

**Systematic Studies of Cosmic-Ray Anisotropy and Energy Spectrum with
IceCube and IceTop**

By
Frank McNally

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

Doctor of Philosophy
(Physics)

at the
UNIVERSITY OF WISCONSIN–MADISON
2015

Date of final oral examination: October 9th, 2015

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

Paolo Desiati, Associate Scientist, Astronomy

Francis Halzen, Professor, Physics

Justin Vandenbroucke, Assistant Professor, Physics

Stefan Westerhoff, Professor, Physics

Ellen Zweibel, Professor, Astronomy

Abstract

Anisotropy in the cosmic-ray arrival direction distribution has been well documented over a large energy range, but its origin remains largely a mystery. In the TeV to PeV energy range, the galactic magnetic field thoroughly scatters cosmic rays, but anisotropy at the part-per-mille level and smaller persists, potentially carrying information about nearby cosmic-ray accelerators and the galactic magnetic field.

The IceCube Neutrino Observatory was the first detector to observe anisotropy at these energies in the Southern sky. This work uses over 300 billion cosmic-ray induced muon events, collected between May 2009 and May 2015 from both the in-ice component of IceCube as well as the surface component, IceTop. The observed global anisotropy features large regions of relative excess and deficit, with amplitudes on the order of 10^{-3} . While a decomposition of the arrival direction distribution into spherical harmonics shows that most of the power is contained in the low-multipole ($\ell \leq 4$) moments, higher-multipole components are found to be statistically significant down to an angular scale of less than 10° , approaching the angular resolution of the detector. Above 100 TeV, a change in the topology of the arrival direction distribution is observed, and the anisotropy is characterized by a wide relative deficit whose amplitude increases with primary energy up to at least 5 PeV, the highest energies currently accessible to IceCube with sufficient event statistics. No time dependence of the large- and small-scale structures is observed in the six-year period covered by this analysis within statistical and systematic uncertainties.

Analysis of the energy spectrum and composition in the PeV energy range as a function of sky position is performed with IceTop data over a five-year period using a likelihood-based

reconstruction. Both the energy spectrum and the composition distribution are found to be consistent with a single source population over declination bands.

This work represents an early attempt at understanding the anisotropy through the study of the spectrum and composition. The high-statistics data set reveals more details on the properties of the anisotropy, potentially able to shed light on the various physical processes responsible for the complex angular structure and energy evolution.

ACKNOWLEDGMENTS

First and foremost, thank you always to my wife, Kathryn, for making me laugh every day.

To my advisor, Stefan Westerhoff, thank you for your patience and calm. I so enjoyed working with you as an undergraduate that when I considered graduate school, there was really only one option in my mind. The WIPAC community has always been wonderful. Thank you to Jon Eisch and Claire Pettersen for getting me started, to Kim Krieger for making me feel at home, to Segev BenZvi for showing me how to code, and to Marcos Santander for being the shining example of what graduate students can hope to be. Thank you to Paolo Desiati for his contagious enthusiasm and gesticulation. Thanks to Dan Fiorino for being my favorite officemate, and Zig and Ian for making sure we didn't get too much done. Thank you to Aimee Lefkow, Renee Lefkow, Ann Austin, and the whole UW Physics department staff, as well as Frances Culwell and the WIPAC staff. Thank you to Jakob, Chris, and Laura, whose experience saved me many times, to Kyle and Zach, for making my time at IceCube that much more fun, and to Josh for teaching me about baseball.

Thank you to my family: your constant, unwavering support allowed me to put off being a real adult for as long as possible. I love you all. Thank you to Greg Devine, whose passion inspired me and many others to pursue this field. A huge thank you to the Carleton physics community, for showing me what I want to do, and to Cindy Blaha in particular for getting me started. My eternal thanks go to Bryan, Eric, and Anna for making Madison my home, and to Annie, Robbie, and Mary for bringing a little home with them. And of course, thank you to the Sour Patch Adults, for playing sports.

It is not until you write something like this that you realize just how many people you have to thank. Thank you to those who I've missed, for understanding. To everyone above and everyone else, thank you so much for making me who I am today.

DISCARD THIS PAGE

TABLE OF CONTENTS

	Page
LIST OF TABLES	vii
LIST OF FIGURES	viii
1 Cosmic Rays	1
1.1 Introduction	1
1.2 Energy Spectrum	2
1.3 Composition	6
1.4 Acceleration Mechanisms	9
1.5 Candidate Sources	15
1.5.1 Galactic Sources	17
1.5.2 Extragalactic Sources	18
1.6 Propagation	19
1.7 Anisotropy	23
1.7.1 Large-Scale Anisotropy	26
1.7.2 Small-Scale Anisotropy	27
1.8 Overview	28
2 Detection	31
2.1 Air Showers	31
2.2 Detection Principle	34
2.2.1 Cherenkov Radiation	36
2.3 IceCube	37
2.3.1 Digital Optical Modules	40
2.3.2 Ice Properties	44
2.3.3 Data Storage and Transfer	45
2.4 IceTop	47
3 Event Reconstruction	51
3.1 IceCube	51

	Page
3.1.1 Angular Reconstruction	51
3.1.2 Energy Estimation	53
3.2 IceTop	56
3.2.1 Angular Reconstruction	56
3.2.2 Likelihood-based Reconstruction (ShowerLLH)	57
4 Analysis Method	70
4.1 Coordinate System	70
4.1.1 Local Coordinates	70
4.1.2 Equatorial Coordinates	71
4.1.3 Coordinate Transformation	73
4.1.4 Solar Coordinate Frame	76
4.1.5 Non-Physical Coordinate Frames	76
4.2 Background Estimation	77
4.2.1 Time-Scrambling	77
4.2.2 Dipole and Quadrupole Subtraction	79
4.3 Mapmaking Procedure	80
4.3.1 HEALPix	80
4.3.2 Relative Intensity and Significance	81
4.4 Power Spectrum	82
4.5 Bayes Factor	83
5 Results	86
5.1 Angular Scales	86
5.1.1 Large-Scale Structure	86
5.1.2 Small-Scale Structure	87
5.1.3 Power Spectrum	93
5.1.4 Systematic Checks	95
5.2 Time Dependence	98
5.2.1 Time Dependence of Large-Scale Anisotropy	98
5.2.2 Time Dependence of Small-Scale Anisotropy	100
5.3 Energy Dependence	102
5.4 Composition and Energy Study	107
6 Conclusions	115
LIST OF REFERENCES	118

DISCARD THIS PAGE

LIST OF TABLES

Table	Page
2.1 IceTop filters and prescales	50
3.1 IceTop simulation	60
3.2 IceTop low-energy simulation	60
3.3 IceTop quality cuts	62
5.1 IceCube data set	87
5.2 Small-scale structure details	91
5.3 Time dependence of large-scale anisotropy	100
5.4 Chemical composition for energy bins	114
5.5 Events by most-likely composition	114

DISCARD THIS PAGE

LIST OF FIGURES

Figure	Page
1.1 Cosmic-ray energy spectrum	3
1.2 Intersecting power law example	5
1.3 Spectrum for various elemental components	7
1.4 Chemical abundances of galactic cosmic rays	8
1.5 Mean logarithmic mass	10
1.6 First-order Fermi acceleration	13
1.7 Hillas Plot	16
1.8 Magnetic field lines in M51	20
1.9 Large-scale anisotropy observed by Tibet-AS γ and IceCube	24
1.10 Small-scale anisotropy as observed by Milagro	28
1.11 Small-scale anisotropy as observed by IceCube	29
2.1 Cosmic-ray air shower development	32
2.2 Muon stopping power	35
2.3 Cherenkov radiation	38
2.4 IceCube schematic	39
2.5 IceCube deployment seasons	41
2.6 DOM schematic	42
2.7 Event rate in IceCube	44

Figure	Page
2.8 Dust absorption in IceCube	46
2.9 IceTop station	48
2.10 Simulated shower development	48
2.11 Vertical equivalent muon in IceTop	50
3.1 Muon event in IceCube	54
3.2 Energy estimation in IceCube	55
3.3 IceTop effective area without cuts	62
3.4 IceTop effective area with quality cuts	63
3.5 Zenith correction	63
3.6 Comparison of reconstructed and true energies in IceTop	65
3.7 Energy and core resolution in IceTop	66
3.8 Composition separation in IceTop	68
3.9 Bayes factor composition test	69
4.1 Local coordinate system	71
4.2 Local acceptance of IceCube	72
4.3 Equatorial coordinate system	74
4.4 Equatorial acceptance of IceCube	75
4.5 Pixelization in HEALPix	80
4.6 Power spectrum for IC59	84
5.1 Relative intensity and significance for large-scale anisotropy	88
5.2 Relative intensity and significance for small-scale anisotropy	90
5.3 Small-scale structure comparison	92

Figure	Page
5.4 Relative intensity and significance for four-hour time scrambling	94
5.5 Power spectrum for IceCube	96
5.6 Relative intensity and significance of solar dipole	96
5.7 One-dimensional projection in right ascension of large-scale structure	97
5.8 Time dependence of large-scale anisotropy	99
5.9 Time dependence of small-scale anisotropy	101
5.10 Power spectrum for residual maps	103
5.11 Relative intensity as a function of energy (Mollweide projection)	104
5.12 Relative intensity as a function of energy (polar projection)	105
5.13 Dipole amplitude and phase as a function of energy	107
5.14 Bayes factor for energy spectrum comparisons	109
5.15 Energy spectrum of deficit compared to the rest of the sky	110
5.16 Relative intensity and significance as a function of most-likely composition . . .	111
5.17 Relative intensity and significance as a function of most-likely composition with minimum energy cut	112
5.18 Bayes factor for composition distribution comparison	113

Chapter 1

Cosmic Rays

1.1 Introduction

In 2012 we celebrated the 100th anniversary of Victor Hess' balloon flight, the event credited with the discovery of cosmic rays [1]. Armed with two electroscopes, he ascended 5300 m in the air during a solar eclipse in an attempt to determine the source of background radiation. It was the belief at the time that the radiation was coming from the Earth, but Hess found that radiation levels increased as he climbed. With the sun eclipsed, this result led him to the conclusion that these ionizing particles must be coming from space.

A century later, many pieces of the puzzle have been filled in. We know that cosmic rays are predominantly made up of ionized nuclei, not the gamma rays that gave them their name. We have measured a spectrum that spans over 12 orders in magnitude in energy, using detectors underwater, underground, on the surface, in the sky, and in space. Yet the question that Victor Hess strove out to discover 100 years ago remains largely unanswered: where do they come from? It is with this question in mind that the following work was done; small steps taken towards understanding.

In this chapter, we provide an introduction to the basic properties of cosmic rays: the shape and features of their energy spectrum, composition, acceleration methods, potential sources, propagation through the galaxy, and observed anisotropy in arrival direction distribution.

1.2 Energy Spectrum

One of the most striking aspects of cosmic rays is their energy spectrum, shown in Fig. 1.1. The flux of cosmic rays as a function of energy, as traditionally taught, follows a nearly unbroken power law of the form $AE^{-\gamma}$, with a spectral index $\gamma \approx 2.7$. Remarkably, it does so over 10 orders of magnitude in energy, and 30 orders of magnitude in flux. This power law behavior begins around 10^{10}eV , where the influence from the solar wind fades. At 10^{11}eV , the Earth sees a flux of cosmic rays at about 1 particle per square meter per second. The spectrum continues with a spectral index of $\gamma = 2.7$ until the first spectral break: the “knee” at about $3 \times 10^{15}\text{eV}$ [2]. At the knee the flux is roughly 1 particle per square meter per year, and begins dropping faster as the spectral index softens to $\gamma = 3.1$. This behavior continues until the second spectral break: the “ankle” at about 10^{19}eV [3]. At the ankle the flux is approximately 1 particle per square kilometer per year, and the spectral index hardens once again to $\gamma = 2.6$.

In 1991, the Fly’s Eye detector in Utah detected the highest-energy cosmic ray every observed, with an estimated energy of $3 \times 10^{20}\text{eV}$ [4]. This cosmic ray was about 40 million times more energetic than the highest energy achievable by the Large Hadron Collider (LHC), with a center-of-mass energy about 50 times as high as what the LHC can currently produce. Cosmic rays of this energy are incredibly rare, with a flux of approximately 1 particle per square kilometer per century.

The highest energy cosmic rays are subject in theory to a cutoff due to the Greisen-Zatsepin-Kuz’min (GZK) limit [6, 7]. The GZK limit describes the energy at which protons become able to interact with the Cosmic Microwave Background (CMB) via the Δ^+ resonance. Protons with energies higher than $\sim 6 \times 10^{19}\text{eV}$ then experience energy loss through one of the following interactions:

$$p + \gamma_{\text{CMB}} \rightarrow \Delta^+ \rightarrow n + \pi^+$$

$$p + \gamma_{\text{CMB}} \rightarrow \Delta^+ \rightarrow p + \pi^0 \ .$$




Figure 1.1 Cosmic-ray energy spectrum as measured with a variety of experiments [5]. The spectrum behaves as a power law, with spectral breaks visible at the “knee” (~ 3 PeV) and “ankle” (~ 4 EeV).

This process does not limit the maximum achievable energy of cosmic-ray accelerators, but instead limits the mean free path of protons above the GZK cutoff to a few tens of Mpc. A suppression of the cosmic-ray flux for energies above 4×10^{19} eV has been observed [8], and while it is consistent with the theory, it is also possible this cutoff is intrinsic to cosmic-ray accelerators.

The hardening of the spectrum around the ankle can be described as the flux from less populous (or more distant), higher-energy sources overtaking the decaying flux from more populous (or nearer), lower-energy sources as they reach the limits of their acceleration methods. This idea is reinforced by the energies of these high-energy cosmic rays. Cosmic rays are all subject to the Lorentz force, and thus are deflected by magnetic field lines. The radius of the circular path followed is known as the Larmor radius L , and can easily be derived from the centripetal force:

$$qv_{\perp}B = \frac{mv_{\perp}^2}{r} . \quad (1.1)$$

Rearranging and replacing $mv = p \approx E/c$ for relativistic particles gives us the following relation:

$$L \approx \frac{E/c}{ZeB} . \quad (1.2)$$

The Larmor radius for particles above the ankle in typical galactic magnetic field strengths of $\sim 2 \mu\text{G}$ [9] is much larger than the size of the galaxy, leading us to believe that cosmic rays above the ankle are extragalactic in origin.

Describing the softening of the spectrum at the knee is more of a challenge. Sources producing a hard spectrum ($\gamma = 2.7$) would never be overtaken by sources producing a softer spectrum ($\gamma = 3.1$). Figure 1.2 shows an exaggerated version of this, in which we can see a harder spectrum (γ_2) overtake a softer one (γ_1), as well as the problems that arise in the opposite case at the intersection of γ_0 and γ_1 . If these were truly spectral indices produced by different source populations, the only way in which this transition could occur was if the decline of the lower-energy (γ_0) sources coincided with the rise of the higher-energy (γ_1) ones. This is incredibly unlikely.




Figure 1.2 Example of three intersecting power laws with various spectral indices. Dashed lines indicate regions where the continuation of a given power law would cause deviation from the observed spectrum. The x- and y-axes represent energy and flux, respectively, in arbitrary units.

A different approach is to treat the softening spectrum not as a continuous power law from a new source population, but instead as the decay of the sources that provided the prior power law. The charged nature of cosmic rays makes this behavior possible. For a source to accelerate a given charged particle over time, it must be capable of preventing that particle's escape. A source of a given size and magnetic field is able to contain relativistic charged particles with energy $E \propto Z$ (described in greater detail in Sec. 1.4). Therefore, the same source can accelerate iron nuclei ($Z = 26$) to energies up to 26 times as high as protons. This mechanism, as suggested by Peters [10], leads naturally to the acceleration of higher- Z nuclei to higher energies, making the decay of a given source population appear more gradual. The knee, in this case, would not be a continuous power law, but would instead be the sum of multiple decaying power laws for elements with increasing Z .

In order to study this hypothesis, experiments capable of observing the knee with excellent energy resolution and composition separation are needed in order to observe the fine structure anticipated. This leads to the study of the knee being one of the most interesting and challenging topics in cosmic-ray astrophysics.

1.3 Composition

The full cosmic-ray spectrum consists of about 85% protons, 12% helium nuclei, 2% heavy nuclei ($Z > 2$), and 1% electrons and positrons. Figure 1.3 shows the flux of a variety of compositions in the GeV to TeV energy range, as measured by balloon-borne (e.g. CREAM) and satellite (e.g. AMS) experiments. These results indicate that protons account for the bulk of the flux up until about 10 TeV, at which point helium becomes dominant. A change in composition as a function of energy is consistent with the Peters cycle mechanism described in Sec. 1.2. If the knee represents the acceleration limit of a source or multiple sources, then the composition should become heavier as energy increases.

Differences in chemical abundances at energies of hundreds of MeV per nucleon are studied to gain insight into cosmic-ray acceleration and propagation processes. Figure 1.4 shows a comparison of solar and cosmic-ray elemental abundances. In general, the abundance




Figure 1.3 Spectrum for selected elemental components as compiled from a variety of balloon-borne and satellite experiments [11]. Energy in the x-axis is given as kinetic energy per particle per nucleus.




Figure 1.4 Relative abundance of isotopes as found in the galactic cosmic rays (*black*) and the solar system (*blue*). Valleys in solar system abundances are less prominent in galactic cosmic rays due to spallation of higher-Z nuclei. This figure represents a collection of work from CRIS [12] and BESS [13], along with a compilation by K. Lodders [14], as presented in ACE News #83 - Oct 6, 2004.

of elements in the cosmic-ray flux matches the abundances found in our solar system, with some notable exceptions for some lighter ($Z = 3 - 5$) and heavier ($Z = 20 - 25$) elements. In these cases, primary cosmic rays propagating to Earth are fragmented into lighter elements via spallation processes in the interstellar medium. Thus, common, stable elements such as carbon ($Z = 6$) and iron ($Z = 26$) produce the observed excesses. The strength of the excesses indicates that galactic cosmic rays have an average age of about 10^7 years, which will be discussed in greater detail in Sec. 1.6.

At energies above 10 TeV, direct detection of cosmic rays becomes unfeasible due to the low rate. Ground-based experiments, however, are capable of detecting the secondary air shower particles produced by collisions of the high-energy cosmic-ray primary with the atmosphere. These showers, described in more detail in Sec. 2.1, do not allow for the same power of composition separation, but several experiments are sensitive to the mean logarithmic mass ($\langle \ln(A) \rangle$) of the cosmic-ray primary. Figure 1.5 shows data from some ground-based experiments, indicating a transition to heavier elements around the knee (*left*), and a transition back to a proton-dominated flux near the ankle (*right*). At energies in excess of 10^{19} eV, the Telescope Array and Pierre Auger collaborations report compositions consistent with proton and iron, respectively [15, 16].

1.4 Acceleration Mechanisms

The extreme energies of cosmic rays required new theoretical models of acceleration. In 1949, Enrico Fermi proposed a model based on stochastic interactions that would naturally produce a power-law spectrum [19]. In this model, randomly moving magnetic clouds reflect charged particles as they travel through interstellar space. In the case that a particle collides head-on with a cloud, the velocity of the particle would be increased by the velocity of the cloud:

$$\Delta E_{head} = \frac{1}{2}m(v_{particle} + v_{cloud})^2 - \frac{1}{2}mv_{particle}^2 \quad . \quad (1.3)$$

In the case that a particle is moving in the same direction as a magnetic cloud, and strikes tail-on, the velocity of the particle would be correspondingly decreased by the velocity of the

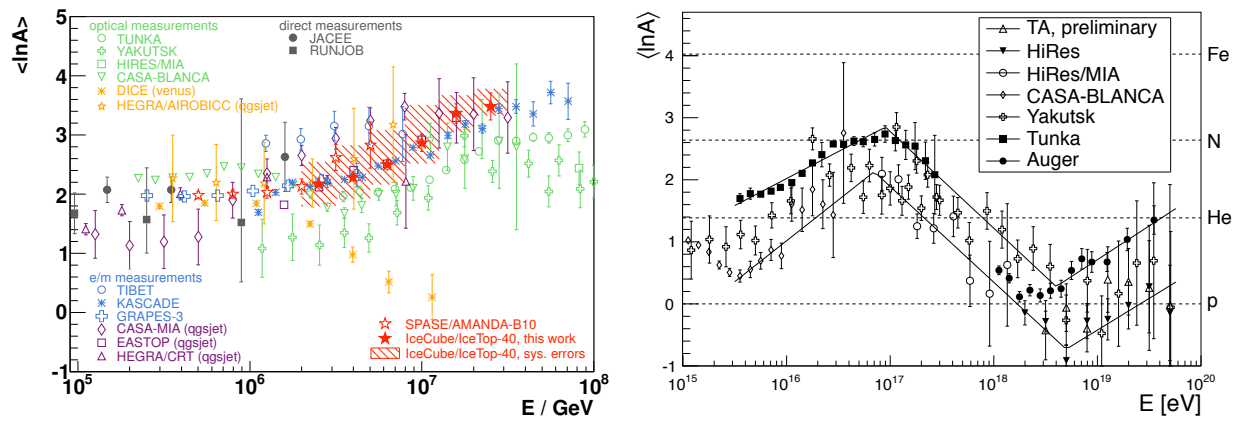


Figure 1.5 Mean logarithmic mass as a function of cosmic-ray primary energy. Measurements near the knee (*left*) [17] are as measured by a variety of direct and indirect detection experiments. Measurements near the ankle (*right*) are strongly influenced by the hadronic model used [18]; here QGSJET01 is used.

cloud:

$$\Delta E_{tail} = \frac{1}{2}m(v_{particle} - v_{cloud})^2 - \frac{1}{2}mv_{particle}^2 . \quad (1.4)$$

The average energy gain for a particle, then, is given by $\Delta E = \Delta E_{head} + \Delta E_{tail} = mv_{cloud}^2$.

This results in a total relative gain of:

$$\frac{\Delta E}{E} = 2 \left(\frac{v_{cloud}}{v_{particle}} \right)^2 . \quad (1.5)$$

This process is referred to as 2nd-order Fermi acceleration, because the fractional energy gain is proportional to the square of the velocities. This model relies on the particles moving much faster than the clouds, so the fractional energy gain ($\propto (v_{cloud}/v_{particle})^2$) is very small per collision, making the process inefficient.

Now, instead of randomly moving magnetic mirrors, consider a strong magnetic shock front moving through the interstellar medium with velocity U , shown in Fig. 1.6a. In front of the shock front is the “downstream” region with density ρ_d , pressure P_d , and temperature T_d . Behind the shock front is the “upstream” region with corresponding density, pressure and temperature ρ_u , P_u , T_u . In the reference frame of the shock, downstream material approaches the shock at $v_d = U$, while gas in the upstream region behind it moves away at velocity v_u . We can determine the ratio of v_u to v_d by using conservation of particle number: $\rho_u v_u = \rho_d v_d$. In strong shock fronts, we expect

$$\rho_u/\rho_d = (c_p/c_v + 1)/(c_p/c_v - 1) , \quad (1.6)$$

where c_p/c_v is the ratio of heat capacities. In the case of fully ionized plasma, $c_p/c_v = 5/3$, which results in $v_d/v_u = 4$. Therefore, as seen in Fig. 1.6b, gas flows away from the shock front at $v_u = 1/4 U$. In the rest frame of the downstream medium, shown in Fig. 1.6c, particles upstream approach the front with velocity $3/4 U$. In the upstream rest frame, shown in Fig. 1.6d, particles downstream also stream towards the region with velocity $3/4 U$. Particles crossing the shock front from either the upstream or downstream regions are likely to have a head-on collision and gain energy.

Defining $V = 3/4 U$, as a particle encounters the shock front, the average energy gain per collision is given by:

$$\begin{aligned}\Delta E &= \frac{1}{2}m(v_{particle} + V)^2 - \frac{1}{2}mv_{particle}^2 \\ &= \frac{1}{2}m(2v_{particle}V + V^2) \ .\end{aligned}\tag{1.7}$$

For $v_{particle} \gg V$, this simplifies to:

$$\Delta E \approx mv_{particle}V \ .\tag{1.8}$$

The relative energy increase resulting from this simplified case is $\Delta E/E \approx 2V/v_{particle}$. A relativistic treatment changes this value to $(4/3)V/c$, but the general result remains the same: this scenario provides fractional energy gains proportional to the velocity of the shock front. Due to the similarity to Fermi's original theory, this method is known as 1st-order Fermi acceleration. Its increased efficiency over the 2nd-order process makes it the preferred method for cosmic-ray acceleration.

To study the spectrum produced by this method, assume the average energy of a particle from the injection population after a collision with the shock front is given by $E = \beta E_0$. Taking P to be the probability that a particle remains in the shock front after one collision, then after k collisions we have the following distributions:

$$N(k) = N_0 P^k \ ,\tag{1.9}$$

$$E(k) = E_0 \beta^k \ .\tag{1.10}$$

These equations can be used to determine the counts as a function of energy by eliminating k :

$$\frac{N(\geq E)}{N_0} = \left(\frac{E}{E_0} \right)^{\ln P / \ln \beta} \ ,\tag{1.11}$$

which leads to:

$$N(E)dE \propto E^{\ln P / \ln \beta - 1} dE \ .\tag{1.12}$$

Next, to solve for the exponential term, we need to estimate the probability that a particle remains in the shock front P . By averaging over angles and velocities, the rate of particles

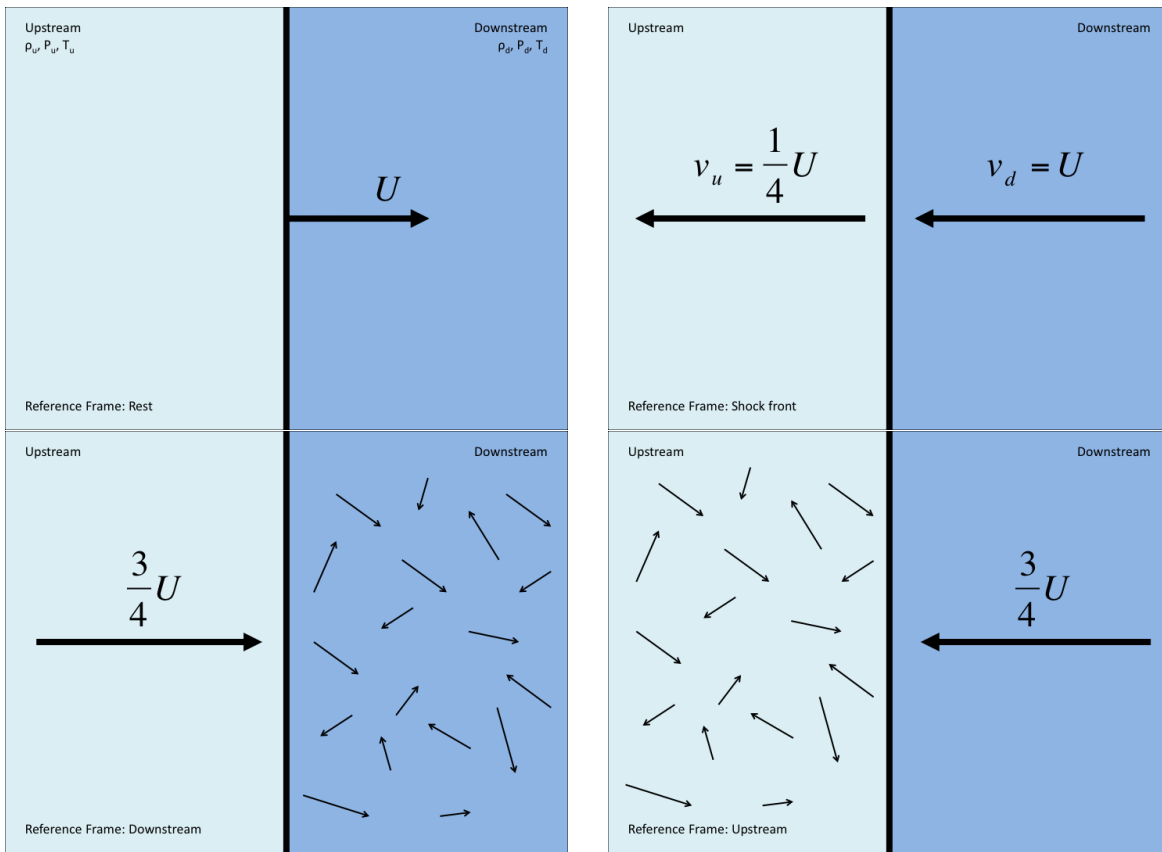


Figure 1.6 First-order Fermi acceleration as shown in the reference frame of (a) an observer at rest, (b) the shock front, (c) the downstream region, and (d) the upstream region. Note that in the reference frames of both the upstream and downstream regions, particles are incoming with $v = \frac{3}{4}U$.

crossing the shock front in either direction is $1/4 nc$, where n is the number density of energetic particles. Particles in the downstream region are rarely lost as they are constantly being swept up in the shock. Particles in the upstream region flow away at a rate of $nv_u = 1/4 nU$. With these two values, the fraction of particles lost is $(1/4)nU/(1/4)nc = U/c$. The fraction of particles remaining after once cycle is

$$P = 1 - \frac{U}{c} . \quad (1.13)$$

Solving for β is easier: we know $\Delta E/E = (4/3)V/c$, so

$$\beta \equiv \frac{E}{E_0} = 1 + \frac{4V}{3c} . \quad (1.14)$$

With $U \ll c$, this gives us

$$\ln P = \ln\left(1 - \frac{U}{c}\right) \simeq -\frac{U}{c} \quad (1.15)$$

$$\ln \beta = \ln\left(1 + \frac{4V}{3c}\right) \simeq \frac{4V}{3c} = \frac{U}{c} . \quad (1.16)$$

The result is $\ln P/\ln \beta = -1$, so

$$N(E)dE \propto E^{-2}dE . \quad (1.17)$$

The resultant spectral index of -2 is close to the measured value of $\gamma \sim 2.7$ mentioned in Sec. 1.2. The difference in values is believed to be a result of the propagation of cosmic rays through the galaxy, and will be explored in Sec. 1.6.

