Software and Computing Challenges

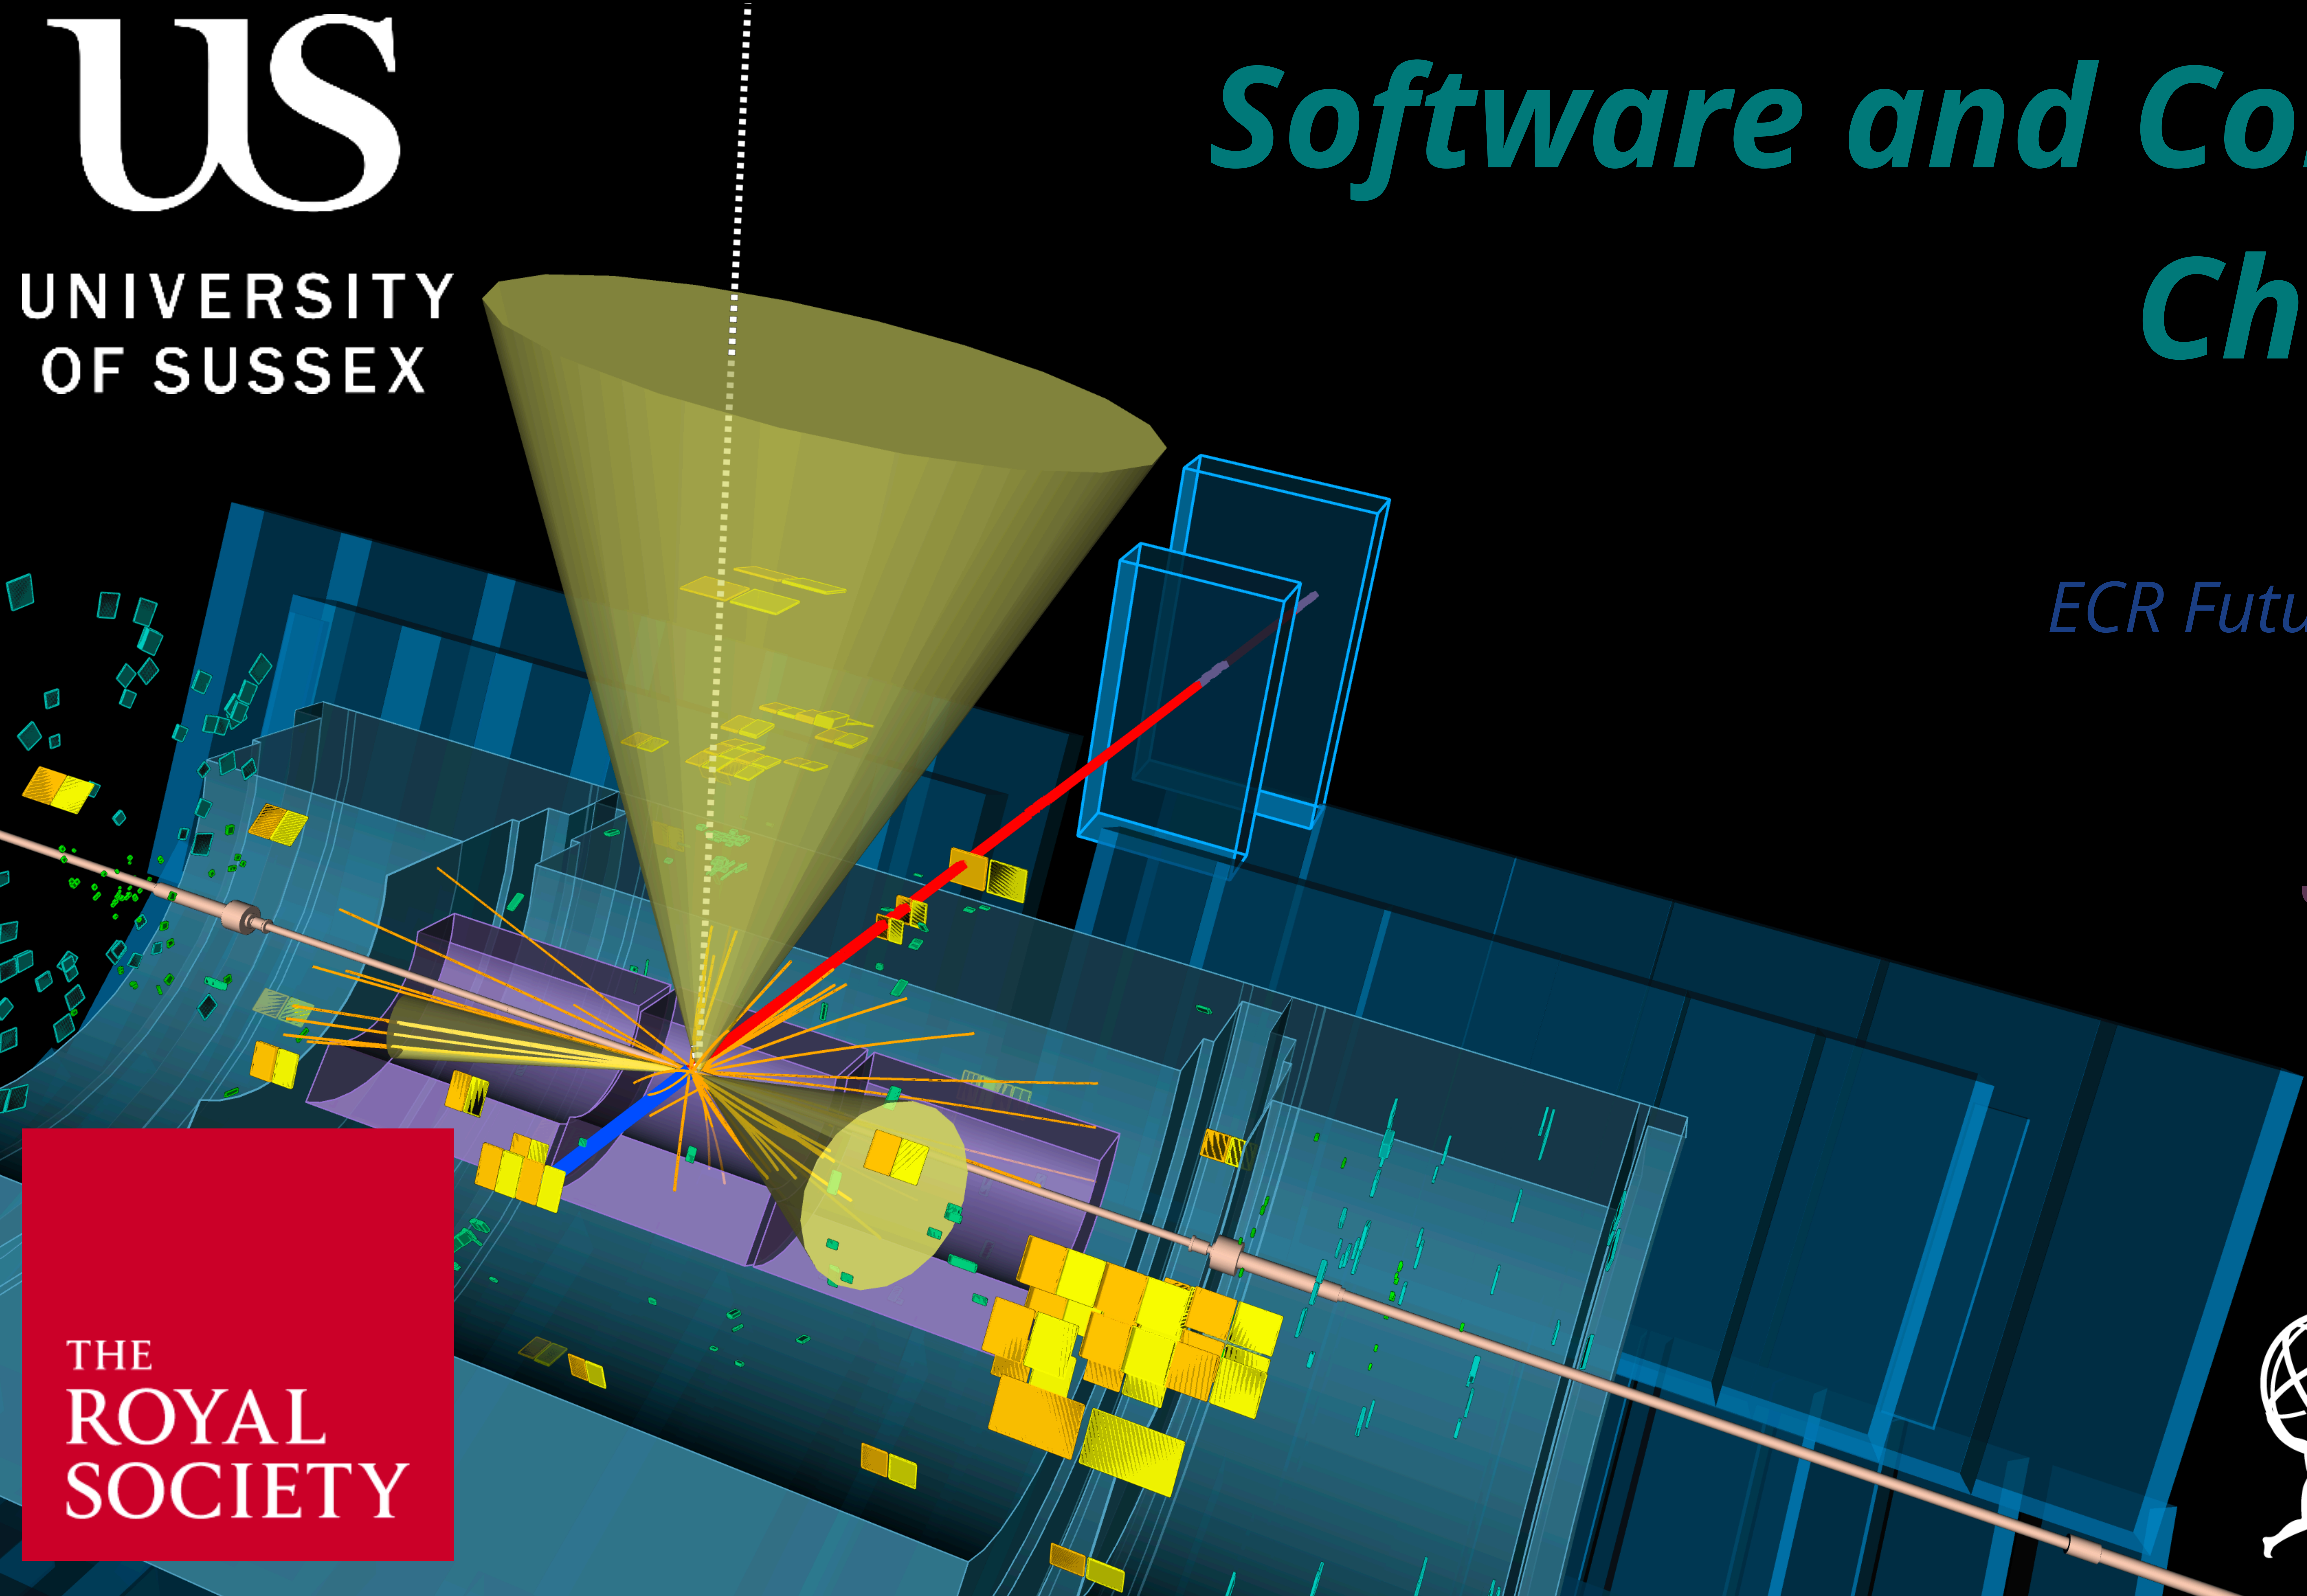

ECR Future Collider Forum
4/11/2022

Josh McFayden

@JoshMcFayden
cern.ch/mcfayden

UNIVERSITY OF SUSSEX

THE ROYAL SOCIETY

ATLAS EXPERIMENT

Hello! 👋

UNIVERSITY OF SUSSEX

- A brief introduction:
  - I'm based at the University of Sussex
  - I have been a member of **ATLAS** for 10+ years.
  - Also work on FASER and R&D for FASER2
  - Currently working on Top & Higgs measurements
  - Previously ATLAS Monte Carlo Production co-coordinator
  - Previously ATLAS Physics Modelling Group co-convener
  - Currently co-convener of the HEP Software Foundation Event Generators WG

Credits

- Slides heavily taken from
  - Graeme Stewart
    - https://indico.cern.ch/event/808335/contributions/3367988/attachments/1843865/3025660/eppsu-software-rd.pdf
    - https://docs.google.com/presentation/d/1lj9HkczcHJOjvTGGG2aCa7vv92NO9utRNxfVhlzaK0U/edit#slide=id.p
  - Caterina Doglioni
    - https://docs.google.com/presentation/d/1SRBNPFj-SxcnHutgKktbUv5KQu2vM_zeKopfuAKA8dY/edit#slide=id.gfac85f9076_0_111
  - Ian Bird
    - https://indico.cern.ch/event/361440/contributions/1775325/attachments/719506/987651/WLCG-FCC.pdf
  - EPSU
    - https://home.cern/sites/default/files/2020-06/2020%20Update%20European%20Strategy.pdf
    - https://home.cern/sites/default/files/2020-06/2020%20Deliberation%20Document%20European%20Strategy.pdf
  - Computational Frontier Report Contribution to Snowmass
    - https://arxiv.org/abs/2210.05822

Introduction

- To some extent the question is: what can we afford?
  - What physics can be done with the computing we can afford?
  - Iterative – evolves as technology and costs evolve
  - Extrapolating computing technology 20 years into the future is not obvious

- Costs
  - CPU and computing itself
  - Storage
  - Networks
  - Compute facilities - Expensive. Is building new facilities still cost-effective?
  - Operations cost
  - Electricity - Was becoming more expensive even before the current energy crisis

How well do we estimate costs?

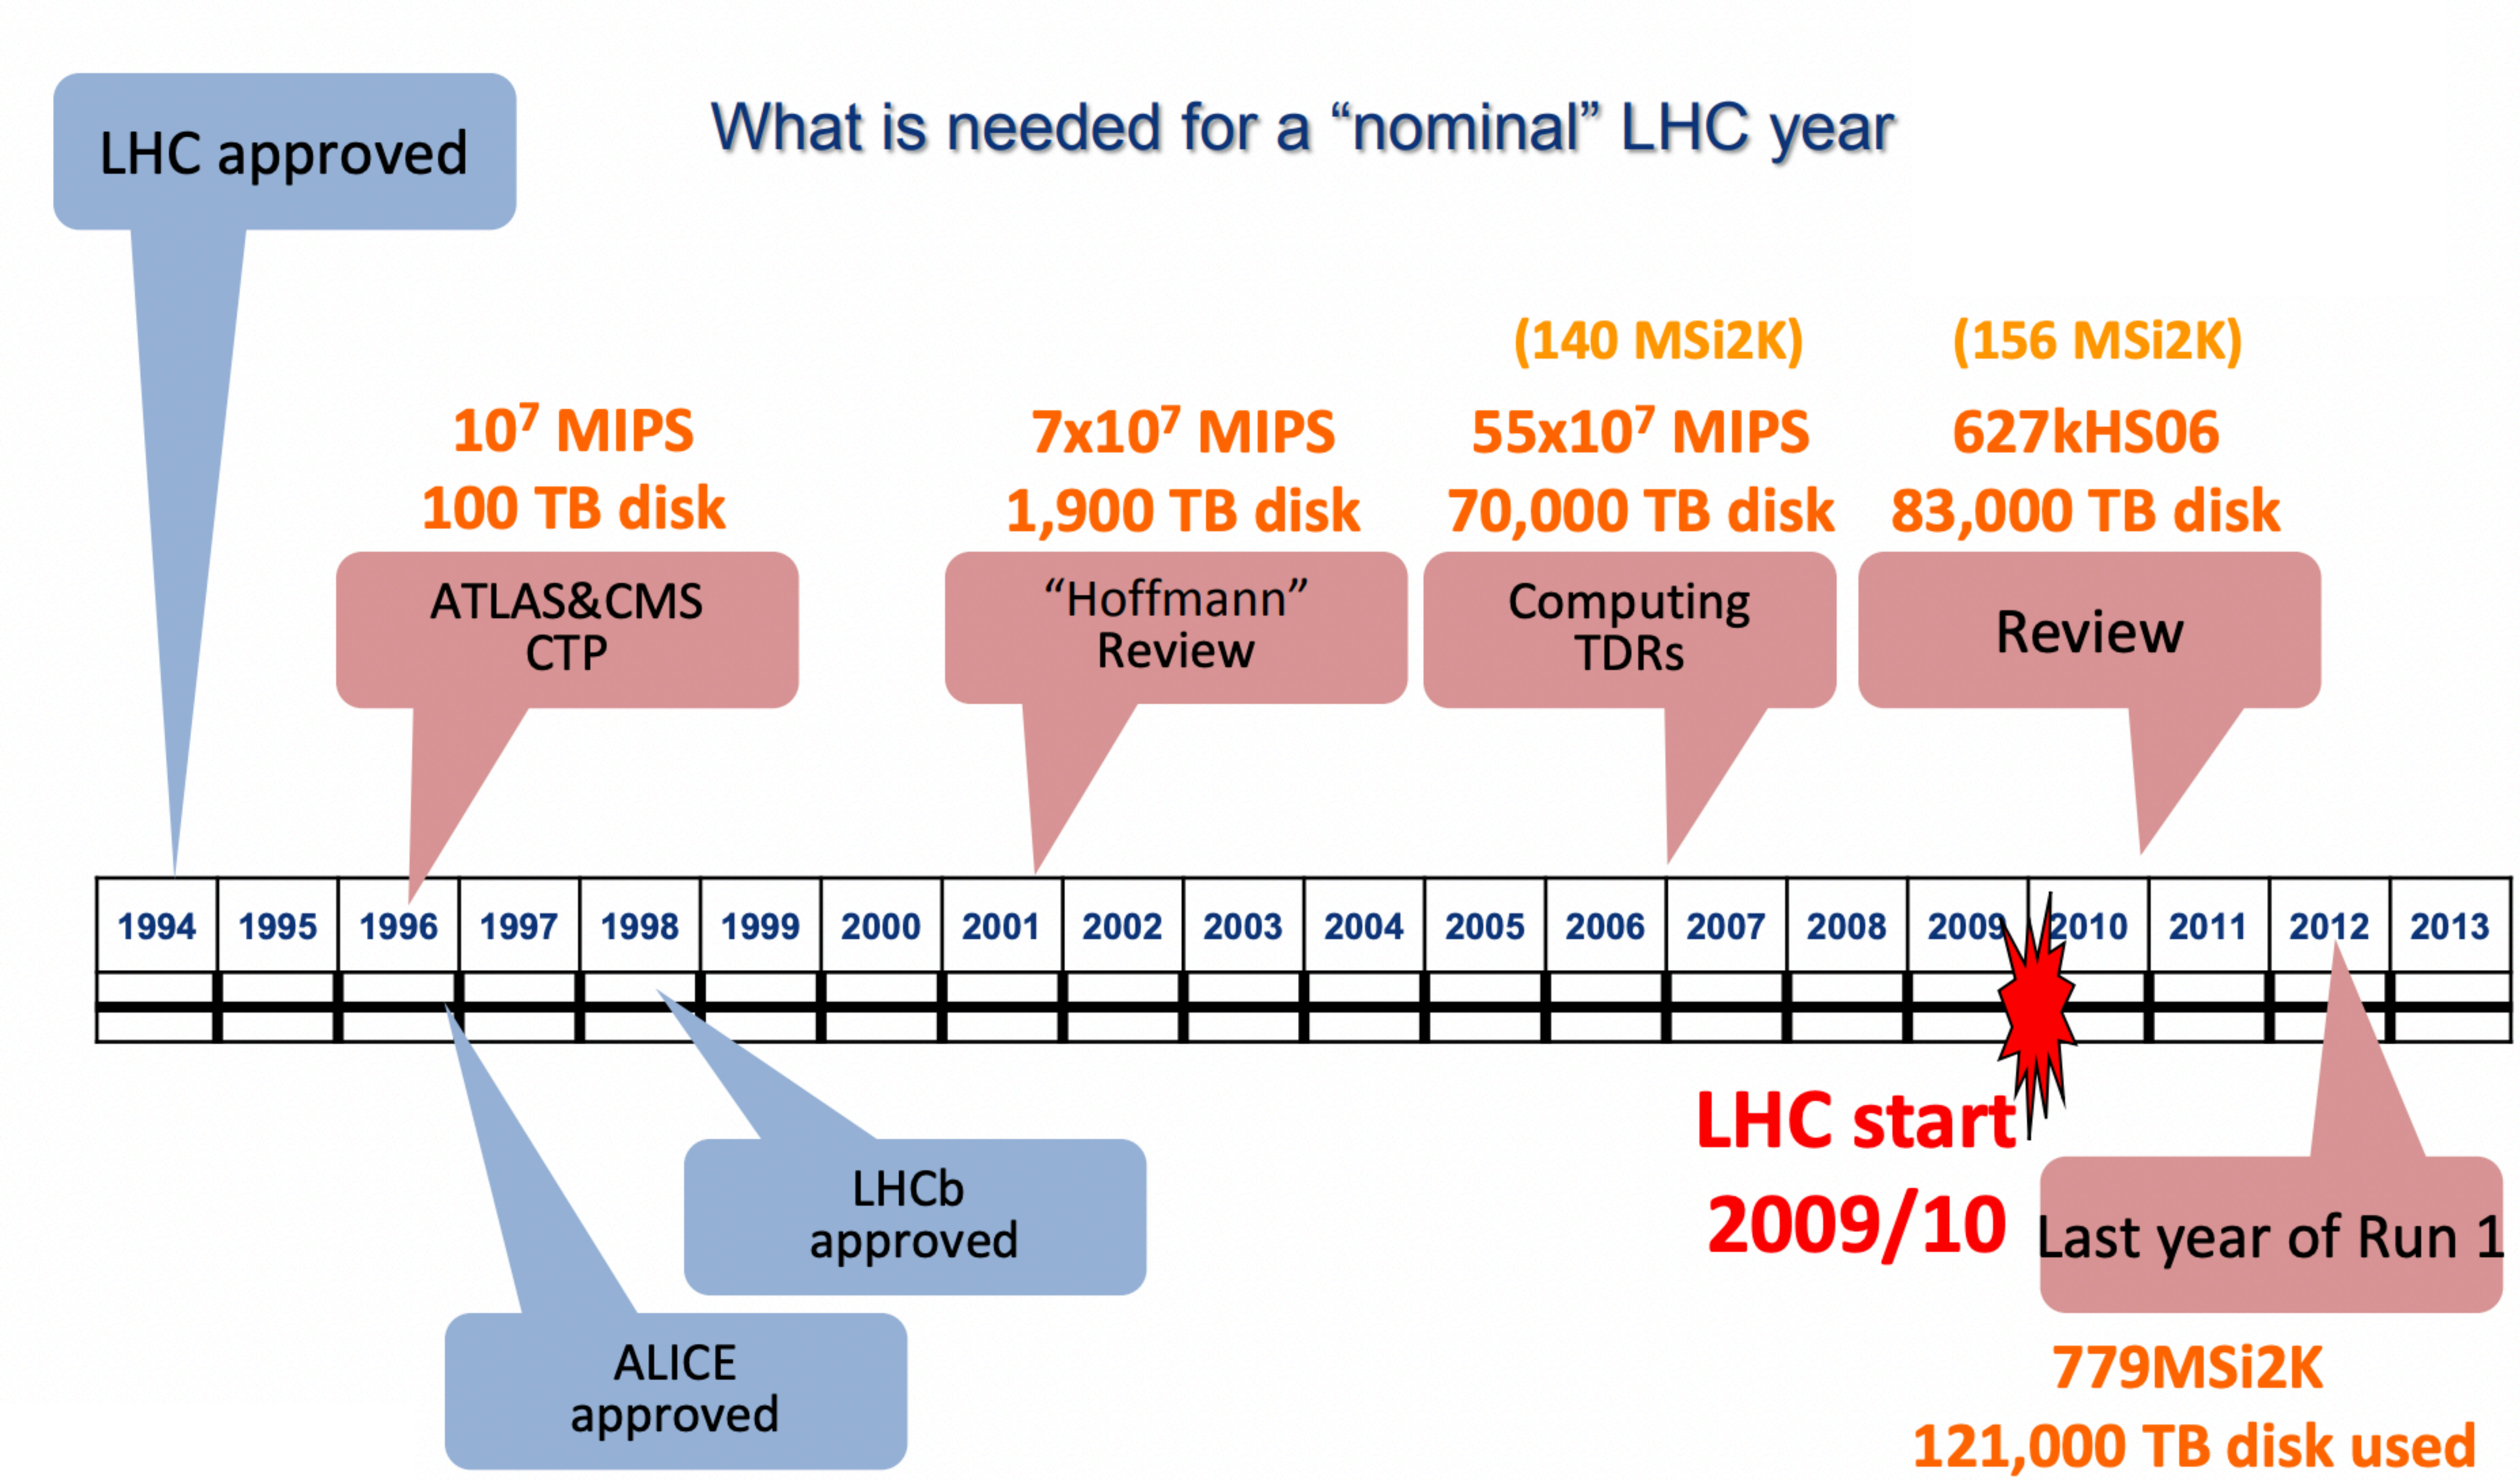

Ian Bird

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
- Some 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 & human effort
- With significant challenges ahead of us to support our developing physics programme

