Direction Reconstruction of IceCube Neutrino Events with Millipede

Alexander Wallace

Thesis submitted for the Degree of Master of Philosophy

Principal Supervisor: Doctor Gary Hill Co-Supervisor: Professor Bruce Dawson

> The University of Adelaide School of Physical Sciences Department of Physics

> > May 2016

Abstract

To conduct neutrino astronomy with the IceCube detector at the South Pole, the direction of the incoming neutrino must be known accurately to within one degree. When a muon neutrino interacts in the ice at the South Pole, it produces a muon which produces Cherenkov light as it travels through the detector. Using the direction of the muon, the direction of the original neutrino can be determined and used for astronomy. Millipede is an algorithm used to numerically determine the properties of the muon track by making predictions about the light signal seen in the detector and checking how this compares to the observed signal using a likelihood maximisation.

With this algorithm, the muon track direction is expected to be resolved to within one degree. However, problems have been encountered with simulated muons where millipede finds a direction which is very different from the true direction or millipede fails to reconstruct the event. After analysis of the likelihood grid scans of some of these events, the problems with millipede seem to be due to the minimiser finding a local minimum in the likelihood surface rather than the desired global minimum. These local minima arise from fluctuations in the likelihood surface. These fluctuations were observed in all dimensions including track position.

The source of these fluctuations was investigated in simulations by first using millipede's predictions as the input waveforms. Poisson fluctuations were then added and produced a less accurate likelihood scan with more fluctuations. Finally, the effect of photomultiplier after-pulses was investigated by removing all signal more than 3μ s after the median time. Removing this signal dramatically improves some of the likelihood scans but many show no change.

After this analysis, the main factors causing these fluctuations in the likelihood surface seem to be a combination of bin-wise fluctuations in the waveform and the presence of after-pulses which are not taken into account by millipede. The after-pulses and other late light seem to be the dominant cause across a range of energies, though generally high energy events, while the fluctuations are the dominant cause for the low energy events.

Declaration of Originality

I, Alexander Wallace, certify that this work contains no material which has been accepted for the award of any other degree or diploma in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library catalogue and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Alexander Wallace

Acknowledgements

I would like to thank my supervisors Gary Hill and Bruce Dawson for their constant guidance and support during the last two years and helping to edit this thesis. I would like to thank Ben Whelan for his extensive help with the technical side of this work.

It has been great working alongside other members of the High Energy Astrophysics Group. In particular I would like to thank the people I shared an office with over the last two years: Mark Aartsen, Sally Robertson, Rebekah Little and Alexander Kyriacou have all been friendly and supportive over the course of this work.

Finally, I would like to thank my parents, Debbie and Steve, for their endless support and encouragement throughout all my studies.

Contents

1	Neutrinos							
	1.1	Background and Discovery						
	1.2	Neutrinos in the Standard Model	2					
	1.3	Neutrino Interactions	3					
		1.3.1 Lepton Family Conservation	4					
	1.4	Summary	5					
2	High Energy Astrophysics							
	2.1	Cosmic Rays	7					
		2.1.1 Cosmic Ray Acceleration	7					
		2.1.2 Cosmic Ray Propagation Through Space	8					
	2.2	Production of Secondary Particles	9					
		2.2.1 Neutral Messenger Particles	10					
	2.3	Particle Detection On Earth	13					
		2.3.1 Cherenkov Radiation	13					
		2.3.2 Detection Methods	14					
	2.4	Summary	16					
3	The	e IceCube Neutrino Observatory	17					
	3.1	Detector Layout	17					
	3.2	Light Signal in DOMs	18					
		3.2.1 Detection of Cherenkov Light	18					
		3.2.2 Signals of Different Flavours	21					
	3.3	Optical Properties of the Ice	23					
	3.4	Discovery of Astrophysical Neutrinos	24					
4	Eve	nt Reconstruction	27					
	4.1	IceCube Coordinate System	27					
		4.1.1 Definition of Vertex	28					
	4.2	Line-Fit	29					
	4.3	SPE and MPE	30					
	4.4	Millipede	32					
		4.4.1 Millipede Time Binning	34					

		4.4.2Millipede Likelihood	36 38
5	Init	ial Testing of Millipede	39
-	5.1	Comparison of Reconstructions	39
	5.2	Grid Scans	11
	-	5.2.1 Healpix Grid	12
		5.2.2 Example Scans of HESE Neutrinos	13
		5.2.3 Testing with HESE Track Event	14
		5.2.4 Scans of Four Simulated Events	16
	5.3	Test of Statistical Errors	18
		5.3.1 1-dimensional Scans and Curve Fitting	18
		5.3.2 Results from multiple events	54
		5.3.3 Test of Overall Smoothness	56
		5.3.4 Comparison to smooth function	58
	5.4	Summary	59
6	Inv	estigation of Likelihood Fluctuations	61
	6.1	Close Grid Scan	52
		6.1.1 Vertex Shifts	52
	6.2	Vertex Scans	38
		6.2.1 Three-dimensional Vertex Scan	38
		6.2.2 Plane Scan	39
	6.3	Fixed Vertex	70
		6.3.1 Fixed Energy Losses	70
		6.3.2 Only Fitting Energy Losses	72
	6.4	Summary	73
7	Pos	sible Causes of Vertex Fluctuations 7	$\mathbf{'5}$
	7.1	Using Millipede Predictions as Input	75
		7.1.1 Adding Poisson Fluctuations	76
	7.2	Attempts to Remove After-pulses	77
		7.2.1 Reconstructions with Time Cut	30
		7.2.2 Tightness of the Minimum	35
		7.2.3 Overall Accuracy of Scans	39
		7.2.4 Comparison of Fluctuations)1
	7.3	Adding After-pulses and Fluctuations)3
		7.3.1 Adding Poisson Fluctuations)3
		7.3.2 Adding After-pulses)5
	7.4	Summary	98
8	Cor	nclusions 10)1
	8.1	Future Work)2

Appendix A Derivations						
A.1 Cosmic Ray Acceleration and Spectrum	. 10					
A.1.1 Second Order fermi Acceleration	10					
A.1.2 First Order Fermi Acceleration	10					
A.1.3 Cosmic Ray Spectrum	. 10					
A.2 Geometric Time for SPE and MPE	. 10					
A.3 Space Angle Formula	. 10					
A.4 Chi Squared Critical Value	. 11					
A.5 Perpendicular Plane Coordinates	. 11					
Appendix B Technical Details of Simulation and Reconstruction 1						
B.1 The IceTray Software	. 11					
B.2 The MC Tree	. 11					
B.3 Using the Millipede Likelihood Function as the Simulation \ldots						
B.4 Added Poisson Fluctuations	. 11					
B.5 Added After-pulses	. 11					
Appendix C Reconstruction Python Code						
C.1 Millipede Free Fit	. 11					
C.2 Grid Scan	. 12					
References						

Chapter 1

Neutrinos

1.1 Background and Discovery

Neutrinos are electrically neutral, almost massless particles which interact only weakly with matter. Due to this, they pass through most matter unaffected. This means there is very little direct evidence of their presence and they are therefore very hard to detect. The existence of the neutrino was first predicted in 1930 by Wolfgang Pauli to explain the perceived violation of conservation laws in radioactive β -decay [1] which is divided into two types: β^+ and β^- decay. In β^- decay, a nucleus has an excess of neutrons and is considered unstable. This causes one of the neutrons to transform into a proton to produce a more stable nucleus. Since the proton has positive electric charge and the neutron is neutral, this process emits a negatively charged electron or β^- particle to conserve the total charge. In β^+ decay, the opposite reaction occurs. A proton transforms into a neutron which emits a positron (or anti-electron) to conserve the positive charge.

These processes were observed by Pauli who measured the "daughter" products of the decay. It soon became clear that certain properties in the reaction were not being conserved. While the electron and positron conserved the total charge, Pauli noticed that the total energy and angular momentum of the parent and daughters were not the same. This led to the theory that the missing energy and angular momentum were being carried away by an as-yet undetected particle. This particle had to be electrically neutral since charge was already accounted for and it had to have a very low interaction probability to explain its lack of detection. This particle was later named the neutrino and denoted by the Greek letter ν . The equations for β^- and β^+ decay then became the following:

$$n \to p + e^- + \bar{\nu}_e \tag{1.1}$$

$$p \to n + e^+ + \nu_e \tag{1.2}$$

where ν_e is the electron neutrino and $\bar{\nu}_e$ is the anti-electron neutrino. It was thought, for many years, that the neutrino was completely undetectable and was

therefore not falsifiable. However, in 1956, Clyde Cowan and Frederick Reines conducted a successful experiment to directly detect the anti-neutrino [2]. In this experiment anti-neutrinos were created through β -decay to interact with a large proton target. Though neutrinos interact very rarely with matter, if there are enough of them present, the probability of at least one interaction increases significantly. When the anti-neutrino interacts with the proton, it does so by a weak interaction to produce a positron and neutron.

$$\bar{\nu}_e + p \to n + e^+ \tag{1.3}$$

This positron then quickly encounters an electron and, as they are each other's anti-particle, they annihilate into pure energy. Since the electron and positron have a known rest energy of 511 keV (511×10^3 electron volts), double this amount of energy is emitted in the form of two gamma rays travelling in opposite directions. A measurement of 511 keV gamma rays can then be used to confirm the annihilation took place. This verified the reaction in equation 1.3 occurred and confirmed the existence of the (anti) neutrino.

1.2 Neutrinos in the Standard Model

The standard model of particle physics is a theory used to describe the different types of subatomic particles and the forces between them. In the standard model, all matter is made up of "elementary" particles which are defined as having no known sub-structure [3]. These particles are divided into two distinct types known as fermions and bosons. The fermions obey laws known as Fermi-Dirac statistics while the bosons obey Bose-Einstein statistics.

In this model, there are four fundamental forces: the gravitational force, the electromagnetic force, the weak nuclear force which causes radioactive decay, and the strong nuclear force which holds the nucleus together. The gravitational force acts on any particle with mass while the electromagnetic force acts on any particle with electric charge. Each of these forces is thought to be carried by a different boson which is exchanged between interacting particles. Three of the forces now have associated bosons whose existence has been confirmed, though the "graviton" which carries gravity has not yet been directly detected and remains hypothetical.

The fermions are themselves divided into quarks and leptons. The quarks can experience all four of the forces while the leptons are unaffected by the strong nuclear force. The standard model is summarized in table 1.1.

	Bosons			
Generation	Ι	II	III	
Quarks	up	charm	top	γ
	down	strange	bottom	g
Leptons	ν_e	$ u_{\mu} $	ν_{τ}	Z^0
	e	μ	au	W^{\pm}

Table 1.1: Table showing particles in the standard model with fermions on the left and bosons on the right.

Here, ν_e, ν_μ and ν_τ are the three types or "flavours" of neutrinos which have corresponding charged leptons, the electron (e), the muon (μ) and the tau (τ). The photon, denoted γ , is a particle of light which carries the electromagnetic force and g denotes the gluon which carries the strong nuclear force and holds quarks together to form hadrons. The neutral Z^0 boson and the charged W^+ and W^- bosons carry the weak nuclear force.

The fermions are divided into three "generations" of matter of increasing mass and instability. The up, charm and top quarks have a charge of +2/3 times the charge of the proton while the down, strange and bottom quarks have a charge of -1/3. Quarks can combine in groups of three to produce composite particles called baryons. Every particle in the standard model has its corresponding anti-particle which has the same mass but opposite electric charge. A quark can join with an anti-quark to form an unstable particle known as a meson. All matter we see in our everyday lives is made up of the first generation of fermions. The proton is a baryon made up of two up quarks and a down quark while the neutron is made of two down quarks and one up quark.

Since leptons don't experience the strong nuclear force, the electron, muon and tau can only be affected by the weak, electromagnetic and gravitational forces, though gravity is by far the weakest. Since the neutrinos have no charge, they are unaffected by the electromagnetic force so can only interact by the weak force and gravity if they have mass. Neutrinos are now believed to have mass but, as with all subatomic particles, it is so low that the effect of gravity is negligible in most cases. This leaves only the weak force having a noticeable effect on the neutrinos.

1.3 Neutrino Interactions

Neutrinos can interact by the weak force in either charged or neutral current interactions. In the neutral current, a Z^0 boson is exchanged between particles

but the particle types remain the same. The neutral current interaction is depicted in a Feynman diagram in figure 1.1.



Figure 1.1: Feynman diagram showing a neutral current interaction between a neutrino and one of the quarks of a neutron.

In the charged current interaction, a W^+ or W^- is exchanged which changes the neutrino into its corresponding charged lepton. This is shown for both the neutrino and anti-neutrino in figure 1.2.



Figure 1.2: Feynman diagram showing a charged current interaction between a neutrino (or anti-neutrino) and a baryon through the exchange of a W boson.

The charged current interaction results in the production of a charged lepton as shown in equation 1.3 which interacts more readily so can be directly detected.

1.3.1 Lepton Family Conservation

When a neutrino interacts by the charged current interaction, quantities such as charge, lepton number and baryon number are always conserved. Lepton or baryon number is simply the number of leptons or baryons before and after the interaction. In other words, if a neutrino (a lepton) interacts with a neutron (a baryon) there will always be one lepton and one baryon left over as the interaction products. However, another number which is almost always conserved is lepton family number. This is a number indicating how many leptons of an individual family $(e, \mu \text{ or } \tau)$ are present before and after the interaction. When the lepton family number is conserved, this means that, if an electron neutrino goes through the interaction in figure 1.2, it always produces an electron rather than a muon or tau. This is summarized below for the charged current interactions:

$$\nu_l + N \to X + l^- \tag{1.4}$$

where N is a nucleon (proton or neutron), X is a baryon whose charge is one unit greater than N and l^- is a charged lepton $(e^-, \mu^- \text{ or } \tau^-)$ with ν_l as its corresponding neutrino. When the anti-neutrinos interact, they produce the positive anti-leptons:

$$\bar{\nu}_l + N \to X + l^+ \tag{1.5}$$

where, this time, X is a baryon whose charge is one unit less than N and l^+ is a charged anti-lepton $(e^+, \mu^+ \text{ or } \tau^+)$.

1.4 Summary

To summarise, the neutrino is electrically neutral and interacts weakly with matter. Although interactions with matter are rare, it is possible, given enough time, to detect the neutrino through these interactions. Since the neutrino is unaffected by magnetic fields, it travels in a straight line until the interaction with matter. This means, when the neutrino is detected, information relating to its origin is preserved.

Chapter 2 High Energy Astrophysics

High energy astrophysics is the study of the most violent and active regions of the universe, investigating the energetic particles produced by objects in these regions. These particles allow a better understanding of these objects which may not be provided by optical light and other low energy radiation.

2.1 Cosmic Rays

High energy astrophysics primarily deals with cosmic rays which are thought to mainly consist of high energy protons but can also be heavier atomic nuclei. These particles travel through space at near the speed of light and can have energies up to the order 10^{20} eV, which is 100 billion times its rest energy. The highest energy cosmic ray ever detected had energy 3.2×10^{20} eV [4].

2.1.1 Cosmic Ray Acceleration

When cosmic rays are produced in an active region such as a supernova remnant, they then have to reach the high energies at which we detect them. The favoured mechanism is known as Fermi acceleration, proposed by Enrico Fermi in 1949. In Fermi's original theory, (known as second order Fermi acceleration) the cosmic ray is accelerated through collisions with interstellar clouds. These clouds act as magnetic mirrors reflecting the particle back and forth [5]. If the cloud is moving towards the particle when the particle collides (a head-on) collision, the particle gains energy from the reflection. If the cloud is moving away from the particle, the particle loses energy in the collision. The individual energy gains from head-on collisions are of the order v/c where v is the speed of the cloud. When the energy gain is averaged, the head-on collisions are slightly favoured so the average is non-zero but, due to the cases where the particle loses energy, the resultant average is now of the order $(v/c)^2$. The average energy gain can be shown (see appendix A.1.1) to be

$$\left\langle \frac{\Delta E}{E} \right\rangle = \frac{4}{3} \left(\frac{v}{c} \right)^2 \tag{2.1}$$

where ΔE is the increase in energy, E is the current energy, and v is the speed of the cloud. Due to the slow speed of the cloud relative to the speed of light, the value of (v/c) will be very small which means, due to this value being squared, the energy gain will be extremely small.

A new mechanism was proposed known as first order Fermi acceleration. In this model, cosmic rays are accelerated across a strong shock front such as those produced in supernova remnants and in the jets of active galaxies. The particle interacts with magnetic irregularities on either side of the shock. If the material is moving very fast (close to the speed of light) which is likely in these active regions, the cosmic ray can gain a large amount of energy with each collision. The fractional energy gain depends on the angle of the collision relative to the motion of the material. The average gain over a round trip can be shown (see appendix A.1.2) to be

$$\left\langle \frac{\Delta E}{E} \right\rangle = \frac{4}{3} \left(\frac{v}{c} \right) \tag{2.2}$$

where v is the speed at which material is ejected from the supernova. Since the value (v/c) isn't raised to the second power as it was in second order acceleration, the average energy gain will not be as small.

With each collision and acceleration, there is a small probability the cosmic ray will escape the acceleration region and travel through space. Therefore, with more crossing across the shock, the individual particles gain energy but the number of particles gaining energy decreases. From this escape probability and the fractional energy gain, the differential energy spectrum can be shown (see appendix A.1.3) to be

$$N(E) \propto E^{-2} \tag{2.3}$$

which predicts the cosmic ray intensity emitted from the source.

2.1.2 Cosmic Ray Propagation Through Space

Magnetic Fields

When cosmic rays travel through space, they encounter galactic and intergalactic magnetic fields which can deflect charged particles. The amount of bending is related to a quantity known as the radius of gyration or gyro-radius [6] and is given by

$$r_g = \frac{p}{ZeB} \tag{2.4}$$

where p is the cosmic ray momentum, Z is the atomic number, e is the charge of an electron and B is the magnetic field strength. For high energy cosmic rays the energy is approximately E = pc so the radius is given by

$$r_g = \frac{E}{ZecB} \tag{2.5}$$

The amount of bending by a magnetic field is inversely proportional to this radius. If the radius is infinite, there is no bending. From the equation, for a given energy, we see that an iron nucleus will be deflected more by a given magnetic field since it has a higher atomic number than a proton which will reduce the gyro-radius. To do cosmic ray astronomy the direction of the cosmic ray's source must be accurately determined. If there is a great deal of bending, the direction of the cosmic ray when it arrives at Earth will be very different from the actual direction to the source.

For this reason, cosmic ray astronomy uses particles which have a large gyro-radius. This is achieved if the cosmic ray has very high energy. For ultra high energy protons (above about 10^{19} eV), the gyro-radius has values of the order 10kpc for galactic magnetic fields and 10Mpc for intergalactic magnetic fields. In these cases, the bending is small enough that particle astronomy can be attempted.

The GZK Effect

When cosmic rays travel through the universe, they can interact with photons from the cosmic microwave background (CMB) which is a low energy background radiation left over from the beginning of the universe. If a cosmic ray photon has energy greater than about 5×10^{19} eV, it can interact with a CMB photon which excites the proton to a Δ^+ baryon which can then decay back to a proton and neutral pion composed of a quark and its anti-quark or a neutron and a positive pion composed of an up quark and anti-down quark:

$$\gamma_{CMB} + p \to \Delta^+ \to p + \pi^0 \tag{2.6}$$

$$\gamma_{CMB} + p \to \Delta^+ \to n + \pi^+ \tag{2.7}$$

In this process, the proton loses energy. This process repeats until the interaction probability becomes insignificant again (i.e. when the proton energy drops below about 5×10^{19} eV.) This energy is known as the Greisen-Zatsepin-Kuz'min or GZK limit which puts an upper limit on the energy of cosmic rays from distant sources which can be detected on Earth [7].

2.2 Production of Secondary Particles

High energy cosmic rays can interact with surrounding material near where they were produced to create secondary particles which also travel out at high energy.

For example, a cosmic ray proton can interact with a background proton or neutron to produce charged or neutral pions [8]:

$$p + (p, n) \to p + (p, n) + (\pi^0, \pi^{\pm})$$
 (2.8)

where the π^- is composed of a down quark and anti-up quark.

Pions are also created when the proton strikes a background photon which elevates the proton to the excited Δ^+ state which decays to a proton or neutron:

$$p + \gamma \to \Delta^+ \to (p, n) + (\pi^0, \pi^+) \tag{2.9}$$

2.2.1 Neutral Messenger Particles

Since cosmic rays can be bent by the magnetic fields and their energy is capped at the GZK limit, this puts an overall limit on the resolution of the cosmic ray direction. It is, therefore, useful to search for neutral particles which originate from the same place as cosmic rays. Since these particles have no charge, they experience no magnetic force and won't be deflected by the magnetic fields allowing them to point directly back to their sources.

The pions produced from cosmic ray interactions are extremely unstable and decay in a matter of nanoseconds by various processes. The neutral pion decays when the quark and anti-quark annihilate to form two gamma rays:

$$\pi^0 \to \gamma + \gamma \tag{2.10}$$

The positive pion decays to an anti-muon and a muon neutrino to conserve lepton family number. Similarly, the negative pion decays to a muon and anti-muon neutrino:

$$\pi^+ \to \mu^+ + \nu_\mu \tag{2.11}$$

$$\pi^- \to \mu^- + \bar{\nu}_\mu \tag{2.12}$$

The muons are also unstable and decay after about 2 microseconds to an electron or positron and corresponding neutrinos:

$$\mu^+ \to e^+ + \nu_e + \bar{\nu}_\mu \tag{2.13}$$

$$\mu^- \to e^- + \bar{\nu}_e + \nu_\mu \tag{2.14}$$

It should be noted that, if the photon in equation 2.9 is a CMB photon, this is the GZK effect. The threshold energy of this proton-photon interaction implies the neutrino will have very high energy. These high energy neutrinos will produce a very large light signal in a detector so would require a large detector volume.

