J. Phys. G: Nucl. Part. Phys. **34** (2007) R359–R400 doi:10.1088/0954-3899/34/10/R01

TOPICAL REVIEW

Cosmic rays: the Second Knee and beyond

Douglas R Bergman[1] **and John W Belz**[2]

[1] Department of Physics and Astronomy, Rutgers—The State University of New Jersey, Piscataway, NJ, USA
[2] Department of Physics, University of Utah, Salt Lake City, UT, USA

E-mail: bergman@physics.rutgers.edu and belz@cosmic.utah.edu

Received 11 April 2007
Published 12 September 2007
Online at stacks.iop.org/JPhysG/34/R359

Abstract

We conduct a review of experimental results on ultra-high energy cosmic rays (UHECRs) including measurements of the features of the spectrum, the composition of the primary particle flux and the search for anisotropy in event arrival direction. We find that while there is a general consensus on the features in the spectrum—the Second Knee, the Ankle and (to a lesser extent) the GZK Cutoff—there is little consensus on the composition of the primaries that accompany these features. This lack of consensus on the composition makes interpretation of the agreed upon features problematic. There is also little direct evidence about potential sources of UHECRs, as early reports of arrival-direction anisotropies have not been confirmed in independent measurements.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Ultra-high energy cosmic rays (UHECRs) are the most energetic form of radiation known to hit the earth. At these energies, above 10^{18} eV, one would like to understand the workings of the astrophysical accelerators which able to produce such high energies, energies many orders of magnitude higher than what is available at terrestrial accelerators. These high energy particles also have, perhaps, much to tell us about the regions in which they were accelerated and the vast spaces through which they passed on their way to us. However, before we can understand, we must measure: what is the energy spectrum, what kind of particles are they, where do they come from?

This will be a review of the latest experimental results from which one may hope to understand UHECRs. There has been significant experimental activity since the last experimental review appeared in this journal (Yoshida and Dai 1998), and there has been some movement towards a consensus on the existence and energies of various features in the

UHECR spectrum. There in not yet, however, a consensus on the best way to interpret those features.

Cosmic rays were discovered by Hess (1912), using the fact that the ionization of air increases with altitude. This cosmic radiation was later directly observed in cloud chambers. Pierre Auger (Auger *et al* 1939) observed coincident hits over a wide area with his detectors, showing that the primary cosmic rays induce a cascade of particles, known as an extensive air shower (EAS), when they encounter the atmosphere. From the numbers of particles involved in these showers, Auger was able to estimate the energy of some showers, and show that some must be very energetic. The phenomenology of EAS was worked out by Heitler (1938) and others.

Several features have been identified or proposed in the spectrum of cosmic rays. The first of these features to be identified was a softening of the spectrum at an energy of about 3×10^{15} eV, an energy below the range we consider in the paper. Since the flux bends *down* at this point, the feature was christened the Knee. When later another downturn at higher energy, about 4×10^{17} eV, was observed, it was naturally called the Second Knee. This is at the lower end of the energy range we consider in this review, the ultra-high energy range.

About a order-of-magnitude higher in energy than the Second Knee, the spectrum becomes harder again. Continuing with the anatomic analogy, this is called the Ankle, because the bend is in the opposite sense than the one seen at the Knee. Finally, at yet another order-of-magnitude higher in energy, one expects a drastic reduction in the flux, the GZK Cutoff (Greisen 1966, Zatsepin and Kuz'min 1966), due to energy losses of the cosmic rays in the cosmic microwave background radiation during their long propagation to us.

The first dedicated experiment to measure the energy and flux of UHECRs was built by Linsley (1963) at Volcano Ranch. Data from Volcano Ranch and other early experiments are covered in other reviews on this subject. In this review, we will cover only those experimental results which were released since the review of Yoshida and Dai (1998) or those which still have a significant impact on the world data set at a given energy. This includes results from Haverah Park, the HiRes Prototype/MIA hybrid experiment, the Akeno 1 km Array, Yakutsk, Fly's Eye, the Akeno Giant Air Shower Array (AGASA), the High Resolution Fly's Eye (HiRes) and the Pierre Auger Observatory.

2. The experiments

The observable characteristics of an UHECR are its energy, the type of particle it is and the direction from which it came. All experiments try to measure these characteristics, directly or indirectly, with various degrees of precision. Since the flux of UHECRs is so low, direct measurement of these properties is impractical, so one must measure the properties of the extensive air shower (EAS) created by the cosmic ray when it enters the atmosphere. In the case of measuring the primary energy and direction, the properties of the EAS are a reasonable proxy for the properties of the primary cosmic ray. All the energy of the primary goes into the shower and most of the shower energy is deposited in the atmosphere. Likewise, the momentum of the primary particle is so much greater than any transverse momentum generated in the shower that the shower points in the same direction as the primary. However, the primary particle type must be inferred from the way the shower develops, which makes it hard to determine the particle type on a shower-by-shower basis even for the best measurements of EAS development.

There are two principal techniques used to detect and measure EASs: measure the density of shower particle at the ground or measure the amount of fluorescence light emitted by the

atmosphere as the shower passes through it. The two techniques can also be used together for hybrid measurements.

Ground arrays sample the shower front at one level, with a sparse array of detectors on the ground. The detectors are typically either slabs of scintillator or tanks of water, which have different sensitivity to the particle components of the shower. Scintillators are primarily sensitive to electrons and photons, but appropriate shielding can allow one to separate out the muon and hadron component as well as differentiate the electrons from the photons. Water tanks, which detect shower particles by the Cerenkov radiation they emit while passing through the detector, are much more sensitive to muons than electromagnetic particles. The shower direction is inferred from the relative timing of the various detector elements as the shower front sweeps across them, but the energy and composition of the shower must be inferred indirectly from the size and shape of the shower footprint. The shower energy is usually determined by the shower density at a given distance from the core (600–1000 m). The energy is roughly proportional to this density and the shower-to-shower fluctuations are reduced at this distance. However, one must take into account the attenuation of the shower as showers at different zenith angles have traversed different amounts of atmosphere. Both the normalization and the attenuation correction contribute to systematic uncertainties and are dependent on shower development modelling. One can gain more information by having separate ground stations which are sensitive to electrons or to muons or to photons, but the model dependence is hard to avoid. On the other hand, shower arrays can run continuously and are mostly independent of the weather.

Fluorescence detectors collect the fluorescence light generated as the shower particles excite the nitrogen in the air. The amount of light produced in this way is proportional to the primary energy, so this technique provides a calorimetric measurement of the shower energy. The direction of shower can be determined from the relative timing of light arriving at the detector, though there are large correlations in fitting between the distance to a shower and the angle it makes with respect to the viewer. A much better determination of the shower geometry is available when one views a shower with two detectors simultaneously which is known as stereo observation (as opposed to monocular observation where one uses timing). In this case, the geometry of the shower is determined uniquely by the two shower-detector planes. The composition of the primary cosmic ray determines the longitudinal development of the shower. This development is directly visible to the fluorescence detector and is often abbreviated by measuring the depth in the atmosphere at which the shower reached its maximum size, $X_{\max}$. The drawback to using a fluorescence detector is that it can only be operated on dark, moonless nights. In addition, one must control the level and variation of aerosols in the atmosphere by choosing an appropriate (desert) site.

We now discuss the various detectors roughly in order of exposure. The exposures of all these experiments is shown as a function of energy in figure 1. In the cases where exposures were not explicitly published by the experimental groups, they have been inferred from the published flux values and uncertainties.

2.1. Haverah Park

Haverah Park was operated from 1962 to 1987 and was the pioneer in using Cerenkov water tanks to sample the shower front at the ground. The array was located at 54.0° N, 1.6° W, near Leeds University, at an atmospheric depth of 1016 g cm^{-2}. The tanks were 2.29 m^2 in area and 1.2 m deep (Lawrence *et al* 1991), with the Cerenkov light collected by a single 5 inch diameter photomultiplier tube suspended so that the photocathode was just in contact with the water. The tanks were then grouped together to make large detection areas able to

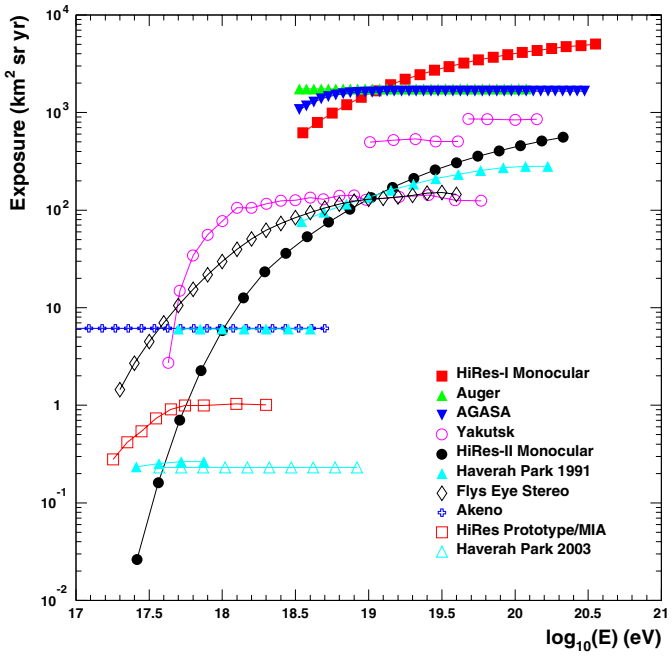

Figure 1. The exposures of all the experiments discussed in the text. The data for the exposures come from the following papers: Haverah Park (Lawrence *et al* 1991, Ave *et al* 2003a), HiRes Prototype/MIA (Abu-Zayyad *et al* 2001), Akeno (Nagano *et al* 1992), Fly's Eye (Bird *et al* 1994), Yakutsk (all arrays) (Egorova *et al* 2004), AGASA (Takeda *et al* 2003), Auger (Pierre Auger Collaboration 2005), HiRes (both monocular measurements) (Abbasi *et al* 2007a).

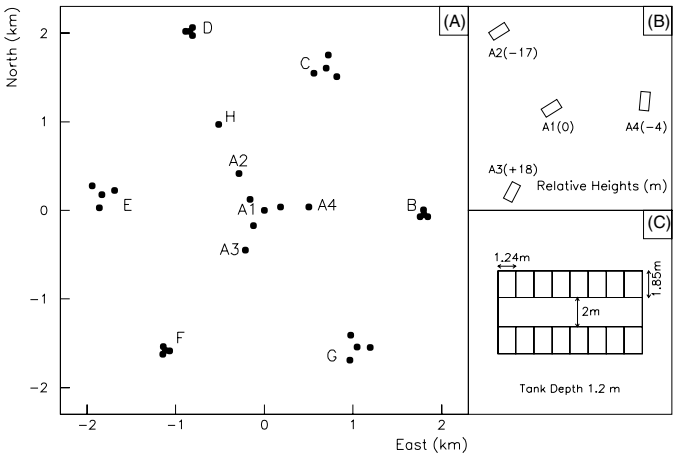

Figure 2. Layout of the Haverah Park array. (A) The whole array. (B) The orientation and relative heights of the detector huts A1–A4. (C) The arrangement of water tanks within one of the four main A-site huts. Figure taken from Ave *et al* (2003b).

detect relatively small fluxes of charged particles. At the centre of the array (see figure 2), a set of 4 × 34 m^2 (A1)–(A4) detectors was used for triggering the rest of the detector. Signals

above threshold (0.3 vem m^{-2}) in the central detector (A1) and at least two of the other three A detectors (at 500 m) were required to form a trigger.

Six groups of tanks (each 4×13.5 m^2) were placed in a ring around the centre tanks. These tanks were used to constrain the core position in large showers, and thus increase the aperture of the entire detector. An infill array of 30 1 m^2 tanks with a spacing of about 150 m was also operated for several years in the area between the central A tanks. Data acquisition consisted of both photography of oscilloscope traces (which allowed for pulse shape and rise time analyses) and digital methods.

Yoshida and Dai (1998) discuss the 'final' analysis of the Haverah Park array (Lawrence *et al* 1991). Since that time, however, the Haverah Park data have been reanalysed using shower simulation code unavailable in 1991. Results from this reanalysis were published in Hinton *et al* (1999), Ave *et al* (2000, 2001, 2002, 2003a, 2003b).

The new analysis uses QGSJet (Kalmykov *et al* 1997) and Corsika (Heck *et al* 1998) to model the showers and GEANT (CERN, Application Software Group 1993) to model the detector response. The reanalysis led to a change in the relation between $\rho(600)$, the density of shower muons at 600 m from the core and the primary energy. The new relation leads to a reduction in the energy of about 30%. A change was also made in the attenuation length used to convert $\rho(600)$ of inclined showers to that of the reference angle. The new attenuation length was lower because only events with $\theta < 45^\circ$ were used (as opposed to $\theta < 60^\circ$ in the original analysis). More vertical showers have a larger electron component, which implies a smaller attenuation length.

2.2. The SUGAR array

The Sydney University Giant Air shower Recorder (SUGAR) was operated in Australia at (30.5° S, 149.6° E), from 1968 to 1979 (Winn *et al* 1986a, 1986b). It consisted of pairs of buried liquid scintillator tanks, with the pair separated by 50 m, on a mile (1600 m) square grid. Each scintillator had an effective area of 6.0 m^2. We only comment on the SUGAR measurements of anisotropy, because of problems with afterpulsing in the photomulitplier tubes (Nagano and Watson 2000). The layout of the detector sites is shown in figure 3.

2.3. The HiRes-Prototype/MIA hybrid

While the HiRes-Prototype/MIA experiment ran for only a relatively short time, and has a relatively small exposure, it was the first experiment to use both fluorescence and ground array measurements simultaneously (Abu-Zayyad *et al* 2000b). This hybrid measurement pointed the way towards later experiments such as Auger. It was formed by the fortuitous juxtaposition of the CASA-MIA array nearing decommissioning and the nascent High Resolution Fly's Eye detector (see section 2.8).

The MIA array was an array of 16 patches of 64 scintillation counters (see figure 4, left). Each counter was 1.9 m $\times$ 1.3 m and buried 3 m below the ground, making the counter primarily sensitive to muons. The total area covered by the array was less than a quarter of a km^2, which limited the total exposure available in hybrid (Borione *et al* 1994).

The HiRes Prototype was an array of 14 fluorescence mirrors arranged so the sky coverage formed a tower overlooking the MIA array, which was 3.5 km to the NE (see figure 4, right). Each HiRes Prototype mirror covered a $16^\circ \times 14^\circ$ of the sky with 256 pixels. Each pixel observed about one square degree on the sky. These mirrors were later rearranged and become the first part of the HiRes detector (section 2.8). The two detectors were located at (40.2° N, 112.8° W), on Dugway Proving Grounds in Utah, USA, at an atmospheric

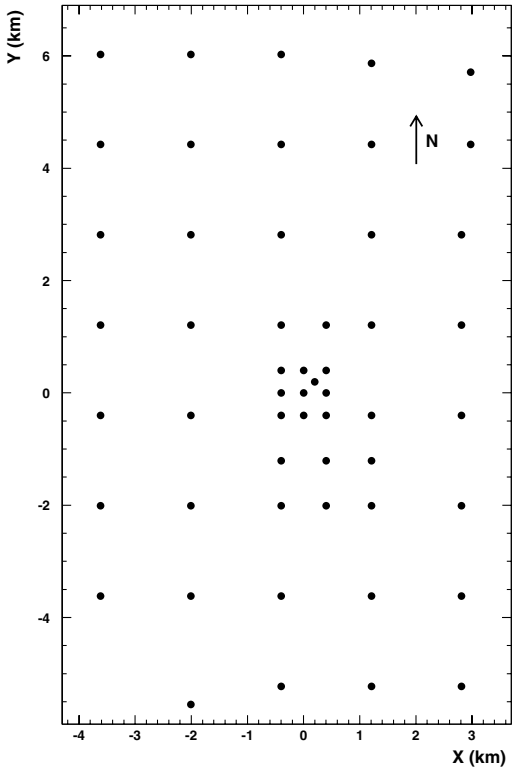

Figure 3. Layout of the SUGAR array. Figure adapted from Winn *et al* (1986b) The origin in the figure is located at (30.5° S, 149.6° E).

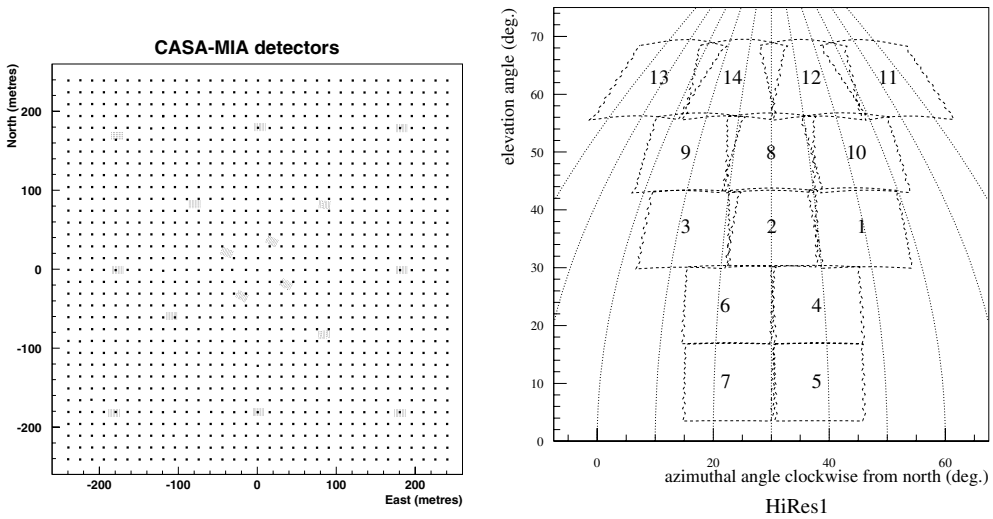

Figure 4. Left: layout of MIA scintillation patches (black rectangles) among the CASA detectors (small squares). Right: arrangement of the HiRes-Prototype mirrors as they view the sky. The MIA array is just below the junction between mirrors 5 and 7 and extends for 4° on either side.

