HEP Software Foundation

Graeme A Stewart, CERN EP-SFT

EIC Software Meeting, Trieste May 2019

HL-LHC and the Intensity Frontier

Our mission:

- Exploit the Higgs for SM and BSM physics
- b, c, tau physics to study BSM and matter/anti-matter
- Dark matter
- Neutrino oscillations and mass
- QGP in heavy ion collisions
- Explore the unknown

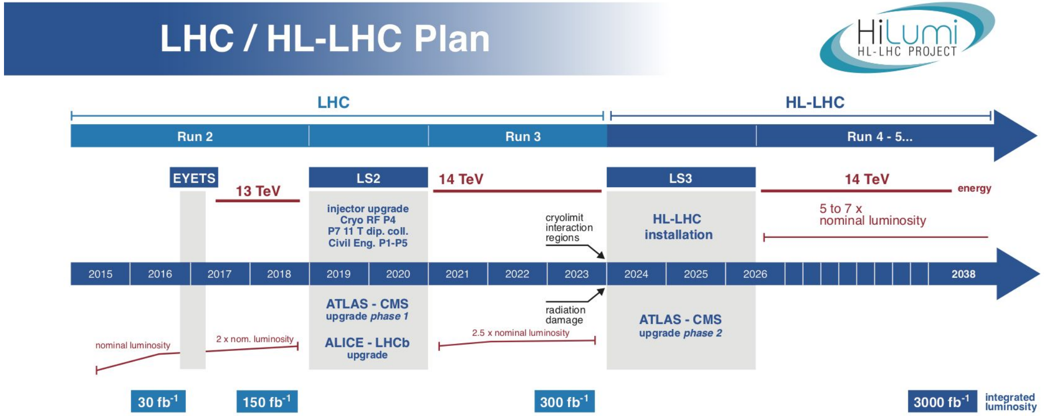

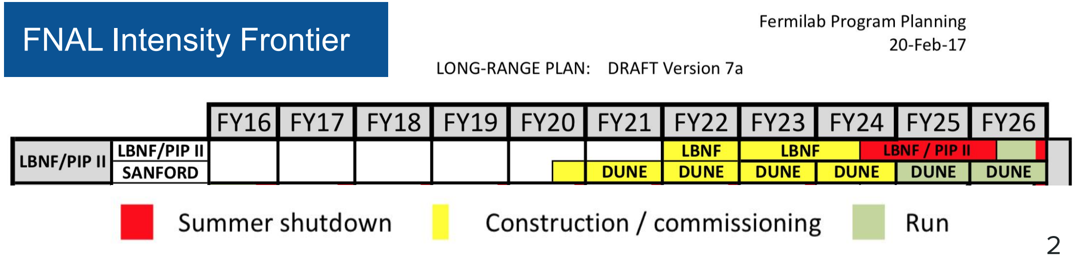

HEP Software and Computing

- High Energy Physics has a vast investment in software
  - Estimated to be around 50M lines of C++
  - Which would cost more than 500M$ to develop commercially
- It is a critical part of our physics production pipeline, from triggering all the way to analysis and final plots as well as simulation
- LHC experiments use about 1M CPU cores every hour of every day, we have around 1000PB of data with 100PB of data transfers per year (10-100Gb links)
  - We are in the exabyte era already
- This is a huge and ongoing cost in hardware and human effort
- With significant challenges ahead of us to support our ongoing physics programme

