Report from the 7th meeting of the IceCube Software and Computing Panel (SCAP) on January 27-28, 2021

L. Bauerdick (chair), P. Couvares, S. Dasu, P. Elmer, M. Livny,
G. Merino, M. Neubauer, T. Schalk, F. Würthwein

March 2021

The IceCube Software and Computing Advisory Panel (SCAP) met with management and members of IceCube Project, Operations, and Collaboration on January 27 and 28, 2021, in a virtual meeting using video conferencing. The panel addressed three specific charge questions shown in the Appendix, together with the meeting agenda.

The SCAP thanks IceCube for well-prepared presentations and accommodating the remote participation format restrictions and a very tight schedule. The panel thanks all participants for the open and frank discussions and the overnight preparation of answers to the panel's questions.

General Remarks

IceCube has been very successful as a unique scientific instrument to explore neutrino astronomy, particularly with establishing the existence of neutrino point sources in recent years. The SCAP welcomes the news of the recent publication of its 10-year all-sky point-source dataset, including an accompanying paper describing the release.

IceCube has transitioned from discovery to precision science, and, to wit, it is set on a path towards creating a catalog of point source observations beyond the 3σ significance level. The required large improvements in sensitivity and resolution to a significant degree are a computational problem, thus the importance of software and computing (S&C) for the scientific program's success. IceCube produces unique high-impact science in astrophysics that justifies an increased effort investment from agencies and the community to guarantee the program's complete success.

The SCAP is satisfied with the IceCube team providing excellent support and services to experimenters and the project, enabling IceCube to produce first-class scientific results while managing resources well. IceCube has successfully addressed several challenges and risks discussed at the meeting within the constraints on budgets and the growing demands for improving and modernizing software and data processing systems.

IceCube will soon begin a new funding cycle for Operations and has asked for a significant increase (+28%) after ten years of flat M&O funding at $7M since its start in 2011. The SCAP heard IceCube

make a strong and convincing case for increased funding to cover increasing costs and support needs in the S&C area.

The SCAP fully **endorses** the stated need for strengthening the Operations team in the area of S&C support, and in particular **the need to hire five additional FTE** to cover essential functions. The requested five additional positions would allow bolstering IceCube's ability

- to continue supporting the expanding science program to deliver high-quality results;
- to modernize and provide long-term support for aging software, workflow, and data systems;
- to adapt the software systems to integrate the IceCube Upgrade, which only has minimal support within the upgrade project and thus falls mainly on M&O;
- to exploit the revolution in calibration expected with new devices deployed in the Upgrade;
- to support modern computing platforms to grow use of HPC platforms;
- to support algorithmic innovations that will increase the detector's scientific reach;
- to make more effective use of GPUs and other computing accelerators and novel technologies, particularly AI/ML.

Without the requested additional efforts and expertise, the SCAP would predict significant risks for the IceCube Operations team to continue their so-far successful program with even baseline operations support.

The SCAP had specific concerns regarding IceCube's long-term planning and overall organizational structures in the S&C area. Those concerns were both regarding the scarcity of expert personnel resources and their ability of top-down planning and decision making to implement and focus efforts towards important longer-term goals.

While the bottom-up process of assessing shorter-term resource needs and allocating available resources seems to be handled well using the monthly ICC meetings, the panel perceived less strength in longer-term planning and strategic vision. This becomes apparent in the lack of robust and detailed longer-term estimates of computing needs, including for the Upgrades. Also, the SCAP perceives insufficient managerial focus on prioritizing resources such that software modernization and forward-looking projects can proceed. IceCube has been largely unable to address longer-term software engineering needs and opportunities while frequently mentioning the increasing necessity of "paying back technical debt" accumulated over the past decade of developing its S&C systems. This perceived weakness could well affect IceCube's ability to fully and timely exploit their invaluable and unique data sets.

