FACULTY OF SCIENCE
UNIVERSITY OF COPENHAGEN

Master Thesis

Eva Brottmann Hansen

Early Atmospheric Muon Rejection with IceCube-PINGU

Academic supervision by Jason Koskinen

July 31, 2016

Abstract

The development of a generic event selection for the Precision IceCube Next Generation Upgrade (PINGU) is presented. PINGU is a planned infill extension to the IceCube Neutrino Observatory at the South Pole, designed to lower the energy threshold of IceCube to $\sim$1 GeV. With the PINGU subarray, IceCube will perform precision measurements of the atmospheric neutrino oscillation parameters, θ_{23} and Δm_{32}, and determine the neutrino mass ordering at 3σ significance within 4 years. The cosmic ray induced muon flux constitutes the largest source of background for neutrino searches with IceCube at the trigger level. To demonstrate the physics reach of PINGU, the detector's ability to reject the this background must be addressed. This thesis presents a set of algorithms, developed to identify and reject events caused by atmospheric muons at the earliest levels of data processing. At the final level, a muon rejection factor of 10^5 is achieved, as is needed to reduce the muon rate below the total neutrino rate.

Acknowledgments

I would like to extend my most sincere thank you! to the IceCube group here at NBI. First and foremost I would like to thank my supervisor, Jason Koskinen, for the confidence he has shown me throughout the project and the enthusiastic guidance I received whenever I needed it. Your interest in my work, expectations to my achievements and nudges to make me present regularly at the working group calls has continually motivated me to work my hardest. Thank you Jason, for the pushes and the support I've received and for introducing me to the original Star Wars movies.

I also owe a great thanks to Michael and Morten, whom I have shared an office with for the last year and a half. In the beginning I hardly knew five terminal commands, but you both sat by my side providing debugging assistance at every error message. I soon enough learnt to deal with error messages myself, and have since profited a great deal from your exhaustive knowledge on all kinds of IceCube related topics. Your delightful company, the afternoon cakes, Friday beers, Monday workout sessions, music breaks, etc. have made me look forward to coming in to the office every single day.

Contents

Chapter 1

Introduction

Of the known particles in the Universe, neutrinos are the second most numerous after photons. Yet, due to their elusive nature, they are still the least understood. The unknowns that we are aware of include the absolute mass scale of neutrinos, the correct mass ordering, whether neutrinos violate charge-parity conservation, and whether they are Dirac or Majorana particles. We know that neutrinos interact through the weak force in one of the three flavor states (ν_μ, ν_e, and ν_τ) determined by the co-participating charged lepton, and that they transition between these states as they propagate. The flavor state transitions are known as *neutrino oscillations* and require that at least two of the neutrinos have non-zero masses. The phenomenon is thereby the only experimental confirmation of the incompleteness of the Standard Model, in which neutrinos are massless.

The idea of neutrino oscillations was prompted by the discovery of oscillations in the neutral kaon system $K^0 \rightleftharpoons \bar{K}^0$ along with the discovery of a second neutrino flavor - the muon neutrino - in 1962 [1]. The hypothesis was supported by consistent observations of a deficit of electron neutrinos from the Sun [2], a mystery which persisted for 30 years. In 1998, Super-Kamiokande presented the first conclusive experimental evidence of muon neutrino disappearance [3]. Four years later, the Solar Neutrino Observatory (SNO) published evidence for conversion of solar electron neutrinos to muon or tau neutrinos [4] by combining their results with the Super-Kamiokande measurements. The spokespersons of the two collaborations were awarded the Nobel Prize for physics in 2015.

Since the discovery of neutrino oscillations, several precision experiments have confirmed the standard oscillation interpretation developed mainly by Pontecorvo, Gribov, Maki, Nakagawa and Sakata in the 1960s and early 1970s [5] [6]; neutrino oscillations arise due to the three flavor states (ν_μ, ν_e, and ν_τ) being mixtures of three mass eigenstates (ν_1, ν_2, and ν_3). Though the square of the mass differences and the ordering of ν_1 and ν_2 have been measured, the complete neutrino mass ordering, i.e. whether the ν_3 mass state is the heaviest or the lightest of the three, has not been established.

Determining the neutrino mass ordering (NMO, also referred to as the neutrino mass hierarchy) is a vital step towards understanding the origin of neutrino masses and fixing their absolute scale. The mechanisms that could generate these tiny neutrino masses require "new physics". Possible explanations include Majorana neutrinos (i.e. neutrinos that are their own anti-particles), the existence of heavy sterile neutrinos, or additional Higgs particles - none of which has so far been observed. The result of the NMO measurement will improve the sensitivity of the searches for these phenomena as well as searches for leptonic charge-parity violation which could play a role for the baryon-antibaryon asymmetry.

The IceCube Neutrino Observatory is the world's largest neutrino telescope, built deep in the glacial ice at the South Pole to detect astrophysical neutrinos at the TeV to PeV scale. With a low energy sub-detector, called DeepCore, IceCube is also sensitive to atmospheric neutrino oscillations at the energy scale of $\mathcal{O}(10\,\mathrm{GeV})$. The current configuration has proven to be competitive with dedicated neutrino oscillation experiments. The proposed infill extension, the Precision IceCube Next Generation Upgrade (PINGU), will lower the energy threshold to a few GeV, targeting signatures of the NMO and precision measurements of the atmospheric neutrino oscillation parameters, θ_{23} and Δm_{32}. The upgrade will also expand the energy range of indirect dark matter searches, improve the sensitivity to supernovae neutrino detection, and enable the first ever neutrino-based tomography of the earth. The scientific case of PINGU, regarding all the above mentioned studies, relies on efficient atmospheric muon rejection which is the focus of this thesis. Specifically, I have developed a generic set of algorithms to identify and separate events caused by atmospheric muons and neutrinos, respectively, at the earliest levels of data processing. Through out the thesis, this is referred to as the *NBI event selection*.

Following this introduction, Chapter 2 offers a brief review of the status of neutrino physics, thereby forming the theoretical foundation and motivation for the PINGU experiment. The third chapter describes the atmospheric neutrino flux which provides the signal for oscillation related studies with PINGU, and the atmospheric muon flux which constitutes the main background for all neutrino observations. The operation principle of the IceCube detector and the proposed PINGU detector is covered in Chapter 4 along with the software tools used to simulate PINGU data. Chapter 5 deals with the software used to reconstruct events in the detector. Chapter 6 is dedicated to the details of a new method of selecting data for analysis with PINGU simulation and eventually data. The thesis concludes with a summary of the results obtained with the NBI event selection and gives an outlook on the future implementation of the event selection in PINGU analyses.

Chapter 2

Neutrinos in the Standard Model and Beyond

The theoretical framework of particle physics, known as the *Standard Model*, was developed over decades and formulated in the 1970's in the form that we know today. It has been immensely successful in providing a self-contained description of most known phenomena in elementary particle physics. The fact that neutrinos have non-zero masses is the only experimentally observed departure from the Standard Model. Expanding the model to account for neutrino masses requires further understanding of the their nature. This chapter briefly covers the theoretical foundation of the physics motivation and detection mechanism of PINGU. We start with a summary of the Standard Model with a focus on the properties of neutrinos. This is followed by a description of the different types of interactions between neutrinos and nucleons at the PINGU energy scale. The third section addresses how neutrino oscillations arise from the mixing of massive states. From there, the prospective physics reach of the PINGU detector is outlined, focusing on how the instrument will contribute to solving the persistent enigmas of the lepton sector of particle physics. Throughout the chapter, natural units are used ($c = \hbar = 1$).

2.1 The Standard Model Boiled Down

The Standard Model of Particle Physics is a relativistic quantum field theory describing the elementary particles that make up all matter and their interactions. The theory is gauge invariant and built on the symmetry group $SU(3) \otimes SU(2) \otimes U(1)$ which describes three of the four fundamental forces of nature, the strong and electroweak forces. The most familiar force in our macroscopic lives, gravity, is not included in the Standard Model. Though this is a fundamental incompleteness of the "theory of almost everything", the influence of gravity is negligible at the minuscule level of elementary particle. The coupling strength of gravity to quarks is $\sim 10^{-37}$ times the strength of the weak force.

The strong and electroweak forces are represented as gauge fields with associated field quanta corresponding to the force mediating particles of integer-spin, called *gauge bosons*. Equivalently, matter fields have associated field quanta, known as *fermions*, that have an intrinsic spin of 1/2.

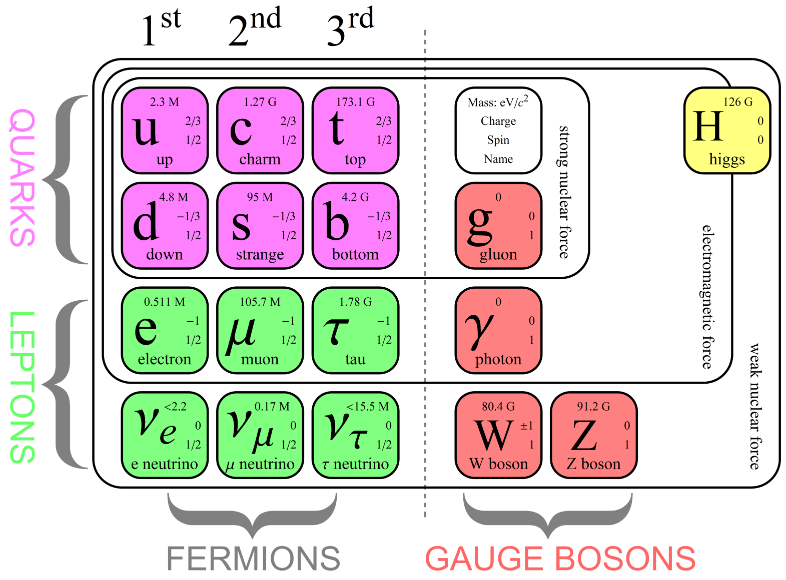

Figure 2.1: Schematic of the standard model of the Standard Model. Taken from [7].

The strong force obeys the $SU(3)$ symmetry which has eight massless generators, called gluons (g). The conserved quantity in strong interactions is the *color* charge, to which gluons couple. As gluons themselves are colored, they have self-interactions which restricts the range of the strong force to the size of a nucleus [8]. Colored objects are subject to "confinement" which only permits colored particles to persist in bound states of zero net color. The strong force has the largest coupling constant of all the forces, which means it is dominant in the presence of color charge. However, the strength of the strong force drops when the distance over which it acts decreases which means that the quarks confined in protons and neutrons are approximately free.

The mediator of the electromagnetic force is the massless photon (γ) which couples to electrical charge with a coupling constant that is about two orders of magnitude weaker than the strong force. Being itself electrically neutral, it has no self-interactions and, therefore, the electromagnetic force is not restricted in range. This means that electromagnetic forces are dominant on macroscopic scales in the presence of electrical charges.

Weak interactions are mediated by the $W^{\pm}$ and Z^0 boson which, unlike photons and gluons, are heavy. The weak boson masses of about 90 GeV limit their range to the subatomic scale. At low momentum transfers, the effective coupling strength is about an order of magnitude smaller than the electromagnetic force, but it increases with energy as the masses become negligible compared to the available energy. Above 100 GeV the electromagnetic and weak force merge into a unified electroweak interaction described by the $SU(2) \otimes U(1)$ gauge group. Independently, the $SU(2)$ group has three generators, one neutral and two oppositely charged, and the $U(1)$ symmetry has one neutral generator - all of which are massless. Spontaneous symmetry breaking of the $SU(2) \otimes U(1)$ group, caused by the Higgs mechanism, gives rise to *electroweak mixing* of the two neutral generators, thereby forming the Z^0 and γ bosons. Likewise, the two charged generators are combined to the two massive charged bosons $W^{\pm}$. The fermions of the Standard Model in turn acquire mass through Yukawa coupling with the Higgs field proportionally to their mass.

The diagram in Figure 2.1, gives an overview of all the Standard Model particles, grouped according to the bosons that they couple to. Fermions are subdivided in to two classes called leptons and quarks, each consisting of six distinct particles. The six quarks, up(u), down(d), charm(c), strange(s), top(t) and bottom(b) are paired in three flavor doublets - three generations - which have identical properties aside from their masses:

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

Quarks have color charge and are, therefore, confined. This means that they can only persist in bound states, *hadrons*, that are collectively color neutral. Quarks also have electrical charge and are, thereby, affected by all the forces.

The lepton group consist of the three charged leptons, the electron (e), muon(μ), and tau (τ), and their corresponding neutral neutrinos ν_e, ν_μ and ν_τ. They are organised in the same manner as the quarks:

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix}, \quad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}, \quad \begin{pmatrix} \tau, \\ \nu_\tau \end{pmatrix} \tag{2.2}$$

The charged leptons interact primarily through the electromagnetic force mediated by the photon γ, whereas the neutrinos have neither electrical nor color charge and therefore only interact via the weak force. According to the Standard Model, neutrinos do not couple to the Higgs Boson and are, in effect, massless. This prediction is consistent with all direct measurements of neutrino masses as well as measurements of neutrino helicities, which show that all neutrinos are left-handed (and anti-neutrinos right-handed). The current strictest limits on the possible neutrino masses, are included in the diagram Figure 2.1 [9]. These are obtained via the energy spectrum of the outgoing electron in β decay by the Mainz [10] and Troitsk [11] experiments. The discovery of neutrino oscillations - first reported by the Super-Kamiokande experiment [3] and the SNO experiment [4] - are currently the only observations not in agreement with the textbook version of the Standard Model. We return to the topic of oscillations in Section 2.3.

It was indicated by experiments at SLAC and CERN [12] that there should be exactly three generations of fermions. As the Z^0 can decay into all fermion-antifermion pairs, provided that the fermion mass is below half the Z^0 mass, measurements of the Z^0 boson's lifetime gives the exact number of fermions with masses below 45 GeV that have weak interactions. Though there could conceivably be quarks with a mass too heavy to affect the Z^0 lifetime, it is hard to imagine that the fourth neutrino would suddenly jump to a mass above 45 GeV. Consequently, any potential additional neutrinos should not have weak couplings: they must either be "sterile" or have non-standard interactions.

2.2 Neutrino Interactions with Matter

While neutrinos can interact with both electrons and nucleons as they propagate through normal matter, the cross section of neutrino-nucleon scattering is by far the larger. Therefore, only neutrino interactions with hadrons will be considered in the following. At the energy range relevant to the present thesis, the total neutrino cross section is very complicated to describe. Below ~1 GeV scattering processes can be understood by considering the nucleon as a coherent object. Above ~20 GeV the neutrino interacts with the constituent quarks individually. PINGU covers an energy region of 1 GeV $<E_\nu<$ 50 GeV,

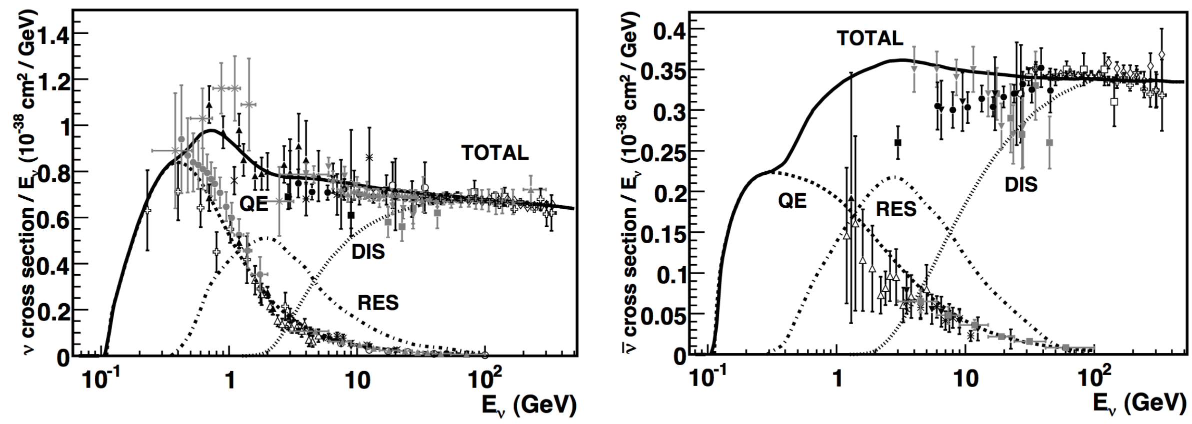

(a) (b)

Figure 2.2: The total neutrino (a) and antineutrino (b) cross section for CC interactions divided by energy is plotted as a function of energy. The three components, quasi-elastic scattering, resonance production, and deep inelastic scattering are also plotted separately [13].

part of which is typically referred to as the "transition region" as it corresponds to the boundary between these two regimes. At neutrino energies of 0.1-20 GeV several interaction mechanisms contribute to the total cross section, and there is no clear way to combine the processes. The interactions fall into three categories [13]:

> *Elastic or quasi-elastic scattering (QE)* is the dominant process below 1 GeV. Neutrinos elastically scatter off a coherent nucleon which stays intact as it recoils. "Quasi-elastic" scattering typically refers to the case where a $W^{\pm}$ boson is exchanged, and the target proton/neutron is converted to a neutron/proton under emission of a charged lepton. In the case of a Z^0 boson exchange, where the nucleon is not converted, the interaction is referred to as "elastic scattering".
>
> *Resonance production (RES)* dominates the cross section around 1-4 GeV. Here the incoming neutrino excites the target nucleon to a short lived baryonic resonance state ($\Delta, N*$), which typically decays to a nucleon and a pion.
>
> *Deep inelastic scattering (DIS)* contributes the most from 10 GeV and up. When sufficient energy is transferred, the constituents of the target nucleon are resolved, and the neutrino scatters off a quark rather than the entire nucleon. The nucleon is completely disintegrated, resulting in a hadronic shower.

The measured contributions from the three scattering processes are plotted in Figure 2.2. The plot shows how QE scattering is dominant below 1 GeV, resonance production takes over from 1-3 GeV, and DIS is dominant above 10 GeV. As predicted for neutrinos scattering off free quarks, the total cross section approaches a linear dependence on energy [14]. The antineutrino cross section is shown to be about half of the neutrino cross section. This discrepancy is due to the helicities of the states involved, as conservation of angular momentum along the interaction axis induces a suppression of backwards antineutrino-quark scattering.

All types of interaction can happen through the exchange of either a $W^{\pm}$ boson, called a charged current (CC), or a Z^0 boson known as a neutral current (NC). In all types of

CC interactions, a charged lepton of the same flavor as the initial neutrino is produced in the final state, whereas in NC interactions a neutrino of the same flavor is found in the initial and final state. Integrating over the full energy region of PINGU, the DIS processes are the most common. The neutral and charged current DIS of a neutrino on a nucleon are illustrated in Figure 2.3.

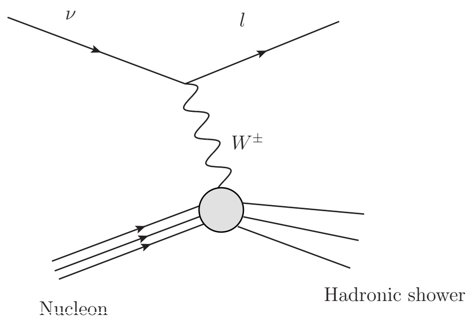

(a) CC: $\nu_l + N \rightarrow l + X$

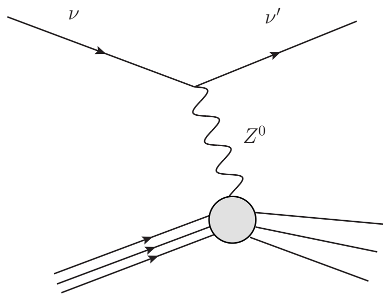

(b) NC: $\nu + N \rightarrow \nu + X$

Figure 2.3: Feynman Diagram of deep inelastic scattering of a neutrino on a nucleon. l denotes the (anti) lepton flavor (e, μ, τ), N is a target nucleon, and X is the remnant of N. The diagrams are drawn using JaxoDraw [15].

In both neutral and charged current interactions, the fragments of the nucleon produce a hadronic cascade. The outgoing neutrino in NC interactions leaves no trace, making it impossible to distinguish the interacting neutrino flavor. In CC interactions, however, the outgoing charged lepton can reveal the flavor of of the incoming neutrino. The signatures of the different neutrino flavor interactions are discussed in Section 4.4.

2.3 Neutrino Oscillations

According to the Standard Model, neutrinos interact in three different flavor states determined by the co-participating charged lepton in the exchange of a $W^{\pm}$ boson. The original formulation of the Standard Model assumes all three neutrino states to be massless. However, observations prove that neutrinos can transform from one flavor to another in flight with a certain probability depending on the energy of the neutrino and the distance traveled. The phenomenon can only occur if neutrinos interact in a different basis than the one in which they propagate, i.e. the eigenfunctions of the Hamiltonian. Furthermore, the Hamiltonian eigenstates must have different non-zero masses which makes neutrino oscillations the only experimental contradiction of Standard Model predictions.

2.3.1 The two-flavor model

To demonstrate how non-zero masses of the propagating states lead to oscillations between the flavor states, we consider the case of only two neutrino types. The simplicity of the solution to the two-flavor problem makes it a useful tool to understand the phenomenology of neutrino oscillations and make crude predictions. The eigenstates of the Hamiltonian, the two "mass states" ν_1 and ν_2, must be linear combinations of the two flavor states, α and β, in which neutrinos interact. In order to preserve orthonormality of the states, we express the flavor mixing as: [1]

$$\nu_1 = \cos(\vartheta)\nu_\beta - \sin(\vartheta)\nu_\alpha \quad , \quad \nu_2 = \sin(\vartheta)\nu_\beta + \cos(\vartheta)\nu_\alpha \tag{2.3}$$

Here ϑ corresponds to a mixing angle that must be determined through experiments. According to the Schrödinger equation, an eigenstate of the Hamiltonian in vacuum with the energy E has the time dependence:

$$\nu(t) = \nu(0)\mathrm{e}^{-iEt} \tag{2.4}$$

Choosing ν_α to be the initial state, such that $\nu_\alpha(0) = 1$ and $\nu_\beta(0) = 0$, the time evolution of the mass states goes as:

$$\nu_1(t) = -\sin(\vartheta)\mathrm{e}^{-iE_1 t} \quad , \quad \nu_2(t) = \cos(\vartheta)\mathrm{e}^{-iE_2 t} \tag{2.5}$$

where $E_{1,2}$ is the energy of the state $\nu_{1,2}$. By substituting ν_1 and ν_2 in Equation 2.3 with the expressions in Equation 2.5 and solving for ν_β, the probability for the electron neutrino to convert into a muon neutrino after a time, t, is obtained:

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta} = |\nu_\beta(t)|^2 &= (\sin(\vartheta)\cos(\vartheta))^2 \left(e^{-iE_2 t} - e^{-iE_1 t}\right)\left(e^{iE_2 t} - e^{iE_1 t}\right) \\ &= \sin^2(2\vartheta)\sin^2\left(\frac{E_2 - E_1}{2} t\right) \end{aligned} \tag{2.6}$$

The probability of measuring a ν_β appears sinusoidal, rendering the term 'neutrino oscillations' rather appropriate. Assuming relativistic neutrinos ($m_i << E_i \Rightarrow v \approx c$), we can express the energy difference $E_2 - E_1$ in terms of the mass difference $m_2 - m_1$ and the time in terms of the distance traveled L and arrive at the relation:

$$P_{\nu_\alpha \to \nu_\beta} = \sin^2(2\vartheta)\sin^2\left(\frac{\Delta m^2}{4E} L\right) , \tag{2.7}$$

where $\Delta m^2 \equiv m_2^2 - m_1^2$. The result of the two-flavor case shows that oscillations only can occur if there is mixing of the flavor states ($\vartheta \neq n\pi$, $n \in \mathbb{Z}$) and a mass difference between the Hamiltonian eigenstates states ($m_2 - m_1 \neq 0$). The maximum oscillation probability is obtained when

$$L_{max\,osc} = \frac{2\pi E}{\Delta m^2} \tag{2.8}$$

Of course, all three neutrino states must be included in the treatment of oscillations to get the accurate probabilities. As it turns out, one mass splitting is significantly larger than the other, making the two-flavor scheme an excellent approximation in many cases. Using the current best fit value of the mass difference $\Delta m_{32}^2 = 2.5 \times 10^{-3}\mathrm{eV}^2$ [9] and converting back to SI-units we get that the maximum probability of the transition $\nu_\mu \to \nu_\tau$ is obtained at combinations of baselines and energy of the order of magnitude $L/E \sim 500$ km/GeV.

[1] The treatment of two-flavor neutrino oscillations in vacuum follows roughly that of Griffiths [8] and Martin and Shaw [16].

2.3.2 The three-flavor model

Extending the neutrino mixing scheme to the three-flavor case complicates the algebra somewhat. The mixing matrix relating the flavor eigenstates to the mass eigenstates becomes a 3×3 complex matrix, known as the *PMNS Matrix*[2].

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.9}$$

The PMNS matrix $\mathbf{U}$, can be parameterised in terms of three mixing angles (ϑ_{12}, ϑ_{13}, ϑ_{23}) and one complex phase (δ_{CP}):

$$\mathbf{U} = \begin{pmatrix} c_{12}c_{13} & s_{12}sc_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & -c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \tag{2.10}$$

where the shorthand notation $s_{ij} = \sin(\vartheta_{ij})$ and $c_{ij} = \cos(\vartheta_{ij})$ is used. The phase factor, δ_{CP}, is related to CP-violation which has yet only been conclusively observed. However, new data from the neutrino beam experiments, T2K and NOνA, presented at the Neutrino 2016 conference, hint at a value near maximal CP-violation $\delta_{CP} = -\frac{\pi}{2}$ [17] [18].

Deviations from unitarity of the PMNS matrix requires the existence additional neutrinos that mix with the three known ones. As mentioned previously, measurements of the Z^0 boson lifetime have ruled out the existence of additional neutrinos which could participate in the weak interactions. Therefore, precision measurements of the three mixing angles are of great importance to establish or reject the existence of sterile neutrinos or neutrinos with non-standard interactions.

