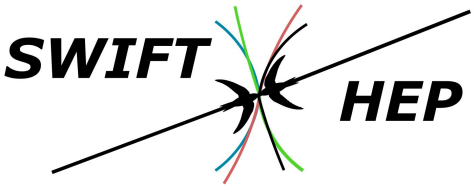


PPAP: Software and Computing

7 July 2023

Dave Britton, Pete Clarke, David Colling, Davide Costanzo, Jon Hays

Outline and introduction (all the content in one slide)

- Software and computing is a “subsystem” of a particle physics experiment
  - Integral part of the experiment. The science we do depends (also) on computing
  - Some of it at the experiment, eg. DAQ, Trigger, prompt processing.
  - Some of it off-site at worldwide centres, eg. offline processing and simulation
  - The physics programme and data rate determines the volume of computing
  - The timeline for computing depends on the experiment timeline and not on the timeline of computing centres availability
  - We need to find space for innovation (GPUs, AI, …) and “blue sky” initiatives (e.g. quantum computing, etc)

- Costs, efficiency, sustainability
  - As any subsystem computing has a cost, and cost/benefit analysis is needed
  - Lots of effort ongoing in reducing cost, more efficient data centres, processors, software, …

- Software and computing for particle physics in the UK
  - As part of a federated international ecosystem
  - GridPP (computing for particle physics), IRIS (federates computing across the STFC portfolio), SWIFT-HEP (common software for particle physics) + experiment dedicated teams
  - Internationally: WLCG, Hep Software Foundation (HSF), Escape, IRIS-HEP(US), CCE (US), …

Software and Computing as part of the experiment

Detector data

Trigger
Real-time system

Experiment operation

Reduced data format

Initial design
Modeling of running exp

Offline processing

User analysis

Simulation (digital twin)

- Physics interaction
- Parton shower
- Particle interactions with material
- Detector response

The computing demand in the next decade

- Demand will increase
  - Driven by detectors with more channels, higher rate triggers
  - This is not just the LHC: DUNE, EIC, SKA, LSST, …
- The working assumption is that compute resources will increase by 10-20% every year for the same cost
- Improvement in efficiency can be achieved by:
  - Using new architecture with lower cost per cycle (such as ARM)
  - More efficient software
  - Co-design the computing to suit the workloads

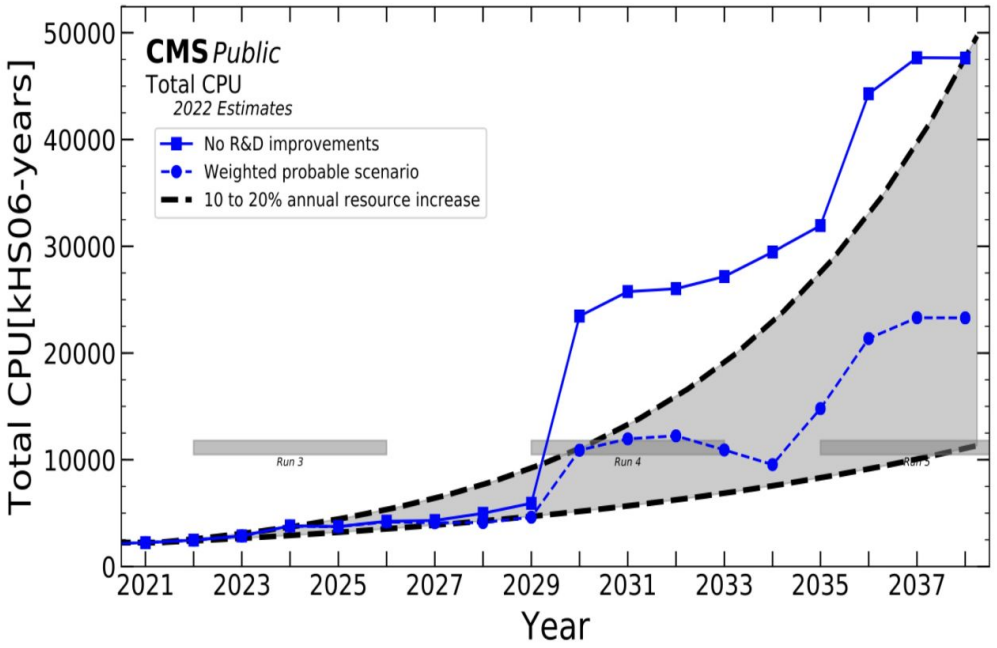

Evolution of computing hardware

- Industry drive in reducing power and cost of compute hardware
  - This is not necessarily linear
  - Two examples: AMD(x86) chips in 2020, ARM_64 happening now
  - Could GPU improve efficiency further?

