IceCube Maintenance and Operations
Fiscal Year 2018 PY3 Mid-Year Report

April 1, 2018 – September 30, 2018

Submittal Date: October 31, 2018

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 FY2018 (PY3) Mid-Year Report is submitted as required by the NSF Cooperative Agreement PLR-1600823. This report covers the 9-month period beginning January 1, 2018 and concluding September 30, 2018. The status information provided in the report covers actual common fund contributions received through September 30, 2018 and the full 86-string IceCube detector (IC86) performance through September 30, 2018.

Table of Contents

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 PY3 installment of $5,250,000 was released to UW–Madison to cover the costs of maintenance and operations during the first three quarters of PY3 (FY2018/FY2019): $798,525 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 $4,451,475 was directed to the IceCube M&O Core account. The second PY3 installment of $1,750,000 is expected to be released to UW–Madison to cover the costs of maintenance and operations during the second half of PY3 (FY2019): $266,175 will be directed to the U.S. Common Fund account, and the remaining $1,483,825 will be directed to the IceCube M&O Core account (Table 1). An additional $291,712 FY2019 funding was awarded to support an IceCube M&O supplemental proposal to further develop talent for STEM research, increase diversity and inclusion in science and related fields, and connect research-based skills and technology to societal challenges.

PY3: FY2018 / FY2019	Funds Awarded to UW for Apr 1, 2018 – March 31, 2019	Funds to Be Awarded to UW for Apr 1, 2018 – March 31, 2019
IceCube M&O Core account	$4,451,475	$1,483,825
U.S. Common Fund account	$798,525	$266,175
TOTAL NSF Funds	**$5,250,000**	**$1,750,000**

Table 1: NSF IceCube M&O Funds – PY3 (FY2018 / FY2019)

Of the IceCube M&O PY3 (FY2018/2019) Core funds, $968,092 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 FY2018/2019 funds for the subawardee institutions.

Institution	Major Responsibilities	Funds
Lawrence Berkeley National Laboratory	DAQ maintenance, computing infrastructure	$91,212
Pennsylvania State University	Computing and data management, simulation production, DAQ maintenance	$70,847
University of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance	$149,265
University of Maryland at College Park	IceTray software framework, online filter, simulation software	$603,369
University of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools	$24,592
Michigan State University	Simulation software, simulation production	$28,807
South Dakota School of Mines and Technology (added in July 2017)	Simulation production and reconstruction	$00.00
Total		**$968,092**

Table 2: IceCube M&O Subawardee Institutions – PY3 (FY2018/2019) 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 six months into the Year 3 of the award, and shows an estimated balance at the end of PY3.

Total awarded funds to the University of Wisconsin (UW) for supporting IceCube M&O from the beginning of PY1 through part of PY3 are $19,318K. With the last PY3 planned installment of $1,750K, the total PY1-3 budget is $21,068K. Total actual cost as of September 30, 2018 is $17,201K and open commitments against purchase orders and subaward agreements are $1,054K. The current balance as of September 30, 2018 is $2,813K. With a projection of $3,070K for the remaining expenses during the final six months of PY3, the estimated negative balance at the end of PY3 is -$257K, which is 1.2% of the PY3 budget (Table 3).

(a)	**(b)**	**(c)**	**(d)= a - b - c**	**(e)**	**(f) = d – e**
YEARS 1-3 Budget Apr.'16-Mar.'19	**Actual Cost To Date** through Sept 30, 2018	**Open Commitments** on Sept 30, 2018	**Current Balance** on Sept 30, 2018	**Remaining Projected Expenses** through Mar. 2019	**End of PY3 Forecast Balance** on Mar. 31, 2019
$21,068K	**$17,201K**	**$1,054K**	**$2,813K**	**$3,070K**	**-$257K**

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, 2018–March 31, 2019, based on v24.0 of the IceCube Institutional Memorandum of Understanding, from May 2018.

	Ph.D. Authors	**Planned Contribution**	**Actual Received**
Total Common Funds	**138**	**$1,883,700**	**$1,329,693**
U.S. Contribution	71	$969,150	$969,150
Non-U.S. Contribution	67	$914,550	$360,543*

Table 4: Planned and Actual CF Contributions for the period of April 1, 2017–March 31, 2018

* The non-U.S. invoicing and contributions are still underway, and it is anticipated that all planned contributions will be fulfilled.

Section II – Maintenance and Operations Status and Performance

Detector Operations and Maintenance

Detector Performance — During the period from February 1 to September 30, 2018, the detector uptime, defined as the fraction of the total time that some portion of IceCube was taking data, was 99.79%, exceeding our target of 99% and close to the maximum possible, given our current data acquisition system. The clean uptime for this period, indicating full-detector analysis-ready data, was 98.21%, exceeding our target of 95%. Historical total and clean uptimes of the detector are shown in Figure 1.

Figure 2 shows a breakdown of the detector time usage over the reporting period. The partial-detector good uptime was 0.9% of the total and includes analysis-ready data with fewer than all 86 strings. Excluded uptime includes maintenance, commissioning, and verification data and required 0.68% of detector time. The unexpected detector downtime was limited to 0.21%.

