Anisotropy of Arrival Directions of Ultrahigh Energy Cosmic Rays

Chad Barrett Finley

Submitted in partial fulfillment of the
requirements for the degree
of Doctor of Philosophy
in the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY
2006

ABSTRACT

Anisotropy of Arrival Directions of Ultrahigh Energy Cosmic Rays

Chad Barrett Finley

This thesis investigates the origins of ultrahigh energy cosmic rays by searching for evidence of small-scale anisotropy in their arrival directions from space. We begin with a short review of what is known about cosmic rays and specifically about the propagation of charged particles through galactic and extragalactic space. We next describe the experimental techniques of detecting these particles by measuring the extensive air showers they produce in the atmosphere, with special focus on the High Resolution Fly's Eye (HiRes) experiment. In stereo operation the HiRes detector observes cosmic ray arrival directions with unprecedented angular resolution, on the order of 0.5° at all energies above 10^{18} eV, opening a new window in particle astronomy. Before analyzing the data produced by this experiment, we examine previously claimed evidence for clustering of cosmic rays seen by the Akeno Giant Air Shower Array (AGASA) experiment, and identify potential biases in the analysis which call the evidence into question. An unbiased method is then developed to search for the same kind of clustering signal and is applied to the HiRes data. We next develop an unbinned maximum likelihood ratio test for point-source searches, a technique that is optimal for data combined from experiments with different angular resolution, and we apply this to the combined set of HiRes and AGASA data. Finally, we investigate a series of recent claims that ultrahigh energy cosmic rays, including those observed by HiRes, are correlated with the positions of bright BL Lacertae objects (a class of active galaxies), and propose tests of these claims that require independent data. We conclude that at present no statistically significant evidence exists for small-scale anisotropy.

Contents

List of Tables

List of Figures

Acknowledgements

I would like first and foremost to thank my advisor, Stefan Westerhoff, for providing such enthusiastic support and encouragement, and offering both the guidance and the independence that made this work possible. Above all, his ability to identify immediately what is essential and what is not, to see the order implicit in still confused or scattered ideas, improved everything I did and set me moving forward time and again. Working together the past few years has been every bit as personally rewarding an experience as it has been an educational one.

I also extend many, many thanks to my collaborators at Columbia: Segev BenZvi, Brian Connolly, and Andrew O'Neill. They offered an unending supply of valuable ideas, helpful criticisms, and outstanding company day in and day out.

I would like to express my gratitude to Cy Hoffman. The importance of his insights, the improvements gained by his editing, and the generosity of his thought and time expended are recognizable and deserving of acknowledgement on page after page of this thesis.

I convey heartfelt thanks to Glennys Farrar and to Gus Sinnis for their support and for always enjoyable and illuminating discussions (including at my thesis defense).

This work would not have been possible without the tireless effort and dedication of the entire HiRes collaboration. I especially would like to acknowledge Charlie Jui,

Michal Seman, Pierre Sokolsky, Benjamin Stokes, and Gordon Thomson for their assistance and their encouragement.

At Columbia, there are many people who have been essential to the positive character of this department; among them I would like especially to thank: Lalla Grimes, the irreplaceable department administrator; Nelson Rivera and Ray Daguilar, who look after the building and the people in it with great spirit; and Jerome Packer, whose forty-four years of service to the department and exceptional dedication to the students make him not just the nuts and bolts but the heart and soul of the university.

I take this opportunity to recognize and to offer enormous thanks to two professors who have advised and encouraged me since my undergraduate days: Allan Blaer and Troels Jorgensen. They are extraordinary professors whose devotion to students is unsurpassed, and they continue to inspire me today.

It gives me great pleasure to express my gratitude to the following people, whose friendship and support over many years have meant so much: Paul Beddoe-Stephens, Shira and Murat Boss-Bicak, David Fisher, Brendan Foster, Ben Greenbaum, Peter Kamali, Ali Kinkhabwala, Thomas Lecky, Adam Lidz and Sonja Pacho-Lidz, Ben Oppenheimer, Rosemary Pitkin, and Niraj Warikoo.

And most of all, I express my deepest gratitude to my family: Mom, Dad, Gail, Susan, Donna, Lori, Shani, Virginia and Charlie. Their enthusiastic encouragement and support from the beginning have made all the difference.

Chapter 1

Introduction: Ultrahigh Energy Cosmic Rays

1.1 The First Ten Decades

At the beginning of the 20th century, experiments in radioactivity indicated that in addition to the radiation emitted by atoms, there existed another mysterious form of radiation, vastly more penetrating than that from any known substance. The radiation appeared to be everywhere, and no amount of shielding could completely prevent it from reaching electroscopes, the basic instrument of the day. Initial attempts to identify the source included measuring whether its intensity diminished with higher altitude above the surface of Earth, but the measurement was difficult and the results inconclusive. The breakthrough finally came in 1912 when Victor Hess, through careful improvements in instrumental sensitivity, found the opposite was in fact the case. In a series of balloon flights, he discovered that the intensity increased with altitude, and that at an altitude of 5 km it was twice that at sea level. Hess had discovered

that Earth is bathed in a continuous shower of cosmic radiation, with a penetrating power unlike that of any radioactive substance[1].

Far from providing a complete answer, however, the discovery of cosmic radiation transformed the previous question into a potentially harder one: what unknown phenomena in space are responsible for these highly energetic rays? Over the next two decades, scientific speculation and public imagination would run wild. A popular summary of Robert Millikan's view was that cosmic rays are the "birth cries" of new matter being created in the depths of space [2]. Despite its quaint sound, Millikan's hypothesis [3, 4] was based on an appealing physical argument: if atomic nuclei were formed out of electrons and ionized hydrogen (at the time, the only elementary building blocks known), then the binding energy would be released in gamma rays; hence the label "cosmic rays," a label which Millikan himself had introduced. However appealing, the gamma ray hypothesis was refuted by Arthur Compton's data collected from sites around the world, which showed that the cosmic radiation intensity depended on latitude in a way consistent with the passage of charged particles through Earth's magnetic field [5]. Millikan's initial reluctance to accept Compton's conclusions and their public clash at an AAAS meeting was high drama and reported on the front page of the New York Times on Dec. 31, 1932 [2].

By this time, the existence of particles coming from space with energies in excess of several GeV was well accepted. It was also generally realized that the radiation being observed in detectors was a secondary product of the primary, higher energy particle arriving from space, which collides with air molecules in the upper atmosphere. What was unexpected, however, was the discovery in 1938 of primary cosmic rays with energies above $10^{15}\,\mathrm{eV}$ — a million times more energetic than previously detected, and a billion times more energetic than radioactive decays.

Small "air showers" of particles due to an initial, higher energy particle had already been detected by observing coincidences between counters spaced up to several meters apart. To detect larger air showers occurring less frequently, the background rate due to chance coincidences must be reduced, which can be done by improving the resolving time of the coincidence circuit, for example. In 1938, Pierre Auger and colleagues undertook coincidence measurements with counters which now had resolving times of microseconds, rather than milliseconds [6]; to their astonishment, they were able to detect coincidences above the background rate even when the counters were spaced hundreds of meters apart. Auger was able to estimate the number of particles in such showers ($\sim 10^6$) and the energy of the primary particles ($\sim 10^{15}$ eV), and further to show that their rate was consistent with a straightforward extrapolation of the power law spectrum of cosmic rays observed at lower energies [7]. Auger concluded:

> One of the consequences of the extension of the energy spectrum of cosmic rays up to 10^{15} eV is that it is actually impossible to imagine a single process able to give to a particle such an energy. It seems much more likely that the charged particles which constitute the primary cosmic radiation acquire their energy along electric fields of very great extension.[7]

Auger's discovery of extensive air showers initiated the transition to the modern era of high energy cosmic ray studies, in which ever larger detectors have been constructed to study rarer, higher energy cosmic rays by measuring the extensive air showers they produce. His conclusion that such high energies were achieved by acceleration over large distances, rather than emerging at once from subatomic bursts, is also consistent with the general picture of cosmic ray origins today. Enrico Fermi proposed the first basic framework of this picture of acceleration via extended interactions with magnetic fields in 1949 [8].

Meanwhile during the 1930's and 1940's, interest in cosmic rays was not as concerned with uncovering their origins as with the incredible particle physics phenomena they revealed [6]. While laboratory accelerators were still developing, cosmic rays were a convenient source of higher energy particle interactions. The positron (1932), the muon (1937), and the pion (1947) were all discovered by studying particle tracks produced by cosmic rays passing through cloud chambers or photographic plates (see e.g. Samios [9]).

However, as accelerators achieved higher energies, the rapidly falling cosmic ray spectrum could not compete as a particle physics laboratory. Continued research focused once again on the nature of cosmic rays themselves. In 1963, John Linsley reported the first evidence of a cosmic ray primary with energy of 10^{20} eV, observed by the Volcano Ranch array of 19 scintillation detectors [10]. At the time, such an observation seemed inevitable; the cosmic ray spectrum appeared as though it could continue to whatever energy (and correspondingly low flux) scientist had the ability and patience to measure [11]. Other than the symbolic significance of this energy, the scientific import lay mainly in the fact that nothing in our galaxy was thought to be capable of accelerating charged particles to such energies, as Linsley noted [10].

The situation changed dramatically with the discovery in 1965 of the cosmic microwave background (CMB). Almost immediately, Fred Hoyle pointed out that high energy electrons would suffer large energy losses propagating through intergalactic space, due to the inverse Compton effect [12]. When he applied the same calculation to the propagation of protons, however, he found that the energy losses would be negligible even up to primary energies of 10^{21} eV. It fell to Greisen [13] and Zatsepin and Kuzmin [14] to show the following year that the CMB would indeed cause substantial energy losses for such highly energetic protons, not because of Compton scattering but

two other processes: pair creation and, in particular, photopion production. Greisen calculated that the energy loss length for a 2×10^{20} eV proton was ~ 10 Mpc, a distance which encompasses only our nearest galactic neighbors and is far short of the distance to the nearest active galaxies. With respect to a straightforward extrapolation of the measured cosmic-ray spectrum, he predicted that "above 10^{20} eV the primary spectrum will steepen abruptly, and the experiments in preparation will at last observe it to have a cosmologically meaningful termination."[13] The suppression effect was so strong that, in light of the total exposure that cosmic ray experiments had up until that time, Greisen commented, "even the one event recorded at 10^{20} eV appears surprising."[13]

The Greisen-Zatsepin-Kuzmin prediction of strong flux suppression above 10^{20} eV—the so-called GZK effect—brings us fully into the modern era of ultrahigh energy cosmic ray research. A succession of detectors—Haverah Park, Yakutsk, Fly's Eye, AGASA, HiRes, Auger—have been constructed since then to try to determine the flux at the highest energies, the composition, and whether any anisotropy of arrival directions can be observed.

It has been noted [6] that versions of the title "The origin of the highest energy cosmic rays" have been used for journal articles for almost a century now. This would indeed be a sobering and depressing observation, were it not for the fact that the meaning of "highest" has been steadily shifting upward, concealing how much in fact has been learned during this time. In the first decades after Hess's discovery of cosmic radiation, it could not be determined what kind of particle (or "ray") it consisted of. While this is still a question for the highest energy cosmic rays today, the composition of cosmic rays up to 10^{14} eV has been directly measured and found to be a proton-dominated mixture of nuclei. Well-developed models of acceleration

of these particles in supernova blast waves are awaiting confirmation by the latest generation of gamma-ray telescopes, which should be able to detect the high energy secondary photons produced by collisions with ambient matter.

Nevertheless at the highest energies observed today, the basic questions would be familiar to Hess and his contemporaries almost a century ago. What is the composition of these particles? Where do they come from? And most importantly, how did they acquire such tremendous energies?

1.2 Modern Overview

The cosmic ray spectrum is shown in Fig. 1.1. For energies below $10^{10}\,\mathrm{eV}$, the flux observed on Earth varies significantly during the eleven-year solar cycle, with maximum intensity of the solar wind excluding cosmic rays more effectively and resulting in a flux minimum [15, pp. 9-11]. Above this energy, the cosmic ray spectrum spans more than ten orders of magnitude in energy, over which the differential flux is well described by a power law:

$$\frac{dN}{dE} \propto E^{-\alpha}. \tag{1.1}$$

The main feature in the spectrum is a steepening of the spectral index from $\alpha \approx 2.7$ to $\alpha \approx 3$ above $10^{15}\,\mathrm{eV}$, known as the "knee." There also appears to be some flattening again in the spectrum above $10^{18}\,\mathrm{eV}$, known as the "ankle."

Considerable consensus now exists that the lower end of the cosmic ray spectrum can be accounted for by the acceleration of charged particles in supernova blast waves in our galaxy. A small fraction of the supernova kinetic energy going into cosmic ray acceleration can account for both the energy density and spectral index of cosmic rays up to energies $\sim 100\,\mathrm{TeV}$ [15]. Direct confirmation of this hypothesis, however, has

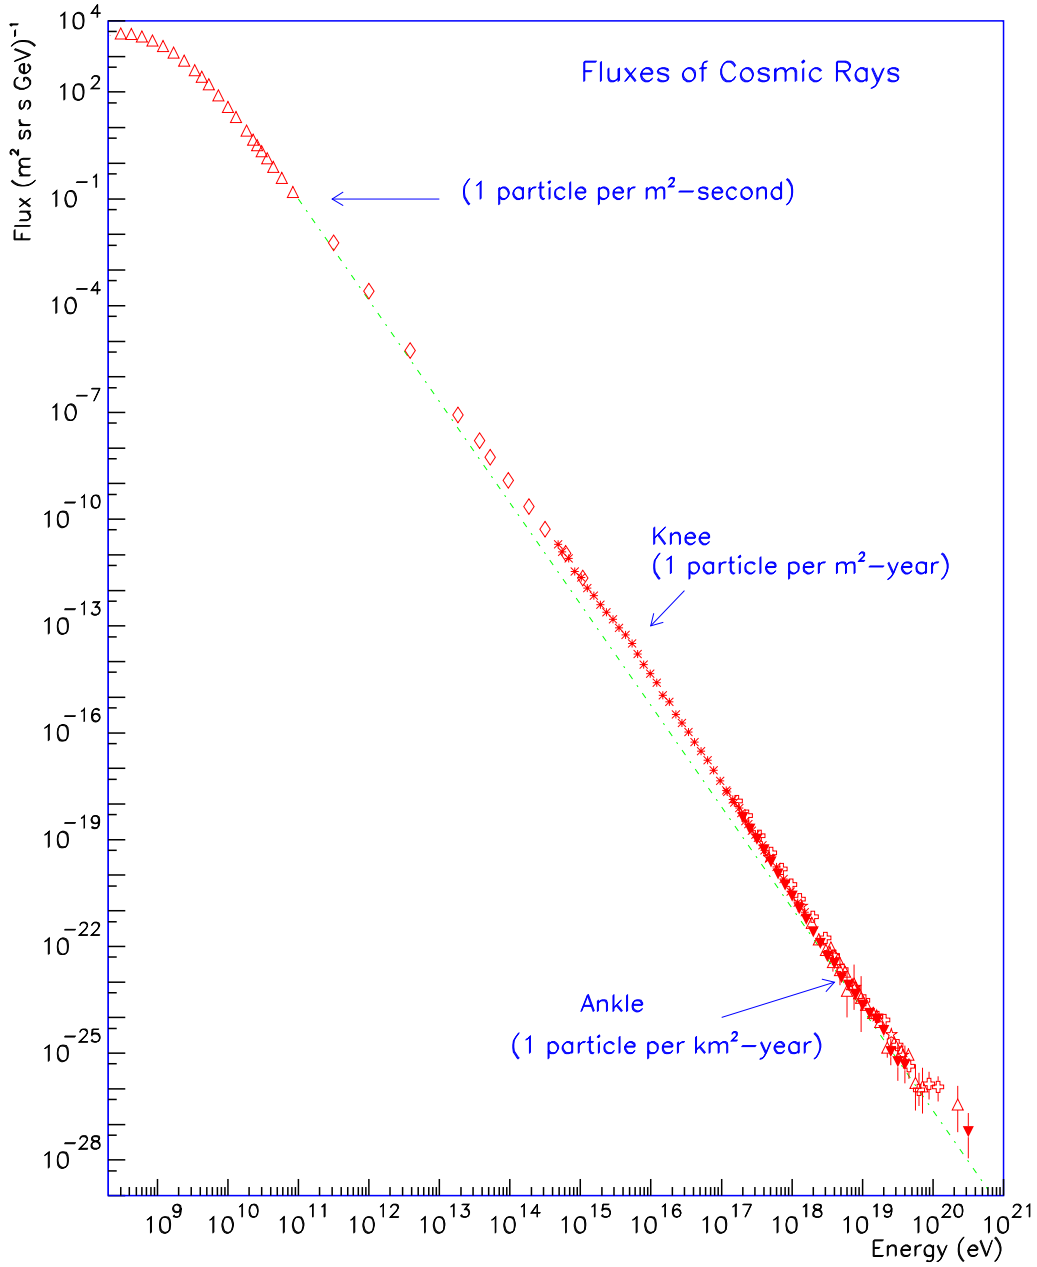

Figure 1.1: Cosmic ray all-particle spectrum. Plot due to S. Swordy, reproduced in Olinto [16].

been difficult. The strongest evidence yet was recently obtained by the High Energy Stereoscopic System (HESS), a high-energy gamma ray telescope which has produced the first resolved image of a supernova remnant in TeV gamma rays, and has found enhanced emission from the region where the shell appears to collide with a molecular cloud, as theories of hadronic acceleration in supernova remnants predict[17].

1.2.1 Acceleration Above 10^{14} eV

The real problem of cosmic ray acceleration begins at energies above 10^{14} eV, which the supernova blast wave model cannot account for, and intensifies with increasing energy. (Coincidentally, this is also the energy range where the diminishing flux finally necessitates air shower experiments rather than direct detection experiments.) The basic constraints that cosmic ray accelerators must meet are summarized in Hillas [18]. Particles can be accelerated in a region of field strength $\sim B$ until their gyroradii $R_{\rm gyro} = E/ZB$ increase beyond the size of the region L, at which point they will escape. Thus the size and magnetic field strength associated with an astrophysical acceleration site determine a maximum acceleration energy $E_{\rm max}$ for a particle of charge Z: $L > 2R_{\rm gyro}$. In fact, the physical properties of the actual acceleration mechanism often require a much larger size. In gradual acceleration models, where the particle scatters off many small field irregularities, the requirement becomes $L > 2R_{\rm gyro}/\beta$, where βc is the characteristic velocity of the magnetic scattering centers. Thus:

$$E_{\rm max} \sim \beta ZBL\,, \tag{1.2}$$

A similar requirement emerges in the context of shock acceleration (where βc is the speed of the shock front), or in direct acceleration by induced electromotive force (where βc is the speed of the conductor). Realistic values of β may be on the order

of $\beta \sim 1/300$ [18].

The severity of the problem of ultrahigh energy cosmic ray acceleration can be seen by locating astrophysical candidates on a plot of L vs. B, and comparing with the combined values of L and B necessary for an object to accelerate protons or iron to energies of 10^{20} eV. This is shown in Fig. 1.2, referred to as a "Hillas" plot after its introduction in Hillas [18]. The top diagonal line indicates the size and magnetic field strength required of an object to accelerate protons to 10^{20} eV, under the realistic assumption of $\beta = 1/300$. Objects below the line are not up to the task; as one can see, all known astrophysical objects are below the line. Even under the extreme assumption of $\beta = 1$, only a few classes of objects appear to be candidates. The problem only becomes more severe when we start to take into account the possible energy loss mechanisms that will compete with acceleration.

1.2.2 Extragalactic Evidence

Several lines of evidence point toward an extragalactic origin for ultrahigh energy cosmic rays. As can be seen from Fig. 1.2, there are virtually no plausible galactic sources, with the possible exception of neutron stars. The gyro-radii of protons and even iron are so large at ultrahigh energies, however, that they cannot be long contained within the galactic magnetic field, and consequently there should be considerable correlation of arrival directions with sources, or at least with the overall source distribution within the galaxy. No such anisotropy has been observed even at 10^{20} eV. Accordingly, the flattening of the cosmic ray spectrum around 10^{18} eV can be seen as a transition from a more rapidly falling galactic flux to an extragalactic flux.

The highest energy cosmic rays $\sim 10^{20}$ eV are still difficult to account for even with

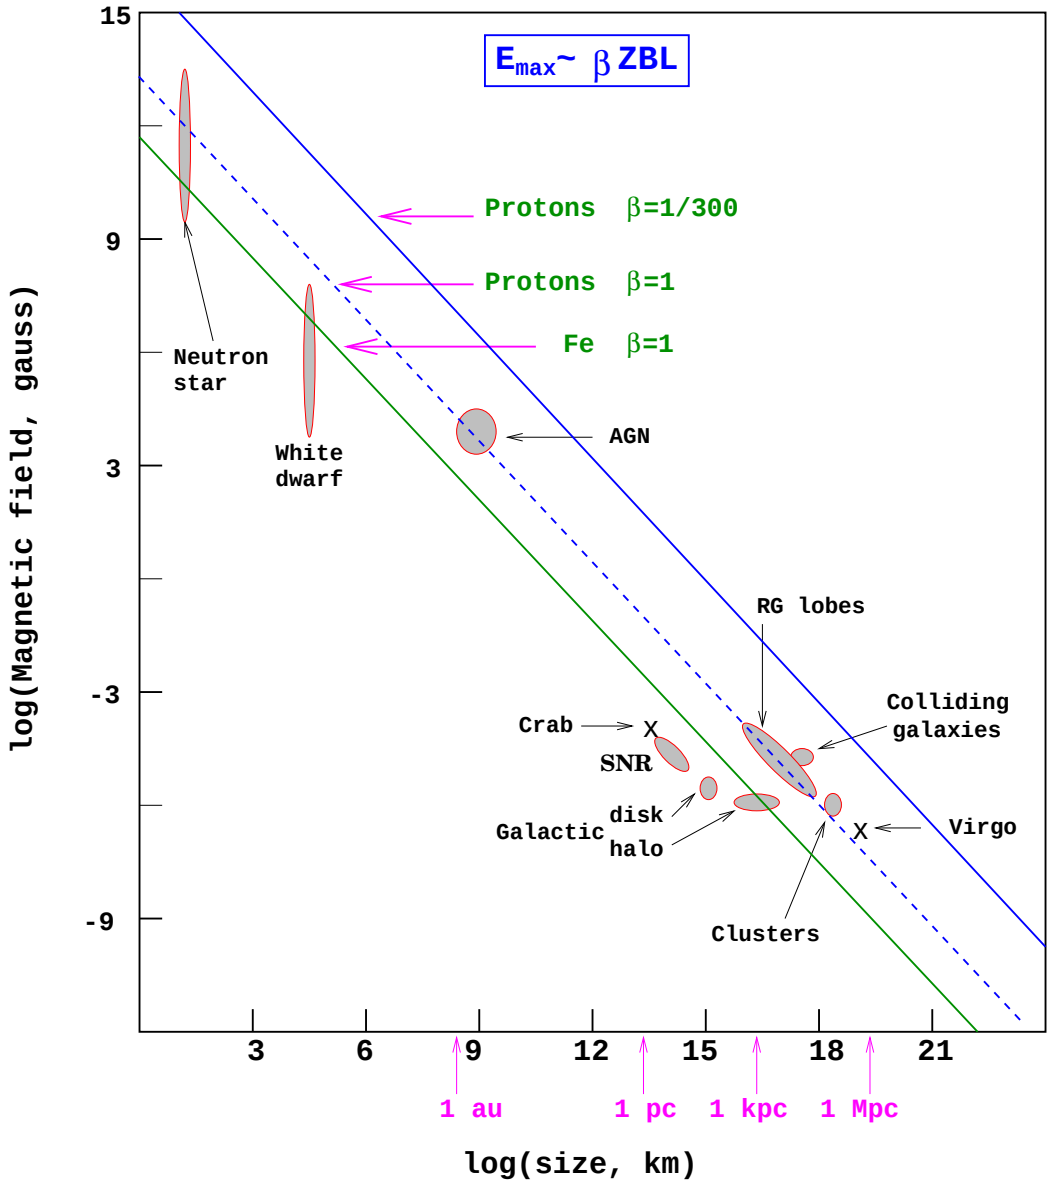

Figure 1.2: Modified Hillas Diagram, showing size and strength of magnetic fields associated with various astrophysical candidates for cosmic ray accelerators. The maximum particle energy which can be contained in an accelerating site is $E_{\max} \sim \beta ZBL$, where βc is the characteristic velocity of the magnetic scattering centers. Objects below one of the indicated diagonal lines cannot accelerate the corresponding particle species to 10^{20} eV. (Plot originally due to Hillas [18], modified here by Morat Boratav and reproduced in Bhattacharjee and Sigl [19].)

extragalactic sources. Compounding the difficulty is the GZK effect, which suppresses the flux above $\sim 10^{20}\,\mathrm{eV}$ from all but the nearest extragalactic neighbors. Fig. 1.3 shows the effect of GZK and related energy losses on a source spectrum with fixed spectral index, as a function of redshift. Evidently, the highest energy cosmic rays must originate in sources closer than redshift $z \sim 0.1$. In fact, since the spectrum above $10^{20}\,\mathrm{eV}$ is dominated by the nearest sources, the spectrum we observe will depend sensitively on the exact distribution of those (possibly rare) sources, and can't be predicted in advance by assuming an averaged, uniform distribution of sources.

1.2.3 "Top-Down" Origin

One radical alternative which should be noted is the suggestion that ultrahigh energy cosmic rays originate in the decay or annihilation of massive particles, presumably left over from the big bang. The class of such models are known as "top-down" scenarios, as opposed to the "bottom-up" scenarios discussed so far where low energy cosmic rays are accelerated to ultrahigh energies. In the proposal of Berezinsky et al. [21], the decay of super-massive (10^{21} to $10^{25}\,\mathrm{eV}$) relic particles, which are hypothesized to comprise a small fraction of the galactic dark matter, would solve both the acceleration problem and the GZK propagation problem simultaneously. However, the model requires fine-tuning so that the lifetime is on the order of the age of the universe and the present density yields the correct flux. There should also be anisotropy in arrival directions due to our off-center location within the galactic dark matter halo; such anisotropy has so far not been observed.

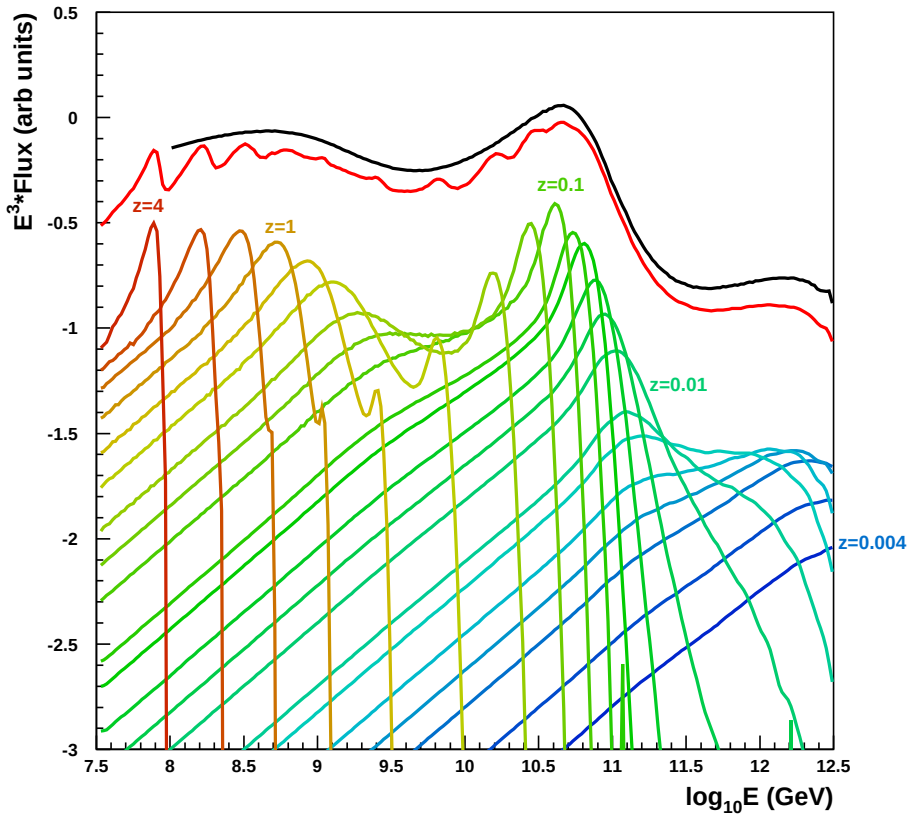

Figure 1.3: Illustration of GZK effect on proton cosmic ray flux (multiplied by E^3 to magnify detail) as a function of energy and distance to sources. Each of the lower curves represents a uniform source distribution on a spherical shell at a different redshift from the observer. (In this example, all sources have the same input spectrum $E^{-\gamma}$ with $\gamma = 2.4$, and the density of sources changes with redshift according to $(1+z)^m$ with $m = 2.5$.) Each curve shows the observed proton spectrum for that shell, taking into account the propagation energy losses due to pion production (GZK effect), pair production, and redshift.
At the top of the figure, the red curve is the combined spectrum for the sum of the shells, and the black curve is the corresponding sum for a finer series of shell spacings. This represents the spectrum predicted for a uniform distribution of sources given the assumed source spectral index γ and source evolution m. From Bergman [20].

1.3 Current Experimental Status

The current best answer to the fundamental question "how many ultrahigh energy cosmic rays are there?" is shown in Fig. 1.4. The cosmic ray spectrum above $10^{18.5}$ eV as measured by the three largest experiments—AGASA, HiRes, and Auger—is shown. The main disagreement between the experiments is an apparent systematic shift between AGASA on the one hand and HiRes and Auger on the other. However, this could be explained by a $\sim 30\%$ shift in the absolute energy scale, consistent with the systematic uncertainties of the experiments. The other point of disagreement is on the shape of the spectrum above $\sim 10^{19.8}$ eV. The AGASA experiment has claimed [22, 23] that the spectrum continues in contradiction to the suppression predicted by GZK. On the other hand, the HiRes experiment has claimed [24–26] that a suppression in the flux is indeed observed, in accordance with the GZK prediction. Disagreeing with this disagreement, it has further been argued by De Marco et al. [27] that the statistical significance of the discrepancy between the two experiments is at the 2σ level, and that only much more data will determine whether the GZK suppression is observed or not. The Auger experiment, still under construction but already on the way to obtaining the largest data sample yet, will hopefully be able to resolve this outstanding issue.

What is agreed upon by all three experiments is the existence of cosmic rays with energies above 10^{20} eV. (A well-reconstructed event above this energy has been reported by Auger, although it could not be included in the spectrum in Fig. 1.4 because it fell outside of the aperture cuts [30].) Whether the number observed above this energy is ultimately consistent with the predicted GZK suppression or not, the fundamental question of how these particles acquire such energy remains unanswered. The GZK effect merely lowers the energy we observe from what it was at the source.

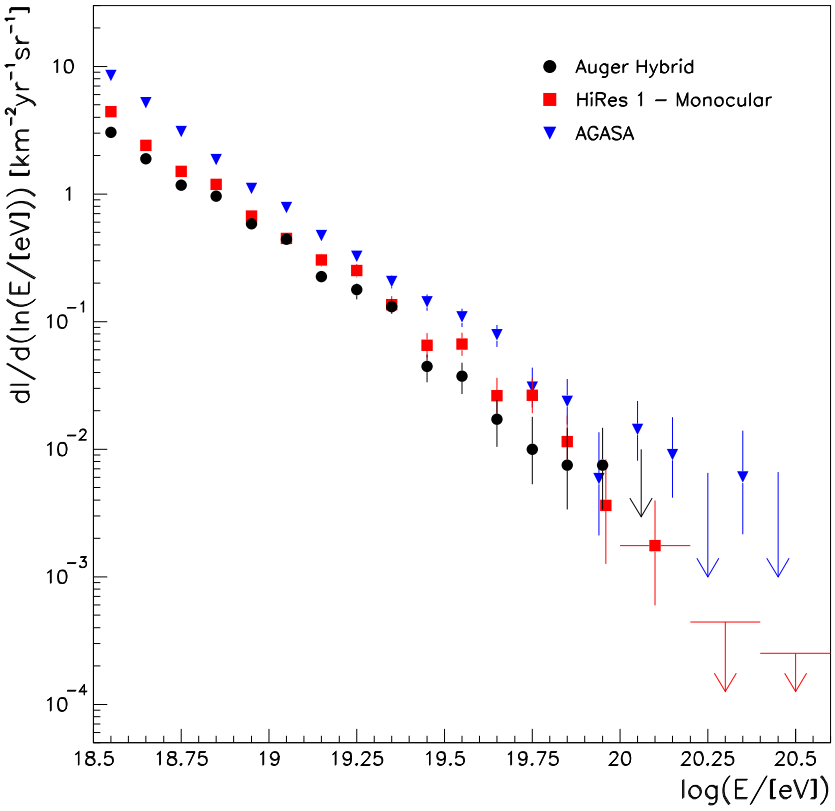

Figure 1.4: Cosmic ray spectrum above $10^{18.5}$ eV as measured by AGASA [23], HiRes I in monocular mode [24], and the Pierre Auger Observatory [28]. (From Westerhoff [29].)

It may be, then, that confirmation of the GZK effect makes the problem of finding the origin of ultrahigh energy cosmic rays only harder.

Chapter 2

Cosmic Ray Astronomy

The only direct way to identify the sources of ultrahigh-energy cosmic rays is to observe cosmic ray arrival directions pointing back toward their sources. Such an observation would go a long way toward answering—if not outright solving—the fundamental question of the origins of the highest energy cosmic rays. Such an observation would also signal the dawn of particle astronomy, opening up a new window on the universe, and the first window not mediated by the electromagnetic spectrum.

In principle, cosmic ray astronomy should be possible. The basic obstacle is that charged cosmic rays will deflect as they pass through the magnetic fields that permeate galactic and extragalactic space. However, higher energy particles will deflect less, so at sufficiently high energies cosmic rays should propagate sufficiently straight that sources (or at least source regions) in the sky can be identified.

The energy threshold and corresponding size of the data sample beyond which cosmic ray astronomy becomes possible is not known. The dominant uncertainties are the distance to the sources, the intervening magnetic fields, and the primary particle composition. Some widely cited estimates are discussed below. However,

the lack of *any* confirmed anisotropy in arrival directions to date already constrains theories of cosmic ray origins.

2.1 Deflections by Magnetic Fields

Charged particles will inevitably experience deflections by magnetic fields as they propagate through space. We can estimate the size of the effect as follows. The gyroradius is given by $R_{\rm gyro} = p/qB_\perp$, for a particle with charge q and relativistic momentum p traveling through a magnetic field with component $B_\perp$ perpendicular to the direction of motion. For ultrahigh energies, $p \approx E/c$. Expressing this in relevant units we have:

$$R_{\rm gyro} \approx 1\,{\rm kpc}\,\frac{1}{Z}\left(\frac{E}{10^{18}\,{\rm eV}}\right)\left(\frac{\mu{\rm G}}{B}\right), \tag{2.1}$$

where Z is in units of elementary charge. To put this in perspective, the galactic magnetic field strength is $\sim$ few μG, and the solar system is $\approx 8\,$kpc from the galactic center. Therefore at energies of $10^{18}\,$eV even protons (the least deflected of charged primaries) will experience large deflections during propagation from anywhere beyond the immediate vicinity of the solar neighborhood.

At higher energies, however, the deflections may become small enough that arrival directions are reasonably correlated with source locations. For a particle traveling a distance S through a uniform magnetic field with perpendicular component B, the angular deflection from its initial trajectory is $\delta = S/R_{\rm gyro}$, or, in units relevant for high energies:

$$\delta \approx 0.5^\circ\; Z\left(\frac{S}{\rm kpc}\right)\left(\frac{B}{\mu{\rm G}}\right)\left(\frac{10^{20}\,{\rm eV}}{E}\right). \tag{2.2}$$

For a given magnetic field strength, a uniform field is the worst case scenario from the point of view of particle deflections. At the other extreme, propagation through

a series of randomly oriented magnetic fields results in a mean deflection of zero, and the important quantity is δ_{rms}. For a particle traveling a distance S through a turbulent field with mean value B_{rms} and correlation length L_c, we have (for $S \gg L_c$):

$$\delta_{rms} = \frac{1}{\sqrt{2}} \frac{qcB_{rms}}{E} \sqrt{SL_c} \tag{2.3}$$

(see Harari et al. [31]). In units relevant for the galaxy:

$$\delta_{rms} \approx 0.1^\circ \; Z \left(\frac{B}{\mu\mathrm{G}}\right) \left(\frac{10^{20}\,\mathrm{eV}}{E}\right) \sqrt{\frac{S}{\mathrm{kpc}}} \sqrt{\frac{L_c}{100\,\mathrm{pc}}}\,. \tag{2.4}$$

For propagation through a turbulent field, the deflections are effectively a random walk with $N = S/L_c$ deflections and the expected $\sqrt{N}$ behavior for the net deviation.

