Chapter 8

Electronics Systems and Software

This chapter of the report deals with the electronics systems and associated software for the experiment, focusing on the ground array. These systems acquire data from the detectors, process the information, generate an event trigger, transport the data to a central site, archive it, and make it accessible to the collaboration. The optical fluorescence detector electronics and the lower levels of the fluorescence detector trigger will not be discussed in this chapter. They are briefly discussed in Chapter 7.

We begin with a discussion of the electronics located at each station, and follow the data flow back to the central site. Interspersed through the text we will introduce the various components of the multi-level triggering scheme and then tie that information together in a section dedicated to triggering. Figure 8.1 shows the block diagram of the station electronics. Each of the blocks will be discussed in more detail in the following sections.

8.1 Station Electronics Packaging

A careful balance must be kept between modularity, to allow each component of the electronics to be optimized, and a requirement for exceptional reliability, which demands a minimum of interconnections and modules. For example, in order to keep the mean failure rate less than 1 station per day per site, a mean time between failures of 5 years for the station electronics system is required.

The electronics logically splits into four main components: front end, station controller, timing, and communications. It will be powered by solar power through DC-DC power converters. To save the power consumption inherent in bus drivers the front end and communications boards will be incorporated into the final design as daughter boards of the station controller or a motherboard. For reliablity we envision that these may be soldered in place. Provision will be included in the final design to allow the insertion of diagnostic modules and for future expansion. For example, the electronics for an air Čerenkov detector could be plugged into one of the spare slots.

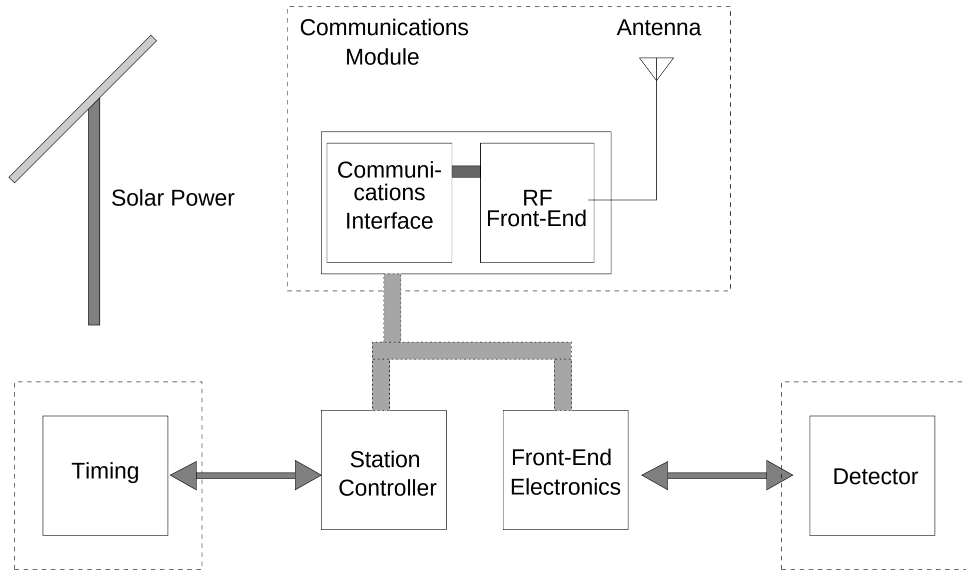

Figure 8.1: Block diagram of ground detector station electronics.

The electronics for each station will be housed in an environmental enclosure in thermal contact with the Čerenkov detector water tank to minimize thermal excursions. The PMT high voltage supplies will be integrated with the tube bases. The solar power system batteries will be housed in a separate enclosure, with a gas-tight electrical connection to the remainder of the electronics. Interconnections between the PMT housings, electronics enclosure, battery enclosure, etc. will be armored to prevent damage by animals and the elements.

8.2 Front End Electronics

The front end electronics at each station must both prepare the raw data for inclusion in a data event, and provide a level 1 trigger. Figure 8.2 shows the block diagram for the water Čerenkov detector front end board.

This board generates level 0 and 1 triggers. These triggers provide a first level of discrimination against single muons and small showers. The cosmic ray particle flux impinging on a 10 m^2 detector is about 2.5 kHz. PMT noise will add another few kHz rate. We will refer to this as the level 0 trigger, which fires any time there is a PMT signal above threshold. The level 1 trigger reduces this several kHz level 0 rate to less than 100 Hz. In practice there will be several level 1 trigger channels, each optimised for particular circumstances, which may be enabled in various combinations via a trigger mask.

There are two main functions of the primary level 1 trigger. One is clipping muon signals to reduce trigger sensitivity to the primary composition. The other is using the integrated time dispersion of the signal to bias against small showers close to the station which deposit signal over a very short time spread. The first function is motivated by the strong dependence of the muon content on the primary species. Suppressing the peak

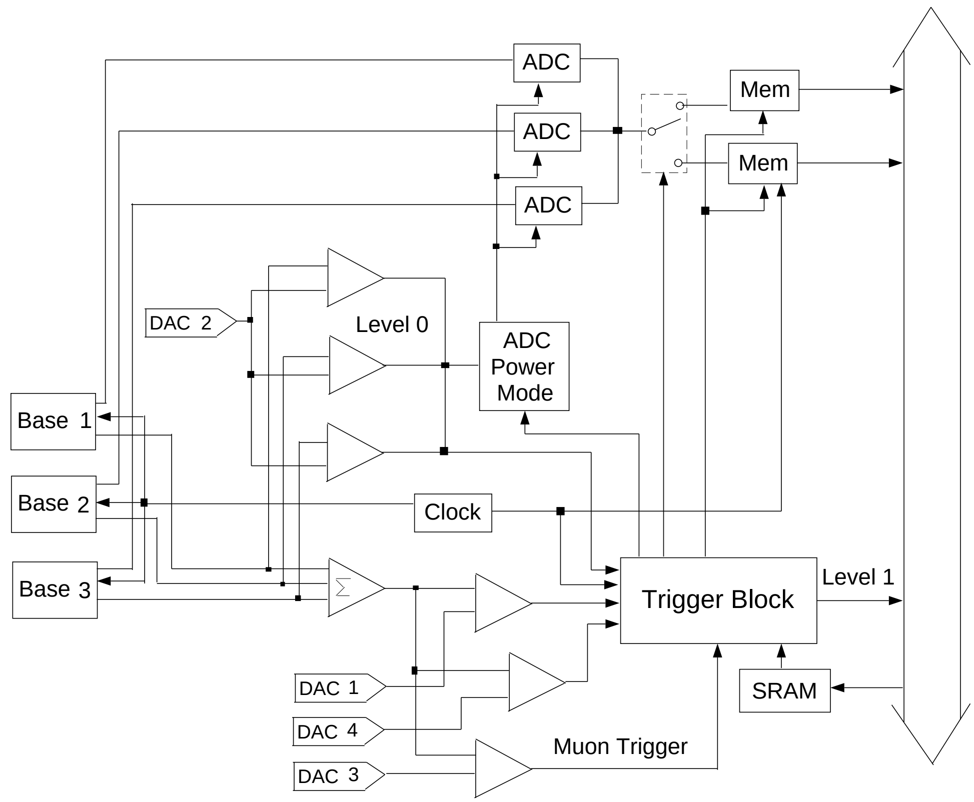

Figure 8.2: Block diagram of a possible front end board.

of the signal pulse generated by a through-going muon before integrating the remaining signal produces a trigger which is less dependent upon the primary particle composition. The second function is motivated by the increase in the time duration of the signal (e.g. the width of the shower front) as one gets farther from the shower core, as indicated in Figure 8.3. Note that the 10-90% rise-time plotted in Figure 8.3 shows little distinction between proton and iron showers, a desirable characteristic for the trigger, whereas in Figure 5.13 the 10-50% rise-time shows a clear distinction. This is because the 10-90% rise-time is dominated by the electromagnetic component of the shower, which arrives late, and depends only weakly upon the primary composition. The 10-50% rise-time, on the other hand, is mainly influenced by the muon content of the shower, which arrives early, and depends strongly upon the primary composition. Requiring that the signal be distributed in time biases the trigger against the abundance of smaller showers which land near the stations. While we want to reject triggers from small showers landing near a station, a variation of this trigger with a larger signal threshold and shorter time duration constraint will be implemented to trigger stations near the core of larger showers.

