A Search for the Origin of Ultra-High Energy Neutrinos with ANITA-4

Luke Batten

A dissertation submitted in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

of

University College London.

Department of Physics and Astronomy

University College London

June 18, 2021

I, Luke Batten, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.

Abstract

The ANtarctic Impulsive Transient Antenna (ANITA) aims to detect Askaryan radio signatures from ultra-high energy (UHE) neutrinos interacting with the Antarctic ice sheet. However, the origin of neutrinos of such high energies is mostly unknown.

This thesis contains a method for finding the potential astrophysical origin of UHE neutrinos, and includes the simulation and analysis used throughout the source search. The thesis also discusses the essential physics, the components of the experiment, and the details of the fourth flight of the ANITA experiment, ANITA-4. Though the source search itself was conducted for this specific flight, the methods presented can be applied to any previous or future flights. As such experiments are not yet sensitive enough to obtain a large dataset of UHE neutrinos, the sub-threshold events detected by ANITA-4 were projected back to the sky. The spatiotemporal proximity of such events were compared to the activity of several objects. Three classes of objects were studied: blazars, gamma-ray bursts (GRBs), and supernovae (SNe). A single object, the flaring blazar known as PKS 1502+106, showed an excess of events pointing back to it, exceeding the 99% confidence level threshold. Such searches and analyses are not yet sensitive enough to detect a high number of ultra-high energy neutrinos, and thus cannot directly point them back to their potential astrophysical origins. However, the sub-threshold analysis conducted in this thesis motivates future source search investigations, especially with the increase in sensitivity of upcoming experiments.

Acknowledgements

This thesis and the work presented herein would not have been made possible without the following individuals and organisations.

First and foremost, I would like to express gratitude to my supervisor, Ryan Nichol, for the opportunity to study at University College London and work on the ANITA experiment. I owe a multitude of thanks to several members of the ANITA Collaboration, especially Linda Cremonesi and Cosmin Deaconu. I extend thanks to Linda for her willingness to help me on all of the projects I was involved in, and particularly for her guidance in the use of Monte Carlo neutrino simulation. I express my thanks to Cosmin, who is truly a fount of knowledge, not only in physics, but computing as well. The ANITA-4 source search would not have been realised if not for all of his assiduous work. I thank Peter Gorham, David Saltzberg, Christian Miki, Brian Hill, Garry Simburger, and many others, for their work and support at the Columbia Scientific Balloon Facility, the Long Duration Balloon Facility, or both! I give thanks to my fellow doctoral students, especially Andrew Ludwig and John Russell. I appreciate Andrew for his analysis work, in particular for his neutrino singlet search that served as a basis for the early part of the source search. I appreciate John Russell for learning the ins and outs of the ANITA experiment with me, whether we be in rural Texas or the harsh continent itself. Cheers, lads.

Thank you to the organisations that helped make this happen. I extend great thanks to the Science and Technology Facilities Council (STFC), the Leverhulme Trust, and Foresters for their support, and the National Science Foundation (SCF) and the National Aeronautics and Space Administration (NASA) for funding the experiment.

I have had great people by my side. As such, I want to offer thanks to my friends and family. Thank you to Marcus, Jamie, and Talia, for your friendship and encouragement, and Matt, Andy, and Sharples, for an enjoyably frostbitten decade. Cheers to the Forgotten Realms community and my local group for keeping me busy (and sane) outside of physics. To my parents, thank you for the undying love and support you have shown me. For all you do, and who you are, I am eternally grateful.

Impact Statement

Experimental particle physics and related fields of study have proven themselves vital to the progression of society. An understanding of the fundamental building blocks of nature has led to the discovery and usage of electricity, medical imaging, and a wealth of other technologies. Detection and analysis of the messengers of light, photons, have led to a greater understanding of our Universe. Such particles are unable to move large distances through space at high energies, meaning they cannot tell us about far-away objects or highly-energetic processes. Particles known as neutrinos, however, can offer this information.

The work presented in this thesis provides the theoretical background of the physics involved and experimental details of the ANtarctic Impulsive Transient Antenna (ANITA) experiment. ANITA aims to find radio signatures of neutrinos in Antarctica. If such particles are found, it may be possible to point them back to their origin. This, in turn, would shed light on the most energetic activities in the Universe and the objects they come from. It would also provide answers to as-of-yet unanswered questions in astroparticle physics, namely how particles are accelerated to such incredibly high speeds. This thesis includes simulation and analysis techniques for pointing potential neutrino events back to their origin, and thus provides a crucial step in the means of achieving this.

Contents

1 Introduction **34**

2 Particle Physics **37**

2.1 The Standard Model 37

2.1.1 Fermions 37

2.1.2 Bosons 39

2.2 Neutrino Physics 39

2.2.1 The Need for the Neutrino 40

2.2.2 Neutrino Oscillations 41

2.2.3 Neutrino Interactions 44

2.2.4 Probing Further 46

3 Ultra-High Energy Astroparticle Physics **47**

3.1 Ultra-High Energy Cosmic Rays 47

3.1.1 Acceleration 47

3.1.2 Spectrum 48

3.1.3 GZK Neutrinos 56

3.1.4 Composition 56

3.1.5 Potential Sources 59

4 The Askaryan Effect **63**

4.1 Induced Electromagnetic Cascade 63

4.2 Cherenkov Radiation 66

4.3 Coherent Radio Emission 68

4.4 Refraction . . . 69
4.5 Detection . . . 71

5 ANITA 73
5.1 Components and Signal Chain . . . 77
5.1.1 Antennas . . . 78
5.1.2 AMPAs . . . 87
5.1.3 Flight Cables . . . 87
5.1.4 iRFCMs . . . 88
5.1.5 Trigger Path . . . 92
5.1.6 Digitisation Path . . . 96
5.1.7 Position and Attitude Determination Systems . . . 98
5.1.8 Instrument Box . . . 100
5.1.9 SIP . . . 100
5.1.10 Power Systems . . . 101
5.1.11 Storage . . . 102
5.2 Pre-Launch . . . 102
5.2.1 Construction . . . 102
5.2.2 Pre-Flight Tests . . . 104
5.3 Launch . . . 105
5.3.1 Preparation and Ascension . . . 105
5.4 Post-Launch . . . 106
5.4.1 Monitoring . . . 106
5.4.2 Flight . . . 109

6 Simulation 110
6.1 icemc . . . 110
6.1.1 Diffuse Neutrino Simulation . . . 111
6.1.2 Source Search Integration . . . 120
6.2 CRPropa . . . 131
6.2.1 Input . . . 131

6.2.2 Modules . . . 133
6.2.3 Flux Spectrum Generation . . . 133

7 Analysis 138
7.1 Pre-Source Search . . . 139
7.1.1 Quality Cuts . . . 139
7.1.2 Thermal Cuts . . . 141
7.2 antarcticR . . . 146
7.2.1 Clustering . . . 148
7.3 Source Search . . . 156
7.3.1 Regression . . . 156
7.3.2 Pointback . . . 161
7.3.3 Proximity . . . 164
7.3.4 Results . . . 173

8 Conclusions 186

A Cosmic ray flux spectra derivations 188
A.1 Second order Fermi acceleration . . . 188
A.1.1 Fractional energy gain . . . 188
A.1.2 Spectrum . . . 190
A.2 First order Fermi acceleration . . . 191
A.2.1 Fractional energy gain . . . 193
A.2.2 Spectrum . . . 194

B Circuit diagrams 197
B.1 Antenna Testing . . . 197
B.1.1 Boresight Tests . . . 197
B.1.2 Pulse Tests . . . 198
B.2 SHORT Board Testing . . . 199
B.2.1 Basic SHORT Board Circuit . . . 199
B.2.2 SHORT Board Circuit with a Noise Source . . . 200

B.2.3 Initial Trigger Path Replication: SHORT Board Circuit with a Hybrid 201

C Break energy calculation **202**

List of Figures

2.1 Initial neutrino interaction processes. 45

3.1 The high energy cosmic ray spectrum, showing the knees, ankle and GZK cut-off. Image from [41]. 49

3.2 A Hillas diagram showing the different types of candidate objects and mechanisms that could produce UHECRs [42]. In order to be able to accelerate cosmic rays to high enough energies, an object must lie above the diagonal lines, each depicting different scenarios. A blue line indicates the combination of magnetic field strength and size needed to accelerate iron nuclei to ultra-high energies, whilst a red line indicates the same for protons. The line associated with iron nuclei is lower than that of protons, as it is easier to accelerate particles of higher charge (proportional to proton number, Z). A solid line represents a shock velocity of $\beta = 1$, whilst the dashed line is the same but for $\beta = 0.01$ (i.e., the combination is more stringent for the dashed line, as a lower shock velocity will require sources of larger magnetic field strength or size to boost the cosmic rays to ultra-high energies). The diagonal lines associated with the Hillas condition assume that $\eta = 1$ from Equation 3.10. Image from [56]. . 52

3.3 The rate of cosmic rays arriving on Earth versus the energy of the cosmic rays. Three features are observed: the two knees and the ankle. As well as the total rate, three separate rates are shown for the individual components as a function of energy. Image modified from [56]. 53

3.4 The observable distance of specific particles as they travel through the Universe. The red region is not observable with protons, and the blue region is not observable with photons. Note that there is no indication of neutrinos on this plot, as they are observable within all energies and distances as shown. Distances of some classes of astrophysical objects are shown, such as those of galaxies and quasi-stellar objects (quasars). Diagram produced by Peter Gorham. . . . 57

3.5 Preliminary Auger composition data, showing the average depth of the shower maximum vs the energy of the shower. Whilst CRs at the low energies of the UHE range appear to be mostly protonic, those at the high end appear to have a mixed composition. Image from [49]. 58

3.6 T_{90} distribution for the first Burst and Transient Source Experiment (BATSE) catalogue. The vertical line shows the division between short-duration and long-duration bursts. Image modified from [55]. . 61

4.1 Interactions involved in the core formation of the electromagnetic cascade. γ^* denotes a virtual photon. 63

4.2 Interactions that decrease the ratio of positrons to electrons in the shower. Note that the atomic electron (bottom left of each Feynman diagram) is at rest. 64

4.3 Interactions that increase the ratio of electrons to positrons in the shower. Only a single diagram for each possible process is shown i.e., the *t-channel* is displayed for the scattering processes. Note that the atomic electron (bottom left of each Feynman diagram) is at rest. 66

4.4 Illustration of how a wavefront is produced from a moving charged particle. The green arrow represents the full path of the charged particle, whereas the black dots represent a single point in time when the spherical waves are emitted. The red rings represent the spherical waves, which overlap in case (b). The blue lines represent the wavefront, and the black lines indicate the direction in which the wavefront is moving. 67

4.5 Side view of the shower axis. The orange circles represent the base of the Cherenkov cone. The figure is slightly tilted into the page to reveal the three-dimensional cone shape. 68

4.6 An illustration of the photon emission at different wavelengths. The black and red waves indicate the wavelengths of the photons. The green translucent lines indicate the particle shower. 69

4.7 An illustration of the polarisation of the emitted photons. The red arrows indicate that Askaryan radiation is radial to the shower axis, giving rise to predominantly vertical polarisation close to the top of the Cherenkov cone. 70

4.8 The full scenario: A neutrino *skims* the ice (shown as the green arrow at an angle relative to the dashed green line) at a small angle relative to the surface. A Cherenkov cone is formed, emitting coherent radiation of radio wavelengths, as these large wavelengths constructively interfere. Predominantly vertically polarised radio waves escape the ice (first thick black line), and are detected by an antenna. In the case of the critical angle (first dotted line from the top of the cone), the refraction angle becomes 90°. The second dotted line shows the case where the angle is so large that total internal reflection occurs. The red lines indicate the directions of polarisation, and the solid cyan line shows the surface of the ice. 71

5.1 The ANITA-4 payload on the launch vehicle, *the Boss*, prior to launch preparation. 74

5.2 An illustration of the setup used to measure the Askaryan effect in ice at the Stanford Linear Accelerator Center (SLAC). The ANITA payload is suspended from a crane hook, and a beam of electrons was fired into an ice target. An electromagnetic shower builds up, and coherent radio waves are emitted, which are then detected by the payload. Image modified from [74]. 75

5.3 Electric field strength vs elevation angle from the beamline of the resulting Askaryan radiation 1 metre from the shower axis for 200-800 MHz (top) and for individual frequency bands (bottom). Dashed lines are the theoretical field strengths for refracted signals for showers of 3×10^{19} eV. The solid line is the same but for electric fields within the ice. The Total Internal Reflection (TIR) cutoff is also shown. Plots from [74]. 76

5.4 An illustration of the polarisation of the emitted photons via an EAS. 77

5.5 ANITA detection scenario, illustrating the neutrino-induced Askaryan effect in ice, which results in radio waves detected by ANITA. Image modified from [71]. 78

5.6 Schematic of the electronics of ANITA-4 from [79]. 79

5.7 Experimental setup for boresight testing. 80

5.8 S parameter plot for vertical co-polarisation for each antenna type. . 81

5.9 Average gain plots for each flight-specific antenna, showing antennas from ANITA-3 and ANITA-4 with VPOL and HPOL components, as well as the reference antenna's gain. The gains of the ANITA antennas were calculated from the S parameters measured from the boresight tests. 83

5.10 Average effective area for ANITA-4 antennas (both VPOL and HPOL shown) as a function of frequency. 83

5.11 Experimental setup for off-boresight pulse testing. 84

5.12 Antenna power spectra (power vs frequency for different azimuths) for a single ANITA-4 antenna (left). The power-azimuth relation for different frequencies is shown (right). The powers shown are relative to the boresight measurement for the same antenna. 85

5.13 AMPA gain (left) and noise temperature (right) vs frequency. The black line indicates the average values of gain and noise temperature at each frequency. 87

5.14 The measurements of group delay for the SFX-500 flight cables. . . 88

5.15 Internals of 1 of the 4 iRFCMs. Displayed are 2 TUFF modules, each of which has 12 RF input channels. ANITA has eight TUFF modules in total, sufficient for 96 RF channels. The TUFF master is used to issue commands to the TUFF boards, such as selecting which notches to use. Image modified from [84]. 89

5.16 The internals of a single TUFF, shown next to a quarter dollar coin for size comparison. Image from [84]. 90

5.17 Fractional masking averaged over 10 minutes for ANITA-3 and ANITA-4. The shaded regions correspond to when the payload was in line-of-sight to the Long Duration Balloon facility. In this case, there is a large amount of anthropogenic noise, and this needs to blocked out via masking. Image from [84]. 91

5.18 Replication of the initial part of the triggering path. A waveform from the generator is passed through an AMPA, then through a hybrid unit. The hybrid converts waveforms of linear polarisations into those of circular polarisations. After this, the signal is sent to the SHORT board and back to the oscilloscope (it has to pass through a SHORT interface, followed by a SHORTmate). 93

5.19 L1 trigger logic for a single antenna in the top ring. CLK represents the clock (4 ns clock edge, 125 MHz). If there is an L0 trigger from either an LCP waveform ($L0_L$), or an L0 trigger from an RCP waveform ($L0_R$), a 4 ns window opens up. If these $L0_L$ and $L0_R$ windows overlap, an L1 trigger is issued. 94

5.20 A screenshot of the oscilloscope with all outputs shown. The peak-to-peak voltages of average spectra are displayed. 95

5.21 L2 trigger logic for a single phi-sector. CLK represents the clock. If there is an L1 trigger from the bottom antenna, a 12 ns window opens up ($L1_T$). If there is an L1 trigger from the middle antenna, an 8 ns window opens up ($L1_M$). If there is an L1 trigger from the top antenna, a 4 ns window opens up ($L1_B$). If these windows overlap, an L2 trigger is issued. 95

5.22 L3 trigger logic for a span of three adjacent phi-sectors. CLK represents the clock. The central phi-sector is represented as $L2_\phi$, and the phi-sectors either side are represented as $L2_{\phi-1}$ and $L2_{\phi+1}$. If two or more L2 trigger windows overlap in adjacent phi-sectors, an L3 trigger is issued. 97

5.23 Comparison between the primary system (ADU5A) and backup system (sun sensor) measurements of heading for three individual runs [87]. 99

5.24 The assembly of ANITA. 103

5.25 Pre-flight tests at CSBF. 104

5.26 ANITA launch preparation setup. The Long Duration Balloon is inflated with helium (left), and a small test balloon (top left) is launched to monitor wind speed. The base of the Long Duration Balloon is attached to the spool vehicle; this balloon is connected to the ANITA payload (with a parachute interconnect), which is on *the Boss* (right). 105

5.27 ANITA launch sequence. When the helium balloon is sufficiently full to carry the payload into the air, it is released from the spool vehicle. The balloon then rises until it is completely vertical above payload. At such a point, ANITA is detached from the launch vehicle, and begins to ascend. 106

5.28 An overview of the AWARE interface. 107

5.29 The full ANITA-4 flight path over the BEDMAP antarctic map, as shown in the AWARE web-monitor. The purple crosses show the highest priority events coming from the ice at a single time near the end of the flight. The white triangles show locations of potential pulser bases, such as the LDB facility and the West Antarctic Ice Sheet (WAIS) Divide. 108

5.30 The ANITA-4 crash site. A team was dispatched to recover the hard drives and other components, such as the AMPAs. Photograph by Christian Miki. 109

6.1 Flowchart of the icemc simulation for diffuse neutrino searches. Image from [91]. 112

6.2 Phase weight diagrams. 114

6.3 Individual simulated source fluences and the quasi-diffuse fluence of prompt neutrinos for GRBs active during ANITA-4. The thin dashed lines represent the individual fluences, whilst the thick solid line represents the quasi-diffuse fluence. 124

6.4 Individual simulated source fluences and the quasi-diffuse fluence of afterglow neutrinos for GRBs active during ANITA-4. The thin dashed lines represent the individual fluences, whilst the thick solid line represents the quasi-diffuse fluence. 125

6.5 Comparison of the simulated quasi-diffuse fluences for prompt (red) and afterglow (black) neutrinos. Due to the much shallower energy scaling of the afterglow neutrino spectra at ultra-high energies, the afterglow neutrinos have a much higher total fluence within ANITA's energy range (which this graph is limited to). Note that the fluence is presented on a linear scale. 125

6.6 Celestial coordinates illustrated: the right ascension (α) and declination (δ) of an object is shown. 127

6.7 Energy flux spectrum for a single source within ANITA-4's detectable energy range. The energy of a neutrino from this source is determined by picking a random energy as distributed by this function. 128

6.8 UHE neutrino spectrum for differing maximum cosmic ray injection energies. 131

6.9 Nucleon and neutrino count vs energy (left). Nucleon and neutrino count scaled with energy squared vs energy (right). 132

6.10 The nucleon flux is normalised to the Auger cosmic ray spectrum at $10^{19.55}$ eV. The neutrino flux is normalised with the same scaling as was done for the nucleon spectrum. 134

6.11 Predicted fluxes (dashed) and experimental limits (bold) for UHE neutrinos. The dashed green line (*) corresponds to the flux simulated with the custom CRPropa-integrated program using the parameters and assumptions as described above. Alongside the simulated flux are a set of cosmogenic neutrino models [119–122], limits set by the ANITA experiment [77, 79] and other experimental limits [123, 124]. The simulated neutrino flux spectrum is similar in shape to the KKSS flux spectrum, but uses a more realistic source power law index, and thus, its flux is lower. The simulated flux is more favourable than the maximum bound set by Kotera, *et al*, and the limits set by Ahlers, *et al*. As such, the simulated spectrum predicts a high flux of ultra-high energy neutrinos that approaches the limits set by ANITA. 137

7.1 An example of a payload blast event. This event is classified as a payload blast due to its large difference in peak-to-peak voltages between antennas in the bottom ring to the top ring. 140

7.2 A West Antarctic Ice Sheet (WAIS) calibration pulse for both the coherently-summed and de-dispersed waveform. The graphs display the vertically-polarised waveforms (blue) in the foreground and the corresponding cross-polarisation (dashed red) in the background. 142

7.3 Fisher discriminant histogram for all vertically-polarised events. . . 144

7.4 Fisher discriminant background and signal efficiency plots: black vertical lines represent cuts considered for the singlet search and green vertical lines represent the cut for the search in this analysis. Note that the cut for this analysis is purposely set to be a lot looser. Due to the nature of the sub-threshold analysis, we accept far more background events, and also capture more signal events. Plot modified from [125]. 145

7.5 Multiple views and the clustering output for a subset of ANITA-4 data using antarcticR. 146

7.6 The interactive online companion of antarcticR. This example shows a high-resolution bedrock map of the Antarctic peninsula with a specific refuge selected. 147

7.7 An example set of data (unclustered). 148

7.8 The reachability plot (generated by the OPTICS algorithm) for our example set of data. The coloured "valleys" represent the closest set of points in the whole dataset in terms of reachability distance, thus these are clustered events. The red valley is significantly more prominent that the green, blue, and cyan valleys, as it contains more events. The black spikes represent the points furthest away from the next closest points, thus they are isolated events. The reachability distance cut, ϵ_{cl}, is set to 200 km, as this selection contains the valleys, i.e., the clusters. Setting a much larger cut would cause the algorithm to group too many outliers into a cluster, and setting it too low would force it to ignore points clearly in a cluster. 149

7.9 An example set of data (clustered). By setting a reachability distance cut, we are able to identify clusters (coloured) and separate them from the desired events, which are isolated (black). 150

7.10 Reachability plot for all events in the pre-clustered dataset. The coloured areas (valleys) represent clusters, whilst the black areas (the spikes, corresponding to events very far away from their closest events) represent non-clustered events. Note that almost all the 887793 events are in clusters, and some of the non-clustered events are very far from the nearest cluster (high reachability distance). Thus, the reachability plot is mostly dominated by such valleys with occasional spikes, and appears stretched. 151

7.11 Pre-clustered events pointed back to the surface. Red events labelled "A" are isolated events, coloured events with other letters are those that clustered. The most active areas are identified, such as the Amundsen-Scott South Pole Station (central green), McMurdo Station (bottom orange), the West Antarctic Ice Sheet (pink), and Mount Vinson (purple). 152

7.12 Pre-clustered events pointing back to the sky. The dense regions correlate to bases, which dominates the entire sky map. Right ascension (ra) and declination (dec) are always in degrees. The z-axis indicates the number of events in a single bin area. 153

7.13 Events recorded when the ANITA payload passed right over Byrd station, projected to the sky. As the payload passes over the station, the declination gets much larger. Right ascension (ra) and declination (dec) are always in degrees. The z-axis indicates the number of events in a single bin area. 153

7.14 WAIS pulses pointed back to the sky. The declination remains constant as the payload did not travel right over the top of the calibration pulsers. Right ascension (ra) and declination (dec) are always in degrees. The z-axis indicates the number of events in a single bin area. 154

7.15 Post-clustered events pointing back to the sky. The distribution of events is as expected, as no dense regions or clusters are observed at all. As only isolated events remain, the scale of the z-axis is now irrelevant (bin areas can only be black or white). To make these events easily visible, the bin areas were increased, so some events appear to be touching, even though this is not the case. Right ascension (ra) and declination (dec) are always in degrees. 155

7.16 Geometric illustration of neutrino-induced radio waves, with elevation angles labelled. θ_ν indicates the neutrino angle, θ_γ indicates the radio-exit angle, and θ_{ch} indicates the Cherenkov angle. 156

7.17 Histograms showing the relationship between d$\phi_{\gamma\nu}$ (top) and $d\theta_{\gamma\nu}$ (bottom) vs linear polarisation angle, l, for a set of diffuse neutrinos. 158

7.18 Node structure of the MTMLP. The input variables and bias node are shown to the left (input is layer 0). The 24 nodes are shown in (the hidden) layer 1, and the output layer contains the two targets of the MLP. 159

7.19 Resolution histograms for diffuse data: The difference between the neutrino angles from regression and from simulation (test data). The standard deviation for the elevation angle is less than 2°, whilst the standard deviation for the azimuthal angle is just above 2°. . . . 160

7.20 Simulation: a map of ANITA's sky exposure to 10^{20} eV neutrinos. Note that the exposure is much higher at low absolute values of declinations, whilst it is uniform thoroughout all values of right ascension. Image from [131]. 161

7.21 Resolution histogram for a single blazar: the difference between neutrino angles from regression and from simulation. 162

7.22 Pointing back to a single source of simulated neutrinos via regression. The distribution is very well contained, with essential no stray events. Right ascension (ra) and declination (dec) are always in degrees. 164

7.23 A sky map of the set of flaring blazars active during ANITA-4. Flat-spectrum radio quasars (FSRQs) are shown in red, BCUs (blazars candidates of uncertain type) are shown in green, and BLLs (BL Lacertae objects) are shown in black. 166

7.24 Pointed back neutrinos overlaid with all blazar sources. The red squares represent single observations of blazars as $4° \times 4°$ bin areas. The distributions represent the locations where the simulated neutrinos point back to. Right ascension (ra) and declination (dec) are always in degrees. The z-axis denotes the amount of events in a bin area. 167

7.25 Example: Construction of the spatial proximity function for a single source. Right ascension (ra) and declination (dec) are always in degrees. 168

7.26 Temporal proximity functions for two blazars. 170

7.27 Normalised cumulative distribution of the spatiotemporal proximity for a single source (green) and a diffuse set (red). 171

7.28 Long-lived source example: an example of a summed proximity distribution for events pointing back to a long-lived source (SN 2016iyl) for the set of alternative universes. 174

7.29 Short-lived source example: an example of the logarithmic summed proximity distribution for events pointing back to a typical transient source (GRB161219A) for the set of alternative universes. 175

7.30 CDF of the logarithmic summed proximity distribution for events pointing back to GRB161214B for the set of alternative universes. Our Universe is represented by the black line, the mean of all alternative universes is represented by the red dashed line, the first significance threshold (CL 95%) is represented by the orange dashed line, and the second significance threshold (CL 99%) is represented by the green dashed line. This GRB is well within expected limits, and is just higher than the mean. 177

7.31 CDF of the summed proximity distribution for events pointing back to PKS 1502+106 for the set of alternative universes. Our Universe is represented by the black line, the mean of all alternative universes is represented by the red dashed line, the first significance threshold (CL 95%) is represented by the orange dashed line, and the second significance threshold (CL 99%) is represented by the green dashed line. This flaring blazar exceeds both limits of the confidence level. . 178

7.32 Fermi All-Sky Variability Analysis (FAVA) high energy light curve for the flare observed during the *first* week (in the direction of PKS 1502+106). The Mission Elapsed Time (MET) is the number of seconds elapsed since a base time (January 1st 2001). The red dashed line indicates the observation week. This plot was generated using the FAVA light curve generator [101]. 180

7.33 Fermi All-Sky Variability Analysis (FAVA) high energy light curve for the flare observed during the *second* week (in the direction of PKS 1502+106). As the measured right-ascension and declination of the localisation of the events differ slightly from week by week, the light curves also differ slightly. The Mission Elapsed Time (MET) is the number of seconds elapsed since a base time (January 1st 2001). The red dashed line indicates the observation week. This plot was generated using the FAVA light curve generator [101]. 181

7.34 Fermi All-Sky Variability Analysis (FAVA) test statistic maps (left is the low-energy photon band, right is the high-energy photon band) for the first flare associated with PKS 1502+106. The corresponding map was generated using the FAVA analysis tools [101]. 182

7.35 Above: events pointing back to the vicinity of PKS 1502+106 (in our Universe) on a spatial proximity graph. Right ascension (ra) and declination (dec) are always in degrees. Below: Histogram of individual event proximities to PKS 1502+106 in our Universe. . . . 183

7.36 Ice radio-exit positions of the highest-priority events from the PKS 1502+106 analysis. 184

7.37 Vertically-polarised (red) and horizontally-polarised (green) waveforms for the triggered channels of event 15594227, an event from the PKS 1502+106 analysis. 185

7.38 The vertically-polarised (red) coherently-summed deconvolved waveform of event 15594227, an event from the PKS 1502+106 analysis. The dashed line in the background corresponds to the cross-polarisation. 185

A.1 Shock scenarios shown in the galactic frame, shock frame, unshocked region frame, and shocked region frame. The green area represents the shocked medium (particles that have been shocked, i.e., downstream). The yellow area represents the shock itself. The red area represents the unshocked region (particles that haven't been shocked yet, i.e., upstream). 193

B.1 Circuit diagram of boresight antenna tests. The antennas were pointed directly towards each other and their gains were calculated. This was done for all ANITA-4 antennas, as well as those from other flights. 197

B.2 A pulse generator was used to send a pulse from one antenna to another. In this case, the receiving antenna's oritentation was altered to see how the power depended on azimuth and elevation. 198

B.3 Circuit diagram of all components of the basic SHORT setup. This setup allowed recording of the voltages output from the SHORTs. To replicate the signal and select certain frequencies, components such as attenutors and bandpass filters were used. Three signals were recorded: the dummy antenna signal, the split pulse, and the SHORT output signal. 199

B.4 Circuit diagram of noise addition to the SHORT setup. A solid state noise source was used to modify the pulse, and was used to check signal-to-noise ratios. 200

B.5 Circuit diagram of the SHORT setup, including a hybrid. A pulse generator and the primary components are used to replicate the antenna signal. As per the actual ANITA trigger system, the waveform then moves through the AMPAs, a hybrid unit, and a SHORT board. 201

List of Tables

7.1 Status of the data at several points during the analysis, including quality cuts, thermal cuts, and clustering cuts. The remaining events are shown in the order the cuts were applied, where the previous cuts have already been applied. For detailed information about each specific stage and the cuts involved, see Sections 7.1.1, 7.1.2, and 7.2.1. 141

7.2 The set of blazars associated with flares during the ANITA-4 flight. The source classes for blazars are flat-spectrum radio quasars (FSRQs), BCUs (blazars candidates of uncertain type), and BLLs (BL Lacertae objects). Note, for blazars that are observed in more than one week, only the first observation of right ascension and declination are shown. 165

7.3 The set of final blazars, along with the test statistic (summed proximity), L, and the results. $\overline{L}$ represents the mean value of the test statistic for the alternative universes, and L_0 is the value of the test statistic in our Universe. 176

7.4 The set of final supernovae, along with the test statistic (summed proximity), L, and the results. $\overline{L}$ represents the mean value of the test statistic for the alternative universes, and L_0 is the value of the test statistic in our Universe. 176

7.5 The set of final gamma-ray bursts, along with the test statistic (logarithm of the summed proximity), L, and the results. 177

7.6 The two flares observed by the Fermi All-Sky Variability Analysis (FAVA) for PKS 1502+106 (3FGL J1504.4+1029) during the ANITA-4 flight. The expected and observed events correspond to the expected and measured high energy gamma ray events from the direction of the above right ascensions and declinations. The expected events correspond to the expected number of gamma-rays, i.e., the background, whereas the observed events number corresponds to the observed number of gamma-rays, i.e., the signal. Sigma corresponds to the the significance of the flux variation. The corresponding information was calculated by using the FAVA analysis tools [101]. 179

Nomenclature

3FGL . . . . . 3rd Fermi Gamma-Ray LAT catalog

4C . . . . . 4th Cambridge catalogue

7C . . . . . 7th Cambridge catalogue

ΛCDM . . . . . Λ Cold Dark Matter[1]

ADC . . . . . Analog-to-Digital Converter

ADU . . . . . Attitude Determination Unit

AGN . . . . . Active Galactic Nucleus (or Nuclei)

AMPA . . . . . Antenna-Mounted Pre-Amplifier

ANITA . . . . . ANtarctic Impulsive Transient Antenna

AWARE . . . . . Active Web for Antarctica Radio Experiments

BATSE . . . . . Burst And Transient Source Experiment

BCU . . . . . Blazar Candidate of Uncertain type

BEDMAP . . . . . Bed Topography of the Antarctic

BLL, BL Lac . . BL Lacertae Object

BZ . . . . . Berezinsky-Zatsepin

CC . . . . . Charged Current

CERN . . . . . European Council for Nuclear Research (Conseil Européen pour la Recherche Nucléaire)

CL . . . . . Confidence Level

CMB . . . . . Cosmic Microwave Background

[1] Λ represents the cosmological constant associated with dark energy.

CP . . . . . Charge Conjugation and Parity

CR . . . . . Cosmic Ray

CRPropa . . . . . Cosmic Ray Propagation Framework

CRUST 2.0 . . . . New Global Crustal Model at 2x2 Degrees

CSBF . . . . . Columbia Scientific Balloon Facility

CTA . . . . . Cherenkov Telescope Array

CW . . . . . Continuous Wave

DAC . . . . . Digital-to-Analog Converter

DAQ . . . . . Data Acquisition Unit

DC . . . . . Direct Current

DIS . . . . . Deep Inelastic Scattering

EAS . . . . . Extensive Air Shower

FAVA . . . . . Fermi All-Sky Variability Analysis

FDA . . . . . Fisher Discriminant Analysis

FLSAT . . . . . Fleet Satellite

FSRQ . . . . . Flat-Spectrum Radio Quasar

G12 . . . . . 12-channel GPS board

GPS . . . . . Global Positioning System

GRB . . . . . Gamma Ray Burst

GZK . . . . . Greisen-Zatsepin-Kuzmin

H . . . . . HESS

HDD . . . . . Hard Disk Drive

HGST . . . . . Hitachi Global Storage Technologies

HPOL . . . . . Horizontally Polarised

icemc . . . . . ice monte-carlo

iRFCM . . . . . internal Radio Frequency Conditioning Module

ISM . . . . . Interstellar Medium (Media)

JSON . . . . . JavaScript Object Notation

K2K . . . . . KEK To Kamioka

LAB3 . . . . . Third-generation LABRADOR

LAT . . . . . Large Area Telescope

L0, L1, L2, L3 . Level 0, 1, 2, 3 (Trigger)

LABRADOR . . Large Analog Bandwidth Recorder And Digitizer with Ordered Readout

LCP . . . . . Left Circular Polarisation

LDB . . . . . Long Duration Balloon

LHC . . . . . Large Hadron Collider

LNA . . . . . Low Noise Amplifier

LOS . . . . . Line-Of-Sight

LPM . . . . . Landau-Pomeranchuk-Migdal

MET . . . . . Mission Elapsed Time

MINOS . . . . . Main Injector Neutrino Oscillation Search

MLP . . . . . Multi-Layer Perceptron

MTMLP ..... Multi-Target Multi-Layer Perceptron

MUOS ..... Mobile User Object System

NASA ..... National Aeronautics and Space Administration

NC ..... Neutral Current

OK ..... Ohio University K[2]

OPTICS ..... Ordering Points To Identify the Clustering Structure

OVV ..... Optically Violent Variable (quasar)

PKS ..... Parkes Observatory

PMNS ..... Pontecorvo-Maki-Nakagawa-Sakata

PMN ..... Parkes-MIT-NRAO

RA ..... Right Ascension

RENO ..... Reactor Experiment for Neutrino Oscillation

RF ..... Radio-Frequency

RCP ..... Right Circular Polarisation

RMS ..... Root Mean Square

ROOT ..... The name of CERN's Object-Orientated Framework[3]

S5 ..... Fifth Survey of Strong Radio Sources

SCA ..... Switched Capacitor Array

SFX ..... Superflexible Coaxial Cable

SHORT ..... SURF High Occupancy RF Trigger

[2]The K represents objects found in the 10th hour of right ascension.

[3]The name ROOT originated from the "roots" of end-users applications. René Brun, creator of the ROOT system, suggested that the acronym could stand for Rapid Object-Oriented Technology.

SIP . . . . . Support Instrument Packages

SLAC . . . . . Standford Linear Accelerator Center

SN, SNe . . . . . Supernova, Supernovae

SNO . . . . . Sudbury Neutrino Observatory

SNR . . . . . Supernova Remnant

SSD . . . . . Solid State Drive

SURF . . . . . Sampling Unit for Radio-Frequency

T2K . . . . . Tokai To Kamioka

TDRSS . . . . . Tracking and Data Relay Satellite System

TIR . . . . . Total Internal Reflection

TMVA . . . . . Toolkit for Multivariate Data Analysis with ROOT

TNS . . . . . Transient Name Server

TUFF . . . . . Tunable Universal Filter Frontend

TURF . . . . . Triggering Unit for Radio-Frequency

UFO . . . . . Ultra-High Frequency Follow-On

UHE . . . . . Ultra-High Energy

UHECR . . . . . Ultra-High Energy Cosmic Ray

UHEN . . . . . Ultra-High Energy Neutrino

UTC . . . . . Coordinated Universal Time

VPOL . . . . . Vertically Polarised

WAIS . . . . . West Antarctic Ice Sheet

Chapter 1

Introduction

It has been over fifty years since the first observation of an ultra-high energy (UHE) particle from the cosmos [1]. A UHE particle has energy greater than 10^{18} eV, and from where, or how, a single particle acquires such an incredibly large amount of energy is still a mystery to this day. Traditionally, astronomy has relied on the quanta of the electromagnetic spectrum, photons, for information about our Universe: from radio waves, to visible light, to gamma rays. Several experiments around the world have been able to deduce information about astrophysical phenomena from photons. However, as the energy of a photon increases beyond approximately 10^{14} eV, the probability of it interacting with the cosmic microwave background (CMB) becomes large. The photon-photon collision gives rise to electrons and positrons ($\gamma + \gamma_{\mathrm{CMB}} \rightarrow e^- + e^+$), limiting its path length. As a UHE photon cannot travel a long distance before pair producing, the Universe becomes opaque to it. Thus, using photons to find out about universal phenomena is inevitably limited at the UHE range. Fortunately, there remain other information-carrying particles.

Cosmic rays, essentially protons or heavier atomic nuclei, speed through the Universe, but do not take linear paths from their origin to the Earth. Due to their electric charge, the paths of these high energy particles are distorted by galactic or intergalactic magnetic fields, making it difficult to trace them back through the Universe. Furthermore, at $\sim 5 \times 10^{19}$ eV, these protons interact with the CMB, producing heavy hadrons known as Delta baryons ($p + \gamma \rightarrow \Delta^+$). The unstable baryons subsequently decay into charged pions and neutrons ($\Delta^+ \rightarrow \pi^+ + n$), or neutral pions

and protons ($\Delta^+ \rightarrow \pi^0 + p$). The charged pions decay to leptons and lepton neutrinos, whereas the neutral particles decay to photons. Thus, cosmic rays either cease to survive and produce other particles, or have a much reduced energy.

There is, however, a particle that can pass through the Universe whilst barely interacting with its radiation-filled background. This particle, the neutrino, only interacts via the weak force[1] since it is a lepton with no electric charge. Due to its low interaction chance, the neutrino avoids the problems that the previous candidates suffer from in the ultra-high energy regime. However, due to this same reason, detection of neutrinos proves more complicated than that of photons and protons.

The ANtarctic Impulsive Transient Antenna (ANITA) aims to observe ultra-high energy neutrinos (UHENs) via their interactions with the Antarctic ice sheet. When a neutrino of considerable energy passes through the ice sheet, it has a chance of forming a particle cascade, in which the resulting particles move as a net negatively charged bunch. Due to this, coherent radio waves are emitted via the Askaryan effect. The radio observatory, ANITA, is suspended in the air by a National Aeronautics and Space Administration (NASA) Long Duration Balloon (LDB) and intends to find neutrino-induced radio signals coming from the ice approximately 40 km beneath it. Observing such ultra-high energy neutrinos would present an opportunity to study the most extreme universal phenomena and processes at the highest possible energies. Together, neutrinos, along with photons, cosmic rays, and gravitational waves, form the basis of multi-messenger astronomy. The most recent experiment, ANITA-4, flew in the austral summer of 2016. Prior to this were three individual flights, as well as the test flight, ANITA-lite.

The contents of this thesis detail my work prior to, during, and following the launch of ANITA-4. This introduction constitutes the first chapter, whilst the following three go over the background material. Chapter 2 introduces the theoretical material of the essential particle physics, including our current knowledge of the Standard Model, and beyond. Neutrino physics is also explained in this chapter.

[1]Note that the neutrino also interacts gravitationally, but as its mass is constrained to be remarkably low (an upper limit of 1.1 eV [2]), its gravitational interaction has negligible effect on its trajectory.

Chapter 3 goes over ultra-high energy astroparticle physics, particularly that of cosmic rays, neutrinos, and the astrophysical sources themselves. Chapter 4 describes how coherent radiation is produced via the Askaryan effect. Chapter 5 introduces the experiment, its physical components, and flight; it also covers the tests I carried out on these invididual components. Chapter 6 goes through the simulation tools, in particular how I modified them for use in my analysis. Chapter 7 details the main analysis I conducted, an astrophysical source search using ANITA-4 data. Finally, I summarise the thesis and draw conclusions from the analysis in Chapter 8.

Chapter 2

Particle Physics

2.1 The Standard Model

The Standard Model is the name of the theory that describes fundamental particles and how they interact [3]. This astonishingly successful theory describes the building blocks of matter, which are represented in the form of fermions, specifically leptons and quarks. It also explains how the electromagnetic, weak nuclear, and strong nuclear forces are mediated by particles known as bosons. The Standard Model is a quantum field theory of the unitary groups: $SU(3) \times SU(2) \times U(1)$. The masses of each of the fundamental particles are generated through the Higgs field[1], whose quantum excitation is known as the Higgs boson [4, 5]. Every particle that makes up the Standard Model has a mirroring particle, which has the same mass, but opposite charge, known as an *anti*-particle (though some particles are their own anti-particle). The current formulation of the Standard Model was formulated in the 1970s, and since then, experiments have confirmed the existence of the top quark [6, 7], the tau neutrino [8], and lastly, the Higgs boson [9, 10].

2.1.1 Fermions

Fermions are particles, whether fundamental or composite, which have half-integer spin (intrinsic angular momentum). Those which are fundamental always have spin $\frac{1}{2}$, and those which are composite have half-integer spin ($\frac{1}{2}$, $\frac{3}{2}$, etc). Fundamental fermions are divided into two groups: quarks and leptons.

[1] It is unknown whether neutrinos gain their mass from the Higgs mechanism or not.

2.1.1.1 Quarks

Quarks are elementary particles, of which there are six types: up (u), down (d), charm (c), strange (s), top (t), and bottom (b). These individual types are referred to as *flavours*. Quarks can be further categorised into up-type (u, c, t) and down-type (d, s, b) quarks, with each *first* member in both sets belonging to the *first* generation, and so on. Up-type quarks have electric charge, $+\frac{2}{3}e$, whereas down-type quarks have $-\frac{1}{3}e$, where e denotes the elementary charge (the electric charge of a proton or the magnitude of the electric charge of an electron). Quarks, however, cannot exist alone due to confinement, and thus find themselves within larger, composite particles, such as baryons and mesons. The first generation of quarks make up the nucleons, protons (uud) and neutrons (udd), in atomic nuclei.

2.1.1.2 Leptons

Leptons make up the remaining fermions in the Standard Model. As per quarks, there are six types. They are also divided into two sets, the electrically charged leptons: electrons (e), muons (μ), taus (τ), and the electrically neutral neutrinos: electron neutrinos (ν_e), muon neutrinos (ν_μ), tau neutrinos (ν_τ). The electrically charged leptons all carry a charge of -e.

In the Standard Model, each lepton has a $lepton\ number$, which is conserved in particle interactions[2].

$$\mu^- \rightarrow e^- + \overline{\nu_e} + \nu_\mu \tag{2.1}$$

In the above case of muon decay, the muon number is conserved due to the production of a muon neutrino, and the electron number is conserved due to the production of both an electron and an anti-electron neutrino. Thus, the total lepton count is preserved.

[2]Note that conservation of lepton number is an "accidental" symmetry of the Standard Model. In some theories beyond the Standard Model, and in neutrino oscillations, individual lepton number is not conserved.

2.1.2 Bosons

The interactions between the elementary particles and forces are mediated by the gauge bosons of the Standard Model: the photon (γ), the gluon (g) and the intermediate vector bosons ($W^{\pm}$, Z^0). Unlike fermions, all bosons have integer spin, with the force-carrying particles being vector bosons (spin 1).

The only gauge bosons in the Standard Model which have mass, the $W^{\pm}$ and Z^0 bosons, govern the interactions of the weak nuclear force, and couple to weak isospin. All elementary fermions interact with the weak nuclear force, and can be grouped into *weak isospin* doublets, in order of their generations.

$$\underbrace{\begin{pmatrix} u \\ d \end{pmatrix}\begin{pmatrix} c \\ s \end{pmatrix}\begin{pmatrix} t \\ b \end{pmatrix}}_{\text{quark doublets}} \quad \underbrace{\begin{pmatrix} \nu_e \\ e \end{pmatrix}\begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}\begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}}_{\text{lepton doublets}}. \tag{2.2}$$

The photon (γ) mediates the interactions of the electromagnetic force, which couples to electric charge. Thus all fermions but the chargeless neutrinos interact with the photon, which is one of the two massless bosons.

The gluon (g) is the remaining massless boson, and the force-carrier of strong nuclear interactions. This boson couples to a quantity known as *colour charge*, with three *colours*: red (r), blue (b) and green (g), which must be conserved. Quarks interact with the strong force, and can take three values of colour. The particles which quarks constitute, however, must be colour *neutral*, from a correct combination of such quarks.

2.2 Neutrino Physics

The Standard Model successfully describes three of the four forces of nature, and the particles that take part in these interactions. However, it does not incorporate gravitational interactions or the presence of dark matter. Particles associated with gravity and those making up dark matter have yet to be observed. The original Standard Model (without modifications) also does not take into account the oscillations of neutrinos.

Of all the massive particles, the neutrino is the lightest. Such particles only interact gravitationally and via the weak force, making the determination of their properties more difficult than other elementary particles.

2.2.1 The Need for the Neutrino

Just prior to the dawn of the 20^{th} century, Henri Becquerel discovered that uranium emitted invisible radiation [11]. Soon after, decays involving the emission of electrons, i.e., beta decays, were intensively studied. In all scenarios, energy must be conserved, and so the same should hold true to these decays. If energy was to be conserved, a discrete spectrum (corresponding to the difference in atomic mass) should be observed. However, in 1914, James Chadwick's initial experiments showed a clear continuous spectrum for beta decays [12]. Thus, it indeed seemed as if the principle of conservation of energy was potentially violated in this scenario.

Enter the neutrino, an elementary particle that is electrically neutral and obeys Fermi-Dirac statistics. The idea of the neutrino[3] was first introduced by Wolfgang Pauli in 1930 to explain the apparent violation of conservation of energy in beta decays [13]. This neutral particle would take away some energy in the interaction, thus conserving energy. Niels Bohr, however, had previously proposed that the same effects would be observed if energy was conserved statistically, i.e., conservation of energy could be violated in some events, but overall, it would still be conserved [14]. Yet, in 1933, experiments by Charles Ellis showed that the beta spectrum had an upper energy bound, and thus ruled out Bohr's suggestion [15]. A year later, Enrico Fermi expanded on Pauli's idea of introducing an additional particle, and provided an effective theory of beta decay[4], which involved not only the physics of nuclear decays, but also the role of the neutrinos within them [16]. Fermi's theory of beta decay was able to accurately calculate reaction rates, though the experimental confirmation of neutrinos would come much later.

Following Fermi's theory of beta decay, there were several methods proposed

[3]The name *neutrino* was conceived by Edoardo Amaldi, meaning "neutral little one", and was later introduced by Enrico Fermi.

[4]Note that this only involved the *electron* neutrino, not the heavier flavours of the other generations.

to detect neutrinos [17, 18], and even preliminary experimental evidence [19, 20]. However, strong evidence came with the Cowan-Reines experiment in 1956 [21]. In this experiment, electron (anti-)neutrinos created in reactor experiments would interact with protons, producing neutrons and positrons. Now, if the positrons were created due to the presence of neutrinos, they would interact with electrons and annihilate to produce two gamma ray photons. A short time later, the resulting neutron would be captured, giving off another photon when the excited nucleus relaxes to the ground state. The successful detection of these joint events provided strong evidence for the existence of neutrinos.

2.2.2 Neutrino Oscillations

As well as in terms of lepton flavour (ν_e, ν_μ, ν_τ), neutrinos can be classified by mass (ν_1, ν_2, ν_3). As such, the flavour eigenstates can be represented by unitary transformations of the mass eigenstates:

$$|\nu_\alpha\rangle = \sum_i U^*_{\alpha i} |\nu_i\rangle \,, \tag{2.3}$$

where U is the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [22, 23], which contains the lepton mixing parameters: the mixing angles between the flavour and mass eigenstates (θ_{ij}), and the CP-violating phase (δ_{CP}), which is only zero if neutrino oscillations obey CP symmetry:

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.4}$$

where $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$. The PMNS matrix has been experimentally shown to not be equivalent to the identity matrix; thus, neutrino flavour eigenstates are not the same as their mass eigenstates. Neutrinos first observed as one flavour eigenstate can later be observed as a different one, i.e., they "oscillate". In order to illustrate how this works, a two-flavour system can be considered, where a neutrino propagates as a superposition of its mass states. Taking this simple scenario in the

form of Equation 2.3, the two-dimensional case of electron-muon neutrino mixing would be represented as:

$$|\nu_\alpha\rangle = \sum_{i=1}^{2} R_{\alpha i} |\nu_i\rangle \,, \tag{2.5}$$

where the PMNS matrix is replaced by a simple rotation matrix, R:

$$R = \begin{pmatrix} \cos(\theta) & \sin(\theta) \\ -\sin(\theta) & \cos(\theta) \end{pmatrix}. \tag{2.6}$$

A neutrino created in flavour state ν_e will thus, initially, have the form:

$$|\nu_e\rangle = R_{\alpha 1} |\nu_1\rangle + R_{\alpha 2} |\nu_2\rangle \,, \tag{2.7}$$

$$= \cos(\theta) |\nu_1\rangle + \sin(\theta) |\nu_2\rangle \,. \tag{2.8}$$

As the neutrino propagates through time ($t > 0$) and space, the mass eigenstates transform with their own associated energies, E_i. After creation, the mass eigenstates approximately evolve as $|\nu_i\rangle \rightarrow e^{iE_i t} |\nu_i\rangle$:

$$|\nu_e(t > 0)\rangle = \cos(\theta) e^{iE_1 t} |\nu_1\rangle + \sin(\theta) e^{iE_2 t} |\nu_2\rangle \,. \tag{2.9}$$

The quantum-mechanical probability of the electron neutrino being observed in the state of a muon neutrino is:

$$P(\nu_e \rightarrow \nu_\mu) = |\, \langle \nu_\mu | \nu_e(t) \rangle \,|^2, \tag{2.10}$$

$$= \left| \Big(-\sin(\theta) \langle \nu_1 | + \cos(\theta) \langle \nu_2 | \Big) \Big(\cos(\theta) e^{iE_1 t} |\nu_1\rangle + \sin(\theta) e^{iE_2 t} |\nu_2\rangle \Big) \right|^2 . \tag{2.11}$$

Expanding and using $\langle \nu_i | \nu_j \rangle = \delta_{ij}$, the cross-terms are eliminated, simplifying the equation:

$$P(\nu_e \rightarrow \nu_\mu) = \left| \sin(\theta)\cos(\theta) \Big(e^{iE_2 t} - e^{iE_1 t} \Big) \right|^2 . \tag{2.12}$$

Applying trigonometric identities and using the exponential form of the sinusoid:

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2(2\theta)\sin^2\left(\frac{(E_2 - E_1)t}{2}\right). \tag{2.13}$$

Considering tiny mass of the neutrino, we can apply the relativistic limit, where the energy carried by the neutrino is much greater than its mass,

$$E_i = \sqrt{p_i^2 + m_i^2} \simeq p_i^2 + \frac{m_i^2}{p_i^2}, \tag{2.14}$$

where p is the neutrino momentum, and m is the neutrino mass. Therefore,

$$(E_2 - E_1)t = \frac{(m_2^2 - m_1^2)L}{2p}, \tag{2.15}$$

where L is the distance travelled. Thus,

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2(2\theta)\sin^2\left(\frac{(m_2^2 - m_1^2)L}{4p}\right), \tag{2.16}$$

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2(2\theta)\sin^2\left(\frac{\Delta m^2 L}{4E}\right), \tag{2.17}$$

where Δm^2 is defined as $m_2^2 - m_1^2$. This equation can be represented in the following form:

$$P(\nu_e \rightarrow \nu_\mu) = \sin^2(2\theta)\sin^2\left(1.27\frac{\Delta m^2 L}{E}\right), \tag{2.18}$$

where Δm is in eV, L is in km, and E is in GeV.

2.2.2.1 Solar neutrino flux

The Homestake experiment of South Dakota was designed to detect electron neutrinos emitted by solar nuclear fusion processes [24, 25]. The experiment gained success in being the first at counting such neutrinos. However, these measurements were approximately three times lower than expected from the astrophysical Standard Solar Model [26], giving rise to the "solar neutrino problem". The problem

persisted for about three decades, until the Sudbury Neutrino Observatory (SNO) provided definitive evidence of solar neutrino flavour oscillations in 2001 [27]. The Homestake experiment had only measured a third of the neutrinos expected, as the electron neutrinos had changed flavour on their course from the Sun to Earth.

2.2.2.2 Measurement of Oscillation Parameters

Since the early experiments in the 1960s, different types of experiments utilising neutrinos from the atmosphere, particle beams, and reactors have calculated the values of the oscillation parameters.

Atmospheric neutrino oscillations have been observed by several experiments, though the first discovery came in 1998 from Super-Kamiokande (Super-K) [28]. Super-K is able to detect neutrinos from all directions, and thus can distinguish between neutrinos produced above the detector (downward-going) and those which have to travel from the other side of the Earth (upward-going). Similarly to the solar neutrino case, there was a deficit of atmospheric neutrinos. The muon neutrinos that are produced from atmospheric interactions have a probability of oscillating into tau neutrinos, which is proportional to the sinusoid of the distance travelled. As upward-going neutrinos had to travel a specific distance (closer to maximising the probability of oscillation) to reach the detector, more muon neutrinos had oscillated into tau neutrinos.

Other experiments, such as MINOS [29], K2K [30], T2K [31], and NOvA [32] have observed neutrino oscillations (such as muon neutrino disappearance) over large baselines. Reactor experiments such as Daya Bay [33], RENO [34], and CHOOZ [35] have been able to measure the disappearance of anti-electron neutrinos.

2.2.3 Neutrino Interactions

When a neutrino moves through a medium, it has a chance of interacting (via the weak force) with the particles that the medium is made up of. A medium consists of atoms, which themselves are composed of neutrons, protons, and electrons. There are several modes of neutrino interaction, though which mechanism is dominant

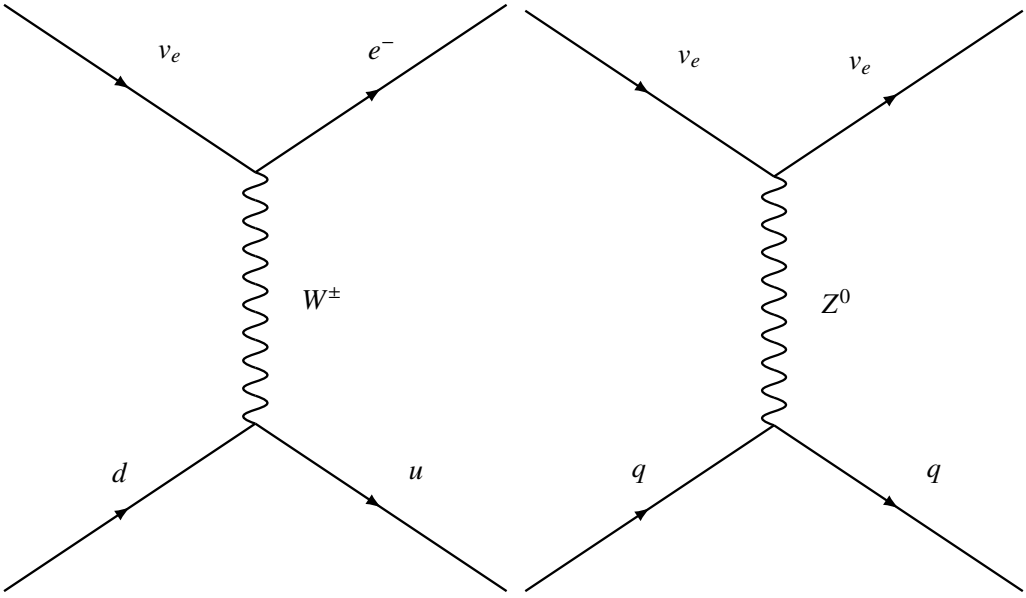

(a) A charged current (CC) neutrino interaction involves the exchange of a $W^\pm$ boson, which produces a charged lepton (such as an electron).

(b) A neutral current (NC) neutrino interaction involves a Z^0 boson, which does not produce a charged lepton, as the neutrino remains.

Figure 2.1: Initial neutrino interaction processes.

depends on the energy of the neutrino. The neutrino can interact with the whole nucleus (as in charged current coherent interactions), with a nucleon (as in quasi-elastic interactions or resonant pion production), or with an electron that orbits the nucleus (as in neutrino-electron scattering). However, at high energies, deep inelastic scattering (DIS) interactions become dominant, and in such a case, the neutrino interacts with a single quark.

Charged current (CC) neutrino-quark scattering occurs when a neutrino interacts with a quark within a nucleon, such as a neutron (made of an up quark and two down quarks). The neutrino interacts with one of these quarks, say a down quark, by exchanging a W boson. In turn, the output of this interaction is an electron and an up quark (one of the down quarks has now become an up quark, such that the neutron is now a proton). See Figure 2.1a for the corresponding Feynman diagram. Similarly, an anti-neutrino can interact with an up quark of a proton, giving rise to a positron and a neutron. The other interaction mode is neutral current (NC), whereby a Z^0 boson is exchanged (and the quark does not change flavour), instead of a W

boson (See Figure 2.1b).

2.2.4 Probing Further

Observations of many yet to be discovered phenomena require large energy scales, those of which that are much higher than those available at present. Current particle accelerator experiments operate at a maximum of $\mathcal{O}(10$ TeV), such as the Large Hadron Collider (LHC), where protons reach energies of approximately 6.5×10^{12} eV [36]. Such experiments cannot accelerate particles to the highest known energies, and thus cannot study how they interact at such energy scales.

Astrophysical objects are thought to be able to produce particles of such energies. Studying these ultra-high energy ($> 10^{18}$ eV) particles presents the opportunity to understand how particles may acquire such high energies, and may even be the key to unlock the door to new physical discoveries.

Chapter 3

Ultra-High Energy Astroparticle Physics

The Earth's atmosphere is constantly bombarded by high energy cosmic rays (CRs), which produce extensive showers of secondary particles. Since their discovery in the early 20th century [37], much has been found out about cosmic rays, however, the ultra-high energy regime of such particles is not completely understood. Although ultra-high energy ($> 10^{18}$ eV) cosmic rays (UHECRs) have been detected [38], it is still unknown where they come from, and how they get to ultra-high energies. Furthermore, the highest energy CRs should interact with the Cosmic Microwave Background (CMB), producing ultra-high energy neutrinos (UHENs) as an end product.

3.1 Ultra-High Energy Cosmic Rays

3.1.1 Acceleration

In order for cosmic rays to gain high, or even ultra-high energies, there needs to be an associated acceleration mechanism. Only the most extreme astophysical phenomena could lead to particles of such high energies.

A logical first step would be to consider charged particles colliding with interstellar media (moving magnetic clouds of gas). The charged particles would reflect due to the presence of the magnetic cloud (in the simplest terms, a "magnetic mirror"), which would either approach or recede. As the particle approaches the mirror,

the probability of a head-on collision will be much larger than a head-tail collision. Thus, the particle will, on average, gain energy. The change in energy will be proportional to the squared velocity of the magnetic mirror, leading to a spectrum with the power law (see Appendix A.1):

$$\frac{dN}{dE} \propto E^{-\left(\frac{1}{\alpha T_{\rm esc}}+1\right)}, \tag{3.1}$$

where N is the number of particles, E is the particle energy, α is a constant that arises from a combination of the velocity of the magnetic mirror and the path length of the particle, and $T_{\rm esc}$ is the escape time.

This is known as second order Fermi acceleration [39], as the energy gain in this scenario is proportional to the mirror velocity squared. However, second order Fermi acceleration suffers from several problems. First and foremost, it does not predict the global power law as seen in the spectra of energetic particles (see Section 3.1.2).

Consider a stellar explosion that results in a shockwave. The temperature and density of the interstellar medium almost instantly changes when hit by the shock. A proton can pass through the shock as the shock approaches it, and can pass back through it again due to being reflected by the tangled magnetic fields. This cycle continues, and each cycle, the proton gains more and more energy, until it escapes the system.

This mechanism is known as first order Fermi acceleration [40]. In this scenario, the energy gain is proportional to the shock velocity (i.e., first order), and results in the global power law spectrum (see Appendix A.2):

$$\frac{dN}{dE} \propto E^{-2}. \tag{3.2}$$

3.1.2 Spectrum

The flux of cosmic rays, as detected on Earth, decreases with increasing particle energy via the power-law,

$$\frac{dN}{dE} \propto E^{-\gamma}, \tag{3.3}$$

where γ is the spectral index (which is typically around 2.7).

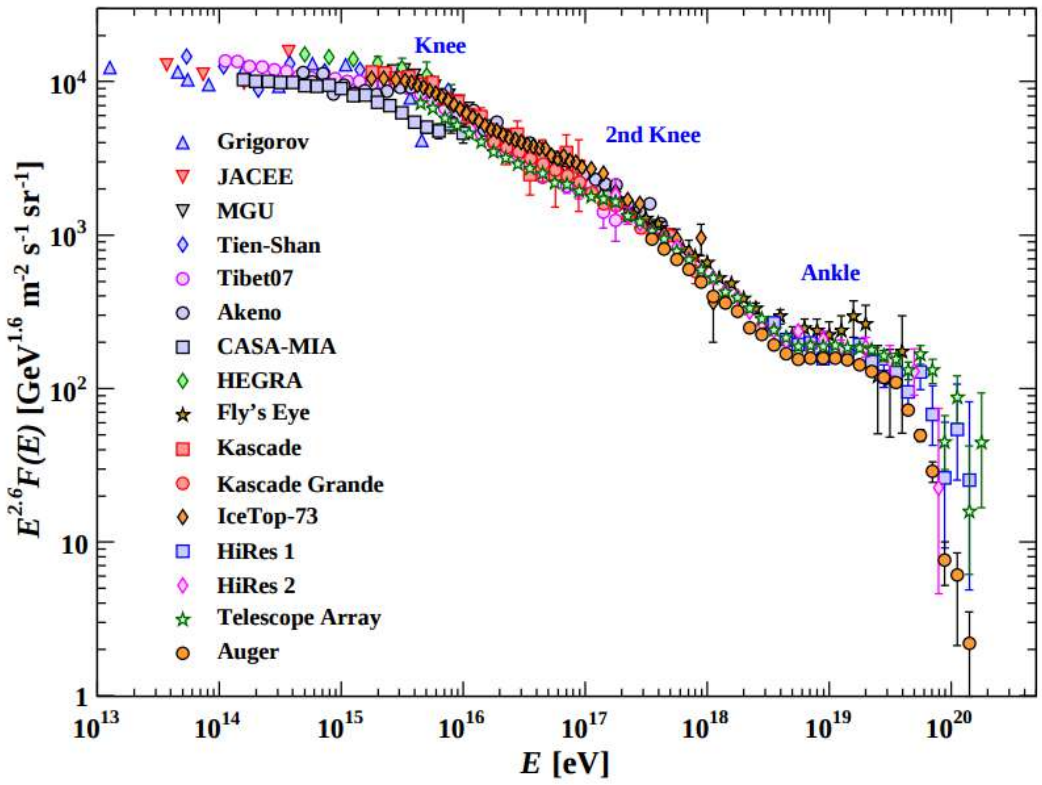

Figure 3.1: The high energy cosmic ray spectrum, showing the knees, ankle and GZK cut-off. Image from [41].

However, at certain energies, the gradient of the spectral slope changes. These observable features (as shown in Figure 3.1) are known as the *knees*, and the *ankle*. At energies even higher than this, the spectrum appears to rapidly fall off. These changes of the spectral index must be due to a change in production, source distribution, and propagation of cosmic rays at certain energies [41].

3.1.2.1 Features

Knees

The cosmic ray spectrum features a *knee* (at $\simeq 3 \times 10^{15}$ eV), as well as a *second knee* or *iron knee* (at $\simeq 10^{17}$ eV).

