IceCube: A Cubic Kilometer Radiation Detector

The IceCube Collaboration

IceCube: A Cubic Kilometer Radiation Detector

Spencer R. Klein, for the IceCube Collaboration*

Abstract— **IceCube is a 1 km^3 neutrino detector now being built at the Amudsen-Scott South Pole Station. It consists of 4800 Digital Optical Modules (DOMs) which detect Cherenkov radiation from the charged particles produced in neutrino interactions. IceCube will observe astrophysical neutrinos with energies above about 100 GeV. IceCube will be able to separate ν_μ, ν_e and ν_τ interactions because of their different topologies. IceCube construction is currently 50% complete.**

Index Terms— `IceCube, Neutrino, Ice`

I. INTRODUCTION

IceCube, shown in Fig. 1, is a 1 km^3 neutrino detector being built to record the interactions produced by astrophysical neutrinos with energies above about 100 GeV [1]. IceCube will observe the Cherenkov radiation from charged particles produced in neutrino interactions, using 4800 optical sensors attached to 80 vertical strings which are deployed in a hexagonal array.

IceCube shares many characteristics with the smaller, laboratory-scale detectors discussed at SORMA. It is a large, segmented tracking calorimeter that measures the energy deposition in segmented volumes of Antarctic ice. It can differentiate between the topologies for electron, muon and tau neutrino interactions. It also has very good timing resolution, which is used to accurately reconstruct muon trajectories and to find the vertices of contained events. The size of IceCube is well matched to the energy scale; a muon with an energy of about 200 GeV travels about 1 km in ice.

II. COSMIC RAYS

A major reason to build IceCube is to find the sources of high-energy cosmic rays [2,3]. Cosmic-rays were first observed almost 100 years ago by Victor Hess. Over the past decades, many experiments have observed the cosmic-ray energy spectrum and composition, from GeV energies up to 3×10^{20} eV. The flux drops rapidly with energy, reaching $1/km^2$/century at the highest energies. Cosmic-rays have a mixed composition containing mostly nuclei from proton to iron, with at most a small fraction of heavier nuclei and photons.

Despite the decades of effort, we still know very little about the origin of cosmic-rays. At energies up to 10^{15} eV, cosmic rays are strongly bent in galactic magnetic fields. They likely originate in our galaxy. Supernovae remnants are the most likely sources. Their strong magnetic fields and shock waves can accelerate charged particles.

Galactic magnetic fields are too weak to confine more energetic particles, which are thought to be primarily extra-galactic. Possible sources are active galactic nuclei (AGNs, galaxies with central supermassive black holes) which emit jets of relativistic particles along their axes. Or, cosmic-rays might be accelerated by the sources of gamma-ray bursts (GRBs). GRBs are believed to originate in the collapse of supermassive stars and/or mergers of black holes and/or neutron stars. Either of these sources may provide appropriate conditions to accelerate nuclei to ultra-relativistic energies.

The most energetic cosmic rays have limited ranges. At energies above about 4×10^{19} eV, cosmic protons are excited by collisions with the 3^0K microwave background radiation, creating a Δ resonance. The decaying Δ emits a lower-energy proton. This energy loss limits the range of more energetic protons to about 100 Megaparsecs. Heavier nuclei are photodissociated by interactions with the microwave background; this leads to a similar range limitation.

Further, all but the most energetic cosmic-rays are bent in the intergalactic magnetic fields and so do not point back toward their origins. At energies above 6×10^{19} eV, bending by interstellar magnetic fields may be tolerable. The Auger collaboration has found evidence that some cosmic-rays may point toward nearby (within 75 Megaparsecs) AGNs [4]. However, the Fly's Eye collaboration does not observe this correlation [5].

In the absence of definitive correlations, we must consider other messengers. TeV photons have been observed from some nearby sources, such as supernovae and some nearby

AGNs. At energies above a few TeV, photons interact with interstellar photons, forming e^+e^- pairs; like protons and heavier nuclei, these photons also have a limited range.

In contrast, neutrinos have very small cross-sections and so can freely travel cosmic distances. They are the only particle able to probe high-energy accelerators out to cosmic distances. Here, we focus on the neutrinos with energies above about 100 GeV which are most relevant for understanding cosmic-ray acceleration. These neutrinos are produced in $\pi^{\pm}$ decays, $\pi^{\pm} \longrightarrow \mu^{\pm} \nu_{\mu}$, followed by $\mu^{\pm} \longrightarrow e^{\pm} \nu_{\mu}\nu_{e}$, producing a 2:1 ratio of ν_{μ}:ν_{e}. IceCube cannot differentiate between ν and anti-ν, so we will lump them together. Over long distances, neutrino oscillations change this 2:1 ratio into a 1:1:1 ratio of ν_{τ}:ν_{μ}:ν_{e}. The charged pions are produced in incidental 'beam-gas' interactions, whereby the nucleons under acceleration interact with either gas or photons present in the accelerator. If cosmic-rays are heavier nuclei, ν_{e} may also be produced by nuclear beta decay of unstable isotopes produced in spallation.

