

Chapter 6

Electronics, data acquisition and database

6.1 Overview

This Chapter describes the systems which digitize the electrical signals from the photodetectors and select events of interest for mass storage and physics analysis. The database is also described in this chapter. It maintains permanent records of the construction, installation and history of the experiment. It is used by all other subsystems and is essential to the offline data analysis.

In addition to digitizing the signals from the photodetectors, the electronics systems supply the high voltages to the photodetectors and provide monitor and control functions. Information from the monitoring systems is regularly recorded in the database to provide a history of the experiment; faults and errors detected by the monitoring systems are logged to the database as and when they occur. The configuration of the database and the types of information recorded are discussed further in Section 6.4.5. Many of the monitoring functions are provided by processors embedded in the main electronics systems and are described with those systems.

The principal function of the electronics is to digitize and record the signals from the photodetectors; the following steps are involved:

- The time and the charge of each pulse from the photodetectors exceeding a (programmable) threshold of 0.3 photoelectrons is measured and digitized as a hit by the front end units.
- The hit information is time-ordered and passed from the front end units to one of several hub crates.
- The hubs combine the data streams from the front end units and pass them to an interface crate.
- The single interface crate receives all the hub data streams and passes them to the central system and trigger farm.

- The central system and the trigger farm together time-order the entire data stream and divide it into partially overlapping time blocks.
- The processors in the trigger farm examine the data in units of single time blocks and apply a software trigger algorithm to separate the interesting events from noise.
- Finally hits from real events are passed to the data acquisition system which provides data storage, run control and user interfaces.

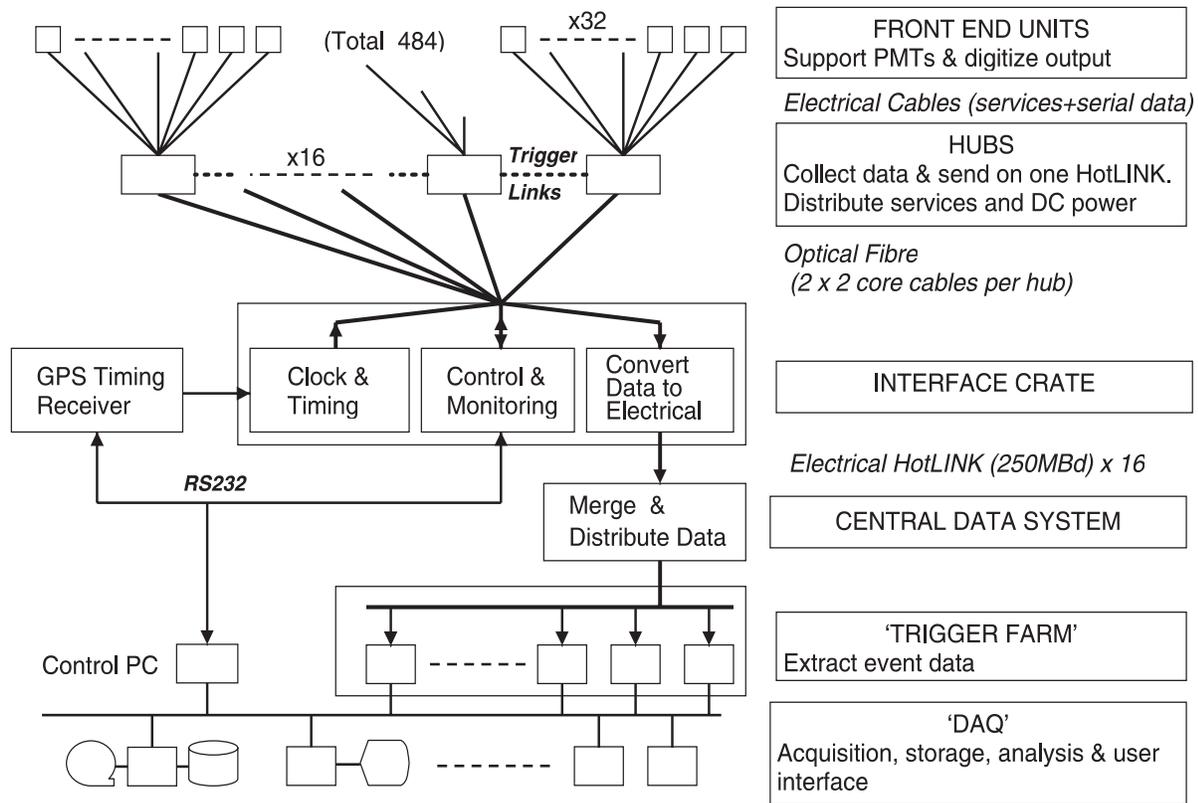


Figure 6.1: Schematic diagram of the principal components and organization of the far detector electronics system.

Figure 6.1 shows the principal components of the far detector electronics system. The near detector system has the same organization but is smaller. The architecture and rate-handling capabilities of the far detector system have been designed to allow the eventual addition of a third supermodule with minimum reorganization.

6.1.1 Physics requirements

The primary physics measurements in the MINOS experiment rely upon the:

1. measurement of the length of events to differentiate between NC and CC neutrino interactions;

2. reconstruction of long muon tracks for momentum measurement by range or magnetic deflection;
3. differentiation between electromagnetic and hadronic showers by their characteristic longitudinal and transverse energy deposition profiles;
4. calorimetric measurement of hadronic and electromagnetic energy; and,
5. recognition of certain simple event topologies characteristic of τ decays.

The provision of timing measurements with sufficient accuracy to determine the direction of a track will allow an extension of the physics capabilities of the far detector to studies of upward-going muons from neutrino interactions in the surrounding rock and atmospheric neutrino interactions in the detector.

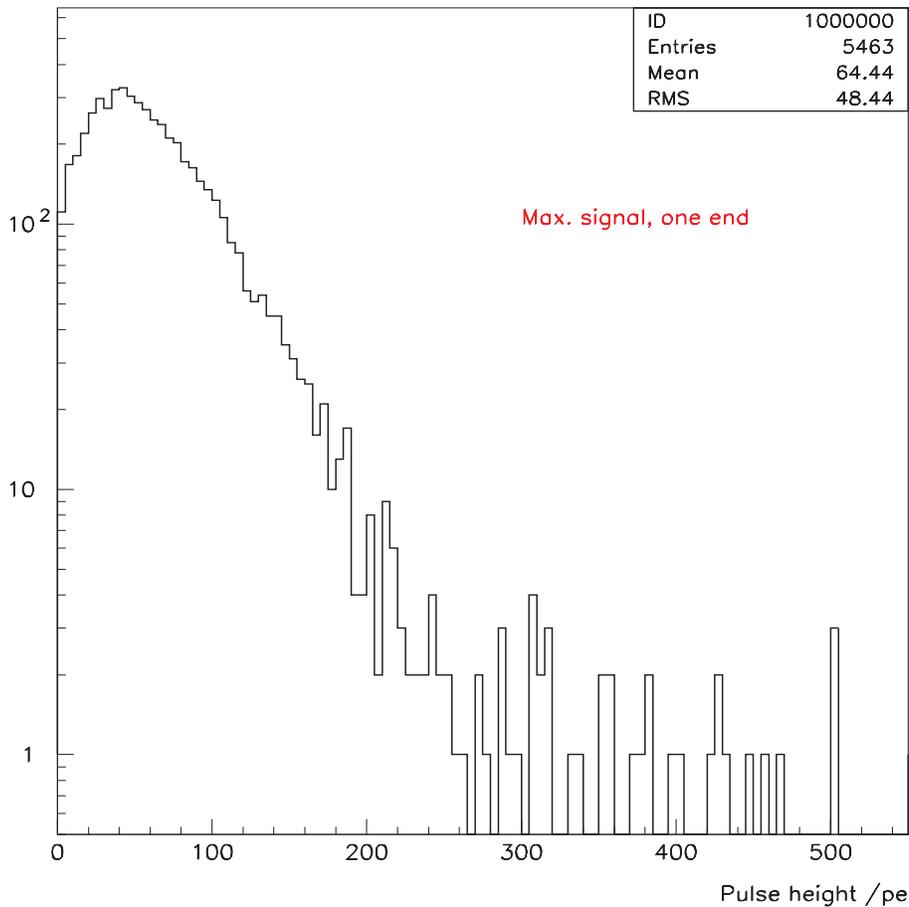


Figure 6.2: Distributions of the maximum pulse heights (in terms of photoelectrons) for interactions of electron neutrinos with the energy spectrum of the highest energy NuMI beam.

The calorimetric energy measurements and the separation of hadronic and electromagnetic showers rely upon the proportionality of light-yield to deposited energy. Figure 6.2 shows distributions of the maximum pulse heights, in terms of photoelectrons, that would

be recorded for electron neutrinos with the same energy spectrum as the muon neutrinos in the highest energy neutrino beam. The maximum signal to be recorded to achieve electromagnetic calorimetry over the entire energy range is equivalent to a 200 – 250 photoelectron signal from the photodetectors. The maximum signal to be recorded by the electronics has been set at the equivalent of 500 photoelectrons for a photodetector channel with normal gain.

Because it is important to maintain high efficiency, the effective thresholds must be set at the equivalent of 0.3 photoelectrons (pe) or less for all channels. The front-end noise must therefore be kept to a small fraction of the single photoelectron signal for a low gain channel. The electronics must have sufficient dynamic range and low noise to handle the maximum signals on normal channels and to resolve the smallest signals of interest from low gain photodetector elements [1].

The event rate of ~ 150 interactions per 1 ms spill in the entire near detector is relatively high although the interactions are distributed uniformly along the detector. The highest instantaneous single-channel rate will be ~ 10 kHz and shaping and dead times must therefore be sufficiently short that the pileup of hits and deadtime losses on a single channel are negligible.

The times of hits must be recorded with sufficient accuracy ($< 1 \mu\text{s}$) to enable events which overlap in one or both views to be resolved using timing information. The ability to flag hits potentially corrupted by pileup from a previous hit on the same channel is desirable, as is the ability to examine the history of the detector for a few microseconds before the time of an event.

Cosmic ray muons will be the main method of cross-calibrating the energy scales of the near and far detectors. It is therefore important that both detectors are able to record cosmic ray muon tracks. The flux of cosmic ray muons at the near detector (~ 300 Hz)[2] is expected to be ample for this purpose; the relatively low muon flux at the far detector (~ 1 Hz) means that the detector must be sensitive for a high fraction of real time.

Radioactive sources will be used as a secondary method of calibration and for setting up detector elements, operating conditions and photodetector gains. The electronics must be able to measure and record the DC current from the photodetectors produced by the sources. Current measurement is also a useful diagnostic tool for the detection of faulty channels.

The absolute rate of beam neutrino interactions at the far detector is extremely low (about 1 mHz) and to control possible backgrounds from cosmic ray interactions it is important to associate an event with the 1 ms spill of the Main Injector. This can be achieved by recording the absolute times of events to $\ll 1$ ms and correlating the event time with the times of machine spills recorded at Fermilab.

Since the thresholds must be low to ensure high efficiency, the singles rates in the two detectors will be determined by radioactivity in the cavern rock and in the detector steel, and the dark-counts from the photodetectors which are predominantly at the level of a single photoelectron. The total counting rate above a threshold of 0.3 photoelectrons due to radioactivity is estimated to be 510 kHz for the far detector and 145 kHz for the near detector[3, 4, 5]. The photodetector dark-count rate has been measured to be ~ 1 Hz per mm^2 of photocathode although it is rather sensitive to temperature. Singles rates are shown in Table 6.1, assuming both the nominal dark-count rates and a worst case of ten times nominal. All the single-channel rates are very low; the higher rates are associated with

Dark count rate per mm ² of photocathode	1 Hz	10 Hz
Far detector maximum rate/channel	88 Hz	230 Hz
Far detector typical rate/channel	28 Hz	170 Hz
Far detector rate/plane	900 Hz	4400 Hz
Far detector total rate	0.9 MHz	4.3 MHz
Near detector maximum rate/channel	140 Hz	290 Hz
Near detector typical rate/channel	21 Hz	160 Hz
Near detector rate/plane	3.2 kHz	24 kHz
Near detector total rate	320 kHz	2.4 MHz

Table 6.1: Expected singles rates in the MINOS detectors due to radioactivity and photocathode dark counts for a threshold which detects 1 photoelectron. 8-fold optical multiplexing is assumed for the per-channel rates.

channels connected to strips at the outside of the detector. Because of the large number of channels the overall total rates are high. Since there is still some uncertainty about the dark-count rates, the electronics has been designed to handle a maximum rate of 20 Mhit/s for the far detector (distributed over three supermodules) and 10 Mhit/s for the near detector, a safety factor of four over the worst-case rates. The design of the back-end of the system can easily be scaled down, with a consequent saving in cost if the rates prove to be less.

An important consideration in the design of the electronics for the far detector is that the power dissipation is limited by the cooling power planned for the Soudan cavern. The design described here is conservative and the power consumption has been kept low (< 25 kW), with the consequent constraint on the speed of the electronics.

Because the experiment will rely upon a detailed comparison of the characteristics of events recorded in the near and far detectors the same electronics will be used for both detectors.

6.1.2 Architecture

The electronics uses a simple continuously-sensitive readout architecture. The times and amplitudes of all signals from the photodetectors above a pre-set threshold are digitized, ‘stamped’ with real-time and sent to processors in a ‘trigger farm’ where a software trigger is implemented. The trigger algorithms which have been studied select events by correlating hits first in time and then space. The use of processors, rather than a hardware trigger, allows great flexibility and event selections could, for example, be made based on the pulse height of hits. The event selected by the processors are then passed on to the data acquisition system for any further analysis or selection and subsequent mass storage.

Such a triggerless readout scheme has the advantage that a large degree of flexibility is retained for the eventual event selection algorithm and does not require the complexities of programmable trigger logic or the distribution of fast trigger signals over a physically large distance. Data links with the necessary bandwidth and the processors necessary to implement such a triggerless system are currently available commercially. They are relatively

Component	Near detector	Far detector	Total	
Front end units	212	484	696	One per 48 photodetector pixels
Hub crates	7	16	23	
Interface crate	1	1	2	
Central system crates	3	4	7	
Trigger farm crate	1	1	2	
Trigger processors	≤ 5	≤ 10	≤ 15	
GPS timing systems	2	1	3	One to record real time of MI beam extraction at Fermilab

Table 6.2: Component numbers for the MINOS electronics systems for the two detectors.

inexpensive and since this is a rapidly developing area of technology it is likely that future commercial developments will reduce costs in the next few years.

