IceCube Maintenance and Operations
Fiscal Year 2019 / 2020 PY4 Mid-Year Report

April 1, 2019 – September 30, 2019

Submittal Date: September 30, 2019

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 FY2019/2020 (PY4) Mid-Year Report is submitted as required by the NSF Cooperative Agreement PLR-1600823. This report covers the 8-month period beginning February 1, 2019 and concluding September 30, 2019. The status information provided in the report covers actual common fund contributions received through September 30, 2019 and the full 86-string IceCube detector (IC86) performance through September 30, 2019.

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 PY4 installment of $5,250,000 was released to UW–Madison to cover the costs of maintenance and operations during the first three quarters of PY4 (FY2019/FY2020): $726,863 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 $4,523,138 was directed to the IceCube M&O Core account. The second PY4 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 PY4 (FY2020): $242,288 will be directed to the U.S. Common Fund account, and the remaining $1,507,713 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.

PY4: FY2019 / FY2020	Funds Awarded to UW for Apr 1, 2019 – March 31, 2020	Funds to Be Awarded to UW for Apr 1, 2018 – March 31, 2019
IceCube M&O Core account	$4,523,138	$1,507,713
U.S. Common Fund account	$726,863	$242,288
TOTAL NSF Funds	**$5,250,000**	**$1,750,000**

Table 1: NSF IceCube M&O Funds – PY4 (FY2019 / FY2020)

Of the IceCube M&O PY4 (FY2019/2020) Core funds, $981,015 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 FY2019/2020 funds for the subawardee institutions.

Institution	Major Responsibilities	Funds
Lawrence Berkeley National Laboratory	DAQ maintenance, computing infrastructure	$92,001
Pennsylvania State University	Computing and data management, simulation production, DAQ maintenance	$71,908
University of Delaware, Bartol Institute	IceTop calibration, monitoring and maintenance	$151,720
University of Maryland at College Park	IceTray software framework, online filter, simulation software	$611,444
University of Alabama at Tuscaloosa	Detector calibration, reconstruction and analysis tools	$24,963
Michigan State University	Simulation software, simulation production	$28,979
South Dakota School of Mines and Technology (added in July 2017)	Simulation production and reconstruction	$00.00
Total		**$981,015**

Table 2: IceCube M&O Subawardee Institutions – PY4 (FY2019/2020) 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 4 of the award, and shows an estimated balance at the end of PY4.

Total awarded funds to the University of Wisconsin (UW) for supporting IceCube M&O from the beginning of PY1 through the end of PY4 are $28,360K (including the supplemental funding of $67,999 in PY2 and $291,712 in PY3). Total actual cost as of September 30, 2019 is $24,458K and open commitments against purchase orders and subaward agreements are $1,034K. The current balance as of September 30, 2019 is $2,867K. With a projection of $3,008K for the remaining expenses during the second half of PY4, the estimated negative balance at the end of PY4 is -$141K, which is 0.5% of the total PY1-PY4 budget (Table 3).

(a)	(b)	(c)	(d)= a - b - c	(e)	(f) = d – e
YEARS 1-4 Budget Apr.'16-Mar.'20	**Actual Cost To Date** through Sept. 30, 2019	**Open Commitments** on Sept. 30, 2019	**Current Balance** on Sept. 30, 2019	**Remaining Projected Expenses** through Mar. 2020	**End of PY4 Forecast Balance** on Mar. 31, 2020
$28,360K	**$24,458K**	**$1,034K**	**$2,867K**	**$3,008K**	**-$141K**

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	**137**	**$1,870,050**	**$1,646,562**
U.S. Contribution	71	$969,150	$969,150
Non-U.S. Contribution	66	$900,900	$677,412*

Table 4: Planned and Actual CF Contributions for the period of April 1, 2018–March 31, 2019

* 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, 2019 to September 15, 2019, the detector uptime, defined as the fraction of the total time that some portion of IceCube was taking data, was 99.82%, 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 98.28%, 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 0.86% 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 0.18%.

