PRINCETON UNIVERSITY

Future Evolution of Computing/Reconstruction in High Energy Physics

David Lange

Princeton University

June 8, 2018

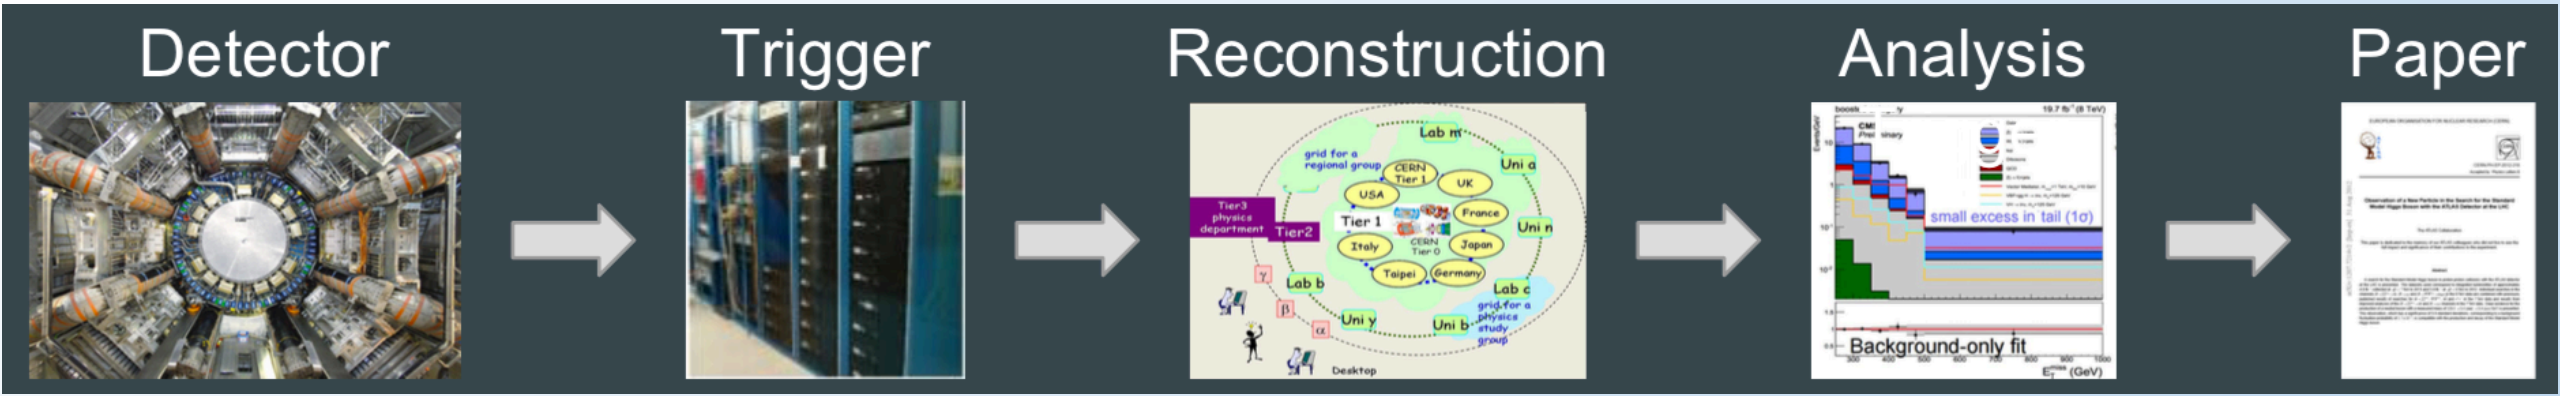

LHC experiments depend on a massive distributed computing infrastructure.

- **Software is a critical part of our physics production pipeline**, from triggering all the way to analysis and final plots as well as simulation
  - Millions of lines of code per experiment supported by numerous other software packages (ROOT, Geant, event generators, etc)

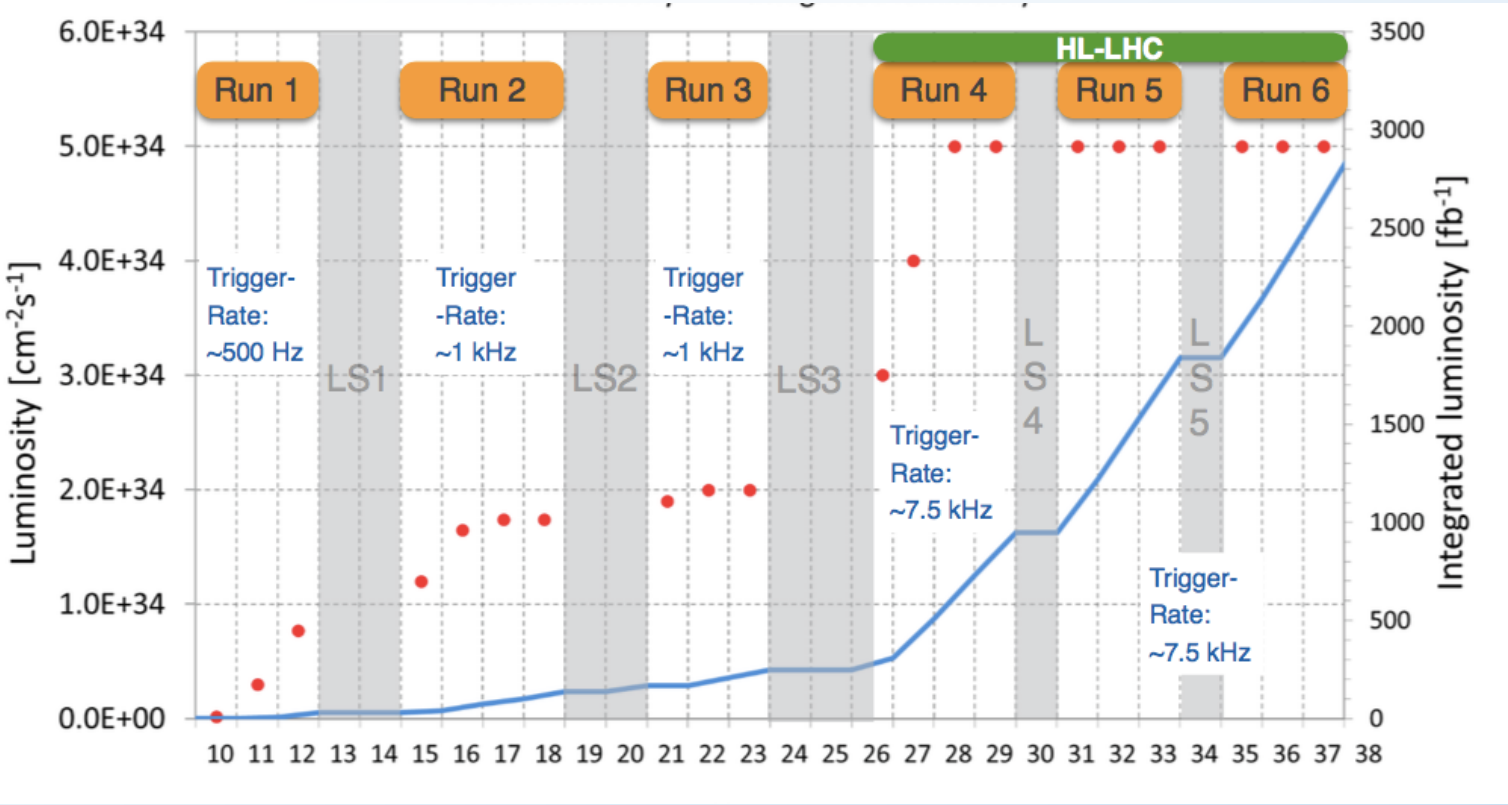

- LHC experiments continuously use about 600k CPU cores and have around 400PB of data stored on disk and 600PB on tape
- HL-LHC brings a dramatic increase in event rates

Physics challenges to CMS/ATLAS software+computing during HL-LHC

- Increase luminosity means more interactions per bunch crossing. This introduces numerous issues:
  1. **Higher detector occupancy**: More sophisticated detector technologies and higher channel count
  2. **Trigger**: Higher rates needed to preserve current physics reach; Use capabilities earlier in the processing chain (e.g., tracking at level-1 trigger) and real-time analysis concepts
  3. **Particle reconstruction**: More difficult to separate patterns means physics impact (eg efficiency vs fake tradeoff) and technical performance (e.g., CPU time) challenges
  4. **Analysis sensitivity**: Searches for lower cross section processes demand higher precision and most robust reconstructed data

