

DUNE Software and Computing Challenges

Paul Laycock for the DUNE Collaboration^{1,*}

¹Brookhaven National Laboratory

Abstract. The DUNE experiment will begin running in the late 2020's. The goals of the experiment include 1) studying neutrino oscillations using a beam of neutrinos from Fermilab in Illinois to the Sanford Underground Research Facility, 2) studying astrophysical neutrino sources and rare processes and 3) understanding the physics of neutrino interactions in matter. The DUNE Far Detector, consisting of four 17 kt LArTPC modules, will produce "events" ranging in size from 6 GBs to more than 100 TBs, posing unique challenges for DUNE software and computing. The data processing algorithms, particularly for raw data, drive the requirements for the future DUNE software framework documented here.

1 What is DUNE

DUNE [1] is both a neutrino oscillation and an astrophysics experiment, the layout of the experiment is shown in figure 1. It will consist of a high intensity neutrino beam, a Near Detector to measure the unoscillated neutrino flux and neutrino scattering physics at Fermilab, and a Far Detector complex 4850 ft below the surface in the Homestake mine at the Sanford Underground Research Facility in Lead, South Dakota. The Far Detector will be sensitive to multiple flavors of neutrinos from the beam, to cosmogenic neutrinos, and to other rare processes [2].

The DUNE beamline will deliver 10 microsecond pulses of muon neutrinos every 0.89 seconds. The beam originates from 120 GeV protons with 1.2 MW power. This is high enough to generate large numbers of interactions/pulse in the Near Detector. Interaction rates at the Far Detector are much lower and require a very large sensitive volume with fiducial mass > 10 kt. The experiment plans to measure the CP-violation sensitive oscillation process $\nu_\mu \rightarrow \nu_e$ for both neutrino and anti-neutrinos at the few % level, unprecedented precision for a neutrino experiment. Clean discrimination of ν_e interactions from ν_μ interactions imposes stringent requirements on the detector technology. The DUNE Far Detector will consist of Liquid Argon Time Projection Chambers (LArTPC) read out both through detection of the ionization and scintillation light. Due to their size and the low background environment deep underground, the Far Detector will also be sensitive to atmospheric neutrinos, to supernovae occurring in our galaxy, to BSM processes, and, at low readout threshold, to solar neutrinos. The Near Detector will be composed of a suite of subsystems designed to fully characterize the flavor and energy of neutrino interactions, monitor the neutrino flux, and minimize systematic uncertainties by minimizing differences between neutrino interactions in the Near and Far Detectors.

*e-mail: laycock@bnl.gov

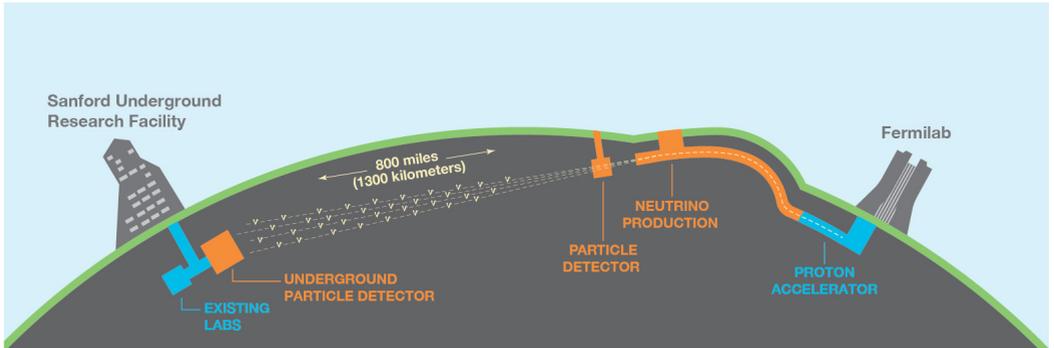


Figure 1. Schematic of the DUNE experiment concept, showing Fermilab (right) producing a high intensity neutrino beam that passes through two particle detectors, one close to the beam source and the other 800 miles away at the Sanford Underground Research Facility.

