HEP Software Foundation

HEP Software Foundation Community White Paper

Graeme Stewart, CERN EP-SFT

3rd Scientific Computing Forum, 2017-10-27

Challenges for the next decade

- High Luminosity LHC brings a huge challenge to software and computing
    - Both rate and complexity rise

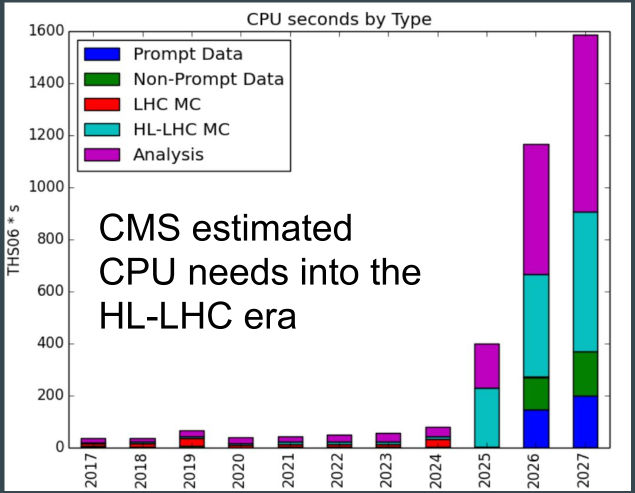

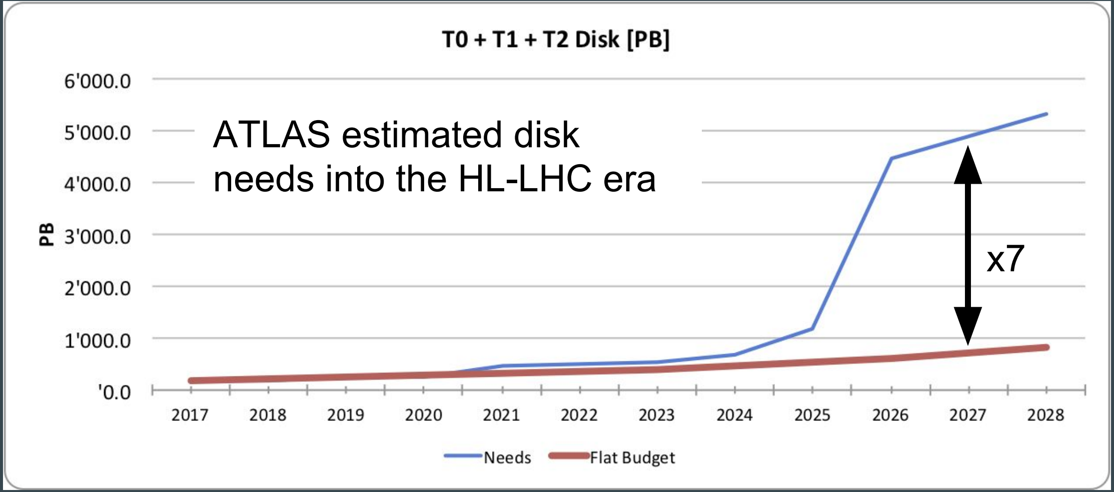

- Not just a simple extrapolation of Run 2
- Resources needed will exceed those from technology evolution alone

The evolving technology landscape

- Single core CPU throughput stalled
- Many/multi core systems are the norm
  - Serial or multi-process processing is under severe memory pressure
- Co-processors now commonplace
  - GPGPUs, FPGAs - greater throughput, far more challenging programming model
- Wide vector registers (up to 512 bit)
- Power a dominant factor
- Storage capacity climbing
  - 100TB disks possible by HL-LHC, but little I/O improvement expected
- Network capacity keeps growing