Fermi acceleration is known as a “bottom-up” process, in which a low-energy particle is accelerated through repeated interactions with its environment. Some more exotic theories propose “top-down” processes, in which extremely high-energy, exotic objects decay, resulting in ultra-high energy cosmic rays among other products. Theorized candidates for these exotic objects include extremely massive dark matter particles, and topological defects left over from the Big Bang. Typically, the decay in these models results in other high-energy byproducts, such as gamma rays and neutrinos. Results constraining these models [20] make the bottom-up acceleration method the more likely of the two.

1.5 Candidate Sources

The simplest approach to determine potential sources of cosmic rays is to calculate the highest energy cosmic rays a given source could magnetically confine. The motion of a charged particle in a magnetic field is given by equating the force on the particle from an external magnetic field to the centripetal force, as in Eq. 1.1. Thus, we can calculate the maximum energy attainable by a source solely through containment principles as:

$$E_{\max} = ZeBR \quad . \quad (1.18)$$

This is known as the Hillas criterion [21]. Figure 1.7 is a Hillas Plot, displaying the magnetic field of potential sources of the highest energy (100 EeV and above) cosmic rays as a function of their size. At the upper left end of Fig. 1.7 we have very small sources with extremely high magnetic fields, like neutron stars and gamma ray bursts. Other larger sources with lesser magnetic fields are also viable potential sources, including active galactic nuclei and colliding galaxies. While the plot also shows that supernova remnants (SNRs) do not have the requisite size and magnetic field strength to produce the highest energy cosmic rays, it does not preclude them from being a source for energies around the knee. Note the separate lines that exist for protons and iron; particles with higher charge are more tightly confined, and therefore able to be accelerated to higher energies by a given source.

The detection of cosmic-ray sources is complicated by their motion in the galactic magnetic field. For ease of calculation, the gyroradius of a relativistic particle can be written as:

$$r_g/\text{meter} = 3.3 \times \frac{(E/\text{GeV})}{Z(B/\text{Tesla})} \quad . \quad (1.19)$$

Given an average galactic magnetic field strength of 50 nT, protons with $E \geq 10^{19}$ eV would have a gyroradius of over 10^{18} m, about the diameter of our galaxy. Direct identification of high-energy sources by searching for correlations in arrival direction is possible for cosmic rays of these extreme energies, but the very low flux makes it challenging. Experiments such as the Telescope Array and the Pierre Auger Observatory attempt the identification through




Figure 1.7 Hillas Plot indicating potential sources of the highest energy cosmic rays based on size and magnetic field strength [22].

this approach, and have massive effective areas — in the case of Auger about 3000 km^2 [23] — to account for the low flux.

An alternative to the detection of cosmic rays is the detection of secondary particles produced by interactions near the source. If there is a high nucleon density near a source, accelerated cosmic rays can interact with the medium, producing neutrinos via meson decay. Neutrinos have no electrical charge and a very small weak cross-section, so they can travel, unperturbed, directly outward from the source, carrying directional and energetic information. The same qualities that make them good messenger particles, however, also make them very difficult to detect. The small cross-section requires detectors with large effective areas of the order of 1 km^3 [24]. Detectors like IceCube and the proposed KM3NET are both capable of this type of source detection.

The interactions that produce neutrinos near the source should produce high-energy gamma rays as well via the decay of neutral pions [25]. Gamma rays also have no charge, but they are subject to absorption, as well as energy losses through pair-production with background photons, significantly reducing their mean free path. The gamma-ray spectrum of a source can indicate whether the neutral pion decays expected near a cosmic-ray accelerator are occurring. Using this method, the Large Area Telescope on the Fermi satellite has discovered hadronic emission in the GeV range from supernova remnants in our galaxy [26]. Other ground-based telescopes, such as VERITAS, HESS, MAGIC, and the newly completed HAWC Observatory continue the search in the TeV range.

1.5.1 Galactic Sources

Supernova remnants (SNRs) provide both the environment and the required energy density to produce the observed cosmic-ray spectrum approaching the knee. Environmentally, the shocks produced by young SNRs are capable of first-order Fermi acceleration, providing the acceleration mechanism. From an energy density standpoint, the connection between SNRs and cosmic rays was first proposed by Baade and Zwicky in 1934 [27]. A typical core collapse supernova ejects masses with a total kinetic energy on the order of 10^{51} ergs.

With an estimated rate of 1 supernova every 30 years, this corresponds to a power output of 10^{42} ergs/s. The observed cosmic-ray energy density in our galaxy is $\sim 1\text{eV}/\text{cm}^3$. This, combined with the volume of our galaxy ($\sim 4 \times 10^{66} \text{cm}^3$) and an average escape time of 6×10^6 years results in a power of 5×10^{40} erg/s needed to maintain a stable state. Therefore, an acceleration efficiency of a few percent would be enough to produce the observed cosmic-ray spectrum.

More recent studies indicate that a young, nearby supernova could produce both the spectrum and composition observed around the knee, and even suggest specific sources [28, 29, 30]. While it is not clear if these particular models are accurate, results from the Fermi satellite linking neutral pion decays with two galactic SNRs are the first evidence for this long-standing general theory [26].

1.5.2 Extragalactic Sources

As the highest-energy cosmic rays are not contained by the galactic magnetic field, we look to extragalactic sources. As Fig. 1.7 shows, some of the highest energy sources in the universe are considered as candidates for the acceleration of ultra-high energy cosmic rays (UHECRs). This includes, but is not limited to, neutron stars, active galactic nuclei (AGNs), gamma ray bursts (GRBs), and colliding galaxies. GRBs, while incredibly energetic, could most likely not accelerate heavy elements: the extremely high number of gamma rays produced in a GRB would photo-disintegrate any nucleus before it could be accelerated to the highest energies. This rules out GRBs if the composition at the highest energies is heavy, as indicated in results from the Pierre Auger collaboration. If the composition remains light, photon-proton interactions should produce many TeV neutrinos at the source. Recent IceCube results show no correlation of TeV neutrinos with observed GRBs, strongly constraining their contribution to the UHECR spectrum [31].

1.6 Propagation

The motion of cosmic rays in our galaxy is anything but straight-forward. After their acceleration at a source, cosmic rays are subject to a wide variety of interactions as they travel towards us. Upon arrival, cosmic rays can carry information on magnetic fields, dust layers, background photons, and a variety of other mediums they may have encountered. This section covers some aspects of the propagation of galactic cosmic rays and the diffusion equation. For more on the propagation of cosmic rays in the galaxy, refer to [32]. The diffusion equation and various approximations are covered in greater detail in [33, 34, 35].

The Milky Way is roughly the shape of a disc, with a radius of ~ 15 kpc and a thickness of ~ 300 pc. A weak magnetic field ($\sim 2 - 6 \mu\text{G}$) acts throughout the disc, with a higher field strength ($\sim 20 - 30 \mu\text{G}$) found in the center. This magnetic field has two components of roughly equal strength: regular and turbulent. The regular component consists of field lines many kpc in size, oriented along the spiral arms. The turbulent component, by contrast, has a characteristic scale on the order of 10-100 pc. It is difficult to measure the magnetic field strength and distribution of our own galaxy, but Fig. 1.8 shows the magnetic field strength and direction in M51 as an example. The magnetic field lines and the cosmic rays that travel along them are interconnected; cosmic rays streaming along field lines produce Alfvén waves, that in turn scatter incident cosmic rays.

As mentioned in Sec. 1.3, we can also study the elemental abundances in the cosmic ray spectra for elements below 1 PeV. As cosmic rays travel through the galaxy, spallation processes cause common, stable nuclei to break into less common elements. The solar elemental abundances are treated as our initial state, so the difference in the ratio between an element and its spallation byproduct (e.g. carbon and boron) compared to the initial state gives a measurement of how much material the cosmic rays traveled through. These measurements indicate that an average cosmic ray travels through a column density of $6 - 10 \text{ g/cm}^2$ before reaching Earth. The confinement time can be calculated using this measurement: given the observed column density, cosmic rays traveling near c through the typical hydrogen density




Figure 1.8 Radio emission in M51 with magnetic field lines overlaid. The strong, regular magnetic field lines follow the spiral arms. Image courtesy of NRAO/AUI.

of the galaxy ($\sim 1 \text{ cm}^{-3}$) propagate for $\sim 10^7$ years. Study of chemical abundances also shows a lesser difference in ratios at higher energies. This simple but important result tells us that higher energy cosmic rays sense less matter, and therefore have shorter confinement times. Since the higher energy cosmic rays spend less time in travel, it also means that the acceleration and propagation processes must be different.

The arrival direction distribution of cosmic rays in the TeV energy range is primarily shaped by diffusion in the galactic magnetic field. Ginzburg and Syrovatskii [36] were the first to analytically write the diffusion equation, which describes the change in number density over time of cosmic rays as a function of energy, position, and particle type:

$$\begin{aligned} \frac{\partial N}{\partial t} = & \nabla \cdot (D_i \nabla N_i) - \frac{\partial}{\partial E} [b_i(E) N_i(E)] - \nabla \cdot \mathbf{u} N_i(E) \\ & + Q_i(E, t) - p_i N_i + \frac{v\rho}{m} \sum_{k \geq i} \int \frac{d\sigma_{i,k}(E, E')}{dE} N_k(E') dE' \quad . \end{aligned} \quad (1.20)$$

The first term of this equation describes the change in the number density due to the influence of diffusion. The diffusion constant D can be written in terms of the diffusive mean free path λ_D and particle velocity v as $D = (1/3)\lambda_D v$. The second term gives the influence of energy loss or gain through particle interactions, with $b_i(E)$ representing the average rate of energy gain per particle. The third term shows the effect of convection with some characteristic velocity \mathbf{u} . The fourth term represents the contribution of accelerating sources. The fifth term covers the loss of nuclei due to spallation processes and nuclear decay, where:

$$p_i = \frac{v_p \rho \sigma_i}{m} + \frac{1}{\gamma \tau_i} = \frac{v_p \rho}{\lambda_i} + \frac{1}{\gamma \tau_i} \quad . \quad (1.21)$$

In this term, λ_i is the mean free path given the density of the medium ρ , and $\gamma \tau_i$ is the lifetime of the nucleus. The sixth and final term represents the inverse of the fifth term: the gain of particles due to cascades and spallation processes applied to heavier nuclei ($k \geq i$).

To study the effect of the diffusion equation on the cosmic-ray spectrum, first consider a simplified model in which we have a uniform distribution of cosmic rays inside a containment volume. Cosmic rays inside the volume have a probability of escape per unit time that depends on their energy, and on average spend a time τ_{esc} before escaping. As each particle

travels through the volume, it is deflected such that τ_{esc} is much greater than the time it would take to travel directly out of the volume, and passes through an average amount of matter given by

$$\lambda_{esc} = \rho v \tau_{esc} \quad . \quad (1.22)$$

This model is known as the “leaky box” approximation.

As discussed by Gaisser [34], in the leaky box model, a delta source Q that injects a spectrum $N_0(E)$ at time t_0 will result in the observed spectrum:

$$N(E, t) = N_0(E) e^{-t/\tau_{esc}} \quad . \quad (1.23)$$

Assuming a steady-state solution while neglecting energy losses and gains through particle interactions (the second term in the diffusion equation) yields the following:

$$\frac{N_i(E)}{\tau_{esc}(E)} = Q_i(E) - \left(\frac{v\rho}{\lambda_i} + \frac{1}{\gamma\tau_i} \right) N_j(E) + \frac{v\rho}{m} \sum_{k>i} \sigma_{i,k} N_k(E) \quad . \quad (1.24)$$

To further simplify this equation, we consider a proton-only spectrum in which the contribution due to the spallation of heavier nuclei is negligible. The proton lifetime is so long that the decay term $1/(\gamma\tau_P)$ can also be neglected. Substituting $\rho v = \lambda_{esc}/\tau_{esc}$ gives us

$$N_P(E) = \frac{Q_P(E)\tau_{esc}(E)}{q + \lambda_{esc}(E)/\lambda_P} \quad . \quad (1.25)$$

Based on an experimentally determined ratio of secondaries to primaries, Gaisser [34] cites a fit of

$$\lambda_{esc} = 10.8\beta \times \left(\frac{4}{R} \right)^\delta \quad , \quad (1.26)$$

where $\beta = vc$ and R is the rigidity (ZE). For protons, $\lambda_P \sim 55\text{g}/\text{cm}^2$, allowing the following simplification since $\lambda_P \gg \lambda_{esc}$:

$$N_P(E) \approx Q_P(E)\tau_{esc}(E) \quad . \quad (1.27)$$

From a combination of Eqs. 1.22 and 1.26, we know $\tau_{esc} \propto \lambda_{esc} \propto E^{-\delta}$. The spectral index produced by the source (α), then, should be the difference between the spectral index observed ($\gamma \simeq 2.7$), and the fit to the secondary to primary ratio ($\delta \simeq 0.6$), or $\alpha \simeq 2.1$. With

a highly simplified model and approach, we arrive at a value close to the source spectrum expected from Fermi acceleration methods, as discussed in Sec. 1.4.

Analytic solutions to the diffusion equation can be found assuming certain symmetries or boundary conditions, but there are also numerical approaches. Software packages such as GALPROP [37] are frequently used to produce numerical solutions. The ability to deal with non-uniform distributions is very important, since, unlike in the leaky box model, the cosmic-ray density varies across the galaxy. The variations in density manifest as the observed anisotropy, which will be discussed in greater detail in Sec. 1.7.

1.7 Anisotropy

Scrambling of cosmic rays in the energy range around the knee in magnetic fields should lead to a completely isotropic sky. However, the study of the arrival direction distribution of cosmic rays has shown that this is not the case. In the last few decades, a number of experiments have provided long-term, statistically significant evidence of a faint anisotropy in the cosmic-ray arrival distribution across six orders of magnitude in energy, from tens of GeV to tens of PeV. The small amplitude of the observed large-scale anisotropy, on the order of 10^{-4} – 10^{-3} , alongside the energy-dependent topology and angular structure, hint at multiple phenomenological contributions to the observations. Muon detectors in the Northern Hemisphere observed cosmic-ray anisotropy from energies of several tens to hundreds of GeV, which is beyond the direct solar wind influence [38, 39, 40, 41]. Various surface arrays and underground detectors reported observations in the TeV energy range [42, 43, 44, 45, 46, 47, 48, 49, 50, 51], in some cases extending up to hundreds of TeV [52, 53]. Recently, the IceCube Neutrino Observatory reported the first observations of cosmic-ray anisotropy in the Southern Hemisphere, with unprecedented event statistics in the energy range between 10 TeV and 1 PeV [54, 55, 56, 57]. Figure 1.9 shows the first observations of this large-scale anisotropy in both the Northern (*top*) and Southern (*bottom*) Hemispheres.

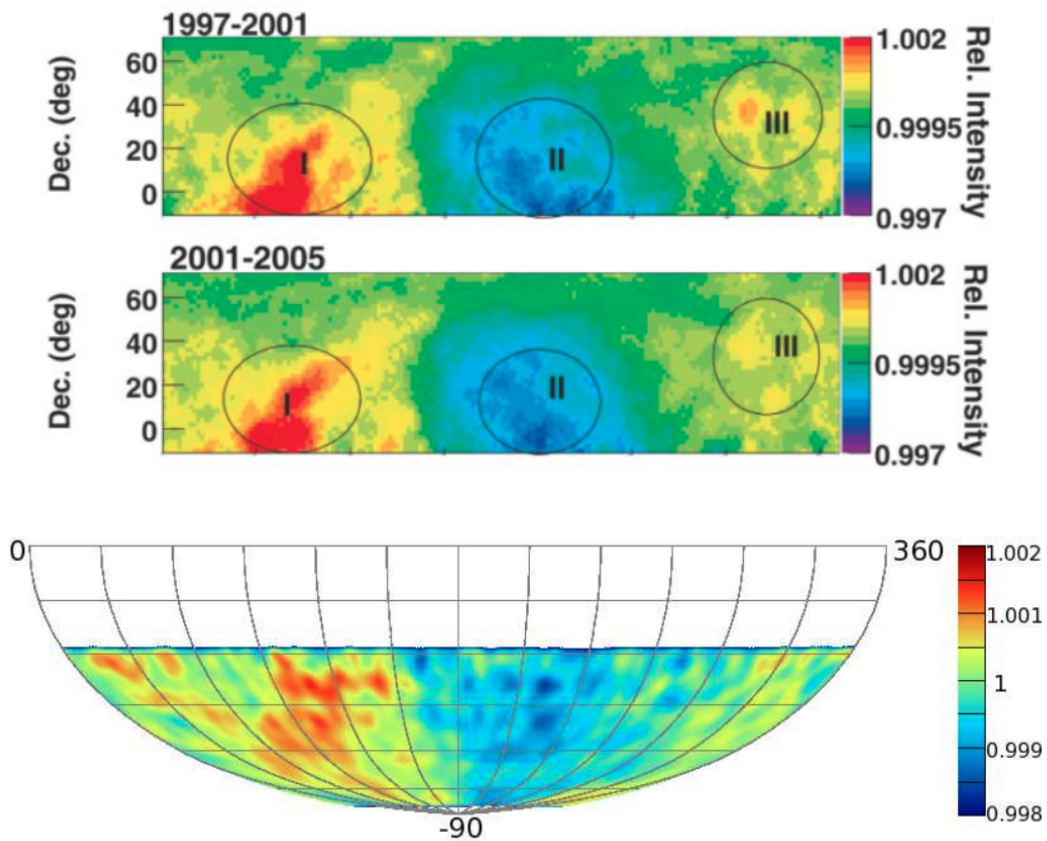


Figure 1.9 The first observations of large-scale anisotropy in equatorial coordinates in the Northern (*top*) and Southern (*bottom*) Hemispheres. The Northern Hemisphere plot is as observed by the Tibet-AS γ experiment over four years of observation [42]. The Southern Hemisphere plot is as observed by IceCube in one year of observation in a 22-string configuration [54].

All the observations show a topologically consistent large-scale anisotropy structure across the sky in celestial coordinates. In the energy interval from 60 GeV to 100 TeV, the cosmic-ray arrival direction distribution shows a wide relative excess in the 30° - 120° range in right ascension, and a deficit in the 150° - 250° range. The degree of this directional asymmetry is found to increase with energy up to about 10 TeV, and to decrease at higher energies up to about 100 TeV. In the 100-300 TeV energy range, a major change in the anisotropy topology is observed by both the EAS-TOP array in the Northern Hemisphere [52], recently confirmed by the Tibet AS γ Array [49], and by IceCube in the Southern Hemisphere [56]. IceCube data show that at an energy in excess of a few hundred TeV, the cosmic-ray arrival distribution seems mostly characterized by a relative deficit at a right ascension of 60° - 120° [56], with amplitude increasing with energy.

It is evident from the observations that the anisotropy cannot be described with a simple dipole, although this is typically used to estimate the amplitude. Instead, a quantitative description of the anisotropy requires a characterization of the distribution over a wide range of angular scales. In particular, the arrival direction distribution can be described as a superposition of spherical harmonic contributions, where most of the power is in the low-multipole ($\ell \leq 4$) terms. A fit using only the low-multipole terms, however, describes the data poorly, indicating that the higher-multipole terms must be accounted for as well. In fact, statistically significant smaller angular scale features, with amplitudes on the order of 10^{-5} – 10^{-4} , have been observed in the TeV energy range by several experiments [58, 59, 55, 46, 51]. More detailed studies with detectors in the Northern Hemisphere have shown that the energy spectrum of the cosmic-ray flux in the most dominant excess region is harder than that of the isotropic cosmic-ray flux [59, 46, 51]. Although further confirmation is needed, this result might indicate that whatever produces the localized region of excess may be responsible for the spectral deviation as well.

The evolution of cosmic-ray anisotropy in energy has also been observed by experiments sensitive to ultra-high-energy particles. The Pierre Auger Observatory, in a search for a dipole and quadrupole component of the cosmic-ray arrival direction distribution, found

that a large shift in the phase of the anisotropy occurs at about 1 EeV as well [60, 61, 62]. The relative excess moves towards the range in right ascension that includes the Galactic anticenter direction, which may be an indication that a new population of extragalactic cosmic rays begins to become dominant. At energies in excess of 1 EeV, cosmic rays are less bent by Galactic and intergalactic magnetic fields, making it possible to transition into the regime of cosmic-ray astronomy [63].

1.7.1 Large-Scale Anisotropy

Compton and Getting first proposed a model for large-scale anisotropy in 1935 [64]. The motion of the solar system as it orbits about the galactic center should induce some anisotropy with respect to the rest frame of cosmic rays, with an excess observed in our direction of motion and a corresponding deficit opposite. The strength of this relative intensity could be calculated using:

$$\frac{\Delta I}{\langle I \rangle} = (\gamma + 2) \frac{v}{c} \cos \rho \quad , \quad (1.28)$$

where γ is the cosmic-ray spectral index, v the orbital speed ($\sim 200 - 240$ km/s [65]), and ρ the angle between \vec{v} and the average direction of cosmic-ray motion [66]. With the observed spectral index and velocity, a dipole is expected with an amplitude of $\sim 0.35\%$. This anisotropy has not been observed, leading to the idea that galactic cosmic rays are co-rotating with the galaxy. A similar dipole anisotropy is also expected due to the Earth's motion around the Sun when the sky is plotted in fixed-sun coordinates. This anisotropy has been observed, and is described in greater detail in Sec. 4.1.4.

One of the earliest analyses of large-scale anisotropy was performed by the Tibet-AS γ experiment. Through the study of the anisotropy over several years, they found that the maps were consistent over periods of strong and weak solar activity [67]. This indicated that the solar influence was not the cause of the observed anisotropy. The authors also found that the large-scale anisotropy showed an energy dependence. The Compton-Getting effect, shown in Eq. 1.28, has no energy dependence, so the anisotropy is also not due to the motion of our solar system through the galactic cosmic-ray medium.

As mentioned in Sec. 1.6, diffusion has an effect on the observed cosmic-ray spectrum. Specifically, the diffusion and injection terms in Eq. 1.20 indicate that the cosmic-ray density changes based on the distance to a source and the time of injection. Given the diffusive transport of cosmic rays through the galactic magnetic field, particles emitted by nearby sources would arrive at Earth in a wide, general excess, leading to a dipole structure with the excess pointing towards the nearby source(s) [68]. The full diffusion equation predicts a dipole anisotropy d , given by

$$d(E) = \frac{3D(E)}{c} \frac{\nabla n(E)}{n(E)}, \quad (1.29)$$

where D is the diffusion coefficient and n the number density of cosmic rays [33]. The amplitude predicted by this approach is in the 10^{-3} to 10^{-2} range, with an energy dependence. This model also predicts a shift in the direction of the dipole related to a shift in sources. At lower energies, the anisotropy might be dominated by a star-forming region in the Orion arm, relatively close to our planet. At higher energies the dipole might shift to more energetic sources near to the center of the galaxy.

Recent work, however, has shown that in the presence of strong regular and turbulent magnetic fields, the ability of the observed anisotropy to point back to sources is negligible [69]. In this work, cosmic-ray propagation is simulated through random realizations of both turbulent and regular galactic magnetic fields. The authors found that in the case of weak turbulence, cosmic rays with energies in the sub-PeV range arrive at the source in a dipole, with the phase pointing along the regular component of the magnetic field. In the case of strong turbulent fields, a dipole anisotropy still tends to arise, but with random phase. In both cases, the dipole phase shows no alignment with the source distribution.

1.7.2 Small-Scale Anisotropy

As mentioned earlier, the anisotropy in the cosmic-ray arrival direction distribution is not well-fit by a dipole. Statistically significant small-scale structures have been found by a variety of experiments. Milagro was the first experiment to observe localized excesses – two structures in the Northern Hemisphere titled Regions A & B, shown in Fig. 1.10. Since




Figure 1.10 Small-scale anisotropy in equatorial coordinates as observed by Milagro [59].

The map is smoothed with a 10° angular radius, and structures larger than 30° in size should not be visible due to the suppression of the lower multipole moments. Regions A and B stand out as the only significant features, other than the Cygnus region which is visible because gamma rays have not been removed from this data.

then, other experiments have confirmed their observations [46, 51], and extended them into the Southern Hemisphere [55], as seen in Fig. 1.11.

Small-scale anisotropy can result from propagation effects in the local magnetic field. The work in [69] indicates that small-scale anisotropy arises naturally from the scattering of a dipole signal. Through scattering, power in low-multipole moments can contribute to power in higher multipoles. Recent work indicates that the resultant spectrum can be predicted analytically through the use of Liouville’s Theorem [70, 71]. In this approach, the small-scale anisotropy reflects turbulence in the local galactic magnetic field, and measurements can provide information on magnetic fields within 10 pc of Earth. Other models suggest that heliospheric effects such as particle re-acceleration due to turbulent magnetic reconnection could produce spectral differences corresponding to local excess regions [72, 73]. Models involving the heliosphere are appealing due to the proximity of Region A to the heliotail.

1.8 Overview

The analysis in this thesis uses 6 years of data – more than 300 billion cosmic-ray events recorded between May 2009 and May 2015 – to study the energy- and time-dependence of the cosmic-ray anisotropy as observed by IceCube. The large size of the data set allows for




Figure 1.11 Relative intensity (*left*) and significance (*right*) of small-scale anisotropy in equatorial coordinates as observed by IceCube [55]. The map is smoothed with a 20° angular radius, and structures larger than 60° in size should not be visible due to the subtraction of the lower multipole moments.

the observation of structures on angular scales approaching the limit of the dataset's angular resolution, as well as a detailed study of the energy-dependence of the anisotropy in the TeV to PeV energy range. At PeV energies, additional data from the IceTop air-shower array is used to provide an independent analysis, with special attention paid to a new, likelihood-based method of reconstructing the cosmic-ray primary energy and composition. Many of the techniques used in this analysis are described in detail in [54, 55, 56] and [57].

Chapter 2

Detection

This work focuses on the analysis of cosmic rays with energies approaching the knee and above. The flux of cosmic rays with energies near the knee is approximately one particle per square meter per year, so detectors with very large effective areas are needed. In addition, the atmosphere is opaque to cosmic rays in this energy range, so direct detection experiments would need to be located at high altitude or in space. A detector in orbit with a square kilometer of effective area would record these high-energy cosmic rays at a rate of 0.03 Hz. Thus, direct detection at these energies is unfeasible. Fortunately, there are other approaches available through indirect detection.

In this chapter, we focus on the basic properties of cosmic-ray air showers, the principles that allow us to detect them, and the details of the IceCube Neutrino Observatory.

2.1 Air Showers

The interaction between a high-energy cosmic ray, written here as a proton, and some absorber nucleus X can be written as:

$$p + X \rightarrow p + Y + \pi^{\pm,0} + K^{\pm,0} + \dots \quad , \quad (2.1)$$

where Y represents the fragmented nucleus. The resultant proton, known as the “leading particle” will go on to form the hadronic component of the shower. The mesons produced go on to form the electromagnetic component of the shower. At energies above the knee, almost any mesons can be created, but pions with little transverse momentum are the dominant component. The products of a cosmic-ray primary with A nucleons and energy E_0 can also




Figure 2.1 Graphical representation of the development of a cosmic-ray air shower. The interaction of the cosmic-ray primary with a nucleus in the atmosphere can produce mesonic ($\pi^{\pm,0}$) and hadronic (N) byproducts. Common interactions and decay processes for these byproducts are shown, most notably the decay of charged pions to muons and neutrinos, the decay of neutral pions to high-energy gamma rays, and the interactions in a hadronic cascade. Source: CERN

be estimated using Eq. 2.1 for A protons, each with energy E_0/A . Figure 2.1 shows the primary channels through which an air shower proceeds. The three main components of the shower – hadronic, muonic, and electromagnetic – are discussed here in greater detail.

The hadronic component of the shower, while rarely detected, strongly influences the size and shape of the shower. In the interaction described in Eq. 2.1, an incident cosmic ray produces a leading particle, a fragmented nucleus, and a variety of mesons. Inelasticity is a measure of how much of the incident particle’s energy was lost to meson production. Typical interactions have an inelasticity of 0.6 [34], meaning that the leading particle retains 40% of the incident energy, more than any other individual component. The remainder of the energy goes mostly to the production of pions, where it is split evenly between the three flavors. The leading particle can then go on to interact again via the same process, reproducing these steps many times before reaching the surface. High transverse momentum in the leading particle can lead to a large spread in the shower’s lateral distribution. This is an important effect for interactions high in the atmosphere, where the thin atmosphere allows for leading particles with small transverse momentum to travel a large distance before interacting.

