System considerations for a digital optical system for a large scale neutrino observatory

H. Kirkham, S.P. Jackson, J.C. Ling, J.Y. Lai, D.T. Liu, E. Hsu, T. Cole
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109
D. Nygren, D. Lowder, G. Przybylski, *Lawrence Berkeley National Laboratory*;
F. Halzen, *University of Wisconsin*;

A prototype **digital optical module** has been constructed for use in conjunction with the **AMANDA neutrino detector** at the **south pole**. The engineering design of the digital module is described, in particular the system aspects of the present design, which was constrained to operate in conjunction with (and adjacent to) existing analog data acquisition hardware. Some results are discussed.

Part 1. BACKGROUND

1.1 Neutrino Astronomy Physics

The neutrino has been central to a number of important discoveries in nuclear physics, particle physics and astrophysics. Because they are not deflected by magnetic fields, and interact only weakly with matter, neutrinos can be used to study objects for which some other types of radiation would be unusable over cosmic distances.

Several types of astronomical objects are expected to be sources of high energy neutrinos, in particular active galactic nuclei (Punch *et al*, 1992; Quinn *et al*, 1996). Much is expected to be revealed about other astronomical processes, such as gamma ray bursters (Dermer and Weiler, 1995) and supernovae (Wischnewski *et al.*, 1995). In addition, a large neutrino telescope might allow confirmation and study of neutrino oscillations (Halzen, 1995).

However, because neutrino interactions are rare, the detectors must be large. Water is a convenient material for a large detector. The idea of using large water-based detectors to detect particle tracks is relatively simple: Čerenkov radiation is emitted by muons created when a high-energy neutrino interacts with matter. This radiation is visible light, so it is transmitted through water, and can be detected. The usual way is with a photomultiplier tube (PMT) housed in a watertight sphere. The combination is called an optical module (OM). These OMs are arranged in an array so that each monitors photons crossing through some volume of the water used for detection.

During the last decades there have been more than a dozen large detector arrays designed for extraterrestrial neutrino detection. All of the current large neutrino detectors use water as the Čerenkov medium except for AMANDA (Antarctic Muon and Neutrino Detector Array) at the south pole. AMANDA (Askebjer *et al*, 1995), located deep in Antarctic ice, is presently the only US detector project. This paper describes the development of a digital optical module to be used in conjunction with the AMANDA project.

1.2 Science Requirements

Planning for a next generation of high-energy neutrino telescope has been under way for some time. The reason for this is that full realization of the scientific opportunities will require a detector scale of 1 km^2 in effective area or 1 km^3 in effective volume (Gaisser *et al*, 1995).

In the long run, the science requirements for a neutrino observatory can be summarized as follows: the detector must furnish the data to allow the reconstruction of any muon tracks in a volume of 1 km^3 with angular uncertainty of the order of one degree. In order to reconstruct tracks with sufficient accuracy, it must be capable of timing the arrival of photons at a given detector in the array within about 1 ns. (The exact timing requirement will depend on the spacing of the detectors in the array.) Each detector should have a dynamic range that allows the detection of single photoelectron events, as well as the measurement of 1000–10 000 photoelectrons together. There is a need to understand possible clusters of pulses arriving at a detector.

Although the waveform of the output is somewhat an artifact of the detector being used, questions about multi-photon events cannot be answered without it. Therefore a goal of the design must be to preserve waveform information, and to separate closely spaced events.

Because the events being sought are rare, it is likely that observation of even a large volume of water must be carried out for many years, perhaps even decades. Because of this, there is a need for a highly reliable observation system, and preferably one that can be repaired and maintained during its lifetime. Meeting these core requirements is a formidable technological task in a large array.

The work described here was done by the Jet Propulsion Laboratory in collaboration with the Lawrence Berkeley National Laboratory. This partnership has its origins in a workshop on neutrino astrophysics technology held at JPL in March 1994. Workshops in June 1994 at Snowmass, Colorado, at LBNL in December that year, and at Arcadia, California in April 1996 helped define the technical areas in which development was thought to be most needed. During this period a group of scientists and engineers from JPL, LBNL and other institutions began to meet regularly at JPL and LBNL. This became known informally as the Local Working Group (LWG).

Through the LWG, and with discretionary funds provided by JPL and its parent organization Caltech, and by LBNL, some initial development work was undertaken in 1995 and 1996. One optical module was sent by the LWG to the pole in the southern summer 1995-6. Two revised OMs were sent early in 1997. Results of this initial work will be presented in this paper.

