

The IceCube Neutrino Observatory: instrumentation and online systems

To cite this article: M.G. Aartsen *et al* 2017 *JINST* **12** P03012

View the [article online](#) for updates and enhancements.

Related content

- [PINGU: A vision for neutrino and particle physics at the South Pole](#)
M G Aartsen, K Abraham, M Ackermann *et al*.
- [IceCube PeV neutrinos and leptophilic dark matter](#)
Marco Chianese
- [2006 December 13 Solar Particle Spectrum](#)
R. Abbasi, M. Ackermann, J. Adams *et al*.

Recent citations

- [Measurements using the inelasticity distribution of multi-TeV neutrino interactions in IceCube](#)
M. G. Aartsen *et al*
- [Constraints on Minute-Scale Transient Astrophysical Neutrino Sources](#)
M. G. Aartsen *et al*
- [STRAW \(STRings for Absorption length in Water\): pathfinder for a neutrino telescope in the deep Pacific Ocean](#)
M. Boehmer *et al*



IOP | ebooks™

Bringing you innovative digital publishing with leading voices to create your essential collection of books in STEM research.

Start exploring the collection - download the first chapter of every title for free.

The IceCube Neutrino Observatory: instrumentation and online systems



The IceCube collaboration

M.G. Aartsen,² M. Ackermann,⁵³ J. Adams,¹⁶ J.A. Aguilar,¹² M. Ahlers,³⁰ M. Ahrens,⁴³ D. Altmann,²⁴ K. Andeen,³³ T. Anderson,⁴⁹ I. Anseau,¹² G. Anton,²⁴ M. Archinger,³² C. Argüelles,¹⁴ R. Auer,³⁰ J. Auffenberg,¹ S. Axani,¹⁴ J. Baccus,³⁰ X. Bai,⁴¹ S. Barnett,³⁰ S.W. Barwick,²⁷ V. Baum,³² R. Bay,⁷ K. Beattie,⁸ J.J. Beatty,^{18,19} J. Becker Tjus,¹⁰ K.-H. Becker,⁵² T. Bendfelt,³⁰ S. BenZvi,⁵⁰ D. Berley,¹⁷ E. Bernardini,⁵³ A. Bernhard,³⁵ D.Z. Besson,²⁸ G. Binder,^{8,7} D. Bindig,⁵² M. Bissok,¹ E. Blaufuss,¹⁷ S. Blot,⁵³ D. Boersma,⁵¹ C. Boehm,⁴³ M. Börner,²¹ F. Bos,¹⁰ D. Bose,⁴⁵ S. Böser,³² O. Botner,⁵¹ A. Bouchta,⁵¹ J. Braun,³⁰ L. Brayeur,¹³ H.-P. Bretz,⁵³ S. Bron,²⁵ A. Burgman,⁵¹ C. Burreson,³⁰ T. Carver,²⁵ M. Casier,¹³ E. Cheung,¹⁷ D. Chirkin,³⁰ A. Christov,²⁵ K. Clark,⁴⁶ L. Classen,³⁶ S. Coenders,³⁵ G.H. Collin,¹⁴ J.M. Conrad,¹⁴ D.F. Cowen,^{49,48} R. Cross,⁵⁰ C. Day,⁸ M. Day,³⁰ J.P.A.M. de André,²² C. De Clercq,¹³ E. del Pino Rosendo,³² H. Dembinski,³⁷ S. De Ridder,²⁶ F. Descamps,²⁶ P. Desiati,³⁰ K.D. de Vries,¹³ G. de Wasseige,¹³ M. de With,⁹ T. DeYoung,²² J.C. Díaz-Vélez,³⁰ V. di Lorenzo,³² H. Dujmovic,⁴⁵ J.P. Dumm,⁴³ M. Dunkman,⁴⁹ B. Eberhardt,³² W.R. Edwards,⁸ T. Ehrhardt,³² B. Eichmann,¹⁰ P. Eller,⁴⁹ S. Euler,⁵¹ P.A. Evenson,³⁷ S. Fahey,³⁰ A.R. Fazely,⁶ J. Feintzeig,³⁰ J. Felde,¹⁷ K. Filimonov,⁷ C. Finley,⁴³ S. Flis,⁴³ C.-C. Fösig,³² A. Franckowiak,⁵³ M. Frère,³⁰ E. Friedman,¹⁷ T. Fuchs,²¹ T.K. Gaisser,³⁷ J. Gallagher,²⁹ L. Gerhardt,^{8,7} K. Ghorbani,³⁰ W. Giang,²³ L. Gladstone,³⁰ T. Glauch,¹ D. Glowacki,³⁰ T. Glüsenskamp,²⁴ A. Goldschmidt,⁸ J.G. Gonzalez,³⁷ D. Grant,²³ Z. Griffith,³⁰ L. Gustafsson,⁵¹ C. Haack,¹ A. Hallgren,⁵¹ F. Halzen,³⁰ E. Hansen,²⁰ T. Hansmann,¹ K. Hanson,³⁰ J. Haugen,³⁰ D. Hebecker,⁹ D. Heereman,¹² K. Helbing,⁵² R. Hellauer,¹⁷ R. Heller,⁵³ S. Hickford,⁵² J. Hignight,²² G.C. Hill,² K.D. Hoffman,¹⁷ R. Hoffmann,⁵² K. Hoshina,^{30,54} F. Huang,⁴⁹ M. Huber,³⁵ P.O. Hulth,^{43,†} K. Hultqvist,⁴³ S. In,⁴⁵ M. Inaba,¹⁵ A. Ishihara,¹⁵ E. Jacobi,⁵³ J. Jacobsen,³⁰ G.S. Japaridze,⁴ M. Jeong,⁴⁵ K. Jero,³⁰ A. Jones,⁸ B.J.P. Jones,¹⁴ J. Joseph,⁸ W. Kang,⁴⁵ A. Kappes,³⁶ T. Karg,⁵³ A. Karle,³⁰

[†]Deceased.

U. Katz,²⁴ M. Kauer,³⁰ A. Keivani,⁴⁹ J.L. Kelley,^{30,*} J. Kemp,¹ A. Kheirandish,³⁰ J. Kim,⁴⁵ M. Kim,⁴⁵ T. Kintscher,⁵³ J. Kiryluk,⁴⁴ N. Kitamura,³⁰ T. Kittler,²⁴ S.R. Klein,^{8,7} S. Kleinfelder,⁸ M. Kleist,³⁰ G. Kohlen,³⁴ R. Koirala,³⁷ H. Kolanoski,⁹ R. Konietz,¹ L. Köpke,³² C. Kopper,²³ S. Kopper,⁵² D.J. Koskinen,²⁰ M. Kowalski,^{9,53} M. Krasberg,³⁰ K. Krings,³⁵ M. Kroll,¹⁰ G. Krückl,³² C. Krüger,³⁰ J. Kunnen,¹³ S. Kunwar,⁵³ N. Kurahashi,⁴⁰ T. Kuwabara,¹⁵ M. Labare,²⁶ K. Laihem,¹ H. Landsman,³⁰ J.L. Lanfranchi,⁴⁹ M.J. Larson,²⁰ F. Lauber,⁵² A. Laudrie,³¹ D. Lennarz,²² H. Leich,⁵³ M. Lesiak-Bzdak,⁴⁴ M. Leuermann,¹ L. Lu,¹⁵ J. Ludwig,⁸ J. Lünemann,¹³ C. Mackenzie,³⁰ J. Madsen,⁴² G. Maggi,¹³ K.B.M. Mahn,²² S. Mancina,³⁰ M. Mandelartz,¹⁰ R. Maruyama,³⁸ K. Mase,¹⁵ H. Matis,⁸ R. Maunu,¹⁷ F. McNally,³⁰ C.P. McParland,⁸ P. Meade,³⁰ K. Meagher,¹² M. Medici,²⁰ M. Meier,²¹ A. Meli,²⁶ T. Menne,²¹ G. Merino,³⁰ T. Meures,¹² S. Miarecki,^{8,7} R.H. Minor,⁸ T. Montaruli,²⁵ M. Moulai,¹⁴ T. Murray,³⁰ R. Nahnauer,⁵³ U. Naumann,⁵² G. Neer,²² M. Newcomb,³⁰ H. Niederhausen,⁴⁴ S.C. Nowicki,²³ D.R. Nygren,⁸ A. Obertacke Pollmann,⁵² A. Olivas,¹⁷ A. O'Murchadha,¹² T. Palczewski,^{8,7} H. Pandya,³⁷ D.V. Pankova,⁴⁹ S. Patton,⁸ P. Peiffer,³² Ö. Penek,¹ J.A. Pepper,⁴⁷ C. Pérez de los Heros,⁵¹ C. Pettersen,³⁰ D. Pieloth,²¹ E. Pinat,¹² P.B. Price,⁷ G.T. Przybylski,⁸ M. Quinnan,⁴⁹ C. Raab,¹² L. Rädcl,¹ M. Rameez,²⁰ K. Rawlins,³ R. Reimann,¹ B. Relethford,⁴⁰ M. Relich,¹⁵ E. Resconi,³⁵ W. Rhode,²¹ M. Richman,⁴⁰ B. Riedel,²³ S. Robertson,² M. Rongen,¹ C. Roucelle,⁸ C. Rott,⁴⁵ T. Ruhe,²¹ D. Ryckbosch,²⁶ D. Rysewyk,²² L. Sabbatini,³⁰ S.E. Sanchez Herrera,²³ A. Sandrock,²¹ J. Sandroos,³² P. Sandstrom,³⁰ S. Sarkar,^{20,39} K. Satalecka,⁵³ P. Schlunder,²¹ T. Schmidt,¹⁷ S. Schoenen,¹ S. Schöneberg,¹⁰ A. Schukraft,¹ L. Schumacher,¹ D. Seckel,³⁷ S. Seunarine,⁴² M. Solarz,⁷ D. Soldin,⁵² M. Song,¹⁷ G.M. Spiczak,⁴² C. Spiering,⁵³ T. Stanev,³⁷ A. Stasik,⁵³ J. Stettner,¹ A. Steuer,³² T. Stezelberger,⁸ R.G. Stokstad,⁸ A. Stöbl,¹⁵ R. Ström,⁵¹ N.L. Strotjohann,⁵³ K.-H. Sulanke,⁵³ G.W. Sullivan,¹⁷ M. Sutherland,¹⁸ H. Taavola,⁵¹ I. Taboada,⁵ J. Tatar,^{8,7} F. Tenholt,¹⁰ S. Ter-Antonyan,⁶ A. Terliuk,⁵³ G. Tešić,⁴⁹ L. Thollander,⁴³ S. Tilav,³⁷ P.A. Toale,⁴⁷ M.N. Tobin,³⁰ S. Toscano,¹³ D. Tosi,³⁰ M. Tselengidou,²⁴ A. Turcati,³⁵ E. Unger,⁵¹ M. Usner,⁵³ J. Vandenbroucke,³⁰ N. van Eijndhoven,¹³ S. Vanheule,²⁶ M. van Rossem,³⁰ J. van Santen,⁵³ M. Vehring,¹ M. Voge,¹¹ E. Vogel,¹ M. Vraeghe,²⁶ D. Wahl,³¹ C. Walck,⁴³ A. Wallace,² M. Wallraff,¹ N. Wandkowsky,³⁰ Ch. Weaver,²³ M.J. Weiss,⁴⁹ C. Wendt,³⁰ S. Westerhoff,³⁰ D. Wharton,³⁰ B.J. Whelan,² S. Wickmann,¹ K. Wiebe,³² C.H. Wiebusch,¹ L. Wille,³⁰ D.R. Williams,^{47,*} L. Wills,⁴⁰ P. Wisniewski,³⁰ M. Wolf,⁴³ T.R. Wood,²³ E. Woolsey,²³ K. Woschnagg,⁷ D.L. Xu,³⁰ X.W. Xu,⁶ Y. Xu,⁴⁴ J.P. Yanez,²³ G. Yodh,²⁷ S. Yoshida¹⁵ and M. Zoll⁴³

¹*III. Physikalisches Institut, RWTH Aachen University, D-52056 Aachen, Germany*

²*Department of Physics, University of Adelaide, Adelaide, 5005, Australia*

³*Dept. of Physics and Astronomy, University of Alaska Anchorage, 3211 Providence Dr., Anchorage, AK 99508, U.S.A.*

⁴*CTSPS, Clark-Atlanta University, Atlanta, GA 30314, U.S.A.*

⁵*School of Physics and Center for Relativistic Astrophysics, Georgia Institute of Technology, Atlanta, GA 30332, U.S.A.*

⁶*Dept. of Physics, Southern University, Baton Rouge, LA 70813, U.S.A.*

⁷*Dept. of Physics, University of California, Berkeley, CA 94720, U.S.A.*

*Corresponding author.

- ⁸*Lawrence Berkeley National Laboratory, Berkeley, CA 94720, U.S.A.*
- ⁹*Institut für Physik, Humboldt-Universität zu Berlin, D-12489 Berlin, Germany*
- ¹⁰*Fakultät für Physik & Astronomie, Ruhr-Universität Bochum, D-44780 Bochum, Germany*
- ¹¹*Physikalisches Institut, Universität Bonn, Nussallee 12, D-53115 Bonn, Germany*
- ¹²*Université Libre de Bruxelles, Science Faculty CP230, B-1050 Brussels, Belgium*
- ¹³*Vrije Universiteit Brussel (VUB), Dienst ELEM, B-1050 Brussels, Belgium*
- ¹⁴*Dept. of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.*
- ¹⁵*Dept. of Physics and Institute for Global Prominent Research, Chiba University, Chiba 263-8522, Japan*
- ¹⁶*Dept. of Physics and Astronomy, University of Canterbury, Private Bag 4800, Christchurch, New Zealand*
- ¹⁷*Dept. of Physics, University of Maryland, College Park, MD 20742, U.S.A.*
- ¹⁸*Dept. of Physics and Center for Cosmology and Astro-Particle Physics, Ohio State University, Columbus, OH 43210, U.S.A.*
- ¹⁹*Dept. of Astronomy, Ohio State University, Columbus, OH 43210, U.S.A.*
- ²⁰*Niels Bohr Institute, University of Copenhagen, DK-2100 Copenhagen, Denmark*
- ²¹*Dept. of Physics, TU Dortmund University, D-44221 Dortmund, Germany*
- ²²*Dept. of Physics and Astronomy, Michigan State University, East Lansing, MI 48824, U.S.A.*
- ²³*Dept. of Physics, University of Alberta, Edmonton, Alberta, Canada T6G 2E1*
- ²⁴*Erlangen Centre for Astroparticle Physics, Friedrich-Alexander-Universität Erlangen-Nürnberg, D-91058 Erlangen, Germany*
- ²⁵*Département de physique nucléaire et corpusculaire, Université de Genève, CH-1211 Genève, Switzerland*
- ²⁶*Dept. of Physics and Astronomy, University of Gent, B-9000 Gent, Belgium*
- ²⁷*Dept. of Physics and Astronomy, University of California, Irvine, CA 92697, U.S.A.*
- ²⁸*Dept. of Physics and Astronomy, University of Kansas, Lawrence, KS 66045, U.S.A.*
- ²⁹*Dept. of Astronomy, University of Wisconsin, Madison, WI 53706, U.S.A.*
- ³⁰*Dept. of Physics and Wisconsin IceCube Particle Astrophysics Center, University of Wisconsin, Madison, WI 53706, U.S.A.*
- ³¹*Physical Sciences Laboratory, University of Wisconsin, Stoughton, WI 53589, U.S.A.*
- ³²*Institute of Physics, University of Mainz, Staudinger Weg 7, D-55099 Mainz, Germany*
- ³³*Department of Physics, Marquette University, Milwaukee, WI, 53201, U.S.A.*
- ³⁴*Université de Mons, 7000 Mons, Belgium*
- ³⁵*Physik-department, Technische Universität München, D-85748 Garching, Germany*
- ³⁶*Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, D-48149 Münster, Germany*
- ³⁷*Bartol Research Institute and Dept. of Physics and Astronomy, University of Delaware, Newark, DE 19716, U.S.A.*
- ³⁸*Dept. of Physics, Yale University, New Haven, CT 06520, U.S.A.*
- ³⁹*Dept. of Physics, University of Oxford, 1 Keble Road, Oxford OX1 3NP, U.K.*
- ⁴⁰*Dept. of Physics, Drexel University, 3141 Chestnut Street, Philadelphia, PA 19104, U.S.A.*
- ⁴¹*Physics Department, South Dakota School of Mines and Technology, Rapid City, SD 57701, U.S.A.*
- ⁴²*Dept. of Physics, University of Wisconsin, River Falls, WI 54022, U.S.A.*
- ⁴³*Oskar Klein Centre and Dept. of Physics, Stockholm University, SE-10691 Stockholm, Sweden*
- ⁴⁴*Dept. of Physics and Astronomy, Stony Brook University, Stony Brook, NY 11794-3800, U.S.A.*
- ⁴⁵*Dept. of Physics, Sungkyunkwan University, Suwon 440-746, Korea*
- ⁴⁶*Dept. of Physics, University of Toronto, Toronto, Ontario, Canada, M5S 1A7*
- ⁴⁷*Dept. of Physics and Astronomy, University of Alabama, Tuscaloosa, AL 35487, U.S.A.*

⁴⁸*Dept. of Astronomy and Astrophysics, Pennsylvania State University, University Park, PA 16802, U.S.A.*

⁴⁹*Dept. of Physics, Pennsylvania State University, University Park, PA 16802, U.S.A.*

⁵⁰*Dept. of Physics and Astronomy, University of Rochester, Rochester, NY 14627, U.S.A.*

⁵¹*Dept. of Physics and Astronomy, Uppsala University, Box 516, S-75120 Uppsala, Sweden*

⁵²*Dept. of Physics, University of Wuppertal, D-42119 Wuppertal, Germany*

⁵³*DESY, D-15735 Zeuthen, Germany*

⁵⁴*Earthquake Research Institute, University of Tokyo, Bunkyo, Tokyo 113-0032, Japan*

E-mail: analysis@icecube.wisc.edu

ABSTRACT: The IceCube Neutrino Observatory is a cubic-kilometer-scale high-energy neutrino detector built into the ice at the South Pole. Construction of IceCube, the largest neutrino detector built to date, was completed in 2011 and enabled the discovery of high-energy astrophysical neutrinos. We describe here the design, production, and calibration of the IceCube digital optical module (DOM), the cable systems, computing hardware, and our methodology for drilling and deployment. We also describe the online triggering and data filtering systems that select candidate neutrino and cosmic ray events for analysis. Due to a rigorous pre-deployment protocol, 98.4% of the DOMs in the deep ice are operating and collecting data. IceCube routinely achieves a detector uptime of 99% by emphasizing software stability and monitoring. Detector operations have been stable since construction was completed, and the detector is expected to operate at least until the end of the next decade.

KEYWORDS: Large detector systems for particle and astroparticle physics; Neutrino detectors; Online farms and online filtering; Trigger concepts and systems (hardware and software)

ARXIV EPRINT: [1612.05093](https://arxiv.org/abs/1612.05093)

Contents

1	Introduction	1
1.1	A functional description of the IceCube observatory	2
1.1.1	IceCube in-ice array	3
1.1.2	DeepCore	4
1.1.3	IceTop	5
1.1.4	The IceCube Laboratory	5
2	The Digital Optical Module	5
2.1	A functional description of the DOM	5
2.2	Components	7
2.2.1	Glass sphere and harness	7
2.2.2	Cable penetrator, cable and connector	8
2.2.3	PMT, gel and magnetic shield	9
2.2.4	High voltage supply and divider	10
2.2.5	Main board and delay board	10
2.2.6	Flasher board	12
2.3	Production and testing	15
2.4	Reliability	17
3	Calibration and stability	18
3.1	DOMCal	19
3.2	Waveform calibration	21
3.3	RAPCal	23
3.4	DOM optical efficiency	26
3.5	Baseline voltage stability	28
3.6	Gain stability	30
3.7	Optical efficiency stability	32
3.8	Dark noise	33
4	Cable systems	37
4.1	Design	37
4.2	Installation	40
5	Drilling and deployment	40
5.1	Hot water drilling	40
5.2	Deployment procedures and instrumentation	43
5.3	Geometry measurement	44
5.3.1	Stage 1 geometry measurement	45
5.3.2	Stage 2 geometry measurement	45
5.3.3	Trilateration validation of DOM coordinates	45

5.3.4	Inclinometers	46
5.4	Hole ice imaging camera	47
6	Online systems	48
6.1	Data flow overview	48
6.2	South Pole System and South Pole Test System	50
6.3	Data readout and timing	51
6.3.1	Communications	51
6.3.2	Master clock system	52
6.3.3	DOR card and driver	52
6.3.4	DOM hit payloads	53
6.4	Data acquisition software	53
6.4.1	StringHub and HitSpool	54
6.4.2	Triggers	54
6.4.3	Event Builder	57
6.4.4	DAQ and DOM configuration	58
6.4.5	Component control	58
6.4.6	Supernova data acquisition system	58
6.4.7	HitSpool request system	59
6.5	Online processing and filtering	60
6.5.1	Overview	60
6.5.2	System design	61
6.5.3	Components	61
6.5.4	Performance	62
6.6	Data handling	63
6.7	IceCube live and remote monitoring	65
6.7.1	LiveControl	65
6.7.2	IceCube live website	65
6.8	Operational performance	66
7	Outlook	68
7.1	IceCube Gen2 and PINGU	69

1 Introduction

In the six decades following the discovery of the neutrino by Cowan and Reines [1], detectors have been realized that have explored the properties and sources of neutrinos. Early developments in the field included radiochemical detection of MeV-scale neutrinos from nuclear fusion in the Sun [2] and the first detections of higher energy atmospheric neutrinos in deep underground telescopes [3, 4]. Advanced atmospheric and solar neutrino observatories provided definitive evidence of neutrino

mass and constrained neutrino mixing parameters [5, 6]. The first neutrinos to be detected from outside the solar system were from Supernova 1987A [7–9]. The search for neutrinos emanating from astrophysical processes led to the development of a new generation of large-scale neutrino detectors.

Atmospheric neutrinos have energies in the GeV–TeV range and result from the interactions of cosmic ray particles with the atmosphere. Cosmic rays in turn are believed to be accelerated in astrophysical objects such as supernova remnants, active galactic nuclei, and gamma ray bursts. These astrophysical acceleration sites should also produce neutrinos up to PeV-scale energies. Neutrinos, because of their weak interaction and electrically neutral character, are useful probes of high-energy phenomena in the Universe. Unlike photons, detecting their origins in astrophysical acceleration sites would unambiguously indicate hadronic acceleration and provide identification of the origins of cosmic rays. Neutrinos arrive undeflected and unscattered upon detection and thus point back to their sources, providing a clear view of the physics deep within shrouded and compact sources. At the most extreme energies, they are the only particles that can reach us from sources at cosmological distances. Notwithstanding the neutrinos seen from SN1987A, a remarkable but extremely rare phenomenon, the detection of neutrinos originating in astrophysical processes outside our solar system requires detector facilities of extreme dimensions to detect the faint fluxes of these weakly-interacting particles.

The DUMAND experiment [10] pioneered the idea of using large-area photomultiplier tubes in the deep ocean to detect high energy neutrinos. Although DUMAND was never realized, experience in operating detectors remotely in a harsh environment led to the successful efforts of Baikal [11], AMANDA [12] and ANTARES [13]. These experiments provided key measurements of the high-energy atmospheric neutrino spectrum, constrained optimistic models of astrophysical neutrino production, and demonstrated the feasibility of the technique. However, the detection of astrophysical neutrinos proved elusive, suggesting that a kilometer-scale array was required to achieve the necessary sensitivity [14–16].

The IceCube Neutrino Observatory is a cubic-kilometer neutrino detector built into the ice at the South Pole. Construction was completed on December 18, 2010, and commissioning was completed in 2011. Its primary scientific objective has been the discovery of astrophysical neutrinos, which was achieved in 2013 [17], and the identification and characterization of their sources. Other science objectives include indirect detection of dark matter, searches for other exotic particles, studies of neutrino oscillation physics, and detection of the neutrino burst from a Galactic core-collapse supernova [18]. A multi-messenger collaboration with optical, X-ray, gamma-ray, radio, and gravitational wave observatories provides multiple windows onto the potential neutrino sources. A key to the success of these initiatives is the reliability and performance of the IceCube instrumentation, as well as the flexibility and stability of the online software systems.

1.1 A functional description of the IceCube observatory

In order to detect neutrinos, IceCube exploits the fact that charged particles resulting from neutrino interactions move through the ice faster than the phase velocity of light in ice, and therefore emit Cherenkov photons. An enormous detection volume is required due to the small interaction cross-section of neutrinos and the extremely low fluxes expected at Earth from astrophysical objects. The ice cap at the South Pole is about three kilometers thick and is a suitable operational site, since it not only offers a large quantity of interaction material but also a medium with excellent optical

qualities. With a Cherenkov photon yield of $O(10^5)$ visible photons per GeV of secondary particle shower energy, the long optical attenuation length of South Pole ice, and large-area PMTs, it is possible to instrument cubic kilometers of ice with a rather sparse spacing of detectors. IceCube is located at Amundsen-Scott South Pole Station, which offers the logistical support required for the construction and operation of the observatory.

The basic detection unit in IceCube is the digital optical module (DOM), covered in detail in section 2. Encapsulated in a glass pressure sphere to withstand the extreme pressure in the deep ice, the main components of a DOM are a 10" PMT, embedded high-voltage generation, a light-emitting diode (LED) Flasher Board, and a Main Board containing the analog and digital processing circuitry for PMT pulses. Detector calibration is described in section 3. Digitized, timestamped PMT signals are sent from the DOMs to a central computing facility at the surface via a cable system described in section 4. Aspects of detector deployment and ice drilling are covered in section 5. An overview of the data flow including DOM readout, event triggering, processing and filtering are presented in section 6, along with data handling, monitoring, and operational performance of the observatory.