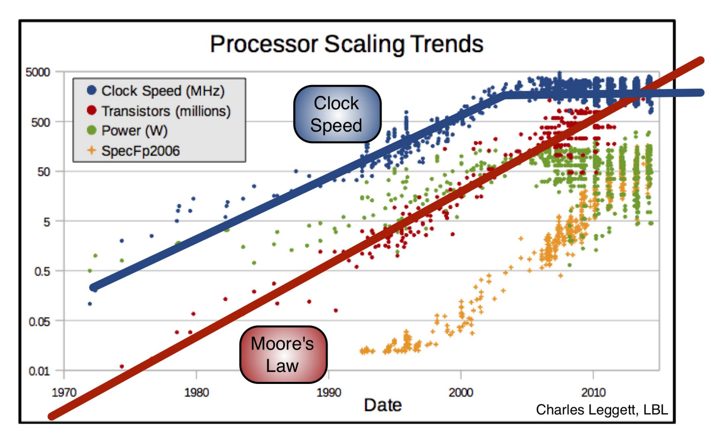

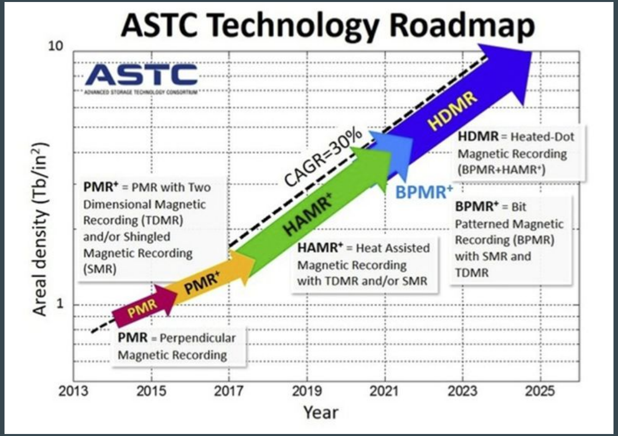

Community white paper inception

- From Spring 2016 discussions, idea started to crystallise at the May 2016 HSF Meeting at LAL
    - *describe a global vision for software and computing for the HL-LHC era and HEP in the 2020s*
- Formal charge from the WLCG in July 2016
    - Anticipate a "software upgrade" in preparation for HL-LHC
    - Identify and prioritize the software research and development investments
        1. to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
        2. to enable new approaches to computing and software that could radically extend the physics reach of the detectors
        3. to ensure the long term sustainability of the software through the lifetime of the HL-LHC

Starting the process

- Started to organise into different working groups at the end of 2016
- Kick-off workshop 23-26 January 2017, San Diego
  - 110 participants, mainly US + CERN
  - 2.5 days of topical working group meetings
  - Extensive notebooks of initial discussions
- Groups held workshops and meetings in the subsequent months
  - Broadening the range of participation
  - Some invited non-HEP experts to participate

Concluding the process

- Workshop in Annecy 26-30 June started to draw the process to a close
    - 90 Participants: 48 US, 42 Europe (of which 20 from CERN)
- 13 working groups presented their status and plans
- Substantial progress on many Community White Paper chapters
    - WGs used the workshop to make further progress on writing
- There was a fair amount of optimism about writing the final Roadmap in August

Editorial board and roadmap document draft

- There was not a lot progress over the summer months
- Set more realistic goals
  - Individual WG chapters by end of September
  - Overall roadmap paper by end of October
- 10 working group chapters available for community review*
  - With a few more in late stages of preparation
- Editorial Board was set up, with the aim of encompassing the breadth of our community
- *First draft* of the text was prepared by a small team within the Editorial Board
  - Released 20 October

- Predrag Buncic (CERN) - Alice contact
- Simone Campana (CERN) - ATLAS contact
- Peter Elmer (Princeton)
- John Harvey (CERN)
- Frank Gaede (DESY) - Linear Collider contact
- Maria Girone (CERN Openlab)
- Roger Jones (Univ. of Lancaster) - UK contact
- Michel Jouvin (LAL Orsay)
- Rob Kutschke (FNAL) - FNAL experiments contact
- Dario Menasce (INFN-Milano) - INFN contact
- Mark Neubauer (U.Illinois Urbana-Champaign)
- Stefan Roiser (CERN) - LHCb contact
- Liz Sexton-Kennedy (FNAL) - CMS contact
- Mike Sokoloff (U.Cincinnati)
- Graeme Stewart (CERN, HSF)
- Jean-Roch Vlimant (Caltech)

