

The architecture of LArSoft toolkit (draft)

The LArSoft project

November 30, 2015

Contents

Introduction	2
Purpose	2
Scope	2
Architecture	2
Overview	2
Simulation	6
Reconstruction	7
Testing	9
Extensibility	9
Data products	9
User code	10
External libraries	13
References	15
Comments	15
Sunday 7:51:39, Ruth Pordes ruth@fnal.gov (“Re: First draft of the architecture description document”)	15

Introduction

Purpose

LArSoft toolkit is designed to enable simulation, reconstruction and physics analysis of data from any detection system based on Liquid Argon TPCs. Its common tools and algorithms render the development and analysis process more uniform across the Experiments. LArSoft is extensible to accommodate evolving Experiments' needs and adoption by new Experiments.

This document provides an overview of the architecture of LArSoft toolkit, including its relationship with the surrounding software environment. The internal flow of the different subsystems is also described. The document intends to capture and convey the significant architectural decisions which have been made on the system, that reflect into the current implementation or drive its development.

Scope

This document describes the architecture of LArSoft to date. It includes description of the communication protocols with libraries it relies on and with packages LArSoft cooperates with. The design and flow of different components is described.

The document is aimed to first-time LArSoft users and developers who want to gain insight of the areas covered by the toolkit, the structure and general organization of the components, and the relationship between components.

It also contains useful information for developers of experiment code who are going to design new experiment code, and it is a suggested introductory reading for developers aiming to add new tools to LArSoft or to extend existing ones.

Finally, it is provided as a snapshot of the current status and principles of LArSoft design, that can be used as a reference for review and as a starting point to define future developments.

This document employs some commonly used LArSoft elements to exemplify flows and connections, but it does not try to exhaustively describe each, or any, of the single elements.

Architecture

Overview

To provide the best solutions for LAr TPC simulation, reconstruction and analysis of data, LArSoft provides a mix of built in tools and algorithms and

interfaces to other existing libraries. Figure 1 illustrates the relation between LArSoft and these libraries.

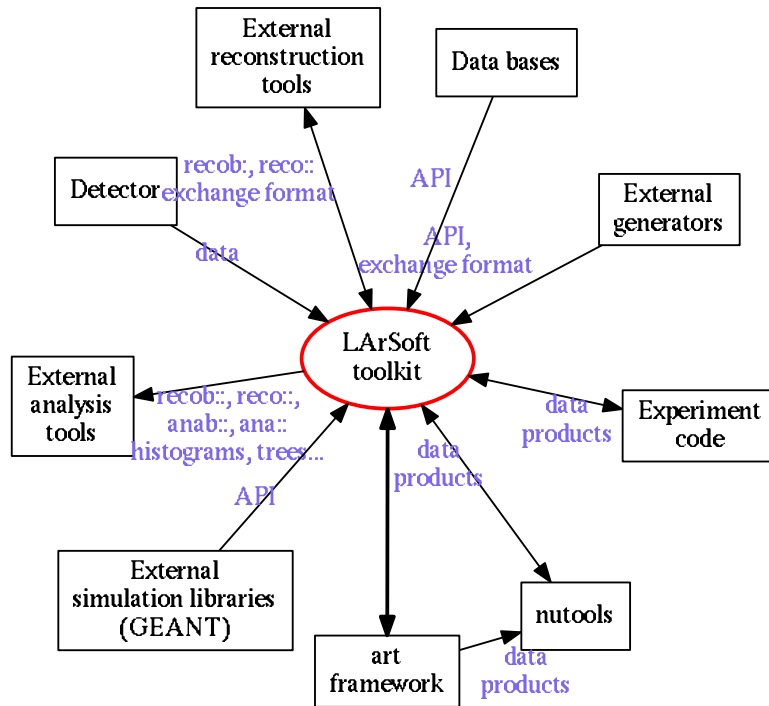


Figure 1: Relationship between LArSoft and other packages and libraries

LArSoft is designed to rely on the *art* framework[ref]. This framework provides an event data model, centralized configuration, data and configuration persistency, management of user algorithm through “modules”, exception handling, and more. LArSoft takes advantage of the libraries *art* framework depends on, by using them directly:

- FHiCL language[ref] to propagate the configuration to its components
- message facility[ref] to regulate text output to console
- CLHEP, ROOT, Boost libraries as needed

LArSoft also shares a common platform, *nutools* [ref], with other neutrino experiments (NOvA). This library provides LArSoft with some basic event display facilities and simulation data structures.

