A Search for Cosmogenic Neutrinos with the Askaryan Radio Array

Jonathan Paul Davies
University College London

Submitted to University College London in fulfilment
of the requirements for the award of the
degree of **Doctor of Philosophy**

October 17, 2014

Declaration

I, Jonathan Paul Davies 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 thesis.

Jonathan Davies _____

Abstract

The Askaryan Radio Array (ARA) is a new experimental effort to develop an array of sub-detectors capable of measuring ultra-high energy neutrino-induced radio pulses in the Antarctic ice sheet. Each sub-detector is able to function as a stand alone neutrino detector, the first of which was installed during the 2011 austral summer. In the following two years a further 3 sub-detectors were installed with updated design and functionality, with more planned over the next few years.

This thesis describes an analysis of the data collected by the first ARA station and presents the results of a search for ultra-high energy neutrinos. No statistically significant evidence for neutrino-induced signals is observed, with no candidate neutrino events. A limit is placed on the flux of ultra-high energy neutrinos and extrapolated to the full 37 station ARA detector operated over a 5 year period.

Acknowledgements

First and foremost I would like to thank my supervisor, Ryan Nichol, for this thesis would not have been possible without his knowledge, help and encouragement. I would also like to thank Mark Lancaster who, along with Ryan, gave me the opportunity and encouragement to study at UCL.

The ARA collaboration provided me with many great collaborators that offered much encouragement and valuable help over the last 3 years, who were also great fun to work with. Special thanks to those I shared my polar exploits with, in particular Hagar Landsman for not only ensuring my return to the North with all my fingers, but also indulging my desire to drive every vehicle under the polar sun.

Thanks to the UCL HEP group, members former and present, many evenings have been lost putting the world to rights at the JB.

Away from Physics I would like to thank my family: my brother Mark, my mother Jane and father Steve for supporting me through the last few years.

Anna has been a constant source of inspiration, picking me up after a hard day and giving so much joy in my life. Here's to the Doctors of Hove.

Finally to my many dear friends, I look forwarding to seeing you all again now that this journey is complete.

Contents

List of Figures

List of Tables

Chapter 1.

Introduction

Over a hundred years ago Victor Hess pioneering work led to the discovery of cosmic rays - charged particles bombarding the Earth's atmosphere [1]. Hess measured that the amount of ionising radiation increases with altitude and deduced that this radiation came from outside the Earth's atmosphere, which had a shielding effect. The discovery of these particles was of fundamental importance in two ways: by observing in detail their interactions a 'zoo' of seemingly more fundamental particles were discovered, which led in turn to the birth of particle physics; furthermore astronomers had a new window other than light through which to observe the Universe.

For over 50 years cosmic rays have been observed with arrival energies in excess of 10^{18} eV, energies well beyond those which can be achieved in terrestrial particle accelerators. However, there is much that is still unknown about the mechanisms that produce them or their sources. At these energies both cosmic rays and photons suffer from horizon problems that limit their use as astrophysical messengers. Ultra-high energy (UHE) neutrinos do not suffer such horizon problems and their observation and study could lead to a new window into the the distant Universe and the possibility of studying particle physics at energies well beyond those produced in the laboratory.

Unlike cosmic rays neutrinos only interact weakly and are not bent by galactic and extra-galactic magnetic fields, meaning that they will point back to their sources. Since neutrino interactions with matter are weak in comparison with other astrophysical messengers they are able to travel cosmological distances uninhibited. The horizon effects that limit the range and flux of the highest energy cosmic rays are also expected to produce a flux of UHE neutrinos whilst in transit, but also at the

sources. Measurement of these neutrinos will provide an insight into some of the unknowns about cosmic ray production mechanisms and source distributions.

The age of neutrino astronomy is edging closer, but the technical challenges of detecting UHE neutrinos are considerable. The UHE cosmic ray flux decreases with energy and it is expected that the neutrino flux will mirror this. The low flux combined with the small interaction cross-sections requires enormous detector volumes.

The Askaryan Radio Array (ARA) is an experiment based at the South Pole, Antarctica designed to measure coherent Cherenkov radio emission from neutrino induced particle cascades in the polar ice sheet. This coherent radio emission, known as Askaryan radiation [2] [3], was measured in a series of experiments at SLAC in a range of dense dielectrics, including ice [4] [5] [6]. Attenuation lengths of radio waves in ice are of the order of 1km, an order of magnitude greater than optical signals, making the possibility of instrumenting large detector volumes a more affordable and achievable prospect.

ARA's current design is for an array of 37 sub-detectors, or stations, each capable of functioning as a stand-alone neutrino detector. Each station consists of a series of deep holes in the ice containing radio antennas. Data is recorded and trigger decisions made by custom electronics and computer for each station. A prototype station, dubbed the TestBed, was deployed in the 2010 Austral summer and collected data for 2 years autonomously. A further 3 stations with upgraded functionality were deployed in the 2012-2013 Austral summer, with further deployments planned over the coming years until ARA reaches its design goal.

This thesis will describe an analysis of data taken over the course of 2011-2012 with the TestBed detector, calibration efforts and development of the data acquisition systems for the new ARA stations.

Chapter 2.

Particle Physics

2.1. The Standard Model of particle physics

The Standard Model of particle physics describes the observed interactions of fundamental particles and has proved incredibly successful. It is, however, incomplete. For one it makes no attempt to include gravity. There are also a number of observed phenomena that the model is unable to explain. One such phenomenon is neutrino oscillations and the associated non-zero masses of the neutrinos.

The Standard Model is a $SU(2) \otimes U(1)$ gauge theory consisting of fermions (quarks and leptons), which are the constituents of matter, and fundamental forces (electromagnetic, weak nuclear and strong nuclear) which are mediated by force carrying bosons.

The fermions, which have spin $+1/2$, are grouped into generations which exhibit similar physical properties. Each fermion also has an associated anti-particle with the same mass but opposite charge.

2.1.1. The quarks

There are 6 quarks grouped into three generations. The quarks each carry fractional electric charge. The up (u), charm (c) and top (t) quarks carry $+2/3$ whilst the down (d), strange (s) and bottom (b) quarks carry -1/3 the charge of the electron.

They are arranged as follows:

$$\begin{pmatrix} u \\ d \end{pmatrix}, \begin{pmatrix} c \\ s \end{pmatrix}, \begin{pmatrix} t \\ b \end{pmatrix}. \tag{2.1}$$

In addition to electric charge, the quarks carry colour charge. Each quark can take on one of three colour charges named red, green and blue (and for anti-quarks anti-red, anti-green and anti-blue). Quarks are not found in isolation due to colour confinement, a property of quantum chromodynamics. When two coloured particles are moved apart, for example a quark and anit-quark ($q\ \overline{q}$), the strong force between them increases linearly with separation. The energy stored in this field will increase to the point where another $q\ \overline{q}$ pair can be formed. The net result is two pairs of $q\ \overline{q}$, each of which is colourless. Instead of existing in isolation quarks combine to form hadrons, colourless objects consisting of two or three quarks. Hadrons with three quarks are known as baryons, typical examples are the proton (u u d) and the neutron (u d d). Two quark hadrons are known as mesons, consisting of a quark and anti-quark, a typical example being the π^+ (u $\overline{\text{d}}$).

2.1.2. The leptons

The leptons are a family of fermions, similar to the quarks, however they do not carry colour charge. They are grouped into three generations analogous to the quarks:

$$\begin{pmatrix} e^- \\ \nu_{\text{e}} \end{pmatrix}, \begin{pmatrix} \mu^- \\ \nu_\mu \end{pmatrix}, \begin{pmatrix} \tau^- \\ \nu_\tau \end{pmatrix}. \tag{2.2}$$

The electron, muon and tau all carry an electric charge of -1. Each of the charged leptons has an associated neutrino, a particle with no electric charge. In the Standard Model neutrinos have no mass and therefore travel at the speed of light.

2.1.3. The bosons

In the Standard Model the interactions of fundamental particles are described through the exchange of particles with integer spin called bosons. Properties of the bosons influence the macroscopic characteristics of the forces. For example the mass of the bosons influences the effective range of the forces.

The Standard Model describes three of the four fundamental forces. Gravity is not described in the Standard Model due in part to its relatively small strength, making its effects negligible compared with other fundamental particle interactions.

The electromagnetic force is mediated through the exchange of a massless spin 1 particle called the photon, γ. The photon couples to particles with non-zero electric charge hence all fermions, other than neutrinos, interact electromagnetically. As the photon has no mass the effective range of the electromagnetic force is infinite. The weak nuclear force is mediated through the charged W^+, W^- and chargeless Z^0 bosons, which couple to weak isospin. All fermions have non-zero weak isospin and so they all interact via the weak force. The $W^\pm$ have masses of 80.45 GeV and the Z^0 has a mass of 91.2 GeV. Although the weak and electromagnetic forces have similar strength the relatively large mass of these bosons makes the force appear weak and short ranged. The strong nuclear force is mediated via 8 massless gluons, g, which couple to particles with colour charge. Since quarks are the only fermions with colour charge they are the only constituents of matter that feel the strong force. Although gluons are massless the properties of quantum chromodynamics mean that the strength of the force increases with separation, leading to the absence of isolated quarks (or coloured objects) in nature.

2.2. The neutrino

As neutrinos are electrically neutral and possess no colour charge they only interact via the weak force, making their detection more difficult than electrically charged particles such as the electron. The existence of the neutrino was first postulated in 1930 by Wolfgang Pauli, who named it the *neutron*, to maintain the conservation of energy and momentum in beta decays. In 1934 Enrico Fermi coined the name

neutrino, meaning little neutral one, after the discovery by James Chadwick of a heavier neutral particle known to this day as the neutron.

The absence of electrical and colour charge for neutrinos means that they are only able to interact via the exchange of the $W^{\pm}$ and Z^0 bosons. Example Feynman diagrams of charged current (CC), involving the exchange of $W^{\pm}$, and neutral current (NC), involving the exchange of Z^0, neutrino interactions are shown in Figure 2.1. By observing charged leptons or hadronic recoils produced in these interactions it is possible to detect neutrinos and study their properties.

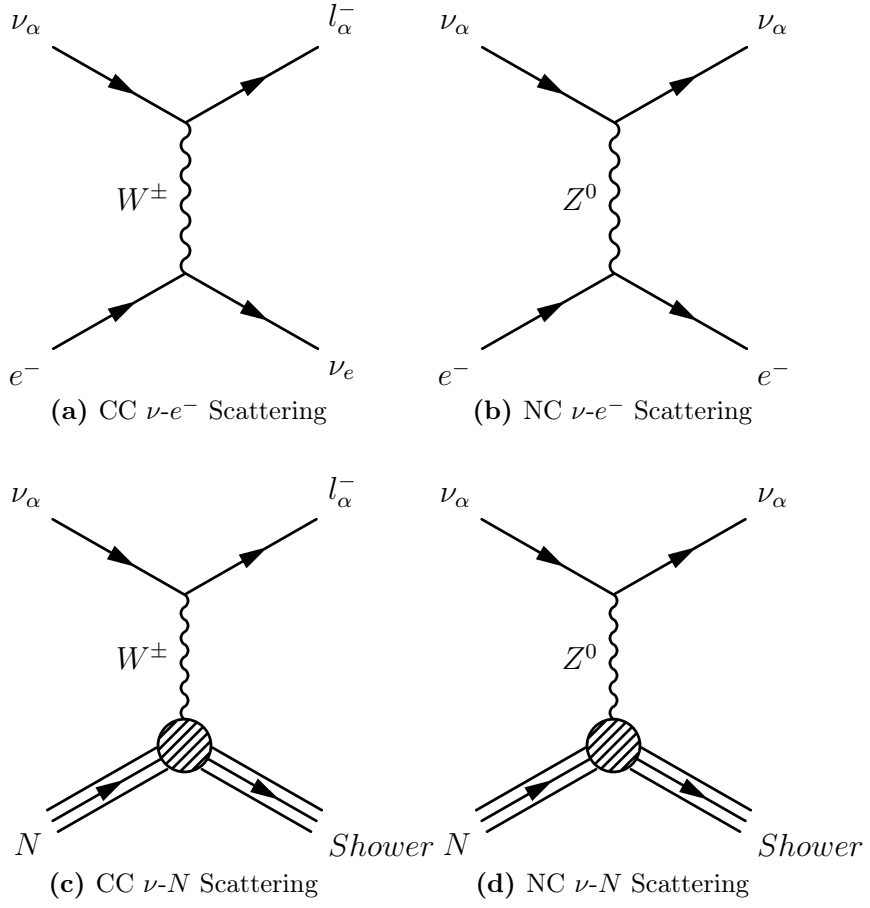

(a) CC ν-e^- Scattering **(b)** NC ν-e^- Scattering

(c) CC ν-N Scattering **(d)** NC ν-N Scattering

Figure 2.1.: Two charged current (CC) and two neutral current (NC) neutrino interactions via which it is possible to detect neutrinos.

In the Standard Model the weak force only couples to left (right) handed (anti-fermions) fermions, meaning that the neutrino, which has no electric or colour charge, can only exist as a left (right) handed particle (anti-particle). This also implies that

neutrinos must be massless in the Standard Model as an additional right handed neutrino would be required to generate mass.

Measurements of neutrino interactions and properties are challenging as these particles only interact via the weak force. Many of these properties are only just becoming accessible to experimental physicists. Observations of neutrino oscillations and non-zero neutrino mass have been made in recent years, both of which are inconsistent with the Standard Model description of neutrinos as massless particles.

2.2.1. Neutrino oscillations

The first experimental evidence for the phenomenon of neutrino oscillations came from the Homestake Experiment in 1968 [7]. The experiment made the first observation of neutrinos from the sun via CC interactions, but observed a deficit of solar ν_{e} compared to theoretical predictions based upon solar models. One possible solution to this problem was provided by neutrino oscillations, in which neutrinos oscillate between flavour states in transit from the sun. A ν_{e} produced in the sun can oscillate to a ν_{μ} or ν_{τ} in transit, both of which could not be observed in the Homestake Experiment since they have energy below the threshold to produce a muon or tau.

Subsequent experiments such as the Sudbury Neutrino Observatory (SNO) [8] and Super Kamiokande [9] provided measurements of the solar and atmospheric neutrino fluxes. In 1998 Super Kamiokande provided the first experimental evidence for atmospheric neutrino oscillations by making observations of the zenith angle dependance of their observed ν_{μ} and ν_{e} flux. Observations of neutrino CC and NC interactions (the latter being sensitive to ν_{e}, ν_{μ} and ν_{τ}) in SNO also provided compelling evidence for the oscillation hypothesis and measurements of the so called 'solar' neutrino oscillation parameters.

The current theoretical understanding of neutrino oscillations is that they are caused by the three known neutrino flavour states being a superposition of three mass states m_1, m_2 and m_3. The relationship between the weak and mass eigenstates is given in Equation (2.3).

$$|\nu_\alpha\rangle = \sum_{i=1,2,3} U_{\alpha i}|\nu_i\rangle \tag{2.3}$$

U is known as the Pontecorvo, Maki, Nakagawa and Sakata (PMNS) matrix [10] [11], where the matrix element $U_{\alpha i}$ gives the relative amplitude of mass eigenstate ν_i ($\nu_i = \nu_1, \nu_2, \nu_3$) found in flavour eigenstate ν_α ($\nu_\alpha = \nu_e, \nu_\mu, \nu_\tau$). Neutrinos can only be produced in weak interactions (as they have no electric or colour charge) and will be in a specific flavour eigenstate which by Equation (2.3) is a superposition of mass eigenstates. Each of the mass eigenstates travels with a different speed, consequently becoming out of phase some time later. A subsequent weak interaction will find the neutrino not in a single flavour state, as it was created, but a superposition. The upshot is that neutrinos can apparently oscillate between flavour states as they travel.

The PMNS matrix is usually decomposed in the following fashion:

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

Where $c_{ij} = cos(\theta_{ij})$, $s_{ij} = sin(\theta_{ij})$ and δ is a CP-violating phase. Factorising the PMNS matrix in this fashion groups the mixing parameters into sets that are probed by different types of experiments.

An illustrative example to demonstrate the phenomenon of neutrino oscillation is to take a two flavour approximation. In such an approximation we have two flavour and two mass states which are related by:

$$\begin{pmatrix} \nu_\alpha \\ \nu_\beta \end{pmatrix} = \begin{pmatrix} cos(\theta) & sin(\theta) \\ -sin(\theta) & cos(\theta) \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \end{pmatrix}. \tag{2.5}$$

A neutrino is produced by a weak interaction in a state ν_α with an energy E, then travels a distance L. At the source we have:

$$|\nu_{(x=0)}\rangle = |\nu_\alpha\rangle = cos(\theta)|\nu_1\rangle + sin(\theta)|\nu_2\rangle \tag{2.6}$$

and after travelling a distance L:

$$|\nu_{(x=L)}\rangle = |\nu_\alpha\rangle = cos(\theta)e^{ip_1L}|\nu_1\rangle + sin(\theta)e^{ip_2L}|\nu_2\rangle. \tag{2.7}$$

Then the probability of the neutrino being in flavour state ν_β at the distance L is given by:

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= |\langle\nu_\beta|\nu_\alpha\rangle|^2 \\ &= |(-sin(\theta)\langle\nu_1| + cos(\theta)\langle\nu_2|) * (cos(\theta)e^{ip_1L}|\nu_1\rangle + sin(\theta)e^{ip_2L}|\nu_2\rangle)|^2 \end{aligned} \tag{2.8}$$

and since

$$\langle\nu_i|\nu_j\rangle = \begin{cases} 1, i = j \\ 0, i \neq j \end{cases} \tag{2.9}$$

we have

$$\begin{aligned} P(\nu_\alpha \to \nu_\beta) &= |sin(\theta)cos(\theta)(-e^{ip_1L} + e^{ip_2L})|^2 \\ &= sin^2(2\theta)sin^2(\frac{(p_1 - p_2)L}{2}). \end{aligned} \tag{2.10}$$

In the limit that $E_i \gg m_i$ we arrive at:

$$P(\nu_\alpha \to \nu_\beta) = sin^2(2\theta)sin^2(\frac{(1.27\Delta m^2_{21})L}{4E}). \tag{2.11}$$

Where $\Delta m^2_{21} = m^2_2 - m^2_1$, L is the distance travelled measured in km and finally E is the neutrino energy in GeV. The result is that a neutrino produced via a weak interaction, in a specific flavour state, has a non-zero probability of being measured in a different flavour state after travelling some distance. This probability oscillates as a function of distance, such that an appropriate choice of propogation length can probe the mixing parameters θ and Δm^2_{21}.

Measurement of oscillation parameters

The initial findings of the SNO experiment have been complemented over recent years with measurements from a series of dedicated neutrino oscillation experiments. SNO made measurements of the 'solar' neutrino oscillation parameters (θ_{12} and Δm^2_{12}) which dominate the electron neutrino survival probability at this L/E. Kamiokande [12] and it's successor Super-Kamiokande [13] make observations of neutrinos produced in cosmic-ray induced pion and kaon decay chains to measure the 'atmospheric' oscillation parameters (θ_{23} and Δm^2_{23}). There are two main categories of additional experiment that contribute to these measurements: accelerator and reactor experiments.

Accelerator experiments involve a neutrino beam that is measured in a near detector close to the beam's origin and then propagated over a long baseline to a far detector. This approach allows for careful selection of the energy and baseline (E and L in Equation (2.11)) to gain the best measurement sensitivity. A number of experiments have used this approach to make precise measurements of the solar and atmospheric mixing parameters over the last decade [14] [15] [16].

Reactor experiments typically involve a detector placed near one or many nuclear reactors and measure the flux of neutrinos produced in the reactor core. Recent results from Daya Bay among other experiments give strong evidence for non-zero θ_{13} [17] [18] [19].

The current measurements of the oscillation parameters are summarised in Table 2.1. Although there is strong experimental evidence for neutrino oscillations there are still notable gaps including the sign of Δm^2_{23} and information about the CP-violating phase δ.

Parameter	Value
Δm^2_{21}	$7.50 \pm 0.20 \times 10^{-5}\text{eV}^2$
$\lvert\Delta m^{32}\rvert$	$2.32^{+0.12}_{-0.08} \times 10^{-3}\text{eV}^2$
$\sin^2\theta_{12}$	0.857 ± 0.024
$\sin^2\theta_{23}$	> 0.95
$\sin^2\theta_{13}$	0.95 ± 0.010

Table 2.1.: The best-fit values derived from a global fit to the current neutrino oscillation data [20].

2.2.2. Neutrino mass

The Standard Model predicts neutrinos to have zero mass, but the phenomenon of neutrino oscillations is best explained by the presence of finite non-zero mass states. Oscillation experiments are most sensitive to differences in the mass states as these directly affect the oscillations. The ordering of these masses from least to most massive is still unknown. Experiments such as NOνA [21] and LBNE [22] will have much improved sensitivity to the mass hierarchy by observing the effects of neutrinos' passage through matter. This is achieved by having longer baselines (L in Equation (2.11)), lower energies and more intense beams than previous neutrino beam experiments.

Oscillation experiments are able to place a lower limit on the heaviest mass state. The measurement of the largest mass splitting $|\Delta^2_{23}|$ combined with the fact that the lightest mass cannot be less than 0 leads to a lower bound on the heaviest active mass state.

Beta decay experiments are sensitive to the neutrino mass via measurements of electron energy in β decays. Tritium (^{3}H), an isotope of hydrogen, can undergo beta decay:

$$^3\text{H} \rightarrow ^3\text{He} + \text{e}^- + \overline{\nu}_\text{e}. \tag{2.12}$$

The energy of the emitted electron follows a β decay spectrum with an end point that depends upon the neutrino mass. In the case that neutrinos are massless the

Parameter	Value	Source
m_1 or m_3	> 0.05eV	Oscillations [20]
$\sum m_i$	$< 0.28 - 0.44$eV	Cosmology [23] [24]
m_β	< 2.0eV	β decay [25] [26]
$\langle m_{\beta\beta} \rangle$	$< 0.11 - 0.25$eV	$0\nu\beta\beta$ [27]

Table 2.2.: Experimental constraints on the neutrino mass.

end point of this spectrum will be equal to the difference in rest mass energy of ^{3}H and $^3\mathrm{He} + \mathrm{e}^-$. Since neutrinos are known to have mass the end point energy of the electron spectrum will be reduced by the neutrino mass. By making measurements of the electron energy it is therefore possible to place constraints on the absolute mass of the neutrino.

Another constraint on neutrino mass comes from cosmological measurements. The most important probe comes from anisotropies in the cosmic microwave background (CMB). A summary of these constraints is found in Table 2.2.

Chapter 3.

Ultra-high Energy Astro-particle Physics

Astro-particle physics encompasses detection of a wide range of particles produced in the Universe at large. The long standing goal of this field is to understand the high energy Universe through observations of these particles, however their detection is extremely challenging. As detector technologies improved astrophysical observations through optical telescopes have been complemented by other frequencies of light including gamma rays, X-rays and infrared frequencies. Recent experiments such as the HESS [28] and VERITAS [29] observatories have extended sensitivity at high energies, which will be further improved with the construction of the Cherenkov Telescope Array [30].

It was realised that electromagnetic radiation was not the only resource available to aid our understanding of astrophysical objects and phenomena. At present there are a number of other astrophysical messengers available to scientists each carrying new and often complementary information about the near and distant Universe. Very quickly it became as interesting to understand these particles themselves and the physics that governs their creation and interactions. The study of charged particles incident upon the Earth's atmosphere heralded a new era in physics in which a myriad of composite and elementary particles were discovered and opened a new discipline in the form of particle physics.

Since the days of Victor Hess and his electroscope experiments [1] much has been learned about the flux and energy of cosmic rays[1] incident on the Earth and many experiments have observed high energy particles of astrophysical origin. By some mechanism distant sources are able to accelerate charged particles to energies in excess of EeV (10^{18}eV), as cosmic rays with these energies have been observed for half a century [31]. These ultra high energy (UHE) cosmic rays can impact on a stationary targets, for example a nucleus in the atmosphere, interacting with centre of mass energies in excess of $\sim$ 100TeV, an order of magnitude higher than achievable with current accelerator technology. For this reason they provide a fascinating glimpse into particle physics in a regime inaccessible to scientists through conventional means.

UHE cosmic rays and other astrophysical messengers carry information about their sources and provide a powerful means to probe the high energy Universe, facilitating the study of extreme conditions in which to test and inform our understanding of physics in the Universe at large. There are, however, limitations to the information that we can learn from many of the messengers used at present. All but the most energetic charged particles will have their trajectories significantly bent by magnetic fields encountered in transit to the Earth. The other main limitation for most messengers are the horizon effects that limit the distances they can travel. UHE gamma-rays will pair produce e^- e^+ off the cosmic microwave background (CMB) preventing them from travelling large distances. Unbound neutrons are unstable with a proper lifetime of around 15 minutes [2], meaning that they will decay in flight producing cosmic rays. Although CMB photons have very low energies the centre of mass energy available when struck by UHE cosmic rays can be sufficient to cause photo-pion production, thus restricting the range of UHE cosmic rays. Neutrinos, on the other hand, do not suffer these horizon effects. Even at ultra-high energies weakly interacting neutrinos have such small cross-sections that they travel effectively unimpeded throughout the Universe.

The horizon effects that limit the information that we can obtain about the most energetic cosmic rays may also produce other messengers detectable on Earth. The interactions with CMB photons are expected to produce a flux of UHE neutrinos which, if detected, would provide vital information that may resolve some of the mysteries surrounding UHE cosmic rays. These cosmic rays will be sufficiently

[1] In this document the term "cosmic ray" will refer to protons or atomic nuclei, not to gamma rays or neutrinos

[2] At 100TeV the neutron lifetime is $\sim$ 3years

boosted that they will point back toward their sources and, since neutrinos are not bent by magnetic fields, it may be possible to identify UHE cosmic ray sources via the associated UHE neutrinos they produce. Furthermore the association of UHE neutrinos with the UHE cosmic ray flux leads to a number of spectral features that may prove vital in distinguishing between possibilities for their sources and acceleration mechanisms.

There are strong links between the flux of UHE neutrinos and UHE cosmic rays. This chapter will discuss the current understanding of the cosmic ray spectrum, the production mechanisms of UHE neutrinos and their detection.

3.1. Cosmic rays

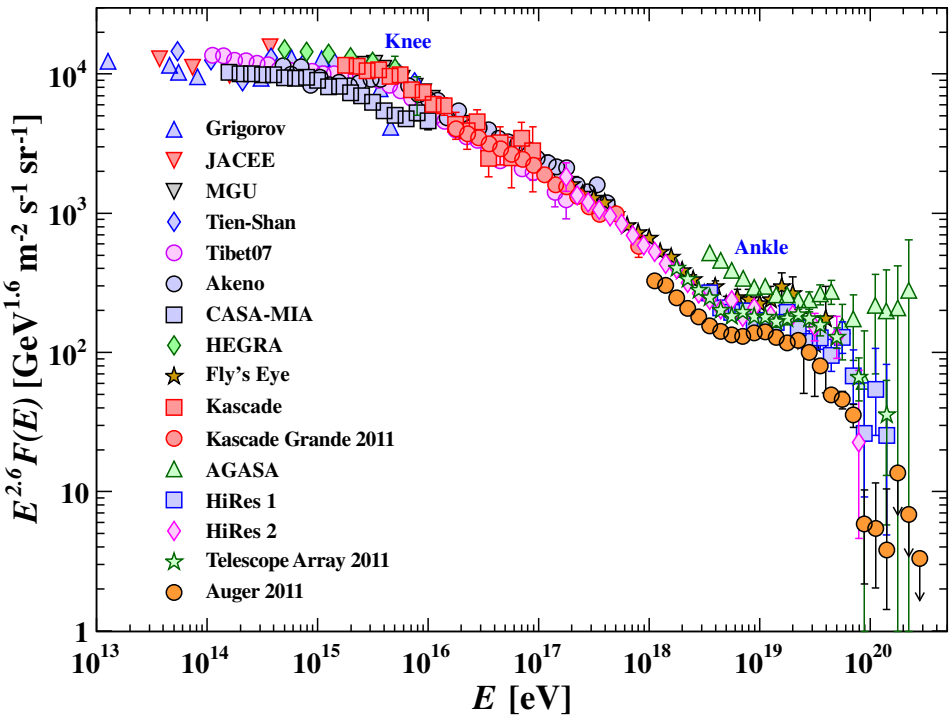

Figure 3.1.: The cosmic ray spectrum for all charged particles from [20]. The spectrum is multiplied by a factor $E^{2.6}$ to highlight some of the key features.

The flux of cosmic rays incident upon the Earth's atmosphere has been measured up to energies of around 10^{20}eV. The spectrum, which is shown in Figure 3.1, is steeply falling and is remarkably well approximated with a power law form $dN/dE \propto E^{-\gamma}$ where γ ranges from 2.7 to 3. There are a number of features in the spectrum that are thought to be due to transitions between different classes of source and acceleration mechanisms. The three main features are: the cosmic ray *knee* around $10^{15.5}$eV, the *ankle* at 3×10^{18}eV, and the *cut-off* above 3×10^{19}eV. The spectral index below the knee is $\gamma = 2.7$, steepening to $\gamma = 3$ between the knee and ankle, at which it returns to a similar index to below the knee. The flux falls off rapidly such that direct detection above 10^{15}eV is very difficult. At these energies the flux drops below tens of particles per m^2 per year and with typical direct detection experiments using balloons or satellites, being $\sim m^2$ in size, the collection of large samples is very difficult.

At these energies and above particles can be observed in arrays that sample the extensive air showers produced by cosmic ray interactions in the atmosphere. In addition it is possible to detect cosmic rays by measurements of nitrogen fluorescence from the shower using fluorescence detectors. The Auger experiment [32] makes use of both methods, with a ground array of Cherenkov detectors and fluorescence detectors sited on the perimeter. Recently the ANITA [33] experiment made observations of cosmic rays via radio pulses originating from the interaction of cosmic ray air showers with the magnetic field in Antarctica. The next flight of ANITA, a balloon-borne experiment, is expected to yield a significantly increased sample of these cosmic rays.

Calculating the primary energy of the cosmic ray in these experiments is challenging for a number of reasons: they involve large numbers of particles and interactions necessitating computer modelling, and the centre of mass energies are well beyond those that can be produced in the laboratory. This means that some extrapolation from experiments such as those at the LHC is needed, introducing uncertainties in particle populations expected in ground level detectors.

The current thinking in the field is that cosmic rays at and below the energy of the knee are produced in galactic astrophysical sources. The two most popular candidate sources are supernova remnants and binary systems. A popular interpretation is that the feature described as the knee is the result of these sources reaching their maximum acceleration energy. Cosmic ray acceleration mechanisms rely on strong magnetic fields. For heavy nuclei the maximum energy acquired in acceleration will

be Z times higher [3]. As we approach the energy of the knee cosmic ray sources cannot accelerate the lightest cosmic rays (protons) to higher energies. The flux is firstly taken over by a population of He nuclei, then by other heavier nuclei in order of charge. This produces a broad feature instead of the sharp one expected from cosmic rays composed of only one type of nucleus.

There are a number of explanations for the feature known as the ankle. One possibility is that it is a result of a higher energy population overtaking a population of lower energy particles, this could be a galactic flux being overtaken by an extra-galactic one. Another possibility is that the dip corresponds to electron positron production caused by interactions between cosmic ray protons and CMB photons[4] [34], which would again rely on the population being extra-galactic in nature due to the large propagation lengths needed for such interactions to take place.

3.2. UHE cosmic rays

At energies in excess of $\sim 6 \times 10^{19}$eV cosmic ray protons will rapidly loose energy in interactions with CMB photons, enhanced by the Δ^+ resonance, shown in Equation (3.1). These interactions lead to a suppression of the high energy tail in the cosmic ray flux. This suppression was first predicted by Greisen [35], Zatsepin and Kuzmin [36] soon after the discovery of the CMB and is known as the GZK cutoff.

$$\mathrm{p} + \gamma_{CMB} \rightarrow \Delta^+ \rightarrow \pi + N \tag{3.1}$$

$$\rightarrow \mathrm{p} + \mathrm{e}^- + \mathrm{e}^+ \tag{3.2}$$

Along with the photo-pion production seen in Equation (3.1) cosmic rays can produce electron positron pairs as in Equation (3.2). The threshold for pair production is about 10^{18}eV with a mean free path ~ 1 Mpc, whereas for photo-pion production the threshold is 6×10^{19}eV with mean free path ~ 6 Mpc. Despite the lower thresholds the energy losses are dominated by photo-pion production as the energy loss per interaction is $\sim 20\%$ compared with only 0.1% for pair production. The GZK cutoff

[3] Where Z is the charge of the nucleus.
[4] $\gamma + \mathrm{p} \rightarrow \mathrm{e}^- + \mathrm{e}^+$

also produces an effective horizon for cosmic ray protons limiting their path length to ~ 100Mpc.

In the case that UHE cosmic rays are heavy nuclei photo-disintegration and pair production become important:

$$A + \gamma_{CMB} \rightarrow (A-1) + N \tag{3.3}$$
$$\rightarrow (A-2) + 2N \tag{3.4}$$
$$\rightarrow A + \mathrm{e}^- + \mathrm{e}^+ \tag{3.5}$$

where A is the mass of the nucleus and N is a secondary nucleon. This process leads to a similar suppression in the cosmic ray flux as in the cosmic ray proton case. Nuclei are still able to undergo the GZK process of Equation (3.1) both before and after photo-disintegration and pair production, but the nucleons have on average $\frac{1}{A}$ of the nuclei's energy. This results in a shift upward in the threshold for photo-pion production, in the case of iron nuclei to $\sim 3.5 \times 10^{21}$eV.

A suppression in the flux of UHE cosmic rays consistent with the GZK cutoff has been observed in both the Auger [37] and HiRes [38] experiments at $10^{19.5}$eV and is shown in Figure 3.2. Although UHE cosmic rays at and after the GZK cutoff have been detected the composition in this regime remains uncertain. Measurements of the depth of shower maximum X_{max}, a parameter that is strongly correlated with cosmic ray composition, are consistent with a move to heavier nuclei at high energies in the Auger experiment [39].

Air showers produced by primaries heavier than protons can be thought of as a superposition of A showers each with energy $\frac{E}{A}$, where A is the mass of the primary. In such a scenario the depth of shower maximum is reduced to the value of a proton with energy $\frac{E}{A}$ and hence heavier primaries give rise to lower measured values of X_{max}. In addition the variation of this quantity, $RMS(X_{max})$, will be reduced due to the presence of A showers, as opposed to a single shower in the case of protons.

Figure 3.3 shows the most recent results measuring these parameters with the Auger experiment. It should be noted that the statistics are limited due to the low fluxes at high energies, and the move to smaller values of X_{max} can equally be

explained by modifications to the interaction cross section at these energies (which are extrapolated from those measured in particle physics experiments).