Neutrinos can also be produced if a neutron is created as shown in equation 2.8 or 2.9. A free neutron is unstable and will decay by β^- decay

as shown in equation 1.1 with a half life of about 10 minutes. This provides another source of production for anti-electron neutrinos. From these decay chains, neutrinos are produced in the same region as cosmic rays and gamma rays so all three can be used to investigate these regions. Gamma rays can also be produced from inverse Compton scattering of electrons. This production mechanism of gamma rays is known as the "leptonic" source as opposed to the "hadronic" source where they are produced by pions.

Both gamma rays and neutrinos have the advantage over cosmic rays that they are unaffected by magnetic fields and can point back to their sources, however, gamma rays have their limitations. Since gamma rays interact more readily with matter and radiation, over cosmological distances, they will be absorbed by the intergalactic medium. Neutrinos, by contrast, have a very low interaction probability. This means neutrinos have the added advantage of pointing back over cosmological distances to sources in the early universe. The propagation of cosmic rays, gamma rays and neutrinos is summarised in figure 2.1.



Figure 2.1: Simple example of particles travelling to Earth. In this example, the cosmic ray path is significantly curved and the gamma ray is eventually absorbed. In this example, shown over a cosmological distance ($\sim 1000 \text{Mpc}$), only the neutrino arrives at Earth and preserves information about the source position.

Neutrino Oscillations

As neutrinos travel through space, it is thought that they "oscillate" between their three flavours: electron, muon and tau. This means that, given enough time, a neutrino of a particular flavour can change into a neutrino of a different flavour. This is due to the mass eigenstates of the neutrinos being different from the flavour eigenstates. When a neutrino is created, it has mass but the mass is so low that it can't be resolved. For this reason, the mass of a single neutrino can be thought of as a superposition of the mass states of the three flavours as shown below:

$$|\nu_{\alpha}\rangle = \sum_{i=1}^{3} U_{\alpha i}^{*} |\nu_{i}\rangle \tag{2.15}$$

and similarly for the anti-neutrinos:

$$|\bar{\nu}_{\alpha}\rangle = \sum_{i=1}^{3} U_{\alpha i} |\bar{\nu}_{i}\rangle \tag{2.16}$$

where $|\nu_{\alpha}\rangle$ is a neutrino flavour eigenstate for a neutrino of definite flavour $(e, \mu \text{ or } \tau)$ and $|\nu_i\rangle$ is a neutrino of definite mass. The coefficient U is a matrix and U^* is the transpose of its complex conjugate [9]. This shows that, while a neutrino can have definite flavour, this does not mean it has definite mass. Over time, the mass eigenstates vary according to the time dependent Schrödinger equation to give a plane wave solution of the form

$$|\nu_i(t)\rangle = e^{-i(E_i t - \vec{p}_i \cdot \vec{x}_i)} |\nu_i(0)\rangle$$
(2.17)

where E_i is the energy of the eigenstate, $\vec{p_i}$ is its momentum and $\vec{x_i}$ is the displacement from the starting point. If the neutrinos are travelling close to the speed of light, this can be simplified to show the oscillation for a neutrino having travelled a distance $L \sim ct$:

$$|\nu_i(L)\rangle = e^{-\frac{im_i^2 L}{2E}} |\nu_i(0)\rangle \tag{2.18}$$

Since each mass eigenstate has different mass m_i , these wave functions propagate at different speeds. As the flavour eigenstates are superpositions of the mass eigenstates, the reverse is also true that the mass eigenstates are superpositions of the flavour eigenstates. Since the mass eigenstates move at different speeds, it is possible for these "waves" to constructively interfere at certain points giving a high probability of observing a specific neutrino flavour which may not be the same as the original neutrino flavour. These probabilities are summarised in the plot in figure 2.2.



Figure 2.2: Plot showing probabilities of observing each flavour given a starting electron neutrino [10].

Given enough distance to travel, the relative abundances of ν_e , ν_{μ} and ν_{τ} arriving at Earth could be quite different from their values at the source. These oscillations are thought to be the most likely source of ν_{τ} .

2.3 Particle Detection On Earth

2.3.1 Cherenkov Radiation

High energy cosmic rays are detected on Earth through the secondary particles they create when they interact with the atmosphere. These interactions produce many charged particles with extremely high energy which means they are travelling very close to the speed of light c. However, when they move through the atmosphere, or water, they are moving through a material where the speed of light is reduced by a factor n, the refractive index. The refractive index is 1.003 for air and 1.33 for water which means, in air, light travels at 99.7% its vacuum speed and in water, it travels 75% of the vacuum speed.

When a charged particle moves through a polarisable material, the particle's electric field causes the molecules in the medium to become polarised. If the particle is travelling slowly through the medium, this disturbance relaxes back into equilibrium as the particle passes. However, if the particle is travelling faster than the response speed of the medium given by c/n, there is not enough time for the molecules to become unpolarised as the particle passes through. This causes the disturbance to remain in the medium and the energy stored radiates out as a coherent shock wave [11], as shown in figure 2.3.



Figure 2.3: Diagram showing light spheres produced by a charged particle, forming a Cherenkov cone. In time t, the light travels a distance of ct/n while the particle travels a larger distance of vt.

This effect is analogous to a sonic boom generated by objects moving faster than sound. For particles travelling at approximately c, the critical angle θ_c in figure 2.3 is given by

$$\cos\theta_c = \frac{1}{n} \tag{2.19}$$

This means, for a denser medium, θ_c has a higher value meaning the cone is more compact and directional.

2.3.2 Detection Methods

Air Showers

Cosmic rays are detected on Earth through the extensive air showers produced when they enter the atmosphere [12]. When a cosmic ray (e.g. a proton) collides with an air molecule, it loses energy and produces secondary particles such as pions which decay to muons through the decay equations outlined above. Due to the high energy of the incident particle, numerous other particles can be created from this energy. By the time they reach the ground, the shower can contain up to 10^{10} particles which are spread out over an area of the order 10 km² for an incident particle of 10^{19} eV. A typical example of a cosmic ray air shower is shown in figure 2.4. The air shower is divided into muonic, hadronic and electromagnetic components.

Cosmic ray detectors such as the Pierre Auger Observatory in Argentina use a large array of water tanks which detect the Cherenkov light produced by these particles as they pass through, providing information about the direction and energy of the shower. Pierre Auger also has fluorescence detectors which detect



Figure 2.4: Diagram showing muonic, hadronic and electromagnetic components of an air shower caused by a primary hadron [13].

the optical light emitted in the air when a particle excites an air molecule and the molecule drops back to its ground state. This allows a study of the shower as it passes through the atmosphere before reaching the Cherenkov tanks.

Gamma rays produce similar air showers when they interact with the atmosphere. The High Energy Stereoscopic System (H.E.S.S.) observatory in Namibia has an array of telescopes which detect the Cherenkov light emitted by the gamma ray air shower as it passes through the atmosphere. Using the signals received at each telescope, the three dimensional structure and direction of the air shower can be determined.

Gamma ray air showers are distinguished from cosmic ray air showers since the shower is dominated by its electromagnetic component. The gamma ray produces an electron-positron pair which produce more gamma rays through bremsstrahlung which then produce more electron-positron pairs. This type of shower has fewer fluctuations and is typically more predictable than the cosmic ray air shower.

Neutrino Detection

Neutrino detection relies on rare neutrino interactions with matter. To ensure a detection of a neutrino, a detector must be very large such that the probability of a neutrino interacting in the detector's volume is increased. A large volume detector is also necessary to observe high energy neutrinos. The interaction products of these neutrinos typically travel large distances within the detector so

a large volume is required to see the entire signal. Neutrino detectors are mainly built underwater or in ice as they are dense but transparent media and light travels only 75% its vacuum speed allowing the detection of Cherenkov light. However, since neutrinos are neutrally charged, they can't polarise the material they move through and can't produce Cherenkov light on their own.

If a neutrino interacts with a proton or neutron in a charged current interaction as shown in figure 1.2, it will produce its corresponding charged partner $(e^-, \mu^- \text{ or } \tau^-)$. Since this new particle has charge and is moving very close to the vacuum speed of light, it can polarise the medium and produce Cherenkov light. The pattern of Cherenkov light seen in the detector can then be used to determine the path of the charged lepton. If these interactions happen at high energy, there will be almost no change in the particle trajectory after the interaction. Therefore, the direction of the charged lepton reveals the direction of the original neutrino.

2.4 Summary

High energy astrophysics has allowed studies of extremely active regions in the universe. However, there are limitations in using cosmic rays and gamma rays to probe these regions. Cosmic rays are charged particles and, therefore, have their paths altered by magnetic fields meaning the cosmic ray detected at Earth may not point back to its source. Gamma rays travel in straight lines but are more readily absorbed by the intergalactic medium making their sources difficult to observe over cosmological distances.

Neutrinos travel in straight lines and pass through most matter unaffected making them ideal particles for pointing back to their sources over large distances. However, to achieve neutrino astronomy, the neutrino must interact in the detector volume. Since neutrinos only interact weakly with matter, they are very difficult to detect. To maximise the probability of a neutrino within a detector, the neutrino detector must have a large volume.

Chapter 3

The IceCube Neutrino Observatory

3.1 Detector Layout

The IceCube detector is a neutrino observatory with a total volume of 1 cubic kilometre located in the deep ice at the geographic South Pole. The ice provides a good target for neutrino interactions and the optical properties allow the detection of Cherenkov light. The detector is an array of 5,160 digital optical modules (DOMs) arranged in 86 strings at a depth of 1.5-2.5 kilometres below the surface where the ice extremely transparent (absorption length ~ 200 metres). On each string, the DOMs are separated by 17 metres and the strings are separated by 125 metres. The DOMs collect the Cherenkov light and, in response, generate an electric current which is measured to calculate the photon intensity [14]. The modern IceCube detector is constructed around its predecessor the Antarctic Muon and Neutrino Detection Array (AMANDA) which was constructed in the late 1990s. AMANDA is located in the corner of the IceCube array and was incorporated into the array in 2007. In 2009, AMANDA was decommissioned and its role is now taken by the Deep Core array [15] which sits at the bottom of the detector and has its DOMs closer together allowing greater sensitivity to low energy events [16]. This layout is shown in the diagram in figure 3.1.

On the surface of the ice there is another detector known as IceTop made up of tanks of ice [17] which also detect Cherenkov light. This provides a better understanding of the origin of the events seen in the detector. If a muon is seen beginning in the "in-ice" detector it has obviously formed from a neutrino interacting in the ice. IceTop, similar to Pierre Auger, can detect cosmic ray air showers. If a muon enters the IceCube detector from the outside and, at the same time, a cosmic ray air shower is detected by IceTop, this implies that the muon was part of the air shower and not the result of a neutrino interaction in the ice.



Figure 3.1: IceCube Detector shown to scale with the Eiffel tower [18].

3.2 Light Signal in DOMs

Due to varying distances between the DOMs and the particle being detected, each DOM in the detector sees a slightly different amount of light. The amount of Cherenkov light seen by certain DOMs and the time the light is detected can reveal information about where the charged lepton was created and its path through the detector. This can then be used to determine the energy and direction of the original neutrino.

3.2.1 Detection of Cherenkov Light

The Cherenkov photons from the charged lepton are detected by a photo-multiplier tube (PMT) inside the DOM. The PMTs used in the IceCube DOMs consist of a negatively charged photo-cathode, a series of 10 dynodes and a positively charged anode. When a Cherenkov photon strikes the photo-cathode, its energy is given to an electron in the photo-cathode material. This "photo-electron" (PE) then accelerates through the PMT towards the first dynode which is an electrode kept at higher positive potential than the cathode. When the electron strikes the dynode, it has gained energy which can then cause multiple electrons to be ejected from the dynode when it collides.



Figure 3.2: Schematic Diagram of an IceCube DOM [19].

These electrons then accelerate to the second dynode which is at a higher potential. This process continues with more electrons being produced along the way until a large number of electrons arrive at the positive anode. The number of electrons arriving at the anode depends on the potential difference between the cathode and anode. For IceCube's PMTs, the anode to cathode voltage is about 1300 V which causes a gain of about 10^7 [20]. This means, for a single incident photo-electron, 10^7 electrons can arrive at the anode, amounting to a strong signal.

The electrons arriving at the anode are measured as an electric current which is proportional to the intensity of the Cherenkov light. The time distribution of this intensity is known as the DOM's waveform. Using the waveforms from all DOMs, the properties of the particle such as direction, position and energy can be determined. This is known as event reconstruction and is typically achieved numerically by predicting the waveforms for given particle parameters and maximising a likelihood function until the predictions closely match the observed waveform. This will be covered in more detail in chapter 4.

However, before the event can be properly analysed from its waveforms, the signal needs to be properly calibrated to determine which parts are due to real photons produced by the event and which are simply characteristics produced in the DOM. When the current is generated, the voltage increase is registered as a digital signal. The digital signal is sampled by four channels: three for the analogue transient waveform digitiser (ATWD) and one for the flash analogue digital converter (FADC.) The ATWD has a sample rate of 300 MHz giving a

sampling interval of 3.3 ns. The ATWD takes 128 measurements giving a total sampling time of about 422 ns. Each ATWD has a ten bit resolution meaning each returns the digital voltage as an integer between 0 and $2^{10}=1024$. For a faint signal, only one or two channels will be active. If either channel's maximum signal value exceeds 768, the third channel is activated as well [21].

The ATWD is typically used for the detection of Cherenkov light which is the first to arrive. This light could originate when the muon is relatively close to the DOM or it could be light emitted earlier in the track and has been scattered. For later detection of Cherenkov light, the FADC, which has a sample rate of 40 MHz or a sampling interval of 25 ns [22], is used. This digitizer takes 256 measurements giving it a longer measuring time of 6.4μ s so it can be used for the broader waveforms. The signal from the four channels is known as the raw data.

The raw data is fed into a wave calibrator which converts the ADC count integers into an actual signal in units of mV. The ATWD and FADC waveforms each have a non-zero baseline which is known for each channel and subtracted by the wave calibrator. Each count in the ADC corresponds to a known voltage gain. The baseline subtracted counts are multiplied by this gain to obtain the corresponding voltage.

The timing of this new waveform is also corrected to allow for known delays due to the signal propagating through the DOM. The waveforms are then corrected for the "droop" in the DOM's transformer [23]. This is due to the transformer acting as a high pass filter which works to attenuate the waveform after the first peak causing the voltage to appear lower than the true value. This is compensated by adding the expected reaction voltage for each point along the waveform. Lastly, the three ATWD channels are combined to give a single ATWD waveform for each DOM. An example of an FADC waveform before and after calibration is shown in figure 3.3.

After this calibrated waveform is produced, it is analysed by a feature extractor which finds the strong peaks, known as pulses, rising above the background noise. The feature extractor will check the values of the waveform for each time bin to find a local maximum above the noise threshold. The slope leading up to this maximum is used to extrapolate a line down to the baseline. The time this line coincides with the baseline is considered to be the pulse's start time. The values in each time bin are summed up until the values drop below the noise threshold or start to rise again after the maximum which defines the beginning of a new pulse [21].

This sum can then determine the number of photo-electrons present in the pulse which is known as the pulse's charge. The pulse width is simply defined as



Figure 3.3: Example plots of an FADC waveform before and after calibration plotted against absolute time. Note the ADC Count has now been converted into a voltage and the timing of the waveform has now been corrected for the propagation delay through the PMT.

the time difference between the beginning and end of the sum. These pulse series are then used by reconstruction algorithms to determine the parameters of the particle.

3.2.2 Signals of Different Flavours

When a neutrino of a particular flavour $(\nu_e, \nu_\mu, \nu_\tau)$ interacts, it will produce a charged lepton of the same flavour as shown in section 1.3.1. Electrons, muons and tau leptons all produce Cherenkov light but the light pattern they produce will depend on the specific flavour. If the incident particle is an electron neutrino (or anti-neutrino) and the charged particle is an electron (or positron) it will interact very readily with the surrounding ice as it has, by far, the lowest mass. This causes the electron to lose its energy to an electromagnetic cascade in the ice, which is seen by the DOMs as a sphere of light expanding outward [24]. As the cascade is made up of particles repeatedly scattering through the ice, it is difficult to determine the direction of the cascade without exact knowledge of the ice properties. These cascades typically have angular resolutions of about 10°-20°.

If the particle is a muon it has a much higher mass and travels through the detector more easily producing Cherenkov light along the way. This produces a track-like structure in the detector which makes it easy to see the direction the particle is travelling. These events can have a reconstruction resolution better





than one degree.

Theoretically, IceCube can detect tau leptons as well which would have been produced by neutrino oscillations. If a tau lepton were detected, the signal should look like a "double bang." This is because the tau is even less stable than the muon and, while passing through the detector, will most likely decay to an electron:

$$\tau^- \to e^- + \bar{\nu}_e + \nu_\tau \tag{3.1}$$

It is theorized that the detector will see a burst of light when the tau is created, a short, faint track as it moves, and a second burst of light caused by the electron. For this double bang to be observed in the detector, the tau must have energy of about 2-20 PeV. The visualisations of the signals from all three flavours are shown in figure 3.4.

3.3 Optical Properties of the Ice

In order to accurately determine the parameters of an event given its light signal, an understanding of how the light propagates in the ice is essential. The IceCube detector was built at a depth sufficient for the ice to be transparent with an absorption length of the order 200 m allowing light to easily propagate. However, even at this depth, light can still be absorbed and scattered through the ice as it travels hundreds of metres to its target DOM. This means, to perform reconstruction, knowledge of how the light will scatter in the ice is essential for predicting waveforms.



Figure 3.5: Plots of the scattering and absorption coefficients as functions of depth and wavelength of light [26]. As shown, blue light of wavelength 400 nm is scattered more than other visible wavelengths with a coefficient of about 0.03 m^{-1} while it is absorbed less with a coefficient of less than 0.005 m^{-1} .

These ice properties were determined by AMANDA and IceCube by creating artificial light sources with DOMs and measuring the scattered light in other DOMs [26].

The distinct peaks in the absorption and scattering shown in figure 3.5 are due to dust layers in the ice. These dust layers are at depths which coincide with IceCube so must be taken into account in any ice model. Cherenkov light is typically in the blue to ultra-violet (short wavelength) region of the spectrum which means that absorption is minimised with respect to longer wavelength red light as shown in figure 3.5. A wavelength around 400 nm has the smallest scattering coefficient around 0.005 m^{-1} . while 600 nm light has a coefficient of 0.1 m^{-1} . However, the blue light has a scattering coefficient of 0.03 m^{-1} while red light has a coefficient of only 0.02 m^{-1} indicating blue light is scattered more than red light.

The ice models used for simulation and reconstruction are updated from time to time as knowledge of the optical properties improves. Simulation and reconstruction use a programming package which incorporates everything that is known about the ice in the detector to accurately predict the paths of all Cherenkov photons as they scatter through the ice. Ice models known as Spice 1 and Spice 2 (Spice is an abbreviation for South Pole Ice) were the earlier ice models while the more recent models of the ice are called Spice Mie and Spice Lea which incorporate a more up to date model of the Mie scattering of light [27, 28].

3.4 Discovery of Astrophysical Neutrinos

Since IceCube directly detects charged leptons rather than neutrinos, it is possible that most of the particles detected actually arise from cosmic ray air showers in the atmosphere rather than neutrino interactions in the ice. One of the advantages of neutrino astronomy is that neutrinos can pass through the whole Earth before reaching the detector. This means that, even though IceCube is at the South Pole, neutrinos from the northern hemisphere can be seen as they come up through the Earth. These are known as "up-going" events whereas the events from the southern hemisphere are "down-going." The up-going events have the extra advantage in that a muon from a cosmic ray air shower on the other side of the Earth won't survive the trip through the planet so, if we see an up-going muon, it can only have been produced by a neutrino. Due to this fact, the detector was originally used only for up-going events [29].

To be confident about down-going events coming from neutrinos, there is a layer at the boundary of the detector called the veto layer which excludes muon events which pass through it from outside. This layer is shown in figure 3.6. If, for example, a muon forms outside the detector before entering the volume, it is unknown whether it came from a neutrino or directly from the atmosphere. If the muon forms in the detector volume, the beginning of the Cherenkov light pattern can be seen which indicates the position of the neutrino interaction. These events where the interaction happens inside the detector are known as starting events. Muon neutrinos from cosmic ray air showers are created by pion decay so will be accompanied by a muon as they enter the detector.

If a neutrino interacts in the detector but is accompanied by a muon which has passed through the veto layer, this neutrino is also vetoed as it is probably atmospheric in origin. For this reason, only the high energy neutrinos are selected as the accompanying muon tracks would be easily seen if present.

Algorithms have been developed which calculate the expected atmospheric



Figure 3.6: Diagram of detector from top and side showing veto region. The 80 m veto layer in the middle is due to the dust layer in the ice. Due to the large amount of absorption in this layer, it is possible a particle entered this layer horizontally without being flagged at the edges [30].

muon and neutrino fluxes at the detector [31]. These predicted fluxes were compared to the observed fluxes as shown in figure 3.7. After vetoing many neutrinos and muons at high energies, the expected atmospheric flux at these energies drops off rapidly as shown in figure 3.7(a). It is also clear that there have been many more events detected at these energies which haven't been vetoed providing the first clear evidence of extra-terrestrial neutrinos.

