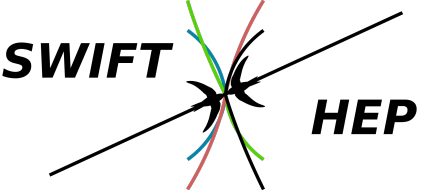

ExaTEPP

The SWIFT-HEP project and UK ~~HPC~~ DRI

23 July 2024

Outline

History and context

SWIFT-HEP and ExaTEPP

A few comments on Digital Research Infrastructure

Outline

History and context

SWIFT-HEP and ExaTEPP

A few comments on Digital Research Infrastructure

UK HEP computing infrastructure

- 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 (Smaller experiments, astronomy, facilities)

- DiRAC: UK high performance computing facility supporting STFC theory users

Physical hardware is only part of the expense. We should remember about the staff effort to support the infrastructure!

UKRI Digital Research Infrastructure:

- Large investments in AI machines, e.g. Isambard AI (which is useful for AI!)
  - … . but mainly GPU based and so not applicable to HEP production computing
- Not necessarily in sync with the needs from the PP community
- A large HPC may not suit PP requirements

Timeline and history

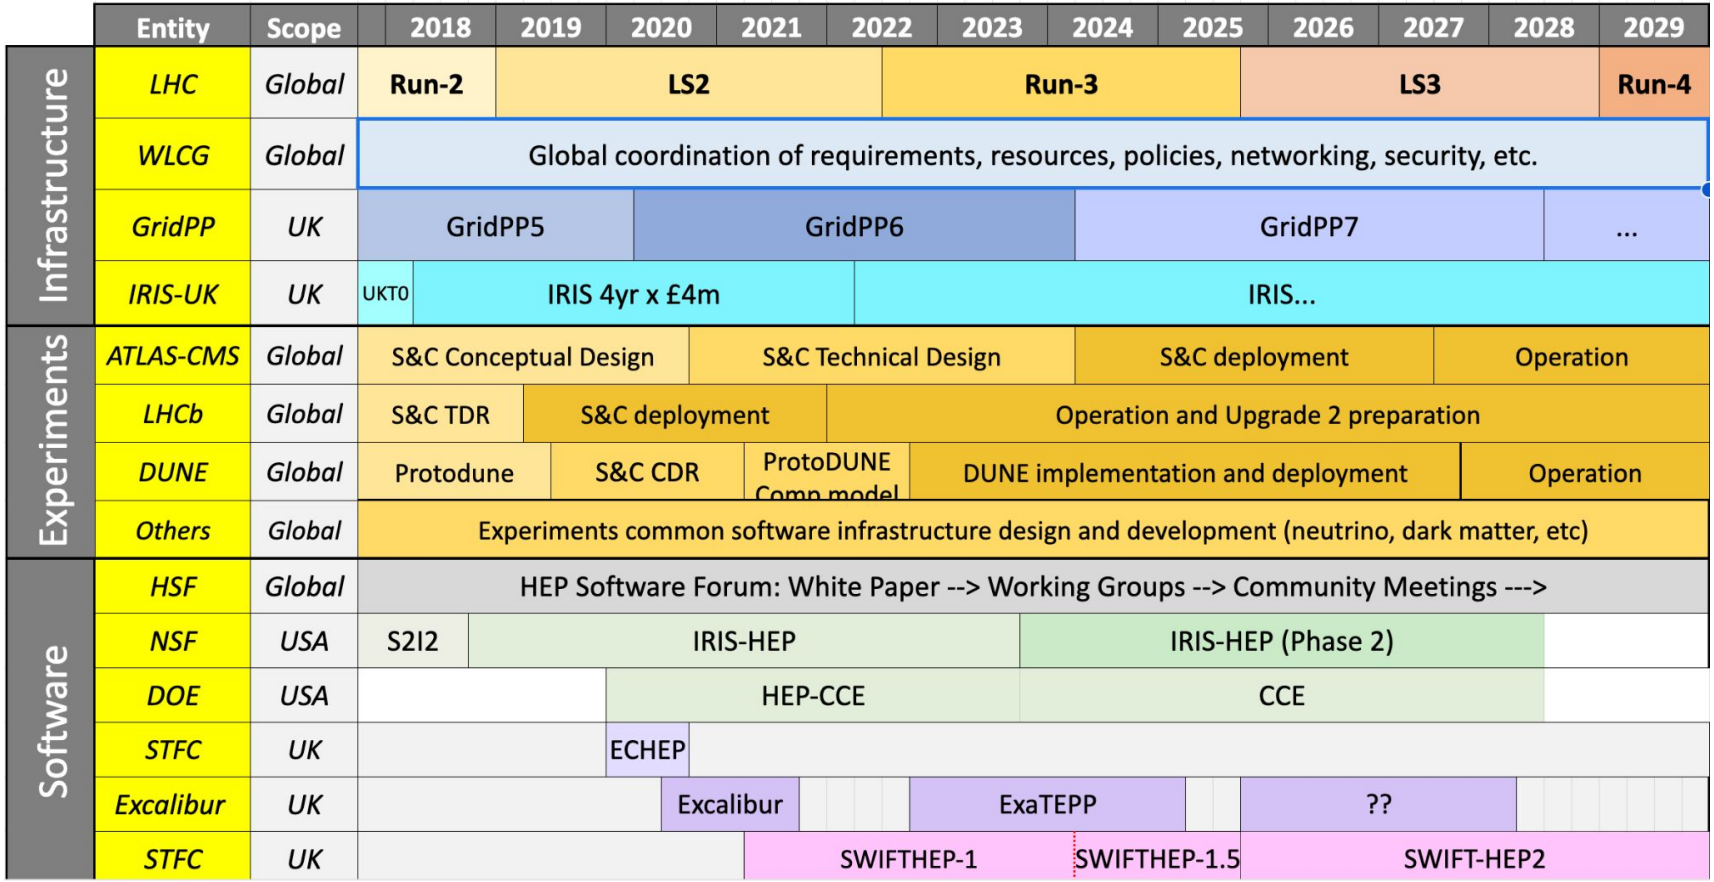

Multi-experiment software development

The UK Particle Physics Technology Advisory Panel (PPTAP) (link) 2021

[34] Likewise within software and computing the UK has significant leadership in a significant number of important areas, including in exploitation of computing accelerators, exploitation of low power compute units, computing operations, enabling software and computing, reconstruction algorithms, software framework development, development of cross- experiment development tools, use of HPC and development of collision simulation/generation programmes.

Layer	Domain	Experiment 1	Experiment 2	Experiment 3
6	Physicists	Analysis code	...	...
5	Experiment Physicists programmer and software engineers	Analysis framework. Simulation, Reconstruction, Calibration Code	...	...
4	Experiment Software Engineers	Software Frameworks	...	...
3	Common Software HSF / SWIFT-HEP	Common software components (Data management, Generators, Geant4, Accelerator integration)		
2	GridPP / WLCG	Middleware infrastructure for Distributed Computing		
1	GridPP / WLCG	Physical Hardware		

Moving Software down the stack

Some results from community surveys

Particle Physics Technology Advisory Panel (PPTAP) 2021
(RTP= Research Technical Professional)

