Part II

Read-out Technology for IceCube

In April 2001, the Interim IceCube Collaboration Board took the decision to change the baseline technology for data transmission from analog optical transmission (see the US proposal from Nov. 1999) to digital transmission. Two competing developments had been underway before, the one called dAOM (digitally controlled Analog Optical Module), the other DOM (Digital Optical Module). This decision was based on successful tests with the DOM string deployed in 2000. These tests have been performed in Jan./Feb.2001 at the Pole. We also followed a recommendation of an experts committee led by Barry Barish (March 2001). Although both techniques meet the technical requirements, the DOM technique was favored, since a) it does not need optical fiber -- leading to a cost saving of 1.8 k$ per OM, b) it does not suffer from the still three times larger rate of fiber/connector failures during refreezing (6-10%), c) it promises an easier time calibration and d) it potentially has a higher dynamic range.

In 1999/2000, the DESY group had focused ist efforts on the development of the dAOM (with the light transmitters supplied by UW and UCI), LBNL led the DOM efforts. A full string (#18) with 41 DOMs was deployed in January 2000. The surface DAQ for the DOM was built by DESY. Due to the extreme time pressure, the LBNL group could not develop and install firmware on the DOM FPGA which is neccessary to demonstrate time stamping, timing resolution and fast data transmission. However, via slow communication it was possible to set the HV, read out temperatures and transmit waveforms at a very low rate. Since all 41 DOM had been equipped with a parallel analog readout via optical fiber as a backup, string #18 could be integrated into the normal Amanda-II DAQ.

We review the history of data transmission techniques in Amanda in section 1, list the physics requrements on the OM readout in section 2 and give a description of the DOM principle in section 3. Section 4 presents the results obtained during the last Pole season which demonstrate the feasibility of the DOM method.

1. Review of Amanda data transmission techniques

The present AMANDA-II detector has been deployed over several seasons, with a stepwise implementation of detectors with improved data transmission techniques.

An array of four strings (AMANDA-B4) was installed in 1995/96, with the OMs connected to the surface by an electrical coaxial cable that transmits both the HV and the analog anode signal; see Figure 1. PMTs are operated at a gain of 10^9 to drive the pulses over 2 km of coaxial cable without in-situ amplification. Hamamatsu R5912-2 PMTs with 14 dynodes are used to achieve the high gain; they require positive high voltage (HV) at the anode. The signal at the surface is picked up via a DC blocking high-pass filter.

The time-offset t_0 is measured by sending a laser-generated light pulse through an optical fiber from the surface to a diffuser ball close to the OM.

The anode signal is considerably attenuated and dispersed by the 2 km cable, with typical time over threshold of 550 nsec and rise time of 180 nsec. A correction for the amplitude-dependent offset of the time when the signal exceeds a fixed threshold (time slewing correction) is applied off-line, resulting in a time jitter of roughly 5 nsec.

Coaxial cables have been replaced by twisted-pair cables for the 216 OMs deployed on six strings (strings 5-10) in 1996/97. Both attenuation and dispersion were significantly reduced compared to coaxial cables. The rise time decreased from 180 nsec to 100 nsec. Some test-OMs were instrumented to transmit analog optical signals through the laser calibration fibers.

In 1997/98 the array was expanded by 3 test strings (strings 11-13) that used optical fibers on all OMs, both for calibration and for analog transmission of the PMT pulses. The main goal was to evaluate this technology under realistic conditions and to compare various optical fibers and optical connectors. Electrical transmission was installed as a backup. The PMTs were operated at a gain near 10^9. Figure 1 shows the functionality of these OMs. The fiber transmitting the calibration signal from the surface penetrates the OM. The light is directed to the diffuser ball via an optical splitter. The PMT anode current is converted into an infrared signal by an LED operating at 1300 nm, which is transmitted to the surface where it arrives essentially without dispersion.

Optical fibers, connectors and penetrators are more vulnerable to the high pressures when the hole refreezes than their electrical counterparts. For electrical connectors, the losses were close to 10% for string 1-4 but, after using penetrators with an appropriate shape, this value was reduced to 3% for strings 5-10. The failures of the optical connections have been, and still are, about 10%, although half of these operate satisfactorily with attenuated amplitudes.

The advantages of the optical analog transmission compared to the electrical analog technique are:

- High bandwidth (rise time of 5 nsec)
- Good double pulse resolution
- Higher dynamic range due to lower HV,
- No correction for amplitude-dependent time slewing needed and, therefore, easier calibration.
- No pickup of electromagnetic noise, no cross-talk
- Operation during drilling operations is possible

The analog fiber optic signal transmission became the default technology for the expansion to AMANDA-II.

Experience with AMANDA II has shown that before establishing analog fiber optical transmission as an adequate technology for IceCube, the following improvements are required: *i)* the light transmitter, whether LED or laser diode, has to be operated with an independent and adjustable bias current, and local electronics in the OM must provide that function, and *ii)* a higher current-to-light conversion factor must be achieved.

This resulted in the dAOM (digitally controlled Analog Optical Module) concept. An alternate technical approach is the DOM (Digital Optical Module), which represents a departure from traditional AMANDA techniques because the digitization of the PMT signal is performed under ice. Figure 2 gives a functional description of both. Both technologies were deployed and tested with the1999/00 completion of AMANDA II.

Figure 2a) shows the dAOM. A DC/DC converter in the OM converts the 60V-supply voltage to the low voltage required to operate a microprocessor and other integrated circuits. The HV is generated by a Cockcroft-Walton generator, which is implemented in the PMT base. Voltages between the cathode and the first dynode, and between the first dynode and the anode, can be set separately. Slow control and reply signals are exchanged over the electrical cable. Also, an electrical analog signal is transmitted as a backup to prevent a complete loss of signal in case of damage to the optical fiber.

Figure 2b) shows the DOM, which makes use of the same HV generator as the dAOM. The waveforms are digitized in the OM and transmitted to the surface via the electrical cable. Determination of the t_0 offset and the synchronization between clocks at the surface and the OM are done using the electrical cable. There is no need for a fiber cable, which results in a savings of approximately 1.8 k$ per channel.

In the 1999/00 season six strings (strings 14-19) with 42 (41) OMs each were deployed. Of the 251 OMs, 189 use an improved passive, analog optical transmission system (with electrical back-up). The PMT anode current that drives the LEDs is now amplified in the OM by a transformer circuit. The HV is supplied from the surface. The rest of the OMs use DC/DC conversion inside the OM. Twenty-three OMs were of the dAOM type, with 10 using LED transmitters and 13 using Laser Diode (LD) transmitters. Forty-one modules are of the DOM type. The DOM modules have been equipped with LEDs and fiber optic read-out for implementation in the AMANDA-II DAQ. Both the dAOM and the DOM are operated at considerably lower voltages than the standard analog OMs. They did not show any significant gain loss during the year 2000. The advantes of the DOM concept are

- Reduced cost due to absence of fiber optics,
- Simpler calibration procedure,
- Higher dynamic range since the signal is digitized before it is attenuated by a fiber,

The weaknesses are a substantial amount of circuitry on the DOM board which must.have a long MTBF.

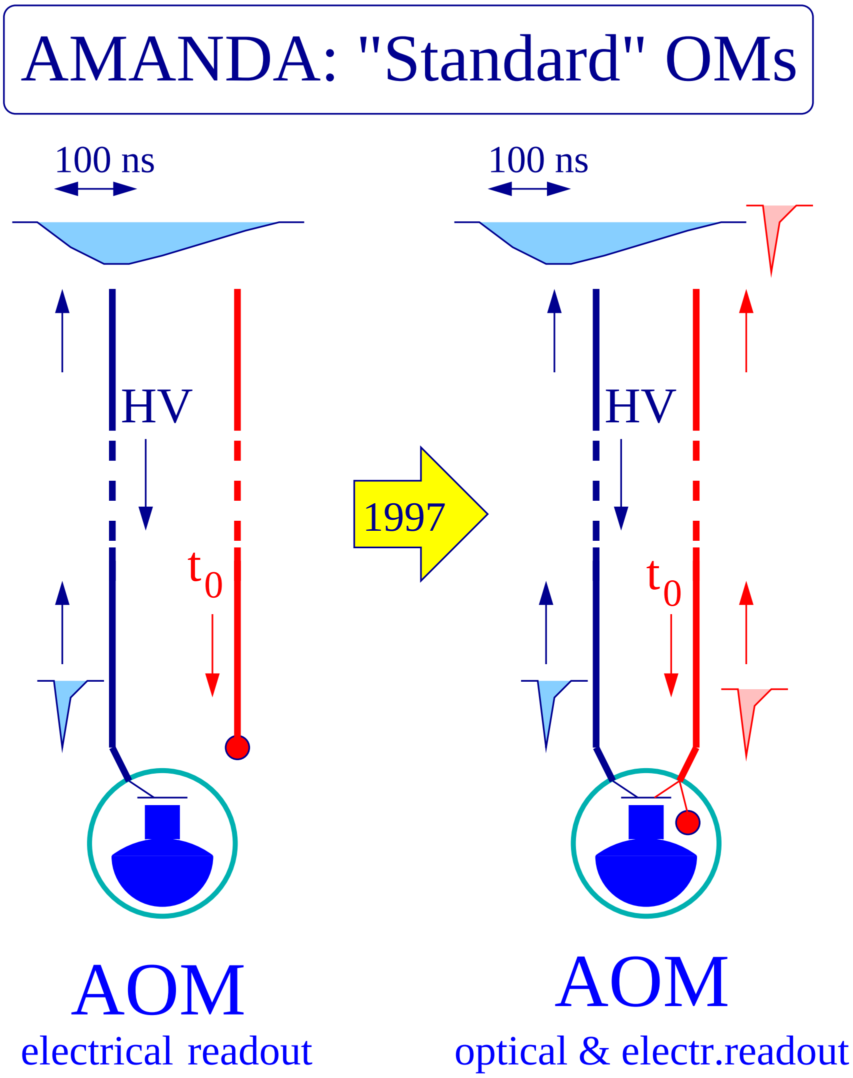

Figure 1: AMANDA Standard OMs: electrical signal transmission (left) and optical signal transmission (right).

