The HEP Software Foundation and its activities

Caterina Doglioni - University of Manchester & Lund University
she/her, @CatDogLUOfM

Much input from: Graeme Stewart's ILCX talk, HSF Coordinators and Community

Introduction to the HEP Software Foundation

Caterina Doglioni - on behalf of the HEP Software Foundation - 2021/11/02 - SwiftHep/Excalibur Workshop

Slide by G. Stewart

HL-LHC, the Intensity Frontier, and beyond

(Some of) our scientific questions:

- Exploit the Higgs for SM and BSM physics
- b, c, tau physics to study BSM and matter/antimatte
- Explore the unknown, e.g. dark matter
- QGP in heavy ion collisions
- Neutrino oscillations and mass

Our tools:

- [present] (HL-)LHC, DUNE, Belle II
- [longer-term] ILC, FCC, CEPC, BEPC
- also connected to HSF:
  - smaller LHC-adjacent experiments (e.g. FASER)
  - nuclear physics experiments (e.g. FAIR, EIC)

FNAL Intensity Frontier

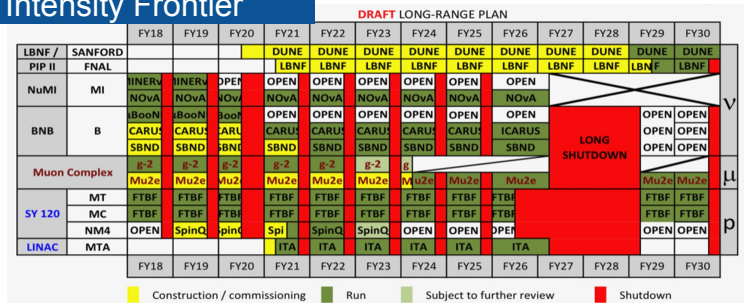

Shared background: HEP Software

- 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 1000PB of data transfers per year (10-100Gb links)
  - We are in the exabyte era already
- This is a *huge* and *ongoing* cost (hardware & human effort)
- With **significant challenges** ahead of us to support our ongoing physics programme

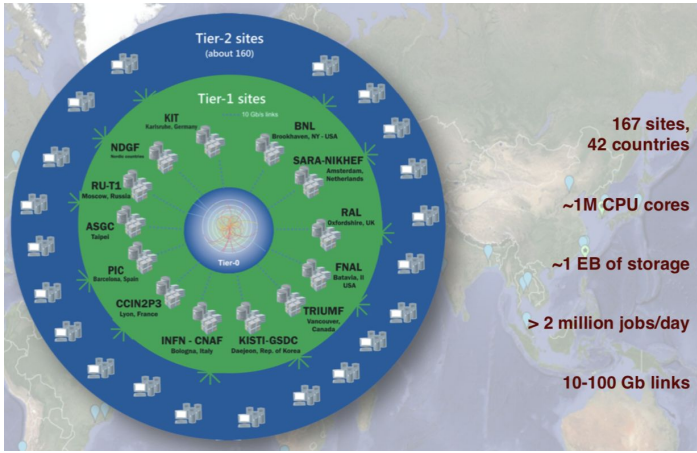

Slide by G. Stewart

Common challenges & needs in the next decade

Increase in transistor density not as fast as increase in demands (and does not mean increase in code performance)