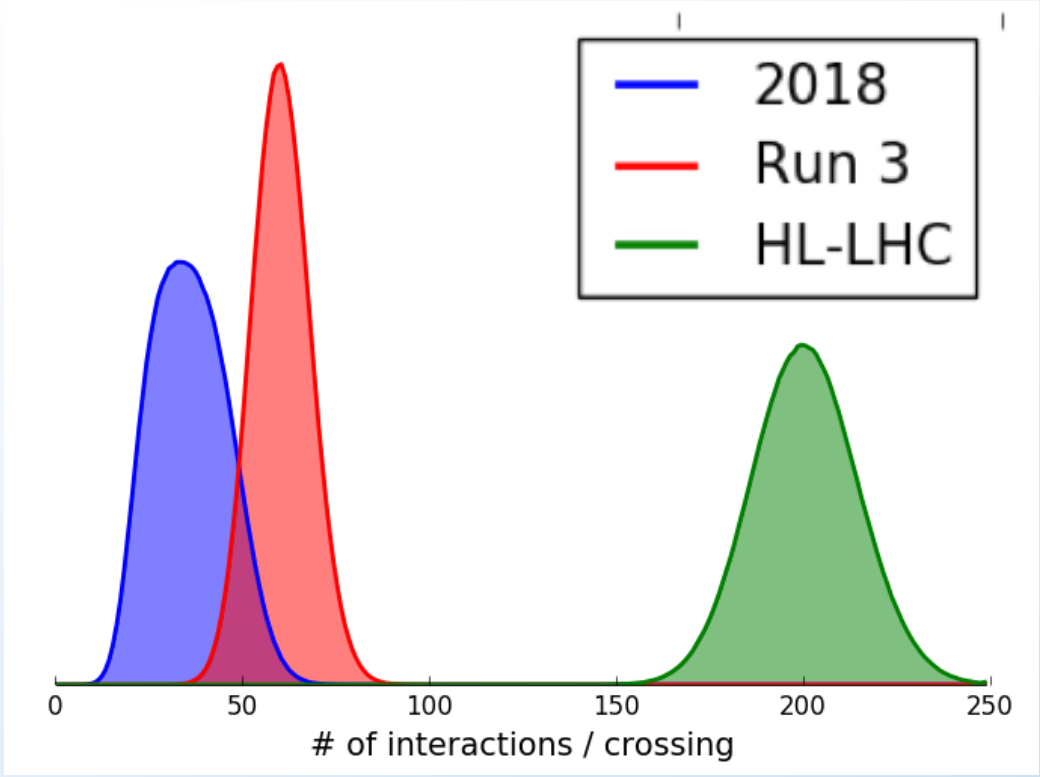

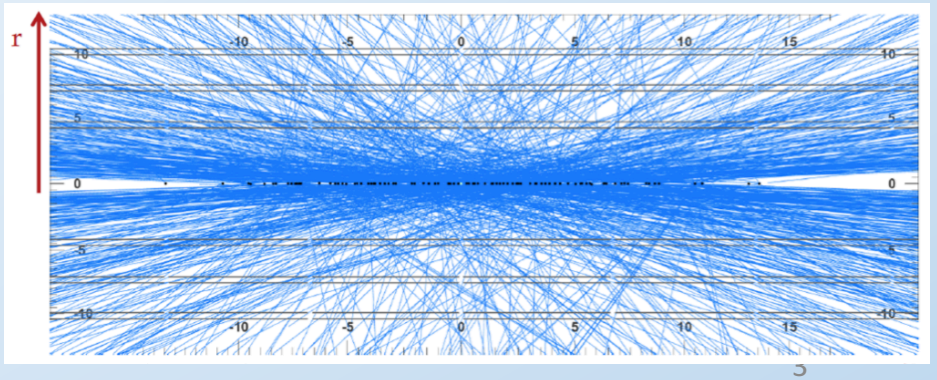

Physics challenges to LHC software + computing

- Software trigger (HLT) and offline event reconstruction algorithms face similar challenges:
  - ~4x increase in pileup
  - ~10x higher projected input rates

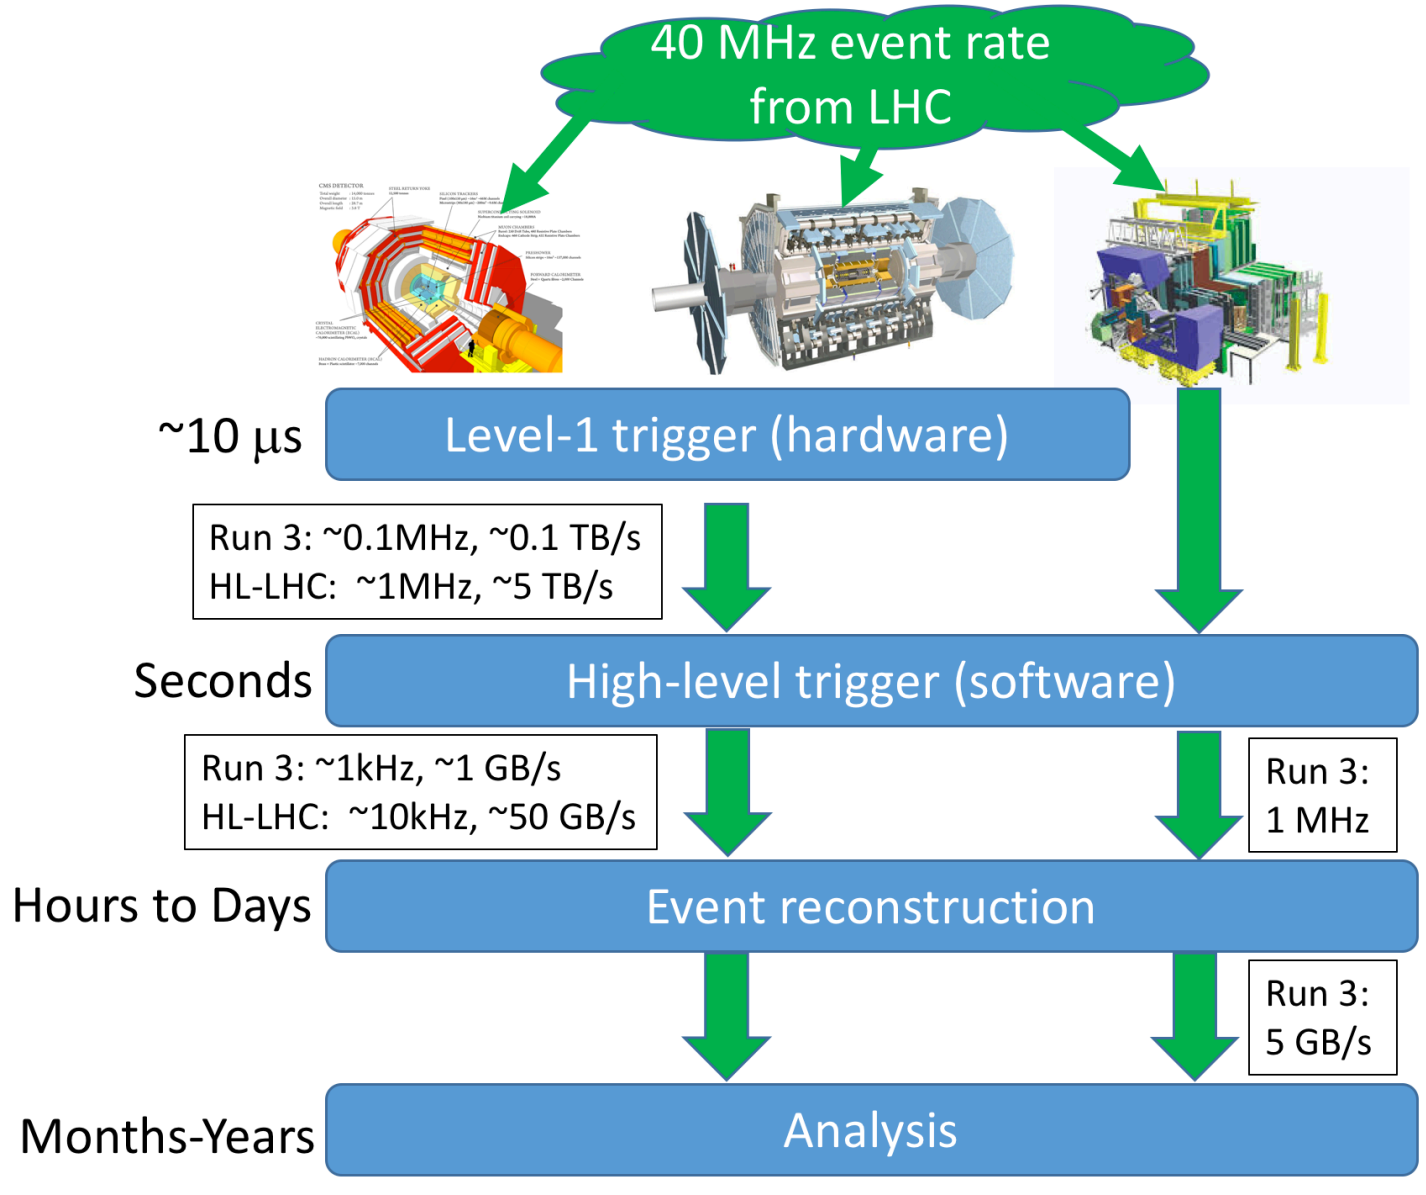

A large increase in resources over a “flat budget” would be needed if we simply scale how these algorithms and facilities work today.