LArSoft interactions include:

- experiment detector data, through customized input modules converting data into LArSoft data classes
- external event generators (e.g., CRY[ref], GENIE[ref]), via API or through HEPEVT exchange format; additional event generators are natively implemented in LArSoft
- GEANT[ref], an external detector simulation library
- data bases, via direct connection or web proxy (*libwda* [ref])
- external reconstruction tools (e.g., *pandora*), through a LArSoft interface
- custom analysis tools, that can use LArSoft data classes directly or tailored data formats produced by custom LArSoft modules
- experiment code, written in the form of LArSoft algorithms, modules and services

LArSoft encompasses a collection of tools that can be roughly groups in the following categories:

1. simulation
 2. reconstruction
 3. analysis
- display

The typical full processing chain (fig. 2) includes a reconstruction and an analysis sequence. For simulated events, a preliminary simulation sequence can be run. Processing chains are defined by the experiments according to their needs. LArSoft inherits the flexibility from the *art* framework, that provides users with the flexibility of choosing and arranging processing modules at will, with the only limitation that each module must be provided with all the information it needs to operate. The same module can also be executed multiple times, with different configurations.

The event display is capable of showing many of the data classes from the simulation and reconstruction steps, and it includes a limited ability of running modules with different configuration at run time.

LArSoft also provides an extensible set of data structures describing objects involved in many levels of the physics analysis, e.g., the time-dependent shape of signal from a photon detector, a simulated neutrino or a reconstructed electromagnetic cascade. The use of these common structures is key to flexibility, allowing to replace and directly compare algorithms that use the same data structures.