+

Memory access times not keeping up with processor performance

CMS Note 2021/001

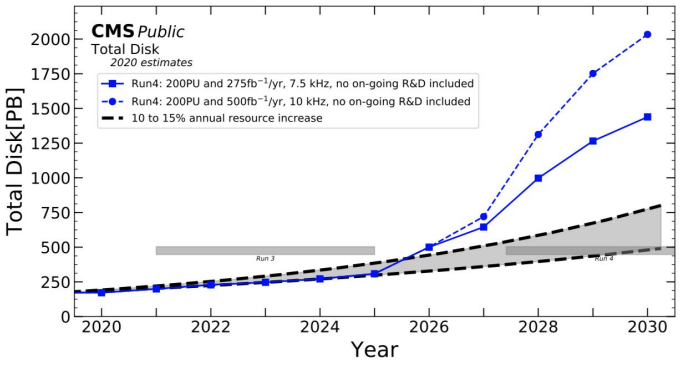

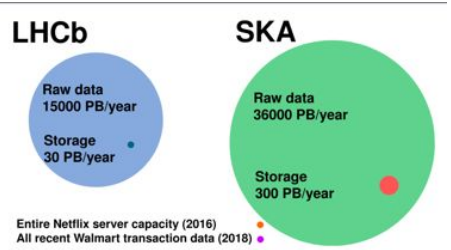

C. Fitzpatrick

ComputingandSoftwarePublicResults < AtlasPublic < TWiki

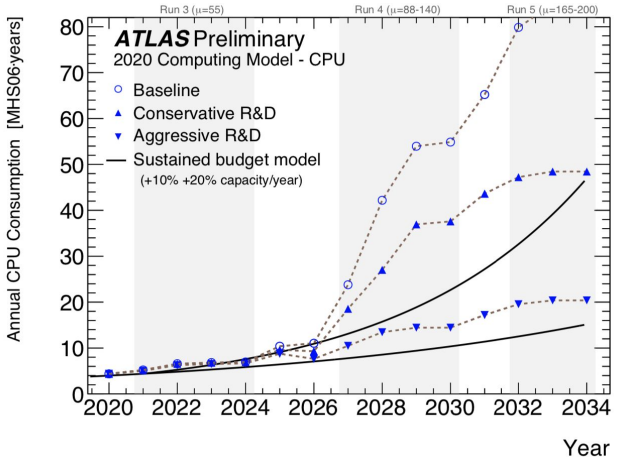

Joint work on software, and **adaptation to new hardware** (HPC, commercial accelerators) can be key to making the most of HL-HLC data and more

Need for joint work recognized in Update of European Strategy of Particle Physics

