NSF S2I2 and the Community Roadmap for Software and Computing

Peter Elmer, Mark Neubauer, and Mike Sokoloff

3 August, 2016

Software and Computing for the HL-LHC Era

The DOE and NSF jointly invest $\approx$ \$35M/year in ATLAS and CMS software and computing, about half in hardware plus operations, about half in software professionals. The LHC funding agencies, worldwide, invest about \$150M/year in these enterprises. In 2014, the LHC experiments used almost 175 PB of tape storage and slightly more disk storage. The event rate anticipated for the HL-LHC era is 100 times greater, and even assuming the experiments significantly reduce the amount of data stored per event, the total size of the datasets will be well into the exabyte scale; they will be constrained primarily by costs and funding levels, not by scientific interest. One long-term goal of a HEP S^2I^2 will be maximizing the return-on-investment to enable break-through scientific discoveries using the HL-LHC detectors.

NSF Scientific Software Innovation Institute (S2I2)

- Part of NSF "Software Infrastructure for Sustained Innovation" (SI2) program, which invests in software as a critical part of the research and education infrastructure
- An NSF S2I2 is a possible path to a "software upgrade" project for HL-LHC in the U.S. Universities
- Investment of $3-5M/year, 5 years, possibility of renewal
- NSF has funded a "conceptualization" (planning) award to prepare a "Strategic Plan" for an eventual S2I2 implementation proposal
- Any plan for an NSF S2I2 must fit in as part of a community roadmap

A Community White Paper

The **Community White Paper (CWP)** will describe a global vision for software and computing for the HL-LHC era. It will discuss issues common to the LHC community and those that are specific to the individual experiments. Many of the topics discussed here will address issues required for a HEP S^2I^2 implementation proposal. These will include

- a broad overview of the grand challenge science of the HL-LHC;
- how new approaches to computing and software can enable and radically extend the physics reach of the detectors;
- **what computing and software research will be required** so that computing and software Technical Design Reports can be prepared several years before Run 4 of the LHC begins; this will include studies of hardware and software architectures and life-cycle processes and costs.
- **identify specific software elements and frameworks** that will be required for the HL-LHC era which can be built and tested during Run 3.
- **organizational issues for the common software and for coordinating research of common interest**, even when the final products will be specific to individual experiments.

NSF S^2I^2 Strategic Plan - I

The separate **Strategic Plan** will identify areas where the U.S. university community can provide leadership and will discuss those issues required for an S^2I^2 which are not (necessarily) relevant to the larger community. Topics will include:

- **where does the U.S. university community already have expertise and important leadership roles;**
- which software elements and frameworks would provide the best educational and training opportunities for students and postdoctoral fellows;
- what types of programs (short courses, short-term fellowships, long-term fellowships, etc.) might enhance the educational reach of an S^2I^2;
- **possible organizational, personnel and management structures and operational processes**;
- **how the investment in an S^2I^2 can be judged and how the investment can be sustained** to assure the scientific goals of the HL-LHC.

NSF S^2I^2 Strategic Plan - II

Using the CWP as a stepping-off point, the Strategic Plan will also address topics specific to proposed U.S. university efforts:

- **the scope of possible projects will be coordinated with other stake-holders, including DOE-funded laboratories** in the U.S. and international collaborators, to assure a coherent community-wide effort;
- the plan will identify possible management structures, mechanisms for continuing assessment, and appropriate education and outreach activities.
- **development, testing and deployment methodologies, validation and verification processes, end usability and interface consideration, and required infrastructure and technologies**;
- the requirements and necessary mechanisms for **human resource development**, including integration of education and training, mentoring of students, postdoctoral fellows as well as software professionals, and proactively addressing diversity and broadening participation;
- **potential risks** including risks associated with establishment and execution, necessary infrastructure and associated technologies, community engagement, and long-term sustainability.

