

**Towards a New Cosmic Ray Composition Measurement in
the Knee**

**A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY**

Scott Patrick Wakely

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY**

September, 1999

Towards a New Cosmic Ray Composition Measurement in the Knee

by Scott Patrick Wakely

Under the supervision of Professor Keith Ruddick

ABSTRACT

Details of a new non-imaging Dual Cerenkov array are presented. The array is operated in coincidence with the Soudan 2 underground tracking calorimeter. It was designed to make multi-parameter measurements of extensive air showers for the purpose of extracting the composition of the primary cosmic rays in the region of the spectral knee 10^{14} eV to 10^{16} eV. The design and operation of the detector are described and results on the composition of cosmic rays in the knee are presented.

Acknowledgements

For years of support, encouragement, understanding, and guidance, I am deeply indebted to my advisor, Keith Ruddick. His deft and subtle collaboration has made this investigation the rewarding experience it has been. Keith's reputation for intelligence, wisdom, and humor is unsurpassed and well-deserved. He is so highly regarded at Minnesota that I once heard a senior faculty member scoff loudly at the absurdity of the understatement, "Keith is one of the best."

I would also like to acknowledge Pete Border for his continuous support through my graduate student career. Among those in the know, Pete is regarded as the most valuable asset in the entire High Energy Physics group. His depth of intellect and breadth of knowledge are matched only by his kindness and generosity. I am profoundly grateful for his help and companionship.

I also acknowledge Vladimir Vassiliev and Leon Mualem, upon whose work the present investigation is built. Both have shown remarkable patience in answering my questions and providing assistance. I would also like to thank Sue Kasahara, who first introduced me to STING some five years ago; her help through the years in understanding the mysteries of Soudan 2 is much appreciated.

I am also grateful for the help of many of the Physics Department personnel; in particular, machinists Brent Anderson, and Carl Johnson, who helped with the construction of the detector components. Thanks to Al Knutson, who provided advice and precious electronics components. Thanks also to Joyce Clarkin, Joanna Cook, and Bobbi Eich for being friends in the administration. Graham Allan warrants a hearty thanks for his friendship and assistance in all matters computer.

The Soudan 2 mine crew helped out in too many capacities to mention. Jim Beatty, in

particular, was a constant source of help and advice. Bill Miller, Jerry Meier, Brian Anderson, Grim Benson, and Whitey Carlson all “did time” working on surface detector efforts; their help has been appreciated.

I thank my fellow graduate students for amusing and supporting me for the last six years: Jake Schmidt, Peter West, Mark Williams, Nina Markovic, Steve Grantz, George Skidmore, Usman Suriono, and Phil Kraus. Without their companionship, this work could never have been completed.

As the husband of a teacher, I think it is appropriate to recognize the impact that a good teacher can have on a life. To that end, I would like to also thank Marcia Trelstad, Craig Stoneberg, Steve Ethen, Gary Prest, and Serge Rudaz for teaching me that the simple act of learning is among life’s most sublime pleasures.

For their support and love, I am grateful to my family and friends, including my in-laws, Jack and Carole Jarpe; my sister, Tricia Wakely; my grandmother and aunt, Barb and LaVonne Billiet; and my friends Yonder Colby and Adam Minter. I am especially thankful to my parents Bonnie and Jim Wakely, to whom this thesis is dedicated. Any success I have achieved, I credit to their love, patience, humor, and understanding.

Finally, I would like to thank my wife Anne and my daughter Rosalie for providing me a reason to leave the lab at night....

Dedication

For Mom and Dad.

Table of Contents

Abstract	i
Acknowledgements	ii
Dedication	iv
List of Tables	ix
List of Figures	x
1 Introduction	1
1.1 Cosmic Rays	1
1.1.1 Spectrum	2
1.1.2 Composition	5
1.1.3 Acceleration	8
1.1.4 Propagation	10
2 Measurement Techniques	12
2.1 Direct Methods	13
2.1.1 Satellite Experiments	13
2.1.2 Balloon Experiments	14
2.2 Indirect Methods	14
2.3 Extensive Air Showers	15
2.3.1 Hadronic Component	15

2.3.2	Electromagnetic Component	17
2.3.3	Muonic Component	19
2.4	Cerenkov Radiation	20
2.4.1	Emission from Single Particles	20
2.4.2	Light Pools	21
2.4.3	Emission from EAS	24
2.4.4	Two Component Model	26
2.5	Dual Cerenkov Technique	31
2.6	Previous Composition Measurements	32
2.6.1	Synopses	32
2.6.2	Comments	36
3	The Soudan 2 Detector	37
3.1	Module Design	38
3.2	Veto Shield	41
4	Dual Cerenkov Array at Soudan	45
4.1	Lower Station Detector	45
4.1.1	Detector Enclosure	48
4.1.2	High Voltage System	50
4.1.3	Safety Interlock Module	52
4.1.4	Microprocessor Control Module	53
4.1.5	Signal Splitter/Current Monitor	54
4.2	Upper Station Detector	55
4.2.1	Detector Enclosure	55
4.2.2	High Voltage System	57
4.3	Photomultiplier Tubes and Bases	57
4.3.1	Design	57
4.3.2	Calibration and Testing	59
4.4	Common Hardware Features	63

4.4.1	Windows and Winston Cones	63
4.4.2	Aperture-Control Logic Unit	64
4.4.3	RS-232/422 Converters	65
4.5	Data Acquisition Hardware	66
4.5.1	Signal Handling	66
4.5.2	Trigger Electronics	68
4.5.3	Skywatcher	71
4.5.4	Other Electronics	71
4.6	Software	72
4.6.1	Primary Control Program	72
4.6.2	Calibration Software	76
4.6.3	Microprocessor Control Program	77
4.6.4	Inter-Integrated Circuit Protocol (I ² C)	78
4.7	Operation	78
5	Data Handling	80
5.1	Data Consolidation	80
5.1.1	Run and Event Matching	81
5.1.2	Event Extraction	81
5.1.3	Event Consolidation	83
5.2	Signal Processing	84
5.2.1	Noise Characteristics	85
5.2.2	Noise Filtering	88
5.2.3	Pulse Matching	91
5.3	Event Characterization	95
6	Analysis	96
6.1	EAS Monte Carlo	96
6.1.1	Hadronic Simulation	97
6.1.2	Muonic Simulation	98

6.1.3	Electromagnetic Simulation	100
6.2	Data Interpretation	101
6.2.1	Energy Determination	102
6.2.2	X_{\max} Determination	103
6.2.3	Iteration of Results	105
6.3	Muon Information	105
7	Results and Conclusions	109
7.1	Results	109
7.2	Comments	111
7.2.1	Limitations	111
7.2.2	Enhancements	112
7.3	Conclusions	114
	Appendix A. Inter-Integrated Circuit Protocol	116
A.1	Hardware	116
A.2	Software	119
	References	125

List of Tables

4.1	Dual Cerenkov Array Specifications.	47
5.1	Data Reduction Summary	95
6.1	Average X_{\max} from Several Interaction Models	108
6.2	Average X_0 from Several Interaction Models	108

List of Figures

1.1	Low Energy Region of the Cosmic Ray Spectrum.	3
1.2	High Energy Region of the Cosmic Ray Spectrum.	4
1.3	Solar and Cosmic Ray Elemental Abundances.	6
2.1	Schematic Depiction of an Extensive Air Shower.	16
2.2	Toy Model of an Electromagnetic Cascade.	18
2.3	Cerenkov Ring.	22
2.4	Cerenkov Emission Graph.	25
2.5	Cerenkov Source Model.	27
2.6	Cerenkov Radiation From a 60° 10 PeV H EAS.	28
2.7	Cerenkov Radiation From a 0° 10 PeV H EAS.	29
2.8	Pulse Height vs. Radius for 4 Vertical Hydrogen Showers.	30
2.9	Normalized Pulse Height vs. Radius for 4 Vertical Hydrogen Showers.	31
2.10	Cerenkov Field Zones.	32
2.11	Mean Primary Mass as Measured by BLANCA.	35
2.12	Summary of Recent Cosmic Ray Composition Measurements.	36
3.1	The Soudan 2 Detector and Veto Shield.	38
3.2	Soudan 2 Module.	39
3.3	Drift Tube Cross Section.	40
3.4	Module Wireplane.	41
3.5	Soudan 2 Cosmic Ray Moon Shadow.	42
3.6	Online Software (STING) Readout	44

4.1	Detector Configuration Map.	46
4.2	Lower Station Detector Enclosure.	49
4.3	Lower Station Detector Enclosure Schematic.	49
4.4	Lower Station High Voltage Module Schematic.	51
4.5	Safety Interlock Module Schematic	52
4.6	Signal Splitter/Current Monitor Schematic.	55
4.7	Upper Station Enclosure Schematic.	56
4.8	PMT Base Schematic.	58
4.9	Sample Calibration Histogram.	60
4.10	PMT Gain Curve and Power Usage.	61
4.11	PMT Linearity Curve.	62
4.12	PMT Assembly Schematic.	63
4.13	Angular Acceptance of Tube with Winston Cone.	64
4.14	Aperture-Control Logic Unit Schematic.	65
4.15	ACLU Switch Logic.	66
4.16	Detector Signal Path.	67
4.17	Trigger Logic.	69
4.18	Instrument Control Interface.	73
4.19	Threshold-Setting Routine.	74
5.1	Event Matching Procedure.	82
5.2	Sample Cerenkov Signal.	84
5.3	FFT of Sample Cerenkov Signal.	85
5.4	FFT of Sample Cerenkov Signal, Log Scale.	86
5.5	A Simulated Cerenkov Pulse.	87
5.6	A Simulated Cerenkov Pulse Power Spectrum.	87
5.7	Simulated Power Spectrum of a Cerenkov Pulse with Noise.	88
5.8	A Poorly Filtered Cerenkov Pulse.	89
5.9	Filtered Cerenkov Signal.	90
5.10	FFT of Filtered Cerenkov Signal.	91

5.11	Pulse Matching Efficiency versus Window Size and Pulse Depth.	92
5.12	Pulse Matching Efficiency versus Pulse Height.	93
5.13	Mean Reconstruction Error versus Window Size and Pulse Height.	94
5.14	Reconstructed Pulse Error versus Pulse Height.	94
6.1	Monte Carlo Programs.	97
6.2	Muon Transmission Probability.	99
6.3	Average Muon Yield.	99
6.4	Energy vs. Pulse Integral for Three Primary Masses.	102
6.5	X_{\max} vs. Pulse Ratio.	104
6.6	Monte Carlo Pseudorapidity Distributions	107
6.7	Monte Carlo Muon Lateral Distributions	107
7.1	Average X_{\max} vs. Particle Energy.	110
A.1	Typical I ² C System Schematic.	117
A.2	I ² C Command Primitives.	118
A.3	I ² C Read Protocol.	119
A.4	I ² C Write Protocol.	119

Chapter 1

Introduction

In this work, the results of an investigation into the composition of primary cosmic rays in the energy region commonly known as the “knee” are reported. This region represents an important crossroads for several different physical processes and observed phenomena. It is at the knee where the spectral index of the particle flux changes. It is at the knee where common acceleration mechanisms reach the limit of their accelerating ability. It is near the knee that modern high energy physics data is exhausted, as particle accelerators reach their maximum center of mass energies. Finally, the knee represents the energy where the rapidly falling flux makes precise, direct measurements of cosmic ray properties untenable.

This chapter contains a brief overview of cosmic rays and experimental techniques for measuring their properties. Chapter 3 contains a description of the Soudan 2 tracking calorimeter, while in Chapter 4 details of the Dual Cerenkov hardware and software may be found. A review of the collected data is in Chapter 5, and the techniques used to analyze it are in Chapter 6. Chapter 7 details the conclusions which can be drawn from the data are detailed.

1.1 Cosmic Rays

Cosmic rays consist of extra-terrestrial nuclei, elementary particles, and gamma rays. Comprising roughly 90% protons, 9% alpha particles, with sparse amounts of heavier nuclei, gamma rays, electrons, and other subatomic particles, cosmic rays impinge on the earth's atmosphere

at a rate of several thousand per square meter per second [1].

Though they have been rigorously investigated since their first observation by Victor Hess during a 1911 balloon flight [2], cosmic rays have defied a comprehensive description. Multiple independent measurements of observable characteristics such as composition and maximum energy remain inconsistent and a source of controversy. Characteristics inferred from those observables, including cosmic ray origins, acceleration mechanisms, and propagation conditions are even less certain. Numerous experimental efforts continue to address these issues.

Cosmic ray energies span a range of over 14 orders of magnitude, with the greatest energies reaching over 2×10^{20} eV [3]. To illustrate the importance of cosmic ray processes with respect to other pervasive galactic phenomena, the following energy comparisons are sometimes invoked: The galactic energy density of cosmic rays is roughly 0.6 eV/cm^3 . This is of the same order as the energy of interstellar magnetic fields (0.25 eV/cm^3), starlight (0.44 eV/cm^3), microwave background radiation (0.25 eV/cm^3) and the turbulent kinetic energy of gases in the interstellar medium (0.31 eV/cm^3) [4]. Clearly, cosmic rays play an important role in the energy balance of the galaxy and an explanation of their origins should be included in any comprehensive description of galactic dynamics.

1.1.1 Spectrum

Cosmic rays are observed at energies from a few times 10^6 eV to over 10^{20} eV. At the low range, this represents energies of the same order as typical nuclear processes. At the high end, this extends well into the macroscopic regime; a particle of 2×10^{20} eV carries the same kinetic energy as a tennis ball traveling 75 mph, roughly the speed of a good serve from a talented amateur.

Figure 1.1 shows the differential flux of particles for four species of primary cosmic ray. The observed flux in the region to the left of the spectral turnover, below roughly ~ 10 GeV/nucleon, is heavily influenced by the activity of the sun. This influence, called “solar modulation”, results from the interaction of the extra-solar cosmic rays with the solar wind. This wind, an energetic plasma with embedded magnetic fields, acts to sweep lower energy particles from the heliosphere [1]. During the 11-year cyclic periods of high solar activity, the increased strength of the solar wind provides a significant barrier to the incident cosmic rays. The effect of the barrier is to




Figure 1.1: Low Energy Region of the Cosmic Ray Spectrum. The differential spectra are from direct measurements at solar minimum. Shown are measurements for hydrogen, helium, carbon, and iron. Adapted from [5] by [6].




Figure 1.2: High Energy Region of the Cosmic Ray Spectrum. Adapted from the Particle Data Group [6]. The data references are \blacktriangle [7], \blacktriangledown [8], \triangle [9], \square [10], \circ [11], \blacksquare [12], \bullet [13], \blacklozenge [14].

suppress, by as much as an order of magnitude, the cosmic ray flux at these low energies.

Above energies of some few hundred GeV per particle, the differential cosmic ray particle flux can be described as a simple power law of the form

$$\frac{dN}{dE} = AE^{-\gamma} \quad (1.1)$$

The spectral index, γ , remains fairly constant over several orders of magnitude at ≈ 2.6 – 2.7 , steepening to ≈ 3.0 – 3.1 above a few times 10^{15} eV. Beyond roughly 10^{19} eV, additional spectral changes occur. Figure 1.2 shows the all-particle differential flux for the nine orders of magnitude above 10^{11} eV. The ordinate, in the standard practice of the field, has been multiplied by a weighting factor of $E^{2.7}$ to accentuate the changes in spectral slope. It is from this view that the physiological nomenclature of the field emerges. Viewing the leftmost stretch of the graph as a “thigh”, the spectral break at 10^{15} eV becomes the “knee”, while the break at 10^{19} eV becomes the “ankle”.

The weighted ordinate acts to disguise what may be the most impressive feature of the cosmic rays. This is the remarkable change in flux over the full, 14 order of magnitude, range of cosmic ray energies. Consider a 1 m^2 , 1 steradian particle detector at the top of the atmosphere. It

can be expected to detect cosmic rays of energy $\geq 10^{11}$ eV at roughly $\sim 10^1$ Hz. However, for cosmic rays of energy $\geq 10^{20}$ eV, this rate plummets precipitously, to $\sim 10^{-16}$ Hz, or one every three hundred million years.

1.1.2 Composition

At low energies ($E < 10^{14}$ eV), and therefore in general (because of the rapidly falling particle flux), cosmic ray elemental abundances roughly match those of the solar system. This includes a relative surplus of even-Z nuclei with respect to odd-Z nuclei, reflective of the greater pairing force in even-Z nuclei.

There are, however, a few notable exceptions to the composition similarity. These include a paucity of Hydrogen and Helium in the cosmic ray flux as well as two regions of surplus, shown in Figure 1.3. Elements just below Carbon, Nitrogen and Oxygen on the periodic table, the so-called “sub-CNO” elements of Lithium, Beryllium, and Boron (Li, Be, B) are overabundant in the cosmic rays by some five orders of magnitude. The elements just below iron on the periodic table, the “sub-Fe” elements, Scandium, Titanium, Vanadium, Chromium, and Manganese (Sc, Ti, V, Cr, Mn) are also overabundant, by between one and three orders of magnitude [1].

There is little agreement over the extent to which the H and He discrepancies should be attributed to an actual scarcity of the elements at the source regions *versus* a form of preferential acceleration/injection mechanism. It should be noted, however, that Hydrogen and Helium have very high first ionization potentials, 13.6 eV and 24.6 eV, respectively. This is of interest because elements with a high ($E > 9.5$ eV) first ionization potential (*e.g.*, the noble gases) are in general underabundant in cosmic rays [16]. This would seem to favor a preferential injection explanation for the H/He absence over differences in source composition.

The sub-CNO and sub-Fe elemental overabundances, however, likely do not reflect the composition of source regions or even the idiosyncrasies of their acceleration and injection mechanisms. The surplus of these elements, none of which are standard products of normal stellar nucleosynthesis, is typically attributed to the spallation, or fragmentation, of the heavier “parent” nuclei (CNO and Fe, respectively) during their propagation from the source regions. Study of the ratios of these parent elements to their spallation by-products offers insight into the properties of the interstellar medium through which the cosmic rays propagate.

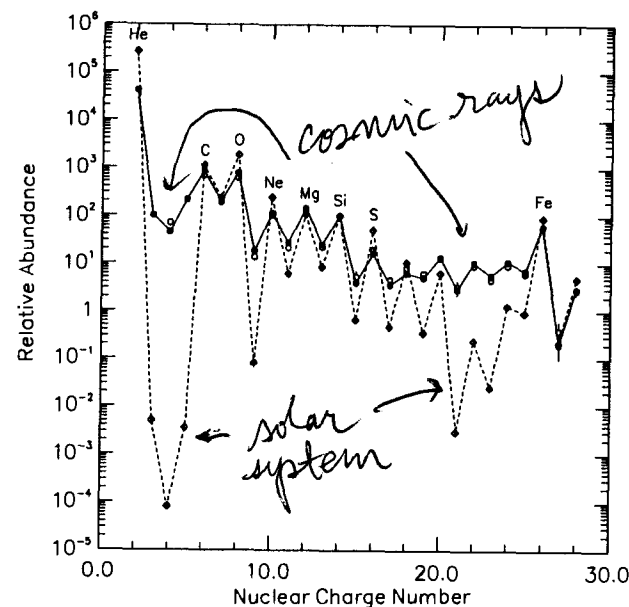


Figure 1.3: Solar and Cosmic Ray Elemental Abundances. The elemental abundances from the solar system (○) are compared to cosmic ray abundances between 70-280 MeV/nucleon (●) and 1000-2000 MeV/nucleon (◐). The normalization is to Silicon at 100. Adapted from [5] by [15].

The break in the energy spectrum above 10^{15} eV, the “knee”, suggests a transition to a new set of source, acceleration, or propagation mechanisms. It is reasonable to expect, therefore, a shift in elemental abundances which reflects the characteristics of the new origins. Unfortunately, near the knee (around 2×10^{15} eV), the data from direct composition techniques are exhausted. The knee represents a *practical threshold* to experimental cosmic rays physics. Because of the rapidly-falling particle flux, direct observation of cosmic rays beyond the knee in satellite or balloon-borne experiments becomes prohibitively expensive in terms of dollars or exposure time.

Beyond these energies, indirect measurements of composition are required. These are measurements which are necessarily more difficult and less certain than the direct variety. Reflecting this, there does not yet appear to be any general agreement on a model of composition beyond the knee. Evidence exists to support a range of models from heavy (increasing mass with increasing energy) [17, 18, 19] to light (decreasing mass with increasing energy) [20, 15, 21]. Mixed models are offered as well [22]. There is currently much activity in the area of high energy composition studies, including that by the surface detector group of the Soudan 2 collaboration.

Somewhat past the knee, at energies of around 10^{17} eV, there appears to be some recent consensus agreement that the composition becomes largely Iron-dominated [23]. The evidence for this conclusion comes from extensive air shower arrays and atmospheric fluorescence techniques. These same methods further suggest a move towards light composition above 10^{19} eV. At those energies, the region of the ankle, there are other reasons to believe the primary cosmic ray flux is all protons. The spectral break itself suggests a transition to new sources while the extreme rigidity of these particles implies these sources are of extragalactic origin.

Extragalactic particles are assumed to have traversed a sizeable quantity of matter in interstellar and intergalactic media *en route* to earth. Therefore, the considerably longer mean interaction length of protons (*versus* Iron) supports a light composition at the highest energies. Furthermore, the universal 2.7 K microwave background radiation field, which at the energies beyond the ankle becomes blue-shifted to relative energies capable of photo-disintegration, limits the travel lifetime of heavier nuclei by fragmenting them into smaller particles. It should be remarked that work at these energies is hampered by extreme scarcity of data.

1.1.3 Acceleration

There are numerous models of particle acceleration posited to describe the observed cosmic ray composition and spectral features. Perhaps the fundamental constraint on a theoretical acceleration mechanism is its ability to reproduce, over a wide range of energies, the power law behavior observed in the cosmic ray energy spectrum. It must also recreate the chemical and isotopic abundances which are measured in the cosmic ray flux. Finally, there should be a means of accelerating the particles to energies in excess of 10^{20} eV. No single theory or mechanism has been found which can recreate all these characteristics over the full range of the cosmic ray energy spectrum. (Nor is there reason to believe that such a mechanism should exist.) Thus, with the breaks in the spectrum as a guide, acceleration mechanisms are generally fit to the three regions of below the knee, above the knee, and beyond the ankle.

Below Knee

A method due to Fermi [24] was the first to reproduce a power law energy spectrum. This method, a stochastic acceleration model, involved the collision of charged particles with magnetic clouds of the interstellar medium. Though the required power law spectrum was obtained, the mechanism was too inefficient to create particles of appropriate energies. The acceleration mechanism was second-order dependent on the (low) cloud velocities and slowed by the sizeable mean interaction times for particle-cloud interactions [25].

A reformulation of Fermi’s original theory leads to what is certainly the most popular current acceleration model. This is “diffusive shock acceleration,” which occurs in the wake of large astrophysical shock waves. It is first-order dependent on the (relativistic) shockwave velocity.

Relativistic supernova shocks expanding into the interstellar medium can provide, via this “first-order Fermi acceleration,” a means by which particles could be accelerated to energies up to $\sim Z \times 10^{14}$ eV (under favorable conditions). Even simplistic models of this supernova remnant (SNR) acceleration can produce a cosmic ray spectral index close to the observed index. The appeal of the model is further enhanced by the approximate parity between galactic supernova power output and total cosmic ray luminosity. With an average energy release of 10^{51} ergs and a rate of approximately one per decade, supernovae feed the galaxy with energy at a rate of $\sim 10^{42}$ erg sec⁻¹, while the total cosmic ray luminosity is $\sim 10^{41}$ erg sec⁻¹ [26].

Thus, a (arguably) modest 10% of the total supernovae power output is needed to accelerate the bulk of the cosmic ray energy spectrum.

