

A Multi-Variate Fit to the Chemical Composition of the Cosmic-Ray Spectrum

By
Jonathan Eisch

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Physics)

at the
UNIVERSITY OF WISCONSIN-MADISON
2014

Date of final oral examination: August 8, 2014

The dissertation is approved by the following members of the Final Oral Committee:

Stefan Westerhoff, Professor, Physics

Albrecht Karle, Professor, Physics

Justin Vandenbroucke, Assistant Professor, Physics

Michael DuVernois, Senior Scientist, Physics

Ellen Zweibel, Professor, Astronomy

A MULTI-VARIATE FIT TO THE CHEMICAL COMPOSITION OF THE COSMIC-RAY SPECTRUM

Jonathan Eisch

Under the supervision of Professor Stefan Westerhoff

At the University of Wisconsin-Madison

Since the discovery of cosmic rays over a century ago, evidence of their origins has remained elusive. Deflected by galactic magnetic fields, the only direct evidence of their origin and propagation remain encoded in their energy distribution and chemical composition. Current models of galactic cosmic rays predict variations of the energy distribution of individual elements in an energy region around 3×10^{15} eV known as the *knee*. This work presents a method to measure the energy distribution of individual elemental groups in the *knee* region and its application to a year of data from the IceCube detector.

The method uses cosmic rays detected by both IceTop, the surface-array component, and the deep-ice component of IceCube during the 2009-2010 operation of the IC-59 detector. IceTop is used to measure the energy and the relative likelihood of the mass composition using the signal from the cosmic-ray induced extensive air shower reaching the surface. IceCube, 1.5 km below the surface, measures the energy of the high-energy bundle of muons created in the very first interactions after the cosmic ray enters the atmosphere. These event distributions are fit by a constrained model derived from detailed simulations of cosmic rays representing five chemical elements. The results of this analysis are evaluated in terms of the theoretical uncertainties in cosmic-ray interactions and seasonal variations in the atmosphere.

The improvements in high-energy cosmic ray hadronic-interaction models informed by this analysis, combined with increased data from subsequent operation of the IceCube detector, could provide crucial limits on the origin of cosmic rays and their propagation through the galaxy.

In the course of developing this method, a number of analysis and statistical techniques were developed to deal with the difficulties inherent in this type of measurement. These include a composition-sensitive air shower reconstruction technique, a method to model simulated event distributions with limited statistics, and a method to optimize and estimate the error on a regularized fit.

Stefan Westerhoff

ACKNOWLEDGMENTS

I dedicate this work to my son, Paxton, whose imminent arrival encouraged countless late nights of hard work, and to my wonderful wife and partner, Vanessa, for her love and encouragement through all the ups and downs of this effort. Without her constant support this work would not have been possible.

My thanks go to my parents, Dan and Connie, for providing all the opportunities for me to discover the world around me and my place in it. Also my siblings for all their support and introducing me to science and computers at an early age, giving me a head start toward my chosen carrier.

I'm grateful to my advisor, Stefan Westerhoff, for his guidance in dealing with such an intractable analysis topic and for his extensive notes on the many revisions of this work. I would like to thank the IceCube group at Madison, who have been wonderful to work with and have made living in Madison such a great experience. I would particularly like to thank Karen Andeen and Albrecht Karle for the thoughtful discussions that lead me to develop many of the methods used in this work.

DISCARD THIS PAGE

TABLE OF CONTENTS

	Page
LIST OF TABLES	v
LIST OF FIGURES	vi
1 Cosmic Rays	1
1.1 Introduction	1
1.2 History of Cosmic Rays	2
1.2.1 Discovery of Cosmic Rays	2
1.2.2 Discovery of Extensive Air Showers	3
1.3 Energy Spectrum	5
1.3.1 The Knee	8
1.4 Composition	9
1.5 Acceleration and Sources	11
1.5.1 Sources	14
1.6 Recent Measurements	17
2 Extensive Air Showers and Their Detection	22
2.1 Initial Interactions	22
2.2 Hadronic Component	23
2.3 Electromagnetic Component	24
2.4 Muons	27
2.5 Extensive Air-Shower Detection Techniques	32
2.5.1 Cherenkov Radiation	32
2.6 Recent Experiments	34
2.6.1 Water-Cherenkov Arrays	34
2.6.2 Ground-based Particle Detectors	35
3 IceCube Detector	36
3.1 Detector Array Geometry	36
3.2 Neutrino Detection	41
3.3 Detection of Muons	46
3.4 Digital Optical Modules	47
3.5 Ice Properties	50
3.6 IceTop Array	52
3.6.1 IceTop Tanks	52
3.6.2 IceTop Tank Response to Air Showers	52
3.6.3 Calibration	55

	Page
4 Simulation	57
4.1 Simulating Extensive Air Showers	57
4.1.1 CORSIKA	58
4.1.2 High Energy Hadronic Interaction Models	59
4.1.3 Atmosphere Models	59
4.2 Resampling	60
4.3 Muon Propagation	60
4.4 Simulating IceCube Detector Response	63
4.4.1 Photonics	63
4.5 Simulating IceTop Detector Response	63
4.6 Simulation Mass-Production	65
5 Reconstruction of Cosmic-Ray Air Showers	66
5.1 IceTop	66
5.1.1 Pulse Extraction	66
5.1.2 Event Cleaning	67
5.1.3 Direction Estimation	67
5.2 Likelihood-based Reconstruction	68
5.2.1 ShowerLLH	69
5.3 IceCube	75
5.3.1 Pulse Extraction	75
5.3.2 Event Cleaning	75
5.3.3 LineFit	75
5.3.4 Likelihood-based Reconstructions	76
5.3.5 MuEx	77
6 Analysis	82
6.1 Discussion of Previous Analyses	82
6.2 Analysis Overview	84
6.3 Data Sample	87
6.4 Event Selection	87
6.5 Building Distributions	88
6.5.1 Simulation Energy Bin Definitions	89
6.6 Simulation Weighting	92
6.7 Fitting Simulation	92
6.7.1 Effective Area	92
6.7.2 Distribution Mean, Variance and Covariance	95
6.7.3 Smoothing Simulation Parameter Distributions	96
6.8 Unfolding Algorithm	107
6.8.1 Least Squares	107
6.8.2 Constrained Least Squares	108
6.9 Performance on a Simulated Spectrum	108
6.10 Error Estimation	114
6.11 Regularization	117
6.11.1 Optimization	121

	Page
7 Results	125
7.1 Seasonal Corrections	125
7.2 Results	135
7.3 Systematic Uncertainties	140
7.3.1 Empirical Data-Simulation Mismatch	140
7.3.2 Hadronic Interaction Models	140
7.4 Summary	149
8 Conclusions	153
8.1 Looking Forward	158
LIST OF REFERENCES	159

DISCARD THIS PAGE

LIST OF TABLES

Table	Page
4.1 CORSIKA minimum kinetic-energy cuts used in simulation	59
4.2 CORSIKA July 1, 1997 South Pole atmosphere parameters	60
6.1 Cosmic-ray filters for IC59.	88
6.2 Quality cuts	88
6.3 Dimensions for unfolding tables	89

DISCARD THIS PAGE

LIST OF FIGURES

Figure	Page
1.1 Rossi's triode coincidence circuit.	3
1.2 Cosmic-ray spectrum as measured by Pierre Auger in 1939.	4
1.3 The cosmic-ray all-particle spectrum as measured by balloon- and ground-based experiments.	6
1.4 The spectrum of selected elemental components of the cosmic-ray flux	7
1.5 The Poly-Gonato model	10
1.6 The H4a model	10
1.7 Model mean natural log of the mass number predictions	11
1.8 Distribution of heavy elements in the SNR Cassiopea A from Chandra.	15
1.9 Possible sites of ultra-high energy cosmic-ray acceleration.	16
1.10 KASCADE-Grande composition result	17
1.11 IT/IC-40 mean natural logarithm of the mass result	19
1.12 IT/IC-40 all-particle spectrum result result	20
1.13 IC-73 composition tests	21
2.1 Electromagnetic particle density	26
2.2 Muon bundle particle energy	28
2.3 Relative number of electrons and high-energy muons	30
2.4 Muon particle density	31
2.5 2-d model of Cherenkov radiation	33
3.1 IceCube array schematic	38
3.2 IC-59 string distances	39
3.3 IceCube angular coordinate system	40

Figure	Page
3.4 Neutrino-nucleon interaction Feynman diagrams	42
3.5 Neutrino-induced cascade event	43
3.6 Neutrino-induced track event	44
3.7 Neutrino-induced tau simulation	45
3.8 Muon energy loss	46
3.9 Digital Optical Module schematic	48
3.10 Dust concentrations in the ice	51
3.11 IceTop tanks during deployment	53
3.12 IceTop tank cross-sectional view	54
3.13 VEM calibration spectrum	56
4.1 Simulated muon range in ice	61
4.2 Ice density profile	62
4.3 GEANT4 based IceTop tank response to muons and electrons	64
5.1 ShowerLLH energy reconstruction performance	70
5.2 Energy reconstruction bias comparison	71
5.3 ShowerLLH primary likelihood performance	73
5.4 Primary likelihood ratio comparison	74
5.5 Geometry used within MuEx	78
5.6 MuEx muonbundle energy loss performance	80
5.7 MuEx muon bundle energy loss comparison	81
6.1 Event distribution for simulation and data	85
6.2 Event distribution for simulation and data (alternate projection)	86
6.3 Gaussian fit to reconstructed energy distribution for one bin of Iron simulation	90
6.4 Parameter space separation for 5 components.	91
6.5 Iron simulation effective area fit.	94
6.6 Smoothed mean values in each observable by simulation bin	97

Figure	Page
6.7 Smoothed variance for each observable in Figure 6.6	98
6.8 Non-degenerate smoothed covariance for each observable in Figure 6.6	99
6.9 Reconstructed energy bias for 5 component simulation	100
6.10 Muon bundle energy loss for 5 component simulation	101
6.11 Average primary likelihood ratio for 5 component simulation	102
6.12 Three-dimensional parameter space for 5 component simulation	103
6.13 Combining statistics from adjacent simulation bins by shifting and rescaling.	105
6.14 Effect of smoothing on a realistic spectrum	106
6.15 Weighted simulated flux	110
6.16 Simulated flux	111
6.17 Simulated mean logarithm of the mass number	112
6.18 Weighted unfolding of a simulated flux	113
6.19 Fit residuals normalized by reconstructed energy-bin for a single unfolding of a simulated flux	113
6.20 Weighted unfolding of a simulated flux with error estimation.	115
6.21 Unfolding of a simulated flux with error estimation.	115
6.22 Unfolded mean logarithm of the mass number of a simulated flux with error estimation	116
6.23 Fit residuals	116
6.24 Weighted unfolded simulated flux with small bins	118
6.25 Weighted unfolded simulated flux with large bins	119
6.26 Regularization optimization.	122
6.27 Weighted regularized unfolding of a simulated flux.	123
6.28 Regularized unfolding of a simulated flux.	123
6.29 Regularized unfolded mean logarithm of the mass number for a simulated flux.	124
7.1 Atmospheric temperature profiles for June and December 2009	126
7.2 Comparison of CORSIKA atmosphere models with satellite measured values	127
7.3 Best-fit offset for reconstructed IceTop energy, by month	129

Figure	Page
7.4 Average monthly surface pressure at the South Pole	130
7.5 Surface pressure and IceTop energy offset correlation	131
7.6 Best-fit offset for reconstructed IceTop likelihood ratio, by month	133
7.7 Best-fit offset for reconstructed In-Ice MuEx Energy, by month	134
7.8 Regularization scan for full year result	136
7.9 Regularized flux result	137
7.10 Weighted regularized flux result	138
7.11 Regularized mean logarithm of the mass number result	139
7.12 Fit residuals	139
7.13 Simulation offset scan 2-dimensional plot	141
7.14 Fit residuals after offset	141
7.15 Fit results following application of best fit offsets	142
7.16 Unfolding an QGSJET-II derived simulated flux	145
7.17 Fit results using QGSJET-II	146
7.18 Unfolding an EPOS-derived simulated flux	147
7.19 Fit results using EPOS	148
7.20 Comparison of all-particle flux for all fit results	150
7.21 Comparison of weighted all-particle flux for all fit results	151
7.22 Comparison of mean logarithm of the mass number	152
8.1 Flux comparison to other analyses	155
8.2 Weighted flux comparison to other analyses	156
8.3 Composition comparison to other analyses	157

Chapter 1

Cosmic Rays

“Unfortunately, I did not have time to study this phenomenon more closely.”

Bruno Rossi on his discovery of extensive air showers, 1933 [1]

1.1 Introduction

Following the discovery of cosmic rays over one-hundred years ago, research into the fundamental nature of these particles led to the discovery of numerous previously unknown particles. With the advent of particle accelerators, physicists no longer had to travel to remote mountain laboratories to study these particles, they could do so in laboratories with experiments they could meticulously control.

Accelerator physics has been such an astounding success that all of the standard theory particles have been found, including most recently the Higgs boson [2, 3]. After decades of building progressively more powerful and expensive accelerators, such as the Large Hadron Collider (LHC), it is not clear where to go next. More powerful accelerators would require new facilities or new technology, and the energies accessible by current accelerators have yielded little signs of new physics.

Now the search for new physics returns to, if it ever really left, the cosmos. The highest energy phenomena in the universe are seen in the extremes of the electromagnetic spectrum, and the subatomic byproducts of cosmic-ray interactions light-years away are being observed by neutrino astronomers, but the direct products of these interactions are still accessible to us only in the form of cosmic rays. Cosmic rays accelerated to energies millions of times greater than could be achieved at the LHC routinely interact in the upper

atmosphere. By measuring the energy, composition and arrival direction¹ of cosmic rays, we hope to learn about their origins and the physical processes behind their acceleration to such high energies.

The recent completion of the IceCube Neutrino Observatory at the South Pole presents a unique opportunity to measure cosmic rays in a new way. With its two sub-detectors, IceTop at the surface of the glacier, and IceCube, deployed deep in the ice, it is able to make simultaneous measurements of the showers of particles initiated by cosmic rays interacting with the upper atmosphere. IceTop measures the shower of particles caused by multiple generations of interactions in the atmosphere, while IceCube detects the energy loss of high energy muons created soon after the initial interaction. The combined information from these two distinct components can be used to measure the composition and energy of the cosmic-ray spectrum, while constraining the particle interaction models used to simulate the cascade of particles they create. As the question of the origin of cosmic rays remains; the snowy, remote laboratories have only become more remote, the simple coincidence circuits replaced by Digital Optical Modules, and the sheets of lead now a mile-thick glacier.

1.2 History of Cosmic Rays

1.2.1 Discovery of Cosmic Rays

Antoine-Henri Becquerel discovered radioactivity in the 1890s. In his further studies of radioactive materials, he found that their presence would discharge an electroscope by causing the air in the electroscope to become conductive. This property was used to measure the level of radioactivity by observing the discharge rate of a nearby electroscope. As more sensitive electroscopes were developed, it was found that they would always slowly discharge, even when heavily shielded and removed from known sources of radiation. The prevailing theory was that the earth had some inescapable trace source of radiation.

To test this theory, Theodor Wulf took a sensitive electroscope to the top of the Eiffel Tower, where the 300 m of radiation-absorbing air below should have blocked out any radiation from the earth. When he found the rate of discharge hardly decreased, he wrote there must be “either another source [of radioactive emissions] in the upper portions of the atmosphere, or their absorption by air is substantially weaker than has hitherto been assumed.” [4]

Victor Hess, a young researcher in Austria, decided the only way to test this ambiguous result was to go even higher. On August 7, 1912, he took three shielded, air-tight electroscopes to over 5000 m in a balloon and found the ionization rate had more than doubled the rate at the ground. [5] This could only mean, as he put it, “that rays of very great penetrating power are entering our atmosphere from above.” Since half

¹At all but the very highest energies, cosmic rays are diverted from their original path by the magnetic fields that permeate all of space, making the energy and composition measurements all the more important.




Figure 1.1: Rossi's triode coincidence circuit. Figure from [6].

of his flights were at night and the others during the day, he could rule out a solar origin for the radiation, and concluded that it must be coming from outer space.

1.2.2 Discovery of Extensive Air Showers

In 1930, Bruno Rossi developed a coincidence circuit (Fig. 1.1) for multiple Geiger-Müller tubes to measure the penetrating nature of cosmic rays. He detected particles that passed through a meter of lead [1] by placing three Geiger tubes in a line with lead blocks between them. These particles would later be identified as muons. Later, when testing the accidental coincidence rate by placing the Geiger tubes separated from each other in a triangle, he found that the rate of coincidence was higher than would be expected by chance, and hypothesized that they must be caused by extensive air showers of particles.

An air shower is produced when a primary cosmic ray interacts with an atom in the atmosphere, producing a cascade of particles. The particles include hadrons, like protons and pions; leptons, like electrons and muons, and photons. Due to the very high momentum of the cosmic ray, these particles travel in nearly the same direction and themselves interact with the atmosphere, again producing more particles and so on through the atmosphere until the initial energy of the cosmic ray has been divided so far that no more particles can be created. Extensive air showers will be covered in more detail in Chapter 2.

Pierre Auger [7], working high in the Alps, set out to measure cosmic-ray induced extensive air showers. Using Geiger tubes separated by up to 300 m (as in Figure 1.2) and lead plates to measure the particle energy, he measured showers estimated to be up to 10^{15} eV. In the conclusion to his 1939 paper, he wrote:

One of the consequences of the extension of the energy spectrum of cosmic rays up to 10^{15} eV is that it is actually impossible to imagine a single process able to give to a particle such an energy. [7]




Figure 1.2: Measurement of the rate of extensive air showers by Pierre Auger using coincidence circuits to measure the shower extent, and lead plate radiators to measure individual particle energy. The figure shows coincident detections above the background rate at distances up to 300 m. Shown in the dashed line is a contemporary theoretical prediction. Figure from [7].

1.3 Energy Spectrum

Cosmic rays have been observed with energies from hundreds of MeV [8] to hundreds of EeV [9]. A recent compilation of measurements is shown in Figure 1.3. The overall spectrum follows a power law

$$\frac{dI}{dE} = I_0 \left(\frac{E}{E_0} \right)^\gamma, \quad (1.1)$$

where I_0 is the flux ² at the energy E_0 , and γ is the power-law index. At the lowest energies, up to 10 GeV, the particle flux slowly increases and is strongly modulated by the solar cycle, and particles are deflected by the earth's magnetic field, resulting in an anisotropic flux at earth that depends on azimuth (the East-West effect) [10] and location [11] (e.g. the South Atlantic Anomaly [12]).

Up to hundreds of TeV, the flux of particles is large enough that it can be measured directly by space- and balloon-based detectors. Modern detectors in this energy range include multiple sensors to simultaneously measure the particle mass and energy, so individual spectra for each element (or isotope) can be measured. By comparing the elemental abundances of the cosmic rays to the abundances of the solar system, we can learn about the propagation of cosmic rays through the galaxy [13]. A recent compilation of elemental spectra is shown in Figure 1.4. For cosmic rays of these energies, the power-law index γ is around -2.7.

At a break in the energy spectrum at 3×10^{15} eV, known as the *knee*, the power-law index steepens to $\gamma \sim 3.1$. The origin of the knee is presently unknown, but current models will be discussed in Section 1.3.1. After the knee, the cosmic-ray flux falls below the level that can be measured by direct detection experiments,³ so ground-based indirect detection methods must be used.

The spectrum hardens again at 3×10^{18} eV to $\gamma \sim -2.6$ at a feature known as the *ankle*. Particles of these energies are believed to be produced outside the galaxy, as their Larmor radius, the radius of the circular path of a charged particle in a magnetic field, is larger than the disk of the galaxy for typical galactic magnetic field strengths of around $5 \mu\text{G}$ [16]. The Larmor radius, L , is given by

$$L = \frac{p}{ZeB} \approx \frac{E/c}{ZeB}, \quad (1.2)$$

where p is the momentum (roughly the energy E divided by the speed of light c), Z is the number of protons in the atom, e is the electron charge, and B is the magnetic field. It is believed that these cosmic rays are of extragalactic origins.

A strong suppression in flux is seen above energies of 5×10^{19} eV. This may be due to the interaction of cosmic rays with the Cosmic Microwave Background (CMB) photons discovered by Penzias and Wilson [17] in 1965. This interaction and the corresponding cutoff was predicted independently by Greisen [18] and

²The flux is defined as the number of particles per area solid-angle time.

³Besides the low flux of these particles, there are physical, logistical and budgetary constraints on a direct particle detector that limit the size of the detector, the magnetic field and the calorimeter, which all limit the energy and mass resolution for the highest energy particles.




Figure 1.3: The cosmic-ray all-particle spectrum as measured by balloon- and ground-based experiments [14]. The spectrum covers 10 orders of magnitude in energy and 30 orders of magnitude in flux, and generally behaves as a power-law with a few features. The lowest energy cosmic rays, those with an energy $< 10^{10}$ eV, are deflected by the earth's and sun's magnetic fields and are modulated by the solar wind. At the *knee*, around 3×10^{15} eV, the spectrum steepens slightly, and at the *ankle*, near 3×10^{18} eV, it flattens again.




Figure 1.4: The spectrum of selected elemental components of the cosmic-ray flux [15]. The elements have been separated for clarity, but all follow the overall power-law trend.

Zatsepin and Kuz'min [19] in 1966. It is now known as the GZK limit. Protons with energies above the GZK limit of about 5×10^{19} eV will interact with the 2.7 K CMB photons in the following interactions:

$$\gamma_{\text{CMB}} + p \rightarrow \Delta^+ \rightarrow p + \pi^0 \quad (1.3)$$

$$\gamma_{\text{CMB}} + p \rightarrow \Delta^+ \rightarrow n + \pi^+. \quad (1.4)$$

However, nuclear photo-disintegration of a heavier mixed composition would produce a similar effect for extragalactic sources [20].

1.3.1 The Knee

The knee of the cosmic-ray spectrum, first reported by Kulikov and Khristiansen [21], at around 3×10^{15} eV, remains a tantalizing mystery in cosmic-ray physics. It represents a change in the spectrum of galactic cosmic rays [22] in an energy range that is inaccessible to direct measurements (and therefore robust composition measurements), though high enough in flux that it can be easily seen in the all-particle spectrum by ground-based experiments. Two of the many proposed origins for the knee [23] are nearby sources and propagation effects.

Using the measured cosmic-ray energy spectrum and measurements of the cosmic-ray composition around the knee, Stanev [24] and others have suggested models where the elements produced in a single young supernova and their subsequent acceleration by the supernova wind and shock are sufficient to explain the energy and composition of cosmic rays through the knee. Others have gone further to suggest specific sources such as the Monogem ring in papers by Thorsett [25] and Wolfendale [26].

As suggested by Peters [27], a knee will naturally be produced by a cosmic-ray accelerator with a finite size, such as a supernova remnant. Such an accelerator can only efficiently accelerate particles up to a maximum energy where the Larmor radius (Equation (1.2)) exceeds the size of the acceleration region. Since the maximum energy is proportional to Z , this maximum energy is spread over a factor of 26 for common elements from hydrogen ($Z = 1$) to iron ($Z = 26$).

The consequence of the knee being the result of a single, nearby source, is that other sources, weaker, further or of a different age, should produce other “knees” in other parts of the energy spectrum. Only high-precision measurements of the spectrum, which so far have not been achieved, would be able to see such fine details. Cosmic-ray measurements also suffer from the fact that experiments are optimized for a certain limited energy range, making features near the limits of their reach indistinguishable from systematic measurement errors.

In other models, such as by Swordy [28], the knee is not the result of a single source, but due to the sum of all accelerators in the galaxy, making the knee a consequence of a rigidity-dependent path length for

escape from the galaxy. The predicted anisotropy and composition through the knee depends on the rate at which particles escape from the galaxy.

All of these models are highly constrained by the all-particle energy spectrum, but the many poorly-constrained parameters and assumptions lead to widely varying predictions about the composition, particularly in the range between the knee and the ankle.

1.4 Composition


At low energies, where direct detection of cosmic rays is possible, the composition has been measured very well, as shown in Figure 1.4.

At energies beyond the ankle, the mean-free path for nuclei other than protons and iron falls below intergalactic distances, so the composition is expected to be primarily a mixture of those two components. This is due to interactions with the Extra-galactic Background Light (EBL) and Cosmic Microwave Background radiation (CMB) [29]. Two large detectors currently measuring the composition in this energy range are Telescope Array, which favors a constant light composition [30], and the Pierre Auger Observatory, which favors a composition that transitions from protons to iron [9]. A more comprehensive overview of composition beyond the ankle can be found in [31].

Between the knee and the ankle, the current measurement results are even less clear. Cosmic-ray models generally predict a transition from lighter to heavier elements, but at these energies only indirect measurements can be made. Interpreting indirect measurements requires detailed simulations, further complicating matters.


A number of cosmic-ray flux models, such as Poly-Gonato by Hörandel [32] (shown in Figure 1.5) and the more recent H4a by Gaisser [33–35] (shown in Figure 1.6), have attempted to predict the cosmic-ray flux and composition. Both of these models start with the directly-measured low-energy flux of individual elements or elemental groups, and extrapolate to higher energies. The main differences at low energy are due to the increased energy limits and statistics of direct measurement experiments over the 10 years between the release of Poly-Gonato and H4a. Each model then imposes a Z dependent spectral break to each element to form a knee. H4a includes a second, harder high-energy component for each elemental group with a break at approximately one order of magnitude higher energy to model the region between the knee and the ankle. All cosmic rays beyond the ankle are modeled by a third proton-only component. Poly-Gonato does not model any components beyond the knee, leading to a large disagreement with observations beyond 100 PeV.

The differences in the models are largely explained by the differing motivation behind them. Poly-Gonato is intended to demonstrate that the knee of the all-particle spectrum can be described by a very simple model, extrapolating low energy measurements of each element with a cut-off at one set rigidity. The choice of a single sharp cut-off rigidity implies that just one source, or an ensemble of identical sources, is



(a) The Poly-Gonato model between 10^5 and 10^8 GeV. (b) 2-The Poly-Gonato model weighted by $E^{2.7}$ for clarity.

Figure 1.5: The Poly-Gonato model by Hörandel [32]. The Poly-Gonato model describes one population of individual elements (here shown as representative groups) with a spectral steepening at a cutoff energy proportional to the nuclear charge. The model does not include higher energy components and only attempts to model the knee region. The colors used for the elements in this plot will be used throughout the rest of this work.




(a) The H4a model between 10^5 and 10^8 GeV. (b) The H4a model weighted by $E^{2.7}$ for clarity.

Figure 1.6: The H4a model by Gaisser [34]. The H4a model defines two populations of Galactic cosmic rays split into 5 component groups; Proton, Helium, Carbon-Nitrogen-Oxygen, Magnesium-Silicon and Iron. The first population is an extrapolation of lower-energy direct measurements. The second population represents a higher energy galactic component, introduced to fit experimental measurements. There is a third proton component to fit the spectrum at extremely high energies.

responsible for all cosmic rays through the knee. In Figure 1.5b, elements up to iron are combined into five groups, but the actual model treats each element separately, which could be observable for a hypothetical direct measurement experiment, but not for indirect detection techniques. Conversely, the H4a model was developed to model the cosmic-ray spectrum at and beyond the knee for simulations of high-energy muons for the IceCube detector. Since shower-to-shower variations are much larger than the variation between elements, only five elemental groups are needed to describe the spectrum. The choice of populations and cut-off rigidities is motivated by fitting the all-particle spectrum as measured from the knee to the highest energies.

Since both models describe the knee with a rigidity-dependent cut-off, each predicts that heavier elements become more dominant through the knee. Figure 1.7 shows the mean natural logarithm of the primary mass predicted by each model as a function of energy. Since H4a was developed with the benefit of better low-energy direct measurements, it should more accurately predict the average mass at low energies. At the high-energy end of the knee, the population of extra-galactic cosmic rays will start to become dominant. Since these are not modeled by Poly-Gonato, it cannot be used to describe the average mass at higher energies.



(a) The mean natural logarithm of the mass for the H4a model. (b) The mean natural logarithm of the mass for the Poly-Gonato model.

Figure 1.7: Predictions of the mean natural logarithm of the cosmic ray primary particle mass number for the H4a and Poly-Gonato models between 10^5 and 10^8 GeV. The Poly-Gonato model predicts a heavier composition across all shown energies, with the difference increasing as the energy approaches 10^8 GeV.

1.5 Acceleration and Sources

It was clear from early on that the power-law spectrum of cosmic rays extended to very high energies. A theory for how particles could reach such energies through stochastic interactions and naturally produce a

power-law spectrum was proposed by Enrico Fermi in 1949 [36], and led the way to the more comprehensive theory we have today.

Fermi's original theory involved *magnetic clouds* wandering the interstellar medium accelerating thermal particles. Imagine a one-dimensional case where we have a particle that is already accelerated to relativistic energies (that is $E_0 \simeq p_0 c$) which is approaching a cloud with velocity v_c in the opposite direction.

In the turbulent magnetic field of a magnetic cloud, the particle's path will become isotropic in the frame of the cloud, and in this frame it will have an energy

$$E_{cloud} = \gamma_c (E_0 + \beta_c p_0), \quad (1.5)$$

where E_{cloud} is the particle energy in the cloud frame, γ_c is the relativistic γ of the cloud ($\gamma = 1/\sqrt{1 - \beta_c^2}$), β_c is the velocity of the cloud ($\beta_c = v_c/c$), and E_0 and p_0 are the initial energy and momentum of the particle. Interactions with a magnetic field are elastic, so if the particle passes through the cloud in the direction it was going, its energy in the original frame will not have changed. If its path is reversed, like a ping-pong ball against a paddle, it will leave the cloud with energy

$$E_1 = \gamma_c (E_{cloud} + \beta_c p_0) \quad (1.6)$$

$$= E_0 \times \gamma_c^2 (1 + \beta_c)^2. \quad (1.7)$$

In the situation that the particle and cloud have velocities in the same direction, the particle would lose the same energy. However, the cloud will encounter more particles head on due to its motion, resulting in overall particle acceleration. In reality, space is not one-dimensional, and integrating over all possible interaction angles, the average energy gain per collision is approximately:

$$\left\langle \frac{\Delta E}{E} \right\rangle \approx \frac{8}{3} \beta_c^2. \quad (1.8)$$

Since the energy gain is proportional to the square of β and was originally proposed by Fermi, this is known as second-order Fermi acceleration. The rate of acceleration depends on the rate of collisions with magnetic clouds. Accounting for all possible geometries for the magnetic field connecting two clouds of distance L , the average time between collisions is $2L/c$, so the average rate of energy increase is:

$$\frac{dE}{dt} = \frac{4}{3} \left(\frac{\beta_c^2 c}{L} \right) E \equiv \alpha E. \quad (1.9)$$

For typical distances between galactic objects of a parsec and cloud velocities of $\beta_c = 10^{-4}$ [37], this acceleration is very slow, with one interaction every few years leading to accretion times on the order of 10^8 years, which is greater than the cosmic-ray escape time for the galaxy. This type of acceleration is considered significant only for turbulent compact objects like supernova remnants.

After accounting for diffusion and loss, as in [37], the energy spectrum is:

$$\frac{dN(E)}{dE} = - \left(1 + \frac{1}{\alpha\tau_{esc}} \right) \frac{N(E)}{E}, \quad (1.10)$$

where α is the rate from Equation (1.9) and τ_{esc} is the characteristic time the particle is in the acceleration region. While this naturally produces a power-law, the index of $1 + \frac{1}{\alpha\tau_{esc}}$ is highly dependent on the geometry and history of the accelerator. Since the cosmic-ray spectrum is so smooth across the entire energy range, it is hard to imagine a cosmic conspiracy of compact accelerators at all energies combining to produce such a featureless spectrum.