Using this background simulation method, IceCube confirmed the detection of 28 high energy starting event (HESE) neutrinos from astrophysical sources from two years of data in 2013. This number was increased to 37 after a third year of data was collected. Although neutrinos travel in straight lines, there is still an intrinsic uncertainty in the arrival directions when they arrive in the detector. This uncertainty is well below one degree for tracks and about 10° for cascades as stated in section 3.2.2. Due to this uncertainty, it is possible that two observed neutrinos coming in slightly different directions originated from the same point if their direction uncertainties overlap. Using this, numerous studies have been done to determine the probability that clustered neutrinos came from the same point, and further, the strength or intensity of the possible source. Using these probabilities, a sky map can be created showing possible clusters of neutrino events represented by a test statistic (TS) which indicates the likelihood that the neutrinos came from the same source.



(a) Neutrino and Muon fluxes plotted against deposited energy.

(b) Neutrino and Muon fluxes against arrival direction.

Figure 3.7: Plots of the predicted neutrino and muon fluxes for different energies and direction compared to actual data [30].



Figure 3.8: Sky Map showing the 37 events from the three year sample in galactic coordinates. The shading represents the test statistic [32].

The sky-map of the test statistic shows what appears to be a significant clustering of neutrino events near the galactic centre which would be a good candidate source of high energy particles. However, this hypothesis was tested by generating a series of random skies to see how clustered the neutrinos could have been by random chance. It was found that 7% of the time, the random clustering is even greater than that shown in figure 3.8. This value is not considered a statistically significant result so this is not considered a point source. So far, none of the observed clustering has been found to be significant enough to indicate any common source for any events.

Chapter 4 Event Reconstruction

Reconstruction is the process of taking raw data from the light signal detected by the DOMs and, from this signal, determining the physical properties of the detected particle such as its energy, direction and position. This is done by computer algorithms which use the parameters of the detector and its DOMs, as well as the optical properties of the ice to determine the most likely particle properties which would produce this signal.

The most basic reconstruction algorithms, such as line-fit, determine the parameters of the event analytically by minimising for the distance between the track and the illuminated DOMs. This provides a first guess of the particle parameters. More advanced algorithms, such as MPE and millipede, calculate these parameters numerically by maximising a likelihood function. All events in the IceCube detector are reconstructed relative to the IceCube coordinate system.

4.1 IceCube Coordinate System

This coordinate system has its origin near the centre of the detector about 2 km below the surface. The y-axis points north of the detector along the prime meridian towards Greenwich. The coordinate system is right handed so the x-axis points 90° clockwise from the y-axis which is along the 90° E line of longitude. The z-axis points straight up from the centre.

The direction of the detected particle is given by its zenith θ and azimuth ϕ . The zenith is the angle between where the particle is coming from and the overhead direction. In other words, an event that is heading straight down will have a zenith of 0 while a horizontal event will have a zenith of 90° ($\pi/2$ radians) and an event travelling straight up will have a zenith of 180° (π radians.) Since the detector is at the South Pole, zenith (in radians) is related to declination δ or

celestial latitude by a simple relation.

$$\delta = \theta - \frac{\pi}{2} \tag{4.1}$$

Azimuth is the horizontal angle and is defined as being zero for the positive x-axis and increases in an anticlockwise direction such that it is 90° ($\pi/2$ radians) for positive y, 180° (π radians) for negative x and 270° ($3\pi/2$ radians) for negative y. Unlike zenith, azimuth has no constant relation to a celestial coordinate since azimuth is fixed relative to the detector and right ascension or celestial longitude completes a revolution every 24 hours. The relation between azimuth and right ascension, therefore, depends on the time of measurement. The coordinate system is illustrated in figure 4.1.



Figure 4.1: Diagram representing coordinates for a down-going event.

4.1.1 Definition of Vertex

In addition to zenith and azimuth, each detected particle has a vertex associated with it which is a point (x,y,z) in space which lies on the particle motion vector. This position defines where the particle track passes in the detector. For example, two muons could be travelling in the same direction which means they have the same (θ, ϕ) coordinates but pass through the detector on opposite sides. This means these muon tracks will have different vertex positions.

In addition to energy and direction, the vertex position is also calculated in the reconstruction. Since the vertex merely defines where the track passes, its position along the track itself can be rather arbitrary. Simulations of muons which store the particles in a Monte Carlo Tree (see appendix B.2) put the muon track vertex at the location of the neutrino interaction. In this case, the vertex could be outside the detector. Reconstruction algorithms generally result in a vertex which is inside the detector volume.

As shown in the visualisation in figure 3.4, the light signal seen in the detector gives an indication of where the particle passes through the detector. Analytic algorithms such as line-fit place the vertex near the beginning of this light pattern by minimising the distance to the illuminated DOMs. Numeric algorithms use this vertex and direction as an initial guess and shift it to maximise the likelihood function. While the vertex could theoretically now be placed at any position along the track and produce the same result, the step size used in the algorithm is generally too small to produce a substantial vertex shift along the track.

4.2 Line-Fit

Most reconstruction algorithms find the particle's direction and position by numerical means. That is, an initial guess particle trajectory is used and its expected signal is compared to the observed signal. The particle parameters are changed until the expected signal matches what is seen in the detector. However, due to the complexity of the problem being solved, it is necessary to have initial guess parameters which are already fairly close to the true parameters (i.e. $< 5^{\circ}$.) The initial guess is generated by analytic means to get an approximate result without being too computationally intensive.

Line-fit is one of the more basic reconstruction algorithms used to provide a first guess for more sophisticated reconstructions. This algorithm ignores the geometry of the Cherenkov cone and simply treats the particle as a straight line in 3-dimensional space. If the particle is at a position $\vec{r_0}$ at time 0 and its velocity is equal to \vec{v} , at time t_i , its position is given by

$$\vec{r} \sim \vec{r_0} + \vec{v}t_i \tag{4.2}$$

If all the DOMs which have registered photon hits have positions given by $\vec{r_i}$ and they receive light at time t_i , the distance between the DOMs and the position of the particle at this time is minimised by minimising the sum of the distances squared given by

$$\chi^2 = \sum_{i=1}^{N_{\text{hits}}} |\vec{r_i} - \vec{r_0} - \vec{v}t_i|^2$$
(4.3)

This can be minimised analytically by differentiating with respect to the free variables \vec{r}_0 and \vec{v} . The minimum occurs for \vec{r}_0 and \vec{v} given by

$$\vec{r}_0 = \langle \vec{r}_i \rangle - \vec{v} \langle t_i \rangle \tag{4.4}$$
$$\vec{v} = \frac{\langle \vec{r}_i t_i \rangle - \langle \vec{r}_i \rangle \langle t_i \rangle}{\langle t_i^2 \rangle - \langle t_i \rangle^2} \tag{4.5}$$

where, for a parameter x_i , $\langle x_i \rangle$ defines the average x_i over all i [33].

From these variables, the particle track parameters can be found. The track vertex position is just given by \vec{r} and its arrival direction can be found by the unit vector in the direction of $-\vec{v}$. The magnitude of \vec{v} is simply taken to be c. However, this method ignores the geometry of the Cherenkov light and any optical properties of the ice. In reality the situation is significantly more complex. This initial guess must now be used by numerical algorithms which take into account the physics of the particle light emission and propagation to the DOMs.

4.3 SPE and MPE

Using the Line-fit solution as a starting track, numerous algorithms exist which take this track and alter its parameters until the most likely set of parameters is found. The reconstruction is achieved by calculating a quantity known as the likelihood for a certain set of track parameters and finding the track for which this quantity is maximised. Likelihood is essentially the probability that the observed signal was produced by given input parameters. In general, likelihood is best described as the product given below:

$$\mathcal{L}(\vec{x}|\vec{a}) = \prod_{i} p(x_i|\vec{a}) \tag{4.6}$$

where $p(x_i|\vec{a})$ is the probability density function for getting the measured charge distribution in time (waveforms) x_i for input track parameters \vec{a} . In the case of muon reconstruction, the input parameters are quantities such as the starting energy, energy losses, direction, time and vertex position of the track.

Early numerical reconstruction algorithms included the single photo-electron (SPE) and multi photo-electron (MPE) methods. These methods take into account the geometry of the Cherenkov cone emitted by the muon and incorporate some knowledge about the optical properties of the ice. As shown in figure 4.2, the muon is travelling through the detector with energy E_0 and velocity unit vector given by \hat{p} . The muon is said to have position $\vec{r_0}$ at time t_0 . The Cherenkov cone is emitted with critical angle θ_c where $\cos\theta_c \sim 1/n$. The distance d is defined as the closest distance between the muon track and the DOM at position $\vec{r_i}$.

To simplify the likelihood calculation for given input parameters $a = (\vec{r}_0, t_0, E_0, \hat{p})$, a new variable is defined known as the time residual $t_{\rm res}$. The time residual is defined as the difference between the time the DOM registers the light and the time the light would have arrived if it had travelled



Figure 4.2: Diagram of muon track and Cherenkov cone [33].

in a straight line without being scattered. Since light is known to scatter off the ice and take an indirect path to the DOM, this time residual will be non-zero in most cases.

The time the light reaches the DOM with no scattering, known as the "geometric time," is determined using the geometry of the Cherenkov cone shown in figure 4.2. The earliest light arriving at the DOM is emitted at the Cherenkov angle θ_c to the muon path. As shown in figure 4.2, the muon travels from position $\vec{r_0}$ at a speed approximately equal to c. The light travels through the ice at a speed c/n where n is the refractive index. By adding the two paths shown in the diagram together, the geometric time can be shown (see appendix A.2) to be

$$t_{\rm geo} = t_0 + \frac{\hat{p} \cdot (\vec{r_i} - \vec{r_0}) + d\tan\theta_c}{c}$$
(4.7)

where \hat{p} is the unit vector in the direction of the muon path. The time residual is given by $t_{\rm res} = t_{\rm hit} - t_{\rm geo}$ where $t_{\rm hit}$ is the time recorded by the DOM after the calibration discussed in section 3.2.1.

Given a time residual for a DOM and knowledge of the optical properties of the ice, it is possible to calculate $p(t_{\text{res},i}|\vec{a})$ which is the probability density function of obtaining this time residual for the *i*-th DOM given guess parameters \vec{a} . These probability density functions are then multiplied together to obtain the likelihood. The probability density $p(t_{\text{res},i}|\vec{a})$ across the waveform is estimated by a "Pandel" function which is a modified gamma distribution detailed in [33]. In the single photo-electron (SPE) method, the likelihood function is simplified such that only one photon is considered for each DOM without altering the Pandel function. In this case, the likelihood function is simply

$$\mathcal{L} = \prod_{i=1}^{n_{\text{hits}}} p_1(t_{\text{res},i} | \vec{a})$$
(4.8)

where n_{hits} is the number of DOMs which see light and $p_1(t_{\text{res},i}|\vec{a})$ is the value of the Pandel function for a single photon in the *i*-th DOM.

This likelihood is calculated for the initial guess parameters given by line-fit or another basic reconstruction algorithm. The particle parameters are then varied and the likelihood is recalculated. This repeats until the parameters which give the maximum are found. The SPE simply uses the value of the unaltered Pandel function for a single photon. A slightly more sophisticated approach is to include the information that only one photon is being considered. This method, known as the multi photo-electron (MPE) method, calculates the probability density of a photon in the *i*-th DOM having $t_{\rm res}$ with the knowledge that it is the first photon in a sample of n_i photons.

The likelihood then becomes

$$\mathcal{L} = \prod_{i=1}^{n_{\text{hits}}} n_i p(t_{\text{res},i} | \vec{a}) (1 - P(t_{\text{res},i}))^{n_i - 1}.$$
(4.9)

where n_i is the number of photons detected for the *i*-th DOM. The probability $P_1(t_{\text{res}})$ is the cumulative distribution of $p(t_{\text{res}}|\vec{a})$ given by [33]:

$$P(t_{\rm res}) = 1 - \int_{t_{\rm res}}^{\infty} p(t|\vec{a}) dt.$$

$$(4.10)$$

Since the MPE method is a more general case as it works for multiple photon counts, this method is assumed to produce a more accurate reconstruction than SPE. However, since both methods are only considering the first detected photon, these methods are generally used as intermediate steps to provide a first guess for a more sophisticated algorithm. They are generally not considered final results.

4.4 Millipede

The millipede algorithm is a method of reconstruction which is primarily used to reconstruct the energies of muons as they pass through the detector [24]. The name "millipede" comes from the method, where the muon track in the detector is divided into segments. In each segment, a certain amount of energy is lost which can create photon later registered by the DOMs. The intensity of light in each DOM can then be thought of as a superposition of all the light emitted at each segment of the track. This is illustrated in figure 4.3.



Figure 4.3: Diagram of a track divided into n losses each emitting light received by m DOMs. Here, E_n is the energy lost at the *n*-th segment and N_k is the number of photons received by the k-th DOM.

The amount of light seen by the k-th DOM is dependent on the ice properties between the DOM and the energy loss segment as indicated earlier. Therefore, without background noise, the number of photons at the k-th DOM is given by

$$N_k = \sum_{i=1}^n \Lambda_i(\vec{r_k}) E_i \tag{4.11}$$

where Λ_i is a coefficient describing how the intensity will change for a given energy loss E_i and a distance vector \vec{r}_k between the loss and the k-th DOM. This parameter incorporates what we know about the optical properties of the ice at this specific location. With background noise ρ , the equation becomes

$$N_k = \rho + \sum_{i=1}^n \Lambda_i(\vec{r_k}) E_i \tag{4.12}$$

This relation can be summarised in the following matrix equation [34]

$$\begin{pmatrix} N_1 - \rho_1 \\ \vdots \\ \vdots \\ N_m - \rho_m \end{pmatrix} = \begin{pmatrix} \Lambda_1(\vec{r_1}) & \vdots & \vdots & \Lambda_n(\vec{r_1}) \\ \vdots & \vdots & \vdots \\ \Lambda_1(\vec{r_m}) & \vdots & \vdots & \Lambda_n(\vec{r_m}) \end{pmatrix} \begin{pmatrix} E_1 \\ \vdots \\ \vdots \\ E_n \end{pmatrix}$$
(4.13)

which can also be given by

$$\vec{N} - \vec{\rho} = \mathbf{\Lambda} \cdot \vec{E}. \tag{4.14}$$

In reality, the DOM signals are also distributed in time. This means that a single DOM can contribute more than one element to \vec{N} where N_k is the photon count in an individual time bin. This means the value of m is now greater than the number of DOMs which adds more variables to the equation.

4.4.1 Millipede Time Binning

To reconstruct an event with millipede, the event waveforms are read into the algorithm as a pulse series map. This pulse series map contains information extracted from the calibrated waveforms similar to the example in figure 3.3 for all DOMs. Each pulse has a given charge which is the number of photo-electrons (PEs) contained in the pulse as well as a start time and a width.



Figure 4.4: Pulse series extracted from the calibrated waveform shown in figure 3.3.

The pulse series map is initially divided up into pulse series for each hit DOM. Millipede needs to take the pulse series map and combine it into a single vector. The width of millipede's time bins are determined by how many photons are allowed in each bin given by the input variable "**PhotonsPerBin**". However, before this step, the waveform is pre-binned for each DOM based on the widths and times of the pulses. This initial binning method creates bins spanning the event time window alternating between a bin spanning a pulse and a bin spanning a gap between pulses. For a bin spanning a pulse, the number of photons in this bin is simply the number of photo-electrons in the pulse. For a bin spanning a gap, the number of photons is 0.



Figure 4.5: Simple example of a pulse series containing two pulses divided into five bins. The widths w_2 and w_4 are the time widths of the first and second pulse respectively. The widths w_1 and w_5 will typically be very large as the time window is comparatively much larger than the pulse series.

As shown in figure 4.5, the bins span the entire time window and are divided by the pulses. Millipede then re-bins the waveform by combining these initial bins until either the photon count exceeds the value of **PhotonsPerBin** or the bin width exceeds 200 ns. However, if the initial bin already has a width greater than 200 ns, this bin stays as it is. This width limit of 200 ns ensures that, for dim signals, a time distribution of photons can still be produced. For example, if the number of photons never reaches **PhotonsPerBin**, the result will be a single bin containing all light surrounded by two empty bins [35]. In this case, information about the relative photon arrival times is lost. The value of 200 ns is optimised to minimise computing time and keep enough information to perform an accurate reconstruction. The comparison between the original pulse series and millipede's binned waveform is shown in figure 4.6. The original pulse series is shown on a negative scale for clarity.

In figure 4.4, the pulse series contained a total of 34 pulses pre-binned into 69 bins. The first bin extends from the start of the time window to about 10000 ns which is the time the pulse series begins. The next bins have combined several pulses each until the width reaches 200 ns. The combining of bins stops if adding another bin will cause the combined bin to exceed 200 ns in width which ensures the bin will stay below this maximum width. For example, in figure 4.5, the first two bins will be combined if $w_1 + w_2 < 200$ ns. The third bin will only be added if $w_1 + w_2 + w_3 < 200$ ns. Otherwise, the third bin will become part of the



Figure 4.6: Pulse series from figure 4.4 (red, shown on negative scale) and pulse series re-binned by millipede. Note the early pulses are close together and are combined to form the wide bins. The later pulses after about 11000 ns are unchanged due to the gaps between them exceeding 200 ns.

next combined bin. This accounts for the very narrow bins at just over 11000 ns, 12300 ns and about 16000 ns seen in figure 4.6. These narrow pulses are bordered by gaps considerably wider than 200 ns, so the bins are not combined and are the same as the initial bins.

This binning produces a vector of photon counts as well as timing information for the bins. One of these vectors is produced for each DOM which has been illuminated by Cherenkov light. The "unhit" DOMs are also used by default and simply use one bin spanning the entire time window and containing no photons. The vectors for all DOMs are then concatenated to form a single vector \vec{N} .

4.4.2 Millipede Likelihood

To calculate the likelihood of a given loss pattern, the predicted photon counts are calculated and compared to the actual signal. This prediction uses the same bins and times as the observed signal and creates a vector of expected photon counts with the same number of elements as \vec{N} . This prediction $\vec{\lambda}$ is calculated using the optical properties of the ice to find the transformation matrix Λ and applying the

matrix equation from equation 4.12. The prediction for the k-th bin is given by

$$\lambda_k = \rho + \sum_{i=1}^n \Lambda_i(\vec{r}_k) E_i \tag{4.15}$$

The values of Λ_i depend on the input parameters and how the track is divided up. This prediction produces the k-th likelihood contribution which follows a Poissonian distribution so is given by

$$\mathcal{L}_k = \frac{\lambda_k^{N_k}}{N_k!} e^{-\lambda_k}.$$
(4.16)

This function is at its maximum when the prediction is equal to the input signal $(\lambda_k = N_k)$. The total likelihood is given by the product the contributions from all bins:

$$\mathcal{L} = \prod_{i=1}^{m} \mathcal{L}_i = \prod_{i=1}^{m} \frac{\lambda_i^{N_i}}{N_i!} e^{-\lambda_i}.$$
(4.17)

In practice, the millipede algorithm calculates the log of the likelihood which is then maximised. Using log laws, the log of the likelihood above is just the sum of the logs of the contributions:

$$\ln \mathcal{L} = \sum_{i=1}^{m} (N_i \ln \lambda_i - \ln(N_i!) - \lambda_i).$$
(4.18)

Due to the complexity of the likelihood function, it is impossible to find the maximum by purely analytic means. Instead, the millipede algorithm is used to find this critical point by numerical methods. In this algorithm, rather than finding the maximum of a function, the negative log likelihood (-LLH) is calculated and a minimum is found. A minimum for the negative log likelihood naturally implies a maximum for the likelihood function. At this point, all the predictions $\vec{\lambda}$ match the signal \vec{N} as closely as possible.

To reconstruct the losses, the predictions λ_i are calculated for an initial guess track known as the "seed" and the observations N_i are extracted from the waveform to calculate the negative log of the likelihood. This is repeated with different energy patterns until the algorithm finds a minimum negative log likelihood (i.e. until it finds the loss pattern and starting energy which best matches the observations.)

The millipede algorithm is also used to reconstruct properties such as direction and vertex position. To perform these reconstructions, initial guesses for these parameters are used to calculate the most likely energy loss pattern as before. The direction and position are then varied by a predefined step size and the energy loss reconstruction is repeated to try to achieve a smaller negative log likelihood. This process continues until a minimum is found.

4.4.3 Issues with Millipede

Millipede has encountered issues in reconstructing simulated events where the reconstructed parameters differ greatly from the simulated parameters or millipede fails to reconstruct the event. Studies have been conducted by members of the IceCube Collaboration to diagnose the problems with the millipede minimiser.

When the event is reconstructed incorrectly, this is thought to be due to fluctuations in the likelihood surface. The set of parameters found by millipede could correspond to a local minimum which is shallower than the true minimum. This can cause the minimiser to become "trapped" in this local minimum which could be a long way off the global minimum [34]. In extreme cases, the minimiser can jump between different local minima with every step which means it has difficulty reconstructing the event at all.

