



2018 LArSoft Work Plan

by: Erica Snider and Katherine Lato
Last updated 12/4/18

Introduction

This document records the LArSoft work plan for 2018. Progress on specific items going forward are recorded in redmine issues with updates to this document given quarterly at Steering Group meetings.

Erica and Katherine discussed 2018 priorities with the Offline leads starting in October of 2017 and through November. 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 were raised by LArSoft, not a particular experiment.

Based on those discussions, LArSoft proposed short-term and long-term priorities outlined below at a Steering Meeting in December 2017. Note, when defining work, we include items that can be done by members of the collaboration. There is a cost associated with making things workable for other experiments, but the benefit is that other experiments develop software that is usable by all experiments. The more this happens, the more all experiments benefit.

Short-term priorities; happening concurrently

These are happening concurrently. The order does not imply priority.

- 1) **Ongoing - Investigate having a new Event Display framework common to all experiments.** Some important features:
 - a) Provides a simple, intuitive means to navigate through large events and multi-TPC detectors.
 - b) Ability to zoom and pan with a data density that matches the screen resolution.
 - c) Provides basic interfaces to art for controlling runs and events, the geometry service, run conditions services. Re-running reconstruction workflows and displaying the result should be possible.

This work does not cover development of drawing routines. It is expected that the core event display users within the experiments will provide this code. Arrangements for this follow-on work will be part of this project.

Tracking this item is via <https://cdcv.sfnal.gov/redmine/issues/19034>

Resources: Erica Snider, Katherine Lato + pre-arranged people from the experiments

3/7/18 status:

[Requirements](#) have been gathered and ranked. These are being reviewed by the community at LArSoft Coordination meetings. The plan for this investigation is documented and tracked in LArSoft issue [19034](#), and is summarized below:

- a) Develop a set of requirements, ranked according to the following categories: essential, strongly-recommended, desired, and possibly-useful.
- b) Get community input on the ranking and final list of requirements at two LArSoft Coordination meetings to ensure that people can give opinions.
- c) Establish criteria to judge between the technical options, and review these with the community.
- d) Identify a list of possible technologies such as Root, Venu, ParaView, QT and other potential technologies and how applicable they are.
- e) Compare technologies. Ideally, write a toy display that demonstrates the requirements, or sufficient subset of critical requirements. Code it up on each platform
- f) Evaluate based on established criteria in task [19037](#)
- g) Take the results to SCD management for review, along with an initial deployment proposal and a request for funding.

5/8/18 status: As reported at the LArSoft Coordination Meeting, the initial set of requirements for a new common event display framework have been approved. These requirements were reviewed at LArSoft Coordination meetings on [April 10th](#) and on [May 8th](#). The next step is to define evaluation criteria and identify potential technical solutions. The plan will be to attempt to perform the evaluation over the summer.

6/7/18 status: No change.

8/3/18 status: No change.

8/31/18 status: Began preliminary investigation of technology options. Noted that during initial search for options, the existing choices (root, Venu, ParaView, Qt) did not appear, suggesting that they are not widely used.

10/11/18 status: Paul Russo is investigating technology options, paying attention to [Eve](#). This is task <https://cdcv.s.fnal.gov/redmine/issues/19038>

11/29/18 status: This item is on the work plan for 2019.

- 2) **Ongoing - Pixel-detectors within LArSoft.** The simulation moves to LArG4 phase 2. The reconstruction piece has two possible scenarios.
 - a) Pattern recognition algorithms shared between pixel and SP/DP LArTPCs. This requires agreeing on a common data format to represent 3D data, whether obtained directly from a pixel detector, or inferred from 2D views in SP/DP detectors. There will likely be implications for the geometry interface, which will then need to support both conventional and pixel-based detectors. Some study will be needed to determine whether these constraints can be met. If code is shared, then support for the development of those algorithms will also be needed.
 - b) No sharing of pattern recognition algorithms between pixel and conventional SP/DP LArTPCs. The code bases in this case diverge, so no further work is required aside from the normal support for data product, service and algorithm development.
<https://cdcv.s.fnal.gov/redmine/issues/19328>

Resources: ??

