



2021 LArSoft Work Plan

Last updated: June 17, 2021

Introduction

This document gives the LArSoft work plan for 2021. 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 October of 2020. The experiments detailed their plans for the next year, the implied requirements for LArSoft, and how LArSoft could help, as well as what the experiments might be able to contribute to LArSoft code. Some items may be raised by LArSoft, not a particular experiment, to address general problems faced by all LArSoft users.

Based on those discussions, LArSoft proposes a plan of work for 2021 along with the relative priority of the various items. The resources responsible for executing parts of the plan is listed in cases where this is known. It should be noted, however, that the LArSoft project team does not have sufficient staff at present to carry out the entire program. Some projects can and should be performed by the experiments. Allowing experiment service credit for contributions to items that benefit an 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.

Latest quarterly changes to the plan and status updates are noted in red below.

Short term priorities for 2021

We define short term priorities as those items that should receive sufficient effort to be substantially completed in CY2021 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. Thread safety, multi-threading, and HPC:

- **Re-architect LArSoft and experiment code to ensure thread safety within critical production workflows** (on-going since 2018)
- **Implement multi-threading where appropriate to improve resource utilization in critical production workflows** (on-going since 2018)
- **Coordinate / collaborate with efforts to use LArSoft on HPC platforms**

Priority: high

Target high-priority production workflows to allow running in multi-threading environments. Focus multi-threading and related efforts toward improving resource utilization on existing production platforms, and enabling execution within the context of on-going work external to the LArSoft Project to run LArSoft on HPC resources. Collaborate with external efforts to run multi-threaded workflows on HPC resources where possible. Known efforts or requests include an SBND project to run LArSoft on Theta at Argonne, and a LArIAT project to run LArG4 on GPUs.

<https://cdcv.s.fnal.gov/redmine/issues/19288>

Resources: Kyle Knoepfel, Saba Sehrish, Mike Wang, in collaboration with HEPReco SciDAC

March Update:

- Updates to the art caching code needed for database access within LArSoft have been completed.
- Toy detector model developed, using as prototyping platform for multi-threading changes to the refactored LArG4 simulation
- Work completed to update TBB usage to comply with current conventions and constructs

June Update:

- Worked to fix some problems found in original release of art caching. Updated art caching code released in art 3.09. LArSoft migration to art 3.09 imminent.
- Significant progress toward preparing services that access databases to use the art caching features. After changes to LArSoft are completed when art 3.09 is integrated, will move to fix experiment DB services.
- Currently working to import a unified API from CMS that will make it easier for LArSoft users to access GPUaaS within jobs.

2. Pixel detectors within LArSoft

- **Re-architect geometry systems to improve integration of multiple readout schemes**
- **Identify and implement common low-level and 3D data products needed**
- **Support development of pixel-based anode simulations**

Priority: high

Work with experiments to develop requirements and design solutions. Expect experiments to carry out a significant fraction of the work to adapt reconstruction and simulation codes.

<https://cdcv.sfnal.gov/redmine/issues/19328>

Resources: Erica Snider, Kyle Knoepfel

March Update:

- Pixel meeting will be scheduled with Tom Junk, Tingjun, Andy Mastbaum to discuss the next steps.

June Update:

- **More discussions required following meetings. To be scheduled.**

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

Priority: high (when the system goes into production)

To migrate LArSoft to the Spack packaging tool, and the SpackDev development system, a standard set of packaging and build tools that have broad HEP community support. Via this migration, address portability and configurability issues raised by experiments with the current build system.

<https://cdcv.sfnal.gov/redmine/projects/spack-planning/issues>

This item requires delivery of the build system, which is pending and expected in 2021. Work will commence when a production release of the system becomes available.

Resources: Patrick Gartung

March Update:

- Work has begun to migrate LArSoft to use cetmodules, a new, Spack/CMake-based build system. The new system is backwards compatible with the current UPS/cetbuildtools-based system, which will serve as a front-end to cet_modules during the transition period. A full migration plan is still under development. Details on Spack at Fermilab can be found in a [presentation by Marc Mengel](#) at the Feb. 23rd LArSoft Coordination. Information on cet_modules can be found in a [presentation by Chris Green](#) at the same meeting.

