Available online at www.sciencedirect.com

SCIENCE DIRECT®

C. R. Physique 5 (2004) 473–482

COMPTES RENDUS PHYSIQUE

Ultimate energy particles in the Universe/Particules d'énergies ultimes dans l'Univers

Directional reconstruction and anisotropy studies

Roger W. Clay *, Bruce R. Dawson, Greg J. Thornton

Department of Physics, University of Adelaide, Adelaide, SA 5005, Australia

Available online 27 April 2004

Presented by Pierre Encrenaz

Abstract

The incident directions of cosmic ray particles at the highest energies are readily measured by modern detectors. Provided that there are sufficient data, the interpretation of such directions should give strong information on the origin of cosmic rays. However, the choice of statistical methods used in such an analysis can be of paramount importance in the reliability of the interpretation. This article summarizes the conclusions which may be drawn from the present data, and proposes methods needed to interpret statistically the large quantities of data which will become available in the near future. ***To cite this article: R.W. Clay et al., C. R. Physique 5 (2004).***

Résumé

Reconstruction de la direction et étude des anisotropies. La direction de provenance des rayons cosmiques, aux énergies les plus élevées, est facilement mesurable par les détecteurs modernes. Avec des statistiques suffisantes, l'interprétation de cette information devrait donner des indications fortes sur l'origine de ces rayons cosmiques. Toutefois, pour que de telles informations soient fiables, le choix des méthodes statistiques est extrêmement important. Nous résumons dans cet article les conclusions que l'on peut tirer à partir des données actuellement disponibles, et nous proposons les méthodes qui nous semblent nécessaires pour l'interprétation statistique des grandes quantités de données qui devraient devenir disponibles dans un futur proche. ***Pour citer cet article : R.W. Clay et al., C. R. Physique 5 (2004).***

Keywords: Direction finding; Anisotropies; Point sources; Statistical methods

Mots-clés : Reconstruction de la direction ; Anisotropies ; Sources ponctuelles ; Méthodes statistiques

1. Introduction

We observe cosmic rays from a location which is asymmetrically placed with respect to most source distributions that one might imagine. For example, we view our galaxy from a location towards the edge of its luminous matter and from within its thin plane. X-ray and gamma-ray sky maps show concentrations of counts within the galactic plane and towards the galactic centre, so one might expect the cosmic ray sky to exhibit similar noticeable structure. Cosmic ray directional structure such as this is known as an anisotropy and has proved remarkably difficult to detect.

It seems likely that our galaxy is incapable of accelerating particles to the highest known energies [1]. However, on a larger astrophysical distance scale, the Greisen–Zatsepin–Kuz'min (GZK) effect [2] is expected to limit our horizon to 100 Mpc at the very most. Within that distance, further structures such as our Local Group of galaxies, the Fanaroff–Riley classified objects,

* Corresponding author.
E-mail addresses: rclay@physics.adelaide.edu.au (R.W. Clay), bdawson@physics.adelaide.edu.au (B.R. Dawson), gthornto@physics.adelaide.edu.au (G.J. Thornton).

1631-0705/$ – see front matter
doi:10.1016/j.crhy.2004.03.010

and the major Virgo and Coma clusters of galaxies indicate specific directions which one might expect to identify in any cosmic ray data-set. These structures are so few that one would be surprised if their overlap resulted in any directional averaging.

To the present time, there is convincing evidence for cosmic ray anisotropy only at rather low energies where huge event numbers have enabled very low (a few parts in 10^{-4}) deviations from isotropy to be detected. This lack of anisotropy has been unexpected. It is presumed to be due to the directional scrambling produced by galactic and intergalactic magnetic fields, although those fields also are expected to have some structure which would be capable of channelling cosmic ray arrival directions (see [3] in this volume). A knowledge of the strength and structure of such fields is therefore a key complement to measurements of the cosmic ray anisotropy. Unfortunately, our knowledge of those fields, and thus their impact on the anisotropy, is currently very poor with the result that starkly differing conclusions may be reached [4,5].

As a general principle, magnetic field induced directional deviations are expected to reduce with energy, provided that the energies for diffusive propagation are exceeded. Thus, one expects directional measurements at the highest energies to be key data in discovering the sources of cosmic rays. Unfortunately, the differential cosmic ray flux reduces roughly with the cube of the energy, and statistical limitations dominate at the highest energies. There is thus a need to consider arguments for optimising the energy range to be used in any anisotropy search. Such optimisation must make assumptions on likely source distances, the charge distribution of the primary particles, and the magnitude and structures of any intervening magnetic fields. A characteristic angular uncertainty is then predictable and analysis procedures can be matched to that expectation. The new Pierre Auger Observatory has an angular resolution which will enable it to make excellent use of any neutral particle (undeflected) directions, and it will have a huge collecting area in two hemispheres which will enable it to study broad scale directional anisotropies with unprecedented sensitivity.

2. Direction finding

The first step in an anisotropy search is to have reliable measurements of cosmic ray arrival directions, preferably with estimates of the measurement uncertainty. The basic techniques for determining the arrival directions of air showers have been discussed by Sommers [6] in this volume. As pointed out there, air shower detectors have traditionally been one of two types, a surface array or a fluorescence detector, but recently the two have come together in the Auger observatory [7] and in plans for the Telescope Array [8] – these are so-called hybrid detectors. In this section we will briefly review these techniques and indicate the typical precision achievable in arrival direction reconstruction.