First-order Fermi acceleration, which describes interactions with astrophysical shocks as produced by supernovae, while not originally proposed by Fermi, follows from second-order Fermi acceleration, but is directly proportional to β .

Following a supernova, ejecta from the expanding remnant races out into the interstellar medium (ISM). A shock forms as the velocity of the remnant v_r exceeds the speed of sound in the ISM. Since the remnant is ionized, the shock velocity is $v_s = 4/3 v_r$. The density of the ISM in the shock ρ_s and downstream (behind the shock) ρ_d are related by the compression ratio

$$R = \frac{\rho_d}{\rho_s} = 4, \quad (1.11)$$

which, by the conservation of matter through the shock $v_s\rho_s = v_d\rho_d$, where v_d is the downstream velocity, gives the relationship:

$$R = 4 = \frac{v_s}{v_d} \quad (1.12)$$

$$v_s = 4v_d. \quad (1.13)$$

Since the upstream (unshocked) ISM is so far unperturbed and isotropic, in the shock frame it has a velocity toward the shock of v_s , while the downstream ISM has a velocity away from the shock of $v_s - v_d$. Thus from the upstream frame, the downstream ISM is approaching with a velocity $v_s - v_d$, and from the downstream frame, the upstream ISM is approaching with the same velocity. In this symmetric situation, accounting for the probabilities of possible interaction angles, the average energy gain per crossing is

$$\left\langle \frac{\Delta E}{E} \right\rangle \approx \frac{4}{3}\beta_s, \quad (1.14)$$

where the relative velocity of the ISM on either side of the shock is $\beta_s = (v_s - v_d)$. As was realized by Bell [38], among others, after accounting for the escape probability, the differential energy spectrum is:

$$N(E)dE \propto E^{-2}dE, \quad (1.15)$$

which is a power law with an index of -2 , independent of details of the origins or geometry of the shock. In contrast to second-order Fermi acceleration, it is much easier to reconcile the cosmic-ray spectrum with

a theory that predicts a constant power-law index, making first-order Fermi acceleration the favored model of cosmic-ray acceleration. Possible sites of shock acceleration are discussed next.

1.5.1 Sources

1.5.1.1 Supernova Remnants

Supernova remnants (SNRs) have long been favored accelerators of galactic cosmic rays [39,40]⁴. Besides being the only known common source of relativistic shocks in the galaxy, with an estimated rate of one supernova every 30 - 100 years [41], a simple energetics argument first proposed by Ginzburg and Syrovatskii [42] motivates their continued popularity. Their argument (here based on [43]): The energy density ρ_E of cosmic rays in the galaxy is measured to be about 10^{-12} erg/cm³. The volume v_{MW} of the Milky Way disk, with a radius of 15 kpc and height of 0.5 kpc, is

$$\pi(15\text{kpc})^2(0.5\text{kpc}) = 353\text{kpc}^3 \approx 10^{67}\text{cm}^3, \quad (1.16)$$

giving the total energy of cosmic rays in the galaxy of 10^{55} erg. The average age of galactic cosmic rays is about 10^7 years (or about 3×10^{14} s), so the total power required to maintain the energy density of cosmic rays is

$$\frac{10^{55}\text{erg}}{3 \times 10^{14}\text{s}} \approx 3 \times 10^{40}\text{erg/s}. \quad (1.17)$$

Three 10-solar-mass supernovae, with ejecta with a velocity of 5×10^8 cm/s per century would yield:

$$\frac{3 \times (10M_{\odot}) \times (5 \times 10^8\text{cm/s})^2}{100\text{years}} = \frac{3 \times (20 \times 10^{33}\text{g}) \times \left(25 \times 10^{16}\frac{\text{cm}^2}{\text{s}^2}\right)}{3 \times 10^9\text{s}} \approx 5 \times 10^{42}\text{erg/s}. \quad (1.18)$$

Even if less than 1% of the kinetic energy of a supernova went into cosmic-ray acceleration, there would still be enough energy to account for the energy density of galactic cosmic rays.

Energetics arguments and descriptions of first-order Fermi acceleration have largely had to stand on their own, as convincing experimental evidence for at least some cosmic-ray acceleration has, until recently, been lacking. The Fermi Gamma-ray Space Telescope observed neutral-pion (π^0) decay from two $\sim 10,000$ year old supernova remnants [44]. Hadronic interactions, such as proton-proton interactions, often create π^0 -mesons, which almost exclusively decay to two γ -rays. From the conservation of momentum and energy, the γ -rays equally split the energy of the π^0 , so each has an energy of $m_{\pi^0}c^2/2 = 67.5\text{MeV}$. Since the π^0 are moving isotropically with respect to the observer, the result is a symmetric *bump* in the γ -ray spectrum centered at 67.5MeV, which is unique to π^0 decays from hadronic interactions. While this result is able to constrain the GeV proton energy distribution at the SNRs, the mechanism for acceleration, and the efficiency of that acceleration, up to and beyond TeV energies is still unknown.

⁴In fact, the 1934 papers by Zwicky and Baade in which they proposed the production mechanisms for neutron stars and cosmic rays were also the first to use the term ‘‘supernova.’’

Observations of SNRs in other bands, such as the X-ray measurements of Cassiopeia A [45] by the Chandra telescope are able to provide other clues to the development of supernova remnants and the environment for possible cosmic-ray acceleration. Using a recent observation of the SNR Cassiopeia A (see Figure 1.8), spectral information was used to fit the distribution of heavy elements in the ejecta cloud of the remnant and their relative abundances⁵. The spatial distributions of elements in an SNR may help in developing a complete model of cosmic-ray acceleration in SNRs.




Figure 1.8: The distribution of heavy elements in the Cassiopeia A supernova remnant, as measured by the Chandra X-ray telescope [45]. The detailed measurement of X-ray emission from SNR such as Cas A allows the measurement of the relative abundances of the elemental components of the ejecta. According to [45], the main component of the ejecta is oxygen, which is spatially correlated in the ejecta with magnesium and neon as shown in Figure 1.8a. Sulfur and silicon (Figure 1.8b) can be found near the edges of the ejecta, along with regions of pure iron(Figure 1.8c).

1.5.1.2 Active Galactic Nuclei

Cosmic rays beyond the ankle, at energies greater than 3×10^{18} eV, have a gyro-radius far too large to be accelerated by SNRs. If they are accelerated slowly by Fermi acceleration, their acceleration sites must either be very large, or have an exceptionally strong magnetic field. Some sources meeting this criteria are shown in Figure 1.9, which plots magnetic field strength versus size of known astrophysical sources. Sources above the solid and dashed lines would be able to contain a proton or iron nucleus long enough for it to be accelerated to 10^{20} eV. Of the potential sources, Active Galactic Nuclei (AGN), making up one tenth of all galaxies, are one of the most promising class of candidates. AGN galaxies are thought to have a supermassive

⁵An unexpected result of the analysis was the concentration of iron at the very edge of the remnant, despite iron being the final element produced in the core of the star before the supernova.




Figure 1.9: This plot [46], by Michael Hillas, relates the size and magnetic field of potential astrophysical accelerators by the gyro-radius for ultra-high energy protons and iron nuclei. While somewhat outdated, this plot, known as the *Hillas Plot*, is a common feature in pedagogical lectures on cosmic rays.

black-hole at their center and are highly luminous in all wavelengths, from radio to the highest energy γ -rays observed. They usually have relativistic outflows, on the tens of kiloparsec scale, that are thought to be a possible site of efficient Fermi acceleration [47].

1.5.1.3 Gamma-Ray Bursts

Gamma Ray Bursts (GRBs) are by far the most energetic events observed, so it would seem fitting that they may be the sources of the most energetic cosmic rays. GRBs are thought to be part of a massive supernova explosion or the merger of binary compact objects like neutron stars. GRB models (for instance [48]) predict jets with Lorentz factors of 100 to 1000, which could allow for very efficient cosmic-ray acceleration. However, recent measurements suggest GRB have an absence of neutrinos [49], which would be naturally produced by hadronic acceleration on such short timescales.

1.6 Recent Measurements

Recent composition measurements around the knee have been summarized by Hörandel [50], but a few results stand out.

A result from KASCADE-Grande [51] (see Figure 1.10) supports a mixed composition around 10 PeV which becomes progressively heavier and iron dominated at 100 PeV before a slight hardening for the light components at 160 PeV. All components other than iron can be described with a single power-law, while iron requires a broken power-law with a break around 60 PeV.

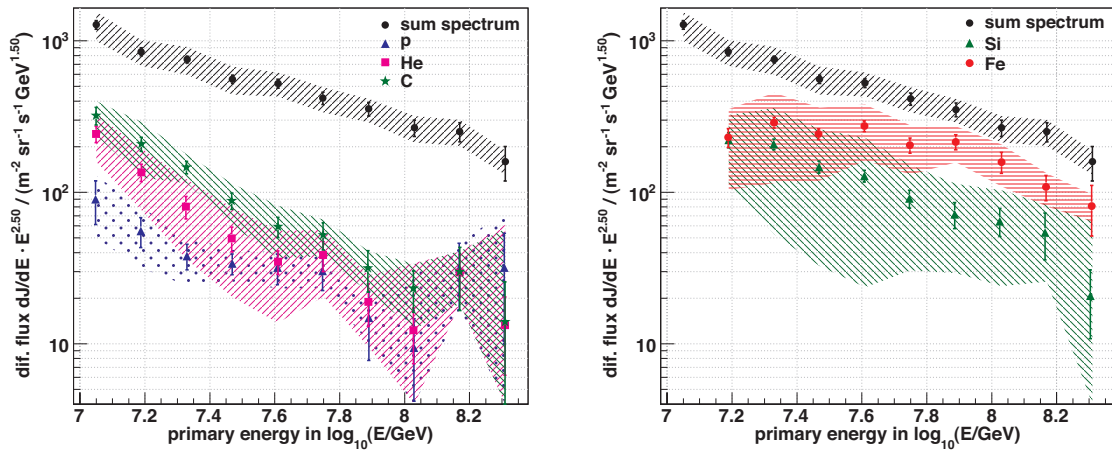


Figure 1.10: The result of unfolding five composition groups from the KASCADE-Grande dataset [51]. On the left, the light composition groups, and on the right, the heavier composition groups. The result shows the progression from a mixed composition to a composition dominated by iron at the highest energies.

An analysis of IceTop and IceCube data from August 2008 [52] also supports a trend toward heavier composition beyond the knee (see Figure 1.11). This publication does not publish spectra for individual elements, but the final data point, at 25 PeV, reports a mean natural logarithm of the mass number A (often written as $\langle \ln(A) \rangle$) of 3.5, which corresponds to a composition of mostly iron or other heavy elements. Figure 1.12 shows the corresponding all-particle spectrum, with a knee at 5 PeV where the power-law index steepens from 2.6 to 3.2.

Another method of determining the spectrum and mass composition of cosmic rays is employed for IceTop data from June 2010 to May 2011 [53]. Instead of fitting the composition, this analysis employs a hypothesis testing method, where a mass composition assumption is used to measure the all-particle spectrum. Figure 1.13 shows the comparison of the three hypothesis compositions. The hypothesis is considered valid when the measured fluxes for all zenith angle bands agree. The authors consider the H4a composition model to agree better than the other two single component compositions tested. However, the fluxes at each zenith angle still does not agree within statistical errors, suggesting further systematic errors. Using the H4a hypothesis, the analysis measures a spectral break around 4 PeV, where the spectral index changes from 2.65 to 3.14, and then another break at 17 PeV, where the index hardens again to 2.90.

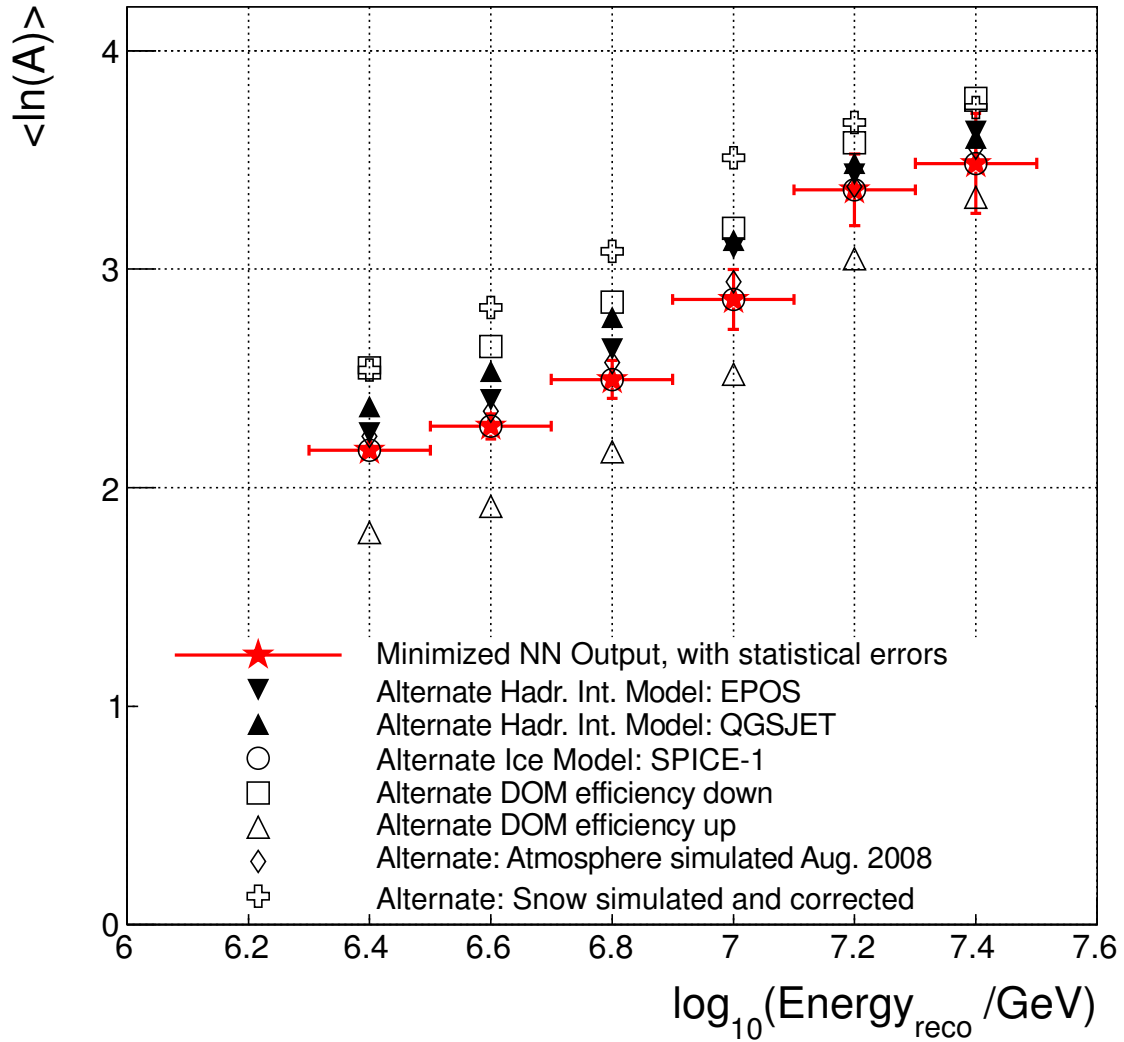


Figure 1.11: The one-month IT/IC-40 composition result [52] shown in the mean natural logarithm of the primary mass. The average cosmic-ray composition was measured as becoming heavier in the region just beyond the knee. The result is also shown for extreme values of potential systematic errors in the measurement.

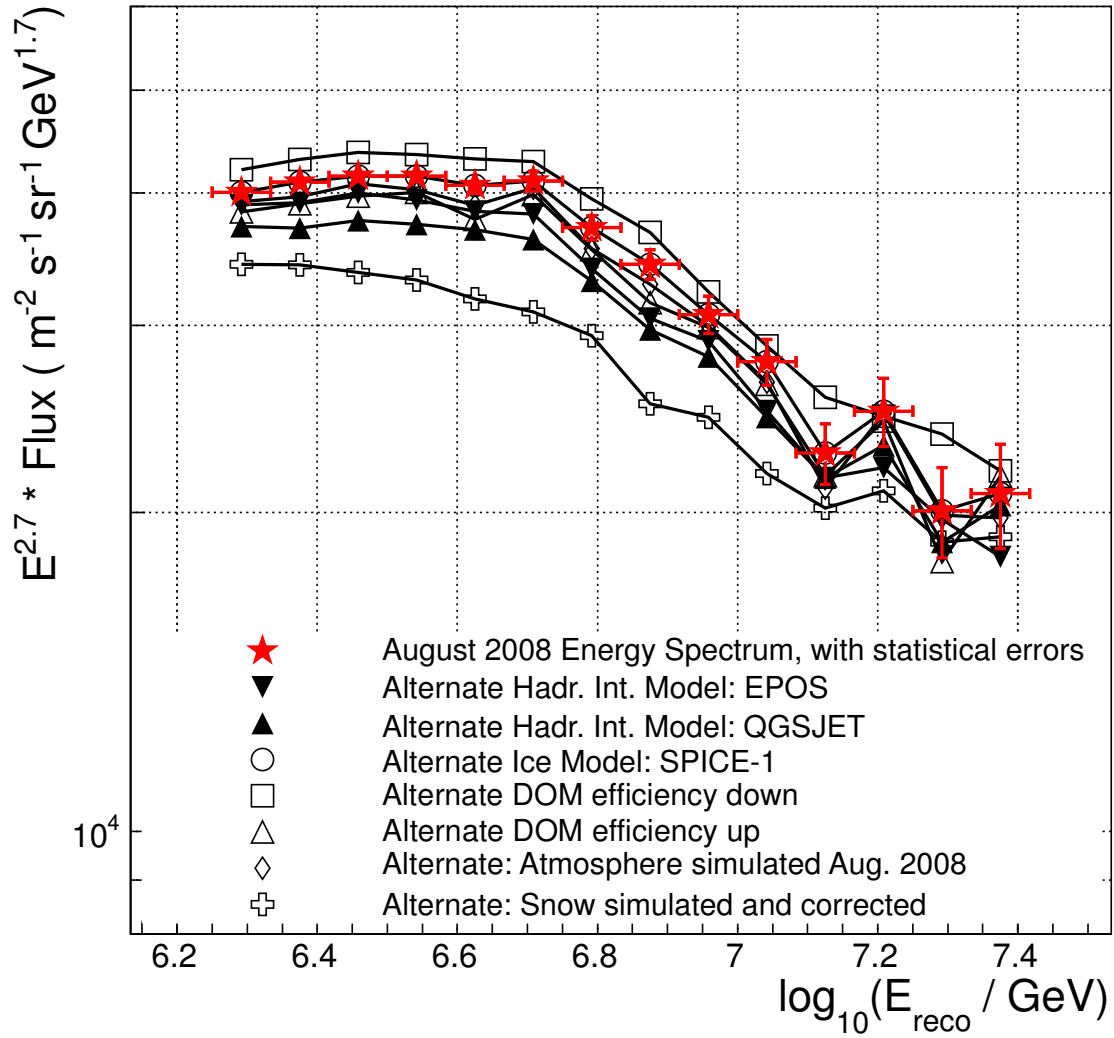





Figure 1.12: The one-month IT/IC-40 all-particle spectrum result [52]. The result can be described as a power-law with a break at 5 PeV. The result is also shown for extreme values of potential systematic errors in the measurement.



(a) Proton.



(b) Iron.



(c) H4a.

Figure 1.13: The all-particle spectrum for one year of data from IceTop-73 in four ranges in reconstructed zenith angle for three composition hypothesis from [53]. Assuming the cosmic-ray spectrum is isotropic, the composition hypothesis that best matches the data will produce the most consistent energy spectrum between zenith ranges. The H4a composition hypothesis is considered by the authors to be the most consistent between zenith ranges compared to the all-proton or all-iron hypotheses. Still, the H4a spectra do not agree within statistical errors and are offset from each other, particularly in the range $6.5 < \log(E/GeV) < 7.5$.

Chapter 2

Extensive Air Showers and Their Detection

“We haven’t got the money, so we’ve got to think!”

Ernest Rutherford (*Attributed*)

Extensive air showers, as introduced earlier in Section 1.2.2, are the result of the interaction of a cosmic ray with the earth’s atmosphere, resulting in a shower of particles. Since the primary cosmic ray cannot be observed, its properties must be determined from the measured properties of the air shower. These measurements are compared to detailed simulation of shower development in the atmosphere in order to determine the properties of the primary cosmic ray. This section gives a brief overview of the major processes of air-shower development and the different components of air showers. A more extensive explanation can be found in Gaisser [54]. This section also presents an overview of air-shower detection methods. Simulation of air showers and theoretical uncertainties are addressed in Chapter 4.

2.1 Initial Interactions

The initial interaction of a cosmic ray with an atomic nucleus in the air starts off the atmospheric particle cascade known as an extensive air shower. For simplicity, the first interaction can be thought of as a proton-proton interaction

$$p + p \rightarrow p + p + N(\pi^+ + \pi^- + \pi_0), \quad (2.1)$$

where a cosmic ray proton interacts with an air nucleon, resulting in a proton continuing on along with a number of pions. Of course the interactants are not always protons, and at these energies nearly any baryons or mesons may be created, and these are dealt with in detailed simulations.

The height of the first interaction depends inversely on the interaction cross-section with air, which depends on the cosmic-ray particle type and energy. The general trend is that the cross section increases with particle energy and atomic mass. The first interaction of a cosmic-ray nucleus with A nucleons and

total energy E_0 can be approximated by the superposition principle as the interaction of A protons each with an energy of E_0/A that interact as in Equation (2.1).

The proton-air cross-section is very difficult to determine for a number of reasons. The cross-section cannot be calculated directly, so it must be inferred from experiment. Since the center of mass energy is far beyond that accessible by accelerators, and individual air-shower interactions cannot be directly observed, the cross-section of the first interaction and all subsequent interactions must be inferred from the statistical development of the air showers themselves. These measurements are very model dependent [55, 56] and depend on assumptions about the composition of the cosmic-ray primaries. For protons around the knee, the mean free path in air is predicted to be around 50 g/cm^2 [55, 56].

The development of the air shower after the initial interaction has three major components: the electromagnetic, hadronic and muonic components. By far the most common secondary particles of the first interaction are pions with very little transverse momentum. Short-lived neutral pions, which decay almost exclusively as

$$\pi^0 \rightarrow \gamma + \gamma, \quad (2.2)$$

will feed the electromagnetic component of the air shower. Charged pions, which have a relatively long lifetime, may interact and continue to feed the hadronic component, or decay, primarily as

$$\pi^+ \rightarrow \mu^+ + \nu_\mu \quad (2.3)$$

$$\pi^- \rightarrow \mu^- + \bar{\nu}_\mu, \quad (2.4)$$

with most of the muons surviving to the surface. Muons with over 300GeV are able to penetrate over a kilometer of ice and reach the IceCube detector.

The following sections will discuss the three air-shower components.

2.2 Hadronic Component

The hadronic component of the air shower is rarely detected directly by ground-based experiments, but as the source of the other components (electromagnetic and muonic), the development of the hadronic shower is a large source of their statistical and systematic uncertainty.

A hadronic shower at the South Pole, traveling through at least 680 g/cm^2 of atmosphere, has undergone some 14 generations by the time it reaches the surface. In each of these generations, an interaction like Equation (2.1) occurs, where the leading proton gets the largest share of the energy and the rest is split equally by the three types of pions. The fraction of energy transferred to the pions is referred to as the inelasticity. Inelasticity suffers from the same measurement difficulties as the cross-section. A typical inelasticity of 0.6 [54] would mean that 20% of the incoming energy would go into neutral pions at each interaction, which

would immediately decay as shown in Equation (2.2) and feed the electromagnetic component of the shower (discussed in Section 2.3).

The other 40% of the energy would go into charged pions, which would either interact again, or decay as in Equation (2.3) to produce muons (discussed in Section 2.4) and neutrinos. The pion cross-section is measured to be about two-thirds of the proton cross-section. The probabilities for a pion to decay or interact are equal when the average distance before decay equals the average distance before interaction

$$\Gamma\tau c = \frac{\lambda}{\rho}, \quad (2.5)$$

where Γ is the Lorentz factor of the pion, τ is its lifetime, λ is the mean free path, and ρ is the density of air.

The hadronic component of the shower, due to the lower Lorentz factors of the particles and the potential for high transverse momentum, is also the main contributor for the variation in the lateral spread of particles in air showers. Since the atmosphere is very thin at high altitudes, even a particle with moderate lateral momentum can travel for several kilometers before interacting, creating spatially separated sub-showers.

2.3 Electromagnetic Component

The electromagnetic component of the air shower, usually initiated by a π^0 as in Equation (2.2), consists of electrons, positrons and photons. The electrons can come from the decay channels

$$\pi^0 \rightarrow \gamma + e^+ + e^- \quad (2.6)$$

$$\pi^0 \rightarrow e^+ + e^- + e^+ + e^-, \quad (2.7)$$

but almost always are from gamma-ray pair production in the presence of a nucleus. The cross-section for pair production at high energies is given by the Bethe-Heitler approximation [57]:

$$\sigma_{pair}(E_\gamma) \approx \alpha Z^2 r_e^2 \ln E_\gamma, \quad (2.8)$$

where Z is the atomic number of the nucleus, r_e is the classical radius of the electron, and α is the fine-structure constant, $\alpha = Ze^2/\hbar c$. For air, the cross section is

$$\sigma_{pair} = 6 \times 10^{-26} \text{cm}^2, \quad (2.9)$$

which corresponds to a mean free path of about 37 g/cm². The energy of the gamma ray is split approximately equally by the electron and positron, which go on to produce more photons through bremsstrahlung.

Bremsstrahlung is the radiation given off when a charged particle is accelerated, in this case an electron or positron accelerated by a charged nucleus in the air, $e^\pm \rightarrow e^\pm + \gamma$. Bremsstrahlung is quantum-mechanically

similar to pair production, which is reflected by the relationship in their cross section

$$\sigma_{brem} = \frac{9}{7}\sigma_{pair}. \quad (2.10)$$

Photons and electrons continue in this way through the atmosphere, with a doubling of the number of particles N at each interaction length λ ,

$$N(d) = 2^{d/\lambda}, \quad (2.11)$$

where d is the atmospheric depth. Each particle has an energy

$$E(d) = \frac{E_0}{N(d)} \quad (2.12)$$

until the critical energy E_c for pair production $E_c = 2m_e c^2$ is reached. At this point, the number of particles has reached a maximum. The atmospheric depth of the shower maximum is known as X_{max} , which is a crucial parameter of air-shower physics, as it is related to shower energy and primary composition. The depth of X_{max} is given by:

$$X_{max} = \lambda \ln \left(\frac{E_0}{E_c} \right) / \ln 2. \quad (2.13)$$

Figure 2.1 shows the lateral distribution of electromagnetic particles reaching ground-level (2834 m a.s.l.) for a simulated 10^{16} eV iron-induced air shower. The particle density is highest at the shower core, which is the location where the projected track of the primary cosmic ray would have hit the ground had it not interacted with the atmosphere. The particle density drops smoothly as a function of lateral distance from the shower core. The lateral distribution of electromagnetic particles is often parameterized by the Nishimura-Kamata-Greisen (NKG) function [58],

$$\rho_{\text{NKG}}(r, s, N_e) = \frac{N_e}{r_M^2} \frac{\Gamma(4.5 - s)}{2\pi\Gamma(s)\Gamma(4.5 - 2s)} \left(\frac{r}{r_M} \right)^{s-2} \left(1 + \frac{r}{r_M} \right)^{s-4.5}, \quad (2.14)$$

where r is the lateral distance from the core, r_M is the Molière radius¹, N_e is the number of electrons, and s is the *shower age*. The shower age is an ad-hoc parameter to describe atmospheric depth of the measurement compared to the atmospheric depth of X_{max} .

Besides producing more photons, electrons and positrons at low energies can also be scattered by an appreciable angle as formalized by Molière [59]. They also produce Cherenkov radiation², which can be detected by specialized ground-based telescopes like VERITAS [60], MAGIC [61] and H.E.S.S. [62].

¹The Molière radius is a property of an interacting medium and is defined as the radius of a cylinder containing 90% of a shower's energy deposition. It is related to the radiation length X_0 and the average nuclear charge Z by $r_m = 0.0265X_0(Z + 1.2)$.

²Cherenkov radiation will be discussed further later in this chapter.




Figure 2.1: The density of electromagnetic particles for a simulated (see Section 4.1.1) 10^7 GeV iron-nucleus induced air shower at an altitude of 2834 m. The particle density is shown as the number of particles per 2.17 m^2 , which is the area of an IceTop water-Cherenkov tank (see Section 3.6.1). The particle density is highest at the shower core and drops off smoothly away from the core. There are a total of 2.15×10^6 electromagnetic particles at the surface. This plot was produced using PyCORSIKA, software developed as part of this analysis, to read a simulation file produced by CORSIKA.

2.4 Muons

Muons do not readily interact with the atmosphere beyond producing Cherenkov radiation from ionization, so they do not contribute to the development of the shower. Since they are created directly by the hadronic component of the shower and are relatively unaffected by the geomagnetic field and intervening matter, they are effective messengers about hadronic interactions at various stages of shower development. At low energies, or where the mean interaction time (the mean free path divided by the speed) is long compared to the muon lifetime, muons can decay before reaching the detector:

$$\begin{aligned}\mu^- &\rightarrow e^- + \bar{\nu}_e + \nu_\mu, \\ \mu^+ &\rightarrow e^+ + \nu_e + \bar{\nu}_\mu.\end{aligned}\tag{2.15}$$

The relative number of high energy muons in an air shower is related to the mass of the cosmic-ray primary. Higher mass primary particles have a higher cross section, and will usually interact higher in the atmosphere. At the same time, from the superposition principle, we know that the products of a high-mass primary will have lower energies on average, and from Equation (2.5) we know that pions have a higher decay probability the lower their energy. These effects lead to an overall larger number of high-energy muons for heavier showers.

Figure 2.2 shows the lateral position of high-energy (at least 500 GeV) muons at ground level (2834 m a.s.l.) for a 10^{16} eV iron-induced air shower. The highest-energy muons and the highest density of lower energy muons are clustered in a small group within several meters of the shower core, with the density of muons falling away at greater distances. Since the highest energy muons were created very early in the shower development (tens of kilometers above the surface) the angle between them is very small, and they will continue to be only meters apart as they reach the IceCube detector, another 1.5km below the surface. These tightly grouped highest-energy muons are referred to as a *muon bundle*.