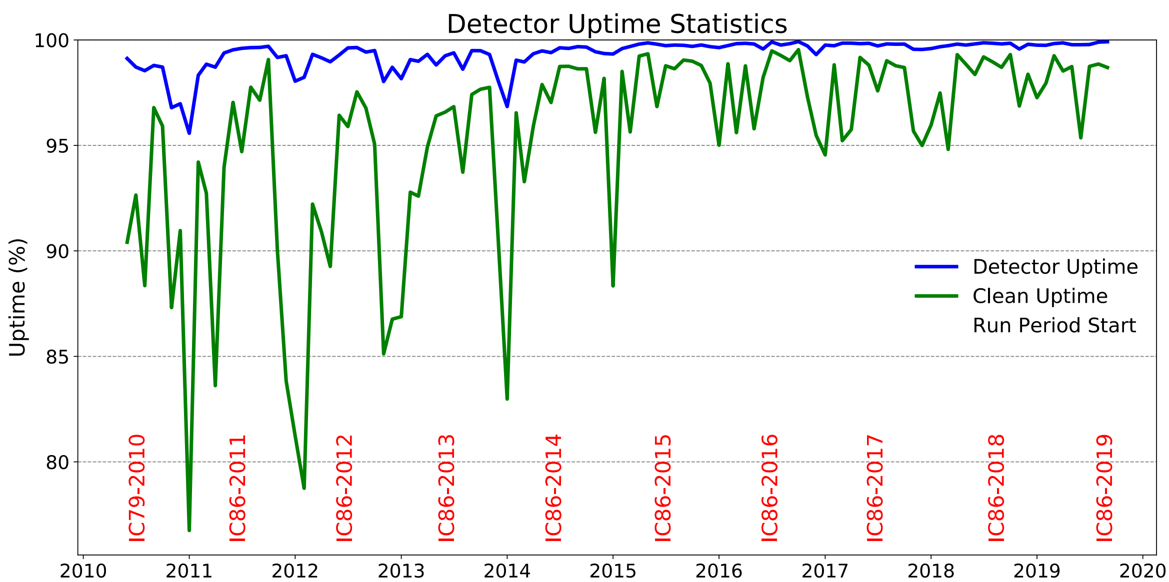

Figure 1: Total IceCube Detector Uptime and Clean Uptime

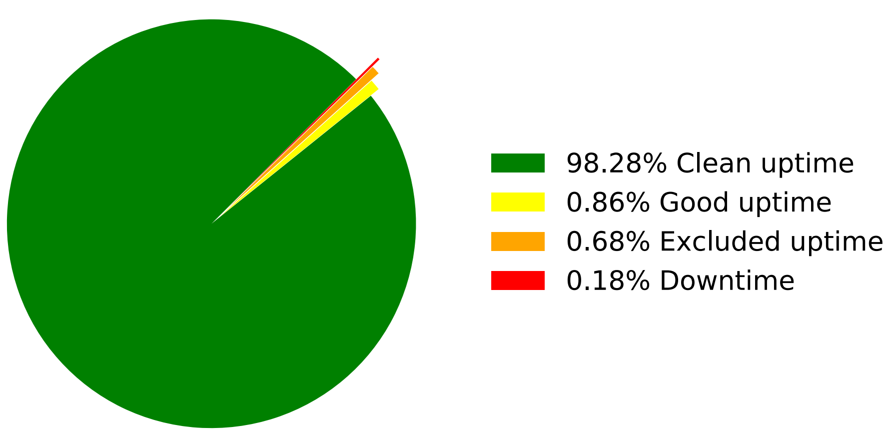

Figure 2: Cumulative IceCube Detector Time Usage, February 1 – September 15, 2019

Hardware Stability — The total number of active DOMs in the data stream is currently 5405 (98.5% of deployed DOMs), plus three DOM-mainboard-based scintillator panels and the IceAct trigger mainboard. No DOMs failed during this reporting period; the previous DOM failures were in December 2018. One previously dead DOM was returned to data-taking (next section).

The failure rate in the commercial Acopian power supplies that supply the DC voltage to the DOMs increased in late 2015 to an unacceptably high level. We transitioned approximately half of the detector to alternate Mean Well MSP-200-48 power supplies during the 2018–19 pole season. As of September 2019, none of the alternate supplies have failed. We plan on switching completely to the Mean Well supplies in the 2019–20 pole season.

The winterovers discovered in March 2019 that approximately 50% of the fans in the DOMHub ATX power supplies have failed over time (Fig. 3, left). This is not necessarily unexpected, as the lifetime of this type of fan is ~8 years. We believe this is likely the cause of an elevated number of ATX module failures over the past several years, as some modules have no working fans. We have identified and tested an alternative fan with a higher-lifetime hydrodynamic bearing (Fig. 3, right) and plan to replace the failed fans this austral summer season. This can be done with minimal detector downtime since the modules are redundant and hot-swappable.

Figure 3: Left: 50% of the fans in the DOMHub ATX power supplies have failed. Right: replacement fan type (Noctua NF-A4x20 FLX) tested successfully at SPTS.

After an extended testing and commissioning campaign over the past two years, IceCube was successfully transitioned in March 2019 to a new master clock, a Spectracom SecureSync. Tests with flasher LEDs using the new clock show that the detector timing resolution is at least as good as with the previous master clock (Fig. 4). The new clock handled the GPS week rollover in April 2019 without incident. The hot spare master clock will also be exchanged with a new Spectracom this coming maintenance season.

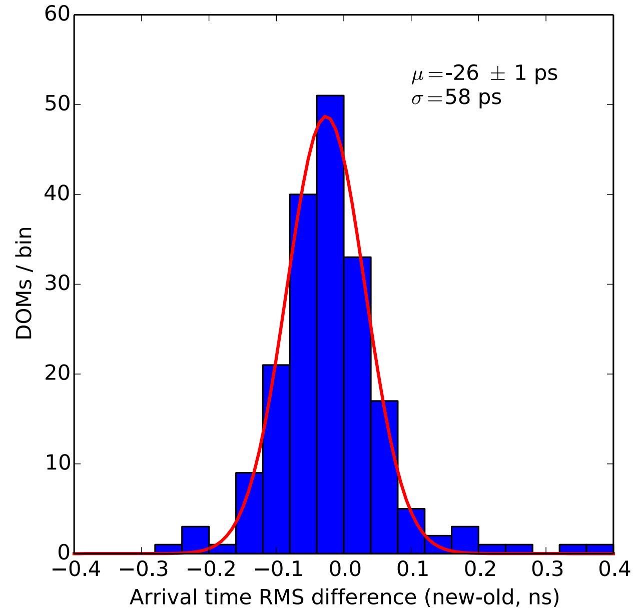

Figure 4: Comparison of flasher light arrival time spread using old and new master clocks. The new master clock shows comparable timing resolution.