An important by-product of the SNR model, however: the production of high-energy gamma rays, has yet to be observed. These photons, which should be created by the decay of neutral pions produced in the acceleration regions, are a compelling piece of evidence which has yet to emerge. Enhanced searches by next-generation gamma ray observatories may help settle this long-standing acceleration issue [26].

Above Knee

Supernovae remnant shock mechanisms are not capable of particle acceleration much beyond the knee region. The spectral break, therefore, likely marks a transition between acceleration sources [27]. The debate over the character of the new sources is ongoing and is hindered by a lack of accurate composition data for the region around and beyond the knee. Above-knee acceleration theories must be capable of providing for acceleration up to energies of 10^{19} eV. (The power requirements, however, are modest: only 0.01% of the upper-knee range.) These theories fall into one of two categories: variations on shock acceleration models and new source models.

The richness of the shockwave acceleration model insures its invocation in numerous forms to describe the cosmic ray flux above the knee. In most instances, secondary acceleration of energetic sub-knee cosmic rays is posited. In these “reacceleration” models, some portion of the pre-accelerated cosmic ray population is increased in energy through interactions with a series of new shockwaves.

Reacceleration models may involve unique configurations of traditional shock mechanisms, or ensembles of large shocks acting in concert. Examples include supernovae shocks expanding into the effluvium of particularly active stars, or multiple clustered supernovae creating localized environments where particles rapidly bounce between multiple fronts [28]. Single galactic objects, however, typically lack the field strength, lifetime, and shock size to provide significant acceleration. Thus, the termination shock of the galactic wind itself is often cited as a possible reacceleration site. Because of its size and extremely long lifetime, it could accelerate particles to energies of some 10^{20} eV [25]. Acceleration in the jets of active galactic nuclei (AGN) has also been explored.

Finally there is the possibility of entirely new cosmic ray sources. Electromotive acceleration near pulsars [29], accretion dynamics near massive bodies in binary systems [25] and effects from coalescing black holes or neutron stars [30] are among the sources which have been investigated.

Ankle

At energies above the ankle, there is little data against which to test the theories of acceleration. However, the rigidity of particles at this energy suggests an extragalactic origin. Furthermore, the limited range of energetic nuclei in the intergalactic medium (the Greisen-Zatsepin-Kuzmin effect, see Section 1.1.4) restricts the maximum distance of the source regions.

Morfill and Jokipii have offered a model wherein heavy nuclei are accelerated at the large scale intergalactic shock fronts [28]. High energy protons, which are thought to dominate the flux at the highest energies, are then created from these nuclei by photo-disintegration processes. This model matches the heavy composition measurements below the ankle as well as the largely light measurements above it. Shock front configurations such as might be found at the boundaries of colliding galaxies could also achieve the required energies.

Mechanisms consistent with the Standard Model may, in fact, be less useful for the description of the highest energy cosmic rays. To address the existence of these particles, increasingly exotic source theories are offered. The collapse or annihilation of cosmological topology defects such as cosmic strings or monopoles are typical of these theories [30]. Clearly, this region of the spectrum lies at the extreme boundaries of what can be explored with current instruments. A next-generation air shower array, The Pierre Auger Observatory, is currently under construction in Argentina to specifically target this energy regime.

1.1.4 Propagation

Propagation models are the final key in reconstructing cosmic ray origins. Acceleration model results must be “propagated” before they can be used. This involves simulating the transit of model output nuclei through the interstellar medium to produce spallation, magnetic deflection and containment effects, which can alter isotopic abundances, isotropy and spectral indices.

In the “Leaky Box” model, cosmic rays diffuse freely inside a fixed confinement volume, reflecting at the boundaries. With every encounter with a boundary, a particle has a finite

probability of escape. This leads to an exponential distribution of particle pathlengths in the confinement volume, which acts to alter the observed spectra and composition of cosmic rays relative to these characteristics at the source [31].

It has been observed that reducing the number of short pathlength trajectories in a model improves the model's agreement with experimental data, especially at low energies [32]. This has led to adaptations of the leaky box model to include either abrupt or gradual cutoffs of the shorter pathlengths. These models, called the Zero-Short and Double-Exponential (Nested Leaky Box) models, respectively, address, in particular, the observed decrease in secondary-primary ratios with increasing energy [33].

Other factors such as radioactive decay *en route* and ionization energy loss must be taken into account by propagation models. At the highest energies, interactions with the microwave background radiation must be included, as they may act to limit the lifetime of high energy nuclei via photo-pion production and nuclear photo-disintegration (Greisen-Zatsepin-Kuzmin or GZK effect) [30].

Chapter 2

Measurement Techniques

The goal of understanding the origin and acceleration processes of cosmic rays rests on the ability to accurately measure the energies and masses of the incoming particles. A determination of the composition of the cosmic rays throughout the spectrum would provide much needed information to constrain and inform the multitude of theories of their origin.

Unfortunately, the properties of the flux conspire with the thickness of the atmosphere to allow *simple* measurements only in a small “window of opportunity” in the particle spectrum. At low energies, the atmosphere acts as a shield, filtering out particles before they can reach ground-based detectors. With a thickness some 1000 g/cm^2 , the atmosphere represents an essentially insurmountable barrier to an incident nucleon which has a typical interaction length of $\lambda_N \leq 70 \text{ g/cm}^2$. At the highest energies, the flux has fallen to the point where appreciable data rates can only be achieved with very large acceptance or very long exposure detectors. Thus, there is only a small region in the energy spectrum where relatively inexpensive, logistically uncomplicated, and somewhat small detectors can make a contribution.

In the face of these difficulties, there have emerged two fundamental methods for the measurement of cosmic rays. The “direct” method, which is applicable to the study of lower-energy particles (below $\sim 10^{15} \text{ eV}$), involves bypassing the atmosphere by placing instruments on high-altitude balloons or satellites. The “indirect” method, which provides access to higher energy particles, makes use of the atmosphere itself as part of the detector, measuring not the primary cosmic rays themselves, but rather the products of their interaction with the atmosphere.

As previously mentioned, the knee roughly represents the crossover point between these two techniques.

2.1 Direct Methods

Obviously the most accurate way to measure the cosmic rays is to bring the detector directly to them. In the direct measurement technique, high altitude weather balloons and satellites are used to convey actual instruments to the top of or above the interfering effects of the atmosphere.

Direct measurement experiments typically contain a series of detectors to allow a simultaneous determination of the charge and energy of incoming cosmic rays on an event-by-event basis. This can be achieved with some form of a stopping calorimeter, in which ionization energy loss curves are used to match charge and energy. Calorimeters require large amounts of material to be effective, however, and are thus typically heavy instruments.

Alternatives often used include plastic scintillators, which provide charge discrimination (z), or a combination of Cerenkov and transition radiation detectors. The effects these detectors measure are sensitive to the charge (z^2) and energy (γ) of a particle, respectively. Since they are purely electromagnetic in origin, they do not require nuclear interactions to operate and large area detectors can be constructed with relatively small mass.

2.1.1 Satellite Experiments

Satellite-borne instruments represent the ultimate in direct detection experiments. They typically have (relatively) long exposure times and are completely free from the interfering effects of the atmosphere. The price of this is reduced acceptance due to the weight and size constraints of the delivery vehicles. Additionally, there are the drawbacks of excessive cost and complicated logistics in terms of deployment and data recovery.

Satellite experiments have performed measurements of the particle flux from the solar modulation region up to some 100 TeV/nucleus. A few examples of these experiments include the very large Heavy Ions in Space (HIIS) instrument, which flew aboard NASA's Long Duration Exposure Facility (LDEF) [34], the Ultra Heavy Nuclei detector on the High Energy Astronomy Observatory Satellite (HEAO-3) [35], the Advanced Composition Explorer (ACE) [36], and the Ulysses satellite. There have also been experiments flown aboard the Space Shuttle, including

the University of Chicago "Egg" [37] experiment, and experiments proposed for deployment on the Space Station (*e.g.*, ACCESS).

2.1.2 Balloon Experiments

Balloon-borne experiments have advantages over satellite experiments in cost, deployability, and size. The price of these advantages is the slight interference of the atmosphere. Though the high-altitude balloons can reach altitudes as high as 50 km, where the atmosphere is only ~ 5 g/cm² thick, the remaining gas acts to complicate the measurement. The complication is essentially that, though the particles incident on the detector can be identified unambiguously, the presence of the atmosphere introduces a small, but non-negligible uncertainty in the source spectrum. That is, it can not be determined whether a detected particle was in fact a primary particle, or rather a product of a previous interaction with the atmosphere. Corrections can be applied, but they introduce an additional model-dependent uncertainty in the final results.

Notable balloon experiments include the Japanese-American Collaborative Emulsion Experiment (JACEE) series [18, 38], the High Energy Anti-matter Telescope (HEAT) [39], the Cosmic AntiParticle Ring Imaging Cherenkov Experiment (CAPRICE) [40], and the proposed Transition Radiation Array for Cosmic Energetic Radiation (TRACER) [41] detector.

2.2 Indirect Methods

The rapidly falling cosmic ray particle flux effectively excludes the direct measurement of particles above $\sim 10^{15}$ eV. The required acceptances and exposure times are simply not yet practically achievable in balloon or satellite-borne experiments. Additionally, ground-based direct experiments are not possible because the atmosphere is too thick for primary particles to traverse.

Fortunately, these same atmospheric shielding effects can be exploited to make *indirect* measurements of the primary cosmic ray particles. This is achieved by studying not the primary particles themselves, but rather the properties of their interactions with the atmosphere. By deploying large-acceptance, long-term detector arrays at the surface of the earth, measurements of these interaction products can be made relatively simply. The principal challenge faced by these experiments is that of correctly *inferring* the properties of the primary cosmic rays from the properties of these interactions, called extensive air showers.

An extensive air shower (EAS) is the cascade of particles which accompanies the collision of a cosmic ray with an atmospheric nucleus. A widely-dispersed field of daughter particles is created which extends deep into the atmosphere. In the case of more energetic primaries, these particles can be detected at ground level. The lateral extent of the air showers increases with cosmic ray energy, allowing ground-based detectors to measure cosmic rays whose trajectories actually lie far outside the physical boundaries of the arrays. In this way the atmosphere becomes an actual boon to experimenters, vastly increasing the effective acceptance beyond that which could be practically achieved with actual instrumentation area. Cosmic rays of the most rare variety, those at the highest energies, can thus be measured by simply building arrays of larger size and exposure time.

A detailed description of extensive air showers follows.

2.3 Extensive Air Showers

At energies upwards of 10^{13} eV, interactions between atmospheric nuclei and incident cosmic rays lead to extensive air showers (EAS). The swarm of particles created in these high energy collisions can extend laterally over hundreds of meters. The structure of the swarm retains much information about the particle which initiated it. Extensive air showers consist of three primary constituents: a hadronic component, a soft electromagnetic component, and a penetrating, muonic component; see Figure 2.1. Cherenkov radiation, a byproduct of the electromagnetic cascade, will be discussed in Section 2.4. The neutrino flux, which is for the most part undetectable, is generally ignored, other than to note that it accounts for some 5% of the total shower energy [16].

2.3.1 Hadronic Component

An extensive air shower begins when a cosmic ray primary interacts strongly in the atmosphere. Iron nuclei, which have an interaction length λ_{int} of approximately 16 g/cm^2 (roughly 33 km above sea level for vertically incident particles) at 10^{16} eV, will interact higher in the atmosphere than protons, which have $\lambda_{int} \approx 60 \text{ g/cm}^2$ (roughly 23 km above sea level). The standard atmosphere itself is 1033 g/cm^2 deep.

The interaction between an atmospheric nucleus and the primary will produce a cascade




Figure 2.1: Schematic Depiction of an Extensive Air Shower.

of high-energy daughter particles, including K 's, π 's, η 's, N 's, and others. Many of these particles will also interact with the atmosphere, sustaining and building the cascade until shower maximum. The shower maximum occurs when the energy of the interaction products fall to the point that decay losses begin to overwhelm interaction losses. Below this energy, the nuclear shower core will attenuate with a characteristic length of around 200 g/cm^2 [25], and will be of little further influence on the air shower as a whole.

The hadronic shower core gives rise to the other components of an air shower. Low energy (*i.e.*, non-interacting) charged mesons may decay into muons, fueling the penetrating portion of the EAS. The electromagnetic cascade is started and continued through the decay of neutral mesons, mostly via $\pi^0 \rightarrow \gamma\gamma$.

2.3.2 Electromagnetic Component

The electromagnetic (EM) shower is the largest component of an extensive air shower, in total particle number as well as in total energy dissipated. The EM shower is really a superposition of multiple electromagnetic cascades, each proceeding from an original energetic gamma ray. In a simplified "toy" model of an EM shower with an initiating gamma of energy E_0 , the cascade proceeds approximately¹ as shown in Figure 2.2.

The initial photon pair-produces after a characteristic length λ_{pair} to create an electron and positron, each of energy $E_0/2$. After a characteristic length λ_{brem} , these leptons will bremsstrahlung to produce photons of average energy $E_0/4$. At high energies, $\lambda_{EM} \equiv \lambda_{pair} \approx \lambda_{brem}$ is essentially the same as a radiation length (37.1 g/cm^2) and so after a distance $n\lambda_{EM}$ into the cascade, there will be 2^n particles [25]. This particle multiplication will continue until the average particle energy falls below a critical energy E_{crit} . Below this energy, (roughly 75 MeV) ionization energy losses overwhelm bremsstrahlung losses, until eventually energy losses become catastrophic, rapidly attenuating particle numbers.

It is possible, therefore, to relate the depth of the shower maximum, as well as the size of

¹A rigorous description of the electromagnetic cascade must take into account other processes, such as ionization energy loss, collision energy losses, and Compton scattering [25].




Figure 2.2: Toy Model of an Electromagnetic Cascade.

shower maximum (*i.e.*, the number of particles) to the energy of the primary particle:

$$X_{max} \propto \ln(E_0/E_{crit}) \quad (2.1)$$

$$N_{max} \propto E_0/E_{crit} \quad (2.2)$$

The depth of the shower maximum is also sensitive to the mass of the initiating particle. Since the electromagnetic component of a shower is fed from the hadronic component, fluctuations in the latter are reflected in the former. Heavy primaries interact earlier in the atmosphere and thus begin feeding the EM shower earlier, resulting in an X_{max} which is relatively higher in the atmosphere. This effect is enhanced in the superposition model, wherein showers initiated by particles of mass A and energy E_0 are taken to behave like a parallel collection of showers initiated by A particles of energy E_0/A . The superposition model leads to several subshowers of lesser energy, each multiplying to the critical energy faster, decreasing the length required for the shower as a whole to reach maximum development.

2.3.3 Muonic Component

The penetrating muonic component of the extensive air shower is fed by the decay of charged mesons. Muon numbers gradually increase to a maximum value, after which they slowly decrease due to ionization energy loss.

The mass of the primary particle determines the multiplicity of the muons in an extensive air shower. Though exact numbers are model dependent, the following crude argument can illustrate that heavy primaries should create more muons at a given energy than light primaries [25].

There are three factors which give rise to higher muon multiplicities by heavier primaries. The first involves simply the fact that higher- A primaries will interact higher in the atmosphere, generating pions in a less dense environment more conducive to decay (producing muons) versus interaction. The second derives purely from the superposition model wherein the lower energy (E/A) subshowers of A particles will produce lower energy pions, which have a higher probability of decay. The third factor stems from the fact that pion multiplicities rise quite weakly with energy:

$$N_\pi(E) \propto E^\alpha \quad \text{with } 0 < \alpha < 0.5 \quad (2.3)$$

but a particle of mass A is treated in the superposition model as simply A particles of energy E_0/A , so

$$N_\pi(E, A) \propto \left[\frac{E}{A} \right]^\alpha \quad (2.4)$$

for one particle. For all A particles:

$$N_\pi(E, A) \propto A \left[\frac{E}{A} \right]^\alpha \quad (2.5)$$

Thus, at fixed energy, a typical ratio of pion multiplicities for heavy primaries to proton primaries will be

$$\frac{N_\pi(A)}{N_\pi(1)} \propto A^{1-\alpha} \quad (2.6)$$

Therefore, if pion production scaled more efficiently with energy, (*i.e.*, greater than linearly) hydrogen would be a more effective muon producer than iron.

2.4 Cerenkov Radiation

The following section explores the behavior of Cerenkov radiation from single emitting particles and from electromagnetic cascades. The phenomena discussed can be used to help identify characteristics of primary cosmic rays from the properties of extensive air showers.

2.4.1 Emission from Single Particles

A charged particle of mass m in a dielectric medium with index of refraction n will generate Cerenkov radiation when its velocity

$$v \geq \frac{c}{n} \quad \left(\beta \geq \frac{1}{n} \right). \quad (2.7)$$

In terms of a minimum, or critical, energy, this is

$$E_{crit} = \frac{mc^2}{\sqrt{1-1/n^2}}. \quad (2.8)$$

Defining the variable $\delta = n - 1$, this becomes

$$E_{crit} = \frac{mc^2}{\sqrt{2\delta}}. \quad (2.9)$$

The minimum energy corresponds to roughly 21 MeV for electrons and 4 TeV for muons at sea-level, and has a dependence on the altitude. This stems from the dependence of the index of refraction on the density of the atmosphere:

$$\delta(z) \equiv n(z) - 1 \propto \rho(z) \quad (2.10)$$

The Cerenkov radiation is emitted in a continuous series of nested cones, at an angle from the particle velocity vector according to the relation

$$\cos \theta = \frac{1}{\beta n} \quad (2.11)$$

The number of photons produced by a particle of charge ze over a path length dx and per unit energy interval of the photons is:

$$\frac{d^2 N}{dx d\lambda} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2}\right) \quad (2.12)$$

Defining $Q(\lambda)$ as the quantum efficiency of the collecting photocathode, it is useful to define the quantity X_{pe} , which is the unit photoelectron radiation length:

$$\frac{1}{X_{pe}} = (2.2 \times 10^{-2}) z^2 \int_0^\infty Q(\lambda) d\frac{1}{\lambda}. \quad (2.13)$$

This allows a recasting of the photon-production distribution rate can as a photoelectron distribution rate, for a particle of energy E :

$$\frac{dN_{pe}}{dX} = \frac{1}{X_{pe}} \left(1 - \left[\frac{E_{crit}}{E}\right]^2\right) \quad (2.14)$$

By inserting the atmospheric and quantum efficiency parameters used in this investigation, the photoelectron emission length X_{pe} is $2.23 \times 10^{-2} \text{ g/cm}^2$.

2.4.2 Light Pools

Figure 2.3 shows the Cerenkov light pool created at ground level by a single non-interacting electron traveling vertically through the atmosphere. The dots on the graph represent effective photoelectrons which would be collected on a bi-alkali photocathode. There are two features of import in the light pool: the central peaked region, and the pile-up near the maximum radius, beyond which there is no radiation.




Figure 2.3: Cerenkov light pool created by a single non-interacting electron traveling vertically through the atmosphere. The dots represent photoelectrons.

The peculiar appearance of the light pool can be explained by a careful consideration of the emission graph shown in Figure 2.4. The central graph in this figure shows how the radius of Cerenkov emission, as observed on the ground, varies with the height of Cerenkov production. The term "radius of emission" refers to the distance from the shower core (at ground level) at which radiation falls. This distance is a function of the height of production, and the opening angle of the Cerenkov emission, which itself is simply dependent on the local atmospheric density. Thus, the graph represents a characterization of the atmosphere. This particular curve is based on a simplified model of the atmosphere which assumes an exponential decrease in density with a scale height of 8 km. The analytical form can be derived as follows:

Proceeding from the Cerenkov angle condition,

$$\cos \theta_c = \frac{1}{\beta n} \quad (2.15)$$

with highly relativistic particles

$$\cos \theta_c \approx \frac{1}{n} \quad (2.16)$$

where n , the local index of refraction, is given by:

$$n = 1 + \kappa\rho. \quad (2.17)$$

Here κ is a constant and ρ represents the exponentially-falling (with height) density model of the atmosphere:

$$\rho = \rho_0 e^{-\frac{H}{8 \text{ km}}} \quad (2.18)$$

Combining Equations 2.16 and 2.17 yields:

$$\cos \theta_c \approx \frac{1}{1 + \kappa\rho}. \quad (2.19)$$

Expanding both sides:

$$1 - \frac{\theta_c^2}{2} + \frac{\theta_c^4}{4} + \dots = 1 - \kappa\rho + \frac{(\kappa\rho)^2}{2} + \dots \quad (2.20)$$

Dropping higher order terms in $\kappa\rho$ and rearranging gives:

$$\theta_c^2 = 2\kappa\rho \quad (2.21)$$

or

$$\theta_c = \sqrt{2\kappa\rho} \quad (2.22)$$

In the limit of small angles, the radius of emission will equal the product of the emission height and the emission angle:

$$R = H\theta_c \quad (2.23)$$

Thus, combining Equations 2.23, 2.22, and 2.18 yields an expression for the radius of emission as a function of the height of emission:

$$R = H\sqrt{2\kappa\rho_0} e^{-\frac{H}{8 \text{ km}}} \quad (2.24)$$

with $\kappa\rho_0 = 2.73 \times 10^{-4}$ and $\sqrt{2\kappa\rho_0} = 2.34 \times 10^{-2}$.

The maximum of this curve represents the maximum radius of emission for a Cerenkov-emitting particle, the so-called Cerenkov radius. In the simple model above, this maximum occurs at $H = 16$ km, for a Cerenkov radius of 138 m. A more realistic atmospheric model would yield a Cerenkov radius of ≈ 140 m.

From the figure, it can be seen that an emitting particle at the top of the atmosphere will fill a region on the ground which is somewhere between the center and the Cerenkov radius. The region it fills is denoted by the smaller graph above the central figure. As the particle progresses deeper and deeper into the atmosphere, it will fill regions on the ground closer and closer to the Cerenkov ring limit. Finally, at the broad, *Jacobian peak* centered on 16 km, most of the emitted Cerenkov photons will fill a very narrow spread of radii, leading to the pile-up effect seen in Figure 2.3. Only after the particle has passed to within 5 km of the surface will any radiation reach the central region of the light pool.

The consequence of this behavior is that heights of emission can be inferred from the patterns of Cerenkov light seen at the surface. Light pools with sharply-peaked central regions will tend to come from particles emitting low in the atmosphere, while prominent Cerenkov rings suggest emission higher in the atmosphere.

2.4.3 Emission from EAS

Section 2.4.2 explains describes the information content in the Cerenkov pool of an emitting particle. However, it refers to single emitting particles traversing the atmosphere, not the vast




Figure 2.4: Cerenkov Emission Graph.

distribution of particles associated with an extensive air shower. Nevertheless, the same type of information remains in the Cerenkov light pool when discussing radiation from EAS. This section describes an interpretation of electromagnetic cascades which illustrates the usefulness of the single-particle behavior.

These results emerge from interpretation of recent Monte Carlo studies of electromagnetic shower behavior performed at Minnesota [42]. The Monte Carlo program, CHESS, performs a full three dimensional simulation of electromagnetic cascades, including the emission of Cerenkov radiation. Further details can be found in Chapter 6.

2.4.4 Two Component Model

The essential conclusion on which the Dual Cerenkov Technique is based is that the Cerenkov-emitting electrons in an electromagnetic cascade can be segregated into two distinct populations, based on their angular deflection from the shower core. These populations are labeled "direct", and "indirect". (NB: this nomenclature has nothing to do with the direct and indirect measurement techniques described in Sections 2.1 and 2.2.)