Tests have been done on simulated events where the initial guess parameters are varied to test the effect on the accuracy of the fit. These tests revealed that, if the initial guess track position is only slightly off the true position, the error margin increases greatly [36]. This means the accuracy of the reconstruction depends on the quality of the initial guess in an issue known as the "seeding problem." This, of course, proves very problematic in the reconstruction of real events where the true parameters are unknown.

Chapter 5 Initial Testing of Millipede

For the first part of this project, the reconstruction was run on a set of simulated muon events generated with the Spice-Mie ice model. In this reconstruction, a sample of track events with energy greater than 20 TeV was selected and reconstructed with a full free scan in which millipede was allowed to minimise for the direction. Unless otherwise stated, the position and energy losses of the track are being fitted as well. The initial guess track for the minimiser is named "PoleMuonLlhFit" which is an early likelihood fit using the SPE method [37]. The previous fit is compared to the more advanced millipede fit.

To properly test the accuracy of both reconstructions, the space angle between their result and the true (Monte Carlo) direction is used (see appendix A.3) as a measure of accuracy.

5.1 Comparison of Reconstructions

The space angle between the PoleMuonLlhFit direction and the true direction was compared to the angle between the millipede direction and the true direction. This is represented in a histogram in figure 5.1 showing the distribution of angular error in the fits.

The histogram shows that, for both fits, the majority of events are fit close to the true direction with a space angle difference less than 1°. Millipede performs slightly better with a greater number of events within 1°. However, there are still many events beyond 2° with some greater than 10° for both fits. This demonstrates that, if the initial guess from PoleMuonLlhFit is far from the true direction (> 10°) millipede has difficulty getting any closer to the true direction. This is possibly due to the likelihood fluctuations creating local minima. These bad fits could also be due to the signal in the DOMs being relatively faint and the track hard to pinpoint. To understand this, the angular differences were plotted



Figure 5.1: Histogram of space angles between fit and true direction for both reconstructions. Note millipede has a higher number of events with small angular error.

in a 2-dimensional histogram against the original muon energy taken from the simulated data as shown in figure 5.2.



Figure 5.2: 2D histograms showing space angle difference compared to the muon energy. Millipede is shown to have its least accurate fits at generally lower energies while PoleMuonLlhFit has its least accurate fits across all energies.

The 2-dimensional histograms reveal that, for millipede, the events which are particularly bad fits are at energies at the lower end of the spectrum of the order 10 TeV. This suggests that millipede is far better at reconstructing bright, high energy events while PoleMuonLlhFit has poor fits across the entire spectrum. Both reconstructions were compared directly in a plot of millipede angular difference against PoleMuonLlhFit difference.



Figure 5.3: Histogram of millipede error against PoleMuonLlhFit error. The line where errors are equal is shown. Millipede has a smaller error in 66.4% of the events.

The comparison plot between the two reconstructions in figure 5.3 shows a slight majority of events have millipede fits closer to the true direction (below dashed line) than PoleMuonLlhFit but there are still many where millipede is a poorer fit (above dashed line.)

5.2 Grid Scans

To better understand these issues with the fitting, likelihood grid scans have been performed which, rather than letting the algorithm find the most likely direction, provide the algorithm with a fixed direction and calculate the likelihood value. For each direction, millipede is used to minimise the negative log likelihood for all parameters (including vertex position) except the direction. This means, for each direction, as the algorithm is running, a track of constant direction is simply being translated in space as shown in figure 5.4 until the most likely position is found. These fits each produce a likelihood value which can then be used to create a map of the entire sky showing the negative log likelihood for all directions.



Figure 5.4: Example of a muon track having its vertex shifted twice while maintaining constant direction. This repeats until the best track position is found.

5.2.1 Healpix Grid

The grid scan is run using the Hierarchical Equal Area and isoLatitude Pixelisation (healpix) grid which is a method to map out a sequence of points on the surface of a unit sphere. The sky is divided into a certain number of pixels each with the same angular size in order to map out the sequence of points as uniformly as possible [38]. A specific grid is defined by a parameter known as N_{side} . This parameter is always a power of 2 (1,2,4,8,...) and is related to the number of pixels in the sky by

$$N_{\rm pix} = 12N_{\rm side}^2 \tag{5.1}$$

This means the simplest map divides the sky into 12 pixels. The next value of $N_{\rm side}$ will divide each of these pixels into four equal pieces giving a 48 pixel sky. The next will be 192 pixels and so on. Each pixel is centred on a specific direction which, in this case, is a zenith and azimuth in IceCube coordinates. With increasing pixel number, the sky is mapped out in bands of constant zenith. The sky map on the healpix grid is displayed with 0° at the top and 180° at the bottom. Azimuth has its zero in the centre and increases to the left. In the context of IceCube coordinates, this means the south celestial pole (directly overhead) is at the top of the sky map.

5.2.2 Example Scans of HESE Neutrinos

Scan for Track Event

An example grid scan created by millipede is shown in figure 5.5 for a muon event in the first high energy starting event (HESE) sample of 28. This scan was produced by the IceCube collaboration following the discovery of astrophysical neutrinos discussed in section 3.4. The colour scale shown is the negative log likelihood value for each direction. The resolution of the likelihood scan is given by the σ value defined as the point where the negative log likelihood increases by 1.15 units from the minimum. This is because the values $-2\ln\mathcal{L}$ follow a χ^2 distribution with two degrees of freedom, zenith and azimuth (see section 5.3.)



(a) Visualisation of a muon event nicknamed "Zoot".



Figure 5.5: Visualisation and sky map of muon event. The skymap shows a clear minimum region in negative log likelihood given by the blue spot centred around 60° zenith and 30° azimuth. This minimum indicates the most likely direction. The actual resolution of this event is 0.359° given by the 1.15 increase in negative log likelihood.

The full sky scan of the muon event shows a relatively narrow minimum region for negative log likelihood. This indicates a sharp "peak" in likelihood for these directions indicating fine angular resolution for the muon (0.359°) .

Scan for Cascade Event

The following millipede grid scan in figure 5.6 was performed on a very high energy ($\sim 1 \text{ PeV}$) cascade event.

The scan of the cascade shows a much broader deep blue region than the muon event which indicates the direction for the cascade has a higher level of uncertainty associated with it. The resolution of this event (1.34°) is coarser that the resolution for the muon, however, it should be noted that the high energy of this event causes it to have better resolution than most cascades.







-log likelihood

Figure 5.6: Visualisation and sky map of cascade event. The sky map shows a broader minimum region than the track centred around 70° zenith and 320° azimuth. This event has resolution 1.34°. This resolution is difficult to see in the scan as the negative log likelihood varies by about 3000 units and σ is determined by an increase of only 1.15.

5.2.3 Testing with HESE Track Event

Millipede was tested on a muon event from the HESE event nicknamed "Animal" using PoleMuonLlhFit as the seed. The visualisation is shown in figure 5.7. For each direction in the scan, millipede was given the seed vertex from PoleMuonLlhFit and allowed to vary the vertex and energy losses to find the minimum. This was performed on a sky with $N_{\rm side}$ of 8 giving 768 pixels shown in figure 5.8. Also shown, is the direction given by millipede when the likelihood is minimised for direction (the "full free fit") represented by the triangle.

Since muons, generally have very precise resolution in their scans, the behaviour of the likelihood surface is better understood by zooming in on this minimum region. Another scan was done which fitted the vertices but only scanned over directions close to the direction given by millipede's full free fit. This region consisted of a $3^{\circ} \times 3^{\circ}$ square centred on the millipede fit direction with an N_{side} of 512 to get a more precise likelihood scan. While the bottom half of the $3^{\circ} \times 3^{\circ}$ scan in figure 5.9 appears to be trending toward smaller negative log likelihood, there are many local minima where the minimiser could get stuck.

Also, the direction fitted by millipede is on the edge of this minimum region rather than in the centre which implies that the global minimum of the scan isn't where millipede has fitted the direction. To properly test the accuracy of millipede with these scans, the scans were run on a selection of simulated muon events.



Figure 5.7: Visualisation of muon event nicknamed "Animal." As shown, the muon is created in the top right and loses much of its energy in the early stages shown by red and orange. It then travels on a downward trajectory through the detector losing small amounts of its energy along the way.



Figure 5.8: Full sky map of Animal with vertex being fitted for each direction. While there is a minimum around the result from the full free fit at 54.2° zenith and 341° azimuth, there are extreme fluctuations seen throughout the scan.



Figure 5.9: $3^{\circ} \times 3^{\circ}$ scan of likelihood surface from figure 5.8. This scan exhibits fluctuations where the likelihood value is changing significantly for very similar directions. This likelihood surface has a resolution 0.46°.

5.2.4 Scans of Four Simulated Events

Four examples were chosen from the sample of simulated events to more closely compare the reconstructions given by Millipede and PoleMuonLlhFit: one where both fits are close to the true direction (within 0.1°), one where millipede is close and PoleMuonLlhFit is far (outside 1°), one where millipede is far and PoleMuonLlhFit is close and one where they are both bad fits. A grid scan was run on each of these events on a $3^{\circ} \times 3^{\circ}$ region centred on the true direction.

The directions given by the millipede free scan and the PoleMuonLlhFit reconstruction were compared to the true direction and the direction corresponding to the minimum negative log likelihood in the scan. Most of these grid scans show a great deal of fluctuations in the millipede likelihood surface. Due to these fluctuations, the scan in figure 5.10(a) where both are close has a minimum further from the true direction than the free scans. In figure 5.10(b), when millipede is close, the scan looks much smoother and, since millipede starts from PoleMuonLlhFit, this shows that millipede has reconstructed this event quite accurately.

When PoleMuonLlhFit is closer, in figure 5.10(c) the minimum from the scan gets closer to the millipede free fit but both are far from the true direction. The likelihood surface also looks relatively smooth for this event but is quite inaccurate. The scan in figure 5.10(d) where they are both bad fits shows the millipede free





The millipede likelihood surface has

minimum is close to the true direction.

minimal fluctuations and the scan

(a) Scan where both fits are close. There is an obvious minimum region of negative log likelihood but there are many local minima.



MC) Directio Millipede Fit MuonLlhFit Scan Minimum 500.00 7400.00 131 7300.00 7200.00 131. 7100.00 enith 132. 7000.00 6900.00 132.5 6800.00 133.0 6700.00

(c) Scan where PoleMuonLlhFit is close. Again, the millipede likelihood surface shows fluctuations and a minimum region which is not centred around the true direction.

(d) Scan where neither fit is close. Millipede likelihood shows huge fluctuations with no obvious minimum.

) 193.5 193.0 Azimuth (degrees) 192.5 192.0

194.5 194.0

Figure 5.10: Grid Scans of four simulated events.

fit direction and PoleMuonLlhFit are close to each other but far from the true direction. The minimum from the scan gets closer to the true direction but the likelihood surface still has a lot of fluctuations.

5.3 Test of Statistical Errors

5.3.1 1-dimensional Scans and Curve Fitting

While the grid scans in figure 5.10 demonstrate fluctuations in the likelihood, there appears to be an overall trend showing a minimum region near the true direction. A minimum region was also observed for the HESE track event shown in section 5.2.3.

A test of the overall accuracy of these scans would be to determine how often the true direction falls within one standard deviation (σ) of the minimum. For a large sample of events, this should happen 68% of the time. The value of σ also gives an indication of the event's resolution. This value can be found using the knowledge that the parameter $-2\ln\mathcal{L}$ follows a Chi squared (χ^2) distribution.

This χ^2 distribution has two degrees of freedom (zenith and azimuth), meaning σ is defined as the distance from the minimum where the value increases by 2.3 units (see appendix A.4). However, since $-2\ln\mathcal{L}$ follows the χ^2 distribution and the parameter displayed on the scans is the negative log likelihood, this critical value is halved. This means the 1σ surface on the scans contains all points within 1.15 units of the minimum.

The scans in previous examples have fluctuations which make it very difficult to find the 1σ surface from these scans alone. To find σ , the negative log likelihood must increase by only 1.15 but the scale of some of these scans is of the order 100 units. Between directions, the negative log likelihood can vary by substantially more than 1.15 units. This could mean that the value of σ is much smaller than the pixel size and the events have very fine resolution (of the order 0.01 degrees) but this is thought to very unlikely. It is more likely that the likelihood fluctuations are hiding a statistical surface. To eliminate these fluctuations for the purpose of finding σ , the likelihood data was fitted to a parabola. This parabola would have an obvious minimum and a 1σ surface around this minimum. These fits were created for the four examples seen in figure 5.10. The parabolas are shown in figure 5.11.

When the likelihood surface is fit to a parabola, it is revealed that, for the cases when millipede is close to the truth, in figures 5.11(a) and 5.11(b), the value of sigma is large and the true direction is contained in the 1σ surface.





(a) Scan where both fits are close. The minimum region is broad so the 1σ surface includes the true direction. This has a σ value of 0.53° .



(c) Scan where PoleMuonLlhFit is close. Likelihood surface is steep producing a small value of σ . 1σ surface doesn't contain true direction. This has a σ value of 0.086° .

(b) Scan where millipede fit is close. There is little change since smooth surface in figure 5.10(b). The 1σ surface includes the true direction. This has a σ value of 0.28°.



(d) Scan where neither fit is close. 1σ surface is small and doesn't contain true direction. This has a σ value of 0.16°.

Figure 5.11: Grid Scans of four simulated events.

When millipede is a poor fit, in figures 5.11(c) and 5.11(d), the steepness of the likelihood surface produces a small σ value so the true direction is not contained in the 1σ surface.

The shape of the likelihood surface can be further investigated by plotting the negative log likelihood against the angle away from the minimum. This helps verify if the overall trend of the likelihood function is parabolic. This was created for the previous four examples as shown in figure 5.12.



Figure 5.12: One-dimensional likelihood scan of four events. Plotting the negative log likelihood against angle away from minimum reveals the likelihood surface follows a parabolic trend away from the minimum. The case where millipede is a good fit in figure 5.12(b) matches a parabola the closest.

As shown in figure 5.12, the underlying behaviour of the likelihood surface appears to be parabolic despite the fluctuations. When millipede is a good fit in figure 5.12(b), the behaviour is closest to a parabola though the other cases exhibit similar behaviour. However, due to the fluctuations, it is still difficult to find the value of 1σ without fitting a parabola.

More Events

These one dimensional scans and fitted parabolas were created for another bright, down-going muon. The visualisation of this event is shown in figure 5.13.



Figure 5.13: Visualisation of Spice-Mie simulated downgoing event.



(a) Fitted parabola for one degree radial scan. 1σ surface around minimum is shown. This has a σ value of 0.18°.



Figure 5.14: 1D and 2D fitted parabolas.

Fitting a parabola to the 1D likelihood surface in figure 5.14(b) reveals that the points in the plot follow this overall trend and the fitted parabola runs through the middle of the points. On the 2D grid in figure 5.14(a), the parabola's minimum is very close to the original scan's minimum. The true direction almost sits on the 1σ surface indicating this is a relatively well reconstructed event. Through fitting a parabola to this event in 1D and 2D, it appears that there is indeed an underlying smooth likelihood surface for this event which is hidden by the fluctuations.

Another event was chosen which has a well defined horizontal track but only has an energy of 30 TeV so creates a much fainter signal. The visualisation is shown in figure 5.15.



Figure 5.15: Visualisation of faint horizontal event.

A grid scan was performed on this event on a region within 1° of the true direction. The scan in figure 5.16 exhibits an elliptical minimum region as demonstrated by the fitted parabola and its 1 σ surface. This shows that the scan is well resolved in zenith but not so well in azimuth. However, the 1 σ surface is quite small in both dimensions causing it to not include the true direction. This is an example of an event whose fluctuations cause an inaccurate likelihood surface as the true direction is outside 1 σ but the millipede fit still manages to find a minimum close to the true direction as shown by the triangle.



Figure 5.16: 1D and 2D fitted parabolas.

This was repeated for a faint event whose visualisation is shown in figure 5.17.



Figure 5.17: Visualisation of low energy event.

The visualisation of this event shows a very faint signal in the detector. As millipede has been shown in figure 5.2(a) to reconstruct these low energy events relatively poorly, this would be a good test of how the energy affects the likelihood surface. Similar to before, this grid scan was run over a 1° radius region centred on the true direction.

Figure 5.18 reveals large fluctuations in the likelihood surface of this event. Since the negative log likelihood can vary by more than two units between adjacent pixels and there is no clear minimum region, it is impossible to find σ without fitting a parabola. The fluctuations make the fitted parabola extremely flat with a



(a) One degree radial likelihood scan. (b) Fitted parabola. This has a σ value of 1.96° .

Figure 5.18: Scans of low energy event.

minimum outside the region of the scan and a large 1σ surface 1.96° in radius. This demonstrates millipede's generally poor angular resolution for faint, low energy events.

5.3.2 Results from multiple events

To test how accurate the scans are overall, this method of parabola fitting was applied to a sample of 431 events simulated with the Spice-Mie model. Each event had a grid scan performed over all directions within 3° of the true direction. This was then fit to a parabola. For each event, the space angle $\Delta\theta$ between the parabola's minimum and the true direction (i.e. the truth's offset angle) was calculated and compared to the value of σ .

These values were plotted in a scatter plot of σ against the offset angle of the truth. The scatter plot shows about half of the 431 events have 1σ surfaces which include the true direction but there are many events which have the true direction well outside 1σ . For statistical uncertainties in the directions, it is expected that about 68% would lie within 1σ . The behaviour of all 431 events as a whole is best described on a cumulative plot of the ratio $\Delta\theta/\sigma$. This plot takes the ratio of the offset angle of the truth to σ for each event and calculates how many events have a smaller value of this ratio. This can then be used to see how many events have true direction within 1σ , 2σ etc. The cumulative plot in figure 5.20 shows that slightly less than half of the events have true directions within 1σ and there still many events (about 20%) outside 1σ . The percentages for the first three multiples of σ are shown in table 5.1.



Figure 5.19: Plot of σ against offset angle of truth. Dashed line marks points where truths sits on 1σ surface. True direction is outside σ for 60% of events.



Figure 5.20: Cumulative plot of 431 events. First three multiples of σ and their corresponding event counts are shown. Also shown is the curve expected for statistical errors.

mumple of 0. 70 expected (number expected)	/0 within (number within)
$\sigma: 68.27\% (295)$	39.68%~(171)
$2\sigma: 95.45\% (412)$	71.93% (310)
$3\sigma: 99.73\% (431)$	81.21% (350)

Multiple of σ : % expected (number expected) | % within (number within)

Table 5.1: Table showing percentages of events within each multiple of σ with the expected percentages from the definition of σ .

The results show there are not enough events within each multiple of σ for these errors to be purely statistical. This suggests that the inaccuracy in the scans are caused by more than just statistical fluctuations.

5.3.3 Test of Overall Smoothness

To understand how severe the fluctuations in the likelihood surface are, the overall smoothness of the scan must be properly quantified. This is found by calculating the second derivatives of the negative log likelihood with respect to direction. For a given 1D slice on the grid scan, the second derivative is approximated by taking three adjacent points and calculating the likelihood difference between points 1 and 2 and the difference between points 2 and 3. This gives the first derivative between these points. To calculate the second derivative, the difference between these first derivatives is calculated:

$$\frac{d^2 L L H_i}{dn^2} = (L L H_{i+1} - L L H_i) - (L L H_i - L L H_{i-1})$$
(5.2)

where n is the pixel number along the 1D slice. However, this parameter must be normalised according to how flat or steep the surface is. This is determined by the second derivative of the parabola p(n) with respect to pixel number which, by definition, is constant across the slice. For all points along the 1D slice, the second derivative of the parabola is subtracted from the second derivative of the likelihood. This difference is squared to give all terms the same sign. The terms are added up over the 1D slice and the square root of the result is calculated. This is normalised by dividing by the number of terms in the sum and the second derivative of the parabola giving a relative RMS value:

$$RMS_{\rm rel} = \frac{\sqrt{\sum_{i=2}^{N-1} \left(\frac{d^2 LLH_i}{dn^2} - \frac{d^2 p}{dn^2}\right)^2}}{(N-2)\frac{d^2 p}{dn^2}}$$
(5.3)

where N is the number of points along the 1D slice, meaning, since each point used in the sum must have one point on either side, N-2 points can be used in the sum. This parameter is added up over all slices across the scan to give a sense of the overall size of the fluctuations. A larger value for $RMS_{\rm rel}$ means the fluctuations are more extreme while a small value means the scan is reasonably



Figure 5.21: Comparison between RMS_{rel} and millipede fit accuracy.

smooth.

This parameter is calculated along lines of constant zenith. The space angle between the millipede fit and the true direction is plotted against $RMS_{\rm rel}$ in figure 5.21 to see how the quality of the scan affects the accuracy of the overall fit. This comparison shows a slight upward trend with poorer fits corresponding to likelihood surfaces with more fluctuations. However, there are still events with relatively large fluctuations which fit reasonably close. $RMS_{\rm rel}$ is also compared to the muon energy to understand which types of events cause some scans to become smoother than others.

When RMS_{rel} is compared to the energy, as shown in figure 5.22, the scans with the most extreme fluctuations are revealed to correspond to some of the lowest energy events. This is similar to the behaviour found in figure 5.2 where the angular error of the fit was larger for the lower energies.



Figure 5.22: Histogram of muon energy against RMS_{rel} .

5.3.4 Comparison to smooth function

For each likelihood scan, the parabola provides a new "best fit" direction given by its minimum. This minimum could be very close to the true direction while the millipede free fit is very different. This would imply that the minimiser is getting stuck in a local minimum whereas the true minimum is found using the grid scan. The space angle $\Delta\theta$ between the parabola minimum and the true direction is plotted against the angle between the millipede free fit and the true direction in figure 5.23. The line where both are equal is shown.