A Nominal Timeline

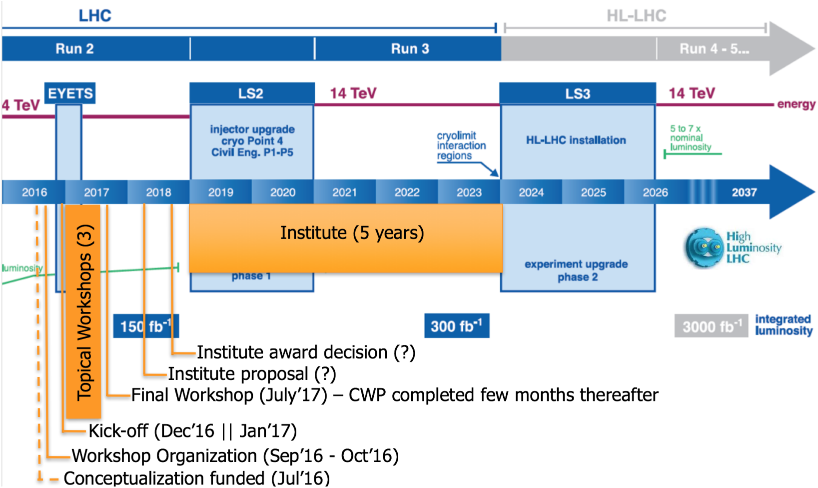

Status

The NSF S2I2 Conceptualization project was funded in July, 2016. We are now beginning to get organized.

The proposal for a general Community Roadmap has been widely discussed with all of the LHC experiments and the HEP Software Foundation (HSF). There is broad support for the idea.

The CWP roadmap plan, to be carried out by HSF, was presented to the LHCC. It fits with the current notion of HL-LHC computing TDRs in $\sim$2019.

WLCG has produced a charge for this CWP to the HSF and the LHC experiments (see separate link) with an aim to complete it by the end of August, 2017.

We are now organizing some dedicated workshops. We will also use sessions at existing meetings when possible.

Some Detector Simulation Issues

The conceptualization proposal identifies **Detector Simulation** as a major focus area for discussion. Challenges related to simulating high pile-up events include

- **CPU resources** for high statistics samples required to compare with real data for preparing triggers;
- **high memory utilization**;
- **flexible simulation strategies** capable of providing a broad spectrum of precision of detector response, from "fast" (parametric) simulation optimized for speed to full simulation in support of high precision measurements and new physics searches.
- software to **emulate upgraded detectors** (including trigger systems) and suppport optimizing their configurations and calibration strategies.
- what **mix of CPUs, GPUs, and/or FPGAs** will provide the best bang-per-buck, including all life-cycle costs.

Goals for today: Expand the list of topics for the simluations working groups to address and begin to think about how the work should be done.

Some Data Access and Management, Workflow, and Resource Management Issues

Data handling systems will need to scale to the multi-exabyte scale during the HL-LHC era. Issues to be addressed will include:

- **analyst access** to data and metadata;
- **workflow management** systems capable of handling millions of jobs running on large numbers of heterogeneous, distributed computing resources;
- tools for measuring and monitoring **networking bandwidth and latency** between resource targets for use in job brokering;
- **software-defined networking** technologies to enable efficient use of resources;
- **event-based data streaming**, as well as dataset-based or file-based;

Goals for today: Expand the list of topics in this area, with an emphasis on the potential role of machine-learning; begin to think about how the working groups might organize themselves to address these issues.

Questions Each Working Group Should Address

- What are the **specific challenges** for the HL-LHC?
- 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 the HL-LHC experiments** or to HEP and nuclear physics experiments more generally?
- How could we organize more effectively both to build common projects and to support experiment specific solutions?
- What is required to **make common software packages sustainable**?

Some Possibly Provocative Questions to Consider

Looking ahead 10 to 20 years, the computing and software landscape will change substantially. Trying to keep up with Moore's Law requires taking advantage of vectorization and many-core or multi-core architectures. "Business-as-usual" cannot keep up with the demands of the higher luminosity and larger number of detector channels project for the HL-LHC era.

- To what extent should the LHC experiments used dedicated hardware ("grid" facilities), national resources (such as those provided by XSEDE centers), opportunistic resources, commercial clouds? What are the **real life-cycle costs**, including hardware acqusition, power, HVAC, space, operating staff?
- To what extent should HEP develop its own software, and to what extent should it **use commercially developed products** and/or services?
- If very high speed networking is required, is it more cost effective to **locate processing cycle resources at a limited number of sites**?
- We are currently developing software for the next generation of hardware. **How should we design algorithms and software in the next 5 years for hardware as we imagine it might exist 10 to 20 years from now**?

Join the Effort !!

Today

- generate 10 **questions for the conceptualization stage** that do not suggest solutions;
- identify 6 areas of **software R&D for common tools** which would benefit from additional funding;
- propose 6 ways to bring **younger members** of the community (graduate students and post-docs) into the process.
- propose 6 ways we could **organize differently** to better achieve our goals (in labs, in universities, nationally, internationally)

Moving forward

- **visit** the web site at `http://s2i2-hep.org/` ;
- **join** the Google Group linked on the right of that page;
- **invite** your colleagues to participate;

Backup slides

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.