DUNE Software and Computing Review Report

July 13th - 14th, 2023

Review Panel: Debbie Bard (NERSC), Simone Campana (CERN), Peter Elmer (Princeton University), Daniel Elvira (FNAL), Rex Tayloe (Indiana University), Yoshi Uchida (Imperial College London), Lynn Wood (PNNL), Frank Wuerthwein (UCSD)

Overall Finding:

The DUNE Collaboration and US DUNE Software and Computing (S&C) Consortium and its leaders have done well in the first steps to identify the elements required for a successful Phase-1 experiment. They have also done well to identify international contributions for hardware resources. However, more work is needed for management to better identify, organize, and assign people to the large number of software and computing tasks required.

Charge Question:

Does the Phase-1 DUNE Software and Computing system meet the needs of the collaboration?

a. What are the milestones necessary to establish the reliability and readiness of the DUNE Software and Computing system?

i. Are the requirements and timing of each milestone clearly defined?

b. What computing elements have international partners agreed to support?

Findings:

Milestones have been presented and are identified in the Phase-1 milestone document and in review talks. They are broken down into the three groups corresponding to the three major detector milestones. In the milestone document, the milestones and timelines are listed. However, the detailed timing, requisite activities and dependencies are not described.

Comments:
US DUNE S&C should evaluate the strategy of contributing to most tasks and consider prioritizing support for US DUNE's high-priority tasks while seeking international support for others. US DUNE S&C may want to consider identifying tools, services, or subprojects needed by DUNE that present the opportunity for international collaborators to support completely.

It is difficult to understand how the specific physics goals of the collaboration map onto the DUNE S&C systems technical requirements. One suggestion is to align software and computing activities and learn from physics needs by defining joint milestones related to the production of a "Physics Book" or other planned benchmarking of analyses before data taking.

Recommendations:
1. US DUNE is encouraged to clearly delineate the US responsibility in the DUNE S&C system beyond hardware contributions. For software activities that may also be supported by international partners, the specific features and attributes that are a sole US responsibility should be identified. This should be done by December 2023.

Charge Question:
For US-supported elements of the Phase-1 DUNE Software and Computing system:

a. What are the requirements of the Phase-1 elements? When must those requirements be met to achieve the Phase-1 needs?
b. What is the current technical readiness?
c. What future developments are necessary to meet the Phase-1 DUNE scientific needs?
d. What opportunity is present in the development of the element?
e. What risks are present in the development of the elements and the mitigation plan to reduce risks?

f. What is the prioritization of the elements?

Findings:
The US is expected to provide 50% of the total computing capacity in compute, disk storage, and tape archival. This expectation was essentially met during the last year.

The existing DUNE software and computing is expected to evolve to a production system that meets the physics requirements for operations without beam by 2028, and with beam by 2031. The first high profile measurement with the former could be a supernova as early as day 1 of no-beam operations. The first high profile measurement for the latter could be the neutrino mass ordering as early as six months after the start of physics data taking with beam in 2031.

The project makes heavy use of common tools from across FNAL and HEP more generally, using DUNE effort to adapt these tools to DUNE needs.

DUNE is currently using the ART framework which is unable to meet the Phase-1 physics goals of the collaboration. The milestone for a replacement framework capable of processing one detector module is first availability in 2027, one year before data taking. Backwards compatibility between ART plugins and the replacement is desired but not guaranteed. R&D effort towards a replacement framework has begun but the development into a mature collaboration-supported element has not.
The US DUNE S&C Consortium is presently working at the level of ~15 FTE. Most of the existing funding is set to expire within the next year. The project estimates that 25 FTEs are needed in 2024, and presented an effort envelope for the next decade or more that is essentially flat.

Comments:
US DUNE S&C Consortium is capable of processing, simulating, and analyzing the existing prototype DUNE detector across a globally distributed computing system, including NERSC. This indicates that they have significant functional mature software, algorithms, services and

computing infrastructure in operations today.

The DUNE computing, disk storage, and tape archival storage planning is adequate for the stage the project is in. The expectation to evolve this plan as knowledge about input and the computing model itself evolves are as expected.