The principal components of the electronics system are shown in Figure 6.1. They are:

1. front end units;
2. hub crates;
3. interface crate;
4. trigger farm hardware;
5. trigger farm and event selection software and,
6. GPS timing system.

The numbers of each component are given in Table 6.2; their functions are now described in detail.

6.1.3 Front end units

A front end unit is a self-contained, screened assembly which houses the electronics to digitize the signals from three 16-channel PMTs, as well as control and monitoring electronics, and power supplies. Thirty two front end units are connected to a hub crate, each by a single cable which provides power, timing and control, and carries the digitized output data in serial form.

The electronics for each channel provides amplification and pulse shaping, with a discriminator controlling a track-and-hold to store the peak of the pulse. As each pulse is detected its time is recorded, and it is entered in a queue to have its peak value digitized. The outputs of 16 channels are multiplexed on to a single ADC which digitizes the stored peaks, taking them in order of their arrival.

One 16-channel sub-assembly serves one 16-channel PMT. Three such sub-units are included in each front end unit. Each outputs data on any pulses at its inputs as a time-ordered stream. The shared part of the front end unit electronics merges these three streams into one, also time ordered, and outputs it over a serial link. The data transmitted for each hit consists of the channel address, the time stamp, the pulse amplitude and status flags which indicate the origin of the hit (data, pedestal, calibration) and possible corruption by pileup.

Other parts of the front end electronics provide essential support functions:

- Calibration: charge injection and phototube current measurement.
- Control and monitoring: make settings and report operating conditions.
- Clock and timing: provide synchronization with the central system.
- Power supply: convert and regulate supplies; supply HV to the photodetectors.

A major consideration in the design has been the need for immunity to electromagnetic interference (e.g., from the welding while detector planes are being assembled). This is especially true because of the wide dynamic range being called for, and the need to detect single photoelectrons. The signals from the photodetectors are low level signals. Inclusion of the photodetectors and all their associated low level electronics in a single well-screened box is of major importance. So too is connection through a single electrical cable, with a consequent absence of ground loops. Potential sources of noise and the strategy to reduce or eliminate them are further discussed in Section 6.1.9

A second important design feature is ease of maintenance. Packaging a set of detectors and their electronics as a single easily exchanged unit is advantageous. A possible bad channel needs only the one unit to be replaced with a new one which has been given a complete functional test in the lab. It is also relatively easy to make such a unit so it can be 'hot swapped' (safely exchanged without powering down any of the system or affecting its operation).

A constraint on the design of the front end units has been the ultimate limitation on power imposed by the problem of cooling the far detector cavern, the need to avoid unduly heating the photocathodes, the desirability of not introducing the further cost and complication of an elaborate cooling system for the electronics and the possibility of future detector upgrades such as a change in the degree of optical multiplexing. The total power dissipation of the far detector electronics has therefore been assumed to be limited to 25 kW for an eventual detector composed of three supermodules with four-fold optical multiplexing. Assuming 16 kW is available for the front end units (leaving 9 kW for the hubs, computers, and ancillary electronics), if the efficiency of the power supplies and distribution is 75% (after any local conversion and regulation) 4 kW is available for the front end units of each supermodule. Allowance for future expansion to three supermodules with optical multiplexing reduced to

four-fold therefore requires a maximum of 170 mW/channel; 200 mW/channel is regarded as the maximum tolerable power dissipation. The present design complies with these figures, including the allowance for possible future expansion. Power dissipation in each front end unit is ~ 10 W including losses in local regulation.

Each front end unit will serve three 16-channel Hamamatsu PMTs, giving a total of 48 channels per unit. This is a compromise between sharing the overheads (monitoring, power supplies, cables and connectors, and mechanics) over an adequate number of channels, without requiring an inconveniently large number of fibers to be concentrated into a single unit. It is also well matched to the far detector, with one unit serving one end of two planes.

The design of the amplitude digitization system is based on shaping the signal with a ~ 400 ns peaking time (which determines the extent of pileup errors), and a worst case dead time of 6 μ s for the last channel to be read after all channels in a group of 16 have been hit simultaneously. These times have been chosen to give a similar effective speed to that possible in a simple triggered system. (Any such system, which does not pipeline the signals in analog form or after unconditional fast digitization, would be limited by the physical extent of the far detector.)

The dynamic range of the ADC system is nominally 14-bit, with full scale corresponding to 500 photoelectrons on high gain pixels. Assuming a 2.5:1 gain spread between pixels and no individual channel trimming, the 1 pe peak from a low gain pixel lies at 13 counts above the pedestal. This is comfortably good resolution. A true 14-bit system might be expected to have a noise level lower than the quantization noise from the ADC, i.e. ≤ 0.3 counts rms. This is clearly not needed for MINOS, especially as detection of pulses is by a hardware discriminator not the ADC. Around 1 count rms noise from the ADC is good enough. Hence, the design is only nominally 14-bit.

The A-D conversion accuracy required is modest: 5% over most of the range. Around 1% is, however, desirable for calibration purposes, which needs some provision for averaging out ADC non-linearities at the low end of the range.

It is crucial to achieve low noise at the discriminator since it is required to detect single photoelectron signals efficiently but not to trigger at an excessively high rate on electronic noise. The finite width ($\sim 100\%$ FWHM) of the 1 pe pulse height distribution from the PMTs means that the discriminator threshold must be set at the equivalent of 0.3 – 0.5 pe for efficient detection of single photoelectrons. The rate that the discriminator fires on electronics noise must be much less than the dark count rate of ~ 1 Hz. The rms noise at the discriminator must therefore be the equivalent of ~ 1 ADC count to give a low enough noise background rate for the lowest gain pixels.

The timing capability of the detector is ultimately limited by the scintillator and wavelength shifting fiber to around 10 ns rms when a single photoelectron is detected although this improves for larger pulses. A timing resolution of ~ 2 ns rms has been demonstrated for signals of > 10 photoelectrons. These factors guide the design of the signal channel up to the discriminator output. While the PMTs are very fast, making the electronics to match conflicts with the need for a wide dynamic range, low power consumption and moderate cost. The present design is conservative and it is expected that it can be improved. There will be variation between channels, due to component variations, of up to 15 ns from the mean. This will be stable and easily corrected using the calibration system. There will also

be some dependence on pulse height; total ‘timing walk’ should be under 10 ns, and could be corrected offline when maximum accuracy is needed. Time is measured on each channel in units of 5 ns which seems the best that can be assured if Field Programmable Gate Arrays (FPGAs) are used; further work, and faster devices, may better this.

The data transfer rate from a front end unit will be ~ 50 khits/s, which is much lower than the peak rate at which it can digitize. This rate is also higher than the average rate per unit of 28 khits/s derived from the very conservative system design criterion of 20 Mbits/s for the whole far detector. The peak rate from the near detector will be high during spills with around 250 hits expected from a unit in one spill. Sufficient buffering is provided to smooth out these bursts of data.

6.1.4 Hubs and interface crate

The hubs provide front end units with DC power, system-wide clock and timing, and connection to the control and monitoring system. The data are passed via the interface crate to the trigger farm. For convenience and economy the data are multiplexed into a single serial stream.

Hubs are placed relatively close to the front end units. Each hub serves up to 32 front end units and is connected by a small number of fiber optic links to a single interface crate, providing a single point of connection to the whole system.

The interface crate connects to a GPS timing receiver, which is the source of clock and timing, and also to the computer which manages the control and monitoring. The interface crate outputs the data from all the front end units into the central trigger farm and data handling system. A particular consideration in selecting this organization is to limit the effect of faults, and to permit their easy localization and repair. It also complies with the requirement to avoid ground loops and their risk of introducing interference.

Data rates. The data rate is assumed to be distributed evenly across the hubs. Each hit will be represented by 6 bytes of data, the time stamp and the channel number both being extended later in the central system to give 8 bytes of data per hit. The link from front end unit to hub must handle 165 kByte/s (or, since it is serial, 1.3 Mbit/s). Serial operation at about 2.5 MHz allows a comfortable margin while at the same time limiting the rate possible from a malfunctioning front end unit so that the system cannot be overloaded. Each hub must be able to send 5 MByte/s of data to the interface crate. The links chosen are capable of 25 MByte/s, giving an ample margin.

The near detector system is designed for the same rate from each front end unit as the far detector. The figure is actually slightly higher, as there is more of a contribution from real events, but the difference is not significant. The same system is thus adequate, with a smaller number of front end units, and the number of hubs reduced proportionately from the far detector.

Control and monitoring. The system provides control and monitoring connections to the front end units. The performance requirements for this are minimal; only low speed is

required. Simplicity of development, ease of fault-finding, and ability to operate without interacting with data taking, are more important criteria.

For this reason control and monitoring is based on 'RS232' type serial links, communicating with microcontrollers in various units throughout the system. Those in the hubs provide for local control and monitoring, and by means of address commands, communications with those in the front ends. The system is entirely independent of the timing and data systems. The monitoring in the front end units also provides a readout of the PMT integrated current which is used for radioactive source calibrations.

6.1.5 Central data system and trigger farm

The purpose of the central data system is to receive data from the front end units (via the interface crate) and to present it to processors in the trigger farm for the selection of events of interest by a software trigger algorithm. The data from the front end units consist of a series of short hit records, one for each digitized photodetector pulse. Each hit record consists of a channel address, a time stamp, the pulse amplitude and status flags. The time stamps are used to sort the stream of hit records from all detector channels into a single time-blocked and time-ordered data stream which is passed to the trigger farm processors.

The nominal hit rate from all sources is ~ 1 Mhit/s but, since there are still some uncertainties associated with the photodetector dark-count rate, the system is very conservatively designed to handle an absolute maximum rate of 20 Mhit/s which also allows for the addition of a third supermodule and a reduction in the degree of optical multiplexing.

The interface crate and the central data system will be adjacent. The serial links between these crates will use short twisted-pair electrical cables. Each link carries a stream of data from a single hub, serving 32 front end units, with each front end unit being uniquely identified in the outgoing data stream.

The architecture of the central data system has been designed to support an eventual three supermodule detector. It is shown in Figure 6.3. It consists of up to six VME crates (3 at the near detector) containing receiver cards and sorter cards. The receiver cards receive data from the interface crate on up to 24 serial links (7 at the near detector). Only four VME crates and 16 serial links are required for two supermodules. The hit data from each link is unpacked into 32 separate data blocks, corresponding to the front end units served by one support module. The receiver card then extends and completes the channel number and time stamp information received from the front end units into 8 bytes of hit information.

The sorter card comprises a VME single-board computer whose function is to assemble the data from the receiver cards into time-ordered trigger time blocks and pass the time-blocked data to the trigger farm.

The output of this VME system is connected to the trigger farm by two 132 Mbyte/s networks. The data rate on each network is not expected to exceed 80 Mbyte/s, leaving considerable reserve capacity over the already conservative design rate. The trigger farm and the central system communicate through the control network which is of the same type as the data network.

Two time block lengths will be used. In the receiver card input stage, the 'front end' time blocks are matched to the front end time-stamp counter wrap-around period of 1.3 ms (18 bits with 5 ns resolution). In the receiver card output stage the front end time blocks

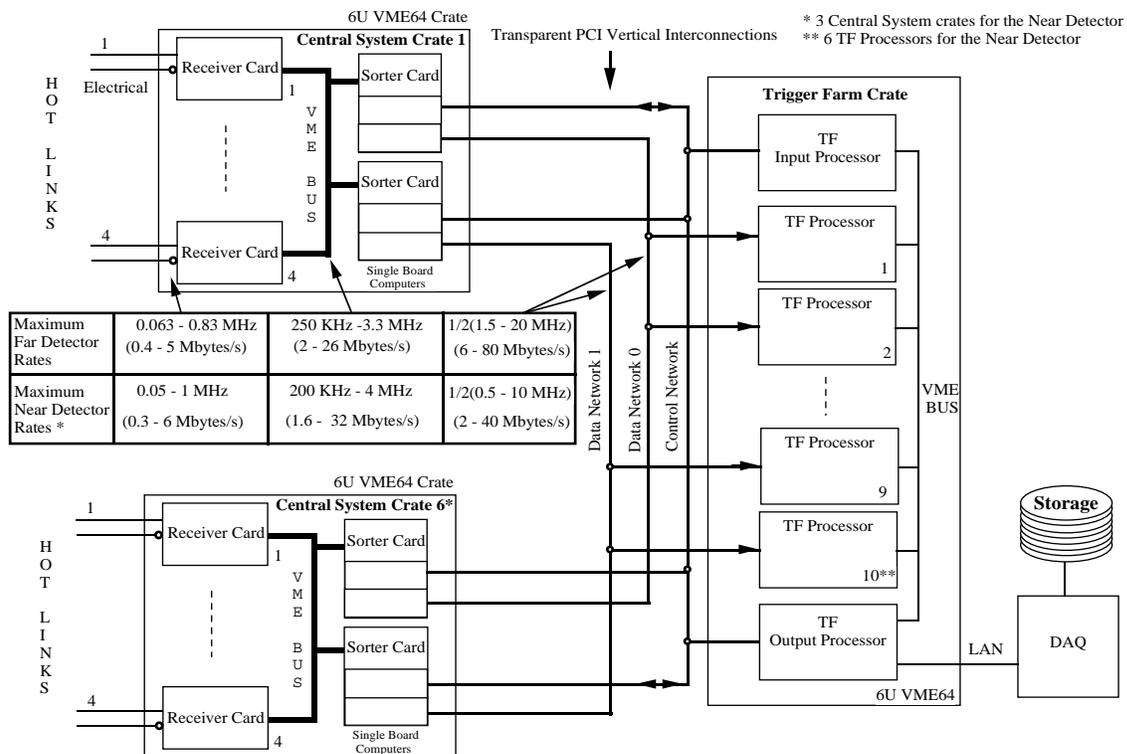


Figure 6.3: Central data system and trigger farm organization.

are assembled in sequence to form ‘trigger’ time blocks. The amount of memory allocated to the trigger time blocks is sufficient to contain the hits from 38 front end time blocks (50 ms) at the highest data rate. These trigger time blocks are then used by all subsequent stages in the system. The number of front end time blocks merged to form a trigger time block will be variable, controlled by software, to give flexibility in varying the average amount of data in the trigger time block in order to optimize system efficiency. Adjacent trigger time blocks will overlap by the period of one front end time block to ensure that there is no inefficiency for events straddling the boundary of blocks. Trigger time blocks are passed over the two data networks to the trigger farm, alternating odd and even numbered blocks.