The origins of the knee features are still widely discussed, though there are

two separate scenarios to explain them. The first of which is the propagation origin, which suggests that very high energy particles are unable to be contained by the magnetic fields of our galaxy. Equating the magnitude of the centripetal force to the Lorentz force for a particle in our galaxy, and assuming that the speed of the cosmic ray is approximately that of light,

$$|q|cB = \frac{mc^2}{r_{\mathrm{L}}}, \tag{3.4}$$

where q is the charge of the particle, c is the speed of light, B is the magnitude of the magnetic field, m is the mass of the partcle, and r_{L} is Larmor radius. The charge is simply the product of the elementary charge, e, and the proton number, Z, and the length is that of the galactic arm, r_{g}:

$$E = ZeBr_{\mathrm{g}}c\,. \tag{3.5}$$

Assuming a background magnetic field of 10^{-9} T, and that the length of the galactic arm is 10^{16} m, for protons (Z is simply 1 in this case):

$$E_{\mathrm{H}} = 1.6\times10^{19}\ \mathrm{C}\cdot10^{-9}\ \mathrm{T}\cdot10^{16}\ \mathrm{m}\cdot3\times10^{8}\ \mathrm{ms}^{-1}, \tag{3.6}$$

$$E_{\mathrm{H}} = 4.8\times10^{-4}\ \mathrm{J}, \tag{3.7}$$

$$E_{\mathrm{H}} = 3\times10^{15}\ \mathrm{eV}\,. \tag{3.8}$$

For iron, the calculation is the same, but Z = 28, and so,

$$E_{\mathrm{Fe}} \simeq 10^{17}\ \mathrm{eV}\,. \tag{3.9}$$

The propagation scheme thus shows that beyond the length of the galactic arm, cosmic rays of high energies can escape the weak galactic magnetic fields, leading to a reduced measured flux on Earth.

The second explanation of the knees comes from a direct limit on the maximum energy a cosmic ray can acquire through acceleration mechanisms (i.e., first

order Fermi acceleration). The shock fronts of objects in our galaxy are unable to accelerate cosmic rays to energies above the knee, hence the steepening of the flux. The Larmor radii of the particles must be smaller than the size of the source itself (R), in order to be accelerated by the source's magnetic field without escaping (B). These factors (i.e., a combination of magnetic field strength, characteristic size, and others) make up the requirements of a source to accelerate cosmic rays to high energy levels [42]:

$$E_{\max} = \eta^{-1} \beta_{\mathrm{sh}} eBR\Gamma, \tag{3.10}$$

where $E_{\max}$ is the maximum energy of a particle accelerated by the source, η is the acceleration efficiency, β_{sh} is the shock velocity, and Γ is Lorentz factor of the relativistic flow.

This combination of factors is known as the "Hillas condition" [42], and is shown diagrammatically in Figure 3.2. Clearly, the Hillas condition shows that it is very challenging for galactic sources to be able to accelerate protons (and iron) to energies beyond the knee (and second knee), due to the high strength of the magnetic fields required. It is believed that Supernova Remnants (SNRs) are responsible for accelerating cosmic rays to energies up to the knee. Even so, there are constraints set on the strength of the magnetic fields of SNRs, and so there is also a constraint on the maximum energy a cosmic ray can acquire from them [43].

Ankle

Following the knee features, the cosmic ray spectrum seems to become flatter at around 5×10^{18} eV. This is thought to be due to the transition from galactic to extragalactic cosmic rays [44, 45]. Figure 3.3 shows how an extragalactic component would contribute to the cosmic ray flux spectrum.

3.1.2.2 The End of the Cosmic Ray Spectrum

At ultra-high energies, the cosmic ray spectrum seems to trail off, to have essentially no observable flux, and this is thought to be due to the Greisen-Zatsepin-Kuzmin (GZK) limit [46, 47].

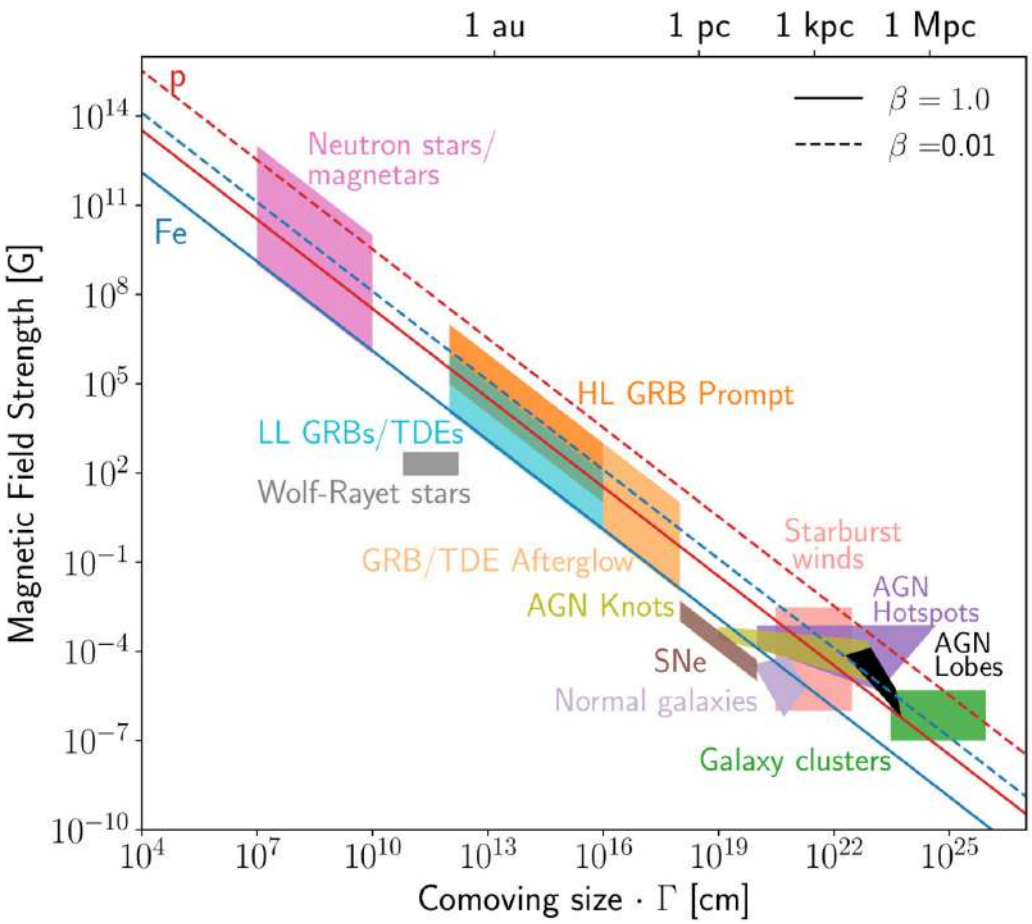

Figure 3.2: A Hillas diagram showing the different types of candidate objects and mechanisms that could produce UHECRs [42]. In order to be able to accelerate cosmic rays to high enough energies, an object must lie above the diagonal lines, each depicting different scenarios. A blue line indicates the combination of magnetic field strength and size needed to accelerate iron nuclei to ultra-high energies, whilst a red line indicates the same for protons. The line associated with iron nuclei is lower than that of protons, as it is easier to accelerate particles of higher charge (proportional to proton number, Z). A solid line represents a shock velocity of $\beta = 1$, whilst the dashed line is the same but for $\beta = 0.01$ (i.e., the combination is more stringent for the dashed line, as a lower shock velocity will require sources of larger magnetic field strength or size to boost the cosmic rays to ultra-high energies). The diagonal lines associated with the Hillas condition assume that $\eta = 1$ from Equation 3.10. Image from [56].

The GZK limit comes from considering interactions of ultra-high energy cosmic rays with cosmic microwave background (CMB) photons, which then produce pions and nucleons via a Δ^+ resonance. The pions and nucleons then decay to many

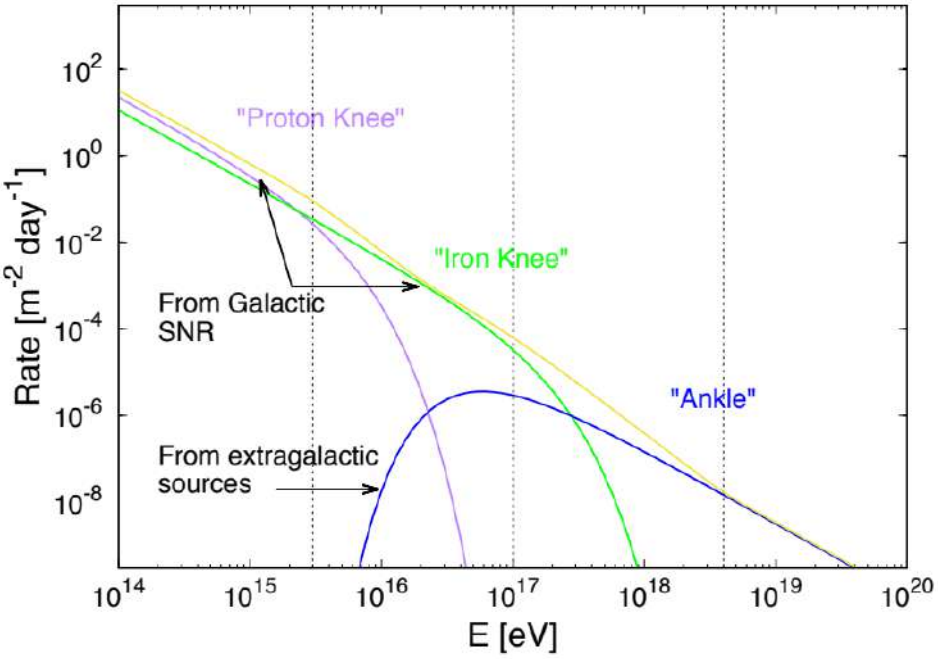

Figure 3.3: The rate of cosmic rays arriving on Earth versus the energy of the cosmic rays. Three features are observed: the two knees and the ankle. As well as the total rate, three separate rates are shown for the individual components as a function of energy. Image modified from [56].

particles, including neutrinos [48]:

$$p_{\mathrm{CR}} + \gamma_{\mathrm{CMB}} \rightarrow \Delta^+ \rightarrow n + \pi^+$$
$$n \rightarrow p + e^- + \overline{\nu_e}$$
$$\pi^+ \rightarrow \mu^+ + \nu_\mu$$
$$\mu^+ \rightarrow e^+ + \nu_e + \overline{\nu_\mu}$$

The Delta baryon can also decay into a proton and a neutral pion, instead of a neutron and a charged pion. Taking into account the energies of the particles in the primary interaction, we can find the energy associated with the GZK limit. First, the square of the 4-momenta of the proton, p_p, and CMB photon, p_γ, is equated to the square of the 4-momentum of the Delta baryon, p_Δ:

$$(p_p + p_\gamma)^2 = p_\Delta^2 = M_\Delta^2 . \tag{3.11}$$

where M_Δ is the mass of the Delta baryon, and the following relation is used, $p_\Delta^2 = M_\Delta^2$. Taking the dot product of the left side of Equation 3.11, and using the relations, $p_p^2 = M_p^2$, where M_p is the mass of the proton, and $p_\gamma^2 = 0$:

$$(p_p + p_\gamma)^2 = p_p^2 + 2p_p \cdot p_\gamma + p_\gamma^2, \tag{3.12}$$

$$= M_p^2 + 2p_p \cdot p_\gamma . \tag{3.13}$$

Now evaluate the dot product:

$$p_p \cdot p_\gamma = E_p E_\gamma - 2\vec{p}_p \cdot \vec{p}_\gamma, \tag{3.14}$$

$$= E_p E_\gamma - |\vec{p}_p||\vec{p}_\gamma|\cos(\theta), \tag{3.15}$$

where E_p is the energy of the proton, E_γ is the energy of the photon, $\vec{p}_p$ is the momentum of the proton, $\vec{p}_\gamma$ is the momentum of the photon, and θ is the associated angle between the proton and photon. Taking that angle to be 180°,

$$p_p \cdot p_\gamma = E_p E_\gamma + |\vec{p}_p||\vec{p}_\gamma| . \tag{3.16}$$

As the mass of the photon is zero, $|\vec{p}_\gamma| = E_\gamma$. Taking the relativistic limit for the proton, its energy will be much higher than its mass, $E_p \gg m_p$, and thus approximately $|\vec{p}_p| = E_p$. Therefore,

$$p_p \cdot p_\gamma = E_p E_\gamma + E_p E_\gamma, \tag{3.17}$$

$$= 2E_p E_\gamma\,. \tag{3.18}$$

Inserting this back into Equation 3.13, we arrive at,

$$(p_p + p_\gamma)^2 = M_p^2 + 4E_p E_\gamma\,. \tag{3.19}$$

Thus, 3.11 becomes,

$$M_p^2 + 4E_p E_\gamma = M_\Delta^2, \tag{3.20}$$

and so,

$$E_p = \frac{M_\Delta^2 - M_p^2}{4E_\gamma}\,. \tag{3.21}$$

Using $E_\gamma = k_B T$, we find the energy cut-off of the cosmic ray spectrum to be,

$$E_p = \frac{M_\Delta^2 - M_p^2}{4k_B T}, \tag{3.22}$$

where k_B is the Boltzmann constant, and T is temperature.

Inserting the values:

$$E_p = \frac{(1232\ \text{MeV})^2 - (938.3\ \text{MeV})^2}{4 \cdot 8.6 \times 10^{-11}\ \text{MeV/K} \cdot 2.7\ \text{K}} \simeq 7 \times 10^{20}\ \text{eV}\,. \tag{3.23}$$

Thus, it is not expected to detect cosmic rays above this energy limit[1], as they

[1]Taking into account the far more energetic CMB photons at the tail of the Maxwell-Boltzmann distribution, the limit becomes approximately 5×10^{19} eV. This increases the denominator of the calculation, leading to a lower total energy.

would readily interact with the CMB and decay to lower energy particles.

3.1.3 GZK Neutrinos

Many of the decay products of UHE proton-photon interactions are electrically-charged, so their paths are influenced by magnetic fields, and thus cannot provide information about the origin of these UHE cosmic rays. As photons start to pair-produce electrons and positrons at $\sim 10^{14}$ eV, the photon horizon is largely limited at these energy scales.

However, cosmogenic neutrinos move through the Universe unattenuated on their way to Earth, and are electrically neutral, so they can trace directly back to the cosmic ray origin. Furthermore, these neutrinos see the Universe as transparent up to an energy of 10^{25} eV (see Figure 3.4) and potentially beyond. Because they stem directly from UHECRs near the GZK cut-off, the neutrinos are referred to as GZK neutrinos[2] [46, 47].

3.1.4 Composition

Detectors such as the Pierre Auger Observatory are able to determine the approximate composition of UHECRs. Arrays of detectors detect the energies of the charged particles in cosmic ray air showers. Combining in-air measurements from flourescence detectors with on-ground measurements allows the calculation of the total energy loss as a function of average depth of the shower maximum, $\langle X_{\max} \rangle$. At any point along the cosmic ray's path, there will be a number of interactions proportional to the density of the medium (g cm^{-3}) and its path length (cm). Heavier nuclei, such as those of iron, will cause a cosmic ray shower higher in the atmosphere,[3] which has a far lower density. Therefore, we expect $\langle X_{\max} \rangle$ to be lower for heavier nuclei, and can use this to determine the composition of cosmic rays. As seen in Figure 3.5, the average depth of the shower maximum seems to decrease

[2] Berezinsky and Zatsepin first realised that ultra-high energy neutrinos must be a decay product of the GZK effect, as charged pions are produced from the decay of a Delta baryon. Thus, these neutrinos are sometimes referred to as cosmogenic (or BZ) neutrinos [48].

[3] The cross section of a cosmic ray with a particle in the atmosphere depends on the nucleon number of the cosmic ray's nucleus (A). The nuclear interaction of a heavy cosmic ray nucleus (high A) can be thought of as a superposition of A individual nucleons each undergoing an interaction with $1/A$ of the cosmic ray's energy. Thus, heavier cosmic rays (iron nuclei), will interact much earlier in the atmosphere than lighter cosmic rays (protons).

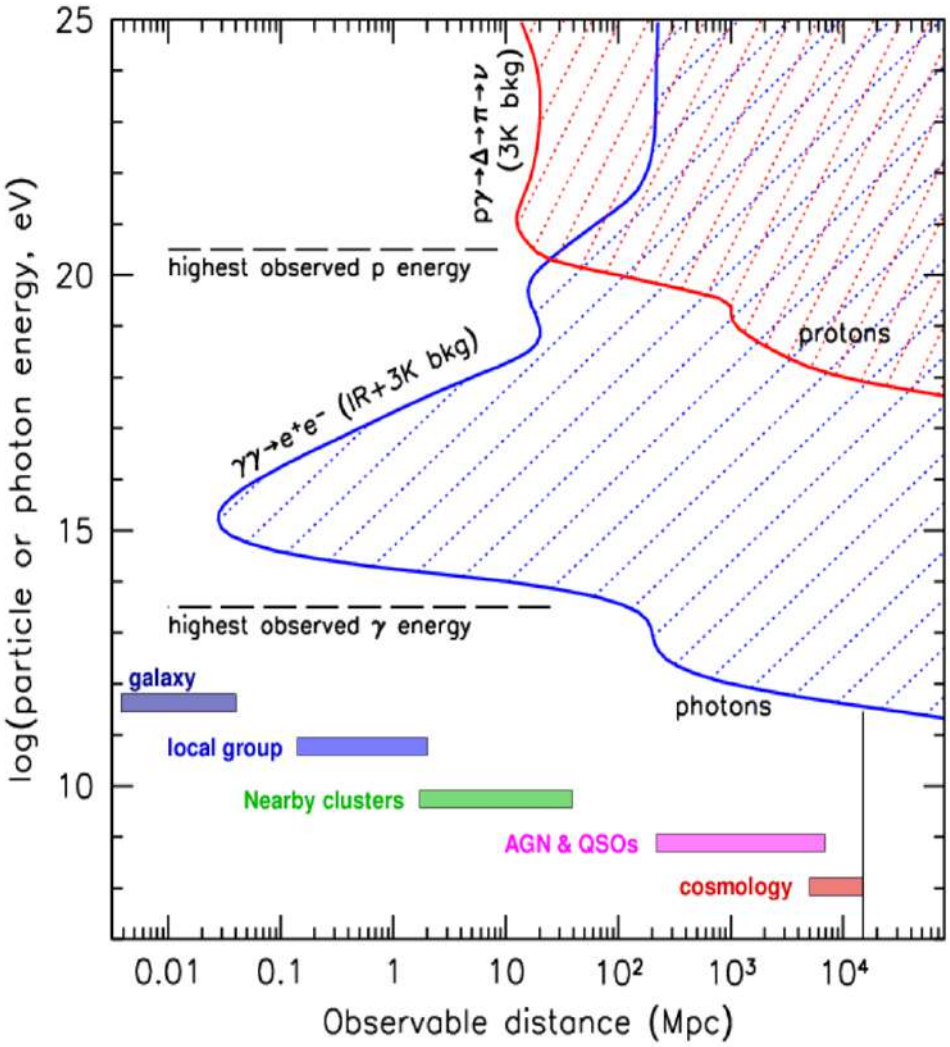

Figure 3.4: The observable distance of specific particles as they travel through the Universe. The red region is not observable with protons, and the blue region is not observable with photons. Note that there is no indication of neutrinos on this plot, as they are observable within all energies and distances as shown. Distances of some classes of astrophysical objects are shown, such as those of galaxies and quasi-stellar objects (quasars). Diagram produced by Peter Gorham.

when the energies become ultra-high, suggesting that cosmic rays have mixed compositions [49].

Cosmic rays that consist of nuclei heavier than hydrogen (protons), can undergo photodisintegration when they interact with the CMB. In this case, instead of producing nucleons and pions that can themselves go on to form neutrinos, nucleons are ejected:

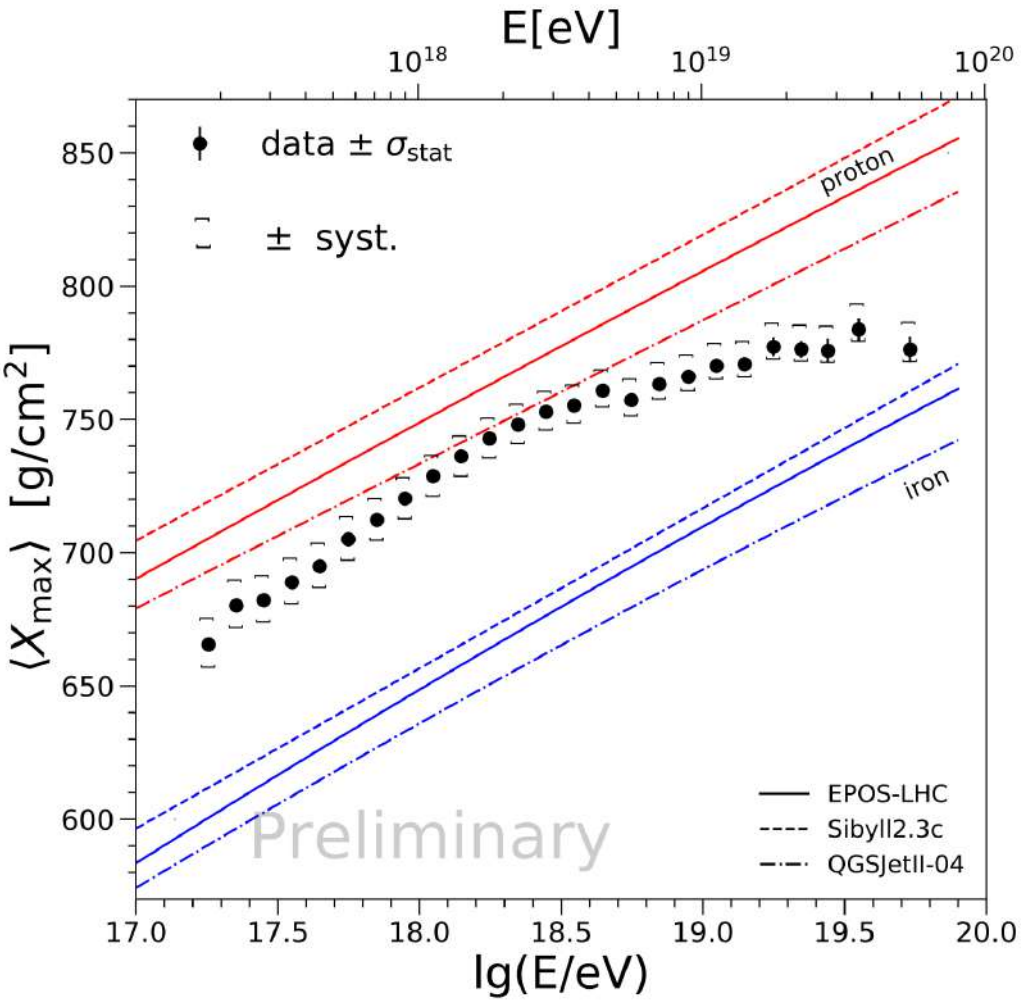

Figure 3.5: Preliminary Auger composition data, showing the average depth of the shower maximum vs the energy of the shower. Whilst CRs at the low energies of the UHE range appear to be mostly protonic, those at the high end appear to have a mixed composition. Image from [49].

$$N_{CR} + \gamma_{CMB} \rightarrow N^*_{CR} + X,$$

where N is the cosmic ray nucleus before interaction, N^* is the remaining cosmic ray nucleus, and X represents the ejected nucleons. In such an interaction, the mass number of the nucleus is decreased. It is still possible for neutrinos to be produced as before, but it is no longer the dominant mechanism. The neutrinos produced by the interactions of heavier nuclei (such as iron) are of lower energies

than those produced by the interactions of protons. However, only heavy nuclei undergo photodisintregration, meaning that this effect does not occur in protons travelling through the cosmos.

3.1.5 Potential Sources

There are many proposals for the source of these UHE cosmic rays, such as blazars, supernovae (SNe), and gamma-ray bursts (GRBs). At the centre of several galaxies are supermassive black holes, which are active galactic nuclei (AGN), and many have relativistic particle jets. Blazars are a class of AGN that have these jets pointed towards Earth—it is thought that certain processes cause protons to gain ultra-high energies [50]. SNe, highly energetic and luminous explosions of stars, could be a potential source of ultra-high energy particles. Due to the properties of neutrinos, it is likely that these particles would be able to escape before the light of the supernovae. GRBs are short-lived, but massively powerful explosions thought to originate from violent astrophysical events, such as the collision of two neutron stars, to form a black hole. GRBs should be able to accelerate high-energy hadrons, which would then cool to produce UHECRs [51]. There are also some other suggestions for the source of UHECRs. These include particles undergoing the Penrose effect [52] (where energy is extracted from a rotating black hole) and decays of dark matter particles [53].

3.1.5.1 Blazars

Quasars, also termed quasi-stellar objects or quasi-stellar radio sources, are a type of extremely luminous AGN. There are several classifications of quasars. BL Lacertae (BL Lac) objects[4] are radio-loud quasars with almost featureless non-thermal emission spectra across all wavelengths. Optically violent variable (OVV) quasars, now known as flat-spectrum radio quasars (FSRQ), on the other hand, are highly variable, and have strong emission lines. Blazars[5] are quasars that have their relativistic particle jets pointing directly towards Earth, thus appearing incredibly bright.

[4]The name, BL Lac, is based on the object of the same name, BL Lacertae. The original BL Lacertae object was first thought to be a variable star, and so was assigned a variable star designation, BL (the names of which follow an archaic convention; BL is not an acronym).

[5]Blazar is a contraction of BL Lac and (OVV) quasar.

These particle jets are formed when accretion disks form around rotating black holes. As the matter of the disk gets closer to a black hole, the particles within it rapidly heat up and ionise, which creates extremely powerful magnetic fields. The ionised matter is then ejected as jets perpendicular to the plane of rotation. As such, these relativistic jets are thought to be a potential source of UHECRs [54].

3.1.5.2 Supernovae

Supernovae (SNe), highly luminous stellar explosions, can occur at the end of a star's life cycle. During the explosion, a shock wave expands into the interstellar medium (ISM). The shock wave compresses and heats up the gas of the ISM, and this compressed mass rapidly moves radially outward. The compressed matter pushes back on the shockwave, forming a reverse shock that propagates into the inner shell. Between the shocked ISM and ejecta, a contact discontinuity is formed, whereby the pressure is approximately constant. The reverse shock then reheats the ejecta. Supernova Remnants (SNRs) are thought to be the origin of galactic CRs, up to the energies associated with the *knee*, and potentially the ankle. However, galactic SNRs are very unlikely to produce cosmic rays with ultra-high energies.

Some types of SNe may be potential sources of UHECRs and their associated neutrinos. Non-relativistic SNe, however, do not have strong enough magnetic fields or large enough physical span to accelerate cosmic rays to sufficiently high energies, even in the scenario of heavier nuclei (See Figure 3.2).

That said, relativistic SNe may contribute to the flux of UHECRs [57–59]. Relativistic supernovae are supernovae that exhibit relativistic ejecta, and are often associated with Gamma Ray Bursts (GRBs).

3.1.5.3 Gamma Ray Bursts

Gamma Ray Bursts (GRBs) are extremely short bursts of high energy gamma rays. These bursts can last from about a hundredth of a second to hundreds of seconds (see Figure 3.6 for a time distribution of GRBs), but are generally classified as short-duration (< 2 s) and long-duration (> 2 s) bursts [55]. When the bursts occur, they are the brightest sources of gamma rays in the observable Universe (about a hundred times brighter than SNe). Short duration bursts are believed to be a result

of two compact objects (such as binary neutron stars or a neutron star and a black hole) colliding. Long duration bursts are thought to arise from red giants collapsing into their own cores. Specifically, the duration refers to the variable, T_{90}, which is different for each GRB. T_{90} represents 90% of the total counts of a GRB, i.e., the duration in which the GRB counts increase from 5 to 95% above the background. After the initial "prompt" burst, an "afterglow" of lower-frequency waves follows. The diverse mechanisms leading to both short and long duration GRBs give rise to a number of potential sites of UHE particle production: the initial explosion of the star or collision of two stellar objects, the resulting jets, or within the reverse shock during the afterglow stage.

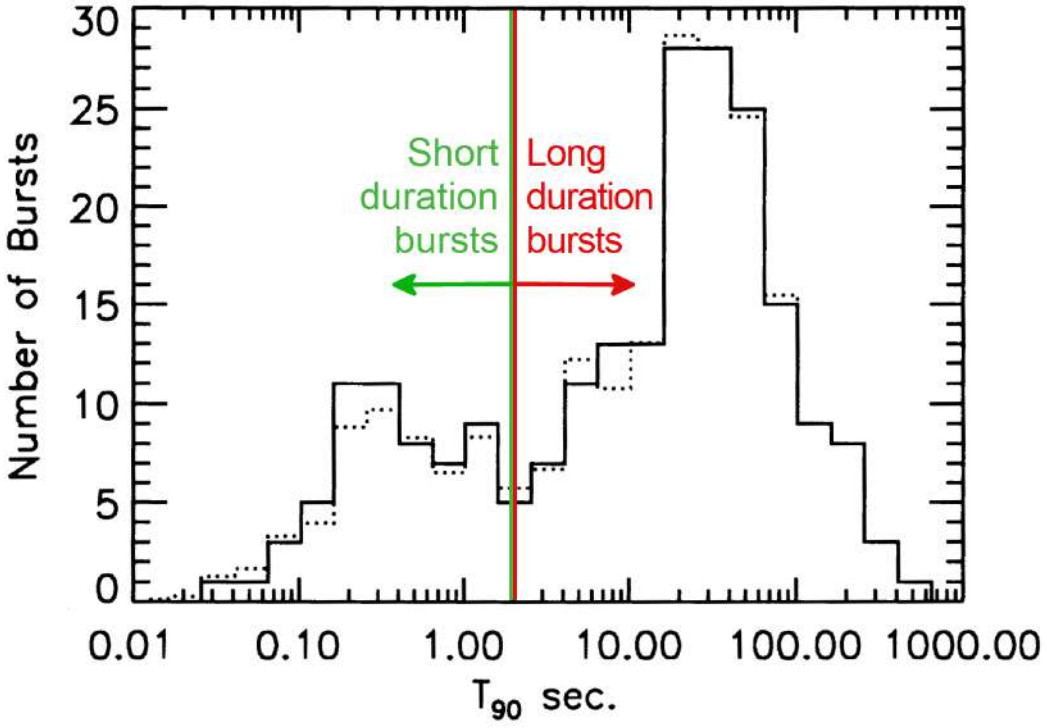

Figure 3.6: T_{90} distribution for the first Burst and Transient Source Experiment (BATSE) catalogue. The vertical line shows the division between short-duration and long-duration bursts. Image modified from [55].

In order to explain the presence and properties of both an initial prompt burst and the resulting afterglow, a GRB can be modelled as a relativistic "fireball" [60, 61]. A fireball, in this case, is essentially an optically thick plasma, whereby the plasma is made up of photons, electrons, and positrons. The highly en-

ergetic (and thus high temperature) plasma is confined to a compact shell, which is the eventual source of the immense energy output. If the fireball is instead made up of a large amount of baryons, almost the entire energy of the fireball would simply be converted into kinetic energy shared between the high number of baryons. Instead, if the fireball is almost completely made of photons, electrons, and positrons, and has only a small baryonic load, the small amount of baryons would be accelerated to ultra-high speeds [62, 63]. The relativistic fireball will expand under its own pressure, and thus decrease in temperature. When the temperature has dropped significantly, the electron-positron pairs annihilate into photons. As a large amount of pairs have now converted into photons, the plasma is no longer opaque, allowing the photons to escape at enormous energies. When emitted particles catch up with those slower than them and collide, energy is converted to radiation from the internal shock. This internal shock gives rise to the initial burst fluence of photons [64–66].

Fireball models also take into account external shocks, whereby the shell collides with something outside of the fireball, such as interstellar media (ISM) [67]. This external shock gives rise to the afterglow photon fluence. However, photons aren't the only particles produced from these extremely luminous explosions. Binary neutron star merger events are predicted to release $\approx 5 \times 10^{46}$ J of energy, mainly in the form of neutrinos and gravitational waves [68]. As it is possible that fireballs could contain a small baryon load, protons or heavier nuclei (UHECRs) could be accelerated to very high velocities. These particles would then go on to interact with the CMB, forming neutrinos in the process.

Chapter 4

The Askaryan Effect

As described in Section 3, ultra-high energy neutrinos (UHENs) can be produced by a range of processes, and subsequently make their way to Earth. One needs to determine how to detect such elusive particles once they arrive.

4.1 Induced Electromagnetic Cascade

When a UHEN moves through a medium, it can interact with the particles that constitute the medium. Specifically, neutrinos can undergo two types of weak in-

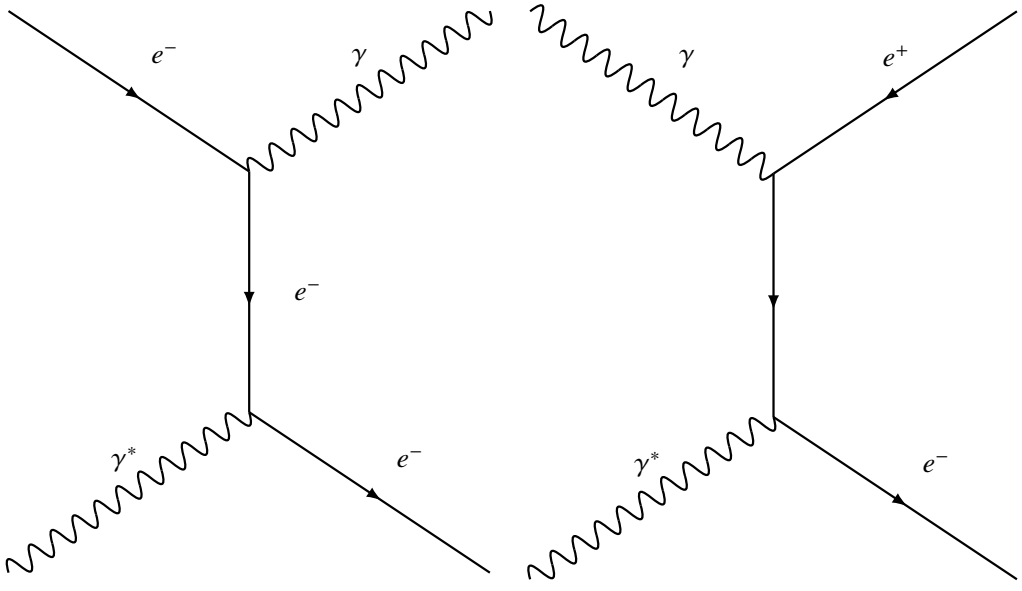

(a) Bremsstrahlung. **(b)** Electron-positron pair production.

Figure 4.1: Interactions involved in the core formation of the electromagnetic cascade. γ^* denotes a virtual photon.

teractions: charged current (CC) and neutral current (NC) (see Section 2.2.3). Both CC and NC interactions can produce initial hadronic showers when the neutrino interacts with a quark, though only electron-neutrino CC interactions can produce an initial electromagnetic cascade. In the case of an NC interaction, a charged lepton cannot be initially produced, but instead a hadronic shower can arise. However, interactions of particles in the hadronic cascade within the medium produce neutral pions (which, in turn, decay to photons that pair produce), which takes a signif-

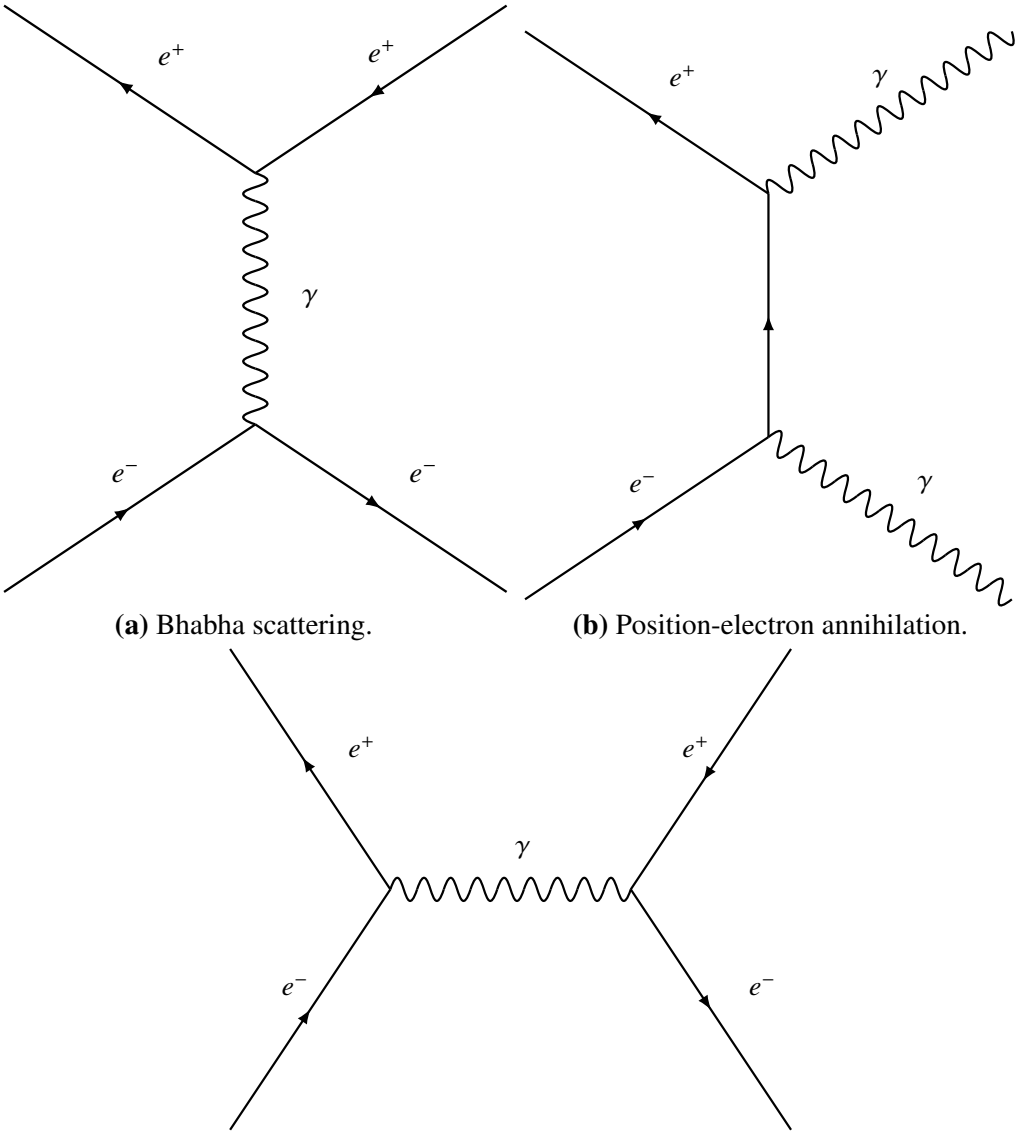

(a) Bhabha scattering.

(b) Position-electron annihilation.

(c) Bhabha annihilation.

Figure 4.2: Interactions that decrease the ratio of positrons to electrons in the shower. Note that the atomic electron (bottom left of each Feynman diagram) is at rest.

icant amount of the energy per interaction. As the incident energy increases, the net energy fraction of the hadronic component that is transferred to pions, and thus the electromagnetic component, increases. Thus, for ultra-high energy scenarios, almost all of hadronic shower's energy ends up in the electromagnetic sector.

As an induced electromagnetic shower arises no matter the starting interaction, we can take a look at how such a shower develops. From an electron-neutrino CC interaction (or from the decay of photons that arise in the hadronic cascade), the particles that are produced, electrons and positrons, go on to develop an electromagnetic shower. The electrons in this shower give off Bremsstrahlung photons, which themselves go on to pair produce to electrons and positrons, increasing the number of particles in the shower (see Figure 4.1). The particles produced can interact with atomic electrons, increasing the number of electrons in the shower relative to the number of positrons. Such interactions are electron-positron annihilation, as well as Bhabha, Møller, and Compton processes.

Electron-positron annihilation is when positrons interact with electrons, giving rise to photons. In this case, the positrons and atomic electrons annihilate, removing the positronic component from the shower. Bhabha scattering liberates an electron from an atom when a positron interacts with the atomic electron itself (by exchanging a photon). See Figure 4.2 for these processes. Due to these interactions, the positrons no longer contribute to the moving charge bundle as much.

The other particles that arise from the above interactions are photons and electrons. As in Bhabha scattering, both Møller (electron-electron scattering) and Compton (photon-electron scattering) processes liberate atomic electrons (see Figure 4.3 for these processes). These electrons are swept into the electromagnetic shower, and thus a net negative charge of about 20% builds up. The positive ions (from which the electrons were liberated) are left behind as the negatively charged particle bunch moves through the medium.

4.2 Cherenkov Radiation

When in a medium, a moving charged particle emits electromagnetic waves spherically. If its speed is less than the speed of light, its wavefronts do not overlap.

Now, consider that the charged particle is highly energetic. In the case that its speed exceeds that in which light travels through the medium, these waves will overlap. The scenario for the particle being slower than, and faster than, the speed of light in a medium is shown in Figure 4.4. The Huygens–Fresnel principle states that every point on a wavefront itself is a source of wavelets. Thus, each of these spherical waves interfere to form a new wavefront. The angle of the emitted coherent radiation with respect to the shower axis is known as the Cherenkov angle, which in turn gives rise to a Cherenkov cone (as shown in Figure 4.5). The Cherenkov angle is defined as:

$$\cos(\theta_{\text{ch}}) = \frac{1}{n\beta}, \tag{4.1}$$

where θ_{ch} is the Cherenkov angle, n is the refractive index of the medium, and β is

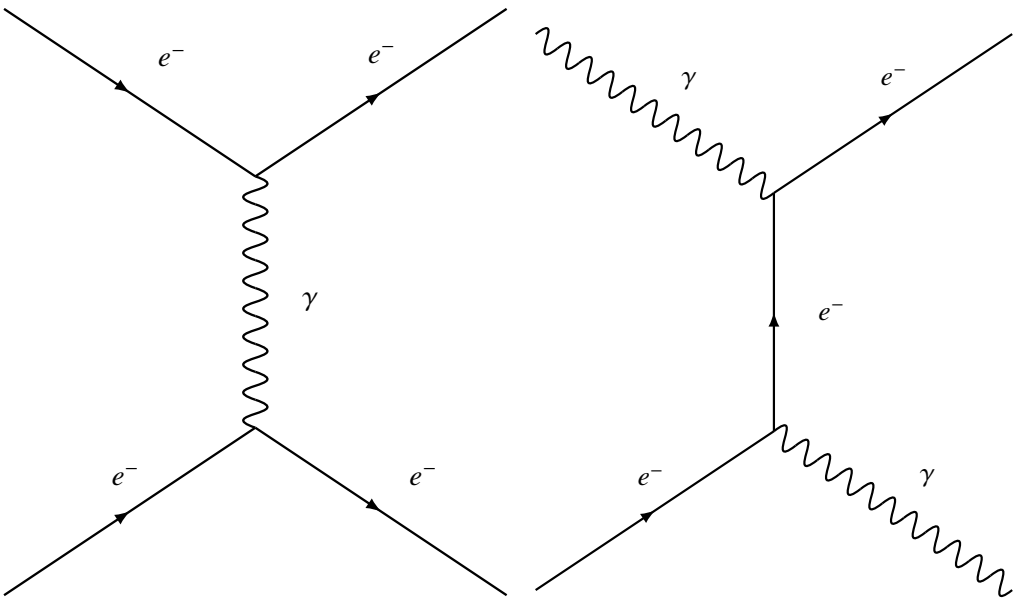

(a) Møller electron-electron scattering. **(b)** Compton electron-photon scattering.

Figure 4.3: Interactions that increase the ratio of electrons to positrons in the shower. Only a single diagram for each possible process is shown i.e., the *t-channel* is displayed for the scattering processes. Note that the atomic electron (bottom left of each Feynman diagram) is at rest.

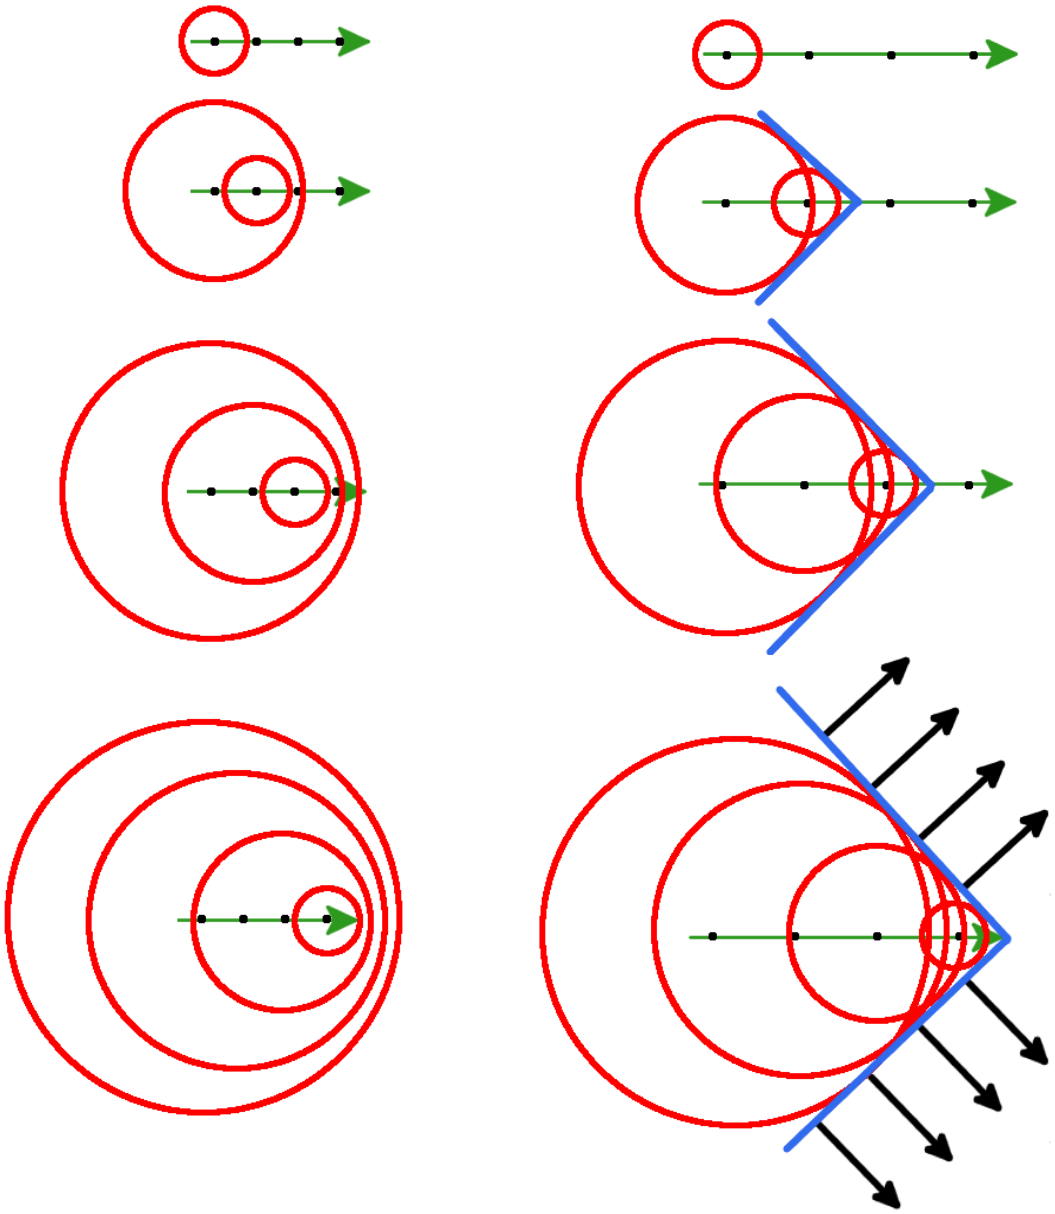

(a) A charged particle travelling slower than the speed of light. There is no interference as the waves do not overlap.

(b) A charged particle travelling faster than the speed of light in a medium. There is interference, a wavefront forms, emitting radiation, which will form in the shape of a cone.

Figure 4.4: Illustration of how a wavefront is produced from a moving charged particle. The green arrow represents the full path of the charged particle, whereas the black dots represent a single point in time when the spherical waves are emitted. The red rings represent the spherical waves, which overlap in case (b). The blue lines represent the wavefront, and the black lines indicate the direction in which the wavefront is moving.

the ratio of the speed of the particle to that of the speed of light.

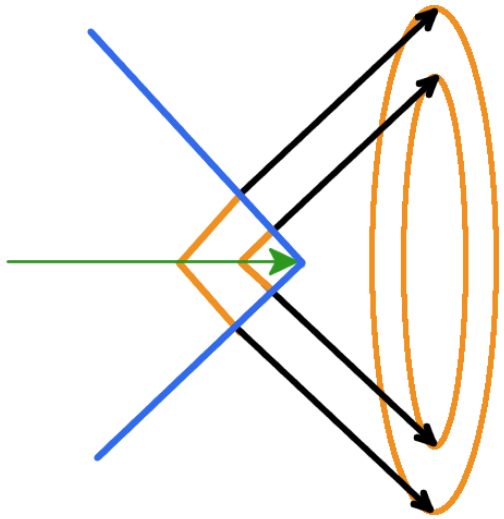

Figure 4.5: Side view of the shower axis. The orange circles represent the base of the Cherenkov cone. The figure is slightly tilted into the page to reveal the three-dimensional cone shape.

4.3 Coherent Radio Emission

In Section 4.1, it was determined that a UHEN can give rise to a net negatively charged particle bunch when moving through a medium. Allowing the rapidly moving particle bunch to be considered as a single particle, this is the same scenario as detailed in Section 4.2. In the case of a UHEN-induced electromagnetic cascade, the size of the particle bunch is a lot smaller than the Cherenkov radio wavelengths from all the shower particles[1], giving rise to constructive interference. The emitted radio waves are thus coherent, i.e., they have a fixed phase difference. The process of a net charge being produced by a neutral particle, in turn producing coherent radio waves, is specifically known as the Askaryan effect [69]. In the case of wavelengths being a lot smaller than the bunch size (such as wavelengths of visible light), the wavelengths are not added coherently, leading to non-coherent emission,

[1] In dense media, such as ice, the shower is characterised in the transverse direction by its Molière radius, which is confined to only a few centimetres. Radio waves are the only part of the electromagnetic spectrum with wavelengths larger than this bunch size.

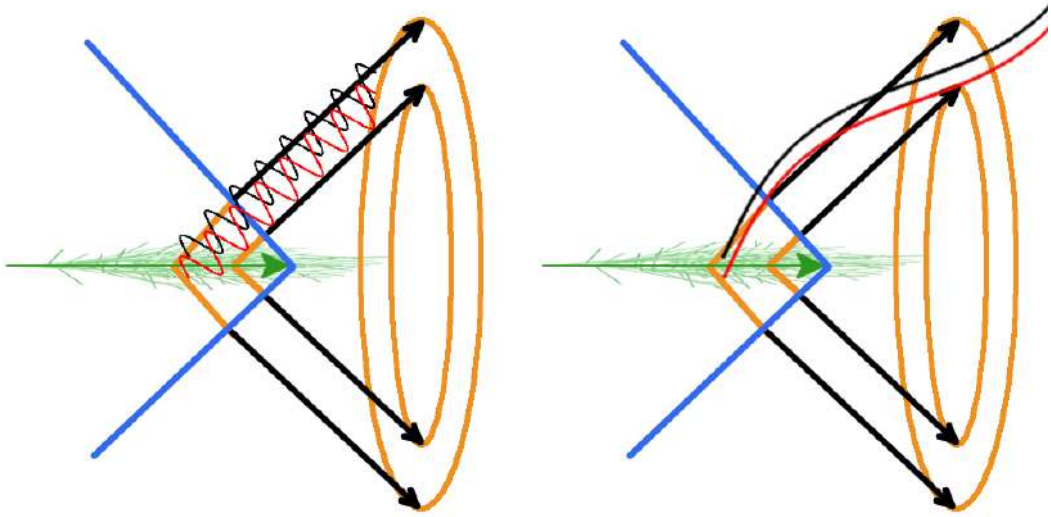

(a) When the wavelengths of the emitted photons are smaller than the bunch size, the phase difference between them is random. This leads to incoherent radiation.

(b) When the wavelengths of the emitted photons are much larger than the bunch size, the phase difference is negligible. This leads to coherent radiation, as the wavelengths constructively interfere.

Figure 4.6: An illustration of the photon emission at different wavelengths. The black and red waves indicate the wavelengths of the photons. The green translucent lines indicate the particle shower.

and thus this coherent Askaryan radiation is not produced (both cases are shown in Figure 4.6).

As the wavefronts move out, coherent Askaryan radiation is emitted, which is radially polarised along the cone. Because the electric field is radial around the shower axis, the upper-most part of the Cherenkov cone will be predominantly vertically polarised (as shown in Figure 4.7). As a net *negative* charge is considered, the electric field lines will point towards the source of the charge.

4.4 Refraction

An *ice-skimming* neutrino can give rise to an electromagnetic cascade, which in turn gives rise to coherent radio waves (Askaryan radiation). The radio waves make their way to the surface of the ice. Those at the top of the Cherenkov cone refract, allowing them to pass through the ice, whilst remaining predominantly vertically polarised. Further to the center of the cone (thus becoming more horizontally po-

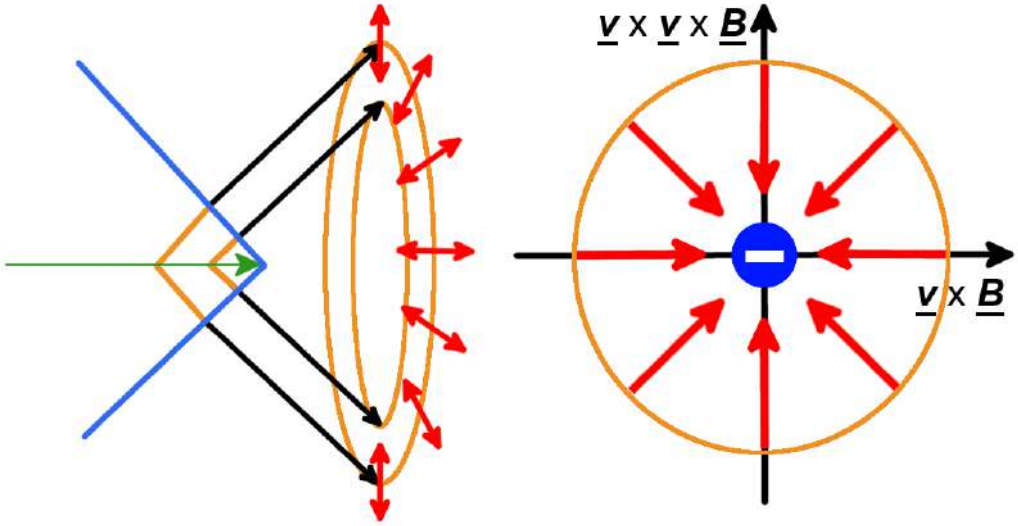

(a) Side view: shower axis goes from left to right. Emitted polarisation is predominantly vertical at the top of the cone. The red arrows indicate the direction of the polarisation. The figure is slightly tilted into the page to reveal the three-dimensional cone shape.

(b) Front view: shower axis points out of the page. It is clear that the radiation is emitted radially, giving rise to predominantly vertical polarisation at the top of the cone. The filled blue circle represents the moving negative charge. The time variation of the moving net negatively charged particle bunch gives rise to a magnetic field. The horizontal axis is aligned with $\underline{\boldsymbol{v}} \times \underline{\boldsymbol{B}}$ and the vertical axis is aligned with $\underline{\boldsymbol{v}} \times \underline{\boldsymbol{v}} \times \underline{\boldsymbol{B}}$, $\underline{\boldsymbol{v}}$ is the direction of the shower, $\underline{\boldsymbol{B}}$ is the magnetic field.

Figure 4.7: An illustration of the polarisation of the emitted photons. The red arrows indicate that Askaryan radiation is radial to the shower axis, giving rise to predominantly vertical polarisation close to the top of the Cherenkov cone.

larised), the angle between the surface normal of the ice and the radio wave becomes larger. A small increase in this angle still allows the photons to refract, but at a certain point, it will reach a critical angle, whereby the angle of refraction becomes 90°. When the angle is larger than the critical angle, the photons undergo total internal reflection, meaning that no light is refracted. The vertically polarised light closer to the bottom of the cone passes back through the ice. See Figure 4.8 for each of these scenarios.

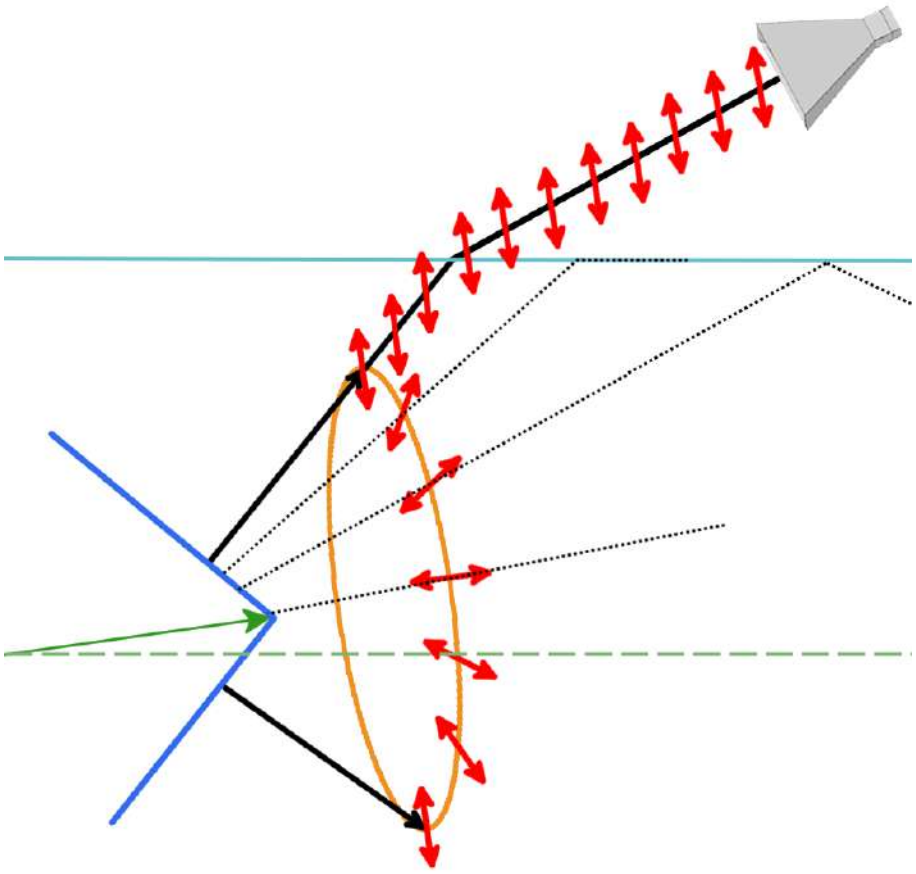

Figure 4.8: The full scenario: A neutrino *skims* the ice (shown as the green arrow at an angle relative to the dashed green line) at a small angle relative to the surface. A Cherenkov cone is formed, emitting coherent radiation of radio wavelengths, as these large wavelengths constructively interfere. Predominantly vertically polarised radio waves escape the ice (first thick black line), and are detected by an antenna. In the case of the critical angle (first dotted line from the top of the cone), the refraction angle becomes 90°. The second dotted line shows the case where the angle is so large that total internal reflection occurs. The red lines indicate the directions of polarisation, and the solid cyan line shows the surface of the ice.

4.5 Detection

The scenario for a neutrino giving rise to coherent vertically polarised radiation has been justified. A radio antenna pointed in a given direction can detect the radio wave, as shown in Figure 4.8.

The electric field of the Askaryan radiation will become attenuated as the signal moves from the ice to its point of detection (see Section 6.1.1.4 for more information). The electric field is calculated as to be,

$$|\epsilon| = |\epsilon_0^{1\mathrm{m}}| \cdot \exp\left(\frac{-d_{\mathrm{ice}}}{l_{\mathrm{ice}}}\right) \cdot \frac{1\mathrm{m}}{d_{\mathrm{ice}} + d_{\mathrm{air}}}, \tag{4.2}$$

where $|\epsilon|$ is the magnitude of the electric field at the detector, $|\epsilon_0^{1\mathrm{m}}|$ is the magnitude of the electric field 1 m from the shower, d_{ice} is the distance travelled in ice, l_{ice} is the attenuation length in ice, and d_{air} is the distance travelled in air.[2]

To detect these neutrino-induced radio waves from all azimuthal directions, an array of radio antennas can be utilised. This radio array would need to observe a large volume of an appropriate dielectric, such as ice, in order to maximise the chance of observing a neutrino interaction. Ideally, the experiment would take place in a radio-quiet environment, in order to avoid a high amount of radio waves produced by modern human activity. The ANtarctic Impulsive Transient Antenna (ANITA) experiment was designed, constructed, and launched in the quietest continent, Antarctica, in an attempt to detect this unique type of coherent radio emission.

[2]Equation 6.18 reduces to this form, as we are using the peak of the Askaryan pulse.

Chapter 5

ANITA

The ANtarctic Impulsive Transient Antenna (ANITA) is an 8-metre-tall balloon-borne radio array, consisting of three rings of quad-ridged horn antennas. The payload (see Figure 5.1) is attached to a Long Duration Balloon (LDB), and once airborne, moves in accordance with the circumpolar winds, approximately 40 km above the Antarctic ice sheet. The ANITA radio array is divided into 3 rings (top, middle, and bottom) and 16 azimuthal *phi-sectors* for a full 360° coverage in azimuth in each of the rings.

ANITA looks for coherent radio emission from neutrinos interacting in the Antarctic ice sheet (see Figure 5.5), via the Askaryan effect [69]. Section 4 describes the interactions and processes involved. As ice is particularly transparent to radio wavelengths, the radio waves produced by the neutrino-induced Askaryan effect can travel through the ice with little attenuation. These radio waves are primarily vertically polarised (VPOL), as only the top part of the Cherenkov cone escapes or traverses on the boundary between the Antarctic ice and the air. After escaping the surface of the ice, the radio waves make their way to the ANITA instrument. The power of the emitted radiation is proportional to the square of the shower energy.