The charged pions produced by the shower weakly decay primarily via the following channels:

$$\begin{aligned}\pi^\pm &\rightarrow \mu^\pm + \nu_\mu \\ \pi^\pm &\rightarrow e^\pm + \nu_e .\end{aligned}$$

Charged pions decay to muons 99.9877% of the time, forming a penetrating component of the shower. Relativistic muons produced will typically reach ground level without significant deflection or decaying, despite their $2.2 \mu\text{s}$ lifetime in their rest frame. Muons with energies greater than 300 MeV have enough energy to travel over 1.5 km in ice and trigger the IceCube detector.

Neutral pions decay to high-energy gamma rays via the electromagnetic force, and can then produce electron-positron pairs via the following process:

$$\pi^0 \rightarrow \gamma + \gamma \rightarrow 2e^+ + 2e^- \quad .$$

Products of this decay can continue to produce high-energy gamma rays, either by the annihilation of positrons with atmospheric electrons, or by electron-induced bremsstrahlung. The resultant high-energy gamma rays can then begin the cycle anew through pair-production. In a simplified model, the number of particles effectively doubles every interaction length (λ) through pair-production, annihilation, and bremsstrahlung. In each generation the energy is evenly spread over the particles produced, allowing the calculation of the energy per particle as a function of depth:

$$\begin{aligned} E(X) &= \frac{E_0}{N(X)} \\ &= \frac{E_0}{2^{X/\lambda}} \quad . \end{aligned} \tag{2.2}$$

This process continues until the average energy is less than the critical energy for pair-production ($E_c = 2m_e c^2 \sim 1 \text{ MeV}$). Therefore, we can determine the depth at which the shower has reached its maximum size as:

$$X_{\max} = \lambda \log_2 \left(\frac{E_0}{E_c} \right) \quad . \tag{2.3}$$

This toy model shows important properties of shower development.

2.2 Detection Principle

As charged leptons pass through matter, they lose energy through both continuous and stochastic energy loss processes. Electrons and muons lose energy continuously via ionization, which results in Cherenkov radiation (see Section 2.2.1 for more details). There are also multiple sources of stochastic energy losses that occur, including bremsstrahlung, pair-production, and photonuclear interactions. Losses due to ionization tend to dominate below a given energy value, giving way to stochastic losses at higher energies.




Figure 2.2 Stopping power of a muon in copper as a function of muon momentum, separated into regimes based on the dominant mechanism of energy loss. Particles in the “Bethe” regime lose energy mostly through ionization. At higher energies radiative loss mechanisms (e.g. bremsstrahlung) dominate. Source: [74]

For electrons, the tipping point occurs around 1 GeV. Above this energy, bremsstrahlung is the dominant process, while below it ionization is the leading source of energy loss. Electrons lose energy at a high rate, which limits their range of propagation to a few meters of water equivalent (mwe). This limited range makes electrons only detectable at or above the surface.

Muons have a much greater ($\sim 2000\times$) mass than electrons, and as bremsstrahlung is inversely proportional to the mass squared, this means their energy losses are much less. The transition point between ionization and stochastic losses for muons is located around 1 TeV. Photon production is proportional to rate of energy loss, which is tied to the slope in Fig. 2.2. Muons with energies below 1 TeV, in the ionization-dominant region, have a roughly constant slope and thus produce similar numbers of photoelectrons. Above 1 TeV, energy losses for muons are primarily stochastic and scale with muon energy. The energy loss rate for these can be written as:

$$-\frac{dE}{dx} \approx a + bE \quad , \quad (2.4)$$

where $a \sim 0.26 \text{ GeV/mwe}$ represents the ionization term, and $b \sim 3.57 \times 10^{-4} \text{ mwe}^{-1}$ is the constant associated with stochastic losses [75]. Using this equation, the average range of a muon with initial energy E_0 can be estimated as:

$$x \approx \frac{1}{b} \ln\left(1 + \frac{b}{a} E_0\right) \quad . \quad (2.5)$$

Using this equation, a muon with an energy of 1 TeV will travel a mean distance of 2.4 km in the ice.

2.2.1 Cherenkov Radiation

When relativistic charged particles move faster than the speed of light in a medium, they produce an effect known as Cherenkov radiation. As a charged particle passes through a medium, it disturbs the orbitals of nearby electrons. The electrons relax into a stable state through the emission of photons. In this way, a pattern of emission follows the path of the muon. If the muon is moving faster than the speed of light in the given material, the

resulting wavefronts will constructively interfere as shown in Fig. 2.3, producing an amplified conical wavefront. This optical analog to a sonic boom is known as Cherenkov radiation. In a medium with index of refraction n , with a particle traveling at $\beta = v/c$, we can use this simple picture to determine the opening angle of the conical wavefront. In some time t the wavefront will travel a distance $c_n t$. In the same time, the particle will travel $vt = \beta ct$. Then the opening angle of the cone can be found using:

$$\cos \theta = \frac{c_n t}{\beta ct} = \frac{1}{n\beta} . \quad (2.6)$$

This effect occurs in air, where the average index of refraction is $n = 1.0003$ at standard temperature and pressure. However, particles must be moving faster than $v > c/n = 0.9997c$ in order for the effect to occur. In water, the average index of refraction is 1.33, meaning particles moving faster than $0.75c$ will produce Cherenkov radiation. The use of a higher index of refraction makes the detection of lesser energetic particles easier. This is why use of a detector medium is so important, and why so many are made with the abundant resource of water, or, in the case of IceCube, ice.

2.3 IceCube

The IceCube Observatory [77], located at the geographic South Pole, comprises a neutrino detector buried in the deep ice (hereon labeled IceCube) and a surface air-shower array (labeled IceTop). Shown in Fig. 2.4, the completed detector is a cubic kilometer in size, consisting of 86 vertical strings with a total of 5,160 optical sensors, called Digital Optical Modules (DOMs) buried between 1,450 and 2,450 m below the surface of the ice. The strings are separated by an average distance of 125 m, and act as support, power supplies, and data transfer mechanisms for the 60 DOMs evenly spaced along each one's length. IceCube also has eight strings that make up the DeepCore detector [78], with smaller DOM and string spacing designed to allow for neutrino studies at lower energies. The detector is proving to be remarkably stable over time, recently approaching 100% livetime.




Figure 2.3 Cherenkov radiation schematic showing the combined wavefront produced by a relativistic charged particle. Light waves produced at the numbered positions constructively interfere along the cone with opening angle θ . Adapted from [76].




Figure 2.4 Schematic of the IceCube Neutrino Observatory, complete with IceTop stations and DeepCore. A predecessor of IceCube, the Antarctic Muon And Neutrino Detector Array (AMANDA II) is also shown [79]. The instrumented volume begins 1450 m below the surface. Source: IceCube.

IceCube was built to detect neutrinos via the production of their charged lepton partners (e^\pm, μ^\pm, τ^\pm). Neutrinos interacting with the Earth — either in the atmosphere, ice, or ground — weakly interact, producing charged leptons that retain energetic and directional information. As these secondary particles pass through the ice, they induce Cherenkov radiation, which is observed in the DOMs. The same method of detection is applicable to cosmic rays: IceCube is also sensitive to downgoing relativistic muons produced by cosmic-ray air showers.

The Antarctic winter is too harsh for construction, so IceCube was built in stages during December and January, the warmest months. For the remainder of the year, IceCube took data in partial detector configurations. The seasons, from IC1 to IC86 are named based on the number of operating strings, and shown in Fig. 2.5. After the detector was completed in 2010, the following seasons are labeled by the year in which they began (IC86-2011, IC86-2012...). The analysis in this work uses six years of data from IC59 to IC86-2014.

2.3.1 Digital Optical Modules

The Digital Optical Module (DOM), shown in Fig. 2.6, is a customized device that is at the heart of the IceCube Observatory. The outside of each DOM consists of a pressure-protective glass sphere, 35 cm in diameter. The glass houses a 10" diameter R7081-02 photomultiplier tube (PMT) made by Hamamatsu photonics, as well as a power supply, light emitting diodes (LEDs) for calibration, and onboard electronics for digitization of PMT signals [80, 81]. The pressure sphere is made of $1/2$ inch-thick borosilicate glass that can withstand pressures up to 690 atm, and is filled with nitrogen gas at $\sim 1/2$ atm. Radioactive decays in the glass cause fluorescence in the ice, and are a major contributor to the dark noise rate (< 500 Hz) observed by the PMT. The photocathode is protected from the Earth's magnetic field by a mu-metal grid. Each DOM consumes approximately 3.5 W of power from the 48 V source sent down the string. The DOMs were built to last; even within their extreme environment, stress tests indicate the survival probability after 15 years is $\sim 94\%$.




Figure 2.5 Top-down view of the IceCube array. Strings deployed for each season are marked in different colors, beginning with IC-1 in yellow on the right edge of the figure. Source: IceCube.




Figure 2.6 Schematic of a Digital Optical Module (DOM) and its primary components.
Source: IceCube.

Each DOM has two types of digitizers onboard for converting the signal from the PMT. The Analog Transient Waveform Digitizer (ATWD) is a custom low-voltage system, featuring 128 switched capacitors that store the analog signal with a 300 MHz sampling rate. Each ATWD has four channels: three for input with different gains at a ratio of 16:2:1/4, and one for digitization. The digitization process takes $29\ \mu\text{s}$, so each DOM has two ATWDs that alternate operation to minimize dead time due to signal readout. The second type of digitizer is the Fast Analog to Digital Converter (FADC). FADCs are commercial digitizers with a 40 MHz sampling rate. In comparison to ATWDs, they have no dead time, providing valuable coverage for long signals, but they do have a limited dynamic range and wide timing resolution.

A DOM's typical response to a single photoelectron (SPE) is an electric pulse with an amplitude of 10 mV, width of 5 ns, and a time resolution of ~ 2 ns. Digitization is started by the Field-Programmable Gate Array (FPGA) when a signal exceeds the discriminator threshold of $\frac{1}{4}$ SPE. A full "DOM launch" consists of waveforms from both the FADC and ATWD, as well as a timestamp. Onboard storage space and communication rate with the surface, however, limit the data rate, so IceCube runs in two kinds of coincidence modes. Hard Local Coincidences (HLCs) require two of the four nearest neighboring DOMs to be hit within a $1\ \mu\text{s}$ coincidence window. HLCs occur at a rate between 3-15 Hz per DOM, and cause the ATWD to be read out. By contrast, Soft Local Coincidences (SLCs) have a rate of around 350 Hz, and only three FADC samples are read out for each. Most SLC hits are noise — note the closeness in rates — but they can occasionally provide valuable missing event information.

Data packets with timestamped HLC and SLC information are sent from the DOMs up strings to computers in the IceCube Lab, near the center of the detector on the surface. These packets are then ordered in time and checked against trigger conditions. The primary trigger for IceCube is Simple Majority Trigger 8 (SMT8), which requires coincident hits in 8 DOMs within $5\ \mu\text{s}$. If this trigger is passed, all locally-coincident hits within a $\pm 10\ \mu\text{s}$ window are recorded, with overlapping time windows merged. The SMT8 trigger has an

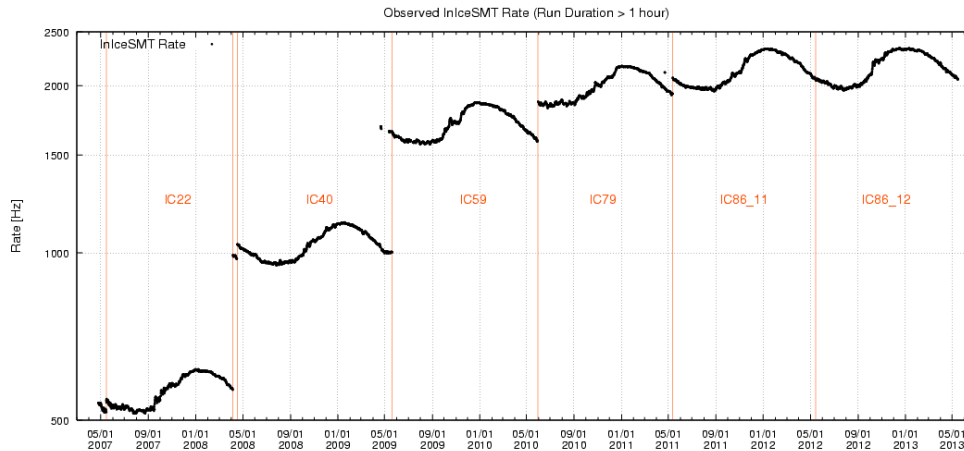


Figure 2.7 Rate over time for various detector configurations. Large discontinuities in rate occur due to the activation of new IceCube strings, and therefore an increase in detector effective area. Variations within a given detector configuration are due to seasonal effects such as average temperature and pressure.

average rate of ~ 2700 Hz for the completed detector, with a seasonal variation of $\pm 10\%$. Figure 2.7 shows the seasonal variation in rate as a function of time, with the jumps in rate corresponding to the beginning of a new season with more active strings. Data packets that pass the SMT8 or other trigger conditions are combined into events that are then sent on to reconstruction and filtering.

2.3.2 Ice Properties

The ice at the South Pole is not a uniform medium: it has been built up layer by layer over 100,000 years. Understanding the properties of the ice is extremely important for IceCube, as it directly affects the accuracy of event reconstructions. The ice properties are described in terms of the absorption and scattering of ultraviolet light, the region where Cherenkov radiation peaks.

Starting at the surface is a 50 m deep layer of compacted snow, called “firn.” The solid ice begins below the firn, but it does not immediately reach the transparency levels desired. Air bubbles embedded in the ice strongly scatter light, leading to a mean free path of ~ 10 m. The scattering from these air bubbles is the dominant factor until ~ 1300 m down. At depths

near IceCube, the air bubbles are compressed to the point that they become air hydrate crystals, with each molecule of air surrounded by molecules of ice. The index of refraction of air hydrates is nearly identical to that of the ice itself, making the air bubbles no longer a source of scattering. At IceCube depths, the properties of the ice can be tested by using LEDs located on the DOMs. In a “flasher run,” light is emitted in a pattern from one DOM. Ice properties can be determined by using the arrival time distribution (see [82] for more). Results from these tests indicate that from 1,450 – 2,450 m, the average scattering length goes from 25 to 70 m, and the absorption length increases from 70 to 200 m.

With air bubbles no longer a major factor, dust becomes the dominant source of scattering and absorption at these depths. The dust occurs in layers; several dust peaks are shown in Fig. 2.8. The strongest of these dust peaks, known simply as the “dust layer,” is approximately 65,000 years old, corresponding with periods of lower temperature in the late Pleistocene. The dust layer is 2,025 m down, and has an absorption length of < 20 m and a scattering length of 5 m. Dust layers in general are parallel to the surface, but towards the bottom of the detector they begin to exhibit a tilt believed to reflect the shape of the bedrock beneath the ice. The ice model used in reconstructions accounts for the fact that the absorption and scattering are not just a function of depth.

2.3.3 Data Storage and Transfer

IceCube’s trigger rate of 2.6 kHz and limited bandwidth available from the South Pole necessitates a method of storing essential data in a compact form. Events that pass a series of filters have their full information stored, but for events like cosmic-ray air showers that dominate the trigger rate, this is not possible. The Data Storage and Transfer (DST) format was created with this in mind.

For every event that triggers the detector, the DST format stores the time, number of triggered DOMs, the angular information from two directional reconstructions, and some other basic parameters. The arrival direction is first estimated using a χ^2 linear-track fit to the DOM hits. Then, using this estimate as a seed, a more complex likelihood-based




Figure 2.8 Peaks in absorption due to dust as a function of depth for a variety of IceCube strings [83]. The average absorption decreases as a function of depth, but various peaks due to dust layers, most notably at 2000 m, are visible. The tilt of the dust layers can also be seen in this plot as the dust features show a dependence on string position.

reconstruction is applied, accounting for aspects of light generation and propagation in the ice. For more information on the angular reconstruction procedures, refer to Section 3.1.1. The directional fits are stored as a pixel number in a HEALpix map with $n_{\text{side}} = 1024$ (see Section 4.3.1), so that both zenith and azimuthal information can be stored in a single number. The data are then compressed and transferred via the South Pole Archival and Data Exchange (SPADE) satellite communication system, which allows about 3 Gb of data per day.

Studies of cosmic-ray anisotropy using IceCube primarily rely on the DST dataset. Simulation studies with the CORSIKA Monte Carlo program¹ [84] indicate a median energy of 20 TeV and a median angular resolution of 3° for the likelihood-based fit.

2.4 IceTop

IceTop is the surface array component of the IceCube Observatory. It consists of 81 surface stations spread over a square kilometer. Each station is located near an IceCube string, so the stations have an average spacing of 125 m in a hexagonal grid. Also similar to IceCube, IceTop was built in phases, and the naming convention reflects the number of stations active during that season. IceTop is described in detail in [85]. The data used in this analysis was taken between May 2009 (IT59) and May 2015 (IT81-2014).

IceTop stations consist of two light-tight tanks, each one 1.8 m in diameter, 1.3 m in height, and filled with transparent ice up to a height of 0.9 m. Each tank hosts two DOMs, identical to the ones used in IceCube (see Section 2.3.1). Figure 2.9 shows an IceTop station during installation (*left*), and the placement of the DOMs inside an IceTop tank (*right*). The DOMs are operated at different gains for an increased dynamic range, with a high-gain DOM at 5×10^6 and a low-gain at 10^5 . IceTop detects showers through the same mechanism as IceCube, but its surface location near the shower maximum, shown in Fig. 2.10, makes it sensitive to the full electromagnetic component of the shower, not just the muonic component.

¹COsmic Ray SIMulations for KAscade: <http://www-ik.fzk.de/corsika/>




Figure 2.9 An IceTop station during installation (*left*), and DOM setup inside an IceTop tank (*right*). Source: IceCube.

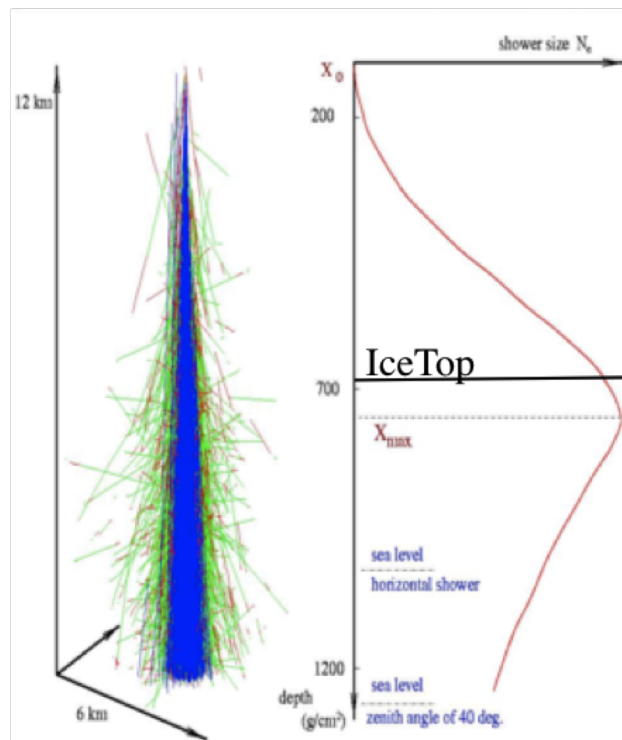


Figure 2.10 Simulated shower development (*left*) and size of shower as a function of atmospheric depth (*right*). IceTop's location near shower maximum is marked.

IceTop has a base trigger rate of 1600 Hz per DOM. To reduce this, the high gain DOMs in each station are operated in Hard Local Coincidence (HLC) mode. Data from the DOMs is read out if both are hit within $\pm 1\mu\text{s}$. The HLC requirement reduces the trigger rate from 1600 Hz per DOM to 3000 Hz for the whole detector. To reduce the trigger rate further, a simple majority trigger (SMT) similar to that of IceCube is applied. The SMT trigger requires at least 6 HLC hits within a time window of $5\mu\text{s}$. As there can be two HLC hits per station, the trigger requires at least 3 stations to be hit. The SMT trigger brings the rate down to 30 Hz.

Tanks in IceTop record signal in units of vertical equivalent muons (VEMs). One VEM corresponds to the charge produced by a single muon entering the tank vertically, and is calculated on a tank-by-tank basis. Since 2009, this calibration has been performed with a single-muon calibration trigger to avoid downtime for separate calibration runs. Figure 2.11 shows the calibration plot for a single tank, with the clear peak from VEMs.

Even with the reduced rate from the SMT trigger, there is limited bandwidth for data transfer from the South Pole via satellite. In IceTop, data is filtered according to the number of stations triggered. The filters used are of the form `IceTopSTAX[_InIceSMT]_YY`, where `X` indicates the minimum number of stations hit, `_InIceSMT` indicates that it also passed the in-ice SMT, and `YY` is the year of operation. Each filter is given a prescale that determines what fraction of events that pass that filter are kept. A prescale of 3 indicates that $1/3$ of the events that pass that filter are kept. Table 2.1 shows the filters and corresponding prescales used in this analysis.




Figure 2.11 Histogram of charge deposited in photoelectrons for a single IceTop DOM taken from experimental data [85]. The curves for the feature from VEMs and the background are fit with a muon fit function for high-gain DOMs as determined from simulation.

Filter	Prescale
IceTopSTA3_09	8
IceTopSTA3_InIceSMT_09	3
IceTopSTA8_09	1
IceTopSTA3_[10-11]	3
IceTopSTA3_InIceSMT_[10-11]	2
IceTopSTA8_[10-11]	1
IceTopSTA3_[12-14]	10
IceTopSTA8_[12-14]	1

Table 2.1 Names and prescales for each filter used in this analysis

Chapter 3

Event Reconstruction

In this chapter we discuss angular reconstruction and energy estimation methods used in both IceCube and IceTop. Special attention is paid to **ShowerLLH**: a project developed for this analysis that reconstructs the most likely cosmic-ray primary core position and energy in IceTop for a variety of composition assumptions.

3.1 IceCube

3.1.1 Angular Reconstruction

In IceCube, the timing and location of DOM hits are used to reconstruct the arrival direction of the muon(s) passing through the ice, and therefore obtain an estimate for the direction of the primary particle.

As a first guess, the event is treated as a plane wave moving through the ice. The basic reconstruction algorithm `linefit` performs a linear track fit by minimizing the sum of the squares of the distances between the track and the hits, ignoring the conical shape of the Cherenkov wavefront (see Sec. 2.2.1) and the scattering and absorption in the ice. The chi-square distribution to be minimized can be written as:

$$\chi^2 = \sum_{i=1}^N (\vec{x}_0 - \vec{x}_i + \vec{v}(t_i - t_0))^2 \quad , \quad (3.1)$$

where \vec{x}_i and t_i represent the position and time of the i^{th} hit, and the particle passes through \vec{x}_0 at time t_0 with velocity \vec{v} . This algorithm is very fast, and results in an angular resolution with an accuracy of a few degrees, which makes it a good seed for more complex reconstructions.

More advanced reconstructions attempt to use more of the information available, notably the shape of the Cherenkov cone and the properties of the ice, to get a more accurate result. Starting with $p(\vec{x}_i, t_i|\vec{a})$ as the probability of observing a hit at position \vec{x}_i and time t_i given a muon track with parameters \vec{a} , likelihood-based reconstructions attempt to maximize

$$\log \mathcal{L} = \sum_{i=1}^N p(\vec{x}_i, t_i|\vec{a}) . \quad (3.2)$$

The probability density function for p is given by the Pandel parameterization [86], an analytic expression that describes the likelihood of a photon traveling some distance d with a given time residual t_{res} using a gamma function:

$$p(t_{res}, d) = \frac{1}{\Gamma(a)ba} x^{a-1} e^{-t_{res}/b} \quad (3.3)$$

$$a = \frac{d}{\lambda}, \quad b = \frac{1}{\tau^{-1} + c/(nX_0)} , \quad (3.4)$$

where X_0 represents the absorption length of light in ice, and λ and τ are characteristic length and time scales, respectively, treated as free parameters. The time residual represents the difference between the arrival time of a given photon and a “direct hit” photon that travels directly from the source to the DOM. This parameterization accounts for the index of refraction n , as well as the scattering properties of the ice via the τ , λ , and X_0 parameters, which are all determined by simulation [87]. The advantage of the Pandel parameterization is that it reasonably approximates a complex solution at a very high speed.

There is no such analytic solution or approximation for the likelihood value, so numerical minimization algorithms such as MINUIT are used to search for the minimum of $-\log \mathcal{L}$. To avoid settling in local minima, the search is performed iteratively with a variety of starting conditions. For each event, anywhere from one to thirty-two pseudo-random seed directions

are chosen from a Sobol sequence in order to increase the probability of finding a global minimum.

A simple likelihood-based reconstruction uses a single photoelectron (SPE) approach to muon track reconstruction [88]. The SPE reconstruction only considers the first photon to arrive at each DOM in its likelihood calculation, ignoring DOMs with no hits. This method improves upon the `linefit` reconstruction it uses as a directional seed. SPE is also fast, allowing a single iteration to be calculated in almost real-time and recorded in the DST dataset. All IceCube cosmic-ray event reconstructions used in this analysis are calculated using the SPE algorithm. Figure 3.1 shows a particle track through IceCube, with the red line indicating the SPE reconstruction.

More advanced reconstruction algorithms — such as the multiple photoelectron (MPE) fit that uses the timing information of all hits instead of just the first one [89] — offer better performance, but are slower. The ability of the SPE reconstruction to run in near-real time, coupled with its reasonable angular resolution, make it an excellent choice for anisotropy analysis.

3.1.2 Energy Estimation

In order to study the anisotropy in cosmic-ray arrival direction as a function of primary energy in IceCube, a similar energy estimation procedure to that of [56] is used. Events are classified using the number of DOMs that collected Cherenkov light, N_{channel} , and the reconstructed zenith angle, θ_{reco} . N_{channel} is used as an energy estimator of the muons detected by IceCube. The reconstructed angle θ_{reco} is considered because at larger zenith angles muons, and therefore the primary cosmic-ray particles, must have higher energy in order to reach and trigger the buried IceCube experiment. Simulation data are used to determine bands in primary particle energy as a function of N_{channel} and the cosine of θ_{reco} , as shown in Figure 3.2. The figure shows that for a given muon energy, events at larger zenith angles are produced by cosmic-ray particles with higher energy.




Figure 3.1 Sample response of IceCube to a muon passing through the detector. The size of each DOM hit indicates the number of photons detected, and the color indicates the arrival time, from early (red) to late (blue). The red line indicates the `linefit` directional reconstruction.




Figure 3.2 Median true energy as a function of the cosine of the reconstructed zenith angle θ_{reco} and the number of DOMs hit in the event, N_{channel} , from simulation.

The treatment of N_{channel} as an energy estimator is dependent on the detector configuration. In addition, procedures such as event splitting have changed over time. To account for these changes, the above process is repeated with simulation for each detector configuration. A B-spline function in N_{channel} and $\cos(\theta_{\text{reco}})$ is used to fit the data in Fig. 3.2 in order to reduce the errors due to limited simulation data statistics. Events are then separated into nine statistically independent data samples, and assigned the mean primary particle energy of the corresponding energy band. The resolution of such a reconstructed energy depends on the detector configuration and energy band, but is on the order of $0.5 \log_{10}(E/\text{GeV})$. This energy resolution is primarily limited by the relatively large fluctuations in the fraction of the total shower energy that is transferred to the muon or muon bundle.

3.2 IceTop

3.2.1 Angular Reconstruction

In IceTop, the timing, location, and the charge deposited at each DOM are all used to reconstruct the arrival direction of the cosmic-ray primary.

Similar to IceCube, the most basic fitter treats the shower front as a plane wave, ignoring the curvature in favor of speed. This reconstruction, called **ShowerPlane**, begins by setting the origin position at the shower's center of gravity: the signal-weighted average of the tank positions. With the shower core set, the time at which a tank is hit as a function of position \vec{x} can be written as

$$t(\vec{x}) = t_0 - \frac{1}{c}(\vec{x} - \vec{x}_0) \cdot \vec{s} \quad , \quad (3.5)$$

where \vec{x}_0 , t_0 are the position of the origin and the time the shower passes through it, respectively, and \vec{s} is the unit vector normal to the shower plane. The goal of the fitter is to minimize the square of the weighted time difference

$$\chi^2 = \sum_i w_i (t_i - t(\vec{x}_i))^2 \quad . \quad (3.6)$$

A constant weight of $w_i = 1/\sigma^2$ is applied, where $\sigma = 5$ ns is the uncertainty in the measured trigger time. Through expansion, the equation becomes

$$\chi^2 = \frac{1}{\sigma^2} \sum_i [t_i - t_0 + (d_{i,x}s_x + d_{i,y}s_y + d_{i,z}s_z)]^2 . \quad (3.7)$$

Since \vec{s} is a unit vector with $\vec{s} = (s_x, s_y, \sqrt{1 - s_x^2 - s_y^2})$, this equation is non-linear. To approximate the solution, `ShowerPlane` is solved in two iterations. In the first, z is set to a fixed value $\langle z \rangle$, an approximation that is valid since the maximum variation in height across IceTop tanks is small (~ 5 m). With fixed z , the program minimizes the equation:

$$\chi^2 = \frac{1}{\sigma^2} \sum_i [t_i - t_0 + d_{i,x}s_x + d_{i,y}s_y]^2 . \quad (3.8)$$

In the second iteration, `ShowerPlane` corrects for tank heights by adjusting the hit times at each tank correcting for the tank heights,

$$t'_i = t_i + \frac{z_i - \langle z \rangle}{c} \cos \theta , \quad (3.9)$$

where θ is the reconstructed zenith angle from the first iteration. The steps from the first iteration are then repeated using the adjusted times to return a final direction.

