



**IceCube Maintenance and Operations
Fiscal Year 2016 / 2017 PY1 Annual Report**

April 1, 2016 – January 31, 2017

Submittal Date: March 1, 2017

University of Wisconsin–Madison

This report is submitted in accordance with the reporting requirements set forth in the IceCube Maintenance and Operations Cooperative Agreement, PLR-1600823.

Foreword

This FY2016 / 2017 (PY1) Annual Report is submitted as required by the NSF Cooperative Agreement PLR-1600823. This report covers the ten-month period beginning April 1, 2016 and concluding January 31, 2017. The status information provided in the report covers actual common fund contributions received through September 30, 2016 and the full 86-string IceCube detector (IC86) performance through January 31, 2017.

Table of Contents

Foreword	2
Section I – Financial/Administrative Performance	4
Section II – Maintenance and Operations Status and Performance	6
Detector Operations and Maintenance	6
Computing and Data Management Services	12
Data Processing and Simulation Services	23
Software	24
Calibration	25
Program Management	26
Section III – Project Governance and Upcoming Events	30

Section I – Financial/Administrative Performance

The University of Wisconsin–Madison is maintaining three separate accounts with supporting charge numbers for collecting IceCube M&O funding and reporting related costs: 1) NSF M&O Core account, 2) U.S. Common Fund account, and 3) Non-U.S. Common Fund account.

The first PY1 installment of \$3,500,000 was released to UW–Madison to cover the costs of maintenance and operations during the first six months of PY1 (FY2016): \$498,225 was directed to the U.S. Common Fund account based on the number of U.S. Ph.D Authors in the last version of the institutional MoU’s, and the remaining \$3,001,775 was directed to the IceCube M&O Core account. The second PY1 installment of \$1,750,000 was released to UW–Madison to cover part of the costs of maintenance and operations during the remaining six months of PY1 (FY2017): \$249,113 was directed to the U.S. Common Fund account, and the remaining \$1,500,888 was directed to the IceCube M&O Core account. The last PY1 (FY2017) installment of \$1,750,000 is planned to be released to cover part of the second half of PY1 (Figure 1).

PY1: FY2016 / FY2017	Funds Awarded to UW for Apr 1, 2016 – Sept 30, 2016	Funds Awarded to UW for Oct 1, 2016 – March 31, 2017	Funds to Be Awarded to UW for Oct 1, 2016 – March 31, 2017
IceCube M&O Core account	\$3,001,775	\$1,500,888	\$1,500,888
U.S. Common Fund account	\$498,225	\$249,113	\$249,113
TOTAL NSF Funds	\$3,500,000	\$1,750,000	\$1,750,000

Table 1: NSF IceCube M&O Funds – PY1 (FY2016 / FY2017)

Of the IceCube M&O PY1 (FY2016/2017) Core funds, \$952,147 were committed to the U.S. subawardee institutions based on their statement of work and budget plan. The institutions submit invoices to receive reimbursement against their actual IceCube M&O costs. Table 2 summarizes M&O responsibilities and total FY2016/2017 funds for the subawardee institutions.

Institution	Major Responsibilities	Funds
Lawrence Berkeley National Laboratory	DAQ maintenance, computing infrastructure	\$82,889
Pennsylvania State University	Computing and data management, simulation production, DAQ maintenance	\$68,771
University of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance	\$162,158
University of Maryland at College Park	IceTray software framework, online filter, simulation software	\$587,577
University of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools	\$23,870
Michigan State University	Simulation software, simulation production	\$26,882
Total		\$952,147

Table 2: IceCube M&O Subawardee Institutions – PY1 (FY2016/2017) Major Responsibilities and Funding

IceCube NSF M&O Award Budget, Actual Cost and Forecast

The current IceCube NSF M&O 5-year award was established in the middle of Federal Fiscal Year 2016, on April 1, 2016. The following table presents the financial status ten months into the Year 1 of the award, and shows an estimated balance at the end of PY1.

Total awarded funds to the University of Wisconsin (UW) for supporting IceCube M&O from the beginning of PY1 through mid-year PY1 are \$5,250K. With the last PY1 planned installment of \$1,750K, the total PY1 budget is \$7,000K. Total actual cost as of January 31, 2017 is \$5,666K and open commitments are \$578K. The current balance as of January 31, 2017 is \$757K. With a projection of \$974K for the remaining expenses during the final two months of PY1, the estimated negative balance at the end of PY1 is -\$217K, which is 3.1% of the PY1 budget (Table 3).

(a)	(b)	(c)	(d)= a - b - c	(e)	(f) = d - e
YEARS 1 Budget Apr.'16-Mar.'17	Actual Cost To Date through Jan. 31, 2017	Open Commitments on Jan. 31, 2017	Current Balance on Jan. 31, 2017	Remaining Projected Expenses through Mar. 2017	End of PY1 Forecast Balance on Mar. 31, 2017
\$7,000K	\$5,666K	\$578K	\$757K	\$974K	-\$217K

Table 3: IceCube NSF M&O Award Budget, Actual Cost and Forecast

IceCube M&O Common Fund Contributions

The IceCube M&O Common Fund was established to enable collaborating institutions to contribute to the costs of maintaining the computing hardware and software required to manage experimental data prior to processing for analysis.

Each institution contributes to the Common Fund, based on the total number of the institution’s Ph.D. authors, at the established rate of \$13,650 per Ph.D. author. The Collaboration updates the Ph.D. author count twice a year before each collaboration meeting in conjunction with the update to the IceCube Memorandum of Understanding for M&O.

The M&O activities identified as appropriate for support from the Common Fund are those core activities that are agreed to be of common necessity for reliable operation of the IceCube detector and computing infrastructure and are listed in the Maintenance & Operations Plan.

Table 4 summarizes the planned and actual Common Fund contributions for the period of April 1, 2016–March 31, 2017, based on v20.0 of the IceCube Institutional Memorandum of Understanding, from April 2016. The final non-U.S. contributions are underway, and it is anticipated that most of the planned contributions will be fulfilled.

	Ph.D. Authors	Planned Contribution	Actual Received
Total Common Funds	139	\$1,917,825	\$1,671,395
U.S. Contribution	78	\$1,064,700	\$1,064,700
Non-U.S. Contribution	61	\$853,125	\$812,007

Table 4: Planned and Actual CF Contributions for the period of April 1, 2016–March 31, 2017

Section II – Maintenance and Operations Status and Performance

Detector Operations and Maintenance

Detector Performance — During the period from April 1, 2016, to February 1, 2017, the detector uptime, defined as the fraction of the total time that some portion of IceCube was taking data, was 99.74%, exceeding our target of 99%. The clean uptime for this period, indicating full-detector analysis-ready data, was 97.73%, exceeding our target of 95%. Historical total and clean uptimes of the detector are shown in Figure 1. The modest decrease in clean uptime in December and January is due to maintenance during the austral summer season.

Figure 2 shows a breakdown of the detector time usage over the reporting period. The partial-detector good uptime was 1.12% of the total and includes analysis-ready data with fewer than all 86 strings. Excluded uptime includes maintenance, commissioning, and verification data and required 0.90% of detector time. The unexpected detector downtime was limited to 0.26%.

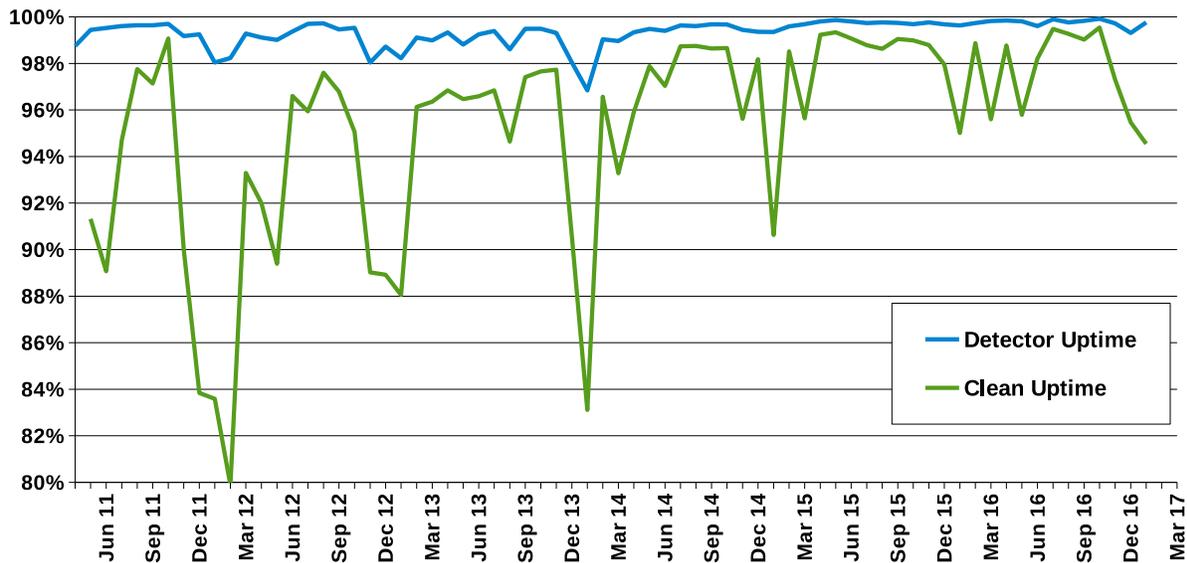


Figure 1: Total IceCube Detector Uptime and Clean Uptime

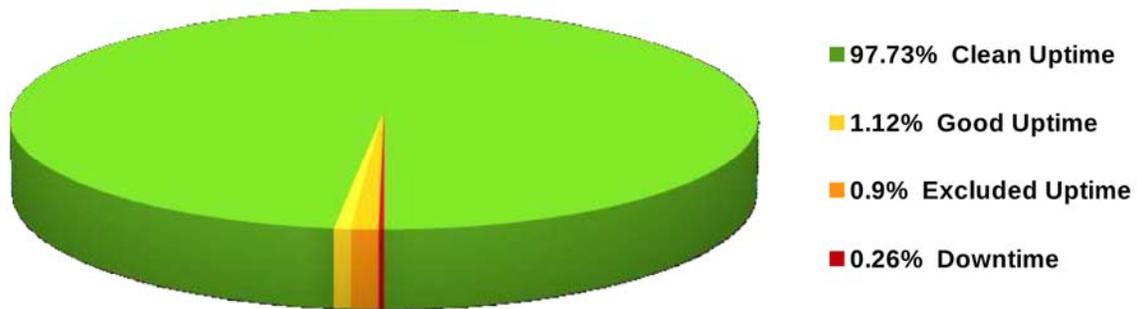


Figure 2: Cumulative IceCube Detector Time Usage, April 1, 2016 – February 1, 2017

Hardware Stability — The last DOM failures (2 DOMs) occurred during a power outage on May 22, 2013. No DOMs have failed during this reporting period. The total number of active DOMs remains 5404 (98.5% of deployed DOMs), plus four scintillator panels and the IceACT trigger mainboard. No custom data acquisition hardware components in the ICL have failed during the reporting period.