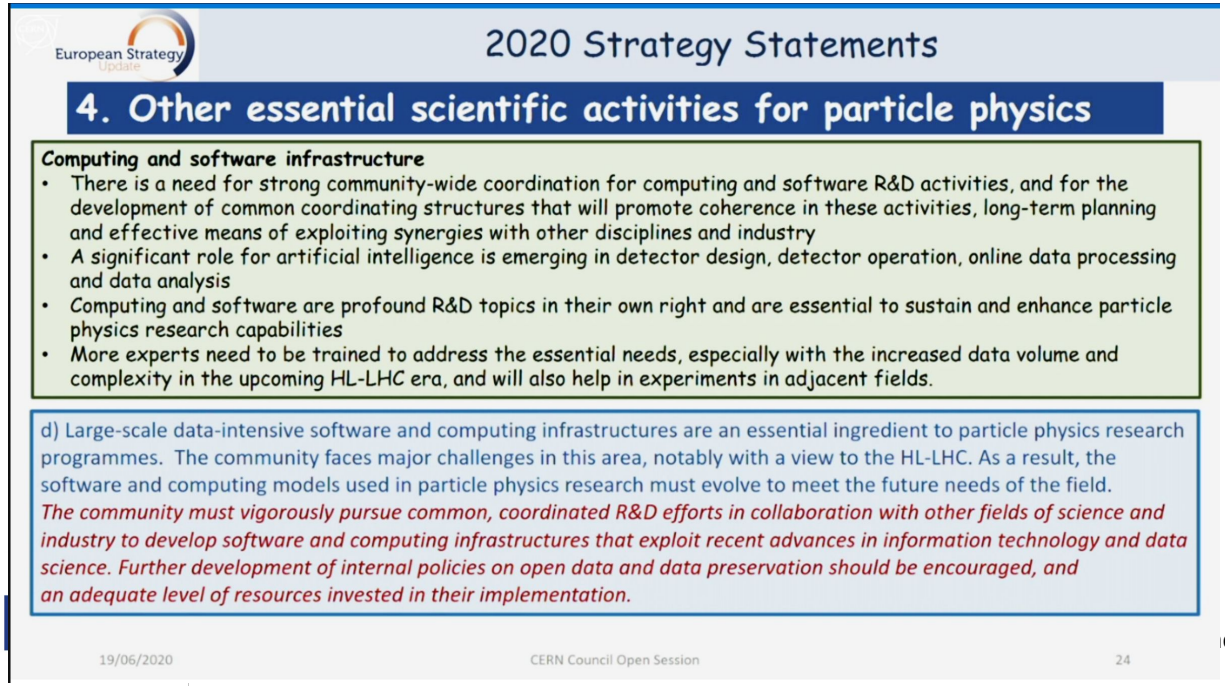

European Strategy Update

2020 Strategy Statements

4. Other essential scientific activities for particle physics

Computing and software infrastructure

- There is a need for strong community-wide coordination for computing and software R&D activities, and for the development of common coordinating structures that will promote coherence in these activities, long-term planning and effective means of exploiting synergies with other disciplines and industry
- A significant role for artificial intelligence is emerging in detector design, detector operation, online data processing and data analysis
- Computing and software are profound R&D topics in their own right and are essential to sustain and enhance particle physics research capabilities
- More experts need to be trained to address the essential needs, especially with the increased data volume and complexity in the upcoming HL-LHC era, and will also help in experiments in adjacent fields.

d) Large-scale data-intensive software and computing infrastructures are an essential ingredient to particle physics research programmes. The community faces major challenges in this area, notably with a view to the HL-LHC. As a result, the software and computing models used in particle physics research must evolve to meet the future needs of the field. *The community must vigorously pursue common, coordinated R&D efforts in collaboration with other fields of science and industry to develop software and computing infrastructures that exploit recent advances in information technology and data science. Further development of internal policies on open data and data preservation should be encouraged, and an adequate level of resources invested in their implementation.*

19/06/2020 CERN Council Open Session 24

European Strategy Update, deliberation document

The upgrades of the ATLAS and CMS experiments have been documented in a series of Technical Design Reports and have been approved, and the international collaborations are gearing up to commission these detectors by 2027, the scheduled start of the first HL-LHC run. The timely delivery of these upgrades is a milestone for the global particle physics community, and the continued allocation of adequate resources is a priority. Based on continued innovations in experimental techniques, the untapped physics that is surely awaiting in the third LHC run and the HL-LHC era can be unlocked. Incorporating emerging new technologies into trigger systems, computing and management of big data, reconstruction algorithms and analysis methods is the path to get the best out of these upcoming datasets.

The HEP Software Foundation [website]

HSF | Working Groups | Activities | Meetings | Communication | Projects & Support | About

HEP Software Foundation

The HEP Software Foundation facilitates cooperation and common efforts in High Energy Physics software and computing internationally.

- Common meetings and community events
  - Organized in Working Groups
- Cross-experiment activities (including training)
- Support to ongoing and future projects
- Our philosophy is bottom up, a.k.a. *do-ocracy*

HSF activities (past/present)

HSF activities

- Initial goal: describe vision for software and computing taking us to 2020s
  - Community whitepaper *"A Roadmap for HEP Software and Computing R&D for the 2020s"* (Comp. Soft. Big. Sci, arXiv) - sub-chapters also submitted to arXiv
  - Stepping stone for funding for Institute for Research and Innovation in Software for High Energy Physics (IRIS-HEP) [US effort]
  - Updates of individual chapters / overall document for LHCC reviews (see later)
- Organization in lightweight coordination team + Working Groups:
  - Event generation, detector simulation, reconstruction, analysis, frameworks, tools and packaging, education and training, Python in HEP
  - Also supporting more lightweight activity areas (e.g. differentiable programming, quantum computing)
- HSF's role: information conduit, meeting point
- Engagement with prioritization efforts (Europe/US)
  - Submission to the Update to the European Strategy of Particle Physics
  - Submission to Snowmass

See talks in **Wednesday sessions** for **connections and synergistic work** in Swift-HEP

Selected working groups: recent activities

Event Generators

First step of all simulated physics

- State-of-the art generators: next-to-leading order → non-negligible computational costs, especially for precision physics
- Clear connection with GPU and HPC [2018 workshop discussion] → priority of WG

arXiv:2004.13687.pdf, see also arXiv:2109.14938

Porting and optimizing generators on GPUs is especially important to be able to exploit modern GPU-based HPCs (such as SUMMIT [163], where 95% of the compute capacity comes from GPUs [164]). Some work

Simulation

More data rates/complex detectors → more simulation

- Liaising with Geant IV community (workhorse)
- New directions:
  - Tracking on GPUs [see talks tomorrow]
  - ML-aware simulation (MLSim meetings)

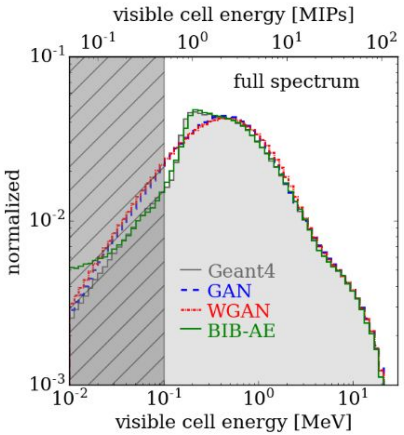

Performance of different ML architectures for photons in ILC Calorimeter, S Diefenbacher et al.

