ATLAS Computing towards HL-LHC

H.O.W. Workshop

18-Mar-2019

James Catmore, Davide Costanzo

- The ATLAS upgrade
- Computing projections for HL-LHC
- The big challenges

HL-LHC computing. Start from an upgraded detector

New trigger system:
up to 10 kHz output rate

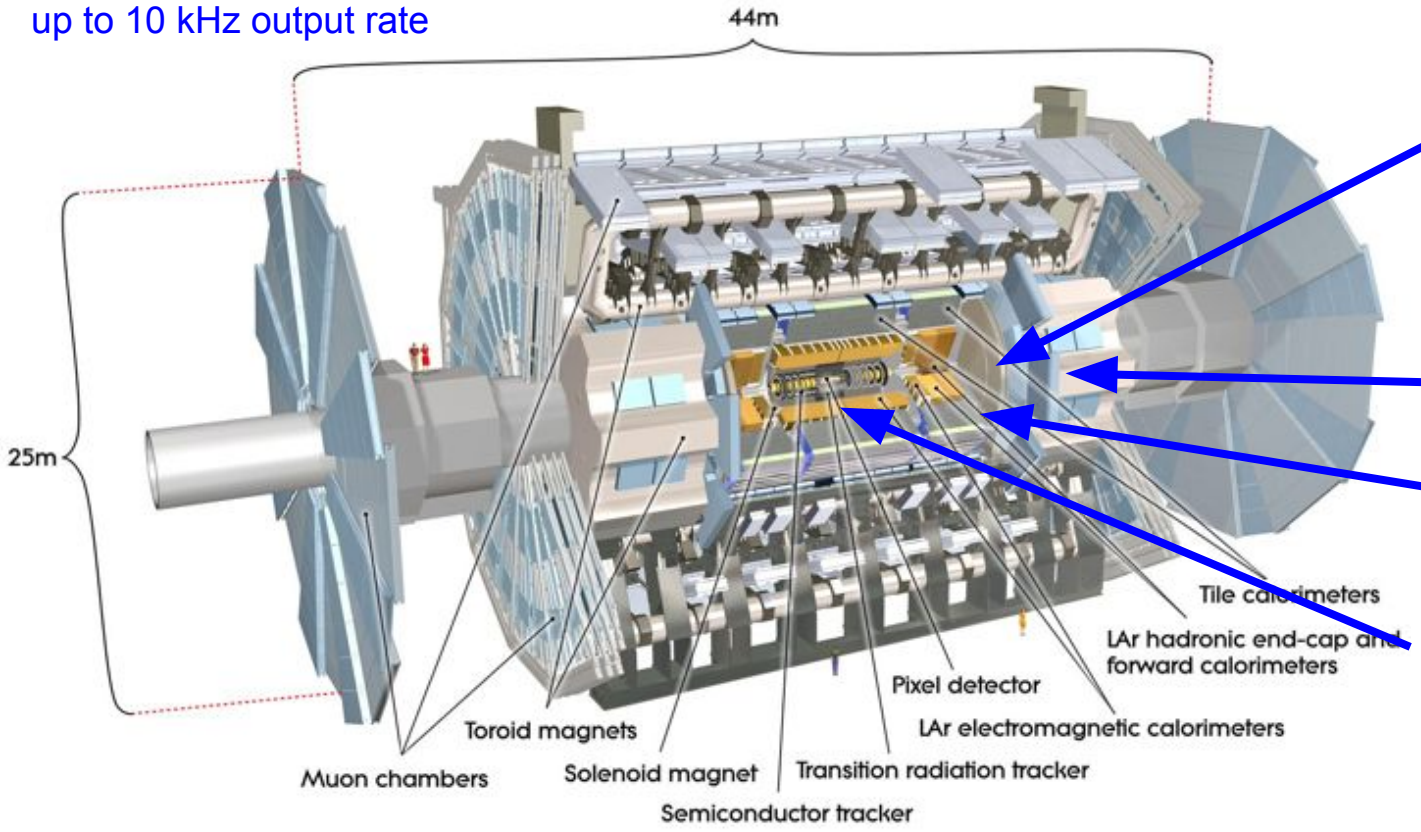

High granularity timing detector

New small wheels

Calorimeter front end electronics

New Inner tracker
(All Silicon technology)