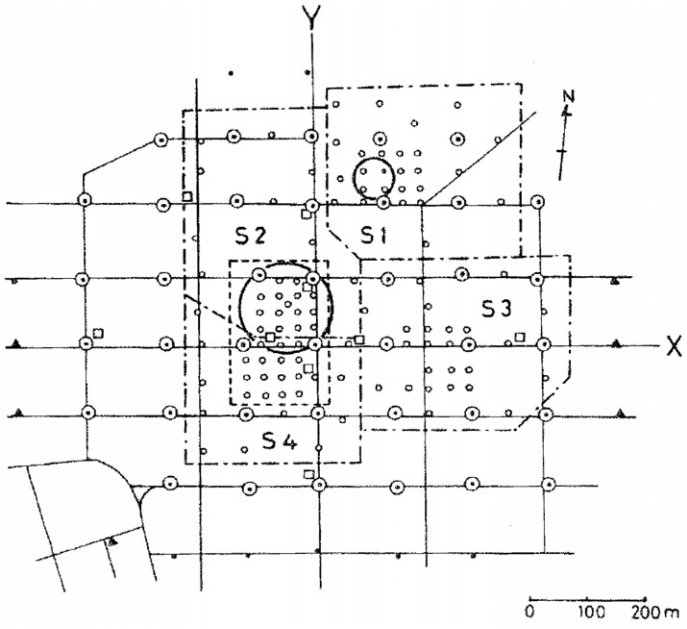

Figure 5. The layout of scintillators in the Akeno 1 km array. Figure taken from Nagano *et al* (1984).

depth of 860 g cm^{-2}. The detectors operated in hybrid mode from August 1993 until May 1996.

2.4. Akeno

The Akeno 1 km^2 array was an array of 156 scintillation counters, each with an area of 1 or 2 m^2. The primary spacing between the counters was 120 m, but three regions had a smaller spacing of 30 m (see figure 5). These closely spaced regions allowed the measurement of showers over a wide range of energies. The total area of the array was about 1 km^2, but a fiducial area of only 700 m $\times$ 600 m was used for the spectrum calculation. The array was located near the Akeno Observatory (35.8° N, 138.5° E) in Japan, at an atmospheric depth of 550 g cm^{-2} (Nagano *et al* 1992), and was operated from the late 1970s through 1990.

2.5. Fly's Eye

While the Fly's Eye experiment was fully covered in Yoshida and Dai (1998), it continues to have significant exposure in some energy ranges above the Second Knee. Fly's Eye was the pioneering fluorescence detector, with full-sky coverage at one site (see figure 6), and about half-sky (with full elevation) coverage from a second site for stereo observation (Baltrusaitis *et al* 1985). The fluorescence light from the EAS was collected by 67 mirrors, each with an area of 1.95 m^2. Each tube viewed a spot on the sky about 5° in diameter. The two sites were separated by a distance of 3.3 km and located at (40.2° N, 112.8° W), on Dugway Proving Grounds in Utah, USA, at an atmospheric depth of 860 g cm^{-2}. The Fly's Eye detector was operated from 1981 through 1992 (Bird *et al* 1994).

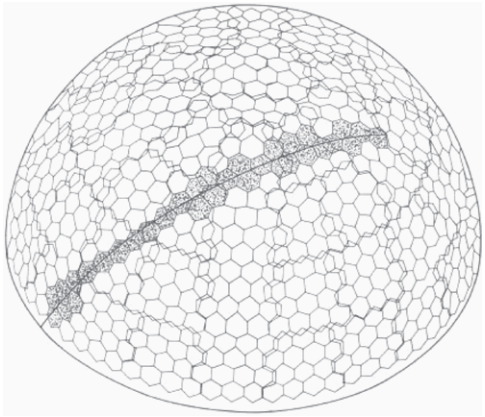

Figure 6. The arrangement of Fly's Eye pixels as projected onto the dome of the sky with the track of an event superimposed. Figure taken from Sokolsky (1989).

2.6. Yakutsk

The Yakutsk array is a set of three nested arrays. At the centre is a closely spaced array of 19 scintillation counters arranged in a hexagonal pattern with 62 m spacing covering an area of 0.026 km^2. Each of these detectors has an area of 0.25 m^2. This is surrounded by more widely spaced counters, each 2 m^2 in area, on triangular grids of spacing 500 m and 1 km. Three counters on the corners of a triangle were required to fire in coincidence to form a trigger for the rest of the array (Afanasiev *et al* 1993, Pravdin *et al* 1999). Prior to 1990, there were 19 counters at a spacing of 500 m, forming 24 trigger triangles and covering an area of 2.5 km^2. The surrounding 1 km array contained 29 counters (with some overlap with the 500 m array), 40 trigger triangles and covering 16 km^2. This layout is shown as the filled circles in figure 7. Between 1990 and 1992, the detectors were rearranged, removing 10 counters from the 1 km array, but adding 18 counters to the 500 m array. This increased the size of the 500 m array substantially at the expense of the 1 km array. In this arrangement, there are 63 triangles in the 500 m array, covering 7.2 km^2, but only 24 triangles in the 1 km array, covering 10 km^2. This layout is shown by the open circles in figure 7 (Ivanov *et al* 2003).

In addition to the scintillators for detecting charged particles, there is an array of Čerenkov detectors, consisting of large photocathode photomultiplier tubes collecting EAS light directly. The PMTs have effecting collecting areas of either 176 or 530 cm^2 (Ivanov *et al* 2003). The positions of these detectors is indicated by open squares in figure 7. Finally, there are a number of buried muon detectors, shown as filled triangles in figure 7.

The Yakutsk array has been in continuous operation since 1970 (Egorova *et al* 2001). It is located at (61.7° N, 129.4° E) (Uchihori *et al* 2000).

2.7. The Akeno Giant Air Shower Array (AGASA)

The Akeno Giant Air Shower Array (AGASA) grew out of the Akeno 1 km^2 Array at the same site (35.8° N, 138.5° E). AGASA is an array of 111 scintillation counters, each with an area of 2.2 m^2, on a roughly square grid with a spacing of about 1 km (see figure 8). It covered a total area of about 100 km^2 (Nagano *et al* 1992, Ohoka *et al* 1997). It was operated from the mid 1980s (as the Akeno 20 km Array; AB in figure 8) through 2004. AGASA was the first detector to have substantial exposure in the region above the expected GZK Cutoff.

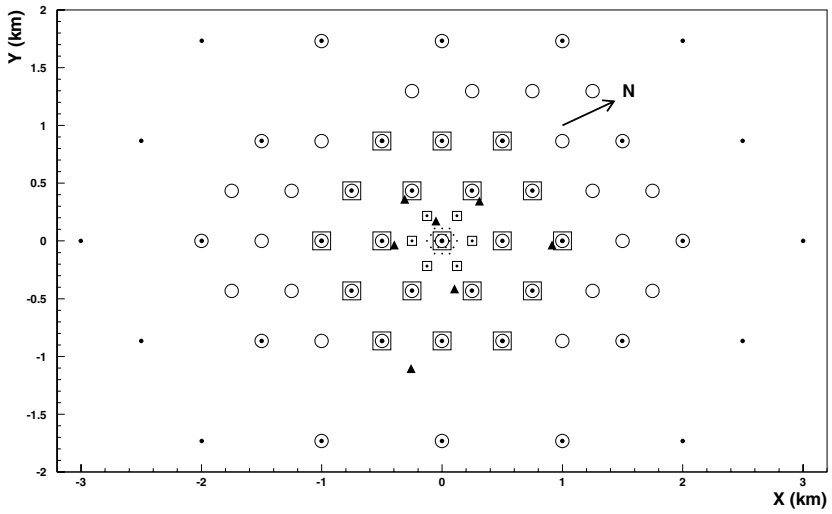

Figure 7. Layout of the Yakutsk array as it existed prior to 1990. Figure adapted from Afanasiev *et al* (1993) and Ivanov *et al* (2003). Filled circles represent the 2 m^2 scintillator detectors in the arrangement before 1990. Open squares represent these detectors after 1992. Open squares represent the position of Čerenkov detectors, while triangles the position of shielded muon detectors.

The signal in any particular counter was digitized by means of a logarithmic amplifier. The charge from the photomultiplier tube was stored on a capacitor, and the capacitor was discharged with given RC time constant. In this way, the time-over-threshold is proportional to the logarithm of the signal size. This system provides a large dynamic range but is susceptible to large errors by coincident signals arriving during the discharge period. The data from all the detectors in a particular branch was read out when a trigger of five or more detectors had signal within 25 μs (Chiba *et al* 1992). The branches were unified in 1995 (Ohoka *et al* 1997), so that coincidences between five detectors between two branches would also trigger the system.

2.8. The High Resolution Fly's Eye (HiRes) detector

The High Resolution Fly's Eye (HiRes) detector was designed using the experience gained in operating the Fly's Eye detector. Using larger mirrors (4 m^2 effective area) and smaller pixels (1°) it was able to increase the usable aperture by a factor of 10 over Fly's Eye. It was deployed on two desert hills separated by 13 km at (40.2° N, 112.8° W) and (40.1° N, 113.0° W), on Dugway Proving Grounds in Utah, USA. HiRes was operated from May 1997 to April 2006.

The two sites allow for stereo observation of events, which gives very good geometrical reconstruction of the showers. HiRes-I was located on the same hill where the full-sky Fly's Eye detector had been located. It operated as a stand-alone site from May 1997 till the end of 1999. It consisted of one ring of mirrors viewing elevation angles from 3° to 17° (see figure 9). The limited elevation coverage limited the aperture for energies below 1×10^{18} eV but not at higher energies. Very energetic events are visible at large distances, where the limited elevation angle covers most of the volume of atmosphere where the shower occurs. At lower energies, the detector cannot see showers as far away, so most of the available atmosphere lies above the portion observed by the detector. This becomes significant for HiRes-I at the

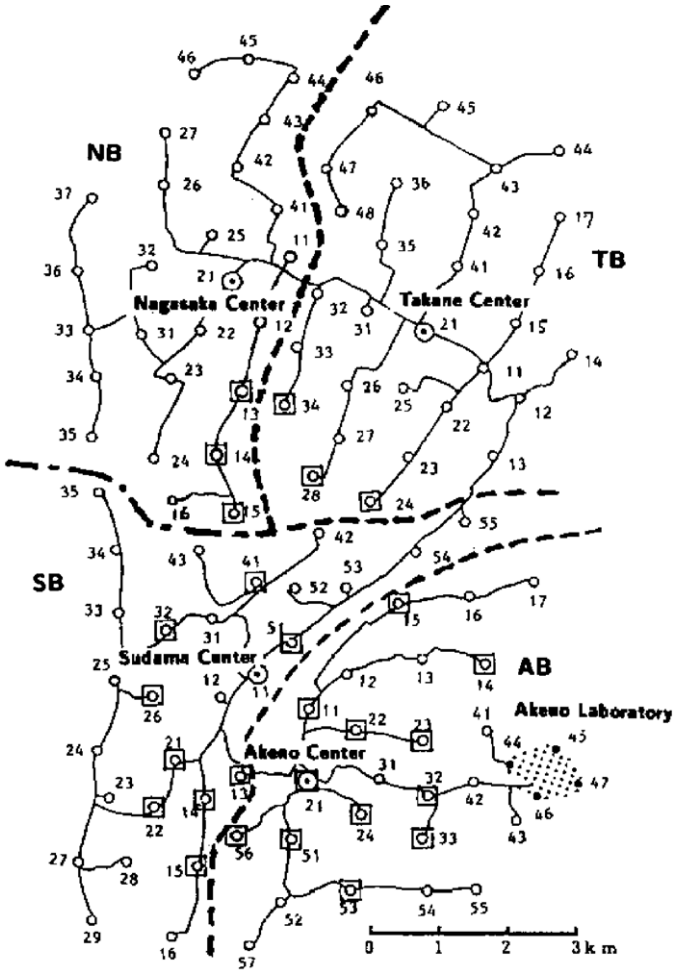

Figure 8. The layout of scintillator detectors in AGASA. Figure taken from Ohoka *et al* (1997). Circles represent the positions of 2.2 m^2 detectors, squares the positions of shielded muon detectors and lines the communication network.

energy given above. HiRes-I recorded the time and pulse height of each tube, and a simple coincidence trigger of three tubes in each of two, adjacent 4×4 clusters of tubes within a mirror, was required to record the event.

The second site, HiRes-II, became operational in December 1999. It consisted of two rings of mirrors viewing elevation angles from 3° to 31° (see figure 9). The larger elevation angle coverage allows a lower threshold of about 2×10^{17} eV. In addition, HiRes-II had a flash ADC (FADC) data acquisition system, which recorded the voltage of each phototube every 100 ns (Boyer *et al* 2002). The sum of the signals in each row and column of tubes was also digitized with an FADC for triggering and recording large signals at low gain. A trigger bit was set whenever three of five consecutive rows or columns were over threshold in coincidence (this coincidence also took into account the signal moving from one row or column to the next in sequence). Two trigger bits, in rows or columns or both, were required to trigger the readout of an event. The FADC system has the feature that light from different tubes can be combined by time.

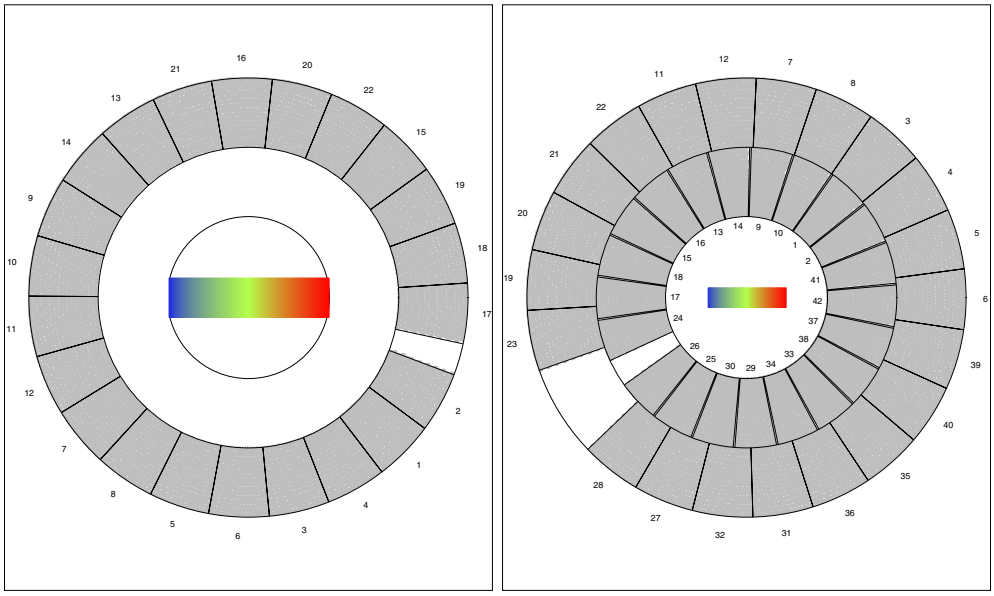

Figure 9. The configuration of the mirrors at HiRes-I (left) and HiRes-II (right). The outer ring covers 3°–17° in elevation, the inner ring 17°–31°. Figure taken from Reil (2002).

With its fine angular resolution, measuring the pointing direction of each PMT becomes important. This measurement can be done by recording the images stars on a screen placed over the face of the PMT cluster (Bergman 2001) or by measuring the change in the noise rate of each individual PMT as stars move in and out of their field of view (Sadowski *et al* 2002).

Because HiRes can look so much further through the atmosphere than Fly's Eye could, it was imperative that there was a mechanism in place to measure the transmission of UV light through the atmosphere. This was accomplished by means of a bistatic LIDAR system (Abbasi *et al* 2006a, 2006c). A steerable laser at each site shoot a planned series of light pulses through the atmosphere, and the light scattered from these pulses was collected by the detector at the other site. These shots look very much like the cosmic rays signals except that they move up through the atmosphere. By measuring the amount of light scattered at various angles one can determine the amount scattering by air molecules (Rayleigh scattering), which is fairly constant, and the amount of scattering by aerosols, which can vary from night to night. The phase function of the scattering by aerosols is also measured. The earliest data collected by HiRes-I were taken before any LIDAR system was in operation. For these data one must rely on average measurements. However, the LIDAR system was in operation for several years, so a good characterization of the atmospheric clarity was obtained. The atmosphere in Dugway is in fact quite clear, given that it is a desert.