The direct component of the cascade comprises high energy particles traveling close to the core of the shower. These particles have scattered from the core at angles less than the Cerenkov emission angle $\theta_c \sim 1^\circ$. Since they all share nearly the same trajectory, they generate light pools on the surface which are largely identical. Because these lateral distributions overlap, the particles essentially act as a single radiator, producing a single light pool, similar to Figure 2.3, but of greater intensity.

As the energetic particles of the direct component traverse the atmosphere, they undergo ionization energy loss and multiple Coulomb scattering. These effects act to distribute the particles to angles much larger than the Cerenkov angle. The resulting population comprises electrons scattered to angles up to $\sim 30^\circ$ from the main shower core. These particles form the indirect component of the cascade. With the extended angular spread of the indirect component, the *quasi-coherence* found in the direct particles is lost, and the subsequent Cerenkov lateral distribution fills a much larger surface area on the ground.

Figure 2.5 demonstrates the two populations with a schematic depiction of the extensive air shower front. This picture of the shower front, as a Cerenkov-radiating body, has emerged from




Figure 2.5: Cerenkov Source Model.

Monte Carlo investigations performed using the CHESH simulation package (See Chapter 6). In this model, the shower front-cum-radiator is essentially a 100 m diameter, 2 m thick body of electrons emitting to two distinct regions on the ground. Note that the indirect radiation fills the both the exterior indirect zone as well as the interior direct zone, but the converse is not true.

Given this characterization of the electromagnetic cascade, it is straightforward to describe the properties of a Cerenkov light pool as emitted from a full extensive air shower. The central, direct, core of the shower fills a region on the ground in essentially the same fashion as a single particle. Therefore, if a large number of direct particles exist at high altitudes (16 km in the toy model, 12 km in a realistic model), there will be a prominent ring feature to the interior portion of the light pool. (An example of this is seen in Figure 2.6). However, if a large number of direct particles survive to near ground level, there will be a prominent central peaked region to the light pool (See Figure 2.7).

As was discussed in Section 2.3.2, the number of particles surviving to any given depth in the atmosphere is largely dependent on the mass and energy of the initial cosmic ray. To review, an increased energy or decreased mass will lead to a deeper X_{\max} . Thus, characterizing the features of the interior region of a Cerenkov light pool can provide information about the mass and energy of a cosmic ray.

The total amount of light emitted by the indirect component of the cascade depends on the total track length of all the emitting particles in the indirect component. Since most of




Figure 2.6: Cerenkov Radiation From a 60° 10 PeV H EAS.




Figure 2.7: Cerenkov Radiation From a 0° 10 PeV H EAS.




Figure 2.8: Pulse Height vs. Radius for 4 Vertical Hydrogen Showers.

the cascade electrons will eventually enter the indirect population, the light output from this component should be nearly linearly dependent on the number of cascade particles. And, as detailed in Section 2.3.2, the number of cascade particles is linearly dependent on the energy of the shower. This behavior can be seen in Figures 2.8 and 2.9, which show a series of Cerenkov lateral distributions. The lateral distributions clearly scale with energy, but remarkably, when normalized by the shower energy, they exhibit a high degree of similarity at the radii beyond 150 m.

This property of the lateral distribution allows an estimation of the shower energy to be made by measuring the pulse height at regions beyond 150 m. Furthermore, this estimation is largely independent of variations in X_{max} . This is because the angular difference between two heights of X_{max} , as viewed from distances typical of the indirect zone (~ 200 m), are much smaller than the characteristic angle of emission (up to 30°) in the indirect zone.

With these features in mind, it is only a slight simplification to separate the Cerenkov light pool into two regions, defined for their primary dependencies. The indirect zone contains a largely energy-dependent light field, while the direct zone contains a largely X_{max} -dependent




Figure 2.9: Normalized Pulse Height vs. Radius for 4 Vertical Hydrogen Showers.

light field. Figure 2.10 shows a schematic representation of such a characterization. An additional 50 m radius zone in the center of the figure, the “Excluded Zone”, identifies a region in the light pool which the Monte Carlo cannot simulate, due to certain simplifications made in the interests of speed.

2.5 Dual Cerenkov Technique

In the current investigation, the properties of the Cerenkov radiation, described above, are exploited to make multi-parameter measurements of extensive air showers.

Assuming that the shower core location is known, the minimum detector configuration for taking advantage of the two-zone Cerenkov characterization is two independent instruments. Showers whose lateral distributions are such that their mass zones overlap one detector and their energy zones overlap the other are candidates for the Dual Cerenkov technique.

To make just these measurements, a pair of atmospheric Cerenkov detectors has been deployed to operate in conjunction with the Soudan 2 tracking calorimeter. The calorimeter, which samples high energy, low-deflection ($\theta < 1^\circ$) underground muons, provides a shower trajectory




Figure 2.10: Cerenkov Field Zones.

which can be used to determine the shower core location on the surface. This information, together with the data from the dual Cerenkov array, should provide a simultaneous estimation of the energy and mass of the primary particle.

2.6 Previous Composition Measurements

The accessibility of the knee region to relatively small air shower arrays and Cerenkov detectors has been exploited by a number of investigators. Figure 2.12 shows a recent summary of several measurements; it offers a glimpse of the large variety of measurements which have been made in the region of the knee. A brief synopsis of a small number of composition experiments follows.

2.6.1 Synopses

Soudan 1989, 1993

In one of the first deep underground/surface coincidence experiments, DasGupta *et al.*[43] deployed a small proportional tube array on the surface above the Soudan 1 experiment. The

results, reported in 1989, indicated a light composition. In 1993, a new version of the experiment was operated, this time in coincidence with the Soudan 2 detector [44]. This investigation, too, suggested a light composition.

Soudan 1997

Another Soudan 2-based cosmic ray investigation was completed in 1997 [15]. This analysis proceeded by examining exclusively the properties of deep underground muon multiplicities measured in the Soudan 2 detector. With this data, Kasahara *et al.* was able to draw conclusions about the incident cosmic ray flux. These results too suggested a largely light composition from 8×10^{14} eV to 1×10^{15} eV

Soudan 1998

In 1998 another experiment at Soudan 2 was concluded. This investigation was the precursor to the current investigation. It used a single atmospheric Cerenkov array operated in coincidence with the Soudan 2 calorimeter [45]. The addition of the Cerenkov data to the muon data allowed an independent measurement of the cosmic ray energy to be made. This greatly constrained the freedom in muon production, theoretically allowing a much more accurate estimate of cosmic ray mass. The results instead revealed serious weaknesses in the Monte Carlo programs used to simulate high energy hadronic interactions in extensive air showers. These results cast doubt on any composition conclusions drawn by experiments using muon multiplicity data.

DICE

The Dual Imaging Cerenkov Experiment, or DICE, comprises two 2 m diameter telescopes each equipped with 256 densely packed phototubes [46]. The telescopes, which are spaced 100 meters apart, use an imaging technique to make independent measurements of X_{\max} for extensive air showers.

Examining the energy range between 10^{14} eV and 10^{16} eV, DICE reports a light composition, tending towards purely proton beyond 5×10^{15} eV. These data are shown in Figure 2.12

KASCADE

The Karlsruhe Shower Core And Array Detector, KASCADE, is situated on the site of Forschungszentrum Karlsruhe, in Karlsruhe, Germany [47, 48]. KASCADE comprises 252 scintillator stations spread over 4×10^4 m². It also features a 20×16 m² finely segmented central hadron calorimeter. Muon detectors cover an area of 150 m², providing 0.5° accuracy in muon tracking.

KASCADE results are derived from simultaneous measurements of muonic, hadronic, and electromagnetic air shower components. Several independent analyses have been performed using various components measured in the detector. Initial results vary with the form of analysis used; considering electrons or hadrons leads to a composition growing heavier with energy, while a “central granularity” analysis suggests a composition growing lighter. The KASCADE collaboration points to deficiencies in the hadronic interaction Monte Carlos as a likely source of the discrepancy.

VULCAN

The VULCAN experiment is a nine-station wide-angle atmospheric Cerenkov array operated at the site of the Antarctic Muon and Neutrino Detector Array (AMANDA) at the South Pole [49]. The South Pole Air Shower Experiment-2 (SPASE-2), a 130 station scintillator array, is also deployed at this site.

Operating these three detectors in coincidence will yield a full multi-parameter measurement of all extensive air shower components. Initial results utilizing only the SPASE-2 and VULCAN instruments suggest a fairly heavy composition remaining constant throughout the knee. The VULCAN collaboration intends to include high energy muon data collected in the AMANDA array in their future analysis.

BLANCA

BLANCA, the Broad Lateral Non-Imaging Cerenkov Array, is a 144 station Cerenkov experiment located in Dugway, Utah, on the site of the Chicago Air Shower Array (CASA) [50]. BLANCA makes estimates of cosmic ray properties by using measurements of Cerenkov lateral distributions. Shower core locations are provided by CASA.




Figure 2.11: Mean Primary Mass as Measured by BLANCA. From [50].

BLANCA has collected a large number of coincident events in the region of the spectral knee. Their most recent results provide greater detail of the knee region than any previous experiment. Figure 2.11 shows their measurement of the mean cosmic ray primary mass throughout the knee region. A clear mass increase is seen prior to the knee, followed by a decrease and another increase beyond the knee. These features remain essentially the same, regardless of which interaction model is used.

HEGRA

Located in the Canary Islands at La Palma, HEGRA, the High Energy Gamma Ray Astronomy experiment, is a $200 \times 200 \text{ m}^2$ extensive air shower array. It consists of 154 scintillator stations, 17 geiger towers, and roughly 100 atmospheric Cerenkov detectors (AIROBICC). A pair of imaging Cerenkov telescopes also operate on the site.

Like BLANCA, the HEGRA experiment uses shower core information gathered from a scintillator array, and Cerenkov lateral distributions measured in the Cerenkov array to make estimates of cosmic ray properties. Recent results from HEGRA [51] indicate a mixed composition growing heavier throughout the knee.




Figure 2.12: Summary of Recent Cosmic Ray Composition Measurements. The hatched region signifies the range of direct measurements. The grey region covers the error range introduced by the hadronic interaction models, VENUS and QGSJET.

From [47].

2.6.2 Comments

The composition in the region of the knee has traditionally been a free-for-all controversy. For every study showing an increase in mass, there has been another study showing a decrease. To some extent, this remains true.

However, compared to the older generation of data, the data from newer experiments appear to be more alike than dissimilar. The results presented at the latest International Cosmic Ray Conference appeared to favor a composition which was mixed, or somewhat heavy prior to the knee, moving towards heavier through the knee. The results of many experiments, BLANCA, in particular, indicate that the instruments are improving greatly and that while the question of the knee remains open, it is perhaps not as wide open as it once was.

Chapter 3

The Soudan 2 Detector

The Soudan 2 detector is a high-resolution tracking calorimeter [52, 53] set in a $10\text{ m} \times 15\text{ m} \times 75\text{ m}$ cavern 710 meters below ground level. The cavern is located on the 27th level of the Soudan mine, the oldest, richest, and deepest mine in the history of the Minnesota iron range. The mine itself is located in the Soudan Underground Mine State Park in the town of Soudan in northern Minnesota. Begun in mid-1988 and completed in November 1993, Soudan 2 has so far collected over five kiloton-years of data, including over 57 million cosmic ray muon events. The detector was originally designed to detect the decay of nucleons [54, 55], specialized for the super-symmetric decay modes (such as $p \rightarrow K^+ \nu$) that water Cerenkov detectors find difficult to extract from background. Soudan 2 has been used in a number of other physics capacities, however, notably muon astronomy efforts [56], monopole flux measurements [57], cosmic ray composition studies [15, 42, 45], and neutrino oscillation investigations [58, 59].

The Soudan 2 detector produces triggers at a rate of roughly 0.45 Hz. Cosmic ray muon events occur at a rate of roughly 0.30 Hz, while natural radioactivity and electronic noise contribute an additional 0.15 Hz.

The detector itself, illustrated in figure 3.1, comprises 224 stacked calorimeter modules surrounded by an active veto shield of proportional tubes. The $1.0\text{ m} \times 1.1\text{ m} \times 2.7\text{ m}$ modules (whose size was limited to the size of the mine elevator cage) are collected in seven “walls” two modules high and eight modules across to create an operating detector $8.0\text{ m} \times 14.0\text{ m} \times 5.4\text{ m}$ in size. Each module is 4.28 metric tons and the total detector mass is 960 metric tons, with




Figure 3.1: The Soudan 2 Detector and Veto Shield.

an average density of 1.6 g/cm^3 .

The active veto shield is composed of over 1500 extruded aluminum proportional tubes. Typically 20 to 25 feet long, these tubes form a $14\text{ m} \times 31\text{ m} \times 10\text{ m}$ box around the central calorimeter. Its primary function is to greatly reduce the background for “contained” nucleon decay events by identifying events which originate from outside the detector fiducial volume.

With an overburden equivalent to 2090 meters of water, the muon energy threshold (for greater than 50% transmission probability) is $\sim 0.7\text{ TeV}$. The bulk of this overburden consists of Lake Vermillion Greenstone ($\rho \sim 2.80\text{ g/cm}^3$), interspersed with pockets of iron ore [15].

3.1 Module Design

Each Soudan 2 detector module, as shown in figure 3.2, consists of 241 1.6 mm thick corrugated steel sheets interposed by 7560 1.5 cm Hytrel (Hytrel 5544) plastic drift tubes. The sheets are stacked to create a honeycomb structure, inside which the tubes rest. The drift tubes, which extend along the global Z axis, are wrapped in copper-bonded $0.125\text{ }\mu\text{m}$ mylar sheets, with 21 conductors spaced along the length of the tubes. The tubes are filled with a gas mixture of 85% Ar, 15% CO_2 . The mylar-embedded copper strips in the center of the module are held at -9 kV while the outermost are grounded. This provides an electric field of 180 V/cm along the axes of all the drift tubes. See Figure 3.3. On both sides of each module, over the mouths of the drift tubes, is a wireplane assembly (Figure 3.4) consisting of 63 vertical $50\text{ }\mu\text{m}$ gold-plated tungsten




Figure 3.2: Soudan 2 Module.




Figure 3.3: Drift Tube Cross Section.

anode wires and 240 horizontal cathode pads, made of copper. The anode wires, which rest at 10 mm from the mouths of the tubes, are held at +2.2 kV. The cathode pads are located 5 mm behind the anode wires. The entire module is sealed with an insulated gas-tight steel skin.

Charged particles passing through a drift tube ionize the drift gas, releasing free electrons. These electrons drift in the module's constant electric field to the ends of the drift tubes where they are presented with the large voltage gradients from the anode wires. Under the influence of this field, the electrons cascade and the resulting avalanche is collected on the anode wire. An image of this collected charge is induced on the cathode and the pulses on both anode and cathode are then amplified and digitized at 200 ns intervals. The drift velocity of the electrons is nominally $0.6 \text{ cm}/\mu\text{s}$ and the maximum drift time is $83.3 \mu\text{s}$.

By matching pulses on an anode-cathode pair, y and x measurements can be obtained. The measured resolution in this plane is 0.38 cm [52]. Measuring the drift time of the electrons provides a z measurement, allowing a full 3-dimensional reconstruction of the original ionizing particle's path. The z-axis resolution is 0.65 cm and the aggregate angular accuracy of the detector is better than 0.25° for muon tracks, as determined by a detection of the moon's cosmic ray "shadow" (See Figure 3.5). Figure 3.6 shows a typical muon event, as displayed in the Soudan 2 event readout software.




Figure 3.4: Module Wireplane.

The starting point for a drift time measurement, T_0 , is the time at which the ionization electrons were deposited in the detector. This time can be determined, for example, by identifying the point at which a ionizing particle passes through the wireplane assembly. When it does so, it deposits electrons in a “blob” across several anode and cathode channels. This generates a unique signature and represents a point at which the trajectory and time are well defined.

3.2 Veto Shield

The active veto shield is a 2040 m² detector consisting of over 1500 extruded aluminum proportional tube modules. Each module contains two stacked layers of four hexagonal tubes to form a unit 0.2 cm wide × 0.1 cm thick × 6–8 m long. The modules are placed in contiguous panels surrounding the main calorimeter, covering some 97% of the total solid angle. The upper and lower faces of the veto shield are further enhanced by an additional layer of tubes, arranged perpendicular to the first, providing x and y locations for through-going particles.

The modules' electronics are grouped in four-tube units by layer and are triggered off the main detector. Preamp boards located directly on tube modules amplify and digitize the signals from the anode wires. Pulses over threshold are read out during a trigger along with the module ID and pulse time. The single-layer random rate due to Compton electrons generated

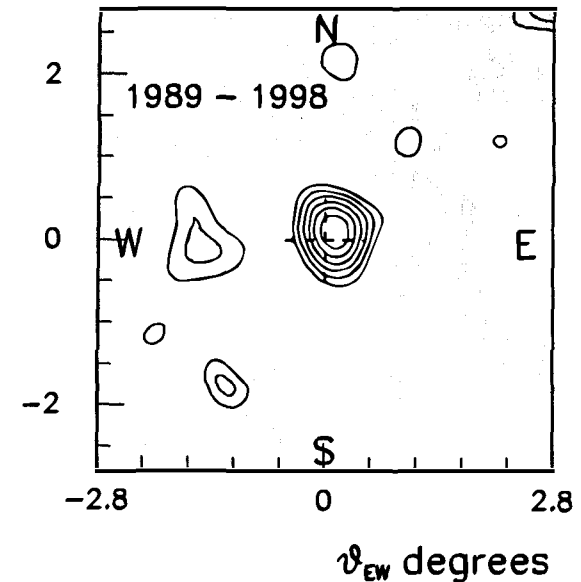


Figure 3.5: Soudan 2 Cosmic Ray Moon Shadow. The lines represent contours of normalized deviation, spaced at intervals $Z = 0.5$ above $Z = 2.0$. Thus the center of the contours represents a 4σ moon shadow

by natural gamma radiation is $\sim 30 \text{ m}^{-2} \text{ s}^{-1}$. This provides an unacceptably high rejection rate, so coincident pulses on both inner and outer module layers are required to generate a veto [58].

The efficiency of the shield as a muon detector has been studied by tracing well-reconstructed muon tracks from the main detector to the shield faces and checking for hits at those locations. Shield efficiencies for single-layer hits have been determined to be 98%. For simultaneous hits in adjacent stacked layers this drops to 92% [15].

The original function of the veto shield was to greatly reduce the background for nucleon decay events by identifying those events which originated from without the detector volume. As a low-cost low maintenance anti-coincidence detector, its resolution is inferior to that of the main calorimeter, especially for events along the axis of a module. Nevertheless, the footprint of the active shield is nearly four times that of the calorimeter. Thus, for muon measurements where the calorimeter size is a limiting factor, the addition of the shield as a primary detector is of great importance [45].

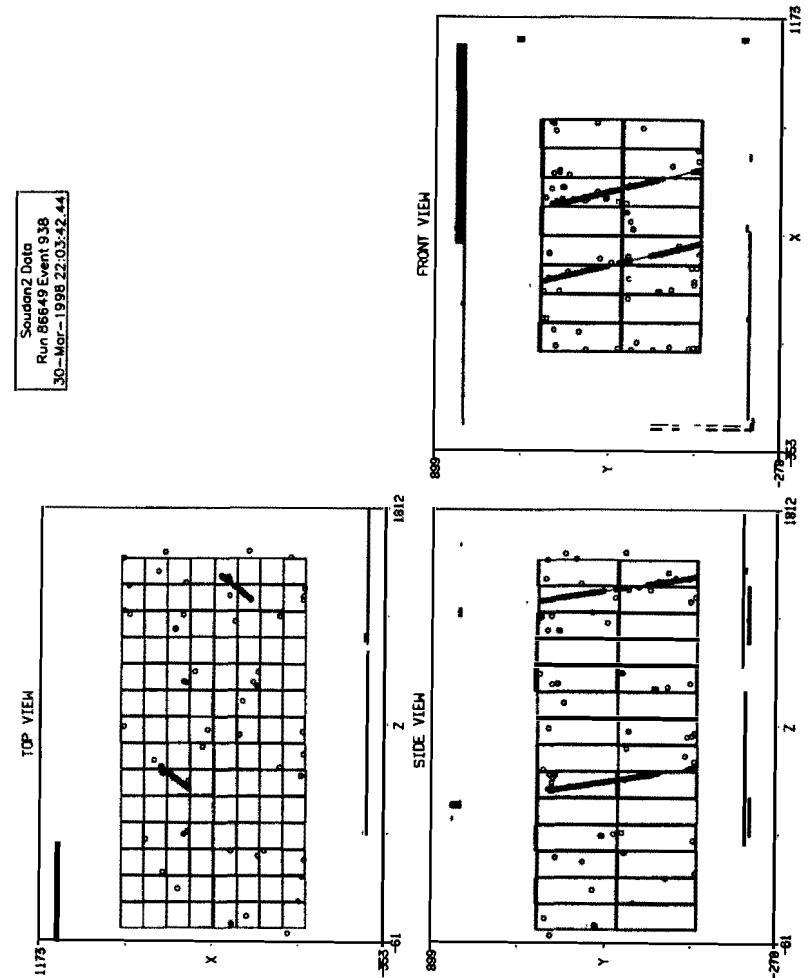


Figure 3.6: Online Software (STING) Readout

Chapter 4

Dual Cerenkov Array at Soudan

The Dual Cerenkov array at Soudan consists of two separate photomultiplier tube (PMT) arrays, named for their relative positions to the electronics trailer. The Upper Station Detector was originally built in the spring of 1993 to operate in coincidence with the Soudan 2 calorimeter to attempt to elucidate the composition of primary cosmic rays in the region of the knee. The Lower Station Detector was conceived in 1996 as an upgrade which might provide additional information lacking in a single array configuration. The two stations use identical photomultiplier tubes, high voltage distribution bases and Winston cone light gathering devices, but differ in their design and configuration.

Figure 4.1 shows the layout of the experiment site. The two Cerenkov detectors are shown in reference to the underground Soudan 2 calorimeter site and the main electronics trailer. The Cerenkov stations are labeled to indicate the boundaries of their mass-sensitive areas as well as their excluded zones within 50 m.

4.1 Lower Station Detector

The Lower Station Detector (LSD), installed in 1997, comprises an aluminum detector enclosure, 12 photomultiplier tubes and all associated detector electronics. It is connected to the primary experiment trailer by a cable run of some 440 feet. The detector control system, located inside the enclosure, consists of a computer controlled high voltage system, a microprocessor controlled




Figure 4.1: Detector Configuration Map.

Location	47.8° N, 92.2° W
Atmospheric Depth	978 g/cm ²
U.S.D. Distance to Soudan 2	745 m
U.S.D. Zenith w.r.t. Soudan 2	13.9°
L.S.D. Distance to Soudan 2	745 m
L.S.D. Zenith w.r.t. Soudan 2	20.9°
Array Spacing	113 m
Tubes/Detector	12 (9870, EMI)
Total Collecting Area	1.2 m ²
PM Operating Gain	0.5 – 2.0 × 10 ⁴
PM Response Time (FWHM)	12 – 19 ns
Trigger Threshold	2.0 – 2.5 × 10 ³ PE
Trigger Integration Time	300 ns
Digitizing Frequency	100 MHz
Digitizer Depth	40.96 μs
Dead Time (Read-Out)	200 ms
Live Time Fraction	~ 94%
Time Resolution (WWVB clock)	1 ms
Soudan 2 Trigger Line Delay	7 μs
Noise Rate per Tube	8 PE/nsec
Maximum Background Current	< 50 μA/tube
Data Rate:	
free running	0.045 – 0.10 Hz
coincidences	4 – 6 per hour
Average Duty Cycle (Winter)	3%

Table 4.1: Dual Cerenkov Array Specifications.

supervisor module, a high voltage safety interlock module, a current monitor, a LeCroy 428F Linear Fan-In/Out, and an Aperture-Control Logic Unit (ACLU), a device used to control the mechanical roof system (See Section 4.4.2).