Sommers has presented the generic surface array technique, where measurements of the shower front arrival time at a number of (at least 3) non-collinear detectors returns the shower arrival direction. In principle this is a simple technique, but there are subtleties which must be optimized for the best result. The sparseness of typical surface arrays is always a problem (e.g., 10 m^2 detectors on a 1.5 km grid for the Auger Project, 2.2 m^2 detectors on a 1 km grid for AGASA), and so is the sparseness of the particles in the shower front itself, especially far from the core. This requires a careful definition of what is meant by the shower front arrival time at a detector, and a proper representation of the fluctuations expected because of the sparse sampling. The large-scale shape of the shower front is also a consideration. A plane shower front has often been assumed in the past, but the need for improved direction resolution has led to better characterization of this shape. In principle the shape depends on the development history of the shower, so it will change slightly from one shower to the next. It is also a function of the shower zenith angle, with more inclined showers having a more plane-like structure.

The AGASA Collaboration handles these issues with parametrization both of the curvature of the shower front, and the fluctuations of the arrival time of the first particle as a function of core distance [9]. They quote arrival direction resolutions for showers with zenith angles less than 45° ranging from 4.2° at $10^{18.5}$ eV to 1.2° at 10^{20} eV (representing the 68% confidence radius) [10]. The Auger Collaboration is investigating the optimal reconstruction procedures, but early estimates of the angular resolution are given in [11]. For a mixture of iron and proton-initiated showers above 10^{19} eV, the 68% confidence limits vary from 1.1° for showers at a zenith angle of 20° to 0.3° for showers at 80°. The impressive improvement at large zenith angles is a result of the foreshortening of the array spacing at such angles, and a flatter and more well defined shower front (at these angles the majority of shower particles are muons which suffer very little scattering in their paths). On the other hand, the arrival direction resolution of gamma-ray initiated showers is less precise because of the lack of a sharp muon front, and ranges from 4.0° at 20° to 1.0° at a zenith angle of 80°.

Fluorescence detectors have a different set of challenges. When a shower is viewed by a single fluorescence site (a monocular or 'mono' observation), the first step in the geometrical reconstruction is to determine the 'shower-detector plane' (SDP), a plane in space defined by the shower axis and a point representing the detector. Then, the arrival time of light at each of the camera pixels is used to determine the orientation of the shower axis within the SDP. The timing fit can be quite robust if there is sufficient curvature in the tangent fitting function [6]. This is equivalent to saying that results are good if the apparent angular velocity of the shower image changes as it crosses the telescope field of view. Unfortunately this is not always observed, especially for events with short angular track length. In that case, a range of possible solutions is presented – for example, a distant shower with a component of velocity towards the detector can give a similar angular velocity to a closer shower moving away from the detector – and these possibilities cannot be distinguished. Here, the mono fit of the arrival direction produces a very elliptical error box on the sky, with the minor axis of the ellipse determined by the quality of the SDP fit (typically $\pm 0.5^\circ$) and with the major axis, determined by the quality of the timing fit, being up to several degrees long.

If more than one fluorescence site observes an air shower, the stereo technique can be applied. In this case the shower axis is simply the intersection of the two shower-detector planes, and resolution in the direction of the axis can be better than 0.5° in both dimensions. The resolution depends somewhat on the 'opening angle' between the two SDP's, but even in the case where this angle is not optimal, resolution approaching the benchmark 0.5° can be achieved by employing additional timing information from synchronized clocks at the two fluorescence sites.

An exciting new technique has had its first experimental tests with the Auger project. Here showers are viewed by surface and fluorescence detectors, and a new 'hybrid' technique is returning excellent reconstruction of the position and direction of the shower axis. The problem mentioned above of degeneracy in the fluorescence detector timing fit can be solved with an additional measurement, and in Auger this means a measurement of the arrival time of the shower at one or more surface detectors. In practice, a composite chi-squared function is minimized, using timing information from both the fluorescence and surface detectors [12]. Complications such as the shape of the shower front can be avoided, since good hybrid reconstruction is possible using the arrival time at the single detector measuring the largest particle density, together with the assumption of a plane shower front. The median angular error for 10^{19} eV showers is expected to be of order 0.4° for mono hybrid events, with little degradation below that energy [13]. Auger is expecting that up to 70% of hybrid events at 10^{19} eV will be viewed in fluorescence stereo, and the prospects with stereo-hybrid reconstruction are for even better angular resolution.

3. Results so far

At the highest energies, the number of experiments which have provided useful anisotropy data is very limited. In the north, the Haverah Park ground array experiment based on water Cherenkov detectors [14] has provided key data. However, that pioneering experiment had limited collecting area. Ongoing experiments, reviewed in [14], are in Japan (the AGASA scintillator-based ground array), the United States (HiRes, an array detecting UV atmospheric nitrogen fluorescence light produced by shower cascades), and Siberia (Yakutsk, a scintillator and air Cherenkov array). These experiments are continuing to provide evidence on both the broad angular scale and possible point source components in the data. High energy data in the south are sparse and so far have only been provided only by the SUGAR array [15]. SUGAR was based on widely-spread muon detectors and pioneered important techniques, but its data are generally not of the quality of those from the other experiments mentioned. However, the southern skies are regarded as particularly important as they contain some very significant potential sources and the Pierre Auger Observatory in Argentina is to search there first, while its northern observatory is being constructed.