2/19/18 status: Need agreement on the data representation with the pixel community which we were hoping to get at the 3D reconstruction workshop. Since that hasn't been scheduled, to move ahead, we propose using the [Space Point Associated Charge in the LArSoft repository](#). There may be additional requirements, and LArSoft is willing to consider them as they arise. We'd like to agree on the information that needs to be carried around. It may not matter if we agree on the truth part.

6/7/18 status: it has proven difficult to schedule a meeting with Alan Bross (and his compatriots), Erica Snider, James Sinclair, Tracy Usher, Brett Viren, Jonathan Asaadi, Jen Raaf, Chris Backhouse, Daniel Dwyer, Kazu Terao, Xin Qian. With DUNE's interest in the software from Alice, it may open an opportunity to move forward.

7/16/18 status: LArSoft would like agreement on a common representation for pixel data versus reconstructed 3D space-points from wires. The project is organizing a meeting in early August with interested parties to discuss whether and how to move forward with collaborating on a common foundation for direct sharing of 3D algorithms between pixel and detectors with wire readouts. Starting with a one-hour video call to determine next steps.

8/3/18 status: At an August 2nd video-conference with Erica Snider, Tracy Usher, Xin Qian, Kazu Terao, Chris Backhouse, Jen Raaf, Hunter Sullivan, Tom Junk, Giuseppe Cerati, people shared what they were doing and where they are in broad terms. We began to discuss how we might proceed to develop a common basis for all of this work. Working on developing an in-person workshop to share and discuss the various 3D reconstruction efforts in liquid and gaseous argon, and ways to organize those efforts around some common set of data structures/organization or other interfaces to share software. Decided to work on a joint document for people to add information to. Shared via a google document which Erica, Katherine and Xin contributed to.

9/4/18 status: Asked people individually to review and revise the [google doc](#).

11/7/18 status: Erica spoke to Hiro Tanaka from SLAC about integrating pixel reconstruction and simulation within LArSoft. He is working on DUNE ND pixels, and is committed to getting the detector reconstruction and simulation into a common framework with DUNE FD. After exchanging ideas, LArSoft provided a list of LArSoft and software experts who could attend a workshop devoted to working on this integration. I also noted other pixel people I have spoken to in relation to this work item. Hiro's focus is more narrow than our previous work on this topic in that he is only interested in pixels. He agreed that it would be good to avoid precluding integrating other 3D sim/reco stakeholders in making choices for pixels. The timeline for this meeting could be before the end of the year.

12/4/18 status: The meeting did not happen in 2018. This item is moving into the 2019 LArSoft work plan.

- 3) **Ongoing - Use of SIMD vectorization to optimize LArSoft algorithms.** Introduce vectorized types to redesign some existing LArSoft data structures and modify existing algorithms, with the goal of exploiting vectorization opportunities to significantly improve performance of existing LArSoft algorithms. This project involves profiling real use case scenarios, redesigning data structures and applying changes to the most performance-critical algorithms, documenting performance comparisons and other studies, preparing reports and writing documentation as needed. Addresses in part the 2017 long-term priority on concurrency, <https://cdcvs.fnal.gov/redmine/issues/17920>

Resources: Guilherme Lima

2/19/18 status: Big loops over waveforms are candidates for this work: DetSim (SimWire), WireCell, CalWire

Guilherme consulted with Brett about WireCell. Brett noted that updating the fftw library to a new release recently produced a factor of 4 improvement in CPU performance. Guilherme surmised that this meant that they are using a vectorized fftw, so is not a candidate.

DetSim and WireCell, must choose one or the other until the LArG4 refactoring is done.

SimWire shows up in Soon's profiling results as a 60% consumer of CPU, verifying this as a good target. Goal at this point is still only to demonstrate on some piece of code that this strategy can have big benefits on performance in places.