The IceCube Neutrino Observatory consists of a subsurface “in-ice” array of DOMs, including the more densely instrumented DeepCore sub-array, and the IceTop surface array. The entire detector uses the same DOM design and associated surface readout. A schematic layout of the observatory is shown in figure 1.

1.1.1 IceCube in-ice array

In order to detect the Cherenkov photons emitted by charged particles traversing the ice, 5160 DOMs are deployed between 1450 m and 2450 m below the surface of the ice on 86 vertical strings. Each string consists of 60 DOMs deployed along a single cable containing twisted copper-wire pairs. The primary in-ice array consists of 78 strings with a vertical separation of the DOMs of 17 m. The strings are deployed within a hexagonal footprint on a triangular grid with 125 m horizontal spacing, instrumenting a volume of one cubic kilometer of ice. This design was chosen in order to meet the primary science requirement of detecting astrophysical neutrinos in the energy range of $O(\text{TeV})$ – $O(\text{PeV})$.

Two different event topologies form the standard signatures of neutrinos in IceCube. Track-like events originate from a charged-current interaction of a high-energy muon neutrino with a nucleus, producing a hadronic shower at the vertex and an outgoing muon that emits Cherenkov light in a cone along its track. The angular resolution for muon tracks and hence the incident neutrino direction is typically 0.6° , confirmed by analysis of the shadows of the Moon and Sun in cosmic rays [19, 20]. Muons with energies above a critical energy, about 1 TeV in ice, predominantly lose energy by radiative processes that exhibit a stochastic behavior with large fluctuations. This results in large variability in the amount of energy deposited for different muons of the same energy. A second class of events are electromagnetic or hadronic showers from interactions of all neutrino flavors, resulting in a more spherical light generation in the detector. Since the total light output of such a shower is directly proportional to its energy, and the showers are often well-contained in the detector, the neutrino energy reconstruction for such events is much more precise than for track-like events. The average *deposited* energy resolution for both event types is about 15% [21].




Figure 1. The IceCube Neutrino Observatory with the in-ice array, its sub-array DeepCore, and the cosmic-ray air shower array IceTop. The different string/station colors represent different deployment seasons.

1.1.2 DeepCore

A subset of in-ice DOMs is deployed deeper than 1750 m with a denser instrumented volume and correspondingly lower energy threshold. This sub-array, DeepCore [22], consists of eight specialized and closely-spaced strings of sensors in the center of the array, along with the seven central standard IceCube strings. The inter-string spacing in DeepCore varies from 41 m to 105 m, with an average spacing of 72 m.

The eight specialized DeepCore strings have a DOM-to-DOM spacing of 7 m for the bottom 50 DOMs, deployed at depths of 2100 m to 2450 m. The remaining 10 DOMs are deployed at depths shallower than 2000 m with a spacing of 10 m to form a veto cap, allowing better rejection of downgoing atmospheric muons. Depths from 2000 m to 2100 m are not instrumented, as the optical scattering and absorption is significantly increased in this region of the ice (the “dust layer” [23]).

Six of the specialized DeepCore strings are fully instrumented with DOMs using PMTs with 35% higher quantum efficiency than the standard IceCube modules. The remaining two specialized strings are equipped with a mixture of standard and higher quantum efficiency DOMs. The denser geometry and increased efficiency result in a lower energy threshold of about 10 GeV, compared to

about 100 GeV for most IceCube analyses. The DeepCore design is optimized for the detection of neutrinos with energies from 10 GeV to 100 GeV, increasing IceCube’s ability to detect atmospheric neutrino oscillations, neutrinos from WIMP dark matter annihilation, and Galactic supernovae [22].

1.1.3 IceTop

The cosmic ray air shower array IceTop [24], located on the surface of the ice at 2835 m above sea level, consists of 162 ice-filled tanks, instrumented with PMTs that detect Cherenkov radiation and arranged in 81 stations on the surface, using approximately the same grid on which the in-ice array is deployed. A denser infill array is formed by the eight stations in the center of IceTop, corresponding to the denser inter-string spacing in DeepCore. Each tank is filled with ice to a height of 0.90 m. The two tanks at each surface station are separated from each other by 10 m. Each tank contains two standard IceCube DOMs, one “high-gain” DOM operated at a PMT gain of 5×10^6 , and one “low-gain” DOM operated at a gain of 10^5 , to increase the dynamic range for air shower detection. Cosmic-ray-initiated air showers are typically spread over a number of stations. The light generated in the tanks by the shower particles (electrons, photons, muons and hadrons) is a measure of the energy deposition of these particles in the tanks. IceTop is sensitive to primary cosmic rays in the energy range of PeV to EeV with an energy resolution of 25% at 2 PeV, improving to 12% above 10 PeV [25]. For the infill array, the energy threshold is lowered to about 100 TeV. The energy range of IceCube as a cosmic-ray detector fully covers the “knee” region of the spectrum and extends to the energies where a transition from Galactic cosmic rays to a population of extra-galactic particles may occur. The IceTop array has additionally been used to study high- p_T muons, PeV gamma rays, and transient events, such as the radiation effects of solar flares. It also serves as a partial veto for the detection of downward-going neutrinos with IceCube.

1.1.4 The IceCube Laboratory

The IceCube Laboratory (ICL), located at the surface in the center of the array, is the central operations building for the experiment (figure 2). Surface cables from the array are routed up two cable towers on either side of the structure and into a server room on the second floor (section 4). The server room houses the South Pole System (section 6.2), including all data acquisition and online filtering computers. Power is supplied by the South Pole Station generators.

The building is temperature-controlled to a target of 18 °C; continuous airflow through the server room is important to avoid overheating. The second floor is shielded against electromagnetic interference with a metal mesh, and the building shield is connected to the shields of each surface cable as they enter from the cable tower bridges. Because of the low humidity, strict protocols to minimize electrostatic discharge when working on ICL equipment are necessary.

2 The Digital Optical Module

2.1 A functional description of the DOM

The DOM is the fundamental light sensor and data acquisition unit for IceCube. It consists of a spherical glass housing containing a downward-facing 10''-diameter PMT [26] and associated circuit boards that allow near-autonomous operation (figure 3). Data acquisition, control, calibration, communication, and low-voltage power conversion are integrated in one annular circuit board




Figure 2. The IceCube Laboratory (ICL) is the central operations building for the experiment. Two cable towers connect the array’s surface cables to readout computers on the second floor.

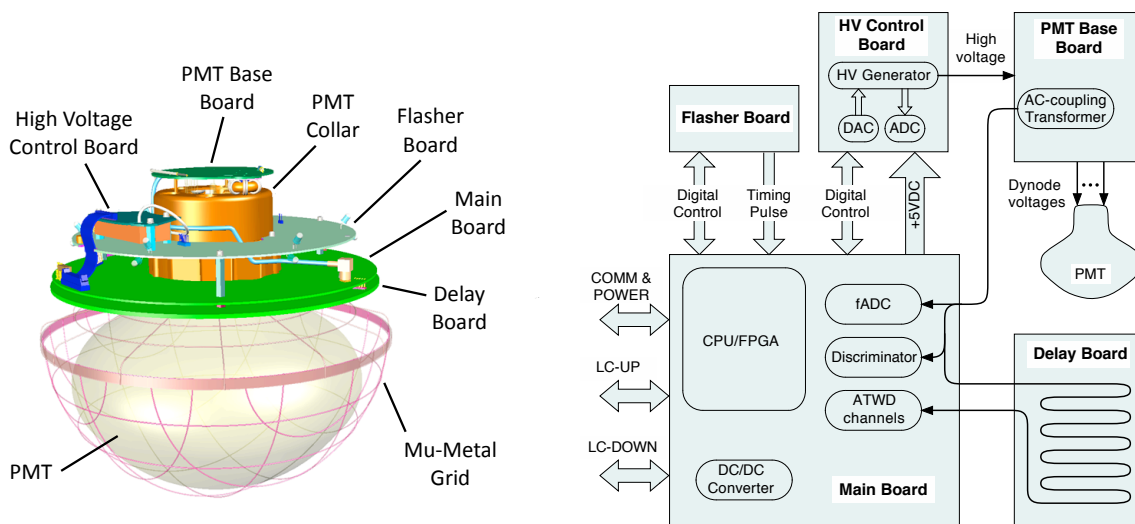


Figure 3. Components of the DOM, showing mechanical layout (left) and functional connections (right).

(Main Board) that fits around the neck of the PMT [27]. Separate circuit boards generate PMT high voltage, interface to the PMT pins, delay PMT signals, and generate calibration light flashes that can reach other DOMs. Key requirements for the DOM include the precise recording of a wide variety of PMT pulse widths and amplitudes with nanosecond time resolution, robustness in a challenging deployment environment, and long-term reliability.

The PMT detects signals from particles interacting in the ice, typically ranging over energies from 10 GeV to 10 PeV and distances up to 500 m away. At a gain of 10^7 (section 2.2.4), corresponding PMT waveforms can have amplitudes from 1 mV up to and beyond the linearity limit of the PMT (~ 2 V) and widths from 12 ns up to around 1500 ns. In order to accommodate such a variety of signals, the DOM includes multiple digitizers with overlapping dynamic range and different sampling speeds (section 2.2.5). Each DOM independently detects individual photons,

starting a recording of the PMT waveform that also includes photons arriving up to $6.4 \mu\text{s}$ later (a “hit”). The hit time is saved along with the waveform shape, allowing the determination of the times of arriving photons relative to this reference. The DOM accumulates such hit data for a period of about 1 s before sending the data up as a block. However, if data readout is interrupted, the DOM can store $\mathcal{O}(10 \text{ s})$ of data before overflowing local memory (16 MB of SDRAM), depending on hit rate. Separately, the PMT hit rate is recorded by the DOM in 1.6384 ms intervals, as a collective increase of all rates could signify detection of many low energy neutrinos in case of a Galactic supernova event (section 6.4.6) [28].

DOMs transmit their data to computers in the ICL over a twisted wire pair that also provides power (section 4). Wire pairs are bundled to form the vertical in-ice cables and the horizontal surface cables. Each wire pair is shared between two DOMs, with data transfers initiated by a surface computer. Separately, dedicated local coincidence (LC) wiring to neighbor DOMs above and below allows quick recognition of neighboring coincident hits, where nearest or next-to-nearest neighbors are hit within a common time window. The time window is configurable and is set to $\pm 1 \mu\text{s}$ for both in-ice and IceTop DOMs. The signals are forwarded from one DOM to the next through the dedicated wiring. The span of the forwarding is software-configurable and is currently set to two for in-ice DOMs, i.e. a DOM signals its neighbor and next-to-nearest neighbor DOMs in both up and down directions along the string. The local coincidence connections for IceTop, which allow coincidences between the two tanks in a station, are described in ref. [24]. Local coincidence hits (“HLC” hits) often have complex PMT waveforms indicating multiple photons detected in each DOM and are therefore saved in full detail; otherwise, the DOM saves abbreviated information appropriate to single photon detection (section 6.3.4).

The DOM is capable of interpreting commands from the surface that specify tasks for configuration, data-taking and transmission, monitoring or self-calibration. Self-calibration functions establish PMT and amplifier gains as well as sampling speed (section 3.1). The RAPCal system (section 3.3) is implemented for tracking each local DOM clock’s offset from universal time, allowing PMT pulses that were independently recorded in many DOMs to be built into events by surface computers.

2.2 Components

2.2.1 Glass sphere and harness

The glass sphere housing has an outer diameter of 13 '' and thickness of 0.5 ''. The spheres protect the inside electronics and PMT against long-term applied pressure of 250 bar (2.6 km-equivalent water depth) as well as temporary overpressure up to 690 bar during the refreezing of melted ice in the drill hole. The housings were produced by Benthos (Falmouth, Massachusetts), based on a design for deep-sea environments but using borosilicate glass from Kopp Glass with very low potassium and other radioactive trace elements that would contribute to the dark noise count rate (section 3.8). Optical transmission was measured in representative glass samples as 93% at 400 nm, decreasing to 50% at 340 nm and 10% at 315 nm (normal incidence). Fresnel reflection is not included in the quoted transmission, since the effect of Fresnel reflection is small in ice, where the refractive index is better matched to the glass.




Figure 4. (Left) DOM as deployed on main in-ice cable, showing cable breakout to the penetrator assembly and the mechanical support system. (Right) Schematic of cable connections for a set of four DOMs serviced by two wire pairs from the surface that carry power and communications. The “T” labels indicate where electrical termination ($140\ \Omega$) is installed in one of two DOMs that share such a wire pair. Other wire pairs are used for bidirectional signaling between neighboring DOMs, in order to check for in-time coincident detections.

All spheres were tested up to 690 bar hydrostatic pressure by the manufacturer. Each was delivered as two hemispheres that mate precisely at the equator and were sealed during assembly (section 2.3). The DOM is held by an aluminum waistband with rubber gaskets against the glass above and below the equator seam. Figure 4 shows how the DOM is attached to the main in-ice cable via a harness of steel rope and a chain that carries the weight load around the DOM. The main cable bends around the DOM, and the DOM axis stays vertically aligned with the string.

2.2.2 Cable penetrator, cable and connector

A penetrator assembly brings three wire pairs out through a 16.3 mm hole in the DOM glass sphere. The wires are routed inside a customized cable, shown in figure 4, and terminate at a pressure-tight, waterproof connector that mates with a similar connector that continues each pair into the main cable. One wire pair carries power and a bidirectional digital communications stream, connecting ultimately with a computer in the IceCube Laboratory building (section 4). The other wires lead to neighboring DOMs directly above and below, carrying LC digital pulses that signify time-correlated hits in nearby DOMs (section 2.2.5).

DOMs were produced in two versions, in which the communications wire pair was either electrically terminated ($140\ \Omega$) or unterminated inside the DOM. The terminated DOM is deployed 17 m below the unterminated one (7 m or 10 m in DeepCore strings) and therefore includes a correspondingly long penetrator assembly cable (figure 4).

The entire penetrator assembly was designed and produced by SEACON Brantner & Associates (El Cajon, California). The part that seals against the DOM glass is a hollow steel bolt that is secured inside the DOM by a nut and spring washer and compresses a fluorosilicone O-ring against the outside surface. The steel part includes additional sealing around the wires that pass through it. Outside the DOM, a plastic shell is molded around the steel and onto the cable jacket. External mechanical features like the penetrator are subject to large stresses during deployment and the refreezing process; a right-angle bend outside the DOM was included for robustness, based on previous experience deploying AMANDA modules.

2.2.3 PMT, gel and magnetic shield

DOMs use the 10''-diameter Hamamatsu R7081-02 PMT, or the corresponding high-quantum-efficiency (HQE) version, Hamamatsu R7081-02MOD, for DeepCore strings. The PMT properties have been measured and described in ref. [26]. The PMT is specified by Hamamatsu for the wavelength range 300 nm–650 nm, with peak quantum efficiency around 25% (34% for HQE) near 390 nm. It features a box-and-line dynode chain with 10 stages, with in-ice DOMs (both standard and HQE) operated at a gain of 10^7 (section 3.1).

The PMT bulb faces downwards in the bottom glass hemisphere, secured in high-strength silicone gel to a depth sufficient to surround the photocathode area. The gel provides both good optical coupling and mechanical support for the whole assembly of PMT and circuit boards. The gel thickness between the PMT envelope and glass sphere is approximately 1 cm. Originally the gel was supplied from General Electric as RTV6136-D1, and later as a similar formulation from Quantum Silicones (Virginia, U.S.A.). It is optically clear with transmission of 97% at 400 nm, 91% at 340 nm, and 65% at 300 nm (normal incidence). The refractive index is 1.41, yielding less than 0.1% reflection as light passes from the sphere glass into the gel and then into the PMT envelope. The characteristics of the cured gel are specified to remain stable in the temperature range -70°C to 45°C . Visual inspection of non-deployed DOMs reveals no indication of cracks (“crazing”) after after more than 10 years, and studies of the long-term optical efficiency of deployed DOMs reveal no measurable aging effects (section 3.7).

To reduce effects of the ambient South Pole magnetic field (550 mG, 17° from vertical) on the PMT collection efficiency, a mu-metal cage surrounds the PMT bulb up to the neck. It was constructed as a wire mesh with typical wire spacing 66 mm and wire diameter 1 mm, blocking about 4% of the incident light, and delivered by ITEP Moscow. Without such a shield, this PMT exhibits 5–10% lower collection efficiency, poorer single photoelectron resolution, and gain variations of 20% depending on azimuthal orientation, for a South Pole magnetic field strength and orientation [29]. With the shield in place, the interior magnetic field is 2.8 times smaller than the external field, pointing mostly along the axis and therefore reducing efficiency by less than 2% for this type of PMT.

Other interior DOM components are held in place by attachment to the PMT, mostly via screws into a molded plastic collar glued around the PMT neck. The PMT Base Board is soldered directly to the PMT pins.

2.2.4 High voltage supply and divider

The PMT high voltage subsystem consists of a resistive voltage divider circuit (PMT Base Board) directly solder-mounted on the PMT and a separate High Voltage Control Board. The High Voltage Control Board includes a DAC and ADC for setting and reading out the PMT high voltage, connected to the Main Board with a digital interface, as well as the high voltage generator itself.

The high voltage generator is a custom encapsulated module (#9730A) designed by EMCO High Voltage (California). The maximum high voltage is 2047 volts, specified for up to 30 μ A current. The voltage setting, in steps of 0.5 V, is controlled by the DAC output, and the actual voltage is monitored via a high-impedance divider and the ADC. The output ripple is less than 1 mV, and stability is better than 1 V RMS. Power consumption of the high voltage supply is less than 300 mW at full load.

The generator output is carried to the PMT Base Board [26] via a high voltage coaxial cable. The voltage divider, designed for low power consumption, presents a total resistive load of 130 M Ω . The PMT is operated with cathode at ground potential, so the anode signal output is AC-coupled using a 1:1 bifilar-wound toroid transformer mounted on the Base Board; this toroid was modified once during DOM production in order to reduce distortion of high-amplitude signals (section 3.2). The transformer secondary is then wired to the Main Board analog input with a coaxial cable. The single photoelectron (SPE) output waveform has been described in ref. [26]. With a 100 Ω load connected to the transformer, and operating at standard PMT gain of 10^7 , the SPE peak voltage before front-end amplification is approximately 8 mV with a FWHM of 7–8 ns. Several effects combine to increase the FWHM of digitized SPE waveforms to \sim 13 ns (peak \sim 5 mV).

2.2.5 Main board and delay board

The Main Board, designed at Lawrence Berkeley National Laboratory, has been described in detail in [27]. Essentially an embedded single-board data-acquisition computer, the Main Board interfaces to other boards as shown in figure 3 and provides many key functions of the DOM, including:

1. Control of all the devices inside the DOM, including the high voltage power supply for the PMT, the flasher board, and various sensors (pressure, temperature, power supply voltage monitor). Also supplies necessary DC power to the subsystems.
2. Digitization of the PMT waveforms, using a custom integrated circuit (ATWD: Analog Transient Waveform Digitizer [27]) and a continuously sampling fast ADC (fADC).
3. Providing computational functions, including PMT gain calibration, compressing digitized waveforms, temporarily storing the data, and creating time-stamped data packets.
4. Communicating with the data acquisition (DAQ) system on the surface.
5. Exchanging timing pulses with the surface DAQ to calibrate the internal DOM clock.
6. Exchanging LC pulses with the adjacent DOMs.
7. Uniquely identifying each DOM by providing a Main Board ID generated from a component device identifier.

8. Providing an onboard adjustable low-intensity optical source for direct calibration of PMT gain and timing, and hosting a daughter board with an adjustable high-intensity optical source for inter-module calibration.

The data flow starting from the PMT is shown in figure 5. PMT waveforms are amplified and compared to a discriminator threshold. Two discriminators are available, the SPE discriminator that is used for in-ice DOMs and typically set to a voltage threshold corresponding to 0.25 PE, and an MPE discriminator used for the larger-amplitude signals in IceTop. A discriminator crossing begins a “launch” of the high-speed waveform capture and digitization. Each DOM is equipped with two ATWD chips, and each chip is provided with three different amplifier gains with nominal values of 16, 2, and 0.25 in order to completely cover the dynamic range of the PMT output (up to 150 mA, or 7.5 V, when saturated). A fourth ATWD channel on each chip is used for calibration inputs and is not normally read out. The ATWD chips are configured to sample the input voltage at 300 Msps and operate by analog storage of waveform samples in switched capacitor arrays of depth 128, followed by a 10-bit digitization [30]. In order to record the waveform starting from before the discriminator threshold crossing, the signal is first routed through the Delay Board. Here a total delay of about 75 ns is accomplished by an approximately 10 m-long, 0.25 mm-wide serpentine copper trace embedded in the dielectric and sandwiched between ground planes. The highest-gain channel is used for most pulses, and lower-gain recordings are also retained as needed when pulses reach 75% of the range of a higher-gain channel, in order to avoid any loss of information due to digital saturation. The mean amplifier gains of all deployed DOMs for high-, medium-, and low-gain channels are 15.7 ± 0.6 , 1.79 ± 0.06 , and 0.21 ± 0.01 respectively.

The ATWD recording duration is 427 ns. This is sufficient for reconstructing light produced within tens of meters of a DOM, but photons from farther away may arrive over a broader time interval due to the optical scattering of the ice. Such distant signals are also lower in amplitude, and the information is captured in the 10-bit 40 Msps fADC. The fADC samples continuously, and the FPGA is programmed to save an interval of $6.4 \mu\text{s}$ after the launch. Its amplifier provides a dynamic range comparable to the high-gain ATWD channel, but has extra pulse shaping to accommodate the lower sampling speed. An example of a digitized waveform with multiple pulses is shown in figure 6.

Every digitizer launch results in a “hit” record. Hits are transferred from the FPGA to SDRAM lookback memory (LBM) via Direct Memory Access (DMA), and the Main Board CPU bundles them and sends them on request to the surface computers. The amount of information included in a hit depends on whether a signal was also detected in one of the neighboring DOMs. In case of an isolated signal (no coincidence), only a time stamp and brief charge summary are sent, and the digitization process is aborted. Conversely, when a nearest or next-to-nearest neighbor DOM also signals a launch within $\pm 1 \mu\text{s}$ (local coincidence), the full waveform is compressed and included in the hit record. The LC signaling operates via digital pulse codes sent on the extra wire pairs described in section 2.2.2.

As explained in ref. [27], two sets of ATWD chips are operated alternately in order to reduce deadtime; the second ATWD is available to launch during the digitization step of the first, after a re-arm delay of 50 ns. Significant deadtime only occurs after two back-to-back launches and depends on how many ATWD channels are digitized, and whether the initial hit had an LC condition. Since the full waveform is not needed in the absence of LC, the digitization can be aborted early, and the ATWD channels can be cleared and reset. The timing sequence for back-to-back hits is shown in figure 7.




Figure 5. Data flow diagram for recording and processing of PMT waveforms in the DOM to form “Hit Records” that are sent to the surface DAQ computers. As shown by dashes, full waveform data are only included when neighbor DOMs report time-coincident signals above the SPE discriminator threshold. Additionally, data from low-gain channels are omitted for waveforms that are within range of higher-gain channels.

The total accumulated deadtime for each individual DOM is measured by counting discriminator crossings when both ATWDs and the fADC are not acquiring data. This deadtime varies seasonally based on the atmospheric muon flux [24]. The median fractional deadtime during a high-rate period for in-ice DOMs is 6.6×10^{-5} , for IceTop low-gain DOMs is 7.2×10^{-6} , and for IceTop high-gain DOMs is 3.2×10^{-3} .

2.2.6 Flasher board

Each DOM contains an LED Flasher Board, which is used to generate light *in situ* for a variety of calibration purposes [23, 31], including:

1. Verifying the timing response of the DOMs throughout the analysis software chain.
2. Measuring the position of the deployed DOMs in ice.
3. Measuring the optical properties of the ice.
4. Verifying the performance of shower reconstruction algorithms in measuring position, direction, and energy.




Figure 6. The same signal sampled in the ATWD (top) and the fADC (bottom): the ATWD recording duration is 427 ns whereas the fADC recording duration is 6.4 μ s. Energy reconstruction in IceCube uses the charge and time recorded in the waveform [21].

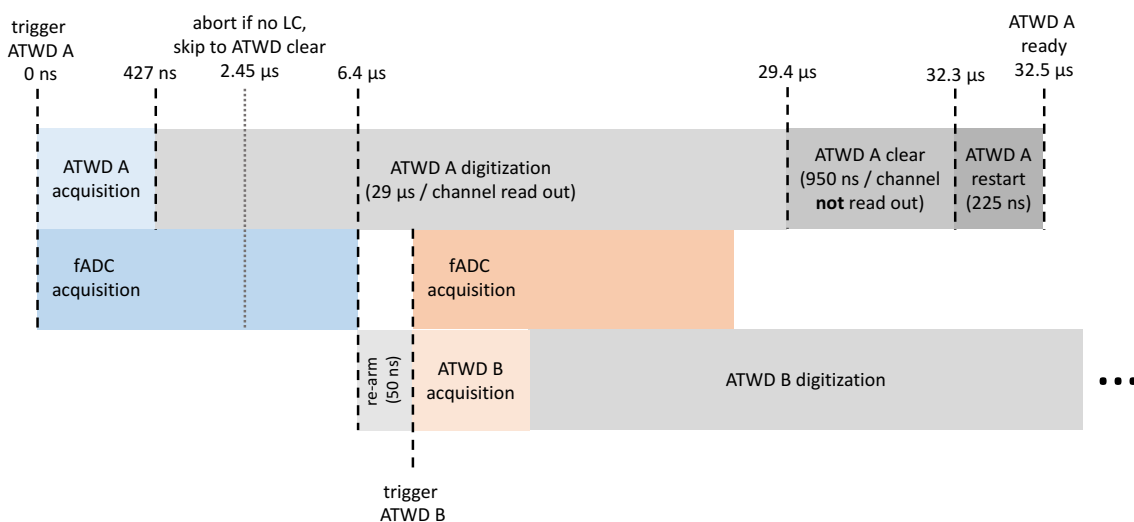


Figure 7. Timing of ATWD and fADC acquisition and associated deadtime, for back-to-back HLC hits with one ATWD gain channel of four digitized and read out. The horizontal (time) axis is not to scale.

