LArSoft Steering Group Meeting Notes, September 13, 2023

Attendees: Mathew Muether, Herbert Greenlee, Thomas Junk, Dominic Brailsford, Angela Fava, Will Foreman, Giseppe Cerati, Erica Snider, Katherine Lato

LArSoft project report:
- Purpose of the meeting is to give experiments spokes-level oversight of the direction of project work to ensure alignment with experiment needs and objectives. Also includes an update on the status of work on the current work plan, which itself is a high level description of what the LArSoft Project support team (SciSoft) are working on each year. We are currently in the process of developing the 2024 work plan by meeting with each experiment in September and October.

- Discussed the current effort level and budget outlook.
  - The addition of two people (each at a fraction of an FTE) to the project this year has had a positive impact on the work on multi-threading and thread safety. One has GPU expertise, which has so far gone underutilized.
  - Looking forward, there is considerable uncertainty about what the FY24 budget will look like. We currently expect the computing operations budget to go down, but do not yet have reason to believe that this will result in a loss of effort available to the project team.

- The work plan with latest status updates is available at: https://indico.fnal.gov/event/61123/
  - The plan lists the things we want to do with the understanding that, due to the limited effort available, we will not be able to come to completion on many items.

  - Question about if everything is listed as "high priority", does that mean no differentiation between them? A: No. The items listed as high priority reflect those that are considered to be the most important to the experiments, so the most important for the team to work on. We generally acknowledge going in, however, that we cannot complete all the work in the high-priority items given the available effort, so further prioritization is applied within those items. We list them all as "high priority" in the plan in part as a statement to management about the work that is needed.

Round Robin:
- ICARUS - Angela Fava:
  - Event display is relevant to us. That remains one of our priorities.

- DUNE - Thomas Junk:
  - We're most concerned about SPACK and AL9 and what must be done to make progress. It's urgent but not directly tied to physics. Interested in decoupling neutrino generators. It affects our ability to do our physics.
    - Erica: Robert Hatcher is the resource there. He had a plan that looked like it would work. It lay dormant for a while, but has received some attention recently. It was going to start with GENIE then all the generators could use the same pattern. Once you get one to work, it should be pretty easy to get them all to work the same way. We hope this gets wrapped up before the end of the year.

- ○ There was a decision to go with SPACK, but the transition has to be owned by people who are eager to develop. Who is leading this work?
    - ■ Person leading the project is Steve White. Mark Mengel and Chris Green are working on it. We (LArSoft) are one of many clients for that. The migration work falls on LArSoft, but there are changes to SPACK itself that are needed to support building of releases for products like art. This has been the holdup.
    - ■ External SPACK team has been reasonably responsive to our needs. We are an official collaborator. Some of the bigger changes took time.

- ● DUNE - Dominic Brailsford:
  - ○ Planning on moving to full-size DUNE FD. Processing and I/O infrastructure changes needed to make this work. Multi-threading is mandatory for the processing side. Data is so big, I/O will be an issue, eg, how to write to disc. We're going to need a proper strategy with external help. Perhaps this is more on SciSoft side.
    - ■ Erica: Some will fall to framework, so not LArSoft. We're keeping an eye on the framework changes needed for this, and if LArSoft code needs to be adapted.
    - ■ Dominic: There will be different technologies for top and bottom TPCs in DUNE FD. For instance, sampling time is different: 200 MHz vs 250 MHz. This is not currently supported.
    - ■ Tom: May affect calorimetry. Need vectors in more places than just time arguments.
    - ■ Erica: We should talk about the details during the one-on-one discussions on the 2024 work plan to see whether there is something that needs to go in the work plan.
- ● DUNE - Will Foreman
  - ○ Working to integrate the blip reco for low energy physics.
  - ○ Erica: Should be able to complete this via the routine code and release management processes. Will monitor in case this requires dedicated project work / effort.

- ● MicroBooNE - Herbert Greenlee:
  - ○ Interested in SPACK migration like everyone. We're still using legacy LArG4. Are we going to be forced to switch to refractored LArG4?
    - ■ Erica: As long as the legacy works for you, that's probably okay as long as no big development is needed. Will not be forced to migrate.
    - ■ Herb: I don't think we'll be doing any major development at the Geant level. There are people looking at different ways of doing light simulation, which I don't think requires refactoring.
  - ○ Herb: We're going to be forced into containerization to keep MCC9 alive.
    - ■ Erica: That is likely true. AL9 will be a problem otherwise, as it implies lots of changes to MCC9.
    - ■ Herb: I don't know if LArSoft can contribute to solving our concerns (re. containerization).
    - ■ Erica: probably mostly SciSoft for that

- SBN - Giuseppi Cerati:
  - Since Geant4 is the step that takes most of the memory, it would be great if there was support for multi-threading. Within the event. Moving to new refactored version of LArG4. Not completed yet. Wasn't as easy as we thought, still in the process of migrating.
    - Erica: Once you get there, you should get multi-threading.
    - Giuseppi: WIthin one event or across events?
    - Erica: Believe so, yes. There are some obvious places where you should be able to do multi-threading.
    - Giuseppi: Complete migration, then see how much resources it takes.
  - Separate machine learning workflow, want output to be available in a CAF format. One can try to patch things after the fact, but an integrated version would be desirable.
    - Erica: That would be a great thing to integrate. I'd love to talk to someone about making that happen.
    - Giuseppi: Should be a generic solution, not tied to a particular workflow.
    - Erica: It would be useful to have a conversation to figure out what a common approach would look like. Giuseppi agrees.