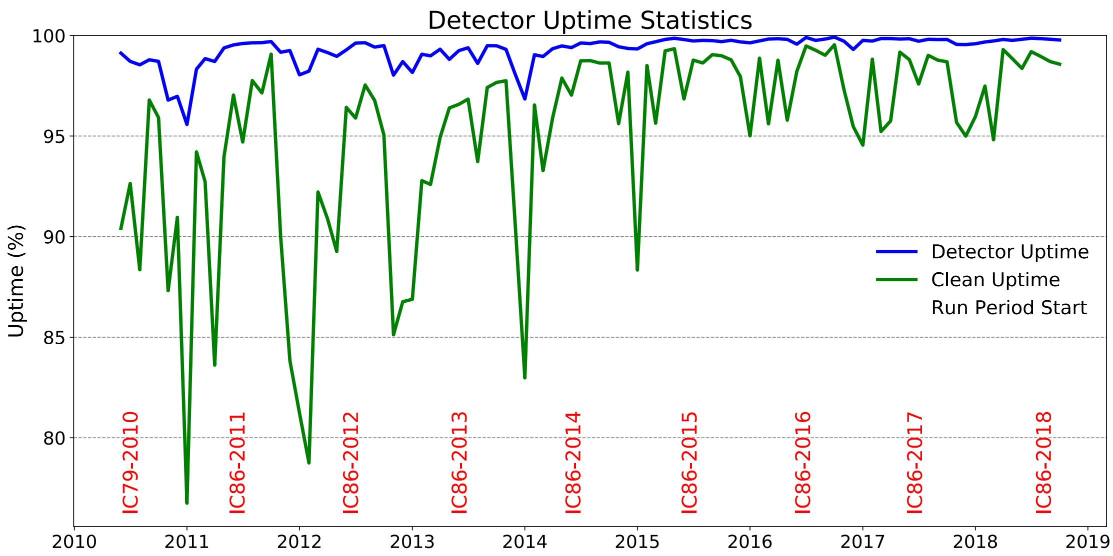

Figure 1: Total IceCube Detector Uptime and Clean Uptime

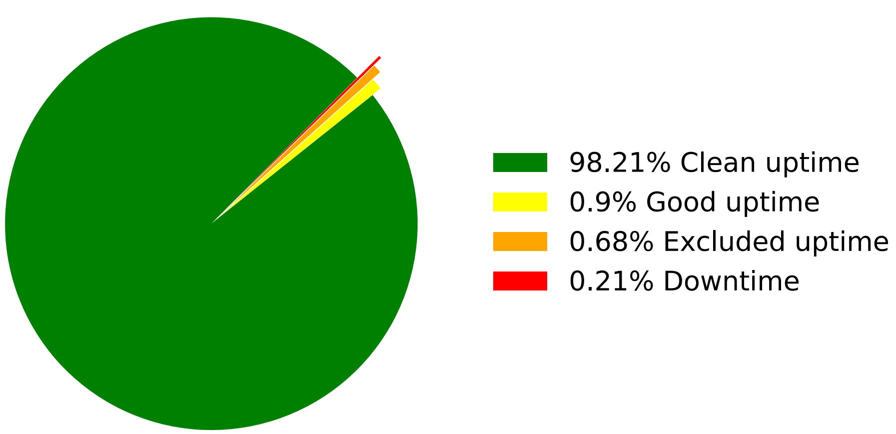

Figure 2: Cumulative IceCube Detector Time Usage, February 1 – September 30, 2018

Hardware Stability — No DOMs have failed during this reporting period. However, the DOM mainboard in one of the 2015–16 prototype scintillator panels failed in August 2018. As the failure mode is similar to a few of the failed in-ice DOMs, we plan to recover the hardware during the 2018–19 austral summer season for diagnosis; this may inform future DOM designs. The total number of active DOMs in the data stream is currently 5406 (98.5% of deployed DOMs), plus three scintillator panels and the IceACT trigger mainboard.

The failure rate in the commercial Acopian power supplies that supply the DC voltage to the DOMs from the DOMHubs increased starting in late 2015. During the 2016–17 austral summer season, we performed a complete replacement of the supplies using an updated model. The failure rate has stabilized; however, it does not appear to be decreasing back to the previous stable levels. During the 2017–18 season, we installed 32 alternate DC power supplies (Mean Well MSP-200-48). So far, none of the alternate supplies have failed, indicating that the failure rate is less than or equal to the Acopian units (Fig. 3). We plan to install 72 more Mean Well units this season.

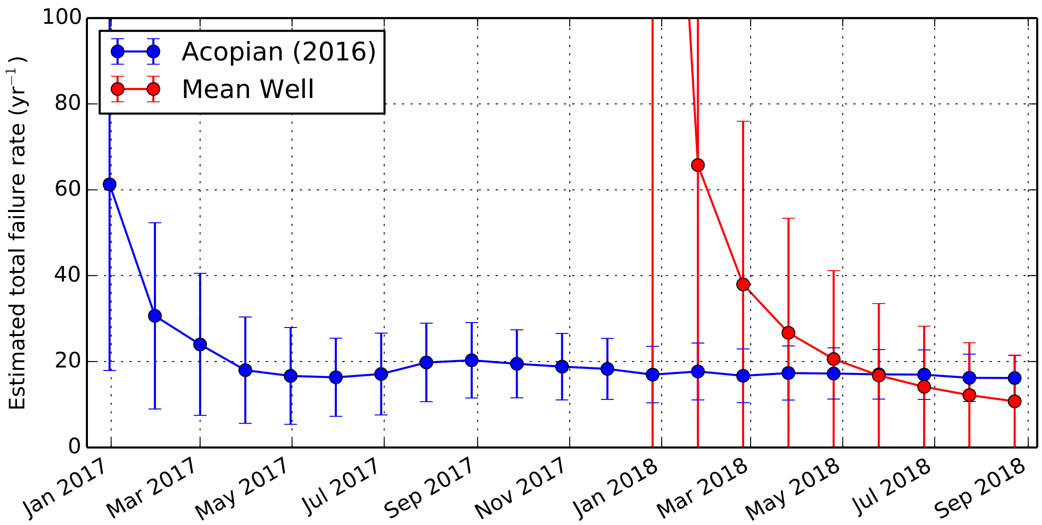

Figure 3: Estimated full-detector failure rate of the two types of DOM power supplies in use. No Mean Well units have failed as of September 2018.

The primary GPS receiver that serves at IceCube's master clock (Symmetricom ET6000) lost satellite lock in early 2018. After a power-cycle, the unit re-locked but intermittently reported degraded synchronization quality. Switching the detector to the hot spare backup clock prevented any significant impact to data-taking. As the existing clocks are unsupported and have unfixed firmware bugs, we are in the process of testing a replacement master clock. All qualification tests so far have been successful, and a transition of the detector to the new clock is planned for 2019.

IC86 Physics Runs — The eighth season of the 86-string physics run, IC86–2018, began on July 10, 2018. Detector settings were updated using the latest yearly DOM calibrations from March 2018. Changes to the trigger and filter settings were minimal. Support for an external trigger from

IceACT was released but not enabled, due to hardware issues with the telescope (see “Surface Detectors”). Due to improvements in handling bad DOMs in the readout software, we were able to recommission four DOMs with a failed partner DOM on the same wire pair, increasing the active DOM count from 5402 to 5406 (out of a total of 5484).

For this reporting period, the average TDRSS daily transfer rate was approximately 70 GB/day. The bandwidth saved from filter optimizations has allowed us to reserve more for HitSpool and calibration data transfer. We are implementing automated HitSpool data captures triggered from the northern hemisphere by LIGO gravitational wave alerts, enabling searches for sub-threshold coincident neutrinos.

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:Tyranena release in April 2018 and subsequent patch releases, which improve start-of-run behavior and fix a bug in the Simple Multiplicity Trigger affecting only the IceACT trigger.

- Finalization of the pDAQ:Urban_Harvest DAQ release planned for October 2018, which improves DAQ–IceCube Live interaction and DOM memory overflow reporting.