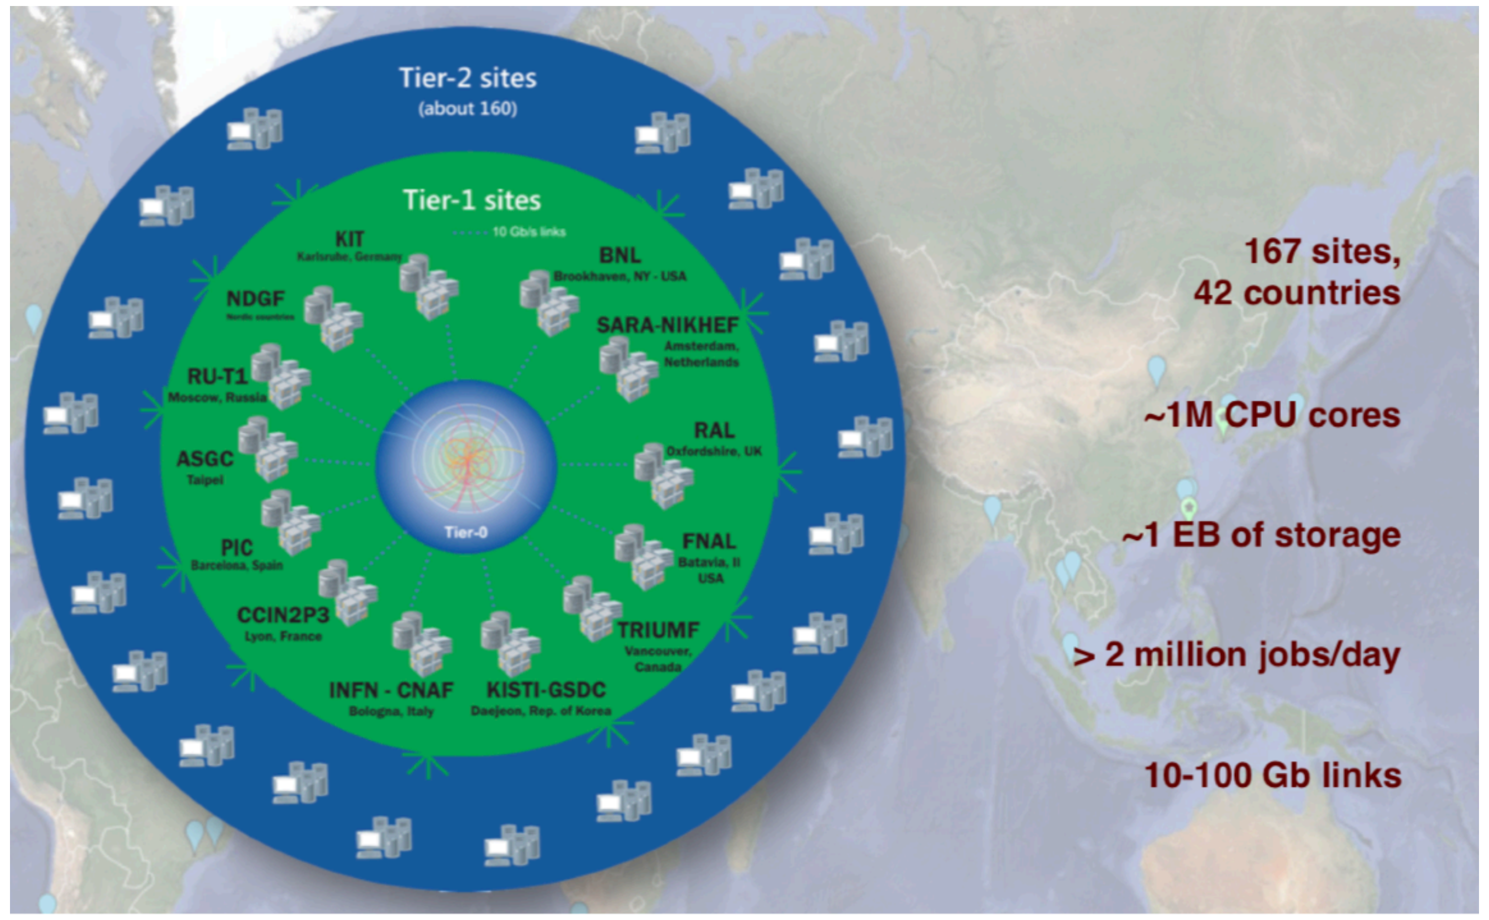

ROOT Data Analysis Framework

athena

ATLAS Experiment main repository for Athena code

DD4hep

Geant 4

HEP.TrkX

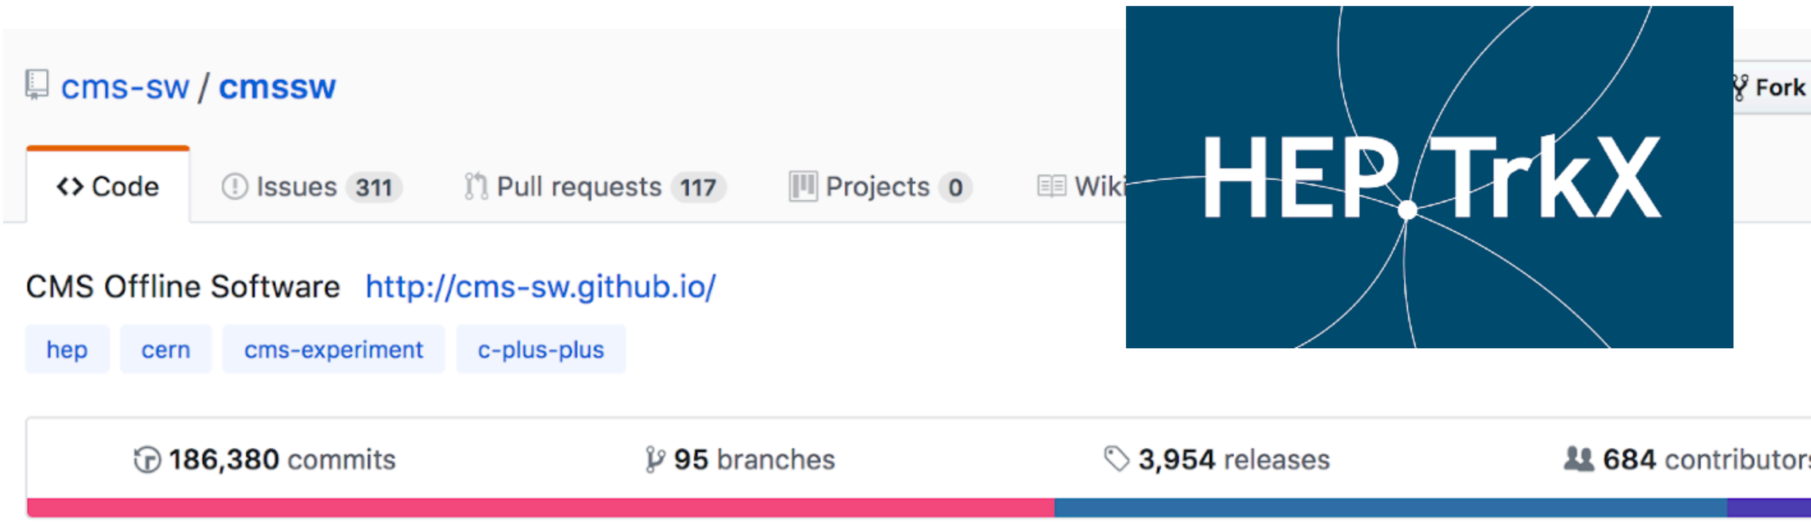

Technology Challenges

- 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
  - 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

K Rupp

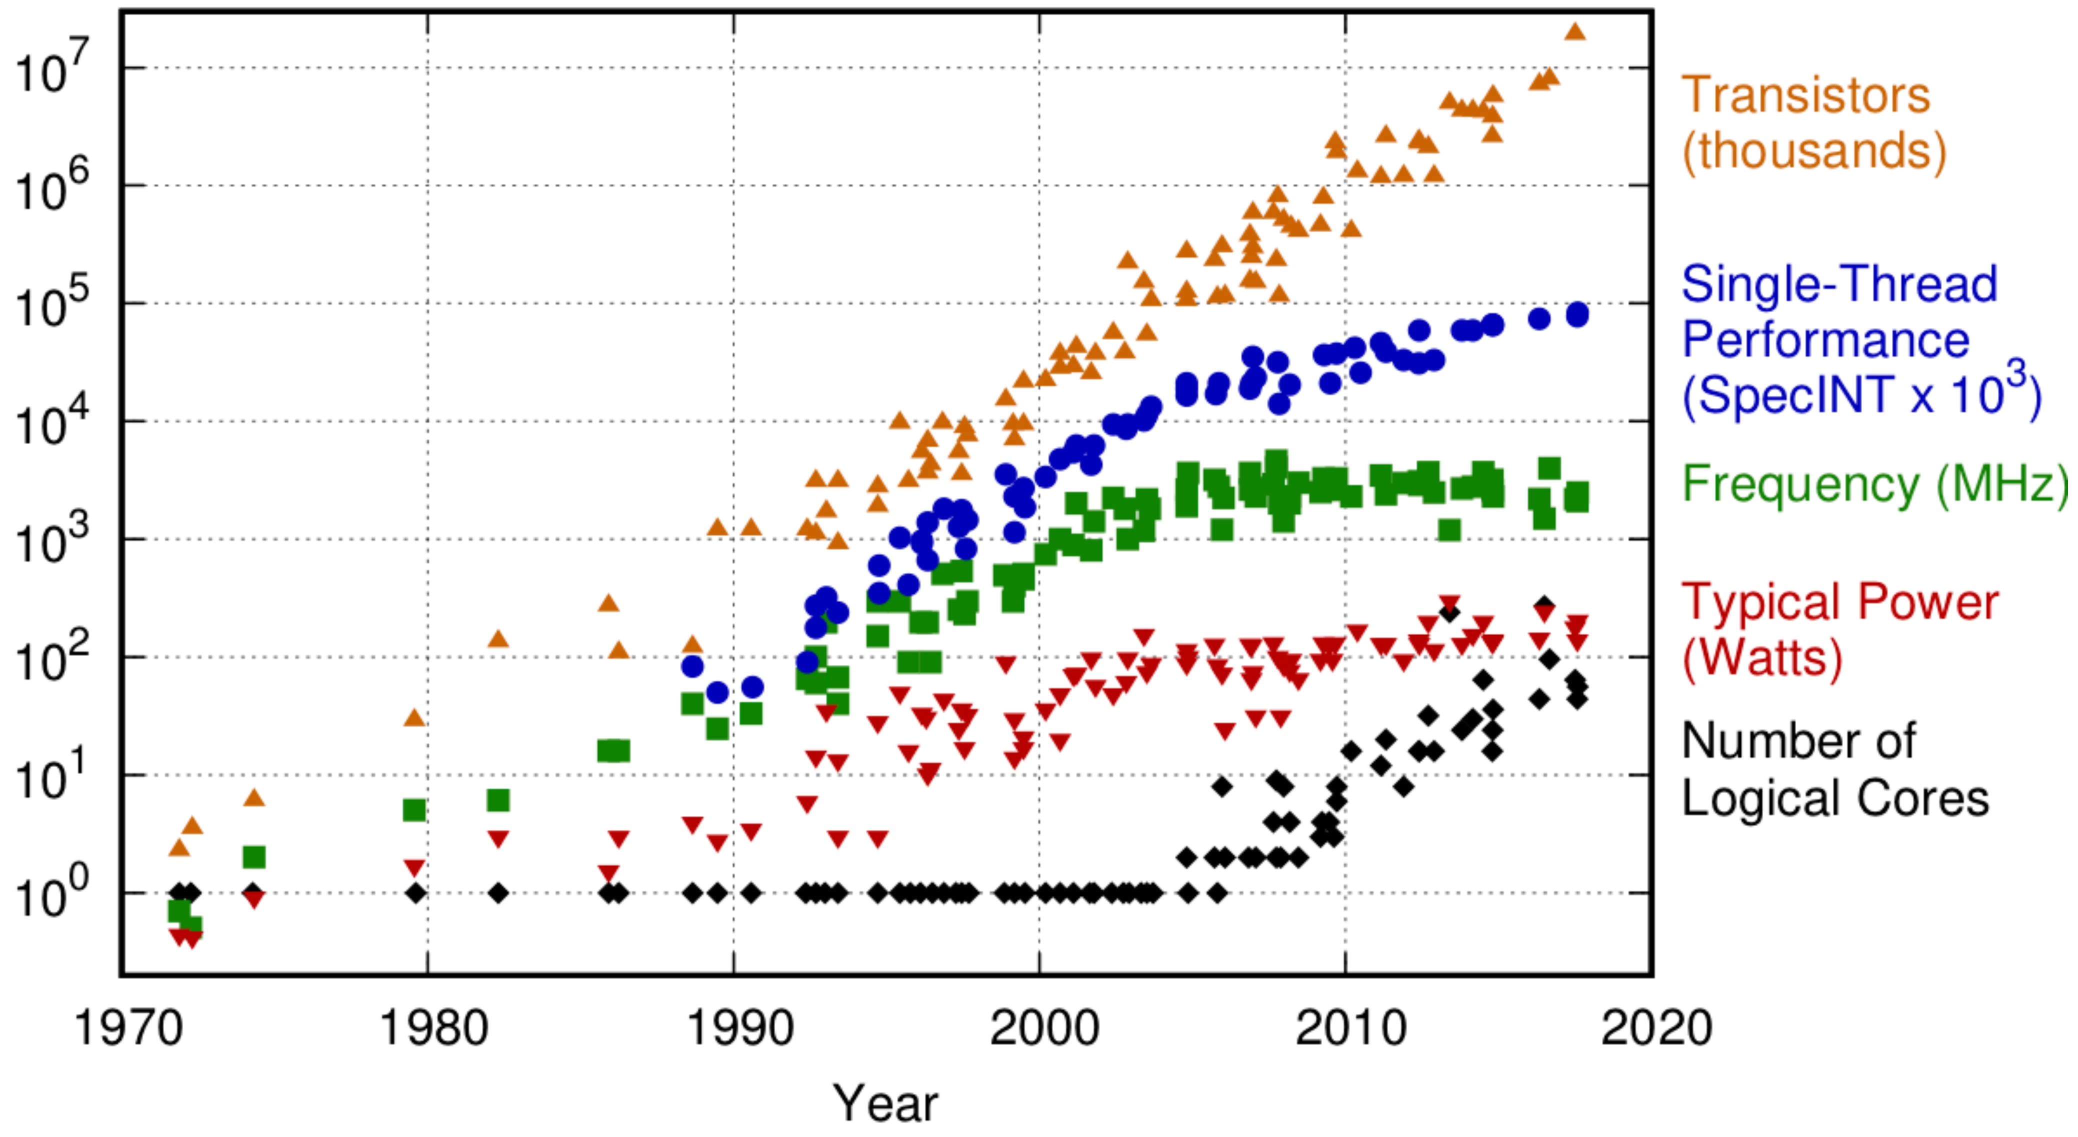

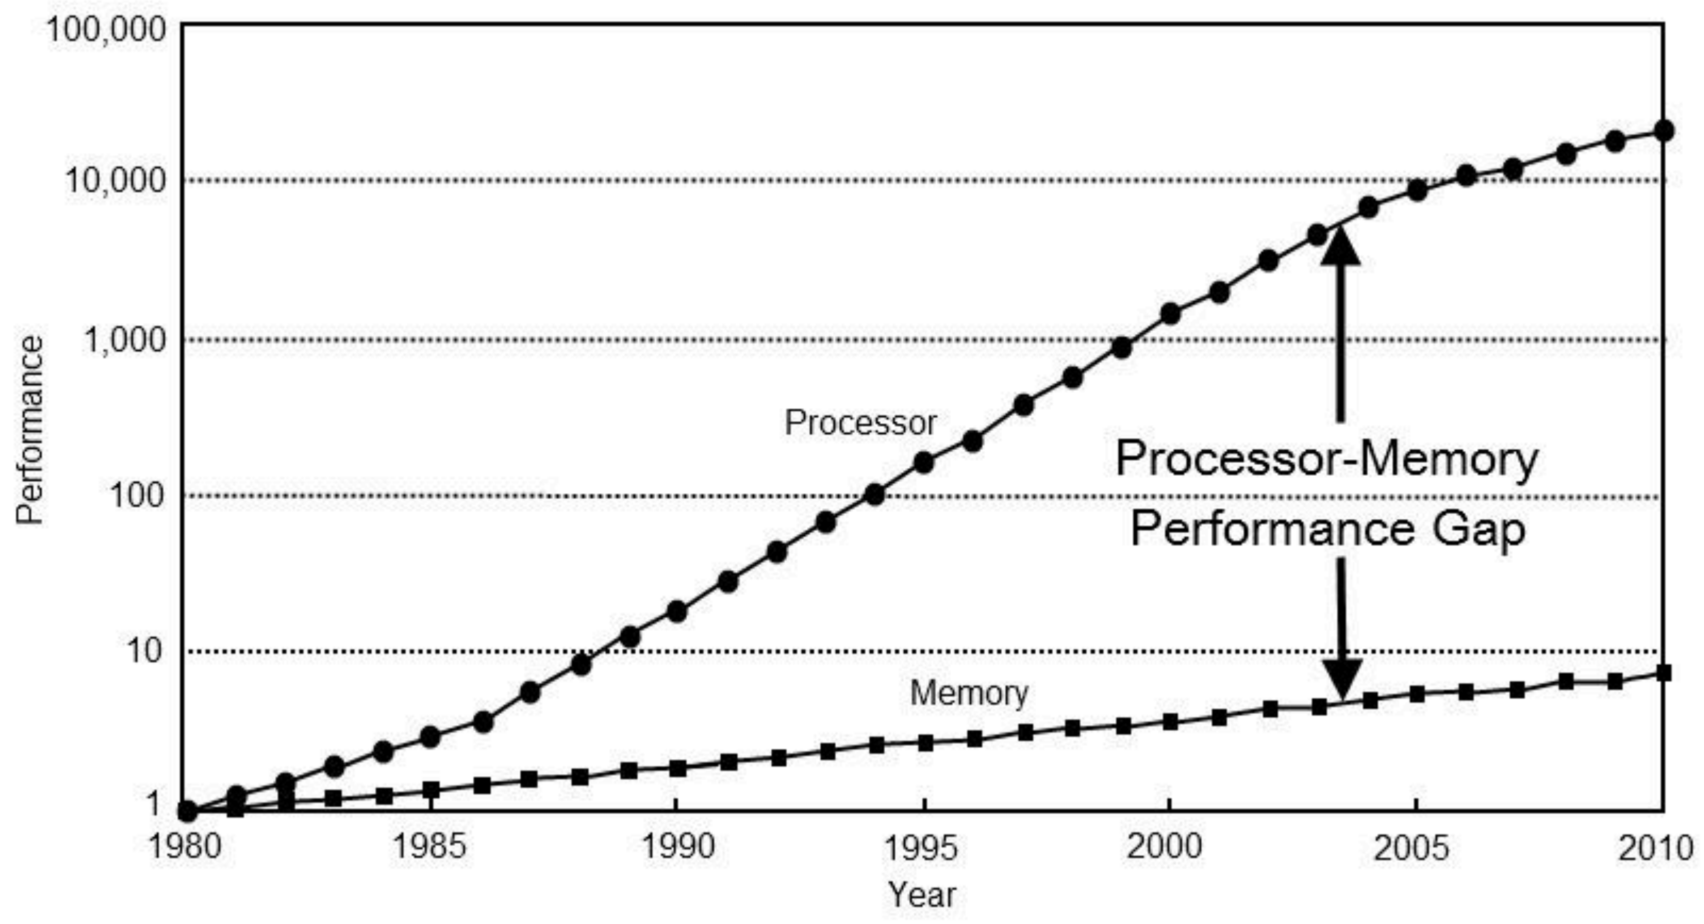

Projections for HL-LHC

- Timeline for Run 3 - Run 5 currently used for projections
- Expect more than doubling of instantaneous luminosity
  - Also increased pile-up

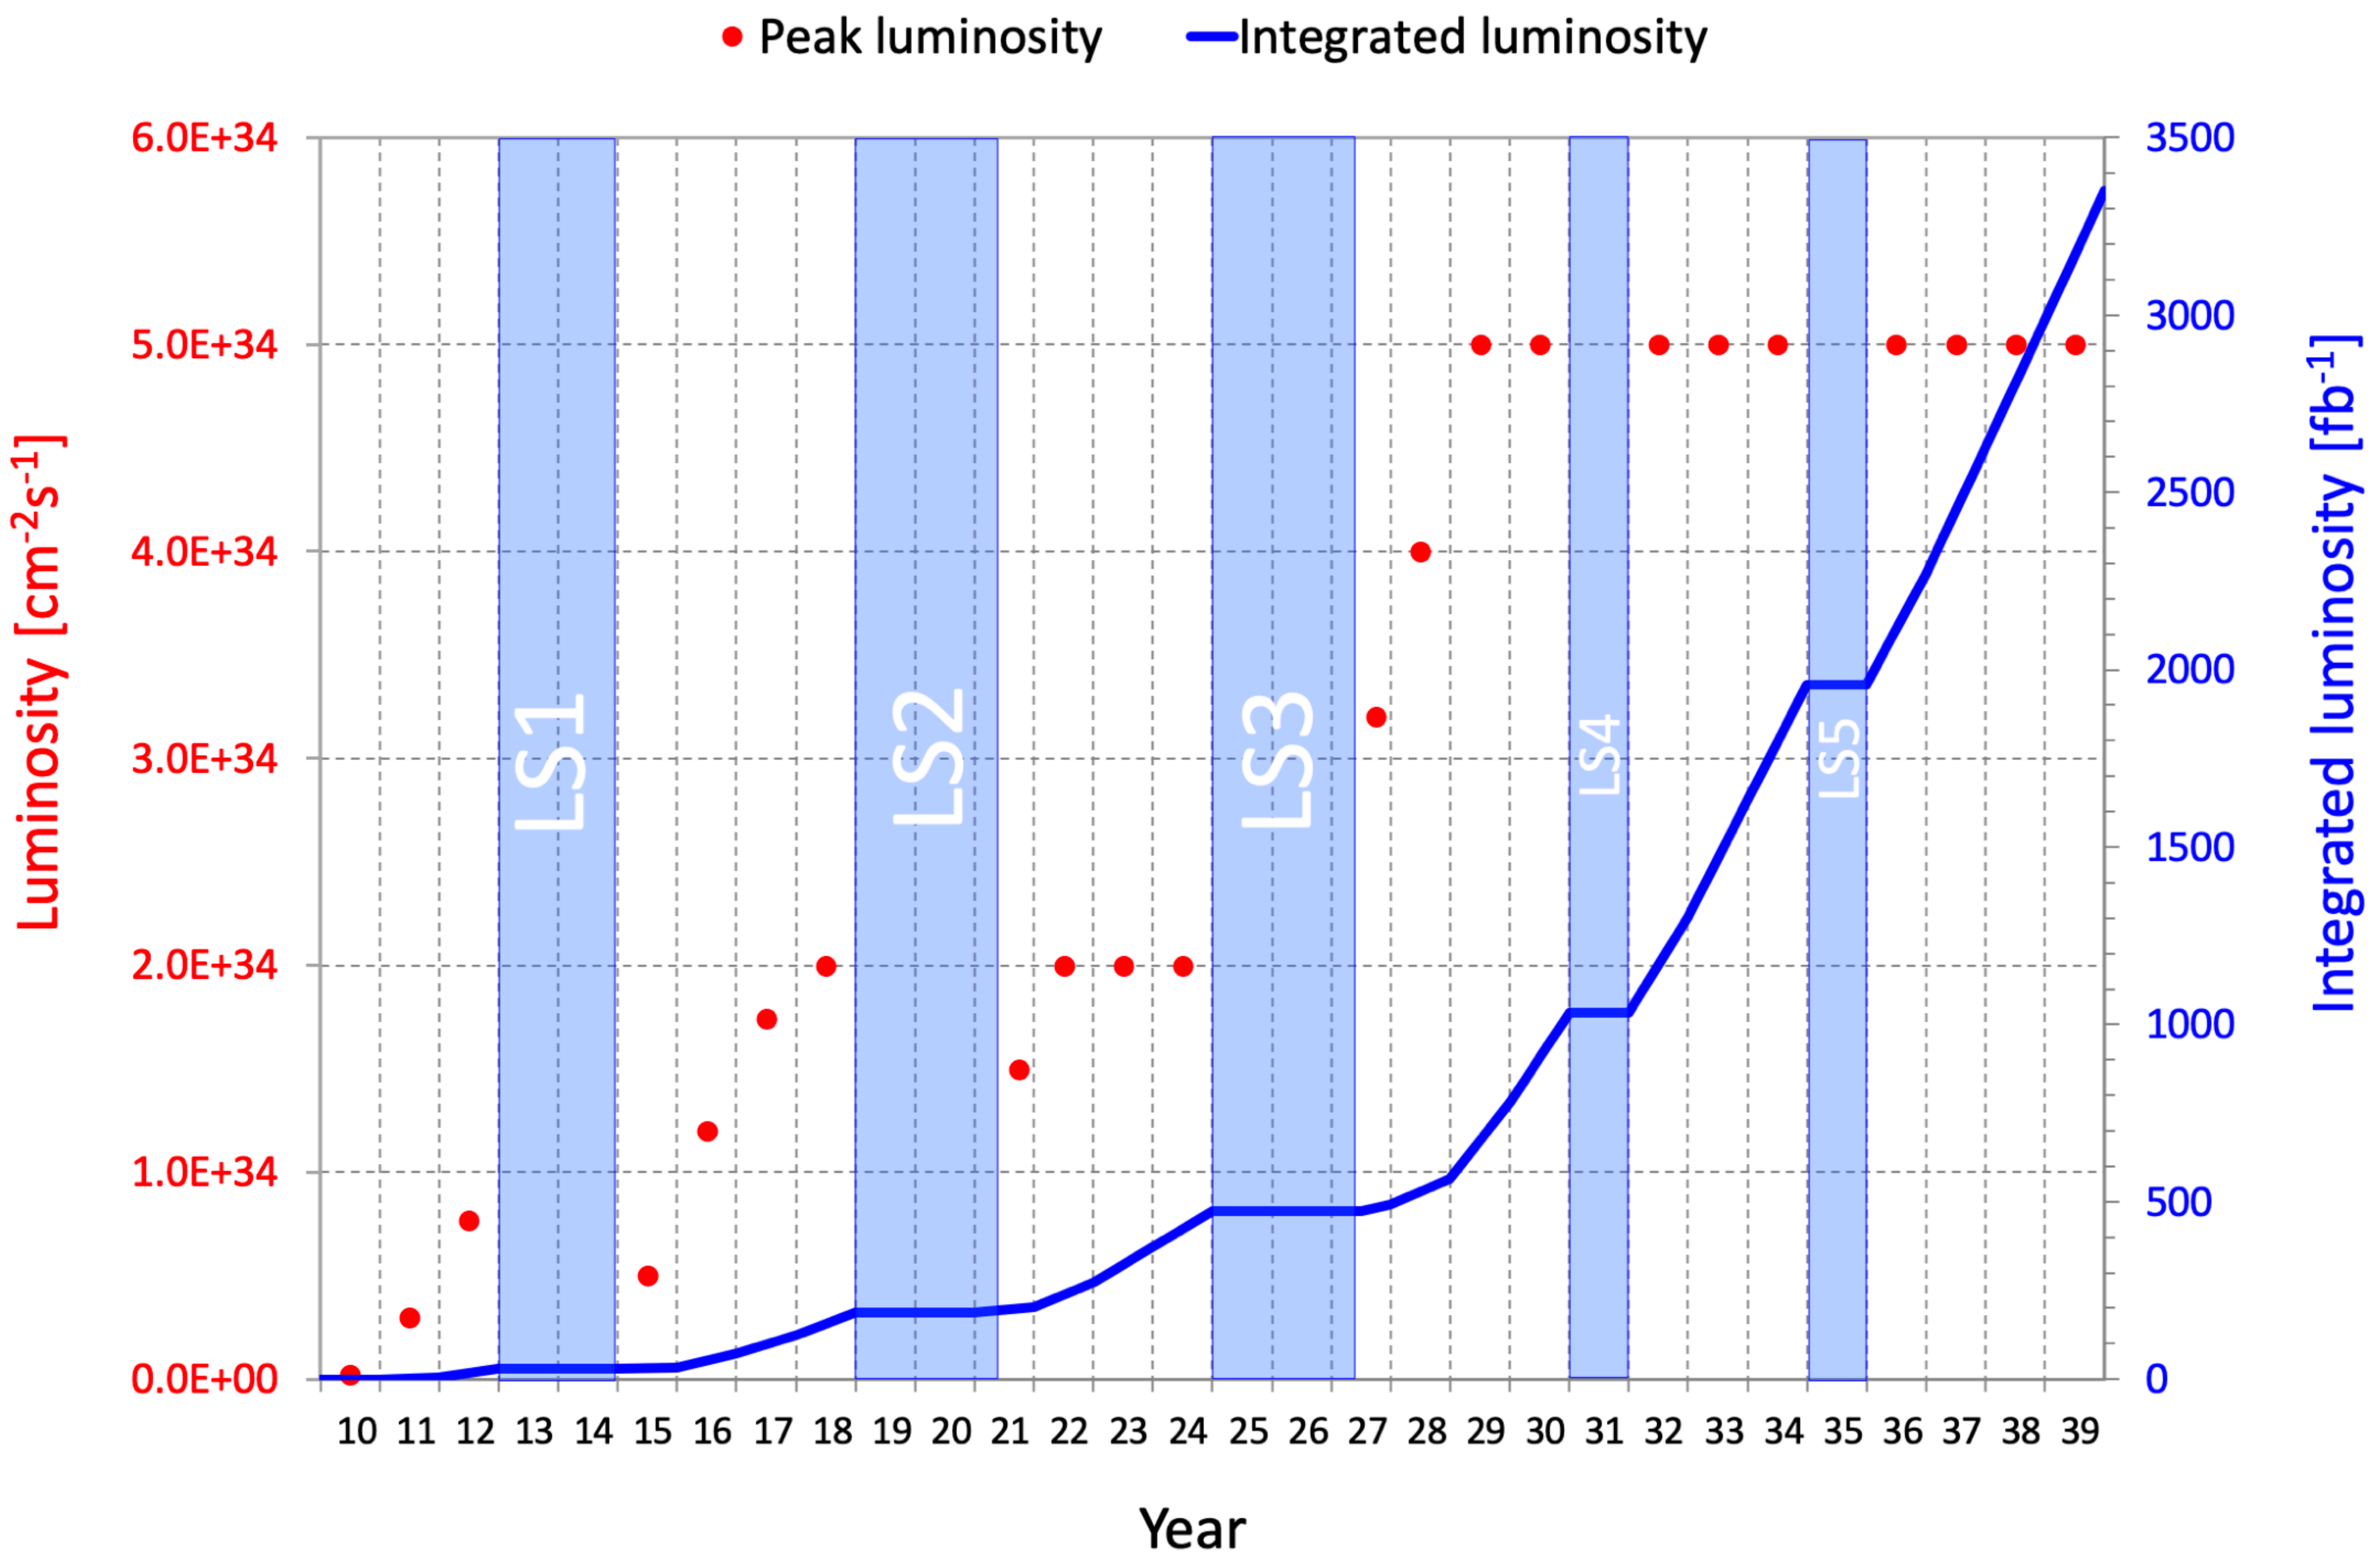

Projections for HL-LHC

- Timeline for Run 3 - Run 5 currently used for projections
- Expect more than doubling of instantaneous luminosity:
  - Also increased pile-up
- Several scenarios for computing model:
  - Baseline
  - Conservative
  - Aggressive