The first observation of the Askaryan effect was reported in 2001, where GeV electrons were fired into a block of silica sand [72]. The Askaryan effect was next observed in rock salt in 2004 [73]. In 2006, the effect was observed in ice whereby bunches of 10^9 electrons, where each electron was of energy $\simeq 28.5$ GeV, were

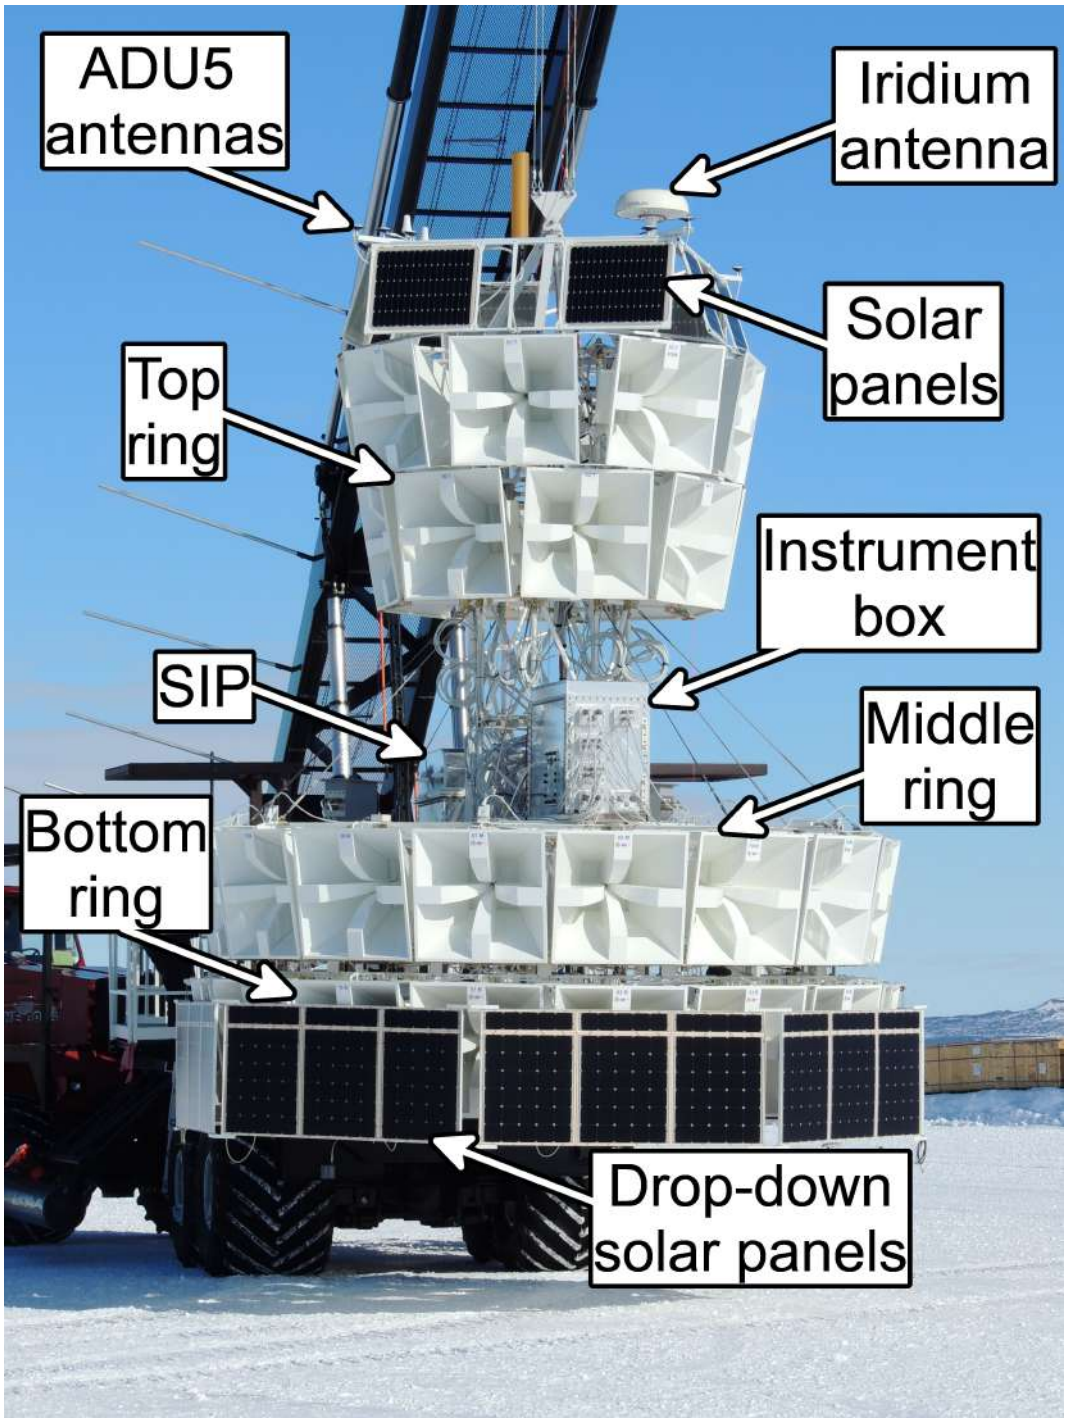

Figure 5.1: The ANITA-4 payload on the launch vehicle, *the Boss*, prior to launch preparation.

fired into a 7.5 metric ton ice target (See Figure 5.2 for the setup). This created showers of total energy approximately 3×10^{19} eV. Coherent radiation was emitted (see Figure 5.3), which was subsequently measured by the ANITA payload [74].

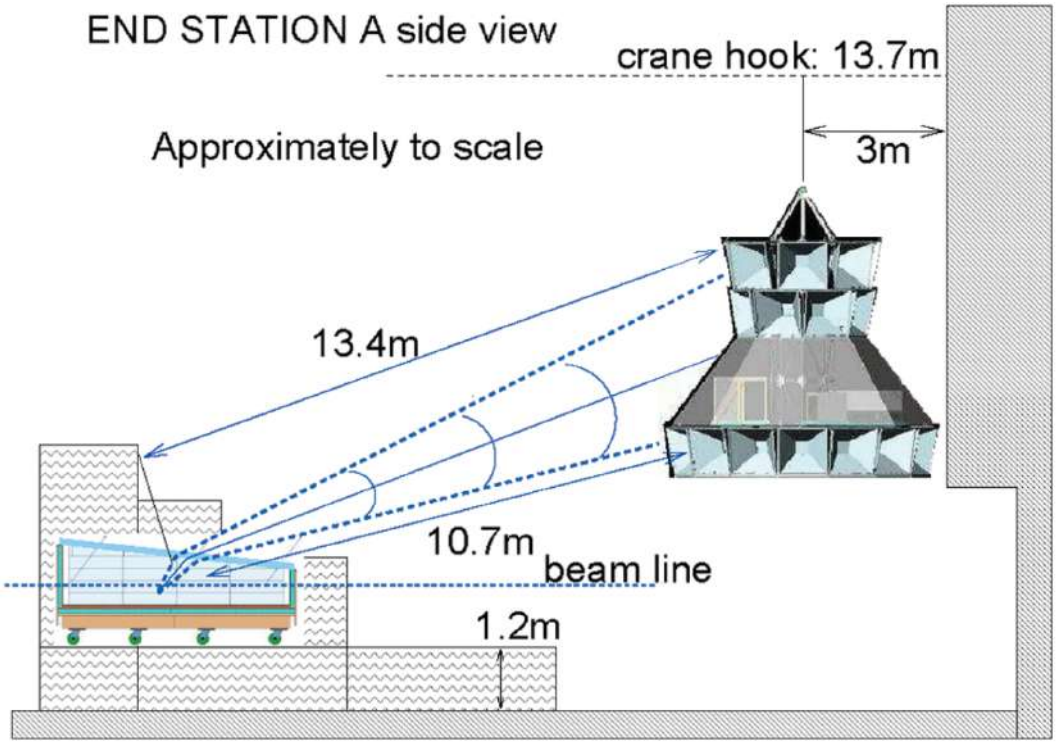

Figure 5.2: An illustration of the setup used to measure the Askaryan effect in ice at the Stanford Linear Accelerator Center (SLAC). The ANITA payload is suspended from a crane hook, and a beam of electrons was fired into an ice target. An electromagnetic shower builds up, and coherent radio waves are emitted, which are then detected by the payload. Image modified from [74].

In addition to detecting neutrinos, ANITA is able to detect cosmic rays. When a cosmic ray strikes the Earth's atmosphere, it can cause an extensive air shower (EAS). As the shower develops, it produces electron-positron pairs. The number of pairs produced is proportional to the shower energy, and these particles (having opposite and equal charges) experience the Lorentz force as they travel through the geomagnetic field. The particle pairs separate and gyrate, which gives off synchrotron radiation. The geomagnetic field is almost vertical at the magnetic poles, and the shower is polarised perpendicular to the vertical magnetic field and the shower axis (see Figure 5.4). As such, the emitted radio waves are mostly horizontally polarised (HPOL).

The first two flights of ANITA, ANITA-1 and ANITA-2, were able to set some of the best limits on GZK neutrino fluxes [75, 76]. ANITA-1 flew for 35 days,

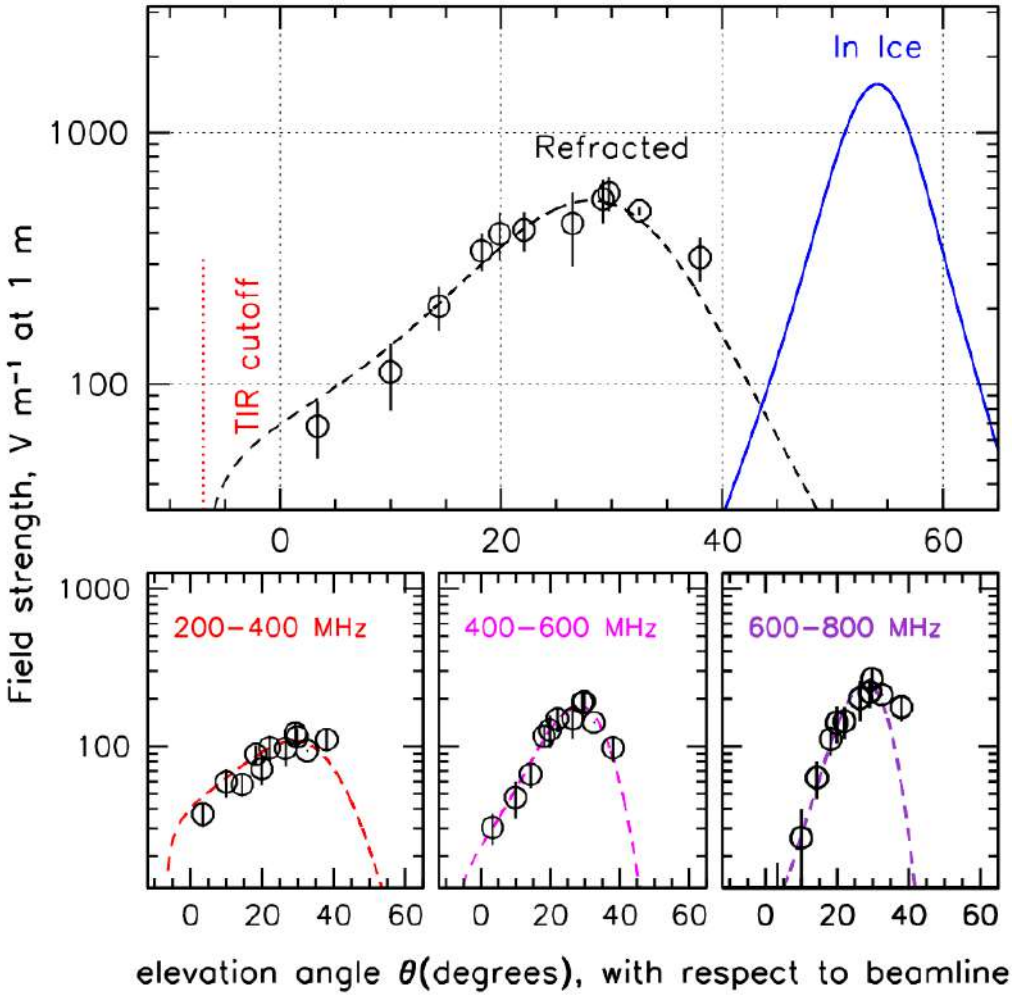

Figure 5.3: Electric field strength vs elevation angle from the beamline of the resulting Askaryan radiation 1 metre from the shower axis for 200-800 MHz (top) and for individual frequency bands (bottom). Dashed lines are the theoretical field strengths for refracted signals for showers of 3 $\times 10^{19}$ eV. The solid line is the same but for electric fields within the ice. The Total Internal Reflection (TIR) cutoff is also shown. Plots from [74].

found 0 neutrino candidates (and 1 anomalous event), but found 16 events consistent with radio impulses from UHECRs. As ANITA-2 was only sensitive to vertical polarisation, it only saw 4 radio emissions from cosmic rays, but saw a surviving possible neutrino candidate.

ANITA-3 was launched in December 2014 and flew for 22 days, and collected about twice the volume of data that ANITA-2 did. It observed 21 events consistent with EASs associated with cosmic rays, as well as an unusual upward-going cosmic-

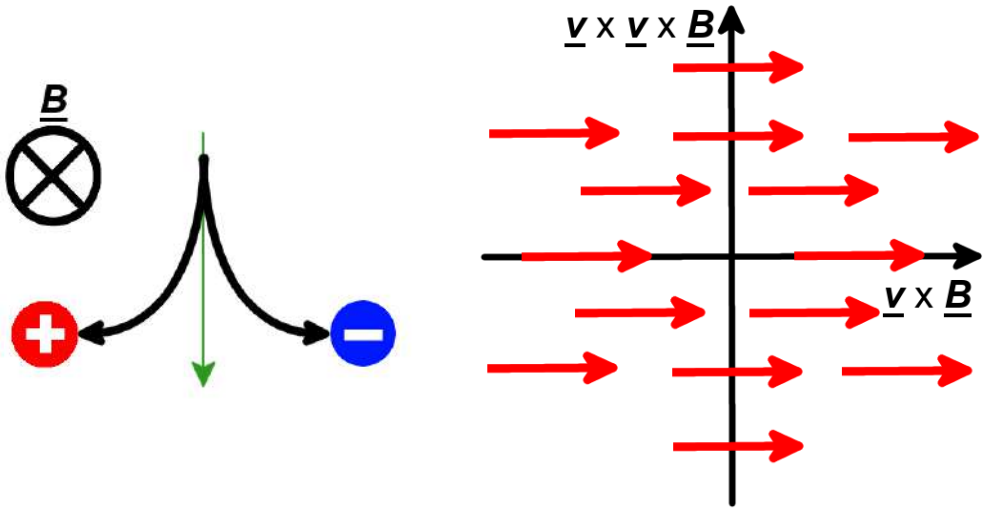

(a) Side view: shower axis goes from top to bottom. The green arrows represents the shower axis, and the magnetic field points into the page.

(b) Front view: shower axis points out of the page. It is clear that the radiation is emitted horizontally. The horizontal axis is aligned with $\underline{\boldsymbol{v}} \times \underline{\boldsymbol{B}}$ and the vertical axis is aligned with $\underline{\boldsymbol{v}} \times \underline{\boldsymbol{v}} \times \underline{\boldsymbol{B}}$, $\underline{\boldsymbol{v}}$ is the direction of the shower, $\underline{\boldsymbol{B}}$ is the magnetic field. The geomagnetic radiation is linearly polarised across the plane perpendicular to the shower axis (due to the direction given by the Lorentz force, $\underline{\boldsymbol{v}} \times \underline{\boldsymbol{B}}$).

Figure 5.4: An illustration of the polarisation of the emitted photons via an EAS.

ray like event [77, 78]. ANITA-4 was built and tested at the NASA Columbia Scientific Balloon Facility (CSBF), and launched at the Long Duration Balloon (LDB) Facility in December 2016, and flew for 28 days.

5.1 Components and Signal Chain

In order to process radio signals from potential science events, components of the ANITA experiment include antennas (to convert the power of the radio wave to a measurable voltage), sufficient amplification, a triggering system, and storage drives, etc. Each step of the testing presented herein, including all graphs produced and photographs taken, were carried out by myself, unless otherwise credited.

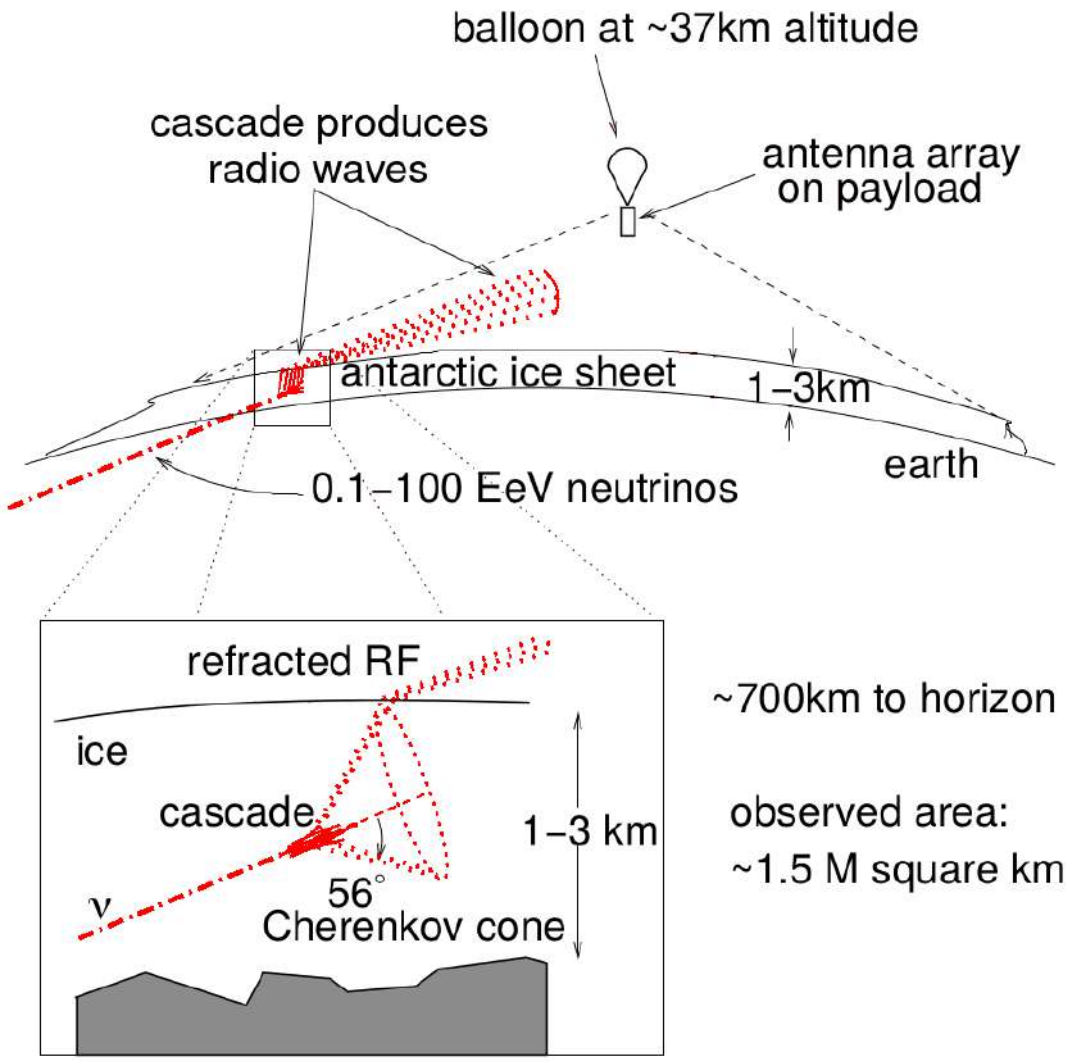

Figure 5.5: ANITA detection scenario, illustrating the neutrino-induced Askaryan effect in ice, which results in radio waves detected by ANITA. Image modified from [71].

5.1.1 Antennas

The ANITA-4 payload is comprised of 48 high-gain quad-ridged horn antennas. The antennas were custom-designed by Antenna Research Associates, Inc., and had two orthogonal feeds corresponding to vertical and horizontal polarisation, for a total of 96 channels. The frequency range for the antennas was 200–1200 MHz, and each had an aperture of 1 m^2. Each antenna is separated from the next antenna of the same ring by 22.5°. The area where most power is radiated, i.e., the beam

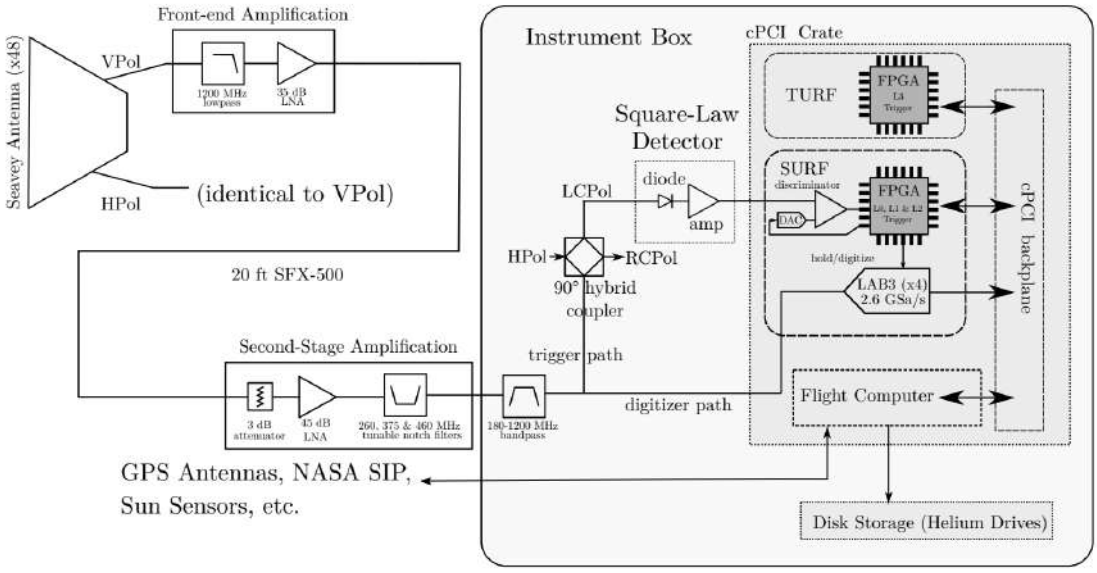

Figure 5.6: Schematic of the electronics of ANITA-4 from [79].

width of each antenna, is approximately 25-30° (in azimuth) from the center at the -3 dB point. Thus, using 16 antennas in a single ring provides 360° coverage to observe incoming radio waves at any azimuthal angle. The antennas are the first part of the signal chain. The signals from each antenna feed is amplified (see Sections 5.1.2 and 5.1.4), notch-filtered (see Section 5.1.4.1), and goes through a trigger system (see Section 5.1.5) that decides if the signal should be kept and digitised (see Section 5.1.6). The full system chain is shown in Figure 5.6.

5.1.1.1 Antenna Tests

Each antenna was tested on boresight (where the antennas directly face each other), and a select few were tested for off-boresight.

Boresight Measurements

Boresight measurements allow one to the calculate scattering parameters, as well as the gain, of the antennas; both are dependent on frequency.

A radio-frequency (RF) network anaylser is used for such tests, with the first port connected to the transmitting antenna. The second port is connected to the the receiving antenna, which is elevated higher than the transmitting antenna, in order to mitigate reflections of radio waves. The setup is shown in Figure 5.7,

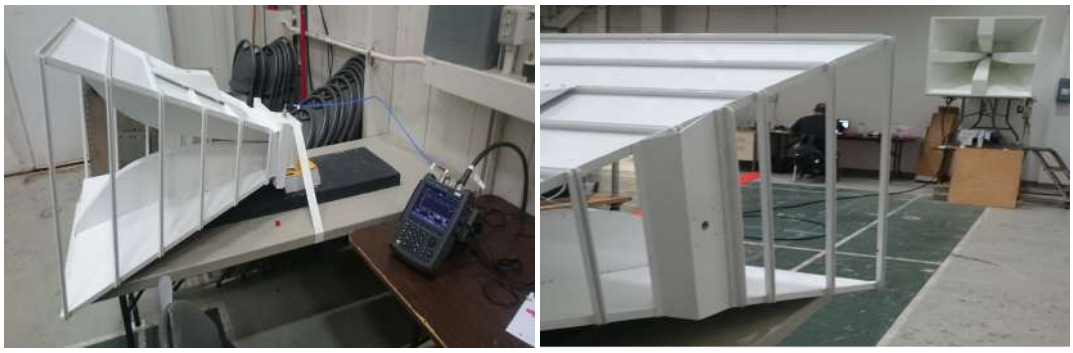

(a) The transmitting antenna and the RF analyser.

(b) The transmitting antenna (closest) and the elevated receiving antenna (furthest).

Figure 5.7: Experimental setup for boresight testing.

and the circuit diagram is shown in Appendix B.1. Radio waves are sent from the transmitting antenna to the receiving antenna. The set of scattering parameters (the transmission and reflection parameters of the antennas in the frequency domain), as well as the gains, are found for both VPOL and HPOL waveforms for each antenna. For each of these, co-polarisation and cross-polarisation measurements are carried out by rotating the receiving unit 90°. Co-polarisation describes where an antenna system sends and receives radio waves of the same polarisation (if sending and receiving VPOL, the measurement is known as vertical co-polarisation), whereas cross-polarisation is where the polarisation of the received signal is orthogonal to the original signal. This brings the total measurements per antenna to be 4 (co- and cross-polarisation for both horizontal and vertical polarisations). In total, the 48 ANITA-4 antennas, 7 ANITA-3 antennas, 2 ANITA-2 antennas, and 1 ANITA-1 antenna were tested.

A network analyser directly measures the parameters of a 2×2 scattering matrix, S, which describes the input-output relationship between its ports. For example, S_{21}, often just denoted S21, represents the power received at antenna 2 relative to the power input at antenna 1. When S21 is 0 dB, this means that all the power from antenna 1 is received by antenna 2. When S21 is lower than 0 dB, some of the power is lost. If S21 has been measured to be -10 dB, this means that only a tenth of the power is received. S22 is indicative of the reflection measurement, as both the input and output ports are the same. As the ANITA-4 antennas were used as receiving

antennas (associated with port 2) only, the spectra of the S21 and S22 parameters were measured for each antenna (see Figure 5.8).

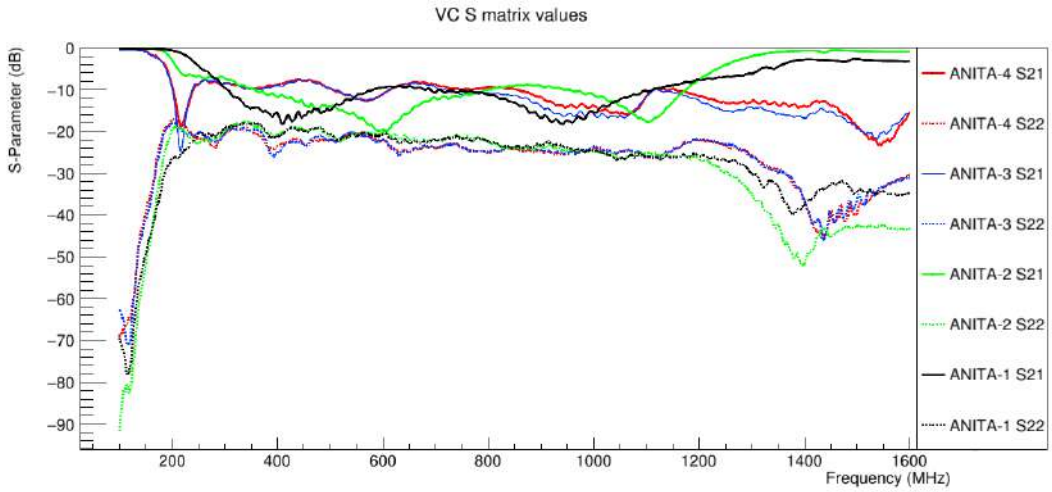

Figure 5.8: S parameter plot for vertical co-polarisation for each antenna type.

The gain of an antenna is essentially its ability to receive radio waves in a given direction, i.e., the ratio of the power received by an antenna to that of an isotropic antenna. The modified Friis equation (Equation 5.1) was used to find these gains. To do this, the gain of the reference antenna, marked at specific points by the manufacturer, was splined into a function. From the modified Friis equation, one can derive each receiving antenna's gain using raw S-matrix parameter values, the reference gain, the distance between the antennas, and frequency (see Equation 5.2).

The modified Friis equation [80] is:

$$\frac{P_\mathrm{r}}{P_\mathrm{t}} = G_\mathrm{t}(\theta_\mathrm{t},\phi_\mathrm{t})G_\mathrm{r}(\theta_\mathrm{r},\phi_\mathrm{r})\left(\frac{\lambda}{4\pi R}\right)^2(1-|\Gamma_\mathrm{t}|^2)(1-|\Gamma_\mathrm{r}|^2)|\mathbf{a_t}\cdot\mathbf{a_r^*}|^2 e^{-\alpha R}, \tag{5.1}$$

where P_r is the power received by the receiving antenna, P_t is the power delivered to the transmitting antenna, G_t is the gain of the transmitting antenna, $\mathrm{G_r}$ is the gain of the receiving antenna, with the angular components of gain denoting the gain in the direction (θ,ϕ) one antenna views the other antenna. λ is the wavelength, R is the distance between the two antennas. Γ_t, Γ_r are the reflection coefficients of the

transmitting and receiving antennas, respectively. $\mathbf{a_t}$, $\mathbf{a_r}$ is the polarisation vectors of the antennas, and α is it the attenuation coefficient of the medium between the antennas.

It is expected that each of the antennas have similar gain profiles across all frequencies that ANITA is sensitive to. The measured and calculated gains of the ANITA antennas are shown in Figure 5.9. When the frequency exceeds ANITA's maximum detectable frequency, the gain largely decreases, as the antennas are no longer reliable waveguides. The frequent small increases and decreases in gain are due to the testing environment (specifically, reflections of the radio waves by the walls of the hangar); in an ideal scenario (such as an anechoic chamber), reflections would be mostly avoided. As the antennas are directly pointed towards each other, we can greatly simplify Equation 5.1. Noting that the gains of the antennas (in the case of boresight measurements) are independent of azimuthal and elevation components ($G(\theta,\phi) = G(0,0) = G$), that the polarisation vectors are parallel, and that the absorption coefficient of the intervening medium is approximately 0, an equation for the gain can be derived:

$$G_{\mathrm{r}} = S_{21} - G_{\mathrm{t}} - 10\log_{10}\left(1 - 10^{\frac{-S_{22}^2}{5}}\right) - 10\log_{10}\left(1 - 10^{\frac{-S_{11}^2}{5}}\right) + 20\log_{10}\left(\frac{4\pi R f}{c}\right), \tag{5.2}$$

where S_{21} is the scattering matrix parameter 2,1 (S_{21} represents the power received at antenna 2 relative to the power input at antenna 1), S_{22}, S_{11} are scattering matrix parameters 2,2 and 1,1 respectively (these parameters represent how much power is reflected from the antennas), f is the frequency, c is the speed of light in a vacuum.

The effective antenna area is a useful quantity for showing how effective an antenna is at receiving power, and is used in calculating power received by the antennas. This quantity scales with both gain and the squared wavelength, meaning the effective area should reduce as frequency increases. It can be found from the gain by using Equation 5.3, and is shown in Figure 5.10.

$$A_{\mathrm{eff}} = \frac{G\lambda^2}{4\pi}, \tag{5.3}$$

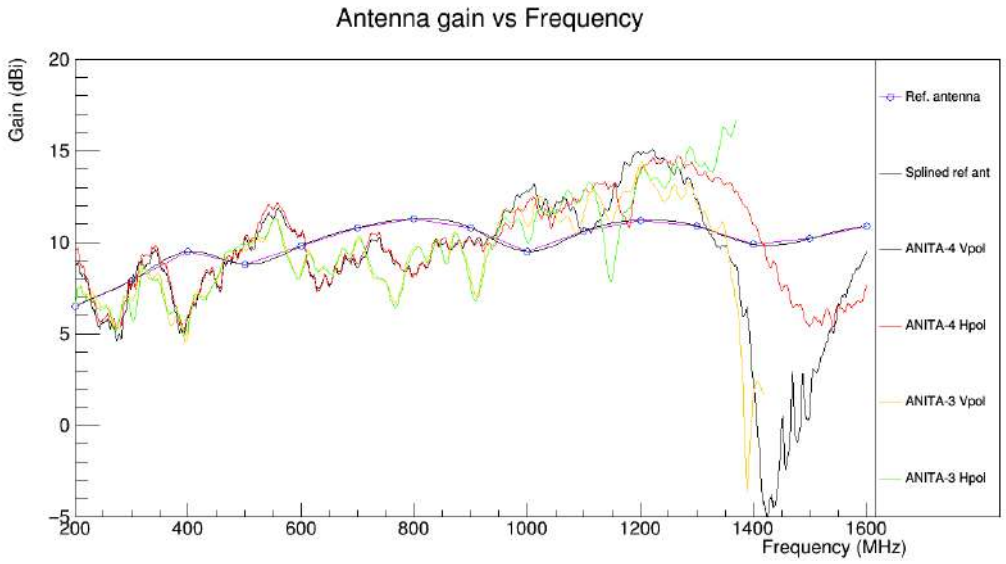

Figure 5.9: Average gain plots for each flight-specific antenna, showing antennas from ANITA-3 and ANITA-4 with VPOL and HPOL components, as well as the reference antenna's gain. The gains of the ANITA antennas were calculated from the S parameters measured from the boresight tests.

where A_{eff} is the effective antenna area, G is the receiving antenna gain.

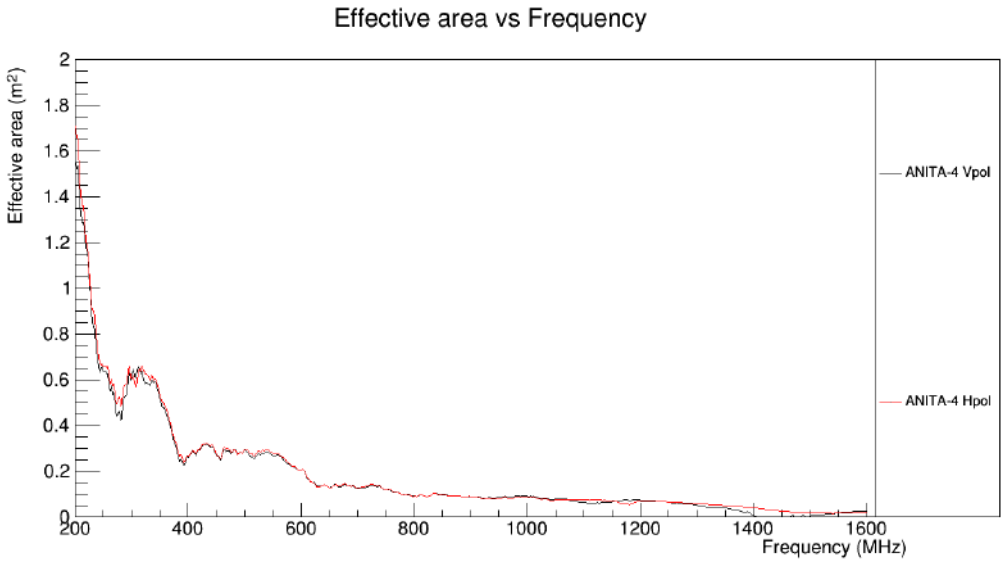

Figure 5.10: Average effective area for ANITA-4 antennas (both VPOL and HPOL shown) as a function of frequency.

Off-Boresight Pulse Measurements

In order to test receiving antennas in the case that they are not directly facing the

(a) Antenna positions for pulsing. The receiving antenna (closest) was adjusted in elevation and azimuth by the adjustable mount. The pulsing antenna (furthest) was mounted on an antenna crate.

(b) The pulsing setup used an impulse generator unit as well as an oscilloscope to read out the data.

Figure 5.11: Experimental setup for off-boresight pulse testing.

transmitting antenna, the antennas are directed away from it, both in elevation and azimuth. The receiving antenna is mounted on a rotation mount that can move in both angles (see Figure 5.11). A Picosecond Pulse Labs impulse generator is used to send a pulse to the receiving antenna (see Appendix B.2 for the corresponding circuit diagram).

Time domain waveforms are recorded, as well as the Fourier-transformed data for each rotation in azimuth and elevation. Due to the design of the mount, it is not possible to achieve elevations greater than +60° or lower than -10°.

Figure 5.12 shows the power distributions at fixed elevation angle ($\theta = 0°$), but altered azimuth, for vertical cross-polarisation data. Figure 5.12 also illustrates that antennas closer to boresight receive more power across 200–1200 MHz: ANITA's detectable radio frequency range.

5.1.1.2 Power Received During Flight

The power received by an antenna is calculated in Equation 5.4:

$$P = |S| \cdot A_{\text{eff}} = \frac{|\epsilon^2|}{2Z_0} \cdot A_{\text{eff}}, \tag{5.4}$$

where P is the power received by the antenna, $|S|$ is the magnitude of the electric field's Poynting vector, A_{eff} is the effective antenna area, ϵ is the electric field measured at the antenna, and Z_0 is the impedance of free space.

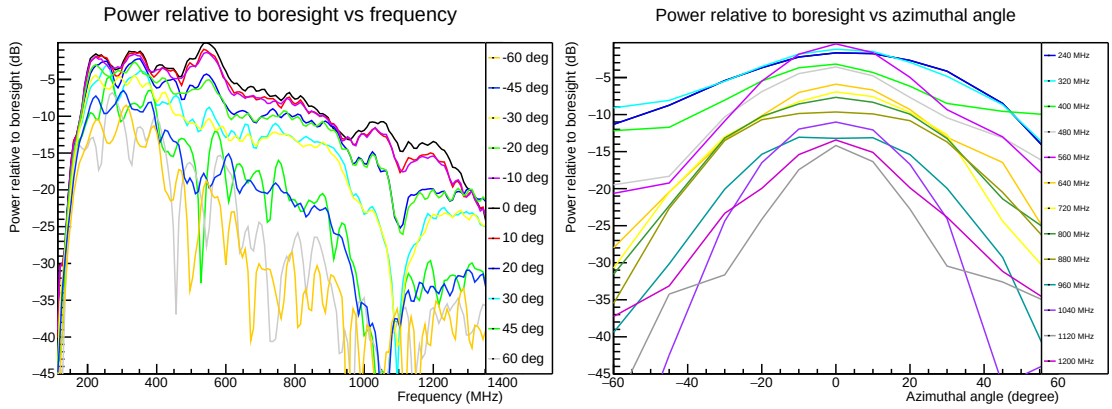

Figure 5.12: Antenna power spectra (power vs frequency for different azimuths) for a single ANITA-4 antenna (left). The power-azimuth relation for different frequencies is shown (right). The powers shown are relative to the boresight measurement for the same antenna.

In order to calculate the power received by an antenna during flight, we need to find the magnitude of the electric field as measured by the payload, $|\epsilon|$. For all frequencies detectable by ANITA, the gain is lowest at around 300 MHz (see Figure 5.9). Using Figure 5.3, the electric field when 1 metre from a shower (of energy 3 $\times$ 10^{19} eV) is $\sim$ 100 Vm^{-1} at 300 MHz. As ANITA is suspended far above the ice, the signal has to propagate both through the ice and air before reaching the payload. This leads to a reduction in the magnitude of electric field, and thus the power received. Using Equation 4.2, and assuming that the distance travelled in ice is equivalent to its attenuation length and that the total path length is 500 km, the electric field as measured at the payload is:

$$|\epsilon| = 100\ \mathrm{Vm}^{-1} \cdot \mathrm{e}^{-1} \cdot \frac{1\mathrm{m}}{500{,}000\ \mathrm{m}} = 7.4 \times 10^{-5}\ \mathrm{Vm}^{-1}\,. \tag{5.5}$$

Using Figure 5.10, the effective area is approximately 0.4 m^2 at 300 Hz. Inserting the magnitude of the electric field from Equation 5.5, effective area, and the impedance of free space into Equation 5.4, the power as measured at the antenna is,

$$P = \frac{(7.4 \times 10^{-5}\ \mathrm{Vm}^{-1})^2}{2 \times 376\Omega} \cdot 0.4\mathrm{m}^2 = 3\ \mathrm{pW}\,. \tag{5.6}$$

The corresponding voltage measured at the antenna is,

$$V = \sqrt{PR} = \sqrt{3\ \text{pW} \cdot 50\Omega} = 12\ \mu\text{V}, \tag{5.7}$$

where V is the voltage correponding to the power received by the antenna (from Equation 5.6), and R is the resistance of the receiver.

The voltage measured by the antenna (as in Equation 5.7) is $\sim 10\ \mu$V, which is rather minute. In fact, it is comparable to voltage fluctuations caused by that of the thermal noise observed by an antenna. The thermal power is calculated to be,

$$P = E_{\text{th}}\Delta f = k_{\text{B}} T \Delta f, \tag{5.8}$$

where E_{th} is the thermal energy received by the antenna, T is the average temperature of ANITA's field of view, and Δf is the frequency range. Thus, the power is,

$$P = 1.38 \times 10^{-23}\text{JK}^{-1} \cdot 150\text{K} \cdot (1200 - 200) \times 10^{-6}\text{s}^{-1} = 2\ \text{pW}. \tag{5.9}$$

Inserting the thermal power into Equation 5.7, the thermal voltage is,

$$V = \sqrt{PR} = \sqrt{2\ \text{pW} \cdot 50\Omega} = 10\ \mu\text{V}. \tag{5.10}$$

When signals later get digitised, the electronics require an input pulse of ~ 10 mV, which is a factor of a 1000 higher. As such, sufficient amplification is needed:

$$A = 20\log_{10}\left(\frac{1\ \text{mV}}{1\ \mu\text{V}}\right)\text{dB} = 60\ \text{dB}, \tag{5.11}$$

where A is the required amplification.

At minimum, there needs to be at least 60 dB of amplification applied to the signal (as calculated in Equation 5.11). Thus, to account for low voltage signals and losses in the digitiser system (from insertion loss and the four-way split detailed in Section 5.1.6), a total of 80 dB of amplification is used. Throughout the signal chain, two stages of amplification are carried out, firstly within the Antenna-Mounted Pre-Amplifiers (AMPAs), and later in the internal Radio Frequency Con-

ditioning Modules (iRFCMs).

5.1.2 AMPAs

The outputs of the antennas are passed onto the AMPAs, which are directly connected to the antenna VPOL and HPOL outputs. Each attached to a single feed, the AMPAs are the front-end amplification units, each containing a 200–1200 MHz bandpass filter followed by a 35 dB Low Noise Amplifier (LNA). A total of 98 AM-

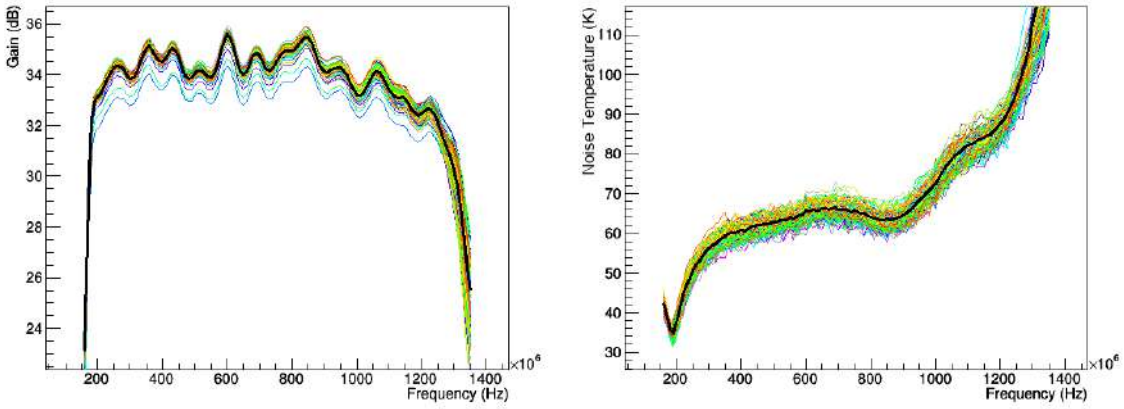

Figure 5.13: AMPA gain (left) and noise temperature (right) vs frequency. The black line indicates the average values of gain and noise temperature at each frequency.

PAs (96 channels and 2 spares) were tested for their gains. They were also tested for their noise temperatures. The noise temperature is essentially the additional noise that the amplifier itself inserts onto a signal before it is amplified. Each AMPA was contained very well within a tight noise temperature envelope. Both gain and noise temperature graphs are shown in Figure 5.13.

5.1.3 Flight Cables

As soon as a signal leaves the AMPA, it travels through a HELIAX Superflexible (SFX-500) Coaxial Cable. SFX-500 cables have polyethylene jackets, and operate in temperatures as low as -50 °C. These cables have a very high shielding effectiveness (110 dB). Each cable needs to be rigorously tested in order to look for defects, and in order to do this, the group delay of each cable was found. The cables with the most similar group delays were used for the experiment. See Figure 5.14 for the distribution of flight cable group delays.

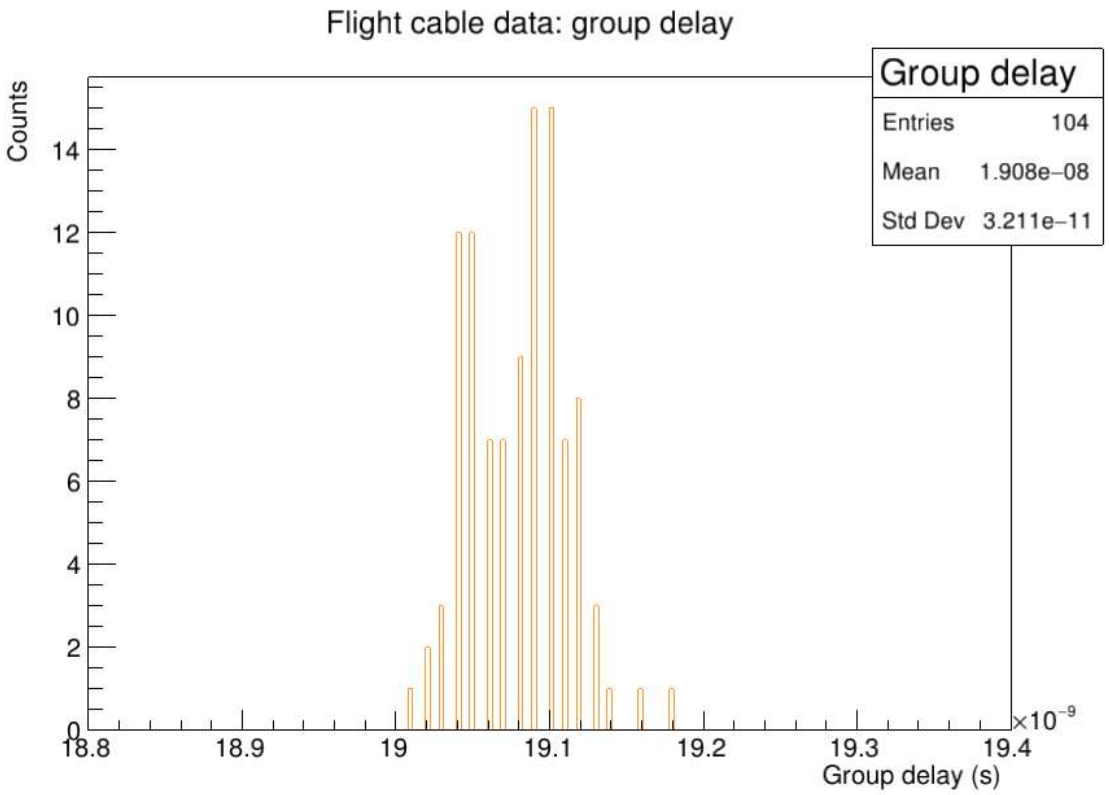

Figure 5.14: The measurements of group delay for the SFX-500 flight cables.

5.1.4 iRFCMs

After reaching the end of the flight cables, the signals go through second stage amplification. This stage involves the iRFCMs, of which there are four, each having 24 input channels. The iRFCMs contain two Tunable Universal Filter Frontend (TUFF) modules each [84]. The internals of an iRFCM are shown in Figure 5.15.

5.1.4.1 TUFFs

Within the iRFCM units are the TUFFs. A TUFF is a computer board with the following components: one bias tee, two amplifiers, a microcontroller, and a set of three tunable notch filters, which are used to reduce anthropogenic narrow-band noise of specific frequencies. A photograph of a TUFF is shown in Figure 5.16. Due to high Continuous Wave (CW) interference throughout the previous flights, each TUFF had three different notch filters, each set to a different default central frequency. Notch 1 was set at 260 MHz, as frequencies around this point are used by

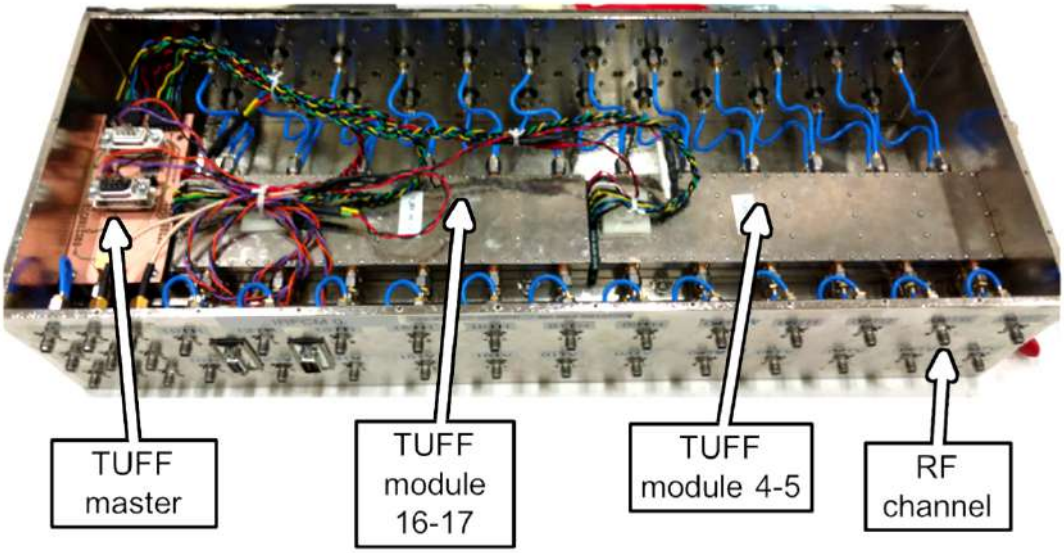

Figure 5.15: Internals of 1 of the 4 iRFCMs. Displayed are 2 TUFF modules, each of which has 12 RF input channels. ANITA has eight TUFF modules in total, sufficient for 96 RF channels. The TUFF master is used to issue commands to the TUFF boards, such as selecting which notches to use. Image modified from [84].

both the Ultra High Frequency Follow-On (UFO) and the Fleet Satellite (FLSAT) systems. The central frequency of Notch 2 was set at 375 MHz, as this is within the range of Mobile User Object System (MUOS) satellite communication systems. Notch 3's default frequency was set to 460 MHz, due to the high CW interference from bases such as the Amundsen-Scott South Pole Station and McMurdo Station. These central notch frequencies are tunable, so can be changed from the defaults, should troublesome CW interference associated with other frequencies arise during flight. Not all notches need be on at the same time, such as when away from major science bases. This method of dealing with frequency-specific CW interference is superior to previous methods.

During periods of high CW interference in previous ANITA flights, the system trigger rate was much higher than its digitisation rate, leading to a large portion of digitisation dead time (see Section 5.1.6.1). In order to deal with large portions of digitisation dead time in ANITA-3, phi-masking was applied. During this previous flight, when ANITA-3 passed over noise-loud bases on the continent, it applied

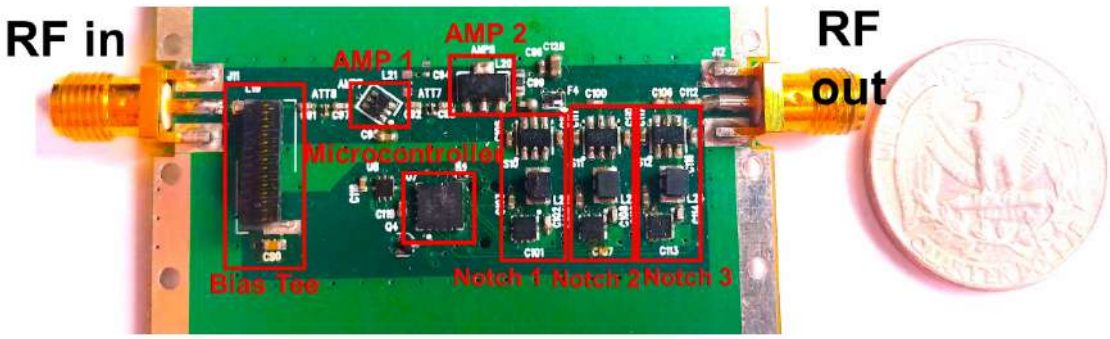

Figure 5.16: The internals of a single TUFF, shown next to a quarter dollar coin for size comparison. Image from [84].

masking, i.e., it could block out (mask) certain phi-sectors from participating in the final level of triggering. This helped to reduce the trigger rate, and thus the digitisation deadtime. However, on average, over half the payload was masked in ANITA-3, meaning the overall instrument deadtime was increased (see Figure 5.17 for a comparison of masking usage between ANITA-3 and ANITA-4). Dead time is the time in which the payload is no longer able record events and potential science signals (i.e., due to masking and / or digitisation dead time). Live time is the the time spent actually recording signals, and an aim of ANITA-4 was to maximum this. Due to the implementation of the TUFF boards, as well as the chosen trigger scheme, in ANITA-4, the digitisation live time percentage increased from 73.7% to 92.3%, and the total live time percentage increased from 31.6% to 91.3%, from ANITA-3 to ANITA-4.

As well as notch-filtering, these boards are used for second-stage amplification to bring the radio signals to the millivolt level. The amplification boost is $\simeq$ 45 dB, making the signals sufficient for digitisation. One iRFCM contains 2 TUFF modules, as well as a TUFF master that is responsible for sending commands to the TUFFs. Each TUFF module has 12 channels.

After the signals are notch-filtered, they exit the iRFCM and are subject to bandpass filtering between 200–1200 MHz. Second stage amplification may be responsible for additional noise beyond the frequency range of the antennas and

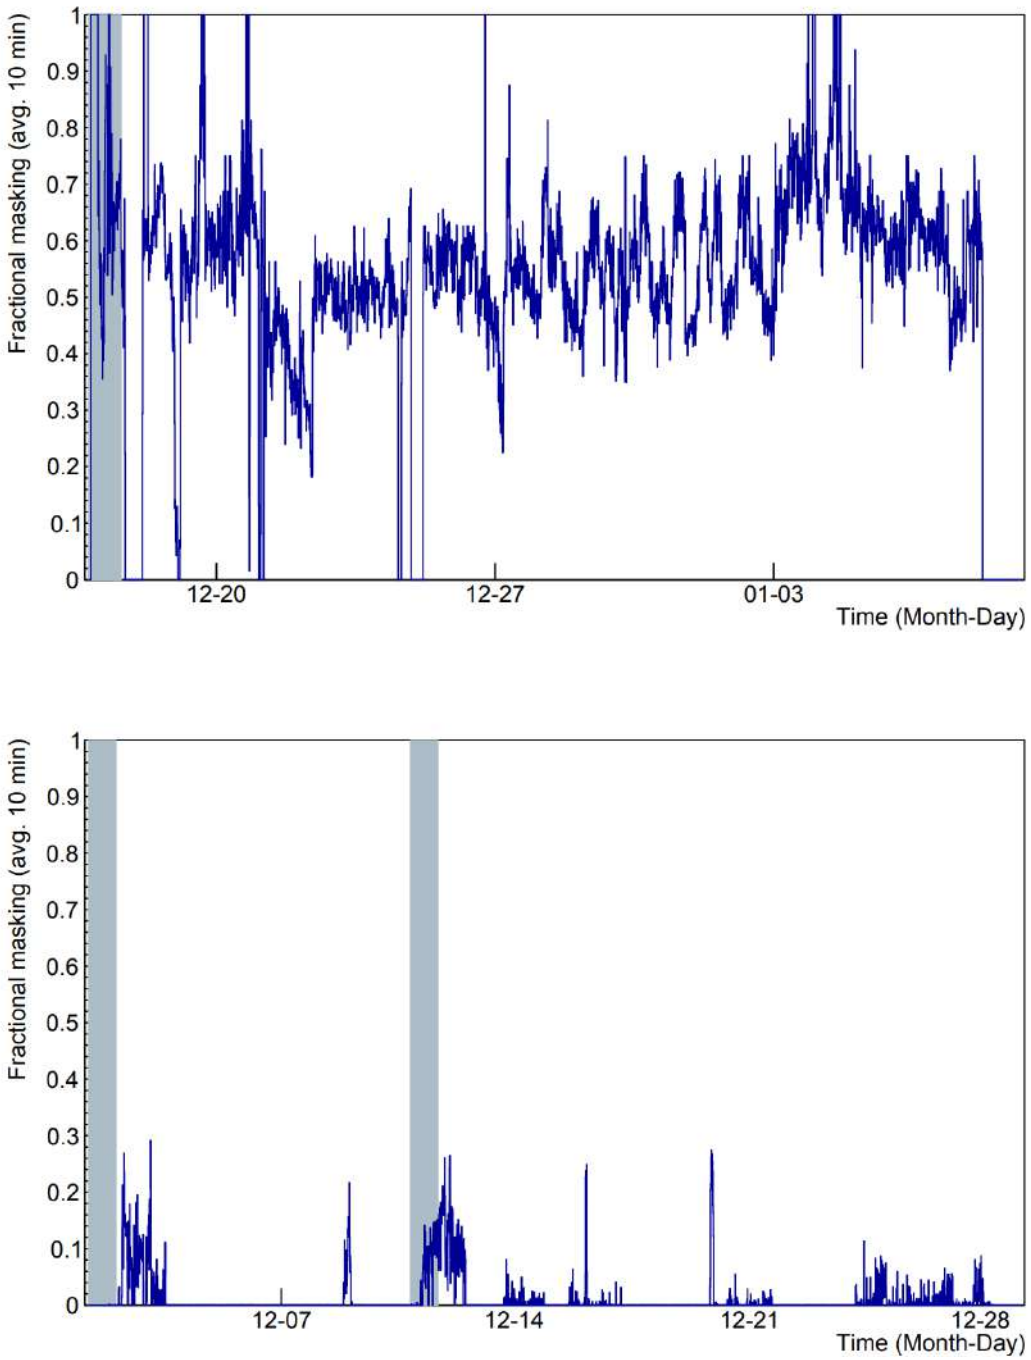

Figure 5.17: Fractional masking averaged over 10 minutes for ANITA-3 and ANITA-4. The shaded regions correspond to when the payload was in line-of-sight to the Long Duration Balloon facility. In this case, there is a large amount of anthropogenic noise, and this needs to blocked out via masking. Image from [84].

frontend amplifiers, so noise at these frequencies are filtered out (as was done to signals after leaving each antenna).

Now that the signals are sufficiently notch-filtered and amplified, the signals pass through a 3 dB splitter into a triggering and digitisation path. The signal in the triggering path leads to a 90° hybrid coupler unit, whilst the signal in the digitisation path leads to a Sampling Unit for Radio Frequency (SURF).

5.1.5 Trigger Path

It is not possible to record and store every single sample of a waveform that ANITA observes. This is due to strict power and storage limitations. Consider recording every sample for ANITA-4: 2.6×10^9 samples are recorded per second, and each sample is 12 bits in size. There are 96 RF channels (both HPOL and VPOL for each of the 48 antennas), so for a flight of 28 days, this would take up over 900 petabytes, an unfeasible amount of data to store. Even if this was possible, the power needed to repeatedly record all wavesforms for a single channel is $\sim$ 10 W. For 96 channels, this would take up almost 1 kW, which far exceeds the system's power constraints (see Section 5.1.10). Thus, ANITA only digitises a small fraction of these events. In order to choose which events to keep, a trigger system is utilised.

Instead of triggering on linearly polarised waveforms, ANITA-4 triggers on circularly polarised waveforms. Linearly polarised signals have equal amounts of both left-circular polarisation (LCP) and right-circular polarisation (RCP). Neutrino-induced events are highly linearly polarised, and thus will be near to equal in the two components of circular polarisation. Unwanted signals, such as those from satellites, are often dominated by *either* LCP or RCP. Requiring the events to pass in *both* LCP and RCP allows us to exclude most of the satellite events. In order to get from linearly polarised waveforms to circularly polarised waveforms, the signals are passed through a 90° hybrid unit (see Equation 5.12). The hybrid units convert VPOL and HPOL waveforms from channels of the same antenna into LCP and RCP waveforms:

$$|L\rangle = \frac{1}{\sqrt{2}}(|H\rangle + i\,|V\rangle), \qquad |R\rangle = \frac{1}{\sqrt{2}}(|H\rangle - i\,|V\rangle), \tag{5.12}$$

where $|L\rangle$ and $|R\rangle$ represent LCP and RCP waveforms and $|H\rangle$ and $|V\rangle$ represent

HPOL and VPOL waveforms.

5.1.5.1 Trigger Logic

ANITA-4 employs a four-level trigger system. In order to reach the next trigger level, the former level must be passed. ANITA-3 had a three-level trigger system, and its levels were referred to as Level 1 (L1), Level 2 (L2), and Level 3 (L3). However, ANITA-4 has an extra initial triggering level, and so this is referred to as the Level 0 (L0) trigger.

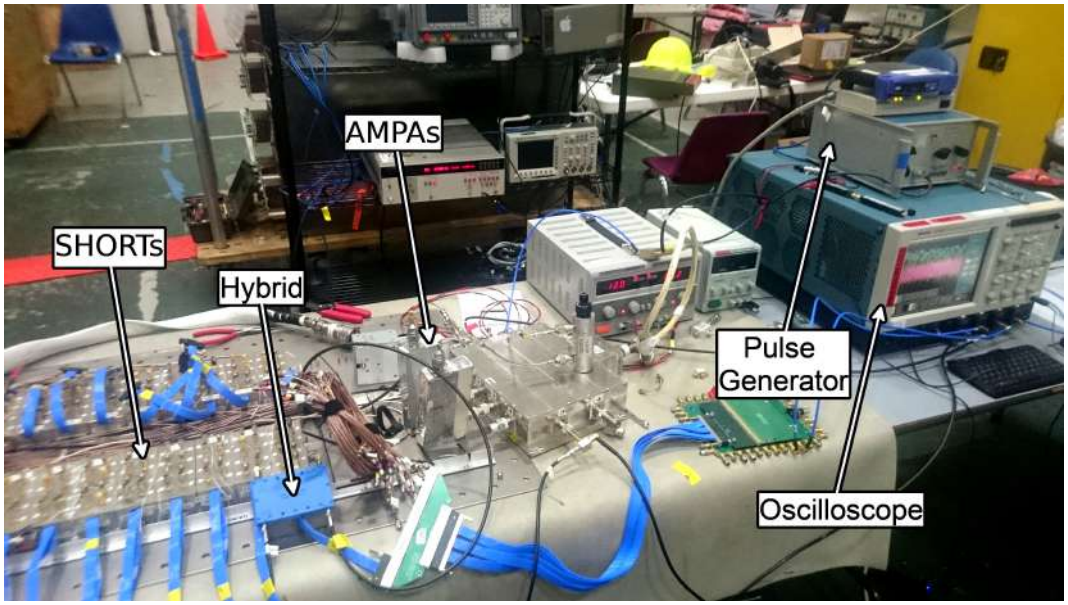

Figure 5.18: Replication of the initial part of the triggering path. A waveform from the generator is passed through an AMPA, then through a hybrid unit. The hybrid converts waveforms of linear polarisations into those of circular polarisations. After this, the signal is sent to the SHORT board and back to the oscilloscope (it has to pass through a SHORT interface, followed by a SHORTmate).

Level 0 Trigger

The LCP and RCP waveforms from the hybrid pass through a SURF High-Occupancy RF Trigger (SHORT). A SHORT consists of a tunnel diode and an LNA, and receives four channels as input. The diode functions as a square law detector, providing a DC voltage output proportional to the power (i.e., proportional to the square of the voltage input) that is continuously integrated with a time constant of 5 ns. The output is then sent to a SURF input and undergoes discrimination. If either

output voltages of the LCP or RCP signal exceed the threshold of the the SURF Digital-to-Analog Converter (DAC), an L0 trigger is issued.

Level 1 Trigger

An L1 trigger is only issued if both LCP *and* RCP outputs from the same hybrid pass the L0 trigger criterion. If only one component of circular polarisation passes the L0 trigger, or neither of them do, an L1 trigger is not issued. If both outputs pass the L0 trigger, they must do with 4 ns of each other, else an L1 trigger is not issued by the SURF. The L1 trigger logic is shown in Figure 5.19.

