



2025 DRAFT LArSoft Work Plan

Last updated: November 13, 2024

Introduction

This document gives the DRAFT LArSoft work plan for 2024. Progress on specific items going forward are given quarterly at Steering Group meetings.

Erica and Katherine discussed priorities with each experiment in a series of meetings in September and October of 2024. The experiments detailed their plans for the next year, the requirements for LArSoft, and how the LArSoft Project Team (“the Project”) could help, as well as what the experiments might be able to contribute to LArSoft code. Some items may be raised by the Project to address general problems faced by all LArSoft users. Major observations from these discussions are listed in Appendix B since they may be of interest to other experiments and the project as a whole.

Based on those discussions, LArSoft proposes a plan of work for 2025 along with relative priorities of the various items. The Project resources responsible for executing parts of the plan are listed in cases where this is known. It should be noted, however, that the Project does not have sufficient staff at present to carry out the full program of work in this plan, or the entirety of the short term or long term programs individually. The Project will collaborate with experiments to prioritize and coordinate the work to best meet experiment needs. We also note that some work related to the plan can be performed by the experiments. Allowing experiment members to receive service credit for contributions to work plan items that benefit their experiment, but that are of a more general nature would help to fill some of this gap in effort. The more this happens, the more all experiments benefit.

1. Short term priorities for 2025

We define short term priorities as those items that should receive sufficient effort to be substantially completed in CY2025 due to the urgency or timeliness of the need, and where the Project will drive the timeline. The meaning of priorities listed is discussed in Appendix A.

1. **Enable execution of critical production workflow elements on GPUs or other accelerators where significant performance gains can be obtained, or where needed for execution at HPC centers.**
 - **Identify target algorithms in consultation with experiments**
 - **Enable execution on GPUs via portability layers, GPU-as-a-service, or other techniques as appropriate**
 - **Coordinate / collaborate with efforts to use LArSoft on HPC platforms.**

Priority: high

This is a multi-year effort that started in 2021 with the introduction of GPU-as-a-service for the execution of AI/ML inferencing tasks in LArSoft. In 2024, work turned to optimizing a specific algorithm and performing studies to measure performance gains of running on a GPU, as well as expanding the set of ML algorithms that could use GPUaaS inferencing to include a general ML reconstruction workflow. Work in 2025 will continue in these directions, focusing on algorithms that operate on high-density data or that otherwise require many parallel calculations, and allowing GPU-enabled Geant4 code, once this is available, as an option for particle tracking and optical photon ray tracing.

Resources: Kyle Knoepfel, Mike Wang, Marc Paterno, Saba Sehrish, Robert Hatcher.

2. **Thread safety and multi-threading to ensure optimized resource utilization on grid or HPC resources for critical production workflows**
 - **Modify or re-architect LArSoft and experiment code as needed to ensure thread safety.**
 - **Implement multi-threading where appropriate to improve resource utilization.**

Priority: high

This is a multi-year effort that started in 2018 working on thread safety and multi-threading to improve resource utilization on existing grid resources, and to prepare LArSoft code to be effective in an HPC environment. A fully thread-safe production workflow LArSoft were completed in 2021. In 2022 and 2023, the focus was on critical production workflows where the introduction of these multi-threading promised improved resource utilization or increased throughput. Work completed during that phase included making production database services thread safe, and making a DUNE super-nova neutrino workflow run in multi-threaded mode.

In 2025, we will continue to identify and address production code where thread safety or multi-threading is needed for efficient execution.

Resources: Kyle Knoepfel, Mike Wang, Marc Paterno, Saba Sehrish.

3. Provide support for a multi-experiment event display capable of dedicated and integrated displays of TPC charge collection and photo-detector systems, along with external detectors such as CRTs, spectrometers, calorimeters, etc.

Priority: high

Decide whether LArSoft will provide support for Titus (or possibly some other) experiment-agnostic framework for event displays. Titus support would entail re-architecting to make proper, extensible use of the LArSoft Geometry system, providing a library of primitives to assist in displaying LArSoft geometry and data product information and in controlling the display, and developing a system of plug-ins for experiment-specific elements. Similar work would apply to a non-Titus solution. Agreements on long-term maintenance and code ownership would be part of this work as well.

4. SPACK - migrate to a new, community-supported build system for art and LArSoft

Priority: high

This is a continuation of a multi-phased project to migrate LArSoft to a Spack-based packaging and build system, which has broad HEP community and industry support. Via this migration, we will address a number of long-standing portability and configurability concerns raised by the experiments that are inherent to the current UPS-based system. The overall strategy centers on minimizing disruption to the experiments by providing backward compatible changes until the final migration. The first major milestone, preparing LArSoft for the transition to a system with a Spack-based back end, was completed in Q4 of 2021. The final phase will end support for UPS. An aggressive education campaign is part of this project.