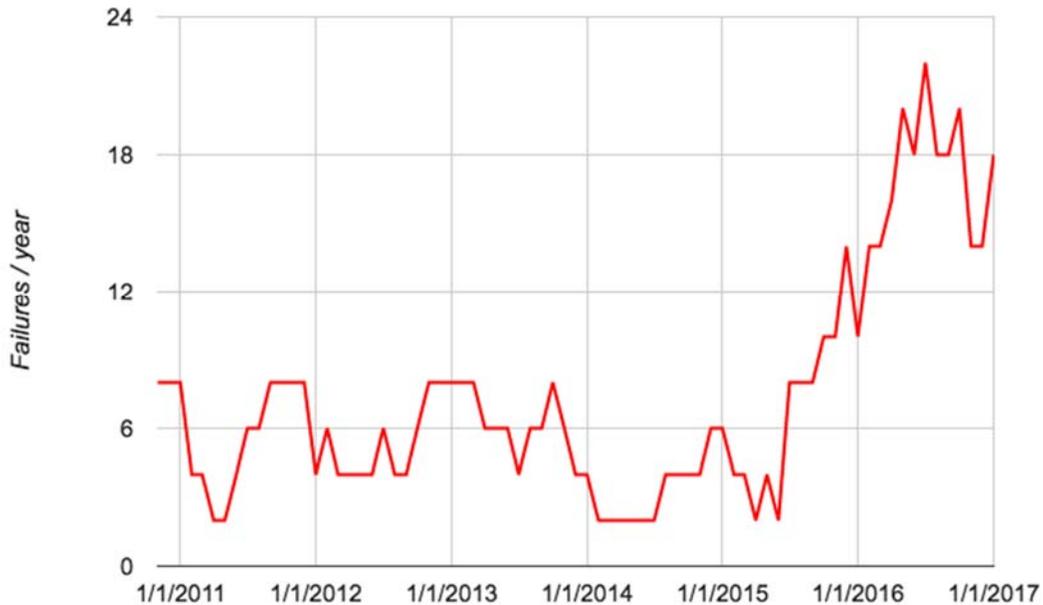


Figure 3: Acopian (DOM) power supply failure rate vs. time (6-month rolling average). The failure rate has been stabilized with a hardware replacement and is expected to decrease.

The failure rate in the commercial Acopian power supplies that supply the DC voltage to the DOMs from the DOMHubs has increased, starting in late 2015 (Figure 3). Failed units from the past two seasons have been sent back to the manufacturer for analysis, leading to a modified design with improved robustness. During the 2016–17 austral summer season, we performed a complete replacement of the supplies using the updated units. A few units suffered early burn-in failure, but the failure rate appears to have stabilized at a lower rate, with one unit failing in 2017 so far.

A subtle failure mode in several of the commercial DOMHub ATX power supplies that caused them to lose their redundancy was discovered by the winterovers in 2016. The non-redundant units were replaced during the 2016–17 summer season.

IC86 Physics Runs — The sixth season of the 86-string physics run, IC86–2016, began on May 20, 2016. Detector settings were updated using the latest yearly DOM calibrations from March 2016. Two new DAQ triggers were added to the configuration: an IceTop infill trigger, using the infill tanks in the core of the array to target low-energy cosmic ray air showers, and a scintillator calibration trigger. Filter changes include a new magnetic monopole search filter, a unified selection for optical and gamma-ray follow-up events, and retirement of the Galactic Center filter.

Starting with IC86–2015, we have implemented changes to the methodology for producing online quasi-real-time alerts. Neutrino candidate events at a rate of 3 mHz are now sent via Iridium satellite, so that neutrino coincident multiplets (and thus candidates for astrophysical transient sources) can be rapidly calculated and distributed in the Northern Hemisphere. This change enables significant flexibility in the type of fast alerts produced by IceCube. With the IC86–2016 release, we have completed this transition, retiring the South Pole analysis components and moving all follow-up analysis to WIPAC.

For this reporting period, the average TDRSS daily transfer rate was approximately 77 GB/day. The bandwidth saved from filter optimizations has allowed us to reserve more for HitSpool data transfer, saving all untriggered IceCube hits within a time period of interest. We have recently added alert mechanisms by which we save HitSpool data around Fermi-LAT solar flares and LIGO gravitational wave alerts; the second LIGO physics run started on 30 November 2016.

Data Acquisition — The IceCube Data Acquisition System (DAQ) has reached a stable state, and consequently the frequency of software releases has slowed to the rate of 3–4 per year. Nevertheless, the DAQ group continues to develop new features and patch bugs. During the reporting period, the following accomplishments are noted:

- Delivery of the pDAQ:New_Glarus release in May 2016, providing support for the new IceTop infill trigger, performance enhancements in the DOMHub components, and delivery of time calibration monitoring quantities to IceCube Live.
- Delivery of the pDAQ:One_Barrel release in September 2016, providing speed improvements to IceCube’s primary trigger algorithm and better cleanup if DOMs drop from data-taking.
- Delivery of the pDAQ:Potosi release in January 2017, which includes performance improvements and improved robustness in the DOMHub components (Figure 4), and lays the ground work for the use of the HitSpool cache as the primary hit stream.
- Work towards the next DAQ release in spring 2017, which will include the separation of the hub-based data processing from the run-based DAQ system to pare detector downtime to the absolute minimum. An improved process communication framework used among the DAQ components is also under development.

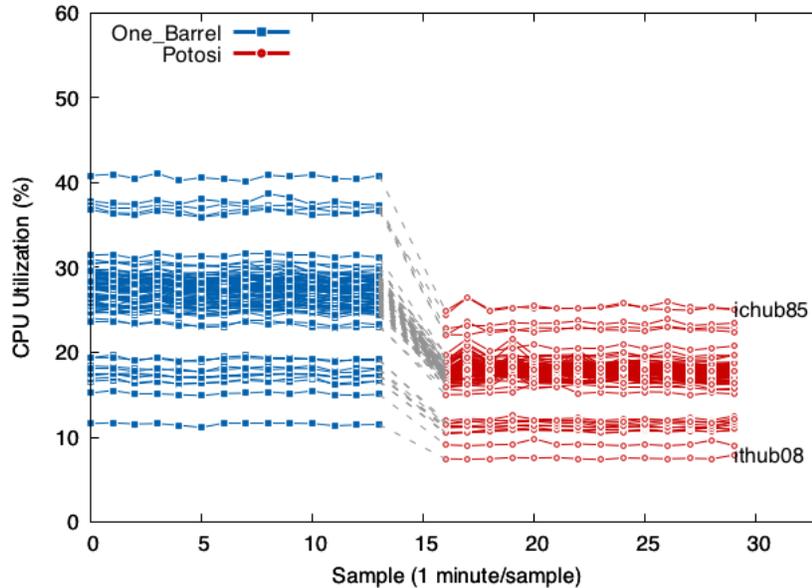


Figure 4: CPU utilization improvement in the 2017 DAQ “Potosi” release (lower is better).

Online Filtering — The online filtering system (“PnF”) performs real-time reconstruction and selection of events collected by the data acquisition system and sends them for transmission north via the data movement system. In addition to the standard release for the IC86–2016 physics run start including the filter changes for the season, three additional releases (V16-07-01, V16-07-02, and V16-07-03) addressed stability issues seen when delivering large quantities of monitoring data for the new I3Moni 2.0 system to IceCube Live. All PnF monitoring quantities are now being delivered to the new system and instability issues have been largely resolved.

Monitoring — The IceCube Run Monitoring system, I3Moni, provides a comprehensive set of tools for assessing and reporting data quality. IceCube collaborators participate in daily monitoring shift duties by reviewing information presented on the web pages and evaluating and reporting the data quality for each run. The original monolithic monitoring system processes data from various SPS subsystems, packages them in files for transfer to the Northern Hemisphere, and reprocesses them in the north for display on the monitoring web pages. In a new monitoring system under development (I3Moni 2.0), all detector subsystems report their data directly to IceCube Live. Major advantages of this new approach include: higher quality of the monitoring alerts; simplicity and easier maintenance; flexibility, modularity, and scalability; faster data presentation to the end user; and a significant improvement in the overall longevity of the system implementation over the lifetime of the experiment.