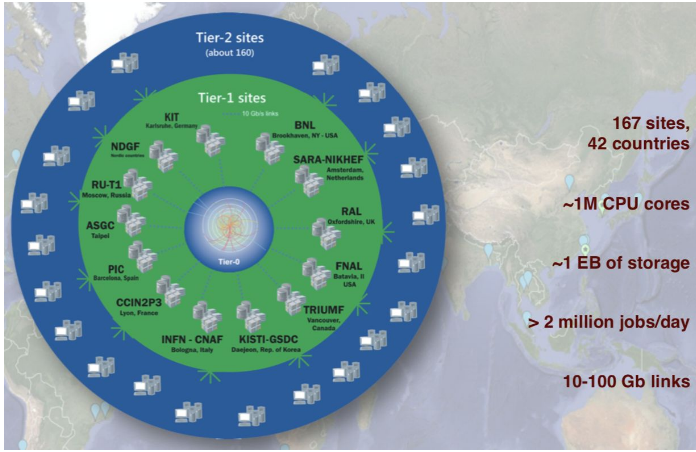

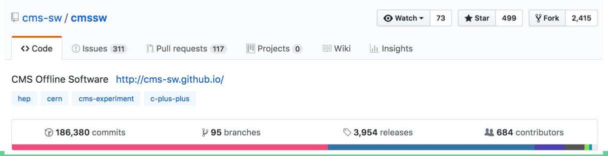

Challenges for the Next Decade

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

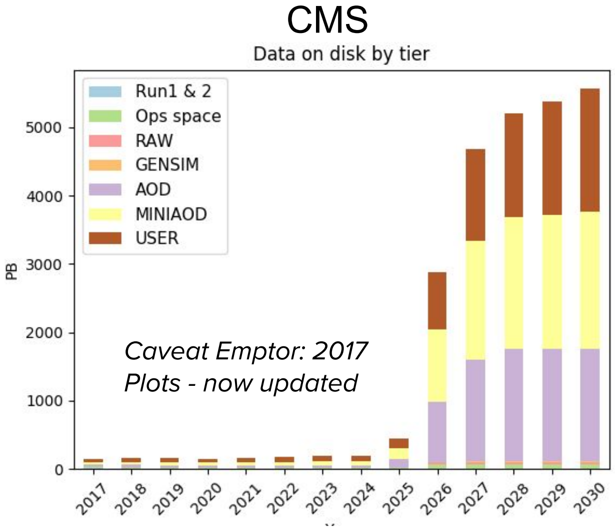

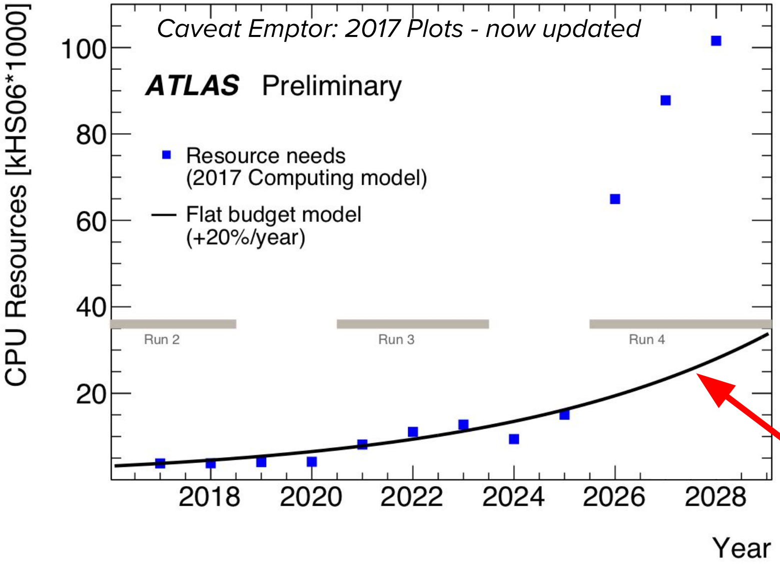

- Not just a simple extrapolation of Run 2 software and computing
  - Resources needed would hugely exceed those from technology evolution alone

This is even probably too optimistic, ~5-10%?

Technology Evolution

- Moore's Law continues to deliver increases in transistor density
  - But, doubling time is lengthening
- Clock speed scaling failed around 2006
  - No longer possible to ramp the clock speed as process size shrinks
  - Leak currents become important source of power consumption
- So we are basically stuck at ~3GHz clocks from the underlying Wm^{-2} limit
  - This is the *Power Wall*
  - Limits the capabilities of serial processing
- Memory access times are now ~100s of clock cycles

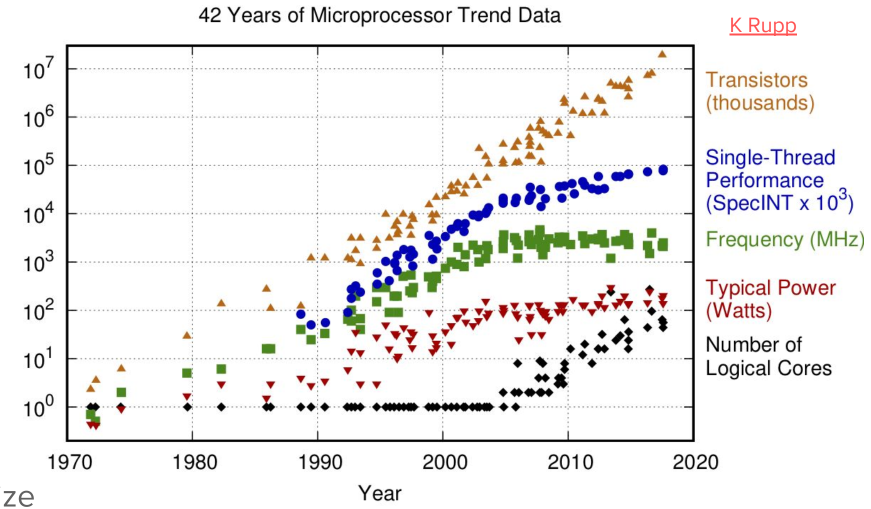

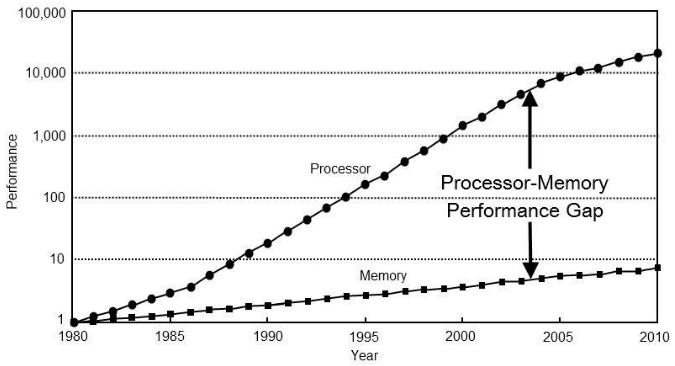

Decreasing Returns over Time

- Conclusion is that diversity of new architectures will only grow
- Best known example is of GPUs

[link]