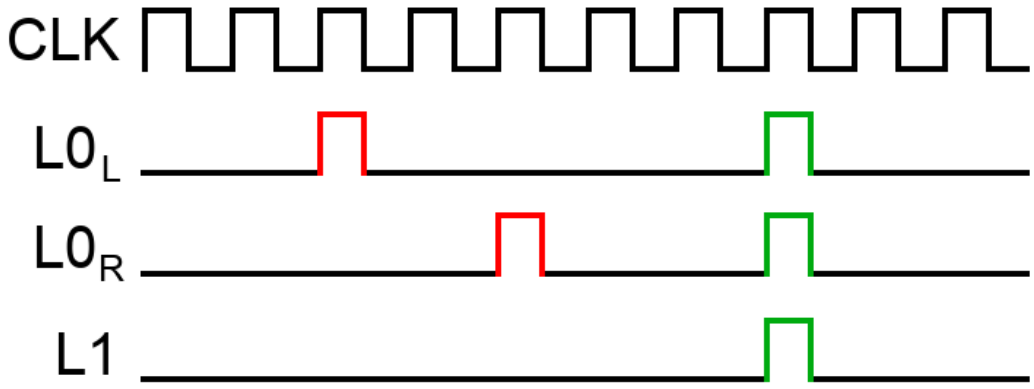

Figure 5.19: L1 trigger logic for a single antenna in the top ring. CLK represents the clock (4 ns clock edge, 125 MHz). If there is an L0 trigger from either an LCP waveform ($L0_L$), or an L0 trigger from an RCP waveform ($L0_R$), a 4 ns window opens up. If these $L0_L$ and $L0_R$ windows overlap, an L1 trigger is issued.

SHORT Testing

In order to test that each SHORT is performing up to standard, the first part of the ANITA signal and triggering chain is replicated. See Figure 5.18 for the experimental setup, as well as Appendix B.2 for circuit diagrams. To mimic a waveform, such as that from a single channel, an impulse generator produces a dummy signal. This is then passed through a splitter, and finally through the SHORT boards. Each of the three corresponding waveform (initial antenna signal output, split pulse, and output waveform) channels are linked up to an oscilloscope (see Figure 5.20). The peak-to-peak voltages ($V_{\text{p-p}}$) are recorded for each SHORT output for 96 channels, to test if there are any channels that produced low voltage responses. During testing, some of the SHORT channels reported low voltages, associated with a particular faulty

port in the SHORT interface. This SHORT interface was not used for the ANITA-4 flight.

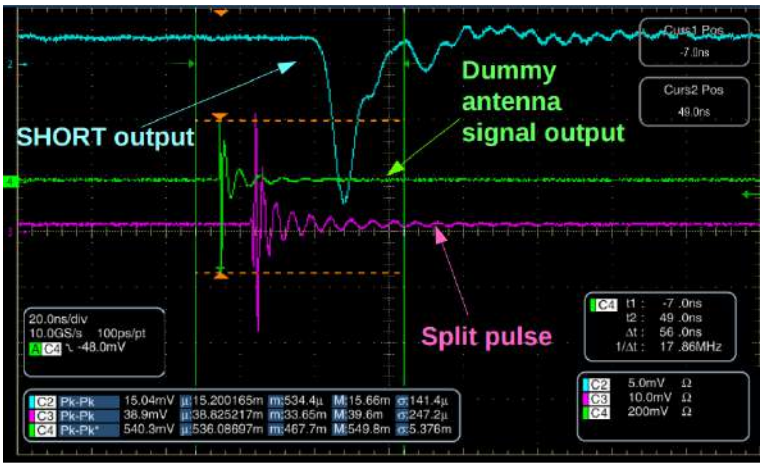

Figure 5.20: A screenshot of the oscilloscope with all outputs shown. The peak-to-peak voltages of average spectra are displayed.

The setup was also modified to include the hybrid units (to test their functionality), as well as a solid-state noise source and a manual step attenuator (to alter signal-to-noise ratio and check efficiencies of the triggering system).

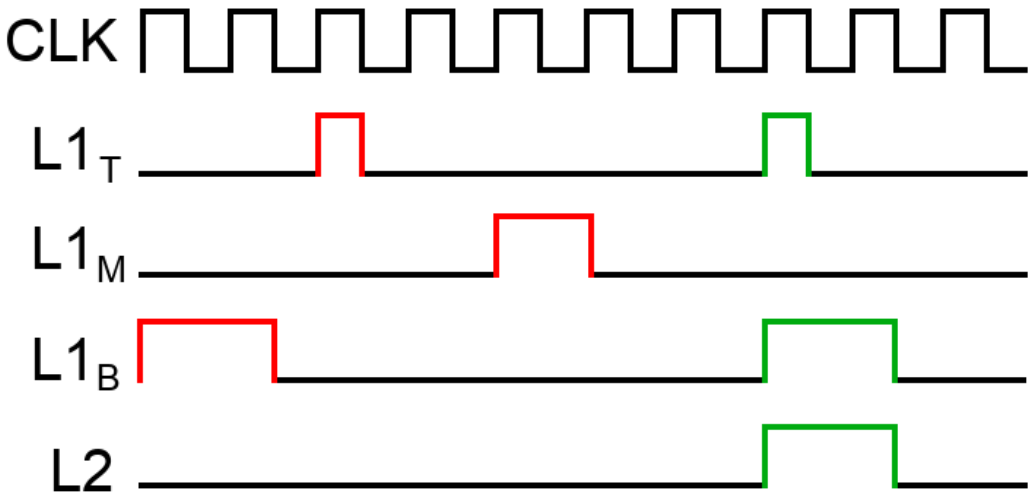

Figure 5.21: L2 trigger logic for a single phi-sector. CLK represents the clock. If there is an L1 trigger from the bottom antenna, a 12 ns window opens up ($L1_T$). If there is an L1 trigger from the middle antenna, an 8 ns window opens up ($L1_M$). If there is an L1 trigger from the top antenna, a 4 ns window opens up ($L1_B$). If these windows overlap, an L2 trigger is issued.

Level 2 Trigger

When an L1 trigger is issued for a single antenna, a check is made to see if a second L1 trigger has also been issued within an antenna of the same phi-sector. If two of these L1 triggers are issued, within a specific time window, an L2 trigger is issued by the SURF, which is then sent to the Triggering Unit for Radio Frequency (TURF). It is also possible for three of L1 triggers to issue an L2 trigger. In order to look for "up-going" signals (those from the Antarctic ice sheet), a timing window scheme is used. Because up-going signals will likely be first recorded by a bottom ring antenna, the allowed window (12 ns) is the longest in this case. The time window is shorter if antennas in the middle ring (8 ns) or top ring (4 ns) activate first. The L2 trigger logic is shown in Figure 5.21.

Level 3 Trigger

The L3 trigger is the global trigger, which finally determines whether an event is kept, or discarded. Unlike the previous trigger levels, the global trigger is not issued by the SURF, but instead by the TURF. The L2 signals are sent from the SURF to the TURF, and if an L2 trigger is activated, a timing window (12 ns) opens up. Now, if an L2 trigger is also issued in an adjacent phi-sector, and within the time window, an L3 trigger is issued by the TURF. The corresponding event is now marked to be recorded, and the TURF commands the SURF to digitise it. The L3 trigger logic is shown in Figure 5.22.

5.1.6 Digitisation Path

Signals travelling through the digitiser path get sent straight to the SURF, as there is no need to convert VPOL and HPOL signals to LCP and RCP, or apply a diode response to the waveforms. This allows the waveforms to be digitised. As there are 96 RF channels, and each SURF can handle 8 RF channels, ANITA-4 requires 12 SURF boards.

Once a single waveform has entered a SURF, it is split by a four-way splitter. Each of the four waveforms now enter a separate third-generation Large Analog Bandwidth Recorder And Digitiser with Ordered Readout (LABRADOR) chip (LAB3) [85]. Each LAB3 chip has 8 input channels (and a single clock channel),

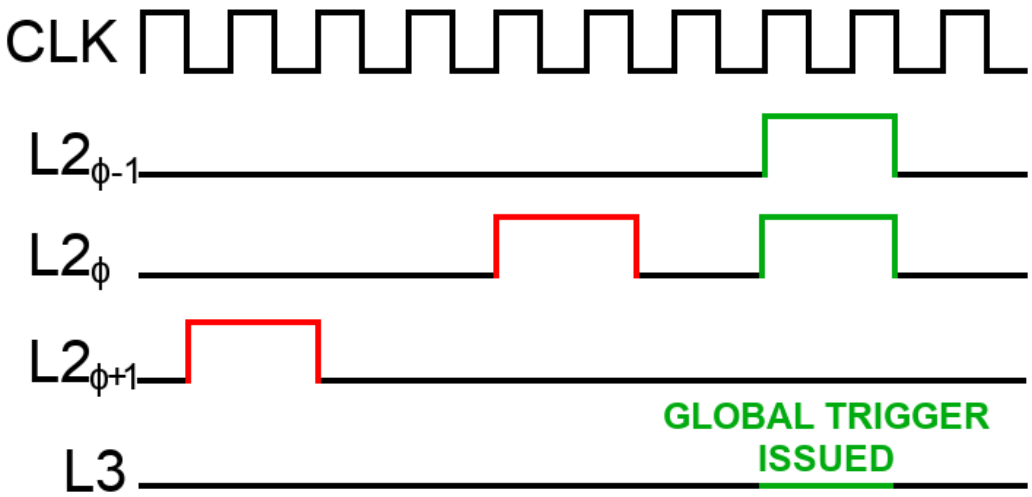

Figure 5.22: L3 trigger logic for a span of three adjacent phi-sectors. CLK represents the clock. The central phi-sector is represented as $L2_{\phi}$, and the phi-sectors either side are represented as $L2_{\phi-1}$ and $L2_{\phi+1}$. If two or more L2 trigger windows overlap in adjacent phi-sectors, an L3 trigger is issued.

and each of these input channels has a 260-capacitor Switched Capacitor Array (SCA). The SCA has a sampling rate of 2.6 GSa/s, and as it has 260 capacitors, it records a signal waveform of length 100 ns. Though sampling is done at 2.6 GSa/s, digitisation is done at a much lower rate ($\sim$ 1 MSa/s) to save power. When a TURF issues an L3 signal, a firmware block known as the Event Controller sends a `Hold` and `Digitise` signal to the SURF. A LAB3 chip within the SURF is commanded to stop sampling and the voltage is recorded by an Analog-to-Digital Converter (ADC), which is part of the LAB3 chip. The digitised event is sent to the flight computer, where it is saved to a storage unit.

5.1.6.1 Digitisation Dead Time

When a global trigger occurs (an event is to be kept), a LAB3 chip stops sampling whilst the associated data is read and digitised. Digitising such information is not instant, so a LAB3 chip needs to spend time "held". One of the other three LAB3 chips in the same SURF can continue digitising events whilst that one is occupied. If, however, the other three LAB3 chips have been commanded to hold and digitise as well, it is not possible to record further events. The duration in which all of these LAB3 chips are held is known as *digitisation dead time*. In particularly high

areas of activity, the global trigger rate can become greater than the rate required to digitise events. [86]

5.1.7 Position and Attitude Determination Systems

As well as recording the waveforms of actual events, ANITA records its position (latitude, longitude, and altitude) and attitude (heading, pitch, and roll) information throughout the flight. The ANITA payload is attached to a Long Duration Balloon (LDB), and once off of the launch vehicle, has no way to alter its own course. When the payload rises in altitude such that it reaches the lower stratosphere, the winds of the Antarctic polar vortex fully determine ANITA's flight path. This makes it impossible to exactly predetermine where ANITA will travel around the Antarctic continent. It is possible that a payload could veer off course and plummet into the depths of the Southern Ocean, causing data to be lost. Having Global Position System (GPS) information allows us to precisely point events to the continent, as well as to terminate the experiment in a recoverable location. It is useful to have precise position and attitude information for several other reasons, both during flight and analysis. As such, ANITA employs not only independent GPS systems, but also multiple backup systems.

5.1.7.1 GPS Antennas

ADU5s

The primary GPS antennas used by ANITA are the fifth-generation Attitude Determination Units (ADU5s). There are two separate ADU5 systems, known as ADU5A and ADU5B; each system has four antennas positioned on each corner of the upper gondola frame (shown in Figure 5.1). Each individual antenna is named after its system (ADU5A, ADU5B) followed by its position (aft, fore, port, and starboard), i.e., ADU5B fore. Each antenna is able to measure the payload's position and the associated time, and the combination of four antennas allows attitude to be measured.

G12

In addition to the eight ADU5 antennas, ANITA utilised another system, known

as the G12 unit. This antenna only reports position, but not attitude. The G12 antenna records this information at the same rate as the ADU5 systems, which is, at maximum, 20 Hz.

5.1.7.2 Backup Position and Altitude Systems

In the case that both ADU5 systems and the G12 unit will stop functioning correctly during flight, several backup systems are also installed. These are the magnetometer, sun sensors, and the accelerometer, all of which were fixed onto the payload deck.

The magnetometer records the magnitude and direction of Earth's magnetic field at a given time. Provided with both an estimate of the altitude and an appropriate model of the geomagnetic field, it becomes possible to approximate the position of the payload.

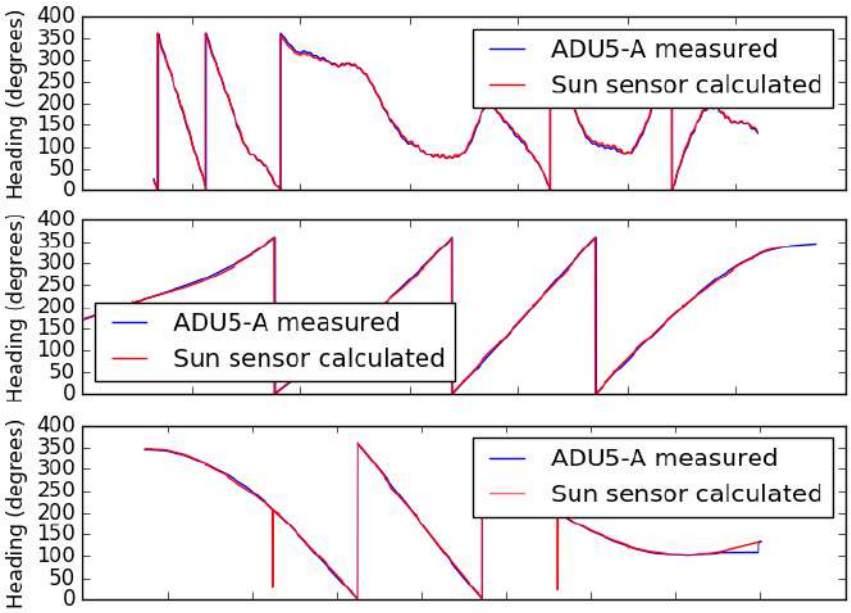

Figure 5.23: Comparison between the primary system (ADU5A) and backup system (sun sensor) measurements of heading for three individual runs [87].

A series of sun sensors can, in addition to the timing information, calculate the heading of the payload. As ANITA flights always take place during the austral summer, the Sun never goes below the horizon. Therefore, we always have a relative position measurement. Comparison between heading determination of the primary

system and the sun sensors are shown in Figure 5.23.

An accelerometer was the final backup system; it can be used for acceleration measurements and orientation determination.

Luckily, ANITA-4 did not have any GPS failures throughout its whole flight, so none of these backup systems were used in analysis, as the information from the primary systems was complete.

5.1.7.3 Minimum Bias Triggers

The primary GPS systems are not solely used for position and attitude determination. In order to search out science signals, ANITA records waveforms of events that have passed the final level of the trigger system. As such, these events are biased. However, these signals are not the only ones recorded by ANITA, and are far from the average event received by the quad-ridged horn antennas. To get an idea of the noise background, a trigger is forced every second, whatever the shape of the waveform. Each forced trigger is known as a *minimum bias trigger*, and the associated events are known as *minimum bias events*. In order to force the triggers to record these events on a regular basis, the ADU5 and G12 units issue a Pulse-Per-Second (PPS) trigger.

5.1.8 Instrument Box

The instrument box holds most of the vital electronics required by ANITA-4. As such, it houses all parts of the triggering system, the digitising system, and the flight computer, all of which are stored in the compact Peripheral Component Interconnect (cPCI) crate. The instrument box also houses the vital disk storage systems and the iRFCMs, as well as a set of temperature sensors. The instrument box is located on top of the payload deck.

5.1.9 SIP

The other large box on the payload deck is the Support Instrument Package (SIP). A SIP is used in every payload launched by NASA's LDB program. The SIP has four communication lines to Earth: 2 Tracking and Data Relay Satellite Systems (TDRSS), known as FAST-TDRSS and SLOW-TDRSS, Iridium, and Line-Of-Sight

(LOS). The former three channels allow one to monitor the payload whilst it is far away from the LDB facility via satellite communications. The LOS channel only functions when the payload is within range of LDB, where ANITA was launched and the payload was actively monitored. LOS has a much higher data rate than all other communication lines, but can only be used when the payload is within 500 km of LDB. However, only approximately 1% of data collected by ANITA can be sent via telemetry.

As well as communication systems, a SIP is fitted with its own housekeeping systems, such as temperature trackers and a GPS. The SIP is enclosed in a Faraday shield (customed designed by the ANITA collaboration) to block electromagnetic interference. Ultimately, the SIP is responsible for the termination of the flight by separating the ANITA payload from the balloon.

5.1.10 Power Systems

In order to keep ANITA-4 alive whilst it is 40 km above the Antarctic continent, the payload was designed to have a low power budget. The required power to keep ANITA in operation is approximately 600 W, which is less than that needed to power a conventional microwave oven. As the flight takes place during the austral summer, the Sun is always in view, and can provide the energy ANITA requires. Making use of this, ANITA-4 used 32 solar panels, arranged into 2 sets. The uppermost set, of which there are 8 panels, is attached to the top of the gondola. The bottommost set of 24 panels has to be deployed during flight, due to the size restraints of the launch vehicle (see Figure 5.1). This is done by a remote signal sent to the SIP, which releases the strings tied to the panels. Once deployed, these solar panels can be used to power the main instrument box.

Although the Sun is always in view, the energy received changes with time due to the daily variation in the Sun's elevation. In order to ensure that there is enough power to keep ANITA working, particularly at times of low incoming solar energy, a battery farm is used. The battery farm consists of 12 rechargeable lead-acid batteries, and the model used is the LC-X series 12V 20Ah Panasonic (LC-X1220P). A charge controller is used to direct where the power goes. When extra

power is needed, power is taken from the batteries and delivered to the the necessary systems. When the solar panels can fully power the system, the excess energy is used to charge these batteries.

5.1.11 Storage

When a single event is flagged for recording, either via passing the global trigger or forced by a minimum bias trigger, the associated waveforms need to be stored. Due to data rate limitations of communication channels, it is impossible to telemeter all of the events. For this reason, events are recorded on-board by several units and two types of storage drives, should an individual board or drive type fail.

The first two drives have the largest storage capacity and are of the Hard Disk Drive (HDD) variety. The HDD model used is the Hitachi Global Storage Technologies (HGST) He8 Ultrastar Helium 3.5″, which has an 8 TB storage capacity. Instead of recording half of the total events each, they store the same copy of the data, should one specific drive fail. The second drives used, of which there are six, are of the Solid State Drive (SSD) variety, and each has 1 TB of storage space.

In case both sets of storage devices fail, or their data becomes corrupted, the highest priority events are telemetered from the payload to the ground (see Section 5.1.9). As it is only possible to telemeter approximately 0.1% of the data taken, a prioritizer was used to establish which events were classed as the highest priority. Each triggered event was subject to the Event Prioritizer, which performed analysis to determine the most signal-like events. To find out more information about the prioritizer, see Ben Strutt's doctoral thesis [88].

After a flight is terminated, the storage devices and other essential parts of the experiment are recovered. For ANITA-4, both the recording of the data and its recovery were successful.

5.2 Pre-Launch

5.2.1 Construction

Once all the electronics and components are tested, they can be installed and fitted. The gondola, the skeletal frame of ANITA, is the first thing that is constructed,

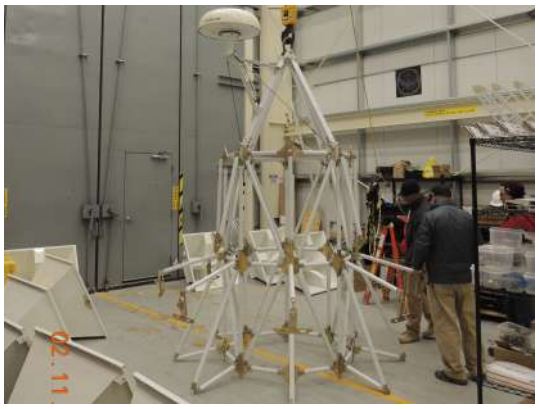

(a) The initial frame of the gondola is constructed and the Iridium antenna is installed.

(b) The top ring antennas are mounted.

(c) A crane hook is used to lift the top ring of the payload so that the rest of the gondola can be constructed. The solar panels are fitted on top.

(d) The neck and middle ring of the gondola are constructed.

(e) The middle ring antennas are mounted. After this, the bottom ring of the gondola is constructed, and the final antennas and solar panels are attached.

(f) The instrument box and SIP are fitted onto the payload deck. ANITA is then loaded onto the launch vehicle: *the Boss*.

Figure 5.24: The assembly of ANITA.

(a) The vacuum test: ANITA is subjected to vacuum conditions within the thermal-vacuum chamber.

(b) The hang test: ANITA is suspended by a crane hook and checked to see if its core systems do not interfere with those of CSBF.

Figure 5.25: Pre-flight tests at CSBF.

and is built from the top-down. When each subsequent part of the gondola is constructed, the antennas and other components are fitted. See Figure 5.24 for the construction of the ANITA payload broken down into six steps.

5.2.2 Pre-Flight Tests

When the payload is fully constructed, it has to undergo a series of tests before it is ready to fly. The first pre-flight test was carried out at CSBF, and involves a thermal-vacuum chamber. During this test, the instrument box is put into a large vacuum chamber. The electronics are subject to vacuum conditions, similar to that which they would experience in the lower stratosphere. In such conditions, it is possible for electronics to overheat due to the changes in external temperature and pressure.

The other pre-flight test is known as the *hang test*, where the ANITA payload is hung from a crane and sent radio pulses from a test antenna. Each antenna channel is monitored to check that pulses are correctly received. This also serves as a compatibility test to make sure that the ANITA systems don't interfere with those of CSBF, and vice versa. The compatibility test was repeated at the LDB Facility in Antarctica. Images of both pre-flight test setups are shown in Figure 5.25.

Figure 5.26: ANITA launch preparation setup. The Long Duration Balloon is inflated with helium (left), and a small test balloon (top left) is launched to monitor wind speed. The base of the Long Duration Balloon is attached to the spool vehicle; this balloon is connected to the ANITA payload (with a parachute interconnect), which is on *the Boss* (right).

5.3 Launch

The launch preparation and actual launch of ANITA-4 were carried out by NASA support personnel. To prepare for launch, the launch vehicle and the spool vehicle are first driven to the launch site. A long protective cloth is spread across the ground to prevent damage to the balloon. The base of the Long Duration Balloon is held by the spool vehicle, whilst the flight crew attaches it to a parachute. The parachute is then attached to the payload, which is held by the launch vehicle.

5.3.1 Preparation and Ascension

The balloon, which is made of ultra-thin (20 micrometres thick) polyethylene film, has helium valves attached to it. It is then filled up with helium until the predetermined amount is reached (Figure 5.26 shows the balloon being inflated prior to launch). The movement of a small test balloon indicates the wind speed. When the wind speed is low and steady, and there is minimal potential for adverse weather conditions to arise during launch and early ascension, the spool is released. Once the balloon string is vertical, the payload itself is separated from the launch vehicle, taking flight. ANITA-4 successfully launched at 8:00 am on the 2nd of December 2016 Coordinated Universal Time (UTC). The launch images are shown in Figure 5.27.

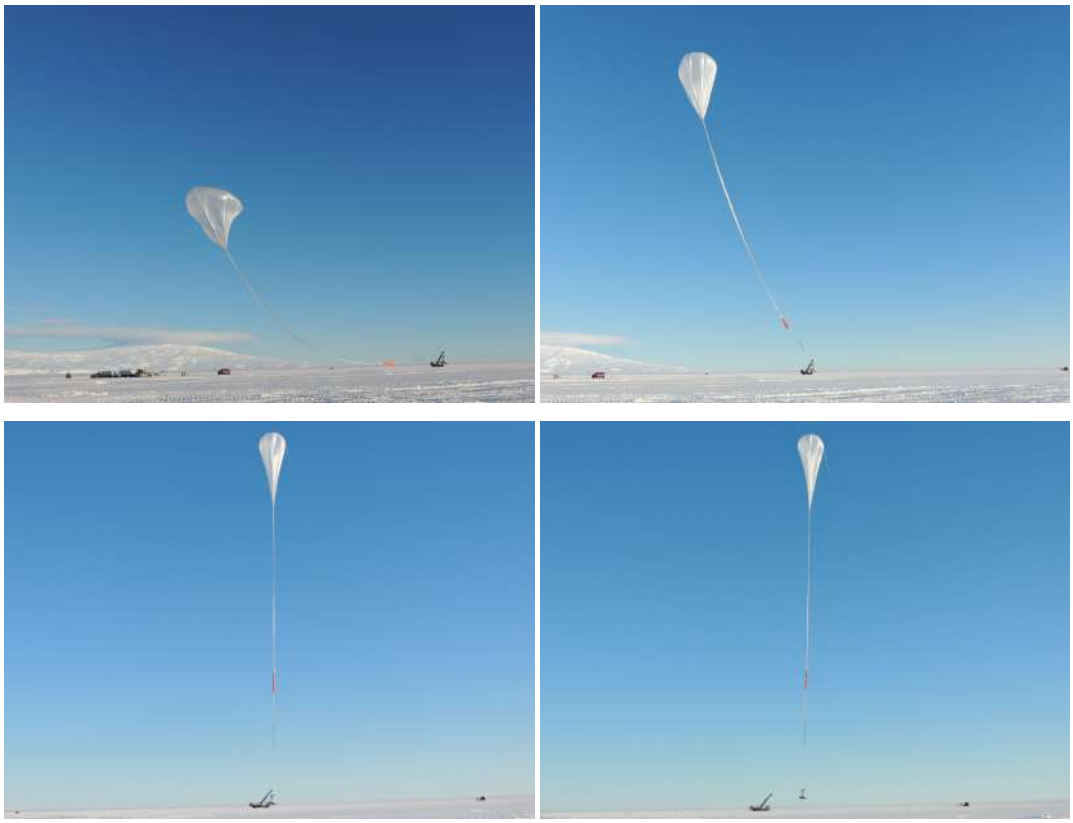

Figure 5.27: ANITA launch sequence. When the helium balloon is sufficiently full to carry the payload into the air, it is released from the spool vehicle. The balloon then rises until it is completely vertical above payload. At such a point, ANITA is detached from the launch vehicle, and begins to ascend.

5.4 Post-Launch

5.4.1 Monitoring

After the successful launch of ANITA-4, it was closely monitored by on-ground systems. The primary system is the Active Web for Antarctica Radio Experiments (AWARE), a program used to read in and plot data from the payload as it moves, and also to look at data from the past. AWARE works by collecting such data via telemetry, and then processing it with `C++ ROOT` programs [89]. The output of these programs are JavaScript Object Notation (JSON) files, which are handled by jQuery and plotted by Flot. The results are then displayed on the AWARE website. AWARE has several useful features for those monitoring the payload during the flight. First of all, it features the waveforms of triggered events, as well as plots of housekeeping data, such as temperature, power, notch filter status, and position

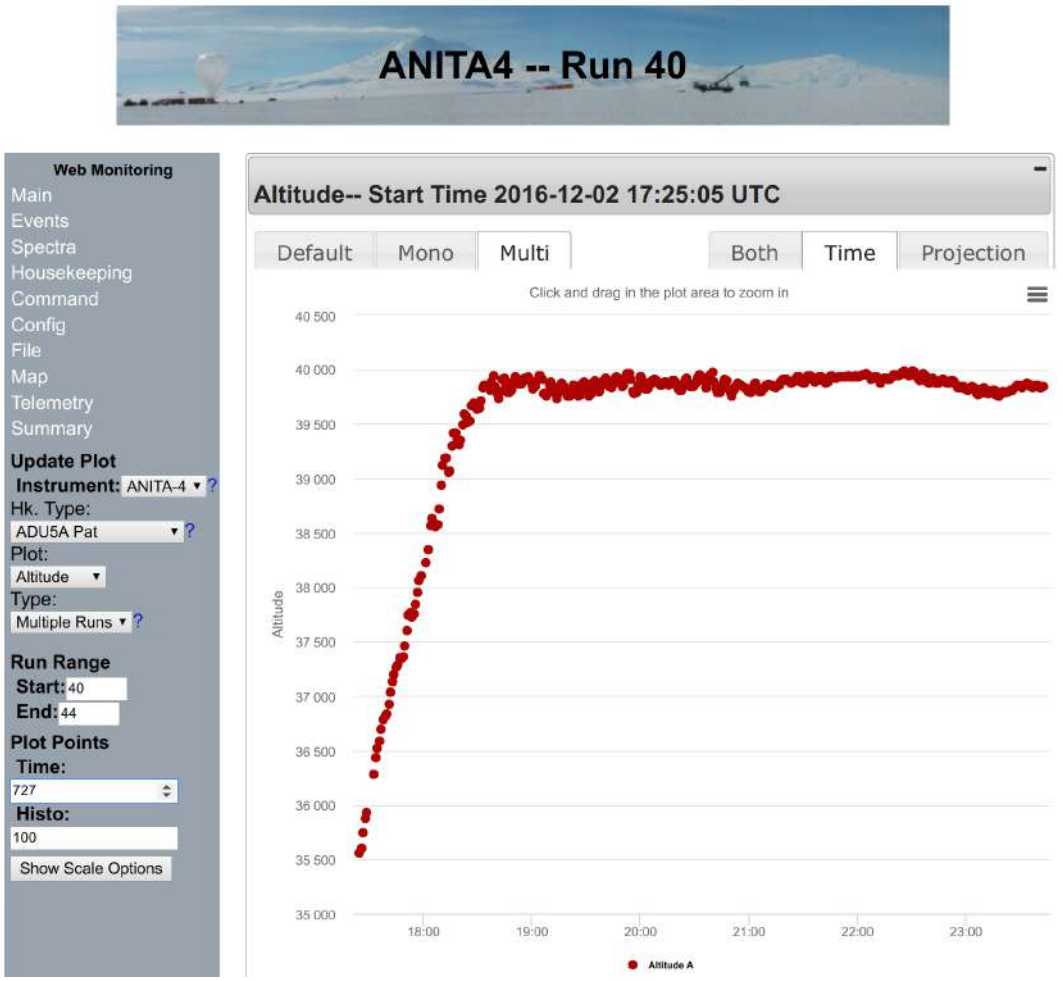

(a) Housekeeping data: shown is the payload's altitude from run 40 to 44, where ANITA reached a maximum height of about 40 km, and then stabilises. Information is from the ADU5A unit, which is processed and then telemetered to the ground.

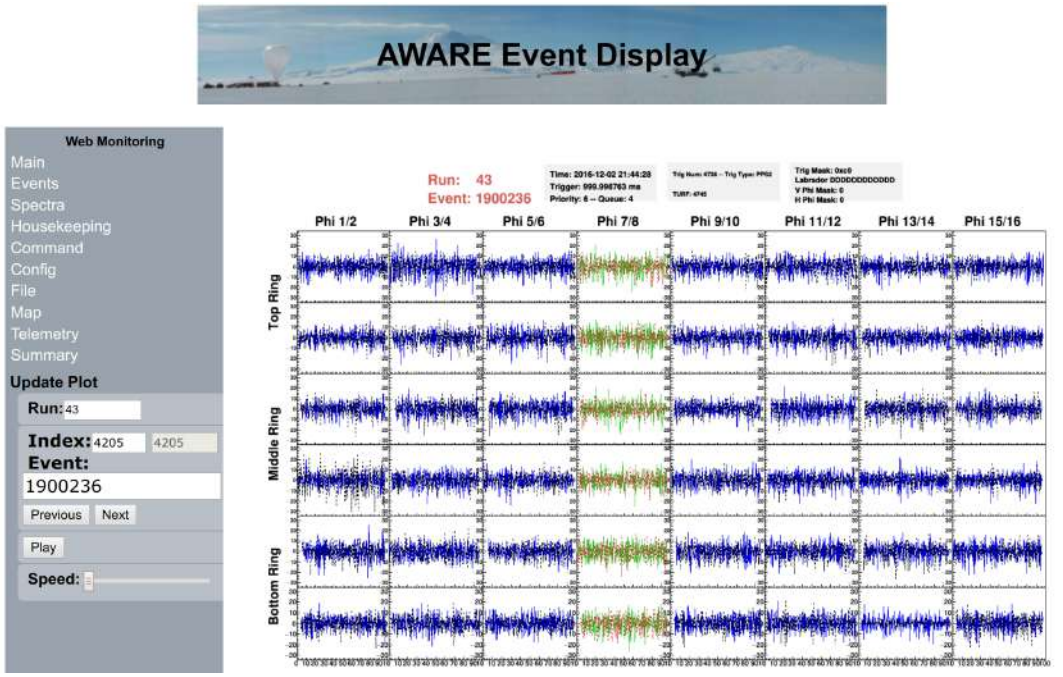

(b) Event display: A low-priority event received when ANITA's altitude became stable.

Figure 5.28: An overview of the AWARE interface.

information (as shown in Figure 5.28. The status of all telemetry systems can be checked from AWARE. AWARE also plots a map of ANITA's flight path in real time, as well as event locations.

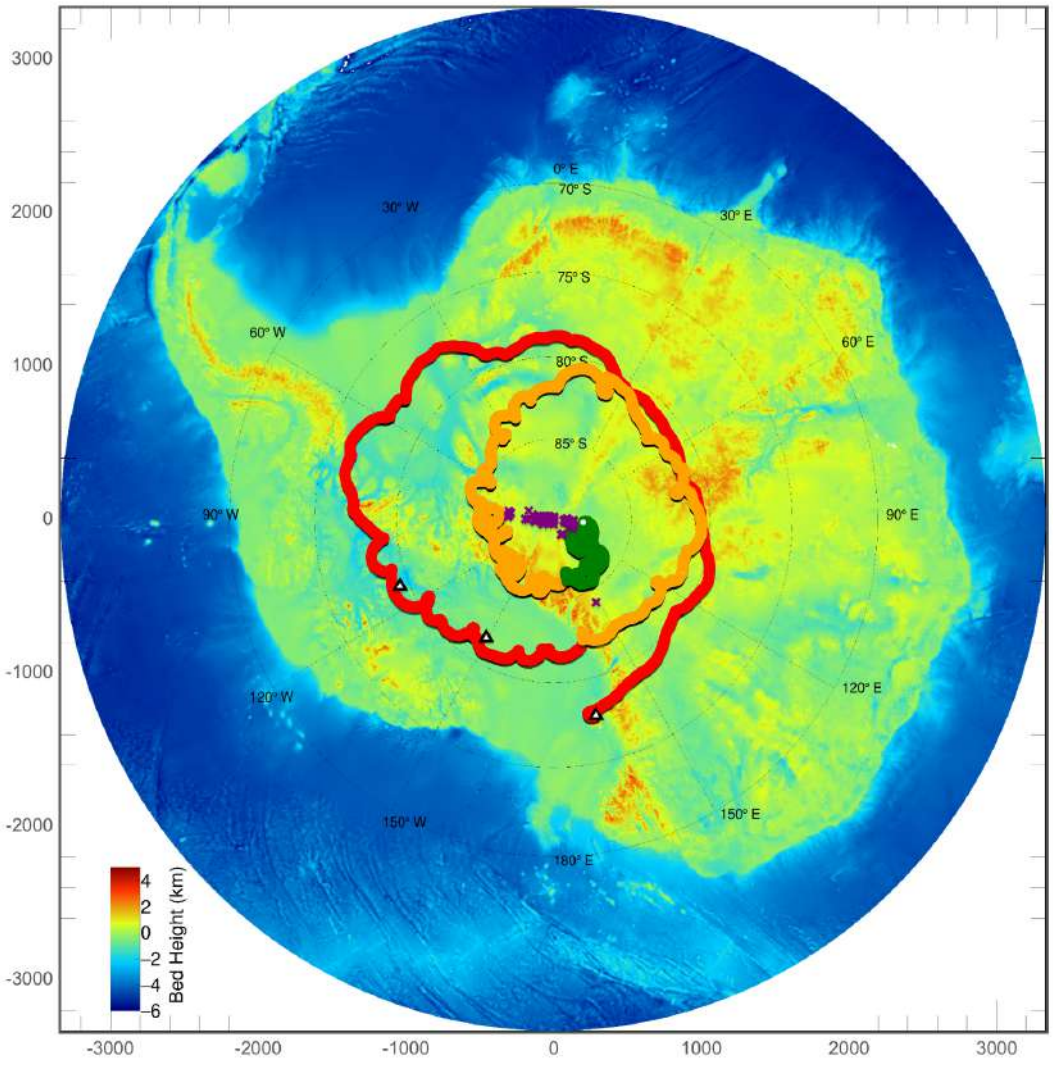

Figure 5.29: The full ANITA-4 flight path over the BEDMAP antarctic map, as shown in the AWARE web-monitor. The purple crosses show the highest priority events coming from the ice at a single time near the end of the flight. The white triangles show locations of potential pulser bases, such as the LDB facility and the West Antarctic Ice Sheet (WAIS) Divide.

A secondary program was also used in addition to AWARE, known as WebANITA. Instead of using `C++`, it is Python-based, and makes use of Bootstrap and HTML5. Using AWARE and WebANITA, the payload was monitored 24/7 by physicists around the world.

5.4.2 Flight

As the payload-carrying balloon increases in altitude, the balloon becomes much larger in volume. An increase in altitude means a decrease in atmospheric pressure, leading to a much greater balloon size. Once ANITA reaches a steady altitude, the balloon becomes approximately 1,100,000 m^3 in volume, the same size as the bowl volume of Wembley Stadium [90]. ANITA-4 reached a steady altitude 5 hours after launch (See Figure 5.28a).

During its time in the air, ANITA completed two loops around Antarctica. Figure 5.29 shows the full flight path in the AWARE viewer: the red (first) loop shows the payload's passage from the LDB facility until the start of the second loop (orange). ANITA then spirals around again until it starts to begin the final loop (green). The flight path was very favourable, with much time spent over dense regions of ice. At 5:00 am on the 30th of December (UTC), the ANITA mission was terminated. The full ANITA-4 flight lasted for almost 28 days.

Figure 5.30: The ANITA-4 crash site. A team was dispatched to recover the hard drives and other components, such as the AMPAs. Photograph by Christian Miki.

Chapter 6

Simulation

Ultra-high energy (UHE) particles, particularly neutrinos, can prove difficult to detect due to their low flux. However, it is possible to simulate the physical processes that involve the production, propagation, and detection of neutrino-induced radio waves, as well as the interactions of the primary neutrinos. A complete simulation of this type allows one to predict what a radio-array such as ANITA should observe from a set of astrophysical neutrinos interacting with the ice sheet. One such simulation is `icemc` [91], which allows the user to simulate Askaryan radiation from UHE neutrinos and the subsequent detection by ANITA, another is `CRPropa` [92], a cosmic-ray propagation framework that can be used to model the UHE neutrino flux.

6.1 icemc

`icemc` is an existing `ROOT`-based `C++` program that uses Monte Carlo techniques to simulate the interactions of neutrinos with the Antarctic ice sheet, as well as the propagation and detection of radio waves. This program simulates neutrinos in the EeV-ZeV energy range, whilst most other frameworks (particularly those used in accelerator experiments) work with lower energy neutrinos, typically in the MeV-TeV (or PeV) range [93, 94]. Instead of tracking the paths and determining the interactions of every single particle in these UHE neutrino-induced showers, `icemc` uses a parameterisation of the shower's electric field. This makes it far easier to determine the peak electric field strength of the Askaryan radiation. In

addition, the response of the ANITA instrument and its full signal chain is built-in, and the simulation generates output files in the same format that is used by existing ANITA software. For these reasons, `icemc` is the primary simulation program used throughout the analysis.

6.1.1 Diffuse Neutrino Simulation

Neutrinos that are not confined to come from a single point in space are known as *diffuse*. The diffuse part of the simulation framework has been built up and improved upon by multiple contributors since its inception prior to the first ANITA flight. In order to simulate the detection of diffuse neutrinos, the `icemc` simulation is broken down into a few steps: setup and event generation, neutrino interaction point and direction determination, signal generation, signal propagation, and signal detection.

6.1.1.1 Setup and Event Generation

In order to simulate ANITA's detection scenario, simulated neutrinos must first be assigned properties before they move through the ice, and the properties of the detector must be set. To do this, the user must specify input options and parameters via a configuration file. The event input/output subset determines the number of neutrinos to be simulated, and the energy scheme is chosen (a monoenergetic set of neutrinos can be simulated, or a flux spectrum can be specified). Several different neutrino energy spectra are available for use in `icemc`, as well as custom spectra, should one desire (see Section 6.2). The payload subset is used to select which payload to use in the simulation, i.e., ANITA-3 or ANITA-4. In addition, the position and altitude of the payload is assigned: the payload can follow a specific flight path or be held static at a certain position along it.

The neutrino interaction within the ice sheet must be modelled, so a model of the Antarctic continent is used. `icemc` uses two independent models of the continent: the high resolution model is the Bed Topography of the Antarctic (BEDMAP) [81], and the lower resolution model is the New Global Crustal Model at 2x2 Degrees (CRUST 2.0) [82]. BEDMAP uses a range of different techniques to map

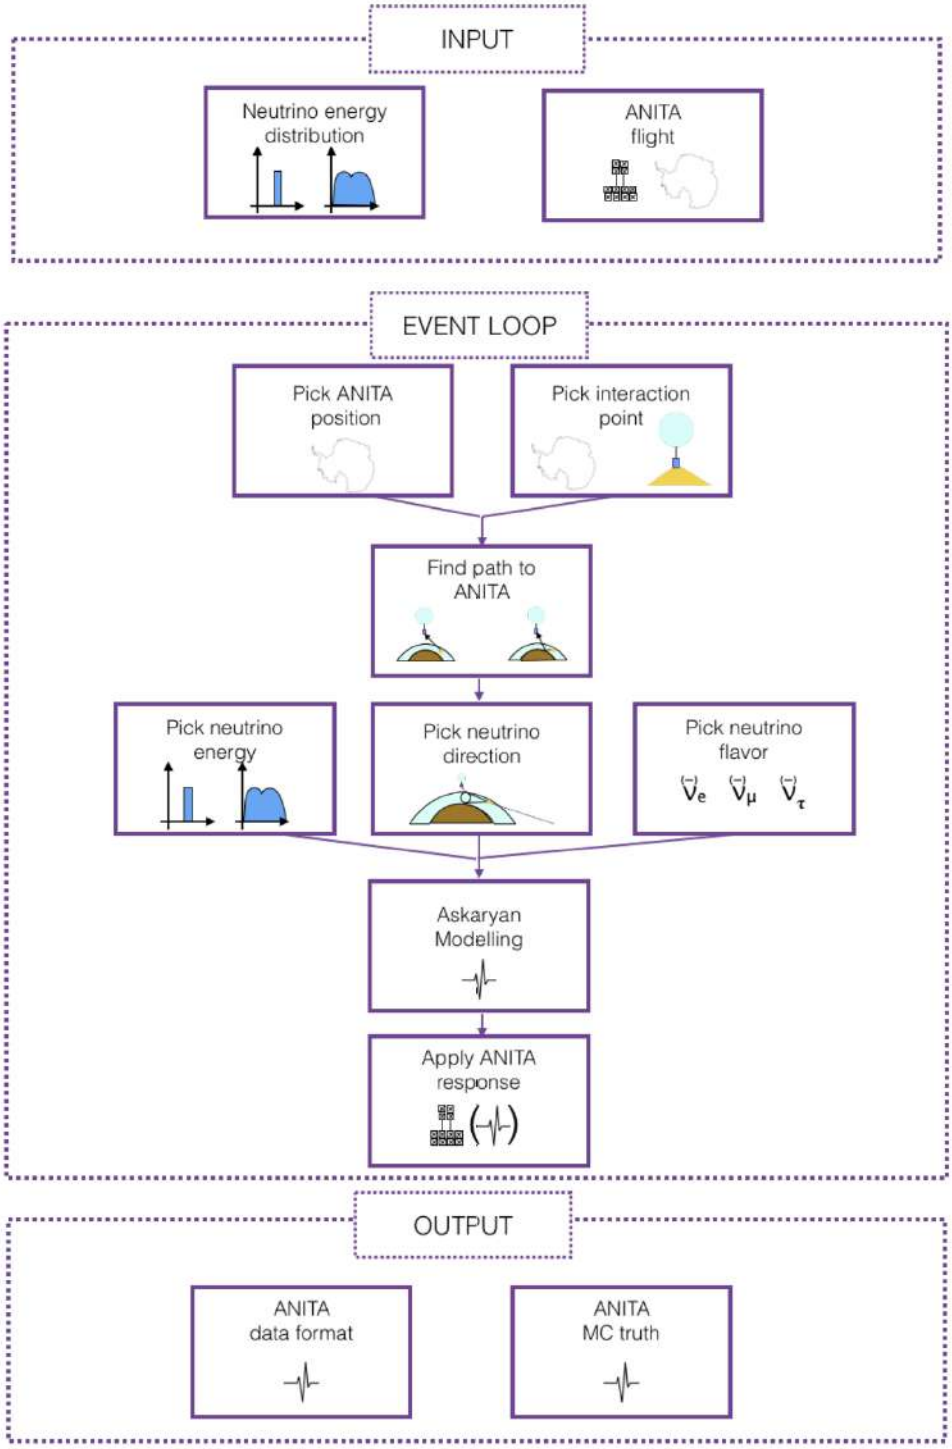

Figure 6.1: Flowchart of the icemc simulation for diffuse neutrino searches. Image from [91].

the thickness of ice, and its nominal spatial resolution is 5 km. CRUST 2.0 uses a $2^\circ \times 2^\circ$ worldwide grid to predict the density of ice, water, sediment, and crust.

Because very large numbers of neutrinos need to be simulated, the (lower resolution, and thus faster) CRUST 2.0 model is used throughout the simulation. The total volume of ice found by `icemc` using the CRUST 2.0 only differs by 1.15% compared with the BEDMAP model. [83] Other settings for the properties of ice can be adjusted in the input files, such as the depth-dependent refractive index.

Further setting subsets include modifications to the system's trigger properties, signal and noise adjustments, interaction properties, and input test scenarios. For ANITA-4, the notch settings of the TUFFs (see Section 5.1.4.1) can also be adjusted.

6.1.1.2 Neutrino Interaction Point and Direction Determination

Now that the detector properties are set up and the neutrino energy scheme is known, the location of the neutrino's interaction within the ice needs to be chosen. If a static payload location is not chosen, a random point along its specified flight path is used per simulated neutrino. If all neutrino interactions within the ice were considered, and all paths were propagated, it would take an extraordinarily long time for the simulation to run. Instead, each event is simulated under conditions where ANITA can potentially see the resulting neutrino-induced Askaryan radiation. For ANITA to actually see the incoming radio wave at this point, the interaction location of the neutrino within the ice is simulated within the horizon of the payload (typically 700 km from the payload). In addition to the interaction position of the neutrino, the neutrino's direction needs to be chosen. The direction of the neutrino will determine the vector of the Cherenkov cone's axis. Again, to see the event, the vector of the Cherenkov cone needs to be close in direction to that of the vector of the interaction point to the payload. A neutrino with its interaction point already determined can only point back to a limited range of points on an annulus on the sky.

Event Weights

Every event is assigned a weight, which takes into account the survival chance of the neutrino and the unused phase space (see the explanation of phase weight below). During analysis, events with the lowest weights can effectively be rejected. Accounting for weights is crucial in calculating the volumetric acceptance, which is itself required to find the projected limits on neutrino flux. The two quantities, the

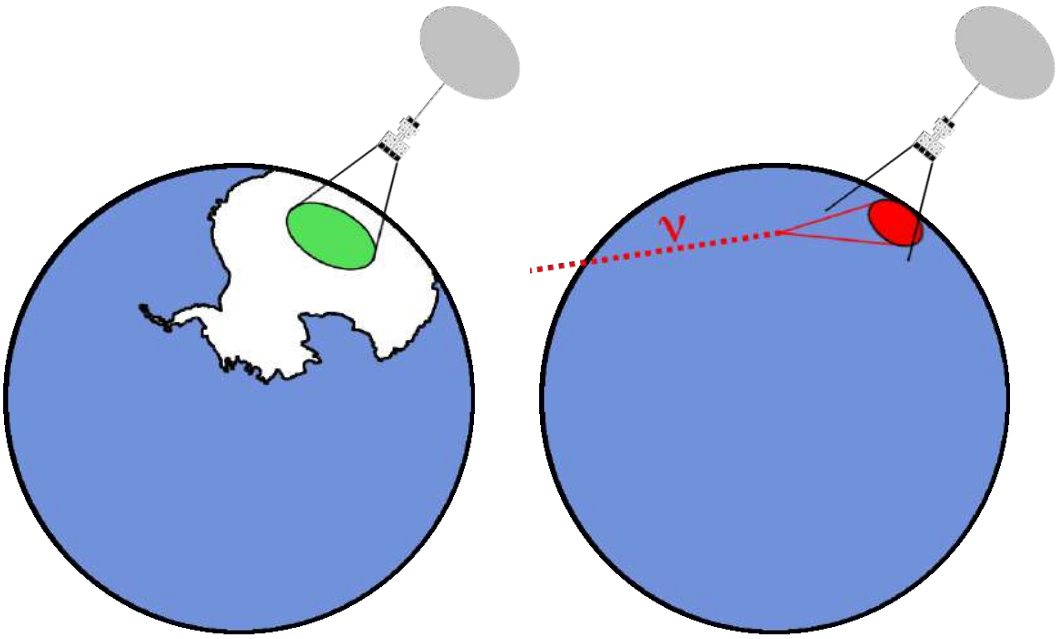

(a) The position weight is calculated by the ratio of the volume of ice used to simulate events (green) to the total volume of ice in Antarctica (white).

(b) The direction weight is calculated by the ratio of the solid angle formed by the Cherenkov cone (red) to that of a sphere (blue).

Figure 6.2: Phase weight diagrams.

neutrino interaction point and the direction of the neutrino, can be used to calculate the *phase weight* of an event, which contributes to the overall weight.

The phase weight is itself the product of two different, independent weights: the position weight, w_{pos}, and the direction weight, w_{dir} (see Figure 6.2 for an illustrative description of the phase weights). The position weight accounts for the fact that `icemc` only simulates neutrinos within a certain volume of ice, i.e., within ANITA's observable range (the horizon), at one single time. Thus, for every event, the position weight is defined as the ratio of the observed ice volume, V, to the total volume of Antarctic ice, V_{tot},

$$w_{\text{pos}} \equiv \frac{V}{V_{\text{tot}}}\,. \tag{6.1}$$

The direction weight accounts for only simulating events with good directional topology, i.e., the axis of the event's Cherenkov cone is close in direction to the line from the interaction point to the payload. Therefore, this weight is calculated as the

ratio of the solid angle formed by the Cherenkov cone, Ω, to the solid angle of a sphere, Ω_{sph},

$$w_{\text{dir}} \equiv \frac{\Omega}{\Omega_{\text{sph}}}\,. \tag{6.2}$$

For an event:

$$\Omega = \int d\Omega, \tag{6.3}$$

$$= \int_{\theta_{\text{ch}}-\theta_{\text{th}}}^{\theta_{\text{ch}}+\theta_{\text{th}}} \sin(\theta) d\theta \int_0^{2\pi} d\phi, \tag{6.4}$$

$$= (-\cos(\theta_{\text{ch}} + \theta_{\text{th}}) + \cos(\theta_{\text{ch}} - \theta_{\text{th}})) \cdot 2\pi, \tag{6.5}$$

where θ_{ch} is the Cherenkov angle, θ_{th} is the threshold angle around the Cherenkov angle (used for sampling θ), and ϕ is the azimuthal angle. Using the trigonometric identities: $-\cos(a) + \cos(b) = -2\sin\left(\frac{a}{2} - \frac{b}{2}\right)\sin\left(\frac{a}{2} + \frac{b}{2}\right)$, and $\sin(-a) = -\sin(a)$, the equation reduces to,

$$\Omega = 4\pi \sin(\theta_{\text{th}}) \sin(\theta_{\text{ch}})\,, \tag{6.6}$$

and therefore the direction weight is,

$$w_{\text{dir}} = \frac{4\pi \sin(\theta_{\text{ch}}) \sin(\theta_{\text{th}})}{4\pi}, \tag{6.7}$$

$$w_{\text{dir}} = \sin(\theta_{\text{ch}}) \sin(\theta_{\text{th}})\,. \tag{6.8}$$

As well as the phase weight, the *absorption weight* needs to be taken into account. The absorption weight, w_{abs}, arises from the chance of the neutrino being absorbed as it moves through the Earth. Depending on its path, a neutrino will move through media of different densities, such as water and bedrock. For every medium the neutrino passes through, there is an associated interaction length, l, calculated by,

$$l_j = \frac{1}{\sigma_j \rho_w}, \tag{6.9}$$

where σ_j is the cross section, and ρ_w is the density of water. j represents the medium/layer the neutrino is passing through. Thus, the probability of it being absorbed in that layer is,

$$P_j = \exp\left(\frac{-d_j}{l_j}\right), \tag{6.10}$$

where d_j is the distance travelled through that layer.

Thus, for the full path of a single neutrino,

$$P = \prod_j P_j = \prod_j \exp\left(\frac{-d_j}{l_j}\right) = \exp\left(-\sum_j \frac{d_j}{l_j}\right), \tag{6.11}$$

and so the absorption weight is,

$$w_{\text{abs}} = \exp\left(-\sum_j \frac{d_j}{l_j}\right). \tag{6.12}$$

Making use of this combination of weights makes sure that only the most favourable paths and interactions are simulated, as only considering events with the best topologies will generate a signal detectable by the payload.

6.1.1.3 Signal Generation

As an ultra-high energy neutrino moves through the ice sheet, it can produce a shower, with electromagnetic and/or hadronic components, giving rise to Askaryan radiation (see Section 4). To simplify things, it is possible to use a parameterisation to find the peak of the Askaryan radiation, instead of modelling the entire cascade. The electric field strength will scale with the energy of the shower, and is dependent on the frequency of the associated Cherenkov radiation. If the peak of the shower is observed at exactly the Cherenkov angle, the electric field strength will be calculated as [95]:

$$\epsilon_{\mathrm{ch}}^{1\mathrm{m}}(E_{\mathrm{sh}}, f) = 2.53 \times 10^{-7} \frac{E_{\mathrm{sh}}}{\mathrm{TeV}} \frac{f}{f_0} \frac{1}{1 + \left(\frac{f}{f_0}\right)^{1.44}}, \tag{6.13}$$

where $\epsilon_{\mathrm{ch}}^{1\mathrm{m}}(E_{\mathrm{sh}}, f)$ is the electric field strength one metre away from the interaction point, E_{sh} is the energy of the shower, f is the frequency of the Cherenkov radiation, and f_0 is the peak frequency at the Cherenkov angle, which is 1.15 GHz. As such, this peak frequency corresponds to the wavelength of the Molière radius of the particle shower.

The angle of the Cherenkov radiation will often not be exactly aligned with the actual view angle, which is the angle between the shower axis and the observer. In cases such as this, the magnitude of the electric field strength will, of course, decrease. In order to account for this, the original parameterisation is modulated with a Gaussian and a sinusoidal function. This makes the equation dependent on view angle [96]:

$$\epsilon_{\mathrm{v}}^{1\mathrm{m}}(E_{\mathrm{sh}}, f, \theta_{\mathrm{v}}) = \epsilon_{\mathrm{ch}}^{1\mathrm{m}}(E_{\mathrm{sh}}, f) \cdot \frac{\sin(\theta_{\mathrm{v}})}{\sin(\theta_{\mathrm{ch}})} \cdot \exp\left[-\left(\frac{\theta_{\mathrm{v}} - \theta\mathrm{ch}}{\Delta\theta}\right)^2\right], \tag{6.14}$$

where $\epsilon_{\mathrm{v}}^{1\mathrm{m}}(E_{\mathrm{sh}}, f, \theta_{\mathrm{v}})$ is the electric field strength observed observed at view angle θ_{v} one metre from the interaction point, θ_{ch} is the Cherenkov angle, and $\Delta\theta$ is the width of the Cherenkov cone. If the view angle is set to the Cherenkov angle, the sinusoidal functions cancel out, and the argument of the exponent becomes zero. In such a case, Equation 6.14 trivially reduces to the Equation 6.13.

The width of the Cherenkov cone for initially electronic showers are different from those that are initially hadronic. Electromagnetic showers have a reduced width due to the Landau-Pomeranchuck-Migdal (LPM) effect [98–100]. As the LPM effect suppresses bremsstrahlung at very high energies, the shower becomes elongated, leading to a reduced width [96]. On the other hand, the calculation of the width in hadronic showers is dependent on how much energy actually contributes to the signal. Hadronic showers deposit some energy into particles such as muons

and neutrinos, which do not contribute to the Askaryan signal, so the width of the shower is reduced [97]. As a result, the electromagnetic and hadronic components of showers are evaluated independently.

Before `icemc` calculates the corresponding electric field strength, it checks that the view angle is within an appropriate range of the Cherenkov angle. If this were not to be taken into account, the electric field strength would have a chance to be largely suppressed, as seen in Equation 6.14, and thus the chance of measuring an associated signal would be extraordinarily small. This contributes towards the weight of each event, and thus the lowly weighted ones will be rejected by construction.

6.1.1.4 Signal Propagation

Once the electric field strength is calculated, the propagation through ice needs to be accounted for. As the signal moves through a medium, its intensity is reduced, i.e., the signal is attenuated, due to the scattering and absorption of photons. This attenuation is an exponential function of the path length, d_{ice}, which is how far the signal has to travel to the ice-air boundary. The path length is scaled to the attenuation length of ice, l_{ice}, reducing the overall intensity of the signal:

$$I = I_0 \exp\left(\frac{-d_{\text{ice}}}{l_{\text{ice}}}\right), \tag{6.15}$$

where I is the intensity of the propagated signal, and I_0 is the original intensity before getting propagated.

Once through the ice, the signal reaches the ice-air boundary. At this point, the signal can either refract through the ice, or surpass a critical point whereby it totally internally reflects (see Section 4.4 for a discussion on this). Only in the case of refraction can a signal be detected by an instrument above the surface of the ice. When such a signal refracts from ice to air, the Fresnel transmission coefficients for the perpendicular and parallel components of the electric field must be taken into account:

$$t_{\perp} = \frac{2\sin(\theta_{\mathrm{t}})\cos(\theta_{\mathrm{i}})}{\sin(\theta_{\mathrm{i}}+\theta_{\mathrm{t}})}, \qquad t_{\parallel} = \frac{t_{\perp}}{\cos(\theta_{\mathrm{i}}-\theta_{\mathrm{t}})}, \tag{6.16}$$

where $t_{\perp}$ is the perpendicular Fresnel refraction coefficient, θ_{t} is the transmission angle, θ_{i} is the incidence angle, and $t_{\parallel}$ is the parallel Fresnel refraction coefficient.

Incorporating these coefficients (Equation 6.16), and taking into account the total path length, the (perpendicular and parallel components of the) electric field is further attenuated:

$$\epsilon_{\perp} \propto t_{\perp} \cdot \frac{1\mathrm{m}}{d_{\mathrm{ice}}+d_{\mathrm{air}}}, \qquad \epsilon_{\parallel} \propto t_{\parallel} \cdot \frac{1\mathrm{m}}{d_{\mathrm{ice}}+d_{\mathrm{air}}}, \tag{6.17}$$

where $\epsilon_{\perp}$ is the perpendicular component of the electric field, $\epsilon_{\parallel}$ is the parallel component of the electric field, and d_{air} is the path length that the signal has travelled through the air.

After the signal has surpassed the ice-air boundary, it propagates through the air until it is detected. Like in ice, this propagation needs to be accounted for. However, the attenuation length of radio waves in air is so large that the argument of the exponential (the same functional form as Equation 6.15) tends to zero. Thus, the attenuation term becomes approximately 1, and is not included in the final formulation.

Including these factors, the electric field strength of the signal when it reaches the detection instrument is shown in Equation 6.18:

$$\epsilon_{\perp,\parallel} = \epsilon_{\mathrm{v}}^{1\mathrm{m}}(E_{\mathrm{sh}}, f, \theta_{\mathrm{v}}) \cdot \exp\left(\frac{-d_{\mathrm{ice}}}{l_{\mathrm{ice}}}\right) \cdot t_{\perp,\parallel} \cdot \frac{1\mathrm{m}}{d_{\mathrm{ice}}+d_{\mathrm{air}}}\,. \tag{6.18}$$

6.1.1.5 Signal Detection

The instrument used to detect the signal is already chosen and set up beforehand (see Section 6.1.1.1). Now that the signal has been produced and has propagated through both the ice and air, the response of the instrument can be calculated.

In order to model the geometry of the payload, photogrammetry and phase-center measurements are used. The electric field of the Askaryan pulse is propagated to the front of the antennas of the ANITA payload. The gains of the antennas

that were measured (See Section 5.1.1) are applied to the incoming signals. The trigger and digitisation processes are taken into account by applying the trigger and digitisation responses as were measured before flight. The signal processing chain for the ANITA-4 instrument, including the trigger logic, is implemented in the same way as was covered in Section 5.1.

6.1.2 Source Search Integration

Previously, `icemc` mainly provided support for *diffuse* neutrino searches, i.e., neutrinos arrive from random allowed positions in space and not from specific objects. My work involves modifying the simulation to handle neutrinos that are produced by specific astrophysical objects. In order to do this, astronomical catalogues and modelling techniques were implemented into the existing simulation.

6.1.2.1 Cataloguing

Astronomical catalogues are compendia of parameters for certain types of astrophysical sources. Each catalogue has information for each associated object, such as the object's celestial location, its redshift, and its time and date of detection. The information available varies largely with different types of objects.

Two separate catalogues were integrated into the simulation for each considered source type. For blazars, the Fermi All-Sky Variability Analysis [101] (FAVA) and the 3rd Fermi Gamma-ray LAT catalogue [102] (3FGL) were considered. For GRBs, the IceCube [103, 104] and Swift [105] catalogues were used. Finally, the Transient Name Server [106] (TNS) and the Open Supernova Catalog [107] were considered for supernovae.

The user of the simulation can select a catalogue to search for a specific object, (such as SN 2016ijk), a specific object subtype, (such as SN Ia), or objects within a certain time range. This is especially useful when only looking for objects during a specific ANITA flight. It is also possible to look for a selection of multiple individual objects, and even define a custom object.

Though catalogues provide necessary parameters, we need to go a step further to determine the properties of the incoming neutrinos in specific cases.

6.1.2.2 Modelling

In order to find the properties of the neutrinos, such as their energy, I take certain theoretical models into account. In the case of GRBs, the fireball model that is used works up to a maximum neutrino energy (which is associated with the maximum shock-accelerated proton energy). For other source types, one can sample a flux distribution (as explained in Sections 6.1.2.3 and 6.1.2.4).

GRB Modelling

GRBs are associated with two types of neutrinos: prompt and afterglow (See Section 3.1.5.3). In order to find out which of the two types of neutrinos contribute to an overall higher flux, a GRB fireball model is utilised [110–112].

First, a ΛCDM flat Universe model is considered, and the Friedmann equation is used [113, 114]:

$$\left(\frac{H(z)}{H_0}\right)^2 = \Omega_\Lambda + \Omega_\mathrm{m}(1+z)^3 + \Omega_\gamma(1+z)^4, \tag{6.19}$$

$$H(z) = H_0\sqrt{\Omega_\Lambda + \Omega_\mathrm{m}(1+z)^3 + \Omega_\gamma(1+z)^4}\,. \tag{6.20}$$

where $H(z)$ is the Hubble parameter at a specific redshift (z), H_0 is the Hubble constant, Ω_γ is the density parameter for radiation, Ω_m is the density parameter for matter, Ω_Λ is the vacuum density parameter (cosmological constant). The values of the density parameters and Hubble's constant are taken from the Planck 2018 results [115]. The catalogues implemented into `icemc` provide several other parameters needed for the calculations, such as the object's redshift. The luminosity distance to each GRB can be calculated as follows:

$$d_L = (1+z)c\int_0^z \frac{dz}{H(z)}, \tag{6.21}$$

where d_L is the luminosity distance, and c is the speed of light.

From this point until the end of the section, I used the methodology from the following references: [116, 117]. However, in the calculations in this thesis, I used newer cosmological parameters and applied the methods to ANITA-4 and its asso-

ciated GRBs.