Processing in the trigger farm is divided into two stages. The first stage is to merge-sort each set of trigger time blocks from the central system crates into a single block of time ordered hits. In the second stage a software trigger is applied to the time ordered data to select events of interest and pass them on to the DAQ system.

The trigger farm consists of a VME crate containing VME single-board computers with dual PCI Mezzanine Cards (PMC) slots for processing and sequencing. Additionally there are high performance PMC based network cards for data transfer. Two single board computers, the input processor and output processor, are dedicated to sequencing. The trigger processing is run on an array of single board computers (trigger farm processors) with a maximum capacity of up to 18 in the crate, although it is anticipated that no more than 10 such processors will be necessary.

Three networks are used between the central system and the trigger crate. The control network allows the farm input processor to identify the time blocks in the sorter cards and

sequence their transfer to allocated processors in the farm. The two remaining networks are used to transfer the data between the central system and the trigger farm. The data networks perform only one-directional block transfers which keeps the transfers efficient. The control network is random access which is less efficient but runs at the rate of the time blocks which is much lower.

The control network connects the input processor to each sorter card. The input processor uses the control network to determine when central system time blocks are available, where they are located in the sorter card's memory and to signal to the sorter cards when the time block data has been transferred to the farm. The data networks are used to transfer time blocks from the sorter cards to the farm trigger processors under the control of the input processor.

The farm trigger processors are divided into two equal groups, one for each of the two data networks. Each data network and trigger farm processor group handles either odd or even time blocks. The trigger farm receives time blocks for the same time period from the central system which are processed by the first free processor available within the appropriate odd or even trigger farm processor group. The input processor is responsible for coordinating the assignment and transfer of time blocks from the sorter cards to the farm processor groups. The output processor is responsible for coordinating the transfer of processed time blocks to the DAQ system. The output processor buffers and reassembles processed time blocks back into strict time order before sending them on to the DAQ system.

6.1.6 Trigger farm and event selection software

The functions of the software running in the trigger farm include:

- data transfer and scheduling software for the input processor;
- triggering software and trigger efficiency monitoring in the trigger processors;
- event collection, sorting and monitoring in the output processor;
- control and data transfer software for the interface to the DAQ system, and
- error detection and recovery.

The event selection (trigger) software is based on time-ordered blocks of data; a block of data (a trigger time block) is DMA transferred to a free processor in the trigger farm under the control of the input processor. These blocks are time synchronized and represent a complete time section of readout from the detector. The input processor ensures that no triggers are lost due to boundary effects by overlapping the time sections passed to the trigger processors; any identical triggers created by this overlap are identified and eliminated by the output processor. The length of the trigger time blocks into which the data are divided is programmable and will be at least 50 ms (depending on the actual singles rate); the blocks will overlap by 1.3 ms. The task of a trigger processor is to:

1. merge the blocks of data into a single time-ordered stream;
2. perform some basic monitoring of raw singles rates (more sophisticated monitoring using triggered events will be performed in the output processor and/or the data acquisition processor);
3. identify events within this time block of hits by applying a trigger algorithm;
4. pass triggered events to the trigger farm output processor for sorting and passing to the data acquisition system and
5. inform the control processor that it is free to accept another block of data.

The trigger algorithm must reduce the data rate into the data acquisition system by removing noise and background from the data stream in real time with minimal loss of events of physics interest. The dominant source of background in the data stream is expected to be the singles rate from radioactivity and photocathode noise, both of which should be relatively easily discriminated from physics events. Three types of information can be used to achieve this:

1. Time structure. Physics events will contain hits that are strongly correlated in time, i.e., the hits will occur within a narrow time gate defined by the propagation time of tracks through the detector and light propagation time in the scintillator.
2. Spatial structure. Physics events will contain spatially clustered and/or contiguously strung hits (showers/tracks) that span a number of detector planes.
3. Pulse height. Photocathode noise will have a lower pulse height (1 pe) than hits from a minimum ionizing particle. Imposing a threshold cut on hits allowed to contribute to the trigger could be a fast and powerful discriminant against this noise. However, in order to maintain independence of design parameters such as phototube gains and scintillator light yield at this stage, no use is made of pulse height in the trigger algorithm described here although a trigger based on summed pulse height could easily be implemented in the trigger processors and OR'ed with other triggers such as the one described below.

Since the trigger processors are required to operate in real time, a simple and fast trigger algorithm is required. The algorithm operates in two stages:

Level 0. A search is made through the readout in time order until a minimum of 3 hits are found to occur within any 50 ns window. This window is wide enough to accommodate expected fluctuations in the relative times of hits before the final timing calibration is applied. Simulation studies of neutrino events in the detector have shown that losses to neutrino events from this requirement are less than 0.5%; these losses consist entirely of neutral current events with very few hits. If this condition is satisfied the second, more stringent, level 1 criterion is tested.

Level 1. Require that at least M planes out of any group of N contiguous planes of the detector have at least one channel hit in the 250 ns following the earliest hit in the above 50 ns gate. A 250 ns window is chosen here as a conservative event size. If the M/N condition is not satisfied the trigger search continues where it left off in level 0. If the condition is satisfied an event trigger is generated and the event is passed to the farm output processor.

At this stage an event is defined as all hits contained within the 250 ns window plus any hits that occurred in a preceding (programmable) time window of $T \mu s$. Such a detector history allows the possibility that the event has been corrupted by pileup, or that hits have been lost due to single-channel deadtime, to be assessed. Although this means that the first $T \mu s$ and last 250 ns of each trigger time block would not be scanned for a trigger, since it could not provide the full readout time span, this generates no trigger inefficiency since the 1.3 ms overlap ensures that triggers in these time regions are found in the preceding or following buffer by another processor. The performance of the trigger algorithm presented here has been timed with a trigger scan of these regions included.

Studies, using the simulation software described in Section 9.2, have shown that this 2-level trigger with $M = 4$ and $N = 5$ excludes a negligible fraction of charged current neutrino events and around 11% of neutral current events, these being events with very few detector hits. Changes to the (level 1) M out of N plane trigger are trivial to implement in this scheme and have no significant impact on the algorithm timings. For example if, following further trigger studies, the 4/5 plane trigger used here were required to be a 3/5, 4/6, ... plane trigger to accept a higher fraction of neutral current events, it could be accommodated easily. The first three columns of Table 6.3 show the level 0 and level 1 noise trigger rates for a 3/5 plane trigger and 4/5 plane trigger as a function of the detector singles rate (assumed to be distributed over three supermodules). There is no difference in processing requirements for the different trigger configurations.

Monitoring of the pre-trigger singles rates will be performed in the trigger processor farm on a channel basis. Channel rates exceeding an acceptable threshold will be transmitted to the output processor and thence on to the online system which will take the appropriate action. Simple monitoring of this kind has very little impact on the processing power required in the farm which is dominated by the data merging and trigger decision.

Prototype trigger processing code for the above algorithm was written in C to estimate the maximum number of processors the trigger farm will require to keep pace with potential noise rates. Noise data, including an assumed 5 kHz double-Compton rate from radioactive decays, were generated at a number of rates up to and including a worst case of 20 MHz. Timing tests were performed on a 300 MHz ALPHA RISC processor (Alpha Server 1000A 5/300, Alpha 21164 chip). The results are shown in Table 6.3. The timings include the histogramming of raw channel singles rates for channel monitoring.

The processing requirements are well within practical limits. Faster processors than a 300 MHz Alpha are available today and improvement in processor speed is likely over the next few years; the provision of sufficient processing power in the trigger farm is not expected to be a problem.

Triggered events are passed to the output processor when they are identified by the trigger processors. They can arrive out of chronological order by an amount defined by the

Singles Rate (MHz)	Trigger Rate (Level 0, kHz)	Trigger Rate (Level 1, Hz) 3/5 planes	Trigger Rate (Level 1, Hz) 4/5 planes	Max Processors
2	4	6	< 1	1
6	70	26	< 1	3
10	290	130	< 1	5
16	1100	750	< 3	8
20	2000	1500	< 5	10

Table 6.3: Trigger rates and processor requirements of the trigger farm. The processing power estimates were obtained using an Alpha Server 1000A 5/300, Alpha 21164 chip processor.

maximum length of time it takes a processor to trigger-analyze one trigger time block. In addition, because trigger time blocks overlap, some events will be found twice. The output processor is required to buffer the events it receives in a rolling buffer, time order them and remove or merge (as appropriate) event duplicates and overlaps before passing them on to the DAQ system.

The event rate to the output processor from beam and cosmic ray muon events will be relatively low. In the far detector the event rate will be dominated by around 1 Hz of cosmic ray muon events (2.4 kByte/s uncompressed). The near detector rate is more complex but will average at approximately 16 Hz of neutrino interactions, 8 Hz muons from ν interactions upstream of the target region and up to 270 Hz cosmic ray muons (virtually all out of spill), giving a total rate of just under 300 Hz equivalent to a data rate of about 240 kByte/s uncompressed. Cosmic ray muon events are essential for cross-calibration of the near and far detectors; the cosmic muon rate at the near detector is quite high and can be reduced by randomly vetoing out-of-spill triggers as part of the near detector trigger algorithm.

6.1.7 GPS timing system

For the analysis of the data it is necessary to know whether events occurred during the 1 ms period of a beam spill. Because of the remote location of the far detector the most practical way to determine this is to log the absolute times of events and spills, and correlate these offline. Each event will therefore be stamped with the real time obtained from the GPS clock.

The clock and timing of the entire system is derived from a GPS timing receiver connected to the interface crate. It is assumed to provide a 10 MHz clock, and 1 Hz pulse outputs, with an RS232 interface for status and time of day information. These two signals are multiplexed into a single signal suitable for distribution to the hubs over a single optical fiber. There they are converted to electrical signals, reformatted, and fanned out to the front end units. Conversion to a different clock frequency is made with Phase Locked Loop (PLL) clock multiplier chips. The design has to maintain minimum timing skew between front end units and therefore requires fast fiber optic transceivers, and ECL electrical transmission (with transformer coupling for common mode noise immunity). It will be easy to check the

equality of fiber lengths using optical time domain reflectometry (TDR) from the interface crate; this might be desirable after repair or replacement. The short hub-to-front-end cables can reasonably be assumed to be constructed to equal length. The electrical propagation delays cannot be so easily checked after installation; the best way to ensure they are nearly equal will be to keep them short.

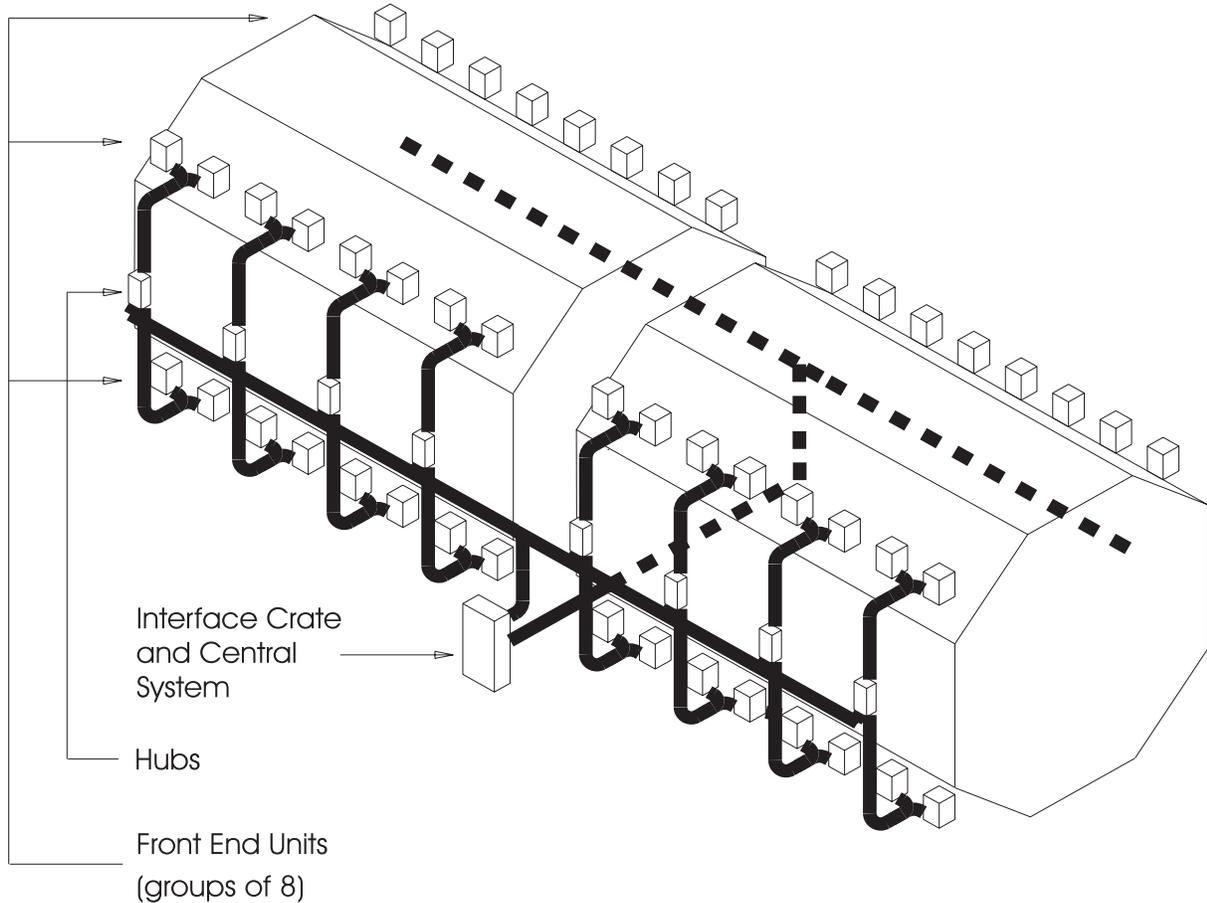


Figure 6.4: Layout of the electronics on the MINOS far detector.