2.2 Galactic Magnetic Field

The existence and approximate strength of a widespread galactic magnetic field (GMF) was first deduced in 1950 by Arnulf Shlüter and Ludwig Biermann, thanks to new observations of, of all things, cosmic rays [32]. The newly-measured cosmic ray spectrum revealed an isotropic distribution of arrival directions up to energies of $10^{16}\,\mathrm{eV}$ and implied a cosmic ray energy density of $0.6\,\mathrm{eV/cm^3}$ in space. It was argued that such a high energy density couldn't apply to intergalactic space; rather, the cosmic rays must be confined within the 2000-lightyear-thick disk of the Milky Way, implying an average magnetic field of about 5μG. This remarkable prediction has subsequently been born out by modern measurements.

The best probes of the GMF arise from the Faraday rotation of radio emission from pulsars in our galaxy and polarized extragalactic radio sources. While the

details of the GMF are hard to reconstruct, the outline of its features are known (See Alvarez-Muniz and Stanev [33] and references therein for the following.) Its principle components are a regular and a turbulent field, of roughly comparable strength ($\sim$ few μG). The regular field has a spiral structure in the galactic plane, related to the spiral arms of the galaxy. There is disagreement, however, about whether the field reverses from one arm to the next (bi-symmetric spiral) or not (axi-symmetric spiral). These two field configurations are shown in Fig. 2.1. There is also disagreement about whether the field reverses or not on crossing the galactic plane, and whether there exists an additional halo or dipole component along the galactic polar axis.

The known features of the GMF set important limits on the regime where charged particle astronomy will be possible. While cosmic rays with energies up to 10^{18} eV are expected to be completely isotropized by the GMF, above 10^{19} and especially 10^{20} eV it is expected that lighter primaries (e.g. protons) can travel large distances through the galaxy and still remain relatively undeflected. Also, the roughly equal strength of the regular and turbulent components of the GMF is advantageous: comparing Eq. 2.2 and Eq. 2.4 for deflections in uniform and turbulent fields, respectively, the S vs. $\sqrt{S}$ dependence on propagation distance indicates that, to first order, deflections of ultrahigh energy cosmic rays will be dominated by the regular component of the GMF alone.

However, the still considerable uncertainties in our knowledge of the GMF limit our ability to make more than order-of-magnitude estimates for the deflections. It is plausible that some directions of propagation through the galaxy correspond to much smaller deflections than others, due to the happenstance of our location and the exact GMF structure. Unfortunately, as shown in Kachelriess et al. [34], calculations of the deflections for different GMF models corresponding to some of the remaining

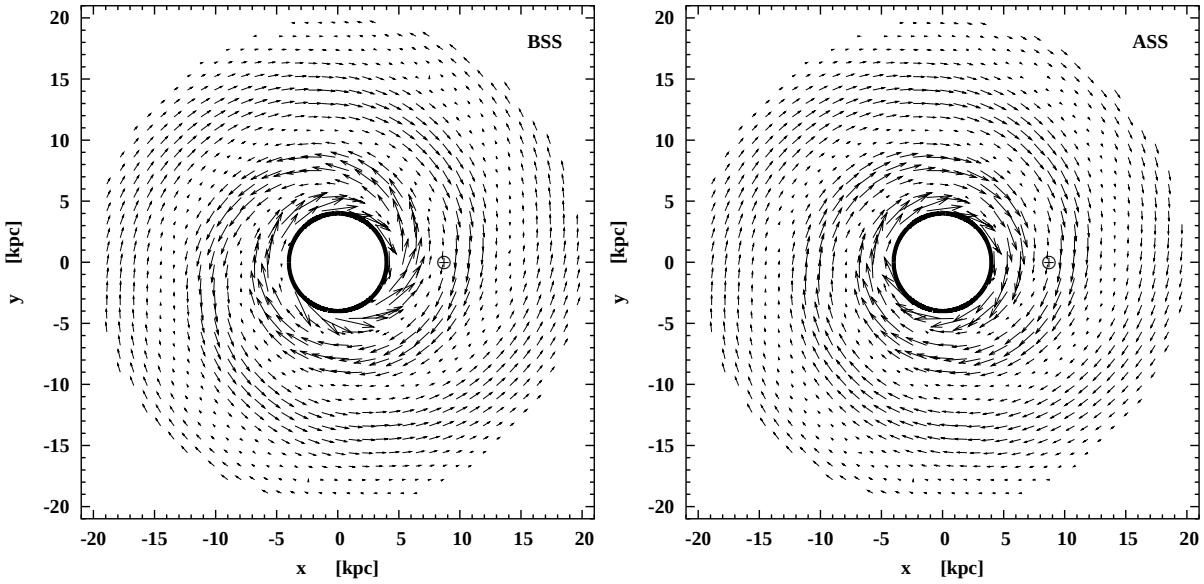

Figure 2.1: Two models of the spiral structure of the galactic magnetic field. Left: Bi-symmetric spiral with field reversals from arm to arm. Right: Axi-symmetric spiral without reversals. The location of the solar system is indicated by the small circle at $r \approx 8\,\mathrm{kpc}$. The magnetic field for $r < 4\,\mathrm{kpc}$ is not shown, for clarity. From Alvarez-Muniz and Stanev [33].

uncertainties lead to drastically different predictions for the small-deflection, "good seeing" directions through the galaxy. The authors considered three models of GMF which differ according to the parity of the GMF across the galactic plane, the inclusion of a dipole field, and the assumptions about the field dependence on the z coordinate out of the plane. The details of the models, labeled TT, HMR, and PS, can be found in Kachelriess et al. [34]. What is mainly of concern here is the difference in predicted deflections between the models, as reproduced in Fig. 2.2. For each direction in the sky (shown in galactic coordinates), the color scale indicates the net deflection of a $4 \times 10^{19}\,\mathrm{eV}$ proton propagated backward through the GMF until it is 50 kpc away from the galactic center. The regions of the sky corresponding to small deflections, indicated by dark colors in the plots, have only a narrow overlap between the three models. Moreover, the true GMF may differ sufficiently from these models that none of the predicted directions of small deflections are accurate.

2.3 Extragalactic Magnetic Field

While the details of the galactic magnetic field are still uncertain, enough is known to indicate that at energies near and above $10^{20}\,\mathrm{eV}$, light cosmic ray primaries like protons should be able to pass through the GMF with minimal ($\lesssim$ few degrees) deflection. In this energy range, then, the promise of cosmic ray astronomy hinges on whether these particles can travel undeflected over cosmological distances as well. This depends on the strength and distribution of extragalactic magnetic fields (EGMF).

Compared to the GMF, the EGMF is much more uncertain. Faraday rotation measurements of polarized sources located behind galaxy clusters indicate fields at the μG level up to 1 Mpc from the center of the cluster [32, 35]. Propagation through

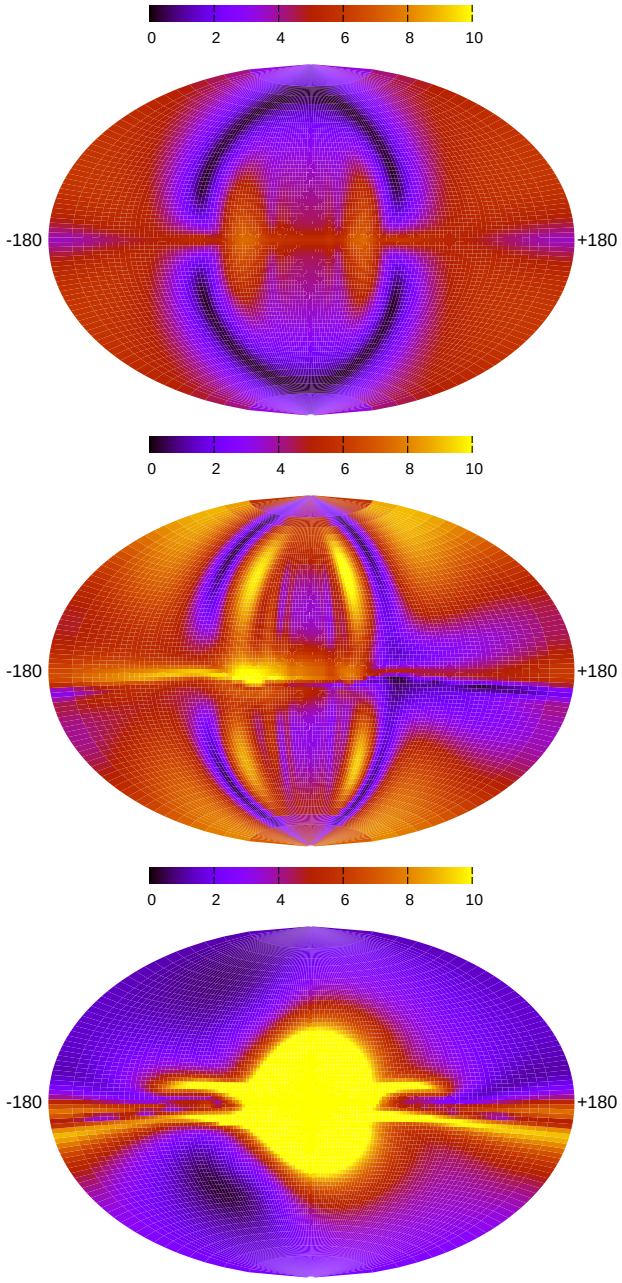

Figure 2.2: Deflections (scale in degrees) for three different models of the GMF: TT (top), HMR (middle), and PS (bottom), described in the text. In each direction of the sky (shown in galactic coordinates), a proton of energy 4×10^{19} eV is propagated backward through the GMF until it is 50 kpc from the galactic center. The total deflection δ of the particle is indicated. From Kachelriess et al. [34].

such fields over these distances would induce large deflections even for 10^{20} eV protons. However, galaxy clusters and filaments constitute only a fraction of the volume of the universe. On much larger scales, only upper limits of $1 - 10\,\mathrm{nG}$ on the field strength exist [36]. Yet even these upper limits are not tight enough to predict that cosmic ray astronomy must be possible. For a worst case scenario, e.g. a 1 nG field coherent across the present horizon, we see from Eq. 2.2 that the corresponding bound on the deflection of a 10^{20} eV proton traveling 100 Mpc is $\delta \lesssim 50^{\circ}$.

Fortunately, the upper bound on the EGMF inferred from Faraday rotation measurements may be far above the actual field strength. Theoretical predictions vary widely, and there is considerable discrepancy in the results of numerical simulations. For example, in Dolag et al. [35] the authors perform a magneto-hydrodynamical simulation using seed fields which are amplified by structure formation. They constrained the simulation so that it reproduced the specific features of large scale structure in our region of the universe, and then studied the deflections of proton cosmic rays traveling toward us within it. Fig. 2.3 shows the size of the deflections as a function of direction in the sky; in Fig. 2.4, the corresponding cumulative distribution of deflections is shown for protons at 4×10^{19} eV for a range of distances. Even for protons at such low energies, deflections of $\lesssim 1^{\circ}$ during travel over 100 Mpc account for 80% of the sky. The low field strength of $\sim 10^{-11}$ G which is found in this simulation for most of the EGMF is a strong boost for the prospects of cosmic ray astronomy at ultrahigh energies.

However, numerical simulations described in Armengaud et al. [37] predict much dimmer prospects. Here, the authors perform a cosmological simulation with magnetic fields generated at the shocks that form during large-scale structure formation.

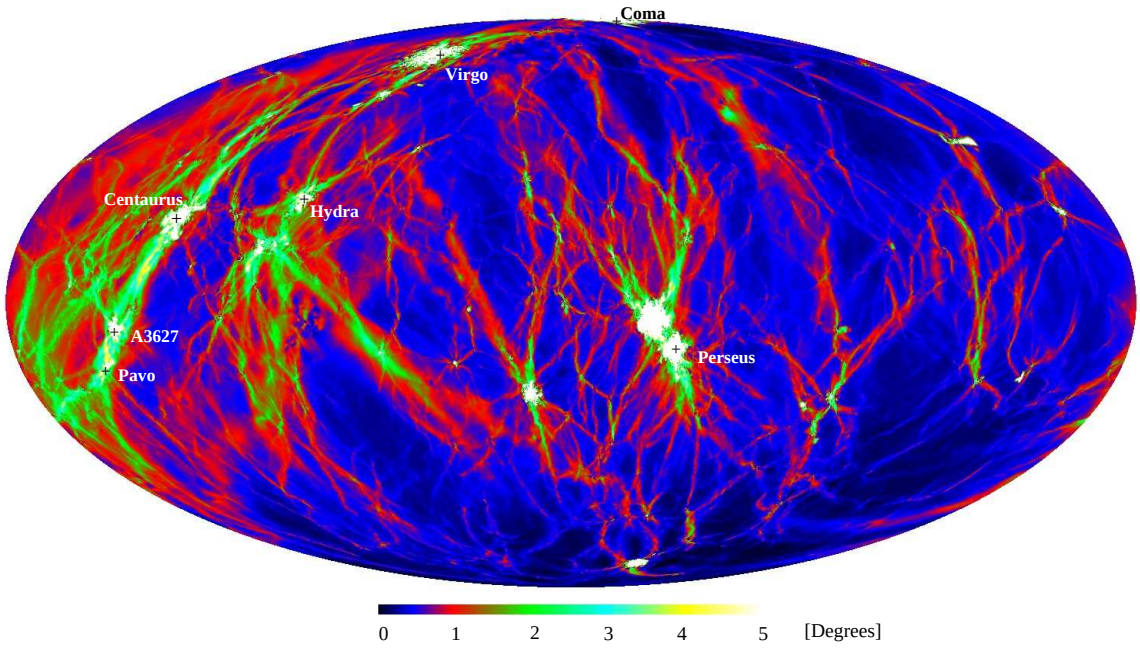

Figure 2.3: Skymap showing the accumulated deflections for a proton with energy 4×10^{19} eV traveling 107 Mpc through the magnetic field structure of a simulation of the Local Universe. (Galactic coordinates, with anti-center in the middle.) From Dolag et al. [35].

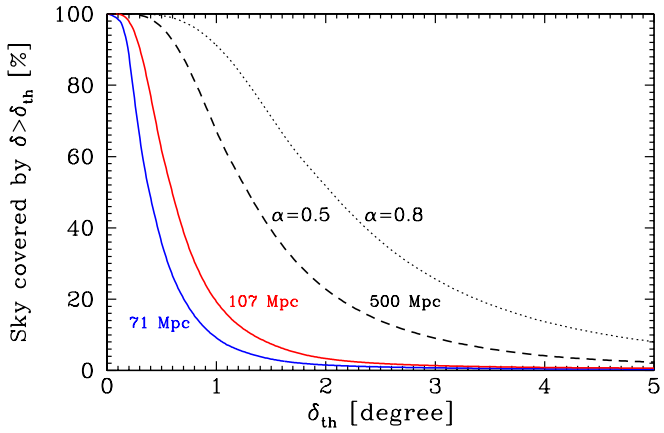

Figure 2.4: Cumulative fraction of the sky with deflection angle larger than $\delta_{\rm th}$ for 4×10^{19} eV protons. Two extrapolations to a distance of 500 Mpc are shown. From Dolag et al. [35].

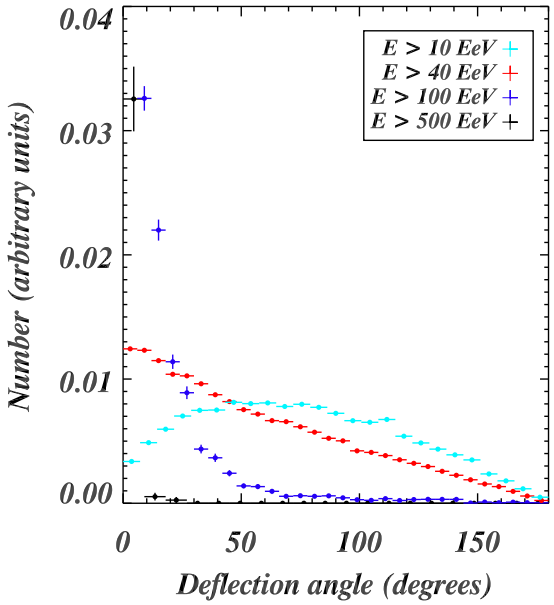

Figure 2.5: Deflection angle distribution for protons from sources with a source density of $2.4 \times 10^{-5}\,\mathrm{Mpc}^{-3}$ (cumulated over many different realizations of source locations). From Armengaud et al. [37].

Source locations are distributed throughout the simulated volume with a source density $2.4 \times 10^{-5}\,\mathrm{Mpc}^{-3}$. In Fig. 2.5, the distribution of deflection angles for protons of various energies is shown. Even at energies of $10^{20}\,\mathrm{eV}$, the typical deflection is tens of degrees. The authors have noted the discrepancy with the predictions in Dolag et al. [35] described above, and conclude that in addition to the model differences, there may be other technical issues involved.

2.4 Neutral Cosmic Ray Primaries

Neutral cosmic ray primaries are immune to deflection by magnetic fields, so their arrival directions might point directly toward sources at the limit of experimental angular resolution. The extent to which this scenario is actually realized depends on what fraction of the cosmic ray flux is neutral, and what type of neutral particles are involved.

If the GZK effect occurs as expected, then a proton-dominated flux will contain a neutron component, due to the process $p + \gamma \rightarrow \Delta \rightarrow \pi^+ + n$. However, these neutrons will decay back to protons with a mean decay length given by:

$$L_{\rm decay} = \gamma c \tau_{\rm decay} \approx 1\,{\rm Mpc} \left(\frac{\rm E}{10^{20}\,{\rm eV}} \right) . \tag{2.5}$$

The neutron decay length is compared with the nucleon-GZK interaction length in Fig. 2.6. Above $\sim 3 \times 10^{20}\,$eV, neutrons will tend to interact (resulting in a proton or neutron again) before decaying, but at lower energies, it is clear that any GZK produced neutrons will rapidly decay back to protons, and the flux of neutrons is expected to be quite small below $10^{20}\,$eV.

Photons also suggest themselves as neutral cosmic ray candidates. The main obstacle to cosmic ray photon propagation is electron-positron pair production due to interactions with CMB and other background photons. The left plot in Fig. 2.7 shows the photon attenuation lengths in the infrared, cosmic microwave, and radio background. At $10^{15}\,$eV, the attenuation length due to the CMB is as small as 10 kpc, but rises again at higher energies, and would be ~ 10 Mpc at ultrahigh energies, if not for the radio background. Unlike the CMB, the radio background is not well known, but as shown in the right plot in Fig. 2.7, for energies below 10^{20} eV even the

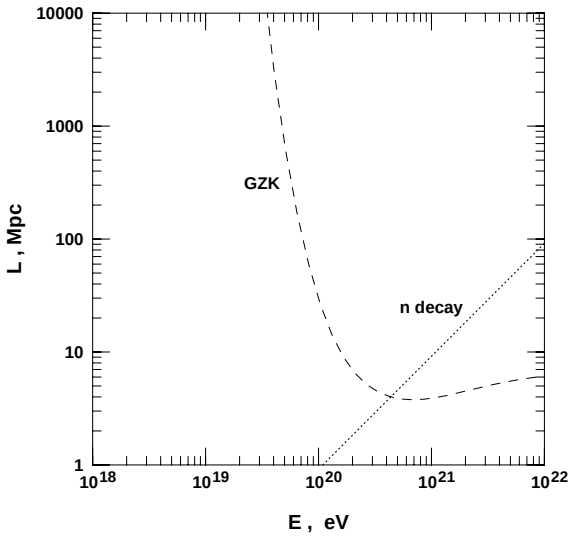

Figure 2.6: Proton/Neutron GZK interaction length (dashed line) and neutron decay length (dotted line) as a function of energy. Adapted from Stanev [11].

complete absence of the radio background would only increase the interaction length by less than an order of magnitude.

After pair production, what happens next depends on the extragalactic magnetic field strength. For fields $\sim 1\,\mathrm{nG}$ or stronger, the electrons lose energy quickly, and as shown in the right plot of Fig. 2.7, the synchrotron energy loss length is much shorter than the photon interaction length. In this scenario, pair production is effectively the end of the ultrahigh energy cosmic ray. However, in the case of much lower magnetic field strengths $\sim 0.01\,\mathrm{nG}$, such as found in the simulations of Dolag et al. [35] described above, synchrotron energy loss is low ($\propto E_e^2 B^2$) and inverse Compton scattering is then possible. This would effectively lengthen the propagation distance of ultrahigh energy photons.

Top-down models (see Sec. 1.2.3) of ultrahigh energy cosmic ray production tend

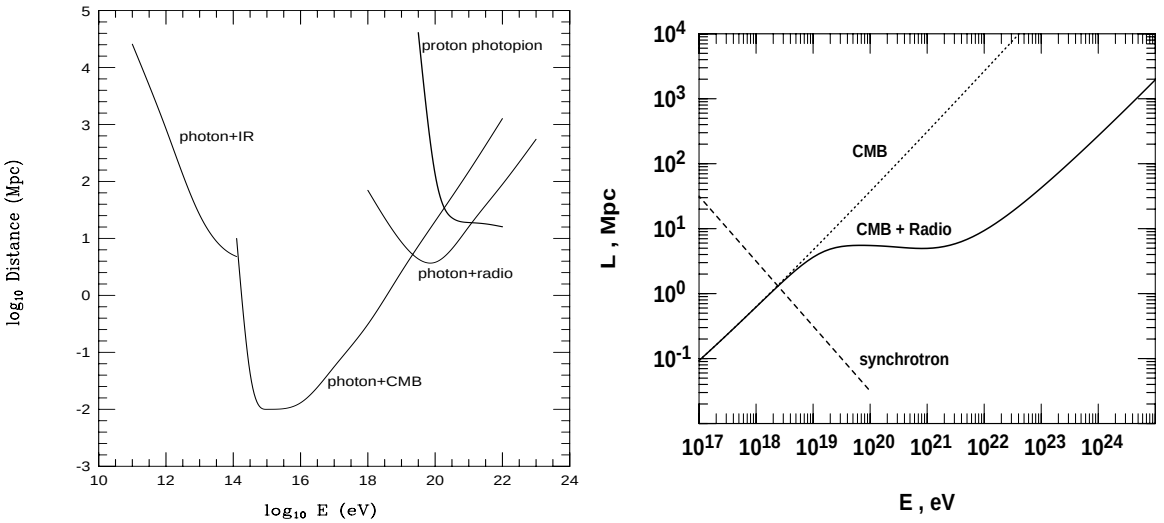

Figure 2.7: *Left*: Photon attenuation lengths in extragalactic space due to the infrared, cosmic microwave, and radio backgrounds. For comparison, the proton GZK attenuation length is also shown. Adapted from Torres and Anchordoqui [38].
Right: Photon interaction length due to CMB only, and due to CMB plus estimated radio background. Also shown is the synchrotron energy loss length of electrons in 1 nG magnetic field. Adapted from Stanev [11].

to generate predominantly photons, but conventional acceleration scenarios also predict a photon component of the flux. Like neutrons, photons are an expected by-product of the GZK effect, from the decay of π^0. Gelmini et al. [39] estimated that these "GZK-photons" would account for between 10^{-4} to 10^{-2} of the total ultrahigh energy nucleon flux above 10^{19} eV, depending on the assumptions of source spectrum, source distribution, and intervening backgrounds.

Neutrinos propagate large distances without deflections or interactions, and "GZK-neutrinos" (from the decay of π^+) are also an expected consequence of the GZK effect. However, these neutrinos would be detected by dedicated neutrino experiments such as ANITA and IceCube; the ultrahigh energy primaries responsible for the air showers observed by cosmic ray detectors have cross sections (determined by the height at which they first interact in the atmosphere) much too large to be neutrinos.

2.5 Timing Correlations

Many of the highest energy astrophysical phenomena suspected of accelerating ultrahigh energy cosmic rays also exhibit rapidly changing behavior on the time scale of hours or even less. These include active galactic nuclei (AGN), gamma-ray bursts (GRB), and supernovae. If cosmic rays were observed to correlate with any of these sources in time as well as direction, it could drastically improve the signal to noise ratio in point-source searches and lead directly to the identification of sources.

To correlate with the light signal, one first must check what delay is incurred by massive particles traveling a distance L compared with light, $\Delta t = L/\beta c - L/c$. For

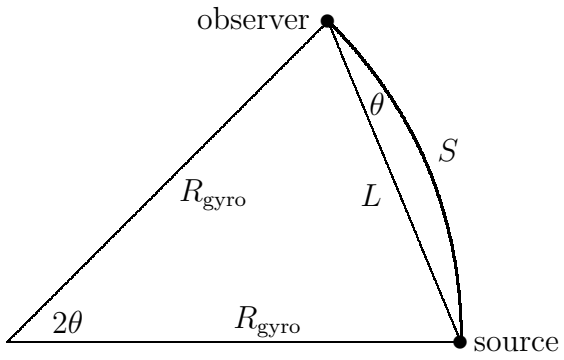

Figure 2.8: Deflected path S vs. straight-line distance L between source and observer.

ultrahigh energies we therefore have $\Delta t \approx L/(2c\gamma^2)$ or:

$$\Delta t \approx 10^{-4}\,\mathrm{s}\left(\frac{L}{\mathrm{Mpc}}\right)\left(\frac{m}{\mathrm{GeV}}\right)^2\left(\frac{10^{18}\,\mathrm{eV}}{E}\right)^2. \tag{2.6}$$

Evidently, at ultrahigh energies, the time delay between light and massive particles propagating rectilinearly would be negligible: even a primary as massive as iron and as low in energy as 10^{18} eV would lag only 30 seconds or so behind light after traveling 100 Mpc.

However, *charged* primaries will not propagate rectilinearly, due to intervening magnetic fields. For a charged particle which undergoes a small deflection on its way from the source, what is the time delay introduced by the deflection? A simple sketch illustrating the case for a uniform magnetic field is shown in Fig. 2.8. For a source located a distance L from the observer, a cosmic ray arriving at an angular separation θ from the line of sight to the source will have traveled along the arc S. The extra path-length $\Delta S = S - L$ is given by

$$\Delta S = L\left(\frac{\theta}{\sin\theta} - 1\right) \approx L\,\frac{\theta^2}{6}. \tag{2.7}$$

Therefore the time delay Δt is:

$$\Delta t = \frac{\Delta S}{c} \approx 60\,\text{days} \left(\frac{L}{\text{kpc}}\right) \left(\frac{\theta}{\text{deg}}\right)^2 . \tag{2.8}$$

For a galactic source, Eq. 2.8 implies that timing information will be of modest help at best. For an extragalactic source, timing information will be effectively useless for charged particles: even a tenth of a degree deflection over 100 Mpc will introduce a delay of hundreds of years between the arrival of the light signal and the cosmic rays.

Evidently, only a neutral particle brings with it potentially useful timing information. And as discussed earlier, any ultrahigh energy photon or neutron which arrives on earth is not likely to have been the same particle during the entire journey from the source. Even a small segment of the trip traversed as a charged particle can turn a fraction of a second delay into weeks.

2.6 Summary

On the pessimistic side: the galactic magnetic field strength is understood well enough that the pointing of even cosmic ray protons back toward their sources seems unlikely except at energies well above $10^{19}\,\text{eV}$. At energies $\sim 10^{20}\,\text{eV}$ proton cosmic ray astronomy seems feasible, *if* extragalactic field strengths are well below the current upper limits. Such low fields are consistent with the results of some cosmological simulations, but not all.

Neutral primaries offer the best prospects for cosmic ray astronomy, but they are expected to be only a small fraction of the total flux. Photons may comprise from 10^{-4} to 10^{-2} of the total nucleonic flux above $10^{19}\,\text{eV}$. Neutrons may comprise a substantial fraction of the flux above $10^{20}\,\text{eV}$, but will still spend on average at least

half of their propagation time as protons and thus incur some deflections. Below 10^{20} eV, the neutron fraction of the flux is expected to be very small.

Time correlations between cosmic rays and high energy astrophysical phenomena holds out the promise of vastly improving the signal to noise ratio in anisotropy searches. Unfortunately, even the smallest of deflections leads to delays from days to millennia between the arrival of the light signal and the cosmic rays. Such studies are still worth performing, but it is difficult to motivate *a priori* any particular search window lasting much longer than the light signal itself.

In short, the uncertainties in the cosmic ray source distribution, intervening magnetic fields, extragalactic radio background, and primary composition are so great that complete isotropy and significant anisotropy are more or less equally plausible.

Chapter 3

Measuring Extensive Air Showers

The extremely low flux of ultrahigh energy cosmic rays (~ 1 per km^2 per year above 10^{19} eV) makes direct observation of these particles infeasible. Fortunately, when these cosmic rays arrive in the upper atmosphere, they initiate extensive air showers of particles which continue to the ground, with billions of particles spread over tens of square kilometers. The study of ultrahigh energy cosmic rays, then, is the study of these showers, and the inferences which can be made about the primary particle. We begin with a brief review of the basic properties of extensive air showers, and then describe the main methods of detection.

3.1 Properties of Extensive Air Showers

When a cosmic ray primary (whether proton, nucleus, or photon) interacts high in the atmosphere, a cascade of particles develops as a thin disk propagating along the original direction of motion. Each of the particles can in turn interact, so the number of particles multiplies, until reaching a maximum and then falling off as more and more of the particles drop below the threshold for further particle production.

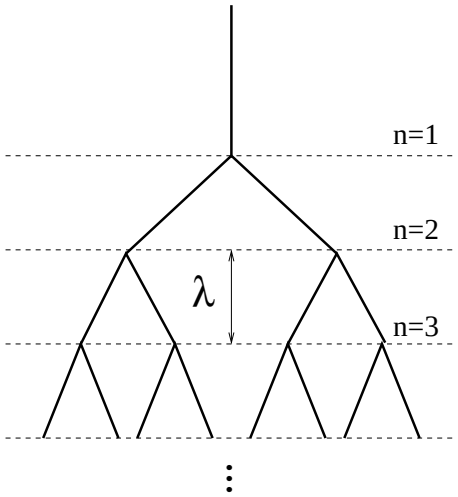

Figure 3.1: Heitler's toy model for the branching process in an air shower cascade. From Pierog et al. [41].

3.1.1 Heitler's Toy Model

The simplest quantitative model of this process is due to Heitler [40] and described in Gaisser [15] as follows.

Consider an incoming particle with energy E_0, as shown at the top of Fig. 3.1, which undergoes an interaction, leading to two particles which split the initial energy equally. After traveling a characteristic collision length λ, each of the particles branches again and the energy is split in two again. At a depth X in the atmosphere (more accurately, the slant-depth, for inclined showers), there have been $n = X/\lambda$ branchings, and the total number of particles $N(X)$ is therefore

$$N(X) = 2^{X/\lambda} \,. \tag{3.1}$$

The branching continues until the energy of each particle equals E_c, the critical energy for the splitting process. The number of particles at this point is $N_{\max} = E_0/E_c$, which

occurs at the corresponding slant-depth

$$X_{\max} = \lambda \frac{\ln(E_0/E_c)}{\ln 2} \,. \tag{3.2}$$

While this is highly simplified, there are two basic features of the model which hold for electromagnetic cascades and approximately for hadronic cascades as well: the number of particles is proportional to the primary energy, $N_{\max} \propto E_0$; and the depth of shower maximum is proportional to the log of the primary energy, $X_{\max} \propto \ln E_0$.

This simple model can be pushed further to describe the difference in showers originating from light and heavy nuclei. In the superposition model, a primary nucleus of mass A and energy E_0 is treated as a superposition of A nucleons of energy E_0/A. Each nucleon produces a shower independently, with corresponding $X_{\max}$:

$$X_{\max} = \lambda \frac{\ln(E_0/A\,E_c)}{\ln 2} \,. \tag{3.3}$$

Thus, iron and proton primaries of the same energy will have different shower developments, with iron primaries tending to produce showers with smaller $X_{\max}$ (i.e., higher in the atmosphere). It should be emphasized, however, that such differences only apply *on average*: in real showers, the fluctuations in $X_{\max}$ from one shower to the next are much larger than the differences due solely to composition. $X_{\max}$ is thus a statistical tool, but not accurate enough determine the primary composition of individual showers.

3.1.2 Components of an Air Shower

What the shower consists of depends on the type of primary particle. Photons generate a predominately electromagnetic cascade: the dominant process for a high energy

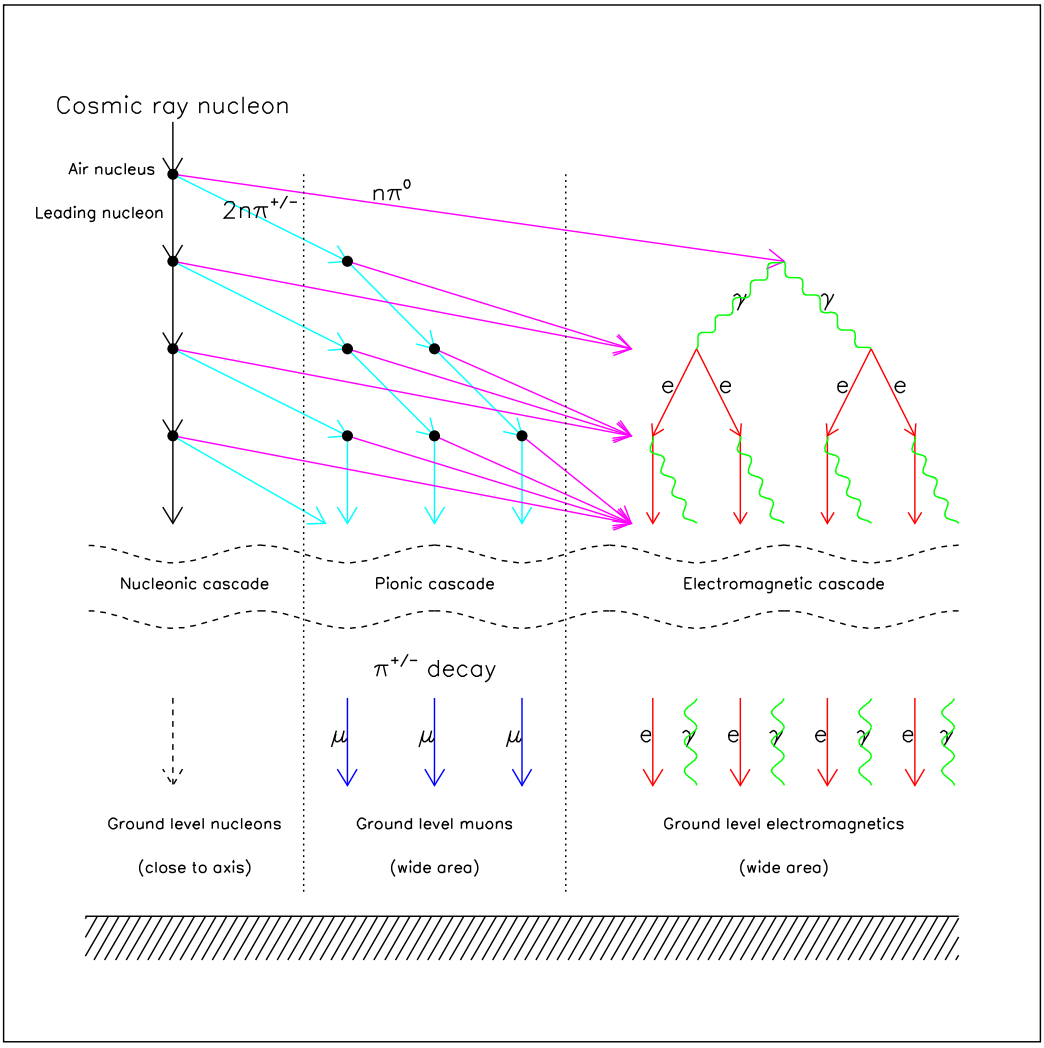

Figure 3.2: Highly simplified schematic diagram indicating basic cascade processes in an extensive air shower. Here, the cosmic ray primary is assumed to be a nucleon, initiating an air shower which can be divided into three categories: the nucleonic cascade; the pionic cascade; and the electromagnetic cascade. Pion decays feed the electromagnetic cascade and also the muon component which reaches ground level. (Adapted from Auger Collaboration [42].)

photon traveling through air is electron-positron pair production, and these particles in turn generate more photons via bremsstrahlung. For nucleonic primaries, there is a hadronic core consisting of nucleons and pions, and the pion decays in turn feed both the electromagnetic component (via photons from neutral pion decays) and the muonic component (via charged pion decays). A highly simplified schematic of these interrelated processes is shown in Fig. 3.2.

Photon-initiated showers are thus expected to be muon-poor compared with nucleon-initiated showers of the same energy, with only 5% of the number of muons [15, p. 197]. This offers another possible way to study composition in addition to $X_{\max}$. Iron showers, in turn, are expected to have 80% more muons than proton showers of the same energy [42]. Because this property is related to the fact that the iron shower develops faster, it is partially correlated with the $X_{\max}$ value, but it provides a somewhat independent handle on the composition and is measured by a completely independent method.