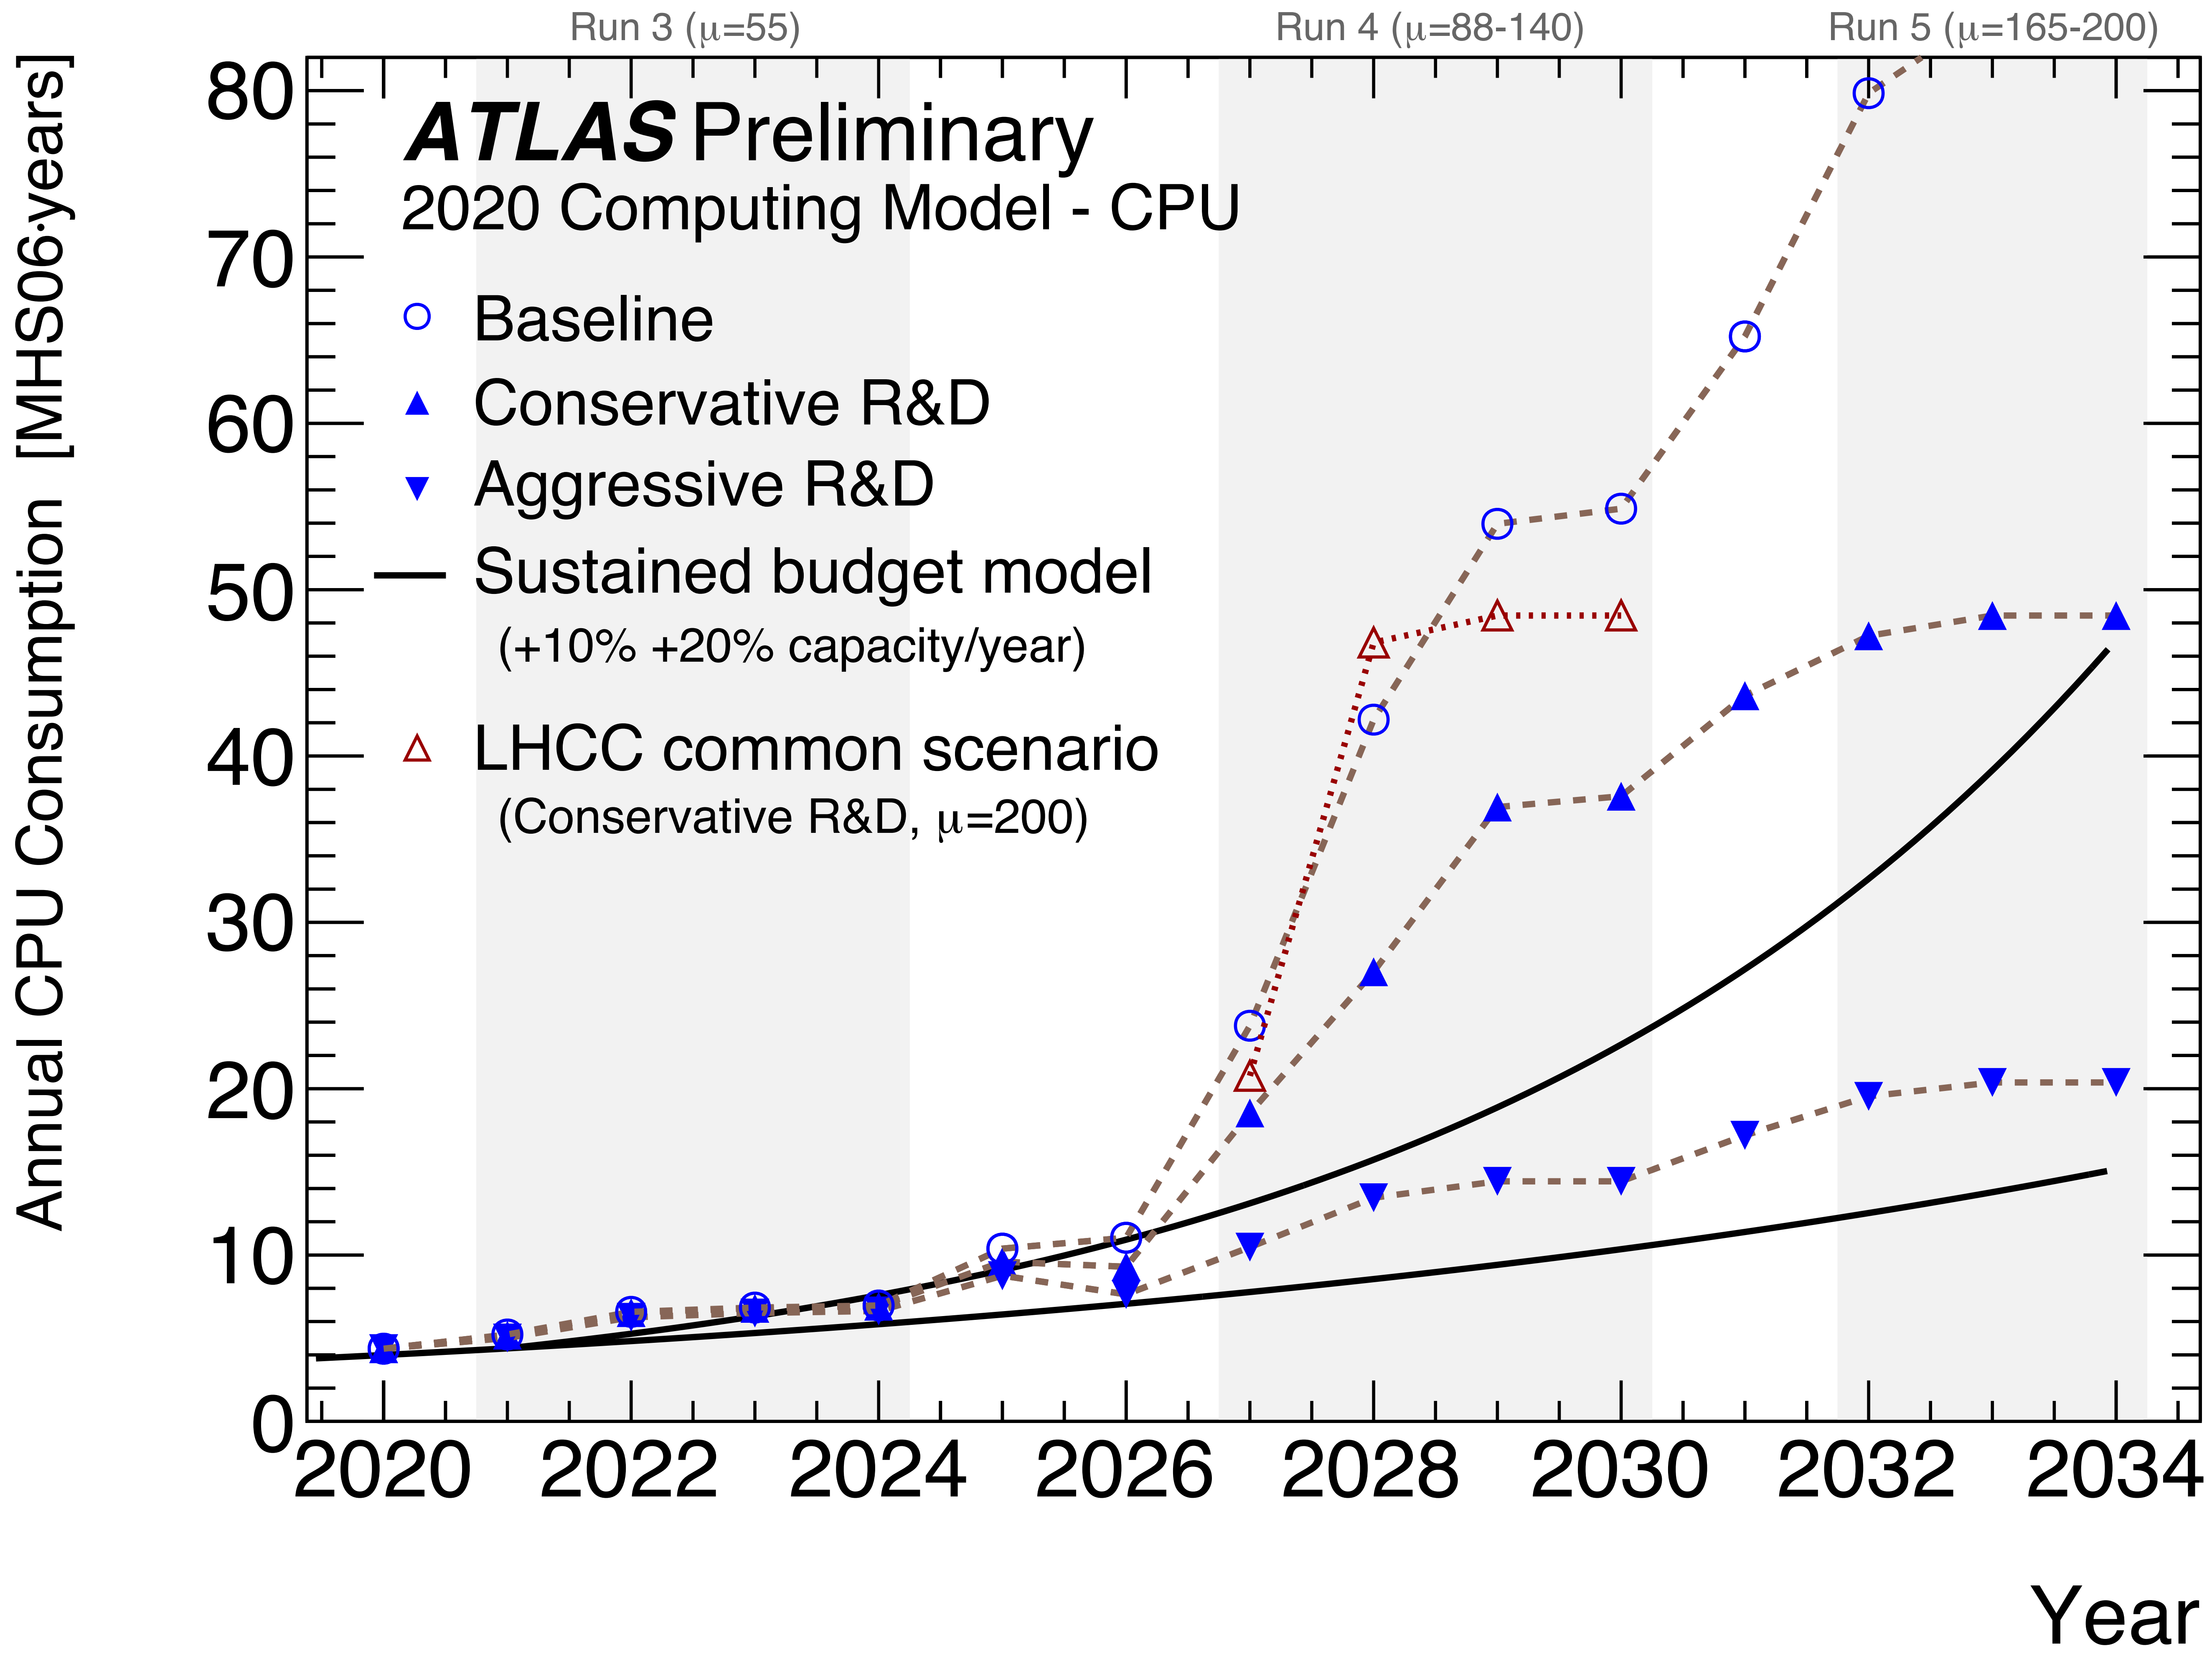

Projections for HL-LHC

- **Baseline**:
  - Assume Run-2 performance (compromise on physics quality)
- **Conservative**:
  - Achieve better physics quality for same CPU time / event as in Run-2
- **Aggressive**:
  - CPU time / event halved, generate 30% (simulate 10%) fewer events

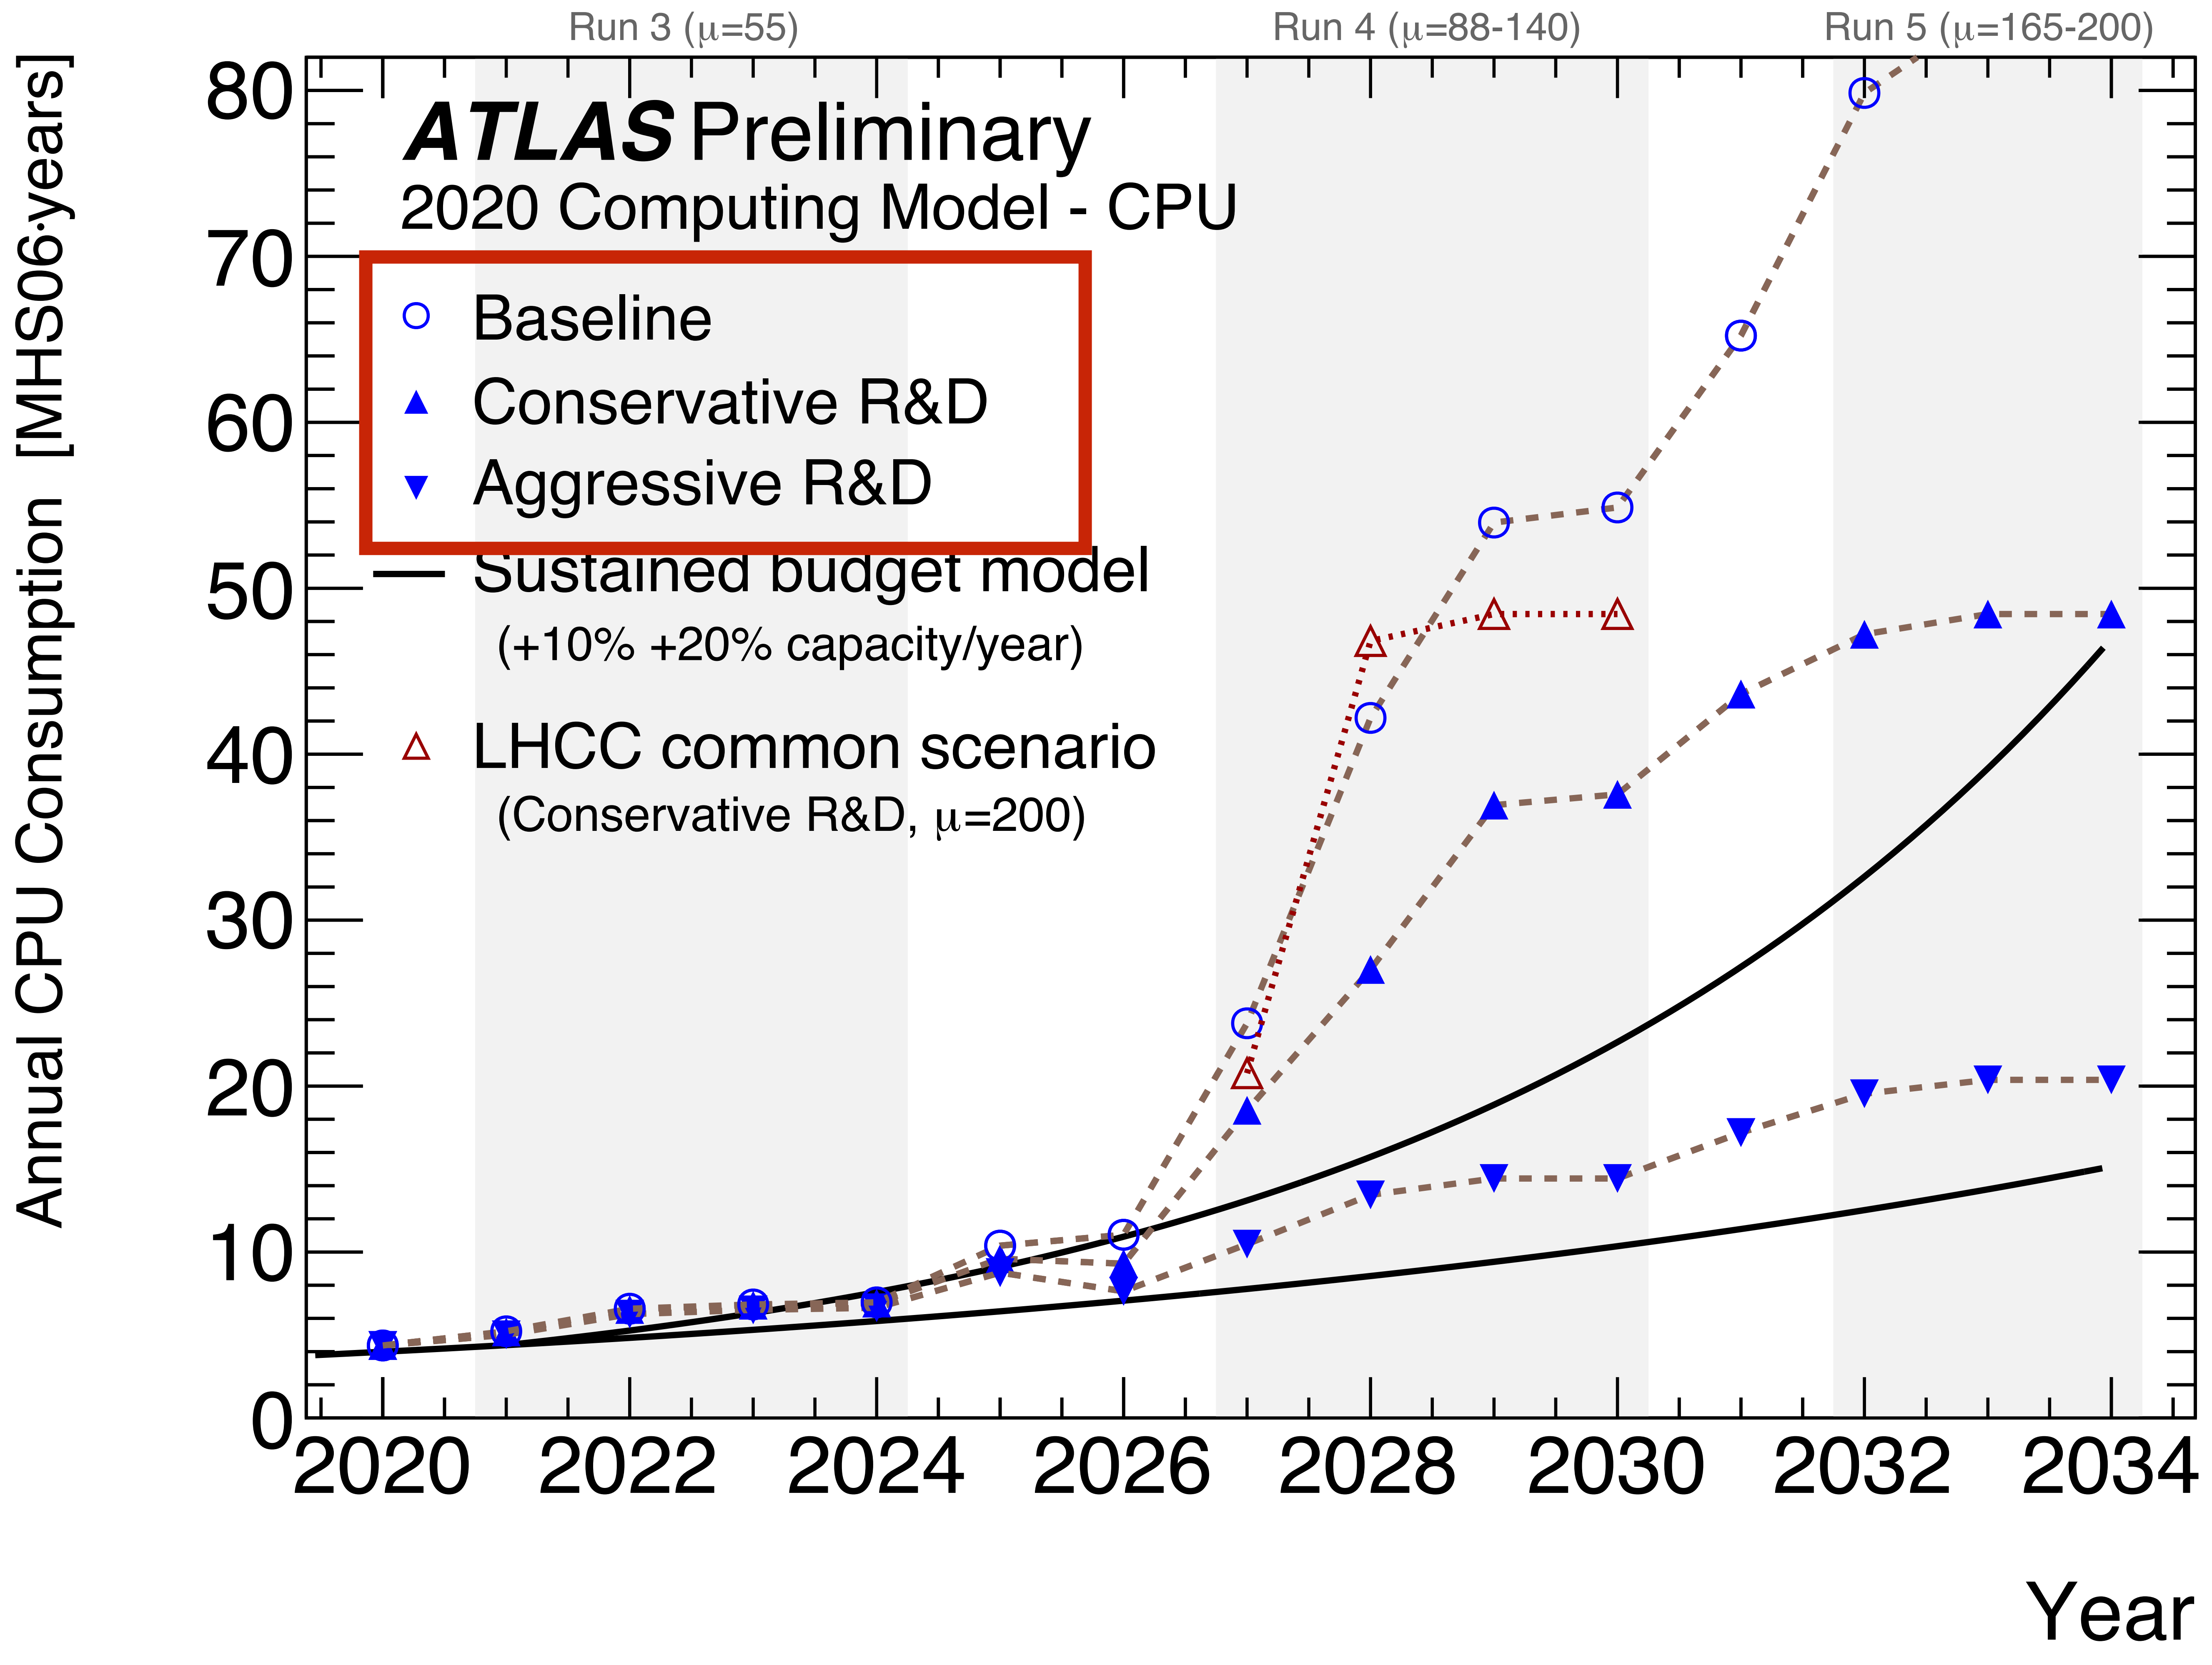

C. Gutschow

Projections for HL-LHC

- **Baseline**:
  - Assume Run-2 performance (compromise on physics quality)
- **Conservative**:
  - Achieve better physics quality for same CPU time / event as in Run-2
- **Aggressive**:
  - CPU time / event halved, generate 30% (simulate 10%) fewer events