The number of electron and high-energy (> 500 GeV) muons are plotted for simulated proton and iron showers in Figure 2.3. The figure clearly shows the benefit of energy reconstruction using the combination of both values. Each set of single energy simulations occupies a wide, overlapping range in both number of electrons and muons, but only a relatively small area in the combined plane. At the same time, while the *average* proton and iron shower at a given energy are well separated, the distributions of events overlap considerably. Considering that protons and iron are the *most* different of the common elements that make up the cosmic-ray spectrum, it is clearly impossible to unequivocally determine the mass and energy from these two values. Even assigning a probability for a certain mass group is impossible without beforehand knowing the relative fluxes of each mass group. The well-defined separation in the electromagnetic and high-energy

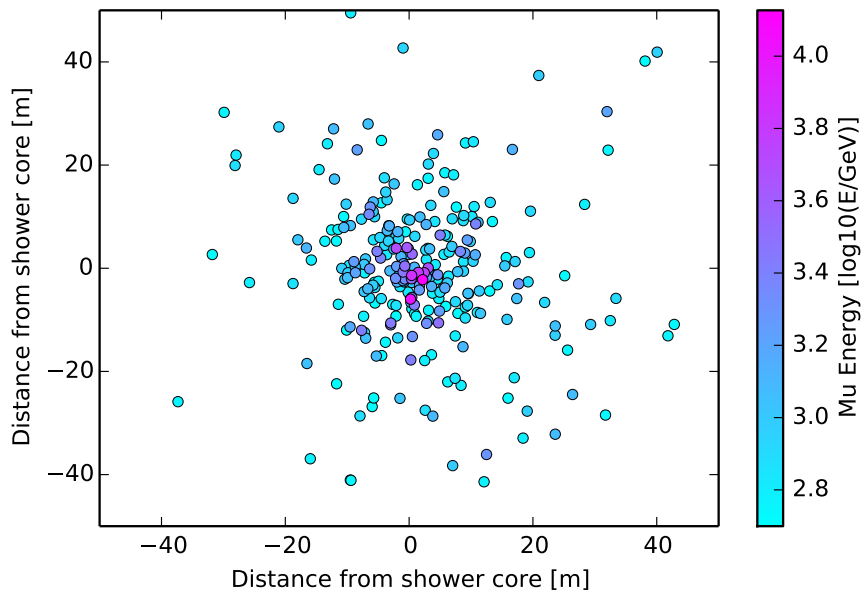


Figure 2.2: High energy (> 500 GeV) muons for a simulated 10^7 GeV iron-nucleus induced air shower at an altitude of 2834 m. The tight bundle of muons, most within about 10 m, and the highest energy muons in an even closer bunch, is called a *muon bundle*. These muons have enough energy to penetrate several kilometers into the ice, and are so close in space and time to be indistinguishable by IceCube. There are 281 muons with more than 500 GeV of energy at the surface.

muon components of event *distributions*, but not individual events, motivates the technique developed for this work, and will be described in detail in Chapter 6.

Lower energy muons, with a typical distribution on the ground shown in Figure 2.4 for a 10^{16} eV iron-induced air shower, are produced by the remaining hadronic component of the air shower at much lower altitudes. While the average distribution falls smoothly with lateral distance from the shower core, the number of muons per IceTop tank area has large fluctuations. Unlike electrons, these muons can produce Cherenkov light as they pass through the entire depth of an IceTop tank, so they can make up the bulk of the observed signal in IceTop far from the shower core.

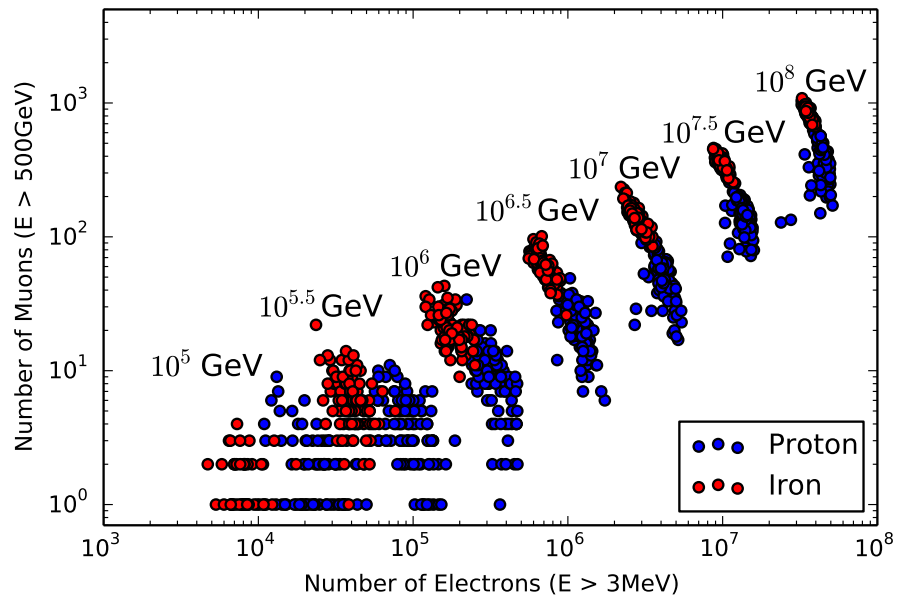


Figure 2.3: The simulated relative number of electrons and high-energy muons for vertical showers of various energies. This plot shows the anti-correlation between the number of muons and the number of electrons caused by fluctuations in the early interactions leading to more or less energy going into the muonic or electromagnetic components of the shower. The larger fluctuations in both muon and electron number can be seen for proton-induced showers, due to the larger variation in the height of the first interaction because of the lower interaction cross-section.




Figure 2.4: The density of muons for a simulated 10^7 GeV iron-nucleus induced air shower at an altitude of 2834 m. As in Figure 2.1, the particle density is shown as the number of particles per 2.17m^2 . The muon density is in general much lower than the electromagnetic particle density, however it does not fall as fast as a function of radial distance from the shower core. Since an individual muon is capable of producing enough Cherenkov light to trigger an IceTop tank, it can make up a significant portion of the detected light far from the shower core. There are 5.2×10^4 muons at the surface.

2.5 Extensive Air-Shower Detection Techniques

There are many methods of detecting extensive air showers. All of them take advantage of the fact that relativistic charged particles radiate energy as they pass through a medium. The methods can be roughly divided by those that detect Cherenkov radiation, and those that directly detect ionization.

2.5.1 Cherenkov Radiation

In the classical pictures, as a charged particle travels through a dielectric medium, it exerts an electrostatic force on the nearby charged particles that make up the medium, briefly electrically polarizing the medium as it passes by. The energy from this interaction will be radiated as an electromagnetic wave traveling out at the speed of light in the medium³, c/n , where n is the refractive index of the medium.

If the charged particle is traveling with a velocity greater than c/n , as is the case for the particles that make up an extensive air shower, the radiation will add coherently as a conical wavefront of light.

Cherenkov radiation is emitted at a specific angle θ_c with respect to the particle's velocity vector. The Cherenkov angle depends on the refractive index in the medium. Figure 2.5 shows a two-dimensional representation of the geometry that defines the conical wavefront and the Cherenkov angle. If the particle is traveling at a speed $\beta = v/c$, then in an arbitrary time t after producing a differential amount of polarization in the medium, the particle will have traveled a distance $t\beta c$, while the radiation from the polarization will have traveled tc/n . Since the time is arbitrary, and c is constant, we see that the angle between the coherent wavefront and the particle direction is given by

$$\cos \theta_c = \frac{1}{\beta n}. \quad (2.16)$$

For example, ice at the South Pole, $n = 1.38$ at 337 nm [63], $\theta_c = 43.6^\circ$, and for air at 10km above the South Pole, $n - 1 = 8.97589 \times 10^{-5}$ at 337 nm [64], $\theta_c = 0.767^\circ$.

The number of Cherenkov photons produced per unit wavelength $d\lambda$ and path length dx is given by the Frank-Tamm equation [65]

$$\frac{d^2 N}{d\lambda dx} = \frac{2\pi\alpha Z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right), \quad (2.17)$$

where α is the fine-structure constant. In ice, this results in about 200 photons/cm between 300 and 600 nm. Since the number of photons is constant with frequency and photon energy is hc/λ , most of the energy is in short wavelength photons. For most materials, the Cherenkov energy spectrum peaks in the ultra-violet, as the wave velocity, c/n , is not constant across all wavelengths. For instance, n is negative for X-rays, excluding them from being produced by the Cherenkov process.

³The speed of light here specifically refers to the phase velocity.




Figure 2.5: A 2-dimensional representation of the geometry that defines the conical wavefront of the Cherenkov radiation produced by a particle through a dielectric medium. In this illustration, the particle is traveling horizontally from the left to right in a medium where $n = 1.5$. The particle starts at the left end of the horizontal line. After a time t , the particle has traveled a distance βtc (where $\beta = v/c$) and the electromagnetic wave it emitted at the start has expanded to a radius of tc/n . Along the particle's path to the right it has continuously emitted electromagnetic radiation (depicted here as two more circles), which constructively interfere at the diagonal line connecting the edges of the circles.

2.6 Recent Experiments

There are a number of ongoing and upcoming cosmic-ray experiments. The following experiments offer examples of different detection techniques and their relative strengths.

2.6.1 Water-Cherenkov Arrays

Water-Cherenkov surface arrays measure the particles of the air-shower cascade reaching the ground with tanks or pools of water. The relativistic charged air-shower particles produce Cherenkov photons while traversing the water which are then detected with photomultiplier tubes. IceTop and IceCube⁴, Milagro [66]⁵, the Pierre Auger Observatory [67], and HAWC [68] are examples of water Cherenkov experiments.

While the name "water-Cherenkov" might seem to imply an identical response to all charged particles, the water serves not only as a detection medium but also an interaction target. Gamma rays will quickly pair produce, and electrons will lose their energy through collisional and radiative losses, which are maximized for electron-electron interactions. The resulting charged particles from the electromagnetic cascade lose energy through Cherenkov radiation. Muons, which are much heavier than electrons, can lose much less of their total energy in muon-electron interactions and typically lose energy only through Cherenkov radiation.

Experiments that focus on detecting primary gamma rays, such as HAWC in Mexico, consisting of water tanks, 7 m in width and 4 m in height, nearly adjacent to each other, take advantage of this difference to veto muon-rich hadronic showers from their signal of muon-poor gamma-ray induced showers. The Pierre Auger Observatory in Argentina, consisting of hundreds of water tanks with 1 km spacing, is designed to detect hadronic showers with a primary energy greater than 10^{18} eV. The detector takes advantage of the penetrating nature and relatively high transverse momentum of GeV muons to measure their lateral density at the ground. The lateral density is used to derive the primary energy by comparison to detailed simulations.

The IceTop detector, consisting of 80 pairs of ice tanks, 1 m in height and 2 m in width, with 125 m spacing, is designed to detect Cherenkov light from both the electromagnetic and muonic component of an air shower. It also uses detailed simulation to derive primary particle quantities from the resulting signal. The IceCube detector, instrumenting the glacial ice from 1500 m to 2500 m below the surface, measures Cherenkov light and more rare stochastic energy loss from TeV muons. By sampling both the high energy muons, created early in the shower development, and the extensive air shower near the shower maximum at the same time, IceTop and IceCube are able to make a more precise measurement of the primary mass and energy than either detector alone.

⁴IceTop uses tanks of ice on the surface of the glacial ice, which is instrumented by the IceCube detector. In the case of IceTop and IceCube, the water is frozen and it may be called ice-Cherenkov.

⁵The Milagro experiment has concluded, but it employed both a pool of water and outrigger tanks of water.

2.6.2 Ground-based Particle Detectors

There are also a number of other methods of charged particle detection used for ground-based detectors. The 6700 m² ARGO-YBJ [69] experiment, at 4300 m above sea level (a.s.l.) in Tibet, uses Resistive Plate Counters (RPCs). RPCs use two parallel plates of high-resistivity plastic separated by a few millimeter gap filled with a quenching gas that is kept at high relative voltage. When a high-energy charged particle ionizes the gas in the gap, the electric discharge is read out by nearby conductive strips. This method can give nanosecond time resolution and millimeter spacial resolution. RPCs give ARGO very good spacial resolution for measuring air showers. While multiple layers of RPCs can be used for particle identification, the ARGO-YBJ experiment uses a single layer, making it unable to reliably distinguish between muons and electrons.

The Grande array, part of the Karlsruhe Shower Core and Array DETector [70] (KASCADE) experiment at 110 m a.s.l. in Karlsruhe, Germany, uses thirty-seven 10 m² unshielded scintillator stations with 137 m spacing. Plastic scintillator is a medium that produces light by luminescence caused by ionizing radiation. It has the advantage over water-Cherenkov detectors of being mechanically stable and producing more light per unit thickness. Since plastic scintillators are relatively thin, they are less sensitive to differences between muon and electron signals. While Grande uses unshielded scintillators, ‘shields’ of steel, lead or snow are sometimes used. A shielded detector will be less sensitive to low energy particles, which are absorbed in the shield, but would be more sensitive to higher energy particles, which cascade in the shield resulting in more particles passing through the scintillator.

Chapter 3

IceCube Detector

“There are technical problems as the cost of making very large area muon detectors is very high and, as the number of muons in the shower is only around 2% of the electron component, complete sampling of the muon component is extremely difficult.”

A. A. Watson, on measuring primary mass
with muons, 2002

The IceCube Neutrino Observatory, located at the South Pole, consists of two main parts: a surface array called IceTop, and the cubic-kilometer deep-ice component, commonly called “IceCube.”

3.1 Detector Array Geometry

IceCube, shown in Figure 3.1, consists of 86 strings of sixty, 10-inch photomultipliers (PMTs) frozen into the polar ice-cap at a depth of 1450 to 2450 m. Eighty of the strings are spaced at 125 m on a hexagonal grid and have their 60 PMTs spaced evenly at 17 m average separation. The remaining six strings comprise a low-energy-optimized sub-detector called DeepCore [71], which is situated at the center of IceCube and has closer spaced strings and PMTs. This analysis uses data taken by a pre-completion configuration of 59 strings and IceTop stations, shown in Figure 3.2, which took data from May 20, 2009 to May 30, 2010.

The IceCube coordinate system [72], depicted in Figure 3.3, is defined with regard to the survey coordinates in use at the South Pole. Survey coordinates are given in Northings, Eastings and elevation in units of the International Foot¹. Northings are in the direction of the prime meridian (or *Grid North* towards Greenwich, England), and Eastings are measured 90° clock-wise from Grid North (in the *Grid East* direction). A reference position near the South Pole Station is defined as 50000' E, 50000'N, to avoid negative numbers. The origin of the IceCube coordinate system is at 46500'E, 52200'N at an elevation of 883.92 m.

¹There are exactly 0.3048 m per International Foot.

The ice surface elevation, as surveyed at the first IceCube string, String 21, is 2832 m, making the ice surface $z = 1948.08$ m in the IceCube coordinate system.




Figure 3.1: The IceCube and IceTop detector as completed. On the surface of the ice, the IceTop array, with each station location depicted by a colored dot, surrounds the IceCube Lab, the central hub for the entire array. Below the surface, there is a 50 m *firn* layer of highly compacted snow and ice crystals above the glacial ice. The IceCube array is instrumented between a depth of 1450 m and 2450 m. DeepCore is a higher density in-fill array in the center of IceCube. The Eiffel Tower is shown for scale.

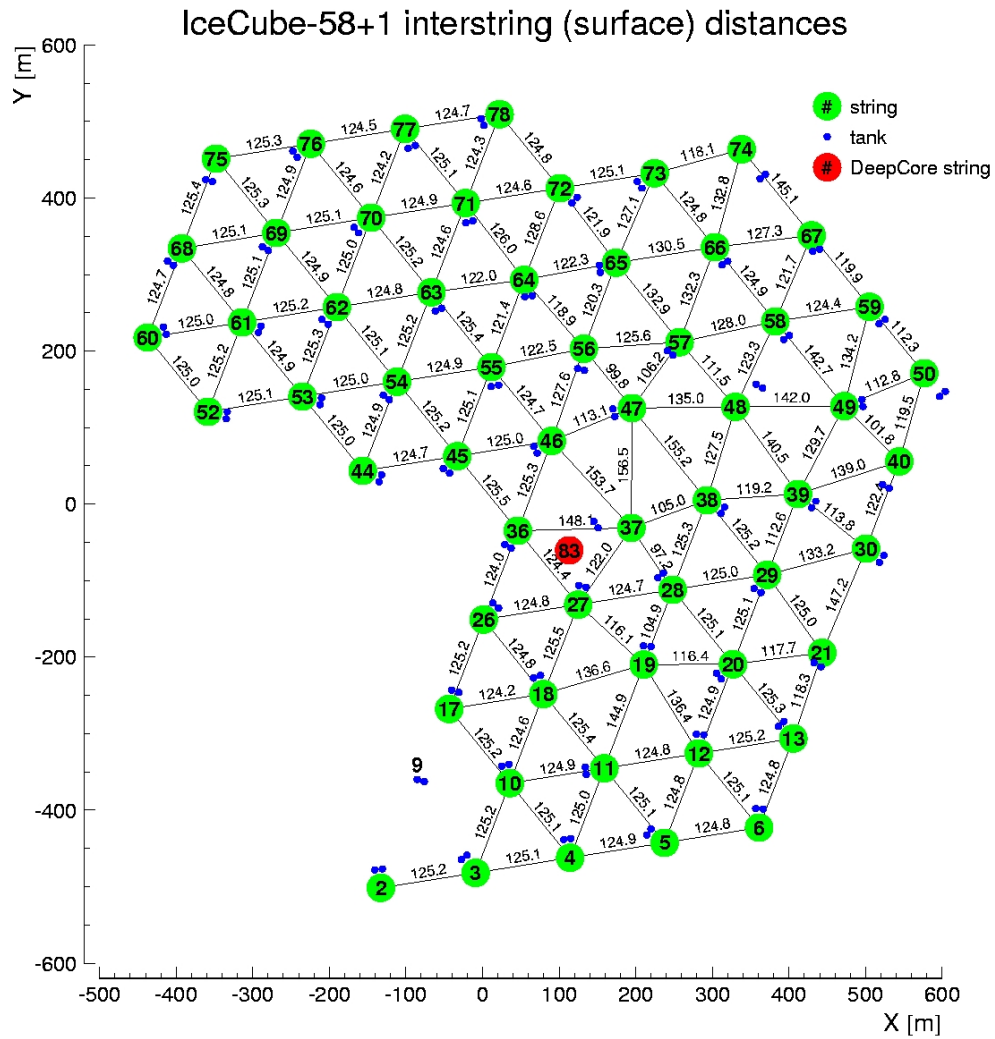


Figure 3.2: The inter-string distances for the IC-59 detector as deployed. This analysis uses data taken using the pre-completion IceCube-59 (IC-59) string detector. The blue dots correspond to the locations of the IceTop tanks that make up each of the 59 IceTop stations. One of the IceCube strings (without a corresponding IceTop station) is a DeepCore string.




Figure 3.3: The IceCube angular coordinate system used for expressing direction. θ is the polar angle with regard to the \hat{z} axis and ϕ describes positive rotation around the \hat{z} axis, measured from the \hat{x} axis. Zenith = $\pi - \theta$, Azimuth = $\pi + \phi$. This image: (*The IceTray Coordinates* [73] by Dusan Turcan. Oil on canvas, 2006)

3.2 Neutrino Detection

The IceCube detector was primarily designed for the detection of cosmogenic neutrinos, a subject covered in detail elsewhere [74], but the design implications are important for this analysis. Neutrinos, due to their neutral charge, can only be detected through their interaction with matter. Neutrino-nucleon interactions (depicted in Figure 3.4) can be either *neutral current* interactions where a proton or neutron exchange a Z boson with a neutrino or antineutrino of any lepton flavor (Figure 3.4 a,b), or *charged current* interactions, resulting in a charged lepton by exchange of a W boson (Figure 3.4 c,d).

High-energy neutral-current interactions result in a nuclear cascade which is nearly indistinguishable from an electromagnetic cascade caused by a charged-current interaction. The notable exception is muon-flavor neutrino charged-current interactions, where the resulting lead muon can carry away the bulk of the energy and travel within a small angle from the direction of the original neutrino. Since the range of muons in ice can be several kilometers, the detectable interaction volume can be increased several times while providing a long lever-arm for direction reconstruction.

IceCube can measure photons given off by the particle cascades (see Figure 3.5) and Cherenkov photons produced when muons traverse the detector (as in Figure 3.6). By measuring the arrival time and number of photons at the photomultipliers, we can reconstruct the arrival direction and (through comparison with detailed simulation) the energy of the primary neutrino. The directional and energy reconstruction qualities of the IceCube detector also make it well-suited to make precision measurements of cosmic-ray muons and their energy loss processes as they travel through the ice. While the direction of neutrino-induced muons can be reconstructed to 1° , neutrino-induced cascades can only be reconstructed to about 15° . Neutrino-induced τ s produce a characteristic signature in the detector of a cascade at the initial interaction, a Cherenkov track, and another cascade when the τ decays [75]. The Cherenkov track length is energy-dependent, only about 50 m per PeV due to the very short lifetime of the τ , making positive identification of charged-current τ neutrino interactions possible for very high energy events. As of this writing, no τ neutrino events have been positively identified by IceCube, but such an event would look similar to the simulation shown in Figure 3.7.




Figure 3.4: Feynman diagrams representing neutral-current (a and b) and charged-current (c and d) neutrino-quark interactions.




Figure 3.5: A cascade event from August 2011. The event was reconstructed as starting within the instrumented volume, with a total deposited energy of $(1040.7^{+131.6}_{-144.4})$ TeV. The small grey dots depict the location of strings of DOMs. The colored dots correspond to triggered DOMs and are colored according to the arrival time of the first photon. They are given a volume proportional to the total number of photons detected. Even though the event direction is not obvious by eye, the precise photon arrival time and distribution constrain the arrival direction to a median angular resolution of 13.2° . Figure from [74].




Figure 3.6: A track event from November 2010. The event was reconstructed as starting within the instrumented volume, with a total deposited energy inside the detector of $(71.4_{-9.0}^{+9.0})$ TeV. The long track of the muon through the detector makes the direction obvious by eye. The event is reconstructed to a median angular resolution of $< 1.2^\circ$. The neutrino energy cannot be well constrained due to the muon leaving the detector, making the total energy deposited only a lower limit on the neutrino energy. Figure from [74].




Figure 3.7: A simulated ν_τ -event in the IceCube-22 detector from [75]. The grey circle in the upper left is the position of the tau creation and the grey circle in the lower right is the location of the tau decay. The small black dots show the positions of IceCube DOMs, and the colored dots show the time of the first photon arrival, with red for the earliest arrivals, progressing through the spectrum to violet for the latest arriving photons. The event shows the characteristic *double bang* signature of two cascades separated by a Cherenkov track.

3.3 Detection of Muons

As muons travel through matter, their energy loss is a combination of continuous and stochastic energy loss processes. Continuous loss is due to ionization, given off as Cherenkov radiation. Stochastic energy loss processes include bremsstrahlung, e^+e^- pair production, and photonuclear interactions. The total energy loss is a sum of these two components and can be approximated as

$$-\frac{dE}{dx} \approx a + bE, \quad (3.1)$$

where $a = 0.26$ GeV/mwe, representing the ionization losses, and $b = 3.57 \times 10^{-4}$ /mwe, representing stochastic losses. MWE stands for “meters water equivalent,” which accounts for the difference in density between ice and water. Figure 3.8 is intended as a schematic to show the regions of the dominant energy loss processes. Since photon production in the ice is proportional to the energy loss, muons with energy less than 1 TeV are nearly indistinguishable. Above the critical energy $E_{\mu c}$ of 1 TeV [76], energy losses are primarily stochastic and scale with muon energy.




Figure 3.8: The stopping power, or average energy loss per cm, for muons in copper as a function of energy. Figure from [15].

3.4 Digital Optical Modules

The fundamental component of the IceCube detector and the IceTop surface array is the Digital Optical Module, or DOM [77]. A DOM is an integrated unit consisting of a PMT, along with timing, high-voltage, calibration, triggering, sampling and communication hardware enclosed in a glass sphere (see the drawing in Figure 3.9).

The PMT is a 10-inch Hamamatsu Photonics photomultiplier [78] (with a high quantum-efficiency photocathode in the case of DeepCore). The digitally controlled high-voltage supply is housed within the DOM, and the AC coupled output is split between the front end for the triggering electronics and a delay board which is then fed into the digitization front end.

The DOM uses two types of digitizers, the Analog Transient Waveform Digitizer (ATWD) and a *Fast* Analog to Digital Converter (FADC)². There are two ATWDs in each DOM, allowing one to be read-out while the other is ready to acquire a new signal. Each ATWD has three input channels used for capturing the input signal, which are fed with three different gain input amplifiers. A fourth ATWD channel is used to digitize the signal from the triggering discriminator. When the ATWD is triggered, each input is fed into a switched analog capacitor storage bank, 128 samples deep, switched at 300 MHz. The FADC, with a 40MHz sampling rate, does not suffer from dead time, but has a limited dynamic range, and the timing resolution is relatively wide compared to typical single photo-electron signals from the PMT.

While each DOM triggers digitization independently, the onboard storage space and the communication rate to the surface limit the amount of data that can be recorded. To limit the data rate, the DOM uses Hard Local Coincidence (HLC) signals from neighboring DOMs to determine if a full digitization should be sent to the surface, or if a much smaller data packet (called Soft Local Coincidence or SLC) should be sent instead. In normal operation, the ATWD is not read-out unless an HLC condition is satisfied, which incurs several microseconds of deadtime for that ATWD. A HLC is satisfied if a neighboring DOM is triggered within 1 microsecond (to account for photon arrival delays due to scattering). An SLC data packet consists of 3 FADC samples for IceCube, or the integrated charge of the FADC sample for IceTop.

Each DOM has a suite of calibration routines built in to independently monitor and calibrate the PMT, digitization circuits, timing and other functions. They also include several calibration LEDs, which can be independently flashed, with their supply voltage pulses digitized, to perform detailed calibration and light propagation measurements in-situ.

Each DOM transfers timestamped HLC and SLC packets to the data acquisition computers located in the IceCube Lab on the surface near the center of the detector. The packets are then time-ordered and

²The FADC here should not be confused with a Flash-ADC.




Figure 3.9: A schematic drawing of a digital optical module (DOM) [77].

checked against trigger conditions to build events. These events are then sent on for reconstruction and filtering to reduce the data rate before being transferred by satellite to the northern hemisphere.

3.5 Ice Properties

Since IceCube uses natural glacial ice deposited over tens of thousands of years as the detection medium, detailed measurements of light scattering and absorption have been carried out. Glacial ice is the most transparent solid known for wavelengths between 200 and 400nm [63], which is the region where the power of Cherenkov radiation peaks, making it an ideal detection medium. At the depths instrumented by IceCube, from about 2500 to 1500m depth, all air bubbles in the ice have been transformed into air hydrate crystals, which have a nearly identical index of refraction as ice, so scattering and absorption are dominated by layers of dust corresponding to stadials³ during the last glacial period.

Dust concentrations in the ice were measured with a specialized device, called the Dust Logger, during detector deployment, and extensive in-situ measurements of the ice properties with the DOMs have been used to fit the scattering coefficient as a function of depth. These two measurements agree very well over the instrumented region, as shown in Figure 3.10. The figure shows the agreement between the scattering profile fit using flashers and the scattering profile measured by the dedicated logging device used during deployment. The disagreement above 1400 m is well outside of the instrumented volume and has very little impact on simulation or reconstruction.

³Stadials are global cold, dry periods, where decreased precipitation and a dry atmosphere allowed a relatively high accumulation of dust.




Figure 3.10: The average dust concentrations as measured by the *dust logger* device used during deployment (red), compared to the effective scattering coefficient b_e fit from in-situ measurements using LED flashers at 405 nm (black). The scattering coefficient determines the mean distance between scatters ($1/b$). The fit scattering agrees quite well with the value from the dust logger data over the instrumented volume. Above and below the detector the scattering coefficients are not well constrained due to the low probability of light scattering back to the detector, but matter very little for the same reason. Figure from [79].

3.6 IceTop Array

IceTop [80] is the surface air shower array component of the IceCube Neutrino Observatory. It has an area of approximately 1 km^2 . It is comprised of 81 stations with an average spacing of 125 m, giving it an energy range of about 100 TeV to 1 EeV, limited by spacing and instrumented area. The array is at a height of about 2834 m a.s.l. with an average atmospheric depth of about 680 g/cm^2 . Since the IceTop array uses the same surface cables as IceCube, the geometry of the array mirrors that of IceCube (see the diagrams in Section 3.1), with the exception of the DeepCore strings, which do not have IceTop counterparts.

3.6.1 IceTop Tanks

Each IceTop station consists of two ice-Cherenkov tanks spaced 10 m apart, each with two DOMs. Figure 3.11 shows two tanks of an IceTop station during deployment. One of the DOMs in each tank is run with the same *gain*, or amplification, as the DOMs deployed in the deep ice and is called the *high gain DOM*. The other DOM is run in a *low gain* mode, to amplify the signal less to avoid saturation for very bright signals and increase the dynamic range of a single tank. The DOMs are connected at the surface junction box so that the high gain DOM in each tank is an HLC neighbor with the two DOMs in the other tank. This allows the individual high gain trigger threshold to be set relatively low, while requiring the HLC condition to maintain a practical data rate. The IceTop DOMs also produce an SLC record, which consists of a timestamp and integrated ATWD charge.

The IceTop tanks themselves (see Figure 3.12) are 1.82 m in diameter, and have an ice depth of 0.90m. The inner surface of the tank is covered with a reflective material, either a Tyvek bag or a zirconium coating. With the Tyvek bag, an average through-going muon produces about 125 photoelectrons in the DOM, half submerged near the middle of the tank. After the ice has started to freeze, the empty space at the top of the tank is filled with perlite to act as a diffuse reflector and rigid insulation. The top cover of the tanks is level with the top of the surrounding snow (which is filled in after installation).

3.6.2 IceTop Tank Response to Air Showers

As an ice-Cherenkov detector, the IceTop tank is sensitive to the diffusely reflected Cherenkov photons caused by particles in the tank. For muons, the interaction length for stochastic energy losses is much longer than the depth of the tank, so the signal is typically proportional to the track length through the tank. For gamma rays and electrons, where the interaction length is only 35 cm, the response is calorimetric up to 50 GeV, where the electromagnetic cascade starts to extend beyond the geometry of the tank. Since the response to muons and electromagnetic particles in the shower cannot be distinguished, the expected response to incident particles is simulated using GEANT4 [81]. Simulation is covered further in Chapter 4.



Figure 3.11: Two IceTop tanks pictured during deployment. The IceTop tanks are placed in a trench, which is filled in after deployment so that the top of each tank is level with the surrounding surface.




Figure 3.12: A cross-sectional view of an IceTop tank. Cherenkov light emitted in the ice is diffusively reflected by the white reflective liner of the tank and detected by the pair of DOMs. The design of the tank and the liner ensure that the amount of light reaching each DOM is essentially equal. In the usual configuration, one DOM is run in *high gain* and is sensitive to as few as one muon, and the other DOM is run in *low gain* mode and only records when the first DOM nears the limit of its dynamic range. Figure from [80].

3.6.3 Calibration

The IceTop detector is calibrated in two ways. The first is the standard DOM self-calibration routine, which determines the PMT gain for single photoelectrons from the dark-noise charge distribution. The second is an IceTop specific trigger to calibrate the tank response to vertical through-going muons, defining a reference charge corresponding to a Vertical Equivalent Muon (VEM).

The VEM calibration is performed using a calibration trigger for IceTop. The firmware in the IceTop DOM has a calibration read-out condition besides the typical HLC condition, where every 8192nd (2^{13}) trigger *without* a local-coincidence is read out and sent to the event builder. The anti-coincidence condition serves to filter out air showers with a surviving electromagnetic component and produces a rich sample of single minimum-ionizing muons. The charge corresponding to the peak of this distribution for a given tank is the VEM charge, shown in Figure 3.13.