IC86 Physics Runs — The ninth season of the 86-string physics run, IC86–2019, began on July 18, 2019. Detector settings were updated using the latest yearly DOM calibrations from March 2019. The external trigger from the roof IceAct telescope was enabled, and a new low-energy neutrino realtime event stream was added. "Pass3" improvements to calibration and waveform unfolding were deferred, as benefits to physics analyses were not sufficiently compelling to outweigh the discontinuity that would be introduced to on-going data analyses.

One previously dead DOM, 30-60 "Rowan", spontaneously started communicating again after a full-detector power cycle in late 2018. The DOM was successfully recalibrated and returned to data-taking for the IC86–2019 run start, after a 13-year hiatus.

We have implementing automated HitSpool data captures triggered from the northern hemisphere by LIGO-VIRGO gravitational wave alerts, enabling searches for sub-threshold coincident neutrinos. Coincident hitspool data around black-hole merger candidate alerts are archived to backup disk, while data around neutron-star merger alerts are sent via satellite using a low-priority "topoff" queue.

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. During the reporting period, the following accomplishments are noted:

- Delivery of four pDAQ:Urban_Harvest patch releases, fixing a bug affecting the IceAct trigger when rates were very low, and adding improved robustness to internal component stalls.

- Work towards the pDAQ:Vintage release, which will incorporate changes necessary for migration to Python version 3 (Python version 2 will no longer be supported past 2020).

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. Filter changes were released to support the IC86–2019 physics run start in July 2019, as well as internal improvements to software component control.

An increased number of component crashes within the PnF system has been observed (Fig. 5). While this only causes a temporary backlog in processing, restarting event filtering requires intervention by the winterovers. To address this, a watchdog system is under development that can detect the crashes and restart the necessary components automatically.

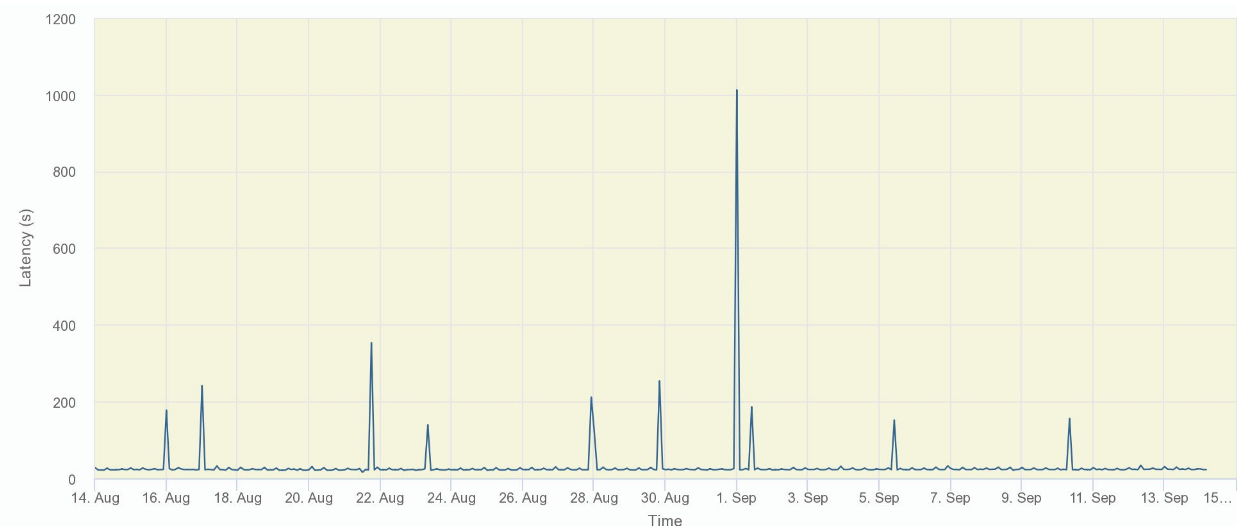

Figure 5: Latency of PnF system over time (lower is better). The spikes are temporary backlogs due to software component stalls.

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

- Live v4.0 "Arcadia" (April 2019): includes an upgrade to Python 2.7 as well as upgrades to the internal database and web framework. This release also includes an improved supernova DAQ (SNDAQ) web page.

- Development toward a late 2019 patch release, which fixes minor bugs around run status updates and event counts as well as adding support for an alert if the detector is running for extended periods in a non-standard configuration.

The uptime of the I3Live experiment control system during the reporting period was approximately 99.997%.