A finding from previous SCAP meetings struck this panel (that has predominantly new membership) as still being mostly valid, to quote from the 2018 SCAP report: *"As in previous occasions, the panel recognizes the unique culture of the collaboration and how it has been facilitating the ongoing stream of scientific discoveries. However, the panel continues to be concerned that some elements of this culture may hinder or negatively impact the ability of the computing enterprise to support the foreseen evolution of IceCube's science."*

While the panel acknowledges that IceCube has successfully addressed some of the previously expressed concerns, this SCAP report aims to address these issues with targeted recommendations for consideration of IceCube management and the collaboration. Like at the last meeting, *"[t]he panel maintains its previous global recommendation that more attention and effort should be devoted*

to developing and defining goals, predicting impact, and tracking progress and trends in all aspects of the computing enterprise, including effective and efficient coordination with the analysis/science activities, realistic accounting for how human resources are utilized, and tracking of data processing and computing resource consumption together with its impact on the science mission."

To renew this recommendation with some urgency seems appropriate and important as IceCube continues to succeed in demonstrating that innovation in data analysis and bringing to bear additional computing resources can push its limits of discovery and extend the reach of its dataset, even after more than a decade of operations.

Comments on the Impact of the COVID Pandemic

The SCAP commends IceCube for managing to continue operations and data analysis productively during the pandemic. The direct cost impact was estimated to be relatively small at less than $100K.

- The SCAP worries about indirect impacts, e.g., on the timing for a smooth transition between M&O CA periods. The renewal proposal was submitted but not yet reviewed, just two months before the end of the current 5-year period. A gap in funding would be unsustainable.
- The forced stop of Upgrade work at the pole etc., will likely have a medium-term impact on M&O, in particular if the Upgrade project needs to be delayed and/or re-baselined. We expect significant related side effects on S&C that IceCube should estimate and alert agencies about. IceCube should arrive at a good understanding of the cost and schedule impact on M&O and S&C when the agencies consider the COVID cost impact on the Upgrade project.

Addressing Immediate S&C Effort Shortfalls

IceCube made a case for requesting increased funding levels for adding five developer positions in the areas of workflows, simulation, visualization, reconstruction, and GPUs.

Comments The case appears strong, and the plan to address the immediate shortfall of effort and expertise is well developed and focussed. The SCAP strongly supports the request to add these 5 FTE to S&C support. The plan for where and how they will deploy this effort is sound. We caution that hiring five people is not an easy task that can be accomplished quickly.

- In making this request to agencies, it might help to frame the new positions in ways emphasizing the additional science opportunities; the need to address mounting risks related to long-term lack of maintenance and accumulated "technical debt;" and making the case where these additions are required to help the full exploitation of the Upgrade; in addition to strengthening ongoing maintenance and increasing the available core efforts.

Recommendation The increase in centrally-managed full-time professional Research Software Engineer effort at the level of 5 FTE is considered essential by the committee to fully and effectively exploit the enormous scientific potential of IceCube. We encourage IceCube management to continue making a strong case in NSF to secure the required funding.

S&C Long-Term Planning and Formulating Strategic Goals

The panel heard that IceCube expects continued improvements in its science reach with additional computing innovation. Detector resolution is limited by systematic effects like the details of ice properties. IceCube has a plan to improve these gradually, but M&O support for these efforts appears resource-limited. Innovative computational approaches, or even brute-force application of computational resources, could significantly improve IceCube as a scientific instrument.

Comments The SCAP believes that the scientific potential of IceCube and the wealth and richness of its scientific outcomes are ultimately determined by its ability to invest in improving and modernizing software and making efficient use of available computing resources.

