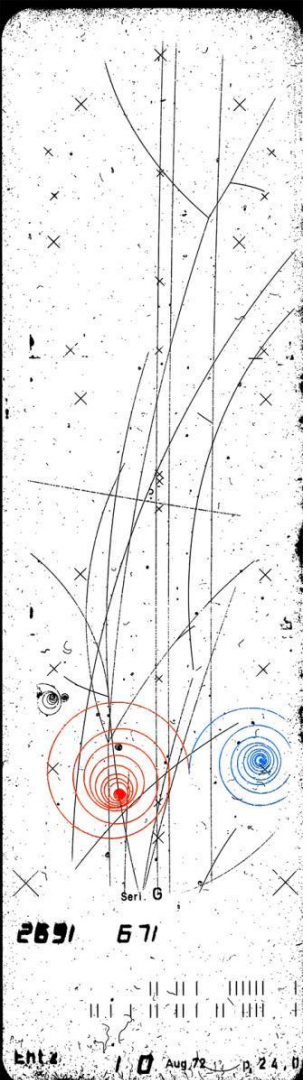

Computing at CERN

Maite Barroso Lopez CERN IT department

CERN School of Computing

What is Physics Computing?

- Yearly input: A few petabytes of data
- Yearly output: A few hundred physics papers
- Data reduction factor of 10^7 to 10^8 !!
- How is it done?

What do we need?

- **Millions of lines of code (C++,Python, …)**
- **Large infrastructure**
  - Customized hardware
  - PC farms
  - Database and storage systems, data management systems
  - Distributed analysis facilities
  - The grid

What happens to the data?:

- **Event filtering, tagging and storage**
- **Calibration, alignment**
- **Event reconstruction**
- **Storage**
- **Event simulation**
- **Physics analysis**

Computing at CERN: the big picture

Data

Detector

Data Acquisition (trigger)

Data Recording and Archive

Data Reconstruction and Derivation

Simulated Data

Event Generation(part. inter.) Simulation(int. w/ det.) Digitization(simul detector resp)

Data Analysis

HPCs

Preservation: Data and Analysis

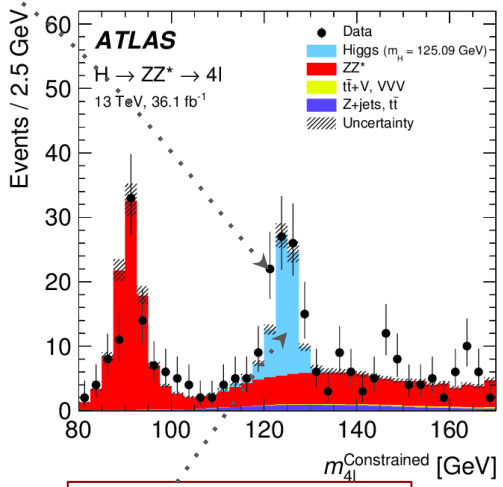

Step by step

- **Each step involves some data reduction**
  - data are discarded (online)
  - data are compressed (offline)
- **In each step the data get closer to be interpretable in**
- **physical terms**
- **Some steps are repeated many times until the output is**
- **satisfactory (offline reprocessing)**

Our Data *Generators*

Experimental physics is about **DATA**!

CERN experiments transfer all data to the Computing Center

Data is stored and ready to be processed, analyzed and shared

A comparison of the yearly data volumes of current and future projects:

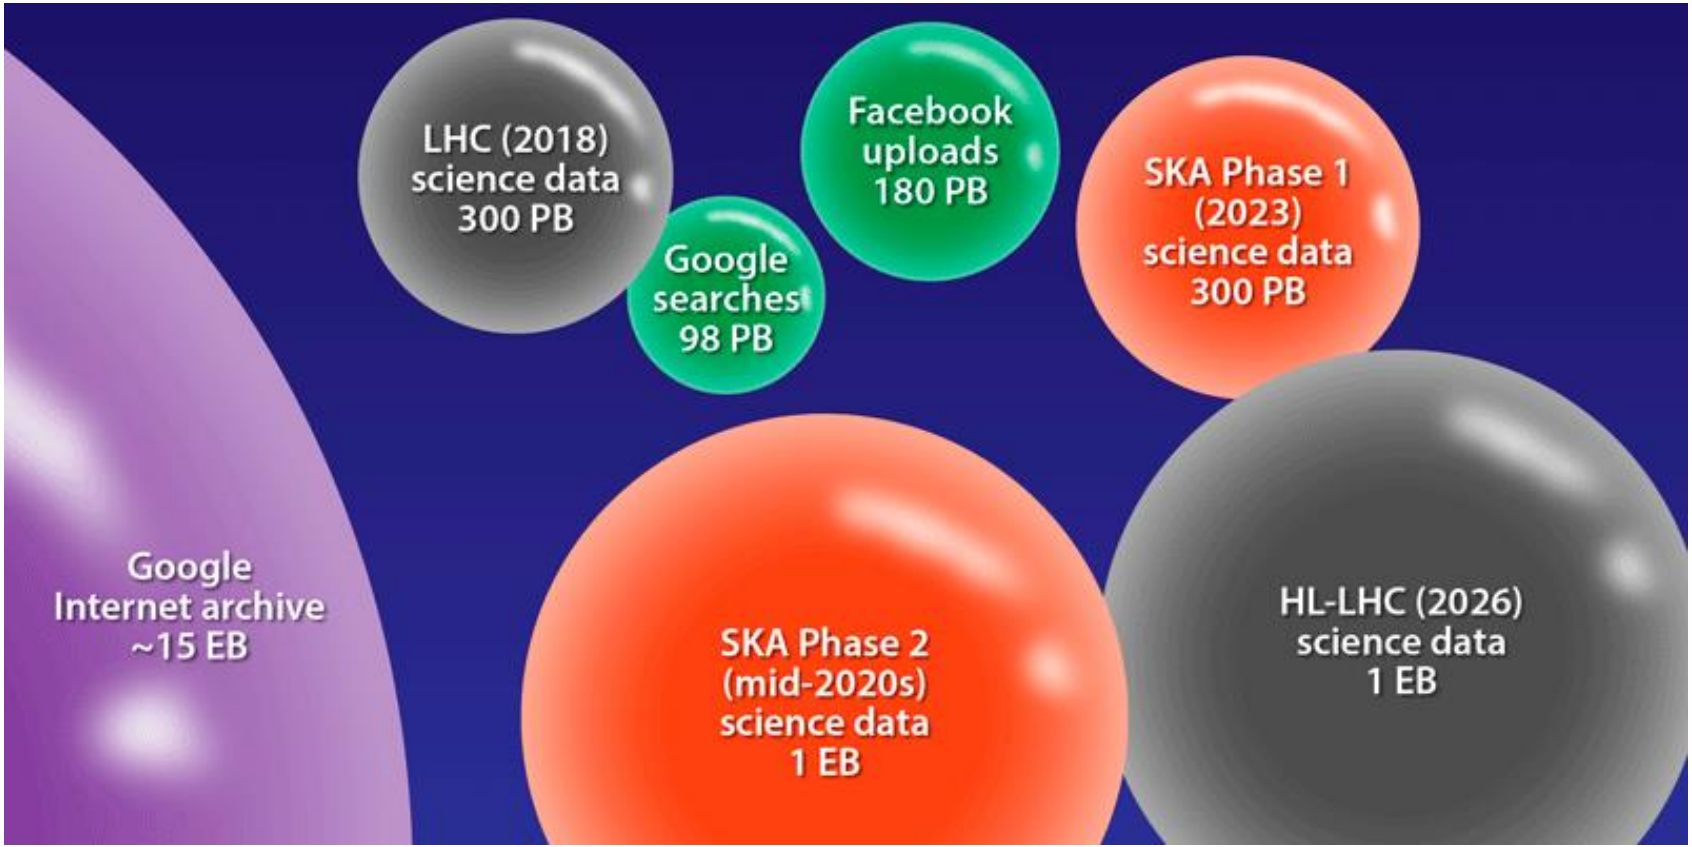

Credits: APS/Alan Stonebraker and V. Gülzow/DESY

LHC Experiments: the *big* 4

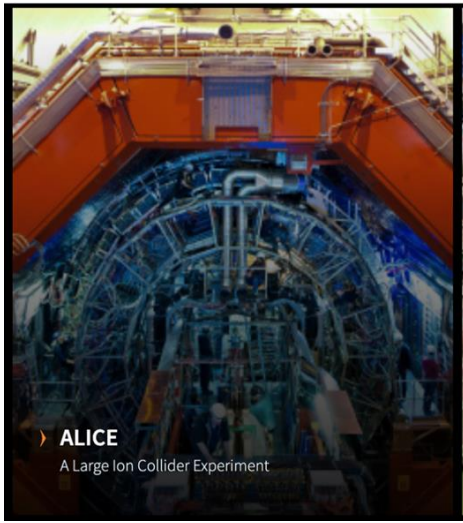

heavy-ion physics, strongly interacting matter at extreme energy densities (quark-gluon plasma)

