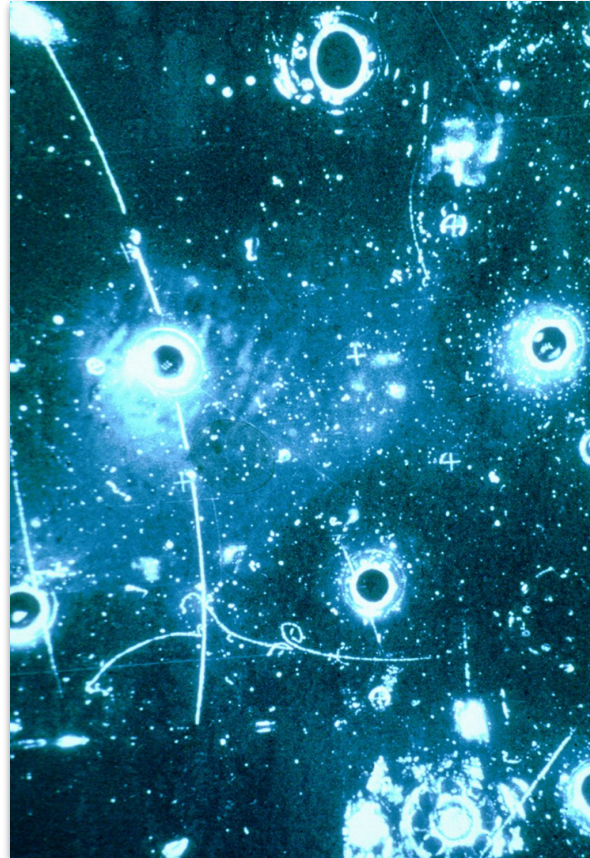# Preparing for the HL-LHC computational challenge
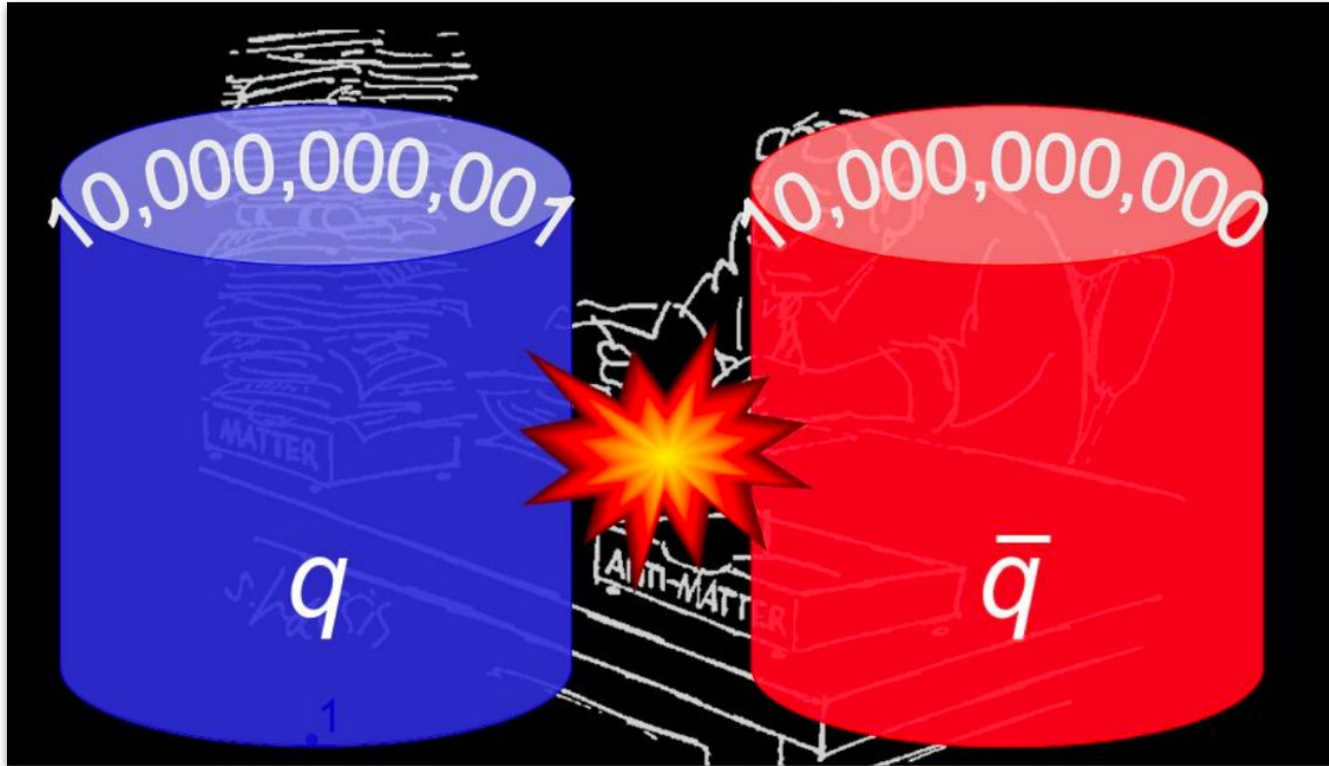
D. Piparo, CERN - tCSC 2021

Outline

In this talk we will introduce some basic concepts related to HEP data processing and analysis workflows, seeing them in action in the context of LHC experiments. We'll also talk about the evolution of the LHC accelerator and experiments. We'll characterise at a high level what are the consequences of those upgrades for the HEP data processing software, in particular in the context of an evolving hardware and computing infrastructure.
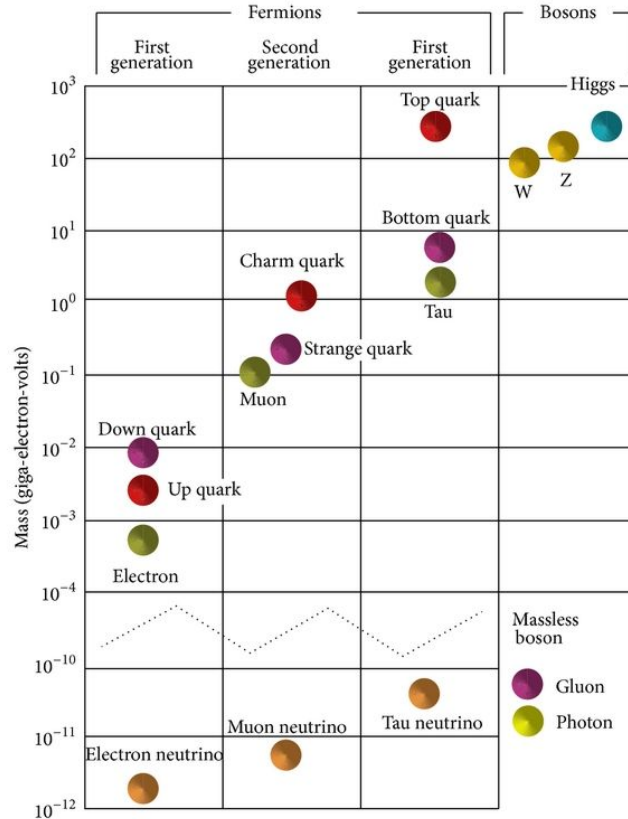
Understanding the Universe