A “muon” trigger based upon the total light collected will be selectable to facilitate calibration of the PMTs on through-going muons. This muon trigger can be set globally or for a specific set of PMTs, to minimize the trigger bias on muon calibration pulse height distributions. These distributions have been nicely demonstrated using a test tank at AGASA. The feasibility of a special ”muon-decay” trigger involving an isolated muon-like signal followed by a delayed electron-like signal is being studied. A random noise trigger will also be incorporated as an additional diagnostic tool.

To record information useful in searching for Gamma Ray Bursts, a scaler will accumulate a count of GRB level 1 triggers (a clean "electron" signal well above PMT noise) occurring during 0.1 s intervals. This count will be read by the station controller and regularly passed along to the central facility, along with other accumulated data.

The level 1 triggers are sent (except for the GRB triggers) to the timing system and station controller via a dedicated interrupt line, causing the time of each level 1 trigger to be latched and generating an interrupt to the station controller to read out the captured waveform. As an option, the level 1 trigger can be set to disable itself upon formation, valuable for triggering on isolated muons as a calibration and diagnostic tool.

The 100 Hz limit of the level 1 trigger is imposed by dead-time considerations in the electronics. What does this mean in terms of the level 1 trigger condition? Using Poisson statistics, a 5 kHz level 0 trigger rate and a 20 μs trigger aperture (this will be justified later in this section), implies a 450 Hz random doubles rate, and a 23 Hz random triples rate of uncorrelated level 0 triggers. The additional rate of level 1 triggers due to correlated level 0 triggers from small showers can easily be reduced to a few Hz or less using the rise-time constraint. This has been measured at Havarah Park using an 8.9 m^2 water Čerenkov detector and a waveform digitizer. The detector outputs were measured (described in terms of rise-time from 10% to 50%) whenever the signal was above 3.7 particles (vertical equivalent muons). The spectrum was strongly peaked towards fast, bandwidth limited, rise-times. A rise-time cut of >100 ns reduces the rate by a factor of 1000. Our simulations exhibit a similar sharp reduction in the trigger rate when a rise-time cut is applied. The 100 Hz design rate limit for the level 1 trigger is comfortably above what is expected for triggers comparable to those used in the simulations.

Our simulations indicate that a dynamic range of $\geq$ 14 bits and a sampling rate of $\geq$ 40 Mhz is advisable to make optimum use of the PMT signals. The large dynamic range facilitates reconstruction of the lateral profile of the showers while the fast sampling facilitates extracting parameters related to the muon component of the shower. Fortunately, only modest precision (8-10 bits) is required and the large dynamic range is not required at the fast sample rate. We are investigating several options, including multiple ADCs with overlapping ranges, analog compression of the signals, and switched capacitor arrays [152] with multi-range digitization. CMOS parts can be used for almost all of the components, minimizing power consumption. The front end board will have buffering for at least 2 events, minimizing the dead time incurred transferring the recorded waveforms to the station controller.

With careful design, power requirements for the front-end electronics can be kept reasonably low. Two factors contribute to this: the timing resolution requirements are modest enough to be met with fast CMOS, and the duty cycle for cosmic ray events is low so that most logic can remain quiescent most of the time. Judicious use of ASIC circuits can also help to reduce power consumption and increase reliability.

How long does the FADC record need to be? A conservative approach would be to

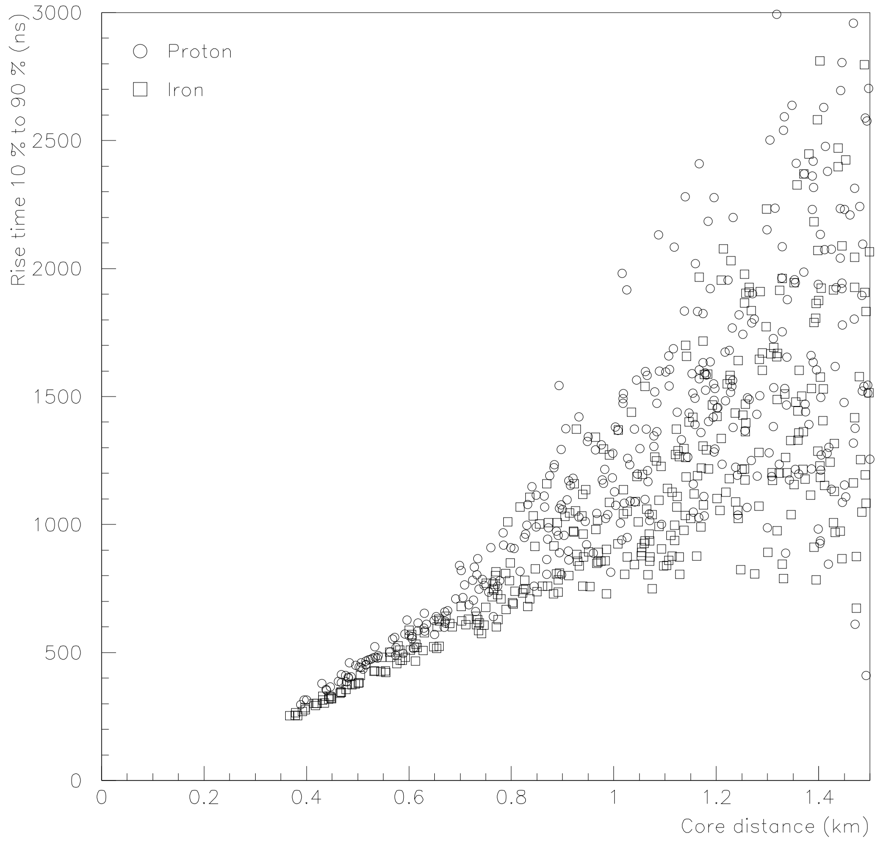

Figure 8.3: Distribution of 10-90% "rise-time" as a function of distance from the shower core for a sample of 100 simulated 10^{19}eV proton and iron showers. (*Cf.* Figure 5.13 where the 10-50% rise-time is displayed.)

have the record long enough to handle the full waveform of the very largest events one might expect to see. If they exist, showers with energies $\geq 10^{21}$ eV might be detected at a rate of 1 per several years. These showers will trigger stations out to about 5 km from the shower core. At this distance from the shower core a 20 μs window is required to catch $\approx$ 95% of the particles. If we run the FADC system for 20 μs each time a $\sim$1/2 photo-electron discriminator level is crossed ($\sim$5 kHz level 0 trigger), the system will be active only 10% of the time. Only the front-end amplifiers, comparators, and the clock crystal need to be operating in the intervening 90% of the time. This can reduce the power consumption of the front-end significantly. For example, the estimated power consumption was reduced from 12 W to 2.6 W in one design exercise. The FADC window will be extended for 15 μs after any level 1 trigger, so that the tail end of the waveforms is not lost even for events in which the FADC was started on a random signal. See Figure 5.10 for an example of the operation of such a trigger in a simulated air shower event.

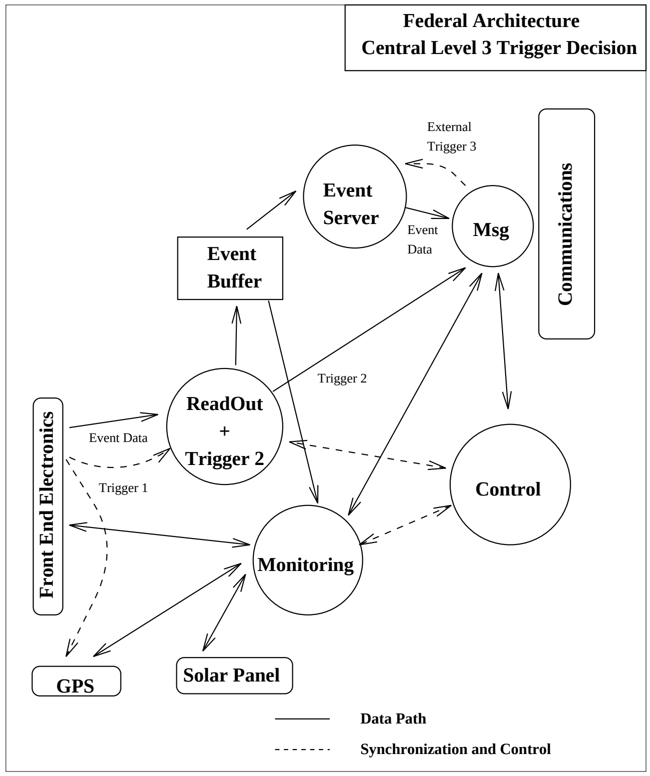

Figure 8.4: Block diagram of station controller code.