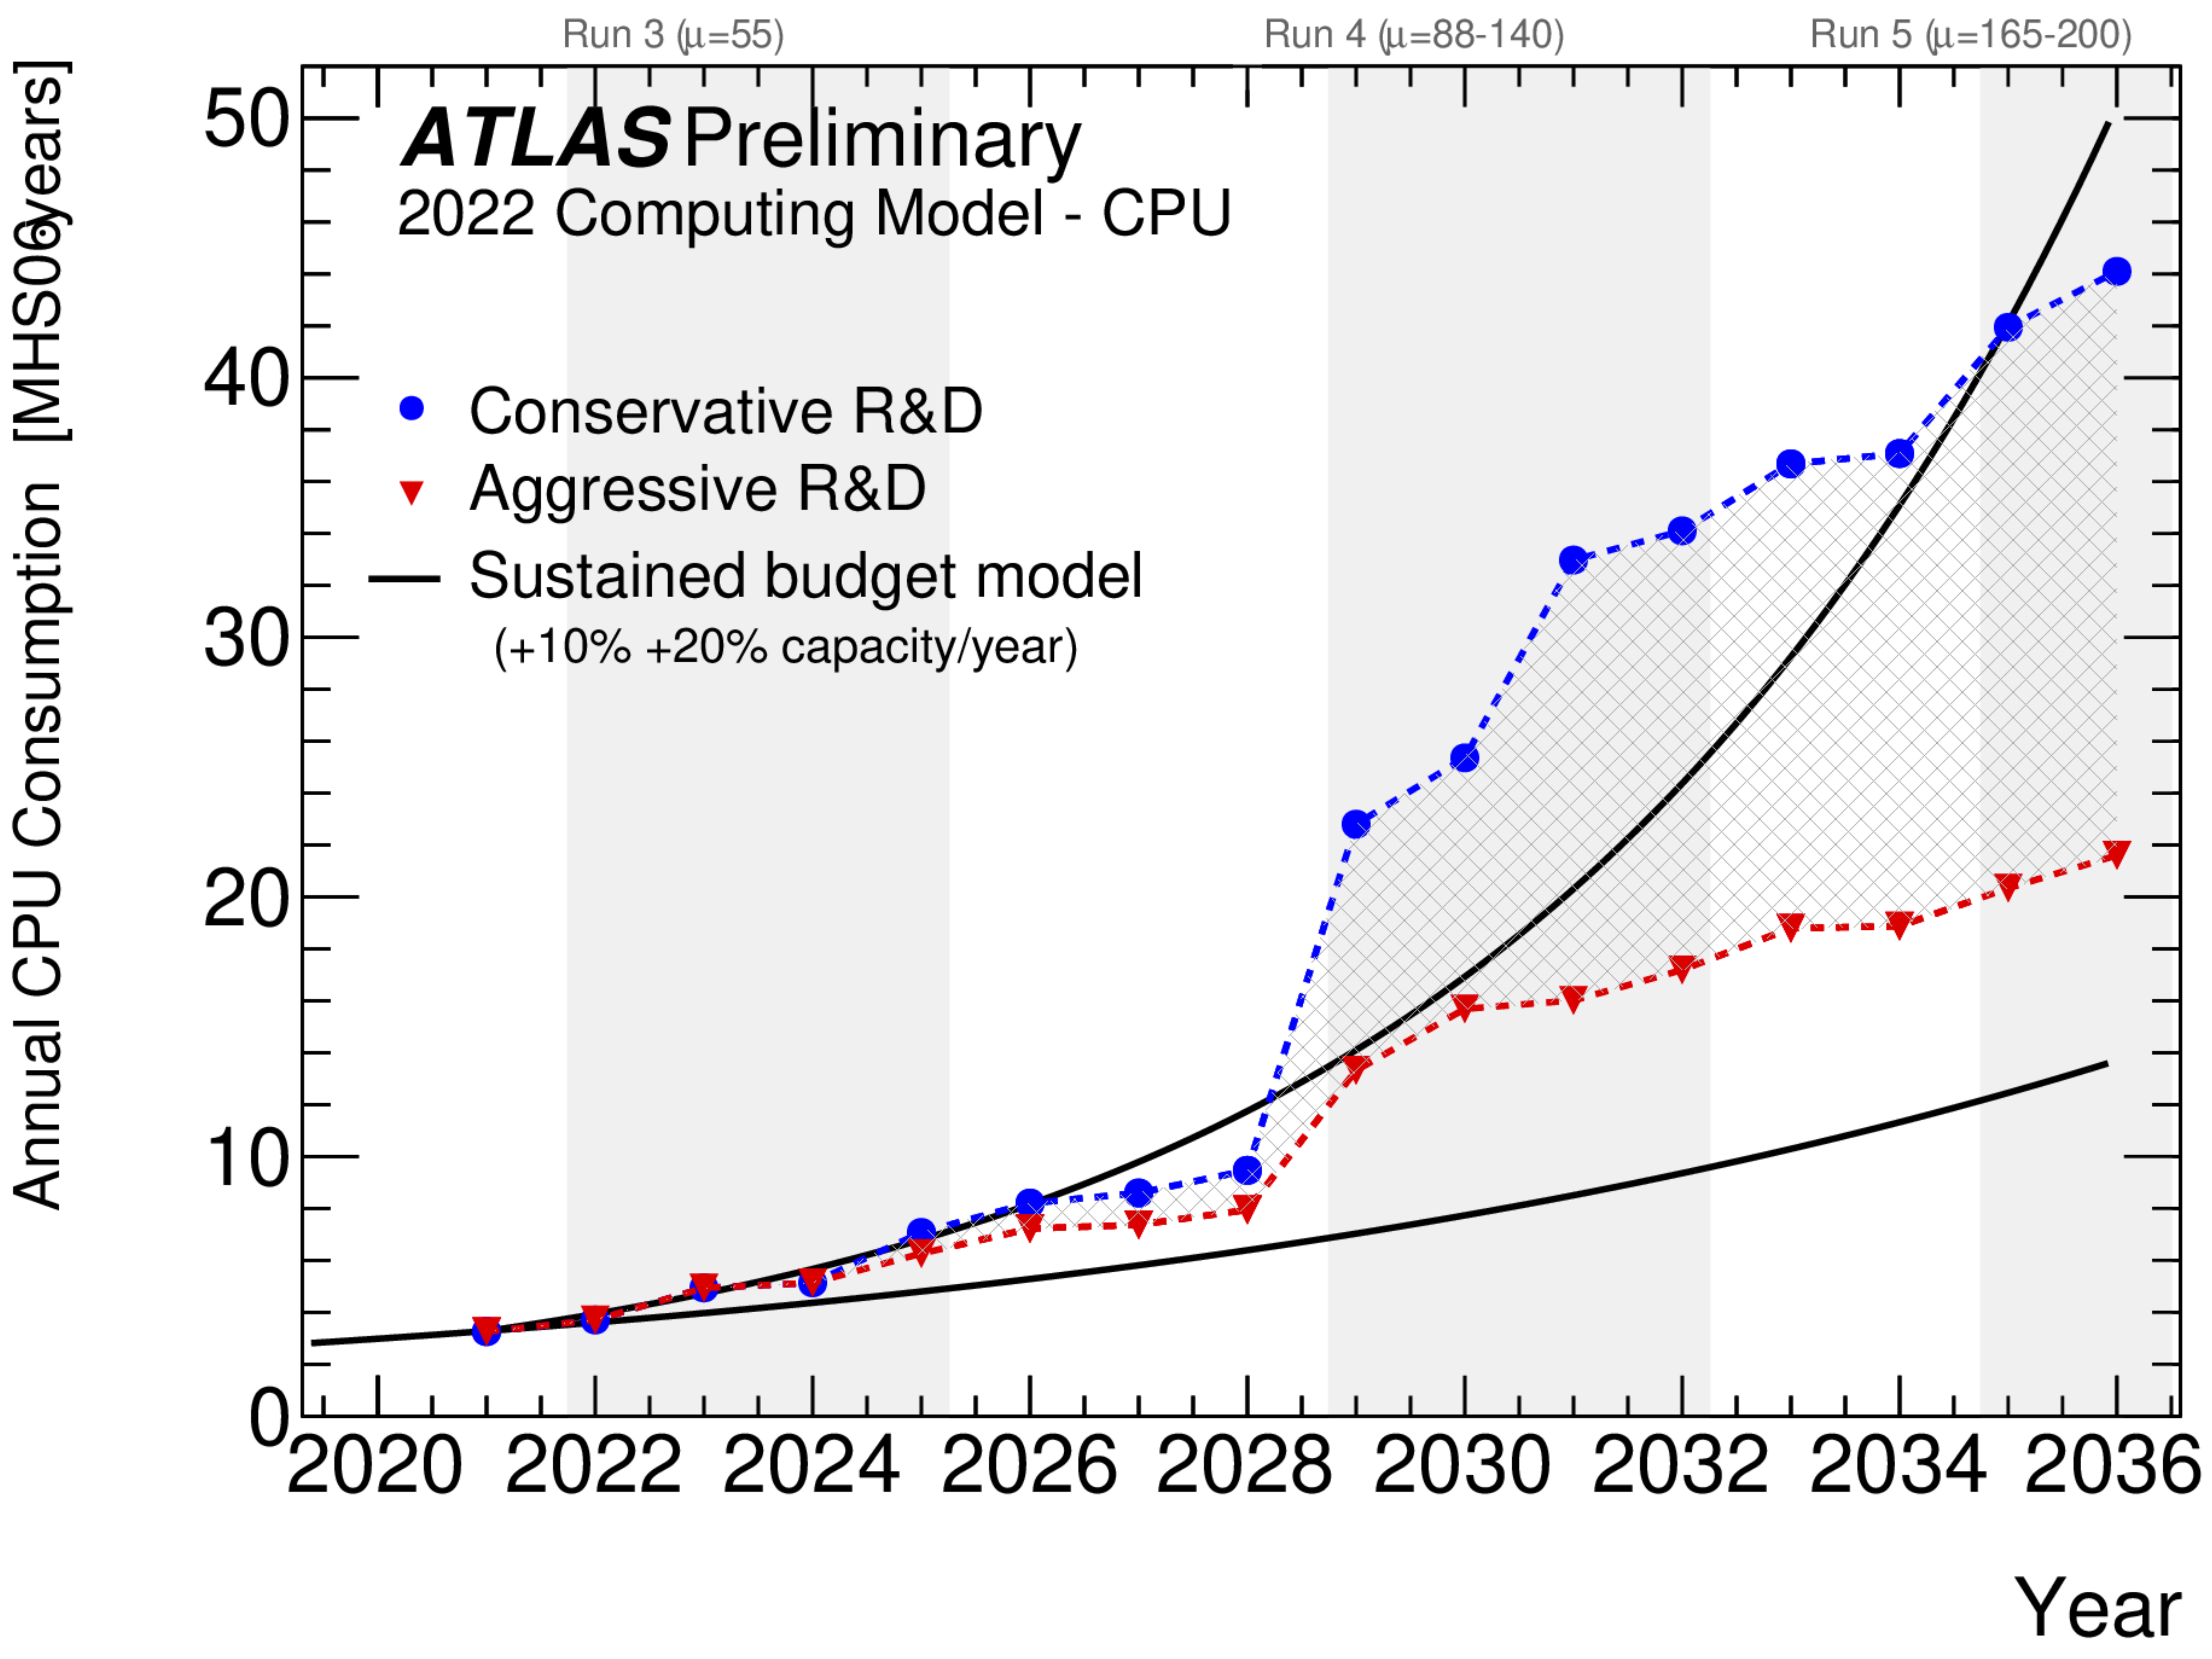

1 | Major developments from the 2013 Strategy

- 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.

Heterogeneity challenge

UNIVERSITY OF SUSSEX

- There are a lot of possible parallel architectures on the market
  - CPUs with multiple cores and wide registers
  - GPUs with many cores; FPGAs - Nvidia (many generations - often significantly different), AMD, Intel
- Many options for coding, both generic and specific:
  - Cuda, TBB, OpenACC, OpenMP, OpenCL (→ Vulcan), alpaka, Kokkos, ...
- Frustratingly no clear winner, mutually exclusive solutions and many niches
  - One option is to isolate the algorithmic code from a ‘wrapper’ that targets a particular device or architecture, hiding details in a lower level library
- Require large effort in software development
  - Adaptation of programs, libraries, workload and data management systems
  - All these have to be developed, commissioned and verified
  - Needs a huge investment of skilled developers
  - Not always available at the required level

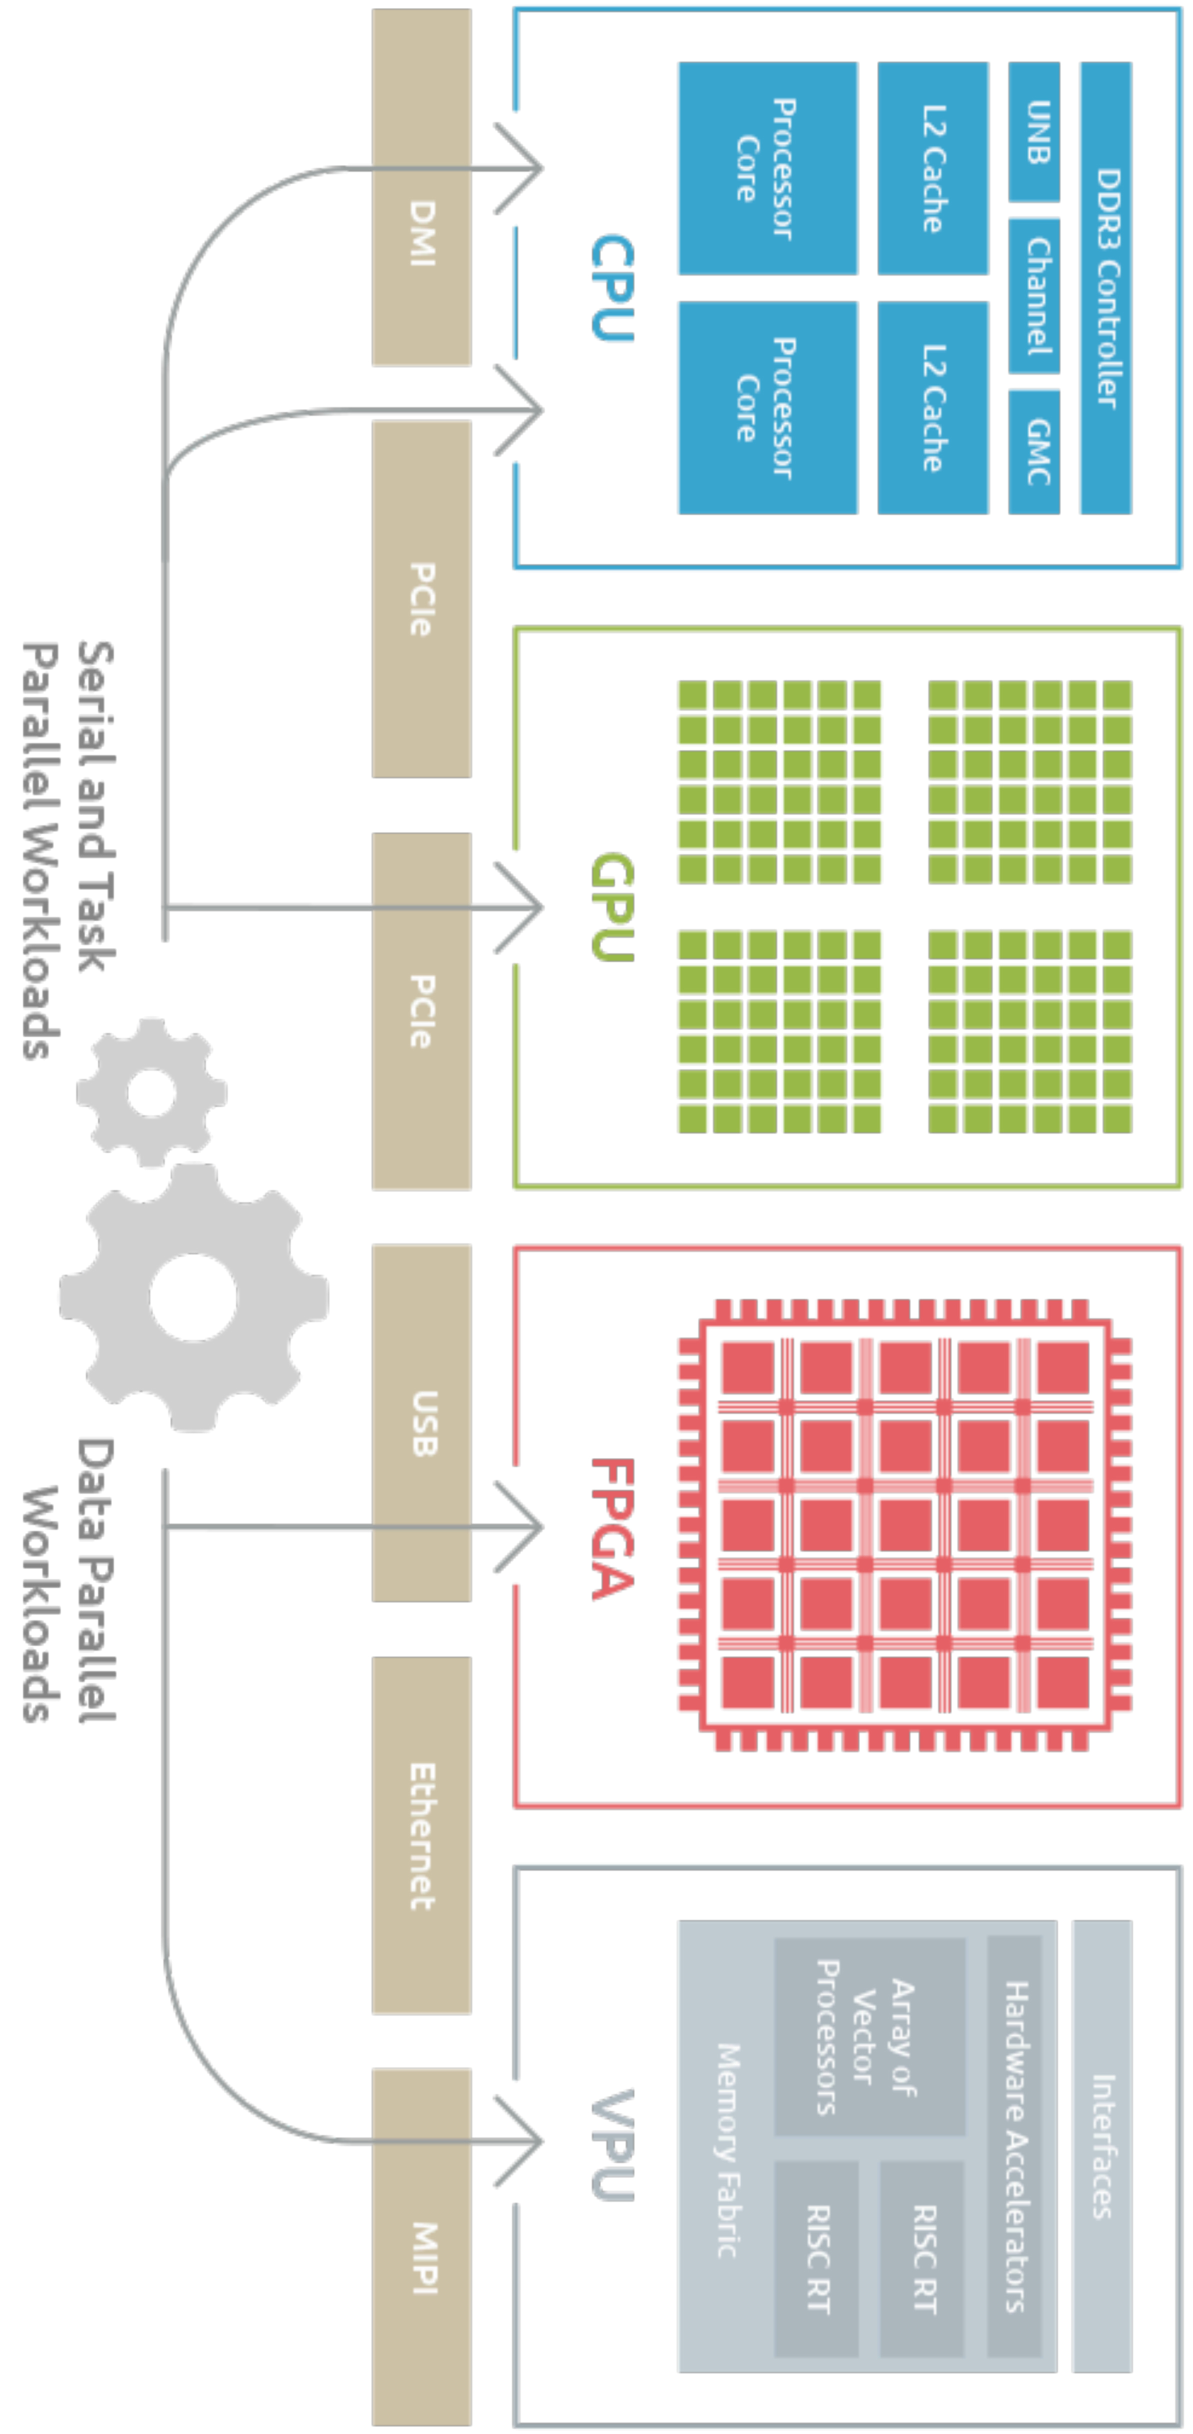

Machine Learning

- Used for many years in HEP
  - Algorithms learn by example (training) how to perform tasks instead of being programmed
- Significant advances in the last years in 'deep learning'
  - Deep means many neural network layers
  - Fast differentiability and use of GPUs
- Rapid development driven by industry
  - Vibrant ecosystem of tools and techniques
  - Highly optimised for modern, specialised hardware