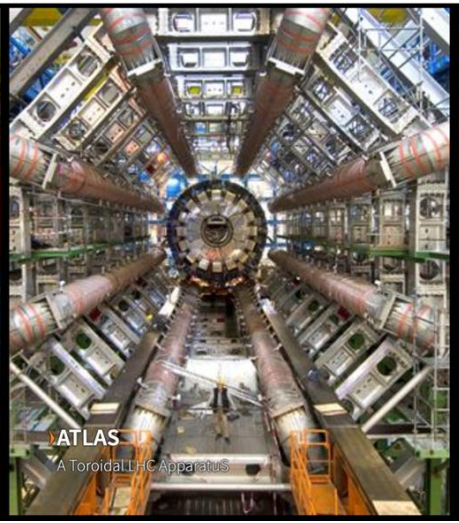

General-purpose:Higgs boson, extra dimensions, dark matter

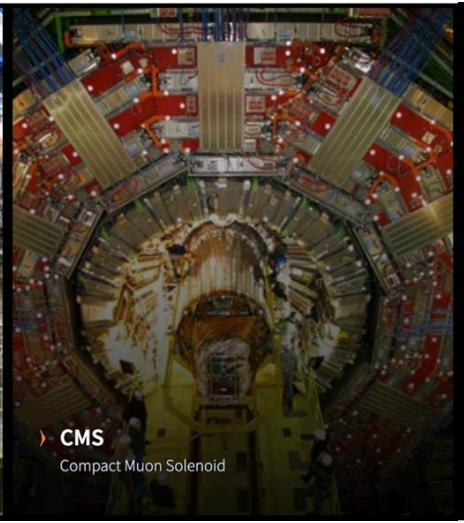

General-purpose:Higgs boson, extra dimensions, dark matter

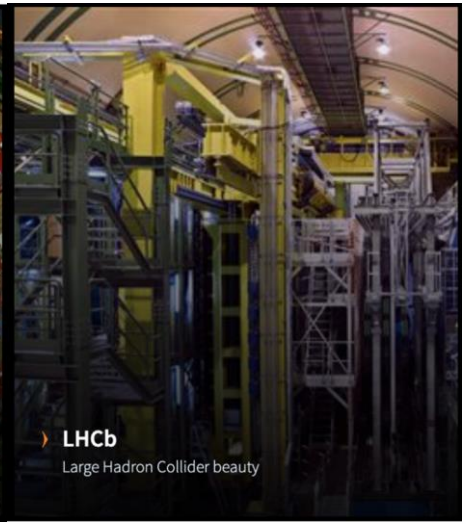

CP violation in the interactions of b-hadrons: matter-antimatter asymmetry of the Universe.

A computing perspective: O(10-100) Pb/year, 0.5M running jobs (continuous), fat and long-haul networks

Plus more experimets

- Fix target (COMPASS, CLOUD, NA64, etc)
- Antimatter (AEGIS, ALPHA, etc)
- Non-accelerator (AMS, CAST, OSCAR)
- Experimental facilities (ISOLDE, Medicis, n_TOF, Neutrino platform)

Astroparticle Physics and Radio Astronomy

Knowledge transfer and expertise sharing among large ESFRIs on scientific computing

Data Recording from remote sources

Computing at CERN: the big picture

Data

Detector

Data Acquisition (trigger)

Data Recording and Archive

Data Reconstruction and Derivation

Simulated Data

Event Generation(part. inter.) Simulation(int. w/ det.) Digitization(simul detector resp)

Data Analysis

HPCs

Preservation: Data and Analysis

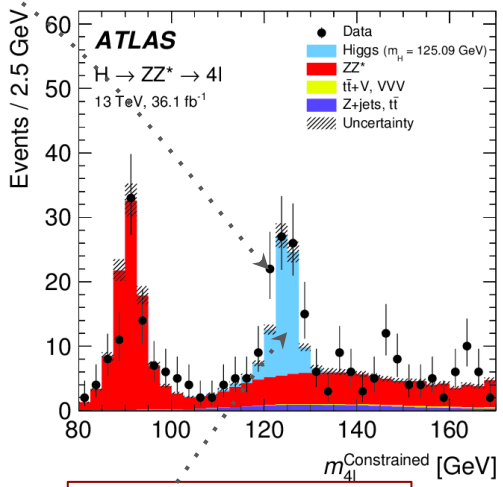

Online versus offline

- **Online**
  - In real time, fast!
  - Decisions are irreversible
  - Data cannot be recovered
- **Offline**
  - From almost real time to long delays
  - Decisions can be reconsidered
  - Data can be reprocessed

Online processing

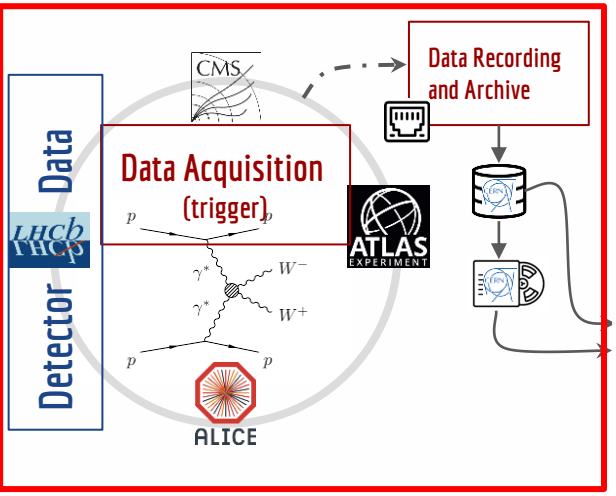

Trigger: event selection

- **Primary (design) collision rate: 40 MHz**
- **Recording rate: a few hundred Hz to kHz**
- **How is this achieved?**
- Multi-level trigger – chain of yes/no decisions
- Very fast first level: (Programmable) hardware
- Slower higher level(s): Software on specialized or
- commodity processors

Has to be reliable

- **Rejected data are lost forever**
- **Continuous monitoring**
- **Do not lose new physics!**
- **Must therefore be open to many different signatures of potentially new physics in the detector system**

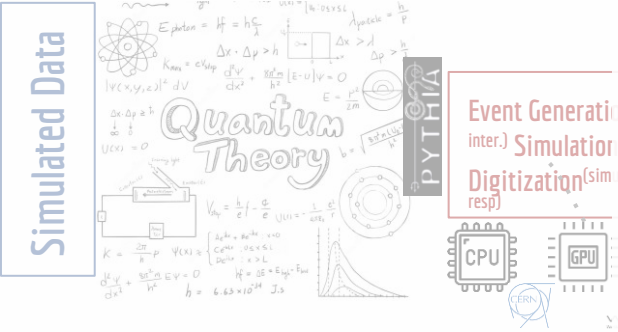

Online processing

Data Acquisition
(trigger)