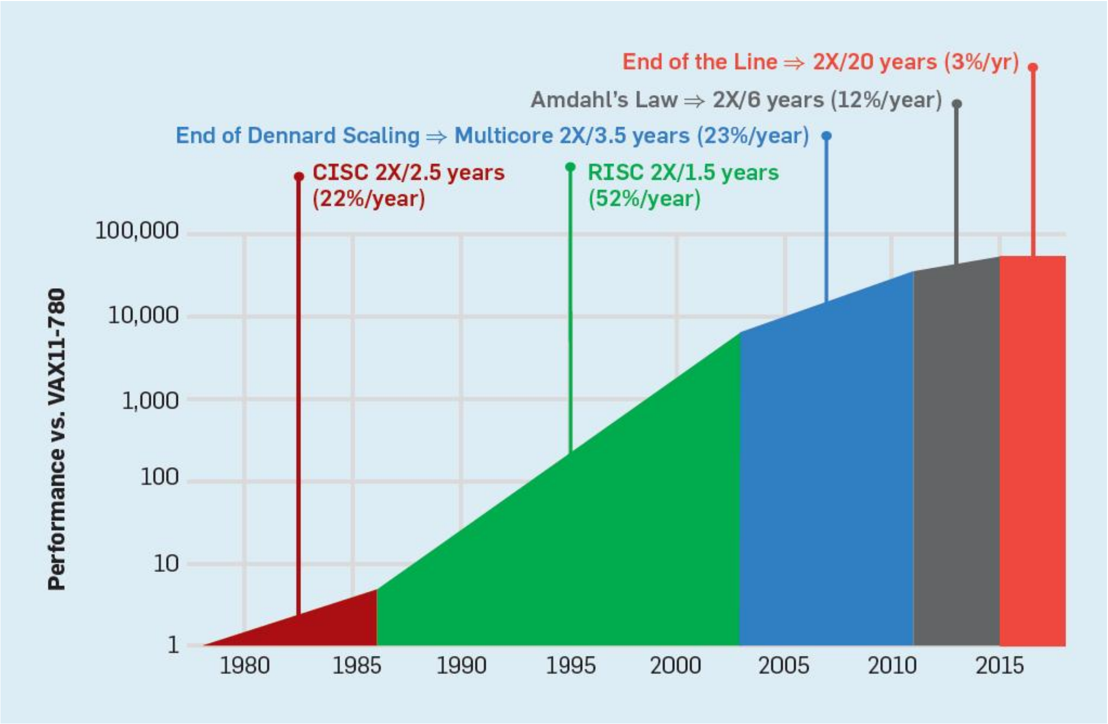

GPUs dedicate far more transistors to arithmetic

Drivers of Technology Evolution

- Low power devices
  - Driven by mobile technology and Internet of Things
- Data centre processing
  - Extremely large clusters running fairly specialist applications
- Machine learning
  - New silicon devices specialised for training machine learning algorithms, particularly low precision calculations
- Exascale computing
  - Not in itself general purpose, but poses many technical problems whose solutions can be general - HEP pushed to use HPC centres, especially in US
- Energy efficiency is a driver for all of these developments
  - Specialist processors would be designed for very specific tasks
  - Chips would be unable to power all transistors at once: dark silicon is unlit when not used

