HSF Community Roadmap for Software and Computing

Peter Elmer - Princeton University

21 June, 2016

A Software “Upgrade” for the HL-LHC?

- The HL-LHC will provide 100 times the current data, with significantly increased data (pileup) and detector complexity.
- Most of the current software, which defines our capabilities, was designed 15-20 years ago: there are many software sustainability challenges
- Estimates of computing needs run faster than Moore’s Law by factors of 3-30(?), but technology change will also make it challenging to exploit Moore’s Law without software evolution.
- Over the past 2 years we have put the HEP Software Foundation (HSF) in place as an umbrella for activities to address these issues.

Where does HSF go from here?

The HEP Software Foundation (HSF) (http://hepsoftwarefoundation.org) was created 1.5 years ago as a means for organizing our community to address the software challenges of future projects like the HL-HLC. An initial set of collaborative activities have begun (see recent HSF workshop at LAL-Orsay). The HSF has the following objectives:

- Catalyze new common projects
- Promote commonality and collaboration in new developments to make the most of limited resources
- Provide a framework for attracting effort and support to S&C common projects (new resources!)
- Provide a structure to set priorities and goals for the work

Where does HSF go from here?

The HSF has demonstrated some initial collaborative activities between people working on different experiments. However what is needed to address the future HEP software/computing challenges (HL-LHC and others) is additional dedicated resources for projects.

There *are* a couple of "common" software-focused projects today which have acquired "new" resources: DIANA-HEP (http://diana-hep.org) and the software work package of AIDA2020. Neither of these was really proposed or funded "as part of" HSF, but they are the kinds of projects we want to foster under the HSF umbrella. How concretely do we go about doing that?

Even more concretely: can we build something resembling a "software upgrade" project for HL-LHC?

Community Roadmap for HEP S&C

As a next step for HSF and as a step towards software upgrades for HL-LHC, we are proposing a **Community Roadmap for HEP Software and Computing**. Such a community white paper should describe a global vision for software and computing for the HL-LHC era and HEP in the 2020s; this should include discussions of elements that are common to the HEP community (LHC community, etc.) as a whole and those that are specific to the individual experiments. It should also discuss the relationship of the common elements to the broader scientific computing communities.

This document could a step in the HL-LHC process over the next years and play a role in discussing possible funding scenaries for a “software upgrade”.

Status

- The CWP roadmap plan was presented and discussed at the HSF meeting at LAL-Orsay. It was generally agreed that this is a necessary next step.
- The proposal for a community roadmap, to be carried out by HSF, was presented to the LHCC.
- Ian Bird is now working with some of us to formulate a charge (for discussion) to kickstart this process.
- Key elements include establishing working groups with the experiments and others, specific charges and actual dates for workshops.

Community Roadmap Process - DRAFT, for discussion

We propose a series of workshops over the next year to build the community roadmap:

- Organization of working groups with experiments by September.
- A half-day (Sun 9 Oct) just before CHEP
- A “kick-off” workshop after CHEP and before Xmas (near CERN?)
- Several dedicated “topical” workshops in Jan-Jun 2017 covering software required in the various areas:
  - Detector Simulation, Triggering, Event Reconstruction and Visualization
  - Data Access and Management, Workflow and Resource Management
  - Physics generators, Data Analysis and Interpretation, Data and Software Preservation
- A final workshop in summer 2017 (near CERN?)

We should build on existing community activities when possible (e.g. DPHEP, Reco Algorithms Forum/CTD, IML).