8.3 Station Controller

The station controller provides the local intelligence at each detector station. It performs a large number of tasks: controlling the PMT HV supplies, monitoring HV, solar panel performance, battery state, etc., reading data from the front end electronics, reading the time tag for each event from the timing system and applying necessary corrections, performing one or more levels of trigger logic, transferring data to the central data acquisition system, and local detector monitoring. In order to perform these tasks, the station controller will use a real time operating system such as OS-9. The station controller board will integrate these functions onto a compact low power board utilizing a highly integrated low power microcontroller, such as the IBM PPC403. We have benchmarked major elements of the station controller code and verified that this processor, which has extremely low power consumption (0.25 W) is fast enough for our purposes, with a comfortable factor of 10 safety margin. Figure 8.4 shows the block diagram for the station controller code.

It will not, in general, be possible to reset a misbehaving station remotely. Therefore, the station controller code must be robust and must monitor itself for satisfactory operation, automatically restarting itself as necessary. Similarly, it will not be feasible to change EPROMs in over 1600 stations scattered over 3,000 km^2 in order to install updated station controller operating code. Operating code and run conditions must be down-loadable through the communications network.

Only a minimal, very thoroughly tested and debugged, core program will be stored in the PROM. The primary function of the core program will be to reset PMT high voltages to a safe state, and to interact with the communications network to download operating code. The operating code will contain modular links so that selected portions of the code can be updated via the communications network, minimizing the frequency that the entire program needs to be transmitted when program updates are installed. The entire program, including operating parameter storage, will reside in less than 0.25 Mb of memory.

The station controller implements the level 2 trigger in software. It analyzes the digitized waveforms from the three PMTs in order to tag muons, determine the total signal magnitude, and characterize the time distribution of the signal. Using this information it generates the level 2 trigger, reducing the trigger rate from ≤ 100 Hz to ≤ 20 Hz per station.

The waveforms are highly compressible. The amount of compression depends upon the amount of low level noise present in the signal and the complexity of the compression algorithm. Our most pessimistic simulations indicate that a compression factor of 3 is easily achievable. Thus 4 Mb of buffer memory can store ~ 100 seconds of level 2 triggers. Implementing this amount of buffer memory is quite feasible, and is comfortably longer than the 10 second maximum trigger formation time.

To minimize the number of boards, many of the miscellaneous interfaces will be incorporated on the station controller board. The station controller board will incorporate sampling ADCs for monitoring the status of the solar power system and PMT high voltages. It will contain DACs to control the PMT high voltage supplies and incorporate a GPS interface with a ≥ 2 event FIFO.

8.4 Timing

Each station needs an accurate time base to 1) match level 2 trigger data with its neighbors, 2) allow shower directions to be reconstructed, and 3) provide a time base for the communications system. In the reference design the timing system is based upon Global Positioning System (GPS) receivers. The Leeds group has demonstrated that 10 ns relative timing can be obtained using inexpensive off-the-shelf scientific GPS receivers. Power consumption is less than 1W.

To allow cross-matching of cosmic ray event times for on-line shower recognition and off-line direction reconstruction it is crucial that data acquisition at each detector station be synchronized. We have simulated the effect of clock synchronization error on the angular resolution of the detector. Our simulations indicate that 15 ns synchronization error will contribute $\sim 0.15°$ to the direction resolution. Given that our intrinsic angular resolution may in some cases be less than $0.5°$ (see Chapter 5), this is a reasonable requirement for the timing system.

The GPS is a US government satellite constellation which continuously disseminates position and time information worldwide by means of coded 1.5 GHz radio broadcast signals. The system is now fully operational and expected to remain so for several decades.

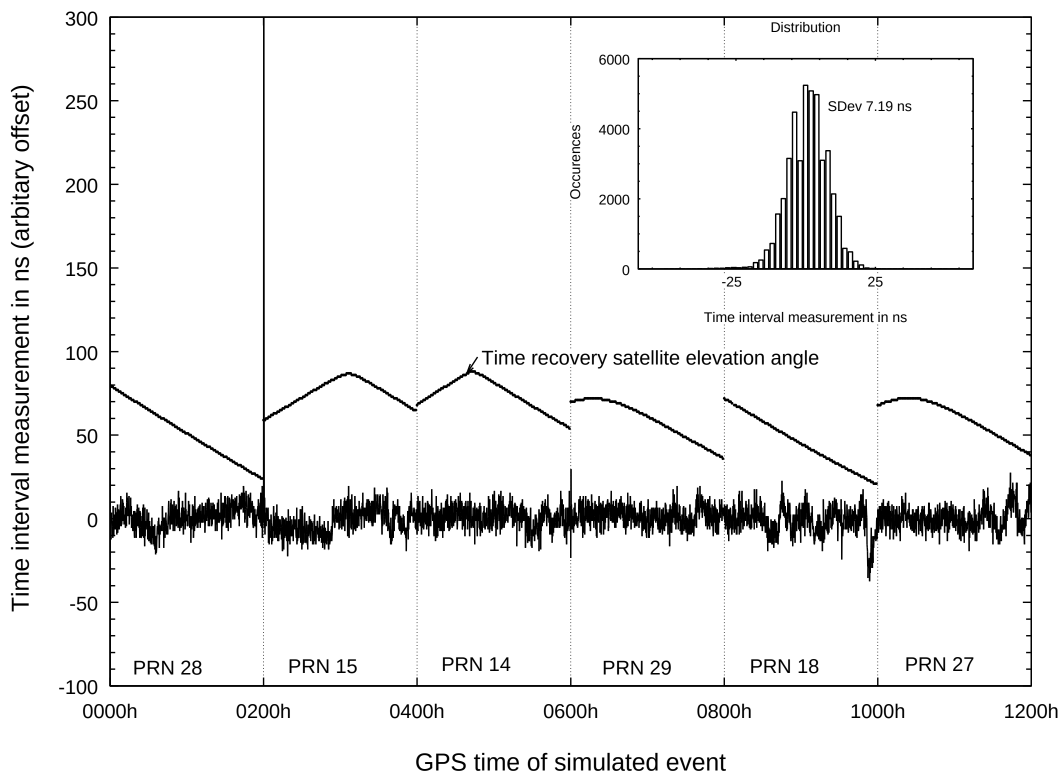

Figure 8.5: Apparent event time difference with the time-tagging systems separated by 11.2 km on the morning of 26/1/95. The upper line of the main graph shows the elevation angle of the satellite being used for time recovery, its PRN code being given at the bottom of the figure. The inset graph shows the distribution of the time interval errors shown in the main graph.

Unrestricted civilian access to the deliberately degraded Standard Positioning Service (SPS) is allowed. By the use of a simple technique known as "scheduled common view" the effect of this degradation can be eliminated when only relative accuracy is required.

The prototype design is very simple and has demonstrated excellent results[149]. The basic technique is to exploit a commercial GPS receiver module, in conjunction with a custom designed 100 MHz twin channel counter latch assembly. These are both interfaced with a computer. In the final design the interface will include a multi-event FIFO for the event trigger latch, and incorporated on the station controller board.

One of the latch triggers is fed from the GPS receiver 1 pps output, and the other from the event trigger source. By reading the latch values as they occur, in conjunction with information from the GPS receiver, the computer can calculate event time stamps by simple linear interpolation. In effect the GPS 1 pps is used to continuously calibrate the local counter time frame.

To remove the effect of the deliberate SPS accuracy degradation the scheduled common view technique referred to above is employed. The control computer instructs the GPS receiver to base the 1 pps output on information from a single satellite, which is periodically changed since the satellites remain visible for less than six hours at a time. Each system in the network uses the same satellite simultaneously by means of a pre-determined tracking list.