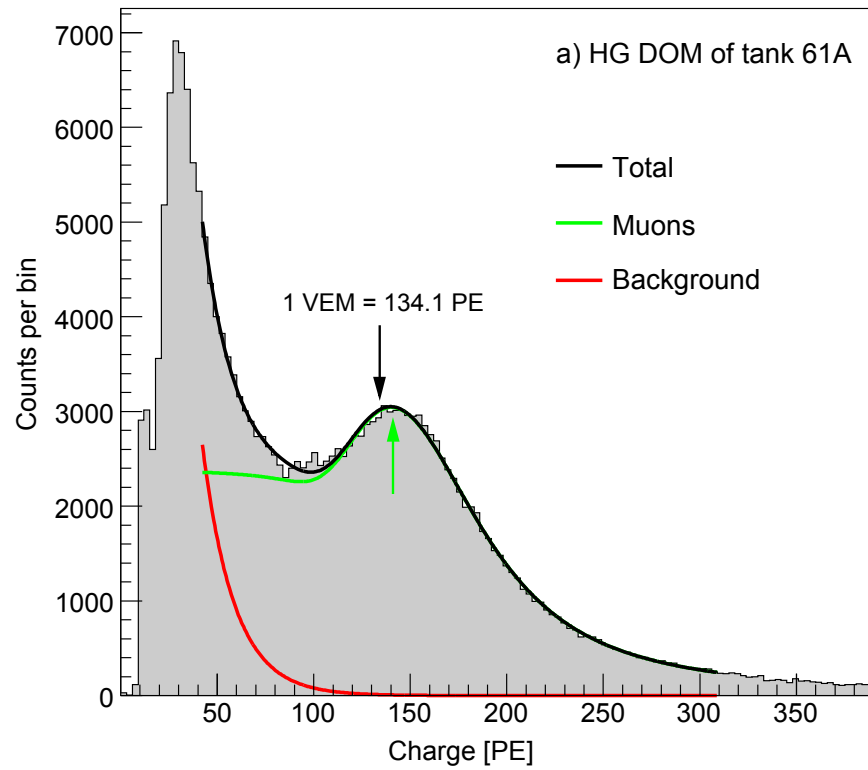


Figure 3.13: The charge spectrum of anti-coincidence VEM calibration triggers for a high gain IceTop DOM. The black fit line is the sum of the background electromagnetic component (fit in red) and the muon calibration signal (fit in green). The definition of the VEM is the peak of the combined distribution (the black arrow) which is slightly below the peak of the muon component (the green arrow). Figure from [80].

Chapter 4

Simulation

I have no data yet. It is a capital mistake to theorize before one has data. Insensibly one begins to twist facts to suit theories instead of theories to suit facts.

Sherlock Holmes, *A Scandal in Bohemia*

The analysis of a modern astroparticle physics experiment does not start with data, it starts with detailed simulation of the detector response to signal and background events. This is especially true for cosmic-ray air shower experiments. Since the particle interactions that determine the development of the extensive air shower are probabilistic, a simple equation cannot describe all possibilities. For instance, a proton cosmic-ray could interact high in the atmosphere resulting in a charged pion carrying almost all of the initial energy. If this pion decays to a muon, the result could be a single muon carrying most of the initial energy, and a very small air-shower. Alternately, instead of a charged pion, a neutral pion could carry most of the initial energy, and decay into two photons, causing an electromagnetic shower with very small hadronic or muon components. While neither of these extreme situations are likely, they must be accounted for in the final analysis. The only way to account for the compound processes that result in these very different showers is by creating extensive libraries of simulated events.

4.1 Simulating Extensive Air Showers

Simulation of extensive air showers and the detector response is handled in several steps. The cosmic-ray primary generation, primary and secondary propagation, and all interactions to the surface are handled by CORSIKA [64], which employs various hadronic and electromagnetic interaction models. The particles that reach the surface are written to a file, which is used as an input for the IceCube simulation software. The IceCube simulation software consists of several modules, which are described in more detail later in this section. The CORSIKA showers are read and sampled around the detector by TopSimulator, which also handles particle propagation, interaction and photon propagation near and within the IceTop detectors

using GEANT4 [81]. Muons are propagated through the IceCube detector using MMC [82] and the resulting photons are propagated to the DOMs using Photonics [83]. The PMT, DOM electronics, triggering and filtering are handled by various software modules that make up the standard IceCube simulation chain called IceSim. All simulation processing is distributed across multiple computer clusters and is managed by the IceProd Framework [84].

4.1.1 CORSIKA

The first step in the simulation process is running CORSIKA [64]. CORSIKA starts by randomly sampling a primary cosmic ray from the input direction and energy ranges. It then propagates the primary cosmic ray through the earth's atmosphere and magnetic field and simulates particle interactions according to the selected hadronic and electromagnetic interaction models¹.

For this analysis, simulations were created to represent five mass groups: proton (H), helium (He), oxygen (O), silicon (Si) and iron (Fe). These elements were selected because they are roughly equally spaced in the logarithm of the mass (number)². Since the mass spacing is smaller than the mass resolution of the analysis, the difference between using ranges of elements or single elements to represent the group is not significant.

As the analysis method requires the core of the air shower to pass through both IceTop and IceCube, we can constrain the simulated distribution of angles and impact parameters to geometries that could reasonably remain in our data sample after event selection. Since IceTop and IceCube have a roughly 1km diameter and the top of IceCube is 1.5km from the surface, the maximum zenith angle of a primary passing through one edge of IceTop and the opposite top edge of IceCube is $\arctan(1/1.5) = 33.69^\circ$. To account for misreconstructed events, we simulated to 40° zenith angle.

Azimuth angles were randomly selected with a uniform probability across all angles. CORSIKA's geometry is defined to have the projection of the primary particle track pass through the center of the observation level x, y plane, so the location of the core of the shower on the detector is selected in a later step (see Section 4.2).

The simulated energy range is $5.0 \leq \log(E/\text{GeV}) \leq 8.0$ where E is the primary energy. For logistical reasons, these were split into 30 energy bins of width 0.1 in $\log(E)$. The simulated differential spectrum is $\frac{dN}{dE} \propto E^{-1}$, so there are an equal number of simulated events in each energy bin. Beyond 100 PeV, the number of secondary particles per shower causes the simulation to take several days per shower, making sufficient statistics prohibitive. At the same time, the number of data events drops off quickly at those energies. Since computer time is a finite resource, efficiency considerations limit the energy range of the simulation.

¹CORSIKA does not handle interactions of neutrinos in the atmosphere, but this is a negligible effect for this analysis.

²The $\ln(A)$ values are 0, 1.39, 2.77, 3.33, 4.02, respectively.

CORSIKA propagates each particle until it either interacts, reaches the observation level (and is recorded in the output file) or has a kinetic energy below a minimum energy threshold specified by the ECUTS option in the steering file. The energy cuts used are shown in Table 4.1.

Particle Type	Minimum Energy (GeV)
hadron	0.05
muon	0.05
electron	0.01
gamma	0.002

Table 4.1: CORSIKA minimum kinetic-energy cuts used in simulation

The single CORSIKA observation level is set to be slightly above the highest IceTop station, at a elevation of 2834 m. This level defines the origin of the coordinate system used in the particle output file. CORSIKA uses an internal coordinate system with the x -axis pointing in the projected direction of magnetic north. In the IceCube coordinate system, magnetic north is at 119° with regard to the x -axis, so the `ARRANG` parameter, the rotation of the detector x -axis with regard to the internal x -axis, in CORSIKA is set to -119° .

4.1.2 High Energy Hadronic Interaction Models

The high-energy hadronic interaction model for simulation is SIBYLL [85]. This is the standard interaction model used in the IceCube collaboration. FLUKA [86, 87] is used as the low-energy (< 80 GeV) hadronic interaction model, which is above the suggested lower limit for SIBYLL (60 GeV [88]) and well below the maximum for FLUKA (20 TeV).

4.1.3 Atmosphere Models

The atmospheric model used for the CORSIKA simulation is an included parameterization of the MSIS-E-90 [89] model for July 1, 1997 [88] (Specified using `ATMOD 12` in the CORSIKA steering file.). The atmosphere is specified in five layers, where the mass overburden $T(h)$, given in g/cm^2 , is modeled as an exponential function of the height h , given in cm above sea level, for the first four layers

$$T(h) = a_i + b_i \cdot e^{-h/c_i} \text{ for } i = 1, \dots, 4, \quad (4.1)$$

and linearly for the fifth layer until reaching the edge of the atmosphere

$$T(h) = a_5 - b_5 \cdot h/c_5. \quad (4.2)$$

The parameters for this atmosphere are shown in Table 4.2.

Layer i	Altitude h (km)	a_i (g/cm ²)	b_i (g/cm ²)	c_i (cm)
1	0...4	-163.331	1183.70	875221.
2	4...10	-65.3713	1108.06	753213.
3	10...4	0.402903	1424.02	545846.
4	40...100	-0.000479198	207.595	793043.
5	> 100	0.00188667	1	$5.9787908 \cdot 10^9$

Table 4.2: CORSIKA July 1, 1997 South Pole atmosphere parameters, as given by [88].

4.2 Resampling

For the detector simulation, each CORSIKA shower core is randomly placed within a 1200 m radius from the origin of the IceCube coordinate system at the observation level just above IceTop. Only particles near IceTop tanks or any muons with sufficient energy to penetrate to IceCube are stored to later propagate and calculate the detector response. Simulating high-energy air showers with CORSIKA is computationally expensive, and it is not known beforehand which impact parameters will trigger the detector or pass quality cuts, so each shower is resampled with a random core position 100 times.

Since the statistical fluctuations due to the limited CORSIKA production are unchanged, this resampling does not increase the statistical fluctuations. For instance, a shower that has a 50% chance of triggering the detector over the sampling region will show up in the final data sample zero or one time if we sampled only once, or some distribution around fifty times if we resampled 100 times. The resulting distribution is the binomial distribution in either case. Accounting for the number of times each CORSIKA shower reappears in the final event sample prevents undue bias, while still making efficient use of limited computing resources. This procedure is described in detail in Section 6.7.1.

4.3 Muon Propagation

Muons with an energy at the surface of 273.0 GeV are propagated through the ice using Muon Monte Carlo (MMC) [82]. MMC is a detailed simulation of muon propagation through water and ice. The 273.0 GeV lower limit was chosen to minimize processing time while still allowing all muons with at least a 0.1% chance of reaching IceCube. Following section 6 of [82], the 99.9th percentile muon range was calculated from simulated vertical muons with energies from 0.1 to 10^{10} GeV as shown in Figure 4.1. The range values were fit with the function

$$x_f = \ln(1 + E_i \cdot b/a)/b \quad (4.3)$$




Figure 4.1: *Left*: Simulated range of 99.9% of muons for discrete energies between 0.1 and 10¹⁰ GeV along with a fit (given in Equation (4.4)). *Right*: the ratio of the fit to the calculated values. Plots from [82].

to determine the values for a and b . These values can then be used to determine an energy cut where only 0.1% of muons will exceed a given range

$$E_{cut}(x) = (e^{bx} - 1) a/b, \quad (4.4)$$

where x is the distance in mwe, and the fit values are $a = 0.212 \frac{\text{GeV}}{\text{mwe}}$ and $b = 0.251 \cdot 10^{-3} \frac{1}{\text{mwe}}$. For IceCube, a 200 m safety margin above the top of the detector was used giving a distance of 1450m – 200m = 1250m. Due to ice bubbles captured in the first 200 m of ice [90], the varying density profile in the first 200 meters (see Figure 4.2) can be treated as 35 fewer meters of solid ice (see Figure 4.2). The nominal ice density is $\rho = 0.9216 \text{ g/cm}^3$, so

$$x_f = (1450\text{m} - 200\text{m} - 35\text{m}) \cdot 0.9216\text{g/cm}^3 = 1119.74\text{mwe}, \quad (4.5)$$

and the energy threshold is

$$\begin{aligned} E_{cut}(1119.74\text{mwe}) &= \left(e^{0.251 \cdot 10^{-3} \cdot 1114.16} - 1 \right) a/b \\ &= 0.32a/b \\ &= 274\text{GeV}. \end{aligned} \quad (4.6)$$

The result of the MMC propagation of muons through the ice is a collection of ionizing muon tracks and cascades corresponding to stochastic energy losses. The light emission from these two types of energy losses is handled by Photonics.




Figure 4.2: Ice density profile for the first 400 m, in black. To simplify simulation, 35 fewer meters of constant density ice is simulated, which corresponds to the same integrated density approximation shown in red.

4.4 Simulating IceCube Detector Response

4.4.1 Photonics

The South Pole ice is exceptionally clear, but the ice still scatters and absorbs photons on their path from emission to detection. Since the actual path of a photon cannot be simply ignored or parameterized [83], individual photons need to be simulated to determine their arrival time distribution and detection probability. Since optical photons are a result of any muon energy loss process in the ice and muons can easily carry a large fraction of the primary cosmic-ray energy, tracing the path of each simulated photon to its eventual detection or absorption (and far more likely, though less usefully, the latter), would be incredibly computationally intensive (even with the latest ray tracing and geometry optimizations). To avoid intensive simulations of each cosmic-ray shower, the *Photonics* [83] software was created.

Photonics takes advantage of the near azimuthal symmetry of the ice (leaving the horizontal ice layers as the only variation) to simulate track segments (ionizing muons) and cascades (stochastic energy losses) of various zenith angles and depths to detectors at an arbitrary relative position. Photonics creates tables of the timing distribution and light yield by simulating each possible geometry to within a set tolerance.

In event simulation, photons at a DOM can be sampled from the photon arrival probability density from nearby track segments and cascades using Photonics table lookups. The resulting photon arrivals are then fed into a simulation of the PMT and DOM electronics, resulting in digitized readouts and triggers, and finally a file in the same format as a standard data file.

4.5 Simulating IceTop Detector Response

As part of the resampling of the CORSIKA showers (discussed in Section 4.2), individual particles landing near IceTop tanks are propagated on-the-fly, and energy losses and photon production are handled by *g4-tankresponse* [91], a GEANT4 [81] based simulation of an IceTop tank (or tables created by parameterizing the same). GEANT4 uses a model of the IceTop tank and the surrounding snow to simulate the passage of particles through, and their interaction with, the nearby environment and the IceTop tank itself. The tank model is a detailed model, including the air above the snow, the snow depth over each tank, the reflective properties of the liner and even the glass pressure-sphere of the DOMs. Figure 4.3 shows the effect of snow above an IceTop tank on the response to muons and electrons. The 1 GeV muons, shown in the left panel, are unaffected by the snow, with the same response for 0 and 50 cm of snow. The 100 MeV electrons deposit some of their energy in the snow above the tank, leaving only 20% for 50 cm of snow (red). Even this detailed treatment only takes about three times the time required by parameterization.

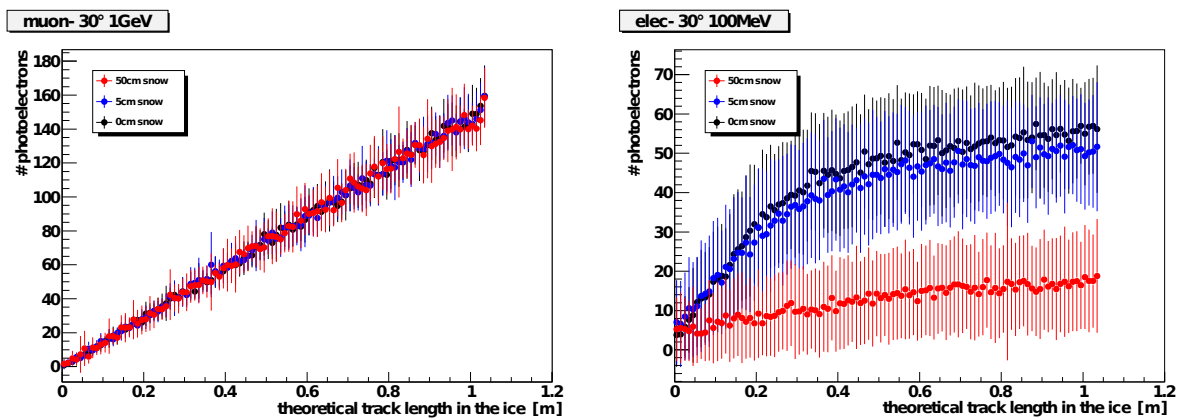


Figure 4.3: Simulated light yield as a function of track lengths within an IceTop tank after passing through various depths of snow. 50 cm of snow is shown in red, 5 cm in blue and no snow in black. *Left*: the light yield for a 1 GeV muon. A 1 GeV muon will only lose a small fraction of its energy passing through even several meters of snow, so the resulting photon production in the IceTop tank is unchanged. *Right*: the light yield of a 100 MeV electron. For 5 cm of snow (blue), only a small amount of the electron-induced cascade occurs outside the tank, and most of the deposited energy is still detected. For 50 cm of snow, the bulk of the cascade will occur outside of the tank, resulting in a much lower response. Plots from [91].

4.6 Simulation Mass-Production

Mass-production of cosmic-ray simulations is a technically difficult problem due to the vast difference in computation and storage requirements across three orders of magnitude in primary energy. A single CORSIKA shower can take from 10 seconds to multiple days to simulate, and the files can range from kilobytes to 20 gigabytes in size. A typical dataset will have tens of thousands of individual files created on multiple clusters on hundreds of nodes. Fortunately, IceProd [84] provides a way to create, document and track datasets. Over the course of this analysis, IceTop simulation was incorporated into the IceProd framework, resulting in improvements to both the simulation and the framework.

Chapter 5

Reconstruction of Cosmic-Ray Air Showers

“If your experiment needs statistics, you ought to have done a better experiment”

Ernest Rutherford

This chapter will discuss the observables and reconstructed properties of cosmic-ray air showers that will be used to measure the cosmic-ray energy spectrum and mass composition. These parameters are reconstructed primary energy and a composition-sensitive likelihood ratio from IceTop, and the reconstructed energy-loss by high-energy muon from IceCube.

5.1 IceTop

Event reconstruction in IceTop, as in all of IceCube, follows a modular, multi-step approach. Specialized software, individually tested and code-reviewed, handles specific tasks in the analysis pipeline. The first step, after the data is decoded and unpacked from the compressed form sent over the satellite is calibration, performed by a module called DOMcalibrator. This step, shared with IceCube analysis, simply converts the raw digitizer values into voltages. Since IceTop is calibrated in VEM, the details of this calibration step are relatively unimportant.

5.1.1 Pulse Extraction

The next step in the analysis chain is TopWaveProcessor. TopWaveProcessor has several functions. First, it selects the highest-gain unsaturated ATWD waveform, then determines the baseline by taking the average of the last 40 samples of the waveform. Since the typical decay time for light deposited in the tank is ~ 30 ns, and the ATWD readout is 426 ns long, all photons will have dissipated by that time. The baseline is subtracted and the remaining pulse is integrated to yield a total charge. This charge is divided by the peak of the single photoelectron charge spectrum to yield a total charge in units of photoelectrons.

The next step of pulse extraction is to determine a time for the pulse. The pulse time is determined by taking the time corresponding to the intersection of the baseline with a line extrapolated from the

steepest slope before the highest peak of the waveform. This choice of algorithm is simply to exclude small uncorrelated noise pulses preceding the pulse corresponding to the triggering air shower.

At this point, each triggering IceTop tank has a pulse, consisting of a time and charge in photoelectrons, for each high-gain DOM and possibly the low-gain DOM.

5.1.2 Event Cleaning

The final step in the calibration of an IceTop event is TopEventBuilder. Because the properties of individual tanks differ, the muon calibration must be applied to the photoelectron charge for each tank. By dividing by the VEM calibration (in units of VEM/photoelectrons, see Section 3.6.3), the tank-to-tank variation can be removed.

Since each tank has both a high-gain and a low-gain DOM sampling photons from the same tank, they carry redundant information. If both DOMs produce a valid pulse, and if the high-gain DOM has a signal of less than 30 VEM, that pulse is the only one reported. However, if the signal is greater than 30 VEM, a pulse with the time from the high-gain DOM and the charge from the low-gain DOM is reported. This hybrid pulse preserves the more precise timing of the high-gain DOM while taking advantage of the dynamic range provided by the low-gain DOM.

There is a chance that more than one air shower will contribute to a single event as determined by the global event builder. As a final step, TopEventBuilder uses the time and spacial location of each pulse to determine which pulses can be causally connected to each other. It then returns causally connected groups of pulses ordered by total charge. For the purpose of this analysis, the largest such group will be the only one considered and will hereafter be referred to as the IceTop event.

5.1.3 Direction Estimation

The direction of the IceTop event is fit by TopRecoPlane by approximating the shower front as a plane, which is roughly valid near the core of the shower. For the best fit direction $\hat{n} = (n_x, n_y, -\sqrt{1 - n_x^2 - n_y^2})$, the χ^2 of the pulse times is minimized. For the average time T_0 and each pulse with a time t_i at location (x_i, y_i, z_{icetop}) ¹, the minimized function is:

$$\chi^2 = \sum_i w_i (t_i^{measured} - t_i^{fit})^2 = \sum_i \frac{(t_i^{measured} - T_0 + \frac{n_x x_i + n_y y_i}{c})^2}{\sigma^2}. \quad (5.1)$$

In a second iteration, the relative height Δh is used to calculate a correction time factor of $\Delta h/c \times \cos \theta$ for greater precision, where θ is the angle from zenith. A constant weight $w_i = 1/\sigma^2$ is used, where $\sigma = 5\text{ns}$.

¹In the first approximation, all tanks are considered to be on a plane at height z_{icetop} .

After the direction fit, the remaining shower parameters are the core position and an estimate of the primary composition and energy. Reconstructing these shower parameters requires a more sophisticated type of analysis discussed in the next section.

5.2 Likelihood-based Reconstruction

The method of maximum likelihood is the estimation of a set of unknown shower parameters \vec{a}_0 from a set of independent observed values $(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n)$. An observed value \vec{x}_i is drawn from a probability density function (PDF)² $f(\dots)$ which is one of a set of probability density functions $\{f(\dots|\vec{a}), \vec{a} \in \vec{A}\}$ from all possible shower parameters.

The joint probability density function for all measurements is

$$\begin{aligned} f(\vec{x}_1, \vec{x}_2, \dots, \vec{x}_n|\vec{a}) &= f(\vec{x}_1|\vec{a}) \times f(\vec{x}_2|\vec{a}) \times \dots \times f(\vec{x}_n|\vec{a}) \\ &= \prod_{i=1}^n f(\vec{x}_i|\vec{a}) . \end{aligned} \quad (5.2)$$

In practice, a likelihood $\mathcal{L}(\dots)$ is calculated, with a hypothesis \hat{a} as an argument of the likelihood,

$$\mathcal{L}(\hat{a}, \vec{x}_1, \vec{x}_2, \dots, \vec{x}_n) = \prod_{i=1}^n f(\vec{x}_i|\vec{a}) . \quad (5.3)$$

For numerical stability, the logarithm of the likelihood function is typically used,

$$\ln \mathcal{L}(\hat{a}, \vec{x}_1, \vec{x}_2, \dots, \vec{x}_n) = \sum_{i=1}^n \ln f(\vec{x}_i|\vec{a}) . \quad (5.4)$$

In general, there is no analytic solution, so $\ln \mathcal{L}$ is numerically maximized with respect to \hat{a} . The relative log-likelihoods in the neighborhood of the global maximum can be used to estimate the error on \hat{a} by relating f to the pdf for a normal distribution. The pdf for a normal distribution is

$$f_{normal}(x, \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} , \quad (5.5)$$

where μ is the mean and σ is the standard deviation. For an arbitrary case set $\mu = 0$ and $\sigma = 1$ then

$$f_{normal}(x, \mu = 0, \sigma = 1) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} . \quad (5.6)$$

For a single x , $\mathcal{L}(x) = f_{normal}(x)$, and we know the maximum likelihood corresponds to the location $x = \mu$, the relative likelihood at $x = \sigma = 1$,

$$\frac{\mathcal{L}(x = 1)}{\mathcal{L}(x = 0)} = \frac{\frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}}}{\frac{1}{\sqrt{2\pi}}} = e^{-\frac{1}{2}} , \quad (5.7)$$

²In this case, the probability density function is a density drawn from a multi-dimensional histogram.

or for the $\ln \mathcal{L}$:

$$\ln \left(\frac{\mathcal{L}(x=1)}{\mathcal{L}(x=0)} \right) = \ln \mathcal{L}(x=1) - \ln \mathcal{L}(x=0) = -\ln 2. \quad (5.8)$$

For an arbitrary $\ln \mathcal{L}$ with a maximum at \hat{x}_0 , any \vec{x} that satisfies the condition

$$\ln \mathcal{L}(\vec{x}) - \ln \mathcal{L}(\hat{x}_0) \geq -\ln 2 \quad (5.9)$$

falls within the 1σ confidence interval.

5.2.1 ShowerLLH

For composition-sensitive IceTop reconstruction, a new likelihood-based reconstruction module, called ShowerLLH, was developed. ShowerLLH uses tables of simulated IceTop tank responses as a function of snow overburden and impact parameter to air showers to define the PDF for proton and iron nuclei showers across a range of directions and energies.

After the shower direction is reconstructed by TopRecoPlane, the direction component of the impact parameter is already set and not varied during the likelihood maximization procedure. The snow overburden is a measured property of each tank, so this is also kept constant. The remaining parameters, core position (the other component of the impact parameter) and energy, are maximized for each tested composition. The likelihood is maximized using a brute-force scan, with three iterative grids for the shower core using energy bins equally spaced in the logarithm of primary energy. After locating the best shower core for each primary composition, an array of likelihoods corresponding to each energy bin are returned along with that position. Through code optimization, the brute-force scan method runs sufficiently fast compared to other steps of the reconstruction, while avoiding having to smooth the likelihood space to avoid local maxima or discontinuous derivatives that confuse common numerical optimizers.

For this analysis, ShowerLLH was run for both proton and iron hypothesis. The reconstructed primary energy value used in this analysis is the average of the most likely proton- and iron-hypothesis energies. Figures 5.1 and 5.2 show the reconstructed primary energy bias, which is the reconstructed energy divided by the true energy, as a function of true energy. The iron simulations (red contours) show consistent performance across the entire energy range, with the most likely value centered on the primary energy, and the tails of the event distribution relatively symmetrically distributed around the mean, with the resolution slowly improving at high energies. At the low energies, below 10^{15} eV, there is a slight bias toward overestimating the primary energy. This is due to selection effects from triggering and quality cuts, as showers that have an upward fluctuation in signal at IceTop are more likely to be included in the final event sample. The performance for proton simulations (blue) is similar, with the effects of shower-to-shower fluctuations more exaggerated due to the smaller interaction cross-section for protons in the atmosphere. Compared to iron, lower energy proton showers are able to trigger the detector and pass the quality cuts due to the higher probability of




Figure 5.1: ShowerLLH energy reconstruction performance. The relative energy bias, the logarithm of the true energy subtracted from the logarithm of the reconstructed energy versus the logarithm of the true energy, is shown for proton (blue) and iron (red contours) simulation.




Figure 5.2: The median logarithm of the true energy subtracted from the logarithm of the reconstructed energy versus the logarithm of the true energy for proton (blue) and iron (red) simulation. The central 68% is shown as vertical error bars.

deeply-penetrating proton primaries causing showers that develop lower in the atmosphere and have an X_{max} closer to the IceTop detector. Similarly, at high energies, deeply penetrating proton primaries induce air-showers that have not yet reached X_{max} by the time they reach IceTop. These showers have a smaller lateral distribution, and are consequently reconstructed with a lower energy.

ShowerLLH is also used to calculate the primary likelihood ratio, which is defined as the logarithm of the ratio of the maximum likelihood value for the proton hypothesis divided by the maximum likelihood value for the iron hypothesis at the best fit core location. Figures 5.3 and 5.4 show the event distribution of this value for proton (blue) and iron (red contours) simulation versus the logarithm of the true energy. At low energies, these values are not well separated, which is largely due to the limited information available at the trigger threshold. At higher energies, where more IceTop stations contribute to the likelihood, the primary likelihood ratio is better separated.




Figure 5.3: ShowerLLH primary likelihood performance. The primary likelihood ratio, described in the text, is shown for proton (blue) and iron (red contours) simulation.




Figure 5.4: The median primary likelihood ratio versus the logarithm of the true energy for proton (blue) and iron (red) simulation. The central 68% is shown as vertical error bars.

5.3 IceCube

5.3.1 Pulse Extraction

In IceCube, several dust layers may be crossed between a photon's emission and detection, so the properties of the ice need to be dealt with in the reconstruction. Since a vast majority of IceCube DOM triggers are single photoelectrons, the base calibration unit for IceCube is the photoelectron. FeatureExtractor is the module that takes a calibrated waveform from DOMcalibrator and extracts a pulse with a charge in photoelectrons with a time. The method is similar to that used for IceTop, with the exception that a single waveform can yield multiple pulses.

5.3.2 Event Cleaning

The IceCube event builder groups causally connected triggers from IceTop and IceCube into single events, accounting for a separation of $1450\text{m}/c = 4850\text{ns}$ for a vertical shower, to $1450\text{m}/(\cos\theta \cdot c) = 5600\text{ns}$ for a shower inclined at $\theta = 30^\circ$. Additionally, due to the large volume of IceCube, it also groups together multiple causally connected IceCube events, as it takes a muon $\sim \sqrt{3 \cdot 1000^2}\text{m}/c = 5793\text{ns}$ to traverse IceCube from corner to corner. As the event builder continues to merge triggers into a single event as long as the coincidence criteria are met, there can be many uncorrelated triggers in a single event. Since the highest energy IceTop trigger is of the most interest, any IceCube trigger correlated with that trigger is the only trigger that should be kept. The module `coincTWC` is responsible for selecting the IceCube pulses causally correlated with the largest IceTop trigger, according to the time constraints as calculated above with an additional safety margin. These pulses are used by further direction and energy reconstructions.

5.3.3 LineFit

LineFit is designed simply as a first guess position and direction for tracks in the ice, in this case, the high energy muon bundle. It treats the Cherenkov light in the detector as a planar (as opposed to conical) wavefront of photons perpendicular to the track at a positive velocity from a vertex point. It employs an analytical solution, to avoid minimizing over a complicated parameter space, and runs very quickly.

Given a pulse with hit time t_i from a DOM with location \vec{r}_i , the particle track can be described in terms of velocity \vec{v} and vertex \vec{r}_0 as

$$\vec{r}_i = \vec{r}_0 + \vec{v}t_i . \quad (5.10)$$

For all pulses in an event, the χ^2 to be minimized is

$$\chi^2 = \sum_{i=1}^n (\vec{r}_i - \vec{r}_0 - \vec{v}t_i)^2 , \quad (5.11)$$

where n is the number of DOMs (or pulses) contributing to the event. By differentiating with respect to the free parameters, r_0 and \vec{v} , the solutions are

$$\vec{r}_0 = \langle \vec{r}_i \rangle - \vec{v} \langle t_i \rangle \quad (5.12)$$

$$\vec{v} = \frac{\langle \vec{r}_i t_i \rangle - \langle \vec{r}_i \rangle \langle t_i \rangle}{\langle t_i^2 \rangle - \langle t_i \rangle^2} . \quad (5.13)$$

5.3.4 Likelihood-based Reconstructions

IceCube also uses a number of likelihood-based reconstructions. The simplest of these are both based on a calculation of a *time residual* from a track. The unscattered Cherenkov wavefront will arrive at a DOM located at \vec{r}_i at a time

$$t_{geo} = t_0 + \frac{\vec{p} \cdot (\vec{r}_i - \vec{r}_0) + d \tan \theta_c}{c} , \quad (5.14)$$

where \vec{p} is the direction vector of the particle, and d is the distance of closest approach between the DOM and the track. The time residual is the difference between the actual hit time and the time of a *direct hit* based on the geometry

$$t_{res} = t_{hit} - t_{geo} . \quad (5.15)$$

The Single-PhotoElectron, or SPE [92], reconstruction considers only the first photon to arrive at each DOM, and ignores all untriggered DOMs. The PDF used is that for *any* photon (see Section 5.3.4.1), but in the case of multiple photons arriving at a DOM, the first photon will typically be scattered less than average. Due to the much larger volume of the detector farther from the track than near, most DOMs have only a single photon, for which the PDF is valid. SPE is used primarily because it is relatively robust to uncertainties in calibration and local ice properties and provides a good seed track for further reconstructions.

The Multiple-PhotoElectron, or MPE [93], reconstruction takes advantage of the fact that the first photon is scattered less than average. For N photons, the MPE PDF is

$$p_N^1(t_{res}) = N p_1(t_{res}) \left(\int_{t_{res}}^{\infty} p_1(t) dt \right)^{N-1} , \quad (5.16)$$

where p_1 is the SPE PDF. For $N = 1$, the MPE PDF is the same as the SPE PDF. For large N , the MPE PDF is sharply peaked towards small t_{res} . MPE has such a narrow PDF for large charges, which are highly sensitive to small mis-estimates of the track location, that it requires a very good seed track, such as that from SPE. Muons from cosmic rays are parallel but not collinear, and the single track assumption breaks down for shower-induced muon bundles. The MPE PDF also does not account for the fact that much of the light produced by a muon is not Cherenkov light, but actually stochastic-induced cascades that emit light more spherically, which breaks down the all-Cherenkov assumption of the MPE PDF. For these reasons, MPE is not used for this analysis, and SPE is used as the seed for further reconstruction.