High energy sources, including GRBs, will emit electromagnetic radiation at a range of wavelengths. In order to represent the total luminosity, the bolometric magnitude of the source is used, as it takes into account each and every wavelength. In order to find a flux, the gamma-ray bolometric energy, which is the total energy released from a gamma-ray burst, needs to be calculated ($E_{\gamma,\mathrm{tot}}$). Taking into account the luminosity distance and parameters provided from the catalogue:

$$E_{\gamma,\mathrm{tot}} = \frac{4\pi d_L^2 S_\gamma}{1+z}, \tag{6.22}$$

where S_γ is the photon fluence.

The total gamma ray energy is used to find the break energies ($E_{\gamma,\mathrm{b}}$) that determine the shape of the photon spectra. The break energy of the photons is related to the total energy via the Ghirlanda relationship [108]:

$$E_{\gamma,\mathrm{b}} = \frac{3\times10^4}{1+z}\left(\frac{E_{\gamma,\mathrm{tot}}}{10^{53}\mathrm{ergs}}\right)^{0.56}\mathrm{GeV}\,. \tag{6.23}$$

Prompt neutrinos

Prompt photons arise from the internal shock of the GRB event from numerous processes (See Section 3.1.5.3). For prompt modelling, the photon spectrum can be approximated as a Band fit [109] with break energy, $E_{\gamma,\mathrm{b}}$.

$$\frac{dN_\gamma}{dE_\gamma} \propto \begin{cases} E_\gamma^{-1} & : E_\gamma < E_{\gamma,\mathrm{b}} \\ E_\gamma^{-2} & : E_\gamma \geq E_{\gamma,\mathrm{b}}, \end{cases} \tag{6.24}$$

where N_γ is the number of photons, E_γ is the energy of the photons.

Neutrinos may arise from these photons interacting with protons (p), producing a Delta baryon (Δ^+). The Delta baryon decays to protons and neutral pions, or neutrons and charged pions. The photons are produced by the decay of these neutral pions. The charged pions then go through a decay chain, where the final products are neutrinos. It is assumed that each of these final particles receive an equal share

of the energy, so a single neutrino receives a quarter of the Delta baryon's energy. Thus, the neutrino break energy ($E_{\nu,\mathrm{b1}}$) is (Appendix C shows the derivation of this break energy):

$$E_{\nu,\mathrm{b1}} = 0.015 \left(\frac{E_{\gamma,\mathrm{b}}}{\mathrm{GeV}}\right)^{-1} \left(\frac{\Gamma_{\mathrm{i}}}{1+z}\right)^{2} \mathrm{GeV}, \tag{6.25}$$

where Γ_{i} is the Lorentz factor of the internal shock.

After this, the neutrino flux is further suppressed by pions rapidly losing energy due to synchrotron radiation, resulting in a second break energy, $E_{\nu,\mathrm{b2}}$.

The corresponding neutrino spectrum is:

$$E_\nu^2 \Phi_\nu \propto \begin{cases} E_\nu & : E_\nu < E_{\nu,\mathrm{b1}} \\ \mathrm{const} & : E_{\nu,\mathrm{b2}} \geq E_\nu \geq E_{\nu,\mathrm{b1}} \\ E_\nu^{-2} & : E_\nu > E_{\nu,\mathrm{b2}}, \end{cases} \tag{6.26}$$

where E_ν is the energy of the neutrinos, Φ_ν is the fluence of the neutrinos, $E_{\nu,\mathrm{b1}}$ is the first break energy of the neutrinos, and $E_{\nu,\mathrm{b2}}$ is the second break energy of the neutrinos. To find the total neutrino fluence, we can stack the resolvable sources, which is known as the *quasi-diffuse* fluence. The prompt neutrino spectra for all GRBs active during ANITA-4 is shown is Figure 6.3.

Afterglow neutrinos

For afterglow modelling, the broken photon energy spectrum is:

$$\frac{dN_\gamma}{dE_\gamma} \propto \begin{cases} E_\gamma^{-\frac{3}{2}} & : E_\gamma < E_{\gamma,\mathrm{b}'} \\ E_\gamma^{-2} & : E_\gamma \geq E_{\gamma,\mathrm{b}'}, \end{cases} \tag{6.27}$$

where $E_{\gamma,\mathrm{b}'}$ is the break energy of the afterglow photons.

The corresponding neutrino spectrum is:

$$E_\nu^2 \Phi_\nu \propto \begin{cases} E_\nu & : E_\nu < E_{\nu,\mathrm{b}'} \\ E_\nu^{\frac{1}{2}} & : E_{\nu,\mathrm{max}} \geq E_\nu \geq E_{\nu,\mathrm{b}'}, \end{cases} \tag{6.28}$$

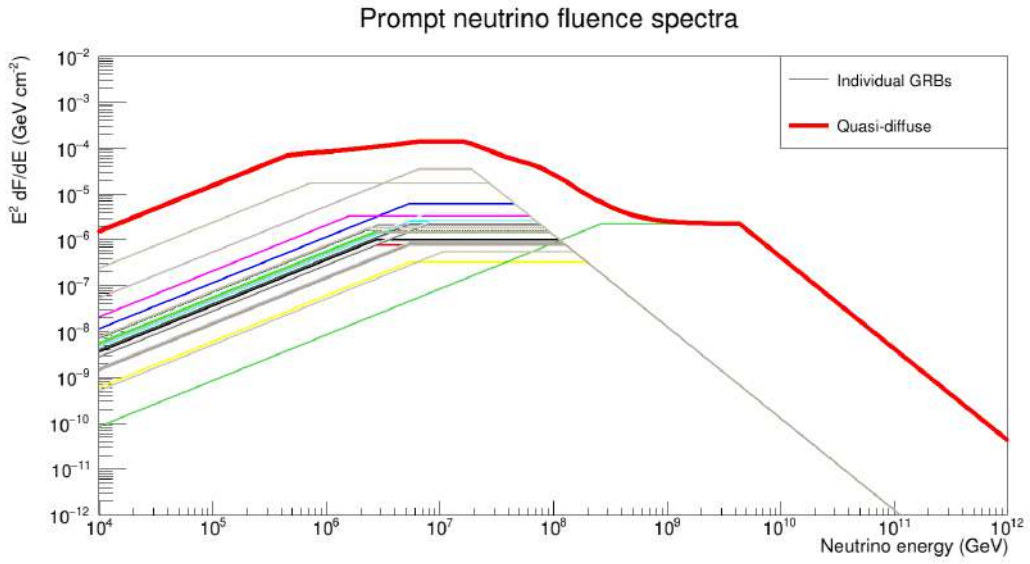

Figure 6.3: Individual simulated source fluences and the quasi-diffuse fluence of prompt neutrinos for GRBs active during ANITA-4. The thin dashed lines represent the individual fluences, whilst the thick solid line represents the quasi-diffuse fluence.

where $E_{\nu,\mathrm{b}'}$ is the break energy of the afterglow neutrinos, and $E_{\nu,\mathrm{max}}$ is the maximum energy of the afterglow neutrinos. Again, we plot the individual sources and the quasi-diffuse spectrum, as in Figure 6.4.

Now that the prompt and afterglow spectra are computed for each GRB, we can compare their quasi-diffuse fluences. Figure 6.5 shows that the quasi-diffuse afterglow neutrino fluence clearly dominates over that of prompt neutrinos. As the afterglow neutrino spectrum has a strict energy cut-off, the associated energy must be found for each GRB.

Still following the aforementioned papers [116,117], the total jet kinetic energy is assumed to be approximately $E_{\mathrm{kin}} = 10E_{\gamma,\mathrm{tot}}$, and for the external shock, the radius, r_{e} is:

$$r_{\mathrm{e}} = \left(\frac{3E_{\mathrm{kin}}}{4\pi m_{\mathrm{p}} c^2 \Gamma_{\mathrm{i}}^2} \right)^{\frac{1}{3}}, \tag{6.29}$$

where m_{p} is the mass of the proton, and Γ_{i} is the internal shock Lorentz factor. We have assumed that the number density of the interstellar medium is 1 cm^{-3} and

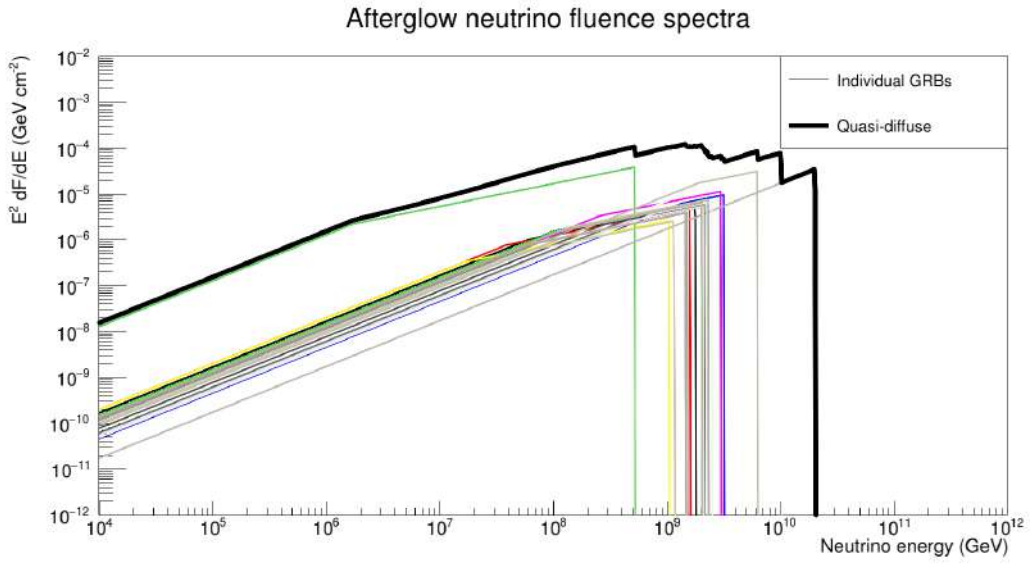

Figure 6.4: Individual simulated source fluences and the quasi-diffuse fluence of afterglow neutrinos for GRBs active during ANITA-4. The thin dashed lines represent the individual fluences, whilst the thick solid line represents the quasi-diffuse fluence.

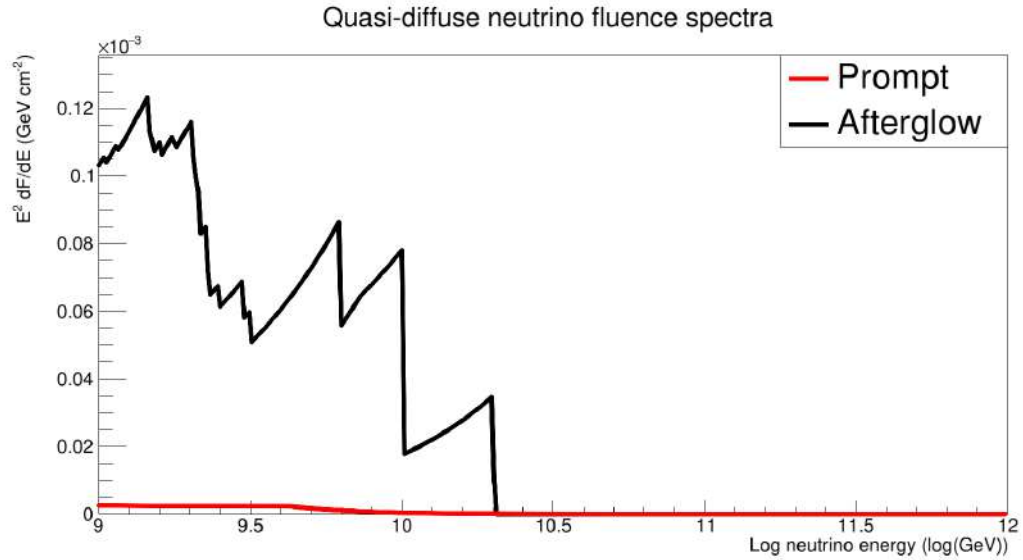

Figure 6.5: Comparison of the simulated quasi-diffuse fluences for prompt (red) and afterglow (black) neutrinos. Due to the much shallower energy scaling of the afterglow neutrino spectra at ultra-high energies, the afterglow neutrinos have a much higher total fluence within ANITA's energy range (which this graph is limited to). Note that the fluence is presented on a linear scale.

the following calculations work with the assumption that the internal shock Lorentz

factor is approximately 300. Both of these assumptions are used throughout the calculations for all GRBs [116, 117].

Now that we have calculated the radius of the external shock, the Lorentz factor of the shock's plasma shell (Γ_e) can be found:

$$\Gamma_e \simeq 195 \left(\frac{E_{kin}}{10^{54}\ \text{ergs}} \right)^{\frac{1}{8}} \left(\frac{T_{90}}{10\ \text{s}} \right)^{-\frac{3}{8}}, \tag{6.30}$$

where T_{90} is the measurement of the duration in which 90% of the GRB's energy is emitted.

The magnetic field (B_e) is calculated to be:

$$B_e = \sqrt{\frac{4\pi m_p c^2}{5}}. \tag{6.31}$$

The reverse shock-accelerated protons go on to interact with photons to form Delta baryons, which decay to form pions and nucleons (as before). The charged pion decays into charged leptons and three lepton neutrinos, and again, we assume they each receive the same energy.

Finally, the maximum shock neutrino energy, $E_{\nu,\max}$, is found to be:

$$E_{\nu,\max} = \frac{f_{p\to\pi}}{4(1+z)}(0.1\Gamma_e B_e r_e), \tag{6.32}$$

and we assume that the proton-to-pion conversion efficiency ($f_{p\to\pi}$) is 0.2.

6.1.2.3 Single Object as Neutrino Origin

Once the objects are selected from the catalogues (and, if necessary, the maximum neutrino energy is found via modelling), `icemc` uses them as origins for the ultra-high energy neutrinos.

Using a single astrophysical object as an origin of neutrinos requires a different procedure than the diffuse neutrino case. To start with, the object type(s) and source name(s) must be specified in the input file. As each object is associated with a catalogue, when an object is picked, its right ascension and declination (see Figure 6.6 for an illustration), redshift, and other properties are found and passed to the

main simulation. Its energy scheme (i.e., monoenergetic or not), as well as the minimum and maximum possible energies of the neutrinos produced by the source, are also specified. Each object has an associated time-windowed exponential flux spectrum, and its input parameters are dependent on the properties of the object. The input parameters are: the start time of the object's activity (t_0), the corresponding end time (t_1), the spectral index (γ), and the maximum neutrino energy ($E_{\max}$).[1]

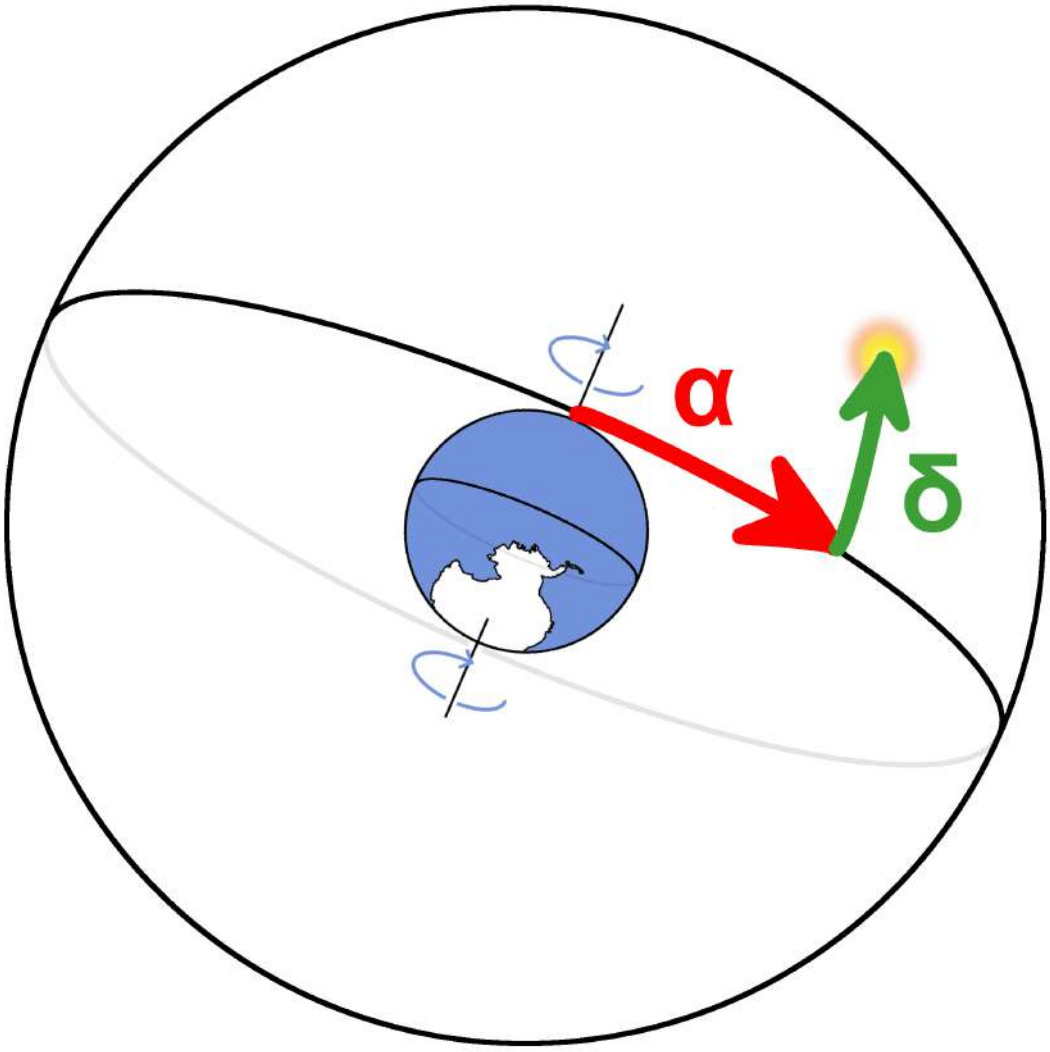

Figure 6.6: Celestial coordinates illustrated: the right ascension (α) and declination (δ) of an object is shown.

Once the source and its flux spectrum are determined, the direction and energy

[1]The function that calculates the flux spectrum also includes the normalised energy and the flux at the normalisation energy E_{n}, F_{n}. However, as the shape of the flux spectrum is used to calculate energies, we only need the parameters which form the shape of the spectrum.

of a neutrino from the source need to be found. If a monoenergetic input spectrum is chosen, the neutrino is simply fixed at that energy, and so no flux spectrum needs to be used (E is simply E_ν). However, if the energy is not confined to a single value, a flux spectrum must be used.

At a certain known flux (F_n) and energy (E_n):

$$A = F_\mathrm{n} {E_\mathrm{n}}^{\gamma}, \tag{6.33}$$

where A is the normalisation factor, and γ is the spectral index. We can use this to work out the flux spectrum for any other object of the same type:

$$F = AE^{-\gamma}, \tag{6.34}$$

where F is the flux and E is the energy. An example of a single source flux is shown in Figure 6.7.

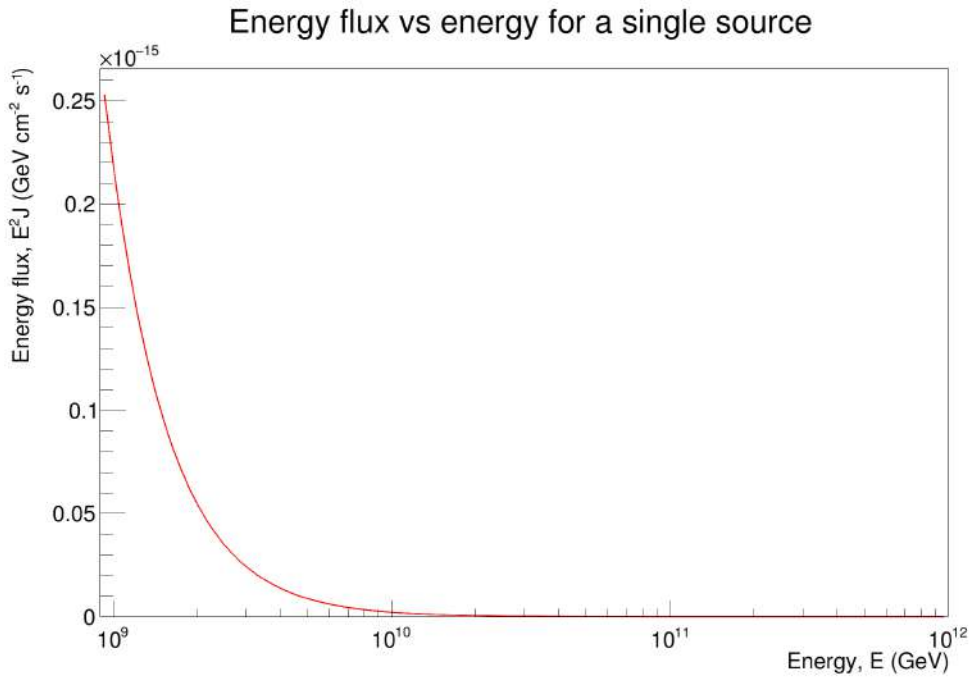

Figure 6.7: Energy flux spectrum for a single source within ANITA-4's detectable energy range. The energy of a neutrino from this source is determined by picking a random energy as distributed by this function.

The flux spectrum is used to pick the energy: any energy within the range can be chosen, but is picked according to the shape of the spectrum.

Now that the energy is determined, the neutrino's direction needs to be found. From the catalogue, we know the origin in celestial coordinates (right ascension and declination) of the object in question. As the instrument and interaction point are already set up, and the location on the flight is chosen, we know where ANITA is, as well as where the neutrino interacts. The rotation of Earth (known from the event time) and right ascension are used to find the hour angle (see Equation 6.35), which tells us where the object is with respect to the Earth's rotation.[2] Taking into account the hour angle, as well as the object's declination, the directional vector of the neutrino, **s**, can be found (as in Equation 6.36):

$$h \equiv t - \alpha, \tag{6.35}$$

$$\mathbf{s} = \cos(h)\cos(\delta)\hat{\mathbf{i}} + \sin(h)\cos(\delta)\hat{\mathbf{j}} + \sin(\delta)\hat{\mathbf{k}}, \tag{6.36}$$

where h is the hour angle, t is the local sidereal time, α is right ascension, **s** is the directional vector, and δ is declination. $\hat{\mathbf{i}}$, $\hat{\mathbf{j}}$, and $\hat{\mathbf{k}}$ are the Cartesian unit vectors.

Using the spectrum, and pointing back to space, the energy and direction of the neutrino from the source can be found, and `icemc` can run as normal.

6.1.2.4 Multiple objects as neutrino origins

As well as simulating individual objects, one might want to simulate a set of objects, such as all the active blazars during ANITA-4. This is useful for determining which sources ANITA would favour observing, should the neutrinos of such sources be detectable. ANITA is not expected to see neutrinos from high-declination sources, or those that have very low fluxes.

For the case of using multiple objects as the origin of neutrinos, the simulation has to take into account each active object's flux. Multiple objects can mean multiple object types, or those of the same type. For each source (or collection of

[2]Event time and the corresponding rotation of Earth can be expressed in local sidereal time. This gives us a measurement of Earth's rotation with respect to celestial objects.

sources) selected, each is assigned a flux and is subject to source type-specific cuts (i.e., a GRB neutrino will have a maximum possible energy). The minimum and maximum allowed energy is assumed to be the same for each neutrino (if it wasn't already set).

When selecting an object from a set to simulate a neutrino from, the total combined flux of all the objects needs to be found. To calculate this, each object is assigned a single value of flux. If the neutrino energy is fixed for an object, the flux is:

$$F = AE_\nu^{-\gamma}\,. \tag{6.37}$$

However, if there is a maximum neutrino energy cut-off, the flux for a single object calculated to be

$$F = A\left(E_{\min}{}^{\frac{1-\gamma}{\gamma-1}} - E_{\max}{}^{\frac{1-\gamma}{\gamma-1}}\right), \tag{6.38}$$

where $E_{\min}$ and $E_{\max}$ are the minimum and maximum neutrino energies, respectively.

For each source in the list of sources, its flux (those from Equations 6.37 and 6.38) is added to the total flux. Source 1 has flux F_1, and so on. Flux bands are naturally produced from this method. The first object has a band between 0 and F_1, whereas when the second object is added (and thus stacked), its band is between F_1 and $F_1 + F_2$, and so on. Thus, an object of a higher flux has a larger flux band. A random number from 0 to the total flux is generated. Whichever flux band the randomly generated flux is in, is the source that is chosen. Now that an object has been picked, we can simply find the associated neutrino energy from its own flux distribution, as we did for the single source case. Whichever source is chosen, its direction is picked the same way as before.

The sources picked in this mode have been weighted by the flux, but there are also other factors that affect whether a neutrino passes or not. At high declinations, even if the neutrino flux is particularly high, it is still very unlikely for the neutrino

to produce a detectable signal, as it is geometrically disfavoured. Thus, the multiple object scenario shows which sources are more likely to be seen by ANITA (an example of which is shown later in Figure 7.24).

6.2 CRPropa

The Cosmic Ray Propagation Framework (`CRPropa`) is a simulation used to propagate ultra-high energy nuclei, from hydrogen to iron, through space. This framework allows for the modification of the associated theoretical parameters, such as the maximum energy of cosmic rays, and the value of the source power law index.

A Python program was created using CRPropa to produce a 1D simulation for generating neutrino flux spectra at ultra-high energies.

6.2.1 Input

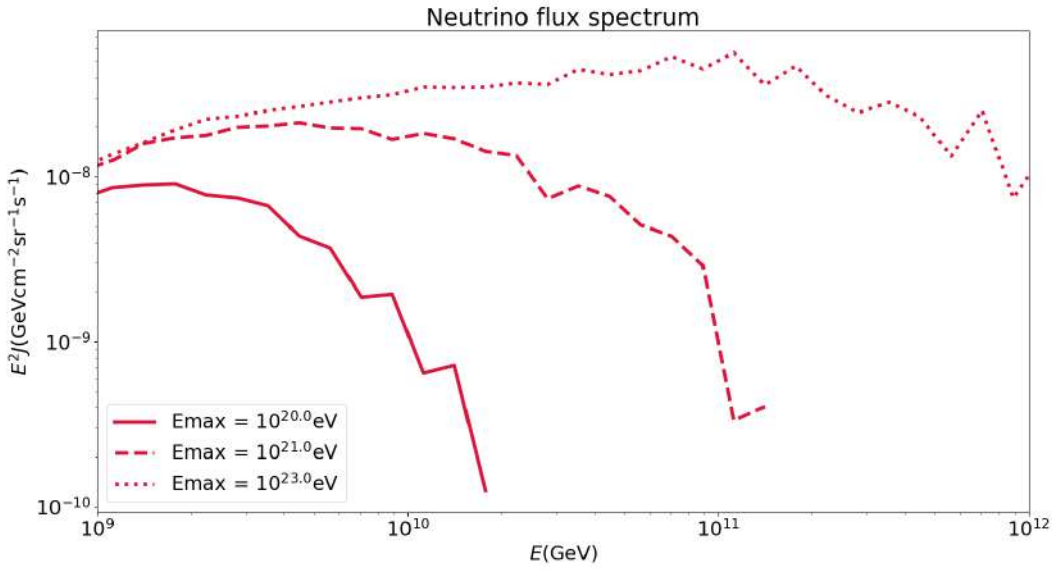

Figure 6.8: UHE neutrino spectrum for differing maximum cosmic ray injection energies.

The user can specify the following input parameters: maximum cosmic ray energy, $E_{\mathrm{p,max}}$, the source power law index, α, and the number of cosmic rays to propagate, N. The user can also specify if they wish to include an exponential cut-off factor to the injection spectrum, as well as a source evolution factor, m. The source evolution factor determines how the emissivity (the product of luminosity

and number density) of the source evolves with redshift. The source evolution scales with $(1+z)^m$. If a source evolution factor is applied, the user can also specify at what redshift the source evolution plateaus, z_S. The default parameters used are $E_{p,max} = 10^{23}$eV, $\alpha = 2$, and $N = 3 \times 10^7$. The cut-off and source-evolution are applied, $m = 3$ and $z_S = 3$. Figure 6.8 shows the effect of changing the highest attainable cosmic ray injection energy.

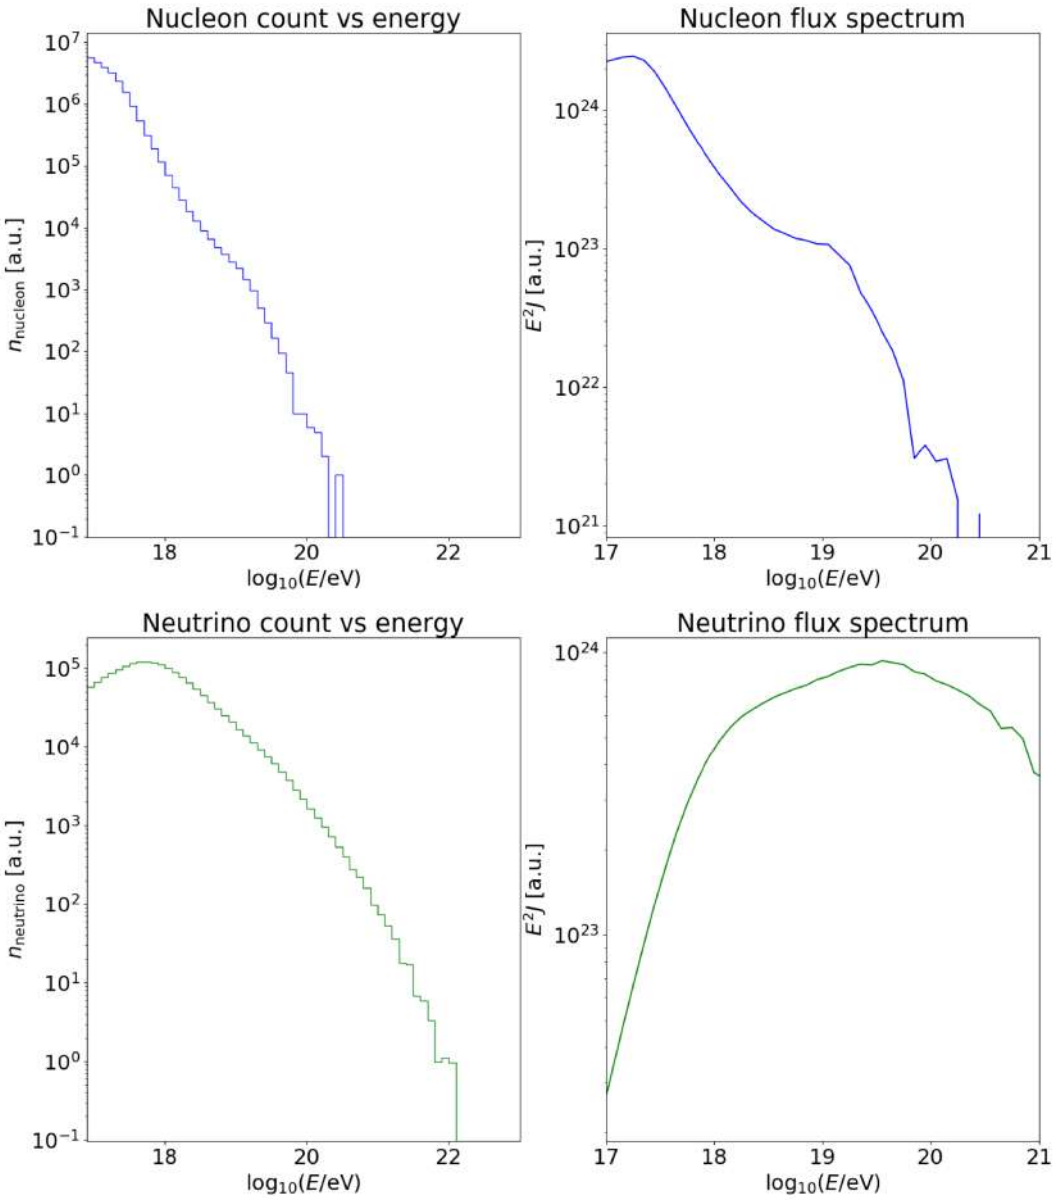

Figure 6.9: Nucleon and neutrino count vs energy (left). Nucleon and neutrino count scaled with energy squared vs energy (right).

6.2.2 Modules

`CRPropa` has a modular design: and thus has different modules for particle propagation, background and interactions, etc.

This program, which is based off of `CRPropa`, uses the `SimplePropagation` module, which is responsible for moving the particles with every step size. The minimum step size was set to 10 kpc, and the maximum was set to 10 Mpc. This module propagates the particles rectilinearly. The `Redshift` module was also used, as it updates the redshifts of the particles as they move through space.

The sources used are uniformly distributed objects with redshifts between 0 and 4, with source evolution applied. Cosmic rays are injected into these sources accordingly to a power-law spectrum. The cosmic ray composition was chosen to be fully protonic, but this is easily adjusted. An exponential cut-off was also applied to the cosmic ray flux:

$$\frac{dN_{\mathrm{p}}}{dE_{\mathrm{p}}} \propto E_{\mathrm{p}}^{-\alpha} \exp\left(\frac{-E_{\mathrm{p}}}{E_{\mathrm{p,max}}}\right), \tag{6.39}$$

where N_{p} is the number of protons, and E_{p} is the proton energy.

After the protons are emitted from the objects, they go through several interactions, and so multiple interaction modules are included. The `PhotoPionProduction` module involves photo-meson production when cosmic rays interact with the CMB. The `Nuclear Decay` module is used when the program is working with nuclei heavier than protons. The other interaction module used was `ElectronPairProduction` for photons interacting with the CMB. Taking the interactions into account, charged nuclei are subject to continuous energy loss. In applicable cases, the secondary neutrinos are also propagated, as well as the cosmic rays. Two final observers are specified: one for cosmic rays, and one for neutrinos.

6.2.3 Flux Spectrum Generation

With the above modules and assumptions, the program can run and generate flux spectra, both for neutrinos and cosmic rays. In the highest possible energy ranges,

where injected protons are, at most, 10^{23} eV, the neutrino flux will surpass that of cosmic rays. The high energy cosmic rays will decay into neutrinos, losing energy in the process. The neutrinos from the first decay of these cosmic rays will continue to move through the Universe unattenuated and are then detected. Once the particles

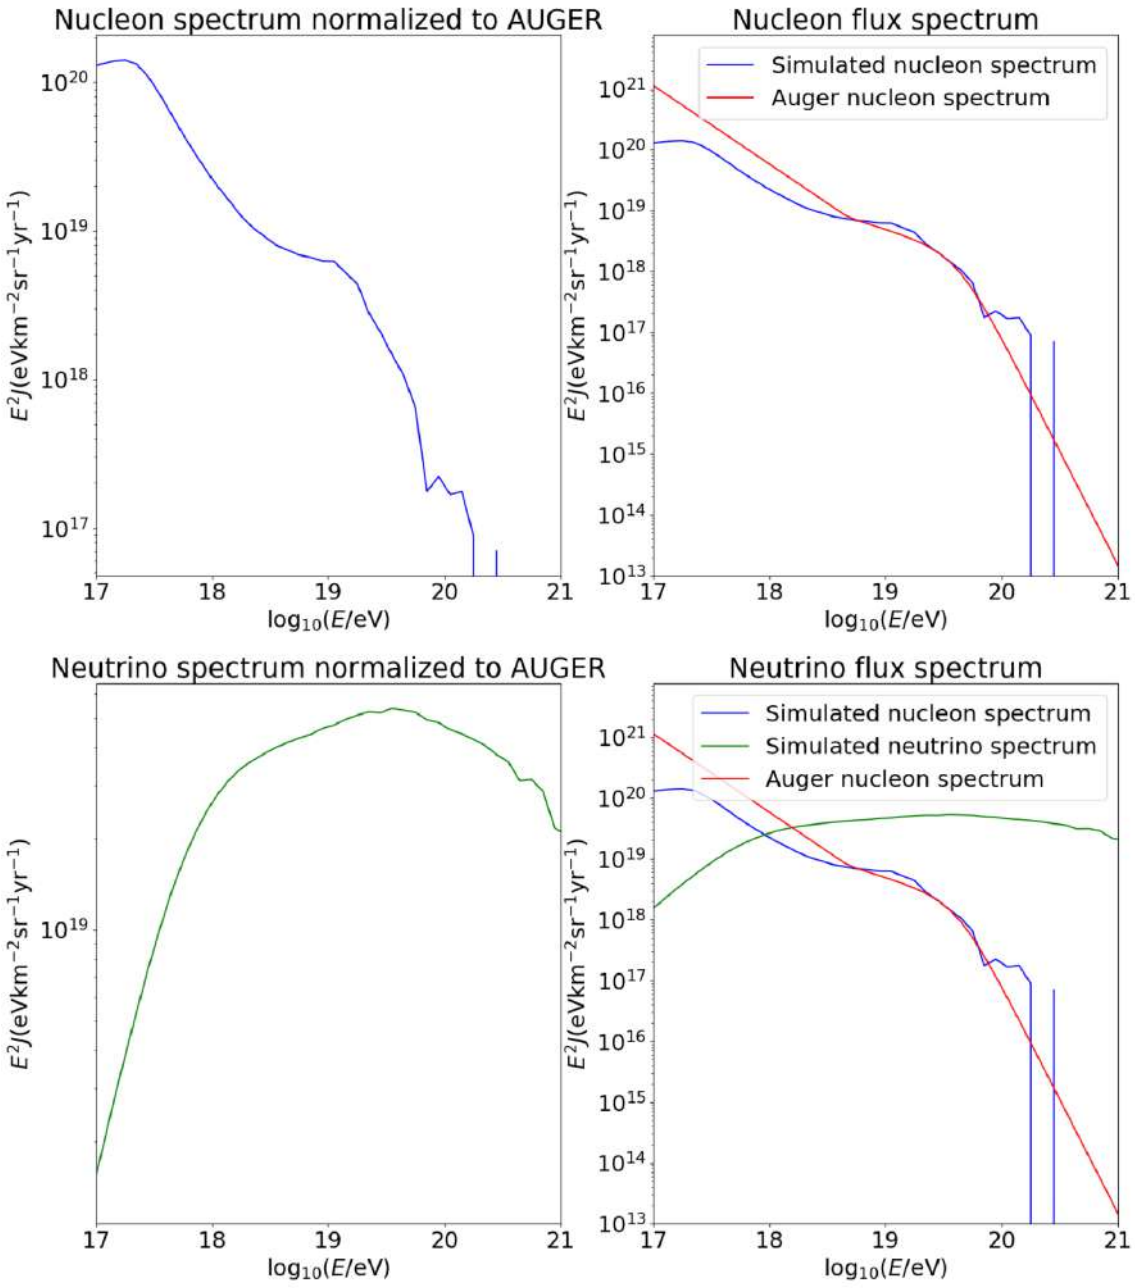

Figure 6.10: The nucleon flux is normalised to the Auger cosmic ray spectrum at $10^{19.55}$ eV. The neutrino flux is normalised with the same scaling as was done for the nucleon spectrum.

have arrived and are counted by the observer modules, the flux spectra are scaled by E^2. The output nucleon and neutrino counts and fluxes are shown in Figure 6.9.

However, the output is in arbitrary units, and we need something to normalise to. Thus, the cosmic ray spectrum is normalised to the Auger UHECR spectrum at $E = 10^{19.55}$ eV [118]. This follows a broken spectrum, as shown in Equation 6.40.

$$J(E) = \begin{cases} J_0 \left(\frac{E}{E_{\text{ankle}}}\right)^{-\gamma_1} & : E < E_{\text{ankle}} \\ J_0 \left(\frac{E}{E_{\text{ankle}}}\right)^{-\gamma_2} \left[1 + \left(\frac{E_{\text{ankle}}}{E_{\text{S}}}\right)^{\Delta\gamma}\right] \left[1 + \left(\frac{E}{E_{\text{S}}}\right)^{\Delta\gamma}\right]^{-1} & : E > E_{\text{ankle}}, \end{cases} \tag{6.40}$$

where $J_0 = 3.3 \times 10^{-19}$ eV^{-1} km^{-2} sr^{-1} yr^{-1}, $E_{\text{ankle}} = 4.82 \times 10^{18}$ eV, $\gamma_1 = 3.29$, $E_{\text{S}} = 4.21 \times 10^{19}$ eV, $\gamma_2 = 2.60$, and $\Delta\gamma = 3.1$.

The nucleon spectrum is normalised to the Auger spectrum (and the neutrino spectrum is scaled in the same manner). This is shown in Figure 6.10.

The neutrino flux spectrum can be converted into the format needed by `icemc`, and thus can be used as a flux input. Overall, the use of `CRPropa` within `icemc` allows one to generate one's own neutrino flux spectrum and adjust the parameters as needed. Taking into account the previous assumptions and using the default value of the variables, this simulation method gives the following neutrino flux spectrum, which is shown alongside theoretical fluxes and experimental limits in Figure 6.11.

The resulting spectrum from this simulation is similar to other predicted flux spectra, especially those that use a proton-only (or proton-dominant) model for the initial cosmic rays. At ANITA's lowest detectable neutrino energy, 10^{18} eV, the simulated flux is very close to the pure proton-induced neutrino flux predicted by Kotera, *et al* [122] (upper bound of grey area in Figure 6.11). However, as the energy increases towards ANITA's maximum detectable neutrino energy, the simulated flux deviates from this bound. This is mainly because the simulation works with a higher maximum proton acceleration energy, and has a smaller value for the power law index. As more protons are produced at higher energies (due to both the increase in maximum proton energy and the shallower power law), a higher flux of

higher energy neutrinos is to be expected.

Overall, the simulations explained in this chapter allow one to generate neutrino flux spectra and the events associated with them. The usage of simulated neutrino events from specific directions is a crucial part of the source search analysis, and is explained in the next chapter.

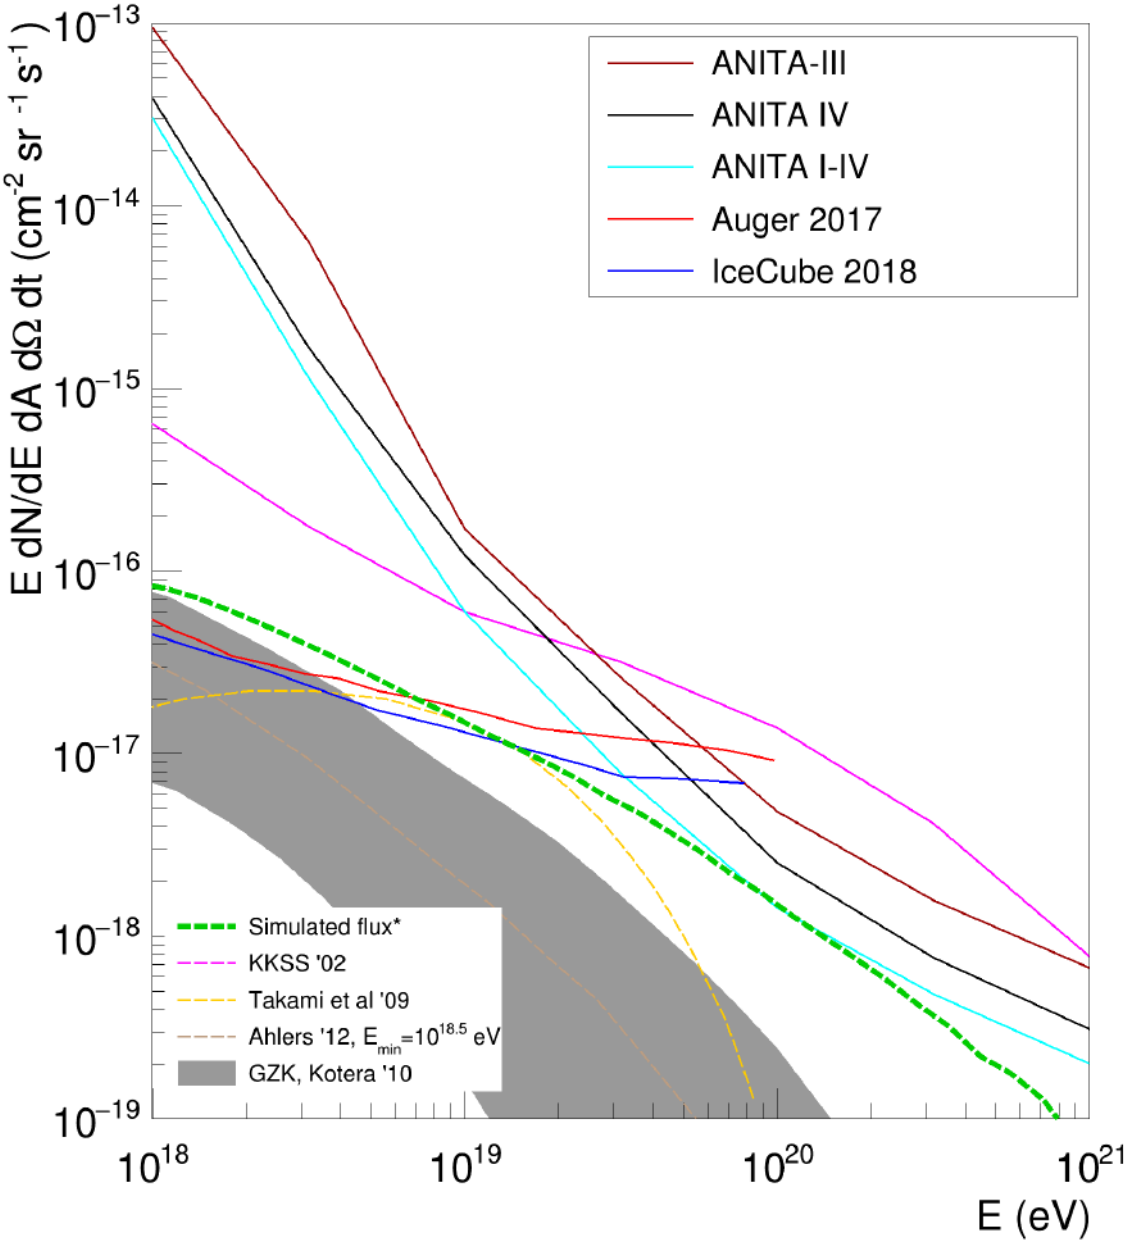

Figure 6.11: Predicted fluxes (dashed) and experimental limits (bold) for UHE neutrinos. The dashed green line (*) corresponds to the flux simulated with the custom CRPropa-integrated program using the parameters and assumptions as described above. Alongside the simulated flux are a set of cosmogenic neutrino models [119–122], limits set by the ANITA experiment [77,79] and other experimental limits [123, 124]. The simulated neutrino flux spectrum is similar in shape to the KKSS flux spectrum, but uses a more realistic source power law index, and thus, its flux is lower. The simulated flux is more favourable than the maximum bound set by Kotera, *et al*, and the limits set by Ahlers, *et al*. As such, the simulated spectrum predicts a high flux of ultra-high energy neutrinos that approaches the limits set by ANITA.

Chapter 7

Analysis

The analysis contained within this thesis is a *source search*—it aims to determine the source origin of ultra-high energy neutrinos. As such, only the initial steps of an ANITA-4 singlet search are repeated [79]. For those interested in the neutrino singlet search, please consult Andrew Ludwig's thesis [125].

The aim of an ANITA neutrino analysis is to find radio signals of neutrino origin. However, two classes of background events are present within the data collected: anthropogenic noise and thermal noise. The class of anthropogenic noise is associated with radio signals produced by human activity, such as radio communications from research stations, but also other devices used across the continent. Thermal noise arises from the continuous motion of free electrons in the Antarctic ice sheet, as well as the amplifiers that are part of the ANITA instrument. Thus, it is critical to separate the signal class from the two background classes. Radio signals of anthropogenic origin can be removed from the sample by finding the positions they exited the ice from. Anthropogenic signals form clear clusters to research stations and areas of high activity, and thus these undesirable events can be removed via clustering methods. Thermal signals can be separated from potential neutrino signals by examining the waveforms of known thermal events and simulated neutrino events. Unlike neutrino-induced radio signals, thermal signals will not be as coherent between antennas. Thus, we can use the qualities of the waveforms from multiple antennas to cut away the thermal noise events.

In order to carry out a source search, however, the analysis uses multiple can-

didates to project back to their potential celestial origins. Thus, the cut used to discriminate between thermal noise and possible neutrino candidates is made a lot looser than a singlet search, allowing "sub-threshold" events from the singlet search to be analysed. In addition, an independent method of loose clustering is used to remove candidates that came from Antarctic bases. The post-clustered data forms the basis of the final event set for this analysis.

These sub-threshold events are pointed-back and transformed to celestial coordinates, under the hypothesis of neutrino origin. Spatial and temporal profiles of known astrophysical objects are found using simulation (see Section 6). These profiles are compared to the time and celestial position of the pointed-back events. In order to assess the background, several alternative universes are simulated whereby the spatiotemporal properties of the objects within are shifted. If the events are significantly higher in proximity to the actual spatiotemporal profile of an object in our Universe, it suggests an excess of ultra-high energy neutrinos come from that source.

7.1 Pre-Source Search

Before the source search can commence, the first steps of the ANITA singet analysis are carried out. This involves performing quality cuts on the ANITA-4 dataset. Quality cuts are necessary to remove events that do not correspond to potential science signals, including payload blasts and digitiser glitches. Only "real" events are used, which means that the recorded radio signals are not forced by a minimum bias trigger (see Section 5.1.7.3) or reconstructed in a nonmeaningful way (such that the elevation angle of the event was outside the range of 60°to -50°).

7.1.1 Quality Cuts

Payloads blasts are a class of unwanted near-field events that show up in both vertical and horizontal polarisations. One class of these blasts can be defined by its fluctuating waveforms in most channels, whereby the absolute value of the voltage rises to a maximum, and then slowly fades out (see Figure 7.1 for an example payload blast). Payload blasts appear most dominant in the bottom and middle an-

tennas. For this reason, if the peak-to-peak voltage ratio between a top antenna and a bottom (or middle) antenna in the same phi-sector is too low, the event is rejected. The second class of these blasts are opposite of the first class of blasts, i.e., they are extremely dominant in the top antennas. These events are rejected, given that the peak-to-peak ratio exceeds a certain point.

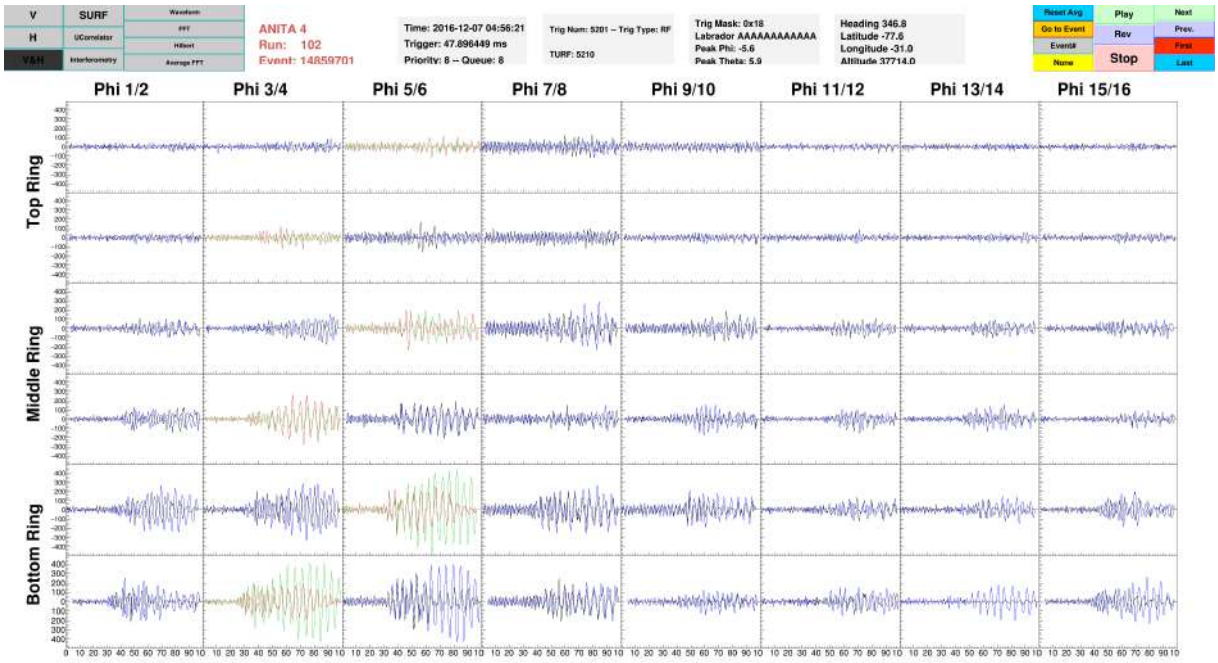

Figure 7.1: An example of a payload blast event. This event is classified as a payload blast due to its large difference in peak-to-peak voltages between antennas in the bottom ring to the top ring.

As payload blasts appear to come from the payload (near-field), the resulting waveforms should not add coherently under the far-field plane wave approximation. We can use this to further reject payload blasts. The coherently-summed waveform of an event is formed from several waveforms. For incoherent signals, the average value of the peak voltage of the Hilbert envelopes[1] from these individual waveforms compared with the average value of the peak voltage of the coherently-summed waveform should be very different. If the difference is too large, the events are rejected.

Digitisers are used in the ANITA experiment to convert triggered waveforms

[1]The Hilbert transform can be used to acquire the analytic signal of a waveform. The Hilbert envelope is the magnitude of this analytic signal.

Status	Events remaining
All data	96307012
Quality cuts	
Non-real event	89822992
Digitizer glitch	89767340
Payload blast	71975141
Thermal cuts	
Fisher discriminant analysis	887793
Clustering cuts	
OPTICS algorithm	917

Table 7.1: Status of the data at several points during the analysis, including quality cuts, thermal cuts, and clustering cuts. The remaining events are shown in the order the cuts were applied, where the previous cuts have already been applied. For detailed information about each specific stage and the cuts involved, see Sections 7.1.1, 7.1.2, and 7.2.1.

into digitally recorded signals. However, the process of digitisation occasionally "glitches", i.e., unusually large voltages arise. Digitiser glitches occur in two modes: single sample and step function glitches. A single sample glitch can be identified as having a huge voltage spike in only a single sample: the other samples remain the same. A step function glitch can be identified as follows: after several voltage-consistent samples, each sample following a certain point gets displaced by a large voltage. The voltage is now consistent, but largely displaced from the previous consistent voltages. After several samples, the voltages return to the originally recorded values. As such, this glitch looks like a step function in the waveform data. To account for these glitches, voltage asymmetries up to 500 mV are flagged and excluded from the final analysis, as are voltages that exceed a single volt. 71,975,141 events remained after all quality cuts. See Table 7.1 for the events remaining after specific cuts.

7.1.2 Thermal Cuts

Up to this point, the ANITA-4 singlet analysis has been replicated. However, this source search and the aforementioned analysis have to differ in their handling of the potential events. A singlet neutrino search specifically looks for only the most isolated events that have very strong neutrino signatures. In order to achieve this,

almost all of the sub-threshold events (those just below the strict neutrino-like event criteria) are cut away. This results in a very small number of events after clustering, typically ∼1,[2] potential final events.

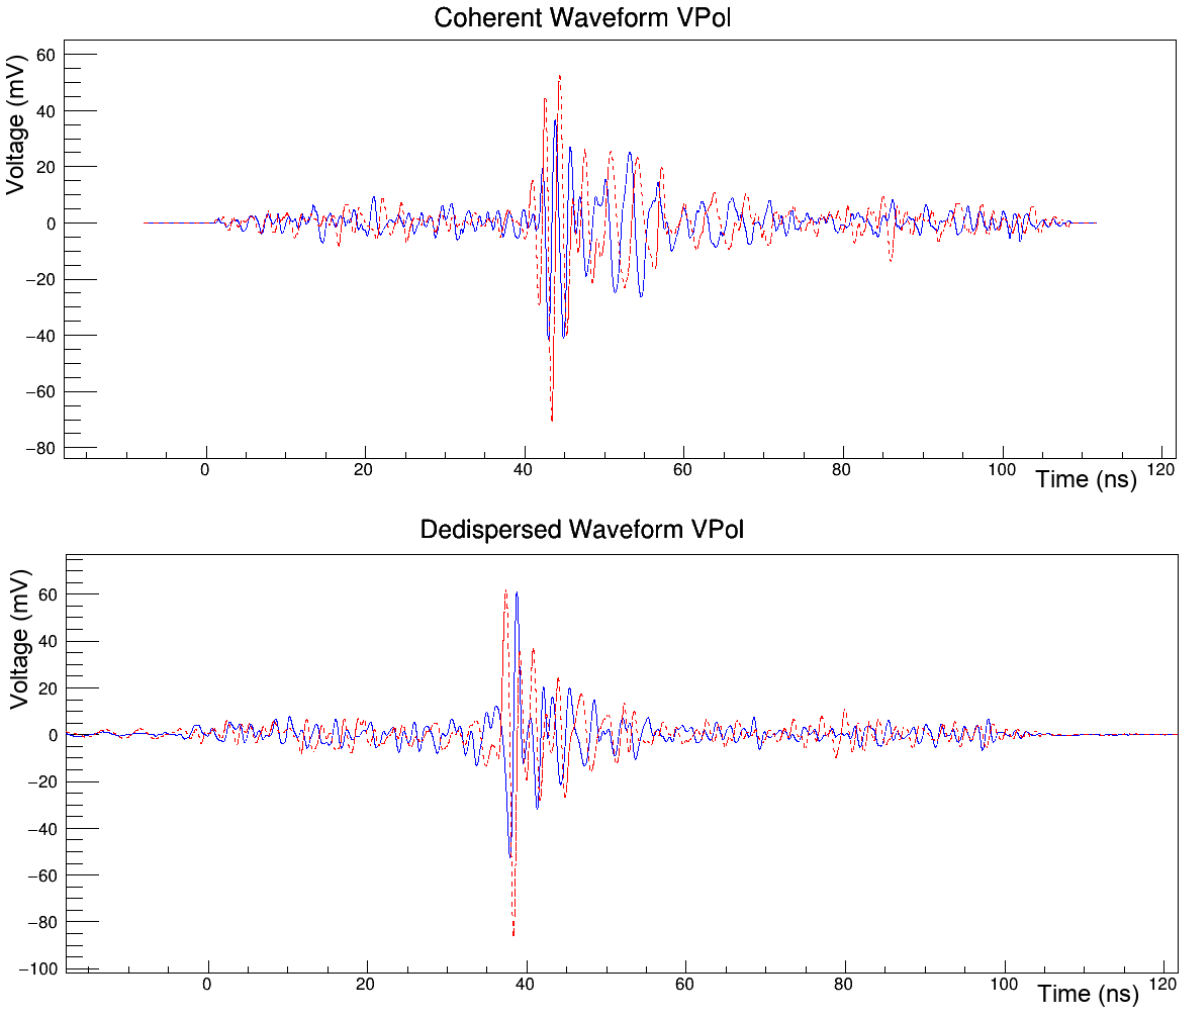

Figure 7.2: A West Antarctic Ice Sheet (WAIS) calibration pulse for both the coherently-summed and de-dispersed waveform. The graphs display the vertically-polarised waveforms (blue) in the foreground and the corresponding cross-polarisation (dashed red) in the background.

Waveforms of radio signals associated with thermal noise and those from neutrino signals are intrinsically different. Instead of just using a single waveform, we can combine information from multiple antennas to find properties useful in distinguishing neutrino-induced radio signals from thermal noise. A coherently-summed waveform is constructed by combining information from antennas belonging to the

[2] ANITA-4 observed a single event on a background estimate of $0.64^{+0.69}_{-0.45}$

five phi-sectors closest in direction to the radio signal. Each waveform is delayed with respect to its position and summed, forming a waveform with a much higher signal-to-noise ratio. In doing so, the signals coherently add to each other, whilst the background thermal noise does not add coherently and essentially cancels out, thus increasing the signal-to-noise ratio. However, the coherently-summed waveform is not a perfect reconstruction of the desired radio signal, as parts of the ANITA instrument itself modify the signal. Thus, de-dispersion is applied to the waveform in order to correct it: this consists of removing the phase (and amplitude) response from the system's output (impulse response). This is done by dividing out the phase distortion only, as the amplitude response is very consistent across ANITA's frequency band. Figure 7.2 shows a coherently-summed waveform and de-dispersed waveform for a radio signal sent from a calibration antenna from the West Antarctic Ice Sheet (WAIS). These "WAIS calibration pulses" are sent from a radio emitter in West Antarctica by the ANITA Collaboration in order to see what a known impulsive signal should look like once processed by the ANITA instrument.

Neutrino-induced radio signals come from a coherent process, and therefore are expected to be highly impulsive. To get an idea about the impulsivity of an event, two variables are introduced: *impulsivity* and *power window gradient*. Impulsivity is a quantitative measure that uses the peak of the Hilbert envelope to assess how the power is distributed along the waveform. Waveforms that have a large proportion of their power close to the peak of the Hilbert envelope have a high impulsivity. Those which have more uniformly distributed power will have a very low impulsivity. In order to focus on the most impulsive parts of the waveform, we can restrict the waveform to a "window" around its peak. To find the power window gradient, we restrict the waveform to 10% of its power and calculate how large the time window is. We then carry out the same method in steps of 10% of the power until 50%, and calculate the time window each time. The power window gradient is the average difference between these time windows.

In order to distinguish thermal noise from neutrino events, ROOT's Toolkit for Multivariate Analysis (TMVA) is used. Specifically, this analysis uses Fisher

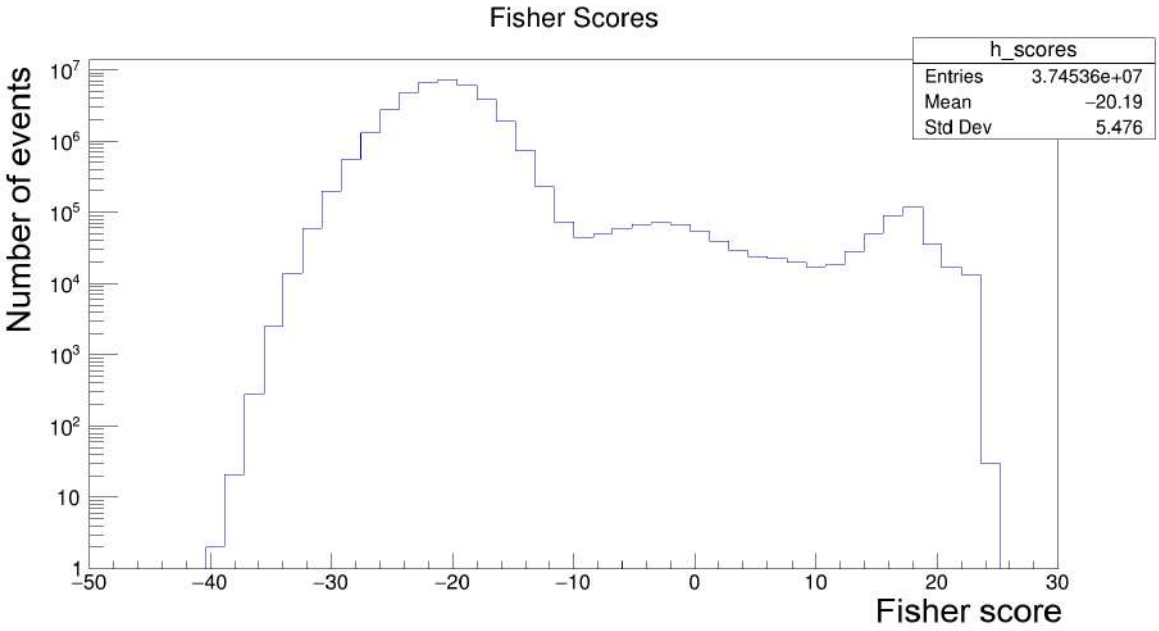

Figure 7.3: Fisher discriminant histogram for all vertically-polarised events.

Discriminant Analysis (FDA). Such a method projects a vector onto a hyperplane that splits the hyperplane in two. As such, this linear classification vector allows one to separate data into two classes, in this case, neutrino-like and thermal noise-like. Nine total input variables are used in combination to discriminate the two event types. These variables are the impulsivity (of both the coherent and de-dispersed waveforms, and the difference between them), linear polarisation fraction (of both the coherent and de-dispersed waveforms, both windowed and non-windowed), the difference between the Hilbert peaks of the coherent and de-dispersed waveforms,[3] and the power window gradient of the de-dispersed waveform.