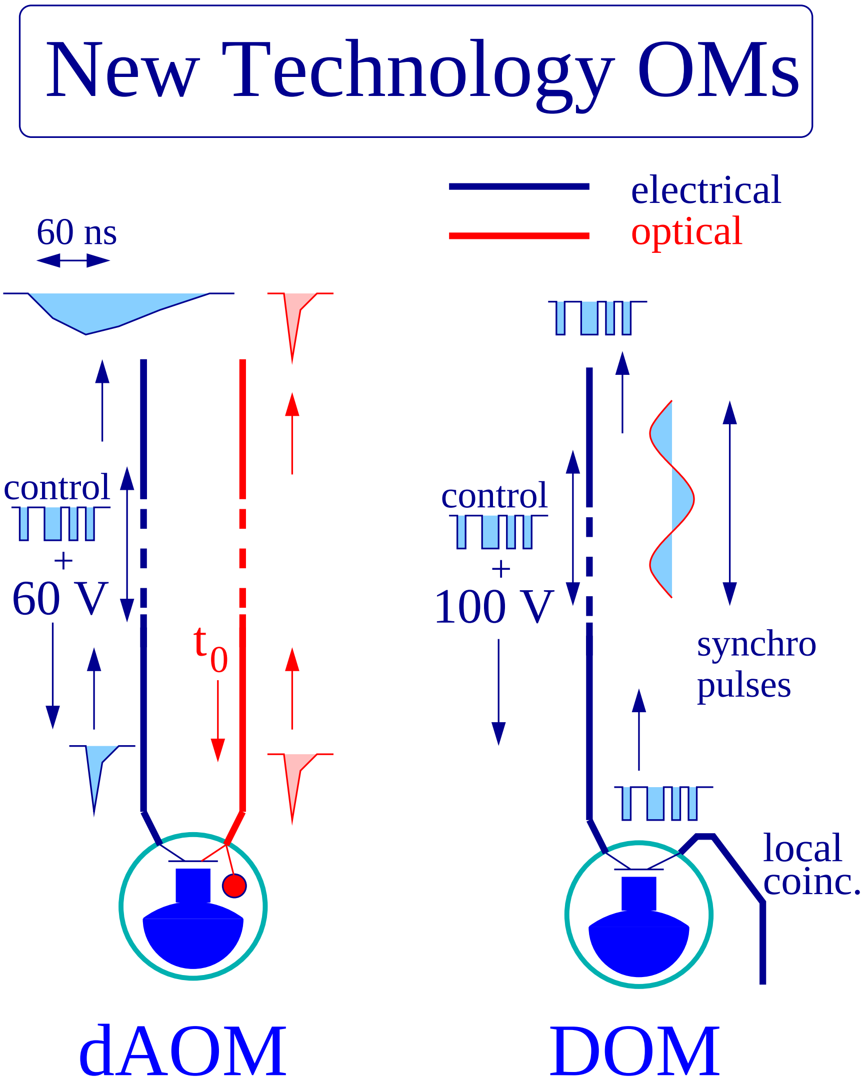

Figure 2: Digitally controlled Analog Optical Module (left) and Digital Optical Module (right).

2. Physics Requirements on Optical Module and Read-out

The physics requirements can be summarized as follows:

- time resolution of $\leq$ 5 nsec rms,
- waveform digitization with 200-320 Msps,
- dynamic range of $\geq$ 200 PE/15 nsec,
- noise $\leq$ 500 Hz.

While extensive simulations of the IceCube detector and its performance have been done, ongoing studies of improved analysis techniques using the operating AMANDA II detector could still suggest changes in these requirements.

Time Resolution: From operation and Monte Carlo simulations of the AMANDA detector we have learned that event rates and angular resolution of atmospheric neutrinos, as well as the efficiency for rejecting the background of atmospheric muons, do not depend on the accuracy of time calibration (σ_{t0}) and on time resolution (σ_t) provided they do not exceed 10 nsec. Deterioration of the detector's performance becomes apparent only above 15 nsec because scattering in the bulk ice dominates the time smearing. For the higher energy TeV—EeV neutrino signals in IceCube, the requirement may even be looser because the average distances over which photons are collected are increased and so are the effects of scattering. For electromagnetic showers, waveform digitization will be important, but a timing accuracy similar to the one required for muon tracks should be adequate.

Still, we consider a resolution of about 5 nsec a resasonable requirement. Better resolution results in better performance in rare cases. For instance, the detection of short muon tracks very close to strings where scattering is negligible, may benefit from better time resolution. Examples are the detection of nearly vertical neutrinos from the annihilation of dark matter particles and the detection of oscillating beams from an accelerator at a Northern site. One may also argue for improved time resolution on the basis of better reconstruction of the direction of electromagnetic showers induced by electron-neutrinos, or to separate close showers in tau-neutrino double-bang events. We note that these physics processes have not yet been simulated in sufficient detail.

Waveform data: Experience with the AMANDA detector for observing electromagnetic showers and high-energy showering muons has demonstrated the significant advantages of using the presently available "reduced" waveform information, like amplitude *and* time-over-threshold (TOT) rather than amplitude alone. Also, Monte-Carlo simulations for IceCube indicate that reconstruction algorithms making use of the full waveform information yield better angular precision, in particular for high energy tracks. Therefore IceCube reconstruction of distant tracks and very large high energy showers will require full waveform digitization. Waveforms should be recorded over a time interval of a few hundred nsec with $\geq$ 200 Msps for good timing and double pulse resolution. Recording beyond this time limit can presumably be done with lower sampling steps, i.e., a few tens of Msps, but this has not been demonstrated. High sampling rates over more than a microsecond would be attractive but is unlikely to be critical.

Dynamic Range and Linearity: The basic requirement is a dynamic range of 200 photoelectrons per 15 nsec. The interval 15 nsec represents the typical time duration of a single photoelectron pulse. This requirement is mostly driven by the reconstruction of energy and direction of very high-energy showers, typically 1 PeV. These produce measurable light over times of several hundreds of nsec. The amplitude and the waveform of the light pulses at the OM carry information on the position, the directionality and the energy of the shower. Figure 3 shows the PMT pulses from a 1 PeV shower detected by AMANDA II OMs at 3 distances between 45 and 167 m. The pulses were taken with an oscilloscope. The light output equivalent to that of a 1 PeV electromagnetic shower was generated by a N2 laser (337 nm).

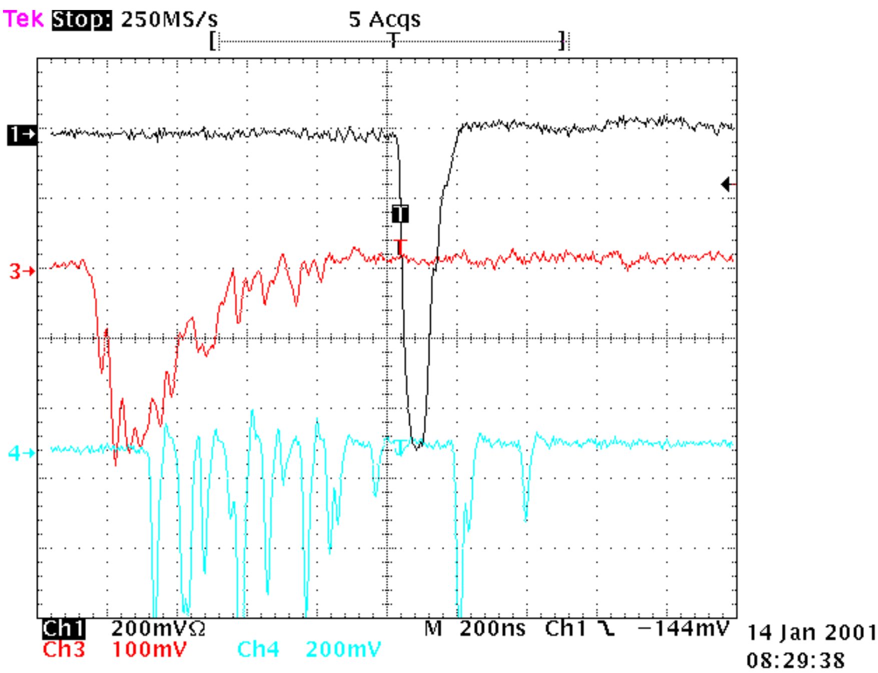

Figure 3: PMT signals from a laser pulse corresponding to a 1 PeV shower detected by AMANDA II optical modules at distances of approximately 45 m, 115 m, and 167 m.

Because of scattering of the photons in the ice, the shape of the photon pulse broadens with increasing distance from the point of emission. For example, at a distance of 60 m (~ half the string spacing in IceCube) a 1 PeV shower will generate a PMT pulse of 1000 photoelectrons with a FWHM of 100 nsec in forward direction of the shower. In backward direction one obtains 200 photoelectrons and a FWHM of 150 nsec.The example illustrates the directionality information contained in these data. Such an event will be observed by hundreds of OMs up to distances of more than 250m from the interaction. The maximum current in this event reaches 200 pe/15 nsec. Because of the greatly enhanced cross-section at 6.4 PeV (Glashow resonance), events at this energy may be preferentially observed, despite a falling spectrum.