If there are, in fact, only three flavors, the PMNS matrix is unitary and can easily be inverted to express the mass states in terms of the flavor states. The oscillation probabilities can be derived following the same principle as in the two-flavor case: by computing the temporal evolution of an eigenstate of the Hamiltonian, where the initial state is be given by Equation 2.9.

2.3.3 The Neutrino Mass Ordering

With three neutrino states there are two independent mass splittings: Δm_{21} and Δm_{32}. Like in the two-flavor case, the mass splittings occur in cosine or sine squared functions, making oscillation probabilities insensitive to the sign of the (squared) mass difference. In other words, the size of the neutrino masses can be derived from vacuum oscillation studies, but not the relative ordering of the mass states. It turns out that we inhabit a universe where one mass difference happens to be significantly larger than the other: $|\Delta m_{32}| >> |\Delta m_{21}|$, but, we do not know the absolute values of the neutrino masses or even their relative ordering. The question of the neutrino mass ordering (NMO) is illustrated in Figure 2.4.

[2]Named after Bruno Pontecorvo, Ziro Maki, Masami Nakagawa, and Shoichi Sakata.

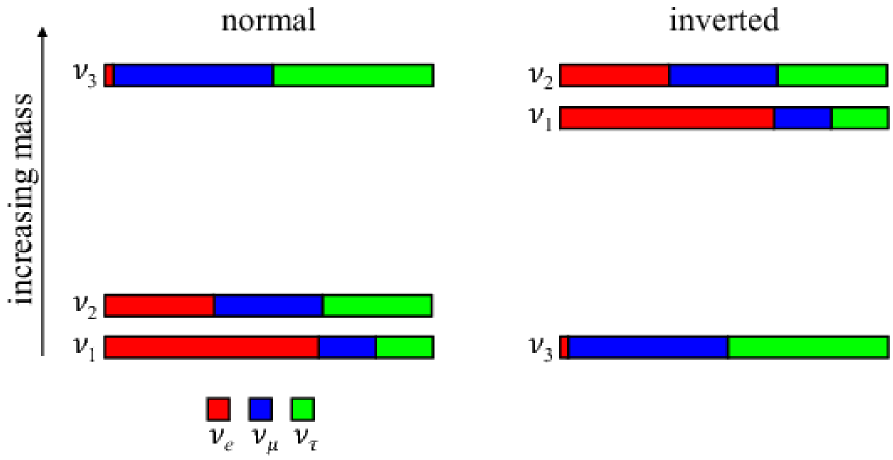

Figure 2.4: The known oscillation parameters allow for two different orderings of the neutrino mass states. Figure modified from [19].

The ordering of neutrino mass states can be extracted if neutrinos travel through sufficient amounts of matter. As electron neutrinos created in nuclear processes in the Sun travel to the surface, the survival probability ($P_{\nu_e \to \nu_e}$) is affected by the electron density depending on the sign of the mass splitting Δm_{21}. Thus, solar neutrino experiments have established the ordering of the first and second mass state ($m_2 > m_1$) through measurements of the transition probability for $\nu_e \to \nu_\mu$.

In the earth, two distinct physical effects influence the neutrino flavor propagation depending on the sign of Δm_{32}. The first is the MSW effect [20] which, due to the electron content in matter, enhances the oscillation probability for either $\nu_\mu \to \nu_e$ or $\bar{\nu}_\mu \to \bar{\nu}_e$, depending on the neutrino mass ordering. The presence of electrons creates an effective potential affecting only the electron neutrino flux, due to CC coherent forward scattering. The same goes for anti-electron neutrinos, but the sign of the potential is opposite. The effective potential introduces a change in the effective mixing angle which enters in the oscillation probability. In 1985 S. P. Mikheyev and A. Y. Smirnov discovered that a resonance exists where the oscillation amplitude approaches one. To fulfil the resonance condition, the sign of the potential and the mass splitting must be the same. The mass ordering can, thereby, be inferred by the observation (or non-observation) of the MSW effect.

The second effect arises from the density transition at the Earth's mantle-core interface where neutrinos passing through this interface can undergo "parametric enhancement" of their oscillation probability [21].

2.4 The Physics Portfolio of PINGU

The Precision IceCube Next Generation Upgrade (PINGU) is a proposed in-fill array for IceCube at the South Pole Station, designed to detect neutrinos in the energy range of roughly 1-10 GeV. Studying atmospheric neutrinos that undergo oscillations as they traverse the Earth, PINGU aims to determine the ordering of the neutrino mass eigenstates and constrain current limits in the mixing parameters θ_{23} and Δ_{32}. The results from PINGU will contribute to illuminate the remaining unknowns in the leptonic sector: the nature of the neutrino (Dirac or Majorana), the extent to which CP symmetry may be violated in the leptonic sector, and possible the presence of a potential sterile neutrino. In addition, PINGU will increase the sensitivity to lower energy neutrinos of astrophysical origin; potentially from accumulated dark matter or supernovae. A very brief overview of the various studies to be conducted with PINGU is provided here:

2.4.1 Precision measurements of oscillation parameters

PINGU aims to measure θ_{23} primarily through the deficit of muon neutrinos in the atmospheric flux. In the three-neutrino framework, at the relevant energies and baselines, ν_μ-disappearance is largely due to transitions of the type $\nu_\mu \rightarrow \nu_\tau$. Those are mainly caused by the splitting of the second and third mass eigenstate. The dependence is evident in the approximate survival probability of the muon neutrino:

$$P_{\nu_\mu \rightarrow \nu_\mu} \approx 1 - \sin^2(2\theta_{23}) \sin^2 \left(\frac{\Delta m^2_{23} L}{4 E_\nu} \right) \tag{2.11}$$

where E_ν is the neutrino energy and L is the distance that the neutrino has travelled. The parameters θ_{23} and Δm_{32} are the least well measured of the mixing parameters. Current global fits pin θ_{23} close to 45$^\text{o}$ [9], suggesting that the third mass eigenstate consists of almost equal contributions from of ν_τ and ν_μ. The case of the exact identity $\theta_{23} = 45^\text{o}$ is referred to as *maximal mixing* and might imply the existence of an unknown fundamental symmetry. In case θ_{23} is not exactly 45$^\text{o}$, determining whether it is slightly higher or lower is of great importance for uncovering the origin of neutrino masses.

2.4.2 The Neutrino Mass Ordering

The ordering of ν_1 and ν_2 is known from solar neutrinos which undergo MSW effect in the sun. Whether the third mass state ν_3 is heavier (normal ordering) or lighter (inverted hierarchy) than the two is yet to be established. The ordering and absolute masses of the neutrinos is of intrinsic interest, but also important for oscillation experiments with sensitivity to leptonic CP violation. Determination of the NMO will influence the interpretation of non-oscillation experiments (neutrino less double beta decay) that are sensitive to the particle nature of the neutrino (Dirac vs. Majorana). Thus the results will indirectly help to test the popular see-saw neutrino mass models and the related mechanism of leptogenesis in the early universe [22]. As mentioned, parametric enhancement and the MSW effect influence the oscillation probabilities of ν and $\bar{\nu}$ differently depending on the ordering of the NMO. The effects become visible around 2-15 GeV, which is seen in the "oscillograms", shown in Figure 2.5. The oscillograms represent the survival probability of a muon neutrino ($P_{\nu_\mu \rightarrow \nu_\mu}$) across a range of neutrino energies and baselines. By lowering the energy threshold to $\sim$1 GeV, PINGU is expected to determine the NMO with a 3σ significance in about 4 years of data taking given the normal ordering scenario [23].

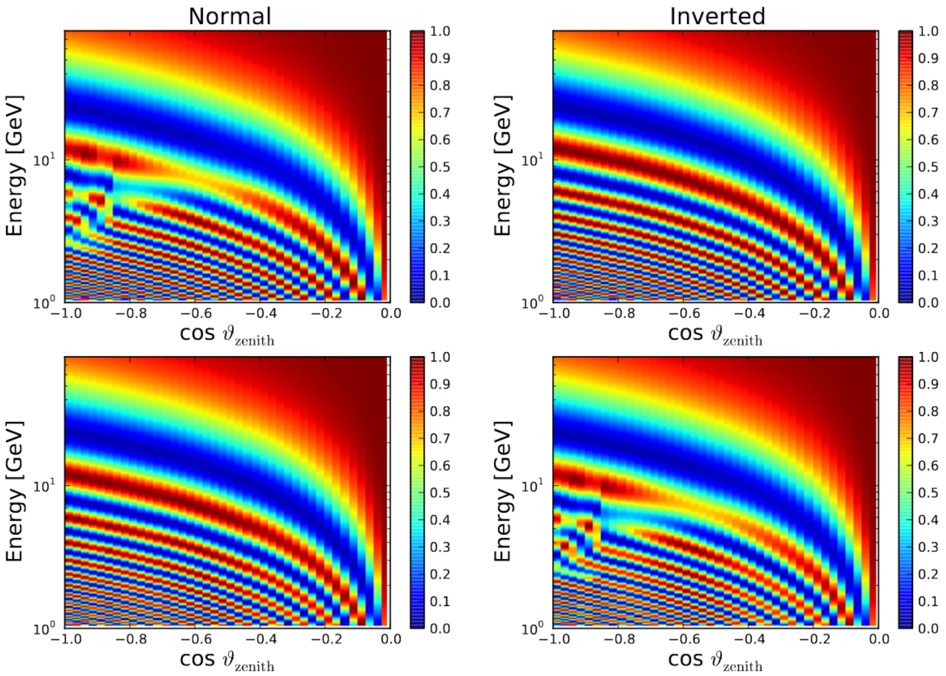

Figure 2.5: The oscillation probability $P_{\nu_\mu \to \nu_\mu}$ as a function of energy and zenith angle for normal and inverted ordering, respectively. Matter effects are seen for neutrinos (top) in the case of normal ordering and for antineutrinos (bottom) if the ordering is inverted. The travel distance is derived from the zenith angle, where $\cos\theta = -1$ corresponds to a path directly through the centre of the Earth.

2.4.3 Tau Appearance

Due to the lack of direct measurements of ν_τ oscillations, unitarity of the PMNS matrix has yet to be established experimentally. PINGU aims specifically at the ill-constrained matrix element $U_{\tau 3}$ through detection of ν_τ-appearance in the atmospheric neutrino flux. The analysis is technically difficult due to the small ν_τ interaction rate relative to the other neutrino flavors as well as the limited ability of PINGU to distinguish neutrino flavors. The size of PINGU will allow for a substantial rate of 3000 CC ν_τ interactions per year [24].

A measured value of $U_{\tau 3} \approx \frac{1}{2}$ would further strengthen the three flavor model and confirm unitarity of the third mass eigenstate. Deviations from $U_{\tau 3} \approx \frac{1}{2}$, on the other hand, would favour an explanation including at least one sterile neutrino or neutrinos with non-standard interactions.

2.4.4 Bonus PINGU Physics

Taking further advantage of the matter effects influencing atmospheric neutrinos that propagate through the Earth, PINGU will perform the first ever neutrino-based Earth tomography. The dependence of neutrino oscillation probabilities on the electron density of the traversed matter allows PINGU to directly probe the chemical composition of the Earth's core which has never been measured.

Lowering the energy threshold with PINGU will enhance current searches for dark matter with IceCube and extend the mass range of testable dark matter candidates. One prominent candidate for dark matter is the Weakly Interacting Massive Particle (WIMP). Given that they interact gravitationally, WIMPs could be thought to accumulate in the Sun, the Earth?s centre, the galactic centre, and dwarf spheroidal galaxies. PINGU will search for an excess in the neutrino flux from WIMPs self-annihilating to neutrinos in the direction of potential dark matter accumulators.

Last but not least, the detector will have enhanced sensitivity to neutrinos bursts from supernovae. These neutrinos have energies much lower than those of relevance to PINGU oscillation studies - $\mathcal{O}(10\,\text{MeV})$ instead of $\mathcal{O}(1\ \text{GeV})$. Therefore, individual neutrino interactions can not be detected. Instead, a core-collapse of a supernova in the Milky Way or a nearby galaxy can be observed as a short-term, coherent increase in the effective noise rate on all DOMs in the detector.

Chapter 3

The Atmospheric Neutrino Source and Background

The PINGU detector will utilise the intense and constant neutrino beam provided by cosmic rays interacting in the atmosphere for the purpose of studying fundamental neutrino properties. This chapter starts with a description of the primary cosmic radiation which leads to an examination of the creation and characteristics of the resulting atmospheric neutrino flux. From there, we study the accompanying atmospheric muon flux which constitutes a main background to any neutrino induced muon search.

3.1 Cosmic Rays

The earth is bombarded by high energy particles known as cosmic radiation arriving from all directions at a rate of $1000(\mathrm{m}^2\mathrm{s})^{-1}$. Today, 100 years after their discovery, the origin of the highest energy cosmic rays is still not understood [25]. They consist mostly of protons (∼90%) and alpha particles (∼9%), but also include a small fraction of heavier nuclei up to uranium as well as elementary particles as electrons, positrons, and neutrinos [26]. The energy spectrum of cosmic rays can be described by a power law of the form

$$\frac{\partial N(E)}{\partial E} \propto E^{-\gamma} \tag{3.1}$$

which spans 12 orders of magnitude from 10^8 eV to 10^{20} eV. The spectrum is relatively featureless, with a spectral index of $\gamma = 2.7$ everywhere except between the "knee" around $E \approx 4$ PeV and "ankle" at $E \approx 10^{18}$ eV. In that region it steepens to $\gamma \approx 3.0$. The all-particle energy spectrum is shown in Figure 3.1.

The primary particles typically interact at a height of about 25 km above sea level with energies far exceeding the binding energy of the air molecules (usually oxygen or nitrogen). The deep inelastic scattering creates a shower of pions and kaons that decay into lighter particles such as muons and neutrinos, see Figure 3.2.

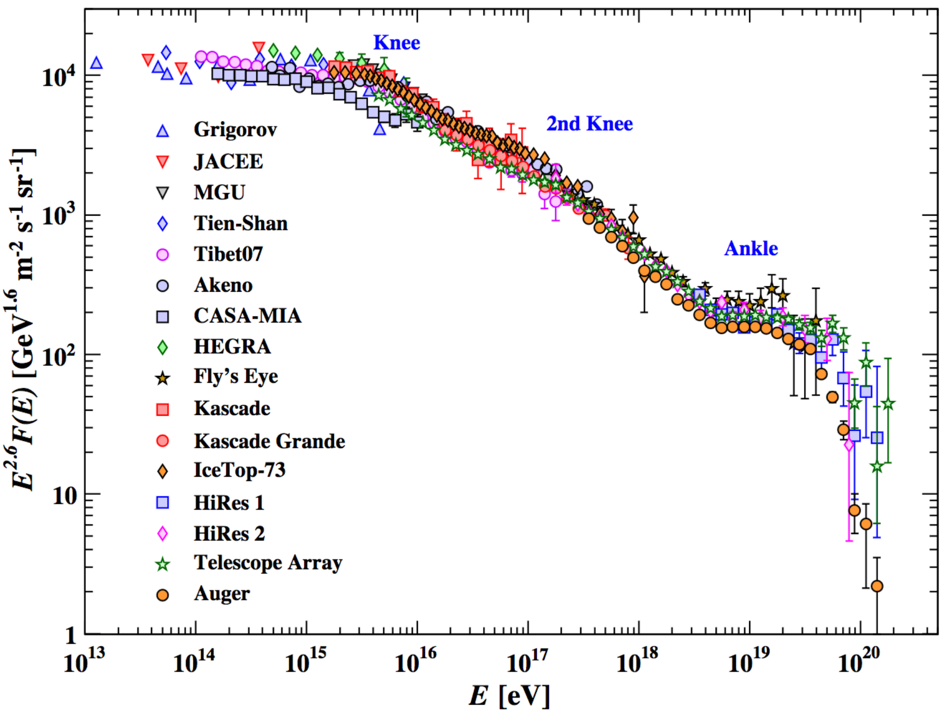

Figure 3.1: Primary cosmic ray spectrum. Figure taken from [27].

A typical interaction chains is:

$$\begin{aligned} p + N \rightarrow X + &\pi^{\pm} \\ &\hookrightarrow \mu^{\pm} + \overset{(-)}{\nu}_{\mu} \\ &\quad \hookrightarrow e^{\pm} + \overset{(-)}{\nu}_{e} + \overset{(-)}{\nu}_{\mu} \end{aligned} \tag{3.2}$$

The fully ionised nuclei, constituting the bulk of cosmic rays, are deflected by galactic magnetic fields, which render their arrival directions mostly randomised. The result is an isotropic flux of secondary electrons, muons and neutrinos reaching the surface of the earth at comparable rates.

Though part of the cosmic ray flux can be attributed to supernovae remnants in our galaxy, sources of the highest energy cosmic rays are still unknown. As mentioned, the primary cosmic rays also have a small neutrino component which, contrary to the charged constituents, points back to their origin. By measuring the arrival directions of primary neutrinos at the TeV to PeV energy scale, IceCube seeks to uncover astrophysical sources of the high energy flux of cosmic rays. For neutrino astronomy, the secondary neutrinos and muons created in the atmosphere are the main sources of background. However, for the studies of fundamental neutrino physics, also conducted with IceCube, the atmospheric neutrino flux constitutes the signals.

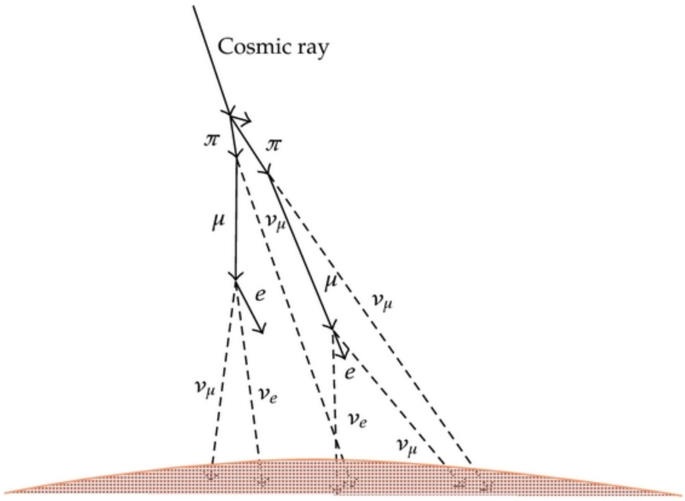

Figure 3.2: Neutrino production in cosmic ray shower. Figure taken from [28].

3.2 The Atmospheric Neutrino Source

The cosmic ray interactions in the atmosphere provide a stable flux of neutrinos with a wide range of baseline-to-energy ratios below $500\frac{\text{km}}{\text{GeV}}$, where signatures of neutrino oscillation become visible. Equation 3.2 indicates that the ν_μ flux is about twice as high as the ν_e flux, since the ν_μ component arises from both the muon production and decay, whereas the ν_e production only comes from the muon decay. However, neutrino oscillations cause a deficit in the ν_μ flux, especially in the up-going direction. Being highly relativistic, the lifetime of secondary particles in the cosmic ray air showers are Lorentz boosted proportionally to their energy, E. Correspondingly, the probability of a secondary particle producing a neutrino as it decays drops with energy. The result is roughly a factor of $1/E$ suppression of the conventional atmospheric neutrino spectrum with respect to the primary cosmic ray spectrum. Several models of the atmospheric neutrino flux exist and they roughly agree on both the absolute flux and flavor type ratios below 100 GeV, where the production is dominated by pions. Comparisons can be found in [29]. Measurements of the atmospheric neutrino flux have been done by IceCube, its predecessor AMANDA, Super-Kamiokande, and Frejus. The results are found to be consistent with the theoretical predictions, but the measurements are not sensitive enough to discriminate between the different flux models, as is shown in Figure 3.3.

The magenta band indicates a predicted component of the atmospheric neutrino flux which has yet to be observed. The, so called, prompt neutrinos are produced by the decay of charmed mesons. Due to the high energy threshold of charmed meson production, prompt neutrinos are only created with energies above 10 TeV. The short lifetimes of charmed meson, compared to pions and kaons, allows them to decay before interacting with an air nucleus. Thus, the energy spectrum of the prompt neutrino flux is expected to follow the primary cosmic ray spectrum, but with a smaller normalization due to the low production probability of charmed mesons in hadron showers [30].

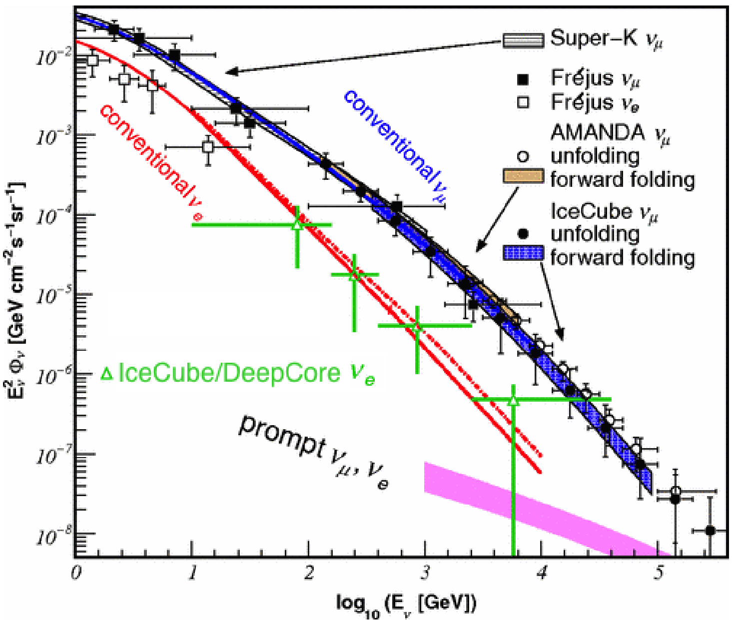

Figure 3.3: The atmospheric muon neutrino and electron neutrino flux. The modelled spectrum of neutrinos from pion decays are shown in the bands labeled "conventional". The figure is modified from [29].

3.3 The Atmospheric Muon Background

Apart from neutrinos, muons above a few hundred GeV are the only products of cosmic radiation which can penetrate deep enough into the Earth to be detected in IceCube. At the surface of the Earth, the integrated muon flux above 1 GeV is around 70 $m^{-2}s^{-1}sr^{-1}$ and the mean energy 4 GeV [31]. As the muons propagate, they deposit energy through continuous ionisation and radiative processes at a rate depending on their energy and the material they cross. Above a few GeV muons are minimum ionising, and deposit about 200 MeV per meter of water (or ice) they traverse. This amounts to an approximate energy loss of 300 GeV for vertical muons in the 1.5 km of ice shielding the IceCube Detector. The muon intensity in water, including contributions from both atmospheric muons and neutrino-induced muons, integrated over energy is shown as a function of depth in Figure 3.4. The contribution from neutrino-induced muons is roughly constant with depth and is seen not to dominate until depths of $\sim$20 km where the intensity plateaus.

In the fiducial region of PINGU, around 2 km below the ice surface, atmospheric muons outnumber neutrino-induced muons by a factor of 10^6. Thus, efficient background rejection is of vital importance for any neutrino induced muon search.

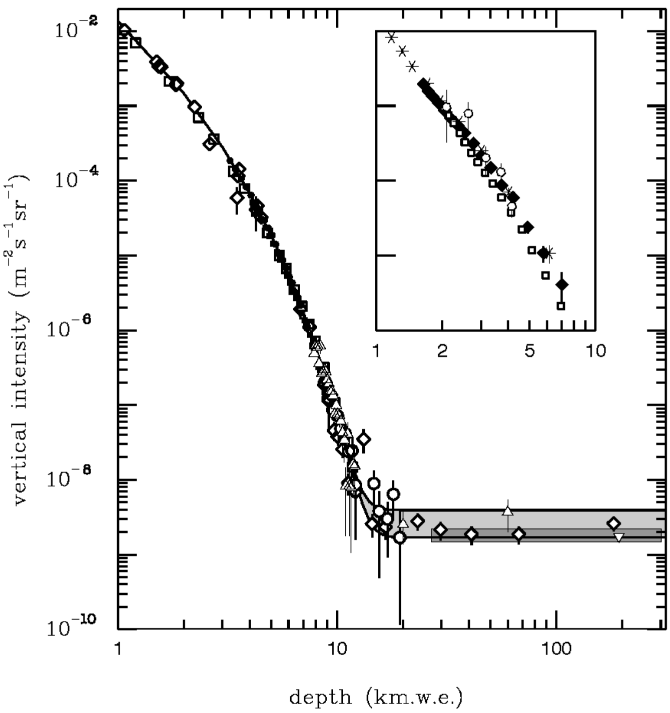

Figure 3.4: The vertical muon intensity as function of depth using the community standard of meters water equivalent [31]. The intensity plateaus where the neutrino-induced muon flux begins to dominate.

Chapter 4

Chapter 4: Neutrino Detection with IceCube-PINGU

The detection principle of the proposed detector upgrade, PINGU, relies widely on the same mechanisms, hardware, and software as its predecessor, IceCube-DeepCore. Hence, Cherenkov based neutrino detection is introduced in the context of IceCube-DeepCore, along with a description of the data acquisition system. Equipped with the necessary terminology, we proceed to explore the proposed PINGU upgrade and the software tools used to simulate PINGU data which also closely align with most DeepCore methods. The final section offers an overview of the different types of events visible in IceCube and the topology by which they can be partially identified.