2.9. The Pierre Auger Observatory

The Pierre Auger Observatory (Auger Collaboration 2004) is a planned, hybrid detector with sites in both the northern and southern hemispheres. The southern detector is currently being deployed in Malargue, Argentina, at (35° S, 69° W). When completed the surface array component of the detector will consist of 1600 water tanks (10 m^2 in area, 1.2 m deep) on a triangular grid, with a separation of 1.2 km. The total area covered by this array will be

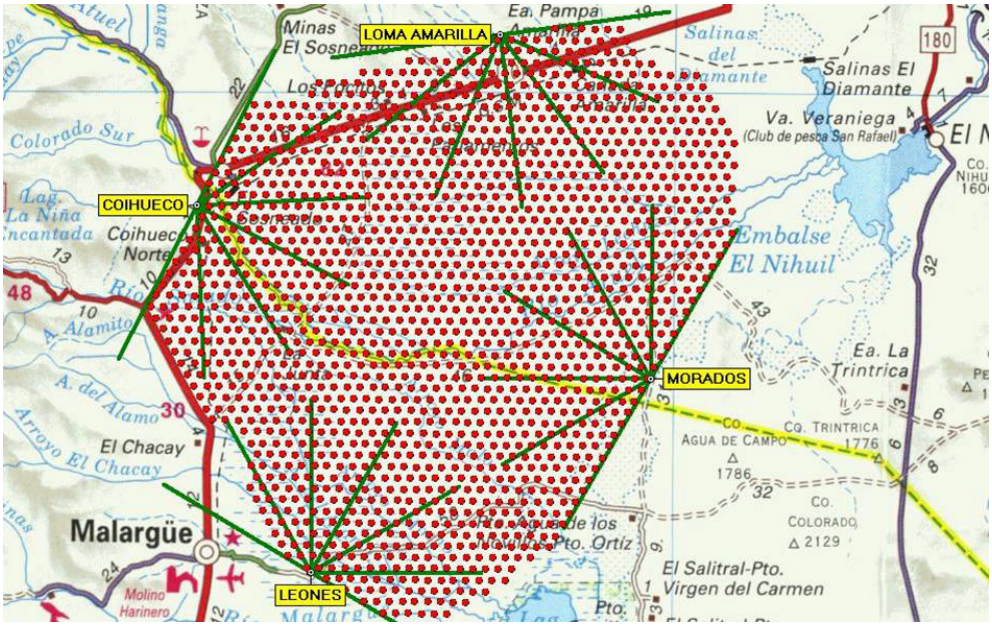

Figure 10. The layout of the Auger surface array detectors (red dots) and the overlooking fluorescence detectors. Figure taken from Mantsch *et al* (2005).

3000 km^2, about 30 times that of AGASA (see figure 10). The surface array will be overlooked by a set of four fluorescence stations placed around the edges of the array and covering its entire area. The fluorescence detectors consist of mirrors viewing 30° in both elevation and azimuth.

3. The UHECR spectrum: flux measurements

When Yoshida and Dai (1998) published their review, there was the suggestion that many experiments saw a hint of the expected GZK Cutoff. That suggestion has become a near certainty with the HiRes claim to have observed the cutoff (Bergman 2007, Abbasi *et al* 2007a). However, the consensus was radically different for several years, when AGASA claimed to have observed as many as 16 events about 100 EeV (Sakaki *et al* 2001). This discrepancy points to the importance of understanding the energy resolution of any experiment when looking for features in the spectrum.

Rather than list the recent measurements of each experiment, we thought it would be more illuminating to discuss in turn each of the features of the flux spectrum, giving the evidence for each feature, and the degree of confidence we have in their existence. The features of the spectrum are best described as changes in the spectral slope. As such, one must understand not only the energy at which the change occurs, but the degree to which we know the spectral slope between features. Because of differences between the normalization of different experiments, it is often difficult to compare the absolute energy of the various features. It is much easier to compare the spectral slopes. The ratios of the energies of the features (the difference in the logarithm of the energies) are also of prime importance as it does not depend on the absolute energy scale. This gives special importance to experiments which can measure the energy of more than one feature.

Table 1. The measured slope parameters and break point energies for the Second Knee.

Experiment	Slope below	Break point $\log_{10}\left(\frac{E}{\text{eV}}\right)$	Slope above
Akeno (Nagano *et al* 1992)	3.02 ± 0.03	17.8	3.24 ± 0.18
Fly's Eye (Bird *et al* 1993)	3.01 ± 0.06	17.6	3.27 ± 0.02
Haverah Park (Ave *et al* 2003a)			3.33 ± 0.04
HiRes (Abbasi *et al* 2007a)			3.32 ± 0.03
Average slopes	3.02 ± 0.03		3.29 ± 0.02

We discuss three features of the ultra-high energy cosmic ray spectrum: the Second Knee, the Ankle and the GZK Cutoff. The Second Knee is a softening of the spectrum (the spectral slope becoming steeper) in the 10^{17} eV decade. The Ankle is a hardening of the spectrum (the spectral slope becoming less steep) in the 10^{18} eV decade. The GZK Cutoff is an expected, drastic reduction in the UHECR flux above $\sim 10^{19.8}$ eV due to photopion production of protons on the cosmic microwave background radiation.

3.1. The Second Knee

Four experiments have shown evidence for the Second Knee: Akeno (Nagano *et al* 1992), Haverah Park (Lawrence *et al* 1991, Ave *et al* 2003a), Fly's Eye (Bird *et al* 1993, Bird *et al* 1994) and the HiRes Prototype/MIA experiment (Abu-Zayyad *et al* 2001). This includes experiments which have measured a *break* in the spectrum in the 10^{17} eV decade, having measured different spectral slopes above and below the break energy. Of these, its unclear whether one should still include the Haverah Park data below the break because of the recent reanalysis (Ave *et al* 2003a). The reanalysis only measures the spectral slope above the break and does not treat data from the infill array used to measure the spectral slope below the break at the lowest energies. It is also difficult to consider the HiRes Prototype/MIA measurement, as there is not a separate measurement of the spectral slopes above and below the Second Knee, and the measurement of the slope below the Second Knee has no associated error.

In addition to the experiments listed above, Yakutsk (Egorova *et al* 2004) has data above the Second Knee (but provides no fit to the spectral slope), while HiRes (Abbasi *et al* 2007a) measures the spectral slope above the Second Knee in its monocular analysis (as does Haverah Park in its recent analysis). These measurements can be used to increase confidence in the knowledge of the spectral slope above the Second Knee. (The latest HiRes measurement in Abbasi *et al* (2007a) includes and supersedes the data shown in several earlier papers (Abbasi *et al* 2004a, 2005a, 2005b). A detailed account of the HiRes systematic uncertainties can be found in (Abbasi *et al* 2007c).)

The measured spectral slopes and break point energies are shown in table 1. There is good agreement between all the experimental results for the spectral slopes, giving very high confidence in the existence of a break, with the difference between spectral having a significance of 7.5σ.

We have performed our own fits to these experimental results. The fits use the binned maximum likelihood method comparing the actual number of events in bin to the expected number given a flux and exposure (Yao *et al* 2006). The likelihood in this method is normalized in such a way as to give a goodness-of-fit metric which approaches the χ^2 distribution in the large statistics limit. The fitting function is a broken power law, where the break points are

Table 2. Our broken power law fits to spectrum measurements in the Second Knee energy range. The fit parameters include a normalization (not shown), slope parameters above and below the break and the break point energy for the Second Knee.

Experiment (reference)	χ^2/DOF	Slope below	Break point $\log_{10}\left(\frac{E}{\text{eV}}\right)$	Slope above
Akeno (Nagano *et al* 1992)	8.3/13	3.04 ± 0.02	17.8 ± 0.2	3.25 ± 0.12
Fly's Eye (Bird *et al* 1993)	13.7/18	3.04 ± 0.05	17.60 ± 0.06	3.27 ± 0.02
HiRes/MIA (Abu-Zayyad *et al* 2001)	2.5/5	3.02	17.6 ± 0.2	3.23 ± 0.14
Haverah Park (Ave *et al* 2003a)	1.4/5			3.32 ± 0.05
Yakutsk T-500 (Egorova *et al* 2004)	45.2/15			3.213 ± 0.012
HiRes (Abbasi *et al* 2007a)	8.55/15			3.26 ± 0.02
Global fit (at Fly's Eye E scale)	109.4/93	3.02 ± 0.01	17.52 ± 0.02	3.235 ± 0.008

allowed to vary as some of the fit parameters. The fitting function is continuous, with only one normalization parameter.

The results from our fits to the experiments which make measurements in the Second Knee energy range are shown in table 2. A few specific comments on the fits are important. The fit to the HiRes Prototype/MIA spectrum required special handling. Abu-Zayyad *et al* (2001) mention that a fit to the first six points of the spectrum give a spectral slope of 3.01. We performed a fit to the same points and found a spectral slope of 3.02 ± 0.11. We then *fixed* this slope and performed a broken power law fit to all the points. In fitting the Haverah Park spectrum, we only fit the lowest seven points, to avoid being biased by the Ankle. For the fits to the Yakutsk T-500 (Egorova *et al* 2004) and HiRes-II Monocular (Abbasi *et al* 2005b, Abbasi *et al* 2007a) spectra, the fit included another break point and spectral index for the Ankle and beyond. These parameters will be discussed in section 3.2. The fit to the HiRes-II spectrum does not include the points below $10^{17.5}$ eV, which agrees with the way the HiRes Collaboration performs its own fits. These spectra along with the results of our fits are shown on the left-hand side of figure 11. The fit power law is only drawn in the energy range that was used in the fitting.

The one recent measurement which does not agree *in shape* with the other measurements in the Second Knee region is that of the Čerenkov array at Yakutsk (Ivanov *et al* 2003, Knurenko *et al* 2007). The Second Knee in this analysis appears at a very high energy, $\sim 10^{18.2}$ eV, which is inconsistent with the standard Yakutsk ground array measurement (Egorova *et al* 2004). We have not fit these data, but it is shown in figure 11.

To combine all the data into a global fit, we scaled each experiment so that the fit fluxes all agree at 10^{18} eV. This may be equivalent to removing several degrees of freedom from the combined fit, but does not necessarily give the best χ^2. We have not adjusted the number of degrees of freedom to account for the shifting. The global Second Knee Fit included all the data above, except the Yakutsk measurement using the Čerenkov Array. The fit does include most points left out of the individual fits, but does no include the highest energy bin of the HiRes-II monocular sample which may be above the GZK Cutoff. The results of the fit

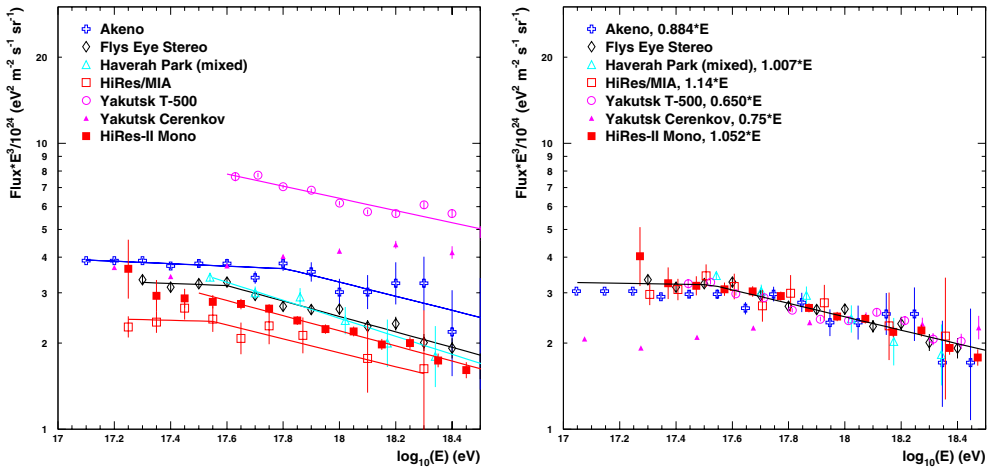

Figure 11. Left: flux measurements in the Second Knee energy range. The shown fits are our calculation. Right: flux measurements in the Second Knee energy range, scaled so that the flux agrees with the Fly's Eye result at 10^{18} eV. The scaled data points were fit to a broken power law spectrum in a global fit, with the result shown.

are also included in table 2. The scaled flux measurements and the result of the global fit are shown on the right-hand side of figure 11.

One can see that there is, in fact, a consensus among the various experiments on the existence of the Second Knee and on the specific spectral slopes at energies above and below the Second Knee. The only major discrepancies are in the energy normalization between the experiments.

3.2. The Ankle

Before considering all the experiments which have observed the Ankle, we note that only one experiment has observed both the Second Knee and the Ankle within a single measurement: Fly's Eye. Observing both features allows a measurement of the ratio of the energies of the two features (or the difference in the logarithm of the energies). The Fly's Eye Collaboration fits their stereo spectrum to a broken power law, with breaks at $10^{17.6}$ eV for the Second Knee and $10^{18.5}$ eV for the Ankle (Bird *et al* 1993, Bird *et al* 1994). These break points were picked by eye and not allowed to float in the fit. In our fit to the Fly's Eye stereo spectrum (see tables 2 and 4), we find the break points at ($\log E$, E in eV) 17.60 ± 0.06 and 18.52 ± 0.09. The difference, $\Delta \log E = 0.92 \pm 0.11$, corresponds to a ratio $E_2/E_1 = 8.3^{+2.5}_{-1.8}$.

In addition to the Fly's Eye experiment, two other experiments have strong observations of the Ankle: Yakutsk (Egorova *et al* 2004) and HiRes (Abbasi *et al* 2005b, 2007a). Of these, only HiRes provides measurements of the spectral slopes. Haverah Park observed the Ankle in its old analysis (Lawrence *et al* 1991), but only the section between the Second Knee and the Ankle was treated in their recent reanalysis (Ave *et al* 2003a). We use only the results from the later analysis, which does provide a spectral slope for energies below the Ankle. AGASA has also observed a break in their spectrum (Takeda *et al* 2003), but at a very high energy, higher than one would expect from other experiments given the overall flux level. The AGASA Collaboration has distanced itself from interpreting the break as a measurement of the Ankle, because the break also coincides with the energy at which the AGASA detector becomes fully efficient (Teshima 2002). We only use AGASA flux measurements above

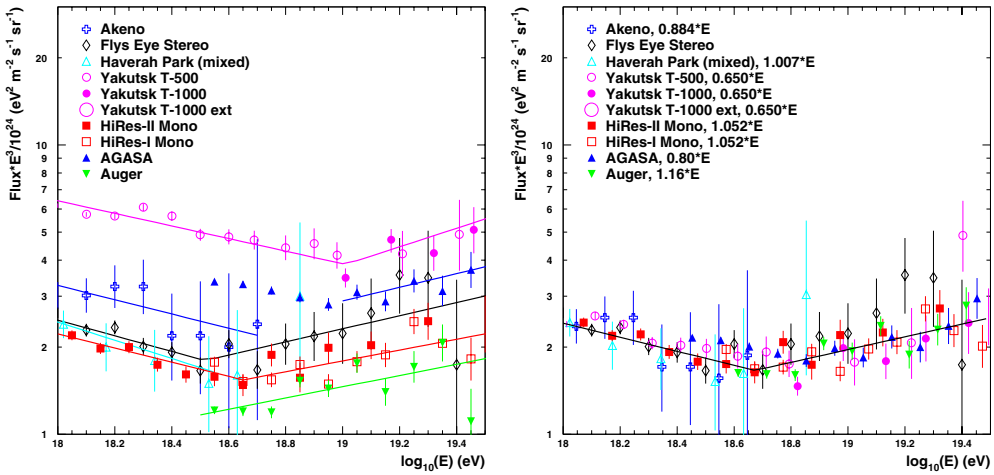

Figure 12. Left: flux measurements in the Ankle energy range. The shown fits are our calculation. Right: flux measurements in the Ankle energy range, scaled so that the flux agrees with the Fly's Eye result at 10^{18} eV (or 10^{19} eV for those experiments with no measurement at 10^{18} eV). The scaled data points were fit to a broken power law spectrum in a global fit, with the result shown.

Table 3. The measured slope parameters and break point energies for the Ankle.

Experiment	Slope below	Break point $\log_{10}\left(\frac{E}{\text{eV}}\right)$	Slope above
Fly's Eye (Bird *et al* 1993)	3.27 ± 0.02	18.5	2.71 ± 0.10
HiRes (Abbasi *et al* 2007a)	3.25 ± 0.01	18.65 ± 0.04	2.81 ± 0.03
Haverah Park (Ave *et al* 2003a)	3.33 ± 0.04		
Average slopes	3.26 ± 0.01		2.80 ± 0.03

10^{19} eV in our fits. The Pierre Auger Observatory also measures the flux above the Ankle (Pierre Auger Collaboration 2005), but does not provide a measurement of the spectral index.

The measured spectral slopes and break points from these experiments are shown in table 3.

As we did for the Second Knee, we have performed our own fits to all the reported spectra. Some of these fits are identical to those listed previously, and the fitting methodology is the same in all cases. Two important differences are that the spectra from both HiRes-I and HiRes-II in monocular mode have been fit in a combined fit, as has the data from all three triggers of the Yakutsk experiment. No scaling was done to match the different measurements from these experiments as the data are presumed to be already matched in energy scale. In both these fits, as well as in fits to AGASA and Auger spectra, measurements above $10^{19.8}$ eV have been excluded, so as not to complicate the fits with data above the GZK Cutoff. In addition, HiRes-II data were only fit above $10^{17.5}$ eV, and AGASA data were only fit above 10^{19} eV. The results of all these fits are displayed on the left-hand side of figure 12. In the figure, on can clearly see that there is a good consensus on the spectral slopes above and below the Ankle energy. The parameters found by the fits are given in table 4.

We again combined all the data into a global fit by scaling energies. The scalings were identical to those used in section 3.1, where there was overlap. Data from HiRes-I were

Table 4. Our broken power law fits to spectrum measurements in the Ankle energy range. The fit parameters include a normalization (not shown), slope parameters above and below the break and the break point energy for the Ankle.

Experiment (reference)	χ^2/DOF	Slope below	Break point $\log_{10}\left(\frac{E}{\text{eV}}\right)$	Slope above
Akeno (Nagano *et al* 1992)	8.3/13	3.25 ± 0.12		
Fly's Eye (Bird *et al* 1993)	13.7/18	3.27 ± 0.02		
Haverah Park (Ave *et al* 2003a)	1.4/5	3.32 ± 0.05		
Yakutsk (Egorova *et al* 2004)	50.3/22	3.22 ± 0.01	19.01 ± 0.07	2.68 ± 0.06
HiRes (Abbasi *et al* 2007a)	29.6/28	3.22 ± 0.03	18.65 ± 0.04	2.81 ± 0.03
AGASA (Takeda *et al* 2003)	6.7/6			2.76 ± 0.08
Auger (Pierre Auger Collaboration 2005)	28.9/11			2.81 ± 0.03
Global fit (at Fly's Eye E scale)	184.2/125	3.242 ± 0.008	18.70 ± 0.02	2.78 ± 0.02

scaled at the same level as data from HiRes-II, since these two data sets are already matched in energy. The same is true of the three Yakutsk data sets. The data from AGASA and Auger are more problematic since they have no connection to a measured flux at 10^{18} eV. For these two experiments, we required that the flux at 10^{19} eV matches the energy scale of the HiRes experiment. This matching was only required to be within ∼5%. While a very tight requirement on match flux could be made at 10^{18} eV, the different position of the fit Ankle break points gives a range of fit flux values at 10^{19} eV. Thus, only a loose matching was made for the AGASA and Auger data. After scaling, the energy of the Ankle in our fits ranges from $10^{18.5}$ eV for Fly's Eye to $10^{18.8}$ eV for Yakutsk. The results of the combined fit to the data (which include a floating break point for the Second Knee) are shown on the right-hand side of figure 12, with the fitted parameters given in table 4.