6/7/18 status - As of mid-April, Guilherme has packaged VecCore for inclusion into LArSoft, and was almost ready to start porting his first vectorization demo back into LArSoft. In general, he still needs guidance from the experiments or the project on which code to work on. We have also discussed a new strategy that involves re-structuring algorithms so that they are intrinsically more vectorizable. He will report status at the LArSoft Coordination Meeting on June 19.

7/16/18 update: [Status update](#) was given. Next code to test on is some DUNE TPC modules - grid machines are quite old, which presents obstacles.

8/3/18 update: Since June, Guilherme has been focused on another project with an Aug 13 deadline. Expect him to devote significant time between Aug 13 and the end of Sept.

Separately, Giuseppe Cerati and Mike Wang have been working on vectorizing the hit-finding code within the context of Giuseppe's SciDAC award. The plan is to make these developments within a test environment, then port it back to LArSoft. Algorithmic changes have already reduced the execution time by a factor of 10.

9/4/18 update: Guilherme has made no progress since the last update. The work by Mike and Giuseppe has concluded and achieved a factor of eight speedup of the hit-finding code, though without the use of vectorization. A new student has agreed to contribute to this work, though we have not yet discussed the details of the work to be completed with their advisor.

9/13/18 update: A new project is being launched to investigate the gains of enabling x86 SIMD vector extensions in selected algorithms and libraries relevant to the DUNE and ProtoDUNE production workflows. The studies may include various types of modifications to some algorithms in order to optimize the gains achieved. We hope the results to be immediately applicable within a LArSoft production environment running on grid resources. The work will be carried out by Agathangelos Stylianidis under supervision of Nektarios Benekos, and in consultation with the LArSoft team and other members of DUNE / ProtoDUNE offline organizations.

10/11/18 update: There was an initial presentation at the [9/25/18 LArSoft Coordination meeting](#).

12/4/18 update: Agathangelos Stylianidis has left research for a job in industry. While Nektarios Benekos has two potential PhD CS students interested, his own status is not clear, so we may not be able to count on him for progress. Guilherme has made no progress and is no longer on the project.

- 4) **Done - Optimization / profiling work.** Profile production LArSoft workflows to identify problem areas and potential solutions. This work includes:
- a) setting up an application to run LArSoft
 - b) identifying a major production workflow that will be the target of the profiling work
 - c) preliminary profiling of LArSoft applications, looking for opportunities for optimization if provided

d) promoting ease of use for profiling tools. --

<https://cdcvs.fnal.gov/redmine/issues/17921>

Resources: Soon Jun

2/19/18 update: LArSoft team has examined profiling results for ProtoDUNE, 35t and FD production workflow chains compiled by Soon Jun Yun. Noted that TensorFlow was the dominant consumer of resources in reconstruction phase, with few other large targets. At the 2/14/18 Offline Leads meeting we asked the experiment to provide a person to review the results in more detail with LArSoft team and identify actionable bottlenecks.

3/7/18 update: Thomas Junk will look at the profiling results (both memory and CPU) available at <https://g4cpt.fnal.gov/LArSoft.html>

From a LArSoft perspective, this work is now almost done. Soon Jun is providing help in easing the use of profiling tools.

6/7/18 update: The results of this work were given to DUNE in March. Since then, two changes have been implemented to address a couple of the most significant problems. The first is in `cheat::BackTracker::TrackIDToSimIDE`, which is called extensively in the mergeana stage of the production processing chain. The change reduced the per event execution time by 10 minutes.

The second change was also in the back-tracker code which associates tracks to particles, and involved a similar architectural problem. The relevant code was completely replaced and moved to a new `ParticleInventory` class as part of a more extensive re-architecture of the back-tracker code.

A third target was identified in the PMA code where the distance between points is calculated. This has not yet been changed.

The profiling also provided information about memory usage, but the experiment has been unable to study this in sufficient detail to identify targets for optimization.

Aside from possible further consultation with the experiment on interpreting results, the LArSoft portion of this work is completed.

- 5) **Done - Error handling policy for LArSoft.** The policy should also prescribe what common conditions constitute an “error” versus a “warning”, etc. An education campaign will then be needed to disseminate this information.