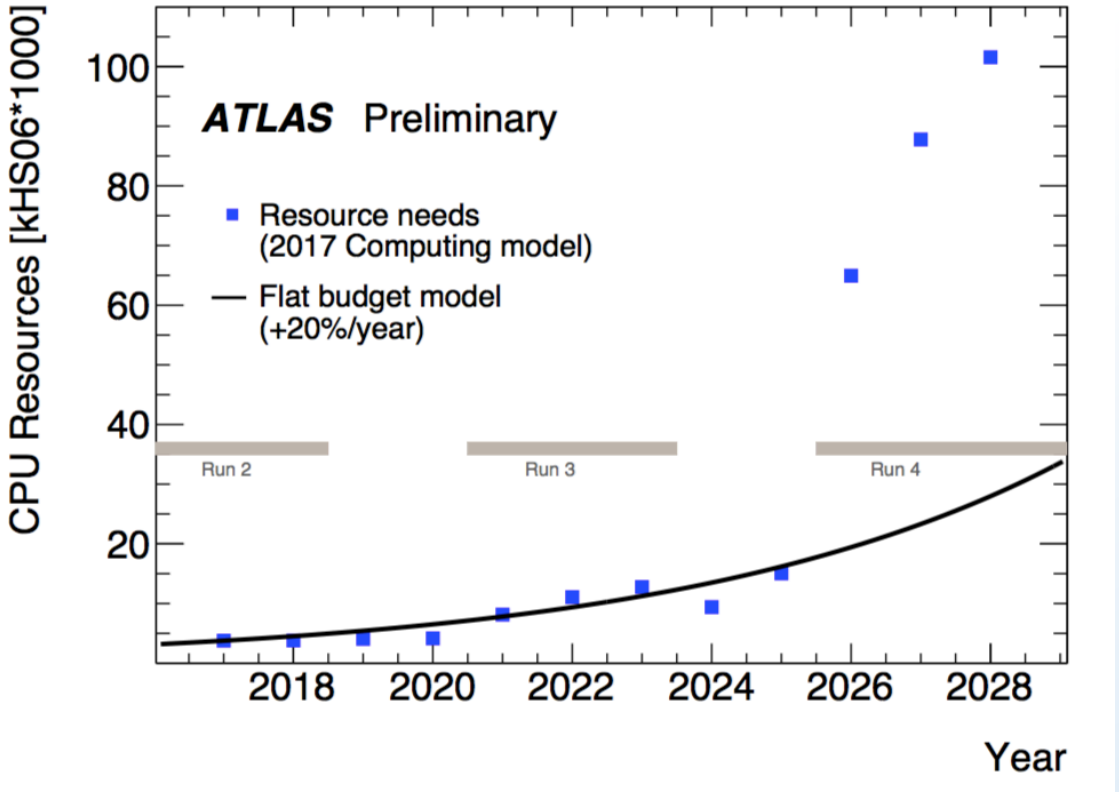

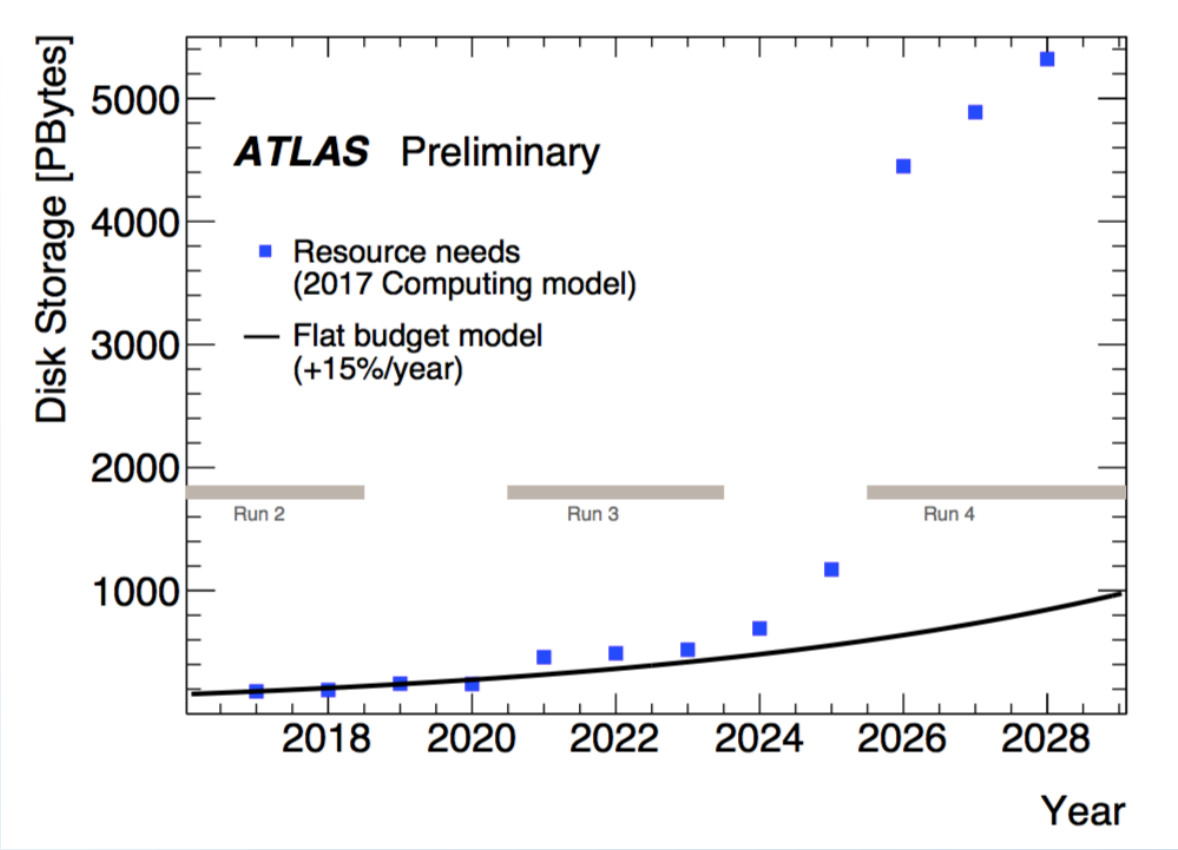

Signal dominated data rates: Experiments move away from hardware triggers and towards real-time analysis ideas

- LHCb will implement a fully software-based trigger for Run 3.
- “Real-time analysis”: Perform analysis directly on the output of the trigger rather than a truly offline reconstruction
  - Keep many more events: Drop the RAW data to gain large factor in data reduction
  - All stages of processing must be ready from the start
    - Requires final calibrations ready in hours
    - No going back to rerun reconstruction algorithms

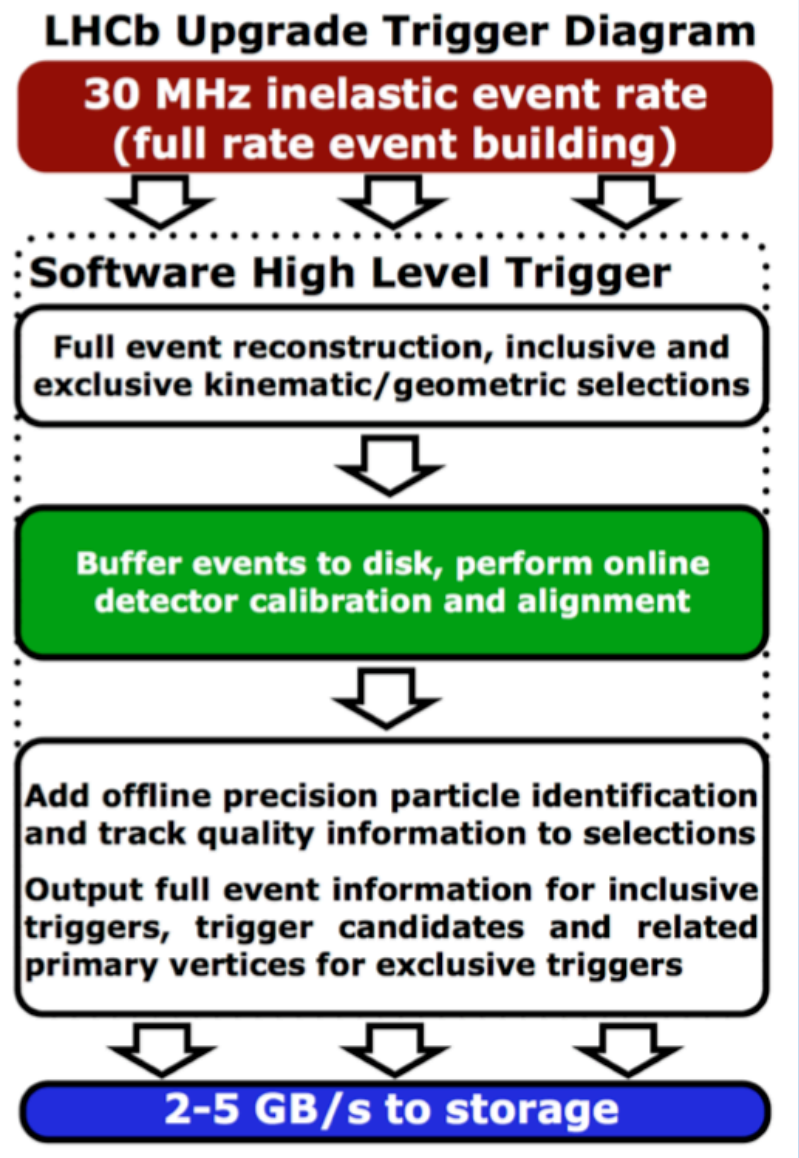

Technology challenges to LHC software and computing

End of processor speedup

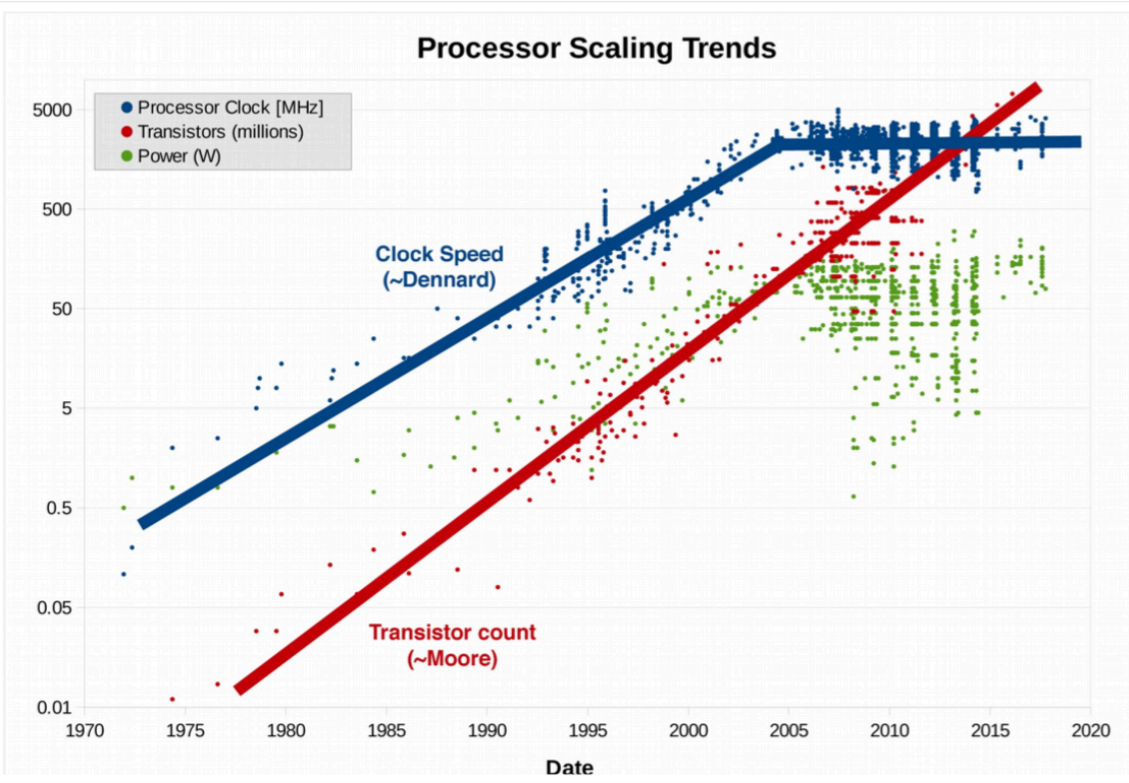

Towards many-core computing

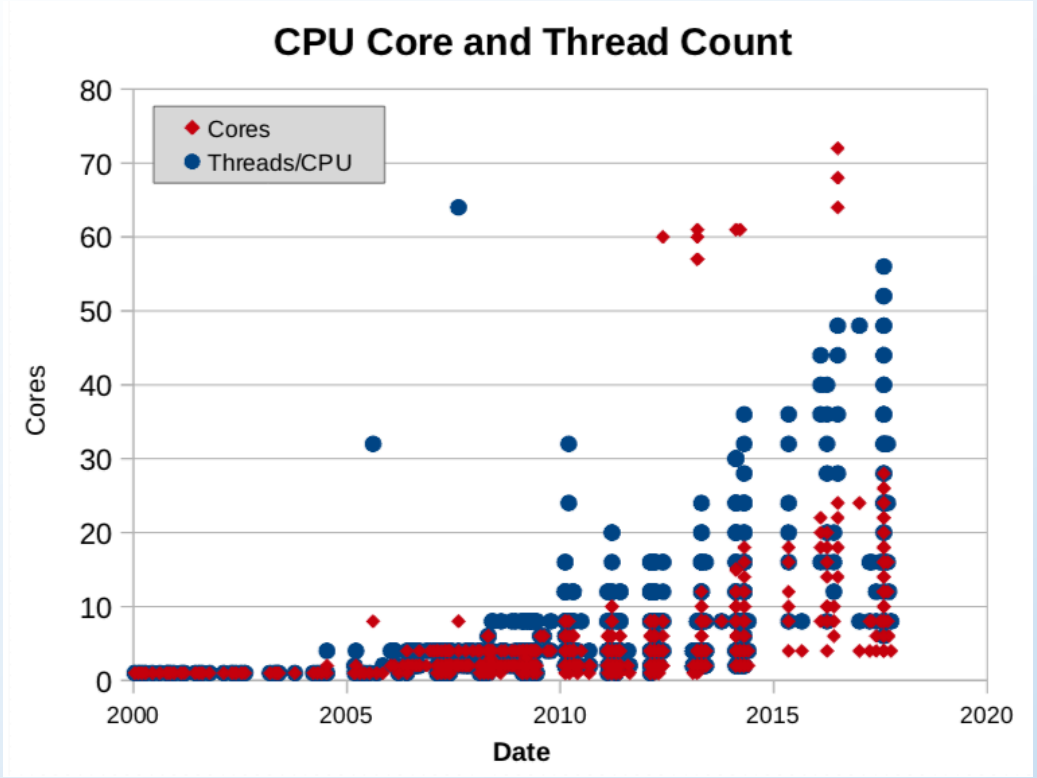

Trend towards more cores and slower memory access. This is directly at odds with ”traditional” HEP applications: Single threaded and memory-heavy