3.3. The GZK Cutoff

Four experiments currently claim to have significant aperture above the expected energy of the GZK Cutoff of 6×10^{19} eV (Greisen 1966, Zatsepin and Kuz'min 1966): AGASA (Takeda *et al* 2003), Yakutsk (Egorova *et al* 2004), HiRes (Abbasi *et al* 2005b, 2007a) and Auger (Pierre Auger Collaboration 2005). Of these AGASA claims to see a continuation of the spectrum above the GZK Cutoff with no change in spectral index. Yakutsk and HiRes on the other hand claim that their spectra are consistent with the presence of the cutoff, and HiRes has claimed an observation of the feature. Auger in its preliminary spectrum makes no claim in either direction. The Haverah Park experiment had previously claimed to see four events above 10^{20} eV (Lawrence *et al* 1991), but upon reanalysis all four of these events were found to be less than 10^{20} eV (Ave *et al* 2003a). (The average energy of the four events is $10^{19.88}$ eV in the reanalysis. This point is displayed in figures 13 and 14, but is not included in any fits.)

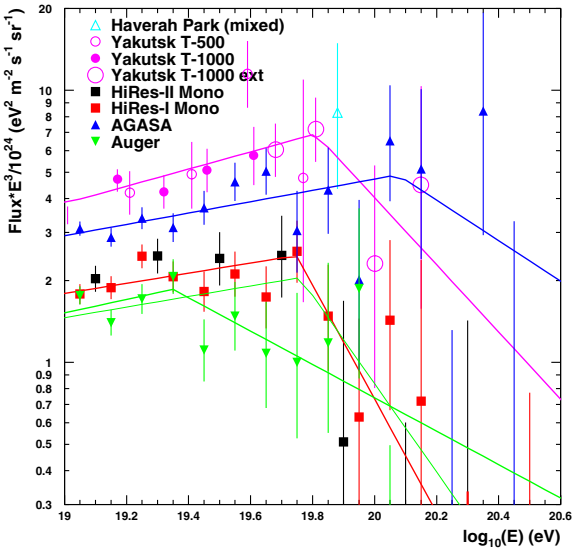

Figure 13. Flux measurements in the GZK Cutoff energy range. The shown fits are our calculation.

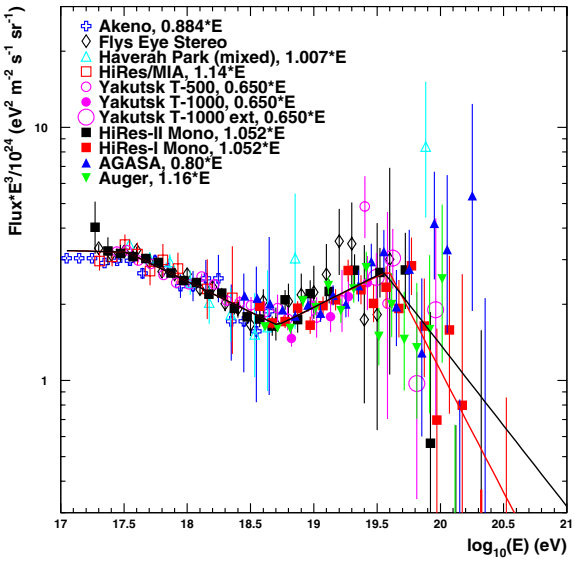

Figure 14. All the flux measurements discussed in this review fit to a broken power law. The black line shows the result of the fit including the AGASA data, while the red line shows the fits without including the AGASA data.

As we have for the other features, we fit each of these results with our own fit. In these fits, the empty bins with finite exposure play a very important role. The binned maximum likelihood method allows one to include these bins in the fit. We first fit each experiment with no allowed break point for the GZK Cutoff (but including one for the Ankle in the case of Yakutsk and HiRes). We then fit allowing a floating break point for the GZK Cutoff. In all

Table 5. Our broken power law fits to spectrum measurements in the GZK Cutoff energy range. The fit parameters include a normalization (not shown), slope parameters above and below the break and the break point energy for the GZK Cutoff. The first fit for each experiment shows the slope and χ^2 for a fit with no allowed GZK Cutoff break. A local χ^2 minimum was found fitting the Auger data with a break point near the HiRes value.

Experiment (reference)	χ^2/DOF	Slope below	Break point $\log_{10}(\frac{E}{\mathrm{eV}})$	Slope above
Yakutsk	55.6/24	2.73 ± 0.06		
(Egorova *et al* 2004)	51.6/22	2.68 ± 0.06	19.81 ± 0.10	4.2 ± 0.9
HiRes	64.3/37	2.88 ± 0.03		
(Abbasi *et al* 2007a)	34.6/35	2.81 ± 0.03	19.75 ± 0.04	5.1 ± 0.7
AGASA	16.1/11	2.81 ± 0.07		
(Takeda *et al* 2003)	15.2/9	2.79 ± 0.07	20.1 ± 0.4	3.7 ± 2.0
Auger	34.8/12	2.82 ± 0.03		
(local minimum)	31.4/10	2.81 ± 0.03	19.8 ± 0.2	4.6 ± 2.5
(Pierre Auger Collaboration 2005)	21.1/10	2.76 ± 0.03	19.35 ± 0.07	3.6 ± 0.3

cases we found a break, but the significances vary. The results of these fits are presented in table 5 and figure 13 (only for the fits with the floating GZK Cutoff break point). No account was taken of overlapping exposure in different measurements for either the Yakutsk or HiRes spectra.

It is clear that HiRes sees a break with a very large significance (the χ^2 falling from 64.3 to 34.6 while adding only two degrees of freedom), whereas none of the other breaks are nearly as significant. The fact that any break point is found in the AGASA data is interesting and results largely from the empty bins at $10^{20.25}$ eV and $10^{20.45}$ eV where one still expects a significant number of events given the AGASA exposure. The local minimum found in fits to the Auger spectrum points, perhaps, to the fact that fitting this region with a broken power law is reaching its limits. Both Auger and HiRes-I spectra seem to indicate a rounded cutoff. This could be from resolution, however, as HiRes-II seems to see a very sharp cutoff.

We end the section on flux measurements with a comprehensive fit to all the data presented above (with the noted exceptions). The fit has three floating break points, four spectral slopes and one normalization. The parameters are listed in table 6. As we did above for the individual fits in the GZK Cutoff energy range, we first fit without a break point for the GZK Cutoff, then with it. The improvement in χ^2 tells us about the significance of the added break point. We also did these fits with and without the AGASA data, data which seem to be at odds with the other experiments in this energy range. Both these fits are shown in figure 14. In either case χ^2 drops by about 41 while adding two degrees of freedom. This corresponds to a bit more than 6σ significance. However, while the preponderance of evidence points to the existence of the GZK Cutoff, the calculated significances are somewhat suspect because of the ad hoc nature of the scalings involved.

A final measure of the significance of the break at the GZK Cutoff energy can be found by comparing observed and expected numbers of events above $10^{19.8}$ in fits with and without the GZK Cutoff break point. With the scaled AGASA data included, one finds 42 events with energies above $10^{19.8}$ eV, where one would expect 85 from a fit with no high energy break point. The expected number drops to 45.4 when one adds the break point. The statistical probability of expecting 85 events and observing 42 or less is 1.8×10^{-7} or just over 5σ. If one removes the AGASA data, one observes 28 events, while expecting 67 with no high

Table 6. Our broken power law fits to all the scaled spectrum measurements. J_{18} is the fit flux at an energy of 10^{18} eV in units of (eV m^2 s sr)$^{-1}$. The break point energies are all in units of eV.

	With AGASA		Without AGASA	
Parameter	No GZK BP	GZK BP	No GZK BP	GZk BP
χ^2/DOF	264.5/156	223.0/154	235.1/141	193.7/139
$J_{18}/10^{24}$	2.44 ± 0.01	2.43 ± 0.01	2.44 ± 0.01	2.43 ± 0.01
γ_1	2.99 ± 0.01	2.99 ± 0.01	2.99 ± 0.01	2.99 ± 0.01
$\log_{10} E_1$	17.52 ± 0.02	17.52 ± 0.02	17.52 ± 0.02	17.52 ± 0.02
γ_2	3.228 ± 0.007	3.230 ± 0.007	3.228 ± 0.007	3.230 ± 0.007
$\log_{10} E_2$	18.66 ± 0.03	18.67 ± 0.02	18.66 ± 0.03	18.69 ± 0.02
γ_3	2.83 ± 0.02	2.74 ± 0.03	2.85 ± 0.02	2.77 ± 0.03
$\log_{10} E_3$		19.56 ± 0.05		19.60 ± 0.05
γ_4		3.64 ± 0.19		4.00 ± 0.26
$N^{\rm obs}_{lE>19.8}$	42	42	28	28
$N^{\rm pred}_{lE>19.8}$	85.0	45.4	67.0	30.1

energy break point or 30.1 with it. The probability of observing 28 events when expecting 67 is 6×10^{-8} or 5.3σ.

In conclusion, there is strong evidence for each of the three acknowledged features in the UHECR spectrum: the Second Knee, the Ankle and the GZK Cutoff. The most compelling evidence for the first two features comes from the large degree of consensus on what the spectral indices are at energies above and below each feature. And where the spectral index changes there *must* be a feature. The exact energy of either feature is not so well in agreement, especially for the Ankle. This comes primarily from the difficulty in determining the absolute calorimetric scale of any experiment. In addition, even after trying to adjust the scales of the different experiments to get identical fluxes at some energy (10^{18} eV in our case), the energy of Ankle still varies by a factor of 2. This may point to the Ankle not being a particularly sharp break.

At this point, only the HiRes experiment has presented compelling evidence of the GZK Cutoff. Moreover, they observe *both* the Ankle and the GZK Cutoff and so can measure the ratio of the two energies ($10^{19.75-18.65} = 12.6$). While Yakutsk and Auger do not have enough data to claim an observation of the Cutoff yet, their data do add to the significance of the observation by HiRes.

4. The UHECR spectrum: composition measurements

Unlike the situation in measuring the flux of cosmic rays, there is little consensus among experiments on the primary composition of cosmic rays. Some of the disagreements, no doubt, come from the difficulty of the measurements and their indirect nature. In order to have an aperture of a large enough area, one is forced to study extensive air showers (EAS), rather than the primary cosmic rays themselves. For measurements of the energy (and therefore the flux) there is a strong correlation between the energy deposited in the atmosphere by the EAS and the kinetic energy of the primary cosmic ray. It is only in the shape of the shower that one can try to get at the composition of the primary cosmic rays.

EASs grow as the primary cosmic ray collides with an atom in the atmosphere producing many secondary particles which divide the energy of the primary between them. The secondaries will also collide with atoms in the atmosphere and form a cascade of more

and more particles (with less and less energy) as the shower develops. When the particles of a shower become too low in energy, they no longer contribute to the growth of the shower, but are rather absorbed by the atmosphere. In this way, the shower grows geometrically for a certain distance through the atmosphere but the growth first begins to slow, and eventually the shower will begin to shrink. The depth in the atmosphere where the shower is at its largest is usually called $X_{\max}$. $X_{\max}$ will depend logarithmically on the energy because that energy can be spread among more particles. The change in the average $X_{\max}$ with $\log E$ is called the elongation rate.

Nuclei heavier than hydrogen (i.e. more than one nucleon) develop, to first order, as if each constituent nucleon created its own shower. This means that one would expect iron with 56 nucleons to generate a shower with a similar $X_{\max}$ to that created by a proton 56 times less energetic. Unfortunately, EAS development is not so simple and there are large fluctuations in $X_{\max}$ from shower-to-shower for a given type of primary. These fluctuations should be smaller for heavier primaries because one is essentially averaging over many showers. At any given energy, however, the distributions of $X_{\max}$ for proton and iron showers overlap by enough that one cannot expect to determine the primary particle type on an event-by-event basis.

Still the average $X_{\max}$ at a given energy and the width of the distribution should tell one something about the average composition of UHECRs at that energy. Unfortunately, there is no consensus on the absolute value of the average $X_{\max}$ to expect for a given primary type: different models give different results at the level of 20% of the difference between protons and iron. Two recent interaction models used extensively are the QGSJet model (Kalmykov *et al* 1997) and Sibyll (Fletcher *et al* 1994). In either case one must also model the development of the shower, using programs such as Corsika (Heck *et al* 1998) or Aires (Sciutto 1999). The models agree better in the determination of the elongation rate, which is expected to be about 60 g cm^{-2}/decade for showers generated by either iron primaries or proton primaries. One would also expect the width of the $X_{\max}$ distribution to be robustly determined by the models, but we have not seen any predictions of this (and only one published experimental measurement (Cassiday *et al* 1990)).

Only fluorescence detectors can determine $X_{\max}$ directly for each shower, though Yakutsk, with its Čerenkov detectors, can make a measurement as well. Four experiments have made measurements of the average $X_{\max}$ as a function of energy: Fly's Eye (Cassiday *et al* 1990, Bird *et al* 1993), Yakutsk (Egorova *et al* 2001), HiRes/MIA (Abu-Zayyad *et al* 2000a, Abu-Zayyad *et al* 2001) and HiRes (in stereo) (Abbasi *et al* 2005d). These measurements are shown in figure 15. While the figure includes lines for the QGSJet expectation for iron and proton primaries (the proton line is higher), one should pay less attention to the absolute position of the lines than to their slope and separation.

We have fit the data of each experiment (and for HiRes/MIA and HiRes stereo combined) for the elongation rate, allowing for possibly one break point in the fit. The results are shown in table 7. These fits are simple χ^2 fits where the break point (if any) is one of the fit parameters and the fit line is required to be continuous in that case. We allow the floating break point for fitting the Fly's Eye data following the fit shown in Bird *et al* (1993) and for the combined HiRes/MIA and HiRes stereo data following what is shown in Abbasi *et al* (2005d), though they did not in fact perform a combined fit.

As we stated at the beginning of this section, there is little consensus in either the average $X_{\max}$ values or in the elongation rates. The Fly's Eye and Yakutsk data seem to indicate a slowly changing composition becoming lighter because the elongation rate is greater than the expected value of $\sim$60 g cm^{-2}/decade for constant composition. The change occurs from $10^{17.5}$ eV to over least 10^{19} eV. They also agree with each other well in the actual average $X_{\max}$ values measured. These results, especially those from Fly's Eye which had wide visibility,

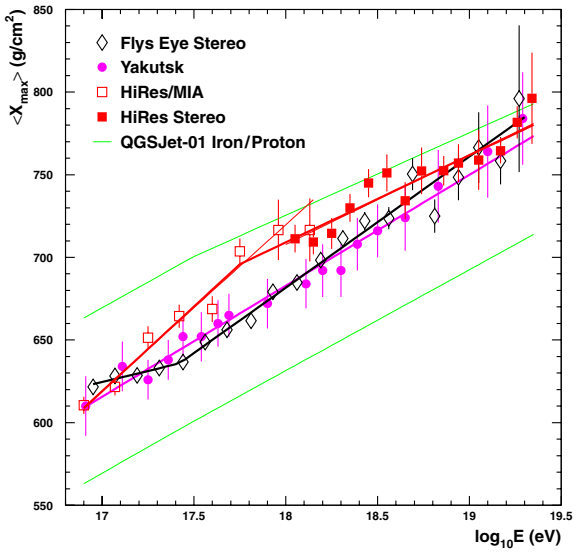

Figure 15. Measurements of the average $X_{\max}$ as a function of energy.

Table 7. Our fits of the elongation rates (and break points) to experimental data of average $X_{\max}$ versus log E. The fit elongation rates have units of g cm^{-2}/decade.

Experiment (reference)	χ^2/DOF	Elongation rate	Break point $\log_{10}\left(\frac{E}{\text{eV}}\right)$	Elongation rate
Yakutsk (Egorova *et al* 2001)	2.6/17			67 ± 6
Fly's Eye (Bird *et al* 1993)	16.6/16	27 ± 14	17.42 ± 0.08	79 ± 3
HiRes/MIA (Abu-Zayyad *et al* 2001)	6.0/6			100 ± 8
HiRes stereo (Abbasi *et al* 2005d)	7.4/12			54 ± 6
HiRes/MIA, HiRes combined fit	12.9/18	103 ± 7	17.75 ± 0.10	53 ± 5

were used to argue for a changing composition in conjunction with the changing spectral slope of the Ankle.

The HiRes Prototype/MIA experiment and the HiRes stereo measurement, however, tell a different story. HiRes/MIA measures a more quickly changing composition, again getting lighter, but one that is complete by an energy of 10^{18} eV. The HiRes stereo measurement confirms this by measuring a constant, light composition above 10^{18} eV. This result supports the explanation of the Ankle as being the result of energy losses of extragalactic protons due to electron-pair creation (Berezinsky and Grigor'eva 1988, Berezinsky *et al* 2006).

It is important to keep in mind that the Fly's Eye measurement quotes a systematic uncertainty in $X_{\max}$ of 20 g cm^{-2} (Cassiday *et al* 1990), so that the difference in the measured average $X_{\max}$ values between the experiments in not inconsistent. In addition, the HiRes Prototype and HiRes proper had much finer angular resolution than Fly's Eye, which translates

into finer measurements of shower profiles. Thus, while the statistical uncertainty in the HiRes/MIA measurements is larger than that of Fly's Eye, the systematic error may be smaller.

