Computing for the DUNE Long-Baseline Neutrino Oscillation Experiment

Heidi Schellman for the DUNE collaboration^{1,2,*}

¹Department of Physics, Oregon State University, Corvallis OR, 97331

²Scientific Computing Division, Fermilab, Batavia, IL, 60555

Abstract. This paper is based on a talk given at Computing in High Energy Physics in Adelaide, South Australia, Australia in November 2019. It is partially intended to explain the context of DUNE Computing for computing specialists.

The Deep Underground Neutrino Experiment (DUNE) collaboration consists of over 180 institutions from 33 countries. The experiment is in preparation now, with commissioning of the first 10kT fiducial volume Liquid Argon TPC expected over the period 2025-2028 and a long data taking run with 4 modules expected from 2029 and beyond.

An active prototyping program is already in place with a short test-beam run with a 700T, 15,360 channel prototype of single-phase readout at the Neutrino Platform at CERN in late 2018 and tests of a similar sized dual-phase detector scheduled for mid-2019. The 2018 test-beam run was a valuable live test of our computing model. The detector produced raw data at rates of up to 2GB/s. These data were stored at full rate on tape at CERN and Fermilab and replicated at sites in the UK and Czech Republic. In total 1.2 PB of raw data from beam and cosmic triggers were produced and reconstructed during the six week test-beam run.

Baseline predictions for the full DUNE detector data, starting in the late 2020's are 30-60 PB of raw data per year. In contrast to traditional HEP computational problems, DUNE's Liquid Argon TPC data consist of simple but very large (many GB) 2D data objects which share many characteristics with astrophysical images. This presents opportunities to use advances in machine learning and pattern recognition as a frontier user of High Performance Computing facilities capable of massively parallel processing.

1 Introduction

The DUNE 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 Homestake mine in Lead, South Dakota, 2) studying astrophysical neutrino sources and rare processes and 3) understanding the physics of neutrino interactions in matter. I will concentrate on the neutrino oscillation and supernova capabilities of the experiment and the ways that they drive computing.

*e-mail: heidi.schellman@oregonstate.edu

The neutrino beam from Fermilab will consist almost entirely of muon-type neutrinos when produced. Neutrinos are known to come in (at least) 3 flavors which can be distinguished by their interactions - electron type neutrinos produce electrons when they interact via charged currents; muon neutrinos, muons; and tau neutrinos, tau particles. But these flavors do not correspond to fixed mass states. All 3 flavors of neutrinos are mixtures of mass states, much as light polarized in the x direction can be considered a superposition of x' and y' polarizations along alternate axes rotated by 45 degrees. When neutrinos propagate through space, it is the mass state that sets their wavelength and if the neutrino goes far enough, the multiple mass states corresponding to the initial flavor state will get out of phase. When the mixture is later probed about its flavor, it may give a different answer than the neutrino that started out. This phenomenon is neutrino oscillation and has been shown to exist in multiple experiments since it was first confirmed in 1998[1].

Figure 1. Illustration of the neutrino flavor and mass states. The mass states are a superposition of the flavor states. Courtesy the particlezoo.net.

DUNE, in particular, wishes to understand the conversion of muon neutrinos created in Illinois into electron neutrinos at a far detector (FD) in the Homestake mine in South Dakota and compare that conversion rate between neutrino and anti-neutrino beams. The location of the FD and energy of the neutrino beam were chosen to maximize the oscillation effect. A difference in the conversion rate for neutrinos and anti-neutrinos could be evidence for matter-antimatter asymmetry in the neutrino sector, a phenomenon called CP violation.

To make these measurements, we need to be able to distinguish electron neutrino interactions appearing in the muon neutrino beam from the dominant muon neutrino interactions one would expect in the absence of oscillations. Doing this requires a very large detector, as neutrino interactions are intrinsically rare, but an extremely fine grained one as well. Noble liquid time projection detectors, which read out large transparent volumes of liquid by drifting electrons from interactions to charge detectors through strong electric fields, have the needed capabilities of extremely large scale and fine-grained resolution. The proposed DUNE far-site detector will instrument four $14 \times 12 \times 58$ meter volumes of Liquid Argon with readout granularity of 0.5 cm. The detectors will be located 4850 ft below the surface to lower the rate of cosmic rays traversing the detector by orders of magnitude and thus allow sensitivity to very low energy solar and astrophysical neutrinos as well as the higher energy neutrinos produced at Fermilab.