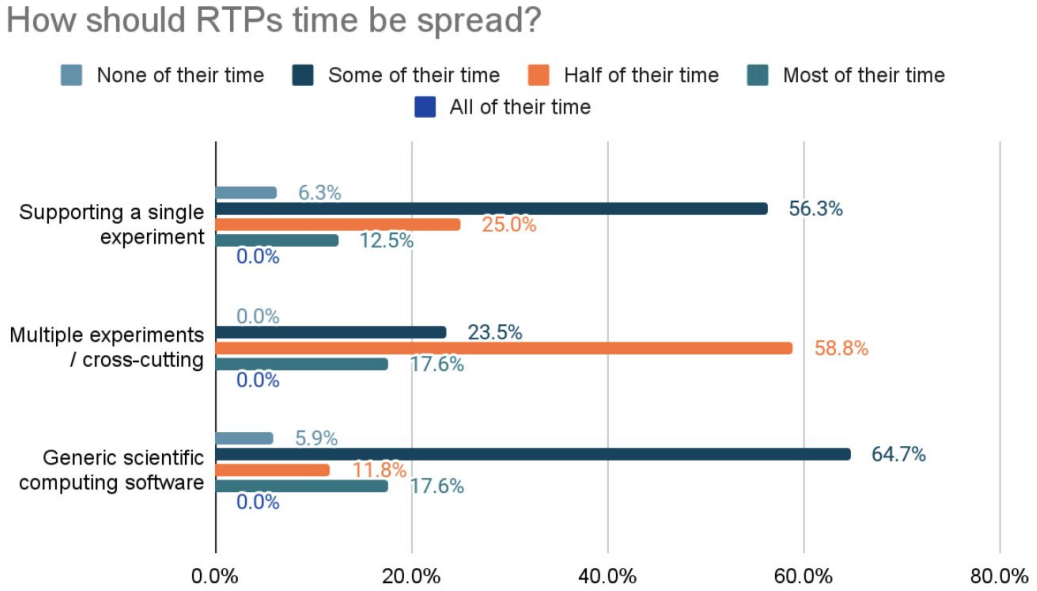

AI-generated particle physics community meeting

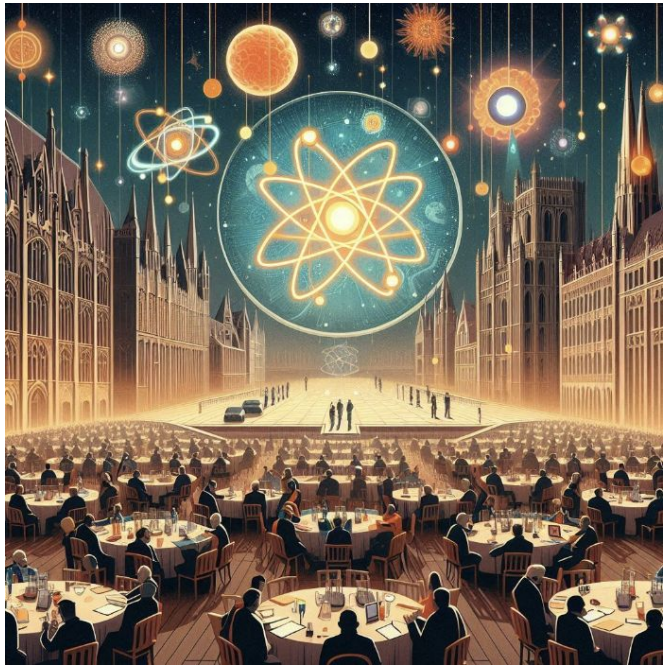

Particle Physics Advisory Panel (PPAP) 2024

Q: Key infrastructure requirements
Most people said computing resources are adequate

More Research Software Engineers are needed

Outline

History and context

SWIFT-HEP and ExaTEPP

A few comments on Digital Research Infrastructure

Current UK HEP software projects

SWIFT-HEP (5 FTE) Software Infrastructure for HEP	ExaTEPP (3 FTE) Exascale software for PP