The comparison between the millipede free fit and the parabola accuracy in figure 5.23 shows many events where the millipede fit is quite far from the true direction while the parabola has been able to find a closer minimum. However, there are also other events where the parabola minimum is further away. For these event scans it is possible the likelihood surface doesn't have an obvious central minimum region like the one shown in figure 5.10(b). This means, when the parabola is fit to the scan, only part of the parabola is covering the three degree radial scan meaning the minimum could be quite far away. In this case, millipede is probably getting stuck in a local minimum which happens to be near the true direction whereas the broader minimum region could be further away.



Figure 5.23: Histogram of angular error of parabola against angular error of millipede's fit.

5.4 Summary

After evaluating the overall errors in millipede fits, it becomes clear that millipede has a greater amount of accuracy than previous reconstruction algorithms. However, there are still many cases when millipede gets the wrong answer when reconstructing the simulated muon events. These errors can be of the order 10° which is well outside the resolution required for astronomy.

After performing grid scans over the millipede likelihood surface, it becomes clear that the surface is not smooth. Many events have likelihood surfaces which have severe fluctuations leading to many local minima. While fluctuations would be expected in any likelihood surface due to imperfections in the model used for the predictions, these fluctuations in millipede's likelihood space have been observed to cause errors in direction fitting which are more apparent at lower energies.

After comparing the accuracy of the likelihood scan to the accuracy of the free fit, it becomes obvious that some events have a shallow likelihood minimum which gives a larger uncertainty even with the parabola fit. The fluctuations in the likelihood surface add to the difficulty of finding the true direction using millipede's free fit. These fluctuations in the likelihood surface now need to be investigated.

Chapter 6

Investigation of Likelihood Fluctuations

Since millipede is best at reconstructing bright, high energy events, an event was chosen which should be easy for millipede to reconstruct accurately. This event would be expected to have minimal fluctuations in the likelihood surface and millipede should be able to find the best direction easily. This would determine if millipede exhibits the same behaviour seen in figure 5.10 (a), (c) and (d) even with a well defined event.

This simulated event has a starting muon energy of 1.09 PeV $(1.09 \times 10^{15} \text{eV})$ and travels along an almost horizontal track.



Figure 6.1: Visualization of high energy event in detector.

The visualization of this event shows numerous DOM hits in the detector from the creation of the muon to the muon exiting the detector.

6.1 Close Grid Scan

This bright event was scanned over a $1^{\circ} \times 1^{\circ}$ grid centred on the true direction. This scan used the position given in the full free millipede scan as the seed and fitted the vertex and energy losses for fixed directions. Since this event is high energy and has a clear track, the likelihood surface should look smooth and have a resolution less than 1°. On this scan, the minimum pixel direction, the true direction and the direction from the free reconstruction are shown.



Figure 6.2: $1^{\circ} \times 1^{\circ}$ grid scan of bright event. Note the extreme likelihood fluctuations. ($RMS_{rel} = 338$.)

The likelihood surface for this event shows similar fluctuations to the previous scans. While the free fit managed to find a direction very close to the true direction, the minimum pixel is further off. Between pixels in the scan, the negative log likelihood varies by more than 1.15 and it is impossible to determine the resolution without fitting a parabola.

6.1.1 Vertex Shifts

Since the vertex is being fitted by millipede between directions in the scan, it is possible that the likelihood could be fluctuating in the vertex space. It is possible that, for almost identical directions, the vertex is being fitted to a very different location. For a smooth scan, we would expect the vertex fit to also follow a coherent pattern, correlated to a change in the track direction. To investigate this further, a correlation study can be done in which the fitted vertex position for each direction in the scan can be compared to that direction's likelihood.

Since the free fit managed to find a direction reasonably close to the true direction in this example event, the fitted vertex for this direction is used as the reference and the fitted vertices for the other directions are compared to this original vertex. The distances between the new fitted vertices and this original position in metres are plotted out over the same region in the sky. Each pixel in this scan maps to a pixel in the scan shown in figure 6.2.



Figure 6.3: $1^{\circ} \times 1^{\circ}$ grid scan of vertex shift. This scan reveals large differences in vertex shift for very similar directions.

The grid scan in figure 6.3 showing the distance between fitted vertices shows fluctuations implying the vertex is fitted to different locations for very similar directions with no obvious pattern. There is a slight correlation where the small vertex shifts, shown in blue, correspond to the small negative log likelihood values. The greatest vertex shifts happen at the top right of figure 6.3 in the same area as the higher values of the negative log likelihood.

These vertex shifts are also plotted against the negative log likelihood for a more direct comparison. This is shown in figure 6.4.



Figure 6.4: Comparison plot of distance from original vertex against negative log likelihood. The plot reveals that the smallest values of negative log likelihood correspond to the smallest and largest vertex shifts.

The plot of distance against negative log likelihood in figure 6.4 reveals two clusters at the lower end of the likelihood scale. This shows some directions near the minimum have small vertex shifts while some have shifts up to 10 metres. This is demonstrated in figure 6.3. However, from this, it is unknown if these shifts are mainly along the particle track, perpendicular to it, or due to a combination of both. For this reason, the shifts in vertex were separated into a component parallel to the direction from the free scan and one perpendicular to this direction. This situation is illustrated in figure 6.5.

The plots from figure 6.3 and figure 6.4 were recreated this time showing the vertex shifts parallel and perpendicular to the fitted direction. The plots for the parallel component, in figure 6.6, reveal behaviour almost identical to the total distance between vertices including the two clusters at the minimum end of the scale. This indicates that the majority of the vertex shift is in the parallel direction.

The perpendicular distance scan in figure 6.7 shows that, for most directions, there is minimal shifting in this direction with most directions shifted less than 3 metres. The comparison plot reveals a more interesting trend where the overall likelihood gets worse if the perpendicular shift is greater. This verifies that it is mainly the shifts in perpendicular rather than parallel direction which affect the likelihood. This would be expected since a vertex shift in this perpendicular



Figure 6.5: Example of a vertex shifting parallel and perpendicular to the original direction where $d^2 = d_{\parallel}^2 + d_{\perp}^2$.

direction would shift the entire track in space which would change the likelihood value. However, there are some directions which obviously break this trend having a low value for negative log likelihood but a large vertex shift. This possibly indicates there is another local minimum some distance away.


(a) $1^{\circ}{\times}1^{\circ}$ grid scan for parallel component of vertex shift.



(b) Comparison plot of parallel component of shift against negative log likelihood. This reveals similar behaviour to the total vertex shift shown in figure 6.4.

Figure 6.6: Grid scan and comparison plot for parallel shift.



(a) $1^{\circ}{\times}1^{\circ}$ grid scan for perpendicular component of vertex shift.



(b) Comparison plot of perpendicular component of shift against negative log likelihood. Contrary to the parallel shift, this shows a clear positive correlation between negative log likelihood and perpendicular vertex shift.

Figure 6.7: Grid scan and comparison plot for perpendicular shift.

6.2 Vertex Scans

To analyse the likelihood function relative to the vertex, scans were performed which fixed the vertex and allowed millipede to minimise for the direction and calculate the likelihood.

6.2.1 Three-dimensional Vertex Scan

The vertices were scanned over a 3-dimensional space in a $10m \times 10m \times 10m$ cube centred on the vertex from the full free scan to get a more complete view of the likelihood surface's vertex dependence. At each point in the cube, the direction was fitted by millipede.



Figure 6.8: Likelihood scan over $10m \times 10m \times 10m$ cube. The fitted muon track is also shown. The muon can be seen entering the cube on the left at $(x, y, z) \sim$ (52, 40, 231) and exiting on the right at $(x', y', z') \sim (62, 40, 230)$. The negative log likelihood shows minima towards the centre of the cube but also shows some fluctuations in vertex space.

The scan shows the likelihood gets considerably worse at the top and bottom of the cube which is expected as these areas correspond to shifts in the perpendicular direction which has been shown to affect the likelihood. There are still some fluctuations in the likelihood between vertices but a clear minimum region towards the centre.

6.2.2 Plane Scan

Since the largest likelihood sensitivity with the vertex shifts seems to be in the plane perpendicular to the track, the vertices were scanned over a 2-dimensional plane. This plane is defined relative to the track from the full free fit. The plane is centred on the fitted vertex from this reconstruction and is perpendicular to the fitted track as shown in figure 6.9.

These positions were then fixed and used by the minimiser to find the most likely direction. Since this particular event is travelling almost horizontally along the x-axis, the perpendicular plane will be almost vertical on the yz plane. The vertex scan was done on this plane for a $2m \times 2m$ square as this is the scale of the majority of vertex shifts seen in figure.



Figure 6.9: Example of a fitted track with $2m \times 2m$ perpendicular plane centred at its vertex.

For this plane scan, as shown in figure 6.9, the plane was centred at the vertex from the full free fit while the initial guess direction was the direction given by the free scan. Each point on the plane is given two coordinates x' (the "horizontal" displacement) and y' *the "vertical" displacement) which map to a vertex in IceCube coordinates (x, y, z) as shown in section A.5. For each point in the plane, the direction is fitted by millipede. The scan is shown in figure 6.10. The likelihood scan shows a clear minimum region centred around the minimum from the free scan. Despite this overall trend toward a well defined minimum, there are still relatively large likelihood fluctuations between adjacent directions.

Scanning over the vertex space shows that the likelihood value depends greatly on vertex position. As expected, shifting the track in a perpendicular direction away from the minimum causes the negative log likelihood to become larger.



Figure 6.10: Plot of negative log likelihood for fixed points on a perpendicular plane, allowing millipede to fit the direction and energy losses. The red cross is the minimum likelihood position. This central minimum region reveals the negative log likelihood increases with perpendicular distance.

These scans also indicate that the likelihood function also has fluctuations in vertex space. Whether the fluctuations occur in vertex space more than direction space must now be investigated.

6.3 Fixed Vertex

6.3.1 Fixed Energy Losses

To check that the likelihood function has fluctuations mainly in vertex space as opposed to direction space, a new scan was created for the event in which everything (including energy losses) is fixed between directions and millipede is only used as a likelihood calculator. In the original simulation, the data is stored in a Monte Carlo (MC) tree where the muon track is accompanied by daughter products which carry energy away from the muon (see appendix B.2.) These products can be seen through the light pattern in the DOMs. Millipede would then reconstruct the energy losses from this observed signal. In this alternative method, the true muon losses are known and fixed keeping their energies and positions along the track constant.

For the grid scan in zenith and azimuth, the track is kept at a constant vertex



Figure 6.11: Example of two tracks with different directions but common vertex and the same energy loss pattern.

and allowed to pivot around this point. As the track moves, its losses stay at the same positions relative to each other and the vertex meaning they move with the track. This is illustrated in figure 6.11. The losses are kept in the same relative true positions on the track for each direction, however, the actual positions of these losses in (x,y,z) must be found for a likelihood calculation. These positions would vary as the track direction changes as shown in figure 6.11. The new position is found for each loss by calculating the distance d between this loss and the fixed vertex. The vector between the vertex and the loss is given by

$$\vec{v} = \pm d(\sin\theta\cos\phi\hat{\mathbf{x}} + \sin\theta\sin\phi\hat{\mathbf{y}} + \cos\theta\hat{\mathbf{z}}) \tag{6.1}$$

Whether +d or -d is used depends on whether the loss is "before" or "after" the vertex on the track. If the loss occurs before the particle reaches the vertex, the vector will be pointing towards (θ, ϕ) so +d is used. If the loss occurs later, the vector is pointing in the opposite direction so -d is used. This vector \vec{v} is then simply added to the vertex position vector to obtain the new position for the loss in x,y and z. These parameters are then used by millipede to calculate the likelihood. This is shown on the same $1^{\circ} \times 1^{\circ}$ grid as before. The vertex was fixed at the true vertex.

The grid scans in figure 6.12 show that the likelihood surface for fixed losses is much smoother than when the vertex and losses were being fitted as shown by the values of RMS_{rel} . However, in both scans there is an extreme jump in the top right. This is more obvious on the scan where the losses and vertex are fixed as it might be partially corrected for by the fitting of the vertex. However, even with this likelihood jump, the overall scan with fixed losses shows a smooth likelihood



(a) Grid scan over zenith and azimuth for fitted energy losses and vertex. $(RMS_{rel} = 338.)$

(b) Grid scan for fixed energy losses and vertex. This scan reveals a smoother surface with a clear minimum. $(RMS_{\rm rel} = 30.)$

Figure 6.12: Grid scan for fitted and fixed losses.

surface with a clear minimum.

6.3.2 Only Fitting Energy Losses

To finally verify that the fluctuations are the most extreme in the vertex space, the scan was run again but this time, millipede was allowed to reconstruct the energy losses for each direction in the scan while keeping the vertex position constant at its true value. Millipede is also allowed to vary the positions of the losses along the track. The region of the scan is the same $1^{\circ} \times 1^{\circ}$ grid as before.

The grid scan in figure 6.13 where energy losses are fitted but vertex is fixed also shows a likelihood surface which is not as smooth as the surface for fixed energy losses from figure 6.12(b) but is still a great deal smoother than when the vertex is being fitted in figure 6.12(a) as shown by the RMS_{rel} values. The scan now shows a clearer minimum which is not too far from the true direction and almost coincides with the free fit direction. This indicates that most of the fluctuations in the likelihood surface arise when the vertex is allowed to float in the fit.



Figure 6.13: Scan over $1^{\circ} \times 1^{\circ}$ grid with fitted energy losses but fixed vertex. The scan minimum now almost coincides with the millipede free fit. The likelihood surface is still smoother than when the vertex is fitted. ($RMS_{rel} = 108.$)

6.4 Summary

The $1^{\circ} \times 1^{\circ}$ scan of a bright muon event reveal extreme fluctuations in the likelihood surface, with respect to zenith and azimuth, which are thought to be the main cause of the problem with millipede's accuracy. Closer investigation of the individual directions of these scans reveals that the vertex is being fitted by millipede to very different locations for almost identical directions.

The main contribution to likelihood differences seems to be vertex shifts in the plane perpendicular to the muon track. While the shift in this plane is comparably small, the negative log likelihood is shown in figure 6.10 to increase with increasing distance from the centre of this plane.

Figures 6.12 and 6.13, when the vertex is fixed, show a smoother likelihood surface indicating that likelihood fluctuations arise mainly when the vertex is allowed to float freely. This suggests that most of the likelihood fluctuations arise mainly in the vertex space rather than direction space. This causes the fitted vertex to fluctuate between directions. The source of these fluctuations in vertex space must now be investigated.

Chapter 7

Possible Causes of Vertex Fluctuations

The likelihood fluctuations in vertex space are most likely due to a combination of an incomplete model of the ice and features of the waveforms in real and simulated data which millipede is not taking into account when calculating its predictions. These features could be fluctuations in the waveform, or secondary pulses created inside the PMT.

7.1 Using Millipede Predictions as Input

To check that these extra waveform features are the most likely cause of fluctuations and there are no internal errors in millipede, the predictions calculated by millipede's likelihood function are used as the input pulses for reconstruction.

In other words, the millipede likelihood function is now used as the simulator. In this method, millipede is given the true (Monte Carlo) direction, vertex position and energy losses for a simulated event and calculates the predicted waveforms. These waveforms are then used as an input signal, creating an event which millipede then reconstructs. If there are no internal errors in millipede, this should produce a smooth likelihood surface with a minimum close to the true direction. This method was tested on the bright event from figure 5.13 using millipede's predictions for the true direction as the simulated DOM waveforms. The one degree radial scans are shown in figure 7.1. Unless otherwise stated, the point on the scan labelled "Millipede Fit" is the result from the full free fit using the original pulses (i.e. before altering the waveform.)

These scans reveal that using millipede predictions as the input pulses creates an extremely smooth likelihood surface after reconstruction (as shown by the



(a) Scan using Original Waveforms as input. This has $RMS_{rel}=45$.

(b) Scan using Millipede's Predicted Waveforms as input. This has $RMS_{\rm rel}{=}5.5.$

Figure 7.1: One degree radial scans using original simulated pulses and millipede's predicted pulses. Note when millipede predictions are used as the input, the likelihood surface becomes smooth.

 $RMS_{\rm rel}$ values) which is centred on the true direction. The minimum of the scan is within one bin of the true direction. This strongly implies that the source of the fitting errors and likelihood fluctuations is differences between the predicted waveforms and the actual pulses. The next test is to artificially add waveform fluctuations to understand how the likelihood surface is affected.

7.1.1 Adding Poisson Fluctuations

After creating these predictions from the true particle parameters, Poisson fluctuations are added by replacing the predicted charge of each pulse with a Poisson deviate which is a random variable sampled from a Poissonian distribution (see appendix B.4.) The original value predicted by millipede is used as the mean of this distribution. The Poisson deviate replaced the millipede prediction for each pulse to introduce waveform fluctuations. The 1° radial scans are compared in figure 7.2.

When the Poisson deviate is added, it is revealed that fluctuations are introduced into the likelihood space as is reflected in the RMS_{rel} values. However, the likelihood surface is still smoother than the result using the standard simulation shown in figure 7.1(a). While this result implies that waveform fluctuations can cause likelihood fluctuations, it also implies that Poisson fluctuations in the true waveform are probably not the only cause of likelihood and vertex fluctuations. The other features in the waveform, which



 $RMS_{rel}=5.5.$

Poisson Fluctuations as input. This has $RMS_{rel}=13.$

Figure 7.2: One degree radial scans using millipede predictions and Poisson deviate.

millipede is apparently not taking into account, must now be investigated.

7.2Attempts to Remove After-pulses

After-pulses are spurious features in a waveform which appear several microseconds after the main pulse and are created by ionisation of residual gases in the PMT. While the PMT requires a vacuum to operate properly, in reality, there are trace amounts of gas such as hydrogen, oxygen and helium contaminating the PMT.

As an electron travels from the photo-cathode to the anode, it accelerates and gains energy. When a high energy electron strikes the gas molecule, it can ionise the gas molecule by ejecting one or more of its electrons. This changes the gas molecule into a positive ion which is drawn by the electric field towards the negative cathode. When the ion strikes the cathode, its energy is given to electrons which accelerate towards the dynodes and anode. This process is shown in the diagrams in figure 7.3.

These electrons create new pulses in the DOM which are weaker than the main pulse. This new signal will also arrive at a later time determined by the potential difference and the mass and charge of the positive ion. Singly ionised helium (He⁺) appears to be the main contributor to after-pulses [40].

For the PMTs used in IceCube at their operating voltage of around 1300 V,



Figure 7.3: Diagrams of PMT showing production of after-pulses. **Top:** The primary electron accelerates towards the positive anode and collides with contaminant gas. **Middle:** The ionised gas accelerates to the negative cathode. **Bottom:** Once the positive ion strikes the photo-cathode, secondary electrons are produced and accelerate towards the anode.

the after-pulses have been found to be around 300ns to 11μ s after the main pulse [20]. The waveform predictions calculated by millipede don't include any after-pulses. However, the after-pulses are present in the observed waveforms for real events and are also included in Monte Carlo simulated events. It is thought that these differences in the predicted and measured waveforms could account for the fluctuations in the likelihood surface [41]. To test this idea, for each DOM in the detector, the median time of the waveform is calculated and all pulses more than 3μ s after this time are removed from the waveform.

To demonstrate the nature of the after-pulses in the waveform, the average pulse is plotted for the bright muon event from figure 5.13. For each DOM, the charge of a given pulse is calculated as a fraction of the total charge in the detector. The time of the pulse relative to the median time is also calculated. The results from all DOMs are collected together to show how the waveform evolves as a function of the offset from the median time. This average pulse is plotted for the original simulation and for millipede's predictions in figure 7.4. The cut at 3μ s is shown by the dashed line.



10 10 10 10 10 of Total Charge 10 10 10 10 10 10 10 10-1 10⁻¹ 10-14 15000 -5000 0 5000 Time after Median (ns)

(a) Waveform from original simulated pulses.
Note the relative increase of charge around 8µs.

(b) Waveform Predicted by Millipede. Here, there is no extreme charge increase at 8μ s.

Figure 7.4: Average waveform for actual pulses and for millipede prediction. In both waveforms, the charge is seen to decrease almost exponentially after the peak around zero. In the millipede predictions, this trend continues to the end of the waveforms but in the actual simulation, the charge starts increasing around 5μ s to form a secondary peak at 8μ s due to the contributions of the after-pulses in all DOMs.

The average waveform reveals a clear bulge around 8μ s after the median which is not present in millipede's prediction which seems to follow an exponential decay. However, this secondary pulse is over 1000 times smaller in amplitude than the main pulse.

7.2.1 Reconstructions with Time Cut

A grid scan was performed on this same event, after removing all pulses more than 3μ s after the median to investigate the effect on the likelihood surface. This is displayed in figure 7.5.



Figure 7.5: One degree radial scans using original pulses and 3μ s cut. Note the 3μ s cut produces a smooth likelihood surface with a much smaller $RMS_{\rm rel}$ value.

As shown in figure 7.5(b), removing pulses more than 3μ s after the median time dramatically improves the shape of the likelihood surface (RMS_{rel} decreases from 45 to 3.1) and even shifts the minimum region closer to the true direction. This gives a strong indication that, despite the very small amplitude of the after-pulses, they seem to have a large effect on the reconstruction. As the vertex shifts in the perpendicular direction were shown in chapter 6 to closely correlate with the negative log likelihood, these shifts should now also produce a smooth pattern. Similar to section 6.1.1, for each direction in the scan, the perpendicular distance between the fitted vertex and the vertex of the minimum negative log likelihood is calculated and displayed in figure 7.6.