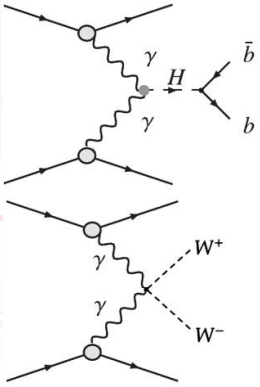

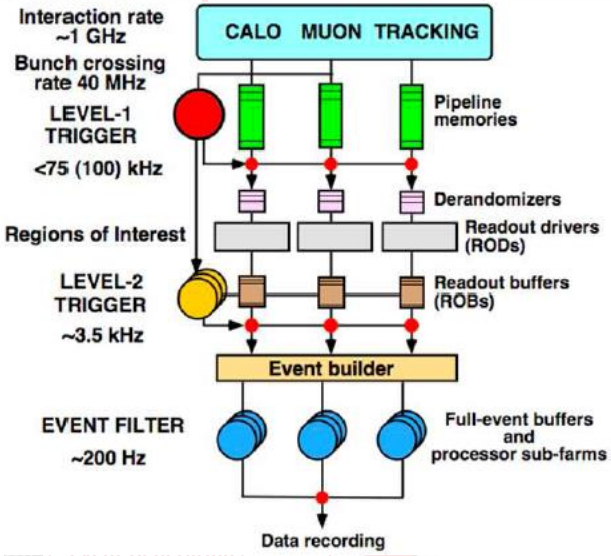

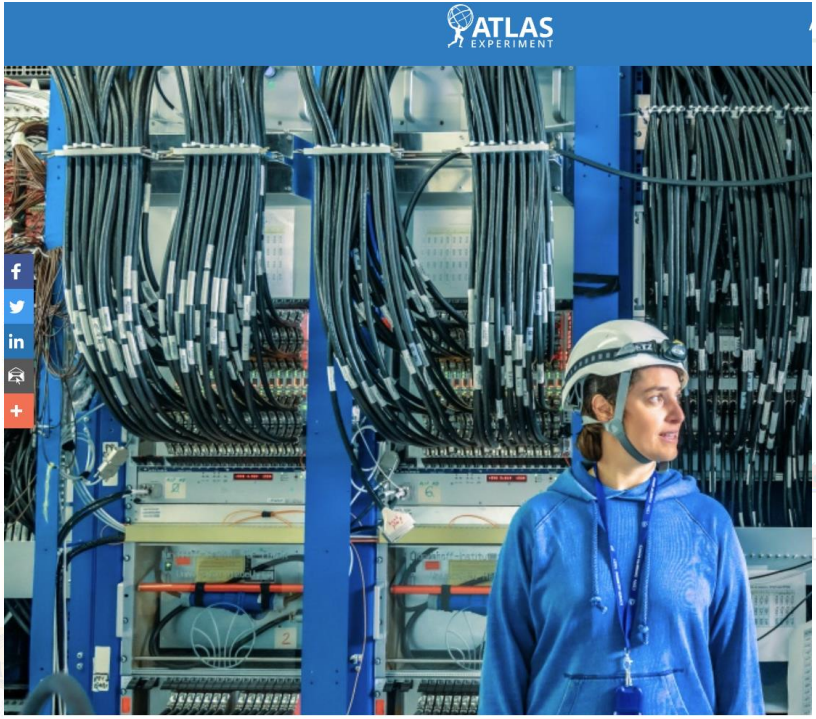

updates > briefing > ATLAS event selection system readies for LHC Run 3

From the Hit to the Bit: DAQ

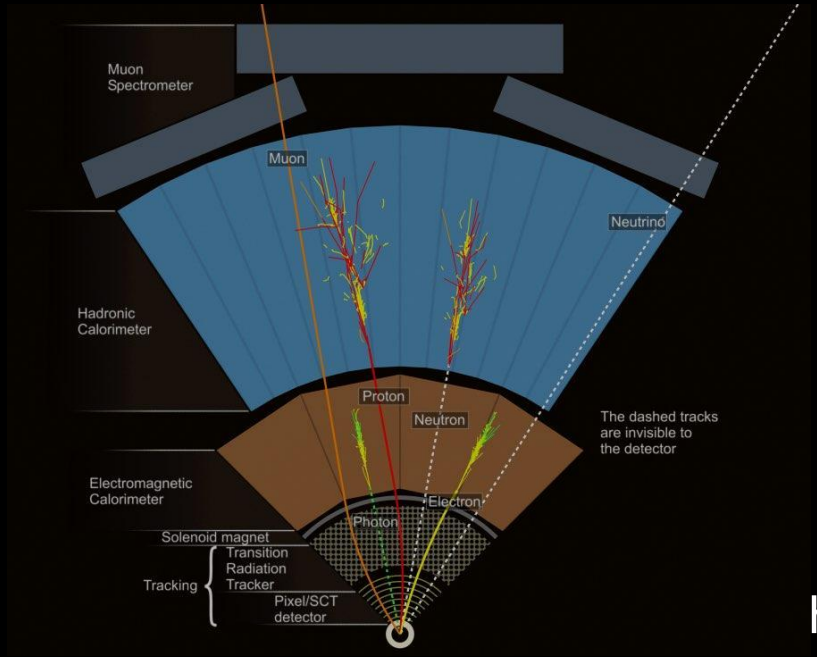

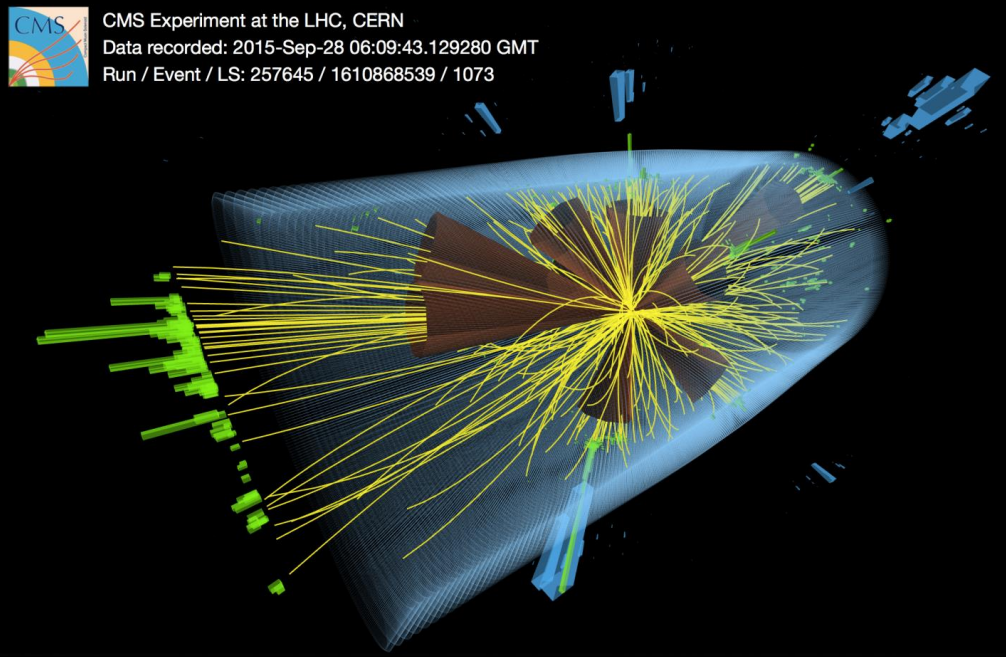

han

40 million pictures a second

Synchronised signals from all detector parts

From the Hit to the Bit: event filtering

- L1: 40 million events per second
  - Fast, simple information
  - Hardware trigger in a few microseconds
- L2: 100 thousand events per second
  - Fast algorithms in local computer farm
  - Software trigger in <1 second
- EF: Few 100 per second recorded for study

