consistent with our background estimate, a limit on neutrino models will be set.

5.2.3 The Calculation of Analysis Parameters

Before events can be analyzed and potentially cut, the raw information collected by ANTIA-III needs to be transformed into several important analysis parameters. For our analysis of ANITA-III data, these parameters are calculated for all 4 polarizations, V-pol, H-pol, LCP, and RCP.

ANITA Data

The majority of the data an ANITA flight collects is stored in ANITA’s payload on disks, until those disks are recovered. Once the disks are recovered the data is reformatted into a ROOT analysis database. The database stores data in two TTree tree objects. A header tree (headTree) stores single valued data, such as event number, trigger time and phi sector masks. An event tree (eventTree) contains uncalibrated time domain waveforms for all 96 data channels and 12 SURF clocks. Other data trees store information about the payloads orientation and position, and data from on board temperature and voltage sensors [47].

Filtering

ANITA-III saw persistent CW noise throughout its flight. One method for dealing with this is by filtering the raw waveforms received from eventTree. The geometric filter developed by Brian Dailey [46] is used to do this. The geometric filter looks at the waveform in the frequency domain and notches out the peak power. CW noise appears as a strong peak in the frequency domain, while the impulsive signals ANITA searches for are spread out in the frequency domain. This approach is some what analogous to ANITA-IV’s TUFFs, discussed in chapter 3. Though the geometric filter acts well after triggering in software, while the TUFFs filter in real time, in hardware, before triggering. The area notched out is filled in to smooth the waveform and prevent artifacts when transforming the signal back into the time domain. The geometric filter also subtracts off the CW phasor from the complex frequency domain waveform to correct the phase for the removal of the CW peak. Figure 5.2 shows an example of a pre filtered waveform in black, and the waveform after filtering in red. After filtering, interferometry is performed on the waveforms.

Interferometry, Finding the Arrival Direction of Signals

Any two of ANITA’s antennas that see the same impulsive signal should have some time delay between the observed signals. It just takes the signal slightly longer to reach one of the antennas than it does to reach the other. The length of that time delay can be

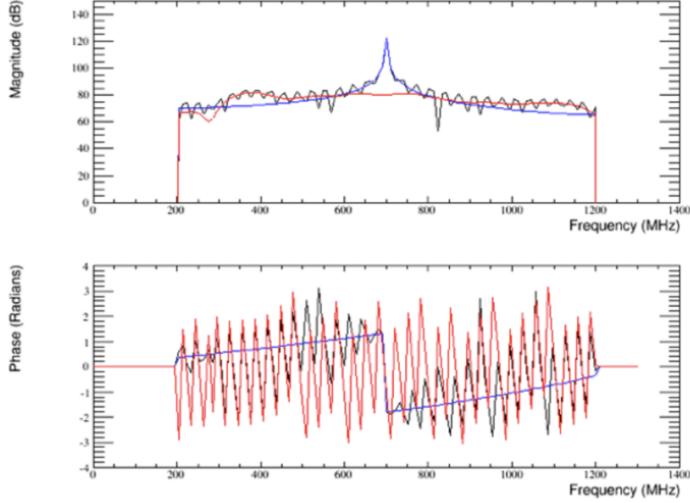


Figure 5.2: Top shows the geometric filter at work removing a CW peak from a frequency domain gain waveform. The waveform is in black, while the waveform after filtering is in red. The geometric filters fit to the CW noise is in blue. Bottom shows the same thing for the frequency domain phase [46].

used to narrow down the direction that the signal originated from, as shown in Figure 5.3. Every additional antenna that sees a signal offers additional pointing information. In a perfect world, with no noise mixed into a signal, just three antennas seeing an impulsive signal would be enough to point back to where that signal came from. In practice however, finding the time delay between antennas is not trivial. ANITA usually observes a strongly impulsive signal with between 6-12 antennas [47]. Using all of the different antenna pairs, the arrival direction is determined as best as possible by summing all of the correlation maps found from individual antenna pairs, and choosing the point with the highest correlation value as the arrival direction. This is the event pointing used in our analysis cuts. Figure 5.4 shows an example of the correlation map from a calibration pulser.

Correlation Peak

The value at the peak of this summed correlation map is the correlation peak used in our analysis cuts.

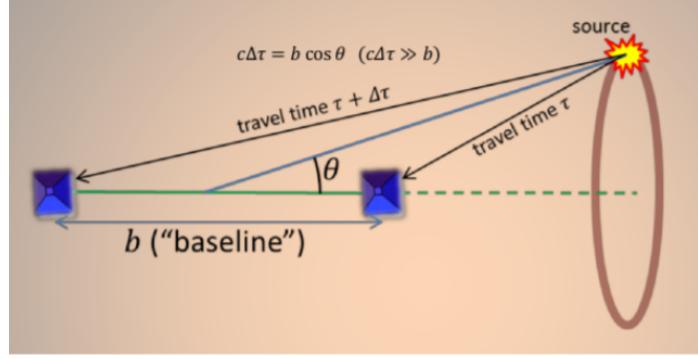


Figure 5.3: Figure displays the circle that can be drawn using two antennas with a baseline b and a time delay of $\Delta\tau$ between the arrival time of the signal [47].

SNR

Once the arrival direction is determined, each antenna pair has a time delay determined from that arrival direction. Coherently summed waveforms are created using those time delays. The signal to noise ratio of an event is calculated from the coherently summed waveforms of that signal. ANITA-III saw impulsive relics after impulsive signals, so the area before the impulsive signal was used to estimate the noise of the signal from a 10ns window. Figure 5.5 shows the 10ns window used to calculate the noise, while the signal is calculated as one half the peak to peak amplitude of the pulse.

Hilbert Peak

If the observed coherently summed waveform is $V(t)$, then the magnitude of the analytic signal of $V(t)$, $A(t)$ is a measure of its instantaneous power [47]. $A(t)$ is defined in equation 5.1

$$A(t) = V(t) + iH(t) \quad (5.1)$$

Where $H(t)$ is the Hilbert transform of $V(t)$. The Hilbert peak is then the peak in the magnitude of $A(t)$.

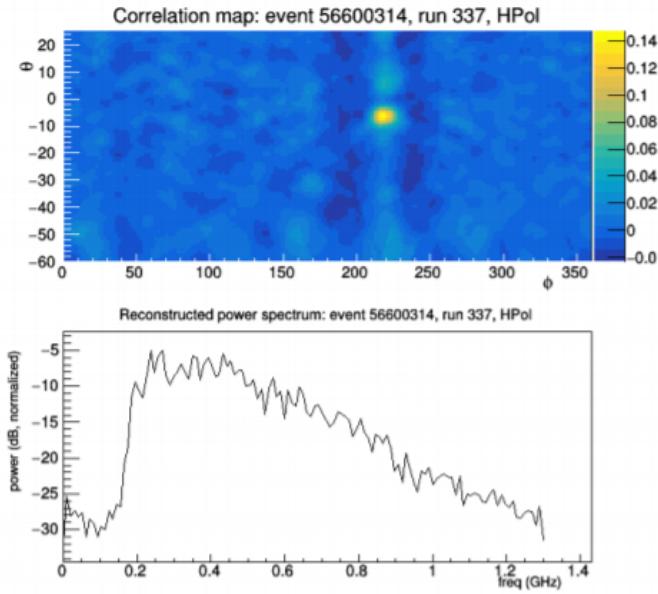


Figure 5.4: The top shows the correlation map from an WAIS calibration pulse. The bottom shows the frequency power spectrum of the coherent sum reconstruction [47].

Healpix bin and Healpix weight

One of the features of the binned analysis is that we spilt events in different Healpix bins. The Healpix bin in which an event falls is determined by tracing its pointing back to the surface of Antarctica. Using functions written in ANITATools by Ryan Nichol and BEDMAP2, a map of the surface of Antarctica, the location at which the event came out of the ice is found. This location in longitude and latitude can be used to find which Healpix bin an event falls in. Events that are very close to the boundary between Healpix bins are assigned a weight based on how much of a one standard deviation error ellipse around the event, based on uncertainty in the events pointing, is inside of the Healpix bin. Events that fall entirely inside of one Healpix bin have a weight of one.

5.2.4 Quality Cuts

Now that I have described our main analysis parameters, and how they are found, we can talk about our data cuts. Quality cuts are the very first stage of cuts. They are used to

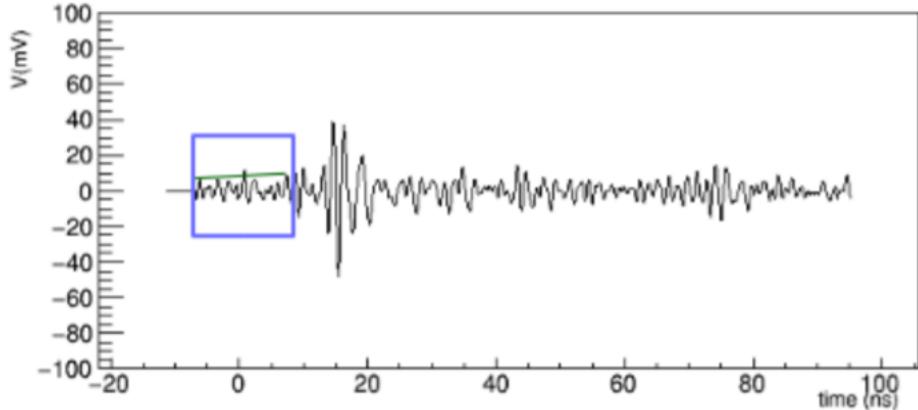


Figure 5.5: Figure shows a waveform from a calibration pulse. The blue box shows the region used to calculate the noise of the waveform, while the green line shows the average value averaged over all waveforms [47].

remove events that prevent effective interferometry. The quality cuts remove about 50% of recorded events while removing about 25% of simulated neutrinos with an SNR greater than 5.0. Table 5.1 shows the events cut in both V and H-pol for all of the quality cuts.

No Trigger Cut

The no trigger cut just requires the recorded event to have a trigger. If the event did not cause a hardware trigger, then it is cut.

Trigger Type Cut

This cut requires the trigger type to be a radio frequency (RF) trigger. Events that are not RF triggers are cut.

SURF Saturation Events

The SURF's operating range only extends up to 1.5 V. If an event's waveform goes beyond that operating range, it can become distorted[38]. Due to this, events with more than three waveforms that exceed 1.5 V are cut.

DC-Offset Events

Some events have waveforms with noticeable DC offsets. This is thought to be due to digitization problems[38]. If an event has a mean value in their waveform of greater than 100 mV in any channel, it is cut.

Short-trace Events

A complete waveform in ANITA-III has 240 data entries or samples. If an event has less than 240 samples for any reason, it is considered incomplete and cut.

Payload Blast Events

Payload blast events are events that appear to be coming from behind ANITA's antennas from the payload itself. They are both impulsive and often have a high SNR. Because of these features, payload blast events can be quite problematic to our analysis. Both this cut and the nadir noise events cut are designed to remove payload blast events. This cut removes events that L3 trigger across 6 or more phi sectors. Payload blasts are able to trigger many channels spread out dramatically in direction because they come from behind the antennas, the direction the individual antennas are facing does not matter much when the source is so near.

Nadir Noise Events

Though ANITA-III does not have a true 'Nadir' ring now that the experiment's bottom ring has been updated to include 16 antennas, like the middle and top rings of past flights, this cut is the same as it was in Abby's and Brian's ANITA-II analysis. Thus the name has not changed. This cut is designed to remove events with significantly more power in the bottom ring than in the top ring, another characteristic of an event very near the antenna. If the maximum peak voltage in the bottom ring, is less than one half the maximum peak voltage in the top ring, the event is cut.

Quality cuts:							
polarization H							
total events processed	6316237	as first cut		as ordered cut		as last cut	
cut-id	description	number	fraction	number	fraction	number	
0	trigger type	0	0.00000	0	0.00000	0	0.00000
1	L3 triggering phisectors	3217924	0.50947	3217924	0.50947	97764	0.01548
2	payload blast	24	0.00000	24	0.00000	1	0.00000
3	waveform saturation	2655	0.00042	1751	0.00028	294	0.00005
4	waveform dc offset	0	0.00000	0	0.00000	0	0.00000
5	short waveforms	77	0.00001	36	0.00001	35	0.00001
6	nadir noise	73234	0.01159	39542	0.00626	39542	0.00626
7	no trigger	3104268	0.49147	0	0.00000	0	0.00000
total events cut:		3259277					
surviving events:		3056960					
polarization V							
total events processed	6316237	as first cut		as ordered cut		as last cut	
cut-id	description	number	fraction	number	fraction	number	
0	trigger type	0	0.00000	0	0.00000	0	0.00000
1	L3 triggering phisectors	3062707	0.48489	3062707	0.48489	96930	0.01535
2	payload blast	216	0.00003	216	0.00003	11	0.00000
3	waveform saturation	2132	0.00034	1861	0.00029	512	0.00008
4	waveform dc offset	0	0.00000	0	0.00000	0	0.00000
5	short waveforms	77	0.00001	41	0.00001	41	0.00001
6	nadir noise	62292	0.00986	42901	0.00679	42901	0.00679
7	no trigger	2957773	0.46828	0	0.00000	0	0.00000
total events cut:		3107726					
surviving events:		3208511					

Table 5.1: This table shows the number of events cut by each of the quality cuts. ‘As first cut’ shows the amount cut by that quality cut if it takes place first. ‘As ordered cut’ shows the amount cut by that quality cut if it takes place in order. ‘As last cut’ shows the amount cut by that quality cut if it takes place as the last quality cut. The number columns are the number of events cut, while the fraction columns show the fraction of events cut.

5.2.5 Stage 1 Analysis Cuts

Stage 1 analysis cuts are the first stage of analysis cuts. They are direction based, removing events that reconstruct in directions we should not see neutrinos from, or directions we expect to see a large amount of noise from. Table 5.2 shows the number of events (which survived the quality cuts) cut by any of the stage 1 analysis cuts. Stage 1 analysis cuts remove about 95% of events while removing about 30% of simulated neutrinos with an SNR greater than 5.0. This allows for much faster data processing in subsequent steps of the analysis.

Solar Reflection Cut

The reflection of the sun off of the ice is a hot spot for noise events [47]. Events that point to within 5 degrees of the sun’s reflection are cut.

Localization to Continent Cut

For this ANITA-III neutrino search, we only expect to see neutrino signals from the ice. Thus events that do not reconstruct back to the continent are cut.

Elevation Angle Cut

ANITA-III’s antennas have a 6dB fall off at 22.5 degrees, and the point to 10 degrees below the horizon. This means any signals arriving from below -35 (slightly more than $22.5 + 10$) degrees should be greatly reduced in power. Many of the events we do see from those angles are misreconstructions. Events that reconstruct to angles above the continent are similarly thought to be misreconstructions [38]. Events that reconstruct to above 6.0 degrees below the horizon, or below 35.0 degrees below the horizon are cut.

Triggering Phi-sector direction Cut

Another sign of a misreconstruction is an event that reconstructs to a phi sector in which it did not cause a L3 trigger. Events that do not trigger in the phi sector they reconstruct to are cut.

Calibration Pulsar Cut

Events originating from WAIS and LDB are cut if their nanosecond timestamp is consistent with the calibration pulsars at those locations.

5.2.6 Stage 2 Analysis Cuts

Stage 2 analysis cuts remove events with characteristics in their SNR, correlation peak or Hilbert peak that are not consistent with neutrino signals. Table 5.3 shows the number of

Analysis cuts (stage 1):						
polarization H						
total events processed		3057121				
cut-id	description		as first cut	as ordered cut	as last cut	
			number fraction	number fraction	number fraction	
2	solar reflection	66370	0.02171	66370	0.02171	20014 0.00655
6	reconstruct to continent	2362321	0.77273	2362295	0.77272	1928 0.00063
11	elevation angle	2533149	0.82861	172600	0.05646	46405 0.01518
29	hardware trigger direction	2193907	0.71764	316375	0.10349	316361 0.10348
18	wais pulse	11830	0.00387	11244	0.00368	11244 0.00368
total events cut:				2928884	0.95805	
surviving events:				128237	0.04195	
polarization V						
total events processed		3208510				
cut-id	description		as first cut	as ordered cut	as last cut	
			number fraction	number fraction	number fraction	
2	solar reflection	28307	0.00882	28307	0.00882	10207 0.00318
6	reconstruct to continent	2565087	0.79946	2565062	0.79946	1563 0.00049
11	elevation angle	2780598	0.86663	216259	0.06740	53821 0.01677
29	hardware trigger direction	2314227	0.72128	280554	0.08744	280545 0.08744
18	wais pulse	631	0.00020	381	0.00012	381 0.00012
total events cut:				3090563	0.96324	
surviving events:				117947	0.03676	

Table 5.2: This table shows the number of events cut by each of the stage 1 analysis cuts. ‘As first cut’ shows the amount cut by that analysis cut if it takes place first. ‘As ordered cut’ shows the amount cut by that analysis cut if it takes place in order. ‘As last cut’ shows the amount cut by that analysis cut if it takes place as the last stage 1 analysis cut. The number columns are the number of events cut, while the fraction columns show the fraction of events cut.

events (which survived the stage 1 analysis cuts) cut by any of the stage 2 analysis cuts. The stage 2 analysis cuts remove about 73% of events, while removing about 25% of simulated neutrinos with an SNR greater than 5.0.

Ratio of Highest Peak Cut

Neutrino signals are expected to be highly impulsive, which should render as a single distinct peak in the correlation map. CW and thermal noise, however, are expected to produce multiple peaks [47]. If the ratio of the second largest to largest peak in the correlation map is less than 0.9, then the event is cut.

Correlation Peak Cut

A highly impulsive event should have a large peak value on its correlation map [47]. Events with a correlation peak value below 0.04 are cut.

Hilbert Peak Cut

Impulsive events should have the majority of their power concentrated over a small window in time. They should also have a high peak power value within that time window. The Hilbert peak is a measure of both of these. Events with a Hilbert peak value below 25 are cut.

Satellite Stripe Cut

Noise from geostationary military satellites is a significant problem for ANITA-III. To counter act this a new analysis cut was developed. It is believed that satellite CW noise boosts the correlation peak values of events and biases their reconstruction towards the satellite's location. This effect can be seen in Figure 5.6. When events are plotted on an event ϕ reconstruction direction vs ANITA's payload longitude plot. distinct stripes in the density of the map are observed. These stripes are strongest when looking at the reconstructed ϕ in the LCP, which is the same as the polarization of the Satellite's CW noise. At some longitudes certain satellites are not observable to ANITA's payload, and thus the stripes associated with those satellites drop out of Figure 5.6.

To remove events effected by this satellite contamination, events which fall within one of these satellite stripes, and are stronger in LCP than in RCP in their correlation peak, and the satellite is visible to ANITA when the event was observed are cut. The ratio of LCP/RCP that is cut on varies between stripes, but for most stripes is approximatly 1.5.

5.2.7 Final Analysis Cuts

The final analysis cuts include any analysis cuts that have been optimized for, or that need to happen after the optimization of other cuts. This includes the binned analysis's main

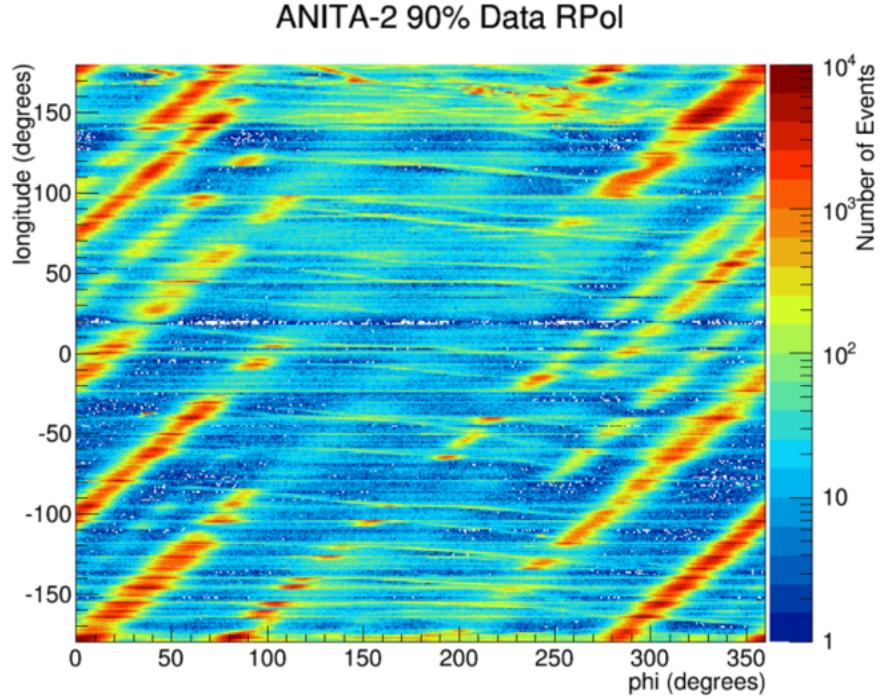


Figure 5.6: Image shows the increased density of events at specific linear combinations of payload longitude and event phi. These over densities appear as stripes in the plot. This effect is believed to be caused by satellite influence. The same effect shown here in ANITA-II's R-pol can be observed in ANITA-III's L-pol. An ANITA-II image is shown here because the stripes are more well defined in the 90% sample than in the 10%. Figure made by Oindree Banerjee.

cut, the linear discriminant cut. Table 5.4 shows the number of events (which survived the stage 2 analysis cuts) cut by any of the final analysis cuts. The final analysis cuts remove about 99.9% of surviving events, while removing only about 10% of simulated neutrinos with an SNR greater than 5.0.

Circular polarization Peak Separation Cut

The threshold for the circularly polarized peak separation cut was optimized for by Sam Stafford, as described in section 5.2.10, for the ANITA-III binned analysis. This cut removes an event if the correlation peak in LCP is more than 46 degrees from its correlation peak in RCP.

Analysis cuts (Stage 2):						
polarization H						
total events processed	124850.0				as first cut	
description	number	fraction	number	fraction	number	fraction
2nd/1st peak ratio	31487.0	0.25220	31487.0	0.25220	6143.0	0.04920
correlation peak	20441.0	0.16372	7770.0	0.06223	1839.0	0.01473
hilbert peak	19463.0	0.15589	6528.0	0.05229	3108.0	0.02489
satallite stripe cut	66279.0	0.53087	40763.0	0.32650	40763.0	0.32650
total events cut:			86548.0	0.69322		
surviving events:			38302.0	0.30678		
polarization V						
total events processed	115338.0				as first cut	
description	number	fraction	number	fraction	number	fraction
2nd/1st peak ratio	41047.0	0.35588	41047.0	0.35588	13119.0	0.11374
correlation peak	25455.0	0.22070	10046.0	0.08710	1383.0	0.01199
hilbert peak	29496.0	0.25574	8630.0	0.07482	3599.0	0.03120
satallite stripe cut	58907.0	0.51073	28829.0	0.24995	28829.0	0.24995
total events cut:			88552.0	0.76776		
surviving events:			26786.0	0.23224		

Table 5.3: This table shows the number of events cut by each of the stage 2 analysis cuts. ‘As first cut’ shows the amount cut by that analysis cut if it takes place first. ‘As ordered cut’ shows the amount cut by that analysis cut if it takes place in order. ‘As last cut’ shows the amount cut by that analysis cut if it takes place as the last stage 2 analysis cut. The number columns are the number of events cut, while the fraction columns show the fraction of events cut.

Circular polarization Peak Strength Cut

Like the circular polarization peak separation cut’s threshold, the threshold for the circular polarization peak strength cut was also optimized for by Sam Stafford, as described in section 5.2.10. An impulsive event in V-pol should have its power split up fairly evenly between LCP and RCP. The cut removes an event if either the LCP or RCP peak is below 0.015. In practice these two circular polarization cuts primarily remove thermal noise. There is no reason for thermal noise to have a correlation in peak reconstruction or strength between any of their polarizations.

Linear Discriminate Cut

The linear discriminate cut is this analysis's main cut. If the linear discriminate value for a given event is below some threshold (usually referred to as the `cutVal`) than the event is cut. The linear discriminate value is a linear combination of the events SNR and correlation peak value, and is defined as shown in equation 5.2

$$\text{Linear discriminate} = \text{SNR} - \text{slope} * \text{Correlation Peak} \quad (5.2)$$

The slope in equation 5.2 was optimized for by Sam Stafford as described in section 5.2.10 and is -6.0. The value for `cutVal` is optimized separately for each Healpix bin, as described in sections 5.2.10 and 5.2.10.

Cut on events in cut Healpix bins

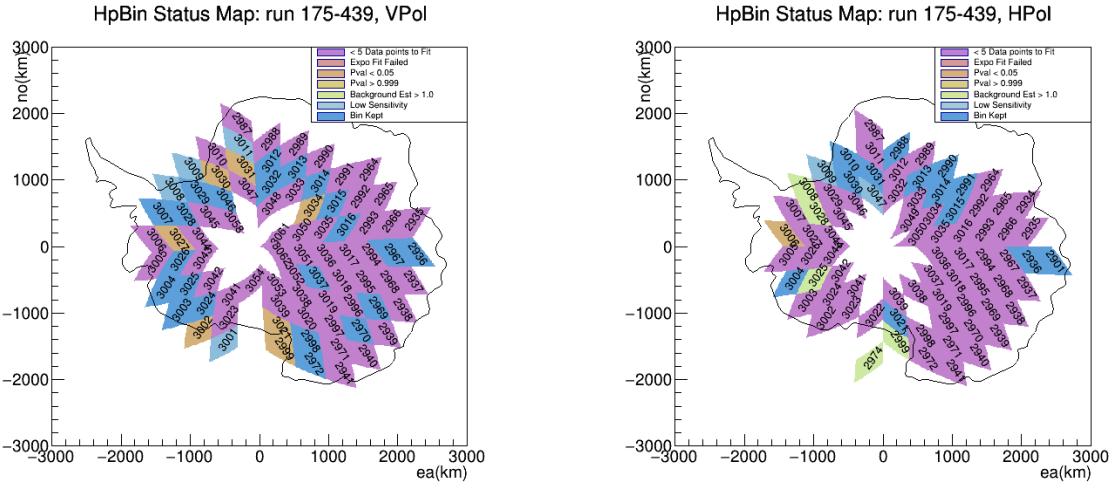
Individual Healpix bins can be cut for several different reasons. If an event is in a Healpix bin that is cut, then that event is also cut.

The first cause for a Healpix bin to be cut is if it does not have at least 5 bins (histogram bins, not Healpix bins) with data to be fit to obtain our background estimate with the process described in section 5.2.8. If we cannot get a background estimate for a Healpix bin it is not possible to use that bin to set a neutrino limit.

The second cause for a Healpix bin to be cut is if the fit found from the process described in section 5.2.8 returns a bad p-value. The process of finding a fit's p-value is described in section 5.2.8. If the p-value for a Healpix bin's fit is less than 0.05, or greater than 0.999 then we do not consider the fit good enough to accurately represent the data, and the Healpix bin is rejected.

The third cause for a Healpix bin to be cut is if the Healpix bin has a background estimate greater than 1.0. If the background estimate for an individual bin is that large, then that Healpix bin will not be able to help us set a lower limit, and is thus cut.

The forth and final cause for a Healpix bin to be cut is if the Healpix bin has very low neutrino sensitivity. We base our estimate for neutrino sensitivity for a Healpix bin



(a) Status of Healpix bins in the V-pol channel analysis.

(b) Status of Healpix bins in the H-pol channel analysis.

Figure 5.7: Blue Healpix bins are kept in the analysis. All others are cut for various reasons identified by their color. The light blue/gray Healpix bins cut for low sensitivity will later be used as a sideband before examining the full 90% data set.

off of how many simulated neutrinos passing all cuts before our linear discriminant cut are coming from that area of the ice. More simulated neutrinos means that area of the ice is more sensitive to neutrinos. The least sensitive 1% of Healpix bins (after all previous Healpix bin cuts) are removed for low sensitivity. These few Healpix bins can be used as a sideband later in the analysis.

Figure 5.7 shows a map of which Healpix bins are kept, and for what reason the other Healpix bins are cut overlaid on top of an outline of Antarctica. 22 Healpix bins are kept in the V-pol channel analysis. 12 Healpix bins are kept in the H-pol channel analysis.

Cut on events with a weight less than 0.5

Events with an event weight of less than 0.5 are cut. An event weight of less than 0.5 corresponds to the event being less than 50% likely to have come from the Healpix bin it is seen passing in. We want passing events to pass in the Healpix bin the are mostly within.

Analysis cuts (Stage 2 + Final):						
polarization H						
total events processed	128237.0	as first cut		as ordered cut		as last cut
description		number	fraction	number	fraction	number
2nd/1st peak ratio	32177.0	0.25092		32177.0	0.25092	3150.0
correlation peak	20685.0	0.16130		7863.0	0.06132	466.0
hilbert peak	19688.0	0.15353		6640.0	0.05178	935.0
satallite stripe cut	68750.0	0.53612		42667.0	0.33272	15458.0
copol peak separation	54880.0	0.42796		6570.0	0.05123	4684.0
copol peak strength	31570.0	0.24618		2048.0	0.01597	2048.0
low event bin weight				1010.2	0.00788	
healpix bin rejected				27688.8	0.21592	
linear discriminant cut				1573.2	0.01227	
total events cut:				128236.6	1.00000	
surviving events:				0.0	0.00000	
polarization V						
total events processed	117947.0	as first cut		as ordered cut		as last cut
description		number	fraction	number	fraction	number
2nd/1st peak ratio	41611.0	0.35279		41611.0	0.35279	9267.0
correlation peak	25749.0	0.21831		10170.0	0.08623	380.0
hilbert peak	29800.0	0.25266		8770.0	0.07436	1153.0
satallite stripe cut	60787.0	0.51538		30162.0	0.25573	13217.0
copol peak separation	48025.0	0.40717		5894.0	0.04997	4842.0
copol peak strength	29635.0	0.25126		2028.0	0.01719	2028.0
low event bin weight				164.6	0.00140	
healpix bin rejected				12096.0	0.10255	
linear discriminant cut				7050.3	0.05978	
total events cut:				117946.2	0.99999	
surviving events:				1.0	0.00001	

Table 5.4: This table shows the number of events cut by each of the final analysis cuts. ‘As first cut’ shows the amount cut by that analysis cut if it takes place first. ‘As ordered cut’ shows the amount cut by that analysis cut if it takes place in order. ‘As last cut’ shows the amount cut by that analysis cut if it takes place as the last final analysis cut. The number columns are the number of events cut, while the fraction columns show the fraction of events cut.