4.1 The IceCube-DeepCore Detector

Encompassing a cubic kilometre of ice, the IceCube Neutrino Observatory is the world's largest neutrino detector. It consists of 5160 detection units buried kilometres deep in the Antarctic ice at the geographical South Pole. Neutrino interactions are detected indirectly through the blue light emitted when a charged particle traverses a medium faster than the local propagation speed of light, a phenomenon known as the Cherenkov Effect [32]. The detection units, called Digital Optical Modules (DOMs, see Section 4.1.3), collect Cherenkov light from charged particles produced by neutrino interactions in the ice. The telescope, though primarily designed to look for astrophysical neutrino sources, is a multi-purpose facility. The surface array, IceTop, serves as a veto and calibration detector. However, it can also measure arrival direction, flux, and composition of cosmic rays with the aim of uncovering their origin. In addition, a more densely instrumented sub-array in the clearest region of the ice, called DeepCore, allows IceCube to study the fundamental properties of neutrinos by exploiting the large atmospheric neutrino flux. The increased resolution of DeepCore along with higher quantum efficiency DOMs effectively lowers the energy threshold of the detector from about 100 GeV to 10 GeV. By incorporating the surrounding IceCube strings as an active veto region, DeepCore enables observations of neutrino oscillations as well as searches for sterile neutrinos and indirect detection of dark matter. An overview of the detector layout is shown in Figure 4.1.

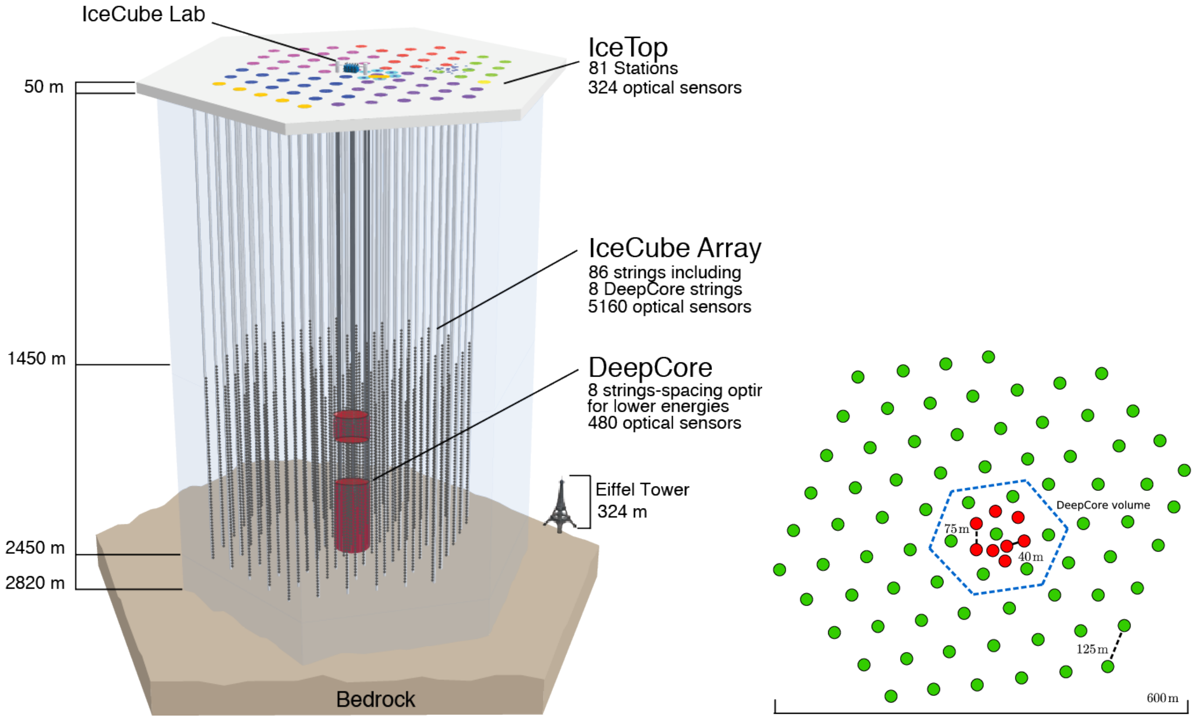

(a) (b)

Figure 4.1: (a) Layout of the IceCube/DeepCore Detector. Each black dot corresponds to a DOM and the red area indicates the region of the DeepCore instrumentation. (b) Overview of the string geometry. DeepCore strings are represented by red dots and the original IceCube array by green dots.

4.1.1 Medium and Location

The tiny neutrino-nucleon cross section of about $10^{-38}\frac{\mathrm{cm}^2}{\mathrm{GeV}}$ imposes certain requirements on the detection medium used for neutrino telescopes. At the GeV to PeV scale and upwards neutrino detection requires a very large detector volume ($\mathcal{O}(\mathrm{km}^3)$) for two reasons: (i) A larger volume increases the probability of seeing an interaction. As the atmospheric neutrino flux decreases with energy, the volume necessary to obtain sufficient statistics increases, correspondingly. (ii) At the high end of the IceCube energy range, the light emission can extend hundreds of metres. A large detection volume allows for more light to be collected resulting in a better estimate of the initial neutrino energy. The demand for a kilometre scale detector volume encourages the use of naturally and abundantly occurring materials with good optical properties. The 3 km deep Antarctic glacier at the geographical South Pole meets all the requirements for both atmospheric and astrophysical neutrino Cherenkov detection. The 1500 m of ice covering the detector provides effective shielding from atmospheric muons and the cold environment suppresses thermionic noise in the electronics. The deep ice is formed through 100,000 years of compression of lower layers of snow by the layers above. The slow process creates an environment of very low radioactivity. At the depths of the instrumentation, 1500 m below the surface, the transparency of the ice far exceeds that achievable for laboratory-grown ice [33]. In the clearest region, the average scattering length is about 50 m for light with a wavelength of 400 nm which is close to the peak of the DOM sensitivity. At these depths, the radioactive contamination ranges from $0.1 - 1 \cdot 10^{-12}$ g of Uranium or Thorium and $0.1 - 1 \cdot 10^{-9}$ g Potassium per 1 g of ice [23].

Though the transparency generally increases with depth, two striking features have to be accounted for when modelling the ice. At a depth of 2000 m to 2100 m, a band of high dust concentration is found, known as the *dust layer*. There, the scattering length drops to about 5m. Below this dust layer, the clearest ice is found, allowing for the most efficient Cherenkov light collection. The other notable characteristic of the detection medium is the ice immediately surrounding the DOMs which has a much larger concentration of bubbles. The *hole ice* is an effect of the deployment method, where holes slightly wider than a DOM are drilled with hot water. After the instrumentation is in place the water refreezes, leaving the optical properties in the holes different from the bulk ice. The scattering length of the hole ice is modelled to about 0.5 m. The ice properties are calibrated using data taken with a "dust logger" device lowered down the holes during the deployment. Calibration data is also taken during operation using LED flashers attached to the DOMs. The ice model is fitted to the data assuming that the optical properties change with depth and are constant in the horizontal plane with the exception of the hole ice.

In addition to the excellent optical and radioactivity properties of the deep ice at the South Pole, the nearby Amundsen-Scott Station provides the necessary infrastructure for running and maintaining the detector. This includes electrical supply, satellite communication, data acquisition, and accommodation for scientist and technical support.

4.1.2 Geometry

The Digital Optical Modules are arranged on a total of 86 strings that each have 60 DOMs attached at depths from 1450 m to 2450 m. The 78 strings that constitute the original IceCube array, form a hexagonal pattern with an average horizontal string distance of 125 m and a vertical DOM spacing of 17 m. The DeepCore fiducial volume includes the seven central "standard IceCube strings" and eight additional "DeepCore strings", placed between the central string and the innermost hexagon with an average string distance of 75 m. On the DeepCore strings, 50 DOMs are placed in the very clear deep ice below the dust layer with a 7 m vertical spacing. The remaining 10 DOMs are placed immediately above the dust layer with a spacing of 10 m to provide additional vetoing against atmospheric muons. An overview of the central detector baseline including a proposed layout of the PINGU strings is shown in Figure 4.3.

4.1.3 Digital Optical Modules

The detection units in IceCube are called Digital Optical Modules (DOMs) and serve the purpose of converting Cherenkov light from the charged particles of a neutrino interaction into electronic signals. The principal component of a DOM is the 25 cm Hamamatsu photomultiplier tube (PMT) and a main board housing most of the electronics. Both are enclosed in a 33 cm glass sphere as shown in Figure 4.2. The PMT is optically coupled to the glass sphere by optical gel, which also provides protection and stability to the PMT. With 10 dynodes yielding a total amplification of 10^7, the DOMs are sensitive to single photons. The Standard IceCube PMTs have a maximum quantum efficiency of about 25% at a wavelength of 390 nm which is close to the peak of the Cherenkov light spectrum. The PMTs used in DeepCore DOMs are similar to the original IceCube DOMs, but have a different cathode material with a roughly 40% higher quantum efficiency. All PMTs are oriented downwards in the direction where IceCube is most sensitive to neutrino oscillations. A small voltage is supplied through the string cable and transformed by a high voltage divider to provide the high voltage needed for the PMT. Each DOM also contains a flasher board with 12 Light Emitting Diodes (LEDs) used for studying ice properties, as mentioned in the previous section.

When a photon is absorbed by the PMT photocathode an electron gets emitted which is then amplified by the series of dynodes. The resulting signal is readout as a change in the voltage of the PMTs anode. When the voltage on the anode exceeds a threshold of 25% of the peak of a single photoelectron, the DOM is said to be *hit*. A hit causes the measured time series of charges (the *waveform*) to be read out by the main board where the waveform is digitised by three different circuits running in parallel; two analogue transient waveform digitisers (ATWDs) and one fast analogue-to-digital converter (fADC). The waveforms from the ATWDs have an integrated timing precision of 3.3 ns and a maximum duration of 426 ns. The high time and charge resolution is necessary for event reconstruction algorithms to determine the possible origin of single unscattered photons to a sub-metre precision. The second ATWD is only needed in case the DOM is hit again while the first ATWD is still digitising. The fADC is available for long signals and is sampled every 25 ns for a maximum duration of 6400 ns.

The high quantum efficiency of the PMTs also entail a high rate of "random noise hits"

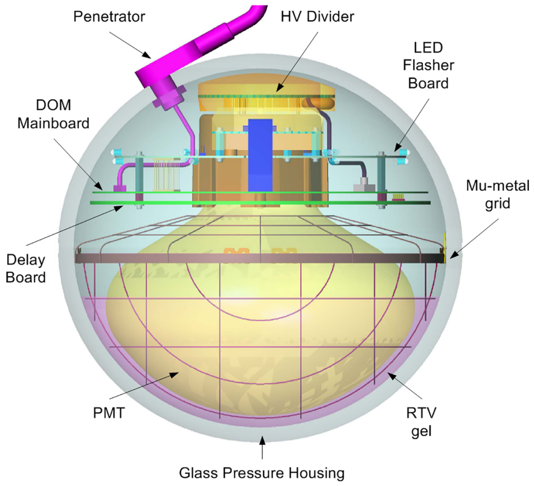

(a)

(b)

Figure 4.2: (a) Schematic illustration of the components in a Digital Optical Module. (b) Photograph of a DOM.

caused by radioactive decays in the glass sphere or thermal electrons from the cathode. Therefore, it is neither feasible or beneficial to read out every hit DOM. To reduce the contribution of noise hits in data, the DOM has a circuitry to check for hits in local coincidence with itself. A discriminator is fired whenever a neighbouring or next-to-neighbouring DOM on the same string is also hit within a time-window of $\pm 1\mu s$. For historical reasons, hits fulfilling this requirement are referred to as "hard local coincidence" (HLC) hits and the rest as "soft local coincidence" (SLC). If the local coincidence requirement is met, the DOM sends signal to the Data Acquisition System (DAQ), where the triggers are formed. Whenever one of the trigger requirements (see Section 5.1.1) are met, both the high and low resolution waveforms of the DOMs in HLC mode are read out for processing. The SLC waveform information is reduced to three fADC bins around the peak value. The low charge and timing resolution of the SLC hits is sufficient to discriminate dim atmospheric muons that slip through the upper part of the detector without causing HLC hits.

4.2 The Future low-energy extension: PINGU

An integral part of the proposed detector upgrade, IceCube-Gen2, is the Precision IceCube Next Generation Upgrade (PINGU). With an even denser infill of optical modules in the DeepCore region, PINGU aims to lower the energy threshold from the current 10 GeV to about 1 GeV. Though the main objectives of PINGU is the determination of the neutrino mass hierarchy and precision measurements of oscillation parameters, the increased efficiency at low energy will also extend the sensitivity range for indirect dark matter searches, and increase sensitivity to supernova neutrino detection. With the sensitivity to matter effects influencing the oscillation probabilities of neutrinos traversing the Earth, PINGU will also enable studies of the composition of the Earth core through neutrino based tomography.

The exact layout of PINGU is not yet settled upon. The Monte Carlo Data used in the work presented in this thesis is simulated with the former baseline geometry, called V36. It consists of the current DeepCore strings and 40 additional strings with an average spacing of 22 m. Each string supports 96 DOMs distributed at a depth of 2150 m to 2450 m with a 3 m spacing. The design of PINGU is currently being optimised to lower the construction expenses without compromising the science reach. The new configurations that are being tested can be deployed in only two seasons. They consist of 20-26 strings with a total of 192 DOMs with a vertical spacing of 1.5 m. The results presented here for the achievable muon background rejection is expected to be reproducible with either of the proposed geometries given suitable reoptimizations of the cut values.

The Standard PINGU DOM (PDOM) is a slightly modified version of the DeepCore DOM. The biggest change is a simplification of the waveform digitisation by using state-of-the art electronic components which were not yet available when DeepCore was constructed. The three components running in parallel in the IceCube and DeepCore DOMs will be replaced by a single analogue-to-digital converter that can be read out continuously.

Another planned improvement of PINGU is a reduction of air bubbles in the hole ice. To avoid the significant decrease in scattering length around the DOMs (see Section 4.1.1) the melted water will be degassed as part of the drilling process.

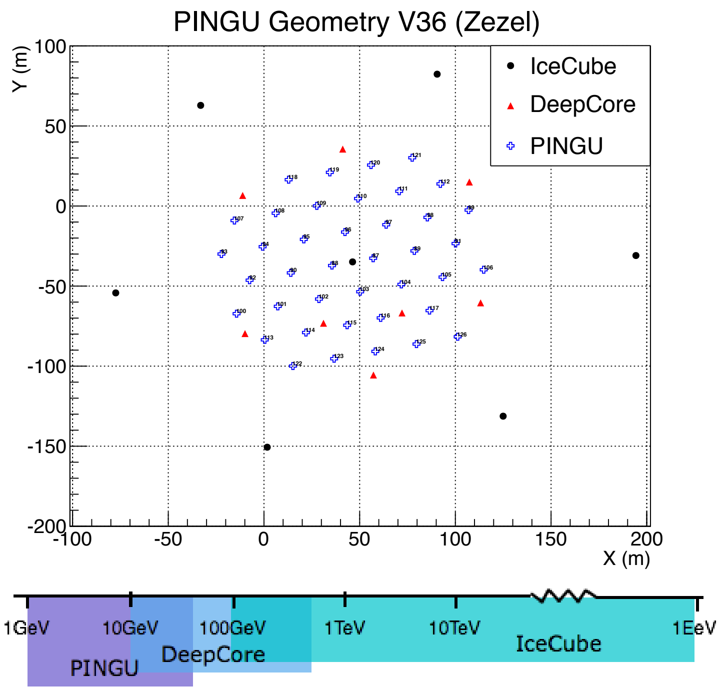

Figure 4.3: An overview of the central strings of the IceCube detector including the former baseline geometry for PINGU. Below, a scetch of the energy range of the individual subdetectors.

4.3 Data Simulation

The investigations presented in this thesis are based on simulated data. The Monte Carlo Software tools developed for IceCube can also be used to simulate the PINGU detector as the detection principal is identical. The detector response is modelled like the DeepCore

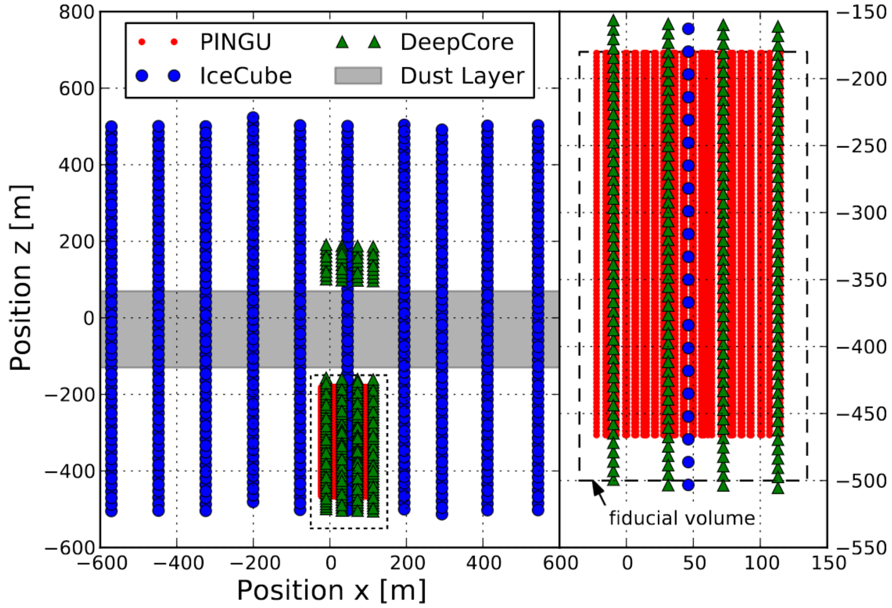

Figure 4.4: Side view of the IceCube Detector, including DeepCore and the proposed PINGU (geometry V36). On the right, a zoomed in figure of the central detector region. The black dashed line indicates the fiducial volume of PINGU.

DOMs without modifications. As the PINGU DOMs mainly differ from the DeepCore DOMs in the digitisation electronics, the data extraction should be functionally the same.

4.3.1 Atmospheric Neutrino Simulation

The first step of the simulation chain is the event generation which is carried out with the GENIE (Generates Events for Neutrino Interaction Experiments)[34] software package. GENIE models the neutrino interactions starting off with an isotropic flux on the surface of a cylindrical generation volume. The energy spectrum follows a user-defined power-law distribution - typically E^{-1}. All neutrinos that pass through the detector volume are forced to interact inside the detector. The type of interaction (see section 2.2) is chosen semi-randomly and the events are weighted according to their probability. The interaction is modelled in detail and all the created particles are stored for further processing. The particle propagation is carried out by GEANT4 [35] and the Cherenkov photons are tracked using CLSim, a GPU-based software which treats depth-dependent photon scattering in IceCube.

4.3.2 Atmospheric Muon Simulation

MuonGun is a toolkit for efficiently simulating atmospheric muons deep in the ice [36]. Instead of simulating full cosmic ray air showers and propagating the muons through the firn and ice, MuonGun constructs a parameterization of the muon flux under the ice and draws muon bundles from the parameterised distribution. The method is less precise than the full shower simulation, but much more efficient, as it allows the user to only generate muons which interact in the detector and have the desired energies. MuonGun generates muons starting on a fixed cylindrical surface. Samples are drawn from a parametrization of the atmospheric muon flux as a function of vertical depth, zenith angle, multiplicity, energy, and for bundles, the distance of each muon from the shower axis. The events used in the for the work presented in this thesis are weighted according to the GaisserH4a flux model [37]. The Muons are propagated through the detector using the IceCube software, Muon Monte Carlo (MMC). Like in the neutrino simulations, Cherenkov photon propagating is done with CLSim.

4.4 Event Topologies

The typical separation between DOMs in PINGU limits the ability to determine the flavor of individual neutrino interactions. The four distinct combinations of neutrino flavor and boson exchange result in either a track-like or cascade-like energy deposition and often a combination of the two. Cascades arise from hadronic and electromagnetic showers which typically overlap and cannot be separately distinguished. They appear as point-like sources with a typical size of the order of one metre or less. Cascade are approximately spherical, but have some directionality due to momentum conservation. Tracks are only created by relativistic muons traversing the detector while emitting Cherenkov radiation. The signatures of the different neutrino flavor interactions are summarised below:

NC all flavor: In neutral current interactions, only the hadronic cascade is detected. The size of the hadronic cascade depends on the momentum transfer and neutrino scattering process. The energy carried away by the outgoing neutrino is not seen in the detector, so the energy reconstructions are biased towards lower energies.

CC ν_μ: If the energy taken by the outgoing muon is sufficient for it to traverse several layers of detection modules, it can create a track pointing away from the interaction vertex. Muons at typical PINGU energies of a few GeV have a range on the order of ten or twenty metres which is comparable to the string spacing. Therefore, muon tracks in PINGU only elongate the spherical cascade topology slightly.

CC ν_e: The outgoing electron induces an electromagnetic shower overlapping the hadronic shower. The energy of the electron is deposited quickly through radiation of high energy photons (Bremsstrahlung) and following $e^\pm$ pair production which, in turn, emit more Bremsstrahlung.

CC ν_τ: With a lifetime of only 3×10^{-13} s, a few GeV tau travels less than a millimetre before it decays. The tau decay has a branching ratio of $\sim 17\%$ into a muon which leaves a track pointing away from the initial hadronic cascade, just like the CC ν_μ events. In such decays, two outgoing neutrinos carry away some of the initial state energy. The remaining $\sim$83% of the τ decays produce either an electromagnetic or hadronic cascade which cannot be independently resolved from the initial hadronic cascade. Only in the rare case of a tau being created with an energy of tens of PeV can the event be uniquely identified as having two cascades connected by a track.

Different analyses may prefer different detection channels depending on the relative importance of the accuracy on energy and direction reconstruction of the incoming neutrino. Directional reconstructions benefits from the longer ‘lever arm’ of a track-like topology and, therefore, the best angular resolution is usually achieved in the CC ν_μ detection channel. High angular resolution is typically a priority for point source and indirect dark matter searches, which look for a small excess of neutrino events in a certain direction. The energy reconstruction quality depends mainly on the fraction of initial neutrino energy that is deposited inside the detection volume, i.e. whether the event is ‘fully contained’. For energy reconstructions an outgoing muon can be a disadvantage as it might proceed out of the detector, leaving a signal that is only partially contained.

Chapter 5

Event Reconstruction

The two key observables to neutrino oscillation measurements are the energy carried by the initial neutrino and the path length it has travelled. The neutrino energy can be estimated from the amount of light created in the event and the topology. In the case of atmospheric neutrinos traversing the earth, the path length can be deduced from the direction of the light deposition. Precise reconstructions of energy and direction are essential to oscillation studies as well as extra-terrestrial neutrino source searches, dark matter induced neutrino searches, and Earth tomography.

In order to obtain a precise reconstruction of the energy and direction of an incident neutrino, a complete account of the event is necessary. For a CC ν_μ event that amounts to a total of eight variables to reconstruct: the position and time of the interaction vertex, the energy of the hadronic cascade, and the length/energy of the outgoing muon. This chapter briefly covers the main steps of the translation from the set of waveforms sent from the DOMs to the full description of the neutrino interaction. Focus is on the software applied to simulated PINGU data which also closely aligns with most DeepCore methods.

5.1 Event Building

The first step of the reconstruction chain is the classification of a set of DOM launches as an “event”. The definition of an (potentially interesting) event is represented by a set of trigger conditions, which determine when the DOMs are read out. For further processing, the raw data, consisting of both HLC and SLC waveforms, is calibrated and individual photon hits and times extracted.

5.1.1 Trigger Conditions

The high quantum efficiency DOMs used in DeepCore have, individually, a noise rate of roughly 600 Hz [38] which is mostly caused by thermal electronic noise or radioactive decay in the glass housing. The high noise rate makes it meaningless to store all hits in local coincidence. Instead, certain trigger requirements have to be met in order to build an event from a set of fADC and ATWD waveforms. The DAQ accommodates various triggers intended for the various purposes of IceCube that are based on the multiplicity of HLC hits in spatial or temporal coincidence [39]. For the PINGU simulations used in this study, a Simple Multiplicity Trigger 3 (SMT3) has been used. The PINGU SMT3 requires at least three HLC hits on PINGU DOMs within a time window of 2.5 µs. When a trigger condition is met, the full detector is read out within a 10 µs time window around the time of the triggering hit. The data from overlapping time windows is combined to an event and stored in an IceCube specific formatted file which is passed on to the Processing and Filtering framework.

5.1.2 Feature Extraction

In the Processing and Filtering (PnF) framework the Feature Extraction tries to break down the waveforms into single photon hits and corresponding time stamps. As the finite time and charge resolution does not always allow individual photons to be identified, the concept of a *pulse* is introduced. A pulse holds a time stamp t, duration Δt, and the integrated charge deposited on a DOM in the time interval $[t; t + \Delta t]$. The charge is normalised to the single photon waveform and measured in units of photoelectrons (PEs). The term 'pulse' is often used interchangeably with 'hit', but has a less direct physical interpretation. The outcome is a *pulse series* (or *hit series*), which contains all the pulses of an event and the DOMs on which they occurred. The pulse series is the object that is used in the filtering and reconstruction algorithms.

5.2 Noise Cleaning

While the trigger only counts the HLC hits, all the non-local coincidence (SLC) hits in the relevant time window are also read out and included in the pulse series produced by the waveform deconvolution. Most of these hits are caused by thermal noise in the electronics or radioactive decay in the DOM glass housing, and muddle the appearance of the physical event. The HLC hit sample represents a very clean pulse series. However, removing all SLC hits from the event engenders a non-negligible loss of 'physics hits' which can be important for the energy reconstruction and for particle identification. Efficient noise cleaning that keeps as many physics hits and as few noise hits as possible is crucial for the event reconstruction, as well as for the muon background rejection. Several hit cleaning algorithms exist, but for most PINGU purposes the two following are used consecutively:

SeededRT cleaning: The algorithm starts out with a subset of seeds that are mostly physics related hits. All hits within a time window (T) and radius (R) from one of the seeds are then added iteratively until no more hits are added. The method is illustrated in Figure 5.1. Only SLC hits close to the initial seeds are kept in the SeededRT (SRT) cleaned pulse series. For the simulated PINGU data, a cluster of HLC hits around the DOM that caused the trigger are used as seeds. The RT-range is set to $R = 40\,\mathrm{m}$ and $T = 250\,\mathrm{ns}$ for PINGU DOMs and $R = 150\,\mathrm{m}$ and $T = 1000\,\mathrm{ns}$ for IceCube DOMs.