Hardware Evolution in a Nutshell

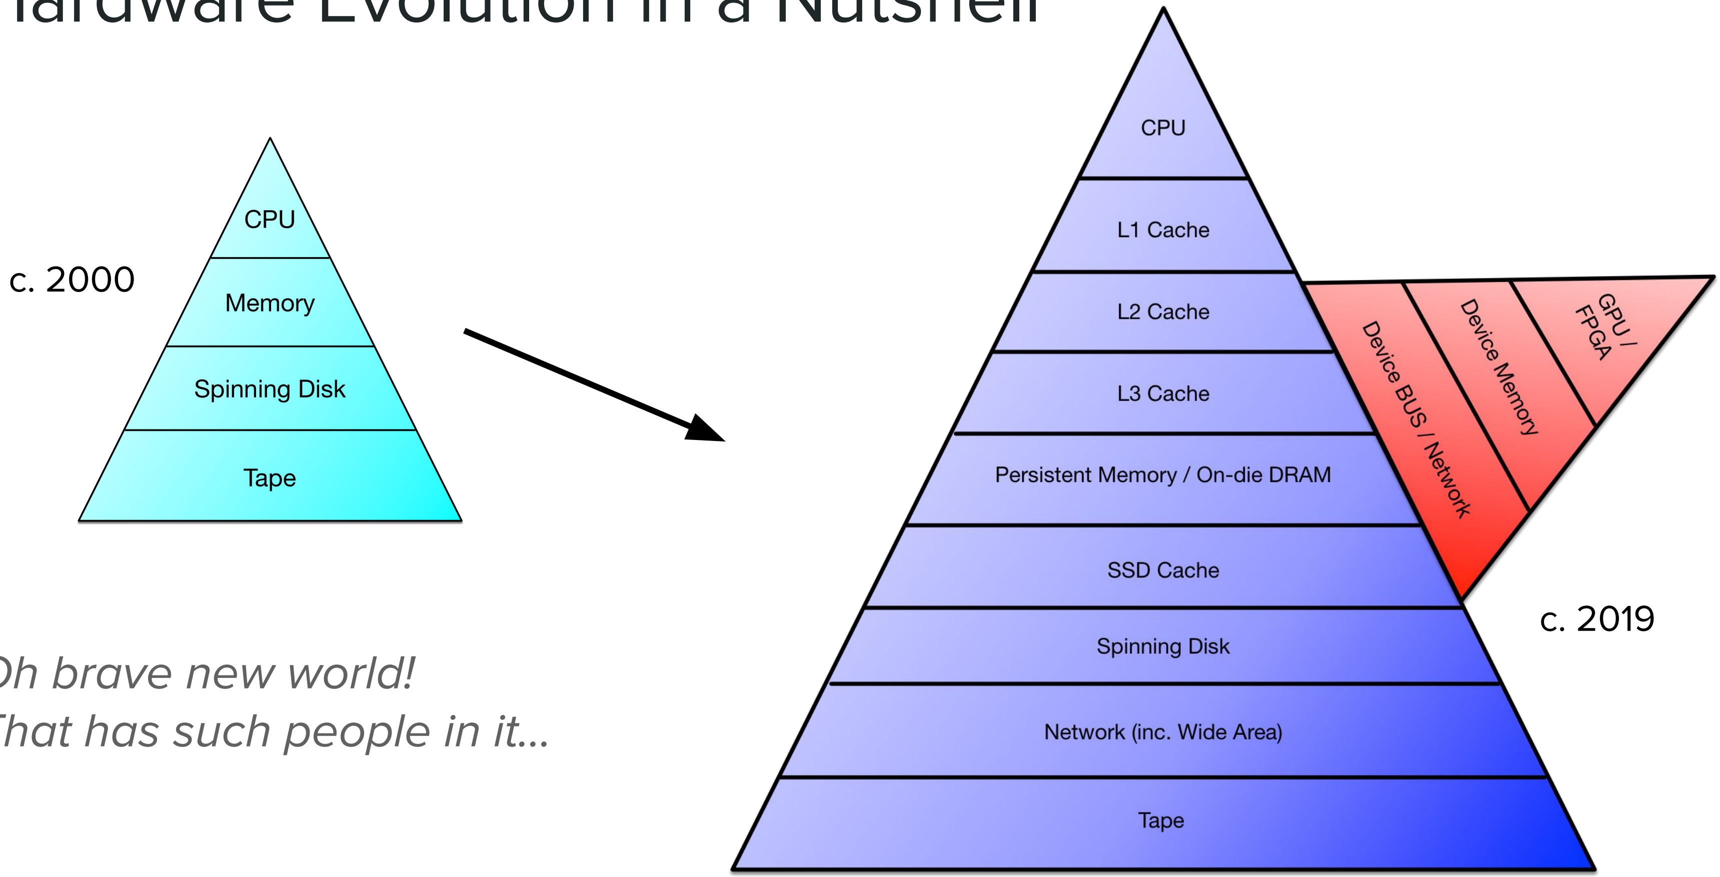

Oh brave new world!
That has such people in it…

Software Challenges for HL-LHC

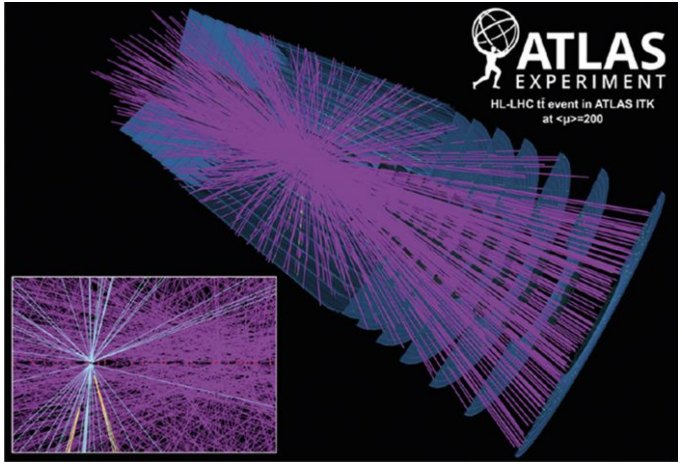

- Pile-up of ~200 ⇒ particularly a challenge for charged particle reconstruction
- With a flat budget, Moore's lawish improvements are the real maximum we can expect on the HW side
- HEP software typically executes one instruction at a time (per thread)
    - Since ~2013 CPU (core) performance increase is due to more internal parallelism
    - x10 with the same HW only achievable if using the full potential of processors
        - Major SW re-engineering required (but rewriting *everything* is not an option)
    - Co-processors like GPUs are of little use until the problem has been solved
- Increased amount of data requires to revise/evolve our computing and data management approaches
    - We must be able to feed our applications with data efficiently
- HL-LHC salvation will come from software improvements, not from hardware

HEP Software Foundation (HSF)

- The LHC experiments, Belle II and DUNE face the same challenges
  - HEP software must evolve to meet these challenges
  - Need to exploit all the expertise available, inside and outside our community, for parallelisation
  - New approaches needed to overcome limitations in today’s code
- Cannot afford any more duplicated efforts
  - Each experiment has its own solution for almost everything (framework, reconstruction algorithms, …)
- The goal of the HSF is to facilitate coordination and common efforts in software and computing across HEP in general
  - *Our philosophy is bottom up, a.k.a. do-ocracy*

Community White Paper Inception

- We wanted to 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
- Long process of 1 year, with many working groups and 2 major workshops

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

- 70 page document
- 13 sections summarising R&D in a variety of technical areas for HEP Software and Computing
  - Almost all major domains of HEP Software and Computing are covered
- 1 section on Training and Careers
- 310 authors from 124 institutions
- https://doi.org/10.1007/s41781-018-0018-8; arXiv:1712.06982

HSF-CWP-2017-01
December 15, 2017

Contents

HSF Working Groups

- The Roadmap established what challenges the community faced
  - But it did not spell out *how* to face them in detail
- HSF had adopted a model of working groups from its earliest days
  - These were open groups of people in the community, motivated enough to organise around a common topic, usually at their own initiative
- This model was a good one for moving forwards on the key topics - simulation, reconstruction and analysis
  - We were a little more formal this time around
    - Call for nominations from the whole community, then search committee
    - Significant engagement from LHC experiments and beyond, e.g. Belle II
- The HSF's role here is one of an information conduit and meeting point
  - Report on interesting and common work being done
  - Forum for technical comments and discussion
  - Encourage cooperation across experiments and regions

Important Practical Matters

- Copyright and Licensing
  - Long neglected inside collaborations
  - Essential to be able to
    - Open source our software
    - Combine with other open source projects
- Copyright
  - Keep as low a number as practicable
  - E.g. © CERN for the benefit of collaboration X
- License
  - Favour liberal licenses for industry collaboration:- LGPL, Apache, MIT
  - Avoid GPL for libraries you want other people to use
- Software Tools
  - Active group promoting best practice for correctness and performance
- Packaging
  - We don't build our experiment software in isolation
  - Need a software stack, incorporating many components from the open source world and HEP community
  - Preference for tools that are not home grown and have a wider support base
  - Spack (LBNL) and Conda actively being prototyped

[HSF-TN-2016-01; HSF-TN-2016-03]

Event Generators

- Event generators are the start of the simulation chain
    - At the LHC Run1 only leading order generators were used
    - Negligable CPU consumption compared with detector simulation - no pressure to optimise
- However, with LHC upgrades coming higher order generators become much more important
    - These are inherently much more costly to run
    - Problems of negative weights can increase hugely the samples needed for weighted event samples
- In addition, the theory community, who develop these codes usually work in small teams
    - Recognition for technical improvements is limited/missing