5.2.8 Background estimate

An crucial component of the binned analysis is obtaining an estimate of the expected background for each healpix bin. In order to do this, we would like to fit a function that matches the distribution of events observed passing our quality cuts and analysis cuts. This background estimate combined with our simulated data set will help us inform where to place our main cut, the linear discriminate cut. Due to its relationship with the linear discriminate cut we represent the data at this stage of the analysis as a histogram in the linear

combination of the variables plotted for the linear discriminant cut, as shown in equation 5.3.

$$LD = \text{SNR} - LD_{slope} * \text{Correlation Peak} \quad (5.3)$$

How LD_{slope} was determined will be addressed in detail in section 5.2.10. The SNR is the signal to noise ratio of the event's peak for the given polarization. The Correlation Peak is the peak correlation value for the given polarization. Figure 5.8 shows an example of the data being fit, along with an exponential fit to the data.

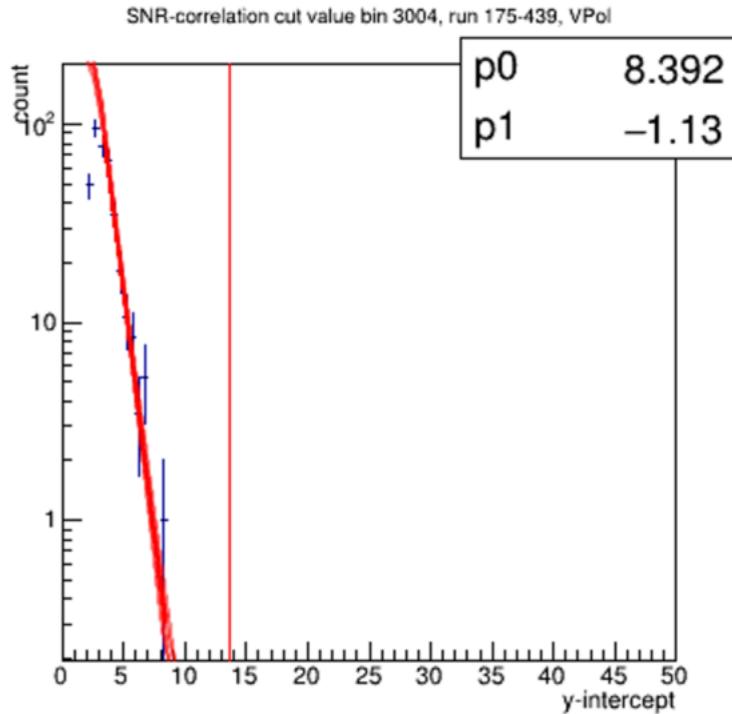


Figure 5.8: Figure shows an example of the LD histogram being fit by an exponential decay function for Healpix bin 3004 in the V-pol channel. Data points are in blue. The best fit is in red. The vertical red line is the optimized y-intercept cut threshold for the linear discriminant cut (before the overall optimization).

The functional form chosen to fit the observed distributions is an exponential decay. During the original ANITA-II Binned Analysis it was observed that the data fell off very

quickly for higher values of LD and the fall off after the peak appeared to be linear on a log plot. Both of these features match the characteristics of an exponential decay.

The data was fit from the second data point after the peak of the histogram, out to include up to 23 histogram bins with zero data. The data was fit by the functional form shown in equation 5.4 where a and b are the two fit parameters, and x is the x axis value, or the linear combination of correlation peak and SNR from equation 5.3.

$$fit(x) = e^{a+b*x} \quad (5.4)$$

The data is fit using the Minuit2 migrad minimization, where the equation being minimized is a negative log likelihood equation assuming Poisson like non integer values.

$$Likelihood = fit(x_i)y_i + y_i * \ln \frac{y_i}{fit(x_i)} \quad (5.5)$$

x_i and y_i represent the x axis, and y axis values for the i th data point as seen in Figure 5.8. The likelihood shown in equation 5.5 has been simplified using the first few terms of an expansion. Values for a , b , σ_a and σ_b are obtained from the migrad fitting algorithm.

Once the best fit has been determined, and the y-intercept for the linear discriminant cut has been chosen, the background is obtained by integrating the best fit from the linear discriminant cut to infinity. Assuming the parameter b has a negative value, which it always should for a decaying exponential, the integral can be simplified down into equation 5.6, where $CutVal$ is the y-intercept of the linear Discriminant cut.

$$\text{Background Estimate} = \frac{-1}{b} * \exp(a + b * CutVal) \quad (5.6)$$

In our ANITA-III Binned Analysis, we obtain our background estimates from a 10% data set, so to obtain a background estimate for the number of events we expect passing, for a given Healpix bin in the 90% data set, we would multiply the result of equation 5.6 by a factor of 9.

Fit Validation: P-values

In order to use our exponential fit as a background estimate, it is important to validate that our data behaves like our model does. To do this we do a p-value test. We generate many sets of random pseudo experiment data by treating the exponential fit as a probability distribution. Each pseudo experiment has the same number of events as the actual fit data does. The likelihood of the original exponential fit fitting this new pseudo data is then found.

$$\text{Fit Log Likelihood} = \sum_i f(x_i)y_i + y_i * \ln \frac{y_i}{f(x_i)} \quad (5.7)$$

Where x_i and y_i are the x and y coordinates of the i th data point, and $f(x)$ is the value returned by the fit at an arbitrary position x . The percentage of pseudo experiments that return a lower log likelihood value than our original fit is then calculated, and that is our p-value. Figure 5.9 shows an example of a histogram created from pseudo experiments' log likelihood values. The red line is the log likelihood value of that Healpix bin's actual data.

Figure 5.10 shows the distribution of p-values found for all Healpix bins with sufficient data to be fit. If our exponential fit is a good model of the data then one would expect a flat distribution of p-values, which appears to be what we observe (for Healpix bins with a p-value greater than 0.05).

5.2.9 Background Estimate Systematic Uncertainty

There is a lot of uncertainty that comes from this background estimate however. First of all, the data is not perfectly fit by an exponential decay. This means there are uncertainties in the fit parameters which can be propagated on into uncertainties in the background estimate.

Additionally because we used a 10% data set to create our background estimate, it is possible that the events in the 90% data set will not fall in exactly the same healpix bins, thus altering our expected number of events in each bin. We try to take account of this by calculating the spillover of events between bins.

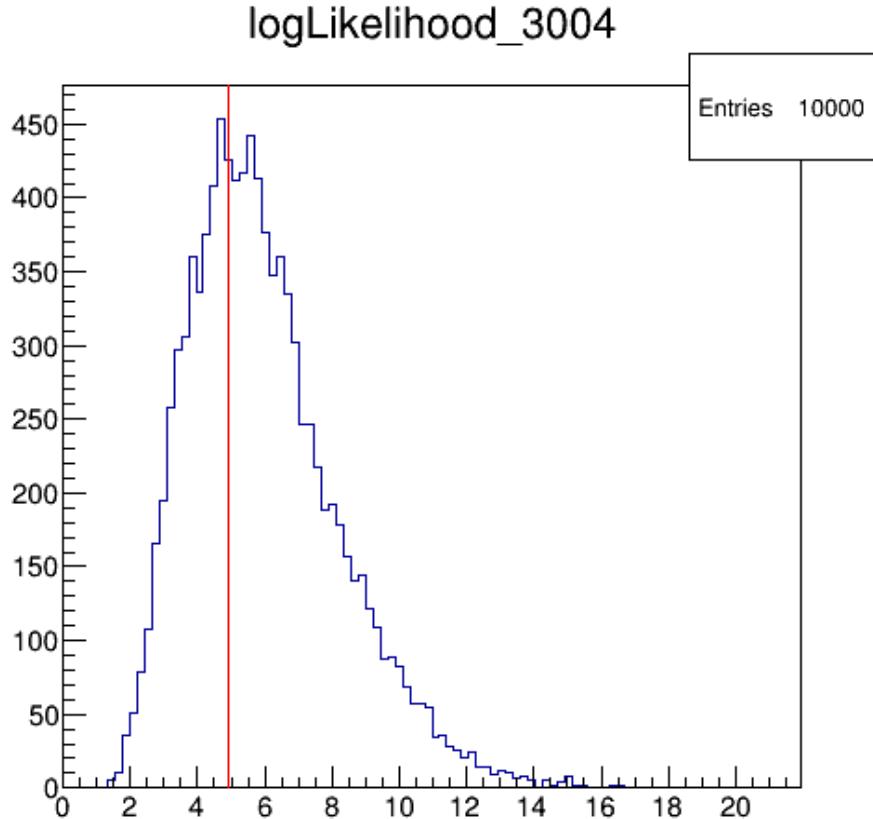


Figure 5.9: Plot shows a histogram of log likelihood values for p-value pseudo experiments. The original fits log likelihood is the red vertical line. This plot is from Healpix bin 3004 in the V-pol channel.

We also attempt to account for any bias created by our specific choice of our fit model. We chose an exponential decay, which has one of the steepest decays of any simple function. To attempt to account for any uncertainty from this choice we compare our background estimates from an exponential fit with the background estimates we would obtain for a power law fit.

Fit Parameter Uncertainty

Using Minuit2 we obtain the uncertainty in both of our fit's parameters. These two parameters however not only correlated, but it is also possible that their uncertainties are not symmetric. To take into account these complications when propagating the parameter

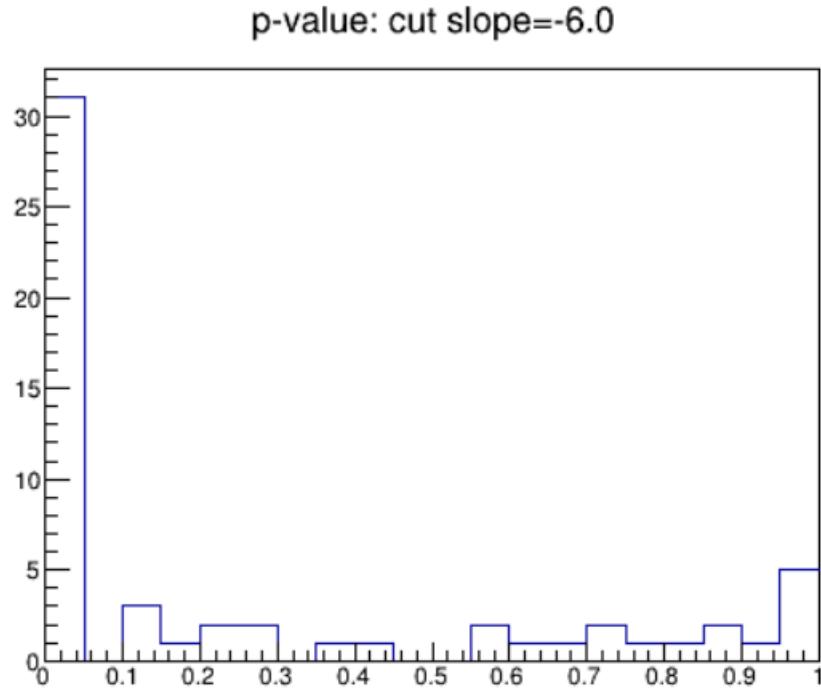


Figure 5.10: Plot shows a histogram of p-values found from different Healpix bins in the V-pol analysis channel. Healpix bins which were not fit due to insufficient data are plotted as having a 0 p-value here, so the first bin should be ignored. The approximately flat nature of this histogram implies that an exponential fit is a good model of our data.

uncertainties to the background estimate, instead of relying on the parameter uncertainties we obtain the one standard deviation error contour from Minuit2. This gives us an error ellipse for the two parameters.

Due to the non-linearity of the relationship between the fit parameters and the background estimate, instead of explicitly propagating the errors, we simply use simulation to find the uncertainty in the background. In order to do this we need to create a two dimensional probability distribution function to draw random parameter values from, taking into account their correlation. Using the error ellipse from Minuit2 we can do this following the procedure outlined in Appendix C.

To calculate the uncertainty due to the parameters we draw N pairs of fit parameters from the two dimensional PDF, and calculate the random background estimate those parameters give. This results in a distribution of N different random background estimates.

To find the one standard deviation error bars the distribution is treated as a CDF. The lower background limit is the value where 15.8% of the background distribution is below that value. Similarly, the upper background limit is the value where 84% of the background distribution is below that value. To turn these limits into background error bars the actual background estimate must be subtracted from them. These error bars are reported but not actually used. The random backgrounds created from this method however are then further modified by the other systematic uncertainties.

Uncertainty due to Spillover of events between Healpix Bins

If every event in the 10% data set represents nine events in the 90% data set, then we expect those other nine events to be distributed around the 10% data set event. This means that some of those nine events could point back to different Healpix bins. In order to account for that possibility in our background estimate we calculate the spillover between bins.

The first step in calculating the spillover is to, for each event (that has passed our quality and analysis cuts) create a grid in θ and ϕ around the events actual θ and ϕ . ANITA-III has an uncertainty in ϕ of approximately 0.5 degrees, and an uncertainty in θ of approximately 0.25 degrees. Uncertainties in θ and ϕ are assumed to be Gaussian. In order to densely sample those uncertainties, a grid spacing of one tenth of those uncertainties is chosen. Grid points extend out to a four σ level. This results in an oval of grid points around the actual event pointing. Each grid point has a probability associated with it. The probability, $P(G, P)$, is the chance that an event at the grid point G could mistakenly be reconstructed to point back to the direction ANITA found, P.

$$P(G, P) = \frac{1}{2\pi * \sigma_\phi * \sigma_\theta} \exp\left(-\frac{(\theta_G - \theta_P)^2}{\sigma_\theta^2} - \frac{(\phi_G - \phi_P)^2}{\sigma_\phi^2}\right) \quad (5.8)$$

Where θ_G and ϕ_G are the θ and ϕ at the grid point G. θ_P and ϕ_P are the θ and ϕ at the event reconstruction P. σ_θ and σ_ϕ are the uncertainty in θ and ϕ .

Each grid point is then ray traced back to the Antarctic continent to find where a signal coming from that direction would have left the ice. This results in a latitude and longitude

for the grid point. Which Healpix bin the grid point falls in is obtained using built in Healpix functions. Each grid point is multiplied by a factor to take into account the area on the earth that that grid point is representing. This is found by converting the position for each grid point into easting and northing, then finding the distance from any individual grid point to the grid points around it. For example, if the grid index for ϕ is i and the index for θ is j , then the factor the grid point G_{ij} would be multiplied by would be:

$$A_{ij} = \frac{\sqrt{(no_{ij} - no_{i-1j})^2 + (ea_{ij} - ea_{i-1j})^2} + \sqrt{(no_{ij} - no_{i+1j})^2 + (ea_{ij} - ea_{i+1j})^2}}{2} \times \frac{\sqrt{(no_{ij} - no_{ij-1})^2 + (ea_{ij} - ea_{ij-1})^2} + \sqrt{(no_{ij} - no_{ij+1})^2 + (ea_{ij} - ea_{ij+1})^2}}{2} \quad (5.9)$$

Where no_{ij} is the northing of grid point G_{ij} , and ea_{ij} is the easting of grid point G_{ij} . If a given grid point does not have a neighboring grid point in one of the directions around it, the above equation is modified to take that into account and base that dimension of A_{ij} off of only one neighboring grid point instead of both its neighbors. Once each grid point for a given event has been multiplied by the area it represents, all of the probabilities are normalized so their total is equal to one. The final probability for any given point grid point G_{ij} is then:

$$\text{Area Normalized } P_{ij} = \frac{P_{ij} * A_{ij}}{\sum_i \sum_j P_{ij} * A_{ij}} \quad (5.10)$$

Where P_{ij} is calculated from equation 5.8 and A_{ij} is calculated from equation 5.9.

The Probabilities from a given event's grid points, that fall within the same Healpix bin, are then summed together and stored in an indexed map. This is done for all events that have passed all cuts before the Linear Discriminate cut. Once all of the events' spillover contributions have been added into the indexed map a final normalization is done. This normalization requires the total probability contribution from any bin to its neighbors and itself to add up to one. After this final normalization, to find the background uncertainty contribution of Bin X to Bin Y, you would multiply Bin X's background estimate by Bin

Healpix bin 2967		Healpix bin 2968		Healpix bin 2969	
Receiving Bin	Cont	Receiving Bin	Cont	Receiving Bin	Cont
2967	0.9904	2968	0.9868	2969	0.9285
2966	0.0037	2995	0.009	2995	0.0583
2994	0.0037	2994	0.0039	2970	0.0075
2936	0.0013	2969	0.0002	2996	0.004
2937	0.0009			2968	0.0013
				2939	0.0005

Table 5.5: This table displays an example of the spillover contributions for Healpix bins 2967, 2968 and 2969. The column labeled Cont is the percent contribution to the Receiving Bin, of the giving bin's background.

X's probability contribution to Bin Y. Table 5.5 shows an example of three Healpix Bin's contributions to themselves and their neighbors.

A positive error bar on a Healpix bin's background estimate is acquired by summing all of the neighboring Healpix bin's background contributions to a given Healpix bin. This value is referred to as the inflow. A negative error bar on a Healpix bin's background estimate is acquired by subtracting a Healpix bin's background contribution to itself from its background estimate. Thus for Healpix bin i , its inflow and outflow are described by equations 5.11 and 5.12

$$Inflow_i = \sum_j Background_j * Contribution_j \text{ to } i \quad (5.11)$$

$$Outflow_i = Background_i (1 - Contribution_i \text{ to } i) \quad (5.12)$$

Fit Choice Uncertainty

For this analysis, the fit model chosen was an exponential decay. It was not however the only option. In order to attempt to take into account any additional systematic uncertainty introduced by the choice of an exponential our LD data is fit by both an exponential decay and a power law fit. A power law can fit our data nearly, or possibly just as well as an exponential but results in a systematically higher background estimate. Figure 5.11 displays

both a power law fit and an exponential decay fit to the same Healpix bin's data.

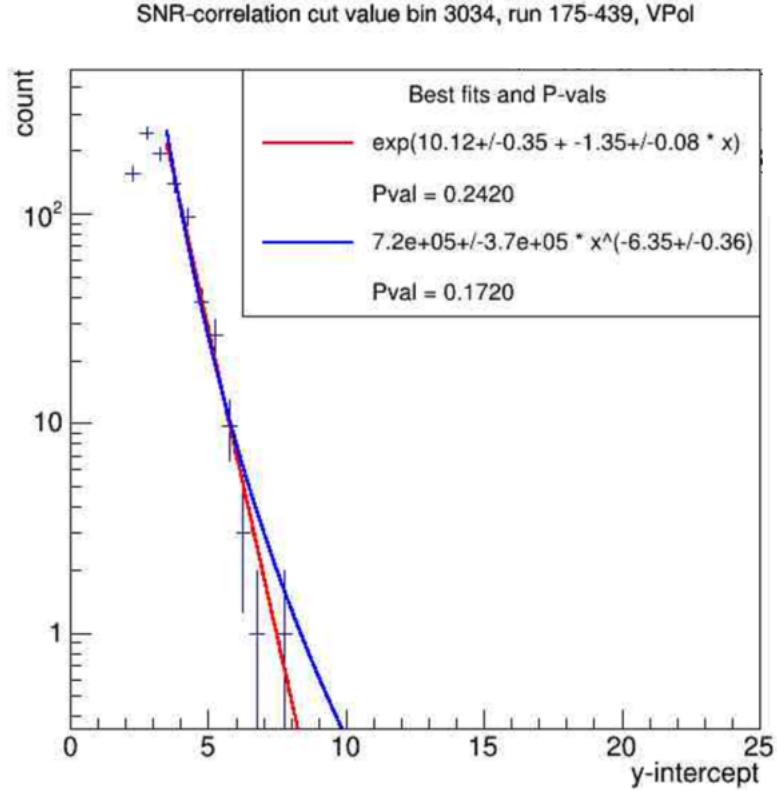


Figure 5.11: Figure displays both an exponential fit and a power law fit on the same axis, fit to the same set of data. The data is in black. The exponential fit is in red. The power law fit is in blue. Both the power law fit, and the exponential fit fit the data well, while also diverging significantly at high values of the linear combination variable.

A background estimate for a given y-intercept of the linear discriminant cut is calculated from the exponential fit as shown in equation 5.6, by integrating the exponential fit from the y-intercept of the linear discriminant cut, CutVal, to infinity. The same thing is done for the power law. The equation for a power law is shown in equation 5.13. The equation for the background estimate from a power law fit is shown in equation 5.14.

$$Pow(x) = a * x^b \quad (5.13)$$

$$\text{Power Law Background Estimate} = \frac{-a}{b+1} CutVal^{b+1} \quad (5.14)$$

Note that equation 5.14 is only valid when $b < 0$, which should always be the case for our data. Once a background estimate is obtained from both fits, the difference between those two fits is taken as a one sided error bar on our background estimate. This is nearly always a positive error bar.

Systematic Uncertainty Values

These three sources of uncertainty are incorporated into our overall background estimate by generating a background distribution including randomized contributions from each of the systematic errors. Thus our random background estimate can be represented as follows.

$$Background_r = background_{est} + param_r + inflow_r - outflow_r + choice_r \quad (5.15)$$

Where $background_{est}$ is the background estimate obtained from integrating the exponential fit above the CutVal. $param_r$ is the randomized error from the fit parameter uncertainty. $inflow_r$ and $outflow_r$ are the randomized error from the inflow and outflow, respectively. $choice_r$ is the randomized error from the choice of fit model uncertainty. $param_r$ is obtained by subtracting the random background calculated with randomized fit exponential fit parameters from the $background_{est}$.

$$Param_r = background_{est} - background(a_r, b_r) \quad (5.16)$$

$inflow_r$ and $outflow_r$ are both obtained by using the build in ROOT function TRANSDOM::Gaus(Mean,Sigma). Both $inflow_r$ and $outflow_r$ must be only positive however, because they are representing one sided error bars. This is handled by taking the absolute

value of the number returned by ROOT's Gaus function.

$$\begin{aligned} inflow_r &= \text{Abs}(\text{Gaus}(0, inflow)) \\ outflow_r &= \text{Abs}(\text{Gaus}(0, outflow)) \end{aligned} \quad (5.17)$$

Where $inflow$ and $outflow$ are obtained as described in subsection 5.2.9. It is possible (though for all relevant Healpix bins very improbable) for $outflow_r$ to be larger than $background_{est}$. If this happens then $outflow_r$ is set to be equal to $background_{est}$.

$choice_r$ is obtained by pulling a random value from a log normal distribution. The mean of the log normal distribution is set to be the average of the background estimates obtained from the exponential fit and the power law fit. The variance of the log normal distribution is set to be one half the difference between the background estimates from the exponential fit and the power law fit. The power law distribution is described by equation 5.18 [49].

$$\text{logNormal}(x, \mu, \sigma) = \frac{1}{x * \sigma * \sqrt{2\pi}} * \exp\left\{-\frac{(\ln x - \mu)^2}{2 * \sigma^2}\right\} \quad (5.18)$$

Where

$$\begin{aligned} \sigma &= \ln \frac{\text{mean}^2}{\sqrt{\text{variance} + \text{mean}^2}} \\ \mu &= \sqrt{\ln \frac{\text{variance}}{\text{mean}^2} + 1} \end{aligned} \quad (5.19)$$

A log normal is used here to attempt to more accurately model the positive definite nature of this source of error. This distribution naturally goes to zero at zero and has an extended tail which is typical of one sided sources of error. Using equation 5.18 in combination with the dartboard method (also known as Monte Carlo Integration) of obtaining a random number from a distribution, $choice_r$ is obtained.

Following this procedure thousands of random backgrounds are created in order to form a distribution for each Healpix bin. An example of such a background distribution is shown in figure 5.12. These distributions are used during our optimization of the linear discriminant cut. The mean of this distribution is taken as our expected background estimate, including systematics. The distribution is transformed into a CDF and the point where 15.8% of the

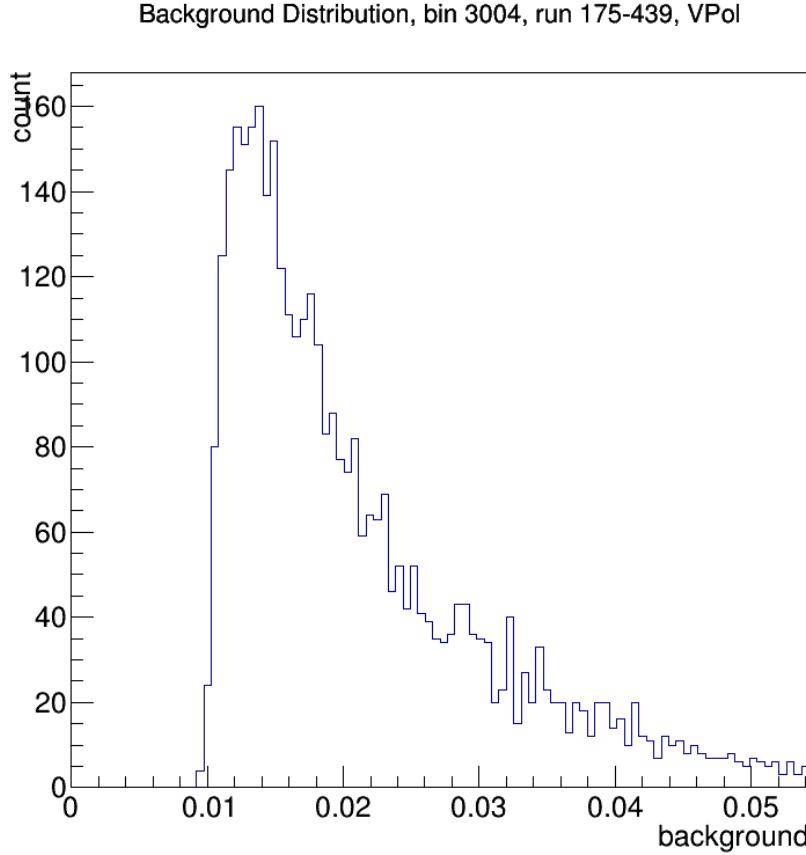


Figure 5.12: Figure shows a typical background distribution generated from including systematic uncertainties. Plot shown is for Healpix bin 3004.

background distribution is below that point is taken as the lower one standard deviation bound. Similarly, the upper bound is the point when 84% of the background distribution is below the data point. These values are reported for all Healpix bins used in our analysis in Tables 5.6 (for V-pol) and 5.7 (for H-pol)

5.2.10 Optimization

We optimize several aspects of our analysis in order to set the best limit on neutrino models possible. Before we can talk about our optimizations however, our figure of merit must be addressed. Our figure of merit is a scale factor ‘C’.

Bin Number	optimized LD cut	Sim events passing	Expo fit Background	Background Estimate	Upper Error limit	Lower Error limit
3015	10.9	0.264	0.032	0.075	0.120	0.055
3037	11.5	0.259	0.067	0.146	0.308	0.096
3014	11.5	0.259	0.064	0.144	0.244	0.103
3013	10.8	0.252	0.060	0.140	0.216	0.104
3016	9.5	0.203	0.024	0.055	0.121	0.036
3012	11.1	0.179	0.057	0.127	0.231	0.089
2969	11.9	0.159	0.043	0.093	0.195	0.061
2967	9.7	0.130	0.025	0.066	0.127	0.045
2998	12.4	0.129	0.024	0.183	0.286	0.110
2970	12.2	0.099	0.057	0.124	0.231	0.085
3028	12.3	0.094	0.088	0.232	0.314	0.180
3029	10.7	0.075	0.068	0.153	0.243	0.112
3032	12.1	0.075	0.081	0.171	0.383	0.112
2936	13.3	0.058	0.114	0.246	0.425	0.174
3026	10.9	0.057	0.056	0.178	0.310	0.115
3024	11.9	0.053	0.034	0.076	0.182	0.049
3025	12.3	0.053	0.029	0.066	0.135	0.045
3046	9.8	0.050	0.023	0.049	0.166	0.030
3004	11.7	0.050	0.064	0.142	0.230	0.103
3003	14.5	0.043	0.095	0.203	0.342	0.144
2972	11.6	0.035	0.001	0.899	1.857	0.273
3007	10	0.028	0.009	0.528	1.062	0.177

Table 5.6: Table shows (for V-pol) the final optimized y -intercept cut, unnormalized number of passing sim events, background estimate from just the exponential fit, the background estimate obtained from the distribution created from the inclusion of systematic uncertainties, as well as the uncertainty on that estimate for all passing Healpix bins in the V-pol analysis channel.

The Scale Factor C

C is multiplied by our simulated neutrino data set to scale the number of neutrinos being predicted by the model we are testing against. We want to be able to reject the model with the smallest amount of neutrino production possible, so that we can set the strongest limit possible. Thus in general we are optimizing by minimizing C . In order to save computer time however, when optimizing our linear discriminate cut's y -intercept value we optimize over an analogous quantity. Our neutrino model is Kotera Max. For a single Healpix bin with a background estimate of b then the likelihood of observing s neutrinos on top of our

Bin Number	optimized LD cut	Sim events passing	Expo fit Background	Background Estimate	Upper Error limit	Lower Error limit
3014	8.6	0.581	0.039	0.086	0.245	0.057
3013	9.7	0.473	0.105	0.196	0.552	0.123
3015	10.9	0.380	0.079	0.152	0.261	0.108
2936	10.3	0.238	0.067	0.127	0.316	0.080
2991	9.8	0.209	0.070	0.133	0.457	0.082
2988	9.7	0.195	0.132	0.243	0.631	0.152
2990	10.2	0.192	0.059	0.109	0.337	0.067
3031	10	0.094	0.040	0.071	0.223	0.043
3030	9.2	0.090	0.033	0.067	0.203	0.041
3004	12.9	0.077	0.082	0.438	0.844	0.218
2901	8.7	0.076	0.211	0.350	1.088	0.212
3010	8.4	0.051	0.139	0.287	0.630	0.187

Table 5.7: Table shows (for H-pol) the final optimized y-intercept cut, unnormalized number of passing sim events, background estimate from just the exponential fit, the background estimate obtained from the distribution created from the inclusion of systematic uncertainties, as well as the uncertainty on that estimate for all passing Healpix bins in the H-pol analysis channel.

background is given by equation 5.20

$$L(s) = \frac{\exp(-(b+s)) * (b+s)^b}{\Gamma(b+1, b)} \quad (5.20)$$

Equation 5.20 is really just the Poisson equation for a predicted count of b and an observation of $b+s$ with a few simple modifications. The Poisson equation's normalization of $b!$ becomes $\Gamma(b+1)$ to account for the fact that our events are weighted and thus not always integers. $\Gamma(b+1)$, the complete gamma function, then becomes $\Gamma(b+1, b)$, the upper incomplete gamma function, to account for the addition of b to s , which alters the integration limits when finding the normalization. What we want to know for a given Healpix bin though is not the probability of observing s neutrinos, but of observing more than $s * C$ neutrinos, $P(> s * C)$.

$$P(> s * C) = \int_{s*C}^{\infty} P(s') ds' \quad (5.21)$$

We want to find C such that the chance of observing more than $s*C$ neutrinos in any

Healpix bin is less than 10%.

$$0.1 = \prod_i P(> s_i * C) \quad (5.22)$$

Where s_i is the unnormalized amount of signal passing bin $_i$'s cuts given our Kotera Max neutrino model.