When the original waveforms are used in the reconstruction, as shown in figure 7.6(a), there is no obvious pattern to how the vertex is fitted as a function of zenith and azimuth. Like the example in section 6.1.1, very similar directions have vertices fitted to very different positions. After the 3μ s cut in figure 7.6(b), the vertex shift shows a clearer minimum region around the centre and adjacent directions are fitted to similar positions. The scan of vertex shift isn't as smooth as the likelihood surface but it definitely shows an improvement. The likelihood is also plotted on a 3D diagram showing the vertices fitted across the



Figure 7.6: One degree radial scans showing perpendicular vertex shift (in metres) on colour scale. Note when the 3μ s cut on after-pulses is applied, the vertex shifts show a clear minimum near the true direction.

scan with negative log likelihood on the colour scale. In these plots, shown in figure 7.7, each point in 3D space is a pixel from the one degree radial scans shown in figure 7.5. The position of each point shows the fitted position for this fixed direction and the colour of the point indicates the negative log likelihood.

The 3D plot, in figure 7.7(a), of the vertices reveals likelihood fluctuations when the original pulses are used. When the 3μ s cut is applied, in figure 7.7(b), the distribution becomes a smoother, elliptical pattern with a clear minimum region in negative log likelihood represented by the blue points on the figure. These points are also represented on a 2D plot in figure 7.8 by projecting them onto the average plane passing through the points. The scan over the average plane shows, for the original waveforms in figure 7.8(a), there is no overall pattern in the negative log likelihood values while for the 3μ s cut in figure 7.8(b), the negative log likelihood has a clear central minimum in vertex space. Also, the spread in fitted vertices is reduced (now all vertices fit within 3 metres.) This reveals that, when the cut is applied, the likelihood surface is smoothed out for vertices as well as direction.



Figure 7.7: 3D Plots showing negative log likelihood for fitted vertices. Each point corresponds to a pixel in the scans in figure 7.5 and the colour scale corresponds to the negative log likelihood. When a 3μ s cut is applied, the vertices form a more coherent pattern with minimum negative log likelihood towards the centre.



Figure 7.8: Plots showing fitted vertex positions on average plane with negative log likelihood on the colour scale. When the 3μ s cut is applied, the vertices become more organised, with higher values (orange-red points) on the outside and a minimum (blue points) near the centre.

1-dimensional Scans

As shown in section 5.3, when the original simulated waveforms are used, the fluctuations in the likelihood surface make it impossible to find σ without fitting a parabola. Since, for this event, the scan is smoother after the 3μ s cut, it is possible that σ can now be found from the likelihood values alone (without fitting a parabola.) This event is now plotted with negative log likelihood against angular separation from the minimum. In these plots, each point is given by a pixel in the scans from figure 7.5. These plots are shown in figure 7.9.



Figure 7.9: One dimensional scans showing negative log likelihood against space angle away from minimum when using the original simulated waveforms and the 3μ s cut on the waveforms. Note that when the 3μ s cut is applied the pattern becomes narrower, more closely resembling a parabola. The value of σ is shown to be about 0.2° which agrees closely with the value from the fitted parabola shown in figure 5.14(a).

The plot with the 3μ s cut reveals that the likelihood function has indeed become closer to a parabola. The fact that the points don't sit perfectly on a parabola, and there is still some spread, implies that the smooth likelihood scan in figure 7.5(b) is not completely symmetric. In other words, the higher values of negative log likelihood (the top curve in figure 7.9(b)) are in the direction in which the negative log likelihood increases more rapidly than for the bottom curve. However, the surface is now smooth enough that it is possible to estimate a σ value without fitting a curve. Simply using the negative log likelihood values produces a σ value of about 0.2° for this event which agrees closely with the value obtained from the parabola shown in figure 5.14(a).

Test with Another Event

This 3μ s cut was applied to another simulated track to see if it has the same smoothing effect on the likelihood surface. The likelihood scans in zenith and azimuth are shown in figure 7.10.



(a) Scan with Original Pulses. This has $RMS_{rel}=50$.

(b) Scan with 3μ s cut on pulses. This has $RMS_{rel}=42$.

3.5

1988.50

1987.00

1985.50

1984.00

1982.50 1981.00

1979.50

1978.00

1976.50

Figure 7.10: One degree radial scans using original waveform and 3μ s cut. Note that when the cut is applied, while the likelihood surface becomes, overall, slightly smoother, there is a region to the left of the true direction where the values suddenly increase.

For this event, as shown in figure 7.10(b), removing pulses more than 3μ s after the median seems to slightly broaden the minimum region of the direction scan and the likelihood surface, overall, is smoother. However, there is now a region where the negative log likelihood suddenly increases with respect to nearby directions. This behaviour wasn't present in the scan from the original pulses which indicates that, by performing this cut at 3μ s, there is potentially useful "late light" which is being cut. The one dimensional scans for this event, in which negative log likelihood is plotted against space angle away from minimum, is shown in figure 7.11.

After the 3μ s cut on this event, the one dimensional likelihood points in figure 7.11(b) spread out with respect to the points in figure 7.11(a). However, it should be noted that the overall trend of the negative log likelihood appears more parabolic than before. The points in figure 7.11(a) seem to follow a linear path while in figure 7.11(b) the maximum and minimum points follow a curved path. However, unlike the event shown in figure 7.9, there is a notable drop in negative log likelihood around 0.3° away from the minimum. This degeneracy means, for this event, a parabola must be fit to the surface to find σ .



Figure 7.11: One dimensional scans showing negative log likelihood against space angle away from minimum when using the original simulated waveforms and the 3μ s cut on the waveforms. Note that when the 3μ s cut is applied the points become more spread out.

These results show that the method of simply removing pulses more than 3μ s after the median time doesn't improve all scans in general and, as shown in figure 7.10, may introduce odd behaviour. However, even with this behaviour, the likelihood surface has become overall smoother for both of these events (as given by the $RMS_{\rm rel}$ values.)

7.2.2 Tightness of the Minimum

The resolution, σ , of a likelihood scan is defined by the points where the negative log likelihood increases by 1.15 units. In order to calculate the value of σ directly from a likelihood scan (without fitting a parabola), the negative log likelihood must vary gradually enough such that the 1.15 level is at least one pixel away from the minimum. In other words, if, when moving one pixel away from the minimum, the negative log likelihood increases by substantially more than 1.15, this is a very tight minimum and the 1σ surface is hidden.

To investigate the tightness of the minima, the average increase in negative log likelihood is calculated for a series of pixels neighbouring the minimum of the scan. This series of pixels is illustrated in the example shown in figure 7.12. If the increase in negative log likelihood is greater than 1.15, the σ surface is buried between the pixels. The standard deviation (spread) in the negative log likelihood values is also calculated. If this spread is also more than 1.15, this confirms the fluctuations are still too great to determine the resolution.



Figure 7.12: Example of the scan minimum and surrounding pixels. The negative log likelihood is calculated for pixels neighbouring the minimum (dashed line) and the minimum negative log likelihood is subtracted to find the increase.

The spread in negative log likelihood values is compared to the average increase for the neighbouring pixels, as shown in figure 7.12, for a sample of 250 events. This is shown in a histogram, in figure 7.13, for reconstructions using the original simulated pulses and using the 3μ s cut pulses.

As shown in the histograms in figure 7.13, there are many events in which the average likelihood increase is more than 1.15 and the likelihood spread is more than 1.15 (outside the dashed lines.) When the original pulse series is used in the reconstruction, 67% of events lie outside the dashed box shown in figure 7.13(a) while, for the 3μ s cut shown in figure 7.13(b), 35% are outside this box. This implies that, when the 3μ s cut is applied, there are more events in which the negative log likelihood increase and spread are small enough for the σ surface to be calculated directly.

To check that there is a 1σ surface which can be found from the negative log likelihood values alone, two more events were selected which had increases, from minimum to adjacent pixels, less than 1.15 when the 3μ s cut is applied. These scans, centred on the true (Monte Carlo) direction, are displayed in figure 7.14 with the negative log likelihood represented by the colour scale. The colour scale is set to only include pixels in which the increase is less than 1.15. Where the pixel colour changes to red indicates where the 1σ surface can be found. Also plotted, is the 1σ surface obtained by fitting a parabola.



Figure 7.13: Comparison histogram of spread in negative log likelihood values against average increase in negative log likelihood for pixels neighbouring the minimum. The dashed lines show the points when either quantity is equal to 1.15. Events with a negative log likelihood increase greater than 1.15 have very tight minima which hide the σ value. When the original pulses are used, about 67% of events have minima which are too tight. After the 3μ s cut, 35% have minima which are too tight.



(a) Scan of event with a relatively broad minimum using the 3μ s cut.

(b) Scan of another event with a broad minimum using the 3μ s cut.

Figure 7.14: Scans of two events in which the increase after one pixel is less than 1.15. The colour scale shows the negative log likelihood increase from the minimum. The transition to red pixels indicates the 1σ surface obtained directly from the likelihood values. For both events, this surface seems to agree with the 1σ surface from the parabola.

As shown in the event scans in figure 7.14, the 1σ surface found directly from

the likelihood surface agrees closely with the surface obtained from the parabola. This implies that, for some events, the minimum is broad enough such that it is possible to estimate the σ values without fitting a parabola. The likelihood surface was also plotted out for an event which had an increase greater than 1.15 for the first pixel to understand the nature of this minimum (whether it is a local or global minimum.) This is shown in figure 7.15. As before, the colour scale shows the negative log likelihood increase and the 1σ surface for the parabola fit is shown.



Figure 7.15: Scan showing increase in negative log likelihood from the minimum. This scan shows a broad minimum region but a deep local minimum where millipede is becoming trapped. The 1σ surface shown is centred around the parabola minimum which demonstrates the parabola fit has correctly found this minimum region.

As can be seen in figure 7.15, the negative log likelihood increases by about 4 units one pixel away from the minimum though there is a broad minimum region shown in light blue. The parabola's 1σ surface is centred on the parabola's minimum. This indicates that the parabola is finding the broad minimum region which is closer to the true direction (the centre of the scan), which implies that the deep minimum found by millipede is false.

According to figure 7.13, once the 3μ s cut is applied, there are still many events in which the minimum is too tight to determine the σ value. Some events still have negative log likelihood increases and spreads of the order 1000. Performing a finer grid scan could reveal that these events simply have very fine resolution. However, as shown in figure 7.15, it is more likely that most of these events have deep local minima due to likelihood fluctuations. A parabola fit is needed to find the broader minimum region shown in figure 7.15. More work could be done to determine why millipede is producing these narrow minima but, in the next section which investigates the scan accuracy, σ is calculated by fitting parabolas.

7.2.3 Overall Accuracy of Scans

Similar to section 5.3, grid scans were performed on 431 simulated muon events with this 3μ s cut on the pulses. As has been shown in figure 7.11 and figure 7.13, there are events where performing the 3μ s cut doesn't produce a smooth enough surface to calculate σ . For this reason, the grid scans of these events were, again, fit to parabolas to find σ . The space angle between the true (Monte Carlo) direction and the minimum of the parabola $\Delta\theta$ was calculated. The ratio $\Delta\theta/\sigma$ is displayed in a cumulative plot in figure 7.16 showing the percentage of events falling within each multiple of σ . This is also summarised in table 7.1.



Figure 7.16: Cumulative plots of $\Delta \theta / \sigma$ for 431 events.

Multiple of σ	% within for Original Waveform	% within for $3\mu s$ cut
σ	39.68%	38.98%
2σ	71.93%	73.55%
3σ	81.21%	83.53%

Table 7.1: Table showing percentages of events within each multiple of σ for both original waveforms and 3μ s cut.

The result, in figure 7.16, shows that performing the 3μ s cut does little to increase the accuracy of the grid scan and, while some scans may become



Figure 7.17: Comparison of $\Delta\theta/\sigma$ between original pulses and 3μ s cut. The dashed line shows where $\Delta\theta/\sigma$ is equal for both original pulses and 3μ s cut. Events below this line have scans of increased accuracy. The result shows about half have increased accuracy while the accuracy gets worse for the other half.

smoother, the minimum is still around the same place. As shown in table 7.1, there are now fewer events within 1σ but more within 2σ though not by much. This is summarised in the comparison histogram shown in figure 7.17. The comparison plot shows some events where the position of the minimum relative to σ improves and some where the minimum gets further away relative to σ . However, most events show little or no change implying this cut has very little effect on the accuracy of the grid scans.

The cut was further compared to the results using the original simulated waveforms by performing a full free fit on a large sample where millipede is used to reconstruct all parameters including the direction to see how the space angle between the fit and true direction improves. This is shown by calculating the space angle between the true direction and the fit for the original pulses and the 3μ s cut and subtracting the cut value from the original value. A positive number means the accuracy has improved. The resultant histogram is shown in figure 7.18.

The histogram of energy against accuracy improvement in figure 7.18 shows that most events have hardly any change after applying the cut and there seems to be no energy dependence on whether or not the accuracy improves. Particularly for the low energy events, there are around the same number of events which have poorer fits as there are which have better fits. This, again, implies that performing



Figure 7.18: Histogram of muon energy against difference between original error and error with 3μ s cut. Note, for almost all events, there is very little change.

the 3μ s cut has little effect on the accuracy.

7.2.4 Comparison of Fluctuations

To check the overall effect on the scans, for both the original waveforms and the 3μ s cut on the original waveforms in the sample of 431, $RMS_{\rm rel}$ was found in the method described in section 5.3.3 for lines of constant zenith. The value of $RMS_{\rm rel}$ was compared to the value given by the original pulses. This comparison is plotted in a histogram shown in figure 7.19. The line of equality is shown where $RMS_{\rm rel}$ doesn't change after the cut. Events which fall below this line have smoother scans after the cut. The comparison in figure 7.19 shows most events have smaller $RMS_{\rm rel}$ values after the cut implying smoother scans though some of these changes may not be obvious by simply looking at the scans.

Finally, the change in RMS_{rel} is compared to the initial (Monte Carlo) energy of the muon to understand which types of events are most affected by this 3μ s cut. This is plotted in a histogram of muon energy against the ratio $(RMS_{rel})_{Original}/(RMS_{rel})_{Cut}$. Values of this ratio greater than one correspond to the 3μ s cut producing a smoother scan than the original pulses. The histogram of muon energy against this ratio is shown in figure 7.20.

This histogram reveals most events (~ 74%) have RMS_{rel} ratios greater than one, meaning, as shown in figure 7.19, most events have smoother scans after the



Figure 7.19: Comparison histogram of RMS_{rel} . Note the majority of events (74%) are below the line of equality indicating that most events produce a smoother scan when the 3μ s cut on the pulses is applied.



Figure 7.20: Comparison histogram of energy against RMS_{rel} ratio. Note most events have minimal change while there are some low energy events where RMS_{rel} increases after the cut, possibly due to potentially useful late light being removed.

 3μ s cut. There are, however, some low energy events and a small number high energy events where $RMS_{\rm rel}$ increases after the cut. This is probably an extreme case of the example shown in figure 7.10 where the cut is removing late, scattered light which could be used in the reconstruction. There are also some low energy events in which the scan is becoming considerably smoother. This could either be due to the removal of after-pulses or fluctuations occurring in the tail of the waveform.

These results have shown that, while the likelihood surfaces and accuracy of some events improve as a result of the 3μ s cut as shown in figure 7.5 when $RMS_{\rm rel}$ dropped from 45 to 3.1, about one third of events show little or no change and some events (about 26%) have scans with more fluctuations when the cut is applied. This demonstrates that, in general the reconstruction can't be improved by simply removing the pulses 3μ s after the median. However, the result from the event shown in figure 7.5 strongly suggests that the after-pulses are the main contributor to the likelihood fluctuations and inaccuracies in the fitting. The effect of after-pulses must now be compared to the effect of Poisson fluctuations to determine which has the greater effect on the results.

7.3 Adding After-pulses and Fluctuations

To check which, between the after-pulses and Poisson fluctuations, has the greatest effect on the likelihood fluctuations and overall accuracy, both after-pulses and fluctuations are added to the millipede predictions found in section 7.1. The fluctuations and after-pulses are applied to the millipede predictions, using methods described in Appendices B.4 and B.5 respectively, for the sample of 431 events. The values of $RMS_{\rm rel}$ are calculated for the scans from these new sets of waveforms as well as the smooth scans using millipede predictions.

7.3.1 Adding Poisson Fluctuations

To understand if the likelihood fluctuations get more extreme when Poisson fluctuations are added, the ratio $(RMS_{rel})_{Poisson}/(RMS_{rel})_{Prediction}$ is calculated to compare RMS_{rel} for the added fluctuations to RMS_{rel} for millipede's predictions. The muon energy is plotted against the RMS ratio for an added Poisson deviate to check which types of events are affected by this change. Plotting the muon energy against the increase in fluctuations in figure 7.21 reveals the events where the Poisson deviate has the greatest effect generally correspond to the lower energies. However, there are some events where the Poisson deviate appears to produce a smoother scan than when millipede predictions are used. This is probably due to millipede failing to reconstruct a dim event and the Poisson fluctuations producing a smoother surface by chance. The full free fit was also



Figure 7.21: Histogram of muon energy against ratio of RMS_{rel} for Poisson fluctuations relative to RMS_{rel} for millipede predictions. Note that most of the events which are affected by Poisson fluctuations are low energy. Some of these have ratios less than one which means, possibly due to a dim signal in the detector, millipede has trouble reconstructing these events and, by chance, the Poisson fluctuations manage to create a smoother surface.



Figure 7.22: Histogram of muon energy against increase in angular error when the Poisson deviate is applied. Note, again, that high energy events are mostly unaffected by the Poisson fluctuations.

run on a larger sample of events with added fluctuations. For each event, the space angle (angular error), α is calculated between the fit and the true direction for both the fit using millipede predictions and the fit using the predictions with added Poisson fluctuations. The muon energy is plotted in a histogram in figure 7.22 against the difference in angular error ($\alpha_{\text{Poisson}} - \alpha_{\text{Prediction}}$.)

This histogram in figure 7.22 shows Poisson fluctuations have a greater effect on the accuracy of low energy events. Some events have angular errors increasing by more than 4° though there are some events where the Poisson fluctuations produce a more accurate fit. Most high energy events are unaffected by the fluctuations.

7.3.2 Adding After-pulses

These same tests were applied to the events which had after-pulses added to the millipede predictions. In this method, the after-pulses are approximated using a probability of after-pulses given an initial number of primary photo-electrons. In other words, the size and number of after-pulses depends on the magnitude of the primary pulse predicted by millipede. This probability is shown in appendix B.5 and defines the probability of a certain number of after-pulses given a number of initial photo-electrons and an ionisation probability. This method was tested first on the event shown in figure 7.5 and the ionisation probability was varied until the magnitude of the after-pulses approximately resembled the increase shown in figure 7.4(a). The relative abundances of the different elements affect the time distribution of after-pulses and are approximated from [40]. The plots of the average pulses are shown in figure 7.23.

As shown in figure 7.23, adding after-pulses to the millipede predictions produces an increase in average charge 1000 times smaller than the main peak. This pulse series is then given to millipede to perform reconstruction. Since the predictions calculated in the likelihood function won't take these new after-pulses into account, the likelihood surface should have more fluctuations. The one degree radial grid scan is shown for this event in figure 7.24.

The scan of this event shows that, when after-pulses are added to the millipede predictions, in figure 7.24(b), the likelihood fluctuations get considerably worse though the minimum is still relatively close to the true direction. While the after-pulses are of a similar magnitude to the original simulation, the scan produced bears little resemblance to the original scan shown in figure 7.5(a). This demonstrates how much the specific behaviour of the after-pulses can affect the likelihood surface. This was tested across the same events from section 7.3.1 to understand how these added after-pulses affect the quality and overall accuracy



(a) Waveform predicted by millipede with no fluctuations or after-pulses.

(b) Millipede prediction with added after-pulses.

Figure 7.23: Average waveform for millipede predictions and predictions with added after-pulses.



Figure 7.24: One degree radial scans using pure millipede predictions and predictions with added after-pulses.



Figure 7.25: Histogram of muon energy against ratio of RMS_{rel} for added after-pulses relative to RMS_{rel} for millipede predictions. Note, in contrast to Poisson fluctuations, the after-pulses have at least a small effect across all energies. There are some dim events where the surface manages to become smoother while the events which have increased fluctuations are generally at high energies.

of the fits. Similar to the tests in figures 7.21 and 7.22, the fluctuations are compared between the result using these added after-pulses and pure millipede predictions as well as the overall accuracy of the fits. The plot of the muon energy against the ratio $(RMS_{rel})_{AfterPulses}/(RMS_{rel})_{Prediction}$ is shown in figure 7.25.

The plot in figure 7.25 reveals a slight trend towards higher energy muons having more fluctuations when the after-pulses are added. This is probably due to the high energy events producing a brighter signal in the detector with a greater number of photo-electrons making after-pulses more likely. The muon energy is plotted in figure 7.26 in a 2D histogram against the increase in error $(\alpha_{\text{AfterPulses}} - \alpha_{\text{Prediction}})$ for this method.

