Proceedings of ICRC 2001: 1237 © Copernicus Gesellschaft 2001

ICRC 2001

The IceCube detector

A. Goldschmidt, for the IceCube collaboration

Institute for Nuclear and Particle Astrophysics Lawrence Berkeley National Laboratory Berkeley, CA 94720

Abstract. IceCube is a proposed detector for high energy neutrinos, to be located at the South Pole. The present design calls for the deployment of 4800 optical modules between depths of 1400 m and 2400 m, positioned on 80 strings of 60 modules each. The enclosed volume will be ~1 km^3. The construction and operation of a detector of this scale at a remote location pose a number of challenges for deployment, data acquisition and data handling. An overview of the plans for each of these is given. In particular, time calibration results obtained with a digital string in the AMANDA array indicate a local timing accuracy of 5 ns can be achieved with IceCube modules.

1. Introduction

IceCube is a kilometer-scale detector for high-energy neutrinos, which is proposed for construction at the South Pole (IceCube Collaboration, 1999). There are many scientific reasons for a neutrino telescope of this large size, the main one being the relatively low event rates predicted for candidate point sources such as Active Galactic Nuclei and Gamma-ray Bursts, and for the diffuse flux of high energy neutrinos expected from aggregates of sources. The prospects for science and the expected performance of IceCube are presented at this conference (see the contributions by C. Spiering and by M. Leuthold, these proceedings). This paper describes some of the technical features of IceCube, its construction and operation. The deployment and reliable operation of a detector this large at a very remote location pose interesting technical challenges in meeting the goal of obtaining the highest quality data. We discuss here the advances in drilling technology, in data collection, and in data handling that are envisioned for IceCube.

There is, of course, a large body of experience available from the construction and operation of AMANDA (see the contributions by S. Barwick and by R. Wischnewski, these proceedings), which provides a sound basis for IceCube. Drilling to depths of 2.4 km, deployment of strings of optical modules, the acquisition of data over km distances, and the observation of atmospheric neutrinos have all been demonstrated with AMANDA over the past seven years at the South Pole. Nevertheless, the construction of a detector with four times as many strings and eight times as many optical modules (OMs) over a comparable period of time requires advances in technology for drilling and deployment. Comparable advances are also required for data acquisition and operation, and the amount of data to be processed will increase by an order of magnitude.

The present plan for IceCube is that it will consist of 80 strings, each with 60 optical modules located between 1400 and 2400 meters depth. The vertical spacing between OMs is 17 meters and the spacing between strings is 125 meters. The exact arrangement of the strings and spacing may change with more detailed optimization studies, but these figures adequately characterize the array for the present.

2. Advances in drilling technology

The melting of holes in the South Pole ice that are 60 cm wide and 2.4 km deep is a major IceCube engineering task, and the fuel requirements for this operation dominate the logistics needs of the experiment. Over the years, the drilling equipment has improved dramatically; larger power and flow-rate availability, and the use of larger power plants and hoses have led to reductions in the time and fuel needed to drill each hole. The present system uses 3.8 cm diameter hoses contained on several reels.

Correspondence to: R.G. Stokstad (rgstokstad@lbl.gov)

*The complete author list is found at the end of this volume.

2.1 Wotan

The upgrade to the current drilling system is called "Wotan.'' Wotan is a 5 MW, single-winch hot water drilling system using a single large 6.4 cm diameter hose that can reach depths of 2400 m in about 40 hours for a total fuel expenditure of 27,000 liters. Wotan is modular in design, being mounted on 13 sleds that can be towed to the required position and connected in a short time. All aspects of the Wotan system will be controlled from one central location.

The principal component of Wotan is the winch. It has a drum diameter of 2.3 m, flange diameter of 3.7 m, and is 5.4 m between flanges. It is large enough to hold 3000 m of 6.4 cm hose on a single drum. This will eliminate the current need to change reels during hole drilling, saving 10 hours per hole. It will be transported to the Pole in two (LC-130 Hercules) flights and assembled in the field. The modular design of the system will reduce the heating plant set-up time from four weeks to less than one week. Keeping major portions of the plumbing in heated spaces will also reduce considerably the chances of system freeze-up, and reduce the need for ethanol from 6500 liters to less than 1000 liters per season.

The Wotan system includes the main heating plant, pre-heating plant and water well, high-pressure pumping plant, water tanks, generators, auxiliary fuel, stand-by generators, workshops, and drill command post. These modular structures are towed from the storage area to the drill site and assembled using "quick connect'' techniques in about 100 hours.