Past Optimizations

During Sam Stafford's analysis of the same 10% ANITA-III data set that is used as a training sample in this analysis three parameters were optimized for simultaneously by minimizing C. The first two parameters optimized for were the circular polarization strength threshold and circular polarization separation threshold used by the C-pol cut's Sam introduced to the Binned Analysis. The third parameter was the linear discriminant cut's slope. This analysis is using the values Sam's analysis optimized those parameters to.

Optimizing the Linear Discriminant Cut

The optimization of the linear discriminant cut's y-intercept (aka cutVal) does not use C. This is because the calculation of C requires all of the Healpix bins and this optimization is done for each bin individually. A similar method is used however. If S_{up} is the amount of signal that we can reject with 90% confidence for a single Healpix bin, and s is the amount of signal predicted by a model, then imagine what it means is $s = S_{up}$. It would mean that we could just barely reject the model. If however, $s > S_{up}$, it means we could not only reject the model that gave us s , but that we could scale s down some by some scale factor C and we would still be able to reject the model. Thus what we want is to maximize how much larger than S_{up} s is, so that we can have as small of a value for C as possible. This means that what we will actually be doing to optimize our linear discriminant cut, per Healpix bin, is maximizing s/S_{up} .

Both the background estimate b and the amount of passing simulated neutrinos s depend on the cutVal. Many different possible values for cutVal are tested. For each one of these

values s and b are calculated. s is just found by counting the number of simulated events that pass the linear discriminate cut (and all previous cuts) for the value of cutVal being tested. In order to take into account the systematic uncertainties in the estimation of b , many different random backgrounds are calculated using equation 5.15. Each of these random backgrounds creates a different likelihood function.

$$L_i(s, b_i) = \frac{\exp(-(b_i + s)) * (b_i + s)^b_i}{\Gamma(b_i + 1, b_i)} \quad (5.23)$$

Where b_i is the i th random background. A smeared likelihood function is created by summing over the different likelihood functions.

$$L_{\text{smeared}}(s) = \frac{1}{N} \sum_{i=0}^N L_i(s, b_i) \quad (5.24)$$

This smeared likelihood function accounts for the effects of the systematic uncertainties in the background estimate. Figure 5.13 shows an example likelihood $L(s)$ in red, calculated with the background estimate with no systematic uncertainties, and an example smeared likelihood $L_{\text{smeared}}(s)$ in blue, calculated using thousands of random backgrounds.

Similar to how C was calculated, S_{up} is now calculated such that only 10% of the likelihood function (now the smeared likelihood function) lies above S_{up}

$$0.1 = \int_{S_{\text{up}}}^{\infty} L_{\text{smeared}}(s') ds' \quad (5.25)$$

Figure 5.14 shows a typical data for s and S_{up} plotted against cutVal . Figure 5.15 shows a typical plot of s/S_{up} plotted against cutVal . As mentioned earlier, the optimal value for cutVal is when s/S_{up} is its maximum. Plots of s/S_{up} vs cutVal generally have a strong peak, as shown in Figure 5.15. Both Sam Stafford's ANITA-III analysis and Brian Dailey's ANITA-II analysis optimized for the best possible cutVal per bin as well, however their analysis did not include systematic uncertainties. In general the addition of systematic uncertainties causes the optimization to find higher linear discriminate cut y-intercept values. This in turn reduces both the background estimates and the chances of background events passing our analysis cuts.

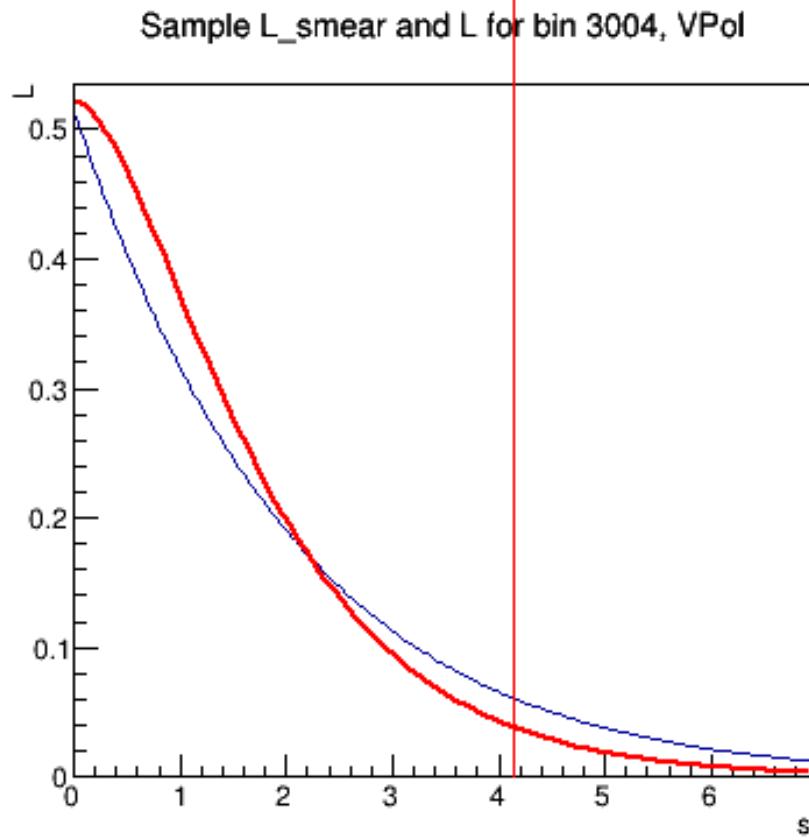


Figure 5.13: Plot shows the effect of smearing one of our likelihood functions with the inclusion of systematic uncertainty. The unsmeared likelihood is in red, while the smeared likelihood function is in blue. The red vertical line is the value of S_{up} found with the smeared likelihood function.

Optimizing the Linear Discriminate Cut over all Healpix Bins

Though we have already optimized each individual Healpix bin's y-intercept value for its linear discriminate cut (or `cutVal`), none of those optimizations take into consideration the bigger picture. We want to optimize for the best overall limit possible, not the best possible individual limits in each Healpix bin. In order to correct for this, we allow all of the passing Healpix bins' `cutVals` to shift up and down together. The overall shift that returns the lowest value for `C` is used, and all Healpix bins' background estimates and number of passing simulated events are reevaluated to reflect this change. This overall `cutVal` optimization happens before Healpix bins are cut for either having too much background

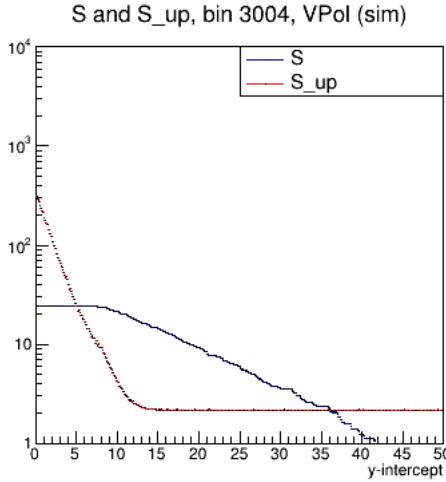


Figure 5.14: Plot displays both s (in blue) and s_{up} (in red) plotted on the same axis. Healpix bin 3004 for the V-pol channel is plotted here.

or low sensitivity.

In most cases this increases the `cutVal` for each bin by less than 0.3, sometimes not changing the final value at all. Increases in the `cutVal` reduces our background estimates and strengthens the linear discriminate cut.

Optimizing for Healpix Orientation

Our initial Healpix bin placement was simply chosen because it was the default. It is possible however that a more ideal orientation exists. If for example a large source of anthropogenic noise was near the boundaries between Healpix bins it could lead to a large number of weighted background events spread between two Healpix bins instead of centered inside of one bin. This would worsen our ability to find neutrinos in both Healpix bins and could cause worse exponential fits for both bins. It would also increase the systematic uncertainty in both bins. Thus it is believed some Healpix orientations could be more ideal than others.

In order to account for these effects the analysis is performed for 100 different varying orientations with different shifts in longitude and latitude. The approximate size of a single Healpix bin is 5.6 degrees wide in longitude, and 5.6 degrees wide in latitude. We sampled a grid of shifts to both longitude and latitude. Ten different shifts were tested for each, with

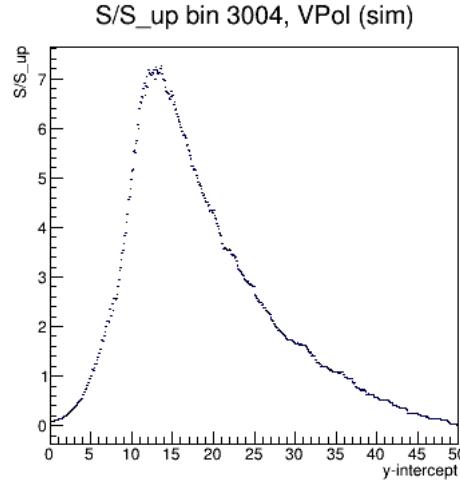


Figure 5.15: Plot displays s/s_{up} . The peak of this plot is this Healpix bin's optimized y-intercept cut value. Healpix bin 3004 for the V-pol channel is plotted here.

step sizes of 0.56 degrees. C was calculated (with the inclusion of systematic uncertainties) for each of the possible 100 orientations and the orientation that yielded the lowest value was selected. Table 5.8 and 5.9 show a color coded table of the values for C retrieved from the different orientations for V-pol and H-pol channels respectively.

		Theta										
		0	1	2	3	4	5	6	7	8	9	
Grid Index		Angular shift	0	-0.56	-1.12	-1.68	-2.24	-2.8	-3.36	-3.92	-4.48	-5.04
Phi	0	0.00606	0.0068	0.012445	0.01349	0.007674	0.006839	0.006918	0.008128	0.007718	0.005922	
	1	0.56	0.006202	0.00606	0.011954	0.009829	0.007499	0.007718	0.006839	0.006131	0.006761	0.005821
	2	1.12	0.006273	0.006607	0.016032	0.012374	0.007898	0.0068	0.00507	0.006494	0.006722	0.006131
	3	1.68	0.006131	0.007852	0.018197	0.009605	0.007718	0.005401	0.005623	0.005754	0.005656	0.008082
	4	2.24	0.005158	0.008561	0.013259	0.00866	0.006918	0.005464	0.006607	0.005888	0.006095	0.008463
	5	2.8	0.006202	0.007898	0.009386	0.008511	0.00861	0.00537	0.00537	0.005721	0.005495	0.009226
	6	3.36	0.0068	0.008511	0.00827	0.009495	0.006166	0.006026	0.004926	0.005188	0.007244	0.009943
	7	3.92	0.007286	0.010411	0.007079	0.007943	0.006202	0.00631	0.005278	0.005788	0.005821	0.007989
	8	4.48	0.006683	0.008128	0.005888	0.005991	0.006273	0.006273	0.005099	0.005591	0.007371	0.008318
	9	5.04	0.006095	0.006918	0.00606	0.00631	0.006569	0.005754	0.005041	0.006569	0.007674	0.006958

Table 5.8: Table shows the values of C returned by varying offsets in longitude and latitude for the Healpix orientation in the V-pol channel. More green and lower values are better. More red and higher values are worse.

		Theta										
Grid Index		0	1	2	3	4	5	6	7	8	9	
Phi	Angular shift	0	0.037368	0.033304	0.038905	0.063096	0.06166	0.039355	0.036728	0.030549	0.025264	0.039582
	0	0.56	0.031441	0.039582	0.034674	0.053089	0.032546	0.044668	0.038681	0.033497	0.028347	0.028184
	1	1.12	0.02884	0.035481	0.037154	0.042658	0.039582	0.045186	0.041448	0.051286	0.027861	0.038019
	2	1.68	0.029343	0.044412	0.037584	0.031081	0.033884	0.043904	0.04217	0.04217	0.032546	0.038019
	3	2.24	0.030903	0.032546	0.032546	0.038238	0.031623	0.04121	0.034674	0.041687	0.048417	0.036099
	4	2.8	0.026303	0.029854	0.034674	0.038459	0.040041	0.035075	0.036517	0.041687	0.052481	0.035075
	5	3.36	0.034277	0.029854	0.030026	0.040738	0.041928	0.033113	0.034674	0.037154	0.077179	0.040272
	6	3.92	0.032359	0.029343	0.029343	0.034277	0.035075	0.039811	0.043401	0.042904	0.050408	0.046238
	7	4.48	0.039129	0.027861	0.027861	0.03694	0.035481	0.040504	0.04217	0.049831	0.057544	0.048978
	8	5.04	0.026915	0.028022	0.028022	0.032359	0.032359	0.028022	0.034475	0.059566	0.059566	0.055271

Table 5.9: Table shows the values of C returned by varying offsets in longitude and latitude for the Healpix orientation in the H-pol channel. More green and lower values are better. More red and higher values are worse.

Chapter 6

DISCUSSION AND CONCLUSIONS

Analysis of the 10% data set is now complete. Cuts are set. The linear discriminant cut has been optimized for all Healpix bins for both the H-pol and V-pol channels. Background estimates for the 90% sample have been estimated. One event has been found passing our 10% sample. The analysis of our low sensitivity Healpix bin sideband, and the 90% sample H-pol channel is underway. A complete unblinding and examination of the 90% sample V-pol channel is forthcoming. This section outlines the results from our 10% data set, our sidebands and our H-pol channel, as well as discussing what comes next.

6.1 Results with the 10% data set

Tables 5.6 and 5.7 show our final background estimates, optimized linear discriminant cut y-intercepts, and the amount of simulated neutrinos passing our final cuts. With our final cuts tuned, we see one V-pol event passing, and 0 H-pol event passing.

6.1.1 Events Passing in our 10% data set

The event passing in our 10% data set shows characteristics of a payload blast event. The events and some key information about each of them can be found in table 6.1. For comparison, Sam Stafford's analysis saw three V-pol and five H-pol events passing in his 10% sample [47]. The addition of the satellite stripe cut removed many of those events. The addition of systematic uncertainties also increased optimized cutVals which created more stringent cuts. More work on the payload blast quality cut should be done in the future.

Event Number	Polarization	Run Number	Healpix Bin	Event Weight
35083936	V	254	3032	1

Table 6.1: Table shows events passing in both H and V-pol analysis channels.

6.2 Analysis Efficiency

Based on our 10% data set we can estimate the efficiency of our analysis compared to Brian Dailey's ANITA-II binned analysis and Abby Vieregg's ANITA-II Clustering analysis. This analysis is currently the only ANITA-III neutrino search with its analysis cuts set. Figure 6.1 shows the efficiency curves for both ANITA-II analyses. Figure 6.2 shows the efficiency curves for passing H-pol Healpix bins. This analysis appears to on average achieves earlier efficiency turn on and a sharper efficiency rise than the ANITA-II analyses.

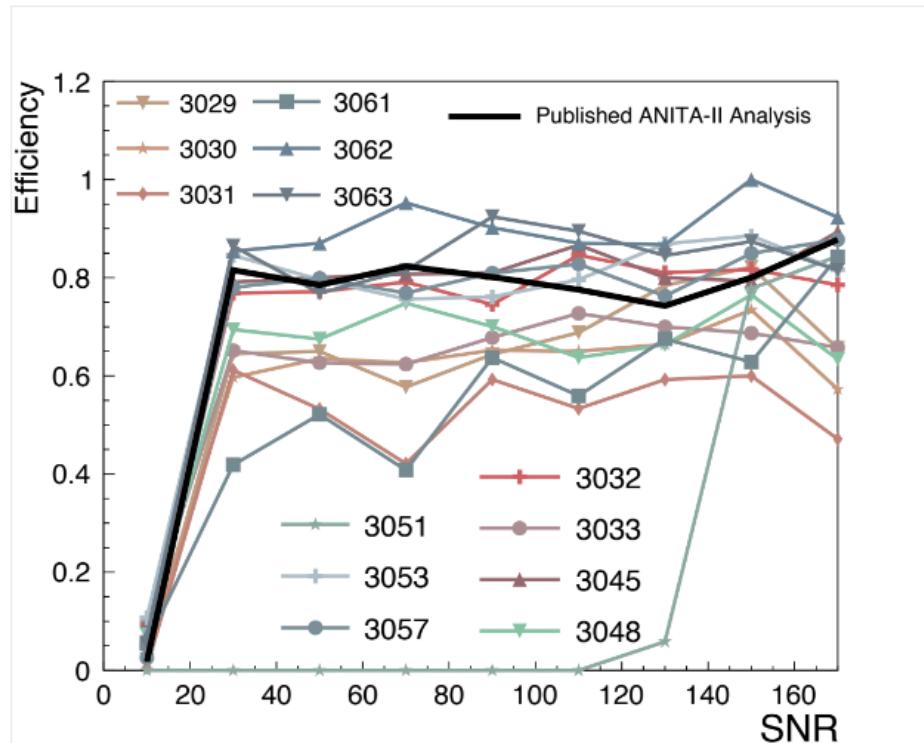


Figure 6.1: In black is the SNR vs efficiency curve for the ANITA-II clustering analysis. Other lines are the SNR vs efficiency curves for individual Healpix bins in the ANITA-II binned analysis. Figure credit to Oindree Banerjee

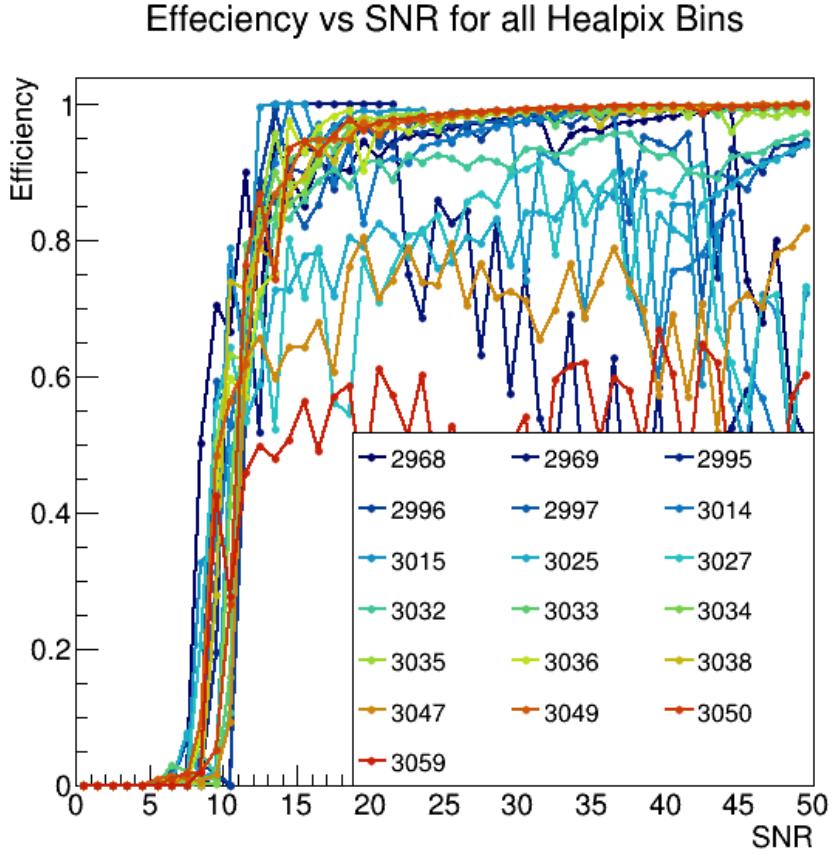


Figure 6.2: Figure shows efficiency vs SNR curves for passing H-pol Healpix bins. Efficiency calculation does not include loss from quality cuts, which should be negligible at high SNR. Data at SNRs > 20 as shown as a rolling average (of 2) to help counteract the lack of data at high SNRs.

6.3 Results from our 90% sideband data set

As mentioned in section 5.2.7, The least sensitive 1% of Healpix bins are set aside to be used as a sideband. We do not expect to see a neutrino in them, because they have such low sensitivity, but they should have typical background events. The sideband bins used are Healpix bins 3009 and 3047 in H-pol, and Healpix bins 3001, 3008, 3009 and 3011 in V-pol. In H-pol, 11 events are observed to pass our LD cut, all in Bin 3009. In V-pol 2 events are observed to pass our LD cut, one in 3001, and one in 3011. Table 6.2 shows all

13 events.

Event Number	Polarization	Run Number	HpBin	Weight	Set Number	LD Value	Latitude	Longitude
36916590	V	266	3011	1	7	9.315624	-77.9071	-6.47056
64175392	V	364	3001	1	1	14.29046	-76.396	-157.773
40579667	H	280	3009	1	4	9.31357	-77.8852	-34.3578
40993619	H	281	3009	1	7	8.883397	-77.869	-34.3293
41049372	H	281	3009	1	9	9.522537	-77.9022	-34.3152
41064920	H	281	3009	1	5	9.829857	-77.9164	-34.2397
41073934	H	281	3009	1	9	10.94168	-77.8822	-34.4221
41103084	H	281	3009	1	7	11.02763	-77.8647	-34.3658
44543143	H	296	3009	1	1	18.56046	-77.887	-34.3355
44600768	H	296	3009	1	5	10.81732	-77.8824	-34.4193
44677846	H	296	3009	1	3	29.27787	-77.9191	-34.6429
44678128	H	296	3009	1	5	15.19691	-77.9262	-34.6244
44678408	H	296	3009	1	5	11.77425	-77.8906	-34.5353

Table 6.2: Table shows Events passing the LD cut in 90% dataset sideband Healpix bins. All H-pol events cluster with one another. Both V-pol events appear to be payload blasts.

All 11 of the H-pol events are found to cluster within 40 km of at least 9 of the other H-pol events. A German research base is known to be in that area of the ice. A clustering cut would remove all 11 H-pol events.

Both of the V-pol events appear to be payload blast like, further evdence that our analysis's payload blast cut needs refinement.

6.4 Results from out 90% H-pol data set

In the full 90% H-pol data set, 6 events are observed to pass. Two of them are observed by other analysis and are believed to be Cosmic Ray events. Three of them are believed to be payload blasts. One does not fall into either of those catagories, but also does not appear to be strongly impulsive. Table 6.3 shows all 6 passing H-pol events. Other ANITA-III analysis have found upwards of 20 CR candidates, our analysis however unfortuently has rather poor sensitivity to H-pol events because so few H-pol Healpix bins pass. Using a

larger training data set or loosening requirements on the amount of data needed to attempt a fit could help this in the future.

Run Number	Event Number	Hp Bin	Data Set	Seen by other analyses	Payload Blast like?
216	23937896	3015	4	No	Yes
236	28876488	2990	7	No	Yes
250	33484995	3013	2	Yes	No
267	36975694	3010	3	No	No
284	41529195	3030	4	Yes	No
429	82011215	2936	4	No	Yes

Table 6.3: Table shows events passing the LD cut in the 90% data set H-pol channel. Two events were observed by other analyses and are cosmic ray candidates. Three are believed to be payload blasts. Once event, 36975694, required further investigation.

6.5 Discussion

Payload blasts are observed passing in all three of our examined data sets. This implies a better cut is needed for them in the future. Though we see no events clustering in our 90% H-pol data set, we see 11 clustering in a single Healpix bin from our H-pol sideband Healpix bin 3009. The German base located there appears to be the second brightest spot in Antarctica for our analysis just before the LD cut. The first being McMurdo Station. This might implicate areas with large production of anthropogenic noise may be insufficiently modeled by a 10% data set using our methods. Clustering Removes these events, and appears to be necessary for our analysis to deal with Healpix bins with very large anthropogenic contributions .

In most cases for Healpix bins examined in both the 10% and 90% data sets, our fits match the observed distributions well. Even in the case of H-pol Healpix bin 3009, a low statistics tail is what emerges to lead to the 11 passing events, while the bulk of the distribution is well modeled by our exponential fit. I think this shows our method of fitting the 10% data set to estimate the 90% data set works well in most cases. Figure 6.3 shows the LD distributions for H-pol Healpix bin 3009 from the 10% and 90% data sets side by side, while Figure 6.4 shows the LD distributions for H-pol Healpix bin 3031 from the 10% and

90% data sets side by side. Healpix bin 3009 is the worst case of our background estimate not working, while Healpix bin 3031 shows a result with a good match. Most other healpix bins' 10% and 90% distributions match more similarly to Healpix bin 3031 than 3009 in this regard.

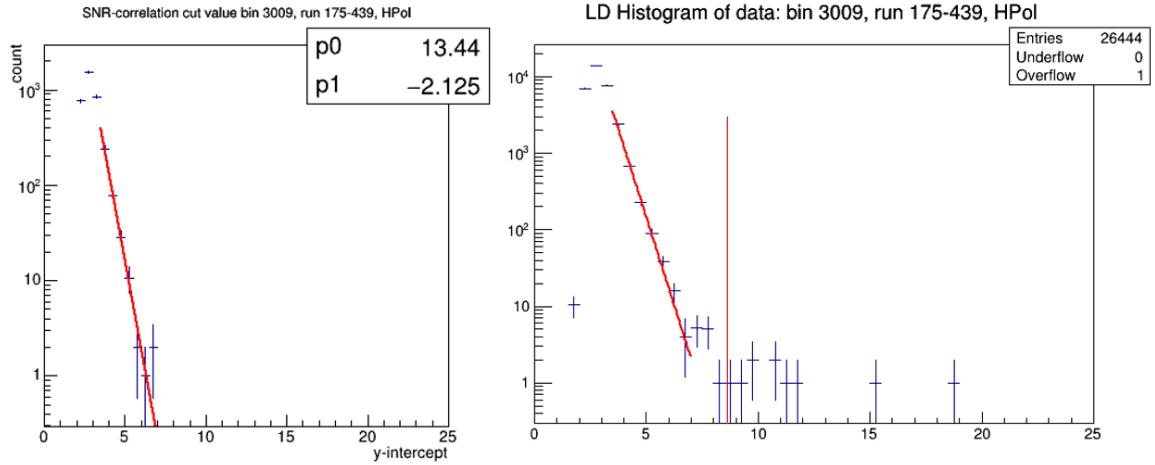


Figure 6.3: The 10% data set and 90% data set LD distributions side by side. The 10% is on the left, while the 90% is on the right. The red diagonal line is the exponential fit to the 10% data. In the 90% data is has been scaled up by a factor of 9. The red vertical line is the LD cut value for this Healpix bin. The 90% data shows a long tail for this bin not observed in the 10% data set.

The excess of passing events that both Brian Dailey and Sam Stafford's analyses observed appears to be mostly resolved. We do have a handful of payload blast like events passing, but overall the addition of systematic uncertainty to our optimization and the satellite stripe cut has served its purpose, and greatly reduced the number of events observed to be passing.

6.6 Future Analysis Improvements

In a future iteration of the binned analysis many additional things can be done to further improve on our results.

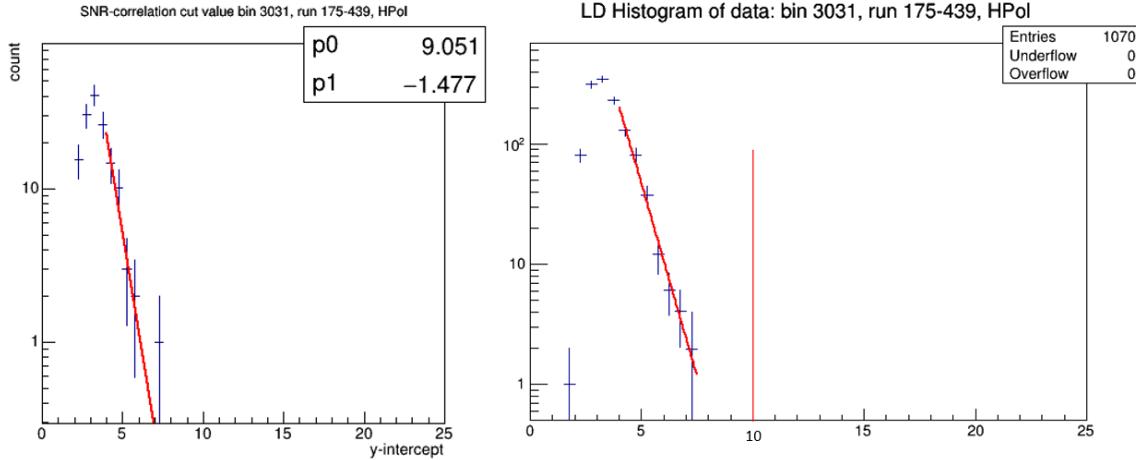


Figure 6.4: The 10% data set and 90% data set LD distributions side by side. The 10% is on the left, while the 90% is on the right. The red diagonal line is the exponential fit to the 10% data. In the 90% data is has been scaled up by a factor of 9. The red vertical line is the LD cut value for this Healpix bin. The 10% and 90% data sets match remarkably well for this Healpix bin.

6.6.1 Improved Payload blast cut

There is evidence that out payload blast cut needs to be improved. All examined data sets show payload blast like events passing.

6.6.2 Definition of SNR

The method we are using to calculate SNR has been shown to be less efficient sense our analysis began. We are using the early part of the waveform to estimate the noise, however in ANITA-III the early part of the waveform is very short. An alternate method has been developed by collaborators, which calculates the SNR from nearby noise events. Switching the entire analysis over to using an alternate definition of SNR could have wide reaching effects, because our primary cut, the linear discriminate cut, is a linear combination of SNR and correlation peak.