Major milestones for 2025 will include (a) achieving Spack readiness for the transition, (b) approval of migration plan for LArSoft and the experiments, (c) completion of the transition to Spack.

Details on Spack at Fermilab can be found in a [presentation by Marc Mengel](#) at the LArSoft Coordination meeting. Further information on the build system can be found in the following presentations by Chris Green: [Feb 23, 2020](#), [Nov 2, 2020](#). Update on [September 2023](#). and an overall plan – [Spack strategy.docx](#) with instructions at <https://fifewiki.fnal.gov/wiki/Spack>

Resources: Patrick Gartung, Chris Green, Kyle Knoepfel, Marc Paterno

5. Neutrino event generator refactoring

Priority: high

Refactor the interface to the GENIE neutrino event generator such that the GENIE version can be selected through the runtime environment configuration. One side effect will be to provide a

template for how to interface other event generators (e.g., Achilles, GIBUU, NuWro) so as to simplify integration and maximize flexibility. The work will commence after the final phase of the Spack migration is completed.

Resources: Robert Hatcher, Steven Gardiner, Liang Liu

This work will follow completion of the Spack migration.

6. **LArSoft infrastructure updates**

Follow-up on issues raised by the experiments that have implications for LArSoft infrastructure and architecture. These include:

- Support for non-planar cathode geometries to facilitate tracking across non-planar cathodes.
- Support for TPC-dependent drift velocities and electron lifetimes.
- Re-factoring hit finding code to allow experiment dependent handling of long pulses
- Review use case of jumpered and staggered readout wires within a single logical TPC.

7. **Review and update / augment LArSoft documentation in key areas as requested by experiments.**

Priority: high

Continuing project to provide improvements covering areas noted by experiments, and new developments in LArSoft capabilities and the surrounding ecosystem. Initiatives will include:

- Maintain and enhance an up-to-date introductory slide deck on LArSoft - <https://indico.fnal.gov/event/49621/#1-see-attached-source-material>, ensuring alignment with other sections of LArSoft.org
- Periodically verify links in GitHub wiki material and correct ones that are invalid.
- Continuously review, enhance and update LArSoft documentation as needed.

Resources: SciSoft team

8. **Provide consultative support to experiments for transitioning to the refactored LArG4 detector simulation framework, or developing simulations for new detectors**

Priority: high

Provide consultative support for efforts within the experiments to migrate to the new simulation framework. Advise and provide guidance in refactoring experiment code to be compliant with the workflows and interfaces within that framework.

Resources: Robert Hatcher for LArG4 / artg4tk, SciSoft team for geometry, other infrastructure

Currently using refactored larg4: ICARUS, SBND (?), DUNE, leaving only MicroBooNE. So this one might be done.

2. Long term and continuing priorities for 2025

We define long term priorities as those that we do not anticipate completing within CY2025 due to the nature of the work involved. Unless explicitly indicated, a project defined as a long-term priority should not be construed as implying it is “low priority”. A “continuing priority” is a set of tasks that by construction represent a continual, recurring stream of work that may or may not span more than a calendar year.

1. Provide technical expertise, advice and leadership in matters related to LArSoft and software development

2. Develop and execute support plans as needed in areas where assistance from the Project has previously been requested by experiments, or advised by the Project.

Priority: per request

- a. Align Wire-cell reconstruction and simulation with LArSoft workflows and interfaces (eg, factoring electron drift from anode response simulations via drifted electrons)
- b. Extend the NuWro direct integration model to include the GiBUU event generator framework and the NEUT event generator. Depends on the GENIE refactoring listed in the short term priorities.
- c. Investigate need for magnetic field map within LAr volume. Design and develop the interface for such a service. -- Make this a named project. Assign architect, etc.
- d. Coordinate between GARSoft and LArSoft as needed to allow GARTPC design and development
 - i. Ensure sharing of data products and geometry (and with other detectors)
 - ii. Ensure sharing of any services, should that be needed (eg, for magnetic field)
 - iii. Possible introduction of non-uniform magnetic field to LAr volume
- e. Participate in discussion and groundwork associated with DUNE migration to a new event processing framework
- f. Education and assistance with optimization of critical production workflows, particularly for the operating SBN experiments.
- g. Provide consulting assistance on the development, implementation and deployment for pixel readout geometry descriptions for the various DUNE pixel detectors.
- h. Provide a solution suitable for running TextFileGen at production-scale

3. Architecture work to facilitate integration of machine learning into LArSoft algorithms

Priority: per request

Possible targets include standard modules to provide translation and downsampling for image-based algorithms. A work plan item can be opened once a set of deliverables is defined.