- Work towards the 2019 releases, which will address lingering component stalls at run starts and optimize the trigger and hit sorting components for multicore architectures.

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 filter changes to support the IC86–2018 physics run start, a new version of the realtime event processing system was released that streamlines event selection and supports additional types of neutrino alerts to a network of multi-messenger observatories.

Detector Monitoring and Experiment Control — Development of IceCube Live, the experiment control and monitoring system is transitioning to a maintenance phase after several major feature releases in 2017. This reporting period has seen one major release with the following highlighted features:

- Live v3.2 (May 2018): Highlights include an improved “billboard” status page (Fig. 4), a faster run history page, and automatic run configuration failover in the case of Acopian power supply failures.

No downtime was detected for the I3Live experiment control system during the reporting period, resulting in an uptime greater than 99.999%.

Figure 4: Improved IceCube Live detector status page.

Supernova System — The supernova data acquisition system (SNDAQ) found that 99.78% of the available data from February 1 to September 30, 2018 met the minimum analysis criteria for run duration and data quality for sending triggers. An additional 0.02% 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.

Efforts to include a data-driven trigger in SNDAQ that is independent of an assumed signal shape are under way. Release progress has been slowed by turnover in Ph.D. students, as SNDAQ has typically not been developed and maintained by M&O software professionals. This model is under re-evaluation.

Surface Detectors — Snow accumulation on the IceTop tanks continues to increase the energy threshold and thus reduces the trigger rate of the surface array by ~10%/year. 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. After a prototype installation in 2015–16, we redesigned the panels and electronics and deployed two new scintillator stations in the 2017–18 season, consisting of seven panels each. Panels can be elevated above the snow surface and subsequently raised (Fig. 5). We are analyzing the data from both stations in order to finalize the hardware design for a complete array. Air shower events coincident with IceTop have been identified, enabling initial cross-calibration efforts (Fig. 6).

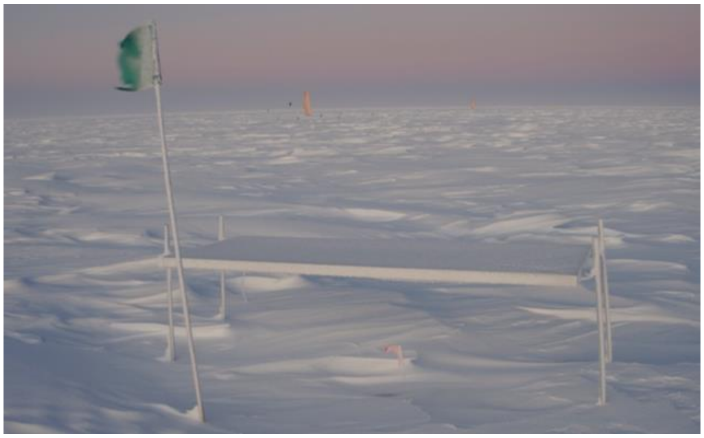

Figure 5: Elevated scintillator panel in operation after the 2018 winter.

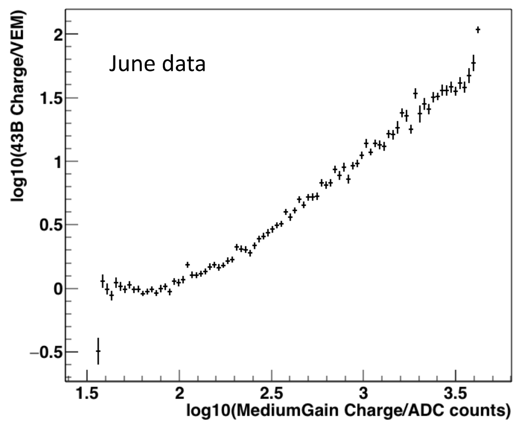

Figure 6: Air shower charge in scintillator panel (x-axis) vs. IceTop Tank 43B (y-axis)

During the 2017–18 season, an upgraded air Cherenkov telescope (IceACT) was installed on the roof of the ICL. The new telescope was commissioned successfully, with plans to trigger IceCube/IceTop readout starting with the IC86-2018 physics run. However, the DAQ system of the telescope failed one week after the start of the winter data-taking season. The telescope was moved into the antenna doghouse but could not be repaired by the winterovers. Modifications to the IceACT DAQ and a lens heating system to address snow accumulation are planned.

The central "FieldHub" of the scintillator stations will be enhanced in the 2018–19 season with a second IceACT telescope, using a stand design validated over the 2018 winter. Additionally, two prototype radio antennas will be connected to the scintillator DAQ electronics, enabling feasibility studies for a cosmic ray radio air shower component of the scintillator array. We are in the process of developing a comprehensive, unified plan regarding all proposed IceTop surface enhancements.

***Personnel* —** Robert Snihur joined WIPAC in April 2018 and is responsible for offline Level2 data processing.

Computing and Data Management Services

South Pole System – This time period started with the hiring process of our WinterOver detector operators for the upcoming season. Over 50 applications were screened and 15 phone interviews conducted. Five candidates were selected for in-person interviews and job offers were made to the two strongest candidates. Two more candidates were selected as a backup (these candidates were also referred to the South Pole Telescope). After successfully PQ-ing both primary candidates training in Madison started on August 1st and continues through mid October when both operators flew to Denver for organized team activities with ASC.

The major activity for the 2018-19 pole season is a partial refresh of the server hardware at the South Pole and the South Pole Test System. The last refresh finished in 2014-15. Several new server platforms were evaluated for performance and maintainability. The Dell R740 platform was selected since it provided the best maintainability and configurability needed at a reasonable price point. Approximately half of the servers at the South Pole will be replaced in the 2018-19 season.

The remaining servers will be replaced in the 2019-20 season. We purchased (40) servers for 2018-19; (19) of those were deployed into the test system, (21) were integrated into the computing cluster for burn in.

Several server configurations were purchased. These configurations accommodate the current needs of the system and also offer some decent performance improvements. This includes machines with more storage capacity to replace fpmaster and one of the JADE data archival servers, with the rest replacing the fpslave machines. The fpslave machines, which handle the majority of the data processing workload, will all be replaced as there is benefit to keeping that computing pool homogenous.

Testing and burn in were successfully concluded in September of 2018. 21 servers were packed up and entered into the NSF's cargo stream for delivery to the South Pole. The replacement will begin in November and is expected to have minimal impact on detector uptime. The previous generation of hardware will remain on station as a fallback in case unexpected problems arise. Those servers will be returned to UW-Madison in the 2019-20 season.

As part of yearly maintenance all systems (old and new) will be updated to the newest firmware and Operating System security patches.