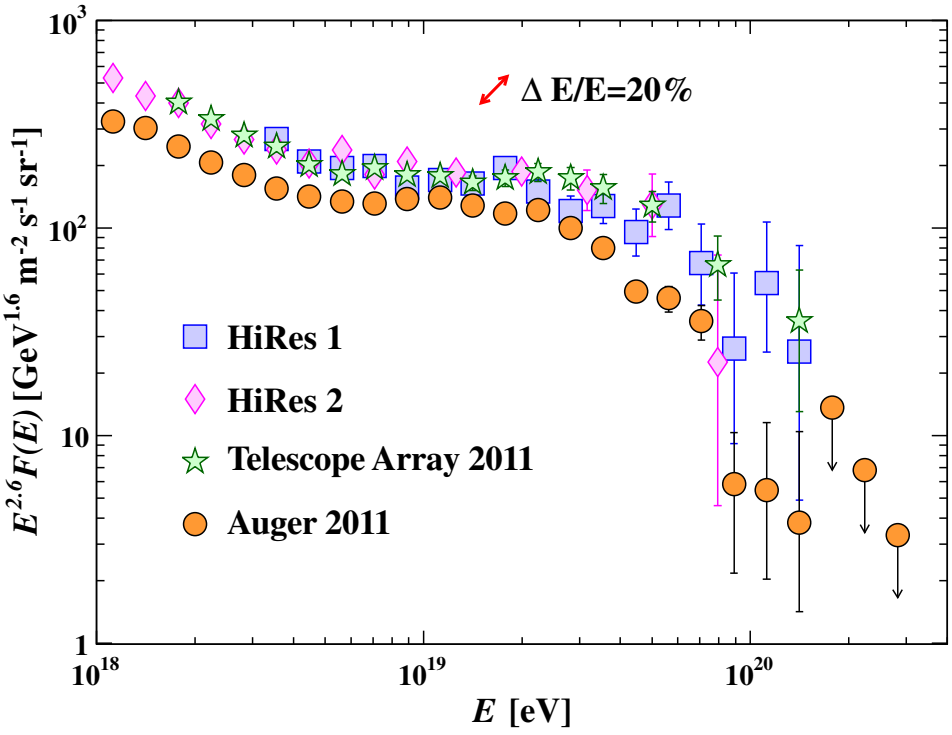

Figure 3.2.: The cosmic ray spectrum in the UHE regime from [20]. The flux is multiplied by $E^{2.6}$ to highlight key features and to aid comparisons with Figure 3.1. Both HiRes and Auger data show features consistent with the cosmic ray ankle.

There remain a number of mysteries surrounding UHE cosmic rays. Due to the GZK cutoff and the associated horizon, their sources must be nearby in cosmological terms. Due to their very high momentum, and hence rigidity, at these energies there should be little deviation from their source by magnetic fields, but to date none have been identified. The energies are so massive that it is very hard to explain how astrophysical objects provide sufficient acceleration, as illustrated in the Hillas plot in Figure 3.4. Complementary information is required to address some of these unanswered questions, information that may be provided by observations of UHE neutrinos.

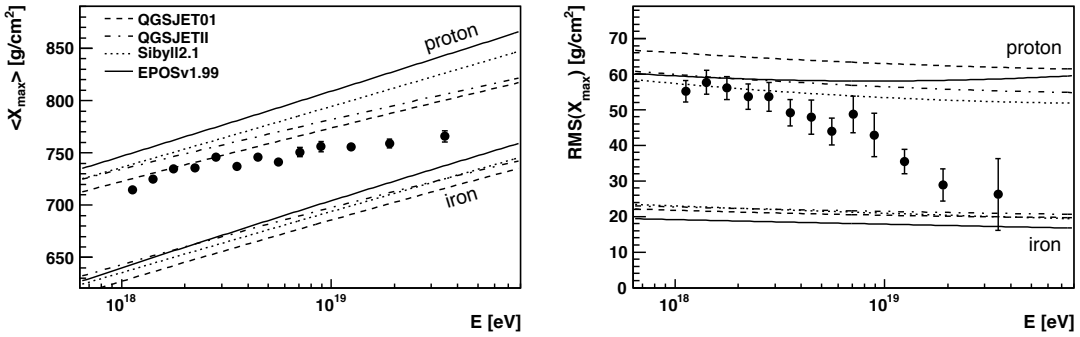

Figure 3.3.: The UHE cosmic ray composition as measured by Auger from [39]. The average depth of shower maximum $\langle X_{max} \rangle$ and the variation in depth of shower maximum $RMS(X_{max})$ are shown as a function of energy from Auger data (points). For comparison the expected values from simulation (lines) are also shown. Both parameters show a trend toward higher mass cosmic rays with increasing energy.

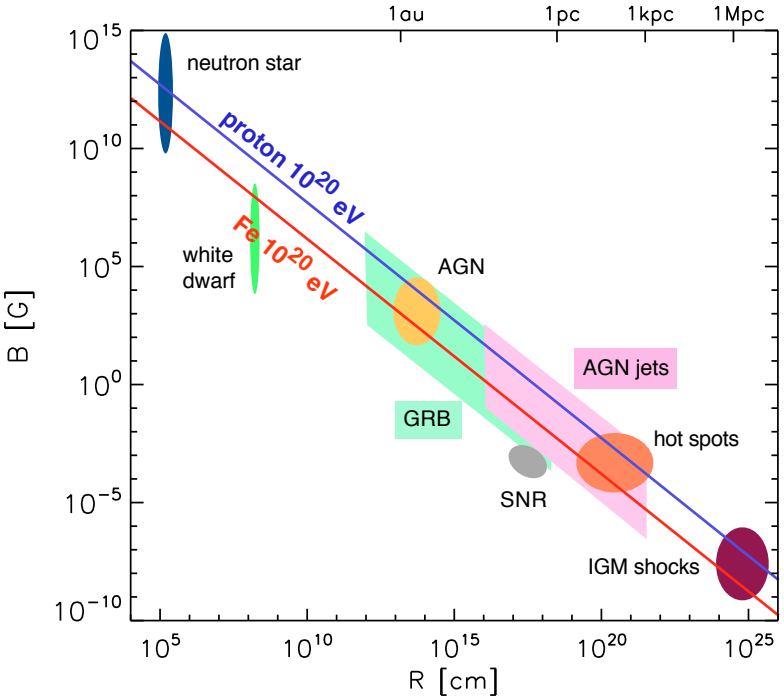

Figure 3.4.: The Hillas plot from [40]. Sources above the red and blue lines are unable to confine (and hence accelerate) iron nuclei to 10^{20}eV and protons to 10^{21}eV via magnetic fields.

3.3. UHE neutrinos

Considering the GZK process in Equation (3.1) it is clear that subsequent decays of charged pions and neutrons will produce a so called 'guaranteed' flux of neutrinos. Beresinsky and Zatsepin [41] were the first to predict such a flux and they are referred to as (BZ) neutrinos. BZ neutrinos are produced in the following decay of charged pions:

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu \\ \mu^+ &\rightarrow \overline{\nu}_\mu + \mathrm{e}^+ + \nu_\mathrm{e}. \\ \mathrm{n} &\rightarrow \mathrm{p} + \mathrm{e}^- + \overline{\nu}_\mathrm{e} \end{aligned} \tag{3.6}$$

The GZK process will produce ν_e, ν_μ and $\overline{\nu}_\mu$ with the associated flux having a ratio of flavour states ν_e: ν_μ: ν_τ of $1 : 2 : 0$. In the case of neutrinos produced in photo-disintegration and neutron decay only $\overline{\nu}_\mathrm{e}$ are produced.

Due to the production mechanism there is strong link between the neutrino and cosmic ray spectra. Observing and measuring the properties of BZ neutrinos would provide otherwise unobtainable information regarding the cosmic ray spectrum. For example it would be possible to infer details of the emission spectra of cosmic rays (such as the maximum energy to which they are accelerated) and information about the sources.

An example of a predicted flux is shown in Figure 3.5. The flux of ν_μ and $\overline{\nu}_\mu$ come exclusively from the decay of charged pions from the photo-pion production mechanisms discussed previously. Since the UHE cosmic ray spectrum is steeply falling most of these interactions will occur close to the threshold nucleon energy of 6×10^{19}eV and transfer approximately 5% of the energy available to neutrinos, resulting in a single peak between 10^{18} and 10^{19} eV in both the electron and muon neutrino fluxes. The lower energy peak between 10^{16} and 10^{17} eV is due to neutron decays producing $\overline{\nu}_\mathrm{e}$. For reference the Waxman-Bachall [42] [43] limit is shown on the same axes. This limit is an upper bound on neutrino fluxes caused by photo-pion production in the sources of UHE cosmic rays and is a commonly used reference flux. Figure 3.6 shows the effects on the neutrino flux resulting from various compositions

for UHE cosmic rays. This shows, as expected, a supression of the higher energy flux of neutrinos (which are comprised mainly of ν_e, ν_μ and $\overline{\nu}_\mu$ from photo-pion production) and associated increase in the lower energy flux (comprised of $\overline{\nu}_e$ from neutron decay) as the comsic ray composition shifts to heavier nuclei.

3.3.1. Particle physics with UHE neutrinos

UHE neutrinos are particularly interesting from a particle physics perspective due to their enormous energies and vast distances over which they propagate. Neutrinos impacting on stationary protons can give rise to centre of mass energies an order of magnitude larger than those available at the LHC. For example a 10^{19} eV neutrino would give rise to ~ 140 TeV being available to produce new particles. Studying their interactions and possible divergences from those expected would be a powerful probe for new physics beyond the standard model in a regime that it is difficult to imagine replicating through current accelerator technology.

If experiments were able to distinguish between neutrino flavours there is the possibility of observing neutrino oscillations on baselines and energy scales previously inaccessible. With the current understanding of neutrino oscillations the flavour composition described in Section 3.3 would be maximally mixed leading to a ratio ν_e: ν_μ: ν_τ close to $1 : 1 : 1$.

Having established the observational and theoretical motivation for detecting UHE neutrinos the practicalities are now discussed.

3.3.2. UHE neutrino detection

As neutrinos only interact via the weak force it is not possible to detect them directly, they must interact via processes shown in Figure 2.1 and observations made of the by-products. At 10^{19} eV the neutrino nucleon interaction cross-section is expected to be in the region of $0.3 \times 10^{-31}\text{cm}^2$ [46] making interactions probable in km^3 detectors. This requires experiments to have large active volumes, too large to be purpose built, and hence naturally occurring bodies are utilised.

Cosmic ray air shower experiments, such as Auger, make use of the Earth's atmosphere as an interaction volume for cosmic rays. They also have sensitivity to

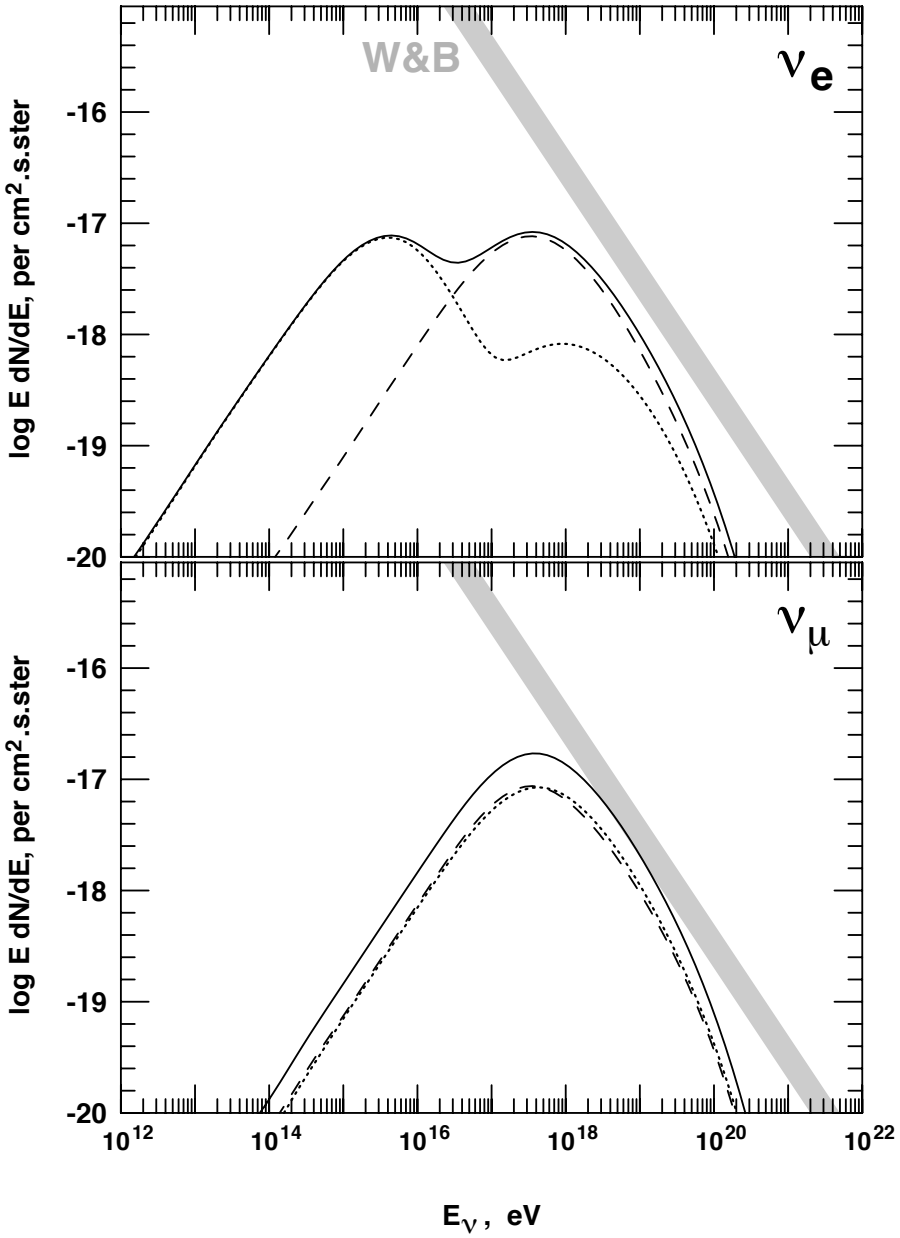

Figure 3.5.: Predicted fluxes of ν_e (top) and ν_μ (bottom) neutrinos from [44]. Dashed lines correspond to neutrino fluxes, dotted lines to anti-neutrino fluxes and the sum total by solid lines.

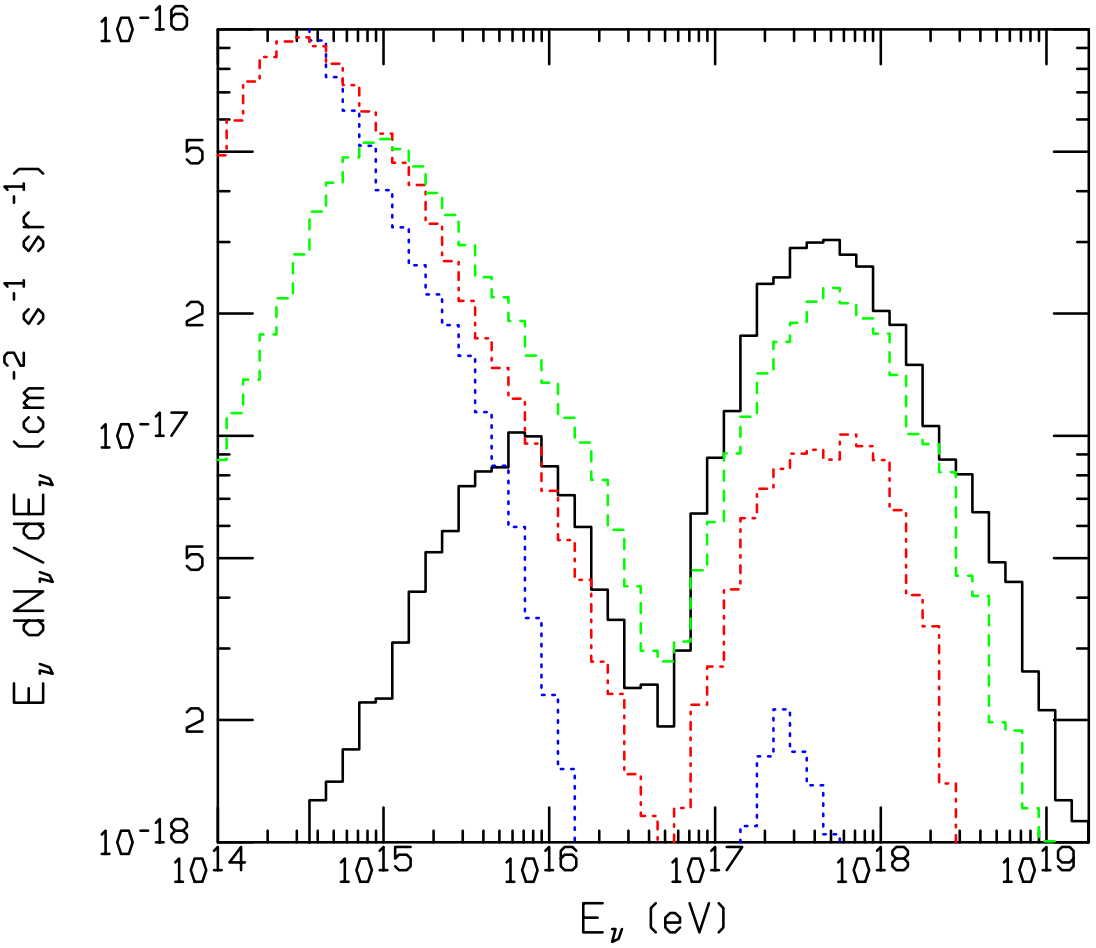

Figure 3.6.: Predictions of BZ neutrino fluxes ($\nu + \overline{\nu}$) for protons (black, solid), ^{4}He (green, dashed), ^{16}O (red, dash-dotted) and ^{56}Fe (blue, dotted) from [45].

neutrino interactions as the volume they see is so large. Distinctions can be made between cosmic ray and neutrino induced air showers, although so far none of the latter have been observed. As cosmic rays (be they protons or heavy nuclei) have much higher interaction cross-sections than neutrinos the corresponding interaction length in the atmosphere is relatively short. By looking for air showers at large zenith angles (i.e. close to the horizontal) the length of atmosphere traversed by any particle interacting close enough to be detected is very large, and therefore highly unlikely to be a cosmic ray. Auger is most sensitive to ν_τ by looking for showers induced by the decay products of a τ after the propagation and interaction of a ν_τ in the Earth. The inferred shower direction is required to be Earth-skimming or coming from a nearby mountain range, with the effective volume set by the decay length of τ at high energies ($\sim$ 10km).

Many experiments rely on detecting neutrinos through the light emitted by their interaction products. In, for example, a charged current interaction between a neutrino and nucleon energy is transferred from the neutrino to a charged lepton, and nucleus, which quickly develops into a shower of charged particles. Due to the high energies involved these particles will be travelling at super-luminal speeds and Cherenkov light is emitted. Experiments such as IceCube [47], ANTARES [48] and the future experiment km3Net [49] make use of naturally occurring materials to produce enormous detectors large enough to be sensitive to the low neutrino fluxes. IceCube, which was built around the previous AMANDA [50] experiment close to the geographic South Pole, uses ice as the interaction medium and observes Cherenkov light using photo-multiplier tubes (PMTs). For ANTARES and km3Net the interaction medium is water in the Mediterranean.

IceCube primarily searches for lower energy neutrinos in the TeV- PeV range. Cosmic ray induced μ are a significant background at the lower end of this range. At these energies searches focus on up-coming neutrinos that have traversed the Earth before entering the detector from below. At high energies a different approach is taken, the outer layer of PMTs is used as a veto allowing for detection of neutrinos incident from the sides and top of the detector. These methods have proved useful in observing neutrinos over a wide range of energies, including the first observation of neutrinos believed to be of extra galactic origin [51] [52] [53].

Despite the success of Cherenkov light based experiments at detecting energetic neutrinos, the attenuation lengths of optical light in naturally occurring media are

such that the technology is prohibitively expensive to scale up to the $\sim$ 100km^3 necessary to observe UHE neutrinos above PeV.

Neutrino interactions in dense media can cause the deposit of thermal energy, via ionisation losses, which could be measured in acoustic detectors. Given sufficiently large attenuation lengths in the chosen medium this method could be used to instrument km^3 scale detectors. To date only proof of principle experiments have been constructed and operated. This method has been used to instrument both ice and water as the interaction and signal propagation media through ACoRNE [54], AMADEUS [55], Lake Baikal [56], SAUND [57] and SPATS [58] experiments.

3.3.3. The Askaryan effect and radio detection

In 1962 Gurgen Askaryan proposed that extremely energetic particle cascades, such as those induced by neutrino interactions, in dense dielectrics could produce coherent radio pulses [2] [3]. Secondary electrons, positrons and gamma rays cause an electromagnetic shower to develop in the medium. Although the incident neutrino carries no electric charge a net excess in charge builds due to the presence of electrons in the medium. A combination of scattering effects cause electrons to be promoted from the medium into the shower, at the same time positrons annihilate with electrons in the medium, resulting in a net negative charge excess of $\sim$ 20%. The charge excess travels at super-luminal speeds through the medium causing Cherenkov emission which adds coherently for wavelengths greater than the shower dimensions. The power emitted scales with the number of particles in the shower squared N^2 for coherent emission, which occurs at frequencies below 1GHz in ice.

This effect was experimentally confirmed in a series of experiments using the SLAC beam and a range of naturally occurring materials, initially in sand [4] and later with salt [5] and ice [6]. In the absence of a controlled source of UHE neutrinos this was achieved by instigating electromagnetic showers using the SLAC beam. Short time duration pulses of GeV photons were fired into a sand target in the first experiment, and in the latter two experiments beams of electrons into salt and ice, resulting in the development of electromagnetic showers over a number of meters. The expected radio emission was measured using radio antennas and exhibited the characteristic broadband frequency content, linear polarisation and coherence, as well as the scaling of emitted power with respect to number of charged particles

shown in Figure 3.7. Figure 3.8 shows an example Askaryan pulse from the SLAC beam tests [5] which is consistent with the predicted bimodal signal with a rise time $\sim$ 100ps.

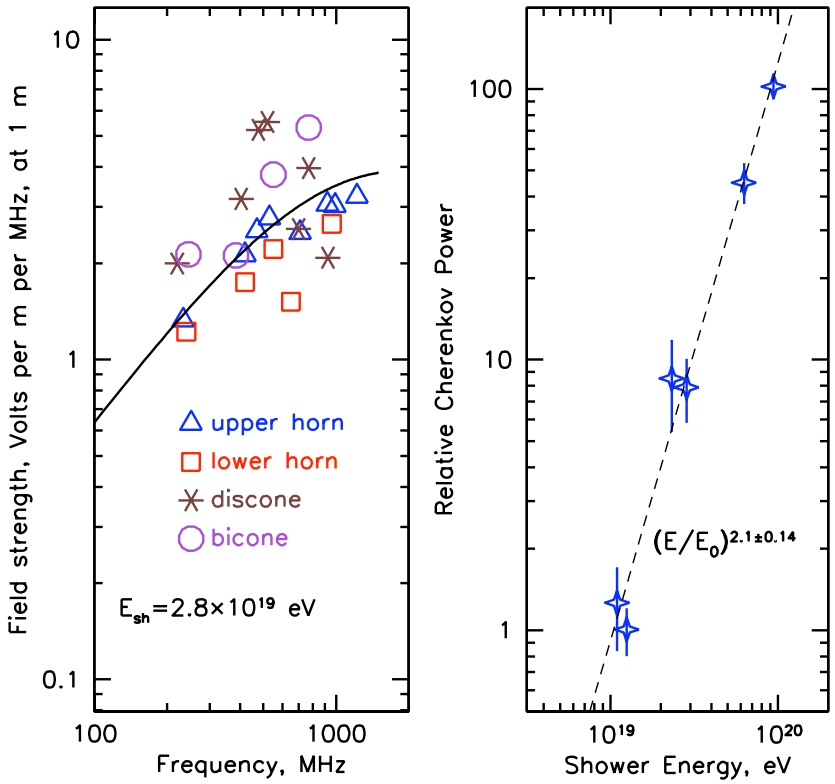

Figure 3.7.: Measurements of Askaryan radiation in ice from [6]. The measured field strength of an Askaryan pulse as a function of frequency is shown on the left as measured by a series of different antenna designs (triangles and squares correspond to horn antennas at the top and bottom of the ANITA instrument [59], whereas stars and circles are from measurements using additional antennas following discone and bicone designs). The observed power as a function of shower energy is shown on the right, demonstrating the expected quadratic dependency.

The dimensions of such showers are determined by the properties of the medium in which they develop. The Moliere radius defines the transverse size of the shower, and hence the charge excess, and is given by:

$$R_M = X_0 \times 21\text{MeV}/E_c \tag{3.7}$$

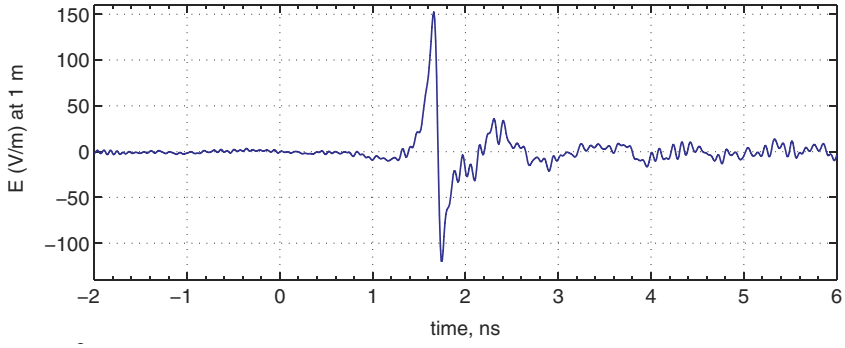

Figure 3.8.: Askaryan pulse field strength as measured in salt from [5].

where E_c is the critical energy and X_0 is the radiation length, both of which depend upon the medium. The radiation length is the distance over which a particle will lose all but $\frac{1}{e}$ of its energy. The critical energy is that where the ionisation loss rate is equal to that due to bremsstrahlung. Ionisation does not produce sufficiently energetic particles to add to the shower, and hence is essentially an energy loss mechanism. Bremsstrahlung, on the other hand, produces photons that can pair produce or promote atomic electrons into the shower, and is a mechanism by which the shower grows. In the case of ice the the radiation length $X_0 \sim 40$cm and the critical energy $E_c \sim 54$MeV, leading to $R_M \sim 10$cm. The shower development occurs over several meters in ice, with $\sim 90\%$ of the charge excess contained within a cylinder R_M wide and ~ 1cm long. This 'pancake' of charge gives rise to coherent emission below 1GHz.

Since the attenuation length of radio signals in ice is an order of magnitude longer than for optical light (~ 1 km [60] [61] versus ~ 100m for optical), detecting UHE neutrinos via Askaryan radiation has the potential of allowing detector volumes large enough to observe BZ neutrinos. The abundance of ice in Antarctica has led to a number of pioneering efforts to detect neutrinos in this manner.

The ANITA [59] [62] experiment consisted of radio antennas mounted an a long duration high altitude balloon that can see millions of km^3 of ice. To date two flights have been made, each lasting ~ 30 days, with a third due for the 2014-2015 austral summer. Due to the large volume of ice that ANITA can observe the experiment is able to place world best limits on the high energy tail of the expected GZK flux. However, the distance from neutrino interaction point leads to low signal to noise ratios, and detailed analysis is needed to remove signals of an anthroprogenic nature.

The live time of such experiments is limited as there is only a short window in the summer in which it is possible to launch and recover the experiment. This also coincides with the peak of human activity on the continent and the associated increased radio background this brings with it.

The Radio Ice Cherenkov Experiment (RICE) [63] was composed of 18 radio antennas buried in the ice close to the south pole. The antennas operated in a frequency range of $100\text{MHz} - 1\text{GHz}$ and were deployed in a 200m wide cuboid 600m above the AMANDA neutrino telescope (which formed the precursor to the IceCube experiment). Installation of the radio antennas brought them closer to potential neutrino signals than ANITA, increasing the signal to noise ratios, however with this comes a decrease in detector volume associated with the geometry and attenuation of signals in the ice.

Two experiments currently under construction aim to have much greater sensitivity to neutrino fluxes in the $10^{17}\text{eV} - 10^{20}\text{eV}$ range. The Askaryan Radio Array (ARA), which will be described in detail in Chapter 4, and the Antarctic Ross Ice-shelf ANtenna Neutrino Array (ARIANNA) [64] will both consist of a large number of radio antennas buried in ice. ARIANNA will be formed of over 900 independently operating stations each of which contains 8 antennas buried in the Ross Ice Shelf, Antarctica. Neutrino induced cascades are detected via radio emission that arrives at the antennas either directly from the shower or indirectly, having reflected off the ice-sea boundary below the ice shelf.

Chapter 4.

The Askaryan Radio Array

Radio detection of neutrinos is a promising experimental method in the search for UHE neutrinos. The likes of the ANITA [59] [62] and RICE [63] experiments have shown that it is feasible to instrument large volumes of ice relatively inexpensively. However, to date no UHE neutrinos have been observed with such experiments and the next generation detectors must find a way of improving sensitivity by an order of magnitude or more in the $10^{17} - 10^{18}$eV region. The Askaryan Radio Array (ARA) is one such experiment that looks to build on the pioneering work by other experiments in the field. ARA will consist of a series of antenna clusters, or stations, buried deep in the ice near the Amundsen-Scott South Pole Station. Each of these antenna clusters will have dedicated triggering and digitisation electronics to enable them to operate as stand-alone neutrino detectors. This design lends itself to a phased installation that will allow the detector volume to grow as sensitivity is required, with the aim of first establishing an UHE neutrino flux and then giving the option to expand to an observatory class detector to make detailed measurements.

By burying the antennas in the ice the signal to noise ratio expected from neutrino signals increases significantly compared with balloon-borne experiments, such as ANITA, leading to a decrease in the energy threshold but with an associated decrease in detector volume due to the geometry and attenuation of radio signals within the ice. In addition the stations are able to operate year round and are not limited to the relatively noisy summer season in which human activity is at its peak. The ANITA experiment has to remove significant amounts of anthroprogenic radio signals from their data set. By choosing a location that is relatively isolated, and by collecting

data year round ARA will be able to greatly reduce the limitations that this places on live time and hence neutrino sensitivity.

The ice sheet upon which the South Pole Station sits is ~ 2.8 km thick and has exceptional radio clarity [60] [61]. Figure 4.1 shows the results of ice attenuation measurements taken at the South Pole which results in an attenuation length estimate of $L_\alpha = 1450^{+300}_{-150}$m at a frequency of 380MHz and temperature of $-50°$C.

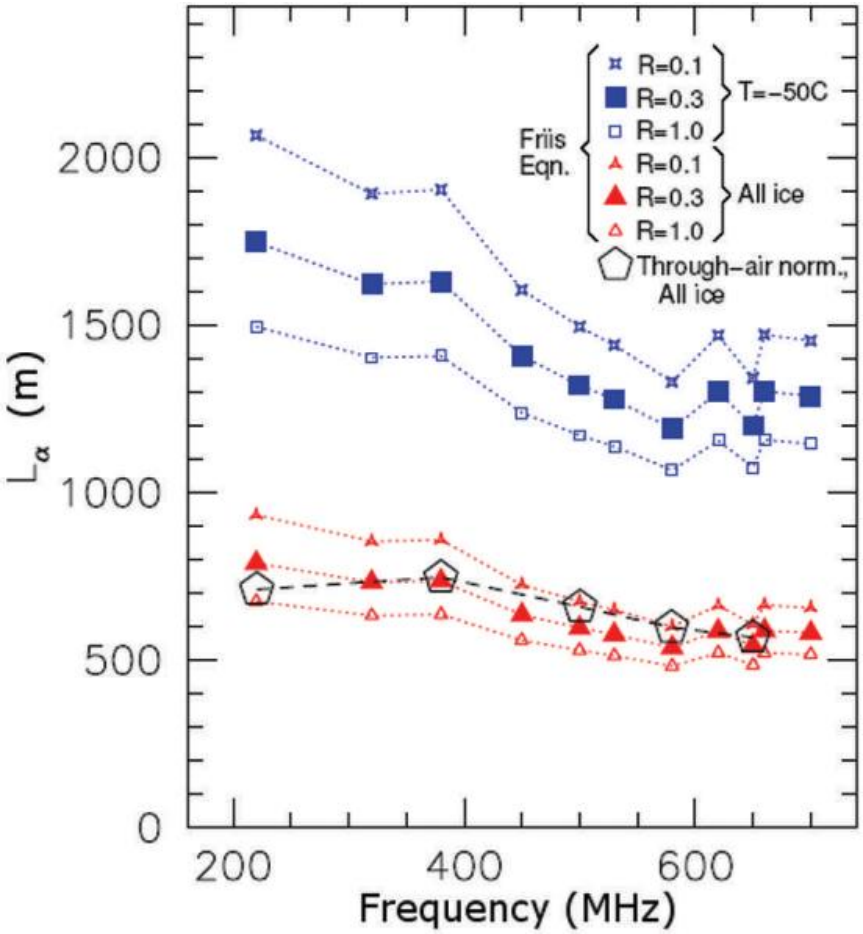

Figure 4.1.: Ice attenuation lengths as a function of frequency from [60]. The lower set of lines correspond to average attenuation lengths under various assumptions for reflectivity of the bedrock below the ice sheet. The open pentagonal symbols are obtained by normalising the transmitted and received signals in the air relative to in the ice. The upper set of lines show derived attenuation lengths taking into account the temperature profile in the ice.

The top layer of ice, known as the firn, consists of compacted snow of lower density than solid ice, leading to a varying index of refraction through this layer. This results in ray-bending and shadowing affects that limit the effective volume of a detector

deployed within it, whilst complicating triggering and reconstruction of incident radio signals' origin. The firn typically extends for ~ 150m, below which the index of refraction changes little [61]. There is a clear benefit to installing antennas below this layer to circumvent these effects, but the costs of drilling rise rapidly with depth. The challenge of drilling wide, deep and dry holes to these depths is not insignificant, although ARA is able to draw on the expertise within the IceCube collaboration which successfully deployed photo-multiplier tubes at depths of $1.5 - 2.5$km [47].

The expected signal from neutrino induced Askaryan radiation is a highly linearly-polarised radio frequency (RF) impulse. This will come in the form of a spherical wave-front, which is well approximated by a plane wave for distant sources. The antennas effectively sample this wave-front and the timing differences between signals received in pairs of antennas can be used to identify both the source direction and the distance from the station. By measuring the polarisation of detected signals it is possible that additional information can be obtained about which part of the Cherenkov cone each antenna has sampled, hence providing another observable with which to determine the event topology. The receive antennas used within a station should have a dipole response and be split into polarisation along two orthogonal directions. This allows for detection of RF signals polarised along an arbitrary direction removing any bias to a particular polarisation.

The impulsive nature of the Askaryan signal, with experimentally observed rise times of < 100ps [4], necessitates high sampling rate digitisation of the analogue signal received by the antennas. Power is at a premium in such remote locations, placing a severe constraint on the consumption of the trigger and data acquisition systems. Application Specific Integrated Circuits (ASICs) are used that meet these challenging requirements.

Askaryan radiation is not limited to production in UHE neutrino induced cascades, but can also be produced from cosmic ray primaries. Typically a cosmic ray will interact in the atmosphere causing a shower of secondary particles to be produced. At sufficiently high incident energy it is possible for the core of this air shower to penetrate the ice and produce Askaryan emission, in a manner analogous to that described for neutrino induced showers. This places a very high energy threshold on the cosmic rays detectable by this mechanism as their air showers not only need to penetrate the ice, but also to contain sufficient energy to produce a detectable signal. These cosmic ray events are further suppressed by their geometry. Down

going cosmic rays are much more likely to initiate a shower that penetrates the ice, but the active volume of ice viewable by ARA is very small for such events.