6.6.3 Size of the Training Data Set

We use 10% of the total data to train our analysis cuts. However it is possible that using more of the data would be more optimal. Having more data to fit will result in more accurate and robust background estimates. What percent of the data should be used for training has not been looked into in detail, and should be investigated.

6.6.4 Rigidity

One concern we have with our method of obtaining our background estimate is the effects a single outlying event can have. It appears that roughly 50% of Healpix bins rejected for having a p-value of less than 0.05 have a single, high linear combination, value. This may mean we are rejecting Healpix bins with candidates in them. One way to deal with this is to always exclude the largest event, and add an additional systematic uncertainty to our background estimate based on how much excluding that event changed the fit. This should prevent any Healpix bins from being rejected by outliers and give us some insight into how much an outlier is effecting our background estimate.

6.6.5 Optimizing for Sensitivity

We currently optimize our analysis to set the best limit on neutrino models assuming we see a number of events passing consistent with our background. In the future, we may want to instead optimize to maximize our chance of detecting a neutrino.

6.6.6 Gerrymandering our Healpix bins

Though right now we are using equal area Healpix bins, that is not an important feature of our analysis. Each bin already has a different sensitivity to neutrinos due to variations in ice thickness and how long ANITA observed that ice. In the future we may want to create more finely spaced Healpix bins, then combine bins with neighbors until we are able to get a usable background estimate from the new, oddly shaped area. This could potentially allow us to keep more of ice we are currently rejecting with our Healpix bin cut.

BIBLIOGRAPHY

- [1] The ANITA Collaboration. Dynamic tunable notch filters for the antarctic impulsive transient antenna (anita). *arXiv eprint*, 2017.
- [2] York G. Donald et al. Multiwavelength astronomy. <http://ecuip.lib.uchicago.edu/multiwavelength-astronomy/index.html>.
- [3] The Pierre Auger Collaboration. The pierre auger cosmic ray observatory. *Nuclear Instruments and Methods in Physics Research A*, 798:172–213, 2015.
- [4] Peter Tinyakov. Latest results from the telescope array. *Nuclear Instruments and Methods in Physics Research A*, 742:29–34, 2014.
- [5] B.P. Abbott et al. Observation of gravitational waves from a binary black hole merger. *Phys. Rev. Lett.*, 116:061102, 2016.
- [6] IceCube Collaboration. First observation of pev-energy neutrinos with icecube. *Phys. Rev. Lett.*, 111:021103, 2013.
- [7] B. T. et al Cleveland. First direct observation of time-reversal violation. *Astrophys. J.*, 496:505, 1998.
- [8] K. S. et al Hirata. Results from one thousand days of real time directional solar neutrino data. *Phys. Rev/ Lett.*, 65:1297–1300, 1990.
- [9] A Suzuki. The 20th anniversary of sn1987a. *Journal of Physics: Conference Series*, 120, 2008.

[10] Bradley W. Carroll and Dale A. Ostlie. *An Introduction to Modern Astrophysics*. Pearson Education Inc., 2007.

[11] Icecube Collaborationi. Evidence for high-energy extraterrestrial neutrinos at the ice-cube detector. *Science*, 342:6161, 2013.

[12] G. T. Zatsepin V. S. Beresinsky. *Phys. Lett. B*, 28:423, 1969.

[13] K. Greisen. *Phys. Rev. Lett.*, 16:48, 1966.

[14] V. A. Kuzmin G. T. Zatsepin. *JETP Lett.*, 4:78, 1966.

[15] James J. Beatty et al. The highest energy cosmic rays. *Annu. Rev. Nucl. Part. Sci.*, 59:319–345, 2009.

[16] et al. Seo E. *Astrophys.J.*, 378:763, 1991.

[17] et al. Grigorov N. *Proc. Int. Cosm. Ray Conf.*, 12, 1971.

[18] et al. Nagano M. *J. Phys. G*, 18:423, 1992.

[19] et al. Antoni T. *Astropart. Phys.*, 24:1, 2005.

[20] J. et al. Abraham. *Phys. Rev. Lett.*, 100:211101, 2008.

[21] Perrone L. *Proc. Int. Cosm. Ray Conf.*, 30:331, 2007.

[22] et al. Takeda M. *Astropart. Phys.*, 19:447, 2003.

[23] R. et al. Abbasi. *Ap. J.*, 684:790, 2008.

[24] Swordy SP. Cronin JW, Gaisser TK. *Sci. Am.*, 276:44, 1997.

[25] Engel et al. Neutrinos from propagation of ultrahigh energy protons. *Phys. Rev. D*, 64, 2001.

[26] Q. A. Askar'yan. Coherent radio emission from cosmic showers in air and in dense media. *JETP Lett.* ??, 48:988–990, 1965.

[27] P. W. Gordham et al. The antarctic impulsive transient antenna ultra-high energy neutrino detector design, performance, and sensitivity for 2006-2007 balloon flight. *Astropart. Phys.*, 32:10–41, 2009.

[28] M. et al. Ackermann. *Ap. J.*, 675:1014, 2008.

[29] I. et al. Kravchenko. *Phys. Rev. D*, 69:082002, 2006.

[30] S. Hoover. A search for ultrahigh-energy neutrinos and measurement of cosmic ray radio emission with the antarctic impulsive transient antenna. *Dissertation*, 2010.

[31] R. J. Protheroe and P. A. Johnson. *Astropart. Phys.*, 4:253, 1996.

[32] O. E. et al. Kalashev. *Phys. Rev. D*, 66:063004, 2002.

[33] O. E. et al. Kalashev. *Phys. Rev. D*, 65:103003, 2002.

[34] C. et al. Aramo. *Astropart. Phys.*, 23:65, 2005.

[35] M. et al. Ave. *Astropart. Phys.*, 23:19, 2005.

[36] Huber P. Barger, V. and D. Marfatia. *Phys. Lett. B*, 642:333, 2006.

[37] P. W. Gorham et al. Observational constraint on the ultra-high energy cosmic neutrino flux from the second flight of the anita experiment. *Phys. Rev. D*, 82, 2010.

[38] Abigail Goodhue Vieregg. The search for astrophysical ultra-high energy neutrinos using radio detection techniques. *Dissertation*, 2010.

[39] N. et al. Lehtinen. *Phys Rev. D*, 69:013008, 2004.

[40] S. W. et al. Barwick. *Phys. Rev. Lett.*, 96:171101, 2006.

[41] P. W. Gorham et al. Characteristics of four upward-pointing cosmic-ray-like events observed with anita. *Phys. Rev. Lett.*, 117, 2016.

[42] Jiwoo Nam. Overview of the fourth flight of the anita experiment. *ICRC 2017 Proceedings*, 2017.

- [43] Ryan Hupe. Investigating the performance of the interferometric trigger for future flights of the antarctic impulsive transient antenna. *Dissertation*, 2015.
- [44] S. Wissel. Overview of the third flight of the anita long-duration balloon payload. *ICRC 2015 Proceedings*, 2015.
- [45] G. Fowles. *Introduction to Modern Optics (2nd ed.)*. Dover, 1989.
- [46] Brian Dailey. Analysis of the second flight of antarctic impulsive transient antenna with a focus on filtering techniques. *Dissertation*, 2017.
- [47] Sam Stafford. A search for ultra-high energy cosmic neutrinos: Data analysis of the antarctic impulse and transient antenna, third flight. *Dissertation*, 2017.
- [48] The HEALPix Team. Healpix, data analysis, simulations and visualization on the sphere. <http://healpix.sourceforge.net/>.
- [49] Eric W. Weisstein. Log normal distribution. MathWorld—A Wolfram Web Resource <http://mathworld.wolfram.com/LogNormalDistribution.htm>.
- [50] K. A. et al. Olive. Review of Particle Physics. *Chin. Phys.*, C38:090001, 2014.

Appendix A

FITTING WITH ROOT'S WEIGHTED LOG LIKELIHOOD OPTION

Root has built in fitting methods that utilize minuit's function fitting routines. This is a wonderful tool when you don't need to care about the details. Root can be rather cryptic if you need to know exactly what it is doing with your data. Here I outline exactly what root does when you tell it to fit with the 'LW' option.

A.1 What 'LW' means to Root and when to use it

Root's fitting option's second input are various fitting flags. 'L' in these flags means that instead of doing a standard chi squared fit to the supplied histograms's data it will do a negative log likelihood fit. The 'W' means that the binned histogram data is weighted. The criterion for passing the 'W' flag is obvious. If your data is weighted, pass the 'W' fit option. One would want to use a negative log likelihood minimization instead of a chi squared when individual bin's in your histogram have low statistics. Negative log likelihood minimizations use Poisson distributions which fit low statistics situations better than the normal distributions that chi squared minimizations rely on.

A.2 What Root does when you pass Fit the 'LW' options

FitUtil::EvaluatePoissonLogL is the function that gets called when you use the Fit function on a histogram in Root. Within EvaluatePoissonLogL are several if statements that alter the equation that is being minimized. If sumW2 is on, it alters the function being minimized.

It is off by default. If off the equation being minimized per bin is:

$$\log L_i = y_i \ln \frac{y_i}{f(x_i)} - y_i + f(x_i) \quad (\text{A.1})$$

where y_i is the total number of event in the i th bin, x_i is the center of the i th bin, and $f(x)$ is the function being fit evaluated at position x . If `sumW2` is on, the equation being minimized becomes:

$$\log L_i = y_i \left(y_i \ln \frac{y_i}{f(x_i)} - y_i + f(x_i) \right) \quad (\text{A.2})$$

The other two if statements alter how $f(x)$ are calculated. `UseBinVolume` adds a normalization to $f(x)$, multiplying it by the bin volume. For example, if you have a 2d histogram with bin widths of 0.5 in one dimension, and 0.75 in the other dimension, then:

$$f'(x) = 0.5 * 0.75 * f(x) \quad (\text{A.3})$$

and $f(x)$ is replaced by $f'(x)$ in the equations for $\log L_i$. For ‘LW’ fits `UseBinVolume` is by default on. `UseBinIntegral` changes $f(x)$ from the function that you are fitting, to the integral of that function evaluated with a lower bound equal to the low edge of the bin, and the upper bound equal to the upper edge of the bin. For ‘LW’ fits `UseBinIntegral` is by default off, but can be turned on by passing the flag ‘I’ along with ‘L’ and ‘W’ to the `Root` fitting function.

Even when doing a negative log likelihood fit, the fit function will return a chi squared value. If you are trying to evaluate the goodness of your fit then this is not the value you want to use. You need to call the value ‘`edm`’ from `TVirtualFitter::GetStats`. This is equal to:

$$edm = LogL = \sum_i \log L_i \quad (\text{A.4})$$

One strange thing that I have noted with my usage of these methods is that the goodness of fit value, ‘`edm`’, seems to return the same value no matter what you have `SumW2` set to.

Whether it is on or off, it returns the value of LogL as if it were off. It seems to still do the minimization with the correct SumW2 setting, however. I believe the reason it returns the same value no matter how you have SumW2 set is so the user has something consistent to compare to other goodness of fit measurements no matter the SumW2 setting.

Appendix B

FITTING A 2D GAUSSIAN TO A STANDARD ERROR ELLIPSE

The standard equation for a Gaussian distribution in one dimension is:

$$G(x) = N * e^{-\frac{(x-x_0)^2}{2\sigma^2}} \quad (\text{B.1})$$

Where N is the Gaussian functions normalization. x_0 is the mean value of x , and σ^2 is the variance of x . This can then be generalized into two dimensions for the case when the two parameters are correlated with one another as:

$$G(x, y) = N * e^{-\frac{1}{2(1-\rho^2)}\left(\frac{(\theta_i - \hat{\theta}_i)^2}{\sigma_i^2} - 2\rho\frac{(\theta_i - \hat{\theta}_i)(\theta_j - \hat{\theta}_j)}{\sigma_i \sigma_j} + \frac{(\theta_j - \hat{\theta}_j)^2}{\sigma_j^2}\right)} \quad (\text{B.2})$$

Where ρ represents the correlation between the parameters θ_i and θ_j , $\hat{\theta}_i$ and $\hat{\theta}_j$ are the mean values of θ_i and θ_j and σ_i^2 and σ_j^2 are the variances of the parameters. Everything we need to know for Equation B.2 can be calculated given an error ellipse for the two parameters θ_i and θ_j .

σ_{inner} in Figure B.1 is the distance to the error ellipse from the center of the ellipse for a vertical line drawn through the center of the ellipse. It can be used to calculate the correlation, ρ , between θ_i and θ_j [50].

$$\rho = \sqrt{1 - \left(\frac{\sigma_{inner}}{\sigma_i}\right)^2} \quad (\text{B.3})$$

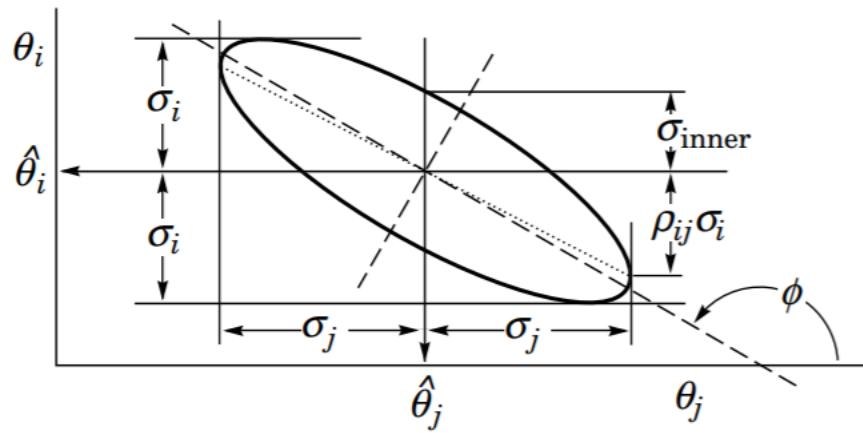


Figure B.1: Figure shows the error ellipse for two correlated parameters [50].

ρ is also be defined as:

$$\rho = \frac{Cov[\theta_i, \theta_j]}{\sigma_i \sigma_j} \quad (B.4)$$

Note that if θ_i and θ_j are correlated, ρ will be positive. If they are anticorrelated then ρ will be negative. ρ varies between -1 and 1 with a value of 0 meaning the two parameters are not correlated and the error ellipse is a circle. Larger values of ρ will result in a more stretched out error ellipse.

Appendix C

STEP BY STEP INSTRUCTIONS ON PROGRAMMING THE TUFF MASTER

Before the TUFF Master can be used to communicate with the TUFFs it must be programmed. This appendix outlines step by step instructions on how to program it.

C.1 Step 1: Preparing the TUFF Master to be programmed

First we need to get the TUFF Master into a state in which it can be programmed. Connect the TUFF Master to a windows computer via a micro USB to USB cable. Switch the switch next to the micro USB port into ‘DEBUG’ mode. (If the TUFF Master is not being programmed it should be in ‘DEVICE’ mode.) Bridge the 3.3 V and PB5 pins using a jumper.

C.2 Step 2: Preparing your computer

Open the software Energia. It can be downloaded for free from <http://energia.nu/download>. Once inside Energia, open ‘tuff_master.ino’. The latest version of the TUFF Master firmware can be downloaded from <https://github.com/barawn/tuff-master> along with a readme that explains the TUFF Master’s basic commands.

C.3 Step 3: Setting up Energia

From the drop down ‘Tools’ menu select ‘Board’ and choose ‘Launchpad (Tiva C) w/tm4c123 (80 MHz)’. If you do not see that option available then select ‘Board Manager...’ and install the ‘Energia Tiva C Boards’.

Energia must also be told which Com port to find the TUFF Master at. Select the drop down ‘Tools’ menu again and select ‘Port’ and choose the COM port the TUFF Master is connected to your computer through. If the TUFF Master is the only device connected to your computer then it should be the only option there.

C.4 Step 4: Open Energia’s Serial Monitor

Open the serial monitor by selecting ‘Serial Monitor’ from the drop down ‘Tools’ menu. The serial monitor prints to the screen communications to the TUFF Master and the TUFF Master’s responses. At the bottom of the screen monitor window set the baud rate to 115200 and select ‘Both NL & CR’ for the line endings.

C.5 Step 5: Programming or uploading the TUFF Master

Click the upload button, which looks like an arrow to the right. This should program the TUFF Master. If this is the first time the TUFF Master is being programmed, in the serial monitor you should see:

```
{“log”: “updating BOOTCFG”}  
{“log”: “BOOTCFG updated”}  
{“log”: “boot irfcm unassigned v#.#”}
```

If the TUFF Master has already been assigned an iRFCM, then you will see that number replacing the unassigned in the last log output.

C.6 Step 6: Assigning the iRFCM Addresss

If the TUFF Master does not have an iRFCM address already, or it needs to be changed then that can be done now in the serial monitor. The serial monitor can be used to send commands to the TUFF Master. Send the following command to assign the TUFF Master to iRFCM N:

```
{“set”:{“addr”:N, “save”:1}}
```

Where N should be replaced with the iRFCM number (generally between 0 and 3). The serial monitor should print the return statement:

```
{“ack”:N}
```

Other commands can also be sent this way.

Appendix D

PLOTS FOR PASSING HEALPIX BINS IN H AND V-POL CHANNELS

The plots here are all passing Healpix bins in both the H-pol and V-pol analysis channels. For each Healpix bin, he differential distribution plots (including the exponential fits), the plots of s and s_{up} , the plots of s/s_{up} , and the plots of the linear discriminant cut are all shown. Plots are ordered from most sensitive Healpix bin to least sensitive.

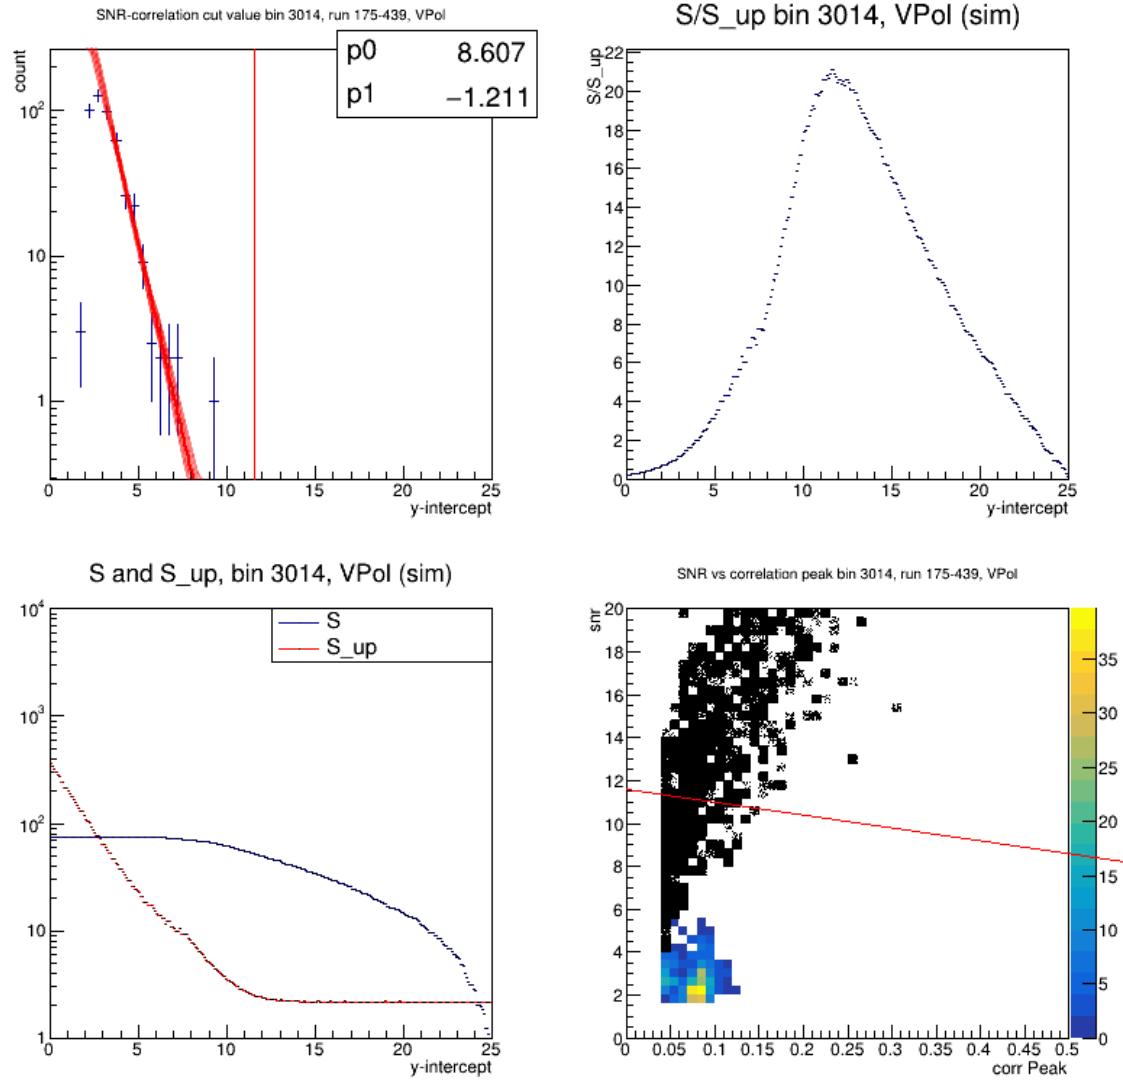


Figure D.1: Optimization plots for Healpix bin 3014 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

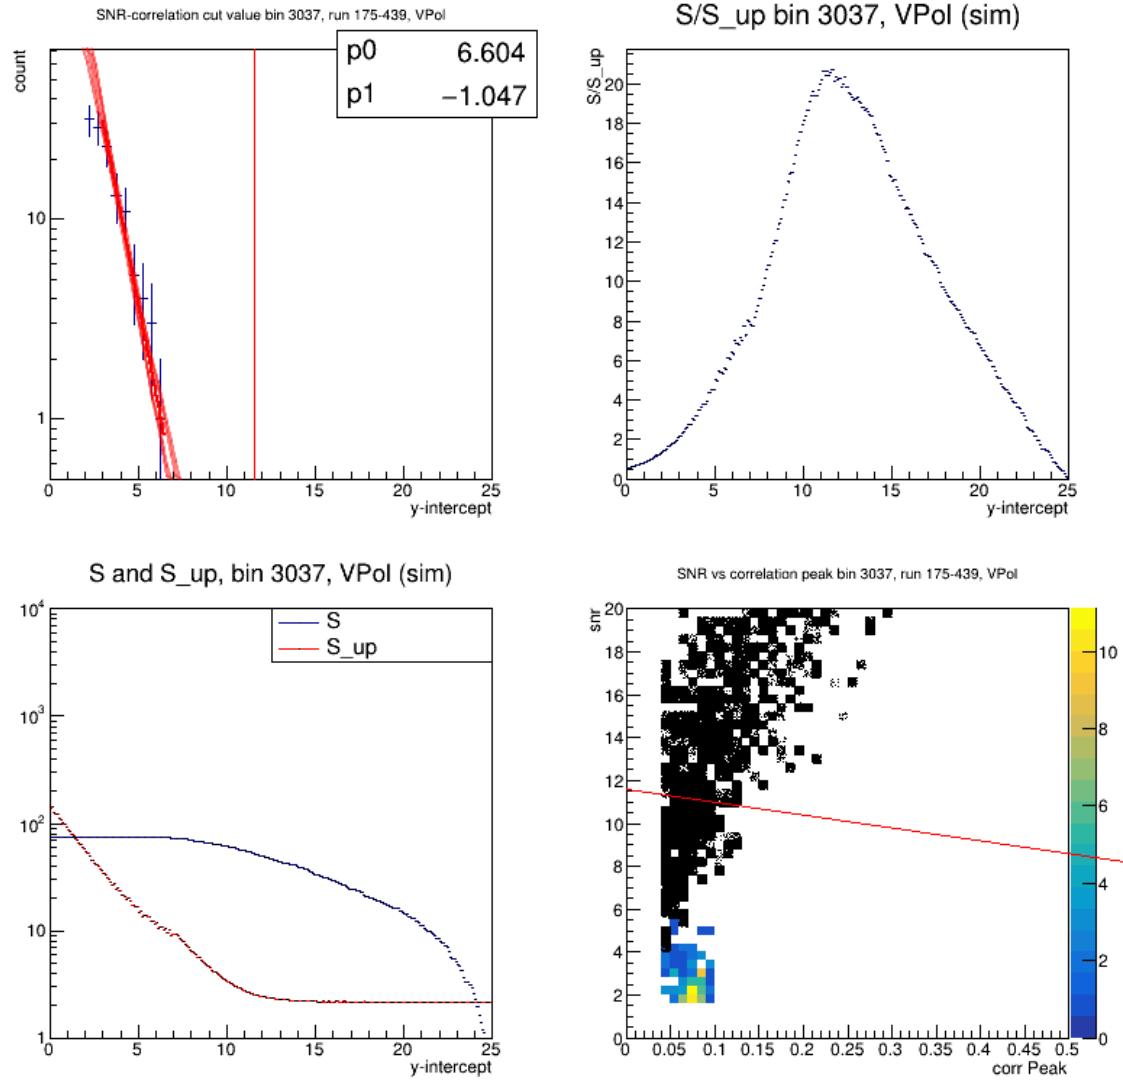


Figure D.2: Optimization plots for Healpix bin 3037 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

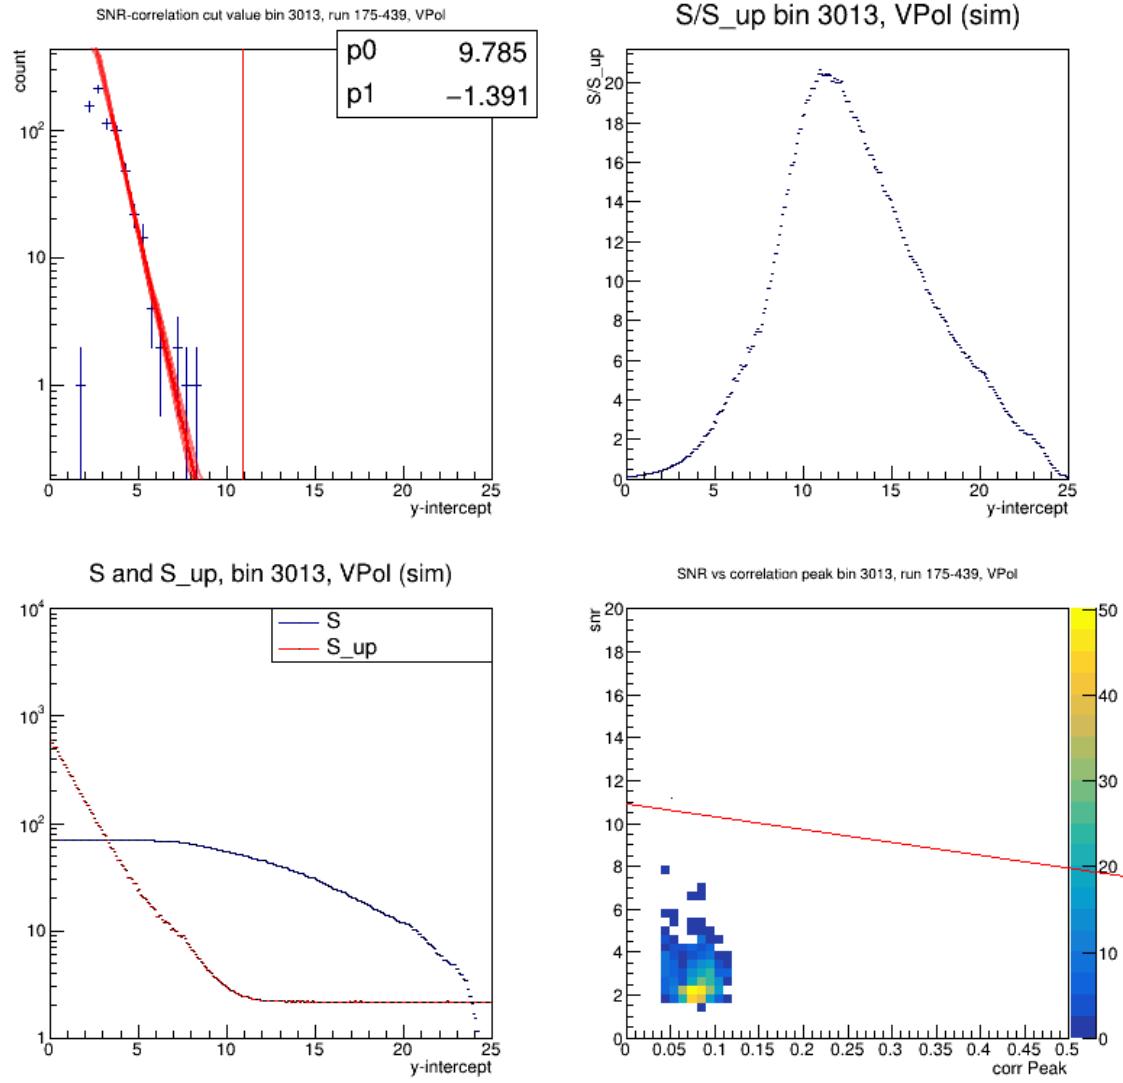


Figure D.3: Optimization plots for Healpix bin 3013 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