This reconstruction results in a median angular resolution of 3° , as calculated using simulation, similar to the angular resolution of the likelihood fit stored in the DST dataset. More detailed angular reconstructions exist — notably `laputop`, a likelihood-based algorithm used in IceTop energy spectra analysis — but these require quality cuts that markedly reduce statistics. The speed, reasonable angular resolution, and low requirements make the `ShowerPlane` reconstruction a good choice for anisotropy studies.

3.2.2 Likelihood-based Reconstruction (ShowerLLH)

A major part of this thesis was the development of a likelihood-based reconstruction for cosmic-ray air showers detected by IceTop. As discussed, IceCube has large statistics, but is limited in two ways. First, the high trigger rate makes the use of the DST format necessary, limiting the information retained per event. Second, IceCube is only sensitive to the muonic

components of air showers. As described in Sec. 3.1.2, there is a significant amount of variation in the fraction of the primary energy that is transferred to the muonic core, placing an upper limit on our energy resolution. IceTop has a much lower rate, but is able to store more information per event, and it sees the entire electromagnetic component of the shower. This allows in theory for superior energy resolution. ShowerLogLikelihood (**ShoweLLH**) was designed to take into account the additional information available — specifically the snow depth and charge at each tank — and return the most likely core position and primary energy, with the special function of being able to do so for a variety of composition assumptions.

3.2.2.1 Method

The **ShoweLLH** reconstruction begins with the building of likelihood tables from simulation. For every active tank in every shower, the appropriate bin in a 4-dimensional histogram (primary energy (E) \times snow depth at tank (S) \times distance from shower core (D) \times charge deposited (C)) is filled. Tanks which are active but not hit are included. Saturated DOMs are placed in the maximum charge bin. The completed tables are then normalized so that each bin ($E \times S \times D \times C$) gives the probability that a primary particle with energy E deposited a charge C at a tank D meters away with S meters of snow on top of it. These tables are built up using single-composition simulations, effectively adding a fifth dimension of composition.

To use these tables for reconstruction, a kind of brute-force method is applied. For each shower we have two measurable quantities: snow depth and charge deposited at each tank. Using these, the goal is to return the most likely composition, core position, and primary energy for the shower. The composition is a special case, and will be treated separately, so for now we will focus on the core position and primary energy. Given a shower from data, the program begins by testing a series of preset coarse grid positions across the array. For each test position, the reconstruction calculates the log-likelihood $\log \mathcal{L}$ that a shower of energy E deposited the observed charges (C_i) at each tank by summing the appropriate bins in the

log-probability table P :

$$\log \mathcal{L}(E) = \sum_{i=1}^{n_{\text{tanks}}} P(E, S_i, D_i, C_i) \quad . \quad (3.10)$$

This process is repeated, returning a likelihood value for every energy bin at the tested core position. By comparing the highest likelihood value for each position, we can find the most likely core position in our coarse grid. This entire process is repeated twice more, each time with a finer grid centered on the previous most likely position. The grid steps are initially 125 m apart, then 20 m, then 5 m. At the end of the process, the most likely core position, primary energy, and likelihood value are returned.

Recall that the likelihood tables are made for single compositions. By repeating the above method for a variety of compositions, ShowerLLH can return the most likely core position and primary energy for a variety of composition assumptions. Because it also returns a likelihood value, it effectively returns a most likely composition as well.

As mentioned in Sec. 3.1.2, air showers are affected by the amount of atmosphere they must pass through, leading to a zenith-dependence in the energy estimation. While capable of including zenith bins in the likelihood tables, the current iteration of ShowerLLH does not do so, as it results in zenith bin migration, affecting our ability to produce skymaps. In place of zenith bins, simulation is used to determine the median energy bias and resolution of ShowerLLH as a function of reconstructed zenith angle. A B-spline function is fit to the result and then used to correct the reconstructed energy values. Section 3.2.2.2 provides more information on the effect of this fit.

The accuracy of ShowerLLH is strongly influenced by the simulation used to build the likelihood tables. Large amounts of simulation are needed to properly fill the 4-dimensional likelihood tables, but the tables are mostly independent of detector configuration. Since more simulation is available for IT73 than IT81, the IT73 simulation sets shown in Table 3.1 were used to make the likelihood tables. The tests performed in Section 3.2.2.2 were run over low-energy simulation in addition to determine the effect of showers outside of the reconstructed energy range. The low-energy simulation is shown in Table 3.2.

ID	Composition	Energy Range [$\log_{10}(E/\text{GeV})$]	Number of Events
7006	H	5.0 – 8.0	3×10^6
7579	H	7.0 – 9.5	1.2×10^6
7241	He	5.0 – 8.0	1×10^6
7263	He	5.0 – 8.0	2×10^6
7791	He	7.0 – 9.5	1.2×10^6
7242	O	5.0 – 8.0	1×10^6
7262	O	5.0 – 8.0	2×10^6
7851	O	7.0 – 9.5	1.2×10^6
7007	Fe	5.0 – 8.0	3×10^6
7784	Fe	7.0 – 9.5	1.2×10^6

Table 3.1 Information on simulation used to build the likelihood tables for this analysis. All of the above simulation datasets throw showers with zenith angles ranging from $0 - 40^\circ$ with an E^{-2} spectrum.

ID	Composition	Energy Range [$\log_{10}(E/\text{GeV})$]	Number of Events
7351	H	4.0 – 6.0	5×10^5
7483	He	4.0 – 6.0	5×10^5
7486	O	4.0 – 6.0	5×10^5
7394	Fe	4.0 – 6.0	5×10^5

Table 3.2 Information on low-energy simulation used to test the reconstruction of low energy events. These simulation datasets were not used in the creation of the likelihood tables.

3.2.2.2 Performance

Before applying `ShowerLLH` to data, a set of quality cuts is developed. Figure 3.3 shows the effective area for the IT73 configuration as a function of true energy with no cuts applied. The goal of the cuts is to improve the energy resolution, and for the detector effective area to reach and maintain a peak value as a function of energy.

The cuts used are adapted from previous work [90]. The cut names and percentage of events cut are shown in Table 3.3. The first cut just insists that `ShowerLLH` ran successfully, effectively ensuring that the plane fit succeeded. The second cut selects events with $\cos(\text{zenith}) \geq 0.8$. The simulation sets used were designed for IceCube-IceTop coincidence analysis, and throw showers in the zenith range of $0^\circ - 40^\circ$. The zenith cut was chosen close to, but not exceeding, the 40° simulation limit. Note that, in theory, `ShowerLLH` could reconstruct showers well beyond this limit, given adequate simulation. The third cut requires the reconstructed position to be inside the detector. The grid search method used in `ShowerLLH` allows for the reconstruction of showers outside the array, but accuracy is greatly improved when reconstructing contained showers. The fourth cut requires that the loudest station is not on the edge of the detector. Along with the third cut, this serves to restrict the sample to contained showers. The fifth and final cut requires the maximum charge deposited to be at least 6 VEM. This helps reduce contamination of the high-energy sample from low-energy showers.

Figure 3.4 shows the effective area of the detector (in the IT73 configuration) with the quality cuts applied. The best-fit line is shown beginning where the curve flattens at about $6.2 \log(E/\text{GeV})$. All following performance-related plots were created using these quality cuts, and begin at the point where the effective area becomes flat.

As mentioned in Section 3.2.2.1, the reconstruction does not account for the zenith-dependence. Figure 3.5 shows a consistent underestimation of the energy for higher reconstructed zenith angles. To adjust, the energy resolution as a function of zenith is fit with a B-spline function. Using this function, we can apply an energy correction as a function of zenith for each event.




Figure 3.3 Effective area of IT73 simulation as a function of true energy with no quality cuts applied.

No.	Cut Name	Cut Percentage	Cumulative Percentage
1.	Reconstruction Passed	88%	88%
2.	Zenith Cut	82%	70%
3.	Containment Cut	75%	61%
4.	Loudest Station not on Edge	71%	50%
5.	Max Charge Cut	50%	27%

Table 3.3 Quality cuts applied to ShowerLLH. Each cut is listed by name, with the percentage of events that pass the cut (“Cut Percentage”), and the percentage of events that pass that cut and all others before it (“Cumulative Percentage”).




Figure 3.4 Effective area of IT73 simulation as a function of true energy with all quality cuts applied. The solid line represents a flat, linear fit to the data above $6.2 \log(E/\text{GeV})$.




Figure 3.5 Energy resolution as a function of reconstructed angle ($\cos(\theta_{\text{reco}})$). Points indicate the median value, with error bars showing 68% containment. The solid line indicates the best-fit B-spline function. The energy resolution is defined as the difference between the logarithms of the reconstructed energy (E_{LLH}) and the true energy (E_{true}).

With the cuts and zenith correction applied, we can now study the performance of **ShowerLLH**. Figure 3.6 shows the probability that a simulated cosmic-ray primary with a energy E_{LLH} is reconstructed with energy E_{reco} . The distribution is peaked along the $E_{\text{LLH}} = E_{\text{true}}$ line. There are some high-energy events – whose cores are typically located outside the detector – that are reconstructed as contained, low-energy showers. This mis-reconstruction is far less damaging, however, than the opposite case where low-energy events contaminate the high-energy sample. High-energy events have a much lower flux than low-energy, so even a small percentage of mis-reconstructed low-energy events could dominate the high-energy bins.

With the shape of the distribution in Fig. 3.6 in mind, we can look at the energy resolution as a function of various parameters. Figures 3.7a, 3.7b, and 3.7c show the median energy resolution as a function of true energy, zenith angle, and distance from the center of the array, respectively. In each of these we can see the success of the reconstruction: with the exception of the few outliers in each case, the reconstructed energy is typically within a bin’s distance ($0.05 \log_{10}(E/\text{GeV})$) of the true energy. In Fig. 3.7b, the beneficial effect of the zenith correction is evident (compare to Fig. 3.5).

In addition to most-likely energy and composition values, **ShowerLLH** also returns a most-likely core position, found using the iterative grid-search method. Figure 3.7d shows the distance between the reconstructed and true core positions as a function of true energy. The finest grid search was performed with a grid spacing of 5m, and the resulting core resolution is primarily $< 10\text{m}$.

3.2.2.3 Composition Sensitivity

As mentioned in Sec. 3.2.2.1, the likelihood values returned by **ShowerLLH** can be used to determine a most likely composition for each event. The power for composition differentiation lies in the lateral distribution of the shower, which is discretized in the charge and distance bins in the likelihood tables. Figure 3.8 shows the distribution of the difference between the iron and proton likelihood values for proton (*blue*) and iron (*red*) simulation, for a given bin




Figure 3.6 Probability (log-scale) that a cosmic-ray primary with energy E_{true} is reconstructed with energy E_{LLH} . This plot includes simulated events with energies below ShowerLLH’s low-energy threshold for reconstruction. None of the low energy events are reconstructed with $\log_{10}(E/\text{GeV}) > 6.2$, the value where the effective area becomes flat. Mis-reconstructions primarily occur in the under-estimation of high-energy events.

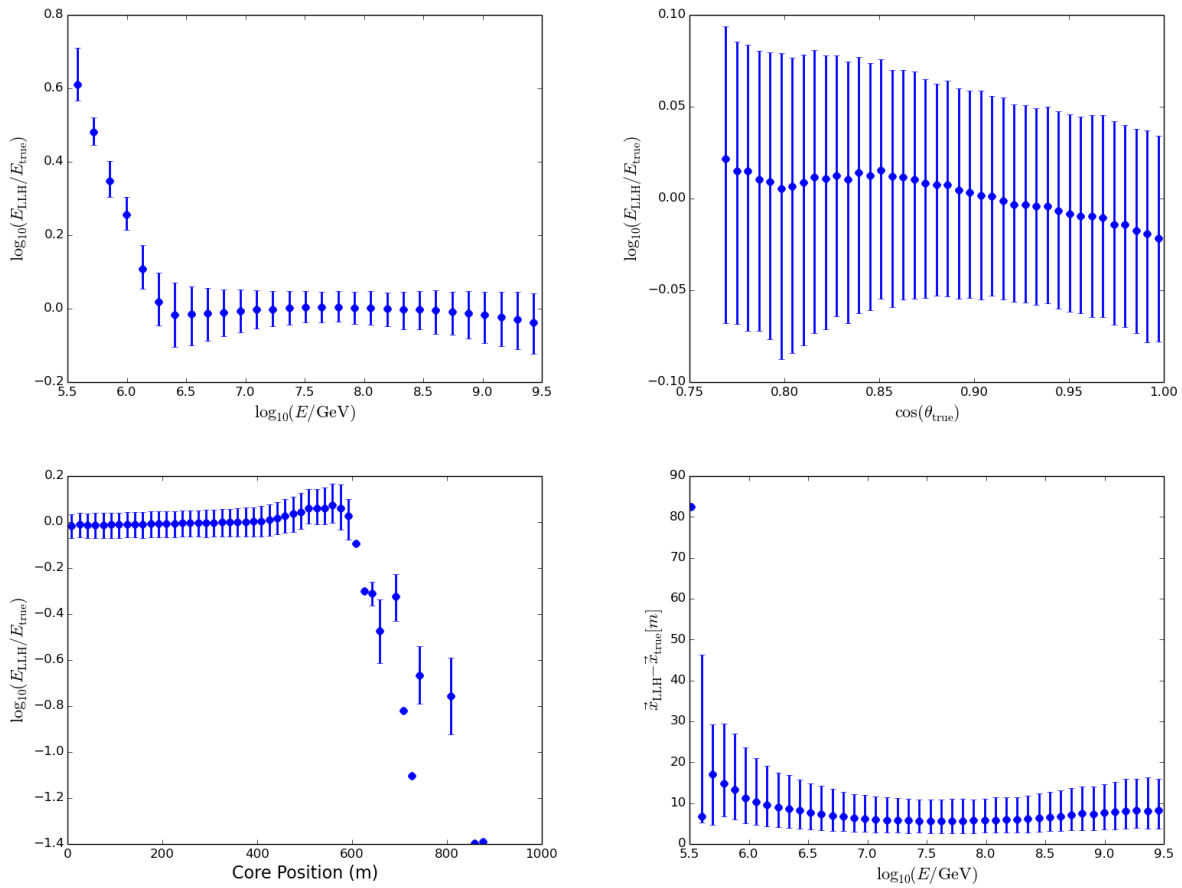


Figure 3.7 ShowerLLH performance with quality cuts applied. The energy resolution is shown as a function of (a) true energy, (b) true zenith angle, and (c) distance from the center of the array. The core resolution is also shown (d) as a function of true energy.

Points indicate the median value with error bars containing 68% of events.

in true energy. This plot illustrates that the most-likely composition returned by **ShowerLLH** does not itself have much value; the iron distribution peaks at 0, indicating iron showers at these energies are equally likely to be reconstructed as protons. The difference between the iron and proton likelihoods, however, has some composition-separation power. The Bayes factor, described in greater detail in Sec. 4.5, is a useful method for determining whether two distributions arise from the same or separate sources. Values of $B_{21} > 1000$ are considered strong evidence against a single-source hypothesis. Figure 3.9 shows the natural log of the Bayes factor calculated by comparing a pure proton sample to one with some percentage of iron, with the line on the plot indicating a Bayes factor of 1000. The calculated Bayes factor reaches the target value at samples composed of approximately 23% iron. The calculation of the Bayes factor is highly dependent on the number of events, and these values were found using $\sim 16,500$ events from simulation, compared to the median bin size of $\sim 37,000$ events in the data sample used for this part of the analysis. This method is not currently powerful enough to do a detailed analysis of composition near the knee, but when applied to large-statistics datasets can provide some composition separation power.




Figure 3.8 Distributions of the difference between iron and proton likelihood values for simulated events with true energy between 6.5 and 6.75 $\log_{10}(E/\text{GeV})$, from simulation.

The blue line corresponds to the distribution for proton simulation, and the red corresponds to iron. The separation of peaks indicates the power of **ShowerLLH** to do composition separation.




Figure 3.9 Natural log of the Bayes factor, calculated by comparing a pure proton sample to one with some percentage of iron. The green lines indicates a Bayes factor of 1000. All simulation was weighted to an $E^{-2.7}$ spectrum and had quality cuts applied. Randomly selected events from an iron distribution were added multiple times. Points indicate the median of the resulting distributions for each iron percentage, and error bars 68% containment.

Chapter 4

Analysis Method

This chapter details the methods used to study the cosmic-ray anisotropy. First, we provide an introduction to the coordinate systems used in Sec. 4.1. Next, the time-scrambling method used to estimate an isotropic background from data and methods of subtracting large-scale structure are detailed in Sec. 4.2. The binning of the sky and calculation of relative intensity and significance are covered in Sec. 4.3. Then, the calculation of the power spectrum from the sky maps is described in Sec. 4.4. Finally, Sec. 4.5 describes the Bayes factor, a useful method of determining whether two distributions share the same source by comparing their shapes.

4.1 Coordinate System

The research presented in this thesis relies heavily on the production and analysis of skymaps. This section focuses on the definition of the coordinate systems used, and the transformations between them.

4.1.1 Local Coordinates

Events that trigger the IceCube detector are reconstructed in the local coordinate system. Local coordinates are fixed to the detector, with the origin near the center of the detector. Figure 4.1 shows the local coordinate system with some key reference points. Zenith is defined as directly overhead. The directions of the International Reference Meridian (sometimes known as the Prime Meridian) and International Date Line are used as grid directions of

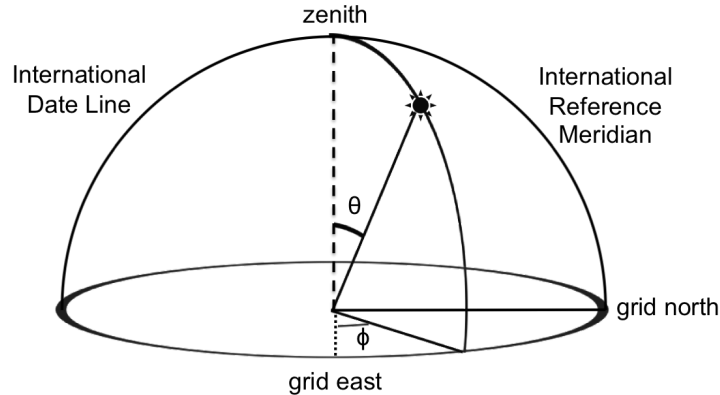


Figure 4.1 Local coordinate system with zenith angle θ and azimuth angle ϕ shown for a sample event. Zenith angle is calculated from directly overhead, and azimuth counterclockwise from grid east. Grid north points in the same direction as the International Reference Meridian.

north and south, respectively. The arrival directions of events are defined in this system using two angles. The zenith angle (θ) is measured from the zenith. The azimuth angle (ϕ) is measured counter-clockwise from grid east. For reference, in this coordinate system, $\theta = 0^\circ$ is directly overhead, $\theta = 90^\circ$ is the horizon, $\phi = 90^\circ$ is the International Reference Meridian, and $\phi = 270^\circ$ the International Date Line.

In local coordinates, the effective area of the detector is not constant as a function of azimuth, due to lanes in the hexagonal lattice. Figure 4.2 shows the event distributions in zenith and azimuth for one day of data.

4.1.2 Equatorial Coordinates

Equatorial coordinates are used to define the locations of distant objects on the celestial sphere. Referred to in this work as the sidereal coordinate frame, equatorial coordinates are a fixed-star coordinate system, so the locations of distant objects do not change over time. The celestial sphere is defined by the projection of Earth's features outward to infinity, namely the equator and poles. The plane in which the Earth orbits around the sun is known as the ecliptic plane, and is separated from the celestial equator by the tilt of the Earth's




Figure 4.2 Local acceptance of six years of IceCube data in zenith (*left*) and azimuth (*right*). Events with zenith angles above 90° are mis-reconstructed and therefore not expected to follow the same distribution. Note the preferred lanes for detection in azimuth that result from the hexagonal lattice shape of the detector.

axis. The intersection of the ecliptic plane and the celestial equator is defined as the vernal point.

Similar to local coordinates, equatorial coordinates use two angles to assign direction, as shown in Fig. 4.3. Declination (δ) is measured from the celestial equator to the source, with values from -90° at the south celestial pole to $+90^\circ$ at the north. Right ascension (α) is measured from the vernal point counter-clockwise along the celestial equator. The equatorial coordinate system treats the tilt of the Earth's axis as a constant value, when in reality the Earth precesses. Fortunately, this change is very slow (precession occurs with a period of $\sim 26,000$ years), so over epochs of 50 years the tilt is effectively constant. In this work, the J2000 epoch is used.

Equatorial coordinates also come with their own time frame, known as sidereal time. One sidereal day is the time it takes for the Earth to rotate 360° . In the time it takes the Earth to perform one rotation, it has also moved in its orbit, so it must rotate slightly more for the sun to return to the same local position. The result is that the solar day of 24 hours is 4 minutes longer than the sidereal day.

In a fixed-star coordinate system, the detector rotates throughout the day. The lanes in effective area found in the local coordinate system are averaged out, leading to a smooth detector acceptance in right ascension, as seen in Fig. 4.4.

4.1.3 Coordinate Transformation

For every reconstructed event, the directional information is stored in local coordinates (θ , ϕ), and the arrival time in Modified Julian Date (MJD). The Julian Date (JD) has been long used by astronomers, and represents the number of days since the beginning of the Julian period on 1/1/4714 BCE. The MJD was introduced in 1957, making a half-day adjustment so each day begins at midnight Greenwich Mean Time:

$$\text{MJD} = \text{JD} - 2400000.5 \quad . \quad (4.1)$$

The coordinate transformation for the IceCube detector from local to equatorial coordinates is made easy by virtue of its location at the South Pole. All conversions are written in




Figure 4.3 Equatorial coordinate system, showing sample lines of constant declination and right ascension. The plane formed by the celestial equator is tilted at an angle with respect to the ecliptic plane due to the tilt of the Earth. Credit: [91].

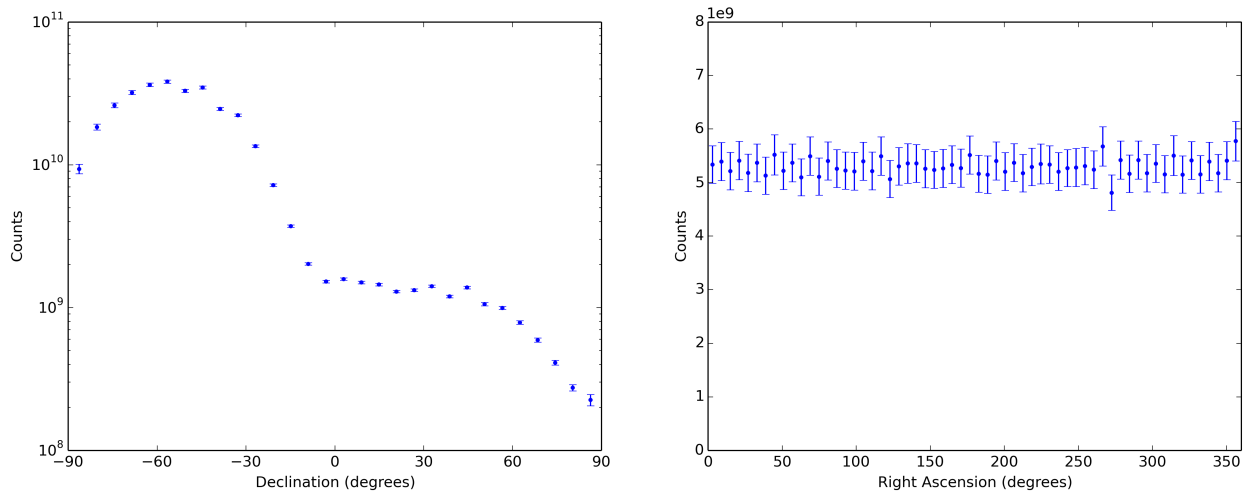


Figure 4.4 Equatorial acceptance of six years of IceCube data in declination (*left*) and right ascension (*right*). Events with declination angles above 0° are mis-reconstructed and therefore not expected to follow the same distribution. The preferred lanes of acceptance present in local coordinates are averaged out over time.

degrees, unless expressly noted otherwise. Conversion of zenith angle to declination is given by $\delta = \theta - 90^\circ$. Calculating right ascension from a given azimuth angle is done with the hour angle (HA) as an intermediary, where $HA = \phi + 90^\circ$. Right ascension can be calculated from the mean sidereal time t_s and HA:

$$\alpha = t_s - \text{HA} \quad . \quad (4.2)$$

The calculation of mean sidereal time from MJD is given by:

$$t_s = 280.46061837 + 1.31850007701 \times 10^7 T + 3.87933 \times 10^{-4} T^2 - 2.58331181 \times 10^{-8} T^3 \quad , \quad (4.3)$$

where

$$T = \frac{\text{MJD} - 51544.5}{36525} \quad . \quad (4.4)$$

Conversions in this analysis were performed using the SLALIB library [92].

4.1.4 Solar Coordinate Frame

As the Earth orbits around the Sun, we observe an excess in the relative intensity of cosmic rays in the direction of motion and a corresponding deficit in the direction opposite to the motion, similar to the Compton-Getting effect mentioned in Sec. 1.7.1. This effect manifests itself as a dipole in the relative intensity when the cosmic-ray arrival directions are plotted using solar time, i.e., in a frame where the position of the Sun is at a fixed location

$$\alpha_{\text{solar}} = \alpha - \alpha_{\odot} \quad . \quad (4.5)$$

In this coordinate system, only a conversion in the right ascension coordinate is applied. The sun does move in declination, but at the South Pole the motion is very gradual and ignored for purposes of this analysis. This solar dipole has been measured previously by IceCube [55, 56], and now serves as a check of the sensitivity of the analysis methods used.

4.1.5 Non-Physical Coordinate Frames

Ideally, the solar dipole should not cause any systematic uncertainties in the analysis of cosmic-ray arrival directions in sidereal time, as any signal in solar time averages to zero over

a year. In practice, however, seasonal variations, gaps in detector uptime, and limited sky coverage can all cause the solar anisotropy to influence the sidereal frame, and vice versa. In order to study this mutual influence, we consider two non-physical time scales: anti- and extended-sidereal time.

In frequency space, solar time has a frequency of 365.24 cycles per year. The sidereal day is roughly four minutes shorter, with a frequency of 366.24 cycles per year. The influence of the solar dipole on the sidereal anisotropy can be estimated from the influence it has on the other side band in frequency-space, i.e., on a frame with 364.24 cycles per year. This is the anti-sidereal frame. No physical signal is expected with a frequency of 364.24 cycles per year, so any significant “signal” that appears in the anti-sidereal frame stems from a modulation of the solar frame, and is equivalent to the systematic effect of the solar frame on the sidereal frame. Frequency-space can also be used to estimate the effect of the sidereal anisotropy on the solar dipole. In this case, the second side band of the sidereal frame, the extended-sidereal time frame with a frequency of 367.24 cycles per year, can be analyzed [93].

4.2 Background Estimation

In order to study the anisotropy in the cosmic-ray arrival direction distribution, we must simulate the detector response to an isotropic sky, known as the *reference map*. As shown in Fig. 4.2, the grid layout of the detector creates preferential lanes for detection. Thus, the detector response to an isotropic sky is not itself isotropic. Ideally, simulations would be used to reproduce the detector response to an isotropic background, but it is difficult to produce simulations with accuracy down to the part-per-mille (or better) level. Instead, the reference map can be calculated directly from data.

4.2.1 Time-Scrambling

Consider a map of all detected events in the sidereal coordinate frame. This data map is anisotropic in nature, partly due to the anisotropy of astrophysical sources, but also partly due to detector effects. Differences in detector acceptance as a function of zenith and

azimuth, gaps in detector uptime, and seasonal variations all affect the observed anisotropy in equatorial coordinates

The time-scrambling method, described in detail in [94], is a common method for estimating a reference map from the local distribution of events. The goal is to separate the anisotropy due to detector response from the anisotropy due to astrophysical sources. The method, in general, involves keeping local coordinates fixed, and assigning times randomly sampled from real arrival times within a given time window. The resultant reference map has events shuffled in right ascension while maintaining the local event distribution.

In greater detail, for every event, 20 events are added to the reference sample with the same arrival direction in the detector’s local coordinate system (θ, ϕ) . Each of these events is then given a new time value randomly selected from a histogram of real arrival times within a chosen time window Δt . They are then converted to equatorial coordinates (δ, α) , and given a weight of $1/20$. The creation of additional events and subsequent weighting serves to reduce statistical fluctuations.