HL-LHC in pictures

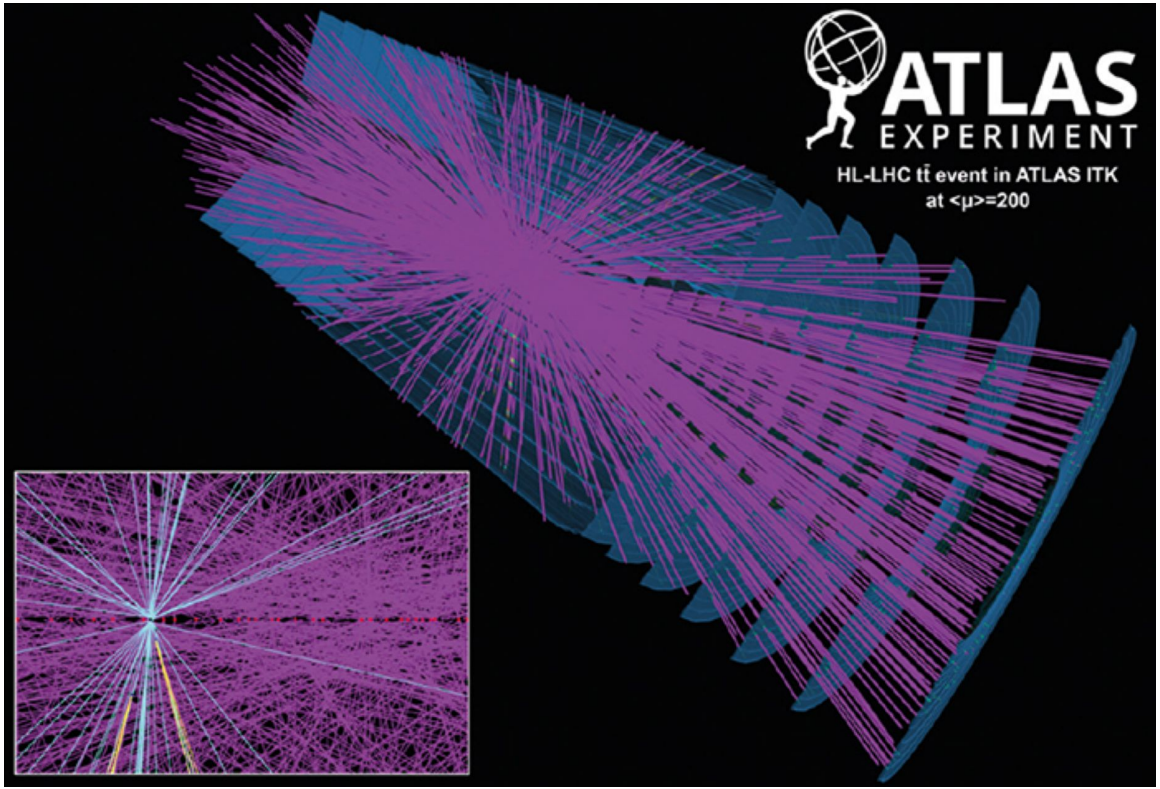

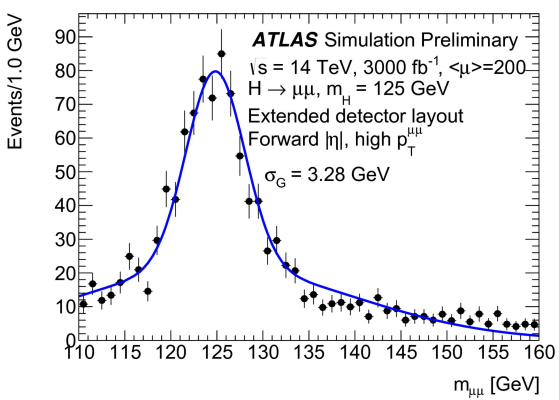

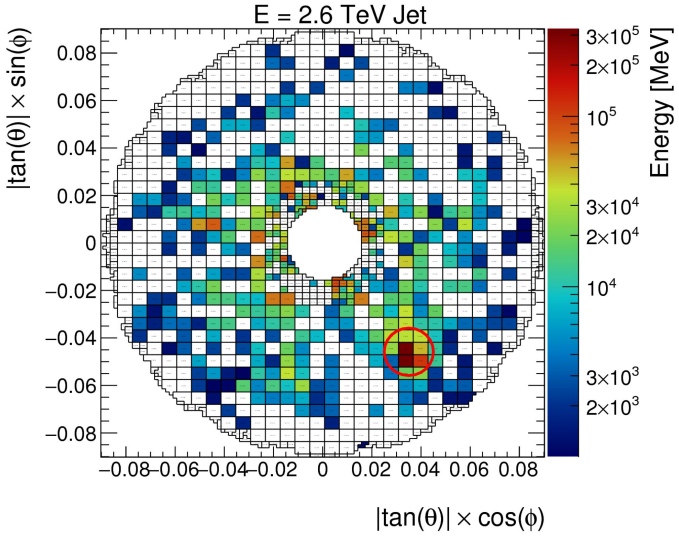

Event processing

- Higher luminosity
  - More interactions per crossing
- Increased event rate
- Bigger and more complex events
  - ITk with >~5B channels
- Desire to do “better” physics