- The SCAP stresses the central support roles for S&C operations for a multi-decade science program. IceCube needs to maintain S&C systems for periods of more than a decade. That requires top-down planning and focussing dedicated software development efforts towards long-term goals.
- To realize its ultimate scientific potential and to achieve the best possible resolution and discovery capabilities, IceCube may need transformative software changes. The SCAP heard about ideas and opportunities for innovative algorithms, for new systematic study procedures, for improving calibration and data processing for the upgraded detector, for opportunities to employ massive computing, like to "brute-force" tracking of individual photons, for innovations in the understanding of the underlying ice model, for overcoming limits in its ability to marshal computing resources, etc.
- Thus, it is essential for the S&C organization and leadership to bring limited resources to bear and find ways to optimize the impact of a relatively small group of software and computing experts on overall scientific productivity, including the long-term needs.
- The SCAP heard that software is not part of the Upgrades, and thus the needs for Upgrade related software efforts should be quantified separately. Needed efforts should not be forgotten or presumed to somehow come out of ongoing M&O efforts.
- We believe that now is the time to transition towards hardening the infrastructure for a long-term sustainable program of observations.
- It appears that at present, IceCube is not yet in a position to quantitatively justify a science reach prediction that could drive a program of software and computing R&D and directly connect technical goals and improvements in computing to science reach.
- Nor is there presently a computing and personnel resource-needs projection for a focussed R&D program that would target specific science goals. Instead, IceCube's effort and funding seem to be driven by what's available rather than justifiably needed to achieve an ambitious and forward-looking science program.
- Potential increases in science reach can be used to quantitatively justify computing resource needs, R&D investments and effort profiles needed for algorithm development including ML/AI, software development goals and required maintenance staffing, etc. This could be done in a way similar to how IceCube argued for and justified the ongoing and planned detector instrumentation upgrades.
- Such a plan and vision that outlines how growing S&C efforts would result in additional physics opportunities can rally the community and the agencies behind it to provide the required efforts and resources.

- A crisply defined R&D and staffing plan that outlines the work on algorithms, software, and computing, with quantitative indications and "flow-down" of personnel and computing resource needs, etc., from the science reach would become particularly important in connection with the Gen2 Upgrade. We saw elements of such a plan, but we did not yet see a coherent picture emerge. We believe that more decisively moving towards this ideal could significantly improve the effort and resources available via a mix of growth of the collaboration, growth in agency funding, and growth in international contributions.

Recommendations Create an S&C strategic plan that allows making a crisp case for IceCube's S&C support needs, flowing down from its science goals and capabilities, and that pro-actively plans for the significant investments needed in the areas of S&C developments, that formulates the science vision and the S&C technology goals to achieve them. The plan should help secure the required additional personnel resources and expertise.

- As part of this planning process, develop a set of quantitative science-reach goals that flow down to needed efforts and resources, thus justified by a detailed estimation of improvement goals. This should include Gen2 instrument, software, calibrations, etc.
- With this vision and strategy, determine how to create an organizational structure that will best support the mission

Reproducibility of Analysis Results

IceCube implemented a process to improve the reproducibility of analysis results that includes traceability of both software and data. The panel congratulates the IceCube team for developing and implementing procedures to improve the traceability and reproducibility of analysis results.

Recommendation Keep the momentum and use the effort invested in developing and documenting the ten-year data release to publish yearly updates to this comprehensive dataset.

Community engagement and participation

The panel acknowledges that members of the IceCube collaboration have participated in a number of proposals in the S&C area with a broader community, as presented during the review. This reflects positively the compelling nature of the IceCube science program and strength of the team. The following recommendation is to go beyond this to increase community engagement and leveraging of developments from outside of IceCube.

Recommendation The IceCube collaboration should be proactive with embedding themselves within the broader community of software development and ML. A proactive engagement with the broader community in S&C is important to position IceCube to affect the software tools that your science critically depends upon, benefit broadly from knowledge and technologies developed by other experiments, and to help in the professional development of your early career researchers (e.g. working collaboratively on modern open source tools from computer science and data science).

Next Meeting

The time since the previous SCAP meeting in June 2018 was rather long, given that several major S&C activities are evolving. IceCube should schedule SCAP meetings as regular checkpoints.

Recommendation We recommend having the next follow-up SCAP meeting in 12 to 18 months.

Addressing the Charge Questions

Charge Question (1)

The last SCAP included 10 recommendations for the project: have we adequately addressed, resolved, or are making suitable progress toward resolution on these recommendations?

Mostly **Yes**, except for recommendations 2018-1 to 2018-3.