In each hadronic interaction, roughly a third of the energy goes into the electromagnetic component. Thus most of the energy of the shower is eventually found in the electromagnetic component, and most of that energy is dissipated by the electrons and positrons through ionization losses. The atmosphere, in effect, is a calorimeter, and the integral of the shower's longitudinal profile is a measure of the primary energy:

$$E_{\mathrm{em}} = \alpha \times \int dX\, N(X) \tag{3.4}$$

where α is the energy loss per unit path-length in the atmosphere and E_{em} is the energy in the electromagnetic component. The total energy E_0 of the primary may be somewhat higher than this, due to energy not observed (e.g. neutrinos or muons which go into the ground); for iron nuclei this missing energy can be up to 15% [42].

3.2 Methods of Detection

There are two main techniques used to observe extensive air showers today, sketched in Fig. 3.3: *ground arrays*, which sample the shower particles arriving at ground level, and *air fluorescence telescopes*, which observe the passage of the shower through the atmosphere via the light emitted from excited air molecules. Because ground arrays sample the lateral distribution of particles at one location along the shower, while air fluorescence telescopes measure the longitudinal development of the shower, these techniques can yield complementary information about the shower. An important difference between the techniques is the fact that air fluorescence observations can only be performed during clear, moonless nights, which restricts the duty cycle of such experiments to $\sim 10\%$, whereas ground arrays can operate at nearly 100% duty cycle.

3.2.1 Ground Array

For observations of ultrahigh energy cosmic rays, a practical spacing for ground array detectors is $\sim 1\,\mathrm{km}$. This yields an adequate sample of the particle density in the shower front at reasonable cost without constructing an overdense array. The lateral distribution of particles as a function of distance from the central axis of the shower, or *shower core*, is shown in Fig. 3.4 for a $10^{19}\,\mathrm{eV}$ proton.

While most of the particles in the shower lie within approximately 100 m of the shower core, a sufficient number are still observed by the detectors surrounding the core location. The timing and signal strength measured by these detectors can then be used to reconstruct the location of the shower core and the arrival direction of the cosmic ray primary. Typically, the more detectors which are hit, the smaller the

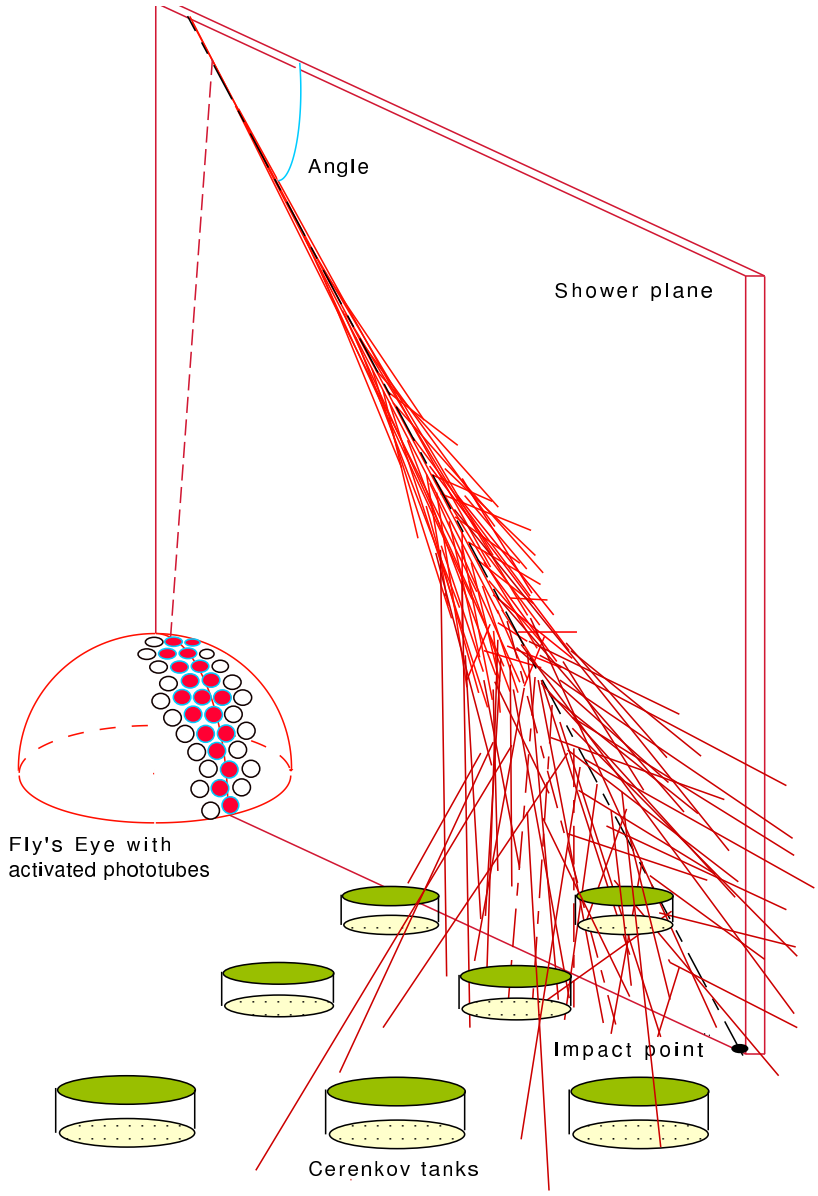

Figure 3.3: Sketch of the two main techniques used today for observing and measuring the properties of extensive air showers. *Air fluorescence telescopes* consist of an array of PMTs—often referred to as a Fly's Eye array—which observe the longitudinal development of the shower by detecting UV light from nitrogen molecules excited by the passage of charged particles. *Ground arrays* sample the lateral distribution of particles arriving on the ground, using either scintillators or water-Cerenkov detectors. Plot due to the Auger Collaboration, reproduced in Westerhoff [29].

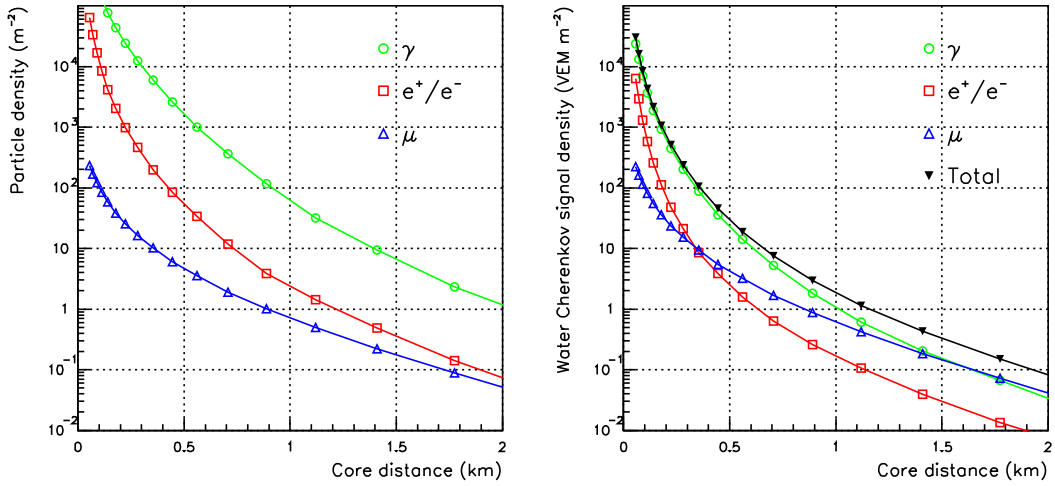

Figure 3.4: Left: Lateral distribution of particle densities at ground level for a simulated 10^{19} eV proton shower. Right: Convolution of particle densities with response of 1.2 m deep water Cerenkov detector. The signal is expressed in units of the signal produced by a muon traversing the detector vertically, hence "vertical equivalent muon" or VEM [6]. From Auger Collaboration [42].

angular uncertainty of the arrival direction. Thus the angular resolution tends to improve for higher energy showers, and to some extent for more inclined showers, which effectively encounter a denser array as seen from the point of view of the shower axis. Typical uncertainties are $\sim 3°$, with $\sim 1°$ achievable under optimal circumstances.

Once the shower core is located, the strength of the signal in each detector serves as a sample of the lateral distribution of the particle density. From this distribution the energy of the shower can be inferred. In practice, the sampled distribution is typically used to estimate a single quantity: the signal density at a fixed distance from the shower core. AGASA, for example, used $S(600)$, the energy deposition density in plastic scintillators at 600 m from the core; Auger, which has a larger spacing between detectors than AGASA, uses $S(1000)$, the time-integrated water Cerenkov

signal that would be measured by a tank 1000 m from the core [28]. The fixed distance is chosen so that the signal's correlation with shower energy has minimal dependence on the type of primary cosmic particle. This distance can be determined by shower simulations. However, because shower simulations are based on extrapolations of hadronic interactions observed at much lower center-of-mass energies in accelerators, the unknown uncertainties in these models propagate into the energy estimation.

The major advantages of ground arrays are the robustness of the detectors and the nearly 100% duty cycle. The chief disadvantage is that the ground array samples the shower at only one stage in its longitudinal development, and the lateral sampling itself is relatively sparse compared to the particle density in the core. The result is that energy estimates in ground arrays rely considerably on the still uncertain extrapolations of particle physics models, and are not as direct as the more calorimetric estimates in the air fluorescence technique.

3.2.2 Air Fluorescence

As the air shower particles travel through the atmosphere, they dissipate energy by exciting and ionizing air molecules. The excited nitrogen molecules fluoresce in the near-UV with an emission spectrum shown in Fig. 3.5. This light is emitted isotropically, allowing the shower to be "seen" as it develops through the atmosphere.

Most importantly, the intensity of the fluorescence light is proportional to the number of charged particles. The fluorescence yield—the number of photons per charged particle per meter—is relatively constant over the range of temperatures and pressures in the atmosphere which the shower passes through, averaging about 4.8 photons per electron [42].

Thus as one observes the shower development, the number of photons detected

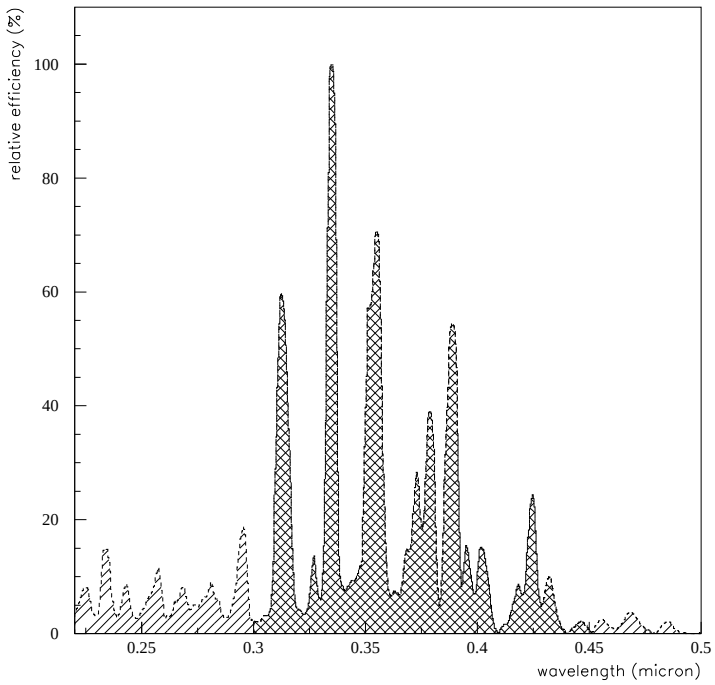

Figure 3.5: Emission spectrum of fluorescence light from the 2P band of molecular nitrogen and the 1N band of the N_2^+ molecular ion. Approximately 80% of the light is emitted between 300 nm and 450 nm. Adapted from Auger Collaboration [42].

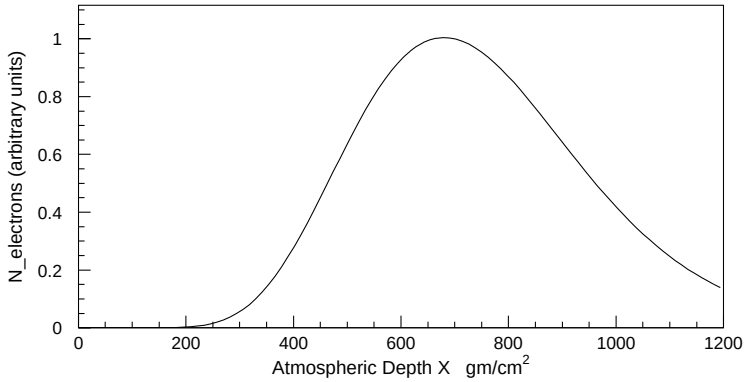

Figure 3.6: Typical form of Gaisser-Hillas shower profile.

at each depth X can be used to infer the number of charged particles $N(X)$. From Eq. 3.4, the integral of the number of particles along the shower path is related to the electromagnetic energy, $E_{\rm em} = \alpha \times \int dX\, N(X)$, where $\alpha = 2.2$ MeV/(g cm^{-2}) is the average energy deposited by a charged particle per g/cm^2 of atmosphere [42].

In practice, simply integrating along the entire shower is generally not possible. One reason is that the detector often only views a partial segment of the shower development. The standard practice is rather to fit the observed shower with a generic model for a shower profile, most commonly the Gaisser-Hillas [43] formulation:

$$N(X) = N_{\rm max} \left(\frac{X - X_0}{X_{\rm max} - X_0} \right)^{\frac{X_{\rm max} - X_0}{\lambda}} \exp\left(\frac{X_{\rm max} - X}{\lambda} \right) , \qquad (3.5)$$

where X_0 is the depth of the first interaction, and a typical value for the characteristic length λ is 70 gm cm^{-2}. A typical Gaisser-Hillas shower profile is plotted in Fig. 3.6.

Moreover, the light received by the detector is not solely fluorescence light: Cerenkov photons are also copiously produced. If the shower travels relatively directly toward the detector, the Cerenkov signal can exceed the fluorescence signal. Even when the shower is not viewed head-on, the Cerenkov photons can be scattered toward the detector, an effect which grows with the number of Cerenkov photons as the shower ages. The effect of adding Cerenkov light to the fluorescence light is illustrated in Fig. 3.7, using a shower orientation with respect to the detector that especially dramatizes the effect.

Scattering processes therefore need to be well understood to correct for their effect. The two mechanisms are Rayleigh (molecular) scattering and Mie (aerosol) scattering. Rayleigh scattering is relatively straightforward to model, as it simply depends on the density of air that the light passes through, and the atmospheric density profile

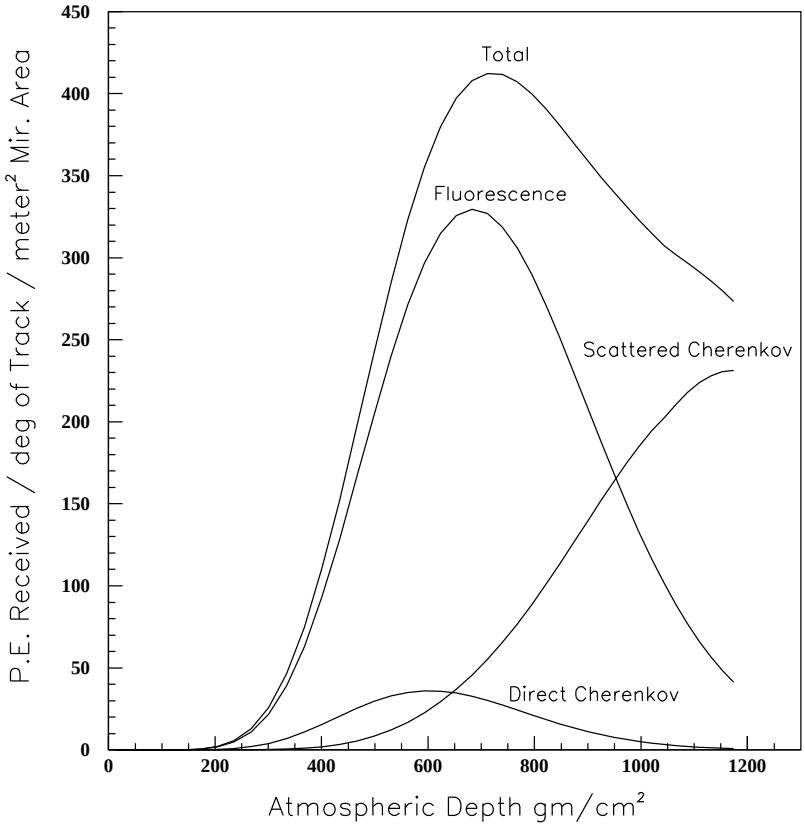

Figure 3.7: Simulation of signal vs. atmospheric depth of light signal (after passing through 300-400 nm filter) for a 10^{19} eV proton shower, oriented toward the detector so that the contribution from Cerenkov photons is enhanced. From Auger Collaboration [42].

is generally stable. Mie scattering, however, depends on the nature and quantity of particulate matter suspended in the air, and on short time scales can undergo large variations with respect to its long-term average. It is therefore necessary to monitor the atmospheric conditions while data is being taken if aerosol scattering is to be correctly accounted for.

While Cerenkov photons are scattered toward the detector, fluorescence photons are also scattered away by the same processes. All of these effects must be accounted for when the number of detected photons is used to infer the number of charged particles in the shower.

The chief advantage of the air fluorescence technique is that, while it depends on laboratory measurements of the nitrogen fluorescence yield and on constant atmospheric monitoring at the site, in the end it achieves a more direct, calorimetric measure of the energy in the shower than a ground array does. The main disadvantage of the air fluorescence technique is that it only works at night, while the moon is down and the weather is clear, resulting in a 10% duty cycle.

There is an additional benefit of the air fluorescence technique not yet mentioned. Because it views the shower as it traverses many kilometers through the atmosphere, this allows for much better angular resolution of arrival directions than is generally achievable with ground arrays. The extent to which this improved angular resolution can convey useful information about the origin of ultrahigh energy cosmic rays is the primary subject of this thesis.

Chapter 4

The HiRes Experiment

4.1 Description of the HiRes Detector

The High Resolution Fly's Eye (HiRes) experiment is located in the western Utah desert at the U.S. Army's Dugway Proving Ground. (113° W longitude, 40° N latitude, vertical atmospheric depth 860 g/cm^2). It consists of two detector sites located on hilltops and separated by 12.6 km. Fig. 4.1 depicts the geographical arrangement, and illustrates the basic advantage of having two sites: stereo observations dramatically improve the ability to pinpoint the shower trajectory.

In order to trigger both detectors, low energy showers must fall between the two sites, as illustrated by the shower in Fig. 4.1. For high energy showers, however, this geometry is less typical. HiRes can detect showers at distances over 30 km; thus the detectors have been built to look out over the desert in all directions, and most high energy showers are seen around the detectors rather than between the two sites.

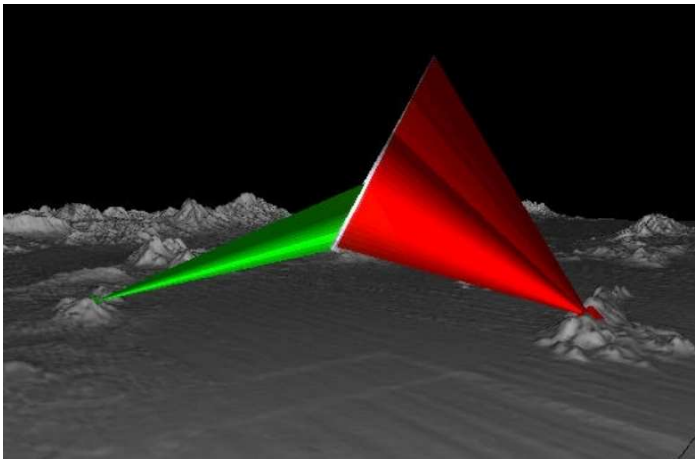

Figure 4.1: Simulation of an extensive air shower as viewed in stereo by the two HiRes detector sites. HiRes-I is on the left, located on top of Five Mile Hill; its view of the shower is indicated in green. HiRes-II is on the right, located on Camel's Back Ridge; its view of the shower is indicated in red. Note that HiRes-II has a second ring of mirrors, allowing an extended view of the early shower development.

4.1.1 HiRes-I and HiRes-II Detectors

The detectors are not identical. The first detector, HiRes-I, consists of a single ring of 22 mirrors looking at elevations between 3° and 17° degrees above the horizon, and a full 360° of azimuthal direction. The second detector, HiRes-II, has two rings of mirrors (42 in total), giving greater elevation coverage, from 3° to 31°.

Each mirror is built from four spherical quadrants and has a total area of 4 m^2. The night sky is imaged onto a cluster of photomultiplier tubes (PMTs) in a 16×16 array, with each tube forming a pixel of about 1°. The cluster is located slightly offset from the focal plane of the mirror in order to optimize the spot-size with respect to the size of the pixels [44], and to make the spot-size more uniform over the whole area of the mirror. The front of the cluster has a glass filter which transmits in the near ultraviolet, between 300-400 nm.

In addition to differences in the viewing geometry, the two detectors differ in data acquisition systems. The HiRes-I detector was completed first, in 1997, and uses a sample-and-hold system measuring the total integrated charge within a 5.6 μs time window. When the trigger logic boards determine from individual tube triggers that a mirror trigger should be generated, the PMT charge and trigger times are digitized and sent to the central DAQ computer [45].

The HiRes-II detector was completed in 1999 and uses a flash analog-to-digital conversion (FADC) data acquisition system. The FADC system continuously digitizes all PMT signals every 100 ns, which get stored in a 820 μs delay buffer. Among the advantages of this configuration are that it allows for adequate delay to make trigger decisions, and for deadtimeless readout of all channels from before to after an event [46].

4.1.2 Calibration

There are several aspects of calibration involved in order to properly reconstruct air shower measurements. The most fundamental is the absolute calibration of the PMTs, which is accomplished using a Roving Xenon Flasher (RXF). This is a portable xenon flash bulb which can be carried from one PMT cluster to the next. At each one, it is placed at the center of the mirror so that it illuminates the cluster directly. The operation is time-consuming, so it is performed approximately once a month (e.g. a good project during a night with bad weather). The RXF serves as a standard candle, with pulse-to-pulse variations $\sim 0.3\,\%$, and stability over a night of operation better than $2\,\%$ [47]. At HiRes-I, the response of the PMTs determines the gain: the slope of the plot of pedestal-subtracted QDC counts vs. number of photoelectrons. At HiRes-II, the response is used to adjust the software-controlled gains to yield one

FADC count per photoelectron [45].

An additional calibration is provided by a central YAG laser at each site, which delivers light to each of the clusters via quartz optical fibers. One bundle of fibers goes to the center of the mirror, where it illuminates the cluster directly. The other is mounted on the side of the cluster so that it first illuminates the mirror, thus providing a way to monitor the mirror reflectivity. The YAG calibration is performed simultaneously for all of the clusters at the beginning and end of each nightly run.

4.1.3 Atmospheric Monitoring

The atmosphere must also be well understood and modeled in order to properly account for the effects of scattered light, which include fluorescence light scattered away from the detector and Cerenkov light scattered toward the detector. Mie scattering by aerosols can vary significantly with elevation and on short time scales, and therefore must be monitored continuously. This is done using two steerable UV lasers, one installed at each site. Over the course of one hour, the lasers fire in a pre-determined series of directions in order to sample adequately the aperture of the detector. The scattered light from the beam emitted by one site is observed by the other site (see Fig. 4.2). These measurements are used to determine the aerosol vertical optical depth, the aerosol horizontal attenuation length at detector level, and the aerosol phase function. A detailed description of this technique is given in Abbasi et al. [48].

The sequence of laser shots can also be interrupted in order to perform a shoot-the-shower measurement. When a high energy candidate event is detected, the laser at one site can shoot at several intervals along the line of sight toward the shower. The other detector can observe the scattered light from this shower-detector plane, giving detailed, timely atmospheric information for the most important events [49].

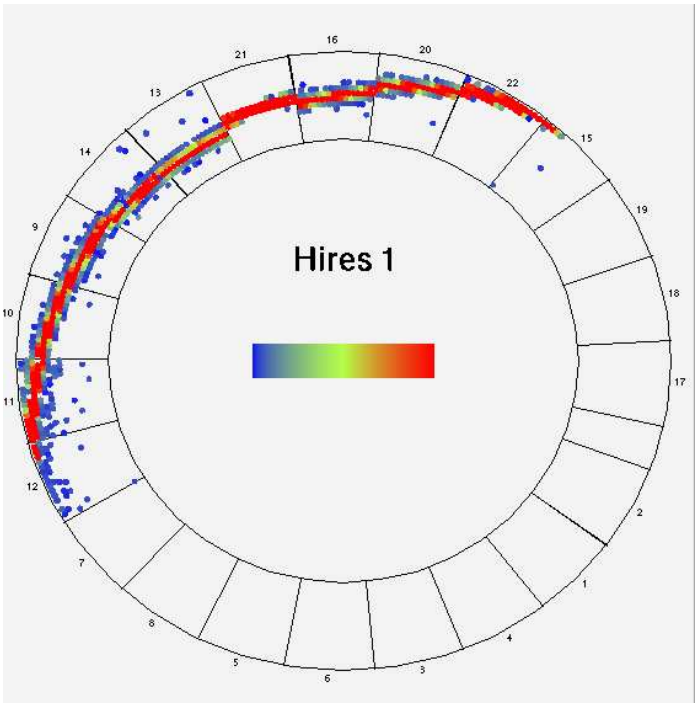

Figure 4.2: Laser Shot: HiRes-I observation of scattered light from a nearly horizontal laser shot emitted by the steerable laser at HiRes-II. The beam starts from the southwest (left), travels past HiRes-I about 450 m to the north (top), and off the northeast (right). Such a long track allows observation of a large range of scattering angles. From Abbasi et al. [48].

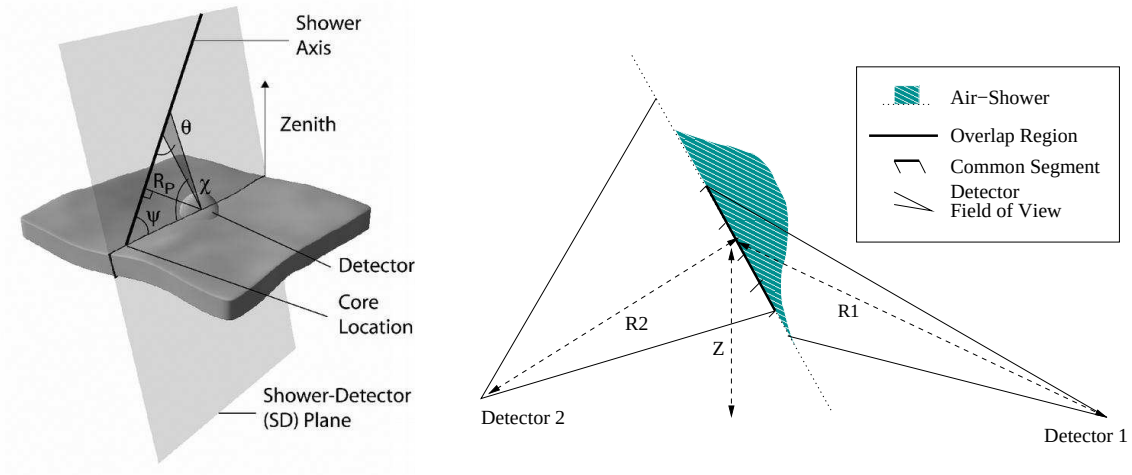

(A) (B)

Figure 4.3: A: One shower-detector plane, defined by a single detector and the air shower axis (from [25]). B: Stereo observation of shower development, defining two intersecting shower-detector planes (from [50]).

4.2 Reconstruction of Extensive Air Showers

The reconstruction of an air shower can be divided into two main components: first, determining the trajectory of the shower, and second determining the profile of the shower development. Because the trajectory is determined primarily from the pointing direction of the PMTs and the time of the PMT signals, reconstructed arrival directions are largely insensitive to the calibration and atmospheric monitoring issues discussed above. These issues are important, however, for determining the shower energy and $X_{\max}$ parameter.

4.2.1 Trajectory Reconstruction

Fig. 4.3 illustrates the basic principles involved in reconstructing the shower trajectory. The left plot shows how the shower axis together with one detector determine

a shower-detector plane. While the HiRes PMT pixel size is 1°, a long track across many PMTs allows determination of the shower-detector plane to an accuracy of 0.2°. Additionally, the orientation of the shower within the shower-detector plane can be constrained by the time sequence of the PMT signals as the shower develops and the image moves across the array. This "monocular" reconstruction estimates the impact parameter $R_{\rm p}$, and the angle ψ which the shower makes with the ground (within the shower-detector plane), both shown in the left plot of Fig. 4.3.

A monocular reconstruction of an air shower event is shown in Fig. 4.4. From the upper two panels it can be seen that the shower track is already clearly distinguishable from the noise hits based on pointing directions alone. The lower left panel shows the timing information, which further distinguishes tubes involved in the shower from random noise. The times t_i of the PMT hits, and corresponding angles χ_i at which the PMTs view the shower in the shower-detector plane (see labels in Fig. 4.3A), are related to the parameters $R_{\rm p}$ and ψ by

$$t_i = T_0 + \frac{R_{\rm p}}{c} \tan\left(\frac{\pi - \psi - \chi_i}{2}\right) , \tag{4.1}$$

where T_0 is the time at which the shower passes through the point of closest approach. (N.B. T_0 is not the time at which light *arrives* at the detector from that point; that occurs at $T_0 + R_{\rm p}/c$.) One therefore wants to find the best fit $R_{\rm p}$, ψ, and T_0 from the PMT timing and pointing directions. The ability to determine three parameters depends on the degree to which the relation between t_i and χ_i is nonlinear. In the lower left plot of Fig. 4.4, the nonlinearity is clearly evident, but the degeneracy can be considerable in shorter tracks, leading to much larger uncertainties in the orientation of the shower within the shower-detector plane.

However, stereo observation of an air shower yields two shower-detector planes,

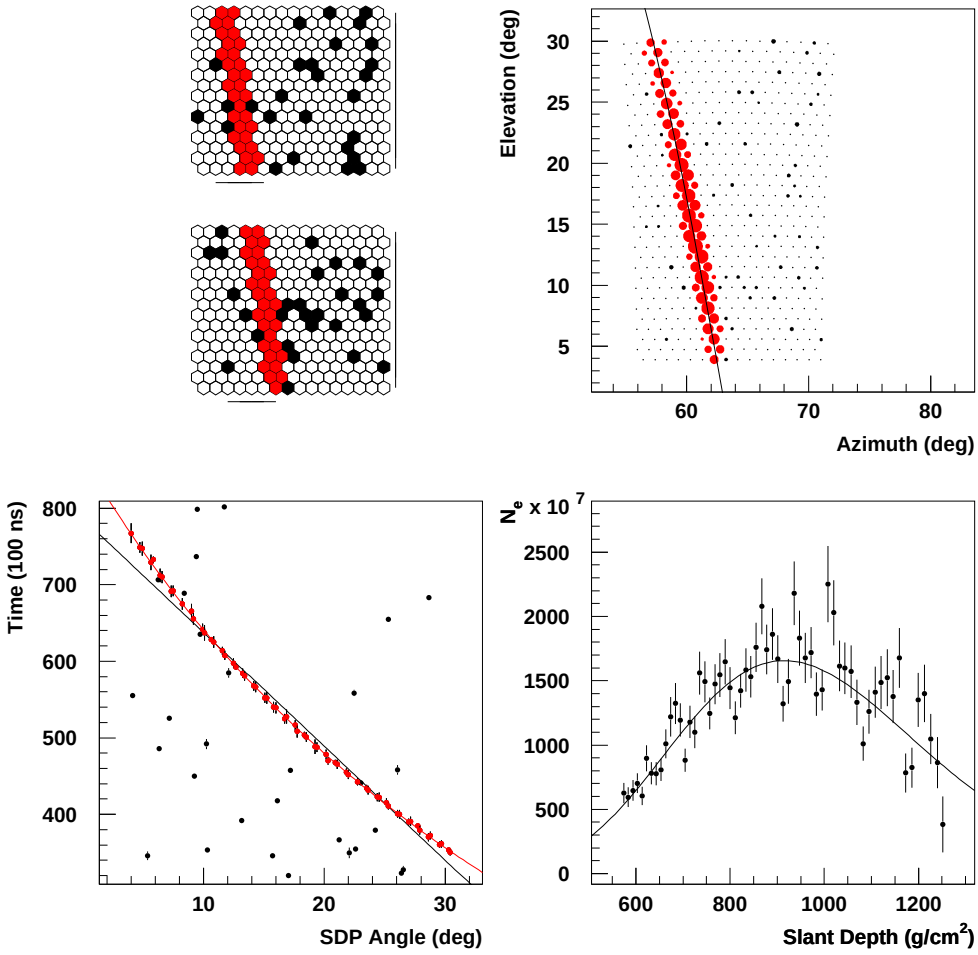

Figure 4.4: HiRes-II Event display and track reconstruction of an event with reconstructed energy of 2.4×10^{19} eV. Upper left panel: the two mirrors which triggered for this event. Upper right panel: the viewing elevation and azimuthal angles of the PMTs, with the fitted shower-detector plane superimposed. Lower left panel: the time of the tube hits in FADC time slices vs. the angle of the tube measured along the track, with a straight-line fit and a time fit superimposed. Lower right panel: reconstruction of the number of charged particles in the shower as a function of slant depth in atmosphere, with a fit to the Gaisser-Hillas model superimposed. From Abu-Zayyad et al. [25].

as shown in Fig. 4.3B. The four variables $(\theta, \phi, R_{\rm p}, \psi)$ defining the shower axis and location can generally be determined from the intersection of these planes alone. In practice the HiRes stereo reconstruction improves on this by performing a global χ^2 minimization using both the pointing and timing information of all of the PMTs to obtain the best estimate of the shower geometry.

4.2.2 Profile Reconstruction

Once the shower trajectory is identified, the PMT signals can be used to estimate the size of the shower (i.e. number of charged particles) as a function of atmospheric depth, which ultimately results in an overall shower profile.

While the signal in a given PMT serves to estimate the number of photons from the shower along a particular segment, these photons are a mixture of fluorescence photons and scattered Cerenkov photons. The mixture depends on the rest of the shower, so it is not straightforward to extract the fluorescence part this way. A different approach is to employ an iterative process. Using a Gaisser-Hillas profile (Eq. 3.5) with an initial estimate for $N_{\rm max}$ and $X_{\rm max}$, a simulated shower is generated and the corresponding PMT signals are determined based on all of the direct and scattered light and the properties of the detector. This is compared with the observed PMT signals, and the process is repeated with new estimates of $N_{\rm max}$ and $X_{\rm max}$ until best fit values are found. Fig. 4.5 shows an example of a stereo-reconstructed profile of an event observed by the HiRes detector, and indicates the separate components of the light signal which correspond to the best fit.

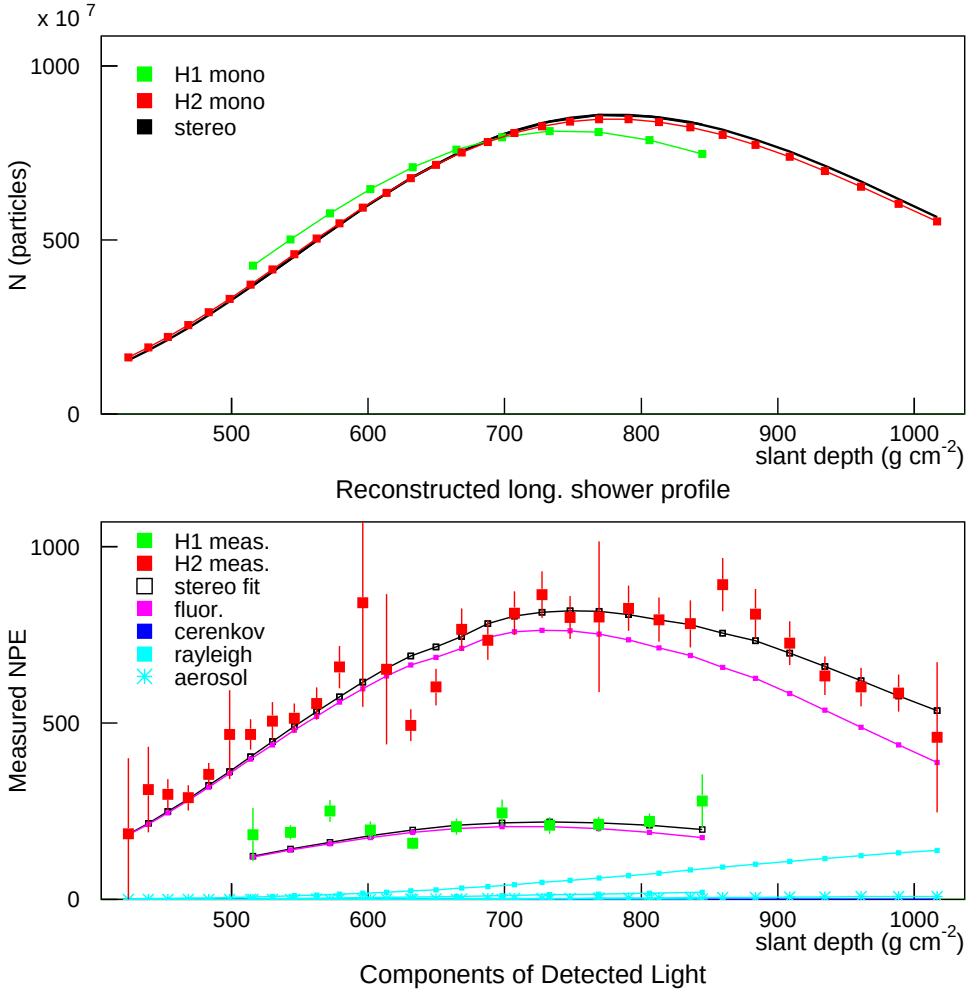

Figure 4.5: Stereo reconstruction of an event observed by the HiRes detector:
Upper panel: Longitudinal profile of the air shower (number of charged particles as a function of depth in atmosphere) as reconstructed by HiRes-I mono, HiRes-II mono, and HiRes stereo.
Lower panel: Measured number of photoelectrons per degree of track per m^2 of effective mirror area, for HiRes-I and HiRes-II (squares with error bars). Also shown is the expected number of photoelectrons given the best stereo fit. The total light signal consists of the individual contributions shown from fluorescence light, direct Cerenkov light, and Rayleigh and Mie (aerosol) scattered light. (In this case, the best fit requires no direct Cerenkov light and only minimal light scattered by aerosols.) This shower was approximately 25 km from HiRes-I and 15 km from HiRes-II, and has a stereo-reconstructed energy of 13 EeV.