© CERN
http://cds.cern.ch/record/39468
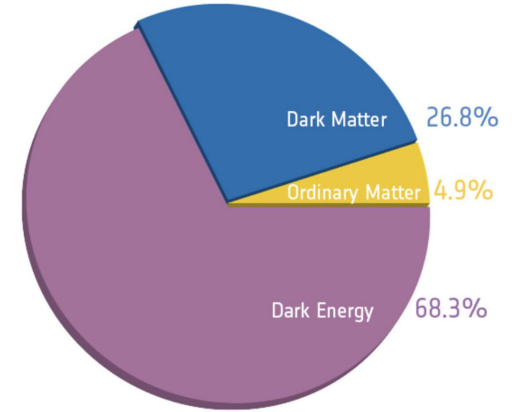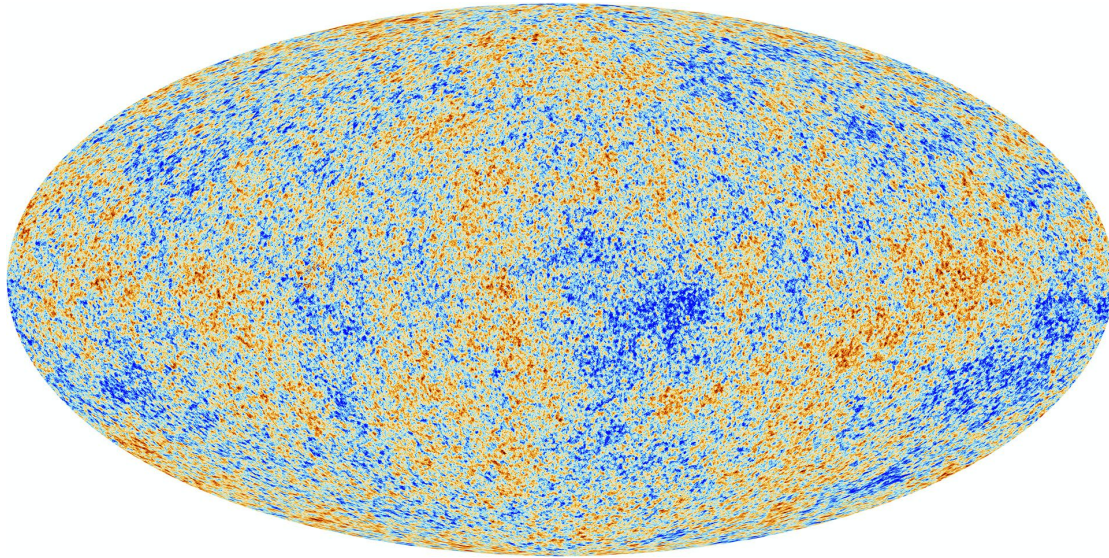
Some Unanswered Questions

A. Hoecker, https://indico.cern.ch/event/34851/

Some Unanswered Questions

Some Unanswered Questions

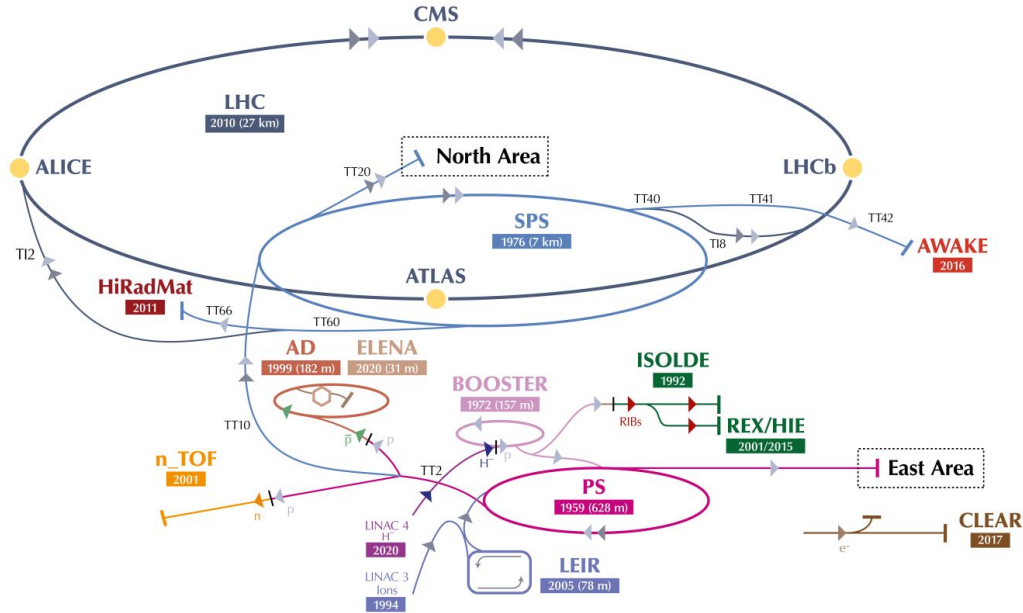**© ESA and The Plank Collaboration**
http://www.esa.int/Science_Exploration/Space_Science/Planck/Planck_reveals_an_almost_perfect_Universe

The CERN accelerator complex
Complexe des accélérateurs du CERN
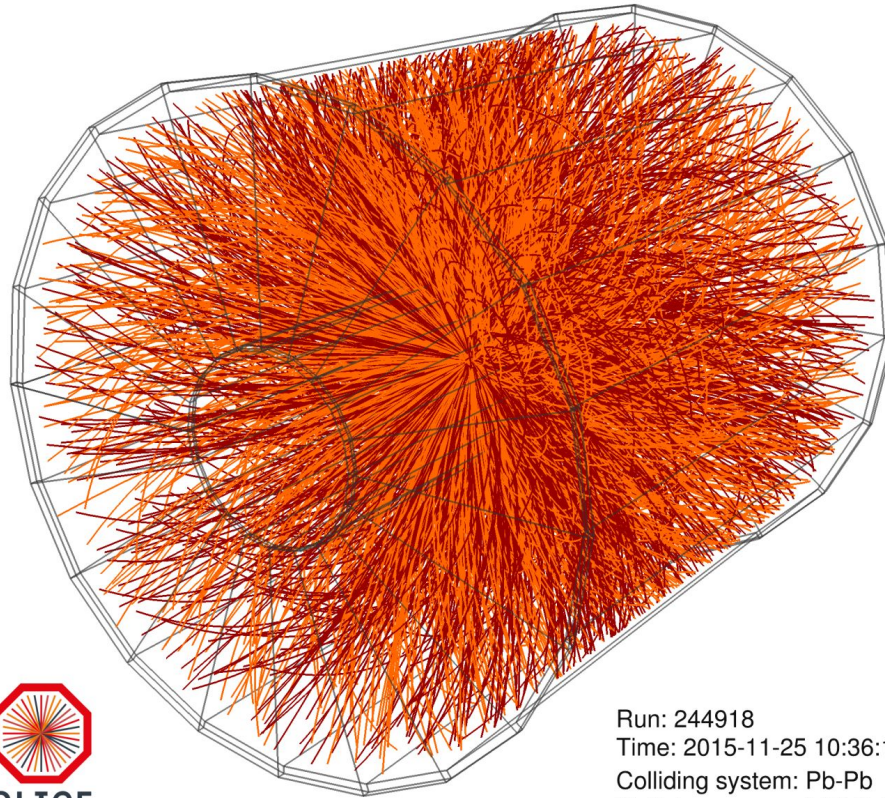
Collide

Run: 244918
Time: 2015-11-25 10:36:18
Colliding system: Pb-Pb
Collision energy: 5.02 TeV

ALICE

CMS Experiment at the LHC, CERN
Data recorded: 2015-Oct-30 19:23:54.631552 GMT
Run / Event / LS: 260424 / 211873064 / 115