We are not alone: Other scientific areas and the broader data science community have quickly increasing needs

Astronomy/Astrophysics: Eg, SKA

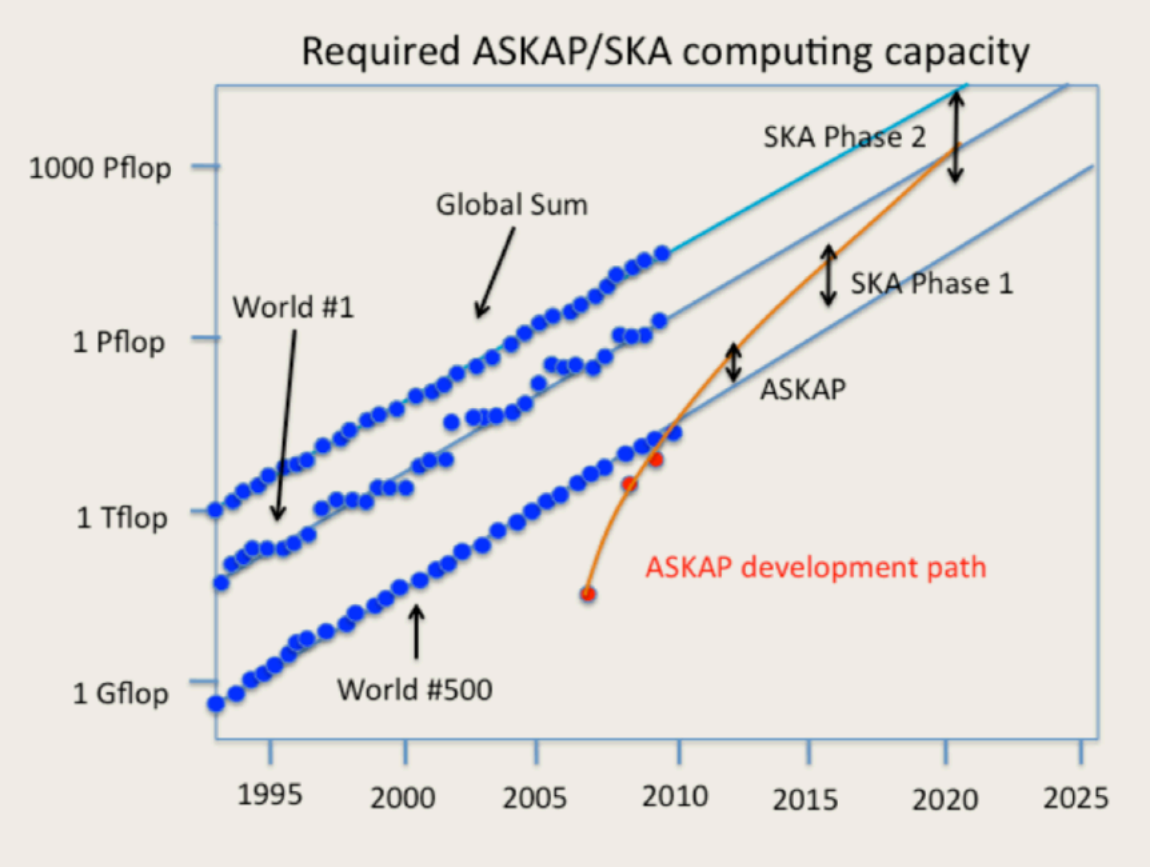

https://arxiv.org/pdf/1101.1355.pdf

Genomics: Eg, DNA sequencing

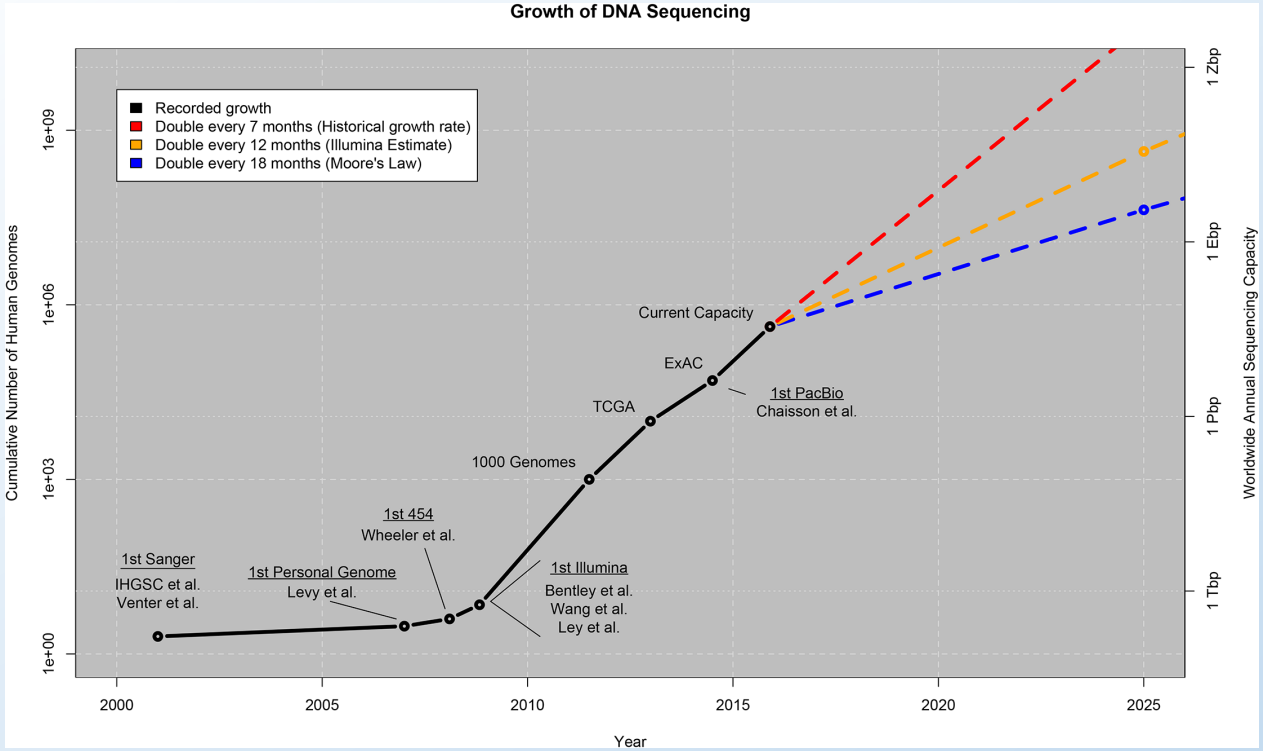

http://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.1002195

Fortunately we are not alone: Other scientific areas and the broader data science community have quickly increasing needs