<https://cdcvs.fnal.gov/redmine/issues/12778>

Resources: Katherine Lato, Erica Snider

3/7/18 - Work on this item was deferred to the second quarter of 2018.

6/7/18 - Since error handling relies on art exceptions, asked the art personnel to update/release [ExceptionInformation](#). Started drafting a policy based on [CMS_SWGuideEdmExceptionUse.pdf](#)

9/6/18 - The LArSoft team drafted an [error handling policy](#) based on the use of the `art::Exception` class. This is being socialized at various meetings.

- 6) **DUNE-specific solution is DONE -- Introduce support for global wire coordinates.** The requirements for this work are documented in <https://cdcvs.fnal.gov/redmine/issues/11522>. At present, DUNE has implemented an ad hoc solution to provide global wire coordinate functionality. This project will provide a native LArSoft solution. Addresses in part the 2017 long-term priority of the TPC / locality interface and optimization in the Geometry service.

Resources: Leigh Whitehead

3/6/18 status: With the departure of Robert Sulej, Leigh Whitehead becomes the new point of contact. He has been working on CNN imagine ID, one of the two types of reconstruction/event selection algorithms that can benefit from this. The version he was using for the CVN isn't quite complete and

generic. He looked for one that Robert Sulej had mentioned, and found the [current version of the code](#). It works in the wire channel / TDC space as required for the image production, but might be useful for standard reconstruction too if it is included at the early steps and continued through the clustering and tracking.

LArSoft goal is to have this "generic" version use the existing LArSoft infrastructure, and then infer a useful interface to make code like that easier to write.

6/15/18 update: The experiment reiterated the importance of this work, and their commitment to making the DUNE-specific solution more generic so that it can become part of the native LArSoft geometry framework. They do not have the resources to work on it at this time, so need to wait until at least September, and in particular, until after ProtoDUNE starts running. We should follow up with them at that time.

8/3/18 update: no change, still waiting until after ProtoDUNE starts running.

9/4/18 still waiting until after ProtoDUNE starts running.

10/23/18 - asked for update since ProtoDUNE has started running.

11/13/18 - asked again for update.

11/29/18 - There were several people who left the project. Mike Wallbank implemented a version of this for DUNE and the 35-ton, but is no longer involved in this work. Leigh has implemented a global wire scheme for the CVN event selection and while it may be useful to generalize it so it can be used by all larsoft reconstruction algorithms, no one in protoDUNE is able to drive this effort at this moment given protoDUNE responsibilities. We will keep the [redmine issue](#), but will no longer be tracking the effort as part of the LArSoft work plan.

- 7) **Ongoing - Architecture-dependent libraries.** Support is needed to allow coexisting builds that include / exclude, support for different computing backends under a given OS flavor.
- a) An immediate goal will be to allow the standard setup procedure to support a generic library as well as one built with avx2, or other vectorization technologies
 - b) More generally, there are a number of areas that might benefit from this feature, such as allowing use of GPU backends when available, or selecting the optimal SIMD instruction set for a given computing node. The types of backend support required needs to be studied, and appropriate tools adapted or developed to allow the required setup procedures.

Resources: ??

3/7/18: Since we are waiting on Spack for this, an issue has not yet been created.

8/3/18: As part of exploring vectorization, Giuseppe and Erica have been working on this issue. After a conversation with Mike Kirby, they will write an email to FIFE support about how to target jobs to particular nodes based on the type of vectorization extensions they support. Plan is to take some of Giuseppe's jobs (after being ported back to LArSoft) and want to demo running some of those. We also discussed alternative approaches to adapting the back-end libraries to the hardware a job is running on that does not depend upon support from the job

submission system. We have not settled on a solution on the environment configuration end, though surmised that an additional ups qualifier might suffice for our tests.

9/4/18 update: we hope to integrate a portion of this work into the vectorization project by testing computing power of Fermigrd as a function of the vector extensions used, and working with Fermigrd operations to allow targeting of jobs to nodes based on the available vector extension in order to complete the tests. Note, this may not work.