The neutrino flux from cosmic-ray accelerators has been estimated by two methods. The first uses the measured cosmic-ray flux and the estimated photon and matter densities at acceleration sites. The second extrapolates the measured TeV photon flux to higher energies, assuming that the photons are from π^0 decay. That leads to an estimate of the number of $\pi^{\pm}$. Both approaches find similar neutrino fluxes, and both lead to a similar conclusion: that a neutrino detector with an area of ~ 1 km^3 is needed to observe neutrinos from astrophysical sources.

III. EARLY LARGE NEUTRINO DETECTORS

For obvious cost reasons, a 1 km^3 detector must use a natural detecting medium. One approach to such a large detector is to search for optical Cherenkov radiation from charged particles produced in neutrino interactions. Three media have been proposed: seawater, freshwater (in a lake), and Antarctic ice. All three have advantages and disadvantages. Water has a very long scattering length but relatively short absorption length. Seawater has high backgrounds from ^{40}K decays and bioluminesence, while the available freshwater lakes suffer from limited size. On the other hand, in ice, the scattering length is shorter than in water, and, once deployed, detector hardware is not recoverable. All three approaches have been pursued. The DUMAND collaboration proposed a large seawater detector back in the 1980's. Currently, the ANTARES, NESTOR and NEMO collaborations are working on detectors in the Mediterranean Sea. A Russian-German collaboration has built a detector in Lake Baikal.

Neutrino detection in ice was pioneered by the AMANDA collaboration. It requires a thick ice sheet, so AMANDA was built at the Amudsen-Scott South Pole station, where the ice is about 2800 m deep. The collaboration drilled holes in the ice using a hot water drill, and lowered strings of optical sensors before the water in the hole refroze.

AMANDA deployed its first string on Christmas Eve 1993, at a depth of 800-1,000 m. It was quickly found that the ice

had a very short scattering length, less than 50 cm. This was explained by small (50 μm) air bubbles in the ice. Fortunately, at the higher pressures present at ice depths greater than 1400 m, these bubbles collapse. With this understood, in 1995-6 AMANDA deployed 4 strings with detectors mounted between 1500 and 2000 m deep. These detectors worked as expected, and AMANDA detected its first neutrinos [6]. This success led to AMANDA-II, which, by 2000 consisted of 19 detector strings holding 677 optical sensors. Since 2000, AMANDA-II has been recording about 1,000 neutrino events per year.
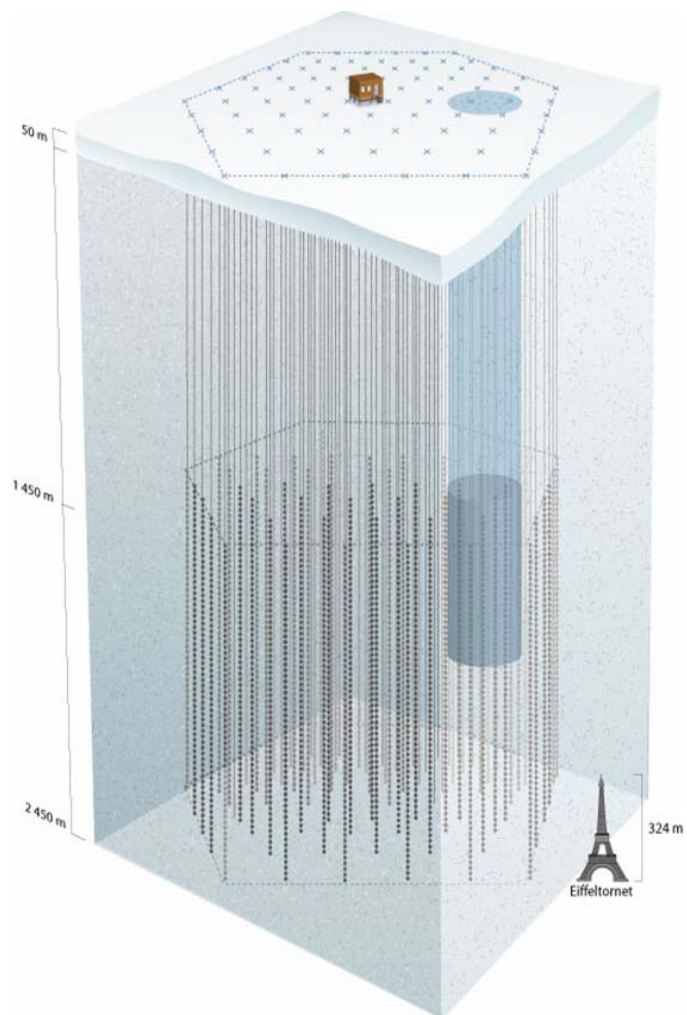
Fig. 1. Schematic of the IceCube detector, showing the 80 strings. The dark cylinder shows the volume of AMANDA.