3.1. Harmonic analysis

Early cosmic ray detectors had very poor angular resolution, often just the range of angles which resulted from increased atmospheric attenuation as the zenith angle increased. Since the zenith scans a full range of Right Ascension every sidereal day, it was natural to count events in intervals (perhaps 15 minutes) of sidereal time, accumulate a data-set of such counts over a few years, and then perform a fourier analysis in local sidereal time (Right Ascension at the zenith). The result was a series of Right Ascension 'harmonics' of the count rate (sky brightness) within a rather broad band of declination centred on the latitude of the detector. The angular resolution was so poor that only the first and second harmonics were usually considered. Those early experiments were at low energies relative to the presently studied upper end of the energy spectrum, characteristically 10^{14} eV to 10^{16} eV, and count rates were typically such that data-sets of a few hundred thousand events were readily accumulated. No evidence for bright point sources was found and harmonic analysis has continued to be a well used technique at all energies, to search for broad scale anisotropy down to low levels.

The harmonic analysis of data has a long history and the technique is well known. It is analogous to analysing a two dimensional random walk with the harmonic amplitude being the sum of N co-planar vectors each having magnitude 2. Linsley [16] summarized a number of important issues concerning the uncertainties in the amplitude and phase of a harmonic vector measured in this way. If the fractional amplitude of the first harmonic amplitude is r, the parameter $k_o = r^2 N/4$ defines the probability of obtaining an amplitude greater than r from a random population through $P(>r) = \mathrm{e}^{-k_o}$. There has been little evidence for any observed anisotropy amplitude deviating significantly from the random expectation, i.e., for $r^2 N/4$ greatly exceeding unity.

With data-sets containing over 10^5 events being common, statistical upper limits to first harmonic amplitudes of below 1% are found. This is a remarkable degree of isotropy. At levels of anisotropy as low as this, issues of the uniformity of exposure of the array in sidereal time become of major concern. These relate to seasonal effects of detector maintenance, seasonal variations

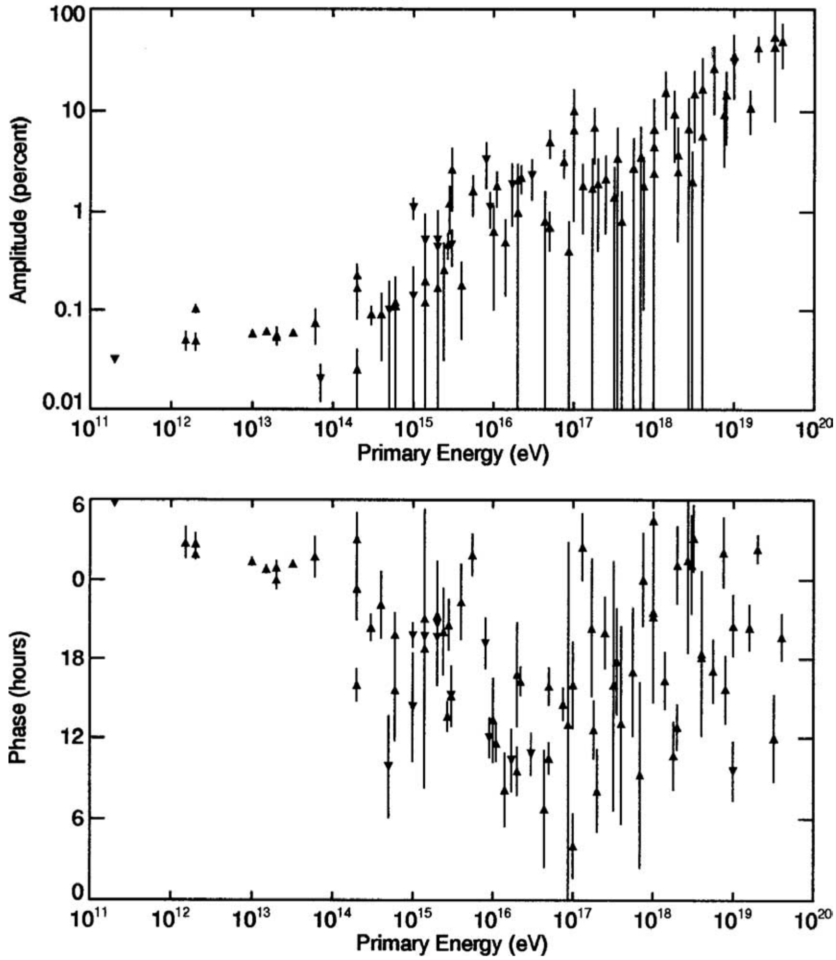

Fig. 1. The energy dependence of the amplitude and phase of the first harmonic of the cosmic ray anisotropy in Right Ascension (from [19]; first published by CSIRO). The error bars follow the analysis described in [16].

in atmospheric properties (such as temperature or pressure which may affect the detector energy threshold or analysis) etc. One cannot guarantee that these will average out over a sidereal year and analyses which seek to identify solar time components in the sidereal time analysis [17,18] are commonly used. Such procedures soon become complex (e.g., [19]).

Fig. 1 indicates the characteristics of the derived cosmic ray first harmonics in Right Ascension over the full range (ten decades) of cosmic ray energy. The data are for broad declination bands set by the array angular acceptances. It appears that there is a general trend of increasing amplitude with energy. This is an observational artefact due to the statistical limitations of cosmic ray data-sets, which measure a cosmic ray flux that reduces by thirty decades over the energy range of this plot. It illustrates the key role of observational statistics in working with such data.

There is no reason to expect that the measurement of anisotropies in terms of Right Ascension would represent the most rational physical analysis. If there is a particular source direction and a broad flux of particles diffuses past us, a dipole anisotropy will result. The possibility of such a flow at 10^{15} eV was discussed by Clay et al. [20]. If an analysis of a dipole anisotropy is in terms of first harmonics of right ascension, there is a directional ambiguity in the data which can then only be resolved by comparable data from another latitude (preferably the other hemisphere) and second harmonic information.