1.3 Engineering Issues

The development of a large neutrino observatory is an engineering project. The prototype described in this paper was viewed as a technology testbed for a future cubic kilometer array. The science requirements, outlined in the previous section, drive the design of the observatory, but the way the various problems are solved involves engineering and engineering management. We next review some of the engineering problems.

First, the need to reproduce short pulses presents a bandwidth problem. (This is examined below). A digital signal, transmitted at slow speed, was viewed as a solution. A design based on digitizing the signal inside the OM, and transmitting the waveform as digital data was therefore contemplated. While this met with scepticism from some in the science community, we think that the success of the prototypes proves the point that a digital system is workable in the polar ice environment.

Second, there was a desire to connect the newly developed system to the existing system at the south pole. While the new hardware could be viewed as a technology testbed, it was important to check it against the AMANDA approach. The difficulty of combining a digital system with an analog system had to be overcome.

1.3.1 The Bandwidth Problem

A PMT waveform has a risetime of just a few ns, so that according to the usual rule that BW $\gg 1/\tau$ (where τ is the pulse width), a bandwidth of much more than 100 MHz would be required to transmit it faithfully.

This is difficult with an analog transmission line. Analog lines introduce loss that varies roughly as the square root of the frequency, with the result that a complex waveform such as the output of a PMT will be distorted at the receiving end of the connection. A PMT pulse of several volts magnitude inserted into the cable will appear at the receiving end as a tiny pulse of just a few mV (or less), with greatly increased risetime and falltime. For example, the commonly-used RG-59 coaxial cable has an attenuation of about 0.4 dB per 100 m at 1 MHz, rising to about 10 dB at 100 MHz. A bandwidth of 100 MHz would be reached with about 30 m of such cable. 3 km of this cable (required to reach around a 1-km cube) would have a loss at 100 MHz of 300 dB. The waveform could not be reconstructed.

A less stringent requirement is the resolution of multiple closely spaced photon hits. The rule of thumb for the resolution (rather than reproduction) of closely spaced rectangular pulses is that the bandwidth must be more than $1/2\tau$, where τ is both the pulse width and separation. If 20 ns PMT pulses 20 ns apart are even to be recognized as separate therefore requires a bandwidth of 25 MHz. A length of about 100 m of RG-59 would be the limit.

Better cables do exist, but none is likely to be an order of magnitude better, and all have the same frequency characteristics. The best one could hope for with an analog circuit of a few km or more would be to resolve pulses that were a few tens of ns apart. The waveforms could not be recovered.

One alternative to communicating the PMT waveshapes would be to use fiber optics. However, this would require that a separate cable be installed for power, and thus an additional penetration of the OM.

While a fiber system could transmit a digital replica of PMT waveforms in real time, a wire-based system could also transmit a digital replica, albeit at a rate slower than real time. This has the advantage of being a more conservative approach than the use of fibers, and is compatible with the delivery of operating power to the OM.

The option we selected was therefore digital, and wire-based.

1.3.2 Connections to AMANDA

While it was necessary to connect the new-technology detectors to AMANDA to verify performance, it was crucial not to affect adversely the performance of the AMANDA system. AMANDA is a system in which all the OMs are connected to a central location by wire—coaxial cable was being used when this work started. The pulses from the PMTs, which have an amplitude of a few volts and risetimes on the order of a few ns, were applied directly to the cable, and thus signalled to the central unit.

An interface between the new LWG OMs and the AMANDA timer would be available. Essentially, the new system used the AMANDA hardware as no more than a simple coincidence detector. If one of the PMTs connected to the central unit detected a "hit," it would start a timer. Other PMT signals occurring within the duration set by the timer would then be recorded, but isolated signals (not representing tracks) would be discarded. The connections are shown in simplified form in Figure 1.