However, despite this success, the limitations of AMANDA were becoming obvious. It was too small, and the technology did not lend itself to easy expansion. The AMANDA optical sensors consisted of photomultipliers with resistive bases in a pressure vessel. High voltage was generated on the surface, and analog signals were returned to the surface. Since AMANDA was a prototype detector, several transmission media were tried: coaxial cables, twisted pairs, and, later, optical fibers. The 2.5 km long coaxial cables and twisted pairs suffered from very limited bandwidth, while the analog optical fibers had a very limited dynamic range. Further, the

system was finicky, and not all of the optical modules survived the high pressures present when the water in the drill holes froze. Finally, AMANDA consumed considerable electrical power and required manpower intensive yearly calibrations. IceCube was designed to avoid these problems.

IV. IceCube hardware

In addition to being much larger than AMANDA, IceCube was designed to be much simpler to deploy, operate and calibrate. When it is complete in 2011, it will consist of 80 strings of photomultipliers, each containing 60 digital optical modules (DOMs). The strings are placed on a 125 m hexagonal grid. DOMs are placed on a string with 17 m spacing, between 1450 m and 2450 m below the surface. The surface electronics are in a counting house located in the center of the array.

Fig. 2. A block diagram of the IceCube main board electronics.

Each string of 60 DOMs is supported by a cable that contains 30 twisted pairs (each pair is connected to two DOMs in parallel), plus a strength member and a protective covering.

In addition to the deeply buried DOMs, the IceCube Observatory includes a surface air shower array known as IceTop. IceTop consists of 160 ice filled tanks with two tanks near the top of each string. Two DOMs are mounted in each tank to detect the Cherenkov photons from charged particles in the air shower. IceTop has a shower threshold of about 300 TeV.

The main task in IceCube construction is drilling holes for the strings of DOMs. This is done with a 5 MW hot-water drill, which generates a stream of 200 gallons/minute of 88^0C water. This water is propelled through a 1.8 cm diameter nozzle at a pressure of 200 pounds/square inch, melting a hole through the ice. Drilling a 2500 m deep, 60 cm diameter hole takes about 40 hours. Deploying a string of DOMs takes about another 12 hours.

Because of the Antarctic weather, the high altitude and the remote location of the South Pole, logistics is a key issue for IceCube. The construction season lasts from roughly

November through mid-February. Everything needed must be flown to the Pole on ski-equipped LC-130 transports planes.

IceCube construction began in 2004/5, when the first string was deployed. In 2005/6, eight additional strings were deployed, and, during 2006, data was taken with nine strings. In 2006/7, thirteen strings were deployed, followed by eighteen in 2007/8, leaving the detector half done.

V. DOM hardware

Each DOM contains a 10" (25 cm) photomultiplier tube and associate electronics in a 35 cm diameter pressure sphere. The electronics includes a Cockroft-Walton high voltage power supply, electronic timing calibration systems, light emitting diodes for photonic calibrations, and a complete data acquisition (DAQ) system. The PMTs are currently run at a gain of 10^7, with typical high voltages of 1300-1500 volts. An average single photoelectron produces a pulse about 10 mV in amplitude and 5 nsec width into the 43 Ω impedance of the DAQ system. The DAQ system is designed to record the arrival time of all detected photoelectrons, with a relative precision of better than 5 nsec, across the entire array.

A block diagram of the DAQ system is shown in Fig. 2. The central elements of the DAQ hardware are two waveform digitization systems, the Analog Transient Waveform Digitizer (ATWD) and the fADC ('fast' ADC). A digitization cycle is initiated by a discriminator trigger; the threshold is set at a voltage corresponding to about 1/4 photoelectron. When this happens, the FPGA will start ATWD and fADC digitization on the next clock edge. To make up for delay in the trigger circuit, the signal goes through a 75 nsec delay line before the digitizers. This delay line limits the system bandwidth to about 100 MHz.

The ATWD digitizer uses a custom switched-capacitor array chip. Each ATWD chip includes four parallel inputs, each with 128 capacitors. When launched, the system acquires data at 200 to 900 megasamples per second (MSPS); IceCube runs the ATWDs at 300 MSPS, providing 400 nsec of recording capacity. Three ATWD channels are run in parallel, with input gains in the ratio of 16:2:1/4, providing more than 14 bits of dynamic range. After acquisition, the voltages on the capacitors are digitized with 128 10-bit Wilkinson ADCs, each multiplexed to the four capacitors which acquire a single time sample. A fourth ATWD input (not shown) is used for electronics calibrations. Each DOM contains two ATWD chips. They are operated in a ping-pong fashion – while one is digitizing, the other is live; this greatly reduces the dead time. The fADC digitizer uses a 10-bit, 40 MSPS commercial ADC chip. When triggered, the system records 256 samples, covering 6.4 µs.