Detector Simulation, Triggering, Event Reconstruction and Visualization

Challenges surrounding high pile-up simulation, including the CPU resources needed for large statistics samples needed to compare with data from high trigger rates, high memory utilization, generation and handling of the large (min-bias) samples needed to achieve accurate description of high pile-up collision events, and a flexible simulation strategy capable of a broad spectrum of precision in the detector response, from “fast” (e.g. parametric) simulation optimized for speed to full simulation in support of precision measurements and new physics searches (e.g. in subtle effects on event kinematics due to the presence of virtual particles at high scale). Software required to emulate upgraded detectors (including the trigger system) and support determination of their optimal configuration and calibration. Software in support of triggering during the HL-LHC, including algorithms for the High-level Trigger, online tracking using GPUs and/or FPGAs, trigger steering, event building, data “parking” (for offline trigger decision), and data flow control systems. New approaches to event reconstruction, in which the processing time depends sensitively on instantaneous luminosity, including advanced algorithms, vectorization, and execution concurrency and frameworks that exploit many-core architectures. In particular, charged particle tracking is expected to dominate the event processing time under high pile-up conditions. Visualization tools, not only in support of upgrade detector configurations and event displays, but also as a research tool for data analysis, education, and outreach using modern tools and technologies for 3D rendering, data and geometry description and cloud environments.

Data Access and Management, Workflow and Resource Management

Data handling systems that scale to the Exabyte level during the HL-LHC era and satisfy the needs of physicists in terms of metadata and data access, distribution, and replication. Increasing availability of very high speed networks removes the need for CPU and data co-location and allows for more extensive use of data access over the wide-area network (WAN), providing failover capabilities, global data namespaces, and caching. Event-based data streaming as complementary to the more traditional dataset-based or file-based data access, which is particularly important for utilizing opportunistic cycles on HPCs, cloud resources, and campus clusters where job eviction is frequent and stochastic. Workflow management systems capable of handling millions of jobs running on a large number of heterogeneous, distributed computing resources, with capabilities including whole-node scheduling, checkpointing, job rebrokering, and volunteer computing. Systems for measurement and monitoring of the networking bandwidth and latency between resource targets and the use of this information in job brokering. Software-defined networking technologies which enable networks to be configurable and schedulable resources for use in the movement of data.

Physics generators, Data Analysis and Interpretation, Data and Software Preservation

There are many theory challenges in the HL-LHC era, among them are improving the precision of SM calculations, better estimation of systematic uncertainties, and elucidation of promising new physics signals for the experiments. Software needed to make connection between observations and theory include matrix element generators, calculation of higher-order QCD corrections, electroweak corrections, parton shower modeling, parton matching schemes, and soft gluon resummation methods. Physics generators that employ concurrency and exploit many-core architectures will play an important role in HL-LHC, as well better sharing of code and processing between LHC experimenters and phenomenologists. Data analysis frameworks that include parallelization, optimized event I/O, data caching, and WAN-based data access. Analysis software that employs advanced algorithms and efficiently utilizes many-core architectures. Tools and technologies for preservation and reuse of data and software, preservation and re-interpretation of physics results, analysis providence and workflow ontologies, analysis capture, and application packaging for platform abstraction. Future software repositories and build platforms that leverage advances in these areas and improved software modularity and quality control that will allow a broader community of people to effectively contribute to software in the HL-LHC era.

Working groups - example questions to address

In addition to addressing issues specific to a given topic, each group should presumably address questions which cut across boundaries, including:

- What are the specific challenges for the HL-LHC (IF, etc.)?
- What opportunities exist to exploit new or advanced algorithms (e.g. deep learning)?
- How can emerging architectures improve the bang-per-buck and what software evolution is needed to exploit them?
- Which problems are specific to individual experiments and which are common to (for example) the HL-LHC experiments or to HEP and nuclear physics experiments more generally?
- What is required to make common software packages sustainable?

What could an explicit CWP process support?

Going back to the subset of HSF goals I listed earlier:

- Catalyze new common projects
- Promote commonality and collaboration in new developments to make the most of limited resources
- Provide a framework for attracting effort and support to S&C common projects (new resources!)
- Provide a structure to set priorities and goals for the work

The CWP process, an eventual CWP and (simultaneously) the pursuit of specific plans/proposals will support precisely these goals.

Conclusion

- What are your thoughts on this CWP/roadmap proposal?
- Do you see opportunities where such a process and document would fit with other national or international efforts?
- Could a community roadmap (help, hinder) things you are already pursuing?
- How do we put together a process to write a CWP over the next year?
- What should the CWP actually look like?
- What opportunities we can create and pursue which leverage both the process of putting together a CWP and the eventual CWP itself?