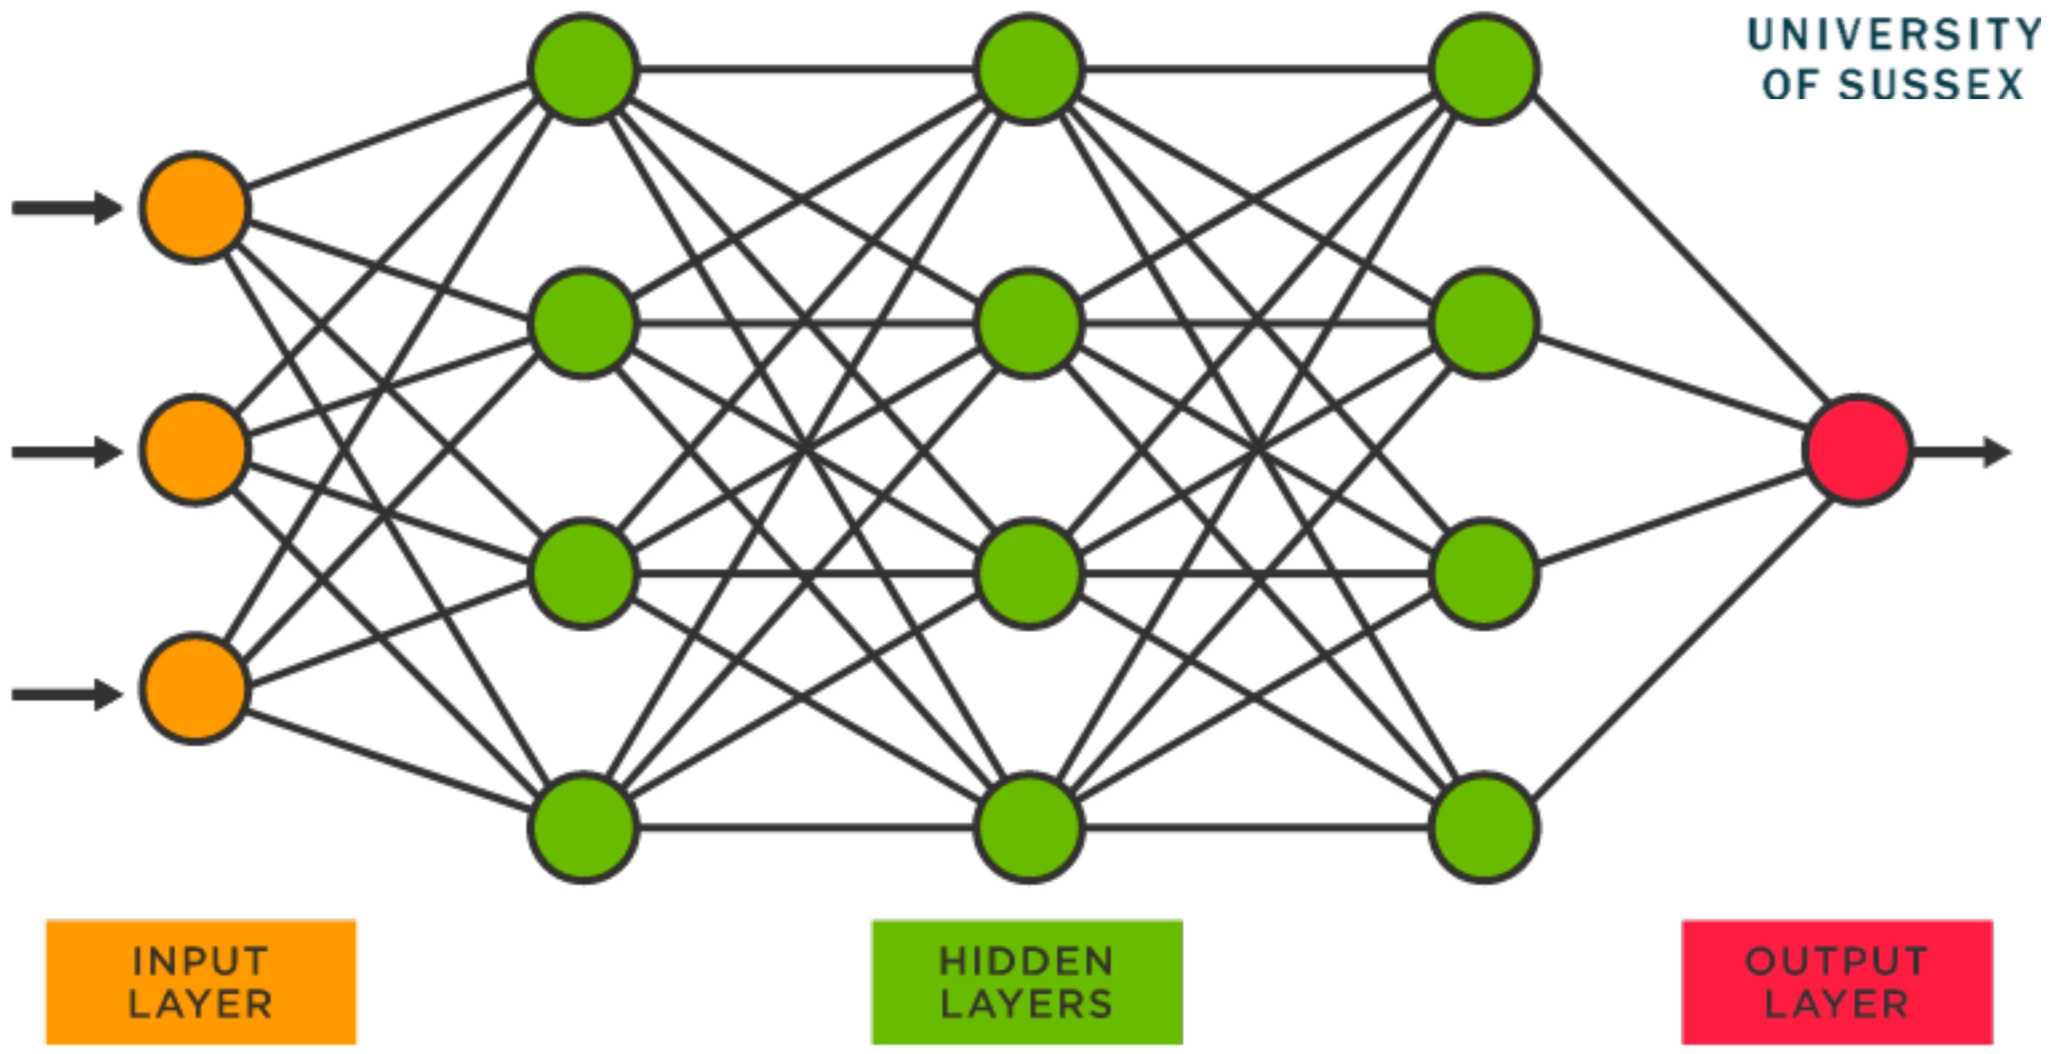

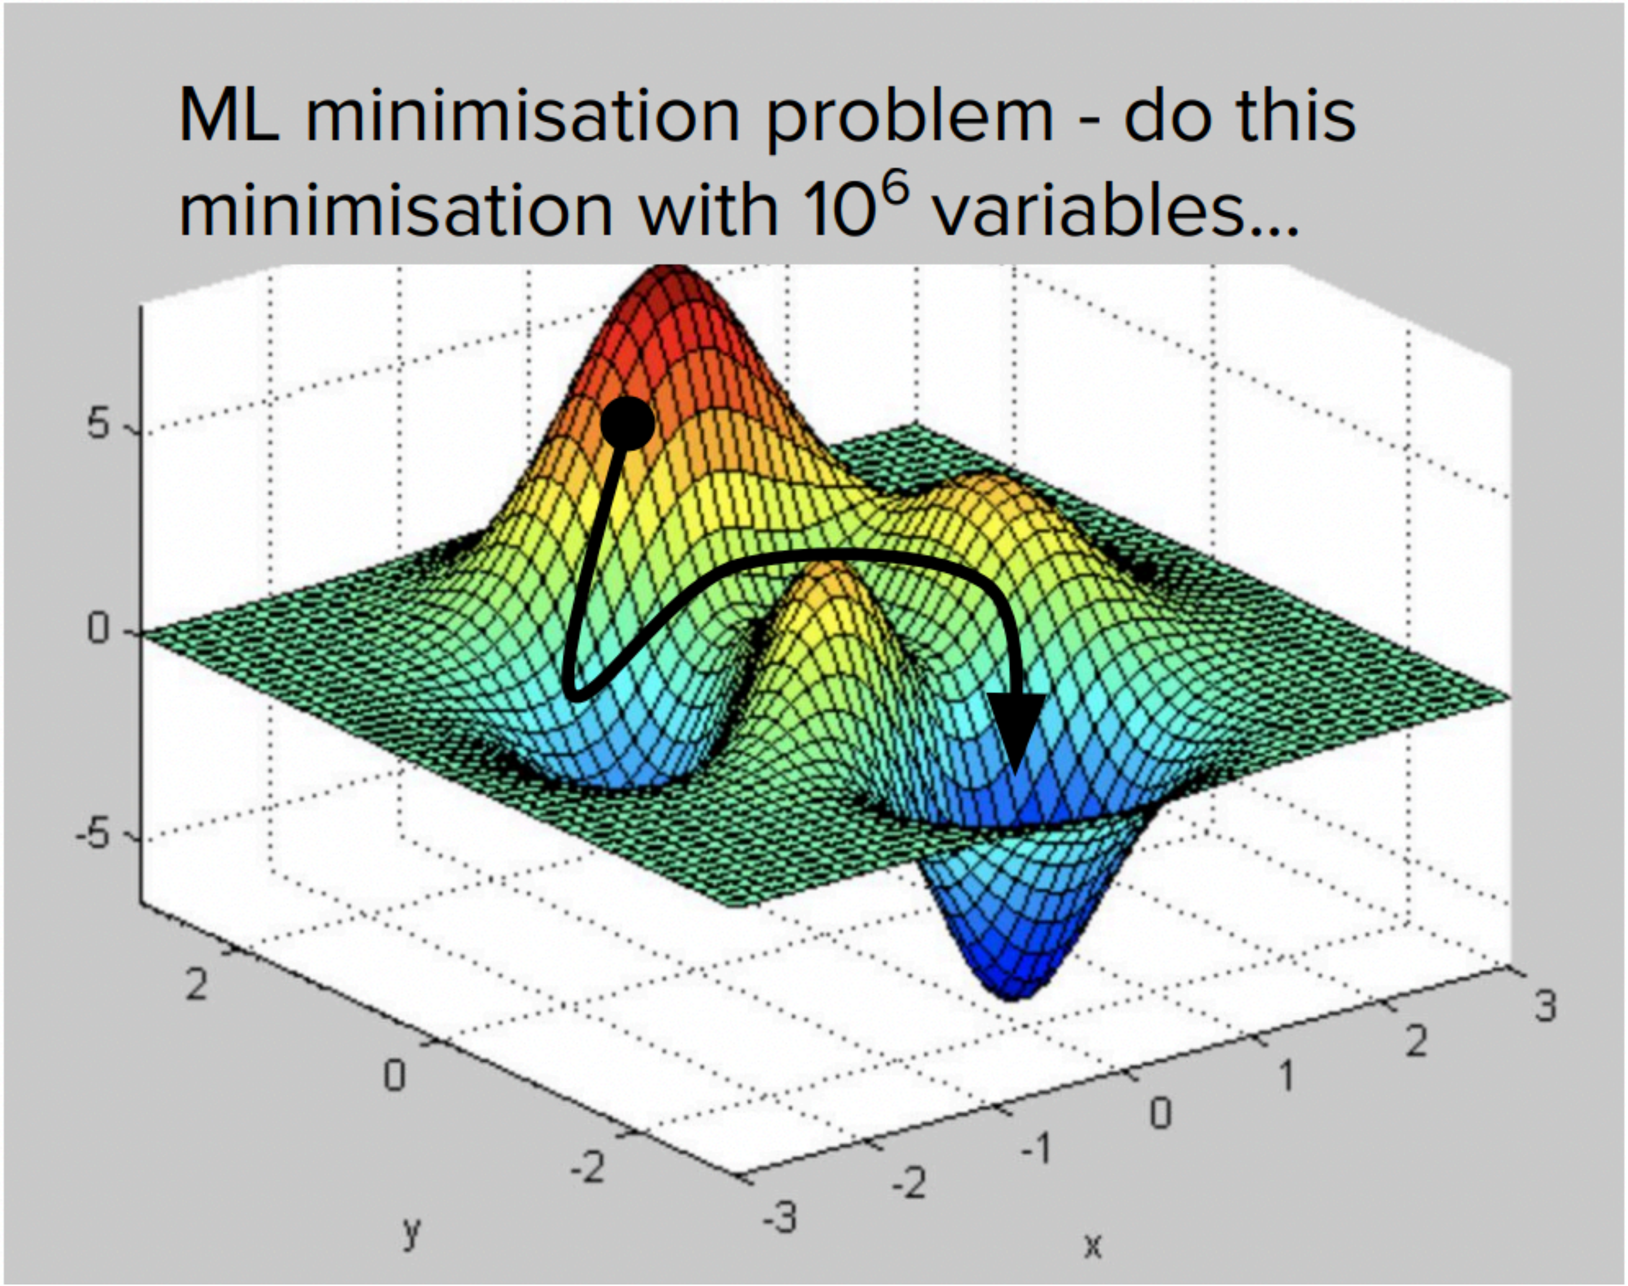

Machine Learning

- Better discrimination
  - Important input for analysis (see improvements with Higgs)
  - Also used at HLT as inference can be fast (N.B. training can be slow!)
  - HEP analogies to image recognition or text processing
- Replace expensive calculations with trained output
  - E.g. calorimeter simulations and other complex physical processes
- There are significant opportunities here
  - Need to combine physics and data science knowledge
  - Field evolves rapidly and we need to deepen our expertise
- Integration into our workflows is not at all settled
  - Resource provision, efficient use, heterogeneity and programming models pose problems

Table 1 | Effect of machine learning on the discovery and study of the Higgs boson

Analysis	Years of data collection	Sensitivity without machine learning	Sensitivity with machine learning	Ratio of *P* values	Additional data required
CMS[24] $H \to \gamma\gamma$	2011–2012	2.2σ, $P = 0.014$	2.7σ, $P = 0.0035$	4.0	51%
ATLAS[43] $H \to \tau^+\tau^-$	2011–2012	2.5σ, $P = 0.0062$	3.4σ, $P = 0.00034$	18	85%
ATLAS[99] $VH \to bb$	2011–2012	1.9σ, $P = 0.029$	2.5σ, $P = 0.0062$	4.7	73%
ATLAS[41] $VH \to bb$	2015–2016	2.8σ, $P = 0.0026$	3.0σ, $P = 0.00135$	1.9	15%
CMS[100] $VH \to bb$	2011–2012	1.4σ, $P = 0.081$	2.1σ, $P = 0.018$	4.5	125%

Machine learning at the energy and intensity frontiers of particle physics,

https://doi.org/10.1038/s41586-018-0361-2

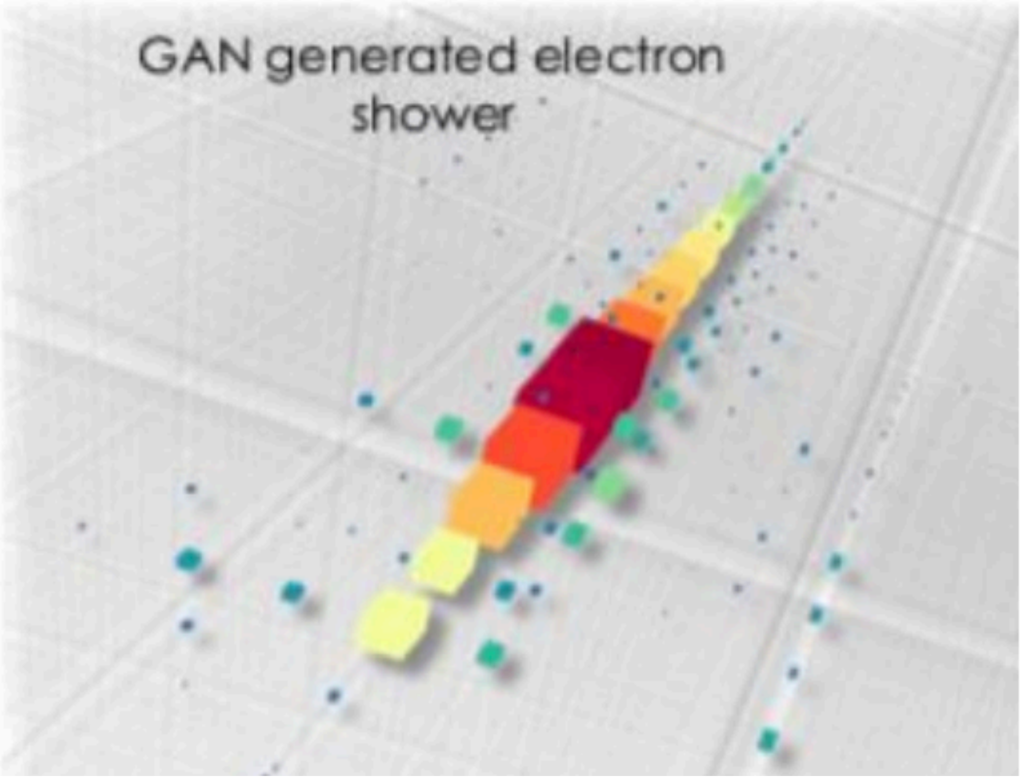

Use of Generative Adversarial Networks to simulate calorimeter showers, trained on G4 events (S. Vallacorsa)

Far Future Ideas

- Intensely active research in e.g. Quantum Computing
  - Europe have invested 1B€ in Quantum Flagship Program; US invest heavily as well (including for HEP)

- Certainly a game changer if engineering of sufficient, stable q-bits can be achieved
  - Rapid progress in the last 5 years, but still far from being practical and useful
  - Even with some spectacular breakthroughs commercialisation would take time

- How should HEP be involved? And at what level?
  - Are these with extra resources or some effort that we dedicate from our pool?
  - Mapping QC to current HEP algorithms? New algorithms enabled by QC? Programmable? Maintainable?

Event Generators

- Starting the simulated events chain from theory
  - Previously a very small part of LHC computing budget (cf. detector simulation), no pressure to optimise
- Increasing use of higher precision (NLO, NNLO, …)
  - Negative weights become a serious problem
  - Greatly increases the CPU budget fraction given to event generation
- Theory community not rewarded generators optimisation
  - Not enough experts and lack incentives to adapt to modern CPU architectures
- Technically these codes are a good target for optimisation
  - A lot of pure maths, floating point intensive
    - No inputs, small outputs
    - Ideal for HPC environments
    - Some parts of the code has been ported to GPU as well (MadGraph)
    - Can we find ways to collaborate with software engineers?

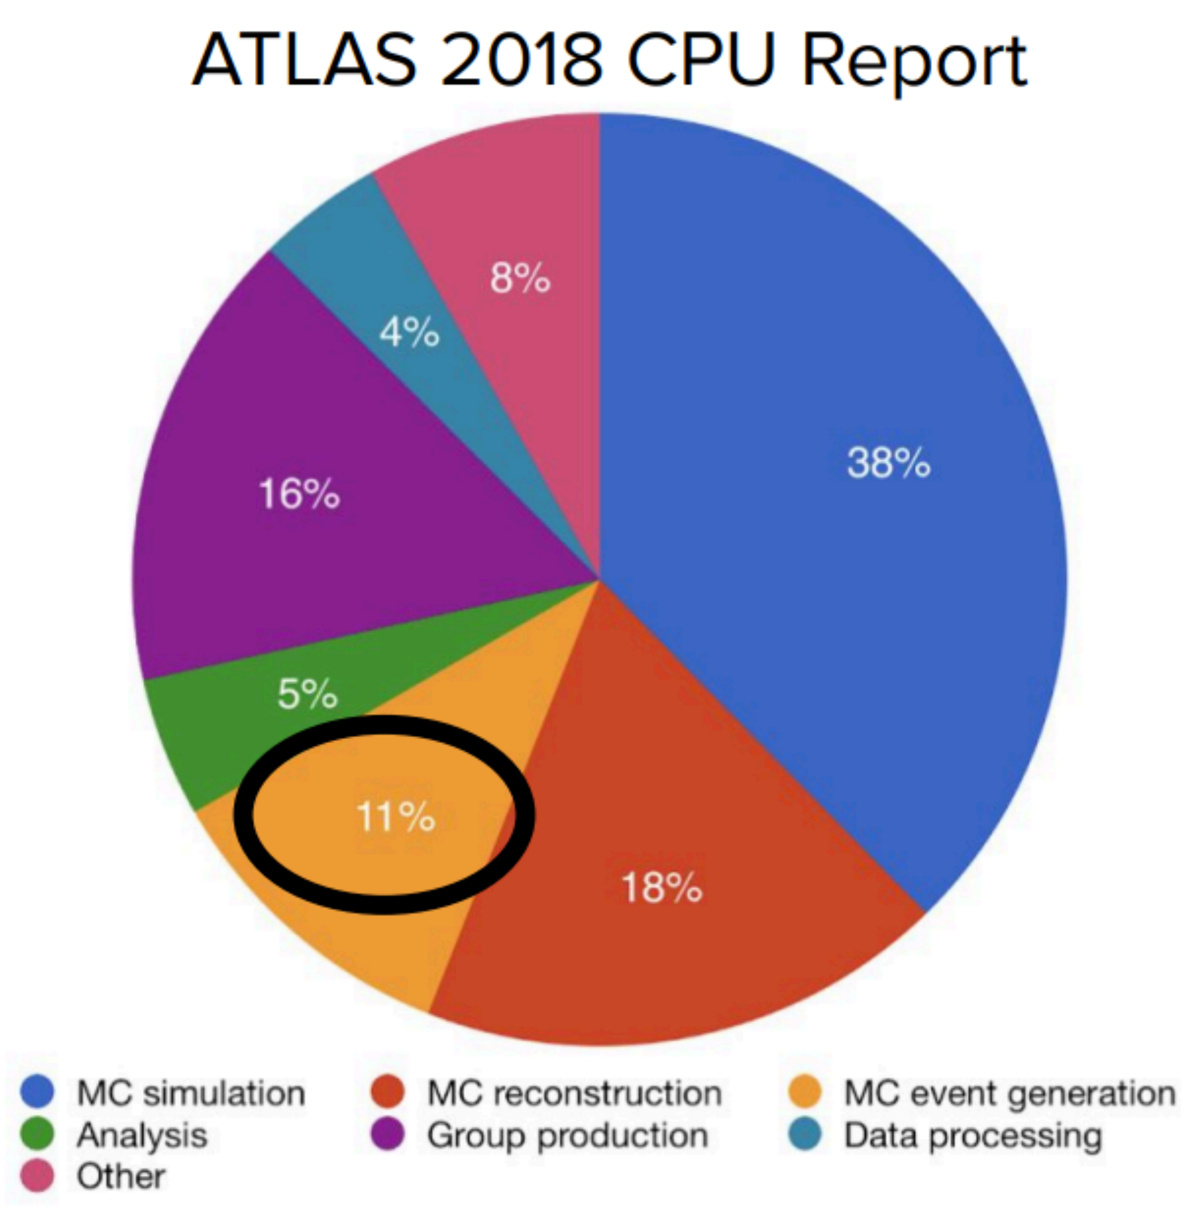

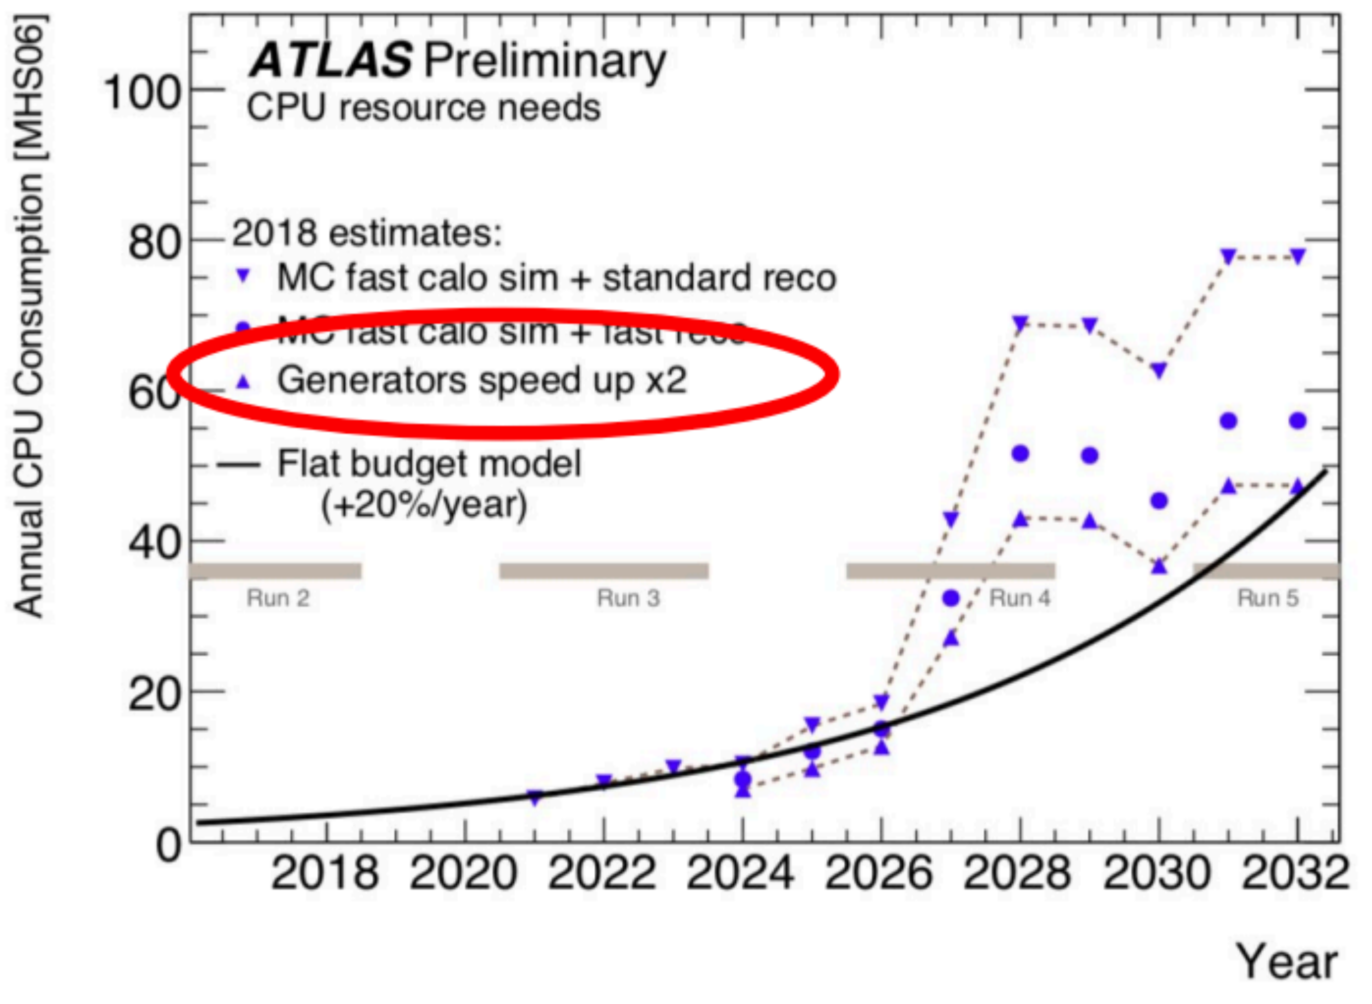

Simulation

- A major consumer of LHC grid resources today
  - Experiments with higher data rates will need to more simulation
- Flat budget scenarios don't give a lot more cycles
  - So need faster simulation
- Technical improvement programme helps
  - GeantV R&D modernises code and introduces vectorisation
  - Serious studies of GPU porting (US Exascale Computing Project)
- Probably not be sufficient to meet future needs
  - Will need to trade off accuracy for speed with approximate and hybrid simulation approaches
- Machine learning techniques are gaining ground
  - Are they are good enough cf. Geant4?
  - Can they be integrated?
  - Lots od recent progress!