4.3 HiRes Data Sample

This section describes the HiRes stereo event sample that is analyzed in this thesis. These events were recorded between the beginning of December 1999 and the end of January 2004, during which there were 2052 hours of observing time. Stereo observation began in the fall of 1999 and was relatively stable by that December. (Data taking was interrupted, however, due to access restrictions at the Dugway Proving Ground beginning in the fall of 2001, and could not be resumed until the following spring, after the peak winter viewing months had passed).

The following quality cuts have been applied in order to extract a sample of well-reconstructed events for the study of arrival directions (other analyses, such as measuring the spectrum, use different cuts):

- Geometry fit: $\chi^2/\mathrm{dof} < 5$
- Shower profile fit: $\chi^2/\mathrm{dof} < 5$
- Zenith angle uncertainty: $\sigma_z < 2^\circ$
- Azimuth angle uncertainty (scaled): $\sigma_a \cdot \sin\theta_z < 2^\circ$
- Energy uncertainty: $\sigma_E/E < 20\%$
- Zenith angle: $\theta_z < 70^\circ$

The last cut, requiring the zenith angle to be less than 70°, has been applied because in simulations of the detector and event reconstruction, the energy and angular errors are seen to grow worse beyond this range (see e.g. Fig. 5.2).

Weather conditions which reduce the quality of the data are cut implicitly by the above requirements, rather than by explicit weather cuts.

After these cuts, we arrive at a data set with 4495 events of all energies, including 271 events above 10^{19} eV. A plot in equatorial coordinates of the arrival directions of these high energy events is shown in Fig. 4.6.

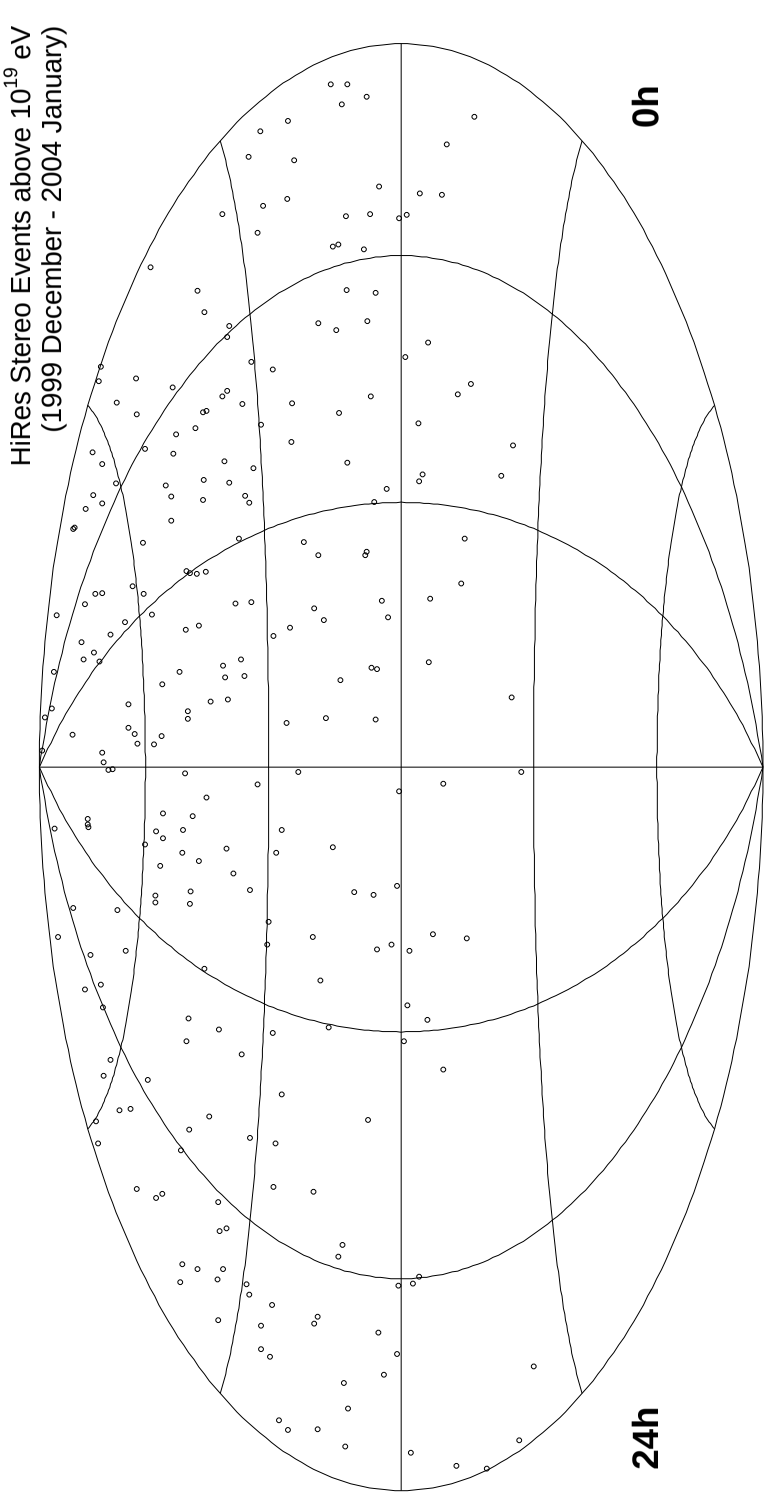

Figure 4.6: Skymap (in equatorial coordinates) of the 271 HiRes stereo events above 10^{19} eV examined in this study.

Chapter 5

HiRes Angular Resolution

The HiRes detector measures cosmic ray arrival directions with a typical uncertainty $\sim 0.5^\circ$ in stereo operation. This angular resolution is unprecedented at ultrahigh energies, and provides a unique opportunity to search for small-scale anisotropy. This chapter describes how the angular resolution is studied using detector simulations, how the uncertainties can be modeled, and how both measurement and systematic errors of the real detector are estimated.

5.1 Angular Resolution of the Simulated Detector

The angular resolution of HiRes is estimated using a full detector simulation of proton showers generated with CORSIKA 6 [51], using QGSJET for the first interaction. For this study, an E^{-3} differential energy spectrum was simulated, so that the distribution of reconstructed energies would be similar to that of the data. The simulated events are reconstructed in the same way as the data, and the same quality cuts described earlier in Ch. 4.3 are applied.

In Fig. 5.1, the distribution of opening angles $\Delta\theta$—the angular difference between

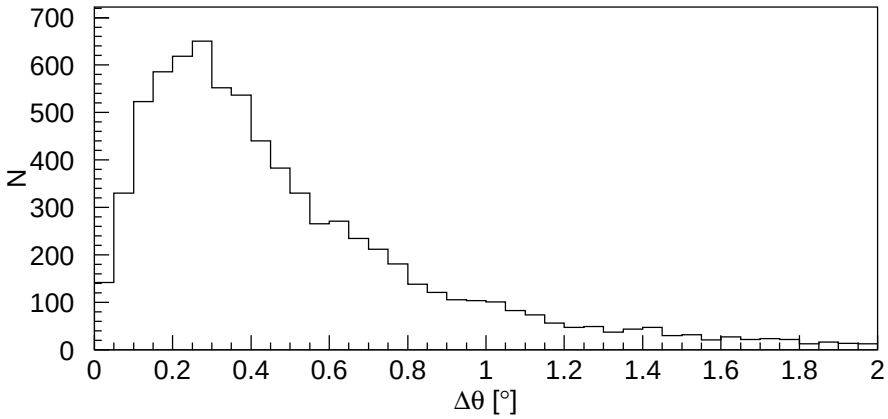

Figure 5.1: The distribution of opening angles $\Delta\theta$ between the true and reconstructed arrival directions in simulated HiRes data above 10^{19} eV. 68 % of events are reconstructed within 0.58° of their true arrival direction.

the simulated (true) and reconstructed arrival directions—is shown for events above 10^{19} eV in the simulated data set. It is found that 68 % of the events are reconstructed within 0.58° of their true arrival direction ($\Delta\theta_{68} \approx 0.6°$).

In Fig. 5.2, the zenith angle dependence of these angular reconstruction errors is shown. This is the only plot where events with $z > 70°$ are included. One can see that the errors grow worse for zenith angles beyond 70°, and also that there are few events reconstructed in this range; hence the cut. In azimuth angle, the angular resolution is essentially constant.

In Fig. 5.3, the energy dependence of the angular reconstruction errors is shown. The opening angle is seen to be essentially constant between 10^{18} eV and 10^{20} eV, and is in fact optimal around $10^{18.5}$ eV. For events below 10^{18} eV, however, the angular errors grow worse, which again corresponds to the energy region where few showers can be reconstructed.

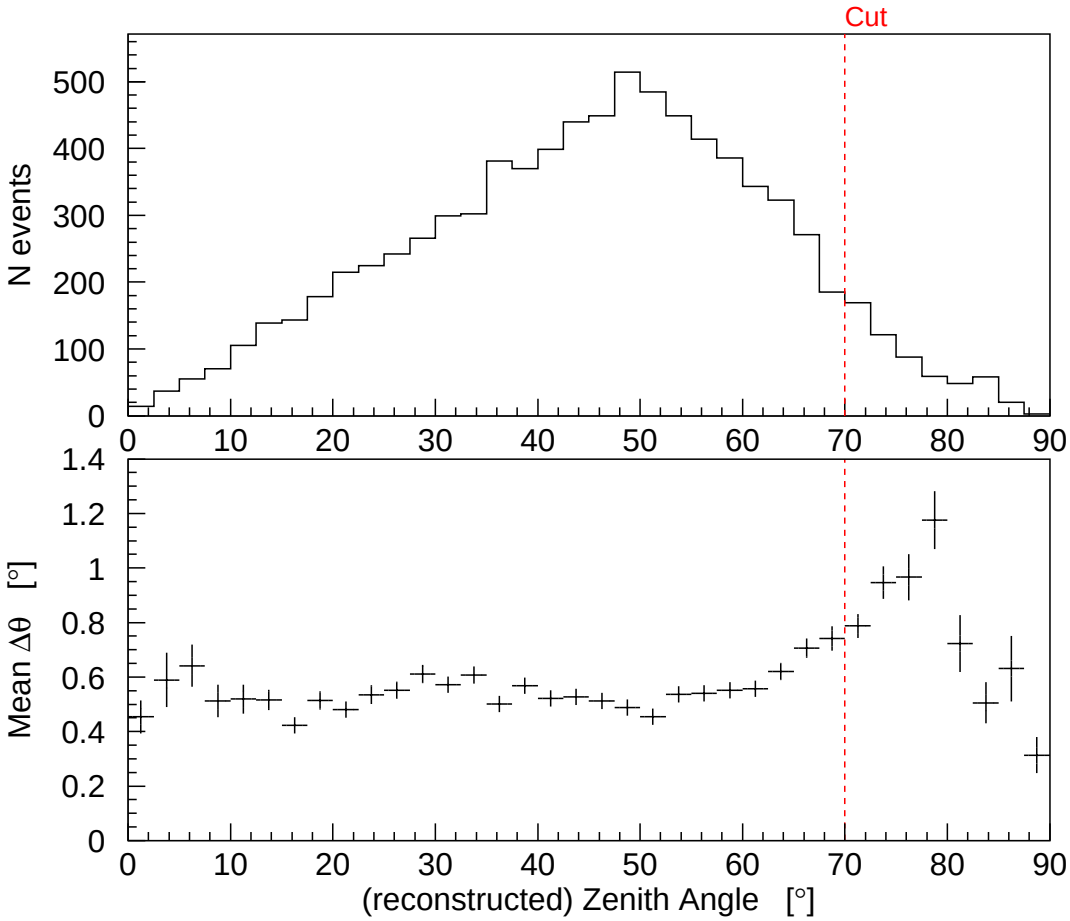

Figure 5.2: *Above:* Zenith angle distribution of reconstructed events above 10^{19} eV in the simulation. Note that the events with zenith angle beyond 70° are cut in the analysis of arrival directions (see Ch. 4.3). *Below:* Mean value of the opening angle distribution as a function of zenith angle, for the same simulated set of events.

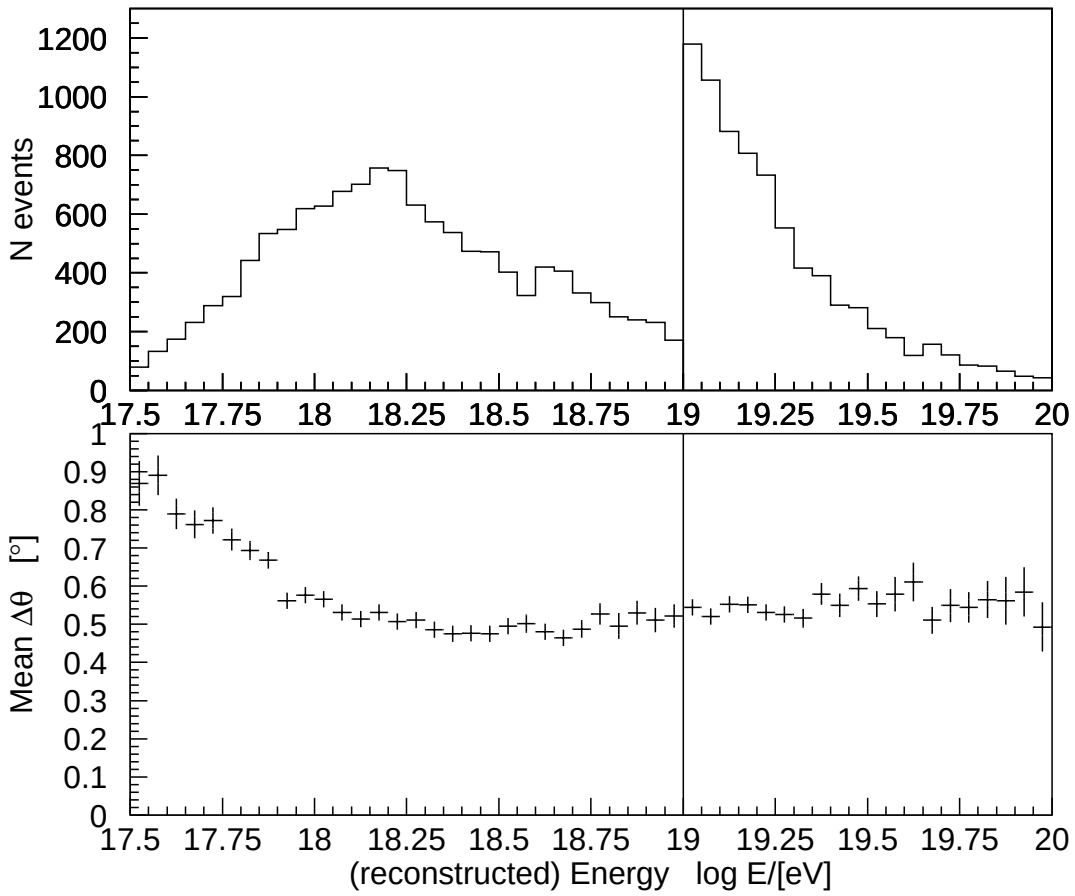

Figure 5.3: *Above:* Energy distribution of reconstructed events in simulation; note that separate simulated event sets are used above and below 10^{19} eV, in order to yield adequate statistics at all energies while maintaining an E^{-3} simulated spectrum for both sets. The drop in the distribution at low energies is due to the fact that the detector is not fully efficient below $10^{18.5}$ eV. *Below:* Mean value of the opening angle distribution as a function of energy, for the events shown above.

5.1.1 Modeling the Errors using Fits to the Error Distribution

The opening angle distribution of simulated events shown in Fig. 5.1 can be used to make a simple model for the angular error distribution expected in the data. In the simplest case, if the point-spread function is assumed to be circular Gaussian (that is, with the same σ in the 'x' and 'y' directions), then the distribution of opening angles $P(\Delta\theta)$ is given by (see Appendix Sec. A.2):

$$P(\Delta\theta) = \frac{\Delta\theta}{\sigma^2} \exp\left(-\frac{(\Delta\theta)^2}{2\sigma^2}\right) . \tag{5.1}$$

It is easy to show (Appendix Sec. A.3) that the value $\Delta\theta_{68}$ which encloses 68.27 % of the distribution is $\Delta\theta_{68} = 1.5152\,\sigma$. For HiRes events above 10^{18} eV, the value is $\Delta\theta_{68} \approx 0.6^\circ$, and therefore if the errors are indeed Gaussian and circular then the corresponding sigma is $\sigma \approx 0.4^\circ$.

In Fig. 5.4, the distribution $P(\Delta\theta)$ for a fixed Gaussian $\sigma = 0.4^\circ$ is compared with the HiRes distribution of opening angles shown earlier in Fig. 5.1, normalizing to the same number of events. One can see that the actual distribution consists of a large population of somewhat better resolved events, and a tail of somewhat worse events. It is possible to fit this distribution well using a sum of two Gaussian opening angle distributions:

$$P_{\rm II}(\Delta\theta) = w\frac{\Delta\theta}{{\sigma_1}^2} \exp\left(-\frac{(\Delta\theta)^2}{2{\sigma_1}^2}\right) + (1-w)\frac{\Delta\theta}{{\sigma_2}^2} \exp\left(-\frac{(\Delta\theta)^2}{2{\sigma_2}^2}\right) , \tag{5.2}$$

where the best fit values are $\sigma_1 = 0.24^\circ$, $\sigma_2 = 0.62^\circ$, and $w = 0.6$. This is shown in Fig. 5.5.

5.1.2 Modeling the Errors Using Individual Uncertainty Estimates

Another method of modeling the angular errors is to make use of the uncertainty estimates resulting from the reconstruction. The stereo trajectory reconstruction provides estimates of the zenith uncertainty σ_z and the azimuth uncertainty σ_a. In the following, the 'scaled-azimuth' ($\widetilde{\sigma_a} = \sigma_a \cdot \sin\theta_z$) will be used for clarity so that both uncertainties are measured in real degrees on the sky.

Therefore for each event there are uncertainty estimates $(\sigma_z)_i$ and $(\widetilde{\sigma_a})_i$, which we use to define the uncertainty σ_i of the ith event:

$$\sigma_i = \left(\frac{(\sigma_z)_i^2 + (\widetilde{\sigma_a})_i^2}{2} \right)^{1/2} . \tag{5.3}$$

In principle, each σ_i should be a better estimate of the ith event's error than could be achieved using a fixed σ value derived from a fit of all event errors. This will be the case to the extent that: 1) the angular errors in the reconstruction are generally circular, and 2) the individual zenith and azimuth errors are indeed Gaussian distributed with the estimated widths $(\sigma_z)_i$ and $(\sigma_a)_i$, respectively.

Fig. 5.6 illustrates a test of this proposition. What is plotted here is the distribution of $(\Delta\theta)_i/\sigma_i$ values for the same simulated set of events as in Fig. 5.4. That is, each opening angle is scaled by its individual angular uncertainty. Scaled this way, the distribution should obey

$$P(s) = s \exp\left(-\frac{s^2}{2}\right) , \tag{5.4}$$

where $s = (\Delta\theta)_i/\sigma_i$ and is distributed with unit variance.

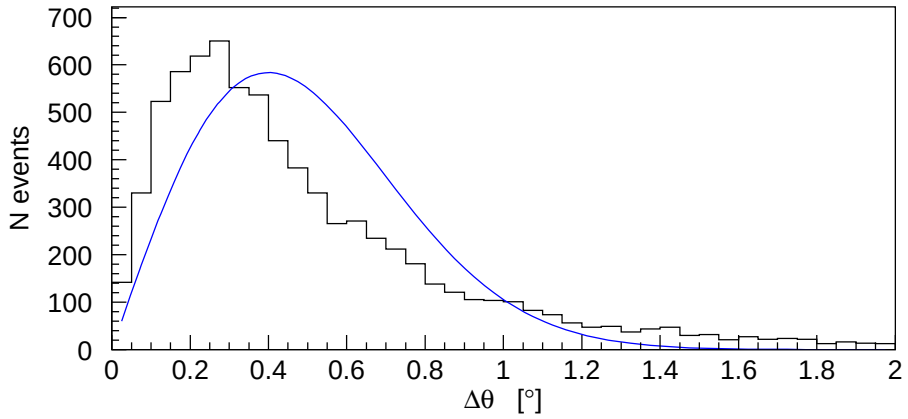

Figure 5.4: The same distribution of opening angles shown in Fig. 5.1 for simulated events above 10^{19} eV. The curve shows the $P(\Delta\theta)$ distribution expected for this number of events, assuming a fixed Gaussian sigma $\sigma = 0.4°$ for the entire sample.

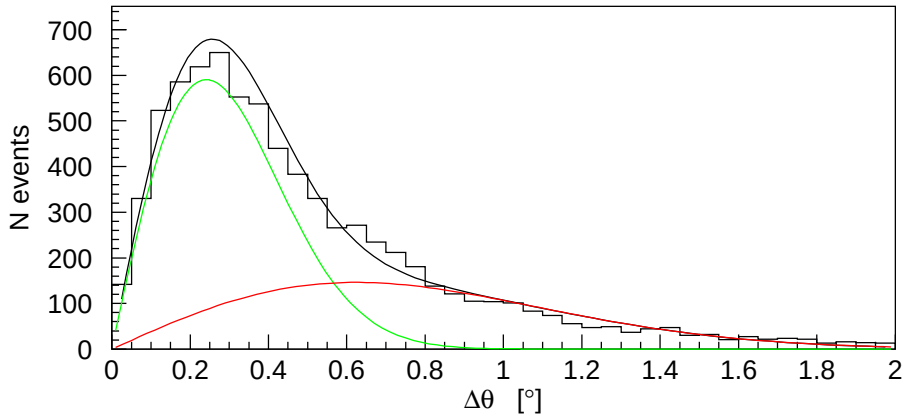

Figure 5.5: The same distribution of opening angles shown in Fig. 5.4, now fit by a sum of two Gaussian distributions $P_{\rm II}(\Delta\theta)$ (Eq. 5.2). The green curve shows the contribution from the best-fit $\sigma_1 = 0.24°$ distribution, and the red curve from the best-fit $\sigma_2 = 0.62°$ distribution, with fractional weights 0.6 and 0.4, respectively. The black curve shows the sum of the distributions.

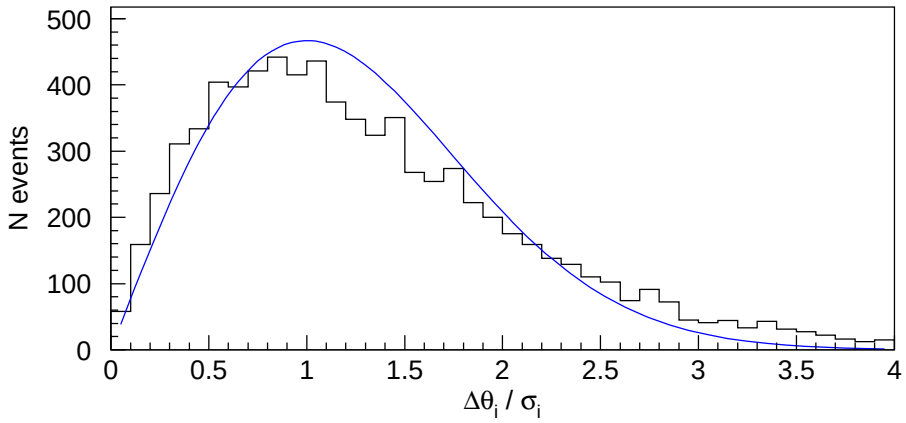

Figure 5.6: The same distribution of opening angles shown in Fig. 5.4, but now each opening angle $(\Delta\theta)_i$ is scaled by its uncertainty σ_i in the reconstruction. The curve shows the distribution expected if the reconstruction errors are circular and the uncertainties are correct Gaussian estimates.

The agreement in Fig. 5.6 compared with Fig. 5.4 indicates that the individual σ_i uncertainties are good Gaussian estimators of the angular reconstruction errors, compared with a fixed uncertainty estimate of $\sigma = 0.4°$. While the match to a double Gaussian in Fig. 5.5 may appear better still, the individual uncertainty estimates σ_i have an important property: maximum likelihood methods can make use of the individual event uncertainties in such a way as to optimize the information extracted. Effectively, more weight is given to the events with the smallest uncertainties. The analysis of correlations between HiRes events and BL Lacertae objects in Ch. 10 applies this model of the individual angular errors.

5.2 Angular Resolution of the Real Detector

5.2.1 Systematic Errors and Detector Alignment

There are two interrelated checks of the accuracy of arrival direction reconstruction for real data. The first method is a direct check of the mirror pointing and alignment of the photomultiplier tubes and clusters. This has been studied using star light in Sadowski et al. [52]. When the image of a UV bright star passes across a PMT, the light signal increases, and this is compared against the expected signal based on the known position of the star in catalogs.

The second method involves reconstruction of the laser tracks which are used for atmospheric monitoring, and serves as a check on both the mirror pointing and the reconstruction process applied to real detector data. There are a variety of fixed and movable lasers used by HiRes in addition to the steerable lasers at each site. Comparison of the reconstructed position of the beam with the actual GPS-measured location provides a test of both the detector alignment and the accuracy of the trajectory reconstruction.

Using both methods, no evidence has been found for any systematic shift as large as or larger than 0.2°. Therefore it is expected that the angular reconstruction error for cosmic ray events is dominated by the statistical measurement uncertainty $\sim 0.5^\circ$.

5.2.2 Comparison of Real and Simulated Event Uncertainties

While the angular error of reconstructed real events cannot be known, the estimated uncertainties for real and simulated events can be compared. In Fig. 5.7, the distribution of σ_i values for the 4495 real HiRes events (all energies) is compared with the σ_i distribution for simulated events (all energies, with simulated E^{-3} energy spectrum).

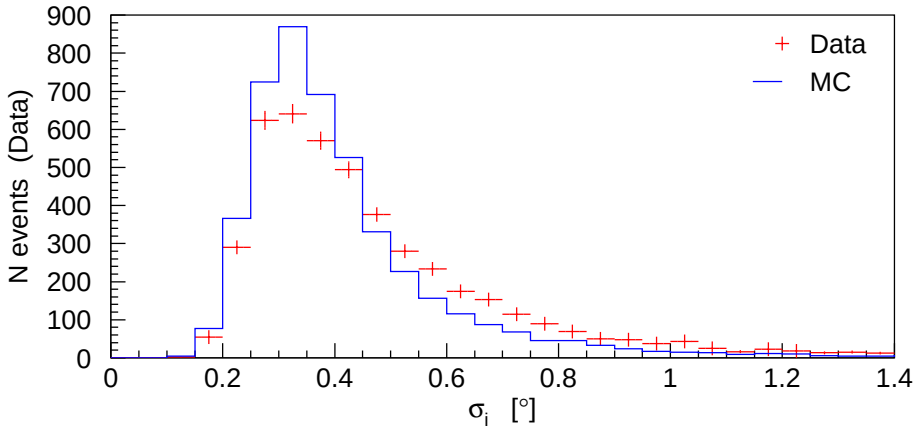

Figure 5.7: The distribution of estimated σ_i values for real HiRes data of all energies (4495 events), compared with σ_i values for simulated events with a similar E^{-3} energy spectrum (number of events normalized to match real events in plot.)

The distributions are similar, with the σ_i values in the real data tending to be slightly larger than in the simulated events (the means of the two distributions are 0.47° and 0.41°, respectively). It comes as no surprise that the real world is somewhat messier than simulation. Small effects from systematic pointing errors discussed above can be expected to contribute to the estimated uncertainty, since geometric irregularities in the pointing make it harder to fit the trajectory consistently, leading to larger uncertainties in the reconstruction.

Another comparison of the estimated uncertainties is shown in Fig. 5.8. The upper plot shows simulated events, comparing the estimated uncertainties and the actual angular errors as a function of energy. To facilitate the comparison, recall that the mean $\langle \Delta\theta \rangle$ of a Gaussian opening angle distribution is given by $\langle \Delta\theta \rangle = \sigma \, (\pi/2)^{1/2}$ (see Appendix Sec. A.2). Therefore, if the opening angles $\Delta\theta$ are distributed according to $P(\Delta\theta)$ (Eq. 5.1) with width $\sigma_{\rm E}$ in each energy bin, where $\sigma_{\rm E} \approx \langle \sigma_i \rangle$, then it should be the case that $\langle \Delta\theta \rangle / 1.253 \approx \langle \sigma_i \rangle$ in each bin. These two quantities are indeed seen

to agree well in the figure, indicating that the uncertainties are good estimates of the angular error (except for a tendency to underestimate the error at energies below 10^{18} eV where few events are reconstructible).

In the lower plot of Fig. 5.8, the σ_i values of the same simulated events are now compared with those of the real data. The two show a similar dependence on energy, with the uncertainties for the real events consistently slightly higher, as expected. Taken together, the two plots increase our confidence that the simulation is a good model of the experiment, and that the angular uncertainties estimated by the reconstruction are reliable.

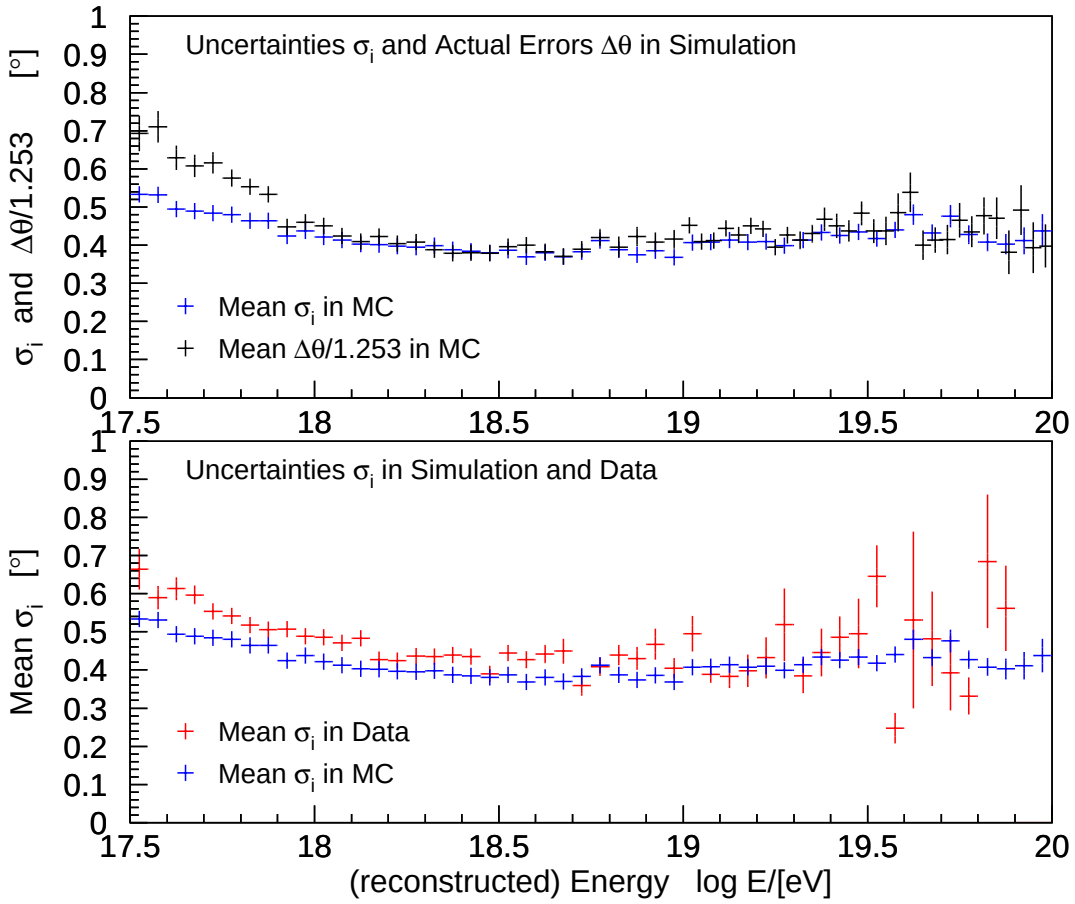

Figure 5.8: *Above:* Mean σ_i as a function of energy for simulated events. As described in the text (Sec. 5.2.2), the actual angular error can be compared with the estimated uncertainty by plotting $\langle \Delta\theta \rangle / 1.253$ and $\langle \sigma_i \rangle$, if the errors are Gaussian distributed according to $\langle \sigma_i \rangle$ in each energy bin. One can see that the estimated uncertainties and actual errors agree well in the simulation, except below 10^{18} eV, where the angular reconstruction error is somewhat worse than estimated (and where few events are reconstructible; see Fig. 5.3).
Below: Mean σ_i for the same set of simulated events as above, compared with the mean σ_i values for the real HiRes data (4495 events). The two sets illustrate the same general dependence on energy, with the mean uncertainty for real events consistently slightly higher than in the simulation. (Note that above $10^{19.5}$ eV, the number of real events is small and comparisons are more difficult.)

Chapter 6

Previous Claims of Small-Scale Anisotropy Observed by AGASA

All major experiments to date have shown that the distribution of arrival directions of ultrahigh energy cosmic rays is, to first order, remarkably isotropic. From time to time, however, evidence for small-angle clustering among the highest energy events has been claimed, most recently and notably by the AGASA (Akeno Giant Air Shower Array) [53] experiment [54–59]. However, estimates of the chance probability of this clustering signal vary from 10^{-2} to 10^{-6} and beyond. Therefore it is essential to assess correctly the significance of this evidence in order to compare it with anisotropy studies of other cosmic ray experiments.

The focus on the highest energy events is well-motivated, since these will presumably be the least deflected after traversing galactic and extragalactic magnetic fields. However, this involves some important details which sometimes go unmentioned. Identifying the highest energy events typically requires making a choice for

the minimum energy E_c that defines the data set. Also, small-scale anisotropy analyses often involve identifying close pairs of events, requiring a choice for the maximum angular separation θ_c that defines a pair. On the one hand, choosing a higher energy threshold E_c should reduce deflections and allow clusters to show up within smaller angular separations θ_c. This holds especially for ground array detectors, such as AGASA and Auger, in which the angular resolution improves at higher energies. On the other hand, as a function of energy E the cosmic ray flux drops faster than E^{-2}, so the statistical power of the available data quickly weakens with higher energy thresholds.

For a precise model of cosmic ray source distributions and Galactic and extragalactic magnetic fields, these competing forces would imply optimal choices for E_c and θ_c to maximize the clustering signal. At present, however, not nearly enough is known about any of these to make *a priori* choices useful. Instead, what is typically done, explicitly or implicitly, is to scan over a range of values for E_c and θ_c, and identify the values which maximize the clustering signal. In this case, the final significance of the result must include a penalty factor for the *a posteriori* cuts arrived at by scanning.

As will be shown, it is the various ways of handling this penalty factor—or, in some cases, the failure to include it at all—which leads to the wide range of significances attached to the AGASA clustering signal.