To ensure uninterrupted data taking for the next year we ordered, tested and shipped over 4,000lbs of cargo to South Pole via NSF's cargo stream. This includes the servers to refresh ~50% of the server systems, and miscellaneous spare parts used during the current winter (hard drives, power supplies, etc.).

To ensure uninterrupted data taking for the next year we ordered, tested and shipped over 4,000lbs of cargo to South Pole via NSF's cargo stream. This includes the servers to refresh ~50% of the server systems, and miscellaneous spare parts used during the current winter (hard drives, power supplies, etc.).

Data Transfer – Data transfer has performed nominally over the past ten months. Between January 2018 and September 2018 a total of 20.7 TiB of data were transferred from the South Pole to UW-Madison via TDRSS, at an average rate of 77.5 GiB/day. Figure 1 shows the daily satellite transfer rate and weekly average satellite transfer rate in GB/day through August 2017. The IC86 filtered physics data are responsible for 95% of the bandwidth usage.

Since September 2016 the JADE software handles all the IceCube data flows: disk archive at the South Pole, satellite transfer to UW-Madison and long term archive to tape libraries at NERSC and DESY. JADE continues to operate smoothly and has been an effective tool for handling a variety of our routine data movement workflows. This has been confirmed over the last year with experience from both the Winterovers and IT staff at UW-Madison.

Between January 2018 and September 2018 a total of 20.7 TiB of data were transferred from the South Pole to UW-Madison via TDRSS, at an average rate of 77.5 GiB/day.

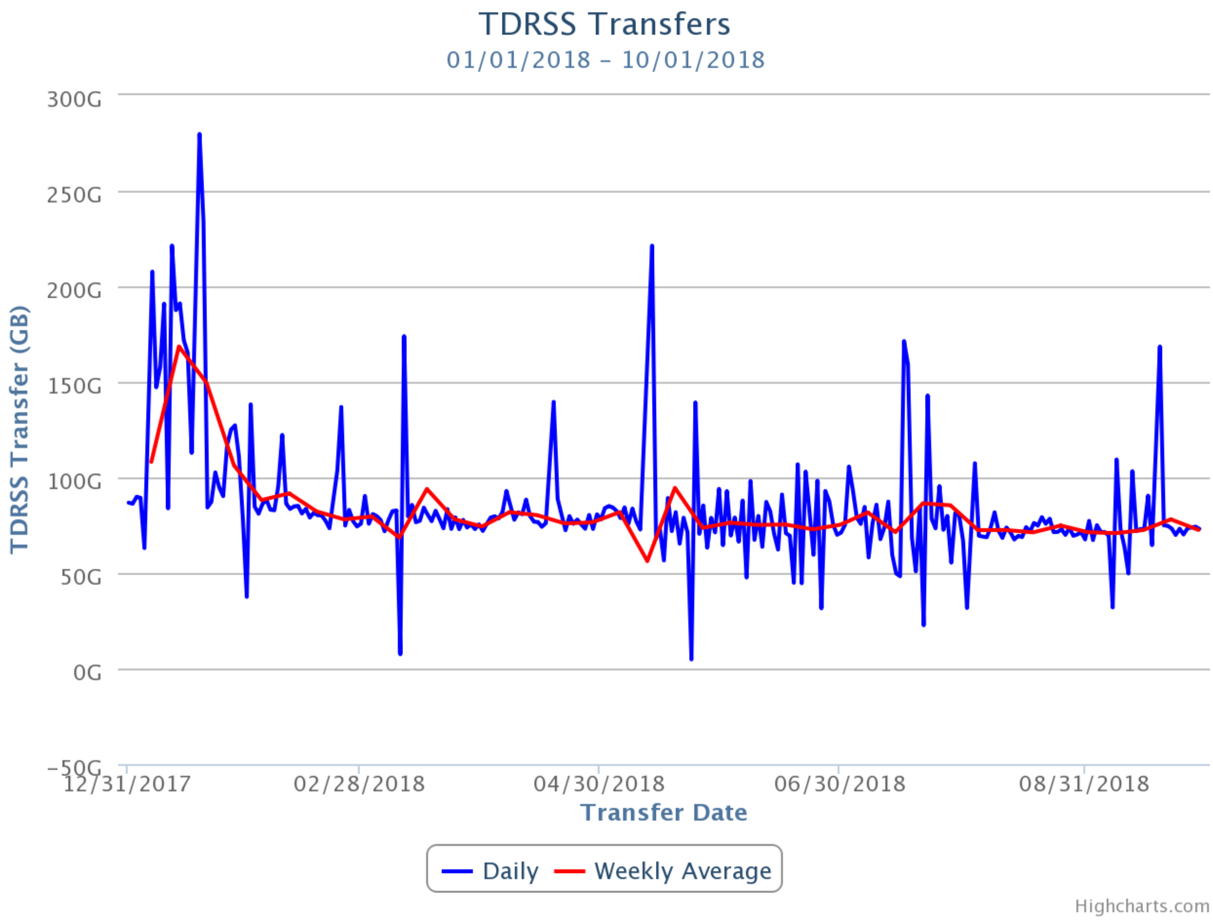

Figure 1: TDRSS Data Transfer Rates, January 1, 2018–September 30, 2018. The daily transferred volumes are shown in blue and, superimposed in red, the weekly average daily rates are also displayed.

Data Archive – The IceCube raw data is archived by writing two copies on independent hard disks. During the reporting period (January 2018 to September 2018) a total of 315 TiB of unique data were archived to disk averaging 1.2 TiB/day.

In May 2018, the set of archival disks containing the raw data taken by IceCube during 2017 was received at UW-Madison. These disks are processed using JADE which now indexes the metadata, bundles the data files into chunks suitable for storage in tape libraries, and replicates the data to the long-term archives at DESY and NERSC.

Beginning in September 2016 we have been regularly transferring archival data to NERSC. At this time, the total volume of data archived at NERSC is 2.8 PB. Figure 7 shows the rate at which data has been archived to NERSC since the start of this service. The plan is to keep this archive stream constantly active while working on further JADE functionality that will allow us to steadily increase the performance and automation of this long term archive data flow.