The I3Moni 2.0 infrastructure for collecting the monitoring data is in place at SPS, and monitoring quantities are now being collected from all major subsystems and displayed on IceCube Live web pages. Since December 2015, the I3Moni 2.0 beta release has been active. The public release was delayed from mid-2016 primarily due to the additional PnF development needed to deliver high-bandwidth monitoring quantities; this issue has been resolved. The new system will be phased into normal monitoring activities in early 2017, and the old system retired at the mid-year IC86–2017 run start.

Experiment Control — Development of IceCube Live, the experiment control and monitoring system, is still quite active. This reporting period has seen one major release with the following highlighted features:

- Live v2.9.3 (May 2016): 31 separate issues and feature requests have been resolved. This release adds tracking of problematic DOMs within IceCube Live, support for script-initiated HitSpool requests from the Northern Hemisphere, and improvements to the beta release of the I3Moni 2.0 system.
- Live v2.9.4 (December 2016): 65 separate issues and feature requests have been resolved, most relating to the I3Moni 2.0 system. A new reliable communication scheme for real-time alerts was added, and code supporting the retired ITS communications system has been removed.

Features planned for the next few releases include: finalization of Moni2.0 into its public release, and creating or improving dedicated monitoring pages for the JADE, SNDAQ, and realtime follow-up subsystems. The uptime for the I3Live experiment control system during the reporting period was 99.996%.

GCD Database System —The database for geometry, calibration, and detector status (GCD) information and associated code has grown difficult to maintain and is not easily extensible, either to new hardware such as the prototype scintillators, or to improved calibration techniques such as the recently-deployed online DOM gain corrections. We are developing a new GCD generation system that uses a new underlying database and rewritten interface code. We plan to have the new system ready for online deployment at the IC86–2017 run start.

Supernova System — The supernova data acquisition system (SNDAQ) found that 99.71% of the available data from April 1, 2016 through February 1, 2017 met the minimum analysis criteria for run duration and data quality for sending triggers. An additional 0.03% of the data is available in short physics runs with less than 10-minute duration. While forming a trigger is not possible in these runs, the data are available for reconstructing a supernova signal.

A new SNDAQ release (2016-09-09) was deployed that continues the cleanup of the build and deployment system. The build system has been completely revised and simplified and is now in line with other IceCube software. Efforts to include a data-driven trigger that is independent of an assumed signal shape are under way.

On July 30, 2016, a malfunctioning DOM caused the muon-subtraction calculation to report a false high-significance supernova alert. While this was quickly confirmed as a likely false alarm, the Wintervers followed the documented procedures and saved relevant secondary data. The error in the significance calculation has been corrected to make it robust against this rare DOM failure mode. Once power-cycled, the problematic DOM behaved normally again.

Surface Detectors — Snow accumulation on the IceTop tanks continues to reduce the trigger rate of the surface array by ~10%/year. Uncertainty in the attenuation of the electromagnetic component of air showers due to snow is the largest systematic uncertainty in IceTop’s cosmic

ray energy spectrum measurement. The snow also complicates IceTop’s cosmic-ray composition measurements, as it makes individual air showers look “heavier” by changing the electromagnetic/muon ratio of particles in the shower.

A decision by the NSF and the support contractor has been taken to stop any further snow management efforts. With this in mind, we have developed a plan to restore the full operational efficiency of the IceCube surface component using plastic scintillator panels on the snow surface above the buried IceTop tanks. By detecting particle showers in coincidence with IceTop and IceCube, the scintillator upgrade allows a) determination of the shower attenuation due to snow as a function of energy and shower zenith angle, and b) restoration of the sensitivity to low-energy showers lost due to snow accumulation on the tanks.

Four prototype scintillators were deployed during the 2015–16 pole season at IceTop stations 12 and 62. These scintillators have been integrated into the normal IceCube data stream and have been taking data continuously since early 2016. Development of a new version of the scintillators is underway, which includes a new digitization and readout system, a different photodetection technology (SiPMs), and a streamlined, lighter housing. We are working to minimize field logistics of any future scintillator installations with a streamlined deployment and trenching plan. Total power consumption of an array covering IceTop is expected to be modest: each station should use 25 Watts and 1500 Watts is anticipated for the complete proposed deployment including ICL hardware power consumption.

A prototype sonic snow measurement sensor was deployed during the 2016/17 Austral Summer Season by G. deWasseige (VUB Brussels) with support from WIPAC technical personnel. The sensor is working properly and is measuring the depth of snow above IceTop tank 37A, taking 1 measurement per day and sending the information in realtime to the IceCube Live system.

Additionally, a prototype air Cherenkov telescope (IceACT) was installed on the IceCube Lab. During the polar night, IceACT can be used to cross-calibrate IceTop by detecting the Cherenkov emission from cosmic ray air showers; it may also prove useful as a supplementary veto technique. After commissioning last austral summer, IceACT was uncovered at sunset and took data during the austral winter. The mainboard trigger system allowed cross-calibration of the time offset between IceTop and IceACT, and coincident events have been identified. Upgrades to the IceACT DAQ computer and a new LED calibration system were installed during the 2016–17 pole season; a camera upgrade has been postponed to the 2017–18 season.

Operational Communications & Real-time Alerts — Communication with the IceCube winterovers, timely delivery of detector monitoring information, and login access to SPS are critical to IceCube’s high-uptime operations. Several technologies are used for this purpose, including ssh/scp, the IMCS e-mail system, and IceCube’s own Iridium modem(s).

We have now developed our own Iridium RUDICS-based transport software (IceCube Messaging System, or I3MS) and have moved monitoring data to our own Iridium modems as of the 2015–16 austral summer season. We retired the Iridium short-burst-data system (ITS) in October 2016 and have added its modem to the I3MS system to increase bandwidth.

This past austral summer season, the contractor installed a new antenna housing “doghouse” on the roof of the ICL. This will alleviate overcrowding and self-interference of antennas in the existing doghouse. We are now monitoring the temperature of the new doghouse over the winter. Assuming no problems are encountered, we will migrate all GPS and Iridium antennas to the new housing in the 2017–18 season.

Personnel — No changes.

Computing and Data Management Services

South Pole System – A significant amount of time was spent on training and (re-)hiring new WinterOvers. Unfortunately, both of our primary candidates had to be replaced late in the training period. Due to the long time scales UTMB requires to review medical information and the lack of communication during this process only one of our replacement WinterOvers could be hired in time to receive basic training in Madison. The second replacement operator had to be trained on site. The current PQ review policies and PQ expiration policies need to be addressed for next year in order to ensure continuity in operator training and availability as well as to help reduce the unnecessary costs.

The month of October was dedicated to finalizing preparations for the upcoming South Pole summer season. Highest priority was assigned to identifying and PQ’ing new alternate WinterOver candidates after both primary candidates had to be replaced with their backups for medical and/or personal reasons. Two suitable alternates from within the collaboration were identified and underwent the physical qualification process in case one of the backup candidates had to be replaced during the summer season.

Several new test machines, virtual and physical, were added to the SPTS to test upcoming installations at Pole (sky camera, snow sensor, ...). This also required the installation and testing of new software components to communicate with the corresponding sensor hardware. A more realistic testing environment was provided for the JADE data archival system (additional client hosts, more archival disks available).

A new monitoring plugin was developed to keep better track of power consumption in the IceCube Lab (ICL). The plugin was successfully tested at the SPTS using the same hardware currently installed at South Pole.

Operating System security patches and bug fix updates were thoroughly tested at SPTS to exclude any potential issues with version mismatches. No issues were identified and all patches and updates were approved for rollout at the production system at Pole during the upcoming season.

Several improvements were made on the monitoring mirror server that allows experts in the North to access real-time monitoring data at the Pole during satellite coverage. The mirror server required several modifications to be compatible with the new Nagios/CheckMK software versions to be installed and all functionality was restored.

On top of the priority list was training of the new WinterOvers who missed most or even all of training in Madison due to their short-notice activation. Experts deploying from Madison spent a significant amount of time in familiarizing the new operators with IceCube hardware, software and operational procedures. The WinterOvers were quick learners and by the end of the season most of the training topics were covered and they felt comfortable operating the detector and maintaining the data center in ICL.

In order to ensure stable detector operation over the course of the winter about 50% of our Uninterruptable Power Supply (UPS) batteries were replaced. The old batteries had reached their vendor recommended replacement age of 4 years. The replacement went smooth and without any interruption to IceCube data taking. The new batteries will maximize the time we can operate the detector in case of power issues on station.

The new Nagios/CheckMK server was deployed at South Pole. Nagios Core was upgraded from version 3.5 to version 4.1, providing a significant performance improvement due to multi-threading ability of the new release. Further on-site testing revealed a bug in the new version that could not be identified during testing in Madison as it related to the tie-in into ASC's VoIP phone system for paging alerts. This required an in-situ upgrade of the Nagios system to a version released in late October. The upgrade went without issues and addressed the problem properly; the new monitoring system is fully operational and performs as expected.

The total number of monitored quantities in the IceCube data center was increased from ~9,300 to over 10,000 - most significantly the addition of detailed data base monitoring (MongoDB).

The old IceCube Teleport System (ITS), a custom written Iridium messaging system based on Iridium's Short Burst Data protocol was replaced with a more efficient, more reliable implementation based on Iridium's new RUDICS technology. The ITS messaging system was decommissioned and all modems migrated over to the RUDICS channel, providing reliable 24/7 data exchange between the ICL at South Pole and the data center in Madison.

The RUDICS iridium installation at the IceCube Lab (ICL) is experiencing electro-magnetic interference as a result of the antennas being located too closely to each other and to other antennas. To address this issue a new antenna box ("doghouse") was installed on the roof of ICL. At the time the experts left the ice the structure was not yet commissioned for IceCube use (fire panel installation pending) but the structure was equipped with environmental monitoring sensors so we can monitor the conditions inside the new structure over the winter. Next summer season we will install mounting rails and grounding gear and migrate our antennas from the old, small antenna box to the new structure to reduce electro-magnetic interference amongst the antennas.