6.1 AGASA Claims of Small-Scale Clustering

The AGASA experiment reported possible clustering in the arrival directions of ultra-high energy cosmic rays as early as 1996 [54], and has updated this data sample and

analysis in several publications [55–58]. The first report of clustering in 1996 identified three pairs of events with angular separation less than 2.5° among the 36 events with energies above 4×10^{19} eV. The corresponding chance probability was found to be 2.9%. It was noted that the angular separation of 2.5° is "nearly consistent with the measurement error ($\sqrt{2} \times 1.6°$)" [54]. The minimum energy of 4×10^{19} eV was justified under the assumption that the GZK cutoff should lead to an accumulation of events around 4×10^{19} eV, and therefore that events above this energy may point back to nearby sources. The values for E_c and θ_c identified in this report set the stage for all analyses which followed.

In 1999, a new publication by AGASA [55] identified a stronger clustering signal using these cuts with an enlarged data set now containing 47 events. The following year, AGASA published an updated list with 57 events above 4×10^{19} eV observed through 2000 May [56]. There is also an additional event below 4×10^{19} eV which was added to the list because it forms another doublet. (This is an unfortunate source of confusion. Following many authors, this extra event is not included in this analysis because it is not clear how many additional events there are between it and 4×10^{19} eV.) Not counting the extra event, there are four doublets and one triplet in this set. A skyplot of these events is shown in Fig. 6.1.

This set was analyzed by Tinyakov and Tkachev [60], who calculated the chance probability as a function of the threshold energy E_c of the data set, while keeping the angular bin size constant at 2.5°. The lowest probability was found to be less than 10^{-4} with $E_c = 4.8 \times 10^{19}$ eV. Since this probability was obtained by scanning over energies, it does not reflect the true significance of the clustering signal. To estimate the correct chance probability, the authors numerically calculated a correction factor by generating 10^3 random sets of events which were then subjected to the same

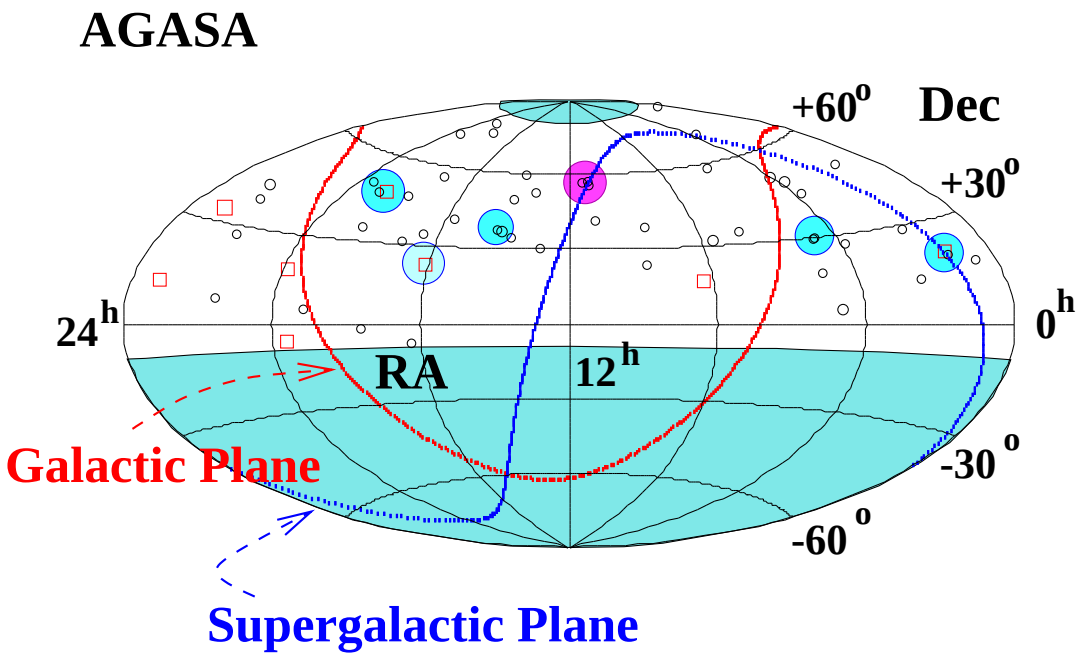

Figure 6.1: Skyplot of arrival directions of the 57 events observed by AGASA through 2000 May. Open circles represent cosmic rays with energies between 4×10^{19} eV and 10^{20} eV. Squares represent cosmic rays with energies above 10^{20} eV. The large blue circles identify pairs of events separated by less than 2.5°; the purple circle identifies a triplet of such events. From Hayashida et al. [56].

scanning in E_c. It was found that 27 (3) random samples had a probability of less than 10^{-2} (10^{-3}), and the authors concluded that the correction factor was of order 3. The final chance probability was given as 3×10^{-4}, considerably lower than the chance probability reported by the AGASA collaboration in the original publications [54, 55].

A similar scan was then performed over the size of the angular bin, *i.e.* the maximum angular distance between events that defines a cluster. The probability shows a minimum at 2.5°, but since this was interpreted as the angular resolution of the experiment, no correction factor was applied to the final chance probability.

In Takeda et al. [57], the AGASA group applied this scanning technique again to a data set which was now reported to include 59 events above 4×10^{19} eV—essentially the same data set as the one published in 2000 [56], though it is unclear whether the one event below the energy cutoff was kept, or whether one or two new events were added. Five doublets and one triplet were reported in the sample. A scan over angular separations was again performed, showing the peak at 2.5°. Performing a scan over energies, the significance of the clustering above 4×10^{19} eV was said to be $4.6\,\sigma$, and above 4.5×10^{19} eV it was said to be in excess of $5\,\sigma$. No statistical penalties were applied for either the energy or angular separation scan.

The most recently published study by AGASA [58] in 2003 recapitulates much of the above analysis. The same 59 events are analyzed, though by the end of July 2002 there were already 72 events above 4×10^{19} eV which had been reported [59]. (The AGASA experiment continued to operate through the end of 2003.) Forgoing a scan over energies, the chance probability for all of the clusters (one triplet + five doublets = eight pairs) in the total set of 59 events is simply reported to be less than 10^{-4}.

6.2 Angular Two-Point Correlation Function

A standard tool for studying anisotropy is the angular two-point correlation function, and we will make use of it here to examine the anisotropy claim of AGASA in more detail. Various estimators for the correlation function exist (see e.g. Landy and Szalay [61]). Here, we define the estimator as

$$w(<\theta) = \frac{n_{\rm p}}{\langle n_{\rm MC}\rangle} - 1 \ , \tag{6.1}$$

where $n_{\rm p}$ is the number of pairs of events in the data sample with angular separation less than θ, and $\langle n_{\rm MC}\rangle$ is the average number of such pairs in simulated isotropic sets with the same number of events and same detector acceptance in right ascension and declination as the data sample. Note that this definition of $w(<\theta)$ is cumulative over angles up to θ. It reveals how the correlation signal varies as a function of the angular threshold for defining a pair.

Fig. 6.2 shows the results of the angular correlation estimate for the 57 published AGASA events with energies above $4 \times 10^{19}\,$eV. The top plot shows the number of pairs of events with angular separation less than θ, for both the data and simulated sets. The middle plot shows the angular correlation estimate $w(<\theta)$. The errors shown are Poisson, $\sqrt{n_{\rm p}}/\langle n_{\rm MC}\rangle$, which give some idea of the deviation from isotropy. However, the variance of the correlation estimator $w(<\theta)$ is in general larger than the Poisson variance (see Landy and Szalay [61]). This can be understood because pair-counts are not an independent Poisson process: when a new event is added to a data set, the likely number of new pairs it creates depends on the number of pairs already present. Instead, the significance of the correlation is estimated by using simulations. The bottom plot shows the fraction $F_{\rm MC}$ of simulated event sets with

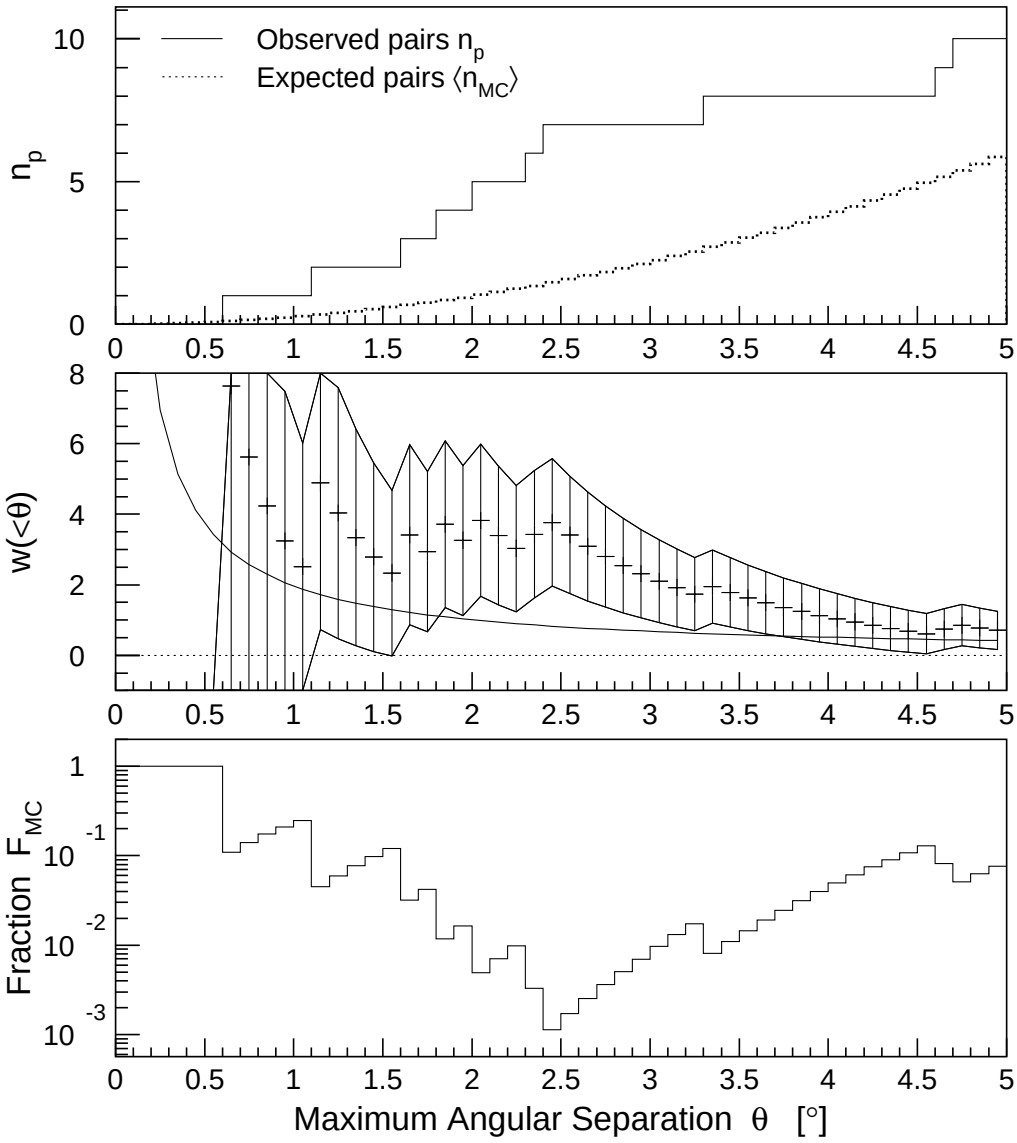

Figure 6.2: Angular two-point correlation estimate w and related quantities for the AGASA data set. *Top:* Number of pairs of events with angular separation less than θ, where n_p is the number observed in the data and $\langle n_{\rm MC} \rangle$ is the mean number of pairs found in simulated sets. *Middle:* Angular correlation estimate $w = n_p/\langle n_{\rm MC} \rangle - 1$ for the AGASA data set. Poisson errors $\sqrt{n_p}/\langle n_{\rm MC} \rangle$ are plotted. The mean for isotropic distributions is 0 (dotted line); the solid curve shows the Poisson error for isotropic distributions. *Bottom:* Fraction of simulated sets with the same or greater number of pairs than the n_p observed in the data.

the same number or more pairs as observed in the real data, n_p.

It can clearly be seen that the most significant deviation from isotropy (i.e. smallest value of $F_{\rm MC}$) in the AGASA data set occurs using a bin size of precisely $\theta_c = 2.5^\circ$. If this exact bin size were chosen *a priori*, and if the data sample itself (that is, the energy threshold $E_c = 4 \times 10^{19}\,\mathrm{eV}$ defining the sample) were chosen *a priori*, then $F_{\rm MC}$ would be the chance probability of the observed correlation.

6.3 Question of Significance of Claims

In evaluating the significance of the clustering signal, it is essential to determine whether the original choices of $E_c = 4 \times 10^{19}\,\mathrm{eV}$ and $\theta_c = 2.5^\circ$ were in fact *a priori*.

Consider what would have been required to formulate such an *a priori* hypothesis. In the case of the angular resolution of the experiment, Monte Carlo simulations can be used to determine the optimal angular size for a cluster search. Such a study needs to take into account that the angular resolution for a ground array depends on a variety of factors. For the AGASA detector, the angular error continues to shrink with increasing energy, as shown in Fig. 6.3. At $10^{20}\,\mathrm{eV}$, AGASA reports $\theta_{err} < 1.2^\circ$ [55]. Eight of the 57 events in the data set are in fact above this energy.

In addition, the angular resolution of ground arrays depends on the alignment of the shower with the detector array. In general the errors will be asymmetric. The ratios of the 68 % and 90 % opening angles seen in Fig. 6.3 are clearly not those of a circular, two-dimensional Gaussian distribution. These complications mean that there is no reasonable justification for the statement that the optimal angle for a cluster search is simply $\sqrt{2} \times 1.6^\circ$.

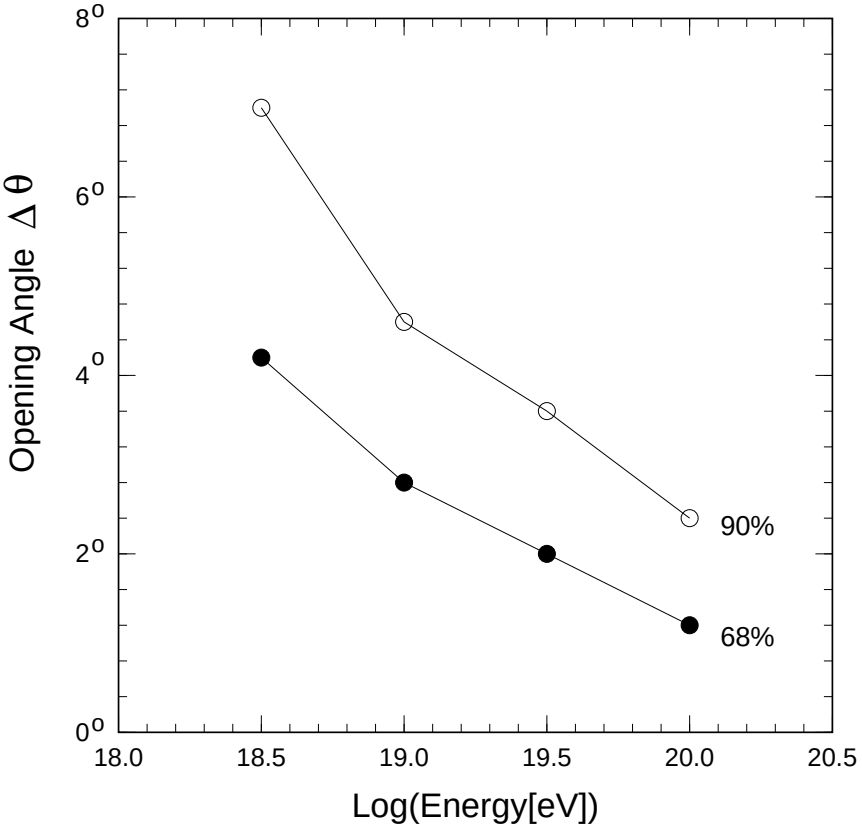

Figure 6.3: Accuracy of arrival direction determination for simulated AGASA events: the angular distances which enclose 68% and 90% of the opening angle distribution (i.e. differences between simulated and reconstructed arrival directions) as a function of energy are shown. From Takeda et al. [55].

Furthermore, the search angle which optimizes the clustering signal will also depend on the expected background of chance clusters. For small data sets, the chance occurrence of a pair is small, and the signal to noise ratio can be optimized with a larger separation angle in the search [62].

In summary, a clustering search which is *a priori* should begin by first using a Monte Carlo simulation to identify the optimal opening angle size. The first clustering paper [54] gives no indication that such a search program was undertaken, nor does it claim that 2.5° is an *a priori* choice. It merely observes that the clustering signal is strongest for $\theta_c = 2.5°$, a value which coincides to some extent (but only approximately, as $\sqrt{2} \times 1.6° = 2.26°$) with the angular resolution of the experiment around the given energy. This makes the value interesting, but not *a priori*.

The scan over threshold energies E_c is motivated on the physical grounds described earlier: higher energy events should be less deflected, and therefore yield a a significant clustering signal on smaller angular scales. It is noteworthy that the original paper [54] does not restrict the analysis to 4×10^{19} eV, but mentions at least two other energy thresholds that were looked at as well (5×10^{19} eV and 6.3×10^{19} eV). This approach is certainly valid. However, it does not constitute an *a priori* search program, which demands a choice for E_c and θ_c prior to examination of the data. Because the values of E_c and θ_c are determined by examining the data, a calculation of the *a priori* probability does not represent the true significance of the observation. Either the cuts must be tested with independent data, or the statistical penalty must be evaluated and included in the calculation of the chance probability.

Chapter 7

Evaluating the AGASA Signal Using an Autocorrelation Scan

In Chapter 6, it was shown that the evidence claimed by AGASA for clustering of cosmic ray arrival directions depended crucially on the choice of the energy threshold defining the data sample and the angular separation defining an event pair. However, it was also shown that these values cannot, at present, be determined *a priori*, and therefore the significance of the claim is called into question.

In this chapter, a general method to address this problem is proposed. A clustering signal among the highest energy events can be best identified by scanning *simultaneously* over energy thresholds and angular separations to find the values for E_c and θ_c which optimize the signal. The chance probability of the signal is determined by counting the number of simulated data sets which yield a stronger signal under an identical scan. With this procedure, the statistical significance is determined without treating *a posteriori* cuts as *a priori* ones. This method, called the autocorrelation scan, is described in Sec. 7.1 and applied to the AGASA data in Sec. 7.2.

However, any analysis of the full AGASA data set remains biased, because the thirty original events that led to the clustering hypothesis in the first place are still included in the data sample. One can avoid this bias by removing the early data and only scanning over the events which have been detected since the original claim. Alternatively, one can test the AGASA clustering hypothesis by applying the original cuts to the newer events directly. Since the cuts are now *a priori*, this test requires no statistical penalty. It has the virtues of being simple and rigorously unbiased. Both tests are described and performed in Sec. 7.3.

7.1 Method: Autocorrelation Scan

One approach to an unbiased small-scale clustering search is to perform an autocorrelation test repeatedly while scanning over energy thresholds and angular separations. Essentially, we consider the set of N events above energy E, count the number of pairs n_p of events separated by less than θ, and evaluate the probability $P(N, \theta)$ of finding this number or more pairs, given N and θ. We repeat this for a range of values for E and θ, and use the smallest probability P_{min} found in the scan to identify the strongest clustering signal. To estimate the true statistical significance of this signal, we perform identical scans over simulated sets of isotropically distributed data, counting the fraction of simulated sets which yield the same or smaller value for P_{min}. This fraction is P_{ch}, our estimate for the significance.

The virtue of this approach is that by letting the energy threshold and angular separation vary, we let the scan itself determine the optimal balance between the better statistics of the low energy data set and the (presumably) smaller angular deflections at high energies. In the case of ground arrays there is an additional

advantage because the angular resolution of the detector typically improves at higher energies, which serves as another competing effect against the larger statistics at lower energies.

It should be noted that, just as in the usual two-point correlation function, higher-order multiplets are counted by the individual number of pairs which they contain. A triplet of events, for example, will be counted as two or three pairs, depending on the individual separations of the three events.

To determine the probabilities $P(N, \theta)$, we generate a large number of simulated data sets (typically 10^7) corresponding to an isotropic distribution of cosmic rays, corrected for the detector variation in relative exposure over the sky. We then construct a table of values P_{MC}, where $P_{MC}(N, \theta, n)$ is the fraction of data sets in which the first N events contain exactly n pairs separated by less than θ. Then the probability $P(N, \theta)$ for observing n_p or more pairs at (N, θ) is simply:

$$P(N, \theta) = \sum_{n=n_p}^{\infty} P_{MC}(N, \theta, n) = 1 - \sum_{n=0}^{n_p - 1} P_{MC}(N, \theta, n). \tag{7.1}$$

For some combination N_c and θ_c, P has a minimum: $P_{min} = P(N_c, \theta_c)$. We identify this as the strongest potential clustering signal. To determine the statistical significance, we perform the same scan over n_{MC} Monte Carlo data sets, finding the minimum probability $P^i_{min} = P^i(N^i_c, \theta^i_c)$ for each trial and counting the number of trials n^*_{MC} for which $P^i_{min} \leq P_{min}$. The significance is finally identified as:

$$P_{ch} = \frac{n^*_{MC}}{n_{MC}}, \tag{7.2}$$

that is, the chance probability of observing the value P_{min} or less in an isotropic distribution.

In practice, rather than use an arbitrary fixed increment ΔE of energy, it is often simpler to rank the events by energy and, starting at the top, decrease the energy threshold one event at a time ($\Delta N = 1$). For the set of N_c events, the lowest energy in the set corresponds to the energy threshold E_c.

7.2 Autocorrelation Scan of the AGASA Data Set

This scan is performed on the published AGASA data above 4×10^{19} eV, which consists of 57 events [56]. To generate Monte Carlo events for determining the probabilities, we follow Tinyakov and Tkachev [60] in using a zenith angle (θ_z) distribution $dn \propto \cos\theta_z \sin\theta_z d\theta_z$, corresponding to geometric acceptance of isotropically distributed cosmic ray arrival directions. We use the same $\theta_z < 45^\circ$ cut as employed by AGASA, and assign uniformly random arrival times, corresponding to the uniform exposure of AGASA in right ascension [56, 57]. We scan over angular separations from 0° to 5° in increments of 0.1°. The results of the scan are shown in Figure 7.1.

The strongest clustering signal is contained within the $N_c = 36$ highest-energy events, where there are $n_p = 6$ pairs separated by less than $\theta_c = 2.5^\circ$. (The energy threshold E_c corresponding to this subset is 4.89×10^{19} eV.) At this spot $P = P_{min} = 8.4 \times 10^{-5}$, that is, 839 out of 10^7 MC data sets had the same or greater number of pairs at the same values for N and θ. This value for P_{min} (not P_{ch}) is essentially the same as the 10^{-4} probability found in Tinyakov and Tkachev [60] for the same energy threshold and angular separation.

To evaluate the significance of this result, we perform the same scan over simulated AGASA data sets and count how many simulated sets have $P_{min}^{MC} \leq P_{min}^{data}$. We find that 3475 out of 10^6 simulated sets meet this condition, implying a chance probability

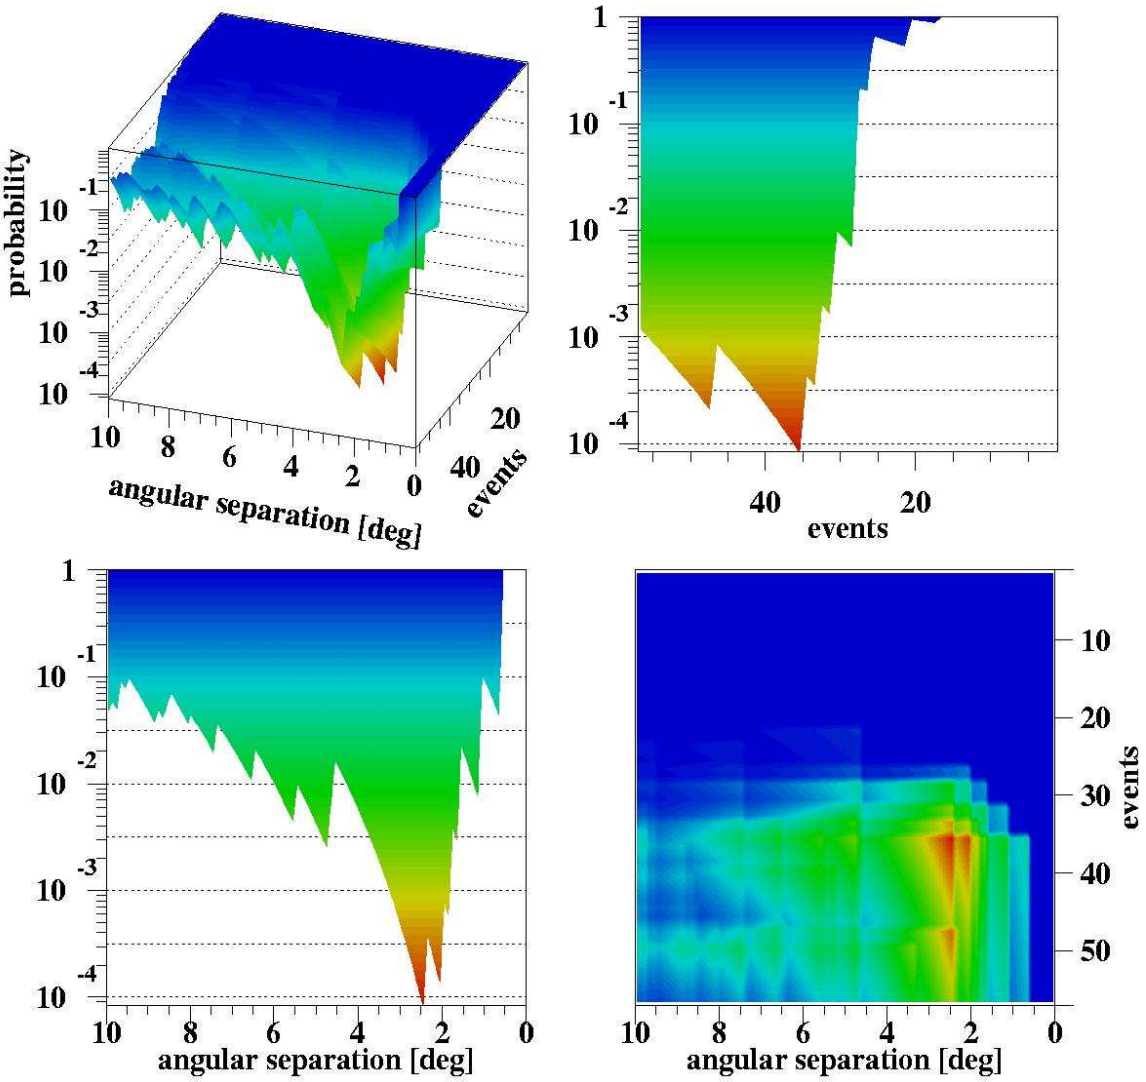

Figure 7.1: Scan of AGASA events above 4×10^{19} eV, shown in four different views. $P_{min} = 8.39 \times 10^{-5}$ and $P_{ch} = 0.3\%$ for the clustering signal at $N_c = 36$, $\theta_c = 2.5°$ with $n_p = 6$ pairs. (N_c corresponds to an energy threshold $E_c = 4.89 \times 10^{19}$ eV.)

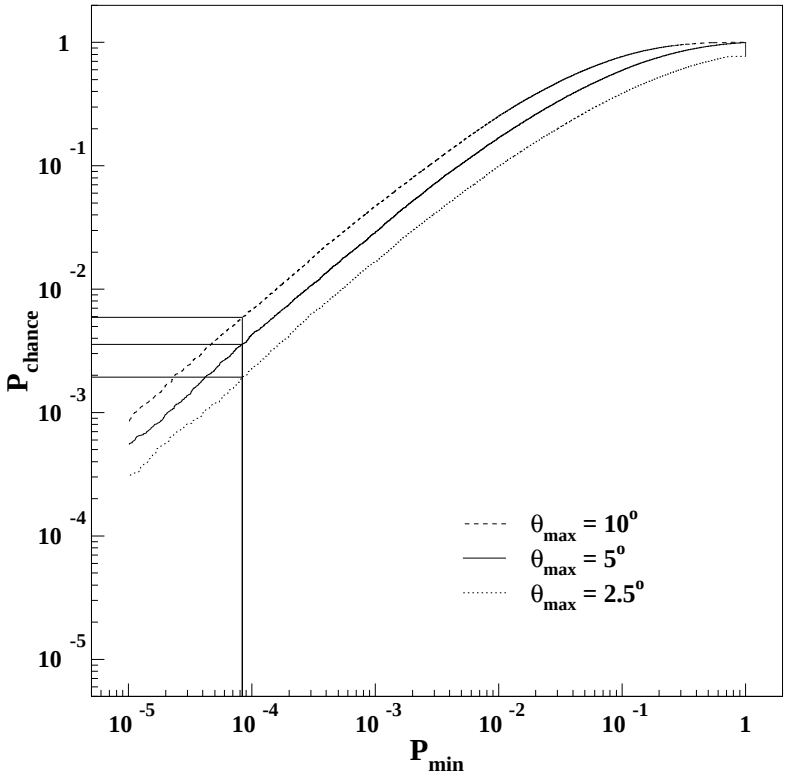

Figure 7.2: P_{ch} as a function of P_{min} for the AGASA data, with $N_{max} = 57$ and three different values of the scan parameter θ_{max}.

of 0.3%. Figure 7.2 illustrates how P_{ch} varies as a function of P_{min} for the simulated AGASA sets.

7.2.1 Robustness of Scan Parameters

In performing the data scan and the simulated scans, it is necessary to choose four parameters which can affect the final result: N_{max}, the total number of events included in the scan; θ_{min} and θ_{max}, the angular extent of the scan; and $\Delta\theta$, the size of the angular binning. In Table 7.1 we show a range of values for these parameters and the

N_{max}	θ_{min}	θ_{max}	$\Delta\theta$	P_{ch}
57	0°	5°	0.1°	3.48×10^{-3}
36	0°	5°	0.1°	2.40×10^{-3}
100	0°	5°	0.1°	5.63×10^{-3}
57	**1**°	5°	0.1°	2.96×10^{-3}
57	**2.5**°	5°	0.1°	2.05×10^{-3}
57	0°	**2.5**°	0.1°	2.00×10^{-3}
57	0°	**10**°	0.1°	5.77×10^{-3}
57	0°	5°	**0.5**°	2.31×10^{-3}
57	0°	5°	**0.02**°	4.05×10^{-3}

Table 7.1: The effect on P_{ch} due to variations in the scan parameters. In each case, P_{ch} was determined using 10^6 Monte Carlo data sets. The top row lists the values used in the text. The parameters which are varied are indicated in bold.

effect on the final value of P_{ch}. We motivate our choices for each of the parameters as follows:

In choosing the extent of the scan in N_{max} and θ_{max}, it is clear that the significance would be biased if one scanned out precisely to the maximum clustering signal and no further. An investigator who reports a clustering signal in a scan can reasonably be expected to have scanned out at least twice as far in search of an even stronger signal; hence a reasonable estimate of P_{ch} should extend N_{max} to $\sim 2 \times N_c$, and further, if N_c is very small. The same can be said for θ_{max} with respect to θ_c. As for θ_{min}, it should be no larger than the best attainable angular resolution; in the present case, one could choose 0° or 1° with little effect on the final probability. Finally, it can be seen in Table 7.1 that reducing the angular bin size $\Delta\theta$ also has a negligible effect at small scales.

We note that over each of the ranges shown in Table 7.1—a factor of three in event number, a factor of 16 in angular area, and a factor of 25 in angular binning—the

chance probability remains within 2×10^{-3} to 6×10^{-3}. Therefore our result does not depend sensitively on these scanning parameters.

7.2.2 Significance Compared with Previous Estimates

The value of 0.3% we calculate for P_{ch} is 10 times larger than that calculated in Tinyakov and Tkachev [60]. Although they use an angular scan to demonstrate that the separation angle 2.5° maximizes the signal, they nevertheless treat the choice as an *a priori* one and make no correction for it.

In Takeda et al. [57], the AGASA collaboration analyzes the same data set and finds that at 4×10^{19} eV the significance of the clustering signal is $4.6\,\sigma$, and that at a slightly higher energy threshold it is "$5\,\sigma$ or more". These results imply chance probabilities of 4.2×10^{-6} and 5.7×10^{-7}, respectively—three to four orders of magnitude lower than the probability we have presented here. This overestimation of the significance of the clustering signal arises in part from the application of Gaussian statistics to a non-Gaussian distribution: these significances are obtained by measuring the excess clustering signal in units of standard deviations, $(N_{obs} - N_{exp})/\Delta N_{exp}$, when in fact this distribution is not Gaussian at all (see Ch. 6.2). Straightforward simulation reveals that assumptions of Gaussianity can vastly overestimate the significance. Having cited Tinyakov and Tkachev [60] and made use of their technique, the authors ignore their warning on exactly this point. Furthermore, they ignore the statistical penalty involved in scanning over energy thresholds, and they do not consider a penalty for the choice of angular separations.

7.3 Unbiased Test of AGASA Clustering Hypothesis

A more rigorous statistical test of the clustering hypothesis can be performed by isolating the data which led to the original cuts of 4×10^{19} eV and 2.5°. Since these values are justified *a posteriori* in conjunction with the observation in 1996 that they lead to a clustering signal [54], they can only be treated as *a priori* for a data set independent of the one which was used to derive them. We can do this by dividing the AGASA data into an "original data set" comprising the events observed through October 1995 which formed the basis of the original clustering claim, and a "new data set" comprising the events which have been observed since then. Using the list of events published in 2000 [56], there are 30 events in the original set and 27 in the new one. [1]

Because the new data set is independent, we can test the original clustering hypothesis directly without the need for any statistical penalties. We simply count the number of pairs of events using $E_c = 4 \times 10^{19}$ eV and $\theta_c = 2.5°$, and we find one pair. The chance probability for one or more pairs among 27 events is 28%.

We can investigate whether there is a better choice of E_c and θ_c for the new data set by performing an autocorrelation scan. Figure 7.3 shows the results of scanning over both the old and new data sets separately. The strongest clustering signal in the original set has a chance probability $P_{ch} = 4.4\%$ and occurs for $\theta_c = 2.4°$ and $E_c = 4.35 \times 10^{19}$ eV (with 3 pairs among the 26 highest energy events, and a minimum probability $P_{min} = 0.33\%$ in this bin). This confirms that the cuts originally selected

[1] In [54] (1996), the original data set is said to contain 36 events above 4×10^{19} eV. However, the lists published in 1999 and 2000 [55, 56] contain only 30 events during this same time period, due to a reevaluation of the energies (according to Uchihori et al. [63]). The three original clusters are present in all sets.

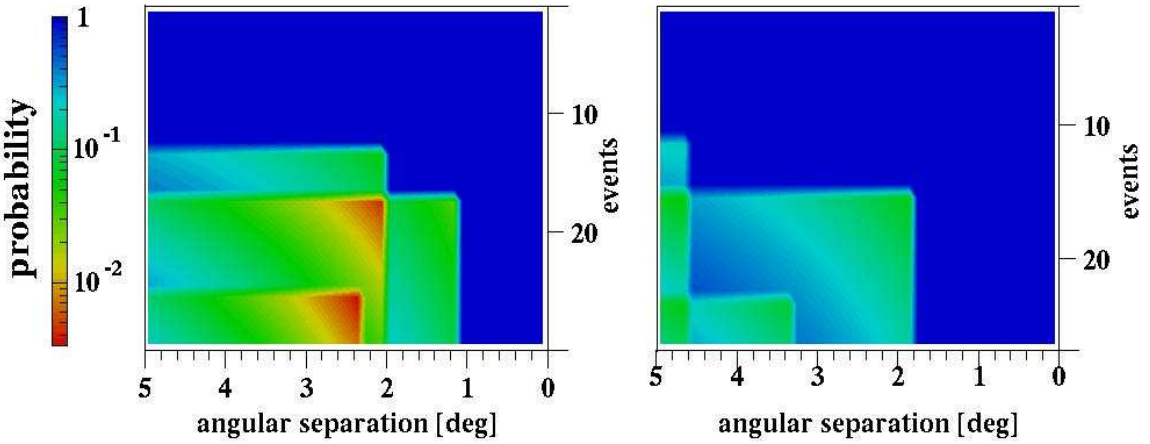

Figure 7.3: Autocorrelation scans for the "original" (left) and "new" (right) AGASA data sets, using October 31, 1995 as the dividing point. The chance probability of the strongest clustering signal in the original data set is 4.4% (at $\theta_c = 2.4°$, $N_c = 26$, $E_c = 4.35 \times 10^{19}$ eV, with $P_{min} = 0.33\%$). In the new data set, the strongest clustering signal has $P_{ch} = 27\%$ (at $\theta_c = 4.7°$, $N_c = 16$, $E_c = 4.97 \times 10^{19}$ eV, with $P_{min} = 5.5\%$).

in 1996 were nearly optimal for that data set. However, when the new events are scanned, there is no hint of clustering at the 2.5° scale or any other angular separation. The "strongest" clustering signal occurs at $\theta_c = 4.7°$ with $P_{ch} = 27\%$.