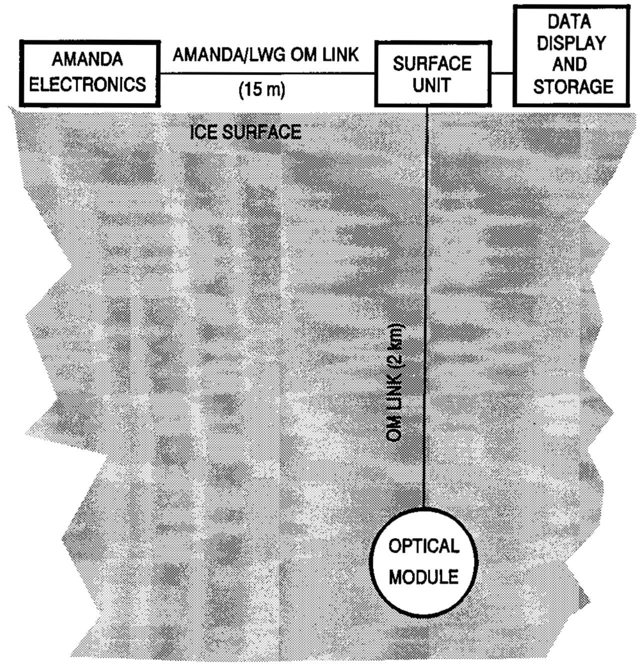

Figure 1
The LWG optical module and its connections to AMANDA

Part 2 System Design

2.1 Overview

The AMANDA approach is to communicate the outputs from all the PMTs to a central location by means of real-time signalling. In order for the new LWG OM to be compatible with AMANDA, it had to furnish a real-time signal. (In this sense, it may not be a precursor for a future large array.) A trigger pulse was sent from the OM as soon as the PMT registered a hit. If the signal passed the timer coincidence test, AMANDA would return a pulse to the LWG OM instructing it to send the data representing the recorded PMT hit. The sequence of events was thus

1. The PMT hit was registered in the optical module, and a fast pulse signifying this was sent immediately to the AMANDA surface station. The waveform was stored in analog form. AMANDA determined photon arrival time from the fast pulse.
2. If AMANDA determined that there was a coincidence with other photons, it would send a command back to the OM, requesting the data.
3. On receipt of this command, the OM would digitize the stored waveform, and transmit the data to the surface.

Inside the LWG OM, a microprocessor controlled the operation of the analog waveform store, and the communication system. Under its control, if a signal was not received from AMANDA within 30 μs (signal round-trip time from the OM to the surface, plus a margin), the analog waveform would be cleared, and the system reset. The microprocessor also controlled the PMT HV supply, to adjust the PMT gain in response to signals from the surface module.

2.2 Optical module

2.2.1 ATWR

The LWG optical module was built around a fast digitizing IC designed at Lawrence Berkeley National Laboratory. (Kleinfelder, 1990). This chip, called the Analog Transient Waveform Recorder (ATWR) captures an input waveform by storing samples of it on capacitors that are switched, one after another in rapid succession, onto a circuit carrying the signal. Figure 2 shows the approach, somewhat simplified.

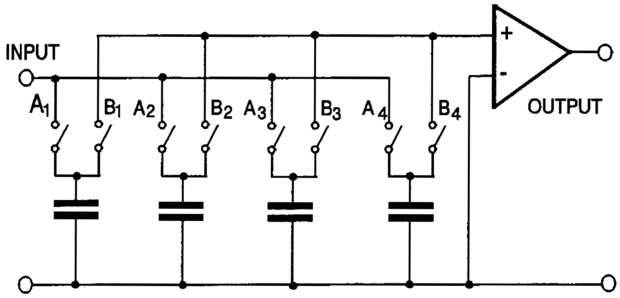

Figure 2
The Analog Transient Waveform Recorder circuit captures information as charge on capacitors.

To record a signal, the switches A are momentarily closed in turn. (Since signals in hardware typically travel at about 0.66*c*, or 2×10^{10} cm s^{-1}, the ATWR dimensions correspond to a few ps, so it can be considered to be a point for the purposes of sampling.) To digitize the signal that has been captured, the switches B are momentarily closed, one at a time, so that the output amplifier can buffer the capacitors (which are on the order of 0.5 pF, and must only see high impedance). The reading out process takes very much longer than the reading in—high impedance electronics tends not to be fast.

The method is not new. Waveform recorders using this principle have existed for years. What is new about the ATWR is that it is on a single chip, with speed and flexibility. The integrated circuit has 6 channels that take 256 samples adjusted, in this application, to occupy a duration of about 100 ns.

2.2.2 OM Circuit

The LWG OM was responsible for capturing a PMT waveform, signalling this capture to AMANDA, digitizing the signal if so instructed, and communicating the result to the surface. All this was accomplished over a pair of wires, which also served to deliver the power to operate the OM. Figure 3 gives an idea of what was involved in the OM.