The Lower Station Detector is located on wooden platform built on bare bedrock at the base of a shale mining detritus mound. Its coordinates are 195.24 m east and 181.36 m south in the global reference frame (whose origin is at the lower south east corner of Soudan 2). Its elevation is 466.19 meters, putting it 692.46 meters above the Soudan 2 global reference frame origin. See Figure 4.1.

4.1.1 Detector Enclosure

The Lower Station Detector enclosure, shown in Figure 4.3, was built to survive the harsh winters of northern Minnesota with very little maintenance. The main body of the enclosure consists of 0.060 inch aluminum sheathing glued and riveted over a 1x1 inch hollow aluminum tube frame. By design, the box is man-portable, very light and very rugged. It features a chain-driven four-pane aluminum roof which connects to a drive shaft driven by a 12 VDC high-torque gear motor (Dayton 1L478). The system removes and replaces the roof in a “tipping the hat” motion. The chain drive consists of several No. 25, 1/4” pitch roller chain sprockets geared to a ratio of 40 to 1. The drive shaft is supported by three 1/2” bore light duty pillow blocks mounted inside the electronics chamber of the enclosure. The chain is standard No. 25, 1/4” stainless steel roller chain. Chain tension is maintained by two fixed yolk roller units fixed on the exterior of the enclosure.

There are two linkage arms which connect the chain drive to the roof unit. Two additional arms provide stability. This design allows the roof to shed snow loads of up to 40 pounds without depositing any of the snow inside the detector itself. The system has proven itself reliable, having cycled several hundred times without a single failure. The time to open or close the roof system is roughly 45 seconds. Magnetic sensors are affixed at two locations on the detector to determine the door positions. These sensors are fed into the motor-control unit, the ACLU.

Measuring 42”x42”x60”, the full enclosure is divided into upper and lower sections. The upper section of the box contains 12 Winston cone-fitted photomultiplier tubes, a borosilicate



Figure 4.2: Lower Station Detector Enclosure.




Figure 4.3: Lower Station Detector Enclosure Schematic.

glass window, and a 12 V power supply to power the window heater. The lower section contains all of the detector electronics, the drive system and motor, and the ACLU. It is accessible through a lockable side access panel. Cables from the main electronics trailer enter through the floor of the enclosure.

The walls of the enclosure are insulated with silver-backed Dow Corning R-Max insulation, allowing waste heat from the electronics to keep the ambient temperature inside the lower section roughly 30 degrees warmer than outside without any additional heating.

4.1.2 High Voltage System

The Lower Station high voltage system consists of three single-width custom NIM modules with four high voltage channels each. The 12 channels can each supply one Watt of power in the range of 0 to -2000 V. Each channel is powered by an EMCO High Voltage DC-DC converter, model C20N. These tiny (0.74 in^3) inexpensive devices feature excellent temperature drift characteristics ($< 50 \text{ PPM}^\circ \text{ C}$) and superior voltage stability ($< 0.002\%$ ripple). The output voltage level for the devices is set in reference to an input control voltage at a ratio of 1 to -400 (i.e., 1000 mV in, -400 V out).

Figure 4.4 shows a schematic view of the high voltage modules. The design is elegant and simple, containing only nine active components in addition to the DC-DC converters. The actual modules are constructed with simple two-layer hand-routed circuit boards.

Philips PCF8591

Apart from the DC-DC converters themselves, the most important components in the high voltage system are the Philips PCF8591 devices. These pieces are single-chip, low-power 8-bit CMOS data acquisition units. They feature four multiplexed analog inputs, one analog output and a two-wire serial interface. One PCF8591 can provide four channels of 8-bit analog-to-digital conversion and one channel of 8-bit digital-to-analog conversion, all programmable via the two-wire I^2C (see below) serial protocol.

The DC-DC converter control voltage is set for each channel by the DAC unit of a PCF8591. This 8-bit DAC unit allows 7.125 V precision on the high voltage output settings. Each high voltage channel is monitored by one of the PCF8591 ADC units via a voltage divider feedback

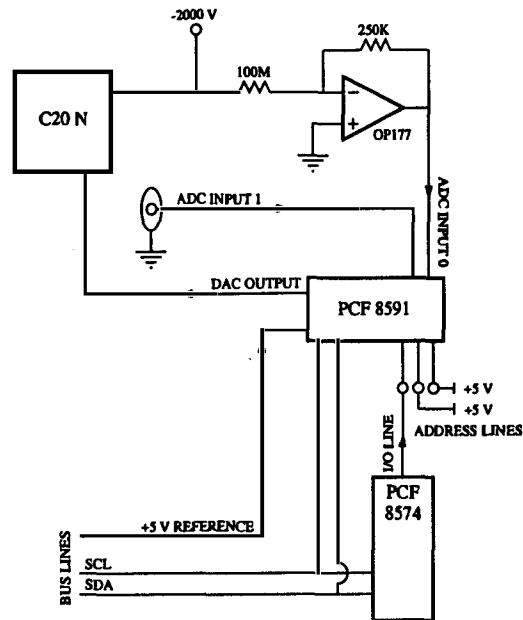


Figure 4.4: Lower Station High Voltage Module Schematic.

circuit. There are four PCF8591s in each high voltage module, one for each DC-DC converter. All of the ADC and DAC functions are referred to a high stability 5 V reference voltage supplied over bus lines by the Safety Interlock Module (See Section 4.1.3). One channel from each ADC unit is routed to a bulkhead lemo connector on the front face of the high voltage module. This channel is normally used to read anode current levels from the external signal splitter module.

The high voltage modules are programmed by a two-wire serial bus using the industrial communications Inter-Integrated Circuit (I^2C) protocol (See Appendix A). Routed through the NIM bin backplane, the bus connects the microprocessor controlled supervisor module to the high voltage and Safety Interlock modules. Making use of an Philips PCF8574 Remote I/O Expander to control an additional addressing line, the system can be expanded to accommodate up to 16 modules, providing 64 independent channels of precision high voltage. Additional bulkhead connectors may also be added to provide access for up to eight channels of additional 8-bit analog-to-digital conversion per module.

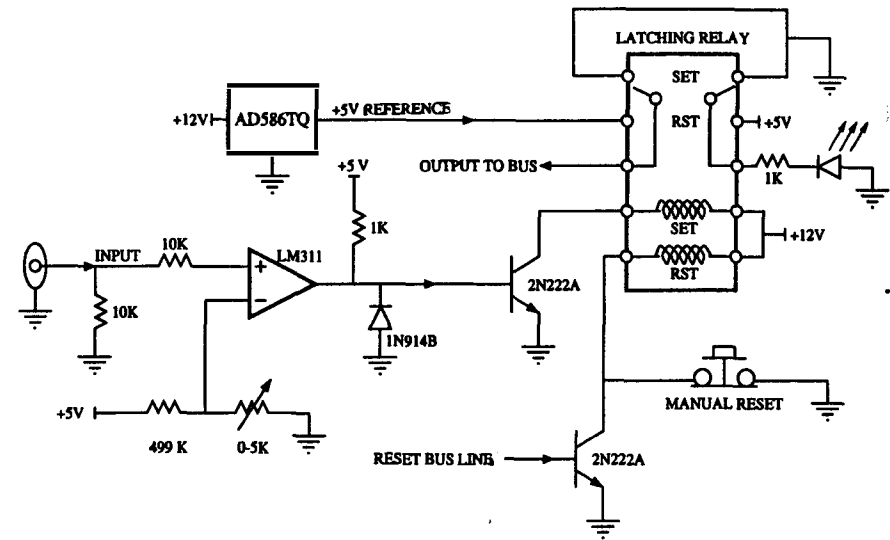


Figure 4.5: Safety Interlock Module Schematic

4.1.3 Safety Interlock Module

The Safety Interlock Module (SIM), Figure 4.5, is a device which provides a final line of defense in the event of a software tripping failure. Should the software system fail to shut down the photomultiplier high voltage in a high current situation, the SIM will do so automatically by grounding the reference voltage which the DAC converters use to program the DC-DC converters.

The SIM sums and monitors with a comparator circuit the voltage levels on its input sockets. When the voltage exceeds an adjustable safety level, a latched relay is tripped, grounding the bus line voltage distributed to the DAC chips. The output of the reference voltage unit is routed to an empty pin to await unit reset. The reference chip, an AD586TQ, is an extreme-precision military-grade temperature-stabilized voltage reference with a maximum drift of 7.5 mV over its operating range of -55°C – 125°C . The status of the SIM, *i.e.*, high or low, is monitored by a comparator circuit in the supervisor module. In the event of a hardware trip, the microprocessor can reset the SIM (by resetting the latched relay over a fourth bus line) and attempt to restart normal operations. The SIM can also be reset manually.

The use of a single reference voltage for all system analog/digital conversions allows for this very simple one-point failsafe feature, while simultaneously assuring constant high voltage settings over all tubes. The high voltage system will not operate correctly without a properly functioning Safety Interlock Module installed somewhere in the custom NIM bin.

4.1.4 Microprocessor Control Module

The heart of the Lower Station is the supervisor unit, the Microprocessor Control Module. This consists of a double-width NIM module, a single-board computer, and various support circuits. The computer, packaged by TERN Electronics as a "MiniDrive LCD Controller", is a 4.8 in x 3.4 in programmable, low-cost 16-bit industrial controller designed for high-reliability embedded applications.

The MiniDrive has a total of 40 I/O lines, including eight 4-bit ADC inputs, 24 bidirectional I/O pins, two high-speed integrated RS232 ports, and seven high voltage outputs. It features an 16-bit 8 MHz NEC V25 microprocessor, 1 MB of battery-backed static RAM, 128 KB of ROM, two 16-bit timers, watchdog circuitry, and power monitoring hardware. The unit costs roughly \$130.

The MiniDrive comes equipped with a MS-DOS work-alike operating system called Datalight ROM-DOS. This OS is fully compatible with MS-DOS 6.0 and executes directly from ROM. It can be directly loaded into ROM by the ACTF firmware utility which also comes with the MiniDrive. Once loaded, the unit can be accessed via serial line and operated like any personal computer.

The microprocessor is an 80386-compatible CPU which runs 16-bit code compiled for this architecture. Applications are compiled on the DAQ computer and the object code uploaded to the unit via XModem protocol. The compiler used in this investigation was Microsoft Visual C++ 1.52.

Apart from the MiniDrive unit, the supervisor module is also equipped with several comparator circuits and a breakout bus to distribute the microprocessor I/O lines. Additionally, there is a digital thermometer unit for monitoring the enclosure temperature and an external thermometer port, for monitoring the ambient temperature.

Operation

The microprocessor performs a number of control and monitoring tasks. First, it acts as controller-in-charge of the I²C bus, which it uses to control the high voltage modules and digital thermometers. Second, it monitors, via a comparator circuit, the status of the DAC reference voltage. If this voltage should drop below the comparator threshold, it indicates that a hardware high voltage trip has occurred, or that the reference voltage IC has malfunctioned. The microprocessor can then attempt to reset the Safety Interlock Module through a bus-distributed reset line. If this fails, there is a likely hardware malfunction and the microprocessor will shut down the Lower Station. Third, the microprocessor must communicate with the Aperture Control Logic Unit (Section 4.4.2) to open/close and monitor the roof system. The final task this unit must perform is to communicate with the primary DAQ computer. It does so through an RS232 interface, over which it receives commands and transmits status updates (see Section 4.4.3).

4.1.5 Signal Splitter/Current Monitor

In order to monitor the anode currents of the photomultiplier tubes, the signals from the bases must be split into AC and DC components. The Signal Splitter/Current Monitor Module (Figure 4.6) consists of a single-width NIM module with 3 columns of 12 front-mounted bulkhead lemo connectors. The leftmost two lemo connectors of each column are AC-coupled, to allow the transmission of all high-frequency signal data. A signal line shunts the DC anode current to a simple transimpedance amplifier which outputs to the third lemo connector. For every 1 μ A of input current, there will be 39.2 mV of voltage across the output port. The value of 39.2 mV corresponds to twice the precision of the PCF8591 ADC input units ($19.6 \text{ mV} = \frac{6000 \text{ mV}}{256 \text{ bins}}$).

The current sum of all the tubes is also provided by the Signal Splitter unit. The outputs of all the transimpedance amplifiers are added in a simple summing amplifier and the resulting voltage placed across a rear-mounted bulkhead lemo for input into the Safety Interlock Module.

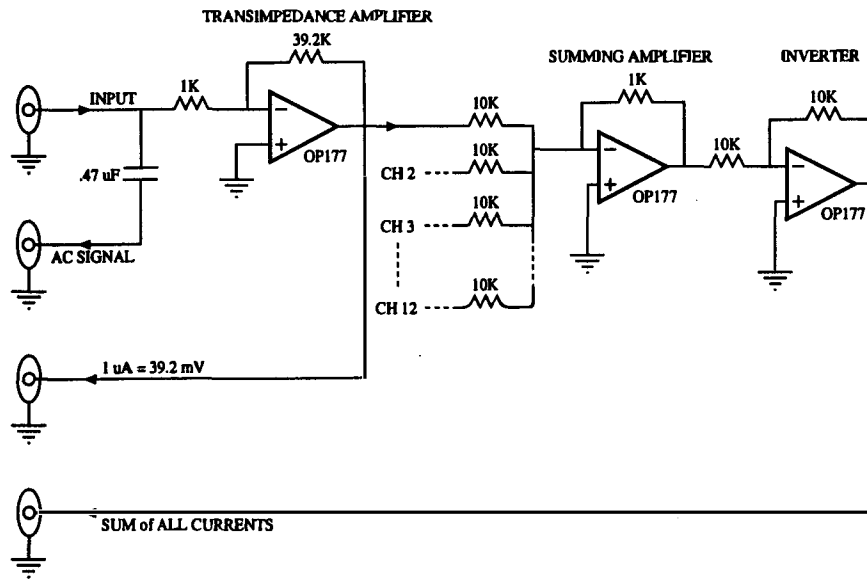


Figure 4.6: Signal Splitter/Current Monitor Schematic.

4.2 Upper Station Detector

The Upper Station Detector is located on a 5 m high tower some 50 m west of the main detector hut. This station, like the Lower Station, consists of a detector enclosure and 12 phototubes. Unlike the Lower Station Detector, the high voltage units and most ancillary electronics for the Upper Station reside in the main electronics trailer. A ~20 m cable conduit contains 12 RG-59 cables for high voltage supply and 12 RG-58 cables for signal return. The cable run extends up the side of the tower, unfortunately making the cables susceptible to radio frequency pickup.

4.2.1 Detector Enclosure

The Upper Station enclosure (Figure 4.7) consists of an 6' long by 2' wide by 4' high wooden frame sheathed with 5/8" painted plywood. The box is topped with a hemi-cylindrical cover which is mounted at the ends on two low profile turntable bearings. The configuration is similar to that found in large pig-roaster units. The cover itself is made of 1/8" textured wallboard with thick support ribs to maintain the curvature. An aluminum brace runs along the top to




Figure 4.7: Upper Station Enclosure Schematic.

provide rigidity along the major axis of the cover.

A 3/8" No. 35 roller sprocket is affixed to the outer half of each bearing holding the cover in place. The sprockets are turned by chains driven by a drive shaft which emerges near the bottom of each side of the enclosure. The shaft is mounted with four pillow blocks to an steel support beam built within the enclosure. A 12 VDC gear motor (same model as in the Lower Station Detector) drives the main shaft. Magnetic sensors are mounted at two locations on the enclosure body to determine the position of the roof system. These sensors are fed into the Aperture Control-Logic Unit.

The lower half of the detector enclosure contains supporting electronics and the drive-shaft system, and is accessible via a hinged access panel. The upper portion of the enclosure contains three boxes with four photomultiplier tubes in each box. These boxes rest side by side within the main enclosure; each is outfitted with an individual protective glass cover. Cables to and from the phototubes pass through the bottoms of these boxes and out the side of the enclosure.

4.2.2 High Voltage System

The Upper Station phototubes are powered by a legacy system inherited from the previous generation Cerenkov experiment at Soudan [45]. This system consists of 3 custom NIM High Voltage modules, with 4 independent channels per module. Each unit can supply up to 3 Watts of high voltage over a range of 0 to -3000 V with an accuracy of 11.72 V (8-bit).

The units are fitted with dual 8-bit analog to digital converters. These devices are used to monitor the anode current of the phototubes (via external transimpedance amplifiers) and the output voltage of the modules themselves. A high-resistance voltage divider reduces the output levels by 1000:1.

The units are controlled by the main data acquisition PC via an ISA bus extender card. Control and readout is performed by simply accessing the range of addresses 0x200-0x3FF. A detailed description of these modules may be found in [45].

4.3 Photomultiplier Tubes and Bases

The Upper and Lower Station Detectors utilize EMI 8970 photomultiplier tubes obtained as surplus from the defunct Harvard-Purdue-Wisconsin proton decay experiment. The tubes are designed for use in water Cerenkov applications; they are hemispherical, 130 mm in diameter with bialkali photocathodes. They are equipped with two focusing electrodes and 11 high gain Cesium-Antimony venetian blind dynodes. The photocathode is sensitive to the wavelength range 290–680 nm, with a broad peak in the quantum efficiency at $\sim 25\%$ between 320 and 400.

4.3.1 Design

The voltage distribution bases were carefully designed to optimize the capabilities of the photomultiplier tubes. The bases, shown schematically in Figure 4.8, were made to use only five active dynode stages with 300 V zener diodes on the final two stages. This allows the final two dynodes to operate at maximum inter-stage voltage throughout the tubes' operating range of 1200–1750 V. This improves linearity and time response.

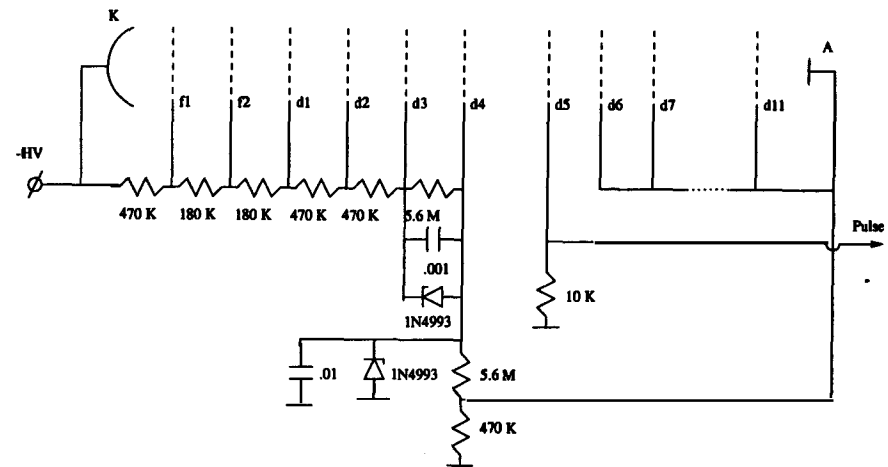


Figure 4.8: PMT Base Schematic.

Gain

The ambient light levels of even the darkest nights are substantial in terms of typical photomultiplier tube environments. On average, the night sky generates 8 photoelectrons per nanosecond on every photocathode of the Cerenkov array. This corresponds to 1.28 nA of photoelectron current. Operating at tube gains much higher than 10^5 would result in excessive currents, leading to an eventual deterioration of anode and dynode effectiveness due to damage of emission surfaces. Thus, the bases were designed to operate with a maximum final dynode current of 50 μA , corresponding to a maximum gain of 4×10^4 under normal ambient light conditions. By way of comparison, a 100 TeV vertical photon-initiated cascade will generate a Cerenkov field which induces a peak current of 50 μA at 80 m from the shower core [42].

Time Response

The previous generation Cerenkov detector at Soudan was designed to attempt to extract mass composition information from the shapes of Cerenkov pulses. Thus a primary requirement of the tubes was that they not significantly alter the waveform of the narrow input Cerenkov pulses.

With standard bases, however, the tubes could be expected to generate unacceptable large

~ 100 ns pulses in response to delta function input pulses. The primary contributor to this pulse width is inter-stage transit-time spread due to the width of the secondary electron velocity spectra (*i.e.*, the spread in the speeds of electrons emitted from dynode stages).

Fortunately, the design decision to limit the number of active dynodes (for gain purposes) also acts to improve timing characteristics. Reducing the number of stages and operating at near maximum inter-stage voltage allows the full-width half-maximum (FWHM) spread of the output pulses to be reduced from 100–200 ns to as low as 12–19 ns. The importance of this legacy constraint is greatly reduced in the present investigation, where only pulse integral information is being used.

Linearity

The linearity of a photomultiplier's output current is limited, in pulse mode, by space charge effects due to high dynode currents. Increasing inter-stage voltages, especially on the later dynodes, is the standard method of improving and extending the linearity of a tube. This benefit, coupled with the improvements in time response, strongly recommends operating the phototubes at maximum inter-stage voltage on last few stages.

4.3.2 Calibration and Testing

Gain

The gain of the photomultiplier tubes was carefully determined using a two step process. First, the relative gains of the phototubes were measured. This was done by collecting (via integrating ADC) histograms of pulse height responses to a standardized blue LED flasher. Such a histogram is shown in Figure 4.9. As a cross-check to the LED pulser measurement, several tubes were recalibrated using a piece of thick scintillator as a light source. Through-going cosmic ray muons provided short, Cerenkov-like pulses for collecting histograms.

Additional response histograms were collected from a single “benchmark” tube at several voltages. This allowed a full gain curve to be constructed. This curve is shown in Figure 4.10, along with the associated tube power spectrum. The total gain of the phototubes varies roughly as $(V-600)^2$. This behavior stems from the double-zener configuration in the base.




Figure 4.9: Sample Calibration Histogram.

After the relative calibrations were complete, the benchmark tube was subjected to independent dark-box pulser tests. In these tests, response histograms were collected with a LeCroy qVt for both the benchmark tube and a Quanticon phototube of known gain. This provided an absolute calibration of the benchmark tube, which, in turn, allowed the response histograms to be converted into absolute gains. Tubes for each detector were chosen from a small pool to minimize the gain differences between them.

Linearity

The linearity of the tubes was tested using a pulsed laser light source and a series of neutral density filters. The results of the test are shown in Figure 4.11. The figure shows the measured value of the peak output pulse current, “anode current”, *versus* the source light intensity, measured in mA, “linear peak current”. A straight line of the graph would indicate perfect linearity.

The results show that the tube response is approximately linear to almost 1.8 mA of peak current, where it deviates 10% from perfect linearity. The average background current, recall, is under $50 \mu\text{A}$ (corresponding a 100 TeV shower at 80 m). A current of 1.8 mA is a full 36 times larger than the typical background rate and corresponds to (via the linearity of Cerenkov pulse





Figure 4.10: PMT Gain Curve and Power Usage.

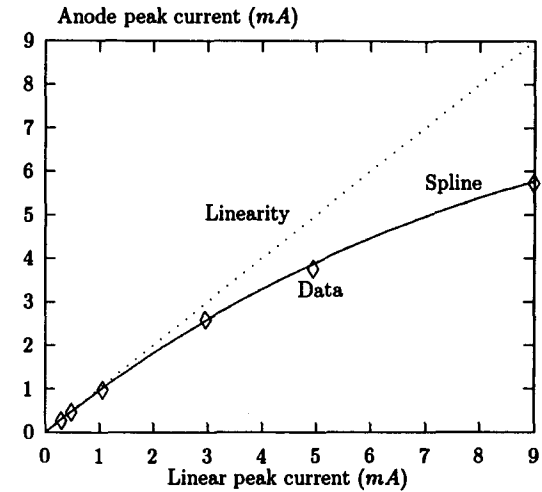


Figure 4.11: PMT Linearity Curve. Laser Output pulse versus measured tube peak current for several intensities of 5 ns laser pulse. The diamonds are data points, the solid line is a fit to the data. The dashed line is the line of true linearity.

height with energy) a shower of $36 \times 100 = 3600$ TeV at 80 m. The linearity deviations at these impact parameters and energies have been taken into consideration in the event reconstruction software.