The non-spherical shape of the initial light pattern can be exploited for good directional and energy reconstruction. While the number of PMTs that saturate increases with higher energies (10 PeV and higher), the total number of PMT that observe a signal grows by an even larger number. Based on evaluations of simulated waveforms, we conclude that a dynamic range of about 200 pe/15nsec would be sufficient for an accurate reconstruction of the relevant event parameters (energy, direction), even though OMs nearest the shower will saturate.

PMTs that are very close to the shower core will saturate at the maximum anode current that a PMT is able to sustain. In the case of the Hamamatsu 10" PMT, which is a prime candidate for IceCube, the anode current is limited to 80 mA, which sets an upper limit on the dynamic range in PE that depends on the gain. The operating conditions of IceCube suggest a gain of ~1 x 10^7, corresponding to a theoretical peak of ~500 PE/15 nsec. It seems reasonable to expect that, with a safety margin, the PMT can sustain anode currents of 200 PE per 15 nsec for intervals of about 100 – 200 nsec.

Recognizing the limitations on the dynamic range set by the PMT, and accepting a saturation of the 10% of the OMs near the shower core, we set the physics requirement for the dynamic range to 200 PE/15 nsec.

Absolute Amplitude Calibration and Stability: The absolute calibration of OMs should be known to better than 5%. The absolute calibration and measurement of the angular acceptance of the instrument will be performed by observing known signals from LEDs and from cosmic ray muon tracks. The measurement of shower profiles requires good calibration of *individual* OMs; see above. The dependence of the muon effective area of the operating AMANDA-B10 detector was shown to be close to, but less than linear (for the Amanda-B10 detector).

Noise Rate, Noise Rate Stability: Ice is a very low noise medium; the noise is dominated by radioactivity introduced by the OM itself. The noise rate in individual OMs is expected to be less than 0.5 kHz. This level of noise is readily eliminated by the use of local coincidences (a pair of OMs that are adjacent or separated by no more than one OM report). Supernova detection is made possible by the low noise rates and becomes impossible if the rates are as high as 2 kHz.

3. The Digital Optical Module

The motivation for the DOM is to capture and record all essential information provided by the PMT for the physical processes initiated by high-energy neutrinos. Data and simulations show that the information content at the PMT ranges from single photons arriving over a span of perhaps a microsecond to hundreds of photons arriving within perhaps 15 nsec, with a continuum of complex patterns between these extremes. The capture of this information requires high speed, high dynamic range waveform recording. To minimize signal degradation by noise sources, waveform capture and conversion from analog to digital format are realized close to the point of signal detection.

Figure 4 shows the block diagram of the DOMs deployed in January 2000. Figure 5 shows the readout scheme, still relying on the backup readout via optical fibers. Figure 6 shows a schematic view of the year-2000 DOM.

Figure 4:

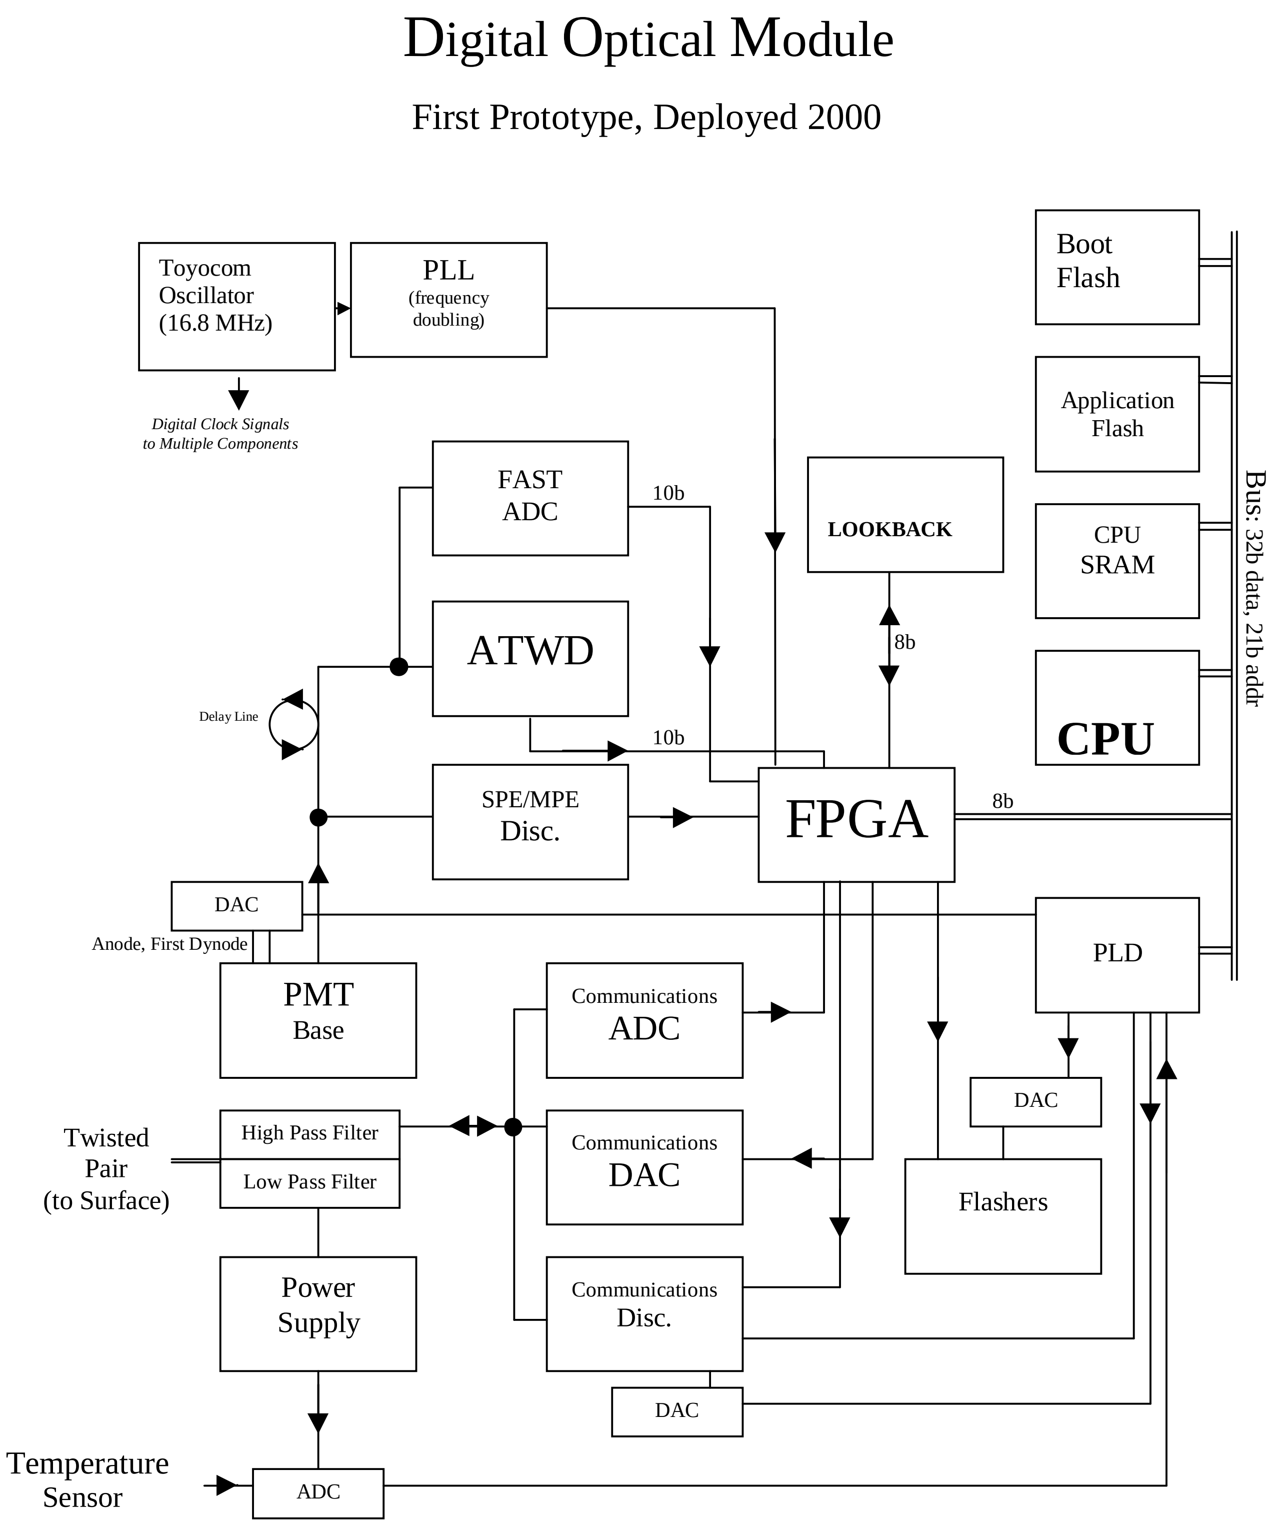
Digital Optical Module
First Prototype, Deployed 2000
Toyocom Oscillator (16.8 MHz)
PLL (frequency doubling)
Digital Clock Signals to Multiple Components
FAST ADC
10b
ATWD
10b
Delay Line
SPE/MPE Disc.
LOOKBACK
8b
FPGA
8b
Boot Flash
Application Flash
CPU SRAM
CPU
Bus: 32b data, 21b addr
DAC
Anode, First Dynode
PMT Base
Communications ADC
PLD
DAC
Twisted Pair (to Surface)
High Pass Filter
Low Pass Filter
Communications DAC
Flashers
Power Supply
Communications Disc.
DAC
Temperature Sensor
ADC