5.3.4.1 The Pandel Function

The PDF used for the SPE reconstruction is based on the Pandel Function [94], which is an analytic estimate for the arrival times of Cherenkov photons from a monochromatic, isotropic point source. It is expressed as a gamma function

$$\begin{aligned}
 p(t_{res}, d) &= \frac{1}{\Gamma(a)b^a} x^{a-1} e^{-t_{res}/b} \\
 b &= 1/(\tau^{-1} + c/(nx_0)) \\
 a &= d/\lambda,
 \end{aligned}
 \tag{5.17}$$

where d is the distance to the track, n is the index of refraction, and τ , λ and x_0 are parameters determined by simulation to describe the average scattering of photons in the ice [95]. When used in reconstruction, the function is convoluted with a Gaussian with a width of $\sigma = 4\text{ns}$ to account for variation in PMT transit time and detector time calibration. A small constant term corresponding to random noise from the PMT and pressure sphere is also added. These two additional factors correct for unphysical narrow arrival time probability very close to tracks and the zero probability for $t_{res} < 0$.

5.3.5 MuEx

Muon energy loss in IceCube is reconstructed using MuEx [96], which like Photorec [97] uses the ice model to reconstruct the photon density at the track from the calibrated pulses from the DOMs. MuEx also refines the direction reconstruction by calculating an average of SPE reconstructions on bootstrapped (sampling with replacement) input pulses. Once the track position is established, the total number of photons detected at each DOM and the DOM position in the ice relative to the track can be used to reconstruct the photon density along the track, and by extension (see Section 3.3), the muon energy.

The photon PDF for cascades with scattering length λ_e and absorption length λ_a at a distance d between the emission point r_i and the detection point r_f , as depicted in Figure 5.5, is

$$\mu(d) = n_0 A \cdot \frac{1}{4\pi d^2} \frac{d/\sqrt{\lambda_a \lambda_e}}{\sinh(d/\sqrt{\lambda_a \lambda_e})} e^{-d/\lambda_a},
 \tag{5.18}$$

and for Cherenkov light emitted at the Cherenkov angle θ_c is

$$\mu(d) = l_0 A \cdot \frac{1}{2\pi \sin \theta_c} e^{-d/\lambda_p} \frac{1}{\sqrt{\lambda_\mu d \tanh \sqrt{d/\lambda_\mu}}},
 \tag{5.19}$$

where

$$\lambda_p = \sqrt{\frac{\lambda_a \lambda_e}{3}},
 \tag{5.20}$$

$$\sqrt{\lambda_\mu} = \frac{\lambda_c}{\sin \theta_c} \sqrt{\frac{2}{\pi \lambda_p}},
 \tag{5.21}$$




Figure 5.5: The geometry used for energy reconstruction in MuEx. Light detected by a DOM at r_f after being emitted at r_i (red) at angle θ_c for Cherenkov emission or at r_i (magenta) for isotropic emission from a cascade. Figure from [96].

and l_0 is the average number of photons emitted per unit distance of track.

Light lost to scattering or absorption is accounted for by incorporating the ice model. Light that is emitted in one ice layer and detected in another is treated as in [98]:

$$\frac{1}{\lambda_{a,e}} = \frac{1}{|\vec{r}_f - \vec{r}_i|} \int_{\vec{r}_i}^{\vec{r}_f} \frac{1}{\lambda_{a,e}(\vec{r})} |d\vec{r}| \quad (5.22)$$

where the effective λ for scattering, e , or absorption, a , is the average over the line between the emission, \vec{r}_i , and detection, \vec{r}_f . For Cherenkov light, the average is taken at the Cherenkov angle θ_c , while for cascades the average is taken perpendicular to the track at the point of closest approach.

Figure 5.6 shows the performance of the muon bundle energy loss reconstruction for proton (blue) and iron (red contours) simulation. The muon bundle energy loss relative to the true energy is plotted versus the logarithm of the true energy. Due to the much larger fluctuations in the first few interactions for protons and the fewer muons per bundle, the muon bundle energy loss distribution for proton primaries is much wider than for iron primaries. Iron primary cosmic rays are more likely to interact high in the atmosphere where charged pions have an increased likelihood to decay into muons. This effect, combined with a greater number of initial nucleons, leads to a greater number of high-energy muons reaching the IceCube detector, and a higher reconstructed muon bundle energy loss. Both primaries have a very linear behavior, with the mean values well separated across the entire energy range, making this a very powerful composition-sensitive parameter.

Figure 5.7 shows the separation of the mean muon bundle energy for proton (blue) and iron (red) simulations, with the rms shown by vertical error bars. The mean values are well separated for most of the distribution except at the lowest energies where selection effects dominate. Since a minimum amount of light is required for triggering and successful reconstruction, only events meeting this minimum will be present in the final sample, leading to an overlap at the lowest energies.




Figure 5.6: MuEx muon bundle energy loss performance. The logarithm of the reconstructed muon bundle energy loss (with the logarithm of the true primary energy subtracted) versus the logarithm of the true energy is shown for proton (blue) and iron (red contours) simulation.




Figure 5.7: The logarithm of the median reconstructed muon bundle energy loss versus the logarithm of the true energy for proton (blue) and iron (red) simulation. The central 68% is shown as vertical error bars.

Chapter 6

Analysis

“The measurement of the mass of the primary cosmic rays is vital for testing any theory of cosmic-ray origin. However, mass determination is the most difficult task facing an air shower physicist.”

A. A. Watson, 2002 [99]

6.1 Discussion of Previous Analyses

The analysis method described in this chapter is in some ways a departure from previous methods used to measure the cosmic-ray spectrum and composition in IceCube. This section will describe the analysis methods used for two recent results which were introduced in Section 1.6.

The analysis of the IceTop 73-station data from June 2010 to May 2011 [53] did not fit a composition, but instead employed a hypothesis-testing method. In the method, simulated events were weighted to a realistic spectrum and histogrammed in a two-dimensional space, the true primary energy and the reconstructed particle density at 125 meters from the reconstructed core of the shower. From that histogram, the average true energy was calculated for each bin in reconstructed particle density. The cosmic-ray spectrum was determined by assigning this energy to all events falling within the corresponding particle density bin. The hypothesis was checked by performing the analysis separately for multiple bins in zenith angle. Since the cosmic-ray flux is nearly isotropic [100], the resulting all-particle spectrum should be the same for all zenith angle bins. Any mismatch may be the result of a mismatch between the data and the simulation (atmosphere, hadronic interaction model, detector model) or weighting (composition, energy spectrum). Using simulation made with the SIBYLL hadronic interaction model with CORSIKA, the analysis found that using an H4a cosmic-ray model resulted in a more consistent flux between zenith bins than an all-proton or all-iron spectrum. All other possible systematic effects were treated as a linear scaling of the

particle density resulting in a shift in the flux, and were not used in the hypothesis test. While the H4a composition model did fit marginally better than the other two compositions tested, the fluxes at each zenith angle did not agree within statistical errors, suggesting a further systematic error. Since the actual cosmic-ray composition was not known and systematic errors were not sufficiently eliminated, the method was not a reliable test of either.

The air shower reconstruction software used by that analysis was very basic. It simply fit the lateral signal distribution in using a parametric curve, and returned the value of that curve evaluated at 125 m as a particle density. A snow-correction factor was applied to the signal at each tank to account for the change in the tank response for electromagnetic particles(see Figure 4.3) with snow depth. The ideal value of the correction factor depends on the ratio of electromagnetic particles to muons. This ratio depends on the primary energy, primary type, and distance from the shower core. Due to the simplicity of the reconstruction, these values are not known when the correction was applied, so a single value was used. ShowerLLH bypasses this complication by simply including these parameters as indices for the signal likelihood tables.

Like the present analysis, the analysis of IceTop and IceCube 40-string data from August 2008 [52] used events measured in both the IceTop and IceCube detectors. The analysis used reconstructions based on likelihood minimizations of parametric equations to fit a parameter from each detector; the air-shower particle density from IceTop and the photon density from muons traveling through IceCube. A neural-network trained on simulation was used to map these values onto primary energy and mass values. The events were then divided into bins by neural-network derived primary energy to determine the all-particle spectrum. The mass-value distribution in each bin was independently fit using templates derived from simulation of five primary masses which determined the average mass result for that bin.


That method had a number of attractive properties, but they came at a cost. The likelihood minimizations of parametric equations used to fit the events were responsible for progressively lower reconstruction efficiency at higher energies. They also made the measurement so sensitive to atmospheric and snow depth changes that only one month of half of the physical detector area could be analyzed. The neural-network mapping enabled a composition-independent energy spectrum to be produced easily but required careful simulation dependent-training and testing. The neural network also had the property that any, potentially unphysical, input will be mapped into a valid combination of energy and mass. Since the energy and mass were fit separately, composition-dependent corrections to the energy could not be applied to refine the energy result and required the use of large energy bins to limit the relative amount of migration, thus reducing the energy resolution of the analysis. Finally, the mass fit was so unstable that one element was removed, two were combined and the result was still too unstable to produce individual spectra. The neural-network guaranteed the result to be within the domain of the input simulation so the extensive systematic error tests only resulted

in a change in the spectrum and composition and gave no indication as to if they improved or worsened the measurement.

Lessons learned from that analysis have informed the design choices for the present analysis. The reconstruction used for IceTop, ShowerLLH, uses tables from realistic air-shower and detector simulations, limiting the parameter space of the fit using a brute-force scan of the parameter space to guarantee fit success. Likewise the reconstruction used for IceCube, MuEx, uses realistic scattering models and always converges. The fit is performed directly on the reconstructed parameters to avoid hiding mismatches between the data and simulation. By using three parameters, the fit is also over-constrained, which provides an additional check on agreement between data and simulation. Since both energy and composition are fit simultaneously, the result is self-consistent at all energies, allowing the use of much smaller energy bins.


6.2 Analysis Overview

This analysis is a simultaneous multi-dimensional fit using three composition-sensitive reconstructed parameters: primary energy and primary likelihood ratio from ShowerLLH (see Section 5.2.1) and muon bundle energy from MuEx (see Section 5.3.5). The three-dimensional distribution of these events from the data sample (see Section 6.3) is fit by the linear sum (see Section 6.8.1) of distributions from simulation binned in primary mass and energy (see Section 6.7). Errors are estimated by Poisson resampling of the data event distribution (see Section 6.10), and fluctuations caused by statistical noise are smoothed out using regularization (see Section 6.11). A graphical representation of the fit is shown in Figure 6.1.



(a) Proton simulation

(b) Iron simulation



(c) Data

Figure 6.1: Event distribution in muon energy loss versus reconstructed primary energy for simulation and data. This analysis uses event distributions from simulation, like those in Figures 6.1a and 6.1b (along with helium, oxygen and silicon), to fit distributions of data events as in Figure 6.1c. The actual fitting uses the full energy range of simulation, five primary masses, and is performed on three-dimensional distributions. These distributions are projected onto an orthogonal plane in Figure 6.2.

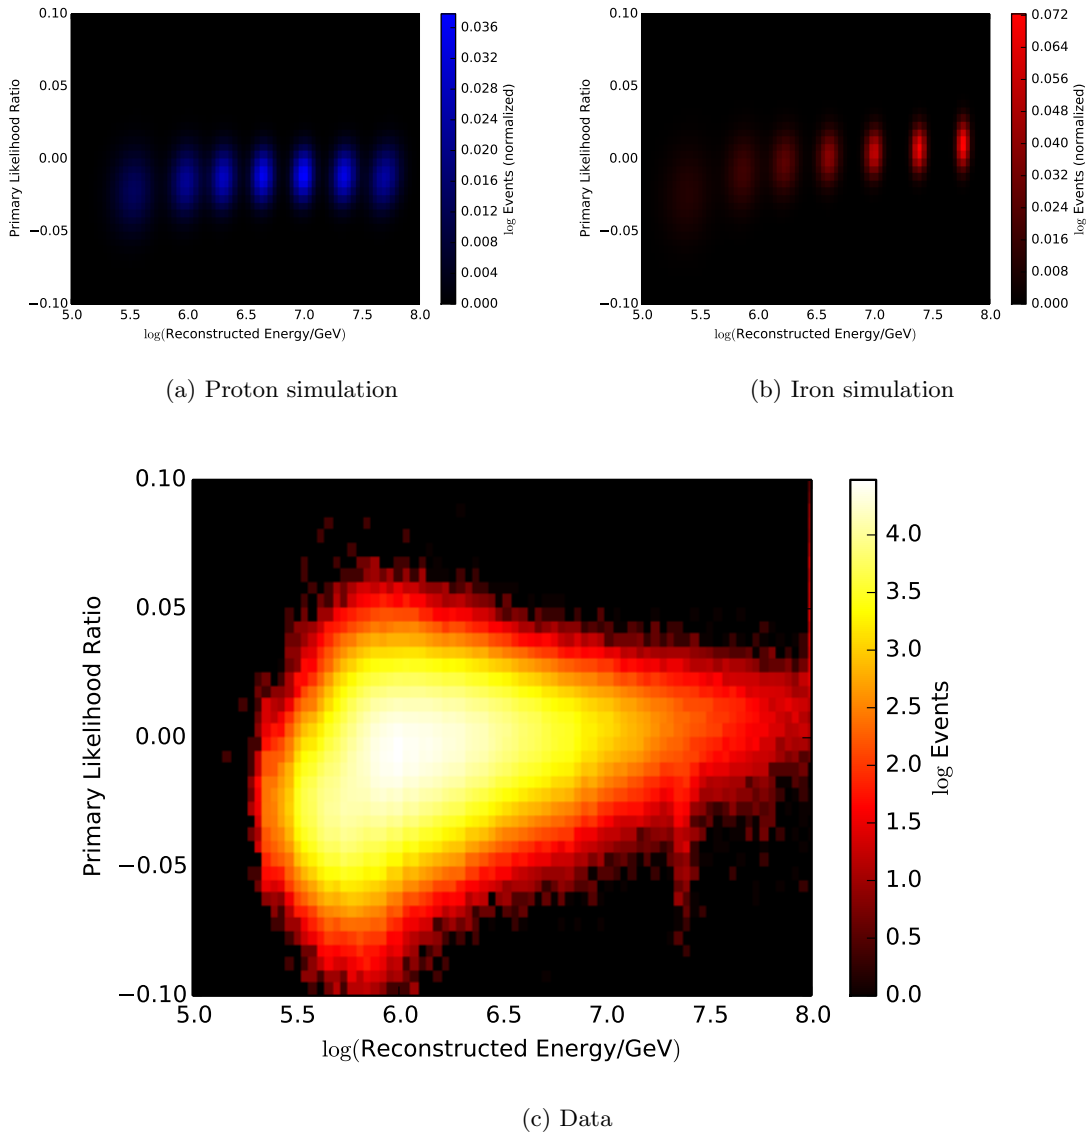


Figure 6.2: Distributions for simulation and data projected onto the primary likelihood ratio versus reconstructed energy plane, but otherwise the same as Figure 6.1.

6.3 Data Sample

The data used in this analysis is from the IC-59 standard physics run configuration. This data, taken from May 20, 2009 until June 1, 2010, includes all IceTop stations and IceCube strings installed by the end of the 2008/2009 South Pole summer season.

All IceCube science data is organized in runs, with a set detector configuration. For the IC-59 configuration, each run is nominally 8 hours long, while some can be shorter as needed for detector operations. Each run is statistically and qualitatively compared to historical runs to ensure detector stability, proper configuration and correct operation of the configured triggers and filters. The comparison includes the overall trigger rate for each sub-detector, reconstruction and filter rates, and the trigger rate for each DOM. The result of this process is the Good Run List, which is compiled by the run coordinator.

The Good Run List includes the date and time of the run start and end, the length of the run, which detectors completely or partially contributed to the run, and any anomalies found in the process of monitoring. Only runs with complete IceTop and IceCube detectors, with no errors found, were used in this analysis. The total time of all runs passing this criteria is 28744488 seconds, or 0.91 years.

In 2009, the IceCube DAQ [101] implemented two types of triggers for physics data. The first is a Simple Majority Trigger (SMT) for both IceCube and IceTop. The IceTop SMT required 6-fold multiplicity and IceCube SMT required 8-fold multiplicity, each in a 5000ns time window. The second is a string trigger, requiring 5 of 7 adjacent DOMs on one string in a time window of 1500ns.

The data rate produced by the IceCube DAQ far exceeds the available bandwidth allotted to IceCube satellite data transfers from the South Pole, so each event is checked against a set of filters and ‘prescales’ to determine if the event should be transferred over the satellite connection. Only every N^{th} event passes a prescale, making the prescale a data reduction factor. The set of raw data sent over the satellite constitutes the “Level 0” data.

6.4 Event Selection

For this analysis, filters defined by the Cosmic-Ray Working Group were used. The filter definitions can be found in Table 6.1. To take advantage of independent measurements of the air shower from IceTop and IceCube, events passing the “IceTopSTA3 InIceSMT” filter were selected to run through subsequent reconstructions. The only difference in triggering and filtering for simulation and data is that simulation has all prescales set to 1, to make the best use of limited statistics.

Table 6.2 describes the quality cuts performed on the data and simulation. The cuts break down into three general groups: quality, containment and agreement. The quality cuts generally ensure that the event was reconstructable and that a meaningful result was obtained. The containment cuts are particularly important

Filter name	Condition	Prescale
1. IceTopSTA3:	IceTopSMT & #Stations ≥ 3	8
2. IceTopSTA8:	IceTopSMT & #Stations ≥ 8	1
3. IceTopSTA3 InIceSMT:	IceTopSTA3 & InIceSMT	3
4. IceTopSTA8 InIceSMT:	IceTopSTA8 & InIceSMT	1
5. InIceSMT IceTopCoincidence:	InIceSMT & #IceTop-DOMs ≥ 1	80

Table 6.1: The cosmic-ray filters for the IC-59 [102] data sample used in this analysis.

Name	Quantity	Criteria
Fit success	ShowerLLH, MuEx	Fit succeeded
ShowerLLH quality	ShowerLLH	Energy uncertainty <i>hard</i> cut
Primary agreement	ShowerLLH	Proton- and iron-hypothesis shower cores less than 15 m apart
IceTop reconstruction containment	ShowerLLH iron result	Nearest IceTop station not on array perimeter
IceTop signal containment	IceTop calibrated pulses	Largest pulse not on array perimeter
MuEx containment	MuEx	Track passes through IceCube Volume
Agreement	ShowerLLH, MuEx	MuEx track passes within 75 meters of ShowerLLH core

Table 6.2: Quality cuts performed on data and simulation reconstruction values.

for IC-59. Due to the concave shape of the array an event can be within some parts of the array while still outside of the instrumented region. The agreement cuts are to ensure that the IceTop reconstruction and ShowerLLH correspond to the same shower.

6.5 Building Distributions

The observables used in this analysis are the primary-particle energy and primary likelihood ratio as reconstructed by ShowerLLH with IceTop data and the muon-bundle energy-loss as reconstructed by MuEx with IceCube data. These three parameters are binned into a three-dimensional histogram with binning and ranges according to Table 6.3.

Events in the data sample are histogrammed separately for each run, and by filter. They are subject to different prescales depending on the filter condition passed (see the prescale column of Table 6.1), so the statistical errors must be accounted for separately.

6.5.1 Simulation Energy Bin Definitions

The simulated events from each of the five primary particles (proton, helium, oxygen, silicon and iron) are split into 80 energy bins, equally spaced in the logarithm of the energy. The simulated energy range is $5.0 \leq \log_{10}(E/\text{GeV}) \leq 8.0$. The initial bin number is 0 and the last bin number is 79. The energy range for bin n is

$$5.0 + \frac{(8.0 - 5.0)n}{80} \leq \log(E/\text{GeV}) \leq 5.0 + \frac{(8.0 - 5.0)(n + 1)}{80} . \quad (6.1)$$


The fit data points are given for the center of each bin. The center of bin n is at

$$5.0 + \frac{(8.0 - 5.0)(n + 0.5)}{80} . \quad (6.2)$$


Simulated events are generated on an E^{-1} spectrum, so there is an equal number of generated showers for each energy bin. Figure 6.3 shows events from bin 45 of the iron simulation histogrammed for each parameter. The distributions were shifted to a mean of zero and scaled to a standard deviation of 1. Each distribution is Gaussian within statistical errors. Figure 6.4 shows the elliptical 1σ -area of each of the five primaries in the primary likelihood and muon bundle energy loss space. These two parameters, along with the reconstructed energy are used to differentiate primary mass composition, so they must be separated in parameter space.

Quantity	Number of bins	Scale	Range	Units
Muon-bundle energy	40	Logarithmic	$10^2 - 10^7$	GeV/m
Primary likelihood ratio	40	Logarithmic	-0.1 - 0.1	-
Primary energy	80	Logarithmic	$10^5 - 10^8$	GeV


Table 6.3: The dimensions of the histograms used for unfolding.



(a) Reconstructed primary energy



(b) Reconstructed muon energy



(c) Primary likelihood ratio

Figure 6.3: Distributions from iron primary simulation with the mean shifted to 0 and the standard deviation rescaled to 1, for true energy bin 45 (around 5 PeV). Also plotted is a best fit Gaussian (red) and a Gaussian with a normalization equal to the number of events and $\mu=0$ and $\sigma=1$ (blue). The Gaussian fit shows overall good agreement with the calculated mean and standard deviation.




Figure 6.4: The separation between the 5 components for energy bin 45. The filled ellipse represents the fit 1σ width in primary likelihood vs muon bundle energy space for each component. Starting in the lower left, the proton ellipse is plotted in blue, helium in cyan, oxygen in green, silicon in magenta, and iron in red.

6.6 Simulation Weighting

In the fitting procedure, a weight is assigned to each primary and energy bin, with each event in that bin receiving the same weight. The simulation was produced with an energy spectrum of E^{-1} , but the cosmic-ray spectrum has been measured to have an energy spectrum $\approx E^{-2.7}$. Without re-weighting, the high energy events in each bin will have a disproportionately large contribution to each histogram. To reduce systematic errors, the simulation is given an additional weight of $E^{-1.7}$ to limit the effect of either a very hard spectrum or a very soft spectrum.

6.7 Fitting Simulation

The full CORSIKA simulations (see Section 4.1.1) required by the analysis are very resource intensive, so simulation statistics cannot reach even a tiny fraction of the data live time. Since this analysis attempts to make the best use of the mass and energy resolution of the detector, the statistics must be separated into relatively small mass and energy bins, which further increases the effect of statistical fluctuations.

The collisions of cosmic rays in the atmosphere lead to complicated interactions, but the average behavior of these interactions vary smoothly and gradually as a function of energy and by mass. Since the behavior of the interactions and the resulting observables change on scales much larger than the bin width, we can use smoothing to remove the noise from the simulated event distributions.

The first step is to fit the overall behavior of the distributions across energy bins. This step is best achieved by either fitting a specific function using prior knowledge of the underlying distribution, or by fitting smooth splines where the underlying distribution is unknown. In either case, correctly treating the underlying statistics is crucial to avoid under- or over-fitting the distribution.

6.7.1 Effective Area

The flux cosmic-ray spectrum at Earth is typically described in terms of a differential intensity. An intensity is defined as the number of events N passing through an area A over a given solid angle Ω and with an energy per nucleus E . The differential intensity is written as

$$\frac{dN}{dE dA d\Omega dt} (\text{GeV}^{-1} \text{m}^{-2} \text{s}^{-1} \text{sr}^{-1}). \quad (6.3)$$

To calculate the intensity for an experiment, it is helpful to combine the experiment-dependent parameters, the area and solid angle into an *effective area*. The effective area is a function of primary mass and energy, which is defined as the equivalent size of a detector which triggers on all events passing through it, and no events passing outside it. Since the effective area is used to account for the acceptance angle, this number is usually much smaller than the physical area of the detector. For this analysis, it is calculated as the product

of the sampling area, the solid angle of generated events, and the probability of an event contributing to a subsample. An effective area can be determined for the subsamples of events that pass triggering, filtering, or specific quality cuts. Figure 6.5 shows the effective area for iron primaries detected in IceTop and IceCube after all quality cuts as a function of energy. The effective area shows three main regions. At low primary energy, the air shower does not have sufficient energy to produce enough particles on the ground to trigger the IceTop detector. At higher primary energies, all showers will produce enough particles to trigger the detector, and the detection efficiency nears 100% for showers that fall within the cut-area for this analysis.

In the intermediate range, the trigger efficiency increases from nearly 0% to nearly 100%. For an air-shower of a given energy, the density of electromagnetic particles at some distance from the shower core is described by a Gaussian distribution in the logarithm of the density, as it is the product of several generations of interactions as the shower passes through the atmosphere. Since each IceTop station has a trigger threshold that corresponds to a certain density of electromagnetic particles, the probability Φ that a shower will trigger a station at a given distance is given by the cumulative distribution function of the normal distribution (replacing the particle number with the logarithm of the particle number),

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-t^2/2} dt, \quad (6.4)$$

which can be written in terms of the error function

$$\text{erf}(x) \equiv \frac{1}{\sqrt{\pi}} \int_{-x}^x e^{-t^2} dt, \quad (6.5)$$

as

$$\Phi(x) = \frac{1}{2} \left[1 + \text{erf} \left(\frac{x}{\sqrt{2}} \right) \right]. \quad (6.6)$$

As an example, Figure 6.5 shows the fit of the Gaussian error function to the effective area as a function of energy for iron simulation.

6.7.1.1 Effective Area Statistics with Oversampling

The most relevant variables determining the probability that an air shower will trigger the detector are the primary energy, mass, and the height of the first interaction point. Since these variables are all set when the CORSIKA simulation is run, the number of CORSIKA showers generated determines the statistical fluctuations for the effective area. Oversampling (see Section 4.5) the showers has the effect of better determining the probability of triggering, but does not increase the CORSIKA statistics.

For each bin of the effective area histogram, the number of times that a given CORSIKA shower occurs in the final sample can be treated as the weight for that CORSIKA shower. Then the value of a histogram




Figure 6.5: Fitting the effective area of iron simulation with a Gaussian error function. *Top*: The number of events passing the quality cuts as a function of energy bin (red), with error bars as described in Section 6.7.1.1. *Bottom-left*: The fit residuals, scaled by the error for each point. *Bottom-right*: The normalized distribution of the fit residuals in blue, compared to a normal distribution with $\mu = 0$ and $\sigma = 1$ in green.

bin B is the sum of the weights of the CORSIKA showers

$$B = \sum_{i=1}^N W_i . \quad (6.7)$$

This sum can be rearranged by grouping the showers by weight. If we have k unique weights, and N_k showers with a weight W_k , the sum can be written as

$$B = \sum_k W_k N_k . \quad (6.8)$$

Since all of the showers are independent, the variances add:

$$\begin{aligned} \text{var}(B) &= \text{var}\left(\sum_k W_k N_k\right) \\ &= \sum_k \text{var}(W_k N_k) \\ &= \sum_k W_k^2 \text{var}(N_k) \\ &= \sum_k W_k^2 N_k . \end{aligned} \quad (6.9)$$

Rewriting this last equation, we can rearrange back to indexing over the individual showers,

$$\text{var}(B) = \sum_{i=1}^N W_k^2 , \quad (6.10)$$

and take the square root to get the error,

$$\sigma_B = \sqrt{\sum_{i=1}^N W_k^2} . \quad (6.11)$$

We see that the error of the number of events in the bin is simply the square root of the sum of the squares of the weights of each shower.

We can also define an equivalent number of events,

$$N_{equ} = \frac{\left(\sum_{i=1}^N W_i\right)^2}{\sum_{i=1}^N W_i^2} , \quad (6.12)$$

which relates the number N of weighted showers with the number N_{equ} of events with weight 1 that would have the same statistical fluctuation.

6.7.2 Distribution Mean, Variance and Covariance

The distribution of simulated events are normally distributed (see Figure 6.3) in each of the three observables: the logarithm of the reconstructed energy, the logarithm of the reconstructed muon bundle energy,

and the mass likelihood ratio. Without making any assumption about the independence of each variable, the distributions can be completely described by the mean in each dimension, and the six non-degenerate values of the nine-element covariance matrix. The errors on these values are easily calculable from the number of events in the sample.

Figure 6.6 shows the mean values for the proton and iron simulations, and Figures 6.7 and 6.8 show the variance and covariance terms (the remaining three covariance terms do not have any additional information). Once the means and the covariance matrix have been calculated for each primary and energy bin, the values are smoothed using splines according to their statistical weight, as each of these parameters are smoothly varying with the primary energy. The green line in each figure is the spline fit to the iron simulation. By fitting each parameter with a spline, the statistical noise, particularly at high energies or low primary mass, is greatly reduced.

Figure 6.9 shows the bias of the smoothed mean reconstructed energy. At low energies, the trigger criteria cause a selection bias, so most events are reconstructed with a higher energy. At high energies, the light primaries have larger statistical fluctuations due to the greater variation in their first interaction height. Since the simulation and tables are limited to a maximum of 10^8 GeV, reconstructed energies for the last primary energy bin can only fluctuate down to lower energies.


Figure 6.10 shows the smoothed mean muon-bundle energy loss. For the lowest energy primary cosmic rays, protons are able to create higher energy muons due to the higher energy per nucleon. At higher energies, the energy per nucleon is high enough that all primaries can readily produce muons capable of reaching IceCube. Showers of low mass primaries typically develop lower in the atmosphere, where the higher air density makes pion decay less likely and fewer muons are produced.

Figure 6.11 shows the smoothed mean primary likelihood ratio. At the lowest energies, few triggering IceTop tanks and large fluctuations make the likelihood ratio similar for most showers. At higher energies, increased statistics make the differences in shower development more apparent and the distributions diverge.


These three parameters: the reconstructed energy, the muon bundle energy loss, and primary likelihood ratio combine to make a three-dimensional distribution with each primary well separated in the parameter space. Figure 6.12 shows the 3-dimensional surface formed by the position of the mean for all three parameters for each primary and energy bin. The smooth progression for each primary is shown by a thick solid line, and points corresponding to the same primary energy are connected with a thin colored line.

6.7.3 Smoothing Simulation Parameter Distributions


We are limited in the amount of simulation statistics we can produce, but to reduce the effect of statistical errors on the final result it is desirable to smooth the distributions from simulation. Once the smooth splines have been calculated as in Section 6.7.2, we can make smoother versions of the simulated distributions



(a) Reconstructed muon energy




(b) Primary likelihood ratio




(c) Reconstructed primary energy


Figure 6.6: Smoothed mean values in each observable by simulation bin. Plotted in blue are the values for proton simulation, in red the iron simulation, and the green line is the best spline fit to the iron simulation. Figure 6.6a shows the fit to the reconstructed muon energy, Figure 6.6b shows the primary likelihood ratio, and Figure 6.6c shows the reconstructed primary energy mean value by bin. See Figure 6.7 for the variance and Figure 6.8 for the covariance between these values.



(a) Reconstructed muon energy variance




(b) Primary likelihood ratio variance




(c) Reconstructed primary energy variance


Figure 6.7: Smoothed variance values for each observable in Figure 6.6. Variance for proton simulation is plotted in blue, and iron simulation is plotted in red. The variance for iron, smoothed with best fit splines, is plotted in green.



(a) Likelihood ratio and muon bundle energy covariance



(b) Primary energy and muon bundle energy covariance



(c) Primary energy and likelihood ratio covariance

Figure 6.8: The non-degenerate terms of the covariance matrix for the three observables plotted in Figure 6.6. The covariances for proton simulation is plotted in blue, and iron simulation is plotted in red. The covariances for iron, smoothed with best fit splines, is plotted in green. The covariance is small compared to the variance (see Figure 6.7).