*Will place final versions on arXiv

First synthesis draft

- 60 page document
- 12 chapters summarising R&D in a variety of areas for HEP Software and Computing
- **Almost all major pieces of HEP Software and Computing are covered**
    - Event Generators should conclude soon; Workload Management might
- Now being reviewed by the CWP participants and the Editorial Board
- Feedback so far is positive, so we anticipate a second draft fairly soon

A Roadmap for HEP Software and Computing R&D for the 2020s

The HEP Software Foundation

Table of Contents

Detector simulation

- Simulating our detectors consumes huge resources today
  - Remains a vital area for HL-LHC and DUNE
- Main R&D topics we need
  - Improved physics models for higher precision at higher energies (HL-LHC and then FCC)
    - Hadronic physics in LAr TPCs needs to be redeveloped
  - Adapting to new computing architectures
    - Can a vectorised transport engine be demonstrated to work in a realistic prototype?
  - Fast simulation - develop a common toolkit for tuning and validation
    - Can we use Machine Learning profitably here?
  - Geometry modeling
    - Easier modelling of complex detectors, targeting new computing architectures

Software trigger and event reconstruction

- Move to software triggers is already a key part of the program for LHCb and ALICE already in Run 3
  - So called ‘real time analysis’ increases signal rates and can make computing much more efficient (storage and CPU)
- Main R&D topics we need
  - Controlling charged particle tracking resource consumption and maintaining performance
    - Do current algorithms’ physics output hold up at pile-up of 200 (or 1000)
    - Can tracking maintain low p_T sensitivity within budget?
  - Improved use of new computing architectures
    - Multi-threaded and vectorised CPU code
    - Extending use of GPGPUs and possibly FPGAs
  - Robust validation techniques when information will be discarded
    - Using modern continuous integration, tackling multiple architectures with reasonable turnaround times

Data analysis and interpretation

- Today we are dominated by many cycles of data reduction
  - Aim is to reduce the input to an analysis down to a manageable quantity that can be cycled over quickly on ~laptop scale resources
  - Key metric is ‘time to insight’
- Main R&D topics we need
  - How to use the latest techniques in data analysis that come from outside HEP?
    - Particularly from the Machine Learning and Data Science domains
    - Need ways to seamlessly interoperate between their data formats and ROOT
      - Python is emerging as the *linga franca* here, thus guaranteeing PyROOT is critical
  - New Analysis Facilities
    - Skimming/slimming cycles consume large resources and can be inefficient
    - Can interactive data analysis clusters be set up?
  - Data and analysis preservation is important

Data management and organisation

- Data storage costs are a major driver for LHC physics today
  - HL-LHC will bring a step change in the quantity of data being acquired by ATLAS and CMS
- Main R&D topics we need
  - Adapt to new needs driven by changing algorithms and data processing needs, e.g,
    - The need for fast access to training datasets for Machine Learning
    - Supporting high granularity access to event data
      - Needed to effectively exploit backfill or opportunistic resources
    - Rapid high throughput access for a future analysis facility
    - Processing sites with small amounts of cache storage
  - Do this profiting from the advances in industry standards and implementations, such as Apache Spark-like clusters (area of continued rapid evolution)
  - Consolidate storage access interfaces and protocols
  - Support efficient hierarchical access to data, from high latency tape and medium latency network

Facilities and distributed computing

- Storage and compute today is provided overwhelmingly from WLCG resources
  - Expected to continue for HL-LHC, but to be strongly influenced by developments in commodity infrastructure as a service (IaaS, commercially this is usually Cloud Computing)
- Main R&D topics we need
  - Understand far better the effective costs involved in delivering computing for HEP
    - This needs to be sensitive to regional variations in funding and direct and indirect costs
      - E.g., smaller sites frequently contribute 'beyond the pledge' resources, power costs and human resources
    - Full model is infeasible, but providing a reasonable gradient analysis for future investment should be possible
      - Should we invest in better network connectivity or in more storage?
  - How to take better advantage of new network and storage technologies (software defined networks, object stores or content addressable networks)
  - Strengthen links to other big data sciences (SKA) and computing science; how to share network resources