Several bugs and issues had been identified on our firewalls over the course of the last year and were addressed with Dell/Sonicwall who provided a firmware upgrade to us. The upgrade did not cause any operational downtime as the two firewalls at South Pole are configured in a redundant setup. All known issues were resolved.

The yearly security patches and bug fixes were installed on all of our Scientific Linux servers.

No issues to report. Detector downtime was minimal as individual services were moved around between nodes by the WinterOvers and on-site experts. The only system suffering a few hours of downtime was SupernovaDAQ which could eventually be revived by experts in the North.

A lot of effort went into migrating our computers system management into a central configuration engine (Puppet). Several key components were moved out of the static kickstart setup into the more dynamic and real-time Puppet tool. This allows operators to manage our compute nodes more easily and more consistently.

All the spare parts used during the previous winter were restocked and all failed components were either replaced or fixed by hardware experts. We should have now sufficient spares again on site to make it through the next season.

All retrograde cargo was packed up and shipped North. This included two copies of all of IceCube/ARA archival data from last year as well as decommissioned and broken equipment. The expected arrival date in Madison is April 2017.

The IceCube M&O South Pole System encountered two significant obstacles to operation during this first project year: (a) the aforementioned difficulties encountered with WinterOver PQ; (b) a cargo shipment sustained visible damage and water infiltration en route to Pole. An investigation was conducted by the contractor into the root cause of the latter incident, and unfortunately, this investigation was unable to determine the breakdown in the contractor's logistic chain. Nevertheless, despite these setbacks, all of our main objectives were achieved.

Data Transfer – Data transfer has performed nominally over the past ten months. Between April 2016 and January 2017 a total of 24.6 TB of data were transferred from the South Pole to UW-Madison via TDRSS, at an average rate of 80.37 GB/day. Figure 6 shows the daily satellite transfer rate and weekly average satellite transfer rate in GB/day through January 2017. The IC86 filtered physics data are responsible for 95% of the bandwidth usage.

In April 2016 we reached an important milestone that was the completion of the replacement of the Northern hemisphere part of the data transfer system. The INGEST software that had been running since 2005 for receiving the data from the SPTR satellite system and storing it in the data warehouse at UW-Madison was replaced by the new JADE software. The new software is much more stable and resilient than its predecessor. This has been confirmed by the last ten months experience from the Winterovers and IT staff at UW-Madison operating the system, which has run smoothly with less maintenance effort.

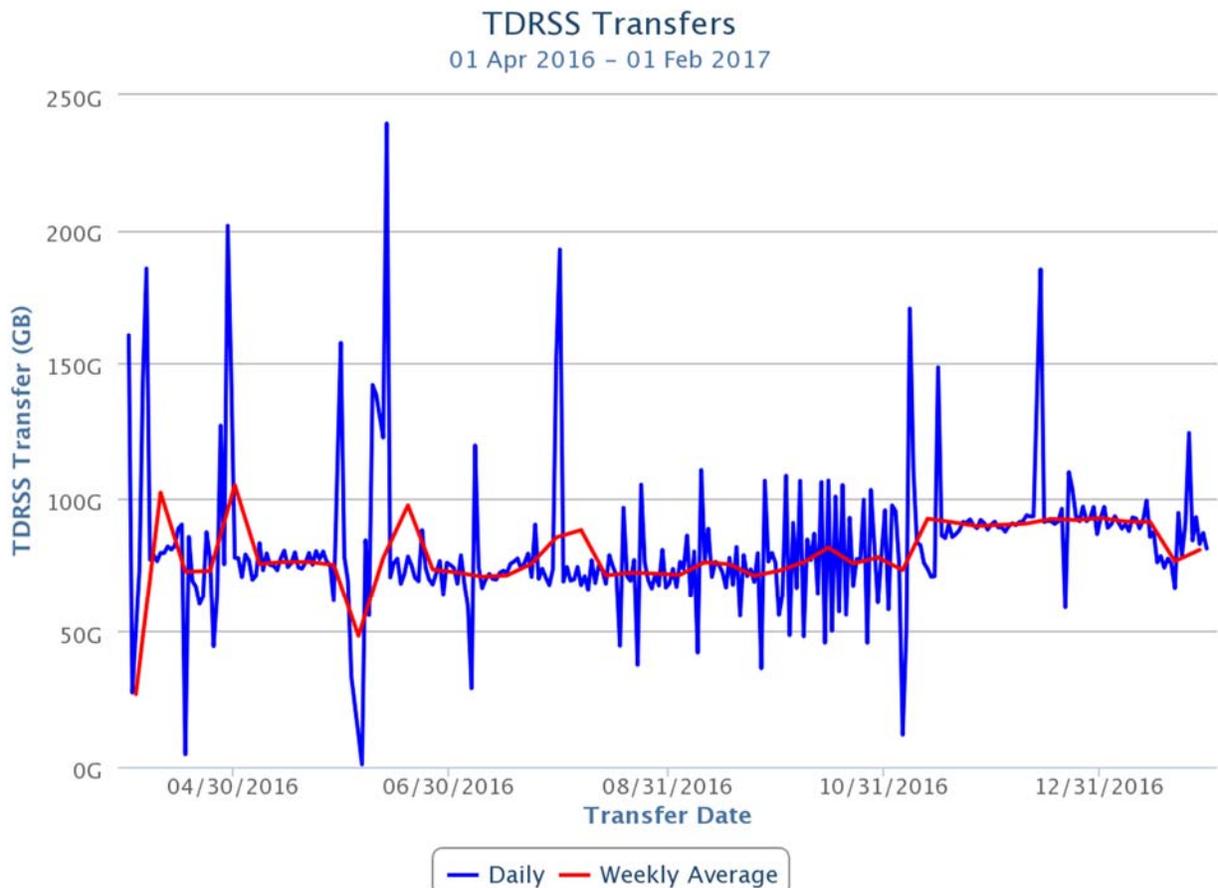


Figure 6: TDRSS Data Transfer Rates, April 1, 2016–September 31, 2016. The daily transferred volumes are shown in blue and, superimposed in red, the weekly average daily rates are also displayed.

Data Archive – The IceCube raw data are archived to two copies on independent hard disks. During the reporting period (April 2016 to January 2017) a total of 369.27 TB of unique data were archived to disk averaging 1.21 TB/day.

In May 2016, we started using the new JADE software for handling the raw data received from the South Pole in archival disk drives. Once JADE has processed the files, their metadata is indexed and the files are ready for being replicated to the long-term archive.

In December 2015, a Memorandum of Understanding was signed between UW-Madison and NERSC/LBNL by which NERSC agreed to provide long-term archive services for the IceCube data until 2019. By implementing the long-term archive functionality using a storage facility external to the UW-Madison data center, we aim for an improved service at lower cost, since large facilities that routinely manage data at the level of hundreds of petabytes benefit of economies of scale that ultimately make the process more efficient and economical.

During the reporting period, additional functionality has been developed as part of the JADE software that allows handling the long-term archive data flows. One of the main requirements driving this development is the need for bundling small files into larger ones of few hundred

Gigabytes size, in order to ensure that the tape drives operate in a high efficiency regime. The first version of JADE that was able to manage long-term archive data flows was available in September. Since then, we have been using it to bundle IceCube data files at UW-Madison and transfer them to NERSC at a rate of up to 10 TB/day. At the time of writing this report, the total volume of data archived at NERSC was 226 TB. This number is still small compared to the total amount of data that we need to archive, so the plan now is to keep this archive stream constantly active while working on further JADE functionality that will allow us to steadily increase the performance.

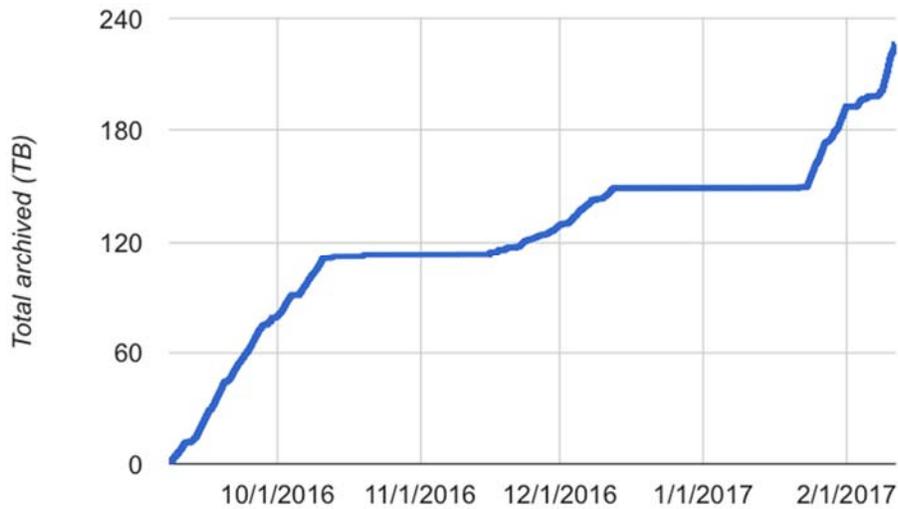


Figure 7: Volume of IceCube data archived at the NERSC HPSS tape facility by the JADE Long Term Archive service as a function of time.

Computing Infrastructure at UW-Madison – The total amount of data stored on disk in the data warehouse at UW-Madison is 5806 TB: 1600 TB for experimental data, 3933 TB for simulation and analysis and 274 TB for user data.

Two Dell Compellent SCv2080 appliances, each containing 168 4TB drives, were purchased during the first half of the reporting period to expand the capacity of the data warehouse. The main need for additional space was mostly to have room to handle the raw data received yearly from the South Pole, which is now ingested into the system as soon as archival hard drives arrive to Madison. The first appliance was brought online in May, and the second one in August. Together they added up about 1024 TB of usable space to the system.