ATLAS 2018 CPU Report

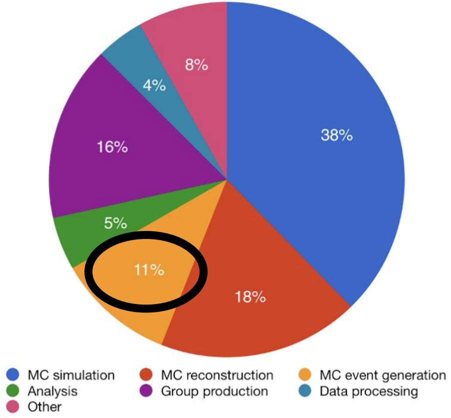

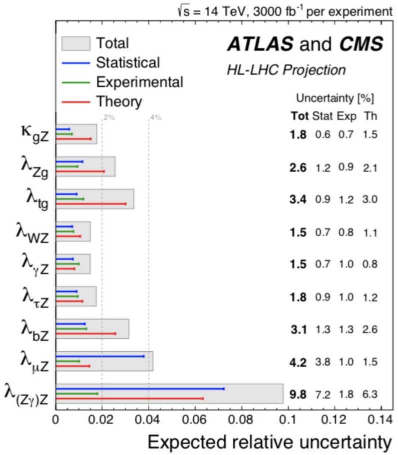

Many electroweak measurement errors dominated by theory (red). B. Hinemann

Event Generators - Technical Improvements

- HSF/LPCC workshop in November brought theory and experiment together to look at computing challenges of event generation
  - This was the first workshop of its kind
- Working group tackling technical challenges
  - Setting a baseline for further comparisons
  - Understanding how to run generators for best efficiency
  - Support for technical improvements (e.g. thread safety)
  - Porting to other architectures
    - Could be very suitable code to do this with (smaller, self contained code bases, numerically intensive)
    - e.g. building on the work done so far in MadGraph with GPUs

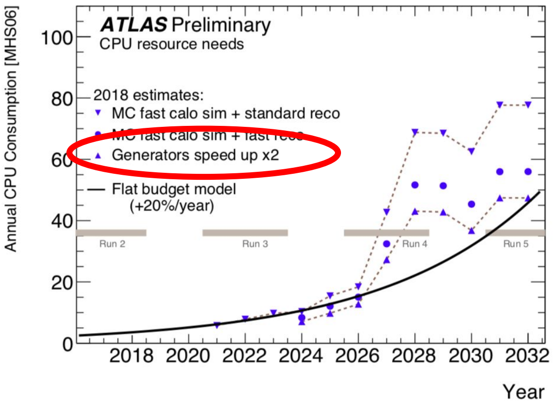

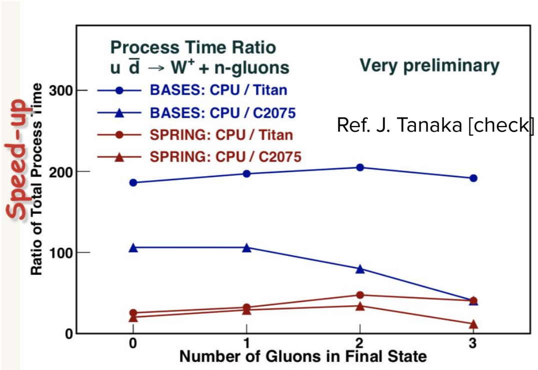

Ref. J. Tanaka [check]

Detector Simulation

- A major consumer of LHC grid resources today
  - Experiments with higher data rates will need to more simulat
- Faster simulation, with no or minimal loss of accuracy, is the goal
  - Range of techniques have been used for a long time (frozen showers, paramtric response)
  - Key point is deciding when it’s good enough for physics
- Machine learning lends itself to problems like this
  - Calorimeter simulations usually targeted
  - Variational Auto Encoders (VAEs) attempt to compress the data down to a ‘latent space’ - can be randomly sampled to generate new events
  - Generative Adverserial Networks (GANs) train two networks, one to generate events, the other to try to classify as real/fake
  - R&D on lifecycle integration into Geant4 is starting

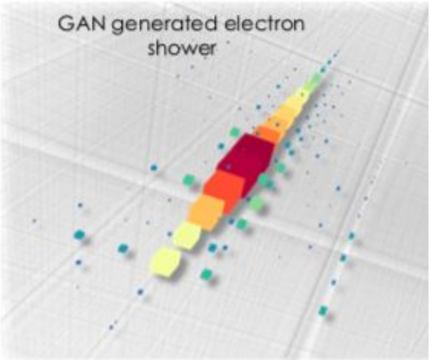

Use of Generative Adversarial Networks to simulate calorimeter showers, trained on G4 events (S. Vallacorsa)

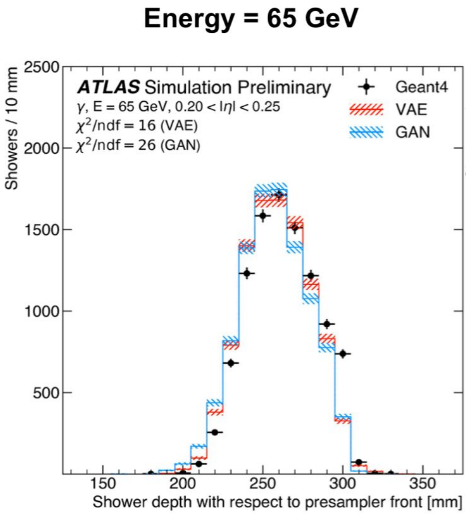

ATLAS VAE and GAN cf. Geant4 simulation [ATL-SOFT-PUB-2018-001.]