The standard Flasher Board is included in every DOM except the “color DOMs” described below. It is an annular board fitted with 12 LEDs (ETG-5UV405-30) specified with output wavelength 405 ± 5 nm. Laboratory measurements with sample DOMs yield a peak at 399 nm and spectral width 14 nm (FWHM) when measured at -20° C, where the peak wavelength is shifted by -1 nm compared to room temperature [23]. The LEDs are arranged in six pairs, evenly spaced around the board with a 60° separation between adjacent pairs. One LED in each pair is pointed downward at an angle of 10.7° ; after refraction through the DOM glass and into the ice, the LED emits light horizontally into the ice. The other LED is tilted upward at an angle of 51.6° ; after refraction the tilted LED emits light upward at an angle of 48° , close to the Cherenkov angle in ice. The angular emission profile of the flasher LEDs was measured in the lab by rotating a PMT connected to an optical fiber pickup around the DOM; the readout angle was recorded using the resistance of a potentiometer at the rotation axis. The angular emission profile of each LED has a FWHM of 30° in air and is modeled as a Gaussian emission profile with $\sigma = 13^\circ$. After refraction through the DOM glass and into the ice, the emission profile is modified to $\sigma = 9.7^\circ$ in the polar direction and 9.8° in the azimuthal direction for the tilted LEDs, and $\sigma = 9.2^\circ$ in the polar direction and 10.1° in the azimuthal direction for the horizontal LEDs. About 10% of the light is emitted outside the Gaussian beam, modeled by a secondary profile proportional to $(1 + \cos \alpha)$, where α is the angle away from the LED axis.

The LEDs are controlled via individual high-speed MOSFET drivers; the flasher circuit diagram is shown in figure 8. The LEDs can be turned on individually or in any combination of the 12, by setting bits in a configuration parameter. The photon output of each LED depends on the width and amplitude of the driving current pulse, which are controlled as common values for all enabled LEDs in each DOM (figure 9). The pulse width parameter controls the width up to a maximum of 70 ns; for sufficiently short driving current pulses the light output narrows to 6 ns (FWHM) with 10% afterglow decaying over 15–20 ns. The brightness parameter (0–127) controls the driving voltage between 4.5 and 15 V, which yields a peak current up to 300 mA through the LED and current-limiting resistor. By varying brightness and width settings as well as the number of LEDs enabled, DOMs can generate flashes from 10^6 to 1.4×10^{11} photons, similar to the total light from neutrino interaction showers between 7 GeV and 1 PeV energy. The low end of this dynamic range requires fine tuning of driving parameters in order to operate LEDs very close to threshold.

The LED current waveforms are recorded in an auxiliary ATWD channel, supplying a rising edge time that also establishes the onset of the optical pulse (after a known turn-on delay). The repetition rate is programmable up to 610 Hz. Although flashers can be operated in multiple DOMs in the same run, the DAQ does not support time-synchronized flashing of LEDs on different DOMs, so coincident flasher events happen only by chance.

Sixteen “color DOMs” (cDOMs) are fitted with multi-wavelength Flasher Boards; 8 are deployed on String 79 in the center of IceCube, and 8 are deployed on String 14 on the edge of IceCube. Each cDOM includes LEDs with nominal wavelengths of 505 nm, 450 nm, 370 nm, and 340 nm. The LEDs are arranged in six pairs as on the standard flasher board, three pairs of 370 nm and 450 nm and three pairs of 340 nm and 505 nm, but all LEDs point outward horizontally. The properties of the LEDs on the standard DOMs and the cDOMs are given in table 1. Differences between the nominal and measured wavelengths are within expectations based on normal LED variation from the manufacturer.

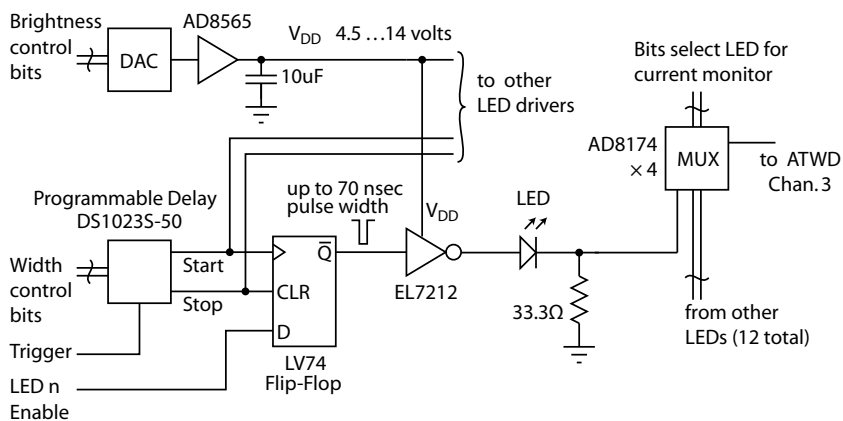


Figure 8. LED flasher circuit diagram for one of twelve LEDs, including current pulse monitor (simplified).




Figure 9. Light output from flasher LED pulses (relative to maximum), depending on brightness parameter (B) and width configuration parameters. Additional dynamic range is available by enabling from 1 to 12 individual LEDs per DOM.

2.3 Production and testing

Approximately 5800 DOMs were built and tested, with approximately 5500 delivered to the South Pole. The DOMs satisfied stringent requirements, needed to ensure reliable operation in the deep ice for at least 20 years. As hot-water drilling was the principal driver of the deployment timeline, the DOM production schedule was structured to supply DOMs as needed and to avoid any inventory shortfall. The production was implemented in a 3-stage approach. Stage 1 was the production of the initial 400 DOMs at three sites: one in the United States at the University of Wisconsin–Madison’s Physical Sciences Lab (PSL, Stoughton, Wisconsin) and two in Europe (DESY, Zeuthen, Germany, and Stockholm University, Sweden). DOM testing was performed at PSL, DESY, and Uppsala University, Sweden. This quantity of DOMs was sufficient to verify production readiness, supply instrumentation for the first year drilling plan, and validate the design after a deployment season.

Table 1. Properties of the standard IceCube flasher LED (tilted (t) and horizontal (h)) and the cDOM LEDs, including wavelength λ , emission FWHW σ in air, DOM polar angular emission FWHM in ice σ_θ , and DOM azimuthal angular emission FWHM in ice σ_ϕ .

LED	nominal λ (nm)	measured λ (nm)	σ air ($^\circ$)	σ_θ ($^\circ$)	σ_ϕ ($^\circ$)
ETG-5UV405-30	405	399	30.0	9.7 (t) 9.2 (h)	9.8 (t) 10.1 (h)
UVTOP335-FW-TO39	340	338	51.0	36.1	42.9
NS370L_5RFS	370	371	55.2	39.1	42.9
LED450-01	450	447	6.8	4.8	5.3
B5-433-B505	505	494	6.4	4.5	4.9

During Stage 2, material and supplies were procured, and another 1000 DOMs were produced and tested. Finally, Stage 3 involved procurements, integration, and testing of the remaining DOMs.

DOM production utilized a formalized process to track each DOM through to the end of the testing process, with each step recorded in a DOM Process Traveler. Technicians were trained and certified to perform DOM integration and test tasks, and each site had separate quality control personnel. Commercial materials were confirmed to be fully tested by the suppliers, and regular visits were made to key vendors. Measurement equipment was calibrated and records maintained that verified traceability to a reference standard. DOM integration took place in an electrostatic discharge (ESD)-, temperature-, and humidity-controlled environment. The introduction of these manufacturing protocols based on electronics industry best practices enabled each production site to work independently yet produce DOMs that performed identically.

DOM integration started with the attachment of the PMT collar to the PMT. The collar provides a mounting point for the electronic boards inside a DOM. The PMT was then mounted into a special jig for precise placement inside the bottom glass hemisphere. In parallel, the mu-metal shield was placed inside the bottom hemisphere, and the gel was mixed and poured into the same hemisphere. The gel was then degassed by placing under a partial vacuum, in order to avoid bubbles and crazing in the gel. After degassing, the PMT was placed in the gel, and the gel was allowed to cure for 48 hours. After curing, the PMT Base Board was soldered onto the leads of the PMT. Separately, the PC Board Stack was assembled by attaching the Delay Board, Main Board, Flasher Board, and High Voltage Control Board together. The Board Stack was then mounted onto the PMT collar, the penetrator assembly was mounted into the top hemisphere, and the two halves of the sphere were joined. With the entire assembly under a bell jar, the spheres were evacuated and backfilled with dry nitrogen, a butyl rubber sealant applied around the seam, and the seam covered with wide plastic tape. The interior gas pressure was reduced to 0.5 bar (at room temperature) so that the seal remains tight even at the assumed minimum ambient South Pole air pressure of 0.64 bar.

As the DOMs are not serviceable after deployment, an extensive testing protocol (Final Acceptance Testing, or FAT) including temperature-cycling and cold-soaking ensured that bad modules and early component failures were identified and repaired before shipping. This testing at pro-




Figure 10. Final Acceptance Testing (FAT) temperature profile, including DOM testing steps performed at each stage.

duction sites was performed in Dark Freezer Labs (DFLs), light-tight walk-in enclosures capable of sustaining temperatures down to -55°C . Main Board and DOM functionality was tested by self-diagnostic software installed on each module. Other tests included gain calibration, dark noise monitoring, LC validation, and booting after cold-soaking. Optical sensitivity, time resolution, and linearity were measured using external light sources fed into the DFLs via optical fibers and diffused over the DOM lower hemispheres at each testing station. Sensitivity was measured primarily with a monochromator-selected white lamp, time resolution was measured with a pulsed diode laser (405 nm), and linearity was measured with a bright pulsed LED source (405 nm).

A typical FAT temperature and testing cycle is shown in figure 10. The initial pass rate of DOMs during FAT was 92%. The primary causes of failures were elevated noise rates detected during the low-temperature soak, functional issues on the Main Board or Flasher Board, and PMT gain instability. The majority of failing DOMs were retested and subsequently passed, while DOMs with serious issues were repaired if possible and retested prior to shipment. After the successful completion of FAT, DOM deployment harnesses were attached (figure 4) and the DOMs individually packed for shipment.

All DOMs were also re-tested at the South Pole before final deployment, to screen out any modules damaged during transit. The self-diagnostic program and dark noise measurements were performed with the DOMs still in their packing boxes, covered with additional light-tight material. Of the approximately 5500 DOMs shipped to South Pole, about 30 (0.5%) were returned to the U.S. after failing on-ice testing.

2.4 Reliability

As of 2016, 5397 of the 5484 deployed DOMs (98.4%) are operating in data-taking mode in the data acquisition system; the remaining 87 DOMs have failed (table 2). We classify DOM failures into two broad categories: failures during deployment and freeze-in, and failures during subsequent operation. The majority of the failures (55) occurred before post-deployment commissioning; we hypothesize that these are primarily attributable to cable failures, water leaks, or freeze-in damage.

Table 2. Number of DOM failures during deployment/freeze-in and after commissioning during detector operation, as well as DOMs with various issues causing them to be operated in a non-standard data-taking mode. The majority of DOMs with non-standard LC settings function normally but have a neighbor with an issue.

DOM failures	N	DOMs in non-standard mode	N
deployment / freeze-in	55	single functional ATWD	12
post-commissioning	32	reduced PMT gain	1
		non-standard LC	158
total	87	total	171

32 DOMs have failed after commissioning, and we include in this count modules on a wire pair taken out of service when the partner DOM on the same pair failed. No particular pattern in the failures is observed, other than they are typically during non-standard operation or an exceptional event: a power outage, calibration run, or flash filesystem upgrade. The most recent two DOMs failed on May 23, 2013, losing communications after a power outage. Diagnosis of DOM failures beyond identifying electrical shorts is challenging.

A total of 171 DOMs have developed issues that affect their data-taking configuration but are still usable. For example, DOMs with a single functional ATWD chip have a higher deadtime but are otherwise good. The LC settings of functional DOMs adjacent on a string to dead DOMs must also be modified. In most cases, local coincidence is disabled in the direction of a dead neighbor DOM, but in a few cases a malfunctioning DOM can still be configured to forward the local coincidence signals along the string even if it is not used in data-taking. These are enumerated in table 2.

We can estimate the surviving fraction of DOMs 25 years after the original deployment, assuming a constant, random failure rate after freeze-in. Specifically, we calculate the Wilson score binomial confidence interval [32] of survival probability using the post-commissioning failure rate of DOMs. The estimated survival fraction as a function of time is shown in figure 11. Currently we estimate the mean failure rate to be $4.1 \pm 1.2 \text{ yr}^{-1}$, resulting in a survival fraction in 2030 of $97.4 \pm 0.3\%$. While this simplified model does not account for an increase in failure rate due to component aging, the recent observed failure rate since detector completion of 1.7 yr^{-1} is significantly lower than the mean predicted rate, since the failure rate during construction was higher. We attribute this to infant mortality and/or to improved operational protocols that minimize the number of DOM power cycles. DOMs are not regularly power-cycled during data-taking but only when required to resolve an intermittent problem.

3 Calibration and stability

Regular calibration of each individual DOM allows translation of the recorded signals into a measurement of incident Cherenkov light at a particular time, the basis of particle event reconstructions. The primary DOM calibration routine is DOMCal (section 3.1), a calibration procedure run yearly on in-ice DOMs and monthly on IceTop DOMs that provides the constants used to convert DOM waveforms into physical units (section 3.2). Global time calibration across the array of DOMs is provided by the RAPCal procedure (section 3.3) that runs during data-taking. The relative and




Figure 11. Actual and predicted fraction of surviving DOMs versus time, based on an assumed constant post-freeze-in failure rate. The dotted lines indicate the central and 95% CL estimates. Increases before 2011 are due to deployments of new strings.

absolute optical efficiency of the DOMs is determined both by laboratory and *in situ* measurements (section 3.4). The stability of these calibrations is also relevant, as not every calibration parameter can be tracked continuously during data-taking (sections 3.5 to 3.7). Understanding the statistical properties and time-dependence of the background “dark noise” is also important for low-energy neutrino analyses, supernova searches, and detector simulation (section 3.8).

The determination of the optical properties of the ice is crucial to the understanding of IceCube data but is beyond the scope of this work. Further details can be found in refs. [23, 33].

3.1 DOMCal

The calibration of the PMT waveforms recorded by the DOMs, i.e. translation of digitized samples into voltage and time, as well as accounting for the gain of the PMT itself, is achieved via DOM-by-DOM calibration constants determined by the DOMs themselves. The primary reference inputs to the DOM calibration software (DOMCal) are a precision electronic pulser circuit providing known charges; a circuit providing a reference DC bias voltage level; the 20 MHz oscillator on the Main Board, used as the timing reference; and single photoelectrons, either from background “dark noise” photons or a low-luminosity LED on the Main Board. The Main Board oscillator is manufactured by Vectron International (model C2560A-0009 TCXO) and has a certified stability of roughly 1×10^{-11} for a sample interval of 5 seconds [27]. The Main Board LED is used to illuminate the PMT on the same DOM, as opposed to the LED flashers described in section 2.2.6 which stimulate other DOMs. This LED produces zero to a few tens of photoelectrons at the PMT [27]. Analysis and fitting of the calibration results is done by the DOMCal software, with the results being transmitted from each DOM to the surface as an XML file of fit parameters.

Because the operating conditions for the in-ice DOMs are so stable, DOMCal is only run once a year on the full detector. IceTop DOMs are calibrated once per month in order to track temperature-dependent changes. Calibration is typically performed on half the detector at a time, the other half remaining in data-taking mode in case of a transient event (such as a supernova) during

calibration. The total calibration run length is a little under three hours. Because the calibration procedure produces light, these runs are excluded from normal analyses.

First, the discriminator circuits used to launch the DOM are calibrated using the electronic pulser circuit, which is capacitively coupled into the analog front-end before the Delay Board [27]. This calibration is later refined using actual PMT waveforms, once the PMT gain is known. Next, the ATWD voltage levels are calibrated by sweeping the input DC bias voltage and recording the corresponding waveforms at each DC level. Because of the slight variations in the ATWD circuits, this calibration provides an individual linear relationship between ATWD counts and voltage for each time sample and gain channel.

The average ATWD offset at zero input voltage, or baseline, is needed for charge integration and can in principle be determined using the previous ATWD calibration. However, in practice, these baselines are extremely sensitive to the operating condition of the DOMs, and since data-taking conditions cannot be exactly replicated while running DOMCal, the baselines used during data-taking are determined instead by using averaged forced triggers taken during a normal run (section 3.5). DOMCal can still use baselines that it measures for numerical charge integration during the calibration procedure.

The highest-gain channels of each ATWD are calibrated using the electronic pulser circuit, and then the gains of the other ATWD channels and the fADC are determined by varying the pulser output and comparing the charges simultaneously recorded in multiple gain channels. This relative calibration is later refined using PMT waveforms stimulated by the Main Board LED.

The tunable ATWD sampling speed is calibrated by digitizing the Main Board oscillator waveform and recording the number of clock cycles as a function of ATWD speed setting; the relationship between configuration setting and sampling speed is fit with a quadratic polynomial. The fADC sampling speed is slaved to the 20 MHz oscillator, which is used as a timing reference, and so is not separately calibrated. The relative timing of the ATWD and fADC waveforms is determined using the electronic pulser circuit; non-linear fitting of the digitized waveforms to known pulse templates is required in order to determine the fADC time offset to the required accuracy. The transit time of the PMT and Delay Board as a function of PMT high voltage is determined by calculating the delay between the digitized current pulse through the Main Board LED and the observed light pulse in the ATWDs.

The PMT gain as a function of high voltage is calibrated using background “dark noise” photons — the charge e prior to amplification is quantized and known. At each voltage level, a histogram of many observed charges is recorded, where the charge is determined by numerical integration of the waveform. Each histogram is fit with the sum of an exponential and a Gaussian function (figure 12). The peak of the Gaussian component is used to determine the amplification of the PMT at each voltage, and a linear fit of $\log_{10}(\text{gain})$ versus $\log_{10}(\text{voltage})$ allows the high voltage of each PMT to be tuned to the desired operating point (10^7 for in-ice DOMs; figure 13). Small corrections (3 – 5%) to the gain of each DOM are determined using charge distributions recorded by the Processing and Filtering system (section 6.5) during normal data-taking. These corrections are used to eliminate a systematic difference in the charge as determined by DOMCal’s numerical integration and the waveform pulse unfolding used in data processing.




Figure 12. Sample SPE charge spectrum at a gain of $10^{7.5}$ as recorded by DOMCal and fit *in situ* with the sum of an exponential and a Gaussian function.




Figure 13. Distribution of PMT high voltages at 10^7 gain for in-ice DOMs.

3.2 Waveform calibration

The deposited charge is the basis for all energy measurements in IceCube. The calibration constants needed are 1) the pedestal, the value the digitizer reads when the input voltage is zero, and 2) the gain, the input voltage change required to increase the readout value by one count. An ATWD waveform consists of readouts from 128 effectively independent digitizers, while an fADC waveform consists of successive outputs of a pipelined digitizer (256 samples). Two calibration constants per digitizer are needed to turn each of these raw readout values (an integral number of ADC counts) into a voltage, from which the deposited charge is calculated.

The baseline is the mean of the pedestals of each sample in a waveform, while the pedestal pattern is the deviation of each sample's pedestal from the common baseline. The fADC has no additional pedestal pattern, but for the ATWD, it is important to distinguish between the pedestals of individual samples and the common baseline of the entire waveform. In order to remove the sample-dependent offset, the DOM subtracts the pedestal pattern from the ATWD waveform before




Figure 14. Averaged beacon waveforms from the fADC (left) and ATWD (right) of an IceCube DOM. The waveforms are an average of approximately 1000 beacon launches. The baseline, which is the mean value of the beacon waveform, is shown as a horizontal red line. Typical raw SPE waveforms are inset.

sending it to the surface. Thereafter, both ATWD and fADC waveforms can be calibrated by first subtracting the common baseline from each sample, then multiplying by the gain. Correct charge measurement and energy reconstruction depends on correct measurement of the baseline, as discussed in section 3.5.

The pedestal pattern is computed by the DAQ software at the beginning of each run by averaging 200 forced-readout waveforms. Accidental single photoelectrons in the individual waveforms are averaged out, but a single coincident bright event can result in an offset in the pedestal average. In order to avoid such light contamination, a second average is computed, and the autocorrelation coefficient of the difference between the two pedestal averages is computed at two different lags. This autocorrelation test detects averages in which a single waveform contains at least a 15 PE pulse (approximately 0.1 PE in the average). If light contamination is detected, the procedure is repeated. The shift between the baseline of the pairs is also calculated to verify that the baseline is stable. This procedure ensures that fewer than 1 DOM in 1000 runs will contain a contaminated baseline.

The baseline is set to about 10% of the maximum value of the digitizer counts in order to capture signals that go below the baseline. Since 2012, the baseline value for each DOM is configured (section 6.4.4) in order to ensure stability. The baseline value differs for each digitizer channel in each DOM, ranging from 112 to 161 counts in the fADC and 109 to 172 counts in the ATWD. The baselines for each digitizer channel in each DOM are measured with beacon hits, forced triggers that are collected during normal data acquisition at a rate of 1.2 Hz per DOM in in-ice DOMs and 4.8 Hz in IceTop DOMs. Beacon waveforms from the fADC and ATWD of a typical DOM are shown in figure 14.

The digitizer gain is measured by DOMCal (section 3.1), a single value for the fADC waveform and a sample-dependent value for the ATWD waveform. The calibrated waveform voltage is then

$$\text{voltage} = (\text{counts} - \text{baseline}) \times \text{gain} \quad (3.1)$$

The waveform start time is then corrected for the transit time, the average time it takes a pulse to propagate through the entire PMT and Delay Board. The transit time correction t_{transit} is dependent on the PMT high voltage V :

$$t_{\text{transit}} = \frac{m}{\sqrt{V}} + b \quad (3.2)$$

where m and b are determined by DOMCal. The typical value of m is $2000 \text{ ns} \cdot \sqrt{V}$, and the typical value of b is 80 ns, which includes the 75 ns delay line of the Main Board. The typical transit time is therefore 130 ns. Waveform start times from the second ATWD chip and the fADC are further corrected for the delay Δt with respect to the first ATWD chip as determined by DOMCal, so the total start time correction is $t_{\text{transit}} + \Delta t$.

Finally, the waveforms are corrected for the effects of droop from the transformer that couples the Main Board to the PMT output. The toroid coupling effectively acts as a high-pass filter on the PMT output that makes the tails of the waveforms “droop” and even undershoot the baseline. This effect is temperature-dependent and is larger at lower temperatures. The droop correction inverts the effect of the high-pass filter and eliminates the undershoot in the waveform tails. This is done by calculating the expected reaction voltage from the toroid at each sample time, and adding the reaction voltage to the calibrated waveform to compensate. The reaction voltage decays according to a temperature-dependent model of the transformer behavior. When a readout contains consecutive launches from the same DOM, the reaction voltage at the end of the last launch is used to correct for the residual droop in the following launch. DOMs use two types of toroid transformers: the “old toroid” with a short time constant that was used in early DOM production, and a “new toroid” with a longer time constant that produces less distortion. Of 5484 DOMs deployed in IceCube, 1204 are the old toroid type, and 4280 are the new toroid type. The full correction is modeled with two time constants, where the DOM’s transient response $\tilde{\delta}(t)$ to an input signal $\delta(t)$ is given by

$$\tilde{\delta}(t) = \delta(t) - N((1 - f)e^{-t/\tau_1} + f e^{-t/\tau_2}) \quad (3.3)$$

where the first time constant τ_1 is given by

$$\tau_1(T) = A + \frac{B}{1 + e^{(-T/C)}} \quad (3.4)$$

In eqs. (3.3) and (3.4), T is the temperature, and the constants A , B , C , N and f were determined empirically with a dedicated analysis. For the old toroids, the second time constant $\tau_2 = 0.75\tau_1$. For the new toroids, τ_2 is 500 ns.

3.3 RAPCal

The Reciprocal Active Pulsing Calibration (RAPCal) is a continuously-running procedure for translating hit timestamps based on the individual free-running DOM clocks to the synchronized surface clocks in the ICL, in order to establish a standard time base for the array to $O(\text{ns})$ accuracy. Subsequently, the ICL-based hit timestamps can then be translated to UTC. The scheme used for the time transfer involves transmission of a pulse to and from each DOM and is shown in figure 15. The base implementation is described in [27]; we describe here the details of the algorithm and validation of the DOM relative time calibration.




Figure 15. Diagram of the RAPCal time calibration scheme (not to scale). Each transmitted and received pair of signals is timestamped (T_{tx}, T_{rx}) in the local clock domain, and by the symmetry of the situation the midpoints T_C between transmitted and received timestamps are synchronous.

While data transmission is paused, a bipolar pulse is generated and sent to each DOM over the power/communications wire pair. After this pulse has been received, a bipolar pulse — having the same shape as the one generated on the surface — is generated in the DOM and sent back to the surface. The pulses are timestamped using the local clock count at each transmission and reception point. The symmetry of the pulses sent in each direction and symmetry of the down- and up- cable transmission enables the time transfer from each free-running DOM clock to the surface clock and hence to UTC, without prior knowledge of the cable length or its dispersive properties.