3. Advances in electronics

The central problem in an array of photosensors like AMANDA or IceCube is to transport the time and amplitude information in a photomultiplier tube signal over distances of ~ 2 km without degradation. The first AMANDA strings used coaxial cable and then twisted copper pair to conduct analog signals directly from the PMT (operated at a gain of 10^9) to the surface. Although these signals were highly dispersed with correspondingly long rise-times, it was still possible to time the arrival of photons with accuracy sufficient to reconstruct the tracks of muons. The next improvement was to use fiber optic cable to transport the analog pulses. While rise times of ~5 ns were attained, the fiber optic connection would occasionally be damaged under the stress of the water freezing in the drill hole after deployment.

An alternative method for preserving signal information is to record the time of arrival of the photon, digitize its waveform in the OM itself and transmit this information using digital communications on copper twisted pair. This is the method that is planned for IceCube, with two OMs per single twisted pair.

3.1 The Digital Optical Module (DOM) system

The DOM system has several key elements. The Analog Transient Waveform Digitizer (ATWD) is a low-power custom chip that captures PMT anode pulses at a sampling rate of 200 - 800 mega-samples per second, with typical operation at 500 MSPS for a total capture time of ~250 ns across 128 samples. Each sample of the waveform is then digitized to 10 bits. A slow speed (33.6 MSPS) 10 bit ADC follows waveforms for longer times, up to ~5 μsec. Timing is provided by a very stable local oscillator that is frequency-doubled to 33.6 MHz. A coarse time stamp is provided at 30 ns intervals by a register incremented by the local oscillator. Digital communication with the surface occurs via ADCs and DACs at each end of the twisted pair. These communications ADCs and DACs also send and receive analog timing pulses, which are used to calibrate the free-running local oscillator in the DOM against a master clock at the surface. Finally, a separate twisted pair cable connecting adjacent DOMs makes it possible to operate the DOMs in a local coincidence mode, which essentially eliminates the ~500 Hz photo-electron noise rate present in each tube.

In January 2000, a string of 41 DOMs was deployed at Pole to develop and test the DOM system. The results reported here were obtained with this test string.

3.2 Time calibration

The calibration of the local oscillator involves determining its frequency and offset relative to the master clock. At regular intervals the transmission of digital data is halted and a timing pulse sent from the surface to the DOM. Comparing the time intervals measured at the surface for pulses sent down with the time intervals measured for their arrival below determines directly the ratio of the oscillator frequencies. This measurement must be repeated at intervals comparable to or less than the time over which the local oscillator may drift by a few ns.

The time offset is determined by sending pulses in both directions and measuring the round-trip time. (With optical fiber, a round-trip measurement can be made by detecting a reflected light pulse. However, with electrical cable, the long cable length, corresponding slow rise time and high attenuation of the reflected pulse require a different method.) The electronic components in the DOM make it possible not only to receive and record a timing pulse, but also to generate a fresh timing pulse in the DOM and send it to the surface. By designing the circuits so that the transmitted and received pulses have the same shape at the surface and below, the time for the pulse to be sent one way will be equal to one half the round trip time, less the "re-transmission" time interval in the DOM. Since there are clocks on the surface and in the DOM whose relative frequencies are