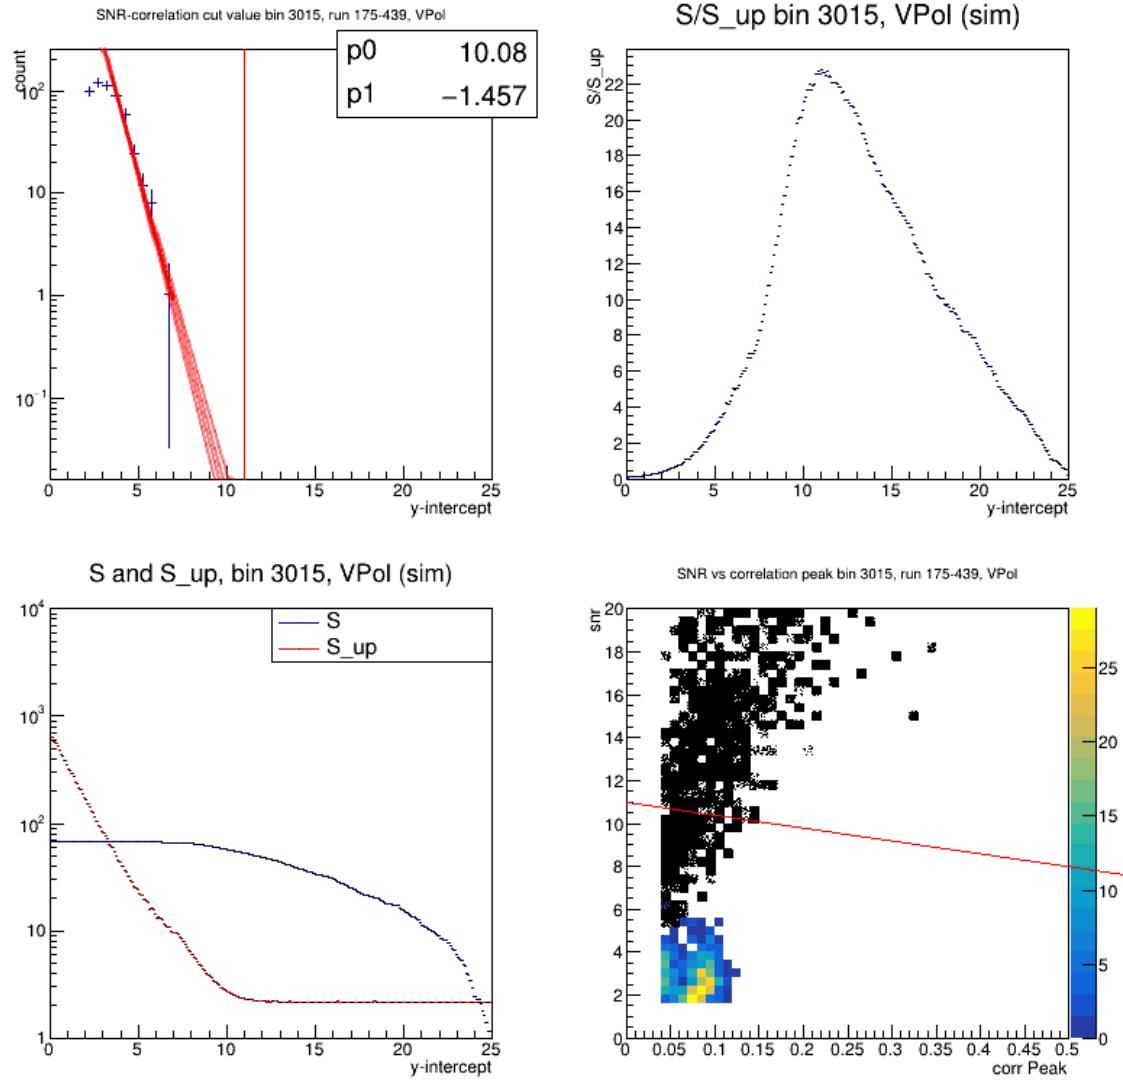


Figure D.4: Optimization plots for Healpix bin 3015 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

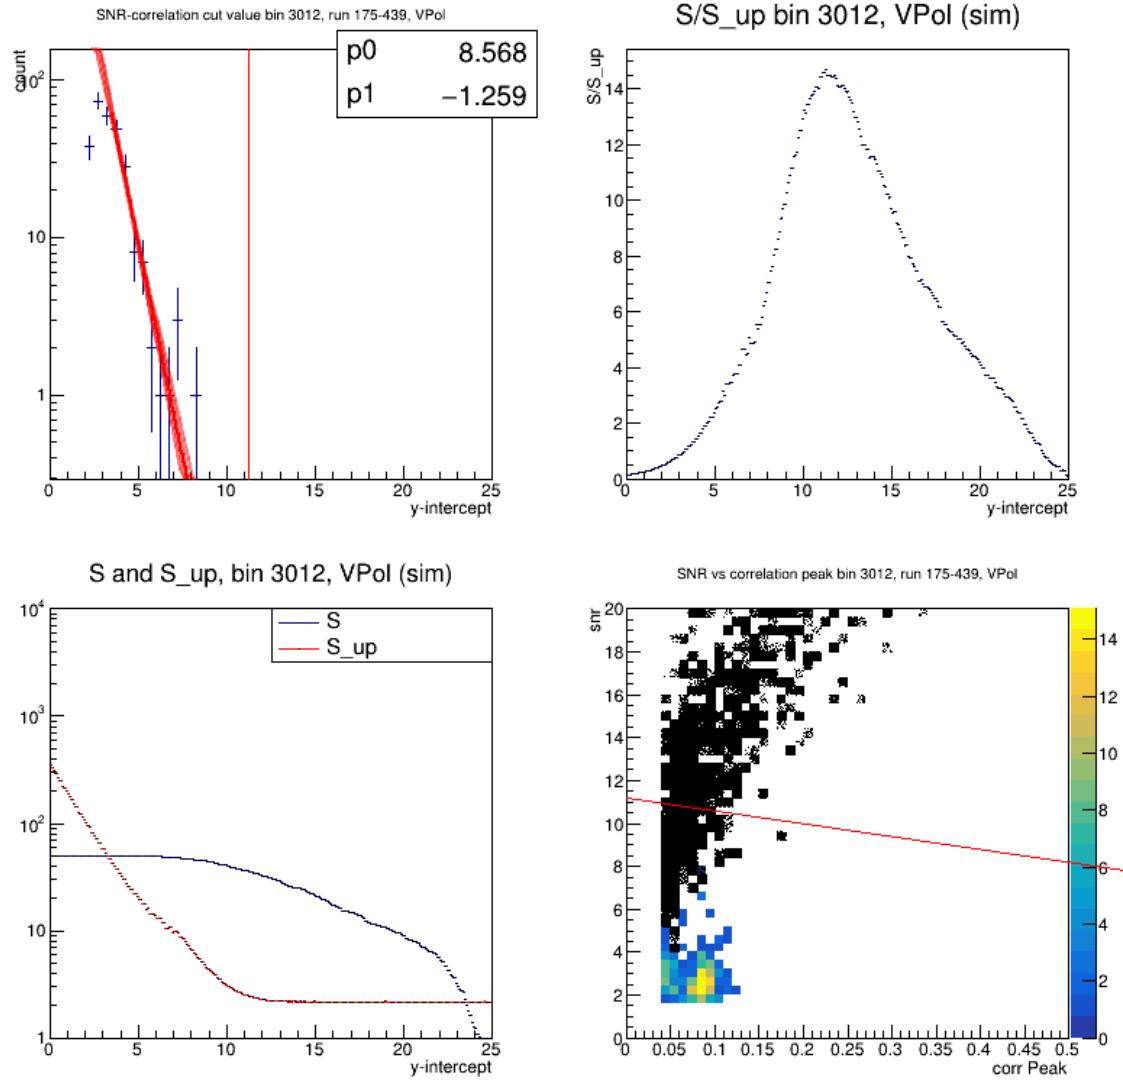


Figure D.5: Optimization plots for Healpix bin 3012 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

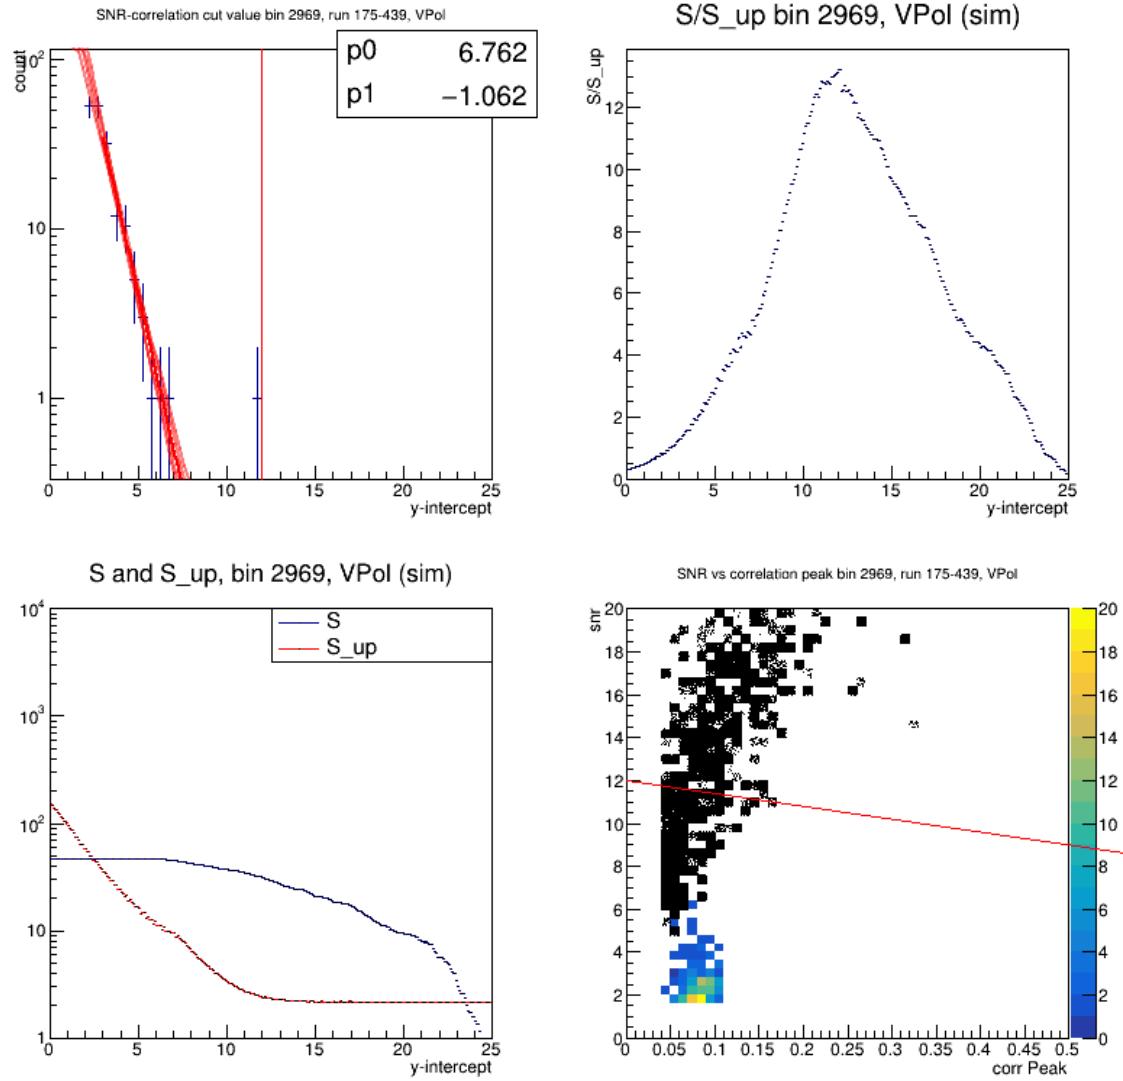


Figure D.6: Optimization plots for Healpix bin 2969 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

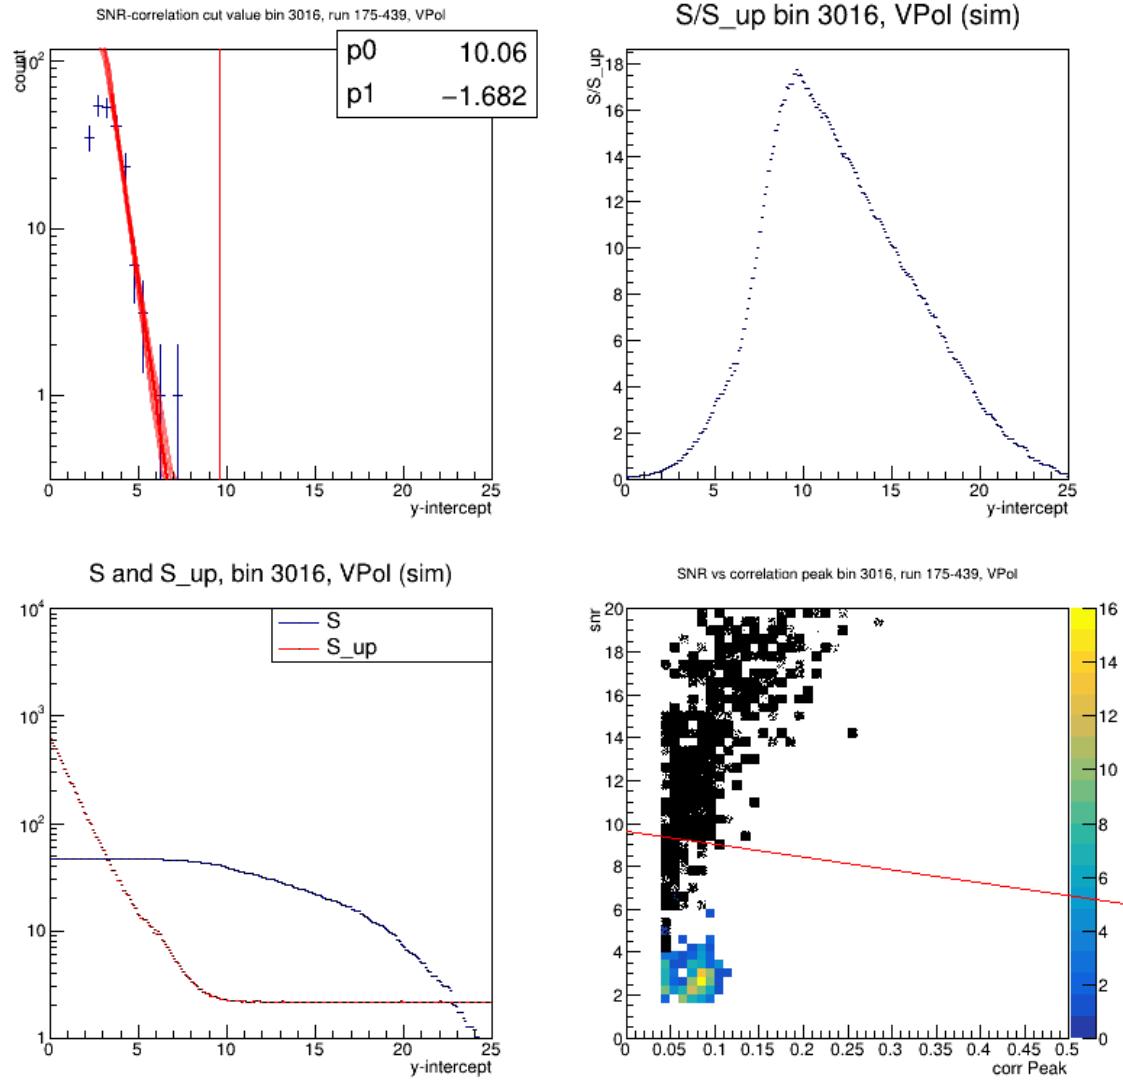


Figure D.7: Optimization plots for Healpix bin 3016 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

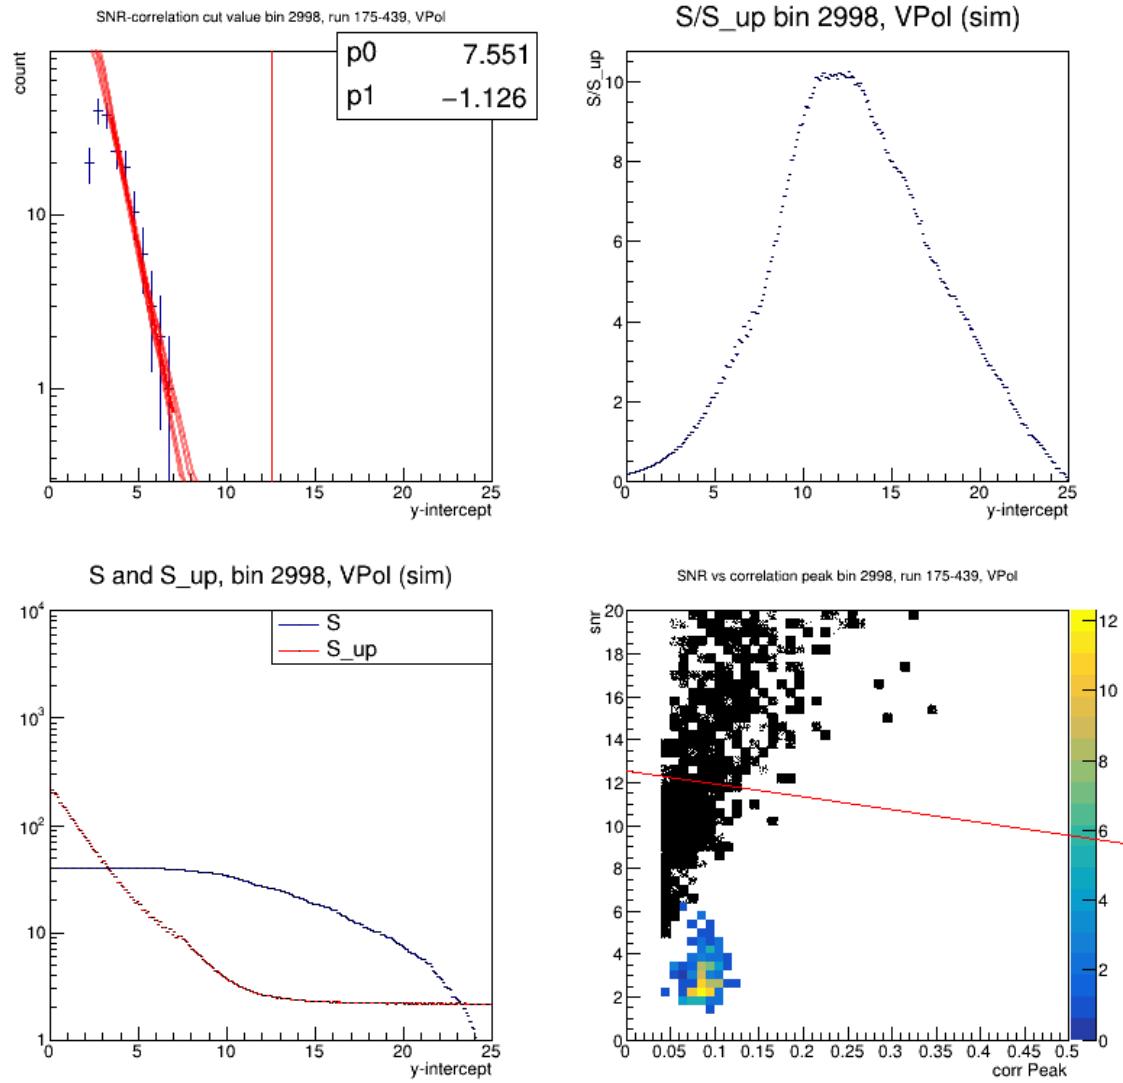


Figure D.8: Optimization plots for Healpix bin 2998 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

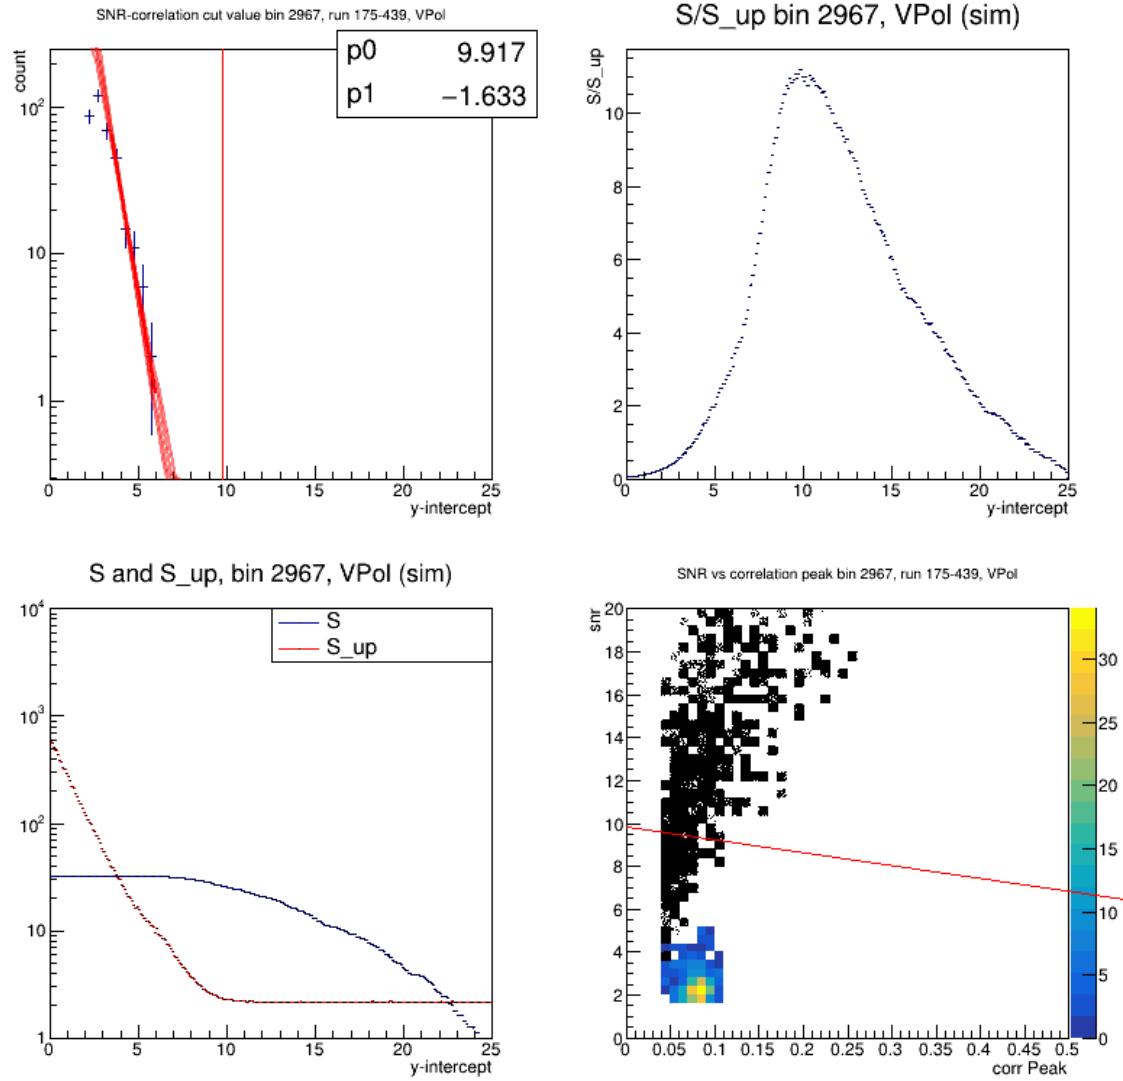


Figure D.9: Optimization plots for Healpix bin 2967 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

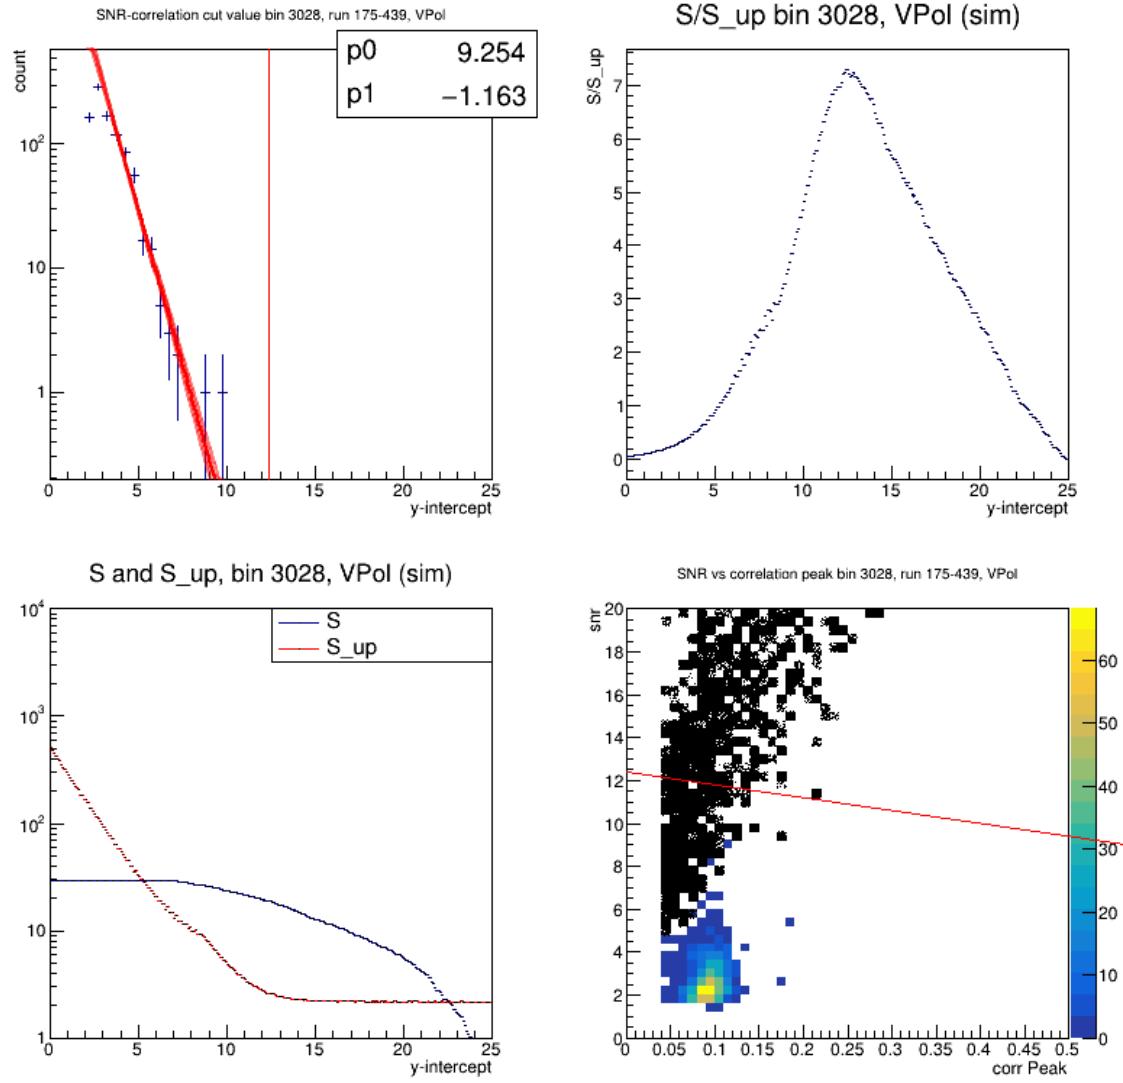


Figure D.10: Optimization plots for Healpix bin 3028 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

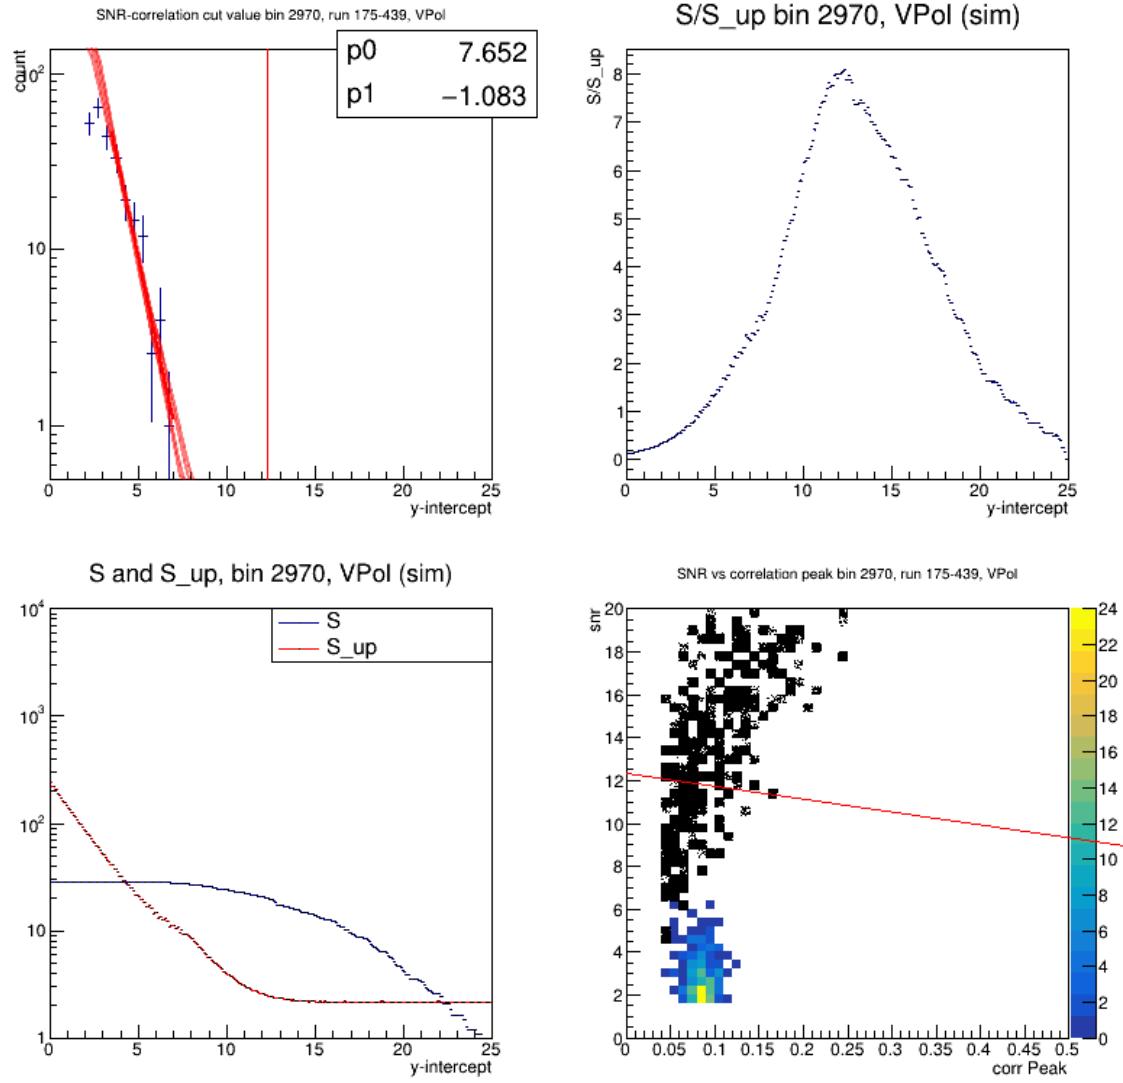


Figure D.11: Optimization plots for Healpix bin 2970 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