Selected working groups: recent activities

Reconstruction & Software triggers

Software that is close to the detector

- **Tracking** is a big resource consumer
  - Support common tools to test new algorithms (e.g. ACTS)
  - Optimize/vectorize
  - Port algorithms to GPUs (e.g. Allen)
  - See afternoon talks
- Ongoing change of paradigm: from offline analysis to **real-time analysis** directly in the trigger (Turbo Stream, Data Scouting, Trigger-Level Analysis)
  - Additionally: use of non-standard workflows for non-standard physics signatures
  - New UK opportunities to collaborate (e.g. ETN graduate network SMARTHEP received a HSF letter of support, coordination in Manchester)
- Upcoming work on **timing** detectors

Data analysis

Software that is close to the physics results

- A common wishlist & ongoing discussions:
  - Compact physics objects for analysis
    - ROOT formats undergoing changes
  - Fast and scalable I/O & analysis software
    - Use of GPUs can help
  - Reinterpretable/reproducible analysis
    - Essential to agree on metadata/languages
- Current focus and discussion avenues:
  - Definition of benchmarks & metadata
  - Test end-to-end workflows, e.g. IRIS-HEP Grand Analysis Challenges
  - Discussions with other communities, e.g.
    - Neutrino, Nuclear physics, dark matter [meetings link]
    - Smaller experiments [see also Snowmass workshop]

Selected working groups: recent activities

Frameworks

Discussion forum on integration of software efforts into experimental frameworks (& more)

- Many software updates during LS2
  - E.g. multithreading rewrite of ATLAS HLT and offline software
  - E.g. introduction of GPU (and software consequences) in CMS and LHCb HLT
- This group serves as a forum to discuss both individual and common issues to foster common approaches
  - Recent topics of interest: pile-up overlay, multithreaded callbacks and monitoring, joint meeting between HSF experts and DUNE on framework requirements

Software Developer Tools & Packaging

Aid development using state-of-the-art tools including building and deploying software stacks

- Most recent meetings about (non-HEP) packaging tool Spack
  - To be used to integrate software ecosystem into HPCs in the US
  - Being tested by Fermilab
  - Survey from 2016 tech paper here
- Very recent discussions: Anaconda changed its Terms of Service (towards more commercial approach) → consequences for HEP software?
  - Possible Terms of Service violations currently being investigated by Fermilab experts

Slide by G. Stewart

Selected working groups: recent activities

PyHEP

The HEP community's broader Python community

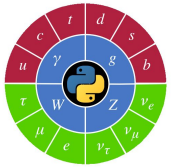

- Started in 2018 at CHEP Sofia, blossomed into two very successful virtual workshops from 2020 onwards
- Huge interest in this area (> 1300 registrants)
  - Topics: data science and ML toolkits
  - Integration with particle physics tools (Coffea, PyROOT, Scikit-HEP, SWAN, zfit)
  - New trends included automatic differentiation [new HSF activity area started before the summer]
- Many talks and tutorials done as notebooks
  - Participants could follow live or use them as offline resources, integrated into Binder
- Everything uploaded to the HSF's Youtube Channel & captioned

Training & Careers

Building skills and recognition of software efforts

- HSF Training Group runs Software Carpentries and other tutorials
  (co-organised between the HSF & IRIS-HEP)
- Highly successful C++ training courses
  (3 this year from HSF and SIDIS, see S. Roiser's talk)
  - Continued curriculum development and sharing material
- Assembling a complete curriculum for training in HEP, using Carpentries templates
  - Following LHCb StarterKit initiative
- Recent paper published on HEP Software Training Challenges

Slide by G. Stewart

Key4HEP: Turnkey Software for Future Colliders

ICHEP conference paper

Challenge: HEP software stacks are wide and deep with many dependencies
Solution: Provide a stand-alone, validated full chain for detector design studies

