Available online at www.sciencedirect.com
SCIENCE DIRECT®
C. R. Physique 5 (2004) 463–472
COMPTES RENDUS PHYSIQUE

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

Extensive air showers and measurement techniques

Paul Sommers

High Energy Astrophysics Institute, Physics Department, University of Utah, Salt Lake City, UT 84112, USA

Available online 24 April 2004

Presented by Pierre Encrenaz

Abstract

A cosmic ray observatory records individual particles using the atmosphere as a transducer and amplifier. Each extremely high energy cosmic ray converts into an air shower that grows to billions of secondary particles. A large cosmic ray observatory senses all such air showers landing in a collection area that spans thousands of square kilometers. Measuring an air shower in detail yields an accurate arrival direction, a good energy estimate, and a likelihood distribution for the mass of the primary cosmic ray. Modern 'hybrid' observatories combine a surface array of particle detectors with telescopes that observe radiation produced by the developing shower front as it traverses the atmosphere. ***To cite this article: P. Sommers, C. R. Physique 5 (2004).***

Résumé

Les techniques de mesure des gerbes atmosphériques. Un observatoire de rayons cosmiques détecte des particules individuelles par leurs interactions avec l'atmosphère, milieu dans lequel le signal est amplifié par un grand nombre de réactions en cascade. Chaque rayon cosmique, s'il a une énergie très élevée, se transforme en une *gerbe atmosphérique* constituée de milliards de particules secondaires. Un grand observatoire de rayons cosmiques détecte de telles gerbes en instrumentant des surfaces de collecte s'étendant sur plusieurs kilomètres-carré. L'étude détaillée d'une gerbe permet de remonter à la direction d'arrivée et à l'énergie du rayon cosmique primaire, et d'obtenir une distribution de vraisemblance quant à sa nature. Les observatoires « hybrides » les plus récents couplent deux types de détection indépendants : un réseau dit « de surface » constitué de détecteurs de particules et des « télescopes » optiques qui collectent la lumière produite par le front de la gerbe à mesure qu'elle se développe dans l'atmosphère. ***Pour citer cet article : P. Sommers, C. R. Physique 5 (2004).***

Keywords: Air showers; Detection techniques; Chemical composition; Longitudinal profile; Gaisser–Hillas function; Superposition model

Mots-clés : Gerbes atmosphériques ; Techniques de détection ; Composition chimique ; Profil longitudinal ; Fonction de Gaisser–Hillas ; Modèle de superposition

1. Introduction

An extremely high energy cosmic ray collides with a nucleus high in the atmosphere. The interaction produces many new energetic particles. Those also collide with air nuclei, and each collision adds a large number of particles to the developing cascade. Some of the produced particles are neutral pions, each one of which immediately decays to a pair of gamma rays. The gamma rays produce $e^{\pm}$ pairs when passing near nuclei. The electrons and positrons re-generate gamma rays via bremsstrahlung, thereby building the electromagnetic cascade. This is an extensive air shower.

E-mail address: sommers@physics.utah.edu (P. Sommers).

1631-0705/$ – see front matter
doi:10.1016/j.crhy.2004.03.009

The number of charged particles in the air shower reaches a *maximum size* $N_{\max}$ that is nearly proportional to the primary energy E. There are billions of charged particles in high energy air showers. The size $N_{\max}$ at shower maximum is approximately equal to $E/(1.6\ \text{GeV})$, although this conversion factor depends slightly on the choice of hadronic interaction model that is adopted to simulate collisions at energies above the reach of collider experiments, and it depends slightly on the atomic mass of the cosmic ray.

The good news is that the atmosphere is a natural transducer and amplifier for individual primary cosmic rays at extremely high energies. The air shower cascades make it possible to measure the flux of these very rare particles. Above 10^{20} eV, the cosmic ray intensity is not more than about 0.5 particles/km^2/sr *per century*. Direct detection of those primary particles above the atmosphere is obviously not feasible. Air shower detectors, on the other hand, can record all arriving particles over detection areas that are thousands of square kilometers.

The bad news is that detecting secondary shower particles (or atmospheric light that they cause) provides only indirect information about the primary cosmic ray. It is not possible to measure directly the primary particle's arrival direction, energy, or mass. We must deduce these quantities indirectly from detailed measurements of the particle's air shower.

Following an overview of air shower physics in Section 2, subsequent sections introduce the techniques used to reconstruct the primary particle's arrival direction, energy, and mass. The reconstruction methods for surface arrays differ from those used for detectors that record the shower development in the atmosphere above the ground. Hybrid analysis procedures exploit these complementary methods when both types of detectors have recorded the same shower.

2. Air shower cascades

Experimental evidence so far indicates that extremely high energy cosmic rays are predominantly atomic nuclei (including protons, i.e., hydrogen nuclei). The cosmic ray initiates its air shower cascade by hadronic interaction with an atomic nucleus in the atmosphere. The hadronic cascade (mostly pions) grows until the energy per pion falls to the level where pions are likely to decay before colliding. In each generation of the hadronic cascade, 1/3 of the energy on average goes to neutral pions which instantly decay to pairs of gamma rays. Each gamma ray develops an electromagnetic subcascade. After n hadronic cascade generations, only $(2/3)^n$ of the total energy remains in the hadronic cascade. The rest has been lost to the cumulative electromagnetic cascade. The $e^{\pm}$ particles eventually dissipate almost all of the original cosmic ray's energy by ionizing atoms along their paths.