Another approach is to decide in advance some candidate poles of dipole anisotropies, and search for those anisotropies directly. An example is the recent analysis by the HiRes collaboration [21]. They have placed limits on the strength of dipoles with poles located at the Galactic Centre, M87 and Centaurus A, for cosmic ray energies above $10^{18.5}$ eV. Such an analysis requires a good understanding of the exposure of the detector over several years of operation so that the sky map expected for a truly isotropic cosmic ray flux can be evaluated. Like several other analyses, the HiRes group have employed a shuffling technique for this purpose (see below).

At slightly lower energies, a little above 10^{18} eV, the AGASA data-set is highly suggestive of a directional anisotropy from the general direction of our Galactic Centre [22]. Those data have characteristics which broadly suggest a dipole component plus a more collimated beam. Supporting evidence have been found in the data-sets of the Fly's Eye experiment [23] and the SUGAR array [24] (the only southern array, prior to the Pierre Auger Observatory, capable of observing the Galactic Centre region with good efficiency). While the SUGAR data support the idea of a collimated component, no data-set has yet provided statistically compelling evidence for such a source. Of course, the Galactic Centre direction would naturally be high on any source search list.

3.2. *Correlation functions and cluster searches*

As we have said, it appears that the broad directional distribution of the cosmic ray beam is close to uniform at all well studied energies. On the other hand it is possible that there may be directional excesses associated with point sources – either because the particles are energetic enough to avoid large deflections in cosmic magnetic fields, or because the particles are neutral (gamma-rays or neutrons). For example, at 10^{19} eV, time dilation allows neutrons to reach us from any region of our galaxy without decay. There is thus good reason to look for small-scale directional excesses, either associated with the directions of predicted sources or within a non-specific search.

The two-point angular autocorrelation function can be used for non-specific searches or a cross correlation function can be used with a second data set. These functions are based on taking the directions in one data-set in sequence and determining the angles between those directions and all the directions in the second data-set (or with the original set for the autocorrelation function). The numbers of angular separations between angles θ and $\theta + \delta\theta$ for a series of values of θ are compared with the expectation for an isotropic cosmic ray flux. That expectation, and the significance of any result in the real data set, may be evaluated using some variant of the shuffling technique (for an example see [25].)

The AGASA experiment has for some time reported possible clustering in arrival directions. Their data-set is the one of the two largest available for high energy cosmic rays, but is still sparse when plotted on a sky map. Intuitively, it would seem very unlikely that there would exist pairs of events in such a data-set that are separated by little more than the array angular measurement error (of the order of 2°). However, three such pairs were found out of a total of 36 events above 4×10^{19} eV in a 1996 analysis [26]. By 2003, in a 59 event set of data, one triplet, and five doubles were found with a claimed chance probability of less than 10^{-4} [27]. Clustering analysis is a particular example of an autocorrelation function search and the chance probability which is derived in any search depends on the number of independent statistical trials which have been carried out. That number may be difficult to determine, and hence become contentious. Finley and Westerhoff [28] provide a critical discussion of the analyses of these AGASA data and the trials involved, following a discussion by Tinyakov and Tkachev [29]. The underlying difficulty with such analyses of trials is that the data are often first viewed and then the analysis is determined. The AGASA 1996 analysis seemed to show clustering based on assumptions of appropriate angular uncertainties. Later analyses include the pre-1996 data-set for which the analysis was not clearly a priori. It is now difficult to determine the number of trials incurred in selecting energy and angular ranges for the analysis, though they were initially selected on an intuitive basis following the viewing of sky maps.

It is reasonable to expect that the highest energy cosmic rays would show some tendency to cluster around the major structures in our local Universe. Those structures are the galactic and supergalactic planes and they contain all the source objects normally considered. We clearly see such planar clustering when we view the Milky Way in visible light, when we view the sky in gamma-rays, and when we plot the positions of the most energetic astrophysical objects. Such a tendency to cluster can be searched for by examining the latitude distribution of cosmic ray arrival directions, either in galactic or super-galactic coordinates. An analysis performed by Stanev et al. [3] appeared to exhibit clustering associated with the super-galactic plane. The status of this result is now unclear as it is not supported by the current AGASA data (see the discussion in [14]) although SUGAR data may be supportive at the highest energies [30]. Such a result, either positive or negative, is a key datum since it would tell us whether ultra-high energy cosmic ray particles are capable of travelling to us undeflected through intergalactic fields. It is disturbing that there is disagreement on this matter.

One can look for clustering within a dataset, but it is also possible to search for correlations between a dataset and a source catalogue. Searches of this kind have been made with physically likely sources such as AGN's and the brightest radio sources, e.g., [31–33]. The results of such searches may be positive but the number of statistical trials involved is hard to quantify (see below).

4. Statistical considerations

Since any detected anisotropy is most likely to be (at least initially) at a marginal statistical significance, it is important to understand both the statistical properties of the distribution of event directions in the case of no anisotropy (i.e., the null hypothesis) and the total number of trials involved in selecting the parameters for all searches undertaken. The former may be calculated by the 'shuffling' technique, which will be discussed later in this section, or by other techniques to calculate the exposure/aperture. The latter requires an a priori determination so that the number of trials is fixed. Where the possible parameter space is large we may need to restrict the search space to limit the number of trials so that we maximize the power of the test. In other words, we may need to optimize the search, but this must be done a priori, based on a consideration of previous results and some assumptions about the likely physics. Such optimization will reduce our sensitivity to the unexpected and it is important to have a mechanism for incorporating suggestive results in a past dataset into the analysis of a future dataset. Results of the initial tests can thus be used to refine the tests themselves, but the modified tests can only be applied in a statistically meaningful way to *subsequent* data sets.