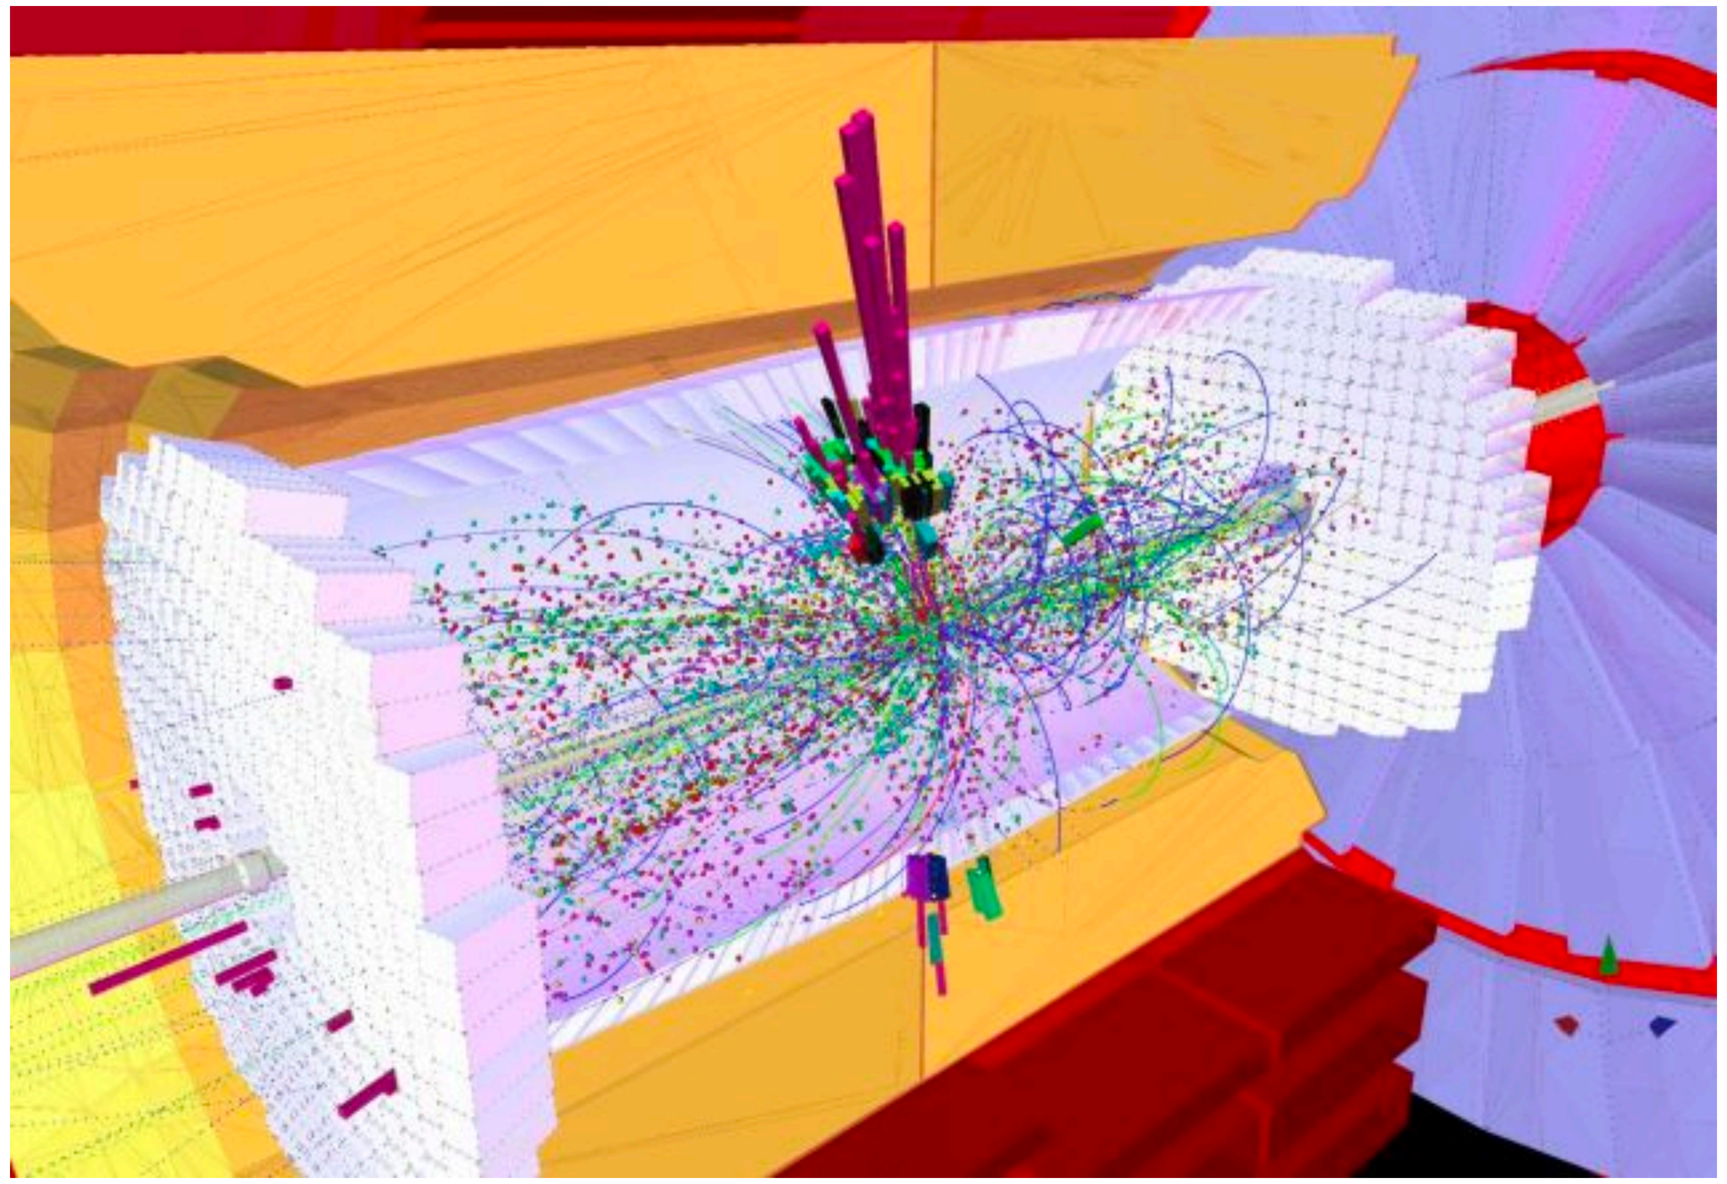

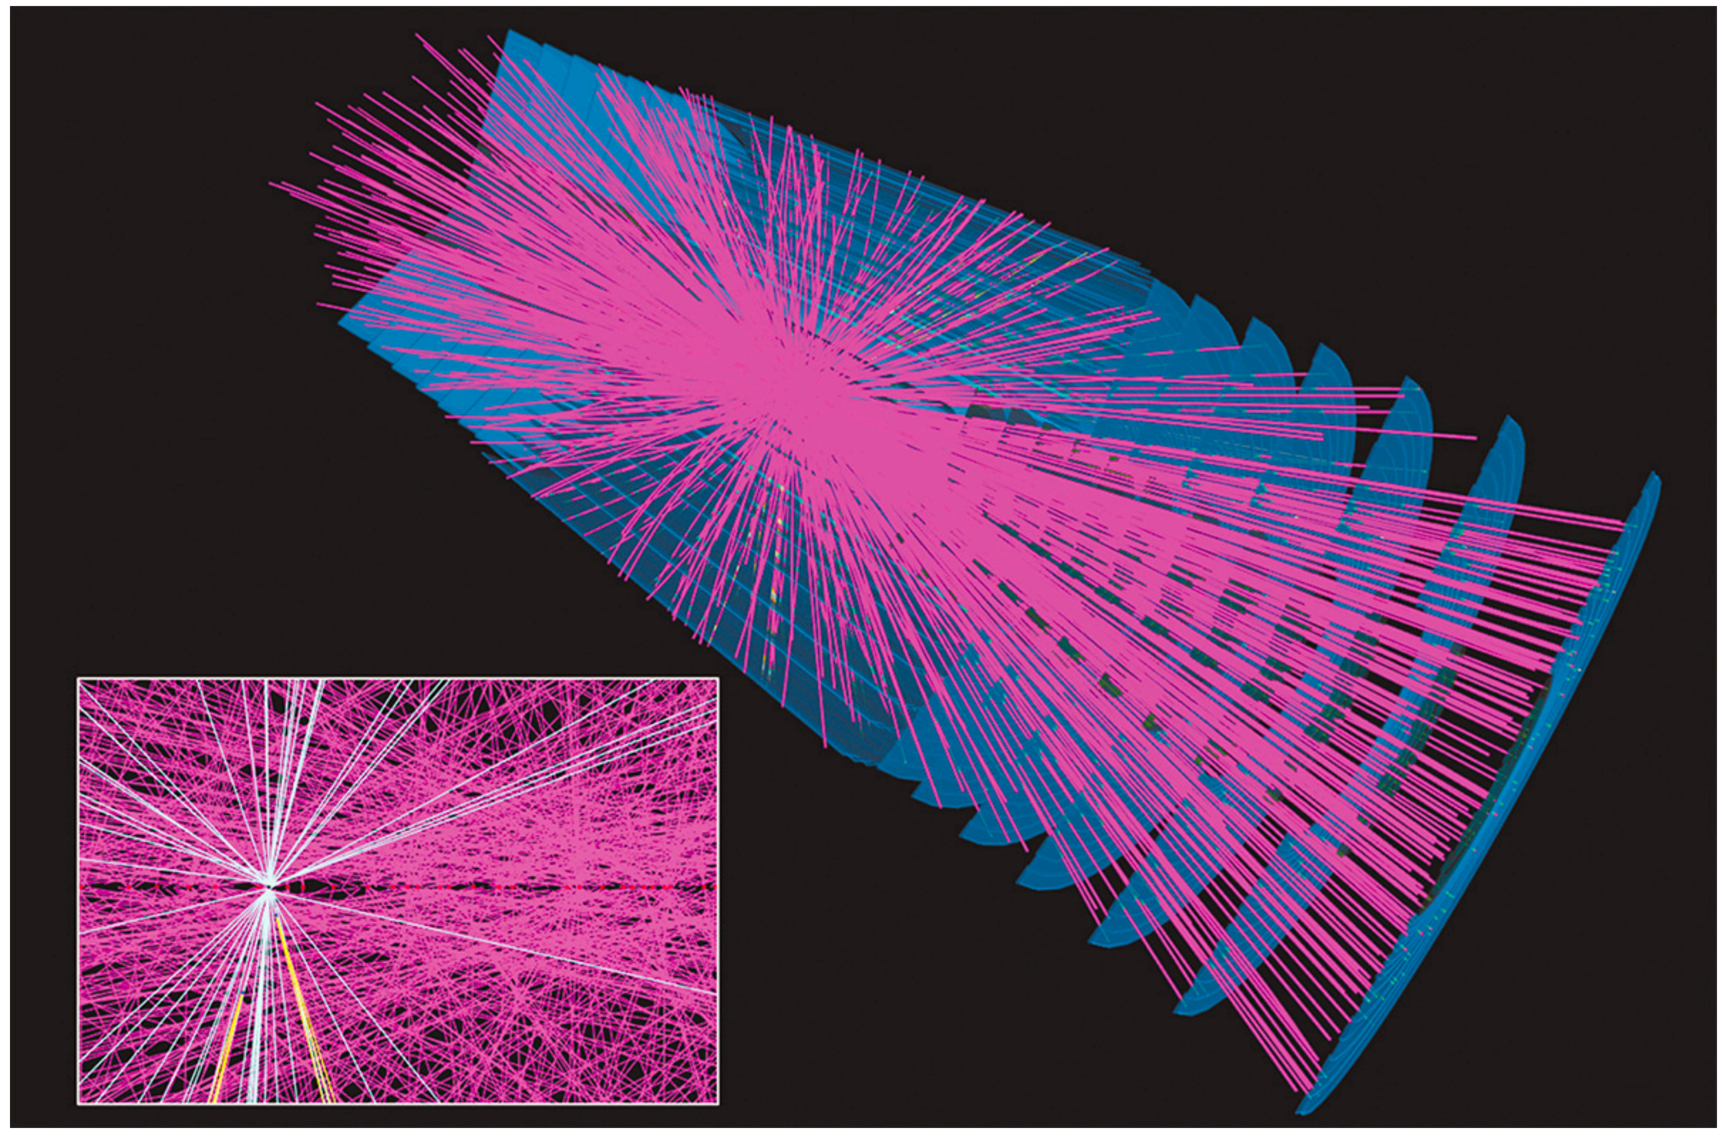

Reconstruction and Software Triggers

- Fully exploiting potential of modern CPUs requires supporting parallel data processing - “multithreading”
  - Needed for memory efficiency
  - Major rewrite of software required for this (at least on ATLAS!)

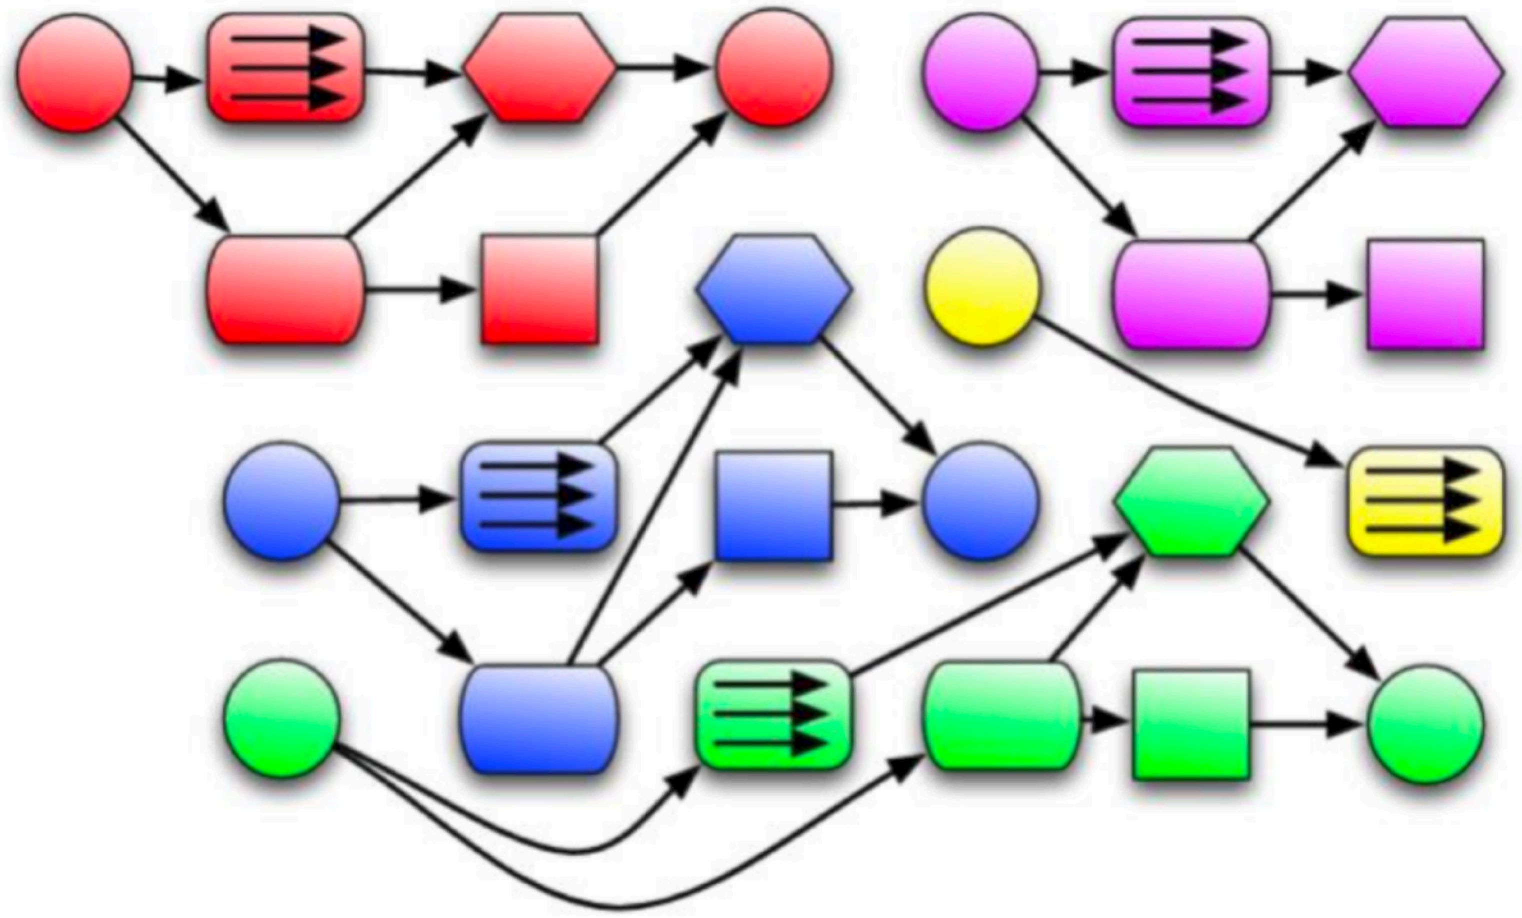

- Exploit concurrency and be as asynchronous as possible

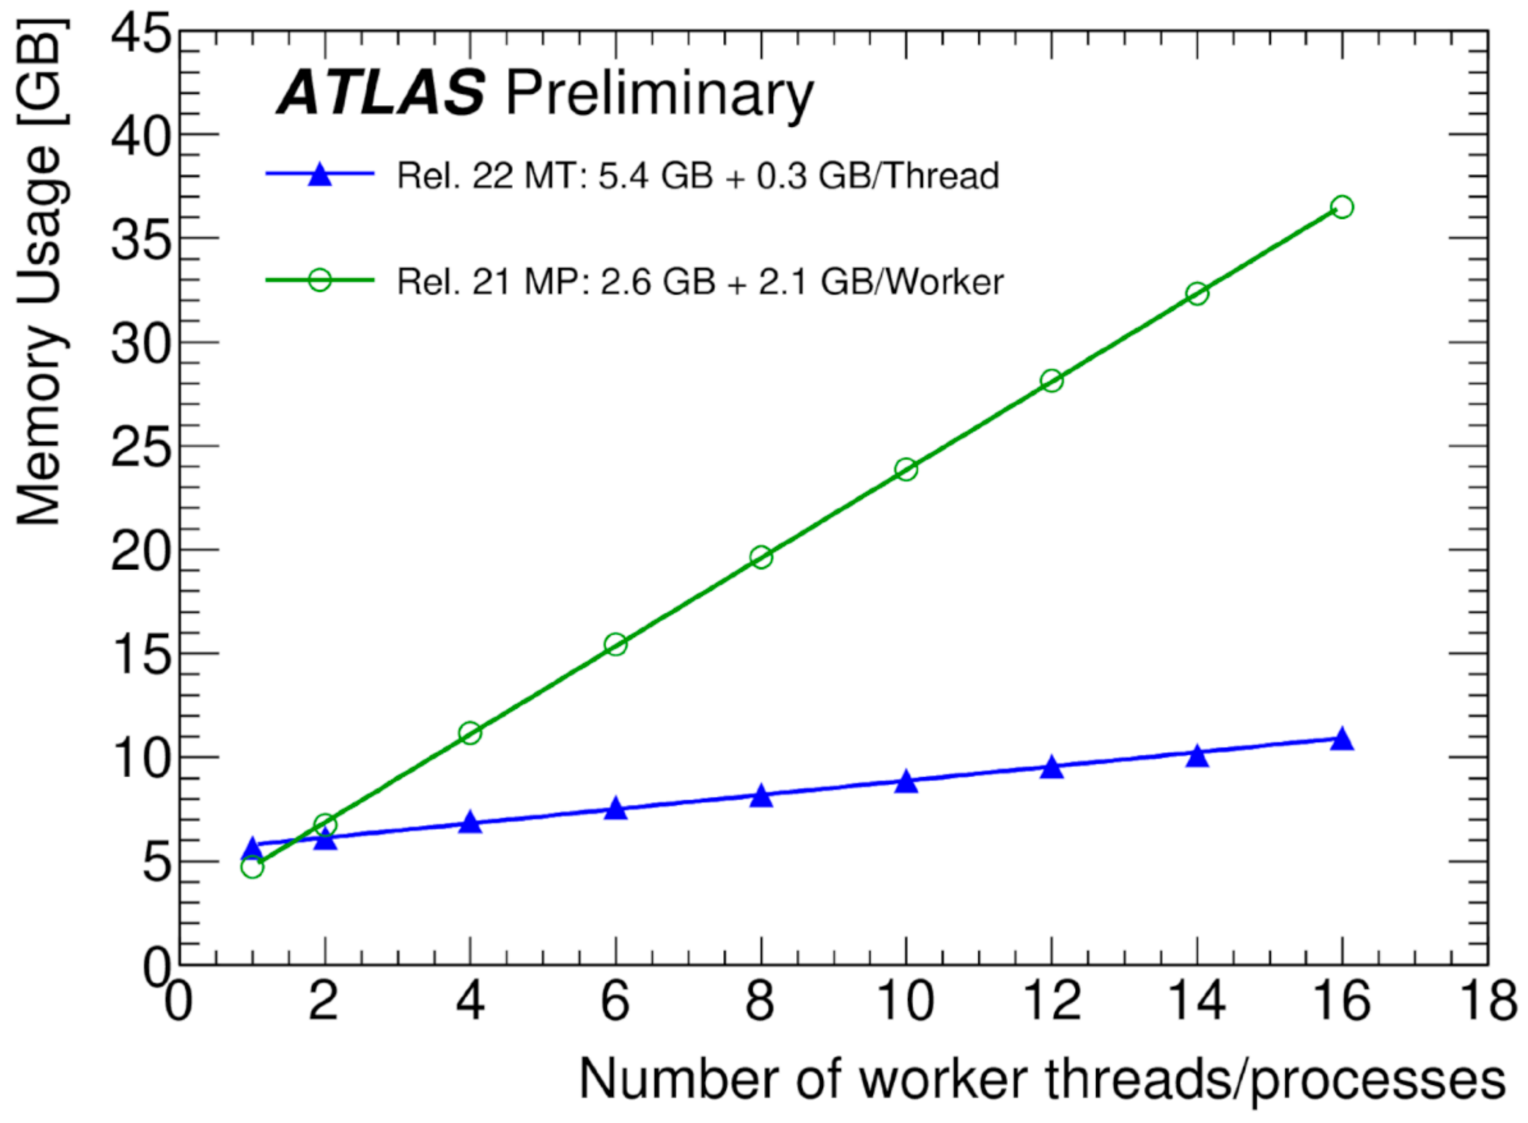

- Transfers between host and device are expensive
  - Port blocks of algorithms, even ones where gain is small

- Even the physics performance can improve when revisiting code!

Reconstruction and Software Triggers

- Hardware triggers no longer sufficient for modern experiments
  - More and more initial reconstruction needs to happen in software
  - Need to deal with tremendous rates and get sufficient discrimination
  - Lots of experimentation with rewriting code for GPUs
  - Orienting the design around the data (optimal layouts) is critical
- Real Time Analysis (HEP Version)
  - 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
- Challenges
  - Buffer for raw events is limited ('real-time' decisions)
  - Calibration needed for final output needs to be ~fast
  - Validation is very important
  - Selectively storing information requires a lot of physics inputs

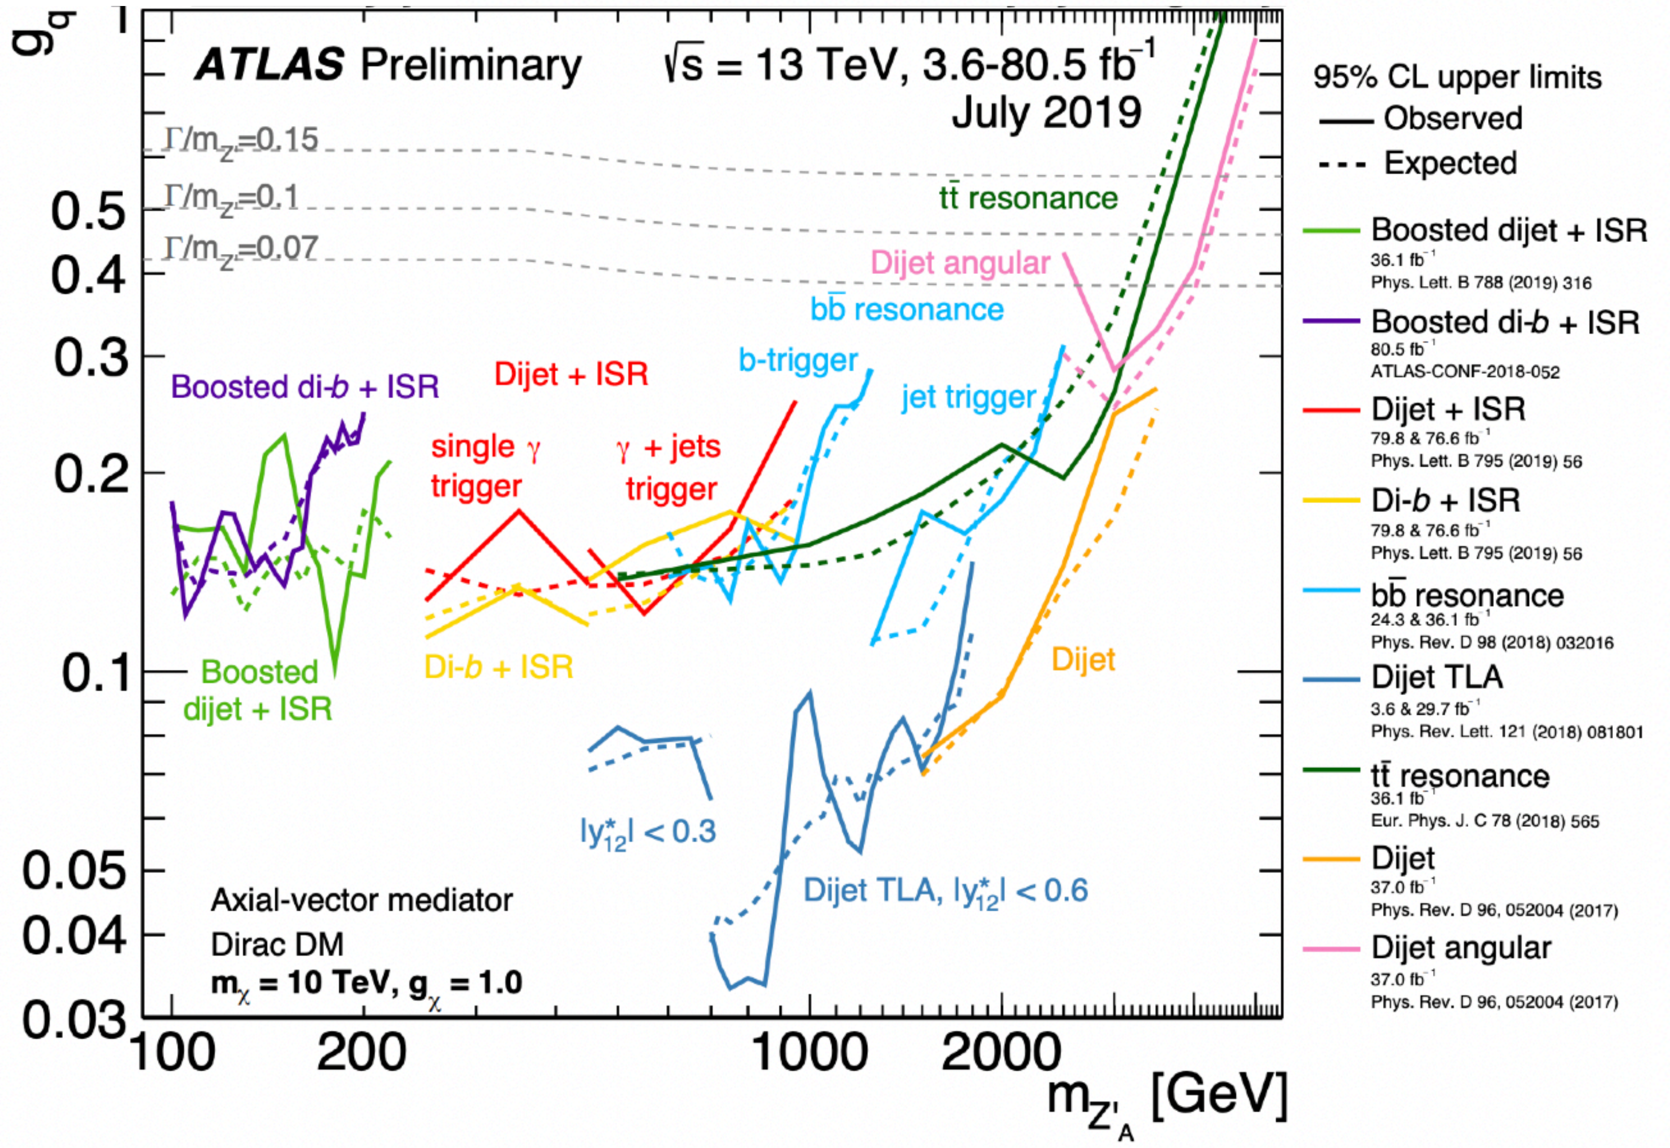

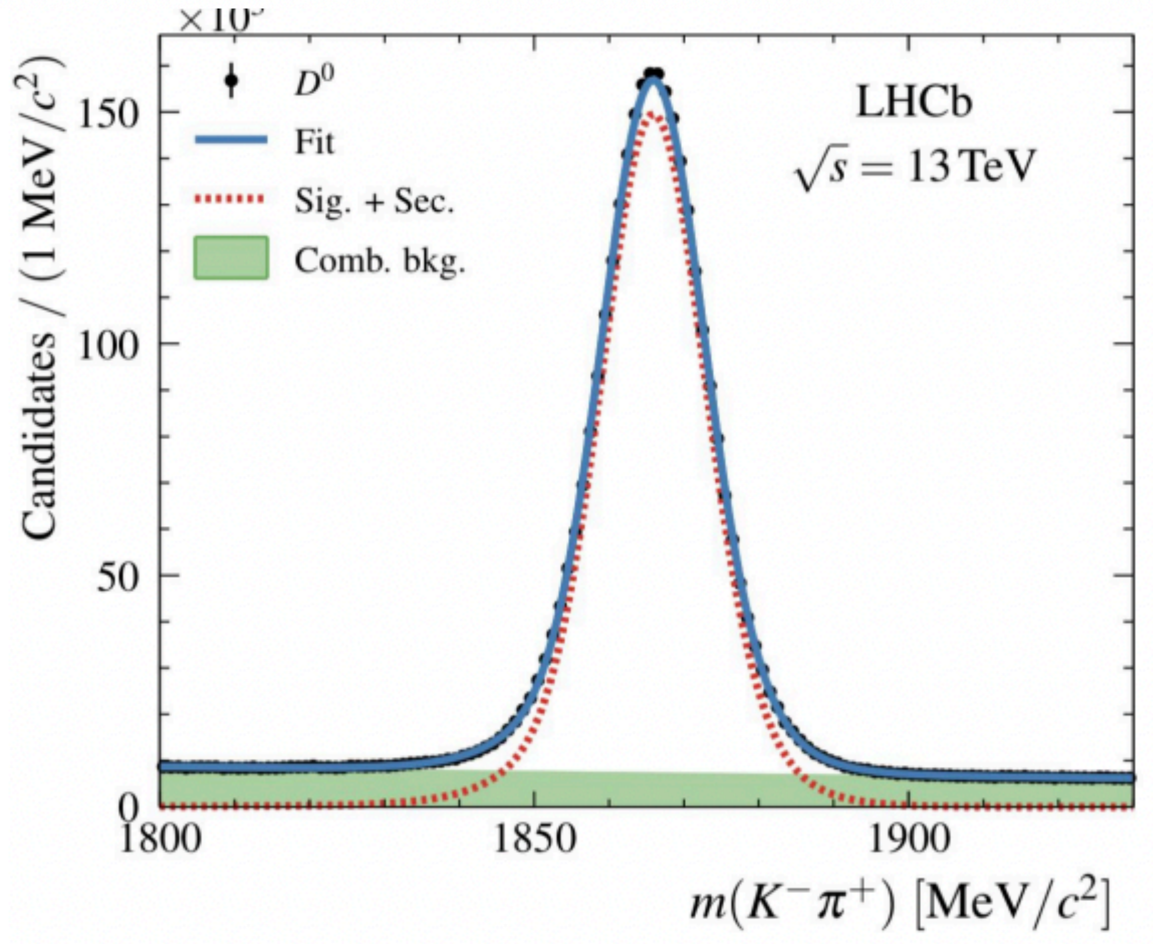

LHCb charm physics analysis using Turbo Stream (arXiv:1510.01707)