As described in previous reports, we suffered a serious incident in our main Lustre filesystem in July that ended up with data loss. Fortunately all of the affected files were simulated or processed data that could be easily re-generated. The impact of this incident was aggravated by few bugs in the Lustre version that we were running in production at the time, version 2.5.3. As a follow-up of this incident, we upgraded all our filesystems to the last Lustre stable version 2.8.0 during the month of November.

One SuperMicro 6048R-E1CR36L server with 36 8TB disks was purchased in October. The plan

for this new storage server is to dedicate it for simulation production activities. In particular, to store the temporary files that simulation tasks generate along the simulation chain. By separating this simulation load from the regular analysis and processing workloads we seek to achieve higher overall performance and stability.

A disk expansion for the data warehouse was purchased in November, targeting additional storage for the upcoming multi-year reprocessing activity known as “Pass2”. The chosen solution consisted of two Dell PowerEdge R430 servers connected to two disk array controllers: one Dell PowerVault MD3460 plus one Dell PowerVault MD3060e expansion, each one holding 60 10TB drives. One of the advantages of this system is that it is almost identical to some of the existing units in the targeted file system, and those have proven to be very reliable. The new systems provide a total usable space of 962 TB.

The IceCube computing cluster at UW-Madison has continued to deliver reliable data processing services. Boosting the GPU computing capacity has been a high priority of the project since the Collaboration decided to use GPUs for the photon propagation part of the simulation chain back in 2012. Direct photon propagation was found to provide the precision required, and it happens to be very well suited for the GPU hardware, running about 100 times faster than in CPUs.

An expansion of the GPU cluster was purchased in September, consisting of seven SuperMicro 4027GR-TR chassis each containing eight Nvidia GTX 1080 GPU cards, two Xeon E5-2637 v4 processors, 64 GB of RAM and 2TB of disk. This cluster has been deployed at the Wisconsin Institutes for Discovery (WID) datacenter and it has replaced the old IceCube GPU cluster that was deployed there in January 2012. The Memorandum of Understanding signed by the WID IT department, the Center for High Throughput Computing (CHTC) and WIPAC by which WIPAC could host one rack of GPU servers at WID in exchange of sharing 30% of the cluster with CHTC is still active and was used to update this cluster. The new servers provide a compute power about five times larger than the old ones, using the same amount of electrical power.

The new GPU expansion was brought online on October 3rd 2016. After this upgrade, the IceCube GPU cluster at UW-Madison has a total of 376 GPU cards.

The focus for the GPU cluster is to provide the required capacity to fulfill the Collaboration direct photon propagation simulation needs. These needs have been estimated to be around 30% higher than the capacity of the GPU cluster at UW-Madison. Additional GPU resources at several IceCube sites, plus XSEDE allocations, allow us to provide that required capacity.

Distributed Computing – In March 2016, a new procedure to formally gather computing pledges from collaborating institutions was started. This data will be collected twice a year as part of the already existing process by which every IceCube institution updates its MoU before the Collaboration week meeting. Institutions that pledge computing resources for IceCube are asked to provide information on the average number of CPUs and GPUs that they commit to provide for IceCube simulation production during the next period. Table 4 shows the computing pledges per institution as of September 2016:

Site	Pledged CPUs	Pledged GPUs
Aachen	83	29
Alabama		6
Canada	1055	41
Brussels		14
Chiba	196	6
Delaware	272	
DESY-ZN	1050	160
Dortmund	2642	10
LBNL	114	
Mainz	24	8
Marquette	60	16
MSU	500	8
UMD	350	24
UW-Madison	700	301
TOTAL	7046	622

Table 4: Computing pledges for simulation production from IceCube Collaboration institutions as of September 2016.

The plan is to implement a feedback planning process by which the numbers from available resources from computing pledges are regularly compared to the simulation production needs and resources used. The goal is to be able to manage more efficiently the global resource utilization and to be able to react to changes in computing needs required to meet IceCube science goals.

A strong focus has been put in the last years to expand the distributed infrastructure and make it more efficient. The main strategy to accomplish this has been to simplify the process for sites to join the IceCube distributed infrastructure, and also to reduce the effort needed to keep sites connected to it. To do this, we have progressively implemented an infrastructure based on Pilot Jobs. Pilot Jobs provide a homogeneous interface to heterogeneous computing resources. Also, they enable more efficient scheduling by delaying the decision of matching resources to payload.

In order to implement this Pilot Job paradigm for the distributed infrastructure IceCube makes use of some of the federation technologies within HTCondor¹. Pilot Jobs in HTCondor are called “glideins” and consist of a specially configured instance of the HTCondor worker node component, which is then submitted as a job to external batch systems.

Several of the sites that provide computing for IceCube are also resource providers for other scientific experiments that make use of distributed computing infrastructures. Thanks to this they already provide a standard (Grid) interface to their batch systems. In these cases we can leverage the standard GlideinWMS infrastructure operated by the Open Science Grid² project for integrating those resources into the central pool at UW-Madison and provide transparent access

¹ <http://research.cs.wisc.edu/htcondor/>

² <https://www.opensciencegrid.org/>

to them via the standard HTCondor tools. The sites that use this mechanism to integrate with the IceCube global workload system are: Aachen, Canada, Brussels, DESY, Dortmund, Wuppertal and Manchester.

Following the negotiations to integrate Manchester in the IceCube distributed infrastructure, in June the UK Grid management (GridPP) officially took the decision to support IceCube across the UK sites and added it to the list³ of UK approved VOs. This is a nice example of the advantages we see of using standard tools and interfaces for building the IceCube distributed system.

Some of the IceCube collaborating institutions that provide access to local computing resources do not have a Grid interface. Instead, access is only possible by means of a local account. To address these cases we developed a lightweight version of a glidein Pilot Job factory that can be deployed as a cron job in the user's account. The codename of this software is "pyGlidein" and it allows us to seamlessly integrate these local cluster resources with the IceCube global workload system so that jobs can run anywhere in a way which is completely transparent for users. The sites that currently use this mechanism are: Canada, Brussels, DESY, Mainz, MSU, Dortmund and Uppsala. There are ongoing efforts at the Delaware, Chiba and LBNL sites to deploy the pyGlidein system.

Beyond the computing capacity provided by IceCube institutions, and the opportunistic access to Grid sites that are open to share their idle capacity, IceCube started exploring the possibility of getting additional computing resources from targeted allocation requests submitted to Supercomputing facilities such as the NSF Extreme Science and Engineering Discovery Environment (XSEDE). In October 2015 a research allocation was submitted to XSEDE that obtained positive reviews and was finally awarded compute time in two GPU-capable systems: Comet⁴, with 5.543.895 Service Units (SU) granted, and Bridges⁵, with 512.665 SU - allocation number TG-PHY150040.

IceCube simulation jobs are running in Comet since March 2016. The use of software distribution via CVMFS and Pilot Job framework based in HTCondor glideins facilitated the quick integration of this system in the IceCube simulation framework. The graph in Figure 8 shows the Comet utilization by IceCube jobs as seen in the XSEDE accounting portal. The steady usage over many months throughout the year shows that we can integrate specialized resources such as XSEDE supercomputers with the IceCube workload system and make use of them in a sustained manner. It should be pointed out though that the Comet system has a limited amount of GPU nodes (36 nodes with two Nvidia K80 each) therefore the total GPU capacity in this machine is quite limited. In particular IceCube cannot fully utilize its allocated 5.543.895 SUs using only GPU time. In order to be able to make maximal usage of the allocated resources in Comet, we considered the possibility of transferring part of the Comet allocation into other systems with larger GPU capacity such as XStream⁶. For this purpose, we requested a

³ https://www.gridpp.ac.uk/wiki/GridPP_approved_VOs

⁴ <https://portal.xsede.org/sdsc-comet>

⁵ <https://portal.xsede.org/psc-bridges>

⁶ <https://portal.xsede.org/stanford-xstream>

supplement test allocation in XStream in order to test the IceCube simulation workload there. The tests uncovered a problem in the Parrot⁷ software that we use for accessing CVMFS in systems that do not provide it natively like XStream. The issue was reported to the Parrot developers in University of Notre Dame who promptly released a patch⁸. By the time we got IceCube GPU jobs successfully running in XStream in December, we were informed that a large transfer from Comet was not possible since XStream was already fully committed.

In order to be able to fully utilize the granted XSEDE allocation, by the end of November we requested a six months extension. The extension was granted but we were informed that only a fraction of the Comet allocation was going to be available for the first half of 2017. Due to that, we also started running CPU jobs in Comet for the last weeks of the year to try to make maximal use of the allocation before it partially expired. This burst of computing activity is clearly seen in the accounting graph in Figure 8.

The Bridges XSEDE system, at the Pittsburgh Supercomputing Center, was not fully commissioned until May 2016. From that moment we were in contact with the local administrators to prepare the system to be usable by IceCube. A special configuration had to be developed to enable outgoing connectivity from the worker nodes as required by the Pilot Job system. This was in place in August, when we started testing the system and adapting IceCube workloads to use it. Bridges was the first site in production using CentOS7 Operating System for simulation, therefore some adaptation work was required. Similarly to what had happened in XStream, the adaptation work uncovered a bug in Parrot that showed up in CentOS7. The issue was reported to the developers and a patch⁹ was produced. We finally started running IceCube simulation regularly in Bridges in November 2016.

As a result of our feedback and interaction, the Cooperative Computing Lab group at University of Notre Dame published a post in the “Community Highlights” section of their web page¹⁰ in February publicizing the IceCube research and our usage of distributed computing infrastructures such as XSEDE.