Trying to measure the composition with a surface detector is more indirect than with fluorescence, and more dependent or shower modelling. In the Haverah Park experiment, the rise time of the shower front as a function of distance from the shower core was used to estimate the composition (Ave *et al* 2003b). One expects that the closer the shower maximum is to the ground (i.e. that larger the value of $X_{\max}$) the longer will be the rise time of the shower front. This is just the result of the range of pathlengths involved. Haverah Park finds a constant, fairly heavy composition between $10^{17.2}$ eV and 10^{18} eV.

In the AGASA experiment, the ratio of number of muons to the number of electrons at 600 m from the shower core was used to estimate the composition (Shinozaki and Teshima 2004, Hayashida *et al* 1995). They find that the composition is getting lighter between $10^{17.5}$ eV and 10^{19} eV, but at a slower rate than observed by Fly's Eye.

The HiRes Prototype/MIA experiment performed a similar measurement to that of AGASA: in addition to measuring the average $X_{\max}$, they also measured the density of muons at 600 m from the shower core. They found a composition which was getting lighter between 10^{17} eV and 10^{18} eV agreeing with the trend in the $X_{\max}$ measurement, but compared to predictions of QGSJet they found a composition heavier than iron even at 10^{18} eV (the situation is even worse using Sibyll). This points to a problem with the current interaction models: they do not produce enough muons. This problem may have been corrected in the new interaction code, EPOS (Pierog and Warner 2006).

5. The UHECR spectrum and composition: synthesis

As one looks at the various measurements of the UHECR flux, one gets the impression that we are beginning to understand the spectrum of cosmic rays at these energies. Even if the experimental measurements do not agree in absolute normalization, there is a consensus on the existence of specific breaks in the spectrum and what the spectral indices are on either side of the breaks. This impression is, however, reversed the moment one begins to consider composition measurements. Not only is there no agreement on where showers of a given energy develop, on average, in the atmosphere, but there is no agreement on the trend. Is the composition of UHECRs getting lighter or staying the same? Experiments disagree.

We have not discussed theoretical or phenomenological models for the acceleration and propagation of UHECRs in this review. They are, of course, crucial for understanding the flux of cosmic rays, but even the best model has trouble dealing with contradictory data. This is especially evident in models of the Ankle region, where two different paradigms for the origin of the Ankle continue to coexist because of the different elongation rate measurements.

To allow for more robust comparisons between data and models one must obtain data those provide more constraints on the models. In measurements of the flux of cosmic rays this means that one would like to have a measurement of two or even all three of the spectrum features within one experiment. This is a difficult undertaking, as the features are roughly each an order-of-magnitude apart in energy, and measuring all three features requires a useful dynamic range of over two and a half orders-of-magnitude. The Telescope Array (TA) experiment (TA Collaboration 2005), which is currently being deployed in Delta, UT, USA, is one experiment which may be able to cover this whole energy range if the TA Low Energy (TALE) extension is funded. In any case, it is important to realize that we must understand more than just the GZK Cutoff.

In measurements of composition, having more constraints also means having a large dynamic range in energy, of course, but it also could mean having more ways, orthogonal ways, of determining the primary particle type. The large hybrid experiments, Auger and TA, will hopefully provide this, by measuring both longitudinal ($X_{\max}$) and lateral distributions for each shower.

6. Anisotropy measurements of UHECRs

While a great deal may be inferred about the sources of UHECR from spectral and composition studies, the evidence thereby obtained is indirect. A complete picture would ideally combine an acceleration model or models with a consistent primary particle arrival-direction distribution.

As we discuss below, the detection of anisotropies in UHECR event arrival directions is a challenging enterprise, which to date has yielded results which are ambiguous at best. But even a completely isotropic cosmic ray sky has something to tell us about source models. For example, at energies below a few $\times 10^{19}$ eV isotropy can be understood as the effect of diffusion by galactic magnetic fields, provided the primary particles are not neutral. At higher energies for which magnetic bending effects are negligible, isotropy indicates that sources are abundant and distributed throughout the sky.

We begin our discussion of anisotropy studies by reviewing the arrival-direction resolutions of the experiments whose data we consider. We will look at searches for both large-scale features, including full-sky dipole moments and small-scale or point-like excesses of UHECR. We conclude with some general comments about lessons learned in anisotropy studies over the past decade.

6.1. Experimental arrival-direction resolutions

It is appropriate that we summarize the arrival-direction resolutions of the various experiments we will consider in this review, namely AGASA, HiRes, Auger, SUGAR and Yakutsk. By 'resolution' we generally mean the size of the opening angle which would contain 68% of events arriving from a point source of cosmic rays. However in some detectors, in particular the monocular fluorescence detectors covered in this review, arrival-direction uncertainties cannot easily be expressed with a single number.

The arrival-direction uncertainties of the AGASA surface detector (SD) array are 'circular' (uniform in two angular dimensions), and decrease with energy, as shown in figure 16 (Takeda *et al* 1999). In this review, we consider AGASA data including events with energies as low as 10^{17} eV. In the literature, angular resolutions of 3° are reported at 10^{18} eV (Hayashida *et al* 1999)) but comparison with figure 16 suggests that a different definition of resolution is being employed.

For the HiRes air fluorescence observatory, there are large differences in angular resolution in monocular and stereo detection modes. Monocular reconstruction by air fluorescence is subject to asymmetric arrival-direction uncertainties. Consider figure 16 (right). The angle of the shower-detector plane (SDP) is generally well reconstructed, for the HiRes-I monocular detector uncertainties in the SDP are parametrized as

$$\sigma_{\mathrm{SDP}} = 88.2^\circ \, \mathrm{e}^{-\Delta\chi/1.9595} + 0.37^\circ \tag{1}$$

where $\Delta\chi$ is the angular tracklength of the shower in degrees. The angle of the track within the SDP, Ψ, is less well reconstructed and is parametrized by

$$\sigma_{\Psi} = 18.4^\circ \, \mathrm{e}^{-\log_{10}(E)/0.690\,85} + 4.1^\circ \tag{2}$$

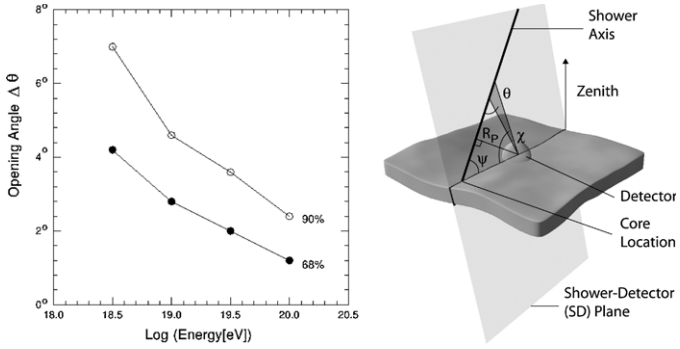

Figure 16. Left: arrival-direction uncertainty for the AGASA experiment. Closed and open circles are the opening angles encompassing 68% and 90% of the data, respectively. Source: Takeda *et al* (1999). Right: the geometry of reconstruction for monocular air fluorescence detector. R_P is the shower's distance of closest approach, Ψ is the angle of the shower within the shower-detector plane (uncertainty parametrized in equation (2)) and χ is the angular tracklength of the shower.

where the energy E is expressed in EeV (10^{18} eV). The parameterizations of equations (1) and (2) are carried out in Stokes (2005).

The HiRes observatory operating in stereo derives its arrival-direction reconstruction from the intersection of two well-measured SDPs. Unlike the case for surface detectors, where higher energy events generally have better hit statistics and hence improved angular resolution, fluorescence observatories tend to see angular resolutions worsen with energy because at higher energies showers are on average farther away. Abbasi *et al* (2004d) report 68% opening angles of 0.57°, 0.61° and 0.69° for showers observed at 10^{19} eV, 4×10^{19} eV and 10^{20} eV, respectively.

Bonifazi *et al* (2005) describe the angular resolution of the Pierre Auger Observatory, for both surface detector (SD) and hybrid (fluorescence and SD) detection modes. The best reconstruction occurs in hybrid mode, in which quoted uncertainties (68% opening angles) of 0.6° are determined by observing Rayleigh-scattered light from laser shots. Angular resolutions for SD-only events are determined by comparing SD reconstruction with the hybrid results, and are found to depend on both the zenith angle (figure 17) and SD station hit multiplicity. Generally, SD angular resolution is found to be better than 2.2° for 3-hit events (approximately $E < 4$ EeV), better than 1.7° for 4-hit events (approximately $3 < E < 10$ EeV) and better than 1.4° for events with multiplicity of 5 or more (approximately $E > 8$ EeV).

The SUGAR air shower detector, consisting of buried liquid scintillator detectors at 50 m spacing, quotes a directional uncertainty of $3^\circ \sec\theta$ for air showers at zenith angle θ (Bellido *et al* 2001). Again, this number corresponds to the space angle that would include 68% of events from a point source. The parameterization should be regarded as an average over energies between $10^{17.9}$ and $10^{18.5}$ eV.

The error in arrival direction for the Yakutsk array is 3° for the primary cosmic rays with energies above 4×10^{19} eV discussed in this review (Uchihori *et al* 2000).

6.2. Large-scale and dipole anisotropy searches

At lower energies, for which magnetic diffusion may be expected to play a major role in determining event arrival directions, the most likely scenario for observing anisotropy is in extended excesses over the sky. (An exception applies to models in which the primary particles

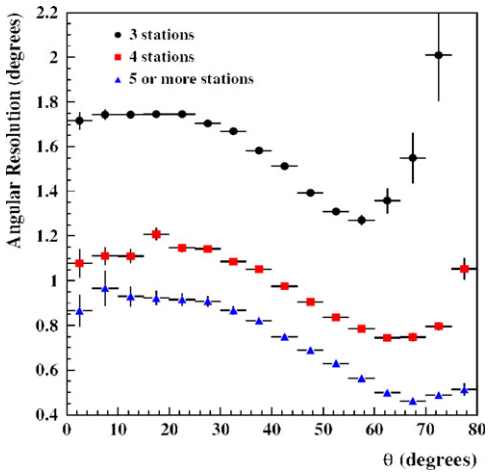

Figure 17. Angular resolution (68% opening angles) for the Auger surface detector (SD) as a function of zenith angle θ for various SD station multiplicities. Circles: three stations or approximately $E < 4$ EeV. Squares: four stations or approximately $3 < E < 10$ EeV. Triangles: more than four stations or approximately $E > 8$ EeV. Source: Bonifazi *et al* (2005).

are neutral; we consider this possibility in section 6.3.4.) In some cases, such excesses may be manifested as dipole moments in the full-sky arrival-direction distribution.

Due to near-uniform exposure of ground arrays in right ascension (RA), the study of harmonic structure—particularly the dipole moment—in RA has historically been an important tool in the search for large-scale anisotropy (Linsley 1975, Sokolsky *et al* 1992). Since in an isotropic sky events will be distributed uniformly in RA (ψ), the dipole fluctuation in intensity I is well parametrized by

$$I(\psi) = I_0(1 + r\cos(\psi - \phi)) \tag{3}$$

where r is known as the harmonic amplitude and ϕ is the phase of the dipole anisotropy.

The variables r and ϕ are commonly extracted by treating them as the magnitude and phase of a Rayleigh vector $\vec{r} = (x, y)$, with elements determined from the n observed events in the data set according to

$$x = \frac{2}{n}\sum_{i=1}^{n}\cos(\psi_i), \qquad y = \frac{2}{n}\sum_{i=1}^{n}\sin(\psi_i) \tag{4}$$

where ψ_i is the RA of the ith event.

The chance probability of observing in an isotropic data set a harmonic amplitude of strength $\geqslant r$ is given by the Rayleigh formula

$$P(\geqslant r) = \mathrm{e}^{-nr^2/4}. \tag{5}$$

As pointed out by Linsley (1975), however, only for $r \gg 2/\sqrt{n}$ (i.e. large anisotropy and/or large n) can r be interpreted as a credible estimator of the true anisotropy.

Hayashida *et al* (1999) report the results of a large-scale anisotropy search using the AGASA data around 1 EeV. First harmonic (dipole) amplitudes in RA are calculated for data by energy threshold, for thresholds between 10^{17} and 10^{20} eV (figure 18), for events having a zenith angle of $\leqslant 60^\circ$. There are approximately 114 000 events in the sample used in this study. Noting the high-significance effect near a threshold of 10^{18} eV, the authors perform logarithmically binned first harmonic analyses and find a harmonic amplitude $r = 4.1\%$ for

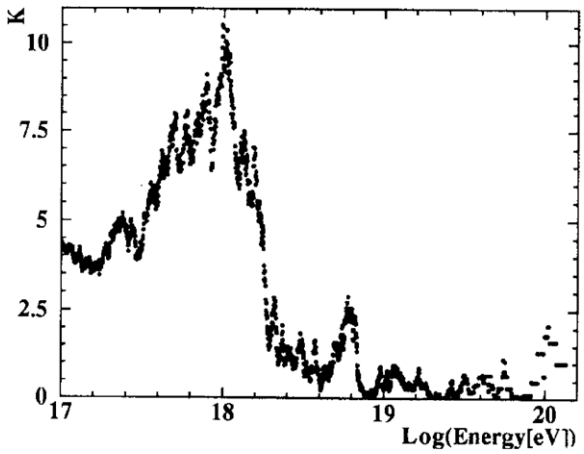

Figure 18. $K = nr^2/4$ for energy thresholds between 10^{17} and 10^{20} eV, for AGASA data taken prior to July 1995. See more complete explanation in the text. Source: Hayashida *et al* (1999).

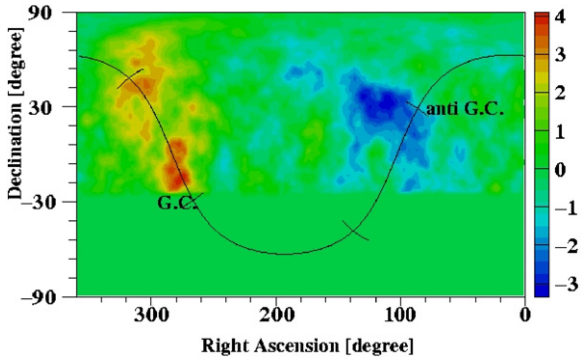

Figure 19. Significance map of excess and deficit events for energies $17.9 < \log E$ (eV) < 18.3. Events are summed within a 20° search circle and compared with expectations from an isotropic background. Source: AGASA web site (AGASA Collaboration 2003).

the $1.0 < E$ (EeV) < 2.0 bin, which contains 18 274 events. The chance probability of this effect in this bin is 0.000 35. Taking into account the number of logarithmic 'energy bins' of 0.5 decade width in the AGASA data (six), the authors arrive at a chance probability of ∼0.21% for this amplitude to be a statistical fluctuation. The authors do not claim to take into account any scan over bin sizes that may have been performed to maximize the observed signal. Considering the significance of fluctuations over the full-sky map (figure 19), Hayashida *et al* (1999) interpret the observed dipole behaviour as evidence for an excess in the direction of the galactic plane: in a 20° radius circle, 308 events are observed where 242.5 are expected (27% excess). Abbasi *et al* (2004c) present the results of searches for dipole anisotropy in the HiRes-I monocular data using a somewhat different technique, because the RA method is limited by its insensitivity to dipoles with predominantly north–south orientation. In this method, specific candidate dipole centres—the galactic centre, Cen A (Farrar and Piran 2000) and M87 (Biermann *et al* 2000)—are identified *a priori* on the basis of theoretical models. Opening angle distributions are plotted and compared to those observed in simulated data sets with dipole anisotropies of varying strengths.

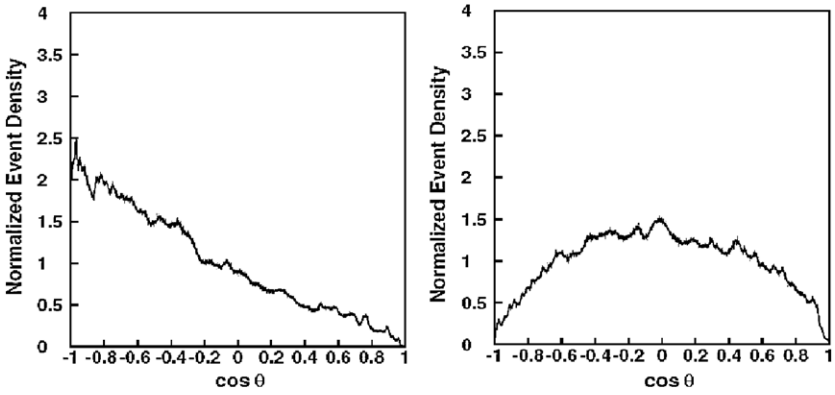

Figure 20. Left: galactic centre dipole function, with angular resolution included (as described in the text) for HiRes-I monocular data set. Right: galactic centre dipole function for simulated data set with $\alpha = 1$. Source: Abbasi *et al* (2004c).

In this analysis, the dipole anisotropy is parametrized as

$$I(\theta) = I_0(1 + \alpha \cos\theta) \tag{6}$$

where θ is the opening angle between a given direction and the global maximum of the distribution, and α is a measure of the strength of the dipole. Note that α is *not* equivalent to the harmonic amplitude r: $\cos\theta$ rather than θ is flat in the case of isotropy under this parameterization, and α is the slope of I versus $\cos\theta$. In general, r is related to α by

$$r = \frac{\pi}{4}\alpha \cos\delta \tag{7}$$

for a dipole with amplitude α centred at declination δ. Also note that there is no phase in this parameterization, the peak of the harmonic being defined as the coordinates of the *a priori* specified source.