One variable to be defined in an anisotropy search is the energy range for the included events. In previous searches it has been usual to place the lower energy limit at about the GZK cut-off energy, so lower limits of $4, 5$ or 6×10^{19} eV have been common. There are two factors underlying this choice, both aimed at selecting a sample that minimizes the angular deflections in the magnetic fields. Super-GZK events are presumed to be preferentially from nearby sources, limiting their path length in the field. Also, the highest energy events will have the smallest deflections, given that the angular deviation of a particle (charge Z) varies as Z/E. The actual energy threshold chosen is usually a compromise between the arguments for a high energy threshold and the statistical requirement that the threshold be low enough to ensure a reasonable size data set. However, a 'the higher the better' approach is not necessarily the best and we will demonstrate later a situation where *lowering* the threshold is the optimal strategy.

For generic tests of isotropy, e.g., harmonic analysis or the 2-point correlation function, the angular scale of the test is implicit in the test itself. One is limited at the large scale by the dipole, and at the small scale by the angular reconstruction resolution. For tests which are more specific to a proposed source the angular scale is very much a free parameter. While the experimental reconstruction resolution again defines a lower limit to what might be useful, any choice of the optimal scale must necessarily involve some assumptions about the propagation of the cosmic rays.

The result of any anisotropy analysis must be tested against the distribution arising if the null hypothesis (i.e., isotropy) is true. Establishing this distribution with sufficient accuracy can itself be a challenge. As we noted above many factors can modify the expected null dataset. The true event arrival direction distribution (in the astronomical frame) is modified by the angular acceptance of the detector and its on-time. Environmental changes may produce energy dependent effects. Even a continuously operating detector will show diurnal and seasonal variations which will produce non-uniformities in the RA distribution of the detected events. For a fluorescence detector, with a typical 10% on-time, and high sensitivity to the weather, the non-uniformity will be particularly large.

The usual approach to determining the array exposure is to use the events themselves as a measure of detector performance. For example, the total event rate over the detector aperture can give a good measure of the detector performance at any time. One implementation of this type of method is the 'shuffling' technique [34] used by the Fly's Eye experiment which, since it uses a fluorescence detector, has a highly variable exposure. In shuffling, a 'background' sample is produced by randomly exchanging each observed event's arrival time with that of another event. The new events then have the same times (and so, on-time distribution) as the real data, but are distributed (semi-)randomly over the detector aperture in a way that reflects the local arrival direction distribution of all events. This process can be repeated to produce a large number of ensembles of 'isotropized' directions on which to evaluate the null hypothesis probability distribution for any anisotropy test [24]. However, such techniques introduce some feedback of any directional excess into the background calculation, thereby reducing the sensitivity. This is likely to be a greater problem where the excess has the same size scale as the detector aperture, such as for a dipole anisotropy.

It is usually necessary to establish a statistically significant anisotropy before we can ask the physics question 'What direction is the excess?' It follows that we should optimize for the detection of anisotropy, and not necessarily for directional accuracy. Below we will show that this may be best served by lowering the energy threshold of the search.

4.1. Optimizing a search – an example

Consider the case of a source for which the *systematic* magnetic deviation of the cosmic rays at earth is zero, but for which the *random* deviation has a characteristic scale of $\theta_m = \theta_{20}/E_{20}$, where E_{20} is $E/10^{20}$ eV and θ_{20} is the characteristic deviation at 10^{20} eV. If, above a minimum energy E, we count the number of events with arrival directions within a circle of radius θ_a and compare this with the expected number of 'background' events B, we will obtain an excess S due to the source. Now let us increase the minimum energy E and consider a change in the radius of the circle to take account of improved reconstruction and/or smaller assumed magnetic deflections at higher energies. Let the radius of the circle change as $\theta_a \propto 1/E^{\eta}$. (We assume that with this changing radius we have the same efficiency for collecting signal events above all values of minimum energy E.) We then have $S \propto F_s$ (where F_s is the source flux), and $B \propto F_b/E^{2\eta}$ (given the energy dependence of the circle radius). Strictly,

S and B are integrals over $E > E_{\mathrm{th}}$, but this does not alter the result for simple power law dependences. The significance of the excess in terms of standard deviations (σ) can be expressed as

$$\sigma = \frac{S}{\sqrt{B}} \propto \frac{F_s}{\sqrt{F_b E^{-2\eta}}}.$$

Using integral spectra, $F_b \propto E^{-\gamma_b}$ and $F_s \propto E^{-\gamma_s}$, this gives $\sigma \propto E^{\rho}$, with

$$\rho = -\gamma_s + \frac{\gamma_b}{2} + \eta.$$

Clearly, if $\rho < 0$ we should decrease the energy threshold to optimise the signal, while $\rho > 0$ argues for a high threshold. Putting in some rough numbers, *viz* $\gamma_b = \gamma_s = 1.7$ and $\eta = 1$ this gives,

$$\sigma \propto E^{+0.15}. \tag{1}$$

A value of $\eta = 1$ might be appropriate when accounting for the reduction in magnetic deflection with energy, but in our evaluation of the appropriate circle size we also need to take into account the detector reconstruction error, especially when searching for small-scale anisotropies. This uncertainty may account for a significant portion of the circle size, and in this case we might expect $\eta < 1$. Thus the significance σ has a weak positive increase with energy at best. Additionally, if we are performing an all-sky search, the number of trials involved in covering the sky goes as $1/\theta_a^2$, which reduces the after-trials significance at higher energies (assuming that $\eta > 0$). This suggests that there is no clear argument in favour of using the highest energy threshold that is practicable, and so the question of the optimum threshold should be carefully investigated.

