

Community White Paper: A Roadmap for HEP Software and Computing

Peter Elmer - Princeton University

13 December, 2016
Computing and Technology Workshop in Vilnius

A Software “Upgrade” for HL-LHC and 2020s HEP?

Looking forward to the next 10 years, we see a number of challenges for HEP software and computing:

- **Scale:** The HL-LHC will integrate 100 times the current data, with significantly increased data (pileup) and detector complexity.
- **Performance/cost:** Estimates of computing needs run faster than Moore’s Law by factors of 3-30
- **Technology/Market evolution:** the return of heterogeneity; technology change will also make it challenging to exploit Moore’s Law without software evolution.
- **Sustainability:** Most of the current software, which defines our capabilities, was designed 15-20 years ago: there are many software sustainability challenges.

Why Software? Software is *the* Cyberinfrastructure




Computer hardware is a consumable.
Software is what we keep, and invest in, over time.

Estimates of Resource Needs for HL-LHC (WLCG)



Data:

- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB



CPU:

- x60 from 2016

Technology at ~20%/year will bring x6-10 in 10-11 years

- ❑ Simple model based on today's computing models, but with expected HL-LHC operating parameters (pile-up, trigger rates, etc.)
- ❑ At least x10 above what is realistic to expect from technology with reasonably constant cost

(Slide from WLCG Workshop Intro, Ian Bird, 8 Oct, 2016)

Back to heterogeneous systems?

Building the worldwide distributed LHC computing grid was largely made possible by the convergence on Linux on (commodity) Intel x86 processors around the year 2000. Building the WLCG at this scale in the heterogeneous workstation era would have been quite difficult. For better or for worse, heterogeneity is returning:

- Diversity of computing processor architectures (general purpose cores vs specialized processors)
- Owned vs commercial/cloud providers
- Some pressure to use systems traditionally designed for other types of applications (e.g. HPC/supercomputer as opposed to HTC/high-throughput systems)
- Possible further commoditizing market pressures (e.g. mobile)

What is software sustainability?

- **Dependent Infrastructure:** Will the infrastructure element continue to provide the same functionality in the future, even when the other parts of the infrastructure on which the element relies change?
- **Collaborative Infrastructure** Can the element be combined with other elements to meet user needs, as both the collaborative elements and the individual elements change?
- **New Users:** Is the functionality and usability of the infrastructure element clearly explained to new users? Do users have a mechanism to ask questions and to learn about the element?
- **Existing Users:** Does the infrastructure element provide the functionality that current users want? Is it modular and adaptable so that it can meet the future needs of the users?
- **Science:** Does it incorporate and implement new science and theory as they develop?

HEP Software Foundation (HSF)

The 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 such as the HL-LHC. 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



HEP Software Ecosystem



IgProf

FroNTier

FairRoot



FastJet

RooStats

EvtGen




Examples, definitely incomplete!



Recent/Nascent Cross-experiment Collaborations

- Experiment frameworks
 - Gaudi, FAIRRoot, CMSSW/Art
- Common Conditions Data Project
 - Discussion/cooperation between ATLAS, Belle II, CMS and LHCb
- Common Software Build and Packaging Tools efforts
 - Working group of HSF comparing HEP and non-HEP solutions
- Cooperation on Reconstruction Software
 - “Connecting the Dots” tracking workshop, HSF sessions
- AIDA2020 (EU funded)
 - DD4hep for detector description, PODIO data model library (LCD, FCC, potentially LHCb)
- DIANA (Data Intensive ANALysis) (NSF Funded)
 - 4-year project on analysis software, including ROOT and its ecosystem

Example collaborative project: DIANA/HEP



The DIANA/HEP project focuses on improving performance, interoperability, and collaborative tools through modifications and additions to ROOT and other packages broadly used by the HEP community.

Website: <http://diana-hep.org>

Google group: <https://groups.google.com/forum/#!forum/diana-hep>

Github: <https://github.com/diana-hep>

Example DIANA talks/activities



dianahep

Team

Activities

DIANA Fellows

Blog

Activities and Research Products

Software products

- [DIANA/HEP organization on GitHub](#)
- [carl docs](#) [DOI: 10.5281/zenodo.47796](#)
- [histogrammar - web page, docs](#)

Papers

- [Carl in the Journal of Open Source Software](#) [JOSS: 10.21105/joss.00011](#)

DIANA Topical Meetings

The DIANA team organizes periodic meetings on various topics of interest, with both invited, external speakers and DIANA team contributions. These and other DIANA meetings can be found in the [INDICO agenda system](#).