This plot, in figure 7.26, of the increase in angular error for added after-pulses shows, similar the result shown in figure 7.25, added after-pulses are shown to affect the accuracy across a range of energies. There is also an overall higher proportion of events where the error has substantially increased (by more than 2°) than when the Poisson deviate was applied. This implies that the after-pulses, and millipede's failure to account for them, are a major factor in the reconstruction errors and could explain the errors and fluctuations seen in the high energy event from section 6.1.



Figure 7.26: Histogram of muon energy against increase in angular error when after-pulses are added. The error increase occurs across a wider range of energies than for the Poisson fluctuations.

7.4 Summary

After testing the reconstruction algorithm with various input waveforms, it becomes clear that the accuracy of the reconstruction and the fluctuations in the likelihood surface depend greatly on the input waveforms given to millipede for reconstruction.

When the millipede predictions are used as the input rather than the original simulated waveforms, the likelihood surface becomes almost perfectly smooth with a minimum which matches the true direction. This shows that the primary cause of the fluctuations in the likelihood surface is not the calculation of the likelihood itself but millipede's failure to accurately predict the waveforms. After allowing these predicted pulses to vary according to a Poisson distribution, the likelihood fluctuations are introduced though are not as extreme as the result with the original simulation. The effect tends to be greater at lower energies.

When the waveform is cut 3μ s after the median time to remove after-pulses, the likelihood surface becomes smoother in most events (about 74%) but shows very little change in about one third of the sample. After adding a Poisson deviate and after-pulses to the signal for a number of events, it was found that many of the low energy events had more fluctuations when the Poisson deviate was applied. When after-pulses were added, the fluctuations became more extreme for most events with a slight trend towards high energy events. This suggests that the cause of the fluctuations in the likelihood surface and millipede's inaccuracy are due to a combination of Poisson fluctuations and after-pulses (or other late light) in the waveforms. The fluctuations in the waveform have little effect on the high energy events while the after-pulses and late light have a noticeable effect across all energies though slightly more for high energies.

Chapter 8 Conclusions

In this work, the accuracy of millipede as a direction reconstruction algorithm was assessed. When run on a sample of simulated events, millipede was found to be a more accurate algorithm than the SPE method which preceded it. However, millipede still showed large inaccuracies for many low energy muon events as well as for some high energy events. These inaccuracies were found to be mainly due to the fluctuations in the likelihood surface creating local minima for millipede to become "trapped."

After investigating these fluctuations in more detail, it became clear that the likelihood fluctuations appear in all dimensions including the vertex space. It was found that, when the vertex is fixed, and millipede is used as a likelihood calculator, the likelihood surface becomes smooth. This implies that the likelihood fluctuations arise when the vertex is allowed to float and is fitted by millipede.

After replacing the usual simulated waveforms with waveforms predicted by millipede's likelihood function for the true direction, the likelihood surface becomes smooth and is centred on the true direction indicating there are no internal errors in millipede. It has been previously suggested that millipede's failure to account for after-pulses in the waveform are contributing to its fluctuations and inaccuracy. Removing all pulses more than 3μ s after the median time has been shown to slightly improve the reconstruction in most cases (about 74%), though in about a third of the events there is little or no change and about 26% of events have poorer fits when the cut is applied.

Adding Poisson deviates and after-pulses to the predictions from millipede's likelihood function is revealed to generally increase the fluctuations and decrease the accuracy. Poisson fluctuations are seen to have a greater effect on low energy events and a very small effect at higher energies. After-pulses have a greater effect than Poisson fluctuations across most energies with a slight trend towards
high energy events. However, it should be noted that this doesn't prove physical after-pulses are the main issue. This simply means late light, which has been unaccounted for in the likelihood function, has the greatest effect. This late light is, most likely caused by after-pulses in the PMT but could also be due to an inaccurate ice model.

8.1 Future Work

Since the late light which is not taken into account seems to be the main cause of the fluctuations and inaccuracies, more could be done to incorporate these features into millipede's likelihood function. Millipede's inaccurate behaviour seems to be due to differences between its likelihood prediction of the waveforms and the Monte Carlo simulated waveforms. Whenever these waveforms are made to be the same, the likelihood fluctuations are strongly reduced as shown in figure 7.1. In reality, there will be after-pulses in the waveform and, currently, millipede's predictions don't take these after-pulses into account.

The late light may be partly due to inaccuracies in the ice model but more could be done to incorporate after-pulses into the millipede predictions. IceCube is interested in high energy neutrino events which will probably produce bright signals with many after-pulses. If the after-pulses are the main cause of this late light, and the millipede predictions were to incorporate these after-pulses, the predictions would more closely match the input pulses. As has been shown in section 7.1, this would produce a smoother likelihood surface which accurately finds the minimum. The increase in accuracy could be examined by comparing the space angles in a histogram similar to figure 5.1.

If the accuracy in the direction fitting for these simulated events improves to well within one degree for a larger range of events, these methods can be applied to real data as it would now be possible to perform neutrino astronomy on a larger sample of events. With more accurate and precise direction reconstruction, it would be possible to perform more accurate correlation studies between neutrinos and cosmic rays as well as more accurate searches for neutrino point sources.

Appendix A Derivations

A.1 Cosmic Ray Acceleration and Spectrum

A.1.1 Second Order fermi Acceleration



Figure A.1: Diagram of particle scattering through cloud with incident angle θ_1 and exit angle θ_2 [42].

The cloud is moving with velocity v along the x-axis. The particle's initial energy in the cloud's frame E'_1 is related to the energy in the lab frame by the Lorentz transform.

$$E_1' = \gamma(E_1 - \beta p_x c) \tag{A.1}$$

where $\beta = v/c$, $\gamma = 1/\sqrt{1-\beta^2}$ and $p_x = |\vec{p_1}|\cos\theta_1$ and is the component of the particle's velocity in the x direction. Since the particle is travelling close to the speed of light, $|\vec{p_1}|c \sim E_1$ so the cloud frame energy can be rewritten as

$$E_1' = \gamma E_1 (1 - \beta \cos \theta_1) \tag{A.2}$$

Similarly, the exit energy in the lab frame can be found by Lorentz transforming out of the cloud frame.

$$E_2 = \gamma E_2' (1 + \beta \cos \theta_2') \tag{A.3}$$

Since the scattering in the clouds is collisionless, there is no energy change in the cloud frame, hence $E'_2 = E'_1$. Energy gain is given by

$$\frac{\Delta E}{E_1} = \frac{E_2 - E_1}{E_1} = \frac{\gamma E_2' (1 + \beta \cos \theta_2')}{E_1} - 1 \tag{A.4}$$

$$\Rightarrow \frac{\Delta E}{E} = \frac{\gamma^2 E_1 (1 + \beta \cos\theta'_2 - \beta \cos\theta_1 - \beta^2 \cos\theta_1 \cos\theta'_2)}{E_1} - 1 \tag{A.5}$$

$$\Rightarrow \left\langle \frac{\Delta E}{E} \right\rangle = \frac{1 + \beta \left\langle \cos\theta_2' \right\rangle - \beta \left\langle \cos\theta_1 \right\rangle - \beta^2 \left\langle \cos\theta_1 \right\rangle \left\langle \cos\theta_2' \right\rangle}{1 - \beta^2} - 1 \tag{A.6}$$

where $\langle \cos\theta_1 \rangle$ and $\langle \cos\theta'_2 \rangle$ are the average values of $\cos\theta_1$ and $\cos\theta_2$ respectively. After many interactions in the cloud, the exit direction is randomised so $\langle \cos\theta'_2 \rangle = 0$ [42]. The value of $\langle \cos\theta_1 \rangle$ is found using the collision rate over all angles. This rate is given by

$$\frac{dn}{d\Omega_1} \propto (1 - \beta \cos\theta_1) \tag{A.7}$$

where $d\Omega_1 \propto d\cos\theta_1$. The value of $\langle \cos\theta_1 \rangle$ is given by the average of $\cos\theta_1$ weighted over the collision rate integrated over all angles.

$$\left\langle \cos\theta_1 \right\rangle = \frac{\int \cos\theta_1 \frac{dn}{d\Omega_1} d\Omega_1}{\int \frac{dn}{d\Omega_1} d\Omega_1} \tag{A.8}$$

$$\Rightarrow \langle \cos\theta_1 \rangle = \frac{\int\limits_{-1}^{1} (\cos\theta_1 - \beta \cos^2\theta_1) d\cos\theta_1}{\int\limits_{-1}^{1} (1 - \beta \cos\theta_1) d\cos\theta_1} = \frac{\frac{1}{2} - \frac{\beta}{3} - \frac{1}{2} - \frac{\beta}{3}}{1 - \frac{\beta}{2} + 1 + \frac{\beta}{2}} = -\frac{\beta}{3}$$
(A.9)

Using this and the fact that $\beta \ll 1$, equation A.6 becomes

$$\Rightarrow \left\langle \frac{\Delta E}{E} \right\rangle = \frac{1 - \beta^2 / 3}{1 + \beta^2} - 1 = \frac{1 + \beta^2 / 3 - (1 - \beta^2)}{1 - \beta^2} \sim \frac{4}{3} \left(\frac{v}{c}\right)^2 \tag{A.10}$$

which is the fractional energy gain shown in equation 2.1.



A.1.2 First Order Fermi Acceleration

Figure A.2: Diagram of particle bouncing back and forth across a shock.

As in second order acceleration,

$$\Rightarrow \left\langle \frac{\Delta E}{E} \right\rangle = \frac{1 + \beta \left\langle \cos\theta_2' \right\rangle - \beta \left\langle \cos\theta_1 \right\rangle - \beta^2 \left\langle \cos\theta_1 \right\rangle \left\langle \cos\theta_2' \right\rangle}{1 - \beta^2} - 1 \qquad (A.11)$$

where $\beta = v_p/c$ and v_p is the speed of the ejected material. As shown in figure A.2, the values of θ_1 are always between 90° and 270° meaning $\cos\theta_1$ is always negative with the most likely collisions happening for $\theta_1 = 180^\circ$. This gives the collision rate equal to

$$\frac{dn}{d\Omega_1} \propto -\cos\theta_1 \tag{A.12}$$

Therefore, the average angle is given by

$$\left\langle \cos\theta_1 \right\rangle = \frac{\int\limits_{-1}^{0} -\cos^2\theta_1 d\cos\theta_1}{\int\limits_{-1}^{0} -\cos\theta_1 d\cos\theta_1} = \frac{\left(\frac{1}{3}\right)}{-\left(\frac{1}{2}\right)} = -\frac{2}{3}$$
(A.13)

The values of θ'_2 are always between 0 and 90° so $\cos\theta'_2$ is always positive and the greatest interaction probability occurs for $\theta'_2 = 0$. This gives

$$\frac{dn}{d\Omega_2'} \propto \cos\theta_2' \tag{A.14}$$

$$\Rightarrow \langle \cos\theta_2' \rangle = \frac{\int_0^1 \cos^2\theta_2' d\cos\theta_2'}{\int_0^1 \cos\theta_2' d\cos\theta_2'} = \frac{\left(\frac{1}{3}\right)}{\left(\frac{1}{2}\right)} = \frac{2}{3}$$
(A.15)

Assuming $\beta \ll 1$, equation A.11 becomes

$$\Rightarrow \left\langle \frac{\Delta E}{E} \right\rangle \sim 1 + \frac{2\beta}{3} + \frac{2\beta}{3} - 1 = \frac{4}{3} \left(\frac{v_p}{c} \right) \tag{A.16}$$

which is the fractional energy gain shown in equation 2.2.

A.1.3 Cosmic Ray Spectrum

The cosmic ray spectrum is derived using the probability of the particle escaping the shock, the fractional energy gain and the compression ratio in the shock. The compression ratio is given by

$$R = \frac{\rho_{\text{shocked}}}{\rho_{\text{unshocked}}} \tag{A.17}$$

where ρ_{shocked} is the density of material "inside" the shock (to the left in figure A.2) and $\rho_{\text{unshocked}}$ is the density outside the shock. The ratio of speed of the shock to the speed of the interstellar medium is given by

$$\frac{v_s}{v_p} = \frac{R}{R-1} \tag{A.18}$$

$$\Rightarrow v_p = \left(1 - \frac{1}{R}\right) v_s. \tag{A.19}$$

The fractional energy gain can then be written in terms of this ratio:

$$\frac{\Delta E}{E} = \frac{4}{3} \frac{v_p}{c} = \frac{4}{3} \left(1 - \frac{1}{R} \right) \frac{v_s}{c}.$$
(A.20)

As the particle travels back and forth across the shock, there is a small probability $P_{\rm esc}$ that it will escape the shock and travel to Earth. This escape probability is given by

$$P_{\rm esc} = \frac{4}{R} \frac{v_s}{v} \tag{A.21}$$

where v is the speed of the cosmic ray and $v \sim c$. The probability of the particle remaining in the shock after k crossings is given by

$$P(\geq k) = (1 - P_{\rm esc})^k.$$
 (A.22)

After k crossings, the energy of the particle is given by

$$E = E_0 \left(1 + \frac{\Delta E}{E} \right)^k \tag{A.23}$$

where E_0 is the initial energy of the particle. This can be rearranges to give

$$k = \frac{\ln(E/E_0)}{\ln(1 + \Delta E/E)}.$$
(A.24)

Since the energy after k crossings, E, increases with k, the number of cosmic rays remaining after at least k crossings will be the number with energy greater than or equal to E. This is known as the integral spectrum and is given by

$$N(\geq E) = N_0 (1 - P_{\rm esc})^k$$
 (A.25)

where N_0 is the initial number of particles. Dividing by N_0 and taking the log of both sides gives

$$\ln(N/N_0) = k \ln(1 - P_{\rm esc}) = \frac{\ln(E/E_0) \ln(1 - P_{\rm esc})}{\ln(1 + \Delta E/E)}.$$
 (A.26)

This implies

$$N(\geq E) = N_0 \left(\frac{E}{E_0}\right)^{-\Gamma} \propto E^{-\Gamma}$$
(A.27)

where Γ is given by

$$\Gamma = -\frac{\ln(1 - P_{\rm esc})}{\ln(1 + \Delta E/E)}.$$
(A.28)

The values $P_{\rm esc}$ and $\Delta E/E$ are assumed to be small meaning the logarithms can be approximated as $\ln(1 - P_{\rm esc}) \sim -P_{\rm esc}$ and $\ln(1 + \Delta E/E) \sim \Delta E/E$. This gives

$$\Gamma = \frac{P_{\text{esc}}}{\Delta E/E} = \frac{\left(\frac{4v_s}{Rv}\right)}{\frac{4}{3}\left(1 - \frac{1}{R}\right)\frac{v_s}{c}}.$$
(A.29)

Since $v \sim c$, this becomes

$$\Gamma = \frac{\left(\frac{4v_s}{Rc}\right)}{\frac{4}{3}\left(1 - \frac{1}{R}\right)\frac{v_s}{c}} = \frac{4v_s}{\frac{4}{3}(R-1)v_s} = \frac{3}{R-1}.$$
(A.30)

For strong shock fronts, the compression ratio is approximately 4, meaning $\Gamma = 1$. Therefore, the integral spectrum is given by

$$N(\geq E) \propto E^{-1}.\tag{A.31}$$

The differential spectrum is found by differentiating $N(\geq E)$ which gives

$$N(E) \propto E^{-2} \tag{A.32}$$

as shown in equation 2.3.

A.2 Geometric Time for SPE and MPE



Figure A.3: Diagram of muon track and Cherenkov cone showing distance muon travels before light is emitted (r_{μ}) and the distance light travels to DOM (r_{γ}) .

The geometric time shown in equation 4.7 is given by t_0 plus the time the muon takes to travel r_{μ} plus the time the photon takes to travel r_{γ} as shown in figure A.3.

$$t_{\rm geo} = t_0 + \frac{r_\mu}{c} + \frac{nr_\gamma}{c} \tag{A.33}$$

where n is the refractive index of the ice, the muon is travelling at speed c and the photon has speed c/n. As shown in figure A.3, the hypotenuse r_{γ} is simply given by

$$r_{\gamma} = \frac{d}{\sin\theta_c}.\tag{A.34}$$

The angle between $(\vec{r_i} - \vec{r_0})$ and \hat{p} is given by α . As \hat{p} is a unit vector, this implies

$$\hat{p} \cdot (\vec{r_i} - \vec{r_0}) = |\vec{r_i} - \vec{r_0}|\cos\alpha.$$
 (A.35)

By the definition of $\cos \alpha$, the side adjacent to α is given by

$$r_{\mu} + \frac{d}{\tan\theta_c} = |\vec{r_i} - \vec{r_0}|\cos\alpha = \hat{p} \cdot (\vec{r_i} - \vec{r_0}).$$
(A.36)

$$\Rightarrow r_{\mu} = \hat{p} \cdot (\vec{r_i} - \vec{r_0}) - \frac{d}{\tan\theta_c}.$$
 (A.37)

$$\Rightarrow t_{\text{geo}} = t_0 + \frac{\hat{p} \cdot (\vec{r_i} - \vec{r_0}) - \frac{d}{\tan\theta_c} + \frac{nd}{\sin\theta_c}}{c}.$$
 (A.38)

By the definition of θ_c where $\cos\theta_c = 1/n$ and the definition of $\tan\theta_c = \sin\theta_c/\cos\theta_c$, this gives

$$\frac{nd}{\sin\theta_c} - \frac{d}{\tan\theta_c} = \frac{d}{\sin\theta_c \cos\theta_c} - \frac{d\cos\theta_c}{\sin\theta_c}$$
(A.39)

$$=\frac{d(1-\cos^2\theta_c)}{\sin\theta_c\cos\theta_c}=\frac{d\sin^2\theta_c}{\sin\theta_c\cos\theta_c}=\frac{d\sin\theta_c}{\cos\theta_c}=d\tan\theta_c.$$
 (A.40)

Finally, substituting into equation A.38:

$$t_{\rm geo} = t_0 + \frac{\hat{p} \cdot (\vec{r_i} - \vec{r_0}) + d \tan \theta_c}{c}$$
(A.41)

as shown in equation 4.7.

A.3 Space Angle Formula

The space angle between two vectors describes how far apart they are on the surface of a sphere. This angle is simply the angle α between two vectors \vec{u} and \vec{v} given by their dot product.

$$\cos\alpha = \vec{u} \cdot \vec{v} \tag{A.42}$$

A unit radial vector in spherical coordinates is given by

$$\vec{u} = \sin\theta \cos\phi \hat{\mathbf{x}} + \sin\theta \sin\phi \hat{\mathbf{y}} + \cos\theta \hat{\mathbf{z}} \tag{A.43}$$

where $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ and $\hat{\mathbf{z}}$ are the unit vectors in the x, y and z directions, θ is the zenith and ϕ is the azimuth as defined before. The unit radial vectors for the true direction and a fit are shown in the example in figure A.4.

The dot product between these two vectors is given by

$$\vec{u} \cdot \vec{v} = \sin\theta_1 \cos\phi_1 \sin\theta_2 \cos\phi_2 + \sin\theta_1 \sin\phi_1 \sin\theta_2 \sin\phi_2 + \cos\theta_1 \cos\theta_2 = \cos\alpha \quad (A.44)$$

This can be simplified to

$$\cos\alpha = \sin\theta_1 \sin\theta_2 \cos(\phi_2 - \phi_1) + \cos\theta_1 \cos\theta_2 \tag{A.45}$$

This formula reflects the nature of spherical coordinates in a way that zenith and azimuth alone can't. For example, if two vectors have the same zenith of 1° but are separated 180° in azimuth, the space angle between them will only be 2° since the azimuths are much closer together at the poles. On the other hand, if these vectors were on the equatorial plane and separated 180° in azimuth, the space angle would be 180° . This formula for space angle was used to asses the accuracy of the fits in chapters 5 and 7.



Figure A.4: Diagram of two vectors in 3-dimensional space with the space angle α between them shown.

A.4 Chi Squared Critical Value

The Chi Squared (χ^2) distribution is a probability density function given by

$$p_x(n) = \frac{\left(\frac{x}{2}\right)^{\frac{n}{2}-1}e^{-\frac{x}{2}}}{2\Gamma(\frac{n}{2})} \tag{A.46}$$

where n is the number of degrees of freedom in the variable x [43, p. 36]. The gamma function of a variable k, $\Gamma(k)$ is equal to the factorial of k, k!, for integer values of k. For the likelihood grid scans, there are two degrees of freedom (zenith and azimuth) so n = 2 which means $\Gamma(\frac{n}{2}) = 1! = 1$. The probability density function for these scans is then

$$p_x(2) = \frac{\left(\frac{x}{2}\right)^{1-1} e^{-\frac{x}{2}}}{2} = \frac{1}{2} e^{-\frac{x}{2}}$$
(A.47)

The probability below some critical value x_c is given by the integral

$$P(x < x^{c}) = \frac{1}{2} \int_{0}^{x_{c}} e^{-\frac{x}{2}} dx = 1 - e^{-\frac{x_{c}}{2}}$$
(A.48)

To find the value of σ , the integral above must be equal to 68% as, in any probability distribution, 68% of the values should lie within 1σ of the mean. This gives the following equation to be solved.

$$1 - e^{-\frac{x_c}{2}} = 0.68 \tag{A.49}$$

$$\Rightarrow e^{-\frac{x_c}{2}} = 0.32 \tag{A.50}$$

$$\Rightarrow x_c = -2\ln(0.32) \sim 2.30. \tag{A.51}$$

This is contrasted with the case for one degree of freedom in which the critical value is about 1. When calculating the σ values in chapters 5 and 7, this critical value for two degrees of freedom is halved since it is $-2\ln\mathcal{L}$ which follows the χ^2 distribution and $-\ln\mathcal{L}$ which is calculated by millipede. Therefore, σ is defined by an increase of 2.30/2 = 1.15 units in negative log likelihood.