Each DOM also contains a 'flasher' board, which has 12 LEDs mounted around its edges. These LEDs are used for a variety of calibrations, measuring light transmission and timing between different DOMs.

The entire system is controlled by a 400k gate Altera Excalibur FPGA, which incorporates an ARM9 hard-core

CPU. The FPGA controls the data acquisition and digitization cycle, compresses (losslessly) and formats data for transmission to the surface, and oversees calibrations.
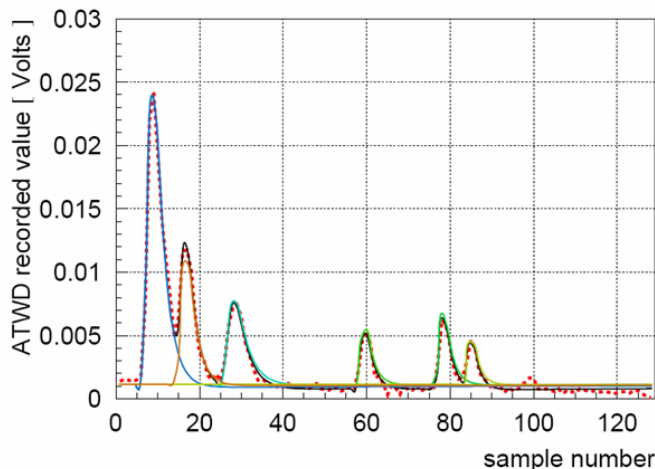

Fig. 3. The ATWD digitizer output from a typical event. Multiple pulses are shown. The waveform is decomposed into a list of photon arrival times, which is used for event reconstruction.

Data is transmitted to the surface via a single twisted pair, which also provides ±48 VDC (96 volts total) power. Each DOM consumes about 3.5 W. The cable also includes local coincidence circuitry, whereby DOMs communicate with their nearest neighbors; they can also pass messages onward. A more robust connector is used than in AMANDA, and a higher fraction of IceCube OMs survive 'freeze-in.' On the surface, the cables are connected to a custom PCI card in a PC; the remainder of the system is off-the-shelf.

IceCube DOMs have several operating modes. They are currently operating in "Hard Local Coincidence" mode: only data is only when a neighbor (either nearest or next-to-nearest) DOM also saw a signal within 1 μs. In "Soft Local Coincidence" mode, an abbreviated 'coarse charge stamp' is saved for isolated hits. It consists of the largest 3 fADC samples out of the first 16 samples. Saved data is formed into packets for transmission to the surface.

The system uses a 40 MHz system clock. Since this clock is used to record the hit times, a precision oscillator is used. The oscillator has frequency stability (Allen variance) of better than $\delta f/f < 10^{-10}$. Despite this accuracy, maintaining the required 5 nsec precision requires frequent synchronization.

Timing calibrations are performed automatically every few seconds (currently once every 3.5 s). During each calibration, the surface electronics sends a timing signal down to each DOM, which waits a few μs until cable reflections die out, and then sends an identical signal to the surface. To maintain the symmetry, both the surface and DOM electronics use identical DACs and ADCs to send and receive signals. With the symmetric setup, transmission times in the two directions are identical. Even though the 3.5 km cable transmission widens the signals to ~1 μs, the transmission time is determined to better than 3 nsec [7]. This accuracy is maintained across the entire array; it has been verified using

