ATLAS and CMS @ HL-LHC Computing models and CPU needs – with a QC bias

Tommaso Boccali (INFN Pisa / CERN)
with inputs by Davide Costanzo (Sheffield / CERN)

Outline

- CMS and ATLAS
  - Design choices, and impact on Computing
- Current LHC Run – back of the envelope computing needs
  - Actual Implementations: Computing Model!
    - With a focus on CPU needs – not storage
- Extrapolations to HL-LHC
  - Current resource estimate needs
  - Mitigations
- Focus on QC
  - What could we use it for? (hopefully we will learn better today!)
  - CPU needs: characterization of the major use cases

CMS and ATLAS specificities

- The two general purpose detectors @ LHC
- Largely different design decisions, but physics potentials have been comparable up to now
- Some typical design decisions:
  - Reconstruction of Charged particles essential to LHC Physics → large **tracking systems**, 10-100+ M acquisition channels
  - Most interesting physics signatures with jets or leptons in the final state
    - calorimeters for jets and electrons, large muon systems taking most of the detectors' volume

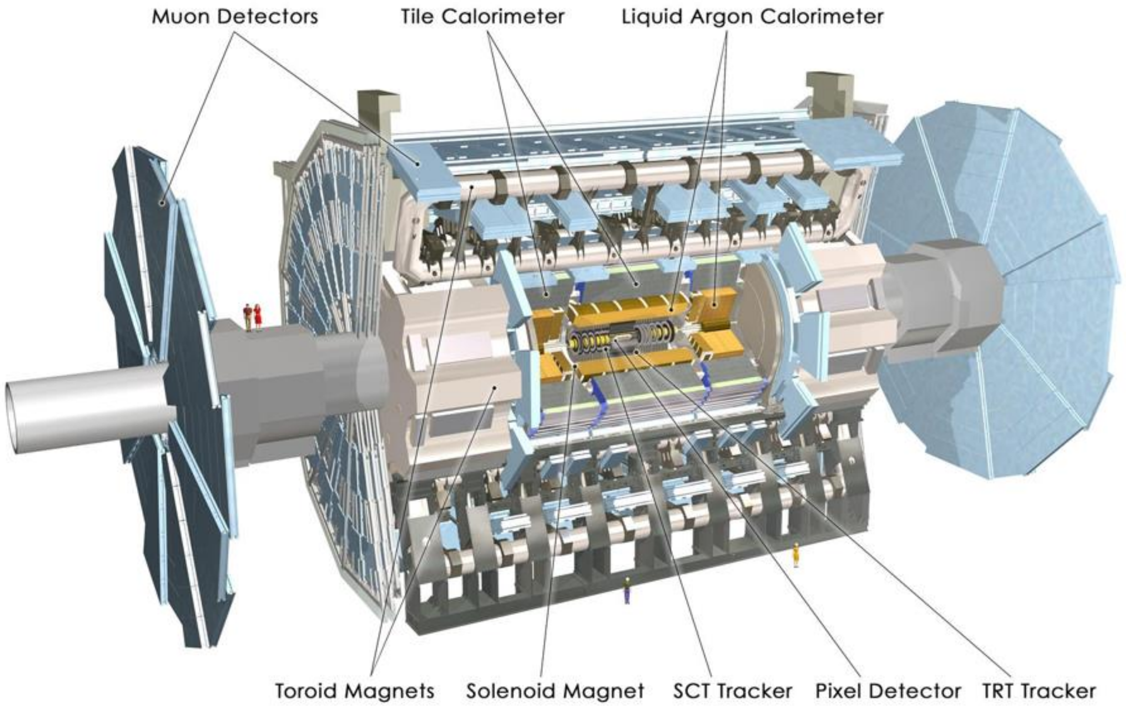

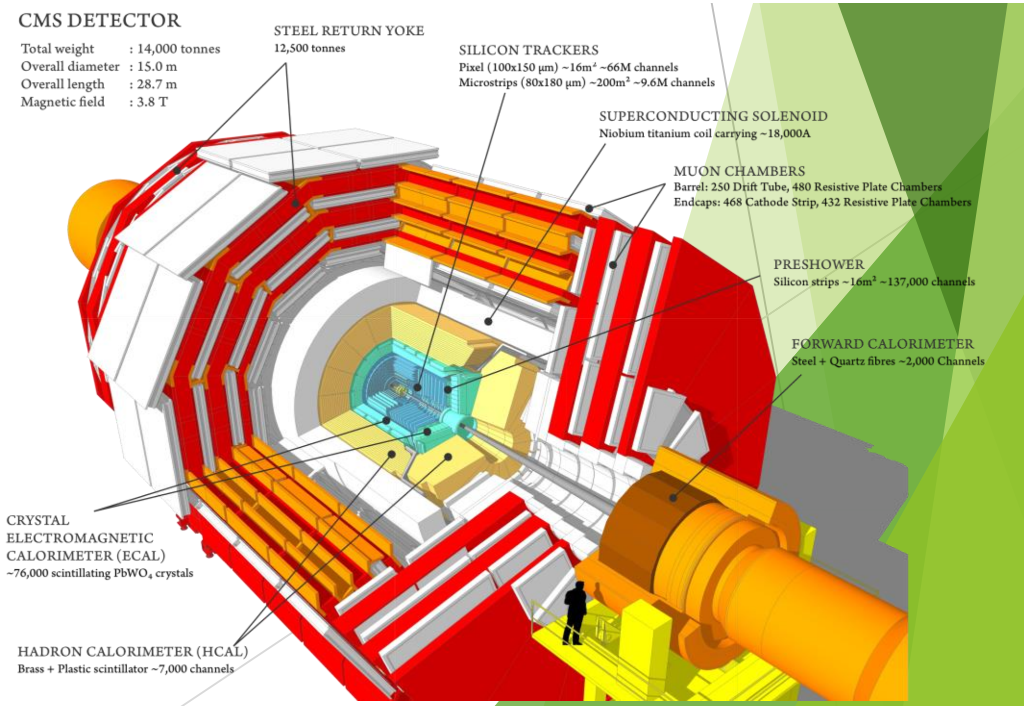

Back of the envelope computing parameters

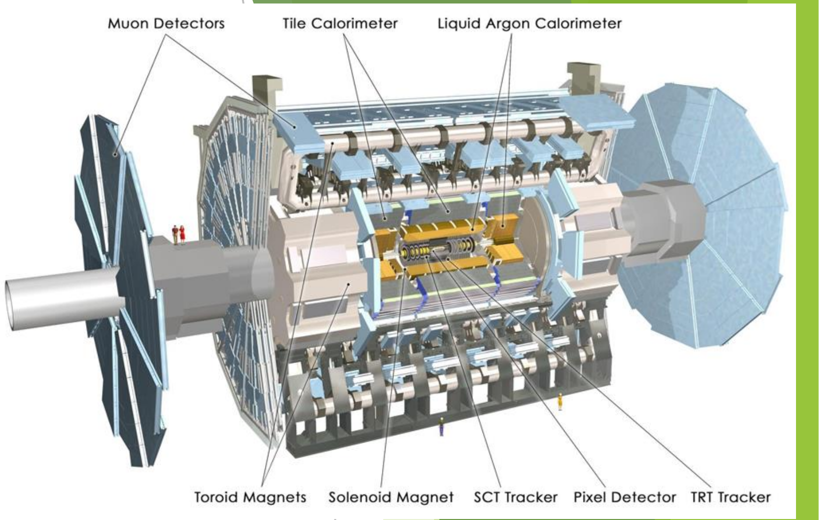