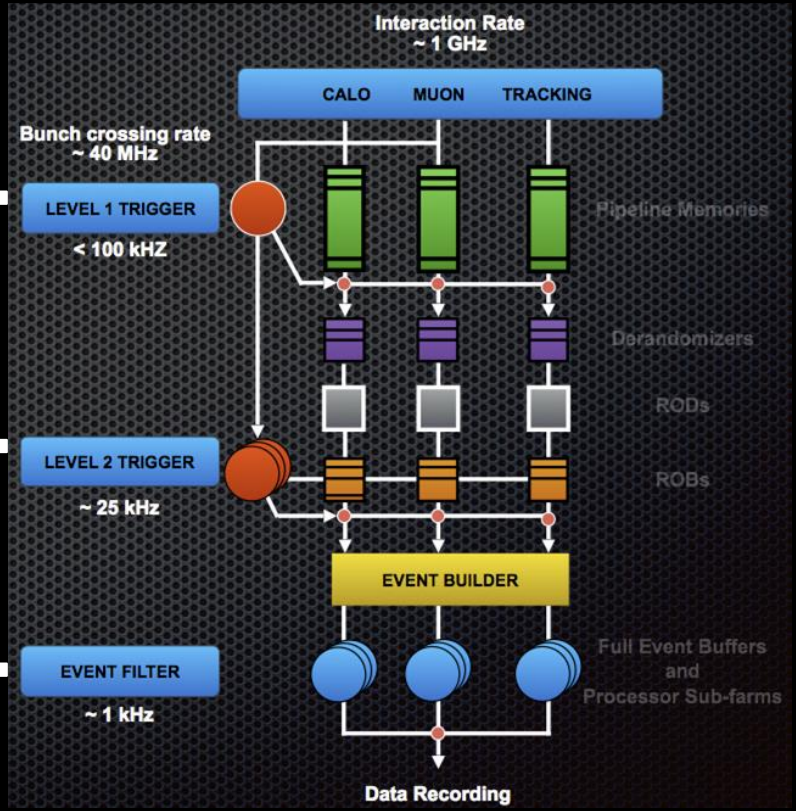

From the Hit to the Bit: event filtering

- L1: 40 million events per second

  1PB/s: few hundred trillion euros/yr !!!
  From 40 MHZ to 100kHz

- L2: 100 thousand events per second
  - Fast algorithms in local computer farm
  - Software trigger in <1 second

- EF: Few 100 per second recorded for study

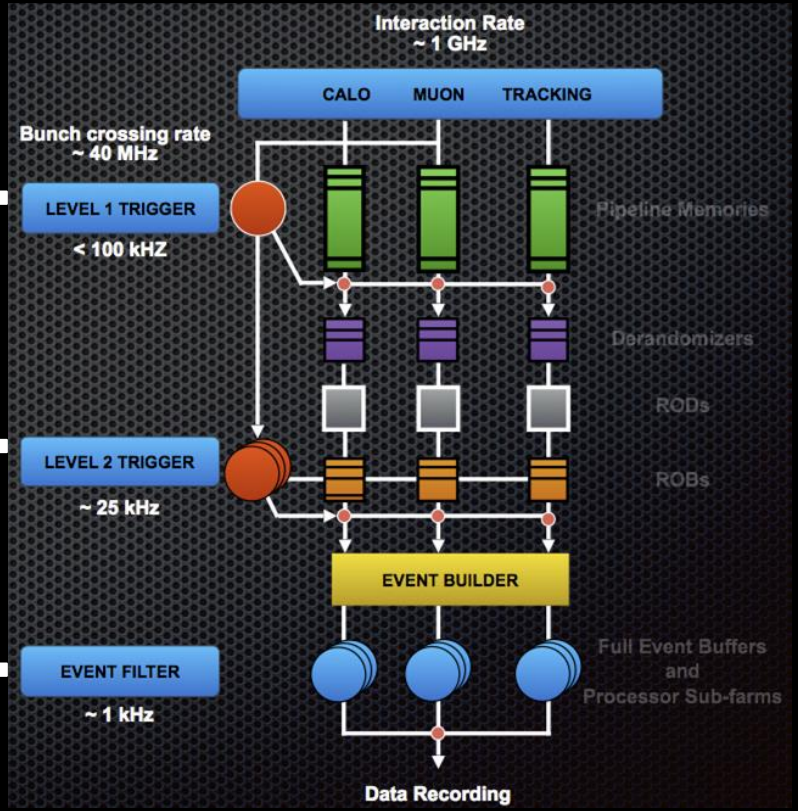

From the Hit to the Bit: event filtering

- L1: 40 million events per second

1PB/s: few hundred trillion euros/yr !!!
From 40 MHZ to 100kHz

-

We keep ~1 event in a million
From 100 kHZ to 25kHz

- Software trigger in <1 second

- EF: Few 100 per second recorded for study

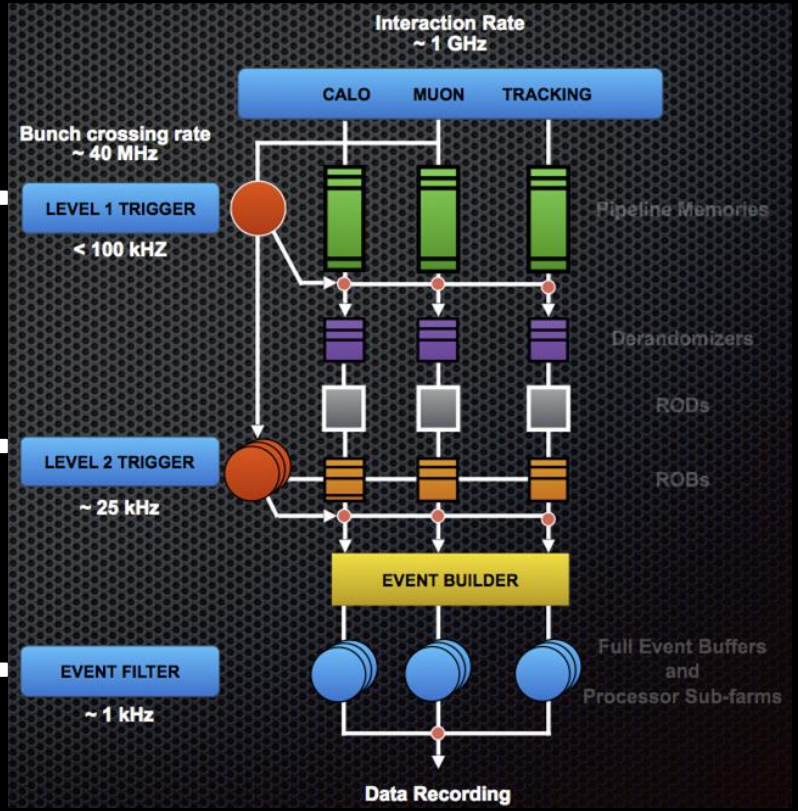

From the Hit to the Bit: event filtering

- L1: 40 million events per second

1PB/s: few hundred trillion euros/yr !!!
From 40 MHZ to 100kHz

-

We keep ~1 event in a million
From 100 kHZ to 25kHz

- Software trigger in <1 second

-

From 25 kHZ to few kHz
Final data rates ~10 GB/s

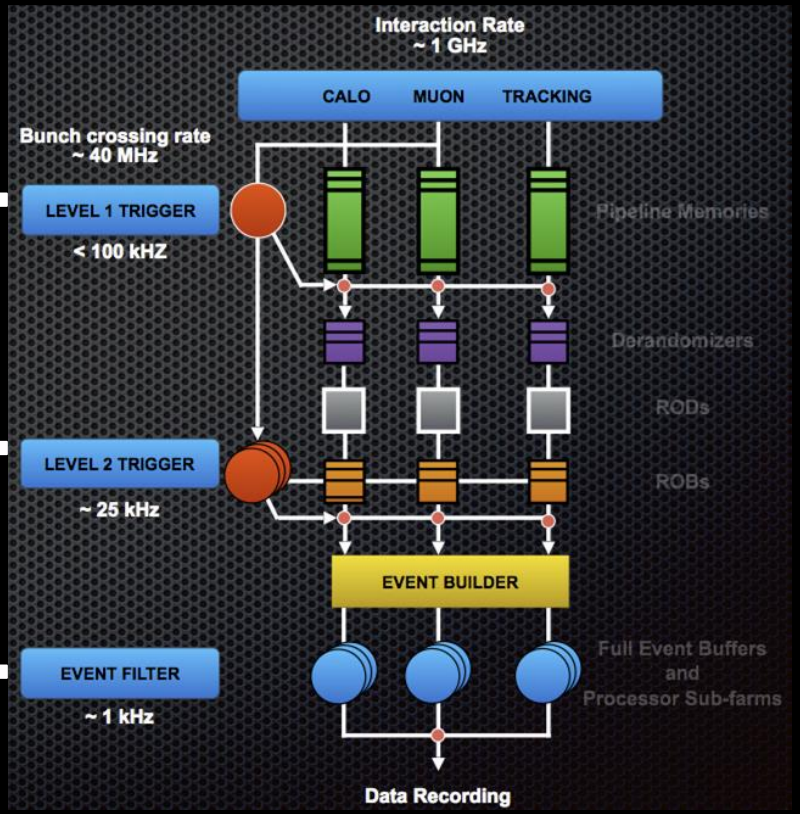

Offline computing

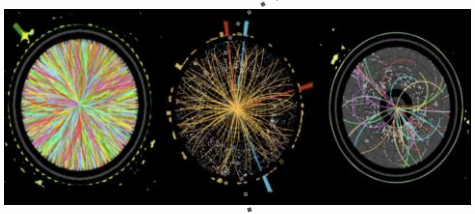

Data Reconstruction and Derivation

Data Analysis

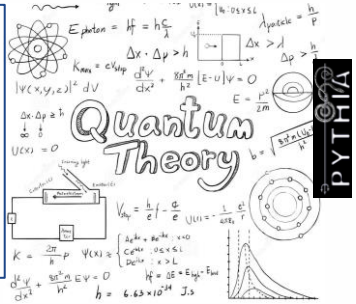

Simulated Data

Event Generation(part. inter.) Simulation(int. w/ det.) Digitization(simul detector resp)