The size of the time window selected determines the sensitivity of the search to features of various angular sizes. Since the Earth rotates 360° in 24 hours, a time window of four hours would make a search sensitive to structures of $4 \text{ hr} / 24 \text{ hr} \times 360^\circ = 60^\circ$ or smaller in size. The size of the time window can intentionally be limited to mask large-scale structures, or limited out of necessity due to changes in detector configuration and/or the environment. The IceCube detector is stable over periods longer than 24 hours, and environmental conditions at the South Pole change very slowly, so in this work a scrambling period of 24 hours is used to make the search sensitive to structures on all angular scales.

It is important to emphasize that scrambling the time of events with a given zenith angle in local detector coordinates is equivalent to randomly modifying the right ascension of the event within the same declination band, with the width of the band determined by the pixelization used (see Sec. 4.3.1). As a result, the residual between the actual arrival direction distribution and the reference maps is sensitive primarily to anisotropy modulations in right ascension. Simulation studies [95] indicate that any large-scale structure is effectively reduced

to its projection onto the equatorial plane, limiting the sensitivity to the determination of the true anisotropy. However, small-scale features, i.e., features well localized in the field of view of the experiment, are not affected by significant deformation. More information on the advantages and disadvantages of this method can be found in [96].

4.2.2 Dipole and Quadrupole Subtraction

Covered in greater detail in Sec. 4.4, large-scale structures in anisotropy can dominate skymaps. To observe small-scale structures, methods of subtracting the large-scale structures are needed. Two methods of subtraction are presented here: time scrambling with windows shorter than 24 hours, and subtraction of a best-fit low-multipole map.

One method for removing the dipole and quadrupole moments was mentioned in Sec. 4.2.1: creating the reference map with time windows of less than 24 hours can mask larger-scale structures. A time scrambling window of 4 hours should only be sensitive to structures 60° or smaller in size, equivalent to masking the dipole and quadrupole components. Unfortunately, analysis of the power spectra of maps produced using this method reveals strong dipole and quadrupole terms (refer to Sec. 4.4 for more details).

A second option is to calculate a two-dimensional fit for the relative intensity map (see Sec. 4.3.2). A directional vector (δ, α) can be written as a unit vector in cartesian coordinates \vec{v} with components (v_x, v_y, v_z) . The fit to relative intensity then solves for the coefficients of the normalized, real Cartesian spherical harmonics for the monopole ($\ell = 0$), dipole ($\ell = 1$), and quadrupole ($\ell = 2$) moments:

$$\begin{aligned} \delta I = & m_0 \frac{1}{2\sqrt{\pi}} + d_{-1} \sqrt{\frac{3}{4\pi}} v_y + d_0 \sqrt{\frac{3}{4\pi}} v_z + d_1 \sqrt{\frac{3}{4\pi}} v_x \\ & + q_{-2} \sqrt{\frac{15}{4\pi}} v_x v_y + q_{-1} \sqrt{\frac{15}{4\pi}} v_y v_z + q_0 \sqrt{\frac{15}{16\pi}} (3v_z^2 - 1) + q_1 \sqrt{\frac{15}{4\pi}} v_x v_z + q_2 \sqrt{\frac{15}{16\pi}} (v_x^2 - v_y^2) . \end{aligned} \quad (4.6)$$

The coefficients m , d , and q correspond to $\ell = 0$, 1, and 2 respectively. The subscript for each coefficient indicates the m value of the term. The minimization of the weighted sum of




Figure 4.5 Pixelization of the sky using HEALPix with (*left to right*) $N_{\text{side}} = 1, 2, 4,$ and $8,$ producing maps with $12, 48, 192,$ and 768 equal-area pixels, respectively. Dots indicate pixel centers, which align in bands of equal-latitude. Source:

<http://healpix.jpl.nasa.gov/>

$(I - \delta I)^2$ is then performed using MINUIT, returning best-fit values for the coefficients. This second method is what was used in the making of Fig. 1.11.

4.3 Mapmaking Procedure

4.3.1 HEALPix

The skymaps in this analysis were produced using HEALPix¹ [97]. As the name suggests, HEALPix divides the unit sphere into pixels of equal solid-angle, with rows of equal latitude. The resolution of a given map is determined by the N_{side} parameter, with the number of pixels $N_{\text{pix}} = 12 \times N_{\text{side}}^2$. N_{side} is a power of 2, ranging from 1 to 1024; the bins for various N_{side} values are shown in Fig. 4.5. For this work, $N_{\text{side}} = 64$ was used, dividing the sky into 49,152 pixels, each with an angular size of 0.84 sq. deg ($2.56 \times 10^{-4} \text{ ster.}$). At this resolution, the pixel size is smaller than the median angular resolution ($\sim 3^\circ$) of the reconstructions used in this work.

In this analysis, all maps undergo a top-hat smoothing procedure in which a single pixel's value is the sum of all pixels within a given angular distance, or smoothing radius. In the case of this data set, the median angular resolution, as found from simulation, is 3° . Therefore,

¹Hierarchical Equal-Area isoLatitude Pixelization: <http://healpix.jpl.nasa.gov>

a smoothing of 5° , roughly equivalent to the optimal bin size for point-source searches [94], is applied to the maps. When statistics are limited, a smoothing of 20° is frequently applied to increase our sensitivity to larger-scale structures.

4.3.2 Relative Intensity and Significance

After the creation of the reference map using the time-scrambling algorithm, relative intensity and significance maps are used to study the anisotropy.

Relative intensity maps display the signal intensity relative to background intensity on a pixel by pixel basis. For pixel i with N_i events and a background (reference map value) of $\langle N \rangle_i$, the relative intensity is given by:

$$\delta I_i = \frac{\Delta N_i}{N_i} = \frac{N_i(\alpha, \delta) - \langle N_i(\alpha, \delta) \rangle}{\langle N_i(\alpha, \delta) \rangle} . \quad (4.7)$$

It is important to note that the error on the relative intensity is large at the limits of the detector exposure because of the small number of events in that region.

The uncertainty in relative intensity must account for the resampling involved in the construction of the reference map. The resampling ratio α_r is given by $\alpha_r \equiv 1/n_{\text{resample}} = 1/20$ in this analysis. To solve for the uncertainty in relative intensity, we can use propagation of error:

$$\sigma_{\delta I_i} = \sqrt{\left(\frac{\partial(\delta I_i)}{\partial(\langle N \rangle_i)}\right)^2 \sigma_{\langle N \rangle_i}^2 + \left(\frac{\partial(\delta I_i)}{\partial(N_i)}\right)^2 \sigma_{N_i}^2} . \quad (4.8)$$

By rewriting $\delta I_i = N_i/\langle N \rangle_i - 1$, this simplifies to:

$$\sigma_{\delta I_i} = \frac{N_i}{\langle N \rangle_i} \sqrt{\frac{1}{N_i} + \frac{\alpha_r}{\langle N \rangle_i}} . \quad (4.9)$$

We also need a way to determine the significance of the difference between the data and reference maps. The statistical significance in this analysis is calculated using the method proposed by Li & Ma [98] — a calculation commonly used when looking for a small signal in samples with a large background. The formula is written for observations of point source (N_{on}) and off-source (N_{off}) events. The ratio of the length of on- to off-source observations is

defined as α_{Li-Ma} , which, in the case of time-scrambling, is equivalent to α_r . The statistical significance S is given by:

$$S = \sqrt{2N_{\text{on}} \ln \left[\left(\frac{1 + \alpha_r}{\alpha_r} \right) \left(\frac{N_{\text{on}}}{N_{\text{on}} + N_{\text{off}}} \right) \right] + 2N_{\text{off}} \ln \left[(1 + \alpha_r) \left(\frac{N_{\text{off}}}{N_{\text{on}} + N_{\text{off}}} \right) \right]} . \quad (4.10)$$

In our adaptation, N_{on} is the number of counts from our data map (N_i). N_{off} is the number of counts from the reference map before weighting ($\langle N \rangle_i / \alpha_r$).

4.4 Power Spectrum

The search for structure at various angular scales can be extended through the study of the power spectrum. This analysis begins with the expansion of the relative intensity map $\delta I = \Delta N / \langle N \rangle$ in a spherical harmonic basis:

$$\delta I(\delta_i, \alpha_i) = \sum_{\ell=1}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_{\ell m}(\delta_i, \alpha_i) , \quad (4.11)$$

where $Y_{\ell m}$ are the Laplacian spherical harmonics and $a_{\ell m}$ the multipole coefficients. Given full sky coverage and 100% detector uptime, the $a_{\ell m}$ coefficients could be estimated with

$$a_{\ell m} \sim \Omega_p \sum_{i=0}^{N_{\text{pix}}} \delta I(\delta_i, \alpha_i) Y_{\ell m}(\delta_i, \alpha_i) , \quad (4.12)$$

where Ω_p is the solid angle subtended by each pixel, which is constant for a given N_{side} in HEALpix. Using these $a_{\ell m}$ values, the power spectrum could be found using

$$C_{\ell} = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{\ell} |a_{\ell m}|^2 . \quad (4.13)$$

The resulting strength of the C_{ℓ} values for a given ℓ indicates structure on the corresponding angular scale. The dipole component ($\ell = 1$) corresponds to structures 180° in size, the quadrupole ($\ell = 2$) to structures 90° in size, and so on.

The primary issue with this approach is that it requires an idealized detector for the Fourier decomposition to work. When analyzing the relative intensity observed by a detector with partial sky coverage and gaps in run time, the $Y_{\ell m}$ spherical harmonics no longer form

an orthonormal basis, resulting in the coupling of ℓ -modes. The work in [99] provides a more detailed discussion on the effects of this coupling, and several ways in which to correct for it.

In this analysis, the `PolSpice`² [100, 101] software package is used to produce the power spectrum. To correct for the partial-sky bias, `PolSpice` calculates C_ℓ values through the two-point correlation function $\xi(\eta)$:

$$\xi(\eta) = \langle \delta I(\delta_i, \alpha_i) \delta I(\delta_j, \alpha_j) \rangle \quad , \quad (4.14)$$

where η is the angle between pixels i and j . The two-point correlation function is related to C_ℓ via the inverse of a Legendre series expansion:

$$C_\ell = 2\pi \int_{-1}^1 \xi(\eta) P_\ell(\cos \eta) d(\cos \eta) \quad . \quad (4.15)$$

`PolSpice` reduces, but does not eliminate, correlation between different ℓ -modes. Other tools such as apodization are applied, but in the final result low- ℓ multipoles with high power can still affect high- ℓ modes. Figure 4.6 shows the power spectrum as calculated in [55], with all events (*blue*), with the best-fit dipole and quadrupole map subtracted (*red*). Differences between the power spectra above $\ell = 2$ are indicative of the effect of the low-multipole moments on higher ℓ -modes.

To estimate errors for a given set of C_ℓ values in this analysis, we create a large number of maps from the power spectrum using the `synfast` function from HEALPIX. The power spectra for each of these maps is then recalculated using `PolSpice`, and its values stored. The error bars shown contain 68% of C_ℓ values (34% in each direction) produced using this method.

4.5 Bayes Factor

The sensitivity of IceTop to the entire electromagnetic component of cosmic-ray air showers allows for energy resolution approaching $\log_{10}(E/\text{GeV}) \sim 0.05$ using `ShowerLLH` (see

²SPatially Inhomogeneous Correlation Estimator for temperature and Polarisation: <http://www2.iap.fr/users/hivon/software/PolSpice/>




Figure 4.6 Power spectra for IC59 data as produced using `PolSpice`, shown for all events (*blue*) and the residual after subtracting the best-fit dipole and quadrupole map (*red*). Light and dark gray bands indicate 68% and 95% containment, respectively, of power spectra generated by isotropic maps. Source: [55]

Sec. 3.2.2.2). With this level of energy resolution, the energy spectrum can be studied as a function of sky position, as binned with HEALPix. The spectrum for any given bin will naturally be different from other bins, due to differences in relative exposure. In order to compare bins, a Bayesian approach is used to compare the shape of a distribution while allowing for differences in exposure, as detailed in [102].

Given two observed distributions \mathcal{F}_1 and \mathcal{F}_2 , the Bayes factor is defined as the ratio of the likelihoods that the observed distributions resulted from one- or two-parent hypotheses, H_1 and H_2 respectively.

$$B_{21} = \frac{P(\mathcal{F}_1, \mathcal{F}_2|H_2)}{P(\mathcal{F}_1, \mathcal{F}_2|H_1)} \quad (4.16)$$

Using Bayes' theorem, we can write the posterior probability of the one-parent hypothesis as

$$\begin{aligned} P(H_1|\mathcal{F}_1, \mathcal{F}_2) &= \frac{P(\mathcal{F}_1, \mathcal{F}_2|H_1)P(H_1)}{P(\mathcal{F}_1, \mathcal{F}_2|H_1)P(H_1) + P(\mathcal{F}_1, \mathcal{F}_2|H_2)P(H_2)} \\ &= \frac{1}{1 + \frac{P(\mathcal{F}_1, \mathcal{F}_2|H_2)P(H_2)}{P(\mathcal{F}_1, \mathcal{F}_2|H_1)P(H_1)}} \\ &= \frac{1}{1 + B_{21} \frac{P(H_2)}{P(H_1)}} \end{aligned} \quad (4.17)$$

In the case that neither model is favored before taking data ($P(H_1) = P(H_2)$), this becomes

$$P(H_1|\mathcal{F}_1, \mathcal{F}_2) = \frac{1}{1 + B_{21}} \quad (4.18)$$

Using this equation, a Bayes factor of $B_{21} > 100$ indicates strong evidence ($> 99\%$ chance) that the distributions compared have different sources. A Bayes factor of $B_{21} < 10^{-2}$ indicates strong evidence ($> 99\%$ chance) that the distributions compared have the same source.

The primary formula used is the calculation of the natural logarithm of the Bayes factor, which compares the shapes of distributions \mathcal{F}_1 and \mathcal{F}_2 , with N_1 and N_2 events, respectively:

$$\begin{aligned} \ln B_{21} &= \sum_{i=1}^N [\ln(\Gamma(\mathcal{F}_{1,i} + 1)) + \ln(\Gamma(\mathcal{F}_{2,i} + 1)) - \ln(\Gamma(\mathcal{F}_{1,i} + \mathcal{F}_{2,i} + 2))] \\ &\quad + \ln(\Gamma(N_1 + N_2 + 2)) - \ln(\Gamma(N_1 + 1)) - \ln(\Gamma(N_2 + 1)) \end{aligned} \quad (4.19)$$

Chapter 5

Results

In this chapter, we will detail the results of observations of cosmic-ray anisotropy as a function of angular scale, time, and energy. Solar, anti- and extended-sidereal timeframes are considered for consistency checks and error estimation. Finally, the results of studies with ShowerLLH are presented, with a focus on the energy spectrum and composition distribution as a function of sky position.

5.1 Angular Scales

As mentioned in Sec. 4.2.2, the anisotropy in the cosmic-ray arrival direction distribution manifests over a large number of angular scales. In this section, we look at the dominant large-scale components, as well as the small-scale structure revealed after their subtraction. The dataset used for this analysis consists of observations over a 6-year time period, as shown in Table 5.1.

5.1.1 Large-Scale Structure

Figure 5.1 shows relative intensity and pretrial significance sky maps for the large-scale structures, as shown in Mollweide (*top*) and polar (*bottom*) projections. All of the maps contain six years of IceCube data, smoothed with a 5° angular radius. The relative intensity map, shown in the left plot of Fig. 5.1, is similar to previously published work based on IC59 data [55] and shows anisotropy at the 10^{-3} level, characterized by a large excess from 30° to 120° and a deficit from 150° to 250° . The corresponding significance of the large-scale

Configuration	Livetime (days)	Number of Events
IC59	339.38 (91.7%)	3.579×10^{10}
IC79	315.76 (92.6%)	4.131×10^{10}
IC86	343.04 (93.0%)	5.906×10^{10}
IC86-II	331.92 (94.0%)	5.630×10^{10}
IC86-III	362.20 (97.9%)	6.214×10^{10}
IC86-IV	369.76 (97.8%)	6.327×10^{10}
Total	2062.06 (94.5%)	3.179×10^{11}
Configuration	Livetime (days)	Number of Events
IT59	338.25 (91.4%)	2.887×10^7
IT73	312.66 (91.7%)	3.690×10^7
IT81	343.04 (93.0%)	3.800×10^7
IT81-II	332.26 (94.1%)	3.713×10^7
IT81-III	361.20 (97.6%)	3.101×10^7
IT81-IV	362.61 (95.9%)	2.810×10^7
Total	2050.01 (94.0%)	1.719×10^8

Table 5.1 Detector configurations and their respective number of events for all years used in this analysis. *IC* indicates IceCube, *IT* IceTop, and the number that follows indicates the number of strings or stations participating in data acquisition.

structure, shown in the right plot of Fig. 5.1, shows the increasingly solid evidence of the observation.

5.1.2 Small-Scale Structure

While the large-scale structure dominates the anisotropy, there is also anisotropy on smaller scales. This structure, with a relative intensity on the order of 10^{-4} , becomes visible after the best-fit dipole and quadrupole are subtracted from the sky map. Figure 5.2 shows the relative intensity of the residual map (*left*), and the corresponding significance map (*right*), as shown in Mollweide (*top*) and polar (*bottom*) projections. These maps show the

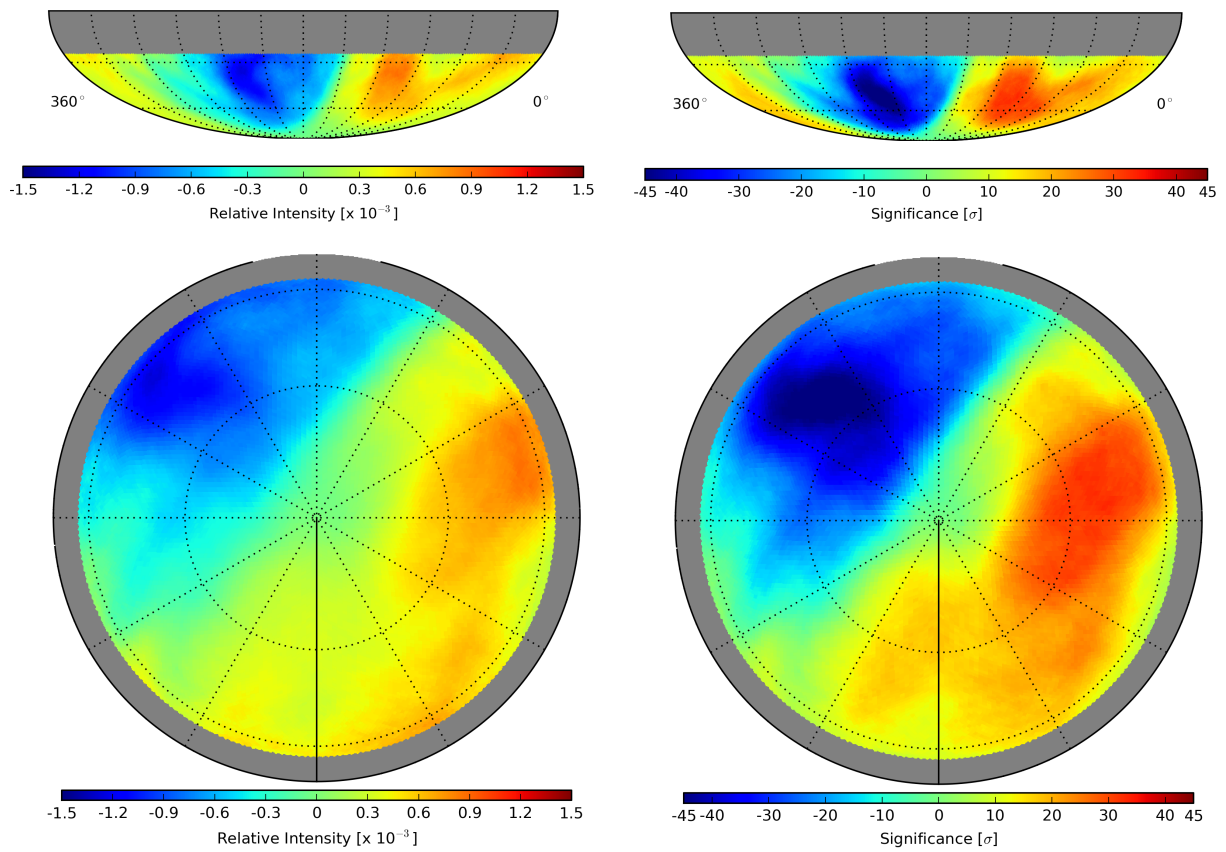


Figure 5.1 Relative intensity (*left*) and pretrial statistical significance (*right*) maps for large-scale structure in Mollweide (*top*) and polar (*bottom*) projections. The maps are in equatorial coordinates and use an angular smoothing radius of 5° .

presence of cosmic-ray anisotropy at TeV energies at angular scales approaching the angular resolution of IceCube for cosmic-ray primaries.

Table 5.2 shows the positions and peak significances of excess and deficit regions with a pretrial significance exceeding 5σ . The regions are numbered to maintain consistency with [55] whenever possible. The significances quoted are pretrial, and any blind search would have to account for the fact that we search for significant excess or deficit regions anywhere in the roughly 10^4 independent bins of the map. However, all but two of the regions listed in Table 5.2 have been previously reported in the analysis of the IC59 data set. In the new data set, which includes the IC59 data set but is a factor of 7 larger, all regions appear with greatly increased significance.

The study of small-scale structures in previously published studies [55] was limited by the statistics available at that time. In order to observe larger structures, a scan in smoothing scale was performed and results were reported for smoothing scales that maximize the significance. With six years of data, this procedure is no longer necessary, as we can now observe small-scale structures at high significance with smoothing radii approaching our angular resolution. Figure 5.3 provides two examples of the increased angular resolution available using the new data set. The left plots show regions as observed using only the IC59 data set with 20° smoothing applied, as in [55]. The same regions are shown in the right plots, using the updated data set and 5° angular smoothing. In both cases, what previously appeared as one region is now observed as two distinct regions, each at high significance. This difference does not appear to be the result of a time-dependence of the small-scale structure, as the same split is visible in the IC59 map with 5° smoothing but not at high enough significance to be previously reported.

5.1.2.1 Four Hour Time-Scrambling

The small-scale structure can also be observed by using a four hour time window in the production of the reference maps, which should mask all structures on angular scales $> 60^\circ$ (see Sec. 4.2.1). The results of this method are shown in Fig. 5.4, and largely correspond

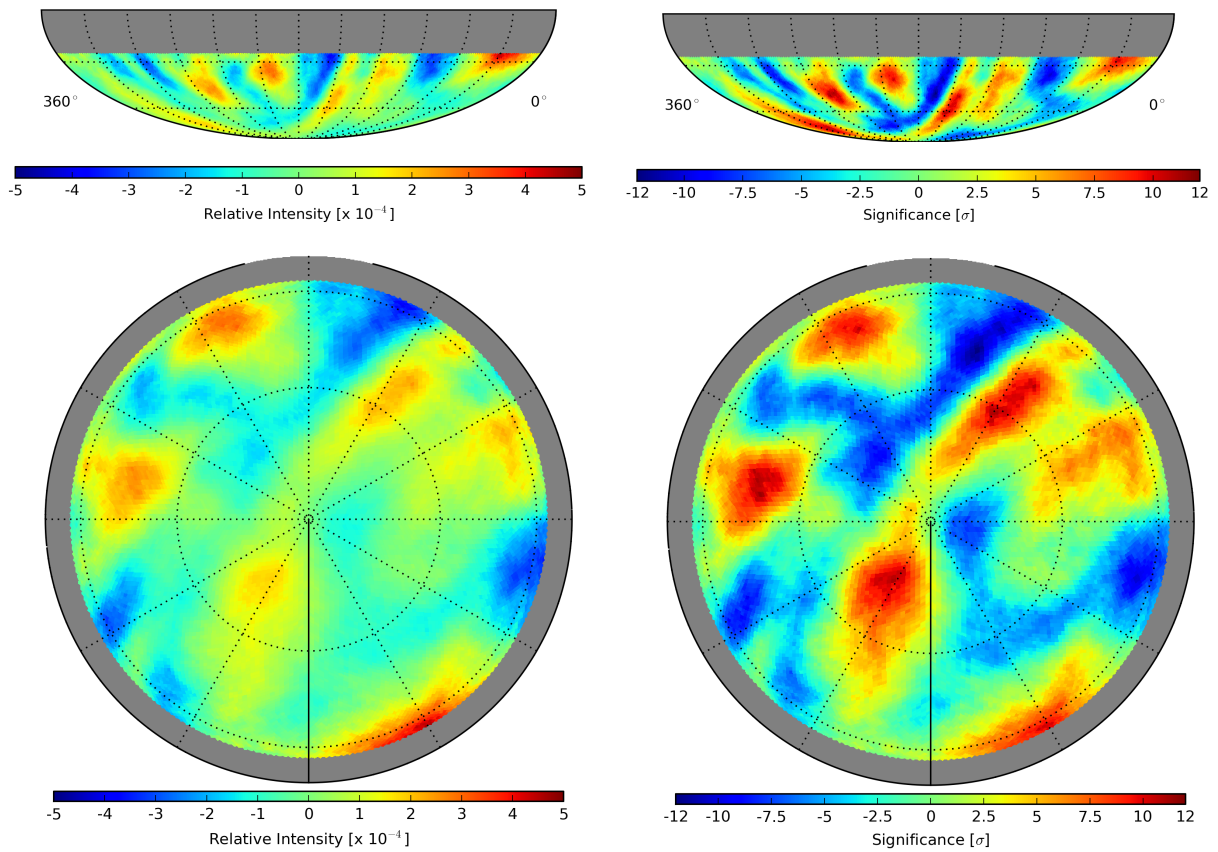


Figure 5.2 Relative intensity (*left*) and pretrial statistical significance (*right*) maps shown after dipole- and quadrupole-subtraction in Mollweide (*top*) and polar (*bottom*) projections. The maps are in equatorial coordinates and use an angular smoothing radius of 5° .

Region	Right Ascension (deg)	Declination (deg)	Peak Significance
1a	$142.5^{+4.9}_{-2.4}$	$-49.7^{+2.3}_{-3.9}$	11.0σ
1b	$110.5^{+5.3}_{-3.5}$	$-55.9^{+5.4}_{-2.3}$	6.9σ
2	$261.0^{+3.4}_{-8.5}$	$-48.9^{+4.7}_{-2.3}$	11.4σ
3	$200.4^{+2.8}_{-1.4}$	$-38.7^{+2.3}_{-2.3}$	10.8σ
4	$327.9^{+11.9}_{-16.8}$	$-74.6^{+4.4}_{-4.4}$	11.0σ
5	$215.6^{+18.7}_{-8.6}$	$-72.4^{+5.2}_{-2.2}$	-9.3σ
6	$74.5^{+4.2}_{-4.2}$	$-36.4^{+5.0}_{-3.8}$	-10.3σ
7a	$317.1^{+4.2}_{-2.1}$	$-38.7^{+5.2}_{-1.5}$	-7.2σ
7b	$292.5^{+1.4}_{-1.4}$	$-41.0^{+1.6}_{-1.6}$	-9.6σ
8	$164.7^{+3.2}_{-2.9}$	$-48.1^{+4.7}_{-3.9}$	-11.9σ
9*	$94.1^{+9.4}_{-40.9}$	$-82.0^{+5.1}_{-2.2}$	-7.9σ
10*	$27.4^{+4.9}_{-1.4}$	$-27.3^{+3.3}_{-2.0}$	10.6σ

Table 5.2 Locations and pretrial peak significance values for the small-scale structures visible after subtracting the best-fit dipole and quadrupole moments. Negative significances indicate deficits in the sky map. Errors indicate the positions of the furthest pixels within 1σ of the peak significance. Regions marked *a* and *b* were previously reported as one region. Regions with an asterisk are new to this analysis.




Figure 5.3 Significance maps in the vicinity of Region 1 (*top*) and Region 7 (*bottom*) (see Table 5.2) as previously published using only data taken with the IC59 configuration ([55]) with 20° smoothing (*left*) and for the full data set used in this analysis with 5° smoothing (*right*). Maps are shown in equatorial coordinates.

to the structures found in Fig. 5.2. This method is not treated as the primary method of dipole-quadrupole subtraction, due to the continued presence of low-multipole moments in the power spectra (see Sec. 5.1.3).

5.1.3 Power Spectrum

The angular power spectrum for the six-year data set is shown in Fig. 5.5. Similar to previous work, it is calculated using `PolSpice` [100, 101], which corrects for systematic effects introduced by partial-sky coverage. The power spectrum is shown before (*blue*) and after (*red*) subtracting the dipole and quadrupole moments from the sky map. In addition, the power spectrum is calculated using a four hour time-scrambling window, illustrating the failure of the method to completely remove low-multipole moments. To calculate the error bars, randomized maps were generated using the C_ℓ values from the original relative intensity map (see Sec. 4.4). The error bars contain 68% of the resultant C_ℓ values from the new maps. The power spectrum confirms the presence of significant structure up to multipoles $\ell \simeq 20$, corresponding to angular scales of less than 10° .