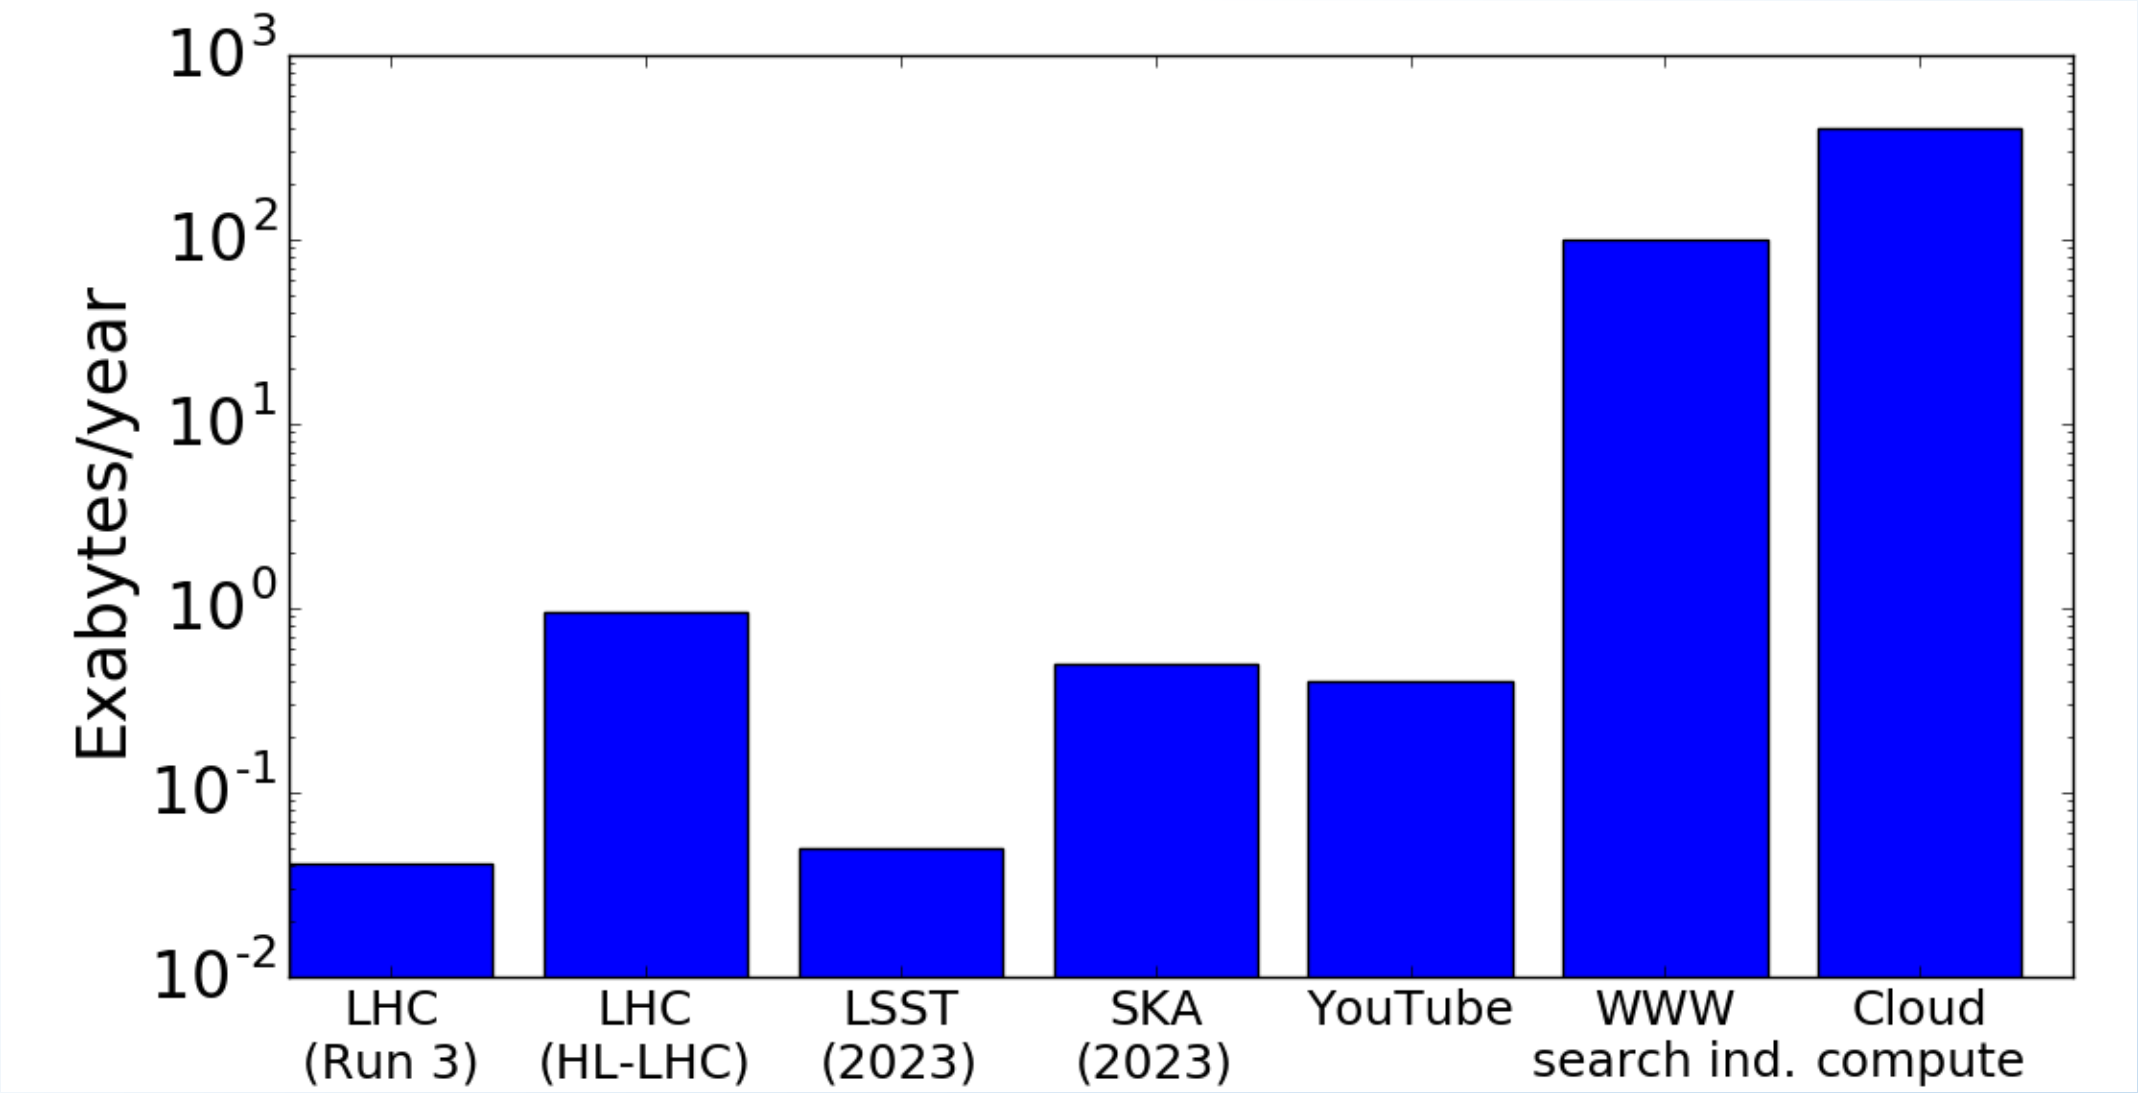

Needed evolution of HEP applications is underway

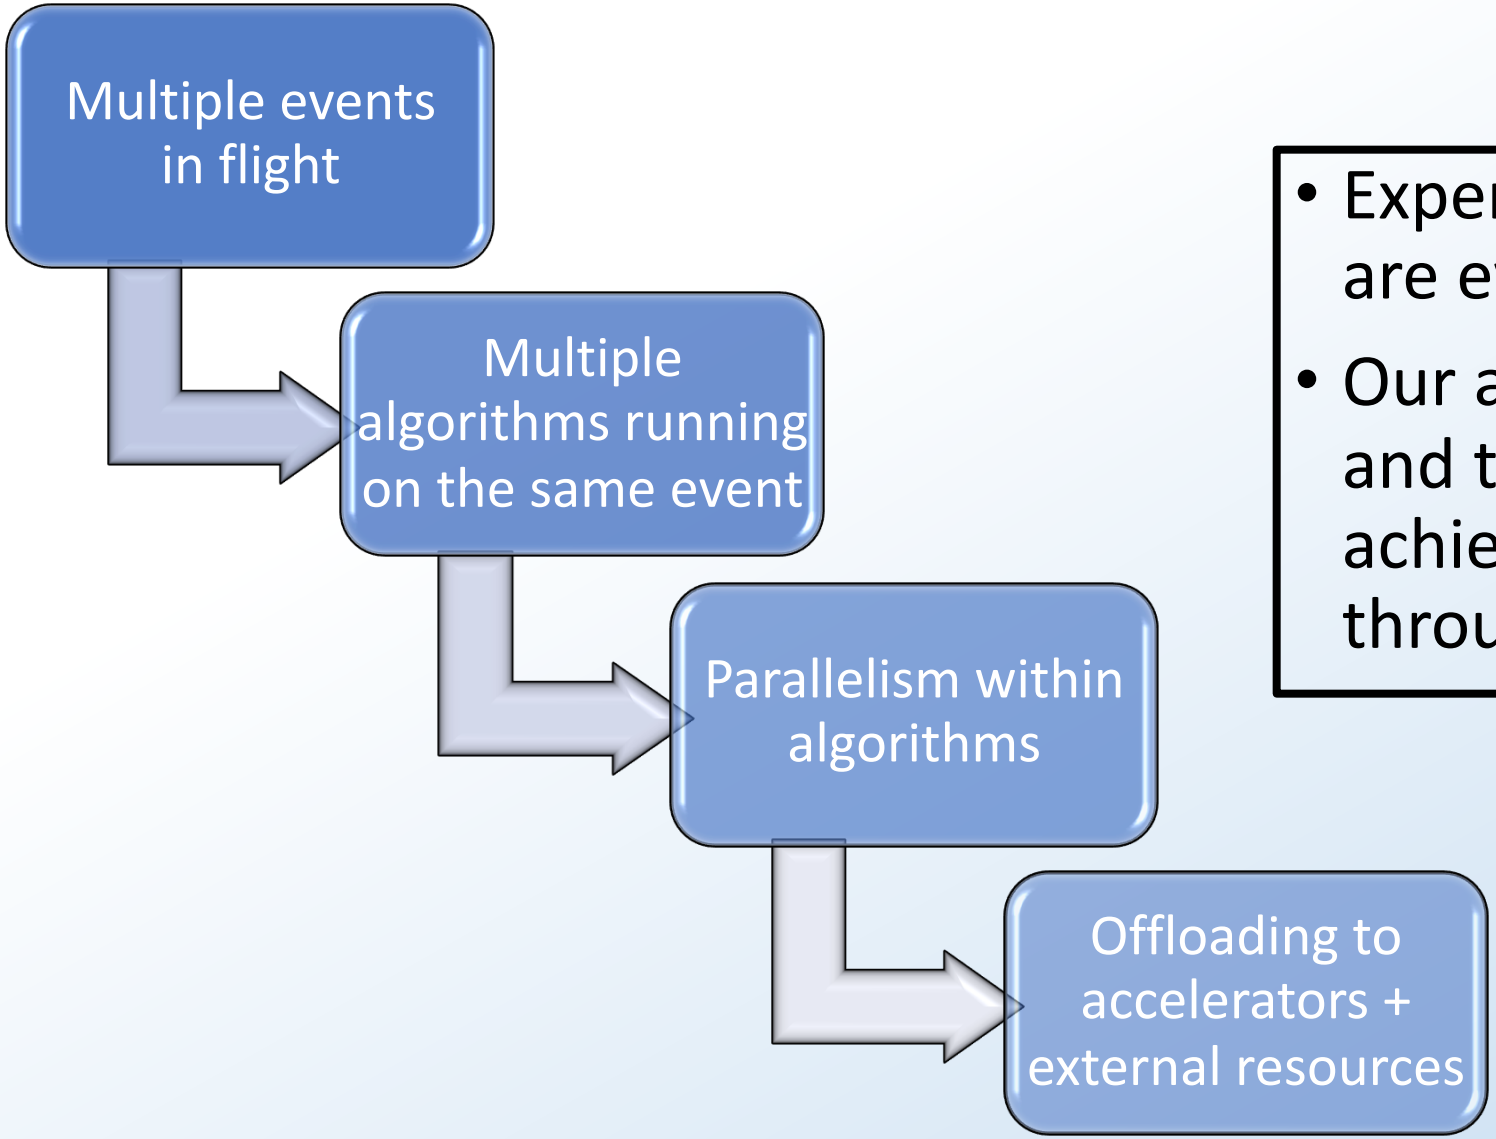

- Experiment frameworks are evolving
- Our algorithms must adapt and take advantage to achieve needed event throughput

Typical event reconstruction application to be parallelized is constructed from many individual algorithms