Sheffield (PI)	Sheffield (PI)
Bristol	Edinburgh
Imperial College	Swansea
Manchester	Warwick
Warwick	Hartree Centre (Daresbury Lab)
UCL	
Rutherford Lab (PPD and SCD divisions)	

What does SWIFT-HEP do?

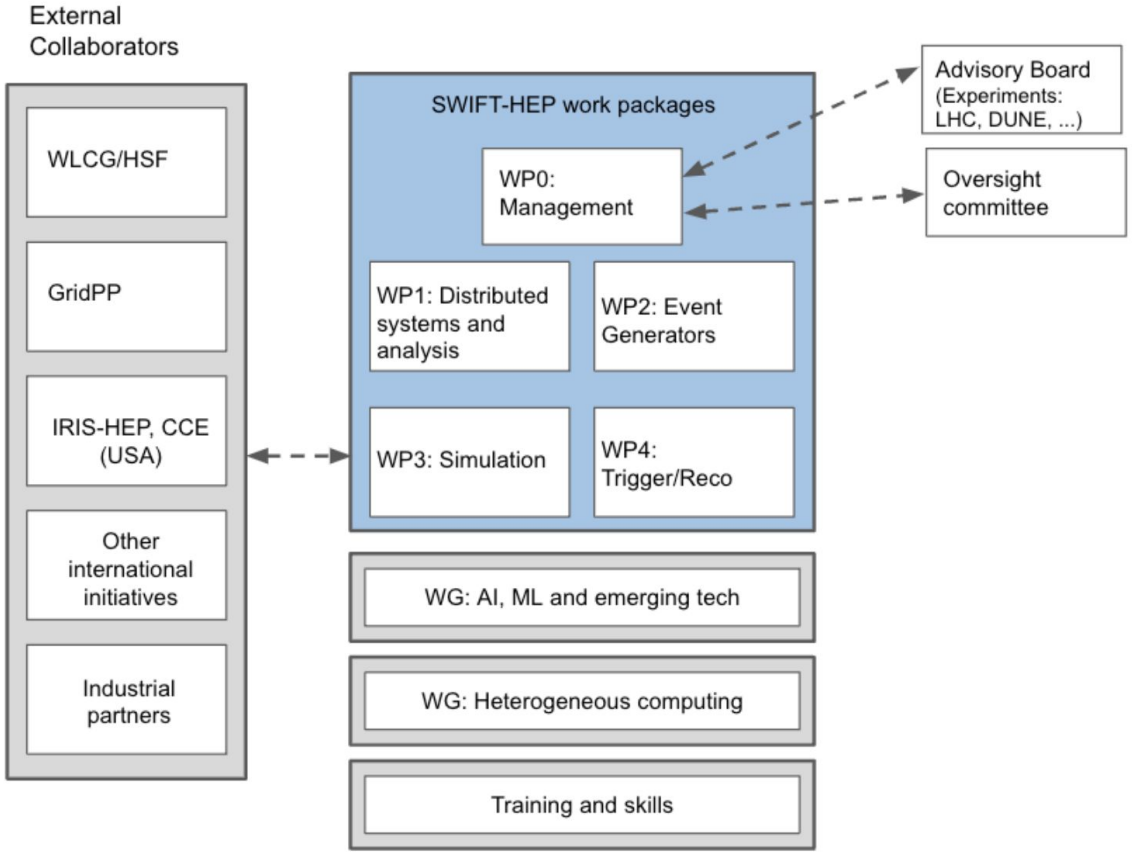

5 FTE. Funded from September 2020 until 2025 (so far) - £2.5M

Part fund people at Universities and national Labs to match with other projects

Overall about 10 people funded directly by SWIFT-HEP

Slow ramp up due to recruitment difficulties (especially WP1 and WP4)