The received dispersed bipolar pulses are digitized by a 20 MHz communications ADC and timestamped using the local clock count. The received pulse timestamp corresponds to a fixed delay after the start of the final pulse recovery. A RAPCal pulse sequence to and from each DOM therefore results in a series of four timestamps $T_{tx}^{\text{DOR}}, T_{rx}^{\text{DOM}}, T_{tx}^{\text{DOM}},$ and T_{rx}^{DOR} (figure 15). These timestamps initially count periods of the DOM’s 40 MHz clock and of the 20 MHz clock of the surface electronics (DOR, section 6.3.2); they are translated into common units for the calculations described below.

Next, in order to improve the precision beyond the sampling speed of the communications ADC, a “fine delay” correction f to T_{rx}^{DOM} and T_{rx}^{DOR} is calculated by interpolating to find the negative-going baseline crossing, with the baseline voltage calculated using the initial samples of the waveform (figure 16):

$$\tilde{T}_{rx} = T_{rx} - f . \quad (3.5)$$




Figure 16. Digitized RAPCal pulse as received by a DOM, after cable dispersion. The zero-crossing of the baseline-subtracted pulse is used as a fine-delay correction f to the received timestamp T_{rx} .

The midpoints $T_C^{\text{DOR,DOM}}$ as shown in figure 15 are then determined, where

$$T_C = \frac{T_{\text{tx}} + \tilde{T}_{\text{rx}}}{2}. \quad (3.6)$$

T_C^{DOM} and T_C^{DOR} then identify a single point in time across the two clock domains.

The next stage of the process is to translate an arbitrary DOM hit timestamp t to UTC. In typical operating conditions, the RAPCal procedure is repeated once per second for each DOM. We use the two nearest RAPCal results before and after t to derive a linear relationship

$$\text{UTC}(t) = (1 + \epsilon) (t - T_C^{\text{DOM}}) + T_C^{\text{DOR}} + \Delta. \quad (3.7)$$

The slope $(1 + \epsilon)$ accounts for drifts in the DOM clocks and is calculated from the T_C pairs of the two neighboring RAPCal results:

$$1 + \epsilon = \frac{T_{C,2}^{\text{DOR}} - T_{C,1}^{\text{DOR}}}{T_{C,2}^{\text{DOM}} - T_{C,1}^{\text{DOM}}}. \quad (3.8)$$

The median magnitude of ϵ is 1.34×10^{-6} . Finally, because the timestamps T^{DOR} count the sub-second time offset into the current UTC second, the UTC time offset Δ of the previous 1-second boundary, provided by the master clock, is added. Details on the master clock system are provided in section 6.3.2.

The stability and repeatability of the calibration is monitored by tracking the cable delay from RAPCal, determined by

$$T_{\text{cable}} = \frac{1}{2} \left((T_{\text{rx}}^{\text{DOR}} - T_{\text{tx}}^{\text{DOR}}) - (1 + \epsilon) (T_{\text{tx}}^{\text{DOM}} - T_{\text{rx}}^{\text{DOM}}) \right). \quad (3.9)$$

A representative distribution of T_{cable} from one DOM over an 8-hour data-taking run is shown in figure 17, with a standard deviation of 0.6 ns. Individual RAPCal measurements in which T_{cable} deviates by more than 10 ns from a moving average are discarded.




Figure 17. Distribution of one-way cable delays from multiple RAPCal measurements on DOM 22-60 (bottom of String 22), shown with a Gaussian fit.

The time calibration through the entire data acquisition and software processing chain is verified using the LED flashers. During commissioning, performed on each string after deployment, all 12 LEDs on each DOM are flashed simultaneously at maximum brightness, and the arrival times of the first photons at the DOM above the flashing DOM are recorded. Given the vertical spacing of 17 m on a standard IceCube string and the group index of refraction of 1.356 in ice at 400 nm [34], the expected light travel time from the flashing DOM to the DOM above is 77 ns. In DeepCore, the DOM vertical spacings are 10 m and 7 m, corresponding to light travel times of 45 ns and 32 ns respectively. The mean light travel time to the DOM above for all flashing DOMs in ice is shown in figure 18. The mean arrival time agrees with the expectation for the DeepCore DOMs. For the standard DOMs, the observed light travel time is about 3 ns longer than the expected light travel time, due to the effects of scattering in the ice over the longer distance. The accuracy of the photon arrival time with respect to the arrival time of any other photon is measured using the difference between arrival times at the two DOMs above the flasher, which eliminates any uncertainty in the flasher source time. Figure 18 shows the distribution of the difference in arrival times at the two DOMs above the flasher for the DOMs with 7 m spacing, as the DOMs with larger spacing are more affected by scattering in ice. The width of the distribution is 1.7 ns, so the measured timing accuracy is $1.7/\sqrt{2}$ ns = 1.2 ns. Muons are also used to verify the time calibration [27], including the absolute time difference between the IceTop surface stations and the in-ice DOMs [35].

3.4 DOM optical efficiency

A baseline value for the photon detection efficiency has been established by combining PMT measurements at 337 nm and a separate model of wavelength- and angle-dependent effects. Absolute sensitivity measurements were performed on 13 IceCube PMTs, using single photons from a primary 337 nm laser beam of calibrated intensity [26]. The results agree well with independent Hamamatsu measurements of sensitivity in the range 270–730 nm, which were then normalized to the 337 nm measurement. The resulting quantum efficiency at the center of the PMT at 390 nm is 25%. A full simulation model of the DOM includes the wavelength dependence of the PMT response,




Figure 18. Left: time from flasher emission to detection at the DOM above for 17 m vertical spacing (red), 10 m vertical spacing (blue) and 7 m vertical spacing (green). The expected light travel times in ice for each distance are marked with vertical lines. Right: arrival time difference (blue) between the two DOMs above the flasher for DOMs with 7 m vertical spacing. The vertical red line denotes expected light travel time in ice, and the mean of a Gaussian fit to the distribution (dashed blue line) is 33 ns, due to the effects of scattering in ice. The width of the distribution is 1.7 ns.




Figure 19. Position dependence of the response of a DOM, normalized to the absolute efficiency at 365 nm. The x-y coordinates measure distance from the center of the DOM.

optical absorption in the DOM glass and gel, discriminator threshold effects, and photocathode non-uniformity. The angular dependence is dominated by the shape of the photocathode and its response variation away from the center, which was measured for 135 bare PMTs [26]. The efficiency was also measured for 16 fully integrated DOMs, including glass and gel; the response of an example DOM is shown in figure 19. The efficiency measurement used the 337 nm nitrogen laser, as well as LEDs at 365 nm, 470 nm, 520 nm, and 572 nm due to the low transparency of the glass at short wavelengths.

The efficiency model as determined from the laboratory measurements was supplemented with *in situ* measurements using Cherenkov light from muons in the ice. The efficiency measured *in situ* includes effects of the cable shadowing the DOM and the refrozen “hole ice” in the vicinity of the DOM. In one study, low-energy muons (median energy 82 GeV) with well-understood light emission were selected to illuminate the downward-facing PMTs in the ice as directly as possible. The number of photons detected at different distances from these muons was then compared to predictions of the simulation model [21]. Based on this and other *in situ* analyses, the central value for the efficiency was adjusted upward by 10% in the simulation, compared to the baseline. For physics analyses, the normalization of the absolute DOM efficiency is generally included as a nuisance parameter with prior uncertainty of 10% that includes other uncertainties related to generation and propagation of the light. Additional laboratory measurements on assembled DOMs are in progress, including wavelength- and angle-dependent effects, and are expected to reduce uncertainties [36].

The absolute calibration at 337 nm was performed at room temperature on a small subset of IceCube PMTs. The relative efficiency of all assembled DOMs was separately measured as part of production testing (section 2.3), using a 405 nm pulsed laser and a system of fibers and diffusers to illuminate DOMs within 50° of the PMT axis. Using this system, the relative efficiency of DeepCore DOMs (high quantum efficiency type) was measured to be higher by a factor of 1.39 compared to standard IceCube DOMs, agreeing well with the manufacturer-specified value of 1.40. Further *in situ* studies using muons yielded a factor of 1.35 [22], an effective value including the Cherenkov spectrum folded with the different wavelength sensitivity curves of the two types of PMTs. The production testing system also established that the efficiency change from room temperature to $-40\text{ }^{\circ}\text{C}$ is less than 1% when gain is maintained at the design value of 10^7 .

3.5 Baseline voltage stability

The beacon hits from which the digitizer baselines are derived (section 3.2) are monitored continuously throughout the year. The average values of the beacon baselines are very stable, with shifts of no more than 0.2 counts per year, corresponding to 0.018 mV in the fADC and 0.025 mV in the high-gain ATWD channel. The baseline shifts from May 2015 to April 2016 are shown in figure 20.

Every year when the detector is recalibrated, adjustments in the SPE discriminator threshold can cause shifts of up to 0.6 counts (0.54 mV) in the fADC baselines, for reasons not completely understood. ATWD baselines are unaffected by the SPE discriminator setting. Correcting recorded waveforms for the effect of transformer coupling as described in section 3.2 has the side effect that small DC offset errors are converted into an apparent PMT current that increases linearly in time over the course of a few microseconds. The effect is stronger for DOMs with old-type toroid transformers, where a baseline error of 0.6 counts can turn an actual deposited charge of 1 PE into a measured charge of over 2 PE with an unphysical time structure. The observed distortion in a simulated single photoelectron charge due to fADC baseline shifts is shown in figure 21. To avoid this problem, whenever the discriminator thresholds are changed, the fADC baselines are re-measured, and the values used for calibration are updated. As long as the discriminator thresholds are unchanged, the baselines are stable to within 0.2 counts as shown above, and no charge distortion is seen at that level of baseline stability.

The baselines can be sensitive to radio-frequency interference (RFI). In 2009, RF signals from the COBRA meteor radar transmitter broadcasting at 46.3 MHz [37] appeared as sinusoidal or

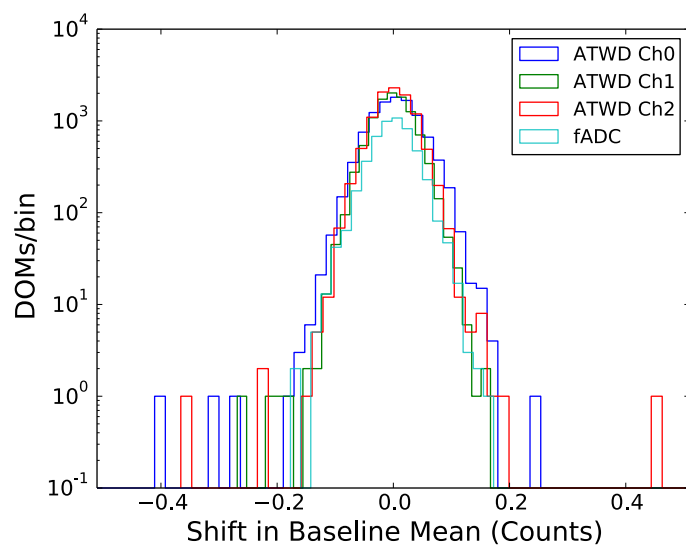


Figure 20. Distribution of shifts in baseline values in all ATWD channels and the fADC for all DOMs from May 2015 to April 2016. The DOM configuration was unchanged during this period.

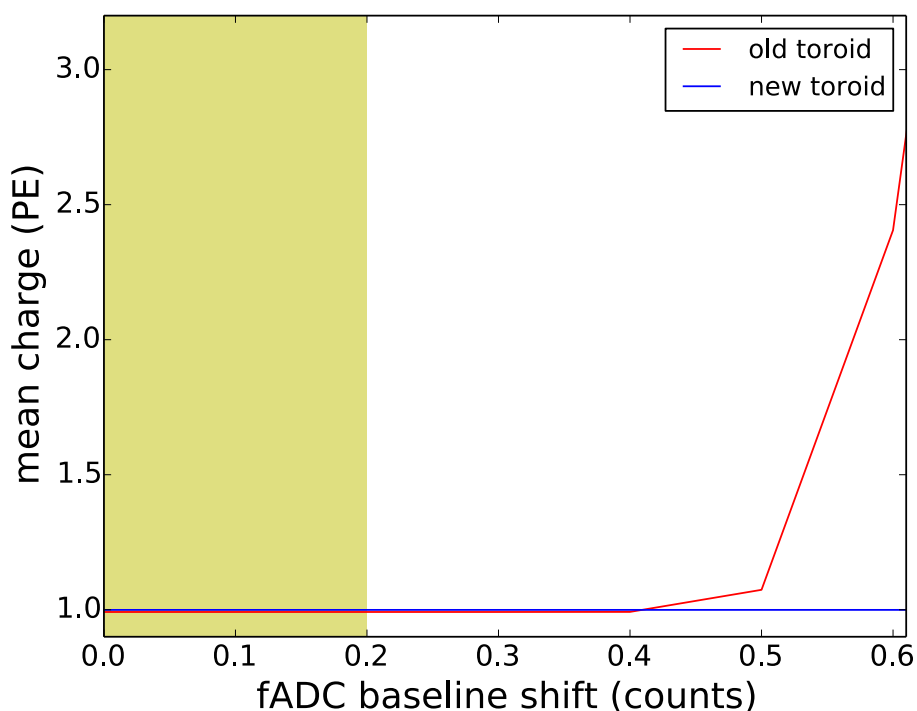


Figure 21. Reconstructed charge from a simulated single photoelectron deposit as a function of fADC baseline shift, for both old- and new-toroid DOMs. The shaded region indicates the observed range of baseline variation from the nominal value in ice; no observable distortion in the charge spectrum is seen at these values.




Figure 22. Fractional ATWD channel amplifier gain shifts, from 2015 to 2016 (left) and 2011 to 2016 (right). Channels 0, 1, and 2 are high gain, medium gain, and low gain respectively.

spiky modulations in the waveform baselines. To mitigate this effect, one DOM was operated without PMT high voltage and used to monitor the RF broadcasts and tag periods of RFI in order to avoid data contamination. Also in 2009, RFI from DOM 68-42 affected nearby DOMs after it failed during DOM calibration and appeared to begin sparking, resulting in sinusoidal distortions in the baselines of neighboring DOMs. The meteor radar transmitter is no longer operating, and investigations into RFI from more recent radar activities at South Pole (SuperDARN [38]) have not revealed any measurable interference, either in DOM baselines or in RAPCal time stability.

3.6 Gain stability

The gain stability of the DOM, or the stability of the amplified charge from converted photons, depends on a number of factors including stability of the PMT high voltage, Main Board channel amplifiers, and the digitizers. We can examine these subsystems using both historical calibration results and by tracking the SPE charge during data-taking.

Variations of the front-end electronic amplifiers or the digitizers themselves can potentially lead to changes in the overall gain of the DOM electronics. The stability is demonstrated by comparing the Main Board channel amplifier gains from sets of calibrations taken from 2011 to 2016 (figure 22). From year to year, the amplifier gain calibration is repeatable to 0.1%, 0.2%, and 0.5% in the high-gain, medium-gain, and low-gain channels respectively. Since detector completion in 2011, a small systematic shift of -0.3% is visible in the low-gain channel, but this is corrected by updating the calibration constants of each DOM.

The stability of the total gain, including both the PMT and front-end electronics, is monitored during data-taking using the single photoelectron peak of the charge distribution on each DOM from cosmic ray muons. A Gaussian + exponential function is fit to the peak region as in figure 23, and the mean of the Gaussian is tracked throughout the year. The threshold, defined as the bin in the histogram for which 99% of the area of the histogram is contained in the sum of all higher bins, is also tracked through the year. The peak position is calibrated to 1 PE and is stable to within 0.01 PE for 95% of all DOMs, as shown in figure 24. Over 99% of DOMs show no measurable




Figure 23. Charge distribution of a typical in-ice DOM. The threshold is marked in red, and a Gaussian + exponential fit to the SPE region is shown in blue. The mean of the Gaussian is used to monitor the gain stability.




Figure 24. Distribution of the mean of the Gaussian fit to the SPE peak (left) and the shift in this value for each in-ice DOM (right) between May 2015 and April 2016.

change in the threshold as long as the discriminator thresholds are unchanged; these settings are only changed once per year.

There are about 12 DOMs that show unpredictable, abrupt shifts in the SPE peak position of 0.05 PE or more. Figure 25 shows the time history of the SPE peak position of one of these DOMs over 4 months. The peak shift corresponds to increases or decreases in the multi-photoelectron (MPE) scaler rate, where the scaler counts the discriminator crossings. This indicates that the SPE peak shift is indeed caused by a change in the DOM gain. However, the SPE scaler rate is stable to within 2%, indicating that the probability to detect single photons is effectively unchanged.

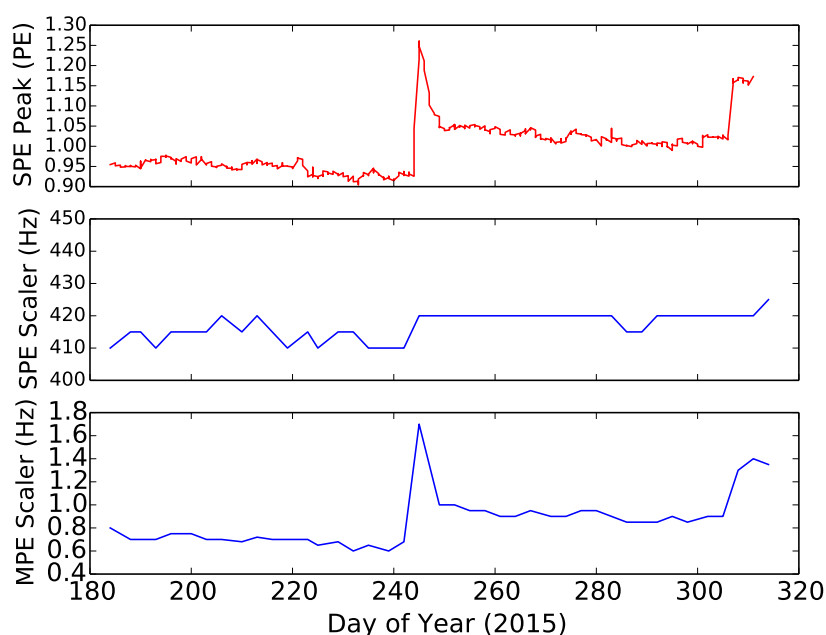


Figure 25. Mean of the Gaussian fit to the SPE peak (top) and the SPE scaler rate (middle) and MPE scaler rate (bottom) from July 2015 to November 2015 for a DOM that shows unpredictable gain shift behavior. The DOM configuration was unchanged during this period.

Long-term stability of the PMTs can be tracked by examining any change in the high voltage vs. gain calibration over time. The fractional change in gain for all in-ice DOMs over a five-year time period is shown in figure 26. For most DOMs, the PMT gain is stable to within $\pm 3\%$ over the time period shown, with a median gain shift of 0.3%, while approximately 1% of DOMs have a gain shift of more than 10%. Any shifts are tracked with regular calibration using the methods previously described.

3.7 Optical efficiency stability

The detector response in IceCube is verified with low energy muons as described in ref. [21]. The detector response is monitored in each run using the track detection probability (TDP) calculated from bright muon tracks with more than 30 hits in IceCube. The muon tracks are reconstructed using the likelihood methods described in ref. [39], but charge and time information from the DOM under study are excluded from the reconstruction. The TDP is defined for each DOM as the ratio of the number of detected tracks within 100 m of the DOM to the total number of tracks within 100 m of the DOM. This ratio depends both on the optical properties of the ice near the DOM and the optical efficiency of the DOM. We do not attempt to separate these effects in the TDP, but rather use the TDP to monitor the overall stability of the detector response. Figure 27 shows the TDP on String 80, which includes both standard and HQE DOMs; the TDP is 20–25% higher for HQE DOMs than for neighboring standard DOMs, whereas the quantum efficiency is 35% higher. The TDP is stable to within 1% since 2012, when the baselines were stabilized by being set in the DAQ configuration. Figure 27 shows the difference in the TDP for all DOMs between a run in 2012 and a run in 2015.




Figure 26. Distribution of the fractional changes in PMT gain from 2011 to 2016 for all in-ice DOMs, at the 2011 operating high voltage for each DOM.




Figure 27. Left: track detection probability (TDP) in String 80, with a mixture of standard DOMs (blue crosses) and HQE DOMs (red circles). The variation with depth is due to the depth-dependent optical properties of the ice. Right: shift in TDP for all in-ice DOMs between 2012 and 2015; standard DOMs are in blue, and HQE DOMs are in red. The mean of the Gaussian fit to the TDP shift in the standard DOMs is -0.1% , and the mean TDP shift in the HQE DOMs is 0.05% .

The detector response stability is also measured with the *in situ* light sources in IceCube. Both the in-ice calibration laser [31] and the flasher LEDs show less than 1% difference in the total charge collected between 2012 and 2015.

3.8 Dark noise

The vast majority of background hits result from dark noise, i.e. effects that lead to the emission of an electron from the cathode of the PMT in the absence of a photon source external to the DOM. Dark noise is a complex phenomenon with numerous possible sources, including thermionic




Figure 28. Time interval between successive hits for all next-to-top layer DOMs (DeepCore excluded). The line is an exponential fit to the Poissonian regime between 7 and 15 ms.

emission, electronic noise, field emission within the PMT, Cherenkov light from radioactive decays, and scintillation / luminescence in the glass of the PMT and pressure sphere. The average total in-ice hit rate is 560 Hz for DOMs with standard PMTs and 780 Hz for high quantum efficiency DOMs. The contribution from cosmic-ray muons, estimated as the in-ice HLC hit rate, is small and decreases with depth from 25 Hz to 5 Hz.

The dark noise can be characterized as a combination of uncorrelated (Poissonian) noise pulses with a rate between 230 Hz and 250 Hz, and a correlated component, with a pulse rate from 280 Hz to 340 Hz. A comparison at low temperature of DOM dark noise to that of a bare PMT suggests that the majority of the noise originates from the glass pressure sphere. Cherenkov light from ^{40}K decays contributes to the uncorrelated noise component, and thus the potassium content of the glass was limited. Measurements of early samples indicated a K_2O concentration of 0.03% by weight, roughly corresponding to 100 Bq of beta decays per sphere.

The correlated noise manifests itself as an overabundance of short time intervals between hits in a single DOM compared to the Poisson expectation (figure 28). The temperature dependence of the noise rate (figure 29) was determined by combining a measured temperature profile of the South Pole ice cap [40] with a fit of the Poissonian expectation of the total dark noise rate to every individual DOM, and was verified in lab measurements. The temperature of the in-ice DOMs, measured on the Main Board, increases with depth from $-31\text{ }^\circ\text{C}$ to $-9\text{ }^\circ\text{C}$ ($10\text{ }^\circ\text{C}$ above the ambient ice temperature).

The short time intervals are clustered in bursts (uninterrupted sequences of time intervals less than 3 ms) with an average number of hits per burst increasing from 3.3 at $-10\text{ }^\circ\text{C}$ to 3.8 at $-30\text{ }^\circ\text{C}$. A study with forced-readout data shows that the phenomenology of the correlated noise component




Figure 29. Dark noise rate of DOMs in ice as a function of temperature, obtained from untriggered HitSpool data. Each data point represents the average of 12 DOM layers from 78 strings (DeepCore excluded).

Table 3. Characteristics of noise components in IceCube DOMs, adapted from ref. [44].

Noise Component	Origin	Distribution	Parameters
afterpulses	PMT	Gaussian	$\mu = 6 \mu\text{s}$, $\sigma = 2 \mu\text{s}$
uncorrelated	thermal noise, radioactive decay	Poissonian	$\lambda \approx 250 \text{ Hz}$
correlated	luminescence (?)	log-normal	$\mu = -6 [\log_{10}(\delta t/\text{sec})]$, $\sigma = 0.85 [\log_{10}(\delta t/\text{sec})]$

in IceCube is in general in good agreement with results reported in ref. [41], but an unambiguous physical explanation is still to be confirmed. Stimulation of glass samples with radioactive sources results in scintillation [42], suggesting that luminescence of the glass triggered by the radioactive decay of ^{40}K and other elements in the pressure sphere is responsible. Naturally-occurring cerium in the glass is a candidate for the active scintillator.

The various sources of dark noise present in a DOM can be separated using the time between successive untriggered hits in HitSpool data (section 6.4.1), as shown in figure 30. An overview of the various noise components and their parameterizations is given in table 3. Late-time correlated afterpulses, a common feature of PMTs, is attributed to ionization of residual gases by electrons that were accelerated between the dynodes. Although afterpulses occur at various time delays, this component is parametrized here with a single average timescale. A noise model incorporating these various sources is used for detector simulation [43].




Figure 30. Histogram of time differences between successive hits from HitSpool data of DOM 15 on string 27 (blue) on a logarithmic scale in order to visualize the different noise components [45]. A fourth subdominant component centered at $100\ \mu\text{s}$ is not parameterized and is still under study.




Figure 31. Mean and standard deviation of the supernova scaler distribution for the entire detector over the course of the first three years of the completed IceCube [46]. Spikes are due to calibration runs which injected light into the detector, and to fluctuations in individual DOMs.

The evolution of the dark noise contribution to the total rate over time was investigated using the supernova scaler stream [28, 46], effectively the summed dark noise rate of the detector with an artificial deadtime applied to reduce the effect of correlated noise (section 6.4.6). A long-term exponential decay in the noise rate is visible; this may be caused by decreasing triboluminescence [47] arising from the initial “freeze-in” of newly deployed DOMs, impurities introduced during the drill process, or a combination of the two effects. The decay is especially recognizable in the standard deviation of the scaler distribution (figure 31), which decreases by 25% over the course of the three years, with the decrease increasing with depth [48]. Changes in the mean rate are initially dominated by the decay component and later by the seasonal variation of atmospheric muons. The noise rate decay is most pronounced when selecting strings that were deployed in the final construction season, as shown in figure 32.