The training data used for the thermal noise dataset are events that reconstructed above the payload (as no neutrino signals would come from above) and passed quality cuts (so as not to include glitch events). Only vertically polarised (VPol) events are used for both the training and testing data, as the analysis operates under the assumption that the candidates are potential neutrinos. The training

[3]The Hilbert peak difference (of the coherent and de-dispersed waveforms) of *non-coherent* waveforms will be centered around zero: applying de-dispersion to the coherently-summed waveform should not change the peak in this case. However, for impulsive coherent signals, we expect there to be a difference, because the de-dispersion reduces the group delays when dividing out the phase distortion. Therefore, the Hilbert peak difference is a decent discrimination variable.

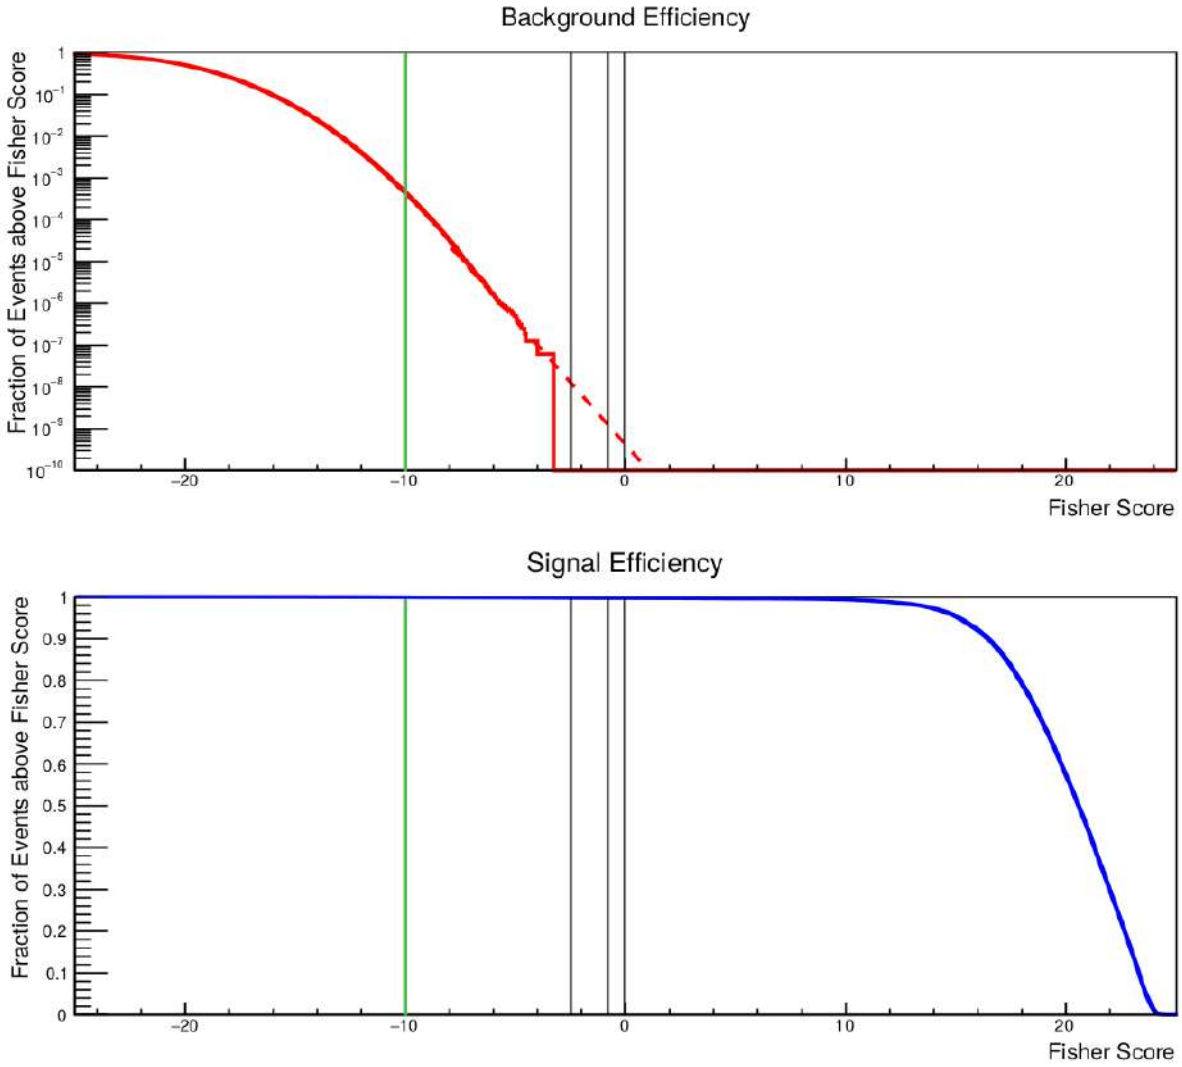

Figure 7.4: Fisher discriminant background and signal efficiency plots: black vertical lines represent cuts considered for the singlet search and green vertical lines represent the cut for the search in this analysis. Note that the cut for this analysis is purposely set to be a lot looser. Due to the nature of the sub-threshold analysis, we accept far more background events, and also capture more signal events. Plot modified from [125].

data used for the neutrino dataset come from the `icemc` simulation program. After the algorithm is trained, the discrimination power between thermal noise-like and neutrino-like events is determined. The strongest discrimination variable is the impulsivity of the dedispersed waveform. When the trained FDA is applied to the test data (post-quality cut ANITA-4 data), the Fisher scores (for VPol events only) form a distribution as shown in Figure 7.3. The efficiency plots are shown in Figure 7.4.

The Fisher score cut was originally set to -1 for the singlet analysis. However,

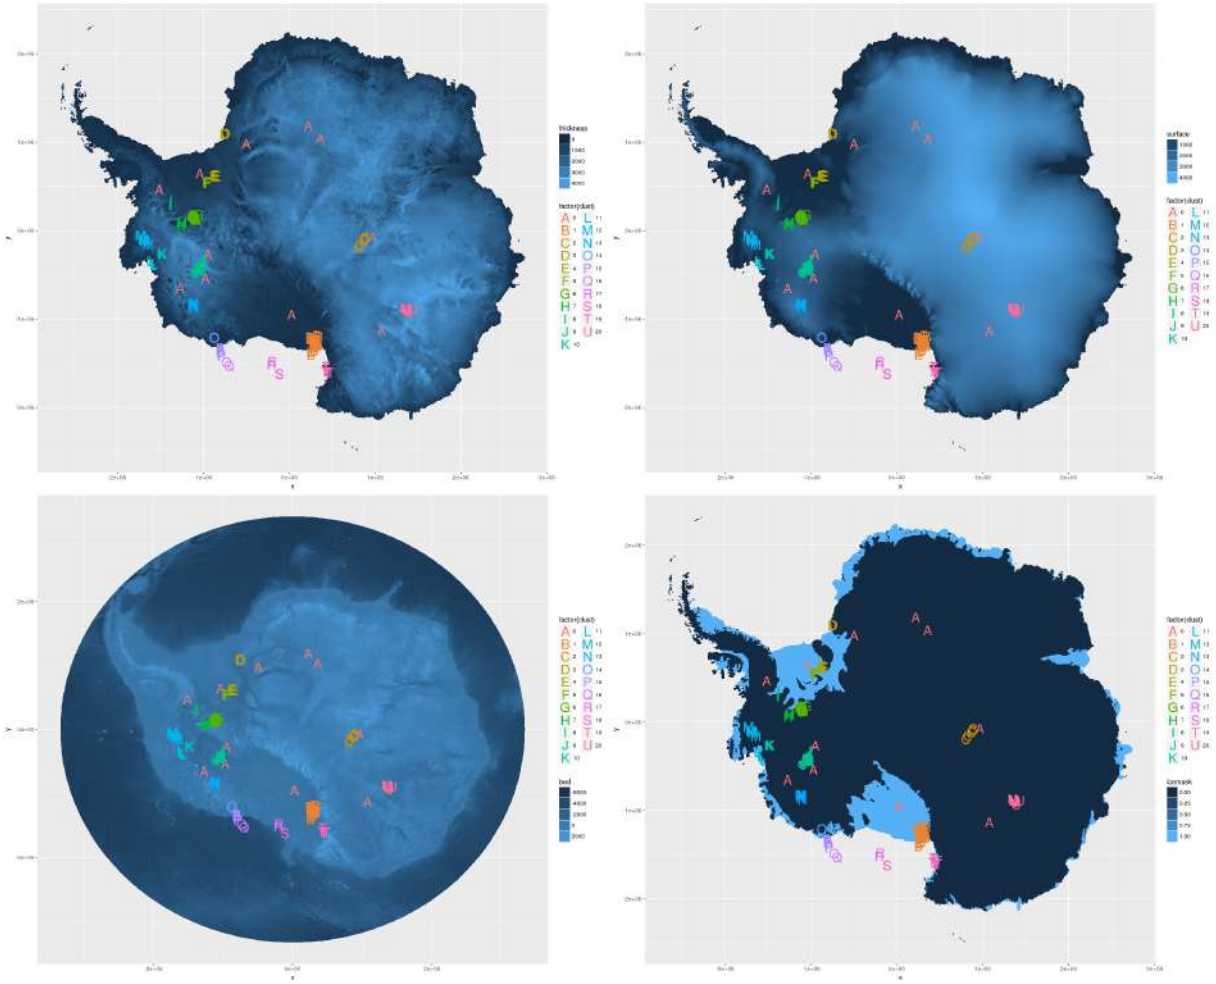

Figure 7.5: Multiple views and the clustering output for a subset of ANITA-4 data using antarcticR.

as we wish to look for more events in the source search, the cut is set to -10. This cut value is chosen because both calibration pulses and simulated neutrinos had events (though very few) in the bins corresponding to a Fisher score of -10, whilst still maintaining fair background leakage. In the source search, more sub-threshold events are allowed to pass, which allows the analysis to point a larger number of events back to their potential astrophysical origins. After thermal cuts, 887,793 events remain. A summary of cuts applied thus far is shown in Table 7.1.

7.2 antarcticR

`antarcticR` is a custom-made R [126] package designed to incorporate clustering and visualisation of events on the Antarctic continent.

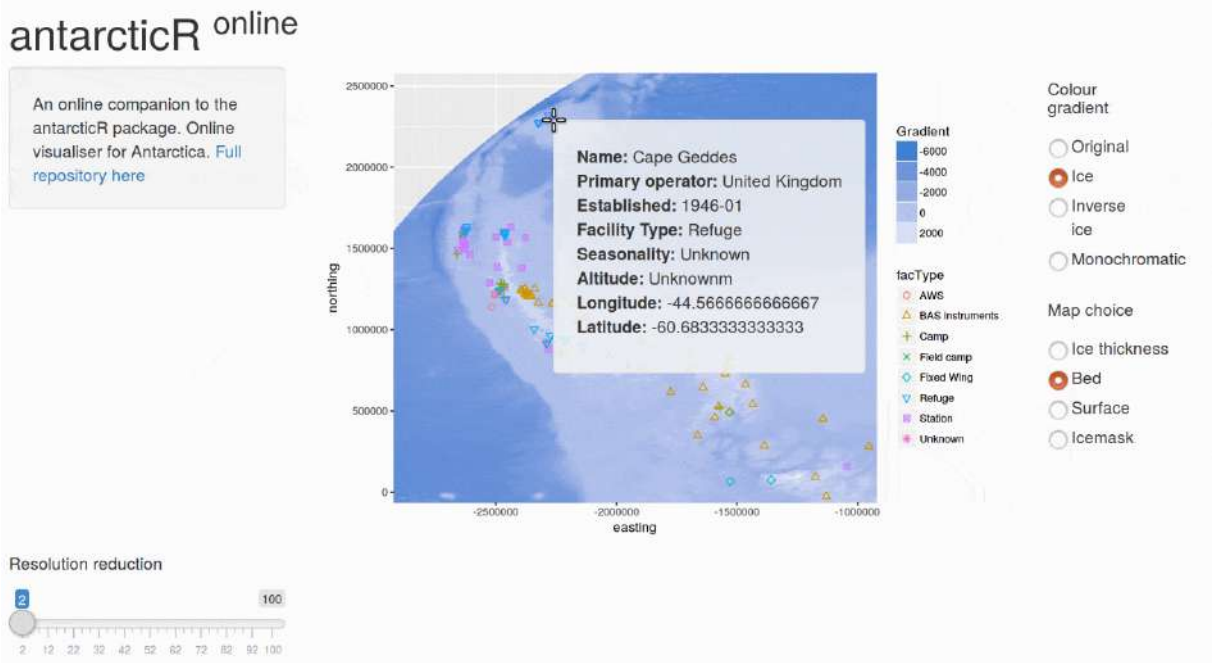

Figure 7.6: The interactive online companion of antarcticR. This example shows a high-resolution bedrock map of the Antarctic peninsula with a specific refuge selected.

This package uses the BEDMAP2 data to map the ice thickness, as well as the bedrock height and other ice qualities, of Antarctica (see Figure 7.5 for an example of event clusters on BEDMAP2 data). The locations of both radio events and scientific bases can be visualised within `antarcticR`. It also contains the most up-to-date Antarctic base lists, including the bases active during both the ANITA-3 and ANITA-4 flight seasons.

`antarcticR` includes multiple functions, such as generating Haversine (shortest distance two points on a sphere) matrices or Cartesian (Euclidean distance) matrices for neutrino events, which are used for mapping data on the Antarctic continent. These distance matrices contain the distances from one event location to another, which is crucial for spatial clustering. The mapping features, such as viewing the BEDMAP data and base lists are available as an interactive standalone or online tool (see Figure 7.6 for an example of the interface), both built using `Shiny` [127]. One can use this to easily find out the seasonality of a desired base, as well as its latitude, longitude, and altitude, and other details.

7.2.1 Clustering

ANITA observes a large amount of anthropogenic noise from research bases, and as such, a significant fraction of recorded events point back to these bases. Events that point to a single extremely active base on the surface will also be reflected in the sky map. However, both the payload and the Earth itself are rotating, meaning the features of the pre-clustered events on the sky map will not appear as a single dense region, but rather dense curved lines. Events are clustered on the surface only in order to attempt to remove events that cluster to these active bases.

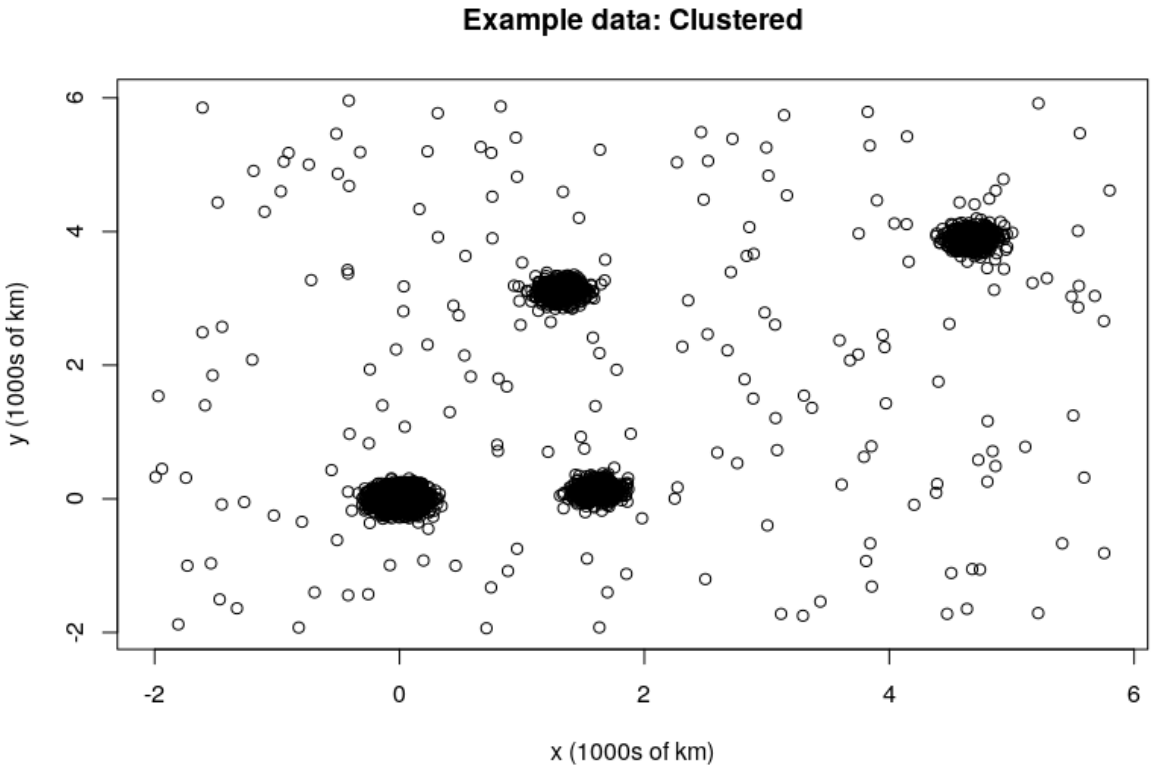

Figure 7.7: An example set of data (unclustered).

`antarcticR` makes use of the Ordering Points To Identify the Clustering Structure (OPTICS) clustering algorithm (within the `dbscan` package [129]) to separate events that are very close to each other on the continent. OPTICS is a density-based clustering algorithm that can identify event clusters of differing densities. This is useful for ruling out both active and semi-active bases in Antarctica, which will differ in spatial density of events.

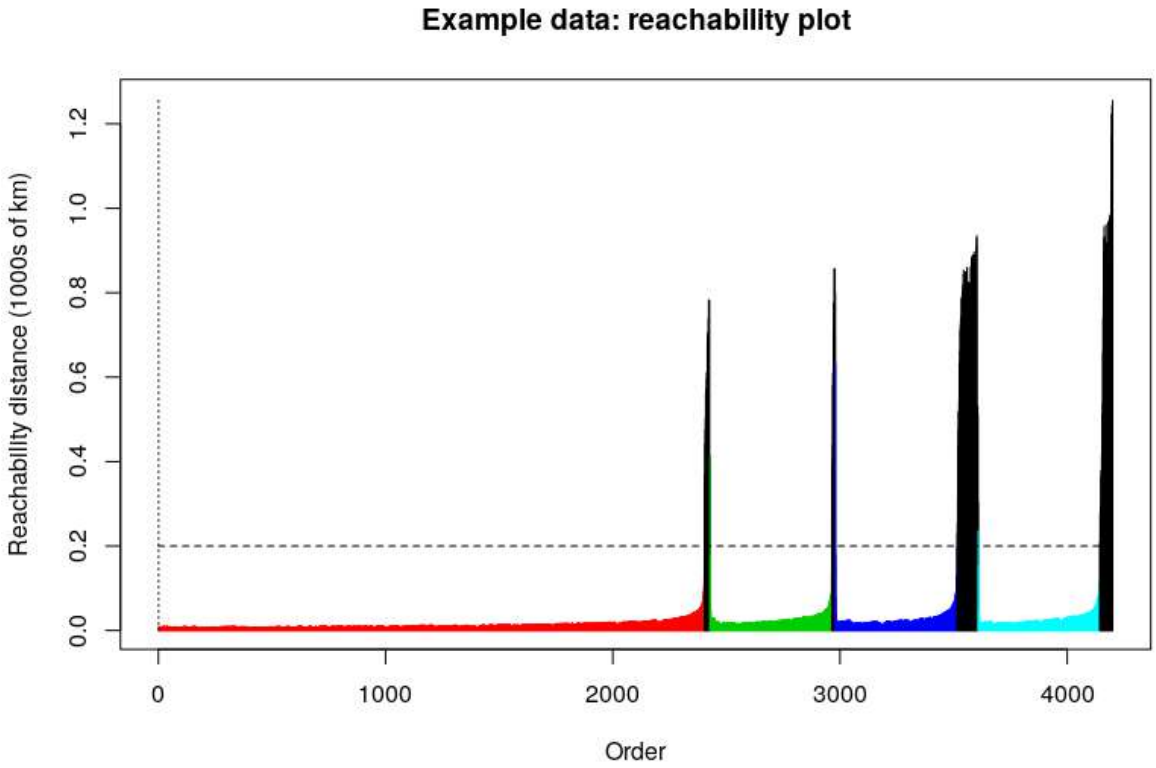

Figure 7.8: The reachability plot (generated by the OPTICS algorithm) for our example set of data. The coloured "valleys" represent the closest set of points in the whole dataset in terms of reachability distance, thus these are clustered events. The red valley is significantly more prominent that the green, blue, and cyan valleys, as it contains more events. The black spikes represent the points furthest away from the next closest points, thus they are isolated events. The reachability distance cut, ϵ_{cl}, is set to 200 km, as this selection contains the valleys, i.e., the clusters. Setting a much larger cut would cause the algorithm to group too many outliers into a cluster, and setting it too low would force it to ignore points clearly in a cluster.

In order for OPTICS to order the points, `antarcticR` extracts the events' longitude-latitude data from `ROOT` and constructs a Haversine matrix. This distance matrix contains the great-circle (Haversine) distances between each of the events. OPTICS uses this to linearly order the events by their distance to each other. Prior to running the clustering algorithm, the user must input two variables, ϵ[4] and N. ϵ is a distance cut-off parameter for core cluster formation. N is the minimum number

[4]Setting a value for ϵ, the maximum core cluster-forming distance, isn't strictly required. In this case, the algorithm just uses the maximum possible distance, which means the program would take an incredibly long time to run. ϵ was generically set to a large value, which was chosen to be 100 km for the real world data.

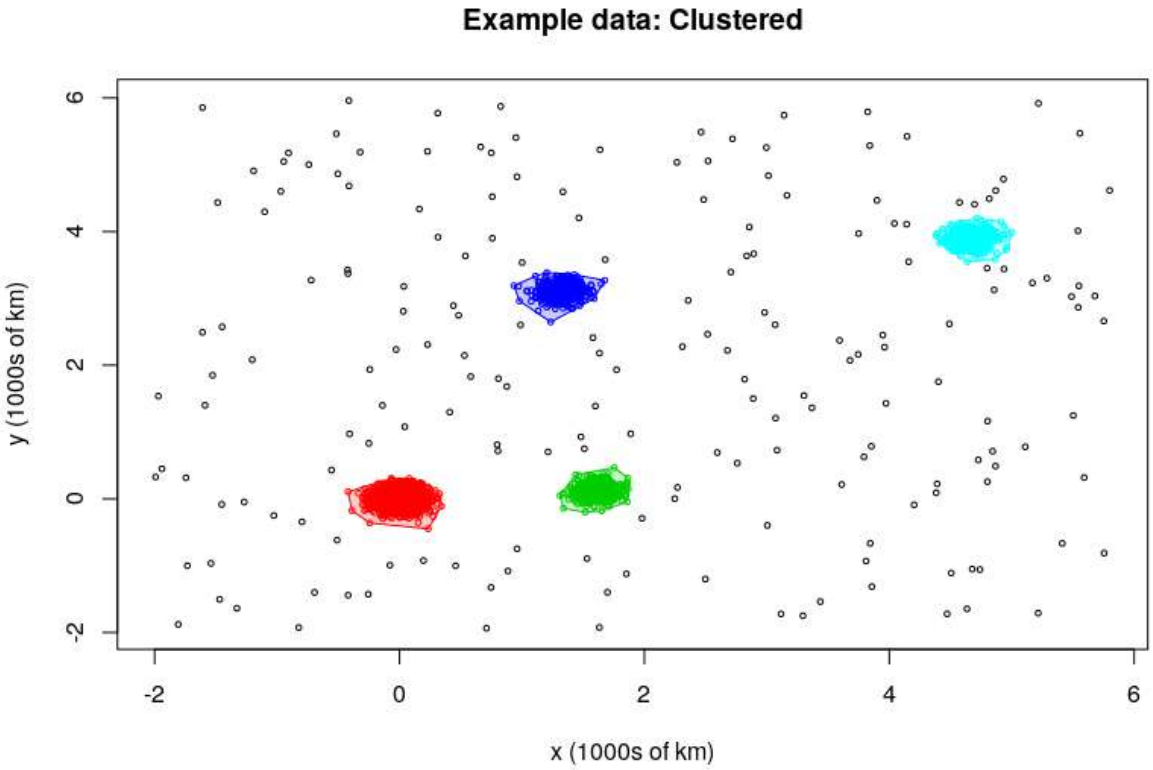

Figure 7.9: An example set of data (clustered). By setting a reachability distance cut, we are able to identify clusters (coloured) and separate them from the desired events, which are isolated (black).

of points required to form a cluster. Thus, any event that has $N-1$ other events within a radius of ϵ is a core event of a cluster. These core events each have a core distance, which is the minimum ϵ required for the event to be classified as a core event. The reachability distance from one event (E_0) to another (E_1) is either the core distance of E_0 or the distance between the points, whichever is larger.

When the clustering algorithm is applied, a reachability plot is constructed. A reachability plot visualises the output of the OPTICS algorithm. Each bin on the x-axis corresponds to one point: it is ordered such that that the spatially closest events are next to each other. The reachability distance is plotted on the y-axis. Events in an event cluster have low reachability distances to their nearest neighbours, and thus events in clusters will form valleys (consistently low reachability distances). The clusters themselves can be extracted by making a reachability distance cut, ϵ_{cl}.

Before applying the OPTICS algorithm to our post-thermal dataset, a simple

example is provided to explain how it works. To start, we generate 4 distributions (similar to the "research bases" ANITA observes), one large and three medium-sized, over a background of events (similar to the isolated sub-threshold candidates ANITA observes), as shown in Figure 7.7. A reachability plot is constructed, as in Figure 7.8, and a cut is set on the reachability distance. This separates the clustered events from the isolated events; the clustered data is then displayed in Figure 7.9.

The algorithm is applied to our post-thermal dataset. For this analysis, ϵ was set to 100 km and N was set to 10. Setting N much lower than 10 would cause a very high number of events to cluster, leaving us with a very small amount of neutrinos to point back to potential astrophysical objects. Setting it higher than 10 would cause less-active bases to be missed and so less clusters would form in general. The reachability plot is constructed (see Figure 7.10), and our ϵ_{cl} is set to 10 km. The clustered events are then simply removed from our dataset.

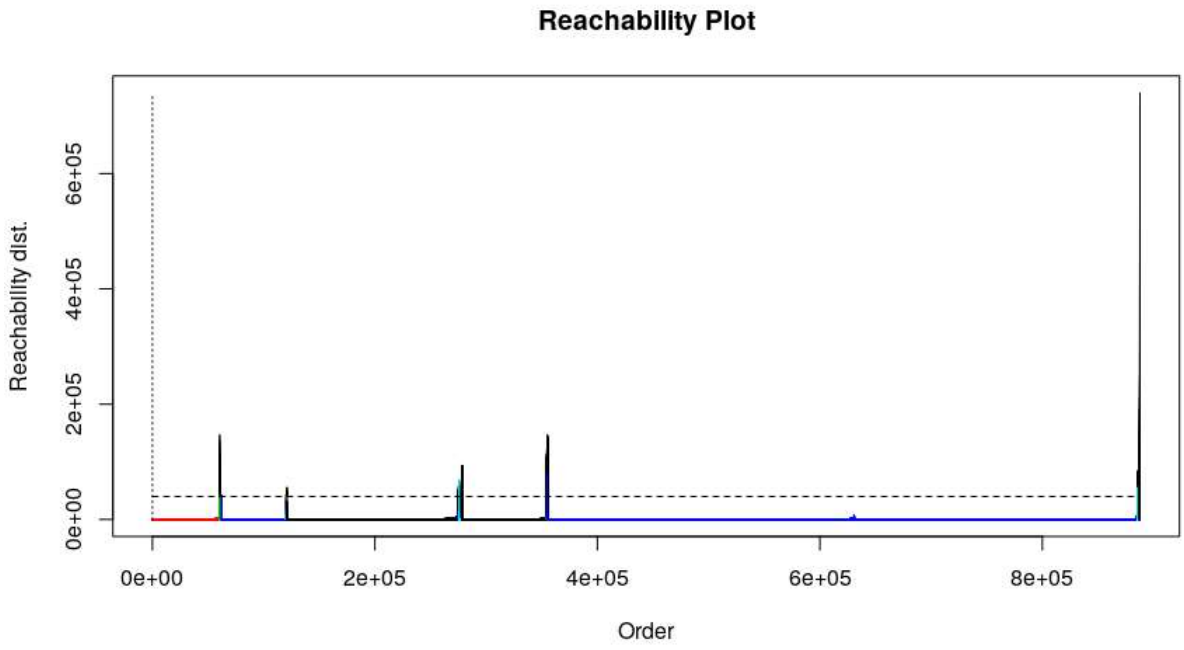

Figure 7.10: Reachability plot for all events in the pre-clustered dataset. The coloured areas (valleys) represent clusters, whilst the black areas (the spikes, corresponding to events very far away from their closest events) represent non-clustered events. Note that almost all the 887793 events are in clusters, and some of the non-clustered events are very far from the nearest cluster (high reachability distance). Thus, the reachability plot is mostly dominated by such valleys with occasional spikes, and appears stretched.

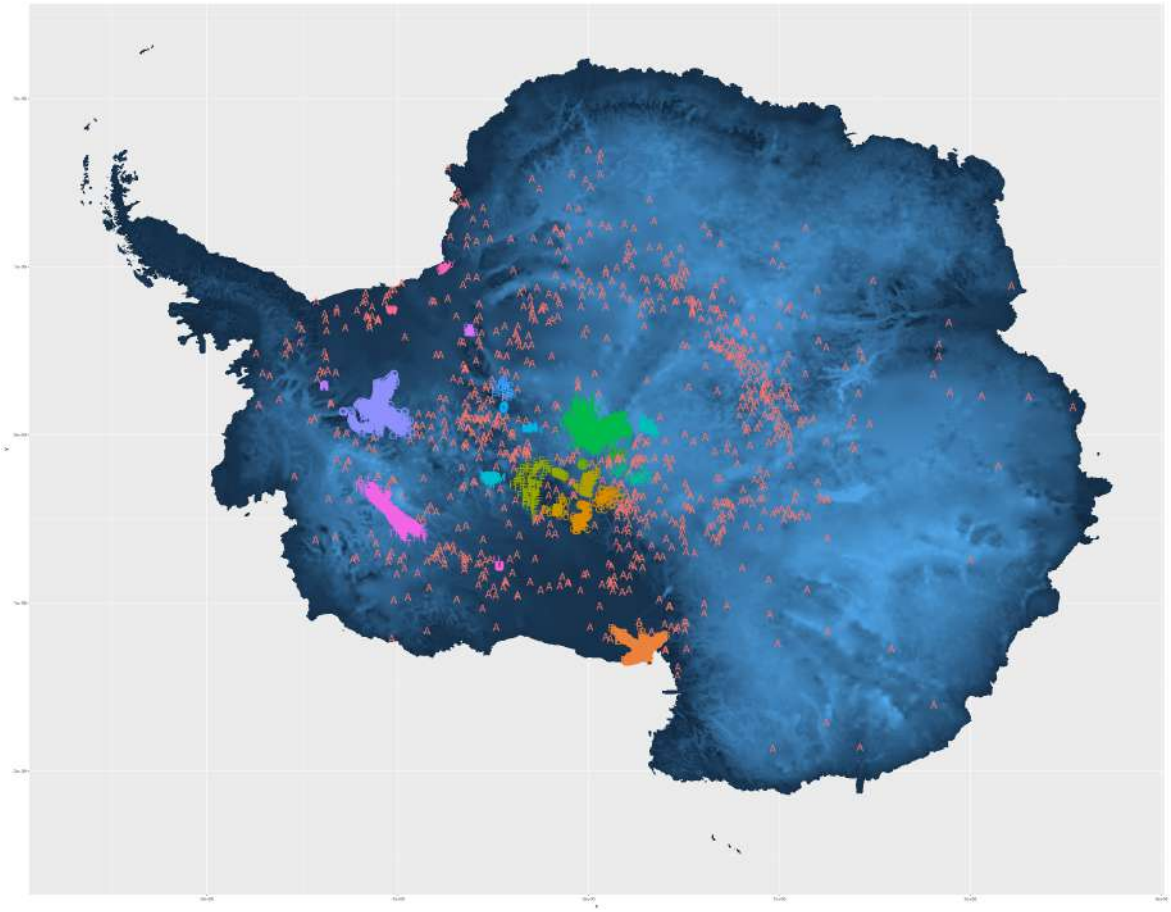

Figure 7.11: Pre-clustered events pointed back to the surface. Red events labelled "A" are isolated events, coloured events with other letters are those that clustered. The most active areas are identified, such as the Amundsen-Scott South Pole Station (central green), McMurdo Station (bottom orange), the West Antarctic Ice Sheet (pink), and Mount Vinson (purple).

The clustering map can be seen in Figure 7.11, and the pre-clustered events projected to the sky are shown in Figure 7.12. The OPTICS algorithm identified 21 clusters from a set of 887793 events, most corresponding to science stations or areas of high activity.

There are several features of the pre-clustered event sky map. The dense regions are each associated with a base. The curling loops from $100° < \alpha < 160°$ are when the payload passes right above Byrd Station (see Figure 7.13). The faint declination-constant curling loops in the background correlate to the WAIS calibration pulses (see Figure 7.14). The events which clustered on the continent were removed from the dataset and the remaining events were projected to a sky map

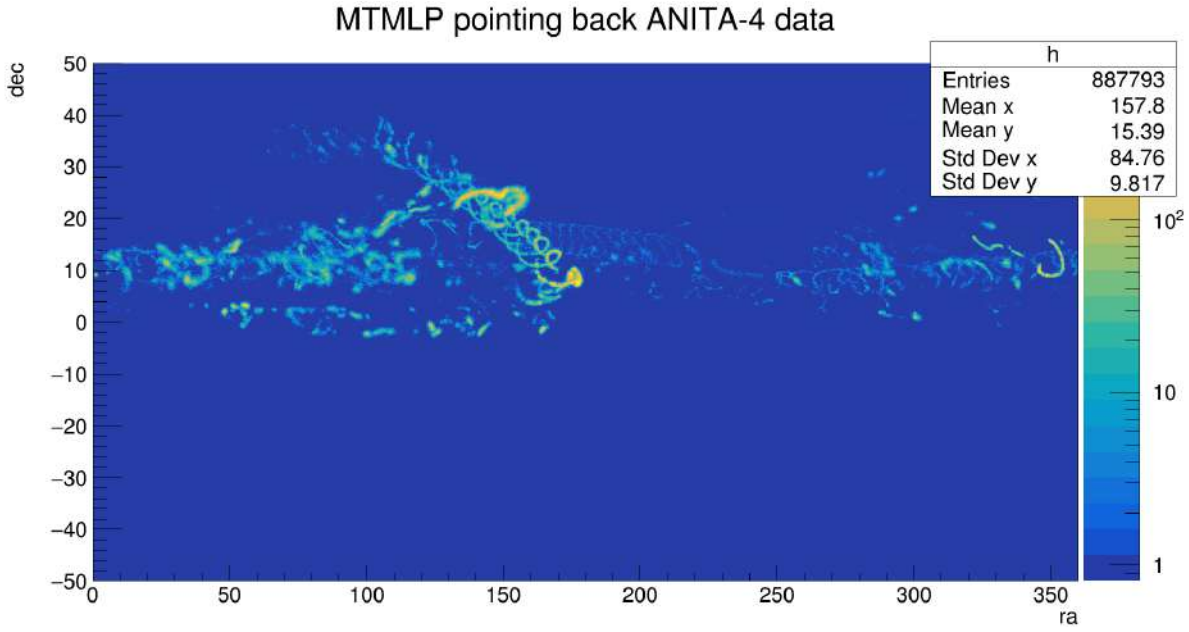

Figure 7.12: Pre-clustered events pointing back to the sky. The dense regions correlate to bases, which dominates the entire sky map. Right ascension (ra) and declination (dec) are always in degrees. The z-axis indicates the number of events in a single bin area.

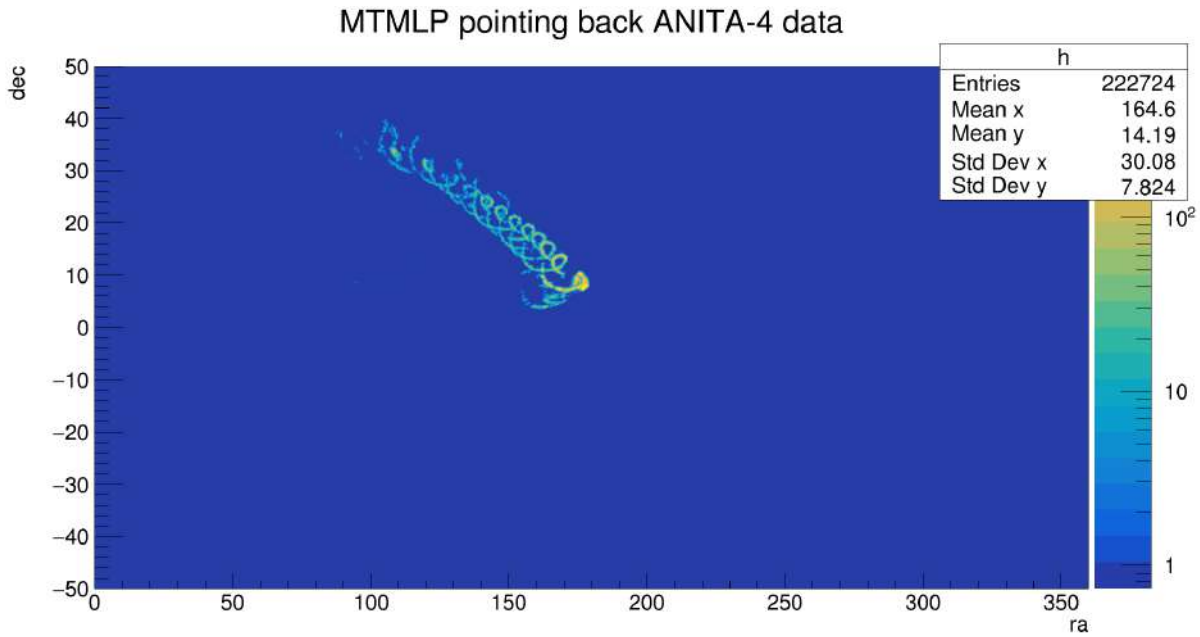

Figure 7.13: Events recorded when the ANITA payload passed right over Byrd station, projected to the sky. As the payload passes over the station, the declination gets much larger. Right ascension (ra) and declination (dec) are always in degrees. The z-axis indicates the number of events in a single bin area.

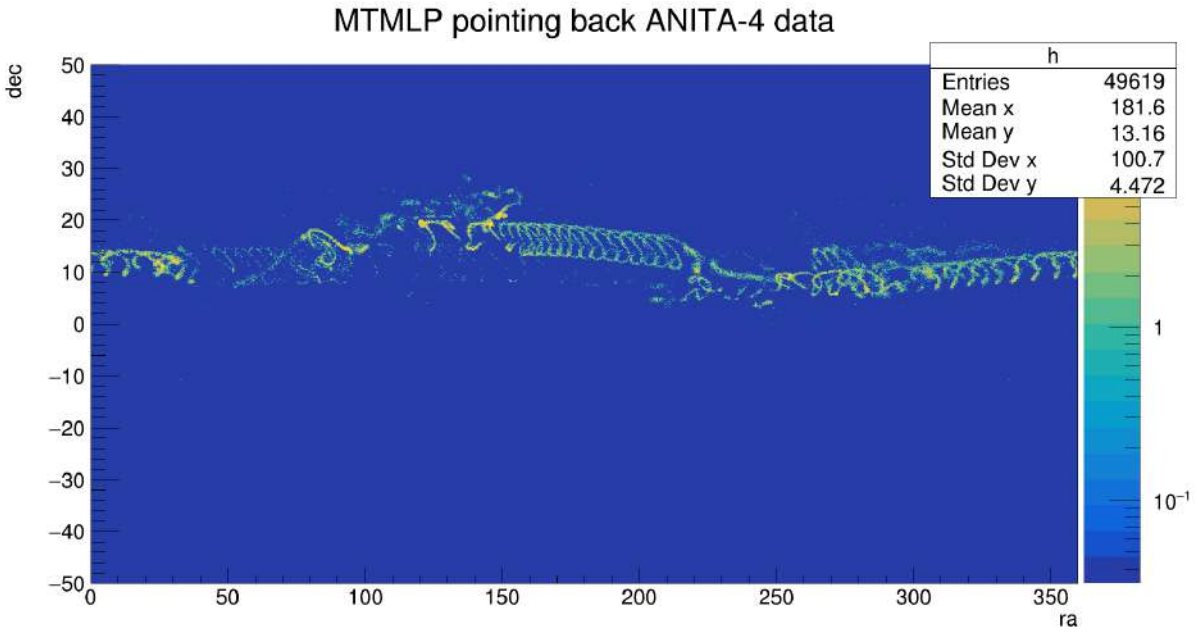

Figure 7.14: WAIS pulses pointed back to the sky. The declination remains constant as the payload did not travel right over the top of the calibration pulsers. Right ascension (ra) and declination (dec) are always in degrees. The z-axis indicates the number of events in a single bin area.

(see Figure 7.15). This was done via regression as explained in the next section. 917 events remained in the final post-clustered dataset. For all of the cuts leading to these remaining events, see Table 7.1.

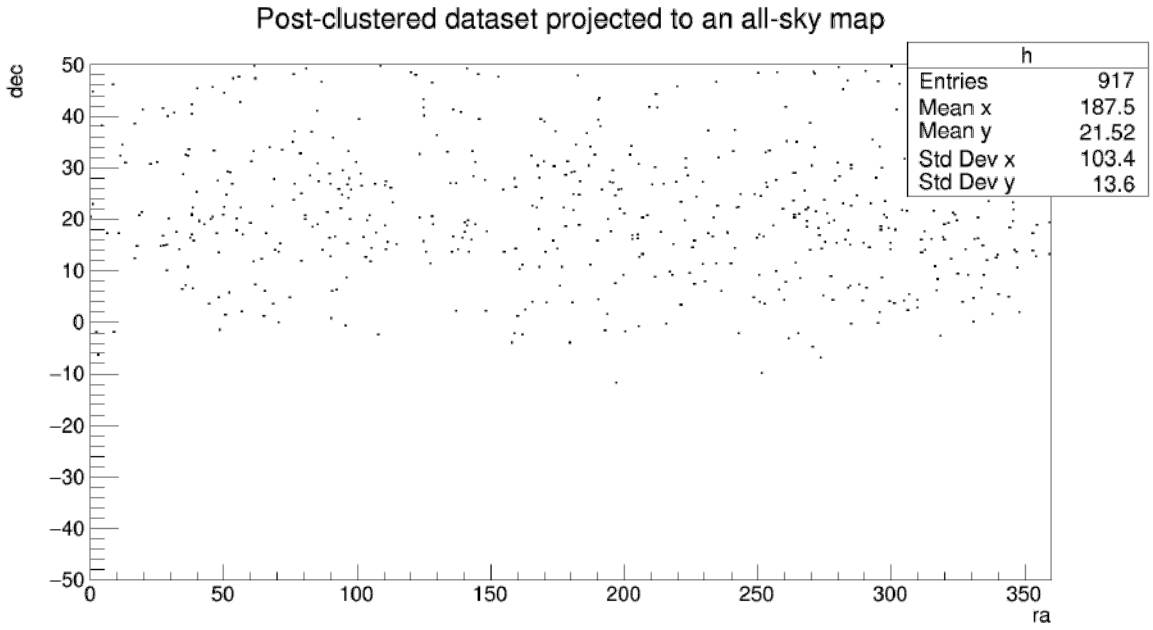

Figure 7.15: Post-clustered events pointing back to the sky. The distribution of events is as expected, as no dense regions or clusters are observed at all. As only isolated events remain, the scale of the z-axis is now irrelevant (bin areas can only be black or white). To make these events easily visible, the bin areas were increased, so some events appear to be touching, even though this is not the case. Right ascension (ra) and declination (dec) are always in degrees.

7.3 Source Search

In Chapter 6, it was explained how the tools to perform source searches were integrated into the simulation framework of `icemc`. Now that a post-clustered event dataset has been acquired, the events need to be pointed back to the sky.

7.3.1 Regression

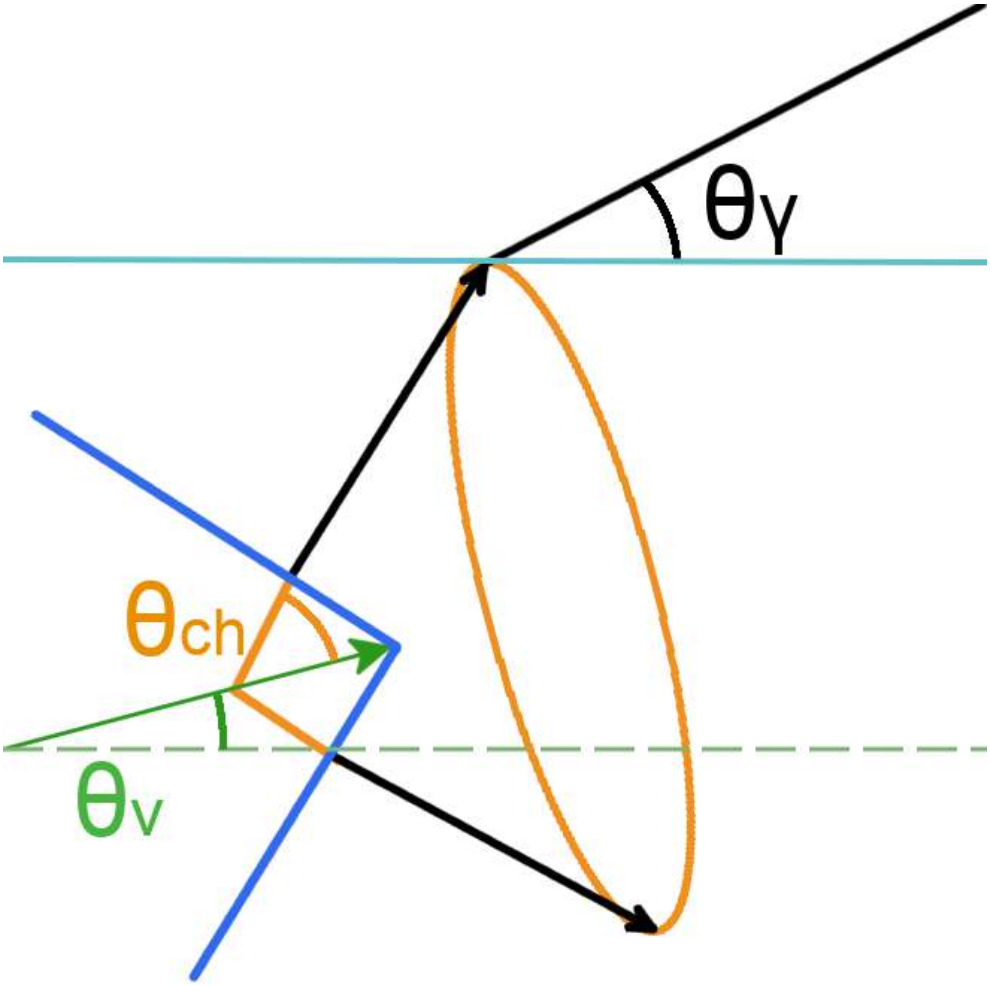

Figure 7.16: Geometric illustration of neutrino-induced radio waves, with elevation angles labelled. θ_ν indicates the neutrino angle, θ_γ indicates the radio-exit angle, and θ_{ch} indicates the Cherenkov angle.

In order to check if there is an excess of events that point to an astrophysical object, we first have to find a method to project the neutrino candidates back to their potential origin. We can't simply trace the neutrino's path back to its origin because we don't have such information readily available. A given radio event, however,

will provide several variables that can be used to reconstruct the neutrino's path. The first and foremost of these is the linear polarisation angle (l) of the event waveform. As detailed in Chapter 6, the polarisation of a radio wave associated with a neutrino will be very specific. A single radio wave will exit the ice at elevation angle, θ_γ, and azimuthal angle, ϕ_γ. Correspondingly, the neutrino's elevation angle is θ_ν, and its azimuthal angle is ϕ_ν. These neutrino angles are what we want to calculate, in order to reconstruct its path, but we only have polarisation and radio wave data. Thus, a pair of key variables, $(d\theta, d\phi)_{\gamma\nu}$, the difference between the radio-exit angles, $(\theta, \phi)_\gamma$, and the neutrino angles $(\theta, \phi)_\nu$ (see Figure 7.16 for a geometric illustration of the elevation angles involved) should form a relationship with linear polarisation angle (see Figure 7.17). As we know the values of the radio-exit angles via interferometry, as well as the polarisation (which gives us the difference of the two angles), it is possible to find the angles associated with the neutrino's direction of travel. This is particularly true for the difference in azimuthal angle, $d\phi_{\gamma\nu}$, due to its very strong dependence on linear polarisation angle (since the signal comes from a radially polarised Cherenkov cone). The difference between the radio-exit and neutrino angles can be calculated as,

$$d\theta_{\gamma\nu} \equiv \theta_\gamma - \theta_\nu, \qquad d\phi_{\gamma\nu} \equiv \phi_\gamma - \phi_\nu \,. \tag{7.1}$$

To find the most likely neutrino pointback angle, we use these three variables, θ_γ, ϕ_γ, and l, as an input to a multi-layer perceptron (MLP). In such a case, we can use a *multi-target* MLP (MTMLP), to find both neutrino point-back angles at the same time. The events used were weighted, as per the event weights provided by `icemc`.

A large set of diffuse neutrinos are simulated in icemc for the ANITA-4 instrument and its flight path. 1.6 billion initial neutrinos are simulated in total (160 runs, 10 million initial neutrinos per run), and processed by the Hypatia astrophysics computing cluster in parallel. Approximately 1 in 1000 events pass the simulation and are detected by the simulated instrument (~1.6 million surviving events). After a weight cut was applied, ~400,000 good neutrino candidates are left.

As the simulation is used for the test dataset, we actually know where the neu-

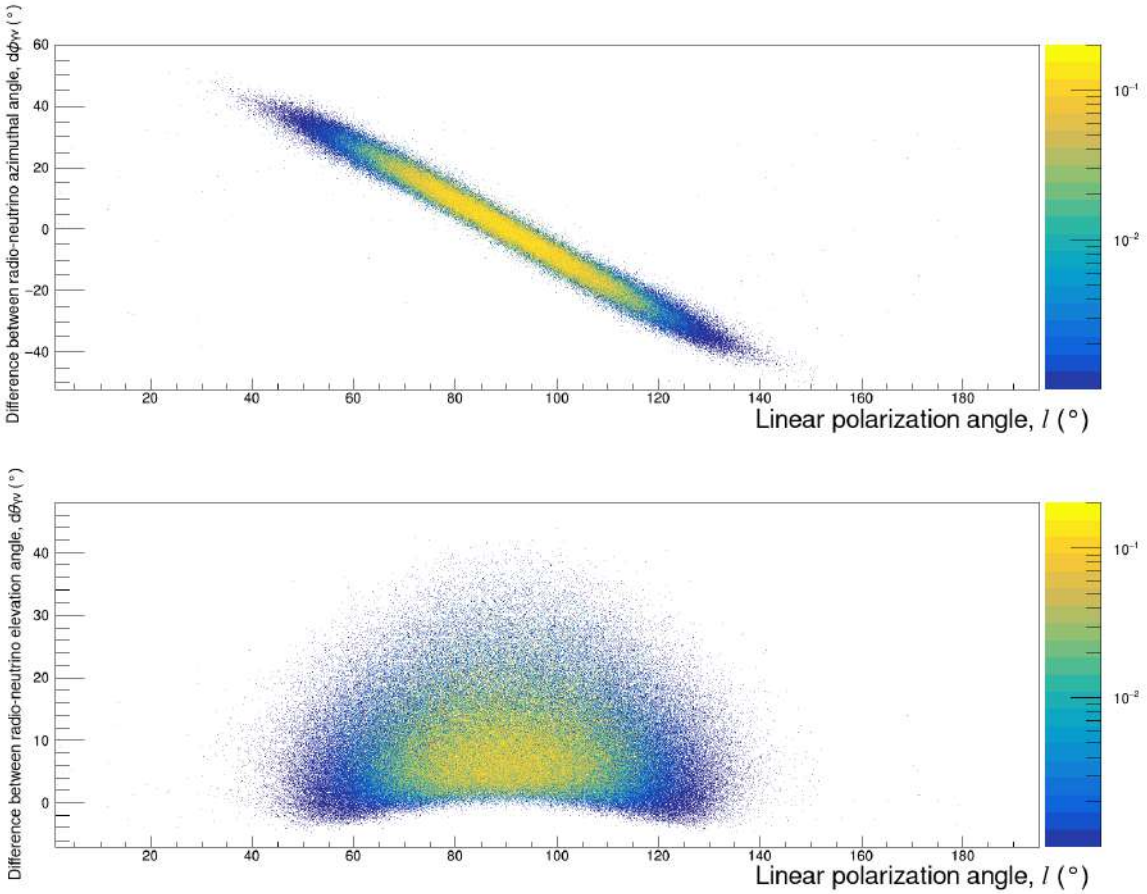

Figure 7.17: Histograms showing the relationship between d$\phi_{\gamma\nu}$ (top) and $d\theta_{\gamma\nu}$ (bottom) vs linear polarisation angle, l, for a set of diffuse neutrinos.

trinos came from, as we used the source origins to simulate them in the first place. Therefore, we can compare the results from the regression and the simulation to assess our accuracy. To verify the results, we use the set of diffuse neutrinos. An analysis framework designed for the ANITA experiment, `UCorrelator`, is used to process the simulation output files [130]. `UCorrelator` processes the waveforms ANITA records, applies sine-subtraction filtering and deconvolution techniques to the waveforms and uses interferometry to find the peak of a theta-phi map. A program was made within `UCorrelator` to process simulated data for ANITA-4. From this, we can extract the peak azimuthal and elevation angle of the radio waves. Three data variables are inserted into the MTMLP to predict the azimuthal and elevation angles of the neutrinos.

The ~1.6M surviving neutrinos are split into two datasets: training and test, in

an 80-20 split. The MTMLP is formed of 3 input variable nodes and a bias node, a single hidden layer, and two outputs (illustrated in Figure 7.18).

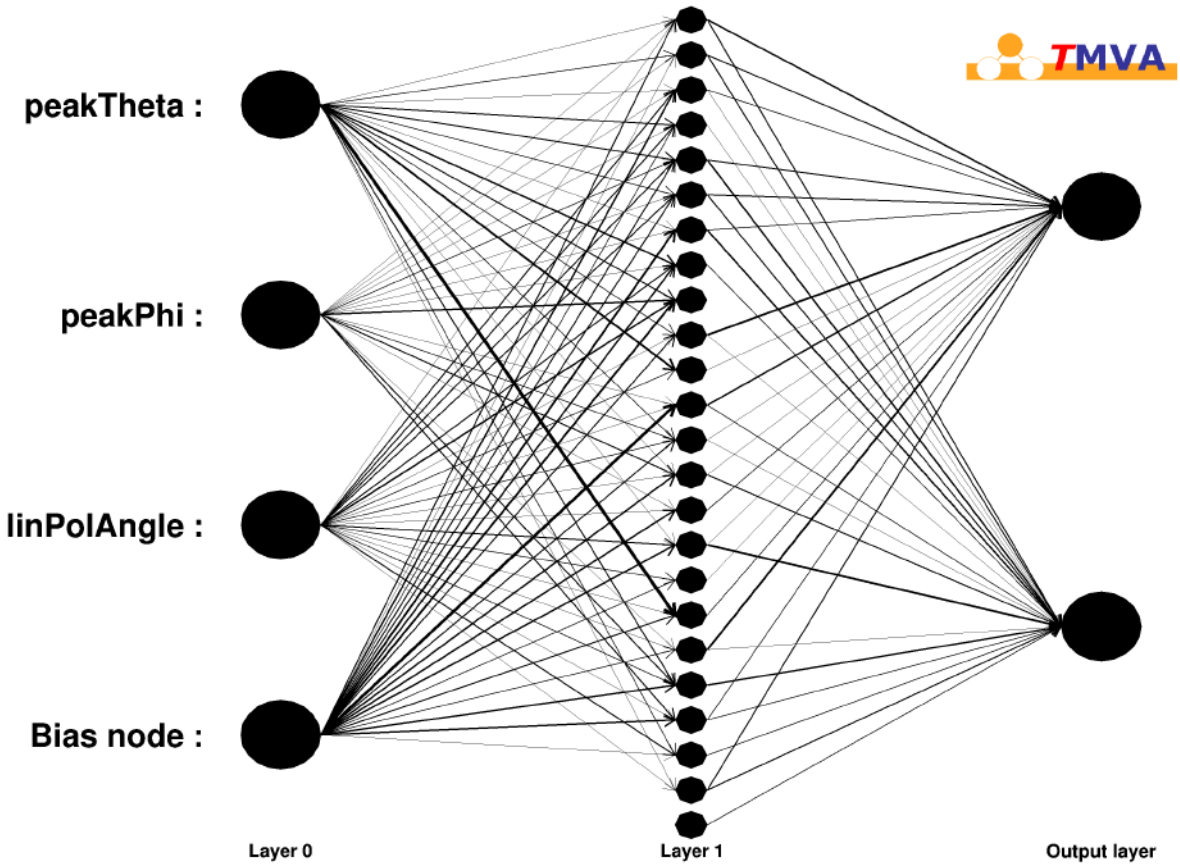

Figure 7.18: Node structure of the MTMLP. The input variables and bias node are shown to the left (input is layer 0). The 24 nodes are shown in (the hidden) layer 1, and the output layer contains the two targets of the MLP.

The test data is then subjected to the MTMLP. To compare the regression results to that of the simulation, we can find the difference in the neutrino angles that they each return. The resolutions for both elevation (R_θ) and azimuthal angle (R_ϕ) are thus defined as,

$$R_\theta \equiv \theta_\nu^{\text{reg}} - \theta_\nu^{\text{sim}}, \qquad R_\phi \equiv \phi_\nu^{\text{reg}} - \phi_\nu^{\text{sim}}. \tag{7.2}$$

The combination of both simulation and regression worked excellently for diffuse neutrinos (see Figure 7.19), as the standard deviations between the actual and predicted angles were only about 2 degrees. What we are really interested in are neutrinos from astrophysical sources, and we must verify that the regression works

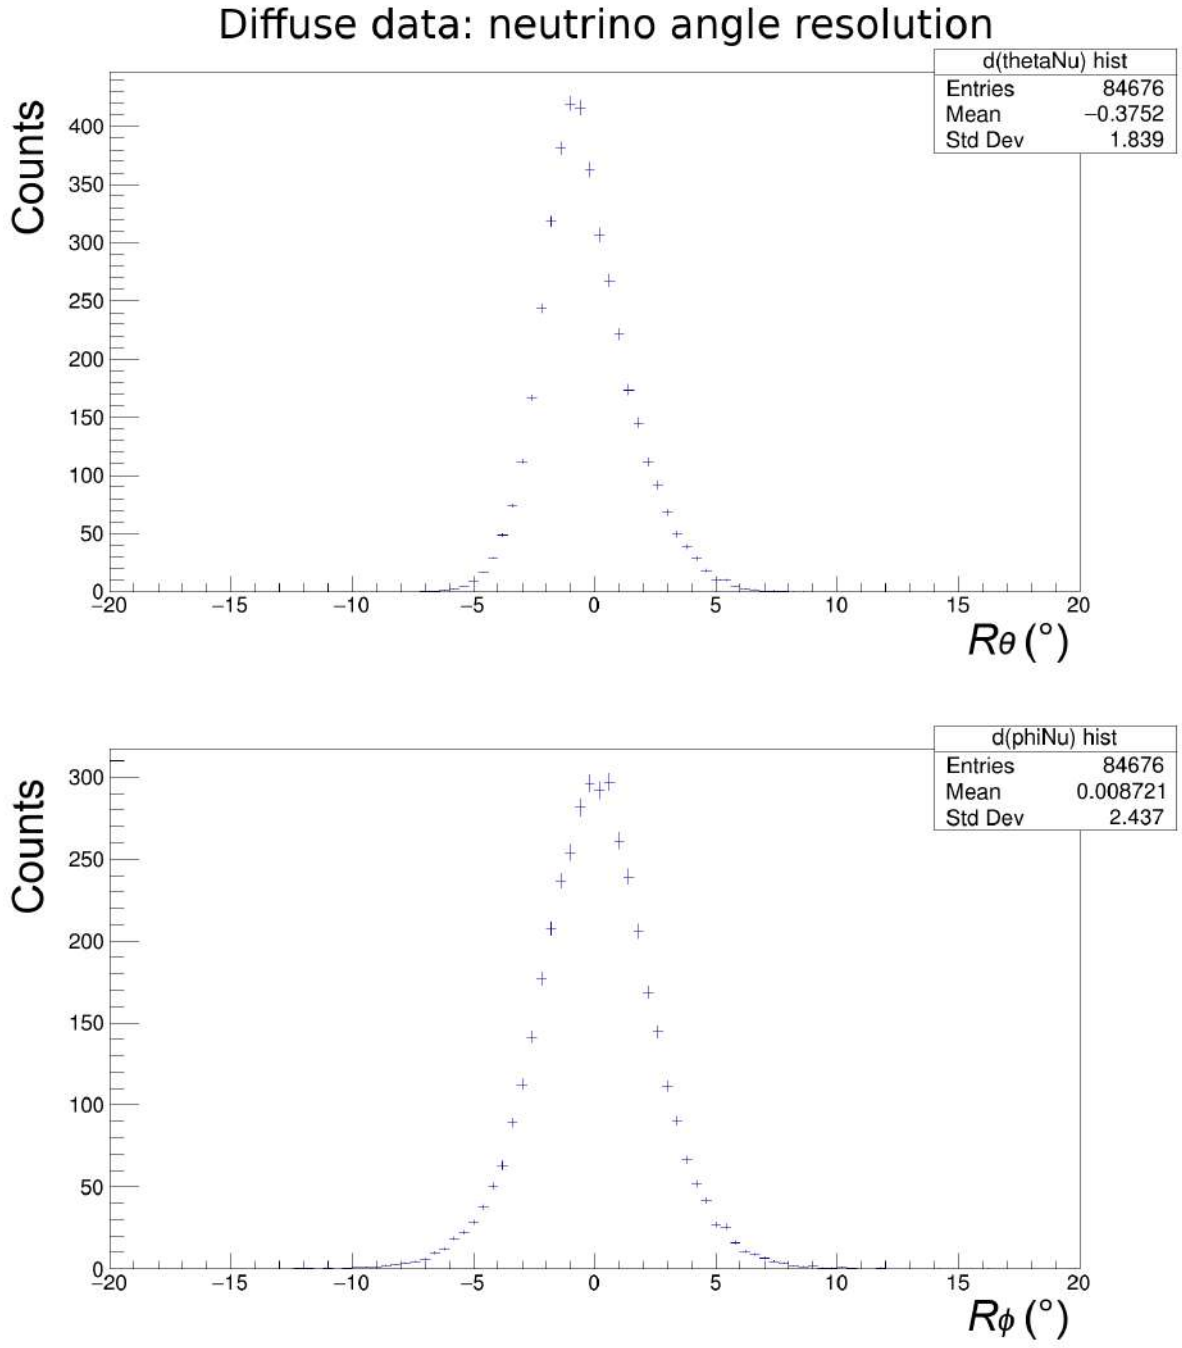

Figure 7.19: Resolution histograms for diffuse data: The difference between the neutrino angles from regression and from simulation (test data). The standard deviation for the elevation angle is less than 2°, whilst the standard deviation for the azimuthal angle is just above 2°.

for these. As such, `icemc` was again used to simulate neutrinos from individual sources. Section 6.1.2 goes over how this was implemented into the simulation. After being simulated, the output files were once again processed with

`UCorrelator`, and fed into TMVA. The events were sent through the MTMLP and we acquired elevation and azimuthal angles for the neutrino. The regression for individual sources were better than expected, with even better resolution than neutrinos from all points in space (diffuse data set). Neutrinos from the simulated diffuse search can come from anywhere in space, but most neutrinos that survive the selection criteria come from objects of low absolute values of declination (if the declination is too obtuse, the neutrino either comes from above or has to make its way through the entire earth and does not skim the Antarctic ice sheet). As such, there are more neutrinos associated with low absolute values of declination (see Figure 7.20), and the data is trained and tested upon these surviving neutrinos. Therefore, we expect the neutrinos from low-declination objects to have a better resolution. See Figure 7.21 for an example for a single blazar. The distributions are very similar for the range of blazars considered.