Figure 6.9: The smoothed mean reconstructed energy bias as a function of primary energy for all five components of the simulation.




Figure 6.10: The smoothed mean muon bundle energy loss as a function of primary energy for all five components of the simulation relative to proton.




Figure 6.11: The smoothed mean primary likelihood ratio for all five components of the simulation as a function of primary energy.





Figure 6.12: The three-dimensional parameter space for the smoothed five component simulation. The smoothed 3-dimensional mean value for each energy bin of each primary is plotted in thick lines. Points with the same primary energy are connected by multi-colored thin lines. This plot shows that combining the information from three parameters at one time breaks degeneracies that may be present for fewer parameters.

by shifting data events from adjacent energy bins, or using the smoothed statistics directly to generate three-dimensional Gaussian distributions.


Assuming only that the descriptive statistics vary smoothly between energy bins, but without making any assumptions about the shape of the distribution itself, one method for smoothing the simulation statistics is to shift the data points from neighboring bins by the shift in the mean of the distribution, and rescale by the difference in the variance, and reweight to the smoothed effective area as shown in Figure 6.13.

However, we can use the Gaussian nature of the distribution, as verified in Figure 6.13, to simplify the procedure, by using the mean and covariance matrix, along with the normalization from the effective area to calculate the smoothed value of each bin in the 3-dimensional distribution directly. Integrating the 3-dimension Gaussian is performed using Fortran code (described in [103]) with a convenient interface from the stats library of SciPy [104].


From Figure 6.3 and Figure 6.13 we can see that the Gaussian description of each of the parameters is approximately valid. Figure 6.14 shows a reconstructed-energy slice in the three-dimensional event distribution for data (center), simulation (left), and simulation after smoothing (right). The simulation has been weighted according to the H4a model for the same live time as the data for illustrative purposes. The smoothed simulation shows reduced statistical noise compared to the simulation before smoothing, and appears much more similar to the data. By smoothing the simulation, we reduce the effect of statistical errors on the fit in order to better identify and correct for systematic errors.



(a) Iron bin 45



(b) with nearest 4 bins



(c) with nearest 10 bins

Figure 6.13: Reconstructed primary energy distributions from iron primary simulation, shifted by subtracting the mean and then rescaled by the variance. In Figure 6.13a only statistics from bin 45 are used. In Figure 6.13b and Figure 6.13c, statistics from the nearest 2 and 5 bins, respectively, in each direction are also shifted and rescaled and added to the distribution, and then reweighted to the original effective area.




Figure 6.14: The primary likelihood versus muon bundle energy distribution of events with a reconstructed energy near 10^7 GeV (bin 56). The simulation is weighted according to the H4a model for the same live time as the data. The smoothed simulation shows a similar amount of bin to bin noise as the data, while the unsmoothed simulated event distribution is dominated by statistical noise.

6.8 Unfolding Algorithm

6.8.1 Least Squares

The three-dimensional histogram of reconstructed parameters from data can be rearranged into a one-dimensional vector \vec{b} . Histograms from reconstructed parameters for simulation corresponding to primary mass and energy bins can be arranged the same way to make up columns of a two-dimensional matrix \mathbf{A} . The corresponding weight vector \vec{x} then describes the weight of each primary mass and energy bin to fit the data, and a vector \vec{n} represents the statistical noise for each bin of \vec{b} . These values can then be arranged into the equation

$$\vec{b} = \mathbf{A}\vec{x} + \vec{n}. \quad (6.13)$$

Approximating the noise as a Gaussian with a mean of 0 and a variance of σ^2 , the noise PDF is

$$p(n_i) = \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{n_i^2}{2\sigma_i^2}}. \quad (6.14)$$

From (6.13), $n_i = b_i - \mathbf{A}\vec{x}$, so the probability of observing \vec{b} with a weight vector \vec{x} and simulation response \mathbf{A} can be written as

$$p(b_i|\mathbf{A}; x_i) = \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{(b_i - \mathbf{A}x_i)^2}{2\sigma_i^2}}. \quad (6.15)$$

The likelihood \mathcal{L} of \vec{x} based on the set of simulation \mathbf{A} and data \vec{b} is

$$\mathcal{L}(\vec{x}) = p(b_i|\mathbf{A}; x_i). \quad (6.16)$$

For histograms with m bins, all m observations are independent, so the overall likelihood of \vec{x} is

$$\begin{aligned} L(\vec{x}) &= \prod_{i=1}^m p(b_i|\mathbf{A}; x_i) \\ &= \prod_{i=1}^m \frac{1}{\sigma_i\sqrt{2\pi}} e^{-\frac{(b_i - \mathbf{A}x_i)^2}{2\sigma_i^2}}. \end{aligned} \quad (6.17)$$

Taking the natural logarithm of the likelihood function, and assuming that the local changes in the variance can be ignored ($\sigma_i = \sigma$)

$$l(\vec{x}) = \ln \mathcal{L}(\vec{x}) \quad (6.18)$$

$$\begin{aligned} &= \ln \left(\prod_{i=1}^m \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(b_i - \mathbf{A}x_i)^2}{2\sigma^2}} \right) \\ &= \sum_{i=1}^m \ln \left(\frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(b_i - \mathbf{A}x_i)^2}{2\sigma^2}} \right) \\ &= \underbrace{m \frac{1}{\sigma\sqrt{2\pi}} - \frac{1}{2\sigma^2}}_{\text{does not depend on } x} \sum_{i=1}^m (b_i - \mathbf{A}x_i)^2. \end{aligned} \quad (6.19)$$

So, a least-squares minimization is equivalent to maximizing the log-likelihood, $l(\vec{x})$,

$$\arg \min_{\vec{x}} \sum_{i=1}^m (b_i - Ax_i)^2 \equiv \arg \max_{\vec{x}} \mathcal{L}(\vec{x}) . \quad (6.20)$$

6.8.2 Constrained Least Squares

It would be easy to find the solution to a least-squared problem using a simple method such as

$$\vec{x} = \left(\mathbf{A}^T \mathbf{A} \right)^{-1} \mathbf{A}^T \vec{b} , \quad (6.21)$$

but this solution, while easy to compute, leaves the domain of the elements of \vec{x} unconstrained. The elements of \vec{x} are linear factors of the flux of cosmic rays, so a negative value would be unphysical. The NNLS (Non-Negative Least Squares) algorithm [105] provides an efficient and mathematically rigorous way to implement this constraint.

The NNLS algorithm solves:

$$\begin{aligned} \arg \min_{\vec{x}} \quad & \left\| \mathbf{A}\vec{x} - \vec{b} \right\|_2^2 \\ \text{subject to} \quad & \vec{x} \in \mathbb{R}_{\leq 0} , \end{aligned} \quad (6.22)$$

where $\vec{x} \in \mathbb{R}^n$, $\vec{b} \in \mathbb{R}^m$, \mathbf{A} is an $m \times n$ matrix and $\|\dots\|_2$ denotes the L_2 norm. The L_2 norm is the square-root of the sum of the square of the absolute value of each element. The number of iterations for the algorithm is provably finite, and in practice the computation time is no more than tens of seconds on contemporary computers for the most complicated situations.

6.9 Performance on a Simulated Spectrum

The method can be tested by using a realistic model of the cosmic-ray spectrum and detector live time to weight the smoothed simulation event distributions from Section 6.7.3. To simulate the effects of limited statistics on the input and measured spectrum, the average contribution from each primary and true-energy bin is resampled according to a Poisson distribution to give an integer number of events in each three-dimensional bin of the simulated event distribution.

The method is tested by applying it to a simulated data set with the energy spectrum and composition taken from the H4a model for the livetime of 0.91 years of IC-59. Figures 6.15 and 6.16 show the model flux with error bars representing the $1\text{-}\sigma$ statistical uncertainty for 0.91 years of lifetime. The mean natural logarithm of the mass number is shown in Figure 6.17. Information is lost in Poisson resampling and even an ideal analysis would not recreate these fluxes perfectly. This simulated flux is unfolded using the method described in Section 6.8.2, and the resulting flux result is shown in Figure 6.18. Systematic errors in the simulation, which are absent in this test by construction, would show correlated residuals in the unfolded

distribution above the level of the statistical noise shown in the fit residuals. The fit residuals, shown in Figure 6.19, are generally featureless, as would be expected when unfolding a simulated flux using the same simulation. In Figure 6.19b, the large fluctuations at the edge of the distribution are due to the very low counts in those bins, and the random noise in the interior of the distribution are due to the statistical noise imposed by Poisson sampling. The large fluctuations at high energies in Figure 6.19a are also due to large relative fluctuations for very low count bins.




Figure 6.15: The simulated H4a flux. The error bars represent the $1\text{-}\sigma$ statistical uncertainty for 0.91 years of lifetime. The points for each primary are slightly offset to the right of the all-particle points for clarity. This plot uses the same color convention used for the rest of this document: all-particle (black), proton (blue), helium (cyan), oxygen (green), silicon (magenta) and iron (red).




Figure 6.16: This unweighted simulated H4a flux. As in Figure 6.15, the error bars represent the $1\text{-}\sigma$ statistical uncertainty for 0.91 years of lifetime. The individual component points are offset to the right to avoid overlapping.




Figure 6.17: The mean logarithm of the mass number of the simulated spectrum shown in Figure 6.15. The error bars represent the $1\text{-}\sigma$ statistical uncertainty for 0.91 years of lifetime.




Figure 6.18: The weighted unfolded flux of the simulation shown in Figure 6.15. The unfolding is performed using the method described in Section 6.8.2 and represents the single best fit to the simulated flux. Since the simulation is weighted to an $E^{-2.7}$ spectrum, this result is a direct output from the NNLS algorithm. The all-particle flux generally follows the model, while the individual component fluxes vary from bin to bin.

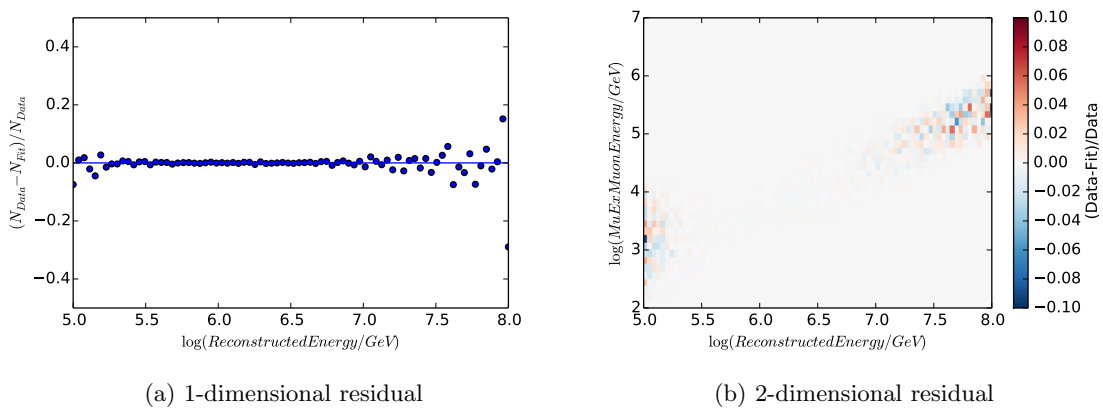


Figure 6.19: The relative fit residuals for the unfolding result for the simulated spectrum shown projected onto the IceTop reconstructed-energy axis in Figure 6.19a and into the muon-energy/IceTop reconstructed-energy plane in Figure 6.19b. A blue line at 0 is shown to guide the eye in Figure 6.19a.

6.10 Error Estimation

The fit is performed on a histogram of the data. Since the number of events in each bin is an independent measurement of the rate of events which satisfy the condition of having reconstructed values between the respective bin edges, the statistical fluctuations follow a Poisson distribution. To estimate the statistical errors of the fit, the fit could be performed on many alternate realizations of the data sample. In this situation, it is not practical to experimentally collect multiple data samples, and the true mean rate of events is not known. However, we can make an upper bound on the statistical errors by Poisson resampling the number of events in each bin using the measured rate of each bin as an estimate of the true rate.

By resampling and unfolding 100 times, best fit and error bars for each element of the result vector \vec{x} can be estimated by returning the median result as the best fit and the central 68 percentile region as the error bars.

Figure 6.20 shows the weighted unfolding result for the simulated H4a spectrum. The unfolding has been performed 100 times, with the median unfolding result represented by the points, and the range of the inner 68 percentiles represented by the error bars. The fluctuations are more pronounced than in Figure 6.18, as the already sampled distribution is used as the mean for the 100 resampled distributions used for estimating the errors. The all-particle points and error bars are determined using the same method on the set of the 100 all-particle results. The error bars show that a wide range of component fluxes could describe the simulated flux. A bin-to-bin oscillation in the all particle spectrum is also visible, while the individual error bars are relatively small. This is due to statistical fluctuations causing alternating over- and under-fitting in adjacent bins due to the overlapping event distributions.




Figure 6.20: The weighted fit result of the simulated flux shown in Figure 6.15.




Figure 6.21: The unweighted unfolded flux shown in Figure 6.20.




Figure 6.22: The mean logarithm of the mass number for the unfolded flux shown in Figure 6.20. The fit shows a general agreement with Figure 6.17, while fluctuating at each bin around the average trend.

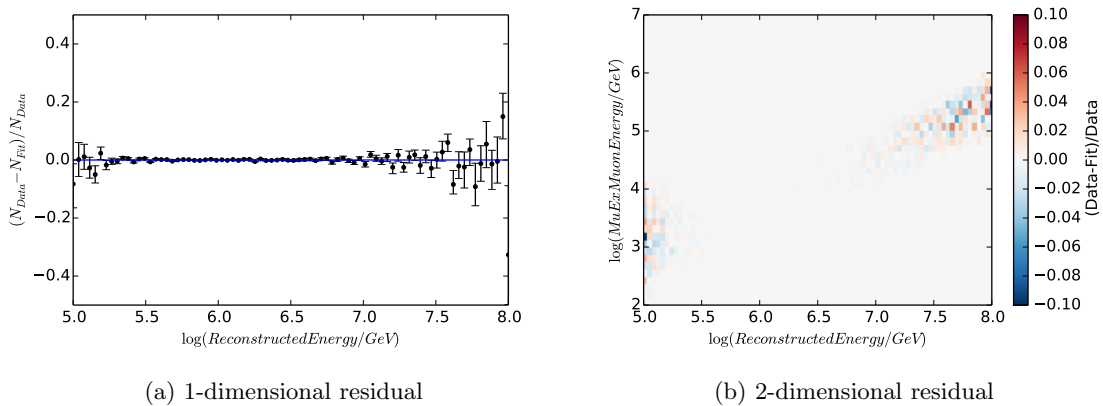


Figure 6.23: The relative fit residuals for the unfolding result for the simulated spectrum shown projected onto the IceTop reconstructed-energy axis in Figure 6.23a and onto the muon-energy/IceTop reconstructed-energy plane in Figure 6.23b.

6.11 Regularization

In Section 6.8.2, the physical constraint of positive flux was introduced to limit the range of solutions to those with a physical interpretation. The solution shown in Figure 6.18 includes another implicit constraint imposed by the energy range represented by each term in the solution (or the width of each bin). Since all events fit by a single term are represented by a single spectral index and normalization, the width of the bin sets the minimum size of spectral features the unfolding is sensitive to. As discussed in Section 6.6, the spectral index of narrow bins is insignificant, but it plays a larger role as the bins are expanded. Fixing the normalization across a wider energy range allows each point to be less influenced by statistical errors, while limiting the ability of the unfolding to fit the features of the data.

Figure 6.24 shows the same unfolding method used in Figure 6.20 except with twice as many bins. The all-particle spectrum varies greatly from bin to bin (by a factor of 8 in some places) to fit the statistical fluctuations imposed by the finite statistics of the simulated spectrum. The variation in the individual spectra is much greater and are omitted from the plot. The all particle spectrum suffers from *oscillation*, where the unfolded spectra will alternate from relatively high to relatively low for adjacent bins. This results from choosing bins smaller than the resolution of the reconstruction.

Figure 6.25 shows the effect of increasing the bin size. The increased bin size reduces the effect of overfitting the statistical noise, but it also limits the ability to measure any small features in the spectrum as predicted by certain models (especially [26]).

The choice of bin size is a simple example of *regularization*, which is the name for mathematical methods to choose a unique best solution for a problem that has many solutions that fit the data equally well. This is essentially the application of Occam's Razor: of the set of solutions that fit the data equally well, the simplest solution is best. This is supported by measurements of the composition at lower energies (see Section 1.3) and models of the spectrum around the knee (see Section 1.4); any change in composition is expected to occur slowly over several orders of magnitude in energy. We can apply this principle to the solution vector \vec{x} by applying a penalty term.

To penalize changes in the slope (the 2nd-derivative) of the flux of each component divided into n energy bins, each component is regularized by a 2nd order Tikhonov penalty matrix [106, sec. 19.5] of dimension $(n - 2) \times n$

$$\mathbf{P} = \begin{bmatrix} -1 & 2 & -1 & 0 & \dots \\ 0 & -1 & 2 & -1 & \dots \\ \vdots & & \ddots & & \dots \\ \dots & 0 & -1 & 2 & -1 \end{bmatrix}. \quad (6.23)$$




Figure 6.24: Unfolding a simulated H4a spectrum using the same method as shown in Figure 6.20, except with twice as many energy bins for each component. The individual spectra are omitted from this plot due to the very large error bars.




Figure 6.25: Unfolding a simulated H4a spectrum using the same method as shown in Figure 6.20, except with half as many energy bins for each component. By reducing the freedom of the unfolding by using larger bins, the result is less susceptible to statistical errors. However, any features that may be present in the spectrum would be washed out by the large bins.

The regularization penalty term competes with the least-squares penalty. We need to account for the fact that the cosmic-ray spectrum is a steeply falling power law, so the penalty matrix is multiplied by a weight matrix of dimension $n \times n$

$$\mathbf{W} = \begin{bmatrix} w_1 & 0 & 0 \\ 0 & w_2 & 0 \\ \vdots & \ddots & \vdots \\ 0 & 0 & w_n \end{bmatrix}. \quad (6.24)$$

The weights are given by summing the $E^{-2.7}$ weighted number of events in the corresponding simulation histogram.

The penalty matrices for each of the m components are combined into a full penalty matrix by stacking them along the diagonal of an $mn \times mn$ square matrix. For instance, for 3 components with 5 energy bins the penalty matrix is

$$\mathbf{P}_{3,5} = \begin{bmatrix} \begin{bmatrix} -1 & 2 & -1 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 \\ 0 & 0 & -1 & 2 & -1 \end{bmatrix} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \begin{bmatrix} -1 & 2 & -1 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 \\ 0 & 0 & -1 & 2 & -1 \end{bmatrix} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \begin{bmatrix} -1 & 2 & -1 & 0 & 0 \\ 0 & -1 & 2 & -1 & 0 \\ 0 & 0 & -1 & 2 & -1 \end{bmatrix} \end{bmatrix}. \quad (6.25)$$

In order to use the NNLS algorithm introduced in Section 6.8.2, we need to format the inputs into the form of

$$\vec{b} \approx \mathbf{A}\vec{x}. \quad (6.26)$$

To incorporate the regularization penalty term, we rewrite the equation as

$$\mathbf{A}^T \vec{b} \approx (\mathbf{A}^T \mathbf{A} + \lambda \mathbf{P}^T \mathbf{W} \mathbf{P}) \vec{x}, \quad (6.27)$$

where λ is the regularization term which needs to be optimized (see Section 6.11.1). We can rewrite (6.22) as

$$\begin{aligned} \arg \min_{\vec{x}} \quad & \left\| (\mathbf{A}^T \mathbf{A} + \lambda \mathbf{P}^T \mathbf{W} \mathbf{P}) \vec{x} - \mathbf{A}^T \vec{b} \right\|_2 \\ \text{subject to} \quad & \vec{x} \in \mathbb{R}_{\leq 0}. \end{aligned} \quad (6.28)$$

The scale for λ is determined by the average value of the solution vector $\langle \vec{x} \rangle$, the Poisson variations of the data \vec{b} , and the actual change in the power-law index of the individual components. The absolute scale

of the solution vector is determined by the absolute scale of the weights for simulated events (which are very small in this case), also making the absolute scale of λ arbitrary (in this case, typically $\lambda \sim 10^{-12}$). Since the change in composition is not known beforehand, the regularization strength must be optimized on the data.

6.11.1 Optimization

If the regularization is too weak, localized Poisson fluctuations cause the fit to oscillate between adjacent energy bins and primary elements in fit solution. If the regularization is too strong, the fit is dominated by the regularization term, and the solution fits a localized fluctuation in one part of the distribution while missing the rest. Both situations lead to a high statistical variance in the fit as estimated by resampling as in Section 6.10. The optimal regularization parameter λ is found by minimizing the sum of the variances of the terms of the fit. Figure 6.26 shows the sum of variances over a wide range of λ . The minimum value of λ is around $\lambda = 10^{-11}$. This value is used for the regularized fits shown in Figures 6.27 to 6.29.

Figures 6.27 and 6.28 show that the all-particle spectrum is well fit with very small estimated errors. The proton and iron spectra fit the input spectra well. They represent the lowest and highest mass primaries and are well constrained by the extreme edges of the distribution. The central component, oxygen, over-fits its input spectrum and that of the surrounding components, helium and silicon. In this case, the analysis is not sensitive to individual fluxes of the central components. However, the mean logarithm of the mass number result (Figure 6.29) does track the input model quite well, with error bars as expected from the statistics of the input spectrum (Figure 6.17).




Figure 6.26: The logarithm of the sum of the variance of the regularized fit points with a minimum around $\log(\lambda) = -11$. Over- or under-regularized fits give a higher variance, so minimizing this term gives the best regularization.




Figure 6.27: The weighted regularized unfolded flux of the simulated flux shown in Figure 6.15. The error bars represent the 68% confidence interval. The lines show the input spectrum before sampling.




Figure 6.28: The unweighted regularized unfolded flux shown in Figure 6.27. The error bars represent the 68% confidence interval.




Figure 6.29: The mean logarithm of the mass number for the regularized unfolded flux shown in Figure 6.27. The error bars represent the 68% confidence interval and the line represents the input H4a model. The composition shows agreement with the input model within the statistical errors. Since each component is regularized individually, the mean logarithm of the mass number is also regularized.

Chapter 7

Results

“Any sufficiently advanced technology is indistinguishable from magic.”

The third of Aurthur C. Clarke’s *laws*

In this chapter, the method developed in Chapter 6 is applied to data from the IC-59 detector. Seasonal variations in the atmosphere are measured and the data is normalized to a single month. Systematic uncertainties resulting from the choice of high-energy hadronic interaction models are studied and applied to the method. Finally, disagreements between the simulation and the data sample are evaluated and discussed.

7.1 Seasonal Corrections

The technique of measuring cosmic rays by their interactions in the atmosphere naturally requires a thorough understanding of the atmosphere and the impact of its dynamics on the measured air shower quantities. Due to the South Pole’s unique location, the primary variation in the atmospheric profile is due to the change in seasons. Figure 7.1 plots the atmospheric temperature as a function of altitude as measured by the AIRS instrument [107] on the Aqua satellite for twice daily measurements in June and December 2009. The measurements are taken at set pressure levels, starting at 1 hPa, leading to bands of constant pressure in the scatter plot. During the dark winter months (June), the atmosphere is overall much colder and denser. In the summer months (December), the atmosphere is less dense and extends further from the surface. The monthly variation is much larger than the daily variation within each of the plotted months.

The simulations generated for this analysis in CORSIKA use a single atmospheric model. This model, number 12 in CORSIKA, is fit from an evaluation of the MSIS-E-90 [89] atmosphere model for July 1, 1997. Figure 7.2 shows a comparison of the atmospheric profile as measured by the AIRS instrument for each day of the IC-59 dataset compared to each of the four South Pole atmosphere models in CORSIKA. The shape of the deviation of atmosphere model 12 from the measured atmosphere matches the MuEx offset shown in Figure 7.7, supporting the interpretation of the offset as an atmospheric correction. To generate this figure, the model pressure was evaluated at the height returned by AIRS for each of its pre-determined pressure




Figure 7.1: The atmospheric temperature as a function of altitude for June and December 2009, as measured by the Atmospheric Infrared Sounder [107] (AIRS) instrument on the Aqua satellite.




Figure 7.2: The relative rms deviation between the pressure at the heights measured by the AIRS instrument compared to the pressure used in CORSIKA atmosphere models 11-14. CORSIKA model 12 was used for the simulation, and is the best model of the atmosphere for months -1 to 3 and 11.

levels. The RMS of the relative difference in these two pressures is plotted for each month and atmosphere model, with the error bars showing the center 68% variation within each month. Atmosphere number 12 is the best model for the most months and shows a relatively low maximum deviation from the measured atmosphere.

Simulations for only one atmosphere are available and the data must be normalized to account for the variation in the atmosphere over the course of the data sample. To determine the effect of atmospheric variations over the course of a year, the data was split into one-month blocks. Assuming the cosmic-ray flux and composition are constant, any statistically significant variations in the event distribution in reconstructed energy, primary likelihood ratio, and muon bundle energy loss between months are caused by the seasonal changes in the atmosphere. Using June 2009 as the reference, the event distribution of each month was shifted in each dimension to compensate for the seasonal changes in the atmosphere. The shift was determined using a likelihood maximization on the binned event distributions, and resulted in a best fit offset for each month in reconstructed energy, primary likelihood ratio and muon bundle energy. The errors on the best-fit were determined as in Section 5.2.

The best-fit offset for IceTop reconstructed energy was fit separately as the energy output of ShowerLLH is not continuous. ShowerLLH returns the energy for the most likely energy bin for both a proton- and iron-hypothesis. The mean of these two values are used as the reconstructed energy, but since the two values are typically the same, the energy distribution is peaked at the values of the energy bins. The peaked energy distribution leads to a stepped likelihood, which is incompatible with automatic optimization algorithms. To estimate the best energy offset, a brute-force scan of offsets was used. The reference distribution was projected into a 1-dimensional distribution in IceTop reconstructed energy. The test distribution was shifted and binned in the same manner, and a Poisson likelihood was calculated in a range of constant effective area. The monthly results in Figure 7.3 are consistent with 0 ($\chi^2/dof = 8.05/13 = 0.62$, $p = 0.84$) and can be compared to the monthly average surface atmospheric pressure (Figure 7.4). The intrinsic uncertainty in the energy offset is demonstrated by the error bars on month 0, which is approximately equal to the quantized energy output of ShowerLLH: 3 decades of energy divided into 80 bins: $3.0/80 = 0.0375$. The offset and atmospheric pressure have an insignificant negative correlation, $r = -0.32$ ($p = 0.29$) as shown in Figure 7.5.

Using the best-fit IceTop energy offset, the best-fit offsets for the other two parameters were found using a joint Poisson likelihood optimization using the same bins as in Section 6.5. The best fit is found by using an affine invariant Markov chain Monte Carlo package called *emcee* [108]. The best-fit offset for the IceTop likelihood ratio (Figure 7.6) is consistent with 0 ($\chi^2/dof = 19.09/13 = 1.47$, $p = 0.12$). The best-fit offset for MuEx energy (Figure 7.7) shows a strong seasonal variation and is overall inconsistent with 0 ($\chi^2/dof = 39.89/12 = 3.07$, $p = 1.4 \times 10^{-4}$). The live time for May 2009 (month -1) is considerably shorter than the other months, leading to the very large statistical uncertainty in the best-fit values. June 2009 is




Figure 7.3: The best fit offset for the IceTop reconstructed primary energy by month. The reference month was June 2009 (month 0), and the error bars correspond to the width of the likelihood distribution around the best fit corresponding to a $1\text{-}\sigma$ interval.




Figure 7.4: Monthly average surface pressure at the South Pole as measured by the AIRS instrument.




Figure 7.5: The measured surface atmospheric pressure and the best fit IceTop energy offset. No significant correlation was found.

included in the fit as a fit quality check. An additional check of the offset optimization was performed using odd and even runs from a single month, resulting in offsets consistent with 0 for all parameters.

While the seasonal change in the muon bundle energy is clear and well established [109], a seasonal change in the primary likelihood ratio is not as clear. Multiple years of data could help support or refute the existence of a seasonal effect, as the change should be roughly periodic, and follow the slight yearly variations in the change in seasons.




Figure 7.6: The best fit IceTop primary likelihood ratio offset by month. The reference month was June 2009 (month 0), and the error bars correspond to the width of the likelihood distribution around the best fit corresponding to a $1\text{-}\sigma$ interval. The distribution of offsets is statistically consistent with 0.




Figure 7.7: The best fit In-Ice MuEx muon bundle energy loss offset by month. The reference month was June 2009 (month 0), and the error bars correspond to the width of the likelihood distribution around the best fit corresponding to a $1\text{-}\sigma$ interval. The offset is statistically inconsistent with 0 and a seasonal dependence is clearly evident.

7.2 Results

Using the monthly offsets described in Section 7.1 to correct for seasonal variations, the full IceTop/IceCube-59 dataset, with a live time of 2.87×10^6 s (or 0.91 years), was fit using the regularization optimization described in Section 6.11.1. The optimization procedure did not result in a clear choice for the regularization parameter, as shown in Figure 7.8, but choosing the weakest regularization supported by the scan results in a choice of $\lambda = 10^{-11.0}$. The resulting spectrum is plotted in Figure 7.9 and weighted by $E^{2.7}$ in Figure 7.10 and the mean logarithm of the mass number is compared to the H4a model [34] in Figure 7.11.

Figure 7.12a shows the relative fit residual (the difference in the number of events in the data and the fit model divided by the number of events in the data) as a function of reconstructed energy. The best fit result has 10% fewer events than the data sample in each bin across a bulk of the entire energy range, which means the simulation-based model is a poor representation of the data. The fit residual pile-up at the highest energies is due to the limited energy reconstruction range of ShowerLLH. The residual around $\log(E/\text{GeV}) = 7.4$ corresponds to a change in the fit composition, and may be an artifact caused by the regularization method. The residuals do not show which parameter, or combination of parameters, is poorly fitting, but the consistent residual implies a systematic mismatch with the simulation across the energy range.

The two-dimensional representation of the fit residuals (Figure 7.12b) shows that the lighter components are systematically under-fit, while the heavier components are systematically over-fit across the energy range. Since the fit is performed in three dimensions, this could be caused by an offset in the model for one or all of the dimensions.




Figure 7.8: The result of the regularization parameter scan as described in Section 6.11.1. The choice of regularization parameter shown here at $\log(\lambda) = -11.0$ is the weakest regularization while still resulting in the least variance in the result.




Figure 7.9: The cosmic-ray spectrum from a 5-component regularized fit. The individual element results are slightly offset from the bin center for clarity.




Figure 7.10: The cosmic-ray spectrum as shown in Figure 7.9 except weighted by $E^{2.7}$ to enhance details in the spectrum.




Figure 7.11: The mean logarithm of the mass number for all particles from the fit shown in Figure 7.9. The overall trend is toward heavier mass primaries at higher energies. The H4a model by Gaisser [34] is included for comparison.