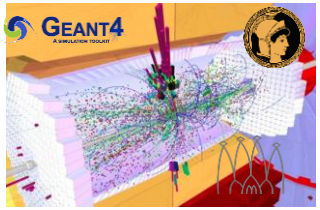

HPCs

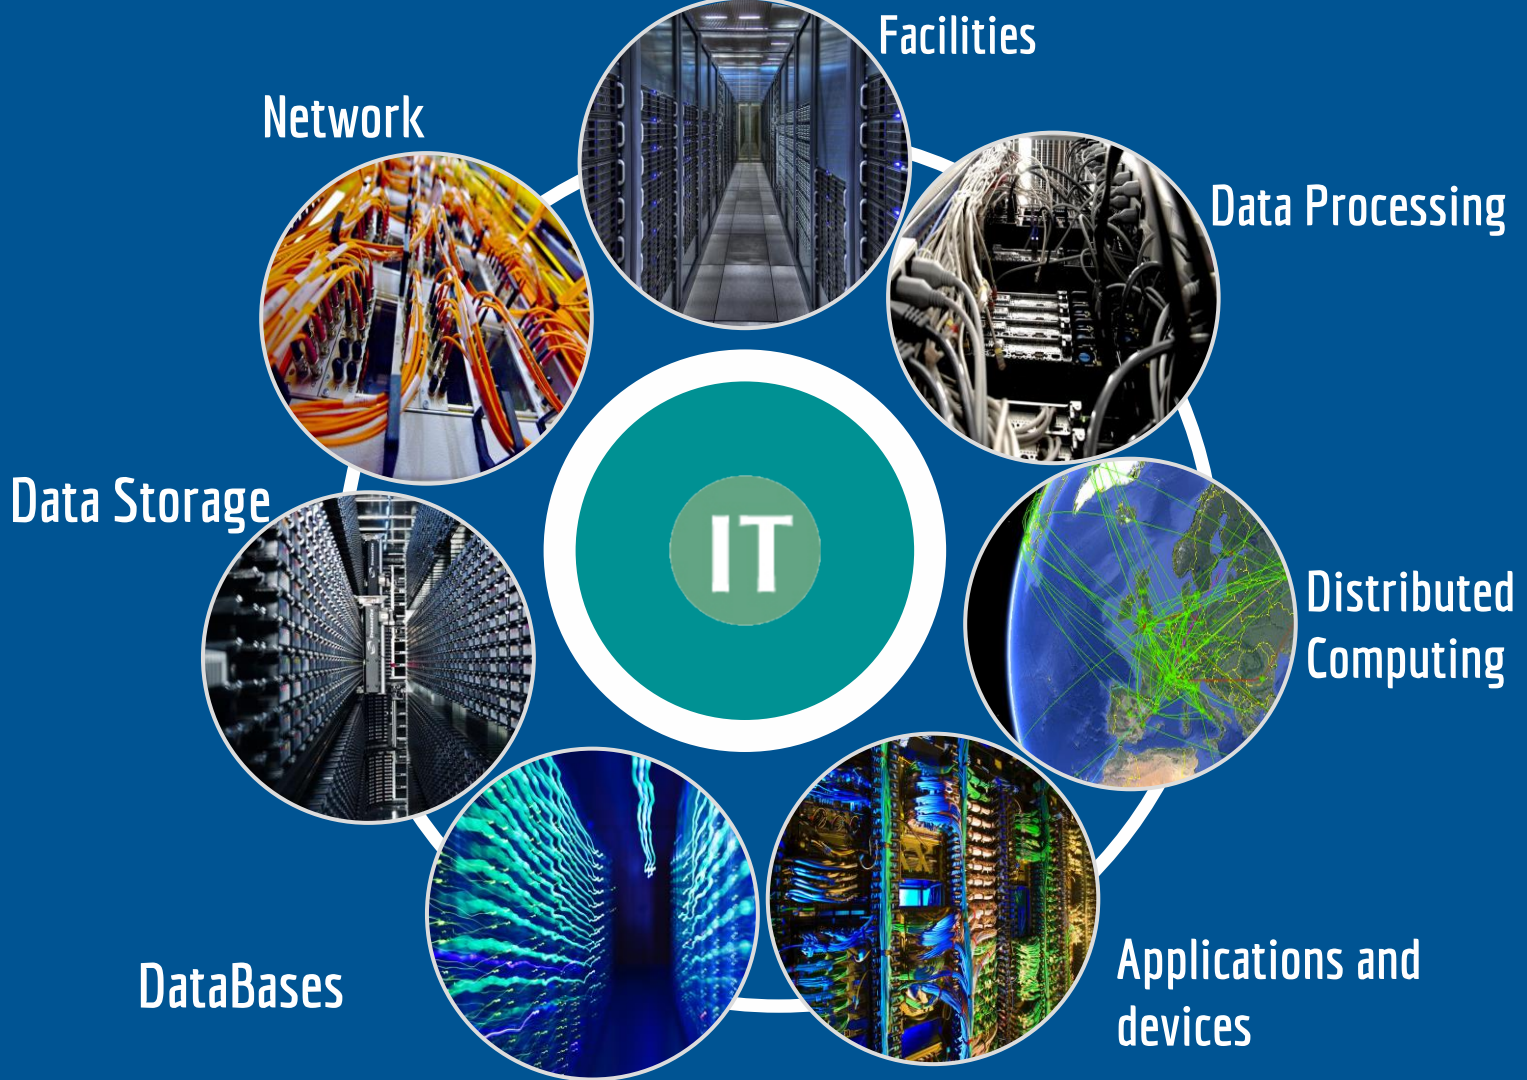
Facilities
Network
Data Processing
Data Storage
IT
Distributed Computing
DataBases
Applications and devices
CERN IT Department

Offline computing

Calibration: Convert raw data to physical quantities

Alignment: Find out precise detector positions

Event reconstruction:

- Reconstruct particle tracks and vertices (interaction points)
- Identify particle types and decays

Simulation:

- Generate artificial events resembling real data as closely as possible
- Needed for background studies, corrections, error estimation, …

Offline: data analysis

- Extract physics signals from
- background
  - Compute masses, cross-sections, branching ratios, discovery limits, …
- Requires sophisticated multivariate techniques
- **Distributed: Physics analysis takes place in many labs all over the world**