muons and artificial light sources. The software keeps track of the timing difference between the in-DOM oscillators and a surface based master clock, and appropriate corrections are applied to the data.
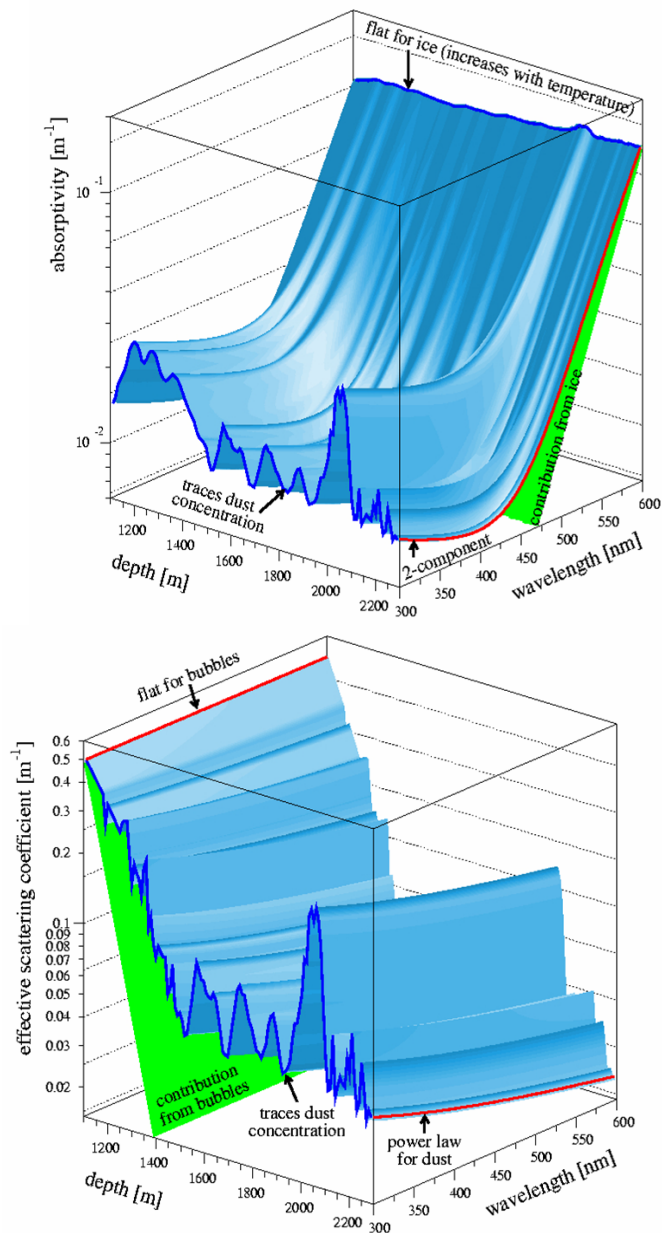

Fig. 4 Absorption (top) and scattering (bottom) lengths of light in South Polar ice, as a function of depth and wavelength.

Amplitude calibrations are done using an LED that is mounted on the main electronics board. It is flashed repeatedly at low intensity (<< 1 photoelectron in the PMT). A charge histogram is accumulated in the FPGA and sent to the surface, where it is fit to find the single photoelectron peak. This is done for a range of high voltages, and the system is then set to the correct HV to give 10^7 PMT gain. These calibrations are extremely stable.

Amplitude linearity calibrations take advantage of the 12 LEDs on the calibration board. The LEDs are flashed
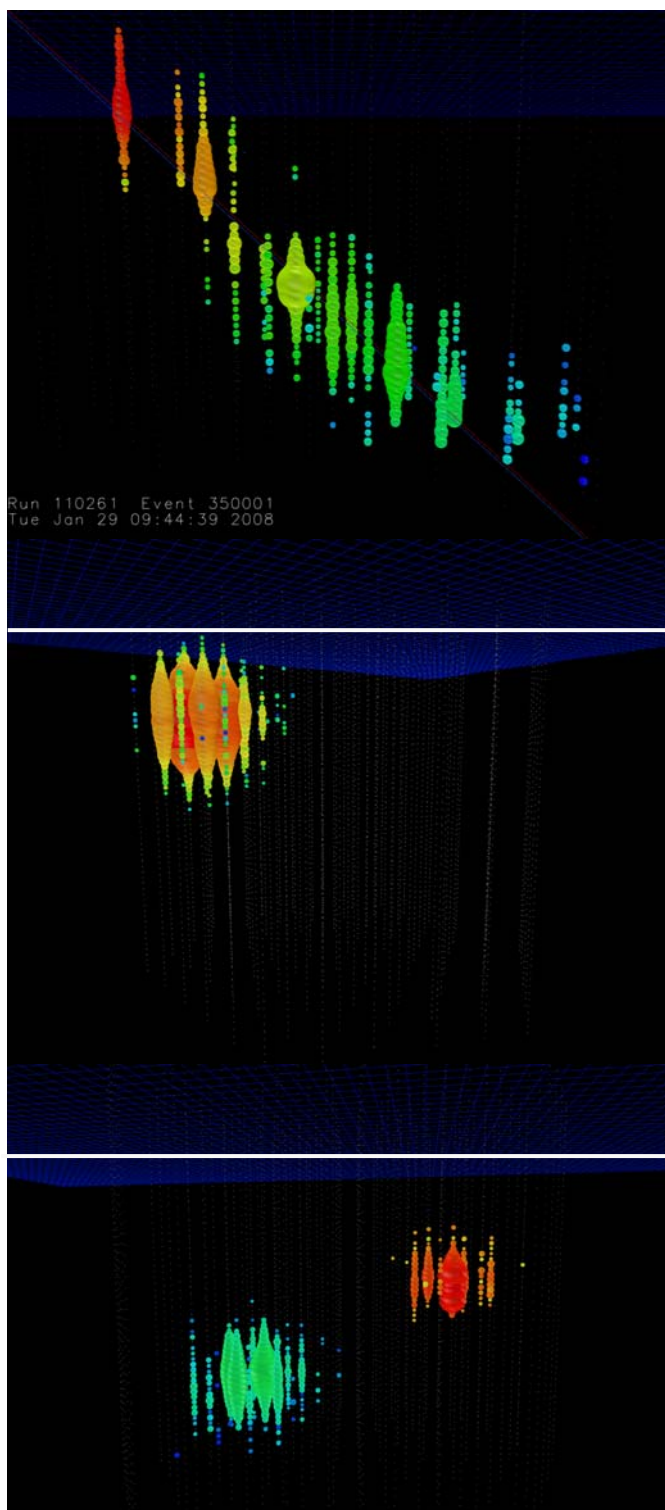
Fig. 5. IceCube event displays for (top) a muon or muon bundle (multiple muons) in IceCube 40 (the 40 string configuration running in 2008), a simulated ν_e (middle) and a simulated ν_τ (bottom). The latter shows the classical 'double-bang' topology. Each dot is from a single struck DOM. The size of the circles indicates the number of detected photons, while the color gives the time, from red (earliest) to blue (latest).