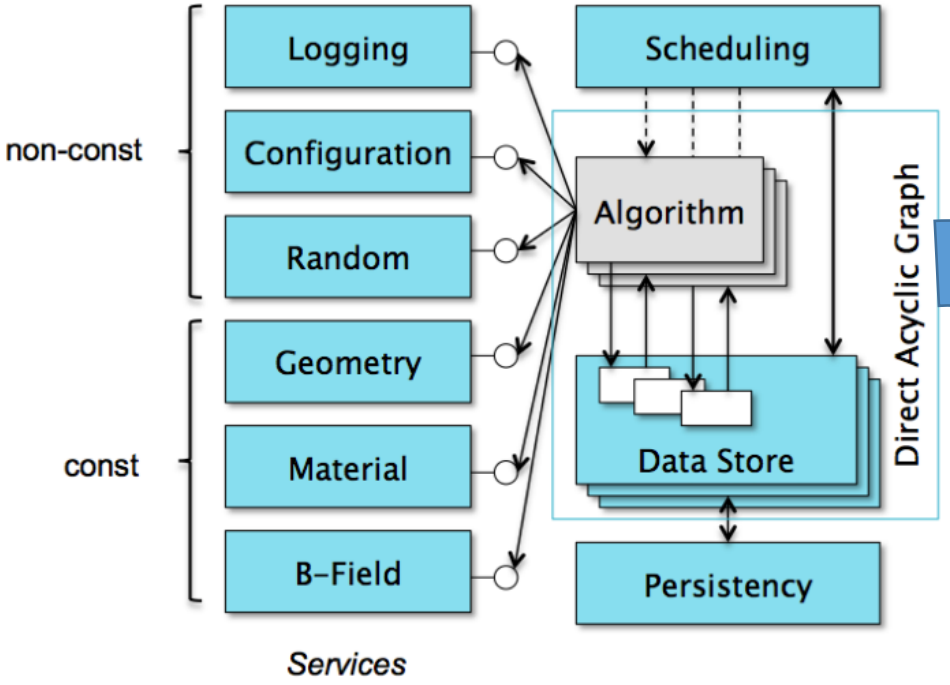

- First critical piece to prepare our algorithms for ~~future~~ present… computing hardware