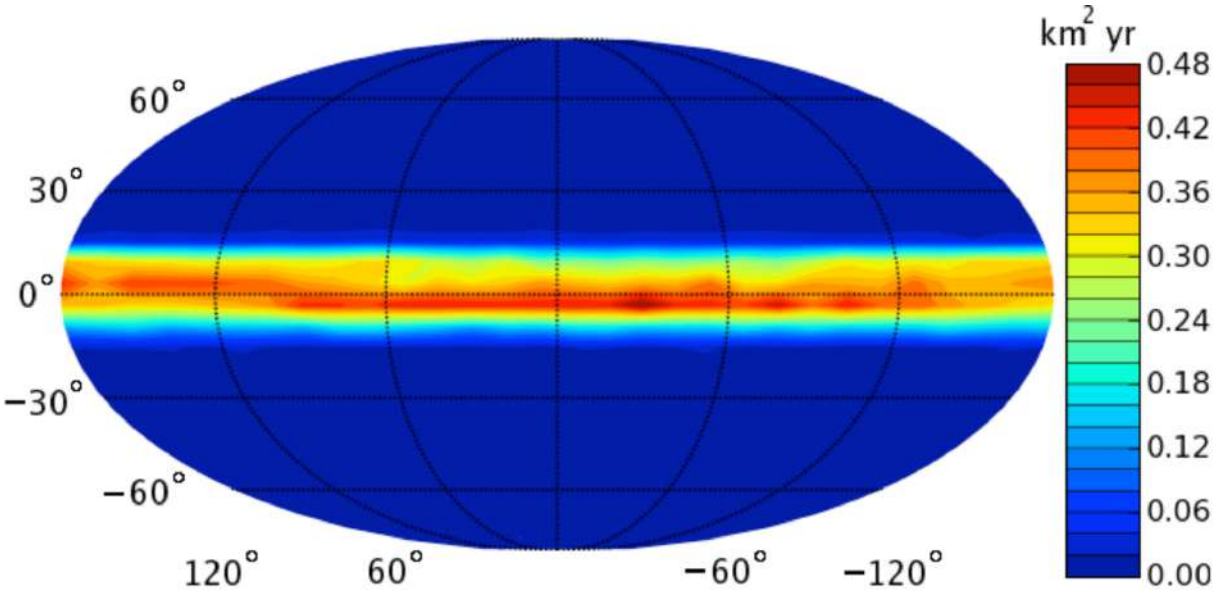

Figure 7.20: Simulation: a map of ANITA's sky exposure to 10^{20} eV neutrinos. Note that the exposure is much higher at low absolute values of declinations, whilst it is uniform thoroughout all values of right ascension. Image from [131].

7.3.2 Pointback

Now that we have confirmed the good agreement between simulation and regression, we can simulate neutrinos from a single blazar, PKS 1502+106

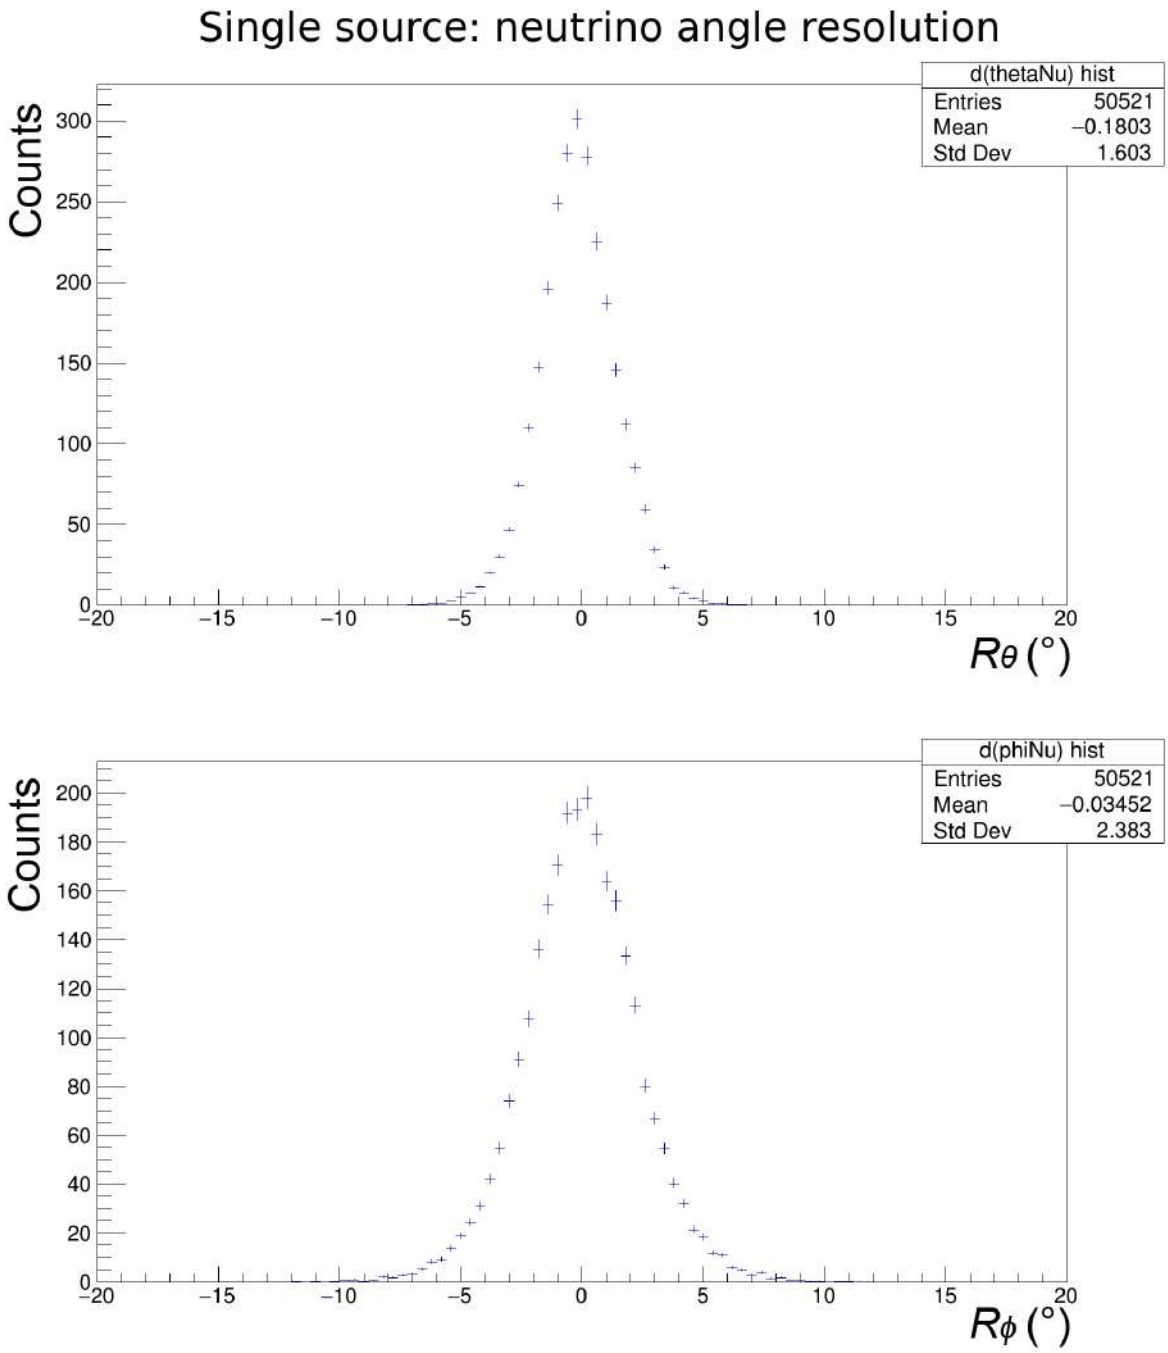

Figure 7.21: Resolution histogram for a single blazar: the difference between neutrino angles from regression and from simulation.

(α = 225.19°, δ = 10.49°), and use the regression to see if neutrinos correctly point back it. To do this, `UsefulAdu5Pat` was used, which is in `anitaEventCorrelator` [132], a set of programs created and maintained by the ANITA Collaboration. When calling UsefulAdu5Pat, the position and attitude

information from the ADU5 GPS units (see Section 5.1.7.1) are available to use. At a certain time within a flight, whether real or simulated, there is an asscociated time, and the parameters that are recorded at such a time.

Each ANITA event has associated radio-exit angles, so we can use these to find the neutrino angles. We now know the point-back angles of the neutrinos, $(\theta, \phi)_\nu$. Azimuth is taken to be north, and the direction ANITA is facing (heading, ψ) is used, so the neutrino point-back angles $(\theta, \phi)'_\nu$ are adjusted to be,

$$\phi'_\nu \equiv \phi_\nu - \psi, \qquad \theta'_\nu \equiv -\theta_\nu \,. \tag{7.3}$$

The declination (δ) and hour angle (h), which tells us where the object is in space when taking into account Earth's rotation, can then be established from these neutrino angles and latitude (lat),

$$\delta = \text{asin}\left(\sin(\text{lat})\sin(\theta'_\nu) + \cos(\text{lat})\cos(\theta'_\nu)\cos(\phi'_\nu)\right) , \tag{7.4}$$

$$h = \text{atan2}\left(\sin(\phi'_\nu), -\cos(\phi'_\nu)\sin(\text{lat}) + \tan(\theta'_\nu)\cos(\text{lat})\right) . \tag{7.5}$$

The (local mean sidereal) time of the event and the hour angle is used to find the right ascension of the object as per Equation 6.35.

Pointing back to the spot where the blazar is in space, we see a well contained distribution of pointed-back neutrino events. The right ascension and declination of the distribution ($\alpha = 225.9° \pm 2.5°$, $\delta = 10.15° \pm 1.64°$) are well within the actual celestial position of PKS 1502+106, as shown in Figure 7.22.

7.3.2.1 Verification

We can also verify how both the regression and simulation perform by simulating neutrinos from a set of sources, and observing where they point back to. If correct, the neutrinos should point back to the most active sources at the smallest absolute values of declination.

Consider the set of flaring blazars (from FAVA), for ANITA-4, as shown in Table 7.2 and visualised on a sky map in Figure 7.23. Notice that only three of these blazars, which are observed to be flaring for more than a week, are within a rea-

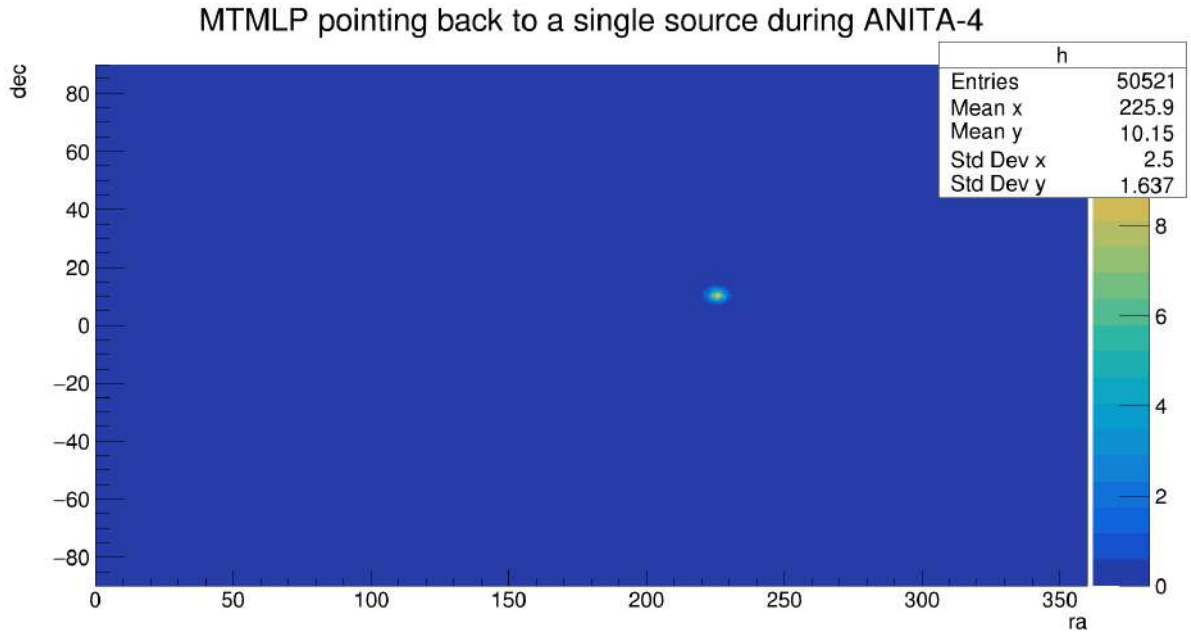

Figure 7.22: Pointing back to a single source of simulated neutrinos via regression. The distribution is very well contained, with essential no stray events. Right ascension (ra) and declination (dec) are always in degrees.

sonable declination range. Using `icemc`, each of the blazars associated with flares during the ANITA-4 flight were set as an initial source of neutrinos. The simulation was then run, and the resulting simulated data were processed, and pointed back to the sky. The results are shown in Figure 7.24. There seem to be three main pointed back distributions, associated with 4C +01.02, CTA 102, and PKS 1502+106. The distribution associated with the location of 4C +01.02 is the most prominent. This is not at all surprising as it has the most acute declination and was observed by FAVA during each week of the ANITA flight.

7.3.3 Proximity

It has been shown that it is possible to point back to a source within good accuracy. However, we want to characterise pointed-back neutrino events by how close they are to the object of interest. Thus, we need some sort of definition of *proximity*, not only in space, but in time as well.

The spatiotemporal proximity, or simply the proximity (P), is defined to be the product of individual spatial (S) and temporal (T) proximities:

Name	Source class	RA (°)	Dec (°)	Observation weeks
PKS 1016-311	FSRQ	154.55	-31.38	1
PKS 2210-25	FSRQ	333.18	-25.32	1
PKS 0521-36	BCU	81.21	-36.13	1
4C +01.02	FSRQ	17.2	1.59	1,2,3,4
CTA 102	FSRQ	338.11	11.74	1,2
S5 1044+71	FSRQ	162	71.73	1,2,3,4
PKS 1502+106	FSRQ	226.21	10.31	1,2
PMN J2234-2656	BCU	338.52	-26.99	1
S5 2007+77	BLL	301.34	77.89	2
PKS B1424-418	FSRQ	217.12	-42.27	2,4
H 2356-309	BLL	0.38	-31.06	3
OK 630	FSRQ	140.39	62.3	3,4
4C +15.05	BCU	31.18	15.31	4
7C 2010+4619	BLL	302.99	46.76	4

Table 7.2: The set of blazars associated with flares during the ANITA-4 flight. The source classes for blazars are flat-spectrum radio quasars (FSRQs), BCUs (blazars candidates of uncertain type), and BLLs (BL Lacertae objects). Note, for blazars that are observed in more than one week, only the first observation of right ascension and declination are shown.

$$P \equiv ST\,. \tag{7.6}$$

However, it is often more useful to work with the logarithm of proximity (ρ):

$$\rho \equiv \log(P) = \log(ST), \tag{7.7}$$

$$= \log(S) + \log(T)\,. \tag{7.8}$$

7.3.3.1 Spatial Proximity

The spatial proximity, S, determines how close an event is to an object in space. In order to assess this, each object has its own spatial proximity function in space, which is dependent on right ascension and declination.

In order to form this function, a few steps are required. Firstly, a fixed number of initial neutrinos is simulated from the source and each neutrino is pointed back to celestial coordinates. This gives a distribution as shown in Figure 7.25. In order to make it easier to fit a function to the distribution, the set of event are smoothed

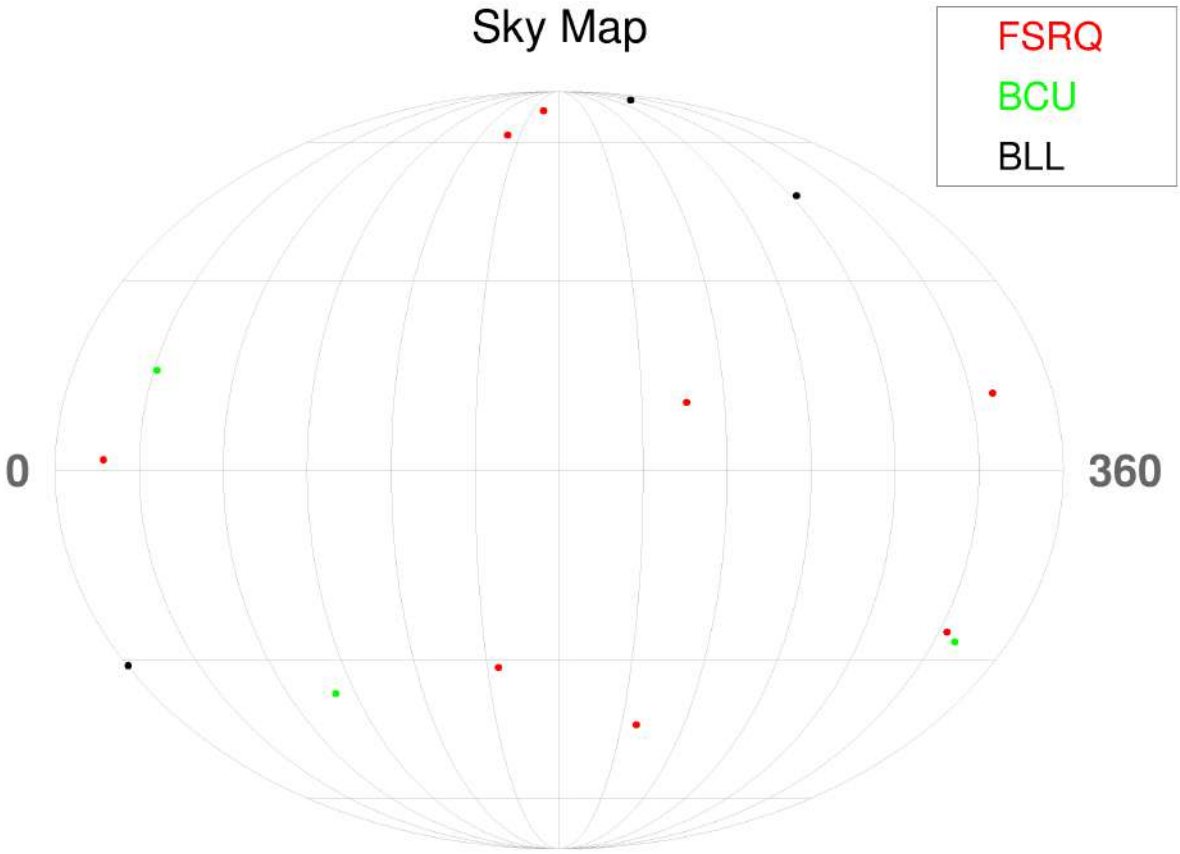

Figure 7.23: A sky map of the set of flaring blazars active during ANITA-4. Flat-spectrum radio quasars (FSRQs) are shown in red, BCUs (blazars candidates of uncertain type) are shown in green, and BLLs (BL Lacertae objects) are shown in black.

with the k5b kernel available in ROOT's smoothing algorithm.

The function we fit to the distribution is a right-ascension-wrapped bivariate Gaussian:

$$S(\alpha,\delta) = \frac{1}{2\pi\sigma_\alpha\sigma_\delta\sqrt{1-\rho^2}} \exp\left\{-\frac{1}{2(1-\rho^2)}\left[\left(\frac{\alpha-\mu_\alpha}{\sigma_\alpha}\right)^2 - 2\rho\left(\frac{\alpha-\mu_\alpha}{\sigma_\alpha}\right)\left(\frac{\delta-\mu_\delta}{\sigma_\delta}\right) + \left(\frac{\delta-\mu_\delta}{\sigma_\delta}\right)^2\right]\right\},$$

where $S(\alpha,\delta)$ is the spatial proximity function, α is right ascension, δ is declination, σ_α is the standard deviation of right ascension, σ_δ is the standard deviation of declination, ρ is the $\alpha-\delta$ correlation, μ_α is the mean of right ascension, and μ_δ is the mean of the declination. An example including typical parameter values for an object is shown in Figure 7.22.

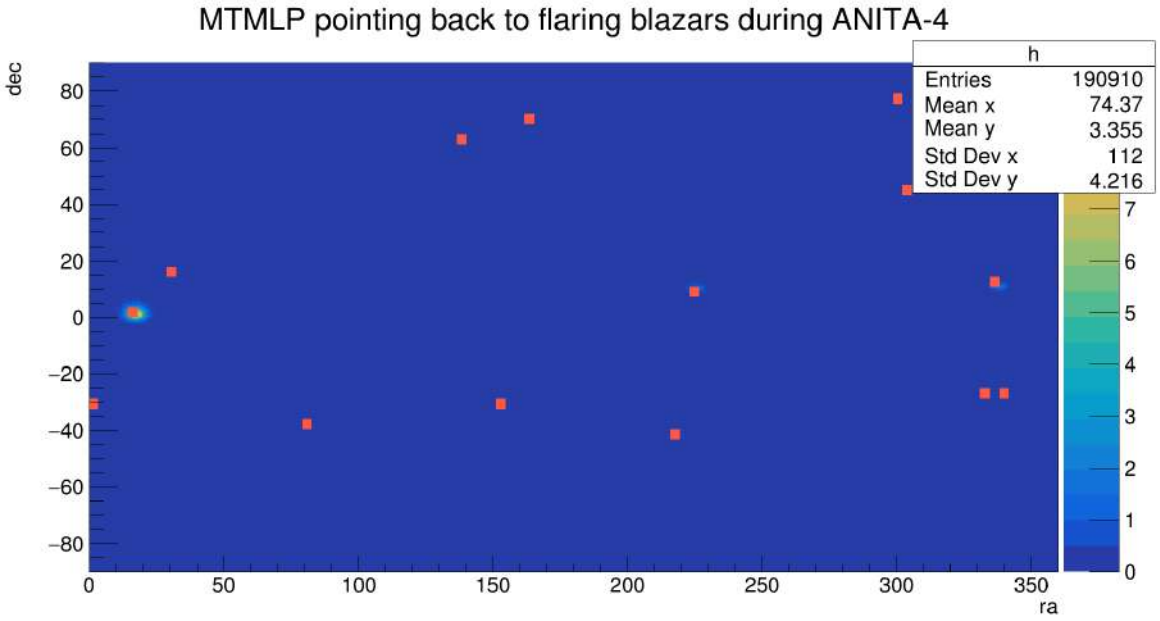

Figure 7.24: Pointed back neutrinos overlaid with all blazar sources. The red squares represent single observations of blazars as $4° \times 4°$ bin areas. The distributions represent the locations where the simulated neutrinos point back to. Right ascension (ra) and declination (dec) are always in degrees. The z-axis denotes the amount of events in a bin area.

Within `icemc`, 160 million initial neutrinos were simulated for each source. Only neutrinos from favourable sources, i.e., those of reasonably low absolute value of declination, will pass the simulation and be detected. If less than 1000 neutrinos pass for a single object, it becomes impossible to form a spatial proximity function. In this case, the disfavoured sources are not included in the final analysis. Though this may seem like a lot of required neutrinos, a fairly high amount is needed to form a spatial proximity distribution that can be fit with a right-ascension wrapped bivariate Gaussian.

Based on the shape of the pointed-back neutrinos, we can find out the distribution's parameters (such as the mean) and use it as a starting point for the fitter. The wrapped bivariate Gaussian is then fitted to the distribution, giving us a function only dependent on right ascension and declination. In order to set the scale of the function so that it will always return a value between 0 and 1, the bivarate Gaussian is normalised, such that the peak is 1. The distributions will be different for every source.

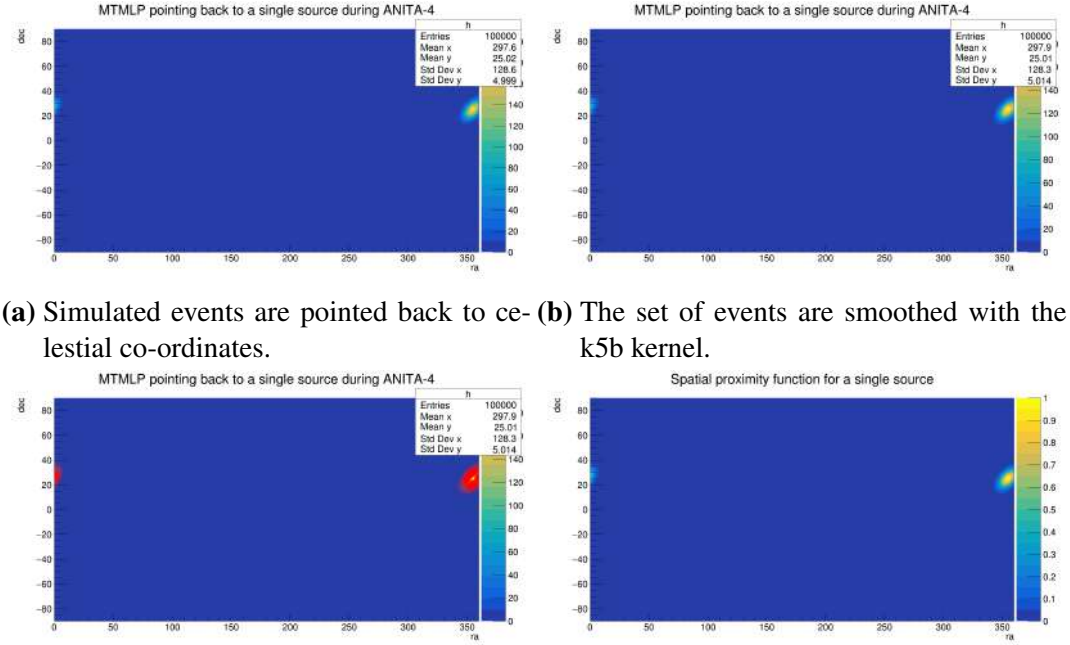

(a) Simulated events are pointed back to celestial co-ordinates.

(b) The set of events are smoothed with the k5b kernel.

(c) A wrapped bivariate Gaussian is fitted to the distribution.

(d) The Gaussian is normalised between 0 and 1 to set scale of S.

Figure 7.25: Example: Construction of the spatial proximity function for a single source. Right ascension (ra) and declination (dec) are always in degrees.

When each event from the data is pointed back to a point in space, its location is input to the spatial proximity function, returning a single spatial proximity.

7.3.3.2 Temporal Proximity

The temporal proximity, T, determines how close an event is to an object in time. Consider the following scheme: An object is only classed as ON ($T = 1$), i.e., temporally active, when it falls within the expected range of times. In all other cases, it will be OFF ($T = 0$). If we use such a scheme, a neutrino that points directly back to an object in space, but is off of the accepted time range, say 0.1 seconds off, will be completely rejected. It is far more logical to work with a scheme whereby the temporal activity starts to fade out with time i.e., the temporal proximity itself is dependent on time, and becomes a continuous variable. Within the expected time range, $T(t)$ can just be assigned as 1. However, when out of the expected time range, the temporal parameter quickly fades out, i.e., T is less than 1, but is still finite. Defining the temporal proximity as such is also very useful for short-lived

events.

How quickly an object temporally fades out, $\frac{dT(t)}{dt}$, is dependent on object type, as well as how long it was on (its duration) during the ANITA-4 flight, d. Each object has a different reported duration, and these can be extracted from the catalogues. An object is reported to be on from t_0 to t_1, and so both a temporal fade-in and fade-out must be incorporated. The temporal proximity function was chosen not to be Gaussian as we know for sure that the object is on within its reported range, thus the distributions have a flat top. As such, in order to account for the temporal activity of different source types, a source type scaling parameter, c, is introduced. This is a constant that essentially determines how long we allow the source to fade in (out) for before (after) the start (end) of its reported duration.[5] The whole temporal function for a single object can be represented as three separate functions:

$$T(t) = \begin{cases} 10^{\frac{c(t-t_0)}{d}} & : t < t_0, \text{ fade-in}, \\ 1 & : t_0 < t < t_1, \text{ reported ON}, \\ 10^{\frac{c(t_1-t)}{d}} & : t > t_1, \text{ fade-out}, \end{cases} \tag{7.9}$$

where d is the duration and c is a source type specific scaling constant.

However, this temporal proximity function does not describe the temporal profiles of all flaring objects well. For example, the blazar, PKS B1424-418, flared on two separate occasions during the ANITA-4 flight. For this, the secondary flare and the between-flare overlap must be established. The secondary flare has times t_0' and t_1', and duration d', and is essentially treated the same as the first flare. Like in the single-flare case, d' is just extracted straight from the catalogue. c is the same as previously calculated as these flares belong to the same source class. The between-

[5] If a parameter such as c was not to be introduced, a large problem would arise. Firstly, objects with very low reported durations, such as GRBs, would not have any significant temporal fade-in and fade-out at all. This would mean that if we detected a GRB just slightly out of its reported time, it would be completed ignored, and we would have missed events associated with this object. Therefore, c must be calculated for each source class. This is done by taking the average duration for objects of a single source class, and finding when the original temporal proximity function drops to 0.1, a tenth of the value for when it was on.

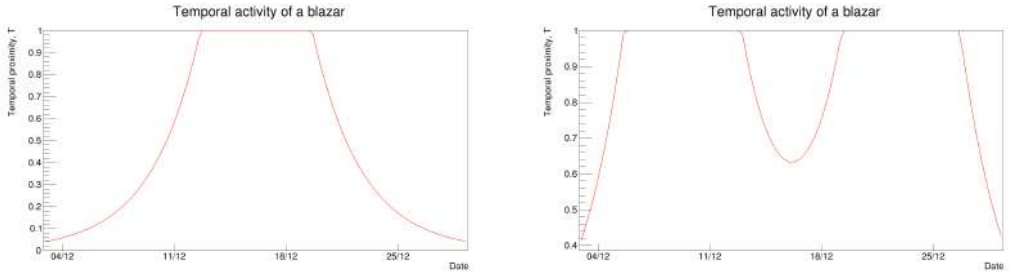

(a) Single-flare blazar. **(b)** Double-flare blazar.

Figure 7.26: Temporal proximity functions for two blazars.

flare overlap is treated as the sum of the fade-out of the first flare and the fade-in of the second flare (but cannot become greater than the maximum value for T, which is 1). See Figure 7.26 for a visualisation for the temporal proximity graph for two sources.

$$T(t) = \begin{cases} 10^{\frac{c(t-t_0)}{d}} & : t < t_0, \text{ flare 1 fade-in}, \\ 1 & : t_0 < t < t_1, \text{ flare 1 reported ON}, \\ 10^{\frac{c(t_1-t)}{d}} + 10^{\frac{c(t-t'_0)}{d'}} & : t_1 < t < t'_0, \text{ between-flare overlap}, \\ 1 & : t'_0 < t < t'_1, \text{ flare 2 reported ON}, \\ 10^{\frac{c(t'_1-t)}{d'}} & : t > t'_1, \text{ flare 2 fade-out}. \end{cases} \tag{7.10}$$

7.3.3.3 Spatiotemporal Proximity

We already know how to calculate the individual proximities, and combining them leads to an expression with three free parameters: right ascension, declination, and event time. The spatiotemporal proximity is thus:

$$P(\alpha, \delta, t) = S(\alpha, \delta)T(t) \tag{7.11}$$

To test how well our definition of proximity fairs, we can simulate a set of diffuse neutrinos, as well as a set of neutrinos from a source. A logarithmic scale is used for both the proximity and the events to demonstrate this. We expect events from individual sources to be as close to zero as possible. Most diffuse events

should have very low log proximities, though we should expect some to come from the source, as diffuse events can come from anywhere. To visualise this, both sets of neutrinos are plotted as a histogram (see Figure 7.27).

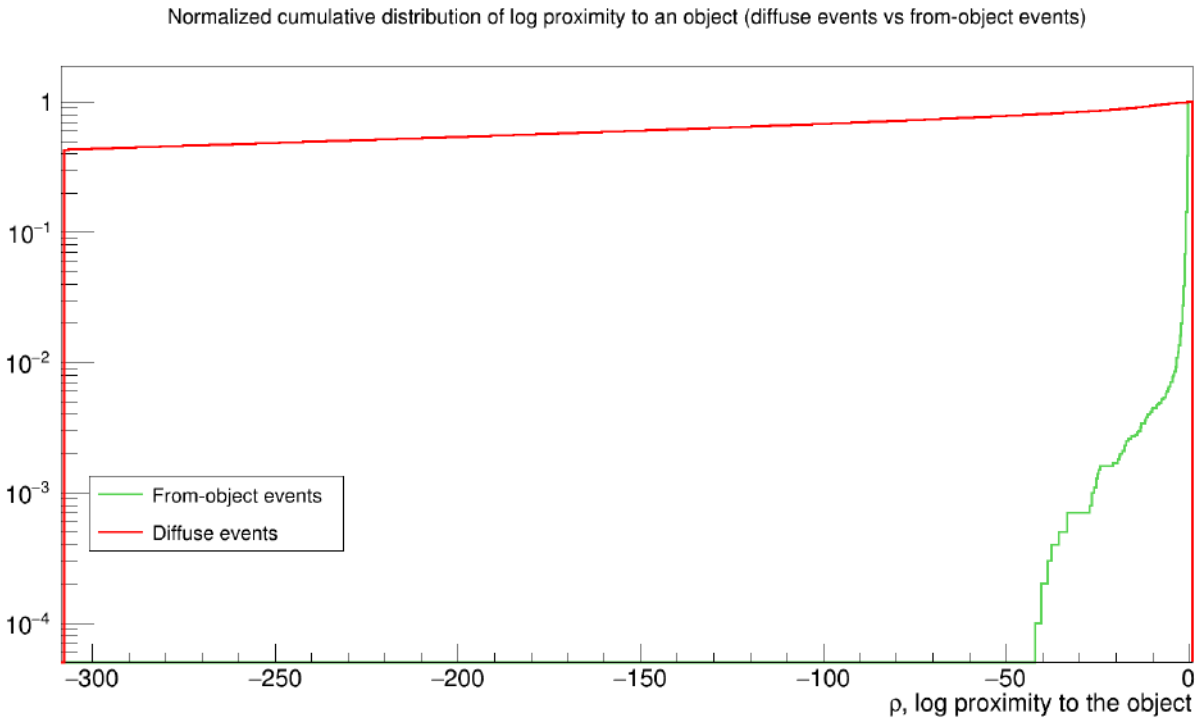

Figure 7.27: Normalised cumulative distribution of the spatiotemporal proximity for a single source (green) and a diffuse set (red).

7.3.3.4 Alternative Universes

Following an observation of an astrophysical object of any type in our Universe, it has fixed position in space, as well as fixed temporal activity (its duration and reported start time are fixed). Consider that several events point back to this object; they each have a high proximity. As the set of post-clustered events is constant, the summed proximity should also be high, as several events point back to the same object. Now consider that the spatiotemporal properties of the object are altered, i.e., the object is in a different place in space, and has a different temporal activity. The events would no longer have a high summed proximity to the object, as the object itself is no longer the same in respect to its observed properties. We would not expect to measure this object with these alternative properties in our Universe. Thus, we denote this as an object of an alternative universe.

In each alternative universe, the object is spatially displaced. As we are already confined to a very narrow declination band, the object is only displaced in right ascension. The spatial proximity distribution is strongly dependent on declination, so adjusting the object in declination would require a huge amount of extra simulations for each different declination. When the spatial properties of the object are altered, we make sure that they are not in range of the object's original spatial proximity distribution. We can do this by restricting the range of right ascension that the altered objects can take. Consider the spatial proximity function for a single object. If an event's proximity to this object is so small, it is considered to be far enough away to not be associated with the object.

As declination is not altered, the spatial proximity function becomes one dimensional. We can evaluate the right ascension of the function equal to the value of a proximity cut, and thus find the difference between the peak right ascension (α_0) and the displaced right ascension ($\Delta\alpha$):

$$\alpha_0 - \Delta\alpha < \alpha < \alpha_0 + \Delta\alpha\,. \tag{7.12}$$

From the above, a spatial exclusion zone has been produced, where an object in an alternative universe cannot lie.

The temporal displacement of the object is source-type dependent, due to the largely variant durations of the object. Blazars can only have a certain start time and duration.[6] For GRBs, as they are extremely short lived, random start times and durations are chosen within their expected range.

It is possible to produce infinite of these alternative universes, i.e., altering the spatial and temporal proximity distributions for each of the objects. We can essentially use this to build up a background, as we shouldn't expect objects of different spatiotemporal properties to give rise to events associated with the object in our Universe. Thus, each object was spatiotemporally displaced within 10,000 alternative universes. A proximity is calculated for each event to the object in question.

[6]The Fermi All-Sky Variability Analysis (FAVA) only reports blazar flares in terms of weeks, i.e., one flare could be active during week 1 and week 2, and so would be reported as flaring for 2 weeks.

For a single universe, the proximity for all events in the ANITA-4 dataset to a single object are summed, giving the summed proximity. Distributions of these summed proximities are produced, one for each object. Each of the summed proximity distributions is then evaluated for the associated object in our Universe, in order to see if there is an excess of events that point back to it.

7.3.4 Results

7.3.4.1 Primer

The analysis is carried out for objects of three types: blazars, gamma-ray bursts, and supernovae. This subsection explains how the test statistic plots are generated and how they can be interpreted.

To understand the test statistic (summed proximity) distributions, please refer to the previous section, which describes the analysis. A very brief outline of the analysis is as follows: the ANITA-4 events are subjected to a set of loose cuts to acquire our final candidates. These final events are then pointed back to their astrophysical origin, under the assumption that they are neutrinos. Now we consider only one object (such as a blazar). A "proximity" is calculated for each event, telling us how proximate it is to the object in space and time. This is then repeated for all the ANITA-4 events, and these are summed up to see how close all the events are to the object. This is what we call a "summed proximity". In order to form a background estimate, we displace the object in space and time, i.e., if the events were proximate to the object beforehand, they most likely no longer are. We say that this object belongs to an "alternative universe".[7] A "summed proximity distribution" is produced for each object. Each bin of this plot represents the summed proximity of the ANITA-4 events to the object in a specific alternative universe.

10,000 alternative universes (where objects are randomised in space and time) are simulated and the ANITA-4 events are pointed back to their origin. In order to

[7] An "alternative universe" can be thought of as follows. Consider our Universe: it has a vast amount of astrophysical objects within it. We know they each have a specific position in space and a temporal profile for when they are active. Now consider that all of these objects are no longer where we know them to be, and they each exhibit different temporal behaviour too. These objects with alternative spatial and temporal profiles constitute an "alternative universe".

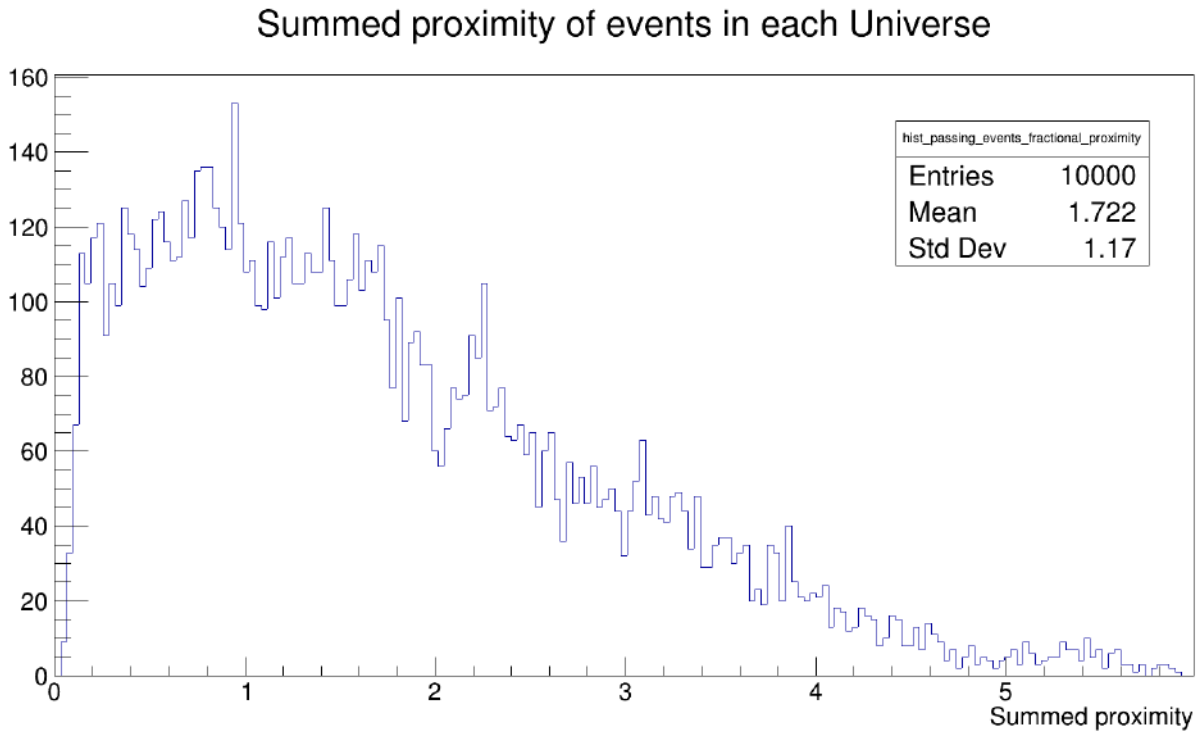

Figure 7.28: Long-lived source example: an example of a summed proximity distribution for events pointing back to a long-lived source (SN 2016iyl) for the set of alternative universes.

determine if the pointed-back events correspond to an object's spatial and temporal activity in *our* Universe, a summed proximity distribution (see Figure 7.28 and 7.29 for examples) is found for the object (either a blazar, supernova, or gamma-ray burst). If the ANITA-4 events point closely back to an object in a specific universe, it will have a high value of summed proximity (multiple events point back to it closely in time and space). If the ANITA-4 events are not proximate to the object in question in a specific universe, it will have a low value of summed proximity. Thus, if the ANITA-4 events actually point back to the object in our Universe, it should have a much higher value of the test statistic (summed proximity) than the mean. For long-lived sources, the summed proximity is our test statistic. For very short-lived sources, the logarithm of the summed proximity is our test statistic instead (as the temporal proximity is extremely low, this is a much better choice of test statistic).

The final candidates for evaluation are shown in the subsections below. The

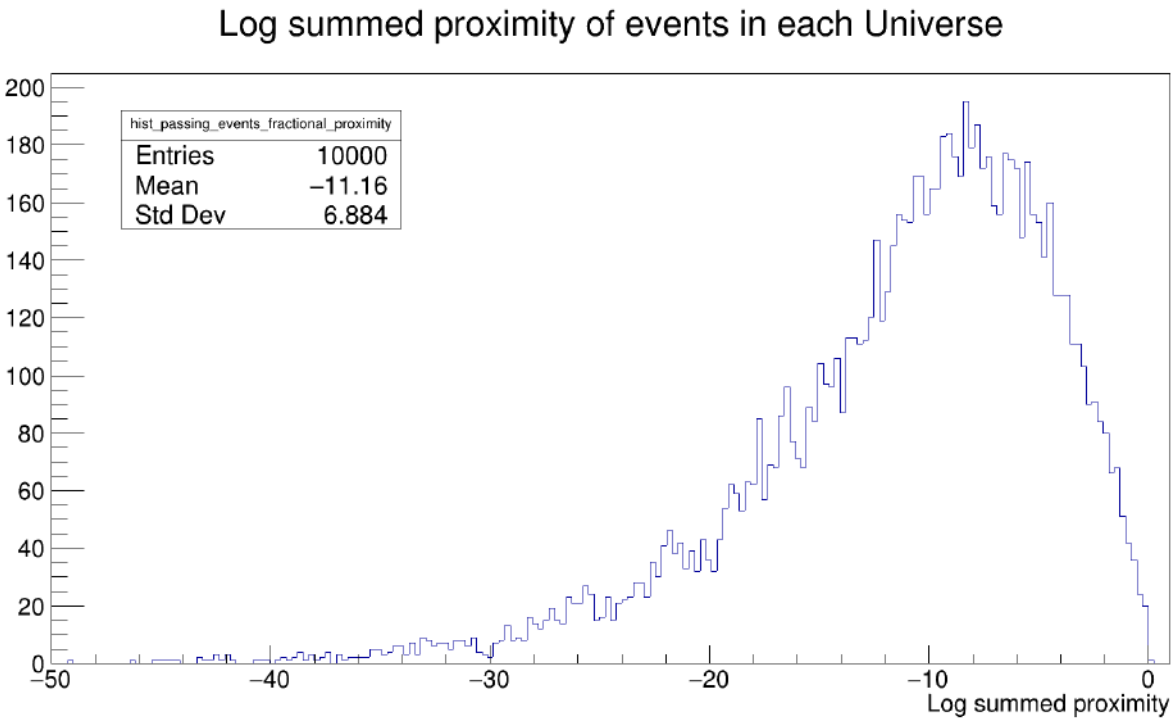

Figure 7.29: Short-lived source example: an example of the logarithmic summed proximity distribution for events pointing back to a typical transient source (GRB161219A) for the set of alternative universes.

cumulative distributions of each of the summed proximity distributions are found. If the cumulative distribution of the test statistic (summed proximity) exceeds the 95% confidence level (CL), i.e., the p-value for our Universe is < 0.05, the results are considered to be of potential significance and will undergo further analysis.

7.3.4.2 Blazar Results

Four blazars were active during ANITA-4 within a reasonable declination range, and enough neutrinos were found from the simulation so that their spatial proximity distributions could be generated. The results are found in Table 7.3.

One one object, PKS 1502+106, appeared to have an excess of events pointing back to it in our Universe (universe 0) in comparison to the vast majority of alternative universes (universes 1 to 10,000), which is discussed in Section 7.3.4.5. Both objects from the 4th Cambridge catalogue fell close to the average summed proximity (no significant excess observed at all), and CTA 102 fell just outside the region of minimum significance.

Name	$\overline{L}$	L_0	p-value	Significant?
PKS 1502+106	1.23	4.59	0.0024	Yes
4C +15.05	2.33	2.33	0.43	No
4C +01.02	0.85	1.11	0.27	No
CTA 102	1.49	3.55	0.052	No

Table 7.3: The set of final blazars, along with the test statistic (summed proximity), L, and the results. $\overline{L}$ represents the mean value of the test statistic for the alternative universes, and L_0 is the value of the test statistic in our Universe.

7.3.4.3 Supernova Results

Name	$\overline{L}$	L_0	p-value	Significant?
SN 2016iyl	1.72	1.67	0.44	No
SN 2016iwk	0.52	0.57	0.37	No
SN 2017hn	1.20	0.22	0.80	No
SN 2017ae	1.69	2.06	0.33	No
SN 2016ixb	0.54	0.32	0.58	No
SN 2016ivt	1.69	0.90	0.66	No
SN 2016ixe	0.14	0.15	0.23	No
SN 2016jdj	0.98	1.08	0.34	No
SN 2016jby	1.78	3.58	0.09	No

Table 7.4: The set of final supernovae, along with the test statistic (summed proximity), L, and the results. $\overline{L}$ represents the mean value of the test statistic for the alternative universes, and L_0 is the value of the test statistic in our Universe.

Ten supernovae passed both the declination cuts and the requirements to form proximity distributions. All of these final object candidates fell within the expected range, and none exceeded the 95% confidence level. The results are summarised in Table 7.4. An example summed proximity distribution for a supernova is shown in Figure 7.28.

7.3.4.4 Gamma-Ray Burst Results

Due to the very short durations of gamma-ray bursts, the temporal proximity will be very low. Thus, the summed (spatiotemporal) proximity of such objects will be far smaller than the previously discussed objects. For this reason, the logarithmic summed proximity, L, is used as a test statistic, as opposed to the summed proximity. An example logarithmic summed proximity distribution for a GRB is shown in Figure 7.29.

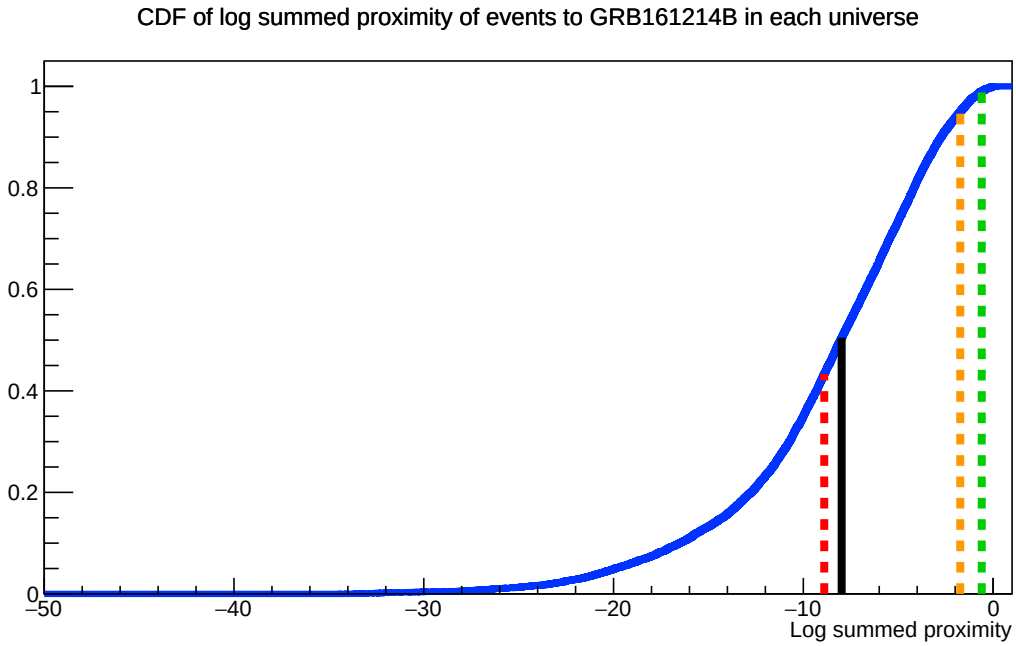

Figure 7.30: CDF of the logarithmic summed proximity distribution for events pointing back to GRB161214B for the set of alternative universes. Our Universe is represented by the black line, the mean of all alternative universes is represented by the red dashed line, the first significance threshold (CL 95%) is represented by the orange dashed line, and the second significance threshold (CL 99%) is represented by the green dashed line. This GRB is well within expected limits, and is just higher than the mean.

Name	$\overline{L}$	L_0	p-value	Significant?
GRB161219A	-11.16	-13.13	0.69	No
GRB161214B	-8.90	-7.98	0.50	No
GRB161214A	-8.84	-4.82	0.24	No
GRB161207813	-10.62	-3.25	0.10	No
GRB161207224	-21.90	-19.21	0.43	No

Table 7.5: The set of final gamma-ray bursts, along with the test statistic (logarithm of the summed proximity), L, and the results.

The analysis was repeated for each gamma-ray burst. In each case, our Universe was finally evaluated against the cumulative distribution function (CDF) for the test statistic in each alternative universe (see Figure 7.30 for GRB161214B). There were no excesses of events pointing back to GRBs in our Universe, indicating that there was no neutrino-based activity from GRBs during the flight of ANITA-4. Table 7.5 contains the results.

7.3.4.5 PKS 1502+106

The only object to have a significant excess of events pointing back to it in our Universe was PKS 1502+106 (see Figure 7.31). PKS 1502+106 is a highly energetic flaring blazar located at (α, δ) = (226.21°, 10.31°), and was active during two observation weeks.

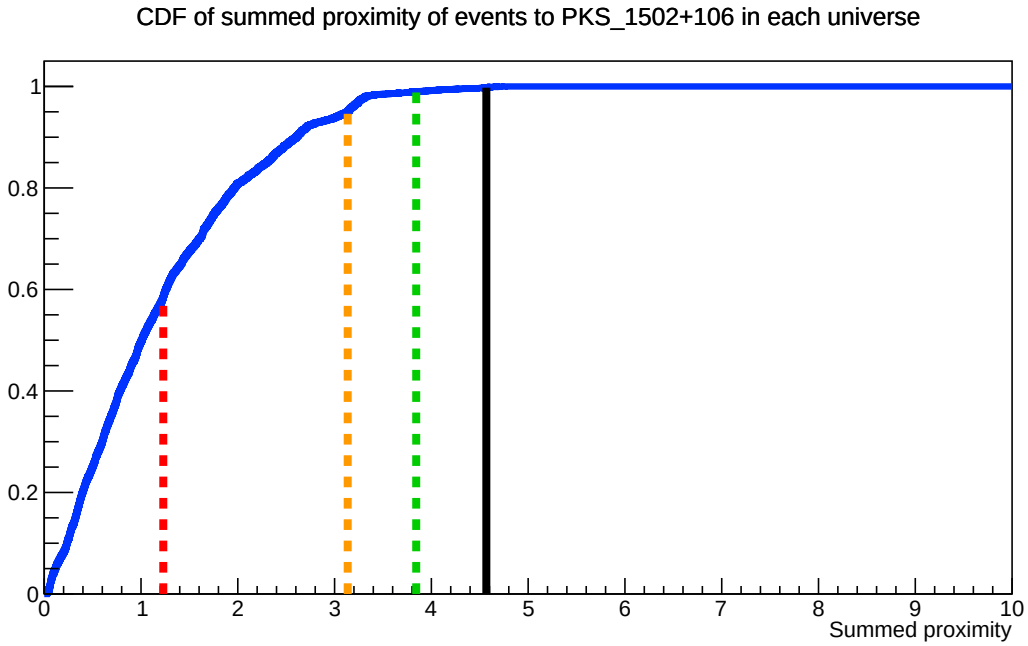

Figure 7.31: CDF of the summed proximity distribution for events pointing back to PKS 1502+106 for the set of alternative universes. Our Universe is represented by the black line, the mean of all alternative universes is represented by the red dashed line, the first significance threshold (CL 95%) is represented by the orange dashed line, and the second significance threshold (CL 99%) is represented by the green dashed line. This flaring blazar exceeds both limits of the confidence level.

According to the Fermi All-Sky Variability Analysis (FAVA), PKS 1502+106 (3FGL J1504.4+1029) was associated with two weekly flares (flare 37 in week 435 and flare 25 in week 436). The first was associated with week 435 (28 Nov to 05 Dec 2016); for ANITA, this was during pre-launch and the early stages of the flight. The second was associated with week 436 (05 to 12 Dec 2016), which was during early-mid flight. Table 7.6 summarises the light curves of the flare during both weeks, and Figures 7.32 and 7.33 display their relative fluxes. Figure 7.34 shows the FAVA "TS" (test statistic) maps of the initial flare. The FAVA TS maps are

the output of a series of likelihood tests to determine the significance of the flares. For more information, see Reference [101]. These maps are displayed in galactic co-ordinates, instead of right ascension (α) and declination (δ). One can find the galactic latitude (b) and galactic longitude (l) as follows,

$$b = \sin^{-1}\Big(\cos(\delta)\cos(\delta_G)\cos(\alpha-\alpha_G)+\sin(\delta)\sin(\delta_G)\Big), \tag{7.13}$$

$$l = l_C - \cos^{-1}\left(\frac{\sin(\delta)\cos(\delta_G)-\cos(\delta)\sin(\delta_G)\cos(\alpha-\alpha_G)}{\cos(b)}\right), \tag{7.14}$$

where α_G and δ_G are the right ascension and declination of the north galactic pole, and l_C is the galactic longitude of the north celestial pole. All calculated terms are scaled by $180°/\pi$ so that they are in degrees. The values used are $\alpha_G = 192.85°$, $\delta_G = 27.13°$, and $l_C = 122.93°$.

Week	Flare number	RA (deg)	Dec (deg)	Expected events	Observed events	Sigma
435	37	226.21	10.31	14.8	43.0	6.07
436	25	225.19	10.49	7.72	22.0	4.36

Table 7.6: The two flares observed by the Fermi All-Sky Variability Analysis (FAVA) for PKS 1502+106 (3FGL J1504.4+1029) during the ANITA-4 flight. The expected and observed events correspond to the expected and measured high energy gamma ray events from the direction of the above right ascensions and declinations. The expected events correspond to the expected number of gamma-rays, i.e., the background, whereas the observed events number corresponds to the observed number of gamma-rays, i.e., the signal. Sigma corresponds to the the significance of the flux variation. The corresponding information was calculated by using the FAVA analysis tools [101].

As our test statistic is a sum, we can look at the individual ANITA-4 events contributing to it (see Figure 7.35). 74% of the summed spatiotemporal proximity came from the 6 most proximate events (out of a total of 917 events). A selection of the top events was analysed. The waveforms were examined and the radio-exit positions of the ice were mapped for each event.

The radio-exit positions of the ANITA-4 events are shown in Figure 7.36. Many of them point back to the nearby mountain range and areas of relatively thin ice, and appear much like thermal noise. Recall that we purposely loosened the cut

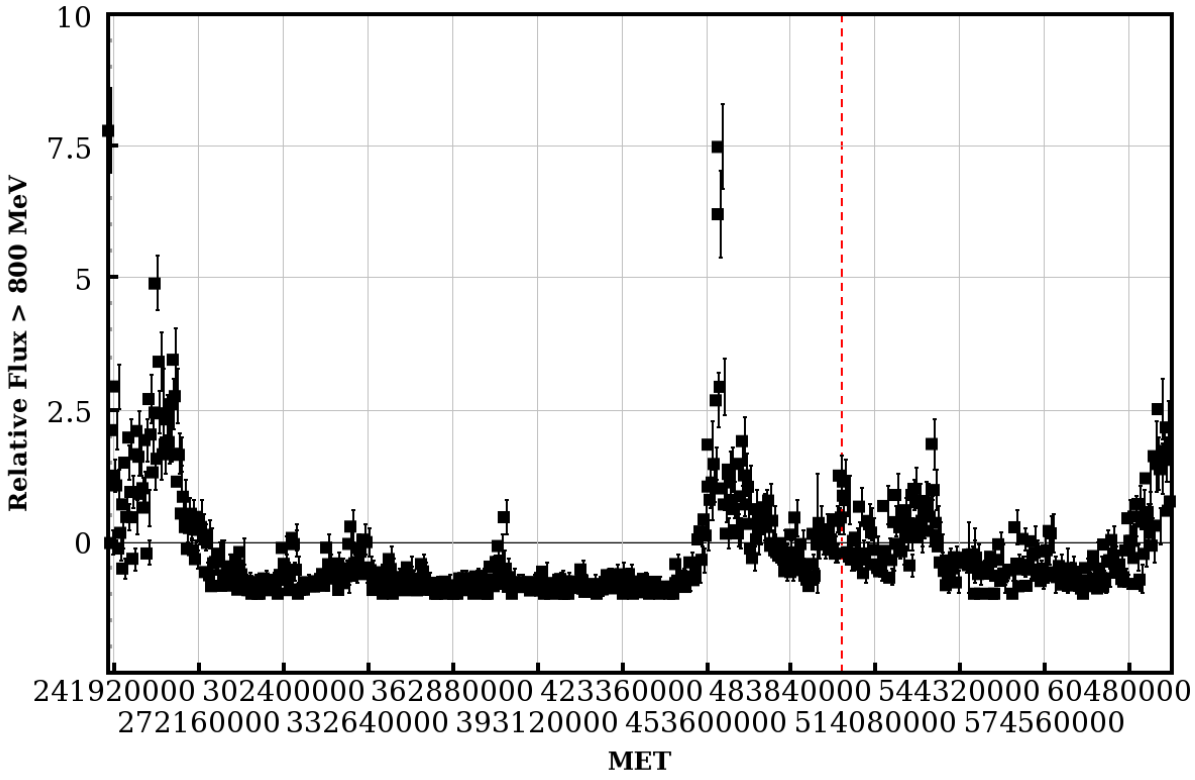

Figure 7.32: Fermi All-Sky Variability Analysis (FAVA) high energy light curve for the flare observed during the *first* week (in the direction of PKS 1502+106). The Mission Elapsed Time (MET) is the number of seconds elapsed since a base time (January 1st 2001). The red dashed line indicates the observation week. This plot was generated using the FAVA light curve generator [101].

to distinguish between thermal noise-like events and neutrino-like events, due to the nature of the analysis. Thus, these waveforms are most likely to be thermal in origin.

Event 15594227 (Figures 7.37 and 7.38) is of particular interest, as it points back to the blazar of interest, has an impulsive waveform, and is fairly isolated. However, its ice-exit position ("upper-most" event in Figure 7.36) is close to Halley Research Station and its airfield, Halley Skiway, indicating that it is likely to be an event of anthropogenic origin. Halley Research Station is a fairly isolated British base, and only a very small amount pre-clustered events point back within ~ 100 km of its location. This station was undergoing relocation during the ANITA-4 flight. [133] On 31st October 2016, a large crack in the Brunt Ice Shelf was discovered close to the station, forcing the crew to move Halley Research Station over 26 km

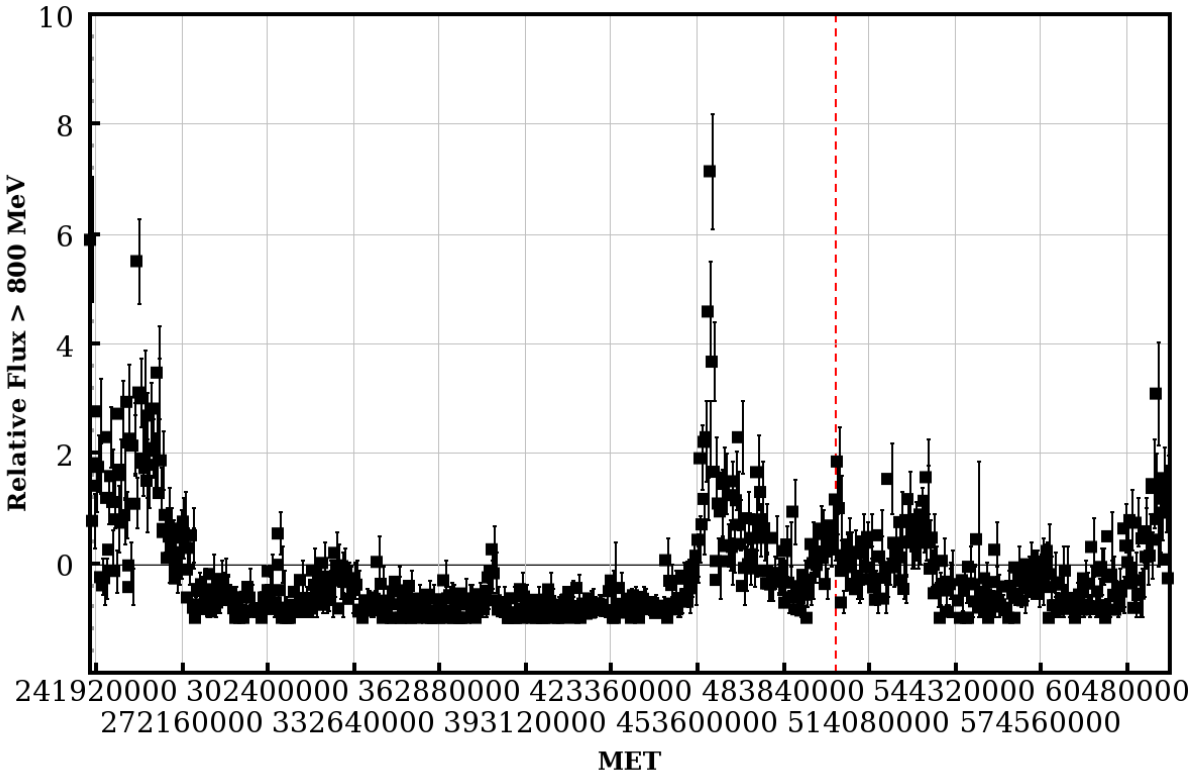

Figure 7.33: Fermi All-Sky Variability Analysis (FAVA) high energy light curve for the flare observed during the *second* week (in the direction of PKS 1502+106). As the measured right-ascension and declination of the localisation of the events differ slightly from week by week, the light curves also differ slightly. The Mission Elapsed Time (MET) is the number of seconds elapsed since a base time (January 1st 2001). The red dashed line indicates the observation week. This plot was generated using the FAVA light curve generator [101].

away. The relocation started a month later, and was mostly complete as of January 2017.