15 m
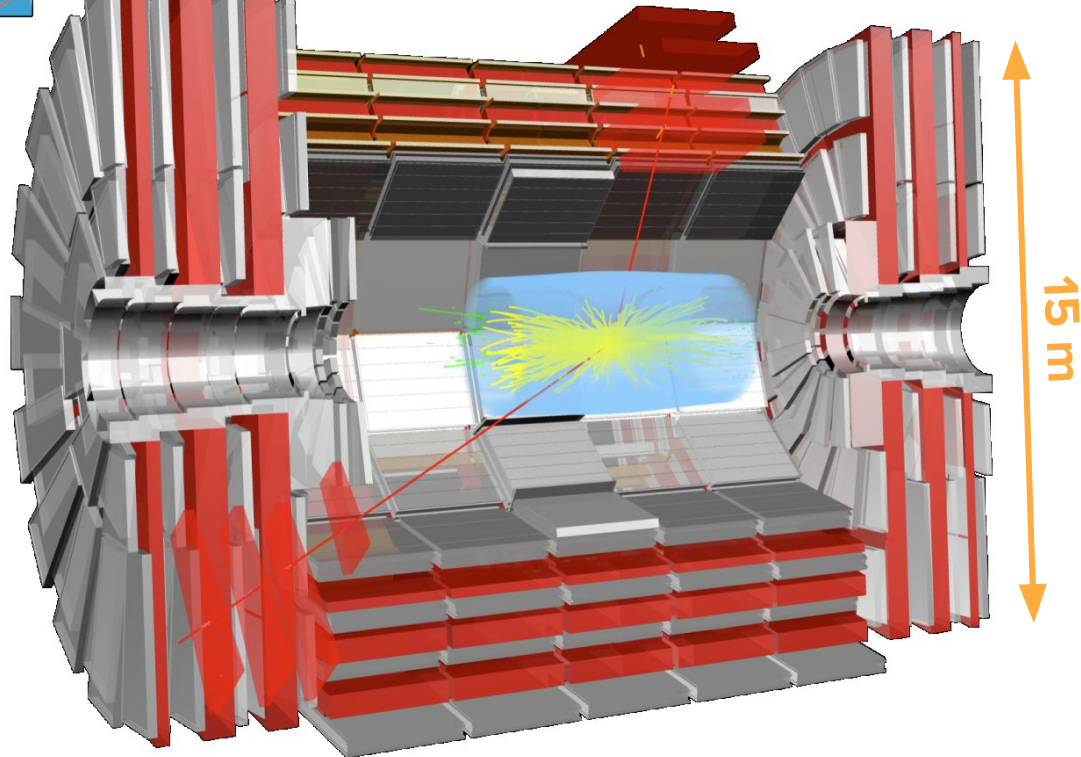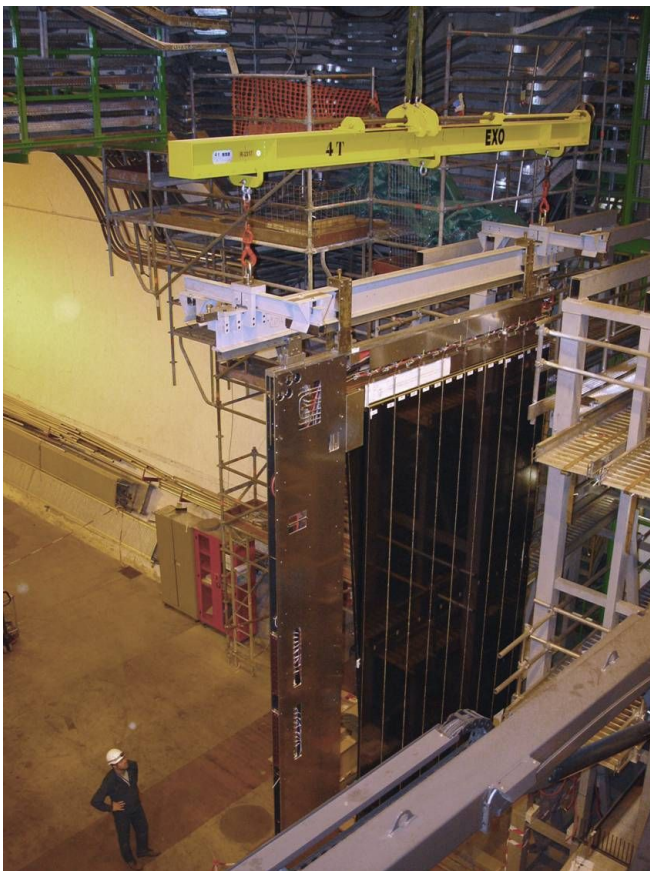
© CERN **LHCb-PHO-OUTTRACK-2006-001-1**