In the case that such cosmic ray induced signals were observed in ARA it would be necessary to separate them from neutrino induced signals. This is likely to be possible through reconstruction of an observed signals' source location. Cosmic ray induced signals would be limited to being produced close to the ice surface as the air shower core cannot propagate far into the ice. No such limitation is present for neutrino induced signals.

4.1. ARA 37

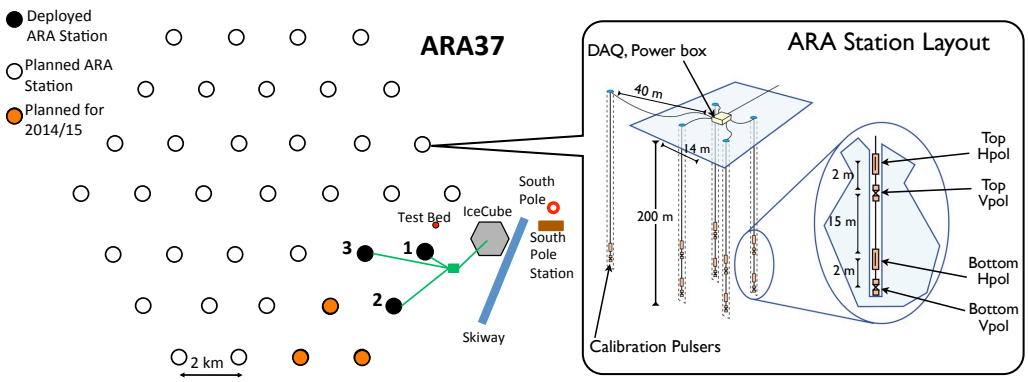

Figure 4.2.: *Left* the ARA-37 layout including TestBed and ARA1-3 provided by Ryan Manu. *Right* schematic for idealised station similar to those deployed as ARA1-3.

The current proposed design for ARA is a hexagonal arrangement of 37 stations named 'ARA37' shown in Figure 4.2. The current design for ARA37 is driven by the aim of observing a flux of UHE neutrinos, and as such this phase will essentially be a counting experiment. The design choices, in terms of number of antennas per station and the geometry of stations within the detector as a whole, are optimised to maximise detector volume. This does not preclude multi-station events, which would provide extra information to determine event details and energies. These events, however, are expected to be rare in comparison with single station events. Once a UHE neutrino flux is established it would be possible to in-fill with stations to improve energy resolution and better distinguish between classes of neutrino event. Tau neutrino interactions, for example, can lead to a secondary decay of a τ

within the detector volume. In addition interactions of ν_μ can produce muons that undergo photonuclear energy losses spatially separated from the initial shower [65], also leading to Askaryan emission. Such 'double-bang' events could produce two particle showers being reconstructed using separate stations.

The modular design enables the individual detectors to be developed as the array is installed, informed by the experience of installation, operation and the data from those previously deployed. A prototype station, known as the 'TestBed', was installed during the austral summer 2010-2011, which is marked in Figure 4.2, and is described in detail in Section 4.2.

4.2. The TestBed

During the austral summer of 2010-2011 a prototype station was deployed in the ice designed to provide a tool with which to asses the suitability of the ice, the radio environment, to aid the development of future stations and provide data with which to conduct the first neutrino analysis. The TestBed consists of 14 horizontally (HPol) and vertically (VPol) polarised [1] antennas deployed in the ice, alongside 2 horizontally polarised surface antennas and a dedicated digitisation, triggering and data recording unit known as the 'DAQ box'. The surface antennas were intended for use in searching for radio emission produced by cosmic ray air showers. The layout of the TestBed is shown diagrammatically in Figure 4.3.

The dedicated hot-water drilling equipment necessary to drill below the firn layer was not yet available limiting the depth of holes to $\sim$ 30m. 6 HPol and 4 VPol were deployed at a depth of 30m, the rest to a depth of $\sim$ 1m. The specifications for various aspects of the TestBed are summarised in Table 4.1 and can be compared to the future stations design shown in Table 4.3.

The TestBed station geometry differs from the design layout for future stations due, in part, to the logistical challenges faced for the initial deployment. The deployed positions and antenna types are summarised in Table 4.2 in a station-centric coordinate system. The coordinate system chosen has $+\hat{x}$ in the direction

[1] Horizontally (Vertically) polarised antennas will be referred to as HPol (VPol) in this document.

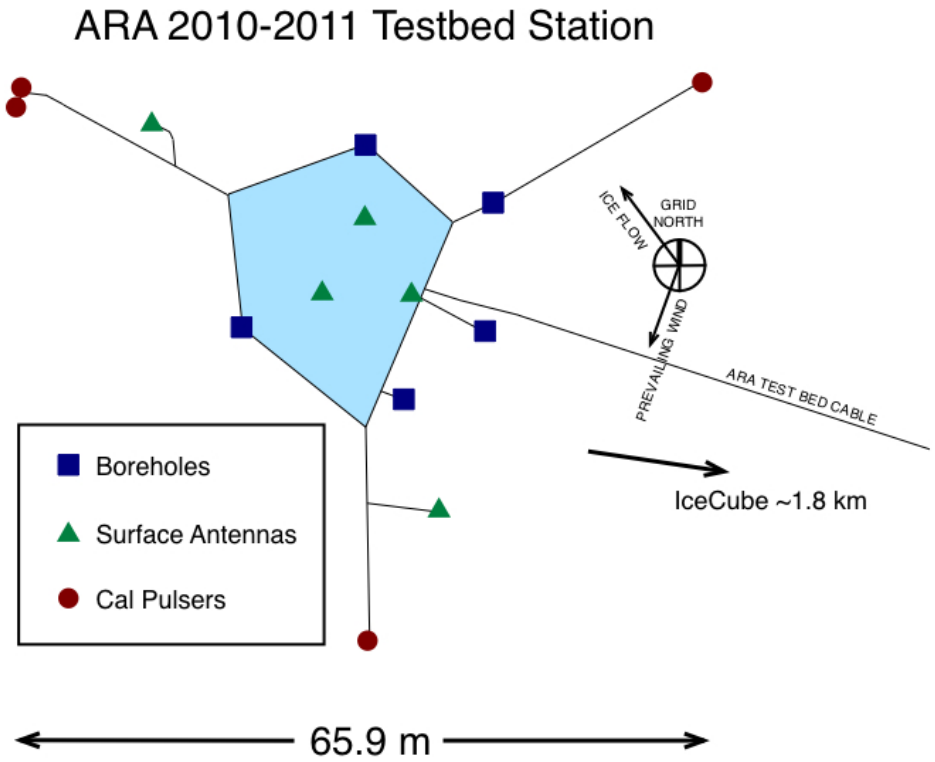

Figure 4.3.: The layout of the TestBed provided by Eugene Hong.

Specified Parameter	**TestBed Station**
Number of VPol antennas	2 near-surface, 4 in ice
VPol antenna type	bicone
VPol antenna bandwidth (MHz)	150-850
Number of HPol antennas	2 near-surface, 6 in ice
HPol antenna type	BSC & QSC
HPol antenna bandwidth (MHz)	250-850
Surface antenna type	fat dipole
Surface antenna bandwidth (MHz)	30-300
Number of surface antennas	2
Number of receive antenna boreholes	4
Borehole depth (m)	30
Vertical antenna configuration	VPol (HPol) above HPol (VPol)
Vertical spacing (m)	5
Approximate geometry	trapezoidal
Approximate radius (m)	10
Number of calibration antenna boreholes	3
Calibration borehole distance from centre (m)	30
Calibration hole geometry	equilateral triangle
Calibration signal type	impulse only
LNA noise figure (K)	< 80
LNA/amplifier dynamic range	30:1
RF amplifier total gain (dB)	> 75

Table 4.1.: TestBed detector specifications.

Hole	Antenna	Position		
		X (m)	Y (m)	Depth (m)
BH 1	BSC, HPol	-8.42	-4.40	20.50
	Bicone, VPol			25.50
BH 2	BSC, HPol	-0.42	-11.13	27.51
	Bicone, VPol			22.51
BH 3	BSC, HPol	9.22	-6.15	22.73
	Bicone, VPol			27.73
BH 5	BSC, HPol	3.02	10.41	30.56
	Bicone, VPol			25.56
BH 6	QSC, HPol	-9.07	3.86	26.41
	QSC, HPol			30.41
S1	Discone, VPol	-2.48	-1.75	1.21
	Batwing, HPol			2.21
S2	Batwing, HPol	4.39	-2.41	1.19
S3	Discone, VPol	1.58	3.80	1.19
S4	Fat Dipole, HPol			
Cal1	HPol	-23.18	17.90	17.50
	VPol			22.50
Cal2	HPol	-2.25	-29.81	34.23
	VPol			29.23
Cal3	HPol	27.58	13.63	1.13
	VPol			1.13

Table 4.2.: TestBed boreholes, antenna types and deployed positions.

of ice flow, the $\hat{x} - \hat{y}$ plane tangent to the earth's geoid shape at the surface and is centred on the south east corner of the DAQ box on the surface of the ice [66].

4.2.1. Signal chain

The TestBed signal chain is shown in figure Figure 4.4. The signal chain consists of radio antennas connected via co-axial cables, and various stages of amplification and filtering, to the DAQ box.

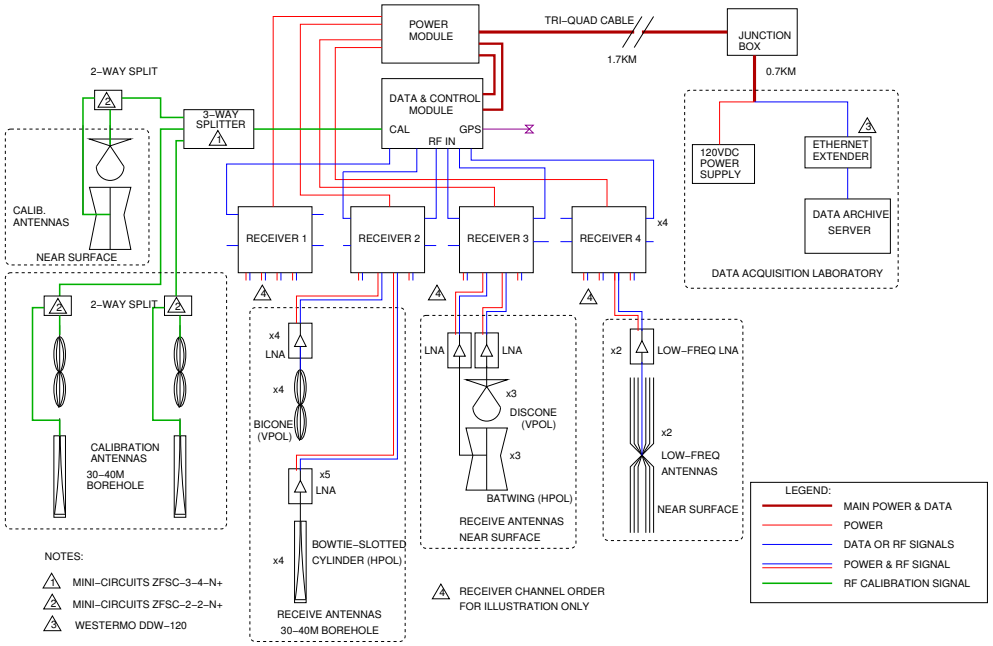

Figure 4.4.: Block diagram of the TestBed signal chain from [67].

Antennas

A number of different antenna types are used both down-hole and near the surface, the signals from which are used for triggering and digitisation in the DAQ box. The main constraints for the antennas are:

- Size limited to fit within a 15cm diameter borehole.
- Design must accommodate feed-through cables which pass signals from lower antennas to the surface.
- Azimuthally symmetric response.
- Sensitive to a single polarisation, VPol or HPol. This requires one design for VPol antennas and another for HPol antennas.

The feasibility of drilling 15cm wide holes to depths of 200m has been demonstrated and the antennas must fit within a cylinder of this size. To minimise the number of holes necessary per station multiple antennas are placed in each borehole, this means that the antenna design must be such that they allow for feed-through cables from lower antennas. Conductive coaxial cables are used and, in order to

maintain azimuthally symmetric response, must be fed through the centre of the antennas.

The necessity for azimuthal symmetry in response is driven by the desire to have no directional bias in sensitivity. Measurements of signal size and the ratio between orthogonal directions of polarisation are important tools with which to constraint event topologies. In order to make the latter measurements antennas must have good response in a single polarisation such that comparisons can be made between VPol and HPol.

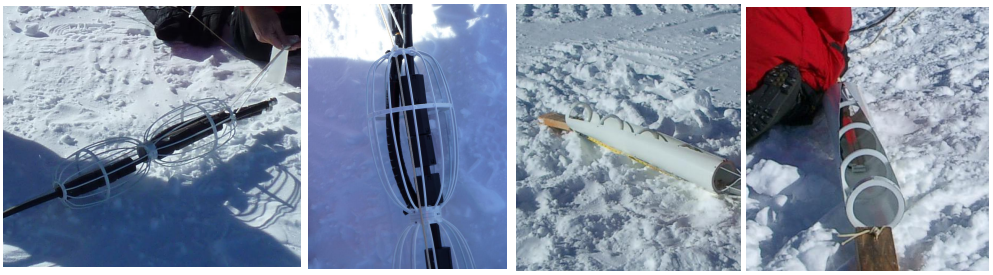

Figure 4.5.: ARA TestBed down-hole antennas from [67]. The left two images are of the bicone VPol antennas, the right two images are of the bowtie-slotted-cylinder HPol antennas.

The VPol antennas used in the boreholes for the TestBed have a wire-frame hollow-centre biconical design shown in Figure 4.5, where the feed region is annular around the pass-through cable. The HPol antenna design is significantly more challenging than that for the VPol antennas and two designs were implemented for testing in the TestBed: a bowtie-slotted-cylinder (BSC) antenna, and a quad-slotted-cylinder (QSC) antenna with internal ferrite loading to lower its frequency response. The design goal was to produce antennas to cover a frequency range of 150MHz to 850MHz. This was achieved with the VPol antennas but proved difficult for both designs of HPol antennas, which struggled to obtain the required response below about 200MHz to 250MHz in ice as shown in Figure 4.6. Although the performance of the QSC HPol antennas was measured to be significantly better than the BSC antennas the latter were deployed in the boreholes due to manufacturing constraints ahead of deployment.

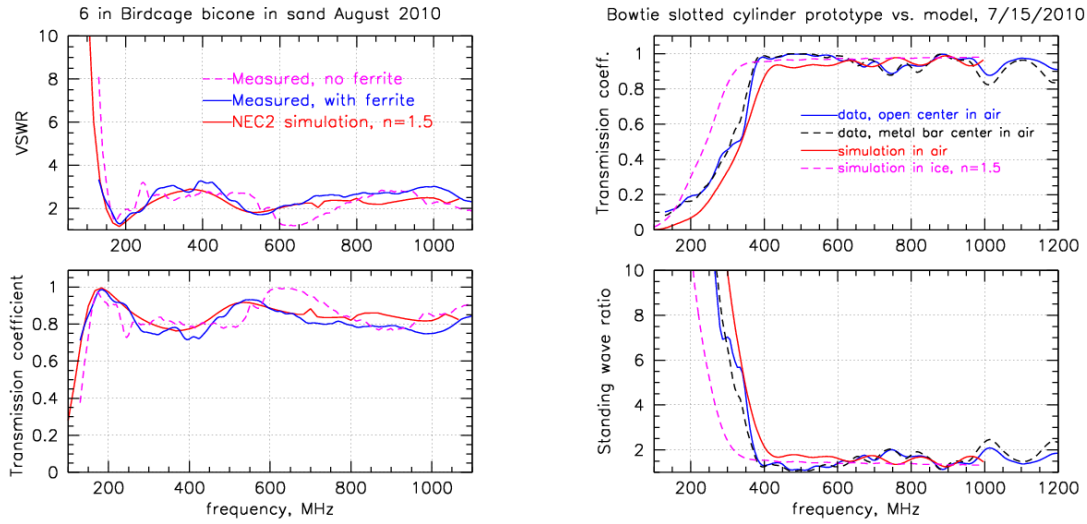

Figure 4.6.: Frequency response of (left) bicone VPol antennas and (right) the bowtie-slotted-cylinder HPol antennas from [67]. The equivalent power transmissivity as a function of frequency (bottom left and top right) shows the expected broadband response in both polarisations. The voltage standing wave ratio (VSWR) is also shown (top left and bottom right).

Amplification, Filtering and Transmission

One of the challenges of deploying receive antennas in the ice is to get signals to the DAQ box with little signal degradation. A number of design choices are made in the signal chain with this in mind. Expected signal sizes are very small and transmission distances from antenna to DAQ box $\sim$ 50m for the TestBed (and $\sim$ 300m for future stations) so several stages of amplification are used to boost the signals before their arrival at the DAQ box. Roughly 1m above the antennas signals are filtered and amplified before transmission to the surface. The filters define the band between 150MHz and 850MHz and notch out a particularly strong in band frequency at 450MHz used for South Pole communications. The output of the initial filters is sent through a low noise amplifier (LNA). The placement of the filters before the first stage of amplification serves two purposes: firstly this lowers the insertion of out of band thermal noise into the signal chain, secondly it prevents strong out of band signals, for example from south pole communications (450MHz), from saturating the

LNA. Large signals input to the LNA can cause it to be pushed into a non-linear mode of operation, leading to compression [2].

Signals are transmitted from down-hole to the surface using coaxial RF cables. Transmission in such a manner causes some signal loss and insertion of noise, however this is not a large effect in the TestBed due to the short distances involved. For future stations in which the transmission distance can be $\sim$ 300m these losses become a real challenge, leading to large decreases in the signal to noise ratio between the output of the LNAs and the surface. In Section 4.3 the modified signal chain for future ARA stations will be discussed.

Measurements, made in the laboratory ahead of installation, of the signal chain noise temperature and total gain are shown in Figure 4.7. The combination of filters and amplifiers contributes approximately 80K of noise, with a further 230K expected from the ambient temperature of the ice seen by the receive antennas. The total system noise is expected to be 310K, and is dominated by the surrounding ice. This leads to an input power of -85dBm at the LNA which is amplified by 80dB by the signal chain ahead of input to the DAQ box for digitisation.

Once a signal has propagated from the antenna through to the DAQ box the signal is split into two paths: one for triggering, and the other for digitisation. These paths are described in the following sections.

4.2.2. Triggering

The power and data rate restrictions from ARA's location mean that it is not possible to continuously record data from the receive antennas. This necessitates the use of a trigger to select interesting events that are then recorded by the digitisation electronics for later analysis. The design of such a trigger must take into account both the signal of interest and the expected background signals. The main backgrounds for ARA are from thermal noise and anthroprogenic signals. ARA's location was chosen to minimise the latter, which is expected to affect a reasonably small fraction of live time for the TestBed and nearly zero for future stations, which are further from noise sources and buried deeper in the ice. The trigger is therefore designed to

[2] Compression in RF amplifiers is a phenomenon whereby linearity of amplifier gain is lost. Large signals are amplified less than small signals leading to a loss of dynamic range

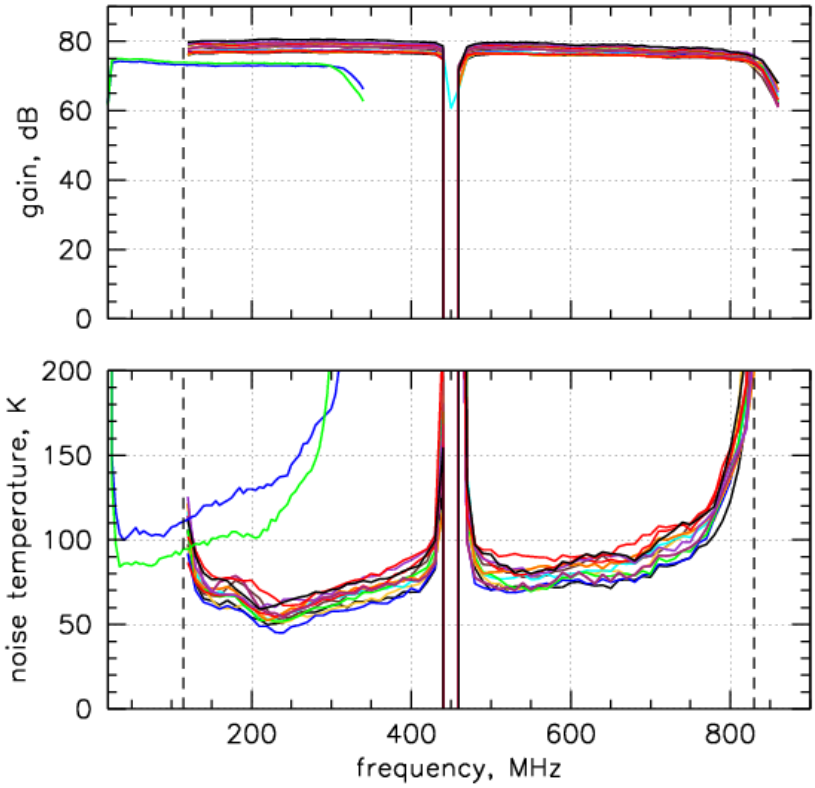

Figure 4.7.: Total gain (top) and noise figure (bottom) for the TestBed signal chain (preamplifiers and receivers) from [67]. The notch filter at 450MHz is clearly visible in both. The gain falls off at high frequencies due to the presence of a low pass filter at 850MHz. The two lines to the left on each figure (green and blue) are for the surface antenna signal chains.

distinguish between temporally correlated RF signals seen in multiple antennas, such as those expected from Askaryan radiation, and random thermal fluctuations, which are much less likely to exhibit temporal correlations. Such a trigger must have a high efficiency of passing neutrino candidate signals whilst rejecting the overwhelming majority of signals due to thermal noise. This is achieved by implementing a time coincidence trigger which requires that excess power is measured in multiple antennas within a short time window. When a trigger condition is met digitisation of the voltage-time waveforms from all the receive antennas is started. These waveforms, along with some additional information regarding the trigger condition and other housekeeping information, are referred to as an event and are centred around the time at which the trigger was asserted.

In the triggering chain of each antenna a coaxial tunnel diode is used to provide a uni-polar signal that is proportional to the RF power integrated over a few nano seconds. These signals are then fed into a discriminator which determines whether the RF power has crossed an an adjustable threshold. The discriminator is implemented via a field programmable gate array (FPGA) using low-voltage differential signal comparators. For each antenna that passes the defined threshold a one-shot [3] is generated in the FPGA, these signals can then in turn be assessed by firmware logic to determine if a trigger condition has been met. In the TestBed the trigger condition is that 3 of the 8 VPol and HPol borehole antennas produce signals that pass threshold within a 100ns window.

The TestBed coincidence trigger leads to a dependence on the individual threshold-crossing rates, or singles rates R_{single} of each antenna caused by thermal noise. In Appendix A it is shown that the rate at which the coincidence trigger is passed, R_{global}, is to first order given by:

$$R_{global} = NC_M^N R_{single}^N t^{N-1} \tag{4.1}$$

for N of M antenna coincidence in a time window t. In the case of the TestBed $N = 3$, $M = 8$ and $t = 100$ns. The choice of window size and number of coincident antennas, N, were informed from simulation of neutrino signals. The trigger efficiency

[3] the one-shot corresponds to a 'trigger passed' signal being asserted for a limited period of time

was measured in the laboratory as a function of signal to noise ratio (SNR) for an input impulsive signal and is shown in Figure 4.8.

For ARA stations limitations on data readout rate from the DAQ and volume written to disk constrains the maximum global trigger rate. Using Equation (4.1) it is possible to calculate the singles rate R_{single} associated with a particular global trigger rate R_{global}. Due to the differences in the response between channels the thresholds are set individually to achieve the desired contribution to the global trigger rate.

The single antenna rates and associated thresholds were read out and recorded by the DAQ at a rate of $\sim$ 0.5Hz. These thresholds, informed by singles rates, were regularly adjusted to ensure that the system triggered near to the desired rate during operation in 2011-2012. In future stations it is intended that the adjustment of thresholds will be handled by a feedback loop, or servo, which would monitor movements in the singles rates and adjust the thresholds accordingly to facilitate autonomous operation of the stations.

There are two types of trigger that cause event readout in the TestBed: 'RF triggers' which are those described above, and 'software triggers' which are taken at approximately 1Hz. Software triggers are instigated by the acquisition software controlling data taking in the DAQ box and are intended to provide a 'minimum bias' sample of events that profile the radio environment over the course of the year.

4.2.3. Digitisation

Signals are continuously sampled using a switched capacitor array until a trigger condition is met. At this point the FPGA initiates a freeze in the analogue sampling and those samples are digitised and readout, taking about 30ms. The sampling and digitisation is performed by 3 functionally identical LABRADOR (Large Analogue Bandwidth Recorder And Digitizer with Ordered Readout) digitisers designed and produced by the Instrument Design Laboratory at the University of Hawaii [68], and is described in detail in Section 5.1.1. The digitisers are operated at a sampling speed of 1GSample/s giving a Nyquist frequency [4] of 500MHz and consist of a switched capacitor array which track and sample the input voltage. The sampling speed is

[4]The Nyquist frequency for a discretely sampled system indicates the maximum resolvable input frequency and is equal to half the sample rate.

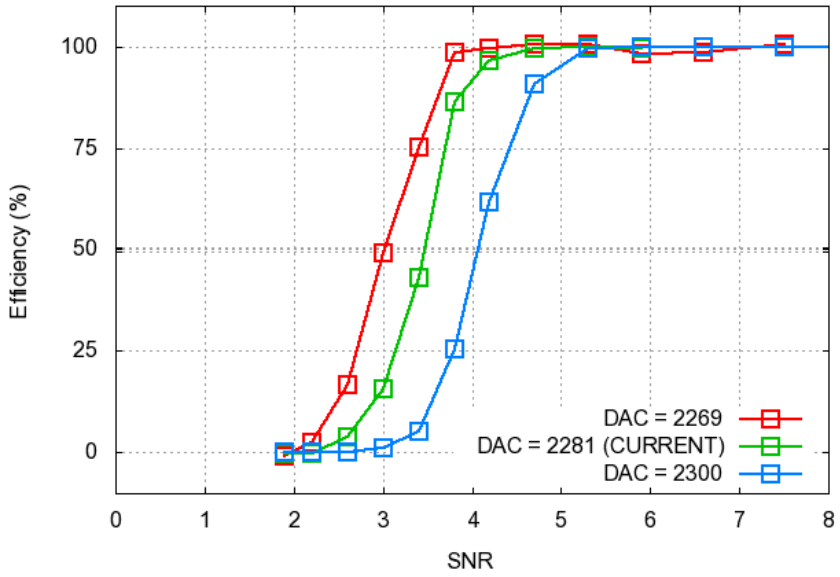

Figure 4.8.: TestBed trigger efficiency measured as a function of voltage signal to noise ratio (SNR) for an impulsive signal from [67]. The SNR is measured with respect to the RMS receiver voltage from baseline thermal noise. The three lines represent the efficiency measurements at different electronically set threshold values for the output of the tunnel diode power detector.

low compared with the maximum frequency expected from Askaryan signals, which can reach 1GHz. In order to make use of the high frequency response of VPol and HPol antennas used in the boreholes these channels were split and fed into two digitiser channels each, with a 0.5ns offset between the two. This interleaving leads to an increased effective sampling rate, with samples being taken every 0.5ns and pushing the Nyquist frequency up to 1GHz. In addition the sampled waveforms for these channels are the same length in time as those for other channels and allow for a readout window > 100ns. The net result is that 8 of the antennas deployed in the boreholes are effectively sampled at 2GSample/s, with the remaining antennas sampled at 1GSample/s.

The digitised waveforms consist of 256 samples per channel and are ~ 256ns long. For the interleaved channels there are 512 samples taken over the same time period. These are readout and packetised by the FPGA before being transferred via a USB connection to software running on a single board computer (SBC).

4.2.4. Calibration systems

In order to assess the functionality of the TestBed and provide an impulsive calibration source a calibration pulser was installed along with the TestBed detector. A GPS-synchronised Rubidium clock was used to trigger the calibration pulser that sends a ~ 250ps duration pulse to antennas deployed 40m from the centre of the detector. The output of the calibration pulser was fed into either a VPol or HPol antenna, of the same design as those deployed in the boreholes, illuminating the TestBed at a rate of 1Hz during data taking. The connection to the transmit antennas requires manual intervention to switch between polarisations, so the pulser was connected to the VPol antenna for the duration of 2011, then switched in the summer season to the HPol antenna for 2012.

4.2.5. Data acquisition and data transfer

Figure 4.9 shows a diagram of the TestBed DAQ, which is housed in a RF shielded aluminium box. The DAQ box sits in a wooden coffin buried in a pit under the surface of the ice. The output of the signal chain from each of the receive antennas is attached to a series of connectors on the outside of the DAQ box. A final stage of

filters is then applied to remove any out of band noise associated with transmission from the boreholes to the surface, before splitting into the digitisation and triggering chains and insertion to the TestBed electronics.

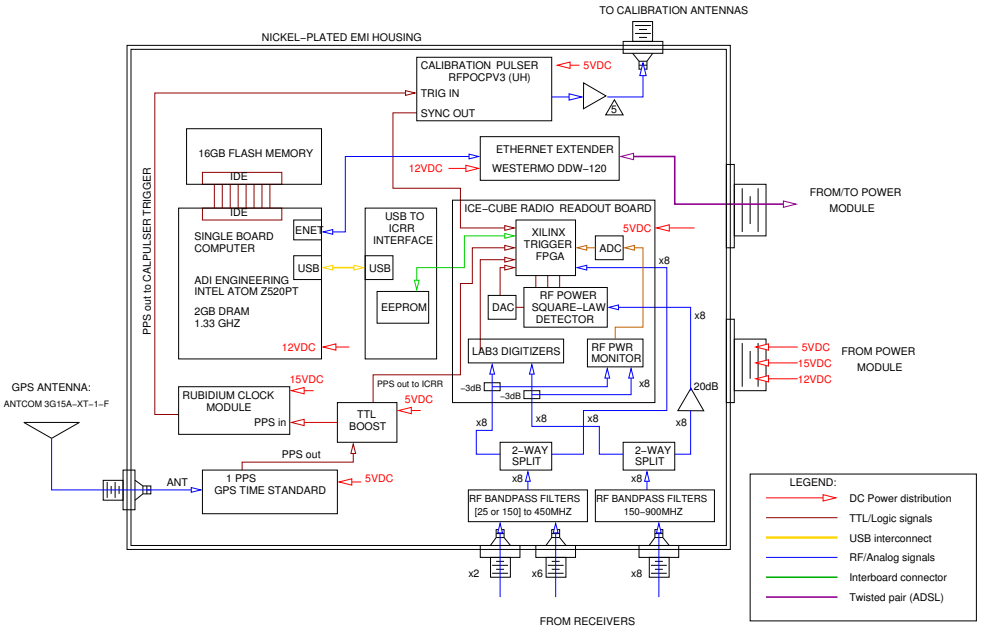

Figure 4.9.: Block diagram of the TestBed DAQ from [67].

The TestBed makes use of a custom digitisation and trigger board - the IceCube Radio Readout (ICRR), a custom USB readout board and a commercial SBC. The ICRR board is shown in Figure 4.10.

The TestBed station, and future stations, have their data taking and operation managed by dedicated software running on a SBC installed in the DAQ box. Due to the extreme conditions in Antarctica these SBCs are cold tested in a climate chamber in the laboratory (as with all of the electronics deployed) prior to deployment to ensure that they are able to function properly at very cold temperatures (the ambient temperature at the south pole falls below $-80°C$ during the winter months).

The data acquisition software is tasked with starting and stopping periods of data taking, or runs, periodically requesting software triggered events, reading out digitised events, reading housekeeping and sensor information and storing all of this data locally. The housekeeping and sensor data consists of useful information to assess the condition of the station and trigger, such as the single antenna trigger rates and ambient temperature. All of the data recorded by the SBC is temporarily

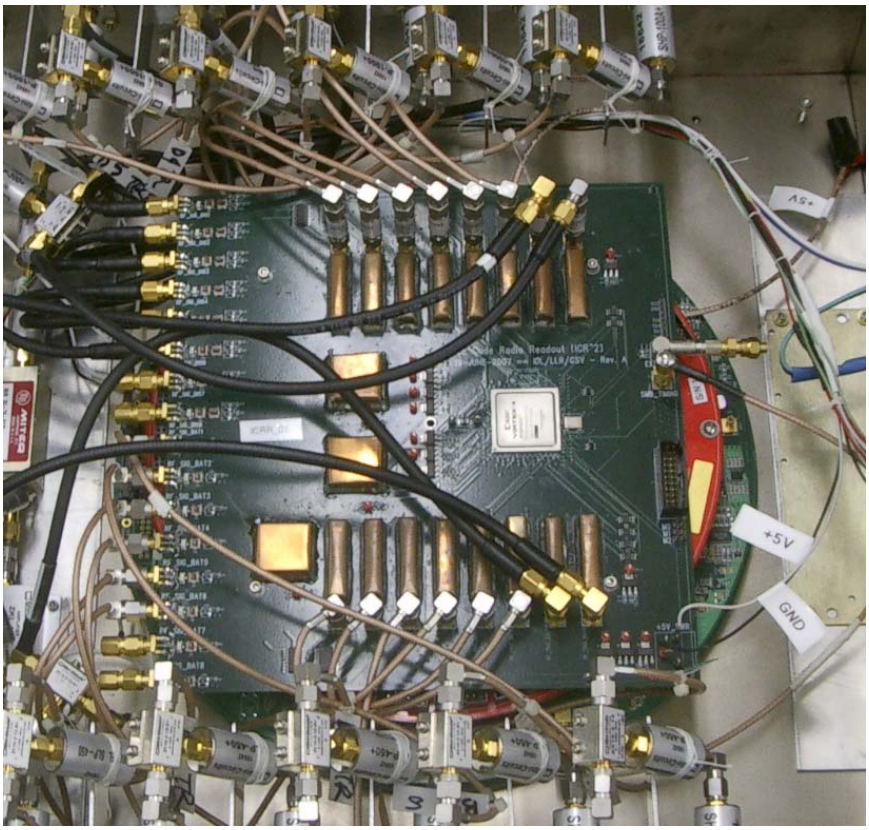

Figure 4.10.: The ICRR board from [69]. On the left the 16 digitisation chain inputs, right the 16 trigger inputs and centre the FPGA. The boards underneath are no longer used in the TestBed.

stored on a flash drive before being transferred to a dedicated server in the IceCube Laboratory. Data transfer occurs over twisted pair wires via an ethernet modem, which is then archived on the receiving server. Approximately $\sim 10\%$ of the data is selected to be sent north via a satellite link to provide a sample to be monitored to assess the detectors' functionality during the course of the year. The full data set was carried by hand on hard disks to a data warehouse in the north during each summer season, when it is possible to move personnel and equipment in and out of the South Pole Station.