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
  - Links to 'Carpentries' being remade (US training projects)
- New areas of challenge
  - Concurrency, accelerators, data science
  - 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
  - Recognition that software is a key part of HEP now

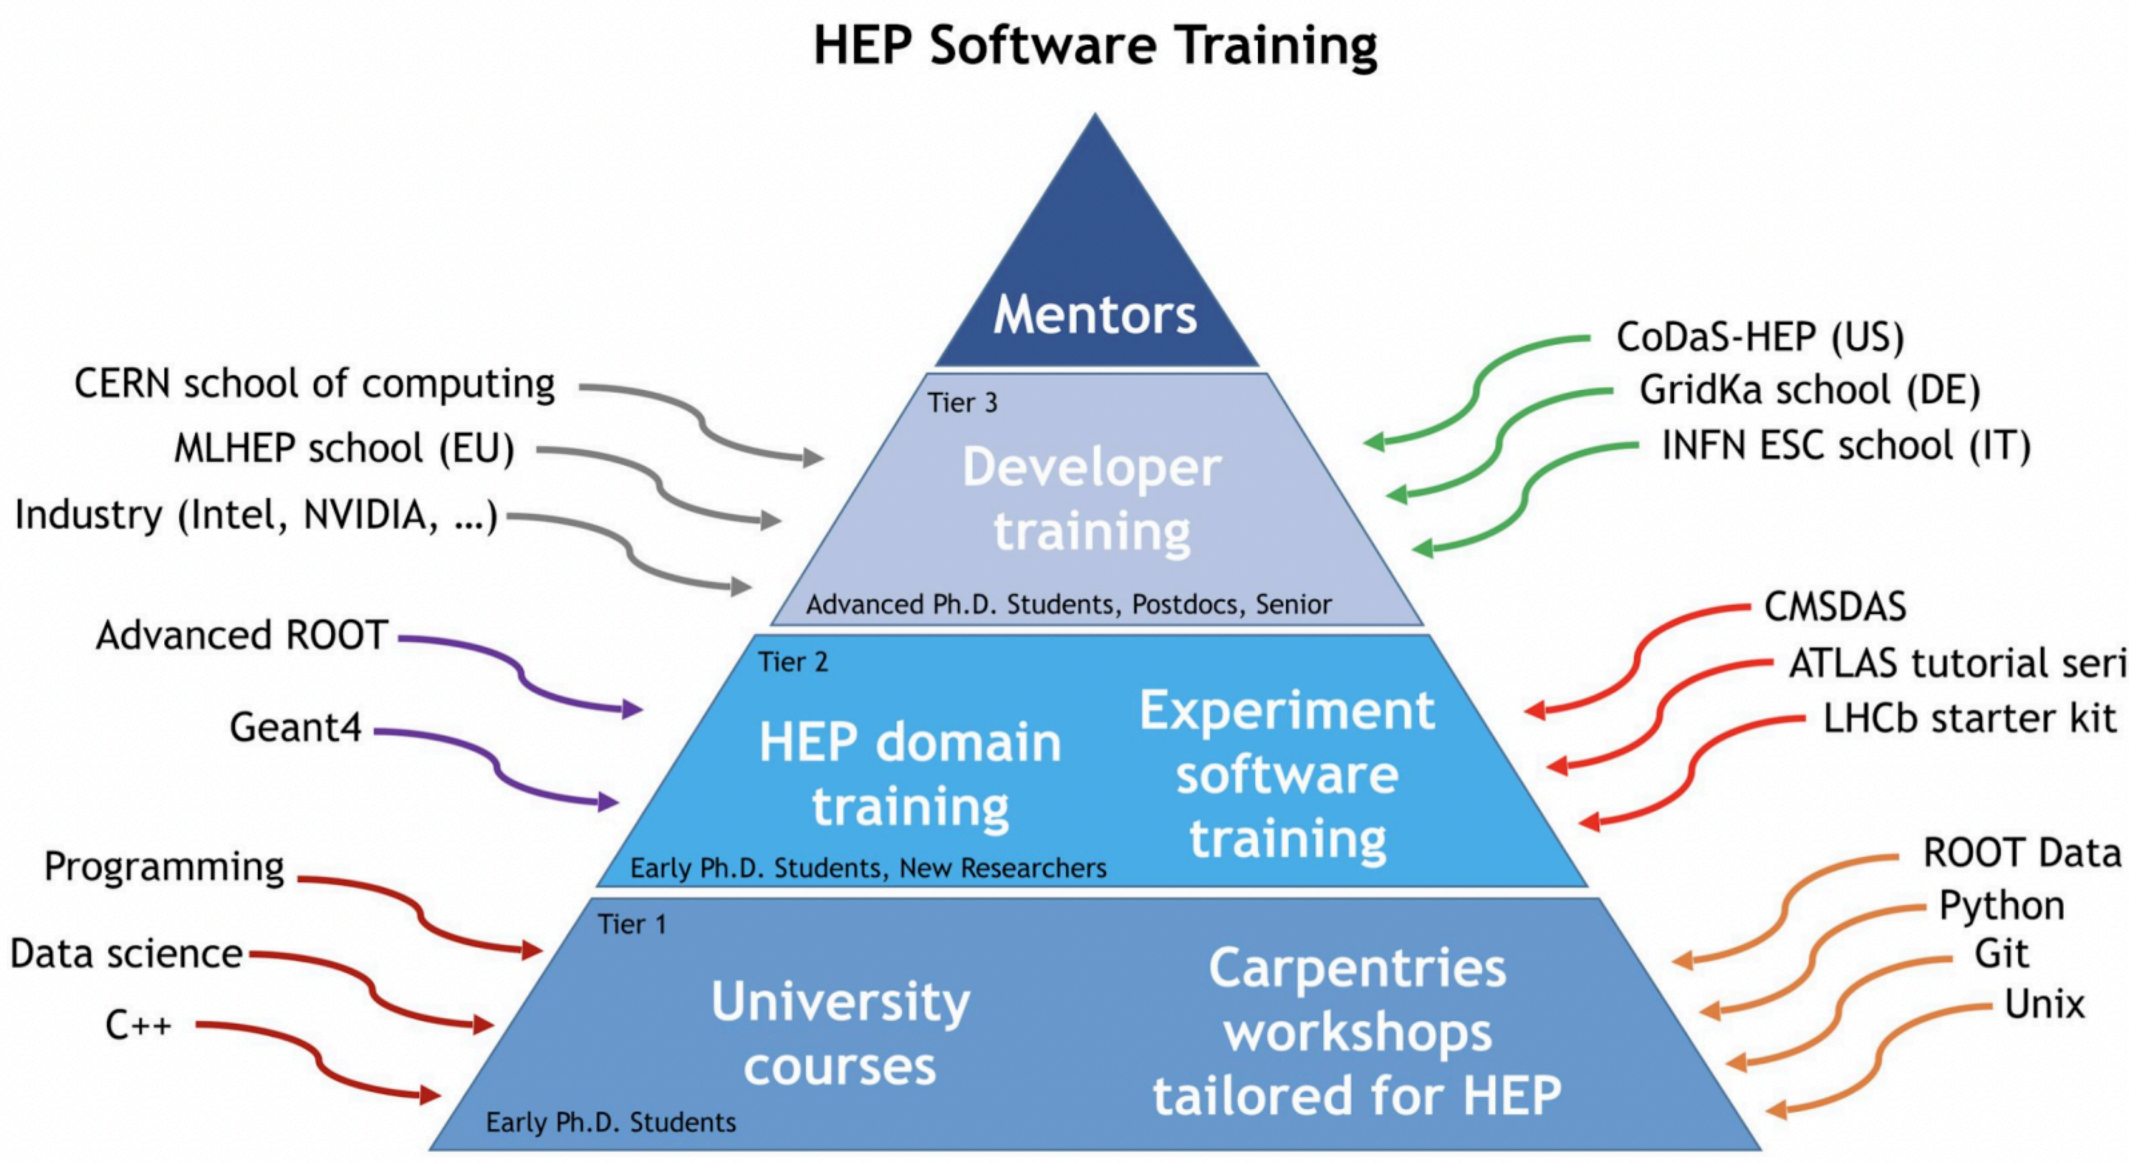

History for ECRs

- Report on the ECFA Early-Career Researchers Debate on the 2020 European Strategy Update for Particle Physics [2002.02837]

Computing and Software Software and computing activities must be recognised not only as means to do physics analyses, but as research that requires a high level of skill.