© CERN **LHCb-PHO-SILTRACK-2008-002-1**

An Onion Structure

© CERN CERN-EX-1301009

Process

Analyse

https://lhcb-public.web.cern.ch/Welcome.html#BsOsc21

Discover!
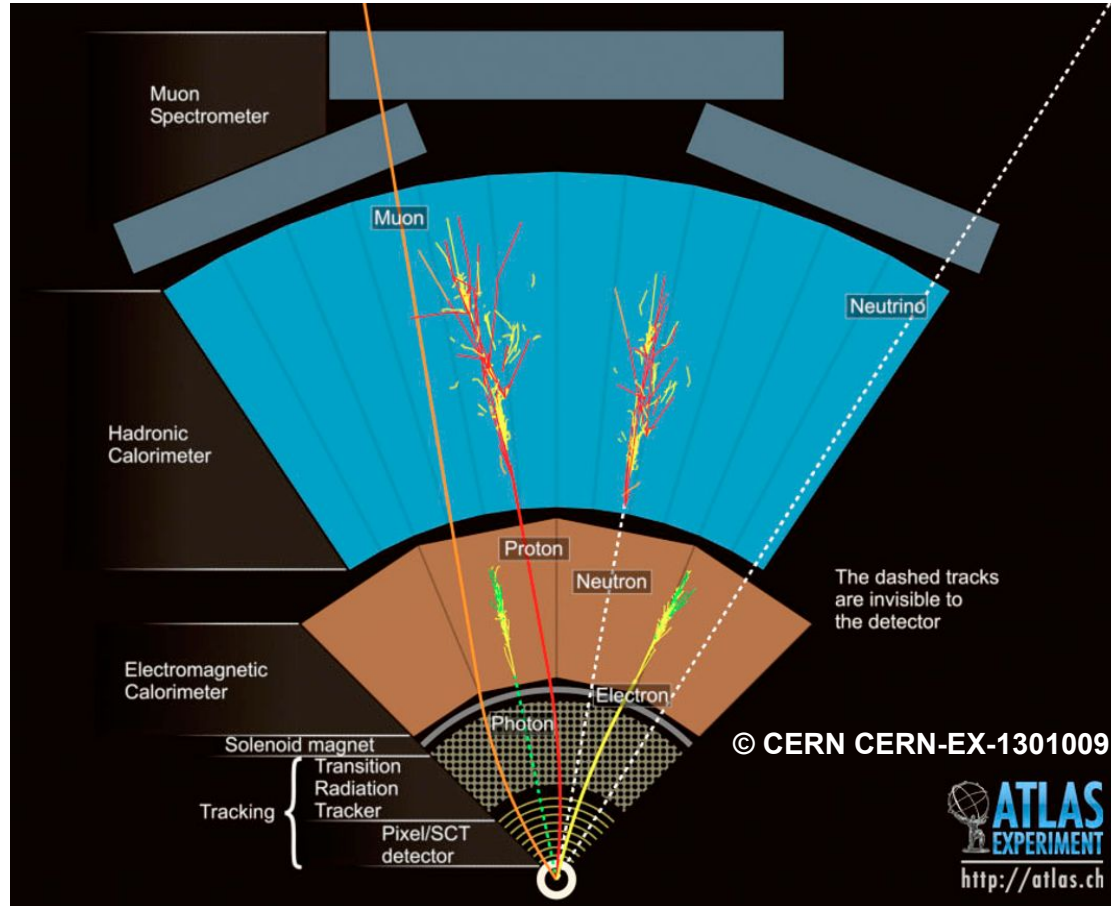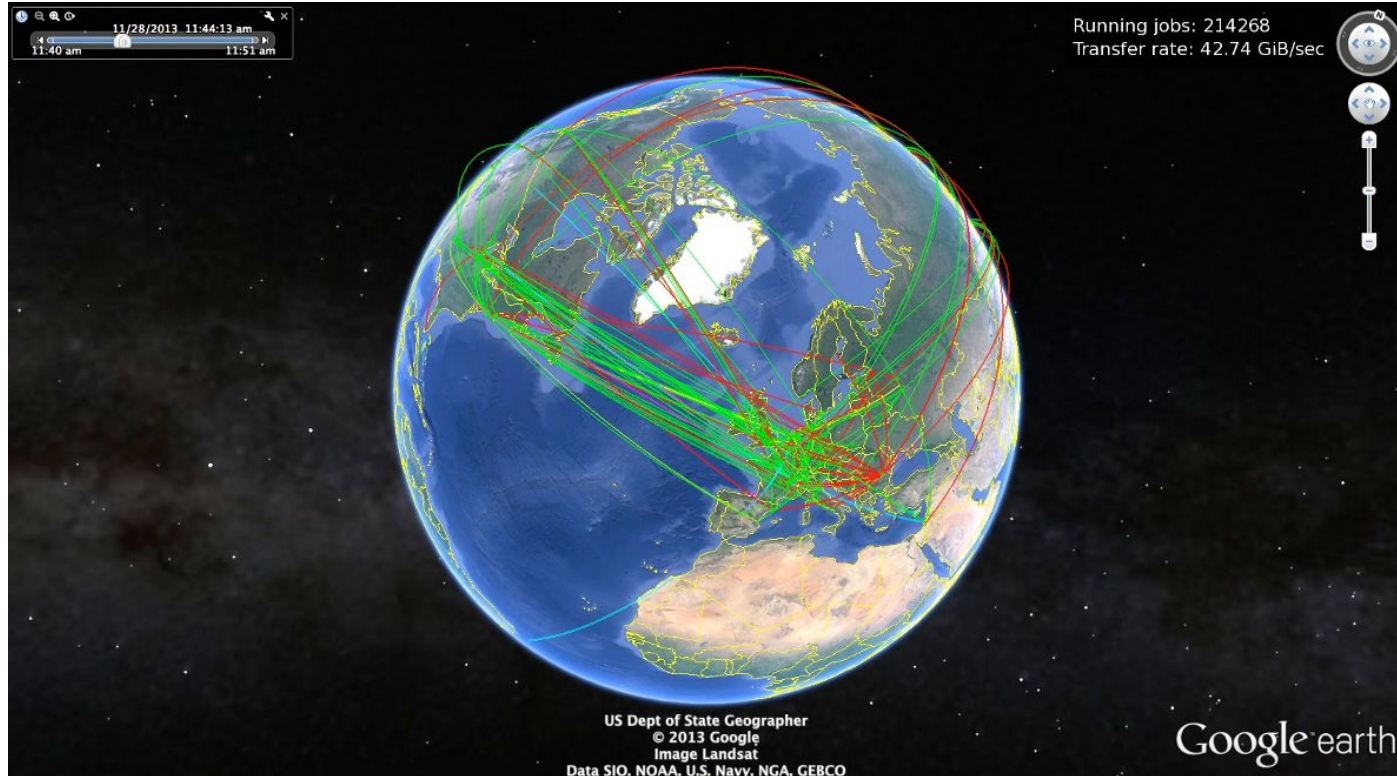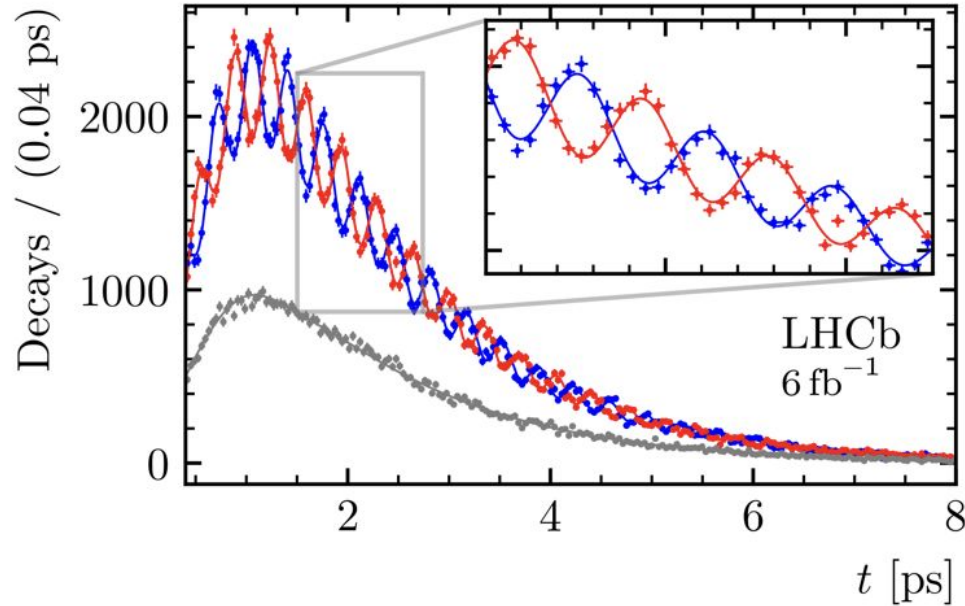
At CERN on 4 July 2012, the ATLAS and CMS collaborations present evidence in the LHC data for a particle consistent with a Higgs boson, the particle linked to the mechanism proposed in the 1960s to give mass to the W, Z and other particles. (Image: Maximilien Brice/Laurent Egli/CERN)

https://home.cern/science/physics/higgs-boson

HEP Analysis Model
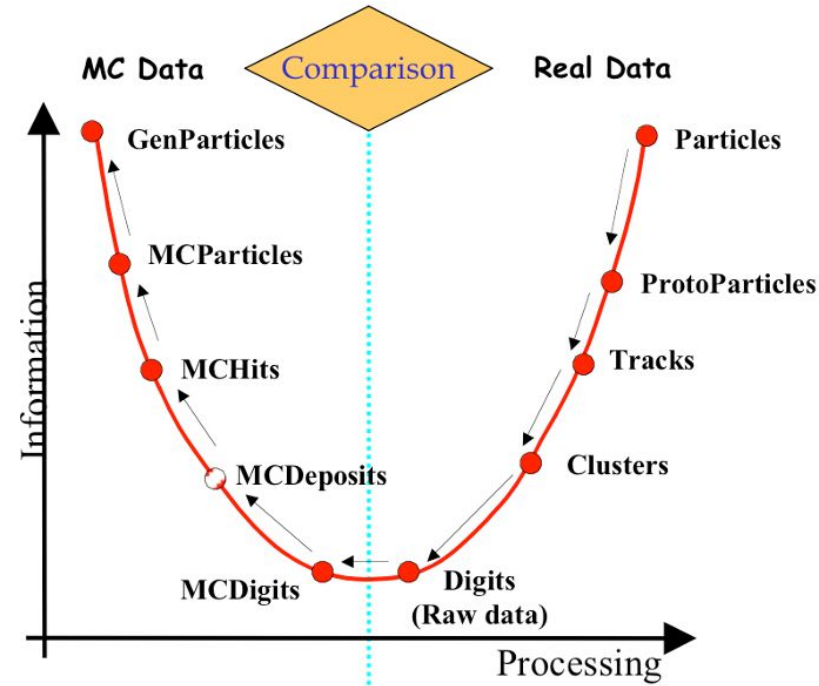
J. Harvey

Monte Carlo simulation follows the evolution of physics processes from collisions to digital signals

Reconstruction "goes back in time" from digital signals to the original particles produced in the collision

Simulation accounts for the usage of substantial computing resources.