Long-Term Exposure Effects

The response of the photomultiplier tubes to extended periods of high anode current was tested prior to this investigation [42]. A continuous $80 \mu\text{A}$ anode current was induced on a tube for 480 hours. Over 200 hours, an exponential decrease in the gain was observed, leveling off at $\sim 60\%$ of the original gain. The supposed cause of this degradation is damage to dynode emission surfaces caused by excessive electron bombardment. After the initial damage is done, there is little further change to the gain of any tube from additional exposure.




Figure 4.12: PMT Assembly Schematic.

4.4 Common Hardware Features

Apart from the photomultiplier tubes themselves, several additional components were used in both of the detector stations. A description of these units follows:

4.4.1 Windows and Winston Cones

To maximize the collection area of the photomultipliers, each tube is fitted with a Winston-like light collection cone [60] of reflective aluminum (Trade name: Alzac, Reynolds Aluminum). The cones are 25 cm in diameter at the top and 12.5 cm at the base. The tube assembly is shown in Figure 4.12.

Figure 4.4.1 shows the angular acceptance of the cone-tube combination. The acceptance is somewhat constant to roughly 45°, after which it rapidly drops to zero by 60°. Normally, such a cone would provide a flat acceptance up to 30°, beyond which it would drop to zero. However, the particular geometry of the PMT assembly (*i.e.*, the hemispherical tube *within* the volume of the cone) modifies the nominal behavior, causing the structure seen in the acceptance curve. The bumps are due to light passing through *two* faces of the photomultiplier.




Figure 4.13: Angular Acceptance of Tube with Winston Cone.

Over the top of the cones rests a 0.125 inch thick window of borosilicate glass. This pane protects the tubes from the elements during data-taking. To prevent vapor condensation on its surface, the windows are fitted with lengths of 0.001 inch gold-plated tungsten wire which form a simple resistive heater of some 40 W. The glass sheeting, made of the same material as the PMT windows, has excellent transmission properties from 400–600 nm, dropping sharply to zero at roughly 290 nm.

Though they nominally offer four times the collection area of bare tubes, the imperfect reflectivity of the aluminum and the transmission features of the windows reduce this amplification factor to three.

4.4.2 Aperture-Control Logic Unit

The operation of the two detector enclosure doors is controlled by the Aperture-Control Logic Units (ACLU), shown in Figure 4.14. The ACLUs regulate the power to the motors which drive these doors. Each enclosure, described above, is fitted with a pair of magnetic proximity switches which report on the location of the door along its cycle. The ACLUs contain standard TTL logic to monitor these switches and route power to the enclosure motors based on the door

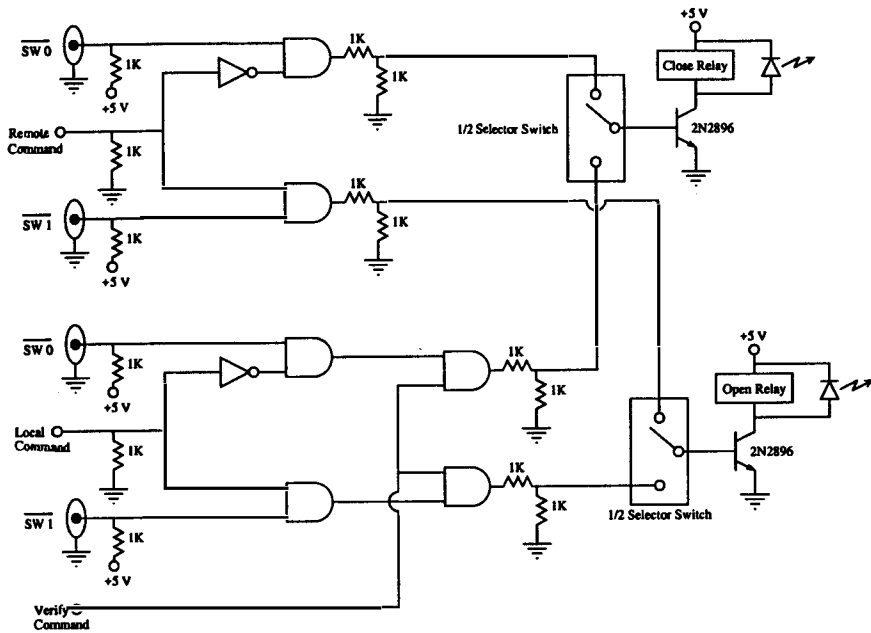


Figure 4.14: Aperture-Control Logic Unit Schematic.

position and input commands. Power is automatically cut off to the motors when the doors are fully opened or closed to prevent damage to the hardware. In the case of a logic failure, fuses insure that the motors will not continue to strain against a closed or open door.

The units can be set to local or remote mode with a front-mounted switch, shown in Figure 4.15. Commands to open or close the doors can be set remotely via TTL level when in remote mode, or by another switch when in local mode. The ACLUs bleed their own power from the 12 V input motor supply.

4.4.3 RS-232/422 Converters

Communication between the main data acquisition computer and the Microprocessor Control Module (Section 4.1.4) in the Lower Station Detector takes place using the RS232 signal standard. However, the maximum cable length supported by the standard is under 50 feet, while the cable run to the Lower Station Detector is over 440 feet long. A pair of RS232/RS422 converters




Figure 4.15: ACLU Switch Logic.

(B&B Electronics 422NOICR) is used to bridge the bulk of this cable distance. These converters transform the RS232 signals from the two computers into optically-isolated, differential RS422 signals. This allows communication at 90kbps over cable lengths of up to 4000 feet.

The units are very noise resistant and serve to protect the sensitive serial inputs of the two computers by optically isolating the long cable run, which is susceptible to lightning strikes. A null-modem converter must be in place at the lower end of the chain, to convert the DCE ("data communications equipment") output of the 422NOICR to a DTE ("data terminal equipment") configuration for input into the Microprocessor Control Module.

4.5 Data Acquisition Hardware

The bulk of the hardware used for acquiring data in the Dual Cerenkov Array can be separated into the categories of signal handling and trigger logic. Each group contains of a number of devices, both commercial and custom. The details of these units follow.

4.5.1 Signal Handling

Signal Path

The signal path of the Cerenkov array obviously begins with the photomultiplier tubes, which absorb photons and produce pulses of electrons. The signals from the tubes in the two detector stations take slightly different routes to their final destination at the waveform digitizing units. See Figure 4.16.

Anode signals from the Upper Station photomultipliers travel along twelve ~20 m RG-58U cables to the main electronics hut. The first portion of cable runs 5 m up the side of the enclosure

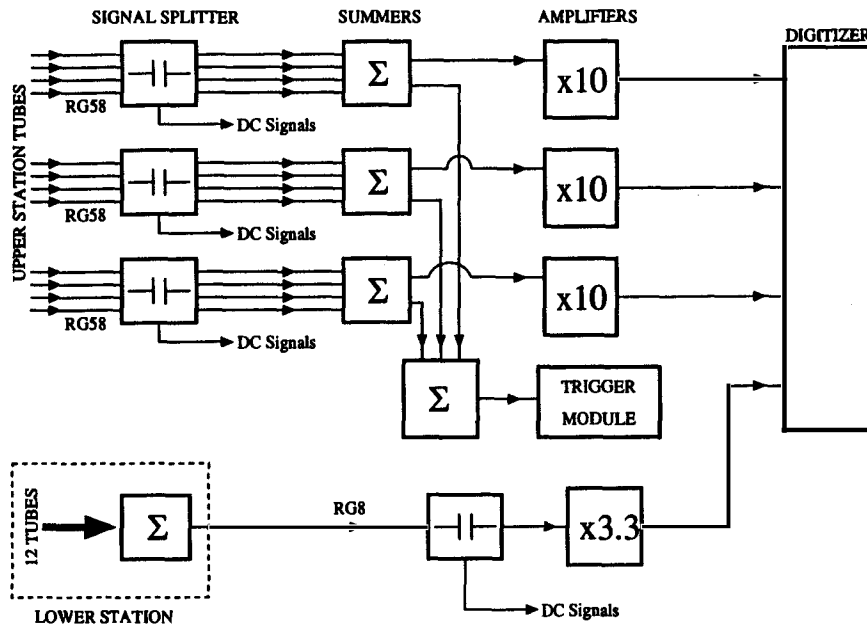


Figure 4.16: Detector Signal Path.

stand where radio frequency pickup is added to the Cerenkov signal. Once at the electronics trailer, the signals enter a custom unit in which the anode DC current is decoupled from each signal and fed into the feedback circuit of the high voltage system. The AC components of the signals are added in groups of four with a LeCroy 428F Linear Fan-In/Fan-Out (Labeled "Summers" in Figure 4.16). The LeCroy units feature 100 MHz input bandwidth and promise a noise contribution $< 750\mu V$.

From each of the three fan-in/fan-out units, an output is routed to a fourth summer. This unit adds the outputs of the three previous modules, yielding the grand sum of all 12 Upper Station tube signals. The summed signal is fed into the trigger electronics, where it is used in the trigger decision.

The second output of each of the three initial summer units is routed to a channel of a LeCroy 612AM amplifier. This variable amplifier unit has an input bandwidth of 140 MHz, and an RMS input noise of $< 50\mu V$. It amplifies the tube signals by $\times 10$ and outputs them to a LeCroy 6841 Waveform Digitizer. Each group of four summed tubes is digitized in a separate

channel of the 3 digitizer units.

To minimize noise pickup and pulse distortion, it was preferred that signals from the more remote Lower Station Detector be carried by a single long cable. Therefore, the outputs of the photomultipliers in the Lower Station Detector are summed on-site with a LeCroy 428F. The 12 phototube signals are added and passed along 135 m of low-loss RG-8 cable. The quality of the cable is such that the signal degradation over this 135 m cable is comparable to that over 20 m of RG-58.

The full Lower Station signal is passed to a signal splitter in the main electronics trailer. The AC component of the signal is passed to a channel of the 612AM and amplified $\times 3.3$. The output is then passed into a channel of the LeCroy 6841. The factor of 3.3 is chosen to equalize the signal amplitudes in the digitizer channels: Upper Station digitizer channels read 4 tubes amplified at $\times 10$ for a nominal total gain of 40; with 12 tubes in the Lower Station channel, an amplification of 3.3 provides parity.

Signal Capture

Three LeCroy 6841 Waveform Digitizers provide the final step in the signal handling process. These CAMAC modules are dual channel, 100 MHz 8-bit digitizers with a depth of 4096 samples ($40.96\mu s$). The Upper Station tubes require 3 input channels, while the Lower Station signal uses another. The two channels of the third digitizer are used to capture other interesting data such as the analog sum of the all Upper Station tubes (*i.e.*, the triggering signal), or spare Upper Station signal lines, which provides a "clean" sample of the RF pickup noise.

In addition to the waveform digitizers, a LeCroy 2249A integrating ADC is available to capture the Cerenkov signals. This provides an independent measure of the integrated tube signals and can be used to verify the accuracy of the software pulse integration routines. Though it is not typically used in normal running, this ADC module is a standard part of the calibration system.

4.5.2 Trigger Electronics

A schematic depiction of the detector trigger logic is shown in Figure 4.17. The heart of the trigger is the Trigger Module. This unit continuously integrates, with 300 ns time constant, the

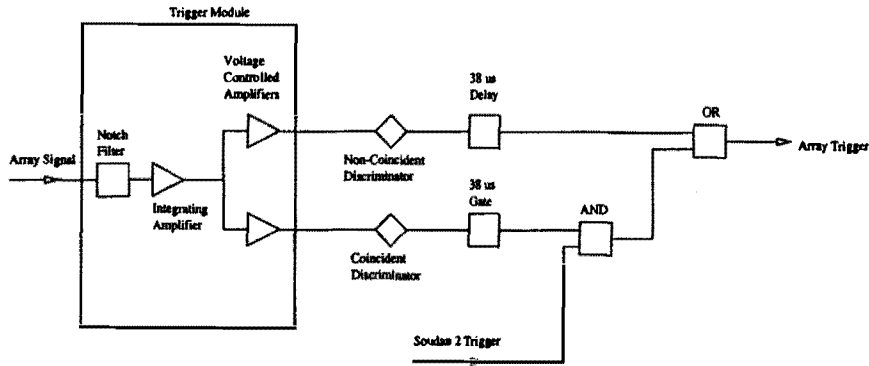


Figure 4.17: Trigger Logic.

summed signal from the Upper Station Detector phototubes. The signal is integrated to insure a trigger which is independent of the input pulse shape. The 300 ns integrating constant is a legacy from the last generation Cerenkov detector; it was set to insure triggering on wide-pulse (> 100 ns) Cerenkov signals.

The initial stage of the Trigger Module is an active notch filter which removes the contaminating 97.9 MHz radio frequency picked up in the elevated signal cables. This filter is not applied to the digitizer signals because they are more accurately cleaned with software techniques.

From the notch filter, the signals are integrated with a 300 ns time constant and the output is split into two channels. Each channel contains an independent voltage-controlled variable-gain amplifier which is programmed by the data acquisition computer via a CAMAC 12-bit DAC module. The outputs of these amplifiers pass through bypass switches and from there to a pair of external discriminator units. The outputs of these discriminators represent "coincident" and "non-coincident" pre-triggers, depending on the channel. The term "pre-trigger" is used to distinguish between the outputs of the initial discriminators and the full array triggers. Pre-triggers must pass additional logic requirements before there can be a true array trigger.

This dual pipeline system allows the detector to operate with two different thresholds. The pre-trigger with the lower threshold, the coincident channel, is used in coincidence with triggers received from the Soudan 2 detector. The higher-threshold, non-coincident channel, is able to generate an array trigger itself. Enabling non-coincident pre-triggers allows the Cerenkov array to run independent of the Soudan 2 detector. This provides periodic detector-state data and

allows purely autonomous atmospheric Cerenkov studies [42].

An additional feature of the Trigger Module, not shown in the figure, is a pair of relay-driven switches which allow the normal amplifier outputs of the module to be bypassed. Instead of firing off the input array signal, pre-triggers can be produced by a pulse generator. This is useful for calibration and detector testing.

The signal from the coincident channel of the Trigger Module passes from the discriminator to a 38 μ s gate unit. This gate is logically AND-ed with the Soudan 2 trigger line, which comes over twisted pair from the Soudan 2 detector itself.

Using a 38 μ s gate allows for significant variation in trigger arrival times from the underground detector. The calorimeter triggers at different times depending on the geometry of the underground muon event in the main detector. Most Soudan 2 triggers arrive at the logic gate within 20 μ s, but significant variation can be accommodated with the gate due to the large 40.96 μ s digitizing depth of the LeCroy 6841.

Under normal conditions, the threshold of the coincidence line is set so that coincidence pre-triggers occur at a rate of roughly 150 Hz. This represents a very low threshold and corresponds to very small Cerenkov pulses. The high pre-trigger rate is chosen to increase the triggering efficiency of true surface-underground coincidences, though it increases the accidental rate as well. Soudan 2 triggers occur at a rate of roughly 0.5 Hz while surface-underground coincidences (using the \sim 150 Hz pre-trigger) have a rate of roughly 4-5 per hour.

For larger Cerenkov pulses, the discriminators from both pre-trigger lines will fire. The output from the higher-threshold, non-coincident, discriminator is logically OR-ed with the output of the coincidence line's AND, described above.

The output of this final OR forms the true detector trigger. To insure that the gated AND has time to fire in the case of a large coincident event (*i.e.*, one that pre-triggers both coincident and non-coincident discriminators), the signal from the non-coincident channel is delayed for 38 μ s before it is passed the OR.

The detector trigger latches the WWVB clock and fires the digitizers. The digitizers then assert a CAMAC Look-At-Me (LAM) signal which causes the computer to begin the readout sequence.

4.5.3 Skywatcher

The Cerenkov array cannot operate in high light levels. Thus, for automatic operation, there must be a means of monitoring the ambient levels of light in the sky. To do this, a photomultiplier tube with a very low gain base is employed. This tube, which rests on the surface of the main electronics hut, is enclosed in a weather-tight enclosure to protect it from the elements. It is fitted with a steeply-sloped window and filament heater to shed snow. The anode current from the tube is read out by a LeCroy 2249A ADC and averaged over 30 seconds to obtain a reliable reading of the light level outside.

To slow the rate of anode and dynode deterioration, the base is designed to draw relatively little current at the operating voltage of 800 V. Despite the low gain, the detector is very sensitive. The ADC input saturates at light levels little brighter than full dark. Moon-lit nights are bright enough to be well above threshold, as are nights when cloud cover reflects light from the nearby town. This provides a quite reliable weather monitor, enabling the detector to automatically remain shut down during poor weather.

4.5.4 Other Electronics

A handful of other hardware components are employed in this experiment. These include a Kinetic Systems 3112 12-bit DAC, which is used to set the amplification levels on the variable gain amplifiers of the Trigger Module. Also, a Kinetic Systems 3063 IGOR (Input Gated Output Register) module, which is used in conjunction with a custom-made breakout module to provide 16 channels of TTL I/O. This is used to control and read-out the ACLU in the Upper Station Enclosure and to set the triggering mode of the Trigger Module.

Additionally, a custom-made CAMAC module is used to interface to the WWVB clock, providing a latched trigger time accurate to 1 ms. This time is read out over the CAMAC dataway by the data acquisition computer and used in the offline event-matching software.

The data acquisition computer, itself, is a custom-built PC-compatible system. It has an inexpensive 233 MHz AMD K6 processor and 64 MB of error-correcting EDO RAM. The operating system is Windows 95. A National Instruments PCI-GPIB card connects the computer to the LeCroy 8901A GPIB-CAMAC Interface module, which allows the computer to send commands and receive data over the CAMAC dataway.

4.6 Software

As with any modern computer-controlled physics experiment, this investigation requires a large amount of software to operate. Over 7000 lines of data acquisition code alone have been written for this experiment. This section contains an overview of several of the most important software components. Some of the following descriptions include rather technical details of the inner workings of the control programs. These are simply included for completeness and may be skipped by the uninterested reader.

There are two central components to the data acquisition software. The bulk of the code consists of the primary control program which runs on the main DAQ PC. This program controls the interface to the acquisition and triggering hardware and coordinates the entire DAQ process. Additional software is resident in the Microprocessor Control Module located in the Lower Station Detector.

4.6.1 Primary Control Program

All detector operations are monitored and controlled by the main data acquisition program, named simply DAQ.EXE. This program, which runs on a personal computer outfitted with MS Windows 95, controls (via a GPIB interface) all DAQ and triggering hardware and acts as the graphical user interface for the array (Shown in Figure 4.18 and Figure 4.19). Additionally, the DAQ program communicates with the Lower Station Microprocessor Control Module via RS232 serial line.

DAQ is a fully object-oriented, multithreaded application written with Microsoft Visual C++ 5.0 (SP3). In addition to the standard Microsoft Foundation Class (MFC) components, DAQ features custom ActiveX controls for data visualization and National Instruments LabWindows libraries for simple data analysis tasks. Serial commands are provided by components from the freeware Windows Foundation Class package.

Program operations are handled by three independent Windows thread objects, the GUI thread, the serial thread and the control thread. The GUI thread is the primary program thread from which the others are spawned. All data objects belong to this thread and it maintains control over all of the elements of the user interface.

The control thread is a so-called worker thread which oversees the main operations of the

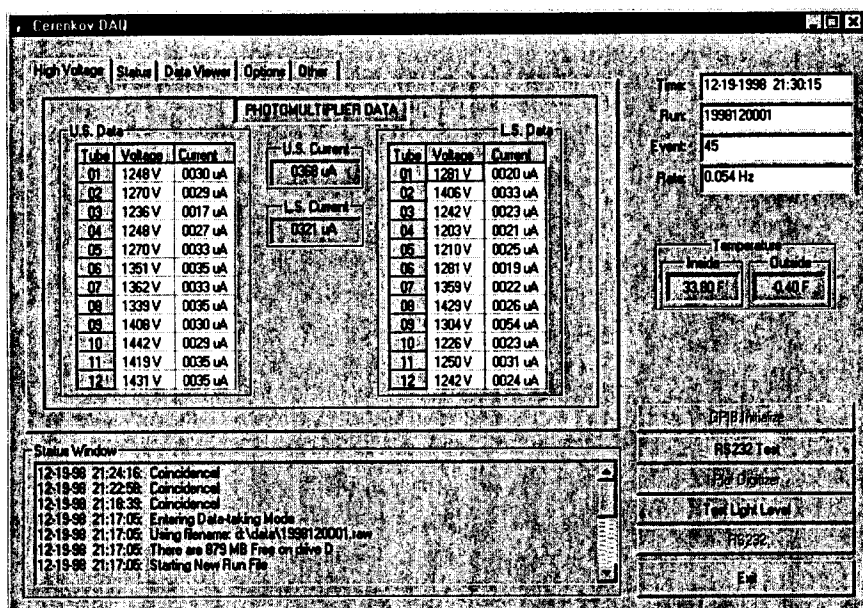


Figure 4.18: The Instrument Control Interface provides information on all detector systems, including tube voltages and currents, event rates, and temperatures. A log window in the lower left corner details recent system events. Buttons in the lower right corner provide access to certain instrument control functions.




Figure 4.19: The threshold-setting interface displays the results of a threshold versus trigger rate graph. The threshold decreases with increasing values of the x axis. To set the appropriate array threshold, a line is fit to this curve, and extrapolated to the desired event rate.

detector, including data readout and outrigger instruction issuing. The control thread consists essentially of a large loop with several nested “if” statements. The “if” statements poll system variables containing the current detector state and command conditions. When the appropriate conditions are met, the subblock of code is executed and the system variables updated accordingly.

Removing control operations to a separate thread allows continuous detector monitoring without the blocking penalty incurred by single-thread polling procedures. Future upgrades of the program should include a complete removal of all polling procedures in favor of pure Windows messaging techniques.

The serial thread performs all operations involving the Lower Station Detector. Certain Lower Station functions such as a detector status update can take up to 2000 ms to perform. A typical readout operation following a trigger, however, takes only 380 ms. Thus, all outrigger operations should be confined to an independent non-blocking thread to increase detector live-time. This thread, once spawned, simply awaits commands in the form of Windows Messages from the control thread. Once received, it performs the tasks of sending a command and waiting for a response before returning to the near-zero-overhead wait state.

Thus, in normal data-taking mode, this thread will asynchronously update system data objects with information such as tube voltages and currents. The control thread, which is responsible for the writing of the output data objects, retrieves the latest update from the outrigger. The detector state information for the outrigger, as stored in every data file, then, can be up to 2000 ms “old”. Due to the time-stability of the detector-state variables (voltages and currents), this is an acceptable tradeoff when the alternative is a two second readout delay for every event.

The use of a common data objects in this multi-threaded configuration requires that all data objects be protected from corruption due to simultaneous access from worker threads. This protection is achieved using simple CCriticalSection Lock and Unlock functions.

Event Storage

After a trigger is registered and the data read out from the electronics, it is stored in a DCEvent (for Dual Cerenkov Event) data object. This is a C++ data object which contains every

detector-state variable, plus data from all instruments, such as digitizers and ADCs. Each object is 28 KB in size. In a departure from past practices, coincident and free-running event objects are identical. This allows for streamlining the DAQ and analysis software. A greatly reduced data rate (~ 0.05 Hz) keeps storage requirements reasonable.