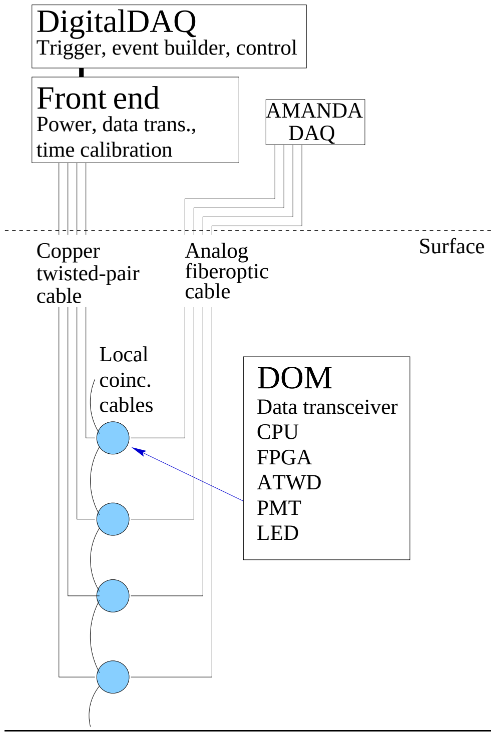

Figure 5: Schematic view of DOM string with connections between various components. Note that this is the year-2000 version with optical backup

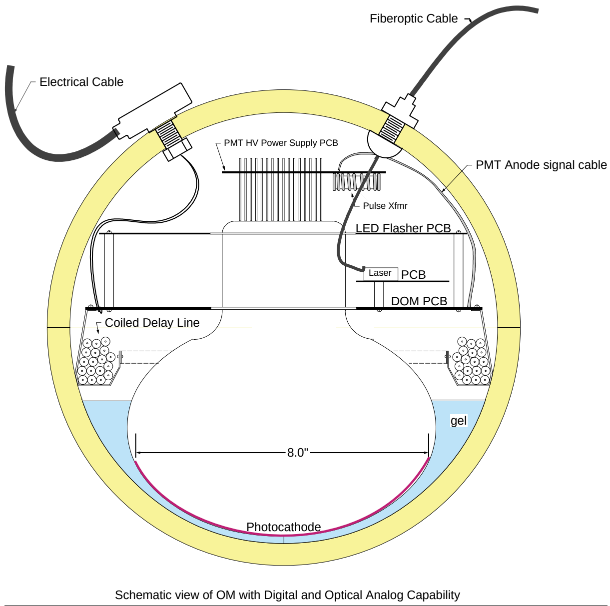

Figure 6: Scale drawing of a year-2000 DOM showing board placement

Waveform capture of every PMT signal over an interval of ~400 nsec is performed at sampling speeds of ~320 Msps. This capability is realized through a custom Application Specific Integrated Circuit (ASIC) designed at LBNL, the Analog Transient Waveform Digitizer (ATWD). The ATWD has four channels that synchronously record four input waveforms facilitating a multi-range linear PMT signal capture for high dynamic range. Analog signal capture is followed by 10-bit simultaneous conversion of all 128 samples/channel by an internal common-ramp Wilkinson ADC. We note the following technical implications of the DOM principle:

a) If the single-channel noise level is close to the intrinsic noise floor imposed by the active analog signal-processing circuitry (i.e., little or no spurious digital/extraneous noise), then the PMT may be operated at the lowest practical gain. Experience indicates that low-gain operation enhances PMT lifetime and may stabilize PMT gain drift behavior.

b) With minimal noise levels, signal conversion to digital data near the point of generation, and low PMT gain, there exists the possibility of achieving the largest practical dynamic range (in units of photoelectrons or PE) without using compression techniques. In the DOM this is realized through the use of multi-range linear signal capture using the PMT anode signal; a dynode signal tap could also be exploited.

c) The 2 – 3 km of twisted pair, which is providing power and control signals to the DOM, can easily support digital data transmission from the DOM at the rates expected. No degradation of data quality is expected to occur during data transmission from the DOM to the surface system.

d) Since all connectivity to the DOM is through the copper-wire twisted pair, reliability may be enhanced, and system costs may be substantially reduced.

e) The measurement and calibration of timing offsets introduced by varying lengths of twisted-pair cable can be performed automatically through software, with no manual interventions needed by pole personnel.

Waveform Capture

The DOM captures waveforms for each PMT “hit”, where a hit is defined as a pulse that exceeds a discriminator threshold held at a small fraction of a single photo-electron (SPE). The waveforms provide accurate information about the SPE pulse shape and its spectrum, which is essential for the determination of the PMT SPE efficiency. The waveforms are also crucial to the analysis of events in which many photoelectrons are liberated. In this class of events, optical scattering in the ice tends to produce a characteristic distribution of arrival times at the PMT, depending mainly on distance, and extending in some cases over hundreds of nsec. For especially energetic EM showers, the overlap of many PE signals in the rising edge of the ensemble signal provides the most sensitive measure of shower directionality. The four channels of the ATWD can provide the rough equivalent of a 14-bit ADC, with sample speeds ranging between 200 MHz and 1 GHz. An active delay line controls the sampling action; no high-speed clocks or high-power dissipation circuit elements are needed. A single DC control current varies the sampling speed.

In addition, a PMT signal is generated with ~30 nsec shaping, and digitized by a dedicated 10-bit FADC operated at 33.6 MHz. This information channel can be used to obtain records longer than the ATWD span, such as may occur during extremely energetic shower events.

Time-Stamping

In addition to PMT signal waveform capture, the DOM must also time-stamp each PMT hit to the required accuracy and resolution. This requires that an appropriate timing signal be available from a local clock in the DOM, which is calibrated (in frequency and offset) against a master clock on the surface. The method adopted for the realization of this capability is perhaps the most unusual feature of the DOM concept. Though unorthodox in concept, the circuit design, components, and operation are completely conventional.

Local Oscillator: In the DOM concept, a stable local oscillator is used to create a local time-base. During waveform capture, PMT hits are time-stamped according to this local time-base. This local oscillator is free-running, rather than phase-locked to the master oscillator, but periodically synchronized.. The reason for this apparently unconventional choice is the operational simplicity made possible by the remarkable stability of the Toyocom oscillator chosen for the DOM. All 50 samples of these oscillators, in tests at LBNL, displayed well-behaved characteristic drifts $\Delta f/f$ of ~ 6×10^{-11} per second.

This level of stability implies that a drift of 1 nsec occurs typically over an interval of ~16 seconds. If measurement accuracy is taken to be about 1 nsec, then a “time-mark” signal from the master clock need appear at the DOM only once every 16 seconds. Because the rate of local clock drift is so small, the time-base relationship between local and master clock may be maintained with a negligible impact on data flow along the cable. The partition of time-base function and data/control flow over the twisted pair is simplified. Transformation from local to global time for the PMT hit data is expected to occur at the surface, although there is no intrinsic reason that the DOM could not carry out the conversion. The frequency of the local oscillator is measured by sending “time synchronization mark” signals at known intervals in one direction, either from surface to the DOM or vice-versa.

Coarse Time-Stamp: The time-stamping process involves both a “coarse” and “fine” time-stamp to reach the desired accuracy and resolution. The normally quiescent ATWD is “launched” by the occurrence of a PMT pulse that exceeds the launch discriminator threshold. The actual ATWD launch is resynchronized to the next local clock edge. The analog signals are delayed by ~60 nsec before entering the ATWD to ensure that the leading edge is clearly visible. The capture of the local clock 32-bit value at the instant of the ATWD launch is the coarse time-stamp for that hit. The coarse time-stamp resolution is the period of the local clock. While the nominal frequency f of the Toyocom oscillators used in the DOM is 16.8 MHz, this basic frequency is doubled to 33.6 MHz, generating the desired internal clock for all DOM operations. The coarse time-stamp is in units of $(2f)^{-1}$ nsec, a period of about 30 nsec.

Fine Time-Stamp: The PMT signal’s leading edge, as recorded in the ATWD, will appear randomly over a ~30 nsec coarse time-stamp interval, since the PMT signal and the local clock are uncorrelated. The position of the leading edge in the ATWD record is the measure of the fine time-stamp. The optimum algorithm to extract the fine time-stamp is yet to be determined, but a simple approach based on leading-edge slope-extrapolation to baseline may be adequate.

Time-mark Distribution: Waveform digitization of an appropriate time-mark signal appears to provide both accurate time measurement and robust error detection. By utilizing a bipolar pulse for the time-mark pulse with step lengths of a few µsec, systematic errors can be largely cancelled out, and distorted pulses detected and rejected. The ADC needed in this case is an ordinary low-power 10-bit device operating at 33.6 MHz. By digitizing the baseline before the bipolar pulse as well as the pulse itself, the zero-crossing point can be found with negligible error from digital or analog offsets; also, distorted pulses can be recognized and rejected. The integrating behavior of the cable contributes a beneficial noise reduction, permitting a significant resolution enhancement when many samples of the waveform are used to find the baseline and zero-crossing point. It is clear that the circuit design must ensure that the edges of the time-mark pulse be very well correlated to the clocks. Sending time mark signals in one direction at regular, precisely known intervals and recording their arrival times is the method by which the frequency of the local clock in each DOM is determined relative to the frequency of the master clock on the surface.