Data processing frameworks

- Experiment software frameworks provide the scaffolding for algorithmic code
  - Currently there are many implementations of frameworks, with some sharing between experiments (e.g., ATLAS and LHCb share Gaudi, Intensity Frontier experiments use art)
  - All of these frameworks are evolving to support concurrency
- Main R&D topics of interest
  - Adaption to new hardware, optimising efficiency and throughput
    - We need the best libraries for this and these will change over time
  - Incorporation of external (co)processing resources, such as GPGPUs
  - Interface with workload management system to deal with the inhomogeneity of processing resources
    - From volunteer computing to HPC job slots with 1000s of nodes
  - Which components can actually be shared and how is that evolution achieved?

Machine learning

- Neural networks and Boosted Decision Trees have been used in HEP for a long time
  - E.g., particle identification algorithms
- More recently the field has been significantly enhanced by new techniques (Deep Neural Networks) and enhanced training methods
  - Very good at dealing with noisy data and huge parameter spaces
  - A lot of interest from our community in these new techniques, in multiple fields
- Main R&D topics of interest
  - Speeding up computationally intensive pieces of our workflows (fast simulation, tracking)
  - Enhancing physics reach by classifying better than our current techniques
  - Improving data compression by learning and retaining only salient features
  - Anomaly detection for detector and computing operations
- However, we do still expect that significant efforts will be required to make effective use of these techniques
  - Good links with the broader Machine Learning and Data Science communities required

Other technical areas of work

Conditions Data

- Growth of alignment and calibration data is usually linear in time
  - Per se, this does not represent a major problem for the HL-LHC
- Opportunities to use modern distributed techniques to solve this problem efficiently and scalably
  - Cacheable blobs accessed via REST
  - CVMFS + Files
  - Git

Visualisation

- Many software products developed for event visualisation
  - Part of the framework, with full access to event and geometry data
  - Standalone as a lightweight solution
- New technologies for rendering displays exist, e.g., WebGL from within a browser

- These areas are both examples of where we can refocus current effort towards common software solutions
- This should improve quality, economise overall effort and help us to adapt to new circumstances

Data, software and analysis preservation

- We seem to be doing well compared to other fields
- Challenge is both to physically preserve bits and to preserve knowledge
  - DPHEP has looked into both
- Knowledge preservation is very challenging
  - Experiment production workflows vary in significant details
  - Variety of different steps are undertaken at the analysis stage, even within experiments
- Need a workflow that can capture this complexity
  - Technology developments that can help are, e.g., containers
- CERN Analysis Preservation Portal forms a good basis for further work
  - Needs to have a low barrier for entry for analysts
  - Can provide an immediate benefit in knowledge transmission within an experiment

Software development, training and careers

- Experiments have modernised their software development models a lot recently
  - Moving to git and CMake as standard components
  - Using social coding sites (gitlab, github) coupled to Continuous Integration
- Additional tools would benefit the community
  - Static analysis of code, refactoring code, performance measures
- Using new tools requires investing in training for the community
  - The more commonality in the tools and techniques, the more training we can share
  - This provides preservation and propagation of knowledge
- Our environment is becoming more complex; we require input from physicists *whose concerns are not primarily in software*
  - Sustainability of these contributions is extremely important
- Recognition of the contribution of our specialists in their careers is extremely important

Community white paper - moving forwards

- Community White Paper process has been a success
  - Engaged more than 250 people and produced more than 300 pages of detailed description in many areas
- Summary roadmap lays out a path forward and identifies the main areas we need to invest in for the future
  - Supporting the HL-LHC Computing TDRs and NSF S2I2 strategic plan
- Current first draft will undergo a process of refinement and conclude in a few months
- HEP Software Foundation has proved its worth in delivering this CWP Roadmap
  - Achieving a *useful* community consensus is not an easy process
- We now need to marshal the R&D efforts in the community, refocusing our current effort and helping to attract new investment in critical areas
  - The challenges are formidable, working together will be the most efficacious way to succeed
  - HSF will play a vital role in spreading knowledge of new initiatives, encouraging collaboration and monitoring progress
  - Workshops planned for next year (with WLCG) and at sessions before CHEP