- Ingredients:
  - Event data model, EDM4hep based on LCIO and FCC-EDM
  - DD4hep for geometry [AIDA2020]
  - Gaudi event processing framework (with *Marlin wrappers*)
  - Packaged and deployed using Spack
  - Fast (Delphes) and full (Geant4) simulation available
  - Documentation on https://key4hep.github.io/key4hep-doc/
- Contributions from ILC, CLIC, FCC and CEPC Communities

Allied with activities like HSF packaging group and discussions on best practices for copyright and licensing

Outline

- Why a HEP Software Foundation?
  - The landscape of software (and hardware)
  - Common challenges and common solutions
- The HSF's past, present and planned activities
  - HSF's Working groups
  - HSF collaborations
  - Recent LHCC review of common software
  - Connections to SwiftHep / Excalibur
- How to collaborate further

Google Summer of Code

- TLDR: remote internships in Open Source Organizations, funded by Google
  - Extremely competitive → great applicants and contributors (COVID notwithstanding)
  - Only 175 hours of coding for this year (roughly half that of previous years)
- Projects for CERN-HSF Organisation
  - 36 proposals, 27 slots granted, 25 students successfully completed the program (93%)
- Highlights [see all projects for 2021]
  - Automatic differentiation via CLAD, data compression using ML
  - GPU abstraction evaluation and open source GPU partitioning, Mesh Cloud
  - Data management plugins for Rucio with Science
  - UK-based common software projects like CONTUR, RIVET...

Documents and work towards LHCC reviews

- LHCC periodically reviews status and plans of HL-LHC software, and communicates outcomes / suggestions to Council [charge of the review]
  - HEP Software Foundation + WLCG liaisons help organise inputs from **common software**
    - This includes *Event Generators, Geant4, Data Organization & Management (DOMA) software components, ROOT, Data Science tools*
  - Experiments / WLCG also submitting separate documents on their own plans
- First review of this kind in May 2020
  - Bottom-up process, with overall editors & responsibles for sections (WG convenors)
  - Open to community feedback
  - Overall Common Software document from first review is on Zenodo, MC WG document is on arXiv:2004.13687 and published in Comput. Soft. Big. Sci
- Second review ongoing as we speak (1-5 November)
  - Agenda can be found on https://indico.cern.ch/event/1058274/

The document submitted to current LHCC review

- Structure of current document (under review):
  - Event Generators *[prepared by HSF]*
  - Detector Simulation
  - Physics Analysis
    - Foundational Tools and ROOT
    - Data Science Tools *[prepared by HSF]*
  - Data organization, Management and Access (DOMA)
- Preparatory meetings organized for each of the areas
  - Content reflects discussion recent developments in all those aspects
    - Members of HSF also participated to meetings with other editorial teams
  - Feedback from broader community requested before submission to LHCC
- Main points/issues from Event Generators & Data Science will be highlighted in the following slides

Chapter editors:

- Event Generators: The HSF Generator Working Group Convenors, Andrea Valassi, Efe Yazgan, Josh McFayden.
- Detector Simulation: The Geant4 Collaboration spokesperson and Deputy Spokesperson, Marc Verderi and Alberto Ribon.
- Foundational Tools and ROOT: The ROOT Project leaders, Axel Naumann and Loranzo Moneta.
- Data Science Tools: The HSF PyHEP Working Group Convenors, Jim Pivarski, Eduardo Rodrigues, Ben Krikler.
- DOMA: WLCG DOMA project leader Mario Lassnig, with Oliver Keeble.

Highlights from event generators document [arXiv:2109.14938]

- Survey of the community and of generators/tools most relevant for HL-LHC
  - Not a single project, no centralised decisions, no dedicated funding →
- Review poses standardized questions to generators, including:
  - Major physics and software updates / bottlenecks
    - Problem of negative event weights in matching still an issue
    - Standardized benchmarking & SW review useful for identifying bottlenecks
  - Plans for migration to GPUs / use of ML / generators in multithreaded frameworks
    - Interest and progress in porting to GPU & CPU vectorization (e.g. MG5aMC)
    - ML not yet widespread but could be useful
  - Is sufficient support available (funding & software)?
    - Funding & career opportunities problematic → loss of know-how
    - HSF (and others) contributed to create collaborations between theorists & SW engineers