The independent data set has less statistical power than the total data set. If we estimate that power by counting the number of all possible pairs (1596) among 57 events, then we find that the original data set contains 27% of those pairs, the independent set contains 22%, and the remaining 51% are "cross" pairs between events in the original and new data sets. If we are careful to avoid contamination by the original cuts, then we can extend the statistical power of this test by including the cross-correlation with the original set.

To do this without contamination by the initial 2.5° cut obtained from the original data set, we replace each of the three doublets in the original set with a single event at each of their averaged positions. We then count the number of autocorrelation pairs

in the independent set, and we now add the number of cross pairs between events in the independent and original data sets. There is one auto pair, as before, and there are two cross pairs. To estimate the chance probability, we generate Monte Carlo data sets of 27 events to replace the independent data set, while holding the original set fixed. We count what fraction of these trials have the same or greater number of auto- and cross-correlation pairs. The chance probability for a total of three or more pairs is found to be 8%.

The unbiased test can be extended to include two more years of data since May 2000 which has been summarized on the AGASA web page [59]. There are a total of 72 events above $E_c = 4 \times 10^{19}\,\mathrm{eV}$ through July 2002, which means the independent set has 42 events and roughly double the statistical power as before. This set adds one new autocorrelation pair to the one already present, and no new cross-correlation pairs. Performing the same analyses as described in the text, the chance probability for two pairs within the independent set of 42 events is 19%. The chance probability for a total of four pairs—within the independent set, and between the independent and original sets—is 12%.

We observe that if the cuts had been *a priori* for the first data set, the chance probability for the three pairs among the first 30 events would be 0.8%. Thus it is the clustering in the first data set which dominates the significance for the total set, despite the fact that the first set by itself represents only a fraction of the total number of possible pairs. This is precisely what is to be expected when a small initial set is used to optimize the cuts.

If we had not modified the initial data set by replacing the doublets with their average positions, then there would have been three cross pairs in the above test, instead of two. The difference is due to the events which form the AGASA triplet.

The chance probability for four or more total pairs in this test would have been 3% rather than 8%. Unfortunately, since it was these doublets in the first set which made the 2.5° cut optimal, they cannot be included in a statistically independent test of the hypothesis. In any case, even this biased test confirms that the significance is dominated by the initial data set.

7.4 Summary

An autocorrelation scan of the published AGASA data set finds a chance probability around 0.3% for the clustering signal previously claimed by AGASA and others.

While the scan is generally unbiased, in the present case the significance is still questionable, because the original events which led to the formulation of the clustering hypothesis in the first place are still included in this data set. Therefore the original claim must be tested in a way that excludes the contribution of the original data to the clustering signal. First we form an independent data set using only the AGASA events observed after the claim. The cuts which were identified originally can now be applied *a priori* in an unbiased test. To increase the statistical power of the test, we include cross-correlations with the original data. Replacing the doublets in the original set with single events to keep this test independent of the original clustering signal, we find a chance probability of 8%.

We conclude that the evidence for clustering in the AGASA data set is weaker than previously claimed, and in fact is consistent with the null hypothesis of isotropically distributed arrival directions at the 8% level.

Chapter 8

HiRes Autocorrelation Scan and Angular Correlation Estimates

With an angular resolution on the order of 0.5°, the HiRes experiment provides a unique opportunity to search for anisotropy of cosmic ray arrival directions on the smallest angular scales yet. In this chapter, we apply the autocorrelation scan described in Ch. 7 to the HiRes stereo data set of events observed between 1999 December and 2004 January (described in detail in Ch. 4.3). The significance of the result in light of the clustering signal reported by AGASA is discussed.

8.1 Autocorrelation Scan: Method

The autocorrelation scan is performed over the total set of $N = 271$ events above 10^{19} eV and over angular separations θ from 0° to 5° in increments $\Delta\theta$ of 0.1°.

These boundaries on the search parameters were chosen *a priori*, based on the AGASA clustering claims and the nature of the HiRes detector. By starting well below the 4×10^{19} eV energy associated with the AGASA clustering signal, the search should

safely encompass the energy region of interest even in the presence of a systematic energy shift of 30% between the two experiments, as suggested by De Marco et al. [27]. Starting at this energy does not appreciably dilute the significance of a clustering signal if one is found at higher energy, since the scan involves repeated searching with successively higher energy thresholds. An additional motivation for starting at 10^{19} eV is the fact that the HiRes angular resolution ($\Delta\theta_{68} \approx 0.6^{\circ}$) is much sharper at this energy than AGASA's ($\Delta\theta_{68} \approx 2.8^{\circ}$, see Fig. 6.3).

While the results inevitably depend on the exact choices for these boundary parameters, the dependence is relatively small (see Ch. 7.2.1 for details and examples).

8.2 Background Estimation

The probabilities calculated in the autocorrelation scan are determined by simulating an isotropic flux of cosmic rays. The simulation is performed as follows:

First, using the full detector simulation described in Ch. 5.1, we generate an isotropic distribution of showers with a differential spectral index $\alpha = -3.0$ in energy, and use the resulting distribution of reconstructed Monte Carlo events to determine the detector acceptance in zenith and azimuth. This results in a local acceptance map representing the detector's exposure to the sky at one instant. Then, this local acceptance map is convolved with the observation times of all of the events, which provide a sample of the detector's relative exposure as a function of time. The result is a skymap representing the detector's exposure to different directions in the sky during the data-taking period.

It should be noted that events of all energies are used, in order to sample adequately the exposure and create a smooth skymap. By using event times to sample the

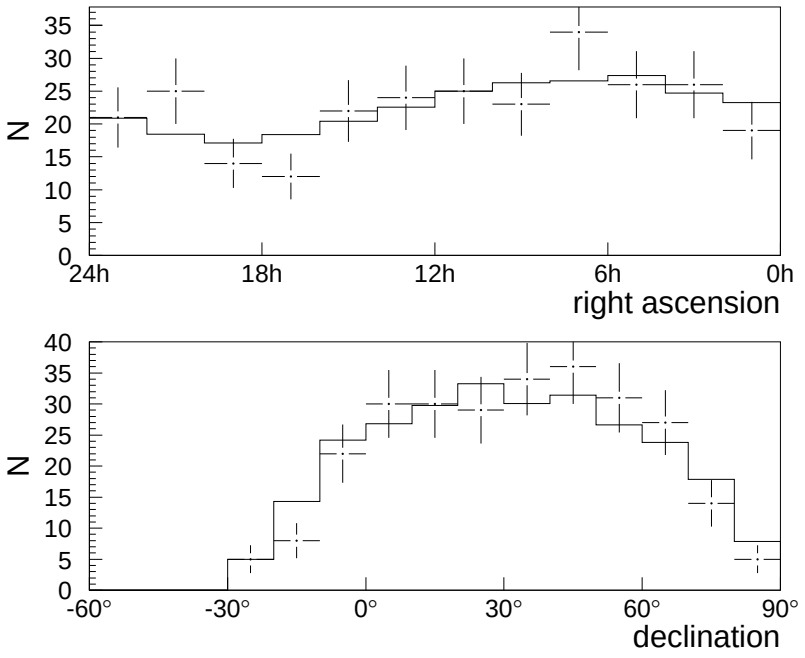

Figure 8.1: Right ascension and declination of events above 10^{19} eV observed from 1999 December through 2004 January. (Data—points with error bars; Simulation—solid line.) For right ascension, $\chi^2/\text{dof} = 0.77$; for declination, $\chi^2/\text{dof} = 0.73$.

detector on-time and exposure, we also are assuming that the distribution of ultrahigh energy cosmic rays in right ascension is isotropic. Fortunately, ground arrays—with their uniform exposure in right ascension and much larger statistics—have confirmed this type of large-scale isotropy well beyond the level required here.

Once the exposure skymap is created, it is used to generate simulated sets of arrival directions. These sets represent an isotropic distribution of cosmic rays as sampled by HiRes. By simulating just the arrival directions directly from the exposure map, one gains tremendously in speed over the full detector simulation. Figure 8.1 compares the distributions of the data and simulated arrival directions in right ascension and declination.

8.3 Sensitivity to Simulated Clustering

To demonstrate the effectiveness of this method and the sensitivity of the HiRes detector, we apply this technique to simulated data with clusters. First, we generate a set of 271 events with the HiRes exposure for isotropic arrival directions described above in Sec. 8.2. We then insert m pairs of events among the N_H highest energy events in the set to simulate clustering above a specific energy threshold.

To create a pair, we pick a point in the sky for the source location and generate two events with arrival directions deviating from the source location according to a Gaussian distribution described below. These artificial cluster positions are chosen at random, but their distribution is forced to reflect the overall exposure of the HiRes detector, so that regions with higher exposure are more likely to contain a cluster. The pair of events is then added to the original isotropic data set, replacing two of the original events in the set. This is repeated until m pairs have been inserted. The set may contain more than m pairs due to chance.

For simplicity, we use a circular Gaussian distribution for the smearing of arrival directions around the source location. The width of the distribution σ_R can be set equal to the angular resolution of the detector, or it can be set to a larger value to simulate additional smearing by magnetic fields. (Note that for the Gaussian distribution $P(\theta) = (\theta/\sigma^2)e^{-\theta^2/2\sigma^2}$, the value $\theta = 1.515\,\sigma$ encloses 68% of the distribution. We therefore define $\sigma_R = 1.515\,\sigma$.)

Table 8.1 shows the results of these simulations using the detector resolution ($\sigma_R = 0.6^\circ$), as well as three times the detector resolution ($\sigma_R = 1.8^\circ$) to simulate additional smearing by magnetic fields. For each choice of N_H, m, and σ_R, we generate 10^4 data sets, and scan them with the procedure described above to find a distribution of values for the significance P_{ch}. The median and 90^{th} percentile values of this

		$\sigma_R = 0.6^\circ$		$\sigma_R = 1.8^\circ$	
N_H [a]	m	median P_{ch}	90% P_{ch}	median P_{ch}	90% P_{ch}
27	2	0.018	0.090	0.13	0.48
	3	2.5×10^{-3}	0.013	0.050	0.25
	4	3.1×10^{-4}	1.5×10^{-3}	0.016	0.11
47	3	0.011	0.067	0.12	0.47
	4	1.9×10^{-3}	0.012	0.059	0.32
	5	3.3×10^{-4}	2.2×10^{-3}	0.029	0.18
89	4	0.016	0.11	0.16	0.59
	6	1.0×10^{-3}	0.012	0.071	0.38
	8	1.1×10^{-4}	7.3×10^{-4}	0.025	0.20

Table 8.1: Results for Simulated Clusters. [a] $N_H = 27$, 47, and 89 events corresponds to simulated clustering above energy thresholds 40 EeV, 28 EeV, and 20 EeV, respectively.

distribution are indicated in Table 8.1.

The table shows, for example, that for a clustering signal on the $\sigma_R = 0.6^\circ$ scale, even three pairs among the 47 highest energy events would typically result in $P_{ch} = 1.1\%$. The table also shows that three such pairs would result in $P_{ch} < 6.7\%$ for 90% of the simulated sets. Thus, an actual value of $P_{ch} > 6.7\%$ could be used to exclude the possibility that sources contributed three such pairs at more than the 90% confidence level.

These results demonstrate the sensitivity to clustering on small angular scales.

8.4 Results

The results of the scan are shown in Figure 8.2. The strongest clustering signal ($P_{min} = 1.9\%$) is observed using the energy threshold $E_c = 1.69 \times 10^{19}$ eV where we observe $n_p = 10$ pairs separated by less than $\theta_c = 2.2^\circ$ within a set of $N_c = 120$ events. By counting the fraction of identical scans over simulated isotropic sets which have

a lower P_{min}, we determine that the statistical significance of this result corresponds to $P_{ch} = 52\%$.

The HiRes stereo data above 10^{19} eV is therefore consistent with the null hypothesis of isotropic arrival directions.

8.5 Comparison with the Earlier AGASA Claim

Comparison with the AGASA clustering result is not straightforward. The HiRes stereo event sample above 4×10^{19} eV is still smaller than AGASA's, though how much smaller depends critically on the level of agreement in absolute energy scale for the two experiments. The possibility of a systematic energy shift of 30% would imply that above the rescaled energy threshold, $(0.7) \times 4 \times 10^{19}\,\mathrm{eV} = 2.8 \times 10^{19}\,\mathrm{eV}$, HiRes has seen 47, rather than 27, events. More importantly, there is the question of how many pairs an independent data set might be expected to contain, given the lack of an obvious source model and the widely varying estimates of the strength of the AGASA clustering. Without assuming a model and source strength, there is no natural way to translate the AGASA observation of five doublets and one triplet separated by less than 2.5° into a meaningful prediction for HiRes.

However, what can be tested using a statistically independent data set is the claim that significant small-scale clustering is a general feature of ultrahigh energy cosmic ray arrival directions. The HiRes stereo data set does not support such a claim. No statistically significant evidence for clustering is observed on any angular scale up to 5° at any energy threshold above 10^{19} eV.

Comparing the observed value of P_{ch} with the values obtained from simulations in Sec. 8.3 (shown in Table 8.1), it can be seen that if the current HiRes data above

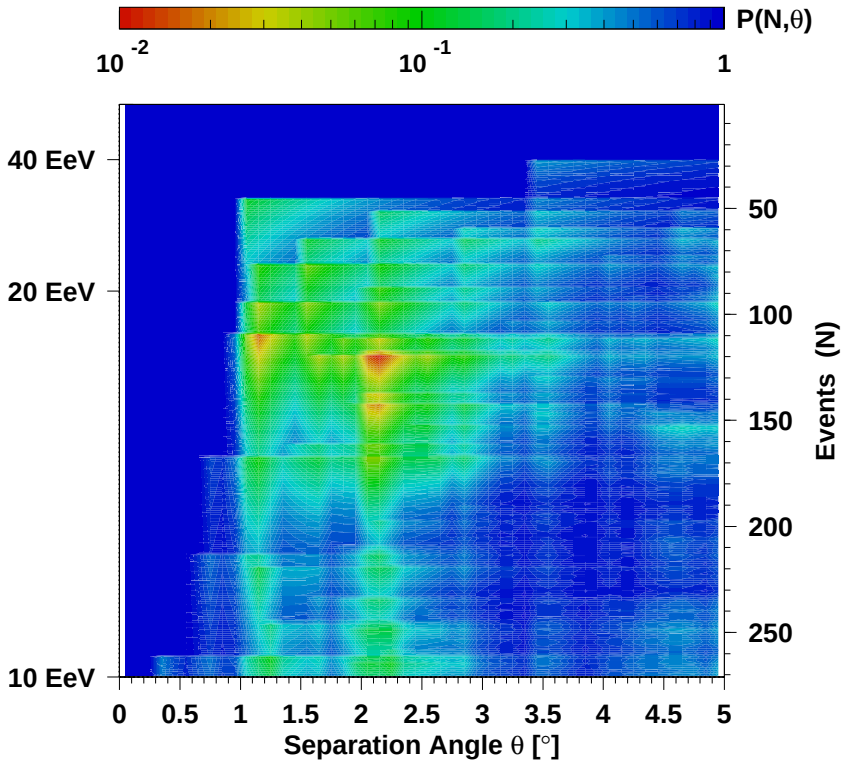

Figure 8.2: Autocorrelation scan of the HiRes data set above 10^{19} eV. $P(N, \theta)$ is the probability of obtaining the same or greater number of pairs as is actually observed in the data using a maximum separation angle θ and searching among the N highest-energy events. These probabilities do not include the statistical penalty due to scanning.

4×10^{19} eV contained two or more pairs of events contributed by compact sources at the angular resolution limit of the detector, then the typical value of P_{ch} would be 0.018 or less, and more than 90% of the time the value of P_{ch} would be much smaller than the observed value of 0.52.

8.6 Combined HiRes-AGASA Angular Correlation Estimate

A related question is whether the combination of the HiRes and AGASA data sets shows a significant clustering signal. Combining the two sets for an autocorrelation scan is unfortunately complicated. The two experiments have different exposures, so determining the probabilities for the number of pairs in a set of N events depends on how many of the events come from AGASA and how many from HiRes. As one performs the scan and lowers the energy threshold, new events get added to the sample from each of the experiments at different rates, and issues of systematic energy shifts between the two experiments become unavoidable.

An alternative approach is simply to add the 27 HiRes events above $E_c = 4 \times 10^{19}$ eV to the 57 AGASA events above the same energy threshold. The angular correlation estimate $w(< \theta)$ described in Ch. 6.2 can then be applied to the combined set, and the result for $\theta_c = 2.5^\circ$ examined. It must be kept in mind that this result will still be biased, because E_c and θ_c were not *a priori* for the AGASA data.

Recall, as seen in Fig. 6.2, that seven pairs of events ($n_p = 7$) are found with angular separation less than 2.5° in the AGASA data set; the expected number of pairs is $\langle n_{\rm MC} \rangle = 1.5$. The corresponding angular correlation estimate is $w(< \theta) = n_p / \langle n_{\rm MC} \rangle - 1 = 3.7$. The fraction of simulated sets with seven or more pairs is 0.1%.

By itself, the set of 27 HiRes events above 4×10^{19} eV has no pairs separated by less than 5°, yielding $w = -1$, but with large uncertainty. However, the combination of the HiRes and AGASA sets substantially increases the statistical power of the two-point correlation function, because for a sample of N events the number of possible pairs is proportional to N^2.

To determine the expected number of pairs $\langle n_{\rm MC} \rangle$, we simulate combined sets with 57 events generated using the AGASA acceptance and 27 generated using the HiRes acceptance. Although these acceptances are not identical, they have a large overlap; the resulting value of $\langle n_{\rm MC} \rangle$ is in fact roughly comparable whether one uses all AGASA events, all HiRes events, or a combination.

The two detectors also have different angular resolutions. While this does not affect simulated isotropic sets, it could affect the correlation estimate for a real clustering signal: a clustering signal of higher significance could appear at smaller angular scales because the HiRes angular resolution is several times sharper than that of AGASA. However, since $w(< \theta)$ includes all pairs with separations less than θ, this effect will not lead to a reduced signal at larger angles. Therefore rather than attempt to estimate the optimal scale for a clustering signal, for the sake of comparison we will continue to evaluate $w(< 2.5°)$ for the combined data set.

Fig. 8.3 shows the results for the 57 AGASA events alone, and for the 84 AGASA and HiRes events evaluated jointly. The addition of the HiRes data brings one new pair with an AGASA event within 2.5°, yielding $n_p = 8$, $n_{\rm MC} = 3.0$, and $w(< 2.5°) = 1.7$. The fraction of simulated sets with eight or more pairs is $\sim 1\%$, but it must be emphasized that this does not represent a chance probability because it includes the same bias in the AGASA data set noted above.

The fact that the angular correlation weakens— not just at 2.5° but at all angular

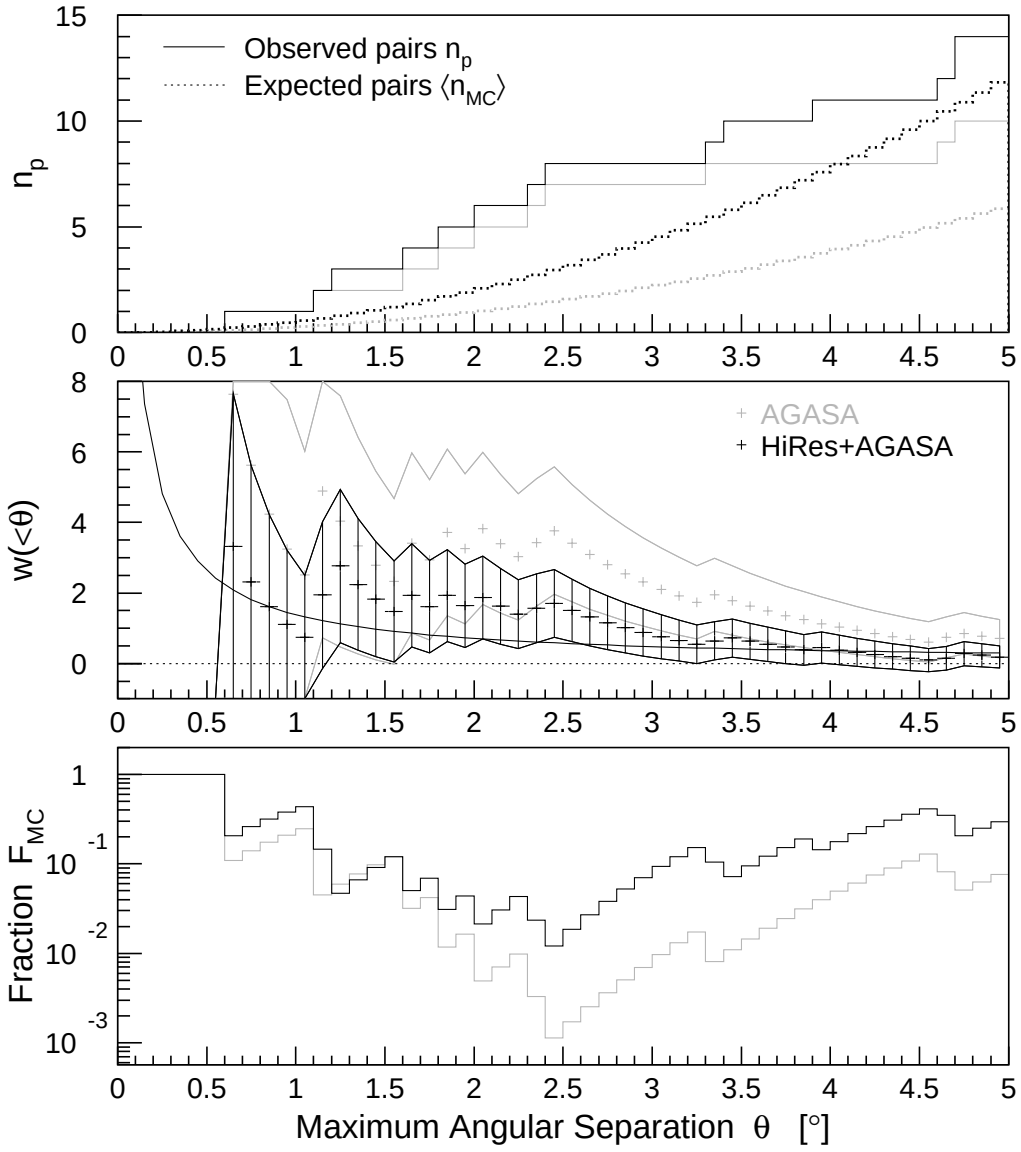

Figure 8.3: Angular two-point correlation estimate w and related quantities for the combined HiRes and AGASA data set (dark), consisting of 84 events with energies above 4×10^{19} eV. (Note the results for the 57 AGASA events alone (light) are reproduced here from Fig. 6.2 for comparison.) *Top:* Number of pairs of events n_p with angular separation less than θ. *Middle:* Angular correlation estimate $w = n_p/\langle n_{\rm MC} \rangle - 1$, with Poisson errors shown. The mean for isotropic distributions is 0 (dotted line); the solid curve shows the Poisson error for isotropic distributions (for the combined data set). *Bottom:* Fraction of simulated sets with the same or greater number of pairs than the n_p observed in the data.

scales—nevertheless offers further evidence that significant clustering is not a general feature of cosmic ray arrival directions. By way of comparison, an observation in the combined data set of $w(< 2.5^\circ) = 3.7$, as seen originally in the AGASA data alone, would have meant the observation of 14 pairs, corresponding to a 10^{-5} deviation from isotropy.

Chapter 9

Maximum Likelihood Point-Source Search

As discussed in the previous chapter, combining the HiRes and AGASA data sets significantly increases the statistical power of the data. However, binning the data by angular separation—as was done with the angular correlation estimate—is a less than optimal solution, because the two experiments have substantially different angular resolution. In choosing a 2.5° bin size, for example, one loses the small angular scale information of the much better resolved HiRes events, and potentially loses a significant signal.

This chapter outlines an alternative method to search for point sources of ultrahigh energy cosmic rays using a maximum likelihood ratio test that does not require a fixed angular bin size. Instead, it uses the probability density function for each individual event. It is thus ideally suited to combining data with different angular errors. The method is applied to the combined set of HiRes stereo and AGASA events above 4.0×10^{19} eV.

9.1 Maximum Likelihood Ratio Test

The hypothesis we wish to test is that several cosmic ray events come from a common point source in the sky. Specifically, we evaluate the likelihood that a certain point in the sky contributes n_s "source" events to an otherwise isotropic distribution of cosmic ray arrival directions. The likelihood of this hypothesis is compared to the null hypothesis that the distribution is purely isotropic ($n_s = 0$), and this likelihood ratio is maximized using n_s as a free parameter. By calculating the likelihood ratio for a dense grid of points covering the sky, we essentially search the entire sky for the most likely position of a source of n_s events. The statistical significance can be estimated by applying the same method to a large set of random isotropic data sets and evaluating what fraction of them have a likelihood ratio which is equal to or larger than the ratio observed in the real data described above.

The maximum likelihood method used here is outlined in Kinnison et al. [64]; a useful and clarifying application to the problem of finding a deficit from the position of the moon in a cosmic ray skymap may be found in Wascko [65]. For a general description of likelihood methods, see for example Meyer [66] and references therein.

9.2 Method

Let a sample of N events consist of n_s source events and $N - n_s$ random background events. Each event is completely described by a set of parameters, which can for example be kinematic variables or, in our case, equatorial coordinates. The parameters for the ith event are combined in a vector $\mathbf{x}_i$, and the probabilities that an event is from source or background depend only on the components of $\mathbf{x}_i$.

Each event is either a source event (with probability Q), or a random background

event (with probability R). Both Q and R are of course functions of the parameters $\mathbf{x}_i$. Assuming that Q and R are properly normalized so that the integral over all possible values for the parameters $\mathbf{x}_i$ is 1, the probability for each event is $n_s \cdot Q + (N - n_s) \cdot R$.

We now define the "partial" probability for event i described by $\mathbf{x}_i$ to occur as

$$P(\mathbf{x}_i) = \frac{n_s}{N} Q(\mathbf{x}_i) + \frac{N - n_s}{N} R(\mathbf{x}_i) \quad , \tag{9.1}$$

and loop over all events to calculate the likelihood

$$\mathcal{L}(n_s) = \prod_{i=1}^{N} P(\mathbf{x}_i) \tag{9.2}$$

for the given set of N events. It is convenient to use the logarithm to avoid small numbers,

$$\ln \mathcal{L}(n_s) = \sum_{i=1}^{N} \ln P(\mathbf{x}_i) \quad . \tag{9.3}$$

The number of source events n_s is allowed to vary and the "best estimate" for n_s is the value that maximizes $\ln \mathcal{L}$. The absolute value of the likelihood function depends on the number of events and is not in itself a very useful quantity.

To simplify the calculations, it is preferable to calculate the likelihood ratio

$$\mathcal{R}(n_s, \mathbf{x}_i) = \frac{\mathcal{L}(n_s, \mathbf{x}_i)}{\mathcal{L}(0, \mathbf{x}_i)} = \prod_{i=1}^{N} \left\{ \frac{n_s}{N} \left(\frac{Q}{R} - 1 \right) + 1 \right\} \quad . \tag{9.4}$$

where $\mathcal{L}(0, \mathbf{x}_i)$ is the likelihood function of the *null hypothesis* ($n_s = 0$). Maximizing $\ln \mathcal{R}$ is equivalent to maximizing $\ln \mathcal{L}$. For a sufficiently large number of events, the parameter $2 \ln \mathcal{R}$ for a likelihood function with q parameters is expected to follow a χ^2-distribution with q degrees of freedom.

In the specific application of this method here, the hypothesis is that more than one event in a skymap originates from a common source. To find the location of such a source, we calculate the likelihood ratio for a dense grid of points on the sky covering the full range of equatorial coordinates (right ascension α and declination δ) accessible to AGASA and HiRes. The source position is essentially treated as a free parameter, as is the number of source events n_s. Searching for the parameters α, δ, and n_s which maximize the likelihood ratio will therefore give us the best estimate for the position of the source and the number of events it contributes. If n_s is close to 1, we can set an upper limit on the number of source events.

For each point on the sky, the probability for a cosmic ray event to come from that position depends on its actual arrival direction in equatorial coordinates. Basically, the closer the event is to the point, the more likely it is to have originated from there. In practice, the probability density also depends on the shape of the errors of the event arrival direction, and so is different for each event.

To review, the vector $\mathbf{x}_i$ contains two components, α_i and δ_i. $Q(\mathbf{x}_i)$ is the normalized probability density for source events to appear at (α_i, δ_i), and $R(\mathbf{x}_i)$ is the probability density for a random background event to appear at this position on the sky. Note that $R(\mathbf{x}_i)$ is different for HiRes and AGASA events due to the different apertures of the experiments. Given Q and R for each event, we can maximize the likelihood ratio for each position on the sky.

Technically, the procedure works as follows. The sky – in this case the region with right ascension $0^\circ < \alpha < 360^\circ$ and declination $-10^\circ < \delta < 80^\circ$ – is divided into bins of size $0.1^\circ/\cos\delta$ and 0.1° in α and δ, respectively. For each bin, we find $\ln\mathcal{R}$ by looping over all events, calculating $P(\mathbf{x}_i)$ with the appropriate normalized probability distributions Q for signal and R for random background. Local maxima

in $\ln \mathcal{R}$ will occur whenever we are close to one or more events. At these points, we find the number of "source events" n_s that maximizes $\ln \mathcal{R}$. Note that while the method is binned, the bin size is small enough so that errors introduced by binning are insignificant.

The likelihood method is ideal in all cases where events have different measurement errors and can also be applied in searches for emission from *a priori* selected source locations.

9.3 Application to HiRes and AGASA Data Sets

The maximum likelihood ratio test is applied to the combined data set of HiRes stereo and AGASA events above 4×10^{19} eV. The HiRes data set and quality cuts are described in Ch. 4.3, and the background estimation is described in Ch. 8.2. The AGASA data set and background estimation are described in Ch. 7.2. The analysis is restricted to the AGASA field of view with $-10^\circ < \delta < 80^\circ$. Within this declination range are 57 AGASA events and 26 HiRes events above the 4×10^{19} eV energy threshold.

9.3.1 Background Probability Function R

The background expectations for HiRes and AGASA are different. The HiRes exposure has a dependence on right ascension. Fig. 9.1 shows the normalized HiRes and AGASA background map in equatorial coordinates.

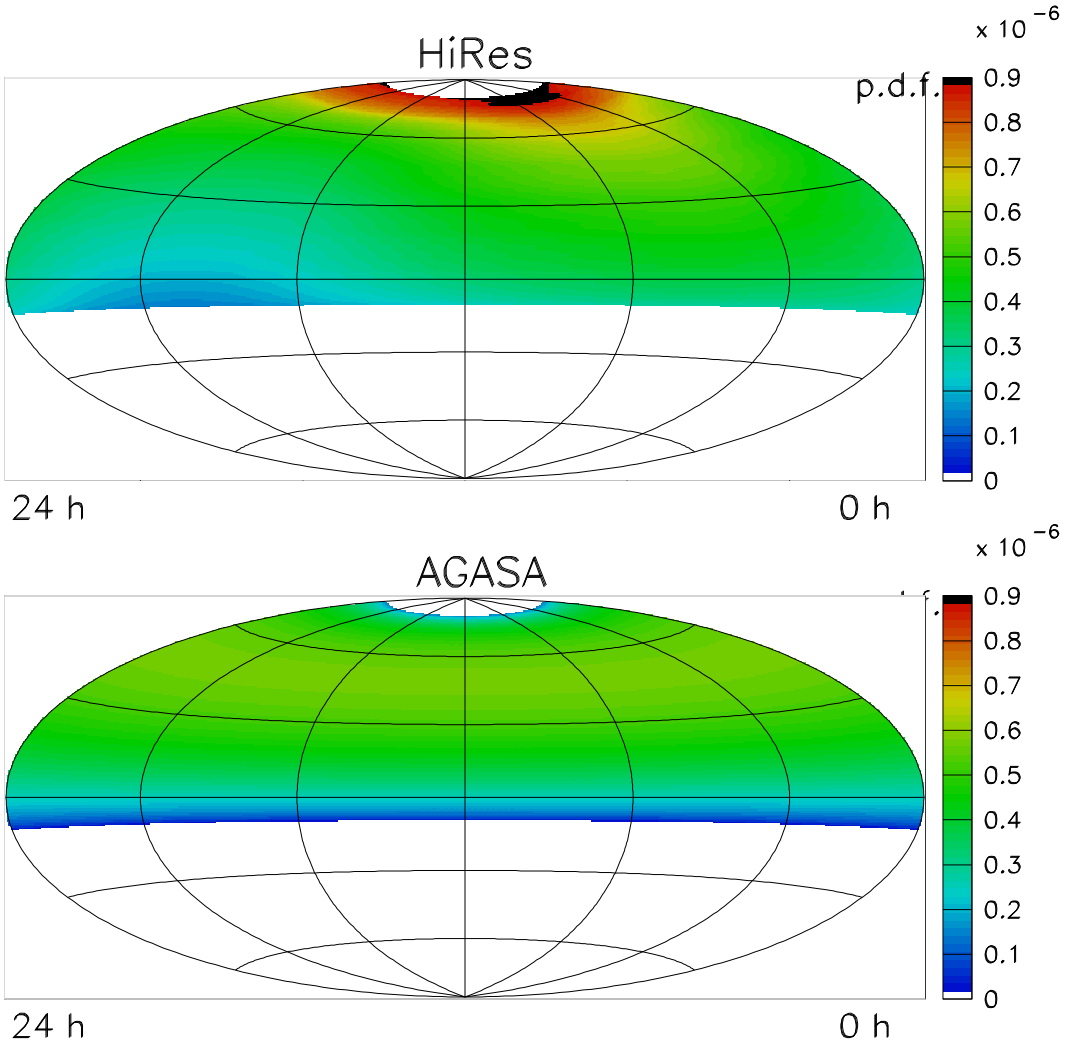

Figure 9.1: HiRes and AGASA Background Maps: Normalized probability densities for random background in equatorial coordinates. The sky is binned in $0.1^\circ/\cos\delta$ and 0.1° bins in right ascension α and declination δ, respectively.

9.3.2 Source Probability Function Q

For the signal probability density function (Q) of the HiRes events, we use a Gaussian function whose width is chosen such that 68 % of the probability density function falls within $\Delta\theta_{68} = 0.6^\circ$ (the HiRes angular resolution, see Ch. 5.1.1). This corresponds to a Gaussian sigma value $\sigma = 0.4^\circ$ in the two-dimensional probability density

$$Q = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{(\Delta\theta)^2}{2\sigma^2}\right) . \tag{9.5}$$

Since the dependence on energy is weak, we use the same value for every HiRes stereo event.

For AGASA, we use the sum of two Gaussian functions chosen such that the 68 % and 90 % opening angle given in Takeda et al. [55] is correctly reproduced (also shown in Fig. 6.3). The width of the Gaussian as a function of energy is taken to have the same energy dependence used in Abbasi et al. [67], given by

$$\sigma_1 = 6.52^\circ - 2.16^\circ \log(E_{\mathrm{EeV}}) \quad \text{and} \tag{9.6}$$

$$\sigma_2 = 3.25^\circ - 1.22^\circ \log(E_{\mathrm{EeV}}) \quad , \tag{9.7}$$

while the total probability density is given by

$$Q = A \left\{ \frac{1}{3} \exp\left(-\frac{(\Delta\theta)^2}{2\sigma_1^2}\right) + \frac{2}{3} \exp\left(-\frac{(\Delta\theta)^2}{2\sigma_2^2}\right) \right\} \quad , \tag{9.8}$$

where A is a normalization factor.

More detailed information on the error shape of individual AGASA events could easily be implemented if it becomes available.

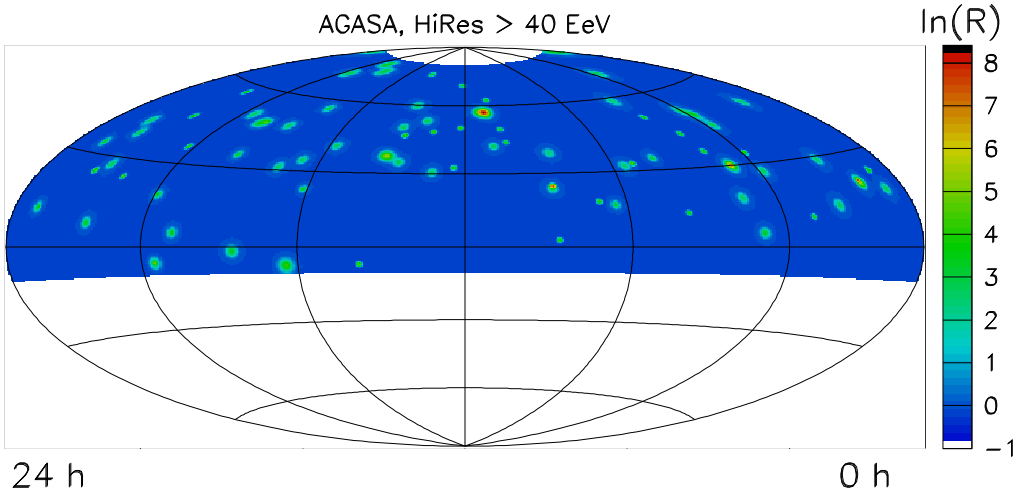

Figure 9.2: Logarithm of the likelihood ratio $\mathcal{R}$ as a function of right ascension and declination for the combined set of AGASA and HiRes events above 4.0×10^{19} eV. Local maxima occur wherever events or clusters of events are located on the sky. The global maximum, *i.e.* the most likely position of a "point source" is at right ascension $\alpha = 169.3°$ and declination $\delta = 57.0°$.