We have explored this line of reasoning more fully using Monte Carlo simulations. For simplicity, we generate events uniformly over a nominal detector exposure. Some additional events are attributed to a point source, and have their arrival directions randomly deflected from the true source direction according to a binormal distribution based on a combination of a 'magnetic field' component, which varies as θ_{20}/E_{20}, and a realistic reconstruction error (based on typical Auger surface detector resolution). The event energies are sampled from a broken power law spectrum with a differential index of -2.7 below, and -5.0 above some break energy, such as may result from a GZK cut-off. We also consider cases where the apparent source position is systematically shifted from the true direction by $\theta_s \propto 1/E$.

The results are analysed by constructing sky maps. In the derivation leading to Eq. (1) above we simply count the number of event directions within a circle of a certain radius centered on the vertices of a grid (we call this the 'circle counting method'). For each vertex we counts a signal S which we compare with the expected background B to calculate a standard deviation $\sigma = (S - B)/\sqrt{B}$. In our simulation we use a more sophisticated technique where we create a sky map by replacing every nominal event direction by a probability distribution function (PDF) for the 'true' direction, and the map is the sum of all the PDFs (e.g., [24]). The 'true' direction may be that of the incident cosmic ray, in which case we would use the event reconstruction error function as the PDF, or it may be the direction to the source, in which case the PDF would also include some estimate of the magnetic deflection.

A true PDF has an integral of unit probability, and because high energy events are generally better reconstructed and/or suffer less deflection in magnetic fields, the PDF of a high energy event will both be narrower and have a higher peak value. We call a PDF with unit integral a 'normalized' PDF. We have also investigated a variant of this method for which the PDF is not normalized to unit integrated probability, but which has a constant value of 1 at the nominal direction of the event – here the integral of the PDF becomes larger for a wider PDF. This variant should produce results which are somewhat analogous to the circle counting method. We call this the 'un-normalized' PDF.

Finally, we introduce some variations of the normalized PDF case, where the test statistic ζ_n at each point on a sky grid is given by

$$\zeta_n = \sum p_i - n\sqrt{\sum p_i^2},$$

where p_i is the PDF of the ith event direction evaluated at the test position, and n is a parameter which has the effect of preferentially reducing the test statistic when the major contribution is by relatively few events. For the unmodified normalized PDF sum ($n = 0$), the sky map has the unfortunate characteristic of sometimes producing a significant peak when two high energy events (those with narrower PDFs) happen to arrive from nearby directions. In that case the power of the test for detecting real anisotropy is poor. We only show results for $n = 1$ and $n = 2$. For the event PDF we use the same binormal distribution that was used to generate the signal event deviations. Although we would not be able do do this for real data this optimization has only a minor affect on the *energy dependence* of the resultant significance.

We find the peak value of the test statistic near the injected source position. To evaluate its significance, we compare that value with a distribution of the global peak values for a large number of data sets with no signal. This gives us a significance that takes into account the number of trials expended in searching over the entire sky.

Here we present some results for a break energy of 2×10^{20} eV (i.e., no GZK cut-off). The assumed PDF is a (normalized or un-normalized) binormal distribution with an energy-dependent standard deviation of $2^\circ/E_{20}$. The samples contain a total of 10^4 events above 10^{19} eV. A signal of approximately 100 events is injected around a mean direction which may be shifted from the true source direction. The signal events are randomly distributed about the mean direction with a binormal distribution with a standard deviation which is also $2^\circ/E_{20}$.

In Fig. 2 we see the results for three cases where the systematic shift of the source direction is set to 0°, $2^\circ/E_{20}$ and $4^\circ/E_{20}$. In the top panel we see that the un-normalized PDF produces a weak increase in significance as the energy threshold is reduced. This is what we expect from Eq. (1) after allowing for the trials factor. The two analyses based on a normalized PDF have the best sensitivity, and both show increasing sensitivity as the energy threshold is lowered. Other simulations show that the relative sensitivity for the two different n values is dependent on the spectra of the signal and background but that $n = 2$ generally gives the better result. (We also note that the $n = 0$ case, not shown here, gives the same general trend of increasing significance at low energies, but not to the same extent as the other two normalized cases).

Fig. 3 shows the shift of the position of the test statistic maximum from the true source direction. For the normalized PDF variations the 'pointing accuracy' is only weakly degraded as the energy threshold is decreased.