4. **Support for integration of community-supported event displays within art and gallery based LArSoft jobs.**
Priority: per request
5. **Work on user support, bug fixes, emergency feature requests, software builds, releases.**
Priority: per request
Resources: SciSoft team
6. **Other topics under consideration to work on include the list of accepted, but not assigned, redmine issues.** These can be found under [accepted redmine issues](#).
Priority: per request

Other considerations on work planning

In response to input received from the experiments, the Project will pursue the following topics as time allows:

- Improving informal channels of communication, which have suffered during the Covid era. To facilitate more efficient communication, community members are encouraged to engage directly with SciSoft team members before and after opening issue tickets. In addition, SciSoft team members understand that greater visibility at experiment software meetings would be helpful.
- Seeking information on recent LArG4 migration experiences, then using this information to enhance the available documentation.
- Seeking information on recent experiences using LArSoft on HPC resources, then using this information to create a common resource for others interested in using HPC.
- Developing a plan to ensure LArSoft is compatible with reading and writing data in HDF5 format, as well as other data formats.
- Tracking progress on DUNE and related community-based data processing framework developments.

Appendix A: Definition of priorities

The priorities listed above have the following meaning. The impacts noted are those negotiated with the experiments.

- High: the most urgent projects that have high impact to the community, and that should be fully staffed at all times until completion
- Medium: a project that has a significant impact to the community, and that should be staffed sufficiently to ensure completion within the calendar year

- Low: a project that does not have immediate impact to the community, could be rolled into the next planning period without imposing undue burdens, and should be staffed after ensuring all other obligations are met.
- Per request: The priority will be set in consultation with the requesting experiment or ticket author.

Appendix B: Major observations from one-on-one meetings with each experiment in September and October of 2024

(Topics in common across multiple experiments)

1. Common

- Event display that is useful in current environment.
- Mutli-threading
- Spack migration
- GPUs for AI/ML

2. DUNE - Met with Tom Junk, Tingjun Yang

- Keep working on the things we're working on.
 - Pixel geometry
 - GPU-ifying code
 - Thread safety
 - Event generator refactoring (eg, Achilles)
- Keep eyes open for what we need on LArSoft side for GPUs.
- DUNE to give a list about Event Display issues if DUNE wants. (Don't have anyone to work on it anyway.)

3. ICARUS - Met with Tracy Usher and Daniele Gibin

- Infrastructure changes.
 - Should have an art tool that finds hits within ROIs. Existing code in the GausHitFinder module breaks long hits into many gaussians. Would prefer to define a different algorithm for ICARUS. So need that code element in a tool.
- Event display
- Multi-threading (not clear new work is needed. Are running first stage of production in MT mode, and possibly a second in the future)
- *** Before we get serious about migrating to new Geant 4, we should have someone talk about the details of the changes that would be relevant for liquid argon detectors.

4. MicroBooNE - Met with Herbert Greenlee.

Summary:

- Assisting with migration to SPACK.
- Assisting with migration to MCC10.

- C. Looking at MicoBooNE simulation to help with decision about sticking with legacy or refactored LArG4

5. **SBN Data/Infrastructure** - Met with Steven Gardiner, Giuseppe Cerati

- A. Running ML / AI, possibly on GPUs. Unclear if other GPU targets are needed outside of simulation
- B. Infrastructure changes to accommodate non-planarity of cathodes. Level of detail needed is unclear, as is the need for detailed Efield description in cathode region.
 - a. Cryostat dependent electron lifetime is also on their list
- C. A clear plan for Spack migration, and a clear picture of what Spack world will look like

6. **SBND** - Met with Andrzej Szelc, Anne Schukraft

- A. Will mount effort to run LArSoft on Polaris at ALCF using same strategy as was used previously with Theta.
 - a. Might have interest in GPU-ified Geant4 or detector simulation
- B. Interested in integrating NuWro and GIBUU into LArSoft using new GENIE model, once that is completed. (after Spack)
- C. Want detailed plan for transition to Spack
- D. Follow-up discussions needed regarding treatment of jumpered wires in geometry. Have a working, albeit cumbersome solution now, so may not need changes.