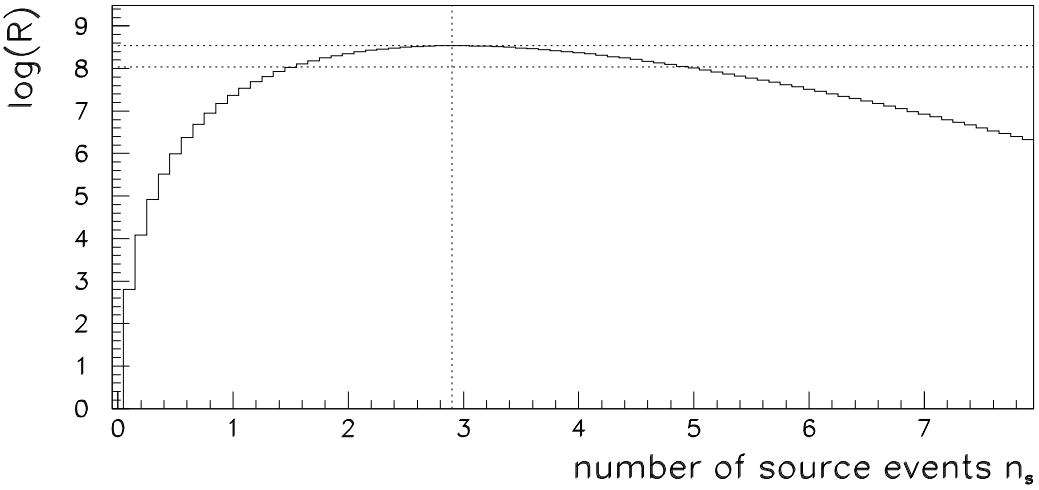

Figure 9.3: Likelihood ratio $\ln \mathcal{R}$ as a function of the number of source events for the position of the maximum.

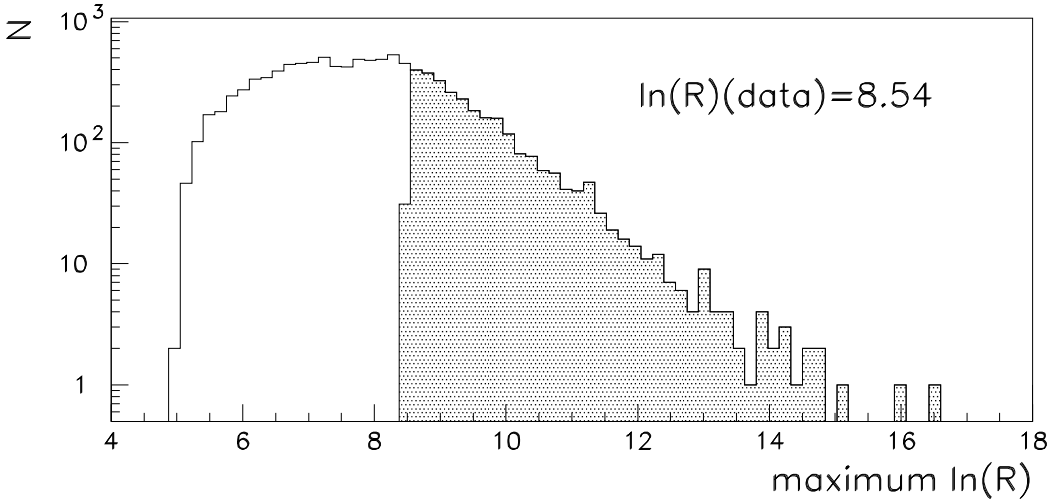

Figure 9.4: Likelihood ratio $\ln \mathcal{R}$ for 10^4 simulated random data sets with the same number of AGASA/HiRes events as the actual data set. The hatched area marks random sets whose $\ln \mathcal{R}$ exceeds the value for the real data set.

9.4 Results

Fig. 9.2 shows the value for $\ln \mathcal{R}$ as a function of right ascension and declination. At each α and δ, the likelihood ratio is shown for the number of source events n_s which maximizes $\ln \mathcal{R}$. One can clearly recognize where each HiRes and AGASA event is located, and one can also recognize the position of multiplets. The point with the largest $\ln \mathcal{R}$ is at right ascension $\alpha = 169.3°$ and declination $\delta = 57.0°$. The corresponding event cluster is a triplet of AGASA events with coordinates (α, δ) and energies E of (1) $(168.5°, 57.6°)$ at $E = 77.6\,\mathrm{EeV}$, (2) $(172.3°, 57.1°)$ at $E = 55.0\,\mathrm{EeV}$, and (3) $(168.3°, 56.0°)$ at $E = 53.5\,\mathrm{EeV}$. This triplet has been described in Takeda et al. [55] and is listed as cluster C2 in Hayashida et al. [56]. The maximum likelihood ratio at this position is $\ln \mathcal{R} = 8.54$ for $n_s = 2.9$. While $\ln \mathcal{R}$ has its global maximum at the location of this cluster, three other doublets, C1 and C4 in Hayashida et al. [56] and a HiRes/AGASA pair at right ascension $\alpha = 146.2°$ and declination $\delta = 24.7°$ have values for $\ln \mathcal{R}$ which are only marginally smaller than 8.54.

The behavior of $\ln \mathcal{R}$ as a function of the fit parameters (number of source events and position of the source) in the vicinity of the maximum gives error estimates for the fit parameters. The $1\,\sigma$ error is given by the interval over which $\ln \mathcal{R}$ drops by 0.5.

Fig. 9.3 shows $\ln \mathcal{R}$ as a function of the number of source events at the position of the maximum. The best estimate of the number of source events is $n_s = 2.9^{+2.0}_{-1.4}$. Similarly we find $\alpha = 169.3° \pm 0.6°$ and $\delta = 57.0° \pm 0.4°$ as the best estimate for the position of the maximum.

The statistical significance of the appearance of a "source" with a given $\ln \mathcal{R}$ in the combined AGASA/HiRes data set can be evaluated using simulated random data sets. The full likelihood analysis is performed for 10^4 random data sets with the same

number of AGASA/HiRes events and the same underlying exposure as the original data set, but isotropic arrival directions. The chance probability for the "source" to appear is then given by the fraction of random data sets which have at least one location with a maximum $\ln \mathcal{R}$ equal or larger than $\ln \mathcal{R} = 8.54$, the value of the maximum in the real data.

Fig. 9.4 shows the distribution of $\ln \mathcal{R}$ for these random data sets. Out of 10^4 simulated data sets, 2793 have a "hot spot" with a likelihood ratio exceeding that of the real data set. The chance probability of the source hypothesis is therefore of the order of 28 %. Consequently, there is no statistically significant evidence for clustering consistent with a point source in the combined data set.

Note that this is *not* simply the chance probability for a triplet, but rather the chance probability for a set of 27 HiRes events and 57 AGASA events to contain a "hot spot" with as high a probability to be a "source" as the triplet. Many of the simulated likelihood ratios larger than 8.54 in Fig. 9.4 are indeed caused by doublets. It is a strength of the likelihood method that it checks for *any* occurrence of strong evidence for a "source" in the data set, regardless of how many events form the underlying multiplets.

9.5 Test of the Method with Simulated Data Sets

The maximum likelihood method is tested by applying it to simulated data sets with sources. These data sets have m events from a common source added to an otherwise isotropic arrival direction distribution. To create such a source, we pick a point in the sky for the source location and generate m events with arrival directions deviating from the source location according to the probability density function of the

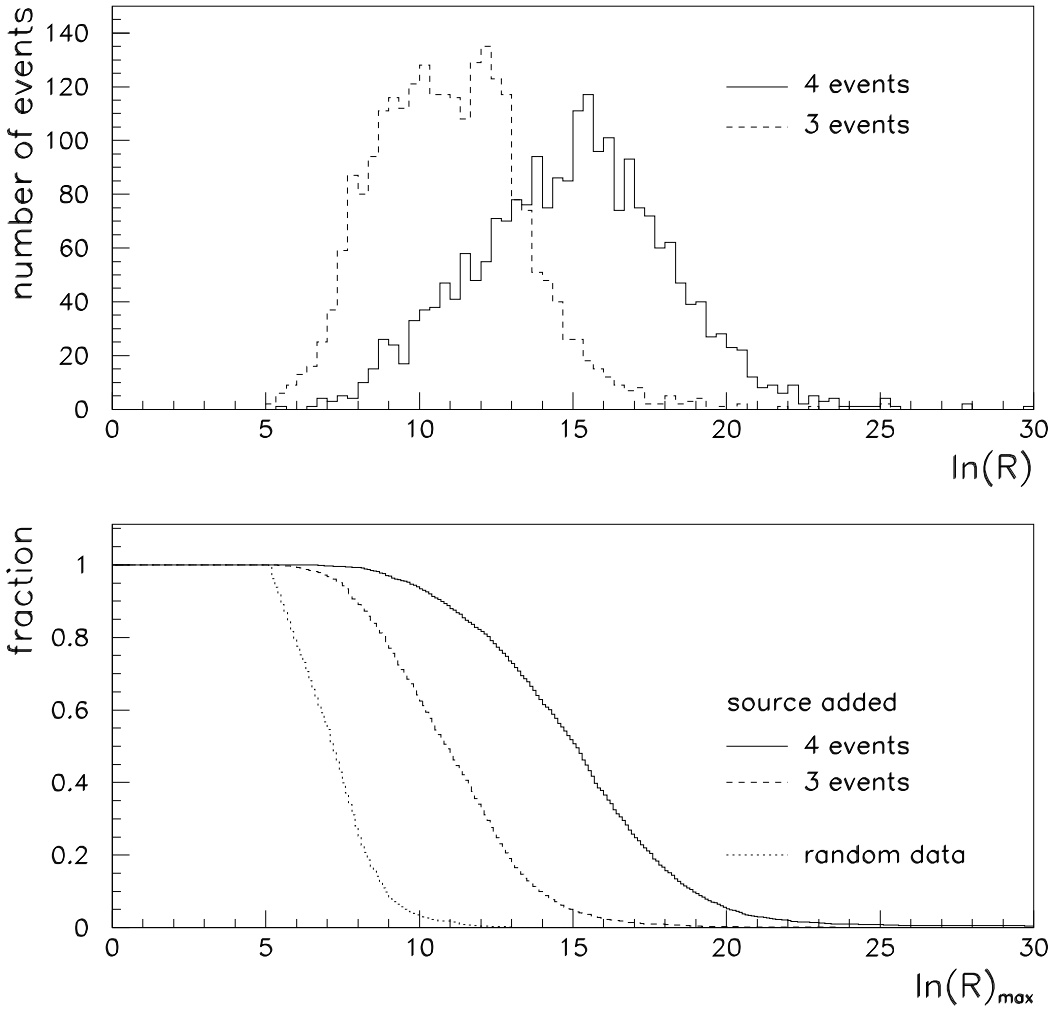

Figure 9.5: *Top:* Likelihood ratio $\ln \mathcal{R}$ for 10^5 simulated random data sets where a 4-event source (solid line) or a 3-event source (dashed line) has been added. *Bottom:* Fraction of simulated data sets with $\ln \mathcal{R} > \ln \mathcal{R}_{max}$ as a function of $\ln \mathcal{R}_{max}$ for simulated random isotropic data sets with no source added (dotted line), a 4-event source added (solid line), and a 3-event source added (dashed line).

individual event as described in Section 2. The source location is chosen at random, but the distribution of locations is forced to reflect the overall exposure of the detector (Fig. 9.1), so that regions with higher exposure are more likely to contain a source. The m source events replace events in the data set, so the total number of events is always 83. There might be additional events close to the source location due to chance.

The full likelihood analysis is applied to these random data sets. Fig. 9.5 (top) shows the likelihood ratio $\ln \mathcal{R}$ for 10^5 random data sets where a source with $m = 4$ (solid line) or $m = 3$ (dashed line) has been added. Fig. 9.5 (bottom) shows the fraction of simulated data sets with $\ln \mathcal{R} > \ln \mathcal{R}_{max}$ as a function of $\ln \mathcal{R}_{max}$ for random isotropic data sets with no source added ($m = 0$, dotted line), a 4-event source added (solid line), and a 3-event source added (dashed line). While there is substantial overlap between the distributions, the plots also show that point sources which add three or more events to an isotropic map are recognized with high efficiency. The medians of the three $\ln \mathcal{R}$-distributions ($m = 0, 3, 4$) are 7.7, 11.6 and 15.2, respectively.

The HiRes and AGASA events in the sample have rather different angular errors, and we expect this to be reflected in the likelihood ratio of the simulated data sets. The $\ln \mathcal{R}$-distribution for clusters that are dominated by HiRes events should have a larger median $\ln \mathcal{R}$ than the data sets dominated by AGASA events. Fig. 9.6 shows the median of the $\ln \mathcal{R}$-distribution as a function of the number of HiRes events that contribute to the cluster. Both for data sets with $m = 3$ and $m = 4$, the median increases if more HiRes events are part of the cluster.

Artificial sources can also be used to test the accuracy with which the position of the point source is reconstructed. Fig. 9.7 shows the difference between the fitted

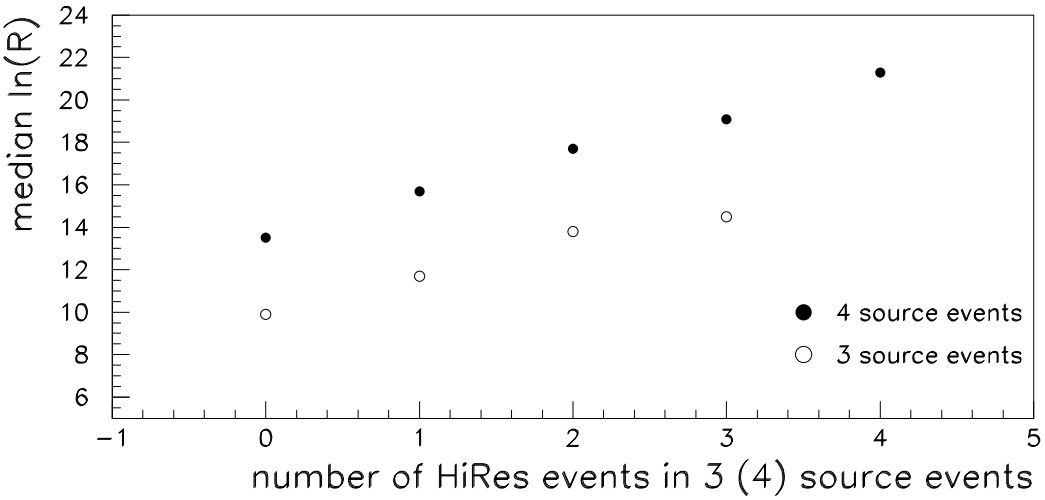

Figure 9.6: Median of the $\ln \mathcal{R}$-distribution for simulated data sets with $m = 3$ and $m = 4$ source events as a function of the number of HiRes events contributing to the source cluster.

and the true right ascension α and declination angle δ for $m = 3$ and $m = 4$. The distributions are fit to a Gaussian function, and the width is in good agreement with the error on the likelihood fit for an individual event obtained from the $(\ln \mathcal{R} - 0.5)-$ method.

9.6 Summary

The most likely position of a point source in the combined data set is a triplet of AGASA events. However, the chance probability for finding an equally good or better "hot spot" in isotropic random data sets is 28 %. Consequently, there is no evidence for a statistically significant hot spot in the combined data set.

The chance probability of the triplet using AGASA data alone has been estimated

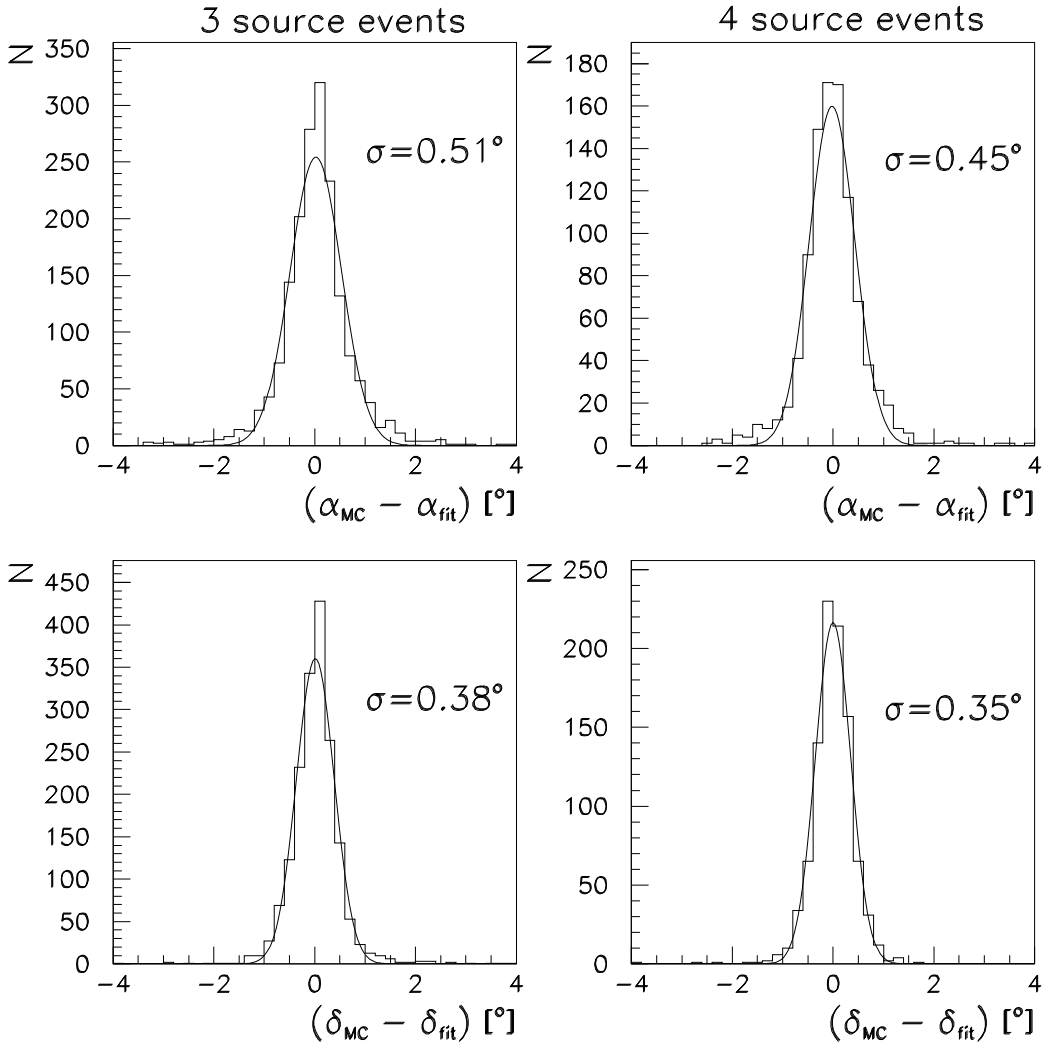

Figure 9.7: Difference between the fitted and true right ascension α and declination δ for 10^5 data sets with an artificial point source added to an isotropic data set for $m = 3$ (left) and $m = 4$ (right) source events. The width of a Gaussian fitted to the distribution is indicated.

in Takeda et al. [57] as being of order 1 %. This estimate is based on a fixed bin size of 2.5°. To test what chance probability an unbinned analysis gives, we repeat the likelihood analysis for the data set comprising only the 57 AGASA events above 4.0×10^{19} eV. The largest likelihood ratio ($\ln \mathcal{R} = 9.66$) appears again near the events forming the triplet, with $\alpha = 169.3^\circ \pm 0.6^\circ$ and $\delta = 57.0^\circ \pm 0.4^\circ$ for $n_s = 2.9^{+2.2}_{-1.7}$.

As before, we evaluate the chance probability for the appearance of a source with $\ln \mathcal{R} = 9.66$ or higher in this data set by analyzing a large number of simulated isotropic data sets. 452 out of 10^4 random data sets have at least one location with a maximum likelihood ratio in excess of 9.66, so the chance probability is 4.5 %. This chance probability is higher than the one given in Takeda et al. [57]. It is important to emphasize that this result is immune to the bias from the fact that the AGASA bin size of 2.5° was chosen *a posteriori* in Hayashida et al. [54] as the bin size that maximizes the clustering signal in the AGASA data set.

Chapter 10

Correlations with BL Lacertae Objects

Cosmic ray arrival directions have long been subjected to extensive searches for correlations with the positions of objects from known astrophysical source classes. Recently, significant correlations between subsets of BL Lacertae objects and cosmic ray arrival directions recorded by the AGASA and Yakutsk experiments have been claimed [68–70].

Searches for correlations with BL Lac objects are well-motivated. BL Lacs are a subclass of blazars, which are active galaxies in which the jet axis happens to point almost directly along the line of sight. The EGRET instrument on board the Compton Gamma Ray Observatory (CGRO) has firmly established blazars as sources of high energy γ-rays above 100 MeV [71], and several BL Lac objects have been observed at TeV energies with ground-based air Cerenkov telescopes [see 72, for a summary]. High energy γ-rays could be by-products of electromagnetic cascades from energy losses associated with the acceleration of ultra–high-energy cosmic rays and their

propagation in intergalactic space [73, 74].

However, the recent claims of correlation between cosmic ray arrival directions and BL Lac objects are controversial. A problematic aspect of the claims is the procedure used to establish correlations and evaluate their statistical significance. Several authors [68–70] explicitly tuned their selection criteria to assemble catalogs that show a maximum correlation with arrival directions of cosmic rays above some energy. An unbiased chance probability for these correlations can then only be arrived at if the claim is tested on a statistically independent data set. Since the available data set is small, this rigorous procedure is often abandoned, and instead an attempt is made to estimate a statistical penalty factor to compensate for the number of trials involved in the tuning. "Hidden" trials, unfortunately, make these *a posteriori* estimates highly unreliable, and claims of BL Lac correlations have been criticized on these grounds [75, 76]. Additionally, in some cases it has been shown that statistically independent data sets do not confirm the correlations [77].

In this chapter, the recent claims of BL Lac - cosmic ray correlations are reviewed, including a very recent claim that used the published HiRes event arrival directions above 10^{19} eV. The maximum likelihood method described in Ch. 9 is modified for the case of multiple source locations. This method is then applied to a search for correlations between BL Lac objects and HiRes stereo events, and the claim based on HiRes events is verified and studied in detail. However, due to the *a posteriori* nature of the search, the significance level cannot be reliably estimated and the correlation must be tested independently before any claim can be made. The precise hypotheses that will be tested in the future with statistically independent data are identified.

10.1 Maximum Likelihood Method

10.1.1 Description

In Ch. 9, an *unbinned* maximum likelihood method for point-source searches has been described. This approach uses the probability density function for each individual event rather than requiring a fixed bin size. Two important advantages of this method are the ability to accommodate events with different errors, and to give weighted sensitivity to angular separations—avoiding the loss of information that follows from choosing an angular separation cut-off. With minor modifications, the same maximum likelihood method can also be used to search for correlations with a specified list of potential sources.

Recall that the premise involved in the maximum likelihood analysis is that the data sample of N events consists of n_s source events which came from some source position(s) in the sky, and $N - n_s$ background events. A background event arrives according to the probability distribution given by the detector exposure to the sky, $R(\mathbf{x})$, where $\mathbf{x}$ are equatorial coordinates. The true arrival direction of a source event is the location of the source $\mathbf{s}$, but the event is observed somewhere near $\mathbf{s}$ according to the probability distribution $Q_i(\mathbf{x}, \mathbf{s})$, where Q_i depends on the angular uncertainty in arrival direction of the ith event.

Because it is not known whether a given event is a source or background event, the probability distribution function (or "partial" probability) for the ith event is a weighted sum of the source and background probability distributions:

$$P_i(\mathbf{x}) = \frac{n_s}{N} Q_i(\mathbf{x}, \mathbf{s}) + \frac{N - n_s}{N} R(\mathbf{x}) \quad . \tag{10.1}$$

10.1.2 Modification for Multiple Sources

Eq. 10.1 describes the distribution of arrival directions under a single-source hypothesis. For a hypothesis with M sources, we must modify Q to include multiple source locations. We will assume for this analysis that the sources have equal luminosity. In this case, we only need to compensate for the varying exposure of the detector to different parts of the sky: the probability for a source event to come from the jth source is proportional to the detector exposure $R(\mathbf{s}_j)$ to the source location $\mathbf{s}_j$. The total source probability distribution Q_i^{tot} for the ith event is then the weighted sum of the individual source probabilities:

$$Q_i^{tot}(\mathbf{x}) = \sum_{j=1}^{M} R(\mathbf{s}_j) Q_i(\mathbf{x}, \mathbf{s}_j) / \sum_{k=1}^{M} R(\mathbf{s}_k) \quad . \tag{10.2}$$

Replacing Q_i in Eq. 10.1 with Q_i^{tot}, we evaluate the partial probability of the ith event at its observed location $\mathbf{x}_i$:

$$P_i(\mathbf{x}_i) = \frac{n_s}{N} Q_i^{tot}(\mathbf{x}_i) + \frac{N - n_s}{N} R(\mathbf{x}_i) \quad . \tag{10.3}$$

The likelihood for the entire set of N events is then given by:

$$\mathcal{L}(n_s) = \prod_{i=1}^{N} P_i(\mathbf{x}_i) \quad . \tag{10.4}$$

The best estimate for the number of events contributed by the sources is determined by finding the value of n_s that maximizes $\mathcal{L}$.

Again, because the value of the likelihood function depends on the number of

events, a more useful quantity than $\mathcal{L}$ is the likelihood ratio $\mathcal{R}$:

$$\mathcal{R}(n_s) = \frac{\mathcal{L}(n_s)}{\mathcal{L}(0)} = \prod_{i=1}^{N} \left\{ \frac{n_s}{N} \left(\frac{Q_i^{tot}(\mathbf{x}_i)}{R(\mathbf{x}_i)} - 1 \right) + 1 \right\} \tag{10.5}$$

where $\mathcal{L}(0)$ is the likelihood function of the *null hypothesis* ($n_s = 0$). In practice, we maximize $\ln \mathcal{R}$, which is equivalent to maximizing $\mathcal{L}$.

10.1.3 Estimating Significance

When the point source or collection of point sources is *a priori*, a useful feature of the maximum likelihood method is that the significance of the resulting $\ln \mathcal{R}$ can be estimated using $\chi^2 = 2 \ln \mathcal{R}$. When n_s is positive, this agrees well with the χ^2 distribution for one degree of freedom. Because n_s corresponds to the *excess* number of events correlating with source positions, a negative best-fit value for n_s will occur whenever there are fewer events near source positions than expected. Negative n_s values are not physically meaningful in the point-source search, but they are useful for evaluating significances. To distinguish an excess in correlations from a deficit, we assign the negative solution $\chi = -\sqrt{2 \ln \mathcal{R}}$ when the best-fit n_s is negative.

We can check the significance estimated from the χ^2 distribution by performing the same likelihood analysis on simulated data sets and ranking them according to their χ values. We will use $\mathcal{F}$ to denote the fraction of simulated, isotropic event sets which yield a value of χ greater than or equal to that of the data.

10.1.4 Implementation

Unlike in Ch. 9 where a common $\sigma = 0.4^\circ$ was used for all HiRes events, in this analysis the individual errors will be used, because events at different energy thresholds are

considered. For the source probability function Q_i we employ a circular Gaussian of width σ_i corresponding to the angular uncertainty of the ith event, as estimated by the stereo event reconstruction (see Ch. 5.1.2). The mean of the angular uncertainty of the HiRes stereo events is slightly larger at lower energies, growing from $\langle\sigma\rangle = 0.44^\circ$ above 10^{18} eV to $\langle\sigma\rangle = 0.60^\circ$ below $10^{17.75}$ eV. This of course is accounted for by the use of individual errors in the maximum likelihood analysis.

The data used in this analysis are the HiRes stereo events observed between December 1999 and January 2004 (see Ch. 4.3). Different energy thresholds are considered, including using no energy threshold, *i.e.* including events down to energies $\sim 10^{17.5}$ eV. Because the geometrical acceptance of the detector is a function of energy, we determine the background probability distribution $R(\mathbf{x})$ in two different ways. For large event samples ($\gtrsim 1000$), we generate a background distribution from a full time-swapping of the data itself: the equatorial coordinates of each event are recalculated using all of the recorded event times, and $R(\mathbf{x})$ is the cumulative map of all of these virtual event locations convolved with a circular Gaussian function for smoothing. For small event samples (e.g. the 271 events above 10^{19} eV) the data set is too sparse to generate a useful time-swapped map. Instead, we rely on a full detector simulation to estimate the local geometrical acceptance, and convolve this with the event times to generate $R(\mathbf{x})$ (see Ch. 8.2 for more details).

10.2 Tests of Previous Correlations Observed with AGASA and Yakutsk Data

We briefly review some past claims of correlations between cosmic ray arrival directions and BL Lacs, and then test these with HiRes data under the same conditions.

All samples of BL Lacs are selected from objects classified as "BL" in Table 2 of the Veron Catalog of Quasars and AGN, 9th or 10th Editions [78, 79].

- *Sample A:* described in Tinyakov and Tkachev [68], contains 22 BL Lacs from the Veron 9th Catalog selected on the basis of optical magnitude ($m < 18$), redshift ($z > 0.1$ or unknown), and 6 cm radio flux ($F_6 > 0.17$ Jy).

- *Sample B:* described in Tinyakov and Tkachev [69], contains 157 BL Lacs from the Veron 10th Catalog with optical magnitude $m < 18$.

- *Sample C:* described in Gorbunov et al. [70], consists of 14 BL Lacs from the Veron 10th catalog that were selected by the authors on the basis of possible association with identified and unidentified γ-ray sources in the Third EGRET Catalog [71].

	Sample (# Objects)	Reference	CR Data and Energy Threshold (EeV)	# Ev.	Binned Analysis Bin Size	Pairs	Prob.	Max. Like. Analysis $\ln \mathcal{R}$	n_s	$\mathcal{F}$
Claim 1:	A (22)	TT01	AGASA>48, Yak.>24	65	2.5°	8	$< 10^{-4}$			
Test:			HiRes > 24	66	2.5°	0	1.00	(0)	(0)	0.75
Claim 2:	B (157)	TT02	AGASA > 40	57	2.5°	12	0.02			
Test:			HiRes > 40	27	2.5°	2	0.78	(0)	(0)	0.26
Claim 3:	C (14)	G02	AGASA>48, Yak.>24	65	2.9°	8	10^{-4}			
Test:			HiRes > 24	66	2.9°	1	0.70	(0)	(0)	0.68

Table 10.1: Previously claimed correlations between BL Lacs and cosmic rays, and tests with independent HiRes data. [TT01] Tinyakov and Tkachev [68]; [TT02] Tinyakov and Tkachev [69]; [G02] Gorbunov et al. [70]. In [TT02] and [G02], the authors also attempt to correct for the deflections of charged primaries by the galactic magnetic field; these results are not considered here.

Table 10.1 shows the correlations originally claimed using these BL Lac samples and cosmic ray data from the AGASA and Yakutsk experiments. The energy thresholds and angular bin sizes vary from analysis to analysis as shown. The results of testing each claim as nearly as possible with an equivalent set of HiRes data are also presented. Both a binned analysis with the originally used bin size and a maximum likelihood analysis using the point spread function of individual HiRes events are performed. In the binned analysis, the number of event-object pairs with angular separation less than the bin size are counted, and the probability for the same or greater number of pairs is evaluated using simulated isotropic event sets. None of the three previous claims of correlations based on other cosmic ray data sets are confirmed by the tests. Each test, in fact, finds a deficit or no excess of HiRes events correlating with BL Lacs, indicated by (0) values for $\ln \mathcal{R}$ and n_s. The fraction $\mathcal{F}$ of simulated sets with stronger correlation than the data is calculated as described above.

In the tests of Claims 1 and 3, the size of the HiRes event sample is comparable to the size of the combined AGASA and Yakutsk event samples. Assuming a Poisson distribution with mean number of event - BL Lac pairs given by the AGASA and Yakutsk results, the observed number of HiRes - BL Lac pairs excludes the claimed correlations at a confidence level greater than 99 % in each case. In the test of Claim 2, the HiRes event sample is smaller than that of AGASA. Here the claimed correlation is excluded at the 90 % confidence level.

10.3 Correlations Observed with HiRes Data

Recently, the published HiRes events above 10^{19} eV were analyzed by Gorbunov et al. [80], and correlations with the BL Lacs of Sample B were claimed at the 10^{-3} level.

The analysis used a fixed bin size of 0.8°, which the authors argued is optimal for a point-source search given the HiRes angular resolution. We verify this analysis by applying the maximum likelihood method to the same data set and source sample, and find $\ln \mathcal{R} = 6.08$ for $n_s = 8.0$; the fraction of simulated isotropic sets with higher $\ln \mathcal{R}$ is $\mathcal{F} = 2 \times 10^{-4}$.

Gorbunov et al. [80] analyzed the entire set of HiRes events above 10^{19} eV because the individual event energies were not published. Therefore this energy threshold was not tuned to maximize correlations with BL Lacs. However, because the original claim [69] was based on AGASA data with energies above 4×10^{19} eV, the correlation in Gorbunov et al. [80] does not confirm a previous claim, but rather represents a new hypothesis. This is demonstrated by the fact that the HiRes data shows no excess correlation with this sample of BL Lacs when the same 4×10^{19} eV energy threshold is used, as indicated in the test of Claim 2 in Table 10.1.

The observed correlation warrants further investigation. We report on extending the analysis to lower energy HiRes data and to the rest of the confirmed BL Lacs in the Veron catalog.

10.3.1 Event Sample: Energy Dependence of Correlations

An important question is whether and how the observed correlation depends on the energy threshold. Figure 10.1 shows the result of the same maximum likelihood analysis above, performed repeatedly using increasing energy thresholds from $10^{18.5}$ eV to 10^{20} eV. The 10^{19} eV threshold corresponding to the published data set is indicated, and it clearly stands out as the threshold that gives a local maximum in the significance of the correlation.

One of the motivations for using an energy threshold in small-scale anisotropy

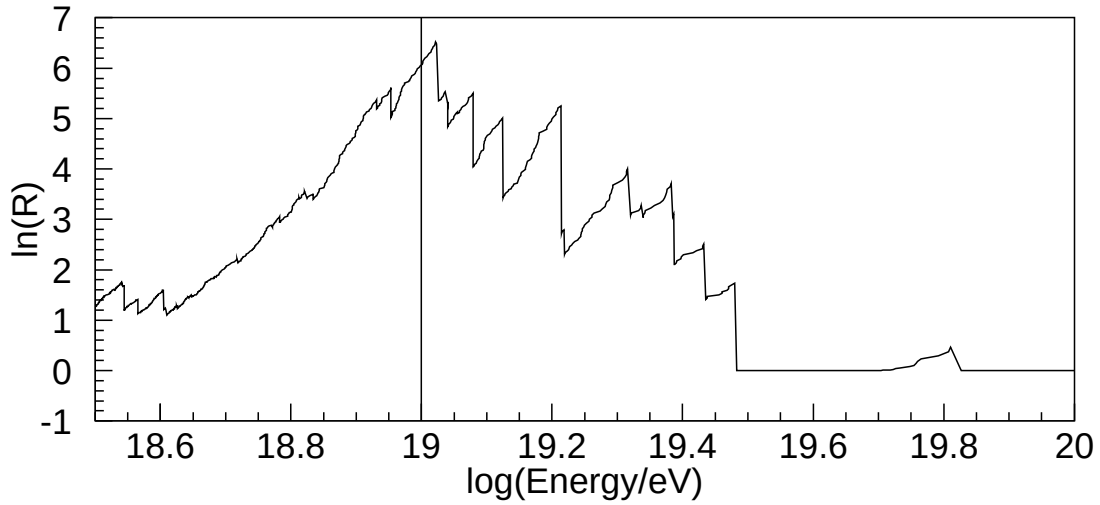

Figure 10.1: $\ln \mathcal{R}$ result as a function of minimum energy threshold of the HiRes data set. The 10^{19} eV energy threshold of the published data is indicated.

searches is that charged cosmic ray primaries are subject to deflection by galactic and extragalactic magnetic fields. The highest energy primaries will be least deflected, and consequently will be the best candidates for correlation studies. However, an *a priori* energy threshold cannot be established, because detailed knowledge of the galactic and extragalactic magnetic fields is lacking.