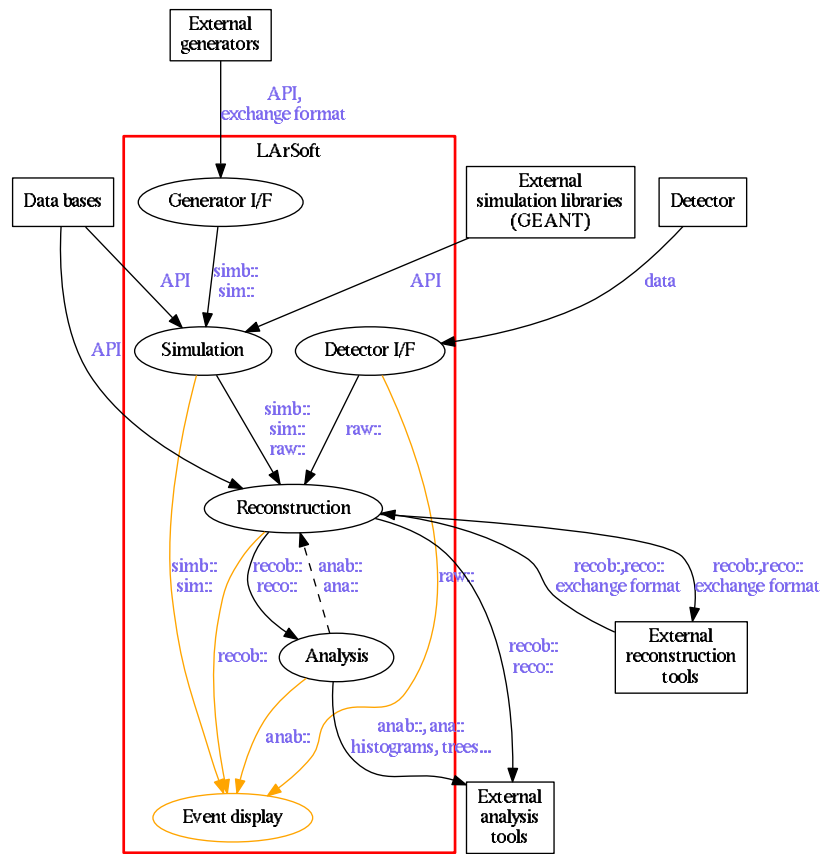


Figure 2: Components of LArSoft and their interaction with external libraries

Simulation

The purpose of LArSoft simulation is to describe a realistic response of the detectors to a known physics event (“truth”). Since the result of the simulation should be equivalent to the output of the detectors, this result is represented by the same data classes. The truth information, not available from the detector, is produced and stored in additional structures.

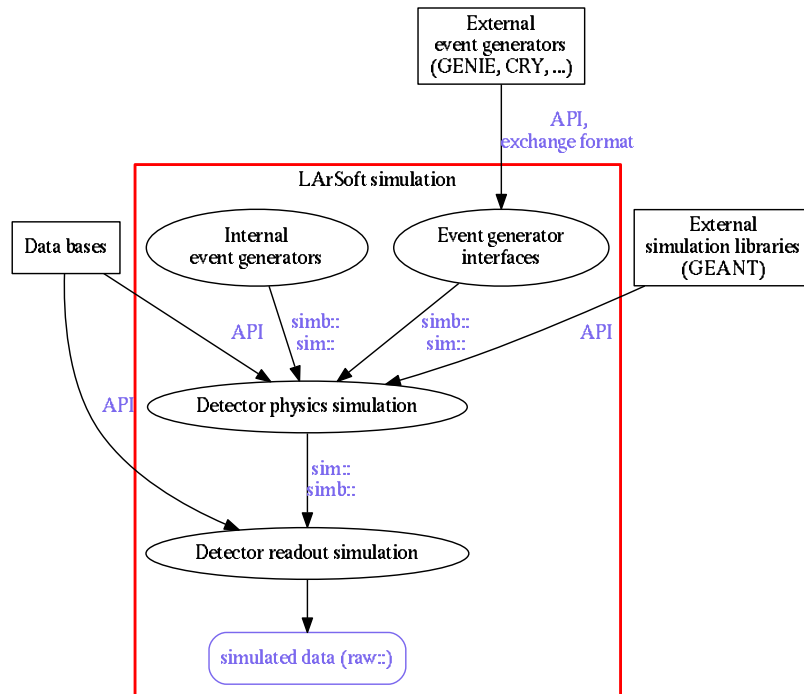


Figure 3: LArSoft simulation flow

The complete simulation chain is summarized in fig. 3. The process is typically described as three steps:

1. event generation
2. detector physics simulation
3. detector readout simulation

The physics event can be generated by an external program or library. LArSoft interfaces directly to GENIE generator (neutrino interactions) and CRY (cosmic

rays). It can also read a generic HEPEVT[ref] format. In addition, LArSoft provides built-in generators to simulate single particles, Argon nucleus decays, and more.

The detector physics simulation includes the interaction of the generated particles with the detector, and the propagation to the readout of produced photons and electrons. This part of the simulation relies on GEANT4 for the interaction of particles with matter. Photon and electron transportation to the readout are implemented in built-in code. Detector parameters (e.g., the intensity of the electric field) can be acquired from the job configuration or from a custom data base.

The last step transforms the physics information, electrons and photons, into digitized detector response, including the simulation of electronics noise and shaping. This is typically implemented with experiment-specific code.

Reconstruction

The reconstruction phase provides standard physics objects to describe the physics event. Reconstruction delivers objects with different level of sophistication and from different steps, as for example hits describing localized charge deposition as detected on a wire, down to a complete hierarchy of three-dimensional tracks. These objects are handed over for further analysis.