In December the Bridges technical team at the Pittsburgh Supercomputing Center contacted us to inform us that the acceptance testing on the new Bridges nodes had finalized so the new Nvidia Tesla P100 GPU were now available for us to use. We quickly confirmed that our code could run in those nodes with no major issues and started using the new more powerful units. In January, Bridges published an article¹¹ in the Inside HPC web to announce their successful upgrade. IceCube was one of the highlighted projects in this article as one “illustrating the potential of HPC data analytics to complement scientific instruments”.

Also in December we were approached by the Comet technical team at the San Diego Supercomputing Center asking us to help them benchmark their new Nvidia Tesla P100 nodes with our simulation code. Our tests helped to uncover an issue with the Nvidia driver version in

⁷ <http://ccl.cse.nd.edu/software/parrot/>

⁸ <https://github.com/cooperative-computing-lab/cctools/issues/1548>

⁹ <https://github.com/cooperative-computing-lab/cctools/issues/1476>

¹⁰ <http://ccl.cse.nd.edu/community/stories/>

¹¹ <http://insidehpc.com/2017/01/bridges-supercomputer-completed/>

the new Comet nodes.

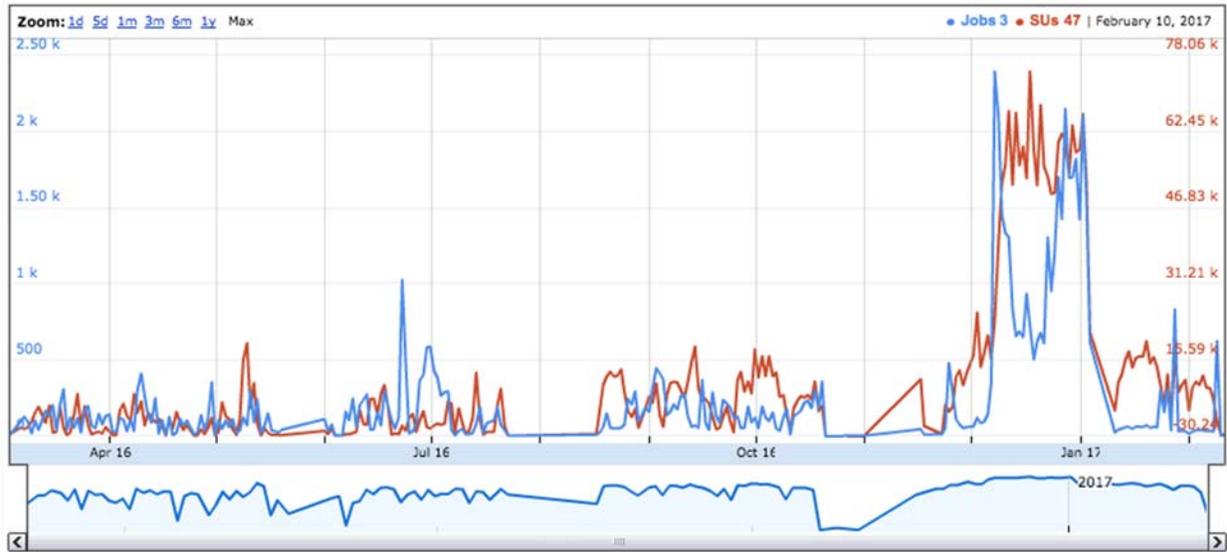


Figure 8: Daily usage accounting for the IceCube allocation TG-PHY150040 in the Comet system. The blue graph indicates Number of Jobs and the red line Service Units. Source: xsede.org portal.

Data Reprocessing – At the end of 2012 the IceCube Collaboration agreed to store the compressed SuperDST as part of the long-term archive of IceCube data. The decision taken was that this change would be implemented from the IC86-2011 run onwards. A server and a partition of the main tape library for input were dedicated to this data reprocessing task. Raw tapes are read to disk and the raw data files processed into SuperDST, which is saved in the data warehouse. The total number of files for seasons IC86-2011, IC86-2012 and IC86-2013 is 695,875; we currently have 26,223 remaining to be processed. The file breakdown per year is as follows: IC86-2011: 10,665 remaining in 63 tapes, IC86-2012: 6,948 remaining in 51 tapes. IC86-2013: 8,610 remaining in 28 tapes. Tape dumping procedures are being integrated with the copy of raw data to NERSC.

Personnel – Ian Saunders left WIPAC on June 2016 to work in Hitachi Data Systems. A new Linux system administrator position to work on the IceCube distributed computing infrastructure was posted and filled in by Vladimir Brik. A new Linux system administrator position to work on the UW-Madison cluster and VM infrastructure was posted and filled in by Alec Sheperd, who started working in the project on January 30th 2017.

Data Release

IceCube is committed to the goal of releasing data to the scientific community. The following links contain data sets produced by AMANDA/IceCube researchers along with a basic description. Due to challenging demands on event reconstruction, background rejection and systematic effects, data will be released after the main analyses are completed and results are published by the IceCube Collaboration.

Since summer 2016, thanks to UW-Madison subscribing to the EZID¹² service we have the capability of issuing persistent identifiers for datasets. These are Digital Object Identifiers (DOI) that follow the DataCite metadata standard¹³. We are in the process of rolling out a process for ensuring that all datasets made public by IceCube have a DOI and use the DataCite metadata standard capability to “link” it to the associated publication, whenever this is applicable. The use of DataCite DOIs to identify IceCube public datasets increases their visibility by making them discoverable in the search.datacite.org portal (see <https://search.datacite.org/works?resource-type-id=dataset&query=icecube>)

Datasets (last release on 15 Nov 2016): <http://icecube.wisc.edu/science/data>

The pages below contain information about the data that were collected and links to the data files.

1. A combined maximum-likelihood analysis of the astrophysical neutrino flux:
. <https://doi.org/10.21234/B4WC7T>
2. Search for point sources with first year of IC86 data:
. <https://doi.org/10.21234/B4159R>
3. Search for sterile neutrinos with one year of IceCube data:
. <http://icecube.wisc.edu/science/data/IC86-sterile-neutrino>
4. The 79-string IceCube search for dark matter:
. <http://icecube.wisc.edu/science/data/ic79-solar-wimp>
5. Observation of Astrophysical Neutrinos in Four Years of IceCube Data:
. <http://icecube.wisc.edu/science/data/HE-nu-2010-2014>
6. Astrophysical muon neutrino flux in the northern sky with 2 years of IceCube data:
. https://icecube.wisc.edu/science/data/HE_NuMu_diffuse
7. IceCube-59: Search for point sources using muon events:
. <https://icecube.wisc.edu/science/data/IC59-point-source>
8. Search for contained neutrino events at energies greater than 1 TeV in 2 years of data:
. http://icecube.wisc.edu/science/data/HEnu_above1tev
9. IceCube Oscillations: 3 years muon neutrino disappearance data:
. http://icecube.wisc.edu/science/data/nu_osc
10. Search for contained neutrino events at energies above 30 TeV in 2 years of data:
. <http://icecube.wisc.edu/science/data/HE-nu-2010-2012>
11. IceCube String 40 Data:
. <http://icecube.wisc.edu/science/data/ic40>
12. IceCube String 22–Solar WIMP Data:

¹² <http://ezid.cdlib.org>

¹³ <http://schema.datacite.org>

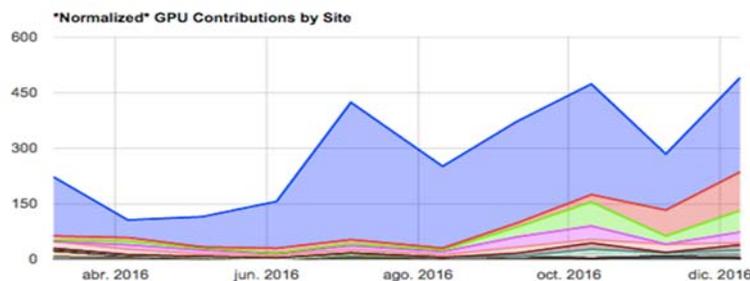
- . <http://icecube.wisc.edu/science/data/ic22-solar-wimp>
13. AMANDA 7 Year Data:
<http://icecube.wisc.edu/science/data/amanda>

Data Processing & Simulation Services

Offline Data Filtering – The data collection for the IC86-2016 season started on May 20 2016. A new compilation of data processing scripts had been previously validated and benchmarked with the data taken during the 24-hour test run using the new configuration. The differences with respect to the IC86-2015 season scripts are minimal. Therefore, we estimate that the resources required for the offline production will be about 750,000 CPU hours on the IceCube cluster at UW-Madison datacenter. 120TB of storage is required to store both the Pole-filtered input data and the output data resulting from the offline production. In the first three months we had to re-process five weeks of Level2 data due to database issues. The cause of this problem will be resolved for next season since we replaced the entire database structure. After this issue no problems have occurred and the data processing is proceeding smoothly. Level2 data are typically available one and a half weeks after data taking.

In preparation for a re-processing (pass2) of season 2010 - 2014, test data has been processed to L2 status (about 10% of the data). The test data is currently under review. Pass2 starts at sDST level in order to apply SPE corrections. This requires a L1 and L2 re-processing. We estimate we will require about 7,850,000 CPU hours and 380TB storage for a complete re-processing to L2. In order to ensure correct re-production of Level2 and Level3 data, additional metadata have been added for each set of run files that include, for instance, software versions for each processing step and personnel in charge.

Additional data validations have been added to detect data value issues and corruption. Replication of all the data at the DESY-Zeuthen collaborating institution is being done in a timely manner. Level3 data production is currently executed for three physics analysis groups. The Level3 data for the current season is usually available eight hours after Level2 completion. This short latency has been realized by fully automating Level3 and partly automating Level2 data processing. The cataloging and bookkeeping for Level3 data is the same as for Level2. The additional required resources to process the pass2 L2 data to L3 has been estimated to 4,000,000 CPU hours and 30TB storage.