All of the event data objects are stored in run files of 200 events. A header object containing run number and start time is prepended to the file and a footer containing end time and total events stored is appended. This procedure is preferred over the use of a global run file object as it provides protection against data file corruption in the case of a computer crash.

The aggregate run files are named for their run number, with a “.raw” extension, *e.g.*, 1999020001.raw. These files are roughly 5.6 MB in size. In the case of a prematurely-ended run (due to software trip, user-intervention, etc.), the run is simply truncated and a footer written. These “shorts” are fully usable by the analysis software.

In the case of a coincident event, the DCEvent object is stored not only in the full run file, but in an individual data file as well. The files are named for their Cerenkov run and event number, with a “.dce” extension (1999020001_0152.dce). A running log of coincidence times (coincidences.log) is maintained. Should the log file or dce event collection become lost or corrupt, all coincidence data can be retrieved simply by scanning through the collected raw files.

The coincidence log, along with the ensemble of dce files forms the input to the event reconstruction code. This code, which runs on the local VMS cluster, combines data from Soudan 2 events and data from the dce events to create single FullEvent data objects.

4.6.2 Calibration Software

The calibration program is essentially a stripped-down version of the main DAQ application, with only a control thread and a GUI thread. Most of the features of the main DAQ program are disabled, but ROOT graphics components and the ability to write to ROOT data objects has been added. When run, the calibration routine allows the user to set voltages on single tubes of the Upper Station Detector. Because the serial thread has been removed, Lower Station commands must be sent manually through a terminal program.

Once the correct voltage is set, the program collects triggers and generates a histogram of

pulse heights read out at the ADC input. The data from the calibration program is plotted to screen and saved to disk as a ROOT histogram named for the tube number, voltage and histogram size (e.g., Tube09_1600V_2000.root).

4.6.3 Microprocessor Control Program

All operations of the Lower Station Detector are controlled and monitored by the Microprocessor Control Module. This module, described previously, features an NEC V25 microprocessor equipped with Datalight ROM-resident DOS. This enables it to run 16-bit code compiled for the x86 architecture. The code for the microprocessor was written using Microsoft Visual C++ 1.52. Hardware driver libraries for primitive I/O operations were supplied by the single board computer manufacturer, TERN, Inc.

The control software consists essentially of an interactive shell which provides access to all hardware manipulation routines. Once started, the software presents a menu of choices and a command prompt indicating readiness to accept instructions. The instructions themselves take the form of simple two-letter patterns. This method of command input allows the Lower Station to be easily controlled by both the DAQ program or by hand, should the need arise.

There are currently 20 commands understood by the Control Module, ranging from a number of high voltage setting options to several levels of detector status update. Most commands involve hardware I/O functions, such as reading/resetting the Safety Interlock Module, reading/setting high voltage levels, or reading/setting the status of the Aperture Control Logic Unit.

Unfortunately, the V25 does not offer bit-level hardware access. Thus every bit-level I/O operation requires a byte-level peek-and-mask procedure. To isolate the messy hardware details, a mini application programming interface (API) was written. This allows for easy hardware manipulation without the need to directly call complicated poke and peek functions. The improved ease of use more than compensates for the slight increase in system overhead; new hardware functions be easily added with a few lines of code.

4.6.4 Inter-Integrated Circuit Protocol (I²C)

All high voltage and temperature-monitoring hardware utilize the Philips I²C (Inter-Integrated Circuit) serial protocol. This two-wire serial bus provides a robust and effective means of communicating with and controlling simple configurations of specialized integrated circuits.

Though many microprocessors have integrated I²C command sets, the V25 does not. This necessitated creating a series of hardware I/O commands to access the I²C serial bus. The details of this software are interesting, but perhaps beyond the scope of this chapter. Instead, they have been relegated to Appendix A.

4.7 Operation

A brief overview of the detector operations flow is included below. The true logic chart followed by the control software is quite sophisticated, and can recover from many hardware malfunctions.

When started, the detector immediately begins monitoring the ambient light level with the Skywatcher system. When the level drops beneath threshold, a command is sent to the Lower Station Control Module and to the Upper Station ACLU. The Lower Station Control Module receives its signal and send the appropriate TTL level to the local ACLU, which opens the door.

The main DAQ PC monitors the levels of the Upper Station door sensors and periodically inquires the Lower Station Control Module about the status of its door. When it has been determined that both doors are opened, a command to run the high voltage is sent to the Command Module and a local voltage-setting routine runs on the main DAQ PC.

After the DAQ PC sets the correct levels on the Upper Station phototubes, it inquires the Microprocessor Control Module about the state of the Lower Station tubes. When they are ready, the DAQ PC begins the threshold-setting routine. This step, which essentially consists of plotting trigger rate versus threshold (see Figure 4.19), ensures that the array triggers are set to a level which is appropriate to the local ambient light conditions. Greater light contamination will result in lower thresholds.

After the thresholds are set, the DAQ program proceeds to collecting data, which it will continue to do until one of two things occur. The first is an excursion in the data rate, either higher than usual, or lower than usual. Either of these are likely consequences of a change in the

ambient light level. The DAQ program will respond by re-running the threshold-setting routine, and continuing to take data. The second event which can halt data-taking is a software trip condition, which occurs when the total current from either detector station exceeds a threshold. In this case, the DAQ program will cut the voltage to the Upper Station tubes and instruct the Command Module to do the same. Finally, it will return to monitoring the ambient light levels in anticipation of subsequent dark periods.

In the case of an approaching dawn, the usual result is a rate excursion followed by a failure of the threshold-setting routine. The routine is designed to fail when reasonable rates cannot be set on the tubes, usually due to excessive light levels or some hardware failure. The DAQ program responds to such a failure by shutting down the high voltage and returning to monitoring the ambient light levels. Note that to reduce wear on the door systems, the doors only close when the ambient light level reaches a level somewhat above the threshold for shutting down the high voltage. This accounts for periodic local light increases which would otherwise cause the doors to open and close repeatedly throughout the night.

Chapter 5

Data Handling

The data for this investigation were gathered during the winters of 1998 and 1999. Each of these data seasons was plagued by extremely poor weather conditions which reduced the number of data-favorable nights. Over these two years, roughly 80 hours of detector on-time were achieved for a duty cycle of under 1%. This is significantly less than could be achieved by even one season of favorable weather. In the 80 live hours, 486 dual Cerenkov events were collected, all of them being nominal coincidences.

In this chapter, the procedures used to prepare this data for final analysis are discussed. These procedures include the matching of coincident events, the conversion of data formats, and the consolidation of data into integrated event files. The details of the signal processing and pulse matching procedures are offered as well. A discussion of event characterizations is included.

5.1 Data Consolidation

Due to the great quantity of Soudan 2 data, combined with the three platform (Cerenkov data: WIN95, Soudan 2 data: VMS, Analysis routines: UNIX) nature of the data pipeline, the task of event matching and consolidation requires several steps. For completeness, the following paragraphs describe in detail the entire process.

As mentioned in Section 4.6.1, all coincident events are logged to the `coincidences.log` text

file on the main DAQ computer. This file contains the information needed to match a coincident event to its counterpart in the Soudan 2 detector, including Cerenkov run number, Cerenkov event number, and Cerenkov trigger time (as obtained from the WWVB clock).

5.1.1 Run and Event Matching

The coincidences.log file is used as input to a hybrid C++/Fortran program, matchrun.exe, which runs on the VMS cluster. This program steps through the log file reading Cerenkov event times. It queries the VMS-resident Soudan 2 run database to determine the Soudan 2 run and tape number corresponding to the specific Cerenkov event time. In the case of a match (which occurs nearly 100% of the time), the Soudan 2 run and tape number are appended to the log entry and written to a new file called runmatches.txt. Matchrun.exe also produces a series of com files which contain the commands necessary for extracting the proper Soudan 2 events from their tape archives.

After the appropriate Soudan 2 run files have been collected from the tape archives, the runmatches.txt file is processed by another hybrid program called matchevt.exe. As with matchrun.exe above, matchevt.exe steps through the runmatches.txt file reading the Cerenkov event times. However, armed with the associated Soudan 2 run number, it opens the appropriate Soudan 2 run file and compares each individual event time in the run to the that of the surface counterpart. A small time jitter in the WWVB clocks is normal, so events which fall within 4 ms of each other are considered matches. (Because the minimum time between Soudan 2 events is ~ 100 ms and the minimum time between Cerenkov events is ~ 400 ms, the event matching is non-ambiguous). About 96% of nominal coincidences are matched to a real Soudan 2 event, meaning there are few spurious triggers over the Soudan 2 trigger line.

5.1.2 Event Extraction

The matchevt.exe program creates three output files. The first, called soap.titles, forms the input to the pickev.exe program, which extracts all matched Soudan 2 events from their associated run files and stores them in a single 4bcb file. This is done simply to reduce storage overhead; each matched Soudan 2 event is only some tens of kilobytes in size while the run file it's contained in is ~ 25 MB.

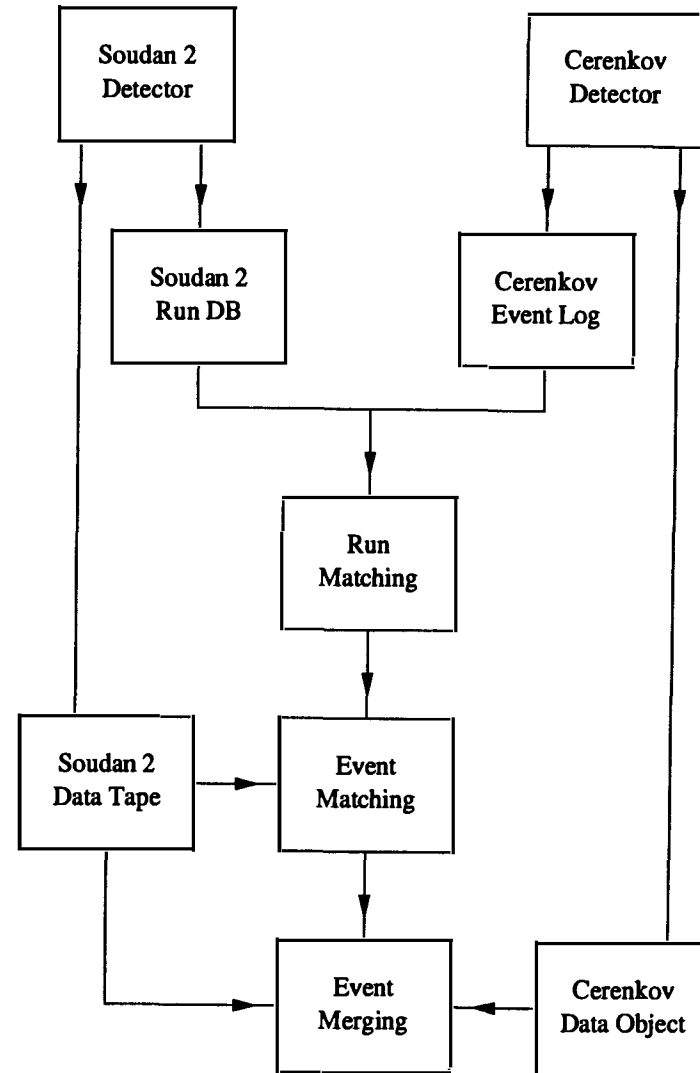


Figure 5.1: Event Matching Procedure.

The second output from `matchevt.exe` is the `evtmatches.txt` file. This file contains run and event numbers for the matched Dual Cerenkov and Soudan 2 events, including the trigger times recorded by each detector. The time difference in two triggers is also written. This file serves as a master event log, containing all the essential matching information about every coincident event. It will be used eventually to match the Cerenkov event objects to the Soudan 2 event objects, whose extraction is described below. The third output file of `matchevt.exe` is a file called `titles.extract`. This file contains a series of SOAP commands which serve as input to the `extract.exe` program, which is described below.

5.1.3 Event Consolidation

The next step in the Soudan 2 event handling process is to extract the event data from the matched Soudan 2 event files. This step is performed as a SOAP job on the VMS cluster. Using the file `titles.extract` as its input, the SOAP job, called `extract.exe`, opens each matched Soudan 2 event file and retrieves the requested data banks stored with the event.

Up to 410 data variables, depending on the number of muon banks found (up to a maximum of 20), are read out and written to a new `s2e` file. The data is written as a raw 410-float (REAL*4) vector to a binary file named for the Soudan 2 run and event number (e.g., `87299_1080.s2e`). These raw data files are used to maintain portability between the VMS platform, where all the Soudan 2 data resides, and the Digital Unix platform, where all analysis takes place.

In the final consolidation step, the `s2e` files produced by `extract.exe` and the log file `evtmatches.txt`, (produced by `matchevt.exe`), are transferred to the Unix cluster, where the Dual Cerenkov `dce` files (direct output of the DAQ program) already reside. On the Unix cluster, the Merge program reads the `evtmatches.txt` file line by line and opens the matched pairs of `s2e` and `dce` files. The `s2e` files are converted “on the fly” to `CS2Event` objects and then, along with the `dce` objects, are consolidated into a composite `TFullEvent` object. This object is written to a file named for the Dual Cerenkov run and event number, with a extension `.evt` (e.g., `1999010012.0123.evt`). The `evt` files contain all the relevant event data from both the Soudan 2 and Dual Cerenkov detectors. They form the input to the analysis routines.




Figure 5.2: Sample Cerenkov Signal.

5.2 Signal Processing

Figure 5.2 shows a typical sample of digitized photomultiplier output from a Cerenkov event. Apart from the actual Cerenkov pulse, located near 250 time bins, the most obvious feature of the sample is the large sinusoidal noise contamination seen throughout. This component originates from a local 97.9 MHz radio station (WEVE, Easy Listening) which is picked up in the elevated Upper Station signal cables (see Section 4.2). Because the 100 MHz waveform digitizers have a Nyquist frequency of 50 MHz, the carrier wave of the radio station is aliased to 2.1 MHz. This corresponds to a wavelength of some 50 bins, as seen in the figure.

Additional features of the data sample include an overall offset from the zero point, (which actually is at 255 on the ordinate), and a series of small one or two bin excursions seen throughout the signal. These small spikes are largely due to rapid fluctuations in the night sky background light inducing changes in the anode current of the phototubes. A component due to instrument noise, particularly the L428F Linear Fan-In/Fan-Out modules, also contributes. The signal offset simply stems from DC anode current induced by the steady light level of the night sky.




Figure 5.3: FFT of Sample Cerenkov Signal.

Clearly, to accurately reconstruct the Cerenkov signal against this background, the removal of the noise features is required. However, removing the noise without affecting the signal is non-trivial and a closer look at the characteristics of the noise will prove helpful.

5.2.1 Noise Characteristics

To describe the process by which the various noise components are removed from the digitized sample, it is helpful to examine the signal in frequency space. Figure 5.3 shows the frequency power spectrum produced by a fast-Fourier transform of the signal in Figure 5.2. The 0-frequency component corresponding to the offset in the signal has been removed, and the remaining features of interest include the huge spike centered on 2.1 MHz, and the relatively constant contribution from all other frequencies. This is shown more clearly in Figure 5.4.

The challenge of the filtering software is to remove the large 2.1 MHz contamination without reducing the spectral power of the actual Cerenkov signal at this frequency. It is not acceptable simply to remove a small band of frequencies about the 2.1 MHz peak. This is due to the nature




Figure 5.4: FFT of Sample Cerenkov Signal, Log Scale.

of the Cerenkov signal's frequency distribution itself. In the time domain, a Cerenkov pulse is typically less than 50 ns in width, approaching a delta function. Thus, in frequency space, a Cerenkov pulse approaches a plane wave; it is a broadband distribution.

Figure 5.5 and 5.6 show a simulated Cerenkov pulse in the time and frequency domains. It can be seen that not only is there significant spectral power near 2 MHz, but the spectrum actually peaks just below this region. In the limit of an infinitely narrow Cerenkov pulse, this will move to a constant offset in frequency space.

With these two components, a model of the digitizer signals in frequency space can be generated. Figure 5.7 shows the power spectrum obtained from a simulated Cerenkov pulse superimposed with a strong 2 MHz contamination. Compare this model with the actual digitizer output spectrum shown in Figure 5.3.

Clearly, the significant contribution of low-frequency components to the overall shape of a Cerenkov pulse requires that care be taken in the filtering process. Simply removing a band of frequencies about 2.1 MHz would not only remove the noise, but a modest portion of the




Figure 5.5: A Simulated Cerenkov Pulse.




Figure 5.6: A Simulated Cerenkov Pulse Power Spectrum.




Figure 5.7: Simulated Power Spectrum of a Cerenkov Pulse with Noise.

Cerenkov bandwidth, as well. The result is a poorly-reconstructed Cerenkov pulse. Figure 5.8 shows the effects of just such a reconstruction.

5.2.2 Noise Filtering

As shown above, the nature of the Cerenkov and noise signals requires some care to be taken in the filtering process. In practice, the filtering is done in a series of four steps, detailed below.

The initial step is to remove the DC offset from the digitizer signal. This is done by applying a fast Fourier transform (FFT) to a 2048 sample selection of the data, centered on the likely location of the Cerenkov pulse. In frequency space, the amplitude and phase of the lowest frequency bin of the signal are set to zero. The inverse transform is then applied and the original signal, minus the offset, is the result.

The data sample is then scanned for windows of “unlikely” signal. Unlikely in this instance simply means that the signal in a group of adjacent bins deviates from the mean of the entire sample by some fixed amount. Technically, the window must exhibit a Z and χ^2 fluctuation which is only 5% likely to occur by chance. This parameter can be tuned to alter the stringency




Figure 5.8: A Poorly Filtered Cerenkov Pulse.

of the filter.

The contents of windows which have been labeled unlikely are then replaced with the average of six bins located 49,50, and 51 bins ahead and behind the replaced bin. 50 bins corresponds to the wavelength of the dominant 2.1 MHz noise contamination. The result of these replacements is to erase peaks in the sample, whether from Cerenkov pulses or other sources. The criteria for bin smoothing are such that it is almost always the Cerenkov pulse alone which gets replaced. What remains after this step is essentially a pure sample, in the time domain, of the digitizer noise background. This includes amplifier noise, night sky fluctuations, and WEVE interference.

The third step in the filtering process is to take the pure noise sample generated above and transform it via FFT to the frequency domain. Once in frequency space, an inverse Bartlett window is applied to the data. This is a linear symmetric 4 MHz selection window centered on the 2.1 MHz WEVE peak. Outside the Bartlett window, all frequencies are set to zero, such that the small selected band comprises mostly the pure radio interference, with small portions of amplifier and night sky noise as well. Note, it should not contain any components of the




Figure 5.9: Filtered Cerenkov Signal.

Cerenkov signal, as they were removed previously.

In the fourth and final step of the filtering process, the frequency representation of the 2.1 MHz noise sample is converted back into the time domain. Here it will essentially be a 2.1 MHz sine wave, broadened slightly by the contamination of neighboring frequencies. This noise wave is then subtracted, in the time domain, from the original digitizer signal. This provides a sample nearly free of oscillation, but with all Cerenkov signals not only intact, but unmodified by frequency leeching.

Figure 5.9 displays the results of applying the preceding filter process to Figure 5.2. Note the absence of the 50 bin sine wave, and the continued presence of the unaltered Cerenkov pulse and night sky fluctuations. Figure 5.10 shows the power spectrum of the same signal. Note the continued presence of spectral power in the 2 MHz range. The amplitude in this region is comparable to that found in the Simulated Cerenkov pulse spectrum of Figure 5.3.




Figure 5.10: FFT of Filtered Cerenkov Signal.

5.2.3 Pulse Matching

Once the digitizer samples have been filtered, they can be scanned for actual Cerenkov pulses. The tradeoff in pulse matching is between accuracy in pulse location-finding and sensitivity to smaller pulses. Pulses may be located more precisely by applying more stringent search criteria, but this necessarily will exclude some smaller pulses which would otherwise be matched. Those pulses which are matched, however, will be reconstructed with very little error. Note that as Cerenkov pulses decrease in size, they begin to look very similar to the spikes associated with night sky light level fluctuations, making them virtually impossible to extract from the sample.

Cerenkov pulses are located in the filtered sample sets in the same fashion used to remove from them in the filtering process: windows of unlikeliness are sought in the sample. However, a different set of search criteria are used. Parameters such as search window size, and unlikeliness cutoff can be set to tune the efficiency of pulse location. A simulation using actual night sky noise data was run to test over 40000 parameter configurations to determine which provided the best combination of pulse-matching efficiency and signal-reconstruction accuracy.




Figure 5.11: Pulse Matching Efficiency versus Window Size and Pulse Depth.

Matching Efficiency

Figure 5.11 shows a small sample of the configuration space comparing pulse matching efficiencies. The variation of pulse-matching efficiency is compared with the size of the search window for a range of input pulses. It can be seen that smaller search windows provide marginally better matching efficiency, though this advantage is minimized as the pulse depth increases. Note that “pulse depth” refers to pulse height for negative pulses.

By analyzing such efficiency graphs, a set of search parameters was chosen for the final analysis. The resulting efficiency vs. pulse height graph is shown in Figure 5.12. It can be seen that 50% of pulses which are six digitizers bins deep are matched accurately. This efficiency rises to nearly 100% for pulses which are twice as large.

Reconstruction Accuracy

Once a pulse is located, the pulse integral is calculated by summing a number of bins before and after the actual peak location. The default values for the pre and post bin summing are two, and six. Accurate location of the peak is important to insure that the entire pulse is included in the integral. This could also be achieved by increasing the size of the integration window. However, since this reduces the signal-to-noise ratio of the result, improving the accuracy in




Figure 5.12: Pulse Matching Efficiency versus Pulse Height.

peak location is preferred.

Figure 5.13 shows the average variation of pulse reconstruction accuracy (measured in percent deviation from actual value) with the size of the search window and pulse height. It can be seen that the average error in reconstructed values increases slightly with the size of the matching window. This reflects the decreased precision in peak location for larger windows. For small pulses, those less than 6 digitizer bins deep, the average reconstruction error exceeds the actual integral value. However, for pulses larger than 8 bins deep, a small search window provides pulse reconstruction which is accurate to 20%. Note that for pulses deeper than 12 bins, the reconstruction accuracy is near perfect.

The accuracy graphs are used to select final reconstruction parameters. The resulting accuracy curve is shown in Figure 5.14, which shows how the reconstructed pulse accuracy varies with pulse depth. The final analysis parameters allow pulses which are deeper than 12 digitizer bins to be identified and reconstructed with near perfect efficiency and accuracy.




Figure 5.13: Mean Reconstruction Error versus Window Size and Pulse Height.




Figure 5.14: Reconstructed Pulse Error versus Pulse Height.

5.3 Event Characterization

Though there are 486 nominal coincident Dual Cerenkov events, not all of these are suitable for use in the Dual Cerenkov technique (See Section 2.5). Recall that there is only a limited area on the surface which is sensitive to mass in one detector and energy in the other. In Figure 4.1, the active areas are the two outer lobes of the mass zones.

To characterize the events, certain data are examined. The initial requirement is that there be measurable Cerenkov pulses in both of the Cerenkov detectors. Because the array triggers off a single detector, a fair number of events contain only one Cerenkov pulse. Additionally, double pulses, when they exist, must occur within a few time bins of each other (corrected for cable delays). This ensures that most accidental coincidences get rejected. Events which pass these tests are called “good”, a label which simply indicates they contain a full set of Cerenkov data.

Unfortunately, the “goodness” criteria are not sufficient to identify candidates for the Dual Cerenkov technique. To test for this, the muon angles from Soudan 2 are used to reconstruct the shower trajectory. This trajectory is then projected back to the surface and the core location is noted. If the core falls within 120 m of one detector, but outside of 50 m (this region is excluded by limitations in the Monte Carlo), and is further than 120 m from the other detector, the event is labeled a “golden” event. Golden events are rated for their geometry and data completeness only; no judgement of the actual data quality is made.