Starting from detector response, either real or simulated, there are many possible patterns of analysis. The more “traditional” one (fig. 4) starts with the calibration of the signals, attempting to suppress noise and revert electronics distortions, and then it proceeds with the reconstruction of charge deposition on a single TPC wire (*hits*), to *cluster* them in groups lying on the same wire plane, and finally with combining clusters from different planes in trajectories (*tracks*) and particle cascades (*showers*), connected by interaction points (*vertices*). The hierarchal connection between them is called a *particle flow*. Many options are implemented in LArSoft for each of these steps, that are interchangeable as they use the same input and output classes.

During any of these steps the detector and data acquisition parameters can be acquired from experiment data bases.

Any external library that utilizes LArSoft data classes to receive inputs and deliver results is also fully interchangeable with the algorithms implemented in LArSoft. A noticeable example is the *pandora* pattern recognition toolkit, that accepts LArSoft hits as input and can present its results in the form of LArSoft clusters, tracks and particle flow objects.

Further common analysis steps are the calibration of the energy deposited in liquid argon by the interacting particles and their identification as specific types (e.g., muons, protons, etc.).

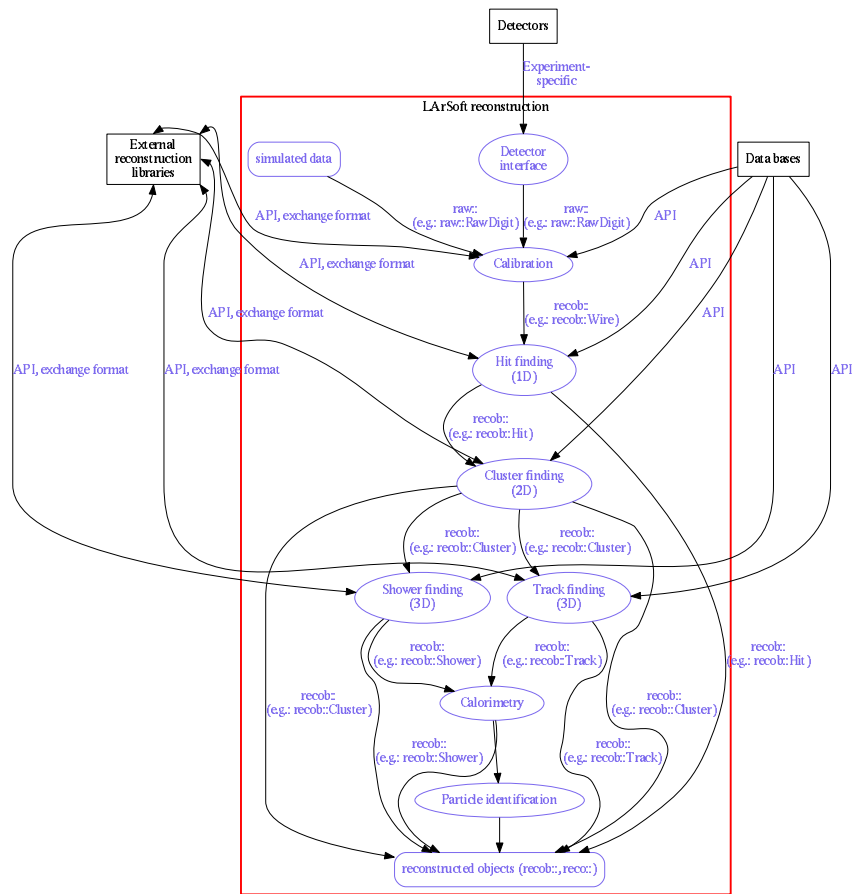


Figure 4: LArSoft “traditional” reconstruction flow

Testing

LArSoft development model allows multiple contributors to modify the code at the same time. This model can create conflicts and dysfunction in the code. Tests are instrumental to the early detection of such defects. LArSoft includes tests at two levels, called *unit tests* and *integration tests*.