Cable Length Timing Offsets

The measurement of the offset (or phase) of the local clock in the DOM requires an additional step, since the distribution of the time-mark signals to the DOMs will include a delay of 8 - 16 microseconds, depending upon the distance from the counting house and the depth of DOM deployment. The time delay associated with signal propagation and recording must be known so that the timing (or calibration of the clock) of each DOM is accurate to <5 nsec relative to a single master clock on the surface. The key idea in the cable delay determination is to send a timing signal not only from the surface to the DOM, but also, after a known delay, an identical signal from the DOM to surface. If the circuits generating and receiving these signals are identical, and the analysis of the signals (e. g., for a zero crossing) are identical on surface and in the DOM, then the cable delay is simply one half the measured sum of the propagation times in each direction, after subtracting the programmed delay in the DOM.

In practice, the DOM must receive, time-stamp and digitize a typical time-mark signal. Next, the DOM will wait for a predetermined number of local clock cycles to let the cable become quiet, and then transmits an identical time-mark signal back up the cable. At the surface, the front-end circuitry must capture, time-stamp, and digitize the return time-mark signal generated by the DOM. As mentioned above, it is essential that the surface front-end circuitry be identical to that in the DOM for this function to display the needed symmetry. The zero-crossing times of the time-mark signals are calculated from the digitized DOM and surface time-mark waveforms. In essence, the arrival times of successive pulses that are sent in the same direction determine the clock frequency; the round-trip time determines the clock offset.

Provided they are the same in the surface and DOM electronics, the propagation delays in the active electronic circuitry do not need to be determined separately - they are effectively included in the measurement of the cable delay. The method is nearly independent of the shape and rise-time of a timing pulse and independent of the algorithm for defining a time reference on that pulse, since the variations in delay associated with these choices are measured as a part of the calibration procedure. In particular, the method is insensitive to rise-time variations due to differing cable lengths.

This "**Reciprocal Active Pulsing**" or RAP method has the additional important advantage in that the procedure can be initiated under software control, and requires only a small fraction of a second to complete. The procedure may be easily repeated as many times as needed to ensure reliability and precision. Commissioning and monitoring are facilitated by this capability.

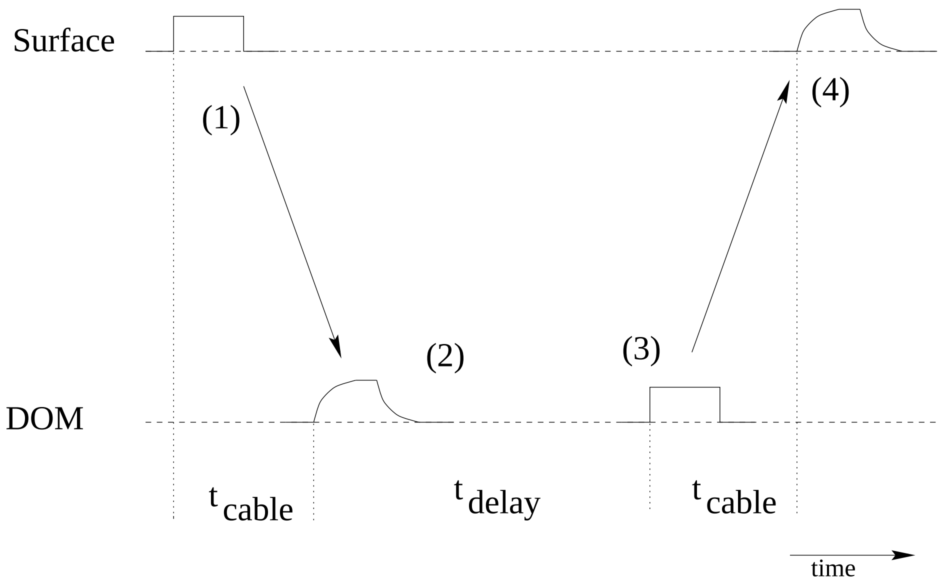

Figure 7: The "reciprocal active pulsing" (RAP) method of mapping local time within a DOM to a master clock time on the surface. (1) Surface sends pulse down the cable at time t=0, and records the master clock time. (2) DOM receives pulse at t_{cable}, and records the local clock time. (3) After a fixed time t_{delay}, the DOM sends a pulse identical to the original surface pulse up the cable. (4) Surface receives this pulse and records master clock time; total time elapsed = $(2 \times t_{cable} + t_{delay})$. t_{cable} can now be calculated, and the master clock time corresponding to the local time recorded in step (2) is known. Electronic hardware and software for transmitting, receiving and analyzing time calibration pulses must be identical at both ends of the cable.

Local Coincidence

This capability has been incorporated into the deployed DOM string, and is intended to reduce the upward flow of noise hit data from the DOM by at least two orders of magnitude. PMT noise hits are "isolated", in the sense that there is no causally related hit in any nearby OMs. Conversely, physical processes of interest produce strongly correlated hits in nearby OMs. As implemented in the DOMs, local coincidence is designed to require either nearest-neighbor or next-to-nearest-neighbor correlations within an adjustable time window, 300–500 nsec. Retaining this capability in the IceCube design has a very small impact on the electronic circuitry. The main impact is the need to provide two-wire links between adjacent DOMs. This connectivity was realized in the deployed DOM string with separate short cables that connect between the main cable and the DOMs. With local coincidence, the DOM penetrator has six wires rather than two. Studies with AMANDA data and simulations indicate that essentially no events are lost by the local coincidence requirement. However, isolated hits are, by design, not transmitted, and all real isolated hits would be lost to analysis. In many AMANDA analyses, isolated hits are commonly eliminated by cuts since they are very likely to be noise hits, so the scientific impact of isolated hits may ultimately be negligible.

At present, it is assumed that a local coincidence requirement is not imposed *a priori*; however, the capability could be implemented and turned on/off at will.

4. Experience with the DOM System from the 00-01 Pole Season

This section is a shortened version of a longer report "Experience with the Digital Optical Module System from the 00-01 Pole Season, A Supplement to the IceCube Technology Document" , prepared at B.Stokstad, LBNL, February 17, 2001. It can be found at http://www-zeuthen.desy.de/nuastro/ice3_proposal

4.1 Summary

The digital test string (String #18) was deployed in January, 2000 and was operated continuously until November, 2000, the start of the 00-01 Pole season. This season provided a crucial opportunity to test and demonstrate additional features of the DOM system under field conditions. The experience gained with the DOM system at Pole this season, documented in the rest of this report, has been very encouraging. There are more measurements and tests that can be done to improve and optimize the performance of the system beyond what has been demonstrated initially during this Pole season. The Digital Optical Model system, so far, is meeting the requirements for IceCube.

The main results obtained are the following:

Time Calibration

The local oscillator in the DOM was calibrated with an overall accuracy of about 5 nanoseconds against a master clock on the surface by sending timing pulses to and from the DOM. A GPS clock and phase-locked-loop provide a universal time base for calibrating the local clocks. The frequency of the local oscillator, measured against the GPS clock, has a short term stability of 2.6 nanoseconds over a 2.6 second interval. Longer-term, unidirectional drift is less than a few nanoseconds over 15 minutes. Time calibration tests were made on all DOMs and the systematic errors in absolute time calibration were estimated by using different combinations of surface testboards and DOMs on the same quad cable and on different cables. The effects of cross talk on time calibration were investigated and appear inconsequential. The calibration measurements are reproducible and the process can be fully automated.

Photon Detection
Photo electron hits in a PMT have been time-stamped. Photons emitted from UV LED flashers in one DOM have been detected by another DOM. Waveforms (single photo electron and multi-hit) are observed and a single photo electron spectrum derived from them. A single (and quick) measurement of dynamic range yielded a value of 150 p.e.

High Speed Communication
The system has been operating for the past year at a reduced data rate (5 kBytes/s) because it was not possible to develop fast communications by the time of the January 2000 deployment. The system itself, however, is designed to operate at the higher communications rates (>20 kBytes/s) that are required to time stamp all hits at the ~1 kHz rate caused by noise photo electrons. This season, high-speed digital communication signals were observed with an oscilloscope at transmission rates up to 1 MHz, exceeding by factors of 2 to 3 the anticipated experimental rates. The nature of these signals and a study of the effects of noise and cross talk indicate that communication at adequate rates should be possible under operating conditions at Pole. Actual demonstration of fast communication and measurement of error rates should be achieved in the next few months once the encoding, decoding, and filtering algorithms have been debugged.

Light Pulser Operation
The arrays of six UV (370 nm) LED light pulsers (flashers) in each DOM were used to perform geometrical calibrations of the AMANDA array and take data for studies of the optical properties of the ice. The flashers are visible at distances up to 200 m away and should be adequate for an IceCube string spacing of 125 m.

Remote Operation
The four testboards (a demo-DAQ) are fully operable over the Internet. Developmental work is continuing at LBNL, restricted only by the availability of satellite communications.

The recent work at Pole was done using a "demo-DAQ" consisting of four testboards that could (each) power and communicate with a single DOM.

4.2 Timing

Intrinsic Time Resolution

There are two main elements in the determination of the frequency and off-set of a local DOM oscillator relative to the master clock on surface. In essence, timing pulses sent in one direction at known time intervals can be used to determine relative frequency, and the round trip time of pulses sent in both directions can determine the offset. The combination of these two is the method of Reciprocal Active Pulsing (RAP).