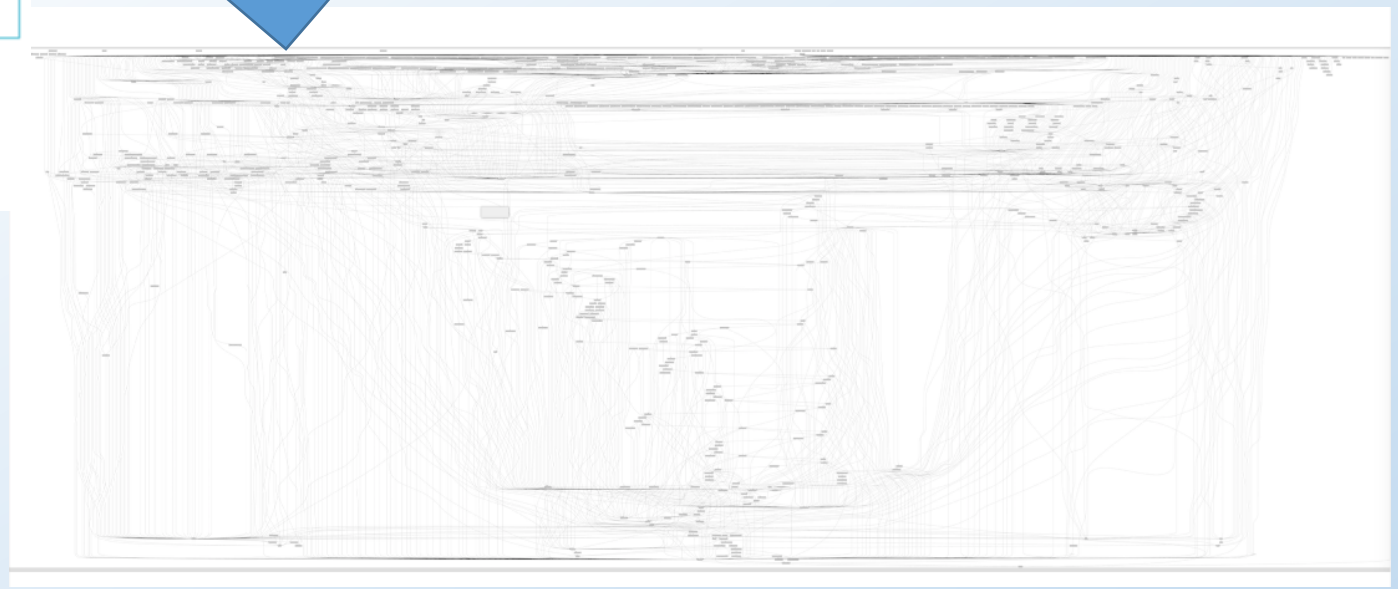

Production frameworks are already able to schedule these workflows efficiently

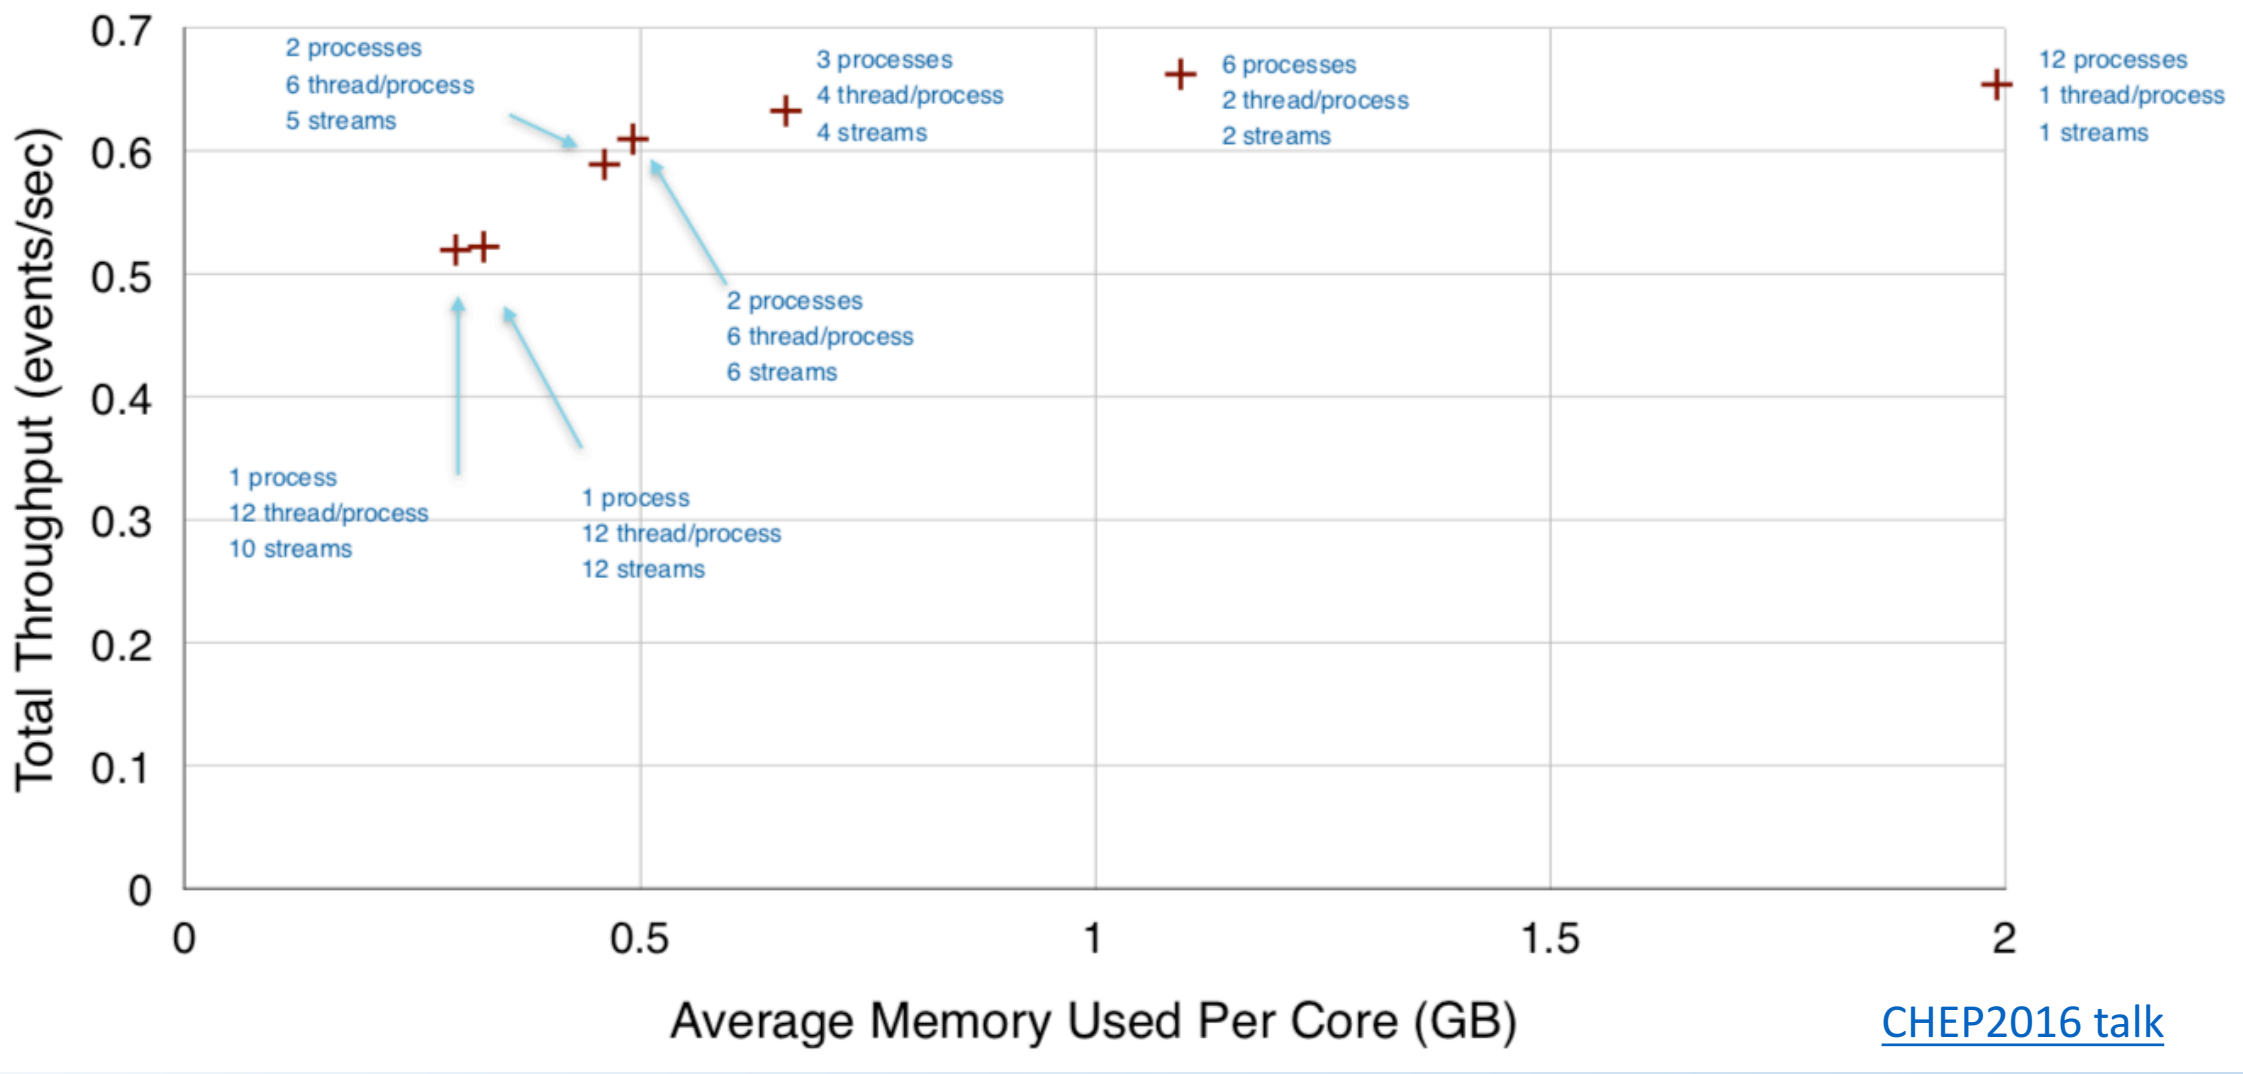

CHEP2016 talk

Event reconstruction and software trigger algorithm R&D

- Enhanced vectorization programming techniques
- Algorithms and data structures to efficiently exploit many-core architectures
- Algorithms and data structures for non-x86 computing architectures (eg, GPUs,FPGAs)
- Enhanced QA/QC for reconstruction techniques
- Real-time analysis
- High precision physics-object reconstruction, identification and measurement techniques
- Fast software trigger and reconstruction algorithms for high-density environments

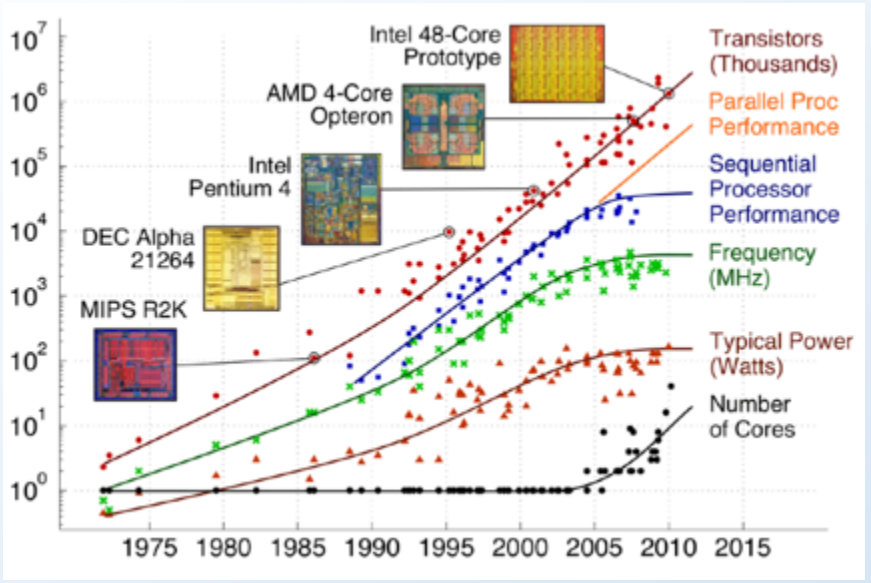

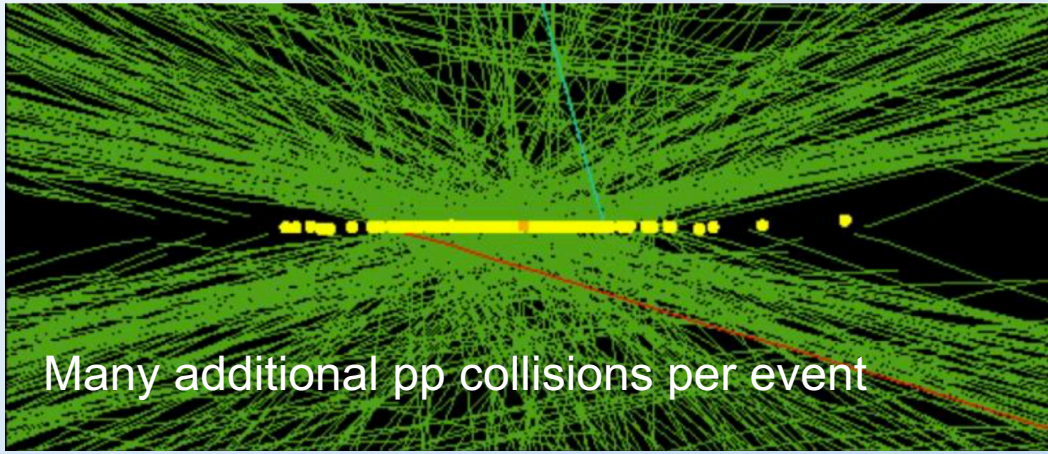

https://arxiv.org/abs/1802.08640

To really gain performance, go beyond thread parallelism on CPU

- Tracking algorithms are a natural target: They are a large fraction of the HEP reconstruction CPU budget
- Challenges:
  - Small matrices mean that usual approaches to vectorizing matrix algebra do not work well
  - Algorithms rely on branch points
  - Using wide vector units (e.g., AVX) often slows the frequency of the CPU. A significant portion of the computational kernel must be vectorized to benefit

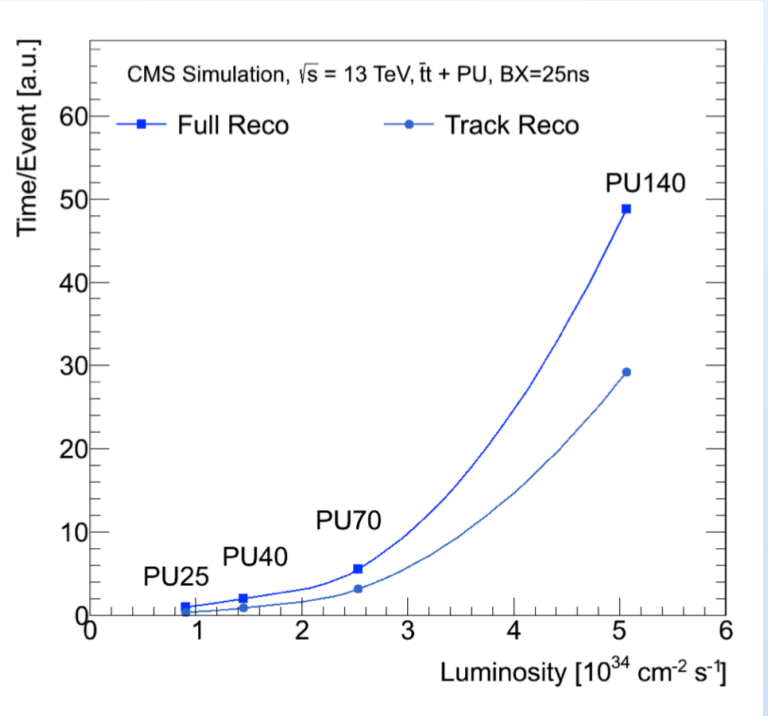

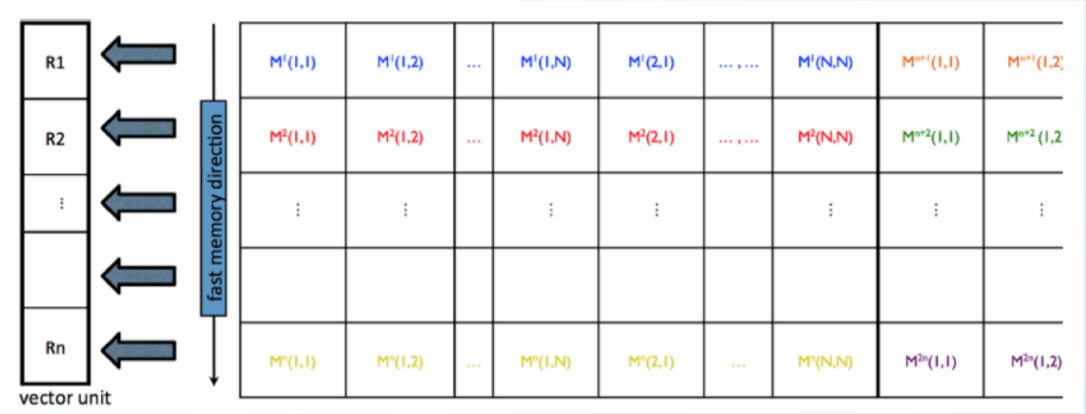

Example project: MatrixPlex project recasts N small matrices to fill a vector unit (or accelerator) while operating on them in sync

- Achieves 2-3x speedup in track building (AVX, AVX-512)

[Matriplex, parallelized kalman fitter project]