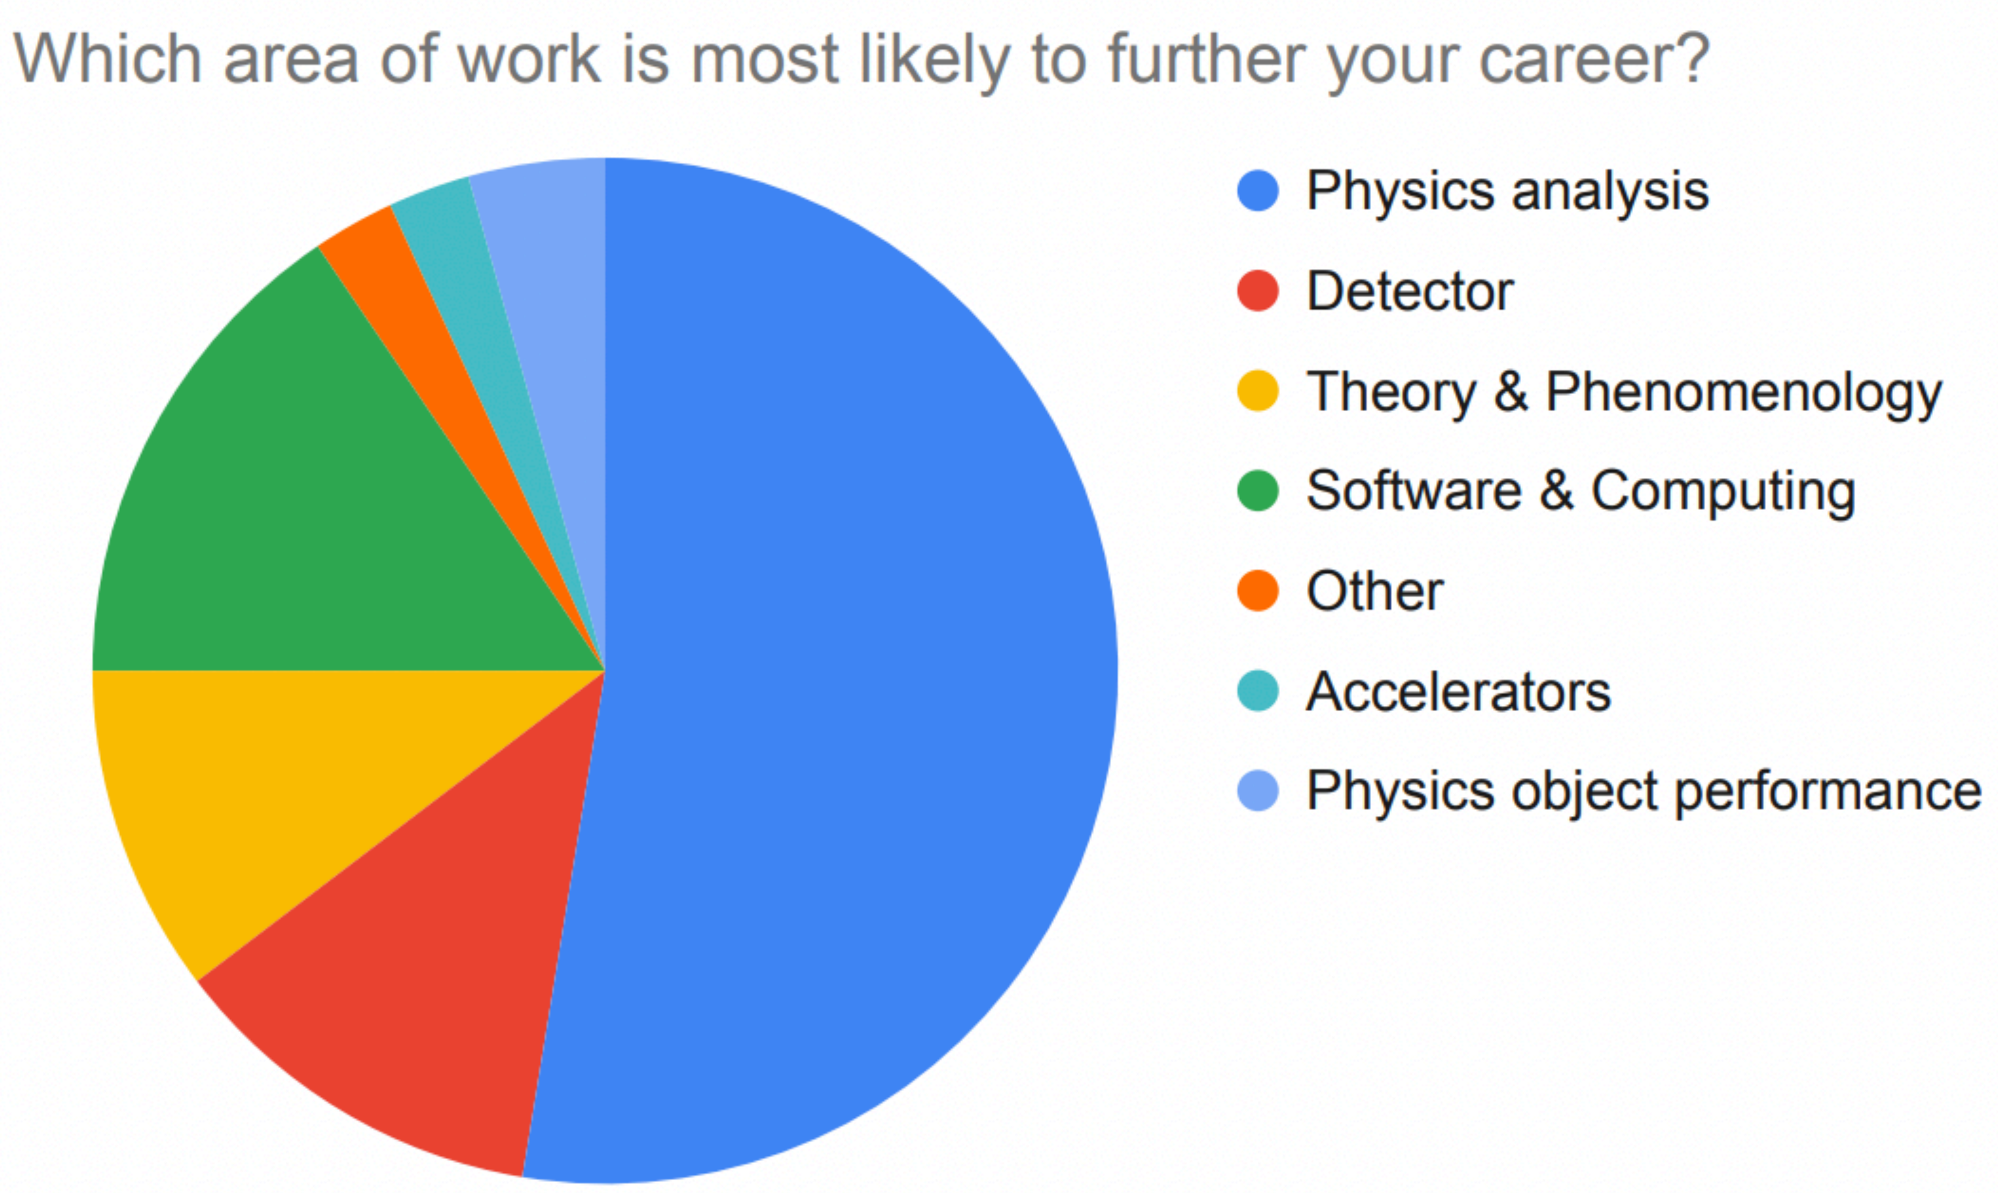

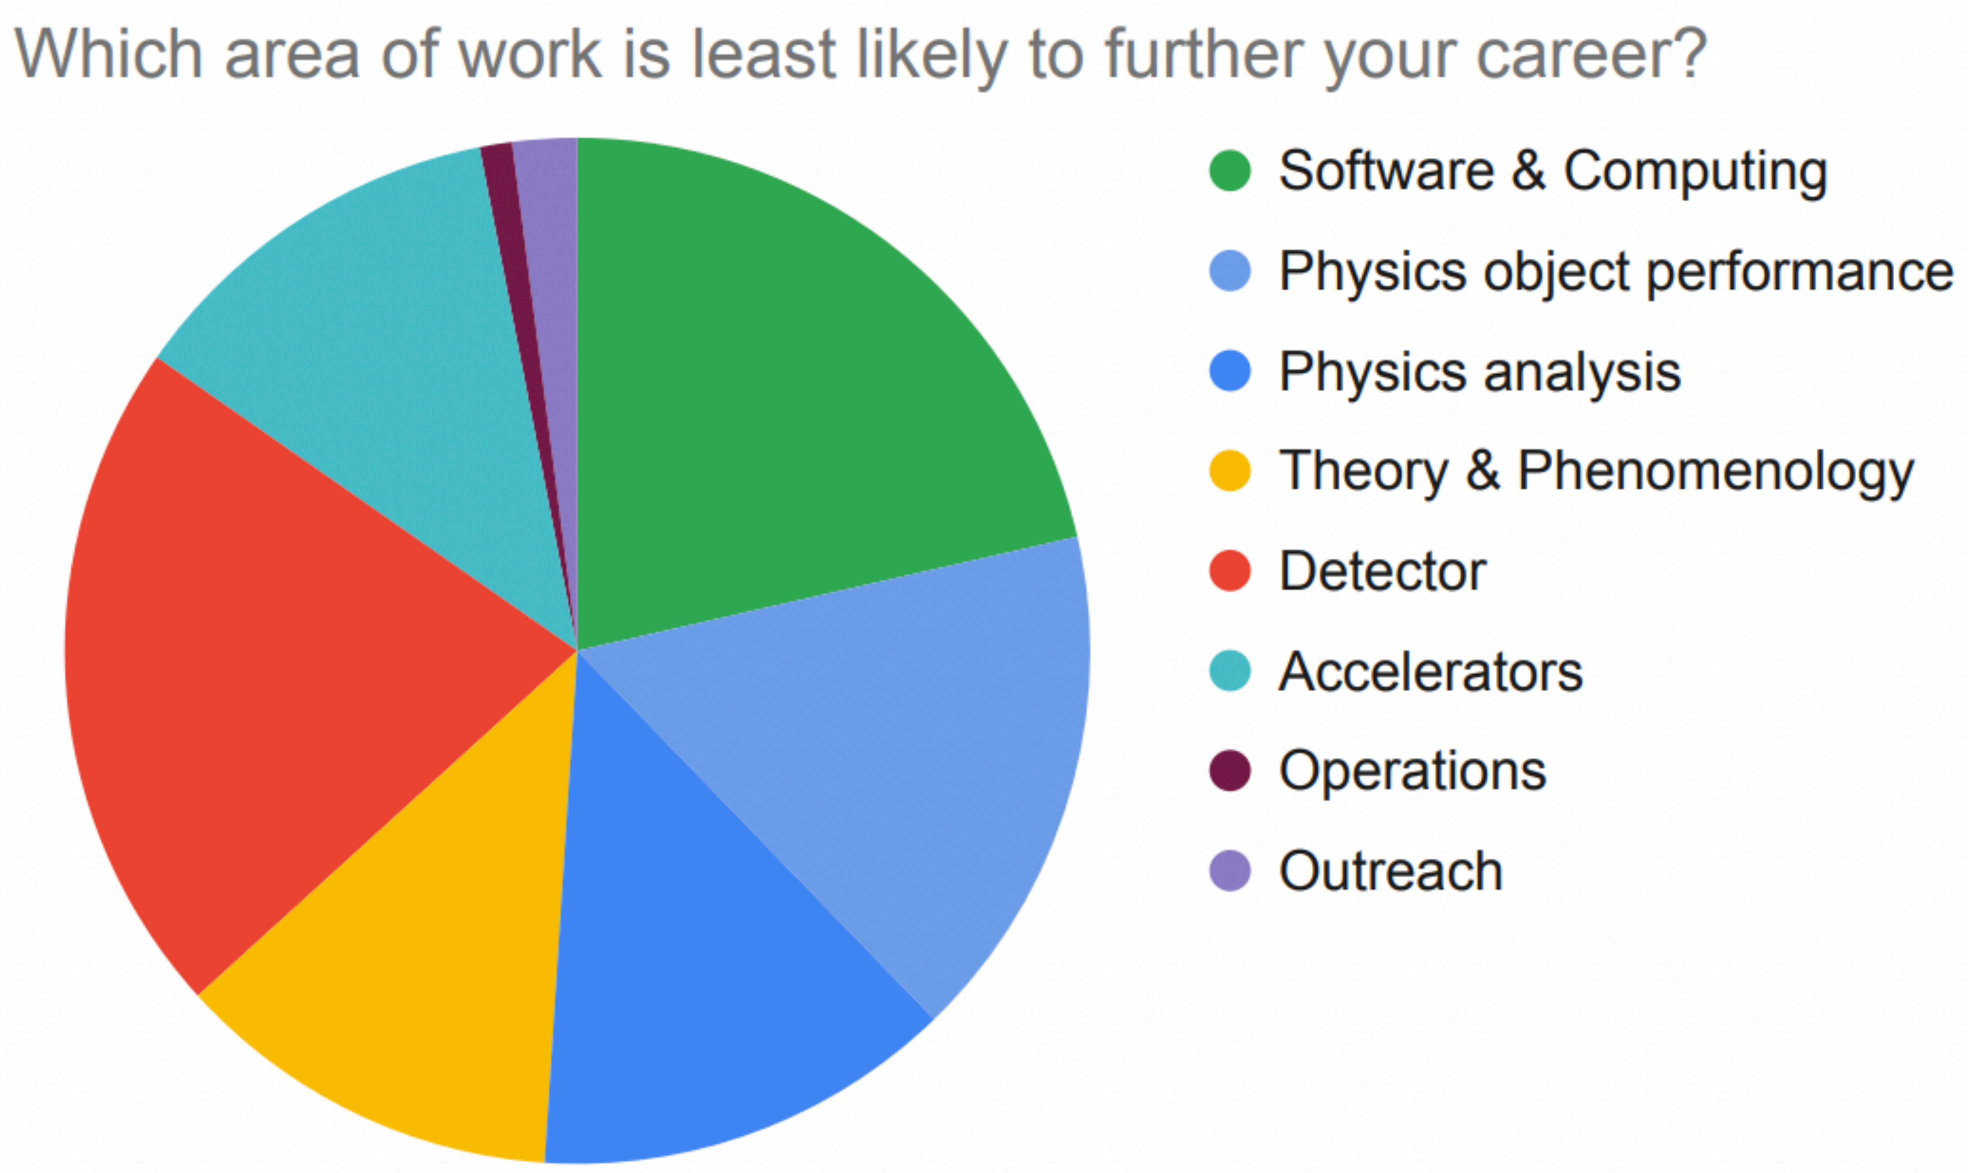

European Strategy and Snowmass

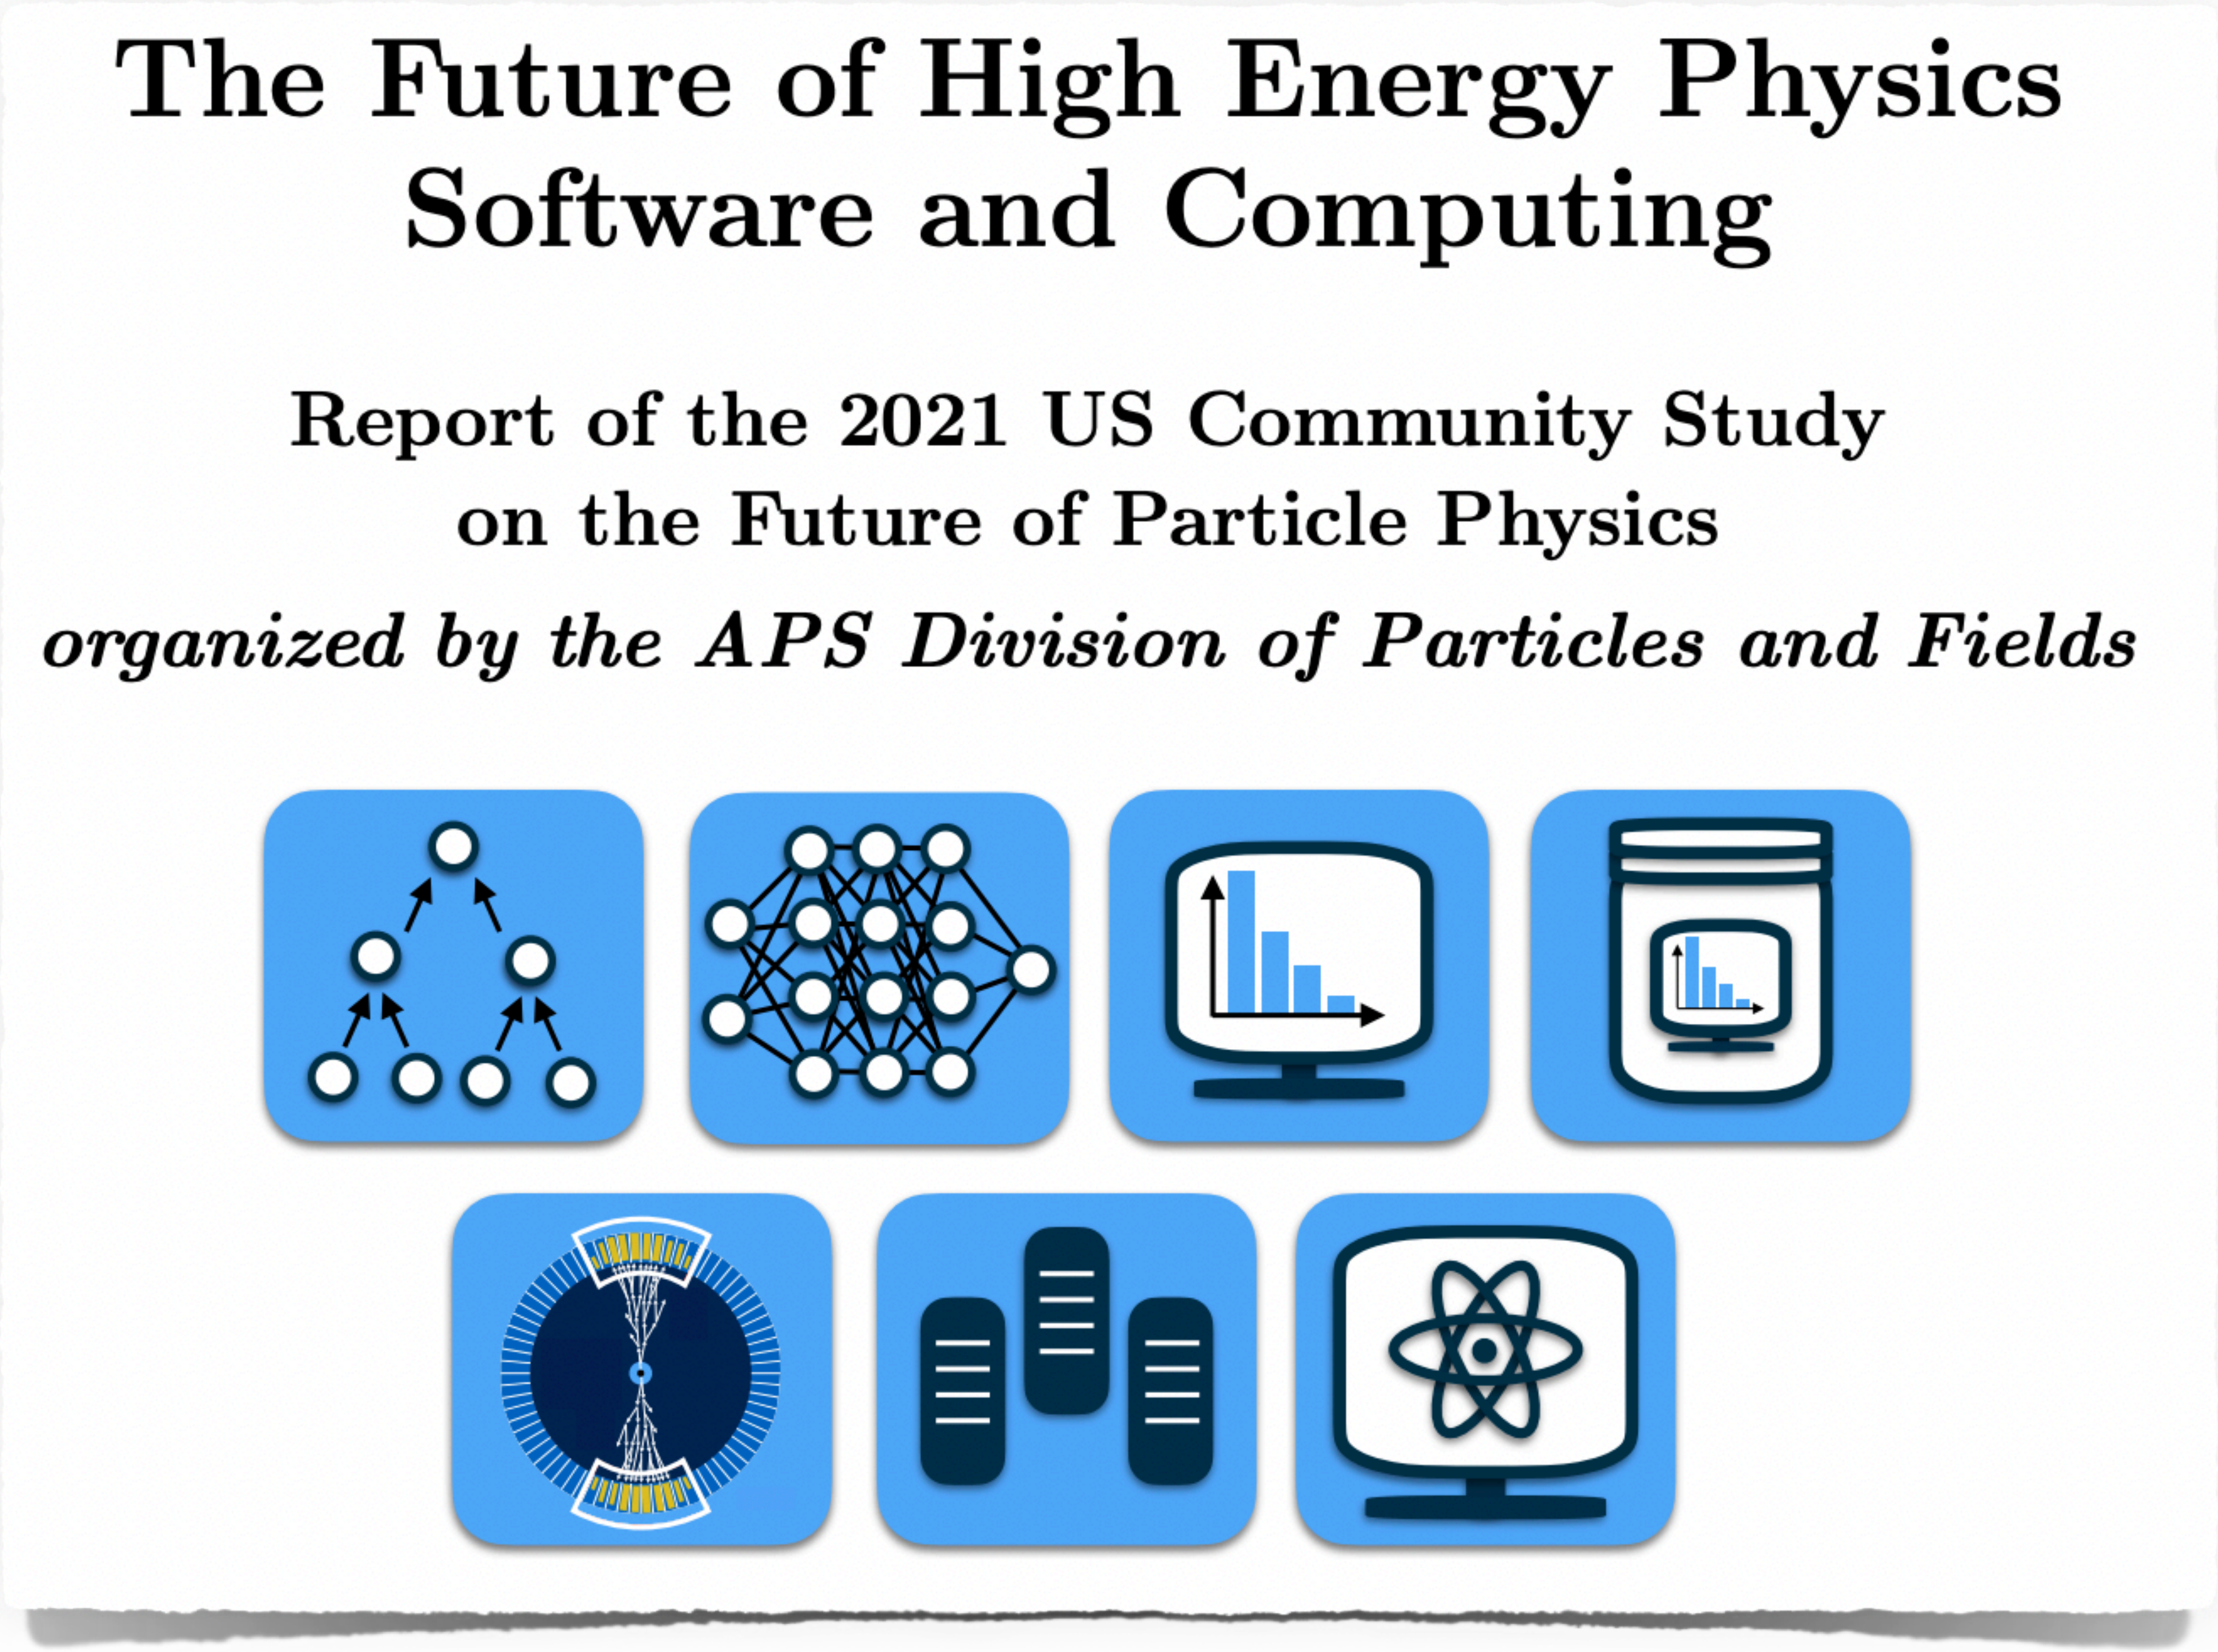

4 | Other essential scientific activities for particle physics

- 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.*

Snowmass | Executive Summary

UNIVERSITY OF SUSSEX

- Through existing, reshaped, and expanded programs, R&D efforts cutting across project or discipline boundaries should be supported from proof of concept to prototype to production.
  - **Computational HEP is a vehicle for cross-cutting R&D. Supporting research in this area at a variety of scales would be broadly impactful.**
  - Examples include S&C for theoretical calculations/generators; cosmological, accelerator, and detector modeling; machine learning methodology and hardware ecosystems; and algorithms and packages across experiment boundaries.

https://arxiv.org/abs/2210.05822

- **Scientific outcomes of particle physics experiments made possible by development of efficient S&C infrastructure.**
- **S&C are profound R&D topics in their own right - essential to sustain & enhance HEP research capabilities.**
- Need for strong community-wide coordination for S&C R&D activities
  - 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.
    - E.G. the HEP Software Foundation & European Science Cluster of Astronomy & Particle physics ESFRI (ESCAPE)
- A significant role for AI is emerging in detector design, detector operation, online data processing & data analysis.
- Important examples of developments in S&C with a large impact on particle physics results are use of multicore CPUs, multithreading and accelerators such as GPUs.
  - Event simulation, event selection and reconstruction and analysis software need to adapt to these and other emerging developments.

Snowmass | Executive Summary

- Support for computing professionals/researchers and physicists to conduct code re-engineering and adaptation will enable us to use heterogeneous resources most effectively.
  - Most HEP software runs on a single computing platform, making it difficult to use the multitude of hardware accelerators and diverse computing resources like cloud, HPC, etc.
  - To satisfy the needs of inherently serial algorithms that are still transitioning towards computing accelerators or are not cost-effective to port, an appropriate level of traditional CPU-based hardware should coexist with more powerful heterogeneous resources.

- **Achievements in S&C development with great impact should be recognised inside & outside the HEP community.**

- **The skills required to perform S&C R&D are a valuable part of the profile of a particle physicist.**
  - Recognition of this will create interesting career opportunities & to retain the engagement of researchers in these topics.

- **More experts need to be trained to address the essential needs**
  - Increasingly important to approach detector design including the impact on computing resources during operation.
  - The community needs to face the challenge of training experts that can bridge the growing gap between these activities.

Snowmass | Executive Summary

- **Strong investment in career development for HEP S&C researchers will ensure future success.**
- Sustainable efforts in computation require continual recruitment and training of the HEP workforce.
  - We need to create an environment that is inclusive, supportive, and welcoming in order to integrate diverse skill sets and experiences.
- Successful training events have been carried out through HEP experiments, institutes/organizations, and growing numbers of university courses.
  - We need to continue and grow these efforts for documentation and training at multiple levels.
- **Faculty/staff positions for physicists with expertise in S&C for HEP are scarce and person-power shortfall in this area endemic.**
  - Funding agencies can catalyze faculty-level appointments in S&C with joint appointments at national laboratories.

Snowmass | Executive Summary

- Long-term development, maintenance, and user support of essential software packages with targeted investment.
- **A new structure is needed to fund modernization, maintenance, and user support of existing tools (grants typically only fund ground-breaking R&D or development of new software).**
- Examples include
  - (i) event generators and simulation tools like Geant4 [2, 3, 4] that do not belong to a particular facility, experiment, or survey,
  - (ii) S&C tools associated with one or more experiments, and
  - (iii) data/software preservation after an experiment has ended.

7 | Environmental and societal impact

- Particle physics, with its fundamental questions and technological innovations, attracts bright young minds.
- Their education and training are crucial for the needs of the field and of society at large. For early-career researchers to thrive, the particle physics community should place strong emphasis on their supervision and training.
- **Additional measures should be taken in large collaborations to increase the recognition of individuals developing and maintaining experiments, computing and software.**
- The particle physics community commits to placing the principles of equality, diversity and inclusion at the heart of all its activities.

Environment & Sustainability

- ESU:
  - The energy efficiency of present and future accelerators, and of computing facilities, is and should remain an area requiring constant attention.
  - A detailed plan for the minimisation of environmental impact and for the saving and re-use of energy should be part of the approval process for any major project. Alternatives to travel should be explored and encouraged.

- ECR 2020:
  - Computing accounts for a few percent of worldwide CO2 emissions.
  - At CERN, computing is minor with respect to accelerator operation, but several other grid sites operate around the globe
  - The ECR delegates urge the HEP community to rethink our data processing needs in terms of sustainable computing.

Summary

- The landscape has shifted significantly in the last decade
  - Concurrency, Accelerators, Heterogeneity, Data Layout, ...
- We are constantly adapting and evolving our legacy software
  - Challenges are not just current experiments, but R&D for future detectors
- Adopting a more radical approach involves committing a lot of effort
  - It really pays off - improved software improves our physics
- We understand the main engineering issues, but not at all problems solved
  - How best to factorise from the specific technologies to avoid lock-in?
- Pyramid of skills and expertise
  - Need a lot of software engineering and physics talent
  - Address training needs
  - Long term career prospects for HEP software experts need to improve
- Huge opportunities for software to improve that we have to grasp
  - Organise around this goal and reach out to industry, software engineers, other sciences

Back-ups

HEP Software Foundation

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

- 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

HSF

HEP Software Foundation

UNIVERSITY OF SUSSEX

Contents

Data Layout and Throughput

- Original HEP C++ Event Data Models were heavily inspired by the Object Oriented paradigm
  - Deep levels of inheritance
  - Access to data through various indirections
  - Scattered objects in memory
- Lacklustre performance was ~hidden by the CPU and we survived LHC start
- In-memory data layout has been improved since then (e.g. ATLAS xAOD)
  - But still hard for the compiler to really figure out what's going on
  - Function calls non-optimal
  - Extensive use of 'internal' EDMs in particular areas, e.g. tracking
- Want to be flexible re. device transfers and offer different persistency options
  - e.g. ALICE Run3 EDM for message passing and the code generation approaches in FCC-hh PODIO EDM generator

CERN computing capacity evolution

Ian Bird

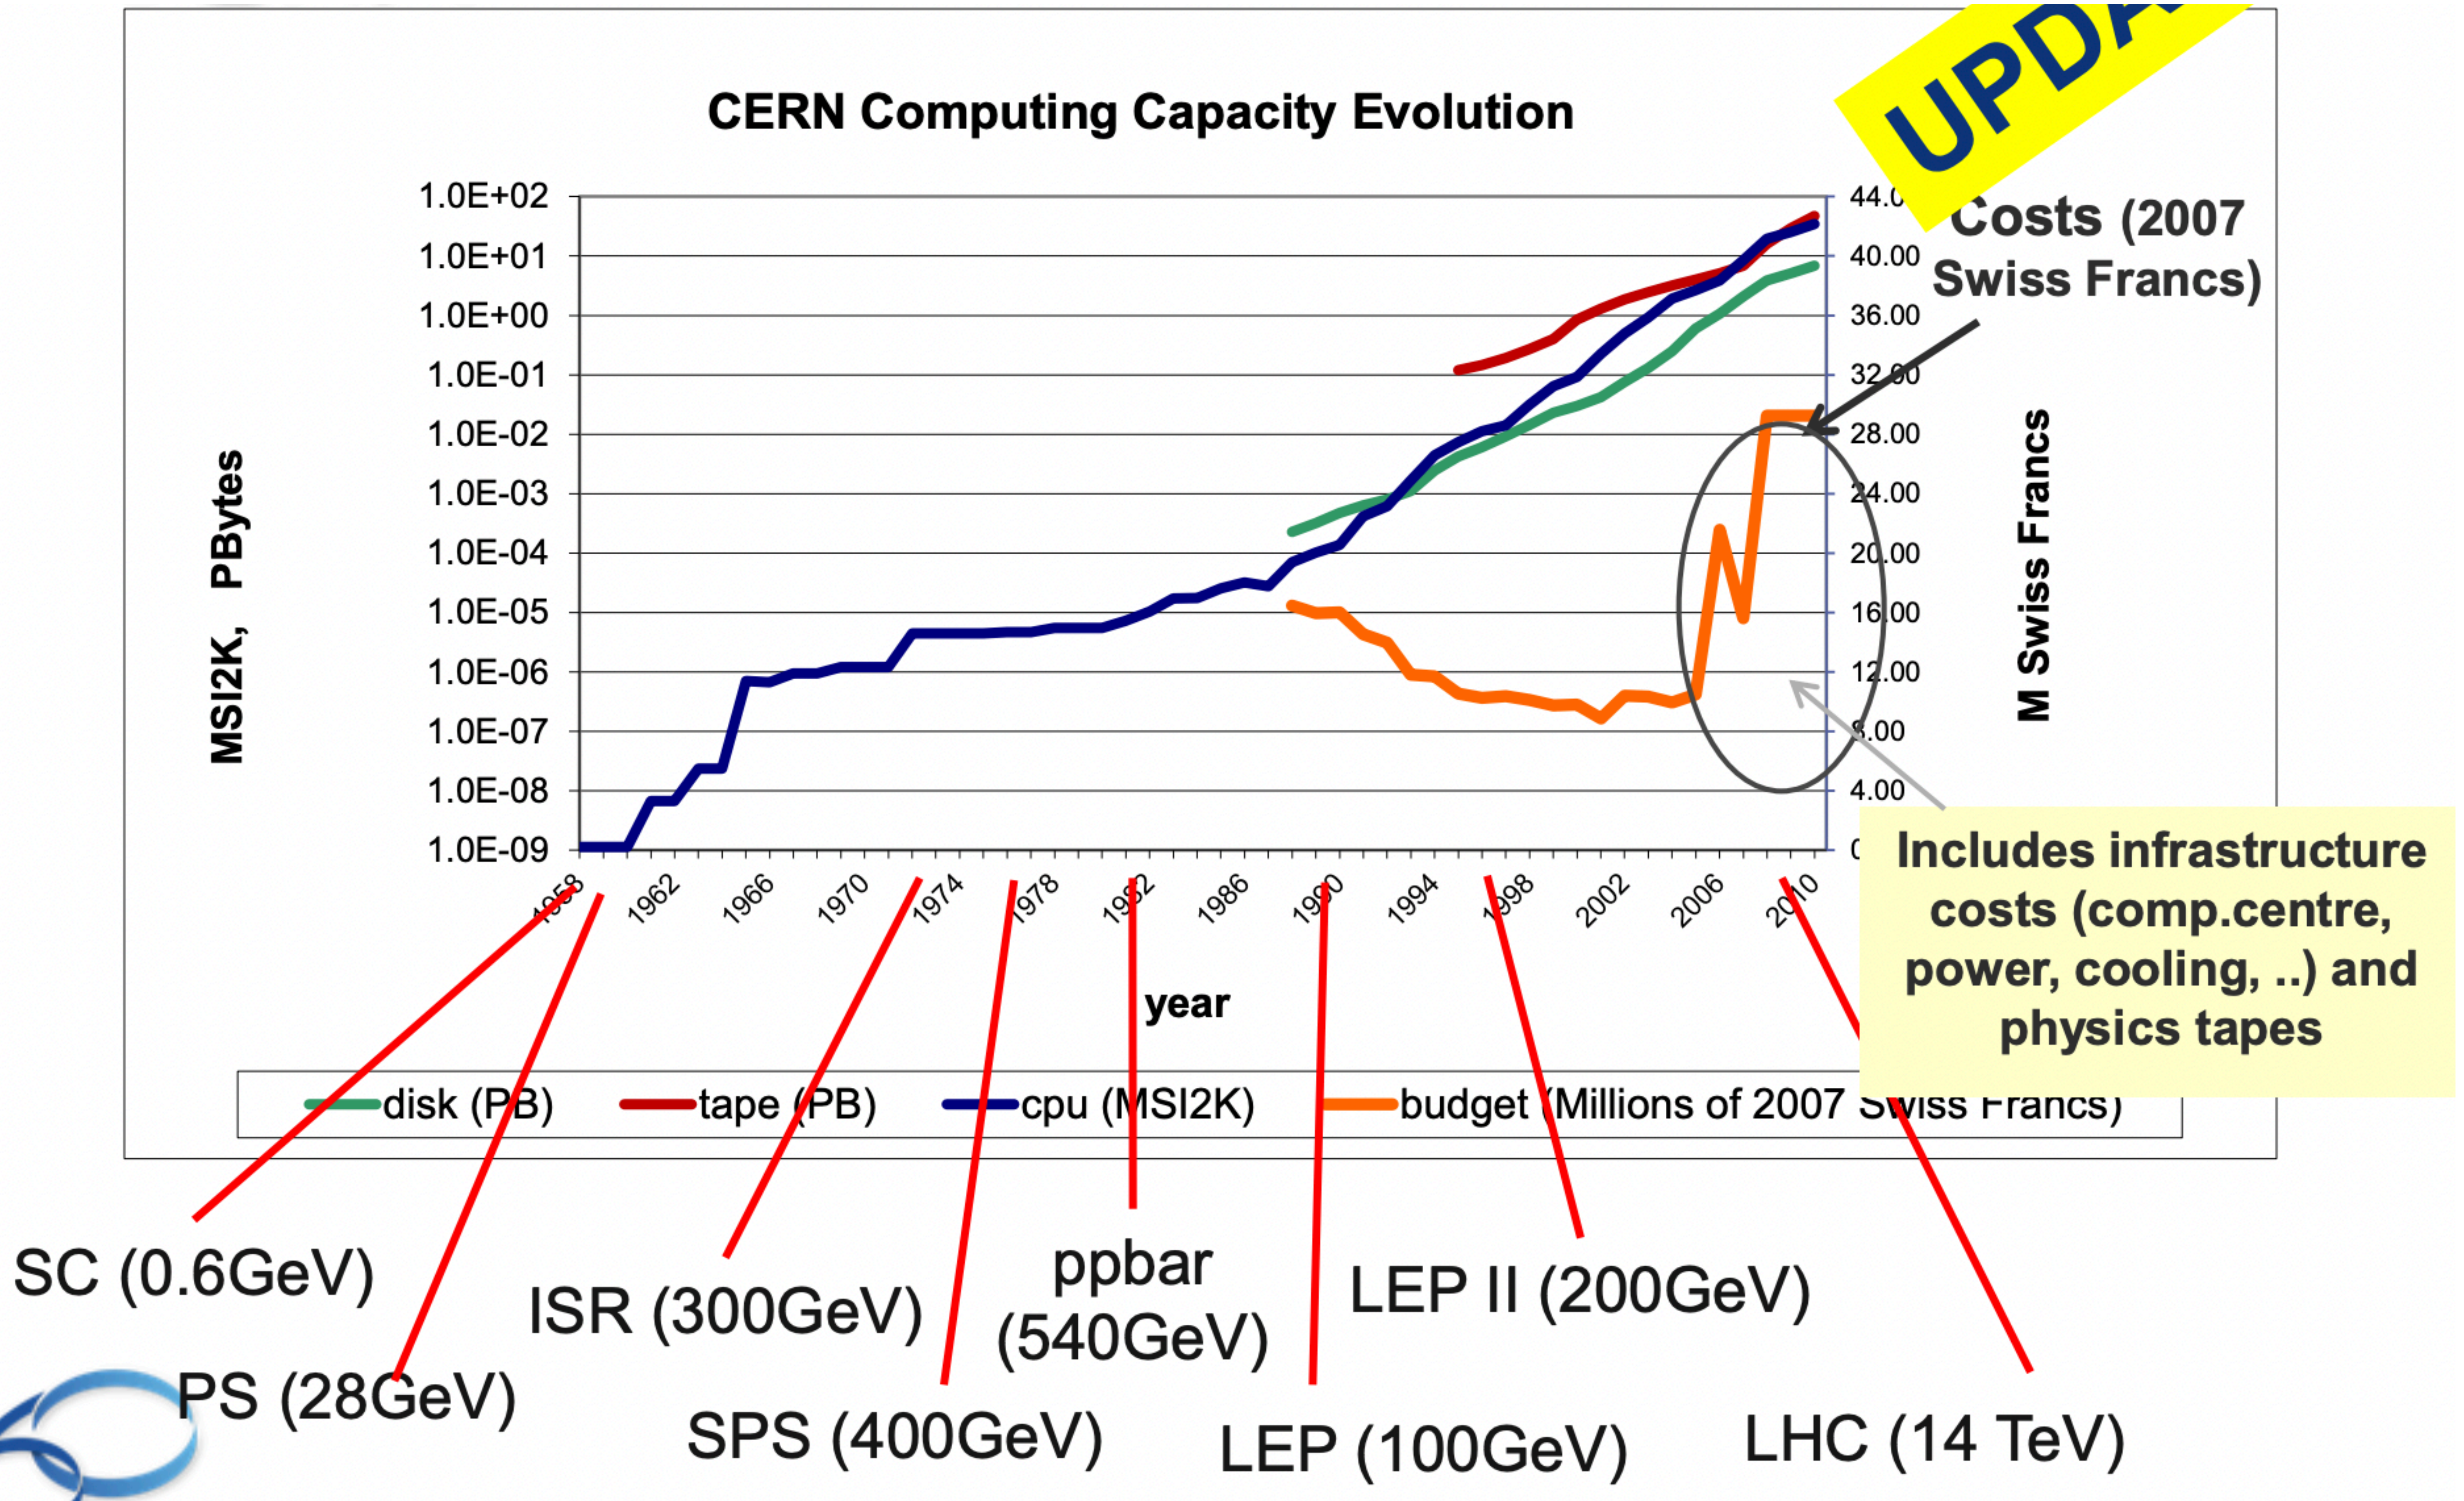

YEARS / ANS CERN

	Generation	Simulation	Reconstruction	DELPHES
Computing unit	$3.5–5.2\cdot10^{10}$	$2.6–3.9\cdot10^{6}$	$5.2–7.8\cdot10^{6}$	$2.4–3.6\cdot10^{10}$
ATLAS equivalent	$3.5–5.2\cdot10^{13}$	$2.6–3.9\cdot10^{9}$	$5.2–7.8\cdot10^{9}$	$2.4–3.6\cdot10^{13}$

https://arxiv.org/pdf/2111.10094.pdf

https://arxiv.org/abs/2002.02837

Optimal Experiment Software

- The Golden Roles
  - Orienting the design around the data (optimal layouts) is critical
  - Bulk data together and exploit concurrency where ever possible
  - Be as asynchronous as possible ○ Framework should hide latency
  - Transfers between host and device are expensive
    - Port blocks of algorithms, even ones where gain is small
  - The physics performance can improve when revisiting code!
    - We have a lot of legacy; revisiting the code oriented to the primary goal simplifies and improves maintainability