Since the dark noise components are not correlated between DOMs, the dark noise rate change is not prominent in local coincidence hits, which are dominated by atmospheric muons (figure 33). Thus, the detector trigger rate as well as many higher-level reconstruction variables are relatively




Figure 32. Mean and standard deviation of the supernova scaler rate of strings deployed in the last deployment season (austral summer of 2010/2011) [46], where changes in the mean rate are still dominated by the decay component. Spikes are due to calibration runs which injected light into the detector, and to fluctuations in individual DOMs.




Figure 33. Mean and standard deviation of the local coincidence hit rate distribution [46]. Changes in dark noise rate do not contribute significantly to changes in the local coincidence hit rate, which is dominated by seasonal and weather-related changes in the atmospheric muon flux.

isolated from the noise rate variations. Nevertheless, the total hit rates for each DOM are tracked over time and updated yearly in a database for simulation and reconstruction purposes. Seasonal variations of these rates are below the 1% level.

4 Cable systems

The IceCube detector may be viewed as a digital network of optical sensors. The IceCube cable systems form the physical backbone of this network, simultaneously supplying power to the DOMs and bi-directional digital communications between the DOMs and the readout hardware at the surface. Copper was chosen over optical fiber at an early stage in the project, based on a previous successful demonstration of the technology with the AMANDA digital-readout optical module string [49] and concerns with mechanical robustness of fiber during freeze-in in the deep ice.

4.1 Design

The cable system comprises the following assemblies: the in-ice cable, IceTop cables, the surface junction box (SJB), the surface cables, and the patch cables (figure 34). The in-ice cable, 2505 m long, is deployed into the ice along with 60 DOMs that are attached to connectors at 30 break-outs spaced 34 m apart. An adjacent pair of DOMs is connected to a distinct twisted wire pair.

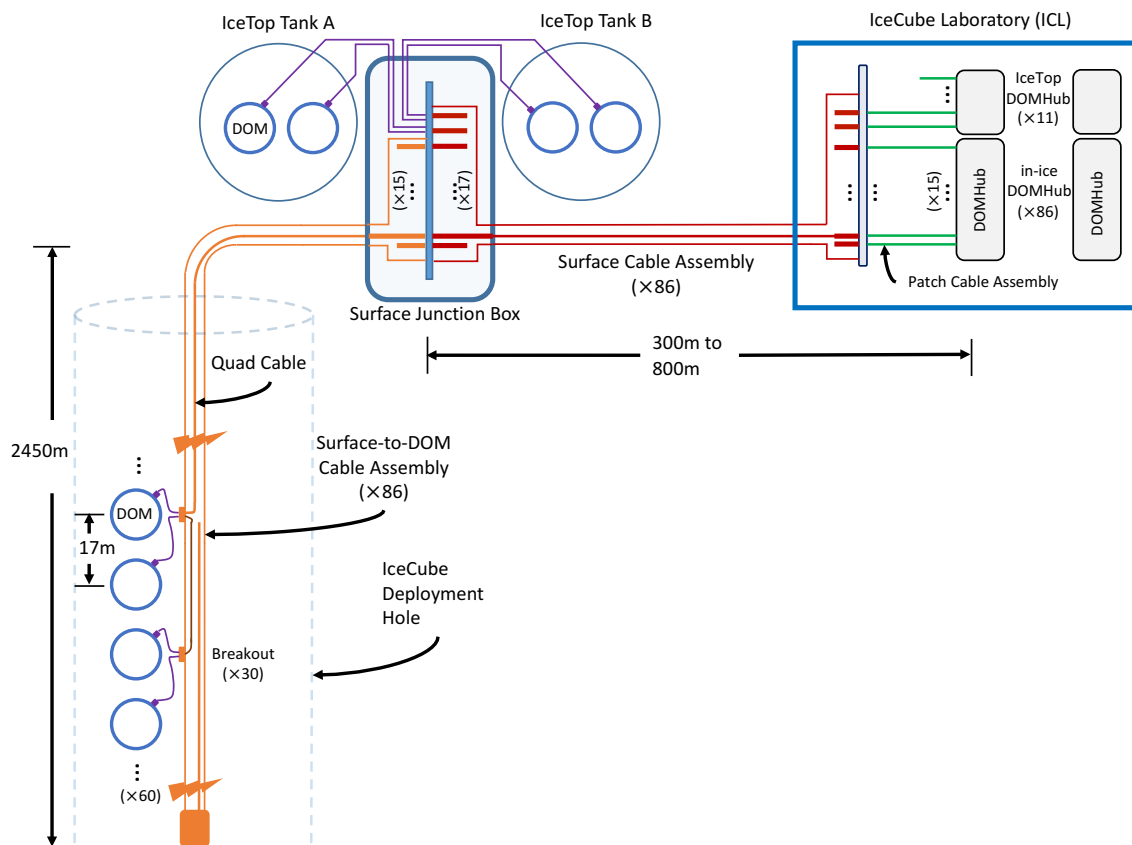


Figure 34. Schematic of the IceCube cable network architecture.


Two wire pairs are combined into four-conductor quad cables meeting stringent electrical performance requirements; the quad arrangement provides enhanced cross-talk immunity and improved mechanical stability during freeze-in compared to a twisted pair.

The 60 DOMs on each cable require a total of 15 quads. An additional 5 quads in the in-ice cable provide for special instrumentation connections, a spare quad, and local coincidence connections between adjacent DOMs. The in-ice cable terminates at the SJB, located just below the snow surface between the IceTop tanks. The SJB is a stainless steel enclosure that houses the in-ice cable and surface cable connections. IceTop cables also terminate and connect to the surface cable in the SJB. The surface cable is trenched 1 m deep into the surface of the snow between the SJB and the ICL. The surface cables vary from 300 m to approximately 800 m in length depending on hole location. The surface cables are pulled up two cable towers into the ICL and terminate at patch panels where the individual quads are separated and connected to patch cables that finally terminate at the DOMHub, the computer that receives DOM signals in the ICL (section 6.2). Each DOMHub services a single in-ice string or 8 IceTop stations. IceTop DOMHubs service fewer DOMs than in-ice DOMHubs due to the higher rates recorded by IceTop DOMs [24].

The IceCube cable system design had to meet a number of stringent mechanical, electrical and functional requirements, due to the extreme conditions of installation at the South Pole. In particular, the in-ice cable was required to withstand hydrostatic pressure of up to 450 bar during

Table 4. Electrical requirements for the cable quads.

Characteristic	Specification	Conditions
Attenuation	≤ 20 dB	At 1.0 MHz end-to-end
Near-end crosstalk suppression	≥ 50 dB intra-quad	At 2 MHz
	≥ 24 dB quad-quad	At 100 MHz
Far-end crosstalk suppression	≥ 30 dB intra-quad	At 2 MHz
	≥ 24 dB quad-quad	At 20 MHz
Differential impedance	$145\Omega \pm 10\Omega$	At 1.0 MHz
Loop resistance	$\leq 160 \Omega$	
Dielectric breakdown voltage	≥ 2000 VDC, $\leq 1 \mu\text{A}$ leakage	

**Figure 35.** In-ice cable cross section, with a nominal 46 mm diameter and mass of 2 kg/m.

deployment. Handling temperatures were specified down to -40°C for all cable components; deployment took place between late November and mid-January each year, and temperatures never dropped below -40°C during these months. The primary electrical requirements for the quad cables are provided in table 4.

The in-ice cable was also required to be less than 50 mm in diameter, weigh less than 2 kg / m, have a minimum static bend radius of 40 cm, carry a maximum static tensile load of 10 kN, and have a breaking strength of 40 kN. During deployment, the maximum tension on the cable was less than 8 kN, taking into account the buoyancy of the DOMs.

The in-ice cable includes 20 quads, 2 polyethylene fillers to maintain structural symmetry, Kevlar outer and inner core strength members, a copper foil shield with drain wires to provide electrical noise isolation, and a water block that prevented water or ice from damaging the symmetry of the cable (figure 35). The surface cable has a similar construction to the in-ice cable with two exceptions: first, the inner Kevlar strength member was replaced with additional conductors needed to service the IceTop tanks; and second, the water block was not required.

A competitive proposal process resulted in two main suppliers for the IceCube cable system. Ericsson Network Technologies (Hudiksvall, Sweden) was chosen to produce the raw cable, and SEACON Brantner & Associates (El Cajon, California) was chosen to manufacture the cable assemblies. The raw cable was manufactured and tested to meet all required specifications prior to

spooling. The in-ice cables were then wound onto custom metal spools, while the shorter surface cables were wound onto wooden spools. SEACON provided the breakouts and connectorized the top of the in-ice cable and both ends of the surface cables. Glass-reinforced-epoxy SEACON XSJJ connectors were used for the in-ice cable to DOM interface; these connectors are water pressure-rated to 690 bar. The 30 breakouts per in-ice cable were installed by slicing open the cable, cutting the correct quads, adding the XSJJ connectors, waterproofing and overmolding the connectors, and then resealing the cable. SEACON also attached 120 Yale Cordage (Saco, Maine) YaleGrips to the in-ice cable that served as mechanical attachment points for the 60 deployed DOMs per cable. The separate quads in each cable were terminated with Military Standard round, metal-shell connectors. After cable assembly was complete, each cable was individually tested for electrical continuity and subjected to high potential (high voltage) testing.

4.2 Installation

Installation of the IceCube Cable System at the South Pole each season was broken into five distinct tasks: surface cable trenching and installation into the ICL and SJB; in-ice cable deployment and connection to the SJB; IceTop freeze control power and communications installation at the ICL and IceTop trench; patch cable installation in the ICL; and testing, troubleshooting and repair after all connections were made. Surface cables were installed early in the season to support IceTop tank freeze control operations, which were necessary to ensure clear crack-free ice in the IceTop tanks [24]. A 1-m-deep trench was mechanically dug between the ICL and each IceTop station. After the cable was placed in the trench, it was pulled into the ICL and connected to an IceTop DOMHub. Later, just prior to filling the tanks with water in the field, the connections were made between the surface cable and the SJB. The IceTop DOMs with their long 17 m penetrator assemblies were connected to the appropriate quad located in the center of the surface cable via the SJB. Two additional cables were connected between the SJB and the tanks that provided power and communications to the IceTop freeze control units. The in-ice cable was prepared for installation while the Enhanced Hot Water Drill was coming up the hole in its reaming phase (section 5.1). After the drill was removed, the end of the in-ice cable was pulled into the drill tower, and DOM deployment commenced (section 5.2). After DOM deployment was complete, the in-ice cable end was taken off the spool and dragged to the IceTop station where its connection to the surface cable in the SJB was made. Finally, patch cables were installed in the ICL between the individual quads and the DOMHub in the ICL. End-to-end string commissioning then commenced.

5 Drilling and deployment

5.1 Hot water drilling

Transforming a cubic kilometer of Antarctic ice into an astrophysical particle detector required drilling 86 boreholes approximately 60 cm in diameter to a depth of 2500 m. Hot water drilling was the only feasible technology to provide rapid access to the deep ice on this scale. Instrumentation was deployed into the water-filled boreholes, becoming frozen in place and optically coupled with the surrounding ice sheet. The 5 MW Enhanced Hot Water Drill (EHWD) was designed and built specifically for this task, capable of producing the required boreholes at a maximum rate of one

Table 5. EHWD System Characteristics.

Specification	Value
Total Power (Thermal + Electrical)	5 (4.7 + 0.3) MW
Maximum Drill Speed	2.2 m/min
Maximum Ream Speed	10 m/min
Water Flow (delivered to main hose)	760 L/min
Water Temperature (delivered to main hose)	88 °C
Water Gauge Pressure (at main pumps)	7600 kPa

Table 6. Average and peak performance for a 24-hour lifetime hole of 2500 m depth. Peak performance corresponds to string 32 from the 2009–2010 drilling season, which had the fastest drilling + reaming time.

Specification	Avg. Value	Peak Value
Total Fuel, ^a AN-8	21,000 L	15,000 L
Time to Drill/Ream	30 hr	27 hr
Hole Production Cycle Time ^b	48 hr	32 hr

^aTotal Fuel includes deep drilling/reaming and firm drilling.

^bHole Production Cycle Time is the elapsed time from start of one hole to start of the next hole.

hole per 48 hours. This section contains an abbreviated description of the drill; a more detailed description can be found in ref. [50].

The procedure involved a drilling phase to create the initial hole, followed by an upward reaming phase to give the hole a targeted diameter profile. During drilling, the location of the drill was recorded with an onboard navigational pack consisting of a Paroscientific pressure sensor (model 8CB4000-I) to measure depth, two liquid pendulums to measure tilt, and a fluxgate compass to measure orientation. The hole diameter was larger than the DOM (35 cm diameter) to compensate for closure from refreezing to provide sufficient time to deploy instrumentation, with contingency time for delays. The elapsed duration from the end of drilling until the hole closes to less than the DOM diameter is referred to as the hole lifetime. IceCube drilling was completed in seven field seasons (approximately 21 months total time). Peak performance occurred in the 2009–2010 season with 20 holes drilled (early November to mid-January). The drill specifications and performance are shown in table 5 and table 6. Figure 36 shows the drilling and reaming time for a typical hole in the 2009–2010 season.

Hot water drilling was not practical in the firm layer, the 50 m-thick surface layer of compressed snow above the ice, because the firm absorbs hot water without melting a hole. An independent firm drill was designed that consisted of a conical drill head wrapped in copper tubing through which hot water circulated at high speed, melting the snow by contact. Hot water drilling commenced after the firm drill completed the first portion of the hole. The firm drill had its own water supply and heater in order to operate in parallel with hot water drilling at other holes. The firm drill reliably achieved a rate of 2 m per hour [50].




Figure 36. Drilling and reaming profile for string 41, from the 2009–2010 deployment season. Zero time is maximum drill depth, so drilling time is negative and reaming time is positive.

The EHWD system was implemented across two separate sites. The Seasonal Equipment Site (SES) provided electricity and a stable supply of hot pressurized water, and the Tower Operations Site (TOS) was where the hole was drilled. The two sites were linked by long cables and insulated hoses. The SES comprised generators, water tanks, pump and heating buildings, a central control building, mechanical and electrical shops, spare parts storage, and the Rodriguez well system (Rodwell), which provides water to compensate for the volume change from ice to water in each hole [51]. Hoses and cables connected SES subsystem buildings together, and wherever necessary custom electrically heated hoses were installed, providing an effective freeze mitigation strategy. The TOS included the drill tower and attached operations building as well as the hose and cable reels. There were two towers and one set of drill reels. After drilling, drill reels were moved to the next hole location, where the second tower had already been staged; the first tower stayed at its existing location to support deployment of the instrumentation. Once deployment had finished, the first tower could be moved to the next location while drilling at the second tower was underway. This leapfrog sequence of the tower structures reduced hole turnover time and allowed nearly continuous drilling operations.

Due to the size and complexity of the SES, it remained stationary throughout each drill season. At the end of the drill season, the SES was decommissioned and repositioned for the following drilling season. The distance between the SES and TOS had a practical limit, referred to as reach, which defined the boundary of a seasonal drilling sector. The maximum reach of the EHWD was 450 m, limited by pressure and voltage drop through the SES-TOS link.

Each drilling season started with a staggered arrival of drill crew members while the SES and TOS were excavated from accumulated snow drift and commissioned. Season startup tasks included SES and TOS warming and hookups, reinstallation of do-not-freeze equipment (such as motor drives, sensors, and some rubber items such as gaskets), generator commissioning, safety checkout and system tuning, initial (“seed”) water delivery to begin the drilling process, and Rodwell development. This phase typically took four weeks.

The production drilling sequence was to drill, ream, and move to the next location. The drill crews worked in three shifts per day of ten people per shift. Independent firm drilling stayed ahead

of deep drilling by at least two holes, and often the Rodwell and first few holes of the season had already been firn-drilled the prior season. Hole production rate was 48 hours per hole on average, and the quickest cycle time was 32 hours. The idle phase between holes was characterized by minimal flow through the system and included regular maintenance tasks and deployment of IceCube instrumentation.

System shutdown would begin approximately two weeks before the end of the austral summer season. Shutdown tasks included flushing the system with propylene glycol, blowing out the plumbing with compressed air, removing do-not-freeze equipment for warm storage, storing the TOS structures and other support equipment, and finally, moving the SES into place for the following season. Due to a strong safety culture and retention of experienced personnel, IceCube had only four reportable drilling-related safety incidents in approximately 52 on-ice person-years.

5.2 Deployment procedures and instrumentation

DOM deployment required 60 DOMs staged in the drill tower, various attachment hardware, pressure sensors, and the in-ice cable on a spool outside the drill tower. The hole diameter was logged before deployment using calipers on the drill head. After verification that the hole was ready for deployment and a pre-deployment safety check was carried out, the in-ice cable was pulled over the top of the tower above the open hole. Four 100-pound weights were connected together and attached via a 2.1 m steel cable to a DOM that had a 17 meter penetrator assembly (see figure 4). The weights and lowermost DOM were attached to the bottom of the in-ice cable via a 15.5 m steel cable. After lowering of the DOM and weights, the next DOM with a 1.8 m penetrator assembly was attached to the in-ice cable. The two DOMs were then electrically connected to the in-ice cable at the breakout. A Paroscientific pressure sensor (model 8CB4000-I) was attached just above the second DOM, and a Keller AG pressure sensor (model 300DS-340-003700) was attached in the middle of the instrumented section of the string. The pressure sensors were read out during deployment to confirm that the string was properly descending into the hole. DOMs with 17 m and 1.8 m penetrator assemblies were alternately attached to the in-ice cable, until all 60 DOMs were attached, after which the remaining 1.5 km of in-ice cable was lowered into the hole (the “drop” phase). The top of the cable was secured by an anchor trenched into the snow near the hole. After the cable was secure, its end was taken off of the spool and connected to the Surface Junction Box (SJB). An example of the time profile of deployment from pressure sensor attachment to string drop is shown in figure 37.

Dust loggers were deployed in selected holes in order to measure the concentration of dust and volcanic ash as a function of depth at various locations in the detector, a critical component of the overall measurement of the optical properties of the ice [23, 52]. A dust logger consists of a 404 nm laser line source emitted horizontally in a fan pattern, paired with an integrated photomultiplier tube and digital photon counter for light detection. Disposable dust loggers were deployed in two holes, and reusable dust loggers were deployed in six holes. Each disposable dust logger was attached to the main cable between the lowermost DOM and the weight stack, and an extension cable ran between the dust logger and the bottom breakout, just above the next-to-lowest DOM. In this mode, the hole was logged once during deployment, and the logger was not retrieved. The reusable dust logger was used to log holes just before deployment, using a separate winch. The reusable dust logger produced two logs for the hole, one in the up direction and one in the down direction, which were used for reciprocal calibration.




Figure 37. Depth of the lowermost DOM vs. time for String 48, deployed in the 2006–2007 season. The final depth is the preliminary Stage 1 depth measured during deployment (section 5.3.1), prior to the Stage 2 corrections derived from LED flasher measurements (section 5.3.2).

5.3 Geometry measurement

The geometry of the detector was determined using drill and survey data during deployment (Stage 1), and then corrected and refined using the LED flashers in ice (Stage 2). The IceCube coordinate system is defined relative to the South Pole grid coordinate system (Northings and Eastings) that moves with the ice sheet. Elevation is defined relative to Mean Sea Level. The IceCube coordinate system origin is located at 46500'E, 52200'N, at an elevation of 2900 ft (883.9 m). This origin is located close to the center of the instrumented volume of IceCube, about 2000 m below the surface of the ice. The y -axis of the IceCube coordinate system is aligned with the Prime Meridian (Grid North), pointing toward Greenwich, U.K. The x -axis of the IceCube coordinate system points 90° clockwise from the y -axis (Grid East). The z -axis is normal to the Earth's surface, pointing upward.

5.3.1 Stage 1 geometry measurement

The (x, y) -coordinates of the string were calculated using the position of the drill tower. Before deployment, when the drill tower was in position, at least three of the tower corners were surveyed from at least one control point. The coordinates for the center of the hole in the tower floor were calculated from the corner coordinates. The string was assumed to be vertical, so the (x, y) -coordinates from the tower were used at all depths on the string. The drill data show deviations of less than 1 m from vertical that are not included in the detector geometry but have been validated with flasher data for select strings, as discussed in section 5.3.3.

The depth of the lowest DOM on the string was calculated using pressure readings from the pressure sensor, converted to depth by correcting for the compressibility of the water in the hole and the ambient air pressure measured before the pressure sensor enters the water. The distance from the tower floor to the water surface was measured with a laser ranger during deployment. The vertical DOM spacings were also measured with a laser ranger aimed down the hole after each DOM attachment. All depths were converted to z -coordinates in the IceCube coordinate system.

5.3.2 Stage 2 geometry measurement

The LED flashers on the DOMs were used to correct the depths of the strings relative to the Stage 1 geometry measurements. The correction was calculated by flashing the horizontal LEDs on a DOM at the surrounding strings and finding the leading edge of the time distribution of the light recorded by the receiving DOM, denoted t_0 . The distance corresponding to the leading edge time is $d = c_{\text{ice}} \cdot t_0$, and the distances for all receiving DOMs are plotted as a function of the vertical distance between the flasher and the receiver, $z' = z_{\text{receiver}} - z_{\text{flasher}}$. The resulting curve describes a hyperbola, $d = \sqrt{D^2 + (z' - \Delta z)^2}$, where D is the horizontal distance between the flasher string and the receiver string, calculated from Stage 1 data, and Δz is the relative offset between the depths of the flashing and receiving string as shown in figure 38. The hyperbolic fit was performed simultaneously on multiple flasher-receiver pairs; in order to avoid bias, this global fit was performed multiple times on randomly selected sets of string pairs. An average offset was calculated for each string and applied as a correction to the Stage 1 z -coordinates of all DOMs on that string. The correction was typically less than 1 m relative to the Stage 1 data, but could be as large as 20 m (larger than the 17 m DOM spacing) in cases where the pressure sensor failed during string deployment before it recorded the final depth reading, resulting in a poor estimate of the Stage 1 depth. The Stage 2 depth correction was less than 3 m for 93% of the strings. The uncertainty on the correction varied from string to string but was typically less than 0.2 m. The uncertainty on the absolute z -coordinate is 1 m, based on timing residual studies from downgoing muons [35].

5.3.3 Trilateration validation of DOM coordinates

The Stage 2 geometry assumes that all strings are vertical. The location data from the drilling process show that the string is not perfectly vertical, although deviations from the vertical are typically less than 1 m. The (x, y) -coordinates from the drill location data at varying depth were validated on the DeepCore strings using a trilateration method. In this method, the 5 DOMs closest to the flasher on each of the three closest strings surrounding the flasher are selected, and a circle of radius $r = \sqrt{d^2 - (z')^2}$ is described around each receiving DOM, where d is the distance between




Figure 38. Left: distance d vs. relative depth z' as measured by flashers for two strings that were deployed at nearly the same depth with no anomalies during deployment. Right: the same plot for a different pair; in this case, the pressure sensor on the flashing string failed during deployment.

the DOM and the flasher calculated from the leading edge time of the received light, and z' is the relative depth of the flashing and receiving DOMs calculated from the method described above. With 15 circles, there are 125 possible combinations of three circles that can be used for trilateration. For each combination, the six intersection points of the three circles are calculated as shown in figure 39, and the (x, y) -position of the flashing DOM is taken to be the centroid of the three innermost intersection points. The final (x, y) -position is the average of the values calculated from each of the 125 combinations of circles. The error bars on the positions are 1σ from a Gaussian fit to all centroid coordinates; the x - and y -coordinates are fitted independently. The measured positions agree with the drill data within the error bars as shown in figure 39. Since the shifts from the nominal (x, y) -coordinates are less than 1 m, and the trilateration measurement is only practical in a few DeepCore strings where the interstring spacing is 40–70 m, these corrections were not applied to the IceCube geometry, which still assumes the vertical string geometry from Stage 2.

5.3.4 Inclinometers

The stability of the East Antarctic ice sheet is one key feature making it suitable for a particle physics installation. However, ice sheets can deform mechanically under the stress of their own weight. To measure the stability of the ice, IceCube installed 50 inclinometers in the deeper sections of the array: two electrolytic inclinometers (Applied Geomechanics, Santa Cruz, CA) housed in ~ 20 -cm-long aluminum pressure modules were installed during the 2007–2008 season, and 48 micro-electromechanical (MEMS) tilt sensors (ADIS16209, Analog Devices) were added to DOMs deployed in seasons 2009–2010 and 2010–2011. This inclinometer array was intended not only for monitoring detector geometry but also as a unique glaciology experiment that permits three-dimensional tracking of deep ice flow and evaluation of complex full-stress models that cannot be effectively tested under laboratory conditions [53].