Supernova System — The supernova data acquisition system (SNDAQ) found that 99.77% of the available data from February 1 to August 31, 2019 met the minimum analysis criteria for run duration and data quality for sending triggers. An additional 0.04% 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 version of the SNDAQ (BT16) was deployed in March 2019 and includes streamlined supernova alert messages to IceCube Live. Additionally, a new system was developed in conjunction with IceCube Live that monitors the status of alerts to the Supernova Neutrino Early Warning System (SNEWS).

Surface Detectors — Two scintillator stations were deployed in the 2017–18 season, consisting of seven panels each. In the 2018–19 season, we successfully demonstrated that the elevated panels can be subsequently raised, as well as enhancing one station with two prototype radio antennas. Measurement of the air shower radio emission provides shower-by-shower energy and mass composition information, significantly enhancing the capabilities of IceTop. The radio antenna readout, however, is currently not functional due to a failure in one set of the surface data acquisition electronics. The other scintillator station is functional and taking commissioning data.

Both the field-deployed and ICL roof IceAct telescopes have successfully taken data during the austral winter and are now covered again for the polar sunrise. The addition of a lens heater has partially, but not completely resolved issues with ice accumulation.

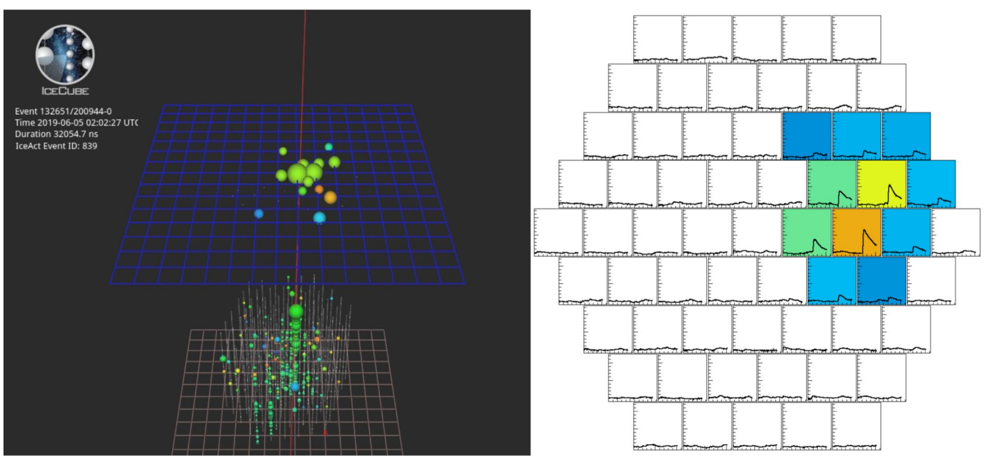

Figure 6: Coincident air shower recorded in IceTop + in-ice array (left) and an IceAct telescope (right), from June 2019.

During the 2019–20 summer season, the prototype scintillator station will be upgraded to the production version of the electronics, and old hardware will be removed. All panels and the central electronics FieldHub will be elevated above the snow surface. The IceAct roof telescope we be refurbished with the same DAQ as the field telescope, and the lens heating system will be upgraded.

A comprehensive, unified plan regarding all proposed IceTop surface enhancements has been reviewed internally by the collaboration. The layout of the surface detectors on the IceTop footprint has been re-optimized to simplify the logistics, in particular by reducing the trenching requirements. The construction schedule will take into account the personnel, cargo, and logistics demands of the IceCube Upgrade project.

South Pole System — Approximately 50% of the South Pole System servers were replaced in the 2018–19 season; performance and hardware stability has been excellent. The remaining servers will be replaced in the 2019–20 season. As key data-taking servers will be replaced, procedures have been developed to minimize detector downtime by swapping software components to alternate machines during the upgrade.

Personnel & Management — An embedded software developer, Jeff Weber, has been hired at WIPAC and started in August 2019. He will initially focus on Upgrade DOM software development but will also contribute in the future to M&O data acquisition maintenance.

Computing and Data Management Services

Data Transfer – Data transfer has performed nominally over the past ten months. Between February 2019 and September 2019 a total of 16 TiB of data were transferred from the South Pole to UW-Madison via TDRSS, at an average rate of 68 GiB/day. Figure 1 shows the daily satellite transfer rate and weekly average satellite transfer rate in 69 GiB/day through September 2019. 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.