- If was even possible to collect all the channels for each and every LHC bur collision, we would end with (~)
  - 40 MHz of events
  - With 100+ M acquisition channels
  - Say 1 byte per channel (very rough approx)
  - → **4 PB/s**

- Clearly unfeasible; what reaches offline realm (difficult to think about QC at online level..) is
  - A strongly **zero suppressed** set of acquisition channels (so down to ~ 1 MB/event; reduction O(100-1000))
  - **1000 Hz** of events selected (reduction 1/40000)
  - The **persistent** events reaching offline are ~ **1 GB/s**

Various levels of Filtering

Offline world
(tapes, disk,…) and then
analysis!

Reality

LHC collisions

Decay of unstable particles

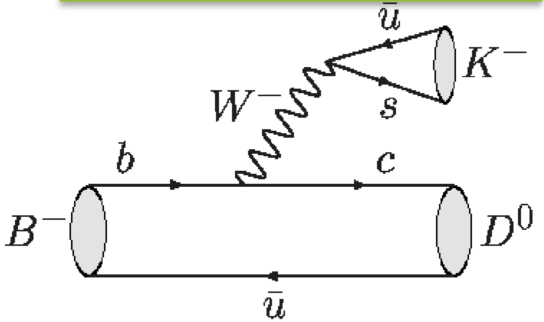

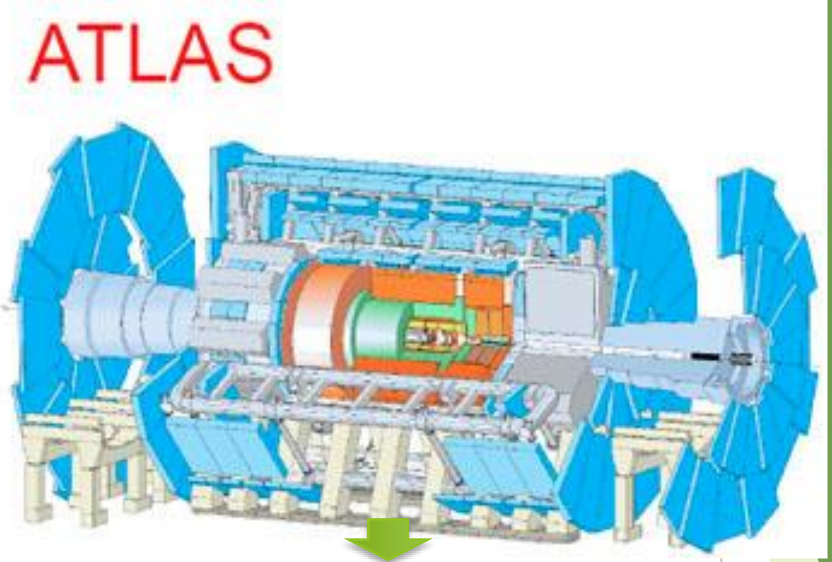

Detector electronics

Trigger (selection) SW

Reconstruction SW

Analysis SW

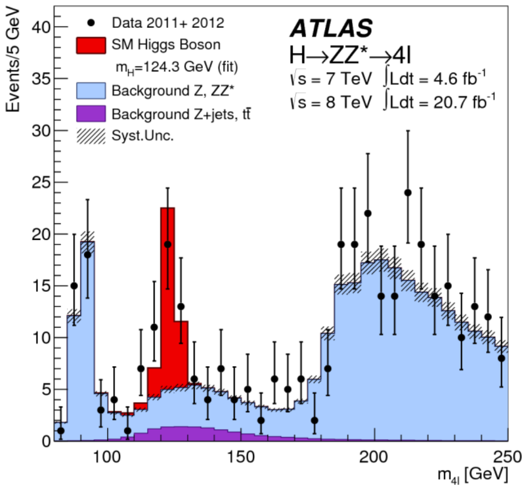

Simulation - all SW

Theoretical model → Simulation of decays of unstable particles → Simulation of interactions particle-detector → Simulation of detector electronics → Trigger Simulation → Reconstruction → Analysis

$$L_{QCD} = \sum_q \overline{\psi}_q \left(i\gamma_\mu D^\mu - m_q \right) \psi_q - \frac{1}{2} Tr\left[\overline{G}_{\mu\nu} \overline{G}^{\mu\nu} \right]$$