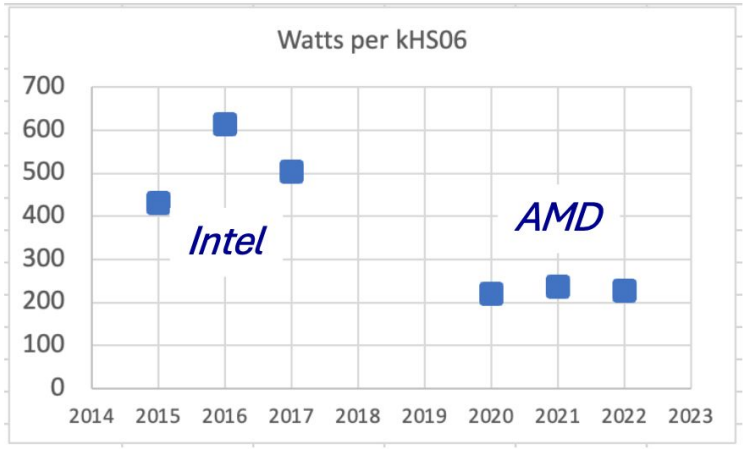

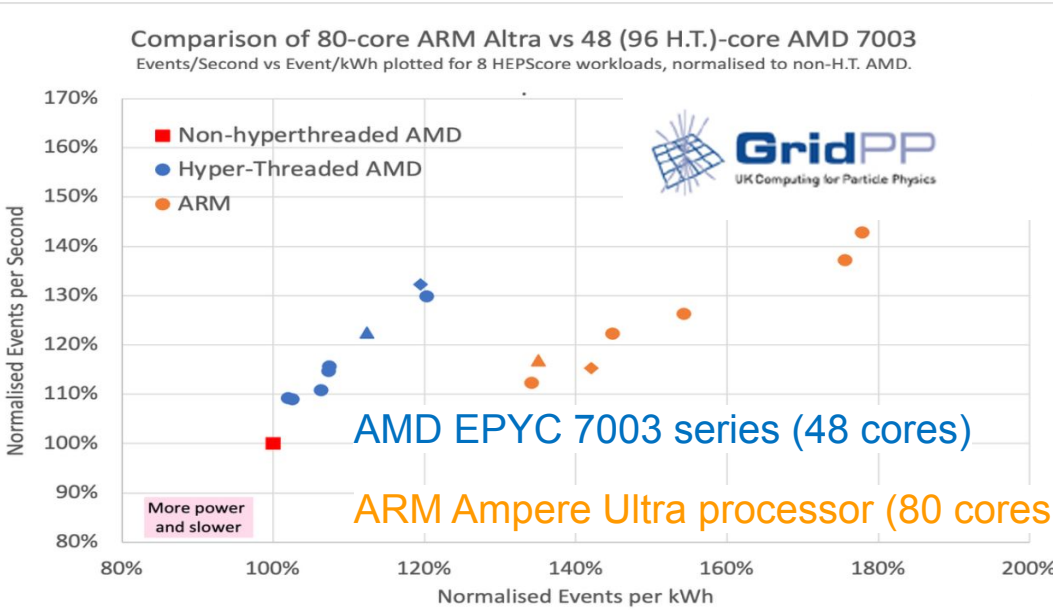

Software improvements. An example from event generators

- Same physics results with less computing (more info)
  - Example from Sherpa (possible thanks to SWIFT-HEP funding)
- Event generators are a key ingredient
  - Higher precision generators needed for high precision physics
  - Not just the LHC, also neutrino physics (e.g. Genie)
- Simulation is a major CPU consumer
  - Continuous optimisation (example in backup)
  - Role of GPUs and AI
- Reconstruction, trigger, analysis
  - More user facing
  - Large compute contributor