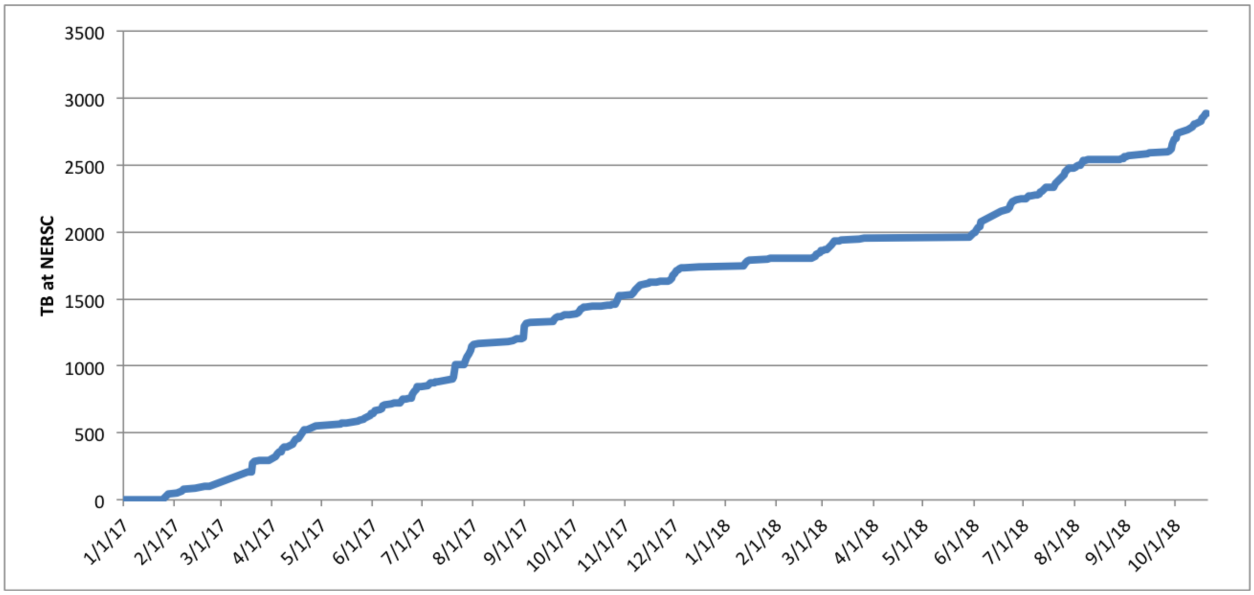

Figure 2: Volume of IceCube data archived at the NERSC tape facility by the JADE Long Term Archive service as a function of time.

Computing Infrastructure at UW-Madison

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 activity since the Collaboration decided to use GPUs for the photon propagation part of the simulation chain in 2012. Direct photon propagation provides the precision required, and is very well suited to GPU hardware, running about 100 times faster than in CPUs. In addition, we have begun the task of decommissioning the data center facility at our 222 West Washington Avenue location, and relocating the equipment to a commercial facility leased by UW-Madison.

The new facility is located approximately seven miles from our 222 West Washington Avenue location. It is a commercial facility offering redundant battery-backed generator power and biometric physical security. UW-Madison leases space, power, and cooling in this facility to support internal needs as well as project needs for the campus as a whole. As such, they have extended the campus high speed networks to the facility which offers us easier integration with the rest of our campus systems without the higher costs typically associated with commercial rates for network transit. In particular, when the project is complete, we will have a full 100 Gb/s network path between this facility and our other facilities on campus.

To facilitate the data center relocation, we purchased and deployed a storage system with 9.6 PB raw capacity (7.2 PB usable) in the new facility. This provided sufficient usable capacity to enable us to consolidate our filesystem infrastructure as well as simplifying the effort of migrating data to the new location. The new equipment was deployed and available in May of 2018. Filesystem configuration, testing, and validation completed at the end of June, and the data transfer began in July. By mid-August, the transfer of 6.0 PB had completed and this now serves as the primary storage for experimental, simulation, and analysis data.

The total amount of data stored on disk in the data warehouse at UW-Madison is 6.5 PB: 2.2 PB for experimental data, 4.0 PB for simulation and analysis and 285 TB for user data.

While storage comprises the largest consumer of physical space, there are a number of other systems which will also be migrated to the new facility. The work to do this will be ongoing and is expected to be complete in the first quarter of 2019. As data and services are migrated to the new facility, hardware that is still under warranty will be relocated. Older hardware that is still usable but out of support will be re-purposed for testing or into the compute cluster as needed. Any old hardware not needed for these purposes will be disposed of.

The other focus has been the continued expansion of the GPU cluster is to provide the capacity to meet the Collaboration direct photon propagation simulation needs. This included the update of 64 GPUs in the UW-Madison GPU cluster to nVidia GTX-1080 cards in March of 2018. Still, the GPU needs have been estimated to be higher than the capacity of the GPU cluster at UW-Madison. Additional GPU resources at several IceCube sites, plus specific supercomputer allocations, allow us to try and reach that required capacity.

Distributed Computing - In March 2016, a new procedure to formally gather computing pledges from collaborating institutions was started. This data is 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 1 shows the computing pledges per institution as of September 2018:

Site	Pledged CPUs	Pledged GPUs
Aachen	27700*	44*
Alabama		6
Alberta	1400	178
Brussels	1000	14
Chiba	196	6
Delaware	272	
DESY-ZN	1400	180
Dortmund	1300*	40*
LBNL	114	
Mainz	1000	300
Marquette	96	16
MSU	500	8
NBI		10
Penn State	3200*	101*
Queen’s		55
Uppsala	10	
UMD	350	112
UTA	50	
UW-Madison	7000	440

Wuppertal	300	
TOTAL (exclusive)	13688	1325
TOTAL (all)	45888	1510

Table 1: Computing pledges from IceCube Collaboration institutions as of September 2018.
*** indicates maximum shared resources, not exclusively for IceCube.**

We are implementing 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 enlarge the distributed infrastructure and make it more efficient. The main strategy to accomplish this has been to try and 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 HTCondor[1]. 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 Grid[2] project for integrating those resources into the central pool at UW-Madison and provide transparent access 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.

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 those sites we have 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, Dortmund, Mainz, Marquette, MSU, Munich, NBI, PSU, UMD and Uppsala. There are ongoing efforts at the Delaware, Chiba and LBNL sites to deploy the pyGlidein system. We continue observing a growing trend both in the number of sites integrated, as well as the computing time delivered.

[1] http://research.cs.wisc.edu/htcondor/
[2] https://www.opensciencegrid.org/

Beyond the computing capacity provided by IceCube institutions, and the opportunistic access to Grid sites that are open to share their idle capacity, IceCube has received additional computing resources from targeted allocation requests submitted to Supercomputing facilities such as the NSF Extreme Science and Engineering Discovery Environment (XSEDE). IceCube submitted a first research allocation request to XSEDE in October 2015 (allocation number TG-PHY150040) that was awarded with compute time in two GPU-capable systems: SDSC Comet[3] and PSC Bridges[4]. The latest renewal for this research allocation, in July 2018, was awarded with compute time in two systems: SDSC Comet with 180,000 SUs and PSC Bridges with 287,251 SUs. IceCube stands out as one of the largest GPU users in XSEDE, and has been acknowledged in several XSEDE press releases.