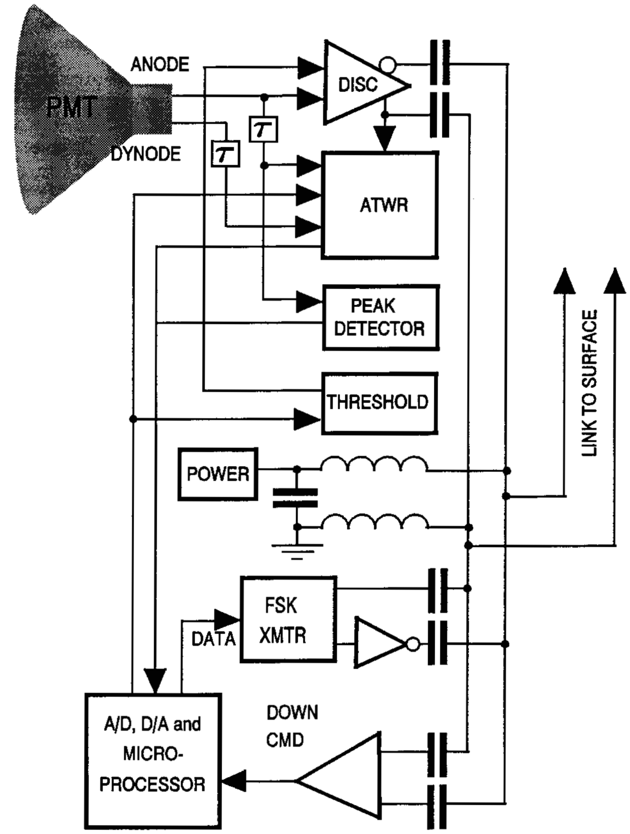

Figure 3
The LWG OM system of analog and digital electronics operating under the control of a microprocessor.

The output of the PMT is connected to three separate systems. First, it drives a discriminator that compares the PMT signal with a threshold set via the microprocessor. If the signal exceeds the threshold, the discriminator generates a pulse that signals AMANDA, and starts the ATWR. The discriminator's delay in the signal to the surface is minimal, and AMANDA can use the pulse for photon timing.

The ATWR does not respond instantaneously, that is, the instruction to capture a waveform must arrive slightly ahead of the waveform of interest. Therefore, the signals of interest—the PMT anode and one of the dynodes—are each connected to a delay line that drives the ATWR. These delayed PMT signals are sampled by the ATWR, which stores the pulses in analog form until a decision is made about whether or not to digitize them.

Finally, the PMT anode pulse is connected to a peak detector that is read by the microprocessor. If the signal is large, there is a possibility that the anode signal will show nonlinearity. In this event, the microprocessor automatically uses the signal from a dynode, sampled by the ATWR at the same time as the anode signal.

2.2.3 Power

The problems of supplying power to a large array has some similarities to the problems of supplying a spacecraft (remote, inaccessible) and a city (physically large, diffuse load). The spacecraft problem is largely one of choosing a source. In the case of the LWG OM, this was not a problem, as power was readily available.

The problem of supplying a city is one of wiring a network to all the places power is needed. For the two LWG OMs for AMANDA, the power had to be delivered over the two wires used for communications. While this was perfectly feasible, it is not a scaling solution. The question will be further addressed in any future work.

Each digital OM used about 375 mW, of which 150 mW was the power consumption of the PMT itself. The transmitter circuitry, which had to drive several km of cable, was the next biggest power consumer. Nevertheless, 375 mW is not considered to be a high level of power: aggregated over a future km^3 array with (say) 10^4 OMs, it amounts to less than 4 kW, including a margin for losses. This could be easily supplied by a photovoltaic array of 100 m^2 area, for example, even allowing for batteries for 24-hour operation.

2.2.4 Time scaling

One of the features of the ATWR is that the sampling speed can be controlled by setting the voltage on one of the IC pins. This allows the flexibility to have the 256 samples occupy a shorter or a longer time even after the OM is in the ice and inaccessible. It also means that the sampling speed is a parameter that should be stored along with the data samples.

In the early designs of the LWG OM, the sampling speed was calculated at the surface by examining the waveform of a clock signal, sampled along with the PMT anode and dynode signals. Later, a software method of deriving the sampling speed inside the OM was adopted. The communication traffic was thus decreased, because instead of transmitting 256 8-bit samples of a waveform, all that was sent was the number of samples in one clock cycle, a single 8-bit number.