The HiRes-I monocular data set used by Abbasi *et al* (2004c) consists of 1526 events observed between May 1997 and February 2003. The analysis technique used consists of histogramming the cosine of the space angle between all events and the maximum of the hypothetical dipole model (figure 20). As described in the previous section, air showers reconstructed by the monocular air fluorescence technique feature highly asymmetric (elliptical) uncertainties. To take this into account, each 'event' is actually represented by a large number of points spread according to the uncertainty functions (equations (1) and (2)), in the production of this histogram.

By comparing the 'dipole functions' thus obtained (specifically, the value of $\langle\cos\theta\rangle$ of the functions) to the dipole functions for simulated data sets with varying α, best fit values for α are obtained. This analysis is performed for each of the three *a priori* source candidates, with the following results: galactic dipole $-0.085 < \alpha < 0.090$ (90% c.l.), Cen A $-0.090 < \alpha < 0.085$ (90% c.l.) and M87 $-0.080 < \alpha < 0.070$ (90% c.l.).

A negative α in this formulation corresponds to a different source model entirely from that of a positive α. Specifically, the maximum of the first harmonic for negative α occurs at 180° from the positive α maximum. While this formulation is perhaps computationally easier, we feel that it muddies the interpretation of confidence limits somewhat and that therefore a formulation with positive-definite α is preferred. Nevertheless, comparisons can be made between this and prior results.

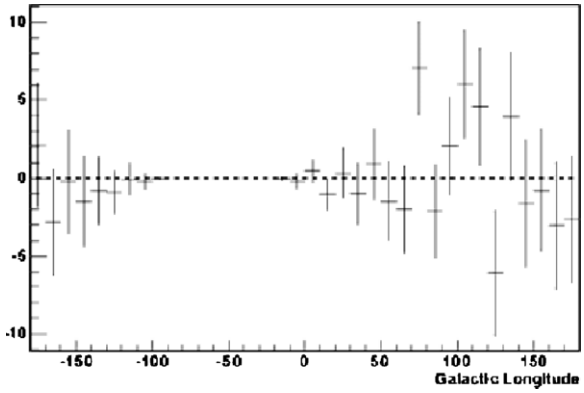

Figure 21. The number of excess HiRes stereo events as a function of galactic longitude for a $\pm 10^\circ$ region around the galactic equator (Sinnis 2005).

The HiRes-I limit on α (galactic dipole) translates to a limit of $-0.058 < r < 0.062$ (90% c.l.) (using equation (7)) for the harmonic amplitude. This result therefore does not exclude the prior AGASA result. In any case, the range of energies probed by the HiRes-I monocular data set does not overlap the energy range in which the AGASA excess was claimed.

The HiRes stereoscopic data set extends to lower energies than HiRes-I monocular, and is thus better suited to compare to the AGASA result. A complete harmonic analysis of the HiRes stereo data has not been reported in the literature, however an analysis of event excesses as a function of galactic longitude is reported by Sinnis (2005).

The data analysed were collected between September 1999 and February 2004, and consist of 4651 events passing quality and weather cuts and possessing a zenith angle less than 70°. 1438 events are in the AGASA energy band between 1 and 2 EeV. Figure 21 shows the distribution in galactic longitude of the number of excess events above background, for a $\pm 10^\circ$ region along the galactic equator. In the 0°–100° longitude region, HiRes measures 59 events against an expected background of 55.8 events, and places a 90% c.l. upper limit of 14.4 events (26%) excess. They conclude that no evidence for galactic centre or Cygnus region excesses exists in the HiRes data, however this result is not inconsistent with the claimed AGASA excess.

The first anisotropy results by the Pierre Auger Observatory are reported in Abraham *et al* (2007). This review aims at providing a direct test of both the AGASA and SUGAR (below, section 6.3.4) excesses in the vicinity of the galactic centre. Figure 22 shows the Auger map of the galactic centre, depicting Li-Ma (Li and Ma 1983) significances. These significances are calculated over windows of 5° radius, for events with energy between $10^{17.9}$ and $10^{18.5}$ eV. With approximately four times AGASA statistics (2116 events) in a 20° circle centred on the AGASA excess location, no significant excess is seen. For the range $10^{17.9} < E < 10^{18.5}$ eV, Auger reports $n_{\rm obs}/n_{\rm exp} = 1.01 \pm 0.02$, in contrast to the 27% excess reported by AGASA. Similar results are also reported for other nearby energy bands, as a check against energy scale discrepancies between the two experiments.

6.3. Searches for small-scale and point-like anisotropy

Point-like arrival-direction excesses may reasonably be expected under two scenarios: charged primaries with sufficient energy that they experience little magnetic deflection in the journey

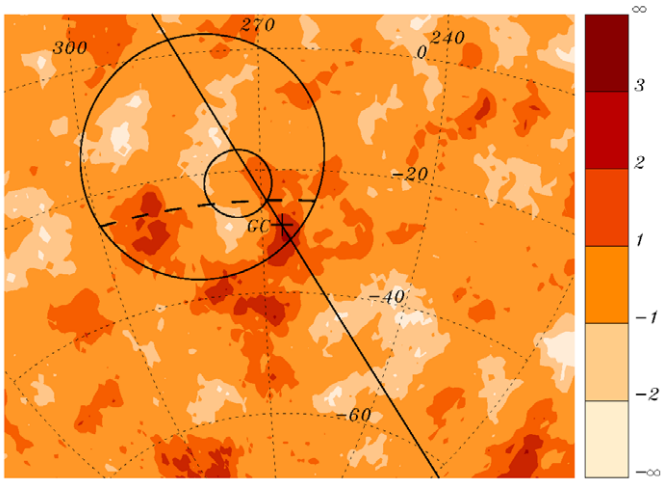

Figure 22. Map of Li–Ma overdensity significances near galactic centre (GC) region of skymap. Overdensities are calculated over windows of 5° radius, for events with energy between $10^{17.9}$ and $10^{18.5}$ eV. The GC is indicated by a cross, lying along the galactic plane (solid line). The regions of the AGASA (large circle) and SUGAR (small circle) excesses are also shown. From Abraham *et al* (2007).

from their source or neutral particles sufficiently long-lived to survive the journey. We consider searches appropriate to both models in the following discussion.

Over the past decade, considerable attention has been focused on reports of arrival-direction 'clustering' at the highest energies by the AGASA experiment. We begin our look at point-like excesses with these reports. As there has been much debate over the statistical significance of the AGASA clustering signal, we will also consider two reanalyses of AGASA data by independent groups. Then, we describe searches for similar effects in other experiments.

We consider the results of searches for point-like excesses at lower energies, motivated by an apparent excess in the galactic centre found in a SUGAR reanalysis and by neutron-source models. Finally, we consider searches for correlation of UHECR with the BL-Lacertae class of Quasi-Stellar Objects (QSOs).

6.3.1. The AGASA clustering signal. 'Clustering' of arrival directions for cosmic rays with energies above 4×10^{19} eV was first reported by AGASA in Hayashida *et al* (1996). In a sample of 36 events, three pairs (labelled C1, C2 and C3) were reported to have an angular separation of less than 2.5° in a data set with arrival-direction uncertainty of about 1.6° (figure 16). Using simulated isotropic data sets, a chance probability (P_{ch}) of such grouping was estimated to be 2.9%.

Additional clustering above 4×10^{19} eV using an expanded data set of 47 events was reported in Takeda *et al* (1999). An additional pair (C4) was found, and cluster C2 was promoted to the status of a 'triplet' with the addition of a third event within 2.5°. By lowering the energy cutoff to 3.89×10^{19} eV (and gaining a 48th event), the authors were able to identify a fifth cluster C5 as well. In this review, the authors estimated the chance probability of observing the same or a greater number of doublets (excluding C5) in an isotropic data set to be 0.32% (counting the triplet as three doublets) and the same or a greater number of triplets in an isotropic data set to be 0.87%.

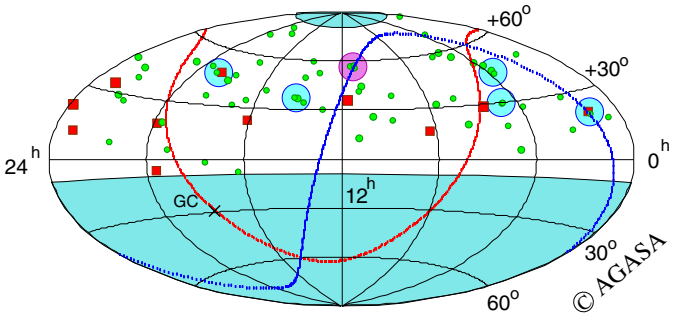

Figure 23. AGASA events with energy above 4×10^{19} eV superimposed on a skymap, in equatorial coordinates. Small circles represent events with energies of (4–10) $\times$ 10^{19} eV, small squares represent events with energies of $\geqslant 10^{20}$ eV. Clusters C1–C7 (C5 excluded) are identified, doublets with larger circles. The curves for the galactic (crossing galactic center, GC) and supergalactic planes are also shown. Figure taken from the AGASA web site (AGASA Collaboration 2003).

A sixth cluster (C6) was reported in Hayashida *et al* (2000), an update of the AGASA energy spectrum and cluster map. This update included arrival-direction information from an additional ten events above 4×10^{19} eV. Using the same technique as Takeda *et al* (1999), the triplet chance probability was revised to 1%. No revised probability was given for the doublets.

At the 27th ICRC, Takeda *et al* (2001) reported an analysis based on a total of 59 events, with five doublets and a triplet. It is not clear if this sample consists of one or two new events added to the Hayashida *et al* (2000) sample: energies below 4×10^{19} eV are not mentioned, and in the most recent skymaps (figure 23) a seventh cluster (C7) replaces C5. The arrival time and energy of one of the C7 events has not been made publicly available, however we have extracted the coordinates from figure 23. The available information on all events comprising the seven AGASA clusters is summarized in table 8.

At the 27th ICRC, AGASA reported $P_{\rm ch} = 0.05\%$ for doublets and $P_{\rm ch} = 1.66\%$ for the triplet. The results of scans of the threshold energy and angular separation criteria were presented. The original selection criteria that energies exceed $>4 \times 10^{19}$ eV and angular separations be less than 2.5° were found to be close to the values giving the most significant clustering signal. Also, the point spread function—representing the observed distribution of events from a point source, due to detector resolution effects—was used to demonstrate that these clusters are consistent with point-like sources at AGASA resolving power. Integral energy spectra were shown for all AGASA events and events comprising the clusters. The clusters are shown to have a relatively hard spectrum with index $-0.8 \pm 0.5(\text{stat}) \pm 0.5(\text{syst})$.

The most recent update of the cluster analysis by AGASA was presented by Teshima *et al* (2003) at the 28th ICRC. This analysis featured the same eight pairs ($8 = 1 \times 3 + 5 \times 1$) as previous analyses (C5 is omitted), and the calculated background was reported as 1.7 pairs. It was also reported that simulation studies indicate a chance probability of observing eight pairs when 1.7 are expected is $<10^{-4}$. This is below the Poisson fluctuation probability of 0.0003. Two-dimensional autocorrelation plots—scans for excesses over separation angles—were also presented, with some correlation of the two coordinates at the 3σ level indicating possible magnetic field deflections consistent with $B_z \sim 0.3$ μG in the galactic halo.

6.3.2. Independent commentary on AGASA clusters. Before turning to comparisons of the AGASA clustering signal with independent experimental measurements, we acknowledge the

Table 8. Clusters of cosmic rays reported by AGASA. Equatorial coordinates have been taken from Hayashida *et al* (2000) and figure 23.

	RA	DEC	Date	UT	$\times 10^{19}$ eV
C1	1^h15^m	21.1°	93/12/03	21:32:47	21.3
	1^h14^m	20.0°	95/10/29	00:32:16	5.07
C2	11^h29^m	57.1°	92/08/01	13:00:47	5.50
	11^h14^m	57.6°	95/01/26	03:27:16	7.76
	11^h13^m	56.0°	98/04/04	20:07:03	5.35
C3	18^h59^m	47.8°	91/04/20	08:24:49	4.35
	18^h45^m	48.3°	94/07/06	20:34:54	13.4
C4	4^h38^m	30.1°	86/01/05	19:31:03	5.47
	4^h41^m	29.9°	95/11/15	04:27:45	4.89
C5[a]	16^h06^m	23.0°	96/01/11	09:01:21	14.4
	15^h58^m	23.7°	97/04/10	02:48:48	3.89
C6	14^h17^m	37.7°	96/12/24	07:36:36	4.97
	14^h08^m	37.1°	00/05/26	18:38:16	4.89
C7[b]	3^h45^m	44.9°	98/10/27	00:45:37	6.11
	3^h41^m	46.6°	Unknown	Unknown	Unknown

[a] Cluster C5, consisting of one event with energy below the typical energy threshold of 4×10^{19} eV, was generally omitted in later publications.
[b] The arrival time and energy of the second event comprising C7 has not been made public.

ongoing debate about the significance of the AGASA result by considering two articles which attempt independent statistical evaluations. Both papers express concern that the choice of energy threshold (and in one case, the angular separation) cut used by AGASA was chosen *a posteriori* in order to maximize the significance of the observed signal.

Tinyakov and Tkachev (2001b) perform a reanalysis of the data above 4×10^{19} eV released in Hayashida *et al* (2000), consisting of four doublets and one triplet. (Clusters C5 and C7 from table 8 are not included in their analysis.) Taking the angular cut of 2.5° as valid based on AGASA's reported angular resolution (Hayashida *et al* 1996, Takeda *et al* 1999) they perform a scan over energy thresholds to determine the cutoff with the lowest chance probability. Since this scan is performed systematically, it can be used to estimate the statistical penalty to be paid for the *de facto*, but heretofore unacknowledged, scan used by AGASA in placing their energy cutoff at 4×10^{19} eV. The results of this scan are shown in figure 24. The minimum chance probability is $<10^{-4}$, however the authors estimate a scanning penalty factor of 3 and thus conclude that the chance probability of the AGASA result is approximately 3×10^{-4}.

Finley and Westerhoff (2004) take the additional step of questioning whether the AGASA angular separation cut was *a priori* justified. The authors point out that the AGASA angular uncertainty is energy dependent (see Takeda *et al* (1999) and figure 16), varying from 2.0° at $10^{19.5}$ to 1.2° at 10^{20} eV for the angular size of the 68% (1σ) error circles. Therefore, the choice of 2.5° as the cutoff was probably chosen so as to maximize the signal in the original AGASA clustering paper. The authors chose to scan over both energy and angular cuts (figure 24) in order to estimate the true significance of the result. While the results of the scan are consistent with the 'most significant energy' threshold reported by Tinyakov and Tkachev (2001b), the penalty for scanning over angles results in a chance probability of 0.35%, approximately ten times larger.

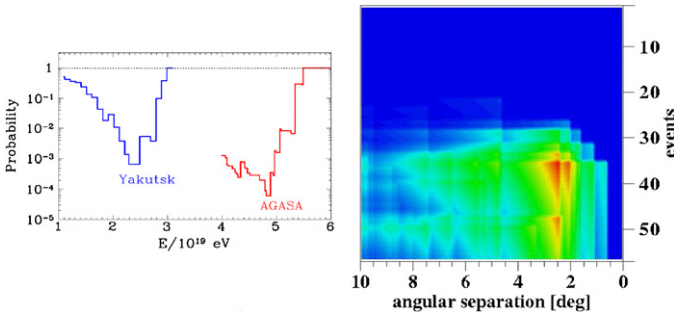

Figure 24. Left: scan of chance probability of observing the same or a greater number of pairs as seen in the Yakutsk and AGASA (Hayashida *et al* 2000) data as a function of energy threshold, for energies exceeding 4×10^{19} eV (Tinyakov and Tkachev 2001b). Right: scan of chance probability of observing the same or a greater number of pairs as seen in the AGASA data as a function of both energy threshold (event number is assigned by energy ordering) and angular separation cut (Finley and Westerhoff 2004). Event number 36 (the cutoff for the minimum chance probability) corresponds to an energy cut of 4.89×10^{19} eV. The colour scale is logarithmic, with red corresponding to chance probabilities of order few $\times 10^{-4}$. In both scans, the 'triplet' C2 is treated as three doublets.

As an additional check on the true sensitivity of the result, Finley and Westerhoff (2004) point out that it is appropriate to perform an independent analysis of AGASA data taken after that used in tuning the energy and angular cuts, i.e. that data for which the 2.5° and 4×10^{19} eV cuts can be regarded as *a priori*. Including just events observed after October 1995 (the end date for events included in the original cluster publication) just one pair is found, with a chance probability of 28%. Since doublet C4 contains one event on either side of the October 1995 cutoff, they then analyse the chance probability in a way that allows for cross-correlations between the two data sets. Allowing for this, the chance probability is 8%.

6.3.3. Search for clustering above 4×10^{19} *eV in other experiments.* It is clear that in order to resolve questions about the validity of the AGASA cluster claims, confirmation or refutation by an independent experiment is required. A data set with statistics comparable to the AGASA sample was obtained by the Yakutsk array over 30 years of running. Through 2003, Yakutsk detected a total of 29 extensive air showers with energy above 4×10^{19} eV. Two doublets are found in these data, with distance between the showers less than 5°. (The angular resolution of the Yakutsk array data is 3° (section 6.1) for air showers at these energies). The chance probability of these clusters is estimated to be 10% (Mikhailov 2004), which we interpret as being consistent with isotropy.

Tinyakov and Tkachev (2001b) include Yakutsk data with energies above 1×10^{19} eV in their reanalysis, described above in section 6.3.2 and figure 24. They find eight doubles (pairs of events within 4°) out of 26 events above a threshold energy of 2.4×10^{19} eV. Accounting for the scanning statistical penalty, they claim a chance probability of 0.002 for this clustering signal.

As of this writing, the data set collected by HiRes contains the best available statistics above 10^{19} eV. Abbasi *et al* (2004b) report on a search for arrival-direction clustering in the HiRes-I monocular data set above $10^{19.5}$ eV (figure 25). The data collected through 2003 contain 52 events with energy greater than $10^{19.5}$ eV, comparable in size to the data set used to establish the AGASA cluster claims.