Data acquisition and processing in a nutshell

- **HEP main data: statistically independent* Events (particle collisions)**
- **Simulation, Reconstruction and Analysis: process "one Event at the time"**
  - **Event-level parallelism** (success of the Grid!)
  - Landscape is changed: advent of parallel data processing frameworks
- **Applications composed of several algorithms to:**
  - **Select and transform measured/simulated "raw" event data into "particles"**
  - **Create simulated "raw" event data** (event generation+simulation of passage of particles through matter+simulation of detector response to such energy depositions)
- **Final result: statistical data (histograms, distributions, etc.)**
  - Typically**: comparison between simulation and data**
- **All of these algorithms:**
  - **Are mainly developed by "Physicists"**
  - May require additional "detector conditions" data (e.g. calibrations, geometry, etc)

* Continuous readout would deserve more details.

Raw data "definition": readout of the ADC of the subdetectors' frontends

Processing time

Data acquisition and processing in a nutshell

Parallelism at all levels
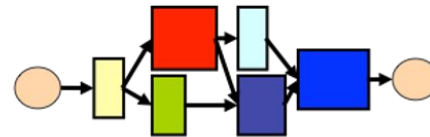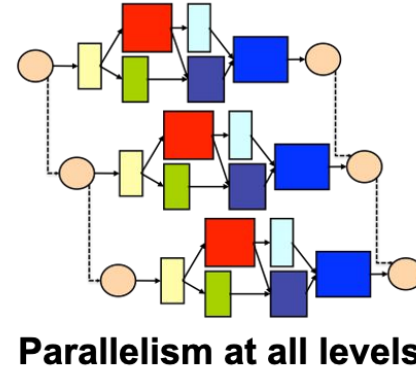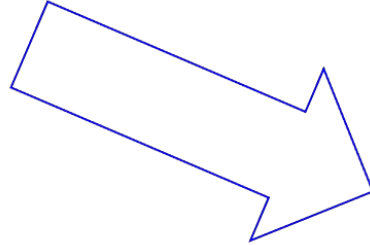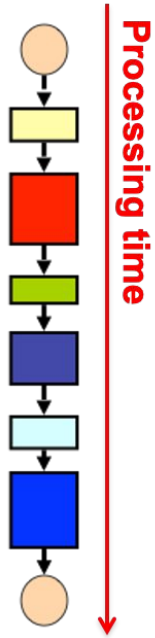
Parallelism within an event

How Much Data?

Experiment	2021 Disk Pledge [PB]*	2021 Tape Pledge [PB]*	2021 Disk+Storage Pledge [PB]*
ALICE	141	146	287
ATLAS	272	336	608
CMS	202	339	541
LHCb	59	109	168
TOTAL	674	930	1412

Numbers approximated to the nearest unit

https://wlcg-cric.cern.ch/core/vopledgereq/listcomp/?year=2017

The Computational Challenges Ahead

How can we achieve all that in the Future?
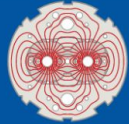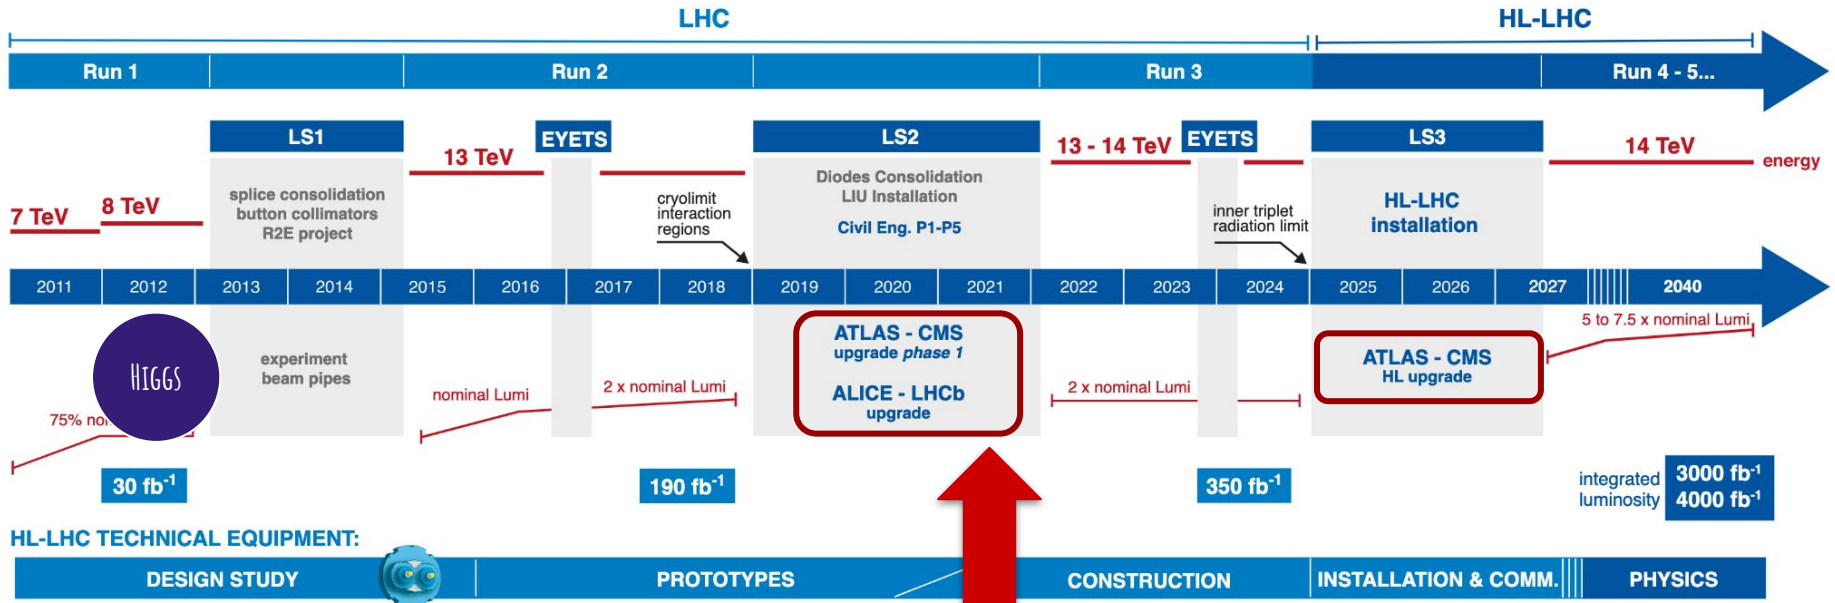
Upgraded detectors + improved accelerator
Alice and LHCb: Upgraded already!

https://project-hl-lhc-industry.web.cern.ch/content/project-schedule

How can we achieve all that in the future?

- **Upgraded detectors and an upgraded LHC**
- **Much more granular measurements, larger RAW events, larger number of objects to be considered during reconstruction**
- All this in the context of rapidly evolving computing platforms!