A.5 Perpendicular Plane Coordinates

For the likelihood scan over a perpendicular plane, shown in section 6.2.2, each vertex on the plane is given two coordinates x' and y' which correspond to a vertex (x, y, z) in IceCube coordinates. The two coordinates, x' and y', are defined by the spherical unit vectors in the azimuthal and zenith directions given by

$$\vec{x}' = x'\hat{\phi} = -x'\sin\phi\hat{\mathbf{x}} + x'\cos\phi\hat{\mathbf{y}} \tag{A.52}$$

$$\vec{y}' = y'\hat{\theta} = y'\cos\theta\cos\phi\hat{\mathbf{x}} + y'\cos\theta\sin\phi\hat{\mathbf{y}} - y'\sin\theta\hat{\mathbf{z}}$$
(A.53)

where θ is the track's zenith and ϕ is the azimuth. The vectors \vec{x}' and \vec{y}' are the horizontal and vertical displacements from the centre of the plane given by (x_0, y_0, z_0) . In the case shown in section 6.2.2, this centre point is the fitted vertex with the minimum negative log likelihood. Each point (x, y, z) on the plane is simply given by

$$(x, y, z) = (x_0, y_0, z_0) + \vec{x}' + \vec{y}'$$
(A.54)

Each coordinate is then calculated from

$$x = x_0 - x'\sin\phi + y'\cos\theta\cos\phi \qquad (A.55)$$

$$y = y_0 + x'\cos\phi + y'\cos\theta\sin\phi \tag{A.56}$$

$$z = z_0 - y' \sin\theta \tag{A.57}$$

These equations were used to translate the two dimensional plane coordinates into IceCube coordinates which are understood by millipede.

Appendix B

Technical Details of Simulation and Reconstruction

In this section, the software used for event reconstruction and simulation will be discussed in more detail. The methods used to produce input waveforms from the millipede likelihood function (as in section 7.1) and adding Poisson fluctuations and after-pulses (as in sections 7.3.1 and 7.3.2) will also be explained.

B.1 The IceTray Software

Event reconstruction for IceCube is achieved with a software framework known as "IceTray" which is written in C++. In IceTray, information about an event is held in a "physics frame" which consists of numerous frame objects containing the event properties such as the raw data, the pulses and any reconstructions. The pulses are saved as an object of type I3RecoPulseSeriesMap. This data type contains a list of all illuminated DOMs and for each DOM, a list of the pulses which make up the waveform. Each pulse has a given time, width (both in ns) and total charge (photo-electron count.) This "pulse series map" is the final stage in data processing which is read into the reconstruction algorithm.

The reconstructions are stored in the frame with the I3Particle data type. This data type contains the physical properties of the reconstructed particle such as the type (e.g. MuMinus), energy (in GeV), arrival direction (θ, ϕ in radians), vertex position (x, y, z in metres) and vertex time (in nanoseconds.)

B.2 The MC Tree

IceCube simulations are performed using what is known as a Monte Carlo (MC) simulation which is commonly used to model complex systems. This simulates a particle defined by a set of parameters describing its velocity, position, energy,

time etc. Eventually, a random process is initiated which has an associated probability based on known physical laws. This will then generate a new set of particle parameters which may experience a different set of processes [44]. If the process is repeated for long enough, this would accurately model a physical system of particles.

A large sample of previously simulated muon-neutrino events was obtained for millipede testing which contained a mix of ν_{μ} and $\bar{\nu}_{\mu}$ events. For these simulated events, the particle track is defined by an MC "Tree" of type I3MCTree which contains the parameters of the original particle and the daughter products which branch off as it loses energy. The tree starts off as a single ν_{μ} or $\bar{\nu}_{\mu}$ stored in an I3Particle. This particle then interacts in the ice by the charged current interaction to produce a μ^- or μ^+ and a hadronic cascade started by the nucleon it encountered.

The muon and hadronic cascade are considered the "daughter" products of the neutrino and they, in turn, interact in the ice to produce their own daughter products. These secondary products include photons which can cause pair production (production of an electron and positron) which will then produce its own signal. The hadronic cascade daughter products can even include muons but these are of such low energy, their signal is very small compared to the original muon. The signal observed in the detector is a combination of all of these interactions but the true muon energy and direction can be extracted from the muon part of the tree.



Figure B.1: Diagram of Monte Carlo Tree of particles produced by a neutrino interaction.

From the tree of particles generated, as well as a knowledge of the optical properties of the ice, the light signal in the DOMs can be calculated. This signal now has the same characteristics as real signal and can be used for direction and energy reconstructions.

B.3 Using the Millipede Likelihood Function as the Simulation

To use the millipede likelihood function as a simulation, the millipede algorithm is given a series of energy losses as a list of particles. Each particle in this list has a given time, energy and position which record how much energy is lost and the location of the loss in space and time.

To create signal from these losses, a new pulse map needs to be defined for the event. In the original simulated waveforms, a pulse map contains the series of pulses in time (each with a start time, width and photo-electron count) for all illuminated DOMs. For the simulation with the millipede likelihood function the map is predefined with a series of pulses, each containing zero charge, for each DOM which defines the time frame in which each DOM is likely to see light. For each DOM a time scale is decided based on how far the DOM is from its nearest loss and when that loss occurred.

For each DOM in the detector, the time the muon passed closest to the DOM is calculated and the loss immediately preceding this time is chosen. The earliest time light can arrive at the DOM is defined as the time it takes to travel along the muon track at the speed of light c and then travel to the DOM at the local speed of light c/n at the Cherenkov angle θ_c . This is shown in figure B.2.

The distances r_{μ} and r_{γ} are defined in the following way: the muon travels a distance r_{μ} along its track until the angle between the track and the distance vector between the muon and the DOM is equal to the Cherenkov angle θ_c . The distance light then needs to travel to the DOM is r_{γ} . If the energy loss occurs at time t = 0, the earliest time a photon can arrive at the DOM is given by

$$t_{\gamma} = \frac{r_{\mu} + nr_{\gamma}}{c} \tag{B.1}$$

The values of r_{μ} and r_{γ} are determined by the relations shown in equations A.37 and A.34 respectively. For each DOM that is expected to see light, a series of zero charge pulses is defined where the pulses are all 25 ns wide and extend 7000 ns from the earliest time defined above. It should be noted that this method of defining the pulse map is for the most general loss patterns. If the true energy losses are used, the pulse map should be similar to the map for the original pulses



Figure B.2: Diagram of muon track showing relative positions of energy loss and DOM.

so in this particular case, the original pulses can be used to define the binning.

Using the pulse map and the positions and times of the losses, millipede can create a response matrix which determines the contributions of each loss toward the photon count for each time bin as mentioned in chapter 4. These contributions are in units of photo-electrons/GeV and are determined by the optical properties of the ice and relative positions of the losses and the DOMs. An example track and loss pattern is shown in the diagram in figure B.3. As in the prediction calculation shown in section 4.4.2, the waveform vector \vec{N} is given by

$$\vec{N} = \Lambda \vec{E} \tag{B.2}$$

where Λ is the response matrix and \vec{E} is the energy loss vector.

This signal vector now contains the waveforms of all DOMs joined into a single vector. However, the point in the vector at which one waveform ends and another begins is recorded in the algorithm. These start and end points are used to convert the signal vector into a map containing the pulse series for each illuminated DOM and identification of the corresponding DOM. The new waveforms are now in a format which can be read back into the millipede reconstruction algorithm.



Figure B.3: Diagram of muon track showing light from losses arriving at DOMs.

B.4 Added Poisson Fluctuations

A Poisson deviate has probability given by

$$P(k) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{B.3}$$

where k is the variable being sampled and λ is the expected value or mean of k. The Poisson deviate is used to add fluctuations to the existing waveform where the mean of the distribution is taken as the original waveform value.

The millipede predictions used in section 7.1 are now used as means in Poisson distributions to produce a new value randomly taken from the distribution. In this method, a random number r_0 is generated between 0 and 1. If the random number is larger than e^{-q} where q is the original charge (photo-electron count), another random number r_1 is generated and multiplied to the first random number. This continues until the resultant product falls below e^{-q} . With larger charge values e^{-q} will be smaller meaning more steps are required to drop below this number. The Poisson deviate is taken as the number of steps required to reach this point [45, p. 293-294]. In other words, the new charge value q_{new} is the smallest value of m such that

$$\prod_{i=0}^{m} r_i < e^{-q} \tag{B.4}$$

where r_i is the *i*-th random number generated between 0 and 1.

B.5 Added After-pulses

To apply after-pulses to the predicted waveform, as was achieved section 7.3.2, the first step is to determine how many after-pulses (if any) there will be. The probability for k after-pulses occurring due to an incident pulse of n photo-electrons is given by a Poissonian [46].

$$P_{\mu}(n,k) = \frac{(n\mu)^k}{k!} e^{-n\mu}$$
(B.5)

where μ is the average number of ionisations per photo-electron. For each pulse in the predicted waveform, a random number is generated and the probability shown above determines how many after-pulses are produced. The value of μ for these tests is chosen to be 0.2.

The time delay between the main pulse and after-pulse depends on the type of gas being ionised and the electric field strength in the PMT. For a pair of dynodes spaced Δx apart, the electric field strength will be $V/\Delta x$ where V is the potential difference. The acceleration of the ion with charge q and mass m is

$$a = \frac{qV}{m\Delta x} \tag{B.6}$$

The distance travelled between dynodes can be expressed as

$$\Delta x = \frac{1}{2}at^2 \tag{B.7}$$

The time for the ion to travel between dynodes starting at rest is then given by

$$t = \sqrt{\frac{2\Delta x}{a}} = \sqrt{\frac{2m(\Delta x)^2}{qV}}$$
(B.8)

The relative abundances of the different ions are chosen based on [40]. In this case, a random number is generated again and, based on where it sits between 0 and 1 and the probability of each ion, this determines the type of ion, and hence the mass and charge to be used to calculate the time delay. The time delay is calculated using equation B.8 and a pulse of 1 photo-electron is placed this amount of time after its corresponding main pulse. This is repeated across all main pulses in all DOMs to produce the new average pulse plot shown in figure 7.23(b).

Appendix C

Reconstruction Python Code

The reconstruction was performed using the reconstruction software "IceRec" in version V14-11-00.

C.1 Millipede Free Fit

The code below shows the calls to the reconstruction functions for the full free fit and the values of the settings.

```
#load spline tables containing ice model
muon_service = photonics_service.I3PhotoSplineService('emu_abs.fits',
 'emu prob.fits', 0)
cascade_service = photonics_service.I3PhotoSplineService
('ems_mie_z20_a10.abs.fits', 'ems_mie_z20_a10.prob.fits', 0)
tray = I3Tray()
#Import file
tray.AddModule('I3Reader', 'reader', FilenameList=[input i3file])
#Define step sizes for vertex and direction
VertexStep = 5*I3Units.m
angleStep = 5*I3Units.deg
tray.AddService('MillipedeLikelihoodFactory', 'millipedellh',
    MuonPhotonicsService=muon_service,
     CascadePhotonicsService=cascade service,
    PhotonsPerBin=100, MuonRegularization=0, ShowerRegularization=0,
    #Set bad or saturated DOMs which will not be included
    ExcludedDOMs=['BadDomsList', 'CalibrationErrata', 'SaturatedDOMs',
     'BrightDOMs'],
    #Set time window and input waveforms
```

```
ReadOutWindow='OfflinePulsesTimeRange',
    Pulses='OfflinePulses')
tray.AddService('I3GSLRandomServiceFactory', 'I3RandomService')
tray.AddService('I3GulliverMinuitFactory', 'minuit',
    MaxIterations=2000, MinuitPrintLevel=1)
tray.AddService('MuMillipedeParametrizationFactory', 'millipedeparam',
#Set step sizes of position, time and direction and spacing of energy losses
    MuonSpacing=0, ShowerSpacing=10, StepT=15*I3Units.ns,
    StepX=VertexStep, StepY=VertexStep, StepZ=VertexStep,
    StepZenith=angleStep, StepAzimuth=angleStep)
#Load first guess track
tray.AddService('I3BasicSeedServiceFactory', 'seed',
FirstGuess='PoleMuonLlhFit',
    TimeShiftType='TNone')
#Perform reconstruction and store result in frame
tray.AddModule('I3SimpleFitter', 'MillipedeFit', SeedService='seed',
    Parametrization='millipedeparam', LogLikelihood='millipedellh',
    Minimizer='minuit')
```

C.2 Grid Scan

To perform a grid scan over zenith and azimuth, for each direction, the result from the free fit is used and the zenith and azimuth are set to new values.

```
def setDirection(frame):
    #Load previously defined directions
    global theta,phi

    #Define new particle
    seedTrack = dataclasses.I3Particle()
    #Set Direction
    seedTrack.dir = dataclasses.I3Direction(theta,phi)
    #Set other parameters to free fit values
    seedTrack.pos = frame['MillipedeFit'].pos
    seedTrack.time = frame['MillipedeFit'].time
    seedTrack.fit_status = frame['MillipedeFit'].fit_status
    #Save result
    frame['forcedDir'] = seedTrack
    return True
```

tray.AddModule(setDirection, 'setdirection')

#Define vertex step
VertexStep = 5*I3Units.m

#Zenith and azimuth are fixed so angle step is zero
angleStep = 0

The reconstruction is then performed as shown in section C.1 using the new parameters from 'forcedDir' as the seed.

References

- [1] Esteban Roulet. Neutrino astrophysics. *astro-ph/0011570*, 2000. Lectures at the ICTP Summer School on Particle Physics and Cosmology.
- [2] F. Reines, C. L. Cowan, F. B. Harrison, A. D. McGuire, and H. W. Kruse. Detection of the free anti-neutrino. *Phys. Rev.*, 117:159–173, 1960.
- [3] Maria J. Herrero. The Higgs System in and Beyond the Standard Model. In Lectures on Particle Physics, Astrophysics and Cosmology, pages 188–252. Springer, 2015.
- [4] D. J. Bird et al. Evidence for correlated changes in the spectrum and composition of cosmic rays at extremely high energies. *Phys. Rev. Lett.*, 71:3401–3404, Nov 1993.
- [5] M. Bustamante et al. High-energy cosmic-ray acceleration. Proceedings of the 5th CERN - Latin-American School of High-Energy Physics, Recinto Quirama (Medellin), Antioquia, Colombia, 2009.
- [6] Antonella Castellina and Fiorenza Donato. Astrophysics of Galactic charged cosmic rays. Planets, Stars and Stellar Systems: Volume 5: Galactic Structure and Stellar Populations, pages 725–788, 2013.
- [7] A. A. Watson. High-energy cosmic rays and the Greisen–Zatsepin–Kuz'min effect. *Reports on Progress in Physics*, 77(3):036901, 2014.
- [8] Ulrich F. Katz and C. Spiering. High-energy neutrino astrophysics: Status and perspectives. Progress in Particle and Nuclear Physics, 67(3):651–704, 2012.
- [9] G. Bellini, L. Ludhova, G. Ranucci, and F. L. Villante. Neutrino oscillations. *Advances in High Energy Physics*, 2014.
- [10] Invisibles. The ceaseless transformation of the three neutrinos, 2014. http://invisibles.eu/outreach/entry/ ceaseless-transformation-three-neutrinos.

- [11] Boris Mikhailovich Bolotovskii. Vavilov-Cherenkov radiation: its discovery and application. *Physics-Uspekhi*, 52(11):1099–1110, 2009.
- [12] A. A. Watson. Extensive Air Showers and Ultra High Energy Cosmic Rays. Lezioni tenute alla" Escuela Mexicana de Astrofísica, 2002.
- [13] Andreas Haungs, Heinigerd Rebel, and Markus Roth. Energy spectrum and mass composition of high-energy cosmic rays. *Reports on Progress in Physics*, 66(7):1145, 2003.
- [14] M. G. Aartsen et al. Search for time-independent neutrino emission from astrophysical sources with 3 yr of IceCube data. *The Astrophysical Journal*, 779(2):132, 2013.
- [15] Tyce DeYoung and IceCube Collaboration. Results from seven years of AMANDA-II. In *Journal of Physics: Conference Series*, volume 136, page 022046. IOP Publishing, 2008.
- [16] Rasha Abbasi et al. The design and performance of IceCube DeepCore. Astroparticle Physics, 35(10):615–624, 2012.
- [17] Rasha Abbasi et al. IceTop: The surface component of IceCube. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 700:188–220, 2013.
- [18] Francis Halzen. IceCube Neutrinos: From GeV to PeV. Nuclear Physics B -Proceedings Supplements, 246–247(0):3 – 17, 2014. Proceedings of the 9th International Symposium on Cosmology and Particle Astrophysics.
- [19] Abraham Achterberg et al. First year performance of the IceCube neutrino telescope. Astroparticle Physics, 26(3):155–173, 2006.
- [20] R. Abbasi et al. Calibration and characterization of the IceCube photomultiplier tube. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 618(1):139–152, 2010.
- [21] Marius Wallraff. Design, implementation and test of a new feature extractor for the icecube neutrino observatory. *Rheinisch-Westfälische Technische Hochschule (RWTH)*, *Aachen*, 2010.
- [22] R. Stokstad. Design and performance of the IceCube electronics. In Proceedings, eleventh Workshop on Electronics for LHC and Future Experiments, Heidelberg, Germany, 12-16 September 2005, page 4, 2005.
- [23] Troy D. Straszheim and the IceCube Collaboration. WaveCalibrator. http://software.icecube.wisc.edu/offline/projects/ WaveCalibrator/index.html.

- [24] M. G. Aartsen et al. Energy reconstruction methods in the IceCube neutrino telescope. *Journal of Instrumentation*, 9(03):P03009, 2014. Note: Section 8 introduces millipede's method of reconstructing energy losses, though the word 'millipede' is not used.
- [25] The IceCube Collaboration. The detection of neutrinos in icecube. https: //masterclass.icecube.wisc.edu/en/learn/detecting-neutrinos.
- [26] Markus Ackermann et al. Optical properties of deep glacial ice at the South Pole. Journal of Geophysical Research: Atmospheres (1984–2012), 111(D13), 2006.
- [27] M. G. Aartsen et al. Measurement of South Pole ice transparency with the IceCube LED calibration system. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 711:73–89, 2013.
- [28] Dmitry Chirkin. Evidence of optical anisotropy of the South Pole ice. In 33rd International Cosmic Ray Conference, The Astroparticle Physics Conference, 2013.
- [29] Anthony M. Brown et al. Latest results from the IceCube experiment. arXiv preprint arXiv:1012.1633, 2010.
- [30] M. G. Aartsen et al. Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector. *Science*, 342(6161):12428562–67, 2013.
- [31] R. Abbasi et al. Measurement of the atmospheric neutrino energy spectrum from 100 GeV to 400 TeV with IceCube. *Physical Review D*, 83(1):19, 2011.
- [32] M. G. Aartsen et al. Observation of high-energy astrophysical neutrinos in three years of IceCube data. *Physical Review letters*, 113(10):101101, 2014.
- [33] J. Ahrens et al. Muon track reconstruction and data selection techniques in AMANDA. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, 524(1):169–194, 2004.
- [34] D. Gora, A. Kappes, D. Altmann, and T. Glunsenkamp. Studies on Millipede. Internal Report icecube/201404001, IceCube Collaboration, 2014.
- [35] Emanuel Jacobi. Release Notes. http://software.icecube.wisc.edu/ icerec_trunk/projects/millipede/release_notes.htmll, June 29, 2013.
- [36] David Altmann and Thorsten Glüsenkamp. Studies for performance in reconstruction and time consumption, 2012. Presented at the IceCube Fall Collaboration Meeting, Aachen, Germany.

- [37] The IceCube Collaboration. 2011 Level2 file structure. https: //wiki.icecube.wisc.edu/index.php/2011_Level2_file_structure.
- [38] Mark R. Calabretta and Boudewijn F. Roukema. Mapping on the healpix grid. Monthly Notices of the Royal Astronomical Society, 381(2):865–872, 2007.
- [39] Nathan Whitehorn and Claudio Kopper. Reconstruction Results Episode 2, 2012. Presented at the IceCube Diffuse Call 14 November 2012.
- [40] P. B. Coates. The origins of afterpulses in photomultipliers. Journal of Physics D: Applied Physics, 6(10):1159, 1973.
- [41] Leif R\u00e4del and Sebastian Schoenen. Update on millipede scans for high-energy events, 2015. Presented at the IceCube Fall Collaboration Meeting, Copenhagen, Denmark.
- [42] M. Kachelriess. Lecture notes on high energy cosmic rays. arXiv preprint arXiv:0801.4376, 2008.
- [43] Christian Walck. Handbook on statistical distributions for experimentalists. University of Stockholm Internal Report SUF-PFY/96-01, available from www.physto.se/~walck, 2007.
- [44] Nicholas Metropolis and Stanislaw Ulam. The monte carlo method. *Journal* of the American statistical association, 44(247):335–341, 1949.
- [45] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. Numerical Recipes in C (2nd Ed.): The Art of Scientific Computing. Cambridge University Press, New York, NY, USA, 1992.
- [46] Julia Haser et al. Afterpulse measurements of R7081 photomultipliers for the Double Chooz experiment. *Journal of Instrumentation*, 8(04):P04029, 2013.