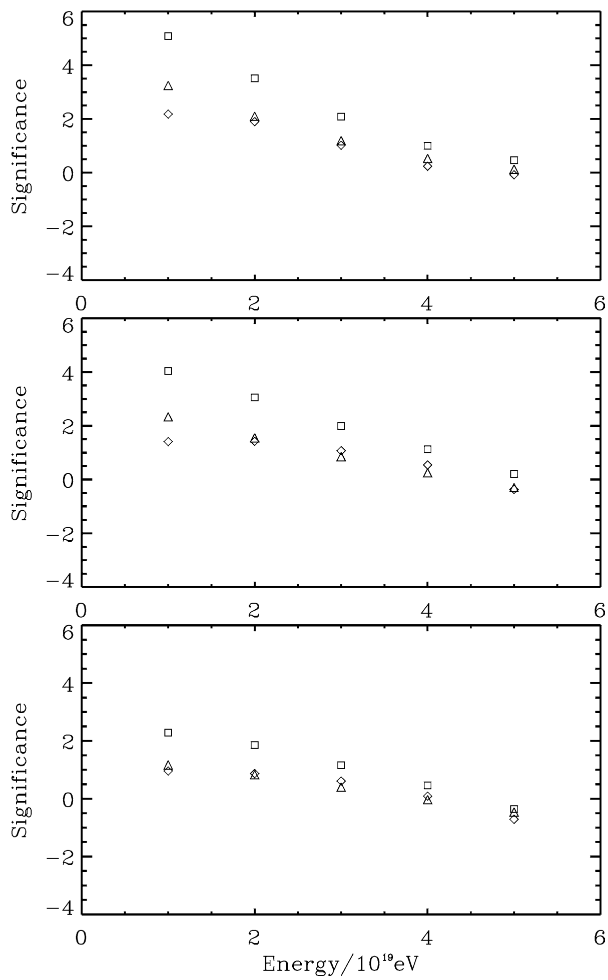

Fig. 2. The average significance of the injected signal for several test statistics as discussed in the text. They are: ◇ – un-normalized; △ – normalized $n = 1$; □ – normalized $n = 2$. The significance is the nominal number of standard deviations by which the value of the test statistic exceeds the average of the global maximum for simulations with no signal. The top panel shows the case where there is no systematic shift of the mean arrival direction of the signal events with energy. In the middle panel the shift is $2^\circ/E_{20}$, while in the bottom panel the shift is $4^\circ/E_{20}$. In all cases the signal events have a random shift about the mean direction that is described by a binormal distribution with a standard deviation of $2^\circ/E_{20}$.

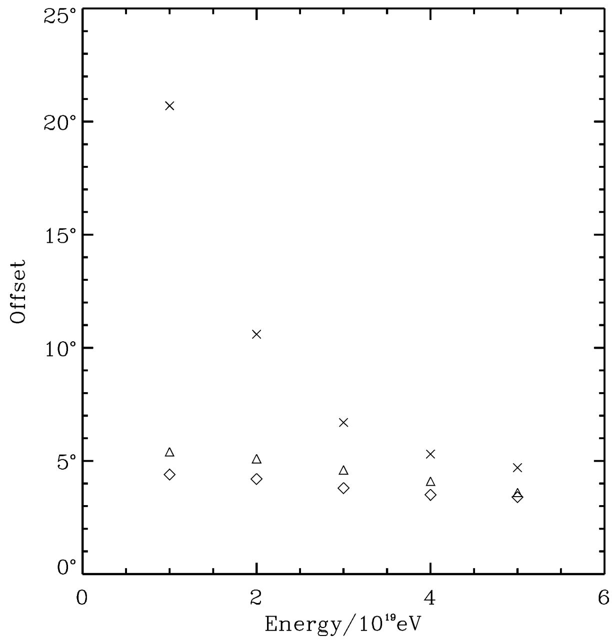

Fig. 3. The mean angular offset of the 'reconstructed' direction for the analysis methods using an un-normalized PDF ($\times$) or a normalized PDF ($\diamond - n = 1$; $\triangle - n = 2$) as a function of the threshold energy. The apparent source direction has a systematic shift of $4^\circ/E_{20}$. The offset for the un-normalized PDF closely follows the mean offset of the signal events.

In conclusion, we see that for the case of a small angular scale excess it is possible to maximize the detection sensitivity by *lowering* the energy threshold to as low as 10^{19} eV while retaining a pointing accuracy normally associated with a much higher threshold.

5. Considerations for a anisotropy analysis prescription

As we noted, the HiRes Collaboration has been careful to define its anisotropy search program a priori [35]. The Pierre Auger Collaboration has also been concerned about the statistical penalties associated with uncontrolled trials and has moved to develop its own a priori method for an ongoing anisotropy search program, such that the number of trials in the analysis of a dataset are known. This is referred to as a 'prescription' [36]. Their process of defining a prescription begins with the selection of a list of possible sources to be searched for in the data-set, including the particle energies, and angular ranges around the sources which are to be accepted, on the basis of array angular response and the expected physics of particle propagation. Decisions then have to be made on the weighting given to the various sources on the list and those weightings are used to allocate fractional probabilities in the search.

As an example, two target sources might be defined but one may be thought to be the more likely source. A conventional level for claiming a positive result would be a total chance probability of the result being less than 0.001. In the prescription, that probability is then partitioned between the two candidate sources such that each is assigned a chance probability for a discovery result, with those two chance probabilities summing to 0.001. Thus, the two source searches might be assigned 0.0007 and 0.0003 for their contributions to any resulting source discovery. The more likely source is given the higher, more accessible, limit to reach and the less preferred source is given a more difficult hurdle, but one which may be achieved in the unexpected circumstance that the latter source proves to be real. This concept can be extended to any number of sources but, of course, each additional source search incurs a penalty which is reflected in a progressive reduction of the prescription probability levels which have to be achieved (the degree of difficulty increases) to claim a positive result.

A prescription includes a definition of the dataset to be analysed and this is the ONLY analysis which can be used on that data-set for discovery purposes. However, the process is progressively iterative. A dataset can still be searched for any source or pattern and, although any discovery could not be claimed (because of there being an unknown number of trials), it can point the way for a prescription to be applied to the next data-set. For the Pierre Auger Project, prescriptions may well be updated on an annual basis.