Initially two prototype systems were constructed which were tested under two circumstances: (a) at the same location and (b) separated by 500 m using a high bandwidth cable link. The time interval measurement error distribution had a standard deviation of 6 ns, with a systematic error of 5 ns. The maximum errors observed during 8 days of operation, with 800,000 test triggers, were +21ns and -35ns.

The same two systems were tested over an 11 km baseline using a microwave link which form part of the MERLIN radio telescope interferometer centered at Jodrell Bank, UK. Figure 8.5 shows the results of this test, in which the time interval measurement had a standard deviation of 7 ns. It was not possible to measure directly the systematic error in this test. However, due to the variable geometry which occurs in the GPS system it is very hard to envisage a fixed source of error scaling with distance of separation. The largest EAS so far observed have footprint diameters of only 5-10 km. Even the largest showers will not have level 2 triggered stations more than about 5 km from the shower core.

Two more sophisticated systems based on the same technique have recently been deployed at the South Pole to allow cross-matching of the SPASE and AMANDA experimental data streams.

8.5 Communications

The communications system is used to correlate data from different detector stations to form an event trigger, transfer event data from detector stations to a central site, and transfer control and diagnostic information bi-directionally between the detector stations and the central site. The bandwidth required of the communications system depends not only upon the station trigger rate, but also on the topology chosen for the data links, and the level of local multi-station coincidence selected. The communications system design is based upon a maximum mean level 2 trigger rate of 20 Hz. We have chosen a wireless RF communication system for the reference design. Since the performance of the communications system is critical for the operation of the experiment, we will devote a number of pages to discussing the design considerations. For additional information refer to Refs [147] and [153].

Fully wired systems (copper or fiber optics) are easily rejected on the basis of cost considerations. For example, the length of copper or fiber (and trenches) needed to connect the stations is 7,200 km per site with 6 connections per station, or 3,600 km with only 3 connections per station (with minimal redundancy in case of link failure). The installed cost of copper or fiber is about $10K/km. This is dominated by the installation cost and depends somewhat on site conditions. Even so, at an estimated cost exceeding $72M, we need not consider this option further. However, the utilization of a fiber backbone may be advantageous. Further studies are ongoing.

8.5.1 Licensing Issues

Any RF transmitter must comply with regulatory requirements in the site countries. This becomes a complex issue, as the regulations unfortunately vary from country to country. For example, in the US several of the ISM (Industrial, Scientific, Medical) bands are attractive to consider for the station transceivers, as these bands allow operation (of type-certified equipment) without purchasing a license for each station. In the US, the 902-928 MHz, 2400-2483.5 MHz, and 5725-5850 MHz bands are the most appropriate. FCC regulations part 15.247 [150] limit the maximum peak transmitter output power to 1 Watt, and the maximum equivalent isotropic radiated power (EIRP) to 4 Watts (36 dBm) in these bands. In addition, the transmitted signal must be spread in bandwidth. The regulations in Argentina are essentially identical to those of the US in those bands. This makes it attractive to consider them for the reference design. A drawback of the ISM bands is that anyone can use them. This is not a problem for a controlled access site where we could restrict the use of other devices which might interfere with our communications system, but is a concern if the general public has access to the site. We will continue to investigate the use of other bands as we explore various options for the communications hardware.

8.5.2 Propagation

A critical parameter for the communication system is the strength of received signals relative to the noise. This will be function of the transmitter power, the transmitter and receiver antenna gains, propagation path loss, and atmospheric attenuation. Fortunately, atmospheric attenuation is less than 1 dB/km at frequencies below 3 GHz even in the heaviest rainfall, and thus is not a serious concern for propagation over 1.5 km distances.

Radio wave propagation in free space is given by the Friis equation[151]. For the case where the transmitting and receiving antennas are aligned and polarization matched, this equation reduces to

$$\frac{P_r}{P_t} = \left(\frac{\lambda}{4\pi D}\right)^2 G_t G_r \tag{8.1}$$

where

$$\begin{aligned}
P_t &= \text{power at output of transmitter,} \\
P_r &= \text{received power at input to receiver,} \\
D &= \text{distance between transmitting and receiving antennas (in meters),} \\
\lambda &= \text{wavelength of the radio signal (in meters),} \\
G_t &= \text{gain of transmitting antenna, and} \\
G_r &= \text{gain of receiving antenna.}
\end{aligned}$$

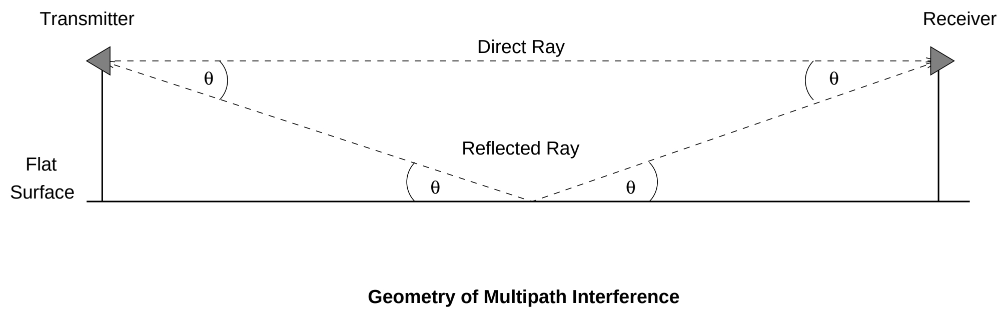

Figure 8.6: Multipath interference of radio waves.

However, reflection of the radio waves from the ground and subsequent interference with the direct waves is an important factor which cannot be neglected. This is depicted in Figure 8.6.

A further complication is that the ground will not be precisely flat, but will have some undulations in it. Figure 8.7 displays the calculated propagation loss versus distance between two 8 m high vertical full wave dipole antennas at 2450 MHz. The shaded area indicates the loss for a perfectly flat terrain. Each solid curve represents a different randomly generated terrain. The envelope does not vary strongly with terrain variations, as long as line of sight between the 2 antennas is maintained. The 8m antenna height is near optimum at 2450 MHz. At lower frequencies the received signal strength continues to increase for higher antennas.

The points where a bad fade occurs are problematic. One effective technique that is often used to combat bad fades is *antenna diversity*. That is, use multiple receiving antennas and pick the one that gives the strongest desired signal. We have simulated a two antenna receiver, where the second antenna is placed vertically below the first. The height of the second antenna is chosen such that the phase of the direct and reflected waves from a flat earth is changed by 90°. The calculated propagation losses are displayed in the right plot in Figure 8.7.

Another technique that is often used to combat fading is *frequency diversity*. That is, spread the signal over multiple frequencies (or a wide bandwidth) to mitigate the effects of a bad fade at one frequency. This will be effective if the path length for the interfering reflected signal is many wavelengths longer than the direct path. This can be the case for reflections from objects not in the plane of the two antennas. However, these effects are not fully modeled in our current simulations, which only include reflections in the plane of the antenna systems.

Spot checks of these simulations were made in 1995 at Woomera, Australia. At 900 Mhz a handful of measurements were made with 1.5 km spacing between antennas at different locations. The measurements agree within a few dB in the mean to that predicted by our simulations. At 2.4 Ghz measurements were only made at 2 locations. Both are ~10 dB below the predicted mean, but are within the envelope of expected path loss. These measurements confirm that there are no gross errors in the simulation, but with only a handful of measurements it is not possible to confirm the details of the predicted path loss distribution. Further radio survey measurements at the selected sites are planned.