Presentations by the DIANA team

- 14 Dec 2016 - Data Plumbing (whiteboard talk), Jim Pivarski, [LPC Coffee Hour](#)
- 5 Dec 2016 - Focus groups: physicist attitudes toward software, Jim Pivarski, [DIANA Meeting](#)
- 14 Nov 2016 - TTreeProcessor: A toy framework for parallel ntuple processing, Brian Bockelman, [DIANA Meeting](#)
- 7 Nov 2016 - Overview of file formats, Jim Pivarski, [DIANA Meeting](#)
- 7 Nov 2016 - Reading ROOT data in Java and Spark, Jim Pivarski, [ROOT Team meeting: follow up with CERN IT](#)
- 30 Oct 2016 - Plotting data on GPUs with Histogrammar, Jim Pivarski, [Parallel Kalman Meeting](#)
- 10 Oct 2016 - Exploring Compression Techniques for ROOT IO, Zhe Zhang [CHEP 2016](#)
- 20 Oct 2016 - Histogrammar: Making Histograms Functional, Jim Pivarski [Chicago Hadoop User's Group \(CHUG\)](#)
- 21 Sep 2016 - Reconstructing Particle Trajectories in High Energy Physics with Xeon and Xeon Phi, David Abdurachmanov, [Intel Xeon Phi User Group Meeting 2016](#)
- 16 Sep 2016 - Expressing Complex Data Aggregations With Histogrammar, [[talk abstract](#), [talk video](#)], Jim Pivarski, [Strange Loop Conference](#)
- 18 Aug 2016 - Computing for Data Analysis, Jim Pivarski, [LPC Physics Forum](#)
- 18 Jul 2016 - Potential Issues with Julia, Jim Pivarski, [DIANA Meeting](#)
- 20 Jun 2016 - High-level analysis scripts with low-level performance, Jim Pivarski, [DIANA Meeting](#)

DIANA Topical Meetings

(<https://indico.cern.ch/category/7192/>)

December 2016

 12 Dec [First Explorations](#) **NEW**

 05 Dec [Software Development, Sustainability and Evolution](#)

November 2016

 14 Nov [Toward Vectorized, Parallel Ntuple processing](#)

 07 Nov [Avro and Numpy for Spark and machine learning](#)

October 2016

 31 Oct [DIANA Meeting - Event generation and fitting on GPUs](#)

September 2016

 19 Sep [DIANA Meeting - Python/ROOT interoperability](#)

July 2016

 25 Jul [DIANA Meeting - Focus groups, surveys and building a roadmap](#)

 18 Jul [DIANA Meeting - The Julia Language](#)

June 2016

 20 Jun [DIANA Meeting - Analysis script language](#)

Defining Longer-term Strategy

- HL-LHC computing requires a major 'software upgrade'
- A Community White Paper (CWP) on the overall strategy and roadmap for software and computing has been proposed
 - Initiated as WLCG charge to the LHC experiments and HSF as a step towards the LHC experiment TDRs in advance of HL-LHC
 - The scope should not be restricted only to HL-LHC
 - Some early software components could be built, tested and used by experiments in LHC Run3
- Organised by the HEP Software Foundation (HSF)
- Paper to be delivered by Summer 2017
- It should play a role in discussing possible funding scenarios for a "software upgrade".

Community White Paper (CWP)

- The CWP should identify and prioritise the software research and development investments required:
 - to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
 - to enable new approaches to computing and software that could radically extend the physics reach of the detectors
 - to ensure the long term sustainability of the software through the lifetime of the HL-LHC
- We need to engage the HEP community in this process through a series of workshops
 - Initiated as an HL-LHC planning process
 - Aiming for a broader participation (LHC, neutrino program, Belle II, linear collider so far)

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 provenance 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.

Likely constraints to fund a “Software Upgrade”

It appears unlikely that significant increases in investments in software will be made by funding agencies purely from particle physics budgets and/or into individual experiments. Other opportunities do perhaps exist, but often imply constraints, for example:

- Investments into software impacting multiple experiments
- Investments into development with impact beyond particle physics
- Investments into development permitting use of computing facilities (e.g. HPC) planned for other non-HEP purposes
- Investments requiring collaborations with Computer Science or Industry

Building the LHC software in use today was possible without too many such constraints. The good news is that the community (with an existing LHC computing system) is better positioned today to make effective progress even with such constraints.

Status

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 ~2019-2020.

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.