Highlights from event generators document [arXiv:2109.14938]

- Opportunities for common projects where contributions would be welcome
  - Modular MC event generator framework
    - Some generators already interoperable, but could be improved
  - Help in porting MC generators to GPUs
  - Profiling and improving computing performance
  - Maintenance and development of “MC Tools”
    - E.g. PDF distribution and implementation (LHAPDF), common builds and packages (GENSER), quick analysis & tuning (Rivet & Professor)...
    - Work ongoing in SWIFT-HEP (C. Gutshow & others)
  - Some of the ore standalone projects could be taken on by RSEs (following example of Swift-Hep)

Highlights from Data Science Tools document

- Survey of the community and of tools
  - Transition to mixed software environment with different strategies and purposes
    - Foundational tools and frameworks (e.g. ROOT)
    - Smaller data-science-like tools exploiting cross-disciplinary components → Data Science
      - Mostly written in Python, but can wrap other languages
      - Some examples are at the IRIS-HEP AGC Tools Workshop starting tomorrow
      - Includes columnar analysis, declarative languages…
    - Productive and regular community discussions under the umbrella of PyHEP
- Challenges:
  - Interoperability of the tools in the entire “ecosystem”
  - Findability of individual tools (contributing vs rewriting)

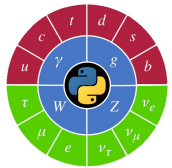

https://scikit-hep.org

Practical points for further collaboration with HSF (beyond what mentioned in LHCC documents)

Slide by G. Stewart

International R&D collaborations

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 USA
  - Analysis systems, innovative algorithms, DOMA
- ErUM-DAT, Helmholtz Institute DE
  - Heterogeneous computing and virtualized environments, machine learning for reconstruction and simulation
- EP R&D, CERN
  - Turnkey software systems, faster simulation, track and calo reconstruction, efficient analysis
- HEP-CCE, DOE USA
  - Portable Parallelization Strategies, I/O Strategy on HPC, Event generators
- AIDAInnova, European Commission EU
  - Turnkey software, track reconstruction, particle flow, ML simulation
- SWIFT-HEP, STFC and ExCALIBUR-HEP, UKRI UK
  - Exascale data management, Event generators, detector simulation on GPUs, FPGA tracking for HLT

Letters of support [see dedicated page on website]

- HSF writes letter of support for projects/proposals aligned with core activities
  - Projects = multi-institute collaborations for software-related activities
  - Recent discussion: could also write letters stating we will collaborate with individual proponents, if proposed activity aligned with Community Whitepaper [see minutes]
  - Procedure in a nutshell (but see minutes/webpage):
    - *Discuss your with the relevant HSF conveners to understand how activities can be aligned*
    - *Write public project summary and clear statement of intent to work within HSF, send request to HSF coord.*
    - *Process needs >> 2 weeks margin so there is enough time to discuss them in HSF*
- What the HSF would like in exchange
  - Nothing! The projects will advance software & computing in HEP in general
  - What would be nice to make progress together:
    - Involvement and presentations at the relevant working group
    - If you deem your project can *identify* with HSF activities, let us know

Current HSF WG Conveners

Event Generators
- *Ben Morgan* (ATLAS+Geant4)
- *Krzysztof Genser* (DUNE+Geant4)
- *Kevin Pedro* (CMS)

Detector Simulation
- *Josh McFayden* (ATLAS)
- *Efe Yazgan* (CMS)
- *Andrea Valassi* (CMS)

Reconstruction and Software Triggers
- *Dorothea vom Bruch* (LHCb)
- *Andreas Salzburger* (ATLAS)
- *David Lange* (CMS)

Data Analysis
- *Allison Hall* (CMS)
- *TJ Khoo* (ATLAS)
- *Nicole Skidmore* (LHCb)

Training
- *Michel Hernandez Villanueva* (Belle II)
- *Meirin Oan Evans* (ATLAS)
- *Sudhir Malik* (CMS)

PyHEP
- *Ben Krikler* (CMS, LZ)
- *Eduardo Rodrigues* (LHCb)
- *Jim Pivarski* (CMS)

Frameworks
- *Chris Jones* (CMS)
- *Kyle Knoepfel* (FNAL)
- *Attila Krasznahorkay* (ATLAS)