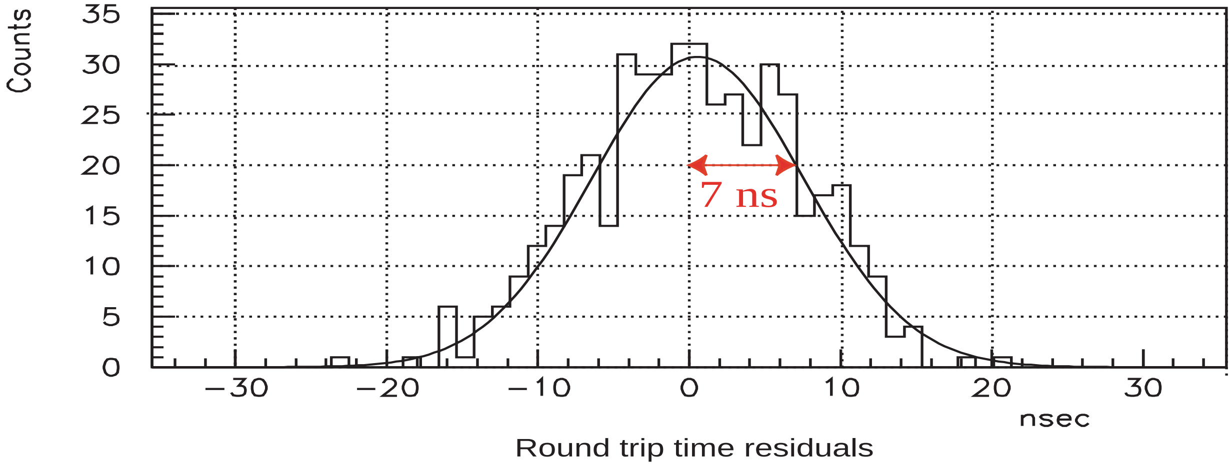

Figure 1. The residual values of round-trip times for time calibration pulses sent from the surface to the DOM, and from the DOM back to the surface. The round trip time is about 20 µsec for a DOM in the middle of the string.

known precisely, measuring the round trip time and the re-transmission delay is straightforward. Since the shapes of the received timing pulses are measured by ADCs on the surface and in the DOMs, one can verify and, if necessary, correct for any differences in the timing pulse shapes. This timing method provides an absolute time offset measurement for the oscillator in each DOM. The electron transit time in the PMT is determined using a single LED pulser mounted near the photocathode.

The residuals of individual round trip times from the mean time give an indication of the achievable timing accuracy. The standard deviation of 7 ns shown in fig. 1 includes contributions from sending pulses down as well as up. The error on the mean round trip time is much smaller than 7 ns. A more detailed analysis indicates that the intrinsic accuracy for a single measurement of a pulse received at the DOM is ~3 ns. Field measurements of systematic errors indicate the robustness of the method and an absolute timing accuracy of about 5 ns. Improvements in accuracy are anticipated.

3.3 Waveform recording.

The ATWD, which has four input channels, captures waveforms from the PMT and digitizes them. An example in which three photons arrived in rapid succession is shown in figure 2. Each bin, or sample, corresponds to an interval of 1.8 ns. Waveforms from the high gain channel and the low gain channel are overlaid. In this example, the first two photons arrived about 30 ns apart. Photons arriving as close as 6 ns will be resolved. In the case of a simple waveform, i.e., one containing only a single photo electron (s.p.e.), only the time and charge information will be transmitted to the surface. Determination whether a waveform is simple or complex will be made in the DOM. For complex waveforms, zero suppression and compression methods will be employed to reduce the bandwidth requirements.

4. System features

IceCube will use 25 cm diameter phototubes with mu-metal shields. Each DOM will be equipped with a set of bright 370 nm LED light pulsers to be used for calibration. These LEDs are so bright their pulses can be seen by modules up to 200 meters away. Adjacent DOMs will be connected via a separate twisted pair cable, which makes possible a level-one hardware trigger in the ice. Requiring a time coincidence (~1 µsec) between nearest or next nearest neighbors practically eliminates the ~500 Hz of s.p.e. noise that would otherwise have to be sent to the surface. Although simulations indicate that such a lowest level trigger would sacrifice little information, the system is nevertheless designed to send up all hits. At the

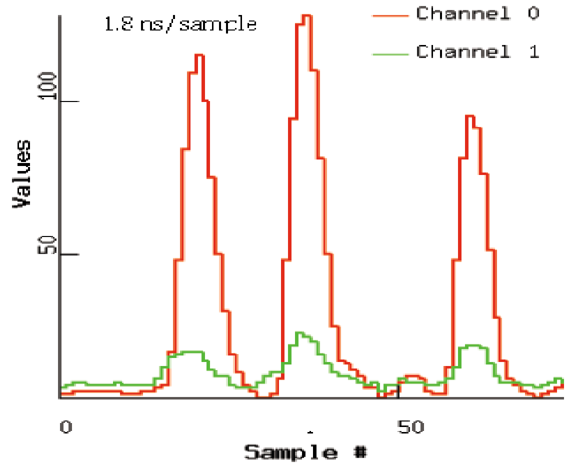

Figure 2. Waveform containing three separate photoelectrons, recorded by the ATWD.

surface, the time stamped hits will be examined by a processor for each string, which will determine whether a "string trigger" has occurred. This trigger might be set, for example, as 3 hits occurring in any 5 contiguous modules.