Figure 2. Electron neutrino appearance signal (top) and background (bottom) as seen in the ArgoNeut experiment[2]. In the true appearance signal, an electron is seen emerging from the primary vertex and then showering. In the background interaction, a muon neutrino enters and produces a final state muon and photons, which propagate some distance before showering.

Figure 3. Left) A far detector cryostat that houses a 10 kT FD module. The figure of a person indicates the scale. Right) A 10kt DUNE FD SP module, showing the alternating 58 m long (into the page), 12 m high anode (A) and cathode (C) planes, 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).

The neutrino beam from Fermilab will be pulsed approximately once/second 24 hrs/day during running periods with of order 15 million pulses per year. Because neutrinos interact extremely rarely, we expect to detect of order 7,500 neutrino interactions/year in each of 4 10 kT detector modules located at the FD site in South Dakota.

Construction of the detector halls and infrastructure for the large 10 kT fiducial volume FD modules is starting now, as are design and construction of detector readout modules. A full technical design report (TDR) for the program has recently been completed and is available in references [3–6]. The DUNE neutrino oscillation experiment will receive beam late in this decade with commissioning of the data acquisition systems for the first far detector module expected to start in 2025-26.

2 ProtoDUNE tests at CERN

Building an experiment of this size requires an extensive period of prototyping. The Argoneut[7], MicroBooNE[8] and ICARUS[9] collaborations have demonstrated the capabilities of large liquid Argon time projection chamber (TPC)s for neutrino detection on scales between 1 and 500 ton fiducial mass. In preparation for the DUNE experiment, a campaign testing proposed DUNE components in 700 ton detectors in the EHN1 hadronic test-beam was launched at CERN in 2018. Both single-phase and dual-phase prototypes were constructed and tested.

2.1 ProtoDUNE-SP

The ProtoDUNE-SP experiment began taking data at CERN in late 2018. ProtoDUNE-SP uses single-phase technology where ionization electrons are collected directly from the liquid argon. The readout system consists of Anode Plane Assemblies (APA)s which each have 3 layers of wires arranged in different directions. Each layer contains 800-1200 wires spaced 0.5 cm apart. Electrons drift from the original interaction in the Argon, through a strong electric field, to the wire planes and induce signals. The location in the plane of hit wires gives one coordinate, the time the signal takes to drift to the wire from the original interaction measures a second coordinate. The third coordinate is derived by combining information from overlaps of signals in the 3 different wire layers. Signals are amplified electronically and then digitized. Figure 4 illustrates the operation of a generic liquid argon time-projection chamber (LArTPC).

Figure 4. Diagram from [10] illustrating the signal formation in a LArTPC with three wire planes [11]. For simplicity, the signal in the first U induction plane is omitted in the illustration.

The ProtoDUNE-SP detector consists of a 700 ton volume of liquid argon with a cathode plane in the center and 3 APAs mounted on each edge of the liquid volume. The drift distance is 3 m with a nominal voltage of 180kV across that distance. Each APA has 2560 channels

and each channel reads out a 12-bit analog-to-digital converter (ADC) every $0.5\ \mu\text{sec}$. For ProtoDUNE-SP the readout time appropriate for a 3 m drift was set to 3 msec, resulting in 6000 12-bit samples per channel. The total data size for six APAs is thus 140 MB with additional header and data from photon and external tagging systems bringing the nominal event size up to around 180 MB. Lossless compression of the TPC readout was implemented in the data acquisition system, resulting in a final compressed event size of around 75 MB.

The test-beam ran at rates of up to 25 Hz over a period of 6 weeks at beam momenta between 0.5 and 7 GeV/c. Time of flight and Cherenkov counters provided beam flavor tagging. Around 8M total ‘physics’ events were written, with around 3M having beam tag information. In total 850 TB of raw test-beam data were written, along with one PB of commissioning and cosmic data. These data were successfully cataloged and written to storage at both CERN and Fermilab at rates of up to 2 GB/sec.