6.1.8 Locations of the electronics components

The physical layout of the electronics for the far detector is shown in Figure 6.4. The front end units are arranged evenly along the length of the far detector, half on each side in two rows along the upper and lower corners. Each side is served by a single row of hubs at an intermediate height. The hubs are spaced about 3.5 m apart, with their front end units 4 to 5 m away. All hub-to-front-end cables are equal in length for equal clock propagation delay. Each hub serves up to 32 front end units, grouped as eight multiplexer box assemblies, two in each row. Some hubs will have slightly fewer units connected to maintain independence between supermodules.

The interface crate is placed in a central position along the detector. It is connected to each hub with fiber optic cables, again of equal length for equal delay. These cables will be

about 40 m long, depending on details of the installation. Excess length can be coiled up, or looped back in the trunking. Each hub has two cables, but they are slender (6 mm × 3 mm) and light: the total weight of each is under 1 kg.

The interface crate outputs one serial data link for each hub to the the central system. The central system crates and trigger farm crate will be housed in two enclosed racks adjacent to the interface crate.

The hubs will be installed on the lower walkway along each side of the far detector. The electrical cables from each hub to its front end units will run in bundles vertically, up and down, before fanning out. Each bundle of 16, assuming a 7 mm diameter, would be about 33 mm overall diameter. Support is required, in the form of trunking, or channel to which they can be tied.

The fiber optic cabling between the hubs and interface crate will run in trunking (or trays) along the length of the detector, and across the far (south) end. The interface crate and the central data system will be at that end on the same level. Fiber cabling is assumed to be with two core 'zip' type cable, as commonly used for patch cords. Individual cables are small and light, and can be factory terminated. They can then be unrolled into the trunking. A 30 mm bend radius is acceptable, so they can be looped back in 75 mm square trunking. Open cable trays might be used, but give less protection. The bulk of the fiber cabling is quite small: the cables along one side of the detector would form a bundle ~20 mm diameter. Excess length will be lost as coils or in the trunking, whichever is most convenient.

The layout of the electronics for the near detector will be similar but more dense since optical multiplexing is not used in the forward section of the detector and only four-fold multiplexing is used in the spectrometer section.

6.1.9 Minimizing potential noise sources

As discussed in Section 6.1.3, the front end electronics must digitize low-level signals. There are a number of potential electrical noise sources at both detector sites, including noise from welding operations, the magnet coils, the lights and high-current electrical apparatus such as pumps, the cranes and ventilation fans, all of which could introduce noise into the front-end of the system, either by radiation, induction or currents flowing in 'ground loops'. Care must be taken both in the design and the installation of the electronics and other detector components to ensure that the effects of these sources of noise are minimized.

The power distribution systems in the Soudan cavern and near detector hall are described elsewhere[6, 7]. The power for the ventilation system, the magnet coils, general use (e.g., welding) will be isolated from the power for the electronics by means of separate transformers. 'Quiet' AC power for the electronics will be supplied from shielded transformers. The power and ground lines from these transformers will run in separate conduits to circuit breakers and thence to receptacles mounted close to the electronics racks. The ground lines for this power will be isolated and run back to a common grounding point at the transformers. This 'quiet power' will be used only for electronics, computers and some test equipment, but not for high power devices such as pumps, large fans or welding apparatus.

Separate power distribution and grounding systems will be provided for the high-current electrical apparatus. The latter will consist of the grid of interconnected reinforcing bars embedded in the concrete floor slabs. Attachment points to the grid will be provided so

that short connections can be made between the apparatus and a low impedance path to the main power ground. The fluorescent lights will have electronic ballasts and filtering to minimize radiated noise, which is especially important as they start to reach the end of their lifetime.

The power supplies for the magnet coils are another potential source of noise. Power for the coils will be supplied from transformers used only for this purpose. Whilst it would be desirable to use linear power supplies for the magnets, the high current (40 kA) for the near detector coil, and efficiency and heat load considerations at the Soudan site, force the use of SCR controlled supplies at both locations. Sufficient filtering of the magnet supplies will be provided to eliminate the high frequency spikes from the SCRs, which are the main concern.

The DC power supplies for the front end units are located in the hub crates. Power-factor corrected, low-noise switching power supplies will be used in order not to introduce line-frequency harmonics on to the quiet power lines. If necessary, AC line conditioners can be used at the power receptacles.

The small size of the photodetectors and sensitive portion of the electronics makes them inefficient antennae at the frequencies of the noise radiated by electric motors and welders. As described in Section 6.1.3, all electrical connections to a front end unit are made by means of a single multi-conductor cable to the hub. This cable will have a ground jacket and run in a shielded cable tray. The front end units will be completely enclosed in aluminium boxes bolted directly to the MUX boxes which house the photomultiplier tubes. The aluminium boxes will be isolated from the steel detector structure to avoid ground loops. The PMTs and the front end electronics will therefore be in a Faraday cage whose only ground connection is back to the hubs and then to the quiet power ground.

During the installation and operation phases care will be taken to ensure that the quiet power grounds are not inadvertently connected to grounds used by other pieces of apparatus or to the detector steel plates and support structures. Tests of this isolation will be part of the installation procedure.

With proper attention paid to the control of the grounding, noise is not anticipated to be a problem. The CDF experiment has successfully used photomultiplier readout while welding was taking place nearby.

6.2 Requirements and performance criteria

The electronics and data acquisition systems for the MINOS detector will operate in a continuously-sensitive triggerless mode. The selection of events for physics analysis and calibration will be performed by software in one or more processors. The primary functions of the system are:

1. To record:
 - (a) neutrino events during the Main Injector spill and
 - (b) cosmic ray events in the time between Main Injector spills.

The recorded data consist of the amplitudes and absolute times (derived from a clock synchronized to the GPS system) of all pulses from the photodetectors above a pro-

grammable threshold.

2. To record calibration pulses from:
 - (a) the charge-injection system for calibration of the electronics system and
 - (b) the laser light calibration system for calibration of the photodetection system.
3. To record, or automatically subtract, ADC pedestal values.
4. To record the DC currents from individual photodetector pixels for:
 - (a) calibration of the combined scintillator and photodetector chain with radioactive sources and,
 - (b) diagnosing faulty photodetectors.
5. To monitor components of the readout and other systems and log fault conditions.
6. To log environmental conditions.

The performance requirements for the electronics systems are shown in Table 6.4.

6.3 Interfaces to other MINOS systems

The electronics and data acquisition systems interact with the other detector systems in three ways:

- connection to the photodetectors;
- location and installation, and
- the exchange of information with other systems.

There are also requirements for power and data links, at both the near and far detector sites.

6.3.1 Connections to the photodetectors

The connection between photodetectors and the electronics is made at the front end units. The constraints are: the front end units must be close to the MUX boxes, which are positioned to keep the clear fibers to the scintillator modules as short as possible; the front end electronics must be well shielded because the signals from the photodetectors are very low level; the power dissipated by the electronics should not heat the photocathodes. It is also desirable that any maintenance required by the electronics should not interfere with the optical connections to the photodetectors and vice versa.

The photodetector-electronics interface is achieved in a composite MUX box-photodetector-front end electronics unit which serves three photodetectors. The front end electronics receives signals from the 16 pixels of a single photodetector (see Chapter 5) via a single mixed-contact D-type connector mounted on an interface unit to the photodetector-MUX box. The

Parameter	Near detector	Far detector	Comments
Spill length	1 ms	1 ms	Neutrino spill The far detector must have > 80% duty cycle for cosmic ray muons out of spill for calibration.
Repetition time	≥ 1.9 s	≥ 1.9 s	Essential for calibration. Beam ν at near detector; radioactivity and PMT dark-current dominate at far detector.
Cosmic muon rate	~ 300 Hz	~ 1 Hz	
Max Instantaneous rate	≤ 10 kHz	< 200 Hz	
Number of channels	9,408	23,232	Assumes PMT gain of 10^6 . Resolve overlaps in near detector; not critical for far detector. For atmospheric ν and upward/downward muon separation in far detector. Pileup and losses in near detector; not critical for far detector.
Threshold	≤ 0.3 pe	≤ 0.3 pe	
Front end rms noise	≤ 0.05 pe	≤ 0.05 pe	
Charge measurement range	0.005 – 80 pC	0.005 – 80 pC	
Digitization accuracy	5%	5%	
Time resolution (1)	≤ 1 μ s	—	
Time resolution (2)	—	5 ns	
Single channel deadtime	≤ 5 μ s	10 μ s	
Additional capabilities			
DC current measurement	50 nA	50 nA	Radioactive calibration source (3×10^5 cps from a source of ~ 1 mCi).
Preamp charge injection Single-channel disable	0.005 – 80 pC	0.005 – 80 pC	Remove bad channels from readout.
GPS-based clock	< 1 μ s	< 1 μ s	Associate events with real-time of MI spill.

Table 6.4: Parameters for the MINOS electronics system.

interface unit is a simple assembly carrying a jumper cable from each photodetector base to the connectors at the rear. The cable allows a connection to be made without exerting force on the photodetector. Each interface unit has one back connector for each of the three photodetectors. The interface unit and the electronics housing are metal boxes to provide immunity from electrical noise pickup. The entire assembly is light-tight.

The high voltage power supplies for the photodetectors are mounted in the front end electronics boxes. These also contain a microprocessor for controlling and monitoring the front end systems and high voltage. Individually controllable high voltages are available to each of the three photodetectors. The high voltage also passes through the interface unit; since less than 1 kV is required the high voltage will use the same connector and jumper cable.

6.3.2 Location and installation

In this Section the location and installation requirements of the electronics systems are described; the installation of the detectors at Soudan and Fermilab is discussed in detail in Chapters 7 and 8.

Location. The physical location of the front end electronics is determined by the location of the photodetectors. Their location, in turn, is constrained by the maximum allowable length of the clear optical fibers. The MUX box-front end unit assemblies are located in a series of 64 short racks placed along the detector at the four 45° faces. Each rack has two shelves with each shelf supporting four MUX box-photodetector-front end electronics units and one rack services 24 photodetectors. At the two upper 45° faces of the detector, the racks are supported on a narrow walkway attached to the detector support structure beams. At the lower corners the racks are supported off the floor.

The rest of the electronics (the hubs and trigger farm) require very little floor space. The hub crates are placed on the intermediate height walkways along the detector, half on each side, and the central system and trigger crates are placed at the south end of the detector. Protected cable runs from the front end units to the hubs and from the hubs, along the detector to the central system are required. There must be room to store extra cable to ensure that all cables are equal in length.

Installation. The electronics is installed incrementally as the detector is assembled. The installation is largely independent of the work done on each steel-scintillator layer. The constraint is that the walkway which supports the front end racks cannot be installed until the steel planes below it are in place. Once the walkway is available, the racks and the MUX box-photodetector-front end electronics units can quickly be added. Each electronics front end box has a single cable which connects it to its hub crate. These are routed to minimize their length and to provide some physical protection.

Once a rack has been installed and connected it is tested, both with the laser calibration system and with the charge injection system to determine that everything is properly connected and operating.

For the commissioning of the electronics for the first supermodule of the far detector, whilst it is being installed, it will be necessary to locate the central system racks at the far end of the detector to avoid interference with mechanical installation work.

6.3.3 Exchanging information with other systems

The data from the trigger farm output processor will be passed to the DAQ system and then to an offline data storage system and, possibly, to offline processing computers. At the near detector site the data are sent directly to the Feynman Computing Center for storage, while at the far detector they are written to a permanent storage medium. A permanent copy of the data will be maintained at the Soudan site and copies distributed to Fermilab and other institutions.

There are also two semi-autonomous systems which can produce signals for the electronics to process. These are the laser calibration system and the radioactive source system, both of which are used for calibration and systems checks. The laser calibration system will produce light pulses in the wavelength shifting fibers which resemble real data pulses. The front end electronics will treat them identically to real data, but depending on the flash patterns, the trigger farm may need to have a different trigger condition to recognize them. The laser calibration data are taken as a separate run so the whole readout chain is notified and the appropriate configuration downloaded.

The radioactive source system produces slowly-varying DC currents from the photodetectors which the electronics measures via the monitoring system. When a certain section of the detector has a source passing through it, it will be necessary to inform the monitoring system to sample the photodetector current for the corresponding set of channels. Any real data taken during this time may be corrupted and therefore the position of the source must be recorded concurrently with the event data.

The monitor and control system has connections to systems other than the electronics readout system. It monitors, among other things, temperature, humidity, magnet currents, and power supplies for all systems. Standard interface modules provide monitor inputs and output signals from the control systems.

6.3.4 Data links

The electronics for the both detectors will require fiber-optic links for timing signals from the GPS receivers to be brought from the surface to the underground detector halls.

A high speed data link between the near detector hall and the Feynman Computing Center (for the transmission of data to be stored), and a link to the Main Injector (to obtain spill timing information) will be required.

6.3.5 Power requirements

The electronics will be connected to the AC power mains. The connection must be to a separate 'quiet' power feed with a single solid ground connection. The electronics has been allocated 75 kW of power at the Soudan site.

6.3.6 Division of responsibilities among subsystems

Item	Subsystem
MUX boxes	Scintillator
Photodetector	Scintillator
Crates	Electronics
Racks	Electronics
Platforms	Near and far installation
Quiet power	Near and far installation
Clean (filtered) magnet power	Magnet steel & coils
Fiber-optics links for GPS timing signals at near detector	Near installation
Fiber-optics links for GPS timing signals in Soudan shaft	Far installation
Data links to Feynman Computing Center	Near installation

Table 6.5: Division of responsibilities among subsystems.

The division of responsibilities among the various subsystems for the provision of items described in this Section is shown in Table 6.5.

6.4 Description of WBS elements

6.4.1 Front end units (WBS 2.3.1)

A front end unit is an assembly which houses the electronics to digitize the times and amplitudes of the pulses from three photodetectors. It contains Field Programmable Gate Array (FPGA) logic to control this process and send serial data downstream to the data acquisition system. As well as supporting these functions it contains the high voltage supplies for the photodetectors and a microcontroller to provide monitor and control functions. A functional diagram of a unit is shown in Figure 6.5.