PYTHIA

GEANT4
A SIMULATION TOOLKIT

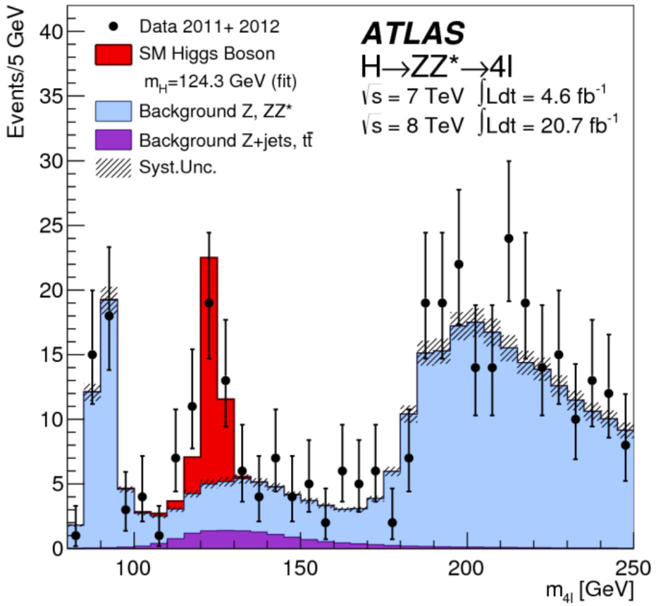

Current (LHC RunII) workflows and computing needs

- Data:
  - Process the 1 GB/s and go from RAW signal to physics objects (Jets, Tracks, Leptons, …)
    - 20-30 sec/ev (1 y ~ 10B events)
    - Considering LHC duty cycle, needs for 20000-30000 computing cores
  - Extract time-dependent calibrations, and reprocess the datasets with them (1-2 times a year…)
  - Perform analysis on the resulting datasets (centrally / at user level)

- Simulation
  - Simulate ~1 simulated event per collected data event
  - "Generate" means:
    - Simulate high energy collisions, and decay the particles
    - Simulate additional interactions due to the LHC bunch structure
    - Simulate selection decisions
    - Simulate the interactions of these with a model of the detector
    - Simulate the front-end electronics
    - Go from the simulated raw signal to to physics objects (Jets, Tracks, Leptons, …)
  - ~ **50-1000 sec/ev**, depending on the model accuracy, precision choices, …
  - Up to **some 100000s** computing cores
  - Perform analysis on the resulting datasets (centrally / at user level)

Resource needs as in 2018

- Resources accessed via WLCG Distributed Computing
  - Initially MONARC hierarchical model, now more “cloudy”

- From WLCG official figures @ 2018 (ATLAS+CMS):
  - ~500k computing cores
  - ~350 PB of disk
  - ~550 PB of tape

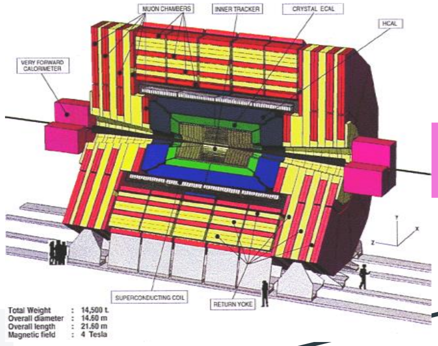

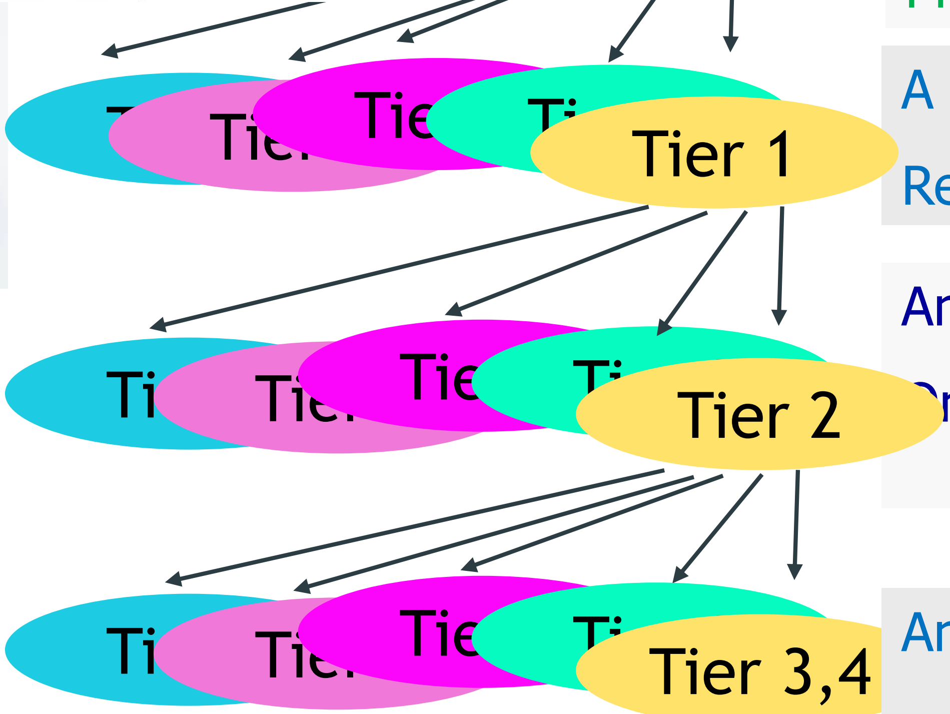

MONARC

CERN

Master copy of RAW data

Fast calibrations

Prompt Reconstruction

A second copy of RAW data (Backup)