Thanks to significant prior effort in the liquid argon (LAr) computing and algorithms community, reconstruction software was ready to go and the first reconstruction pass began soon after data taking started and was complete within two weeks of the end of data taking. Those results were extremely useful in demonstrating the capabilities of the detector and summarized in Volume II of the TDR[4]. A second pass, with improved treatment of instrumental effects ranging from stuck bits to 2-D deconvolution to correction for space charge effects was completed in late 2019. Figure 5 illustrates the signal processing stage of reconstruction, where raw ADC signals have noise and stuck bits removed and are then deconvoluted to yield gaussian hit candidates. Figures 6 and 7 illustrate full pattern recognition and event reconstruction.

While LArTPCs benefit from fine granularity and a uniform detector medium, diffusion, argon purity, fluid flow and the build up of space charge in the active medium can all introduce distortions into the detector response. These effects have all been simulated and tested in the ProtoDUNE-SP data.

Compressed raw input event records were of order 75 MB in size and took 500-600 seconds to reconstruct, of which around 180 sec was signal processing and the remainder high level reconstruction dominated by 40-60 cosmic rays per readout. Memory footprints ranged between 2.5 and 4 GB. Output event record sizes were reduced to 22 MB by dropping the raw waveforms after hit finding. Reconstruction campaigns took of order 4-6 weeks (similar to the original data taking) and utilized up to 15,000 cores on Open Science Grid (OSG) and Worldwide LHC Computing Grid (WLCG) resources. Job submission was done through the POMS[12] job management system developed at Fermilab. POMS supports submissions to FNAL dedicated resources and selected OSG and WLCG sites. Figure 8 shows the distribution of wall hours used for reconstruction in 2019.

For reconstruction, data were streamed via `xrootd`[13] from dcache storage at Fermilab to the remote sites. Despite individual processing jobs taking 15-30 hrs to complete, network interruptions rarely caused job failures.

2.2 ProtoDUNE-DP

The ProtoDUNE-DP detector began taking data using cosmic rays in August 2019. Thanks to preceding data challenges, those data have been successfully integrated into the full data cataloging and reconstruction chains and are now being reconstructed as they arrive. The ProtoDUNE-DP technology locates the readout systems above a thin layer of argon gas above the liquid argon surface. This gas layer allows an external electric field to accelerate the electrons and produce gas amplification. The result is a substantial increase in signal-to-noise in the resulting signals, at the cost of longer electron drifts from the bottom of the liquid volume. Figure 9 illustrates early data from ProtoDUNE-DP.

Figure 5. Comparison of raw (left) and deconvolved induction U-plane signals (right) before and after the signal processing procedure from a ProtoDUNE-SP event. The bipolar shape with red (blue) color representing positive (negative) signals is converted to the unipolar shape after the 2D deconvolution.

Figure 6. 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.

Figure 7. Pandora [14] reconstruction of cosmic rays and beam interaction in a ProtoDUNE-SP event. The left side of the figure shows the full detector volume with all interactions, including cosmic rays and the right side shows the identified beam interaction.

Figure 8. Reconstruction processing distribution amongst sites for DUNE production in calendar 2019. The inner circle shows national contributions while the outer circle shows individual site contributions.

Figure 9. Cosmic ray data from the dual phase prototype

2.3 Conclusions from prototype tests

ProtoDUNE prototype runs are ongoing and will continue through beam tests in 2021-22 at European Organization for Nuclear Research (CERN). Data cataloging, movement and storage techniques were tested before the start of the ProtoDUNE-SP and ProtoDUNE-DP runs and were able to handle the full rate of the experiments. Reconstruction algorithms were also in place on time and were able to produce early results that led to increased understanding of the detector and improved calibrations for a second iteration. These tests also identified some deficiencies in our infrastructure, including incomplete schemes for the transmission of configuration and conditions information between hardware operations and offline computing. The test-beam runs have been extremely valuable in allowing us to determine which variables are important to transmit and in designing improved systems for gathering and storing that information.

An additional run of both ProtoDUNE-SP and ProtoDUNE-DP is planned for 2021 and 2022, allowing further developing and testing of our computing infrastructure before the full detector comes online in the late 2020's.

3 On to full DUNE

The full DUNE FD will begin with one single phase module to be installed at Homestake starting in the middle of this decade. High intensity neutrino and anti-neutrino beams should arrive after a year or so of commissioning of the detector and Long-Baseline Neutrino Facility (LBNF) beamline. This first module will largely resemble a scaled up version of ProtoDUNE-SP with 150 APAs distributed 2-deep at the center and long edges of the cryostat. The argon volume will be $15 \times 14 \times 62 \text{ m}^3$ with a fiducial mass of 10kT. Table 1 summarizes the expected event rates and data volumes for one such detector module. Additional detector modules, likely one dual-phase, another single-phase and one with novel technology will be added. For now, we assume that data volumes and rates coming from other technologies will be less than or equal to the single-phase values.

