International efforts on efficient computing - HEP Software Foundation

Graeme A Stewart, CERN EP-SFT

Efficient Computing for High Energy Physics, Edinburgh 2020-02-17

Future HEP Facilities and Challenges

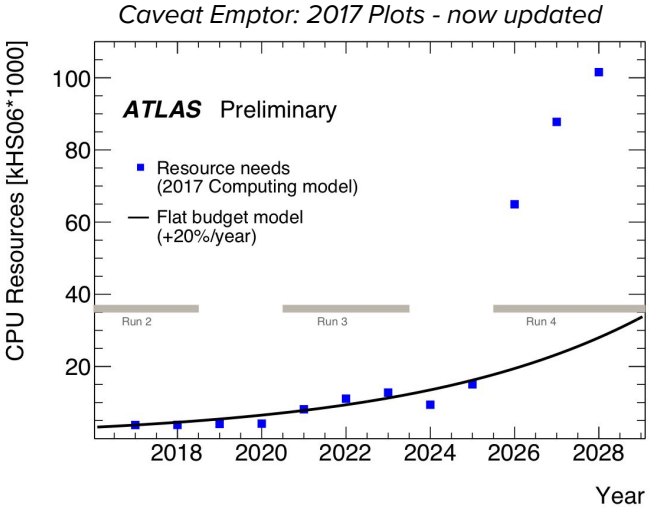

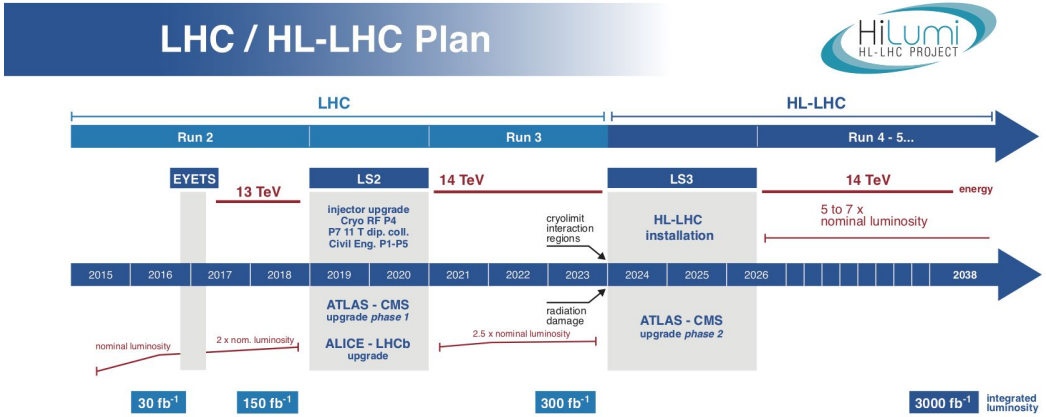

- HL-LHC brings a huge challenge to software and computing
  - Both rate and complexity rise
- Post LHC future now also under intense discussion (CLIC, ILC, FCC, CEPC)