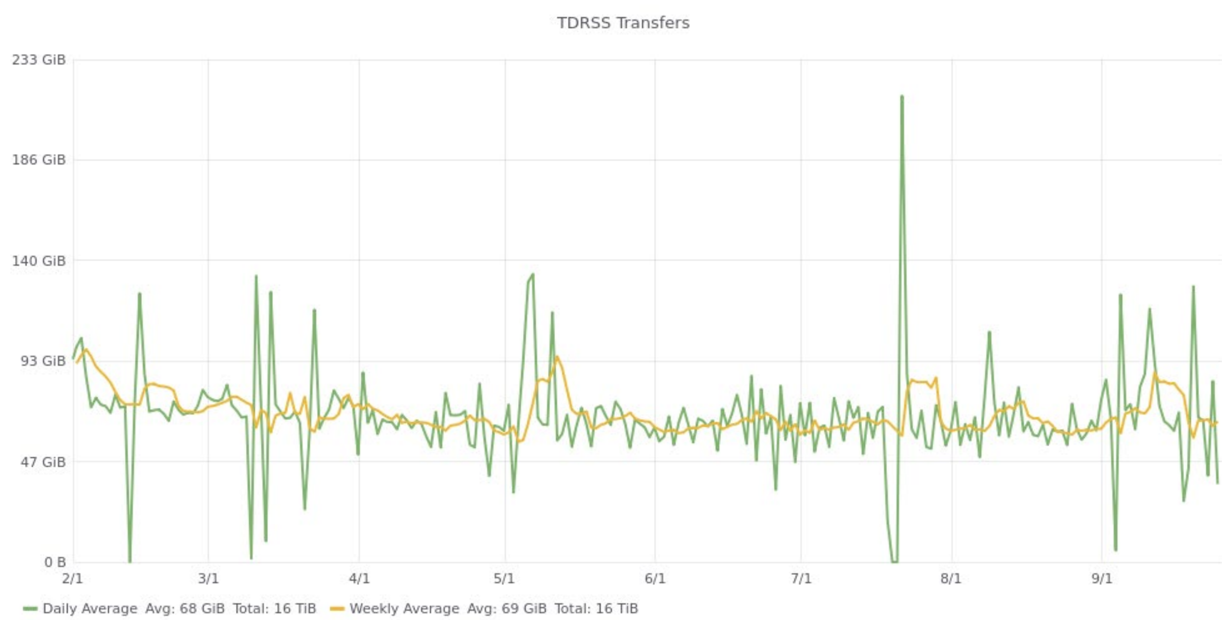

Figure 1: TDRSS Data Transfer Rates, February 1, 2019–September 25, 2019. The daily transferred volumes are shown in green and, superimposed in yellow, the weekly average daily rates are also displayed.

Data Archive – The IceCube raw data is archived by writing two copies on independent hard disks. During the reporting period (February2019 to September 2019) a total of 241 TiB of unique data were archived to disk averaging 1 TiB/day.

In May 2019, the set of archival disks containing the raw data taken the previous season 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 3.2 PB. Figure 2 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.

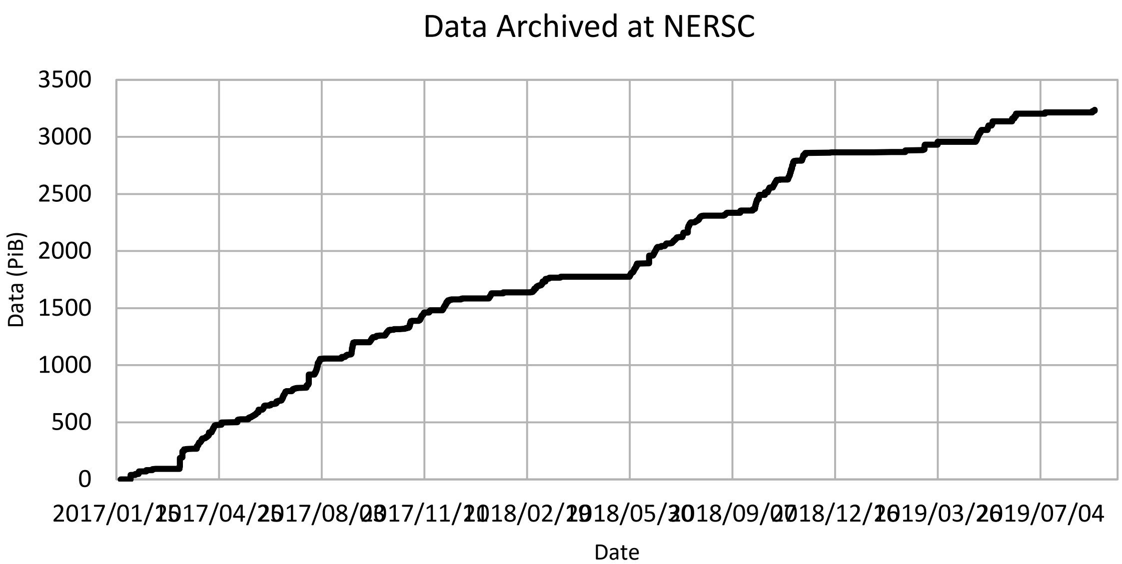

Figure 2: Volume of IceCube data archived at the NERSC tape facility by the JADE Long Term Archive service as a function of time.

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 been 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 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 2019. 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 6.8 PB: 2.1 PB for experimental data, 4.4 PB for simulation and analysis and 297 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 by the end 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 moved 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. 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. IceCube has submitted proposal to Phase I of Internet2's Exploring Clouds for Acceleration of Science (E-CAS) project for $100,000 in credits with Amazon Web Services and Google Cloud Platform and 0.4 FTE. We are also exploring other proposals for dedicated GPU compute infrastructure.