Re-reconstructions with better calibrations

Analysis Activity

On average dimensioned to help ~ 50 physicists in their analysis activities

Anything smaller, from University clusters to your laptop

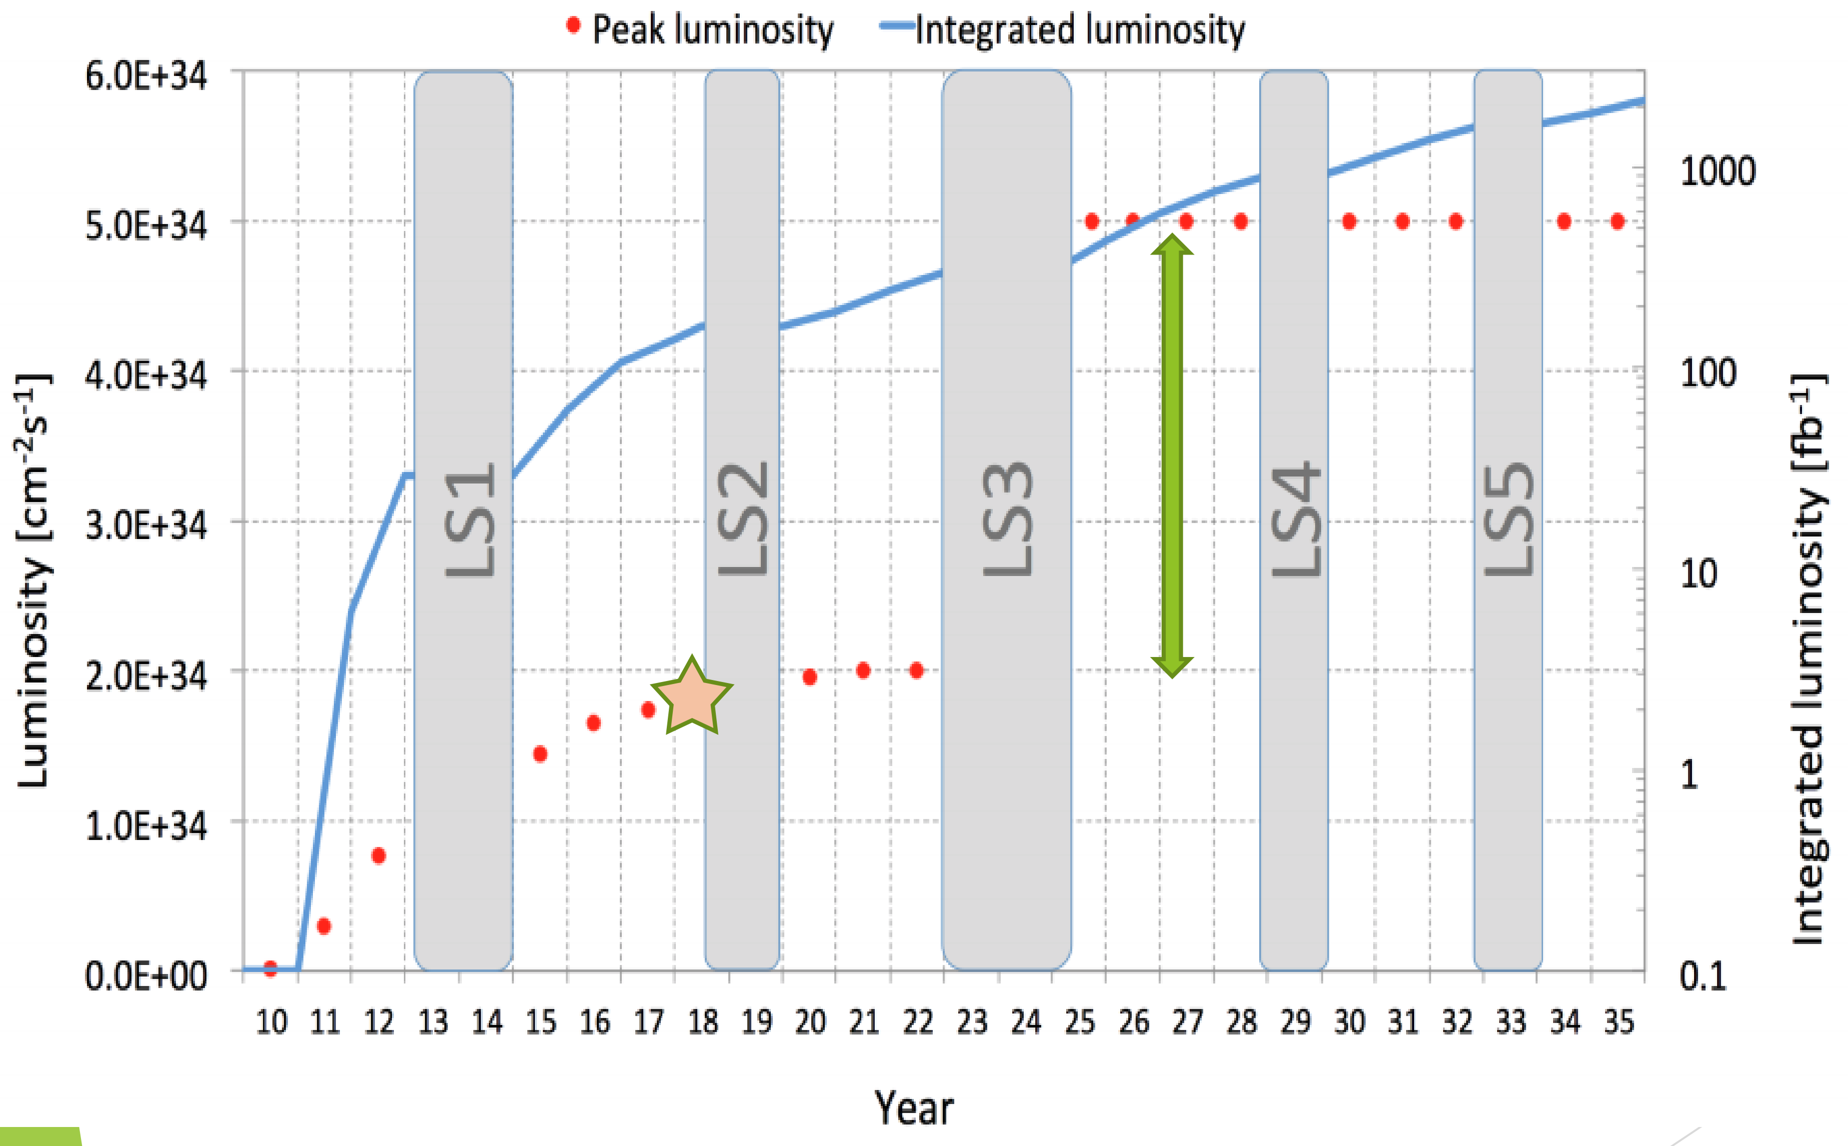

Slightly out of date, but assume red points level are ~ to the event complexity

“we collected ~ 5% of LHC events”

Resource evolution up to 2023

- ATLAS and CMS do not plan any overhaul of their computing model / needs before 2026
- RunIII (2021-2023) needs will be accommodated via adiabatic changes, and resource needs are compatible with a moderate technology trend ("Moore's law" at +15-20%/y)
- Currently no need for breakthroughs
  - Even GPUs, FPGAs, are not on the critical path, much less QC