Unit tests exercise a limited part of the system, typically a single algorithm. Ideally a unit test for an algorithm should test all the functions of that algorithm. In practice, tests for complex algorithms tend to set up and test a few known typical cases.

Integration tests involve the framework and one or more processing modules. These tests can reproduce real user scenarios, for example a part of the official processing chain of an experiment, and they can compare new and historical results. LArSoft tools allow these tests to be run at any time, and a standard suite of tests is meant to be automatically and periodically run.

Extensibility

The extensibility of LArSoft is largely based on the underlying framework, *art*. The *art* framework processes physics event independently, executing on each of them a sequence of modules. The framework also provides a list of global “services” that modules can rely on. Examples of services implemented by LArSoft include the description of detector geometry and channel mapping, the set of detector configuration parameters, and access to TPC channel quality information.

Our description focuses on extensibility in terms of new persistable data structures, of new algorithms implemented in LArSoft and of using external libraries.

Data products

LArSoft provides a basic set of persistable data classes. Each class is associated to a simple concept and a set of related quantities. For example, `raw::RawDigit` describes the raw data as read from a TPC channel; `recob::Cluster` describes a set of hits observed on a wire plane; `anab::Calorimetry` contains information about calibrated energy of a track.

A *data product* is a class that:

- is simple: contains just data and trivial logic to access it; more complex elaborations belong to algorithms
- contains only members from a small selected libraries: C++ standard library is highly recommended; ROOT classes are also accepted

- is not polymorphic

Limitations to ROOT I/O system impose restrictions on the types of allowed data members, e.g., on the set of supported C++11 containers. Relations between data products are expressed by *associations*. Associations are data products provided by *art* that can relate a data product, or an element of it, to another element from another data product. Examples of use in LArSoft include the association between a reconstructed hit and the calibrated signal it's reconstructed from, and between a cluster and all the hits that constitute it.

Data products have a fundamental structural role: they act as messages to be exchanged between algorithms. As such, they are also the format in which most of the results are saved. This allows to arbitrary split the processing chain in multiple sequences of jobs.

User code

Algorithms constitute, together with data products, the heart of LArSoft, and the ability for user to add their own algorithm is central to its design. In fact, LArSoft algorithms differ from users' algorithms only in the judgment that their purpose is considered of wider interest than just for the single user. Indeed, most of the algorithms in LArSoft were written by users to solve a specific problem, and then adopted into the common toolkit. LArSoft encourages users to produce algorithms that perform correctly on any liquid argon detector, and to integrate them into LArSoft itself.

The preferred model for algorithm structure is represented in fig. 5. We refer to it as *factorization* model. The underlying principle it is that the algorithm must be independently testable and portable, using the minimal set of necessary dependences. This also allows for the algorithms to be used in contexts where the *art* framework is not available, provided that some other system supplies equivalent functionalities as, and only when, needed. The model is made of two layers:

1. the algorithm, in the form of a class that
 - is configurable with FHiCL parameter sets
 - consumes LArSoft data products as input
 - produces LArSoft data products as output
 - has the minimal convenient set of dependencies
 - elaborates a single event or part of an event at a time
2. a module for the *art* framework, that:
 - owns and manages the lifetime of one or more algorithm classes