With the aim to continue exploring the possibilities to expand the pool of available computing resources for IceCube, in April 2017 we requested a test allocation for the TITAN[5] supercomputer, at the DOE Oak Ridge National Laboratory. We were awarded 1 million TITAN-hours via the "Director's Discretionary Allocation" program to test the feasibility of running IceCube payloads in that system. The TITAN supercomputer is a very attractive resource because of its large number of GPUs: 18,688 NVIDIA Tesla K20. However, its runtime environment is very different from the standard HTC clusters where IceCube usually runs. We have an ongoing project to try to investigate the best options to integrate TITAN with the IceCube computing infrastructure in an efficient way. One fairly labor intensive way has been tested, and more automated approaches are being explored.

In order to integrate all these heterogeneous infrastructures, we strongly rely on the HTCondor software and the various services that the Open Science Grid (OSG) project has built and operates around it. We continue being active in the OSG and HTCondor communities by participating in discussions and workshops. During the reporting period, IceCube members made oral contributions to the OSG all hands meeting in Salt Lake City, and the HTCondor workshop in Madison. We also gained a seat on the OSG council in March 2018, OSG's governing body, for our continued usage and contributions to OSG.

Personnel

We have had significant personnel changes in 2018. Heath Skarlupka returned to the private sector in May of 2018. Our web developer, Chad Sebranek took a new position on campus in August of 2018. David Schultz was hired into a new role as software lead developer and team lead. Finally, Gonzalo Merino took a new position at PIC and left IceCube in August of 2018. A search was conducted for a new IceCube Computing Lead. Benedikt Riedel will assume this role beginning December 2018.

[3] https://portal.xsede.org/sdsc-comet
[4] https://portal.xsede.org/psc-bridges
[5] https://www.olcf.ornl.gov/titan/

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 is 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 EZID[6] service we have the capability of issuing persistent identifiers for datasets. These are Digital Object Identifiers (DOI) that follow the DataCite metadata standard[7]. 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 12 Jul 2018): http://icecube.wisc.edu/science/data
The pages below contain information about the data that were collected and links to the data files.

1. IceCube data from 2008 to 2017 related to analysis of TXS 0506+056
   - https://doi.org/10.21234/B4QG92
2. IceCube catalog of alert events up through IceCube-170922A
   - https://doi.org/10.21234/B4KS6S
3. Measurement of atmospheric neutrino oscillations with three years of data from the full sky
   - https://doi.org/10.21234/B4105H
4. A combined maximum-likelihood analysis of the astrophysical neutrino flux:
   - https://doi.org/10.21234/B4WC7T
5. Search for point sources with first year of IC86 data:
   - https://doi.org/10.21234/B4159R
6. Search for sterile neutrinos with one year of IceCube data:
   - http://icecube.wisc.edu/science/data/IC86-sterile-neutrino
7. The 79-string IceCube search for dark matter:
   - http://icecube.wisc.edu/science/data/ic79-solar-wimp
8. Observation of Astrophysical Neutrinos in Four Years of IceCube Data:
   - http://icecube.wisc.edu/science/data/HE-nu-2010-2014
9. Astrophysical muon neutrino flux in the northern sky with 2 years of IceCube data:
   - https://icecube.wisc.edu/science/data/HE_NuMu_diffuse
10. IceCube-59: Search for point sources using muon events:
    - https://icecube.wisc.edu/science/data/IC59-point-source
11. Search for contained neutrino events at energies greater than 1 TeV in 2 years of data:

[6] http://ezid.cdlib.org
[7] http://schema.datacite.org

   - http://icecube.wisc.edu/science/data/HEnu_above1tev
12. IceCube Oscillations: 3 years muon neutrino disappearance data:
   - http://icecube.wisc.edu/science/data/nu_osc
13. 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
14. IceCube String 40 Data:
   - http://icecube.wisc.edu/science/data/ic40
15. IceCube String 22–Solar WIMP Data:
   - http://icecube.wisc.edu/science/data/ic22-solar-wimp
16. AMANDA 7 Year Data:

http://icecube.wisc.edu/science/data/amanda

Data Processing and Simulation Services

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. Now that all tapes from the Pole are in hand, we plan to complete the last 10% of this reprocessing task. The total number of files for seasons IC86-2011, IC86-2012 and IC86-2013 is 695,875; we have 67,812 remaining to be processed. The file breakdown per year is as follows: IC86-2011: 221,687 already processed out of 236,611. IC86-2012: 215,934 processed out of 222,952. IC86-2013: 190,442 processed out of 236,312. About 58,000 of the remaining files are in about 100 tapes; the rest are spread over about another 350. Tape dumping procedures are being integrated with the copy of raw data to NERSC.

Offline Data Filtering – The data collection for the IC86-2018 season started on July 10, 2018. 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 IC86-2017 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. 100 TB of storage is required to store both the Pole-filtered input data and the output data resulting from the offline production. We were able to reduce the required storage by utilizing a more efficient compression algorithm in offline production. Since season start we are using a new database structure at pole and in Madison for offline production. The data processing is proceeding smoothly and no major issues occurred. Level2 data are typically available one and a half weeks after data taking.

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.

The re-processing (pass2) has started on June 1st, 2017, and completed in August of 2018. Seven years (2010 - 2016) are currently re-processed. Four years start at sDST level (2011 - 2014) and three years at raw data. Starting at raw data was required for 2010 since sDST data was not available. Since sDST data for 2015 and 2016 has already been SPE corrected, a re-processing of sDST data was required in order to apply the latest SPE fits as we perform for the other seasons.

The reprocessing of pass2 utilized 10,905,951 CPU hours and 520 TB storage for sDST and Level2 data. An additional 2,000,000 CPU hours and 30 TB storage were required process the pass2 Level2 data to Level3.

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 is now being stored on disk in DESY and reused for different detector configurations in order to reduce GPU requirements. This transition to the 2016 configuration was done in conjunction with a switch to IceSim 5 which contains improvements in memory and GPU utilization in addition to previous improvements to correlated noise generation, Earth modeling, and lepton propagation. Current simulations are running on IceSim 5.1.3 with further improvements and bug fixes for various modules. 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 organizes 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. A full transition to IceProd 2 was completed during the Spring 2017 Collaboration Meeting. Production throughput on IceProd2 has continually increased due to incorporation of an increasing number of dedicated and opportunistic resources and a number of code optimizations. A new set of monitoring tools is currently being developed in order to keep track of efficiency and further optimizations.