Figure 40 shows six years of data from 42 of the DOM-embedded MEMS tilt sensors and eight years of data from the two electrolytic inclinometers. Measurements are started after 400 days in




Figure 39. Left: example of three circles drawn around strings receiving flasher data from string 80. The inner intersection points of the circles are marked with triangles, and the centroid of these points, marked with a black dot, is the fitted flasher position. Right: results of the trilateration fit to String 85, one of the DeepCore strings in the center of the detector. The shifts in the x - and y -coordinates are shown as a function of depth, with the drill data shown as a gray line.

ice to avoid drift due to initial settling of the inclinometers. The electrolytic inclinometer at 2455 m (86% ice sheet depth) was installed at the bottom of String 68. For String 45, the drill descended an additional ~ 100 meters in order to deploy an electrolytic inclinometer attached to a 100-pound weight at 2540 m (90% ice sheet depth) using an extension cable. Data points are long-term average DOM inclination in degrees per year, with error bars indicating the standard deviation from trend. The MEMS sensors have higher noise than the electrolytic inclinometers, and aging tests have indicated a long-term drift of $\sim 0.02\text{--}0.03$ m/s², corresponding to $\sim 0.01\text{--}0.02$ degrees per year [54]. The MEMS readings are consistent with $\lesssim 0.01$ degrees of tilt per year with no apparent depth dependence. The deep electrolytic sensor at 2540 m has shown a persistent 0.07–0.08 degrees of tilt per year since installation (shear of $\tan(0.075^\circ) = 0.0013$ per year).

In most glaciers at sufficiently great depth, ice strain undergoes a transition from compression-dominated to shear-dominated with small-scale folding at centimeter to meter scales [55, 56]. This transition depth depends on temperature [40], grain size, and impurities, and is associated with a strong single-maximum polycrystalline ice fabric [57]. Tilt sensors within the IceCube instrumented volume, at depths < 2450 m, show essentially no movement. Profiles of the atmospheric dust embedded in the ice measured with dust loggers [58] show no indication of folding and appear undisturbed over the full IceCube depth.

5.4 Hole ice imaging camera

An imaging camera system, developed at Stockholm University, Sweden, was deployed on String 80 in the final construction season in order to monitor the freeze-in process and optical properties of the drill hole. The system consists of two video cameras housed in separate glass pressure spheres 5.8 m apart, at a depth of 2455 m between the bottom DOM and the string weights. The cameras

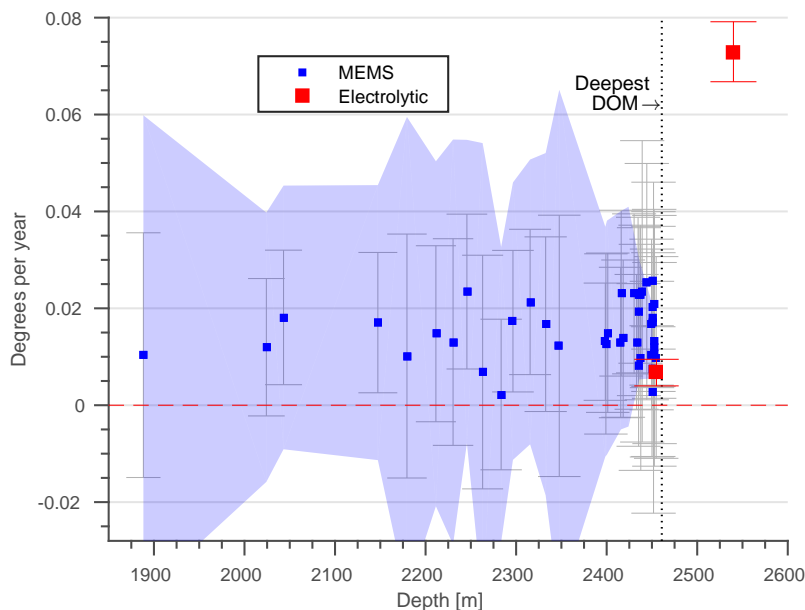


Figure 40. Long term average inclination readings from two electrolytic (red) and 42 MEMS (blue) tilt sensors installed in 2007–2011. Error bars show standard deviation from trend; shaded area is MEMS 95% confidence interval in 15 m bins, not including the quoted drift from aging tests. The reading at 2540 m indicates increasing strain below the IceCube instrumented volume.

can be rotated to point in multiple directions, with some limitations due to mechanical constraints. Each camera is equipped with four LED lamps and three lasers (red, blue, and green). Camera heaters, lights, and movement can be controlled remotely using a dedicated system in the ICL.

The camera system observed that the drill hole was completely refrozen after 15 days. The refrozen hole ice consists of a clear outer layer and a central core of about 16 cm diameter that has a much smaller scattering length than the bulk ice [59]. The optical properties of this hole ice are still under active investigation. No long-term changes have been observed with subsequent camera operations.

6 Online systems

The IceCube online systems comprise both the software and hardware at the detector site responsible for data acquisition, event selection, monitoring, and data storage and movement. As one of the goals of IceCube operations is to maximize the fraction of time the detector is sensitive to neutrino interactions (“uptime”), the online systems are modular so that failures in one particular component do not necessarily prevent the continuation of basic data acquisition. Additionally, all systems are monitored with a combination of custom-designed and industry-standard tools so that detector operators can be alerted in case of abnormal conditions.

6.1 Data flow overview

The online data flow consists of a number of steps of data reduction and selection in the progression from photon detection in the ice to candidate physics event selection, along with associated secondary monitoring and data streams. An overview of the data flow is shown in figure 41.




Figure 41. Data flow in the primary IceCube online systems. See details on each component in the text.

DOM hits are mostly due to dark noise. The first step in data reduction takes place in the DOM, using the Local Coincidence (LC) condition described in section 2.1. Hits that meet the LC criteria are flagged as Hard Local Coincidence (HLC hits) and include a full payload of digitized waveforms, while isolated non-LC hits are flagged as Soft Local Coincidence (SLC) hits and are compressed more aggressively, with only a timestamp and minimal amplitude / charge information transmitted (section 6.3.4).

All DOM hits are read out to dedicated computers on the surface (section 6.2) by the data acquisition system (DAQ). The next level of data selection is the formation of software triggers by the DAQ system. HLC hits across the detector are examined for temporal and in some cases spatial patterns that suggest a common causal relationship. A number of different trigger algorithms run in parallel, described in section 6.4.2. All hits (both HLC and SLC) within a window around the trigger are combined into events, the fundamental output of the DAQ. The event rate varies seasonally with the atmospheric muon flux [24] from 2.5 kHz to 2.9 kHz, with a median rate of 2.7 kHz, and the total DAQ data rate is approximately 1 TB/day.

The DAQ also produces secondary streams that include time calibration, monitoring, and DOM scaler data. The scaler data report the hit rate of each DOM and are used in the supernova data acquisition system (section 6.4.6). The time calibration and monitoring streams are used to monitor the health and quality of the data-taking runs.

The DAQ event data are then processed further with approximately 25 filters in order to select a subset of events (about 15%) to transfer over satellite to the Northern Hemisphere (section 6.5). Each filter, typically designed to select events useful for a particular physics analysis, is run over all events using a computing cluster in the ICL. Because of limitations both on total computing power

Table 7. Breakdown of computing equipment at SPS, indicating number of machines used for each task.

Component	#
DOMHubs	97
Other data acquisition	4
Monitoring	3
Event filtering	24
System infrastructure	8
Other	6

and bounds on the processing time of each event, only fast directional and energy reconstructions are used. This Processing and Filtering (PnF) system is also responsible for applying up-to-date calibration constants to the DAQ data. All processed events, even those not selected by the online filters, are archived locally.

A dedicated system for data movement, JADE, handles the local archival storage to tape or disk and the handoff of satellite data (section 6.6). This includes not only primary data streams but also monitoring data, calibration runs, and other data streams. Low-latency communications for experiment control and real-time monitoring are provided by the IceCube Messaging System (I3MS). Experiment control and detector monitoring are handled by the IceCube Live software system, described in section 6.7.

Data-taking runs are arbitrarily divided into 8-hour periods and assigned a unique run number; data acquisition need not actually pause during the run transition. Detector configuration parameters that affect physics analyses are changed at most once per year (typically in May), indicating the start of a new “physics run”.

6.2 South Pole System and South Pole Test System

The South Pole System (SPS) comprises 19 racks of computing and network hardware that run the various online systems described in this section. The DOM surface cables are connected via passive patch panels to custom 4U computers, called DOMHubs, one DOMHub per in-ice string and 11 additional hubs for IceTop. The remaining servers, including those for higher-level data acquisition, event filtering, detector monitoring, and core infrastructure, are currently 2U Dell PowerEdge R720 servers running Scientific Linux (table 7). The servers are typically upgraded every three to four years. The custom hardware components in the DOMHubs are replaced with spares as failures warrant, and the disks and single-board computer were upgraded in 2013–14.

Each of the DOMHubs is an industrial computer chassis with custom components for DOM power, timing, and communication. A low-power single-board computer communicates with 8 custom PCI DOM Readout (DOR) cards via an industry-standard backplane. An ATX power supply with two redundant modules powers the DOMHub, while two 48 VDC Acopian power supplies, mounted and connected in series inside the chassis to double the voltage to 96 V, supply power to the DOMs. The DOM power is switched and monitored by the DOR cards and is controlled

by software. Another PCI card, the DOMHub Service Board (DSB), is responsible for GPS timing fanout (section 6.3.2).

The SPS computers are connected via switches in each rack that provide redundant connections to a 10–20 Gbps network backbone. The DOMHubs are connected to the rack switches with two bonded 1 Gbps links. Typical network I/O during data-taking for the DAQ Event Builder (section 6.4.3) is about 240 Mbps in each direction. The PnF Central Server sees 200 Mbps in and 640 Mbps out; the output stream is significantly higher than the input as PnF distributes the events to filtering clients and generates multiple output streams (section 6.5).

Redundancy and continuous monitoring of SPS is one of the keys to a high detector livetime (section 6.8). Nagios monitoring software detects and flags problems, including issues with DOM power and communication on the DOMHubs. Severe problems impacting data-taking result in a page to the IceCube winterover personnel via the station’s Land Mobile Radio (LMR) system. A dedicated powered spare server can replace any failed DAQ node, and spare PnF filtering clients can be started to increase throughput in case of a data filtering backlog. SPS hardware is also connected to uninterruptible power supplies (UPS) in order to continue data-taking through station power outages of up to 15 minutes.

Total power usage of the detector and online systems, including computing servers, is approximately 53 kW. The majority is consumed by the DOMs, with an average power consumption of 5.7 W each, including power supply efficiency and transmission losses, for a total of 30.6 kW. The DOMHubs require 128 W each, not including the DOM power, for a total of 12.4 kW. The computing servers consume approximately 200–300 W each, depending on configuration. Most of the of the remaining power is used by the PnF Filter Clients: 20 servers of 300 W each for a total of 6 kW.

A scaled-down version of SPS, the South Pole Test System (SPTS) located in Madison, Wisconsin, U.S.A., allows testing and validation of both hardware and software in the Northern Hemisphere before rollout to SPS. Servers and DOMHubs identical to those at SPS, along with a small number of DOMs in chest freezers, are used in the test system. Although the number of DOMs available is much smaller than in the real detector, recent software improvements allow the “replay” of pre-recorded raw SPS hit data on SPTS systems, providing a data stream to higher-level DAQ and PnF components identical to SPS. Another test system includes a full-length in-ice cable and is used primarily for validation of DOM communications and timing.

6.3 Data readout and timing

While the low-level communications and timing systems of the DOM are described in detail in ref. [27], we review those here in the broader context of the online systems.

6.3.1 Communications

Digital communication between the DOR card and DOM occurs via copper twisted pairs, with two DOMs per pair on the in-ice cable, and one IceTop DOM per pair for increased bandwidth (IceTop hit rates can exceed 3 kHz). The physical layer signaling uses on-off keying with bipolar pulses. The protocol is a custom packet-based scheme. Each packet is assigned a sequence number, and all received packets are acknowledged if the sequence number is correct. Each packet also contains a cyclic redundancy checksum to detect transmission errors. Out-of-sequence packets




Figure 42. Master Clock fanout system, from GPS receiver to DOR cards in each DOMHub.

received are ignored, and non-acknowledged packets are retransmitted. The total bandwidth of the communication channel is 720 kbps per twisted pair.

Messaging is managed from the surface, in that the DOR requests data from each DOM in turn; only one DOM per pair can transmit at a time. Communication is paused once per second to perform a timing calibration (RAPCal; section 3.3); this enables time transfer of DOM clock to DOR clock for every DOM.

6.3.2 Master clock system

The DOR clocks themselves are synchronized to UTC via an active fanout system from a single Symmetricom ET6000 GPS receiver with a temperature-stabilized 10 MHz oscillator, also known as the Master Clock. The fanout tree is shown in figure 42. The 10 MHz output, a 1 Hz output, and a serial time string indicating the UTC date and time are distributed to the DOMHubs via a series of fanouts, using shielded, delay-matched twisted-pair cables. Within the DOMHub, the DSB card continues the fanout via short delay-matched patch cables to each DOR card. The local 20 MHz clocks of each DOR card are phase-locked to the distributed 10 MHz signal.

To avoid the Master Clock being a single-point failure for the detector, a hot spare receiver, using its own GPS antenna and powered through a separate UPS, is continuously active and satellite-locked in case of problems with the primary.

6.3.3 DOR card and driver

Each DOR card is connected to up to 8 DOMs, with 8 DOR cards in a DOMHub. The DOR card controls the 96 VDC power supply to the DOMs and modulates the communications signaling on

top of this DC level; the DOMs can accept input voltages from 40 V to 120 V. Dedicated circuitry monitors the current draw and voltage levels on each twisted pair. A “firmware fuse” can disable power if the current draw deviates from programmable maximum or minimum levels, and this mechanism is supplemented with standard physical fuses.

The software interface to the DOR card, and thus to the DOMs, is provided with a custom Linux device driver. Access to DOR card functions, including DOM power control, communication statistics, RAPCal, and current / voltage monitoring, is facilitated using the Linux `/proc` filesystem interface. Data transfer from the cards to the single-board computer is achieved via DMA over the PCI bus. The driver provides a device file for each DOM for read/write access by higher-level software.

6.3.4 DOM hit payloads

The content of the DOM hit payloads transmitted to the surface depends on whether local coincidence was satisfied, i.e. whether the hit was flagged as HLC or SLC. The DOM Main Board ID and the timestamp of the hit in DOM clock counts are always transmitted, along with trigger and LC flags.

For HLC hits, the digitized ATWD and fADC waveforms are transmitted. Waveforms from lower-gain ATWD channels are only included if the signal amplitude in the higher-gain channel exceeds 75% of the digitizer range. The waveforms are compressed losslessly in the DOM using a delta-compression algorithm that encodes the difference between subsequent samples. The difference values are packed into words of length 1, 2, 3, 6, or 11 bits depending on magnitude, and special values in the bitstream are used to transition between different word lengths.

For both HLC and SLC hits, a chargestamp is included that provides an estimate of the amplitude/charge even if, as in the SLC case, the full waveform is not transmitted. For in-ice DOMs, the chargestamp consists of three samples of the fADC waveform centered around the peak value, along with the peak sample number. For IceTop DOMs, the chargestamp is the sum of all samples of the ATWD waveform, after pedestal subtraction.

6.4 Data acquisition software

IceCube’s data acquisition (DAQ) system is a set of software components running on the DOMHubs and dedicated servers in the ICL. These components are shown in figure 41 and include StringHub, Trigger, Event Builder, Secondary Builder, and a Command and Control server. The DAQ is responsible for detecting patterns of hits in the detector likely to be caused by particle interactions and storing these collections of hits as events.

Hits are read continuously from the DOMs by the StringHub components running on each DOMHub, and a minimal representation of each HLC hit is forwarded to the Trigger components (either the in-ice or IceTop Trigger.) The Trigger components apply a configurable set of algorithms to the hit stream and form windows around interesting temporal and/or spatial patterns. These time windows are collected by the Global Trigger and used to form non-overlapping trigger requests by merging subtriggers as needed, ensuring that the same hit doesn’t appear in multiple events. The merged trigger requests are used by the Event Builder component as templates to gather the complete hit data from each StringHub and assemble the final events.

6.4.1 StringHub and HitSpool

The StringHub software component that runs on each DOMHub is responsible for reading all available data from each of its connected DOMs each second and passing that data onto the downstream consumers. It also saves all hits to a local “HitSpool” on-disk cache and queues them in an in-memory cache to service future requests from the Event Builder for full waveform data.

The StringHub component is divided into two logical pieces: the front end is called Omicron, and the back end is the Sender. Omicron controls all of the connected DOMs, forwarding any non-physics data (calibration, monitoring) to its downstream consumers and sorting the hits from all DOMs into a single time-ordered stream before passing them to the Sender.

Omicron is also responsible for translating DOM hit times into UTC-compatible “DAQ time”, which counts the number of 0.1-ns periods since the UTC start of the year (including leap seconds). The translation uses the RAPCal procedure as described in section 3.3, performed for each DOM every second.

The Sender caches SLC and HLC hits in memory, then forwards a condensed version of each HLC hit to the appropriate local Trigger. Each condensed HLC hit record contains the hit time, a DOM identifier, and the trigger mode. After the Trigger components have determined interesting time intervals, the Event Builder requests each interval from the Sender which returns a list of all hits within the interval and prunes all older hits from the in-memory hit cache after each interval.

One core requirement of the DAQ is that each component operates on a time-ordered stream of data. The DAQ uses its “Splicer” to accomplish this. The Splicer is an object that gathers all input streams during the setup phase at the beginning of a data-taking run; no inputs can be added once started. Each stream pushes new data onto a “tail”, and the Splicer merges the data from all streams into a single sorted output stream. When a stream is closed, it issues an end-of-stream marker that causes the Splicer to ignore all further data. Details of the Splicer algorithm can be found in ref. [60].

As hits move from Omicron to the Sender, they are written to the HitSpool disk cache. These files are written in a circular order so that the newest hits overwrite the oldest data. The files are catalogued in a SQLite database to aid in fast retrieval of raw hit data.

One limitation of the current design is that it only reads data when the full DAQ system is running, so the detector is essentially “off” during certain hardware failures or the periodic full restarts of the system that occur every 32 hours. A future enhancement will split the StringHub into several independent pieces to eliminate these brief pauses. The front end (Omicron) will be moved to a daemon that continuously writes data (including secondary, non-physics data and other metadata) to the disk cache. Part of the back end (Sender) will become a simple HitSpool client that reads data from the disk cache and sends it to the downstream consumers, while another simple component will listen for requested hit readout time intervals from the Event Builder and return lists of hits taken from the HitSpool.

6.4.2 Triggers

The DAQ trigger algorithms look for clusters of HLC hits in space and time that could indicate light due to a particle interaction in the detector, as opposed to uncorrelated dark noise. An algorithm searches for a given multiplicity of HLC hits, possibly with an additional geometric requirement, within a trigger time window. The time scale of the trigger window is set by the light travel time




Figure 43. Schematic representation of triggers using spatial coincidences. Shaded circles represent HLC-hit DOMs. Left: Volume Trigger. Right: String Trigger.

in ice and the geometry requirement involved. Longer readout windows are appended before and after the trigger windows to save early and late hits with the events.

Triggers are generally restricted to a subset of DOMs, such as all in-ice DOMs, IceTop DOMs, or DeepCore DOMs. The algorithms run in parallel over all hits in the DOM set, and then overlapping triggers are merged. The various trigger algorithms are described below, and a summary of the algorithm parameter settings is found in table 8. Trigger settings are changed at most once per year.

The fundamental trigger for IceCube, IceTop, and DeepCore is the Simple Multiplicity Trigger (SMT). The SMT requires N or more HLC hits within a sliding time window of several μs , without any locality conditions. Once the multiplicity condition is met, the trigger is extended until there is a time period of the length of the initial trigger window without any HLC hits from the relevant DOM set. The multiplicity value N is tuned to the energy threshold of the sub-detector, which fundamentally is set by the string or tank spacing.

Other triggers use a lower multiplicity threshold by adding constraints on the HLC hit topology. The time windows for these triggers are based upon the size of the locality volume. The Volume Trigger defines a cylinder of fixed size around each hit DOM and requires a given multiplicity within this cylinder (figure 43a); this allows IceCube to trigger on localized low-energy events that do not satisfy the SMT condition. The Volume Trigger has an additional simple multiplicity parameter that fires the trigger when a certain number of hits is reached, regardless of any spatial restrictions; this prevents the trigger algorithm from slowing down when the detector has triggered already from the primary SMT. The String Trigger requires a certain number of hits within a span of DOMs along a single string (figure 43b); this allows one to trigger on low-energy muons that pass vertically through the detector.

IceCube can detect hypothetical subrelativistic heavy particles such as magnetic monopoles that may catalyze nucleon decays along their trajectory [61]. However, because these particles may travel at velocities less than $0.01c$, the time windows used in the standard triggers are too short. A dedicated Slow Particle (SLOP) trigger has thus been developed to search for slow track-like particle signatures.

The SLOP trigger operates in several stages. The HLC hits, which by design occur at least in pairs along a string, are cleaned by removing pairs that are proximate in time ($\Delta t < T_{\text{prox}}$); T_{prox}




Figure 44. Geometry of a SLOP trigger triplet of HLC pairs.

is tuned to remove most hits from particles traveling near c , such as muons. For all parameters, the trigger algorithm considers the time and position of the first hit within each HLC pair. Next, triplets of HLC pairs within a time window T_{\max} are formed. The geometry of each triplet formed (figure 44) must satisfy track-like conditions: the largest inner angle α of the triangle formed by the HLC pairs must be greater than α_{\min} , and the “velocities” along the triangle sides must be consistent. Specifically, the normalized inverted velocity difference v_{rel} , defined as

$$v_{\text{rel}} = \frac{|\Delta v_{\text{inverse}}|}{\bar{v}_{\text{inverse}}} = 3 \cdot \frac{|\frac{1}{v_{23}} - \frac{1}{v_{12}}|}{\frac{1}{v_{12}} + \frac{1}{v_{23}} + \frac{1}{v_{13}}} \quad (6.1)$$

where $v_{ij} = \Delta x_{ij}/\Delta t_{ij}$, must be less than or equal to a predefined maximum value v_{rel}^{\max} . Finally, the total number of track-like triplets must be greater than or equal to N_{triplet} , set to 5, and all of these track-like triplets must overlap in time.

Other special-purpose triggers exist to collect minimum bias data of various sorts. The Fixed-Rate Trigger (FRT) reads out 10 ms of hit data from the full detector at fixed intervals. This is especially useful for studies of DOM noise. The Calibration Trigger selects a particular type of hit such as special IceTop non-HLC hits that have full waveform readout, and promotes them to a trigger. The Calibration Trigger can also be configured to include all events due to LED flashers in cases where flasher operations require disabling standard triggers. Finally, a Minimum Bias trigger can select one of every N HLC hits and promote this hit to a trigger, adding readout windows as usual; currently an IceTop Minimum Bias trigger with a prescale factor N of 10000 is active.

Many events will satisfy more than one of the trigger conditions, sometimes multiple times. In order to avoid overlapping events, possibly containing the same DOM hits, the triggers and their associated readout windows are merged, while retaining information about the separate triggers. The merged trigger is referred to as the Global Trigger.

Each trigger has defined readout windows around the trigger window; all hits from the full detector, including those DOM sets not involved in the trigger, are requested from the StringHub components and built into events. For the DOM set involved in an in-ice trigger, the readout windows are appended at each end of the trigger window, while for other DOM sets, the readout windows are centered around the trigger start time. Readout windows around IceTop triggers are global and include hits from all other DOM sets before and after the trigger window. The union of overlapping readout windows defines an event (figure 45). Long events such as SLOP or FRT

Table 8. Trigger parameters (as of May 2016) and typical trigger rates of each algorithm. Most rates vary seasonally with the atmospheric muon flux. The merged event rate varies from 2.5 to 2.9 kHz.

Trigger	DOM set	N HLC hits	Window (μs)	Topology	Rate (Hz)
SMT	in-ice	8	5	—	2100
SMT	DeepCore	3	2.5	—	250
SMT	IceTop	6	5	—	25
Volume	in-ice	4	1	cylinder ($r=175\text{m}$, $h=75\text{m}$)	3700
Volume	IceTop infill	4	0.2	cylinder ($r=60\text{m}$, $h=10\text{m}$)	4
String	in-ice	5	1.5	7 adjacent vertical DOMs	2200
SLOP	in-ice	$N_{\text{triplet}} = 5$	$T_{\text{prox}} = 2.5$, $T_{\text{min}} = 0$, $T_{\text{max}} = 500$	$\alpha_{\text{min}} = 140^\circ$, $v_{\text{rel}}^{\text{max}} = 0.5$	12
FRT	all	—	—	—	0.003

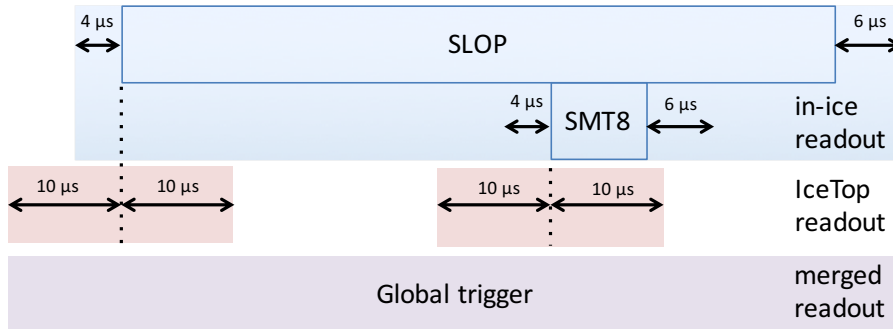


Figure 45. In-ice, IceTop, and merged readout windows for a long event satisfying SLOP and SMT8 triggers.

triggers typically contain several causally independent “physics” events; these typically are re-split before reconstruction and analysis.

6.4.3 Event Builder

The Event Builder receives requests from the Global Trigger, extracts the individual readout windows, and sends them to the appropriate subset of the StringHubs. The StringHubs each send back a list of all hits within the window. When all StringHubs have returned a list of hits, these are bundled with the trigger data into an event.