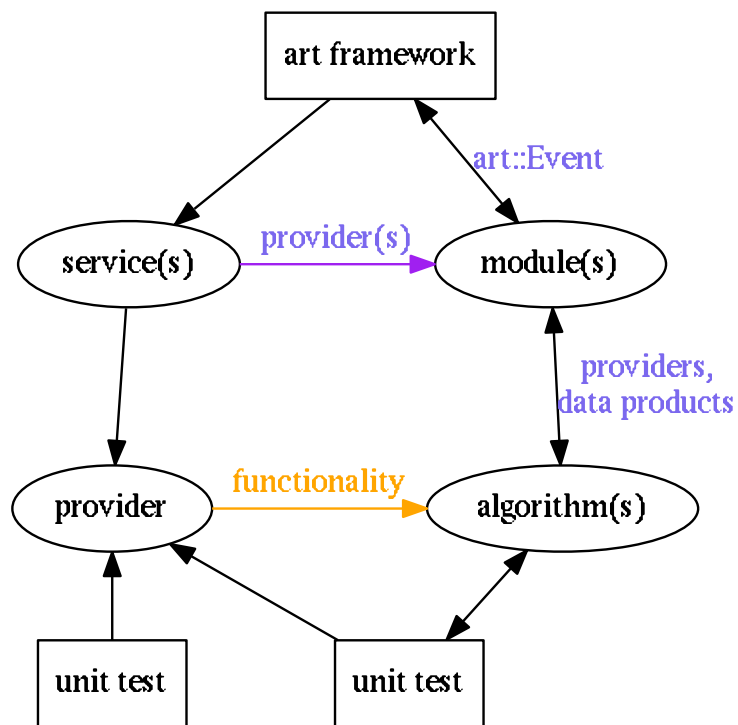


Figure 5: LArSoft algorithm and service model

- provides the algorithm(s) with the configuration, the data products and the information it needs to operate
- delivers algorithm output to the *art* framework

Since algorithms often rely on services, the services also need to follow the same factorization model and be split in:

1. a *service provider*, in the form of a class that:
 - is configurable with FHiCL parameter sets
 - has the minimal convenient set of dependencies
 - provides the actual functionalities
2. a service for the *art* framework, that:
 - owns and manages the lifetime of its service provider
 - provides modules with a pointer to the provider
 - when relevant, propagates messages from the framework (e.g., the beginning of a new run) to the provider

The module is also responsible of communicating to its algorithms which service providers to use. Algorithms exclusively interact with service providers rather than with *art* services.

Other important guidelines for the development of algorithms are:

- interoperability: they should document their assumptions in detail, and correctly perform on any detector if possible
- modularity: each algorithm should perform a single task; complex tasks can be performed by hierarchies of algorithms
- maintainability: they should come with complete documentation and proper tests

Figure 5 shows that if algorithms are not framework-dependent, their unit test can also be framework-independent. Therefore, not only those algorithms can be developed in a simplified, framework-unaware environment, but they can also be tested in that same development environment. In other words, the full development cycle, of which testing is an integral part, can seamlessly happen in the same environment.



Figure 6: Interaction between LArSoft and an external library

External libraries

We call “external” any library that does not depend on LArSoft, with the possible exception of its data products. Examples in this category are GENIE, GEANT4, and *pandora*.

LArSoft’s modularity can accommodate contributions from external libraries into its workflow (fig. 6). The preferred way is to use directly the external library via its interface. This requires an additional interface module between LArSoft and the library, in charge of converting the LArSoft data products into a format digestible by the external library, configuring and driving it, and extracting and converting the results into a set of LArSoft data products.

This is exemplified in the interaction between LArSoft and *pandora* (fig. 7): *pandora* uses its own data classes for input hits, particle flow results and geometry specification. A base module exists that reads LArSoft hits, converts them into *pandora*’s, translates geometry information, and recreates out of *pandora* particle flow objects LArSoft clusters, tracks, vertices, and more.

This approach has relevant advantages: it can be fairly fast; it allows a precise translation of information; it provides the greatest control on the flow within the library; it defines and tracks the configuration of the external library. Its greatest drawback is the need for the LArSoft interface to depend on the external library. If this limitation is not acceptable, a more independent communication channel can be established via exchange files. In this case, LArSoft interface translates data products into a neutral format, possibly based solely on ROOT objects or on a textual representation, and back into data products. The external library is in charge of performing the equivalent operations with the library data format. This is for example the generic communication mechanism with event generators that support HEPEVT format. The strong decoupling comes at the price of a fragmented execution chain and the burden of additional configuration consistency control, for example to ensure that a consistent geometry was used for the information (re)entering LArSoft.