2 Why is DUNE different?

The DUNE Far Detector will eventually be comprised of four very large 17 kt modules each with a fiducial mass of at least 10 kt. Schematics of one Far Detector cryostat and module are shown in figure 2. The size of the sensitive volume combined with the readout times needed to extract physics signals drive the fundamental DUNE software and computing challenges. The Near Detector will be much more similar to modern fixed-target and collider experiments, with the major issue being combination of data from a diverse set of detector technologies in the presence of multiple beam and cosmic ray interactions in a single beam spill.

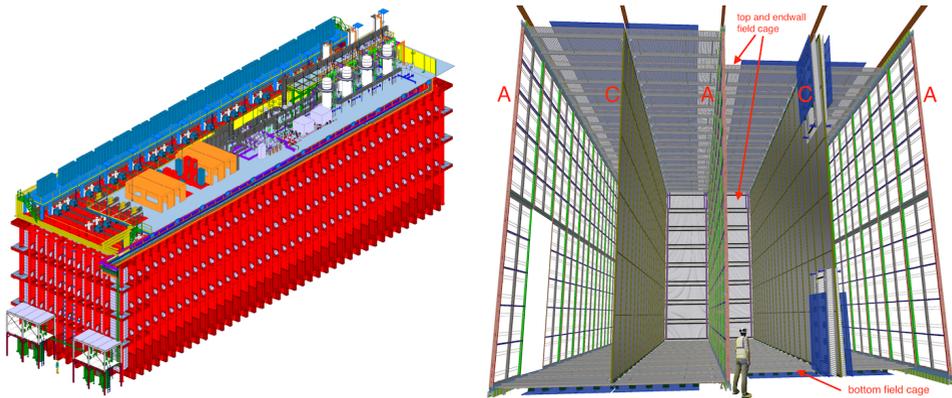


Figure 2. Left) A Far Detector cryostat that houses a 10 kt Far Detector module. The figure of a person indicates the scale. Right) A 17 kt DUNE Far Detector SP module, showing the alternating 58 m long (into the page), 12 m high anode (A) and cathode (C) planes that are 3.5 m apart, as well as the field cage that surrounds the drift regions between the anode and cathode planes. The blank area on the left side was added to show the profile of a single Anode Plane Assembly (APA).

2.1 DUNE “events” are large

Liquid Argon TPC technology is capable of measuring particle trajectories on sub-cm scales over volumes of the order of 10^4m^3 . Electrons from ionization caused by charge particles are drifted over distances of several meters to multiple wire planes or pixel detectors and read out in 500 ns time bins for 5-10 ms. The final DUNE Far Detector modules are expected to have 1.5M readout channels. Technologies currently being considered do not use gas-amplification, leading to a requirement of sensitivity to a few thousand electrons/time slice. A single beam induced readout of a Far Detector module will be 6 GB of uncompressed data. Sensitivity to the details of neutrino events and rare astrophysical sources will require very careful application of any real-time data reduction. To date, lossless zero-suppression has been shown to reduce readout sizes by a factor of 2-3.

The potential of observing supernovae and low energy astrophysical neutrinos adds additional challenges. Astrophysical neutrinos from supernovae and the sun have energies close to the lower energy threshold for detection and will require more sophisticated signal processing than the higher energy beam interaction neutrinos. Supernovae are extremely rare, occurring a few times/century in our galaxy. The DUNE detector could expect to see thousands of low energy anti-neutrino interactions from a supernova in our galaxy. Those interactions could occur over a period of up to 100 s, requiring storing and processing 20,000 standard readouts. A full uncompressed readout of one 17 kt Far Detector module would be 115 TB in size. In addition to the processing and storage requirements, a supernova candidate will require maintenance of coherence across its full geometric and time span.