6.4.1.1 Channel electronics

Introduction. Each input from the photodetector has a separate channel of electronics, whose main function is to capture the peak value and time of each pulse. The signal path of a single front end channel is shown in Fig 6.6. After preamplification and shaping a fast branch of the signal path feeds a discriminator that detects the occurrence of the pulse and causes it to be sampled at its peak; this also permits the time of the signal to be recorded. In addition there are circuits for calibration. One is provided for charge injection, the other for measurement of the phototube DC output. The design described is simple and conventional. It is largely conditioned by the dual requirements for low power consumption and wide dynamic range. The first forces operation from 5 V supplies. It also means the design must be kept simple. It is expected that modest elaboration will prove necessary after initial prototyping, and allowance has been made for this.

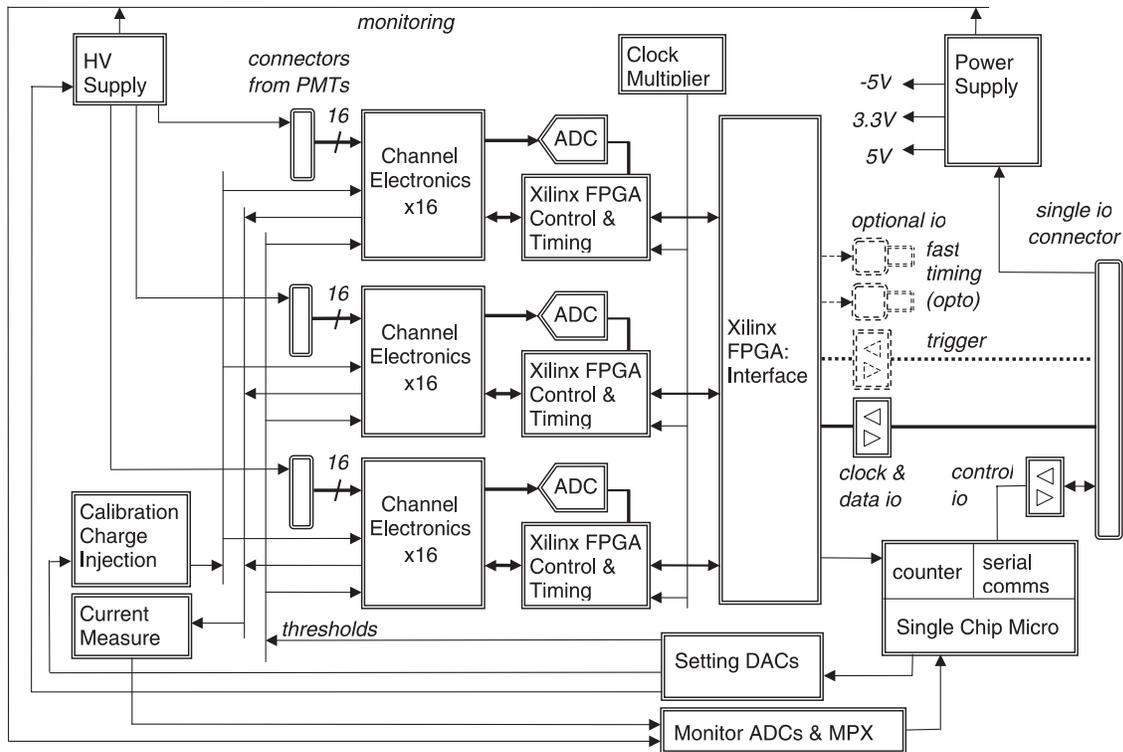


Figure 6.5: Block diagram of one front end unit serving 48 photodetector channels.

Gain and shaping.

Input shaping. A passive input network converts the pulse of charge from the PMT to a somewhat slower voltage pulse. The maximum output pulse amplitude will be a few hundred millivolts: low enough to not cause excessive crosstalk through capacitive coupling to other anodes in the PMT, but large enough to give a good signal to noise ratio. Protection against overvoltages is included.

Preamplifier. An op-amp is used as a preamp, taking a voltage input. It provides modest gain (about 8) to buffer the signal and lift it well above noise, while keeping within the output amplitude possible with 5 V supplies. A little extra high frequency roll-off suppresses noise from its input. It needs ~ 100 MHz gain-bandwidth, and low noise: the AD8005 chip has been used.

The output is a pulse with a peaking time of ~ 100 ns, amplitude from 2.5 mV for 1 photoelectron from a low gain PMT channel, to 3 V for 500 photoelectrons and a high gain channel. The input may be changed to a charge sensitive amplifier, once a satisfactory design can be completed. This would give a slightly better performance at similar cost and complexity.

Slow shaper and buffer. The relatively fast signal from the preamp is reshaped to give a peaking time of ~ 400 ns, and buffered to drive the sampling circuits. This uses another

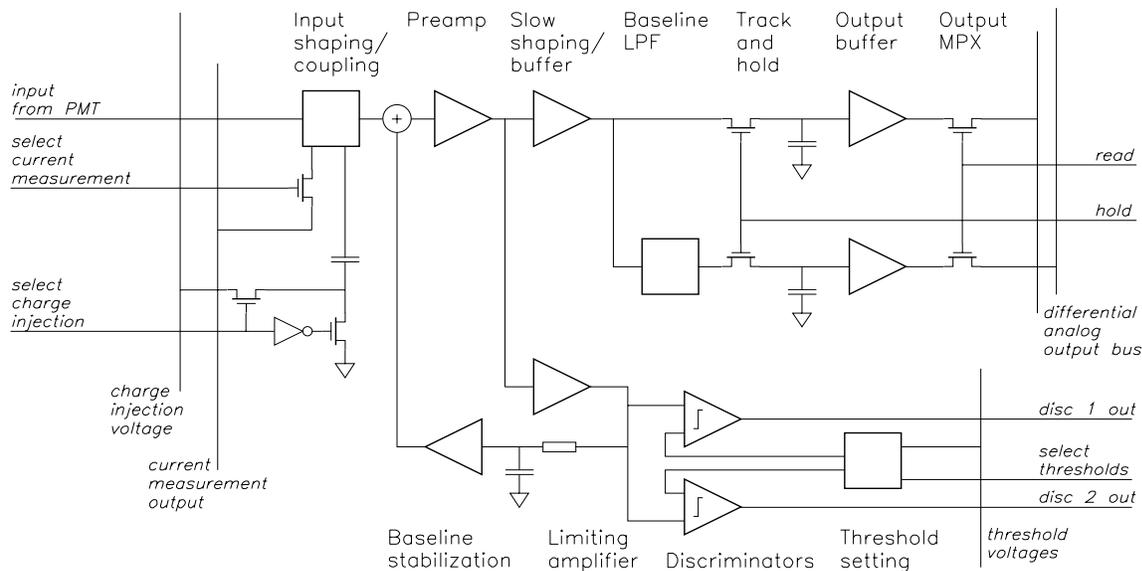


Figure 6.6: Signal path of a single front end channel.

op-amp, and has to have a low output impedance.

Peak sampling.

Baseline tracking. The signal path is DC coupled and there can be a significant offset voltage and, possibly, low frequency (e.g., 60 Hz) interference. The mean baseline of the signal is tracked by splitting the signal and passing one branch of it through a low-pass filter. Both the unfiltered and the filtered signal outputs are sampled and passed differentially to the ADC thereby rejecting DC and low frequency signals.

Track and hold. The low pass filtered and unfiltered signals are fed through two sections of a switch IC to hold capacitors. The voltages on these track the inputs until the switches are turned off. The instantaneous voltages are then held for subsequent readout until the switches are turned back on. The performance of such a circuit is dominated by the switch. The limiting factor is injection of charge when the switch is turned off; the differential arrangement used makes this balance as far as possible. It also makes it necessary to use as large hold capacitors as will give acceptable speed, and to provide low driving impedance. Performance is better the slower the system; the 400 ns shaping is fast enough to make this a critical part in attaining the required dynamic range. The DG611 D/CMOS quad switch is used.

Output buffers and switches. The hold capacitors are read out through unity gain FET op-amp buffers, which are relatively non-critical parts; their low bias current is needed to avoid the stored voltage on the hold capacitors drifting too rapidly. Finally, two more

sections of the quad switch are used between the buffers and an analog output bus (still differential). This allows the output from one channel to be selected for digitization.

Discriminator system.

Limiting amplifier. The output from the preamp also feeds the discriminators. Since fast comparators need their threshold to be set at 20 mV or more, the minimum pulse height at this point has to be ~ 50 mV. Additional gain of about 20 is therefore required, and large pulses will saturate. The amplifier is thus designed to limit, and recover from this condition with minimal delay. This has been designed around a fast op-amp and needs extra components for limiting. The amplifier is fed with a sum of the preamp output and the inverted slow shaper output, thus forming a bipolar shaped signal which allows the comparator to act as a type of constant fraction discriminator.

Baseline stabilization loop. The threshold, at its lowest setting, corresponds to only 0.1 mV referred to the preamp input. This is too low to ignore offset in the amplifiers. The problem is circumvented by adding a low frequency feedback loop from the output of the limiting amplifier to the preamp input. This is a standard form of baseline stabilization. It uses a cheap FET input op-amp configured as an integrator to give the long time constant desired. It also adds limiting on the input to the amplifier; this is not strictly needed, but is a refinement that greatly reduces baseline shifts produced by large pulses at high rates.

Comparators and threshold setting. Two comparators are provided; one is an auxiliary to allow a higher threshold and could be used for an input to a trigger system. The MAX907, a very low power dual comparator, is used. A faster part could be substituted at the cost of increased power consumption. Individual channel threshold setting is provided to take account of variations in PMT channel gain. This may be done by selecting between four possible voltages derived from upper and lower values set by DACs for a group of channels. This is a coarse setting, but seems sufficient.

Calibration.

Charge injection. The input to each channel includes provision for charge injection. It has a capacitor, which can be switched between ground (normal operation) and a charge injection line. This line receives a variable voltage step by the common driver circuit. It is possible to inject charge into one or more selected channels.

Current measurement. The input network is arranged to exhibit a moderately high impedance to DC. The currents from the PMT during source calibration generate easily measurable voltages which can be switched with CMOS switches to a common output line for measurement. The network isolates the signal pulses from this DC path. Coupling to the preamp is AC, so the amplifier bias current can not affect the measurement.

6.4.1.2 Measurement circuits

Digitization.

ADC multiplexing. It is not practical at present to use one ADC per channel to obtain 14-bit resolution. Even though each ADC could be quite slow, they are too expensive and need too much power. Much better results are obtainable sharing one faster converter among several channels. This multiplexing is done by using the output switches on each front end channel to put the held voltage on a common pair of lines for measurement; these lines feed through a buffer amplifier to the ADC. Sixteen channels are multiplexed to one ADC. Suitable 14-bit ADCs have recently become available: the Analog Devices AD924x family. The 3 MS/s AD9243 is used here.

Mode of operation. Operation starts with the channel control logic noting that a hit has been detected on a channel, holding its peak voltage, and recording its time. The control logic of the channel mezzanine continuously scans the status of all its channels, comparing times and selecting the earliest with a hit. As soon as the ADC is free, the output switches of this channel are turned on and its stored voltage is converted. The channel control logic is told the hit has been processed, the track-and-hold reverts to track, and the channel is readied for the next hit. The data from the conversion are paired with the recorded times and channel numbers, and output to the control logic on the main board. There the streams of data from the three channel mezzanines are merged, keeping the data in time order. The data are then passed back to the hub with sufficient buffering to smooth the flow.

Sweeping. The ADC, like all fast, high resolution parts, has significant nonlinearities. Measurements made over a narrow range cannot be relied upon to much better than one count. This is an issue for calibration using muons which measure only over the low part of the ADC range. To overcome this an offset voltage is added into the ADC input. This is derived from a DAC set through the control system, and is able to move the zero over approximately 10% of the range of the ADC. This allows adjustment so low level signals do not use any point in the range of the ADC that proves to be particularly bad; commonly, the differential linearity of ADCs is much worse than average at a just a few points in the range. It would also permit the more sophisticated approach of gradually sweeping the zero position to average out nonlinearities.

Timing.

Channel timing. Timing with 5 ns resolution is provided for each channel. It is implemented in the FPGA which contains all the logic for a group of 16 channels. It uses a 100 MHz clock to drive a 4-bit Gray code counter, with transitions on both clock edges. The outputs from this are distributed to latches, one for each channel. Because the clock and fast counter outputs are not widely distributed on the chip, power dissipation remains low. Note that present FPGAs are capable of being driven by and counting 200 MHz clocks; the present design is slightly conservative.

The Gray code counter provides fast timing, but with only an 80 ns wrap-around period. A separate counter and set of latches, driven by a slower clock which is also used to synchronize the inputs to be timed, provides low resolution timing with a wrap period long enough to avoid ambiguities. Logic combines the two parts of the time into a single binary coded output (18 bits long).

Plane timing output. A high resolution timing output will be provided for groups of channels, rather than each separately. This will use a separate very fast amplifier and discriminator serving a group of 8 channels, inductively coupled to all the channel inputs in the group, to sense the leading edge of the first photoelectron of an event. The discriminator outputs of the groups of eight channels belonging to a plane will be OR'ed and taken to fast optical fiber outputs to allow connection to external TDCs, without risk of interference to the operation of the system. The outputs will be reconfigurable to work in a test mode, allowing operation of the unit to be monitored without interference and without involving other parts of the electronics system.

Clock. The high frequency clock is generated by a clock synthesizer IC, with an internal VCO and dividers forming a frequency multiplying PLL. This gives a flexible choice of clock speed (up to at least 160 MHz, at present), and maintains synchronism to the system clock without having to distribute an inconveniently high frequency. It will be located on the main board.

Control logic.

Functions. The basic function of the logic on the channel mezzanines is to time and sequence the hits, convert their amplitudes, buffer the hit data, and output it in time order to the main board. The FPGA there merges the three streams keeping the hits in time order, and transmits the output stream back to a hub.

An important secondary function of the logic is to detect, as far as is practical, hits that can be detected separately but which cannot be completely or accurately handled by the system. For example, the occurrence of a second hit during the dead period of a channel after a first has been registered can be detected, and indicated by flag.