Charged pions produce air shower muons when they decay. The number of shower muons depends on the amount of energy that is left in the hadronic cascade when pion energies have dropped to the level where decay is more likely than collision. If this happens after relatively few cascade generations, then copious muon production occurs. If the reduction of pion energies takes relatively many generations, then more of the energy will have been lost from the hadronic cascade to the electromagnetic cascade, and meager muon production occurs.

A cascade initiated by an iron nucleus develops like a superposition of 56 nucleons, each with 1/56 of the primary energy. In effect, this jump starts the cascade, and pions get down to energies where they can decay to muons before the electromagnetic cascade has drained too much energy from the hadronic cascade. An iron shower therefore typically has more muons than a proton shower of the same total energy. Moreover, the superposition of 56 lower energy subshowers reaches its maximum size higher in the atmosphere than a proton shower of the same total energy. Statistical determinations of the primary mass distribution (chemical composition) exploit these differences between heavy and light nuclei: heavy nucleus showers produce more muons and they reach maximum size higher in the atmosphere. Sections 2.4 and 6 explore these differences in more detail. See also [1].

2.1. Electromagnetic showers

The decay of π^0 mesons into gamma rays eventually transfers most of the primary cosmic ray's energy to the electromagnetic cascade. Each gamma ray converts to an $e^{\pm}$ pair. The electrons and positrons create new gamma rays by bremsstrahlung. The *radiation length* X_0 is the grammage path length in which their energies attenuate by the factor $1/e$, which is approximately the same as the $1/e$ attenuation of a gamma-ray beam due to pair production. In air, the radiation length X_0 is about 37 g/cm^2. The electromagnetic cascade grows via pair production and bremsstrahlung.

Heitler's heuristic picture [2] of the electromagnetic cascade gives intuitive understanding of its essential properties. One imagines the cascade developing by a sequence of generations. At each generation, every existing gamma ray converts to an $e^{\pm}$ pair, while each existing electron or positron produces a gamma ray in addition to itself. Every generation therefore doubles the number of cascade particles. The grammage interval for each generation is $X_0 \times \ln(2)$, i.e., the path over which the energy of any one particle is expected to be reduced by 1/2. The process continues until the average particle energy is reduced to the *critical energy* below which charged particles lose their energy by ionizing atoms in less than one radiation length. Since ionization

energy loss is about 2.2 MeV/g/cm^2, this critical energy E_c is (2.2 MeV/g/cm^2) × (37 g/cm^2) = 81 MeV. When the particle energies have fallen to E_c, the shower size is $N_{\max} = E/E_c$. The number of generations n needed to reach this maximum size depends on the total energy E. Since the number of particles doubles at each generation, one has at maximum, $2^n = N_{\max} = E/E_c$, so $n = \ln(E/E_c)/\ln(2)$. The maximum size occurs at a slant depth $X_{\max} = n \times X_0 \times \ln(2) = X_0 \times \ln(E/E_c)$ measured along the shower axis from the original particle's point of origin.

Rigorous treatments show that this heuristic model gives the correct depth of maximum ($X_{\max}$) for each energy. In particular, the depth of maximum $X_{\max}$ for electromagnetic cascades increases by $X_0 \times \ln(10)$ for each decade of increase in energy E. This *elongation rate* of 85 g/cm^2/decade is greater than what is expected for air showers that are fed by hadronic cascades.

The *Greisen formula* [3] accurately gives the electromagnetic longitudinal profile, i.e., the number of charged particles as a function of depth:

$$N_e = \frac{0.31}{\sqrt{T_{\max}}}\, e^T\, s^{-3T/2}.$$

Here T is the atmospheric slant depth measured in radiation lengths ($T = X/X_0$) from the original particle's point of origin, $T_{\max} \equiv \ln(E/E_c)$, and s is the *shower age*: $s \equiv 3T/(T + 2T_{\max})$. Many shower properties are well parametrized by the shower age. All positive values of atmospheric depth occur in the range $0 < s < 3$, and shower maximum occurs at shower age $s = 1$.

2.2. Hadronic shower longitudinal profiles

Extremely high energy nuclei produce air showers that are similar to electromagnetic air showers of equal energy. The difference is that there is a hadronic cascade (consisting mostly of pions) which spawns electromagnetic subshowers through $\pi^0 \to \gamma\gamma$ decays. The electromagnetic cascade gains energy at the expense of the hadronic cascade until the low-energy charged pions decay and give away the remaining hadronic energy to muons and neutrinos. Less than 10% of the primary energy escapes the electromagnetic cascade that way. The shower development at high energies tends to be faster than electromagnetic (Greisen formula) development due to the high inelasticity and multiplicity of hadronic interactions that distribute the primary energy among many particles faster than occurs in electromagnetic cascades described in the previous section.

Unlike electromagnetic showers, it is not yet possible to model hadronic cascade development with confidence at the highest energies. Laboratory experiments have studied particle collisions only up to energies of about 2 TeV in the center of mass. A 10^{20} eV proton colliding with a nucleon in the air is a center-of-mass collision energy of 450 TeV. Moreover, air shower development depends strongly on particle production in the very forward direction that is poorly studied in collider experiments.

In principle, detailed air shower measurements offer an exciting opportunity to constrain hadronic interaction physics at energies beyond the reach of accelerator experiments. The challenge is to glean definitive results without a priori knowledge of the primary particle types and their energies.

