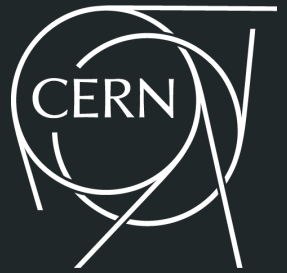

Software and Computing R&D: International Perspective

Graeme A Stewart

PPTAP Computing & Software Roadmap Workshop, 2021-07-19

HEP Software and Computing

- Critical part of our physics production pipeline, from triggering, through production, to analysis and final plots
- >~50 millions of lines of code, mainly C++, a lot of Python
  - Commercial development cost ~500M CHF
- Significant pieces of software are already shared by most experiments:
  - Event generators, Geant4, ROOT plus WLCG and computing software, like Rucio
- LHC experiments use
  - 1M CPU cores every hour of every day
  - Store 1000PB of data (600/400PB tape/disk split)
  - 100PB of data transfers per year (10-100Gb links)
- This is a huge, ongoing cost in hardware and human effort
- With significant challenges ahead of us to support our developing physics programme
  - Trigger rates go to ~x10 for ATLAS and CMS in Run4

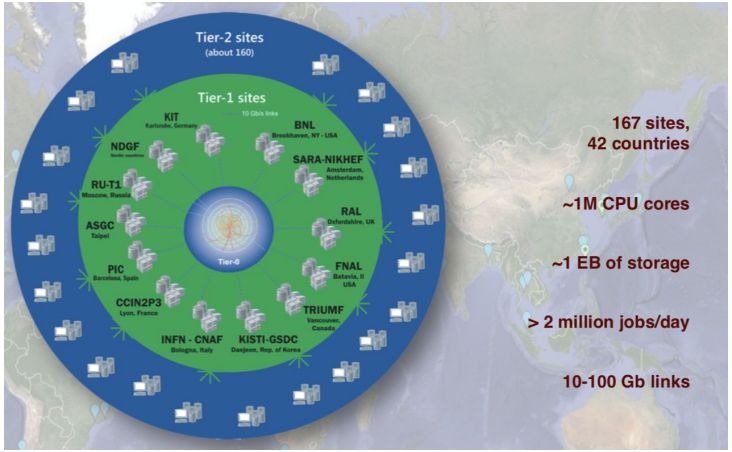

Technology Challenges

- 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 - stuck at ~3Ghz
  - Leak currents become important source of power consumption
  - Memory access times are now ~100s of clock cycles
    - Poor data layouts are catastrophic for software performance
- From a CPU x86_64 monoculture we must evolve towards heterogeneous computing
  - Certainly including GPUs
  - FPGAs, TPUs, etc. also may play a role

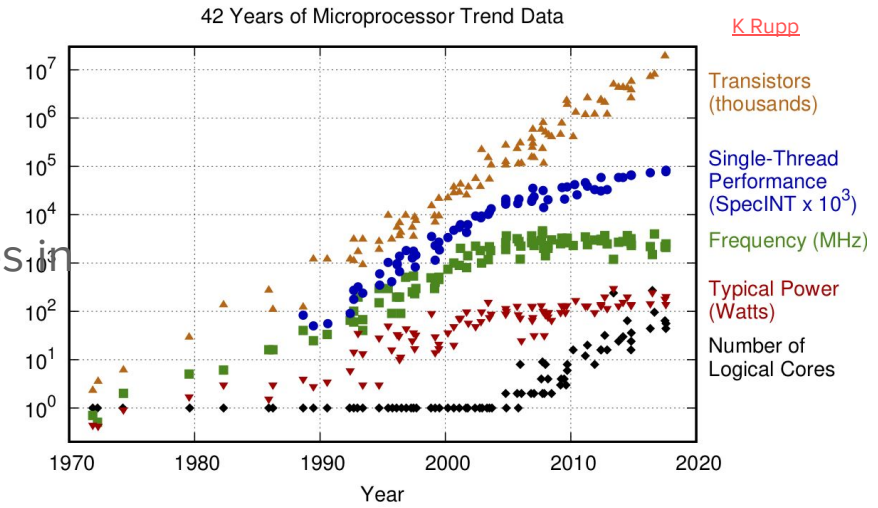

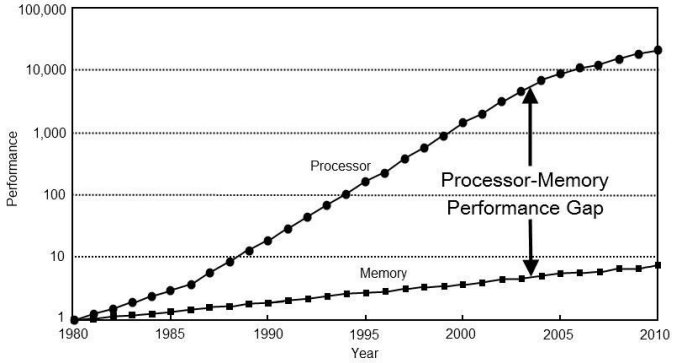

HEP Software Foundation and the Community White Paper

- Process started in 2015 to try to map out a path for software
  - HSF formally charged by WLCG to
    - Anticipate a "software upgrade" in preparation for HL-LHC
    - Identify and prioritize the software research and development investments
      - 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
- Two workshops (San Diego and Annecy)
  - *Active working groups* who took charge of particular chapters of the paper
    - Including holding their own topical workshops
  - Process helped greatly by financial seeding of NSF via DIANA-HEP (pre pandemic travel!)