The detectors should be sensitive to neutrino interactions and radioactive decays above an energy threshold of order 5 MeV. Unambiguous triggering may require a somewhat higher

threshold to avoid false triggers due to ^{39}Ar decays, but beam interactions in the 500-10,000 MeV range should have almost perfect detection efficiency. Sophisticated triggering algorithms should also allow standalone detection of astrophysical sources, including higher energy solar neutrinos and supernova candidates.

The data rates will be dominated by 4,500 cosmic rays expected per module/day. These events are vital for monitoring and aligning the detector. The next most significant source of events will be calibration campaigns with radioactive and neutron sources and lasers. In all cases, the goal is to gather data from the full volume of the detector with as fine a granularity as possible.

Beam interactions themselves are expected to be quite rare, occurring in only 1/2000 beam gates. Extraction of oscillation parameters will require both the powerful electron background rejection, discussed in the previous section, and precise calibration of the energy scale of the experiment, hence the much larger calibration samples.

Process	Rate/module	size/instance	size/module/year
Beam event	41/day	6 GB	47 TB/year
Cosmic rays	4,500/day	6 GB	9.7 PB/year
Supernova trigger	1/month	115 TB	1.4 PB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			12.9 PB/year

Table 1. Data sizes and rates for different processes in each far detector module. Uncompressed data sizes are given. As readouts will be self-triggering an extended 5.4 ms readout window is used instead of the 3ms for the triggered ProtoDUNE-SP runs. We assume beam uptime of 50% and 100% uptime for non-beam science. These numbers are derived from references [15] and [16].

Overall, bottoms-up estimates yield data volumes of around 13 PB/year/module. Lossless compression should reduce this volume and additional modules will likely increase these rates. A maximum rate of 30PB/year across all modules and modes of operation has been specified. We will note that 30 PB/year is an average of 1.3 GB/sec, less than the rates already demonstrated for protoDUNE acquisition and storage. In principle, at 2.5 CPU sec/MB of compressed input, 2000-3000 cores could keep up with these data rates but this throughput must be maintained over many years. In addition, supernova candidates may require bursts of much higher acquisition and processing rates. Table 2 summarizes the computational characteristics expected for FD data.

3.1 Supernova candidates

Supernova candidates pose a unique problem for data acquisition and reconstruction. Supernova physics in DUNE is discussed in some detail in the TDR[4] and only summarized here. A classic core-collapse supernova 10 kpc away would be expected to yield around 3,000 charged-current electron neutrino interactions across 4 detector modules. The oscillation physics is not fully understand and can result in significant modulations of the event rates for different neutrino types over the few tens of seconds of the burst. DUNE’s fine-grained tracking should allow significant pointing power with the most optimistic scenario of four modules and high electron neutrino fraction yielding pointing resolutions of less than 5 degrees. Figure 10 illustrates simulated signatures of supernova neutrino interactions in the far detector. The ability to produce a reasonably fast pointing signal would be extremely valuable to optical astronomers doing followup, especially if a supernova was in a region where dust

Figure 10. Left) a charged current interaction of a 30 MeV energy electron neutrino in the DUNE Far Detector. Right) a neutral current excitation and de-excitation of an Ar nucleus by a 10 MeV neutrino.

Quantity		Value	Explanation
Far Detector Beam:			
	Single APA readout	41.5 MB	Uncompressed 5.4 ms
	Single APA readout	16.6 MB	×2.5 compression
	APAs per module	150	
	Full module readout	6.22 GB	Uncompressed 5.4 ms
	Beam rep. rate	0.83 Hz	Untriggered
Signal processing	CPU time/APA	40 sec	from MC/ProtoDUNE
Signal processing	CPU time/input MB	2.5 sec/MB	compressed input
	Memory footprint/APA	0.5-1 GB	ProtoDUNE experience
Supernova:			
	Single channel readout	300 MB	Uncompressed 100 s
	Four module readout	460 TB	Uncompressed 100 s
	Trigger rate	1 per month	(assumption)

Table 2. Useful quantities for computing estimates for single-phase (SP) readout. For sparse FD events, the pattern recognition phase, which scales with occupancy is expected to be substantially faster than the signal processing phase which scales with detector size.

masks the primary optical signal. The need to be alert to supernovae and to quickly transfer and process the data imposes significant requirements on triggering, data transfer and reconstruction beyond those imposed by the more regular beam-based oscillation physics. For example, a compressed supernova readout of all four modules will be of order 184 TB in size and take a minimum of 4 hrs to transfer over a 100 Gbs network, and then take of order 130,000 CPU-hrs for signal processing at present speeds. If processing takes the same time as transfer, a peak of 30,000 cores would be needed.