The *Gaisser–Hillas functional form* [4] has proved to be effective in fitting measured air shower developments and also simulated air shower developments resulting from various hadronic models (QGSJET, Sibyll, etc.) with variable primary masses. The form is

$$F(x) = F_{\max}\left(\frac{x}{w}\right)^w e^{w-x}.$$

Here x is shower depth measured in units of a typical interaction length λ (in g/cm^2) relative to a reference depth X_0, so $x \equiv (X - X_0)/\lambda$. (The symbol X_0 is unfortunately in standard use for the electromagnetic radiation length as well as this Gaisser–Hillas parameter. Its meaning must be determined from its context.) The Gaisser–Hillas profile 'width' w is the difference between $X_{\max}$ and X_0 measured in units of λ: $w \equiv (X_{\max} - X_0)/\lambda$. The four parameters $(F_{\max}, X_{\max}, X_0, \lambda)$ provide ample size and shape freedom for fitting longitudinal profiles. For measured showers lacking rich profile data, it is sensible to fix λ (e.g., $\lambda = 70$ g/cm^2) and do a 3-parameter fit. The profile amplitude $F_{\max}$, depth of maximum $X_{\max}$, and the single profile width parameter $X_{\max} - X_0$ are adequate in such cases. Freeing the 4th parameter λ is useful only if the data admit detailed fitting of the profile's rise and fall separately. Showers measured only near $X_{\max}$ might best be fitted using only the two parameters $F_{\max}$ and $X_{\max}$, setting $X_0 = 0$ as well as $\lambda = 70$ g/cm^2 (or other preferred fixed values). If the data cannot adequately determine the profile width, an unreasonable result may occur by including X_0 amongst the fitted parameters. The X_0 parameter should not be interpreted as the depth of first interaction. QGSJET and Sibyll simulations, like real measured showers, give longitudinal profiles that are typically best fit with negative X_0 values.

The longitudinal profile of a shower normally means the number of charged particles $N_e(X)$ as a function of atmospheric depth. In that case $F_{\max}$ is the shower size at maximum $N_{\max}$. A fluorescence light measurement as a function of depth relates more closely to the energy deposition rate $\mathrm{d}E/\mathrm{d}X$, and the integral of that function gives directly the total energy deposited in the atmosphere. 'Longitudinal profile' now frequently pertains to the energy deposition rate $F \equiv \mathrm{d}E/\mathrm{d}X$ as a function of atmospheric depth.

The integrated Gaisser–Hillas function has a closed form in terms of the standard gamma function:

$$\int_0^\infty F_{\max}\left(\frac{x}{w}\right)^w e^{w-x}\lambda\,\mathrm{d}x = F_{\max}\lambda\left(\frac{e}{w}\right)^w \int_0^\infty x^w e^{-x}\,\mathrm{d}x = F_{\max}\lambda\left(\frac{e}{w}\right)^w \Gamma(w+1).$$

If $F(X)$ is the function $\mathrm{d}E/\mathrm{d}X$, then this integral function gives the total electromagnetic shower energy. If $F(X)$ is the number of charged particles $N_e(X)$, then multiplying the integral by 2.2 MeV/g/cm^2 (the typical energy deposition rate per particle) gives the electromagnetic shower energy approximately.

2.3. Lateral distribution of shower particles

Energetic secondary hadrons have transverse momenta that are very small compared to their longitudinal momenta. They travel close to the shower axis and essentially parallel to it. The gamma rays from neutral pion decays are similarly directed along the shower axis. Electromagnetic particles spread out from the axis primarily by multiple Coulomb scattering of electrons and positrons. For that analysis, distance from the axis is most naturally measured as a multiple of the *Moliere unit*: $R_M \equiv X_0 \times (E_s/E_c)$ [5], where the 'scale energy' E_s is obtained from the fine structure constant and electron rest mass energy by $E_s = \sqrt{4\pi/\alpha}\cdot m_e c^2 \approx 21$ MeV. (X_0 again denotes the radiation length in air.) Using $E_c = (2.2\ \mathrm{MeV/g/cm^2})\cdot X_0$, the Moliere unit is $R_M = 21/2.2 = 9.5\ \mathrm{g/cm^2}$. In the following formulas, r is measured in Moliere units. The distance in meters should be obtained by dividing by the air density at an altitude about two radiation lengths back along the shower axis from the surface [6]. The Moliere radius two radiation lengths above the Auger altitude is just over 100 meters.

For pure electromagnetic showers, the *NKG function* [7,3] gives the charged particle density as a function of r with dependence on shower age s:

$$\rho(r) = \frac{N(s)}{R_M^2}\frac{r^{s-2}(1+r)^{s-4.5}}{2\pi B(s, 4.5-2s)},$$

where $N(s)$ is the total number of charged particles at age s and $B(x, y) \equiv \Gamma(x)\Gamma(y)/\Gamma(x+y)$ is the standard 'beta function'.

There is no standard lateral distribution functional form for muon density. Muons are certainly distributed more broadly than electromagnetic particles, and their number does not decrease rapidly as the shower grows old. They come directly from the hadronic cascade, where unknown properties of hadronic interactions at the highest energies can affect the longitudinal development. This induces some model dependence also in the expected lateral distribution of muons. Even in the context of a single hadronic interaction model, the shape of the lateral distribution function may depend on the detector's muon energy threshold as well as the distance from the point of maximum muon production to the core at ground level.