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 2019:

Institution	CPU	GPU
Aachen	10	0

Alberta	1200	144
Brussels	400	14
DESY	1400	180
Dortmund	400	48
LBNL	114	0
Marquette	96	16
MSU	750	100
NBI	0	10
Queen’s/King’s College	0	55
UMD	350	36
UW-Madison	6000	400
Wuppertal	500	64

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

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

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, LBNL, Mainz, Marquette, Munich, NBI, and UMD. There are ongoing efforts at the Delaware, and Chiba 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.

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 2019, was awarded with compute time in two systems: SDSC Comet with 150,000 SUs and PSC Bridges with 150,00 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, IceCube is taking part a demonstration at Supercomputing 2019 funded through an EAGER award (NSF OAC 1941481) from the NSF OAC. This demonstration is meant to present ability of public cloud resources to replicate current state-of-the-art high-performance computing.

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 at Jefferson Lab in Virginia, and the HTCondor workshop in Madison.

Personnel

We have had two personnel changes in 2019. Davy Mayer was hired in May 2019 as a web developer. Eric Evans was hired in September 2019 to work on analysis software and monitoring infrastructure.

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

[3] https://portal.xsede.org/sdsc-comet
[4] https://portal.xsede.org/psc-bridges

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[5] service we have the capability of issuing persistent identifiers for datasets. These are Digital Object Identifiers (DOI) that follow the DataCite metadata standard[6]. 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:
    - 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

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

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 **–** The data re-processing (pass2) started on June 1st, 2017, and completed in August of 2018 to (a) unify the multi-year data set; (b) profit from improvements in our understanding of low-level DOM calibration parameters. 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.

More recent improvements in calibration and feature extraction prompted an evaluation for the need to implement such changes in the online filter and subsequently reprocess previous data a second time. However, after careful studies it was determined that there was not sufficient evidence that a pass3 would result in improved science at this time. The question will be revisited before next season, after we have made more detailed studies on the matter.

Offline Data Filtering **–** The data collection for the IC86-2019 season started on July 18, 2019. 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 IC86-2018 season scripts are minimal. An effort was done to clean up filters reconstructions and libraries no longer needed in offline reconstruction resulting in a 36% reduction of CPU utilization and a comparable reduction in memory requirements. We therefore estimate that the resources required for the offline production will be about 480,000 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. Additional savings in storage resulted from a switch to a more efficient compression algorithm in the 2017 season. As of the start of the 2018 season. We switched to a new database structure at pole and in Madison for offline production. The transition went smoothly, with no significant issues. 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.

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 pass previous trigger and filter configurations included in pass2. These include 2012, 2013, 2014, 2015 and 2016, as well as more recent configurations such as 2017 and 2018. As with previous productions, direct generation of Level 2 background simulation data is used to reduce storage space requirements. The 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 6.1.1 with further improvements and bug fixes for various modules. Some new features in IceSim 6 include, individually calibrated PMT waveforms, optimized event resampling for low-energy background simulation as well as new, more detailed ice models along with further improvements in GPU utilization. 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 a one-size-fits-all approach. New strategies are currently under development for simulating and determining the impact of systematic uncertainties in our understanding of ice properties, hole-ice and DOM sensitivity. A new campaign will begin once the necessary tools are in place.

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 – Ric Evans joined the simulation production team in September.

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:

- Program management software is now being used to manage the M&O software effort;
- New release plan which includes WG Tech Leads and the ICC was developed;
- Included the latest muon propagation project (PROPOSAL) into our standard stack;
- Included a server based photon propagation scheme in production to decrease GPU idle;
- Simulation support for systematic features, critical for oscillation analyses (deployed);
- Reducing memory usage in simulation production (deployment-ready);
- Improved support for individual users in production job submission (deployed);
- Increasing real-time production efficiency (on-going);
- The software coordinator ran a mini-bootcamp on Deep Learning;
- Developed a system to run extensive system tests nightly (deployed).
- Significant progress towards a new background simulation model was made and is close to deployment. This model is designed 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

Using single-flashing LED data collected in 2018, we continue to refine measurements of the South Pole glacial ice properties, which are critical input to all IceCube analyses. The azimuthal pointing of all 12 LEDs on each DOM are now calibrated to better than 1°, such that only the intensity of each LED is unfolded in the ice model fit. Using this new flasher data, the depth-dependent scattering and absorption coefficients have been reevaluated along with their errors, taking into account all known systematic uncertainties pertaining to the LED emission and local ice properties near the DOMs (e.g. hole ice). The resulting uncertainties on bulk ice absorption and scattering have been reduced from ~10% down to ~5%.

For many years we have observed an anisotropy in light propagation that is correlated to the direction of the glacial flow axis in IceCube. Previous attempts to model this anisotropy have relied on a scaling of the scattering function strength according to the direction of propagating photons. This model resulted in an improved description of the total photon count on DOMs, but a poor agreement in the leading edge of the time profile due to the modification of the scattering strength. In the last six months significant efforts have been spent on understanding and modelling photon transport while taking into account the microscopic properties of ice, which is a birefringent polycrystal with an optical axis. This means that the average size and orientation of ice grains influences the light dispersion in a direction dependent way. These are currently treated as free parameters until the SPICEcore collaboration finishes their measurements of the ice core they collected. Preliminary attempts to model photon transport taking into account the ice grain properties have resulted in a better description of the observed light pattern in both charge and time. Given these promising results, efforts are now focused on trying to improve the effective parameterisations and speed up the fitting process. We are in close contact with the SPICEcore collaboration so that when their measurements of the depth-dependent ice grain size and c-axis distribution are known they can be included directly in the model for the available depths.