Detector Simulation

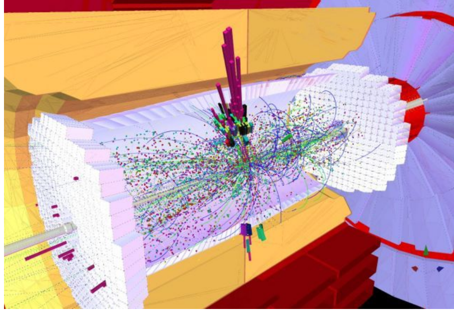

- Technical improvement programme helps (and helps *everyone*)
- GeantV R&D modernises code and introduces vectorisation
  - Speed-ups observed
  - Vectorisation introduces some gains
  - Code modernisation seems to help a lot
    - Reduce complexity and layers of object orientation
- Geant4 now have a new R&D working group that will take studies forward
- Some studies of running Geant4 on GPUs have begun
  - US Exascale Computing Project is funding this
  - Motivated by the next generation of US supercomputers that target exaflop
    - 90% of FLOP capacity in GPUs

DD4hep, DDG4, DDcond

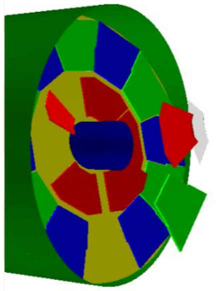

- Detector description toolkit aimed at the full lifecycle of an experiment
    - Conceptualisation, Construction, Operations
- Geometry consisting of a ‘tree’ of detector elements, defined in a single place
    - Simulation (GDML)
    - Reconstruction geometry
    - Analysis extensions

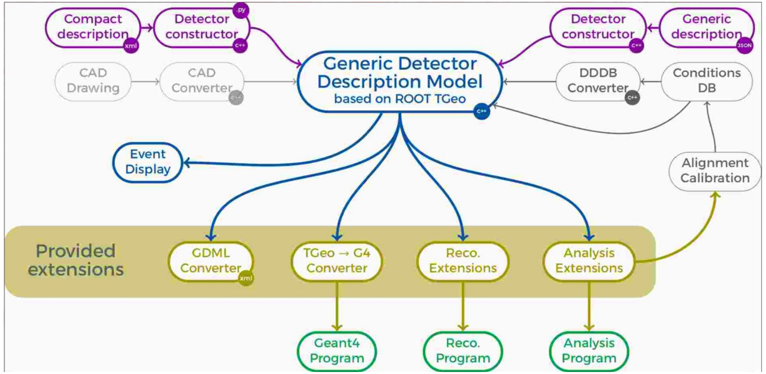

- DDG4
    - Hooks for user actions to generate detector response
- DDcond
    - Shifts detector elements from ideal position
    - Supports IoVs efficiently without locking (allows multi-threaded reconstruction across IoVs)
- Used by ILC, CLIC, CMS, LHCb, FCC, CALICE, SCTF

This project has received funding from the European Union’s Horizon 2020 Research and Innovation programme under Grant Agreement no. 654168.

AIDA 2020

Reconstruction and Software Triggers

- Hardware triggers no longer sufficient for modern experiments
  - More and more initial reconstruction needs to happen in software
- Close to the machine, need to deal with tremendous rates and get sufficient discrimination
  - Pressure to break with legacy code is high
  - Lots of experimentation with rewriting code for GPUs
  - ALICE have ported a lot of reconstruction to GPUs and also improved the algorithms a lot
  - CMS Patatrack project has improved physics performance as well
    - Revisiting old code helps!
- Lessons learned keep data model simple, bulk data, be asynchonous, minimise data transfers

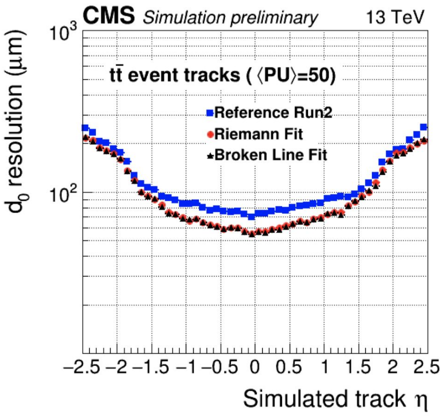

(a) d_0 resolution vs η

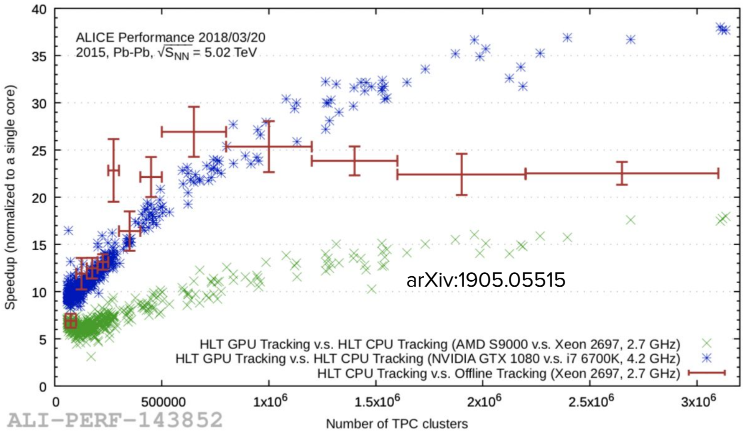

arXiv:1905.05515

Reconstruction and Software Triggers

- Real Time Analysis (HEP Version)
  - Design a system that can produce analysis useful outputs as part of the trigger decision
    - If this captures the most useful information from the event, can dispense with raw information
  - *This is a way to fit more physics into the budget*
- LHCb Turbo Stream has been introduced in Run2 and will be dominant in Run3
- Whole ALICE data reduction scheme is based around keeping 'useful' parts of events (no more binary trigger)
- ATLAS and CMS have schemes under development for special handling of samples for which full raw data is unaffordable

Persistence method	Average event size (kB)
Turbo	7
Selective persistence	16
Complete persistence	48
Raw event	69

LHCb Run2 Turbo took 25% of events for only 10% of bandwidth

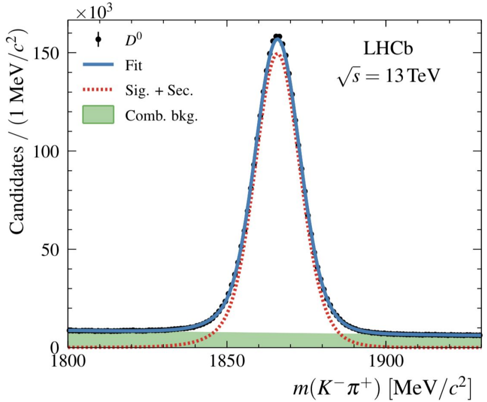

LHCb charm physics analysis using Turbo Stream (arXiv:1510.01707)