June Update:

- **Work on preparing LArSoft for completion of the first phase of the transition to Spack, which provides a Spack-based build back-end that supports the use of UPS and the**

current build interface on the front end. The first phase will be completed with the migration to art 3.09.

- Produced near complete set of draft user documentation

4. **Provide support to experiments in transition to the refactored LArG4 simulation framework**

Priority: high

Provide active, 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: Hans Wenzel for LArG4 / artg4tk, SciSoft team for geometry, other infrastructure

March update:

- No activity to date

June Update:

- No activity this quarter

5. **Neutrino event generator refactoring**

Priority: high

Refactor Genie neutrino event generator initialization in order to decouple LArSoft from Genie. This work will provide an exemplar for how to interface other event generators to LArSoft in a manner that simplifies integration while maximizing flexibility.

Resources: Robert Hatcher

March update:

- No activity to date

June Update:

- No activity this quarter

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

Priority: high

Address areas noted by experiments as in need of significant improvement:

- Foundational principles and best practices for new developers. Reinforcing framework / algorithm separation should be the central message of this documentation.
- Running and developing LArSoft on non-SLF or other non-standard platforms using containers
- Migrating and developing new simulations within the refactored LArG4 framework

- Guidance on writing more efficient and maintainable code
- Adding a wiki page on Event Displays

March update:

- Added wiki page on Event Displays.
- Work progressing to update the base documents on larsoft.org and in the LArSoft wiki used as introductory material, and as reference for design and coding. This work includes
 - Major update to [LArSoft repository and dependencies page](#) to reflect changed dependencies.
 - Consolidating and updating previous workshop and tutorial presentations into a single slide deck to simplify maintenance.

June Update:

- Added information on multi-threading to larsoft.org and to the single slide deck to ensure the LArSoft design principles match reality.
- Added information about the design principle to separate algorithms and services and framework to the slide deck.
- Changed code releases and distribution information to the new way of doing this with github in the slide deck.
- Completed first draft of updated slide deck, and began circulating for comment

7. Support integration of community-supported event displays for use within art and gallery based LArSoft jobs.

Priority: medium

Consult with on-going efforts by the experiments to develop and deploy event display solutions suited to their needs, to assist with integrating those tools into LArSoft and gallery environments. This work can include consulting in re-architecting components of those event displays to enhance usability, extensibility and maintainability. Project details to be defined in consultation with the experiments.

Resources: SciSoft team

March update:

- No activity, except adding a description of [Event Displays](#) options with links to further information.

June Update:

- No activity this quarter

Long term and continuing priorities for 2021

We define long term priorities as those that we do not anticipate completing within CY2021 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. Places where 2021 support from LArSoft team has been requested or advised

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. Geometry work needed to address different wire readout schemes (e.g., cases where ICARUS breaks implicit assumptions: non-vertical collection wires, existence of horizontal wires, multiple channels per geometric wire)
- c. Extend the NuWro direct integration model to include the GiBUU event generator framework and the NEUT event generator.
- d. 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.
- e. Coordinate between GArSoft and LArSoft
 - 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

Learned Aug 2020 that GAr detector is not yet part of the ND baseline design, and is currently competing against other design options.

- f. SBND data reduction strategies
- g. Participate in the discussion of how the DUNE data streams (SN + possibly others/all) will be handled within *art* / LArSoft. May involve moving away from root-based I/O.

2. Architecture work to support a common framework for data preparation

Priority: per request

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

June update:

- *See GPUaaS update under “Short term priorities”, item (1).*

4. **Ongoing work on bug fixes, emergency feature requests, building software, helping users.**

Priority: per request

Resources: SciSoft team

March update:

- Tickets since Jan 1, 2021
 - 10 opened
 - 10 closed

June update:

- *Tickets since Apr 1, 2021*
 - *8 opened*
 - *5 closed, 2 in progress, 1 new*
- *Request from SBN to provide specific end-user support for profiling*
 - *Developed a plan to address concerns*
 - *As part of the plan, are preparing a tutorial on using profiling tools to investigate and address both CPU and memory issues.*

5. **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

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.