2.2.5 Communications

The link from the base station to the digital optical module has many functions. In the down direction the link carries the power for the LWG OM, and commands for the processor (for example to adjust the PMT voltage). In the up direction, the link transfers the pulse that signifies capture of a PMT pulse, and data describing that pulse.

Managing the traffic on this hybrid link is a shared responsibility. Both ends of the link are capable of transmitting, and therefore of jamming the other end. The way it is supposed to work is this: the link is silent (apart from transferring power) until a PMT waveform is captured. This event causes a pulse to be sent from the LWG OM to the surface, Figure 4. If the AMANDA system (not shown in the diagram) determines that the pulse might be interesting, it sends a down command to the OM to digitize and transmit the data.

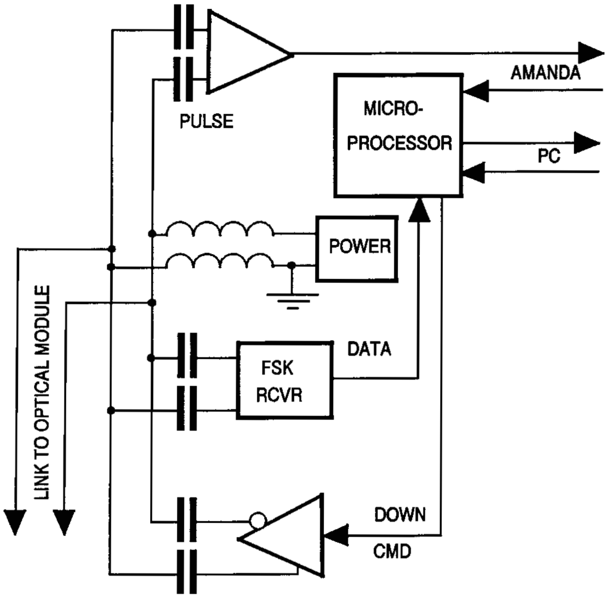

Figure 4
The DOM surface system, where power and commands are sent to the OM, and the returned data are received and stored.

Once the digitize command is received, the OM is essentially dead (unable to capture more pulses) until the data digitization is finished, a time of about 3 ms. The link is tied up for a further 10 ms or so transmitting the waveform data.

The OM that was sent to the south pole at the end of 1995 did not return any data. As part of the checkout procedure after arrival, it was energized on the surface, and tested. It seemed to have failed.

A system autopsy conducted later at JPL suggested the following scenario. In order to avoid causing problems with AMANDA, for the first test the AMANDA timer function had been replaced by a separate test circuit that *always* instructed the OM to digitize hits. What had not been recognized was that the south pole is not a quiet electromagnetic environment. The link from the OM to the surface unit was susceptible to interference from the drilling operations, and the test circuit was fooled into thinking there were many more hits than there really were. The combination of interference and a continuous data stream could not be handled at the surface, and the system looked broken. When the changes that had been made at the pole in an unsuccessful attempt to fix the problem were reversed at JPL, the system was foind to be working. Ironically, a connection to AMANDA might well have cured the problem, but this had not been attempted.

The susceptibility to external interference was reduced in the version sent a year later. One change was the use of frequency-shift-keying (FSK) rather than baseband signalling. This does two things to help. First, it moves the signalling to a higher frequency. Man-made noise sources usually have the characteristic that their amplitude falls with frequency (they are less than perfect impulse generators). Thus, a move to high frequency improves the signal-to-noise ratio. How much improvement is not known, since we know little about the noise experienced at the pole except that it is generated by SCR-controlled drills. Second, the use of FSK (and synchronous detection) improves the noise rejection properties of the link.

However, after these revisions had been made, the LWG was informed that AMANDA had decided to use twisted-pair wires instead of coax from the surface to the OMs. Cross-talk on the cables meant that the LWG OM signalling was likely to interfere with the wideband high-gain amplifiers used by AMANDA. The signal dispersion on the cable to the AMANDA detectors at the surface means that the received PMT pulses have a spectrum that goes down to about 1 MHz. This was originally chosen by the JPL team as one of the frequencies of the FSK communications. Sinusoidal FSK signals at 1 MHz could be misinterpreted by AMANDA as fast pulses from the analog modules.