String triggers will be combined to form global triggers. Since the information at the surface is digital from the start, it is possible to use internet protocols for communication between string processor and front-end CPUs, each of which serves a small number of DOMs. The event builder will combine single hits with complex waveforms to form events that will be passed to local processors for online filtering and short-term storage.

Automatic time and amplitude calibration is necessary for a large array that can be accessed easily only during the Antarctic summer. Since the DOMs are controlled completely by software, this will be designed into the system from the beginning. Upgrades and improvements in the firmware and software will be possible as needed, since all firmware except the most elementary boot application can be downloaded to each DOM. Control of the system will be via Ethernet, such that local and remote (northern hemisphere) control can function equivalently. This should make it possible to operate IceCube during the Antarctic winter with the same level of staffing now employed by AMANDA.

Just as the air shower array SPASE has been a unique asset for AMANDA (see the contribution by X. Bai, these proceedings), the proposed surface array, IceTop, will provide cross checks of the geometry calibration, absolute pointing accuracy and angular resolution of IceCube. In addition, energy deposited by tagged muon bundles in air shower cores will provide an external source for energy calibration. By detecting air showers with energies from < 1 PeV to > 1 EeV, the surface array will also have the potential to veto some potential backgrounds for large (>PeV) neutrino events. IceTop will consist of a set of 80 frozen-water tanks, each ~7 m^2 in area and ~ 1 m deep, located at the top of each IceCube string. Three DOMs located in each tank will have local coincidence triggering. The data from the IceTop DOMs will be handled in much the same way as the data from a set of DOMs on an IceCube string or strings. The IceTop "string processor(s)'' will establish and build air shower events that will be part of the IceCube data flow.

5. Data handling

IceCube will generate a large amount of data. At a trigger rate of 1000 Hz from cosmic ray muons and an estimated 1.5 KB of information per event, the yearly total from the trigger is 50 TB. Events arising from neutrino interactions will be fewer by a factor of ~10^6. Processing of data will occur at the South Pole and in the northern hemisphere. How this processing is divided will be governed by a number of factors in addition to the total amount of data. Satellite transmission from the Pole currently has a bandwidth of ~ 6 GB/day, with high-bandwidth data transmission occurring only for about five hours each day when the TDRS satellite is above the horizon. Interruptions in satellite transmission may occur several times a year. Data stored on magnetic tape is hand-carried from the Pole between November and February when flights are possible. Thus, it is clear that a substantial fraction of the down-going cosmic ray muons will have to be discarded by filters operating at Pole, and that all neutrino candidate events should be transmitted via satellite.

Limited satellite communications also affect the ability to look for time coincidences between events in IceCube and events reported by other detectors, Gamma Ray Bursts (GRB) being the prime example as these occur at a total rate of about two per day for both hemispheres. It is desirable therefore to store all data, including down-going muons, for ~2 days so that it is possible to select events in a short time window of, say, ten minutes during which a GRB occurred. To accommodate this, a disc-resident circular event buffer of ~400 GB capacity is envisioned. Participation in a "supernova watch" requires that notification be sent from Pole to the watch network in as short a time as possible. Some form of 24hr/7day low-bandwidth communication (such as a global satellite telephone) is needed for this.

Data handling at Pole requires a processor farm of ~25 CPUs to keep up with the 1 kHz trigger rate, assuming ~25 ms/event processing time/CPU for this first filter. In the northern hemisphere, full reconstruction of muon tracks for events passing the first filter will take ~750 ms/event and require a processor farm with ~100 CPUs.

6. Summary

The construction and operation of IceCube at the South Pole require advances in drilling, detector technology, and data handling beyond the achievements that have produced AMANDA, which is presently the largest operating neutrino telescope. However, experience with AMANDA and the R&D for IceCube indicate that these advances will be made and demonstrated in time for first deployments during the 2003-2004 austral summer.

7. References

IceCube Collaboration, "IceCube: a Kilometer-Scale Neutrino Observatory, A Proposal to the National Science Foundation," November, 1999

C. Spiering, et al., "Science Potential of the IceCube Detector," these proceedings.

M. Leuthold, et al., "Performance Studies for the IceCube Detector," these proceedings.

S. Barwick, et al., "Scientific Potential of the AMANDA-II High Energy Neutrino Detector," these proceedings

R. Wischnewski, et al., "Performance of the AMANDA-II Detector," these proceedings

X. Bai, et al., Calibration and Survey of AMANDA with SPASE," these proceedings.