A breakdown of the 486 nominal coincident events follows. After the event matching and characterization process, fewer than 10% of the original events were found to be suitable for the Dual Cerenkov Technique. Table 5.1 shows the data reduction efficiencies.

Nominal Coincidences	486	100.0%
Matched to S2 Events	466	95.9%
Good Events	100	20.6%
Golden Events	45	9.3%

Table 5.1: Data Reduction Summary

Chapter 6

Analysis

There is a great complexity in the mechanics of extensive air showers. The dependence of measured quantities on actual primary cosmic ray properties (*e.g.*, energy, charge, mass) is too complicated to characterize with simple analytical models. Thus, to extract meaningful results from raw detector data, one must examine that data in the context of a full Monte Carlo computer simulation.

In this chapter, the details of the data analysis procedure are provided. A description of the extensive air shower simulation and the subsequent interpretation of actual event data is offered. A discussion of the limited use of deep underground muon information is also included.

6.1 EAS Monte Carlo

The key feature to all indirect cosmic ray measurements is their dependence on some form of particle interaction simulation. In order to understand the significance of signals gathered in any variety of extensive air shower detector, a concept of the air shower behavior must be developed. This is typically achieved through computer simulation of all or some of the particle interactions comprising the shower.

The central challenge in creating these Monte Carlo programs is to find an appropriate balance between calculation accuracy and calculation speed. A typical EAS of 10^{15} eV may contain over 10^6 particles. To create a fully complete simulation, each of these particles must be

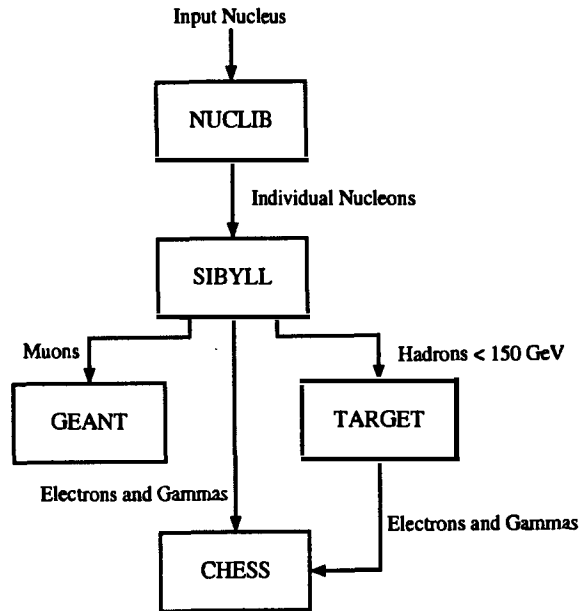


Figure 6.1: Monte Carlo Flow Chart. The figure shows how the Monte Carlo packages interact to simulate full extensive air showers.

tracked along its full trajectory, simulating every interaction it undergoes, and every daughter particle it generates. Clearly, this is a phenomenal computational task.

For the current investigation, the primary components of the extensive air shower which must be simulated are the hadronic, the electromagnetic and the muonic. A number of separate programs are used together to make the calculations. Their configuration is shown schematically in Figure 6.1. A brief description of various packages follows; for a more detailed explanation, see [42].

6.1.1 Hadronic Simulation

In this investigation, EAS hadronic interactions are simulated by a series of codes, depending on the energy range of the simulated hadron. Most of these simulations are common components used in some form throughout the astroparticle industry.

The full EAS simulation begins with an incident cosmic ray interacting with an atmospheric

nucleus. The initial interaction is handled by the NUCLIB [61] component of the HEMAS [62] package. The HEMAS package, which has been modified to include geomagnetic effects, controls the overall structure of the shower simulation.

After NUCLIB simulates the fragmentation of the incident cosmic ray, the high energy daughter particles are passed to the SIBYLL [63] code. This package simulates hadronic interactions down to particle energies of 200 GeV. Hadronic particles of energy less than 200 GeV are handled by the TARGET [64] routine. At energies below 5 GeV, daughter hadrons are no longer tracked. At these energies, they preferentially decay before interacting, and their further contribution to either the Cerenkov light field or deep underground muon flux is essentially zero. The treatment of muons and charged leptons generated in this stage is discussed below. Gammas produced in these steps are passed into the electromagnetic simulator, CHESS. Neutrinos are registered only as an energy loss.

It should be noted that SIBYLL, along with all high energy hadronic interaction Monte Carlos, relies on extended extrapolations of characteristics measured in low energy fixed-target acceleration experiments. The accuracy of these theoretical extensions has only recently been subject to empirical verification. These results, which have a bearing on the reliability of certain Monte Carlo based conclusions, are addressed in [45].

6.1.2 Muonic Simulation

Muons generated in the hadronic simulations are propagated through the atmosphere to the surface using the GEANT [65] package. In the current investigation, little use is made of the simulated muon component of the air shower (See Section 6.3). For detailed deep underground muon studies, an additional step is introduced to propagate the muons, again using GEANT, through the rock overburden to the Soudan 2 detector itself. Figures 6.2 and 6.3 show results on muon transmission probability and average muon yield from just such studies.

An excellent and very detailed discussion of muon propagation simulations and their application to cosmic ray investigations can be found in [15].




Figure 6.2: Transmission probability versus energy for muons through a 700 m overburden of standard rock. From [45].

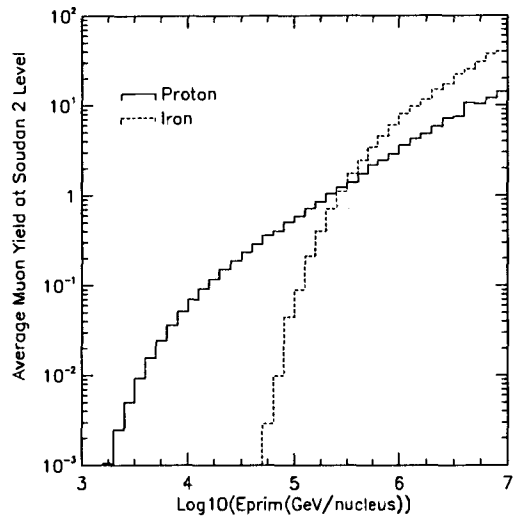


Figure 6.3: Average muon yield at Soudan 2 depth versus primary energy for protons and iron primaries. From [15].

6.1.3 Electromagnetic Simulation

Electromagnetic and Cerenkov EAS components are simulated using CHESH, the Cherenkov and Electromagnetic Shower Simulator [42]. This code is a fully three-dimensional electromagnetic cascade Monte Carlo targeted to the energy range of 10^{14} – 10^{16} eV. CHESH, which was developed at the University of Minnesota, provides very rapid and accurate simulations of electromagnetic shower development and Cerenkov lateral distributions.

CHESH uses a high degree of specialization and careful programming techniques to achieve a level of speed some 60 times faster than competing codes. There is also a parameterized extension of CHESH which simplifies shower mechanics by simulating only two particle populations (the direct and indirect populations discussed in Section 2.4.4). The parameterized CHESH is a full 60 times faster than the full CHESH simulation. This speed comes at the cost of reduced applicability; Cerenkov yields at points within 50 m of the shower core cannot be accurately simulated. The majority of simulations used in this investigation have been performed using the parameterized version of CHESH.

The decision to use the in-house Monte Carlo CHESH, versus the more established CORSIKA [66] code was motivated by a number of issues. First, CORSIKA should be slower than CHESH for all but the most extreme of thinning configurations. While using the thinning feature can greatly speed up the simulations, the techniques employed by CORSIKA are performed at the collection level, rather than the emission level. Thus, heavy thinning likely has a measurable effect of the simulated light yields. CHESH uses a more natural form of thinning at the photoelectron level which retains fluctuations inherent in the shower development.

Secondly, using CHESH allows the bulk of the previous Monte Carlo simulations run at Minnesota to be used in the current investigation as well. This includes several CPU-years worth of acceptance simulations calculated for the single Cerenkov station, but usable by the dual Cerenkov array. CHESH's ease of use and larger user base at Minnesota also contributed to the decision to use it.

6.2 Data Interpretation

The final interpretation of the event data is performed by generating multiple extensive air showers, simulating the response of the detectors to those showers, and comparing the measured data to the simulations. Typically, this would require simulating virtually the entire phase space of allowable cosmic ray properties. This includes a broad span of zenith and azimuth angles, multiple primary masses, and numerous energies.

For a moderately detailed simulation, the full volume in phase space contains over 200,000 parameter configurations. Thus, due to limited computing resources, only a few repetitions (< 10) of each configuration can be run. This makes conclusions based on these calculations susceptible to statistical fluctuations. Even with the reduced number of repetitions, such a simulation would require up to six months of calculation time on a cluster of ten Alpha workstations.

Note that such a simulation, though it covers the full *span* of the phase space, necessarily simulates only a coarse sampling of points *within* it. Most parameters are continuous throughout their allowable range, while the Monte Carlo can simulate only a limited number of discrete points in this range. Thus, for phase space points between the calculated points, interpolations of simulated values must be made, which may introduce uncertainty into the results.

The current investigation is a low-rate experiment, which, by the geometry of the detector arrangement, samples only a small portion of the full phase space. Thus, it has been possible to simulate only those points in phase space which directly correspond to actual data gathered by the detector. That is, showers are simulated on an event-by-event basis.

By simulating only the volumes of phase space corresponding to angles for which events have been gathered, a vast reduction in calculation time can be achieved. Furthermore, an improvement in accuracy is achieved by removing the need to interpolate results. Fixing the simulated angles to those seen in actual events measured in Soudan 2 leaves only energy and mass as free parameters. The great reduction in phase space volume allows many more repetitions of each configuration point to be simulated.

In practice, the angles associated with all the “golden” events (see Section 5.3), are collected and noted. For each event, seven energies and three masses are simulated 50 times to form the test sample.




Figure 6.4: Energy vs. Pulse Integral for Three Primary Masses.

6.2.1 Energy Determination

The energy of cosmic ray events is determined by examining the size of the Cerenkov pulse measured in the “Energy Zone” detector. As detailed in Section 2.4.4, the pulse height in the furthest portions of the Cerenkov light pool is proportional primarily to the energy of the primary cosmic ray.

Figure 6.4 shows the variation of Cerenkov pulse integral (in the “Energy Zone” detector) with the energy of the primary particle for a particular event geometry. The data points represent the average pulse height¹ seen in 50 repetitions of the simulated showers. The error bars represent the RMS deviation of the pulse heights. The pulse integrals for three primary masses are shown; from bottom to top, they are Hydrogen, Magnesium, and Iron. This figure represents data collected from 1050 simulated air showers.

Reconstructed event energies are generated from this data simply by matching the measured

¹The term “pulse height” is used here, perhaps sloppily, to refer simply to the integrated signal seen in the Cerenkov detector

pulse integral to the simulated energy on the graph. Clearly, the differing mass curves imply that selecting the proper energy requires some *a priori* knowledge of the mass. To overcome this, an iterative process is used. In the first iteration, a mass of 1 is assumed. This will be refined once X_{max} information is obtained in the process described below (Section 6.2.3).

There are a few features of interest in this figure. The first is the larger size of the error bars in the Hydrogen curve (lowest), versus the Iron curve (topmost). This is presumably a manifestation of Iron showers' propensity towards smaller fluctuations in their development. This is typically discussed in the context of variations of X_{max} , but the behavior clearly extends to the production of Cerenkov light as well.

Also of interest, though likely not detectable by eye, is the nature of the scaling of the pulse integrals with energy. It was stated in Section 2.4.4 that a consequence of the toy model of electromagnetic cascades was that pulse height in the "indirect" or "Energy" zone of the Cerenkov pool scaled linearly with the energy of the particle. However, the best fit function to the Monte Carlo data is actually a *power law* in energy,

$$\rho(E) = CE^\alpha \quad (6.1)$$

with an index $0.85 \lesssim \alpha \lesssim 0.95$. Though small, this deviation from the idealized behavior would show up as a systematic energy shift if not properly identified and corrected for.

6.2.2 X_{max} Determination

The depth in the atmosphere of maximum shower development, X_{max} , is a sensitive indicator of primary mass. Likewise, the slope of the Cerenkov lateral distribution in the "direct", or " X_{max} " zone (see Section 2.4.4) has been shown to provide an estimate of X_{max} . However, the pulse height at any fixed radius from the core is a function of both the distribution slope as well as the overall normalization. This normalization is simply the global intensity of the radiation field, which is a function of the shower energy. Thus, light pool structure varies with X_{max} , but scales with energy. Therefore, a single-point measurement is not adequate to experimentally determine the slope in the interior region. This would require *a priori* knowledge of the normalization. In other words, to fix the slope of a line, two measurements are required.

Fortunately, the pulse height in the outer region of the lateral distribution provides a measure of the overall normalization. The pulse height in this region is a stable indicator of shower energy.




Figure 6.5: X_{max} vs. Pulse Ratio.

Therefore, X_{max} can be extracted from the event data in a fashion similar to that used for calculating the energy. Instead of directly examining the interior pulse height, the ratio of pulse heights (Interior/Exterior) is used. Figure 6.5 shows how X_{max} varies with the pulse ratio for a series of simulated showers. There are 50 simulated showers for 21 configurations of energy and mass (seven energies and three masses). The data points represent the average values of X_{max} and pulse ratio for those 50 showers, while the error bars represent the RMS deviations in each quantity.

There is a clear correlation between the pulse ratio and the associated X_{max} of the shower. This feature is exploited to extract X_{max} from the event data by directly comparing measured ratios to the simulated results.

The huge size of the errors bars on the simulated points illustrates very well a fundamental limitation of direct cosmic ray investigations: the stochastic nature of air showers. The error bars simply represent the large range of outcome conditions which can arise from identical input conditions. This means that even a perfect event-by-event shower analysis would be incapable of providing useful information on the cosmic ray flux in the absence of a large statistical sample.

The relative importance of individual showers is simply too small given the size of the natural fluctuations in their development.

6.2.3 Iteration of Results

As seen in Figure 6.4 and discussed in Section 6.2.1, there is a clear mass dependence in the “Energy Zone” pulse height from showers of the same energies. This requires that, for an accurate calculation of shower energy, some knowledge of primary mass must be in hand. However, an estimate of the primary mass requires information not only on X_{\max} , but on energy as well. This relationship defines the elongation rate equation:

$$X_{\max} = C_0 + C_1 \log_{10} \left(\frac{E}{A} \right) \quad (6.2)$$

In such a situation, iterative techniques are helpful. Beginning with an initial guess of the mass, an energy is calculated. From this energy and the calculated X_{\max} , an estimate of mass is produced. This mass is then used to get an improved energy estimate, which ultimately leads again to an improved mass estimate. After a few such iterations, the energy has typically converged to better than 1 TeV. The final correction to the energy is typically on the order of 20%.

6.3 Muon Information

The initial design goal of the Dual Cerenkov array at Soudan was to make a unique multi-parameter air shower measurement combining atmospheric Cerenkov radiation and deep underground muon fluxes. If successful, such a measurement would yield not only an accurate estimate of the primary cosmic ray energy, but multiple independent estimates of the mass as well.

Unfortunately, recent results [45] from the single station Cerenkov experiment at Soudan call the resolving power of the underground muon data into question. The problem resides in the Monte Carlo software used to simulate the production of high energy muons in the extensive air showers.

After a lengthy and thorough investigation, it was found that none of the popular high-energy hadronic interaction codes generate transverse momentum distributions which fit the

data observed in nature. These distributions, which are used to determine the final angle of daughter particles following hadronic collisions and decays, determine the properties of the deep underground lateral muon distributions. These underground muon lateral distributions were measured to be much wider than the Monte Carlo predict. Conclusions based on the simulated showers would imply cosmic ray compositions which are substantially heavier than iron.

The root of the discrepancy likely lies in the fact that modern Monte Carlo must rely on greatly extended extrapolations of directly-measured transverse momentum distributions. While the characteristic angle of dispersion for the high energy muons in a shower is less than 0.1° , corresponding to pseudorapidities near 8 (pseudorapidity, $\eta \equiv -\ln \tan \frac{\theta}{2}$), data from accelerator investigations is exhausted at $\eta \sim 3$.

Figure 6.6 shows the distributions from several hadronic interaction Monte Carlo for $\bar{p}p$ interactions at 1.8 TeV. This energy in the center of mass frame corresponds to an energy of 1.73 PeV in the lab frame ($E_{lab} = \frac{E_{cm}^2 - 2m^2}{2m}$), roughly the location of the spectral knee. Clearly, there is a large gap between the last experimental data points and the pseudorapidities relevant to this investigation. It should also be noted that, while there is some spread in the modeled distributions between $\eta = 3$ and $\eta = 6.5$, past $\eta = 7$, the models more or less converge to similar values. The end result of this is that the models tend to generate muon distributions which are more alike than different. This is shown in Figure 6.7 which displays the lateral distribution of muons at 100 m above sea level in 10^{14} and 10^{15} eV p-induced showers. The maximum variation in muon densities is roughly a factor of two up to 2 km from the shower core, with SIBYLL yielding the lowest values. These distributions include all generated muons, not merely high-energy muons.

Because of these deficiencies, and because there exist no Monte Carlo which offer alternative results, a sobering conclusion must be drawn. Namely that, while muon data is essential for locating the core of the extensive air shower in this Cerenkov study, as an independent estimator of primary mass, it is not useful. Until accelerator data can extend well beyond its current limits, Monte Carlo hadronic interaction models may be of limited value in cosmic ray composition studies.

It should be noted that electromagnetic shower properties, such as X_{\max} and Cerenkov light yield, have been shown to be, for the most part, insensitive to differences in the hadron interaction code. Table 6.1 shows X_{\max} calculations from five different hadronic interaction




Figure 6.6: Pseudorapidity distributions for $\bar{p}p$ at 1.8 TeV from several hadronic interaction models, as calculated by [67]. Data points are from [68].




Figure 6.7: Lateral Muon distributions for 10^{14} and 10^{15} eV proton-induced showers as simulated by several codes [67].

Model	Hydrogen			Iron		
	X_{\max}		E. Rate	X_{\max}		E. Rate
	10^{14} eV	10^{15} eV		10^{14} eV	10^{15} eV	
VENUS	503	574	71	355	439	84
QGSJET	504	576	72	355	442	87
SIBYLL	519	592	73	362	458	96
HDPM	521	599	78	352	444	91
DPMJET	492	560	68	353	428	75

Table 6.1: Average X_{\max} calculated by several interaction models for iron primaries at 100 TeV and 1000 TeV. Units are all g/cm^2 , from [67].

Model	Hydrogen		Iron	
	100 TeV	1000 TeV	100 TeV	1000 TeV
HDPM	79.4 ± 83.8	66.0 ± 66.7	12.7 ± 12.5	13.0 ± 12.0
VENUS	77.3 ± 78.7	71.1 ± 72.2	12.6 ± 13.6	12.6 ± 11.9
SIBYLL	75.9 ± 77.2	69.6 ± 68.6	12.5 ± 13.0	12.5 ± 11.7
QGSJET	80.3 ± 88.6	73.6 ± 69.4	14.0 ± 13.3	14.0 ± 12.7
DPMJET	75.2 ± 75.4	77.1 ± 82.5	12.8 ± 12.8	12.6 ± 12.5

Table 6.2: Average X_0 , or point of first interaction calculated by several interaction models for hydrogen and iron primaries at 100 TeV and 1000 TeV. Units are all g/cm^2 , from [67].

Monte Carlos. The results, which are for hydrogen and iron-induced showers at 100 and 1000 TeV display remarkable similarity.

Likewise, Table 6.2 shows the values of X_0 , the point of first interaction, as calculated by the same five hadronic interaction Monte Carlos. This value has a strong effect on how the subsequent shower develops. Clearly, there is little variation between these models.

Chapter 7

Results and Conclusions

The final analysis of the Dual Cerenkov data was performed by applying the selection routines of Section 5.3, and processing the remaining events with the procedures of Chapter 6. The results of these steps are discussed in this chapter, along with a general discussion of some issues impacting the detector. Finally, a summary of the achieved tasks and goals of this investigation are presented.

7.1 Results

The output of the analysis routine is pairs of reconstructed X_{\max} and energy values. These data were subjected to some modest data-quality cuts before being binned and plotted. Figure 7.1 shows the final results of the analysis. The points with error bars signify the average measured X_{\max} found in a particular energy bin. The hollow crosses represent direct measurements of the cosmic ray flux, gathered from [48]. The diagonal lines mark the simulated elongation rates for Iron and Hydrogen primaries.

The measured results are entirely consistent with a flux between Hydrogen and Iron, an important initial quality check. Because these elements bracket the range of normal stellar nucleosynthesis products, they are thought to comprise the bulk of the cosmic ray flux. Only in the highest energy bin is the average X_{\max} further than one standard deviation from this range (it is less than two, however).




Figure 7.1: Average X_{\max} vs. Particle Energy. The hollow crosses represent direct measurements of the cosmic ray flux, gathered from [48]. The diagonal lines mark the simulated elongation rates for Iron and Hydrogen primaries.

Because of the limited data available for this analysis, any conclusions beyond simple claims of consistency are unwarranted. The lowest energy bins in the X_{\max} vs. energy graph contain only a handful of events each, while the higher energy bins contain only one or two events. The statistical fluctuations inherent in air shower development simply require that more data be collected to make meaningful conclusions in an experiment of this sort. A piece of evidence which might possibly have offset the low data rate is the underground muon information. Unfortunately, this data has been shown to be less helpful than hoped, due to the deficiencies identified in the Monte Carlos.

These factors combined have forced this analysis to provide what should be considered preliminary results. Simply, additional data must be collected to realize the potential of this experiment. With more data will come reduced uncertainties and an increased energy threshold. With improvements to the hadronic Monte Carlos, a refined analysis could potentially provide a unique insight into the composition problem. Nevertheless, to the present time, the principal product of this investigation has been a sophisticated and operational detector, a viable analysis technique, and the potential to offer an innovative analysis of a long-standing problem.

7.2 Comments

Through the course of designing, building and operating this experiment, a number of limitations and possible enhancements have been identified. Though some have been mentioned previously, they are again detailed here for completeness.

7.2.1 Limitations

Muon Data

Clearly, the greatest disappointment of the present experiment has been the inability to make further use of the deep underground muon information. While the muon data has the potential to make this investigation unique, the deficiencies of the hadronic interaction Monte Carlo effectively remove this possibility. However, it should be noted that this detector is currently the only detector in the world which has the possibility of addressing these very issues. Thus, while the goal of a unique measure of cosmic ray composition may be currently untenable, a very different, but equally important, goal of improving the current state of hadronic interaction Monte Carlo may be possible.

The possibility of using the present and future data to help inform the current hadronic interaction Monte Carlo should be considered seriously.

Core Location

The efficacy of the dual Cerenkov technique depends closely on the ability to accurately locate the core of the extensive air shower being measured. Unlike air shower arrays, which make high multiplicity electromagnetic cascade measurements, shower cores at Soudan are located by tracing often only a single muon track back to the surface.

This represents a source of uncertainty in the lateral distribution calculations. The issue is essentially one of scale: the typical lateral extent of the deep underground muon fields is roughly the same size as the Soudan 2 detector. Thus, for any given shower, there is no way to know if the full muon core has been contained. If it has, the centroid of the muon hits likely forms a good estimate of the original shower trajectory. If it has not, then the estimate will be less good. An excellent discussion of how limited sampling area affects observed muon multiplicities

can be found in [45].