(this is the major difference with ALICE and LHCb, which will experience problems also on this time scale)

… **but then "we have a problem" →**

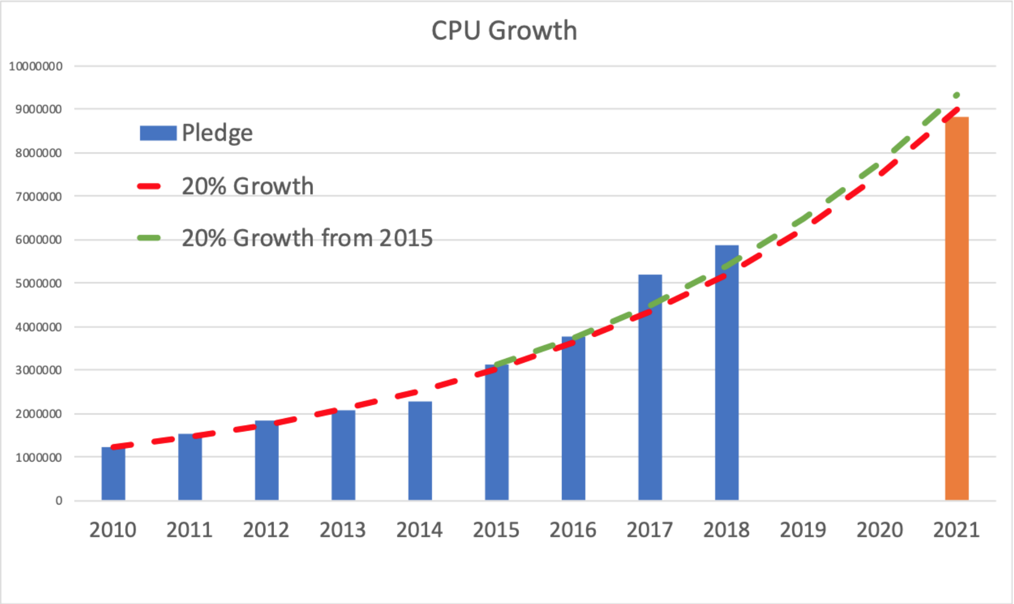

RunIV aka PhaseII aka HL-LHC (2026+)

- LHC drastic change in parameters
  - “event complexity” (pile-up) up a factor ~6x (35→200)
- Experiments in a more challenging environment
  - Upgraded detectors in the more complex environment (~5x bigger events)
    - And CPU time scales more than linearly with complexity – see later
  - Physics demands more events selected (up to 10x)

- Overall: out-of-the-box estimate **~50-100x wrt to previous data taking periods**
- Wait-and-see approach (trust the technology to close the gap) would need a factor 50 Moore’s like improvement in 8y
  - **Optimistically, we can get a 6x (not even clear recently!)**