Comments The SCAP believes that the appointment of a Global Computing Coordinator, as previously recommended, would allow IceCube to address the outstanding issues better.

We perceive that S&C lacks in the preparation of concrete actions and priorities. The job of a Computing Coordinator is not clearly assigned and defined, leaving it unclear who is responsible to overall collect and prepare S&C-related information, do the analysis of actions needed to address issues, and develop a concrete plan of action, to be brought to the ICC for approval.

We also perceive that S&C lacks a set of well-defined, quantified, and prioritized long-term science goals that would inform software development and computing resource needs. We miss the long-term vision that would shape short- to medium-term decision making. We think this is even more importantly needed now as the planning for the Gen2 Upgrade progresses.

The SCAP notes that S&C coordination and leadership changes were only gradual, took a while to implement, and did not go far enough and should be completed.

- The SCAP recognizes that IceCube made progress in implementing a process to assess bottom-up software and computing resource need. It was good to see an example of detailed resource needs estimation for the Oscillation analysis, which however accounts for only less than 10% of all computing resources. This should be adopted as a standard procedure to be followed by all the WGs.
- However, implementing a bottom-up resource needs assessment process will not solve many of the issues mentioned, particularly the slow progress on future planning and the perceived lack of concrete long-term strategic vision.
- We heard that the monthly ICC meetings work well to communicate needs and efforts and fairly allocate and track resources. An SOW-based approach was implemented where IceCube institutes pledge contributed resources, mostly physicist efforts.
- The SCAP is not convinced that this also effectively addresses other essential management functions, like **accounting** for provided volunteer efforts, **personnel planning** for needed expertise, and **creating a strategic vision** and **formulating technology goals** to address the

longer-term needs for Upgrade, to resolve "technical debt," and to successfully make a case for urgently needed additional efforts and expertise.
- The existing process seems insufficient to coordinate and prioritize and to help big-picture developments on the way, like decreasing the time from data to publication, or making progress to "rethink and redesign cleanly" the data processing/filters, etc.
- Indeed, we heard about examples of desired longer-term projects that move forward too slowly or not at all, like the re-working of filters and data flows from the pole, and the need to catch up on technical debt came up frequently. We also heard about the science community's interests in more substantial infrastructure and engineering support, particularly for ML/AI, all of which require planning across all areas and will need more central leadership, prioritization, and decision making.
- The panel concludes that stronger leadership and clearer lines of responsibilities would be advantageous, to deal with limited resources, to help establish and implement longer-term goals and strategies, to allocate precious resources to pay down "technical debt," thereby creating a sound technology basis for future transformative improvements, and to formulate the plans to achieve these goals.

Recommendations: IceCube should watch and assess quantitatively to what extent collaborators deliver on agreements, the level and efficacy of efforts that are actually being made available, and if this approach covers required roles and expertise sufficiently and in a sustainable way.

The SCAP renews the previous recommendation for a general Computing Coordinator.

- The Computing Coordinator should immediately start and lead the efforts to develop a long-term vision and strategy for IceCube S&C that addresses future IceCube scientific program's computing and data requirements and addresses the changing technology landscape. The support of IceCube management and agencies will be vital.
- This plan should include a strategy for addressing future hardware and technology choices, a mechanism to prioritize and fund advanced support and service developments, and an outline of necessary R&D.
- To implement this vision and strategy, determine how to create an S&C organizational structure within IceCube that augments the successful bottom-up communication structures and the ICC, which will best support the science mission. An organization with clearer boundaries may be beneficial.

Charge Question (2)

Are our core software resources supported centrally by Operations reasonable and in line with similar facilities?

While the quality of service and resources provided to the experiment seem excellent, the overall assessment is **No**.