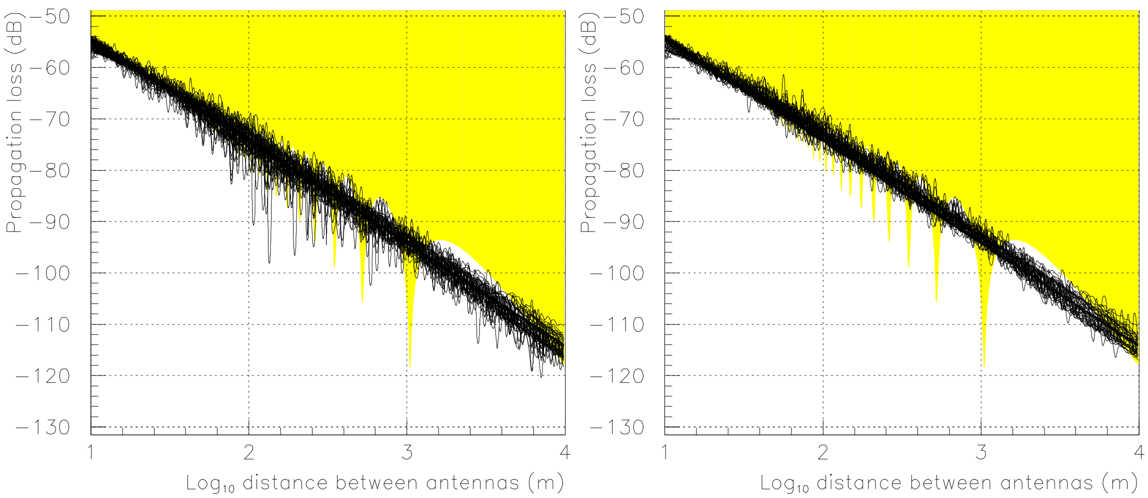

Figure 8.7: Calculated 2450 Mhz radio wave propagation loss for 8 m high full wave dipole antennas. The shaded curve is for an ideal flat plane. The solid lines show simulation results for 100 different random terrains with up to ± 4 m hills). The left plot depicts the case where one receiving antenna is used. The right plot depicts the case where the strongest signal from one of two vertically separated antennas is used.

The distribution of received signal strength in communications systems where there is a strong direct signal in the presence of a number of reflected signals (called multipath) can be shown to follow a Rician distribution. S_i, the received signal strength (Watts) from antenna i, has the (Rician) probability density function

$$f_1(s) = \frac{1}{2\sigma^2} \exp\left(-\frac{s + \alpha_i^2}{2\sigma^2}\right) I_0\left(\frac{\sqrt{s}\alpha_i}{\sigma^2}\right) \tag{8.2}$$

The distribution function is

$$F_1(s) = 1 - Q\left(\frac{\alpha_i}{\sigma}, \frac{\sqrt{s}}{\sigma}\right)$$

where $Q(a,b)$ is called the Marcum Q function and is given by

$$Q(a,b) = \int_b^\infty \exp\left(-\frac{a^2 + x^2}{2}\right) I_0(ax)dx.$$

The receiver measures S_1 and S_2 to determine which antenna has better signal-to-noise ratio (assuming the noise power and signal distributions are the same at both antennas). If we let $S = \max(S_1, S_2)$ then the density of S is

$$f(s) = 2f_1(s)F_1(s) \tag{8.3}$$

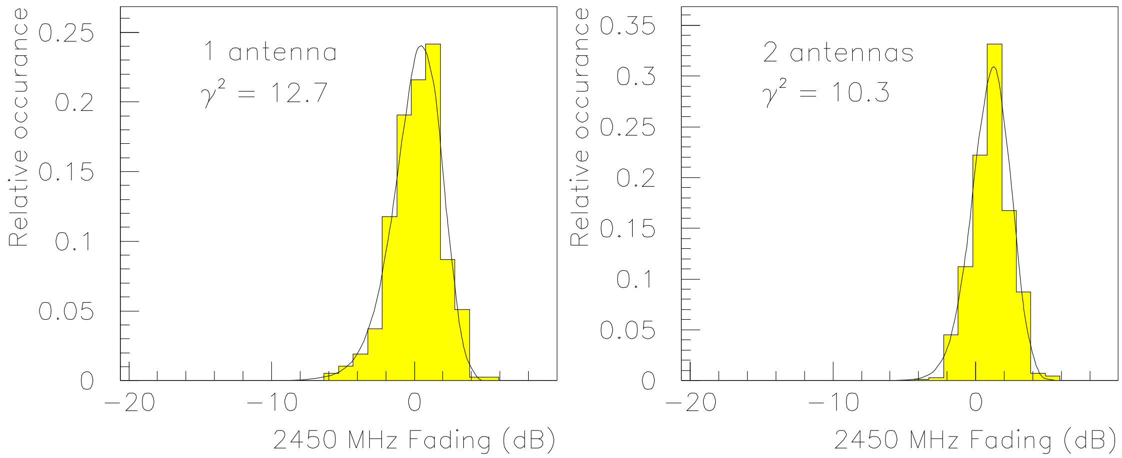

Figure 8.8: Calculated distribution of received signal strength relative to free space for 2450 Mhz radio wave propagation between 8m high full wave dipole antennas separated by 1.5 km. The histograms display the simulation results for 100 random terrains. The curve in the left figure is a least squares fit to Eqn. 8.2. The curve in the right figure is a fit to Eqn. 8.3.

The ratio $\gamma^2 = \alpha^2/(2\sigma^2)$, which is the ratio of received power in the direct path to that in the faded path, is important when determining the performance. The total average received energy is $\Gamma = E(\alpha_i^2 + 2\sigma^2)$, where E is the received energy in the direct path. The error probability shown in later figures is plotted as a function of Γ/N_0, where N_0 is the noise power density.

The calculated distributions of propagation losses are displayed in Figure 8.8, along with least squares fits to $f_1(s)$ (left plot) and $f(s)$ (right plot). As expected by the assumption that the signals at the two receiving antennas have the same probability distribution, the fits produce the same value of γ^2 within errors, with a mean value of $\gamma^2 = 11.5$.

8.5.3 Performance Analysis of Frequency-Hopped Spread Spectrum.

In this section we will analyze the link level performance of a particular data communications system to demonstrate the feasibility of our communications links, using as an example the Plessey DE6003 transceiver [148]. This is a frequency-hopped 2.4 GHz ISM band spread spectrum transceiver with two antenna inputs. A typical frequency-hopped spread spectrum system consists of a data source followed by an error control encoder. The error control encoder adds redundant bits to the data stream to protect against errors. The encoded data is then modulated using binary frequency-shift-keying (BFSK). The modulated signal is frequency-hopped over a set of frequencies. One goal in transmitting information at different frequencies is that the fading statistics will be different at different frequencies. Thus a given multipath situation will yield a particular fading statistic at one frequency but a different (and independent) fading statistic at another frequency. If a small percentage of the frequencies result in poor performance then an error correcting code will be able to correct the errors that occur when transmitting at those frequencies. The receiver consists of a dehopper/demodulator followed by the error control decoder.

There are several types of codes that could be used. These include convolutional codes and block codes. The convolutional codes are easier to decode with reliability information from the channel (eg. the received signal strength). The block codes have a larger variety of code rates (number of information symbols per channel symbol) available and are slightly easier to analyze. We will only consider the use of Reed-Solomon block codes (used in CD players) in this report.

In a block code (such as a Reed-Solomon code) a number of information symbols k is protected by transmitting not only the information symbols but in addition transmitting $n-k$ redundant symbols. These redundant symbols depend on the information symbols so that if an error occurs during transmission and several symbols are received erroneously the correct information can still be determined. Reed-Solomon codes operate on non-binary symbols typically of size 32 to 256 bits. These usually are called bytes. A non-binary symbol is formed by grouping together bits. For example, in a Reed-Solomon code with symbol size 32, each symbol consists of sequences of 5 bits. The encoder groups information symbols into blocks for the purpose of encoding. For the Reed-Solomon codes considered here $k \cdot m$ bits are grouped together. First m bits are grouped and called a code symbol. Based on k information symbols, $n-k$ redundant symbols are determined. The codeword then consists of the k information symbols followed by the $n-k$ redundant symbols. The decoder can correct any pattern of $e = \lfloor (n-k)/2 \rfloor$ errors.

The codeword generated by the Reed-Solomon encoder can be transmitted using BFSK. Usually there are a large number of bits transmitted during each frequency hop. However, a well designed system would transmit only one code symbol per hop. In order to transmit many bits per hop, and only one code symbol per hop, interleaving is required. Interleaving works as follows: instead of generating a single codeword at a time, L codewords are generated (from different data symbols) simultaneously. Then only one code symbol from each of the L codewords is transmitted during a particular hop. Thus there are $L \cdot m$ bits transmitted per hop.