Events are written to a temporary file. When the temporary file reaches a preset configurable size, it is renamed to a standard unique name. When the PnF system sees a new file, it accepts it for processing and filtering (section 6.5). The total latency from detection of photons at the DOMs to DAQ events written to disk is approximately five seconds.

6.4.4 DAQ and DOM configuration

The configuration of the DAQ is managed by two sets of XML files: a cluster configuration file and a hierarchical tree of run configuration files.

The cluster configuration file contains system-level settings used to launch the DAQ, such as component host servers, startup paths, command-line options, etc. Components (other than StringHub) can easily be moved to different hosts for troubleshooting, load balancing, and maintenance.

Run configuration files list the trigger and DOM configuration files to be used for taking data. The trigger configuration file specifies configuration parameters for all trigger components (in-ice, IceTop, and global) used in a run. These include the list of algorithms run by each trigger component, along with readout window sizes and any other configurable parameters (multiplicity threshold, trigger period, prescale factor, etc.).

DOM configuration files (one per hub) list all DOMs that contribute to the data-taking run. All configuration parameters for each DOM are specified, including PMT high voltages, ATWD operating parameters, discriminator thresholds, local coincidence settings, baselines and others.

Run configuration files (including trigger and DOM files) are versioned and frozen once used for data-taking. All relevant configuration parameters are also stored in a database for use in analysis.

An additional geometry XML file contains the (x, y, z) and (string, position) coordinates of the DOMs, needed by the Trigger components. The DOMs entries are indexed by their unique Main Board ID. This ensures that cabling changes on a DOMHub do not result in changes in data-taking or errors in the geometry.

6.4.5 Component control

The DAQ components are managed by a single “command-and-control” daemon, CnCServer, that manages and monitors components and acts as the main external interface to the DAQ. It uses a standard component interface to query and control the components, and a separate interface for components to expose internal data used for monitoring the health of the detector or for debugging purposes.

CnCServer dynamically discovers the detector components during a launch phase, and instructs them to connect to each other as needed. Using the run configuration files, it then distributes each component configuration appropriately. The components are then started to begin a data-taking run. When a run is in progress, CnCServer regularly checks that components are still active and that data are flowing between components.

6.4.6 Supernova data acquisition system

The IceCube DAQ has a parallel triggering and analysis pathway designed specifically for the detection of the many $O(10)$ MeV neutrinos from a Galactic core-collapse supernova. In the case of such an event, these neutrinos will produce interactions in the detector that, individually, are too dim to trigger the standard DAQ, but because of their high number, can cause a coherent rise in the individual hit rates of the DOMs [28].

Each DOM monitors its hit rate and sends a stream of binned counts, using a bin width of 1.6384 ms (2^{16} clock cycles at 40 MHz). An artificial deadtime of 250 μ s is applied after each hit to

Table 9. HitSpool data-requesting services and request characteristics. The SNDAQ quantity ξ is related to the statistical significance of the rate fluctuation.

Client	Trigger Threshold	Time Window	Request Length	Raw Data Size	Frequency
SNDAQ	$7.65 \leq \xi < 10$	$[-30 \text{ s}, +60 \text{ s}]$	90 s	15 GB	0.5/week
	$\xi \geq 10$	$[\pm 250 \text{ s}]$	500 s	85 GB	0.0006/week
HESE	1500 PE	$[\pm 0.5 \text{ s}]$	1 s	175 MB	4/day
	6000 PE				1/month
Solar Flare	Fermi-LAT significant event	$[\pm 30 \text{ min}]$	1 h	600 GB	7/year

reduce the impact of correlated hits (section 3.8). Each StringHub collects the rate stream of each DOM, supplies UTC-based timestamps, and forwards the streams to the Secondary Builder.

The supernova DAQ (SNDAQ) system receives the Secondary Builder stream, rebins the individual DOM rates, and monitors the sum of rates over several timescales for a significant rise. This analysis is described in detail in ref. [28]. One complication is that light deposition from cosmic-ray muons distorts the significance calculation. To correct for this, the trigger rate of the standard DAQ is continuously sent to SNDAQ, and any significant alert is corrected [62]. At a high significance threshold, the capture of all untriggered data around the alert time is initiated using the HitSpool system (section 6.4.7), and the Supernova Neutrino Early Warning System (SNEWS) [63] is notified. SNDAQ latency of approximately 7 minutes is dominated by the sliding window algorithm used to determine average DOM rates.

6.4.7 HitSpool request system

In the event of a significant transient event, subsystems such as SNDAQ can request all untriggered DOM hits from the detector in a particular time interval by sending requests to a HitSpool Request daemon. Presently, the HitSpool Request System has three clients; their basic characteristics are described in table 9. The central daemon passes the request on to every DOMHub, where hits in the requested time interval are gathered and forwarded to a “sender” component. The hits are then bundled and transferred to the Northern Hemisphere for further analysis.

The time windows of SNDAQ HitSpool data requests are based on the statistical significance of the alert and are shown in table 9. The online High Energy Starting Event (HESE) analysis system requests HitSpool data from a symmetrical time window of 1 s around events with a total deposited charge of greater than 1500 PE. The recently implemented HitSpool client for solar flare analyses is triggered externally by significant Fermi Large Area Telescope (LAT) events [64] and requests HitSpool data from a symmetrical time window of one hour around the trigger time. Unlike the other two clients, these data sets are not transferred over the satellite due to their size but are stored locally on disk, with transmission over satellite only pursued in extraordinary cases.

6.5 Online processing and filtering

6.5.1 Overview

The online Processing and Filtering (PnF) system handles all triggered events collected by the DAQ and reduces the data volume to a level that can be accommodated in our satellite bandwidth allocation (about 100 GB/day). This treatment includes application of calibration constants, event characterization and selection, extraction of data quality monitoring information, generation of realtime alerts for events of astrophysical interest, and creation of data files and metadata information for long-term archiving. The PnF system is a custom software suite that utilizes about 20 standard, multi-processor servers located in the SPS computing cluster.

The first step in the analysis of triggered events is the calibration of the digitized DOM waveforms, as described in section 3.2. The geometry, calibration, and detector status (GCD) information needed to process the data is stored in a database. Next, each DOM's waveform is deconvolved using the known DOM response to single photons to extract the light arrival time and amplitude information [21]. This series of time and amplitude light arrival information for each DOM is the basis for event reconstruction and characterization. PnF encodes this information in a compact data format known as the Super Data Storage and Transfer format (SuperDST); this format uses only 9% of the storage size of the full waveform information. The encoding does introduce a small level of discretization error to the data, measured to be 1.1 ns in time and 0.04 PE in charge, smaller than the calibration uncertainties on these values. Any DOM readout whose SuperDST information is found not to be a good representation of the original waveform, or sees large amounts of light, also has the full waveform data saved in addition to the SuperDST record.

Each event is then characterized with a series of reconstruction algorithms that attempt to match the observed patterns of recorded light in the SuperDST with known patterns of light from track and shower event hypotheses [21]. The reconstructed vertex position, direction, energy, and the goodness-of-fit are used to select interesting events by various filter selections. The filter criteria are set by the collaboration each year and are tuned to select events of interest to specific analyses. Each year there are about 25 filter selections in operation; as of 2016, approximately 15% of all triggered events are selected by one or more filters. Some of these filters are designed to search for neutrino events of wide astrophysical interest to the scientific community and trigger alerts that are distributed to followup observatories worldwide [65, 66].

The PnF system also extracts and aggregates data quality and monitoring information from the data during processing. This information includes stability and quality information from the DOM waveform and calibration process, rates of DOM readouts, and rates and stability information for all detector triggers and filters. This information is aggregated for each data segment and reported to the IceCube Live monitoring system (section 6.7).

Finally, the PnF system generates several types of files for satellite transmission and for long term archival. These include:

1. Filtered data files containing events selected by the online filter selections. These events generally only include the SuperDST version of the DOM information and results from the online event reconstructions. The files are queued for satellite transmission to the IceCube data center in the Northern Hemisphere by the data handling system.

2. SuperDST data files containing the SuperDST version of DOM readout information for all triggered events as well as summary information from the online filtering process. This file set is intended as the long-term archive version of IceCube data.
3. Raw data files containing all uncalibrated waveforms from all DOMs for every event.

During normal operations, the DAQ produces a raw data output of ~ 1 TB per day, resulting in a raw data file archive of the same size. The SuperDST and filtered data archive, after data compression, are ~ 170 GB/day and ~ 90 GB/day, respectively.

6.5.2 System design

The PnF system uses a modular design based on a central master server node and a scalable number of data processing client nodes. The central master server handles data distribution and aggregation tasks (requesting data blocks from the DAQ, collating event SuperDST, reconstruction, filter, and monitoring information and writing data files), while the clients handle the per-event processing tasks (event calibration, reconstruction, analysis, and filtering).

The system is built upon the IceCube analysis software framework, IceTray [67], allowing standard IceCube algorithms to be used online without modifications. The system uses the Common Object Request Broker Architecture (CORBA) system as a means for controlling, supervising and interconnecting the modular portions of the system. Specialized CORBA classes allow data to stream from one component to another using IceTray formats. Clients can also be easily added and removed as needed to meet the processing load.

6.5.3 Components

The flow of triggered event data in the PnF system is shown in figure 46. Standard components include:

1. DAQ Dispatch, a process to pick up event data from the DAQ data cache and forward to the Central Server components.
2. Central Servers, which receive data from the DAQ Dispatch event source, distribute events to and record results returning from the PnF client farm, and send events to Writer components. Typically there are four servers in operation.
3. Filter Clients, where the core calibration, reconstruction and filtering processes are applied to each triggered event. Up to 500 of these clients operate in parallel to filter events in real time.
4. GCD Dispatch, a database caching system to prevent the Filter Client processes from overwhelming the GCD database at run transitions.
5. File Writers, responsible for creation of files and metadata for the data archive. There is one writer component for each file type created.
6. Online Writer, responsible for extracting event reconstruction and filter information from the data for events of astrophysical interest and sending this information out in real time via the IceCube Live alert system.

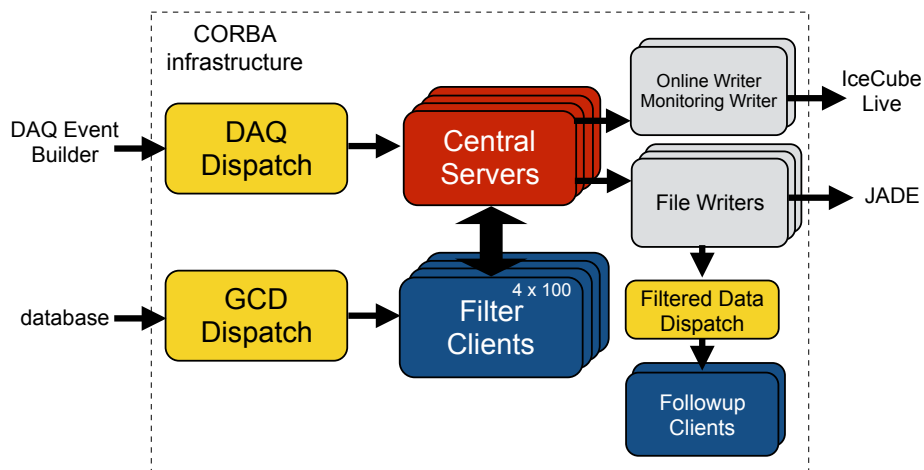


Figure 46. Internal components of the PnF system. Arrows show the flow of data within the system.

7. Monitoring Writer, responsible for aggregating per-event monitoring information, creating histograms, and forwarding to the IceCube Live monitoring system.
8. Filtered Data Dispatch and FollowUp Clients, responsible for looking for bursts of neutrino events on timescales from 100 seconds up to 3 weeks in duration. Any significant burst of neutrinos found generates alerts sent to partner observatories worldwide.

6.5.4 Performance

The PnF system is designed to filter triggered events as quickly as possible after collection by the data acquisition system. A key performance metric is processing system latency, defined as the duration of time between the DAQ trigger and the completion of event processing and filtering. A representative latency history for the system is shown in figure 47, showing typical system latencies of about 20 seconds.

The filter selections used have been relatively stable over several years of operation of the completed IceCube detector, with most seeing only minor updates at the start of each season. The majority of physics analyses derive from a small set of core filters, including:

1. A muon track filter that searches for high-quality track events from all directions. Up-going events for all triggered energies are selected, while only high-energy down-going tracks are selected to avoid the large background of down-going atmospheric muons at lower energies. These selected events are frequently used as the input to point source and transient neutrino searches.
2. A shower event filter that searches for events producing large energy depositions in or near the instrumented volume. These selected events are frequently used as the input to searches for high-energy shower events arising from atmospheric and astrophysical neutrinos.
3. A high-charge filter that searches for any event depositing a large amount of energy leading to a recorded charge of ≥ 1000 photoelectrons in the instrumented volume. While having a large




Figure 47. Typical PnF system latency for a 24-hour period. The latency is defined as the time between DAQ trigger time and time when the online filtering processing is complete. The spikes in latency correspond to DAQ run transitions, when geometry, calibration, and detector status information is updated and distributed to the filtering clients.

overlap with the muon track and shower filters at high energies, this filter targets the highest energy neutrino events of all types. The selected events are used as inputs to searches for high-energy astrophysical and cosmogenic neutrinos as well as for relativistic magnetic monopoles.

4. Cosmic ray filters that search for extended air-shower events in IceTop. The selected events are used as inputs to analyses targeting the flux, spectrum, and composition of the primary cosmic rays observed in the Earth’s atmosphere.
5. A DeepCore contained event filter that searches for contained, lower-energy neutrino events (in the range of 10–100 GeV) from atmospheric neutrino interactions that are contained within the more densely instrumented DeepCore region. The selected events are used as inputs to analyses that measure neutrino oscillation effects and search for indirect signatures of dark matter.

Other filters are employed for more specialized searches, as well as for minimum bias selections.

6.6 Data handling

The bulk of South Pole Station data traffic is handled by geosynchronous satellite links. Due to the location, only geosynchronous satellites with steeply inclined orbits reach far enough above the horizon to establish a link. For a given satellite, this link provides four to six hours of communications once per sidereal day. Multiple satellites are currently utilized by the U.S. Antarctic Program, providing a window of about 12 hours of connectivity with bandwidth of 250 Mbps for uni-directional data transfer and bandwidth of 5 Mbps for bi-directional internet connectivity. For the remainder of the day, Iridium communication satellites allow limited voice and data connectivity and provide up to 2.4 kbps of bandwidth per modem.

IceCube incorporates Iridium modems into two separate systems, the legacy IceCube Teleport System (ITS) and the IceCube Messaging System (I3MS). ITS uses the Iridium Short Burst Data

mode to send short messages of 1.8 kB or smaller with a typical latency (transmission time) of 30 seconds. Messages may either originate or terminate at the ITS Iridium modem at the South Pole. Messages also contain a recipient ID indicating the intended host to receive the message, allowing a many-to-many communications infrastructure between systems running at the South Pole and systems in the Northern Hemisphere. ITS was retired in 2016.

The newer IceCube Messaging System (I3MS), deployed in 2015, incorporates multiple Iridium modems and uses the Iridium RUDICS data mode, providing a 2.4 kbit/s bidirectional serial stream per modem and a minimum latency of about 1.5 seconds. I3MS runs as a daemon on both ends of the link, accepts messages via the ZeroMQ distributed messaging protocol, and transports those messages across the link based on message priority and fair sharing of bandwidth among all users. I3MS message recipients listen for messages using ZeroMQ publish-subscribe (PUB-SUB), allowing a given message to be sent to multiple recipients. I3MS also provides low-bandwidth secure shell (ssh) connectivity to the South Pole, allowing off-site operators access to SPS in the case of detector issues.

Data handling is provided by three servers running the Java Archival and Data Exchange (JADE) software. JADE is a recent Java-based reimplementation and expansion of earlier software, the South Pole Archival and Data Exchange (SPADE). JADE has four primary tasks: data pickup, archiving, satellite transmission, and real-time transmission. The three servers operate independently of one another, and each is capable of separately handling the nominal data volume; thus, data handling can continue seamlessly in case of hardware failure or maintenance.

JADE is configured with a number of input data streams, each consisting of a data server, a dropbox directory, and a filename pattern. The data stream dropbox directories are checked on a regular basis for new files. File completion is indicated by the producer creating a matching semaphore file. For each file, a checksum calculated on the data server is compared to a checksum calculated on the JADE server. After validation, the original data file is removed from the pickup location.

Files are then routed according to the configuration of their data stream, either transmitted via satellite link or archived locally. Archival data were formerly written to Linear Tape Open (LTO) tapes; the tape system was retired in 2015, and archival data are now written to disks. All of the archival data are buffered on the server until the storage medium is full. In case of media failure, the buffered files can be immediately written to new archival media with a single command. Two copies of each archival data stream are saved, and the disks are transported from the South Pole to the IceCube data center each year. Archival data are not regularly reprocessed but are retained indefinitely in case of an issue with the filtered data streams or a significant improvement in the low-level calibration.

Data streams intended for satellite transfer are queued separately. JADE combines multiple smaller files or splits large files to create ~ 1 GB bundles, allowing satellite link operators to manage the daily data transmission. A configurable number of bundles is then transferred to the satellite relay server. If satellite transmission is temporarily interrupted, the excess bundles are staged on the JADE server.

Small files (<50 KB) with high priority are sent via the I3MS Iridium link. In cases where the real-time link is not available, I3MS will queue the messages to be sent when the link becomes available. All I3MS messages are also sent to JADE to send via the geosynchronous satellite link to ensure delivery if the Iridium link should be unavailable for an extended period of time.

6.7 IceCube live and remote monitoring

IceCube operations are controlled and monitored centrally by IceCube Live. IceCube Live consists of two major components: LiveControl, responsible for controlling data-taking operations and collecting monitoring data, and the IceCube Live website, responsible for processing and storing monitoring data as well as presenting this data in webpages and plots that characterize the state of the detector.

6.7.1 LiveControl

LiveControl is responsible for controlling the state of DAQ and PnF, starting and stopping data-taking runs, and recording the parameters of these runs. Human operators typically control the detector and check basic detector status using a command-line interface to the LiveControl daemon. Standard operation is to request a run start, supplying a DAQ run configuration file. LiveControl then records the run number, configuration, start time, etc. and sends a request for DAQ to begin data-taking. After data-taking commences successfully, LiveControl waits a specified amount of time, generally eight hours, then stops the current run and automatically starts a new run using the same configuration.

This cycle continues until stopped by a user request or a run fails. In case of failure, LiveControl attempts to restart data-taking by starting a new run. Occasionally a hardware failure occurs, and a new run cannot be started with the supplied configuration because requested DOMs are unpowered or temporarily unable to communicate. In this case, LiveControl cycles through predefined partial-detector configurations in an attempt to exclude problematic DOMs. This results in taking data with fewer than the full number of available strings, but it greatly reduces the chance of a prolonged complete outage where no IceCube data are recorded.

A secondary function of LiveControl is the collection, processing, and forwarding of monitoring data from DAQ, PnF, and other components. The associated JavaScript Object Notation (JSON) data are forwarded to LiveControl using the ZeroMQ protocol and queued internally for processing. Certain monitoring quantities indicate serious problems with the detector, e.g. the PnF latency is too high. LiveControl maintains a database of critical monitoring quantities and raises an alert if the value is out of the specified range or hasn't been received in a specified amount of time. The alert usually includes an email to parties responsible for the affected subsystem and, for serious problems, triggers an automated page to winterover operators. Alerts generated by controlled components such as DAQ or PnF may also generate emails or pages. All monitoring data are forwarded to the IceCube Live website for further processing and display.

6.7.2 IceCube live website

Two operational copies of the IceCube Live website exist: one inside the IceCube network at the South Pole, and one in the Northern Hemisphere. Monitoring data reach the northern website based on relative priority and using both geosynchronous and Iridium data transport, summarized in table 10.

Messages reaching the website are processed by a database server daemon (DBServer). Messages also may contain directives requesting DBServer to send email, by specifying email recipients and content, or requesting that the monitoring message be published using the ZeroMQ PUB-SUB framework, allowing the message to be passed to an external process. The IceCube Live website

Table 10. Typical data volume and latencies for IceCube Live monitoring messages.

Priority	Transport System	Messages/day	Data/day	Latency
1	I3MS (Iridium)	10,000	1 MB	1 min.
2	I3MS (Iridium)	150,000	5 MB	1–5 min.
3	JADE (Geosynchronous)	300,000	100 MB	1 day

itself uses the Django web framework and contains pages that create graphical views of monitoring data stored in the database. These pages include a front page displaying active alerts and plots of event rates and processing latencies from the previous few hours (figure 48), and a page for each run that displays start time, stop time, and other essential data. The run page contains low-level diagnostic data that include e.g. DOM charge histograms, digitizer baselines, DOM occupancy, etc., and are used to diagnose any problems that occurred during the run and to determine if the run can be used in physics analysis. Historical monitoring data are available from runs back to 2011.

Finally, the IceCube Live website in the Northern Hemisphere can transmit messages to LiveControl at SPS using the Iridium systems. This capability is used to link the commercial Slack chat service to a chat web page in IceCube Live, allowing the IceCube winterover operators to communicate with experts in the Northern Hemisphere during periods with no geosynchronous satellite connectivity. This connection also provides a limited capability to control the detector, allowing operators in the Northern Hemisphere to start/stop components or remotely issue HitSpool requests.

6.8 Operational performance

Detector operational uptime is highly valued, in order to remain sensitive to rare astrophysical transient events. Many redundancies and failsafes have been implemented, allowing an average detector uptime of greater than 99%. The detector uptime measures the fraction of total time that some portion of the detector is taking data; sources of downtime include the transitions between data-taking runs, power outages, and DAQ interruptions due to software or hardware issues. All DOMHubs and servers are equipped with redundant power supplies; these are in turn connected to redundant UPSes. This backup battery power allows the detector to remain fully operational for about 15 minutes in the event of a power outage.

Industry-standard Nagios monitoring software tracks the SPS hardware status and provides the link between LiveControl and the paging system. In the event of a hardware or software failure that interrupts data-taking, a partial-detector configuration can often be started automatically in parallel with the paging alert. Within minutes, winterover operators start an optimal detector configuration excluding only affected string(s), and then proceed to diagnose and fix the issue. Software issues can often be addressed via network access from the South Pole Station, while hardware failures may require a visit to the ICL for repairs.

The recovery of data from all but the last few minutes of runs that fail, and the recent implementation of continuous data-taking, have improved detector stability and decreased detector downtime. These features contribute to an average “clean uptime” in recent years of 97–98% of full-detector, analysis-ready data, exceeding our target of 95%. Historical total detector uptime and clean uptime are shown in figure 49.

Detector Status

Data Acquisition

Current run: **128214** (5h:00m:07s)
 Run config: sps-IC86-2016-infill-iceact-V259
 DAQ release: New_Glarus5_14222:143282
 Total events: 46107357
 Active DOMs: 5408
 Light mode: dark

Control Details

pdaq **RUNNING**

Other Components


DB	RUNNING	stop
I3DAQDispatch	RUNNING	stop
I3MoniDomMon	RUNNING	stop
I3MoniDomSn	RUNNING	stop
I3MoniDomTcal	RUNNING	stop
I3MoniMover	RUNNING	stop
I3MoniPhysA	RUNNING	stop
PFDSTWriter	RUNNING	stop
PFFiltWriter	RUNNING	stop
PFMoniWriter	RUNNING	stop
PFOneWriter	RUNNING	stop
PFRawWriter	RUNNING	stop
PFSer1	RUNNING	stop
PFSer2	RUNNING	stop
PFSer3	RUNNING	stop
PFSer4	RUNNING	stop
sndaq	RUNNING	stop
TFRateMonitor	RUNNING	stop
uptimer	RUNNING	stop

Page WOs


Currently Watched Alerts

multirunfail	OK
runfail	OK
maxdt_sndaq_latency=900	OK
throttledNoisyService	OK
ICL overtemp max2	OK
/mnt/data/pdaqlo...n.tar file count	OK
DB in RUNNING state	OK
Detector not taking data	OK
Excessive run failure rate	OK
FPMaster Disk Usage Alarm	OK
I3DAQDispatch in RUNNING state	OK
ICL overtemp max1	OK
ICL overtemp min2	OK
ICL temperature too high	OK
Lots of LBM overflows	OK
Max SDAQ Latency	OK
Max time between SDAQ alerts	OK
Max WXGoose 3 Temp	OK
Max WXGoose 3 Temp (pages)	OK
Max WXGoose 6 Temp	OK
Maximum Missing DOMs	OK
Min WXGoose 1 Temp	OK
Min WXGoose 6 Temp	OK
PFDSTWriter in RUNNING state	OK
PFFiltWriter in RUNNING state	OK
PFRawWriter in RUNNING state	OK
PFSer1 in RUNNING state	OK
PFSer2 in RUNNING state	OK
PFSer3 in RUNNING state	OK
PFSer4 in RUNNING state	OK
PnF latency too high	OK
PnF rate too low	OK
SERIOUS SN alert triggered!	OK
Supernova DAQ state check	OK
Test Alert	OK
TFRateMonitor in RUNNING state	OK


pDAQ Event Rate



Active DOMs vs Expected DOMs




PnF Latency



SDAQ Processing Latency



LBM Overflows



2017 JINST 12 P03012

Figure 48. The IceCube Live website front page, showing near-realtime detector status. One DOM has dropped out of data-taking in the run shown. This “operator” view includes controls to start/stop components and to page the winterovers.




Figure 49. Detector uptime each month since the start of full-detector operation. “Clean uptime” indicates periods of analysis-ready, 86-string data.