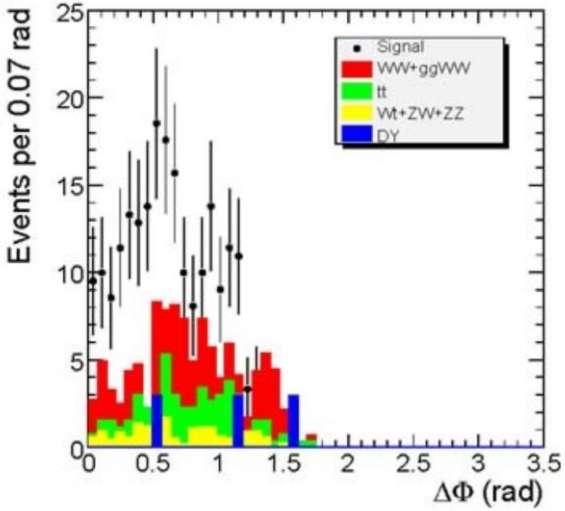

The Worldwide LHc computing Gri

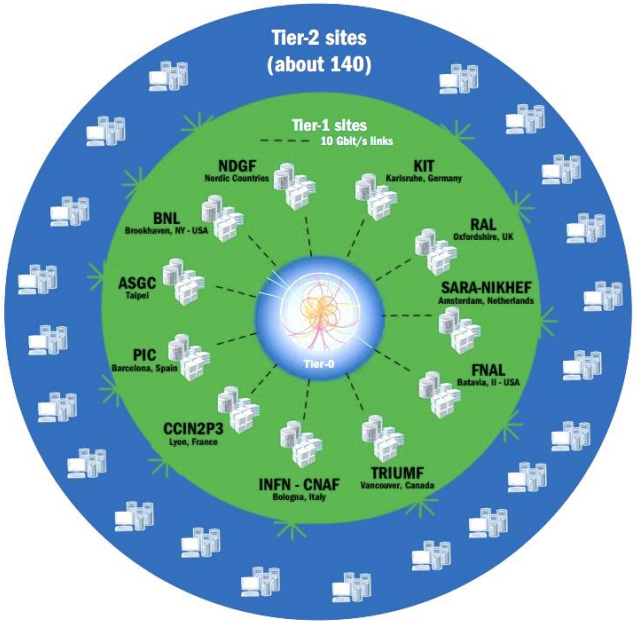

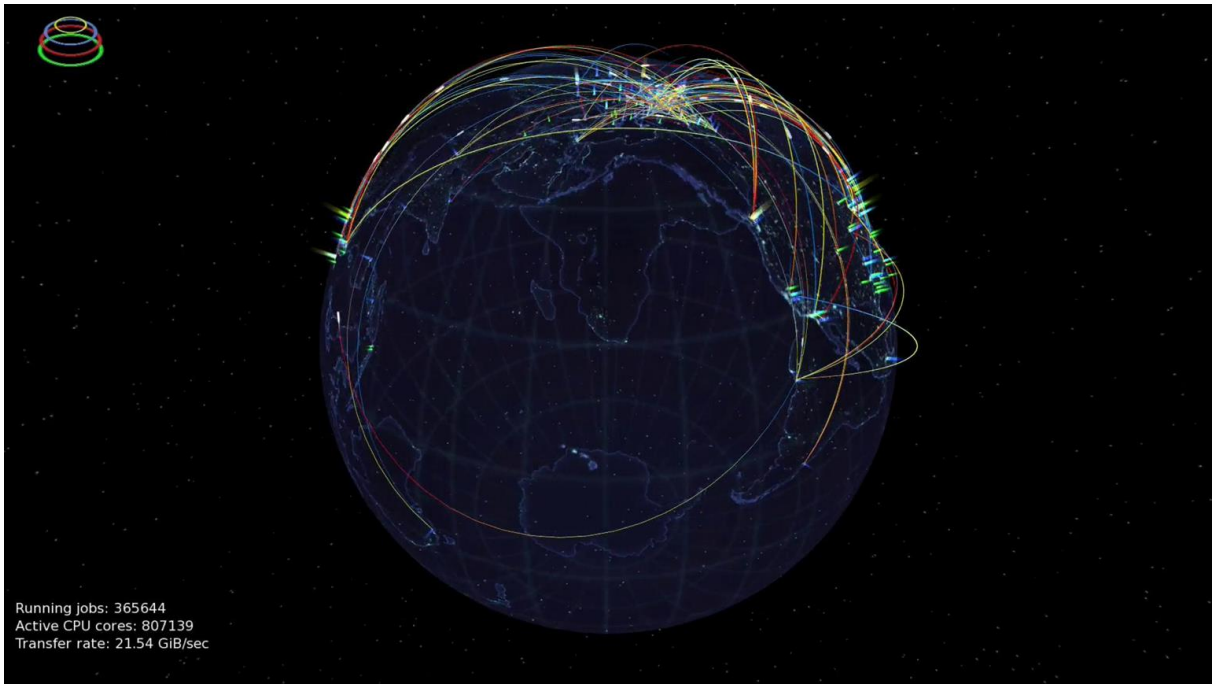

WLCG provides global computing resources to store, distribute and analyse the LHC data

CERN School of Computing

The LHC Computing Grid

- Global collaboration of more than 170 computing centers in 36 countries
- Four-tiered model
- Data storage and analysis infrastructure
- O(10^5) CPUs
- O(100) PByte disk storage (tiers 0 and 1)

The Worldwide LHc computing Gri

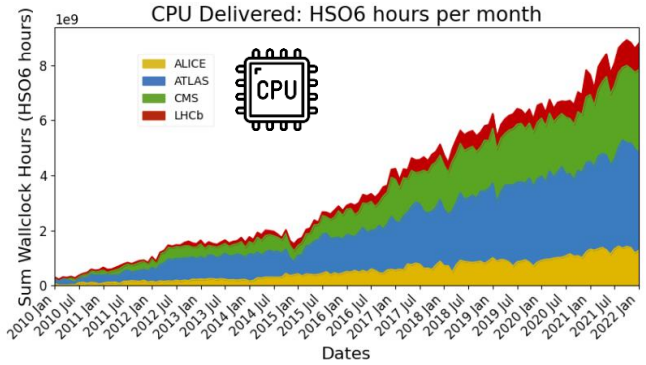

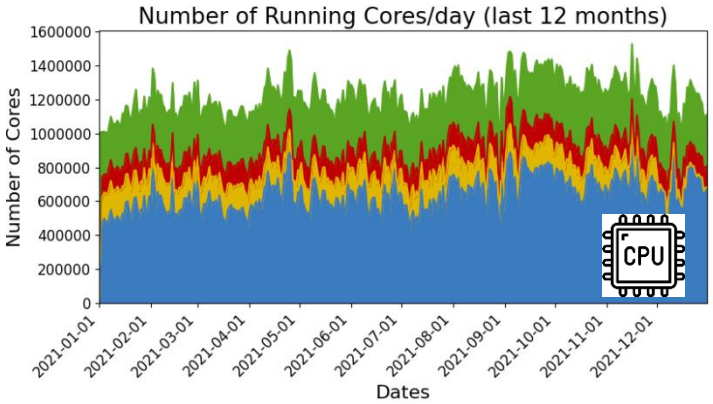

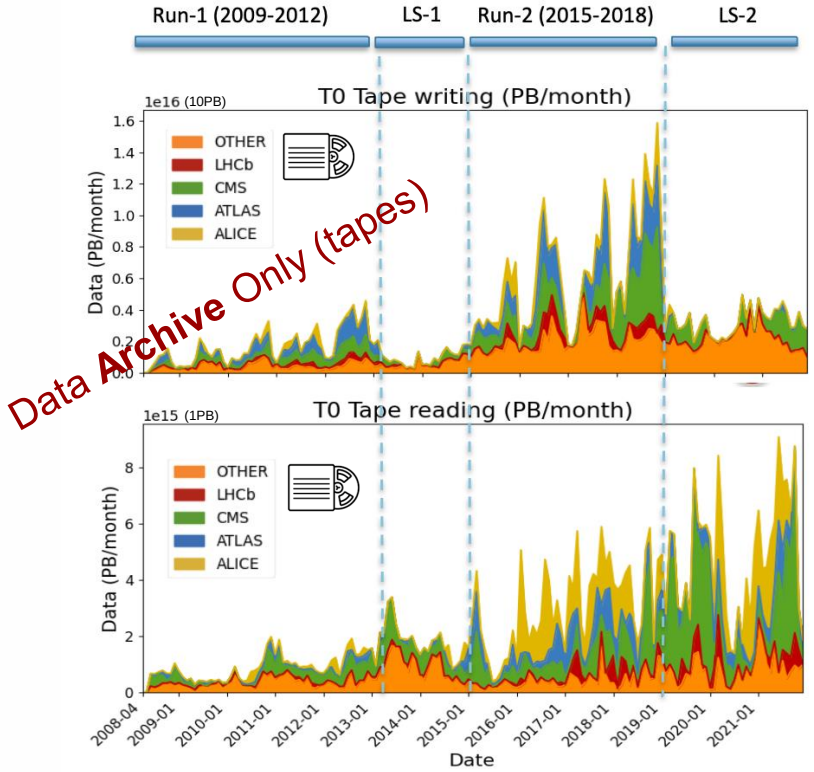

Data flow in ATLAS

200 Hz
RAW: ~1.6 MB/evt

Event Summary Data (ESD): ~ 1.5MB/evt
Analysis Object Data (AOD): ~ 150 kB/evt
Derived data (dESD, dAOD, NTUP, …)
distributed over grid