It is very important to know what the correlation is between the amplitude of the fading at two different frequencies. If the fading amplitudes are independent then error control coding can significantly improve the performance by correcting the errors due to the occasional bad fade at a particular frequency. If, on the other hand, the fading amplitudes are highly correlated between different frequencies, then error control coding will not be effective in reducing the error probability. As we do not yet have survey measurements or simulation results that can accurately estimate this correlation, we consider the two extreme cases: 1) different frequencies are independently faded and 2) different frequencies are identically faded (completely correlated).

The conditional error probability for making an error in determining the value of a received data bit, given a particular signal strength, is given by

$$P_e(S) = \frac{1}{2} \exp - \frac{ES}{2N_0}$$

The average error probability is given by

$$\bar{P}_e = \int_0^\infty f(s) P_e(s) ds$$

For coding purposes we are interested in the probability of error in a code symbol which consists of m consecutive bits transmitted at the same frequency and having the same fading value (not only the statistics are the same but the realization is the same). The conditional symbol error probability is given by

$$P_s(S) = 1 - (1 - P_e(S))^m$$

The average symbol error probability is given by

$$\bar{P}_s = \int_0^\infty f(s) P_s(s) ds$$

Given that the different code symbols are transmitted on different hops, the probability that a Reed-Solomon codeword is decoded incorrectly is the probability that there are more errors than can be corrected by the code. The error correcting capability is $e = \lfloor (n-k)/2 \rfloor$. The probability of error is then

$$\bar{P} = \sum_{l=e+1}^{n} \binom{n}{l} \bar{P}_s^l (1 - \bar{P}_s)^{n-l}.$$

It is usually the case that a Reed-Solomon decoder will fail to decode when there are more errors than can be corrected, rather than put out an incorrect codeword. This means that decoder errors will (with high probability) be detected.

Finally we consider the case where all code symbols experience the exact same fading amplitude. For this case the codeword error probability is given by

$$\bar{P} = \int_0^\infty f(s) \sum_{l=e+1}^{n} \binom{n}{l} P_s(s)^l (1 - P_s(s))^{n-l} ds.$$

Later we show the large difference between coding for independently faded channels and identically faded channels.

8.5.4 Link Budget

In this section we consider the link budget for a system using a 20 dBm (100 mW) transmitter with stations separated by 1.5km. The transmitter and receiver have antennas which each

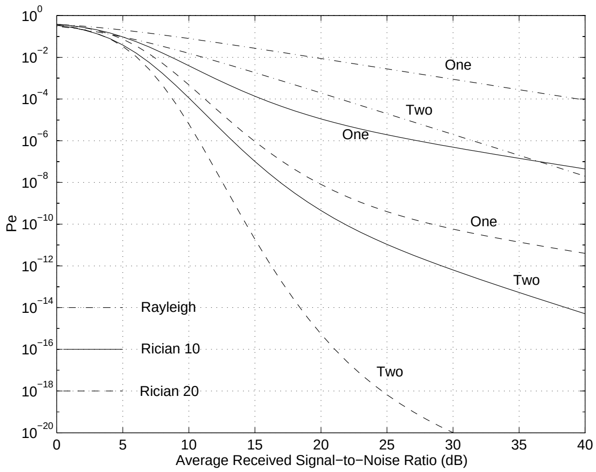

Figure 8.9: Bit error probability for fading without coding with one and two antennas.

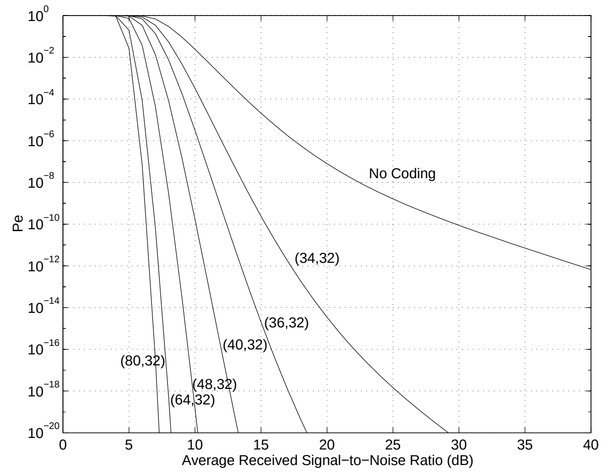

Figure 8.10: Block error probability for Rician fading with coding and symbols independently faded ($\gamma^2 = 10.0$).

provide 3 dB antenna gain. The loss due to free space at frequency 2.45 GHz is 103.7 dB. Thus the transmitted signal power is reduced by 97.7 dB. For a 100 mW transmitter the received power in the direct path is 16.8×10^{-12} W. The data rate considered (that used in the Plessey DE 6003) is 625kbps. From this the (direct path) energy per bit is given by $16.8 \times 10^{-12}/(625 \times 10^3) = 2.7 \times 10^{-17}$ joules/bit. In addition there is the energy from the multipath components. Based on simulations the ratio of the energy received from direct path and multipath components ($\gamma^2 = \alpha^2/(2\sigma^2)$) is 11.5. The total received energy is then $\Gamma = 2.7 \times 10^{-17}(1 + 1/11.5) = 2.9 \times 10^{-17}$ joules/bit. The noise level (from the DE6003 data sheet) is $N_0 B_{IF} = -103$dBm, where B_{IF} is the bandwidth of the IF filter (1.5 MHz for the DE6003). Thus the noise power density $N_0 = 3.3 \times 10^{-20}$ W/Hz. The average received signal-to-noise ratio is $\Gamma/N_0 = 29.5$dB. There are potential implementation losses which reduce the effective signal-to-noise ratio. Plessey estimates these to be of the order of 5dB for the DE6003. Thus the effective received signal-to-noise ratio of 24 dB may be appropriate (in the absence of other interfering signals). Although our simulations indicate $\gamma^2 \approx 11.5$, we will use a more conservative $\gamma^2 = 10$ below.

Interference from other transmitting stations must also be considered, as this may raise the noise. We are considering using a 7 slot time division multiple access (TDMA) scheme with about 237 stations transmitting simultaneously. In a frequency hopping system with 79 or more adjacent channels (like the DE6003) there would be three stations transmitting on each channel at the same time, and an additional six stations transmitting on neighboring channels. With fixed station positions we can arrange the frequency hopping sequence to maximize the distance between interfering stations. In that case, interference from other transmitting stations will typically be 20 dB below the desired signal.

The bit error probability for an uncoded system with one or two antennas is plotted against the average received signal-to-noise ratio in Figure 8.9. In the case of two antennas the receiver is assumed to be able to select whichever antenna has the largest signal. It is clear from this figure that significant gains can be achieved by using two antennas. The three sets of curves correspond to the case of no direct path (Rayleigh, pessimistic), a direct path and faded path with $\gamma^2 = 10$ (Rician 10, realistic) and with $\gamma^2 = 20$ (Rician 20, optimistic).

The codes we considered all operate with the same number of information bits (256) grouped into 32 8-bit bytes. In the uncoded case we just transmit these 32 symbols without any extra redundant bits. The codes have different block lengths n but the same number of information symbols (32). The (48,32) code has $n - k = 16$ and can correct 8 errors in the block of 48. The (64,32) code has $n - k = 32$ and can correct 16 errors in a block of 64. In Figure 8.10 the block error probability is shown for the case where each symbol is independently faded, and the power in the direct path is ten times the power of the faded path (Rician 10). In this case the block error rate may be as low as 1 part in 10^{14}. However, if the fading is identical at different frequencies, the error coding is much less effective at reducing the error probability, reducing it by only a factor of two.

Radio surveys of the selected sites will permit us to refine these calculations, but at present it appears that commercial transceivers like the Plessey DE6003 will provide a satisfactory block error rate level when combined with error correcting codes.

Thus we are comfortable that a radio communications network poses no intrinsic tchnical roadblocks. The actual implementation will undoubtedly differ in the details. In the next section we will consider in more general terms the requirements for Auger.

8.5.5 Network Topology

Once we establish reliable radio links which provide a path via one or more hops between each station and the control center we have, in principle, the means to transfer the information necessary to operate the project. This is trigger information, data packets, software downloads, etc. The network topology and protocols are primarily implemented in software that operates on the communications hardware. It will be advantageous to conform to international standards, such as the ISO/OSI 7 layer model, as much as possible.