8) **Done - LArG4 re-factoring work.**

- a) Phase 1: Continuing with existing plan to deliver equivalent functionality to the current code, but re-structured so as to de-couple various models from LArG4 and, where possible, from Geant4, and to expose the result of energy deposition by Geant4, and more generally by upstream particle interaction simulation.
- b) Phase 2: Abstract the anode simulation model, which is required to handle conventional single-phase, dual-phase, and pixel based detectors within the common framework.

<https://cdcvns.fnal.gov/redmine/issues/14454>

Resources: Hans Wenzel and William Seligman for phase 1.

2/21/18 There is a May 1 deadline from MicroBooNE for incorporating this code into MCC9, their next major MC challenge / production processing run. Hans expects that the first stage of work (which meets MicroBooNE's needs) will be completed on that time-scale, though there appears to be some risk in meeting the deadline. Bill Seligman, who is working on the experiment-side component, is ready to test the downstream changes. The re-factoring is currently blocking Wirecell integration testing, so MicroBooNE is proceeding with a temporary backup plan that takes considerable shortcuts in achieving the re-factoring so that testing can take place.

3/7/18 update - Hans Wenzel gave a [LArG4 refactoring status report](#) at the Feb. 27 LArSoft Coordination Meeting. The presentation highlighted the relationship with the Geant4 Collaboration, and discussed various new features of the GDML configuration, such as defining step limits for specific volumes. The G4 collaboration has implemented our requests for a new interface at the energy deposition level, eliminated performance bottlenecks, and improved cross sections e.g. for kaons. An output data translation layer is still required. No time estimates or task lists for completion were provided, though Hans informally mentioned that the changes related to Geant4 will be completed on the timescale of a few days to a week. The downstream code is ready for testing, although experiments will need to update detector-specific GDML files in order to participate.

5/14/18 update - LArG4 refactoring: Initial testing of a fully re-factored LArG4 and Geant4 within a stand-alone testing environment has been completed. Further testing requires integrating the re-factored LArG4 into LArSoft. This work is now under way. Two of the three downstream modules needed to complete the re-factoring have also been completed, an electron drift module written by Bill Seligman, and a fast photon propagation module, written by Wes Ketchum. Some work may be required to ensure that the necessary workflows are easily adapted to the needs of all experiments. The next phase of this project will address light from the Large Electron Multiplier in dual-phase LArTPCs. A later phase will address the issue of integrating multiple external detector types into the simulation, as will be required for DUNE ND. We note that the current re-factoring introduces some of the necessary concepts into the geometry.

8/3/18 Update: The first stage of Phase I of the LArG4 re-factoring was deployed on July 31 as LArSoft v07_00_00. This version does not include support for AuxDet's, so has partial functionality relative to the full Phase I specification. The missing functionality will be included

in a second stage deployment expected soon. Details on migrating to the new code can be found on the LArSoft wiki in the [LarG4 Migration Notes wiki page](#). Hans Wenzel discussed details of the newly re-factored code at the July 31st LArSoft Coordination Meeting: see [Status and deployment plan for LarG4 refactoring](#). The discussion at that meeting suggested a few areas where follow-up is required.

9/4/18 update: Phase I deployment remains incomplete due to missing documentation and an AuxDet handler, but should work for MicroBooNE's MCC9 release. Work on Phase 2 is proceeding -- abstracting the anode region simulation and handling of "S2" light generated in the dual-phase electron extraction and amplification region.

- The S2 light work is being performed by DUNE collaborators, and is almost completed
- Erica Snider met with Paul Russo and DUNE developers Christoph Alt, Bea Oregui and Paul Russo to refine the design and interfaces for the anode region abstraction. The newly adjusted design will accommodate dual-phase, single-phase and pixel detector workflows as configurable elements without encumbering any detector with features or elements not essential for that detector. Some interface work has already been completed.
- The Phase 2 work can be staged so as to accommodate the on-going integration of the Wire-Cell based electron drift and signal induction models for MicroBooNE.

10/23/18 update: Hans Wenzel provided an update on the [Status of LArG4 refactoring](#) at the October 23rd LArSoft Coordination meeting.