4.3. ARA 1-3

The first full station, ARA1, was deployed in 2011-2012 at a depth of 100 m, with two further stations ARA2 and ARA3 deployed during the 2012-2013 season at a depth of 200 m. The updated design follows that illustrated in Figure 4.2 with 4 boreholes each containing 4 receive antennas, and a further 2 boreholes containing transmit antennas to illuminate the station with calibration signals. The updated specifications for the stations are summarised in Table 4.3 and can be compared with those in Table 4.1.

The successful deployment of antennas at much greater depth than the TestBed greatly reduces the ray-bending problems caused by the changing index of refraction in the firn layer of ice (approximately the first $\sim$ 150m of ice below the surface) [61]. This success, however, necessitates a different method of transmitting signals from down-hole to the surface, since the losses over these distances are large using RF coaxial cables. The output of the receive antennas are passed through the initial filters and LNAs as in the TestBed. The output of the LNAs undergo a further stage of amplification before being converted into an optical signal. This optical signal is then transmitted to the surface over fiber cables, which suffer from much lower losses and noise insertion levels than coaxial RF cables. Each borehole, hosting 4 receive antennas, has 1 Down-hole Transmission Module (DTM) unit that performs the RF to optical conversion and amplification. At the surface the signal is converted back to RF via a Fiber Optical Amplification Module (FOAM), which also further boosts the signal with $\sim$ 40dB of amplification, before the insertion into the DAQ box.

Specified Parameter	**Full Station**
Number of VPol antennas	8
VPol antenna type	bicone
VPol antenna bandwidth (MHz)	150-850
Number of HPol antennas	8
HPol antenna type	QSC
HPol antenna bandwidth (MHz)	200-800
Surface antenna type	fat dipole
Surface antenna bandwidth (MHz)	30-300
Number of surface antennas	4
Number of receive antenna boreholes	4
Borehole depth (m)	200
Vertical antenna configuration	VPol (HPol) above HPol (VPol)
Vertical spacing (m)	20
Approximate geometry	trapezoidal
Approximate radius (m)	10
Number of calibration antenna boreholes	2
Calibration borehole distance from centre (m)	30
Calibration hole geometry	facing two sides
Calibration signal type	impulse and noise
LNA noise figure (K)	< 80
LNA/amplifier dynamic range	30:1
RF amplifier total gain (dB)	> 75

Table 4.3.: ARA1-3 detector specifications.

A new DAQ system was designed and implemented for the full ARA stations based on the use of a new digitiser, again designed and supplied by the Instrument and Design Laboratory at the University of Hawaii. The Ice Ray Sampler (IRS2) application specific integrated circuit (ASIC) was designed to address the design goals of ARA. The IRS2 features higher sampling speeds across all of the input channels (3.2GSample/s) to make use of the high frequency response of the deep antennas and to better resolve the high frequency content in Askaryan signals. The IRS2 also features a large analogue sample buffer, approximately 10µs long, which allows for increased trigger time windows than the LABRADOR digitiser used in the TestBed. This is an important feature since the transit time of RF signals across the detector increases significantly from the TestBed to the new stations, due to the increased distance between receive antennas. This buffer is also deep enough to enable future multi-station triggers, which could be used to further determine event topologies.

The triggering systems were improved to offer finer adjustment of threshold levels for the individual antennas that contribute to RF triggers. For the TestBed the threshold tuning was a time consuming process performed by a user over low bandwidth satellite links to the South Pole Station. The increased granularity of threshold setting was used to implement automatic threshold setting via a feed-back, or servo, loop. The individual antenna trigger rates are read out by the DAQ at 1Hz and used to inform new threshold settings. The continuous adjustment of the thresholds means that the DAQ is able to operate with long-term stability of RF trigger rate, without the need for regular user intervention.

The calibration systems were also changed significantly. Two calibration pulsers were deployed, one in each of the calibration boreholes. Each of the two pulsers are attached to a pair of antennas, 1 VPol and 1 HPol, providing 4 antennas with which to illuminate each station. The DAQ is linked to these pulsers via communications cables that allow switching between the output antennas and attenuation of the pulser output. These features were integrated into the DAQ software to allow calibration pulser settings to be set at the start of each run of data taking, allowing much greater flexibility than in the TestBed, where the setup could only be changed by physical means during the austral summer.

Details of work undertaken in developing and installing the updated DAQ systems for ARA1-3 are discussed in Section 5.2.

Chapter 5.

Calibration and data acquisition development

Development and calibration of the data acquisition systems forms a crucial part of building a detector sensitive to neutrino signals that could make the first measurements of ultra high energy (UHE) neutrino interactions in the antarctic ice. Askaryan radiation from UHE neutrino-induced particle cascades results in short duration radio frequency (RF) signals. These impulses typically extend over < 100ps and result from coherent emission of frequencies up to 1GHz. In order to make measurements of such signals it is not only necessary to use antennas with the necessary bandwidth, but also to be able to sample and digitise their output with appropriate fidelity.

This section will describe the need for, and details of, a number of timing calibrations that were performed on the digitisation chain in the TestBed. In addition to optimising the TestBed detector for data analysis, work was undertaken in developing and installing the first 3 full ARA stations. These stations will eventually form part of a much larger array of sub-detectors that is projected to reach world leading sensitivity to UHE neutrino fluxes.

5.1. TestBed calibration

5.1.1. TestBed timing calibrations

The majority of recorded signals that trigger the ARA stations are expected to be from thermal or anthropogenic sources and are considered backgrounds to detection of neutrino signals. Timing differences between signals being observed in multiple antennas can be used to reconstruct the physical location of their source. In Section 6 it will be shown that reconstruction can provide a powerful tool with which to separate interesting signals from these backgrounds. The need to reject backgrounds in this manner, combined with the signal shape expected from Askaryan radiation, sets a requirement for large bandwidth, high sampling rate digitisation of the signals received in the detector's antennas. The remote location and associated power restrictions preclude the use of commercially available digitisers and necessitate the use of an application specific integrated circuit (ASIC) digitiser.

The TestBed digitises waveforms using three LABRADOR (Large Analogue Bandwidth Recorder And Digitizer with Ordered Readout) ASIC chips designed at the Instrument and Design Laboratory, University of Hawaii [68]. Each chip has 9 parallel RF input channels that are continuously sampled. These samples are digitised in place once read out is initiated by the data acquisition (DAQ). Sampling, digitisation and readout of samples is controlled by signals from a field programmable gate array (FPGA) in the DAQ. A number of timing related calibrations are necessary due to the properties of the LABRADOR chip and its implementation in the TestBed.

Figure 5.1 shows the architecture of the LABRADOR chip used in the TestBed. On each chip a switched capacitor array (SCA) of 260 sampling capacitors per channel is used to sample the input voltages. Sampling is controlled by a write signal, which is common to all 9 channels, driven by an externally provided clock. When a stop signal is asserted by the DAQ the samples are digitised in place using a Wilkinson analogue to digital converter (ADC), which occurs in parallel for all samples on all channels. Once digitisation is complete the digitised voltage values are requested by the DAQ and read out in sequence on a 12-bit data bus. The FPGA then processes these values for packetisation and recording as an event for later analysis. The digitisation and readout stage takes $\sim 50\mu s$ in the LABRADOR

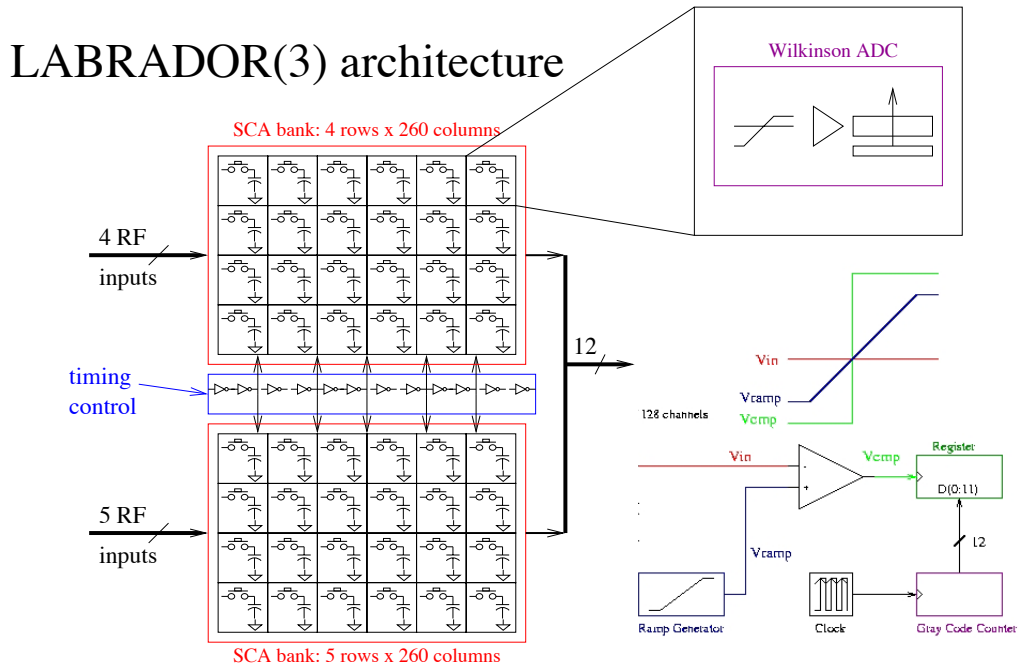

Figure 5.1.: The architecture of the LABRADOR digitiser from [68]. 9 RF input channels are sampled in parallel using a common timing control.

chip. This means that the settle time of the analogue voltage samples in the SCA has a negligible affect on the readout values.

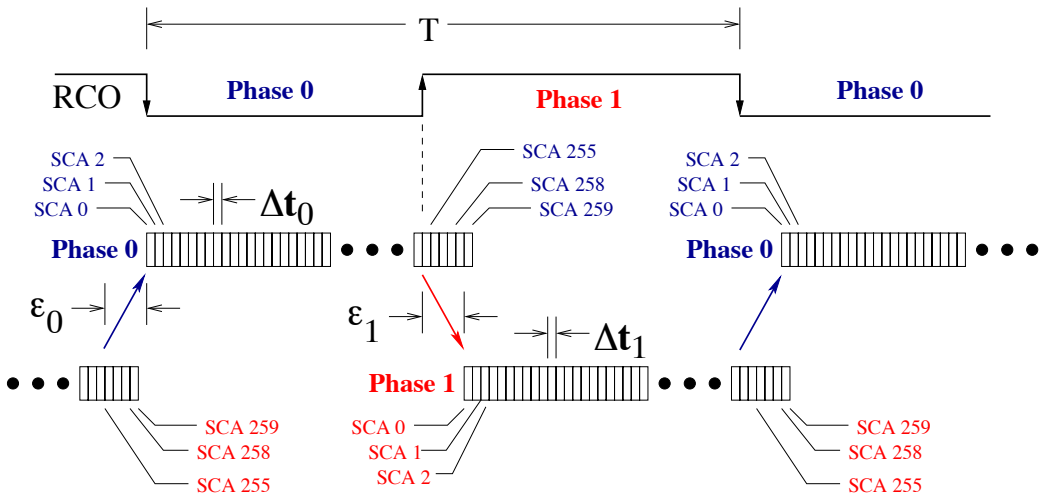

Figure 5.2.: A schematic of the sampling within the LABRADOR digitiser from [68].

To achieve continuous sampling, voltage samples are stored in the SCA as the write signal propagates along the array. The time taken for the write signal to

propagate between adjacent samples varies from sample to sample and from digitiser to digitiser. These 'inter-sample' times are estimated from known input signals and a calibration applied to correct the time base of recorded waveforms. The voltage sampling is illustrated in Figure 5.2. The write pointer signal propagates along the array in one of two phases, Ripple Carry Out (RCO) 0 and RCO1. The transit time of the write signal between adjacent samples differs between these two phases, hence calibrated inter-sample times must be calculated for each.

In order to continue sampling once the write signal reaches the end of the SCA the write pointer wraps back to the first sample, at which time the RCO phase changes. This wrapping means the SCA functions as a ring, continuously sampling voltage values and over-writing the oldest sample taken. The wrapping of the write pointer is illustrated in Figure 5.3. When a stop request is received, to initiate digitisation and read out, the chip issues a 'HITBUS' causing the chip to stop sampling. Due to the changing state of the chip as they are latched the HITBUS removes around 8 samples from the waveform as the stored values are inaccurate. The capacitor immediately after the HITBUS is the earliest sample in the resulting waveform, and the capacitor immediately before the HITBUS is the latest sample.

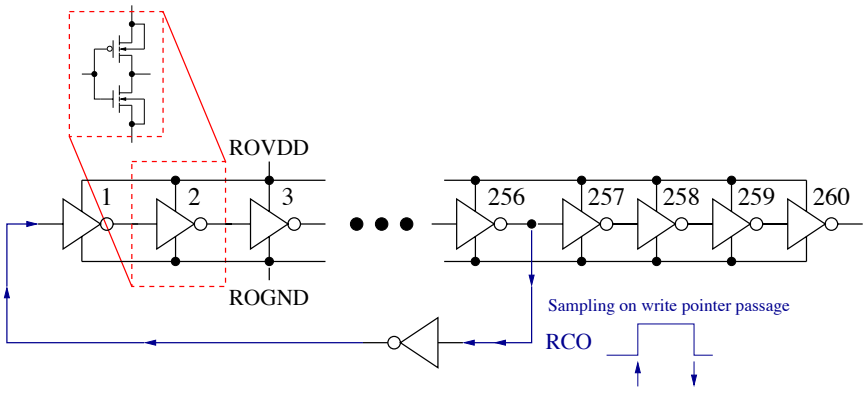

Figure 5.3.: A schematic the write pointer wrap around in the LABRADOR digitiser from [68]. As the write pointer returns to the first sample additional tail samples are taken in order to avoid a gap in sampling.

Although the physical distance travelled when wrapping the write pointer is small the need to change phase causes a much larger time delay than between adjacent samples. To avoid time gaps in the resulting waveforms, the write pointer begins wrapping at the 255^{th} sample. The remaining 4 tail samples, 256-259, are then used to fill in the resulting gap in time between the 255^{th} and 0^{th} samples. This is

illustrated in Figure 5.3. The time taken for the write pointer to wrap, referred to as the 'wrap-around', is another quantity that needs to be estimated and accounted for in calibration, which again varies from chip to chip and between RCO phases.

One of the features of the LABRADOR chip is the ability to operate at a range of sampling speeds. For the TestBed implementation it is operated at a nominal sampling speed of 1GSample/s which corresponds to a Nyquist frequency of 500MHz. The borehole antennas have good response to frequencies up to 850MHz, which means that frequency content would be lost if these antennas had their signals input to a single channel each. In order to retain this high frequency content each of these antennas has its signal split and inserted into a pair of input channels, one offset from the other by 0.5ns. This interleaving process leads to an effective doubling of the sampling speed for the borehole antennas, however the offset needs to be calibrated out in order to accurately reconstruct the original input waveforms.

As well as ensuring the timing information, on a per antenna basis, is as accurate as possible it is also necessary to correct for any offsets between channels. Due to the use of a common write signal each of the channels on a single chip will be aligned once the inter-sample, wrap-around and interleave calibrations are carried out. This is not true for channels on different chips and must be addressed in the final timing calibration, inter-chip jitter removal. Since timing differences between similar signals in pairs of antennas serves as the basis of event reconstruction this calibration is an integral part of distinguishing between noise sources and potential neutrino induced signals.

Inter-sample times

The inter-sample times are the time elapsed between recorded samples of an input signal, corresponding to the time taken for the LABRADOR digitiser's write pointer to propagate between samples. By adjusting control voltages to the chip the global sampling speed can be adjusted, which acts by effecting a change to each of the inter-sample times. Using the selected operating parameters of the chips in the TestBed DAQ a series of dedicated calibration runs were taken in the laboratory prior to deployment in Antarctica. Precise frequency sinusoidal inputs were provided to the DAQ box and a series of runs taken, comprising of many thousands of events

The inter-sample times were calculated by measurements of these sine waves, input to each RF channel on the chip, using a method informed from [70]. In this technique the number of times an input sine wave crosses zero voltage between a pair of adjacent samples is counted. The time between a pair of samples can then be estimated using:

$$\langle \Delta t_{i,i+1} \rangle = \frac{T_{input} \langle NZC_{i,i+1} \rangle}{2N_{i,i+1}} \tag{5.1}$$

where $\langle \Delta t_{i,i+1} \rangle$ is the estimated time between sample i and sample $i+1$, $NZC_{i,i+1}$ is the number of times zero voltage was crossed between these samples, $N_{i,i+1}$ is the number of events used and T_{input} the period of the input sine wave. Since the time between adjacent samples differs between the two RCO phases separate estimates are made for each phase. This method can be applied to each pair of samples as follows:

- The voltage values of each waveform are zero-meaned
- The RCO phase of each pair of adjacent samples is then determined as described in Section 5.1.1
- The total number of times zero is crossed between each pair of samples in each RCO phase is counted over the run (i.e. the voltage goes from positive to negative, or visa versa, between a pair of samples)
- The total number of zero-crossings is then divided by the total number of events for which that pair of samples were active in each RCO phase
- This fractional zero-crossing occupancy ($\frac{NZC_{i,i+1}}{N_{i,i+1}}$) is then scaled by the half-period of the input sine-wave according to Equation (5.1), resulting in an estimate for the inter-sample times

The resulting distribution of inter-sample times is shown in Figure 5.4.

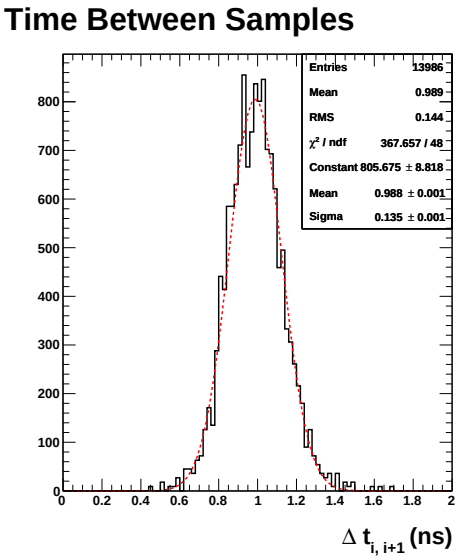

(a) All RCO phases

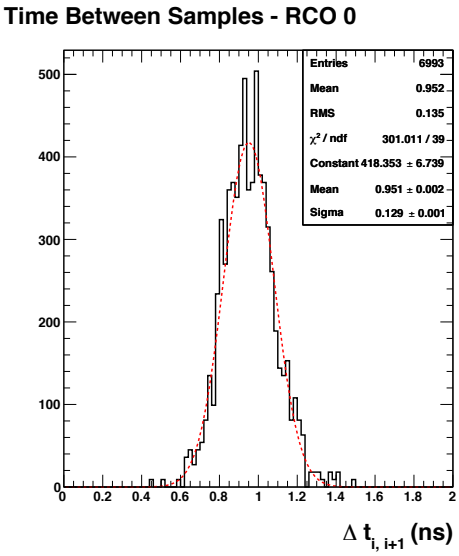

(b) RCO 0

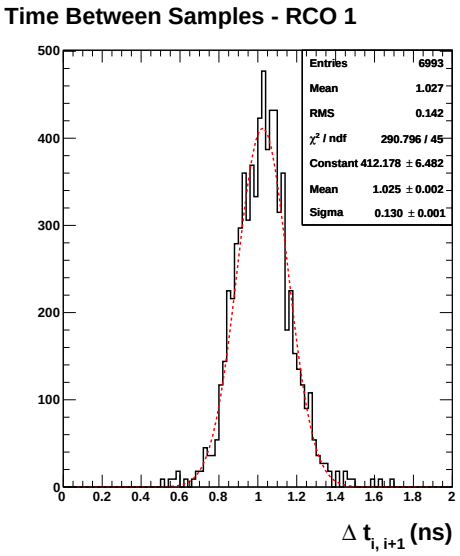

(c) RCO 1

Figure 5.4.: Calculated inter-sample times for all 3 TestBed LABRADOR digitisers. In (a) both RCO phases are included, whereas (b) and (c) show the distributions for RCO phase 0 and 1 respectively.

RCO phase determination

A complication to the calibration procedure is that the RCO phase of the LABRADOR chip at the time of a trigger being asserted is not always recorded correctly. When a trigger occurs and a stop request is sent from the DAQ to the LABRADOR digitisers the RCO phase at this point in time is read out along with the event. In the case that the latest sample in the waveform is in the first 20 samples in the SCA the RCO phase was found to be inaccurate. Since the average sampling speed differs by $\sim 8\%$ between the two RCO phases it is possible to estimate which phase the chip was in using a reference clock that is inserted into 1 of the 9 RF inputs on each chip. The average period of the clock is calculated for both possible RCO phases, and the one that is closest to the nominal period is selected and assigned to the waveforms digitised by each chip.

Wrap-around times

The time taken to wrap between the end and beginning of the SCA, denoted ε_0 and ε_1 for the two RCO phases, is considerably larger than the inter-sample times. Rather than measuring the scaled zero-crossing occupancy, sine waves are fitted to the waveform before and after the wrap-around. By aligning the phase of the two fitted sine waves it is possible to estimate the wrap-around time. The full procedure is as follows:

- The voltage values of each waveform are zero-meaned
- Only events containing at least 20 samples before and after the wrap-around are selected
- Calibrated inter-sample times, calculated as described in Section 5.1.1, are applied to the waveform
- Sine waves are fitted to the last 20 samples before the wrap-around, and first 20 samples after
- The fitted sine wave after the wrap-around is moved in time until the phase matches that before the wrap-around

- The wrap-around time is then taken to be the time between sample 255 and sample 0, and is averaged over a calibration run

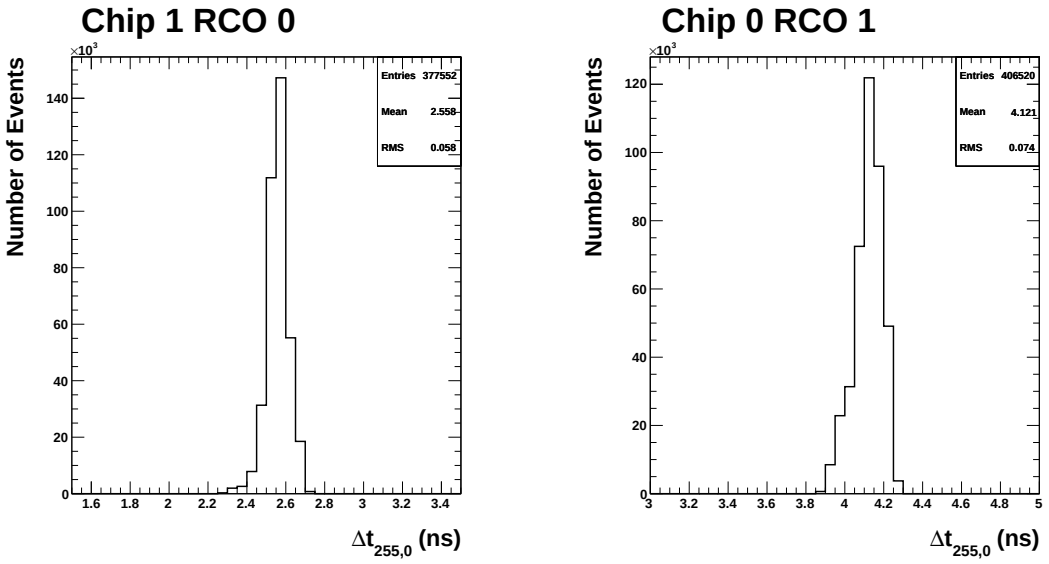

(a) RCO 0 **(b)** RCO 1

Figure 5.5.: Estimated wrap-around times for all events in a calibration run for TestBed LABRADOR chip 0 and both RCO phases. The average value is taken over all events in the run.

Interleaving

To increase the effective sampling rate whilst maintaining the time window of the borehole antennas the signals from each are fed into two digitiser channels with one delayed relative to the other by ~ 0.5ns (approximately half a sample) [1]. The resulting waveforms are then produced by interleaving the two digitised channels for each of these antennas in the offline software. The 8 deep antennas have their signals split in this way across pairs of channels on the first 2 chips. The third chip is operated in the normal manner with each channel supplied by a single antenna.

A calibration is performed to estimate the offset between each pair of channels that have their signals split, such that the resulting interleaved waveforms are as

[1] Since the digitisation of stored voltage values takes a relatively long time ($\sim 50\mu$s) the settle time of the individual samples has no influence on the digitised values.

accurate as possible. This is achieved by attempting to align the phase of input sine waves between these pairs of channels as follows:

- The voltage values of each waveform are zero-meaned
- Inter-sample and wrap-around calibrations are applied to each channel's waveforms
- A sine wave is fitted to each pair of split channels
- The phase offset between each pair of fitted sine waves is recorded
- The interleave offset for a pair of channels is taken to be the mean phase offset over a run of calibration data

The interleave offsets for two pairs of channels are shown in Figure 5.6. The calculated offsets are close to half the average inter-sample times resulting in an increase in Nyquist cut-off from 500MHz to 1GHz for the borehole antennas.

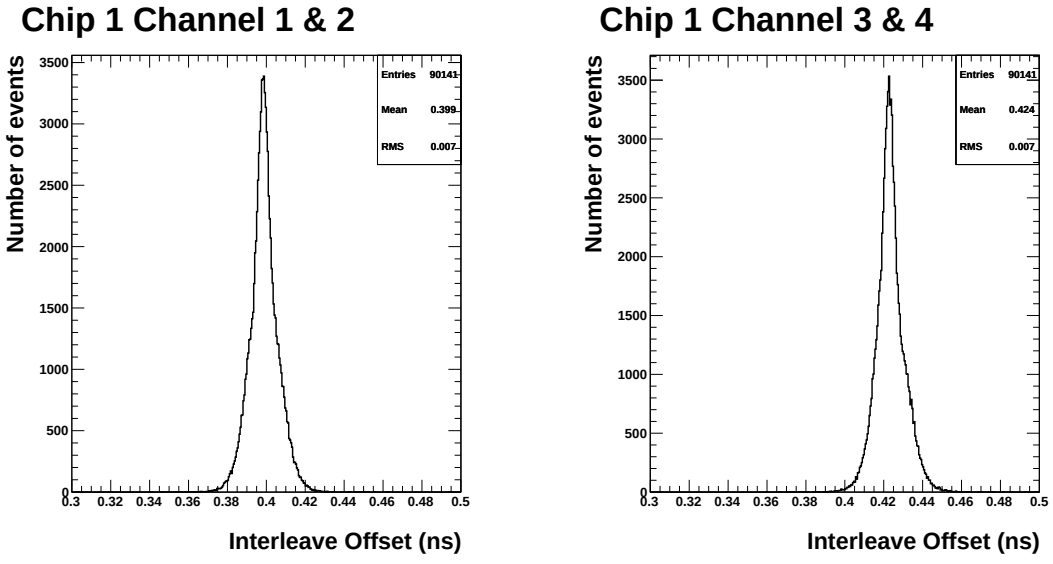

(a) Chip 1 Pair 1 **(b)** Chip 1 Pair 2

Figure 5.6.: Calculated offsets between 2 pairs of channels on LABRADOR chip 1 in the TestBed.

Inter-chip jitter correction

The three LABRADOR digitiser chips used in the TestBed are operated in parallel using the same control signals. Each of the 3 digitisers has a clock signal input into one RF channel that is digitised along with the inputs from the signal chain. These clock signals are used to remove event-to-event jitter of the trigger signals that start digitisation and readout of the SCAs. Using the clock channel it is possible to estimate the phase of the clock signal in each LABRADOR chip. A timing correction is calculated to align the phase of the clocks across all three digitisers and is then applied to all waveforms.

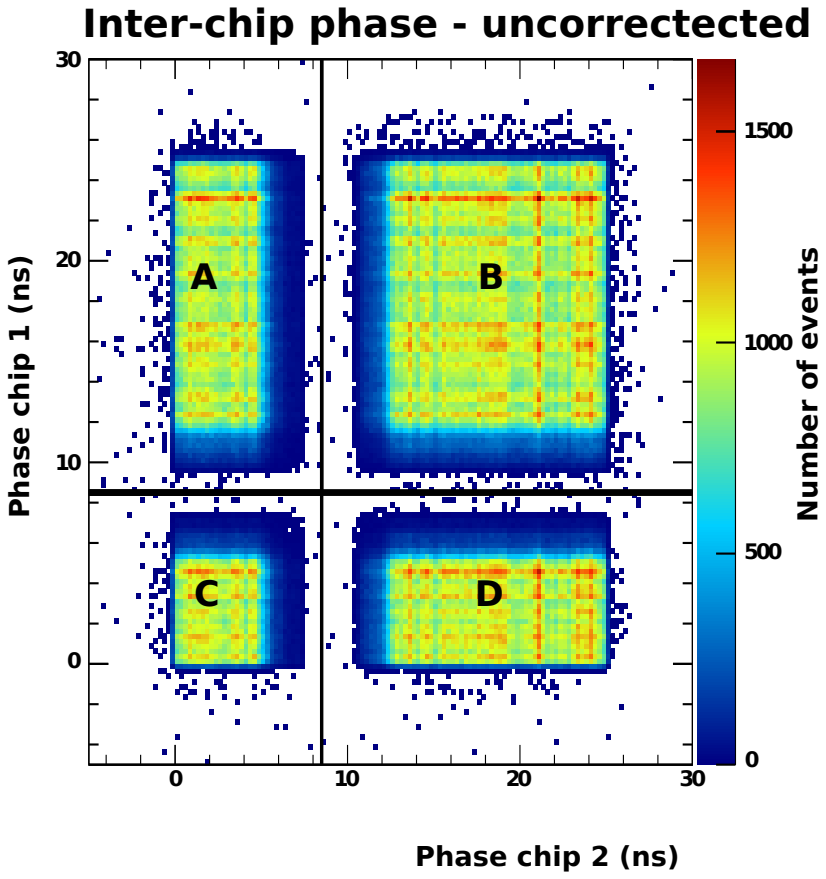

Figure 5.7.: The measured clock channel phase relative to the first sample in LABRADOR chips 1 and 2. Events in region A are interpreted as needing 25ns (1 clock period) added to the phase in chip 2, in region D 25ns added to the phase in chip 1 and in region C 25ns added to the phase in both chips 1 and 2.

In Figure 5.7 the measured phase of the clock channel, relative to the first sample, in two LABRADOR chips are plotted against each other. A feature of the TestBed implementation of the LABRADOR digitiser is that there are regions in this parameter space that are under-populated [2]. This feature is used to resolve period ambiguity between the two chips which would otherwise result in an additional ±25ns (1 clock period) offset between the time base in waveforms recorded in different LABRADOR chips.

Calibration results

In Section 6.4.1 a method will be described for reconstruction which is reliant upon precise measurements of the timing offsets between the same signal being observed in multiple antennas. The cumulative effect of the timing calibrations described can be seen in Figure 5.8 where timing offsets between a pair of antennas are calculated using correlation techniques.

The timing calibration with the largest effect is the inter-chip jitter removal, which corrects for offsets up to 25ns between waveforms recorded on different LABRADOR chips. A shift in the measured mean time offset between the antennas' waveforms is caused by application of cable delays, which account for the differing lengths of cable that signals must propagate through when travelling from the receive antenna to the digitiser. After application of all timing calibrations the resolution on time offsets between antennas is $\sim$ 100ps.

5.1.2. Voltage reversal

Close to finalisation of the analysis described in Chapter 6 a reversal of voltage values recorded in two of the horizontally polarised (HPol) antennas' waveforms was identified. This reversal is illustrated in Figure 5.9 where hundreds of calibration pulser signals originating from the HPol calibration pulser are averaged to remove thermal noise. There is a clear inversion of voltage values for antennas 3 and 4 compared with antennas 1 and 2. It is believed that this flip originates from

[2]This feature in timing offset between channels is believed to be due to a firmware mechanism that attempts to align the clocks between LABRADOR digitisers.

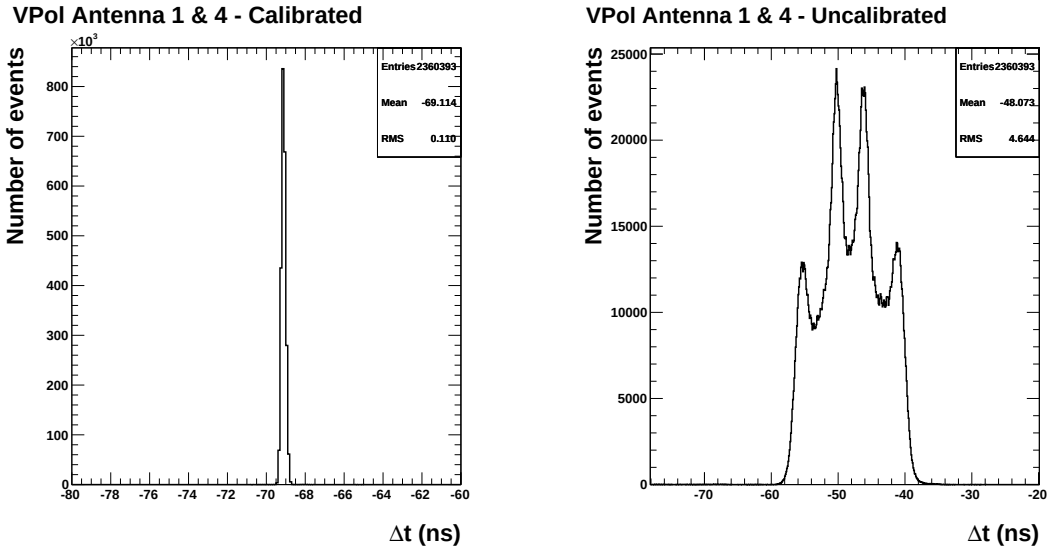

(a) Calibrated waveforms **(b)** Uncalibrated waveforms

Figure 5.8.: Timing offsets for a month of calibration pulser events for (a) timing calibrated and (b) uncalibrated waveforms. The antennas used are vertically polarised (VPol) receiving signals from a calibration pulser connected to a VPol transmit antenna buried close to the TestBed detector.

installation of the antennas upside down relative to antennas 1 and 2, or from inversion of voltages during amplification in the signal chain.

Checks were made on the vertically polarised (VPol) antenna's waveforms from VPol calibration pulser events. No such inversion was found for these antennas. A correction was applied in the analysis described in Chapter 6 that switches positive for negative (and visa versa) voltages for the affected HPol antennas, resulting in a significant improvement in the reconstruction of calibration signals. The effect on measured timing offsets for calibration pulser waveforms is shown in Figure 5.10, where the application of a correction for the voltage flip results in the removal of a second peak in the distribution.