Critical parameters for the communications network are the user data rate and the allowable station to central collection point latency. The data rate from each station is dominated by level 2 trigger packets and monitoring data. Our 20 Hz level 2 trigger rate limit requires 480 bits/second mean rate from each and every station. Adding a 220 bits/second estimate for monitoring data gives us a 700 bits/second user data rate transmitting from each station. It should be noted that the bandwidth required to read out triggered events is totally neglible for these considerations. The bit rate for data from the central collection point to each station is dominated by by the constraint that the (infrequent) software downloads should not take an excessively long time. 600 bits/second is adequate here, but this is admittedly a softer number than the requirement for the other direction. The communications system impelementation should have provisions for upgrades that increase the data bandwidth by a factor of 2 to allow the addition of later enhancements to the ground array. Additional rate is required for packetizing and forward error correcting coding. In many schemes additional headroom is required to avoid packet loss due to congestion. These latter are highly dependent upon the details of the implementation, hence are difficult to specify in general.

There is a trade off between maximum network latency and the amount of memory required in the ground array stations and fluorescence detector electronics to buffer data while awaiting a trigger decision. A 5 second maximum latency specification in each direction yields buffer memory requirements which can be reasonably achieved in both the ground stations and the fluorescence detector. It can also be met by many of the communications schemes we have studied.

It is desirable for most stations to have line-of-site communication with its six nearest neighbors. In many implementations, this provides alternative communication links in case of station failures, a prerequisite for a fault-tolerant network. Occasional obstacles (e.g. hills) can be tolerated.

We have considered two general classes for the network topology: the "flat" or "fraternal" architecture where the level 3 trigger decision is performed by the station controllers; and the "hierarchical" or "federal" architecture where local stations send all level 2 trigger data to a central system via concentrators.

Because of the possibility of "exotic" physics which we may not fully anticipate at the present time, there is a strong motivation to process as many of the trigger levels as possible at a central point, where information from the entire array can be included in the trigger pattern recognition. Therefore, we have chosen the federal network topology. One example of a hierarchical network topology is depicted in Figure 8.11.

In each lowest level cell the concentration may be either by point-to-point links or TDMA global collecting down to the six "first ring" stations surrounding the local base station (see Figure 8.12).

Higher bandwidth is required as the data is concentrated towards the center. This is satisfied with a hierarchy of directional microwave links or fiber backbones. Other implementations of this basic scheme are currently under study.

The concentrator network will not only carry the data from the ground array stations, but also provide links between fluorescence eyes, and between the eyes and the control center.

8.6 Solar Power

The ground array stations will operate from solar power (photovoltaic, PV). The power budget for each station shown in Figure 8.13. The PV system must provide regulated power to the station electronics with minimal maintenance over the expected lifetime of the experiment. Each station will have its own PV array, battery bank, charging regulator, low battery voltage disconnect, and control/monitoring system as shown in Figure 8.13. The orientation of the PV array will be fixed at an angle which maximizes the solar exposure during the winter.

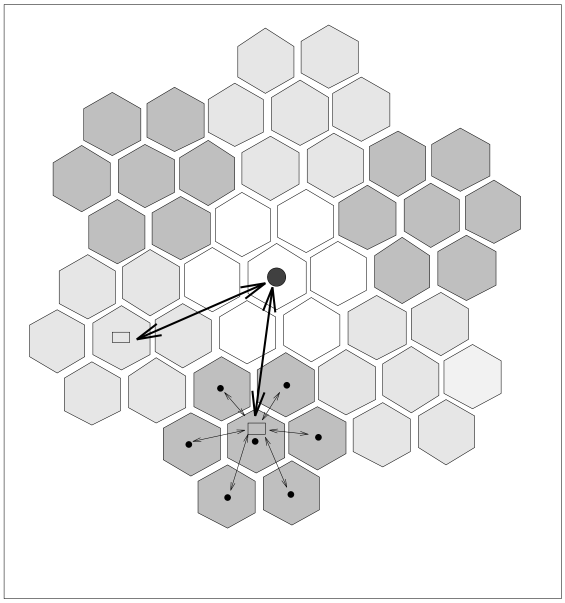

Figure 8.11: Cell configuration in the 2-level hierarchy. Each of the small hexagons represents one cell of stations depicted in Fig. 8.12. A base station at the center of each such cell communicates with a microwave repeater at the center of each cluster of 6 cells. These repeaters in turn communicate with microwave tranceivers at the center of the array.

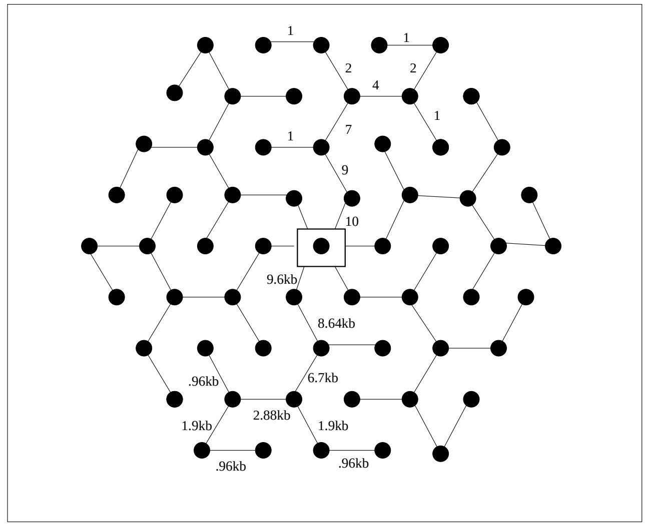

Figure 8.12: Communication links and data flow rates in a level-2 mini-cell. Each dot represents a detector station. The lines indicate the radio communication links between stations. At the center is a base station which communicates with a microwave repeater (see Fig. 8.11).

The number of required solar panels and size of the battery bank are estimated from a knowledge of the daily power requirements, the estimated availability of solar radiation at the site, and the percentage of down time due to insufficient power that can be tolerated. The statistical nature of the event data means that pushing for much less than 2% outage due to insufficient solar power is probably not warranted if it adds substantial cost to the system. However, as it turns out, this goal appears to be easily achieved.

For planning purposes we have sized the system for a site with the solar radiation characteristics of Dugway, Utah. Both of the selected sites have better characteristics than this. A simple but conservative analysis has led us to propose the use of two 53 peak Watt (Wp) silicon solar panel arrays for power collection and two 12 volt gel-cell lead-acid batteries for power storage. This provides sufficient Wp capacity such that with average insolation at Dugway (minimum of 3 hours Wp mid-winter) the batteries never discharge to 50%. The battery capacity is sufficient to power the system for more than 10 days without sun, discharging to 50%. The technology of both of these components is mature and their commercial use is common.

Although the outage probability has not yet been accurately estimated, this system is expected to perform much better than the design goal of less than 2% outage probability per station. A more detailed analysis of the system sizing and outage probability will be

Station Power Budget	
System	Power Rqd (W)
High voltage	1.0
Front end	2.6
Station controller	2.5
GPS receiver	0.85
Communications transceiver	0.83
Transceiver interface	0.5
Total	8.3
Converter losses	×1.2
Total	9.9

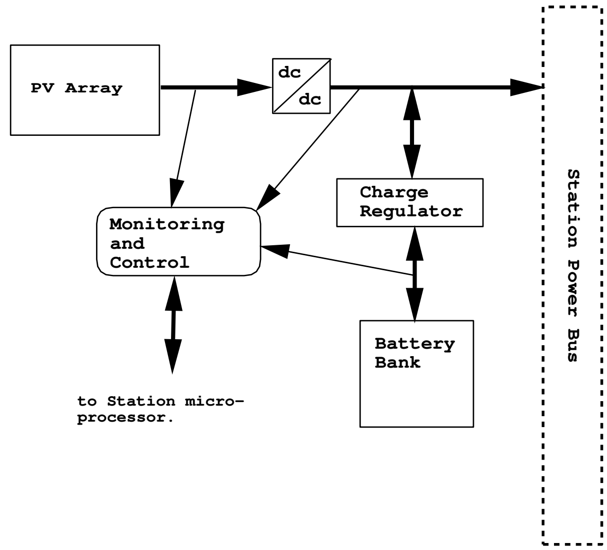

Figure 8.13: Block diagram of the solar power system for each station

performed using accumulated insolation data for the two sites in due course.