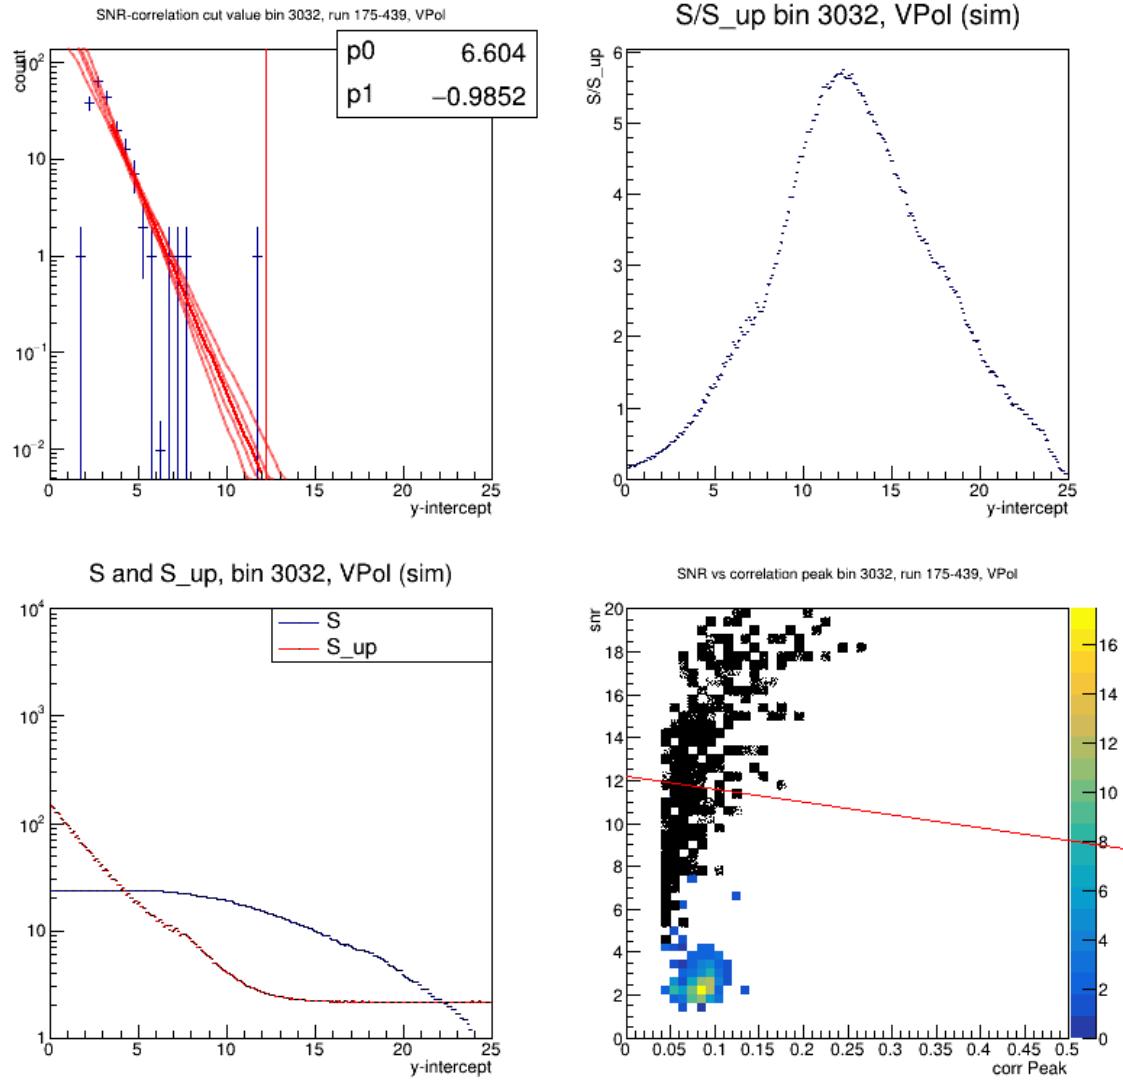


Figure D.12: Optimization plots for Healpix bin 3032 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

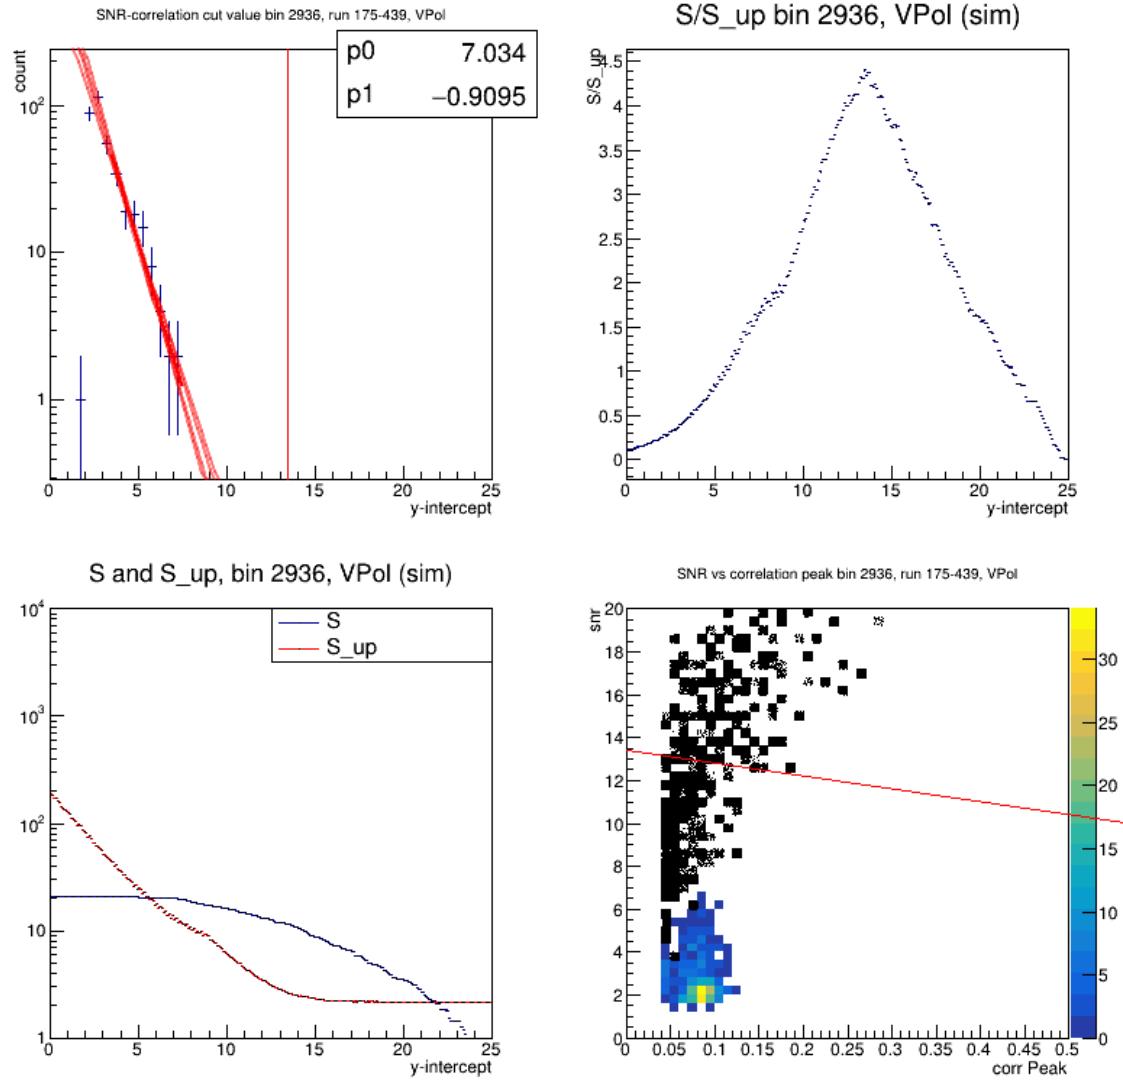


Figure D.13: Optimization plots for Healpix bin 2936 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

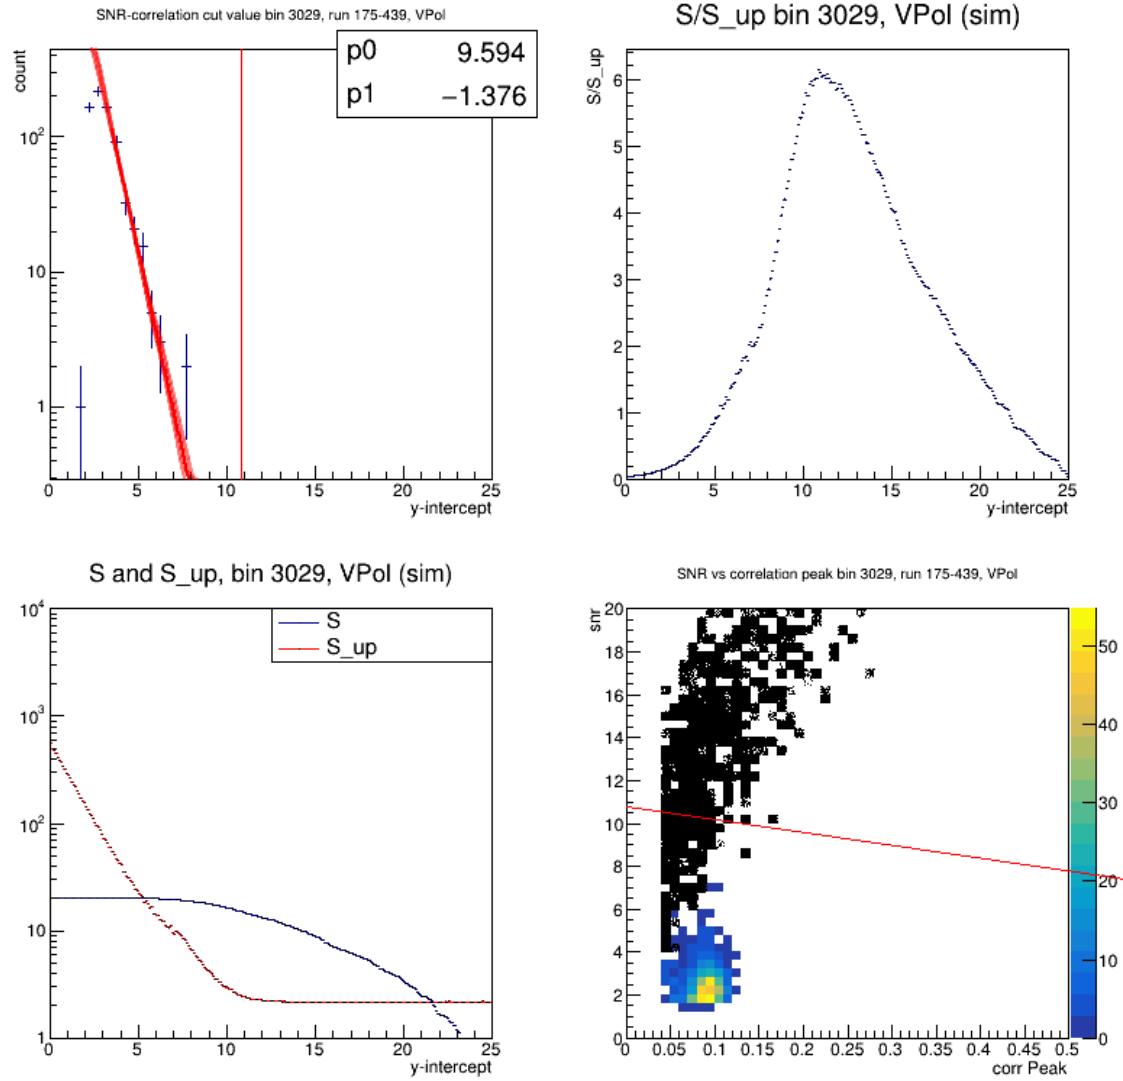


Figure D.14: Optimization plots for Healpix bin 3029 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

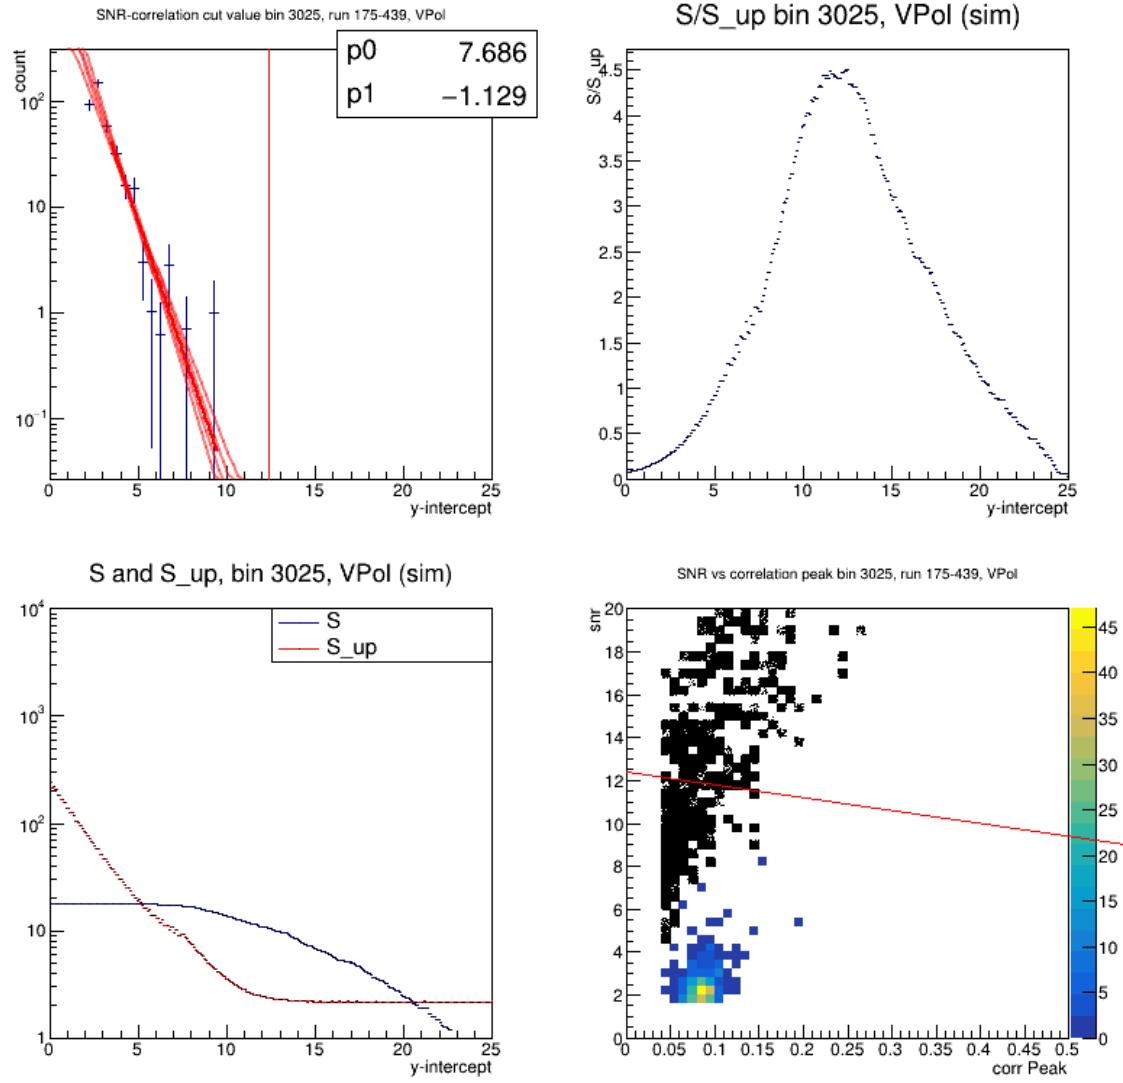


Figure D.15: Optimization plots for Healpix bin 3025 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

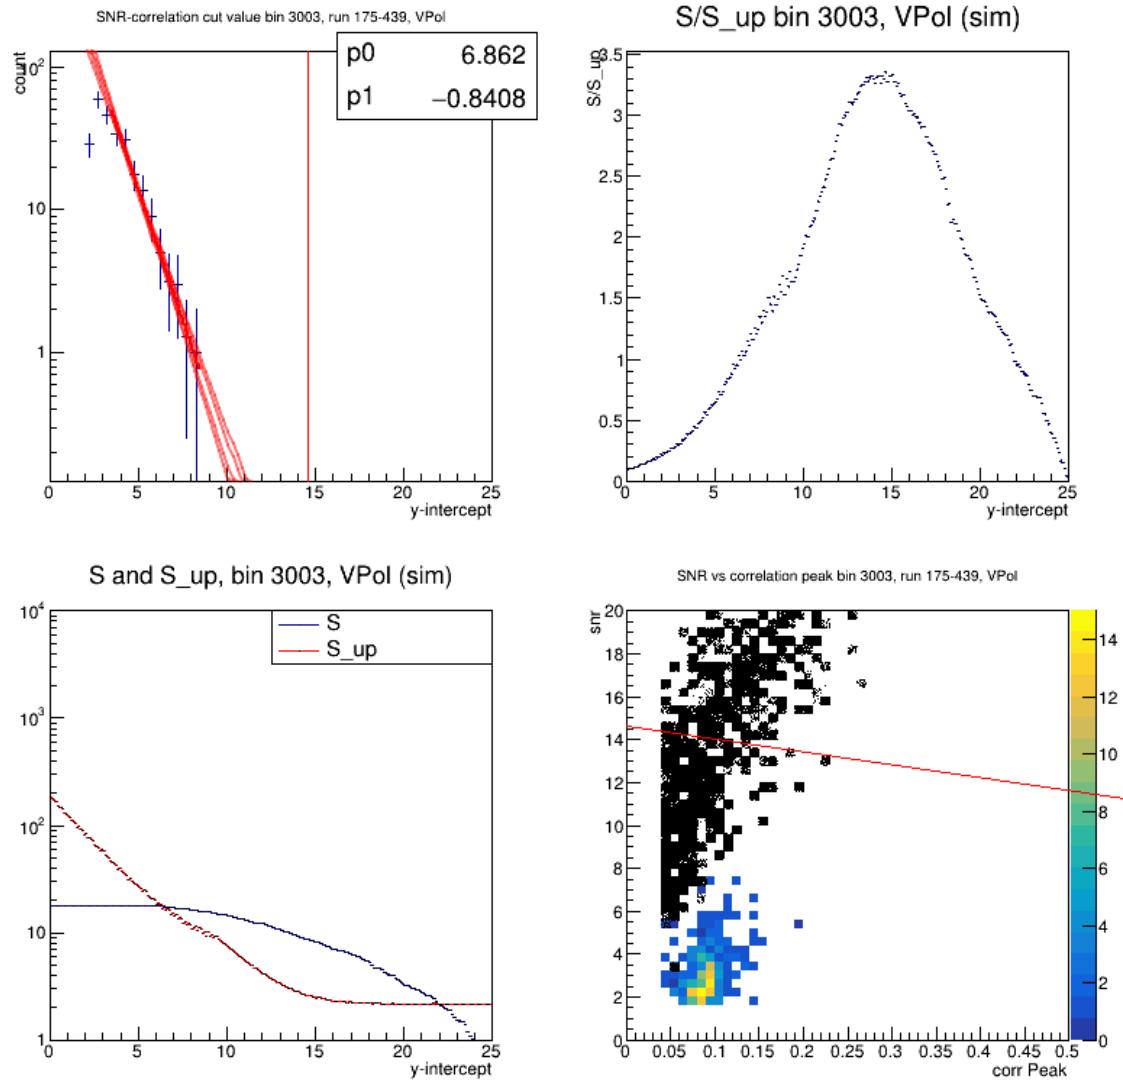


Figure D.16: Optimization plots for Healpix bin 3003 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

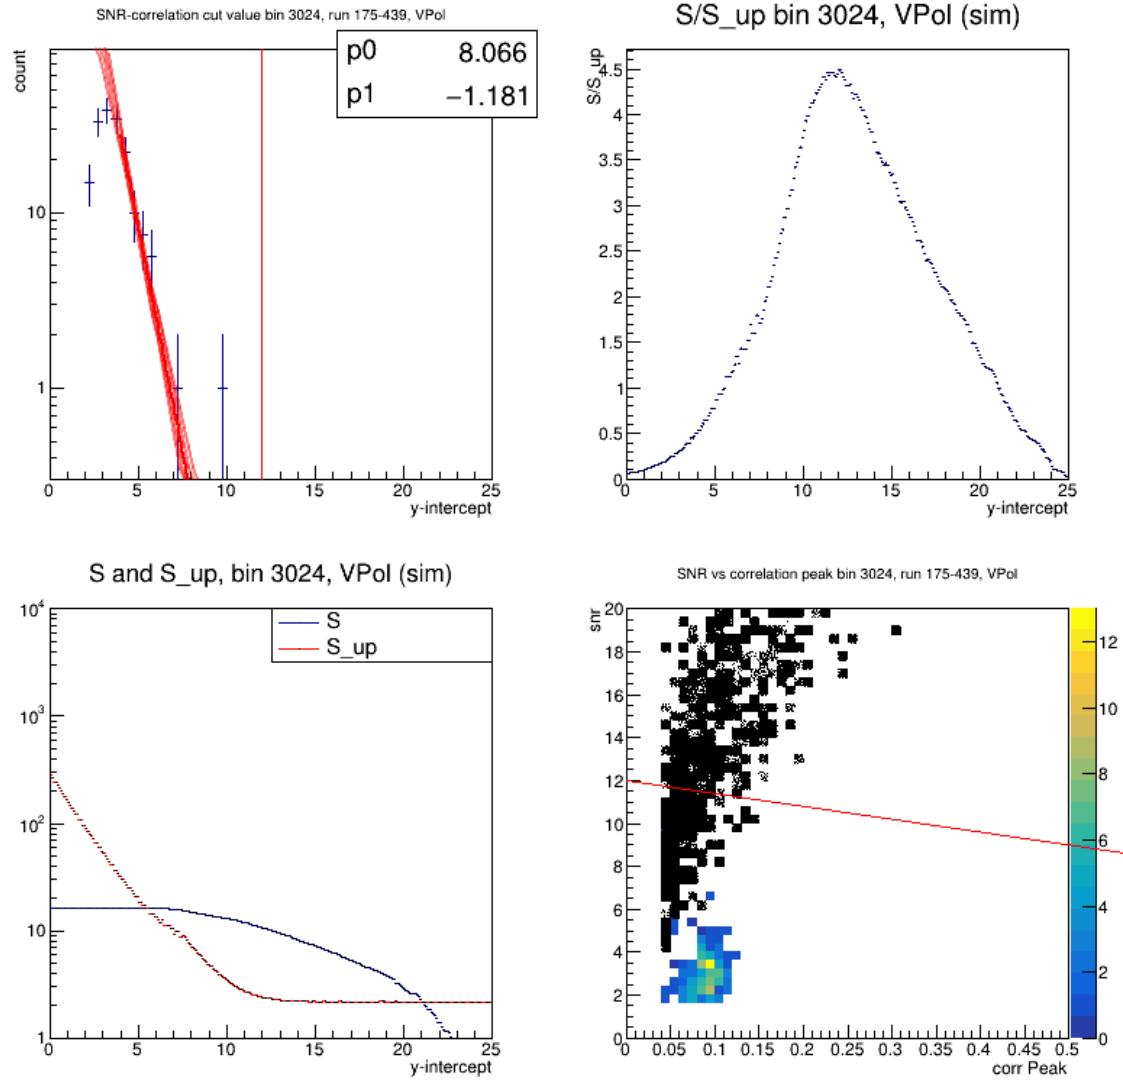


Figure D.17: Optimization plots for Healpix bin 3024 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

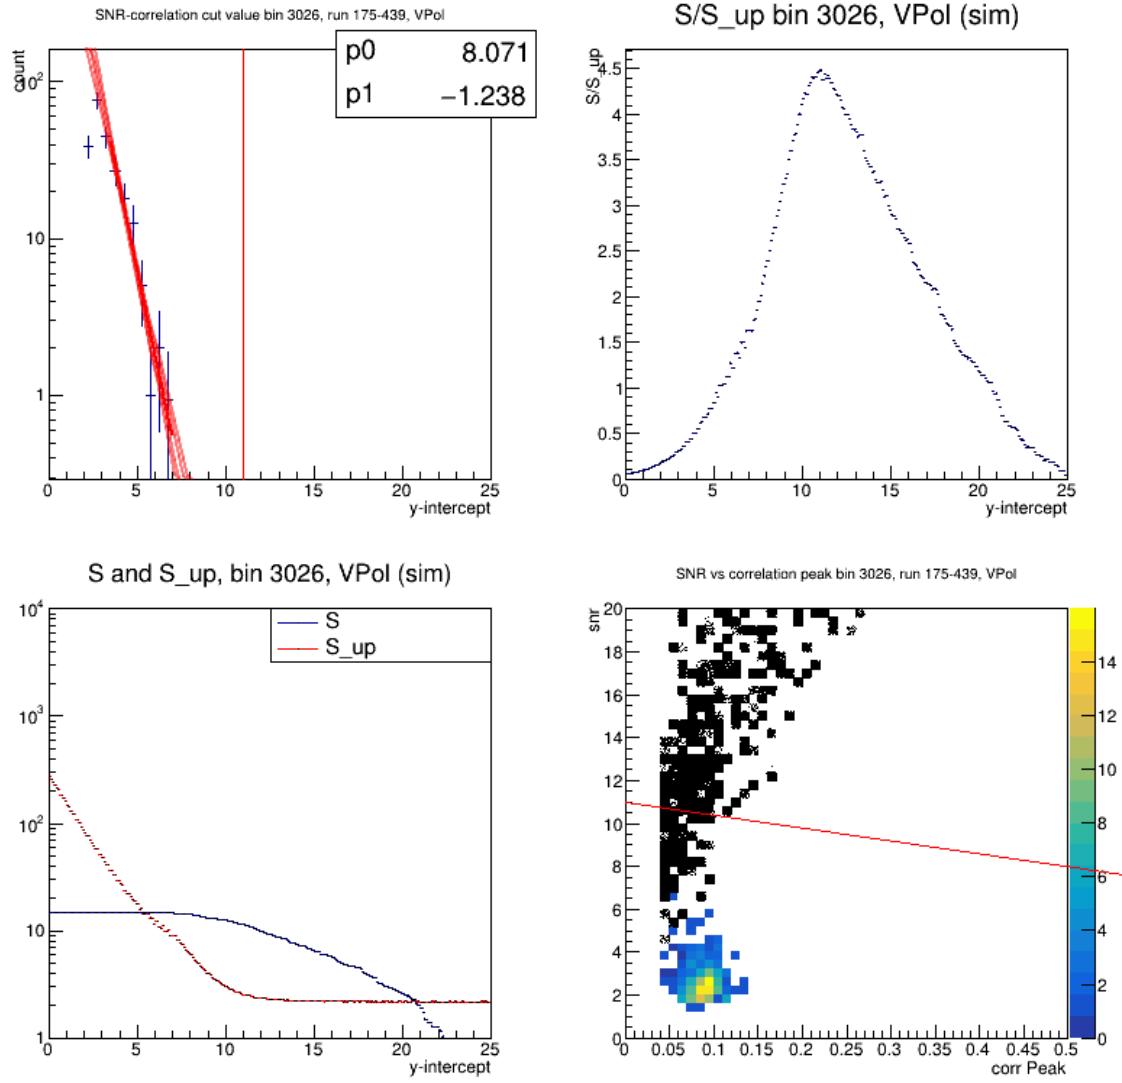


Figure D.18: Optimization plots for Healpix bin 3026 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

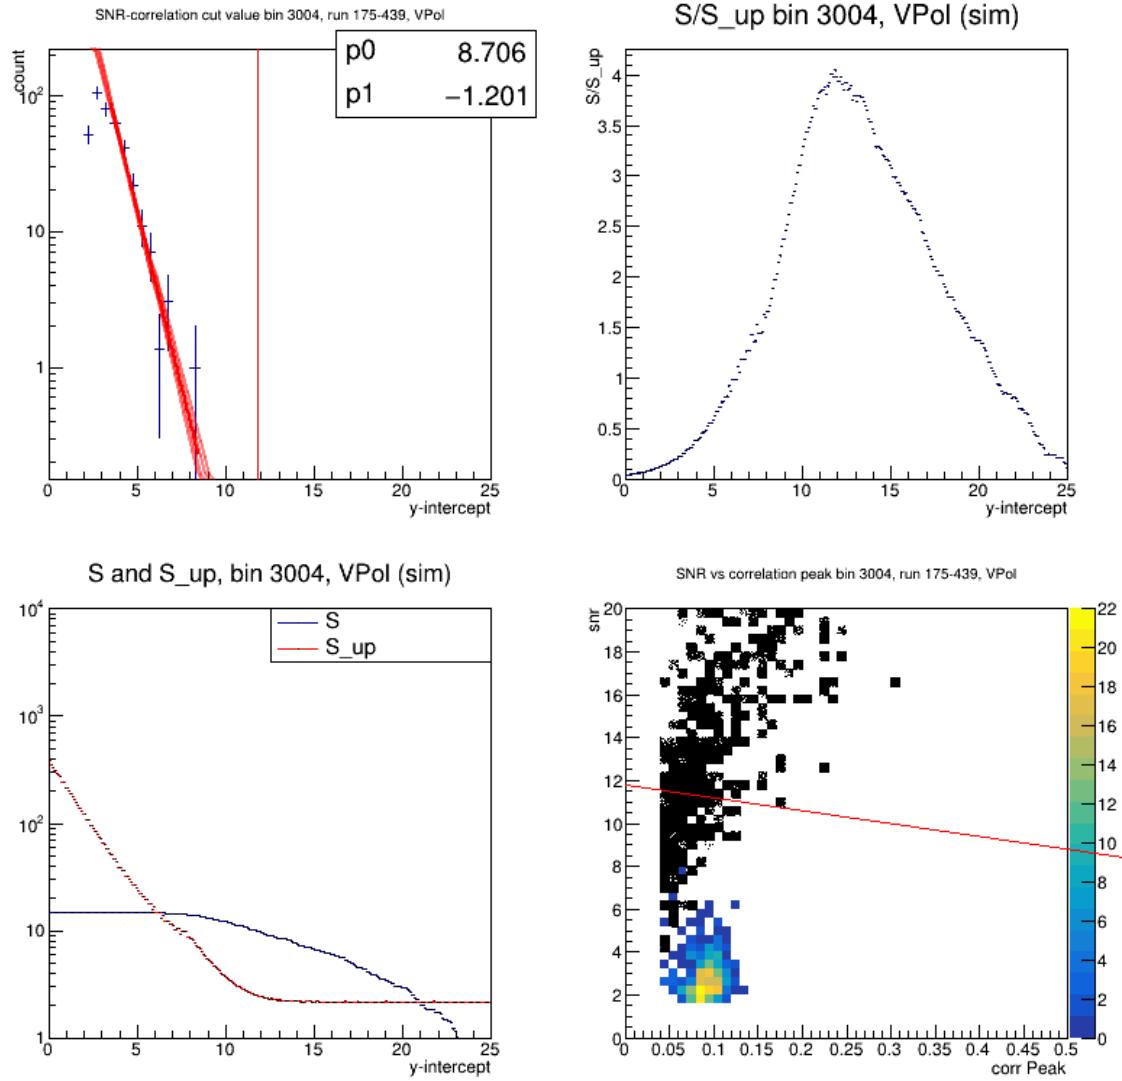


Figure D.19: Optimization plots for Healpix bin 3004 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