Ultimate HL-LHC performance

https://espace.cern.ch/HiLumi/WP2/Wiki/HL-LHC%20Parameters.aspx

HL-LHC upgrades have to be paired to an equally ambitious upgrade of data processing software, computing tools and services.
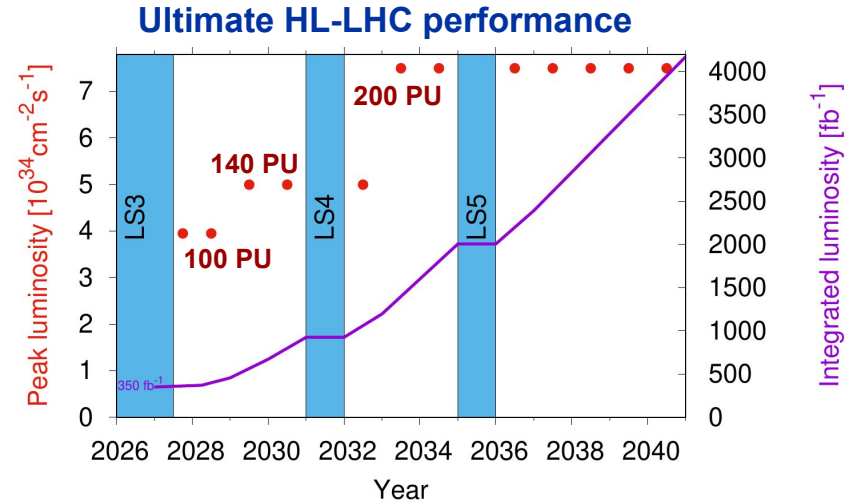
By delivering more integrated luminosity, more parasitic collisions per bunch crossing are generated, more complexity to be handled by our algorithms!

Vertices

78 Primary Vertices: expect up to 200 in the HL-LHC era!!
And no, complexity does not scale linearly with # parasitic collisions :-(
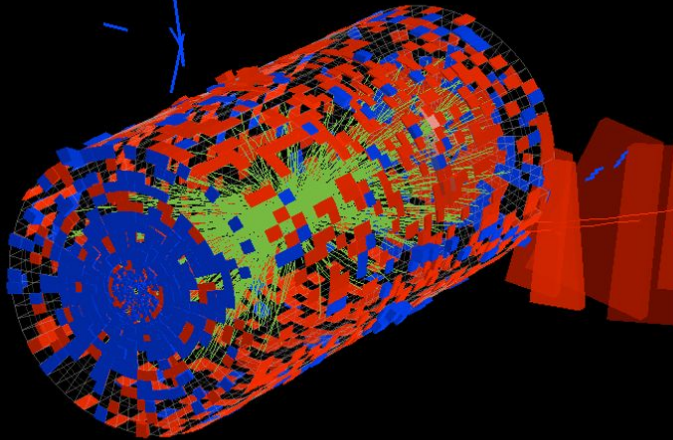
Image credit: Andre Holzner

https://cms.cern/news/reconstructing-multitude-particle-tracks-within-cms
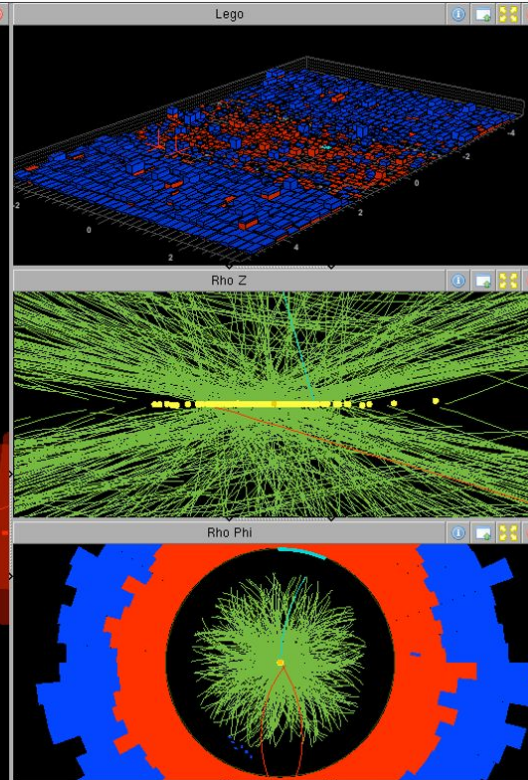
A Large Ion Collider Experiment
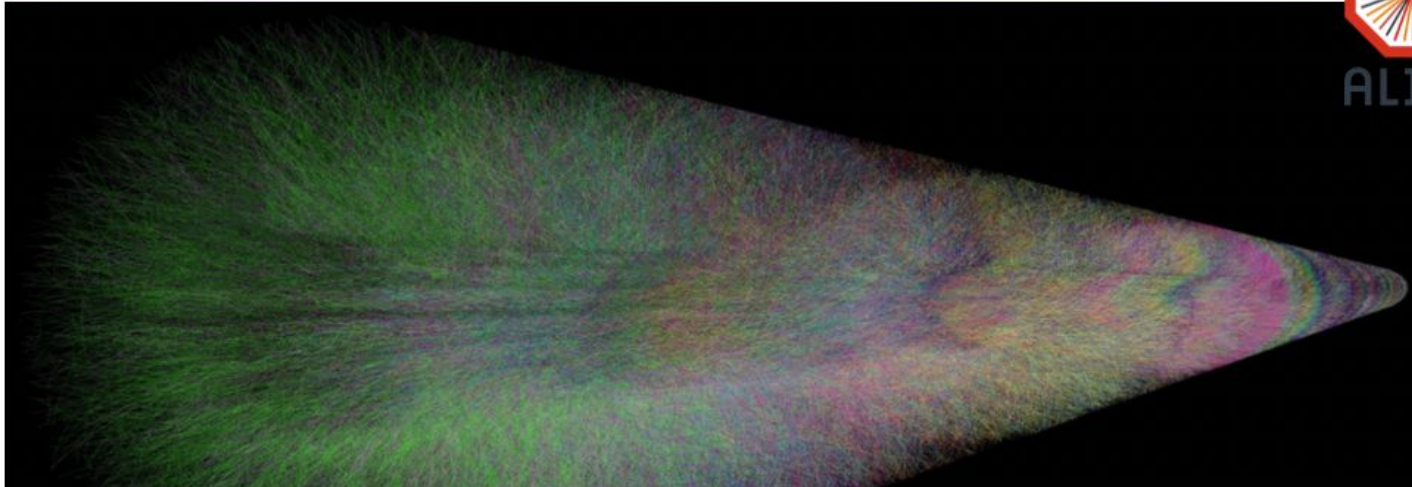
Visualization of **2 ms of 50 kHZ Pb-Pb** data as expected in the ALICE TPC in LHC Run 3
- Basic processing unit: **Time Frame** (~10 ms of data ~500 collisions @ 50 kHZ Pb-Pb)
- Whole TF reconstructed in one shot
- In absence of triggers (reference for drift-time estimate) z position of clusters is not defined