Calibration

CERN Analysis Facility

Tier-0

Data Recording to tape
First Pass Processing

10 Tier-1 centers
RAW data copy on tape
Analysis data on disk
Reprocessing

Tier-1

Tier-1

Tier-1

37 Tier-2 centers
Analysis data on disk
User Analysis

Tier-2

Tier-2

Tier-2

The last LHC Run: data recording

- LHC Experiments recorded 88 Petabytes of data in 2018 (15.8PB only in November)
- The LHC data is aggregated at the CERN data centre to be stored, processed and distributed

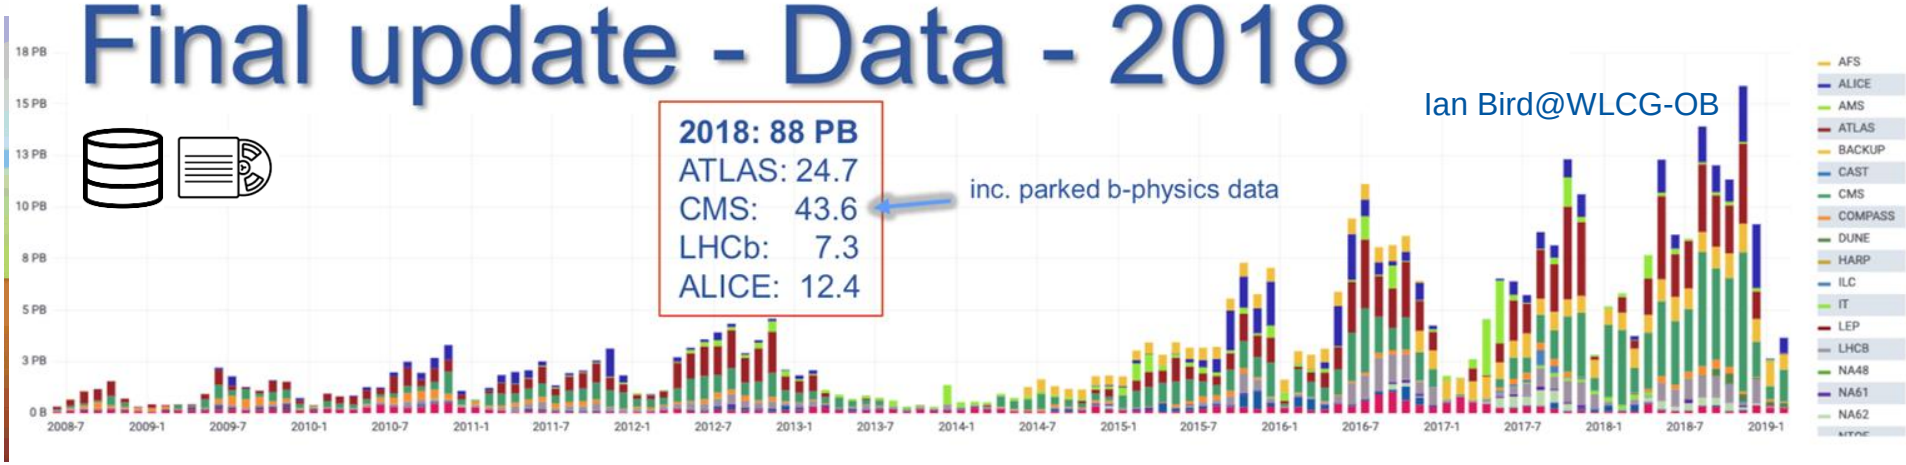

2025 LHC Data Taking

Daily average data rate from experiments to T0 (Gb/s)

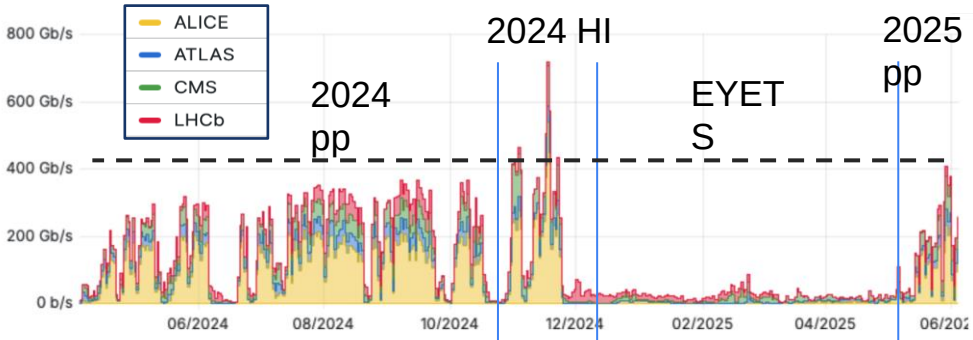

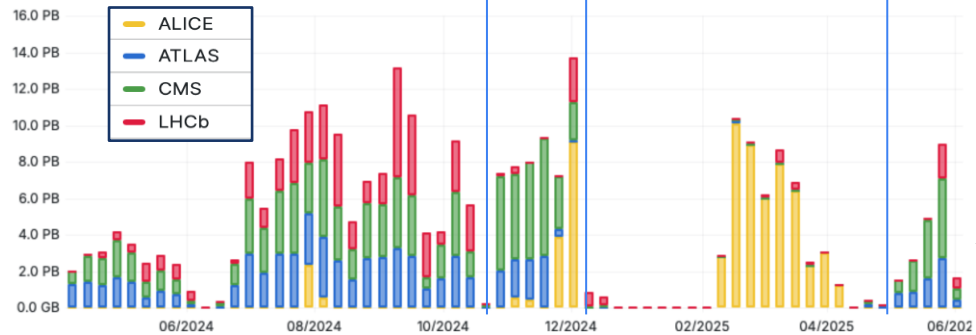

Recorded data into the CERN tape archive (PB/week)

The 2025 LHC data taking started

Data rates higher than in 2024 for the same period, inline with expectations

CERN and generally WLCG services running smoothly

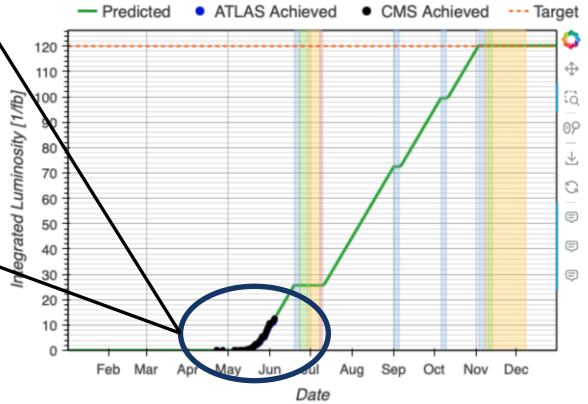

LHC data processing

Number of cores in use (WLCG pledge + opportunistic)

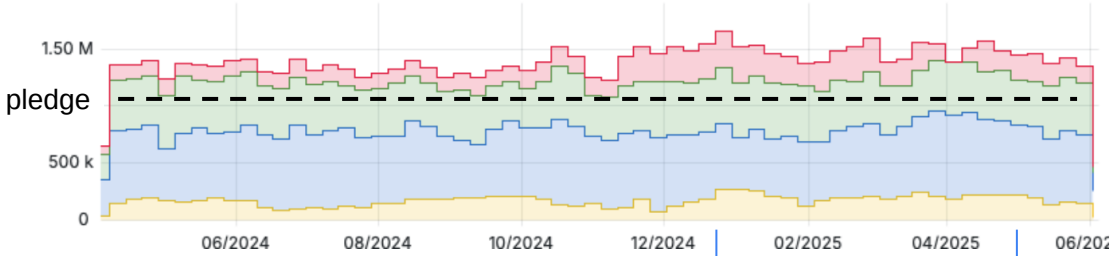

Experiments leveraging all WLCG available cores …

… plus ~50% extra opportunistic capacity (including HLTs and HPCs)