The comparison of the angular power spectrum before and after subtraction of the best-fit dipole and quadrupole moments serves as an estimate of the systematic error caused by the partial-sky coverage. Because of this limitation, the different multipole moments are no longer independent, and while `PolSpice` tries to mitigate the effect of coupling between multipole moments, a significant coupling between the low- ℓ modes remains. This is apparent in the strong reduction in the power of the $\ell = 3$ and $\ell = 4$ multipoles after the dipole and quadrupole subtraction. The systematic error on the $\ell = 3$ and $\ell = 4$ multipoles is therefore large. For multipoles $\ell \geq 6$, the distortion is much smaller and the spectra agree within uncertainties. For these moments, the systematic errors on the C_ℓ are therefore at most of the same order as the statistical errors.

In the unsubtracted power spectrum, the uncertainty in the lower multipole moments causes the C_ℓ value for $\ell = 5$ to be negative — a result of `PolSpice`'s calculation of the C_ℓ values through the use of the two-point autocorrelation function. Simulations using

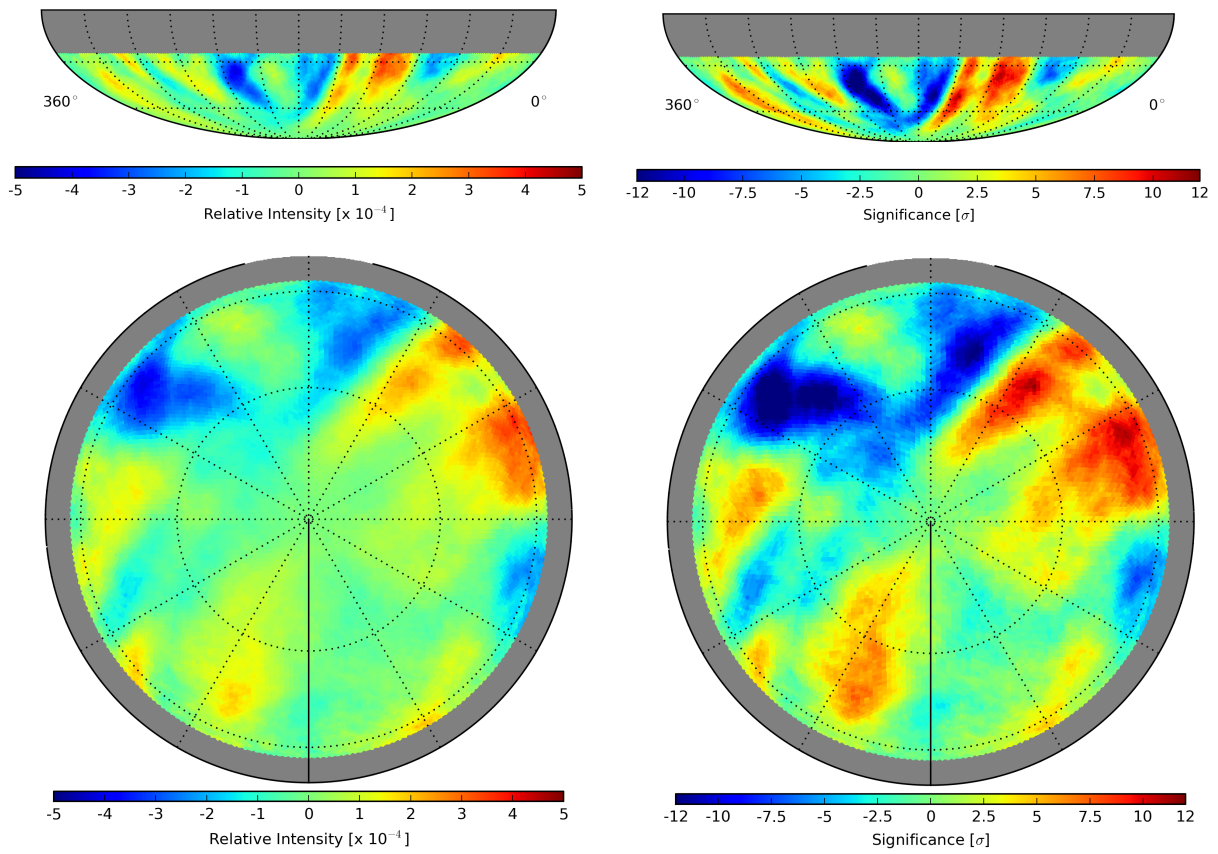


Figure 5.4 Relative intensity (*left*) and pretrial statistical significance (*right*) maps produced using a four hour time-scrambling window, shown in Mollweide (*top*) and polar (*bottom*) projections. The maps are in equatorial coordinates and use an angular smoothing radius of 5° .

artificial sky maps with strong dipole components indicate that this behavior is typical for the weighting and apodization used in this analysis and is another indication of the coupling between low- ℓ multipoles.

5.1.4 Systematic Checks

In [54], several sources of systematic bias are considered, including detector geometry and livetime, nonuniform exposure to different regions of the sky, and seasonal variations in atmospheric conditions. The location of the IceCube detector continues to minimize the effect of some of these sources; the southern celestial sky is fully visible at all times, and seasonal variations are slow and accounted for in the estimation of the reference level. The checks performed in that previous analysis continue to hold, and the detector livetime has improved on average, as seen in Table 5.1. In this section we expand on one possible source of systematic bias that the increased data set allows us to study in more detail: the possible influence of the solar dipole on the sidereal signal and vice versa.

5.1.4.1 Solar Dipole Analysis

As the Earth orbits around the Sun, we observe an excess in the relative intensity of cosmic rays in the direction of motion and a corresponding deficit in the direction opposite to the motion. This effect manifests itself as a dipole in the relative intensity when the cosmic-ray arrival directions are plotted using solar time, i.e., in a frame where the position of the Sun is at a fixed location (see Sec. 4.1.4). This solar dipole has been measured previously [55, 56] and now serves as a check of the consistency and reliability of the analysis methods used.

Figure 5.6 shows the solar dipole as seen in both IceCube and IceTop. The statistical advantage of IceCube produces a significant dipole. In comparison, IceTop’s limited statistics make the dipole less significant. The relative intensity map (*bottom left*) has a strong feature near $\alpha - \alpha_{\odot} = 0$ at the upper edge of the map. This feature is augmented by the decreased




Figure 5.5 Angular power spectra for the relative intensity map for six years of IceCube data. Blue and red points show the power spectrum before and after the subtraction of the best-fit dipole and quadrupole terms from the relative intensity map. Green points show the power spectrum calculated from maps made with a four hour time-scrambling window. Error bars are statistical (see the text for a discussion of systematic errors). The gray bands indicate the 68% (dark) and 95% (light) spread in the C_ℓ for a large sample of isotropic data sets.

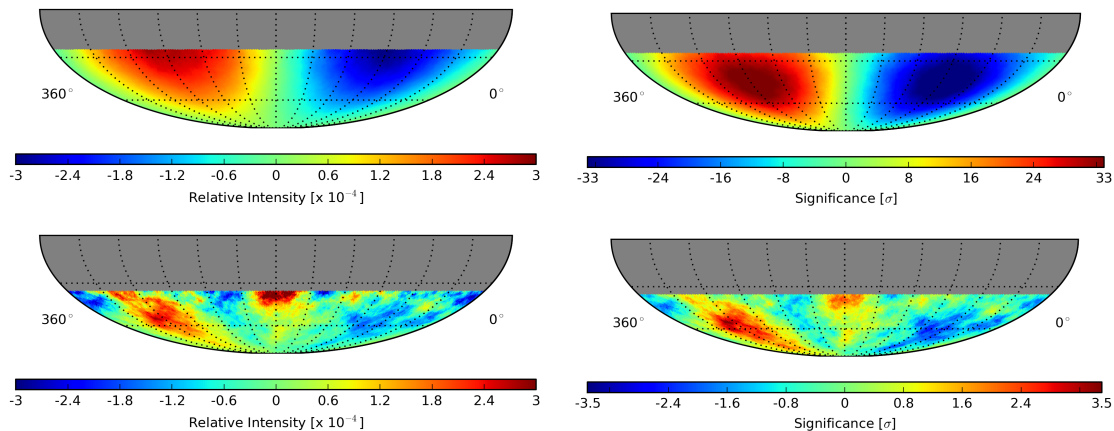


Figure 5.6 Relative intensity (*left*) and significance (*right*) maps for IceCube (*top*) and IceTop (*bottom*) data in solar time with 20° angular smoothing applied. The latitude coordinate still represents declination δ , while the longitude coordinate represents the difference of the right ascension of the event and the right ascension of the Sun ($\alpha - \alpha_\odot$).

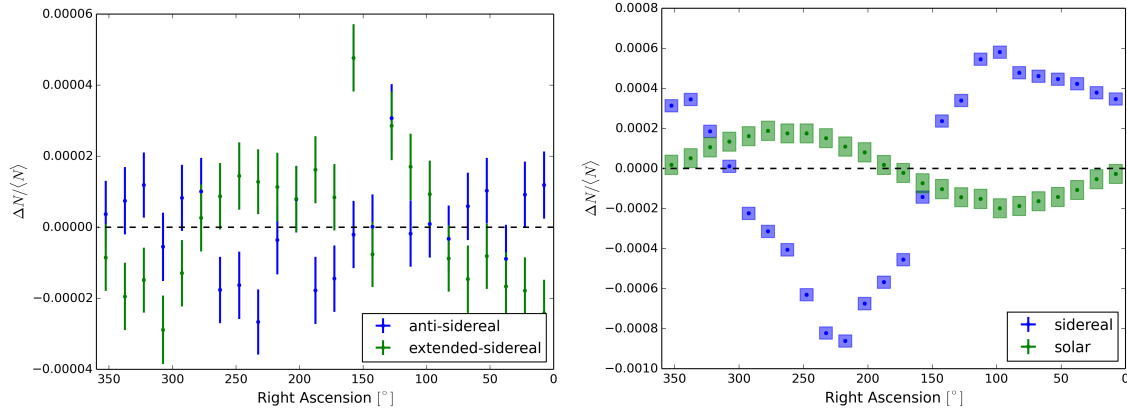


Figure 5.7 Projection of relative intensity as a function of right ascension in anti-sidereal and extended-sidereal time (*left*) and solar and sidereal time (*right*). Error boxes for the solar and sidereal projection indicate systematic errors. Note that both the anti-sidereal and extended-sidereal frames are non-physical, so “right ascension” has no physical meaning in these frames. For the solar frame, the right ascension axis shows the difference between the right ascension of the event and the right ascension of the Sun.

statistics towards the edge of the map, and is diminished in the significance map (*bottom right*).

Ideally, the solar dipole should not cause any systematic uncertainties in the analysis of cosmic-ray arrival directions in sidereal time, as any signal in solar time averages to zero over a year. In practice, however, seasonal variations in the solar dipole can manifest themselves as an anisotropy in the sidereal time frame and vice versa. In order to study this mutual influence, we consider two nonphysical time scales: anti- and extended-sidereal time.

5.1.4.2 Non-Physical Time Analysis

As mentioned in Sec. 4.1.5, detector effects such as seasonal variation, gaps in uptime, and limited sky coverage can all lead to the mutual influence of the sidereal and solar signals. The anti- and extended-sidereal frames are studied in order to estimate the influence of solar signals on the sidereal frame and vice versa, respectively.

Figure 5.7 (*left*) shows the projection of the relative intensity in anti-sidereal and extended-sidereal time. Note that the amplitude of these projections is an order of magnitude smaller

than the amplitude of the solar dipole and the sidereal anisotropy (*right*), indicating that the effect of the solar on the sidereal frame and vice versa is small. We use the maximum amplitude of the relative intensity in the anti- and extended-sidereal frames as a conservative estimate for the systematic error in the sidereal and solar frame resulting from the other.

With systematic errors, we can study the one-dimensional projection of the relative intensity in the sidereal and solar frames, as shown in Figure 5.7 (*right*). In this system, the Sun is located at 0° , so the maximum of relative intensity is at an angle of 270° , in the direction of the Earth's motion. The solar dipole is well measured with six years of IceCube data, and the fit with a dipole results in a χ^2 -probability of 0.38 ($\chi^2 = 20.41$ for 23 degrees of freedom). The amplitude of the sidereal anisotropy is larger, but is not well described by a dipole. Statistical errors are shown, but are very small due to the high statistics of the data set.

5.2 Time Dependence

The IceCube data used in this analysis was recorded over a period of six years and therefore also allows for a study of the stability of the anisotropy over this time period. An observed time-modulation of the anisotropy, in particular one that coincides with the 11-year solar cycle, could be evidence for a heliospheric influence on the observations. Time-dependent studies have been performed previously by several experiments, with contradictory results. Milagro [45] reported a steady increase in the amplitude of their deficit region over a seven-year time period (2000-2007). Both the Tibet [67] and ARGO-YBJ [46] experiments, however, have observed no significant variation in the anisotropy.

5.2.1 Time Dependence of Large-Scale Anisotropy

Figure 5.8 shows the one-dimensional projection of the anisotropy for each detector configuration used in this analysis, each one corresponding to approximately a year of data (see Table 5.1). The shaded regions represent systematic errors determined by calculating the




Figure 5.8 Projection of relative intensity for all declinations as a function of right ascension for each configuration of the IceCube detector from IC59 to the fourth year of IC86. The shaded areas indicate systematic errors, calculated using the anti-sidereal frame for each year independently.

Configuration	χ^2	N_{dof}	p -value
IC59	23.68	23	0.42
IC79	23.69	23	0.42
IC86	21.61	23	0.54
IC86-II	14.90	23	0.90
IC86-III	25.86	23	0.31
IC86-IV	35.06	23	0.05

Table 5.3 Results of a χ^2 test comparing the relative intensity profile (see Figure 5.8) for each year to the collective ensemble. The table provides χ^2 values, number of degrees of freedom, N_{dof} , and corresponding p -values.

maximum amplitude of the signal in the anti-sidereal time frame. Systematic errors are estimated separately for each detector configuration.

Within errors, the large-scale structure is stable over the data-taking period considered here. Table 5.3 shows the χ^2 values calculated by comparing each year to the collective ensemble. The resulting p -values are consistent with random fluctuations, indicating there is no time dependence over the period of this study. A study of the stability over a period of twelve years (2000-2012) using data recorded with the AMANDA and IceCube detectors [103] with the same method also did not find evidence for a time dependence of the structure.

5.2.2 Time Dependence of Small-Scale Anisotropy

To study the time dependence of the small-scale structure, we analyze the relative intensity of the excess and deficit regions listed in Table 5.2 as a function of time. For the location of the regions, we use the values determined from the full six-year data set. Figure 5.9 shows the relative intensity for each detector configuration, i.e., as a function of time, for each region. Also shown is the average value as determined from the analysis of the full six-year data set. The error bars on the data points and the error band on the average indicate statistical uncertainties only, but we list the average flux, including statistical and systematic errors, for each region in the figures. The systematic errors for the individual years have similar

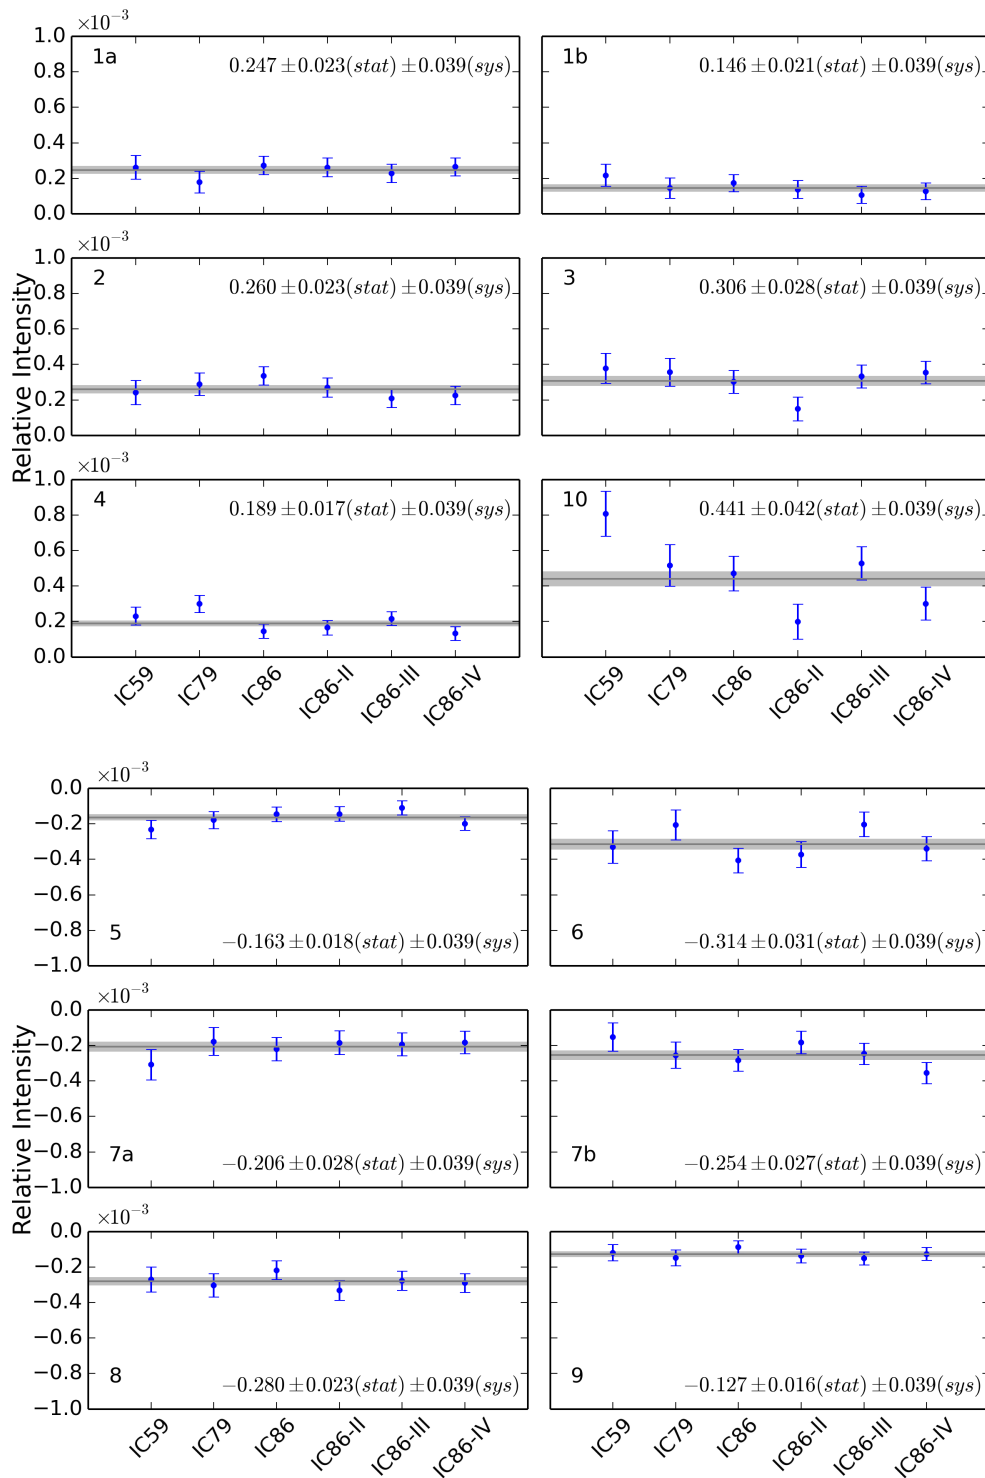


Figure 5.9 Relative intensity as a function of detector configuration for the excess and deficit regions of the small-scale structure listed in Table 5.2. The horizontal lines indicate the six-year average flux. Error bars and bands are statistical, but the relative intensity, including statistical and systematic errors, is given for each region.

values. The relative intensity at the excess and deficit regions of the small-scale structure is constant within errors for the time period covered by this analysis.

As an additional test of the stability of the small-scale map, we subtract the relative intensity sky map of the full six-year data set from the sky map of each individual detector configuration, i.e., of each of the six years of data, and calculate the angular power spectrum of the residual maps. Figure 5.10 shows that for each year, the power spectrum of the resulting maps are, within errors, compatible with isotropy, indicating that there are no significant differences between the maps for individual years and the average. The small-scale anisotropy, like the large-scale anisotropy, is constant over the time period covered by this analysis.

5.3 Energy Dependence

To study the energy dependence of the cosmic-ray anisotropy, we split the data into the nine energy bins described in Sec. 3.1.2. This results in a sequence of maps with increasing median energy, starting with 13 TeV for the lowest-energy bin to 5.3 PeV for the highest-energy bin. The sky maps in relative intensity for all nine energy bins in equatorial coordinates are shown in Fig. 5.11. In addition to the nine maps based on IceCube data, we also show the IceTop map with its median energy of 1.6 PeV. Because of the reduced statistics in these maps, we have applied a top-hat smoothing procedure with a smoothing radius of 20° , improving the sensitivity to larger structure. Note that the relative intensity scale for these plots is identical for energies up to 1.4 PeV, after which it switches to a different scale to account for the strong increase in relative intensity.

The maps clearly indicate a strong energy dependence of the global anisotropy. The large excess from 30° to 120° and deficit from 150° to 250° that dominate the sky map at lower energies gradually disappears above 50 TeV. Above 100 TeV a change in the topology is observed, and at higher energies, the anisotropy is characterized by a wide relative deficit from 30° to 120° , with an amplitude increasing with energy up to at least 5 PeV, the highest energies currently accessible to IceCube. To further illustrate the phase change, the relative

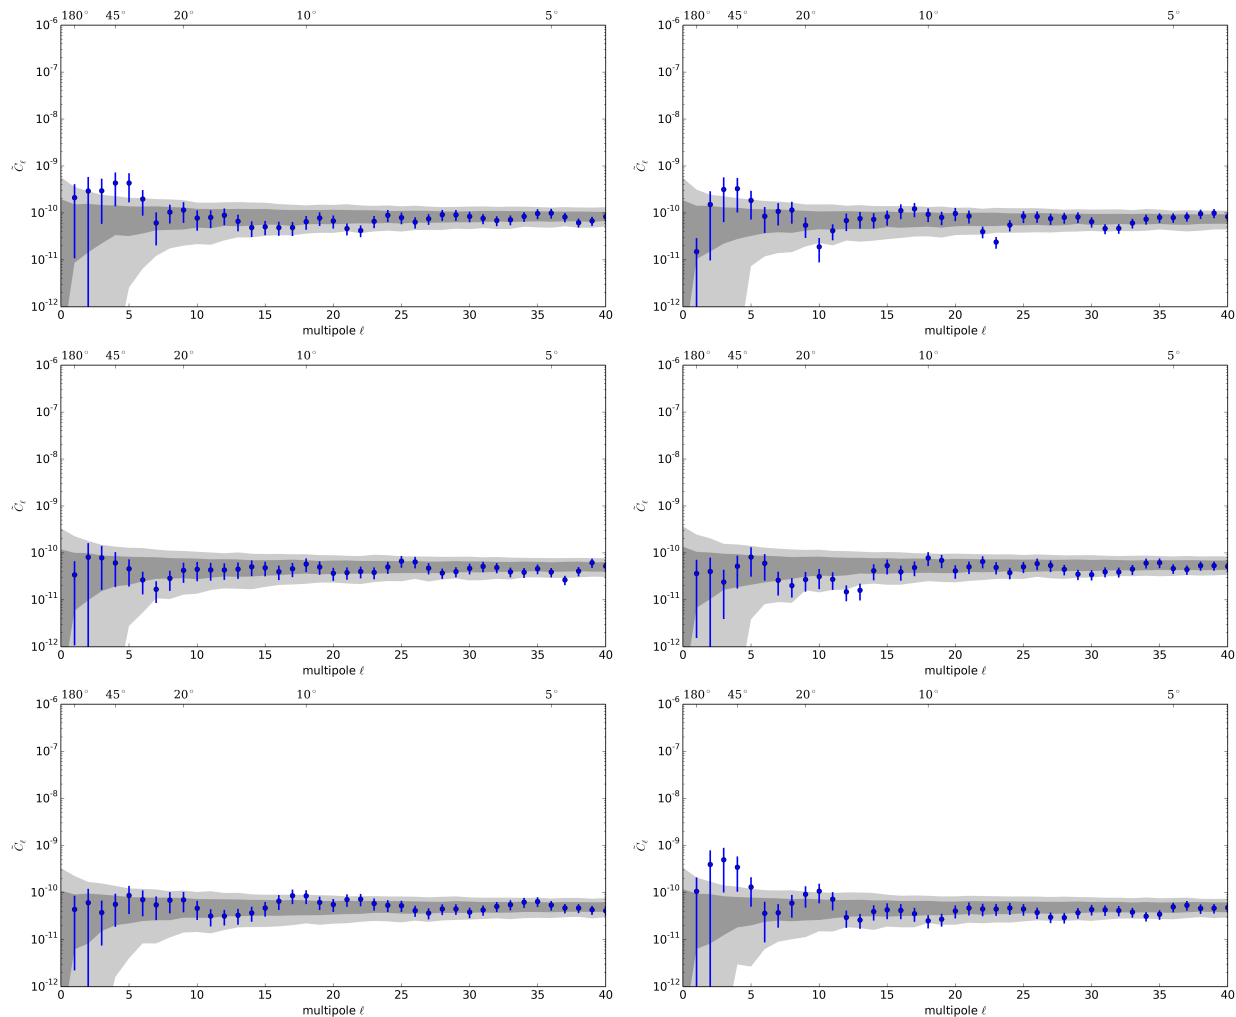


Figure 5.10 Power spectra calculated for residual maps produced by subtracting the six-year relative intensity map from the relative intensity map for each detector season, from IC59 (*upper left*) to IC86-IV (*lower right*). Error bars are statistical (see the text for a discussion of systematic errors). The gray bands indicate the 68% (dark) and 95% (light) spread in the C_ℓ for a large sample of isotropic data sets calculated for each detector configuration independently.

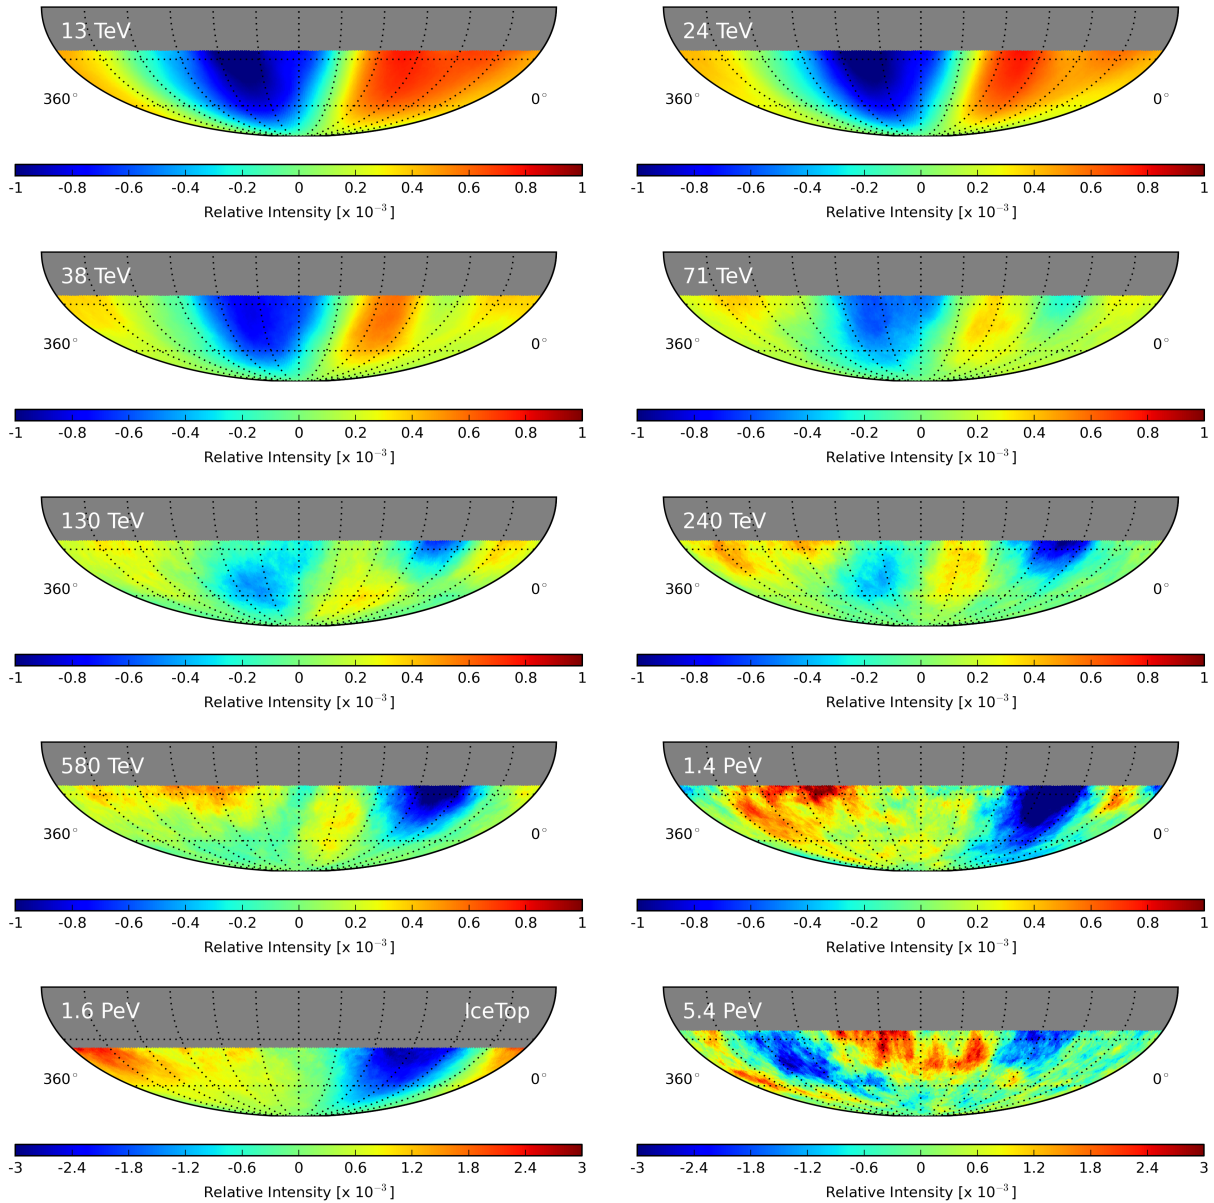


Figure 5.11 Relative intensity maps in equatorial coordinates for the energy bins described in Sec. 3.1.2. The median energy of the data shown in each map is indicated in the upper left. Maps have been smoothed with a 20° smoothing radius. The final two maps are shown on a different relative intensity scale. The map at 1.6 PeV in the lower left panel is based on IceTop data, all other maps show IceCube data.