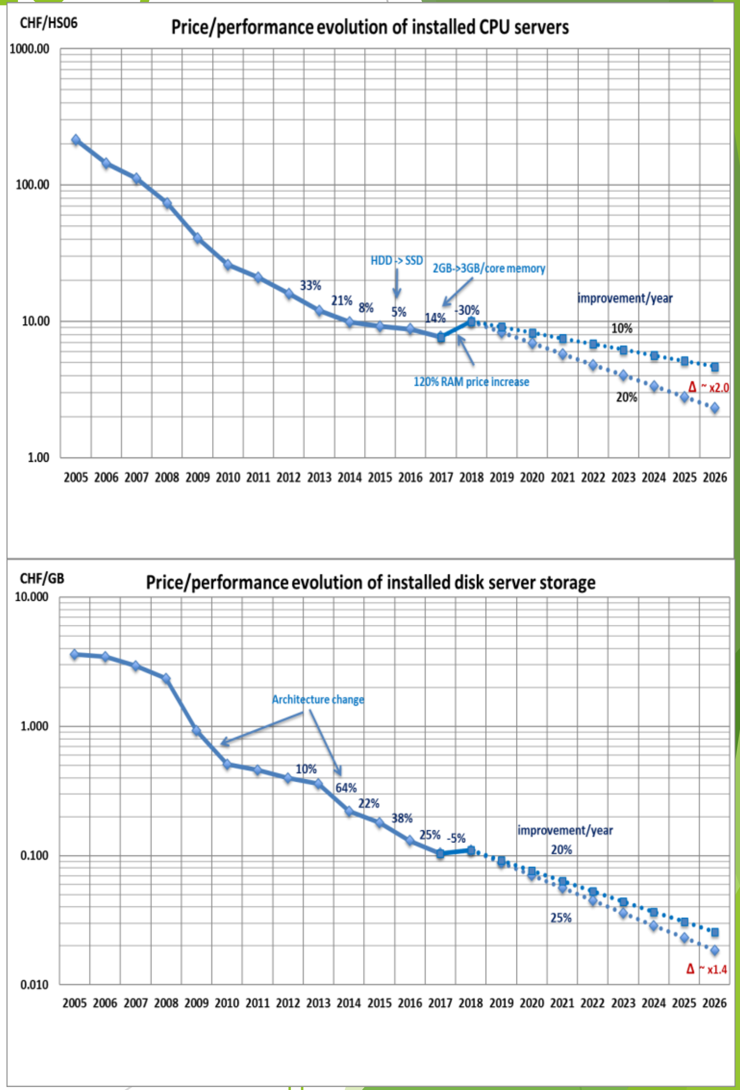

(B. Panzer, CERN)

Last known estimates for 2027 (with already a lot of cuts)

- CPU:
  - If we stay with plain old CPUs (think of Intel Xeons), and assume more and more computing cores with roughly today's speed
    - ~15M Cores needed per experiment
- Disk:
    - ~3 EB per experiment
- Tape:
    - ~10 EB per experiment

- There are differences between the 2 experiments estimates, but mostly due different R&D paths.

- **Take home message for this venue: we are OFF by ~5x on CPU power when considering Moore's law**

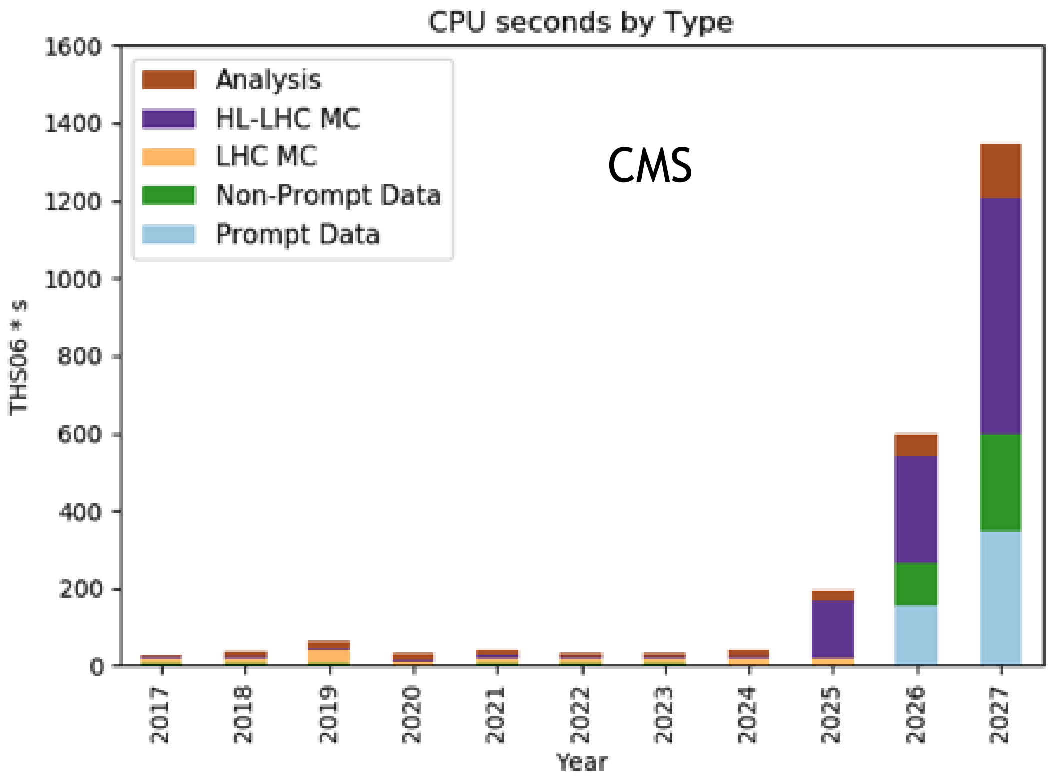

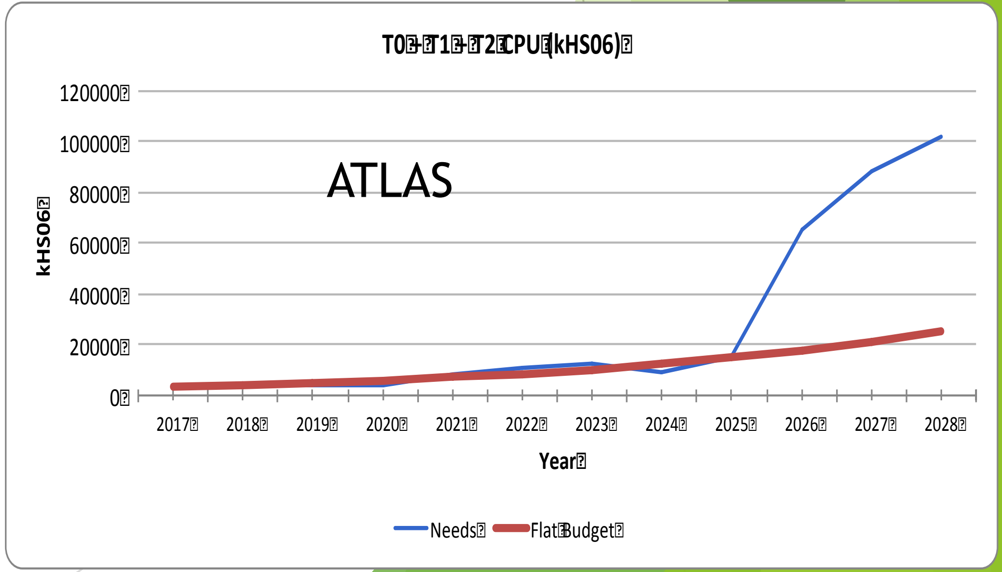

So, how to gain back the 5x?