Figure 8 illustrates the characteristic time resolution of the method. The results shown there are obtained as follows: Communications are halted, the line goes quiet, a timing signal is sent from the testboard to the DOM, and digitized by the communications ADC. Following a time interval of about 60 microseconds (precisely determined by the local oscillator in the DOM) a fresh timing pulse is generated *in the DOM* and sent to surface.

The testboard receives and digitizes this timing pulse. Later, during a digital communications transmission, the time-stamped and digitized timing pulse (received earlier by the DOM) is sent to surface. This round-trip measurement is repeated at intervals of about two seconds. The shapes of the received digitized pulses are fit (off line) to determine the time reference mark. A correction is made for the measured difference in frequency between the surface and local clocks so that the time intervals between successive sent pulses and between successive received pulses can be compared.

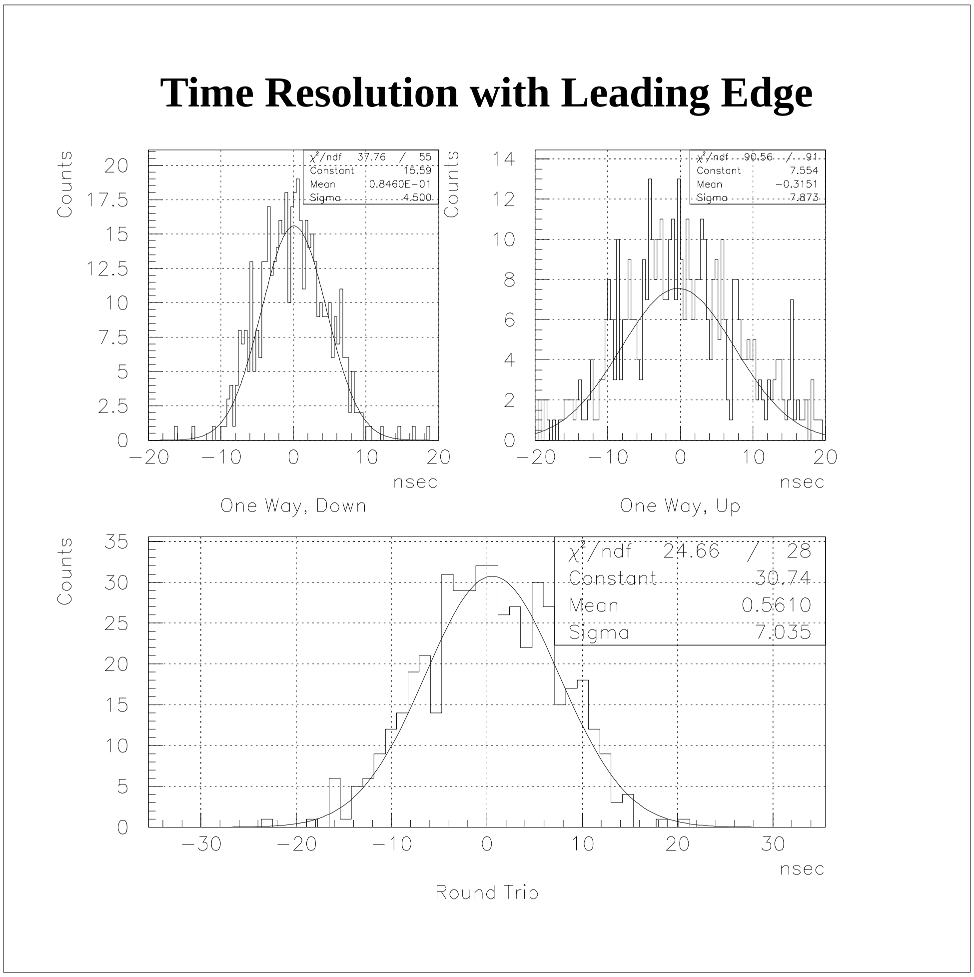

Figure 8: The residual deviations of time intervals from an expected value or an average value, as measured for pulses sent down, sent up, and for a round-trip.

The lower panel in Figure 8 labeled "Round Trip" represents the residuals from the mean of the round-trip signal propagation times. The RMS deviation is 7 ns, and has contributions from the determinations of the receiving clock times in the DOM and in the testboard. These separate contributions are indicated by the other two panels in Fig. 8. These "one-way" residuals are determined by analyzing the receiving clock times for pulses sent in the same direction at two second intervals. The residual in this case is the difference between the time the pulse arrived and the predicted time of arrival based on the elapsed time since the previous pulse was sent. That means that two time measurements are represented in one data point. The intrinsic time resolution for a single pulse sent in one direction is therefore smaller than the RMS value here by a factor of $\sqrt{2}$. There are several reasons for this difference, one of which is increased electrical noise at the surface.

The accuracy with which a PMT pulse can be time-stamped is essentially the intrinsic resolution for receiving a pulse. This is of the order of 3 ns for a pulse received at the DOM.

Most of these measurements were made with the PMT HV off. It was then discovered that the time resolution degrades when HV is on. The additional noise contribution from having the HV on is about 3.5 ns (to be added in quadrature with the intrinsic time resolution). We suspect that the HV DC-DC converter on the PMT base injects noise into the system; this is a matter that needs further study, as one would want to eliminate or at least minimize this source of noise in the next design. In the unlikely event that this degradation of resolution should prove unavoidable, the system would still meet the requirements for physics, as a 5 ns absolute timing accuracy is acceptable.

Clock Frequency

The ratio of the frequency of the DOM clock to the master clock is determined by comparing the number of "ticks" of the surface clock and of the DOM clock for a series of pulses sent in one direction (in this case, down). The ratio of the number of ticks on surface and below for a given interval between time signals gives the ratio of the frequencies of the two clocks. Figure 9 shows the quantity 1-f(DOM)/f(GPS) determined for a number of consecutive time intervals, each 2.6 seconds long during a measurement lasting 15 minutes. The mean value of $-1.1 \cdot 10^{-6}$ represents the (typical) one ppm difference between the Toyocom oscillator frequency and the 16.8 MHz frequency of the GPS-driven clock. The distribution of the frequency ratios about the mean value has a sigma of $1.3 \cdot 10^{-9}$. This corresponds to a RMS "clock drift" of 3.4 ns over a 2.6 second interval.

There does not seem (to the naked eye) to be any evidence for a longer term, i.e. extending over 15 minutes, unidirectional drift exceeding ~2 ns. Such a drift would appear as a non-Gaussian shape for the distribution. This clock stability is typical for DOMs that have been under power for at least one-half hour. This is the time it takes, after turn on, for the local oscillator to stabilize to this level. Clearly, the Toyocom oscillator, in the temperature-controlled environment of the Antarctic ice cap, is a remarkably stable device.

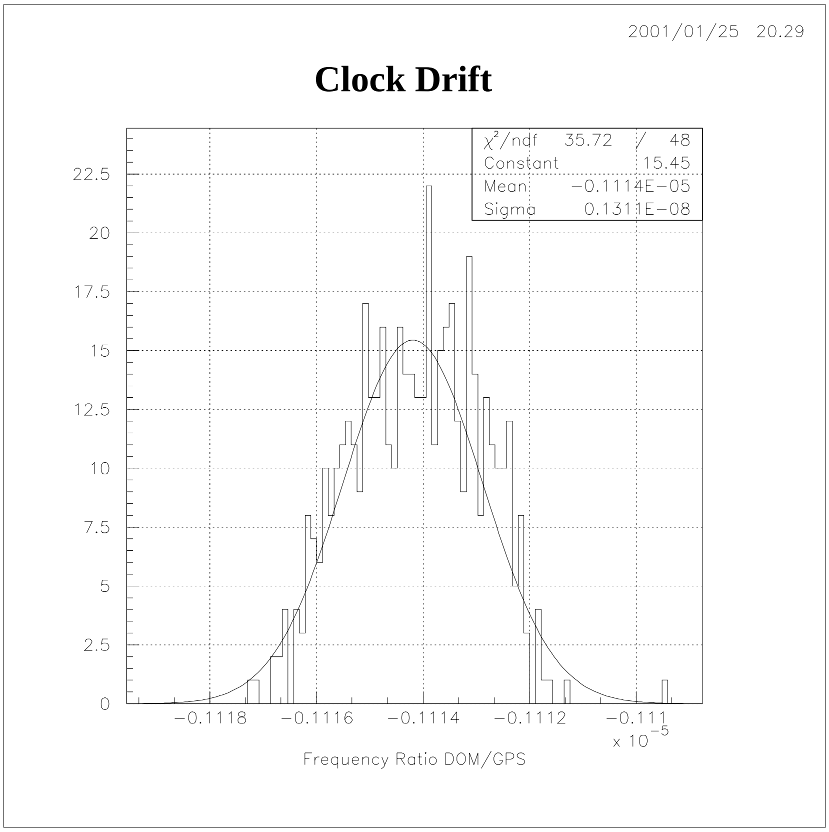

Figure 9 : Clock frequency ratio measurements made over a 15 minute interval. The abscissa is the deviation of the frequency ratio f(DOM)/f(GPS) from unity.

Clock Offset

The clock offset is determined in part by measuring the round-trip time for a pulse to travel to the DOM and for an identical pulse to travel from DOM to testboard. As explained in section 3, the one-way time (needed to determine the clock offset) is one-half the round trip time *provided* the generation, transmission, reception, and analysis of timing pulses is the same for pulses sent in either direction, i.e., there must be an up-down directional symmetry.