6. Conclusion

Observations of the highest energy cosmic rays have demonstrated that Nature is capable of accelerating particles to energies well beyond those attainable in man-made accelerators and, also, well beyond those predicted by our present source modelling. A key to discovering the real sources of cosmic rays is likely to be through careful observation of their arrival directions. This objective has so far been frustrated as the cosmic ray beam has proved to be isotropic beyond any expectations, even at the highest energies. It has become clear that, in moving forward, directional studies will require an unprecedented degree of care in their statistical design, as we have indicated above. A complete understanding will require the combination of such studies with studies of the strength and structure of galactic and inter-galactic magnetic fields. Additionally, the modelling of possible sources is needed in order that we can suitably plan anisotropy searches at these unexpectedly low levels. We need to keep in mind that propagation times, in some intergalactic magnetic field models, may exceed source lifetimes and the sources of presently observed cosmic rays may be unknown because they have passed their active phase.

References

[1] R.J. Protheroe, R.W. Clay, Pub. Astron. Soc. Aust. (2004) in press, astro-ph/0311466.
[2] K. Greisen, Phys. Rev. Lett. 16 (1966) 748;
G.T. Zatsepin, V.A. Kuz'min, Pis'ma Zh. Eksp. Teor. Fiz. 4 (1966) 114;
G.T. Zatsepin, V.A. Kuz'min, JETP Lett. 4 (1966) 78.
[3] T. Stanev, Propagation of ultra high energy cosmic rays, C. R. Physique (2004) in press.
[4] K. Dolag, D. Grasso, V. Springel, I. Tkachev, astro-ph/0310902.
[5] G. Sigl, F. Miniati, T. Ensslin, astro-ph/0401084.
[6] P. Sommers, Extensive air showers and measurement techniques, C. R. Physique (2004) in press.
[7] Auger Collaboration, Nucl. Instr. Methods, in press.
[8] M. Fukishima, et al., The telescope array experiment: an overview and physics aims, in: Proc. 28th Int. Cosmic Ray Conf., Tsukuba, Japan, 2003, p. 1025, http://www-rccn.icrr.u-tokyo.ac.jp/icrc2003/proceedings_pdf.html.
[9] S. Yoshida, et al., Astropart. Phys. 3 (1995) 105.
[10] M. Takeda, et al., Astrophys. J. 522 (1999) 225.
[11] M. Ave, J. Lloyd-Evans, A.A. Watson, in: Proc. of 27th Int. Cosmic Ray Conf., vol. 2, Hamburg, 2001, p. 707.
[12] B.R. Dawson, H.Y. Dai, P. Sommers, S. Yoshida, Astropart. Phys. 5 (1996) 239.
[13] B.R. Dawson, P. Sommers, in: Proc. of 27th Int. Cosmic Ray Conf., vol. 2, Hamburg, 2001, p. 714.
[14] M. Nagano, A.A. Watson, Rev. Mod. Phys. 72 (2000) 689.
[15] M.M. Winn, et al., J. Phys. G 12 (1986) 653.
[16] J. Linsley, Phys. Rev. Lett. 34 (1975) 1530;
J. Linsley, in: Proc. 14th Int. Cosmic Ray Conf., vol. 2, Munich, 1975, p. 592.
[17] F.J.M. Farley, J.R. Storey, Proc. Phys. Soc. A 67 (1954) 996.
[18] F.J.M. Farley, J.R. Storey, Proc. Phys. Soc. B 70 (1957) 840.
[19] A.G.K. Smith, R.W. Clay, Aust. J. Phys. 50 (1997) 827.
[20] R.W. Clay, et al., Publ. Ast. Soc. Aust. 15 (1998) 208.
[21] R. Abbasi, et al., The HiRes Collaboration, Astropart. Phys., in press, astro-ph/0309457.
[22] N. Hayashida, et al., Astropart. Phys. 10 (1999) 303.
[23] D.J. Bird, et al., Astrophys. J. 511 (1999) 739.
[24] J.A. Bellido, R.W. Clay, B.R. Dawson, M. Johnston-Hollitt, Astropart. Phys. 15 (2001) 167.
[25] L.J. Kewley, R.W. Clay, B.R. Dawson, Astropart. Phys. 5 (1996) 69.
[26] N. Hayashida, et al., PRL 77 (1996) 1000.
[27] M. Teshima, et al., in: Proc. 28th Int. Cosmic Ray Conf., Tsukuba, Japan, 2003, p. 437, HE1.3, http://www-rccn.icrr.u-tokyo.ac.jp/icrc2003/proceedings_pdf.html.
[28] C.B. Finley, S. Westerhoff, Astropart. Phys., submitted for publication, astro-ph/0309159.
[29] P.G. Tinyakov, I.I. Tkachev, JETP Lett. 74 (2001) 1, astro-ph/0102101.
[30] R.W. Clay, in: Proc. 28th Int. Cosmic Ray Conf., Tsukuba, Japan, 2003, p. 429, HE1.3.
[31] R.W. Clay, B.R. Dawson, L. Kewley, M. Johnston-Hollitt, Pub. Astron. Soc. Aust. 17 (3) (2000) 207.
[32] P.G. Tinyakov, I.I. Tkachev, JETP Lett. 74 (2001) 445.
[33] F. Ferrer, S. Sarkar, Phys. Rev. D 67 (2003) 103005.
[34] G.L. Cassiday, et al., Nucl. Phys. B Proc. Suppl. A 14 (1990) 291.
[35] J.A. Bellido, et al., in: Proc. 27th Int. Cosmic Ray Conf., vol. 2, Hamburg, 2001, p. 364.
[36] R.W. Clay, et al., in: Proc. 28th Int. Cosmic Ray Conf., Tsukuba, Japan, 2003, p. 421, HE1.3, http://www-rccn.icrr.u-tokyo.ac.jp/icrc2003/proceedings_pdf.html.