Once recorded, these data will need to be transferred from the Homestake mine in South Dakota to Fermilab. Normal beam and calibration operations will generate data rates of a few GB/s but timely access to a supernova readout will require rates of 10 GB/s. Including calibrations, beam events, and astrophysical sources, DUNE will record between 10-30 PB of data/year, which represents a modest computing challenge in comparison to the high luminosity LHC. However, it is not obvious that the large nature of DUNE "events" is well-aligned with the combination of software frameworks and computing architectures used in HEP today.

2.2 Signal Processing

Data from large LArTPC's is much more uniform than data from complex collider detectors. Each of the Anode readout planes shown in figure 2 consists of three wire planes each with different wire orientations, as shown in figure 3. The waveforms read out from these three wire planes constitute the bulk of the Far Detector raw data.

Thanks to previous neutrino experiments and the prototype DUNE experiment at CERN, ProtoDUNE [3], the algorithmic treatment of waveform data to reconstruct the original distribution of ionization, referred to as signal processing, is well advanced [4, 5]. The main challenge for a software framework comes from handling the large quantities of data which will be 6 GB of uncompressed data for a nominal Far Detector readout. Coupled with the memory needed to store intermediate and output data products, it is clear that this nominal "event" size is already problematic compared to the 2 GB RAM found on an average HEP compute node, and that is before considering the extended readout case needed for supernova physics. Thankfully, as each APA is a physically distinct region of the Far Detector, these large events can be naturally broken down into smaller chunks of data. Thanks to its size, a supernova "event" also needs to be sliced in the time dimension with care taken that it must be possible to study the overlaps between these time slices. Provided that the interdependence of these spatial and time slices can be neglected for the most compute-hungry algorithms, this

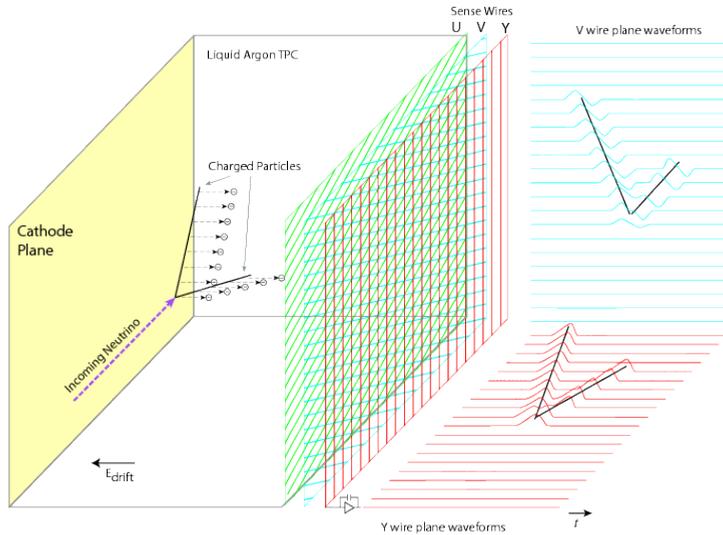


Figure 3. Diagram illustrating the signal formation in a LArTPC with three wire planes. Here the Cathode plane is on the left, with an Anode plane consisting of the three wire planes shown centrally. The waveforms induced on the wire planes by the charged particles produced by the incoming neutrino are shown on the right; for simplicity, the signal in the first U induction plane is omitted in the illustration.

problem looks well-suited for massively parallel computing architectures. Equally, the typical high-throughput compute architectures used in HEP should be amenable to this work so long as an appropriate software framework that can efficiently orchestrate DUNE workflows can be realised.

2.3 DUNE Reconstruction

Following signal processing, the Far Detector data effectively forms images as shown in figure 4. With a suitable amount of innovation and physics insight, such data is very compatible with novel applications of industry-standard image-processing techniques and is also well-suited for massively parallel computing architectures. The difference in data volumes between the raw waveform data and the data after signal-processing, depending on noise suppression, will likely be somewhere between one and two orders of magnitude.