individually, and then together, providing a ladder of light amplitudes that can be used to map out the saturation curve.

One other critical requirement for the IceCube hardware is reliability. Once deployed, it is impossible to repair a DOM, so the system was designed for very high reliability. About 98% of the DOMs survive deployment and freeze-in completely; another 1% are impaired, but usable (usually, they have lost their local coincidence connections). Post-freeze-in, reliability has been excellent, and the estimated 15-year survival probability is 94%.

VI. THE ICE IN ICECUBE

The ice surrounding the DOMs is a critical part of IceCube. Both absorption and scattering are significant. Both are strongly affected by impurities in the ice. These impurities are a reflection of the impurities in the air when the ice was first laid down as snow.

This happened over roughly the last 100,000 years. Because of variations in the long-term dust level in the atmosphere during this period, as well as the occasional volcanic eruption, the impurity concentrations are depth dependent.

Much effort has gone into measuring the optical properties of the ice, using artificial light sources and in-situ measurements. In AMANDA and IceCube, studies have been done using LEDs and lasers that emit at a variety of wavelengths. By measuring the arrival time distributions of photons at different distances from a light source, it is possible to measure both the attenuation length and scattering length of the light. These measurements, although useful, suffer from a limited resolution in depth [8].

Higher resolution depth-dependence measurements of the ice properties come from a 'dust logger' which is lowered down a water-filled hole immediately after drilling. The dust logger shines a thin beam of light into the ice, and measures the reflected light [9]. This provides a measure of the ice properties on a depth scale given by the width of the emitted beam – a few mm.

Figure 4 shows the measured ice absorption and scattering distances, as a function of depth and wavelength. The effect of air bubbles at shallower depths is clearly visible, along with the effects of the broad dust peaks and the native absorption in the ice. Not visible are the very narrow peaks due to thin layers of dust produced by volcanoes.

VII. EVENT RECONSTRUCTION

Data collected by the DOMs and sent to the surface is time-sorted, combined into a single stream, and then monitored by a software trigger. IceCube uses two trigger criteria and may add a third. The main trigger is based on multiplicity; it selects time intervals where eight DOMs (with local coincidences) fired within 5 μs. This collects most of the neutrino events. In 2008, a string trigger was added; it selects time intervals when five out of seven adjacent DOMs fired within 1.5 μs. This trigger has improved sensitivity for low energy events, especially upward going

muons. A third, 'topological' trigger is also under consideration; it will be optimized for low-energy horizontal muons. When a trigger occurs, all data within ±10 μs trigger window is saved, becoming an event. If multiple trigger windows overlap, then all of the data from the ORed time intervals are saved as a single event.

Triggered data is reconstructed by an on-line filter system and selected events are transmitted via satellite to the Northern hemisphere. The filters use simple criteria, 'first-guess' reconstruction algorithms and simplified maximum likelihood fitting. Current filters select upward going muons, cascades (ν_e, ν_τ and all-flavor neutral current interactions), extremely high energy events, starting and stopping events, and air showers seen in IceTop. Currently, about 6% of the events are selected by these filters, comprising about 32.5 Gbytes/day. The remainder of the data is stored on tapes at the South Pole station. The tapes are carried north during the austral summer.

VIII.   EVENT RECONSTRUCTION

In the Northern hemisphere, events are reconstructed using maximum-likelihood fitting techniques. Events are fitted to templates representing different decay modes. Figure 5 shows examples of three different interaction topologies in IceCube [2].