About 0.2% of the loss in clean uptime is due to the failed portions of runs that are not usable for analysis. There is around 1% of clean uptime loss due to runs not using the full-detector configuration. This occurs when certain components of the detector are excluded from the run configuration during required repairs and maintenance; these partial-detector runs are still useful for analyses that have less strict requirements on the active detector volume or in the case of an exceptional astrophysical transient event. There is approximately a 1% loss of clean uptime due to maintenance, commissioning, and verification runs, and short runs that are less than 10 minutes in duration. A breakdown of total detector time by run period is shown in figure 50.

7 Outlook

IceCube’s online systems continually evolve in order to enhance the capabilities and performance of the detector and meet new science needs. For example, the 2016 physics run start included a new low-energy IceTop trigger, a new filter designed to search for moderate-velocity magnetic monopoles, and a transition to a more flexible realtime alert system that transmits candidate neutrino events at a rate of 3 mHz via I3MS to a followup analysis server in the Northern Hemisphere. Specialized filter streams for rare events likely to be astrophysical in origin have also been added within the past two years.

In addition to online software improvements and regular computing upgrades, a number of calibration and maintenance hardware activities are planned. For example, in order to mitigate the effects of continued snow accumulation on the IceTop tanks — increased energy threshold and reduced sensitivity to the electromagnetic air shower component — a small number of scintillator-based detector panels have been installed at existing IceTop stations, and development of additional panels is underway.

The stability of the primary triggers and filters and a high detector uptime enable continued detection of additional astrophysical neutrino candidates. Searches for neutrino point sources are

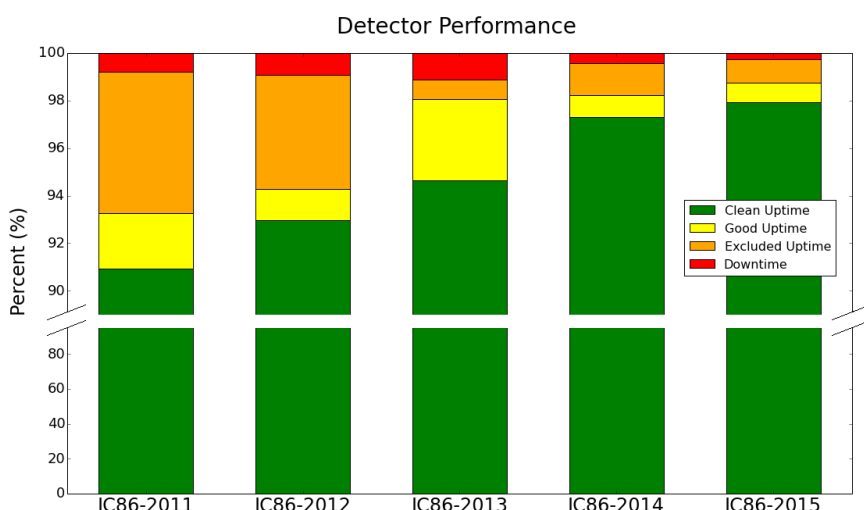


Figure 50. Detector performance breakdown by run period. Clean uptime indicates pristine, full-detector data; “good uptime” includes partial-detector runs that are otherwise healthy; and “excluded uptime” indicates runs that are declared unusable for analysis.

ongoing, and a robust realtime followup campaign facilitates multi-wavelength observations of reconstructed neutrino source directions, with the eventual goal of definitive identification of the astrophysical hadron accelerators.

7.1 IceCube Gen2 and PINGU

As neutrino astronomy matures, designs for next-generation detectors are underway. Neutrino detectors under construction or planned for the Northern Hemisphere include KM3NeT in the Mediterranean [68] and GVD in Lake Baikal [69].

IceCube Gen2 is an experiment under design for the South Pole consisting of a high-energy in-ice array, a surface air shower array with additional veto capabilities, a low-energy in-ice infill array (the Precision IceCube Next Generation Upgrade, or PINGU), and potentially a shallow sub-surface array of radio antennas [70]. The high-energy array features an increased string spacing that allows instrumentation of up to 10 times larger effective target volume relative to IceCube; the scientific focus will be the detection and characterization of astrophysical neutrino sources at the PeV energy scale. The PINGU sub-array [71] will densely instrument 6 Mton of ice in the center of DeepCore, enabling precision neutrino oscillation measurements down to the GeV energy range, determination of the neutrino mass ordering, and dark matter searches at energies above a few GeV. Updated calibration devices will be deployed in the new boreholes in order to measure the optical properties of the ice more precisely and improve event reconstruction. An extended surface array, potentially several times larger in diameter than the high-energy array, will be used both for cosmic ray studies and to veto downgoing atmospheric muons and neutrinos. Radio-frequency detection of ultra-high-energy neutrino-induced showers in the ice is a relatively recent technique which shows considerable promise to achieve effective target volumes of about 100 times IceCube at 10^{18} eV where neutrinos originating from scattering of ultra-high-energy cosmic rays on the cosmic microwave background are expected [72]. The combination of optical and radio-frequency

technologies offers the possibility of a broadband neutrino observatory with diverse astrophysics and particle physics science capabilities.

The baseline sensor design for the array is a modernized version of the IceCube DOM [71]. Mechanical components such as the glass sphere and penetrator, as well as the high-quantum-efficiency PMT, remain unchanged, while the triggering and digitization electronics are being redesigned. Alternative sensor designs are also under study that increase photocathode area, photon collection, angular coverage, and/or directional resolution. Slimmer, more cylindrical profiles may allow smaller hole diameters, decreasing drilling time and fuel usage. The recent development of the capability to deliver cargo and fuel to the station via overland traverse rather than aircraft will reduce fuel costs. This is part of an overall effort to reduce logistical support requirements compared to IceCube construction.

The IceCube detector has achieved “first light” in neutrino astronomy and has the capability to continue operating at least until the end of the next decade, supporting a diverse neutrino physics and astrophysics program and providing unique datasets to the scientific community. The IceCube Gen2 facility will continue this legacy and contribute to further discoveries in neutrino astronomy and multi-messenger astrophysics.

Acknowledgments

We acknowledge the support from the following agencies: U.S. National Science Foundation — Office of Polar Programs, U.S. National Science Foundation — Physics Division, University of Wisconsin Alumni Research Foundation, the Grid Laboratory Of Wisconsin (GLOW) grid infrastructure at the University of Wisconsin — Madison, the Open Science Grid (OSG) grid infrastructure; U.S. Department of Energy, and National Energy Research Scientific Computing Center, the Louisiana Optical Network Initiative (LONI) grid computing resources; Natural Sciences and Engineering Research Council of Canada, WestGrid and Compute/Calcul Canada; Swedish Research Council, Swedish Polar Research Secretariat, Swedish National Infrastructure for Computing (SNIC), and Knut and Alice Wallenberg Foundation, Sweden; German Ministry for Education and Research (BMBF), Deutsche Forschungsgemeinschaft (DFG), Helmholtz Alliance for Astroparticle Physics (HAP), Research Department of Plasmas with Complex Interactions (Bochum), Germany; Fund for Scientific Research (FNRS-FWO), FWO Odysseus programme, Flanders Institute to encourage scientific and technological research in industry (IWT), Belgian Federal Science Policy Office (Belspo); University of Oxford, United Kingdom; Marsden Fund, New Zealand; Australian Research Council; Japan Society for Promotion of Science (JSPS); the Swiss National Science Foundation (SNSF), Switzerland; National Research Foundation of Korea (NRF); Villum Fonden, Danish National Research Foundation (DNRF), Denmark.

IceCube acronym list

ADC	analog-to-digital converter
ATWD	Analog Transient Waveform Digitizer; front-end signal digitizers in the DOM
cDOM	color (LED) DOM
CORBA	Common Object Request Broker Architecture; software process communications middleware used in PnF
DAQ	data acquisition
DESY	Deutsches Elektronen-Synchrotron national research center, Germany
DFL	dark freezer lab used for DOM testing
DMA	direct memory access
DOM	digital optical module; IceCube’s primary light sensor
DOMCal	DOM calibration software
DOMHub	DOM readout computer in the ICL
DOR	DOM readout card used to power and communicate with the DOMs
DSB	DOMHub service board providing GPS clock signals to the DOR cards
EHWD	Enhanced Hot Water Drill used to drill IceCube holes
ESD	electrostatic discharge
fADC	fast ADC; front-end signal digitizer in the DOM
FAT	final acceptance testing of the DOMs
FHWM	full width half maximum
FPGA	field-programmable gate array
FRT	fixed-rate trigger; a software trigger algorithm used in the DAQ
GCD	geometry, calibration, and detector status information needed to process IceCube data
GPS	Global Positioning System; primary time source for IceCube
HESE	high-energy starting event; a bright neutrino event candidate first visible inside the detector
HLC	hard local coincidence, as in “HLC hits”; DOM hits satisfying this condition are read out with full waveform information
HQE	high quantum efficiency (PMT)
I3MS	IceCube Messaging System; enables 24/7 connectivity to the South Pole
ICL	IceCube Laboratory; central building for IceCube systems at the South Pole
ITS	IceCube Teleport System; legacy communications system used to provide 24/7 connectivity to the South Pole
JADE	Java Archival and Data Exchange; data-handling software system
JSON	JavaScript Object Notation; a lightweight data-interchange format
LBM	lookback memory, the SDRAM hit buffer on each DOM

LC	Local coincidence of DOMs on a string or between tanks in an IceTop station; see also HLC and SLC
LED	light-emitting diode
LTO	linear tape open storage media format
MPE	multiple photoelectron, as in “MPE discriminator”
PCI	Peripheral Component Interconnect; a computer bus
PE	photoelectron; also a unit of amplified charge resulting from one converted photon incident on a PMT
PINGU	Precision IceCube Next Generation Upgrade
PMT	photomultiplier tube
PnF	processing and filtering online software system
PSL	Physical Sciences Laboratory, Stoughton, Wisconsin, U.S.A
RAPCal	reciprocal active pulsing calibration; time calibration scheme used to synchronize the DOMs
RFI	radio-frequency interference
SJB	surface junction box; passive connection of the in-ice cable, surface cable, and IceTop tanks
SES	Seasonal Equipment Site used to provide power and hot water for drilling operations
SLC	soft local coincidence, as in “SLC hits”; SLC hits do not satisfy an LC condition and are read out with a timestamp and minimal amplitude / charge information
SLOP	slow particle, as in “SLOP trigger”; a software trigger algorithm used in the DAQ
SMT	simple multiplicity trigger; a software trigger algorithm used in the DAQ
SNDAQ	supernova data acquisition software
SNEWS	Supernova Neutrino Early Warning System
SPADE	South Pole Archival and Data Exchange; legacy data-handling software system
SPE	single photoelectron, as in “SPE discriminator”
SPS	South Pole System, the IceCube computing and networking systems at the detector site
SPTS	South Pole Test System, a scaled-down version of SPS for testing in the Northern Hemisphere
SuperDST	Super Data Storage and Transfer compressed event format
TDP	track detection probability
TOS	Tower Operations Site used at each hole for drilling and string deployment
UPS	uninterruptible power supply
UTC	Coordinated Universal Time, the time standard used in IceCube

References

- [1] F. Reines, C.L. Cowan, F.B. Harrison, A.D. McGuire and H.W. Kruse, *Detection of the free anti-neutrino*, *Phys. Rev.* **117** (1960) 159.
- [2] B.T. Cleveland et al., *Measurement of the solar electron neutrino flux with the Homestake chlorine detector*, *Astrophys. J.* **496** (1998) 505.
- [3] C.V. Achar et al., *Detection of muons produced by cosmic ray neutrinos deep underground*, *Phys. Lett.* **18** (1965) 196.
- [4] F. Reines et al., *Evidence for high-energy cosmic ray neutrino interactions*, *Phys. Rev. Lett.* **15** (1965) 429.
- [5] SUPER-KAMIOKANDE collaboration, Y. Ashie et al., *Evidence for an oscillatory signature in atmospheric neutrino oscillation*, *Phys. Rev. Lett.* **93** (2004) 101801 [[hep-ex/0404034](#)].
- [6] SNO collaboration, Q.R. Ahmad et al., *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, *Phys. Rev. Lett.* **89** (2002) 011301 [[nucl-ex/0204008](#)].
- [7] KAMIOKANDE-II collaboration, K. Hirata et al., *Observation of a neutrino burst from the supernova SN1987a*, *Phys. Rev. Lett.* **58** (1987) 1490.
- [8] R.M. Bionta et al., *Observation of a neutrino burst in coincidence with supernova SN1987a in the Large Magellanic Cloud*, *Phys. Rev. Lett.* **58** (1987) 1494.
- [9] A.I. Afonin et al., *$\bar{\nu}_e$ spectra at two distances from the reactor of the Rovno nuclear power plant: search for oscillations*, *JETP Lett.* **45** (1987) 247 [*Pisma Zh. Eksp. Teor. Fiz.* **45** (1987) 201].
- [10] DUMAND collaboration, J. Babson et al., *Cosmic ray muons in the deep ocean*, *Phys. Rev. D* **42** (1990) 3613.
- [11] BAIKAL collaboration, I.A. Belolaptikov et al., *The Baikal underwater neutrino telescope: design, performance and first results*, *Astropart. Phys.* **7** (1997) 263.
- [12] E. Andres et al., *The AMANDA neutrino telescope: principle of operation and first results*, *Astropart. Phys.* **13** (2000) 1 [[astro-ph/9906203](#)].
- [13] ANTARES collaboration, M. Ageron et al., *ANTARES: the first undersea neutrino telescope*, *Nucl. Instrum. Meth. A* **656** (2011) 11 [[arXiv:1104.1607](#)].
- [14] F. Halzen and D. Hooper, *High-energy neutrino astronomy: the cosmic ray connection*, *Rept. Prog. Phys.* **65** (2002) 1025 [[astro-ph/0204527](#)].
- [15] J.G. Learned and K. Mannheim, *High-energy neutrino astrophysics*, *Ann. Rev. Nucl. Part. Sci.* **50** (2000) 679.
- [16] T.K. Gaisser, F. Halzen and T. Stanev, *Particle astrophysics with high-energy neutrinos*, *Phys. Rept.* **258** (1995) 173 [*Erratum ibid.* **271** (1996) 355] [[hep-ph/9410384](#)].
- [17] ICECUBE collaboration, M.G. Aartsen et al., *Evidence for high-energy extraterrestrial neutrinos at the IceCube detector*, *Science* **342** (2013) 947.
- [18] F. Halzen and S.R. Klein, *IceCube: an instrument for neutrino astronomy*, *Rev. Sci. Instrum.* **81** (2010) 081101 [[arXiv:1007.1247](#)].
- [19] ICECUBE collaboration, M.G. Aartsen et al., *Observation of the cosmic-ray shadow of the moon with IceCube*, *Phys. Rev. D* **89** (2014) 102004 [[arXiv:1305.6811](#)].

- [20] ICECUBE collaboration, F. Bos et al., *Observation of the cosmic-ray shadow of the moon and sun with IceCube*, *ASTRA Proc.* **2** (2015) 5.
- [21] ICECUBE collaboration, M.G. Aartsen et al., *Energy reconstruction methods in the IceCube neutrino telescope*, *2014 JINST* **9** P03009 [[arXiv:1311.4767](#)].
- [22] ICECUBE collaboration, R. Abbasi et al., *The design and performance of IceCube DeepCore*, *Astropart. Phys.* **35** (2012) 615 [[arXiv:1109.6096](#)].
- [23] ICECUBE collaboration, M.G. Aartsen et al., *Measurement of South Pole ice transparency with the IceCube LED calibration system*, *Nucl. Instrum. Meth. A* **711** (2013) 73 [[arXiv:1301.5361](#)].
- [24] ICECUBE collaboration, R. Abbasi et al., *IceTop: the surface component of IceCube*, *Nucl. Instrum. Meth. A* **700** (2013) 188 [[arXiv:1207.6326](#)].
- [25] ICECUBE collaboration, M.G. Aartsen et al., *Measurement of the cosmic ray energy spectrum with IceTop-73*, *Phys. Rev. D* **88** (2013) 042004 [[arXiv:1307.3795](#)].
- [26] ICECUBE collaboration, R. Abbasi et al., *Calibration and characterization of the IceCube photomultiplier tube*, *Nucl. Instrum. Meth. A* **618** (2010) 139 [[arXiv:1002.2442](#)].
- [27] ICECUBE collaboration, R. Abbasi et al., *The IceCube data acquisition system: signal capture, digitization and timestamping*, *Nucl. Instrum. Meth. A* **601** (2009) 294 [[arXiv:0810.4930](#)].
- [28] ICECUBE collaboration, R. Abbasi et al., *IceCube sensitivity for low-energy neutrinos from nearby supernovae*, *Astron. Astrophys.* **535** (2011) A109 [*Erratum ibid.* **563** (2014) C1] [[arXiv:1108.0171](#)].
- [29] E. Calvo et al., *Characterization of large area photomultipliers under low magnetic fields: design and performances of the magnetic shielding for the Double CHOOZ neutrino experiment*, *Nucl. Instrum. Meth. A* **621** (2010) 222 [[arXiv:0905.3246](#)].
- [30] S. Kleinfelder, *Advanced transient waveform digitizers*, *Proc. SPIE* **4858** (2003) 316.
- [31] ICECUBE collaboration, J. Kiryluk et al., *IceCube performance with artificial light sources: the road to cascade analyses*, in *The IceCube collaboration: contributions to the 30th International Cosmic Ray Conference (ICRC 2007)*, Merida Yucatan Mexico, (2007), pg. 1233 [[arXiv:0711.0353](#)].
- [32] E.B. Wilson, *Probable inference, the law of succession, and statistical inference*, *J. Amer. Statist. Assoc.* **22** (1927) 209.
- [33] ICECUBE collaboration, M.G. Aartsen et al., *The IceCube neutrino observatory part VI: ice properties, reconstruction and future developments*, in *Proceedings of the 33rd International Cosmic Ray Conference*, Rio de Janeiro Brazil, (2013) [[arXiv:1309.7010](#)].
- [34] P.B. Price and K. Woschnagg, *Role of group and phase velocity in high-energy neutrino observatories*, *Astropart. Phys.* **15** (2001) 97 [[hep-ex/0008001](#)].
- [35] ICECUBE collaboration, A. Achterberg et al., *First year performance of the IceCube neutrino telescope*, *Astropart. Phys.* **26** (2006) 155 [[astro-ph/0604450](#)].
- [36] ICECUBE collaboration, D. Tosi and C. Wendt, *Calibrating photon detection efficiency in IceCube*, *PoS(TIPP2014)* 157 [[arXiv:1502.03102](#)].
- [37] E.M. Lau et al., *Statistical characterization of the meteor trail distribution at the South Pole as seen by a VHF interferometric meteor radar*, *Radio Sci.* **41** (2006) RS4007.
- [38] R.A. Makarevich, V. Forsythe and A. Kellerman, *Electric field control of E region coherent echoes: evidence from radar observations at the South Pole*, *J. Geophys. Res. Space* **120** (2015) 2148.

- [39] AMANDA collaboration, J. Ahrens et al., *Muon track reconstruction and data selection techniques in AMANDA*, *Nucl. Instrum. Meth. A* **524** (2004) 169 [[astro-ph/0407044](#)].
- [40] P.B. Price et al., *Temperature profile for glacial ice at the South Pole: implications for life in a nearby subglacial lake*, *Proc. Nat. Acad. Sci.* **99** (2002) 7844.
- [41] H.O. Meyer, *Spontaneous electron emission from a cold surface*, *Europhys. Lett.* **89** (2010) 58001.
- [42] AMANDA collaboration, K. Helbing et al., *Light emission in AMANDA pressure spheres*, Tech. Rep. [AMANDA-IR/20030701](#), (2003).
- [43] M. Larson, *Simulation and identification of non-poissonian noise triggers in the IceCube neutrino detector*, *Master's thesis*, University of Alabama, Tuscaloosa AL U.S.A., (2013).
- [44] N. Stanisha, *Characterization of Low-dt Non-poisson Noise in the Icecube Neutrino Detector*, *Bachelor's thesis*, Pennsylvania State University, State College PA U.S.A., (2014).
- [45] D. Heereman von Zuydtwyck, *HitSpooling: an improvement for the supernova neutrino detection system in IceCube*, *Ph.D. thesis*, Université libre de Bruxelles, Bruxelles Belgium, (2015).
- [46] B. Riedel, *Modeling and understanding supernova signals in the IceCube neutrino observatory*, *Ph.D. thesis*, University of Wisconsin, Madison WI U.S.A., (2014).
- [47] Y. Mizuno and T. Mizuno, *Photon emission accompanying deformation and fracture of ice*, *Can. J. Phys.* **81** (2003) 71.
- [48] A. Piegsa, *Supernova-Detektion mit dem IceCube-Neutrino teleskop* (in German), *Ph.D. thesis*, Johannes Gutenberg-Universität, Mainz Germany, (2009).
- [49] AMANDA collaboration, M. Ackermann et al., *The IceCube prototype string in AMANDA*, *Nucl. Instrum. Meth. A* **556** (2006) 169 [[astro-ph/0601397](#)].
- [50] T. Benson et al., *IceCube enhanced hot water drill functional description*, *Ann. Glaciol.* **55** (2014) 105.
- [51] R. Schmitt and R. Rodriguez, *Glacier water supply and sewage disposal systems*, in *Symposium on Antarctic Logistics, 13–17 August 1962, Boulder CO U.S.A.*, National Academy of Sciences — National Research Council, Washington DC U.S.A., (1963), pg. 329.
- [52] N.E. Bramall et al., *A deep high-resolution optical log of dust, ash, and stratigraphy in South Pole glacial ice*, *Geophys. Res. Lett.* **32** (2005) L21815.
- [53] F. Pattyn, *A new three-dimensional higher-order thermomechanical ice sheet model: basic sensitivity, ice stream development, and ice flow across subglacial lakes*, *J. Geophys. Res.* **108** (2003) 2382.
- [54] Analog Devices Inc., personal communication, 27–28 August 2013.
- [55] M. Montagnat et al., *Fabric along the NEEM ice core, Greenland, and its comparison with GRIP and NGRIP ice cores*, *Cryosph.* **8** (2014) 1129.
- [56] D. Jansen et al., *Small-scale disturbances in the stratigraphy of the NEEM ice core: observations and numerical model simulations*, *Cryosph.* **10** (2016) 359.
- [57] K. Cuffey and W. Paterson, *The physics of glaciers*, 4th edition, Butterworth-Heinemann/Elsevier, Burlington MA U.S.A., (2010).
- [58] ICECUBE collaboration, M.G. Aartsen et al., *South Pole glacial climate reconstruction from multi-borehole laser particulate stratigraphy*, *J. Glaciol.* **59** (2013) 1117.
- [59] ICECUBE collaboration, M. Rongen et al., *Measuring the optical properties of IceCube drill holes*, *EPJ Web Conf.* **116** (2016) 06011.

- [60] ICECUBE collaboration, J.L. Kelley, *Event triggering in the IceCube data acquisition system*, *AIP Conf. Proc.* **1630** (2014) 154.
- [61] ICECUBE collaboration, M.G. Aartsen et al., *Search for non-relativistic magnetic monopoles with IceCube*, *Eur. Phys. J. C* **74** (2014) 2938 [[arXiv:1402.3460](#)].
- [62] ICECUBE collaboration, M.G. Aartsen et al., *Recent improvements in the detection of supernovae with the IceCube observatory*, in *Proceedings of Science: 34th International Cosmic Ray Conference*, The Hague Netherlands, (2015) [[PoS\(ICRC2015\)1096](#)].
- [63] P. Antonioli et al., *SNEWS: the SuperNova Early Warning System*, *New J. Phys.* **6** (2004) 114.
- [64] FERMI-LAT collaboration, M. Pesce-Rollins et al., *Fermi Large Area Telescope observations of high-energy gamma-ray emission from behind-the-limb solar flares*, in *Proceedings of Science: 34th International Cosmic Ray Conference*, The Hague Netherlands, (2015) [[PoS\(ICRC2015\)128](#)] [[arXiv:1507.04303](#)].
- [65] ICECUBE and ROTSE collaborations, R. Abbasi et al., *Searching for soft relativistic jets in core-collapse supernovae with the IceCube optical follow-up program*, *Astron. Astrophys.* **539** (2012) A60 [[arXiv:1111.7030](#)].
- [66] ICECUBE, PAN-STARRS1 SCIENCE CONSORTIUM, SWIFT and PTF collaborations, M.G. Aartsen et al., *The detection of a SN II_n in optical follow-up observations of IceCube neutrino events*, *Astrophys. J.* **811** (2015) 52 [[arXiv:1506.03115](#)].
- [67] T. DeYoung et al., *IceTray: a software framework for IceCube*, in *Computing in high energy physics and nuclear physics. Proceedings, Conference, CHEP'04*, Interlaken Switzerland, 27 September–1 October 2004, [pg. 463](#).
- [68] KM3NET collaboration, S. Adrian-Martinez et al., *Letter of intent for KM3NeT 2.0*, *J. Phys.* **G 43** (2016) 084001 [[arXiv:1601.07459](#)].
- [69] BAIKAL collaboration, A.D. Avrorin et al., *The prototyping/early construction phase of the BAIKAL-GVD project*, *Nucl. Instrum. Meth. A* **742** (2014) 82 [[arXiv:1308.1833](#)].
- [70] ICECUBE collaboration, M.G. Aartsen et al., *IceCube-Gen2: a vision for the future of neutrino astronomy in Antarctica*, [arXiv:1412.5106](#).
- [71] ICECUBE PINGU collaboration, M.G. Aartsen et al., *Letter of intent: the Precision IceCube Next Generation Upgrade (PINGU)*, [arXiv:1401.2046](#).
- [72] ARA collaboration, P. Allison et al., *Performance of two Askaryan Radio Array stations and first results in the search for ultrahigh energy neutrinos*, *Phys. Rev. D* **93** (2016) 082003 [[arXiv:1507.08991](#)].