Software improvements. An example from event generators

- Event generators are a key ingredient at the LHC
  - Higher precision generators needed for high precision physics
  - Not just the LHC, also neutrino physics (e.g. Genie)

UCL, Glasgow, Durham
Chris Gutschow et al,
EPJ C82 (2022) 12

- Same physics results with less computing (more info)
  - Example from Sherpa (possible thanks to SWIFT-HEP funding)

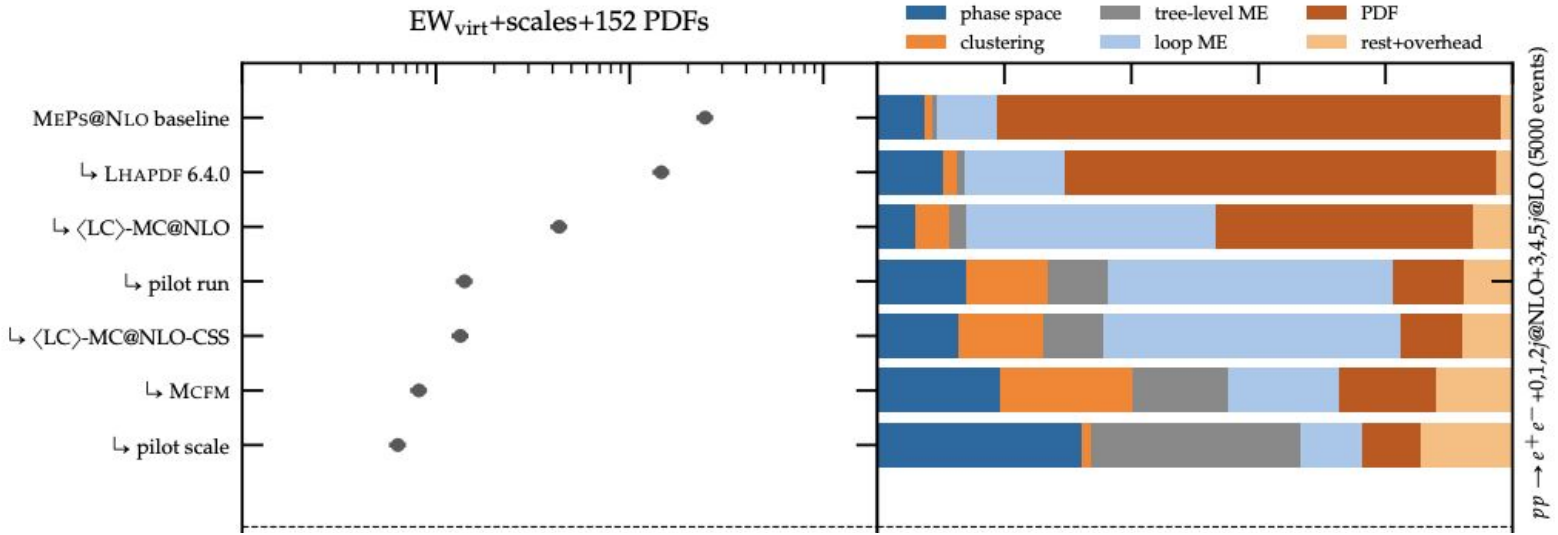

Simulation and GPUs (SWIFT-HEP + ExaTEPP)

Warwick, Sheffield, Manchester

A significant fraction of WLCG CPU time is spent simulating events (and a lot of it is EM showers)

Great collaboration with **Celeritas** (and AdePT) - **See next talk**
Development of common interfaces CelerAdePT (Ben Morgan)
Recent hackathon with Seth and Elliott

Deliverable:
API, data structures, and workflows for implementing GPU-based scoring in AdePT/Celeritas
Working in collaboration with ATLAS (Hackathon in October)

Work initiated in evaluating Mitsuba for optical photons

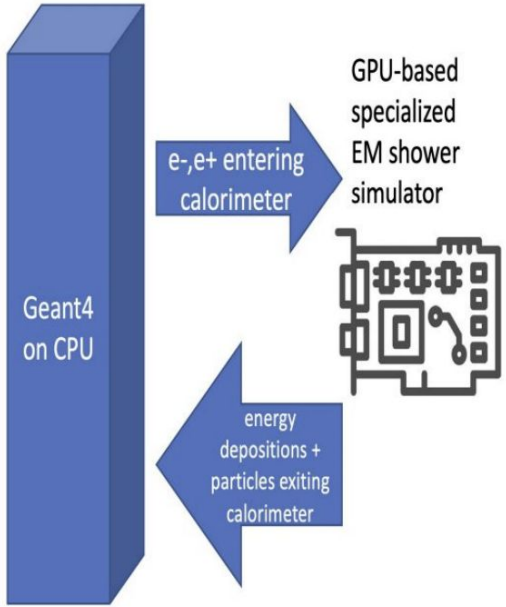

Trigger/Reco work package

Run HLT software on FPGAs (S. Harper et al, mostly for CMS ECAL now) RAL, Warwick, Sussex
Very early R&D, requires mastering HLS coding