S. Piano, https://agenda.infn.it/event/26182

Offline projected computing needs in the HL-LHC era

CERN-LHCC-2020-015

CMS-NOTE-2021-001
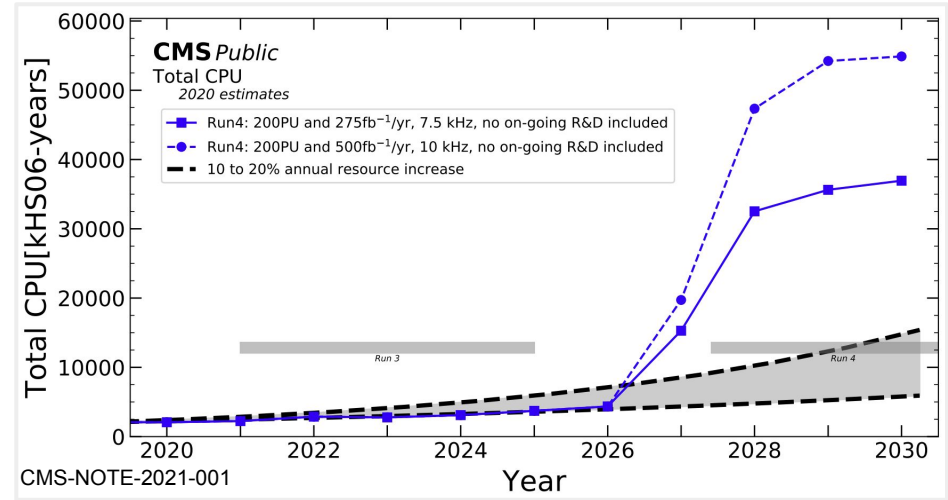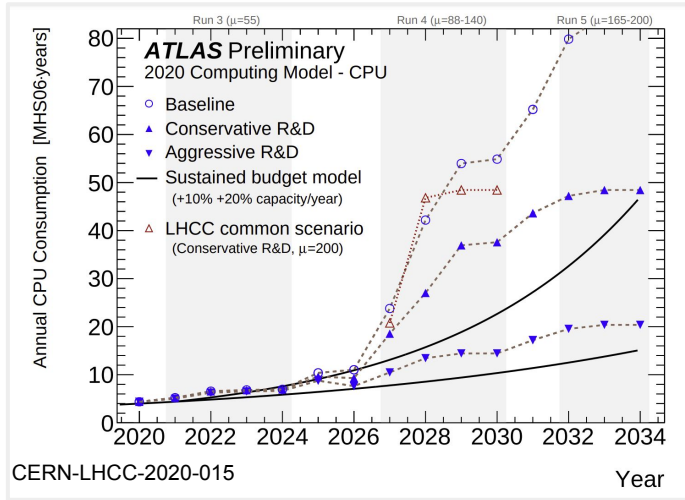
Online rates for upgraded detectors
- ALICE: 1 kHz   → 50 kHZ
- ATLAS: 1 kHz   → 10 kHz
- CMS: 1 kHz   → 7.5 kHz
- LHCb: 12.5 kHz → 30 MHz

What a big gap between projected needs and availability!

It can be reduced with talented and curious physicists-developers :) **A single person can make the difference!**

Can technology alone come to rescue?

Price/performance evolution of installed CPU servers (CERN)

B. Panzer

Samsung DIMM 16GB, DDR4-2133, CL15-15-15, reg ECC
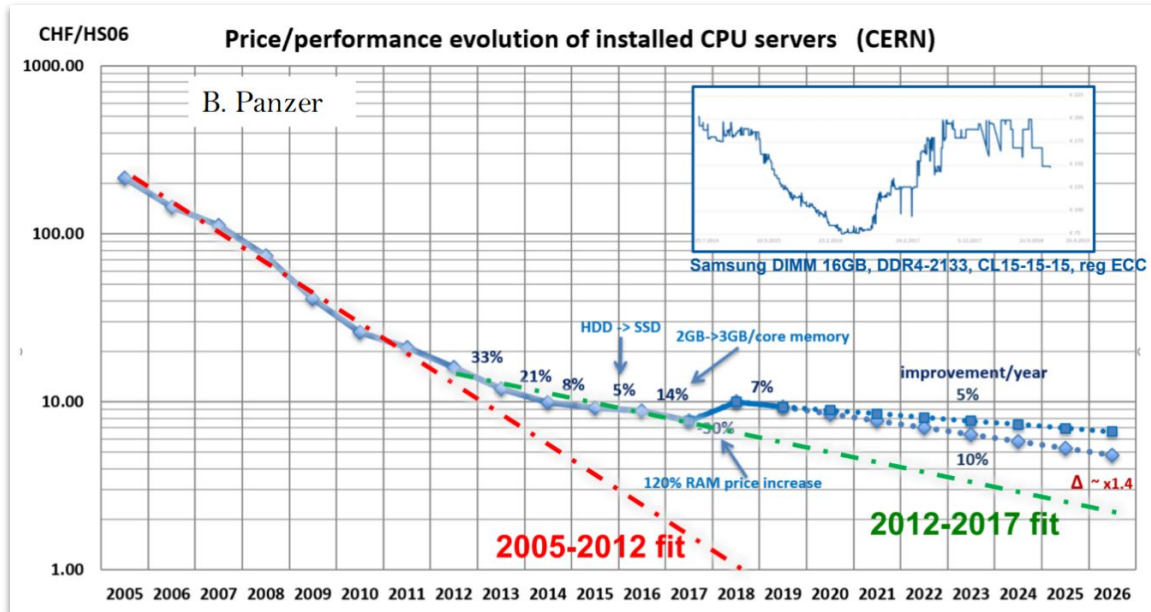
- Predictions affected by large uncertainties
- Cannot count on technology alone
- Fierce competition among CPU manufacturers
- In the HL-LHC era, **computing power might not come only from CPUs, but also from accelerators**

Parallelism 1/2

40 Years of Microprocessor Trend Data

Legend:
- Transistors (thousands)
- Single-Thread Performance (SpecINT x 10^3)
- Frequency (MHz)
- Typical Power (Watts) — "Power wall"
- Number of Logical Cores
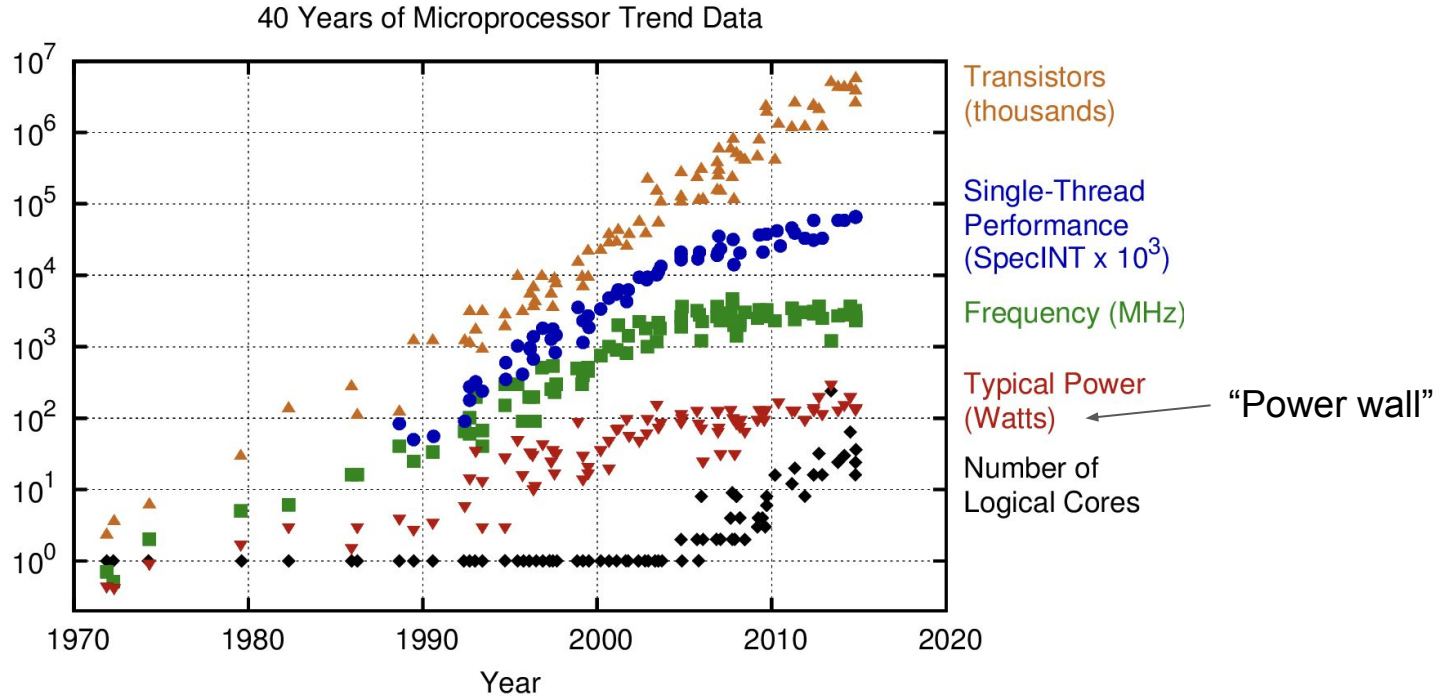
Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2015 by K. Rupp