4 Comments and Conclusions

This discussion has centered on the acquisition and fast processing of raw data from novel and extremely large liquid argon time projection chambers. Many other computing challenges lie ahead but were beyond the scope of this paper. These include

Simulation: particle propagation in liquid argon is reasonably fast to simulate as there are not complicated volume boundaries to cross but simulating electron drift trajectories (and scintillation light trajectories) in a diffusive, electron absorbing, moving medium immersed in a non-uniform electric field remains a challenging computational challenge.

Near detectors: a suite of near detectors are needed to characterize the neutrino beam as it originates at Fermilab. These detectors are still being developed but will introduce a large number of differing detector technologies. While individual interactions are likely to be much smaller than readouts of the far detectors, the beam cycle is of order 1 Hz and each readout will contain multiple cosmic ray and beam interactions.

Data analysis: The small (order 100) group of ProtoDUNE-SP and ProtoDUNE-DP and DUNE developers and analyzers have successfully analyzed the beam and cosmic ray data and performed simulations needed to produce the physics sections of the TDR. We expect analysis of the full experiment to involve many more individuals and much more data. A campaign of training for new users and design of a suite of efficient analysis tools is needed. We have initial prototypes based on NOvA and MicroBooNE analysis.

Fortunately, DUNE is able to take advantage of the huge and heroic developments in software and computing made for the Intensity Frontier and LHC experiments over the past decade. We have demonstrated that, even with preliminary versions of our tools and algorithms, we can quickly reconstruct and analyze data from large liquid argon TPC's at full rate. We look forward to an exciting and fruitful next decade.

Acknowledgements

Schellman's research is supported by the US National Science Foundation under grant 1806849 and through a joint appointment with the Fermilab Scientific Computing Division. Fermilab is a U.S. Department of Energy, Office of Science, HEP User Facility managed by Fermi Research Alliance, LLC (FRA), acting under Contract No. DE-AC02-07CH11359.

References

- [1] T. Kajita, Reports on Progress in Physics **69**, 1607 (2006)
- [2] R. Acciarri et al. (ArgoNeuT), Phys. Rev. **D95**, 072005 (2017), 1610.04102
- [3] B. Abi et al. (DUNE) (2020), 2002.02967
- [4] B. Abi et al. (DUNE) (2020), 2002.03005
- [5] B. Abi et al. (DUNE) (2020), 2002.03008
- [6] B. Abi et al. (DUNE) (2020), 2002.03010
- [7] R. Acciarri et al. (ArgoNeuT), Phys. Rev. **D99**, 012002 (2019), 1810.06502
- [8] R. Acciarri et al., Journal of Instrumentation **12**, P02017 (2017)
- [9] F. Varanini (ICARUS), EPJ Web Conf. **164**, 07017 (2017)
- [10] R. Acciarri et al. (MicroBooNE), JINST **12**, P08003 (2017), 1705.07341
- [11] R. Acciarri et al. (MicroBooNE), JINST **12**, P02017 (2017), 1612.05824

-
- [12] *Production operations management system*, {https://cdcvs.fnal.gov/redmine/projects/prod_mgmt_db}
 - [13] G. Behrmann, D. Ozerov, T. Zangerl, J. Phys. Conf. Ser. **331**, 052021 (2011)
 - [14] R. Acciarri et al. (MicroBooNE), Eur. Phys. J. **C78**, 82 (2018), 1708.03135
 - [15] Schellman, H., Tech. rep. (2019), <https://docs.dunescience.org/cgi-bin/private/ShowDocument?docid=16028&version=1>
 - [16] Klein, Joshua R., Tech. rep. (2019), <https://docs.dunescience.org/cgi-bin/private/ShowDocument?docid=16028&version=1>