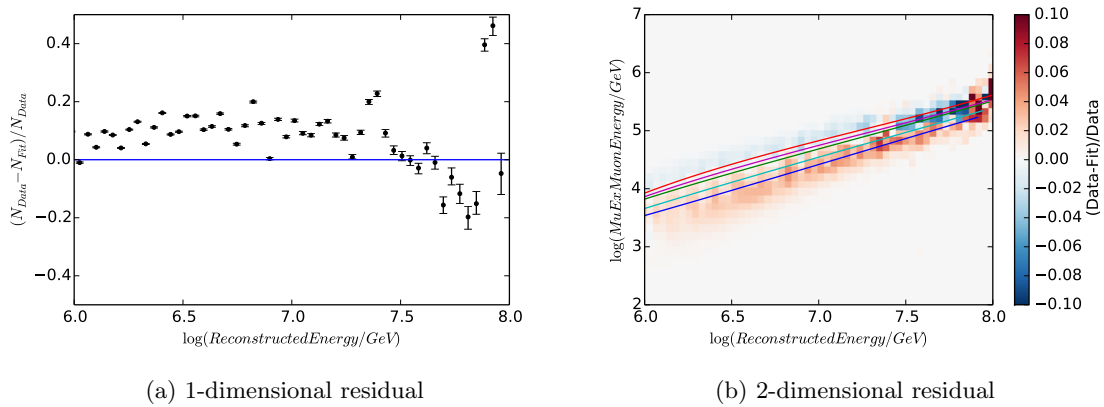


Figure 7.12: The relative fit residuals shown projected onto the IceTop reconstructed-energy axis in Figure 7.12a and into the muon-energy/IceTop reconstructed-energy plane in Figure 7.12b. The range plotted exceeds the range plotted in Figure 7.9 to show the possible influence from outside energy bins. A blue line at 0 is shown to guide the eye in Figure 7.12a.

7.3 Systematic Uncertainties

7.3.1 Empirical Data-Simulation Mismatch

The consistent residual as shown in Figure 7.12 indicates that a constant systematic shift in the simulation may be responsible for the mismatch between the data and simulation. To test this hypothesis, a Markov-chain Monte Carlo fit was used to find a linear shift in the measured quantities that would result in the minimal relative fit residual. For this fit, a hypothesis offset vector $\vec{R} = [x, y]$ is chosen, where x is the offset in muon energy loss and y is the offset in IceTop primary likelihood. Using data from June 2009, the offset is applied to the data before being histogrammed. This histogram is fit using the unregularized procedure outlined in Section 6.8.1. The value to be maximized by the fit \mathcal{L} is the negative relative residual in each bin having 10 or more events, which can be written as

$$\mathcal{L} = - \sum_i \begin{cases} \frac{\|n_{model} - n_{sim}\|}{n_{model}} & n_{model} \geq 10 \\ 0 & n_{model} < 10, \end{cases} \quad (7.1)$$

where the sum is over bin number i , n_{sim} is the number of events in the fit simulation and n_{model} is the number of events in the shifted data. The likelihood is only evaluated for bins having enough events so the standard error on the mean is small enough to make a meaningful contribution to the likelihood.

After applying the offset, the regularization strength scan becomes much more strongly peaked, as can be seen in Figure 7.15a. This implies that the tension in the fit caused by a mismatch between the data and simulation has been reduced. Figure 7.15 shows the results of the fit after applying the offset. Although the composition still trends towards heavier primaries at the highest energy, the relative ratios of primaries is very different. In this result, the heaviest composition around $\log(E/\text{GeV}) = 7.7$ is made up almost exclusively of oxygen, while in the original result, a similar $\langle \ln(A) \rangle$ was the result of a mix of helium and iron.

The maximum \mathcal{L}_{best} was found at $\vec{R}_{best} = [0.040, -0.006]$, with $\mathcal{L}_{best} = -130$. The no-offset case, $\vec{R}_{[0,0]} = [0, 0]$ is highly disfavored with $\mathcal{L}_{[0,0]} = -369$. The interpolated 2-dimensional scan (Figure 7.13) shows that the likelihood space does not strongly favor an offset in muon energy loss, while it strongly favors an offset in the IceTop primary likelihood ratio. The value of the most likely offset in primary likelihood ratio is more than 3 times the greatest variation seen in the monthly offset search in Figure 7.6. This would suggest that the offset cannot be described by the most extreme changes in the atmosphere over the course of a year. While the total flux is fit better by this shifted model, the physical interpretation is unclear.

7.3.2 Hadronic Interaction Models

Hadronic interaction models seek to accurately describe the products of cosmic-ray primary and air-shower secondary particles interacting in the atmosphere. These models determine how the air shower


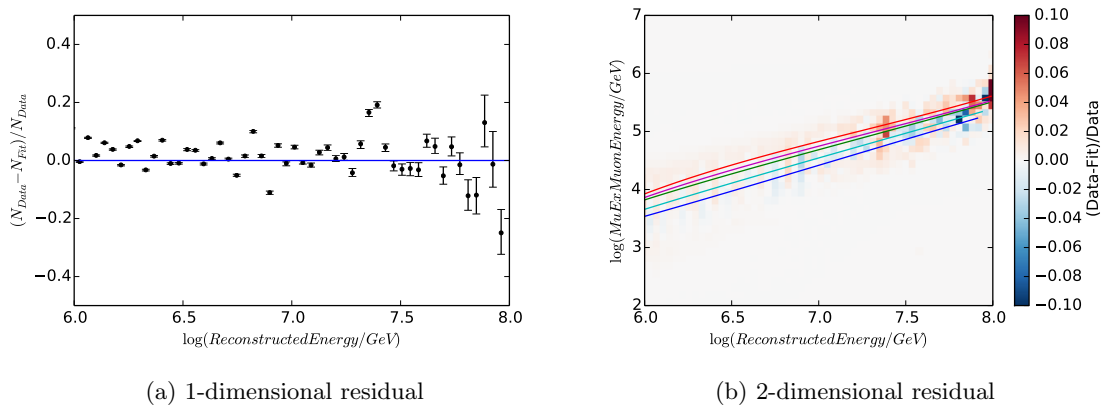


Figure 7.13: The likelihood space for offset in MuEx muon bundle energy loss and IceTop primary likelihood ratio. A green dot is at the point of the most likely offset. See Section 7.3.1 for a full description.



(a) 1-dimensional residual

(b) 2-dimensional residual

Figure 7.14: The relative fit residuals as shown in Figure 7.12 after applying the offset as described in Section 7.3.1.

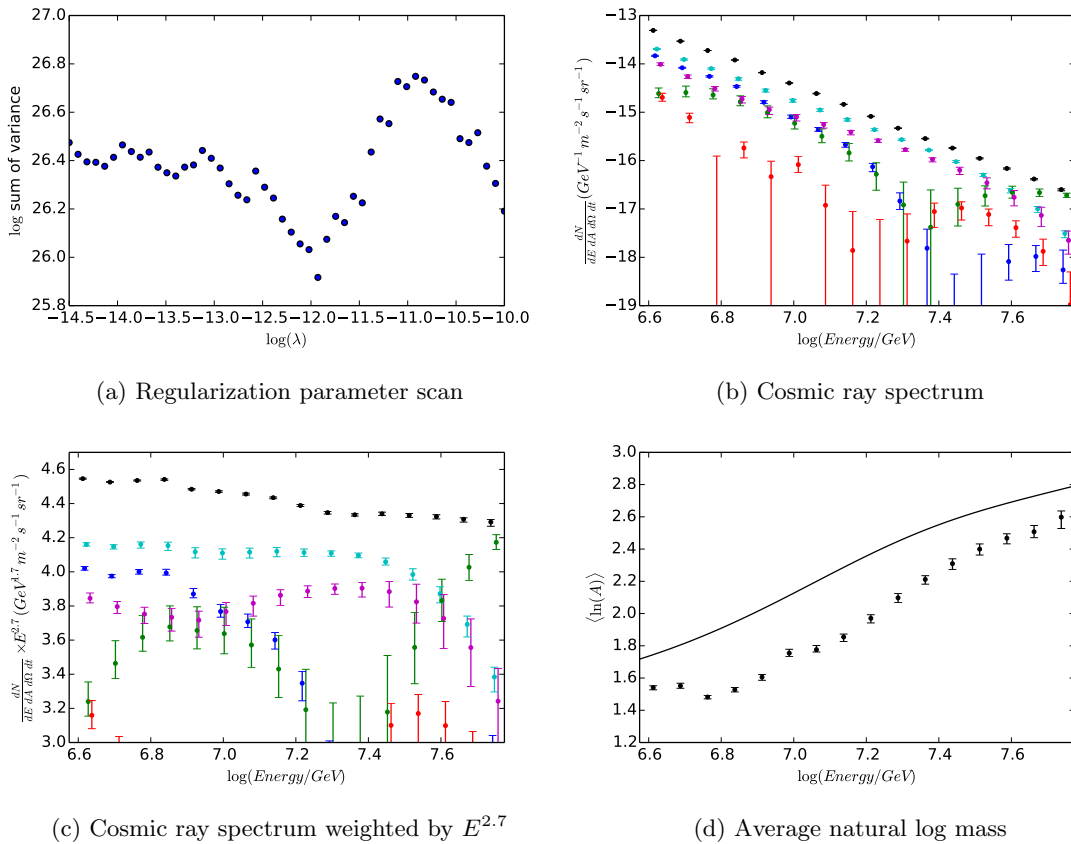


Figure 7.15: The fit results after applying the best fit offsets found in this section. The best fit regularization parameter after applying the offset is $\log(\lambda) = -11.96$.

develops in the atmosphere and how the total energy of the air shower is distributed between the hadronic, electromagnetic, and muonic components by the time it reaches the ground. These models are used at energies beyond those accessible by terrestrial accelerators and are largely extrapolations of lower energy measurements guided by cosmic-ray measurements and intuition. How the observations and assumptions are weighted and implemented leads to differences in the predictions of each hadronic interaction model. In this section, the effect of using the QGSJET-II [110] and EPOS [111] interaction models will be tested.

Simulating extensive air showers is very computationally intensive, regenerating the entire simulation with multiple interaction models is thus prohibitive. Such datasets are also unnecessary, as the model parameters vary slowly across orders of magnitude in primary energy. A targeted set of simulated showers were created to characterize the relative changes in the distribution of reconstructed parameters for each tested hadronic interaction model.

For this test, a *typical* shower geometry was chosen to be simulated multiple times for a set of primary energies and mass. The *typical* shower geometry was chosen by selecting the event with the lowest sum of IceTop energy rms and moun bundle energy loss rms from the set of iron simulation passing all data selection cuts with a logarithm of the true energy in GeV between 6.9 and 7.0. The selected event passed the cuts for 5 of 100 resamples and had a zenith angle of 3.79° and an azimuth angle of 323.2° . The shower was sampled with the core at $x = -44.0$ m, $y = 229.0$ m, which is well inside the detector geometry.

The resulting event distributions in reconstructed energy, muon bundle energy loss and primary likelihood ratio were calculated as in Section 6.7.2. SIBYLL was used to create the original distributions and the differences in the mean, variance and covariance as compared to SIBYLL were calculated. These differences were interpolated across energy and mass groups, and applied to the smoothed values (Section 6.7.3) to create new binned tables for fitting.

These tables are first used to demonstrate the sensitivity of the method to a mismatch in hadronic interaction models. By simulating an H4a spectrum with one hadronic interaction model, and fitting it with another, the resulting fit residuals can be seen systematically deviating from zero, unlike those in Chapter 6 where the same model was used. The tables are then used to fit the data. The ability of each model to describe the data can then be evaluated by comparing the fit residuals to each other and to the tests on the simulated spectrum.

7.3.2.1 QGSJET-II

One of the recent high-energy hadronic interaction models simulated is QGSJET-II [110]. Figure 7.16 shows the SIBYLL-based unfolding of an event distribution of QGSJET-II simulation weighted by the H4a model. The individual spectra are not well fit, with the iron component consistently over-represented and other components under-represented, as shown in Figure 7.16c. The fit residual, Figure 7.16e, shows a

small systematic slope across the energy range, which is indicative of a mismatch between the QGSJET-II and SIBYLL models.

The QGSJET-II simulation fit to the complete dataset results in an ambiguous result of the regularization optimization algorithm, preferring regularization so strong to overwhelm the data (Figure 7.17a). The 1-dimensional (Figure 7.17e) and 2-dimensional fit residuals (Figure 7.17f) show clear structure and an overall failure to fit the data distribution, especially at the highest energies. Viewed in this context, the fit plots show the results of regularization in the case of a generally failing fit (e.g. Figure 7.17c and Figure 7.17d.).

7.3.2.2 EPOS

The EPOS LHC (v3400) [111] high-energy hadronic interaction model, taking into account recent LHC and RHIC measurements, was also used to create simulation tables. Figure 7.18 shows the SIBYLL-based unfolding of an event distribution of EPOS simulation weighted by the H4a model. The regularization optimization scan in Figure 7.18a shows no clear minimum, suggesting that the fit variances are relatively high across the whole range. Choosing a regularization of $\log(\lambda) = -12$ to be consistent with the QGSJET-II test, the results shown in Figures 7.18b to 7.18d are far from the input model, and show large statistical uncertainty. Figure 7.18e shows a large fit residual across almost the entire energy range, and Figure 7.18f shows that one region of the parameter space is being consistently under-fit while another region is consistently over-fit. This effect is due to the designed inflexibility of the fit. It is unable to move the event distributions and only scale their relative magnitudes. Since the EPOS and SIBYLL simulations are offset in the parameter space, this results in a large fit residual. Similar fit residuals were seen in the initial SIBYLL fit of the data, as in Figure 7.12, and it is reasonable to conclude that SIBYLL may not be an accurate model of high-energy hadronic interactions.

Unfolding the data with EPOS simulation is more successful. The regularization scan (Figure 7.19a) was successful, and returned an optimal regularization of $\log(\lambda) = -12.0$, consistent with the value used for the fit using the ad-hoc offset. This value was used for the results as shown in the remaining plots of Figure 7.19. The residual plots (Figure 7.19e and Figure 7.19f) show small residuals, similar to those for the ad-hoc offsets for the SIBYLL simulation in Figure 7.15. However, there is no agreement between EPOS and the offset-SIBYLL in the details of the elemental composition (Figure 7.19c and Figure 7.15c), or in the measurement of the mean logarithm of the mass number.

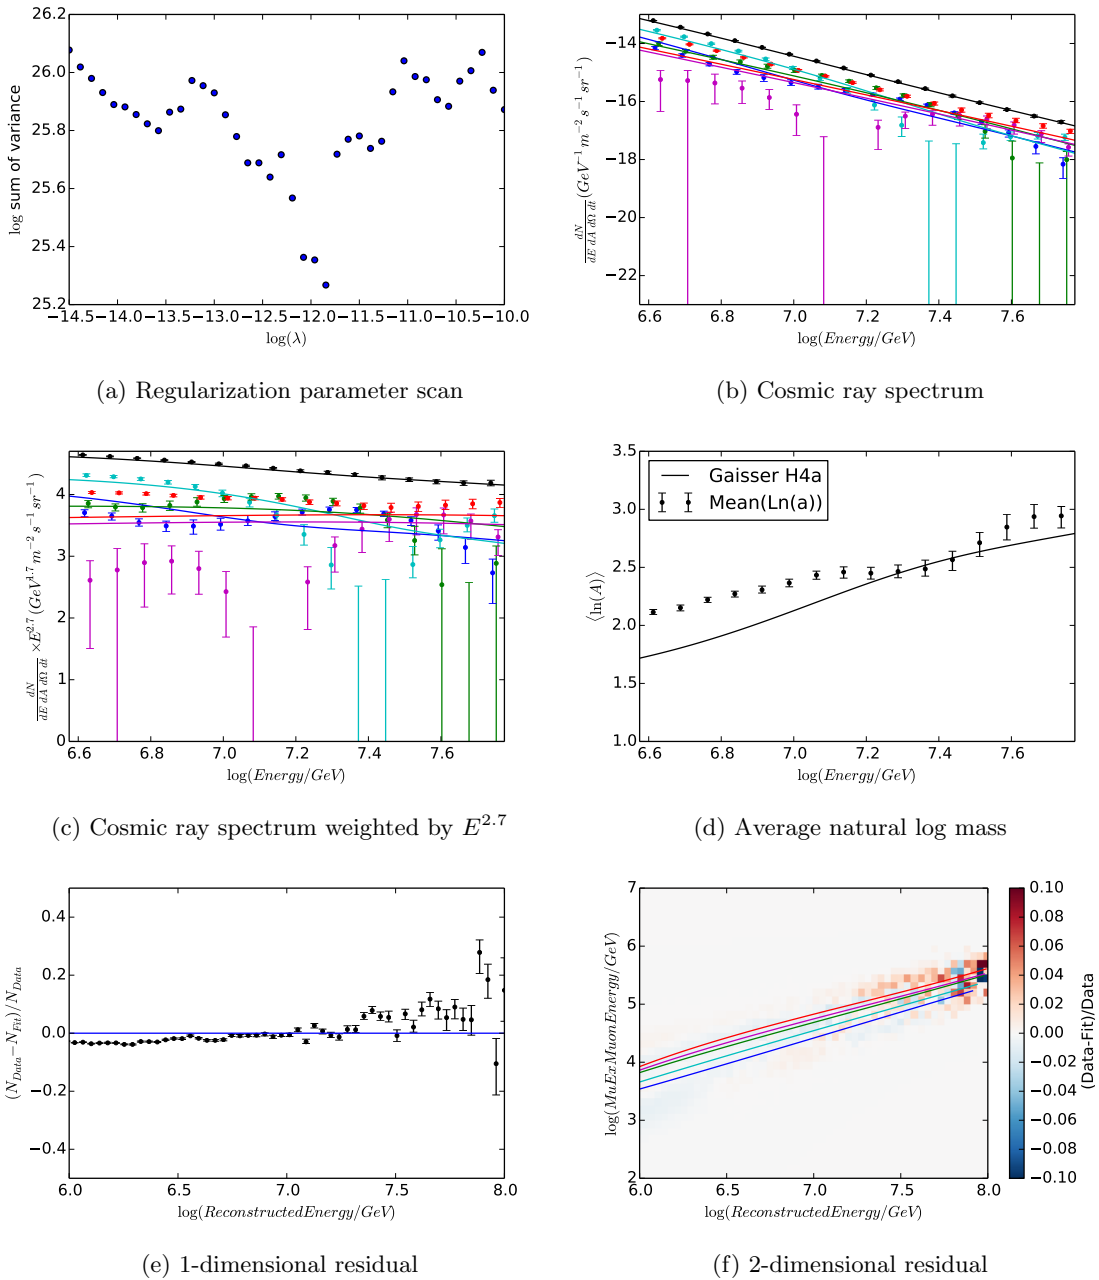


Figure 7.16: The fit results for an QGSJET-II derived simulated H4a flux unfolded with the standard SIBYLL-based simulation. The optimized regularization is $\log(\lambda) = -12$, which is a factor of 10 weaker than for the SIBYLL-derived simulated flux. This value was used for Figures 7.16b to 7.16f. The intermediate elements, making up a small percentage of the flux have very large relative errors. Even with the weaker regularization, Figure 7.16e shows a systematic offset in the fit residuals, negative at lower energies and higher at the highest energies.

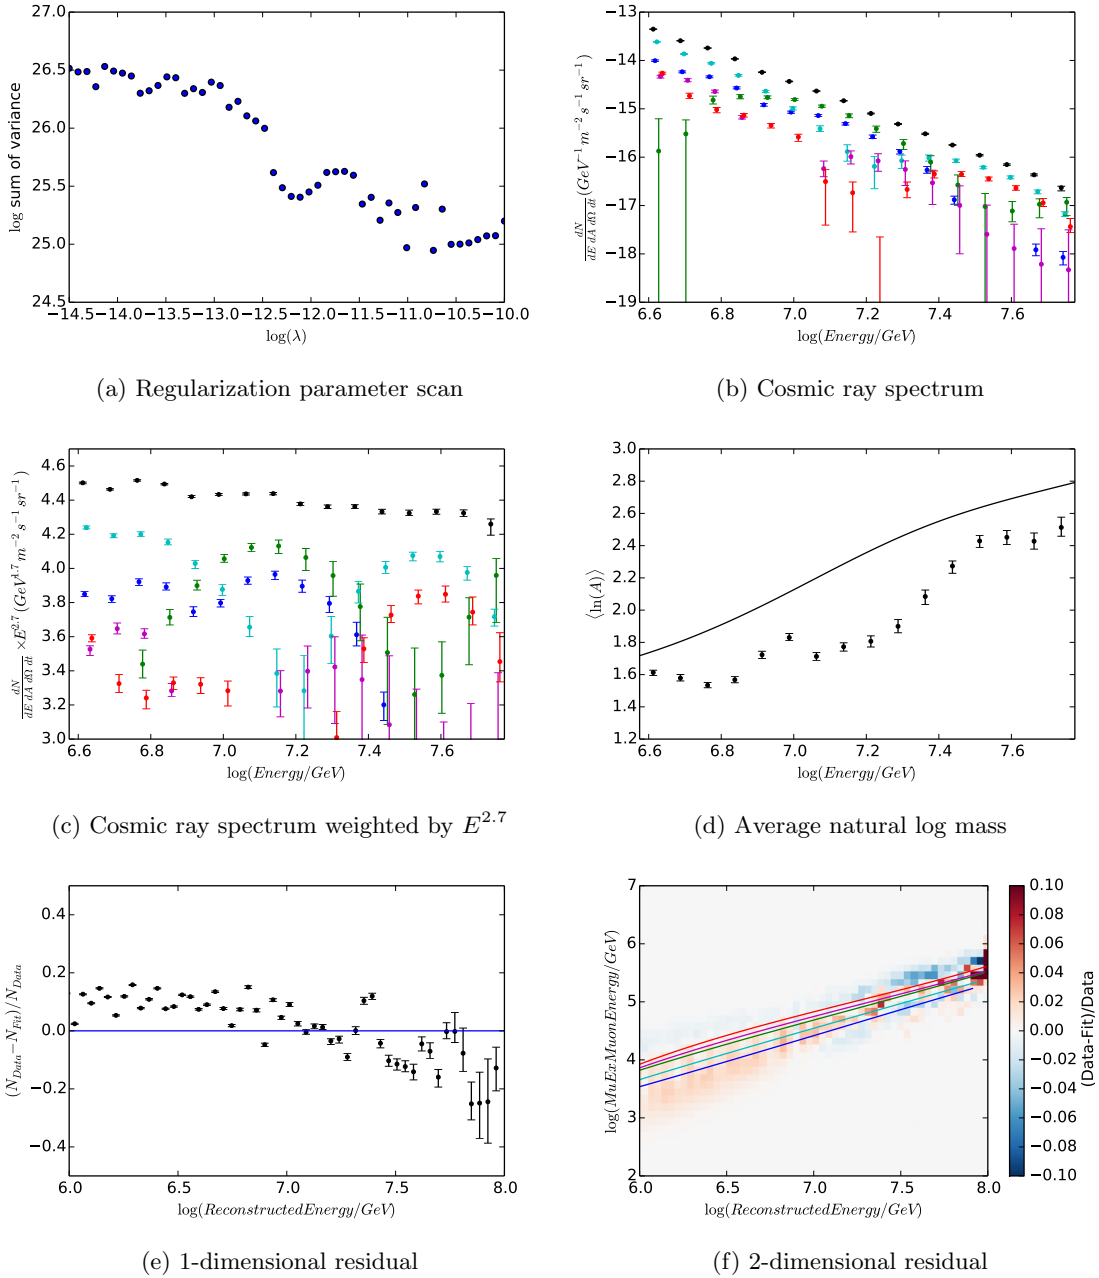


Figure 7.17: The fit results for simulations using the QGSJET-II high-energy interaction model. The regularization of $\log(\lambda) = -12.5$ is maintained for comparison to the original result in Section 7.2.

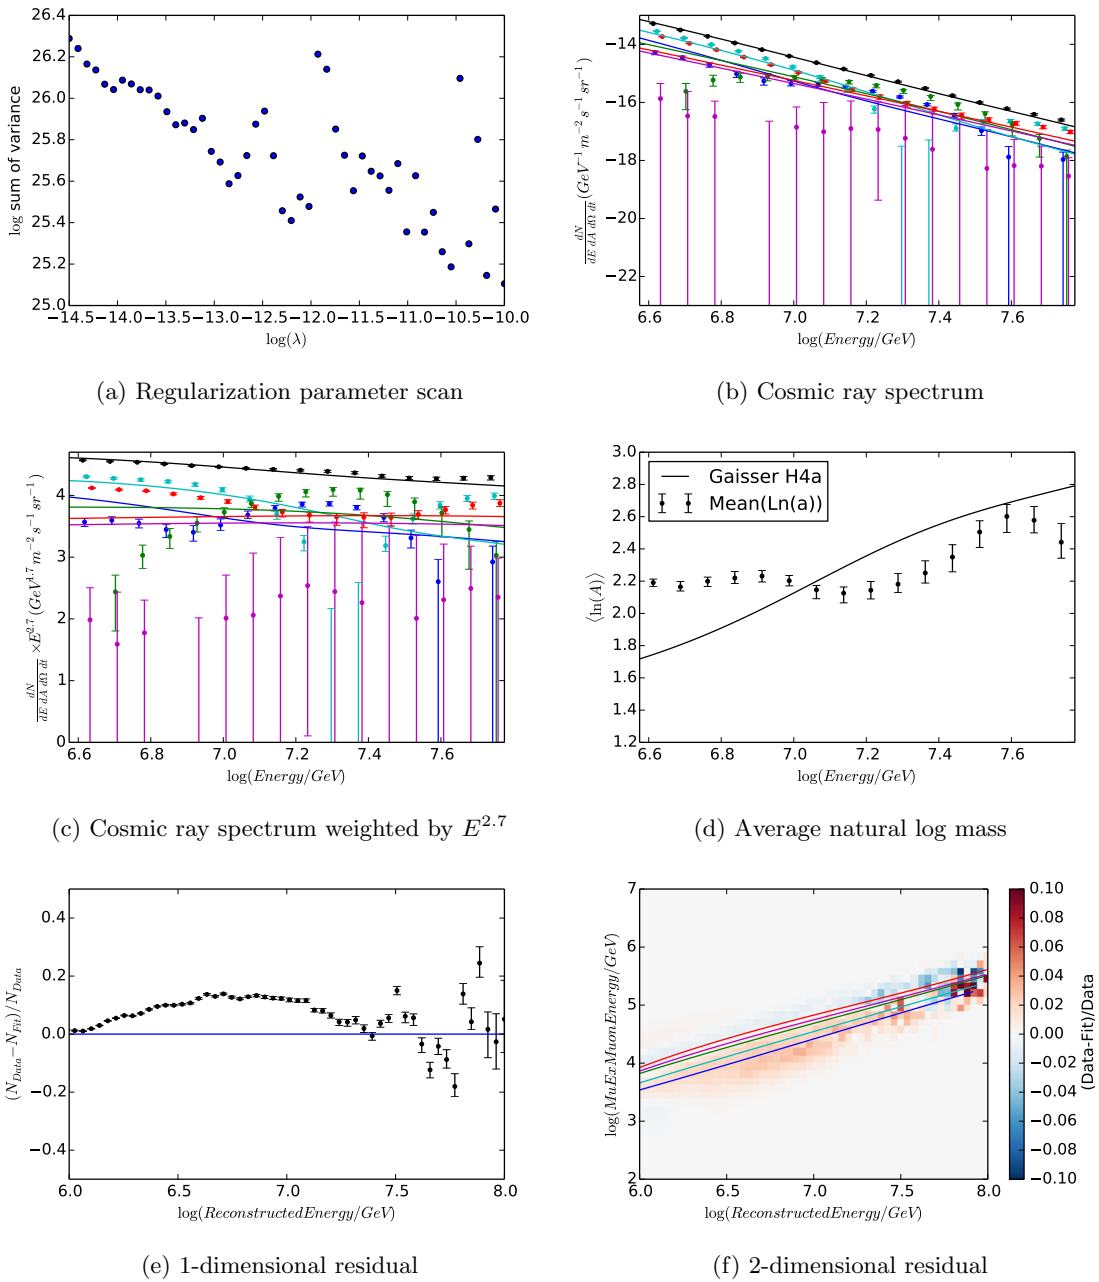


Figure 7.18: The fit results for an EPOS-derived simulated H4a flux unfolded with the standard SIBYLL-based simulation. The regularization optimization scan in Figure 7.18a shows no clear minimum, so a regularization of $\log(\lambda) = -12.0$ was used to be consistent with the same test for QGSJET-II in Figure 7.16. Across the entire range, the unfolded spectrum bears little resemblance to the input spectrum, which is reflected in the residuals shown in Figures 7.18e and 7.18f.

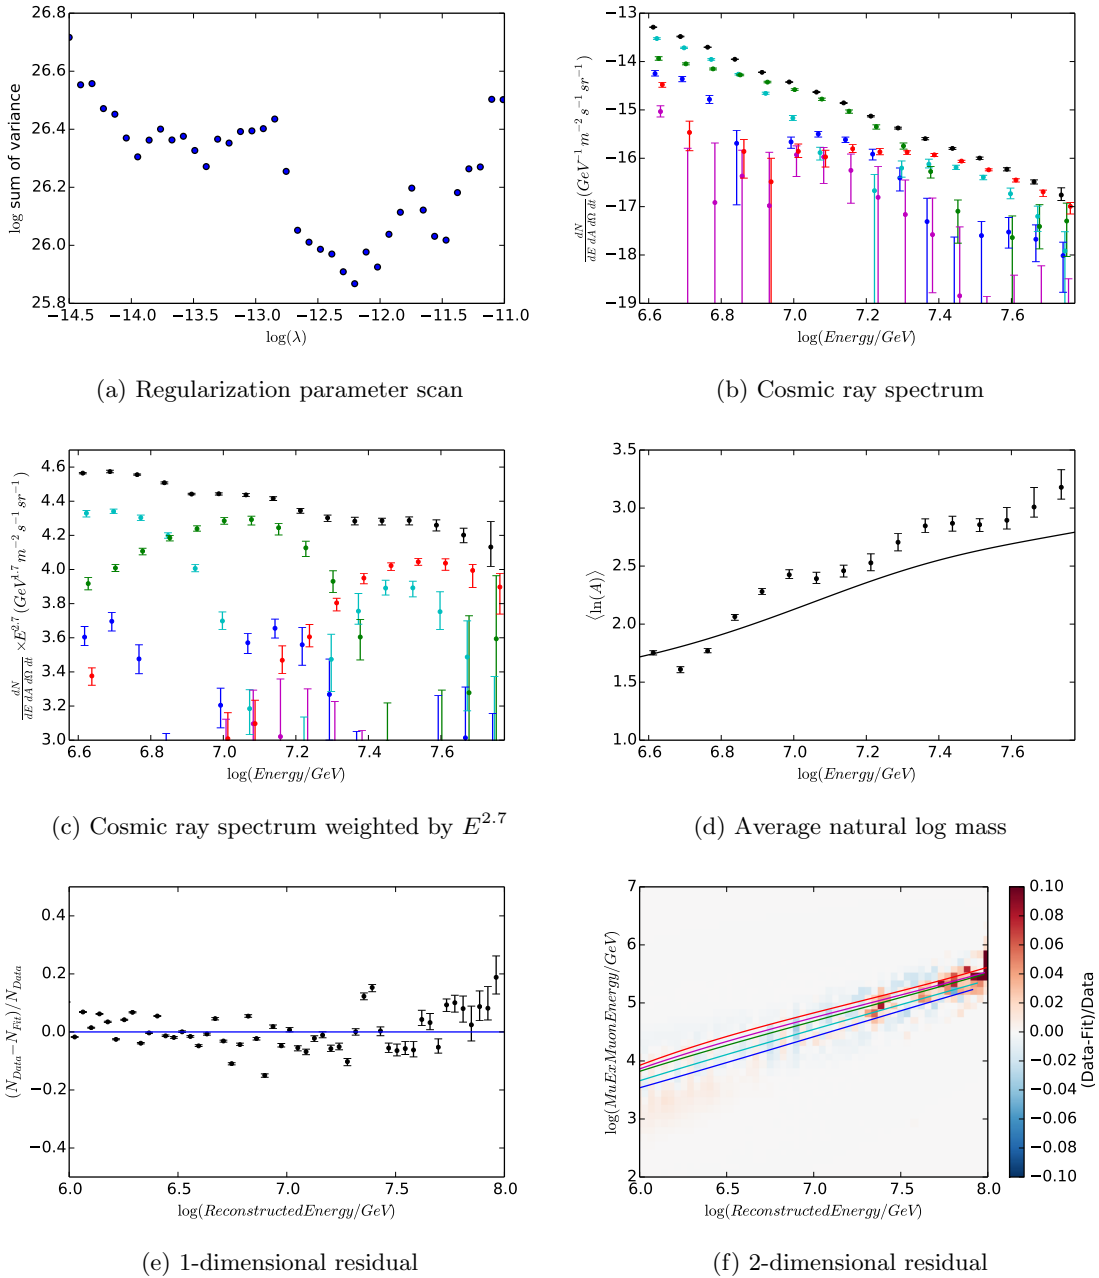


Figure 7.19: The fit results for simulations using the EPOS high-energy interaction model. The optimized regularization is consistent with $\log(\lambda) = -12.5$, the original result in Section 7.2, which is the value used for Figures 7.19b to 7.19f.