Time-Window cleaning: Depending on the purpose of the hit cleaning, either a static or a dynamic time window is used. In the latter case, the pulse series is scanned to find the time window containing the most amount of charge. All hits outside that time

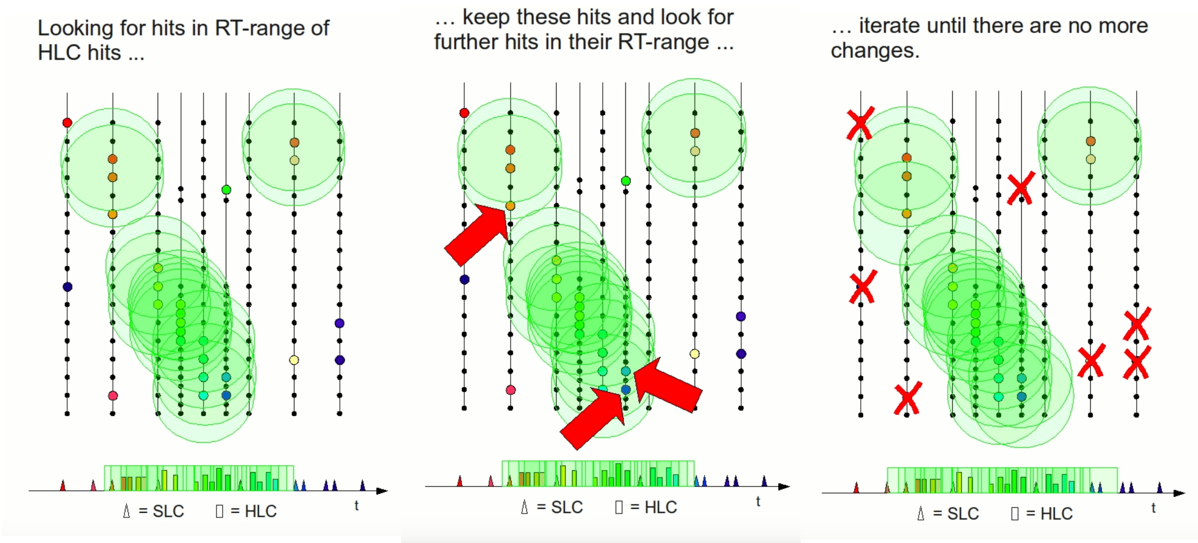

Figure 5.1: Illustration of the Seeded RT-cleaning, taken from [40].

window are removed. With Static Time-Window (TW) cleaning, a time window around one or more relevant triggers is defined and all hits outside it are removed. With PINGU simulations a time window of 6 µs is used for dynamic TW cleaning and a time window of 3.5 µs before and 4 µs after the PINGU SMT3 trigger is applied for static TW cleaning.

5.3 Reconstruction Algorithms

Once the event has been built and noise hits have been mostly removed, several reconstruction algorithms can be applied which either assume the event to have a track-like topology, a cascade-like, or a combination of both. At the energy scale relevant in PINGU, muon neutrinos typically deposit a large fraction of the transferred energy in the hadronic cascade. With PINGU events, the best performance is obtained from a reconstruction algorithm allowing for a combination of a track and a cascade. Here I present the methods that are referred to later in the text.

5.3.1 CLast

Likelihood based reconstructions typically require a good "seed", i.e. a first guess of the event properties. For cascade-like events, the algorithm *CLast* is used to provide an estimate of the vertex position, energy, and direction. The vertex position is determined as the charged weighted average of the positions of all hit DOMs in the TW-SRT cleaned pulse series. The time is set from a number of test times to be the earliest one resulting in four direct hits. A direct hit is defined as a hit that could be caused by an unscattered photon from the interaction position and time in regard. The cascade direction is established from the tensor of inertia of the hit positions, where the eigenvector of the smallest eigenvalue points along the elongation of the pattern.

5.3.2 Monopod

Monopod is a likelihood based reconstruction algorithm, that assumes a cascade-like energy deposition. Using the interaction time and position from CLast as a seed, Monopod receives the expected number of photons per unit energy at each DOM from immense six-dimensional tables, called Photonics tables. To avoid doing the expensive simulation of photon propagation for each event guess, these are done in advance for all possible depths and directions, and the results stored as Photonics tables. The Poisson likelihood can then be calculated from the expected number of photons and maximised to find the most probable values of the event parameters. At PINGU energies, the track of a possible outgoing muon will typically have a length of ten to twenty metres which is comparable to the string distance. Therefore an outgoing muon from a charged current ν_μ interaction only deform the topology of the hadronic cascade a little, and Monopod gives a sufficiently reliable estimate of the vertex position for coarse data filtering purposes.

5.3.3 HybridReco/MultiNest

The most sophisticated reconstruction algorithm used for PINGU analyses is the HybridReco/MultiNest, usually referred to just as MultiNest. The reconstruction process consists of two parts, HybridReco which is responsible for the event hypothesis and MultiNest which is the minimizer. In contrast to CLast and Monopod, HybridReco includes the possibility of an outgoing muon in the event hypothesis. A muon track introduces three additional variables to the problem; the direction defined by two angles and length of the track. Fortunately, the direction of the muon track only differs by a few degrees from the direction of the hadronic cascade at typical PINGU energies which is far below the angular resolution of the detector. Therefore, the cascade and track are assumed to be aligned. This reduces the additional parameters to the outgoing muon energy E_μ, which is proportional to the length of the track, since the muon is minimum ionising. Yet the complexity of the likelihood space is dramatically increased as there are now two different sources from which any hit could originate, namely, the hadronic cascade and the muon. The multimodal nested sampling algorithm, MultiNest [41], provides a robust method for finding the global maximum in such a likelihood space with many local maxima. In order to reduce the computation costs of event reconstructions, limits are set on the coordinates of the vertex position. Forming a square of 100 m by 100 m the MultiNest limits enclose the full PINGU volume which has a radius of roughly 80 m. Because events at the final analysis level are generally required to interact within the instrumented PINGU area, the limits on MultiNest should not affect the reconstructions of the final event sample. The boundaries of the MultiNest parameter space are shown along with the overview of the PINGU detector geometry and the requirement on the interaction vertex position in Figure 5.2.

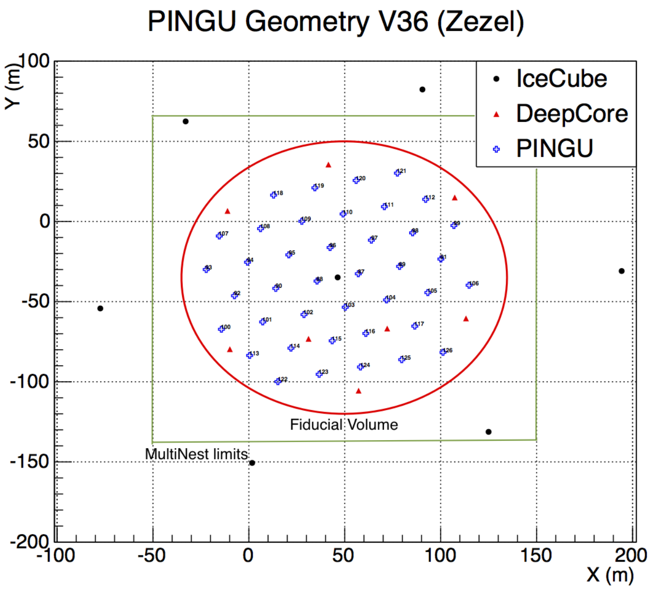

Figure 5.2: The PINGU V36 geometry with the approximate horizontal boundary of the fiducial volume (see section 6.1.3) and MultiNest limits.

Chapter 6

Generic Data Selection with PINGU

In this chapter I present a generic data selection scheme for PINGU, called the NBI events selection. The method accommodates the practical issues of data processing at the South Pole while reducing the atmospheric muon contamination by five orders of magnitude at the final analysis level. The results demonstrate that the atmospheric muon flux will not be a substantial background at the analysis level.

The first section introduces the basic principles and constraints of atmospheric muon rejection in low-energy neutrino studies with IceCube. A short review of the event selection formerly applied to all PINGU analyses, the PINGU Standard Processing (PSP), is given in the second section. The aim is to motivate the development of a new scheme and provide a baseline to which the performance of the NBI event selection can be compared. The third section covers the variables and performance of each individual level of the NBI events selection. Throughout the section, intermediate results are compared to the corresponding level of PINGU Standard Processing. Finally, a few remarks are given on some supplementary work proceeding in other groups which is related to the prospective implementation of the NBI event selection.

6.1 General Remarks on Atmospheric Muon Rejection

In the full IceCube detector atmospheric muons outnumber atmospheric neutrinos roughly by a factor of 10^6 at the trigger level. The DeepCore SMT3 trigger reduces this number by about an order of magnitude. With the current geometry and trigger settings, a comparable ratio is expected in PINGU. The goal of an atmospheric muon filter must therefore be a background rejection factor of at least 10^5 to reduce the atmospheric muon rate below the rate of atmospheric neutrino events.

6.1.1 Practical Limitations

The location of IceCube poses some constraints on the event classification methods used for data processing. All daily data transmission to the northern hemisphere happens via satellite which has a limited bandwidth. Therefore, the first part of the data processing must run at the South Pole in order to reduce the amount of data prior to transmission. At the South Pole, computing and electricity resources are limited which means that the filtering algorithms must be as computationally light weight and as efficient as possible.

Events that pass Step1 will be transmitted from the South Pole via satellite to the IceCube central repository in the United States, allowing for more sophisticated reconstructions. However, the offline processing has often had to be redone several times in IceCube for full years of data. Therefore, it is preferred to reduce the amount of data as much as possible using lightweight algorithms before running computing intensive reconstructions. The constraints on computing resources available for data processing, online as well as offline, must be accommodated in the event selection used for analysing the prospective sensitivities of PINGU.

6.1.2 Muon Vetoing

Atmospheric muons at the GeV scale deposit energy continuously along their trajectory all the way through the ice, thereby creating long, down-going tracks starting in the top of detector. The signature of an atmospheric muon differs from that of low energy neutrino-induced events that are generally more cascade-like, as described in Section 4.4. Unlike muons, neutrinos can have any direction and start anywhere in the detector. However, at the energy scale of $\mathcal{O}(10\,\mathrm{GeV})$ and below, neutrinos are far more likely to cause a PINGU trigger if the interaction happens within or close to the PINGU region. The basic principle of the early muon rejection is, therefore, to define an outer 'veto volume' and an inner 'fiducial volume', and reject all events seemingly starting outside the fiducial volume. This approach has proven to be powerful with DeepCore, where the online muon filter employed in the data stream, obtains a muon rejection of about 90% relative to the trigger level at a muon neutrino efficiency above 99% by using the entire exterior detector volume as an active veto region [38].

Figure 6.1 shows the raw data of an atmospheric muon and a ν_μ event before any cleaning algorithms have been applied to the pulse series. The example illustrates that simply rejecting events based on the number of hit DOMs in the veto region will not provide the needed background rejection at an acceptable ν signal efficiency. Typically, several supplementary variables are constructed to reject atmospheric muon background, based on the spatial and temporal distribution of hits in the veto and fiducial region. It was mentioned in Section 5.2, muon veto techniques generally rely on good noise cleaning. Sometimes a single isolated hit on an outer DOM can reveal the track of an incoming muon. If it is mistaken for noise and removed, the muon will appear as starting within the detector. Similarly, a kept noise hit in a neutrino event can be misidentified as the trace of an incoming muon and cause the neutrino to be vetoed.

6.1.3 Event Quality

The physics portfolio of PINGU relies not only on efficient muon rejection, but also on the detector's ability to accurately reconstruct the direction and energy of the initial neutrino. The reconstruction quality of events at the GeV scale depends on the 'containment' of the

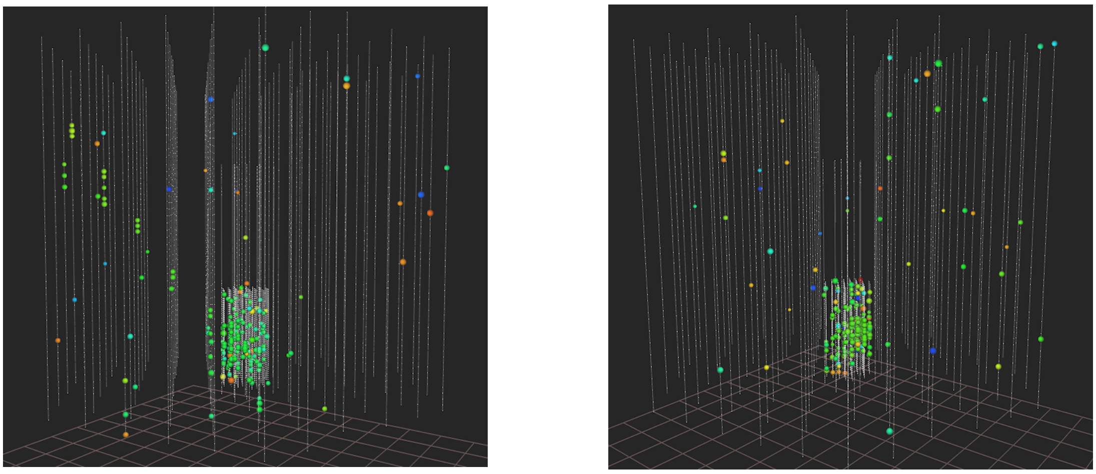

(a) Typical muon event in PINGU.

(b) Typical neutrino event in PINGU

Figure 6.1: Event displays of the two fundamental types of events that we want to distinguish between. In (a) the typical signature of an atmospheric muon leaving a clear track is shown. (b) shows the signature of a 10 GeV muon neutrino starting in the PINGU fiducial volume. Colors correspond to the event time scale where red indicates early hit times and blue represents late hit times.

event, i.e. the fraction of energy deposited inside the densely instrumented PINGU region, and the number of hit DOMs. Collectively, characteristics related to the reconstruction quality are referred to as the ‘quality’ of an event.

The necessary event quality depends highly on the analysis. Some analyses select only fully contained events, by requiring both the start and end point of the reconstructed event to be found inside the fiducial volume. Due to the absence of a clear outgoing muon track in most low energy PINGU events, the containment requirement can be reduced to the containment of the vertex.

Vetoing events that start outside the fiducial volume serves two purposes: it reduces the atmospheric muon contamination and it enhances the quality of the neutrino events in the final sample. The fiducial volume should optimize the signal-to-bakground ratio within a region determined by the reconstruction quality.

6.2 PINGU Standard Processing

For the simulated data samples used in all current analyses of PINGU, a two step event selection, called PINGU Standard Processing (PSP) is applied. The first step is aimed at removing the most obvious atmospheric muon background. The second step rejects all events with a reconstructed vertex outside of the PINGU fiducial volume with the two fold purpose of filtering out atmospheric muons and neutrinos that are likely to be badly reconstructed. The PSP methods were developed before enough muon simulation had been produced to test and optimise the cuts. Therefore, the variables and cut values were based mostly on DeepCore experience.

6.2.1 PSP Step1

The first step of the PINGU Standard Processing exploits the different spacial and temporal distribution of charge for atmospheric muon and low energy neutrino events. It is based on three topological variables proven to have good separation power with DeepCore, and two variables related to the position of the interaction vertex which is reconstructed with CLast and Monopod.

- **Time$_{90\%}$ < 2 µs**: Events caused by atmospheric muons typically last longer than those caused by low energy neutrinos. Atmospheric muons traversing the entire detector deposit energy steadily over the course of 3000 ns. The outgoing muon from a CC ν_μ interaction of $\mathcal{O}(10\,\text{GeV})$ has a much shorter track length, and cascades in general last only for a few ns. Therefore, the time taken to accumulate 90% of the charge in a cleaned pulse-series is required to be less than 2 µs for the event to pass.[1]

- **C2QR6 > 0.5** : Following the same logic regarding the time evolution of energy deposition, the fraction of all charge deposited in the first 600 ns can also be used as a discriminating variable. The two first hits are excluded from the calculation, as they can be noise hits, making neutrino events appear longer in time. Thus the variable is defined as:

$$\text{C2QR6} = \frac{\sum_{t_3}^{t_3+600\text{ns}} Q_i}{\sum Q_i} \tag{6.1}$$

where the subscript i sums over the pulses in the event between the time of the third pulse t_3 and 600 ns later, and Q_i is the charge of the ith pulse. The fraction of charge in the event deposited during the first 600 ns, must be larger than 0.5 to pass.

- **zTravel > −30 m** : For a down-going atmospheric muon track, early pulses are expected to be found higher in the detector than the late pulses. The 'down-goingness' of events is indicated by subtracting the average z-coordinate of the first quartile of hit DOMs from the z-coordinate of each hit DOM and taking the mean:

$$\text{zTravel} = \frac{1}{N}\sum_i (z_i - \widetilde{z_{1/4}}) \quad , \quad \widetilde{z_{1/4}} = \frac{1}{N/4}\sum_{i=0}^{N/4} z_i \tag{6.2}$$

where N is the number of pulses in the event and z_i is the z-coordinate of the ith pulse. For atmospheric muons the average z of the first quartile should be relatively high compared to the rest of the hits, thus the 'zTravel' value should be negative. Events with $zTravel > -30\,\text{m}$ pass the cut.

- **$\mathbf{z_{CLast}}$ < −200 m** : Atmospheric muon tracks typically appear to be starting high up in the detector, if not above it. The vertical position of the interaction vertex reconstructed with CLast is therefore required to be below $z = -200\,\text{m}$ in the IceCube coordinate system.

[1] A more detailed argument is given in Section 6.3.1.1, as this variable also occurs in the NBI event selection.

- $\mathbf{r_{Monopod} < 95\,m}$: A cut is also placed on the radial distance from the vertex reconstructed with Monopod to the centralmost IceCube string. Selecting events with a vertex closer than 95 m further reduces the atmospheric muon rate and ensures that only potentially high-quality events are passed on to the full reconstruction in Step2. To reduce processing time, Monopod only runs on events that pass the four cuts described above .

6.2.2 PSP Step2

As the most obvious atmospheric muons are rejected by the Step1 cuts, Step2 makes use of the full HybridReco/MultiNest event reconstruction. All the cuts concern the containment of the interaction vertex with the aim of rejecting the remaining background events and selecting high quality neutrino events.

- $\mathbf{-500m < z_{MultiNest} < -180m}$: The upper limit of the PINGU fiducial volume rejects down-going events starting in the upper half of the IceCube detector as well as events originating from far below the instrumented volume. The upper boundary is moved slightly upwards compared to Step1 ($z_{CLast} < -200$), as the straight down-going muon contamination is expected to be significantly reduced by Step 1. Occasionally, events reconstruct below the instrumentation. Those reconstructions are unreliable and discarded by the lower boundary of the PIGNU fiducial volume.

- $\mathbf{r_{MultiNest} < 85m}$: To ensure that the events used in the precision measurements of neutrino oscillation parameters have reliable zenith angle reconstructions, a strict radial containment cut is imposed on the vertex. The radial containment cut also targets the more horizontal atmospheric muons which enter from the side of the detector and might therefore have survived the Step1 cuts.

- $\mathbf{cos(\theta_{MultiNest})>0}$**:** The cut on the zenith angle is optional as the various analyses done with PINGU look for a signal in different directions. This is discussed in greater detail in Section 6.3.3.

6.2.3 PSP Performance

The rates of atmospheric muon and neutrino events achieved with PINGU Standard Processing are shown in Table 6.1 and also illustrated in Figure 6.2. Testing the PSP Step1 performance on atmospheric muon background files reveals several issues. Most remarkable is the computing cost of the first step, which will have to run at the South Pole on real data. The PSP Step1 takes on average 2.8 seconds per atmospheric muon event to run, including reading the file and writing it out. With a total trigger rate expected to be a few hundred Hz, the computing power needed exceeds the computer logistics limitations.

The two most time consuming components of the Step1 code are the reconstructions, CLast and Monopod, taking on average 0.015 seconds and 60 seconds per event, respectively. This means CLast can reconstruct muon events at a rate of 67 Hz and Monopod

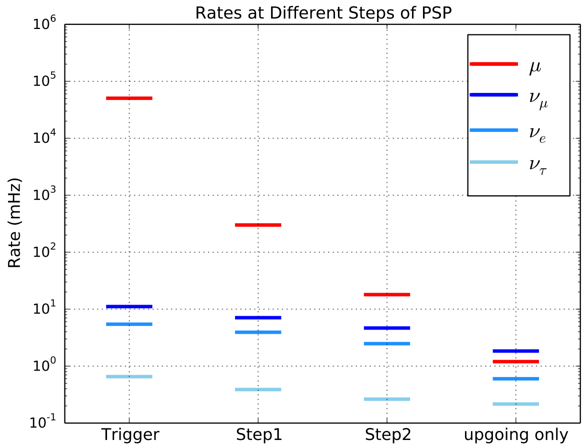

Figure 6.2: Rates at different levels of PINGU Standard Processing. The neutrino rates include both CC and NC interactions.

Rates (mHz)	MuonGun	Genie ν_μ	Genie ν_e	Genie ν_τ
Trigger level	$50 \cdot 10^3$	11.1	5.44	0.654
PSP Step1	299	6.88	3.36	0.388
PSP Step2	17.8	4.66	2.62	0.278
PSP Step2 Upgoing	1.13	1.90	1.33	0.223

Table 6.1: Atmospheric muon neutrino rates after each step of PINGU Standard Processing.

Module	MuonGun rate	Processing speed
CLast	50 Hz	67 Hz
Monopod	1.7 Hz	0.017 Hz
MultiNest	300 mHz	0.2 mHz

Table 6.2: For each of the three reconstruction algorithms used in PSP, the processing speed (in events per second) is compared to the atmospheric muon rate which is reconstructed by that algorithm.

0.017 Hz. Table 6.2 presents the time consumption of the three reconstruction methods used in PSP based on a trial data sample of $3 \cdot 10^5$ events.

Monopod only reconstructs the events that pass the four other cuts of Step 1, which is on average 3.4% of all the triggered events. This leaves a MuonGun rate of 1.7 Hz to be processed at a rate of 0.017 Hz. Due to the limited electricity and computing resources at the South Pole, we do not expect to have 100 computers available for the first step of PINGU data processing. The rates suggests that Monopod is not feasible to use in an online filter at the South Pole.

Furthermore, Table 6.2 shows that the muon rate after Step1 is almost 300 mHz and that MultiNest takes about 1 hour pr. event. Assuming a capacity of about 1000 CPUs available for the second step of PINGU processing, it would be possible to do the reconstruction in real time. However, that also means that if the processing needs to be redone, 1000 computers will be occupied for a year in order to reprocess one year of data.

In conclusion, the PSP scheme does not provide a realistic measure of the background reduction achievable with PINGU. Apart from the method being unrealistic, the resulting background reduction is also not sufficient. The current sensitivity estimate of PINGU to both the neutrino mass ordering and the earth tomography are based on the ansatz that the muon contamination will be negligible at analysis level. Figure 6.2 shows that the muon-to-neutrino ratio is around 1/3 for the up-going sample and 2/1 for all directions, clearly inconsistent with that assumption.

6.3 NBI Event Selection

The basic processing of PINGU data is meant to be a generic event selection applied to all data. Depending on the measurement of interest, whether it be ν_τ-appearance or the neutrino mass ordering, the signal is located at different regions in zenith and energy space. Therefore, this event selection is not designed to favor any signal region or neutrino flavor, but seeks to optimise the sensitivity to all measurements by maximising the quantity

$$S = \frac{N_\nu}{\sqrt{N_\nu + N_\mu + N_{noise}}} \tag{6.3}$$

where N_ν is the total neutrino rate, N_μ the muon rate, and N_{noise} the noise rate.

That being said, the quality of events is also considered in the evaluation of the individual variables and passing criteria. By regarding low quality neutrino events as dispensable, stricter cuts can be justified to achieve more background rejection. The NBI event selection consists of three steps in total:

Step 1) Removes the most obvious atmospheric muons in order to reduce the data rate for satellite transmission. The variables used for Step1 are fast enough to compute that they can run online at the South Pole.

Step 2) Targets atmospheric muons with dim tracks that are removed by the noise hit cleaning and reduces the muon rate sufficiently to run sophisticated likelihood based event reconstructions.

Step 3) Identically to PSP Step2, incoming events are rejected based on the reconstructed vertex position from MultiNest, but the radius of the PINGU fiducial volume is expanded slightly. This reduces the atmospheric muon background contamination further and ensures the reconstruction quality of events used in analyses.

The variables and cut values were fixed using a sample of simulated atmospheric muon and ν_μ events that fulfil the conditions of a PINGU SMT3 trigger. The GENIE ν_μ simulation corresponds to about 100 days of detector livetime with an energy range of 1-80 GeV. The muon events are generated with an energy range of 160-500 GeV and the event sample corresponds to about 2.5 days of detector livetime.

6.3.1 NBI Step1: Obvious muon rejection

The first step of the NBI Event Selection refrains from likelihood based event reconstructions by quantifying several differences between the topology of a typical atmospheric muon event and a centrally starting neutrino event. An ensemble of one dimensional (or “straight”) cuts is used as a classifier, in order to make the method robust against changes in the trigger conditions and detector geometry. Though certain machine learning techniques can offer a better separation of signal and background events, the agreement between data and the Monte Carlo simulations is not expected to be sufficient at the trigger level to use more involved classifiers. By applying straight cuts at the first step, the predicted rates should roughly apply given changes in the data acquisition or minor deficiencies in the Monte Carlo simulations.