Due to the limitations of the analysis (in order to keep enough events so that they can pointed back to their origin, many cuts were loosened), this event was deemed to be isolated. Taking into account the location of event 15594227, which is close to a research station and points to a thin area of ice, it is unlikely that this event is associated with an UHE neutrino.

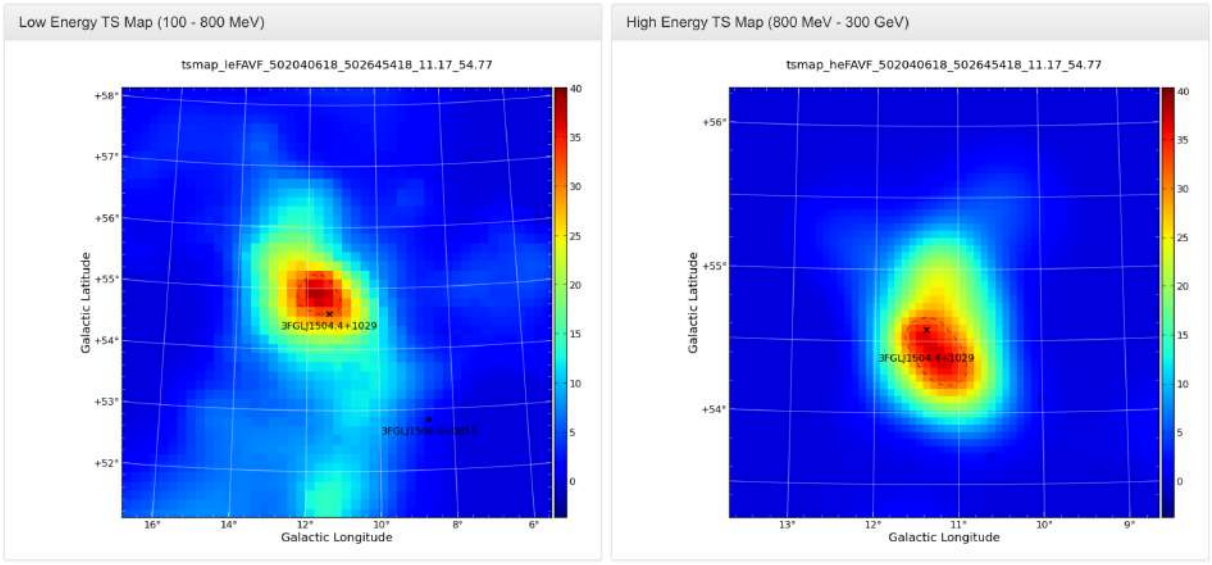

Figure 7.34: Fermi All-Sky Variability Analysis (FAVA) test statistic maps (left is the low-energy photon band, right is the high-energy photon band) for the first flare associated with PKS 1502+106. The corresponding map was generated using the FAVA analysis tools [101].

7.3.4.6 Results Summary

There were no significant excesses of events that pointed to any of the supernovae or gamma-ray bursts. There were also no excesses of events that pointed back to three of the four blazars. However, one particular blazar, PKS 1502+106, had a significant excess (> 99% CL) of ANITA-4 events pointing back to it, matching both its spatial and temporal profiles. The test statistic for events pointing back to this specific blazar ($L_0 = 4.59$) far exceeded the expected level ($\overline{L} = 1.23$).

The most proximate events pointing to PKS 1502+106 were examined. Many of these events appeared thermal in origin, showing very little impulsivity, if any. One event, however, event 15594227, showed a highly-impulsive vertically-polarised waveform. One would either expect it to be associated with an ultra-high energy neutrino, or an anthropogenic signal. It was found that, although isolated after all the cuts, the event pointed back to a research station that was undergoing relocation at the time. As the ice is particularly thin in this region, the chance of the observed radio wave originating from an ultra-high energy neutrino interacting with the ice is exceptionally low. Thus, this event is excluded on the basis that it is likely anthropgenic in origin. Conclusively, this analysis does not find any outstanding

excesses of neutrino-induced events that point back to the considered astrophysical objects.

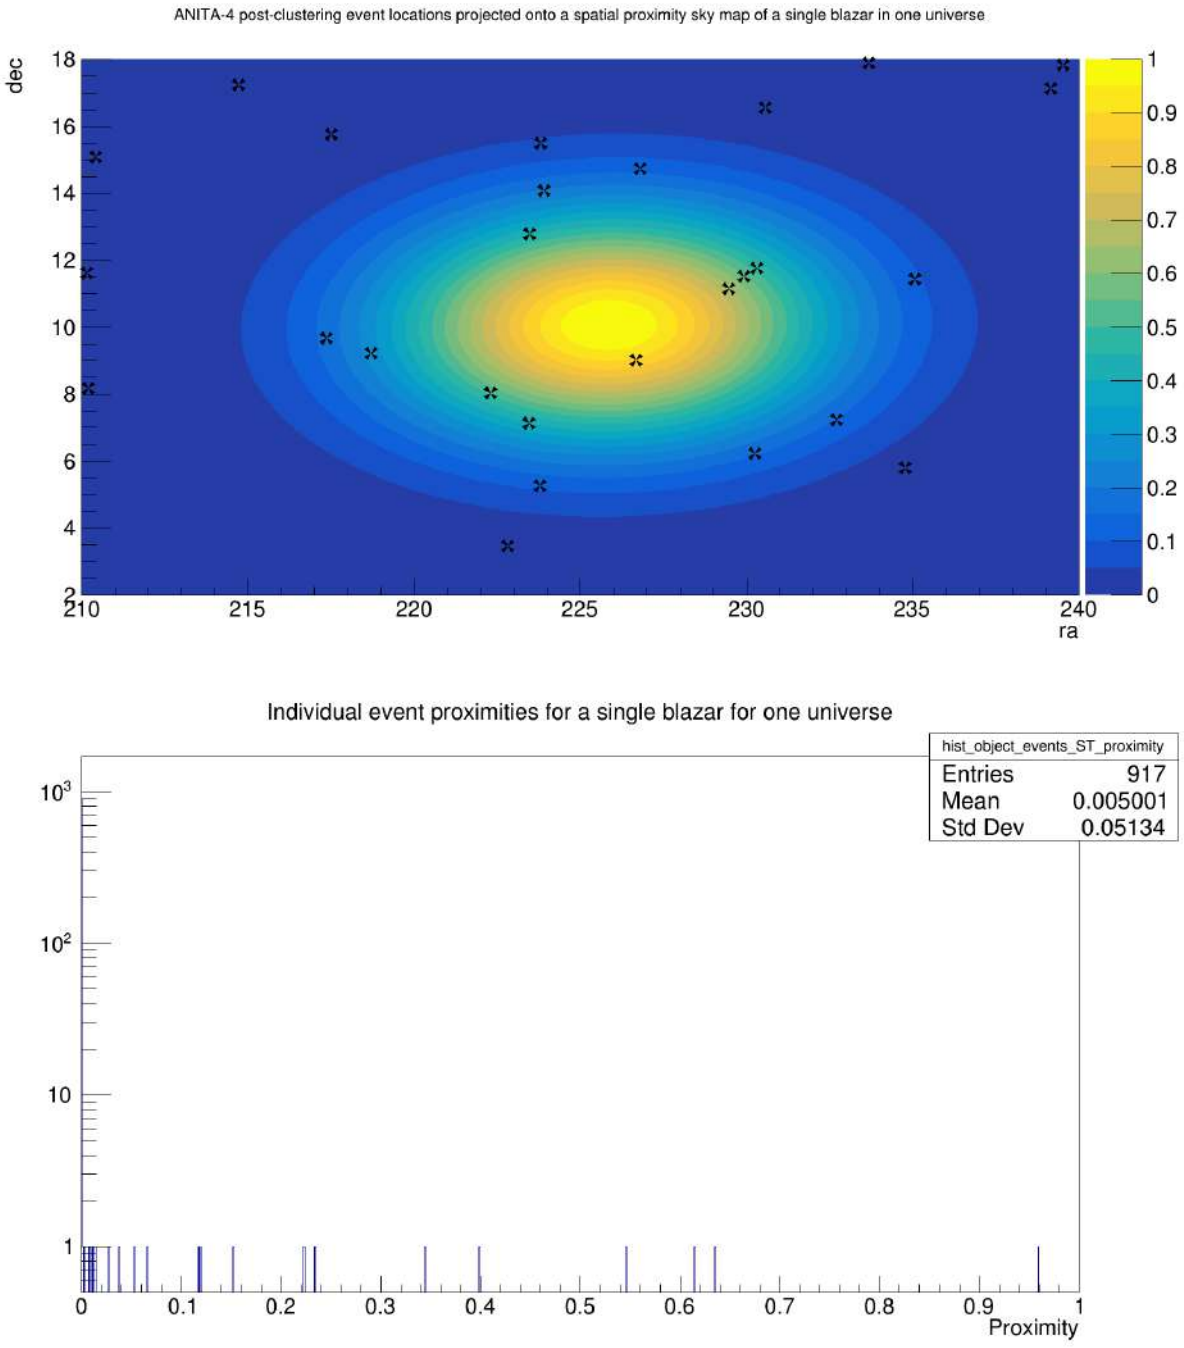

Figure 7.35: Above: events pointing back to the vicinity of PKS 1502+106 (in our Universe) on a spatial proximity graph. Right ascension (ra) and declination (dec) are always in degrees. Below: Histogram of individual event proximities to PKS 1502+106 in our Universe.

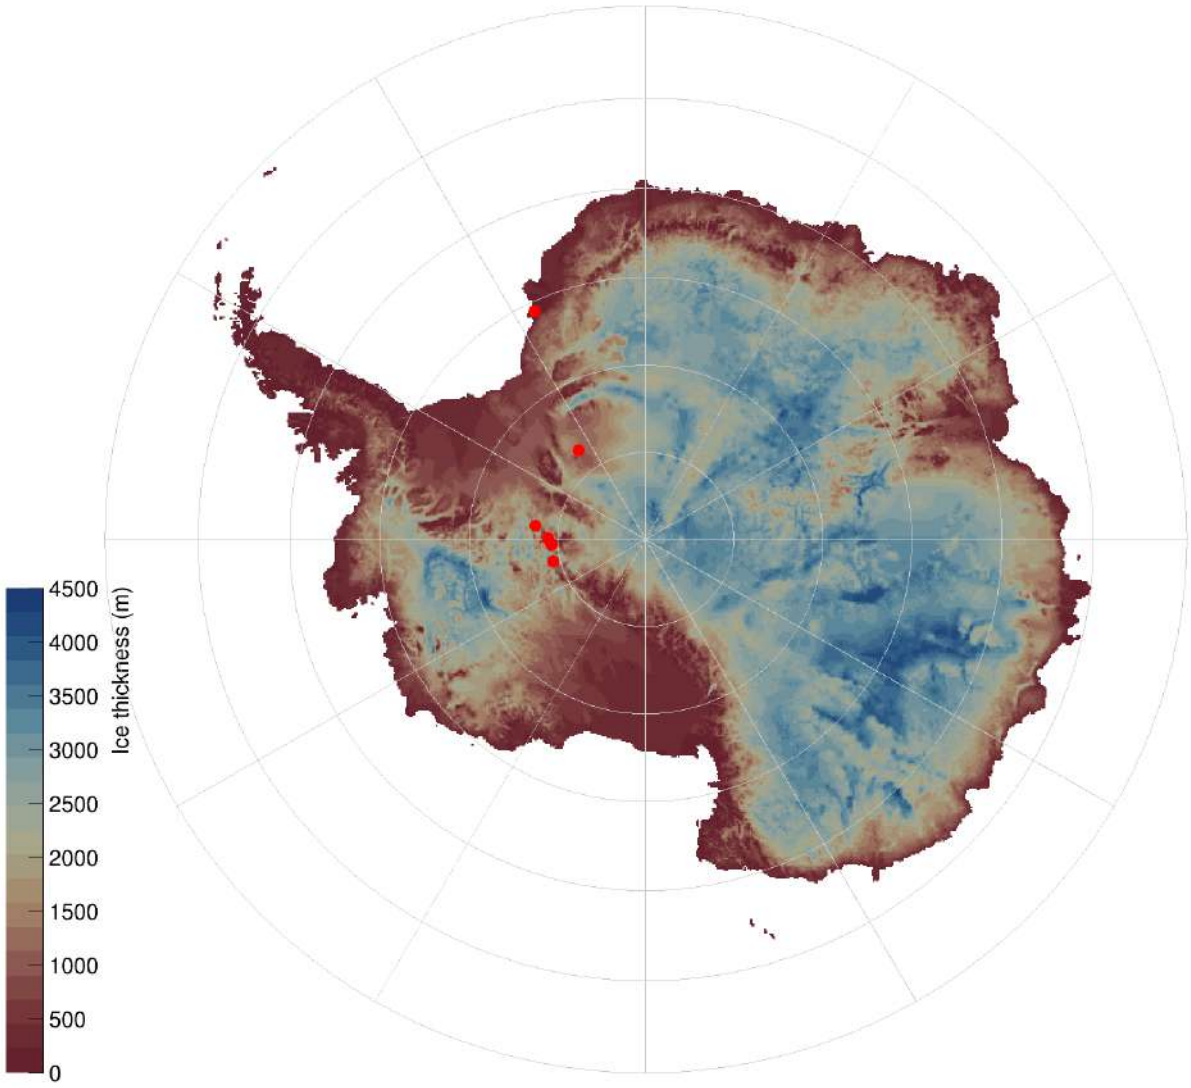

Figure 7.36: Ice radio-exit positions of the highest-priority events from the PKS 1502+106 analysis.

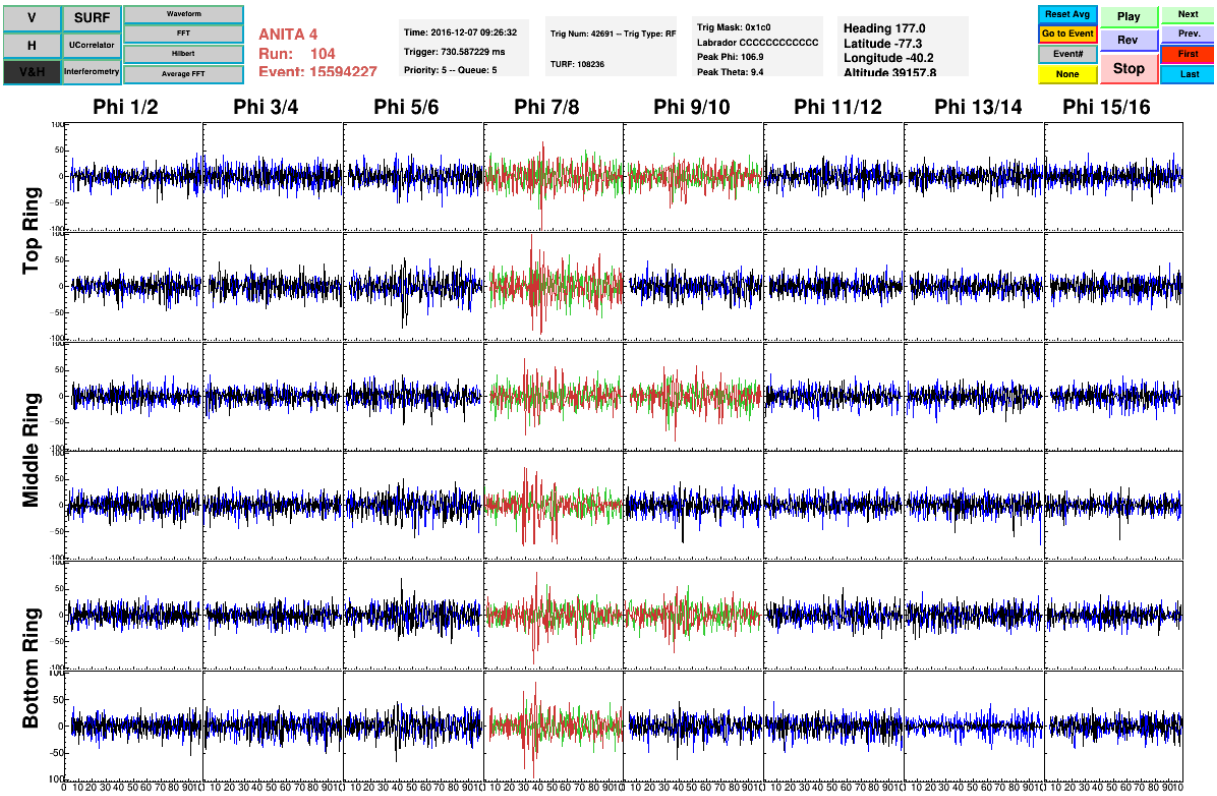

Figure 7.37: Vertically-polarised (red) and horizontally-polarised (green) waveforms for the triggered channels of event 15594227, an event from the PKS 1502+106 analysis.

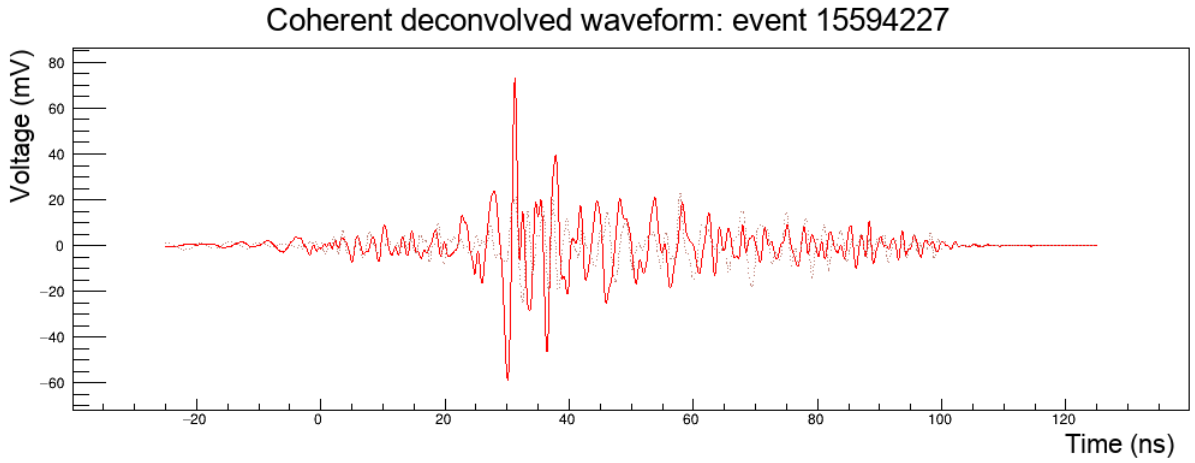

Figure 7.38: The vertically-polarised (red) coherently-summed deconvolved waveform of event 15594227, an event from the PKS 1502+106 analysis. The dashed line in the background corresponds to the cross-polarisation.

Chapter 8

Conclusions

The astrophysical origin of ultra-high energy neutrinos was searched for using data from the fourth flight of the ANITA experiment. Three classes of sources were used for this: blazars, supernovae, and gamma-ray bursts. A set of final events from the ANITA-4 dataset, under the assumption of neutrino origin, were pointed back to the sky to see if their spatial and temporal profiles matched known sources.

The combination of the simulation as described in Chapter 6 and the analysis performed in Chapter 7 provides a method for determining the origin of potential ultra-high energy neutrino-like events. This particular analysis did not find an excess of such events pointing back to any supernova or gamma-ray burst candidates. However, it tentatively found that the highly-energetic flaring blazar, PKS 1502+106, had a significant excess of events (p-value of 0.0024) pointing back to it. This flaring object was active during the ANITA-4 flight, and is within the ideal declination range that the instrument can observe. In addition, this blazar was recently associated with the IceCube event IC-190730A, a 300 TeV neutrino alert. Many of the sub-threshold events, however, appear to have similar profiles to thermal noise, and some of the events may be associated with anthropogenic activity. Even so, it is remarkable that these events pass both thermal and clustering cuts, and are strongly associated with both the spatial and temporal profile of a highly active blazar.

Current experiments are unable to consistently detect a high number of ultra-high energy neutrinos due to the extremely low cross-section of such particles. Fu-

ture experiments, particularly next-generation balloon-borne experiments such as PUEO, aim to be more sensitive to neutrinos of ultra-high energies. More frequent and longer flights would allow such experiments to search for these particles and link them with the reported activity of astrophysical objects. Future analyses could use the simulation and analysis techniques presented here as a means to identify the potential astrophysical origins of ultra-high energy neutrinos.

Appendix A

Cosmic ray flux spectra derivations

A.1 Second order Fermi acceleration

Consider the acceleration that arises by an outgoing charged particle (say, a proton) being reflected by the moving magnetic field of the interstellar medium (ISM). The proton moves at velocity, v, and collides with the ISM at an incoming angle, θ_{in}, which itself is moving at velocity, V, either head-on (the direction of the ISM velocity is opposite to that of the proton), or head-to-tail (the direction of the ISM velocity is the same of the proton).

A.1.1 Fractional energy gain

The probability of the is collision proportional to

$$P_{\text{head-on}} = v + V\cos(\theta_{\text{in}}), \tag{A.1}$$

for a head-on collision, and

$$P_{\text{head-to-tail}} = v - V\cos(\theta_{\text{in}}), \tag{A.2}$$

for a head-to-tail collision.

After the particle has entered the ISM, it will be reflected by the cloud at a random angle, θ_{out}, and exit the medium.

In the frame of the ISM (primed), the Lorentz transform for the incoming particle energy (E_{in}) is:

$$E'_{\rm in} = \gamma_{\rm ISM}\left(E_{\rm in} - p_{\rm in}V\cos(\theta_{\rm in})\right), \tag{A.3}$$

where $\gamma_{\rm ISM}$ is the Lorentz factor of the ISM (magnetic cloud), and $p_{\rm in}$ is the momentum of the particle. The galactic frame is denoted as unprimed. Rearranging, we get:

$$E'_{\rm in} = \gamma_{\rm ISM}E_{\rm in}\left(1 - \frac{p_{\rm in}V}{E_{\rm in}}\cos(\theta_{\rm in})\right). \tag{A.4}$$

Substituting in $E_{\rm in} = p_{\rm in}c$, where c is the speed of light, and the ratio of the cloud speed to that of light, $\beta_{\rm ISM} = \frac{V}{c}$:

$$E'_{\rm in} = \gamma_{\rm ISM}E_{\rm in}\left(1 - \beta_{\rm ISM}\cos(\theta_{\rm in})\right). \tag{A.5}$$

Inside the ISM, it is assumed that the particle does not lose or gain energy, $E'_{\rm in} = E'_{\rm out}$, and undergoes reflection, such that $p'_{\rm in} = -p'_{\rm out}$. Now, using the inverse Lorentz transformation for the outgoing particle's energy, $E_{\rm out}$:

$$E_{\rm out} = \gamma_{\rm ISM}E'_{\rm out}(1 + \beta_{\rm ISM}\cos(\theta'_{\rm out})). \tag{A.6}$$

Equations A.5 and A.6 can be combined to show:

$$E_{\rm out} = \gamma^2_{\rm ISM}E_{\rm in}(1 - \beta\cos(\theta_{\rm in}))(1 + \beta_{\rm ISM}\cos(\theta'_{\rm out})). \tag{A.7}$$

The fractional energy increase of the particle can be found by using $\gamma^2_{\rm ISM} = \frac{1}{1-\beta^2_{\rm ISM}}$:

$$\frac{\Delta E}{E} = \frac{E_{\rm out} - E_{\rm in}}{E_{\rm in}} = \frac{\beta^2_{\rm ISM} - \beta_{\rm ISM}\cos(\theta_{\rm in}) + \beta_{\rm ISM}\cos(\theta_{\rm out}) + \beta^2_{\rm ISM}\cos(\theta_{\rm in})\cos(\theta'_{\rm out})}{1 - \beta^2_{\rm ISM}}. \tag{A.8}$$

To acquire the *average* energy increase, we need to find the average over Equation A.8. As stated before, the particles are reflected in a random direction due to the tangled magnetic fields within the cloud. Therefore, many of the terms become

zero (as $\langle\cos(\theta_{\rm out})\rangle = 0$). However, for the incoming particle, the probability of an encounter will be proportional to the relative velocity between the particle and the cloud:

$$\langle\cos(\theta_{\rm in})\rangle = \frac{\int_{-1}^{1}\cos(\theta_{\rm in})(v - V\cos(\theta_{\rm in}))\mathrm{d}(\cos(\theta_{\rm in}))}{\int_{-1}^{1}(v - V\cos(\theta_{\rm in})\mathrm{d}(\cos(\theta{\rm in}))} = \frac{(\frac{-2V}{3})}{2v}\,. \tag{A.9}$$

As the particle is ultra-relativistic:

$$\langle\cos(\theta_{\rm in})\rangle = \frac{-V}{3c} = \frac{-\beta_{\rm ISM}}{3}\,. \tag{A.10}$$

By substituting Equation A.10 into the remaining terms of Equation A.8, and noting that $1-\beta_{\rm ISM}^2 \simeq 1$, we arrive at an average fractional energy gain of:

$$\frac{\Delta E}{E} = \frac{\beta_{\rm ISM}^2 + \beta_{\rm ISM}^2/3}{1-\beta_{\rm ISM}^2} \simeq \frac{4\beta_{\rm ISM}^2}{3}\,. \tag{A.11}$$

This shows that the fractional energy increase is proportional to the *square* of the interstellar cloud velocity. Hence, this is known as *second order* Fermi acceleration.

A.1.2 Spectrum

From this acceleration scheme, it should be possible to derive the spectrum of the accelerated protons. If second order Fermi acceleration is responsible for accelerating cosmic rays, the resulting spectrum will be a global power law in energy i.e., the exponent does not depend on any additional variables.

The average time between collisions will be $t_{\rm coll} = \frac{L}{c}$, where L is the mean free path. Thus, we can calculate the average rate of energy change to be,

$$\frac{dE}{dt} \simeq \frac{\Delta E}{t_{\rm coll}} = \frac{E}{t_{\rm coll}}\frac{\Delta E}{E} = \frac{c}{L}\frac{\Delta E}{E}E\,. \tag{A.12}$$

Using the fractional energy gain derived in the previous section, the average energy rate becomes,

$$\frac{dE}{dt} = \frac{4\beta_{\rm ISM}^2 c}{3L}E = \alpha E, \tag{A.13}$$

where α is $\frac{4\beta_{\mathrm{ISM}}^2 c}{3L}$.

Thus, the rate of energy increase upon a collision is proportional to the particle energy. Consider that the particles undergo convection (transport by bulk motion). In this case, we can use the advection equation with an additional term, which represents how many particles escape the acceleration region of the system:

$$\frac{dn(E)}{dt} = \frac{d}{dE}[b(E)n(E)] - \frac{n(E)}{T_{\mathrm{esc}}}, \tag{A.14}$$

where $n(E) = \frac{dN}{dE}$, $b(E)$ is the energy loss term, -$\frac{dE}{dt}$, and T_{esc} is the escape time.

Assuming that the particle flow is steady, the first term becomes zero. The energy gain was calculated in Equation A.13. It is now inserted as energy *loss*, so $\frac{dE}{dt} = -\alpha E$:

$$\frac{d}{dE}[\alpha E n(E)] = -\frac{n(E)}{T_{\mathrm{esc}}}. \tag{A.15}$$

Applying the product rule and rearranging,

$$\frac{dn(E)}{dE} = -\frac{n(E)}{E}\left(\frac{1}{\alpha T_{\mathrm{esc}}} + 1\right), \tag{A.16}$$

then solving the equation:

$$n(E) \propto E^{-\left(\frac{1}{\alpha T_{\mathrm{esc}}}+1\right)}. \tag{A.17}$$

We have arrived at the spectrum of particles from second-order Fermi acceleration. It correctly predicts a power law, however, the power law is not global, as the exponent itself depends on the exact configuation of the system. As such, different regions of space would accelerate particles with wildly different spectra, and the sum of the contributions would not return the power law as seen from observations. Thus, a different method of acceleration is needed.

A.2 First order Fermi acceleration

The previous acceleration scheme fails to predict the global power law spectrum for cosmic rays. Astrophysical shocks (from objects such as AGNs and SNe), which are

non-relativistic, could cause relativistic particles to bounce between the upstream (unshocked) and downstream (shocked) regions. In this case, they would gain energy in every collision, and potentially give rise to the power law we expect.

Assuming the gas is monatomic, between the upstream and downstream regions, certain quantities will be conserved. In the frame of the shock, continiuty conditions are applied the flow rate of mass,

$$\frac{dm_{\mathrm{u}}}{dt} = \frac{dm_{\mathrm{d}}}{dt}, \tag{A.18}$$

$$\rho_{\mathrm{u}} v_{\mathrm{u}} = \rho_{\mathrm{d}} v_{\mathrm{d}}, \tag{A.19}$$

where ρ is the density, v is the velocity. u and d represent the upstream and downstream regions of the shock, respectively.

The momentum relation is:

$$\rho_{\mathrm{u}} v_{\mathrm{u}}^2 = P_{\mathrm{d}} + \rho_{\mathrm{d}} v_{\mathrm{d}}^2, \tag{A.20}$$

where P is the pressure.

The energy relation is:

$$\frac{1}{2}\rho_{\mathrm{u}} v_{\mathrm{u}}^3 = v_{\mathrm{d}} P_{\mathrm{d}} + \frac{1}{2}\rho_{\mathrm{d}} v_{\mathrm{d}}^3 - \frac{3}{2} v_{\mathrm{d}} P_{\mathrm{d}}\,. \tag{A.21}$$

Cancelling out the downstream pressure from Equations A.20 and A.21:

$$\frac{\rho_{\mathrm{d}}^2}{\rho_{\mathrm{u}}^2} v_{\mathrm{u}}^3 - 5\frac{\rho_{\mathrm{d}}}{\rho_{\mathrm{u}}} v_{\mathrm{u}}^3 + 4\rho_{\mathrm{u}}^3 = 0\,. \tag{A.22}$$

Using Equation A.19, this rearranges to give two solutions. The trivial solution is:

$$v_{\mathrm{d}} = v_{\mathrm{u}}, \tag{A.23}$$

and the non-trivial solution is:

$$v_{\mathrm{d}} = \frac{v_{\mathrm{u}}}{4}\,. \tag{A.24}$$

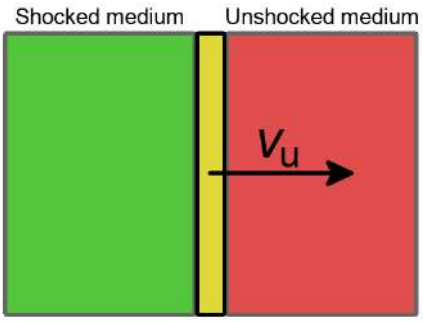

(a) Galactic (observer) frame.

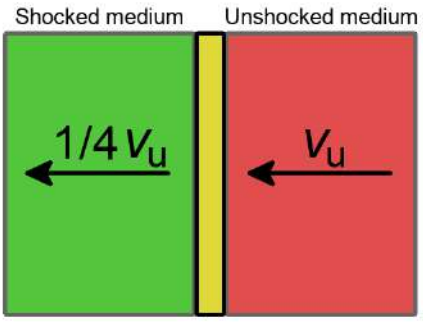

(b) Shock frame.

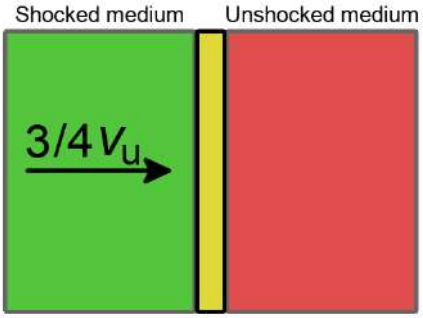

(c) Unshocked region frame.

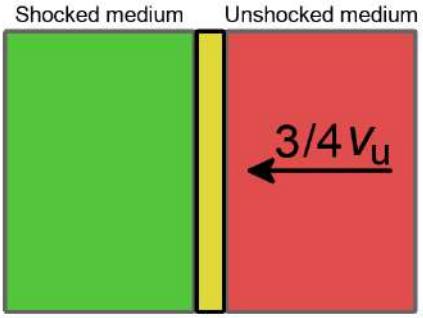

(d) Shocked region frame.

Figure A.1: Shock scenarios shown in the galactic frame, shock frame, unshocked region frame, and shocked region frame. The green area represents the shocked medium (particles that have been shocked, i.e., downstream). The yellow area represents the shock itself. The red area represents the unshocked region (particles that haven't been shocked yet, i.e., upstream).

This is illustrated in Figure A.1.

A.2.1 Fractional energy gain

A particle in the unshocked region will see the plasma approaching it at velocity $V = \frac{3v_u}{4}$ from the shocked region. The energy of the particle in the shocked region (E) will be:

$$E' = \gamma_S \left(E + pV\cos(\theta)\right), \tag{A.25}$$

where γ_S is the Lorentz factor of the shock, p is the particle momentum, θ is the angle of incidence between the particle and the shock.

Taking into account that the shock is non-relativistic,

$$\Delta E = E' - E = pV\cos(\theta), \tag{A.26}$$

and so

$$\frac{\Delta E}{E} = \frac{V}{c}\cos(\theta)\,. \tag{A.27}$$

The incoming particle will have a chance to cross the shock, with the probability being proportional to $\sin(\theta)\mathrm{d}\theta$. One also needs to take into account the rate at which they approach the shock front, which will be proportional to $\cos(\theta)$. Thus, the overall probability is:

$$P(\theta) = 2\sin(\theta)\cos(\theta)d\theta, \tag{A.28}$$

and so the *average* fractional energy gain is

$$\frac{\Delta E}{E} = \frac{V}{c}\int_0^{\frac{\pi}{2}} 2\sin(\theta)\cos^2(\theta)d\theta = \frac{2V}{3c} = \frac{2\beta}{3}\,. \tag{A.29}$$

This shows that the fractional energy increase is proportional to the shock velocity. Hence, this is known as *first order* Fermi acceleration.

A.2.2 Spectrum

We have a new scheme of acceleration that shows a different fractional energy increase, so we can attempt to derive the cosmic ray spectrum once more.

Before a collision, there are N_0 particles with energy E_0. After each collision, the average energy of the particle will increase ($E = \eta E_0$). However, the particle has a chance of escaping the acceleration region, with a probability ψ. Therefore, after k times passing the shock, there will be N particles,

$$N = N_0\psi^k, \tag{A.30}$$

with energy E,

$$E = E_0\eta^k\,. \tag{A.31}$$

The number of times crossing the shock can be cancelled out, giving:

$$\frac{N}{N_0} = \frac{E}{E_0}^{\frac{\ln(\psi)}{\ln(\eta)}}. \tag{A.32}$$

Taking the differential form:

$$n(E) \propto E^{-\gamma}, \tag{A.33}$$

where the exponent, γ, is $1 - \frac{\ln(\psi)}{\ln(\eta)}$. Note that here, γ only represents the exponent of the power law, *not* the Lorentz factor.

We already found the average fractional energy gain per collision, so the average fractional energy gain per cycle is:

$$\frac{\Delta E}{E} = \frac{E_0 - E}{E} = \frac{E_0}{E} - 1 = \frac{4\beta}{3}, \tag{A.34}$$

and so

$$\ln(\eta) = \ln\left(\frac{E}{E_0}\right) = \ln\left(1 + \frac{4\beta}{3}\right) \simeq \frac{v_\mathrm{u}}{c}. \tag{A.35}$$

The fraction of particles that are being swept away from the shock is simply $\psi = 1 - \frac{v_\mathrm{u}}{c}$, so,

$$\ln(\psi) = \ln\left(1 - \frac{v_\mathrm{u}}{c}\right) \simeq \frac{-v_\mathrm{u}}{c}. \tag{A.36}$$

Clearly,

$$\frac{\ln(\psi)}{\ln(\eta)} = -1. \tag{A.37}$$

Therefore, we find the exponent of the spectrum to be,

$$\gamma = 1 - \frac{\ln(\psi)}{\ln(\eta)} = 2, \tag{A.38}$$

and so,

$$n(E) \propto E^{-2}. \tag{A.39}$$

This spectral form clearly doesn't depend on any external parameters, and so the acceleration method is able to predict a global power-law spectrum.

One can go further by accounting for the component of particles which escape upstream and eventually propagate through the Galaxy to be detected on Earth. This further reduces the exponent of the derived cosmic ray spectrum, matching observations.

More generally, the cosmic ray spectrum is:

$$n(E) \propto E^{-\gamma}, \tag{A.40}$$

where γ is approximately 2.7.

Appendix B

Circuit diagrams

B.1 Antenna Testing

B.1.1 Boresight Tests

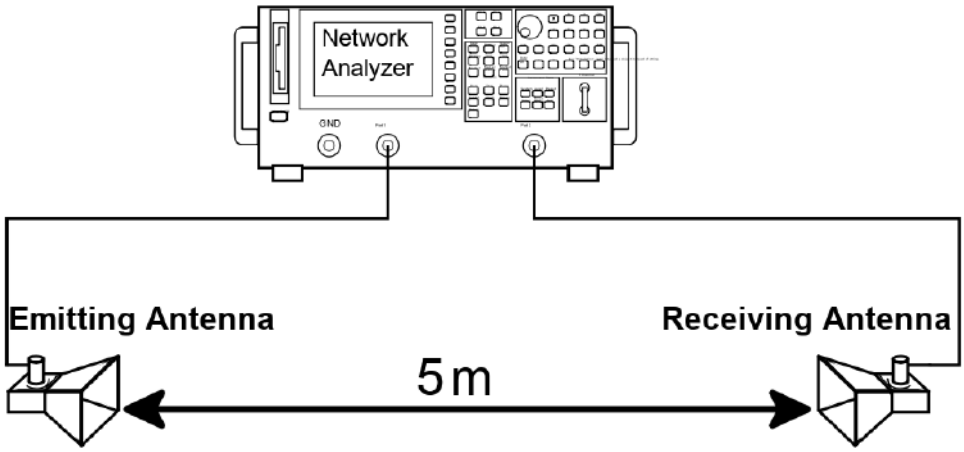

Figure B.1: Circuit diagram of boresight antenna tests. The antennas were pointed directly towards each other and their gains were calculated. This was done for all ANITA-4 antennas, as well as those from other flights.

B.1.2 Pulse Tests

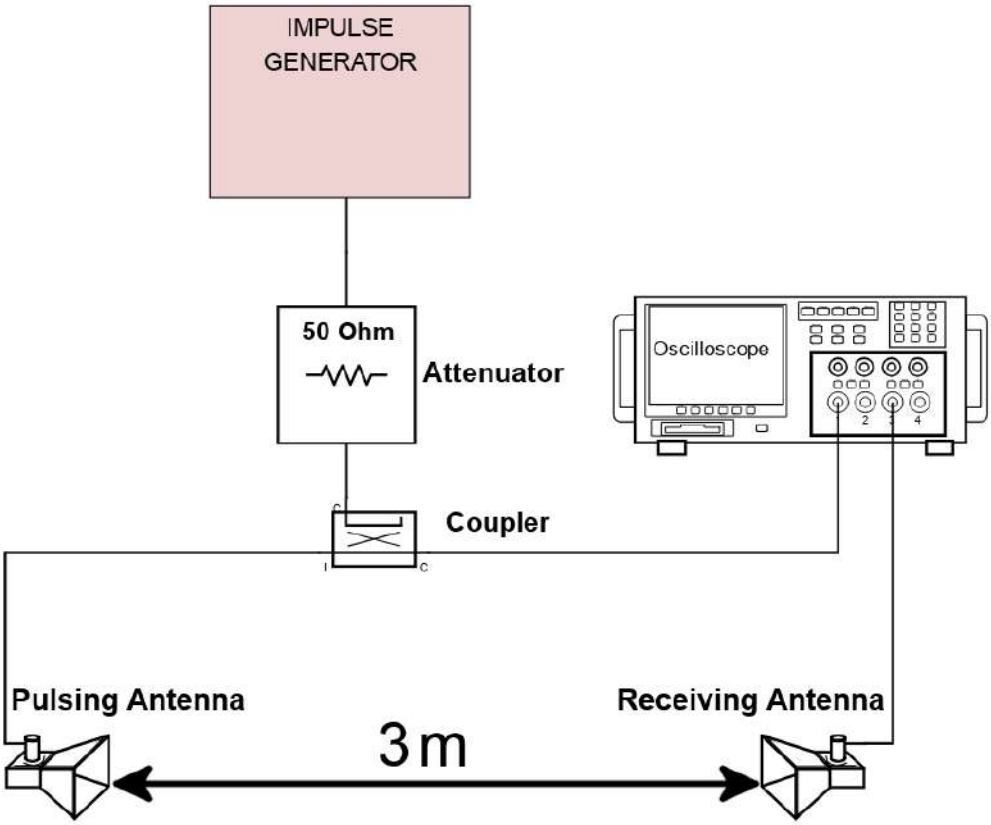

Figure B.2: A pulse generator was used to send a pulse from one antenna to another. In this case, the receiving antenna's oritentation was altered to see how the power depended on azimuth and elevation.

B.2 SHORT Board Testing

B.2.1 Basic SHORT Board Circuit

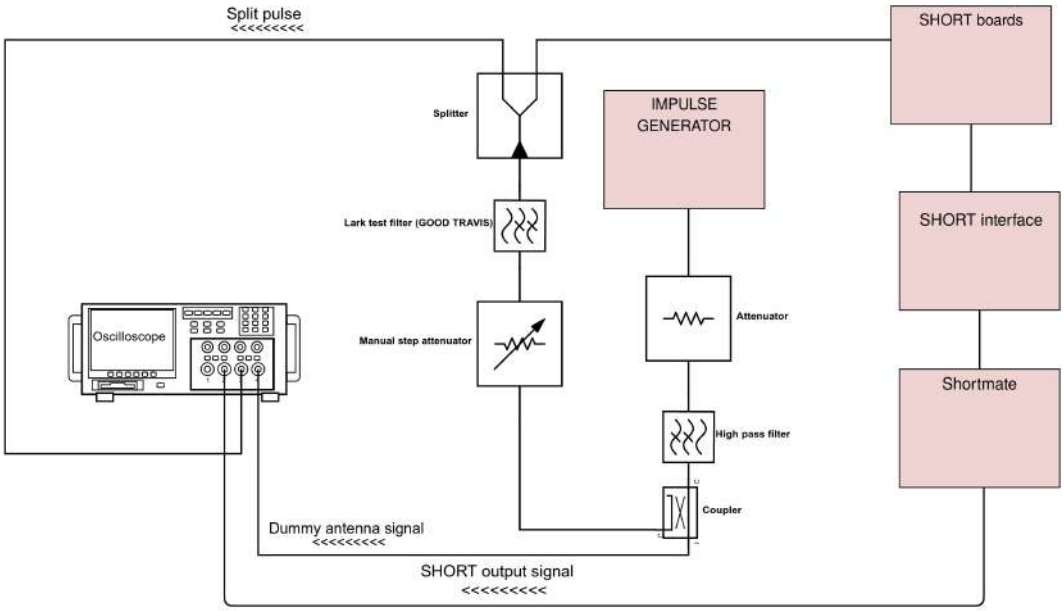

Figure B.3: Circuit diagram of all components of the basic SHORT setup. This setup allowed recording of the voltages output from the SHORTs. To replicate the signal and select certain frequencies, components such as attenutors and band-pass filters were used. Three signals were recorded: the dummy antenna signal, the split pulse, and the SHORT output signal.

B.2.2 SHORT Board Circuit with a Noise Source

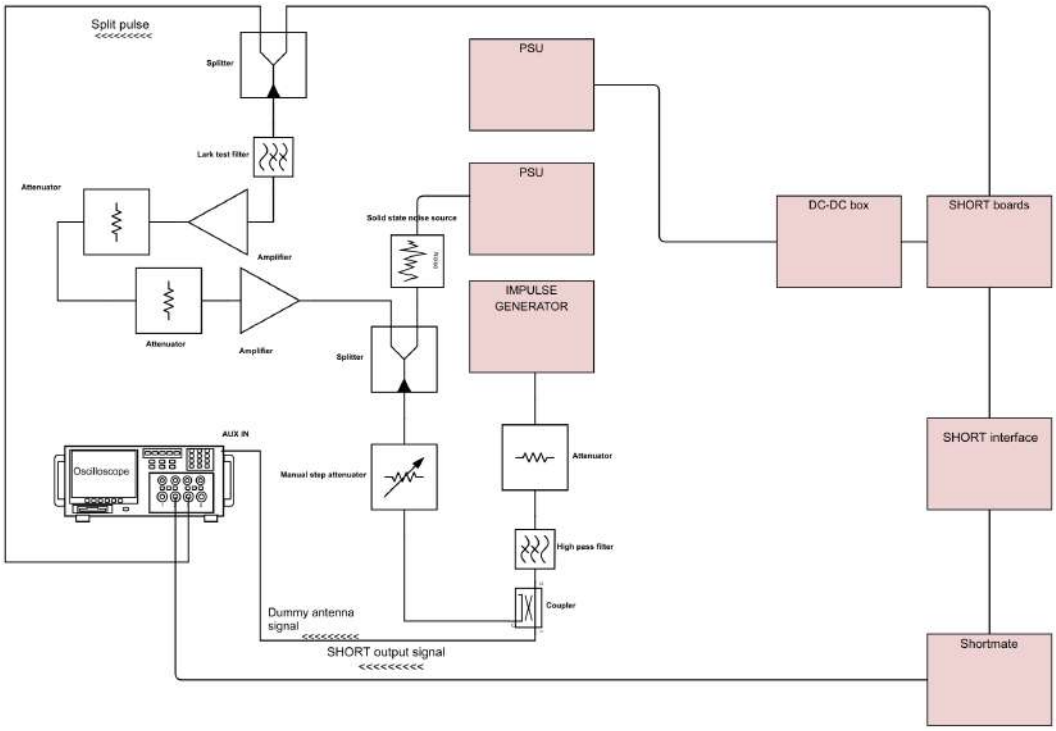

Figure B.4: Circuit diagram of noise addition to the SHORT setup. A solid state noise source was used to modify the pulse, and was used to check signal-to-noise ratios.

B.2.3 Initial Trigger Path Replication: SHORT Board Circuit with a Hybrid

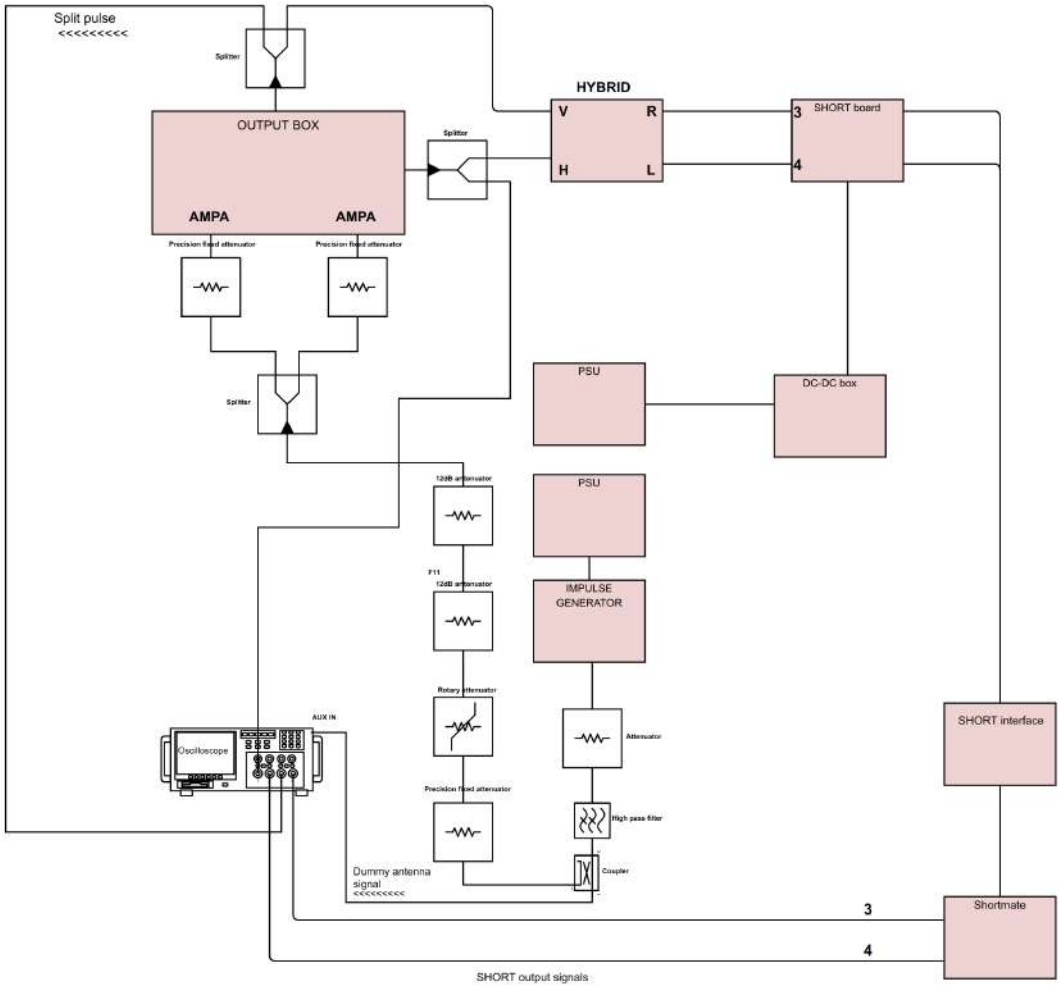

Figure B.5: Circuit diagram of the SHORT setup, including a hybrid. A pulse generator and the primary components are used to replicate the antenna signal. As per the actual ANITA trigger system, the waveform then moves through the AMPAs, a hybrid unit, and a SHORT board.

Appendix C

Break energy calculation

High energy neutrinos may arise from protons (p) interacting with photons (γ), producing Δ^+ baryons, providing that in the rest frame (denoted by primes) of the plasma:

$$E'_{\gamma,b}E'_p \simeq 0.3\ \mathrm{GeV}^2, \tag{C.1}$$

where E is the energy, and b represents the photon break energy.

In the process, the Delta baryon comes away with $0.2E'_p$.

$$E'_p = 0.3\left(\frac{E'_{\gamma,b}}{\mathrm{GeV}}\right)^{-1}\mathrm{GeV}, \tag{C.2}$$

and so,

$$E'_\Delta = 0.2E'_p = 0.06\left(\frac{E'_{\gamma,b}}{\mathrm{GeV}}\right)^{-1}\mathrm{GeV}\,. \tag{C.3}$$

The Delta baryon decays to protons and neutral pions or neutrons and charged pions. The charged pions then go through a decay chain, where the final product is one electron and three neutrinos. It is assumed that each of these final particles receive an equal share of the energy, so a single neutrino receives a quarter of the Delta baryon's energy:

$$E'_{\nu,b1} = 0.25E'_\Delta = 0.015\left(\frac{E'_{\gamma,b}}{\mathrm{GeV}}\right)^{-1}\mathrm{GeV}, \tag{C.4}$$

where $b1$ represents the first neutrino break energy.

And so in the lab frame:

$$E_{\nu,b1} = 0.015 \left(\frac{E_{\gamma,b}}{\mathrm{GeV}}\right)^{-1} \left(\frac{\Gamma_i}{1+z}\right)^2 \mathrm{GeV}, \tag{C.5}$$

where Γ is the Lorentz factor, and z is the redshift. The full neutrino flux spectrum is below:

$$\frac{dN_\gamma}{dE_\gamma} \propto \begin{cases} E_\gamma^{-1} & : E_\gamma < E_{\gamma,b} \\ E_\gamma^{-2} & : E_\gamma \geq E_{\gamma,b} \end{cases} \tag{C.6}$$

$$E_\nu^2 \frac{dN_\nu}{dE_\nu} \propto \begin{cases} E_\nu & : E_\nu < E_{\nu,b1} \\ \mathrm{const} & : E_{\nu,b2} > E_\nu \geq E_{\nu,b1} \\ E_\nu^{-2} & : E_\nu > E_{\nu,b2} \end{cases} \tag{C.7}$$

Bibliography

[1] J. Lindsey, Phys. Rev. Lett. **10** 146 (1963)

[2] M. Aker, *et al*. Phys. Rev. Lett. **123** 221802 (2019)

[3] M. Tanabashi, *et al*. Phys. Rev. D **98** 3 (2018)

[4] P. Higgs, Phys. Lett. **12** 132 (1964)

[5] P. Higgs, Phys. Rev. Lett. **13** 508 (1964)

[6] F. Abe, *et al*. Phys. Rev. Lett. **74** 2626 (1995)

[7] S. Abachi, *et al*. Phys. Rev. Lett. **74** 2632 (1995)

[8] K. Kodama, *et al*. Phys. Lett. B **504** 218 (2001)

[9] G. Aad, *et al*. Phys. Lett. B **716** 1 (2012)

[10] S. Chatrchyan, *et al*. Phys. Lett. B **716** 30 (2012)

[11] A. Becquerel, Compt. Rend. Math. **122** 420 (1896)

[12] J. Chadwick, Verh. Phys. Gesell. **16** 383 (1914)

[13] W. Pauli, Pauli Letter Collection (1930)

[14] N. Bohr, Niels Bohr Collected Works (1928)

[15] C. Ellis and N. Mott, Proc. Roy. Soc. **141** 502 (1933)

[16] E. Fermi, Z. Phys. **88** 161 (1934)

[17] G. Wang, Phys. Rev. **61** 97 (1942)

[18] C. Cowan and F. Reines, Phys. Rev. **90** 492 (1953)

[19] H. Crane and J. Halpern, Phys. Rev. **53** 789 (1938)

[20] C. Cowan and F. Reines, Phys. Rev. **92** 830 (1953)

[21] C. Cowan, *et al*. Science **124** 103 (1956)

[22] B. Pontecorvo, Zh. Eksp. Teor. Fiz. **34** 247 (1957)

[23] Z. Maki, M. Nakagawa, and Shoi. Sakata, Prog. Theor. Phys. **28** 870 (1962)

[24] R. Davis, D. Harmer, and K. Hoffman, Phys. Rev. Lett. **20** 1205 (1968)

[25] B. Cleveland, *et al*. Astrophys. J. **496** 505 (1998)

[26] D. Guenther, *et al*. Astrophys. J. **387** 372 (1992)

[27] Q. Ahmad, *et al*. Phys. Rev. Lett. **87** 071301 (2001)

[28] Y. Fukuda, *et al*. Phys. Rev. Lett. **81** 1562 (1998)

[29] D. Michael, *et al*. Phys. Rev. Lett. **97** 191801 (2006)

[30] M. Ahn, *et al*. Phys. Rev. Lett. **90** 041801 (2003)

[31] K. Abe, *et al*. Phys. Rev. Lett. **107** 041801 (2011)

[32] P. Adamson, *et al*. Phys. Rev. Lett. **116** 151806 (2016)

[33] F. An, *et al*. Phys. Rev. Lett. **108** 171803 (2012)

[34] J. Ahn, *et al*. Phys. Rev. Lett. **108** 191802 (2012)

[35] M. Apollonio, *et al*. Phys. Lett. B **466** 415 (1999)

[36] L. Evans and P. Bryant, J. Instrum. **3** S08001 (2008)

[37] V. Hess, Phys. Z. **13** 1084 (1912)

[38] L. Watson, *et al*. Mon. Not. Roy. Astron. Soc. **418** 206 (2011)

[39] E. Fermi, Phys. Rev. **75** 1169 (1949)

[40] E. Fermi, Astrophys. J. **119** 1 (1954)

[41] K. Olive *et al*. Chin. Phys. C **38** 090001 (2014)

[42] A. Hillas, *et al*. Ann. Rev. Astron. Astrophys. **22** 425 (1984)

[43] A. Bell, *et al*. Mon. Not. Roy. Astron. Soc. **431** 415 (2013)

[44] R. Aloisio, *et al*. Astropart. Phys. **27** 76 (2007)

[45] J. Heinze, *et al*. Astrophys. J. **825** 2 122 (2016)

[46] K. Greisen, Phys. Rev. Lett. **16** 748 (1966)

[47] G. Zatsepin and V. Kuzmin, J. Exp. Theor. Phys. Lett. **4** 78 (1966)

[48] V. Berezinsky and G. Zatsepin, Phys. Lett. B **28** 423 (1969)

[49] A. Castellina, *et al*. PoS ICRC2019 004 (2020)

[50] V. Berezinsky, *et al*. Phys. Rev. D **74** 043005 (2006)

[51] M. Vietri, Astrophys. J. **453** 883 (1995)

[52] E. Berti, *et al*. Phys. Lett. **114** 251103 (2015)

[53] V. Berezinsky, *et al*. Phys. Rev. Lett. **79** 4302 (1997)

[54] R. Blandford, *et al*. Ann. Rev. Astron. Astrophys. **57** 467 (2019)

[55] C. Kouveliotou, *et al*. Astrophys. J. **413** L101 (1993)

[56] R. Batista, *et al*. Front. Astron. Space Sci. **6** 23 (2019)

[57] X. Wang, *et al*. Phys. Rev. D **76** 083009 (2007)

[58] S. Chakraborti, *et al*. Nature Commun. **2** 175 (2011)

[59] R. Liu and X. Wang, Astrophys. J. **746** 40 (2012)

[60] B. Paczynski, Astrophys. J. **308** L43 (1986)

[61] J. Goodman, Astrophys. J. **308** L47 (1986)

[62] A. Shemi and T. Piran, Astrophys. J. **365** L55 (1990)

[63] B. Paczynski, Astrophys. J. **363** 218 (1990)

[64] R. Narayan, *et al*. Astrophys. J. **395** L83 (1992)

[65] M. Rees and P. Meszaros, Astrophys. J. **430** L93 (1994)

[66] B. Paczynski and G. Xu, Astrophys. J. **427** 708 (1994)

[67] M. Rees and P. Meszaros, Mon. Not. Roy. Astron. Soc. **258** 41 (1992)

[68] J. Clark and D. Eardley, Astrophys. J. **215** 311 (1977)

[69] G. Askaryan, J. Exp. Theor. Phys. **14** 441 (1962)

[70] G. Askaryan, J. Exp. Theor. Phys. **21** 658 (1965)

[71] P. Miocinovic, *et al*. eConf **C041213** 2516 (2004)

[72] D. Saltzberg, *et al*. Phys. Rev. Lett. **86** 2802 (2001)

[73] P. Gorham, *et al*. Phys. Rev. D **72** 023005 (2005)

[74] P. Gorham, *et al*. Phys. Rev. Lett. **99** 171101 (2007)

[75] P. Gorham, *et al*. Phys. Rev. Lett. **103** 051103 (2009)

[76] P. Gorham, *et al*. Phys. Rev. D **82** 022004 (2010), Erratum: Phys. Rev. D **85** 049901 (2012)

[77] P. Gorham, *et al*. Phys. Rev. D **98** 022001 (2018)

[78] P. Gorham, *et al*. Phys. Rev. Lett. D **121** 161102 (2018)

[79] P. Gorham, *et al*. Phys. Rev. **99** 122001 (2019)

[80] H. Friis, Proc. IRE. **34** 254 (1946)

[81] M. Lythe, *et al*. J. Geophys. Res. Solid Earth **106** 11335 (2001)

[82] C. Bassin, *et al*. Trans. Am. Geophys. Union. **81** F897 (2000)

[83] J. Ferringo, R. Williams, Jr., and A. Fox. U.S. Geological Survey Fact Sheet 01702 2 (2005)

[84] P. Allison, *et al*. Nucl. Instrum. Methods Phys. Res. A **894** 47 (2018)

[85] G. Varner, *et al*. ANITA Flight Firmware Document

[86] G. Varner, *et al*. Nucl. Instrum. Methods Phys. Res. A **583** 447 (2007)

[87] J. Russell, ANITA Internal Note 684 (2016)

[88] B. Strutt, PhD Thesis (2017)

[89] R. Brun and F. Rademakers, Nucl. Inst. Methods Phys. Res. A **389** 81 (1997)

[90] Archived Wembley Stadium Construction Document (2010)

[91] L. Cremonesi, *et al*. J. Instrum. **14** P08011 (2019)

[92] R. Batista, *et al.* J. Cosmol. Astropart. Phys. **1605** 038 (2016)

[93] C. Andreopoulos, *et al.* Nucl. Instr. Meth. Phys. Res. A **614** 87 (2010)

[94] Y. Hayato, Nucl. Phys. B Proc. Suppl. **112** 171 (2002)

[95] J. Alvarez-Muniz, *et al*. Phys. Rev. D **62** 063001 (2000)

[96] J. Alvarez-Muniz, *et al*. Phys. Rev. D **74** 023007 (2006)

[97] J. Alvarez-Muniz and E. Zas, Phys. Lett. B **411** 218 (1997)

[98] L. Landau and I. Pomeranchuk, Dokl. Akad. Nauk Ser. Fiz. **92** 535 (1953)

[99] L. Landau and I. Pomeranchuk, Dokl. Akad. Nauk Ser. Fiz. **92** 735 (1953)

[100] A. Migdal, Phys. Rev. **103** 1811 (1956)

[101] S. Abdollahi, *et al*. Astrophys. J. **846** 34 1 (2017)

[102] F. Acero, *et al*. Astrophys. J. Suppl. **218** 23 2 (2015)

[103] P. Coppin, IceCube Collaboration. `GRBWeb`. `https://icecube.wisc.edu/~grbweb_public` (2020)

[104] J. Aguilar, *et al*. Proc. ICRC. **235** 8 (2011)

[105] A. Lien, *et al*. Astrophys. J. **829** 7 1 (2016)

[106] International Astronomical Union. Transient Name Server. `https://wis-tns.weizmann.ac.il/` (2020)

[107] J. Guillochon, *et al*. Astrophys. J. **835** 64 (2016)

[108] G. Ghirlanda, G. Ghisellini and C. Firmani, Mon. Not. R. Astron. Soc. **361** L10 (2005)

[109] D. Band, *et al*. Astrophys. J. **413** 281 (1993)

[110] E. Waxman and J. Bahcall, Phys. Rev. Lett. **78** 2292 (1997)

[111] E. Waxman and J. Bahcall, Phys. Rev. D **59** 023002 (1998)

[112] E. Waxman and J. Bahcall, Astrophys. J. **541** 707 (2000)

[113] A. Friedmann, Z. Phys. **10** 377 (1922)

[114] A. Friedmann, Z. Phys. **21** 326 (1924)

[115] N. Aghanim, *et al*. Planck 2018 results. VI. (2018)

[116] J. Rachen and P. Meszaros, Phys. Rev. D **58** 123005 (1998)

[117] D. Besson, *et al*. Astropart. Phys. **56** 376 (2007)

[118] I. Valino, *et al*. PoS ICRC2015 271 (2016)

[119] O. Kalashev, *et al*. Phys. Rev. D **66** 063004 (2002)

[120] H. Takami, *et al*. Astropart. Phys. **31** 201 (2009)

[121] M. Ahlers and F. Halzen, Phys. Rev. D **86** 083010 (2012)

[122] K. Kotera, *et al*. J. Cosmol. Astropart. Phys. **10** 013 (2010)

[123] M. Aartsen, *et al*. Phys. Rev. D **98** 062003 (2018)

[124] E. Zas, *et al*. PoS ICRC2017 972 (2018)

[125] A. Ludwig, PhD Thesis (2019)

[126] R Core Team, R Foundation for Statistical Computing (2019)

[127] W. Chang, *et al*. Shiny: Web Application Framework for R (2019)

[128] M. Ankerst, *et al*. SIGMOD Rec. **28** 2 (1999)

[129] M. Hahsler, M. Piekenbrock and D. Doran, J. Stat. Softw. **91** 1 (2019)

[130] C. Deaconu, ANITA Collaboration. `UCorrelator`. `https://github.com/anitaNeutrino/UCorrelator` (2020)

[131] P. Gorham, *et al*. Astropart. Phys. **32** 10 (2009)

[132] R. Nichol, ANITA Collaboration. `anitaEventCorrelator`. `https://github.com/anitaNeutrino/anitaEventCorrelator` (2020)

[133] British Antarctic Survey, Official Press Release (2016)