Simulation production GPU usage over time normalized to a GTX680 GPU.

Simulation – The production of IC86 Monte Carlo simulations of the IC86-2012 detector configuration concluded in October of 2016. A new production of Monte Carlo simulations has

since begun with the IC86-2016 detector configuration. This configuration is representative of previous trigger and filter configurations from 2012, 2013, 2014 and 2015 as well as 2016. As with previous productions, direct generation of Level 2 background simulation data is used to reduce storage space requirements. However, as part of the new production plan, intermediate photon-propagated data will be stored on disk and reused for different detector configurations in order to reduce GPU requirements. This transition also includes a move to a new release of the IceCube simulation software release, IceSim 5. IceSim 5 contains improvements in memory and GPU utilization in addition to previous improvements to correlated noise generation, Earth modeling, and lepton propagation. All current simulations are now based on direct photon propagation using GPUs or a hybrid of GPU and splined photon PDF tables for high-energy events. Direct photon propagation is currently done on dedicated GPU hardware located at several IceCube Collaboration sites and through opportunistic grid computing where the number of such resources continues to grow.

The simulation production team is now organizing periodic workshops to explore better and more efficient ways of meeting the simulation needs of the analyzers. This includes both software improvements as well as new strategies and providing the tools to generate targeted simulations optimized for individual analyses instead of an one-size-fits-all approach.

The centralized production of Monte Carlo simulations has moved away from running separate instances of IceProd to a single central instance that relies on GlideIns running at satellite sites. Production has been transitioning to a newly redesigned simulation scheduling system IceProd2. Production on IceProd2 has began ramping up after some of the major bugs have been identified and fixed. We expect a full transition by the Spring 2017 Collaboration Meeting. This will include a new set of monitoring tools in order to keep track of efficiency and further optimizations.

IceCube Software Coordination

The software systems spanning the IceCube Neutrino Observatory, from embedded data acquisition code to high-level scientific data processing, benefit from concerted efforts to manage their complexity. In addition to providing comprehensive guidance for the development and maintenance of the software, the IceCube Software Coordinator, Alex Olivas, works in conjunction with the IceCube Coordination Committee, the IceCube Maintenance and Operations Leads, the Analysis Coordinator, and the Working Group Leads to respond to current operational and analysis needs and to plan for anticipated evolution of the IceCube software systems. In the last year, software working group leads have been appointed to the following groups: core software, simulation, reconstruction, science support, and infrastructure. Continuing efforts are underway to ensure the software group is optimizing in-kind contributions to the development and maintenance of IceCube's physics software stack.

The IceCube collaboration contributes software development labor via the biannual MoU updates. Software code sprints are organized monthly with the software developers to tackle topical issues. Progress is tracked, among other means, by tracking open software tickets tied to monthly milestones. An overview graphic of milestones from PY1 is shown in Figure 9.

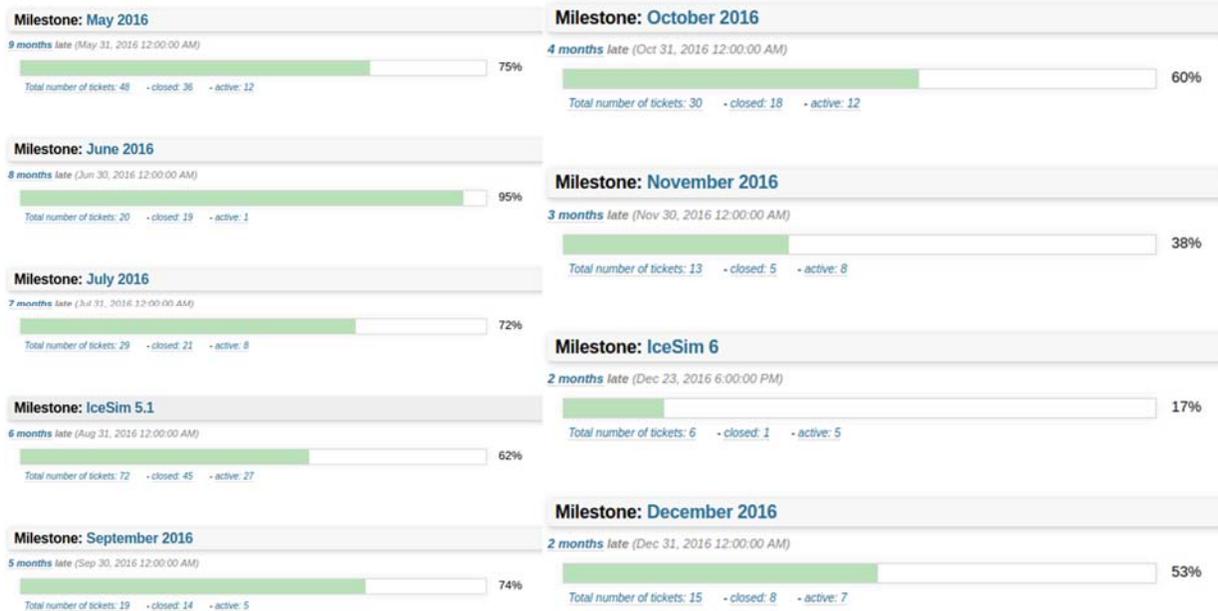


Figure 9: IceCube software milestones and tickets for 2016. 85 remaining open tickets have been reassigned to future milestones.

Calibration

The 4–6% DOM-by-DOM gain corrections introduced in IC86–2015 were updated for the IC86–2016 physics run and are working well in the online PnF system. Because this shift is observable in some physics analyses, we are investigating a “Pass 2” reprocessing of data taken before 2015 to add these corrections. The gain corrections have been extracted from historical minimum bias data, and a 10% data sample from each year since 2011 has been reprocessed and is being investigated by the working groups. The data from past years shows the stability of the position of the average SPE peak is within 1%, which is also reported in the IceCube Instrumentation and Online Systems Paper (<https://arxiv.org/abs/1612.05093>) which was completed and submitted for publication in December 2016. Figure 10 shows the distribution of the mean SPE peak position for each DOM for the FADC and ATWD digitizers.

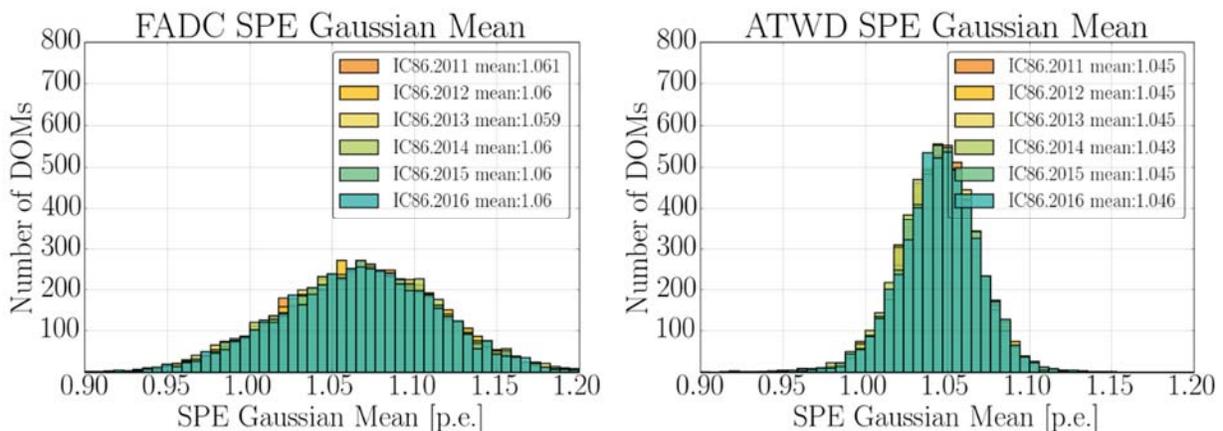


Figure 10: Histogram of the average SPE peak position for each DOM from 2011 to 2016 for the FADC digitizer (left) and the ATWD digitizer (right).

We continue to refine our model of the optical properties of the ice. A new set of LED flasher data was collected in January 2017 using single-LED low brightness data in the DeepCore strings which have the closest horizontal spacing available in IceCube (about 40 m). This data will be used to investigate the scattering in the ice within a couple of scattering lengths, as opposed to the all-detector data used in our previously published ice models, where the typical horizontal spacing is 120 m. A new run with the Swedish camera in the center of IceCube confirmed the continued presence of the bubble column in the refrozen hole ice; the Swedish camera hardware has begun to degrade noticeably, but models indicate that the bubble column should persist beyond the lifetime of IceCube.

Program Management

Management & Administration – The primary management and administration effort is to ensure that tasks are properly defined and assigned, that the resources needed to perform each task are available when needed, and that resource efficiency is tracked to accomplish the task requirements and achieve IceCube’s scientific objectives. Efforts include:

- A complete re-baseline of the IceCube M&O Work Breakdown Structure to reflect the structure of the principal resource coordination entity, the IceCube Coordination Committee.
- The FY2016-FY2021 M&O Plan was submitted in June 2016.
- The detailed M&O Memorandum of Understanding (MoU) addressing responsibilities of each collaborating institution was revised for the collaboration meeting in Stony Brook, April 16-22, 2016.