11/13/18 update: After a close-out meeting with management on 11/8/18, this project had four small items left. Two were documentation, estimated two days. One data product class that needed to be reviewed. It is already in use, so major changes aren't expected. And the last item was closing out the redmine issues that the restructure was supposed to fix, to ensure that it did fix them.

9) **Ongoing - SPACK - new build system for art and LArSoft** (deferred from 2017).

Purpose:

- To migrate to a standard set of build tools that have broad community support
- To address portability and configurability issues raised by experiments with the current build system
- To allow continued full support of Mac OSX as a LArSoft development platform within the context of Apple's System Integrity Protection system

<https://cdcvs.fnal.gov/redmine/projects/spack-planning/issues>

Resources: Chris Green,

2/19/18 update: Milestones are being tracked under a Spack Planning instance in redmine -

https://cdcvs.fnal.gov/redmine/projects/spack-planning/issues?set_filter=1

<http://cdcvs.fnal.gov/redmine/issues/15313>

6/7/18 update: Chris Green presented [Status of Spack Development / Migration](#) at the May 22nd LArSoft Coordination meeting. Working on a Minimum Viable Product which is not a release, but to allow experiments to experiment with it.

8/3/18 update: The Minimum Viable Product was not completed by August 3rd as planned due to changes needed in SPACK dev. Chris Green will present [status at the August 14th LArSoft Coordination Meeting](#).

9/4/18 update: Chris Green delivered the Minimal Viable Product on Aug 29. The goal for the MVP is to provide a set of features that are sufficient to allow the testing needed to determine if a full migration will be successful.

10/23/18 update: Status update for October is in [HSF packaging meeting presentation](#). Also working on reformulating SpackDev as a Spack extension, which should relieve the performance issues associated with the current need for multiple concretization operations.

10) Ongoing - Re-architecture of art services in LArSoft to ensure thread safety.

Gianluca Petrillo, Erica Snider and Jason Stock met with art team leader Kyle Knoepfel in December 2017 to continue the discussion of art service use cases within LArSoft in the context of thread safety. The discussion concluded with an agreement on how to re-architect a certain class of services (such as those that tag event-by-event transients in various electronics channels, and the calculation of event trigger times to be used in the global clock service) such that they are stateless, and therefore inherently thread safe, while still providing the capability of on-the-fly updates. The LArSoft team will work on an execution plan for this project.

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

Resources: Erica Snider, Mike Wang, Paul Russo

2/26/18 update: Need to know what LArSoft has to do to be ready. Have had meetings to discuss the models with the art team and LArSoft.

3/7/18 - Have identified a person to work on this, and will develop a plan over the next couple of weeks. The plan will be presented at a LArSoft Coordination Meeting.

5/8/18 - The person above was pulled away for a priority project elsewhere. We expect them to return to this project around the end of June.

6/7/18 - art v3, which supports multi-threading, has been released. Need education from art team on what multi-threading in art looks like, how we should use it, etc. Information is available at: <https://cdcv.s.fnal.gov/redmine/projects/art/wiki/Wiki#Multithreaded-processing-as-of-art-3> This development leads to the natural follow-on project to this item, which is to develop a test that maps the memory usage vs the number of threads.

8/3/18 - Erica met with Mike Wang and Paul Russo on July 31 to discuss thread safety in LArSoft. We discussed the general strategy of storing service state in the event, the complication of the LArSoft service model that prohibits service providers from knowing about the event, and various details related to the procedure for ensuring that code is thread safe. Paul was involved with some of the underlying work within art to make services thread safe, so will assist Mike with design issues that may arise. Mike's time allocation for LArSoft work is 50%. He will start after returning from vacation, and wrapping up some minor things related to his work for Giuseppe.

9/4/18 - no status update

11/8/18 update - Michael Wang presented [Status and plans for multi-threading in LArSoft](#) at the November 6th LArSoft Coordination meeting. His work will focus first on the deconvolution code in LArSoft, and will enable multiple events to be deconvolved in parallel.