Traccc: integrating the Alpaka framework (S. Martin-Haugh, R. Cross, et al)
Cuda vs Alpaka-cuda are similar
HIP still under development

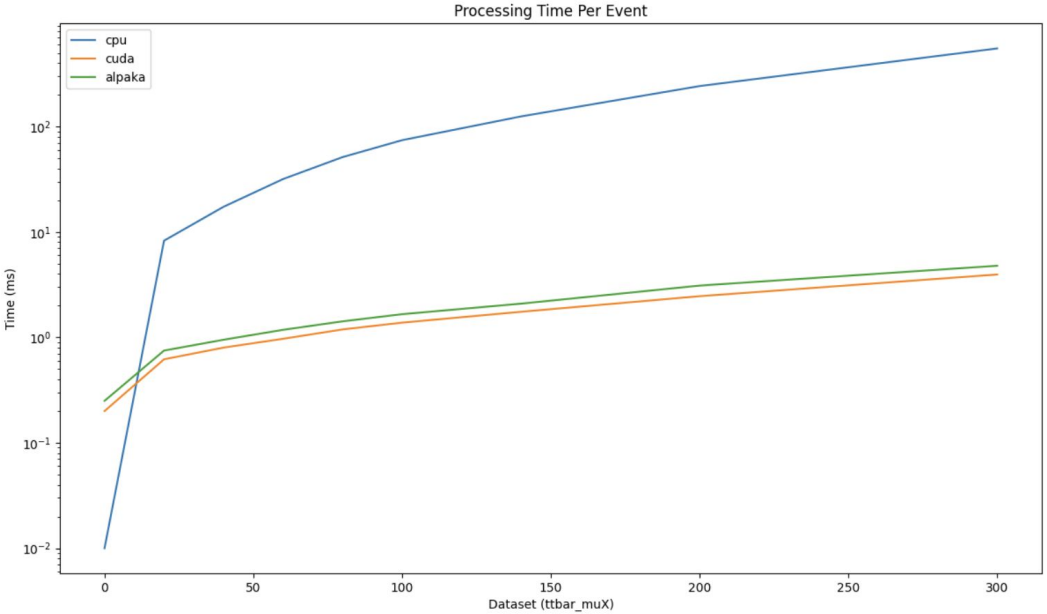

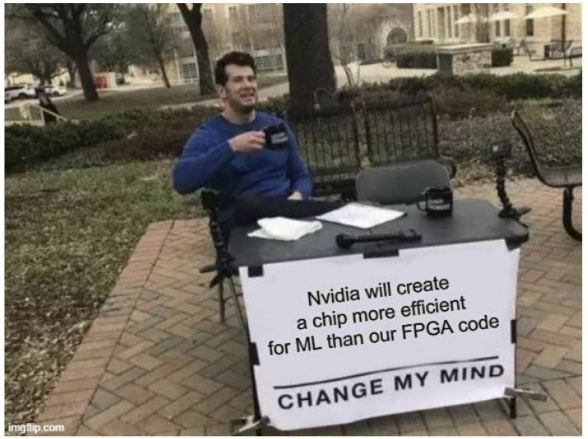

Image from Sam Harper

Analysis grand challenge on UK infrastructure

Bristol, RAL, Imperial, Brunel

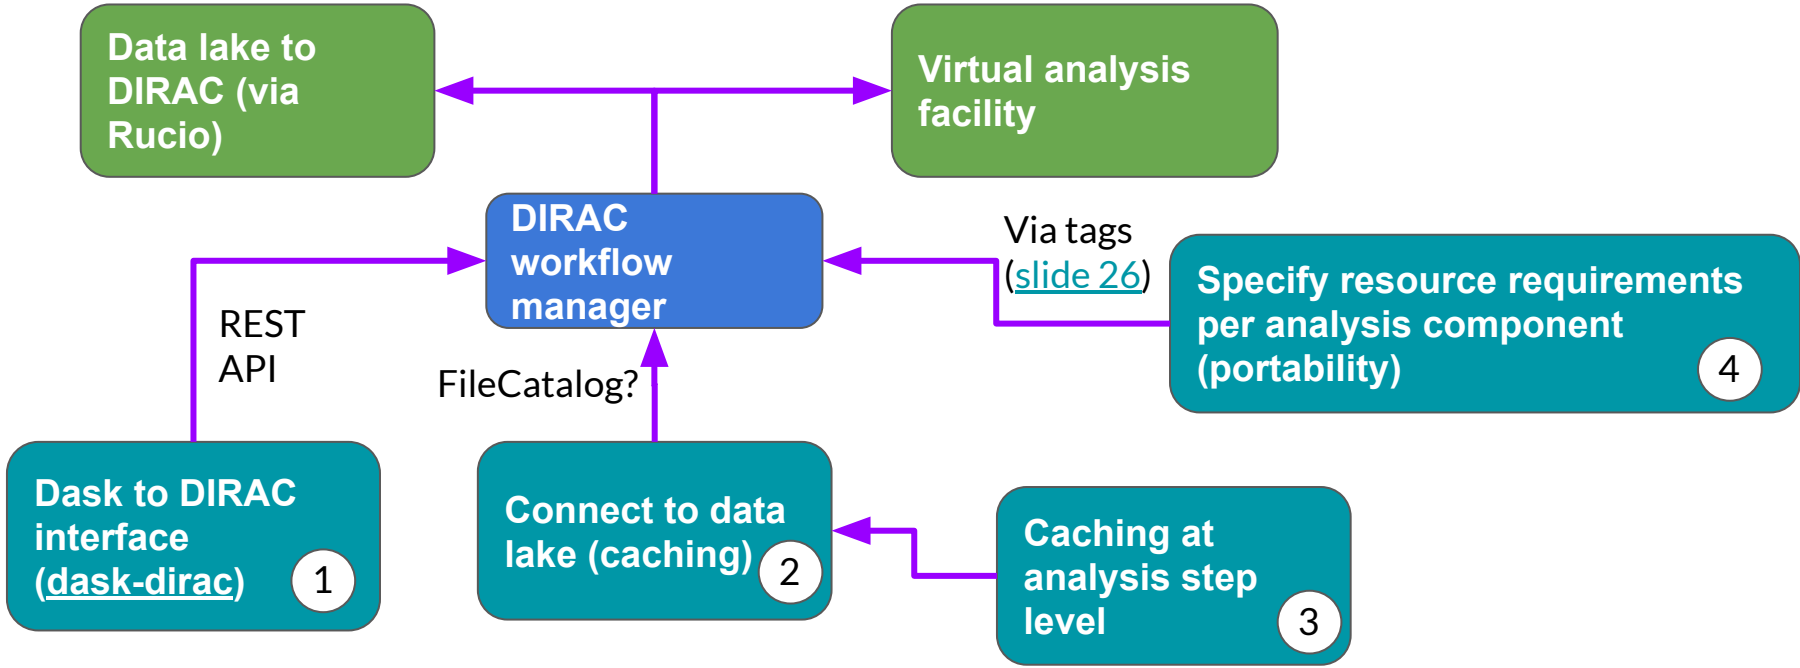

Closest example of what we want to achieve: Dask-based Distributed Analysis Facility (kubernetes slides)

ExaTEPP example: Quantum field theory on a lattice

Solve lattice field theories on a 4D lattice of points

Lots of matrix algebra. Large lattices require large GPU calculations across multiple nodes (**co-design as part of DiRAC**)