https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/

Parallelism 2/2

- Expressing parallelism at the data and task level will be key

- Clock speed of CPUs basically ceased to develop

- Moore's law still valid: more and more transistors in our processors

- These materialised as:

  - Vector units: allow to apply the same instruction on multiple data ("data parallelism")

  - CPU-cores: allow to formulate the problem in tasks executed concurrently ("task parallelism")

Heterogeneity

- For two decades HEP benefited from a very uniform hardware landscape, dominated by the x86 CPU architecture

- {Data, High Performance} Computing centres feature today various CPUs, x86 but also ARM and IBM Power

- And heterogeneous architectures, too: most of the computing power is expressed through accelerators, e.g. GPUs

- Challenging devices to program, radically different way of thinking to the memory, e.g. offload calculations, asynchronous traffic between host and device

High Performance Computing 1/2

- HPC: an asset to address scientific challenges and increase the competitiveness of industry
- Substantial national and supra-national investments have been made in this area
- HPCs are likely to be part of the future HEP computing infrastructure
- HPCs approaching the Exascale (10^{18} floating point operations per second capacity)
  - A veritable race, the first exascale HPC will be there well before Run 4
- How much is an EFlop?
  - LHC needs today the equivalent of ~30 PFlops
  - A single Exascale system enough to process 30 "today" LHC
  - A single Exascale system could process the whole HL-LHC with no software improvement or computing model change
  - Of course unrealistic, but just to give you an idea
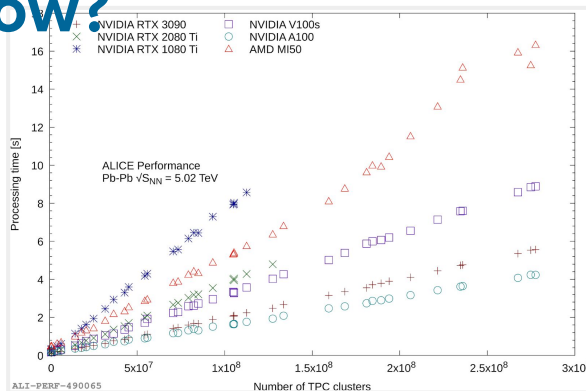
High Performance Computing 2/2

- There are hurdles to overcome to use HPCs for HEP
  - HEP and HPC: **language spoken by experts can be different**
  - **Data access** (access, bandwidth, caches ...): HEP has data processing applications (HTC)
  - **Submission of tasks** (MPI vs Batch systems vs proprietary systems)
  - **HPC: Handful computational kernels VS HEP: thousands of small kernels**
  - Environment **less open than Grid** one (OS, access policies, ...)
  - **Node configuration** (low RAM/Disk, ...)
  - Primary **architecture** (x86_64, Power9, ARM, proprietary, ...)
  - Relationships between Tier-{1,2} providers and **HEP experiments are decades long**
- **And virtually all HPCs offer accelerators, e.g. GPUs**
  - Not being able to use GPUs might prevent HEP to even access HPCs because of existing policies

Solving those hurdles requires also investing in people

GPUs: Where are we now?

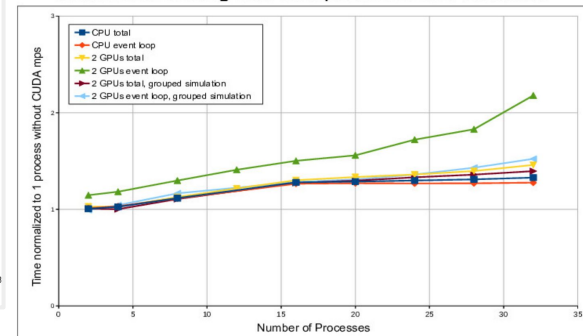Alice, continuous TPC readout; LHCb, fully software-based triggers

- ALICE leveraging AMD GPUs for real-time processing, online data reduction and compression
- LHCb's Allen framework using ~200 NVIDIA GPUs for HLT1 to handle 30 MHz input rate
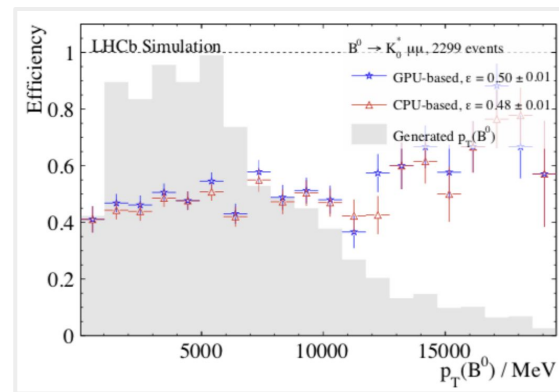
CMS using NVIDIA GPUs for part of the HLT sequence

- Higher throughput and equivalent results to CPU
- Equal or better physics performance
- Will run at HLT in 2022
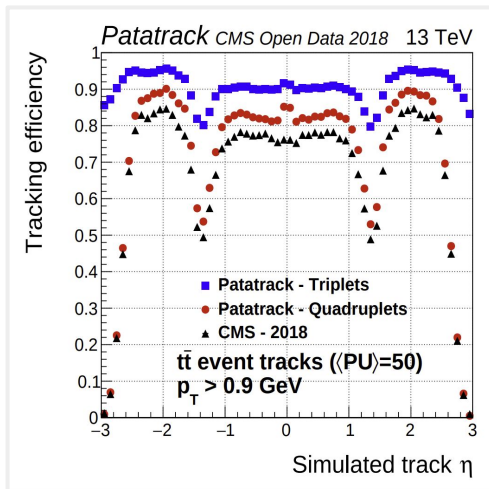
ATLAS

- Promising LAr Calorimeter reconstruction porting to GPU targeted code

FastCaloSim Scaling With Multiple Concurrent Processes

Patatrack CMS Open Data 2018   13 TeV

Bocci, Kortelainen, Innocente, Panteleo, Rovere (2020), arxiv:2008.13461, *Bozzi* (2021) Web seminar , *Rohr* (2020) arxiv:2006.04158, Leggett (2021) vCHEP

How to evolve from here

- A bright future of HEP ahead of us: many questions to answer, discoveries made accessible thanks to the HL-LHC program
- More events, more complex, at a finer granularity
  - Processing and analysing them is a challenge for us
  - To be faced with a constantly evolving hardware landscape
- Two elements which help succeeding: parallelism and heterogeneity
- Interesting times for motivated scientists-engineers!
  - A single person can make the difference
- A lot has been achieved, but a lot still needs to be done
- This tCSC is an opportunity for getting nearer to the solution, as a community