Tools and Packaging
- *Marc Paterno* (Fermilab, DUNE, new)
- *Mircho Rozodov* (CERN, CMS, new)
- *Serhan Mete* (ANL, ATLAS, continuing)

Blue: UK institute affiliation

- Very different wrt an experiment convener role
  - More of a *facilitator* role:
  - Organize WG meetings
  - Encourage exchange of ideas and techniques & stimulate discussion
  - Give platform for unsolved/work in progress problems
  - Document meetings with minutes and notes
  - Contribute to Common Software reviews
- Upcoming call for 2022
  - Some conveners will step down
  - Nominate, or self-nominate

Opportunities for collaboration: training

- Training can be cross-advertised / attended by organized by local institutes can be advertised via HSF
  - Get in touch with the HSF training WG conveners (emails on WG website)
- Very successful GPU workshop held by Manchester/Lancaster/Sheffield in 2020
- More training work ongoing joint between SIDIS and HSF, focused on accelerators
  - See S. Roiser's talk tomorrow

Other opportunities for collaboration

- **Join the meetings**: voices from the community are how bottom-up activities develop
- Software Institute for Data Intensive Science (**SIDIS**)
  - COST Action just submitted → see Stefan Roiser's presentation tomorrow
- Contribute to **common (software) projects** with HSF involvement
  - Examples from personal ex-convenor experience: TrackX, ACTS
- Near future: **ESCAPE / EOSC-Future Test Science Projects**
  - DM-TSP: End-to-end workflows for dark matter experiments to help building the European Open Science Cloud
  - Work just started, discussion with HSF (especially Analysis WG) will happen early next year
  - See more information at this link and in these proceedings

Slide by G. Stewart

Conclusions

- 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 has been fostering much more a shared vision
  - This encourages diverse R&D!
  - There is now real success in attracting funding to this area
  - Recognised links to other main players (WLCG, LHCC, IRIS-HEP, EPPSU, SWIFT-HEP, Experiments)
- There is a lot of work we need to do for upgrades and future projects
  - HSF offers an excellent place to present work, discuss successes (and disappointments!) and to help lead community activities
    - WG convener nominations welcome for 2022

Thanks for your attention!

With some overlap with HSF (we received a letter of support):

about the SMARTHEP ITN

about machine learning & hybrid architectures for real-time analysis,

Host institute: the University of Manchester

Backup slides

Slide by G. Stewart

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
- Workshops and Annual meetings
  - Not quite sure what our 2022 plans are, but we will do something!
- Propose a new activity area
  - The HSF is there to help gather interest

HEP Software Foundation

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

- Differentiable Computing
- Season of Docs
- Google Summer of Code
- intelligent Data Delivery Service
- Licensing
- Quantum Computing
- Reviews
- Visualisation