Personnel – Kevin Meagher joined the simulation production team replacing David Delventhal who left in the summer 2017.

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 infra-structure. 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 seasonally (i.e. 4 times per year) with the software developers to prepare for software releases. Progress is tracked, among other means, by tracking open software tickets tied to seasonal milestones. The IceCube software group has several major projects, labeled as 'on-going' that are nearing completion:

- Simulation support for systematic features, critical for oscillation analyses (on-going);
- Reducing memory usage in simulation production (on-going);
- Improved support for individual users in production job submission (on-going);
- Increased real-time production efficiency (on-going);
- The software coordinator ran a mini-bootcamp on GPGPU computing;
- A system to display histograms generated in mass production was deployed.
- A new simulation model is currently under development to achieve more efficient simulations by sampling important parameter space instead of brute force methods to simulate easily identified background cosmic-ray showers. This dynamic-stack CORSIKA framework provides a realistic path to achieve a rate of simulation production comparable to that of data taking.

Calibration

In the period from November 2017 to January 2018 a large set of LED calibration data was collected. The calibration runs were configured such that only one of the twelve LEDs in all of the DOMs flashed at a time. Using the observed light pattern, the pointing of each of the twelve LEDs on every DOM has now been reconstructed to a high precision as shown in Figure 1. Previously, the unknown pointing and intensity of each LED contributed substantially to the overall error in the fitting of ice model parameters. Now that the pointing is known, this new data sample has superseded previous ones, and a recalibration of all relevant ice model parameters is underway. This includes measurements of the effective scattering and absorption coefficients, the anisotropy of these ice optical properties in azimuth (correlated to the glacial tilt and flow axes), and hole ice properties. In addition, at the time of deployment, each DOM was attached to the main cable in exactly the same way, such that the cable is 90 degrees off from LED number 7 on the flasher board. It follows then that the relative position of the main cable, which shadows the photocathode in an azimuthally asymmetric way, is now also known for all DOMs. The cable position has been included into a database, and can now be used in IceCube simulations.

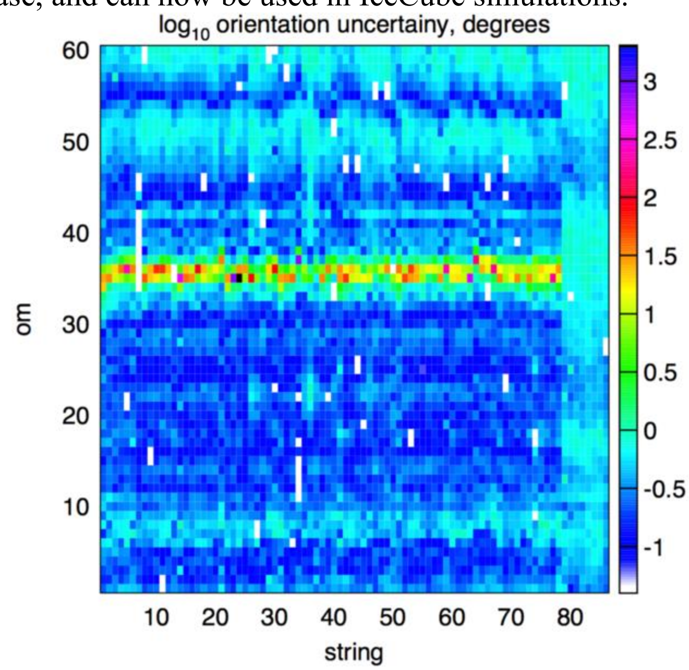

Figure 1: The error on the reconstructed position of LED 7 in degrees for all IceCube DOMs. The uncertainty is better than 1º for most DOMs, except in the dust layer where the optical properties are poor.

Data was also collected with the last movable camera system in late January 2018. The images show the persistence of the bubble column with strong scattering properties. The camera was also pointed at the direction of neighboring strings in attempt to visualize LEDs flashing in DOMs approximately 40 m away. After this data taking period, the motor on this last camera appears to have failed, with the field of view pointing into the waistband of the pressure sphere. Post-processing of the collected images have so far failed to localize the flashing LEDs, but more sophisticated noise subtraction algorithms are currently under investigation. Additionally, a simulation program for modeling the response of the this camera system has been developed, and will be used to extract more quantitative measurements from the collection of images than was previously possible.

Significant improvements have been made in the modeling of individual DOM response to single photo-electrons (SPE). A large sample of SPE waveforms have been collected, and their charge reconstructed. This charge distribution is fit using a convolution of a Gaussian and two exponential functions, yielding much better agreement between the data and the model than previously achieved with simpler models. Significant variations in the SPE charge distribution are observed DOM-to-DOM, which can have important consequences for the detection and reconstruction of low energy (GeV-scale) physics events, which often produce only single-photon hits in a few DOMs. The individual fit parameters for each DOM's SPE charge distribution are now stored in a database and used to model the charge response in IceCube simulations, improving agreement with real data. A publication summarizing this analysis and results is in progress.

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 PY3 M&O Plan will be submitted in November 2018.
- The detailed M&O Memorandum of Understanding (MoU) addressing responsibilities of each collaborating institution was revised for the collaboration meeting in Stockholm, Sweden, September 24-28, 2018.

IceCube M&O – PY3 (FY2018/2019) Milestones Status:

Milestone	**Month**
Revise the Institutional Memorandum of Understanding (MOU v24.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration Meeting	**May 2018**
Report on Scientific Results at the Spring Collaboration Meeting	**May 8-12, 2018**
Revise the Institutional Memorandum of Understanding (MOU v25.0) - Statement of Work and Ph.D. Authors head count for the fall collaboration meeting	**September 2018**
Report on Scientific Results at the Fall Collaboration Meeting	**Sept. 24-28, 2018**
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 2018**
Submit for NSF approval, a revised IceCube Maintenance and Operations Plan (M&OP) and send the approved plan to non-U.S. IOFG members.	November 2018
Revise the Institutional Memorandum of Understanding (MOU v26.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration meeting	April 2019

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)/Communications – The IceCube Collaboration has had significant outcomes from their efforts, organized around four main themes:

1) Reaching motivated high school students and teachers through IceCube Masterclasses and the University of Wisconsin-River Falls' (UWRF) and University of Rochester's (UR) Upward Bound programs, and creating cross-disciplinary opportunities through the LAB^3 project

2) Providing intensive research experiences for teachers (in collaboration with PolarTREC) and for undergraduate students (NSF science grants and Research Experiences for Undergraduates (REU) funding)

3) Engaging the public through web and print resources, graphic design, a new IceCube comic, live talks, virtual reality gaming, and displays.