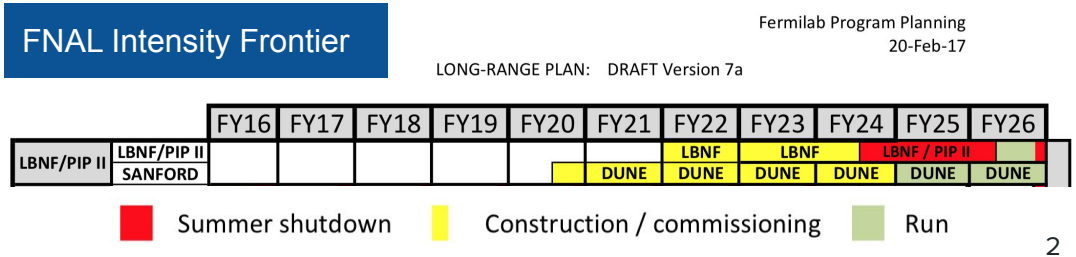

HEP Software and Computing

- 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

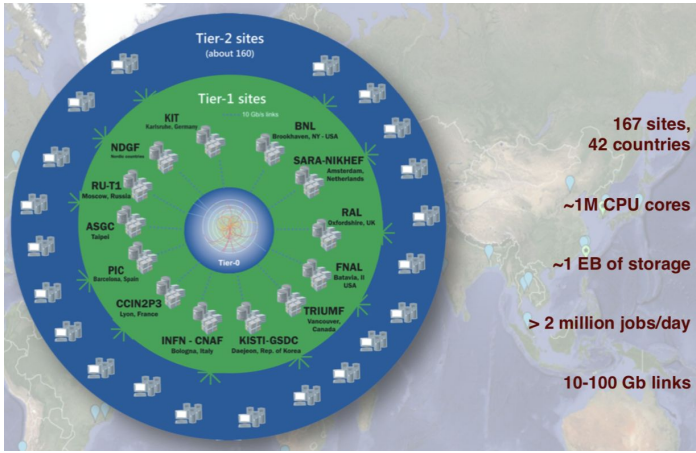

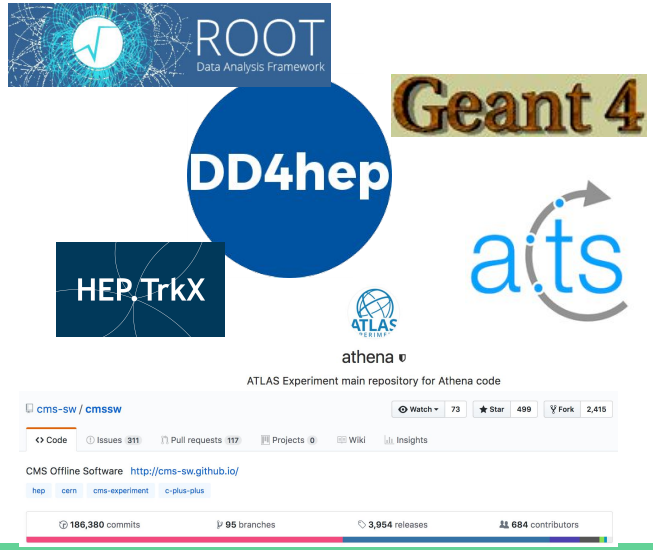

Technology Evolution

- 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

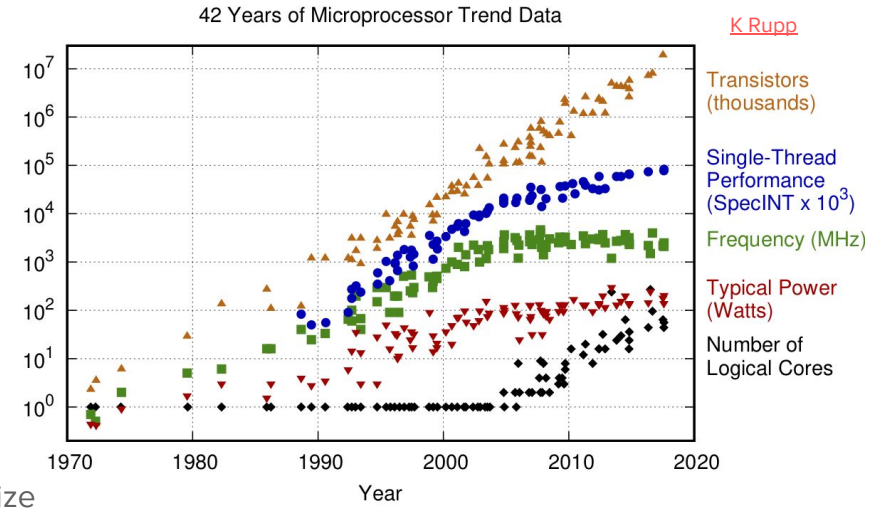

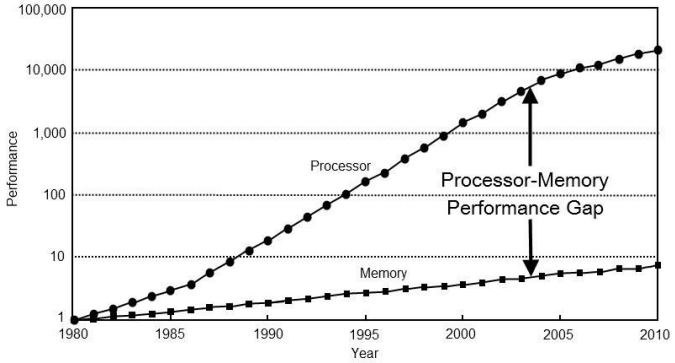

Decreasing Returns over Time

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

[link]