Analysis

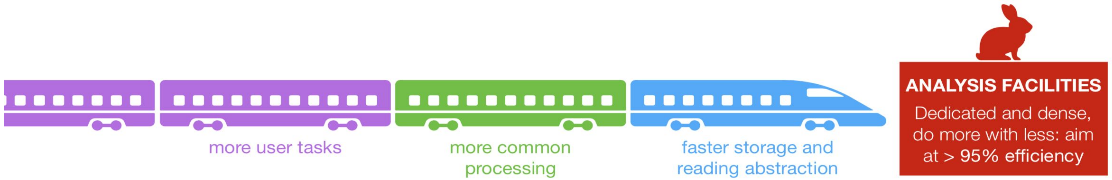

- Scaling for analysis level data also a huge challenge for all LHC experiments
- Efficient use of analysis data can come with combining many analyses as carriages in a train like model (pioneered by PHENIX then ALICE)
  - Also goes well with techniques like tape carousels (ATLAS scheme for rotating primary AOD data from tape systems into a disk buffer)
  - Interest in *analysis clusters*, specialised for analysis operations over the generic grid resources (WLCG/HSF pre-CHEP workshop 2-3 November)
- Reducing volume of data needed helps hugely
  - CMS ~1kB nanoAOD makes a vast difference to analysis efficiency and “papers per petabyte”
  - Smaller EDM is easier to make efficient
  - Requires analyst agreement on corrections, scale factors, etc.
    - However the alternative is perhaps that your analysis never gets done

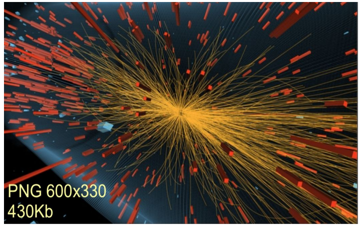

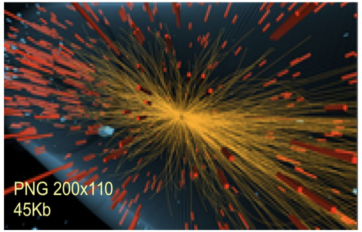

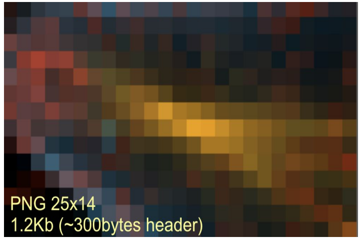

Analysis

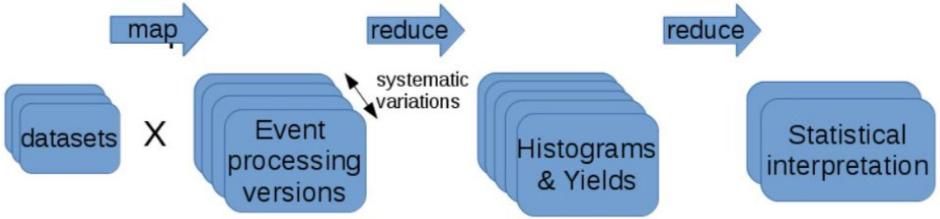

- Improve analysis ergonomics - how the user interacts with the system to express their analysis
  - Streamline common tasks
    - Handle all input datasets; Corrections and systematics
    - Compute per event and accumulate; Statistical interpretations
  - Declarative models, building on ROOT’s RDataFrame
    - Say *what*, not *how* and let the backend optimise
    - E.g. split and merge, GPU execution
- Notebook like interfaces gain ground, as do containers - lots of high level Python
- Interest in data science tools and machine learning is significant for this community - inspiring new approaches (e.g. uproot, awkward array, Coffea, scikit-hep)
  - This is an ecosystem into which HEP can contribute
  - PyHEP Working Group coordinates activities in this area

Many analysis frameworks, multiple per experiment, not well generalised

```
# * Jet select/cleaning against loose leptons , jet pt > 25 , jet id
flow.DefaultConfig(jetPtCut=25,jetIdCut=0,jetPUIdCut=0)
flow.SubCollection("CleanJet","Jet",'''
 Jet_pt > jetPtCut &&
 Jet_jetId > jetIdCut &&
 Jet_puId > jetPUIdCut &&
 (Jet_LeptonIdx==-1 || Jet_LeptonDr > 0.3)
''')
```

A. Rizzi, NAIL prototype

Frameworks and Integration

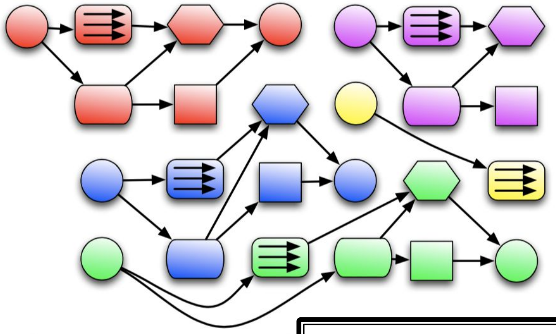

Cartoon of a single job, processing multiple events (colours) through different modules (shapes)

- Increasingly heterogeneous world requires advanced software support infrastructure
  - Software frameworks support use of different devices as well as insulate developers from many of the details of concurrency and threading models
    - Adapt to the new heterogeneous landscape
    - Latency hiding is critical to maintaining throughout
  - Framework development has traditionally been quite fragmented, but new experiments should offer a chance to increase convergence
    - Better to start off together than try to re-converge later (iLCSoft, LArSoft examples of success, albeit without concurrency; Gaudi for LHCb, ATLAS)
    - ALFA for ALICE and FAIR experiments
- New HSF working group being established now (draft mandate)