Detector Quantity

It should be clear that the number of atmospheric Cerenkov detectors used in current experiment is the bare minimum required by the technique. This was a practical necessity borne of limited manpower and resources. However, it effectively limits both the acceptance and resolving power of such an experiment. Traditional Cerenkov arrays typically employ over 100 stations. Increasing the number of stations increases the effective area of the experiment and also provides multi-point measurements of the lateral distributions of each shower. This acts to reduce systematic errors in the analysis.

Detector Location

One of the factors limiting the data rate in this experiment is the placement of the Cerenkov detector stations. Figure 4.1 shows the location of the detector stations, along with their associated mass zone boundaries. Recall that the Dual Cerenkov technique requires that a shower core fall within the mass zone of one detector and outside the mass zone of the other detector. Clearly, this can only hold true in the upper and lower lobes of the respective mass zones. The overlapping region between the two detectors, which constitutes a fairly large fraction of the total mass zone area, is useless for the technique.

To fully exploit the minimal two station arrangement, the detectors should be located more than twice the nominal mass zone radius (120 m) from each other. This removes any mass zone overlap and maximizes the total active detector area. Any future additions to the array should be located with this fundamental constraint in mind.

7.2.2 Enhancements

Event-by-Event Calibration

A complete calibration system using blue LED pulsers was designed and prototyped for this experiment. A pulser was to rest inside the Winston cone of each tube in the array, providing a calibration reading for each event gathered. However, limited resources prevented its implementation.

There is a clear benefit to having a reliable tube-by-tube calibration for every data event. Provided that long-term stability and piece-to-piece light production parity can be maintained, this would represent a very important upgrade to the detector.

Additional Data

Though strictly not a detector enhancement, the collection of more data will be an important step in improving the quality of the detector. While systematic errors have been identified and accounted for wherever possible, certain systematic effects will only emerge in the presence of more data.

As an example, events whose shower cores are equidistant from both detectors provide an excellent opportunity to test the relative gains of the two detectors. This has become a standard technique in larger arrays of detectors. However, because these events are quite rare in the Dual Cerenkov configuration, there has been little opportunity to make use of it. Note that a systematic shift in the station gains will directly affect calculations of X_{\max} , which relies on the ratio of pulse heights seen in these detectors.

Filtering

The use of an ultraviolet-pass filtering system on the photomultiplier tubes would theoretically have a two-fold benefit. The first would be to reduce the contribution of night-sky background (which rises with longer wavelengths) to the DC anode current. This would allow an increase in operating gain, which could help retain more lower energy events. In a single-station configuration, this would also allow more high energy events located further from the detector to be collected. This is of limited usefulness in the dual station configuration, however. In this case, the geometric requirements for "golden" events (*i.e.*, the shower core must fall within the appropriate lateral distribution zones) rule out any increases in collection area gained by lowering the energy threshold.

A more significant benefit of UV-filters would be the ability to operate in the presence of significant moonlight. Moonlight currently limits operations to, at best, two weeks of every month. Even if filtering could only push this to three weeks per month, the overall gain in exposure time would be significant.

7.3 Conclusions

The goal of the present investigation has been to make a new determination of the cosmic ray composition in the region of the spectral knee by measuring multiple components of extensive air showers. The simultaneous detection of stereo atmospheric Cerenkov radiation and deep underground muon fluxes would yield a unique approach to this long-standing problem.

Towards this goal, the following objectives have been addressed:

- A new dual atmospheric Cerenkov array has been designed, constructed, and operated at the Soudan 2 laboratory. The new detector has proven to be robust and reliable through two seasons of data gathering.
- A method of event-by-event shower analysis has been developed. This method can help reduce systematic errors by reducing the volume of the phase space which must be simulated for full detector Monte Carlos. This allows for an increase in the number of shower repetitions per configuration. Use of iterative techniques in the calculation of energy and X_{\max} help reduce known sources of systematic error.
- An analysis of the preliminary detector data has been performed. This analysis indicates that the array operates as expected, yielding results consistent with physical bounds as well as direct measurements of the cosmic ray flux. This indicates that the detector is well poised to make meaningful composition measurements in the near future.
- A full analysis of the event data, which would utilize the deep underground muon fluxes in addition to the Cerenkov lateral distributions has not been performed. Limitations in current hadronic interaction Monte Carlos make this information less useful at this time. However, should these deficiencies be addressed and corrected in the future, the data can be reanalyzed
- The possibility for providing unique insight into the recently discovered hadronic interaction Monte Carlo deficiencies exists. This detector has the means to provide high pseudo-rapidity data which is currently beyond the capabilities of accelerator-based experiments.

The Dual Cerenkov array at Soudan is currently fully operational. The software, hardware and analysis techniques have all been developed; preliminary results indicate that the instrument

works as designed. Though poor weather has limited the quantity of results to date, the detector stands ready to significantly add to the total amount of collected data. A single favorable season of weather could easily quadruple the current data set, extending the upper energy threshold and improving the statistics of the full set. While this detector can never, by its very nature, compete with vastly larger, fully-manned, Cerenkov experiments such as BLANCA [50], it does demonstrate that real physics can emerge from modest endeavors.

Appendix A

Inter-Integrated Circuit Protocol


The Inter-Integrated Circuit Protocol (I^2C) is a two-wire serial protocol used for communication between specialized integrated circuit devices. Developed by Philips in the early 80's to help maximize circuit simplicity and efficiency, I^2C is now usable with over 1000 compatible devices, including ADC's, DAC's, calendar/clock units, and more.

Using I^2C technology on the Lower Station hardware allowed for enhanced functionality, including modularity and robustness benefits, as well as rapid development cycle time. Using specialized multi-function devices with integrated serial capabilities enabled the control and high voltage hardware to be designed with the minimum number of components, making these units very stable and virtually immune to damage.

A.1 Hardware

There are several categories of device which can participate in the I^2C bus. These include:

1. A device that transmits signals onto the I^2C -bus is the transmitter
2. A device that receives signals is the receiver
3. A device that controls signal transfers on the line in addition to controlling the clock frequency is the master
4. A device that is controlled by the master is the slave.

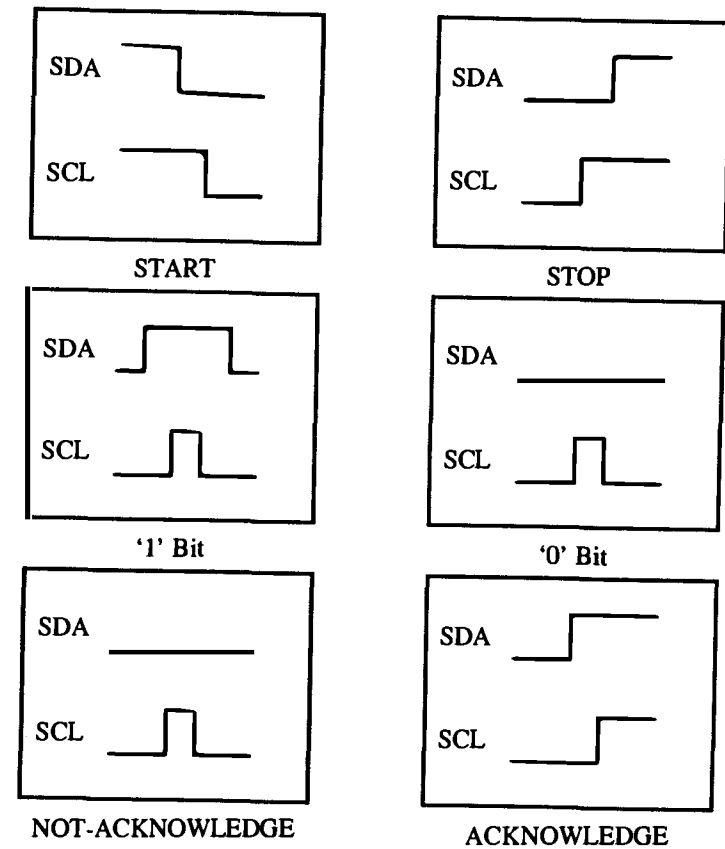
Figure A.1: Typical I²C System Schematic.



The master can transmit or receive signals to or from a slave, respectively, or it can control signal transfers between two slaves, where one is the transmitter and the other is the receiver. It is possible to combine several masters, in addition to several slaves, onto an I²C -bus to form a multimaster system. If more than one master simultaneously tries to control the line, an arbitration procedure decides which master gets priority. The maximum number of devices connected to the bus is dictated by the maximum allowable capacitance on the lines, 400 pF, and the protocol's addressing limit of 16k; typical device capacitance is 10 pF.

Physically, the I²C bus consists of one ground line and two bidirectional signal lines, the Serial Data (SDA) and Serial Clock (SCL) lines. A typical system will consist of at least one microcontroller and one or more peripheral devices connected to the SDA and SCL lines. These data lines themselves are typically connected via pull-up resistor to the digital device supply voltage, V_{DD} (Figure A.1). The logic levels on the signal lines are not fixed, and depend on the associated V_{DD} level.

Each I²C device has four fixed and three programmable address bits (which are set by holding device pins high or low). Thus, only eight identical devices can be connected to the same bus. However, I/O devices can be used to control the levels on the addressing lines, allowing opening additional address space. In the case of the high voltage modules for the Cerenkov array, 12 identical DAC converters were required. Each module was therefore equipped with a Philips PCF8574 IC which toggled the DAC address lines, allowing up to 16 identical units on the bus. The final, eighth, bit of the address byte is a command bit setting the device into either a read or write mode.

The I²C command set comprises six bus-state primitives which are combined to create all command and communication procedures. These are shown in Figure A.2. Though the details

Figure A.2: I²C Command Primitives.

Figure A.3: I²C Read Protocol.Figure A.4: I²C Write Protocol.

of the command procedures vary with each device, they typically proceed with the master issuing an initial START condition, followed by an ADDRESS byte. In the case that the eighth address bit was a WRITE bit, the addressed slave transmits an ACKNOWLEDGE condition, and then DATA bytes are issued one at a time by the master. An ACKNOWLEDGE condition is sent from the slave after each data byte. The absence of an ACKNOWLEDGE condition from the slave, called a NOT-ACKNOWLEDGE, follows the final data bit. A STOP condition from the master completes the transaction.

In the case of a READ command during the address byte, the slave will immediately transmit data bytes following its address acknowledgement. The master transmits ACKNOWLEDGE conditions following every data byte until the last, when it sends a NOT-ACKNOWLEDGE and finally, a STOP. See Figures A.3 and A.4

A.2 Software

The I²C master controller used in the Lower Station hardware bin is the V25 microprocessor. Because this microprocessor does not have any integrated I²C capabilities, the required hardware I/O routines had to be programmed into the command software by hand. Lacking bit-level I/O

access for this architecture, all routines had to be performed at the byte level.

The lowest level of software control is supplied by the two functions `sda()` and `scl()`. Each takes an integer as an argument, and sets the associated pin low on 0 and high for 1. The SDA and SCL lines of the primary I²C bus were assigned to pins 2 and 1 of the second V25 port.

```

void sda(int state){
    unsigned char prev,newer;
    prev = peekb(0xfff0,0x10);
    if (state == 0){
        newer = prev & SDAMASK;    //    & sets lo
    } else {
        newer = prev | ~SDAMASK;    //    |~ sets hi
    }
    pokeb(0xfff0,0x10,newer);
}

```

```

void scl(int state){
    unsigned char prev,newer;
    prev = peekb(0xfff0,0x10);
    if (state == 0){
        newer = prev & SCLMASK;    //    & sets lo
    } else {
        newer = prev | ~SCLMASK;    //    |~ sets hi
    }
    pokeb(0xfff0,0x10,newer);
}

```

Using these two basic functions, the entire range of the I²C command primitive library (Figure A.2) can be generated.

```

void send_start(void){
    scl(1);
    sda(1);
}

```

```

    sda(0);
    scl(0);
    sda(1);
}

void send_stop(void){
    sda(0);
    scl(1);
    sda(1);
}

void send_0(void){
    sda(0);
    scl(1);
    scl(0);
    sda(1);
}

void send_1(void){
    sda(1);
    scl(1);
    scl(0);
}

void send_ack(void){
    sda(0);
    scl(1);
    scl(0);
    sda(0);
}

```

More advanced activities, such as sending and receiving bytes can be achieved using these

functions. However, to further abstract from the low level details as well as significantly speed execution, these functions have been pre-constructed as well.

```

unsigned char get_byte(void){
    unsigned char dat,tmp;
    int i,j;

    pokeb(0xfff0,0x11,0x12); //using pokes shrinks ack time
                             //from 15 ms to .3 ms

    for (i=0;i<8;i++){
        scl(1); //set scl high
        tmp = port_rd(2); //get byte from port
        j=(tmp>>1)&1; //get bit 2 from byte
        dat=(dat<<1)|j; //add bit to output byte
        scl(0); //set scl low
    }
    pokeb(0xfff0,0x11,0x10); //using pokes shrinks ack time
                             //from 15 ms to .3 ms

    sda(1);
    return dat;
}

int get_ack(void){
    unsigned char dat;
    int j;

    sda(1);

    //remember to leave P24 as input
    // too. (0x12)

    pokeb(0xfff0,0x11,0x12); //using pokes shrinks ack time
                             //from 15 ms to .3 ms
}

```

```

scl(1);                //set scl high

j=(port_rd(2)>>1)&1;

scl(0);                //set scl low

pokeb(0xffff,0x11,0x10); //using pokes shrinks ack time
                        //from 15 ms to .3 ms

return j;
}

int send_byte(unsigned char dat){
    int i,j,test = 1;
    int count = 0;

    while(test){
        for (i=0;i<8;i++){
            j=(dat>>7)&1;
            dat=dat<<1;
            if (j==1){send_1();}
            if (j==0){send_0();}
        }

        test = get_ack();
        count++;
        if (count > 4){
            printf("5 Ack Failures!\n");
            ACK_FLAG = 0;
            test = 0;
            return 0;
        }
    }
}

```

```

    }
}
return 1;
}

```

With this simple series of routines in place, any of the devices on the I²C bus can be commanded and controlled with a simple set of function calls. Setting the voltage on a tube, for instance, is as simple as issuing the following commands to the appropriate PCF8591 (DAC) device:

```

send_start();
send_byte(<DAC Address>);
send_byte(0x44);           //command to set DAC
send_byte(<DAC Level>);
send_stop();

```

References

- [1] T.K. Gaisser. *Cosmic Rays and Particle Physics*. Cambridge University Press, 1990.
- [2] V.F. Hess. *Physics Z.*, 13:1084, 1912.
- [3] N. Hayashida *et al.* Observation of a very energetic cosmic ray well beyond the predicted 2.7-k cutoff in the primary energy spectrum. *Phys. Rev. Lett.*, 73:3491–3494, 1994.
- [4] C.J. Waddington. University of MN Phys 5401 Class Notes, 1995.
- [5] J.A. Simpson *et al.* *Annual Rev. Nucl. Part. Sci.*, 33:323, 1983.
- [6] C. Caso *et al.* Review of particle physics. *Eur. Phys. J.*, C3:1, 1998.
- [7] N.L. Grigorov *et al.* *Proceedings of the 12th ICRC (Hobart)*, 2:206, 1971.
- [8] K. Asakimori *et al.* *Proceedings of the 23rd ICRC (Calgary)*, 2:25, 1993.
- [9] T.V. Danilova *et al.* *Proceedings of the 15th ICRC (Plovdiv)*, 8:129, 1977.
- [10] Yu. A. Fomin *et al.* *Proceedings of the 22nd ICRC (Dublin)*, 2:85, 1991.
- [11] M. Nagano *et al.* *J. Phys.*, G10:1295, 1984.
- [12] D.J. Bird *et al.* *Astrophys. J.*, 441:144, 1995.
- [13] S. Yoshida *et al.* *Astropart. Phys.*, 3:105, 1995.
- [14] M.A. Lawrence *et al.* *J. Phys.*, G17:773, 1991.
- [15] S.M. Kasahara. *A Study of Cosmic Ray Composition in the Knee Region using Multiple Muon Events in the Soudan 2 Detector*. PhD thesis, University of Minnesota, 1997.
- [16] M.S. Longair. *High Energy Astrophysics Second Edition*. Cambridge University Press, 1992.
- [17] S. Swordy. *Proceedings of the 24th ICRC (Rome)*, 2:697, 1995.
- [18] K. Asakimori *et al.* *Proceedings of the 24th ICRC (Rome)*, 2:707, 1995.
- [19] G.B. Kristiansen *et al.* *Proceedings of the 24th ICRC (Rome)*, 2:772, 1995.
- [20] J. Kempa *et al.* *Proceedings of the 24th ICRC (Rome)*, 2:681, 1995.
- [21] Boothby. Dice. *Proceedings of the 25th ICRC (Durban)*, 1997.
- [22] A. Ahlen *et al.* *Proceedings of the 24th ICRC (Rome)*, 2:710, 1995.
- [23] A.A. Watson. Charged cosmic rays above 1 tev. *Proceedings of the 25th ICRC (Durban)*, 8:257, 1997.
- [24] E. Fermi. *Physical Review*, 75:1169, 1949.
- [25] Pierre Sokolsky. *Introduction to Ultrahigh Energy Cosmic Ray Physics*. Addison-Wesley Publishing Company, Inc., 1989.
- [26] Rene A. Ong. Very high energy gamma-ray astronomy. *Physics Reports*, 305:93, 1998.
- [27] M. Teshima. *23rd ICRC Inv., Rap., and Highlight papers (Calgary)*, page 257, 1993.
- [28] E.G. Berezhko. Cosmic ray sources. *Proceedings of the 25th ICRC (Durban)*, 8:281, 1997.
- [29] A.R. Bell. *Proceedings of the 22nd ICRC (Dublin)*, 2:420, 1991.
- [30] V.S. Ptuskin. *Proceedings of the 24th ICRC (Rome)*, 8:755, 1995.
- [31] R. Cowsik *et al.* *Can. J. of Phys.*, 46:S646, 1968.
- [32] M.M. Shapiro and R. Silberberg. *Ann. Rev. Nucl. Sci.*, 20:323, 1970.
- [33] R. Cowsik and L.W. Wilson. *Proceedings of the 13th ICRC (Denver)*, 1:500, 1973.
- [34] J. H. Adams *et al.*. Proceedings of the first ldef post-retrieval symposium. (nasa cp-3134), NASA, 1991.
- [35] W. R. Binns *et al.* The uh nuclei cosmic ray detector on the third high-energy astronomy observatory. *Nucl. Instrum. Meth.*, 185:415, 1981.
- [36] E. C. Stone *et al.*. The advanced composition explorer. *Space Science Reviews*, 86:1, 1998.
- [37] S. P. Swordy, J. L'Heureux, D. Muller, and P. Meyer. Measurements of x-ray transition radiation from plastic fibers. *Nucl. Instrum. Meth.*, 193:591, 1982.
- [38] T. H. Burnett *et al.*. Photon and helium energy spectra above 1-tev for primary cosmic rays. *Phys. Rev. Lett.*, 51:1010, 1983.
- [39] S. W. Barwick *et al.*. The high-energy antimatter telescope (heat): an instrument for the study of cosmic-ray positrons. *Nucl. Inst. Meth*, 400:34, 1997.
- [40] G. Barbiellini *et al.*. The 3he/4he ratio with the caprice apparatus. *Proc. of the 25th ICRC (Durban)*, OG 5.2.4, 1997.
- [41] D. Mueller *et al.*. Towards the knee: Direct measurement of the mosmic ray composition with electronic detectors. *Adv. Space Res.*, 19:719, 1997.
- [42] V.V. Vassiliev. *A Study of Cherenkov Radiation from Atmospheric Cascades*. PhD thesis, University of Minnesota, 1997.
- [43] U. Das Gupta *et al.* *Phys. Rev. D*, 45:1459, 1992.
- [44] N.P. Longley. *Ultra High Energy Cosmic Ray Composition*. PhD thesis, University of Minnesota, 1993.

- [45] L.M. Mualem. *A Measurement of Cosmic Ray Composition using Cherenkov Radiation and Deep Underground Muons*. PhD thesis, University of Minnesota, 1998.
- [46] S.P. Swordy and D.B. Kieda. *Proceedings of the 26th ICRC (Salt Lake City)*, 3:144, 1999.
- [47] J.R. Horandel *et al.* Cosmic-ray mass composition in the peV region estimated from the hadronic component of eas. *Proc. 16th European Cosmic Ray Symposium*, HE 8.3, 1998.
- [48] K.H. Kampert *et al.* *Proceedings of the 26th ICRC (Salt Lake City)*, 3:159, 1999.
- [49] J.E. Dickinson *et al.* *Proceedings of the 26th ICRC (Salt Lake City)*, 3:136, 1999.
- [50] L.F. Fortson *et al.* *Proceedings of the 26th ICRC (Salt Lake City)*, 3:125, 1999.
- [51] A. Rohring *et al.* *Proceedings of the 26th ICRC (Salt Lake City)*, 3:152, 1999.
- [52] W.W.M. Allison *et al.* The SOUDAN2 detector: The design and construction of the tracking calorimeter modules. *Nucl. Inst. Meth. A*, A376:36, 1996.
- [53] W.W.M. Allison *et al.* The SOUDAN2 detector: The operation and performance of the tracking calorimeter modules. *Nucl. Inst. Meth. A*, A381:385, 1996.
- [54] W.A. Mann *et al.* The atmospheric flux muon-neutrino / electron-neutrino anomaly as manifestation of proton decay $p \rightarrow e + \text{neutrino}$. *Phys. Letters B*, 291:200, 1992.
- [55] W.W.M. Allison *et al.* Search for the proton decay mode proton to neutrino $k +$ in soudan-2. *Phys. Letters B*, 427:217, 1998.
- [56] M. Thomson *et al.* The observation of underground muons from the direction of cygnus x-3e january 1991 radio flare. *Phys. Letters B*, 269:220, 1991.
- [57] J. Thron *et al.* A search for magnetic monopoles with the soudan-2 detector. *Phys. Rev. D*, 46:4846, 1992.
- [58] W.W.M. Allison *et al.* Measurement of the atmospheric neutrino flavor composition in soudan-2. *Phys. Letters B*, 391:491, 1997.
- [59] W.W.M. Allison *et al.* The atmospheric neutrino flavor ratio from a 3.9 fiducial kiloton-year exposure of soudan 2. *Phys. Letters B*, 449:137, 1999.
- [60] H. Hinterberger and R. Winston. Efficient Light Coupler for Threshold Cerenkov Counters. *Review of Scientific Instruments*, 37:1094, 1966.
- [61] J. Engel *et al.* Nucleus-nucleus collisions and interpretation of cosmic-ray cascades. *Physical Review D*, 46:5013, 1992.
- [62] C. Forti *et al.* Simulation of atmospheric cascades and deep-underground muons. *Physical Review D*, 42:3668, 1990.
- [63] R.S. Fletcher *et al.* SIBYLL: An event generator for simulation of high energy cosmic ray cascades. *Physical Review D*, 50:5710, 1994.
- [64] T.K. Gaisser and T. Stanev. In *Proceedings of the 18th International Cosmic Ray Conference*, volume 3, page 190 and 174, 1983.
- [65] R. Brun *et al.* GEANT3 User's Guide. report dd/ee/84-1, CERN, 1986.

- [66] J. Knapp and D. Heck. Extensive Air Shower Simulation with CORSIKA: A User's Manual. Report kfk 5196b, Kernforschungszentrum Karlsruhe, 1994.
- [67] J. Knapp *et al.* Comparison of hadronic interaction models used in air shower simulations and their influence on shower development and observables. *Forschungszentrum Karlsruhe*, FZKA 5828, 1996.
- [68] F. Abe *et al.* Pseudorapidity distributions of charged particles produced in $\bar{p}p$ interactions at $\sqrt{s} = 630$ and 1800 *gev*. *Physical Review D*, 41:2330, 1990.