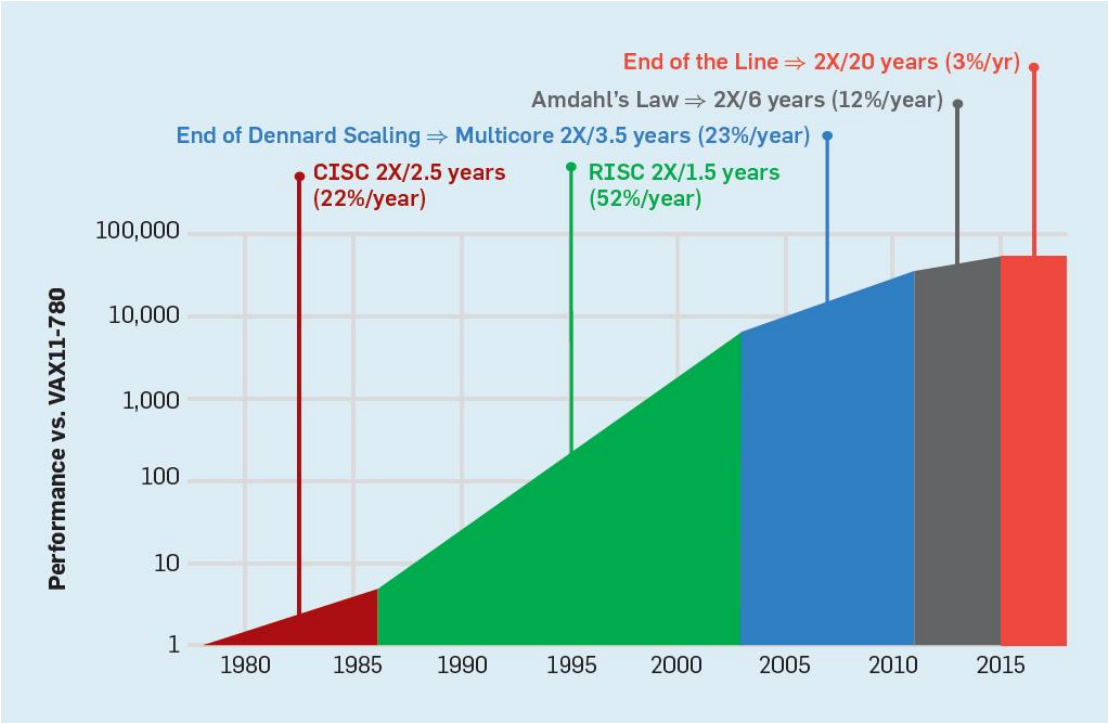

GPUs dedicate far more transistors to arithmetic

Hardware Evolution in a Nutshell

Ye olde processing model

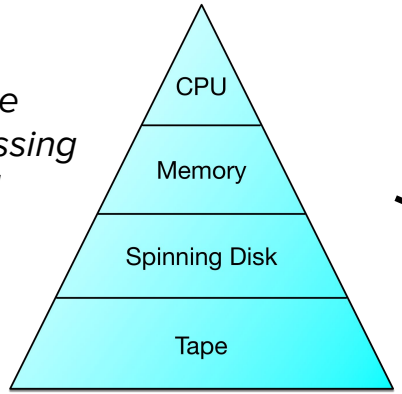

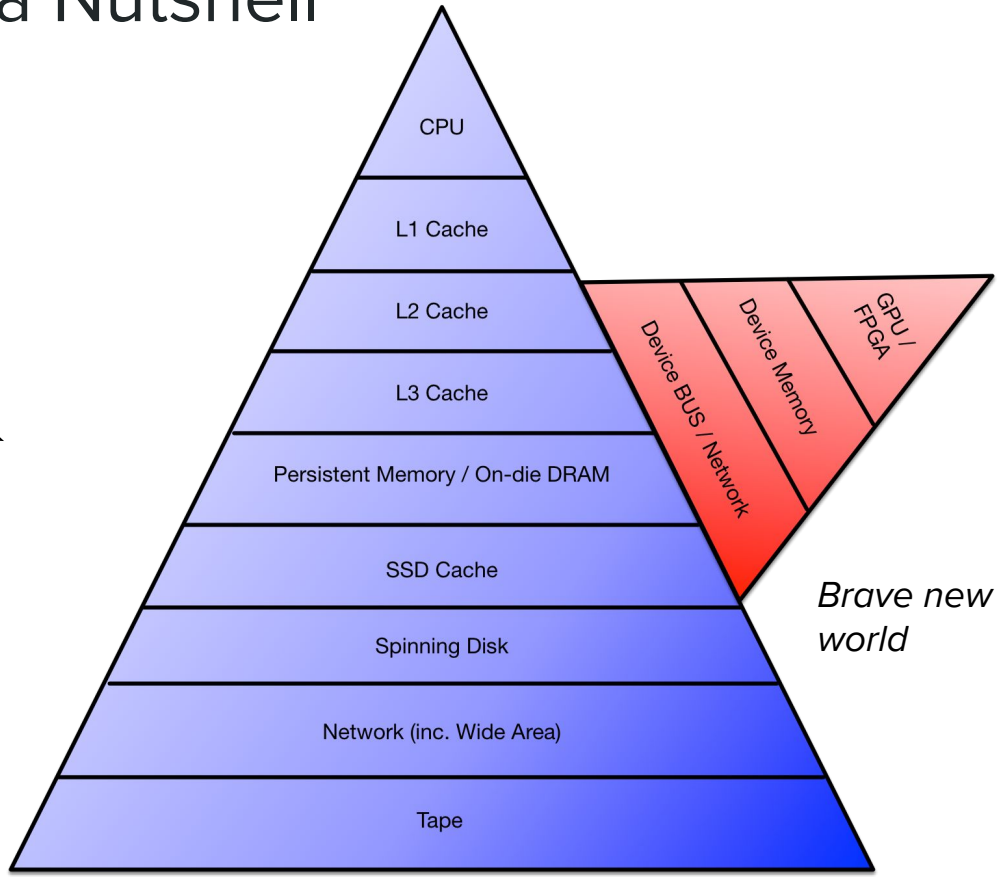

Brave new world

"We're approaching the limits of computer power – we need new programmers now"
John Naughton, Guardian

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, ...)
  - New experiments should not be starting from scratch, but building on best-of-breed
- 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

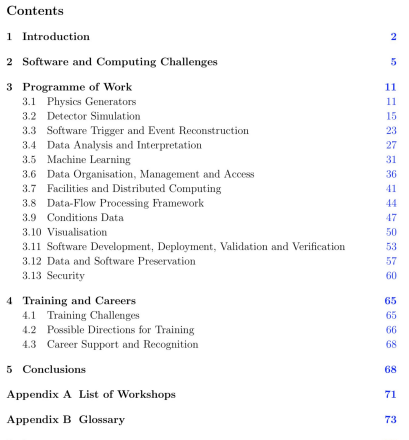
Contents

- 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
    i. to achieve improvements in software efficiency, scalability and performance and to make use of the advances in CPU, storage and network technologies
    ii. to enable new approaches to computing and software that could radically extend the physics reach of the detectors
    iii. 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
- Community engagement: 310 authors from 124 institutes, 14 chapters
- Published in *Computing and Software for Big Science*, https://doi.org/10.1007/s41781-018-0018-8 (and on arXiv)

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 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
  - We setup new working groups for Detector Simulation, Reconstruction and Software Triggers, and Data Analysis
  - HSF Coordination group setup a search committee, whole community could nominate convenors
- 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
- This model was a real success and was expanded this year to all working groups