7.4 Summary

Figure 7.20 shows the all-particle flux result obtained using the standard SIBYLL simulation, SIBYLL with an ad-hoc shift to the data, and the systematics studies done with QGSJET-II and EPOS. Figure 7.21 shows the same as a weighted flux to enhance differences. The results agree within 10%, with the fits diverging above 25 PeV, which corresponds to a change in composition as seen in the mean natural logarithm of the mass number, Figure 7.22. The $\langle \ln(A) \rangle$ shows a trend for each primary from lighter composition at low energies to heavier composition at higher energies, with each fit exhibiting approximately the same slope.

The individual spectra do not agree between each five-component fit. As the hadronic interaction models do not agree in general, this is to be expected and confirmed by the simulation tests using different interaction models. Since all components are fit simultaneously, any systematic residuals in the two-dimensional residual plots are a sign of the over- or under-fitting of at least one component, which will effect all other components.




Figure 7.20: The all-particle spectrum results for the fit with SIBYLL (cyan), EPOS (green), QGSJET-II (blue) and SIBYLL with a best fit offset (red).




Figure 7.21: The weighted all-particle spectrum results for the fit with SIBYLL (cyan), EPOS (green), QGSJET-II (blue) and SIBYLL with a best fit offset (red).




Figure 7.22: The mean logarithm of the mass number results for the fit with SIBYLL (cyan), EPOS (green), QGSJET-II (blue) and SIBYLL with a best fit offset (red).

Chapter 8

Conclusions

“It’s more fun to arrive at a conclusion than to justify it.”

Malcolm Forbes

A measurement of the cosmic-ray spectrum and composition at PeV energies using IceTop and IceCube detectors was presented. By measuring the high-energy muons created early in the cosmic-ray air-shower development using IceCube and the extensive air shower at its maximum using IceTop, the combined detectors provide a promising opportunity for indirectly measuring the mass composition of cosmic rays. The measurement uses a simultaneous fit to three parameters: the primary energy as determined by IceTop, the muon bundle energy loss as measured by IceCube and a primary likelihood ratio from the lateral distribution of particles as measured by IceTop. The method was verified on a simulated spectrum and composition, using these parameters to fit with distributions from simulations. The measurement of the cosmic-ray spectrum agrees roughly with previous measurements. The measured change in composition, from light to heavy is also in agreement with model predictions. The method was designed to easily identify disagreement between data and simulation, and the simulation was found to be a poor description of the data.

Reliable simulation must accurately describe all aspects of air-shower interactions and detector response. The cosmic-ray composition cannot currently be considered a known input quantity, so the only way to determine the fidelity of simulation to nature is to try all possible combinations of primary flux, and compare the distribution of observables. Using standard IceCube simulations employing the SIBYLL hadronic interaction model, a satisfactory fit was not obtained. The fit had a 10% residual across energy bins and a systematic offset visible in the 2-dimensional muon bundle energy loss versus reconstructed energy residual plot. The residuals were centered around 0 by adding an ad-hoc offset to the IceTop primary likelihood, but there is no clear physical interpretation of this offset. Unfolding with the QGSJET-II hadronic interaction model introduced a slope in the fit residuals, suggesting the model’s behavior above 10^7 GeV starts to deviate from the data in an energy dependent way. Unfolding with the EPOS hadronic interaction model resulted in

the smallest absolute residuals, similar to the ad-hoc offset, making it the best physically-motivated model to fit the data.

The method presented here demonstrates an ability to measure the chemical composition of a simulated spectrum within statistical errors when using the same hadronic-interaction model for the unfolding. When a different hadronic-interaction model is used to fit a simulated spectrum, the discrepancy is seen in the one- and two-dimensional fit residuals, providing a good diagnostic for fit quality and differences in air shower physics.

Figure 8.1 and Figure 8.2 compare the fit presented in this work to previous analyses performed with the IceCube detector. The blue band represents the range of the fits using SIBYLL, SIBYLL with the ad-hoc shift to the data, and the systematics studies done with QGSJET-II and EPOS.




Figure 8.1: The cosmic-ray flux as measured by this analysis is shown in a blue band, enclosing the extent of the error bars for each fit shown in Figure 7.20. The flux result from the IceTop-73 and the IceCube/IceTop-40 analysis are shown in green and blue, respectively.




Figure 8.2: The weighted cosmic-ray flux as measured by this analysis is shown in a blue band, enclosing the extent of the error bars for each fit shown in Figure 7.21. The flux result from the IceTop-73 and the IceCube/IceTop-40 analysis are shown in green and blue, respectively.




Figure 8.3: The mean log mass result from this analysis is shown as a blue band, enclosing the extent of the error bars for each fit shown in Figure 7.22. The IceCube/IceTop-40 result is shown in blue, and the H4a model prediction is shown as a black line.

8.1 Looking Forward

The data used in this analysis were collected with the pre-completion IceCube-59 detector. Since that time, four more years of data have been taken and the IceTop and IceCube arrays have been completed, greatly increasing the detector area. The addition of DeepCore and an in-fill IceTop station lower the energy threshold for both detectors. The use of this additional data would greatly increase the statistical power of this analysis at high energies and better study agreement with direct detection experiments at low energies.

Concurrent with the development of this analysis, ShowerLLH has seen significant improvements that were not available in the version used here. Additional simulations have increased the energy range to higher energies and the likelihood tables can now be fit and used without binning to reduce statistical fluctuations in the likelihood space and improve energy resolution.

Extending this analysis to higher energies requires simulation of higher energy air-showers, which is very computationally intensive. The method of smoothing event distributions from simulations showed the reconstructed quantities varied smoothly and slowly over wide energy ranges, suggesting that it could be used to interpolate between a sparser set of simulations, reducing the computational requirements. Simulations using atmosphere profiles from other seasons would also help confirm the validity of the seasonal correction method used in this analysis.

The EPOS hadronic-interaction model fit the data the best of the three models without the need for any ad-hoc corrections, suggesting it is currently the best model for cosmic-ray interactions in the atmosphere. However, the tables used for the test were derived by applying interpolated differences from EPOS event distributions to values fit to SIBYLL simulations. Only a full simulation using EPOS can provide a reliable evaluation of its relative performance for describing high-energy hadronic air shower interactions.

A more speculative but potentially powerful extension to this analysis would be to include non-coincident composition-sensitive measurements. This analysis performs a simultaneous fit to the number of events in each independent bin of a three-dimensional parameter space. Since each bin is independent, there is no requirement that the bins be from the same event distribution, only that they measure the same cosmic-ray spectrum. By adding other composition sensitive event distributions, such as the particle density versus zenith angle as in the IceTop-73 analysis, or muon bundle energy loss in IceCube versus slant depth and zenith angle, statistical uncertainties at high and low energies could be reduced. Adding inhomogeneous datasets could also reduce covariance between terms of the fit by breaking degeneracies that may be present in individual datasets.

LIST OF REFERENCES

- [1] G. W. Clark. *Biographical Memoirs*, volume 75, chapter "Bruno Benedetto Rossi", pages 310–341. The National Academies Press, 1998. URL: http://www.nap.edu/openbook.php?record_id=9649.
- [2] Georges Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys.Lett.*, B716:1–29, 2012. [arXiv:1207.7214](https://arxiv.org/abs/1207.7214), [doi:10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- [3] Serguei Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys.Lett.*, B716:30–61, 2012. [arXiv:1207.7235](https://arxiv.org/abs/1207.7235), [doi:10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021).
- [4] R. P. Crease. *The Second Creation: Makers of the Revolution in Twentieth-Century Physics*. Rutgers University Press, 1996.
- [5] V. F. Hess. Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten. *Physik. Zeitschr.*, 13:1084–1091, 1912.
- [6] B. Rossi. *Moments in the Life of a Scientist*. Cambridge University Press, 1990.
- [7] P. Auger, P. Ehrenfest, R. Maze, J. Daudin, Robley, and A. Fréon. Extensive Cosmic-Ray Showers. *Reviews of Modern Physics*, 11:288–291, July–October 1939.
- [8] M. Aguilar et al. The Alpha Magnetic Spectrometer (AMS) on the International Space Station: Part I - results from the test flight on the space shuttle. *Physics Reports*, 366:331–405, August 2002. [doi:10.1016/S0370-1573\(02\)00013-3](https://doi.org/10.1016/S0370-1573(02)00013-3).
- [9] J. Abraham et al. Measurement of the energy spectrum of cosmic rays above 10^{18} eV using the Pierre Auger Observatory. *Physics Letters B*, 685:239–246, March 2010. [arXiv:1002.1975](https://arxiv.org/abs/1002.1975), [doi:10.1016/j.physletb.2010.02.013](https://doi.org/10.1016/j.physletb.2010.02.013).
- [10] T. H. Johnson. The Azimuthal Asymmetry of the Cosmic Radiation. *Phys. Rev.*, 43:834–835, May 1933. URL: <http://link.aps.org/doi/10.1103/PhysRev.43.834>, [doi:10.1103/PhysRev.43.834](https://doi.org/10.1103/PhysRev.43.834).
- [11] W. C. Barber. East-West Asymmetry and Latitude Effect of Cosmic Rays at Altitudes up to 33,000 Feet. *Phys. Rev.*, 75:590–599, Feb 1949. URL: <http://link.aps.org/doi/10.1103/PhysRev.75.590>, [doi:10.1103/PhysRev.75.590](https://doi.org/10.1103/PhysRev.75.590).
- [12] C. E. McIlwain. Coordinates for Mapping the Distribution of Magnetically Trapped Particles. *J. Geophys. Res.*, 66:3681–3691, November 1961. [doi:10.1029/JZ066i011p03681](https://doi.org/10.1029/JZ066i011p03681).
- [13] M. A. Duvernois. Galactic Cosmic Ray Source Elemental Composition. In S. S. Holt and G. Sonneborn, editors, *Cosmic Abundances*, volume 99 of *Astronomical Society of the Pacific Conference Series*, page 385, 1996.

- [14] J. J. Beatty and S. Westerhoff. The Highest-Energy Cosmic Rays. *Annual Review of Nuclear and Particle Science*, 59(1):319–345, 2009. doi:10.1146/annurev.nucl.58.110707.171154.
- [15] J. Beringer et al. Review of Particle Physics. *Physical Review D*, 86(1):010001, 2012.
- [16] R. Jansson and G. R. Farrar. The Galactic Magnetic Field. *Astrophysical Journal Letters*, 761:L11, December 2012. arXiv:1210.7820, doi:10.1088/2041-8205/761/1/L11.
- [17] A. A. Penzias and R. W. Wilson. A Measurement of Excess Antenna Temperature at 4080 Mc/s. *Astrophysical Journal*, 142:419–421, July 1965. doi:10.1086/148307.
- [18] K. Greisen. End to the Cosmic-Ray Spectrum? *Physical Review Letters*, 16:748–750, April 1966. doi:10.1103/PhysRevLett.16.748.
- [19] G. T. Zatsepin and V. A. Kuz'min. Upper Limit of the Spectrum of Cosmic Rays. *Soviet Journal of Experimental and Theoretical Physics Letters*, 4:78, August 1966.
- [20] N. Globus, D. Allard, and E. Parizot. Propagation of high-energy cosmic rays in extragalactic turbulent magnetic fields: resulting energy spectrum and composition. *Astronomy and Astrophysics*, 479:97–110, February 2008. arXiv:0709.1541, doi:10.1051/0004-6361:20078653.
- [21] G. V. Kulikov and G.B. Khristiansen. On the Size Spectrum of Extensive Air Showers. *Sov. Phys. JETP*, 35(8):441–444, 1959.
- [22] A. W. Strong, I. V. Moskalenko, and V. S. Ptuskin. Cosmic-Ray Propagation and Interactions in the Galaxy. *Annual Review of Nuclear and Particle Science*, 57:285–327, November 2007. arXiv:astro-ph/0701517, doi:10.1146/annurev.nucl.57.090506.123011.
- [23] J. R. Hörandel. Models of the knee in the energy spectrum of cosmic rays. *Astroparticle Physics*, 21(3):241–265, 2004.
- [24] T. Stanev, P. L. Biermann, and T. K. Gaisser. Cosmic Rays IV. The spectrum and chemical composition above 10^4 GeV. *Arxiv preprint astro-ph/9303006*, 1993.
- [25] S. E. Thorsett, R. A. Benjamin, W. F. Brisken, A. Golden, and W. M. Goss. Pulsar PSR b0656+14, the Monogem Ring, and the Origin of the ‘Knee’ in the Primary Cosmic-Ray Spectrum. *Astrophys.J.*, 592:L71–L74, 2003. arXiv:astro-ph/0306462, doi:10.1086/377682.
- [26] A. D. Erlykin and A. W. Wolfendale. Cosmic rays and the Monogem supernova remnant. *Astroparticle Physics*, 22:47–63, October 2004. arXiv:astro-ph/0404530, doi:10.1016/j.astropartphys.2004.05.009.
- [27] B. Peters. Primary Cosmic Radiation and Extensive Air Showers. *Il Nuovo Cimento*, 22(4):800–819, 1961. URL: <http://dx.doi.org/10.1007/BF02783106>, doi:10.1007/BF02783106.
- [28] S. P. Swordy. Expectations for Cosmic Ray Composition Changes in the Region 10^{14} to 10^{16} eV. In *International Cosmic Ray Conference*, volume 2 of *International Cosmic Ray Conference*, page 697, 1995.
- [29] F. W. Stecker. Photodisintegration of Ultrahigh-Energy Cosmic Rays by the Universal Radiation Field. *Phys. Rev.*, 180:1264–1266, Apr 1969. URL: <http://link.aps.org/doi/10.1103/PhysRev.180.1264>, doi:10.1103/PhysRev.180.1264.

- [30] Y. Tsunesada and for the Telescope Array Collaboration. Highlights from Telescope Array. *ArXiv e-prints*, November 2011. arXiv:1111.2507.
- [31] R.-Y. Liu, A. M. Taylor, M. Lemoine, X.-Y. Wang, and E. Waxman. Constraints on the Source of Ultra-high-energy Cosmic Rays Using Anisotropy versus Chemical Composition. *The Astrophysics Journal*, 776:88, October 2013. arXiv:1308.5699, doi:10.1088/0004-637X/776/2/88.
- [32] J. R. Hörandel. On the knee in the energy spectrum of cosmic rays. *Astroparticle Physics*, 19:193–220, May 2003. arXiv:astro-ph/0210453, doi:10.1016/S0927-6505(02)00198-6.
- [33] T. K. Gaisser. Spectrum of cosmic-ray nucleons, kaon production, and the atmospheric muon charge ratio. *Astroparticle Physics*, 35:801–806, July 2012. arXiv:1111.6675, doi:10.1016/j.astropartphys.2012.02.010.
- [34] T. K. Gaisser. Cosmic ray spectrum and composition \geq 100 TeV. Technical report, IceCube, 2011.
- [35] T. K. Gaisser, T. Stanev, and S. Tilav. Cosmic ray energy spectrum from measurements of air showers. *Frontiers of Physics*, 8:748–758, December 2013. arXiv:1303.3565, doi:10.1007/s11467-013-0319-7.
- [36] E. Fermi. On the Origin of the Cosmic Radiation. *Physical Review*, 75(8):1169, 1949.
- [37] S. Longair. *High Energy Astrophysics: Stars, the galaxy and the interstellar medium*. Number v. 2 in High Energy Astrophysics. Cambridge University Press, 1994. URL: <http://books.google.com/books?id=x1TMSMpzD2UC>.
- [38] A. R. Bell. The acceleration of cosmic rays in shock fronts. I. *Monthly Notices of the Royal Astronomical Society*, 182:147–156, January 1978.
- [39] W. Baade and F. Zwicky. On Super-novae. *Proceedings of the National Academy of Science*, 20:254–259, May 1934. doi:10.1073/pnas.20.5.254.
- [40] W. Baade and F. Zwicky. Cosmic Rays from Super-novae. *Proceedings of the National Academy of Science*, 20:259–263, May 1934. doi:10.1073/pnas.20.5.259.
- [41] R. Diehl, H. Halloin, K. Kretschmer, G. G. Lichti, V. Schönfelder, A. W. Strong, A. von Kienlin, W. Wang, P. Jean, J. Knödseder, J.-P. Roques, G. Weidenspointner, S. Schanne, D. H. Hartmann, C. Winkler, and C. Wunderer. Radioactive ^{26}Al from massive stars in the Galaxy. *Nature*, 439:45–47, January 2006. arXiv:astro-ph/0601015, doi:10.1038/nature04364.
- [42] V. L. Ginzburg and S. I. Syrovatskii. *The Origin of Cosmic Rays*. Pergamon Press and Macmillan Company, 1964.
- [43] T. Stanev. *High Energy Cosmic Rays*. Springer Praxis Books. Springer, 2010. URL: <http://books.google.com/books?id=1y9YEHafBMC>.
- [44] M. Ackermann et al. Detection of the Characteristic Pion-Decay Signature in Supernova Remnants. *Science*, 339:807–811, February 2013. arXiv:1302.3307, doi:10.1126/science.1231160.
- [45] U. Hwang and J. M. Laming. A Chandra X-Ray Survey of Ejecta in the Cassiopeia A Supernova Remnant. *Astrophysical Journal*, 746:130, February 2012. arXiv:1111.7316, doi:10.1088/0004-637X/746/2/130.

- [46] A. M. Hillas. The Origin of Ultra-High-Energy Cosmic Rays. *Annual Review of Astronomy and Astrophysics*, 22:425–444, 1984. doi:10.1146/annurev.aa.22.090184.002233.
- [47] F. M. Rieger, V. Bosch-Ramon, and P. Duffy. Fermi acceleration in astrophysical jets. *Astrophysics and Space Science*, 309:119–125, jun 2007.
- [48] D. Guetta, D. Hooper, J. Alvarez-Muñiz, F. Halzen, and E. Reuveni. Neutrinos from individual gamma-ray bursts in the BATSE catalog. *Astroparticle Physics*, 20:429–455, January 2004. arXiv:astro-ph/0302524, doi:10.1016/S0927-6505(03)00211-1.
- [49] R. Abbasi et al. An Absence of Neutrinos Associated with Cosmic-Ray Acceleration in Gamma-Ray Bursts. *Nature*, 484:351–354, 2012.
- [50] J. R. Hörandel. The composition of cosmic rays at the knee. In J. F. Ormes, editor, *American Institute of Physics Conference Series*, volume 1516 of *American Institute of Physics Conference Series*, pages 185–194, February 2013. arXiv:1212.0739, doi:10.1063/1.4792566.
- [51] W.D. Apel, J. C. Arteaga-Velázquez, et al. KASCADE-Grande measurements of energy spectra for elemental groups of cosmic rays. *Astroparticle Physics*, 47(0):54 – 66, 2013. URL: <http://www.sciencedirect.com/science/article/pii/S0927650513000996>, doi:<http://dx.doi.org/10.1016/j.astropartphys.2013.06.004>.
- [52] R. Abbasi et al. Cosmic Ray Composition and Energy Spectrum from 1–30 PeV Using the 40-String Configuration of IceTop and IceCube. *Astroparticle Physics*, 2012.
- [53] M. G. Aartsen et al. Measurement of the cosmic ray energy spectrum with IceTop-73. *Physical Review D*, 88(4):042004, August 2013. doi:10.1103/PhysRevD.88.042004.
- [54] T. K. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, 1991.
- [55] R. Ulrich, J. Blümer, R. Engel, F. Schüssler, and M. Unger. On the measurement of the proton-air cross section using longitudinal shower profiles. *Nuclear Physics B Proceedings Supplements*, 175:121–124, January 2008. arXiv:astro-ph/0612205, doi:10.1016/j.nuclphysbps.2007.10.019.
- [56] R. Ulrich, J. Blümer, R. Engel, F. Schüssler, and M. Unger. On the measurement of the proton-air cross section using air shower data. *New Journal of Physics*, 11(6):065018, 2009. URL: <http://stacks.iop.org/1367-2630/11/i=6/a=065018>.
- [57] H. A. Bethe and L. C. Maximon. Theory of Bremsstrahlung and Pair Production. I. Differential Cross Section. *Physical Review*, 93:768–784, February 1954. doi:10.1103/PhysRev.93.768.
- [58] K. Greisen. Cosmic Ray Showers. *Annual Review of Nuclear and Particle Science*, 10:63–108, 1960. doi:10.1146/annurev.ns.10.120160.000431.
- [59] G. Molière. Theorie der Streuung schneller geladener Teilchen I. Einzelstreuung am abgeschirmten Coulomb-Feld. *Zeitschrift Naturforschung Teil A*, 2:133, 1947.
- [60] J. Holder. VERITAS: Status and Highlights. *International Cosmic Ray Conference*, 12:137, 2011. arXiv:1111.1225, doi:10.7529/ICRC2011/V12/H11.
- [61] J. Cortina, F. Goebel, T. Schweizer, and for the MAGIC Collaboration. Technical Performance of the MAGIC Telescopes. *ArXiv e-prints*, July 2009. arXiv:0907.1211.

- [62] J. A. Hinton and For the HESS Collaboration. The status of the HESS project. *New Astronomy Reviews*, 48:331–337, April 2004. [arXiv:astro-ph/0403052](#), [doi:10.1016/j.newar.2003.12.004](#).
- [63] M. Ackermann et al. Optical properties of deep glacial ice at the South Pole. *Journal of Geophysical Research*, 111(D13):D13203, 2006.
- [64] D. Heck, J. Knapp, JN Capdevielle, G. Schatz, T. Thouw, et al. *CORSIKA: A Monte Carlo code to simulate extensive air showers*, volume 6019. FZKA, 1998.
- [65] I. E. Tamm and I. M. Frank. Coherent Radiation of Fast Electron in Medium. *Doklady Akademii Nauk SSSR*, 14:107–113, 1937.
- [66] R. Atkins et al. TeV Gamma-Ray Survey of the Northern Hemisphere Sky Using the Milagro Observatory. *Astrophysical Journal*, 608:680–685, June 2004. [doi:10.1086/420880](#).
- [67] I. Allekotte et al. The surface detector system of the Pierre Auger Observatory. *Nuclear Instruments and Methods in Physics Research A*, 586:409–420, March 2008. [doi:10.1016/j.nima.2007.12.016](#).
- [68] A. U. Abeysekara and other. Sensitivity of the high altitude water Cherenkov detector to sources of multi-TeV gamma rays. *Astroparticle Physics*, 50:26–32, December 2013. [arXiv:1306.5800](#), [doi:10.1016/j.astropartphys.2013.08.002](#).
- [69] G. Aielli et al. The Status of the ARGO Experiment at YBJ. *Nuclear Physics B Proceedings Supplements*, 166:96–102, April 2007. [doi:10.1016/j.nuclphysbps.2006.12.074](#).
- [70] J. C. Arteaga-Velázquez et al. The KASCADE-Grande experiment: measurements of the all-particle energy spectrum of cosmic rays. *ArXiv e-prints*, September 2010. [arXiv:1009.4716](#).
- [71] R. Abbasi et al. The design and performance of IceCube DeepCore. *Astroparticle Physics*, 2012.
- [72] IceCube Collaboration. Coordinate System [online]. URL: https://wiki.icecube.wisc.edu/index.php/Coordinate_system [cited April 2014].
- [73] IceCube Collaboration. I3Direction [online]. URL: <https://wiki.icecube.wisc.edu/index.php/I3Direction> [cited April 2014].
- [74] M. G. Aartsen et al. Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector. *Science*, 342, November 2013. [arXiv:1311.5238](#), [doi:10.1126/science.1242856](#).
- [75] R. Abbasi, Y. Abdou, T. Abu-Zayyad, M. Ackermann, J. Adams, J. A. Aguilar, M. Ahlers, D. Altmann, K. Andeen, J. Auffenberg, and et al. Search for ultrahigh-energy tau neutrinos with IceCube. *Physical Review D*, 86(2):022005, July 2012. [arXiv:1202.4564](#), [doi:10.1103/PhysRevD.86.022005](#).
- [76] Donald E Groom, Nikolai V Mokhov, and Sergei I Striganov. Muon stopping power and range tables 10 MeV–100 TeV. *Atomic Data and Nuclear Data Tables*, 78(2):183–356, 2001.
- [77] R. Abbasi et al. The IceCube data acquisition system: Signal capture, digitization, and timestamping. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 601(3):294–316, 2009.
- [78] 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.

- [79] D. Chirkin. Study of South Pole ice transparency with IceCube flashers. *International Cosmic Ray Conference*, 4:161, 2011.
- [80] R. Abbasi et al. IceTop: The surface component of IceCube. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 2012.
- [81] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, et al. GEANT4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.
- [82] D. Chirkin and W. Rhode. Propagating leptons through matter with Muon Monte Carlo (MMC). *arXiv preprint hep-ph/0407075*, 2004.
- [83] J. Lundberg, P. Miočinović, K. Woschnagg, T. Burgess, J. Adams, S. Hundertmark, P. Desiati, and P. Niessen. Light tracking through ice and water—Scattering and absorption in heterogeneous media with PHOTONICS. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 581(3):619–631, 2007.
- [84] M. G. Aartsen et al. The IceProd Framework: Distributed Data Processing for the IceCube Neutrino Observatory. *ArXiv e-prints*, November 2013. [arXiv:1311.5904](https://arxiv.org/abs/1311.5904).
- [85] R. S. Fletcher, T. K. Gaisser, P. Lipari, and T. Stanev. SIBYLL: An event generator for simulation of high energy cosmic ray cascades. *Physical Review D*, 50(9):5710, 1994.
- [86] A. Ferrari, P. R. Sala, A. Fasso, and J. Ranft. Fluka. *CERN-library in: http://fluka.web.cern.ch/fluka*, 2005.
- [87] G. Battistoni, F. Cerutti, A. Fassò, A. Ferrari, S. Muraro, J. Ranft, S. Roesler, and P. R. Sala. The FLUKA code: description and benchmarking. In M. Albrow and R. Raja, editors, *Hadronic Shower Simulation Workshop*, volume 896 of *American Institute of Physics Conference Series*, pages 31–49, March 2007. doi:10.1063/1.2720455.
- [88] D. Heck and T. Pierog. *Extensive Air Shower Simulation with CORSIKA: A User’s Guide*. Institut für Kernphysik Forschungszentrum Karlsruhe. URL: <http://www-ik.fzk.de/~corsika/usersguide/usersguide.pdf> [cited March 2013].
- [89] A. E. Hedin. Extension of the MSIS thermosphere model into the middle and lower atmosphere. *Journal of Geophysical Research: Space Physics (1978–2012)*, 96(A2):1159–1172, 1991.
- [90] D. Chirkin. dCORSIKA update. Technical report, IceCube, 2002. URL: <http://icecube.wisc.edu/~dima/work/BKP/DCS/REPORT8/paper.pdf>.
- [91] F. Kislak, T. Melzig, A. Van Overloop, and T. Waldenmaier. IceTop tank simulation using g4-tankresponse. Technical report, IceCube, 2012.
- [92] J. Ahrens et al. Muon track reconstruction and data selection techniques in AMANDA. *Nuclear Instruments and Methods in Physics Research A*, 524:169–194, May 2004. [arXiv:astro-ph/0407044](https://arxiv.org/abs/astro-ph/0407044), doi:10.1016/j.nima.2004.01.065.
- [93] G. Japaridze and M. Ribordy. Realistic arrival time distribution from an isotropic light source. *ArXiv Astrophysics e-prints*, June 2005. [arXiv:astro-ph/0506136](https://arxiv.org/abs/astro-ph/0506136).

- [94] D. Pandel. Bestimmung von Wasser- und Detektorparametern und Rekonstruktion von Myonen bis 100 TeV mit dem Baikal-Neutrino teleskop NT-72. Diploma thesis, Humboldt-Universität zu Berlin, 1996.
- [95] D. Chirkin. New implementation of convoluted MPE pandel function. Technical Report icecube/201101006, IceCube, 2011. URL: http://internal.icecube.wisc.edu/reports/data/icecube/2011/01/006/icecube_201101006_v1.pdf.
- [96] D. Chirkin. Photon propagation approximations and energy reconstruction. Technical Report icecube/201102007, IceCube, 2011. URL: http://internal.icecube.wisc.edu/reports/data/icecube/2011/02/007/icecube_201102007_v1.pdf.
- [97] S. Grullon, D. Boersma, and G. Hill. Photonics-based Log-Likelihood Reconstruction in IceCube. Technical Report icecube/200807001-v3, IceCube, 2008. URL: http://internal.icecube.wisc.edu/reports/data/icecube/2008/07/001/icecube_200807001_v3.pdf.
- [98] M. Ribordy and G. Japaridze. Reconstruction PDF in an Inhomogeneous Medium. Technical Report 20050301, Amanda, 2007.
- [99] A. A. Watson. Extensive Air Shower and Ultra High Energy Cosmic Rays. Lectures given at a Summer School in Mexico, 2002. URL: <http://www.ast.leeds.ac.uk/Auger/augerthesis/mexlects3.pdf>.
- [100] M. G. Aartsen et al. Observation of Cosmic-Ray Anisotropy with the IceTop Air Shower Array. *Astrophysical Journal*, 765:55, March 2013. arXiv:1210.5278, doi:10.1088/0004-637X/765/1/55.
- [101] IceCube Collaboration. 2009 DAQ Configuration Parameters. Technical report, IceCube, 2009. URL: <https://docushare.icecube.wisc.edu/dsweb/Get/Document-47863/2009DAQconfiguration-parameters.pdf>.
- [102] Tilo Waldenmaier for the Cosmic Ray Working Group. IC59 Cosmic Ray Filter Request for 2009. Technical report, IceCube, 2009. URL: https://docushare.icecube.wisc.edu/dsweb/Get/Document-48187/2009_TFT_IceTop_FilterRequest.pdf.
- [103] A. Genz. Numerical computation of multivariate normal probabilities. *Journal of computational and graphical statistics*, 1(2):141–149, 1992.
- [104] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL: <http://www.scipy.org/>.
- [105] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems*. Prentice-Hall, 1974.
- [106] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge University Press, New York, NY, USA, 3 edition, 2007.
- [107] E. T. Olsen. *AIRS/AMSU/HSB Version 6 Data Release User Guide*. NASA Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, 1.2 April 2013.
- [108] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman. emcee: The MCMC Hammer. *Publications of the Astronomical Society of the Pacific*, 125:306–312, March 2013. arXiv:1202.3665, doi:10.1086/670067.
- [109] P. Desiati. Seasonal Variations of High Energy Cosmic Ray Muons Observed by the IceCube Observatory as a Probe of Kaon/Pion Ratio. *International Cosmic Ray Conference*, 1:78, 2011. doi:10.7529/ICRC2011/V01/0662.

- [110] S. Ostapchenko. Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model. *Phys.Rev.*, D83:014018, 2011. [arXiv:1010.1869](#), [doi:10.1103/PhysRevD.83.014018](#).
- [111] T. Pierog, I. Karpenko, J. M. Katzy, E. Yatsenko, and K. Werner. EPOS LHC: test of collective hadronization with LHC data. *ArXiv e-prints*, June 2013. [arXiv:1306.0121](#).