The scope of the calibration group activities has increased to incorporate review of recent SPICEcore logging activities and sensitivity studies for IceCube Upgrade. In November of 2018 several loggers were deployed in the SPICEcore bore hole that can measure the UV attention of light in the ice, the luminescence of ice due to beta irradiation and testing Upgrade camera prototypes. The data analysis from these loggers is still on-going, with additional deployments expected this year at pole. However, even with preliminary results we have found that the photon yield from luminescence is measurable at the level of about 1 photon per MeV (in beta energy), as shown in **Figure 1**, and should be included in future simulations for a more accurate energy scale calibration.

In preparation for the IceCube Upgrade, many studies are underway to determine the sensitivity of the new devices (LEDs, POCAM, Pencil Beam, etc.) to various ice and optical module systematic uncertainties. These studies are presented in the calibration group, and have been valuable in determining the requirements for many of the new devices, such as flasher emission profiles, wavelengths and timing characteristics. These studies will continue to mature and develop prior to the Upgrade installation so that calibration procedures using each of the new devices will be well defined and documented ahead of *in situ* operation, along with their expected performance.

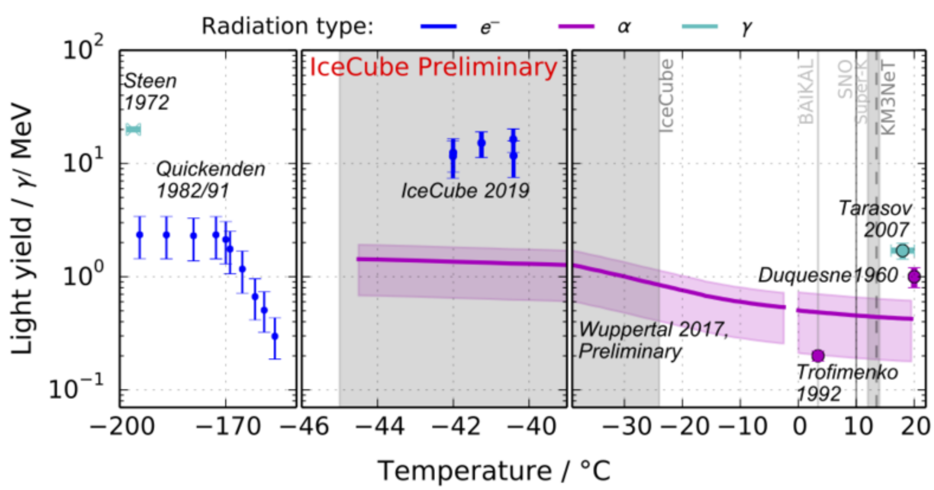

Figure 1: Light yield (photons/MeV) from luminescence in ice. The blue data labelled IceCube 2019 (middle panel) represent the first in situ measurements at the South Pole. Other measurements are made in the laboratory with very different temperature, pressure and impurity properties.

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 was submitted in January 2019 and PY4 M&O Plan will be submitted before the end of November 2019.
- The detailed M&O Memorandum of Understanding (MoU) addressing responsibilities of each collaborating institution was revised for the collaboration meeting in Chiba, Japan, September 15-20, 2019.

IceCube M&O – PY4 (FY2019/2020) Milestones Status:

Milestone	Month
Revise the Institutional Memorandum of Understanding (MOU v26.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration Meeting	**May 2019**
Report on Scientific Results at the Spring Collaboration Meeting	**Apr 30 – May 4, 2019**
Revise the Institutional Memorandum of Understanding (MOU v27.0) - Statement of Work and Ph.D. Authors head count for the fall collaboration meeting	**September 2019**
Report on Scientific Results at the Fall Collaboration Meeting	**Sept. 15-20, 2019**
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 2019**
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 2019**
Revise the Institutional Memorandum of Understanding (MOU v28.0) - Statement of Work and Ph.D. Authors head count for the spring collaboration meeting	April 2020

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 internships, IceCube Masterclasses, and the University of Wisconsin-River Falls' (UWRF) and University of Rochester's (UR) Upward Bound programs
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) to increase STEM awareness

3) *Engaging the public through various means,* including web and print resources, graphic designs, an IceCube comic, live talks, virtual reality gaming, and displays
4) *Developing and implementing communication skills and diversity workshops*, held semiannually in conjunction with IceCube Collaboration meetings

The 6th annual IceCube Masterclass held in January and March 2019 had some weather challenges that forced the South Dakota School of Mines and Technology and UWRF to cancel. New US participants included Marquette University and Mercer University, who hosted in conjunction with Georgia Tech. The UWRF REU program provided summer research experiences to six students, five women and one man. Both UWRF and UR provided IceCube-science–inspired 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.