Slide by G. Stewart

Common challenges

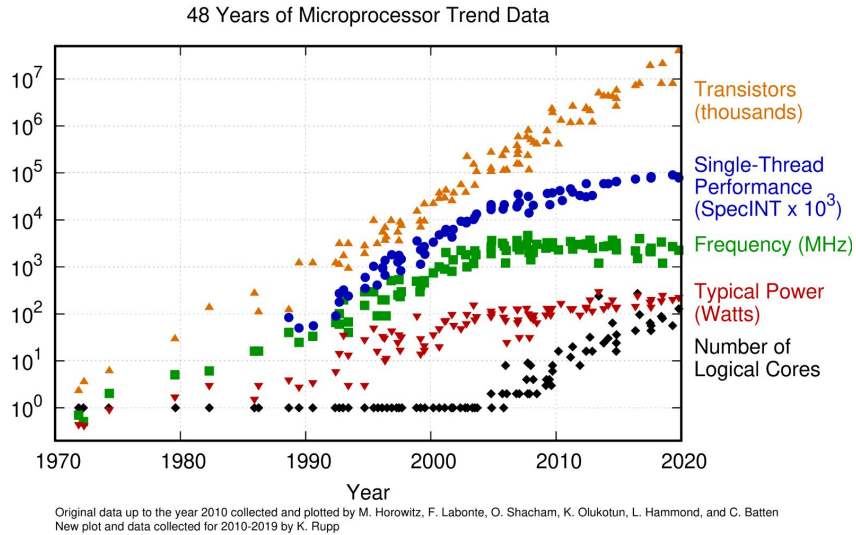

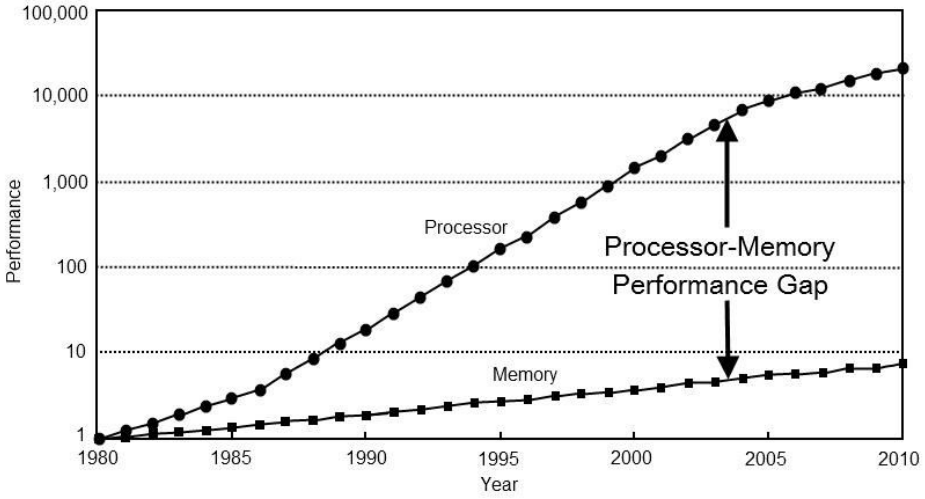

Increase in transistor density not as fast as increase in demands

Memory access times not keeping up with processor performance

Evolution of hardware: can we exploit it, and how?

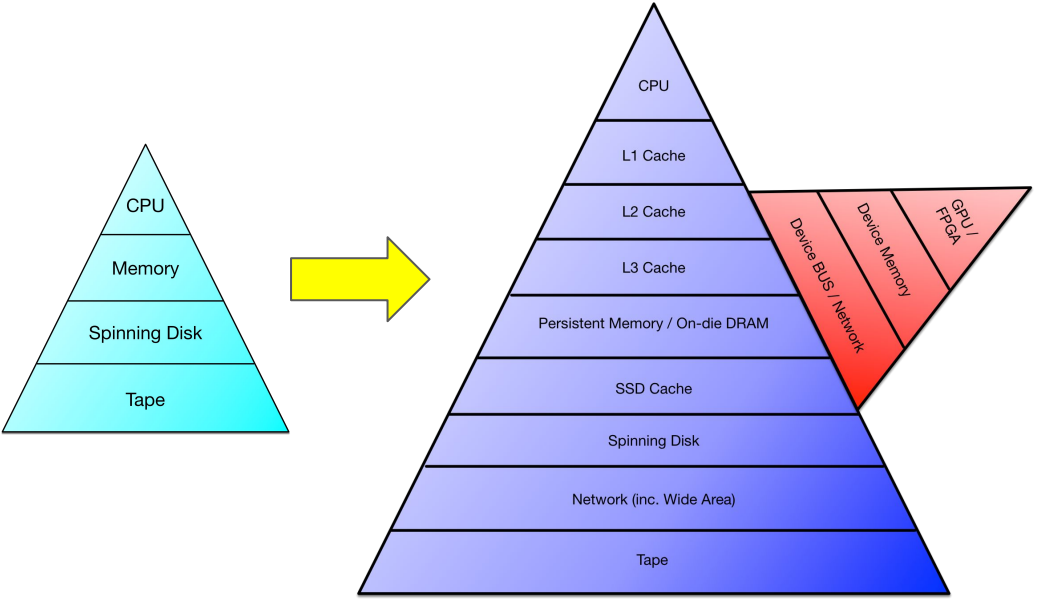

Diagrams by G. Stewart

High Performance Computing / commercial computing accelerators have recently become available → **what do we need** to use them effectively in HEP?

https://excalibur.ac.uk

Exascale Computing ALgorithms & Infrastructures Benefiting UK Research (ExCALIBUR)

ExCALIBUR is a UK research programme that aims to deliver the next generation of high-performance simulation software for the highest-priority fields in UK research. It started in October 2019 and will run through until March 2025, redesigning high priority computer codes and algorithms to meet the demands of both advancing technology and UK research.