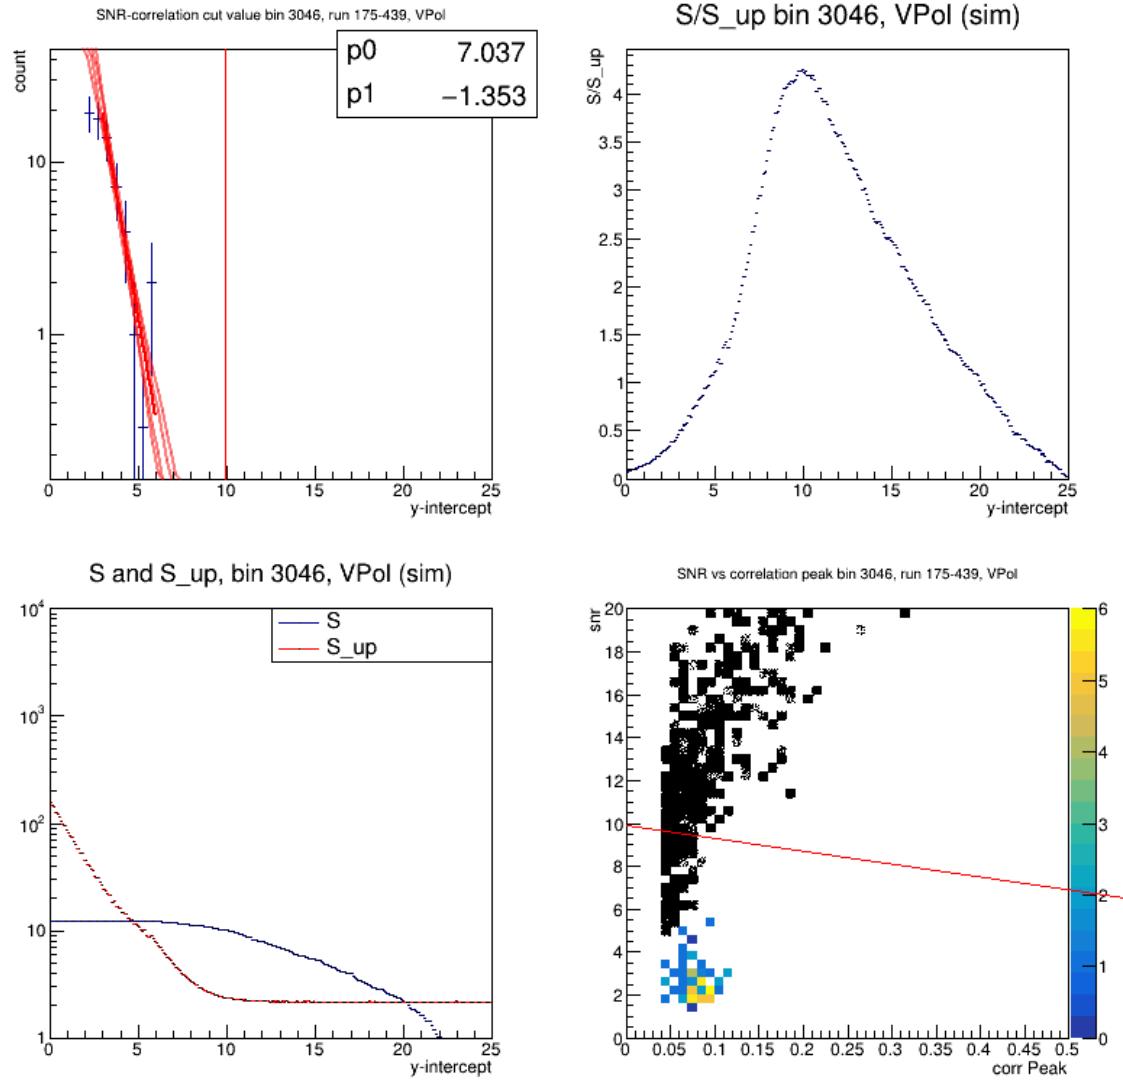


Figure D.20: Optimization plots for Healpix bin 3046 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

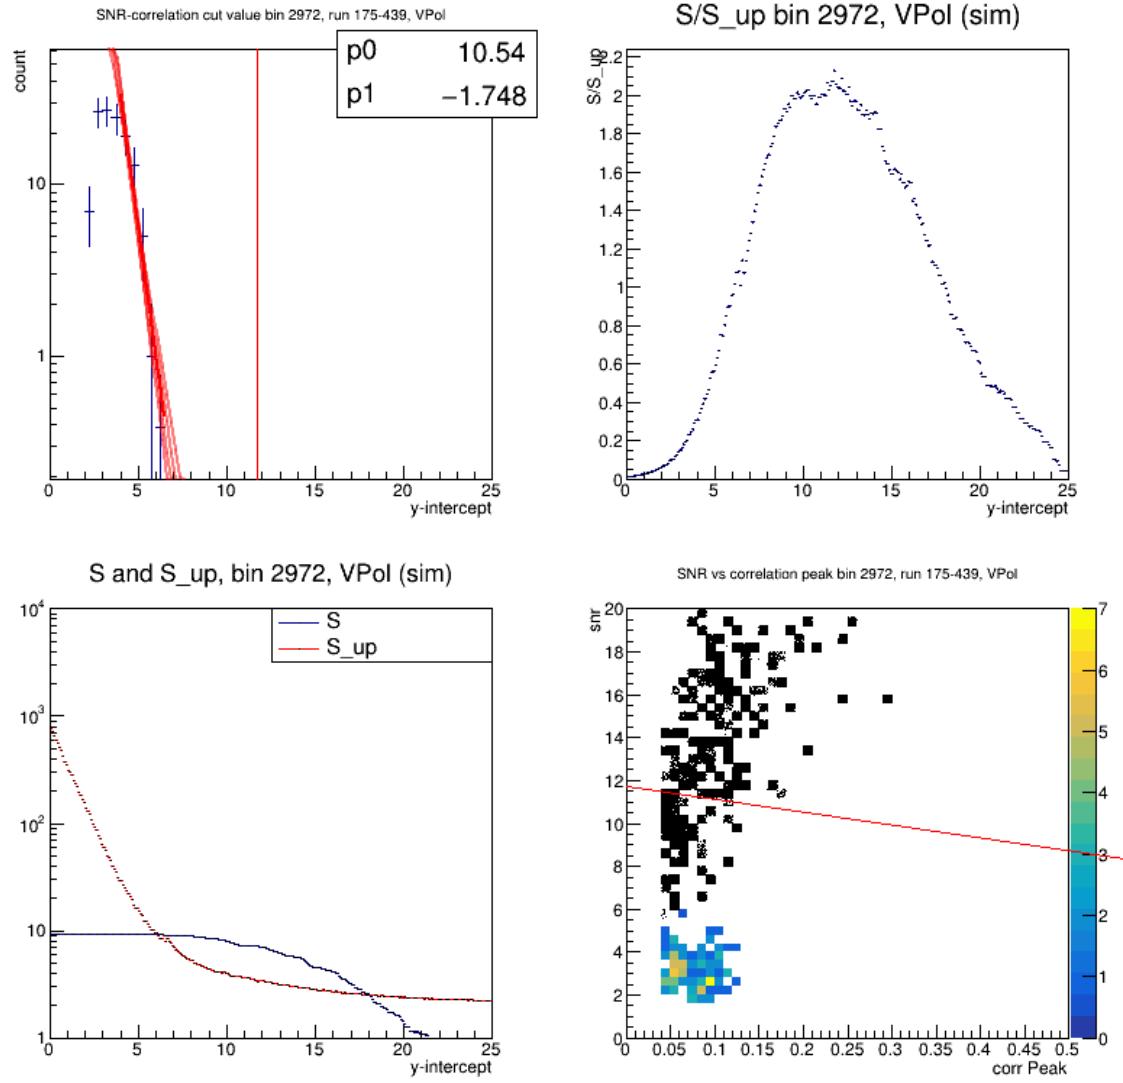


Figure D.21: Optimization plots for Healpix bin 2972 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

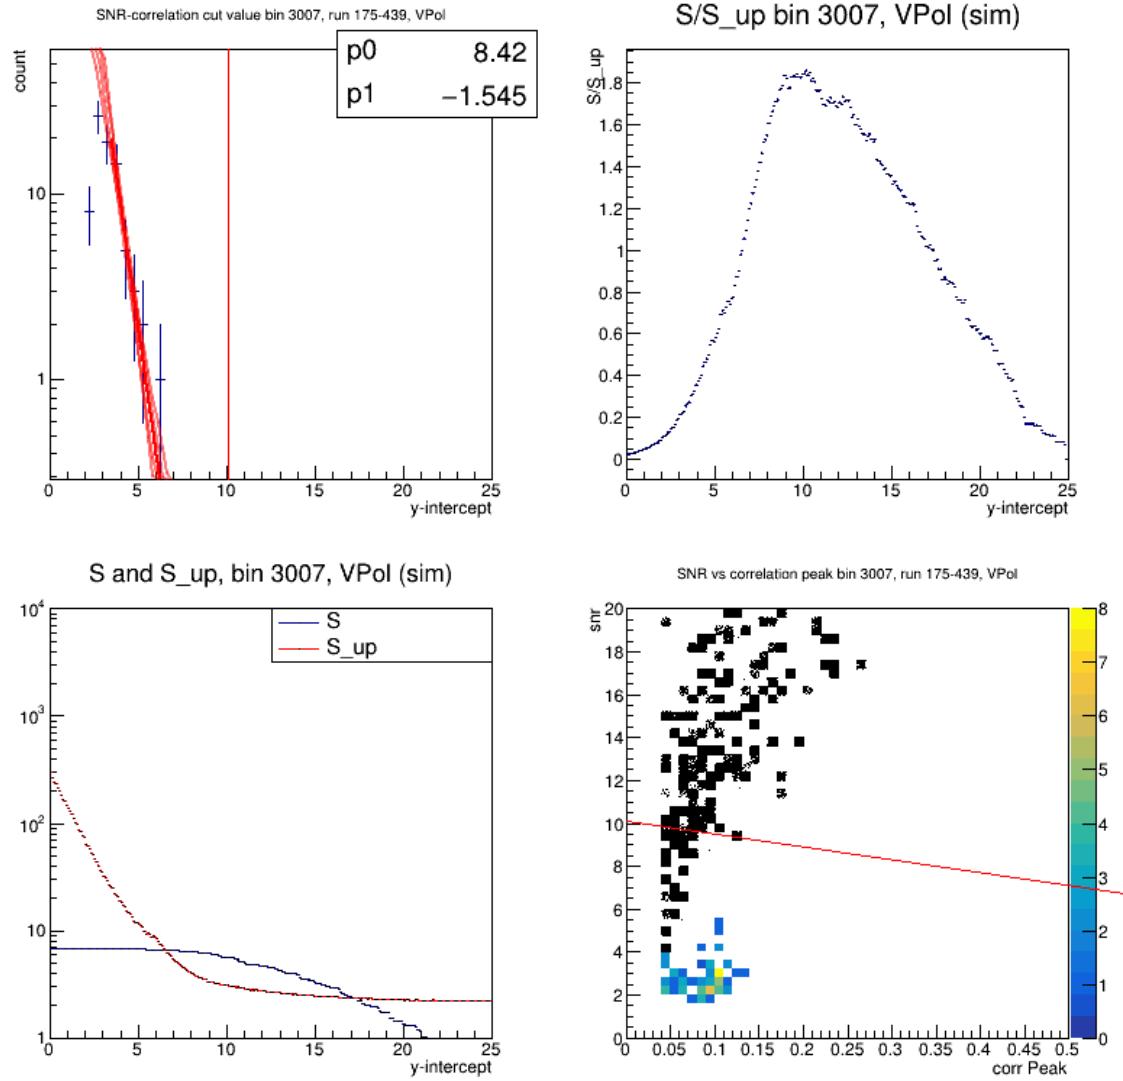


Figure D.22: Optimization plots for Healpix bin 3007 in V-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

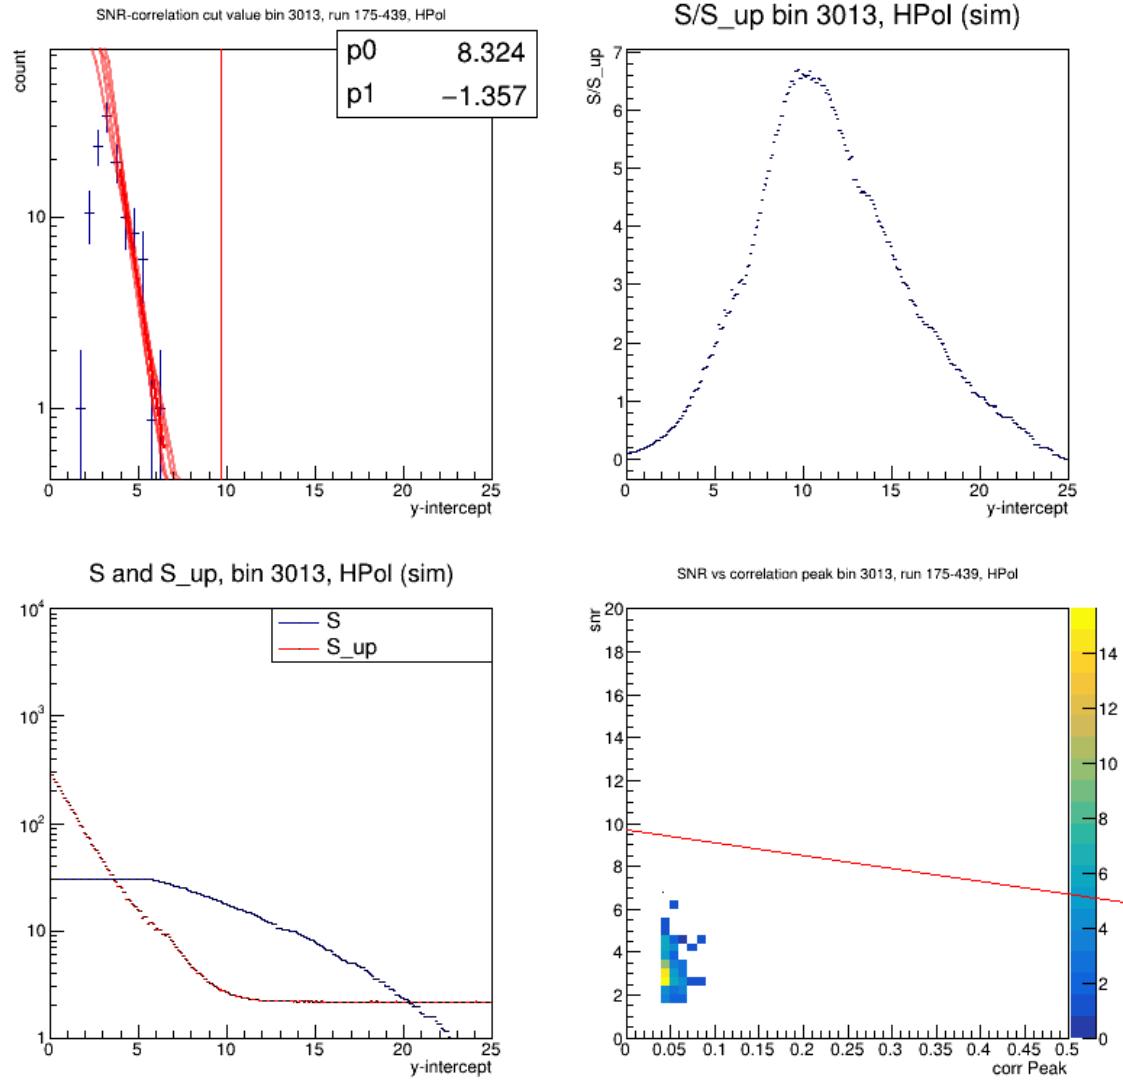


Figure D.23: Optimization plots for Healpix bin 3013 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

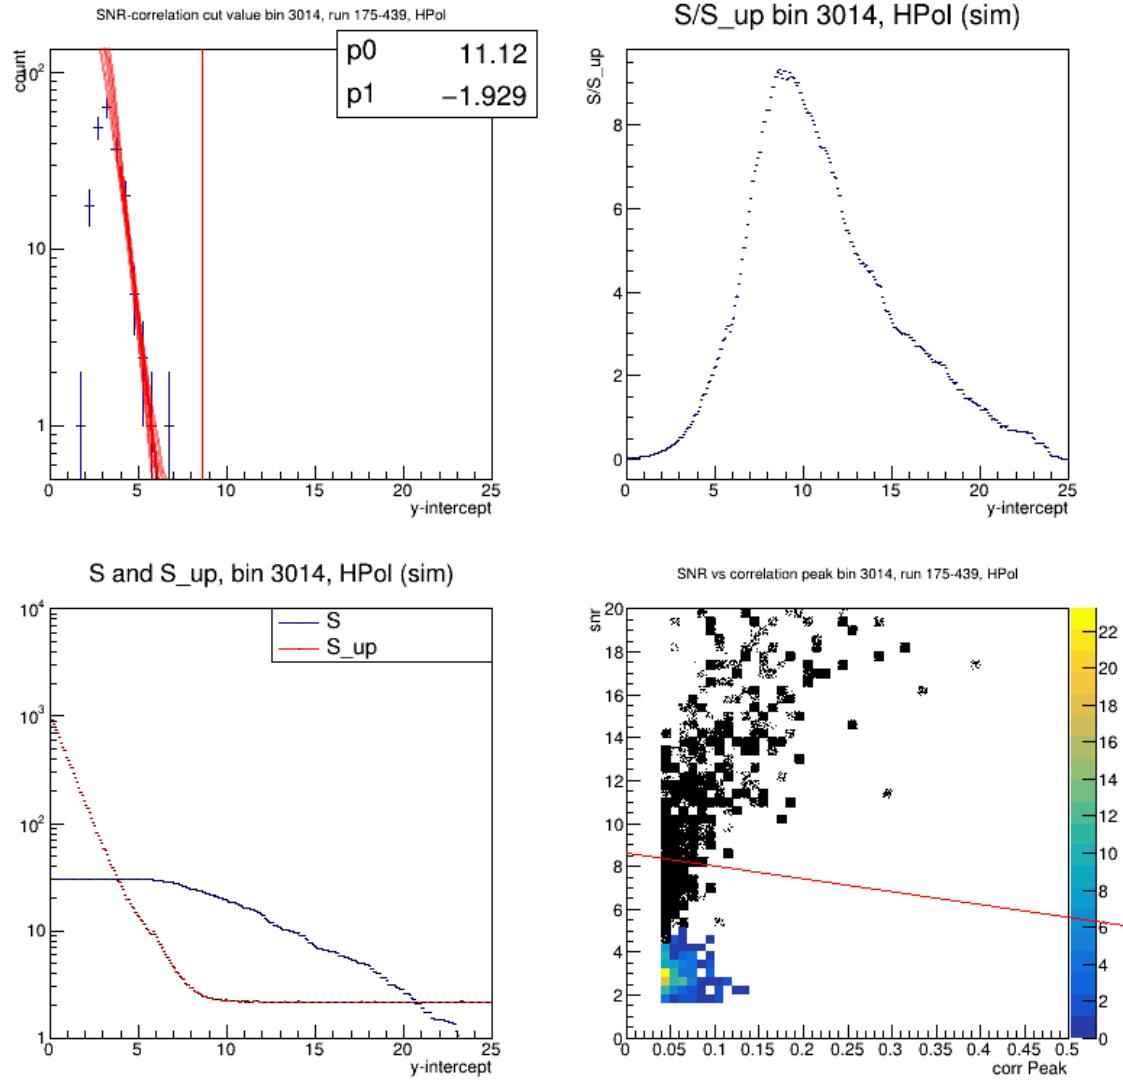


Figure D.24: Optimization plots for Healpix bin 3014 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

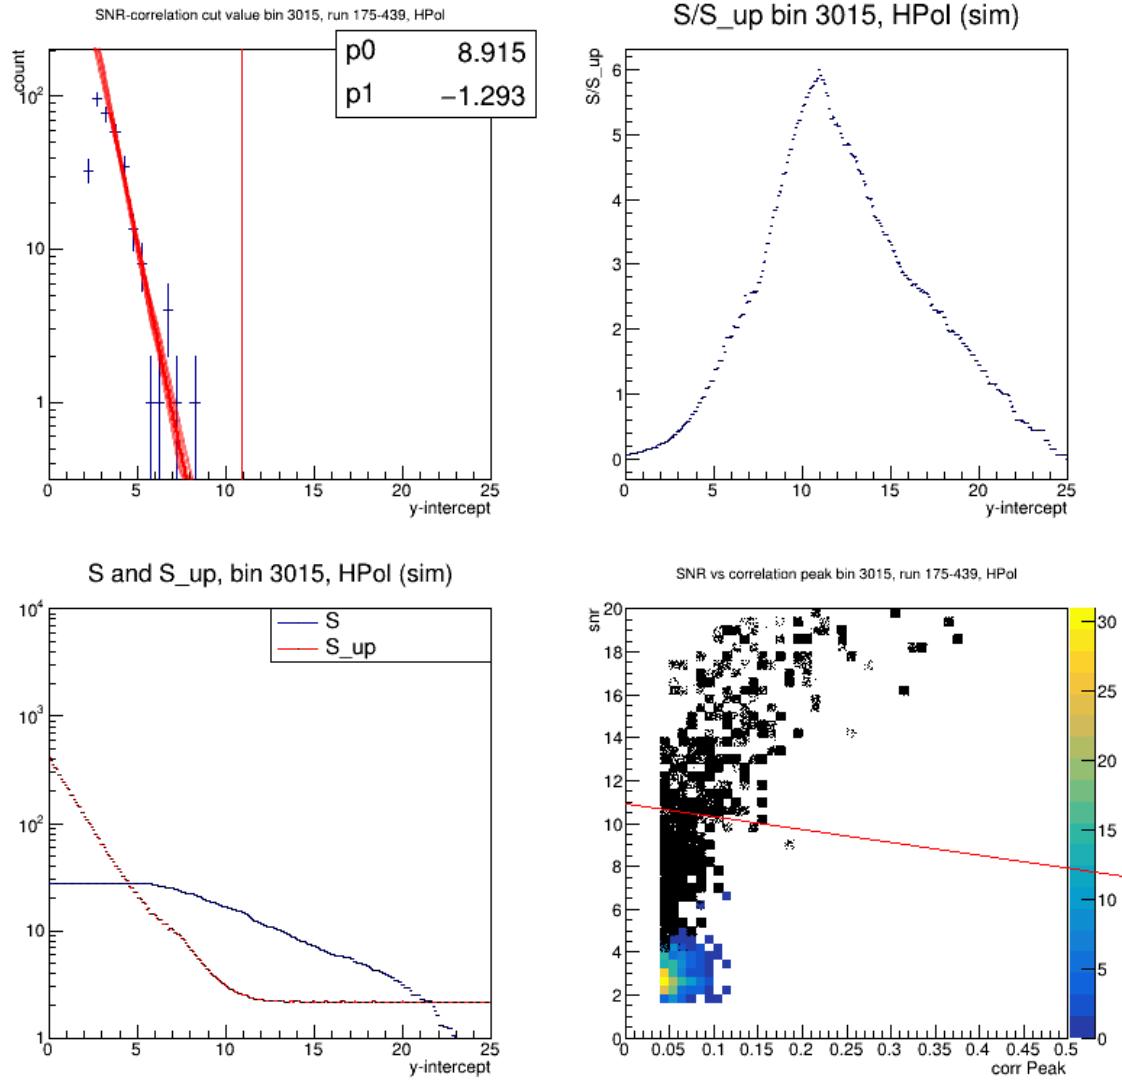


Figure D.25: Optimization plots for Healpix bin 3015 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

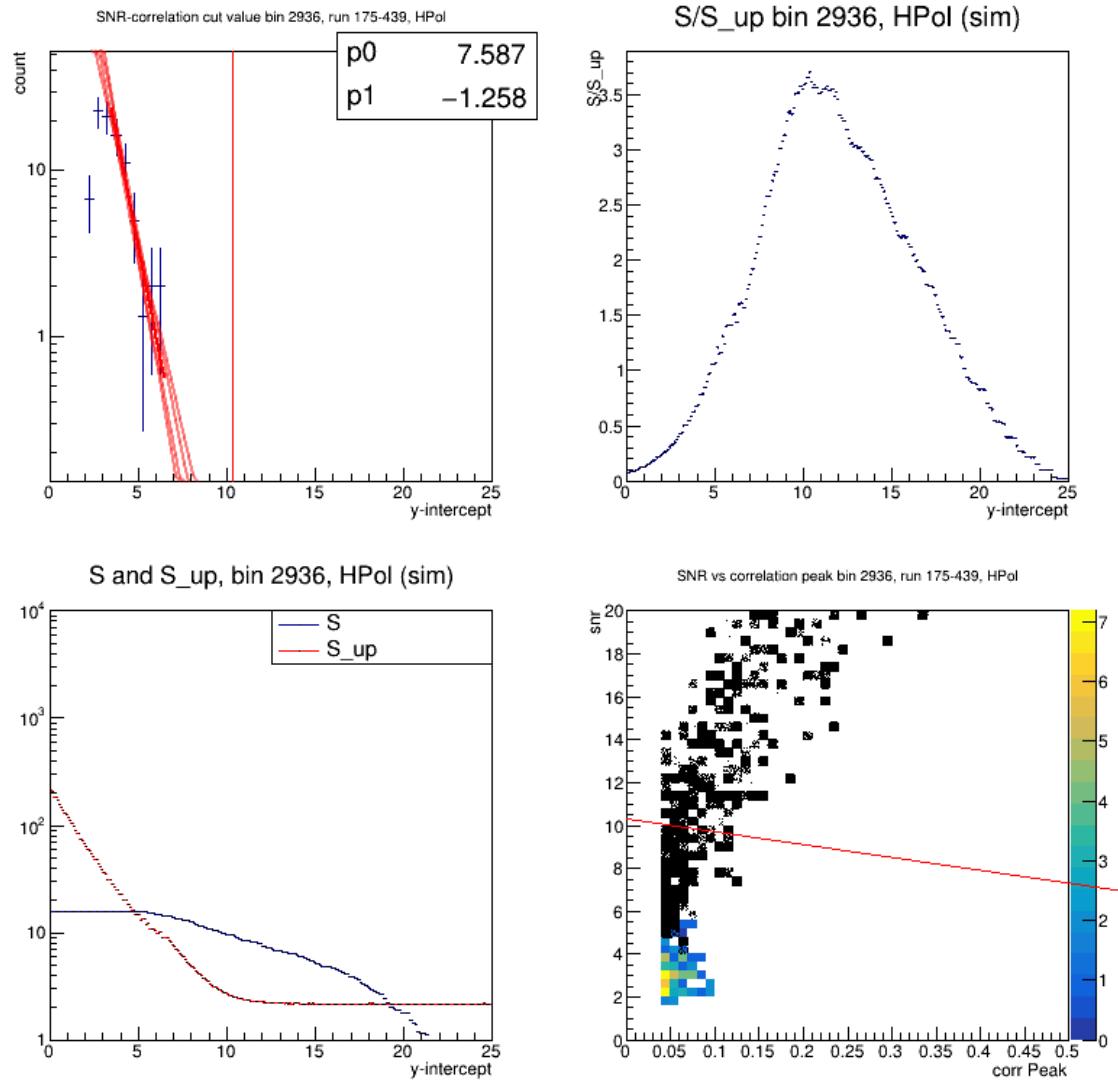


Figure D.26: Optimization plots for Healpix bin 2936 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