6.3.1.1 Step1 Variables

Step1 incorporates ten variables in total. Six variables are designed to identify atmospheric muons based on the activity in the veto region. Three variables exploit the longer duration and track of a typical atmospheric muon event compared to low energy neutrino event. One variable is included to identify and remove events caused purely by noise in the DOMs.

The motivation and description of each variable is given below and the distributions of each variable for atmospheric muons and muon neutrinos are shown in Figure 6.6.

Pure Noise Events

The high quantum efficiency of the DOMs used in DeepCore result in a dark noise rate of roughly 600 Hz [38]. Assuming the same sensitivity and trigger settings for PINGU, the chance of an SMT3 trigger to be caused by coincident noise hits on neighbouring DOMs is non-negligible. As the signature of a low energy neutrino interaction appears similar to that of a cluster of noise hits, the removal of noise events is unavoidably associated with a loss of neutrino events. Therefore, the need to eliminate noise events from data must be accounted for in the predicted neutrino efficiency of PINGU.

- **NoiseEngine:** The NoiseEngine algorithm scans the cleaned pulse series with a sliding time window of 750 ns to find the period with the maximum number of hits. Hit pairs satisfying a constraint ([0.1; 1.0] m/ns) on the possible velocity of a particle causing both hits are identified and mapped on to a sphere of 192 directional bins in zenith and azimuth. A minimum of one bin with more than three hit pairs is required for the event to pass. The apparent velocity range is intentionally wide to keep as many physics hits as possible. The passing conditions of NoiseEngine are fulfilled by 88% of neutrinos and 92% of muons. The settings have not been optimised for PINGU. They align with the standard DeepCore settings, except for the number of directional bins which is four times higher due to the higher number of DOMs in the configuration.

Incoming muon tagging with trigger and HLC logic

The bulk of atmospheric muons are easy to identify by the large amount of light deposited outside the PINGU volume. In most cases incoming muons leave at least one HLC hit and often enough to cause an SMT8 trigger before entering PINGU. These muons are tagged by two veto algorithms, that combined reject 97% of atmospheric muons while keeping 96% of all flavor neutrino events after the PINGU SMT3.

- **Trigger Veto:** A simple and very efficient method to identify incoming muons is provided by the time ordering of the different triggers in IceCube. Most atmospheric muons will cause an SMT8 trigger as they enter the IceCube instrumented volume and a possible PINGU SMT3 trigger will have a later timestamp. By requiring that the first trigger of the event has to be a PINGU SMT3 the muon rate is reduced by 95% while 98% of the neutrinos are kept. The neutrino initiated events that fail this cut start outside the PINGU volume.

- **HLC Veto:** To identify the muons that do not cause an SMT8 trigger, an algorithm is run to search for HLC hits in the veto region that are consistent with a hypothesis of an incoming muon. This method has proven to have good rejection power with DeepCore, where it is used alone as the online filter. The algorithm estimates the "centre of gravity" (COG) and time of the event from the average position and time of all hits inside the PINGU volume. The COG is then improved by removing all hits outside one standard deviation from the average time and recalculating the COG. Events with more than one HLC hit in the veto region consistent with a particle travelling at the speed of light from the HLC DOM and to the COG are discarded. This reduces the atmospheric muon rate further by 39% while keeping 98% of the remaining neutrinos. A schematic illustration is shown in Figure 6.3 along with the distribution of the number of HLC veto hits for atmospheric muons and muon neutrinos, respectively.

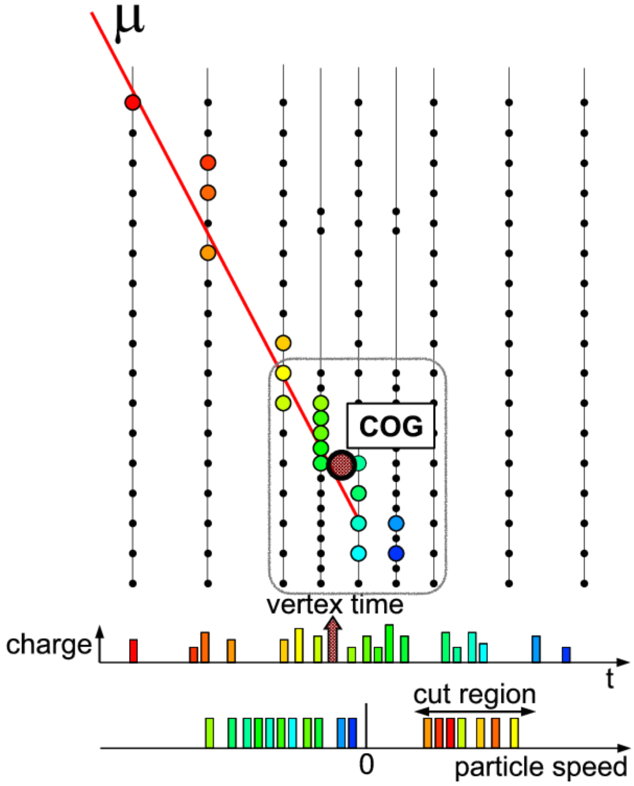

(a) Schematic illustration of the HLC Veto filter.

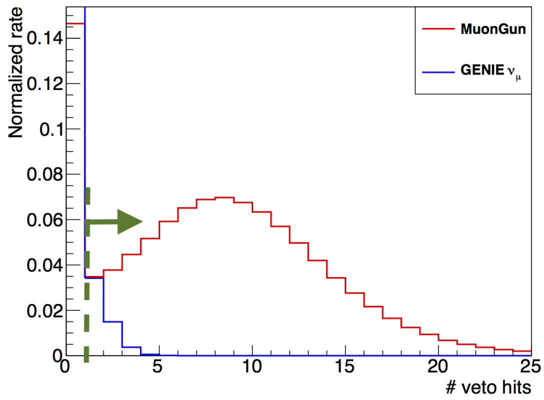

(b) Normalised distribution of the total number of HLC veto hits.

Figure 6.3: In (a) an example of a down-going muon that would be vetoed by the HLC veto. DOMs are represented by the black dots and the colored circles represent hits. The figure is taken from [38]. The distribution of HLC veto hits for ν_μ event and atmospheric muons is shown in (b). All events to the right of the green dashed line are removed.

The trigger veto and the HLC veto are responsible for the most dramatic reduction of the data rate. Though the majority of muons are removed by the veto cuts, the resulting muon rate is still ~100 times larger than the neutrino rate (see Table 6.4). Most of the muons that pass the trigger and HLC veto either slip through the 'empty corridors' in the IceCube grid (illustrated in Figure 6.4) or arrive vertically in between strings. To bring the background rate down further, seven variables are computed to catch the atmospheric muons that do not cause a veto trigger or an HLC veto hit. The rejection power of each variable is evaluated after the trigger veto, HLC veto, and NoiseEngine cut are applied. The distributions of all seven variables is shown in Figure 6.6 for atmospheric muons and muon neutrinos (CC+NC) that pass the veto cuts.

Rate of energy deposition

It can take about 3000 ns for a muon travelling at the speed of light to propagate through the full IceCube detector, while generating Cherenkov light constantly. In comparison, neutrino-induced hadronic and electromagnetic cascades yield more light per unit length, but last only a few nano seconds. Since neutrinos at the PINGU energy scale often deposit a large fraction of energy in the initial cascade, the temporal distribution of collected charge is different for neutrino interactions versus atmospheric muons. At the PINGU energy scale, the longest lasting events are CC ν_μ interactions of ~50 GeV. The outgoing

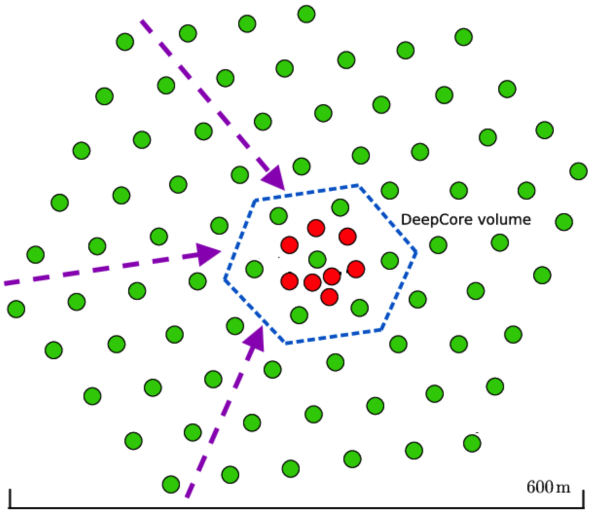

Figure 6.4: Examples of 'corridors' in the IceCube grid where incoming muons can sneak through without leaving veto hits.

muon with an energy $E < 50\,\mathrm{GeV}$ will be stopped well within 800 ns. [2] Three different variables exploit the cascade-like topology of the low energy neutrino interactions:

- **Time$_{75\%}$ < 900 ns:** The time taken to accumulate 75% of all charge in the event should be smaller for neutrino events than for atmospheric muons. The variable is calculated from the time-window-seededRT (TWSRT) cleaned pulse series, setting the time t=0 at the first hit. A range of values for the charge fraction were tested from 50% to 90% to get the best separation between ν_μ events and atmospheric muons. The discriminating power for each value of the charge fraction was evaluated at a 98% signal efficiency. The best separation was found by requiring 75% of all charge to be accumulated within 900 ns.

- **C2QR6 > 0.6:** Based on the same reasoning, the ratio of charge accumulated in the first 600 ns to all the charge in an event should be close to 1 for neutrino interactions and anywhere between 0.1 and 1 for a through-going muon depending on its direction (i.e. the length of the track). As in PSP Step1, the ratio is computed from the TWSRT pulses, ignoring the first two hits which might be noise. Events with a charge ratio below 0.6 are rejected.

- **iLineFit Speed > 0.25 m/ns:** This variable does not directly use the temporal distribution of charge collection, but rather the lack of long muon tracks in low energy neutrino events. Assuming a track-like topology, the event is fit to a straight line with the simple reconstruction, Improved Line Fit. The fit provides an estimate of the speed of the particle creating the track. The spherical nature of the hadronic cascade causes the speed reconstruction to be shifted downwards, whereas the long track of an atmospheric muon allows for a more accurate speed reconstruction close

[2]The mean inelasticity in a deep inelastic neutrino-nucleon scattering for a neutrino energy of 50 GeV is $\langle y \rangle = 0.5$ [42]. A 25 GeV muon deposits 0.2 GeV per metre, such that all energy is deposited within $\frac{25\,\mathrm{GeV}}{0.2\,\mathrm{GeV/m}\cdot 3\cdot 10^8\,\mathrm{m/s}} = 400\,\mathrm{ns}$.

to the speed of light. Events with a reconstructed speed lower than 0.25 m/ns pass the cut.

Hits in upper part of detector

An obvious consequence of the down-going-ness of atmospheric muons is that the event usually starts in the upper part of the detector. Neutrinos, on the other hand, can interact with equal probability in the whole detector, but are unlikely to trigger PINGU if they do not interact close to the fiducial volume at energies below 50 GeV. Thus, atmospheric muons also generally leave more hits above the PINGU volume than low energy neutrinos.

- **Charge above -180 < 12 PE:** Based on an uncleaned pulse series, all charge deposited above the PINGU instrumentation - corresponding to z=-180 m in IceCube coordinates - within a time window of [0, 2000] ns before the PINGU trigger is summed. The length of the time window is set to cover all possible hits from an atmospheric muon. Events with more than 12 PE are discarded.

- **z-coordinate of first hit < -180 m:** To minimize the probability of the first hit being a noise hit, the very clean TWSRT pulse series is used for this variable. Atmospheric muon events that pass the trigger and HLC veto cuts typically have only sparse hits in the veto region, which are often removed by the SeededRT hit cleaning. Therefore, the distribution of the z-coordinate of the first hit in a TWSRT pulse series peaks around the top of the PINGU instrumentation for atmospheric muons that pass the other veto cuts. For neutrino induced events, the first hit in the TWSRT pulses is more equally distributed across the vertical range of PINGU. Events are required to have their first hit in the TWSRT pulse series recorded below $z = -180\,\mathrm{m}$.

Hits in veto region

If the early hits from an atmospheric muon are too far apart and therefore removed by the hit cleaning algorithms, these will not contribute to the time duration of the event. Atmospheric muons that arrive nearly horizontally will not leave many hits in the upper detector layers, and might pass the cuts described above. To target the remaining atmospheric muons, two algorithms are incorporated in the event selection which look for SLC hits anywhere in the veto volume close enough to the PINGU trigger in time and space to reveal an incoming track.

- **Hits in RT cluster < 6:** The algorithm looks for a cluster of hits in the veto region by splitting the SeededRT cleaned pulse series into a PINGU pulse map and a veto pulse map. For each pulse in the veto pulse map the SeededRT cleaning is redone with that single hit as a seed. If a cluster in the veto region has more than 6 hit DOMs occurring before the PINGU trigger, the event is rejected.

- **Isolated veto charge < 16:** The uncleaned pulse map is used to find isolated hits on the veto DOMs that could be causally connected to the PINGU trigger. Such hits are identified from their spatial and temporal distance from the hit that caused the PINGU trigger. The rationale is depicted in the left plot of Figure 6.5, where hits before the trigger are assigned a positive time difference and hits after a negative time difference.

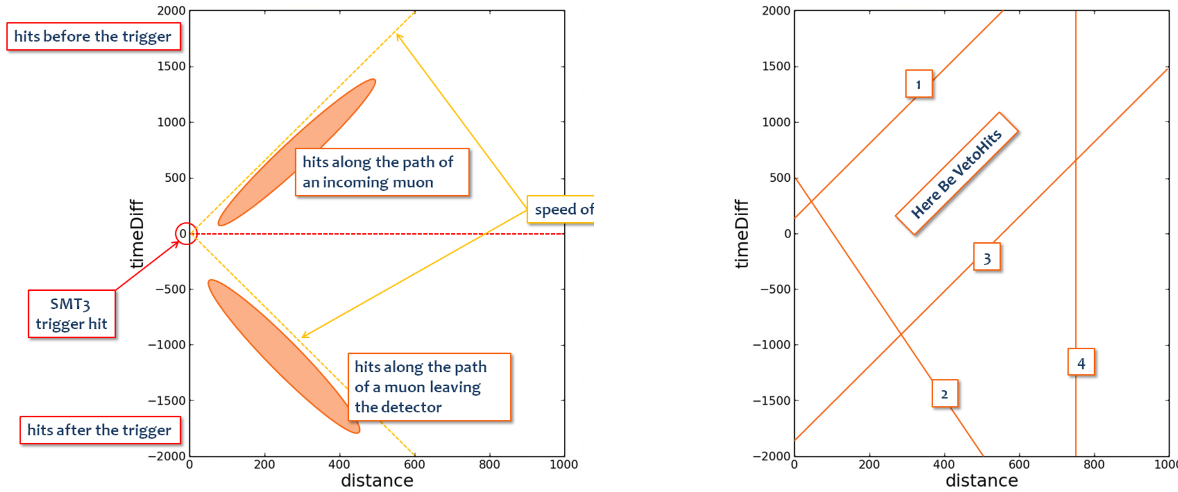

Figure 6.5: Schematic representation of the veto charge cut and motivation. The lines on the right plot represents the “veto region”, where hits from an incoming muon are expected.

The algorithm defines a “veto region” in the time-distance plane, shown in the right plot of Figure 6.5. Hits that are consistent with the hypothesis of an incoming particle, travelling at the speed of light, and causing the trigger are expected to be found in the veto region. The slope of line 1 and 3 corresponds to the speed of light. Line 3 is shifted 2 μs to account for late hits from scattered photons. Line 2 marks the limit of a possible hit from an outgoing muon. Hits to the right of line 4 occur more than 750 m from the trigger and are not taken into consideration. The cut requires the total charge from hits within this veto region to be less than 16 PE.

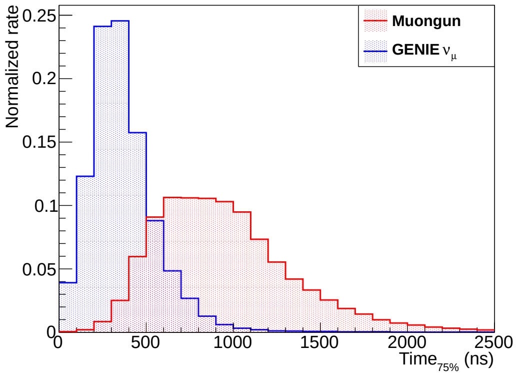

(a) Time to accumulate 75% charge

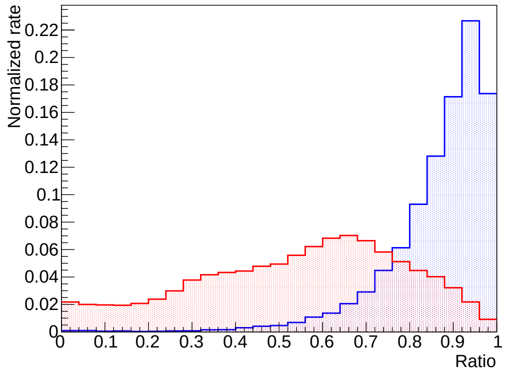

(b) C2QR6

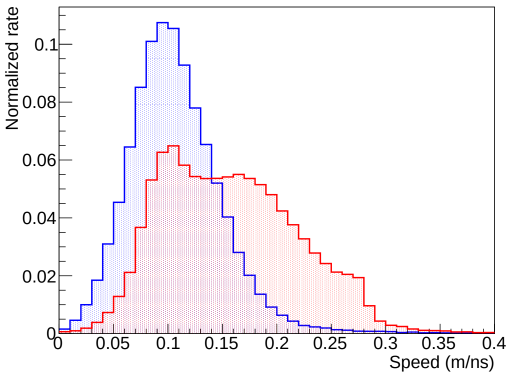

(c) iLineFit Speed

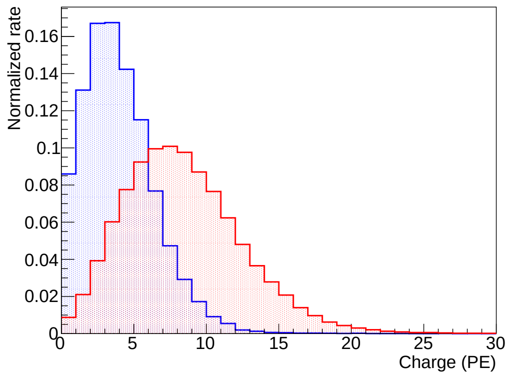

(d) Charge above -180m

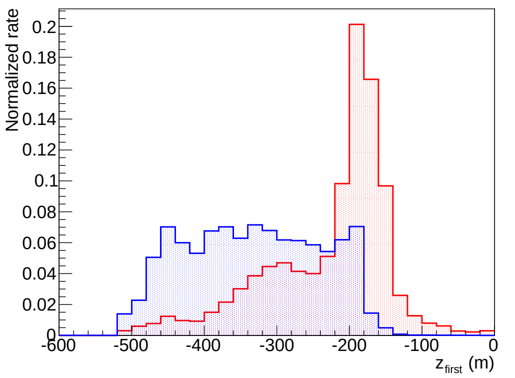

(e) z-coordinate of first hit

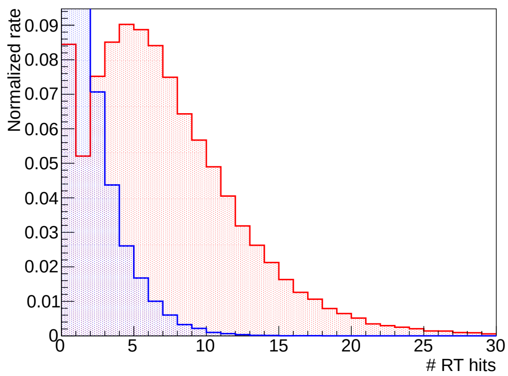

(f) Hits in RT Cluster

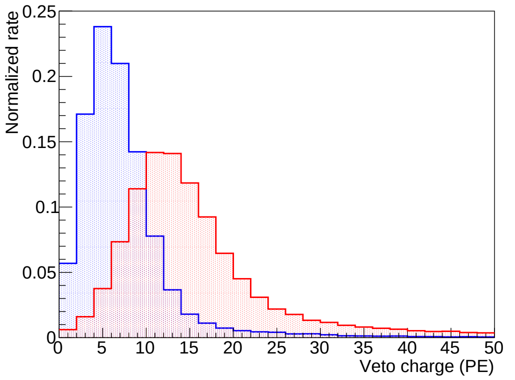

(g) Isolated veto hits

Figure 6.6: Distributions of the the NBI Step1 variables for atmospheric muon events (MuonGun), shown in red, and ν_μ events (GENIE) in blue.

6.3.1.2 Choice of Variables and Cut Optimization

The NBI Step1 variables were picked from a selection of test variables, that were mainly developed for low-energy DeepCore analyses. The choice of variables and cut values is based on the separation between ν_μ events and atmospheric muons. As muon neutrinos give rise to both track-like and cascade-like events, the event selection does not favor any particular event topology. The set of passing conditions were optimised "by hand", examining a few variables at a time, starting with the ones that had the most discriminating power.

For each variable, the muon rejection power was evaluated at a ν_μ efficiency of 98% after the application of the trigger veto, HLC veto, and NoiseEngine. Four variables were found to have strong discriminating power in the first iteration of testing: $\text{Time}_{75\%}$, RTCluster Hits, qAbove180, and iLineFit Speed. The cuts were set on each variable individually around values corresponding to a ν_μ efficiency of 95-99%, depending on the quality of the failing ν_μ events.

From the descriptions of the variables, it is clear that several of them are likely to be correlated. To avoid using redundant variables, the "mock-up sensitivity" (Equation 6.3) was calculated four times - each time applying three of the four cuts and alternating which variable to leave out. The sensitivities obtained with the different combinations of three cuts were compared to that obtained with all four cuts combined, thereby checking that each individual cut had a positive effect on the sensitivity. If one variable was redundant, the sensitivity would be the same whether that cut was omitted or all four cuts were applied.

After applying the first four cuts, three variables still looked informative. Criteria on those were inferred in the same manner by tweaking each cut value, individually, around a 95-99% ν_μ efficiency. To check if the last three variables were redundant, the sensitivity was computed after the full Step1 cut ensemble, and again omitting one of the three variables at a time.

The numerical method for rectangular cut optimisation, provided by the ROOT Toolkit for Multivariate analysis (TMVA), can be competitive with more involved classifiers when given a few variables with good signal-to-background separation, but it typically underperforms with many weak and correlated variables. Attempts were made to optimize a few of the variables at a time with the TMVA method, but eventually the best sensitivity was obtained by manually tuning the cut values. Further efforts were not put into numerical cut optimisation, because additional fine tuning will not be relevant given a change of geometry, trigger settings, pulse cleaning settings, etc.

6.3.1.3 Step1 Performance

Applying the Step1 cut ensemble reduces the total data rate from 50 Hz to about 250 mHz which is more than sufficient to transmit via satellite. The rates following each step of the NBI event selection are shown in Table 6.3 along with the rates obtained with PSP Step1. The results show that a 25% lower muon rate and a 20% higher neutrino rate are obtained with NBI Step1 relative to PSP Step1.

Rates (mHz)	μ_{atmo}	ν_μ	ν_e	ν_τ
Trigger level	$50.4 \cdot 10^3$	11.1	5.72	0.676
NBI Step1	224	8.32	4.43	0.503
PSP Step1	299	6.88	3.36	0.388

Table 6.3: Rates after the first step of NBI Event Selection. The PSP Step1 rates are also included for comparisons.

Variable	Cut value	μ_{atmo}	ν_μ
Trigger Veto	True	5.3 %	98 %
HLC Veto	True	3.3 %	96 %
NoiseEngine	True	3.1 %	85 %
$\text{Time}_{75\%}$	< 900 ns	1.7 %	83 %
iLineFit Speed	< 0.25 m/ns	1.4 %	82 %
qAbove180	< 12 PE	1.2 %	81 %
RTCluster Hits	< 6	0.85 %	80 %
C2QR6	> 0.6	0.72 %	79 %
Isolated veto Charge	$< 16PE$	0.56 %	74 %
z of first hit	< -180m	0.44 %	74 %

Table 6.4: Summary of the Step1 cuts and their individual efficiency. The efficiencies are relative to the rate at trigger level.

All variables and passing conditions are summarised in Table 6.4 including their individual signal efficiency and fraction of accepted muons. The NoiseEngine algorithm is responsible for the biggest loss of neutrinos. Despite this, NoiseEngine is necessary to include in the event selection, as it has shown to significantly improve the agreement between data and Monte Carlo simulations at low energies with DeepCore [43]. To check that NoiseEngine does not remove high quality events, the distribution of the number of DOMs in the TWSRT pulse series, N_{DOM}, is shown in Figure 6.7 for ν_μ events that fail the NoiseEngine criteria. The majority of the lost events produce less than six hits and are, therefore, likely to yield unreliable reconstructions.