The HSF has begun the process of organizing working groups, engaging HEP beyond the LHC experiments and planning for dedicated workshops. Sessions at existing meetings can also be used when possible.

Community Roadmap Process

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

- Initial presentation and organization began at the WLCG workshop and CHEP in October, 2016
 - We are currently developing the WG charges and attracting interested individuals to the WGs
- A “kick-off” workshop at UC San Diego on 23-26 Jan 2017
 - Start real work after a few months post-CHEP gestation in the WGs
 - Discussions on more controversial topics, find path to consensus
 - Develop plans and responsibilities for delivering white paper by summer 2017
- Possible “topical” workshops between Jan-Jun 2017, building on existing community activities when possible (e.g. DPHEP, Reco Algorithms Forum/CTD, IML)
- A final workshop in summer 2017 (in Europe, near CERN?)

HSF Community White Paper Workshop at SDSC

- The kick-off workshop for the HSF CWP process will be on 23-26 January 2017 at SDSC/UCSD
- <http://indico.cern.ch/event/570249/>
- People from many HEP experiments (LHC and beyond)
- A key element in the “community process” to form a consensus
- We are looking for opportunities to introduce new ideas into the HEP discussions
- An early strawman agenda is posted and is evolving quickly
- Registration is open (see URL above)

Possible routes to a “Software Upgrade”

- If we are aiming at a larger “software upgrade” project towards the HL-LHC, an additional ingredient is to find (or liberate/reallocate) the resources to realize this roadmap.
- We need both initial exploratory R&D and eventual development projects!
- In the US, both the NSF and the DOE have at least the notion of eventual resources and/or organization for new common projects in HEP (NSF: SI2, DOE: HEP CCE)
 - The US NSF has funded a “conceptualization” (planning) project with a possible path towards a “Software Institute” (<http://s2i2-hep.org>)
 - The US DOE has seeded the “Center for Computing Excellence” with some initial resources. (<http://hepcce.org>)
- We hope that a clear community roadmap will bring these and other partners together for an HL-LHC software upgrade.

Practicalities: HSF Google Groups

The following Google Groups are relevant:

- Group for discussion of Community White Paper
 - <https://groups.google.com/forum/#!forum/hsf-community-white-paper>
- General announcement group for community messages (low traffic)
 - <https://groups.google.com/forum/#!forum/hep-sw-comp>
- Community Discussion list
 - <https://groups.google.com/forum/#!forum/hep-sf-forum>
- Specific group for US NSF Software Institute Conceptualization
 - <https://groups.google.com/forum/#!forum/s2i2-hep>

Practicalities: Charges for CWP Working Groups

- Over the next weeks we will be formulating charges for the CWP working groups.
- Templates for drafting these charges are in google docs.
- The overall WLCG HSF charge and links to individual WG charge google docs (in preparation) can be found at:
 - <http://hepsoftwarefoundation.org/cwp.html>
- To view and/edit these charges you will need to be subscribed to:
 - <https://groups.google.com/forum/#!forum/hsf-community-white-paper>
- These google docs are just to draft the charges and allow people to self-organize and start discussions. Eventual proper documents (e.g. in latex) can switch elsewhere (e.g. github).

Practicalities: Possible Working Groups

Detector Simulation	full and fast simulations, hi-pileup environments
Triggering	algorithms, GPUs and/or FPGAs
Event Reconstruction	new approaches to event reconstruction
Visualization	tools for data analysis, education, and outreach
Data Access and Management	scaling to the exabyte level
Workflow and Resource Management	millions of jobs in heterogenous systems
Physics generators	better models, better precision, code optimisations
Data Analysis and Interpretation	efficient use of many-core, modern techniques
Data and Software Preservation	preservation and reuse of data and software
Software Development, Deployment and Validation/Verification	improved modularity and quality, contribution
Computing Models, Facilities, Distributed Computing	range of possible models, costing
Various Aspects of Technical Evolution (Software Tools, Hardware)	
Security and Access Control	
Careers, Staffing and Training	perhaps in a separate concurrent white paper
Machine Learning	
Conditions Database	
Event Processing Frameworks	

More details in links at <http://hepsoftwarefoundation.org/cwp.html>

Summary

- The LHC and wider HEP communities are executing a planning process for medium/long term software and computing R&D
- The process should produce the “Community White Paper”, a roadmap for HEP Software and Computing for the 2020s
- In parallel we would like to investigate funding and collaboration possibilities for eventual projects and activities
- Please consider participating in the CWP process, both the WGs and the workshops starting from the HSF workshop SDSC/UCSD in January 2017