Findings

- The data taking/online part of software operations appears to be sufficiently staffed (15 FTE)

and ready to tackle the Upgrade support. In comparison, the simulation/offline part suffers from quite severe under-staffing (7 FTE). Given the recent experience from other large data experiments, it is reasonable to expect a request for additional offline infrastructure support at this stage of the experiment (see below). Upgrades are introducing additional needs.
- IceCube showed several examples of important tasks, like resource utilization improvements, or software maintainability, waiting for person-power to make progress.
- Reliance on short-term graduate students and postdocs for "proof-of-concept" is acknowledged. Engineered products, however, require more expert support and will need to be delivered by the core software group.
- The group realizes that there is a significant waste of computing resources in the way some resources are being used, and on occasion, there are delays to publication.
- The core software team activities are organized and tracked using a ticketing system. Code sprints are organized regularly as a way to concentrate effort with the goal to close tickets and get things done. There are plans to in future track critical tickets in quarterly reports.
- In some cases, IceCube is losing on opportunities (e.g., cooperation with NVidia) due to not having the expert capacity to exploit them.

Comments

- Support for the collaboration from the operations team and core software support is successful. It keeps the systems afloat, and ongoing data processing and analysis efforts are productive at an adequate level. However, we saw that the core software support is really short on resources. Long-term developments to improve productivity (like re-working the filters) and forward-looking initiatives (like reconstruction software improvements) can not be supported adequately. IceCube's reliance on mostly short-term "volunteer" efforts from postdocs and grad students, while the software engineering experts are bound up with the day-by-day efforts, causes problems.

- The collaboration is lacking effort in many places in software to take advantage of possible improvements expediently. Several examples of delays from the identification of a possible improvement to deployment in operations were presented.

- There seems to be excessive reliance on in-kind student/postdoc effort to realize new ideas. A more substantial core software group should be more directly involved in directing any new central developments.

- The regular boot camps for software training that the core software team has organized for many years are working very well. The panel sees this as an effective tool to train the new students and postdocs to write code that is manageable and useful to the overall collaboration.

- IceCube has identified areas that require additional software engineering resources to address these shortcomings. Additional support for GPU usage, especially for scientist applications, ML & AI, is a reasonable request.

- The SCAP notes that other facilities, including the LHC and LIGO, have dedicated efforts targeting Upgrade S&C development, software modernization, essential R&D, etc., within the M&O efforts.

The requested additional 5 FTE should go a long way to resolve many of these issues. Adding them brings IceCub S&C to about equal the number of FTE of similar facilities.

Recommendations

- We recommend that the core M&O offline software effort, including the additionally requested FTE, be organized as a team effort, avoiding siloed expertise as much as possible.
- The core M&O software team should continue to seek further opening their activities to the broader community. Interact with experts from other facilities dealing with similar issues and actively participate in joint software projects where appropriate. Examples of areas where collaboration could be beneficial are the long-term evolution of the software framework (serialization libraries in particular, etc.) and the workflow management system (richer metadata, data popularity, storage tiering, etc.).
- The committee encourages the IceCube team to readily implement their plan to report on critical software tickets regularly and to use these reports as a tool to implement a transparent prioritization process. We think this will ultimately shorten the time-to-production for important issues.

Charge Question (3)

Is our computing model overall (software and hardware) sufficient to handle [the] additional […] scope of IceCube Upgrade and IceCube Gen2?

Likely **Yes**, but this still needs more attention.

Findings It was stated that the collaboration expects the IceCube Upgrade and IceCube Gen2 not to require any fundamental changes to the computing model (software and hardware).

Presently there are no quantitative estimates of resource needs for IceCube Gen2. IceCube Upgrade needs were presented as roughly a 10-20% increase over current annual production resources.

Comments The IceCube computing model has the architectural potential to scale in the SCAP's view, but it requires work to achieve the scaling and planning of how to approach this.