In Figure 6.8 is shown the efficiency of the combined NBI Step1 cuts as a function of the main observables, energy and zenith, as well the vertical and radial position of the interaction vertex which is relevant for the event quality. For simplicity I define a coordinate system shifted in the x,y-plane of the IceCube coordinates, such that the origin is on the centralmost PINGU+DeepCore string[3]:

$$x = x_{IC} - 46.3\text{m} \tag{6.4}$$
$$y = y_{IC} + 34.9\text{m} \tag{6.5}$$
$$z = z_{IC} \tag{6.6}$$

For CC events, the Step1 cuts mildly favor low energy neutrinos, as high energy events are more likely to extend into the veto region. The efficiency drops from about 80% at 5 GeV to about 60% at 50 GeV for CC ν_μ. Neutral current events have a worse efficiency

[3]The central string in the DeepCore-PINGU configuration is number 36

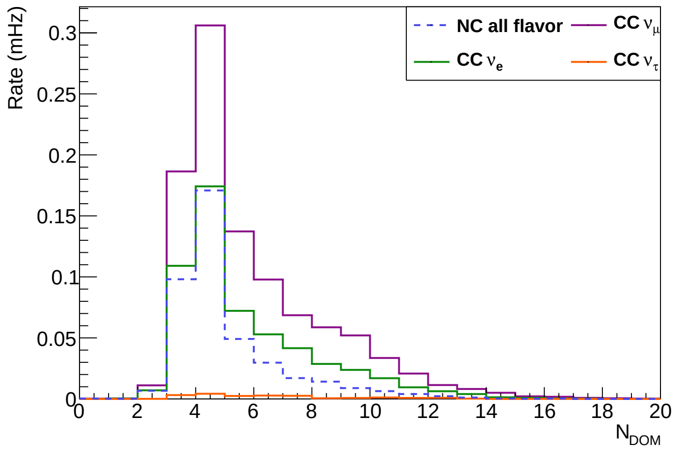

Figure 6.7: Number of hit DOMs in TWSRT pulse series for ν_μ events that fail the NoiseEngine criteria, but pass the other Step1 cuts.

at low energies. Because of the 'missing' energy carried away by the outgoing neutrino, NC events generally have a smaller light yield relative to charge current events of the same energy. Therefore, NC events of a few GeV are more likely to be discarded by the NoiseEngine algorithm than CC events. The ν_τ efficiency drops to zero at zero energy due to the threshold energy of tau production of $E_{\nu_\tau} \approx 3.5\,\text{GeV}$. As is illustrated with the dashed lines, central events in both radius and z are favored. The efficiency is roughly between 70% and 90% for all flavors with an interaction vertex in the PINGU fiducial volume. At large vertex radii, CC ν_μ events have a better efficiency, as the outgoing muon is more likely to extend far enough to cause a PINGU trigger without producing veto hits than a cascade.

Because the current sensitivity predictions for PINGU are based on the PSP event selection, the NBI even selection must keep approximately the same number of neutrino events within the fiducial volume. To compare the vertex positions of the ν_μ events passing the NBI Step1 to those passing PSP Step1, the ratio of passing events as a function of the radial and vertical coordinates (r, z) is plotted in Figure 6.9. The green dashed line indicates the PINGU fiducial volume. Although PSP Step1 mainly relies on containment cuts on the reconstructed vertex position from Monopod, the resulting number of events with contained *true* vertices is similar for the two event selection. The spatial ratio distribution shows that the integrated rate over the full fiducial volume is, in fact, slightly higher following NBI Step1 than for PSP Step1.

6.3.2 NBI Step2: Sneaky Muon Rejection

As it is the case for PSP Step1, the rate following NBI Step1 of 250 mHz is too high to reconstruct with MultiNest. Reconstructing one year of data would occupy for a year ∼1000 of the ∼6000 CPUs available for all IceCube data processing. Reducing the data rate by roughly an order of magnitude would bring down the time required to reconstruct one year of data to a little more than a month. Thereby, significant delays of physics results could be avoided.

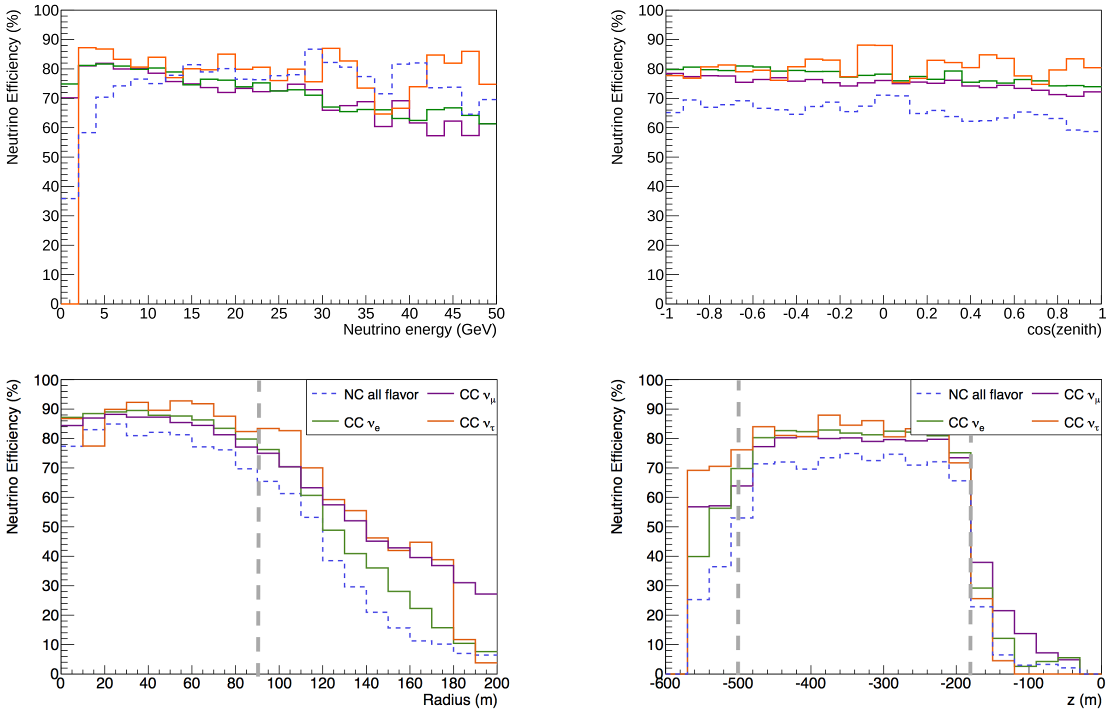

Figure 6.8: Efficiency of NBI Step1 relative to the trigger level. The grey dashed lines indicate the boundaries of the PINGU fiducial volume.

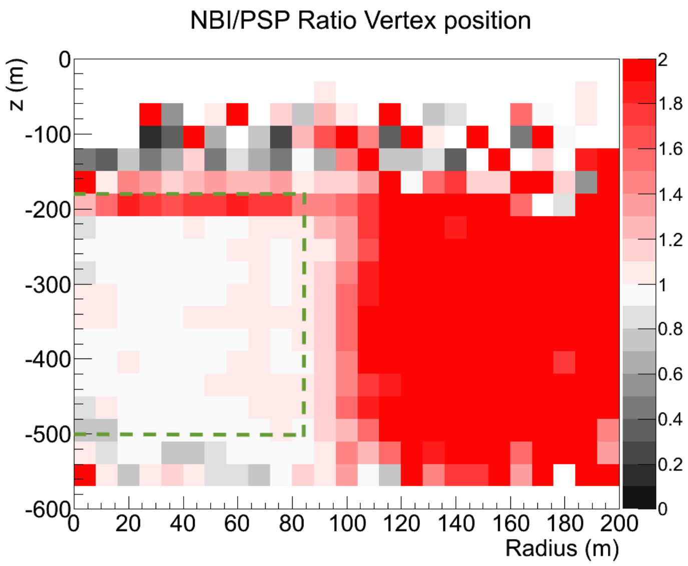

Figure 6.9: Spatial ratio distribution of ν_μ events passing NBI Step1 to PSP Step1. The radius is measured from the central PINGU-DeepCore string.

The second step of the NBI processing reduces the muon rate by an order of magnitude using only trigger level observables, similarly to NBI Step1. Most of the atmospheric muons that pass the Step1 cuts are harder to identify because they leave just a few hits along their tracks caused by scattered photons. The variables devised for Step2 are, therefore, more weakly separating than at Step1 and machine learning techniques are necessary to exploit the information. A Boosted Decision Tree (BDT), is used to construct a single highly discriminating variable - the BDT Score.

A decision tree is a binary tree structured classifier which splits the data into a number of nodes to be classified as either signal or background by applying sequential cuts. Decision trees are easy to visualise and interpret, but also sensitive to statistical fluctuations in the training sample. A Boosted Decision Tree is created from a 'forest' of decision trees, created by reweighting events that are misclassified. Boosting increases the separation power and stability of the classifier [44]. BDTs are known to work well 'out of the box' and require little tuning, even in multivariate analyses with many weakly discriminating and correlated variables, but require a high level of agreement between data and the simulations used for training. The Toolkit for Multivariate Data Analysis with ROOT (TMVA) provides a BDT classifier, which randomly picks events to use for training and testing, respectively, from the list of files given. Therefore, a separate set of files is used for all checks and calculations done subsequent to the BDT cut. The files used for training the BDT include CC and NC ν_μ and $\bar{\nu}_\mu$ events.

6.3.2.1 Step2 Hit Cleaning

Design studies for DeepCore showed that atmospheric muon events with sparsely distributed hits along their trajectory limit the performance of veto algorithms [**?**]. Isolated hits far away from the first trigger are easily misidentified as noise hits by the SRT cleaning which can make an incoming muon appear identical to a neutrino event starting inside the detector. Because the NBI Step1 is based on similar veto techniques as the DeepCore Veto, the same is true for PINGU.

Studying the event displays of a subsample of atmospheric muon events that pass NBI Step1 shows that, in most cases, the track hits are either partially or completely removed by the SRT cleaning. The variables constructed for NBI Step2, are derived from a less strict pulse selection in order to keep every trace of an incoming muon. A static time-window cleaning is applied to the pulse series, preserving all hits from 3500 ns before the PINGU trigger to 500 ns after the trigger. The time-window is asymmetric around the trigger because the muon track is expected to appear prior to the trigger. This method of noise cleaning will be referred to as Time-Window-Before cleaning (TWB) and is illustrated in Figure 6.10. To calculate the variables that specifically look for track-like behaviour in the veto region, the PINGU hits are excluded from the TWB pulse series.

The event display of an atmospheric muon event after the TWB and TWSRT cleaning is shown in Figure 6.11. In the case of this particular muon, a clear track appears in the TWB pulses by which the incoming muon can be vetoed. The track is not present after the TWSRT cleaning, and the pulse map appears identical to that of a neutrino event starting on an outer PINGU DOM.

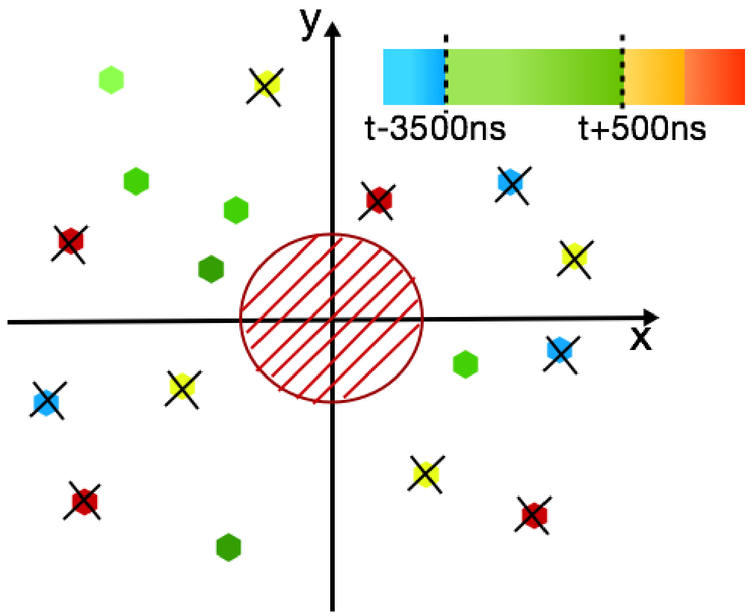

Figure 6.10: Schematic of the pulse selection used to calculate the variables devised for NBI Step2. The time scale is represented by colors shown in the color bar, where t is the time of the PINGU trigger. The central area is shaded, to indicate that PINGU hits can be omitted from the TWB pulse series and only the veto hits considered.

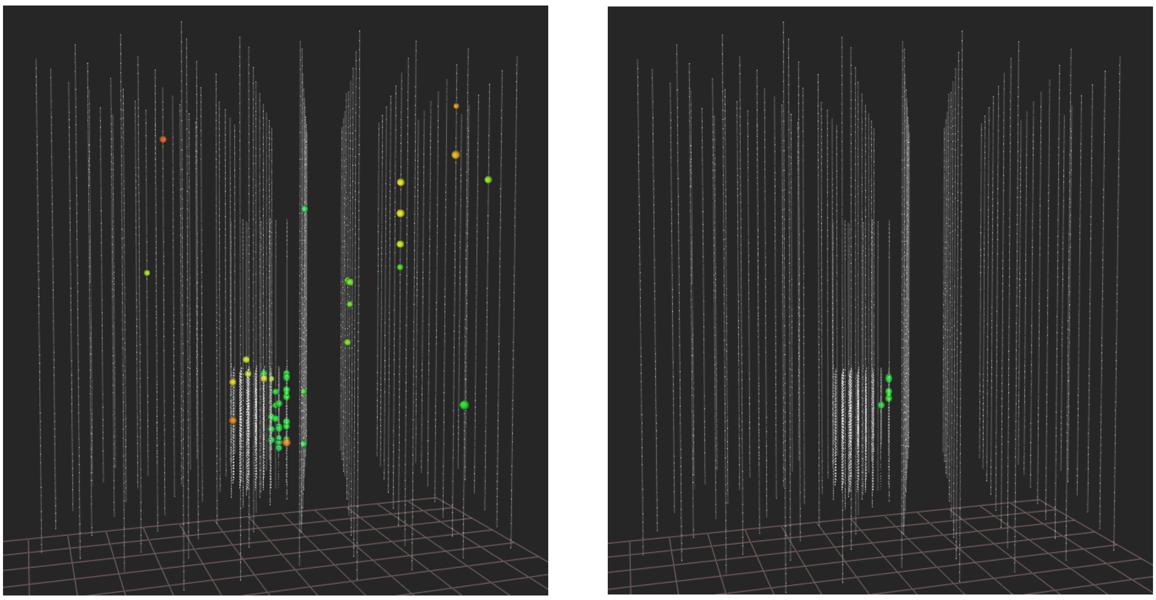

Figure 6.11: Events displays of the same muon event after TWB and TWSRT cleaning respectively. The example demonstrates that the clear track which is present in a time window before the PINGU trigger, is absent after the TWSRT cleaning.

6.3.2.2 NBI Step2 Variables

The variables devised for Step2 of the NBI event selection make use of the fact that the veto hits in a TWB pulse series are randomly distributed for neutrino induced events, whereas they cluster around the track of atmospheric muon events. Step2 combines six variables in a BDT: four variables constructed to identify clusters in the TWB pulse series and two variables recycled from Step1 that still have rejection power. Furthermore, a straight cut is placed on the number of hit DOMs, to ensure the quality of the events passed to MultiNest at Step3 is high.

- **TrackEngine:** The NoiseEngine algorithm, described in section 6.3.1.1, is typically used to identify pure noise triggers by looking for several hits in the same direction, but was originally designed for the opposite purpose of finding tracks. For NBI Step2 TrackEngine requires slightly changed settings; hit pairs have to satisfy a narrower velocity constraint ([0.25, 0.45] m/ns) and the length of the time-window is set to 3500 ns, which should cover all the hits from a muon travelling at the speed of light all the way through the detector (∼1 km). TrackEngine is given only the veto hits from the TWB pulses, and returns the maximum number of hit pairs found in a single (zenith,azimuth)-bin.

- **Average distance between consecutive veto DOMs:** This variable exploits the spatial distribution of veto DOMs in the TWB pulse series. From the list of veto DOMs, ordered according to the time of their first pulse, the distance between each DOM and the previous one is computed and the average distance used as a discriminating variable.

- **Average z-position of hits:** The average vertical position is calculated from all hits (PINGU + veto) in the TWB pulse series. The randomly distributed veto hits in the ν_μ TWB pulse series have have only slightly lower z-coordinates on average than the more down-going atmospheric muon tracks. As the PINGU hits generally constitute a bigger fraction of the TWB pulse series for neutrino events than for atmospheric muons, including the PINGU hits in the calculation of the average-z coordinate improves the separation between ν_μ events and atmospheric muons.

- **Length of position vector sum:** Summing the position vector of all the hits in the TWB pulse series indicates whether the hits are randomly distributed or have some common directionality. If several hits are found in the the same direction, their coordinates add up, whereas randomly distributed hit positions on average cancel out. For simplicity only the x and y coordinates of each hit is used; adding the z-coordinate does not provide better separation. Again, the origin $(x, y) = (0, 0)$ is set at the central PINGU+DeepCore string (see eq. 6.4). The length of summed position vectors is scaled with the number of hits in the pulse series.

$$L_{sum} = \frac{1}{N}\sqrt{\left(\sum_i x_i\right)^2 + \left(\sum_i y_i\right)^2} \tag{6.7}$$

Here N is the number of hits in the pulse series and the subscript i sums over the pulses in the pulse series.

- **Isolated veto charge:** Though the variable is used in Step1 it has rejection power after applying the full set of Step1 cuts. As described in section 6.3.1.1 the number of isolated hits causally related to the trigger is found from the uncleaned pulse series.

- **Time$_{50\%}$:** Using the time duration of an event requires a clean pulse series if the time $t = 0$ is defined by the first hit. When the first hit is caused by noise, the time taken to accumulate 50% of charge in the event can appear much longer than the true value. Therefore, $Time_{50\%}$ is computed from the TWSRT pulse series, as in NBI Step1. A range of values for the charge fraction were tested from 25% to 90% and the best BDT performance was found using the time to accumulate 50% percent of charge in the event.

- **$\mathbf{N_{DOM} > 4}$** : Events with few DOMs are harder to reconstruct and can preferably be removed before running costly reconstruction algorithms. The passing criterion of $N_{DOM} > 4$ is set based on the N_{DOM} distribution for events that pass a preliminary cut from a BDT trained with the six above listed variables. Passing muon events have a spike around $N_{DOM} = 4$ while the ν_μ distribution does not. The final BDT is trained on events that fulfill the N_{DOM} cut. Figure 6.13a shows the N_{DOM} distribution for atmospheric muons and ν_μ events that pass the final BDT cut. The additional cut on $N_{DOM} > 4$ reduces the muon rate further by 5 % and the total neutrino rate by 0.5%.

  To check that 'interesting' neutrino events are not discarded by requiring $N_{DOM} > 4$, Figure 6.13b shows the energy distribution of the ν_μ events that pass the BDT cut, but not the $N_{DOM} > 4$ cut. The energy distribution in the upper plot has a sharp peak at 1 GeV and from the ratio plot below, we see that it corresponds to, maximally, a few percent additional loss.

The distributions of the six variables combined in the BDT are shown in Figure 6.12. The BDT was trained with the above listed six variables on atmospheric muon and ν_μ events that pass the cut on N_{DOM}. The resulting distributions of BDT score for atmospheric muons and ν_μ is shown in Figure 6.14.

For the purpose of optimising the BDT cut, all events that pass NBI Step1 were reconstructed with MultiNest and the relation between BDT score and reconstructed interaction vertex position examined. Figure 6.15 shows the two-dimensional distribution of the BDT score and radial distance of the reconstructed vertex from the central string for atmospheric muon and ν_μ events that satisfy $N_{DOM} > 4$ and $-500\,\text{m} < z_{reco} < -180\,\text{m}$. As expected, atmospheric muons reconstruct further out than the muon neutrinos. The distribution shows that by placing the BDT cut as low as possible with respect to the computational limitations on reconstruction ability, most central neutrinos are kept and the remaining muons will be discarded by the containment cuts.

6.3.2.3 Step2 Performance

The second step of the NBI event selection reduces the atmospheric muon rate to 21.7 mHz, with a total data rate of 33.7 mHz. Not only does Step2 limit the processing time required for Step3 by an order of magnitude, the BDT cut also removes events that would not have been rejected at the final level which is seen in Figure 6.15.

In addition, a few percent more neutrinos are kept than with PSP Step1. The rates at each level are listed in Table 6.5. Averaged over the three neutrino flavors, the NBI Step2 cuts keep roughly 90 % of the neutrinos that pass NBI Step1. The muon and neutrino

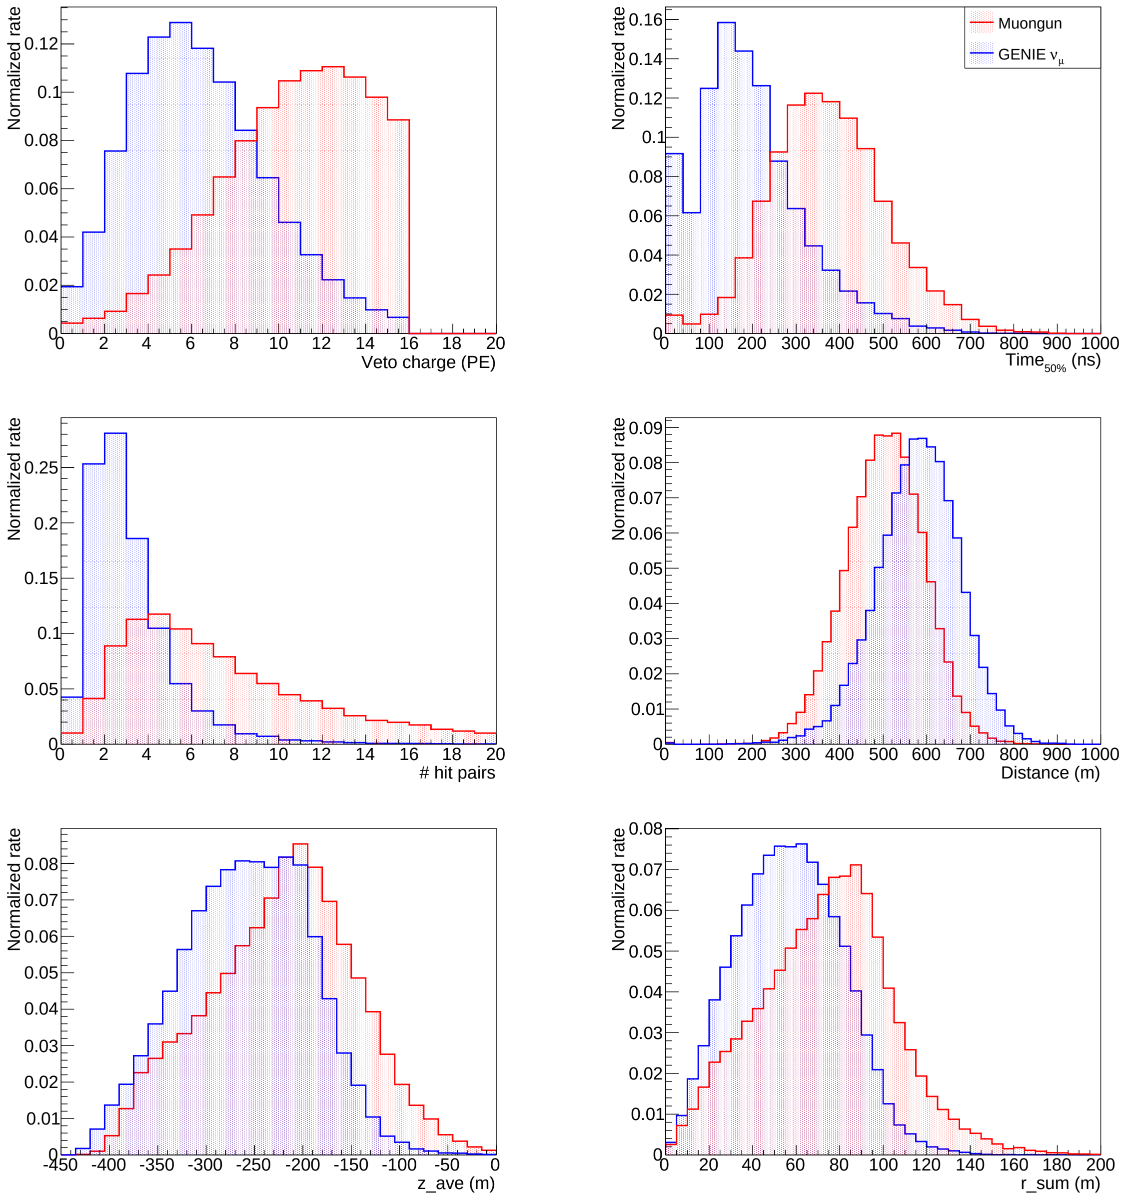

Figure 6.12: Distributions of the Step2 variables used to train the Boosted Decision Tree for events that pass the $N_{DOM} > 4$ cut. Atmospheric muon events are shown in red and the muon neutrino distributions in blue. For "the veto", the impact of the Step1 cut is seen at 16 PE.

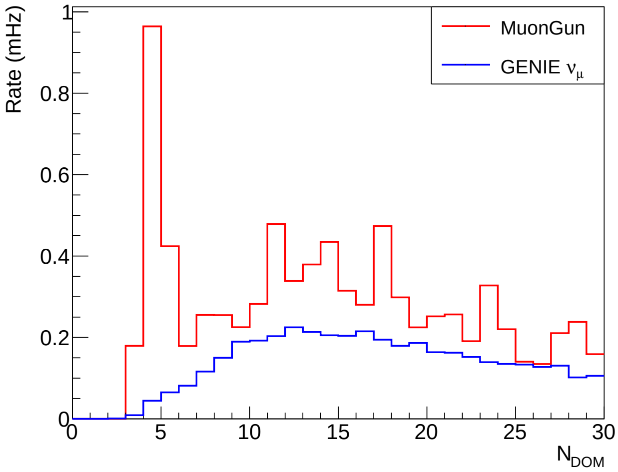

(a)

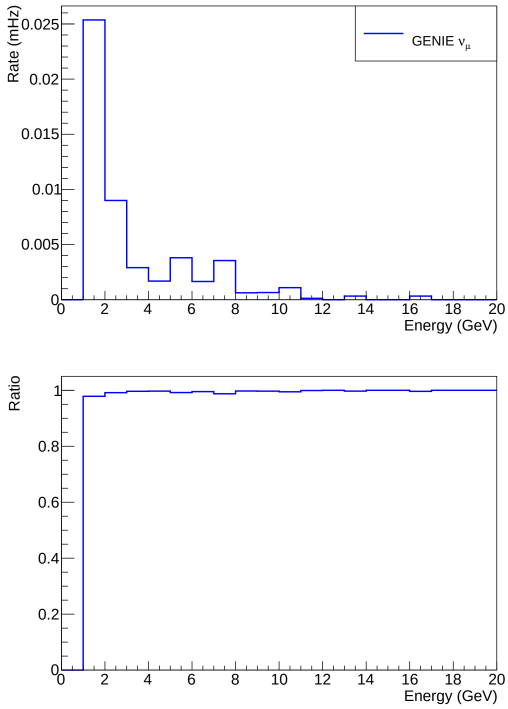

(b)

Figure 6.13: (a) Distribution of N_{DOM} from the TWSRT pulse series following the final cut on the BDT score. Atmospheric muons (MuonGun) are shown in red and muon neutrinos (GENIE) in blue. (b) Energy distribution of ν_μ with $N_{DOM} < 5$ (top) and ν_μ efficiency of the N_{DOM} cut as function of energy (bottom).