Figure 5 (top) shows an actual (data) muon or muon bundle (group of parallel muons from an air shower). The tracks are visible over 1 km. This long lever arm allows for good directional reconstruction, better than 1 degree. Of course, for shorter tracks, the resolution degrades. It is also possible to estimate muon energies by either the length of their tracks, or by measuring the specific energy loss; at energies above 1 TeV, muon energy loss (dE/dx) is proportional to the muon energy.

Figure 5 (middle) shows a simulated ν_e interaction which produces a compact deposition of energy; this is known as a 'cascade.' Cascades are also produced by neutral current neutrino interactions and low-energy (below 1 PeV) ν_τ interactions. Although there is very little directional information, cascade energies may be determined to within a factor of 2.

Figure 5 (bottom) shows a simulated few-PeV ν_τ interaction forming a classic 'double-bang' topology. The interaction produces one cascade when the ν_τ interacts. That interaction produces a τ, which, at PeV energies, can travel hundreds of meters before decaying. The second cascade comes when the τ decays. Several other τ decay modes are under study in IceCube.

Other topologies are also being studied. For example, a ν_μ interacting in the detector will produce a hadronic shower from the struck nucleus, in addition to the μ track. Muons can also stop in the detector.

Of course, the most common events are downward going muons produced in cosmic-ray air showers. These events outnumber neutrino induced events by about 500,000:1. Rejection of this background is a significant difficulty which must be dealt with in event reconstruction.

Events are reconstructed by fitting them to one of these hypotheses. The starting points for these fits are the results of 'first guess' methods. For muons, the main first guess method fits a moving plane to the launch times in the DOMs [10]. For a muon, the plane should have a velocity near the speed of light. An alternate approach uses the measured charge deposition to the 'long axis' in events such as in Fig. 5 (top).
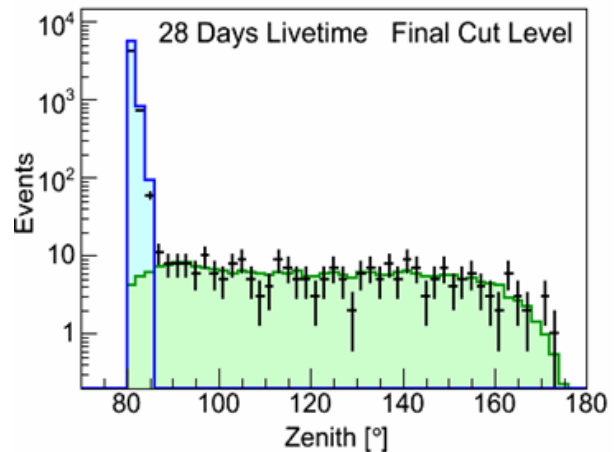

Fig. 6. The azimuthal angle for downward-going, or near downward-going muons in IceCube 40, after tight cuts, compared with the results of cosmic ray muons (blue) and neutrinos (green) simulations. The coincident muon background is largely eliminated (4 downward going events expected) and not shown here.

The maximum likelihood fitter finds the likelihood for different track positions and directions, and, optionally, energy. To do this, it uses functions which model the light propagation, giving the probability distribution for a photon radiated from a track with a given orientation to reach a DOM at a given perpendicular distance and orientation as a function of time. These functions are precalculated using a simulation that tracks photons through the ice, and stored in a 7-dimensional histogram [11]. One of the dimensions is depth, incorporating the depth dependence of the optical properties of the ice.

Because of the high rate of downward going muons, it is not enough to select events with the most likely reconstruction as upward going [12]. Fairly stringent cuts must be applied to eliminate tracks with reasonable likelihoods for being downward going. This can be done by cutting on the estimated errors from the likelihood fit, which can act as a stand-in for the depth of the minimum in the likelihood function. Alternately, one can perform a Bayesian reconstruction, weighting fits to different zenith angles by the relative size of the signal in that direction (effectively requiring that the upward going hypothesis be much more likely).

IceCube is big enough that there is also a significant background due to coincident muon events, whereby two (or more) muons from independent cosmic-ray air showers traverse the detector in the course of one event. Specific

algorithms have been developed to find and reject these events.

After these cuts, a relatively clean sample of well-reconstructed neutrino events remains, as is shown in Fig. 6. There remains an irreducible background of atmospheric neutrinos produced by cosmic-ray air showers in the northern hemisphere. In 1 year (about 319 live days) of IC40 data, we expect about 5,000 atmospheric ν_μ interactions. The atmospheric ν_e background is about two orders of magnitude lower and the atmospheric ν_τ background is almost absent.

The lower backgrounds make the two latter channels attractive avenues to search for extraterrestrial neutrinos. In searches for point sources of neutrinos, off-source regions are used to directly measure the background level [13]. Diffuse neutrino analyses use the fact that the energy spectrum of the atmospheric neutrinos is much softer than for extra-terrestrial neutrinos; by selecting high energy events, one can largely remove the atmospheric background [14].