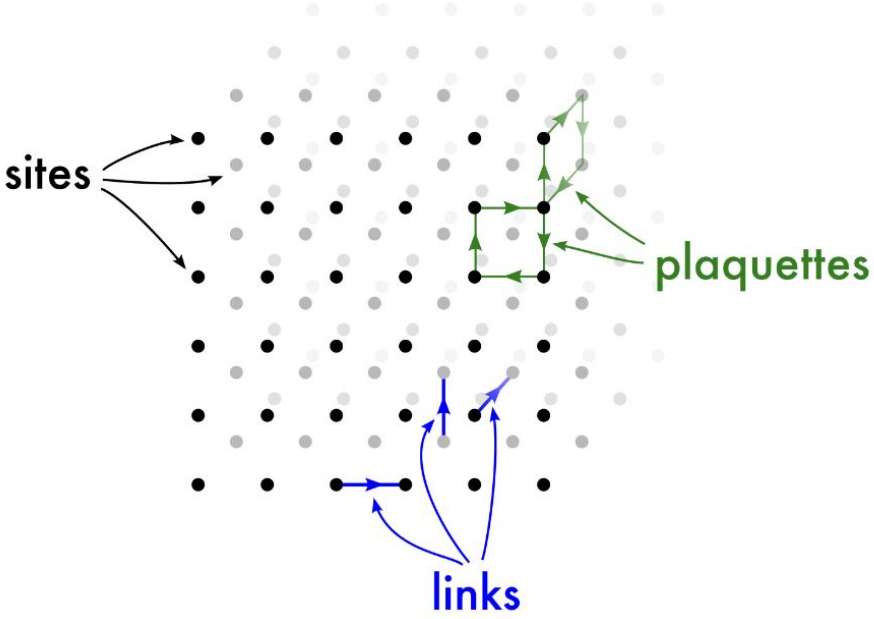

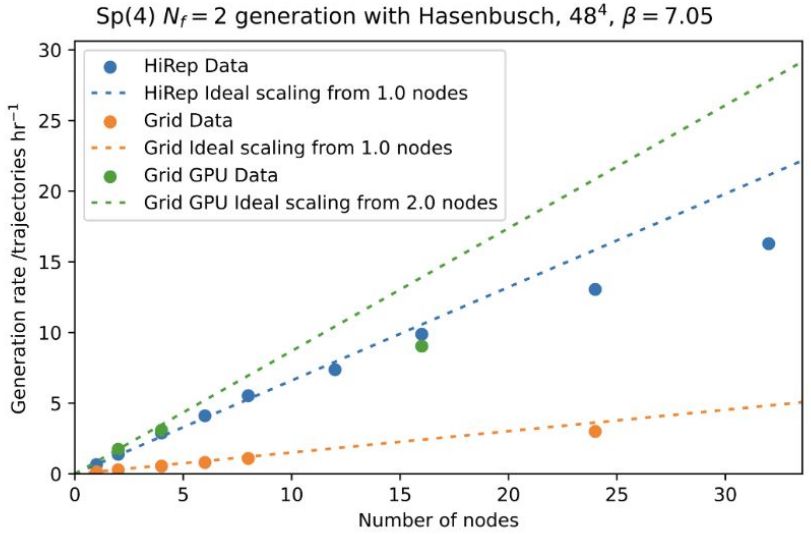

The DiRAC HPC Facility

HPC for theoretical astrophysics, cosmology, particle & nuclear physics in the UK

Memory Intensive "COSMA8" (Durham)

DELL EMC

- 528 TB RAM
- Large-scale cosmological simulations

Extreme Scaling "Tursa" (Edinburgh)

Atos

- GPU-based system
- Large lattice-QCD simulations

Data Intensive "DiaL" (Leicester)

- Heterogeneous architecture for complex simulation and modelling workflows

Hewlett Packard Enterprise

Data Intensive "CSD3" (Cambridge)

DELL EMC

- Heterogeneous architecture for complex simulation and modelling workflows

Project Office (UCL)

Ongoing phase. SWIFT-HEP 1.5 (April 2024 - September 2025)

D1.1	SSD storage to accelerate Data Transfer	Month 1-6
D1.2	SSD storage to accelerate Data processing	Month 6-12
D1.3	Benchmarking new workflows on different QoS	Month 6-18
D1.4	Deployment of an analysis facility on GridPP infrastr	Month 6-18
D2.7	Review of performance bottlenecks in EvGen Including interface formats (e.g. LHE, HepMC)	Month 1-4
D2.8	Develop a flexible, efficient and compressed event format and generator interface code release (e.g. HepMC)	Month 4-18
D3.4	Integration of GPU-based EM transport and scoring in ATLAS simulation with performance metrics	Month 1-18
D3.5	API, data structures, and workflows for implementing GPU-based scoring in AdePT/Celeritas	Month 1-18
D3.6	Customisable optical photon propagation with Mitsuba fully integrated in Geant4.	Month 1-18
D4.7	Report on the performance of the ECAL algorithm using HLS language	Month 1-18
D4.8	Full functionality of traccc within ACTS using Alpaka	Month 1-12
D4.9	Prototype developments using SYCL/oneAPI	Month 12-18

Planning: SWIFT-HEP 2.0 (October 25 to ??)

We submitted a “Statement of Interest” at the end of April (3 pages)

- **3 scenarios**: minimum (5 FTE, 0.6M£/year), core (8 FTE 1M£/year), desired (16 FTE 2M£/year)
- Reviewed by Science Board in June
- We were then asked to define the minimum scenario for 2 years (Oct 2025 - Sep 2027)
- Recommendation to request funding as part of the Digital Research Infrastructure initiative **(DRI)** - See following slides

7 work packages (core/desired), 4 work packages (minimum)

- WP1 Distributed systems and analysis
- WP2 Event generators
- WP3 Simulation of particle interactions with matter
- WP4 Reconstruction algorithms
- WP5 Benchmarking and sustainability
- WP6 Software frameworks for future experiments
- WP7 Artificial intelligence tools and emerging technologies

AI work not very well defined. Starting a working group on AI (kick off in October)

Outline

History and context

SWIFT-HEP and ExaTEPP

A few comments on Digital Research Infrastructure

Digital Research Infrastructure - UKRI

The ExCALIBUR programme

A UK programme to develop Exascale software algorithms.
(Ending in March 2025, a new programme expected in October 25)
ExaTEPP: Exascale software for Theoretical and Experimental Particle Physics