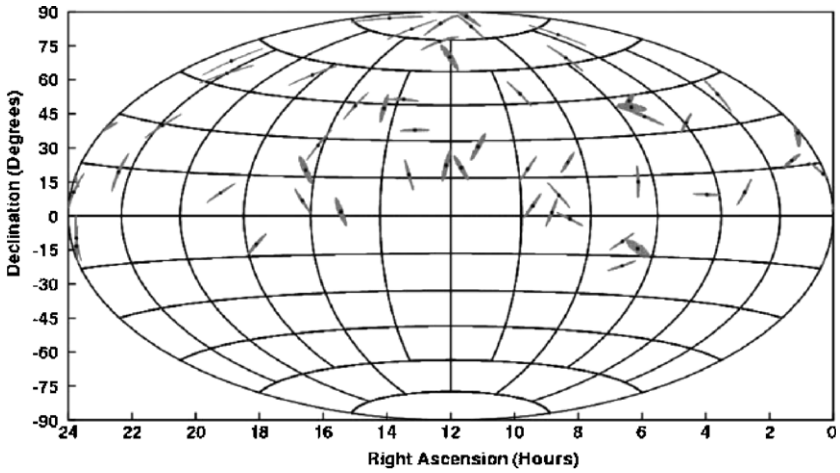

Figure 25. The arrival directions of the HiRes-I monocular events with reconstructed energies above $10^{19.5}$ eV and their 1σ angular resolution. Figure from Abbasi *et al* (2004b).

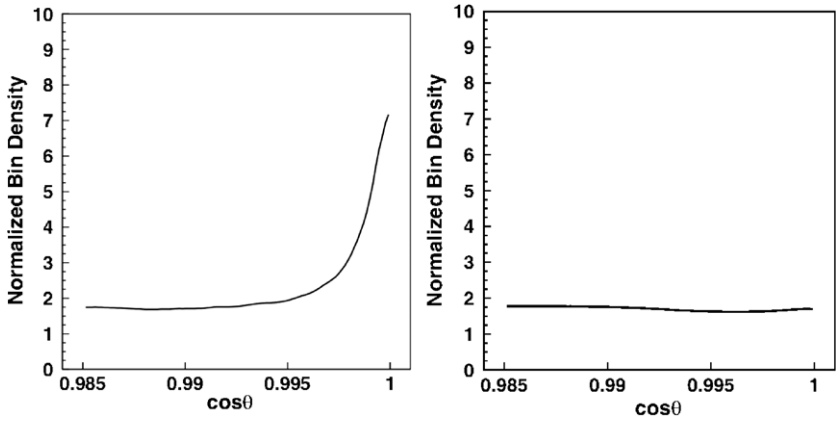

Figure 26. Left: autocorrelation function for AGASA data above 4×10^{19} eV (Hayashida *et al* 2000). Right: autocorrelation function for HiRes-I monocular data above $10^{19.5}$ eV. The upturn of the autocorrelation function at small opening angles is indicative of arrival-direction clustering. Both analyses performed in Abbasi *et al* (2004b).

This analysis is done by calculating an autocorrelation function, effectively a scan over event separation angles. In the monocular analysis, no scan is performed over energy threshold. Rather, the threshold is fixed at $10^{19.5}$ eV. As in the full-sky dipole analysis described above, each event is represented by a set of points distributed according to the angular resolution functions (equations (1) and (2)) in order to account for the effects of asymmetric error ellipses. For each point for a given event, the cosine of the opening angle with every point from all other events in the skymap is histogrammed. The results are shown in figure 26.

The significance of this result is established by comparison of the mean value of the $\cos\theta$ distribution for $\theta < 10^{\circ}$ with the same quantity for simulated data sets with varying numbers of clusters. The authors conclude that the HiRes-I data are consistent with no arrival-direction clustering and place a 90% c.l. upper limit of 3.5 doublets above background for the HiRes-I monocular data set.

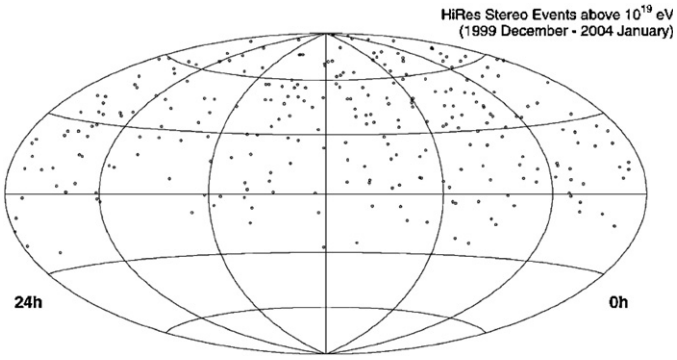

Figure 27. Skymap (equatorial coordinates) of 271 HiRes stereo events above 10^{19} eV. Shown is the typical error radius of 0.6°. From Abbasi *et al* (2004d).

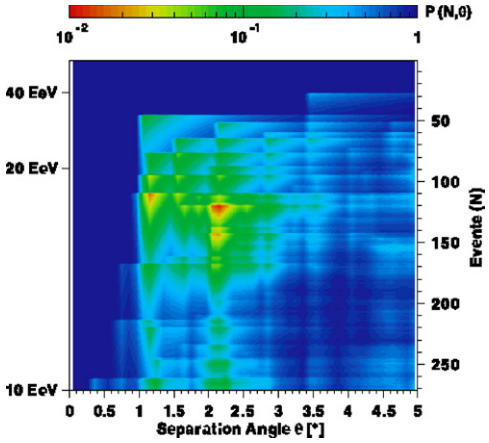

Figure 28. Autocorrelation scan of the HiRes data set above 10^{19} eV. $P(N, \theta)$ is the probability of obtaining the same or greater number of pairs as is actually observed in the data using a maximum separation angle θ and searching among the N highest energy events. These probabilities do not include the statistical penalty due to scanning. From Abbasi *et al* (2004d).

The HiRes stereo data set, while statistically smaller than the monocular data set, features significantly smaller, symmetric uncertainties in event arrival direction (figure 27). The search for arrival-direction clustering in the HiRes stereo data was carried out by Abbasi *et al* (2004d), using a procedure essentially identical to the energy-opening angle scan conducted by Finley and Westerhoff (2004) (section 6.3.2).

The strongest clustering signal, $P_{\min} = 1.9\%$, was observed at an energy threshold $E_c = 1.69 \times 10^{19}$ eV (figure 28). The signal observed corresponded to $n_p = 10$ pairs separated by less than $\theta_c = 2.2^\circ$ within a set of $N_c = 120$ events. Taking into account the statistical penalty for performing the scan, a statistical significance $P_{\text{ch}} = 52\%$ is attached to the clustering signal in the HiRes stereo data. The authors conclude that their data are consistent with the null hypothesis for arrival-direction clustering above that expected by chance.

We conclude our discussion of tests of the AGASA clustering signal with a noteworthy AGASA/HiRes stereo combined data set analysis. Abbasi *et al* (2005c) report the results of

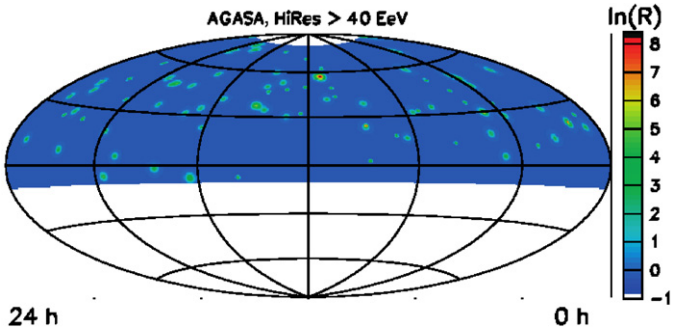

Figure 29. Likelihood ratio $\ln \mathcal{R}$, maximized with respect to n_s (number of source events), as a function of RA and DEC of the source position for the combined set of AGASA and HiRes events in excess of 4×10^{19} eV. Local maxima occur wherever events or clusters of evens are located on the sky. The global maximum, i.e., the most likely position of a point source is at RA $= 169.3^\circ$ and DEC $= 57.0^\circ$. From Abbasi *et al* (2005c).

an unbinned maximum likelihood analysis of the combined data sets above 4×10^{19} eV. The method is as follows.

Define $Q_i(x_i, x_s)$ as the probability for an event observed at coordinate x_i to have a true arrival direction at x_s, and $R_i(x)$ as the probability distribution for the event to be observed anywhere in the sky. R_i depends on the detector acceptance and exposure. The probability associated with a given event, under the point-source hypothesis, is the weighted sum P_i of the source and background probabilities:

$$P_i(x, x_s) = \frac{n_s}{N} Q_i(x, x_s) + \frac{N - n_s}{N} R_i(x). \tag{8}$$

The product of P_i for all events gives the likelihood L for a particular choice of the number of source events n_s. The best estimate for n_s is the value which maximizes L:

$$L(n_s, x_s) = \prod_{i=1}^{N} P_i(x, x_s). \tag{9}$$

The authors actually maximize $\ln \mathcal{R}$, the log of the ratio of the likelihood of n_s relative to the likelihood of the null hypothesis $n_s = 0$,

$$\ln \mathcal{R} = \ln \frac{L(n_s, x_s)}{L(0, x_s)}, \tag{10}$$

where $\ln \mathcal{R}$ is the measure of deviation from the null hypothesis of no source events.

Figure 29 shows the distribution of the likelihood ratio $\ln \mathcal{R}$ for the combined AGASA/HiRes stereo data sets. The significance of fluctuations in this skymap is determined by scanning over Monte Carlo data sets and counting the fraction with $\ln \mathcal{R}_{\mathrm{MC}} > \ln \mathcal{R}_{\mathrm{DATA}}$. The highest value of $\ln \mathcal{R} = 8.54$, corresponding to $n_s = 2.9$, is at the location of the AGASA triplet. The fraction of Monte Carlo sets with greater $\ln R$ is 28%. The authors conclude that no significant point source is found in the combined set of HiRes stereo and AGASA events above 40 EeV.

6.3.4. Search for point-like excesses at lower energy. Although at energies substantially lower than 10^{19} eV magnetic deflections are expected to diffuse arrival directions for charged particles, there are factors motivating the search for point-like excesses. In a reanalysis of

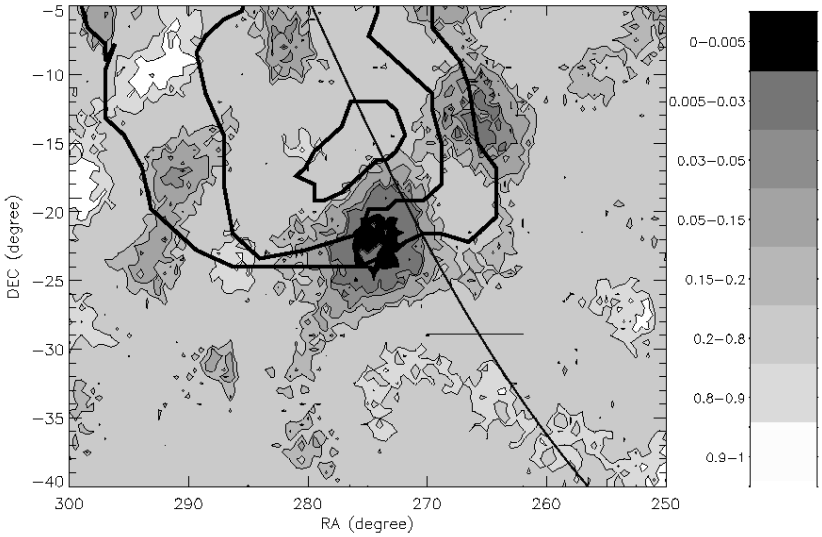

Figure 30. Significance of excess detected in the galactic centre region, in reanalysis of SUGAR data (Bellido *et al* 2001). Contours represent the chance probability of SUGAR detecting the observed density or greater. The galactic plane and galactic centre are indicated by the thin curved line and cross, respectively, while the thick lines represent 2σ, 3σ and 4σ contours from AGASA excess (Hayashida *et al* 1999). The central excess is point-like on the scale of the SUGAR angular resolution.

data from the SUGAR array designed to test the apparent AGASA excess in this region (section 6.2), Bellido *et al* (2001) observed evidence for a point source 7.5° from the galactic centre. Motivated in part by this observation, several theoretical models have predicted fluxes of neutrons from compact sources, in particular the galactic centre (Medina-Tanco and Watson 2001, Bossa *et al* 2003, Biermann *et al* 2004, Crocker *et al* 2005, Aharonian and Neronov 2005, Grasso and Maccione 2005).

The data used in the SUGAR reanalysis consist of 3732 events in the energy range $10^{17.9} \rightarrow 10^{18.5}$ eV, roughly corresponding to the range in which AGASA observes their most significant excess.

The analysis is carried out using a skymap technique, in which events are represented on a map by a Gaussian function equal to their directional uncertainty. The significance of fluctuations is determined by comparing the skymap for the data with simulated skymaps generated by the 'shuffling' technique. In the shuffling technique, most appropriate for ground array experiments with near-uniform exposure in right ascension (RA), simulated skymaps are generated by pairing a real arrival time from one event with a zenith and azimuth angle from another event in the same set, and repeating until a new data set is filled. The resultant 'significance map', in the vicinity of the galactic centre, is shown in figure 30.

A localized excess (point-like within the SUGAR resolution) is observed in the data, centred at RA 274° DEC −22°, close to the position of the AGASA excess and the galactic centre. No excess is observed from the true centre of the galaxy, although SUGAR has a clear view of this region. The excess has an estimated chance probability $P_{\rm ch} = 0.5\%$ and corresponds to a point-source flux of (2.7 ± 0.9) km^{-2} yr^{-1} between $10^{17.9}$ and $10^{18.5}$ eV.

Two recent submissions describe further point-source searches, one each in the northern and southern hemispheres. Abbasi *et al* (2007b) describe a point-source search in northern

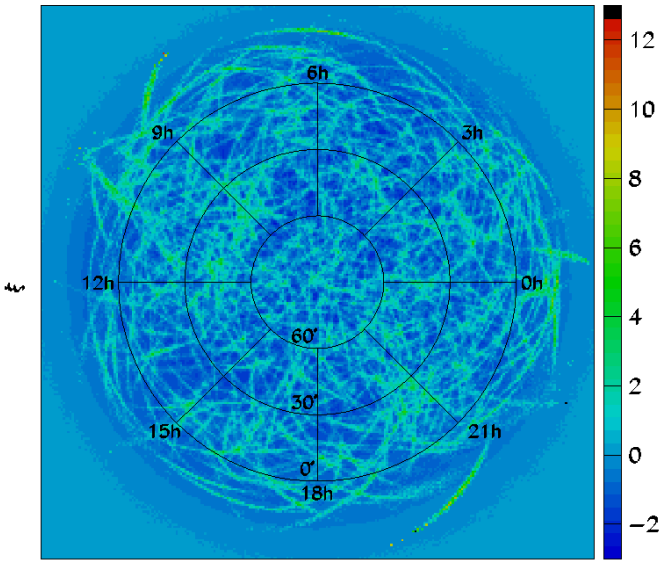

Figure 31. Significance map, plotted in polar projection, equatorial coordinates, for the HiRes-I monocular data set. The quantity plotted is $\xi = (N_{\mathrm{DATA}} - \langle N_{\mathrm{MC}} \rangle)/\sigma_{\mathrm{MC}}$, where each event is represented by an asymmetric probability density function (equations (1) and (2)). From Abbasi *et al* (2007b).

hemisphere using the HiRes-I monocular data above $10^{18.5}$ eV. In the HiRes technique, a binned sensitivity map (figure 31) is created from the data (1525 events) and Monte Carlo simulated data sets. Each 'event' in this map is represented by a probability density function determined from equations (1) and (2).

To search for point sources, Abbasi *et al* (2007b) look for localized excesses by calculating the fraction of 'significant' bins within a 2.5° search circle. Simulated sources are used to tune selection criteria, to maximize sensitivity to true point sources while rejecting false positives due to fluctuations in the isotropic background. This tuning is performed 'blind', i.e. without looking at the actual data. The selection criteria are then applied to the data and no significant point-like excesses are found.

HiRes places an upper limit of 0.8 $\mathrm{km}^{-2}\ \mathrm{yr}^{-1}$ (90% c.l.) on the flux from point-like sources of hadronic cosmic rays in the northern hemisphere. This is below the flux reported by the SUGAR experiment for the point-like source near the galactic centre.

The Auger galactic centre studies reported in Abraham *et al* (2007) and described above in section 6.2 also contain a direct test of the reported SUGAR excess. As shown in figure 22, Auger sees no significant excess in the SUGAR signal region for events with energy between $10^{17.9}$ and $10^{18.5}$ eV. Using a 'semi-analytic' technique to evaluate significances, Auger places a flux limit of $\Phi^{95} = \kappa\ 0.13\ \mathrm{km}^{-2}\ \mathrm{yr}^{-1}$ (95% c.l.) on point-like neutron sources in the vicinity of the galactic centre, where κ is a correction factor for overall flux normalization: adopting the HiRes (AGASA) spectral normalization at 3 EeV corresponds to $\kappa = 1.2$ ($\kappa = 2$). In any case, this number is well below the point-like excess flux reported by SUGAR.

6.3.5. Correlations of UHECR with BL-Lacertae objects. BL-Lacertae (BL-Lac) objects are a class of blazars, active galactic nuclei in which the jet axis points towards the Earth.

It is known that blazars are sources of gamma rays with energies ranging from >100 MeV (Hartman *et al* 1999) to over a TeV (Aharonian *et al* 2006). Such high-energy photons are possible byproducts of UHECR acceleration, and hence BL-Lac objects are an obvious choice for correlation studies with UHECRs.