Laying Good Foundations - Some 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 actively being prototyped (link to Key4hep project in EP R&D)

[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
  - Negligible CPU consumption compared with detector simulation
- 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
- HSF/LPCC Workshop in 2018 spawned a new working group

ATLAS 2018 CPU Report

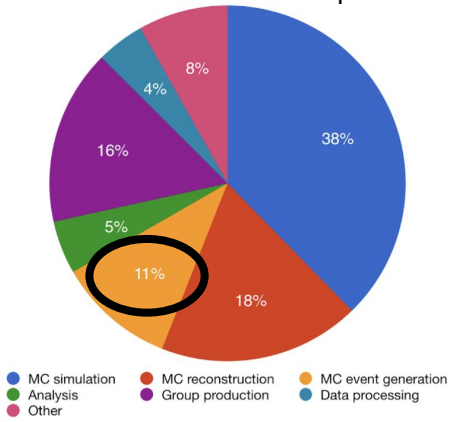

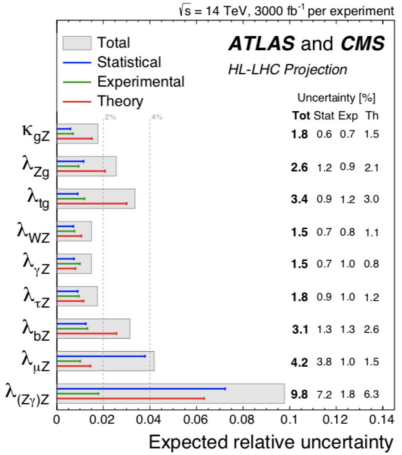

Many electroweak measurement errors dominated by theory (red). B. Hinemann

Event Generators - Technical Improvements and Porting

- 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)
- New Architectures
  - Original port of some elements of MadGraph to GPUs by Japanese Group at KEK
  - Work reinvigorated at ANL over summer 2019
  - HSF WG now actively discussing with MadGraph authors how to continue this
  - Interest also in Sherpa (heavily used by ATLAS)

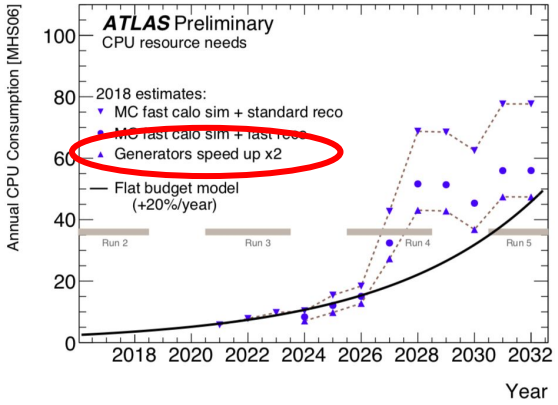

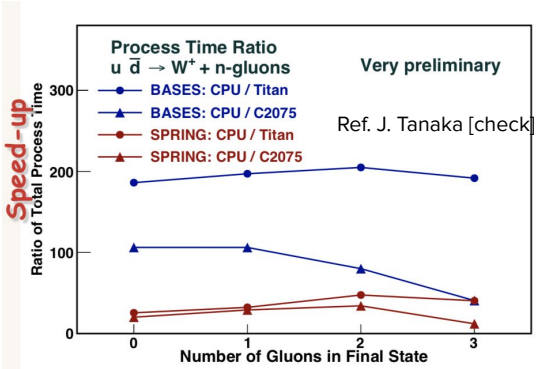

Detector Simulation

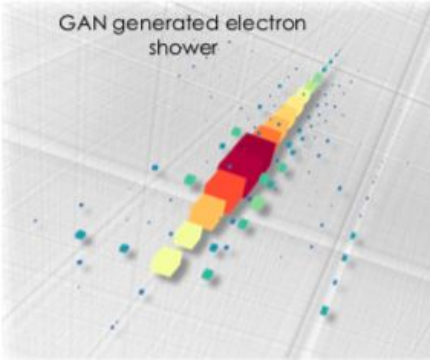

Use of Generative Adversarial Networks to simulate calorimeter showers, trained on G4 events (S. Vallacorsa)