As a fix, the FSK frequencies were decreased by a factor of 10, so that the highest frequency of the digital communication link was an order of magnitude lower than the expected spectrum of the PMT signals. Filtering could then be installed in the lines going into AMANDA, if needed to further reduce the effect of the cross-talk.

An additional result of this change is that the dead-time was larger than it otherwise would have been. It was the desire to minimize this dead-time that resulted in a number of software solutions to reducing the amount of data sent.

The data sent from the LWG OM to the surface are shown in Figure 5. There is a header, two telemetry blocks, the data body, and a trailer.

The header begins with the base-10 number 85: this generates the sequence 01010101, which aids the synchronizing of the FSK system. This is followed by the number 254, which generates the sequence 1111111110 and produces a recognizable signature for debugging purposes. The header also contains a counter that informs the surface system how many samples of the CPU clock were measured during the ATWR acquisition period, the peak level of the signal, and the data count. The clock pulse indication was necessitated by the desire to know the sample rate while at the same time to minimize the data transmitted—in earlier versions the complete clock waveform was sent. A longitudinal parity check (LPC) completes the header. (LPCs were used on all up- and down-data sequences.)

HEADER

85	254	CPULSE	PEAK	N DATA	HDR LPC

TELEMETRY 1

RAW LV	TEMP	HV SENSE	SMPL SPD	N DATA	TLM1 LPC

TELEMETRY 2

OFFSET	POP GRP	PCKT CNT	ANODE SAT	FLAGS	TLM2 LPC

DATA BODY

d_0	d_1	d_2	.....

.....	d_{n-2}	d_{n-1}

TRAILER

5	65	175	1	DATA LPC	TRAIL LPC

Figure 5
The LWG OM up-data consists of a header, two telemetry sequences and a variable-length data sequence, as well as a trailer

The telemetry blocks describe the performance of the OM. In the first block, the voltage at the receiving end of the cable from the surface is monitored, as well as the internal temperature, and the PMT high voltage supply. The sample speed is echoed back to the surface to confirm a command received. The size of the data block is also sent. (Actually, this is sent twice because there was no provision for retransmission in the event of a failed LPC.)

The second telemetry block includes the population group (see below for explanation), the packet count, and an indication of whether the anode signal had saturated (in which case the dynode signal would be a more faithful reproduction). There were 6 flags in the flag byte: bit 0 was used to turn the peak detection on and off, and bit 1 forced the data to be anode data or dynode data if the peak detector was off. Bits 2 and 3 were used to enable the leading baseline and the trailing baseline to be suppressed. Bit 4 enabled more or less frequent parity checks: LPCs could be set to be calculated every 5 data values or once for the entire data body. This flexibility was thought necessary because of the uncertainty about the noise in the environment. Bit 5 allowed the packet trailer to be omitted.

The remainder of the up-data sequence consists of the data describing the PMT waveform, and a trailer. As with the header, the trailer was somewhat redundant: it included arbitrary (fixed) numbers for debugging assistance and to help frame the stored data.

The overhead associated with the data is relatively large, because the data block was so reduced in size to accommodate the lower FSK frequencies. There are 24 non-data bytes, including the two telemetry sequences. The data typically occupied about 50 bytes.

This low number was achieved as follows. The PMT pulses were divided at the OM into two populations. The first 1023 pulses in a sequence of 1024 were population A. The last pulse was classified as population B.

Population B data were transmitted to the surface uncompressed. Population A data were compressed in a number of ways. Either anode or dynode data were transmitted, but not both. The number of samples used was reduced by calculating a running average in the OM, and sending only half the sample number. The value U_j of the signal at sample j was replaced by

$$U_j' = \frac{U_j + \frac{U_{j-1}}{2} + \frac{U_{j+1}}{2}}{2}.$$

The next value transmitted would correspond to U_{j+2}, so that all samples were used equally. In addition, the 2 least significant bits were truncated from the data, so that only six bits/sample were sent. Code to eliminate long sequences of data that correspond only to the baseline following the pulse was also developed, but was not activated in the 1996-7 version.

Population B pulses were characterized by sending the first 64 anode samples **and** the first 64 dynode samples, without compression. This was done to allow comparison of anode and dynode signals.

Compared to the traffic on the 1-MHz version of the communication link, these measures achieved a decrease by about a factor of 10, commensurate with the decrease in carrier frequency. As a result, the dead-time on the link was not worsened appreciably.