5.2. ARA 1-3 DAQ development

During this PhD work was carried out developing the DAQ for the first 3 full ARA stations, ARA 1-3. The DAQ installed for the TestBed differs significantly from

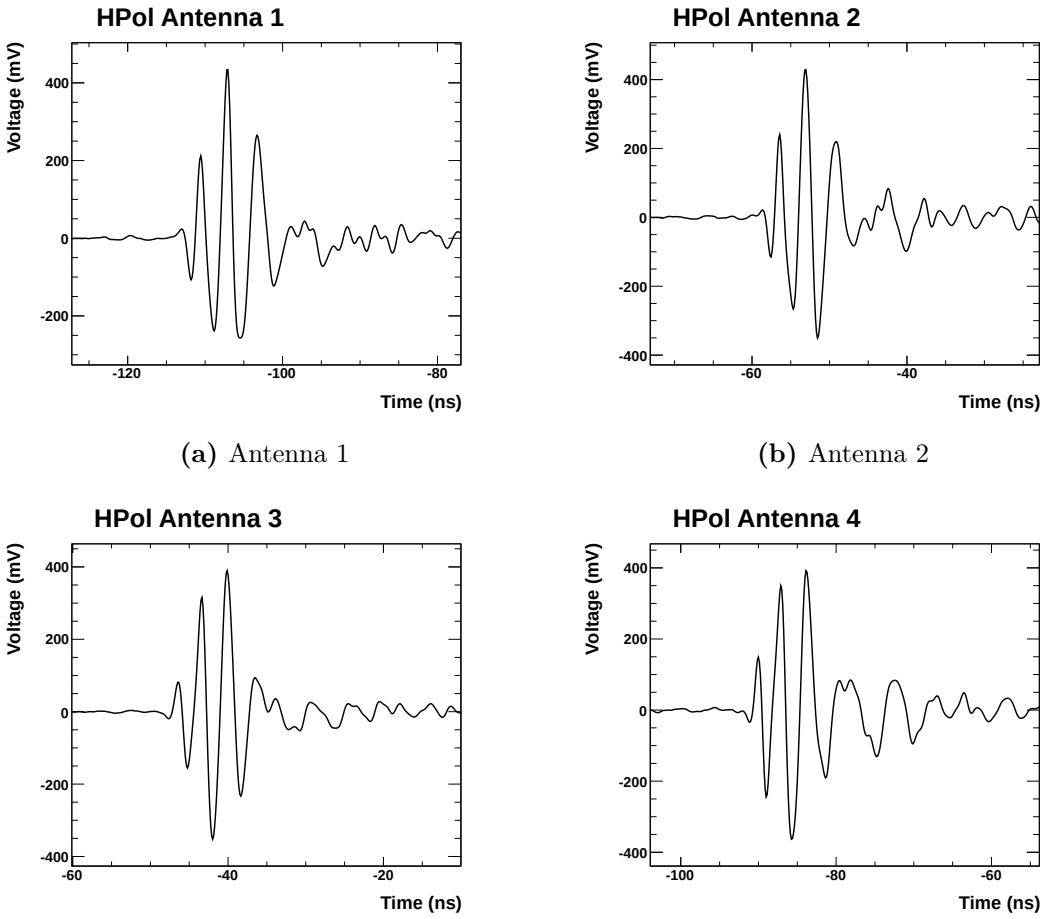

(a) Antenna 1

(b) Antenna 2

(c) Antenna 3

(d) Antenna 4

Figure 5.9.: Averaged waveforms from HPol borehole antennas. The signal captured is from the HPol calibration pulser operated in 2012. There is a clear inversion of voltage values in antennas 3 and 4 compared with antennas 1 and 2. The differing positions in time of the peak voltage is due to the position of the calibration pulser. This relative timing can be used to reconstruct the source location.

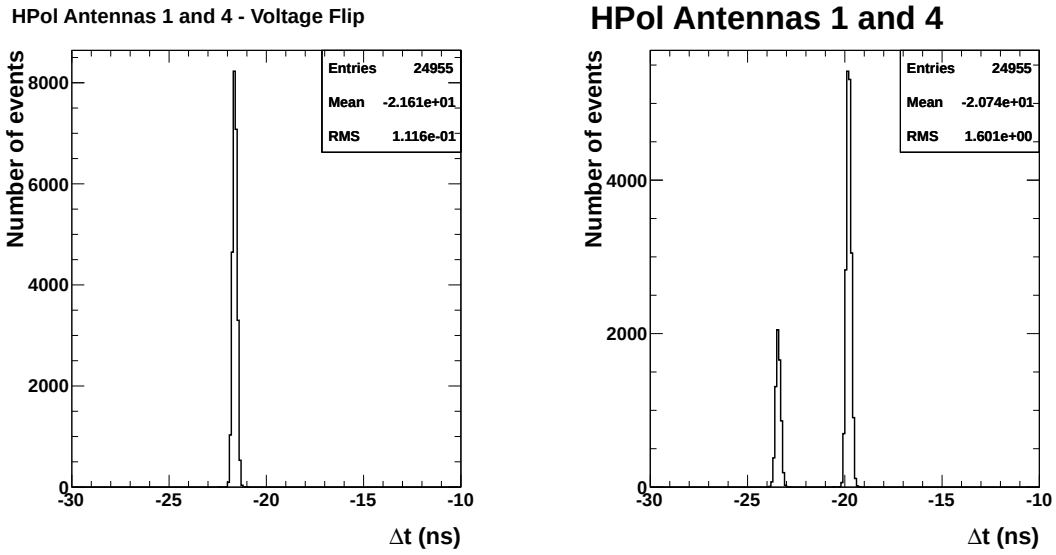

(a) With voltage reversal correction

(b) Without voltage reversal correction

Figure 5.10.: The offset between HPol calibration pulser signals recorded in HPol antennas 1 and 4. The offset is measured by taking the time-offset corresponding to maximum correlation between the received signals. The offset is shown with (a) and without (b) correction for the voltage reversal found in antennas 3 and 4.

that developed and deployed for ARA 1-3. The most prominent change is the use of a new digitiser, the Ice Ray Sampler (IRS2) [71], in place of the LABRADOR digitiser. The IRS2 features increased sampling speed, up to 3.2GSample/s, and a large analogue sample depth of $\sim$ 10µs. These features mean that higher fidelity waveforms can be recorded and trigger conditions can be made on a much longer time-scale than in the TestBed. To accommodate the new digitiser a new suite of electronics boards were also designed and developed, along with associated firmware and software to manage data taking and event read out.

5.2.1. IRS2 testing and calibration

The IRS2 provides 8 radio frequency (RF) inputs and is operated at a sampling speed of 3.2GSample/s, placing the Nyquist frequency well above the maximum anticipated from neutrino induced signals. In addition to the increased sampling speed the IRS2 features analogue buffering of sampled signals. The IRS2 contains 2 sets of 64 capacitor sampling cells (referred to as 'blocks') along with 512 analogue storage blocks for buffering sampled signals. Once every 20ns one of the 2 blocks of sampling cells is transferred to the storage buffer whilst sampling of the RF signal continues in the other block. When a trigger condition is met the relevant analogue storage blocks are requested, digitised and read out, whilst sampling and buffering to the remaining storage blocks continues.

A number of iterations of the IRS digitiser were tested in the laboratory to ascertain their functionality and suitability to use in the new ARA stations. The revision used in ARA 1-3, the IRS2B, was selected as the best functioning among the chips tested. Along with finding, setting and testing the optimal operating parameters for the IRS digitiser preliminary calibration work was undertaken to correct the time base of the waveforms digitised in the IRS2 chip. The results of the latest iteration of timing calibrations are illustrated in Figure 5.11, where an uncalibrated and calibrated waveform are shown for a sine wave input into one of the IRS2 channels.

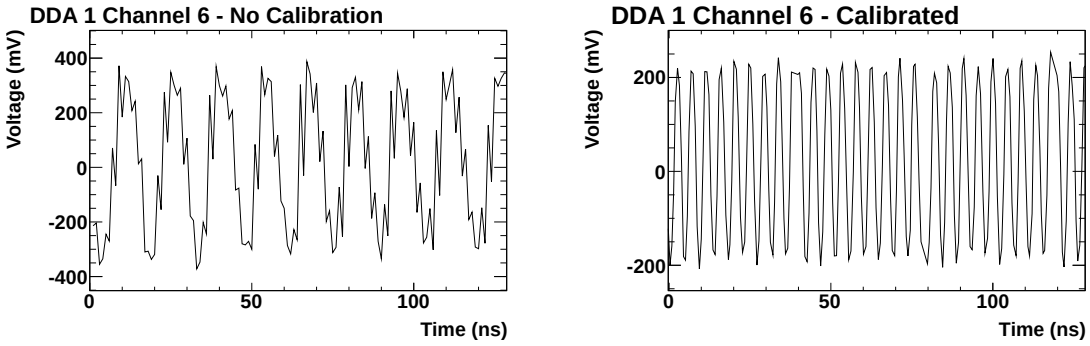

(a) Uncalibrated waveform **(b)** Calibrated waveform

Figure 5.11.: The results of applying the latest iteration of timing calibrations to a 215MHz sine wave input to a channel on the IRS2 chip. The timing calibration was produced by Thomas Meures.

5.2.2. Other DAQ work

Testing and development work was undertaken ahead of deployment for ARA 1-3. Among other activities this included:

- Development and integration of offline software to accommodate new data formats from the ARA 1-3 DAQ systems
- Writing the event read out format and development of the online data acquisition software used to control data taking in the DAQ
- Writing and testing firmware implemented on a USB micro-controller that provides the data connection and control interface between the single board computer, running the online acquisition software, and the field programmable gate array (FPGA) that controls digitisation and triggering in the DAQ
- Testing and developing the trigger subsystems

During the 2012-2013 austral summer I travelled to the South Pole to install and commission the DAQs for ARA 2 and ARA 3. After installation of the new systems in the ice work was undertaken in setting up the data flow from each station to a server located in the IceCube Laboratory. Once data flow was established systems were put in place to allow a subset of data to be transferred via satellite to a data warehouse in the northern hemisphere.

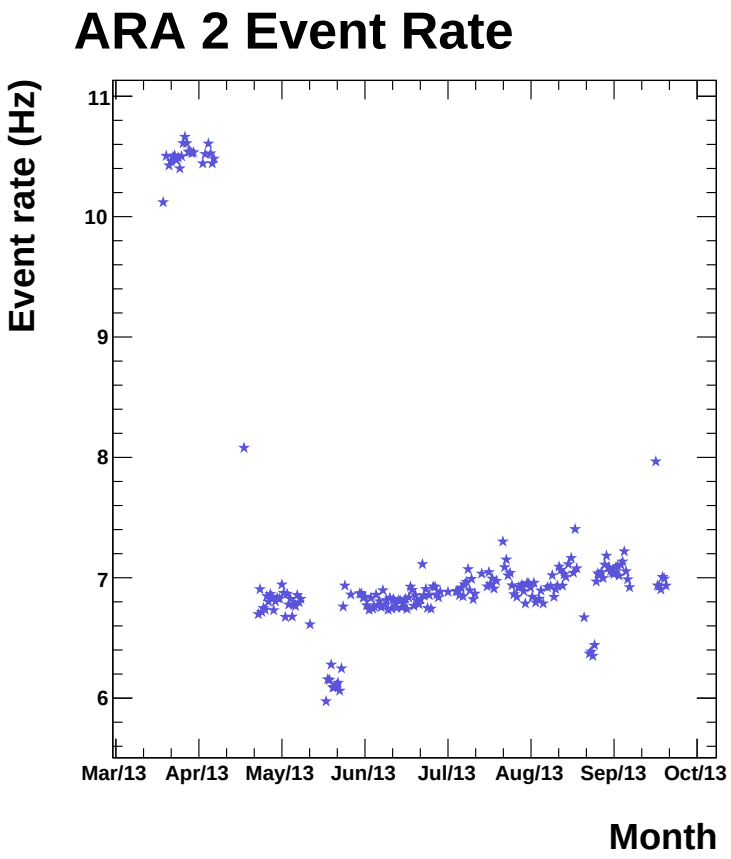

Figure 5.12.: Event rate in ARA2 during 2013.

Upon completion of the installation and deployment activities systems were developed to allow monitoring and controlling of the stations from the north during the course of the austral winter. As part of this work I took responsibility for monitoring and controlling station operations for the duration of 2013, with little down time experienced in ARA 2 and ARA 3 during the year. The event rate between April and September 2013 for ARA 2 is shown in Figure 5.12. Due to disk space limits the trigger threshold was changed at the beginning of May. The effect of this change was to lower the event rate to a manageable level for the remainder of the year. In addition there are two shorter drops in event rate, one in late May and the other in September. These correspond to periods of running where the calibration pulsers (which run at 1Hz) were not in operation.

Due to an electronics issue the DAQ installed at ARA 1 was unable to take data throughout the year and stopped functioning around a month after the summer

season closed. The issues encountered with ARA 1 exemplify the challenges faced in developing and maintaining detectors in such a remote location and in such harsh conditions. These problems are tempered by the success of ARA 2 and ARA 3, showing long term stability of the new DAQ systems over the course of 2013.

Chapter 6.

TestBed Data Analysis

The TestBed is a prototype installed with the intention of showing that a much larger detector will be able to reach the necessary sensitivity to make the first observation of ultra high energy (UHE) neutrinos. During the course of 2011 and 2012 the TestBed recorded around 480 million events. Simulations of the TestBed and constraints provided by a number of experiments inform us that very few, if any, of these events will contain signals caused by neutrino interactions in the ice.

The analysis described in this chapter will show that a single array of radio antennas can be used to reach a reasonable sensitivity to UHE neutrino fluxes that, under realistic assumptions, scales to a projected world leading sensitivity for the full 37 station detector.

6.1. Analysis approach

The TestBed is the first ARA station deployed in the ice around the geographic South Pole. This analysis of the data recorded over its 2 year operating period aims to demonstrate that it is possible to identify neutrino-induced signals within this data set. Almost all of the recorded events will not contain neutrino signals and are therefore considered backgrounds to this analysis. In order to identify the small number of signal events it is necessary to reduce the expected background levels to of order 0.1 events remaining after analysis. This requires background rejection levels better than 1 event in a billion passing, which is a considerable challenge.

The expected sources of background events falls into two categories: firstly random fluctuations in the radio frequency (RF) environment caused by the thermal conditions in and around the TestBed, and secondly RF noise associated with human activity and infrastructure in the vicinity of the station. These backgrounds, which are referred to as thermal noise and anthropogenic signals from this point forwards, have differing characteristics that can be used to separate them. Some anthropogenic sources are expected to produce signals that follow closely the expectations of neutrino-induced events. An important goal of any analysis of TestBed data is to show that this background is at a manageable level that will not impact heavily on a much larger array's sensitivity.

These categories of events have very different properties so an analysis approach is taken that aims to remove them in two stages. The first stage is to remove thermal noise events by requiring coherence between signals in multiple antennas consistent with a common physical source location. Due to the random nature of thermal events it is possible to use reconstruction and coherence tests to remove them from the data set, leaving only directional events. The majority of the remaining events are expected to be of anthropogenic origin associated with infrastructure and human activity around the South Pole Station, hence a second stage is concerned with removing such backgrounds. These signals are likely to exhibit both spatial and temporal correlations not expected from a sample of neutrino events. Therefore it is possible to remove such anthropogenic signals by rejecting reconstructed source locations associated with infrastructure or repeated signals and finally by removing events from time periods associated with high levels of human activity.

In addition to directional anthropogenic signals another category of events is expected in the TestBed resulting from the operation of narrow-band communications equipment. Events containing such carrier wave (CW) signals are likely to pass coherence tests intended to remove thermal events but will not reconstruct to their origin, making spatial correlation tests of limited use. Reconstruction of incident signals relies on using time offsets between signals derived from the correlation between them. For CW signals period ambiguities give rise to a number of possible offsets between a pair of antennas, each as likely as the other, resulting in an inability to clearly identify the source location. Some fraction of these events will miss-reconstruct to physical locations in the ice so it is necessary to remove them by other means. By considering the frequency content of such events it will be shown that it is

possible to separate these signals from neutrino-like events due to their narrow-band nature.

In Section 6.2 and Section 6.3 the data sets and blinding strategy for this analysis will be outlined. A series of tools were developed for use in this analysis in order to reconstruct events, identify those with coherent signals between multiple antennas and to remove CW contamination from the data set. These tools will be described in Section 6.4. A series of cuts are then identified and chosen to separate signals from the various categories of background events. Firstly data quality cuts are described in Section 6.5, then the thermal noise and anthropogenic signal cuts are described in Section 6.6 and Section 6.7 respectively.

Any events that pass all of these cuts are considered to be neutrino candidates. A deliberate choice was made to omit any signal shape cuts in order to minimise dependencies on theoretical calculations of Askaryan signals and assumptions made about radio propagation in the ice.

6.2. Blinding

The TestBed operated with very little downtime for the duration of 2011 and 2012 with the full data set archived to disk at the South Pole. A limited bandwidth satellite link was available to transfer approximately 10% of this data back to the northern hemisphere data warehouse with a latency of a few days. This 'filtered' data set, which was randomly chosen, was used to assess operation and conditions in the TestBed such that significant data quality and downtime issues could be addressed during the winter season. The disks containing the full data set from 2011 and 2012 were hand carried back to the northern hemisphere during the 2011-2012 and 2012-2013 summer seasons. The data was subsequently stored in the data warehouse, hosted at the University of Wisconsin, and used as the basis of this analysis.

A blinding strategy was agreed by the experiment to limit any bias in analyses performed on the data. This was implemented by splitting the analysis data set into two samples: a 'burn sample' intended to be used for training the analysis and associated cuts, and a '90% sample' containing the remaining events. Once the experiment was in agreement with the methods, tools and cuts implemented on the

'burn sample' of data the analysis was unblinded and run over the '90% sample', with any events passing the cuts considered as neutrino candidates.

6.3. Data types

There are two types of event recorded and made available for analysis:

- **RF triggered events** - events in which the trigger condition is met requiring 3 of 8 antennas to have passed a power threshold within a 110ns window.
- **Minimum bias events** - recorded at a rate of approximately 0.5Hz these events are force triggers where the RF trigger condition has not been met and are a reflection of the background conditions in which the TestBed was operating.

The RF triggered events are further sub-divided into two classes:

- **Calibration pulser events** - these events are RF triggered events that are identified as containing signals from a number of calibration antennas by timing information.
- **Non-calibration pulser events** - the remaining set of RF triggered events after calibration pulser events have been removed.

The non-calibration pulser events recorded by the TestBed form the majority of the data set and are those to be analysed and assessed to see whether they contain neutrino induced signals. Minimum bias events were recorded at regular intervals by the DAQ system in order to provide a picture of the radio environment during the course of data taking. As they do not require a physics trigger condition these events are in the main thermal in nature. There are, however, periods for which repetitive or long duration anthropogenic noise sources are active contaminating this thermal sample of events.

During the course of data taking two calibration sources were active: a VPol calibration pulser in 2011, and a HPol calibration pulser in 2012. The pulsers were fired at a rate of 1Hz using a pulse-per-second (PPS), provided by a Global Positioning System (GPS) unit in the DAQ, as a trigger. The result is that these

events are identifiable via trigger timing information and removed from the '90% sample'. These calibration pulser events were used in training the reconstruction methods as well as a number of the cuts used in this analysis. Being the only controlled in-ice source of impulsive radio signals calibration pulser events are used as a proxy for neutrino induced signals and provide an invaluable cross-check for analysis tools and cuts designed to identify impulsive signals.

The two data samples defined in the blinding strategy were comprised of a combination of the event types described above:

- **burn sample** - Sample of events intended for analysis tools and cut development. Comprised of all calibration pulser events, all minimum bias events and 10% of non-calibration pulser events.
- **90% sample** - The remaining 90% of non-calibration pulser events not made available in the burn sample.

A third sample of events consisting of simulated neutrino signals were produced using the official ARA simulation AraSim [72]. This sample was produced with fixed incident neutrino energies in half decade intervals between 10^{17}eV and 10^{21}eV and was used to inform cuts and help with development of analysis tools such as source reconstruction. No simulations of cosmic ray signals were analysed as part of this analysis.

AraSim simulates neutrino signals in three stages: the neutrino propagation and interaction, RF signal transmission through the ice to the TestBed antennas and finally the output of the digitisation and triggering paths within the TestBed detector. The resulting data set contains digitised waveforms that would be expected to be recorded by the TestBed given knowledge of the TestBed detector and neutrino interactions.

Neutrinos are produced with isotropic arrival directions within a uniform volume around the centre of the TestBed. Each event is simulated independently and assigned a weight representing the probability that the neutrino would have reached the interaction point without being absorbed in the Earth. The shower development and RF Cherenkov emission are modelled before signals are propagated to the TestBed using a ray-tracing algorithm to determine the path taken through ice which, due to the changing index of refraction, leads to ray bending effects resulting in zero,

one or two rays arriving at the antennas. Finally the signals are propagated through a model of the triggering and digitisation chains within the TestBed, which are based upon measurements taken in the laboratory ahead of installation in the ice.

In addition to simulating the expected neutrino signals thermal noise is added to the events. The noise model is derived from minimum bias events recorded in the TestBed and tuned to match the observed noise levels and trigger rates.

6.4. Analysis tools

The TestBed records waveforms from both vertically polarised (VPol) and horizontally polarised (HPol) antennas for each event and these polarisations are treated separately in analysis. This is in part due to the differences in signal chain and hence noise levels between the two. In addition it is expected that the majority of anthropogenic signals will be limited to a single polarisation. Once initial data-quality checks have been applied to the events they are split into two 'half events', one comprising of the signals in the VPol antennas and the other similarly from HPol antennas, that are propagated through the remaining cuts in order to assess whether or not they are to be considered a neutrino event.

This section will outline the analysis tools that are used on the analysis data set. They were designed to be used with waveforms of a single polarisation and each of the tools described are applied to both the VPol and HPol antennas separately.

6.4.1. Event Reconstruction

Directional reconstruction of incident radio signals recorded in the TestBed is a powerful tool with which to identify and remove backgrounds, whilst providing invaluable information about possible neutrino events. Interferometric techniques have been successfully used in similar experiments, for example the ANITA experiment [73] [70] [74] [75], that also record high fidelity radio waveforms. There are, however, a number of obstacles to robust reconstruction using interferometry in the TestBed.

The small number of antennas (4 deep antennas in each polarisation) limit the measured strength of any preferred source direction, which would be re-enforced with

more antennas. The varying index of refraction in the firn layer (typically extending $\sim$ 150m into the ice close to the South Pole) causes ray-bending effects that distort the true arrival direction. And finally the uncertainty in the deployed positions of the receive and calibration antennas in the ice can further complicate reconstruction.

A new reconstruction method was developed for this analysis, based upon calculating timing offsets between waveforms that maximise correlation. This is achieved by creating a coherently summed wave (CSW) where individual antenna waveforms, scaled by the total number of waveforms, are added offset relative to one another. The CSW is given by:

$$\mathrm{CSW}(t) = \frac{1}{N}\sum_{i}\psi_i(t + \Delta t_{1,i}) \tag{6.1}$$

$$\tag{6.2}$$

where N is the total number of antennas, ψ_i is the time domain waveform for the i^{th} antenna and $\Delta t_{1,i}$ is the time offset between antenna 1 and antenna i. These offsets are calculated in a manner such that the correlation between the resulting CSW and the original waveforms is maximised, as measured by the normalised cross-correlation $C_{1,2}$:

$$C_{1,2}(\Delta t) = \frac{\psi_1 \star \psi_2}{\sigma_1\sigma_2} \tag{6.3}$$

where $\psi_{1/2}$ are the time domain waveforms, Δt the time offset between them and $\sigma_{1/2}$ are the waveform's RMS. For discretely sampled waveforms containing a limited number of samples the cross-correlation $\psi_1 \star \psi_2$ is given by:

$$\psi_1 \star \psi_2(\Delta t) = \sum_{n}^{N} V_1(t_n)V_2(t_n - \Delta t) \tag{6.4}$$

where the voltages $V_{1/2}$ outside of waveform window are taken to be zero. In order to improve the timing resolution obtained from correlations the individual antenna waveforms are up-sampled, using interpolation to an evenly sampled time-base, so that the time between adjacent samples is 0.25ns and the corresponding offsets Δt used in Equation (6.4) are the same. A simple algorithm is implemented to find the set of timing offsets that maximise the correlation between each waveform and the resulting CSW. The timing differences between pairs of antennas found using this method are shown in Figure 6.1 for a selection of calibration pulser events a width of $\sim$ 150ps.

The VPol calibration pulser signal, which was used during 2011, is a very bright source causing compression in some of the digitised waveforms and resulting in the main peak and side lobes being of a similar size. When the CSW is formed either the main peak or side lobes of the waveforms are lined up, leading to a double peak structure in the time difference between antenna's 1 and 3.

In Figure 6.2 the calculated time differences are applied to offset individual waveforms. These offset waveforms are summed to form a CSW also shown for a calibration pulser event, noise event and simulated neutrino event. The CSWs formed in this manner differ significantly between thermal noise waveforms and signal type waveforms (i.e. the calibration pulser and simulated neutrino events). Two main features are apparent: firstly the size of the peak in the coherently summed wave is much greater for signal type events, and secondly the resulting CSW looks very similar to the shape of the individual waveforms for signals of interest, which is not the case for noise events. The similarity of signals in impulsive events, such as the calibration pulser and simulated neutrino events, leads to constructive interference in the CSW. The noise, on the other hand, adds destructively leading to an increased signal to noise ratio in the former. The larger signal sizes and correlation between individual antennas and the CSW can be used to aid discrimination between signals and backgrounds in an analysis.

The timing difference between pairs of antennas holds information about the arrival direction of the radio signal and are checked against those calculated from a simple ice model. The model takes the index of refraction to be a constant, equal to the measured value at the average depth of the borehole antennas. A pseudo-χ^2 is computed for a series of trial source locations in 1 degree bins in θ, ϕ and logarithmically spaced bins in radial distance R:

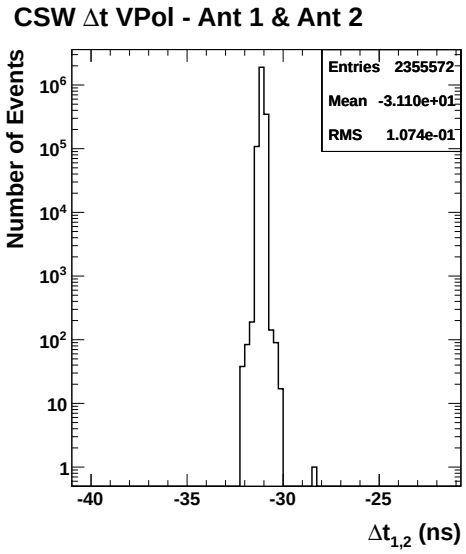

(a) Antenna 1 and Antenna 2

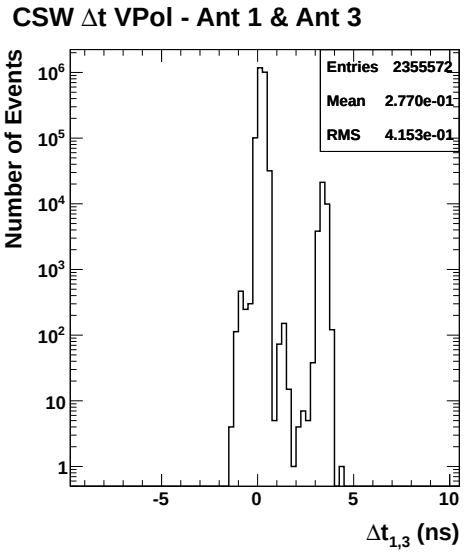

(b) Antenna 1 and Antenna 3

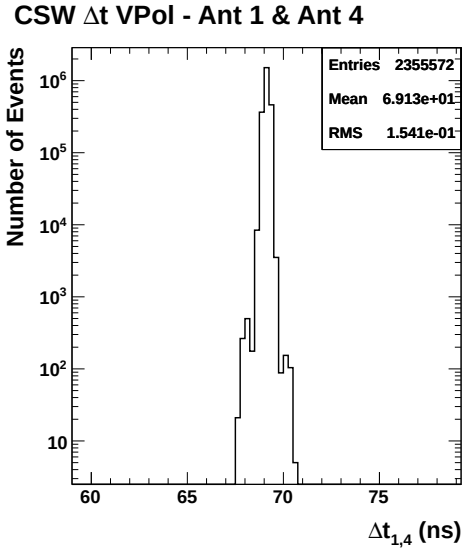

(c) Antenna 1 and Antenna 4

Figure 6.1.: CSW measured time offsets between pairs of VPol antennas for calibration pulser events.

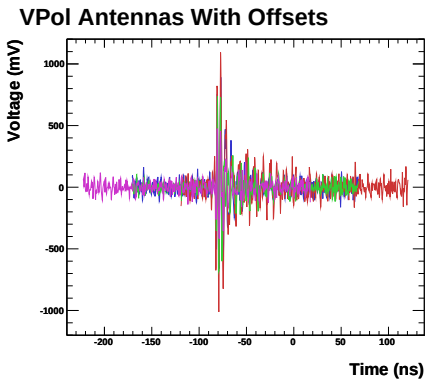

(a) Individual waveforms calibration pulser event

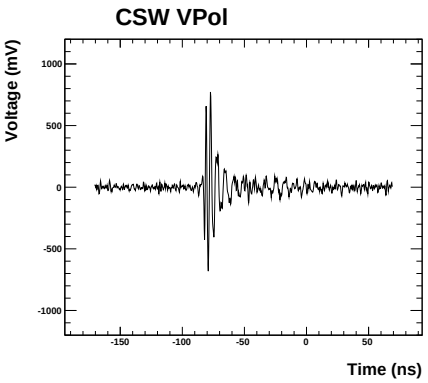

(b) CSW calibration pulser event

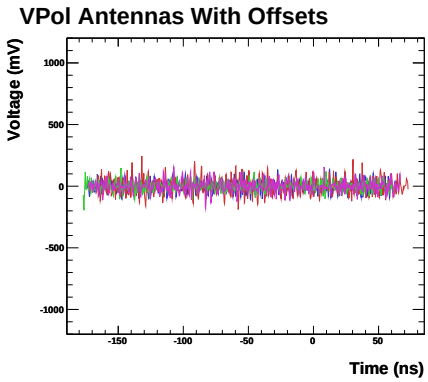

(c) Individual waveforms noise event

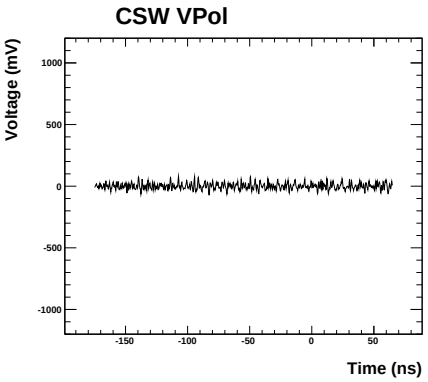

(d) CSW noise event

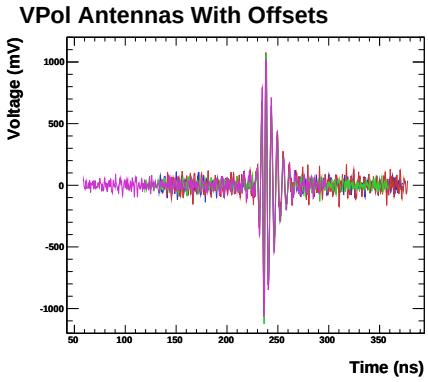

(e) Simulated neutrino event

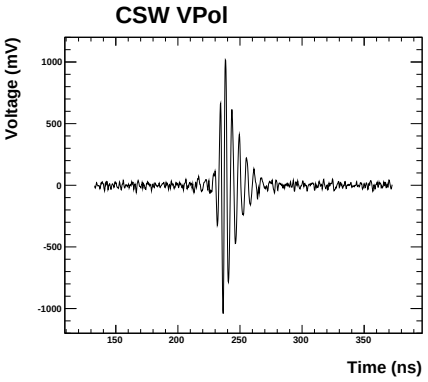

(f) Simulated neutrino event

Figure 6.2.: (a), (c) and (e) show individual antenna waveforms aligned in time using the calculated offsets that maximise the correlation between antennas and the CSW. The blue, red, green and magenta lines show antennas 1, 2, 3 and 4 respectively. (b), (d) and (f) show the resulting coherently summed wave, where the time aligned waveforms are summed and scaled by the total number of antennas.

$$\text{pseudo-}\chi^2 = \sum_i \frac{(\Delta t_{1,i}^{meas} - \Delta t_{1,i}^{exp})^2}{\sigma^2} \tag{6.5}$$

where $\Delta t_{1,i}^{meas}$ is the measured offset between antenna 1 and antenna i, $\Delta t_{1,i}^{exp}$ is the calculated offset expected from a trial source location and σ is taken to be 1ns (which is similar to the timing uncertainty expected given uncertainty on the antenna positions). The reconstructed location is that which minimises pseudo-χ^2 and hence corresponds to the most likely physical location given the measured time offsets. Calculating pseudo-χ^2 values in all possible θ, ϕ and R bins is computationally intensive and a time consuming process, so an algorithm was developed to only search a subset and identify the best fit location.

This method has the benefit of using the rich information contained within the digitised waveforms (correlation techniques result in precision of $\sim$ 150ps resolution in timing differences between pairs of antennas) as well as providing a parameter that describes the goodness of fit in pseudo-χ^2 upon which a cut can be placed. As thermal signals will have random measured offsets between antennas the preferred source location will, in general, have a relatively large pseudo-χ^2 value associated with it, thus a requirement for good reconstruction will additionally reject a large number of thermal events.

Figure 6.3 shows reconstruction maps for a series of event types, consisting of calculated pseudo-χ^2 values for points on the surface of a sphere at the same radial distance as the best fit location. CSW reconstruction is attempted for each polarisation in every event, with the best fit location as well as corresponding pseudo-χ^2 value recorded. The results of ray-bending in the ice are visible in the reconstruction map for a simulated neutrino event in Figure 6.3(c), where the reconstructed elevation angle is above the true neutrino interaction point. This effect is much smaller for the callibration pulser event in Figure 6.3(a) as the calibration pulser is much closer to the receive antennas and is located at approximately the same depth. Figure 6.4 shows the residuals (measured minus expected) for azimuth and elevation angle for an ensemble of calibration pulser events and for neutrino events simulated with AraSim. Reconstruction of azimuth and elevation is $\sim 1^\circ$,

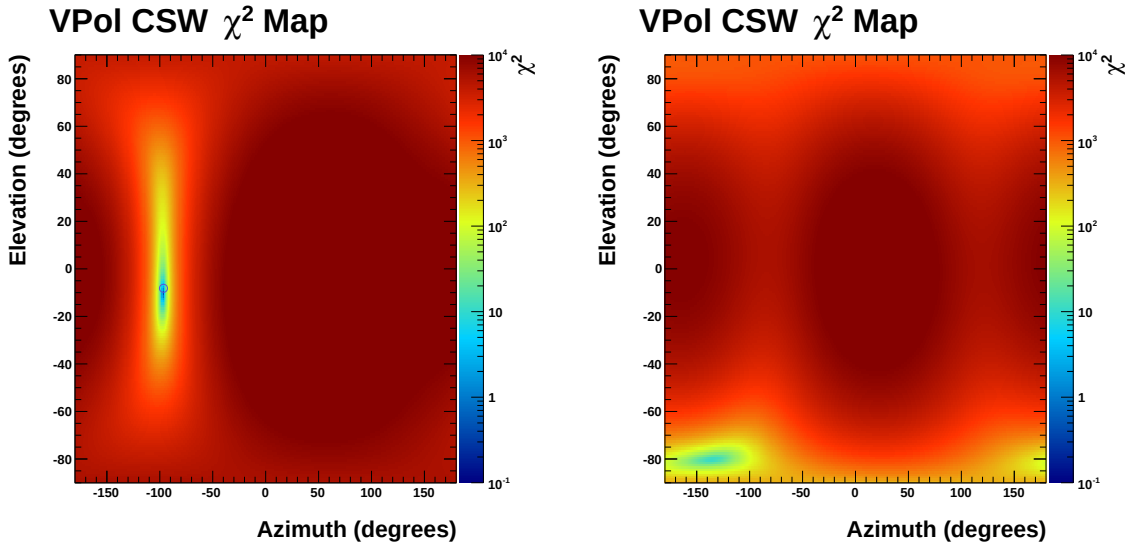

(a) Pseudo-χ^2 map of calibration pulser event

(b) Pseudo-χ^2 map of noise event

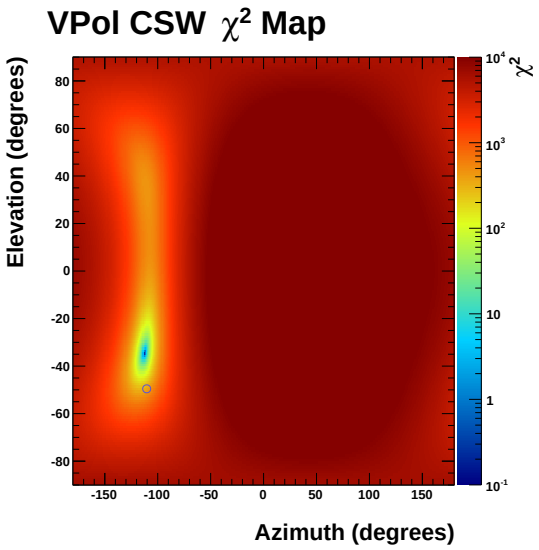

(c) Pseudo-χ^2 map of simulated neutrino event

Figure 6.3.: Reconstruction maps of calculated pseudo-χ^2 values for (a) the same calibration pulser event in Figure 6.2, (b) a thermal noise event and (c) a simulated neutrino event. The circles in (a) and (c) indicate the true source location. The offset between reconstructed and true neutrino interaction point in (c) is due to the ray-bending effects in the ice. The signal type events (a) and (c) have much lower pseudo-χ^2 values than the noise event.

although the ray-bending effects in the ice cause elevation angles to be systematically shifted.