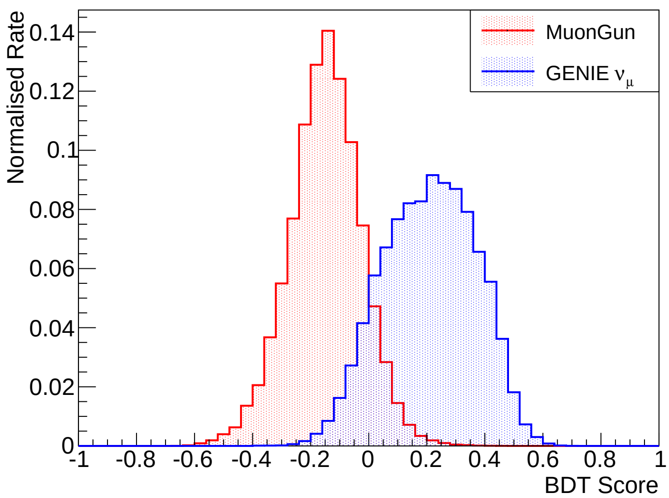

Figure 6.14: Final BDT score distribution after cut on number of DOMs in the TWSRT pulse series.

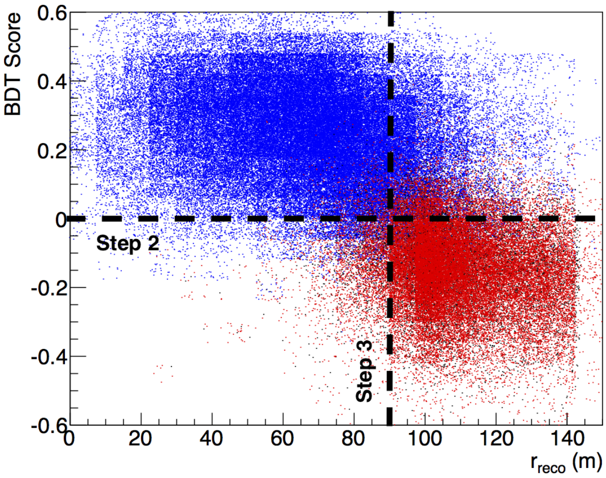

(a) All directions

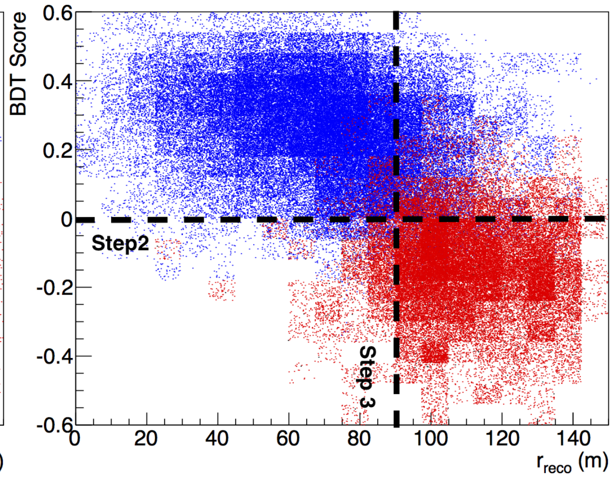

(b) Up-going only

Figure 6.15: Events with $-500\,\text{m} < z_{reco} - 180\,\text{m}$ and $N_{DOMs} > 4$ distributed in BDT score and radius from the central string for atmospheric muons (red) and muon neutrinos (blue).

Rates (mHz)	μ_{atmo}	ν_μ	ν_{e}	ν_τ
NBI Step2	21.7	7.41	4.12	0.472
PSP Step1	299	6.88	3.36	0.388

Table 6.5: Rates following the NBI Step2 cuts. The PSP Step1 rates are also included for comparisons.

efficiencies of the individual Step2 cuts are shown in Table 6.6, where the ratios are relative to the rate following NBI Step1.

Figure 6.16 shows the neutrino efficiencies of Step2 as a function of the energy and direction of the incoming neutrino, as well as the vertical and radial position of the interaction vertex. Again, the ratios are relative to the Step1 rate as we are interested in studying the effect of Step2 alone. Lost events are evenly distributed in zenith angle for all neutrino flavors. The efficiency decreases with energy for all flavors, as the incoming neutrino events are more likely to have a high energy. These high energy events are likely to be recovered in other data streams in IceCube. Within the PINGU fiducial volume, the efficiencies are about 90% or more. The spikes at high radius and z are statistical artefacts due to the very low statistics in those regions - less than five events per bin after the Step2 cuts.

Variable	Cut value	μ_{atmo}	ν_μ	ν_e	ν_τ
N_{DOM}	> 4	97.7 %	99.2 %	99.3 %	98.8 %
BDT score	> 0.0	9.69 %	89.1 %	93.0 %	93.8 %

Table 6.6: Efficiency of each cut of the NBI Step2 event selection. Efficiencies are relative to the rate after NBI Step1.

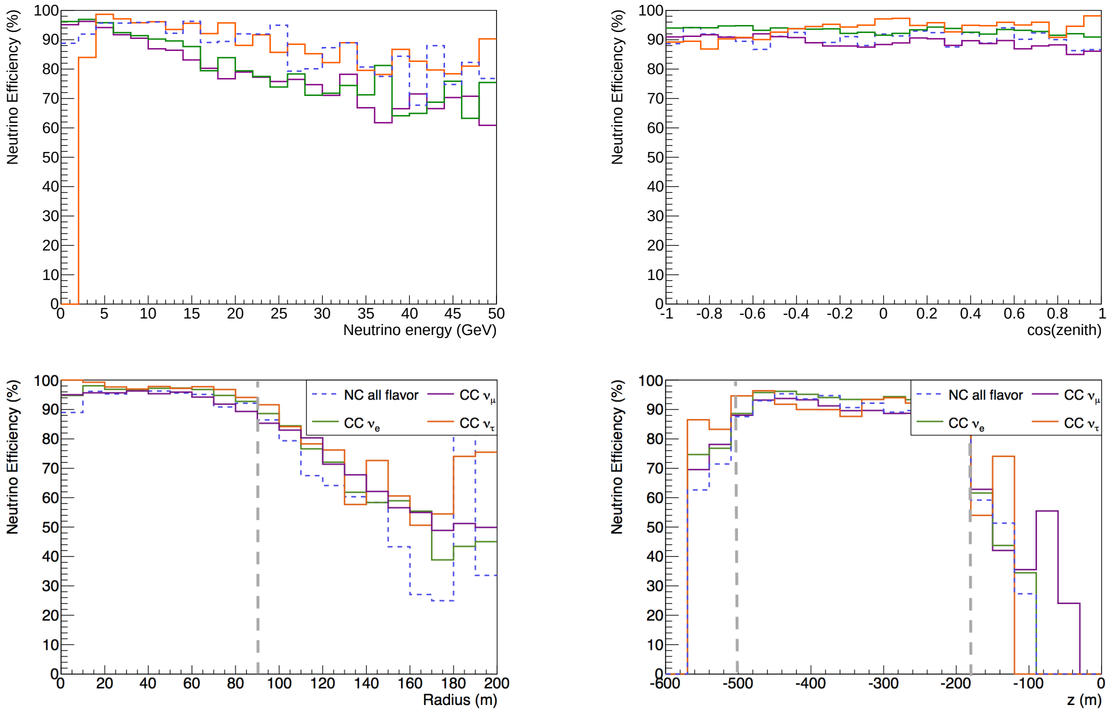

Figure 6.16: Efficiency of NBI Step2 relative to the rates after NBI Step1. The grey dashed lines indicate the boundaries of the PINGU fiducial volume.

To compare the functionality of the NBI event selection with PINGU Standard Processing, the spatial distribution of the surviving neutrinos are compared. Figure 6.17 shows the ratio of the number of ν_μ events passing NBI Step2 to PSP Step1 distributed in vertical and radial position (r, z) of the interaction vertex. The majority of events gained with NBI compared to PSP interact outside the PINGU fiducial volume, but the integrated rates inside (marked by the green dashed lines) are comparable.

6.3.3 NBI Step3: Quality Cuts

The third step of the NBI event selection is a set of containment cuts almost identical to PSP Step2. The containment cuts are designed to ensure a high reconstruction quality of the events at the analysis level, as well as provide further atmospheric muon background rejection. Events are required to have a reconstructed interaction vertex within the PINGU fiducial volume. To determine the optimal range of the fiducial volume, the angular resolution is studied at different vertex radii.

6.3.3.1 Step3 Variables

Just like PSP Step2, the NBI Step3 cuts only regard the containment of the reconstructed vertex position. All reconstructed values in the following refer to the HybridReco/MultiNest reconstruction.

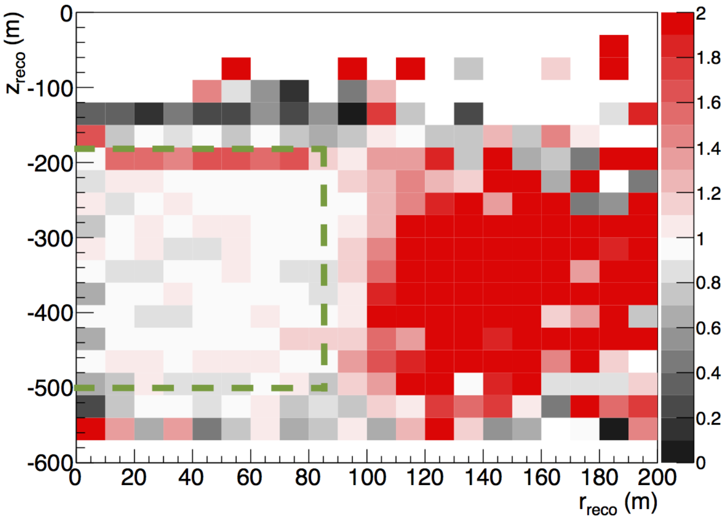

Figure 6.17: Spatial distribution of ν_μ events passing NBI Step2 relative to PSP Step1. The radius is measured from the central IceCube string and the z-coordinate refers to the IceCube coordinate system

- $\mathbf{-500\,m < z_{vertex} < -180\,m}$: The vertical vertex containment cuts are identical to the ones in PSP, and the reasoning behind them the same. The upper limit is at the exact depth of the topmost PINGU DOMs, and is mainly aimed at additional muon background rejection. The lower boundary is meant to remove events with unreliable reconstructions.

- $\mathbf{r_{vertex} < 90\,m}$: The radial containment cut also discards both atmospheric muons and low quality neutrinos. The optimal radius of the fiducial PINGU volume depends on the reconstruction quality of neutrinos at different radii and where the atmospheric muon events reconstruct. In order to determine how the angular reconstruction changes with the radius of the PINGU fiducial volume, the ν_μ zenith resolution is plotted as a function of energy for four different radius bands, shown in Figure 6.18. The radius bands are illustrated along with an overview of the PINGU geometry.

As expected, the quality is worse for interaction vertices reconstructing at higher radii. The resolution bands indicate that the ν_μ reconstruction quality could be sufficient up to a radius of about 100 m, given the current priors on the parameters, x and y, in the MultiNest likelihood space. It would require rerunning the full analyses to determine whether reclaiming neutrinos events reconstructing in a band from 85 m to 100 m would increase the final sensitivities. Therefore, it is reasonable not to discard those events in the generic data processing. However, expanding the radius of the PINGU fiducial volume also increases the muon rate, as seen in Figure 6.15. Therefore, the final radius of 90 m is optimised as the one between 85 m and 100 m that maximises the mock-up sensitivity, defined in Equation 6.3.

- **cos(zenith)>0:** (optional) The highest neutrino purity of the event sample is obtained in the up-going region which is a general preference for all PINGU oscil-

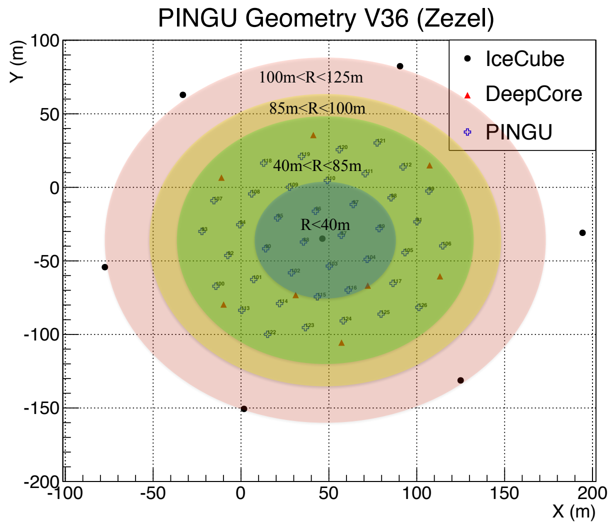

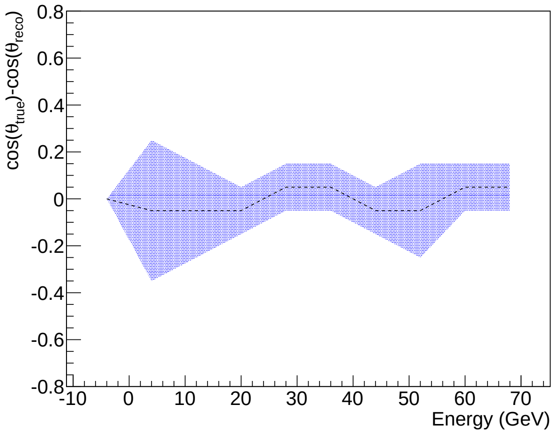

(a) R<40m

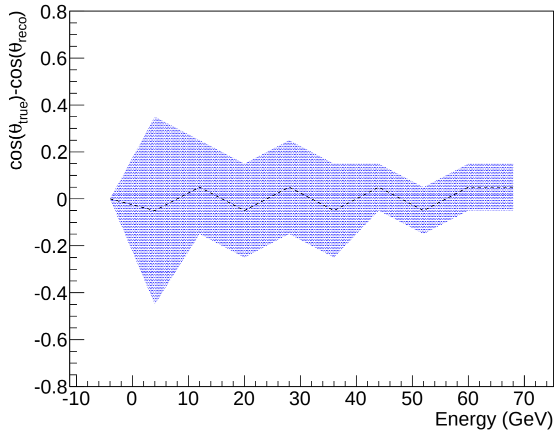

(b) 40m<R<85m

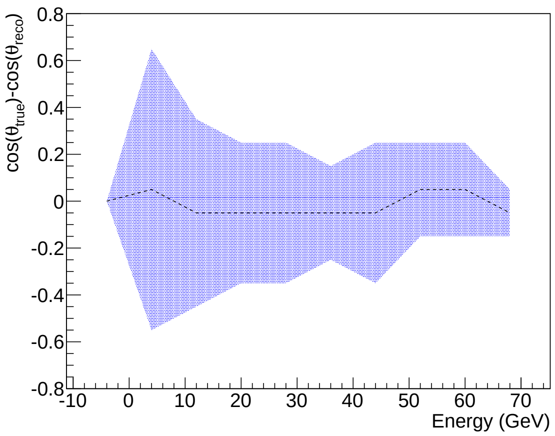

(c) 85m<R<100m

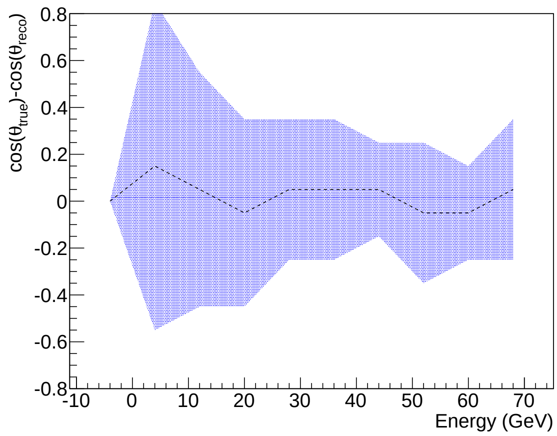

(d) 100<R<125m

Figure 6.18: The zenith resolution of the detector for interaction vertices in at different radius bands. The dashed line shows the median resolution as a function of true energy. The blue error bands cover 68% of the energy bin content. The top plot shows the radial bands overlaying the PINGU string geometry for reference.

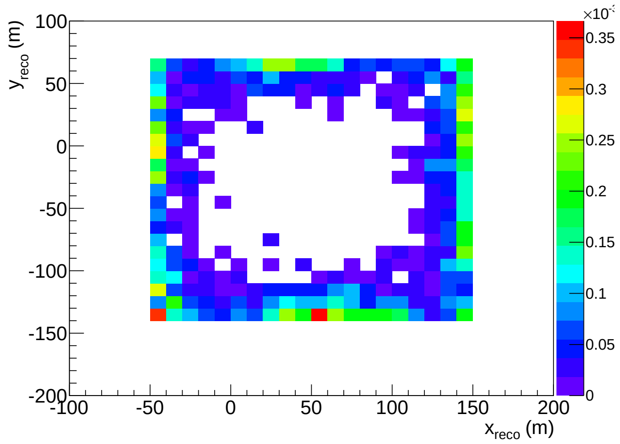

Figure 6.19: Radial distribution of reconstructed vertices for atmospheric muon events that pass NBI Step2 and fulfil $-500\,\text{m} < z_{reco} < -180\,\text{m}$. The coordinates refer to the IceCube coordinate system.

lation analyses. For the mass hierarchy and Earth tomography measurements, only neutrinos that have undergone oscillations in the matter potential of Earth are of interest. This means that only events that are reconstructed as up-going are selected in the final data samples. However, searches for dark matter self-annihilating in the galactic centre can not be as picky, as they look for a neutrino excess in a direction of astrophysical dark matter accumulators above the horizon. The down-going neutrino sample is also used as a control region, e.g. for the normalisation of the atmospheric neutrino flux or to estimate the remaining background contamination at the final cut level. Consequently, the directional cut is provided as an option.

6.3.3.2 A Side Note: The MultiNest Boundaries

The radius of the PINGU fiducial volume is currently constrained by the bounds on the likelihood space that MultiNest maximises (see Figure 5.2). Event reconstructions that converge to the edge of the MultiNest limits are generally less trustworthy. It suggests that the true vertex is somewhere outside of the MultiNest limits. The fact that the angular resolution of the MultiNest reconstruction is good up until a radius where the MultiNest boundaries are tangent to the fiducial volume, implies that the reconstruction quality of events with a true vertex radius slightly above 100 m could improve if the MultiNest boundaries were widened. Thereby, the PINGU fiducial volume could also be expanded. Furthermore, Figure 6.19 shows that the atmospheric muon reconstructions are pushed to MultiNest boundary, indicating that the muon vertices would move further out if the boundary was widened. Ergo, if the MultiNest boundaries were expanded, the PINGU fiducial volume could potentially also increase without loss of muon rejection power.

6.3.3.3 Step3 Performance

The containment cuts applied in NBI Step3 reduce the muon rate by more than 90 % compared to the rate following NBI Step2 at the expense of about 35 % neutrinos. The

Variable	Cut value	μ_{atmo}	ν_μ	ν_e	ν_τ
$\mathrm{z}_{\mathrm{vertex}}$	< -180m	87.1 %	96.6 %	82.0 %	97.0 %
$\mathrm{z}_{\mathrm{vertex}}$	> -500m	78.8 %	88.4 %	75.7 %	86.0 %
$\mathrm{r}_{\mathrm{vertex}}$	< 90m	9.01 %	66.5 %	63.3 %	65.3 %
cos(zenith)	< 0	1.35 %	28.6 %	32.8 %	53.2 %

Table 6.7: Summary of the NBI Step3 event selection. Percentages are relative to the rate after NBI Step2.

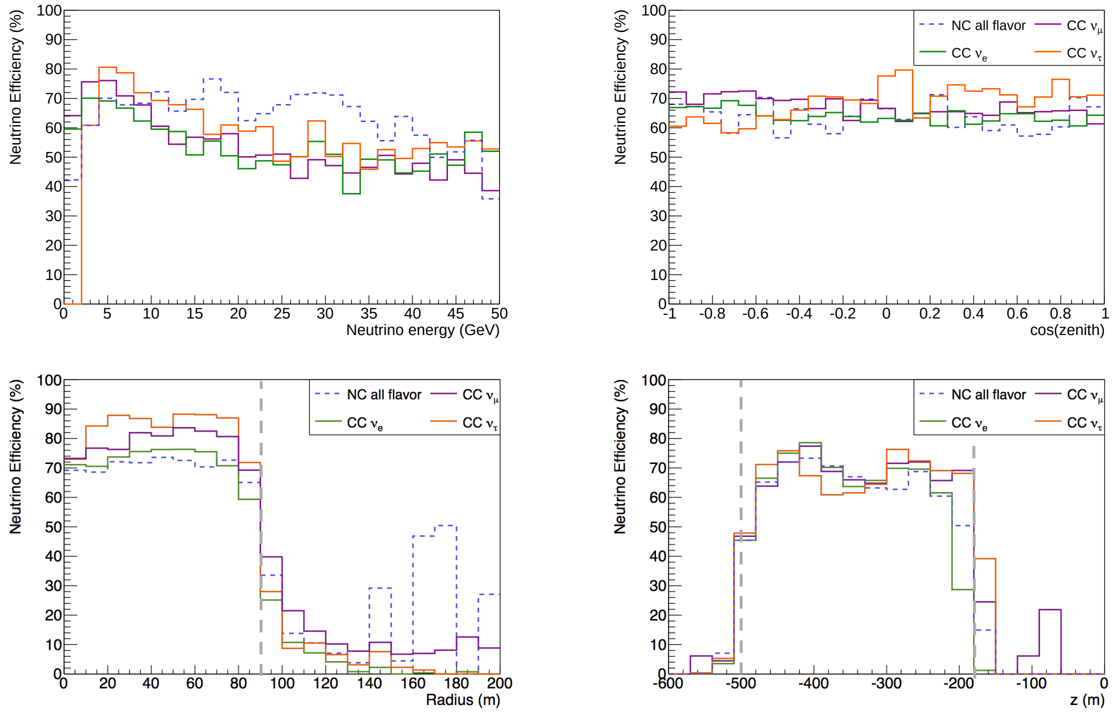

Figure 6.20: Efficiency of NBI Step3 relative to the rates after Step2 as a function of incident neutrino energy, direction, and vertex position.

efficiency of the Step3 cuts, separately, is shown in Table 6.7. The radial containment cut is responsible for the most atmospheric muon rejection. The relatively high muon acceptance of the upper containment cut implies that the previous two steps have successfully removed the atmospheric muons which would have reconstructed in the upper part of the detector.

The combined efficiency of the Step3 cuts are shown as a function of the true zenith and energy of the incoming neutrino as well as radius and z-coordinate of the interaction vertex in Figure 6.20. Again, the neutrino efficiencies are all relative to the rates after NBI Step2, as we are interested in studying the effect of Step3 alone. The efficiency plots for NBI Step3 show the same behavior as the two previous steps. Lost neutrinos are evenly distributed in zenith angle and across the PINGU fiducial volume. For all flavors, the efficiency drops from about 70% at low energies to 50% at higher energies, as incoming events are more likely to have a high energy.

Finally, to compare the spatial distribution of ν_μ events passing the NBI Step3 and PSP Step2, the ratio of the two (r,z)-distributions is plotted in Figure 6.21. The NBI

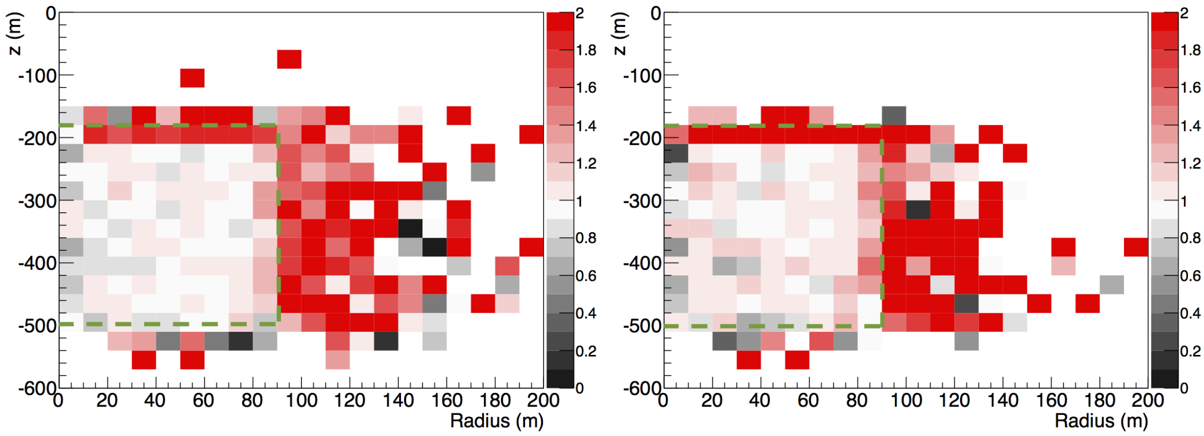

(a) All directions (b) Up-going only

Figure 6.21: Spatial distribution of ν_μ events passing NBI Step3 relative to PSP Step2. The radius is measured from the central IceCube string and the z-coordinate refers to the IceCube coordinate system. The green dashed lines indicated the boundaries of the PINGU fiducial volume.