8.7 Data Acquisition and Archives

The collection of data at the central site will be handled by Unix workstations. Figure 8.14 shows the block diagram for the control center data acquisition (DAQ) code.

The philosophy of the DAQ is the following. The ground array and the fluorescence detector have their separate acquisition hardware and computers. For each detector, there are two servers: one for the data acquisition (trigger, event-building) and archiving, the other for the monitoring, calibration and slow control tasks. The computing resources required are modest by present day standards and can easily be handled by low cost workstations. Each would have ~128 Mb of central memory, a ~4 Gb system disk, ~8 Gb of temporary data storage space, and an output device (e.g. DAT tape or CD-ROM) where the data is permanently and locally stored. Software and licences were also considered (including FORTRAN, PASCAL and C compilers).

A relatively high speed Internet access (>1 Mbit/s) to the site will be provided via fiber optic link or satellite link. This will allow collaborators to monitor the experiment from anywhere in the world and transfer moderate amounts of data. The full data sets will be transferred over this link to a primary distribution center in each hemisphere which will make the data available to the collaboration.

The data acquisition hardware can only be outlined at present. We envisage a few VME crates. Absolute time reference will be obtained from GPS.

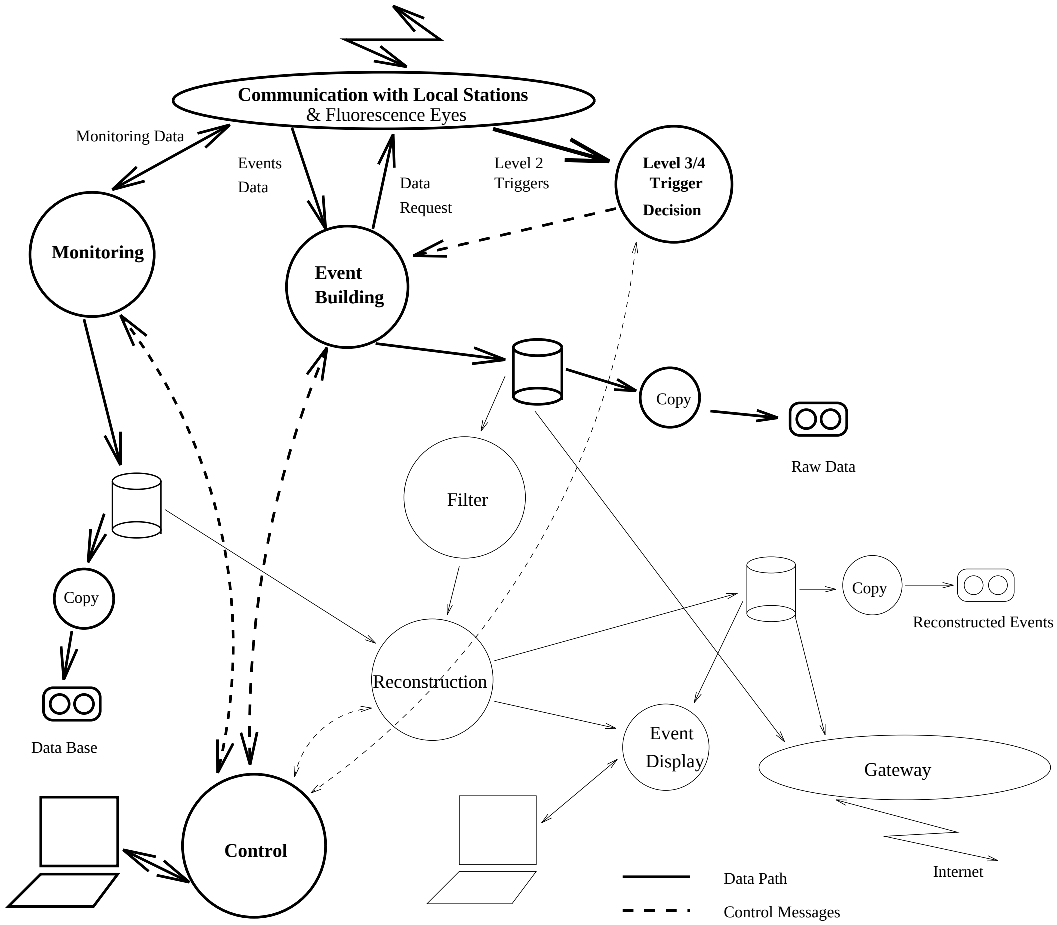

Figure 8.14: Block diagram of central station data flow.

8.8 Trigger

We conclude this chapter with a discussion of the trigger, which must select interesting physics events from a large background of random coincidences of small air showers. In order to accomplish this we have adopted a multi-level trigger design which is distributed both between hardware and software, and geographically. Table 8.1 summarizes the trigger levels and associated rates at each level for the ground array. This rate is compared to the rate of $\geq 10^{19}$eV showers.

The level 1 trigger selects potentially interesting data from a background of several kHz of single muons, small showers, and PMT dark noise. This trigger is implemented in hardware on the front end board (see Section 8.2).

As it will take about 1 ms to transfer the FADC waveforms to the station controller memory, negligible ($< 1\%$) dead time will be incurred if trigger rate has been reduced to less than 100 Hz by the level 1 trigger and the FADC board is double buffered. The station controller then applies software algorithms to analyze the FADC waveforms to refine the level 1 trigger and produce the level 2 trigger. The level 2 trigger rate will be less than 20

Trigger level	Where	How	Condition	Rate	Triggers /Shower	Randoms (%)
0	Station	Hardware	Single μ + small showers + PMT noise	$\approx$ 5 kHz per station	8×10^{10}	100
1	Station	Hardware	Limited/integrated energy	100 Hz per station	1.6×10^9	100
2	Station	Software	Level 1 + FADC info	20 Hz per station	3×10^8	100
3	Central	Software	3 fold level 2 in 20μsec window	0.2 Hz	1.5×10^3	100
4a	Central	Software	Level 3 + 1 additional level 2 in 1μsec window	15 / day	1.5	33
4	Central	Software	Level 3 + 2 additional level 2 in 1μsec window	10 / day	1	0.7

Table 8.1: Ground array trigger summary. In the last 2 columns the trigger rate is compared to the rate of $\geq 10^{19}$eV showers. Level 3-4 rates are for the whole array.

Hz. This rate limit is imposed by the communications system, and is high enough that we will not be cutting out interesting physics events. Level 2 triggers are transmitted to the control center for further trigger processing.

The level 3 trigger nominally correlates information from nearby stations (e.g. 3 in any 7 station cell within a 20 μs coincidence window) or nearby stations plus an external level 2 trigger (e.g. optical fluorescence plus two stations). However, since the trigger is formed centrally in software, other variations can be implemented which are more sensitive to possible exotic phenomena. The level 3 trigger contains the minimum information required to reconstruct a shower. For example, a shower direction cannot be reconstructed from less than three ground stations. However, the background from random coincidences of single muons and small showers (estimated using binomial statistics and the 20 Hz level 2 rate constraint) is 1500 times larger than the rate of $\geq 10^{19}$eV showers. Thus, level 3 triggers are in general not useful for physics analysis without additional constraints. Nevertheless, due to their low rate (0.2 Hz per site), all level 3 triggers will be archived and made available for further offline analysis. Once a level 3 trigger is formed the ground array and the fluorescence eyes are informed and the data from level 2 triggered stations is read out.

One should note that there is a high degree of symmetry in the trigger formation between the ground arrays and the fluorescence eyes. Level 2 and higher triggers from all components are incorporated into all higher level trigger decisions. Thus there are variations of the level

3 and higher triggers which contain only ground array data, contain a mixture of ground array and fluourscence detector data, or contain only data from 1 or more fluorescence eyes. For clarity, table 8.1 includes only purely ground array triggers.

The level 4 trigger further reduces the background contamination. This will be used as the basis for the on-line analysis of data. For the ground array a level 3 trigger, plus 2 additional level 2 triggers (from the 19 station cell centered on the level 3 trigger) occurring within a 1 μs window around the shower front defined by the level 3 trigger will reduce the background to less than 1% of the triggers. In Table 8.1 we note (in line 4a) that one additional level 2 trigger is not sufficient in this scenario; 1/3 of the triggers would be random coincidences.