1. Very easy solution: decrease some parameters, like selection rate. **If 5x less data collected → problem solved**
   - With a large price on Physics
   - It is like buying a Ferrari and using only first gear not to pay gasoline. Not too smart
2. Try approaches which **preserve** physics
   1. Be smarter (fewer reprocessings, less simulation, smaller data formats, tune a good fast simulation, …)
   2. Use cheaper technologies (than CPUs) – GPUs, FPGAs etc seem to offer more “event throughput per $”
      1. … But they need a rewrite of the code base (~10M lines of code per exp)
      2. … But they need programming skills not present in today’s (average) physicist community
      3. Killer application seems to be DL (training?)
3. Anything more at the edge of technology?
   - **We are here to learn if QC could be a solution to some specific parts of the problem**

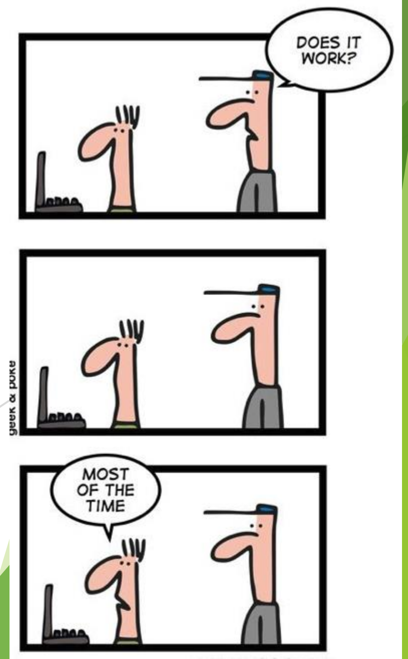

What is the bulk of our processing in 2018?

- **RunII (2018):**
  - Generators range between 1% and 10% of the total CPU needs;
    - Difference depends on the perturbative level (LO, NLO, NNLO), different choices on the market, …
  - Geant4 is currently the most demanding application
    - CMS: ½ of the CPU time for a simulated event
    - ATLAS: >50%
  - Physics Object Reconstruction is 30-40% of the CPU budget
  - Analysis depends critically on the experiment decisions
    - Some 10-30% of the overall budget

- But scaling with event complexity (so to 2026) is largely different

Generators and Geant4 do not scale with complexity; the total time scales with the # of events to be processed

Reconstruction scales with the # of events processed, and scales more than linearly with the event complexity

Analysis scales with the # of events, and mildly with their complexity

So, by 2026….

- We can expect **Reconstruction** (on Data and Monte Carlo events) to be the dominant user of CPU cycles; **Geant4 simulation** following
- **Generation** will scale from today’s values only if we start to need more precise simulations
  - LO → NLO → NNLO → … ?
  - V+ (1,2,3,4,5… N) Jets
  - The negative weights problem? A huge increase in resources if they are not solved
- We can expect the need to have sizeable Fast (“simplified”/”parametrized”/DL) Simulation; but this could clash with the need of more precise measurements

What are the typical algorithms doing?

- **Generation** is the simulation of a single particle collision, hence it has some modelling of a quantum system (be it via explicit matrix element calculation, or sequential steps, …)
  - Currently, done via approximations (perturbative orders, resummations, more and more loops and legs, …)
- **Simulation** in **Geant4** is mostly a transport problem, in which subsequent interactions particle/matter take place
  - Some of them only drive to some energy loss, some others to decays / hard processes, …
  - The more particles / volumes, the more time
- **Reconstruction** is an algorithmic problem, in general most of the time is spent in combinatorial algorithms (nested *for* loops)
  - Searching for doublets, triplets, quadruplets not atypical (N^2, N^3, N^4 …)
- **Analysis** is … anything!
  - In general, there is a selection step followed by a minimization (likelihood, …) step

- How can we solve the resource problems?