Rates (mHz)	μ_{atmo}	ν_μ	ν_e	ν_τ
NBI Step3	1.97	4.93	2.61	0.298
PSP Step2	17.8	4.66	2.62	0.278
NBI Step3 Upgoing	0.292	2.12	1.35	0.241
PSP Step2 Upgoing	1.13	1.90	1.33	0.223

Table 6.8: Rates following the NBI Step3 cuts. The PSP Step2 rates are also included for comparisons.

event selection keeps more ν_μ events in the veto region at the final level. Inside the PINGU region, the integrated ν_μ rates are approximately equal, but the atmospheric muon contamination is much lower for the NBI events. Table 6.8 shows the final rates for both event selections, with and without the cut on zenith angle. Compared to PINGU Standard Processing, the NBI event selection yields a muon rate which is 9 times lower in all directions and 4 times lower in the 'only up-going' case. The resulting rates of each step are also illustrated in Figure 6.22.

6.3.4 NBI Event Selection Summary

The NBI Event Selection provides an efficient and realistic method for rejecting atmospheric muons while keeping neutrinos of all types that interact in the fiducial volume of PINGU. The event selection capitalises on the principles of atmospheric muon vetoing with DeepCore. Some variables are taken from DeepCore low energy analyses and others are devised specifically for PINGU. All the variables and cut values are listed in Table 6.9.

	Variable		Passing Criterium
Step1	Trigger Veto		True
	HLC Veto		True
	NoiseEngine		True
	TimeToSum75		< 900 ns
	C2QR6		> 0.6
	Speed from iLineFit		< 0.25 m/ns
	qAbove180		> 12 PE
	$z_{\text{first hit}}$		< -180 m
	RT CLuster Hits		< 6
	Causal Veto Hits		< 16
Step2	TimeToSum50	BDT Score	> 0.0
	Causal Veto Hits		
	TrackEngine Veto Hit Pairs		
	Mean Distance		
	Mean z		
	Position Vector Sum		
	N_{DOM}		>4
Step3	z_{vertex}		> -500 m
	z_{vertex}		< -180 m
	r_{vertex}		< 90 m

Table 6.9: Summary of the NBI Event Selection. Events that pass all criteria are kept for analysis.

Level	μ_{atmo}	ν_μ	ν_e	ν_τ
Step1	0.44 %	74.3 %	77.4 %	74.4 %
Step2	0.043 %	66.2 %	72.0 %	69.8 %
Step3	$3.9 \cdot 10^{-3}$ %	44.0 %	45.6 %	44.1%
Step3 upgoing	$5.8 \cdot 10^{-4}$ %	18.9 %	23.6 %	35.5 %

Table 6.10: Total efficiencies of each step of the NBI event selection. The efficiencies are relative to the SMT3 trigger level.

Rates (mHz)	μ_{atmo}	ν_μ	ν_e	ν_τ
Trigger level	$50.4 \cdot 10^3$	11.1	5.72	0.676
NBI Step1	224	8.32	4.43	0.503
NBI Step2	21.7	7.41	4.12	0.472
NBI Step3	1.97	4.93	2.61	0.298
NBI Step3 Upgoing	0.292	2.12	1.35	0.241

Table 6.11: Rates at the different steps of NBI event selection.

Table 6.10 summarises the efficiency of the three processing levels. The containment cuts are responsible for a large fraction of the lost neutrinos, but significantly enhance the reconstruction reliability. A muon rejection factor of over 10^5 versus trigger level is obtained in the up-going region. The rejection power is estimated from simulated atmospheric muon events with energies less than 500 GeV. As higher energy muons are generally easier to identify, the muon rejection factor is expected to improve for the full range of energies.

Table 6.11 shows the rates of atmospheric muon events and neutrino induced events for the three flavors separately. At the final level the atmospheric muon rate is reduced to 1/4 of the all flavor neutrino rate for all directions and 1/12 in the up-going direction.

6.3.4.1 Description of the Final Sample

The final sample and the efficiency of all three steps combined (excluding the zenith cut in Step3) is described in Figure 6.23. The features of the efficiencies that have been discussed for the individual steps are also visible for the efficiencies of the three steps combined. The event selection is slightly biased towards lower energies which is an effect of the veto logic; high energy starting events are more likely to extend into the veto region, and incoming events are more likely to have high energy. The efficiencies are relatively stable across zenith angles, and the bulk of the atmospheric neutrinos of all flavors interact within the PINGU fiducial volume.

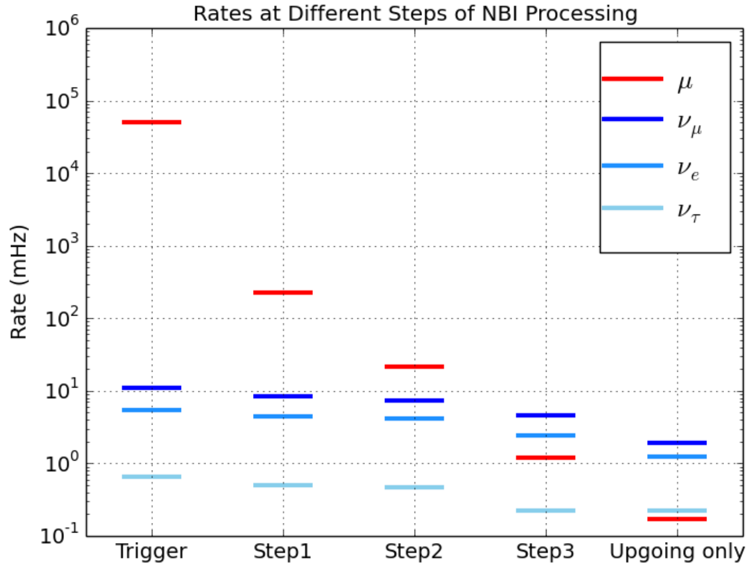

Figure 6.22: Rates at the each level of the NBI event selection. The neutrino rates include both CC and NC events.

6.4 Outlook

To get a full perspective of the performance and adequacy of the NBI Event Selection, the effect of the remaining atmospheric muon background on the final sensitivities must be estimated. The most thorough measure would be to rerun the existing full analyses with the NBI event selection, but seeing as the atmospheric muon background has not yet been included in any analyses, the analysis tools do not currently support the background simulation files. The software tool used for most PINGU analyses, PINGU Simulation Analysis Tool (PISA), is currently being updated to a version suited to include different types of background. The new version is imminent, but due to the time constraint on this thesis, it is not possible to include the sensitivities of a full analysis with the new event selection.

As an indication of how the atmospheric muon background might effect the prospective sensitivity of PINGU to the neutrino mass ordering, the location in energy and zenith of the muons that pass all three steps of the NBI event selection is shown in Figure 6.24 along with the approximate location of the expected NMO signal. The majority of passing muons are found in the down-going region. The ones that do end up in the final up-going sample mainly reconstruct at higher energies and zenith angles than the NMO signal, indicating that they will not influence the current sensitivities greatly.

Work is currently going into optimizing an algorithm for PINGU which, specifically, targets muons that enter through the ‘corridors’ that were shown in Figure 6.4. Once this is done and combined with the NBI Event Selection, and PISA 3.0 is complete, we have a full picture of how the atmospheric muon background affects the final sensitivities.

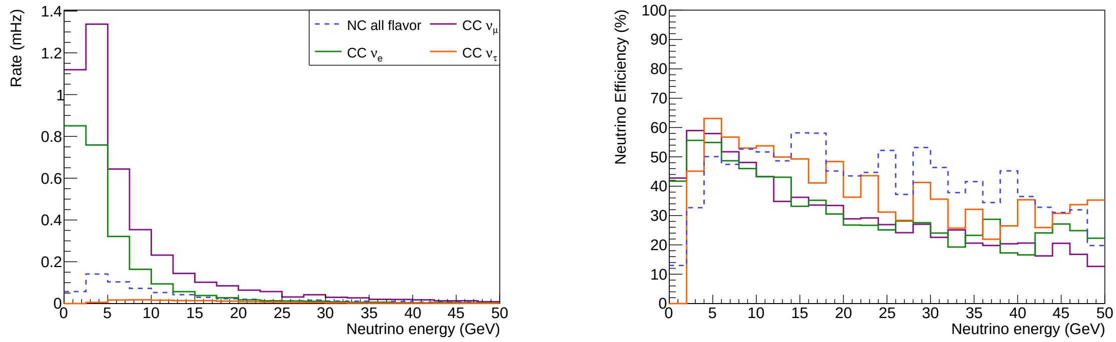

(a) Energy

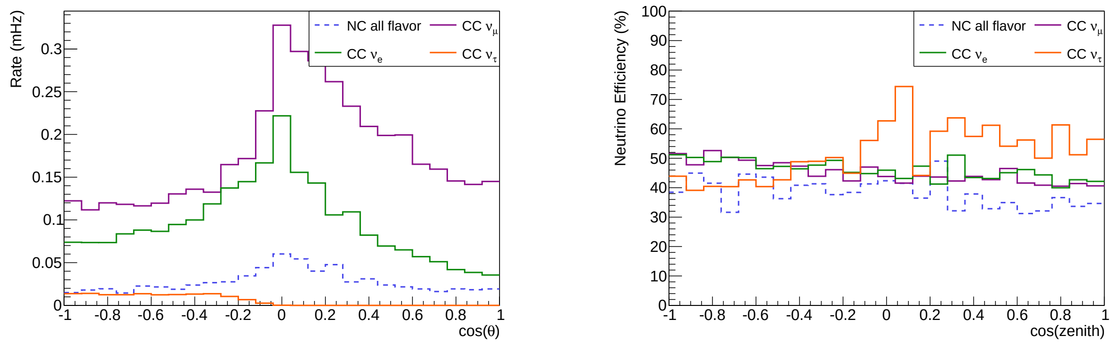

(b) Zenith angle θ

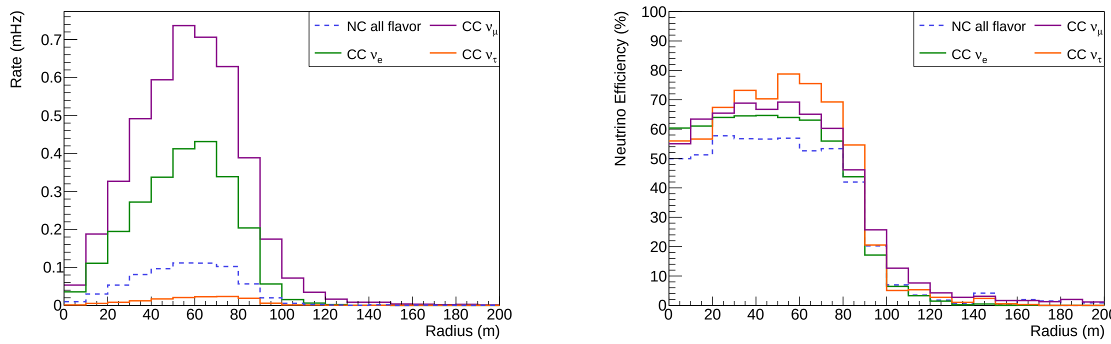

(c) Radial distance of interaction vertex from from center

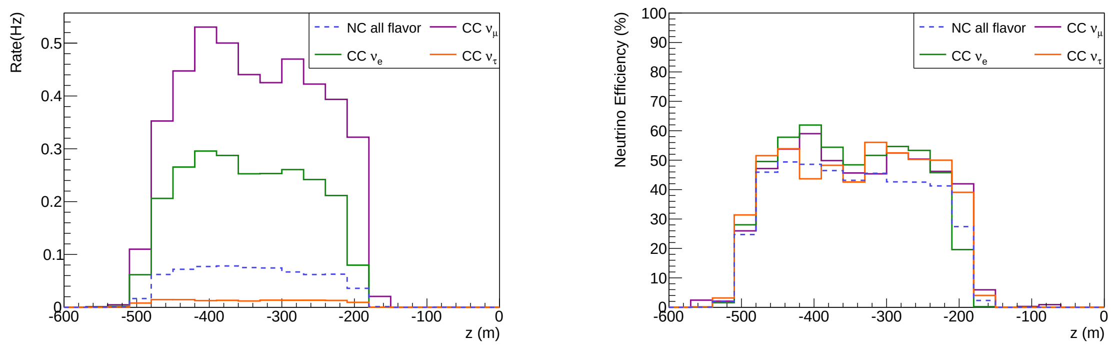

(d) z-coordinate of interaction vertex

Figure 6.23: True distributions at final level. Efficiency of NBI Step3 relative to the trigger level.

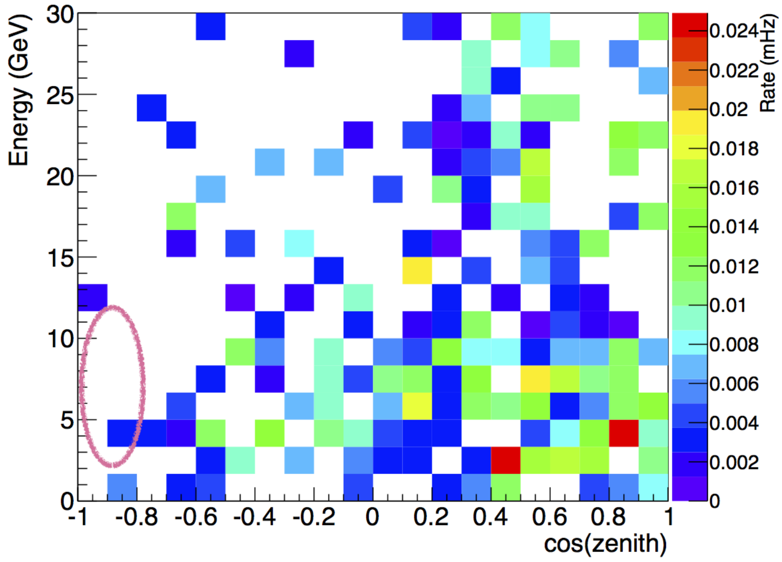

Figure 6.24: The distribution of final level muons in reconstructed zenith and energy. The location of the expected signal is sketched as a pink ellipse.

Chapter 7

Concluding Remarks

In this thesis the development of an analysis independent data selection for the planned PINGU experiment has been presented. The NBI event selection reduces the atmospheric muon background efficiently and can be employed in the PINGU data stream when the detector is built. In a three step procedure, a background rejection factor of more than 10^5 relative to the trigger level is achieved in the up-going direction. Thereby, the atmospheric muon background rate is reduced well below the total neutrino rate. Though the best neutrino purity is obtained in the up-going direction, the neutrino efficiency and muon rejection is sufficient across the full zenith and energy range for the event selection to be applicable to all PINGU analyses.

Current estimates of the prospective sensitivities of PINGU are based on a computing intensive event selection which will not be feasible to run with actual data. The NBI event selection provides an efficient and minimally computing intensive method to obtain equivalent neutrino rates and a lower atmospheric muon rate than PINGU Standard Processing.

The presented event selection accommodates the practical restrictions associated with computing resources and data transmission from the South Pole. Using lightweight software, feasible to run online at the South Pole, the first processing step reduces the atmospheric muon background sufficiently to transmit the data stream to the IceCube central repository via satellite. With a Boosted Decision Tree based on trigger level observables, the second step reduces the atmospheric muon background further, and allows for a computing intensive reconstruction algorithm to be employed in the third step.

The methods used in the NBI event selection are designed to be robust against changes in the detector layout, trigger conditions, and minor deficiencies Monte Carlo Simulations. With the suitable reoptimization of the cut values, the results are expected to be similar for the final detector geometry.

Bibliography

[1] G. Danby, J-M. Gaillard, K. Goulianos, L. M. Lederman, N. Mistry, M. Schwartz, and J. Steinberger. Observation of high-energy neutrino reactions and the existence of two kinds of neutrinos. *Phys. Rev. Lett.*, 9:36–44, Jul 1962. doi: 10.1103/PhysRevLett.9.36. URL http://link.aps.org/doi/10.1103/PhysRevLett.9.36.

[2] B. T. Cleveland, Timothy Daily, Raymond Davis, Jr., James R. Distel, Kenneth Lande, C. K. Lee, Paul S. Wildenhain, and Jack Ullman. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophys. J.*, 496: 505–526, 1998. doi: 10.1086/305343.

[3] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti-neutrino disappearance. *Phys. Rev. Lett.*, 90:021802, 2003. doi: 10.1103/PhysRevLett.90.021802.

[4] R.L. Helmer. First results from the sudbury neutrino observatory. *Nuclear Physics B - Proceedings Supplements*, 111(1):122 – 127, 2002. ISSN 0920-5632. doi: http://dx.doi.org/10.1016/S0920-5632(02)01693-6. URL http://www.sciencedirect.com/science/article/pii/S0920563202016936.

[5] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962. doi: 10.1143/PTP.28.870. URL http://ptp.oxfordjournals.org/content/28/5/870.abstract.

[6] V. Gribov and B. Pontecorvo. Neutrino astronomy and lepton charge. *Physics Letters B*, 28(7):493 – 496, 1969. ISSN 0370-2693. doi: http://dx.doi.org/10.1016/0370-2693(69)90525-5. URL http://www.sciencedirect.com/science/article/pii/0370269369905255.

[7] David Galbraith. The standard model of the standard model, 2012. URL http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/.

[8] David Griffiths. *Introduction to Elementary Particle Physics.* Wiley-VCH, 2008.

[9] K.A. Olive et al. (Particle Data Group). Leptons. In *Chin. Phys. C*, volume 38, page 090001. Chinese Physical Society, 2015. URL http://pdg.lbl.gov/2015/tables/rpp2015-sum-leptons.pdf.

[10] Ch et al Kraus. Final results from phase ii of the mainz neutrino mass searchin tritium β decay. *The European Physical Journal C - Particles and Fields*, 40(4): 447–468, 2005. ISSN 1434-6052. doi: 10.1140/epjc/s2005-02139-7. URL http://dx.doi.org/10.1140/epjc/s2005-02139-7.

[11] V. N. Aseev, A. I. Belesev, A. I. Berlev, E. V. Geraskin, A. A. Golubev, N. A. Likhovid, V. M. Lobashev, A. A. Nozik, V. S. Pantuev, V. I. Parfenov, A. K. Skasyrskaya, F. V. Tkachov, and S. V. Zadorozhny. Upper limit on the electron antineutrino mass from the troitsk experiment. *Phys. Rev. D*, 84:112003, Dec 2011. doi: 10.1103/PhysRevD.84.112003. URL `http://link.aps.org/doi/10.1103/PhysRevD.84.112003`.

[12] S. Schael et al. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006. doi: 10.1016/j.physrep.2005.12.006.

[13] J. A. Formaggio and G. P. Zeller. From eV to EeV: Neutrino Cross Sections Across Energy Scales. *Rev. Mod. Phys.*, 84:1307, 2012. doi: 10.1103/RevModPhys.84.1307.

[14] Kevin Scott McFarland. Neutrino Interactions. In *Neutrinos in particle physics, astrophysics and cosmology. Proceedings, 61st Scottish Universities Summer School in Physics, SUSSP61, St. Andrews, UK, August 8-23, 2006*, pages 65–90, 2008. URL `https://inspirehep.net/record/784224/files/arXiv:0804.3899.pdf`.

[15] D. Binosi and L. Theußl. Jaxodraw: A graphical user interface for drawing feynman diagrams. *Computer Physics Communications*, 161(1-2):76 – 86, 2004. ISSN 0010-4655. doi: http://dx.doi.org/10.1016/j.cpc.2004.05.001. URL `http://www.sciencedirect.com/science/article/pii/S0010465504002115`.

[16] B.R. Martin and G. Shaw. *Particle Physics*. Wiley, 3 edition, 2008.

[17] on behalf of the NOvA Collaboration Patricia Vahle. New results from nova. Presented at the 27th annual Neutrino Conference, London, 2016. URL `http://nova-docdb.fnal.gov/cgi-bin/RetrieveFile?docid=15688&filename=vahle-nova-neutrino2016.pdf&version=15`.

[18] on behalf of the T2K Collaboration H. A. Tanaka. T2k: Latest results. Presented at the 27th annual Neutrino Conference, London, 2016. URL `http://www.t2k.org/docs/talk/236/slides`.

[19] R. B. Patterson. Prospects for Measurement of the Neutrino Mass Hierarchy. *Ann. Rev. Nucl. Part. Sci.*, 65:177–192, 2015. doi: 10.1146/annurev-nucl-102014-021916.

[20] A. Yu. Smirnov. The MSW effect and matter effects in neutrino oscillations. *Phys. Scripta*, T121:57–64, 2005. doi: 10.1088/0031-8949/2005/T121/008.

[21] Q. Y. Liu, S. P. Mikheyev, and A. Yu. Smirnov. Parametric resonance in oscillations of atmospheric neutrinos? *Phys. Lett.*, B440:319–326, 1998. doi: 10.1016/S0370-2693(98)01102-2.

[22] Chee Sheng Fong, Enrico Nardi, and Antonio Riotto. Leptogenesis in the Universe. *Adv. High Energy Phys.*, 2012:158303, 2012. doi: 10.1155/2012/158303.

[23] M. G. Aartsen et al. Letter of Intent: The Precision IceCube Next Generation Upgrade (PINGU), 2014.

[24] IceCube Gen2. PINGU: A Vision for Neutrino and Particle Physics at the South Pole. 2016.

[25] Jörg R. Hörandel. Cosmic-ray composition and its relation to shock acceleration by supernova remnants. *Advances in Space Research*, 41(3):442 – 463, 2008.

ISSN 0273-1177. doi: http://dx.doi.org/10.1016/j.asr.2007.06.008. URL http://www.sciencedirect.com/science/article/pii/S0273117707006308.

[26] P. H. Fowler, R. N. F. Walker, M. R. W. Masheder, R. T. Moses, and A. Worley. Ultra-heavy cosmic ray studies with ariel vi. *Nature*, 291(5810):45–47, 05 1981. URL http://dx.doi.org/10.1038/291045a0.

[27] S.P. Sword. The energy spectra and anisotropies of cosmic rays. *Space Science Reviews*, 99:85–94, 2001. doi: 10.1023/A:1013828611730.

[28] Takaai Kajita. Atmospheric neutrinos. *Advances in High Energy Physics*, 2012, 2012. doi: 10.1155/2012/504715.

[29] M. Honda, T. Kajita, K. Kasahara, S. Midorikawa, and T. Sanuki. Calculation of atmospheric neutrino flux using the interaction model calibrated with atmospheric muon data. *Phys. Rev. D*, 75:043006, Feb 2007. doi: 10.1103/PhysRevD.75.043006. URL http://link.aps.org/doi/10.1103/PhysRevD.75.043006.

[30] Anne Schukraft. A view of prompt atmospheric neutrinos with IceCube. *Nucl. Phys. Proc. Suppl.*, 237-238:266–268, 2013. doi: 10.1016/j.nuclphysbps.2013.04.105.

[31] P.K.F. Grieder. *Cosmic Rays at Earth*. Elsevier Science, 2001.

[32] R. E. Mikkelsen, T. Poulsen, U. I. Uggerhøj, and S. R. Klein. Characteristics of Cherenkov radiation in naturally occurring ice. *Phys. Rev.*, D93(5):053006, 2016. doi: 10.1103/PhysRevD.93.053006.

[33] M. Ackermann et al. Optical properties of deep glacial ice at the south pole. *Journal of Geophysical Research: Atmospheres*, 111(D13), 2006. ISSN 2156-2202. doi: 10.1029/2005JD006687. URL http://dx.doi.org/10.1029/2005JD006687. D13203.

[34] M. Alam et al. GENIE Production Release 2.10.0. 2015.

[35] S. Agostinelli et al. GEANT4: A Simulation toolkit. *Nucl. Instrum. Meth.*, A506: 250–303, 2003. doi: 10.1016/S0168-9002(03)01368-8.

[36] IceCube Collaboration. Muongun, 2016. URL http://software.icecube.wisc.edu/simulation_trunk/projects/MuonGun/index.html.

[37] Thomas K. Gaisser, Todor Stanev, and Serap Tilav. Cosmic Ray Energy Spectrum from Measurements of Air Showers. *Front. Phys.(Beijing)*, 8:748–758, 2013. doi: 10.1007/s11467-013-0319-7.

[38] R. Abbasi et al. The Design and Performance of IceCube DeepCore. *Astropart. Phys.*, 35:615–624, 2012. doi: 10.1016/j.astropartphys.2012.01.004.

[39] J. L. Kelleyand IceCube Collaboration. Event triggering in the icecube data acquisition system. In *API Conference Proceedings*, volume 1630, 2014.

[40] O. Schultz. Slc hit cleaning, 2009. URL https://wiki.icecube.wisc.edu/index.php/SLC_hit_cleaning.

[41] F. Feroz, M. P. Hobson, and M. Bridges. Multinest: an efficient and robust bayesian inference tool for cosmology and particle physics. *Mon. Not. Roy. Astron. Soc.*, 398: 1601–1614, 2009. doi: 10.1111/j.1365-2966.2009.14548.x.

[42] Raj Gandhi, Chris Quigg, Mary Hall Reno, and Ina Sarcevic. Ultrahigh-energy neutrino interactions. *Astropart. Phys.*, 5:81–110, 1996. doi: 10.1016/0927-6505(96)00008-4.

[43] Michael Larson. Simulation and identification of non-poissonian noise triggers in the icecube neutrino detector. Master's thesis, The University of Alabama, 2013.

[44] A. Hoecker et al. TMVA - Toolkit for Multivariate Data Analysis. *ArXiv Physics e-prints*, March 2007.

List of Figures

List of Tables