Memory speeds have not kept up with CPUs

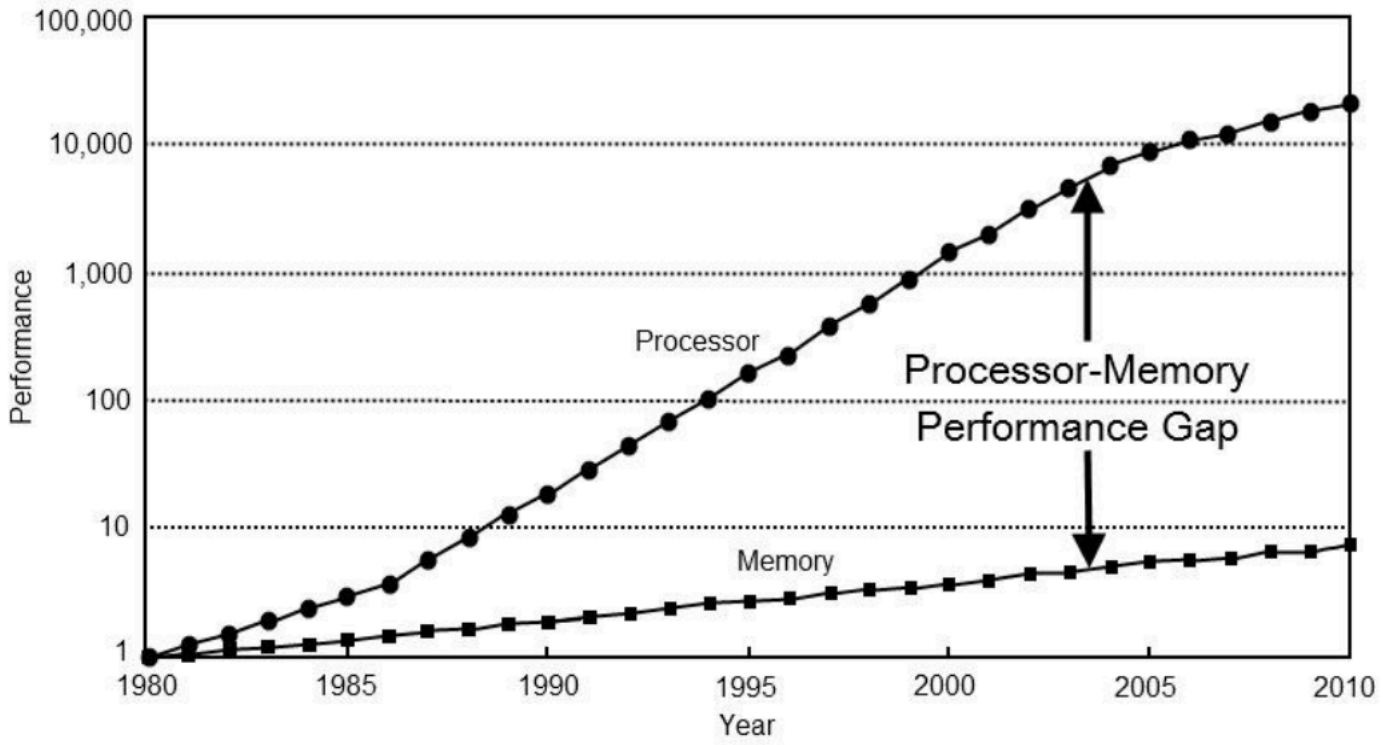

Processor-Memory speed evolution

Consequences

- Rise of larger on-chip cache memory
- Large latency for cache misses in algorithms
- Potentially large gains for using “data aware” programming methods
  - Eg, “Structs of Arrays” vs “Arrays of Structs”

Heterogeneous computing – incorporating accelerators into reconstruction workflows

- Opportunity
  - Industry trending to specialized hardware
  - FPGAs are one example that have become much easier to program in recent years
- Challenge: Recasting algorithms to make efficient use of accelerators

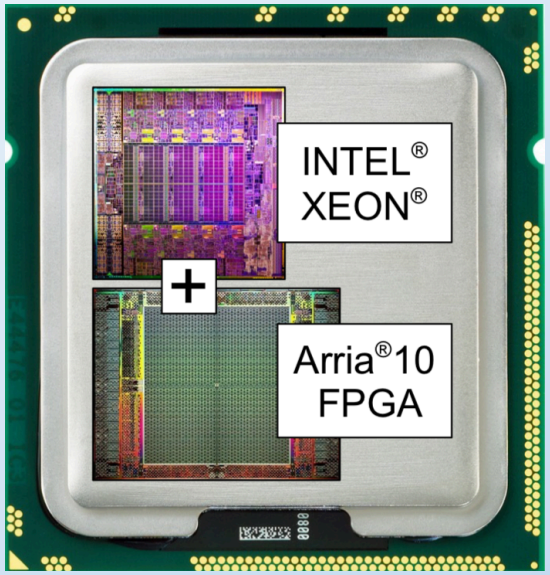

- Managing data movement essential for effective use
  - Massive parallelism within the event
  - Independence from thread ordering in algorithms
  - Avoid unneeded data transfers and transformations
  - Simple data formats optimized for parallel memory access

Example: Cellular automaton for track seeding

Reconstruction algorithms full of potential Machine Learning applications

- HEP has used neural networks and other machine learning techniques for a long time.
- New techniques have brought new capabilities for dealing with noisy data
- Will training deep learning networks become the primary activity of HEP researchers?

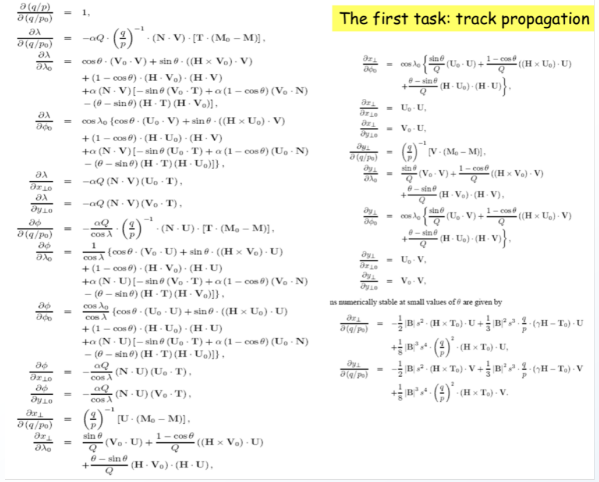

?

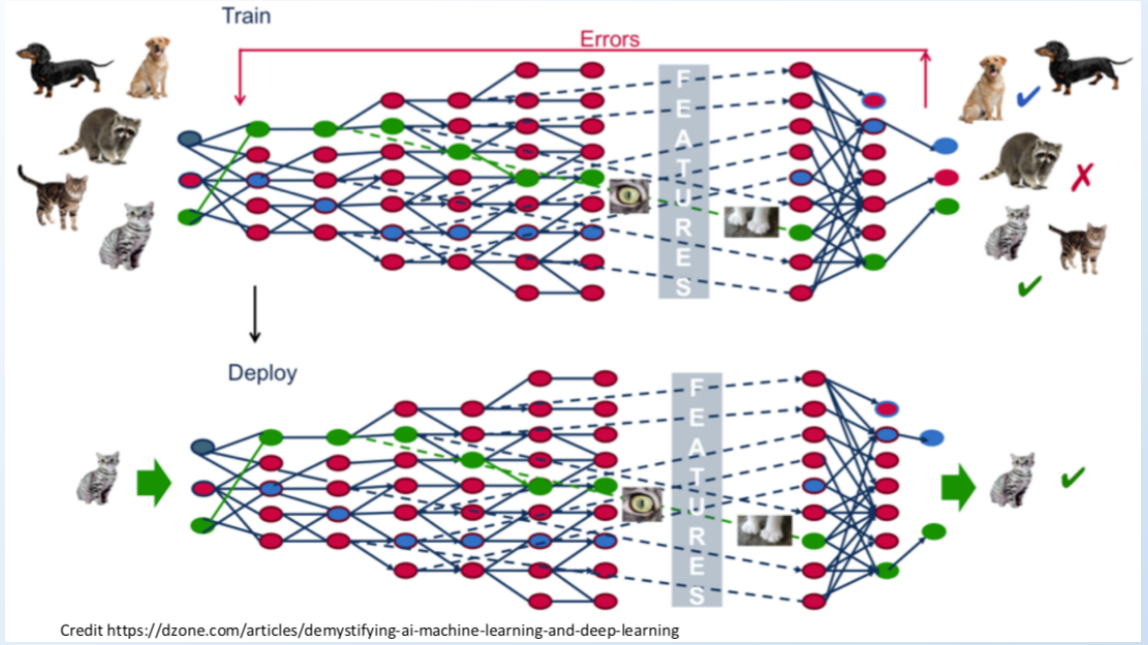

Credit https://dzone.com/articles/demystifying-ai-machine-learning-and-deep-learning

- New approaches bring opportunity for physics-driven choices for which approach is best for the HL-LHC algorithm stack

2016-2017: Development of a “Community Whitepaper” on needed R&D towards software for HL-LHC

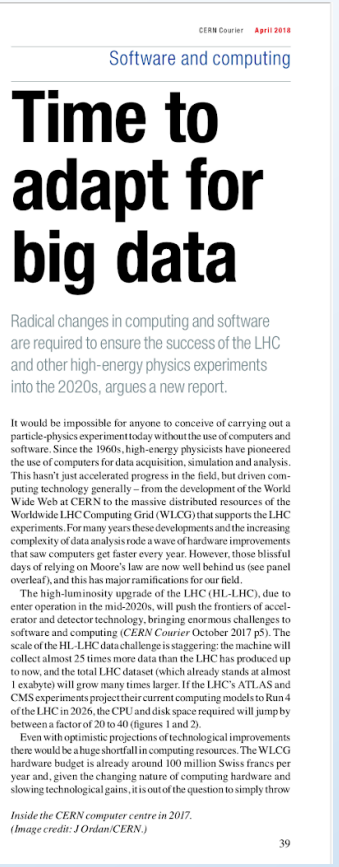

CERN Courier April 2018

Software and computing

Time to adapt for big data

Radical changes in computing and software are required to ensure the success of the LHC and other high-energy physics experiments into the 2020s, argues a new report.

It would be impossible for anyone to conceive of carrying out a particle-physics experiment today without the use of computers and software. Since the 1960s, high-energy physicists have pioneered the use of computers for data acquisition, simulation and analysis. This hasn't just accelerated progress in the field, but driven computing technology generally – from the development of the World Wide Web at CERN to the massive distributed resources of the Worldwide LHC Computing Grid (WLCG) that supports the LHC experiments. For many years these developments and the increasing complexity of data analysis rode a wave of hardware improvements that saw computers get faster every year. However, those blissful days of relying on Moore's law are now well behind us (see panel overleaf), and this has major ramifications for our field.

The high-luminosity upgrade of the LHC (HL-LHC), due to enter operation in the mid-2020s, will push the frontiers of accelerator and detector technology, bringing enormous challenges to software and computing (*CERN Courier* October 2017 p5). The scale of the HL-LHC data challenge is staggering: the machine will collect almost 25 times more data than the LHC has produced up to now, and the total LHC dataset (which already stands at almost 1 exabyte) will grow many times larger. If the LHC's ATLAS and CMS experiments project their current computing models to Run 4 of the LHC in 2026, the CPU and disk space required will jump by between a factor of 20 to 40 (figures 1 and 2).

Even with optimistic projections of technological improvements there would be a huge shortfall in computing resources. The WLCG hardware budget is already around 100 million Swiss francs per year and, given the changing nature of computing hardware and slowing technological gains, it is out of the question to simply throw

Inside the CERN computer centre in 2017. (Image credit: J Ordan/CERN.)

39

[ArXiV link to CWP]

[ArXiV link to trigger/reco CWP]

- Charge from the WLCG in July 2016
  - 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

Conclusion

- Expect HL-LHC reconstruction applications to look very different from traditional HEP or even LHC Run 2 applications
  - No choice but to adapt given the scale of the resource need to capture the exciting HL-LHC physics program
  - Educating our community on opportunities for evolution is critical
- The best approach is still unknown, so its an exciting time to work on new ideas and new concepts

Dilbert 2012-04-17