2.3 Surface System

At the surface, the wires to the LWG OM are connected to a power and communications box. Power (low-voltage dc) and down-commands are transmitted to the OM. Up-data are received and passed on. These various actions take place under the control of a microprocessor. There is also a PC programmed to store and display the received data (see Fig 4).

Part 3 Results

Two LWG optical modules were installed in the ice at depths of about 2 km in January 1997. After some initial problems with cable connections, they were operated beginning in March 1997. One of the modules (OM 1) failed in July 1997, but the other continued operating stably until it was turned off in December 1997.

Figure 6 shows a sample waveform returned from the DOMs, showing narrow, well defined pulses. The apparent coarseness in the DOM data is an artifact of slow sampling, not the result of quantization. Though the returned samples are effectively only 6-bit samples (1 part in 64 resolution), the A/D convertor stage was adjusted to read an output from the ATWR of just a few volts, so that the resolution is about 20 mV.

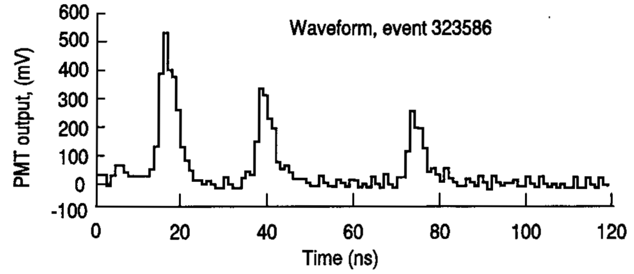

Figure 6
Sample signal returned from the digital OM

Sampling speed is a compromise between detail and number of pulses recorded. At twice the sampling speed, one of the pulses in Figure 6 would have been missed.

The next two figures show distributions obtained from approximately one hour of DOM data, and demonstrate more quantitatively the system's ability to resolve and separate single photoelectron pulses.

Figure 7 shows the distribution of pulse areas, in units of V ns. This plot is intended to determine the single photoelectron level for the DOMs; to avoid threshold bias in this determination, the initial pulse in each waveform is excluded from this plot. The shallow valley and broad peak at 1.5–2 V ns is characteristic of the measured single photoelectron peaks expected for the Hamamatsu PMTs used in the DOMs; the peak to valley ratios are typical of those measured for other PMTs in the AMANDA array.

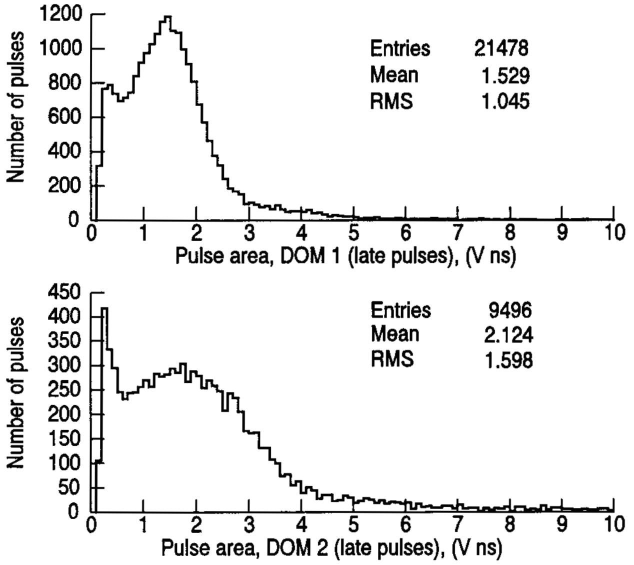

Figure 7
Distribution of pulse areas

A simple pulse-finding algorithm was used to find all pulses at least 50 mV in height and at least 3 samples wide in the data. For each pulse, the pulse width, height, and area were calculated; the time between each pulse and the next pulse in the waveform (leading edge to leading edge) were also determined.

Figure 8 shows the leading edge to leading edge time difference between the first and second pulse in each waveform. This plot quantitatively demonstrates that the DOMs efficiently separate pulses as close as 10 nsec apart.