Figure 10.1 indicates that most of the correlation comes from events with energies between 10^{19} eV and $10^{19.5}$ eV. At these energies, it is generally assumed that the Galactic magnetic field will deflect a proton primary by many degrees; nuclei will be deflected even more. In spite of this, the correlations are consistent with the $\sim 0.5^{\circ}$ scale of the detector angular resolution. This might imply that the correlated primary cosmic rays are neutral, thus removing the motivation for restricting the analysis to events with energies above some arbitrary threshold. A search for correlations with the entire HiRes stereo data set of 4495 events at all energies is justified.

Applying the analysis to the entire data set and Sample B gives $\ln \mathcal{R} = 6.16$ for

Source Sample (# Obj.)	All Energies	$E > 10\,\mathrm{EeV}$
"BL" (157)	2×10^{-4}	2×10^{-4}
"HP" (47)	0.3	6×10^{-3}
"BL"+"HP" (204)	5×10^{-4}	10^{-5}

Table 10.2: HiRes — BL Lac Correlation Results: Fraction $\mathcal{F}$ of simulated HiRes sets with stronger correlation signal. Correlations are with confirmed BL Lacs in Table 2 of the Veron 10th Catalog [79], classified as either "BL" or "HP," with $m < 18$.

$n_s = 31$, with $\mathcal{F} = 2 \times 10^{-4}$. This of course includes the effect of the correlated events above $10^{19}\,\mathrm{eV}$; for the independent sample of 4224 events below $10^{19}\,\mathrm{eV}$, we find $\ln \mathcal{R} = 3.10$ for $n_s = 22$, with $\mathcal{F} = 6 \times 10^{-3}$.

10.3.2 Source Sample

The Sample B of BL Lacs discussed above consists of 157 confirmed BL Lacs in the 10th Veron Catalog with optical magnitude $m < 18$ and which are classified as "BL." The rest of the confirmed BL Lacs are classified as "HP" (high polarization). It is natural to ask about these objects as well. Indeed, of the six blazars which have confirmed detections in γ-rays at TeV energies, half are classified as "HP" and half as "BL" in the Veron catalog.

The same cut on optical magnitude $m < 18$ as in Sample B is applied to the "HP" objects, and arrive at a sample of 47 objects. The result of the maximum likelihood analysis applied to this independent sample of BL Lacs and the HiRes events above $10^{19}\,\mathrm{eV}$ is $\ln \mathcal{R} = 3.13$ for $n_s = 3.0$, with $\mathcal{F} = 6 \times 10^{-3}$. We also perform the same analysis on the events below $10^{19}\,\mathrm{eV}$. No excess is found.

The results for HiRes events of energy above $10^{19}\,\mathrm{eV}$ and all energies with BL Lacs classified as "BL," "HP," or "BL" and "HP" combined are summarized in Table 10.2.

The equivalent analyses have been performed on the same classes of BL Lacs with

$m \geq 18$: no excess correlation is found in any of these cases. It is apparent from these results that the $m < 18$ cut which was identified in Tinyakov and Tkachev [69] as optimal for AGASA also isolates the BL Lac objects which show excess correlations with HiRes events. Under the BL Lac source hypothesis, of course, it is not unreasonable to expect the closer and more luminous objects to contribute more strongly. However, since the Veron catalog is not a uniform sample of BL Lac objects, the interpretation of this cut may involve a more complicated interplay of selection effects from the underlying surveys which make up the catalog.[1]

10.3.3 TeV Blazars

Among the closest and brightest of the "BL" and "HP" BL Lacs are six which are confirmed sources of TeV γ-rays [72]. Five of these, shown in Table 10.3, are high in the northern sky and well within the field of view of HiRes. We perform the maximum likelihood analysis on this set of objects using all of the HiRes data, and find $\ln \mathcal{R} = 4.78$ for $n_s = 5.6$ with $\mathcal{F} = 10^{-3}$. For just the HiRes events above $10^{19}\,\mathrm{eV}$, the result is $\ln \mathcal{R} = 6.15$ for $n_s = 2.0$ with $\mathcal{F} = 2 \times 10^{-4}$. In Table 10.3, we show the results for treating each blazar in turn as a single source hypothesis.

10.4 Summary

Previous claims of correlations between BL Lacs and cosmic rays which were based on AGASA and Yakutsk data have been tested. Using the independent HiRes stereo data

[1] The Veron catalog strives to be "complete" only in the sense of a complete survey of the literature and catalog of all known BL Lacs; it does not represent an unbiased statistical sample of BL Lacs in any way [78, 79]. This does not exclude the possibility of using subsets of the catalog to identify correlations with cosmic rays, but it means that any inferences about the BL Lacs based on such correlations may be highly biased and simply an artifact of the underlying combination of different surveys in the catalog.

TeV Blazars			Correlation Results		
Name	z	V Mag	n_s	$\ln \mathcal{R}$	$\mathcal{F}$
Mrk 421	0.03	12.9	0.3	0.04	0.2
H1426+428	0.13	16.5	(0)[a]	(0)	0.4
Mrk 501	0.03	13.8	3.3	5.27	6×10^{-4}
1ES1959+650	0.05	12.8	2.0	2.87	8×10^{-3}
1ES2344+514	0.04	15.5	(0)[b]	(0)	0.7
Combined Set			5.6	4.78	10^{-3}

Table 10.3: TeV Blazar Correlation Results with HiRes events (all energies). [a] No excess: $n_s < 3.5$ at 90% confidence level. [b] No excess: $n_s < 2.4$ at 90% confidence level.

Source Sample (# Obj.)	All Energies	$E > 10\,\mathrm{EeV}$
"BL" Objects, $m < 18$ (157)	2×10^{-4}	2×10^{-4}
Confirmed BL Lacs, $m < 18$ (204)	5×10^{-4}	10^{-5}
Confirmed TeV Blazars (6)	10^{-3}	2×10^{-4}

Table 10.4: HiRes — BL Lac Correlation Summary: Fraction $\mathcal{F}$ of simulated HiRes sets with stronger correlation signal. All samples are contained within Table 2 of the Veron 10th Catalog. The samples overlap and are *not* independent: "Confirmed BL Lacs" combines "BL" and "HP" classified BL Lacs; TeV Blazars are a subset of the confirmed BL Lacs.

set, these correlation claims are excluded at the 99 % (Claims 1 and 3 in Table 10.1) or 90 % confidence level (Claim 2).

However, the observation by Gorbunov et al. [80] that the set of HiRes stereo events with energies above 10^{19} eV shows an excess of events correlated with confirmed BL Lacs marked as "BL" in the Veron 10th Catalog has been verified. It must be emphasized that the observed correlation does not confirm a previous claim, because it requires a lower energy threshold. It can only be confirmed with new data.

We have explored the extension of the analysis to 1) HiRes events of all energies, and 2) the rest of the confirmed BL Lacs (labeled "HP") in the Veron 10th Catalog. In

each case, correlations at the significance level of $\sim 0.5\%$ are found. While statistically independent from the above result, these are not strictly tests of that claim. However, the combination offers well-defined hypotheses which can be tested with new data.

The results of combining the analysis of low and high energy events and "BL" and "HP" BL Lacs are summarized in Table 10.4. Also shown are the results for HiRes events and the subset of BL Lacs which are confirmed sources of TeV γ-rays.

The analyses described here have only been performed on the data recorded through 2004 January. The HiRes detector continued observations through the end of 2006 March. The independent sample of data since 2004 January is approximately 70 % of the size of the sample analyzed here. The ongoing analysis of this sample will provide an opportunity to test the correlations in Table 10.4. It should be noted that while the correlation signals appear stronger for the events above 10^{19} eV, a conservative approach which includes consideration of the entire data set will avoid the possibility that a real correlation has been "over-tuned" by an arbitrary threshold and is missed in a future analysis.

As mentioned earlier, real correlations on the scale of the detector angular resolution would suggest neutral cosmic ray primaries for these events, or at least that the primaries were neutral during significant portions of their journey through galactic and extragalactic magnetic fields. Primaries such as neutrons and photons are problematic, however, because of short mean free paths ($\sim$ a few Mpc) at these energies. By observing the longitudinal development of the shower, HiRes has some sensitivity to the primary particle type. Showers induced by photons, for example, tend to develop lower in the atmosphere than those induced by nucleons. Future analyses should address the question whether the correlated HiRes stereo events are compatible with gamma-induced showers.

Chapter 11

Conclusions

There have been many predictions that on small angular scales the arrival directions of ultrahigh energy cosmic rays should exhibit measurable anisotropy. Recent results from the AGASA experiment have indeed claimed this is the case. The present work has shown, however, that the significance of those results was overstated. Analyzing the much higher resolution data set from the HiRes experiment, separately and in combination with the AGASA data set, we have demonstrated that there is no statistically significant evidence for small-scale anisotropy in the arrival directions of ultrahigh energy cosmic rays at present.

Hints of anisotropy persist, necessarily just beyond the level detectable with current statistics. Investigating these with the addition of new data should be a priority. To that end, three analysis methods have been presented in this thesis which can maximize the information extracted from cosmic ray arrival directions while minimizing statistical bias. These are an autocorrelation scan for clustering searches, a maximum likelihood point-source search, and a maximum likelihood test for cross correlation with catalogs of astronomical objects. In particular, the point-source search and the

test of cross-correlation with BL Lacertae objects await the remaining HiRes data sample, which extends through the beginning of 2006 April. The analysis of this data is ongoing.

It should nevertheless be emphasized that isotropy, too, is a useful observation. It has long been known that on large angular scales the arrival directions of ultrahigh energy cosmic rays are remarkably isotropic, and this lack of directionality is informative. It is the absence of any correlations with our galaxy that serves as the best indication that ultrahigh energy cosmic rays originate far outside the Milky Way. Continued observation of isotropy on small angular scales at the highest energies promises to inform us as well.

Bibliography

[1] H. Pleijel. Nobel presentation speech, Dec. 10 1936.

[2] William L. Laurence. Millikan Retorts Hotly to Compton in Cosmic Ray Clash. *New York Times*, page 1, Dec. 31 1932.

[3] R. A. Millikan and G. H. Cameron. The Origin of the Cosmic Rays. *Phys. Rev.*, 32:533, 1928.

[4] J. W. Cronin. Cosmic rays: The most energetic particles in the universe. *Rev. Mod. Phys.*, 71:S165–S172, 1999.

[5] A. H. Compton. A Geographic Study of Cosmic Rays. *Phys. Rev.*, 43:387, 1933.

[6] A. A. Watson. Cosmic rays of the highest energies. *Contemp. Phys.*, 43:181–195, 2002.

[7] Pierre Auger, P. Ehrenfest, R. Maze, J. Daudin, and A. Freon Robley. Extensive cosmic-ray showers. *Rev. Mod. Phys.*, 11:288–291, 1939.

[8] Enrico Fermi. On the Origin of the Cosmic Radiation. *Phys. Rev.*, 75:1169–1174, 1949.

[9] N. P. Samios. Stable particles. *Nucl. Phys. Proc. Suppl.*, 36:83–96, 1994.

[10] John Linsley. Evidence for a primary cosmic-ray particle with energy 10^{20} eV. *Phys. Rev. Lett.*, 10:146–148, 1963.

[11] Todor Stanev. Ultra high energy cosmic rays. *ECONF*, C040802:L020, 2004.

[12] F. Hoyle. Origin of Cosmic X Rays. *Phys. Rev. Lett.*, 15:131–132, 1965.

[13] Kenneth Greisen. End to the cosmic ray spectrum? *Phys. Rev. Lett.*, 16:748–750, 1966.

[14] G. T. Zatsepin and V. A. Kuzmin. Upper limit of the spectrum of cosmic rays. *JETP Lett.*, 4:78–80, 1966.

[15] T. K. Gaisser. *Cosmic rays and particle physics*. Cambridge Univ. Press, 1990.

[16] Angela V. Olinto. Highest energy cosmic rays. *AIP Conf. Proc.*, 745:48–59, 2005.

[17] F. A. Aharonian et al. High-energy particle acceleration in the shell of a supernova remnant. *Nature*, 432:75–77, 2004.

[18] A. M. Hillas. The origin of ultrahigh-energy cosmic rays. *Ann. Rev. Astron. Astrophys.*, 22:425–444, 1984.

[19] Pijushpani Bhattacharjee and Gunter Sigl. Origin and propagation of extremely high energy cosmic rays. *Phys. Rept.*, 327:109–247, 2000.

[20] D. R. Bergman. Fitting the HiRes spectra. 2005. To appear in the proceedings of the 29th International Cosmic Ray Conference, Pune, India.

[21] V. Berezinsky, M. Kachelriess, and A. Vilenkin. Ultra-high energy cosmic rays without GZK cutoff. *Phys. Rev. Lett.*, 79:4302–4305, 1997.

[22] M. Takeda et al. Extension of the cosmic-ray energy spectrum beyond the predicted Greisen-Zatsepin-Kuzmin cutoff. *Phys. Rev. Lett.*, 81:1163–1166, 1998.

[23] Masahiro Takeda et al. Energy determination in the Akeno Giant Air Shower Array experiment. *Astropart. Phys.*, 19:447–462, 2003.

[24] R. U. Abbasi et al. Measurement of the flux of ultrahigh energy cosmic rays from monocular observations by the High Resolution Fly's Eye experiment. *Phys. Rev. Lett.*, 92:151101, 2004.

[25] T. Abu-Zayyad et al. Measurement of the spectrum of UHE cosmic rays by the FADC detector of the HiRes experiment. *Astropart. Phys.*, 23:157–174, 2005.

[26] R. U. Abbasi et al. Observation of the ankle and evidence for a high-energy break in the cosmic ray spectrum. *Phys. Lett.*, B619:271–280, 2005.

[27] Daniel De Marco, Pasquale Blasi, and Angela V. Olinto. On the statistical significance of the GZK feature in the spectrum of ultra high energy cosmic rays. *Astropart. Phys.*, 20:53–65, 2003.

[28] Paul Sommers. First estimate of the primary cosmic ray energy spectrum above 3-EeV from the Pierre Auger observatory. 2005. To appear in the proceedings of the 29th International Cosmic Ray Conference (ICRC 2005), Pune, India.

[29] Stefan Westerhoff. Experimental ultra-high-energy cosmic ray physics. 2005. To appear in the proceedings of 22nd International Symposium on Lepton-Photon Interactions at High Energy (LP 2005), Uppsala, Sweden, 30 Jun - 5 Jul 2005.

[30] Paul M. Mantsch. The Pierre Auger Observatory: Progress and first results.

2005. To appear in the proceedings of the 29th International Cosmic Ray Conference (ICRC 2005), Pune, India.

[31] Diego Harari, Silvia Mollerach, Esteban Roulet, and Federico Sanchez. Lensing of ultra-high energy cosmic rays in turbulent magnetic fields. *JHEP*, 03:045, 2002.

[32] P. P. Kronberg. Intergalactic magnetic fields. *Phys. Today*, 55N12:40–46, 2002.

[33] Jaime Alvarez-Muniz and Todor Stanev. The large scale structure of the galactic magnetic field and high energy cosmic ray anisotropy. 2005.

[34] M. Kachelriess, Pasquale Dario Serpico, and M. Teshima. The galactic magnetic field as spectrograph for ultra-high energy cosmic rays. 2005.

[35] Klaus Dolag, Dario Grasso, Volker Springel, and Igor Tkachev. Mapping deflections of Ultra-High Energy Cosmic Rays in Constrained Simulations of Extragalactic Magnetic Fields. *JETP Lett.*, 79:583–587, 2004.

[36] Pasquale Blasi, Scott Burles, and Angela V. Olinto. Cosmological Magnetic Fields Limits in an Inhomogeneous Universe. *Astrophys. J.*, 514:L79–L82, 1999.

[37] Eric Armengaud, Guenter Sigl, and Francesco Miniati. Ultrahigh energy nuclei propagation in a structured, magnetized universe. *Phys. Rev.*, D72:043009, 2005.

[38] Diego F. Torres and Luis A. Anchordoqui. Astrophysical origins of ultrahigh energy cosmic rays. *Rept. Prog. Phys.*, 67:1663–1730, 2004.

[39] Graciela Gelmini, Oleg Kalashev, and Dmitry V. Semikoz. GZK photons as ultra high energy cosmic rays. 2005.

[40] W. Heitler. *Quantum Theory of Radiation*. Oxford Univ. Press, 1944.

[41] Tanguy Pierog, R. Engel, and D. Heck. Impact of uncertainties in hadron production on air-shower predictions. 2006. To appear in the proceedings of International Conference on Interconnection between High Energy Physics and Astroparticle Physics: From Colliders to Cosmic Rays, Prague, Czech Republic, 7-13 Sep 2005.

[42] Auger Collaboration. The Pierre Auger Observatory Design Report, Second Ed. Mar. 1997. `http://www.auger.org/admin/DesignReport/`.

[43] T. K. Gaisser and A. M. Hillas. Reliability of the method of constant intensity cuts for reconstructing the average development of vertical showers. In *International Cosmic Ray Conference*, pages 353–357, 1978.

[44] K. M. Simpson. *Studies of Cosmic Ray Composition Using a Hybrid Fluorescence Detector*. PhD thesis, Univ. Adelaide, Adelaide, 2001.

[45] Gregory Clark Archbold. *A Study of the composition of ultrahigh-energy cosmic rays using the High Resolution Fly's Eye*. PhD thesis, Univ. Utah, 2002. UMI-30-54514.

[46] J. H. Boyer, B. C. Knapp, E. J. Mannel, and M. Seman. FADC-based DAQ for HiRes Fly's Eye. *Nucl. Instrum. Meth.*, A482:457–474, 2002.

[47] J. N. Matthews and S. B. Thomas. The absolute calibration of the HiRes detectors. *Proceedings 27th International Cosmic Ray Conference (ICRC)*, pages 350–353, 2001.

[48] R. Abbasi et al. Techniques for measuring atmospheric aerosols at the High Resolution Fly's Eye experiment. *Astropart. Phys.*, 25:74–83, 2006.

[49] J. Boyer, E. Mannel, and L. Wiencke. 'Shoot the shower': Probing atmospheric clarity of the shower / detector plane at HiRes. 2003. Prepared for 28th International Cosmic Ray Conferences (ICRC 2003), Tsukuba, Japan, 31 Jul - 7 Aug 2003.

[50] R. Abbasi et al. A measurement of time-averaged aerosol optical depth using air-showers observed in stereo by HiRes. *Astropart. Phys.*, 25:93–97, 2006.

[51] D. Heck, G. Schatz, T. Thouw, J. Knapp, and J. N. Capdevielle. CORSIKA: A Monte Carlo code to simulate extensive air showers. 1998. FZKA-6019.

[52] P. A. Sadowski et al. Geometry and optics calibration for air fluorescence detectors using star light. *Astropart. Phys.*, 18:237–248, 2002.

[53] N. Chiba et al. Akeno Giant Air Shower Array (AGASA) covering 100 km^2 area. *Nucl. Instrum. Meth.*, A311:338–349, 1992.

[54] N. Hayashida et al. Possible clustering of the most energetic cosmic rays within a limited space angle observed by the Akeno Giant Air Shower Array. *Phys. Rev. Lett.*, 77:1000–1003, 1996.

[55] M. Takeda et al. Small-scale anisotropy of cosmic rays above 10^{19} eV observed with the Akeno Giant Air Shower Array. *Astrophys. J.*, 522:225–237, 1999.

[56] N. Hayashida et al. Updated AGASA event list above 4×10^{19} eV. *Astron. J.*, 120:2190, 2000.

[57] M. Takeda et al. Clusters of cosmic rays above 10^{19} eV observed with AGASA. *Proceedings 27th International Cosmic Ray Conference (ICRC)*, pages 341–344, 2001.

[58] M. Teshima et al. The arrival direction distribution of extremely high energy cosmic rays observed by AGASA. *Proceedings 28th International Cosmic Ray Conference (ICRC)*, pages 437–440, 2003.

[59] AGASA Collaboration. AGASA results, Jun. 2003. `http://www-akeno.icrr.u-tokyo.ac.jp/AGASA/results.html`.

[60] P. G. Tinyakov and I. I. Tkachev. Correlation function of ultra-high energy cosmic rays favors point sources. *JETP Lett.*, 74:1–5, 2001.

[61] Stephen D. Landy and Alexander S. Szalay. Bias and variance of angular correlation functions. *Astrophys. J.*, 412:64, 1993.

[62] D. E. Alexandreas et al. Point source search techniques in ultrahigh-energy gamma- ray astronomy. *Nucl. Instrum. Meth.*, A328:570–577, 1993.

[63] Y. Uchihori et al. Cluster analysis of extremely high energy cosmic rays in the northern sky. *Astropart. Phys.*, 13:151–160, 2000.

[64] W. W. Kinnison et al. A Search for mu+ $\rightarrow$ e+ gamma. *Phys. Rev.*, D25:2846, 1982.

[65] Morgan O. Wascko. *Study of the Shadow of the Moon in Very High Energy Cosmic Rays with the Milagrito Water Cherenkov Detector*. PhD thesis, Univ. California, Riverside, 2001.

[66] S. L. Meyer. *Data Analysis for Scientists and Engineers.* John Wiley and Sons, Inc., 1975.

[67] R. U. Abbasi et al. A search for arrival direction clustering in the HiRes-I monocular data above $10^{19.5}$ eV. *Astropart. Phys.*, 22:139–149, 2004.

[68] P. G. Tinyakov and I. I. Tkachev. BL Lacertae are sources of the observed ultra-high energy cosmic rays. *JETP Lett.*, 74:445–448, 2001.

[69] P. G. Tinyakov and I. I. Tkachev. Tracing protons through the galactic magnetic field: A clue for charge composition of ultrahigh-energy cosmic rays. *Astropart. Phys.*, 18:165–172, 2002.

[70] D. S. Gorbunov, P. G. Tinyakov, I. I. Tkachev, and S. V. Troitsky. Evidence for a connection between gamma-ray and highest- energy cosmic ray emissions by BL Lacs. *Astrophys. J.*, 577:L93, 2002.

[71] R. C Hartman et al. The Third EGRET catalog of high-energy gamma-ray sources. *Astrophys. J. Suppl.*, 123:79, 1999.

[72] D. Horan and T. C. Weekes. Extragalactic sources of TeV gamma rays: A summary. *New Astron. Rev.*, 48:527–535, 2004.

[73] V. S. Berezinskii, S. V. Bulanov, V. A. Dogiel, and V. S. Ptuskin. *Astrophysics of cosmic rays.* Amsterdam: North-Holland, 1990, edited by Ginzburg, V.L., 1990.

[74] Paolo S. Coppi and Felix A. Aharonian. Constraints on the very high energy emissivity of the universe from the diffuse GeV gamma-ray background. *Astrophys. J.*, 487:L9–L12, 1997.

[75] N. W. Evans, F. Ferrer, and Subir Sarkar. The clustering of ultra-high energy cosmic rays and their sources. *Phys. Rev.*, D67:103005, 2003.

[76] Boris E. Stern and Juri Poutanen. Blind search for the real sample: Application to the origin of ultra-high energy cosmic rays. *Astrophys. J.*, 623:L33–L36, 2005.

[77] Diego F. Torres, Stephen Reucroft, Olaf Reimer, and Luis A. Anchordoqui. On the cross correlation between the arrival direction of ultra-high energy cosmic rays, BL Lacertae, and EGRET detections: A new way to identify EGRET sources. *Astrophys. J.*, 595:L13–L16, 2003.

[78] M.-P. Veron-Cetty and P. Veron. *A catalogue of quasars and active nuclei*. 9th ed. Garching: ESO Scientific Report no. 19., 2000.

[79] M.-P. Veron-Cetty and P. Veron. A catalogue of quasars and active nuclei: 10th edition. *A&A*, 374:92–94, July 2001.

[80] D. S. Gorbunov, P. G. Tinyakov, I. I. Tkachev, and S. V. Troitsky. Testing the correlations between ultra-high-energy cosmic rays and BL Lac type objects with HiRes stereoscopic data. *JETP Lett.*, 80:145–148, 2004.

[81] N. I. Fisher, T. Lewis, and B. J. J. Embleton. *Statistical analysis of spherical data*. Cambridge: University Press, 1987.

Appendix A

Gaussian Distributions in Two Dimensions

This appendix collects a number of useful definitions and some straightforward results for two-dimensional Gaussian distributions and related functions. The results hold strictly in flat, two dimensional space, but can be reasonably applied to the sphere when the distribution is sharply concentrated, as has been the case for the small-scale anisotropy studies in this thesis (i.e. point-spread functions and angular separations $\lesssim 5^\circ$). In the event this is not the case, we point the reader toward an example of a Gaussian-like function constructed for spherical distributions.

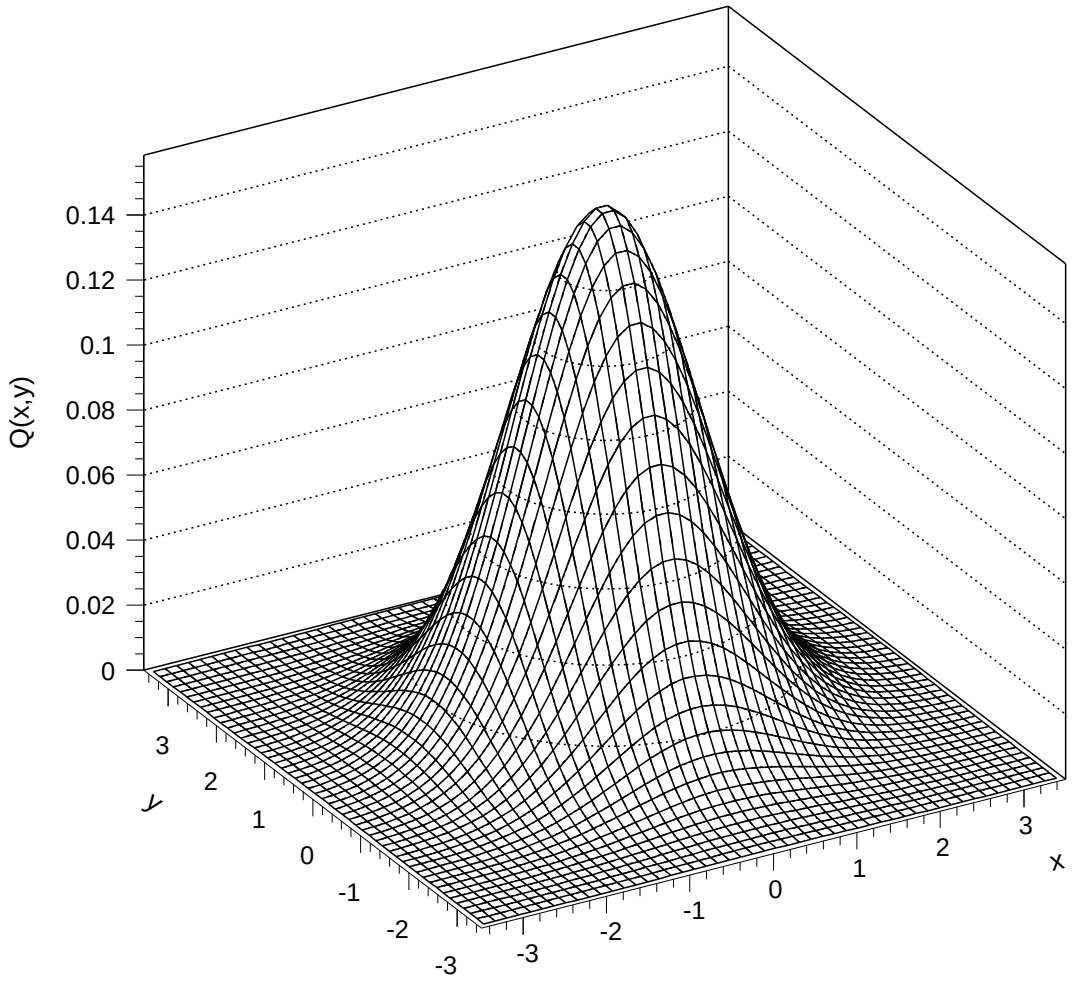

Figure A.1: The two-dimensional Gaussian distribution $Q(x, y)\, dA$, with x and y in units of σ.

A.1 $Q(x,y)$: A Gaussian Point-Spread Function

Working in the x-y plane, a two dimensional Gaussian distribution (representing a detector's point-spread function, for example) is given by:

$$Q(x,y)\,dx\,dy = \frac{1}{\sigma_x\sqrt{2\pi}}\,\exp\left(-\frac{x^2}{2{\sigma_x}^2}\right) \times \frac{1}{\sigma_y\sqrt{2\pi}}\,\exp\left(-\frac{y^2}{2{\sigma_y}^2}\right)\,dx\,dy\,, \qquad \text{(A.1)}$$

which is illustrated in Fig. A.1. If the distribution is circular (as shown in the figure), we can let $\sigma \equiv \sigma_x = \sigma_y$ and $r^2 \equiv x^2 + y^2$ and write

$$Q(r)\,dA = \frac{1}{2\pi\sigma^2}\,\exp\left(-\frac{r^2}{2\sigma^2}\right)\,dA\,. \qquad \text{(A.2)}$$

Note that while there is no ϕ dependence because of the circular symmetry, this is still the two-dimensional probability density function, as indicated by the surface element dA.

In the maximum likelihood analyses of Ch. 9 and Ch. 10, $Q(r)$ is the form of the source probability distribution, with r being replaced by the angular distance $\Delta\theta$ between the source location and the event arrival direction.

A.2 $P(r)$: Opening Angle Distribution for Gaussian Errors

The one-dimensional distribution of r values, $P(r)dr$, for the function $Q(r)dA$ is given by:

$$P(r)dr \;=\; \int_0^{2\pi} Q(r)\,2\pi r\,d\phi\;dr$$

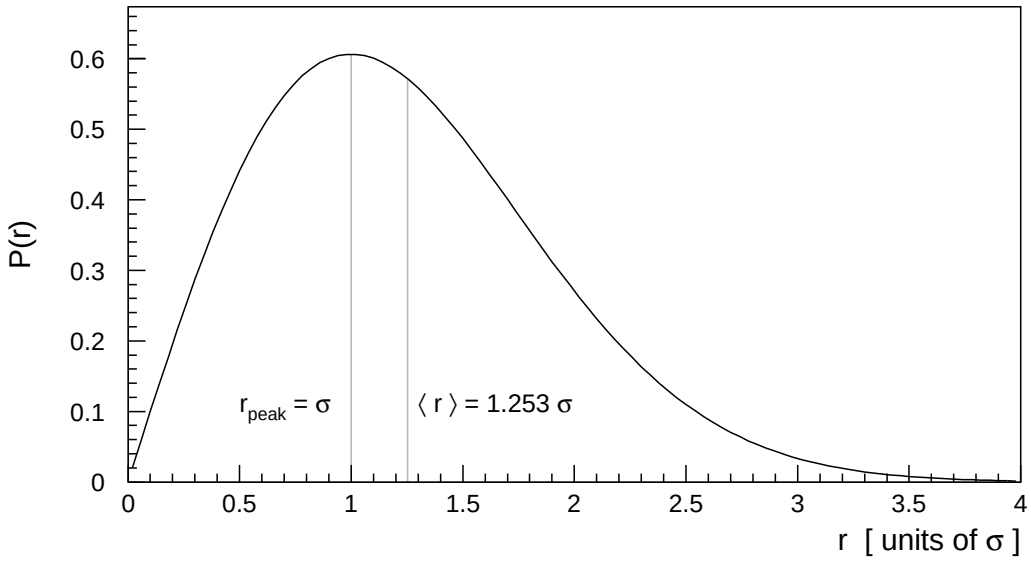

Figure A.2: The distribution $P(r)dr$ of r values for the two-dimensional Gaussian function $Q(r)\,dA$.

$$P(r)dr \;=\; \frac{r}{\sigma^2}\exp\left(-\frac{r^2}{2\sigma^2}\right)dr\,, \tag{A.3}$$

which is shown in Fig. A.2.

Two useful values for this distribution are the location of the maximum of P, $r_{\rm peak} = \sigma$, and the mean value of r:

$$\begin{aligned}\langle r\rangle &= \int_0^\infty rP(r)dr \\ \langle r\rangle &= \sqrt{\frac{\pi}{2}}\,\sigma \\ \langle r\rangle &\approx 1.2533\,\sigma\ .\end{aligned} \tag{A.4}$$

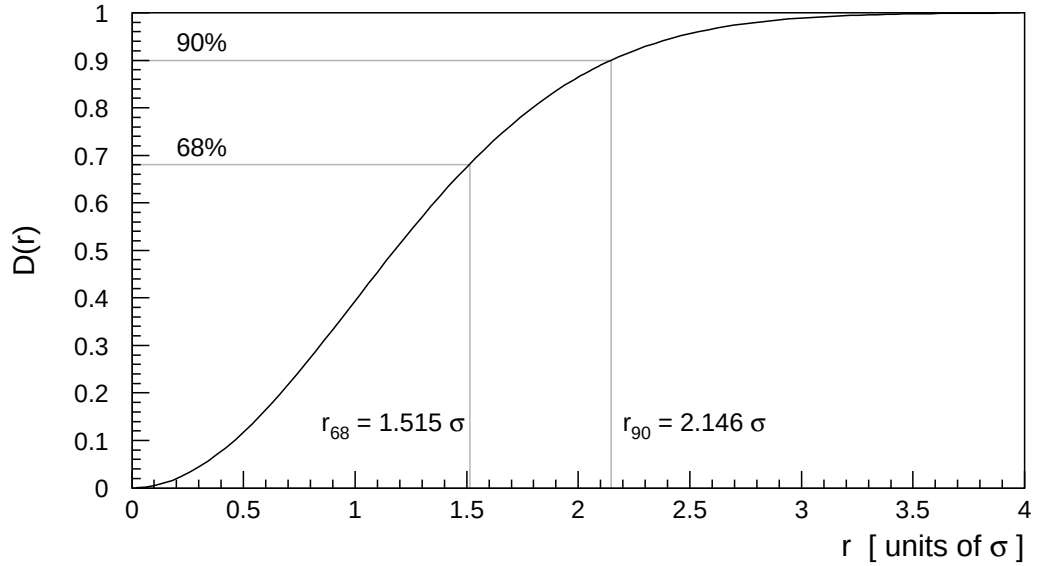

Figure A.3: The cumulative distribution $D(r)$ of the probability density function $P(r)dr$.

A.3 $D(r)$: the Cumulative Distribution of $P(r)$

We are also frequently interested in $D(r)$, the cumulative distribution function of $P(r)$:

$$\begin{aligned} D(r) &= \int_0^r P(r')dr' \\ &= 1 - \exp\left(-\frac{r^2}{2\sigma^2}\right), \end{aligned} \tag{A.5}$$

which is shown in Fig. A.3.

A few useful values of r and $D(r)$ are given in Table A.1. In particular, the value of r which encloses 68% of the distribution is given by $r_{68} = 1.5152\,\sigma$.

Angular resolution is often described in terms of the quantity $\Delta\theta_{68}$, that is, the

r	$D(r)$
$1.5152\,\sigma$	0.6827
$2.1460\,\sigma$	0.90
$2.4860\,\sigma$	0.9545
$3.4393\,\sigma$	0.9973

Table A.1: Values of r and the corresponding fraction of the $P(r)$ distribution they enclose (indicated by $D(r)$, the cumulative distribution of $P(r)$).

angle which encloses 68% of the distribution of angular reconstruction errors for Monte Carlo events. This is a useful definition whether or not the point-spread function is actually Gaussian. The relation $\sigma = \Delta\theta_{68}/1.5152$, however, obviously assumes Gaussianity.

A.4 Distributions on a Sphere

In applying the $Q(x, y)$ distribution and related functions to data on a sphere, we assume that the distribution is concentrated in a small angular region and that curvature can be neglected. If, however, one were interested in a Gaussian-like distribution spread over an arbitrarily large part of the sphere, this approximation would not be valid. In such cases, one can use e.g. the Fisher density [81]:

$$Q(\theta)d\Omega = N \exp\left(k \cos\theta\right)\, d\Omega \tag{A.6}$$

where θ is the angular separation between the center of the distribution and the point in question, and N is the normalization factor given by

$$N \;=\; \left(\int_0^{\pi} \int_0^{2\pi} \exp\left(k \cos\theta\right)\, d\theta\, \sin\theta\, d\phi \right)^{-1}$$

$$N = \frac{k}{4\pi \sinh k}\,. \tag{A.7}$$

Note that larger values of k correspond to more sharply concentrated distributions. If we then use the small-angle approximation $\cos\theta \approx 1 - \theta^2/2$ for such cases, we find

$$Q(\theta)\, d\Omega \approx \frac{k \exp(k)}{4\pi \sinh k}\, \exp\left(-\frac{k\,\theta^2}{2}\right)\, d\Omega \tag{A.8}$$

which, setting $k = 1/\sigma^2$, approaches in the large k limit the flat Gaussian distribution $Q(r)$ in Eq. A.2.