The ray-bending becomes more pronounced for sources at large radial distances from the TestBed and, as there is more target volume at these large distances, the majority of simulated neutrino events suffer from these shifts. This effect is shown in Figure 6.5 where events coming from elevation angles close to horizontal are reconstructed above their origin. It is possible to calculate a transfer function to correct the shift in reconstructed elevation angle between true and reconstruction source location, which is also shown. This is an effective mapping between the ice models used for simulation of RF propagation in the ice and used in the reconstruction. The elevation residuals for simulated neutrino events in Figure 6.4 have such a correction applied and show that the systematic shift is successfully removed.

6.4.2. Carrier wave removal

The frequency band for the TestBed ($150 - 850$MHz for VPol and $250 - 850$MHz for HPol) contains a number of frequencies used for communications at the South Pole. A particularly strong carrier signal used by the South Pole Station at 450MHz is removed by a notch filter placed between the antenna output and low noise amplifiers down-hole, but a number of other transmission frequencies remain unfiltered. Due to the TestBed's remote location and the relatively low levels of human activity (which are largely restricted to the summer season) CW signals are not present in the vast majority of RF triggers and minimum bias events. However, the presence of CW in an event can cause miss-reconstruction of the incident signal and, in some cases, mimic the properties that distinguish neutrino signals from thermal noise. For this reason these contaminated events constitute a background for any neutrino search, although the properties of such signals mean that it is possible to distinguish them from other interesting events.

CW signals are characterised by large unthermal amplitudes in a small range of frequencies, compared with smaller broadband excesses in impulsive, neutrino-like, signals as power is spread over a larger range of frequencies. This section will describe how the frequency content in the VPol and HPol antenna waveforms is used to identify such events and, as the fraction of these contaminated events is small, remove them from the analysis sample.

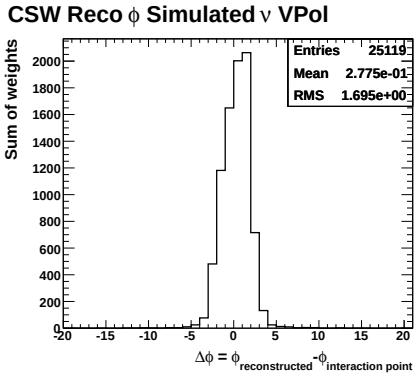

(a) Simulated neutrino reconstruction azimuth

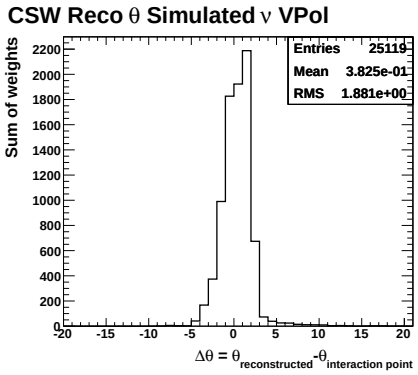

(b) Simulated neutrino reconstruction elevation

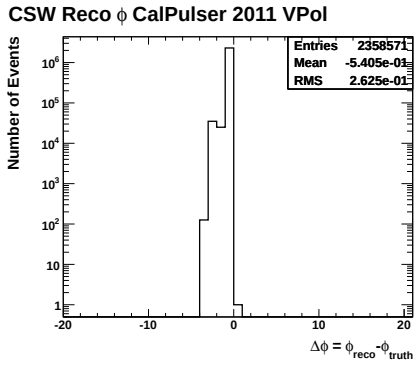

(c) Calibration pulser reconstruction azimuth

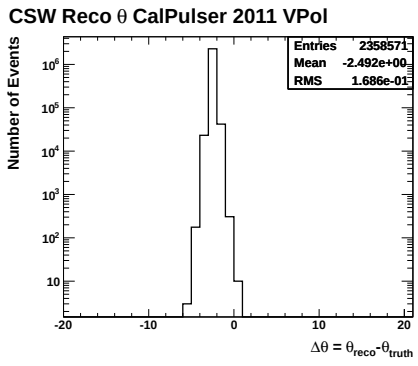

(d) Calibration pulser reconstruction elevation

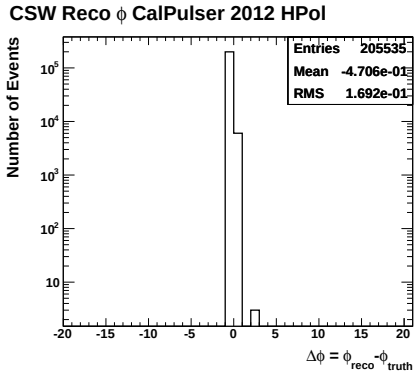

(e) Calibration pulser reconstruction azimuth

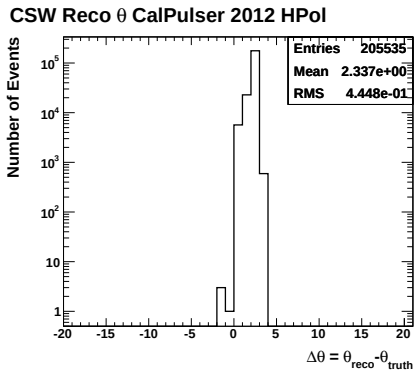

(f) Calibration pulser reconstruction elevation

Figure 6.4.: Residuals for reconstructed source direction azimuth (ϕ) and elevation (θ). For simulated neutrino events the source location is taken to be the neutrino interaction point. Due to ray-bending effects a correction factor is applied for simulated neutrino events to translate the reconstructed elevation angle to the line of sight to the source.

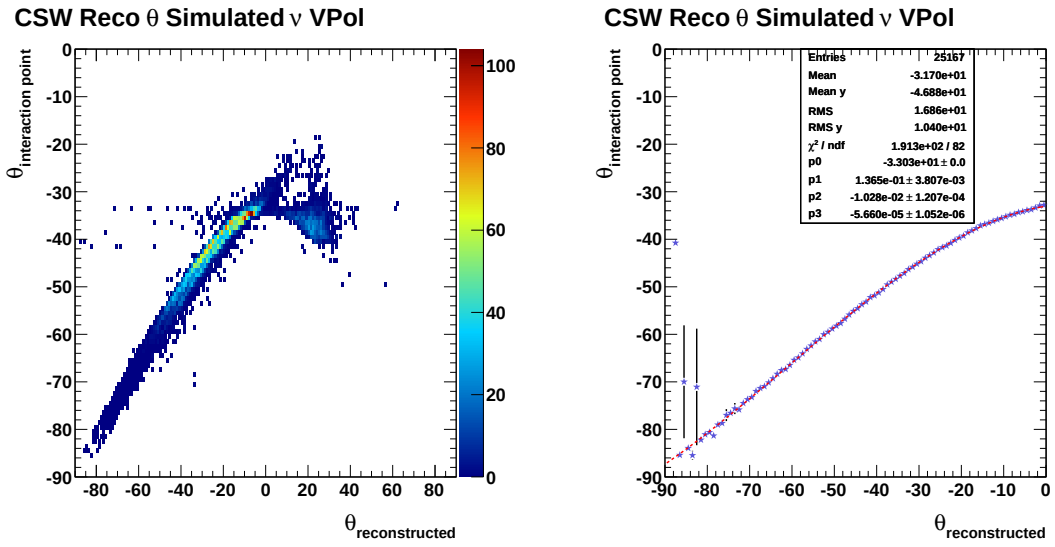

(a) Interaction point θ versus reconstructed θ

(b) Reconstructed θ correction

Figure 6.5.: Reconstruction of simulated neutrino elevation angles. In (a) the neutrino interaction point θ is shown as a function of all reconstructed angles, which exhibits strong ray-bending effects close to and above horizontal angles ($\theta > 0$). In (b) a profile is taken for events that reconstruct downward ($\theta < 0$) and a correction function fitted to the data.

Prior to analysis of an event a thermal noise baseline is generated using minimum bias data, being the purest sample of thermal noise available for analysis. This baseline is taken to be a reflection of the thermal conditions in the TestBed during a run. The frequency content of an event can then be compared with this baseline to assess the level of unthermal power and whether this power is confined to a small range of frequencies as expected from CW events.

Thermal noise baselines

The frequency domain amplitudes of thermal noise are Rayleigh distributed:

$$\text{Rayleigh p.d.f.} = \frac{A}{\sigma^2} e^{\frac{-A^2}{2\sigma^2}} \tag{6.6}$$

where A is the amplitude at frequency f and σ can be used to characterise the distribution. When creating a baseline for a run the Fourier transform is taken of the time domain waveforms for each minimum bias event. The resulting frequency domain amplitudes are histogrammed per frequency per antenna and fitted with Rayleigh distributions.

Figure 6.6 shows one such Rayleigh fit for the case of a run containing mainly thermal events, and for a run containing a fraction of CW contaminated events. For each run the Rayleigh fit is characterised by a series of σ values, which are calculated per frequency per antenna, which are recorded along with the average power per frequency per antenna. Some runs contain a fraction of CW events large enough to distort the Rayleigh distributions with large unthermal tails in the amplitude distribution, making them unusable as a representation of the thermal conditions, and their baselines are rejected as 'bad' (the dashed line in Figure 6.6 shows such a run, the resulting Rayleigh fit would differ considerably from that derived from the solid line, even though these runs were taken on the same day with similar thermal conditions).

Due to the narrow band nature of CW signals it is possible to identify bad baselines by looking for spikes in the average power spectra shown in Figure 6.8 or σ distributions, although the former was found to have better discriminating power.

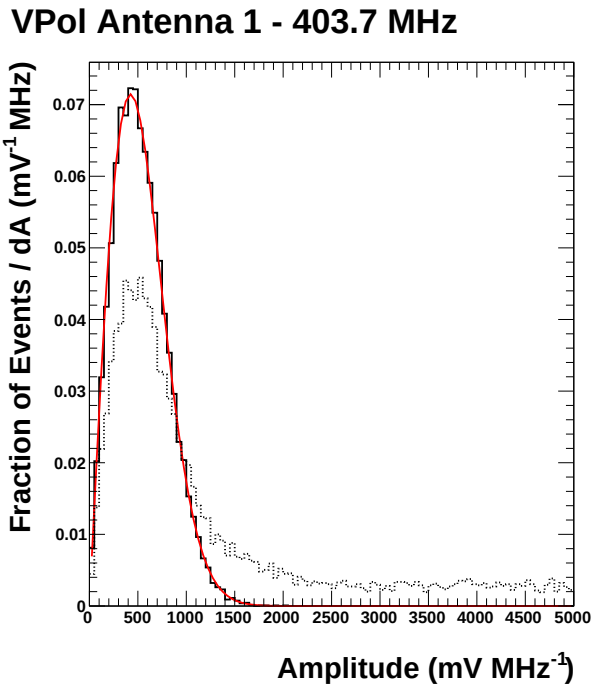

Figure 6.6.: Thermal noise amplitudes in VPol antenna 1 at 403.7 MHz for minimum bias data taken on 20th May 2011 (black) and a Rayleigh distribution fit to the data (red). The dashed black line is the same histogram but this time populated from a run containing a known CW source operating at $\sim$ 403MHz.

These spikes are identified by looking for the second derivative of the average power spectrum falling below a threshold, indicating the presence of a spike. Any run that contains at least one borehole antenna with a spike in the average power spectrum is rejected.

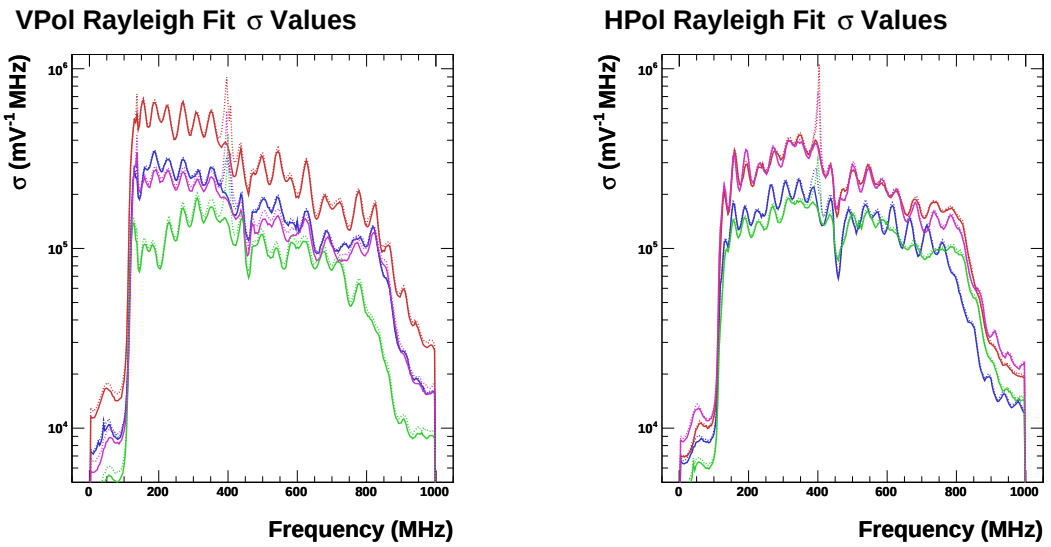

(a) VPol Antennas **(b)** HPol Antennas

Figure 6.7.: Rayleigh fit derived σ values for two different baselines. The solid lines are for a baseline calculated from a thermal noise sample, and the dashed lines for a baseline containing a known CW source operating at 403MHz.

The baselines calculated in this manner are used to summarise the thermal conditions of the TestBed, examples are shown in Figure 6.7. The baseline that an event is compared to should therefore be chosen to be from a period with similar thermal conditions. In the main this will be the baseline calculated from the run in which the event was recorded, but where this baseline contains a large amount of CW contamination (hence the baseline is rejected) the closest run in time is used. The variation of baseline calculated σ values as a function of time is shown in Figure 6.7. The thermal conditions change slowly compared with the rate at which runs are taken, hence no scaling is needed to account for variations in the conditions.

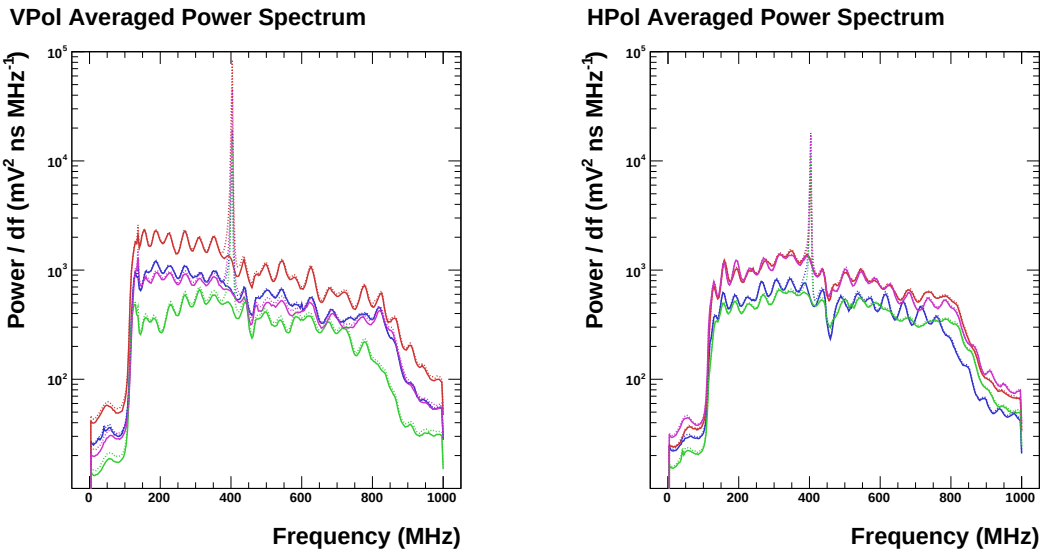

(a) VPol Antennas **(b)** HPol Antennas

Figure 6.8.: Averaged power spectra for (solid lines) a run containing largely thermal events, and (dashed lines) a run containing a CW source. Bad runs, such as that summarised by the dashed lines, are identified by spikes in the averaged power spectra characterised by the second derivative falling below a threshold.

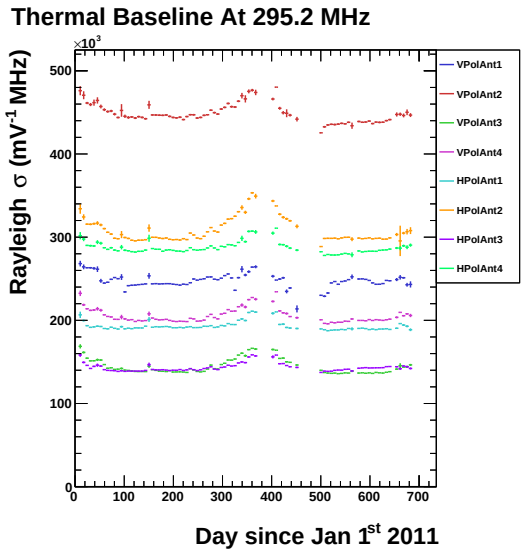

Figure 6.9.: Rayleigh σ values, averaged over a 7 day period, are shown for all antennas at 295.2MHz for 2011 and 2012.

Identifying CW contaminated events

Each event has it's frequency content considered to identify CW signals. Discrete Fourier transforms are taken of each event's time domain waveforms. Using the Rayleigh distributions from the chosen baseline it is possible to estimate the probability of the measured frequency domain amplitudes. These probabilities are calculated using the Rayleigh σ values, which are taken from the relevant baseline, as follows:

$$P(A \geq A_{meas}) = \int_{A_{meas}}^{\infty} \frac{A}{\sigma^2} e^{\frac{-A^2}{2\sigma^2}} dA \tag{6.7}$$

$$= e^{\frac{-A_{meas}^2}{2\sigma^2}} \tag{6.8}$$

where A_{meas} is the measured amplitude of a given frequency for the waveform in question. $P(A \geq A_{meas})$ is then a measure of the probability of thermal fluctuations giving rise to the measured amplitude, and can be interpreted as measuring how 'unthermal' that amplitude is. Most CW transmitters have a single polarisation and will be seen in multiple antennas, a feature that can be utilised to improve discrimination. In order to push down the threshold CW signal size to which the filter is sensitive the antennas in each polarisation are grouped together and considered separately. In each polarisation the product of probabilities for the borehole antennas is taken:

$$P_{prod} = \prod P_i = \prod P(A_i \geq A_{i,meas}) \tag{6.9}$$

where P_i is the probability for antenna i. Any frequencies that have P_{prod} less than some threshold are considered to be in excess. The use of this product was found to have greater sensitivity than various other methods attempted, including taking antennas individually or in combination (for example requiring that 3 out of 4 antennas see an excess at the same frequency).

Example waveforms from VPol antenna 2 are shown in Figure 6.10 for a thermal, CW contaminated and simulated neutrino event, along with the corresponding

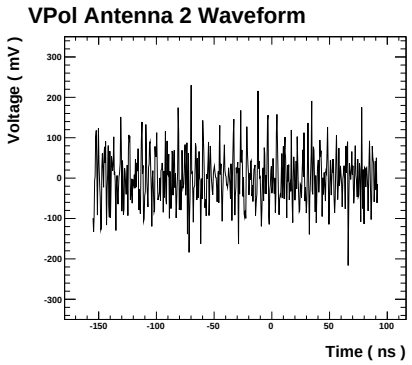

(a) Noise Waveform

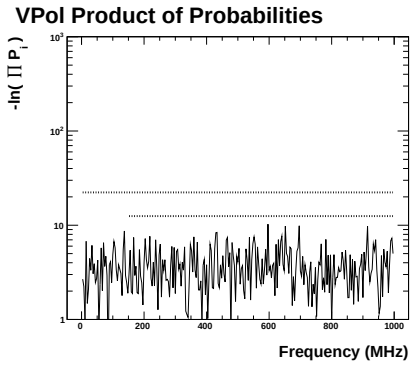

(b) Noise Probability Spectrum

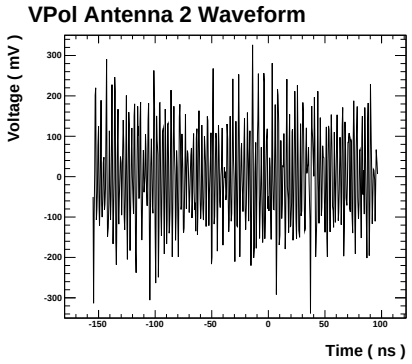

(c) CW Contaminated Waveform

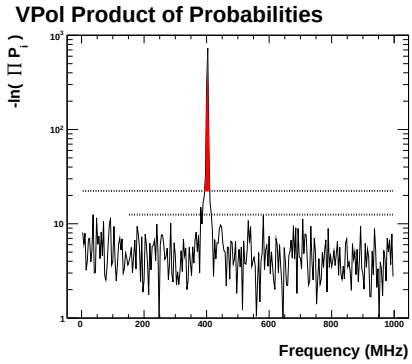

(d) CW Probability Spectrum

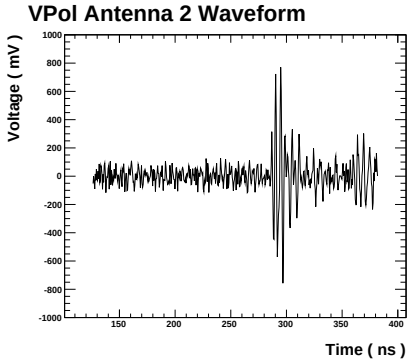

(e) Simulated Neutrino Waveform

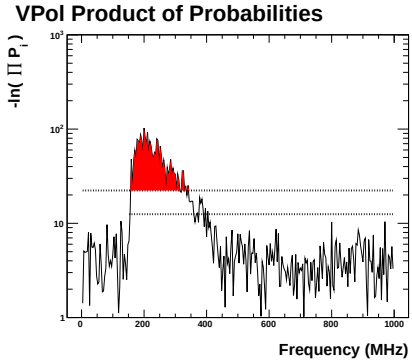

(f) Simulated Neutrino Probability Spectrum

Figure 6.10.: Example waveforms from VPol antenna 2 showing a thermal noise event, CW contaminated event and simulated neutrino event. Also shown are the product of probabilities distribution for all VPol antennas for these events. The dashed lines show the two probability thresholds, with any frequencies passing this threshold regarded as being well in excess of thermal noise levels. The upper threshold is to identify non-thermal excesses and the lower to identify broadband signals.

product of probabilities. Any CW removal must be able to reject contaminated events whilst passing broadband signals, such as those induced by neutrino interactions or from calibration sources. Two thresholds are used to identify CW events: one used to identify non-thermal excesses, and a second, lower threshold, to check whether an excess is broadband in nature. The former, which is set such that 1 in 10^6 thermal noise events will be have at least one frequency in excess, is used to identify events that may contain CW. If an event has no frequencies passing this threshold in either polarisation it passes the CW check. Any events that have 1 or more frequencies passing this threshold are then checked using the lower broadband threshold, which is set at a value of -12.5, informed by checks against calibration pulser events. If the number of frequencies in excess of this lower threshold, which is referred to as 'totalBins', is greater than or equal to 20 then the event is classed as having a broadband signal and is passed. The choice of cut value for the lower threshold was informed from studies using antennuated calibration pulser signals as a proxy for weak neutrino induced events.

In Figure 6.10 the CW contaminated event is clearly identified by a narrow range of frequencies passing both thresholds, and the event is rejected. Figure 6.11 shows how this product of probabilities spectrum changes over the course of a run for calibration events and for non-calibration events. Here the product of probabilities is averaged for all events over 1 minute intervals. Also shown are the same parameter but only for frequencies that are identified as being in excess. A weather balloon launch is clearly visible in these spectra, resulting in a narrow range of frequencies exhibiting highly unthermal amplitudes in the VPol antennas, whereas calibration pulser events show a broad range of frequencies with non-thermal amplitudes.

In addition to the narrow range of frequencies exhibiting non-thermal power the size of these excesses can be used to identify particularly bright events. A parameter, 'minProb', is used to quantify the size of the biggest non-thermal excess in the probability spectra of an event. This parameter is taken to be the minimum value of $\ln(\prod P_i)$ from the probability spectrum for an event. Since CW signals transmit all of their power in a narrow range of frequencies minProb can take much more negative values (and hence larger unthermal amplitudes) than for broadband impulsive signals. As waveforms are formed of a discrete set of time domain samples the Fourier transform is also comprised of a discrete set of frequencies. As a result particularly large amplitude CW signals can lead to leakage into multiple frequencies,

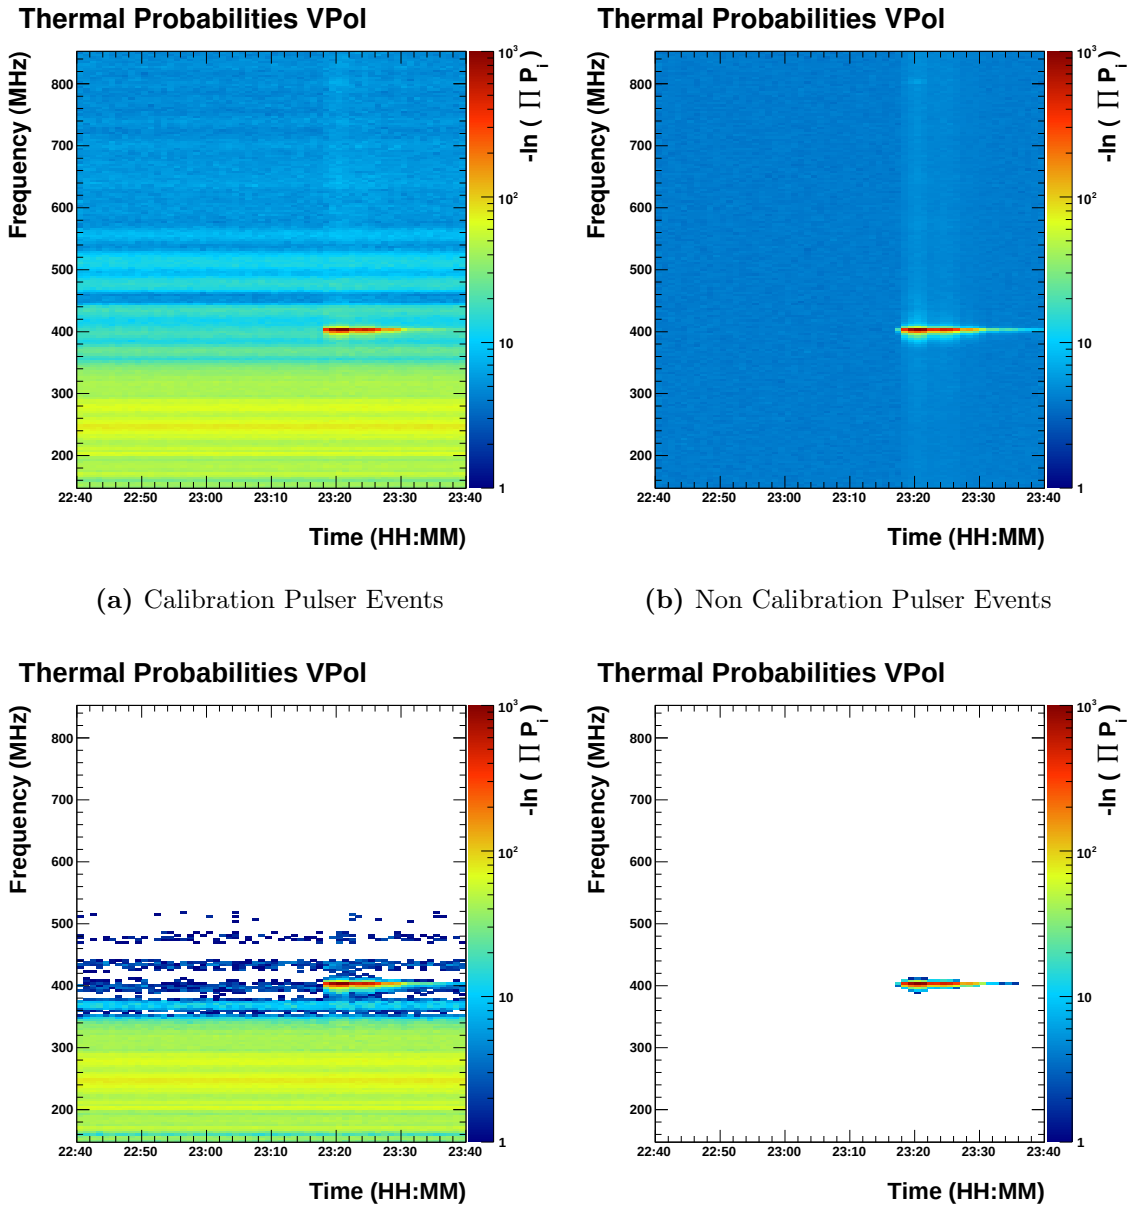

(a) Calibration Pulser Events

(b) Non Calibration Pulser Events

(c) Calibration Pulser Events

(d) Non Calibration Pulser Events

Figure 6.11.: The probability spectra for events averaged over 1 minute periods for (a) calibration pulser events (b) for non-calibration pulser events. A weather balloon launch at 23:20 is clearly visible as the turn on of a CW signal at 403MHz in both figures. The calibration pulser events show a broad range of frequencies having excess power for the duration of the run.

as the actual frequency is not one of those in the Fourier transform. In such a case CW signals can lead to totalBins approaching, or passing, the cut threshold of 20. A cut is placed on minProb to remove such events, removing any event where minProb is less than -400. This cut value is safely removed from those observed in calibration pulser events as illustrated in Figure 6.12.

6.5. Data quality cuts

A number of data quality cuts were implemented to remove various types of corrupted events. Approximately 1 month of data from 2012 were affected by a digitiser operation issue which resulted in waveforms with fewer than the usual number of time domain samples. These events were not recoverable and were thrown away using a minimum number of samples cut.

Another class of corrupted events were ones that had been corrupted during data processing, resulting in voltage-time waveforms with samples that had their voltage values corrupted.

A final class of bad quality events had a systematic and unphysical ramp up of voltage common to all digitised channels. These events were removed via an algorithm that checked the validity of data from each digitiser's clock channel.

6.6. Thermal cuts

The analysis of TestBed data aims to remove all thermal events with a small number of cuts derived from the CSW, and CSW reconstruction described in Section 6.4.1. As the expected number of candidate neutrino events is very small the cuts implemented must reject all thermal events with a high level of confidence, with the target number of thermal events passing cuts around 0.1 in the full data sample.

The minimum bias data collected during the analysis period was used as a sample of thermal events. Due to the periodic presence of CW contamination cuts were applied to the thermal sample to remove contaminated events as described in Section 6.4.2.

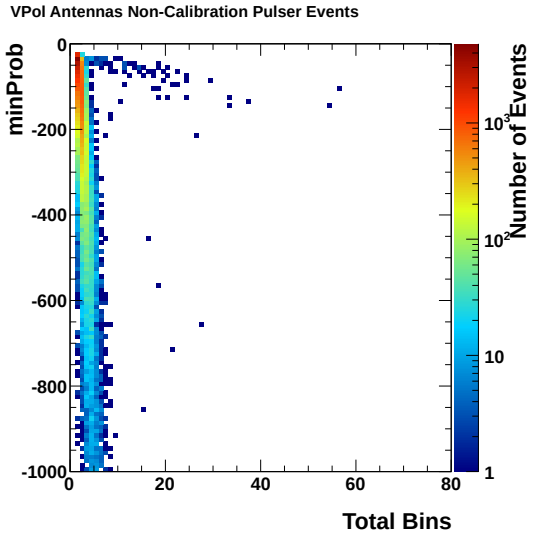

(a) Non-calibration pulser events

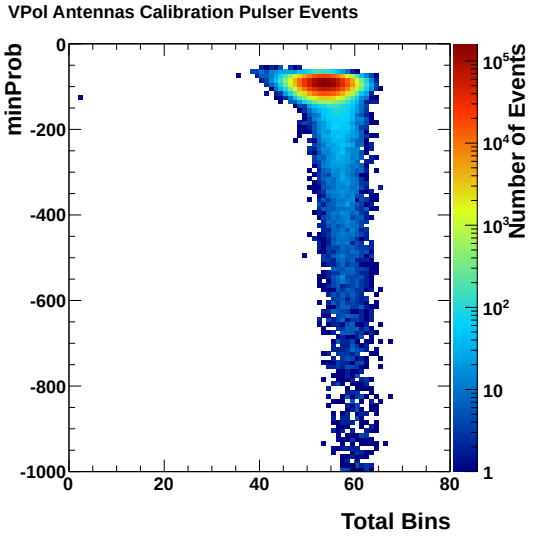

(b) Calibration pulser events

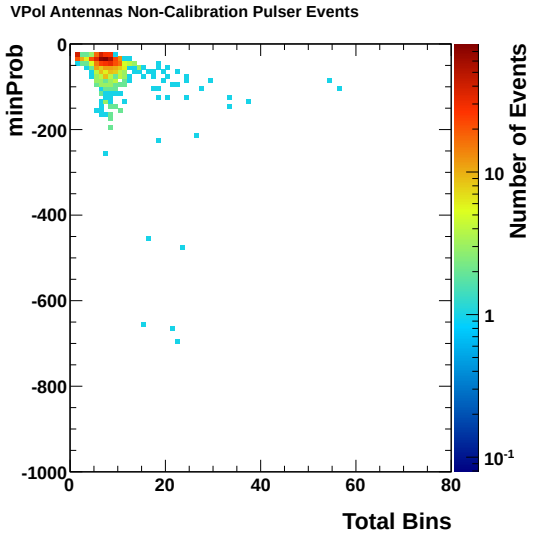

(c) Non-calibration pulser events

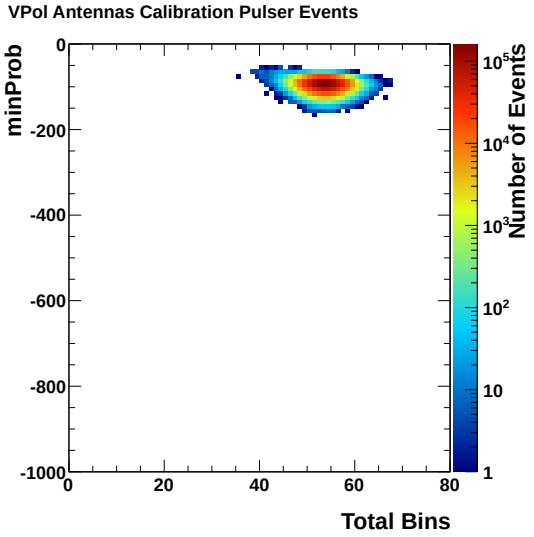

(d) Calibration pulser events

Figure 6.12.: CW parameters minProb (minimum value of $\ln(\prod P_i)$ in probability spectra) and totalBins (total number of frequency bins that are identified to have non-thermal amplitudes) are shown for calibration pulser and non-calibration pulser events. The top row shows these parameters for the full range of in band frequencies, and the bottom panel excludes a range of frequencies around 403MHz which are used by the weather balloon. The large vertical tails in the top row are clearly due to the presence of CW signals from the weather balloon and are used to inform cuts on minProb and totalBins.