Over 1000 BL-Lac objects have been catalogued as of this writing. See for example Véron-Cetty and Véron (2006). This is far greater than the world sample of UHECRs which have been observed with sufficient energy that their arrival directions might be correlated with distant point-like objects: below a few $\times 10^{19}$ eV, trajectories of charged particles are expected to be significantly perturbed by interactions with galactic and extragalactic magnetic fields, obscuring the primary cosmic ray's point of origin.

Such a wealth of potential sources comes with a potential trap, however. By choosing subsets of BL-Lac source objects and varying cut criteria on UHECR events, one may readily find correlations but at the same time accumulate sampling penalties which make the true significance difficult or impossible to know. The authors of the earliest papers reporting correlations of BL-Lacs with UHECRs from AGASA and Yakutsk data sets (Tinyakov and Tkachev 2001a, Tinyakov and Tkachev 2002, Gorbunov *et al* 2002) perform such explicit tuning, and then apply a 'penalty factor' to estimate true statistical significances. Evans *et al* (2003) compare the same data sets to an untuned BL-Lac catalogue and observe no effect. From this they conclude that the previously reported signal was 'entirely due to selection effects', a claim that Tinyakov and Tkachev (2004) dispute.

The soundest use of the claims made using tuned data sets is as a guideline for *a priori* searches in statistically independent data. In reanalysis of data from Haverah Park and Volcano Ranch, Torres *et al* (2003) found no support for the claims of Tinyakov and Tkachev (2001a), but did not specifically address correlations with the BL-Lac subset used by Gorbunov *et al* (2002). The HiRes stereo data above 24 EeV (Abbasi *et al* 2006b) do not support either of these claims.

The most recent claims have been made with 271 events above 10 EeV in the HiRes stereo sample. Gorbunov *et al* (2004) compare these events to 157 BL Lacs from the 10th Veron Catalog (Véron-Cetty and Véron 2000, Véron-Cetty and Véron 2001) having optical magnitude $m < 18$. In a 0.8° binned analysis, they find 10 BL-Lac/UHECR pairs with a chance probability of 0.001 after application of a penalty factor. An unbinned maximum likelihood analysis of the same data by the HiRes Collaboration (Abbasi *et al* 2006b) concurs with the Gorbunov *et al* (2004) conclusions.

Due to the energy threshold used in this analysis, primary cosmic rays originating from distant QSOs would almost certainly need to be neutral to maintain their arrival-direction correlation after passing through galactic and extragalactic magnetic fields. But neutrons possess a lifetime of 3×10^{12} s at 10 EeV and so cannot have originated more than 100 kpc from Earth. The mean-free path of photons at these energies is of the order of a few Mpc. Thus, these results, if true, are particularly intriguing. We await the results of statistically independent comparisons from the next generation of UHECR observatories.

6.4. Summary of anisotropy results

Over the past decade, the search for anisotropy in the arrival directions of the highest energy cosmic rays has been pursued aggressively. During that time, suggestions of statistically significant arrival-direction excesses on all angular scales have invariably failed the crucial test of confirmation by independent data sets. We have not yet arrived at the era of cosmic ray astronomy.

At lower energies (below 10^{19} eV), in spite of substantial statistics there is no confirmed large-scale structure to the cosmic ray sky. While AGASA has seen evidence for an enhanced flux from the galactic centre, the HiRes and Auger observatories have seen no such evidence in data sets with comparable event counts. Tantalizing evidence in the reanalysis of SUGAR data suggested a low-energy point-like excess near the galactic centre, prompting the promulgation of models featuring compact sources of neutrons. But the Auger observatory has set flux limits well below the SUGAR observation level in the galactic centre region, and HiRes has set new constraints on neutron sources in the northern hemisphere.

At higher energies, above a few $\times 10^{19}$ eV, even charged particles should experience negligible deflections by galactic and extragalactic magnetic fields. Consequently, the arrival directions of these primary cosmic rays should maintain their correlation with their sources, leading to detectable small-scale anisotropies. Exciting reports of arrival-direction clustering by the AGASA experiment have not been confirmed by HiRes. And evidence for UHECR correlations with the plausible source candidate BL-Lacertae objects has not persisted in independent data sets.

One important lesson for the community to take from the past decade of anisotropy studies is just how susceptible we are to fooling ourselves with statistics. The more ways one partitions the data, and the more one tunes signal selection criteria, the more chances one has to observe a *fluctuation* and be tempted to call it an *effect*. We are encouraged to note that a discussion is emerging about the pitfalls of hidden trials and the proper fair usage of experimental data, which we hope will improve the efficiency and credibility of anisotropy studies in the future.

Acknowledgments

This work is supported by the US National Science Foundation under grants PHY-0305516 and PHY-0636162. The authors thank M Kirn, L Scott, B Stokes, G Thomson, I Tkachev and G Yodh for useful discussion and advice.

References

Abbasi R U *et al* 2004a *Phys. Rev. Lett.* **92** 151101 (*Preprint* astro-ph/0208243)
Abbasi R U *et al* 2004b *Astropart. Phys.* **22** 139–49 (*Preprint* astro-ph/0404366)
Abbasi R U *et al* 2004c *Astropart. Phys.* **21** 111–23 (*Preprint* astro-ph/0309457)
Abbasi R U *et al* 2004d *Astrophys. J.* **610** L73–6 (*Preprint* asto-ph/0404137)
Abbasi R U *et al* 2005a *Astropart. Phys.* **23** 157–74 (*Preprint* astro-ph/0208301)
Abbasi R U *et al* 2005b *Phys. Lett.* B **619** 271–80 (*Preprint* astro-ph/0501317)
Abbasi R U *et al* 2005c *Astrophys. J.* **623** 164–70 (*Preprint* astro-ph/0412617)
Abbasi R U *et al* 2005d *Astrophys. J.* **622** 910–26 (*Preprint* astro-ph/0407622)
Abbasi R U *et al* 2006a *Astropart. Phys.* **25** 93–7 (*Preprint* astro-ph/0601091)
Abbasi R U *et al* 2006b *Astrophys. J.* **636** 680–4 (*Preprint* astro-ph/0507120)
Abbasi R U *et al* 2006c *Astropart. Phys.* **25** 74–83 (*Preprint* astro-ph/0512423)
Abbasi R U *et al* 2007a *Phys. Rev. Lett.* (*Preprint* astro-ph/0703099) submitted
Abbasi R U *et al* 2007b *Astropart. Phys.* (*Preprint* astro-ph/0507663)
Abbasi R U *et al* 2007c *Astropart. Phys.* (*Preprint* astro-ph/0607094)
Abraham J *et al* 2007 *Astropart. Phys.* **27** 244–53 (*Preprint* astro-ph/0607382)
Abu-Zayyad T *et al* 2000a *Phys. Rev. Lett.* **84** 4276–9
Abu-Zayyad T *et al* 2000b *Nucl. Instrum. Methods* A **450** 253–69
Abu-Zayyad T *et al* 2001 *Astrophys. J.* **557** 686–99 (*Preprint* astro-ph/0010652)
Afanasiev B N *et al* 1993 *Proc. Tokyo Workshop on the Techniques for the Study of the Extremely High Energy Cosmic Rays* ed M Nagano p 35
AGASA Collaboration 2003 http://www-akeno.icrr.u-tokyo.ac.jp/AGASA
Aharonian F *et al* 2006 *Astron. Astrophys.* **455** 461–6 (*Preprint* astro-ph/0607569)

Aharonian F and Neronov A 2005 *Astrophys. J.* **619** 306–13 (*Preprint* astro-ph/0408303)
Auger Collaboration 2004 *Nucl. Instrum. Methods* A **523** 50–95
Auger P *et al* 1939 *Rev. Mod. Phys.* **11** 288–41
Ave M *et al* 2000 *Phys. Rev. Lett.* **85** 2244–7 (*Preprint* astro-ph/0007386)
Ave M *et al* 2001 *Proc. 27th ICRC (Hamburg)* pp 381–4
Ave M *et al* 2002 *Phys. Rev.* D **65** 063007–1–19 (*Preprint* astro-ph/0110613)
Ave M *et al* 2003a *Astropart. Phys.* **19** 47–60 (*Preprint* astro-ph/0112253)
Ave M *et al* 2003b *Astropart. Phys.* **19** 61–75 (*Preprint* astro-ph/0203150)
Baltrusaitis R M *et al* 1985 *Nucl. Instrum. Methods* A **240** 410–28
Bellido J A *et al* 2001 *Astropart. Phys.* **15** 167–75 (*Preprint* astro-ph/0009039)
Berezinsky V, Gazizov A and Grigorieva S 2006 *Phys. Rev.* D **74** 0403005 (*Preprint* hep-ph/0204357)
Berezinsky V S and Grigor'eva S I 1988 *Astron. Astrophys.* **199** 1–12
Bergman D R 2001 *Proc. 27th ICRC (Hamburg)* pp 639–40
Bergman D R 2007 *Nucl. Phys.* B **165** 19–26 (*Preprint* astro-ph/0609453)
Biermann P L, Ahn E J, G M T and Stanev T 2000 Origin of the highest energy cosmic rays *Preprint* astro-ph/0008063
Biermann P L *et al* 2004 *Astrophys. J.* **604** L29–32 (*Preprint* astro-ph/0401140)
Bird D J *et al* 1993 *Phys. Rev. Lett.* **71** 3401–4
Bird D J *et al* 1994 *Astrophys. J.* **424** 491–502
Bonifazi C *et al* 2005 *Proc. 29th ICRC (Pune)* vol 7 pp 17–20
Borione A *et al* 1994 *Nucl. Instrum. Methods* A **346** 329–52
Bossa M, Mollerach S and Roulet E 2003 *J. Phys. G: Nucl. Part. Phys.* **29** 1409–21 (*Preprint* astro-ph/0304023)
Boyer J H *et al* 2002 *Nucl. Instrum. Methods* A **482** 457–74
Cassiday G L *et al* 1990 *Astrophys. J.* **356** 669–74
CERN, Application Software Group 1993 *Geant: Detector Description and Simulation Tool Technical Report* W5013 CERN, Available from http://wwwasdoc.web.cern.ch
Chiba N *et al* 1992 *Nucl. Instrum. Methods* A **311** 338–49
Crocker R M *et al* 2005 *Astrophys. J.* **622** 892–909 (*Preprint* astro-ph/0408183)
Egorova V P *et al* 2001 *J. Phys. Soc. Japan* **70** (Suppl. B) 9–14
Egorova V P *et al* 2004 *Proc. Cosmic Ray International Seminar, Catania 2004* vol 136 pp 3–11
Evans N W, Ferrer F and Sarkar S 2003 *Phys. Rev.* D **67** 103005 (*Preprint* astro-ph/0212533)
Farrar G R and Piran T 2000 Deducing the source of ultrahigh energy cosmic rays *Preprint* astro-ph/0010370
Finley C B and Westerhoff S 2004 *Astropart. Phys.* **21** 259–367 (*Preprint* astro-ph/0309159)
Fletcher R S *et al* 1994 *Phys. Rev.* D **50** 5710–31
Gorbunov D S, Tinyakov P G, Tkachev I I and Troitsky S V 2002 *Astrophys. J.* **577** L93–6 (*Preprint* astro-ph/0204360)
Gorbunov D S, Tinyakov P G, Tkachev I I and Troitsky S V 2004 *J. Exp. Theor. Phys. Lett.* **80** 145–8 (*Preprint* astro-ph/0406654)
Grasso D and Maccione L 2005 *Astropart. Phys.* **24** 273–88 (*Preprint* astro-ph/0504323)
Greisen K 1966 *Phys. Rev. Lett.* **16** 748–50
Hartman R C *et al* 1999 *Astrophys. J. Suppl. Ser.* **123** 79–202
Hayashida N *et al* 1995 *J. Phys. G: Nucl. Part. Phys.* **21** 1101–19
Hayashida N *et al* 1996 *Phys. Rev. Lett.* 1000–3
Hayashida N *et al* 1999 *Astropart. Phys.* **10** 303–11 (*Preprint* astro-ph/9807045)
Hayashida N *et al* 2000 Updated AGASA event list above 4×10^{19} eV *Preprint* astro-ph/0008102
Heck D *et al* 1998 Corsika: a Monte Carlo code to simulate extensive air showers *Technical Report* FZKA 6019 (Forschungszentrum Karlsruhe)
Heitler W 1938 *Proc. R. Soc.* A **166** 529
Hess V F 1912 *Z. Phys.* **13** 1084
Hinton J A *et al* 1999 *Proc. 26th ICRC (Salt Lake City)* vol 3 pp 288–391
Ivanov A A, Knurenko S P and Sleptsov I Y 2003 *Nucl. Phys.* B **122** 226–30 (*Preprint* astro-ph/0305053)
Kalmykov N N, Ostapchenko S S and Pavlov A I 1997 *Nucl. Phys.* B **52b** 17–28
Knurenko S P *et al* 2007 *Nucl. Phys.* B (*Preprint* astro-ph/0611871)
Lawrence M A, Reid R J O and Watson A A 1991 *J. Phys. G: Nucl. Part. Phys.* **17** 733–57
Li T P and Ma Y Q 1983 *Astrophys. J.* **272** 317–24
Linsley J 1963 *Phys. Rev. Lett.* **10** 146–8
Linsley J 1975 *Phys. Rev. Lett.* **34** 1530–3
Mantsch P *et al* 2005 *Proc. 29th ICRC (Pune)* vol 10 pp 115–24
Medina-Tanco G A and Watson A A 2001 *Proc. 27th ICRC (Hamburg)* pp 531–4
Mikhailov A A 2004 *Proc. 28th RCRC (Moscow)* p 1306 (*Preprint* astro-ph/0403231)

Nagano M *et al* 1984 *J. Phys. Soc. Japan* **53** 1667–81
Nagano M *et al* 1992 *J. Phys. G: Nucl. Part. Phys.* **18** 423–42
Nagano M and Watson A A 2000 *Rev. Mod. Phys.* **72** 689–732
Ohoka H *et al* 1997 *Nucl. Instrum. Methods* A **385** 268–76
Pierog T and Warner K 2006 New facts about muon production in extended air shower simulations *Preprint* astro-ph/0611311
Pierre Auger Collaboration 2005 *Proc. 29th ICRC (Pune)* usa-sommers-P-abs1-he14-oral
Pravdin M I *et al* 1999 *Proc. 26th ICRC (Salt Lake City)* vol 3 pp 292–5
Reil K 2002 The energy spectrum of ultra high energy cosmic rays measured by the High Resolution Fly's Eye detectors *PhD Thesis* Utah, USA
Sadowski P A *et al* 2002 *Astropart. Phys.* **18** 237–48
Sakaki N *et al* 2001 *Proc. 27th ICRC (Hamburg)* pp 333–6
Sciutto S J 1999 Aires: a system for air shower simulations *Preprint* astro-ph/9911331
Shinozaki K and Teshima M 2004 *Proc. Cosmic Ray International Seminar, Catania 2004* vol 136 pp 18–27
Sinnis G 2005 *Proc. 29th ICRC (Pune)* vol 7 pp 377–80
Sokolsky P 1989 *Introduction to Ultrahigh Energy Cosmic Ray Physics* (Reading, MA: Addison-Wesley)
Sokolsky P, Sommers P and Dawson B R 1992 *Phys. Rep.* **271** 225–77
Stokes B T 2005 A search for anisotropy in the arrival directions of ultra high energy cosmic rays observed by the high resolution fly's eye detector *PhD Thesis* Utah, USA
TA Collaboration 2005 *Proc. 29th ICRC (Pune)* vol 8 pp 141–4
Takeda M *et al* 1999 *Astrophys. J.* **522** 225–37 (*Preprint* astro-ph/9902239)
Takeda M *et al* 2001 *Proc. 27th ICRC (Hamburg)* pp 341–4
Takeda M *et al* 2003 *Astropart. Phys.* **19** 447–62 (*Preprint* astro-ph/0209422)
Teshima M 2002 personal communication (From public remarks made at the 2002 Aspen Winter Conference on High Energy Radiation from Space)
Teshima M *et al* 2003 *Proc. 28th ICRC (Tsukuba)* pp 437–40
Tinyakov P G and Tkachev I I 2001a *J. Exp. Theor. Phys. Lett.* **74** 499–503 (*Preprint* astro-ph/0102476)
Tinyakov P G and Tkachev I I 2001b *J. Exp. Theor. Phys. Lett.* **74** 1–5 (*Preprint* astro-ph/0102101)
Tinyakov P G and Tkachev I I 2002 *Astropart. Phys.* **18** 165–72 (*Preprint* astro-ph/0111305)
Tinyakov P G and Tkachev I I 2004 *Phys. Rev.* D **69** 128301 (*Preprint* astro-ph/0301336)
Torres D F *et al* 2003 *Astrophys. J.* **595** L13–6 (*Preprint* astro-ph/0307079)
Uchihori Y *et al* 2000 *Astropart. Phys.* **13** 151–60 (*Preprint* astro-ph/9908193)
Véron-Cetty M P and Véron P 2000 A catalogue of quasars and active nuclei *Technical Report* 19 ESO
Véron-Cetty M P and Véron P 2001 *Astron. Astrophys.* **374** 92–4
Véron-Cetty M P and Véron P 2006 *Astron. Astrophys.* **455** 773–7
Winn M M *et al* 1986a *J. Phys. G: Nucl. Phys.* **12** 675–86
Winn M M *et al* 1986b *J. Phys. G: Nucl. Phys.* **12** 653–74
Yao W M *et al* 2006 *J. Phys. G: Nucl. Part. Phys.* **33** 302
Yoshida S and Dai H 1998 *J. Phys. G: Nucl. Part. Phys.* **24** 905–38 (*Preprint* astro-ph/9802294)
Zatsepin G T and Kuz'min V A 1966 *J. Exp. Theor. Phys. Lett.* **4** 78–80
Zatsepin G T and Kuz'min V A 1966 *J. Exp. Theor. Phys. Lett.* **4** 114–7