A lateral distribution functional form is azimuthally symmetric and pertains to the perpendicular distance from the shower axis. Especially for inclined showers (far from vertical), the density on the ground will not be azimuthally symmetric about the core location because (i) the distance from the core along the ground is different from the perpendicular distance to the shower axis; (ii) the shower age at ground level differs for different azimuths; and (iii) charged particle deflection by the geomagnetic field disrupts azimuthal symmetry.

2.4. Muons and $X_{\max}$ as indicators of atomic mass

The discussion in this section exploits the *superposition model*: an air shower initiated by a nucleus of atomic mass A and energy E behaves like the superposition of A proton air showers, each with energy E/A. To be perfectly clear, the picture for an iron nucleus is that 56 protons enter the atmosphere together and interact independently, *not* that the iron nucleus splits into 56 nucleons in a first interaction. The superposition model is not the ultimate in realism. It does, however, give quite accurate predictions and provides a conceptual way to understand how expected values for air shower parameters like depth of maximum $X_{\max}$ and total muon number N_μ should depend on the atomic mass of the primary cosmic ray.

As the sum of many equivalent subshowers, any shower variable is expected to have less fluctuation for heavy nuclei than for protons. In particular, the $X_{\max}$ for a heavy nucleus shower is statistically stabilized near the expected $X_{\max}$ for each of its subshowers. The distribution of $X_{\max}$ values for a pure composition of iron nuclei should therefore be much narrower than the distribution of $X_{\max}$ values for a pure proton composition. The shape of the $X_{\max}$ distribution at one energy can thus be an important handle on composition, even when interaction model uncertainty may preclude any clear inference from the mean value of that measured distribution. Similar remarks apply to the expected distribution for N_μ or the charged particle density at a given core distance. More fluctuation is expected in proton showers than in iron showers. Every shower variable is stabilized for heavy nuclei by statistical averaging that is implicit in the superposition model.

2.4.1. $X_{\max}$ dependence on primary mass

The *elongation rate for protons*, like the elongation rate for electromagnetic showers (Section 2.1), is the change in mean $X_{\max}$ per decade in energy. For two different primary proton energies, E_0 and E_1, the elongation rate κ gives the expected difference in $X_{\max}$ values as

$$X^p_{\max}(E_0) - X^p_{\max}(E_1) = \kappa \log_{10} \frac{E_0}{E_1}.$$

In accordance with the superposition model, a primary nucleus of energy E_0 with A nucleons behaves as a superposition of A proton showers, each with energy $E_1 = E_0/A$. The mean $X_{\max}$ for proton showers of energy E_0 should therefore differ from the mean $X_{\max}$ for primaries of energy E_0 and mass A by

$$X^p_{\max} - X^A_{\max} = X^p_{\max}(E_0) - X^p_{\max}(E_1) = \kappa \log_{10} A.$$

In particular, the air shower of an iron primary ($A = 56$) should reach maximum size higher in the atmosphere by 1.75κ than a proton shower of equal energy. The high energy hadronic cross section and inelasticity govern the proton elongation rate κ. Models having elongation rates in the range 45–55 g/cm^2/decade at high energy yield values for the expected $X_{\max}$ difference between protons and iron in the range 80–100 g/cm^2.

2.4.2. Muon number dependence on primary mass

Using the superposition model, it is only necessary to understand how muon production depends on primary energy for proton showers in order to calculate the expected difference in muon number for a primary of mass A and a primary proton of the same energy. Let $\beta \equiv d \ln N^p_\mu / d \ln E$ be the *muon power index* for protons defined at any primary energy E, so $N^p_\mu = \alpha E^\beta$ near that energy E for some constant α. According to the superposition model, the number of muons expected from a shower of energy E and mass A is

$$N^A_\mu = A \times \alpha (E/A)^\beta = A^{1-\beta} N^p_\mu.$$

To evaluate the iron/proton muon ratio $N^{\rm Fe}_\mu / N^p_\mu = 56^{1-\beta}$ we need to determine β at extremely high energies (between E/56 and E). The particle multiplicity in high energy collisions governs that muon power index for protons, as follows. After n generations of hadronic interactions, the energy remaining in the hadronic cascade is $(2/3)^n E$, since 1/3 of the hadronic energy is lost to the electromagnetic cascade by neutral pion decay at each generation. Muons are produced when pion energies are reduced to an energy E_d where decay is more likely than interaction. The number of produced muons is roughly the remaining hadronic energy divided by this energy per decaying pion:

$$N_\mu = \left(\frac{2}{3}\right)^{n_d} \frac{E}{E_d},$$

where n_d is the number of generations needed to degrade pion energies down to E_d. This gives an expression for β:

$$\beta \equiv \frac{d \ln N_\mu}{d \ln E} = 1 + \ln\left(\frac{2}{3}\right) \frac{d n_d}{d \ln E}.$$

The number of generations n_d increases with primary energy E. To a first approximation, the primary energy increases by the factor $(3/2)N_{\rm ch}$ when n_d increases by 1. (Here $N_{\rm ch}$ is the charged particle multiplicity and the factor 3/2 is because energy is shared also with neutral pions.) It follows that $\ln E$ increases by $\ln(\frac{3}{2}N_{\rm ch})$ for a unit increase in n_d, where $N_{\rm ch}$ is the charged particle multiplicity at the primary energy E. Then

$$\frac{\Delta \ln E}{\Delta n_d} = \ln\left(\frac{3}{2} N_{\rm ch}\right) \quad \Rightarrow \quad \frac{d n_d}{d \ln E} = \frac{1}{\ln(\frac{3}{2}N_{\rm ch})},$$

$$\beta = 1 + \frac{\ln(2/3)}{\ln(\frac{3}{2}N_{\rm ch})} = \frac{\ln(N_{\rm ch})}{\ln(N_{\rm ch}) + \ln(3/2)}.$$