Breakdown of CPU budget in $t\bar{t}$+jets

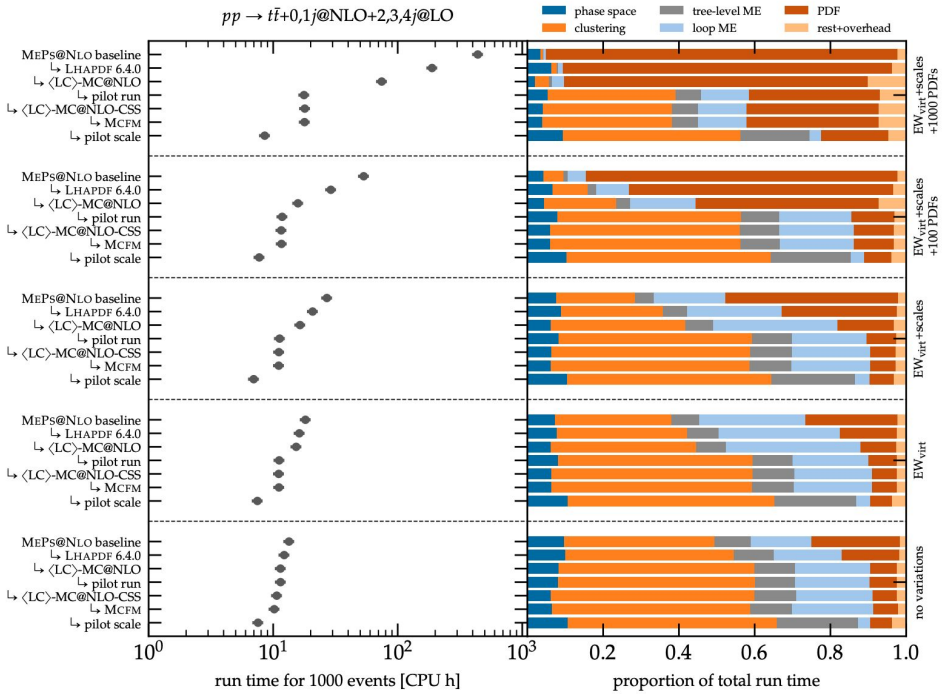

Software and Computing in UK particle physics

- Users interact with the e-Infrastructure from their institutes
  - Very important role of local computing infrastructure (e.g. Tier-3s)

- e-Infrastructure based on a federation concept:
  - Wikipedia: Federated architecture […] allows interoperability and information sharing between semi-autonomous de-centrally organized lines of business, information technology systems and applications.

- STFC projects:
  - **GridPP**: Provides the compute, disk, tape needed to operate PP experiments.
    As well as the software fabric and support needed to run the experiment payloads.
  - **IRIS**: set up to bring coherence, collaboration, coordination and sharing across all of STFC
  - **SWIFT-HEP**: development of common software for PP experiments
- STFC e-Infrastructure Advisory Group (SEAG)
  - Reporting to both science boards. Should interact with other advisory panels

- UKRI projects:
  - ExCALIBUR towards exascale systems (ExaTEPP as a PP working group + Hartree + DiRAC)
  - Many discussions about exascale

Strategic review of particle physics - Comments

- e-Infrastructure is in constant evolution:
  - Enable a growing users base to do science
  - Satisfy the evolving needs and requirements of the experiments
  - Is part of an international community
  - Follow the advancements in hardware and software technology: driven by industry

- The “strategic review of particle physics” makes some comments on computing.
  - There is some misunderstanding, and they appear not to be aware of IRIS or SwiftHEP
  - Most of the things recommended are already done, in particular the common project for experiment software efficiency. (SWIFT-HEP)
  - “Efficiency” may mean many things and require constant investments:
    - Shared and fully used resources (no empty CPUs)
    - Power efficiency through work UK leads in WLCG including ARM introduction
    - Software efficiency gains - the low hanging fruit has been done
    - Joined up computing across STFC to share resources and expertise where it makes
  - Engagement with UKRI DRI: This is happening already, for more than 6 years
    - A much larger community, with many different viewpoints

Outline and introduction (all the content in one slide)

- Software and computing is a “subsystem” of a particle physics experiment
  - Integral part of the experiment. The science we do depends (also) on computing
  - Some of it at the experiment, eg. DAQ, Trigger, prompt processing.
  - Some of it off-site at worldwide centres, eg. offline processing and simulation
  - The timeline for computing depends on the experiment timeline and not on the timeline of computing centres availability
  - We need to find space for innovation (GPUs, AI, …) and “blue sky” initiatives (e.g. quantum computing, etc)