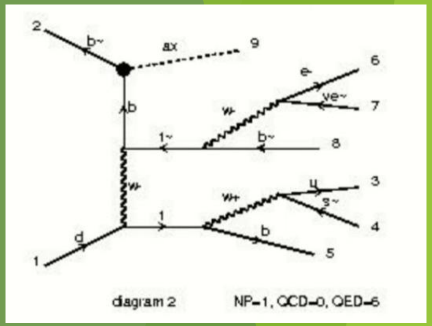

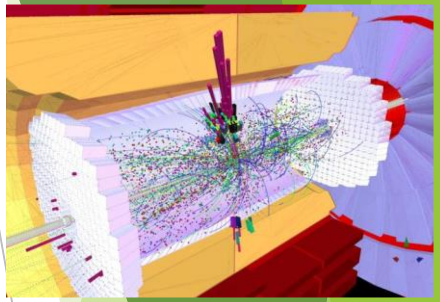

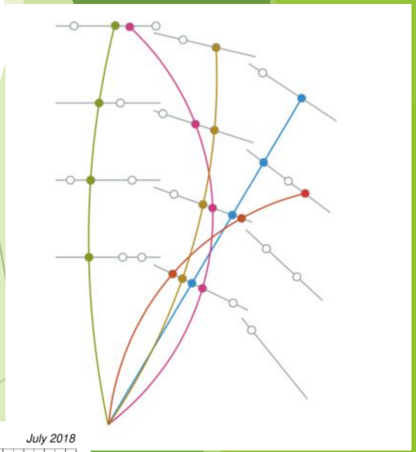

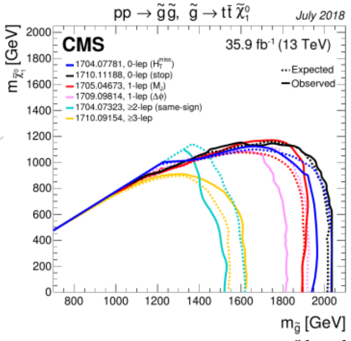

Some approaches possible

- Technical :
  - Buy computers with better performance / price ratios
    - GPUs, FPGAs, ASICs
  - Write more performant software for standard CPU
    - Vectorization, AVX, …
- Operations
  - Use smaller data samples
  - Avoid reprocessings by “doing it right” the first time
  - …
- Political
  - Get more money and buy more computers
  - Access more resources
    - HPC, Volunteer
- Physics
  - Take less data
  - Cancel some parts of the physics program (low pT physics, B physics,…)
  - Delay processing during Long shutdown (== process 2022 events after 2024)

This is the current masterplan

- Try preferentially to explore solution not impacting physics and not requiring more money
- Use the 8 years from now to 2026 to
  - Be prepared to use heterogeneous computing architectures, allowing to
    - Use the best performance/price ratio at any moment, following market
    - Enlarge the basis of potential resources (more HPC centers, more farms, more clusters, …)
  - Better understand analysis models, and reduce the needs for MC, processings, calibration steps, …
- Is this enough?
  - Who knows for sure …
  - In the communities, you can feel a mild optimism though …

Is QC another “weapon” we should study?

- **Disclaimer: we are here mostly in order to LEARN; our understanding of QC possibilities is not necessarily adequate**

- Bird’s eye understanding:
  - **Quantum simulation** could in principle take the place of algorithmic generators, at least for some specific processes
  - **Quantum computing** could be used in principle for generic minimizations, or in order to speed up combinatorial algorithms
    - Or in principle ANY algorithm via a Grover approach

The data problem

- From what we described up to now, it is clear that we need to access / crunch / move large data amount during our processing; while typical QC examples are the factorization of a prime number (~ **0 bandwidth**)

- **Example**: expect a 5 MB event to be processed in Reconstruction (in 2026), in some 50 seconds → if bandwidth less than 100 kB/s, that is the limitation

- **Which is the bandwidth we can expect from a quantum computer?**
  - Even if processing is fast (say quantum tracking), what if it takes 10 min to create the initial state?

- **Can QC be imagined applicable to algorithms which operate on real data/simulation at all?**
  - For these, the above bandwidth requirements raise by orders of magnitude (same 5 MB event to be processed in 500 msec)

Finding minima …

- Finding minima (local / global, more or less approximate, …) is today **not** the largest part of our workflows
  - Yes, we minimize likelihoods at analysis level, we fit functions, … but most of our code is algorithmic
- Still, in case a usable universal QC minimizer would appear, we can *think* of many applications
  - ML/DL can be seen as a minimization
  - Some Tracking algorithms do minimize (DAF)
  - …
- **Is this the simplest utilization which could have a real impact?**
  - And have a low learning curve for us?

Overall

- A reasonable approach could be
  - We honestly **do not think** we can count on QS/QC as a **mission critical tool** for HL-LHC …
  - … but equally, we **cannot be caught unprepared** in the eventuality of a technology / theory breakthrough
- We are sure we can find in our Collaborations interest in following / experimenting / studying QS/QC matters
  - If we are given some initial **guidance**
  - If we are given **access** to emulators / real systems
  - You are seeing today some examples of such activities by single / small groups
  - **Which is the best way to scale activities?**

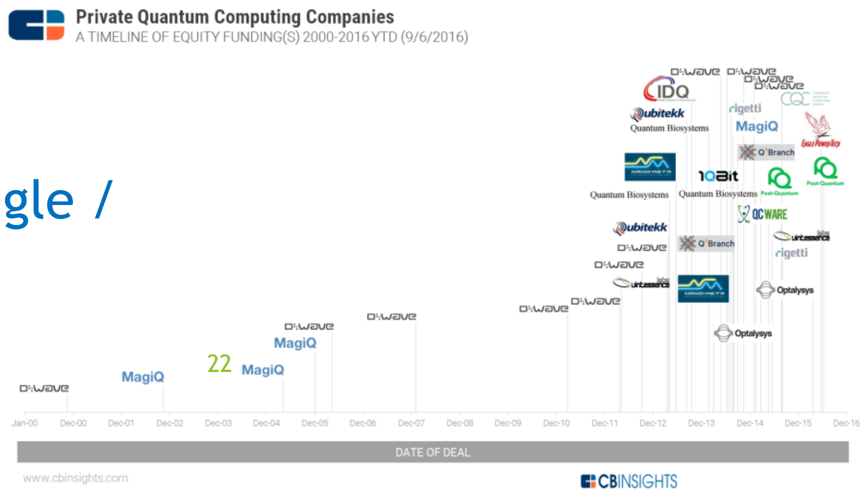

Conclusions

- The Computing Needs for ATLAS and CMS, in a time frame of today + 10y, are exceptionally difficult to meet
- There are directions and R&D programs on “novel” technologies; something even tried (and shown today) on QC
- While we **will not** put QC in the baseline approach for HL-LHC computing, we would appreciate
  - Staying informed
  - Being able to get experience, understanding, ability to run (simple) tests
  - Accessing emulators / real / realistic systems, with a guidance