4) Developing and implementing semiannual communication skills and diversity workshops held in conjunction with IceCube Collaboration meetings

The 5th annual IceCube Masterclass was held in March 2018 and had over 300 participants at 20 locations, including 8 in the US. Three applicants to the UWRF REU programs mentioned taking the IceCube Masterclass a few years earlier as a major motivator for their continued astrophysics research. Both UWRF and UR provided IceCube inspired science summer enrichment courses for their respective Upward Bound programs. Upward Bound provides additional mentoring and skill building activities for low-income/first generation high school students to help prepare them for post-secondary school success. APS and Dane County Arts Board funding was secured to support the LAB^3: Arts + Literature + Physics project. Collaborative artwork inspired by physics was produced by six teams consisting a scientist, a writer, an artist, and 3-4 high school students who worked together for about four months, culminating with a public exhibit in Madison.

Chris Wendt, one of four IceCube scientists who mentored high school students, at the ARTs3 gallery opening.

Science teacher Lesley Anderson from High Tech High in Chula Vista, CA deployed deploy to the South Pole in December 2017-january 2018 with IceCube through the PolarTREC program. She produced a number of engaging videos available on her PolarTREC page (www.polartrec.com/member/lesley-anderson) showing what is like to live and work at the South Pole. Teacher Eric Muhs, who has worked with IceCube and its predecessor AMANDA with the UWRF Upward Bound program for 15 years, will deploy in the upcoming season. UWRF's NSF astrophysics REU program through NSF selected six students for ten-week summer 2018 research experiences, including attending the IceCube software and science boot camp held at WIPAC. Multiple IceCube institutions also supported research opportunities for undergraduates. One student from the 2017 UWRF IRES program attended the IceCube Collaboration meeting in May, 2018, in Atlanta and presented poster on his research---five other posters from his fellow IRES colleagues were also displayed.

The public press conference at NSF in July 12, 2018 held in conjunction with the publication of two papers on neutrinos associated with the Blazar TXS 0506+056 was an opportunity to celebrate achieving a major goal of the IceCube project---detecting the first point source of neutrinos. The press conference was promoted on social media (the top Tweet had over 100, 000 impressions) and was supported by an amazing collection of graphics and videos produced in-house and by partners from other major agencies (NASA, DESY, etc.) inside and outside the US. IceCube worked with colleagues at the Wisconsin Institutes of Discovery and Field Day Labs who produced a virtual reality experience for the Oculus Rift system. It allows users to start at the IceCube Laboratory at the South Pole and follow the path of neutrino back to its source, a black hole billions of light-years away. The game, which takes about 5 minutes to play, has been a big hit at a number of venues including the City of Science day at the 2018 World Science Festival in New York City, where thousands visited the IceCube booth.

IceCube continues to receive multiple request for talks per week, and works hard to provide speakers for all opportunities. IceCube collaborators had a large presence at the Polar 2018 joint meeting of the Scientific Committee on Antarctic Research and the International Arctic Science Committee held in Davos, Switzerland. IceCube had more than a dozen presentations (talks and posters) in total covering science, communication, and outreach, and have collaboration members playing leadership roles in the SCAR Astronomy and Astrophysics from Antarctica Scientific Research Program.

The IceCube Communications office manages press and other communications activities for both the neutrino observatory and the IceCube Collaboration. During the last months, strategies to increase the impact of our activities have focused on the launch of the comic series "Rosie&Gibbs" that targets younger audiences with content related to both IceCube science and the adventures of working in Antarctica. We keep also increasing the production of multimedia content—graphics and videos, which through social networks has increased the reach of IceCube communications from a few thousands to tens of thousands on an average week and with peaks of hundreds of thousands associated with big announcements.

In 2015, IceCube launched a professional development program with a strong focus on communications and diversity. Twice per year, during the spring and fall collaboration meetings, IceCubers have the possibility to participate in a communications training and/or in a workshop to discuss or share best practices to increase diversity and inclusion in IceCube and related fields. Both events are attended by 25-30 people on average and are highly valued by the target communities: early career researchers and women and allies. A strategy to make these efforts sustainable and broaden the topics on diversity and inclusion have led to an INCLUDES program.

A supplemental award to the IceCube Maintenance and Operations grant will expand the NSF Includes network by developing a Multimessenger@Includes activity. The goal is to build on the initiatives of the IceCube Neutrino Observatory and its partners in multimessenger astronomy to further develop talent for STEM research, increase diversity and inclusion in science and related fields, and connect research-based skills and technology to societal challenges. led by the IceCube Diveristy Task Force. The Multimessenger@INCLUDES activities will be twofold: 1) implement collaboration-wide efforts to help IceCube advance as an inclusive and open community that excels in research, fosters STEM careers in academia and industry, and engages local communities; 2) establish alliances with other multimessenger astronomy collaborations and launch a new Multimessenger Diversity Network that will identify common goals and develop shared strategies to broaden participation in STEM, and become a member of the main INCLUDES National Network. This is a 30 month project that started July 31, 2018.

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 2018 for the Fall collaboration meeting in Stockholm, Sweden (v25.0), and the next revision (v26.0) will be posted in May 2019 at the Spring collaboration meeting in Madison, WI.

IceCube Collaborating Institutions

Following the September 2018 Fall collaboration meeting, the Mercer University with Dr. Frank McNally as the institutional lead joined the IceCube Collaboration.

As of October 2018, the IceCube Collaboration consists of 49 institutions in 12 countries (26 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 – Atlanta, GA	May 8-12, 2018
Software and Computing Advisory Panel Meeting – Madison, WI	June 4-5, 2018
IceCube Fall Collaboration Meeting – Stockholm, Sweden	September 24-28, 2018
International Oversight and Finance Group – Stockholm, Sweden	September 28, 2018
ICNO M&O Mid-Term Review	January 7-8, 2019

Acronym List

CPU	Central Processing Unit
CVMFS	CernVM-Filesystem
DAQ	Data Acquisition System
DOM	Digital Optical Module
E&O	Education and Outreach
GPU	Graphical Processing Unit
I3Moni	IceCube Run Monitoring system
IC86	The 86-string IceCube Array completed Dec 2010
IceACT	IceCube Air Cherenkov Telescope
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
PQ	Physical Qualification
SNDAQ	Supernova Data Acquisition System
SPE	Single photoelectron
SPS	South Pole System
SuperDST/sDST	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