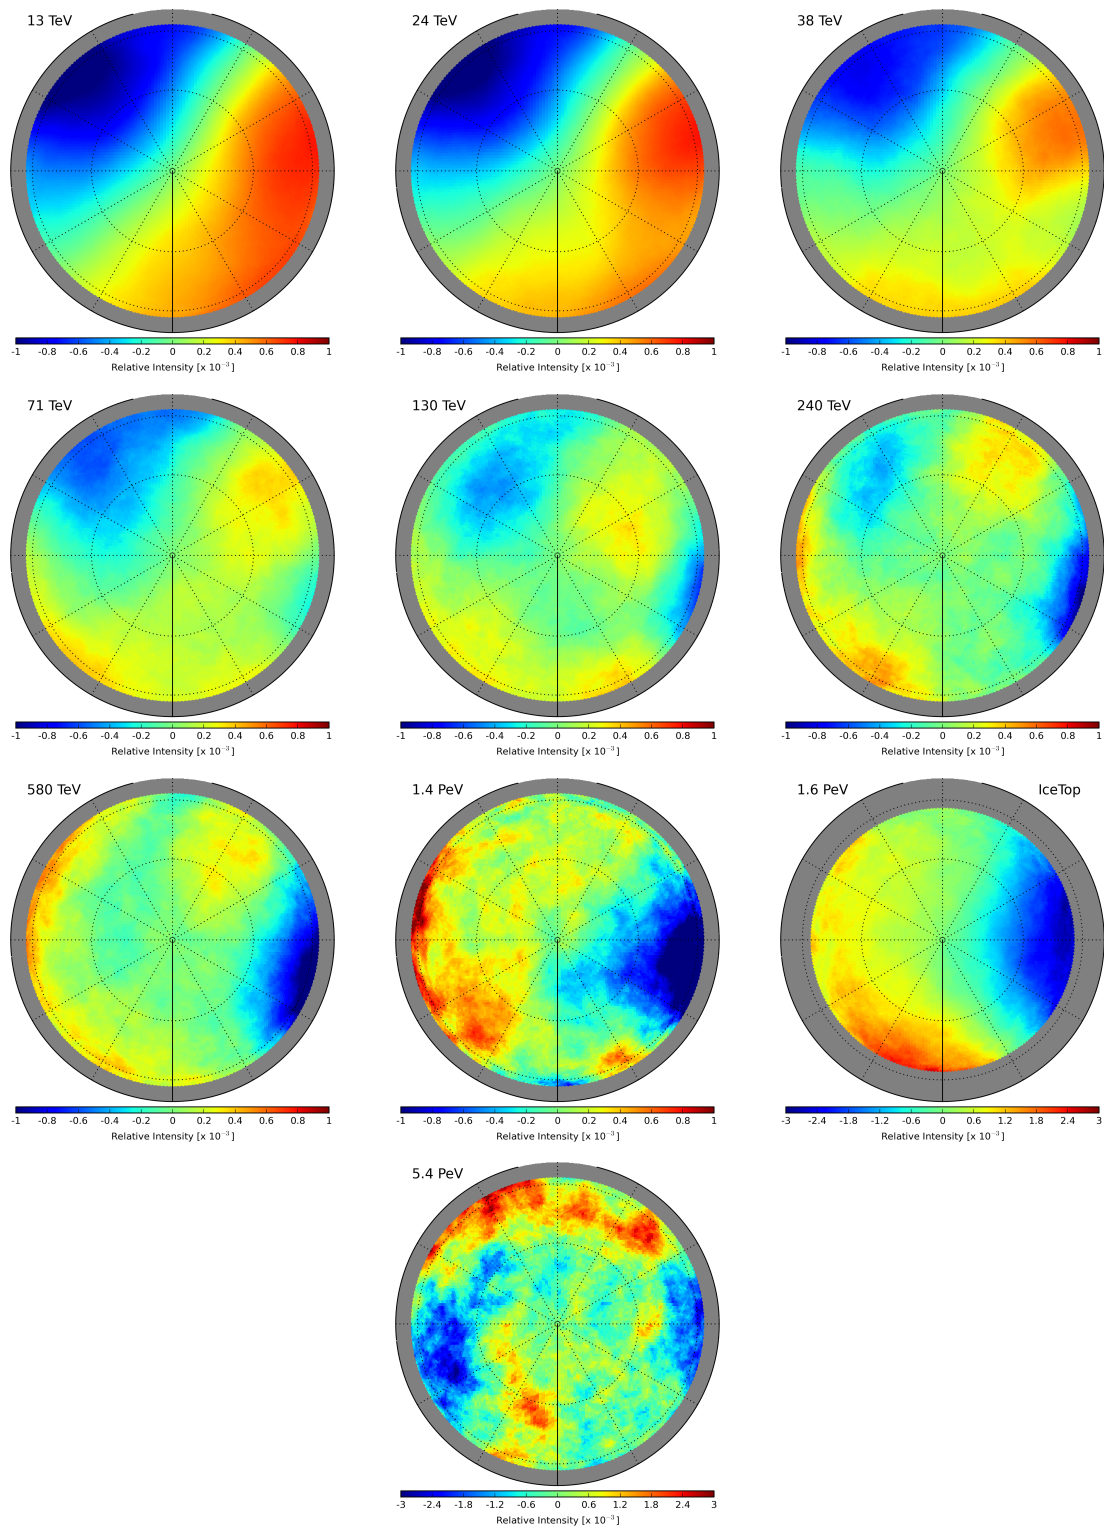


Figure 5.12 Relative intensity maps in polar coordinates for the energy bins described in Sec. 3.1.2. The median energy of the data shown in each map is indicated in the upper left. Maps have been smoothed with a 20° smoothing radius. The final two maps are shown on a different relative intensity scale. The map at 1.6 PeV is based on IceTop data, all other maps show IceCube data.

intensity sky maps are shown in polar coordinates in Fig. 5.12. It is important to note that the time-scrambling method used to calculate the reference map decreases in sensitivity as we approach the polar regions. This effect is clearly visible in Fig. 5.12, where the relative intensity approaches zero at the pole for each map, but is not indicative of the topology of the true anisotropy.

Because of the poor energy resolution, which causes each sky map shown to have a wide energy distribution, it is difficult to determine precisely where the transition in anisotropy occurs, as its slow transition over several maps could be due to overlapping energy distributions or the actual mechanism of change. To illustrate the energy dependence of the phase and strength of the anisotropy, we show in Fig. 5.13 amplitude (*left*) and ϕ -component of phase (*right*) of the dipole moment as a function of energy. Both values are calculated by fitting the full set of harmonic functions to a one-dimensional projection of the relative intensity onto right ascension. We fit this projection rather than the full sky map because the two-dimensional fit of spherical harmonics to the map is difficult to perform with a limited field of view. As a result of the method we apply to generate the reference map, the sky map will in any case only show the projection of any dipole component, so the one-dimensional fit is sufficient to study the energy dependence of the dominant dipole.

The red data points in Fig. 5.13 are based on the IceTop data. While the phase agrees well with that of the IceCube data at similar energies, the amplitude of the anisotropy is larger for the IceTop data than for any IceCube energy bin. Given the large uncertainty on the relative intensity and energy, the discrepancy is currently not statistically significant. If additional data confirms this discrepancy, a possible explanation could be the different chemical composition of the IceCube and IceTop data sets. Table 5.4 shows the relative composition of cosmic rays detected in IceCube and IceTop according to simulation, based on a primary cosmic-ray composition according to the model by Hörandel [104]. For IceCube, we list the composition for all nine energy bins. Elements are grouped in four main categories with increasing mass number as described in the caption. The simulation indicates that the data set recorded by IceTop is composed of 34% protons and 12% heavy elements.




Figure 5.13 Amplitude (*left*) and ϕ -component of phase (*right*) of the dipole moment of the projection of the relative intensity map onto right ascension for the nine energy bins of IceCube (*blue*) and for IceTop (*red*). The projections were fit to the full set of harmonic functions, but only the dipole is reported here. Data points indicate the median energy of each energy bin, with error bars on the energy showing the 68% containment interval.

At a comparable median energy, in the second-highest energy bin, the data set recorded by IceCube is composed of 24% protons and 21% heavy elements. If the anisotropy is predominantly caused by protons, the lighter composition of the IceTop data could lead to a stronger dipole amplitude.

5.4 Composition and Energy Study

A major goal of this analysis is to study spectral differences as a function of sky position, capitalizing on the advanced reconstruction procedures allowed by the IceTop detector. In this section, ShowerLLH is run over 5 years of data, from IT73 to IT81-2014, with all quality cuts applied (see Sec. 3.2.2.2). Thus, for each event, we have a most-likely energy value and composition. Note that the quality cuts limit the zenith range to $\cos\theta_{\text{reco}} \geq 0.8$, so the mask is increased compared to other parts of this analysis.

Figure 5.14 shows the natural log of the Bayes factor ($\ln B_{21}$) in equatorial coordinates. In order to eliminate the natural dependence of the energy spectrum on declination, the Bayes factor is calculated by comparing each pixel's energy spectrum to the average of its

declination band. The result is still dominated by the pole, but, based on the Bayes factor criteria, indicates that the spectra across the sky is consistent with a one-source hypothesis. In an effort to show structure at higher declination angles, the polar region out to 10° is masked in the right plot of Fig. 5.14.

To test a simplified case, the spectrum of the deficit can be compared to the spectrum of the rest of the sky. Figure 5.15 shows that the two spectra are consistent within errors. The resulting Bayes factor is $\ln B_{21} = -0.17$, which is still consistent with a one-source hypothesis.

`ShowerLLH` also returns a most-likely composition for each event. Figure 5.16 shows the relative intensity (*left*) and significance (*right*) in equatorial coordinates for events whose most likely composition is proton, helium, oxygen, and iron (*top to bottom*). Figure 5.17 follows the same format, but an additional cut has been applied, requiring $\log_{10}(E/\text{GeV}) > 6.2$ for the minimum reconstructed energy. In all of these maps the structure is fairly consistent, with the deficit from $30^\circ - 120^\circ$ as the most significant structure in the sky. Table 5.5 lists the number of counts for each map.

The likelihood distribution can also be studied by classifying reconstructed cosmic-ray primaries as heavier or lighter. Using the Bayes factor, we can study the distribution of the difference between iron and proton likelihood values ($D = f_{\text{LLH}} - p_{\text{LLH}}$) returned by `ShowerLLH`. The upper plots of Fig. 5.18 show $\ln B_{21}$ for each pixel on the sky as compared to the average of its declination band without (*left*) and with (*right*) a mask applied at the pole. Similar to Fig. 5.14, top-hat smoothing with an angular radius of 20° was applied. These maps are also consistent with a one-source hypothesis. As the upper-right plot shows, there is even less structure than in the energy spectra plots. Even though there is no visible structure, since the Bayes factor just indicates how different two spectra are but not in what way, we can still look at the deviation of the mean likelihood distribution value for each pixel from the average of its declination band. The result, shown in the bottom plot of Fig. 5.18, is a very slight difference that is primarily visible around the edges of the plot, where statistics are lower.

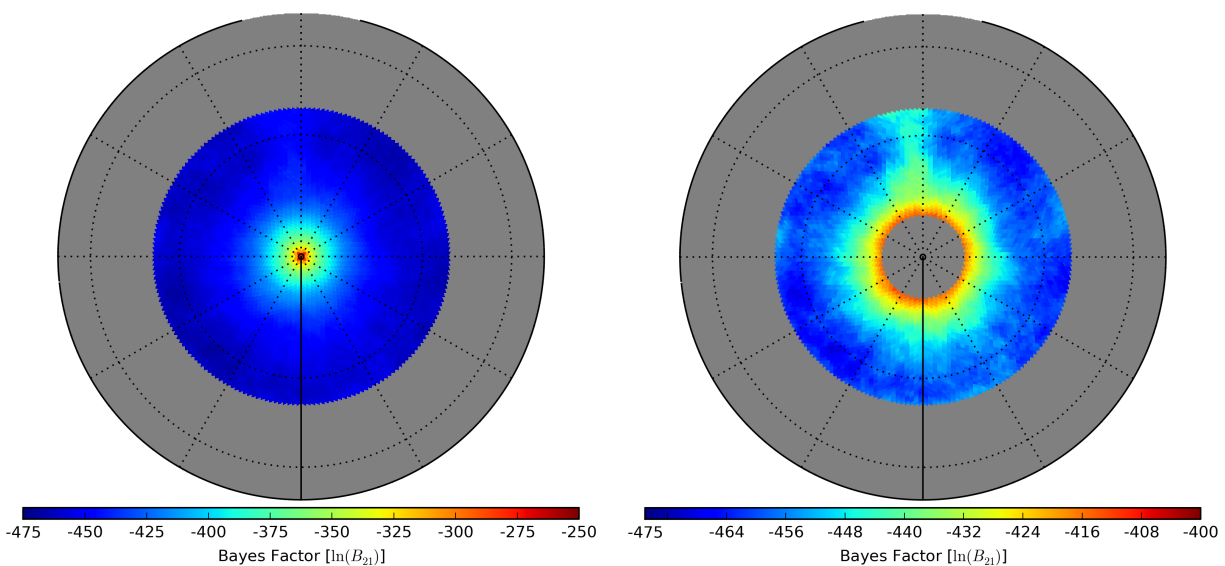


Figure 5.14 Natural log of Bayes factor in equatorial coordinates calculated by comparing each pixel's energy spectrum with its declination-band average. A mask at $\delta = -80^\circ$ is applied in the right plot to show structure outside of the pole. Top-hat smoothing with a 20° angular radius is applied, similar to Sec. 5.3.




Figure 5.15 Normalized energy spectrum for deficit ($\alpha = 60^\circ - \alpha = 120^\circ$), compared to the rest of the sky. A top-hat smoothing with an angular radius of 20° has been applied to all pixels. The resultant value of $\ln B_{21}$ is -0.17 at this sky position, indicating that the spectra are consistent with a one-source hypothesis.

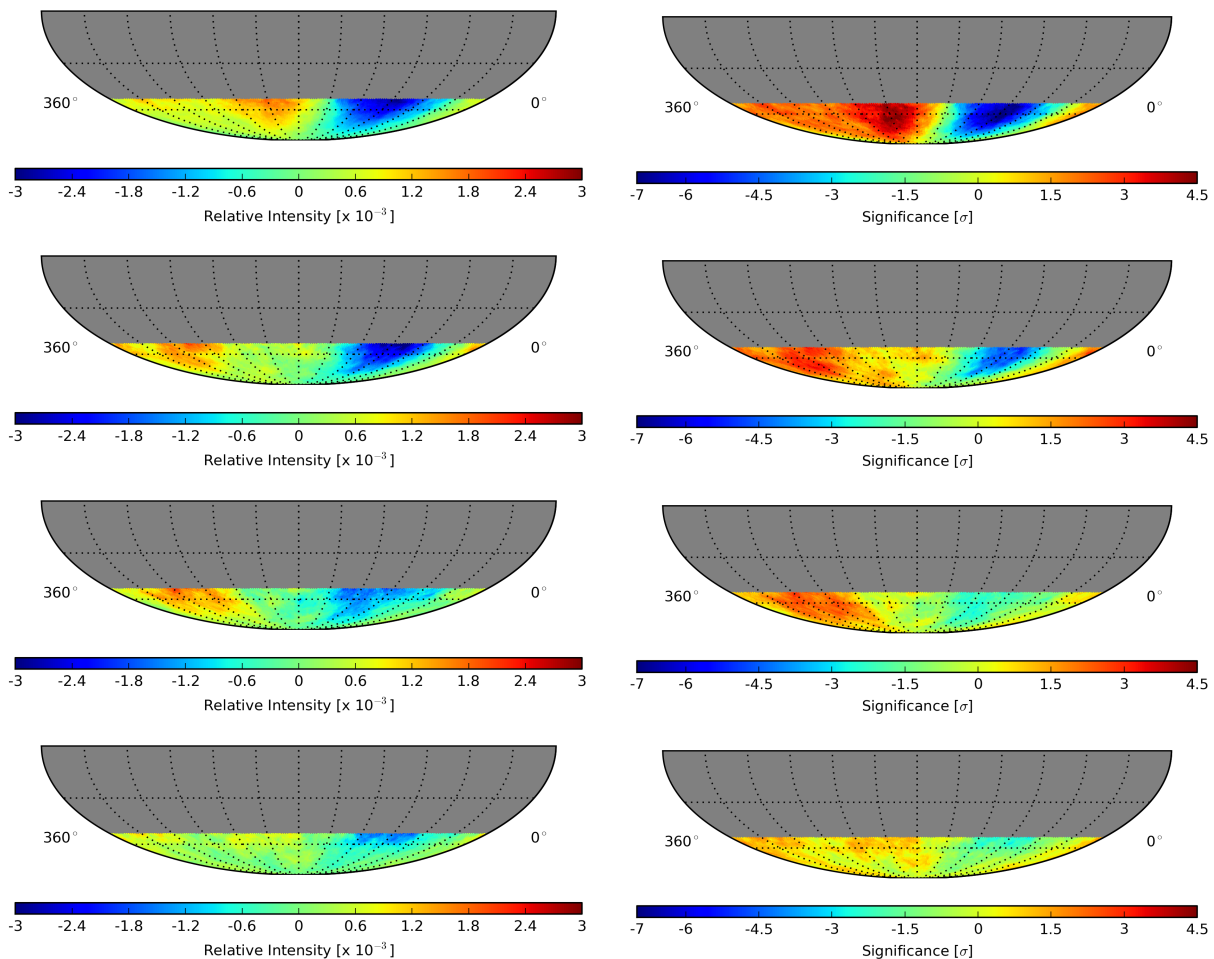


Figure 5.16 Relative intensity (*left*) and pre-trial significance (*right*) in equatorial coordinates for most-likely compositions as returned by ShowerLLH without any cut in reconstructed energy. Shown (*top to bottom*) are proton, helium, oxygen, and iron. Maps have been smoothed with a 20° smoothing radius.

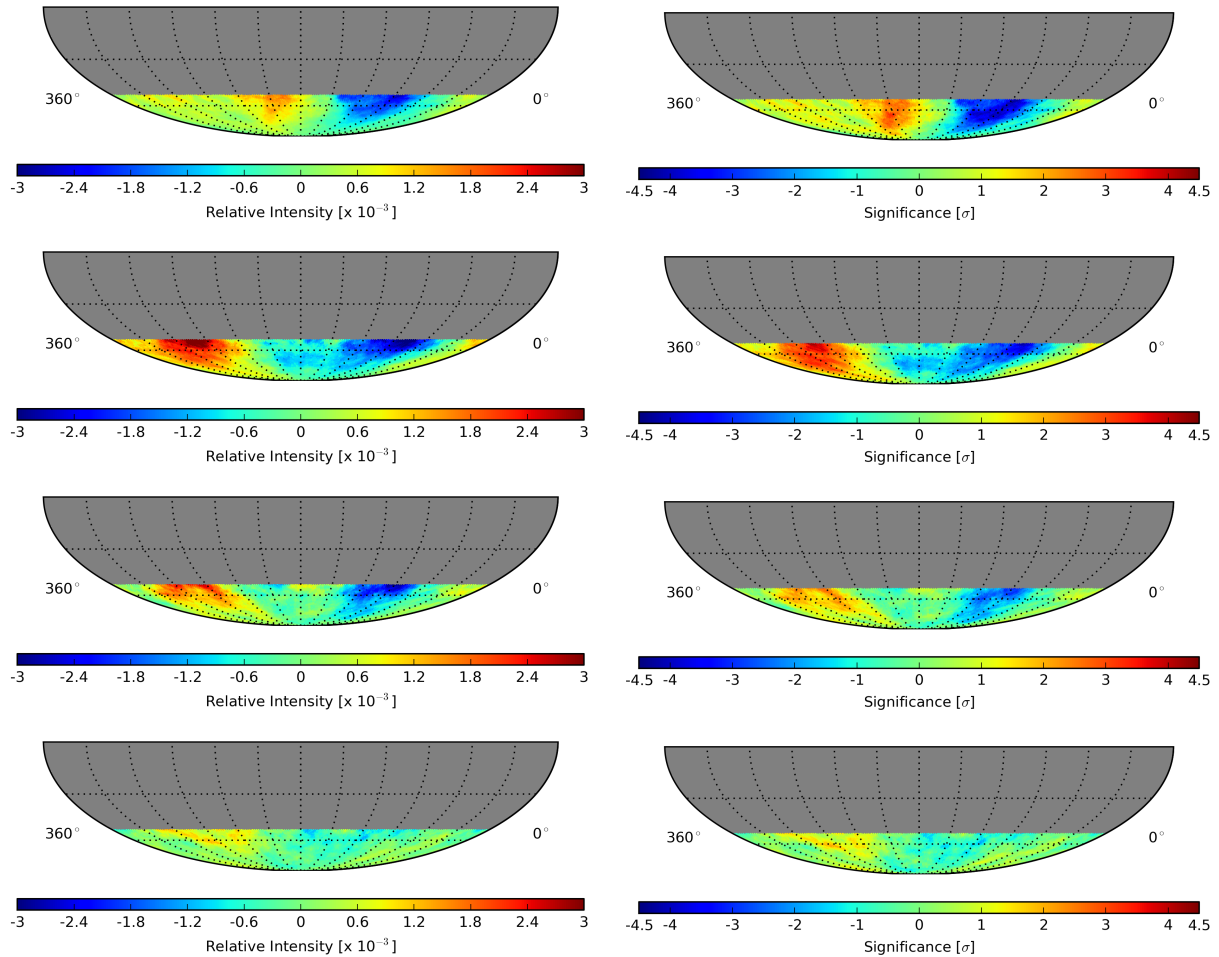


Figure 5.17 Relative intensity (*left*) and pre-trial significance (*right*) in equatorial coordinates for most-likely compositions as returned by ShowerLLH with a minimum reconstructed energy of $\log_{10}(E/\text{GeV}) > 6.2$. Shown (*top to bottom*) are proton, helium, oxygen, and iron. Maps have been smoothed with a 20° smoothing radius.




Figure 5.18 Analysis of the distributions resulting from the difference in iron and proton likelihood values ($D = f_{\text{LLH}} - p_{\text{LLH}}$) as a function of sky position. Bayes factors (*upper-left, upper-right*) calculated by comparing each pixel's energy spectrum with its declination band average. Difference from the mean (*bottom*) calculated by comparing the mean value of the individual pixel's likelihood distribution to its declination band average.

$\log_{10}(E_{\text{median}}/\text{GeV})$	P	He	CNO	Fe
4.12	0.74	0.21	0.04	0.01
4.38	0.70	0.23	0.06	0.01
4.58	0.67	0.25	0.07	0.02
4.85	0.61	0.27	0.09	0.03
5.12	0.54	0.28	0.12	0.05
5.38	0.46	0.29	0.16	0.09
5.77	0.35	0.30	0.21	0.14
6.13	0.24	0.28	0.26	0.21
6.21*	0.34	0.30	0.24	0.12
6.73	0.17	0.18	0.28	0.37

Table 5.4 Chemical composition of each IceCube energy bin and for IceTop as determined by simulation using a primary chemical composition following [104]. Listed is the relative fraction of the composition group in the detected cosmic-ray flux. The composition groups are as follows: P : ${}^1_1\text{H}$, He : ${}^4_2\text{He} - {}^{11}_5\text{B}$, CNO : ${}^{12}_6\text{C} - {}^{39}_{19}\text{K}$, Fe : ${}^{40}_{20}\text{Ca} - {}^{56}_{26}\text{Fe}$. The IceTop bin is marked with an asterisk.

Composition	Counts	Counts ($\log_{10}(E/\text{GeV}) > 6.2$)
Proton	4.53×10^7	2.21×10^7
Helium	2.13×10^7	1.05×10^7
Oxygen	1.53×10^7	8.90×10^6
Iron	1.99×10^7	1.34×10^7
Total	1.02×10^8	5.48×10^7

Table 5.5 Most-likely compositions returned by ShowerLLH and their respective number of events for all years used in this analysis. Counts are reported with and without a minimum energy cut.

Chapter 6

Conclusions

The analysis of 300 billion cosmic-ray events recorded between May 2009 and May 2015 shows anisotropy in the arrival direction distribution consistent with previously published IceCube results [55, 56]. The increased statistics of this data set allow for observation of the small-scale structure at a level approaching the angular resolution. The resulting sky map shows separate structures not resolved in previous analyses, as well as two new regions, an excess and a deficit, observed with high statistical significance.

In addition, a detailed study of the evolution of the anisotropy as a function of energy in the TeV to PeV range shows a strong dependence of the amplitude and the topology of the anisotropy on energy. This analysis extends our previous work [55, 57] and confirms that the anisotropy changes rather dramatically between 130 TeV and 240 TeV; the phase of a best-fit dipole shifts from 50° to around 200° in right ascension. At energies below this shift, the amplitude of the best-fit dipole decreases. Above the shift, it increases again, up to the highest energies currently accessible to IceCube.

A study of the time dependence of the global anisotropy and the small-scale structure over the six-year period covered by this analysis reveals no significant change with time. This result is consistent with previous studies in the Northern and Southern Hemispheres [46, 67, 103].

A search for differences in the energy spectrum as a function of sky position was performed using a detailed energy reconstruction procedure with IceTop, revealing a spectrum

consistent across declination bands. In addition, the likelihood values for different composition assumptions returned by the reconstruction were used to search for a composition-dependence as a function of sky position. This search also indicated consistency across declination bands, with no significant areas of heavier or lighter composition.

The source of the cosmic-ray anisotropy remains unknown. The large-scale anisotropy may be qualitatively explained on the basis of homogeneous and isotropic diffusive propagation of cosmic rays in the Milky Way from stochastically distributed sources. Such discrete sources are responsible for a density gradient of cosmic rays, which causes a dipole anisotropy. Numerical studies show that it is possible to find a given realization of Galactic source distribution that explains the observed non-monotonic energy dependence of the anisotropy amplitude, although for an ensemble of realizations the mean amplitude is overestimated compared to observations [68, 105, 106, 107, 108].

It is known that transport of cosmic rays in magnetic fields is anisotropic, even for large magnetic perturbations compared to the regular mean field [109, 110, 111, 112, 113] and that propagation perpendicular to the local magnetic field direction is slower than in the parallel direction. The misalignment between the regular magnetic field and the cosmic-ray density gradient decreases perpendicular diffusion. This makes it possible to explain the observed smaller amplitudes of the anisotropy [114, 115, 69], but it also means that the anisotropy does not point in the direction of any particular nearby source.

The fact that cosmic-ray anisotropy is not a simple dipole, nor well fit solely by lower-multipole terms, seems to suggest that other transport processes might be important as well. For instance, drift diffusion driven by a gradient of cosmic-ray density in the local interstellar medium, producing a bidirectional anisotropy, was considered by [58] and [116]. Since cosmic rays up to 10 TeV energy scale have a gyroradius approaching the thickness of the heliosphere (see e.g. [117, 118]) an effect from heliospheric magnetic fields should be expected at some degree. Such an effect may be responsible for a redistribution of TeV cosmic ray arrival direction distribution [119, 120, 121].

The small-scale angular features in the cosmic-ray anisotropy may be produced by propagation effects in the local magnetic field. Resonant scattering processes in heliospheric magnetic perturbations or particle re-acceleration in the inner tail from turbulent magnetic reconnection [72, 73] might explain the harder energy spectrum apparent in some localized excess regions. On the other hand, such an observation is also compatible with cosmic rays interacting with interstellar magnetic turbulence, that produce perturbations in the arrival direction distribution of an anisotropic distribution of cosmic-ray particles within the scattering mean free path. Such perturbations may be observed as stochastic localized excess or deficit regions [122, 123], and the corresponding angular power spectrum can be analytically predicted from Liouville's Theorem [70, 71], or simply as an effect of global cosmic ray redistribution due to escape from the propagation region [124].