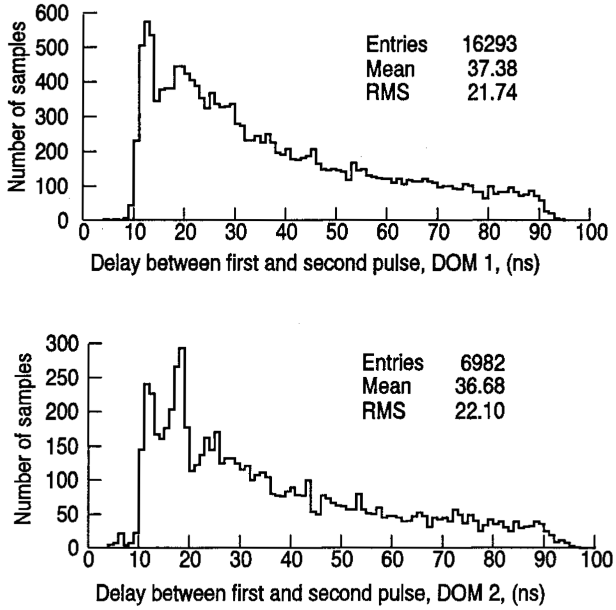

Figure 8
Distribution of pulse separation

Our analysis shows that the objectives of being able to reconstruct short PMT pulses, and to separate closely-spaced events, have been met.

Part 4 FUTURE WORK

The ATWR-AMANDA combined approach suffers two drawbacks. First, there is a need for a direct connection between the OM and the AMANDA central station, a solution that does not scale well. Second, the use of the fast pulse from the OM to the surface is incompatible with other uses for the wires—such as sending the digitized data. There is therefore dead-time while the data are being sent.

Both these problems could be solved. The point-to-point communications approach could be replaced with a much more cost-effective network approach. An all-digital system with time-stamped data would overcome the dead-time problem. However, the LWG has no immediate plans for future work. Available funding has been exhausted. Nevertheless, with the success of the digital approach, one cannot but speculate about a future km^3 array, and the issues of scale that arise.

Part 5 Conclusions

Prototype optical modules were constructed to demonstrate digital techniques applied to a neutrino observatory. They have been successfully tested in conjunction with the AMANDA neutrino detector at the south pole. The new detector used digital communications to return high-quality waveforms of the output of the photomultiplier tube. At the heart of the optical module was a new analog waveform capture IC that could sample the PMT waveform. The dynamic range is larger than for a conventional optical module because not only the anode signal, but also the waveform on a dynode, are returned.

The apparent simplicity of a simple point-to-point approach for data acquisition, communications and power has been replaced by a much more flexible and therefore more powerful digital approach. In two prototypes, the system flexibility allowed for a solution to potentially serious cross-talk problems without compromising data quality.

The networking possibilities afforded by the digital communications link will result in significant cost savings. Similarly, the simplifications in the power system that result from the processing power of the OM imply that the total system cost is much lower than the simpler direct wiring approach that has been favored in neutrino astronomy so far. The success of this development augurs well for the feasibility of the approach extended to a future very large array.

References

Askebjer, P., Barwick, S.W., Berstrom, L., Bouchta, A and Carius, S., (1995), Optical Properties of the South Pole Ice at Depths between 0.8 and 1 Kilometer, *Science* **267**, 1147–1150

Dermer, C.D., and Weiler, T.J., (1995) Cosmological Models of Gamma-Ray Bursts, Astrophysics and Space Science **231** (1-2) pp 377–388

Gaisser, T.K. Halzen, F., and Stanev, T., (1995) Particle Astrophysics with High-Energy Neutrinos *Phys. Rep.* **258**, 174

Halzen, F., (1995) The Case for a Kilometer-scale High-Energy Neutrino Detector *Nucl. Phys. B (Proc. Suppl.)* **38**, pp. 472–483

Punch, M., Akerlof, C.W., Cawley, M.f., Chantell, M., Fegan, D.J., Fennel, S., Gaidos, J.A., Hagan, J., and Hillas, A.M., (1992) Detection of TeV Photons from the Active Galaxy Markarian-421, *Nature* **358,** pp. 477–478

Quinn, J., Akerlof, C.W., Biller, S., Buckley, J., Carter-Lewis, D.A., Cawley, M.F., and Catanese, M., (1996) Detection of Gamma Rays with E>300 GeV from Markarian 501. Astrophysical Journal, **456** (2) Part 2: pp. L83–L86.

Wischnewski, R. *et al.,* (1995) The AMANDA Collaboration *Proc 24th Intern Cosmic Ray Conf* **1**, 661

Acknowledgements

This document was prepared by the Jet Propulsion Laboratory, California Institute of Technology, for the National Aeronautics and Space Administration.