6.6.1. Pseudo-χ^2 cut

For each event CSW reconstruction is performed using the VPol and HPol antennas' waveforms separately. As described in Section 6.4.1 a pseudo-χ^2 value is calculated for the best fit source location during the reconstruction, and this value is used to measure the quality of reconstruction. The reconstruction in each polarisation is classified as 'good' when the best fit pseudo-χ^2 value is below a value of two, and the event's polarisation is passed.

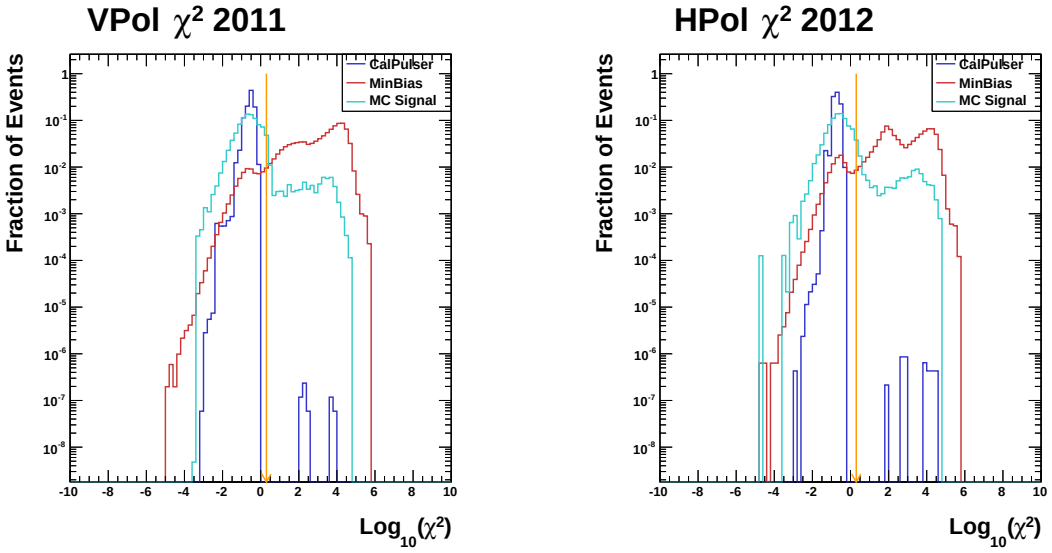

(a) VPol 2011 **(b)** HPol 2012

Figure 6.13.: Best fit pseudo-χ^2 values for minimum bias, calibration and simulated neutrino events. Events are passed when the pseudo-χ^2 falls below the chosen cut value of 2, marked by the orange arrow.

Figure 6.13 shows the pseudo-χ^2 values for each polarisation for a sample of thermal, calibration pulser and simulated neutrino events. Only $\sim 10^{-7}$ calibration pulser events have a best bit pseudo-χ^2 in the same polarisation greater than 1 which was used to inform the chosen cut value.

6.6.2. Powherence cut

Along with being reconstructable, neutrino events are expected to exhibit non-thermal levels of power and coherence between antennas, which can be used to separate

them from thermal noise events. The peak absolute voltage in the CSW, named 'CSWPeakVoltage', is used a measure of the power in the event.

In order to measure the coherence between antennas a cross-correlation is taken between each antenna and the CSW minus that antenna:

$$\mathrm{CSW}(t) = \frac{1}{N}\sum_{i}\psi_i(t+\Delta t_{1,i}) \tag{6.10}$$

$$\mathrm{CSW}^{j}_{N-1}(t) = \frac{1}{N-1}\sum_{j\neq i}\psi_i(t+\Delta t_{1,i}) \tag{6.11}$$

$$\mathrm{C}^{j}_{N-1}(\Delta t) = \mathrm{CSW}^{j}_{N-1} \star \psi_j \tag{6.12}$$

where $\psi_j(t)$ is the time domain waveform of antenna j, $\Delta t_{1,i}$ is the time offset between antenna 1 and antenna i, $\mathrm{CSW}^{j}_{N-1}(t)$ is the CSW minus antenna j and $\mathrm{C}^{j}_{N-1}(t)$ is the resulting cross-correlation waveform. For each cross-correlation waveform the maximum correlation value is taken and these values are summed to form a variable 'sumCorrVals', which gives a measure of the correlation between antennas in the CSW.

The best discrimination between calibration pulser events (as a proxy for neutrino signals) and thermal noise events was found by taking a linear combination of CSWPeakVoltage and sumCorrVals. The resulting parameter, dubbed 'powherence', should have low values for thermal events and high values for signal type events.

In Figure 6.14 we see the powherence value calculated in each polarisation for the thermal data set alongside calibration pulser, simulated neutrino events and the chosen cut value. The relative brightness of the 2011 VPol calibration pulser can be seen clearly in the large values of powherence produced compared with the 2012 HPol calibration pulser.

The chosen cut value was informed from the distribution of the powherence parameter for minimum bias events, which is shown in Figure 6.19. The cut was placed such that the expected number of thermal noise events passing in the full sample would be approximately 0.1 events for each polarisation.

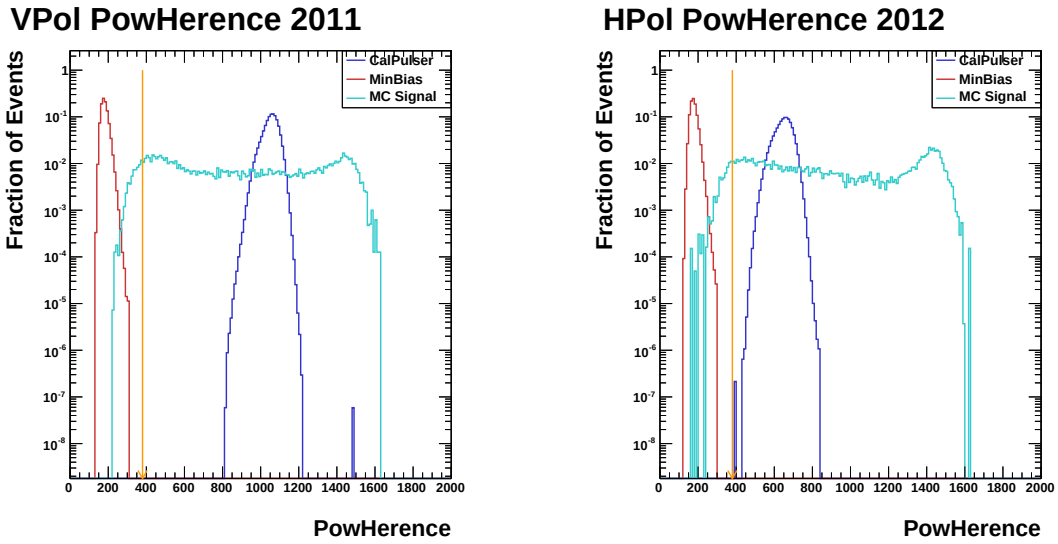

(a) VPol 2011 **(b)** HPol 2012

Figure 6.14.: The composite parameter powherence is shown for a sample of thermal, calibration pulser and simulated neutrino events after the application of the pseudo-χ^2 cut, along with the chosen cut value in orange. Events that have a powherence > 380 are passed as being signal like events.

6.7. Anthropogenic cuts

After applying thermal noise cuts the analysis sample should contain only signal-like events, however a large fraction (if not all) of these are expected to be associated with human activity and infrastructure at the South Pole. These events are likely to be repetitive both in reconstructed direction and in time, meaning they can be identified and removed by considering associated variables.

The first set of cuts are based on the CSW reconstructed source location and are referred to as 'geometry cuts', the second set are to identify noisy periods within the analysis data set and are known as the 'goodTimes' cuts.

6.7.1. Geometry cuts

A series of conservative geometry cuts are placed on an event's reconstructed source location designed to remove a large fraction of noise events:

- **Downward pointing cut** All events that reconstruct to elevation angles above 40° are rejected [1]
- **South Pole infrastructure cut** All events reconstructing within a 50° wide region in azimuth associated with the IceCube laboratory and South Pole Station are rejected
- **Calibration pulser cut** All events that reconstruct to within 5° of any calibration pulser's azimuth are rejected
- **Wind turbine cut** All events reconstructing to within 5° of the wind turbine's azimuth are rejected.

These cuts were informed from the locations of a variety of infrastructure items in the TestBed coordinate system, which are shown in Figure 6.15 along with the reconstructed locations of all events passing thermal and CW cuts. The reconstructed positions are clearly not uniform and show excesses close to the chosen cuts.

6.7.2. Good times cut

The final set of cuts termed 'goodTimes' criteria reject days of the year that have large numbers of non-thermal events in an attempt to remove repetitive noise associated with human activity. The months of December and January are masked off, and all events within them rejected, as they correspond to the busiest periods of the summer season. Human activity is at its peak during these times involving frequent use of snow-mobiles, radio communication devices and a plethora of other potential noise sources. In order to identify other 'noisy' periods of the year the sample of events passing thermal cuts are used to assess activity levels. The total number of events passing thermal cuts over a three day period is considered. A day is masked off, and all non-thermal events from that day rejected, when the sum of events in that day, the previous and next day passes a threshold of 14. The final cut value was informed from the distribution of events passing thermal cuts in the full analysis sample.

Figure 6.16 shows the distribution of non-thermal events as a function of day of the year for both 2011 and 2012. For 2011 the numbers of non-thermal events

[1] RF signals from above ice sources are refracted when incident upon the air-ice boundary. Any events that reconstructs to an elevation angle greater than the critical angle ($\sim 45°$) is likely to be above the ice.

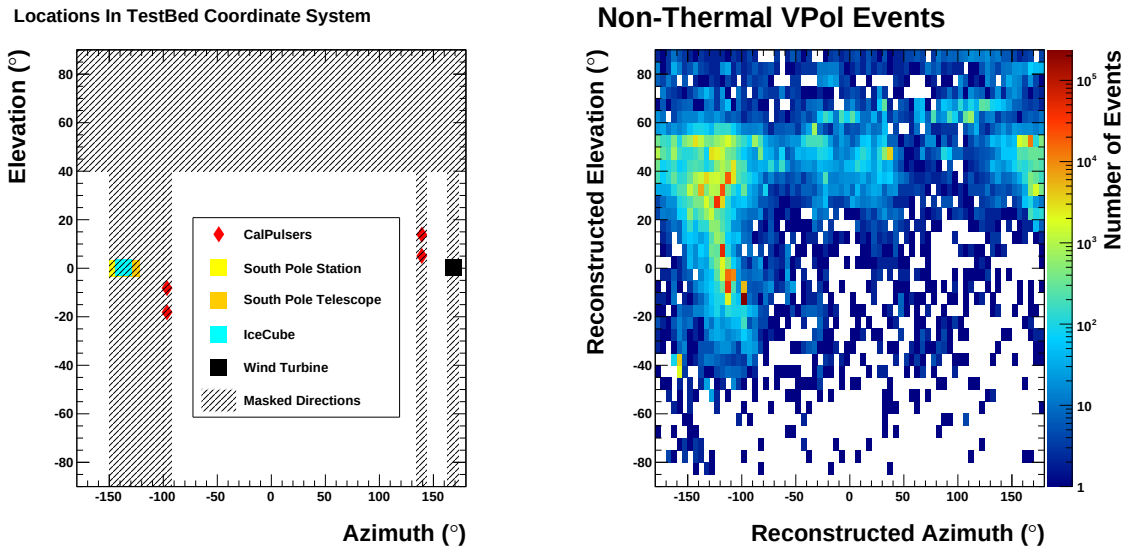

(a) Geometry cuts and infrastructure locations

(b) VPol events passing thermal cuts

Figure 6.15.: The location of a variety of infrastructure items and chosen geometry cuts are shown in (a), along with the reconstructed azimuth and elevation angles for all events passing CW and thermal cuts in (b). The shaded regions in (a) indicate the reconstructed directions that are rejected by geometry cuts.

significantly increase around the beginning and end of the year, with limited activity during the quiet austral winter. 2012, on the other hand, exhibits relatively few periods in which there are very few non-thermal events. Future analysis of these events is expected to be able to categorise them and enable cuts to be designed to remove a large fraction of them.

6.8. Cut summary

The motivation and details of each of the cuts used in this analysis have been described in previous sections, a summary is shown in Table 6.1.

Some of these cuts use a combination of parameters derived from the data, Table 6.2 summarises how these variables are created.

Cut	Description	Cut value
pseudo-χ^2		< 2 pass
powherence	linear combination of CSW-PeakVoltage and SumCorrVals	> 380 pass
CW total bins	non-thermal excess	$= 0$ pass
CW total bins broad band check	lower threshold to check that an excess is broadband	> 20 pass
CWMinProb	minimum value of $\ln(\prod P_i)$	< -400 fail
Geometry	downward pointing cut	$\theta > 40$ fail
Geometry	South Pole infrastructure cut	50° wide region in ϕ
Geometry	Calibration pulser cut	$\pm 5^\circ$ region in ϕ
Geometry	Wind turbine cut	$\pm 5^\circ$ region in ϕ
goodTimes	Summer months cut	Reject events from December and January
goodTimes	Noisy times cut	Reject day when total events passing thermal cuts in previous current and next day > 14

Table 6.1.: Summary of the cuts used in this analysis.

Variable	Description / Value
pseudo-χ^2	Measure of goodness of reconstruction
powherence	$\alpha + \text{CSWPeakVoltage} + \beta * \text{SumCorrVals}$
α	1.143
β	1.571×10^3
CSWPeakVoltage	Peak voltage in the coherently summed wave (CSW)
SumCorrVals	Sum of CorrVals for antennas in CSW.
CorrVals	Peak correlation value of antenna with CSW minus that antenna
CWMinProb	minimum value of $\ln \prod P_i$

Table 6.2.: Summary of the cuts used in this analysis.

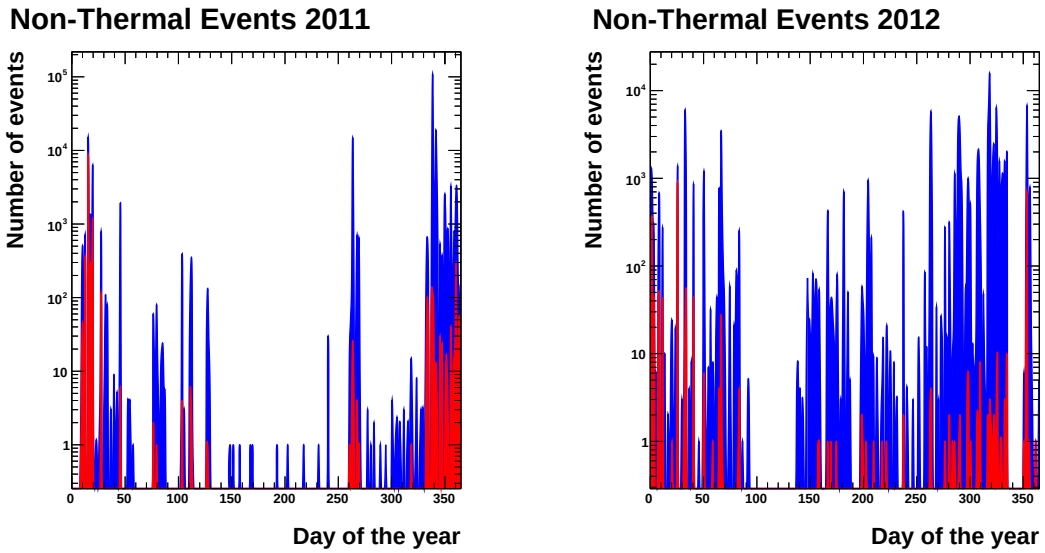

(a) 2011 non-thermal events

(b) 2012 non-thermal events

Figure 6.16.: Shown in blue are the number of events per day passing thermal cuts and in red the subset of these events also passing geometry cuts.

6.9. Cut results

The final analysis cuts and the result of their application to the analysis data sample are detailed in Table 6.3 resulting in no events passing all cuts. The thermal cuts, pseudo-χ^2 and powherence, remove the majority of events illustrating that the main background in the TestBed is of a thermal origin. Carrier wave contaminated events appear to make up a large fraction of the remaining events, with daily weather balloon launches expected to contribute a significant number of events year round.

The observation of no candidate neutrino events needs to be placed into context. In order to do so it is necessary to assess two important factors: firstly the efficiency of detecting simulated neutrinos using these cuts and secondly an estimation of the expected number of background events. The background estimation is discussed in Section 6.11, and the analysis efficiency in Section 6.10.

Total	3.77×10^8		
Cut	Number passing either polarisation		
Data Quality	3.67×10^8		
pseudo-χ^2	3.51×10^8		
	VPol		
cut	passing in order	rejected if last cut	rejected if first cut
pseudo-χ^2	3.49×10^6	-	-
powherence	1.20×10^6	318	2.29×10^6
CW totalBins	4.42×10^5	1.69×10^3	2.37×10^6
CW minProb	1.74×10^5	217	2.12×10^6
Geometry	7.12×10^3	318	2.76×10^6
goodTimes	0	7.12×10^3	2.81×10^6
	HPol		
cut	passing in order	rejected if last cut	rejected if first cut
pseudo-χ^2	1.20×10^6	-	-
powherence	5.14×10^5	144	6.89×10^5
CW totalBins	4.31×10^5	44	4.46×10^5
CW minProb	1.47×10^5	4.47×10^3	1.00×10^6
Geometry	9.15×10^3	183	9.03×10^5
goodTimes	0	9.15×10^3	9.63×10^5

Table 6.3.: Analysis cuts applied to the '90% sample'.

6.10. Analysis efficiency

The efficiency of passing candidate neutrino events for each cut individually and in turn is a fundamental factor in calculating an observed flux or limit using the TestBed data. Dedicated simulated neutrino data sets were generated using the official ARA simulation, AraSim, at fixed neutrino energies. The simulation accounts for the interaction of incident neutrinos in the ice, the emission of Askaryan radiation, the propagation of this radiation through the ice and the full TestBed trigger and digitisation chains. As such these data sets are the best representation available of what neutrino signals will look like in the TestBed data set and can be used to test the effectiveness of the cuts developed and described previously.

The analysis passing efficiency is shown in Figure 6.17 as a function of neutrino energy and also signal to noise ratio (SNR), which is calculated as follows:

$$\mathrm{SNR} = \frac{|V_{peak}|}{\mathrm{RMS}} \tag{6.13}$$

where $|V_{peak}|$ is the absolute value of the peak voltage of each waveform. SNR values are calculated for all antennas and the smallest value in each polarisation compared, the larger of the two is then used as a measure of signal strength (as the RF impulse in events can have an arbitrary polarisation angle, so the brighter of the two polarisations is used).

In addition to the simulated neutrino data sets the analysis cuts were checked against calibration pulser events in two ways. The first was using the calibration pulser waveforms as recorded in the TestBed to check whether any of the non-calibration pulser specific cuts (i.e. excluding the calibration pulser geometry cuts) rejected a large fraction of these events. The second was to take the average of many calibration pulser waveforms, the averaging used to remove event-to-event noise, then add them to minimum bias waveforms with a variety of scalings applied. These events were used as a data driven simulation of attenuated calibration pulser waveforms, with the underlying minimum bias event used as thermal noise. This second category of events were again used as a cross-check for the selected cut values and their efficiencies.

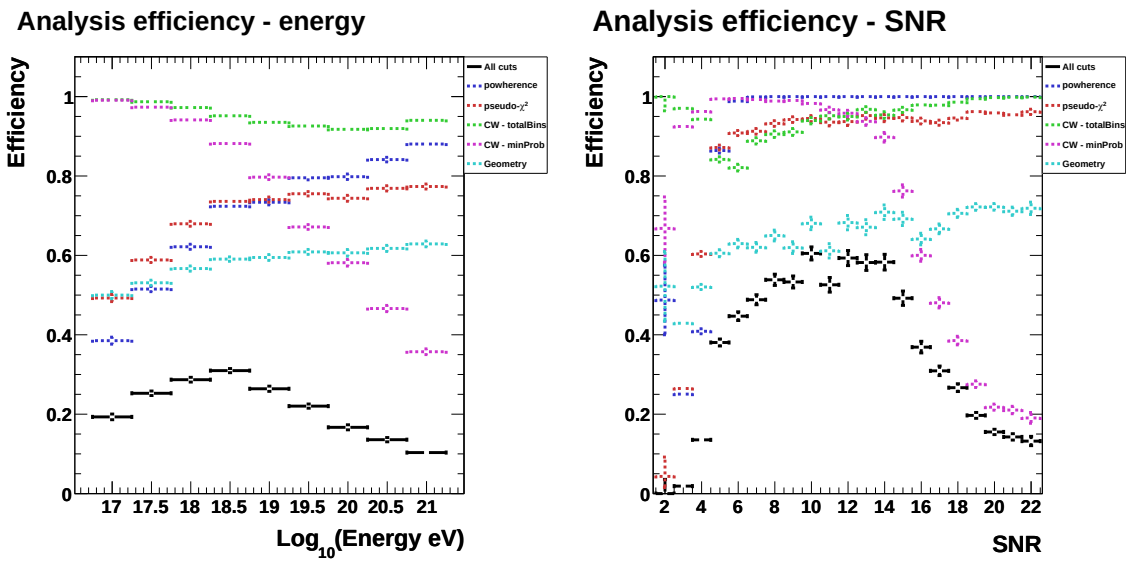

(a) Analysis efficiency as a function of neutrino energy

(b) Analysis efficiency as a function of SNR

Figure 6.17.: Analysis efficiency for simulated neutrino events. The solid black line shows the efficiency after applying all cuts and the dashed lines show efficiencies for individually applied cuts.

The efficiencies in Figure 6.17 are calculated for application of cuts individually (dashed lines) and for all cuts applied. The goodTimes cut is not accounted for as it acts as a scaling of the detector live time. The overall analysis efficiency ranges between 15% and 30% across the range of energies simulated, whereas it peaks at 60% for moderately large values of SNR. Two factors influence the low efficiency as a function of neutrino energy: firstly the simulation produces large numbers of events with very small SNR which mainly consist of upward fluctuations in the thermal noise, and secondly compression in the digitisation chain leads to a drop off in efficiency at high SNR.

The drop off in efficiency at high SNR is driven by very large frequency domain amplitudes being produced when particularly bright signals are propagated through the simulated signal chain. As a result the CW minProb cut removes a large fraction of these very bright events. A possible cause for this effect is incomplete modelling of compression in the signal chain. When large voltages are input to amplifiers in the signal chain they can push them into a non-linear mode of operation, causing smaller amplification of large voltages than small ones. It was found that the simulation does

not account for this effect in the manner observed in bright events in the TestBed data set, for example the 2011 VPol calibration pulser.

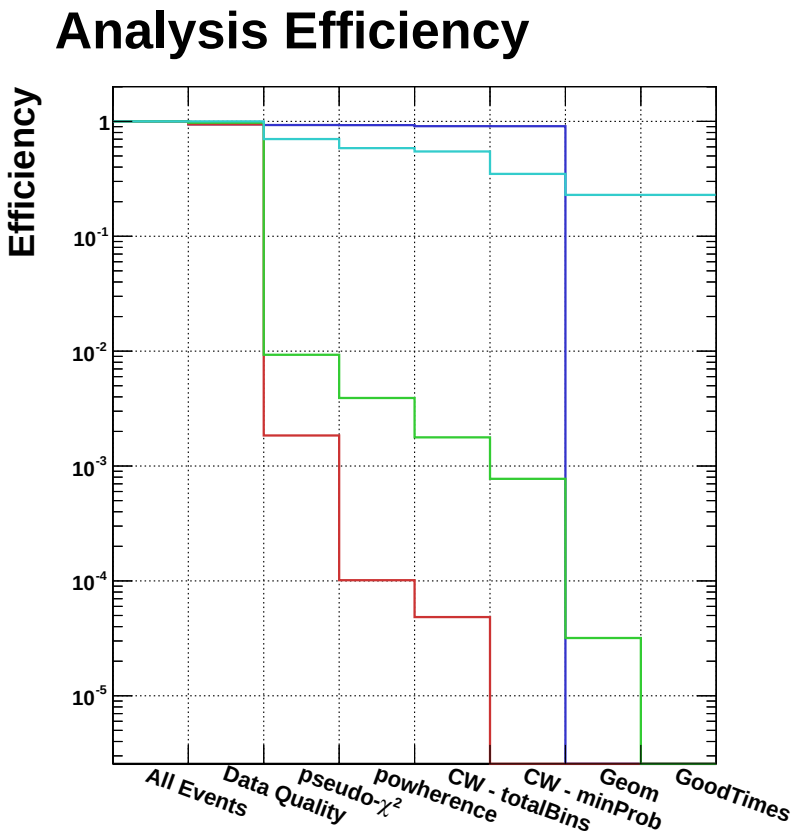

Figure 6.18.: The passing efficiency of analysis cuts applied in turn for various data types. The data types are: calibration pulser events (blue), simulated neutrino events (grey), non-calibration pulser events in the '90% sample' (green) and minimum bias events from the 'burn sample' (red).

The cumulative effect on passing efficiency is shown for a variety of data types in Figure 6.18. The largest drop in efficiency for passing calibration pulser events comes from the data quality issues suffered for $\sim$ 1 month in 2012, affecting 5.8% of these events. Of the events not suffering data quality issues the analysis is 96.3% efficient up to the application of geometry and good times cuts, the first of which specifically removes events reconstructing to the calibration pulser locations.

6.11. Background estimation

Due to the differences in signal type it is necessary to estimate backgrounds separately for thermal noise and anthropogenic events. In order to estimate the expected number of events that pass all cuts the distributions for events failing particular cuts, specific to signal type, were considered and extrapolated from fits to data.

For thermal noise events the value of powherence was considered the most relevant parameter and proved to follow closely an exponentially falling distribution. Two samples of events passing all but the powherence cuts were assessed: the first being all minimum bias events and the second all non-calibration pulser events, in both cases pseudo-χ^2 and anthropogenic cuts are applied. The former sample of events is the purest sample of thermal noise since no trigger condition is met for these events, however a scaling must be applied when accounting for the increase in statistics between this data set and the full set of non-calibration pulser events. The distribution of powherence in VPol and HPol and fits to the falling edge approaching the cut value are shown in Figure 6.19 for the minimum bias sample.

Extrapolating the fits to each of the thermal samples of events provides two estimates for the number of events passing the powherence cut and errors associated with the fits themselves. Since fits to both samples were in good agreement to data the average of the two was taken as the estimate for the background, with the error taken to be the difference between the two. The result is an expected number of background events of a thermal origin to be $0.33^{+0.15}_{-0.15}$.

For the case of anthropogenic signals it is much more difficult to ascertain the expected background levels since these events are from an unidentified mixture of sources, which are not necessarily expected to follow well described distributions. In order to estimate this type of background some of the geometric cuts, designed to remove anthropogenic signals, were relaxed to see whether events failing these cuts could be used to extrapolate past the cut values. Figure 6.20 shows events failing the geometric cut around the azimuthal direction of the IceCube Laboratory and South Pole Station.

A fit was made to the azimuthal distribution of VPol events failing the geometric cut and extrapolated past the cut point at -150. The same procedure was attempted with the results of HPol reconstruction and cuts, but proved difficult to extrapolate so

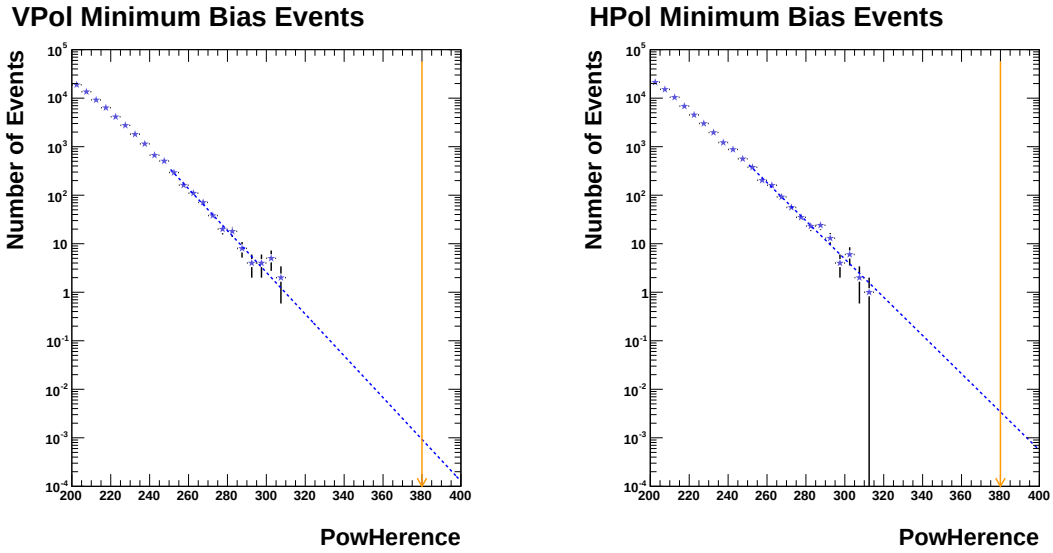

(a) VPol minimum bias events **(b)** HPol minimum bias events

Figure 6.19.: Thermal noise sample taken from the minimum bias data set is fitted to an exponential function for VPol and HPol events. This is then used to extrapolate beyond the chosen cut value (orange line) to estimate the expected number of thermal events passing the cut in the analysis data sample.

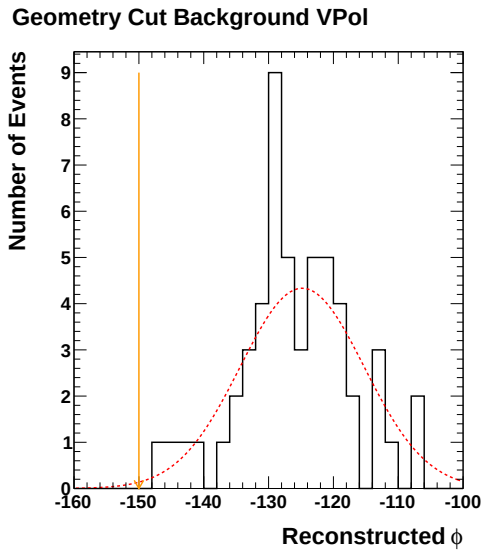

Figure 6.20.: VPol events failing the IceCube Laboratory and South Pole geometry cut via their reconstructed azimuth. The geometric cut is indicated by the orange line at -150 and by the vertical shaded region between 100° and 150° in Figure 6.15 (a).

the estimate for VPol events was also used for HPol. The result of the anthropogenic background estimation was $0.44^{+0.60}_{-0.32}$. Taken in combination with the estimation of thermal backgrounds, the final estimate for the number of background events passing analysis cuts was $0.77^{+0.31}_{-0.43}$.

Chapter 7.

Results

After application of all cuts to the analysis data set (the '90% sample' described in Section 6.2) there are no surviving events and hence no neutrino candidates. Therefore there is no evidence found for a flux of ultra-high energy (UHE) neutrinos detected in the TestBed, hence a model independent limit can be calculated. Upon relaxing the anthropogenic geometry and goodTimes cuts 4 periods of heightened solar activity, caused by solar flares, were identified and detected in the TestBed. The events associated with this solar activity are discussed in Section 7.1, followed by the calculatation of an UHE neutrino limit.

7.1. Solar flare events

The TestBed analysis detected $\sim 6 \times 10^3$ events, that reconstructed close to the location of the sun, associated with 4 solar flare events. Since these events reconstructed to locations above the ice, they failed the geometric cuts in the TestBed analysis. These events are, however, of interest since they are consistent with originating from a distant source and can be used to assess the performance of signal reconstruction. The first solar flare event observed occurred on 15^{th} February 2011, the remaining 3 events were in 2012 on 5^{th} March, 7^{th} March and 21^{st} of November respectively.