The logic also has to monitor its own operation. Most important, it has to detect if operation goes outside normal parameters. It should detect if, for example, buffers become fuller than is anticipated, thereby giving warning of problems, as well as detecting when overflow actually occurs. Similarly, it should check that when times with a short cyclic period have to be extended to longer ones, that no errors are occurring (this is a statistical check only). It also has to make the basic check that its internal time counters remain synchronized with the system time; as long as this is true no action is needed, but a reset has to be made and reported if a deviation is found and confirmed.

Special modes. Various other special modes are required. Some are already included in the design; others are likely to be wanted. Most important are those for measuring the ADC pedestals. This may be done during normal operation by causing a low rate of

conversions to occur at random in the absence of a detected pulse. The data from these are flagged, but otherwise pass through the normal path. If this method is not wanted it can be disabled. It is also possible to command sampling to occur with no pulse detected; data can be returned over the usual path suitably flagged, or via the control system. The various options will be selectable.

The charge injection system is also controllable. It uses the normal operating mode to acquire the data, though this may be modified to gate it to select just the wanted signal. It can however also allow variations to permit more extensive testing of the system. For example, varying the delay from trigger to sampling, and timing sampling from the charge injection rather than the detected pulse.

Diagnostic modes are also needed. For example, to generate dummy data. This allows independent checkout of the later stages of the system, and simplifies diagnosis of data handling errors. It is also easy to generate controlled data rates, continuous and burst, to check the system. Charge injection can also be used for this: and it can be extended so operating across multiple units it can mimic events and test processing algorithms.

A major advantage of putting the logic in FPGAs is that it will usually be possible to add such features after the basic hardware is developed.

Implementation. All the logic for 16 channels are implemented in one Xilinx FPGA. This includes the timing, and buffering the data output. The configuration ROM for the FPGA is on the main board, so the same configuration is applied to all three mezzanines.

Calibration.

Charge injection. Each channel has a charge injection capacitor and switching; a group of these is driven by a single low impedance source of a voltage step, the amplitude of which is accurately variable over a wide range.

The driver uses a bipolar differential pair to rapidly switch the tail current between two loads of $\sim 50 \Omega$. The tail current is set by a precision feedback loop from a DAC driven by the control system. The current consumption is reduced by turning this current off when charge injection is not in use.

To get good accuracy ($\sim 1\%$) over a wide range in amplitude requires careful neutralization of stray coupling capacitances and may require two drivers for high and low parts of the range. A fast risetime and very accurate timing should be possible, allowing reliable calibration of channel timing.

Source current measurement. The front ends can measure the DC current from the PMTs. This is mainly used during the radioactive source calibration mode, but it also allows effective monitoring of the tube leakage currents which will be useful for diagnostic purposes. Each channel can multiplex a voltage representing its DC input onto a single line for measurement. This will be done using a precision op-amp, a low pass filter, and a low-speed high-resolution ADC connected to the control and monitoring microprocessor. This could be the ADC used for other monitoring functions but a separate ADC is used for

simplicity: the cost can be very low. High resolution is not strictly needed, but allows for considerable variation in the strength of radioactive sources that might be used.

Laser calibration system. The laser calibration system is used to inject light pulses into the the wavelength shifting fibers and to test the readout chain from the WLS fibers, through the PMTs, to the front-end electronics. The laser system can inject light pulses of varying amplitude, and so allows us to measure the linearity of the entire scintillator system and its readout chain on a regular basis. In addition, the short duration of the laser pulses allows us to calibrate the relative timing and time walk of each detector channel.

6.4.1.3 Ancillary functions

Monitor and control.

Concepts. A wide variety of monitoring and control functions is required. In normal operation settings have to be made, such as thresholds. It is also desirable to monitor conditions such as supply voltages, and status signals that are not conveniently handled through the main data path. There are also special operating modes to be supported, such as PMT current measurement and pedestal measurement.

These requirements are sufficiently complex that handling them independently by a microcontroller is the easiest and most flexible way. The required data rates are low, so communication over the serial interface included on the chip is adequate. A major advantage of this approach is that development of the monitoring and control is essentially independent of the signal handling.

Microcontroller. The microcontroller should be a mid-range single chip microprocessor. Speed is not important. It would have a few kBytes of program memory, and a few hundred bytes of RAM. Most of its I/O will be through local serial interfaces to ADC, DACs, and the various FPGAs, so a large number of pins will not be required: a 28 pin device is probably adequate. Choice is largely a matter of convenience: for example, the inclusion of a high-resolution ADC in the PIC14C000 which would avoid the need for an external part, or flash program memory allowing in-circuit reprogramming during development.

Required functions. Many control and monitoring functions will be required. A partial list of required functions includes:

1. setting operating mode (triggerless, or various triggered);
2. disabling unwanted or bad channels;
3. initiating pedestal measurements;
4. setting discriminator thresholds and selecting which to apply to each channel;
5. setting ADC offsets (statically, or with a slow sweep);
6. setting tube voltage(s);

7. reinitializing unit hardware;
8. selecting charge injection mode; setting pulse amplitude and channels to inject; and initiating operation;
9. selecting channel to measure tube current, and making measurement;
10. checking incoming and local supply voltages;
11. checking tube voltages;
12. checking tube currents;
13. checking unit temperature;
14. measuring count rate from a selected channel or channels, and
15. reporting status flags in the logic indicating incipient problems or actual errors.

Clock and timing. A local synchronized clock and timing mark must be derived from the system timing. It is adequate that the relation between these is stable, but it desirable that they have a minimal offset. It is also necessary that they are free from glitches. The clock and timing mark signals are transmitted differentially from the hub over two pairs in the cable connecting the hub to the front end unit.

Power supplies.

Low voltage. The multiple low voltage supplies needed within a front end unit are assumed to be derived from a nominal 48 V input through a DC/DC converter and additional regulators. The unit will draw this power from the hub to which it is attached.

High voltage. The design of the front end units includes local generation of the high voltage for the PMTs which offers considerable advantages. It is safe, it allows units to be 'hot swapped', it does not introduce ground loops, it limits high voltage faults to the single unit, and saves cabling. A low power DC/DC converter, with an added filter, and an output stabilization feedback loop regulating the input voltage is used to generate the HV. The load is essentially static, the tube currents being negligible compared with the divider chain current. Individual adjustment of each tube voltage will be provided.

Interface. The only connection between the front end unit and the rest of the electronics is by a single cable to a hub. The use of a single cable avoids ground loops and injects minimum interference into both the front end unit, and the signals carried over the cable. Even so, for reliability the signals are differential and should have good immunity to common mode interference. A screened twisted-pair cable up to 6 m long, with 25-way D-subconnectors, is used to connect a front end unit to a hub. The functions of the conductors are:

1. Data output. The data are output serially using one pair of conductors. The average rates are so low that there is no technical difficulty, and buffering handles the peaks.
2. Power to the front end unit is supplied from the hub using a single pair of conductors in conjunction with a separate 0 V connection.
3. Clock and timing signals require two pairs of conductors; one for the 10 MHz clock, the other for the timing mark.
4. Control and monitoring use two pairs, as separate serial links. Rates are low and a separate link for each direction to the control microprocessor is simple and cheap.

The remaining pairs of conductors in the cable could be used to carry trigger signals if necessary.

6.4.1.4 Mechanical construction

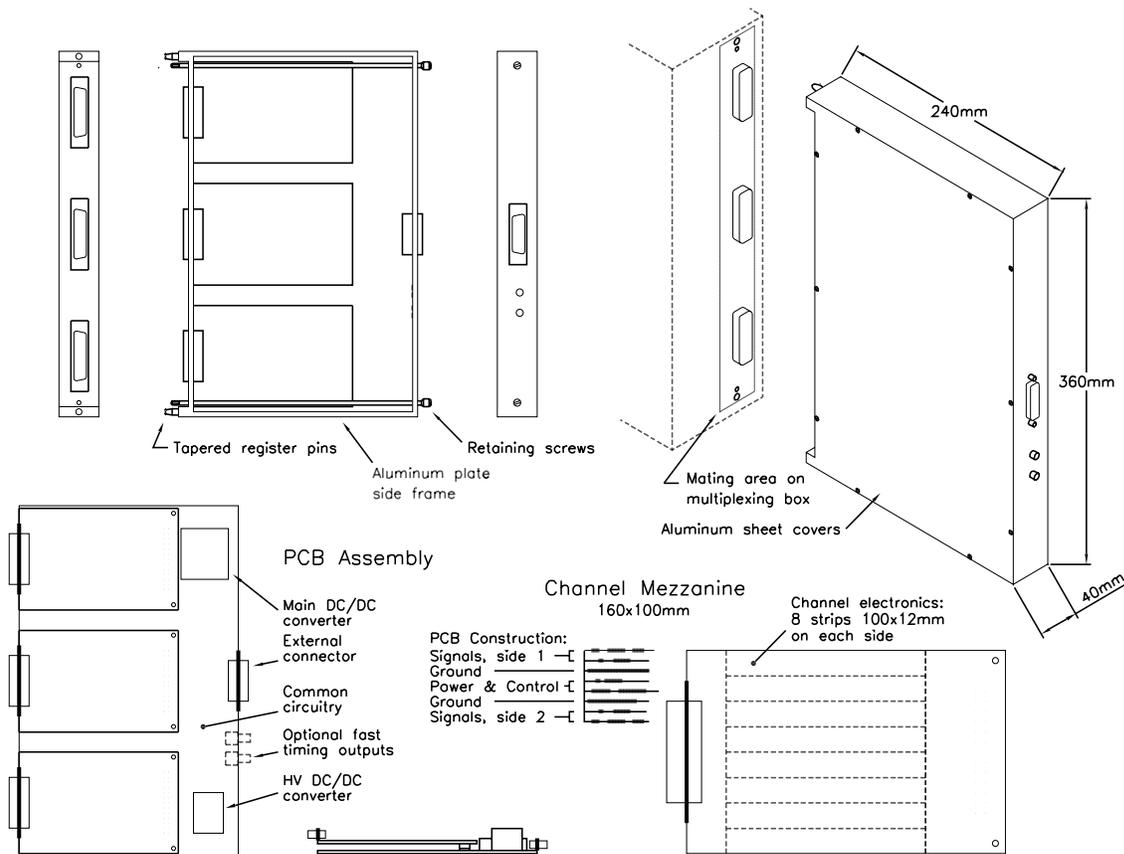


Figure 6.7: Mechanical construction of the front end unit electronics.

The mechanical construction of a front end unit is shown in Figure 6.7. Details are still not finalized, and it is hoped to reduce size and cost, although this would not affect the general organization described.

Electronics. The electronics for each group of 16 channels is built as a separate channel mezzanine board. Three such boards plug on to a main board with the common parts of the system. Each mezzanine has an external connector to one PMT assembly; all other external connections are to the main board.

Channel mezzanine. The channel mezzanine is an 8-layer, 100×160 mm printed circuit board. Construction is entirely surface mount (except connectors), using both sides. Eight channels are laid out on each side, occupying strips approximately 12×90 mm, with the remaining length of the board used for the connectors, the ADC, logic, and other parts common to all channels. The signals are carried on the outer two layers on each side. Layers 3 and 6 are ground planes, and layers 4 and 5 are for power and control signals. Component density is low enough that through vias should suffice. Efficient automatic assembly should be possible.

At the front of the board is a mixed-contact D connector which provides the 16 external signal connections from the PMT, as well as supplying high voltage via a high voltage insert. It is fastened directly to the case, and supports the front of the mezzanine.

Main board. The main board is a 4-layer, 340×200 mm printed circuit board. The mezzanines occupy the larger part. Its own circuitry is in a strip along the back, with the IO connector in the middle of the rear long edge. Power supply components are well clear of the sensitive input circuits, and are electrically screened.

Mechanics. The electronics is housed in a shallow case without ventilation. The case is a little bigger than the main board, and about 50 mm deep. It consists of a frame made of four lengths of 6 mm thick aluminum, to each side of which a cover plate is screwed. Connector apertures and fixings are machined in the front and rear of the frame.

The unit assembles vertically to the rear of the associated MUX box. Two tapered register pins on the front of the unit mate with holes in the MUX box to guide the connectors into alignment. The unit is held in place by two long clamping screws passing through the entire depth of the box and operated from the rear. The front of the unit fits securely against mating areas on the MUX box-PMT combination, the whole forming a rigid assembly.

Environmental. The closed front end case will be heated by the power dissipated inside; the outside surface temperature is predicted to rise by about 8 K above ambient. The thermal coupling to the PMTs is poor enough that the electronics should raise the temperature of the photocathodes by no more than 2 K. The air temperature within the case will rise by about 15 K, which causes no problems for the electronics.

The case, while closed, will not be sealed. The unit could thus be affected by high humidity, which is a possible problem at the near detector. The internal temperature rise helps, and a conformal coating will be applied to the high voltage circuitry. There should be no problem if the ambient relative humidity remains below 90%, given that the temperature is relatively stable.

6.4.2 Hubs and interface crate (WBS 2.3.2)

6.4.2.1 Hubs

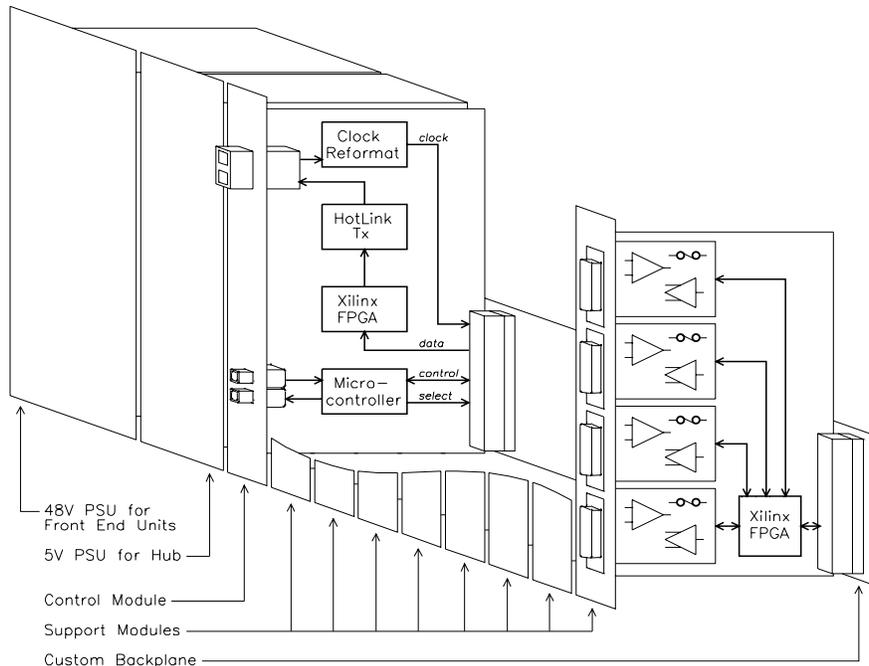


Figure 6.8: Diagram of a hub unit showing the major components.