Using common tools from across FNAL/HEP and adding DUNE adaptations is a good strategy for sustainable software development. However, dependencies from and requirements for external software deserves more attention. US DUNE S&C may want to work with external developers of the tools DUNE depends on to understand and negotiate plans and deadlines and, if necessary, make the in-kind contributions to ensure tool capabilities are available for DUNE on time.

There are essential tools, such as visualization and end-analysis tools, that do not currently appear in DUNE S&C planning. It would be beneficial to conduct a complete survey of required software and develop a plan to identify responsible institutions within DUNE, to ensure all requirements are considered.

Beyond the functionality of the existing systems and the integration validation via data challenges, it was not clear how physics goals map onto technical requirements for the algorithms, software, and systems and their development. We are thus unable to answer the charge question regarding requirements for the Phase 1 minimal system, and what technical developments are necessary to meet the opportunities for early physics in either 2028 or 2031.

User developed software (e.g. physics analyses) that would run on a post-reconstruction end-analysis framework is presented as an area where bottom-up contributions from the collaboration are awaited, rather than being defined or provided for as part of the framework. However, the way in which such contributions should be made, including protocols to follow and the essential software interfaces (e.g. to use common math tools) which must be adhered to, should be clearly defined as part of the DUNE S&C Consortium interface with the DUNE Physics Consortium.

The development of the DUNE S&C framework appears to be on the critical path to delivering a system that can facilitate the physics goals of the collaboration. The current milestones do not provide time for redevelopment of elements that will rely on the DUNE framework, or refinement from physics analysis requirements and machine learning workflows that may arrive when DUNE collaborators begin widespread use.

If funding for the increase in effort is delayed, US DUNE S&C expects a large bulge in staffing in a later fiscal year needed to catch up and deliver the current system in time for the start of Phase 1 data taking in 2028. This planned response to increase effort in later years may not be able to meet the milestones and timelines necessary to deliver the planned system. The US funding agencies are encouraged to work with US DUNE to meet the project’s software and computing staffing needs and US DUNE S&C Management are encouraged to consider how to respond to effort shortfalls based on the prioritized list of projects, in a way that minimizes the impact to the DUNE physics goals.

Recommendations:

2. We recommend the project re-prioritize or add effort such that a flow down, of how DUNE Phase-1 physics goals flow into technical requirements on software & computing, can be produced. Such requirements should become the basis for future reviews, and be refined as progress is made. US DUNE S&C should clearly identify the minimal functioning system required to support DUNE physics needs by November 2023.

3. US DUNE S&C management should develop a plan for how to respond to future software tasks and projects that arise but are not currently identified, by February 2024.

4. US DUNE S&C Management should prioritize developing the technical

requirements of the software framework and establish a detailed plan to develop a framework capable of meeting the collaboration's physics goals, with sufficient time that integration and iteration on the development cycle can be successfully completed. The plan for any necessary remaining R&D and framework-focused software development should be identified by November 2023.

5. US DUNE S&C has done a good job of identifying upcoming total staffing needs and should develop a more detailed plan to meet the staffing needs by January 2024. This plan should include contingencies to respond to various staffing levels.

Charge Question:
How are the US-supported computing elements managed?

a. How is progress of the elements tracked and assessed?
b. How does the US DUNE Software and Computing management respond to opportunity and risk?
c. Are US DUNE Software and Computing responsibilities well-coordinated with international collaborators' responsibilities and other related computing efforts?
d. Are management lines of authority and available resources clear?

Findings:
US DUNE S&C has worked hard to develop relationships with international partners, especially the close effort with the UK which can serve as a model for other relationships.

We were provided an impressive list of software that DUNE is leveraging from FNAL, the HEP community, and the broader scientific research community and industry.

Data Management and Production are organized in teams with members from all DUNE regions (US and international), with large use of collaborative tools. US and non-US technical leads in the DUNE S&C

Consortium both report to the consortium lead.

There is currently no formal agreement for computing institutional commitments. For hardware there is an informal pledging process in place and it is planned to have a more formal oversight through the Resources Review Board. For personpower, there is a table maintained by resource coordination recording the FTE effort in each area including software and computing.

Risks linked to dependencies from external (non-DUNE) components are partially understood and so is the mitigation strategy. A small number of components are under the sole responsibility of non-US DUNE institutes.