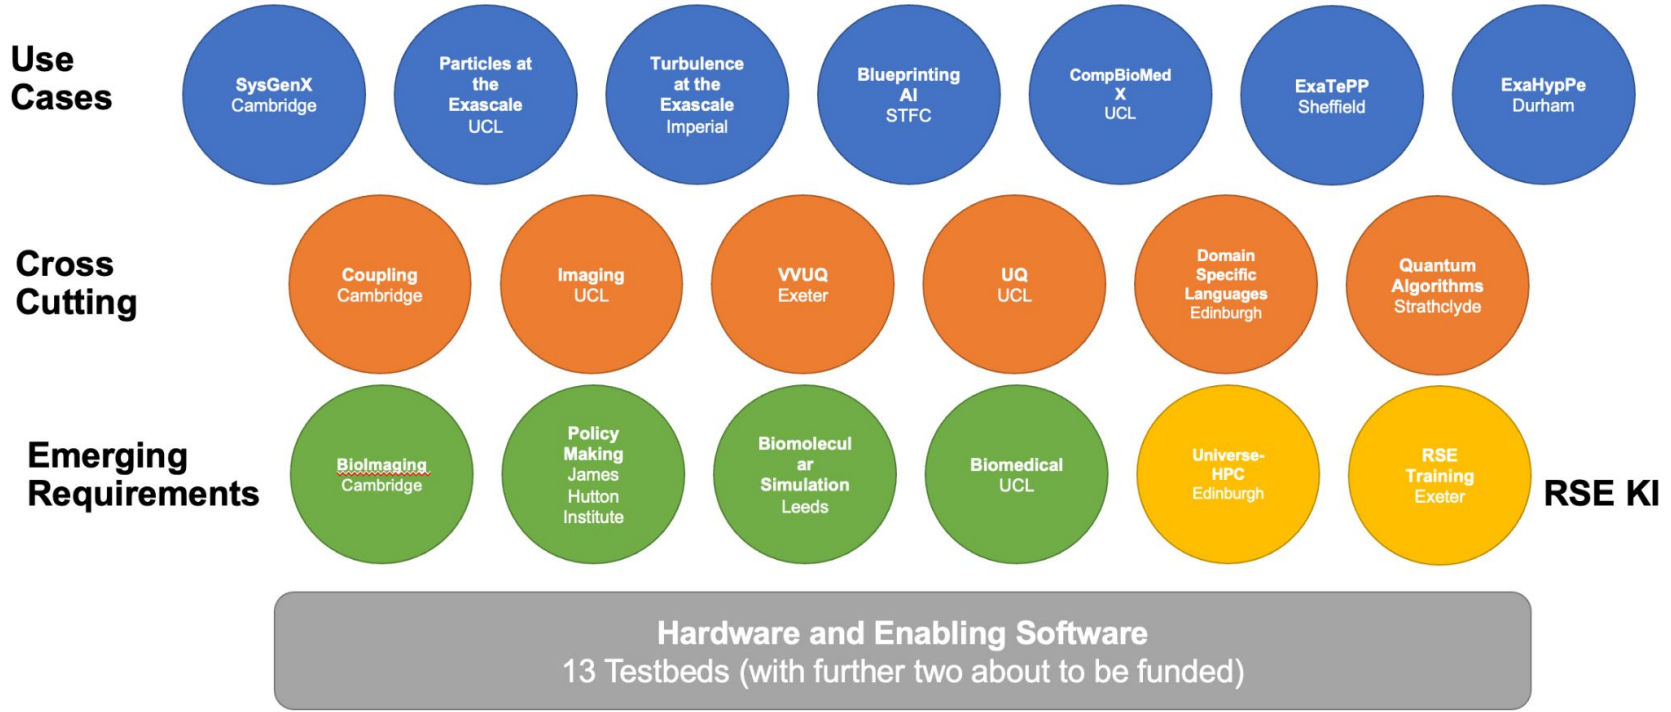

Things I would like to learn this week

New ideas on how to improve the way we organise things

- I like the “all-hands” concept (though it requires parallel sessions)
- Project management and how to monitor deliverables
- How to go from R&D to experiment deployment? Who is in charge?
- …

How do you define AI and how to introduce it in HEP workflows?

- In a useful and meaningful way
- In the light of large “AI machines” being commissioned

HPCs (or large scale computing) and their limitations and usage

- Interaction with other non-HEP communities
- Large HPCs being commissioned in the UK (e.g. IsambardAI in Bristol)

Approaches to training, skills, career development

Conclusions and lessons learned

It's all about people

- Recruiting and retaining staff is the most difficult part of the job
- We can only recruit integer people for longer term (possibly "open ended")
- People have careers and they have options outside HEP and research
- Research Software Engineers (RSE) pools are a good working model provided we can retain (and fund!) people for longer periods of time

The distributed model is our way of work

- Mixture of staff at universities, national labs (RAL PPD and SCD)

Money is always tight. There is no magic money tree

- HEP needs software and computing. This is now loud and clear. (But it's tensioned against other needs!)
- Lot of effort to write proposals, reviews, documents. Better project management?

"Thou shall use GPUs" (they are the solution to all problems, right?)

There are other communities out there

- With different problems and ways of working. We need to continue making our case