3 DUNE Software Framework Requirements

Modern HEP software frameworks are in the process of addressing the increasing heterogeneity of computing architectures that are redefining the computing landscape, work that is closely followed by the HEP Software Foundation [6, 7]. The DUNE collaboration is keen to leverage this work to minimise the development effort needed for its own software framework. Nevertheless, given the unique challenges that DUNE data pose, the collaboration assembled a task force to define the requirements for its software framework. Some of their findings are detailed in the following.

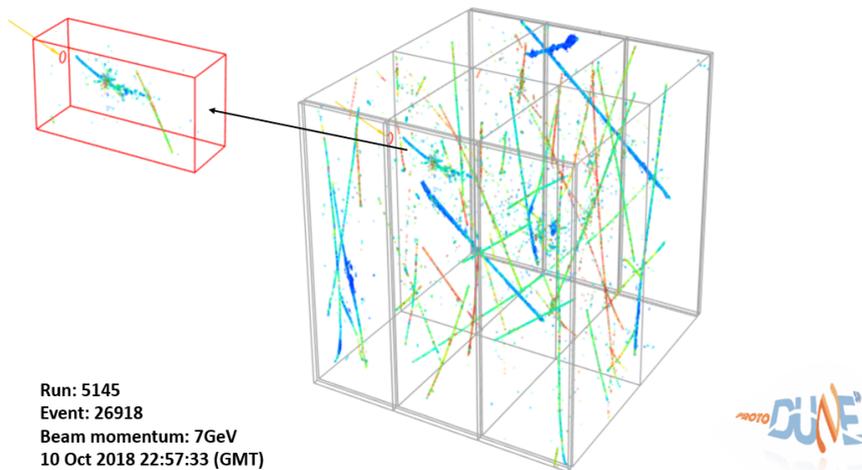


Figure 4. The ProtoDUNE SP detector (gray box) showing the direction of the particle beam (yellow line on the very far left) and the outlines of the six APAs. Cosmic rays can be seen throughout the white box, while the red box highlights the beam region of interest with an interaction of the 7 GeV beam. The 3D points are obtained using the Space Point Solver reconstruction algorithm.

3.1 Data and I/O layer

Following the discussion on raw data processing, the DUNE software framework must have the ability to change its unit of data processing with extreme flexibility and ensure that memory is treated as a precious commodity. It must be possible to correlate these units to DUNE data-taking "events" for exposure accounting, and experiment conditions. Partial reading of data is another strong requirement needed to keep memory usage under control. There must be no assumptions on the data model, the framework must separate data and algorithms and further separate the persistent data representation from the in-memory representation seen by those algorithms. DUNE data is well suited to massively parallel processing compute architectures, so the framework will need to have a flexible I/O layer supporting multiple persistent data formats as different formats may be preferred at different stages of data processing. The sparsity of DUNE data also imply that compression will play an important role in controlling the storage footprint, especially for raw data, and the sparsity of physics signals further emphasise support for data reduction (skimming, slimming and thinning) in the framework. The framework needs to be able to handle the reading and writing of parallel data streams, and navigate the association (if any) between these streams with a critical need to allow experiment code to mix simulation and (overlay) data.

3.2 Concurrency

Many aspects of DUNE data processing are well suited to concurrency solutions and the framework should be able to offload work to different compute architectures efficiently, facilitate access to co-processors for developers and schedule thread-safe and non-thread-safe work. It must be possible to derive the input data requirements for any algorithm in order to define the sequence of all algorithms needed for a particular use case, and it must be possible to only configure those components required for that use case.

3.3 Reproducibility and provenance

As previously stated, the framework must ensure that memory is treated as a precious commodity, implying that intermediate data products cannot occupy memory beyond their useful lifetime. Nevertheless, reproducibility is a key requirement of any scientific tool and the framework must provide a full provenance chain for any and all data products which must include enough information to reproduce identically every persistent data product. By definition, the chain will also need to include sufficient information to reproduce the transient data passed between algorithm modules, even though the product is not persisted in the final output. It is also highly desirable that the framework broker access to random number generators and seeds in order to guarantee reproducibility. All of the preceding considerations imply that the software framework will need a very robust configuration system capable of handling the requirements in a consistent, coherent, and systematically reproducible way.