IceCube M&O – PY1 (FY2016/2017) Milestones Status:

Milestone	Month
Revise the Institutional Memorandum of Understanding (MOU v20.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration Meeting	April 2016
Report on Scientific Results at the Spring Collaboration Meeting	April 16-22, 2016
Submit for NSF approval, a revised IceCube Maintenance and Operations Plan (M&OP) and send the approved plan to non-U.S. IOFG members.	June 2016
Revise the Institutional Memorandum of Understanding (MOU v21.0) - Statement of Work and Ph.D. Authors head count for the fall collaboration meeting	September 2016
Report on Scientific Results at the Fall Collaboration Meeting	Sept. 26-30, 2016
Submit for NSF approval a mid-year report which describes progress made and work accomplished based on objectives and milestones in the approved annual M&O Plan.	October 2016
Revise the Institutional Memorandum of Understanding (MOU v22.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration meeting	April 2017

Engineering, Science & Technical Support – Ongoing support for the IceCube detector continues with the maintenance and operation of the South Pole Systems, the South Pole Test System, and the Cable Test System. The latter two systems are located at the University of Wisconsin–Madison and enable the development of new detector functionality as well as investigations into various operational issues, such as communication disruptions and

electromagnetic interference. Technical support provides for coordination, communication, and assessment of impacts of activities carried out by external groups engaged in experiments or potential experiments at the South Pole. The IceCube detector performance continues to improve as we restore individual DOMs to the array at a faster rate than problem DOMs are removed during normal operations.

Education & Outreach (E&O) – The IceCube Collaboration continues to make progress on E&O efforts organized around four main themes:

- 1) Reaching motivated high school students and teachers through IceCube Masterclasses
- 2) Providing intensive research experiences for teachers (in collaboration with PolarTREC) and for undergraduate students (NSF science grants, International Research Experience for Students (IRES), and Research Experiences for Undergraduates (REU) funding)
- 3) Engaging the public through web and print resources, graphic design, webcasts with IceCube staff at the Pole, and displays
- 4) Developing and implementing semiannual communication skills workshops held in conjunction with IceCube Collaboration meetings

New initiatives underway include citizen science and immersive and interactive learning projects to engage the public in STEM research and related activities. The citizen science project, encouraged by the M&O review panel to increase the reach of IceCube's E&O activities, will boost public participation while advancing scientific research. This project involves Zooniverse activities that will allow amateur scientists to find interesting neutrino events that aren't identified with current computer algorithms. Activities developed by 2016 summer students at UW–Madison and UW–River Falls were advanced by students of the WIPAC fall internship. The project is being refined and will be made public in the coming months. Work is also progressing on the NSF Advancing Informal STEM Learning funded project, a joint effort between WIPAC and the Wisconsin Institutes for Discovery to create interactive and immersive learning programs based on touch tables and virtual reality devices. The groups will compare game-based learning strategies for different audiences using these technologies and further develop current IceCube virtual reality experiences for larger audiences. This two-year project is a pilot grant for larger projects involving research in Antarctica.

The IceCube Masterclasses continue to be promoted within the collaboration, as well as externally at national and international levels. The 2017 edition, on March 8, 12 and 22, includes 15 institutions. The IceCube cosmic ray masterclass was held in conjunction with a special lecture series, Recent Discoveries in Particle and Astroparticle Physics, at Sungkyunkwan University at Suwon, Korea, on August 18, 2016. The IceCube Masterclasses were also described in a talk at the Scientific Committee on Antarctic Research meeting in Kuala Lumpur, Malaysia, on August 22, 2016, and at the April APS meeting in Washington, D.C., January 30, 2017. A workshop featuring the IceCube Masterclasses has been accepted and will be offered at the 2017 summer meeting of the American Association of Physics Teachers in Cincinnati, Ohio.

Physics teacher Kate Miller, from Washington-Lee High School in Arlington, VA, had a highly successful South Pole deployment in January 2017 with IceCube through the PolarTREC program. Prior to deployment, Kate worked with former IceCube PolarTREC teacher Liz Ratliff

and AMANDA TEA teacher Eric Muhs to develop and deliver a nine-day math-science enrichment course that included IceCube research for the UW–River Falls Upward Bound program (July 11-22, 2016). IceCube PolarTREC teacher Armando Caussade, 2014-15 season, published a book on his experiences with IceCube, *A Puerto Rican in the South Pole*.

Five undergraduate students (two women, one of whom is African American) worked at Stockholm University for ten weeks in the summer of 2016, supported by NSF IRES funding. The students are from UWRF, UW–Madison, and the University of Minnesota. The UWRF astrophysics NSF REU program supported six students (two men, one of whom is African American, four students total from two-year colleges) from over 60 applicants for ten-week summer 2016 research experiences, including attending the WIPAC IceCube software and science boot camp. Multiple IceCube institutions also supported undergraduates. Two of the students from the 2015 UWRF IRES program attended the IceCube Collaboration meeting in April, 2016, at Stony Brook, NY. Applications for the 2017 IRES and REU programs are under review.

The IceCube scale model, with one colored LED for each of the 5,160 DOMs and audio mapping of the data, was seen by thousands of people at the World Science Festival in New York City on June 5, 2016. It was also featured in a New York Times Science Facebook Live event. More transportable 1m x 1m x 1m models are under construction at York University with art/science professor Mark-David Hosale as well as at the University of Maryland and at Stockholm and Drexel Universities. The model at Drexel University was completed as a part of their first high school internship program.

During the 2016-17 polar season, IceCube partnered with the Hergé Foundation to use the Tintin character to promote IceCube science to followers of this comics series. Among readers and followers of the IceCube and Tintin websites and social networks, an audience of over 5,000 people was reached.

The E&O team continues to work closely with the communication team. Science news summaries of IceCube publications, written at a level accessible to science-literate but nonexpert audiences, are produced and posted regularly on the IceCube website and highlighted on social media. Fifteen research news articles were published in the review period along with seven project updates. Results of the IceCube sterile neutrino search and neutrino oscillation studies were featured in press releases at a few IceCube institutions.

A new series of videos and multimedia resources was launched, including a video about the sterile neutrino results and latest searches for cosmogenic neutrinos. The videos were watched over 14,000 times during the last few months. Media mentions, tracked since January 2016, include over 250 news pieces appearing in more than 20 different countries, with over 60% of those in the US. In addition to national and international coverage in the media, we have documented local news mentions in twelve different U.S. states and in Puerto Rico.

A new IceCube video is currently in production, to be shown at the APS March Meeting (about 10,000 participants) and other venues. This five-minute video highlights the success of the IceCube construction and summarizes the main contributions of IceCube to neutrino astronomy.

The communication training program targeting PhD students and postdocs, launched at the spring 2015 IceCube Collaboration meeting in Madison, continues to be held during the biannual

collaboration meetings. The spring training in Stony Brook featured an interactive presentation from The Story Collider group focusing on using narratives to explain research, connecting to audiences in meaningful ways, and engaging in cultural conversations about science. At Mainz University in September, participants had podcast training in which they thought about storytelling as a way to share their research and learned about the equipment needed to produce podcasts.

Finally, the IceCube Diversity and Inclusion task force was launched in September 2016. This new initiative is hosted by the E&O and communications teams and includes several other researchers and staff. They have volunteered their time to lead and promote activities to increase diversity within the IceCube Collaboration, including fostering a more inclusive and supportive work environment for all and contributing to the advancement of a more diverse student and workforce population in STEM fields, especially in the US. WIPAC has partnered with the AAAS Community Engagement Fellows program and will host a Diversity, Inclusion and Career Mentoring Fellow during 2017. This fellow, who is cofunded by UW–Madison and AAAS (50%), will create diversity metrics, share best practices with the IceCube community, and promote mentoring training and networking opportunities for IceCube researchers targeting key milestones in their careers.

Section III – Project Governance and Upcoming Events

The detailed M&O institutional responsibilities and Ph.D. author head count is revised twice a year at the time of the IceCube Collaboration meetings. This is formally approved as part of the institutional Memorandum of Understanding (MoU) documentation. The MoU was last revised in September 2016 for the Fall collaboration meeting in Mainz, Germany (v21.0), and the next revision (v22.0) will be posted in April 2017 at the Spring collaboration meeting in Madison, WI.

IceCube Collaborating Institutions

Following the April 2016 Spring collaboration meeting, Universität Münster with Dr. Alexander Kappes as the institutional lead, and SNOLAB with Dr. Ken Clark as the institutional lead, were approved as full members of the IceCube Collaboration. University of Toronto left IceCube after Dr. Ken Clark moved from Toronto to SNOLAB.

After the September 2016 Fall collaboration meeting, the University of Texas at Arlington with Dr. Benjamin Jones as the institutional lead joined the IceCube Collaboration, and the University of Mons left IceCube.

As of January 2017, the IceCube Collaboration consists of 48 institutions in 12 countries (25 U.S. and Canada, 19 Europe and 4 Asia Pacific).

The list of current IceCube collaborating institutions can be found on:

<http://icecube.wisc.edu/collaboration/collaborators>

IceCube Major Meetings and Events

IceCube Spring Collaboration Meeting – Madison, WI
 IceCube Fall Collaboration Meeting – Berlin, Germany

May 1–5, 2017
 October 2–7, 2017

Acronym List

CnV	Calibration and Verification
CPU	Central Processing Unit
CVMFS	CernVM-Filesystem
DAQ	Data Acquisition System
DOM	Digital Optical Module
E&O	Education and Outreach
I3Moni	IceCube Run Monitoring system
IceCube Live	The system that integrates control of all of the detector’s critical subsystems; also “I3Live”
IceTray	IceCube core analysis software framework, part of the IceCube core software library
MoU	Memorandum of Understanding between UW–Madison and all collaborating institutions
PMT	Photomultiplier Tube
PnF	Processing and Filtering
SNDAQ	Supernova Data Acquisition System
SPE	Single photoelectron
SPS	South Pole System
SuperDST	Super Data Storage and Transfer, a highly compressed IceCube data format
TDRSS	Tracking and Data Relay Satellite System, a network of communications satellites
TFT Board	Trigger Filter and Transmit Board
WIPAC	Wisconsin IceCube Particle Astrophysics Center