- 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, parametric 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) and Generative Adversarial Networks (GANs)
    - This is probably *not as easy as we thought* - traditional parametric approaches are hard to beat
  - R&D on lifecycle integration into Geant4 is starting

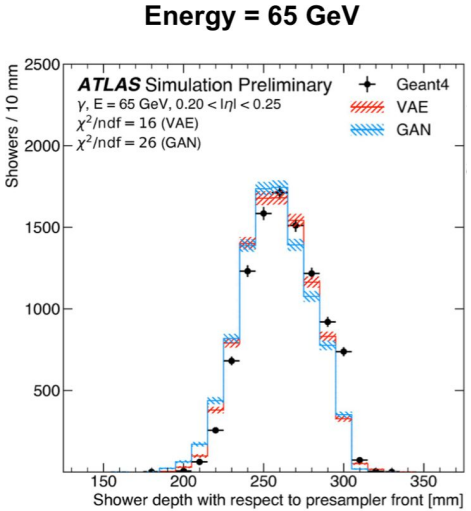

ATLAS VAE and GAN cf. Geant4 simulation [ATL-SOFT-PUB-2018-001.]

Detector Simulation

Processor	GeantV	GeantV-vec	Geant4	G4/GV	G4/GV-vec
SSE4-2.3-15	4457	4333	6627	1.49	1.53
AVX-2.0-15	2621	2331	4938	1.88	2.12
AVX2-2.4-35	1628	1530	2182	1.34	1.43
AVX2-2.5-28	1186	1275	1875	1.58	1.47

GeantV and Geant4 execution speeds for EM physics with and without vectorisation (Andrei Gheata)

- Technical improvement programme helps (and helps *everyone*)
- GeantV R&D modernises code and introduces vectorisation
  - Speed-ups observed (VecCore and VecGeom *backported* to Geant4)
  - Vectorisation introduces small gains, due to costs of “basketisation”
  - Code modernisation seems to help a lot
- Geant4 have an R&D working group that will take studies forward as well as the HSF WG
- Increasingly GPUs are seen as the real target (some architecture convergence)
  - US Exascale Computing Project has funded some work, ongoing at FNAL and Oak Ridge
  - This is a hard problem, but also a relatively green-field area
    - Set realistic expectations
    - Establish a programme of R&D activities (Lund workshop)

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 developments rewriting code for GPUs
    - Physics can get better!
  - Lessons learned: keep data model simple, bulk data, be asynchonous, minimise data transfers
- And design your detector taking software and computing into account
  - ATLAS ITk reconstruction is faster at 200 pile-up than current reconstruction at 60 (aka Don't Panic!)

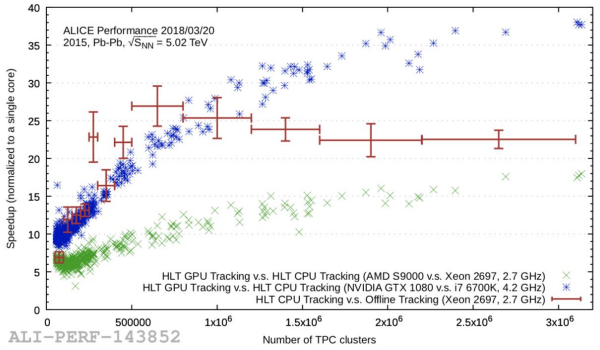

arXiv:1905.05515

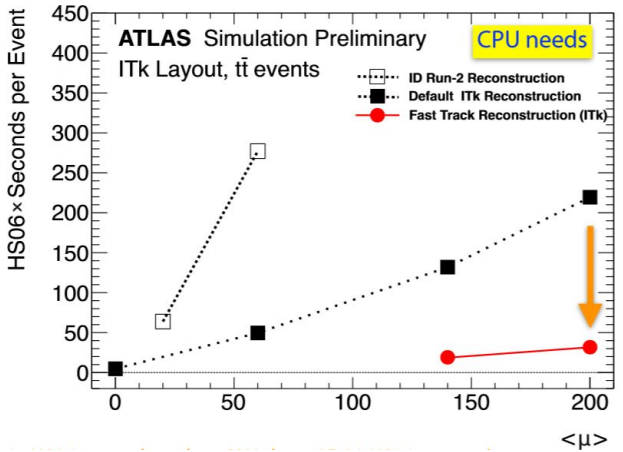

Reconstruction and Software Triggers

- 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)
  - O2 ➡ Online/Offline Data Reduction Farm
- ATLAS and CMS have schemes under development for special handling of samples for which full raw data is unaffordable (aka. data scouting)