11) Done -- Support for transitioning code to art-independent repositories, making them available to run in external frameworks such as gallery, in order to meet the growing demand to run LArSoft code in art-independent contexts.

2/19/18 status:

- a) LArSoft has [a written policy for managing this code](#).

- b) Repositories for the code have been provided.
 - i) <https://cdcv.sfnal.gov/redmine/issues/18302>
- c) LArSoft has to move code to these repositories after we move to art 2.10.
 - i) <https://cdcv.sfnal.gov/redmine/issues/17194>
 - ii) <https://cdcv.sfnal.gov/redmine/issues/18281>
 - iii) <https://cdcv.sfnal.gov/redmine/issues/17179>
- d) Once done, there will be a presentation at a Coordination Meeting.

Resources: Gianluca Petrillo, Lynn Garren, Erica Snider, Saba Sehrish.

4/10/18 status

With Gianluca's planned departure, the relevant tickets ([#17179](#) and [#18281](#)) have been assigned to Saba Sehrish.

6/7/18 status:

This work is almost done. There will be a presentation soon at a Coordination Meeting.

8/3/18 status:

Planning on 8/14/18 presentation.

9/4/18 status:

August presentation didn't happen, but remaining work is expected to complete soon.

10/11/18 Status

Presentation at the [October 9 LArSoft Coordination meeting](#). If experiments do not see a ticket for some algorithm code that they want to use inside gallery environment, they are to submit a ticket, like for other items. So this work item is closed.

- 12) (Added June 14, 2018) **Ongoing - Migrate to GitHub and pull-requests.** Investigate moving the LArSoft central repositories from Redmine to GitHub, along with those collaboration services offered by GitHub that make the migration cost effective. Transition the current model merging to develop by developers and the Code Management team to a pull-request system where authorized experiment personnel handle requests after mandatory CI testing prior to merging. Provide all necessary infrastructure needed to support this system

7/12/18 update: Met with line management to discuss this transition. Adam agreed to talk to Liz about having CMS experts share their experience with the migration and the use of GitHub. Decided it makes sense to look at this from art and LArSoft perspective at the same time.

9/4/18 Determined that migration of Redmine wiki to GitHub is feasible. Decided to include migration to a pull-request system as part of the project in order to alleviate the workload on code management resources with a more scalable approach.

Socialized this decision at the Aug 28 LArSoft Coordination Meeting and discussed in more detail at the Aug 30 Offline Leads Meeting.

Longer-term priorities

We welcome input from the Steering Group to help shape the long-term priorities of LArSoft.

A) Concurrency -

- Purpose:

- To use multi-threading to address memory usage issues and provide flexibility in resource utilization - working with art on this, starting with understanding thread safety issues in the LArSoft use cases for art services.
- To introduce vectorization to those components of the code where speed improvements could be obtained, and to make use of currently unutilized resources - see short-term item above
- Resources: LArSoft team and (proposed) experiment effort

B) Architectural changes to define algorithm interfaces for layered algorithms.

- Purpose: Provide a well-defined set of interfaces for layering algorithms within a framework that allows for run-time configurability below the level of art modules, e.g., the art “tool” currently under development
 - The primary target for this change are the algorithms used to generate, process and deconvolve raw signals from the detector
 - A major thrust is to provide a set of tools that will allow algorithms that are currently detector dependent to be generalized and incorporated into the core LArSoft code
- Resources: The LArSoft team and (proposed) experiment effort.
 - DUNE has already contributed effort to a similar project

C) TPC topology / locality in the Geometry service.

- Purpose: To provide a mechanism to efficiently determine the TPCs adjacent to a given TPC, as would be needed to follow tracks and showers that cross TPC boundaries
- Resources: The LArSoft team and (proposed) experiment effort
- <http://cdcv.s.fnal.gov/redmine/issues/9818> -- requested by DUNE
- 11/15/17 - asking DUNE for priority call on this.
- 9/13/18 - DUNE can provide a method to give neighboring TPCs. DUNE will check to see if this is still needed. If not, this item will be dropped.

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](#).