A hub is a single crate which distributes services and DC power to up to 32 front end units, and accepts data from them for transmission to the central system. It contains a number of support modules and a control module and will be connected to the interface crate of the central system with two 2-core optical fibers. The data are sent out over one fiber, in the form of a 250 MBd serial link. A second provides a low-skew clock and timing signal. The other pair provide monitoring and control. No elaborate processing is provided. The data are merely concentrated onto the one output link for convenience, but not otherwise manipulated. Monitoring and control is likewise simple, though here local decoding of address commands allows individual front end units to be selected.

Figure 6.8 shows a diagram of a hub crate. It is a single 6U crate, though the height might be increased to 9U with power supplies in a 3U subdivision, if support for 64 rather than 32 front end units is required. All modules are VME compatible, using rows *a* and *c* only of the P2 connector position for signals, though this is not required in the system as designed.

Support module. The support module has four panel connectors for cables to front end units. It distributes 48 V power to these through protective Polyswitch fuses, and the clock through ECL drivers and coupling transformers. Differential drivers and receivers are

provided for the signal lines. These connect to a small Xilinx FPGA. This reduces the four serial data streams to one multiplexed stream which is clocked at 10 MHz and passed to the control module. It also combines the four serial control and monitoring connections into one, using address lines from the control module; this includes broadcasting, and activity monitoring. Finally, it provides the logic required when the triggered operating mode is selected.

Control module. The control module has a fast fiber optic transceiver taking duplex SC connectors. The input is a multiplexed clock and timing signal which the module reformats for distribution through the support modules to the front end units. It may convert the clock frequency. The output carries the data from all the front end units in one serial stream. It is driven by a Cypress HotLINK transmitter chip at 250 MBd. This in turn is driven by a Xilinx FPGA which takes the serial data from all the support modules. The data from each front end unit remains in time order but the 32 streams are combined without regard for their relative timing. A second fiber optical transceiver, using separate low speed devices, also accepts a duplex SC connector to provide monitoring and control. It connects to a microcontroller, as well as to the support modules. The microcontroller implements local control functions, in particular decoding and applying address commands, passing addressing information to the support modules. It also provides monitoring for the hub itself. This is independent of the main data path, except that some control and monitoring functions will be implemented.

Hub crate and ancillaries. A standard 6U subrack is fitted with a 48 V, 360 W plug-in power supply for the front end units, a second power supply for local power, and a simple custom backplane. The backplane distributes power, and connects the support modules to the control module.

6.4.2.2 Interface crate

A single crate provides an interface to all the hubs on a detector. It distributes clock and control, and receives data and monitoring information. The function of the crate is primarily the conversion and redistribution of signals.

The interface crate connects directly to each hub; fiber optic cables are used for electrical isolation. All are of equal length to avoid introducing skew in the clock distribution. One cable (two core, 'zip' type) with duplex SC connectors handles clock and data; a second handles the control and monitoring.

The interface crate passes the data, as electrical high speed serial links, to the central data handling system. It provides control and monitoring connections as electrical low speed links (RS232) to the computer managing these functions. It accepts the clock and pulse outputs from the GPS timing receiver and uses these to generate system timing. The timing receiver also has an RS232 link to the control computer, providing status and time of day information.

Clock and data interface module. A 6U board provides 8 fiber optic transceivers for the clock and data cables from 8 hubs. These are high speed (266 MBd or better) for low clock

skew and jitter. For purposes of costing, the 1.25 GBd HFBR5305 is assumed. The board takes the clock and pulse outputs from the timing receiver, and produces a multiplexed signal for distribution. These inputs will go to only one board, with provision for slaving others to it so as to minimize skew. Clock frequency multiplication will be included if necessary. The received data is transformer coupled onto the output, to give best common-mode noise immunity, and protect against accidental damage. The electrical connections are on the rear connector (arranging this so the board is compatible with a VME system, in case this should ever be desirable).

Control and monitoring interface module. This is of similar physical form (and again VME compatible), but with low speed optical transmitters and receivers in place of the transceivers. An RS232 electrical IO is on the rear connector. It is assumed that each link will be connected to a separate serial port in the control computer (so the board might be connected 1:1 to an octal serial PC card), but multiplexing so it can be connected to a single port will be included, and this will be extended to work across a set of boards. A microcontroller to provide local monitoring functions will be included, and this will have some spare inputs available for monitoring conditions off the board.

Crate assembly. For the far detector, three of each type of board are housed in a 6U VME format subrack, together with a plug-in supply for their power. No backplane is required; power is distributed by front-panel patching connectors. The assembly is similar for the near detector but only two of each board are required.

6.4.3 Central data system and trigger farm (WBS 2.3.3)

6.4.3.1 Hardware organization

The central data system for the far detector consists of four central receiver crates and one trigger farm crate; only 3 receiver crates are required for the near detector. The receiver crates receive and sort hit data from the interface crate and pass them to the trigger farm for software event selection, via dual high speed (132 Mbyte/sec) data links. A third, similar, link is used for communication between the trigger farm and receiver crates. Only two links (one for data and one for control) will be required if the total detector singles rate is less than 10 MHz.

Central receiver crates. The central receiver crates are small (seven slot) 6U VME crates which contain receiver and sorter cards. They are connected by high speed data networks to the trigger farm crate. Data from the receiver cards are available on the VME backplane and commercial off-the-shelf processors will be used for time sorting the hits and triggering.

Receiver cards. A receiver card is a custom-built unit containing the buffering and logic to prepare the data from the front end units for sorting into trigger time blocks by the sorter cards. The number of hits stored in these time blocks depends on the detector data rates and flexibility to change the size of the trigger time blocks is included.

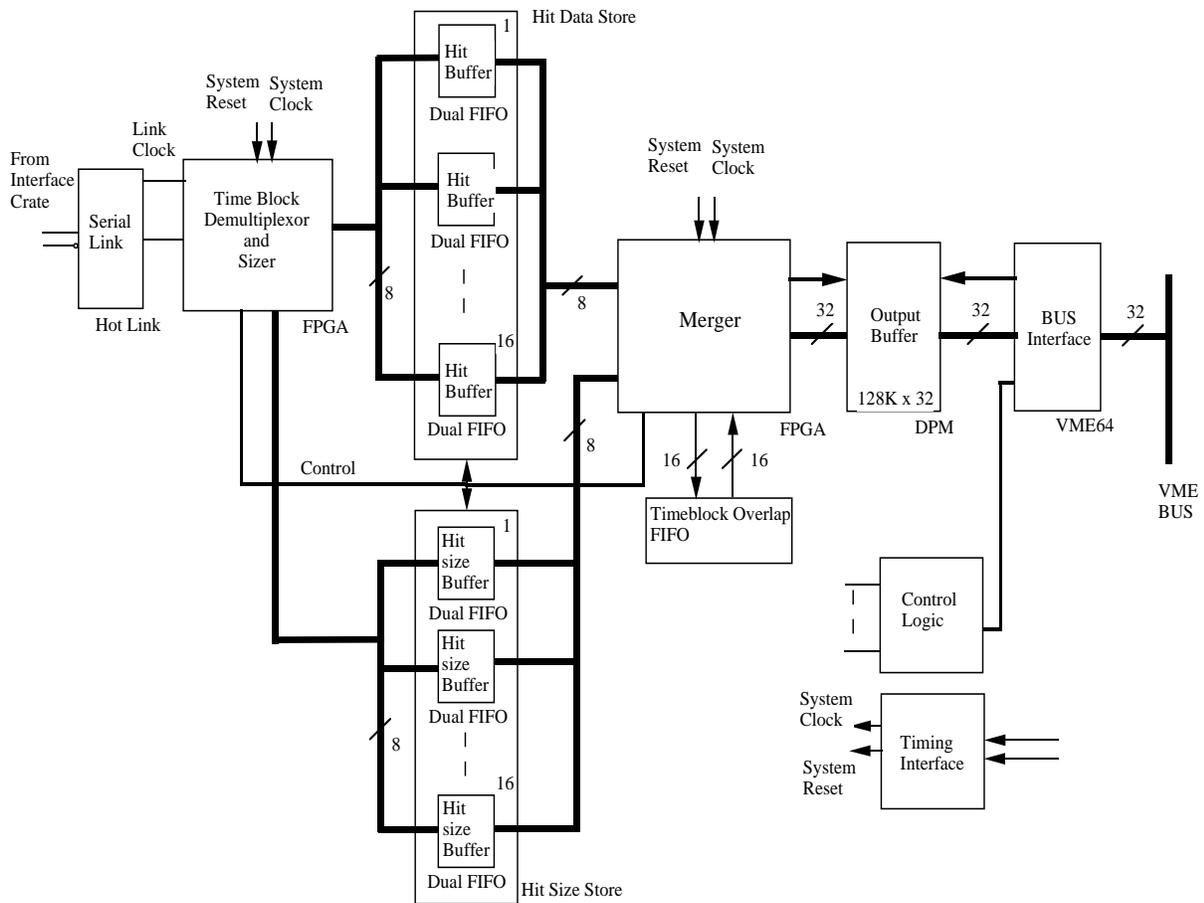


Figure 6.9: Block diagram of a receiver card.

A block diagram of a receiver card is shown in Figure 6.9. Each receiver card is a 6U VME card which receives a single cable from the interface crate. Each input cable is a short serial electrical link based on commercially available transparent point-to-point links (HotLINK). The data from the interface crate come from 32 front end units whose data have been packed together. The rate of this data stream will not exceed 5 Mbyte/s which is well below the limit of the HotLINK at 33 Mbyte/s.

The serial data are converted into bytes from the HotLINK receiver. Each front end unit data stream is processed by the time block demultiplexer/sizer which is a Field Programmable Gate Array (FPGA) connected to the HotLINK. The data are written into a hit buffer FIFO according to which front end unit it originated from. At the same time this FPGA will measure the length of the hit data and store this in a second bank of FIFOs, the hit-size buffers.

The two banks of hit and hit-size FIFOs are connected to a second merger FPGA which extends the hit data with the full time stamp and channel identity, as well as making time blocks of predetermined size for the trigger farm.

Trigger time blocks are made by first reading the hit sizes from the hit size FIFOs and then extracting the corresponding hit data from the hit FIFOs. At the same time 32 bit wide data words are constructed which are more suited for the VME64 bus at the back. The merger FPGA will also periodically preserve one of the front end time blocks to overlap with the next trigger time block.

A hit output from the merger FPGA will be in two 32 bit wide words which are stored into the output buffer dual-port memory. The same process of sizing the length of the time blocks is performed by the merger FPGA; their lengths are stored in the output buffer.

Sorter card. The receiver cards present the sorter cards with time blocks of data from 128 front end units. The time blocks contain hits from any of these units for a fixed period of time ready for time ordering. Time sorting is required to reduce the processing power requirements of the trigger farm processors. Each sorter card is a VME single board computer and the sorting is performed by software.

Each sorter card contains two main software algorithms: a data-sequencing algorithm and a merge-sort algorithm. The merge-sort algorithm merges the receiver card time blocks into a single trigger time block with the hits in time order. The data-sequencing algorithm manages the input/output transfers, calls the merge-sort algorithm and manages the onboard memory buffers.

Two sorter cards are used in each central system receiver crate, one for handling odd, the other for handling even trigger time block numbers. One of the sorter cards will be a master and the other a slave. The master data sequencing code is responsible for coordinating the allocation and availability of the time blocks to the trigger farm.

Each receiver card time block consists of 32 sub time-blocks, one from each front end unit, each of a different length. The data-sequencing extracts the lengths and base address pointers from each of the four receiver cards and transfers the data from the receiver cards to the sorter card input buffer.

The merge-sort algorithm then merges the 128 time blocks from the four receiver cards into a single output time block where all the hits will have been sorted in time order. The

merge-sort algorithm returns the length of the merged block, which is the sum of all the receiver card block lengths, to the trigger farm.

The software required to implement the sorting is quite simple. The processing time for these algorithms should be very similar to the ones used and benchmarked for the trigger farm.

Trigger farm. The operation of the trigger farm has been described in Section 6.1.5. The hardware comprises a VME crate containing up to 20 VME single-board computers each with dual PCI Mezzanine Cards and a high-performance PMC-based network card for data transfer. Two of the single-board computers, the input processor and output processors, are dedicated to sequencing. The remaining processors are used for the selection of events to pass to the data acquisition system. Only one trigger processor will be required if the nominal far detector singles rate of 0.9 MHit/s is achieved; no more than 10 would be required to handle the maximum design rate of 20 Mhit/s.

Trigger farm - receiver crate networks. A maximum of three 132 MByte/s PCI-PCI transparent bridge networks will be used to connect the central receiver crates to the trigger farm. Two of these are used for high speed data transfer; the third is used for controlling the data transfers. The CES PCI vertical interconnect bus has been assumed in the cost estimate.

6.4.3.2 Software algorithms and organization

The software algorithms that will be used in the trigger farm have been described in Section 6.1.6. The algorithms are simple and efficient and should not require extensive software development effort.

6.4.3.3 Interface to the data acquisition system

The interface from the trigger farm to the data acquisition system will be from the output processor Ethernet port using standard TCP/IP protocols.

6.4.4 Data acquisition system (WBS 2.3.4)

The main function of the data acquisition system (DAQ) is to receive events from the detector electronics and record them on a mass storage device. Because the event rates are low, the requirements of the MINOS data acquisition system are not demanding. The central system and the trigger farm build events and present them to the DAQ.

The Fermilab DART system[8], which is flexible, supported and tested, will be used. The anticipated data rates, which are described in Section 6.4.3.2 and Chapter 9, are low and well within its capabilities. DART will run on a PC and an associated cluster of PCs linked by Ethernet will provide real-time data analysis and diagnostic capabilities.

DART provides a graphical user interface which allows the operator conveniently to control data taking runs. In addition it provides an online event buffer from which separate

programs on other PCs in the cluster can access and process events as desired to make (e.g.) online displays of events or monitor detector behaviour.