Resource needs estimates lack information to evaluate need-resource gaps, as well as the impact of R&D.

Comments:
Leadership of the US DUNE S&C efforts have done well to coalesce the US activities to this point. We emphasize the importance of developing an appropriate level of formal management for software and computing, including dependencies, decision-making and risk characterization, to enable the experiment to maintain an appropriate level of control and overview.

The US DUNE S&C Consortium is encouraged to explore and identify the most efficient funding method and work with the funding agencies to develop the appropriate plan.

We suggest that the DUNE S&C Consortium add to resource projection plots the impact of R&D in reducing system needs and mitigating risks, as well as computing budget assumptions as a baseline to measure potential resource gaps. Such additions would give future reviews a sense of the level of certainty at which the needs are known, and indicate the impact of R&D on these projections.

Recommendations:
6. A more formal management structure appropriate to a collaboration of this size should be developed by October 2023. This structure should ensure progress tracking, assessment of opportunities and risks, coordination and communication with other US DUNE consortia and international collaborators, and clearly define management lines of authority. Additional responsibilities and needs are detailed in Recommendations to other questions.

Charge Question:
What is the strategy for training, mentoring, and retaining a diverse pool of early-career scientists with sufficient technical expertise to maintain the US-supported elements over the life of the DUNE collaboration?

Findings:
An innovative and creative scheme is presented for leveraging the needs of DUNE and the opportunities it represents in order to recruit and train a broad range of early-career scientists and students. The strategy for mentorship, especially of a one-to-one nature, and retention in general, appears to be somewhat less clearly fleshed-out.

The inclusion of smaller institutions to leverage their expertise was noted as a distinct advantage to the US DUNE S&C Consortium.

There is a core team of experts working on critical areas of software and computing that are now hired through the DOE grant. There is currently no long-term plan to retain this expertise.

Spending that is explicitly targeted towards specific younger cohorts is proposed: summer internships; undergraduate research positions; graduate student secondments to labs; postdoc ‘pipeline’; and funded recruitment visits.

Comments:
The relatively small amounts of funding that are sought to support

engagement with students and early career scientists will have a very high positive impact, both in terms of PIER as well as for the DUNE project in general, and should be pursued vigorously.

A survey of needed expertise and skills would be valuable in long-term success of the project.

The PIER Plan includes many valuable descriptions of principles and implementation ideas, but as a document, contains much room for improvement in style and organization. The individual institutional contributions are somewhat varied in quality, and it would be helpful if the institutions, and DUNE as a whole, could improve their plans by learning from each other's contributions.

Recommendations:
7. Bring in systematic tracking mechanisms for the progression of individuals at all stages; from undergraduates who are present at recruiting visits and research positions, to graduate students and postdocs who receive training, to quantify the impact of the strategies being pursued and to allow them to be improved. Similarly, the tracking of the number of active instructors, and the development of new instructors, should be instituted such that their number and quality grows in proportion to the needs of the collaboration, rather than burdening the same individuals more and more. The means to track both of these should be identified by January 2024.

8. In event of a funding shortfall, US DUNE S&C should prioritize maintaining the support for high impact PIER and engagement activities.

Charge Question:
What is the process and timeline to develop the DUNE collaboration's software and scientific data Public Access Plan?

Findings:
A draft plan for Public Data Access was presented and appears to be well planned and is based on other experiments' experiences and practices

(e.g. microBooNE).

Comments:
Public data access was presented with data at 4 levels. The plan for level 1-3 data is solid, well founded and looks consistent with the FAIR principles and, likely, new DOE policy. It is also good practice for good science. Level 4 is a bit different in that it is hard to understand the need for the public to have data at the rawest level. The example from microBooNE seems a good way to satisfy the requirements, without too much wasted collaboration effort.

A near term opportunity for exploring the draft plan in practice would be the ProtoDUNE data taking and analysis.

Recommendations:
9. This data access plan should be further pursued to refine and to finalize with the DUNE collaboration. The collaboration should work to ensure that the policy will be consistent with DOE guidance on the Nelson memo on public data access[1] prior to February 2024.

[1] https://www.whitehouse.gov/wp-content/uploads/2022/08/08-2022-OSTP-Public-access-Memo.pdf