For moderate collision energies, $N_{\rm ch} \sim 10$ and $\beta \approx 0.85$, giving $N^{\rm Fe}_\mu / N^p_\mu = 1.8$. At extremely high energies, however, the multiplicity is much higher, in the range 200 to 600. For $N_{\rm ch} = 200$, one gets $\beta = 0.93$ and $N^{\rm Fe}_\mu / N^p_\mu = 1.32$, while $N_{\rm ch} = 600$ gives $\beta = 0.94$ and $N^{\rm Fe}_\mu / N^p_\mu = 1.27$. Very high multiplicity yields the following result: *there are approximately 30% more muons in iron showers than in proton showers of equal energy, with only weak dependence on model differences in the actual high multiplicity value.*

3. Types of detectors

There are two general classes of air shower detectors: those that sample the flux of secondary particles at ground level and those that record radiation from the shower front as it traverses the atmosphere. Surface arrays include muon detector arrays (e.g., SUGAR), scintillator arrays (e.g., Volcano Ranch, Yakutsk, AGASA, and the Telescope Array – TA – project), and water Cherenkov tank arrays (e.g., Haverah Park and Auger). Atmospheric detectors can measure the longitudinal development of the shower. Those include nitrogen fluorescence detectors (e.g., Fly's Eye, HiRes, Auger, TA and the satellite projects EUSO, and OWL), air Cherenkov detectors (e.g., the Yakutsk calibration system), and perhaps advanced radio frequency antenna arrays (e.g., the LOPES development array at Karlsruhe [8]).

The two classes of detector types have complementary strengths and weaknesses. Contemporary cosmic ray observatories now under construction (Auger and TA) have hybrid designs in which some showers can be measured simultaneously by both detector types on clear dark nights. The following sections indicate the methods that are used by surface detectors (SDs), fluorescence detectors (FDs), and hybrid detectors to measure cosmic ray properties: arrival directions, energies, and atomic masses.

4. Measuring arrival directions

Surface arrays determine the arrival direction by recording the arrival time of the shower front at 3 or more non-collinear stations on the ground. The method is conceptually simple assuming the shower front to be a perfect plane. Any pair of stations A and B determine the arrival direction cosine along the direction from A to B as $c(t_A - t_B)/\overline{AB}$, where c is the speed of light, t_A and t_B are the trigger times for stations A and B, respectively, and $\overline{AB}$ is the distance between them. Two independent direction cosines determine a unique arrival direction in the hemisphere above the plane of the detectors.

The shower front is actually a curved surface, not a plane. The trigger times of 3 or more stations give the geometry by chi-square minimization, using the expected relative arrival times based on a realistic curved shower front moving at the speed of light. Those expected times depend not only on the arrival direction but also on the core position, so the core must first be determined from the relative station signal amplitudes.

Geometric reconstruction is quite different for a fluorescence detector eye, which sees an air shower as a spot of light that moves downward through the atmosphere at the speed of light. The track of the spot's center defines a great circle in the direction space of the eye which, together with the eye's location, determines the *shower-detector plane* (SDP). The SDP is the plane that contains the line of the shower axis and the point where the FD resides. If two eyes at different locations record the same shower, then the shower axis must lie in both SDPs. The intersection of the planes determines the shower axis (provided the two planes are not the same). This is the *stereo* reconstruction method. The angular resolution depends on the accuracy of determining the SDPs and on the opening angle between them. The SDP accuracy is better for longer tracks and smaller pixels.

Data from a single FD eye together with the trigger time(s) from one or more SD stations yield a hybrid geometric reconstruction that offers better accuracy than stereo reconstruction. After determining the SDP, 'geometric reconstruction' means identifying the shower axis within the SDP together with the time when the shower front passes some point on that axis. To understand the hybrid method, imagine that you know precisely the angular velocity of the track as the spot passes the center of some particular pixel. If somebody were to tell you the distance to the shower axis at that point of the track, you could calculate when the light was emitted from that point of the axis and (using the measured angular velocity) what angle the axis makes with the pixel's viewing direction. The geometric reconstruction is therefore complete if you are told that one distance to the axis. Since you may not be able to determine that distance accurately from the FD data alone, you can try all possible distance hypotheses. Each one gives a unique geometric reconstruction and consequently a unique prediction for the trigger time of any SD station on the ground. The trigger time of any ground station can then identify which distance hypothesis is correct and thereby select the true geometry. This timing method typically picks out the axis within the SDP with less uncertainty than if the axis is determined by a second (stereoscopic) SDP. For stereo hybrid events, there are timing determinations in two independent planes, providing reconstruction accuracy that is superior to both stereo reconstruction and monocular hybrid reconstruction.

See [9] for specific methods of reconstructing air shower arrival directions as well as techniques for extracting astrophysics results from a set of arrival directions.

5. Measuring cosmic ray energies

Conceptually, energy determination by a fluorescence detector is straightforward. The amount of emitted fluorescence light is proportional to the energy losses of the charged particles. Measuring the fluorescence emission from the full shower

development should yield the total electromagnetic shower energy. It is a robust calorimetric measurement. The only dependence on hadronic model or composition is in the small fraction of primary energy that is assumed to escape the hadronic cascade as muons and neutrinos rather than being transferred to the electromagnetic cascade. Simulations suggest that this fraction is approximately 5% for proton showers and 10% for iron showers. By assuming 7.5%, the error due to ignorance of the primary particle should not be more than about 2.5%. Still, it is important to recognize that air shower measurements are an indirect method for determining the energy of a cosmic ray, and there is *some* systematic model uncertainty that is difficult to quantify. The fraction of the cosmic ray energy not dissipated electromagnetically also fluctuates shower-to-shower, especially for protons.