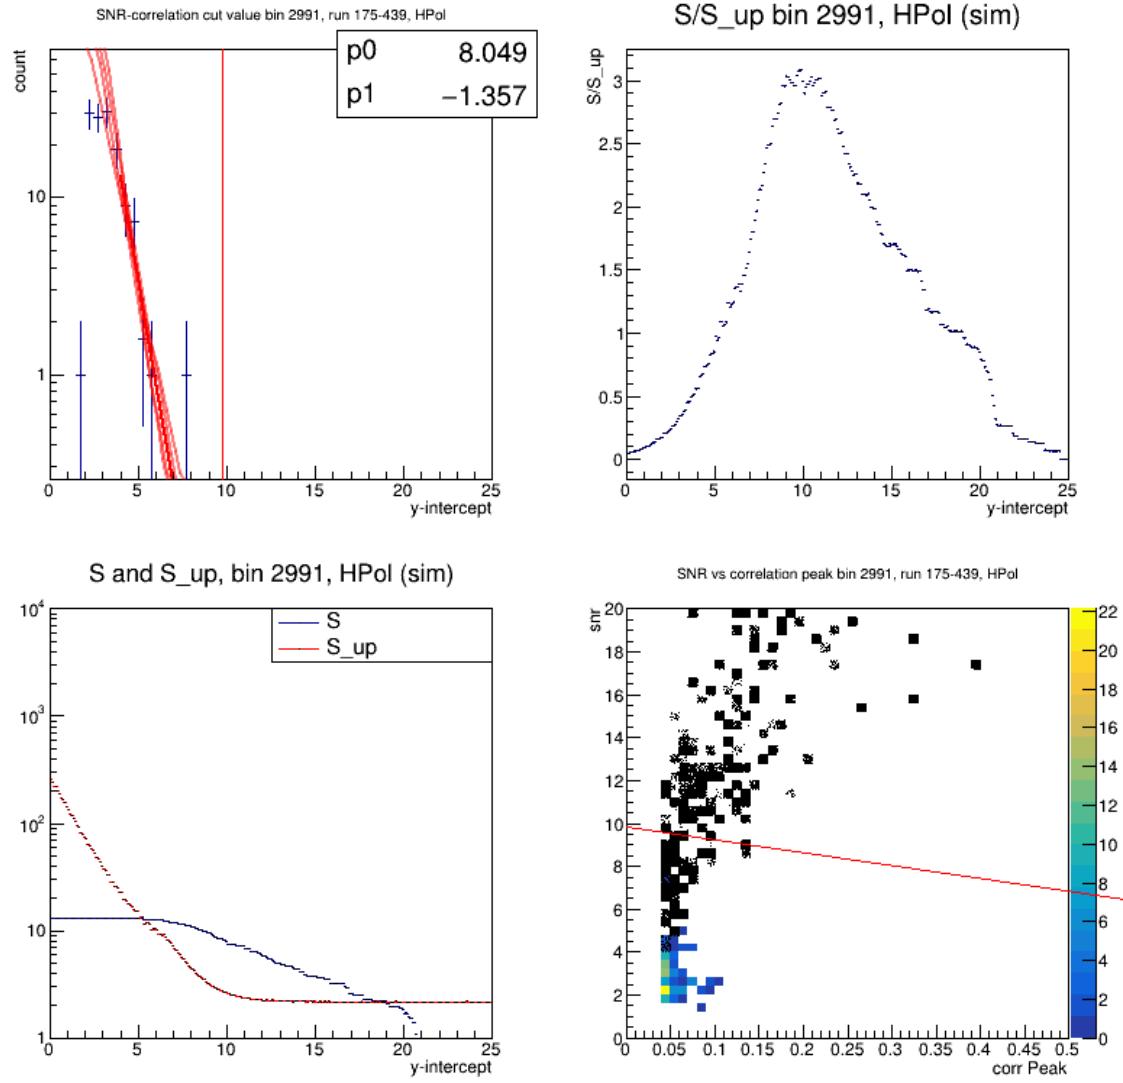


Figure D.27: Optimization plots for Healpix bin 2991 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

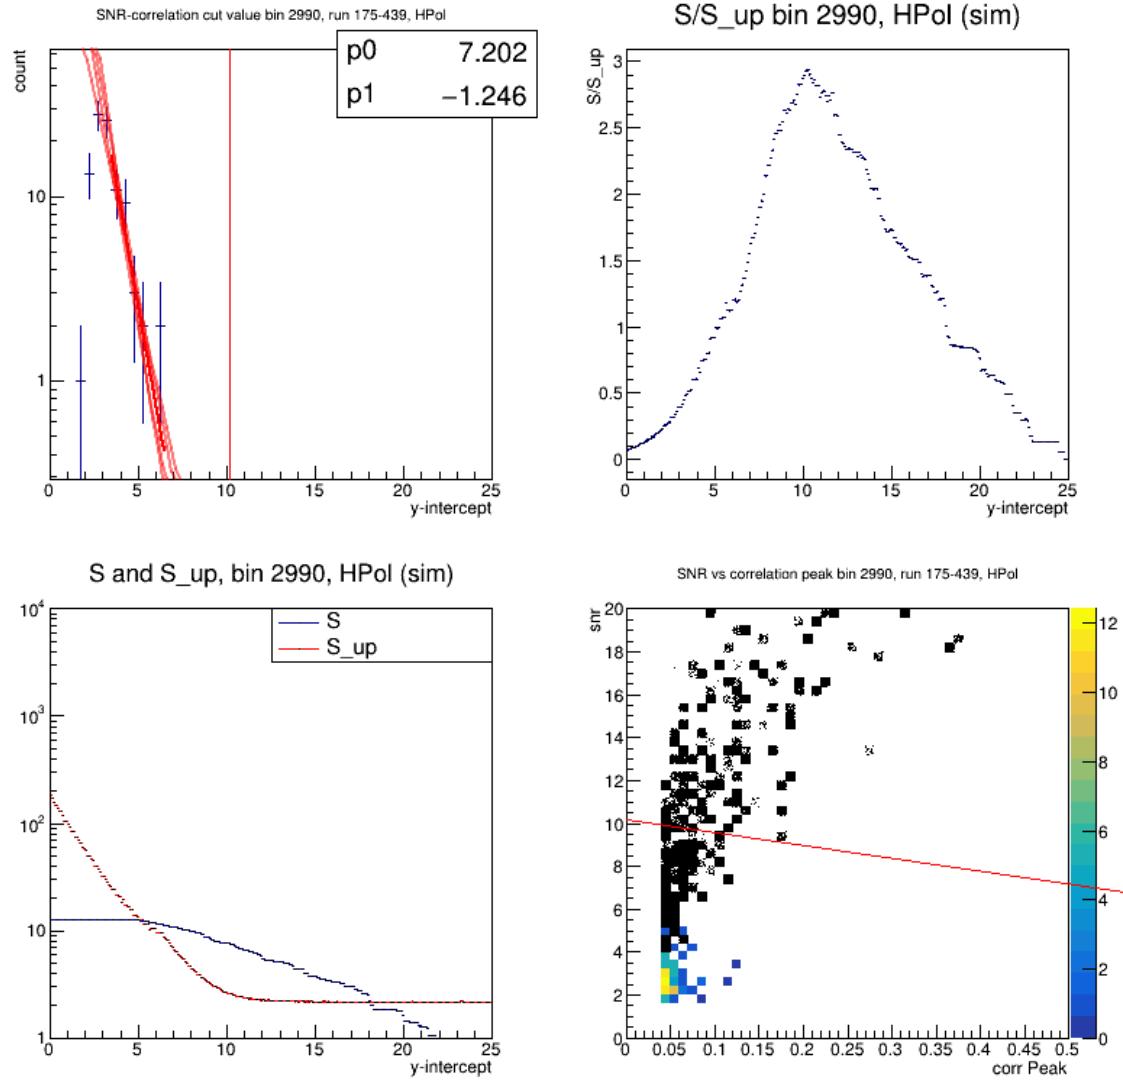


Figure D.28: Optimization plots for Healpix bin 2990 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

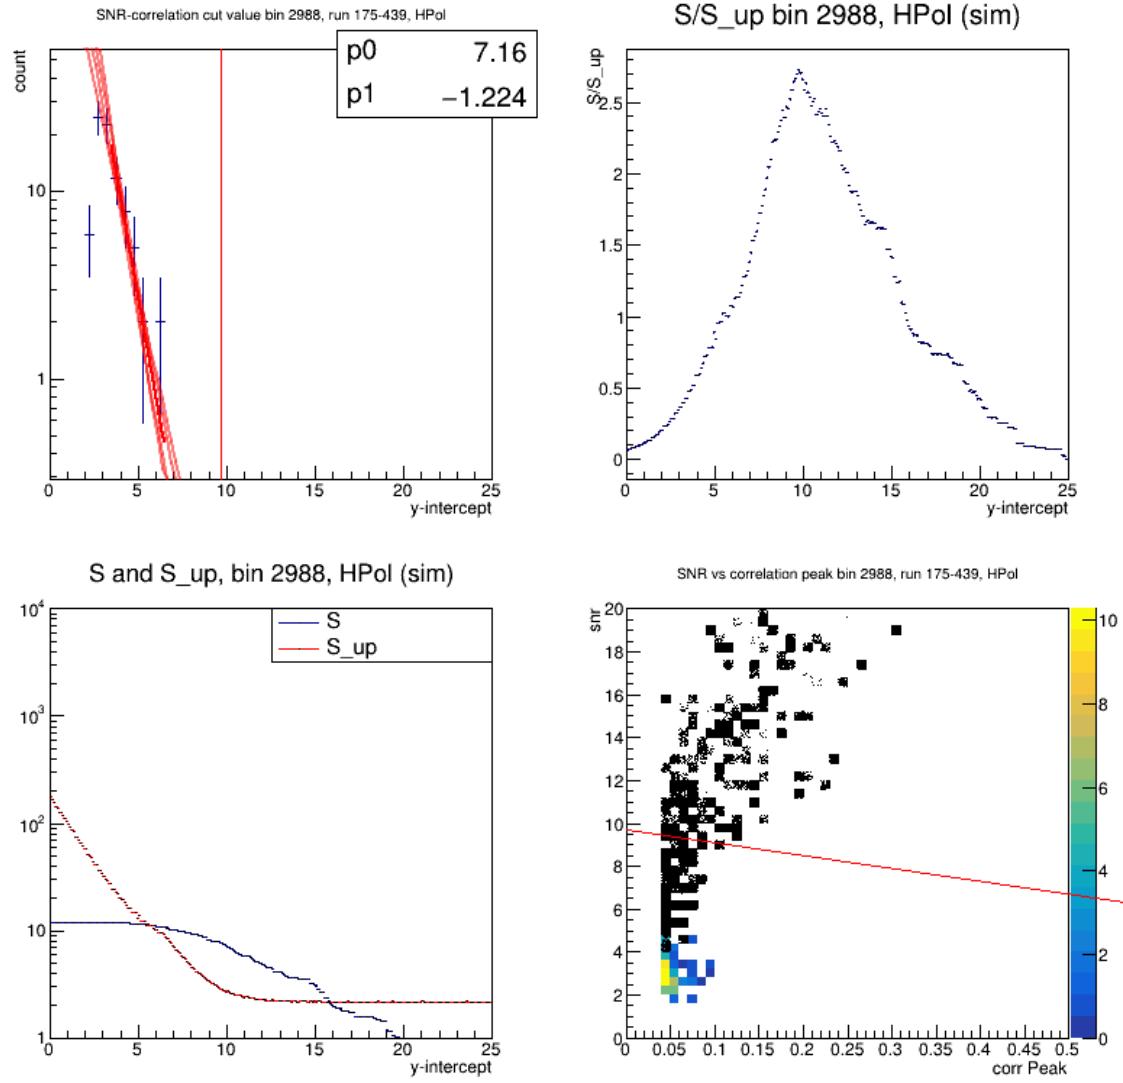


Figure D.29: Optimization plots for Healpix bin 2988 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

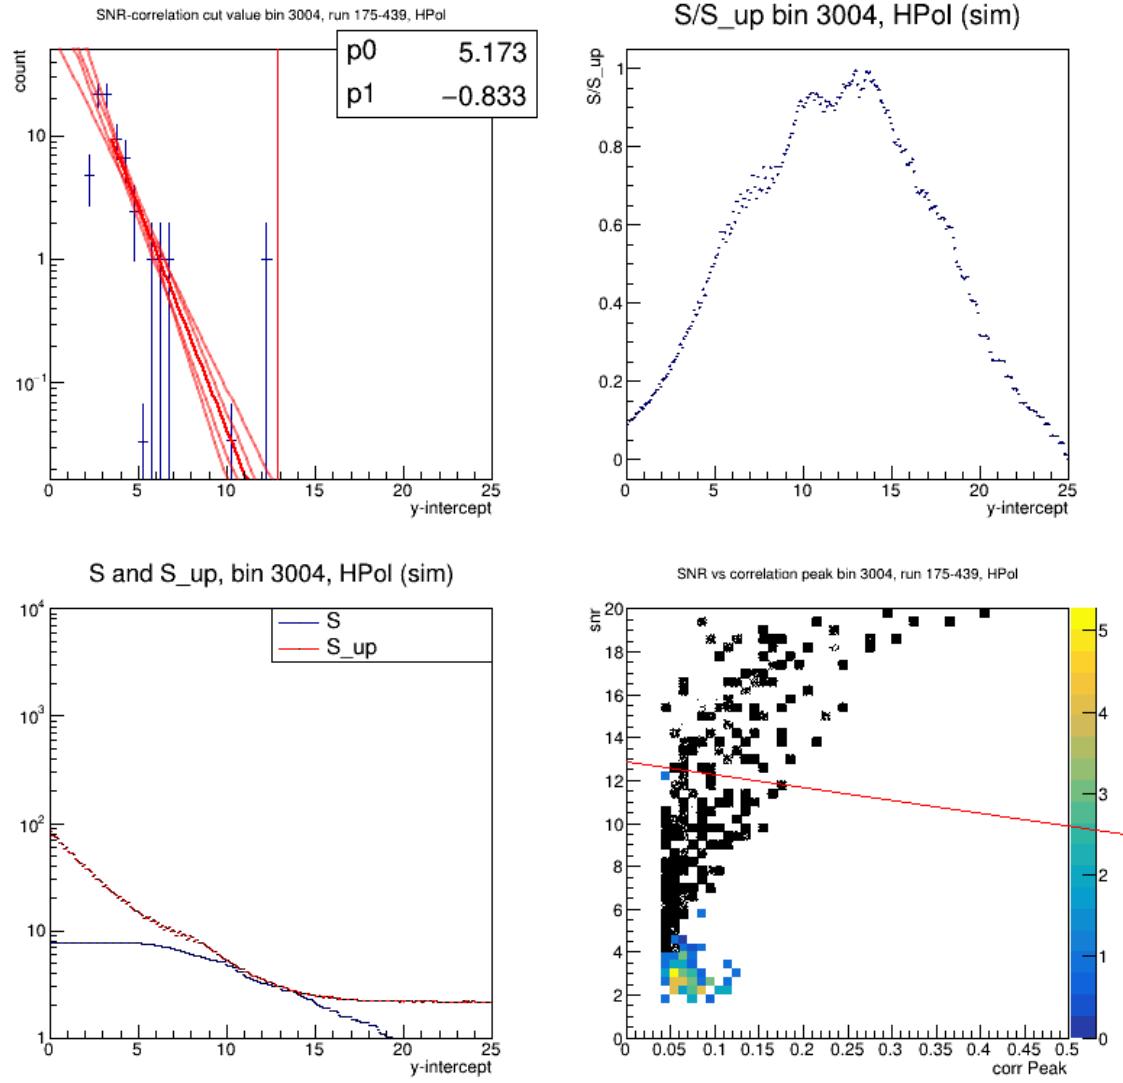


Figure D.30: Optimization plots for Healpix bin 3004 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

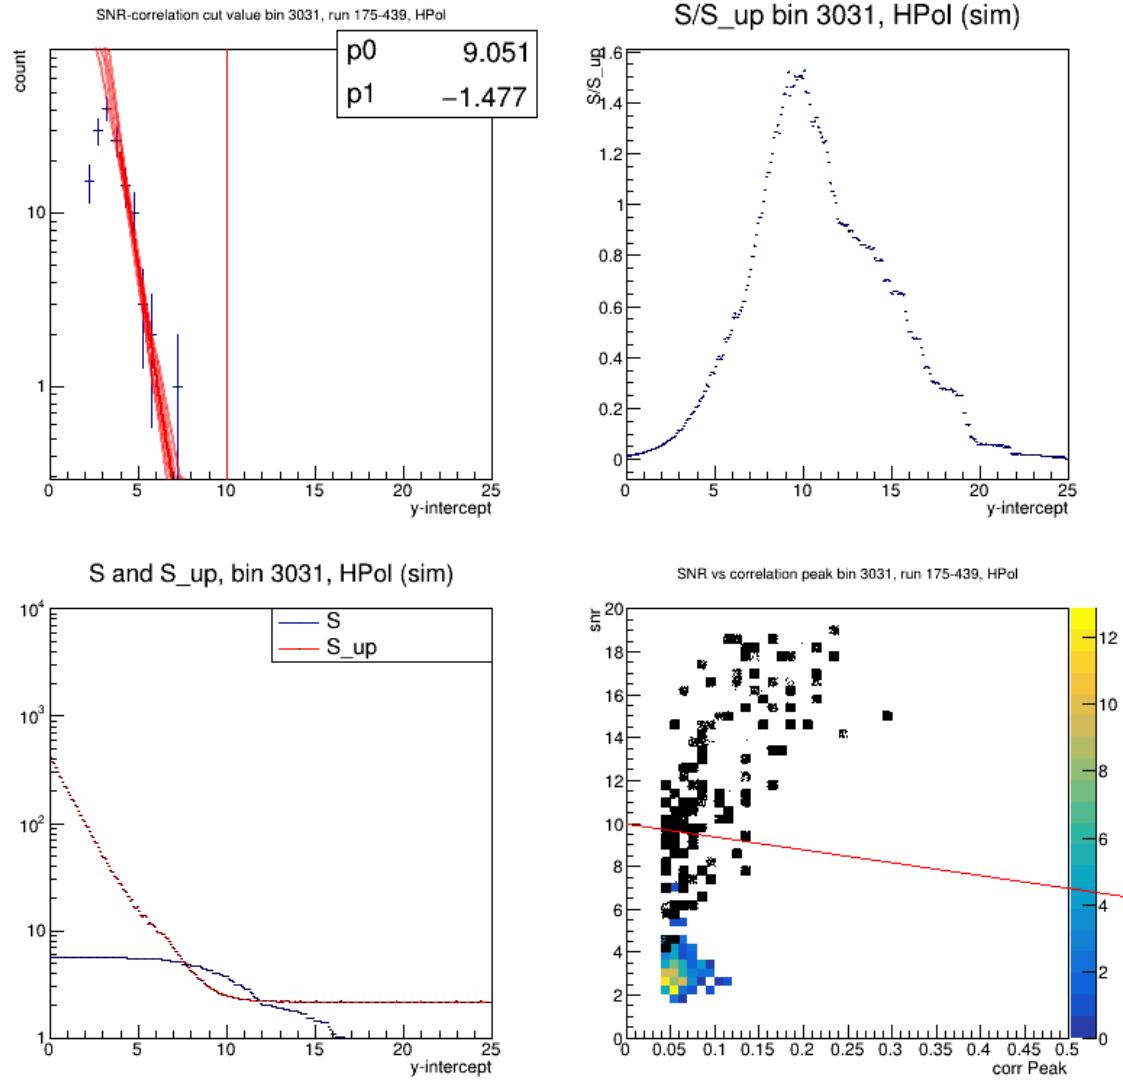


Figure D.31: Optimization plots for Healpix bin 3031 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

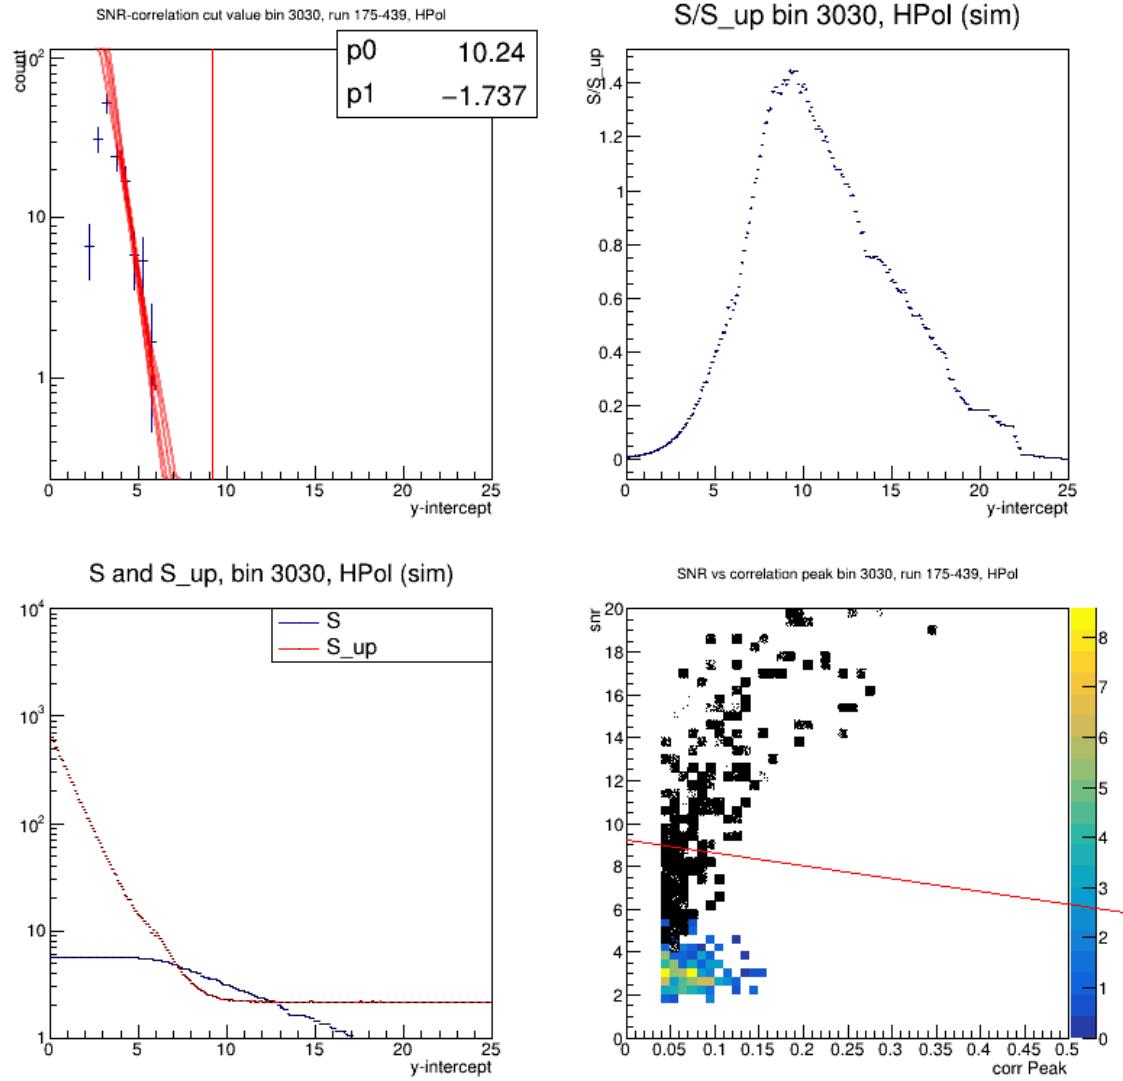


Figure D.32: Optimization plots for Healpix bin 3030 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

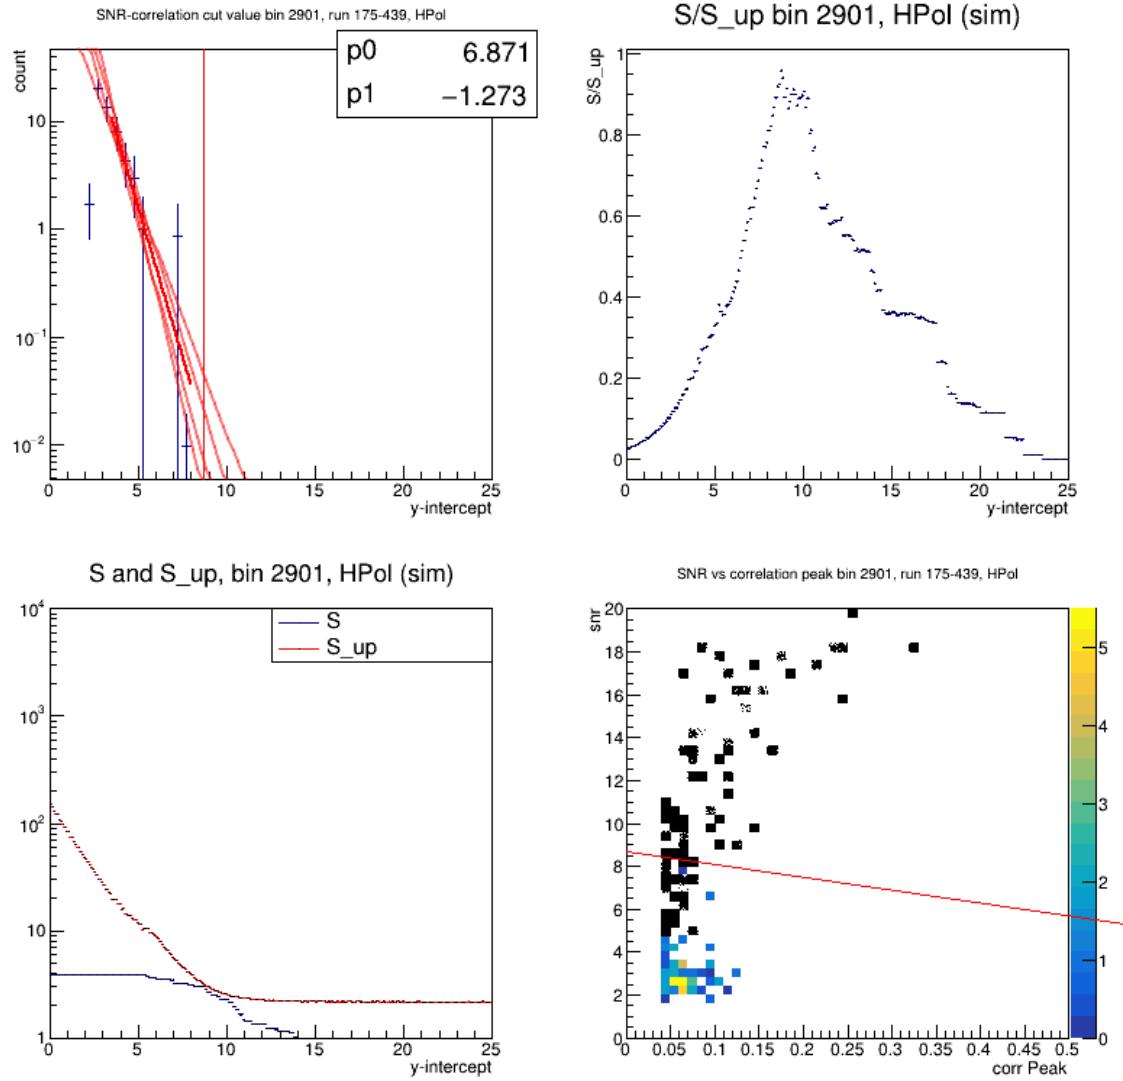


Figure D.33: Optimization plots for Healpix bin 2901 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminant cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.

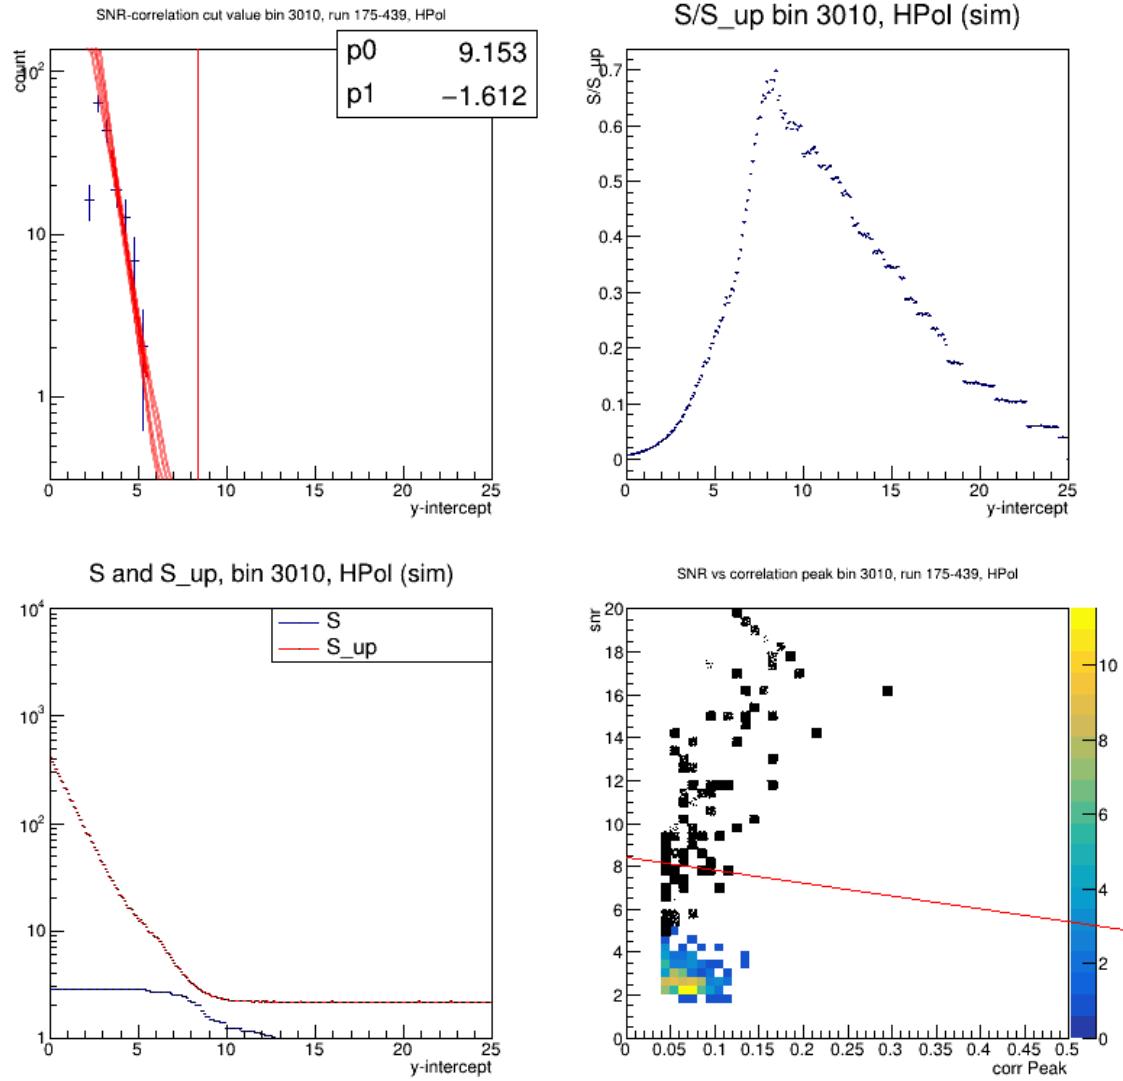


Figure D.34: Optimization plots for Healpix bin 3010 in H-pol. Top Left: Histogram of linear combination data fit by an exponential (in red). Red vertical line is optimized y-intercept of the LD cut. Bottom Left: Plot of passing (unnormalized) simulated neutrino signal (in blue) and S_{up} in red. Top Right: Plot of s/s_{up} . Bottom Right: Plot of correlation peak vs SNR. Linear discriminate cut plotted on top in Red. Black data points are simulated neutrinos, while the colored data points are noise.