- Costs, efficiency, sustainability
  - As any subsystem computing has a cost, and cost/benefit analysis is needed
  - Lots of effort ongoing in reducing cost, more efficient data centres, processors, software, …

- Software and computing for particle physics in the UK
  - As part of a federated international ecosystem
  - GridPP (computing for particle physics), IRIS (support non PP experiments, including Astro), SWIFT-HEP (common software for particle physics) + experiment dedicated teams
  - Internationally: WLCG, Hep Software Foundation (HSF), Escape, IRIS-HEP(US), CCE (US), …

Backup

The fruit is not hanging that low

- SWIFT-HEP operates on “common” software
- The experiments have a very important role to play
  - Most CPU time on WLCG spent on Simulation
  - Example from ACAT talk

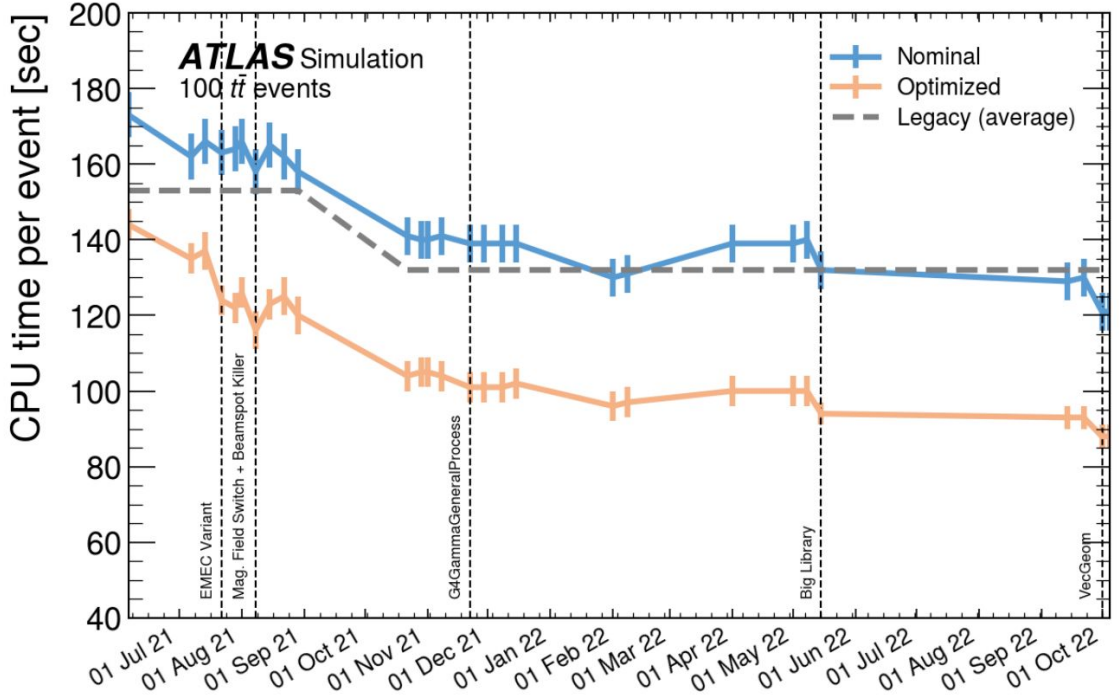

Why (primarily) Federation and not Consolidation?

- Federation allows composition of new solutions out of existing investment (You can only rearrange the building blocks if you have the building blocks).
- Federation enables decision making to be devolved "down" the hierarchy to where it best sits, improving choices and protecting against domination of one community or voice to the detriment of the rest.
- Federation can empower communities in a way that consolidation does not. All these elements become particularly important as the scale grows.
- Federation encourages diversity, of ideas, solutions, and people. It can protect against "group think" and stagnation, and can provide resilience against single points of failure – both geographical and technological.
- Federation enables low risk evaluation and testing of "future" technologies, in particular where they are driven by specific well motivated communities that would otherwise be overlooked or dismissed by a large scale operation with a consolidated approach.
- Federation allows smaller operations to benefit from the full scale of the federation. E.g., security, identity management, accounting and allocation; but also in the building of larger communities to share ideas and solutions.
- Federation allows leveraging of local resources that otherwise would not be available.