In the near future, additional experimental information on the anisotropy will continue to come from detectors with good energy resolution and ability to determine the chemical composition of the cosmic-ray flux, such as the IceTop air-shower array. While this analysis did not find evidence of a difference in the energy spectra or composition of different regions in the sky at PeV energies, the reconstruction method used was new and refinements can be made. With its current levels of detector stability, there will hopefully be many more years of IceTop data that can help shed light on this mystery.

LIST OF REFERENCES

- [1] V. F. Hess. Über Beobachtungen der durchdringenden Strahlung bei sieben Freiballonfahrten. *Phys. Z.*, 13:1084–1091, 1912.
- [2] R. Abbasi et al. All-Particle Cosmic Ray Energy Spectrum measured with 26 IceTop Stations. *Astropart. Phys.*, 44(0):40 – 58, 2013.
doi:<http://dx.doi.org/10.1016/j.astropartphys.2013.01.016>.
- [3] J. Abraham et al. Measurement of the Energy Spectrum of Cosmic Rays above 10^{18} eV using the Pierre Auger Observatory. *Phys. Lett. B*, 685:239–246, March 2010.
arXiv:1002.1975, doi:10.1016/j.physletb.2010.02.013.
- [4] D. J. Bird et al. Detection of a Cosmic Ray with Measured Energy well beyond the Expected Spectral Cutoff due to Cosmic Microwave Radiation. *Astrophysical Journal*, 441:144–150, March 1995. arXiv:astro-ph/9410067, doi:10.1086/175344.
- [5] J. J. Beatty and S. Westerhoff. The Highest-Energy Cosmic Rays. *Annu. Rev. Astron. Astrophys.*, 59:319–345, November 2009.
doi:10.1146/annurev.nucl.58.110707.171154.
- [6] K. Greisen. End to the Cosmic-Ray Spectrum? *Physical Review Letters*, 16:748–750, April 1966. doi:10.1103/PhysRevLett.16.748.
- [7] G. T. Zatsepin and V. A. Kuz'min. Upper Limit of the Spectrum of Cosmic Rays. *Soviet J. Exp. Theor. Phys.*, 4:78, August 1966.
- [8] J. Abraham et al. Observation of the Suppression of the Flux of Cosmic Rays above 4×10^{19} eV. *Physical Review Letters*, 101(6):061101, August 2008.
arXiv:0806.4302, doi:10.1103/PhysRevLett.101.061101.
- [9] X. H. Sun, W. Reich, A. Waelkens, and T. A. Enßlin. Radio Observational Constraints on Galactic 3D-Emission Models. *Astron. Astrophys.*, 477:573–592, January 2008. arXiv:0711.1572, doi:10.1051/0004-6361:20078671.
- [10] 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.

- [11] K. A. Olive et al. The Review of Particle Physics. *Chin. Phys. C*, 38(090001), 2014.
- [12] E. C. Stone et al. The Cosmic-Ray Isotope Spectrometer for the Advanced Composition Explorer. *Space Science Reviews*, 86:285–356, July 1998. doi:10.1023/A:1005075813033.
- [13] T. Sanuki et al. Precise Measurement of Cosmic-Ray Proton and Helium Spectra with the BESS Spectrometer. *Astrophysical Journal*, 545:1135–1142, December 2000. arXiv:astro-ph/0002481, doi:10.1086/317873.
- [14] K. Lodders. Solar System Abundances and Condensation Temperatures of the Elements. *Astrophysical Journal*, 591:1220–1247, July 2003. doi:10.1086/375492.
- [15] D. Ikeda et al. Recent Results of the Energy Spectrum and Mass Composition from Telescope Array Fluorescence Detector. *J. Phys. Conf. Ser.*, 409(1):012097, 2013.
- [16] J. Abraham et al. Measurement of the Depth of Maximum of Extensive Air Showers above 10^{18} eV. *Physical Review Letters*, 104(9):091101, March 2010. arXiv:1002.0699, doi:10.1103/PhysRevLett.104.091101.
- [17] 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*, 42:15–32, February 2013. arXiv:1207.3455, doi:10.1016/j.astropartphys.2012.11.003.
- [18] K.-H. Kampert and M. Unger. Measurements of the Cosmic Ray Composition with Air Shower Experiments. *Astroparticle Physics*, 35:660–678, May 2012. arXiv:1201.0018, doi:10.1016/j.astropartphys.2012.02.004.
- [19] E. Fermi. On the Origin of the Cosmic Radiation. *Phys. Rev.*, 75:1169–1174, Apr 1949. doi:10.1103/PhysRev.75.1169.
- [20] J. Abraham et al. Upper Limit on the Cosmic-Ray Photon Fraction at EeV Energies from the Pierre Auger Observatory. *Astropart. Phys.*, 31:399–406, July 2009. arXiv:0903.1127, doi:10.1016/j.astropartphys.2009.04.003.
- [21] A. M. Hillas. The Origin of Ultra-High-Energy Cosmic Rays. *Annu. Rev. Astron. Astrophys.*, 22:425–444, 1984. doi:10.1146/annurev.aa.22.090184.002233.
- [22] F. Fraschetti. On the acceleration of Ultra-High-Energy Cosmic Rays. *Royal Society of London Philosophical Transactions Series A*, 366:4417–4428, December 2008. arXiv:0809.3057, doi:10.1098/rsta.2008.0204.
- [23] J. Abraham et al. Properties and Performance of the Prototype Instrument for the Pierre Auger Observatory. *Nucl. Instrum. Methods A*, 523:50–95, May 2004. doi:10.1016/j.nima.2003.12.012.

- [24] L. A. Anchordoqui and T. Montaruli. In Search of Extraterrestrial High-Energy Neutrinos. *Annu. Rev. Nucl. Part. Sci.*, 60:129–162, November 2010. arXiv:0912.1035, doi:10.1146/annurev.nucl.012809.104551.
- [25] D. Paneque. Experimental Gamma-Ray Astronomy. *J. Phys. Conf. Ser.*, 375(5):052020, July 2012. doi:10.1088/1742-6596/375/1/052020.
- [26] 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.
- [27] 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.
- [28] T. Stanev, P. L. Biermann, and T. K. Gaisser. Cosmic rays. IV. The spectrum and chemical composition above 10^4 GeV. *Astron. Astrophys.*, 274:902, July 1993. arXiv:astro-ph/9303006.
- [29] 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. *Astrophysical Journal Letters*, 592:L71–L73, August 2003. arXiv:astro-ph/0306462, doi:10.1086/377682.
- [30] 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.
- [31] R. Abbasi et al. An Absence of Neutrinos Associated with Cosmic-Ray Acceleration in γ -Ray Bursts. *Nature*, 484:351–354, April 2012. arXiv:1204.4219, doi:10.1038/nature11068.
- [32] A. W. Strong, I. V. Moskalenko, and V. S. Ptuskin. Cosmic-Ray Propagation and Interactions in the Galaxy. *Annu. Rev. Nucl. Part. Sci.*, 57:285–327, November 2007. arXiv:astro-ph/0701517, doi:10.1146/annurev.nucl.57.090506.123011.
- [33] V. S. Berezhinskiĭ, S. V. Bulanov, V. A. Dogiel, V. S. Ptuskin, and V.L. Ginzburg. *Astrophysics of Cosmic Rays*. North-Holland, 1990.
- [34] T. K. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990.
- [35] T. Stanev. *High Energy Cosmic Rays*. Springer-Praxis books in astrophysics and astronomy. Springer, 2004.
- [36] V. L. Ginzburg and S. I. Syrovatskii. *The Origin of Cosmic Rays*. Gordon and Breach, 1964.

- [37] A. W. Strong et al. The GALPROP Cosmic-Ray Propagation Code. In *Proc. of the 31st ICRC*, Łódź, Poland, July 2009. ID 0902. [arXiv:0907.0559](#).
- [38] K. Nagashima, K. Fujimoto, and R. M. Jacklyn. Galactic and heliotail-in anisotropies of cosmic rays as the origin of sidereal daily variation in the energy region $< 10^4$ gev. *Journal of Geophysical Research*, 103:17429–17440, August 1998. doi:10.1029/98JA01105.
- [39] D. L. Hall et al. Gaussian analysis of two hemisphere observations of galactic cosmic ray sidereal anisotropies. *Journal of Geophysical Research*, 104:6737–6750, April 1999. doi:10.1029/1998JA900107.
- [40] K. Munakata et al. Solar Cycle Dependence of the Diurnal Anisotropy of 0.6 TeV Cosmic-ray Intensity Observed with the Matsushiro Underground Muon Detector. *Astrophysical Journal*, 712:1100–1106, April 2010. [arXiv:0911.1165](#), doi:10.1088/0004-637X/712/2/1100.
- [41] K. Munakata, M. Kozai, C. Kato, and J. Kóta. Long-term variation of the solar diurnal anisotropy of galactic cosmic rays observed with the nagoya multi-directional muon detector. *The Astrophysical Journal*, 791(1):22, 2014. URL: <http://stacks.iop.org/0004-637X/791/i=1/a=22>.
- [42] M. Amenomori et al. Large-Scale Sidereal Anisotropy of Galactic Cosmic-Ray Intensity Observed by the Tibet Air Shower Array. *Astrophysical Journal Letters*, 626:L29–L32, June 2005. [arXiv:astro-ph/0505114](#), doi:10.1086/431582.
- [43] M. Amenomori et al. Anisotropy and Corotation of Galactic Cosmic Rays. *Science*, 314:439–443, October 2006. [arXiv:astro-ph/0610671](#), doi:10.1126/science.1131702.
- [44] G. Guillian et al. Observation of the Anisotropy of 10 TeV Primary Cosmic Ray Nuclei Flux with the Super-Kamiokande-I Detector. *Physical Review D*, 75(6):062003, March 2007. [arXiv:astro-ph/0508468](#), doi:10.1103/PhysRevD.75.062003.
- [45] A.A. Abdo et al. The Large-Scale Cosmic-Ray Anisotropy as Observed with Milagro. *Astrophysical Journal*, 698:2121–2130, June 2009. [arXiv:0806.2293](#), doi:10.1088/0004-637X/698/2/2121.
- [46] B. Bartoli et al. Medium scale anisotropy in the TeV cosmic ray flux observed by ARGO-YBJ. *Physical Review D*, 88(8):082001, October 2013. doi:10.1103/PhysRevD.88.082001.
- [47] B. Bartoli et al. Argo-ybj observation of the large-scale cosmic ray anisotropy during the solar minimum between cycles 23 and 24. *The Astrophysical Journal*, 809(1):90, 2015. URL: <http://stacks.iop.org/0004-637X/809/i=1/a=90>.

- [48] T. Saito et al. Time Dependence of Loss-Cone Amplitude measured with the Tibet Air-Shower Array. *International Cosmic Ray Conference*, 1:62, 2011. doi:10.7529/ICRC2011/V01/0379.
- [49] Z. Feng et al. Northern sky Galactic Cosmic Ray anisotropy between 10-1000 TeV with the Tibet Air Shower Array. *International Cosmic Ray Conference*, 2015.
- [50] J. De Jong et al. Observations of Large Scale Sidereal Anisotropy in 1 and 11 TeV cosmic rays from the MINOS experiment. *International Cosmic Ray Conference*, 4:46, 2011. arXiv:1201.2621, doi:10.7529/ICRC2011/V04/1185.
- [51] A.U. Abeysekara et al. Observation of Small-scale Anisotropy in the Arrival Direction Distribution of TeV Cosmic Rays with HAWC. *Astrophysical Journal*, 796:108, December 2014. arXiv:1408.4805, doi:10.1088/0004-637X/796/2/108.
- [52] M. Aglietta et al. Evolution of the Cosmic-Ray Anisotropy Above 10^{14} eV. *Astrophysical Journal Letters*, 692:L130–L133, February 2009. arXiv:0901.2740, doi:10.1088/0004-637X/692/2/L130.
- [53] A. Chiavassa et al. Studies of the cosmic ray spectrum and large scale anisotropies with the KASCADE-Grande experiment. *Journal of Physics Conference Series*, 531(1):012001, August 2014. doi:10.1088/1742-6596/531/1/012001.
- [54] R.U. Abbasi et al. Measurement of the Anisotropy of Cosmic-ray Arrival Directions with IceCube. *Astrophysical Journal Letters*, 718:L194–L198, August 2010. arXiv:1005.2960, doi:10.1088/2041-8205/718/2/L194.
- [55] R.U. Abbasi et al. Observation of Anisotropy in the Arrival Directions of Galactic Cosmic Rays at Multiple Angular Scales with IceCube. *Astrophysical Journal*, 740:16, October 2011. arXiv:1105.2326, doi:10.1088/0004-637X/740/1/16.
- [56] R.U. Abbasi et al. Observation of Anisotropy in the Galactic Cosmic-Ray Arrival Directions at 400 TeV with IceCube. *Astrophysical Journal*, 746:33, February 2012. arXiv:1109.1017, doi:10.1088/0004-637X/746/1/33.
- [57] 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.
- [58] M. Amenomori et al. Implication of the sidereal anisotropy of ~ 5 TeV cosmic ray intensity observed with the Tibet III air shower array. In D. Shaikh and G. P. Zank, editors, *Turbulence and Nonlinear Processes in Astrophysical Plasmas*, volume 932 of *American Institute of Physics Conference Series*, pages 283–289, August 2007. doi:10.1063/1.2778976.

- [59] A.A. Abdo et al. Discovery of Localized Regions of Excess 10-TeV Cosmic Rays. *Physical Review Letters*, 101(22):221101, November 2008. arXiv:0801.3827, doi:10.1103/PhysRevLett.101.221101.
- [60] P. Abreu et al. Search for first harmonic modulation in the right ascension distribution of cosmic rays detected at the Pierre Auger Observatory. *Astropart. Phys.*, 34:627–639, March 2011. arXiv:1103.2721, doi:10.1016/j.astropartphys.2010.12.007.
- [61] P. Abreu et al. Large-scale Distribution of Arrival Directions of Cosmic Rays Detected Above 10^{18} eV at the Pierre Auger Observatory. *Astrophysical Journal Supplements*, 203:34, December 2012. arXiv:1210.3736, doi:10.1088/0067-0049/203/2/34.
- [62] P. Abreu et al. Constraints on the Origin of Cosmic Rays above 10^{18} eV from Large-scale Anisotropy Searches in Data of the Pierre Auger Observatory. *Astrophysical Journal Letters*, 762:L13, January 2013. arXiv:1212.3083, doi:10.1088/2041-8205/762/1/L13.
- [63] R.U. Abbasi et al. Indications of Intermediate-Scale Anisotropy of Cosmic Rays with Energy Greater Than 57 EeV in the Northern Sky Measured with the Surface Detector of the Telescope Array Experiment. *Astrophysical Journal Letters*, 790:L21, 2014. arXiv:1404.5890, doi:10.1088/2041-8205/790/2/L21.
- [64] A. H. Compton and I. A. Getting. An Apparent Effect of Galactic Rotation on the Intensity of Cosmic Rays. *Phys. Rev.*, 47:817–821, Jun 1935. doi:10.1103/PhysRev.47.817.
- [65] S. R. Majewski. Precision Astrometry, Galactic Mergers, Halo Substructure and Local Dark Matter. In W. J. Jin, I. Platais, and M. A. C. Perryman, editors, *IAU Symposium*, volume 248 of *IAU Symposium*, pages 450–457, July 2008. arXiv:0801.4927, doi:10.1017/S1743921308019790.
- [66] L. Gleeson and W. Axford. The Compton-Getting effect. *Astrophys. Space Sci.*, 2:431–437, 1968. doi:10.1007/BF02175919.
- [67] M. Amenomori et al. On Temporal Variations of the Multi-TeV Cosmic Ray Anisotropy Using the Tibet III Air Shower Array. *Astrophysical Journal*, 711:119–124, March 2010. arXiv:1001.2646, doi:10.1088/0004-637X/711/1/119.
- [68] A. D. Erlykin and A. W. Wolfendale. The Anisotropy of Galactic Cosmic Rays as a Product of Stochastic Supernova Explosions. *Astropart. Phys.*, 25:183–194, 2006. arXiv:astro-ph/0601290, doi:10.1016/j.astropartphys.2006.01.003.

- [69] P. Mertsch and S. Funk. Solution to the Cosmic Ray Anisotropy Problem. *Physical Review Letters*, 114(2):021101, January 2015. arXiv:1408.3630, doi:10.1103/PhysRevLett.114.021101.
- [70] M. Ahlers. Anomalous Anisotropies of Cosmic Rays from Turbulent Magnetic Fields. *Physical Review Letters*, 112(2):021101, January 2014. arXiv:1310.5712, doi:10.1103/PhysRevLett.112.021101.
- [71] M. Ahlers and P. Mertsch. Small-Scale Anisotropies of Cosmic Rays from Relative Diffusion. *ArXiv e-prints*, June 2015. arXiv:1506.05488.
- [72] A. Lazarian and P. Desiati. Magnetic Reconnection as the Cause of Cosmic Ray Excess from the Heliospheric Tail. *Astrophysical Journal*, 722:188–196, October 2010. arXiv:1008.1981, doi:10.1088/0004-637X/722/1/188.
- [73] P. Desiati and A. Lazarian. Cosmic rays and stochastic magnetic reconnection in the heliotail. *Nonlinear Processes in Geophysics*, 19:351–364, June 2012. arXiv:1205.3783, doi:10.5194/npg-19-351-2012.
- [74] W.-M. Yao, C. Amsler, D. Asner, et al. Review of Particle Physics. *Journal of Physics G*, 33:1+, 2006. URL: <http://pdg.lbl.gov>.
- [75] D. Chirkin and W. Rhode. Propagating Leptons through Matter with Muon Monte Carlo (MMC). *ArXiv e-prints*, July 2004. arXiv:hep-ph/0407075.
- [76] <http://skullsinthestars.com/2009/11/20/reversing-optical-shockwaves-using-metamaterials/>.
- [77] A. Achterberg et al. First year performance of the IceCube neutrino telescope. *Astropart. Phys.*, 26:155–173, October 2006. arXiv:astro-ph/0604450, doi:10.1016/j.astropartphys.2006.06.007.
- [78] R. Abbasi et al. The Design and Performance of IceCube DeepCore. *Astropart. Phys.*, 35:615–624, May 2012. arXiv:1109.6096, doi:10.1016/j.astropartphys.2012.01.004.
- [79] E. Andres et al. The AMANDA Neutrino Telescope: Principle of Operation and First Results. *Astropart. Phys.*, 13:1–20, March 2000. arXiv:astro-ph/9906203, doi:10.1016/S0927-6505(99)00092-4.
- [80] R. Abbasi et al. The IceCube Data Acquisition System: Signal Capture, Digitization, and Timestamping. *Nucl. Instrum. Methods A*, 601:294–316, 2009. arXiv:0810.4930, doi:10.1016/j.nima.2009.01.001.
- [81] R. Abbasi et al. Calibration and Characterization of the IceCube Photomultiplier Tube. *Nucl. Instrum. Methods A*, 618:139–152, 2010. arXiv:1002.2442, doi:10.1016/j.nima.2010.03.102.

- [82] M. G. Aartsen et al. Measurement of South Pole Ice Transparency with the IceCube LED Calibration System. *Nucl. Instrum. Methods A*, 711(21):73 – 89, May 2013. arXiv:1301.5361, doi:10.1016/j.nima.2013.01.054.
- [83] M. G. Aartsen et al. South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy. *Journal of Glaciology*, 59:1117–1128, 2013. doi:10.3189/2013JoG13J068.
- [84] D. Heck et al. CORSIKA: A Monte Carlo Code to Simulate Extensive Air Showers. *Forschungszentrum Karlsruhe Report FZKA 6019*, 1998. URL: <http://www-ik.fzk.de/corsika/>.
- [85] R. Abbasi et al. IceTop: The Surface Component of IceCube. *Nucl. Instrum. Methods A*, 700:188–220, February 2013. arXiv:1207.6326, doi:10.1016/j.nima.2012.10.067.
- [86] D. Pandel. Bestimmung von Wasser- und Detektorparametern und Rekonstruktion von Myonen bis 100 TeV mit dem Baikal-Neutrinoobservatorium NT-72. Diploma Thesis, Humboldt-Universität zu Berlin, 1996.
- [87] D. Chirkin. New Implementation of Convolved MPE Pandel Function. *Technical Report icecube/201101006*, January 2011. arXiv:http://internal.icecube.wisc.edu/reports/data/icecube/2011/01/006/icecube_201101006_v1.pdf.
- [88] 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, doi:10.1016/j.nima.2004.01.065.
- [89] G. Japaridze and M. Ribordy. Realistic arrival time distribution from an isotropic light source. *ArXiv Astrophysics e-prints*, June 2005. arXiv:astro-ph/0506136.
- [90] M. G. Aartsen, R. Abbasi, Y. Abdou, 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.
- [91] M. Belloni and T. Timberlake. Available at <http://ftparmy.com/135137-equatorial-coordinates-model.html>.
- [92] P. Wallace. The SLALIB library. *ASP Conf. Ser.*, 61:481, 1994. URL: <http://star-www.rl.ac.uk/star/docs/sun67.htx/sun67.html>.
- [93] F. J. M Farley and J. R. Storey. The Sidereal Correlation of Extensive Air Showers. *Proc. Phys. Soc. Sect. A*, 67(11):996, 1954. URL: <http://stacks.iop.org/0370-1298/67/i=11/a=306>.

- [94] D.E. Alexandreas et al. Point Source Search Techniques in Ultra High Energy Gamma Ray Astronomy. *Nucl. Instr. Meth. Phys. Res. A*, 328:570–577, May 1993. doi:10.1016/0168-9002(93)90677-A.
- [95] M. Santander. Observation of Cosmic-Ray Anisotropy at TeV and PeV Energies in the Southern Sky. *PhD Thesis, University of Wisconsin-Madison*, 2013.
- [96] R. Iuppa and G. Di Sciascio. Time-Average-Based Methods for Multi-angular Scale Analysis of Cosmic-Ray Data. *Astrophysical Journal*, 766:96, April 2013. arXiv:1301.1833, doi:10.1088/0004-637X/766/2/96.
- [97] K. M. Górski et al. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *Astrophysical Journal*, 622:759–771, April 2005. arXiv:astro-ph/0409513, doi:10.1086/427976.
- [98] T.-P. Li and Y.-Q. Ma. Analysis methods for results in gamma-ray astronomy. *Astrophysical Journal*, 272:317–324, September 1983. doi:10.1086/161295.
- [99] R. Ansari and C. Magneville. Partial CMB maps: bias removal and optimal binning of the angular power spectrum. *Monthly Notices of the Royal Astronomical Society*, 405:1421–1430, July 2010. arXiv:0910.4623, doi:10.1111/j.1365-2966.2010.16561.x.
- [100] I. Szapudi, S. Prunet, D. Pogosyan, A. S. Szalay, and J. R. Bond. Fast Cosmic Microwave Background Analyses via Correlation Functions. *Astrophysical Journal Letters*, 548:L115–L118, February 2001. doi:10.1086/319105.
- [101] G. Chon, A. Challinor, S. Prunet, E. Hivon, and I. Szapudi. Fast estimation of polarization power spectra using correlation functions. *Monthly Notices of the Royal Astronomical Society*, 350:914–926, May 2004. arXiv:astro-ph/0303414, doi:10.1111/j.1365-2966.2004.07737.x.
- [102] S. Y. BenZvi, B. M. Connolly, C. G. Pfendner, and S. Westerhoff. A Bayesian Approach to Comparing Cosmic Ray Energy Spectra. *Astrophysical Journal*, 738:82, September 2011. arXiv:1106.1392, doi:10.1088/0004-637X/738/1/82.
- [103] M. G. Aartsen et al. The IceCube Neutrino Observatory Part III: Cosmic Rays. *International Cosmic Ray Conference*, September 2013. arXiv:1309.7006.
- [104] Joerg R. Hörandel. On the knee in the energy spectrum of cosmic rays. *Astropart. Phys.*, 19:193–220, 2003. arXiv:astro-ph/0210453, doi:10.1016/S0927-6505(02)00198-6.
- [105] P. Blasi and E. Amato. Diffusive propagation of cosmic rays from supernova remnants in the Galaxy. II: anisotropy. *J. Cosmol. Astropart. Phys.*, 1:11, January 2012. arXiv:1105.4529, doi:10.1088/1475-7516/2012/01/011.

- [106] V. Ptuskin. Propagation of galactic cosmic rays. *Astropart. Phys.*, 39:44–51, December 2012. doi:10.1016/j.astropartphys.2011.11.004.
- [107] M. Pohl and D. Eichler. Understanding TeV-band Cosmic-Ray Anisotropy. *Astrophysical Journal*, 766:4, March 2013. arXiv:1208.5338, doi:10.1088/0004-637X/766/1/4.
- [108] L. G. Sveshnikova, O. N. Strelnikova, and V. S. Ptuskin. Spectrum and anisotropy of cosmic rays at TeV-PeV-energies and contribution of nearby sources. *Astropart. Phys.*, 50:33–46, December 2013. arXiv:1301.2028, doi:10.1016/j.astropartphys.2013.08.007.
- [109] J. Giacalone and J. R. Jokipii. The Transport of Cosmic Rays across a Turbulent Magnetic Field. *Astrophysical Journal*, 520:204, 1999.
- [110] A. Shalchi. *Nonlinear Cosmic Ray Diffusion Theories*, volume 362. Springer-Verlag Berlin Heidelberg, 2009.
- [111] R. C. Tautz. A Note on Perpendicular Scattering Lengths. *Astrophysical Journal*, 703:1294, 2009.
- [112] Paolo Desiati and Ellen G. Zweibel. The Transport of Cosmic Rays Across Magnetic Fieldlines. *Astrophysical Journal*, 791(1):51, 2014. arXiv:1402.1475, doi:10.1088/0004-637X/791/1/51.
- [113] A. Shalchi. Perpendicular diffusion of energetic particles in collisionless plasmas. *Physics of Plasmas*, 22:010704, 2015.
- [114] F. Effenberger, H. Fichtner, K. Scherer, and I. Büsching. Anisotropic diffusion of Galactic cosmic ray protons and their steady-state azimuthal distribution. *Astron. Astrophys.*, 547:A120, November 2012. arXiv:1210.1423, doi:10.1051/0004-6361/201220203.
- [115] R. Kumar and D. Eichler. Large-scale Anisotropy of TeV-band Cosmic Rays. *Astrophysical Journal*, 785:129, April 2014. doi:10.1088/0004-637X/785/2/129.
- [116] Y. Mizoguchi, K. Munakata, M. Takita, and J. Kota. The sidereal anisotropy of multi-TeV cosmic rays in an expanding Local Interstellar Cloud. *ArXiv e-prints*, September 2009. arXiv:0909.1029.
- [117] N. V. Pogorelov, G. P. Zank, and T. Ogino. Three-dimensional Features of the Outer Heliosphere due to Coupling between the Interstellar and Interplanetary Magnetic Fields. II. The Presence of Neutral Hydrogen Atoms. *Astrophysical Journal*, 644:1299–1316, June 2006. doi:10.1086/503703.

- [118] N. V. Pogorelov, S. N. Borovikov, G. P. Zank, and T. Ogino. Three-Dimensional Features of the Outer Heliosphere Due to Coupling Between the Interstellar and Interplanetary Magnetic Fields. III. The Effects of Solar Rotation and Activity Cycle. *Astrophysical Journal*, 696:1478–1490, May 2009. doi:10.1088/0004-637X/696/2/1478.
- [119] P. Desiati and A. Lazarian. Anisotropy of TeV Cosmic Rays and Outer Heliospheric Boundaries. *Astrophysical Journal*, 762:44, January 2013. arXiv:1111.3075, doi:10.1088/0004-637X/762/1/44.
- [120] N. A. Schwadron et al. Global Anisotropies in TeV Cosmic Rays Related to the Sun’s Local Galactic Environment from IBEX. *Science*, 343:988–990, February 2014. doi:10.1126/science.1245026.
- [121] M. Zhang, P. Zuo, and N. Pogorelov. Heliospheric Influence on the Anisotropy of TeV Cosmic Rays. *Astrophysical Journal*, 790:5, July 2014. doi:10.1088/0004-637X/790/1/5.
- [122] G. Giacinti and G. Sigl. Local Magnetic Turbulence and TeV-PeV Cosmic Ray Anisotropies. *Physical Review Letters*, 109(7):071101, August 2012. arXiv:1111.2536, doi:10.1103/PhysRevLett.109.071101.
- [123] P. L. Biermann, J. Becker Tjus, E.-S. Seo, and M. Mandelartz. Cosmic-Ray Transport and Anisotropies. *Astrophysical Journal*, 768:124, May 2013. arXiv:1206.0828, doi:10.1088/0004-637X/768/2/124.
- [124] R. Rettig and M. Pohl. *International Cosmic Ray Conference*, 2015.