The determination of the time offset enables the determination of an effective electrical length (or pulse transit time) for the twisted pair cable and associated electronics. A rough verification of the timing method can be made by comparing the deduced electrical lengths with the nominal spacing of the DOMs along the cable. There are two DOMs attached to each twisted quad, and both DOMs are attached with the same length "pigtail" connector (12 m long). Thus, in addition to observing the change in round trip time for different quad cables, one can check whether both DOMs on the same quad yield the same electrical length. Figure 10 shows the measured one way electrical length (in microseconds) obtained for the DOMs on 20 different quad cables. A rough linear dependence is observed. In most cases there are several measurements at each data point (on the two DOMs of the quad, or several measurements on the same DOM) but the points cannot be distinguished when plotted on a microsecond scale

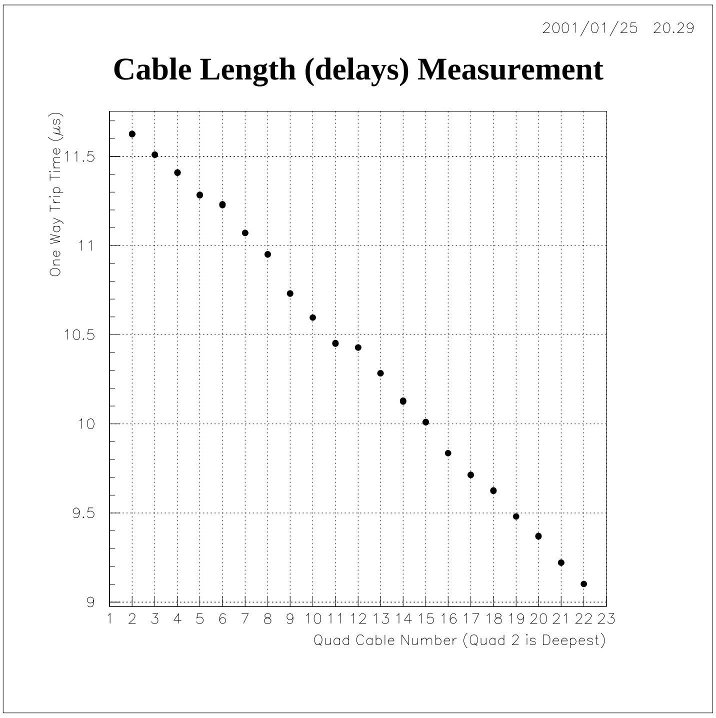

Figure 10: The propagation time for a pulse to travel from the surface to the DOM versus the cable lengths on String 18. The difference in length between successive quad cables is 24 m.

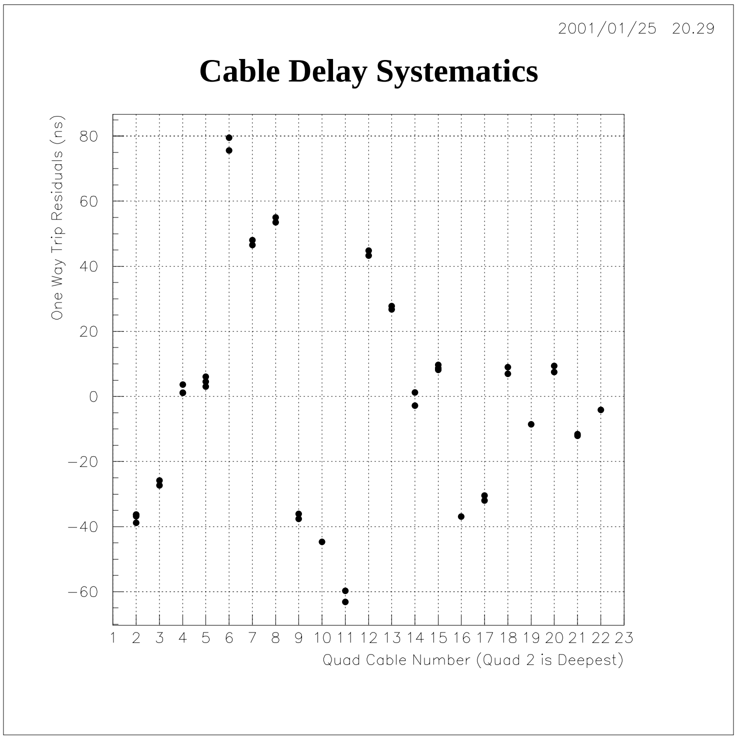

Figure 11: The residuals of the data and the straight line fit in Figure 10. Note that the ordinate here is in nanoseconds.

Fitting a straight line to the data in Figure 10 and determining the residuals of each measurement produces the results shown in Figure 11, where the ordinate is now in nanoseconds rather than microseconds. The first point to note is that where more than one data point exists for a quad, the points lie very close to each other - usually less than 5 ns apart. Measurements with the same DOM–testboard pair made several days apart yield the same value to within 5 ns. Measurements made with different testboards on the same DOM yield similarly close values. Measurements made with the same testboard on different DOMs on the same quad cable yield values within typically 5 ns. This indicates an upper limit on the experimental uncertainty of the method, since there could still be some physical differences in length even for two twisted pairs in the same quad. **It is clear that the RAP timing method is sufficiently accurate and robust for IceCube.**

It is important to note that the above study of cable transit times is only an illustration of the method, since the cable length is not the quantity of interest for reconstructing muon tracks. The effective electrical length will depend on the rise time of the timing pulse (therefore on cable length) and on the choice of the reference mark (leading edge or crossover). What is needed for tracking muons is the time calibration of the local clock - its offset and frequency ratio. The frequency ratios are independent of pulse rise time and choice of time mark. To the extent that up-down symmetry in hardware is achieved, the offset is also independent of these quantities. These demonstrations at Pole indicate that we are already able to provide a universal time stamp that is accurate to 5 ns or better.

For IceCube operation the time calibration measurements will be repeated automatically at time intervals of somewhere between a few seconds and a few tens of seconds. The optimum interval can be determined only after more detailed studies of the nature of the frequency drifts (short term period, long term period). The time it takes to do the calibration measurements is so small that, even for calibrations as often as every few seconds, it would be possible to shut down communications on an entire string (while still accumulating data in each DOM) in order to do time calibrations, one module at a time - with no loss of information.

4.3 Fast Communication

The present communication rate is about 5 kBytes/s and was adopted in 1999 because there was not enough time available to implement fast communications before deployment, even though the DOM and DAQ included the necessary hardware. For operation in a mode in which all hits are sent to surface, the estimated rate at which digital data must be transmitted to surface is 10 kBytes/s for a single DOM, if it is operating at a 1000 Hz noise rate. Having two DOMs on a single twisted pair requires a rate of 20 kBytes/s or 160 kbit/s. In order for data to be transmitted at this or higher rates without introduction of unacceptable error rates, the effects of noise and cross talk must be sufficiently small. The best demonstration of this, of course, is to send data at various rates and measure the error rate.

The DOM FPGA was programmed to encode a string of characters corresponding to about 1 kByte in total length and send that to surface at varying rates. These signals were received at the testboard and observed on a scope. However, the firmware for decoding could not be completed at Pole, which made it impossible to measure error rates. (Given the importance of demonstrating the *timing* capability of the DOM, no time could be devoted to debugging this firmware for fast communication before leaving for Pole.) The quality of the signals observed at surface and the relative absence of cross talk indicate to us that, once the encoding/decoding firmware is in place, fast communication will be demonstrated.

The demonstration that will be the proof of high-speed communication involves debugging the 8b-10b encoding/decoding firmware at LBNL, remote downloading of new firmware to both testboards and DOMs in ice, running of communications tests, and measuring error rates. This can be done from LBNL via the Internet - we hope to accomplish it in the next few months. In the meantime, data will continue to be accumulated by both the DAQ and the testboards with the slow communication protocol. For the future, also other protocols like HDLC are discussed.

4.4 Photon Detection

<u>Waveforms</u>

Sample waveforms from the high-gain channel of the Analog Transient Waveform Digitizer (ATWD) are shown in Figure 13. The top panel shows the pulse from a single photo electron. The obvious ringing that follows the pulse is due to underdamped shaping circuitry in a stage of amplification. (This was introduced by a last-minute gain change before deployment, forced by the need to optimize the relative gains for the fiber optic and electrical (ATWD) outputs.) This ringing has been eliminated by a simple digital filter, which was applied off line and which could even be incorporated in the DOM. However, in future designs the the ringing should not be present at all (as is the case for the KamLAND front end electronics, which also uses the ATWD).

The lower panel shows a case where a second photo electron occurred about 40 ns after the first one. Occasionally, quite dramatic examples of multiple photo-electron wave forms are observed.

<u>Single Photo Electron Spectrum</u>

Integration of the area of the s.p.e pulse produces the charge spectra shown in Figure 14. Spectra for hits recorded simultaneously in both the high gain (x5) and low gain (x1) channels are shown.