- We note there are significant legacy components, like, e.g., the persistency model, which pose scaling risks that need to be addressed.
- Experience of other facilities like LHC shows that every significant increase in scale requires dedicated effort and comes with new problems that need to be understood and addressed.
- IceCube indeed presented several areas for possible significant changes, e.g., with the opportunity to move more, maybe all, data recording off-pole, using more compression, and sending all data across the satellite link. These ideas represent changes in the computing model and involve risks that IceCube needs to evaluate carefully.
- The lack of a science flow-down to the hardware/software requirements makes prioritization decisions difficult considering the resource-limited environment.
- Not just the computing model but the whole software infrastructure and facility might require considerable changes; this should be studied. IceCube is largely running legacy software sys-

tems, while it has transitioned to be a major data processing and analysis facility that supports a large international and diverse scientific community. This bears risks that should drive the need for modernization to a well-structured and long-term maintainable software infrastructure.
- Assigning sufficient efforts to these areas seems hard to converge on in IceCube's bottom-up approach to decision making. IceCube needs more strategic planning, prioritization, and a clearly defined decision process.

The IceCube high-energy astrophysics program seems ready for a fundamental shift from showing evidence that sources exist to high volume detailed measurement of sources that do exist. IceCube is unique in this science, and the world is ready for it, given the high impact and expectations from MMA. This might provide a unique opportunity for IceCube for substantial growth in effort and resources.

The Upgrades should present opportunities to grow the collaboration, the M&O budget, and the international in-kind contributions. A quantitatively justified vision of what's scientifically possible could stimulate the national and international community, as well as the funding agencies, thus leading to significant growth in effort and resources.

Recommendation Perform a quantitative estimate of resource needs for Gen2. Resource needs should start from quantifying the possible improvement in science; from there, they should flow down to the planned instrumentation upgrades, as well as to any software, reconstruction, and calibration improvements; and then flow down further to needed increases in effort and computing resources to fully exploit the IceCube Gen2 science potential.

Appendix: Charge to Committee

Background The IceCube Software & Computing Advisory Panel (SCAP) was established in 2009 and is composed of experts in the fields of software development and scientific computing. The SCAP advises the IceCube Spokesperson(s) and Director on the following topics: on-line computing; on-line and off-line data processing and filtering; off-line computing facilities; and, simulations and analysis support. A written report is provided following each meeting.

Charge

- The last SCAP included 10 recommendations for the project: have we adequately addressed, resolved, or are making suitable progress toward resolution on these recommendations?
- Are our core software resources supported centrally by Operations reasonable and in line with similar facilities?
- Is our computing model overall (software *and* hardware) sufficient to handle additional the scope of IceCube Upgrade and IceCube Gen2?

Current Membership Lothar Bauerdick(Fermilab, chair), Peter Couvares (Caltec), Sridhara Dasu (UW-Madison), Peter Elmer (Princeton), Miron Livny (UW-Madison), Gonzalo Merino (PIC), Mark Neubauer (UIUC), Terry L Schalk (UCSD), Frank Würthwein (UCSD)

Agenda All meetings were conducted using zoom video conferencing.

`Wednesday 27 January 2021`

10:00 – 10:10	Welcome	(K. Hanson)
10:10 – 10:35	IceCube Science Highlights and Priorities	(I. Taboada)
10:45 – 11:05	IceCube M&O, the Upgrade, and Gen2	(K. Hanson)
11:15 – 11:30	Break	
11:30 – 11:45	SCAP replies #1: Coordination	(P. Desiati)
11:50 – 12:05	SCAP replies #2: Infrastructure	(B. Riedel)
12:10 – 12:25	SCAP replies #3: IceProd Scalability	(D. Schultz)
12:30 – 13:00	(Lunch) Break	
13:00 – 13:15	Data Processing	(N. Kurahashi Neilson)
13:20 – 13:35	IceCube Software Coordination	(A. Olivas)
13:40 – 13:55	IceCube Simulation Production	(J.C Diaz-Velez)
13:55 – 14:15	Discussion	

`Thursday 28 January 2021`

10:00 – 10:30	Open Discussion	
10:30 – 10:45	IceCube Data Releases	(I. Taboada)
10:50 – 11:20	IceCube Upgrade / Gen2 Software	(E. Blaufuss)
11:30 – 11:50	ML / AI Future Direction	(C. Kopper)
12:00 – 13:30	Break and SCAP panel private session	
13:30 – 14:00	Closeout / Draft Panel Report and Discussion	