Implementing this conceptually simple calorimetric method encounters numerous difficulties:

- The optical clarity of the atmosphere is variable because of changes in the aerosol density and aerosol composition. This makes it problematic to infer the amount of emitted light based on the observed flux. Detailed atmospheric monitoring can, in principle, overcome this difficulty.
- The absolute fluorescence efficiency is not accurately known. This causes uncertainty when getting the amount of energy deposited from the amount of fluorescence light produced.
- The full longitudinal development is never observed. The FD records only the portion of the shower development that is above ground level and large enough to produce a detectable light flux at the detector. Some extrapolation using a fitted functional form is needed to account for the parts of the development that are not measured.
- Direct Cherenkov light may contaminate the fluorescence signal. Although the Cherenkov emission angle is only about 1 degree in air, the shower electron directions can have quite large angles with respect to the shower axis. The Cherenkov part of the received signal is non-negligible at least out to 20 degrees from the axis direction. The direct Cherenkov light needs to be modeled accurately.
- Light scattered to the FD from the intense forward Cherenkov beam contaminates the fluorescence signal to some extent at all viewing angles. Scattered Cherenkov light has a different spectrum from direct Cherenkov light, and both differ from the direct fluorescence light. Detailed atmospheric monitoring is needed to model the contamination correctly.
- Scattering of fluorescence light and multiple scattering of Cherenkov light cause a kind of halo around the instantaneous point of origin on the shower axis. The amount of collected light depends on the solid angle of acceptance, exceeding what is expected from the point source by an amount that must be modeled.

Numerous laboratory experiments are determining the absolute fluorescence emission spectrum and its dependence on electron energy and air temperature and pressure [10]. Those measurements should reduce this systematic uncertainty to 5–10%. The other itemized difficulties should result in energy errors that are more random than systematic.

Dawson [11] has implemented an iterative energy determination method that is meant to minimize reliance on any assumed functional form for the longitudinal development. He first assumes all of the signal is from fluorescence light and infers an energy deposition from the flux measured in each time bin by itself. A fitted Gaisser–Hillas curve is then used to estimate the Cherenkov beam and the consequent Cherenkov contamination of the fluorescence light in every time bin. The estimated Cherenkov signal is then subtracted from the data and the procedure is repeated, iterating until the Cherenkov-subtracted data do not change between successive iterations. The resulting longitudinal profile of fluorescence light leads to a profile of energy deposition with negligible bias caused by the functional form that is used to estimate the Cherenkov contamination in the iterative procedure. The final longitudinal profile is not a smooth functional fit; it retains the fluctuations that occur in the recorded data. A fitted Gaisser–Hillas function may play another minor role, however, in providing the best estimate of the profile in its unobserved end portions for the purpose of integrating the energy deposition profile in its entirety to get the total electromagnetic shower energy.

The FD quasi-calorimetric energy measurement can provide an important calibration for SD energy measurements which otherwise rely on shower simulations. Shower simulations are necessarily uncertain in their treatment of hadronic interactions at energies that have not been studied in collider experiments.

Simulations show that the signal collected in SD stations far from the core is approximately proportional to the total shower energy. In the case of the Auger array, the signal (s_{1000}) deposited in a water tank 1000 meters from the core is taken to be proportional to shower energy. At that distance, the longitudinal development of particle density reaches its maximum value near ground level for a large range of zenith angles. Since a smooth function changes very little near its maximum value, this method is relatively insensitive to fluctuations in shower development.

Fig. 1 shows the longitudinal development profiles for the NKG electromagnetic particle density at 7 different distances from the shower axis. The curves are analytic: the total shower size as a function of depth is taken to be a Gaisser–Hillas development curve (Section 2.2) with $N_{\max} = 6 \times 10^9$, $X_{\max} = 800$ g/cm^2, $X_0 = 0$ and $\lambda = 70$ g/cm^2; and the lateral distribution at each depth (hence shower age) is given by the NKG function (Section 2.3). The different curves correspond to different fixed distances from the axis measured in Moliere units, using the single conversion to meters that pertains just