Bilingual science educator Jocelyn Argueta was selected to deploy with IceCube to the South Pole in the upcoming season in partnership with the PolarTREC program. Argueta has a nationally touring one-person science show as Jargie the Science Girl where she presents science concepts in a fun and engaging fashion for young audiences. Multiple IceCube institutions also supported research opportunities for undergraduates including Aachen University in Germany who hosted a UWRF undergraduate for 10 weeks in summer 2019.

University of Delaware IceCube scientist Serap Tilav at the World Science Festival.

Portraits at the ICRC2019 gallery show.

WIPAC organized the IceCube booth June 2, 2019, at the World Science Festival in New York City with volunteers from Columbia, Delaware, Stony Brook and WIPAC. WIPAC also participated in several significant events in Madison, including a two-day IceCube exhibit at Science Expeditions in April on the UW campus that drew thousands of visitors, a two-day event for UW alumni and their grandchildren that introduced IceCube science in an accessible way, and a day of interactive IceCube activities at the Saturday Science program on July 13, 2019. In conjunction with the 2019 International Cosmic Ray Conference held on the UW campus, WIPAC co-sponsored and curated a gallery show with York University Art-Science Associate Professor Mark-David Hosale and UW Art Associate Professor Faisal Abdu'Allah that explored what we know and how we know it and took a look at who does cosmic ray research. A particularly rewarding aspect of this event was working with local teens who assisted with taking the portraits of ICRC2019 participants, which were added to the exhibit during the conference.

Local and national media covered an IceCube DOM that was sent on request in June 2019 to be included in the Smithsonian Air and Space Museum. The IceCube virtual reality experience developed for the Oculus Rift system continues to be a big hit. It has been utilized by thousands of people in the last six months at venues including the NSF Large Facilities Workshop, South by Southwest, the World Science Festival, UW's Science Expeditions, and the 2019 Council of Managers of National Antarctic Program (COMNAP) workshop in Bulgaria.

IceCube receives multiple requests for talks each week and works hard to provide speakers for all opportunities. IceCube had a large presence at ICRC2019, with about 80 presentations in total, including posters on the gallery show and on a collaboration between NSF Artist and Writer recipient Donald Fortescue and IceCube collaborator Gwenhael De Wasseige to convert IceCube data to music by associating each IceCube string with a string on a piano. IceCube collaborators play leading roles on national and international organizing committees including the previously mentioned hosting of ICRC2019 as well as leadership roles in the SCAR Astronomy and Astrophysics from Antarctica Scientific Research Program.

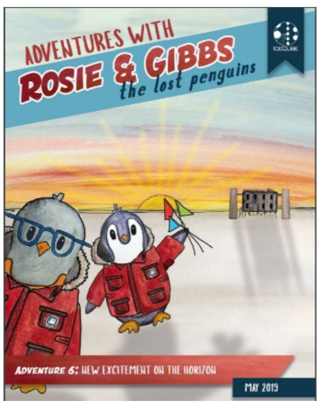

Rosie and Gibbs 6th issue discusses the IceCube Upgrade.

The IceCube communication office manages press and other communication activities for both the neutrino observatory and the IceCube Collaboration. A sixth issue of the IceCube comic series "Rosie & Gibbs" was published in May 2019 that featured the IceCube Upgrade as well as a compilation of all six issues. We continue to produce multimedia content for social networks, which has increased the reach of IceCube communication from a few thousand to tens of thousands on an average week, with peaks of hundreds of thousands associated with big announcements.

IceCube Collaboration meetings include professional development opportunities to improve skills and help create a more inclusive environment. Science journalist Angela Posada-Swafford presented a communication workshop and participated in the evening women's networking event at the Madison IceCube Collaboration meeting in May, which also included an interactive workshop on interviewing tips. At the Chiba IceCube Collaboration meeting, there was an evening gathering for LGBTQ+ and allies, a women and allies networking event, and a communication workshop that included presentations by IceCube collaborator Lu Lu on her augmented reality project to display IceCube events and by anime artist Dr. Yuki Akimoto who also attended the women and allies event.

The Multimessenger Diversity Network funded by an M&O supplemental award 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 has been active. Representatives from the initial four collaborations met for training in Madison on March 9 and 10, 2019. An Astro2020 APC White Paper: "Pursuing diversity, equity, and inclusion in multimessenger astronomy collaborations over the coming decade" was produced. The number of participating collaborations grew by three to seven total. IceCube participated in social media campaigns supporting Pride month and LGBTSTEM Day on July 5, 2019.

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 2019 for the Fall collaboration meeting in Chiba, Japan (v27.0), and the next revision (v28.0) will be posted in May 2019 at the Spring collaboration meeting in Brussels, Belgium.

IceCube Collaborating Institutions

As of September 2019, the IceCube Collaboration consists of 52 institutions in 12 countries (29 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

ICNO M&O Mid-Term Review-Madison, WI	March 11, 2019
IceCube Spring Collaboration Meeting – Madison, WI	April 30 – May 4, 2019
IceCube Fall Collaboration Meeting – Chiba, Japan	September 15-20, 2019
International Oversight and Finance Group – Chiba, Japan	September 20, 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