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

HSF-CWP-2017-01
December 15, 2017

- 70 page document [1712.06982; doi:10.1007/s41781-018-0018-8].
  - 13 sections summarising R&D in a variety of technical areas for HEP Software and Computing
  - 1 section on Training and Careers
  - 310 authors from 124 institutions
- The outcome of the CWP was
  - A strong argument for R&D funding in computing and software, with an emphasis on common projects
  - Establishment of HSF working groups to promote exchange of ideas and cooperation between experiments and projects

Contents

Overview of R&D Projects

HSF does not itself seek funding, but supports bids to funding agencies. Many of these projects received a letter of support or collaboration from the HSF:

- IRIS-HEP, NSF
- ErUM-DAT, Helmholtz Institute
- EP R&D, CERN
- HEP-CCE, DOE
- AIDAInnova, European Commission
- *SWIFT-HEP, STFC*
- *ExCALIBER-HEP, UKRI*

I will give an overview of the non-UK R&D projects here - errors and omissions are my fault!

Note that there is also a continuing process of R&D in existing long running projects, particularly Geant4 and ROOT, where collaborations happen and new ones are very welcome

IRIS-HEP

- US National Science Foundation funded for 5 years from 2019
  - Grew out of conceptualisation forerunner project DIANA-HEP
- Some main areas of interest:
  - Analysis systems
    - Support for Data Science ecosystem tools (uproot, awkward array, etc.)
    - Development of columnar analysis systems, Coffea
  - Data Organisation Access and Management (DOMA)
    - Modernised data delivery methods
    - Storage, caching and end-to-end analysis, Coffea-casa
  - Innovative Algorithms
    - Efficient parallel tracking algorithms, accelerators, machine learning
- Plus:
  - Blueprint meetings, Training, Scalable Systems Lab, OSG, Grand Challenges

ErUM Data IDT

- Innovative Digital Technologies for Research on Universe and Matter (IDT-UM)
  - Funded for 3 years from October 2019
- Active areas of R&D:
  - Technologies for heterogeneous computing and virtualized environments
    - Cloud resource integration, containers, data caches
    - Dynamic resource provisioning, COBaID-TARDIS
  - Reconstruction and Simulation via Machine learning techniques
    - ATLAS L1, Air-Showers, CBM Experiment
    - Generative models for calorimeter simulation (vCHEP and [2102.12491])
  - Tracking and tagging
    - Model datasets [ML4Jets talk] and work in ACTS
- New funding call recently announced, 15-20M€ over 3 years, broader than HEP
- Ideas to organise the German community in this area as well,

CERN EP R&D for Detector Technologies

- R&D programme started in 2020 for 5 years, covering a wide range of detector technologies - including software
- Software task areas:
  - Key4hep
    - Turnkey software stack for detector development studies (simulation, reconstruction and analysis), including accelerator friendly data model generator PODIO [HSF WS talk]
  - Faster Simulation
    - Machine learning based simulation of calorimeters, and deep integration into Geant4 [talk]
  - Reconstruction at high pile-up, with work on tracking (ACTS) [talk]; and high granularity calorimeters [talk]
  - High performance analysis, next generation ROOT data format, RNTuple and integration into object stores [talk]

HEP Center for Computational Excellence

- DOE funded programme, 3 years from 2020
- Large motivation to develop solutions that allow HEP to take advantage of current and upcoming supercomputer facilities
- Research areas (see last meeting for a good overview):
  - Portable Parallelization Strategies - heterogeneous computing APIs
    - Porting of 3 HEP applications (Wirecell, FastCaloSim, Patatrack) to different heterogeneous backends (CUDA, Kokkos, OneAPI) [vCHEP]
  - I/O Strategies
    - HDF5 data format for HEP
    - I/O performance measurements and optimisations
  - Event Generators
    - Re-engineering of Sherpa

AIDAInnova

- Advancement and Innovation for Detectors at Accelerators, EU Funded project from April 2021 for 4 years
  - Follow-on project to AIDA2020 and AIDA
- Software work package, WP12 (good summary at kick-off meeting)
  - Task 12.2. Turnkey Software (DESY, CERN, INFN)
    - Integrated Turnkey Software Stack (Key4hep); Data model toolkit for modern hardware (PODIO); Digitisation extensions for geometry (DDDigi); R&D study on frameworks for heterogeneous resources
  - Task 12.3. Simulation (CERN, DESY, IJCLab, Manchester)
    - Fast simulation techniques integrated into Geant4
    - Machine learning based calorimeter simulation toolkit for training and inference
  - Task 12.4. Track Reconstruction (IJCLab, CERN, INFN)
    - Complete track reconstruction chain with ACTS composable algorithms; portable version of ACTS algorithms, for heterogeneous computing; Machine learning reconstruction algorithm for MPGD detectors
  - Task 12.5. Particle Flow Reconstruction (Warwick, Cambridge, CNRS, INFN, Sussex) - w. Pandora integration
    - Advanced PFA algorithms for DUNE detectors using new readout technologies; PFA algorithm with particle ID for dual-readout calorimeters; Optimised APRIL PFA algorithm for hadronic jets

Summary

- Long standing recognition in the HEP community that software and computing is a challenge for the future
  - Driven by physics programme and technology
- HEP Software Foundation has established itself as an umbrella that can help coordination and communication across HEP
  - Starting with marshalling the community for the CWP Roadmap
  - Continuing through the work of its WGs and role in, e.g., LHCC HL-LHC software and computing review
    - Could be a case for now revisiting that process more generally
- Many R&D projects have started, with more of an emphasis on common software (including broader than HEP in a few cases)
  - Initiatives like Software Institute for Data Intensive Science are also starting
- There is still a lot of very relevant work to do and it's heartening to see projects like Rucio and ACTS mature and develop a real international community behind them