IX. FUTURE PLANS

IceCube completion is scheduled for 2011. In addition to the 80 baseline strings, we are also developing a Deep Core infill array. Deep Core will consist of 6 additional strings with a smaller, 72 m grid spacing. The DOMs will use new high-quantum-efficiency phototubes which will be spaced every 7 m in the deepest, clearest 350 m of ice. In addition, the rest of IceCube will serve as a veto region surrounding Deep Core, allowing for the rejection of cosmic-ray muons and other non-contained backgrounds. The higher granularity, improved optical sensitivity and surrounding veto will give Deep Core a much lower threshold than IceCube, in the 10 GeV region.

IceCube collaborators are also studying prototype radio and acoustic neutrino detectors. These are sensitive to coherent radio-Cherenkov emission from neutrino-induced electromagnetic and hadronic showers and the shock-wave produced by local heating from neutrino induced showers respectively. The radio and acoustic signals should have much larger absorption lengths than light, so these techniques might be usable to build a much larger (100 km^3) array than IceCube. However, because of the technique used, the array would have a much higher energy threshold (10^{17} eV?) than IceCube.

X. CONCLUSIONS AND FUTURE PLANS

The 1 km^3 IceCube neutrino observatory detects Cherenkov radiation from charged particles produced in neutrino interactions. With its 4800 digital optical modules, IceCube acts like a tracking calorimeter, recording the pattern of energy deposition in the ice. Each DOM includes a complete data acquisition system. IceCube construction is 50%

complete and the system is working well with very high reliability.

The segmentation gives IceCube the capacity to separate the different topologies from ν_μ, ν_e and ν_τ interactions. We have developed reconstruction methods that effectively separate upward going muons from ν_μ interactions from the much more intense cosmic ray muon background. These methods achieve an angular resolution of better than 1 degree for long tracks.

REFERENCES

[1]   IceCube Collaboration, "Preliminary Design Document," unpublished. http://www.icecube.wisc.edu/science/publications/pdd/

[2]   F. Halzen and S. R. Klein, "Astronomy and Astrophysics with Neutrinos," *Physics Today*, vol. 61, no 5. pp. 29-35. May 2008.

[3]   F. Halzen and D. Hooper, "High-energy neutrino astronomy: The Cosmic ray connection," Rept. Prog. Phys. vol. 65, pp. 1025-1078, 2002.

[4]   J. Abraham *et al.* (the Pierre Auger Collaboration), "Correlation of the Highest-Energy Cosmic Rays with Nearby Extragalactic Objects," *Science*, vol. 318, pp. 938-943, 2007.

[5]   R. U. Abbasi *et al.* (The Flys Eye Collaboration), "Correlations between HiRes Stereo Events and Active Galactic Nuclei," preprint arXiv:0804.0382, April, 2008.

[6]   E. Andres *et al.* (the AMANDA Collaboration), "Observation of high-energy neutrinos using Cherenkov detectors embedded deep in Antarctic ice," *Nature*, vol. 410, pp. 441-443 (2001).

[7]   A. Achterberg *et al.* (the IceCube Collaboration), "First year performance of the IceCube neutrino telescope," Astropart. Phys., vol. 26, pp. 155-173, 2006.

[8]   M. Ackermann *et al.*, "Optical properties of deep glacial ice at the South Pole, *J. Geophys. Res*. vol. 111, pp. D13203, 2006.

[9]   N. E. Bramall, R. C. Bay, K. Woschnagg, R. A. Rohde and P. B. Price, "A deep high-resolution optical log of dust, ash and statigraphy in South Pole glacial ice," *Geophys. Res. Lett*., vol. 32, pp. L21815 .

[10]   J. Ahrens *et al*. (the AMANDA Collaboration), "Muon track reconstruction and data selection techniques in AMANDA," *Nucl. Instrum. Meth*. vol. A524, pp. 169-194, 2004.

[11]   J. Lundberg *et al*., "Light tracking through ice and water – Scattering and absorption in heterogenous media with PHOTONICS," *Nucl. Instrum. and Meth*. vol. A581, pp. 619-631, 2007.

[12]   A. Achterberg *et al.* (the IceCube Collaboration), "The IceCube Collaboration: Contributions to the 30th International Cosmic Ray Conference (ICRC 2007)," unpublished preprint arXiv:0711.0353, November 2007.

[13]   A. Achterberg *et al*., (the IceCube Collaboration), "Multi-year search for a diffuse flux of muon neutrinos with AMANDA-II," *Phys. Rev.* vol. D76, pp. 042008, 2007.

[14]   A. Achterberg *et al*. (the IceCube Collaboration), "Five years of searches for point sources of astrophysical neutrinos with the AMANDA-II neutrino telescope," *Phys. Rev.* vol. D75, pp. 102001, 2007.