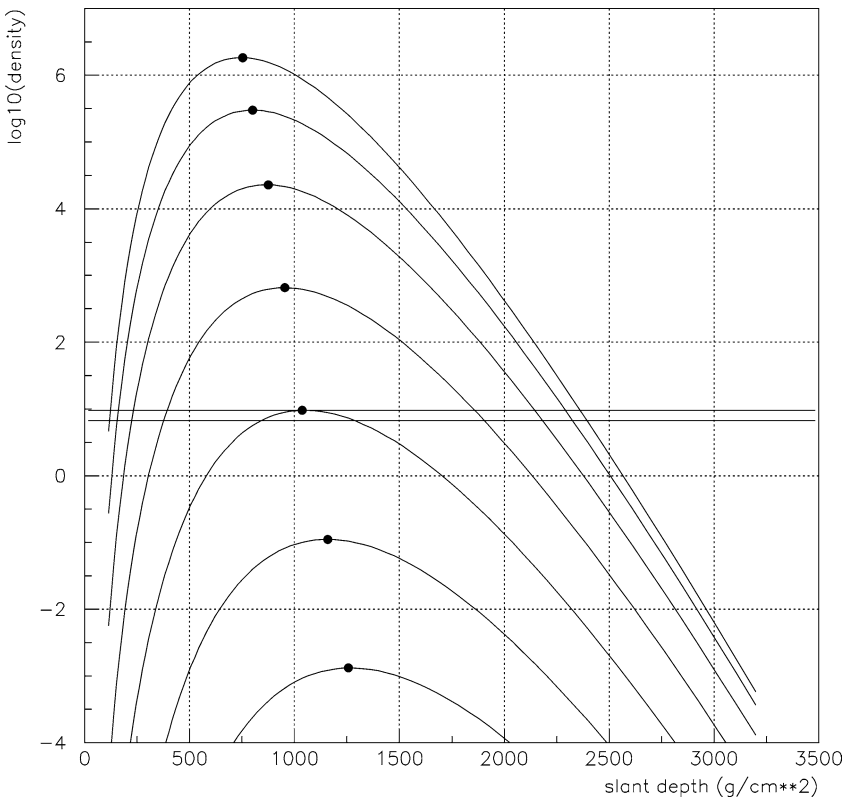

Fig. 1. Longitudinal development curves for the NKG electromagnetic particle density at 7 fixed distances from the shower axis. The distances are fixed in Moliere units but converted to meters using the Auger detector altitude. Top to bottom, the distances (in surface meters) to the axis are 10, 32, 100, 320, 1000, 3200, 10000. Note that the depth of maximum increases with distance from the axis. Two horizontal lines mark the maximum of s_{1000} curve and the value 30% lower. The density remains within 30% of its maximum from 850 to 1300 g/cm^2. The units are approximate number of vertical equivalent muons per 10 m^2 for this shower of $N_{\max} = 6 \times 10^9$.

above the Auger ground level. Each density longitudinal profile has a dot indicating its maximum. You can see that the depth of maximum increases with Moliere distance from the core. Although the total shower size reaches maximum at 800 g/cm^2, the density s_{1000} (measured at the Moliere radius that corresponds to one kilometer at ground level) reaches maximum at more than 1000 g/cm^2. A detector at vertical depth 850 g/cm^2 (Auger) or 920 g/cm^2 (AGASA) is therefore well positioned to be near the s_{1000} maximum for zenith angles at least up to 45 degrees. The horizontal solid lines in the figure mark the maximum of the s_{1000} curve and the value that is 30% less. You can see that the density stays within 30% of its maximum from 850 g/cm^2 to 1300 g/cm^2. Fluctuations in shower development (and even the systematic difference between protons and iron) shift the maximum by only about 100 g/cm^2 or less.

Knowing the slant depth of a shower's measurement, the expected longitudinal development profile of s_{1000} allows one to correct the measured s_{1000} to what it likely was at maximum. For large zenith angles, these correction factors become dangerously large for scintillator arrays. For water Cherenkov detectors, the longitudinal profiles fall much slower with slant depth than the (electromagnetic) NKG behavior plotted in Fig. 1. For water Cherenkov arrays, therefore, relatively small corrections to s_{1000} are needed even at large zenith angles to give the ground parameter (s_{1000} at its maximum), which is proportional to energy.

None of the triggered ground stations is likely to be exactly 1000 meters from the axis. Interpolation using numerous stations closer and farther from the axis gives the estimate for the signal density at 1000 m. Rather than using linear interpolation, one fits an empirical average lateral distribution functional form and takes as s_{1000} the value of that fitted function at 1000 meters from the axis.

A statistical correlation of s_{1000} with shower energy measured calorimetrically by the FD yields the conversion factor from s_{1000} to energy without relying on shower simulations that use untested hadronic interaction models.

See [12] for details about cosmic ray energy determinations from air shower measurements.

6. Probing the primary mass distribution

The indirect measurement of a cosmic ray by its air shower makes it impossible to measure the mass of any individual particle. Shower-to-shower fluctuations in longitudinal development can make an air shower produced by a particle of one mass indistinguishable from an air shower initiated by a particle of a different mass. As described in Section 2.4, however, each mass value leads to its own expected values for the shower depth of maximum and number of muons. While it is not feasible to determine the masses of individual cosmic rays, measurements of $X_{\max}$ and/or N_μ provide important statistical information about the primary mass distribution. Although both of them relate to the speed of shower development, fluctuations in $X_{\max}$ are not highly correlated with fluctuations in N_μ. For example, a fluctuation in depth of first interaction adds directly to a shower's $X_{\max}$, but it has little impact on the shower's N_μ.

A hybrid detector has special composition sensitivity. The FD determines the electromagnetic shower energy and the depth of maximum $X_{\max}$. Measurements by a water Cherenkov SD are then especially valuable in determining the primary mass. As shown in the article [1] in this volume, the main SD parameter s_{1000} and the FD $X_{\max}$ together powerfully constrain the primary mass when the energy and zenith angles are known. The lateral steepness, shower front curvature, and signal pulse shapes offer additional constraints.

Combining also a scintillator array with a water Cherenkov array and fluorescence detectors should allow an even more powerful probe of the primary mass distribution. The scintillator array measures the surface electromagnetic particle density, allowing that to be subtracted from the water Cherenkov signal to isolate the contribution due to muons.