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

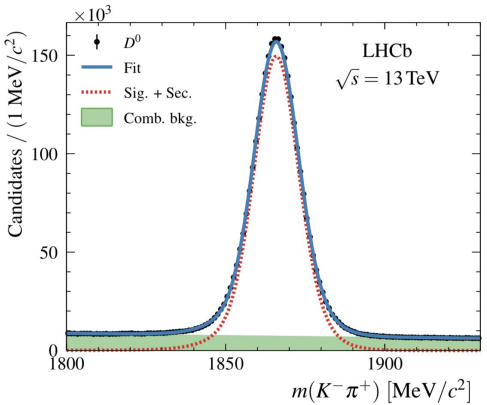

LHCb charm physics analysis using Turbo Stream (arXiv:1510.01707)

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”
- Declarative models, building on ROOT’s RDataFrame
  - Say what, not how and let the backend optimise
  - E.g. split and merge, GPU execution

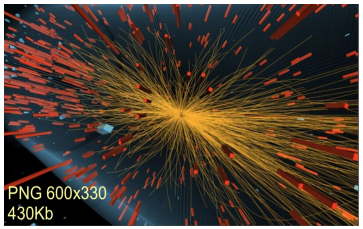

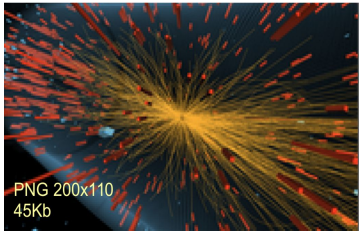

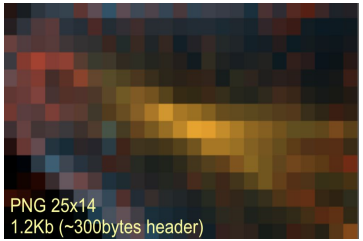

PyHEP (“Python in HEP”) and New Approaches

- Python is ever more popular in Particle Physics
- Impressive developments of a Python scientific ecosystem for HEP in the last 2 years
- With strong links to the general scientific ecosystem
  - Interest in *data science* tools and *machine learning* is significant for this growing community
- Inspiring new approaches for data analysis
  - Exploiting modern approaches - declarative programming, heterogeneous resources, etc.
  - This is an ecosystem into which HEP can, and does, contribute
  - Knowledge transfer goes both ways
  - Various projects under development, inter-communicating
- Yearly PyHEP workshops have been a success
  - This year co-located with SciPy 2020