The DAQ receives events from the output processor in the trigger farm. The output processor single board computer runs the VxWorks real-time operating system and passes the data to the DAQ PC via the Ethernet LAN using conventional TCP/IP protocols.

The event data will be treated differently at the two detector sites. At Fermilab the data will be sent directly to the Feynman Computing Center and stored in the tape vault; at the Soudan site the data will be written locally to a mass storage medium and then copied and distributed to other institutions, including Fermilab, for processing.

Monitoring and control functions will be provided by PCs in the cluster and use custom software, as appropriate to the specific functions required.

6.4.5 Database systems (WBS 2.3.5)

The databases will maintain a record of the construction history of the detector components and the state of the detector during its operating life. Some of the information will be recorded only once, or very infrequently; some will be recorded regularly to enable the performance of the detector to be tracked and calibrated. Information will be recorded either manually or will be derived automatically via software from the various monitoring subsystems (such as the embedded processors in the front end units and hubs) and the data acquisition system.

The information in the databases must be available in many places simultaneously; it will be used to ensure the smooth construction, installation and operation of the detector and will also be vital to the offline data analysis.

Examples of the information that will be recorded in the database, together with the level of detail, the source of the information and the approximate frequency of recording, are given in Table 6.6. The list is by no means exhaustive and some experience will be necessary before the exact frequency at which some of the information is to be recorded can be determined.

As well as recording information about the hardware of the experiment the database will also be used to log book-keeping information such as run start and stop times, numbers of events and operator comments.

The information will be made available to the data processing computers by a single master database at Fermilab containing all the data and satellite synchronized copies at Soudan and the near detector. The satellites will primarily service the trigger farms and first-level offline processors, which will limit their usefulness to queries from remote locations. Physicists working elsewhere will either query the master database or make database copies for their own sites.

Some of the information in the database will be generated and entered remotely, for example at the scintillator factory or the extrusion plant. Since these locations may not have reliable network access, the data will be entered into a PC running mSQL in Linux or Microsoft Access and subsequently downloaded to the master database by dumping it to a flat file and reading it into the master.

The database will store the modification level and repair histories of the various electronics boards. It will be possible to consult the database for details of switch settings and board placement if items need to be replaced. The interface will be via a World Wide Web browser

Item	Information	Level	Source	Frequency
Construction				
Steel plates	mass	plate	installation	once
	thickness	plate	installation	once
	magnetic properties	plate	mill	once
Scintillator	QC tests (light yield)	batch	extrusion factory	once
WLS Fibers	QC tests (light yield)	batch	module factory	once
Scintillator modules	dimensions	module	module factory	once
	mass	module	module factory	once
Photodetectors	QE, uniformity & gain	pixel		once
	operating HV	unit		once
Electronics	modification level	board	factory	as req'd.
	switch settings	board	factory/test rig	as req'd
	calibration constants	channel	factory/test rig	once
Installation and survey				
Steel plates	positions	plate		once
	alignment			once
Scintillator modules	positions	module		once
	alignment			once
Scintillator modules	light yield	strip	source calibration	once
Photodetector	gain	pixel	light flasher	once
Electronics	calibration constants	channel	pulse injection	once
Operation				
Scintillator	calibration constants	strip	muons (software)	weekly
Photodetector	gain	pixel	light flasher	daily
	replacement history	unit		as req'd.
Electronics	calibration constants	channel	pulse injection	daily
	disabled channels	channel	software	as req'd.
	FPGA & ROM programs	unit	software	as req'd.
	Trigger farm software		software	as req'd.
	high & low voltages	each PSU	software	hourly
	repair history	board	manual	as req.
	Magnet	current	each PSU	monitoring systems
field		plate	monitoring systems	hourly
General	environmental conditions (e.g. temperature)		monitoring system	as req'd.

Table 6.6: Examples of the information that will be recorded in the database during the different phases of the experiment.

looking at a web page with a Java applet running a JDBC connection to the database, which will not only find the correct switch settings, but also invoke a picture of the board with the switches set correctly. The same applet will also request the serial numbers of the new and old boards, and put them in the database. Having the same program with the switch settings request the serial numbers will ensure that the new numbers will be entered immediately, correctly and promptly.

The size of the database is estimated to be about 4 GBytes after five years running. Over half the space will be occupied by calibration constants determined from throughgoing muons. The relationships between the individual data items is not sufficiently complicated to warrant an object-oriented database with its attendant expense and possible small company problems. The Oracle database system, which is a commercial, supported product, provides all the required functions. Fermilab has a great deal of experience with Oracle, which is used extensively on site. The Electronics and DAQ cost estimate assumes the use of Oracle as the MINOS database.

6.4.6 Auxiliary systems (WBS 2.3.6)

6.4.6.1 Monitoring

Monitoring functions (for high and low voltages, etc.) are built into the front end units, the hubs and the interface crates, and have been costed with these components. Monitoring of the central system crates uses the Canbus system and is included in the cost estimate. Other control and monitoring functions will be provided by simple custom or commercial standalone units linked to a control PC by Ethernet. One PC at each detector site is required to support the monitoring systems.

6.4.6.2 Absolute timing

The complete far- and near-detector electronics systems must be synchronized to absolute time. The absolute times of the Main Injector spills at Fermilab must also be recorded to allow events (especially in the far detector) to be associated with the MI neutrino beam. GPS timing receivers are the preferred source of absolute time and routinely give accuracies of 100 to 200 ns for moderate cost units.

GPS timing receivers will be used in all three locations: the far detector at Soudan, the near detector and the Main Injector. At Fermilab we assume the timing receiver output can be used by a spare channel of an existing timing module to make the spill time available from their computers over the Internet.

At both the near- and far detector sites the GPS receivers (which require minimal space) will be located on the surface to eliminate the need for low-loss coaxial feeders to be installed in the shafts. Timing signals from the receivers will be sent to the interface crates over fiber-optic cables installed in the shafts.

Distribution of the system clock derived from the GPS system is by way of the interface crate and the hubs; the costs of clock distribution have been included in the costs of these subsystems.

Communications between Soudan and Fermilab are assumed to be by Internet. Timing

is only a minor use of the communications links, and the fiber optics cables are provided by the near and far detector installation tasks.

6.5 Future optimization and engineering

All components of the electronics system will require prototyping and optimization. The areas where this will be most necessary are the front ends and the central system; the goal will be to reduce costs and improve performance.

6.5.1 Front ends

The design described provides a conservative basis for the estimation of costs and schedules; it is not finished at the component level. The design is considered usable for full scale production, but we expect that further work will lead to higher performance and lower cost. Implementation of the channel electronics in an ASIC may prove worthwhile; this is not, however, part of the baseline design presently proposed.

The front end electronics described is specific to the baseline detector design: i.e., Hamamatsu 16 channel PMTs, and 23,000 channels at the far detector and 9,400 channels at the near detector. This system has, however, been designed to allow for some changes that are being considered.

The power dissipation has been kept low enough to allow for possible upgrading from eight-fold to four-fold multiplexing on the far detector, doubling the number of channels to 46,000 (or even 70,000 with the addition of a third supermodule).

The timing performance of the detector is limited by the scintillator and the wavelength shifting fibers to ~ 10 ns rms for a single photoelectron signal. In the current design of the electronics, time is measured on each channel with a resolution of 5 ns, which allows a relatively simple design using FPGAs; a high resolution timing output is provided for groups of channels. The inclusion of high resolution timing on a per-plane basis internal to the front end units by the use of an existing TDC chip is anticipated.

If there is a substantial increase in the light yield from the scintillator it may be worthwhile to improve the single-channel timing resolution. A modest improvement may be possible with the design as described: a factor of two seems as if it may be possible, but this would need to be demonstrated.

Two 16-channel units, each with one M16 phototube, are currently under construction. Although they are considerably simpler than a full front end unit they are designed to be used to demonstrate that the required noise performance and noise immunity (e.g., to welding noise) can be achieved with the proposed scheme, to study timing performance and to study the effects of fringe magnetic fields on the DC – DC converters. The experience gained with them will allow the final front-end design to be optimised; in particular these units have exchangeable mezzanine cards to allow details of the critical analogue first stages to be studied.

6.5.2 Hubs and interface crate

The hub-interface crate system uses well-established technologies. Some prototyping and optimization will, nevertheless, be required. In particular, the scheme for clock distribution needs to be proven and optimised.

6.5.3 Central system

Development and prototyping will be needed to ensure that the desired performance requirements of the central system can be met, and to confirm the estimates of the size of the final system. The development program for the central data system and the trigger farm will make use of a prototype, scaled-down version of the hardware, the network connection and the software to develop and study the performance of the system.

Most of the hardware used will be commercial and, by developing a scaled-down version of the receiver hardware (as PMCs on the microprocessor boards), the real-time aspects of the system will be studied. With some added functions this test receiver card can also be used to act as a generator of simulated data for subsequent receiver card tests. The network performance will also be studied before the final system is assembled.

6.5.4 Trigger farm

Ongoing physics simulation studies will be performed to optimize the trigger algorithm as the detector design evolves. In addition, the special requirements for cosmic ray and calibration triggers will be investigated and the effect of overlapping events in the near detector on trigger algorithms will be studied.

Further analysis and design of software for the trigger farm will be undertaken leading to the production of prototype code, which will be tested and evaluated on a microprocessor system using simulated data as input. Simulation of the trigger farm will allow us to investigate communication, control and scheduling issues in a realistic environment and to ensure that processing, bandwidth and error recovery requirements are understood. This will require DART to be installed on a PC and interfaced to the simulated trigger farm output processor which, in turn, will allow us to acquire experience in the operation of DART in the MINOS context and in adapting it to our specific needs.

6.5.5 Triggered operation

The baseline electronics system is designed to be adaptable to operating in a simple triggered mode if this proves to be necessary.

Assuming a trigger output for each detector plane is available from the front end units, an ' M of N planes' hardware trigger can easily be implemented in the hubs, which are located in a single row along each side of the detector. A trivial addition to the support modules in the hubs would be needed to pass plane trigger outputs between adjacent boards. The simple logic required would be added to that already present on the Xilinx FPGA.

It would be necessary to pass plane trigger outputs between adjacent hubs to allow a M/N type trigger to be made fully efficient for planes at the boundaries between hubs. Only

$N - 1$ signals, probably less than 8, would have to be passed. Direct electrical transmission as differential signals over the short cable required would be acceptable; the addition of a simple interface module to buffer the signals and bring them to a connector is all that is required. The hub is built so these trigger modules merely have to be plugged in if operation in a triggered mode becomes necessary. In this scheme a trigger from each side of the detector is formed separately. The two triggers would be combined before being distributed to the front end units. It is impractical to combine the outputs from both ends of each plane before building the trigger. Such a system would take a large amount of extra cabling, and considerably increase installation costs.

The FPGAs on the front end unit mezzanine boards and main boards can be programmed to operate in a triggered mode. The ready reprogrammability of these devices allows details to evolve as requirements become clearer. They would allow 'per plane' trigger outputs to be formed, which would be passed back through the connection to a hub. They would also accept the global trigger passed back in the same way. The sampling of peak values and recording of times of hits would be unchanged, but now, if no trigger were received within a preset time (say 750 ns) the recorded values would be discarded; voltages would be converted and data output only if there were a trigger signal within a preset time.

A second way a triggered system could operate would be to use slower shaping, and only sample the peak of the photodetector pulses on receipt of a trigger. All channels could then be read out, and it would not be necessary for pulses to have been detected totally reliably by the discriminators on each channel. This strategy would require a long enough shaping time to allow a global trigger signal to be formed, passed back to the front end units and still to sample the peak of the pulse with sufficient accuracy.

Variations would be possible within the triggered mode. Different logic to form the trigger outputs could be selected; times, such as how long pulses are stretched, could be varied. There could also be another input to the trigger system; a signal summed across all channels of one plane could be formed and provided with an additional discriminator to permit a 'total energy' trigger. It would not work well unless the spread in channel gains is corrected, and would not be practical at the near detector where one plane is spread across many front end units.

6.5.6 Test beam studies

MINOS test beam work will start in 1999 and will continue off and on for several years. Initially it will be directed to studying the scintillators being used for calorimetry. The final version of the electronics will not be needed, or available, for the initial tests; commercial digitizing and readout electronics will be sufficient to obtain an understanding of the behaviour of the scintillator and photodetector system. Later runs will be used to test modifications of the scintillator-photodetector system.

The most important part of the electronics system to test is the front end, where photodetector pulses are digitized. The back end (hubs to trigger farm) of the system can be sufficiently well simulated, and tested in the laboratory. Some channels of preproduction front-end electronics will be used with the test beam module for a combined system test in order to identify any problems before full-scale production is started.

The final running, which might take place as late as 2002, will be made to obtain a full

calibration of the complete MINOS detector and readout system. It will be important for these runs to use production electronics in order to understand the calibration and resolution achieved in the final detector system, including the electronics.

These schedules are compatible with the planned program of electronics development.

Chapter 6 References

- [1] J.H. Cobb, "Dynamic range requirements of the MINOS electronics," April 1998, Fermilab report NuMI-L-368.
- [2] J.H. Cobb, "Cosmic ray muon rates in the MINOS near detector," December 1996, Fermilab report NuMI-L-232.
- [3] K. Ruddick, "Underground particle fluxes in the Soudan Mine," September 1996, Fermilab report NuMI-L-210.
- [4] K. Ruddick, "Counting rates in the MINOS plastic scintillator detector due to rock and steel radioactivity," October 1997, Fermilab report NuMI-L-313.
- [5] J.H. Cobb, "Singles rates in the MINOS detectors," July 1998, Fermilab report NuMI-L-372.
- [6] The University of Minnesota, CNA Consulting Engineers, Ericksen-Ellison Associates, Inc., and Miller-Dunwiddie, Inc., "MINOS Far Detector Laboratory Technical Design Report (Including Basis of Estimate & WBS) for Cavern Construction, Cavern Outfitting & Detector Outfitting," October 1998, Fermilab report NuMI-L-263.
- [7] The Fermilab NuMI Group, "NuMI Facility Technical Design Report," Chapter 5, October 1998, Fermilab report NuMI-346.
- [8] Information about the DART system can be found at the URL <http://www-dart.fnal.gov:8000>.