Figure 7.1 shows, for one solar flare period, the averaged frequency domain power relative to a baseline taken from a nearby run, and is typical of the 4 vertically polarise (VPol) antennas during this period. Broadband power excesses are seen in the 4 VPol antennas, which seem to be associated with increased X-ray fluxes as

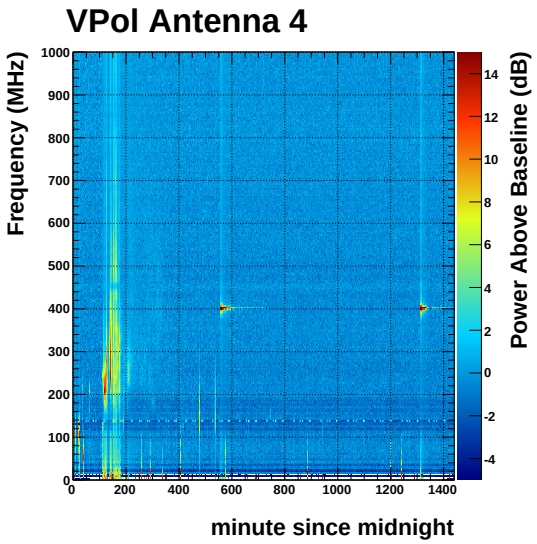

(a) Power above baseline

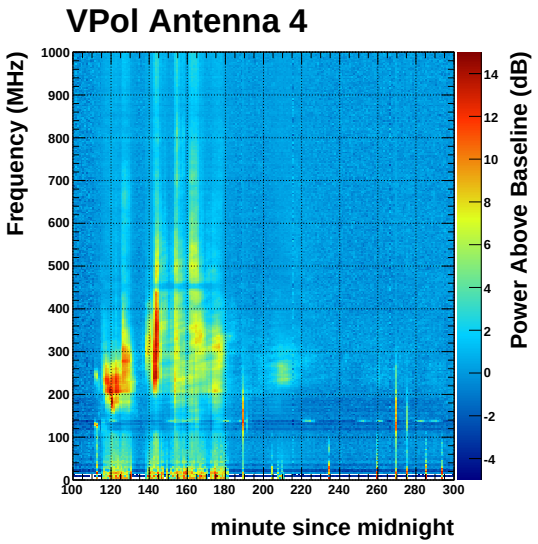

(b) Power above baseline

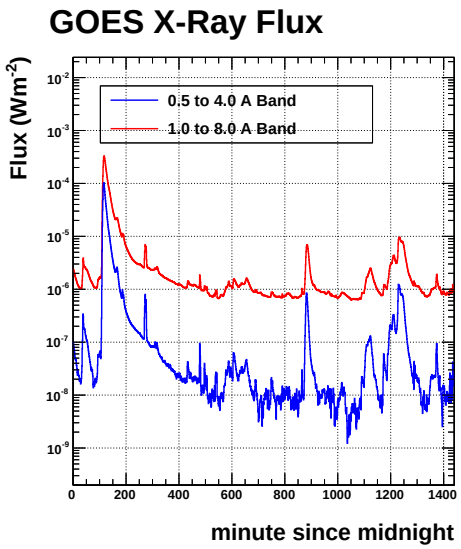

(c) GOES X-ray flux

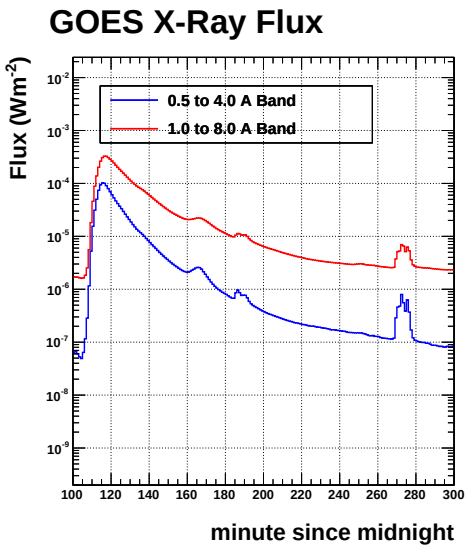

(d) GOES X-ray flux

Figure 7.1.: The solar flare event on 15$^{\text{th}}$ February 2011 as observed in the TestBed. The average frequency domain power in one VPol antenna, taken over a minute period, is compared to a baseline produced from a nearby run in (a) and (b). This is compared to the X-ray flux measured by the GOES satellite system [76] (c) and (d), divided into two bands: $(0.5 - 4.0) \times 10^{-10}$m band (blue), and $(1.0 - 8.0) \times 10^{-10}$m band (red). (a) and (c) show the these quantities over the course of a day, whereas (b) and (d) focus on the period of largest activity.

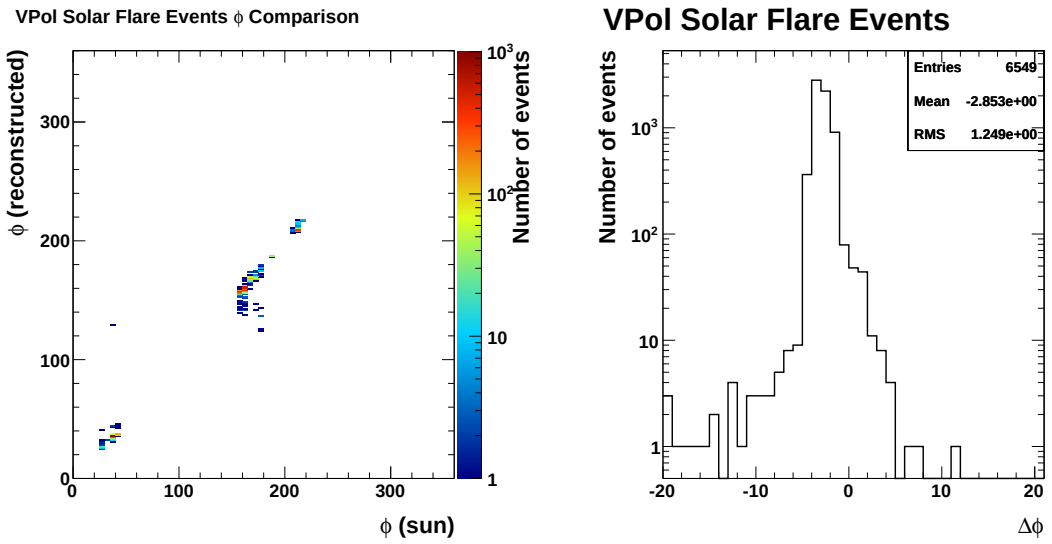

(a) Reconstructed ϕ versus sun ϕ **(b)** ϕ residuals

Figure 7.2.: Comparison of reconstructed azimuth of solar flare events with the calculated position of the sun.

observed by the GOES satellite system [76]. There is a clear coincidence of elevated power levels in the VPol antennas with the largest peak in the GOES X-ray flux at 120 minutes past midnight, however the correlation with subsequent peaks in the GOES flux is not so clear. This is typical of the periods of heightened solar activity that correlate with events observed in the TestBed. Although the small scale time structure of the GOES data is not fully replicated in the recorded events all of these periods do see elevated solar activity and events that reconstruct back to the sun's location in the sky. Inspection of the events' waveforms did not show obvious impulsive signals, however these events exhibit similar waveforms in multiple antennas resulting in non-thermal correlations and the ability to reconstruct their source.

An attempt was made to use these events to determine the reconstruction's accuracy, since the sun's position can be calculated given the time the events were recorded. Figure 7.2 shows the result of comparisons between the azimuth of the sun and reconstructed events.

7.2. Live Time

In order to calculate the sensitivity to a neutrino flux the integrated live time for the analysis data set must be calculated. Due to limitations in the TestBed DAQ there is a period after a trigger is asserted, and event readout initiated, for which it is not possible to form another trigger. This is due in part to the finite digitisation and readout time of the LAB3 digitisers utilised in the TestBed, however this corresponds to only 50µs per event [1]. The majority of the dead-time in the TestBed is understood to be associated with event readout through the DAQ firmware and software, and is rate dependent.

The TestBed electronics monitors dead-time via a 10MHz clock. The dead-time is calculated by taking the number of clock cycles for which the trigger is unavailable in each GPS second, which can be scaled by the clock frequency to recover the fractional dead-time for each second of operation. The total live time for the analysis data set can then be calculated by taking the difference between the number of GPS seconds for which the TestBed operated and the sum of independent dead times recorded.

Having obtained an integrated live time for the entire analysis data set a final adjustment must be made to take into account the goodTimes cut described in Section 6.7.2 and periods for which a significant fraction of events are rejected due to data quality issues. The goodTimes cut masks off periods of the year that are associated with elevated levels of human activity and non-thermal events passing analysis cuts. All runs that are masked off by this cut have their integrated live time removed from the final live time measurement. Similarly any days of the year that have a sizeable fraction of events suffering from data quality issues have their integrated live time removed from the total.

The fractional live time during each day of the year, as well as the resulting integrated live time are shown in Figure 7.3. The integrated live time for 2011 shows a number of features. The first two are the lack of live time between days 0-30 and 335-365 caused by the good times cut, which masks off all events from the first and last month of the year. The same feature is present in the 2012 integrated live time. There is also a period between days 260-270 in 2011 which is masked off, again by the good times cut, but this time due to the large number of events

[1] At a trigger rate of 2Hz this would correspond to a fractional dead-time of 0.01%.

passing the thermal cuts during this period. In 2012 there are other periods in which there is no accumulated live time, notably between days 95-130 and from day 260 onward. The former period is due to a data quality issue which rendered waveform data unusable. The latter corresponds to a large period of time in which there were elevated trigger rates in the TestBed, which corresponds to the drop in fractional live time in Figure 7.3 (d), and large numbers of events passing thermal cuts. This lead to this large period of data being rejected by the good times cuts.

After accounting for data quality issues and applying the goodTimes criteria the integrated live time is 167.4 days for 2011 and 50.2 days for 2012. The resulting live time for the full analysis data set is calculated to be 217.6 days.

7.3. Effective area

The effective area is calculated for the TestBed from simulated neutrino data sets produced using AraSim as follows. Firstly the effective volume, V_{eff}, is calculated for each energy bin using a large sample of simulated events ($\sim 10^6$):

$$V_{eff} = \frac{V_{cylinder}}{N} \sum_{i=1}^{N_{passed}} w_i \tag{7.1}$$

where $V_{cylinder}$ is the volume of a cylinder centred on the detector in which events are simulated, N is the total number of events thrown and $\sum_{i=1}^{N_{passed}} w_i$ is the sum of weights for triggered events. By throwing neutrinos within a cylinder around the detector the computational burden of producing simulation data sets is greatly reduced. The data sets analysed were produced using cylinders of $3-4$km radii, with larger radii chosen for higher energy neutrinos, well above the attenuation length of radio waves in the ice [77].

The effective volume is then converted into an effective area, A_{eff}, by using the approximation:

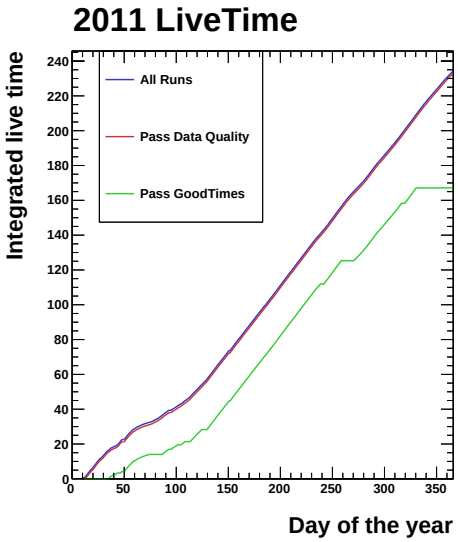

(a) Live time for 2011

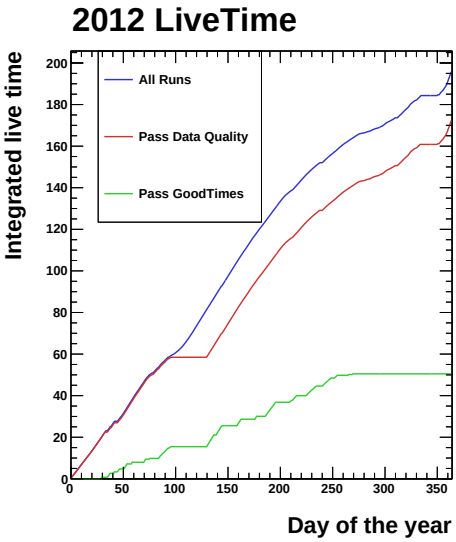

(b) Live time for 2012

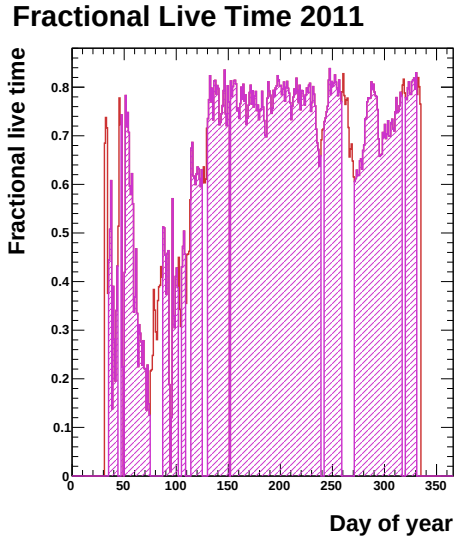

(c) Live time for 2011

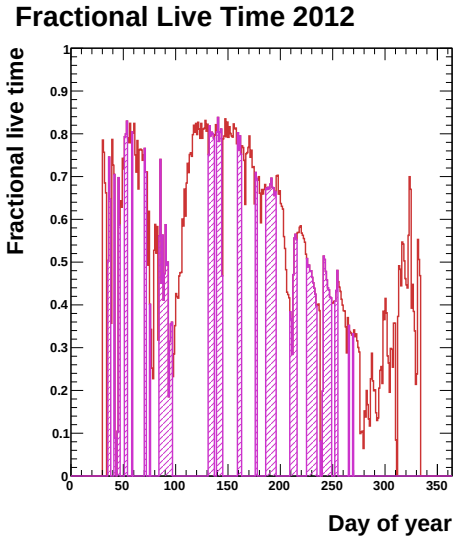

(d) Live time for 2012

Figure 7.3.: Integrated live time for 2011 and 2012. Three lines are shown for each year: live time for all runs, live time for all runs with the majority of events passing data quality checks and finally live time for all runs passing both data quality and goodTimes criteria. Fractional live time is also shown for all runs (red) and for those passing data quality and goodTimes criteria (shaded magenta).

$$A_{eff} \approx \frac{V_{eff}}{l_{int}} \tag{7.2}$$

where l_{int} is the neutrino interaction length in ice [77].

7.4. Neutrino flux limit

An upper limit can be placed on the neutrino flux using the following equation:

$$E_\nu F(E_\nu) \leq \frac{N_{90}}{4\pi A_{eff}(E_\nu) T_{live} \varepsilon(E_\nu) \ln(10)} \tag{7.3}$$

where T_{live} is the total live time of the TestBed analysis data set, $\varepsilon(E_\nu)$ is the energy dependant analysis efficiency, N_{90} is the 90% upper confidence limit on the number of neutrino events excluded by the analysis, which is taken as 2.3 from Poisson statistics, and the factor ln(10) accounts for log-scale energy binning [77]. The flux is defined as follows:

$$F(E_\nu) = \frac{dN}{dE_\nu d\Omega dt dA} \tag{7.4}$$

The live time and effective area are calculated as discussed in Section 7.2 and Section 7.3, and the analysis efficiency taken from application of cuts to simulated neutrino events shown in Figure 6.17.

The resulting UHE neutrino flux limit is shown in Figure 7.4 along with limits from other experiments and a range of theoretical predictions. In addition to the limit calculated from this analysis of two years of TestBed data a projected sensitivity is shown for ARA-37 using a modified simulation [2] . A number of factors lead to the much improved sensitivity reached by the full array. First and foremost is the

[2] At the time of analysis data sets were not available for neutrino energies above 10^{20}eV for the projected ARA-37 sensitivity.

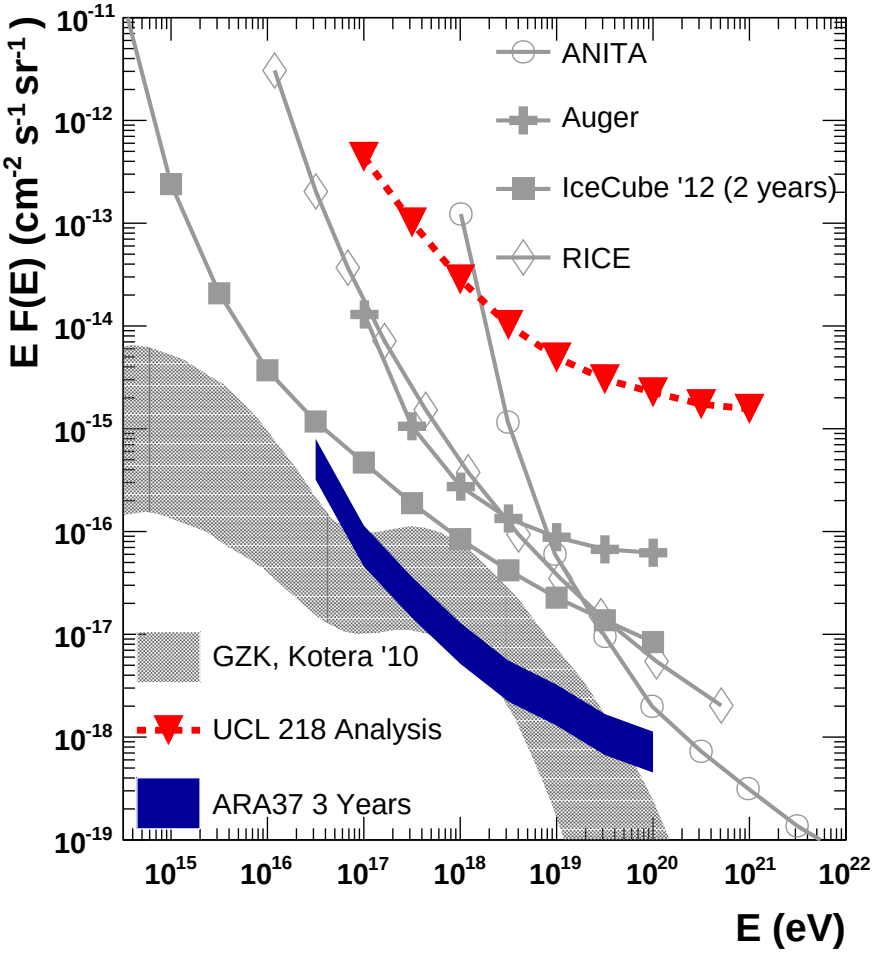

Figure 7.4.: TestBed analysis neutrino flux limit from this analysis 'UCL 218', along with limits from ANITA [62], Auger [78], RICE [79] and IceCube [80]. The shaded band indicates a range of flux predictions from [81] using a variety of assumptions about sources and production mechanisms.

increase in number of stations from one to thirty seven, each acting as a stand-alone neutrino detector, operated over a longer period. Deployment of stations at $\sim$ 200m depths has been demonstrated over the 2012-2013 season and all further stations are expected to deployed at similar depths. By placing the antennas well below the firn layer of ice (which extends $\sim$ 150m below the surface in the vicinity of the south pole), in which the changing index of refraction leads to ray-bending and shadowing effects, an increase in effective volume and area per station of the order ~ 3 over the TestBed is expected.

The expected level of anthroprogenic backgrounds is also anticipated to reduce greatly as the vast majority of future stations will be much further from the south pole station and associated human activity. In addition the receive antennas being placed much deeper in the ice will lead to a reduction in the signal size of above ice noise sources in each station. As well as the reduction in signal size the increased number of receive antennas, from four per polarisation to eight, is expected to significantly improve background identification and rejection. Anthropogenic signal removal via masking off noisy periods of time, which is performed via the goodTimes cut described in Section 6.7.2, leads to large periods of time being removed from the analysis sample and greatly reduces the total detector live time in this analysis. The expected reduction in number anthropogenic signals associated with location will be complemented with better characterisation of such signals and removal as the radio environment becomes better understood.

The re-designed data acquisition electronics and software (DAQ) will also lead to improvements in sensitivity. The updated DAQ allows for triggered event buffering such that the expected dead time associated with event readout described in Section 7.2 will be effectively nil.

Chapter 8.

Conclusions

The analysis of data from 2011-2012 taken with the ARA prototype TestBed detector yields a model independent limit on the UHE neutrino diffuse flux between energies of $10^{17} - 10^{21}$eV shown in Figure 7.4. No neutrino candidate events were found on an expected background of $0.77^{+0.31}_{-0.43}$. Upon relaxing some of the cuts, designed to remove anthropogenic signals, a large number of events reconstructing to the Sun were observed, correlated with increased levels solar activity. Although these signals are not of an impulsive nature they demonstrate the TestBed's ability to detect sources other than calibration signals.

Work was undertaken in developing and delivering the data acquisition systems for the first full ARA stations, which were successfully deployed during 2012-2013. These new stations have much improved functionality over the TestBed and are expected to have significantly increased sensitivity to neutrino fluxes. Under reasonable assumptions the limit derived from analysis of TestBed data was extrapolated to a full 37 station array operating over a 3 year period. The projected sensitivity for the full array is an order of magnitude better than current experiments and should be able to make the first observations of GZK neutrinos when completed.

In the event of making an observation of GZK neutrinos with the full ARA array there are a number of improvements that would be desirable to determine both the energy and direction of propagation of neutrino primaries. Energy measurements would serve as a means to confirm that the observed events are produced by ultra-high energy neutrinos and could be used to make some determination of the energy dependent flux. Directionality of incident neutrino candidates would be of great use

in ruling out possible background sources as well as being used to determine whether there are any preferential directions of origin for these particles.

A detailed calibration study would be necessary to determine the relationship between the measured RF power received in antennas and the energy of the incident neutrino. It could be possible to make measurements in both the laboratory, and the field, of the power radiated by calibration pulsers. These studies could be either linked to experiments conducted at SLAC that measured Askaryan emission [6] , or to dedicated beam studies using ARA equipment. Once deployed these calibration sources could be used to set the scale of received power in the receive antennas in ARA stations, which in turn can be linked to the energy contained in the particle cascade and hence incident neutrino energy. This could be achieved with the current design for calibration antennas local to each station, however with additional work before and after installation to determine the power radiated by these sources. In addition it would be of great use to have sources that are able to illuminate all ARA stations simultaneously, being used as a common, well known source with which to calibrate. The preferred solution would be to deploy high power antennas at the centre of the array at great depth such that they are able to illuminate even the farthest stations.

In order to better determine incident neutrino directions detailed studies would be necessary to determine the received power scales between the horizontally and vertically polarised antennas. This would be needed as the ratio of power between polarisations can be used to determine which part of the Cherenkov cone was sampled in a neutrino event, and hence would be able to constrain the event topology and incident neutrino direction. In addition it would be of great use to better determine the positions of receive antennas in ARA stations, as this would lead to improvements in reconstruction of incident RF signal direction and distance from the station. Both of these improvements could be made with more calibration data from well known pulsers illuminating the ARA stations. A calibration system that would allow for movement of calibration antennas vertically within a calibration hole would be advantageous. This would allow for a large increase in the number of calibration source locations without the expense of additional drilling. This could be achieved with an electronically controlled system that raises or lowers the antenna by known distances, however this would require a large amount of effort to develop and deploy reliably.

With these suggested improvements ARA 37 would provide an excellent window through which to make the first observation of GZK neutrinos.

Appendix A.

Thermal Coincidence Trigger

The ARA trigger system looks at the rate at which N of M antennas have power above a threshold in a time window t. A power threshold is set for each antenna which results in a rate of threshold passing R_{single} from thermal noise. The thresholds are set such that they acheive the same R_{single} rate for each of the antennas contributing to the trigger. This section will describe how the global trigger rate (the rate at which the trigger condition is met) can be determined from the singles rate R_{single}.

A.1. Single antenna, N hits

Let us first consider the case of a single antenna which has "hits" at the rate Γ, where a hit is the power received in the antenna crossing the threshold. The rate for N hits is then given by the hit rate multiplied by the probability that the hit is followed by $N-1$ additional hits within the coincidence window t:

$$\Gamma_N = \Gamma P_{N-1} \tag{A.1}$$

The probability of $N-1$ hits, P_{N-1}, is given by Poisson statistics. The expected number of hits in a coincidence window t is:

$$\mu = \Gamma t \tag{A.2}$$

$$\tag{A.3}$$

hence

$$P_N = e^{-\mu}\frac{\mu^N}{N!} \tag{A.4}$$

so we have

$$\Gamma_N = \Gamma e^{-\mu}\frac{\mu^{N-1}}{(N-1)!} \tag{A.5}$$

A.2. M antennas, N with one or more hits

The rate for single hits among M antennas is

$$\Gamma_1 = M\Gamma \tag{A.6}$$

The rate at which we have N out of M antennas with one or more hits is the rate at which single hits occur (Γ) multiplied by the number of ways of choosing $N-1$ antennas, weighted by the probability that $N-1$ antennas have one or more hits and the rest have none. The probability for zero hits in an antenna is:

$$P_0 = e^{-\mu} \tag{A.7}$$

and the probability for non-zero hits is therefore $1 - P_0 = 1 - e^{-\mu}$. Combining the combinatoric term (from the ways of choosing $N-1$ antennas) and the probability factors we arrive at:

$$\begin{aligned} P_{N-1} &= \frac{(M-1)!}{(M-N)!(N-1)!}(1-P_0)^{N-1}(P_0)^{M-N} \\ &= C_{M-1}^{N-1}(1-P_0)^{N-1}(\frac{1}{P_0})^{N-1}(P_0)^{M-1} \\ &= C_{M-1}^{N-1}(\frac{1}{P_0}-1)^{N-1}(P_0)^{M-1} \\ &= C_{M-1}^{N-1}(e^{\mu}-1)^{N-1}(P_0)^{M-1} \end{aligned} \tag{A.8}$$

So the rate at which the trigger condition is met is then:

$$\Gamma_N = M\Gamma P_{N-1} = M\Gamma C_{M-1}^{N-1}(e^{\mu}-1)^{N-1}(P_0)^{M-1} \tag{A.9}$$

The trigger condition is also met when more than N out of M antennas, so a full treatment would sum the contributions from N of M, $N+1$ of M, ..., M of M antennas. We will only consider the N of M term for now, as the other terms can be found by changing the value of N.

If the hits are rare then we have $\mu \ll 1$, then $e^{\mu}-1 \to \mu = \Gamma t$ and $P_0 \to 1$. So we have:

$$\begin{aligned} \Gamma_N &= M\Gamma\frac{(M-1)!}{(M-N)!(N-1)!}\Gamma^{N-1}t^{N-1} \\ &= N\frac{M(M-1)!}{(M-N)!N(N-1)!}\Gamma^N t^{N-1} \\ &= NC_M^N\Gamma^N t^{N-1} \end{aligned} \tag{A.10}$$

Now using the nomenclature from the ARA trigger systems we have $R_{global} = \Gamma_N$ and that $\Gamma = R_{single}$ we have:

$$R_{global} = NC_M^N R_{single}^N t^{N-1} \tag{A.11}$$

(A.12)

It is clear that as the number of coincident antennas N increases the rate falls by a factor of $R_{single}t$. In the case of the TestBed the singles rate $R_{single} \sim 10$kHz and the time window $t = 100$ns so the contributions fall by a factor of a thousand. Equation (A.12) is therefore a good leading order approximation to the trigger rate expected from thermal noise in ARA stations using such a coincidence trigger.

Bibliography

[1] V. Hess, *Nobel Lectures, Physics 1922-1941* (Elsevier Publishing Company, 1965).

[2] G. A. Askaryan, JETP **21**, 658 (1965).

[3] G. A. Askaryan, JETP **14**, 441 (1962).

[4] D. Saltzberg *et al.*, Phys. Rev. Lett. **86**, 2802 (2001).

[5] P. W. Gorham *et al.*, Phys. Rev. D **72**, 023002 (2005).

[6] ANITA Collaboration, P. W. Gorham *et al.*, Phys. Rev. Lett. **99**, 171101 (2007).

[7] B. T. Cleveland *et al.*, The Astrophysical Journal **496**, 505 (1998).

[8] SNO Collaboration, Q. Ahmad *et al.*, Phys.Rev.Lett. **89**, 011301 (2002), nucl-ex/0204008.

[9] (Super-Kamiokande Collaboration), Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998).

[10] B. Pontecorvo, Soviet Journal of Experimental and Theoretical Physics **26**, 984 (1968).

[11] Z. Maki, M. Nakagawa, and S. Sakata, Progress of Theoretical Physics **28**, 870 (1962).

[12] S. Hatakeyama *et al.*, Phys. Rev. Lett. **81**, 2016 (1998).

[13] Super-Kamiokande Collaboration, Y. Ashie *et al.*, Phys. Rev. D **71**, 112005 (2005).

[14] K2K Collaboration, M. H. Ahn *et al.*, Phys. Rev. D **74**, 072003 (2006).

[15] MINOS Collaboration, P. Adamson *et al.*, Phys. Rev. Lett. **101**, 131802 (2008).

[16] T2K Collaboration, K. Abe *et al.*, Phys. Rev. Lett. **107**, 041801 (2011).

[17] F. P. An *et al.*, Phys. Rev. Lett. **108**, 171803 (2012).

[18] Double Chooz Collaboration, Y. Abe *et al.*, Phys. Rev. Lett. **108**, 131801 (2012).

[19] RENO Collaboration, J. K. Ahn *et al.*, Phys. Rev. Lett. **108**, 191802 (2012).

[20] Particle Data Group, J. Beringer *et al.*, Phys.Rev. **D86**, 010001 (2012).

[21] NOvA Collaboration, D. Ayres *et al.*, (2004), hep-ex/0503053.

[22] LBNE Collaboration, C. Adams *et al.*, (2013), hep-ex/1307.7335.

[23] WMAP Collaboration, G. Hinshaw *et al.*, (2012), astro-ph/1212.5226.

[24] S. A. Thomas, F. B. Abdalla, and O. Lahav, Phys. Rev. Lett. **105**, 031301 (2010).

[25] C. Kraus *et al.*, Eur.Phys.J. **C40**, 447 (2005), hep-ex/0412056.

[26] V. N. Aseev *et al.*, Phys. Rev. D **84**, 112003 (2011).

[27] P. S. B. Dev, S. Goswami, M. Mitra, and W. Rodejohann, Phys. Rev. D **88**, 091301 (2013).

[28] J. Hinton, New Astronomy Reviews **48**, 331 (2004), 2nd {VERITAS} Symposium on the Astrophysics of Extragalactic Sources.

[29] J. Holder *et al.*, AIP Conf.Proc. **1085**, 657 (2009), astro-ph/0810.0474.

[30] M. Actis *et al.*, Experimental Astronomy **32**, 193 (2011), astro-ph/1008.3703.

[31] J. Linsley, Phys. Rev. Lett. **10**, 146 (1963).

[32] J. Abraham *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **523**, 50 (2004).

[33] S. Hoover *et al.*, Phys. Rev. Lett. **105**, 151101 (2010).

[34] V. Berezinsky, A. Gazizov, and S. Grigorieva, Phys. Rev. D **74**, 043005 (2006).

[35] K. Greisen, Phys. Rev. Lett. **16**, 748 (1966).

[36] G. T. Zatsepin and V. A. Kuzmin, JETP Lett. **4**, 78 (1966).

[37] Pierre Auger Collaboration, A. Aab *et al.*, (2013), astro-ph/1307.5059.

[38] High Resolution Fly's Eye Collaboration, R. U. Abbasi *et al.*, Phys. Rev. Lett. **100**, 101101 (2008).

[39] Pierre Auger Observatory, J. Abraham *et al.*, Phys. Rev. Lett. **104**, 091101 (2010), astro-ph/1002.0699.

[40] K. Kotera and A. V. Olinto, araa **49**, 119 (2011), astro-ph/1101.4256.

[41] V. S. Beresinsky and G. T. Zatsepin, Phys. Lett. **B28**, 423 (1969).

[42] E. Waxman and J. N. Bahcall, Phys.Rev. **D59**, 023002 (1999), hep-ph/9807282.

[43] J. N. Bahcall and E. Waxman, Phys.Rev. **D64**, 023002 (2001), hep-ph/9902383.

[44] R. Engel, D. Seckel, and T. Stanev, Phys. Rev. D **64**, 093010 (2001).

[45] D. Hooper, A. Taylor, and S. Sarkar, Astroparticle Physics **23**, 11 (2005).

[46] R. Gandhi, C. Quigg, M. H. Reno, and I. Sarcevic, Astroparticle Physics **5**, 81 (1996).

[47] F. Halzen and S. R. Klein, Review of Scientific Instruments **81**, 081101 (2010), astro-ph/1007.1247.

[48] J. A. Aguilar *et al.*, Physics Letters B **696**, 16 (2011), astro-ph/1011.3772.

[49] KM3NeT Collaboration, M. de Jong, Nucl.Instrum.Meth. **A623**, 445 (2010).

[50] E. Andres *et al.*, Astropart.Phys. **13**, 1 (2000), astro-ph/9906203.

[51] IceCube Collaboration, M. G. Aartsen *et al.*, Phys. Rev. Lett. **111**, 021103 (2013).

[52] IceCube, M. Aartsen *et al.*, Science **342**, 1242856 (2013), 1311.5238.

[53] IceCube Collaboration, M. Aartsen *et al.*, (2014), astro-ph/1405.5303.

[54] S. Danaher and the ACoRNE Collaboration), Journal of Physics: Conference Series **81**, 012011 (2007).

[55] R. Lahmann, Nuclear Instruments and Methods in Physics Research Section A:

Accelerators, Spectrometers, Detectors and Associated Equipment **604**, S158 (2009), {ARENA} 2008.

[56] V. Aynutdinov *et al.*, ArXiv e-prints (2009), astro-ph/0910.0678.

[57] N. Kurahashi, J. Vandenbroucke, and G. Gratta, Phys.Rev. **D82**, 073006 (2010), hep-ex/1007.5517.

[58] S. Boeser *et al.*, (2008), astro-ph/0807.4676.

[59] ANITA Collaboration, P. W. Gorham *et al.*, Phys. Rev. Lett. **103**, 051103 (2009).

[60] S. Barwick, D. Besson, P. Gorham, and D. Saltzberg, Journal of Glaciology **51**, 231 (2005).

[61] I. Kravchenko, D. Besson, and J. Meyers, Journal of Glaciology **50**, 522 (2004).

[62] P. W. Gorham *et al.*, Phys. Rev. D **82**, 022004 (2010).

[63] I. Kravchenko *et al.*, Astroparticle Physics **19**, 15 (2003).

[64] S. R. Klein, IEEE Transactions on Nuclear Science **60**, 637 (2013), astro-ph/1207.3846.

[65] B. Mercurio, *A simulation of the ANITA experiment with a focus on secondary interactions*, PhD thesis, Ohio State University, 2009.

[66] A. Connolly, Internal ARA note: ELOG 91.

[67] P. Allison *et al.*, Astroparticle Physics **35**, 457 (2012).

[68] G. Varner *et al.*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **583**, 447 (2007).

[69] P. Allison, International Cosmic Ray Conference **4**, 350 (2011).

[70] A. Vieregg, *A Search for Astrophysical Ultrahigh-Energy Neutrinos Using Radio Detection Techniques*, PhD thesis, University of California, Los Angeles, 2010.

[71] G. Varner, IRS2 Datasheet.

[72] E. Hong, Internal ARA note: DocDB 776.

[73] S. Hoover, *A Search for Ultrahigh-Energy Neutrinos and Measurement of Cosmic Ray Radio Emission with the Antarctic Impulsive Transient Antenna*, PhD thesis, University of California, Los Angeles, 2010.

[74] M. Mottram, *A Search for Ultra-high Energy Neutrinos and Cosmic-Rays with ANITA-2*, PhD thesis, University College London, 2012.

[75] A. Romero-Wolf, *Search for Ultra High Energy Astrophysical Neutrinos with the ANITA Experiment*, PhD thesis, University of Hawai'i, 2010.

[76] P. C. Chamberlin, T. N. Woods, F. G. Eparvier, and A. R. Jones, Next generation x-ray sensor (XRS) for the NOAA GOES-R satellite series, in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, , Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 7438, 2009.

[77] ARA Collaboration, P. Allison *et al.*, (2014), 1404.5285.

[78] The Pierre Auger, J. Abraham *et al.*, (2009), astro-ph/0906.2347.

[79] I. Kravchenko *et al.*, Phys. Rev. D **73**, 082002 (2006).

[80] IceCube Collaboration, M. G. Aartsen *et al.*, Phys. Rev. D **88**, 112008 (2013).

[81] K. Kotera, D. Allard, and A. V. Olinto, JCAP **10**, 13 (2010), astro-ph/1009.1382.