3.4 Analysis

Machine learning techniques are already heavily used in analysis of ProtoDUNE data, see for example [8, 9] in these proceedings. The framework should therefore give special attention to machine learning inference in the design, both to allow simple exchanges of inference backends and to record the provenance of those backends and all necessary versioning information. Finally, the framework should be able to work with both Near Detector and Far Detector data on an equal footing, and within the same job.

4 DUNE computing timeline

The DUNE collaboration is in the process of writing a conceptual design report (CDR) that will describe the key elements of DUNE software and computing, including the software framework and computing model. The DUNE collaboration is taking the general approach that existing exascale solutions will be evaluated for adoption against the computing requirements laid out in the CDR. DUNE plans to lean heavily on existing exascale solutions, with Rucio [10] already being used as the data management system for ProtoDUNE. Some elements of the requirements are well understood, not least the networking requirements to transfer the Far Detector data from the Sanford Underground Research Facility to Fermilab for offline data processing. Other elements of the computing model are still under development, such as the network requirements to distribute data from storage elements (SEs) to compute elements (CEs) for analysis. At this time, DUNE has added distributed computing resources across the globe using the same infrastructure as that of the LHC experiments, i.e. the WLCG [11], OSG, GlideInWMS, and HTCondor. DUNE computing experts will work together with colleagues in the LHC community and beyond to collaborate on solutions for those challenges that have not yet been solved.

The DUNE Computing Model is based upon a flatter architecture of compute elements compared with the original tiered structure of the LHC experiments. One aspect of the computing model is the late binding between workload management and CEs where workflows are sent to available CE resources regardless of data location. This approach, which allows sites with ~no storage to provide effective compute, is dependent upon reliable, consistent-bandwidth networking to enable streaming of data from the closest SE¹. This model has been used by DUNE to project its resource needs in terms of CEs and SEs, but a full assessment of networking needs beyond the initial storage and distribution of raw data from the DAQ has

¹The "closest SE" here refers to distance when viewed from the network perspective.

yet to be developed. Nevertheless, this model was used successfully in the processing and simulation of the ProtoDUNE I beam data which has now been published in the ProtoDUNE-SP Performance paper [12]. Indeed, the operation of ProtoDUNE II, scheduled to begin in summer 2022, will be used to stress test the solutions laid out in the CDR well in advance of DUNE data-taking, and help to identify further challenges that lay ahead.

References

- [1] B. Abi et al. (DUNE), JINST **15**, T08008 (2020), 2002.02967
- [2] B. Abi et al. (DUNE) (2020), 2002.03005
- [3] B. Abi et al. (DUNE) (2017), 1706.07081
- [4] <http://larsoft.org>
- [5] X. Qian, B. Viren, and C. Zhang, <https://www.phy.bnl.gov/wire-cell/>
- [6] J. Albrecht et al. (HEP Software Foundation), Comput. Softw. Big Sci. **3**, 7 (2019), 1712.06982
- [7] P. Calafiura et al. (2018), 1812.07861
- [8] H. Yu, Z. Dong, K. Knoepfel, M. Lin, B. Viren, and K. Yu, CHEP 2021, these proceedings
- [9] M. Rossi, S. Vallecorsa, CHEP 2021, these proceedings
- [10] M. Barisits et al., Comput. Softw. Big Sci. **3**, 11 (2019), 1902.09857
- [11] I. Bird et al. (2014)
- [12] B. Abi, A.A. Abud, R. Acciarri, M. Acero, G. Adamov, M. Adamowski, D. Adams, P. Adrien, M. Adinolfi, Z. Ahmad et al., Journal of Instrumentation **15**, P12004–P12004 (2020)