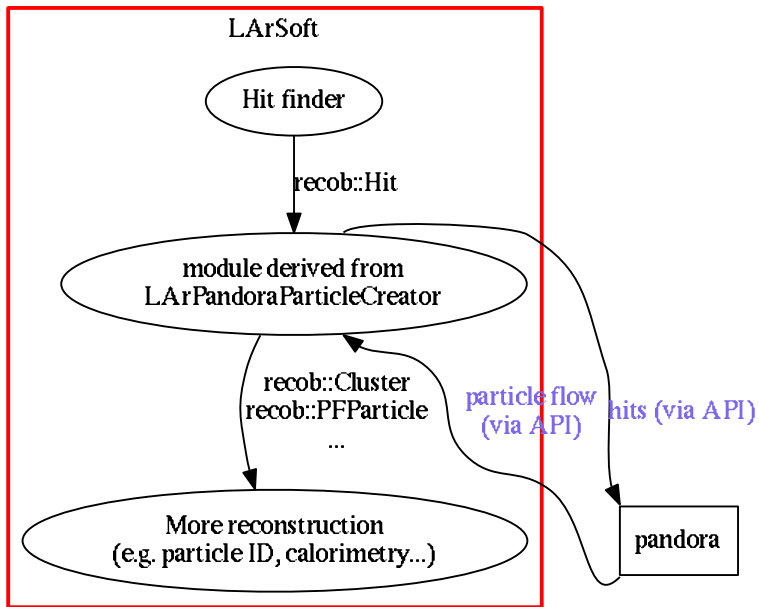


Figure 7: Interaction between LArSoft and *pandora*

References

[ref] *art* framework
[ref] FHiCL language
[ref] message facility
[ref] *nutools*
[ref] CRY
[ref] GENIE
[ref] GEANT4
[ref] *libwda*
[ref] HEPEVT format

Comments

Sunday 7:51:39, Ruth Pordes ruth@fnal.gov (“Re: First draft of the architecture description document”)

[Q 001] does the scope include human interfaces as well as software?

[A 001.1] [GP] I thought mostly not, but I am not completely sure what human interface includes. To be clarified. (*TODO*)

[Q 002] *nutools* event display facility and simulation data structures – still does not make sense to me. Is Visualization one special kind of analysis or does Larsoft have specific interfaces to it?

[A 002.1] [GP] Visualization is a special kind of analysis. But our event display crosses the border with its (limited) ability to *interactively* reprocess the input.

[Q 003] page 4 – can components of the chain be re-executed during a single pass?-

[A 003.1] [GP] I have added a couple of sentences in the previous-to-last paragraph of Architecture > Overview section, that I hope give an answer. The answer is very much in the features of *art*, that I have not covered at all in this text. Should we? (*TODO*)

[Q 004] does event display have a specific meaning - I'll include it in the Requirements glossary – it is different from a generalized visualization and I presume the definition should explain this? Also, if the event display is in *nutools* it is not part of larsoft??? Can we share a glossary in some fashion?

[A 004.1] [GP]

[Q 005] Figure 2. You explicitly mean Detector not DAQ ? Does/shoud daq show up somewhere

[A 005.1] [GP] in practice DAQ products is what we communicate with. It doesn't have to be only that, but I guess that is it effectively what happens. (TODO)

[Q 006] a Fluka interface is in the works with integration hoped for before the end of Dec. Can you include a sentence on this interface?

[A 006.1] [GP] Erica, confirm? (TODO)

[Q 007] page 10. Unit test. These are important. These are not the only tests. I don't see them referred to and perhaps some more specifics might be useful?

[A 007.1] [GP] I added a section about testing. I have added a few words also at the point Ruth specified (at the end of "User code" section). I think it would be good to add a "test" block in one of the high-level diagrams, but I can't figure out where (probably in 2, but how?). Or maybe we have to add a *development model* section?