Training and Careers

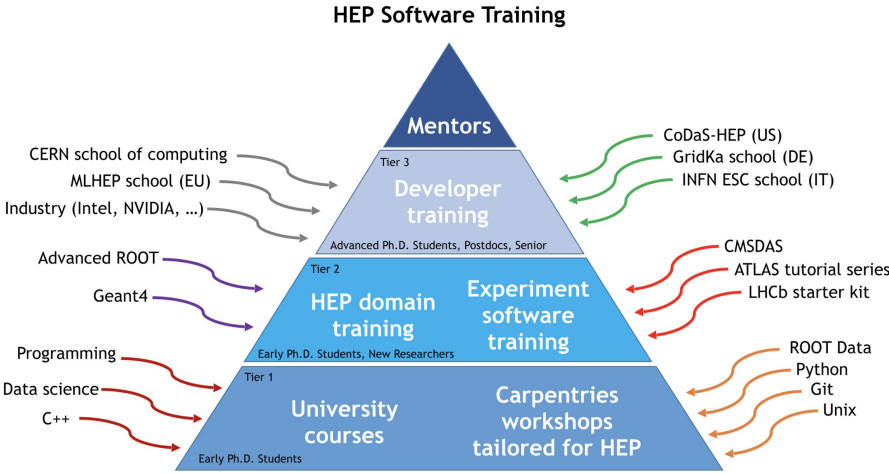

- Many new skills are needed for today's software developers and users
- Base has relatively uniform demands
  - Any common components help us
- LHCb StarterKit initiative taken up by several experiments, sharing training material
  - We ran a Software Carpentries tutorial at CERN last year, another coming this March (co-organised with IRIS-HEP)
- New areas of challenge
  - Concurrency, accelerators, data science (upcoming: oneAPI training from openlab, Alpaka hackathon from openlab/HSF, possible CUDA bootcamp via openlab)
  - 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...

Directions for International Efforts

- Particle physics is in inherently international effort, with an excellent tradition of cooperation in many different domains
  - Detector R&D, Experiments, WLCG, Common Software
- But we have also had incoherent approaches and duplication
- HEP Software Foundation tries to foster much more the shared vision
  - This encourages diverse R&D!
  - **Annual workshop with WLCG is a community focal point** - come to Lund

There is clearly success in attracting funding to this area: IRIS-UK, IRIS-HEP, ErUM-DATA IDT, CERN EP R&D, AIDAinnova; links to other sciences and software engineers via SIDIS. Though overall we believe it is **not yet enough**!

Backup

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

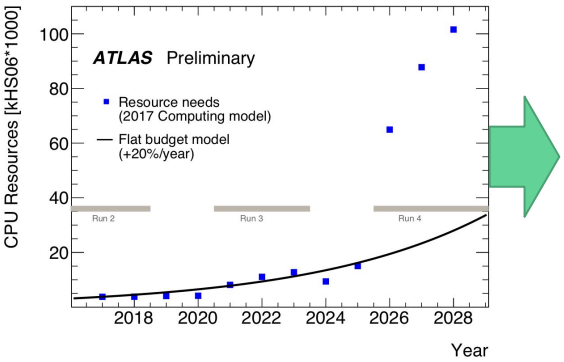

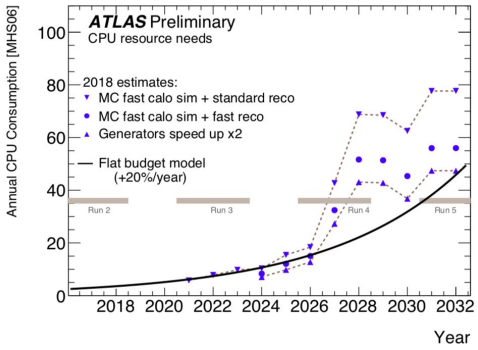

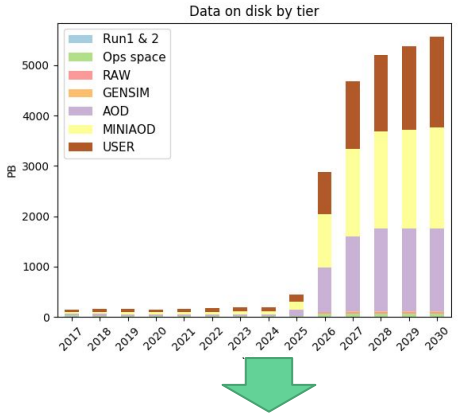

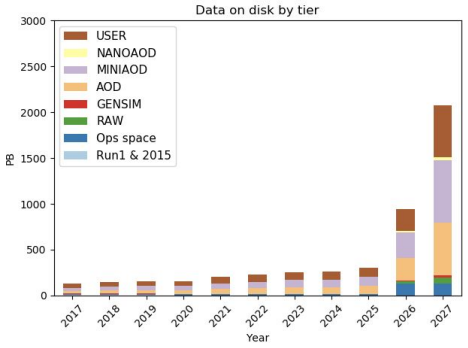

Getting Involved with HSF...

- 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
  - Established tradition of a joint meeting with WLCG each Year
  - This year we will be in Lund, 11-15 May
- Propose a new activity area
  - The HSF is there to help gather interest

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