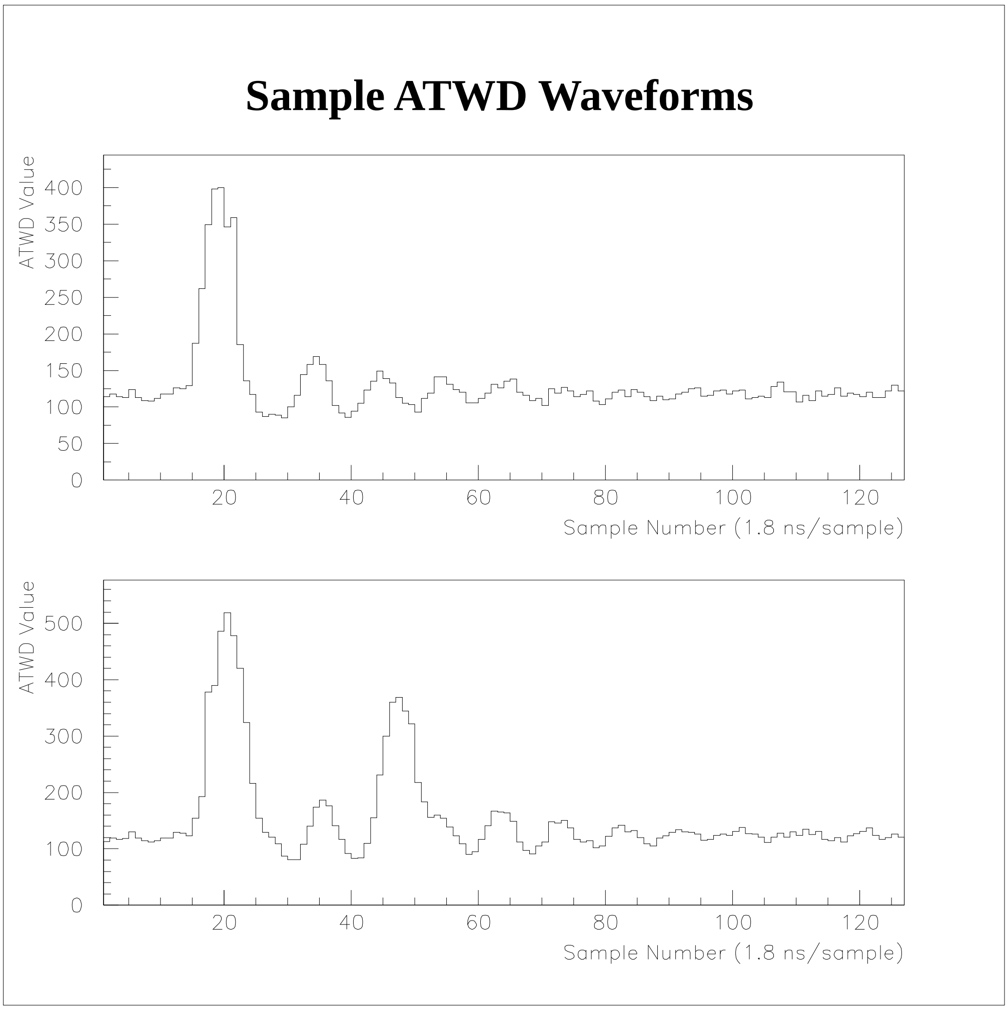

Figure 12: Photomultiplier waveforms recorded in the high-gain channel of the ATWD for: (upper panel) a single photoelectron and (lower panel) two photoelectrons separated by about 40 ns.

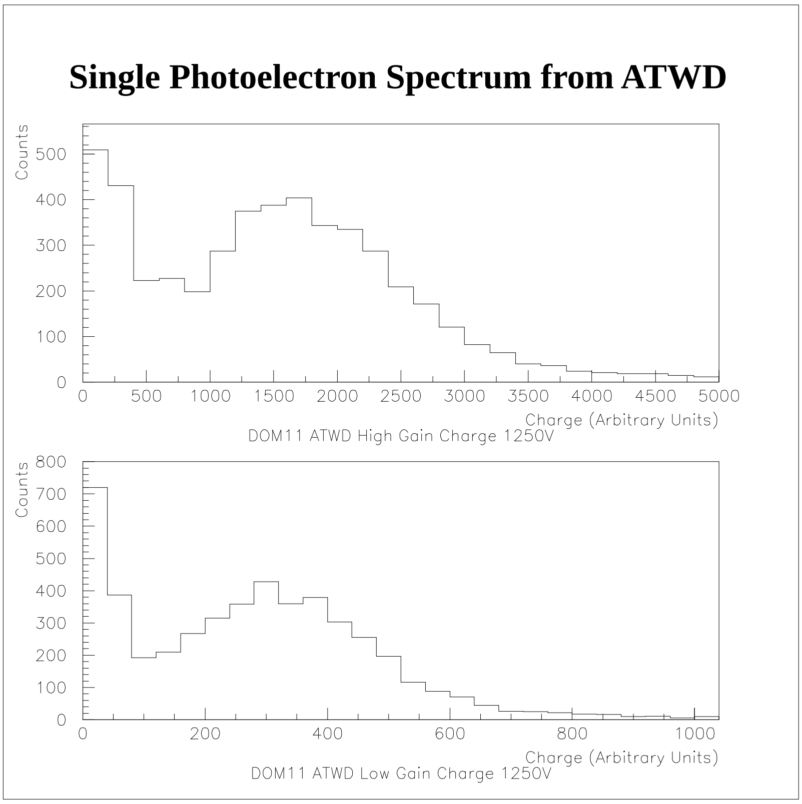

Figure 14. The pulse-charge distribution obtained from the area of the s.p.e. signals in the ATWD. Upper panel: high gain channel. Lower panel: low gain channel. The gains of the two channels differ by a factor of 5.

Dynamic Range

Measurements of the dynamic range of the various analog fiber optic systems were made using a YAG laser during this Pole season. The voltage on the same DOM for which the waveforms and s.p.e. spectra are shown above was lowered from 1250 V to 1150 V, where the peak in the s.p.e. spectrum was still adequate to permit a calibration in terms of photo electrons. Data were then accumulated in both high and low gain ATWD channels over a broad range of laser light intensity. The upper limit of linearity in pulse height for the ATWD in the low-gain channel appeared to be around 150 p.e. This was a quick measurement, on one DOM, but it suffices to show that achieving adequate dynamic range with the DOM will not be a problem. There is the possibility of having one more stage of gain in the third ATWD channel, and there is also the fast ADC (see the block diagram in Figure 4), which will record charge over a longer time than the ATWD. All these can be used to obtain a larger dynamic range than seen in this one quick measurement.

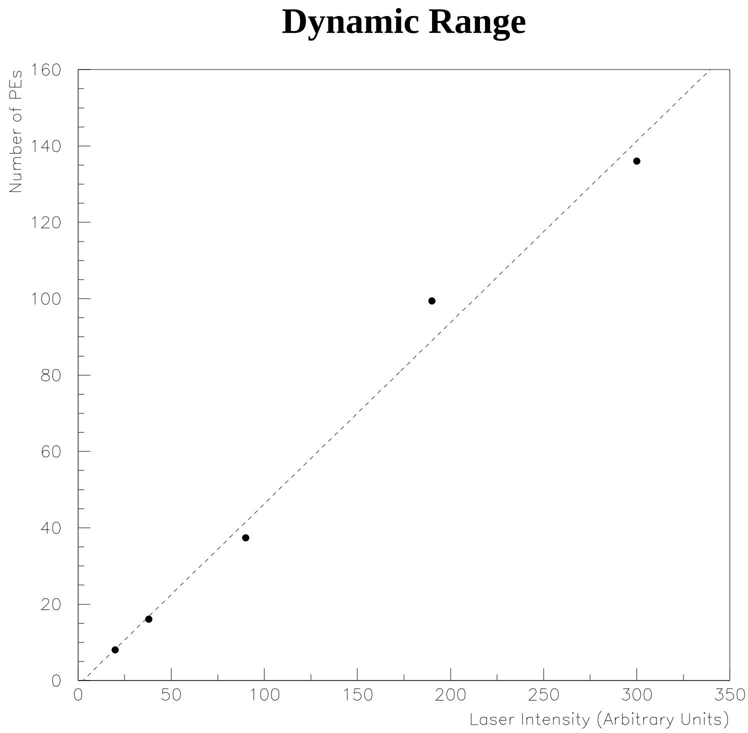

Figure 15: The charge signal in the ATWD as a function of the intensity of laser signal. The full range corresponds to approximately 150 photo electrons.

4.5 Operational Status of the DOM System at Close of the Season

Of the 41 DOMs originally deployed, 40 can communicate with surface. The one exception is a DOM that has an intermittent electrical connection (of undiagnosed origin). 37 DOMs can deliver digital waveforms to surface. The same number of DOMs also deliver an optical analog signal to the AMANDA DAQ via the fiber optic cables that were included on the digital test string. Three DOMs exhibit indications that a PMT base may have failed.

5. Future Work

The collaboration plans a series of tests and proofs to be performed in 2001. These tests will be finalized by the full digital read-out of string #18 and the reconstruction of muons with the help of string #18. We envisage the following steps:

Demonstration of Functionality:

a) Demonstrate measurement of photo-electron transit time in the PMT. This will use the single, collimated LED mounted on the DOM board

b) Complete the demonstration of high speed communication. Measure error ratesas a function of communications speed.

b) Demonstrate local coincidence: Operate a pair of adjacent DOMs connected to a pair of testboards in local coincidence mode. Measure coincidence rate as function of coincidence time window, etc.

Optimization of Performance

a) Investigate origin of extra noise when PMT HV is on.
b) Try different timing pulse shapes and fit algorithms to improve time resolution
c) Characterize frequency drift in Toyocom oscillators.
d) Determine optimum interval between time calibration measurements.
e) Introduce master Toyocom oscillator and calibrate against GPS clock.

Physics

a) Detect down-going muons.
b) Use GPS-based universal time stamp to correlate hits/ events in testboards with events AMANDA.
c) Record and analyze waveforms to determine frequency of complex wave forms.
d) Search for high energy events and record waveforms.

Most of this program has to be realized within the next seven months. It requires that engineering and software efforts as well as cost are shared among the collaboration members.

6. Conclusion

The DOM system, a prototype of what could be used for IceCube, has been shown to perform at a level which is already adequate to meet most of the important requirements. We expect that work in the near future will show that all of the requirements can be met, and that many can be exceeded.