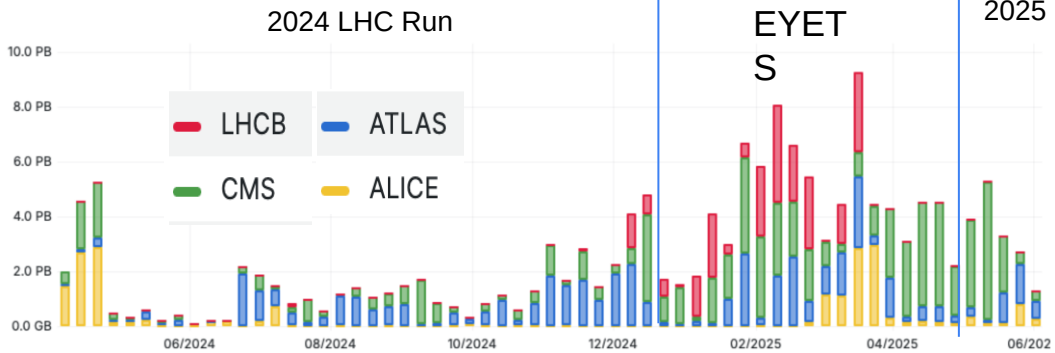

Intense reprocessing campaigns during EYETS recalling data from tape systems

Volume of data recalled from CERN tape (PB/week)

HL-LHC: a computing challenge

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!

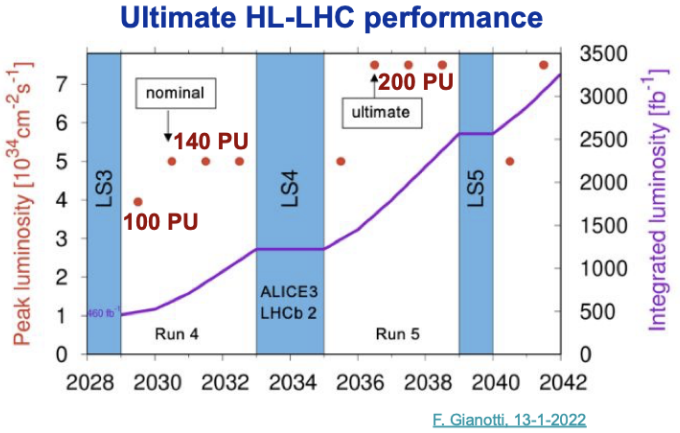

HL-LHC upgrades have to be paired to an equally ambitious upgrade of data processing software, computing tools and services.

By delivering more integrated luminosity, more parasitic collisions per bunch crossing are generated, more complexity to be handled by our algorithms!

Offline projected computing needs in the HL-LHC era

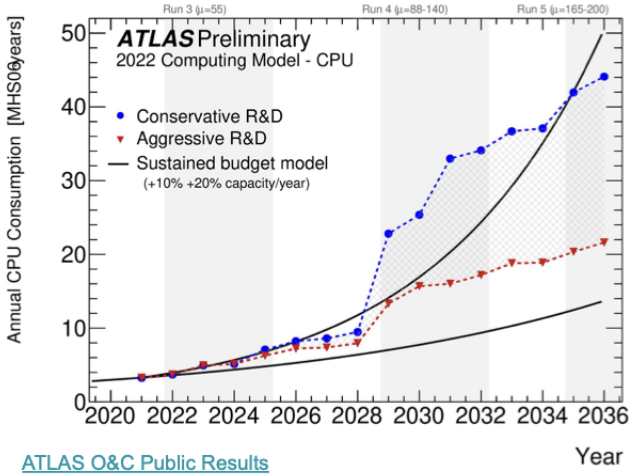

ATLAS O&C Public Results

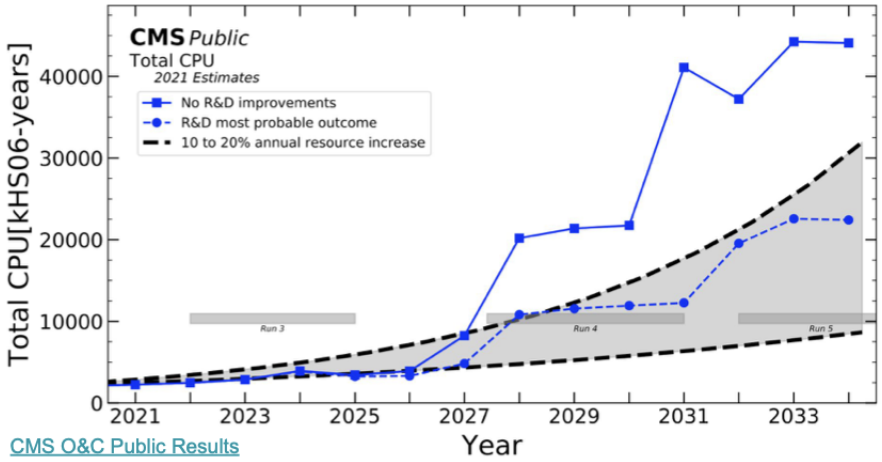

CMS O&C Public Results

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!**

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 RISC-V in the future?)
- 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
- Not embracing heterogeneous platforms could mean making HEP data processing “legacy computing”

CERN Data Center

- Built in the 70s on the CERN site (Meyrin,Geneva)
  - 3.5 MW for equipment
- Hardware generally based on commodity
  - 10,000 servers, providing 500,000 processor cores (14.000 VMs)
  - 120,000 disk drives providing 0.6EB disk space (1EB=1000PB)
  - 40,000 tapes drives, providing 0.5EB capacity (1EB=1000PB)

Green IT

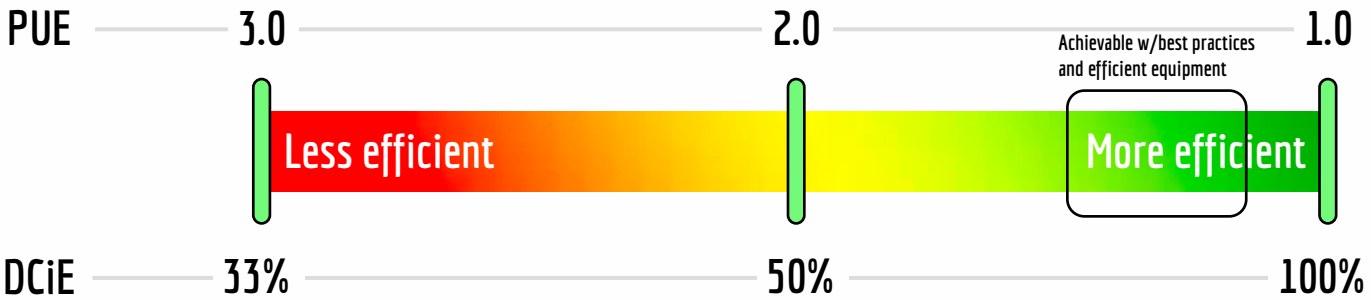

Power Usage Effectiveness (PUE): Total Facility Energy / IT Facility Energy

Take-away (1)

- Computing is instrumental for science. Detectors and sensors evolving: growing IT demands.
  - LHC raw data rates are PB/s scale but lowered to GB/s after data *filtering*
- CERN remains largest scientific repository in the world
- Scientific data already at the **Exabyte scale**:
  - 1EB = 1.000PB = 1.000.000 TB = 1.000.000.000 GB

(1TB is your computer) (100 GB is your smartphone)

Take-away (2)

- 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 lot has been achieved, but a lot still needs to be done

Credits

- Xavier Espinal (material from previous years)
- Danilo Pipparo (tCSC 2024)
- Simone Campana (WLCG material)
- Arnulf Quadt (CERN School of Computing 2024)