Determining the primary mass distribution is challenging due to the indirect method of measuring cosmic rays by their air showers. There are many measurable parameters that correlate with the primary mass, however, and multi-parameter analyses may provide the best sensitivity. Neural networks are a special kind of multiparameter analysis. A neural net can provide a mass likelihood distribution for each measured shower, based on its multiparameter training with simulated showers.

See [1] for specific methods to estimate the primary mass distribution. That article also describes how to identify a photon component of the composition and any neutrino primaries.

7. Discussion

The methods presented in this overview are of the bottom-up type in which algorithms determine the cosmic ray geometry, energy, and mass likelihood from parameters fitted to data from an air shower measurement. The alternative 'top–down' approach is to use simulations to match the observed data. The primary energy, mass, and geometry are varied in order to find simulated detector data that best match the actual raw data. This method is computationally intensive, model-dependent, and perhaps vulnerable to fluctuations in real data and in the simulations. It has the advantage, however, of matching details of the detector data rather than basing conclusions on parameters that are supposed to summarize the observed data. The top–down approach is an important complement to conventional bottom-up determinations.

Any detector of high energy cosmic rays makes its measurements of each primary particle indirectly by recording properties of the particle's air shower. There are important limitations to the reliability of each type of detector by itself. Systematic errors can be identified and corrected by comparing measurements made with different types of detectors. An air fluorescence detector's measurement of energy does not depend on zenith angle, so it can check that a surface array's energy measurements do not vary systematically with zenith angle. On the other hand, a uniform surface array has translational invariance that can be used to check that an air fluorescence detector does not introduce a systematic energy error with shower distance from the eye, due for example to systematic error in atmospheric attenuation corrections. The SD constancy in time also affords a test for time variability in FD energy estimations.

In principle, measurable air shower properties can determine accurately the arrival direction of a high-energy cosmic ray. In practice, common air shower detectors achieve accuracy on the order of 1 degree. That should be adequate for charged particles that are necessarily deflected that much by galactic and/or extragalactic magnetic fields. If superb angular resolution becomes a priority, e.g., because there are tight clusters of neutral particle arrival directions, then stereoscopic hybrid measurements offer a way to achieve resolutions of 0.1 degree or better.

Energy resolution is limited partly by Poisson fluctuations in the detector samplings, partly by uncertainty about the primary particle masses, partly by uncertainty in the high-energy hadronic interaction modeling, and partly by atmospheric uncertainties. Atmospheric variability affects fluorescence detectors more than surface arrays, whereas hadronic interaction uncertainties and the unknown primary masses are more problematic for surface arrays. Energy resolutions of 5%–20% should be possible for either type, and hybrid measurements can improve the accuracy. It is impossible to be completely confident of an energy measurement, since prompt production of weakly interacting particles can conceivably steal energy undetectably from the first interactions.

Measuring the primary particle masses is especially challenging. The first interactions of an energetic proton can distribute its energy to many hadronic subshowers, accelerating the shower development in a way that is comparable to that of a 56-nucleon shower of a primary iron nucleus. On average, a proton shower has a depth of maximum that is at least 80 $\mathrm{g/cm}^2$ deeper than an average iron shower of the same energy. The distributions for iron $X_{\max}$ and proton $X_{\max}$ overlap, however. Similarly, an iron shower is expected to produce 30% more muons, but the proton and iron distributions for N_μ overlap. Jointly measuring the electromagnetic longitudinal development with an atmospheric FD and the muon density with SD arrays offers the best way to derive a likelihood distribution for the primary mass of each cosmic ray and to determine properties of the overall chemical composition. Hybrid detectors are crucial in this regard.

A surface array records cosmic rays with constant energy-independent aperture. Working together with a fluorescence detector, the observatory assigns a reliable energy and a maximum-likelihood atomic mass to each shower. The result is a histogram of shower counts over the two-dimensional (E, A) space. Sliced one way, this is the energy spectrum for different mass groups. Sliced the other way, it is the chemical composition as a function of energy. Projected onto one axis, it is the all-particle energy spectrum. Projected onto the other axis, it is the all-energy chemical composition.

References

[1] P. Billoir, P. Sommers, Identification of the primary cosmic ray, C. R. Physique (2004), in press.
[2] M.A. Longair, High Energy Astrophysics, vol. 1, Cambridge University Press, 1992, pp. 118–122;
W. Heitler, The Quantum Theory of Radiation, Dover, New York, 1984.
[3] K. Greisen, Progr. Cosmic Rays 3 (1956) 1.
[4] T. Gaisser, M. Hillas, in: Proc. 15h ICRC, vol. 8, Plovdiv, 1977, p. 353.
[5] Particle Data Group, Review of particle physics, Phys. Rev. D 66 (2002) 010001.
[6] K. Greisen, Ann. Rev. Nucl. Sci. 10 (1960) 63.
[7] K. Kamata, J. Nishimura, Prog. Theor. Phys. (Kyoto) Suppl. 6 (1958) 93.
[8] http://www.mpifr-bonn.mpg.de/staff/hfalcke/LOPES.
[9] R.W. Clay, B.R. Dawson, G.J. Thornton, Directional reconstruction and anisotropy studies, C. R. Physique (2004), in press.
[10] http://www.auger.de/events/air-light-03/.
[11] B. Dawson, M. Debes, P. Sommers, Shower profile reconstruction with engineering array FD data, Pierre Auger technical note GAP-2002-016, http://www.auger.org/admin/GAP_Notes/.
[12] S. Yoshida, Energy determination of trans-EeV cosmic rays, C. R. Physique (2004), in press.