Training and Careers

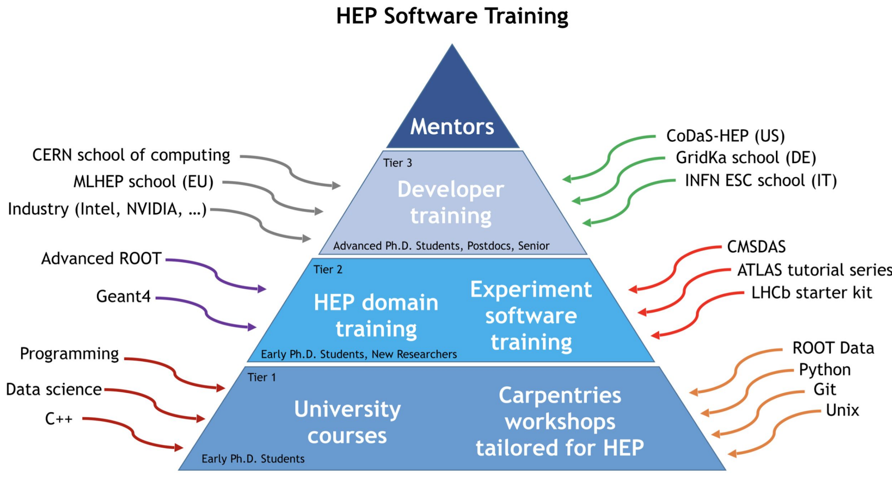

- Many new skills are needed for today’s software developers and users
- Base has relatively uniform demands
  - Any common components help us
- LHCb StarterKit initiative taken up by several experiments, sharing training material
  - Links to ‘Carpentries’ being remade (US training projects)
- New areas of challenge
  - Concurrency, accelerators, data science
  - Need to foster new C++ expertise (unlikely to be replaced soon as our core language, but needs to be modernised)
- Careers area for HEP software experts is an area of great concern
  - Need a functioning career path that retains skills and rewards passing them on
  - Recognition that software is a key part of HEP now

Instrumentation and Computing

Xinchou Lou, Brigitte Vachon
Scientific Secretaries: Emilia Leogrande, Rogers Jones

Meeting the HL-LHC Challenge

- Already since the Roadmap was written experiments have made great progress in meeting the HL-LHC challenge
  - Bad software, is extremely expensive
  - Good and clever software allows much more physics to fit in the budget

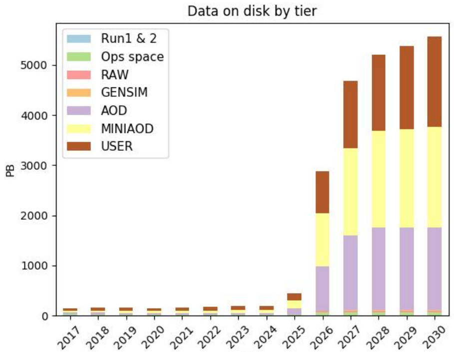

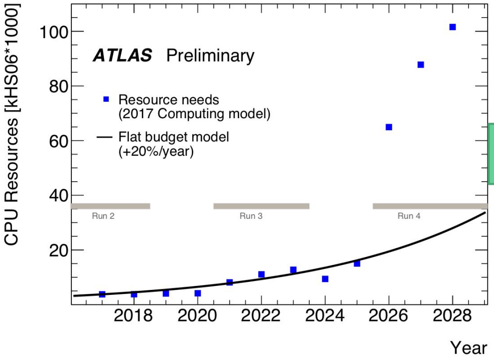

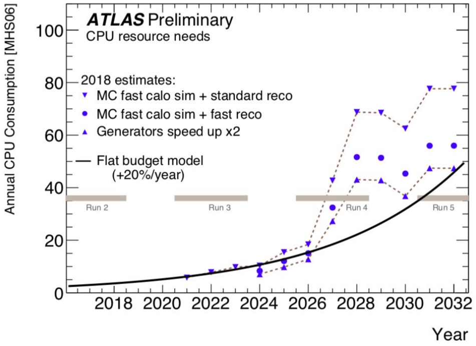

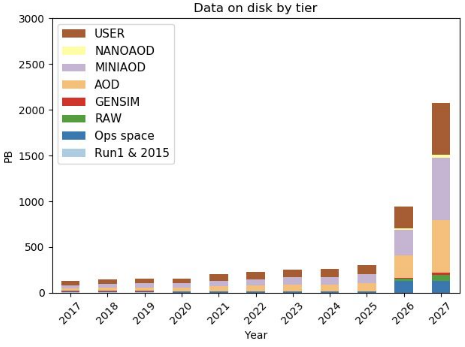

Outlook

- We have a wide ranging and ambitious physics programme in HEP and in associated disciplines
- Our experiments are highly data intensive now and require high quality software and computing
- The landscape for software is becoming ever more challenging
- Working together on common problems is not only the best use of our resources, our funding agencies will mandate it
- HSF is now well established to help HEP achieve that goal
  - Roadmap and active working groups in key areas
- From the start, working with allied sciences has been an ambition

We would be delighted to do that with the EIC software developers

Inventory of a few common HEP software packages

- Generators: All experiment neutral
- Simulation: Geant4
- Geometry: DD4hep (plus DDsim, DDG4)
- Data model: PODIO (EDM generator), LCIO
- Analysis: ROOT, Coffea, Many other Python packages (scikit-hep, uproot, …)
- Reconstruction: Pandora Particle Flow, A Common Tracking Software
- Frameworks: Gaudi, art, Marlin, O2
- Data Management: Rucio
- Workload Mangement: DIRAC
- Software Distribution: CVMFS

Getting Involved...

- Join the HSF Forum, hsf-forum@gmail.com
  - Few messages a week with updates, jobs, items of interest
  - Owned by the community - please just post items of relevance
- Join a working group, https://hepsoftwarefoundation.org/what_are_WGs.html
  - Follow the group's meetings and discussions
  - Suggest a meeting topic
- Indico Main Page
- Annual meetings
  - Establishing a tradition of a joint meeting with WLCG each Year (next short meeting pre-CHEP, November)
- Propose a new activity area
  - The HSF is there to help gather interest

- Data Analysis
- Detector Simulation
- Frameworks
- Physics Generators
- Packaging
- PyHEP - Python in HEP
- Quantum Computing
- Reconstruction and Software Triggers
- Software/Developer Tools
- Training
- Visualization