Plot with new optimisation will be shown at CDT

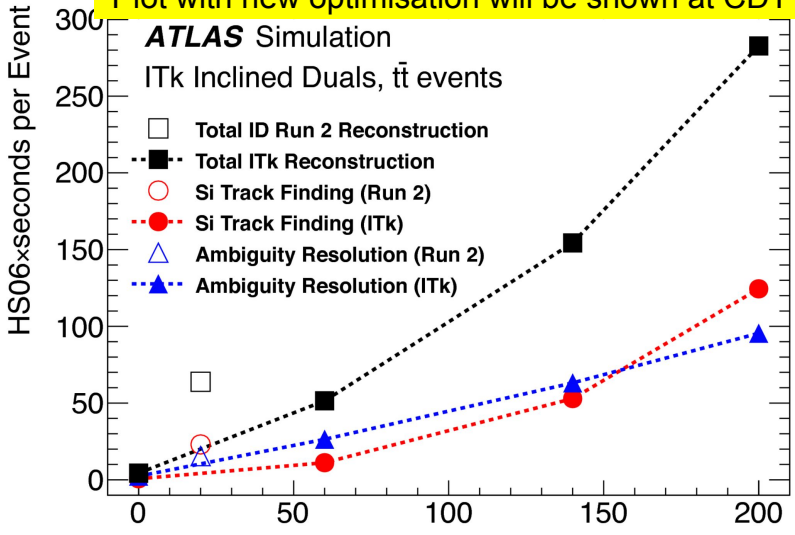

- Better physics performance
  - improved algorithms
- Better CPU efficiency
  - better software engineering
- Reconstruction session on Wed
- Multithreaded running at the event level
  - Exploit parallelism
  - Technology watch session on Tuesday
- Use of diverse hardware architectures
  - Better infrastructure
  - Paolo’s ACAT talk
  - Later talks by Graeme, Charles
  - Session on SW on accelerators on Tue

Challenge: Skills shortage, Training
(Session on Thursday)

CPU projections for HL-LHC

Fast vs Full simulation:

Run 3: 50% of simulation with fast sim

Run 4: 75% of simulation with fast sim

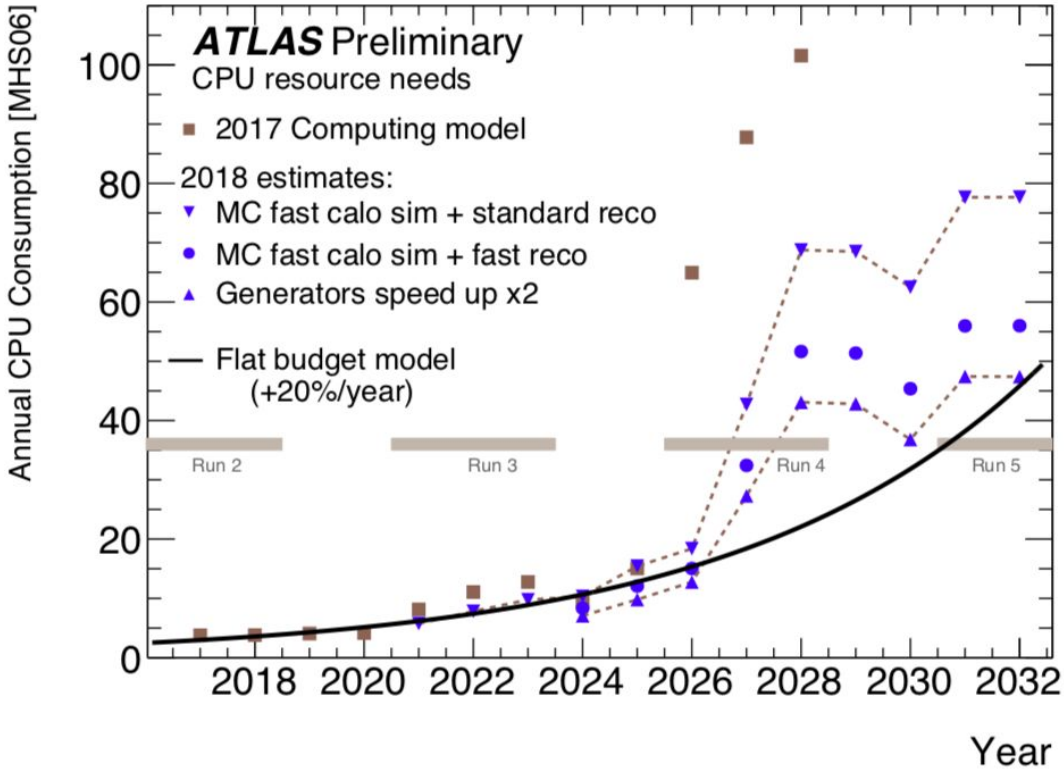

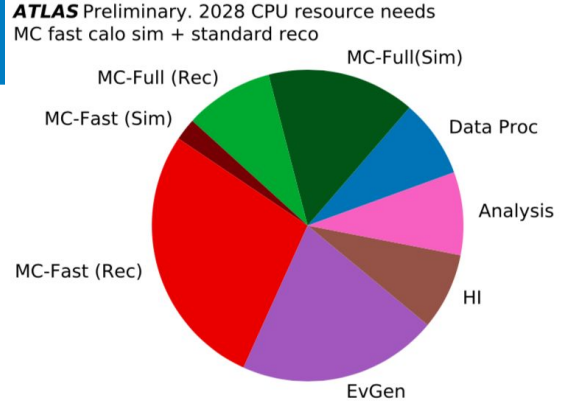

⅔ of global CPU time for simulation

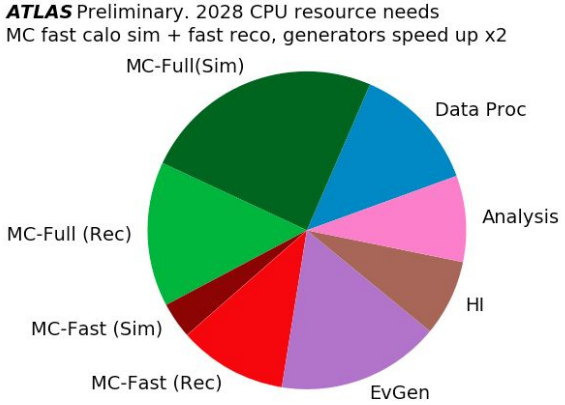

Event generation crucial

(generators workshop last Nov)

Simulation evolution: From full sim to fast chain

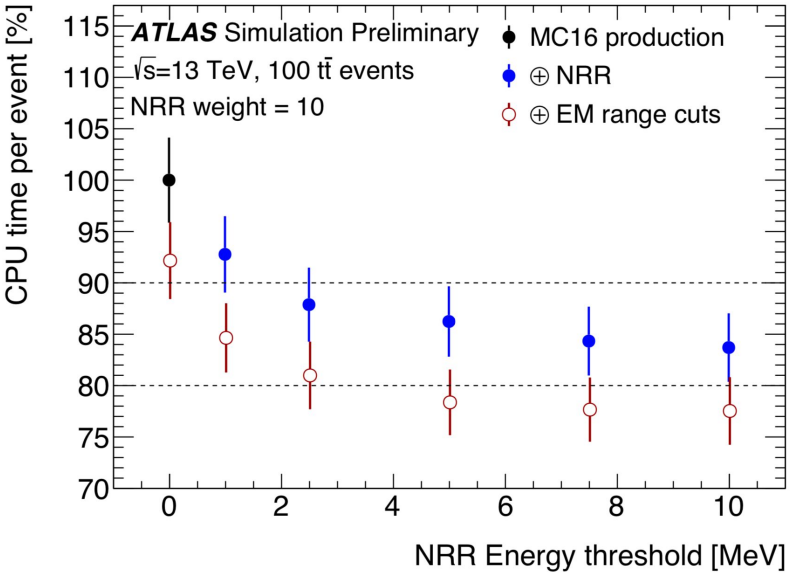

- At least ¼ of CPU to be used for Full simulation:
  - Tuning and improvement of simulation very important
  - Example:Neutron Russian Roulette and EM range
- Fast Chain as a key ingredient
  - Validation as “good for physics” is a major challenge
- New ideas are needed, eg GAN
- Simulation session on Thursday

Challenge: Do better with less CPU → FastChain

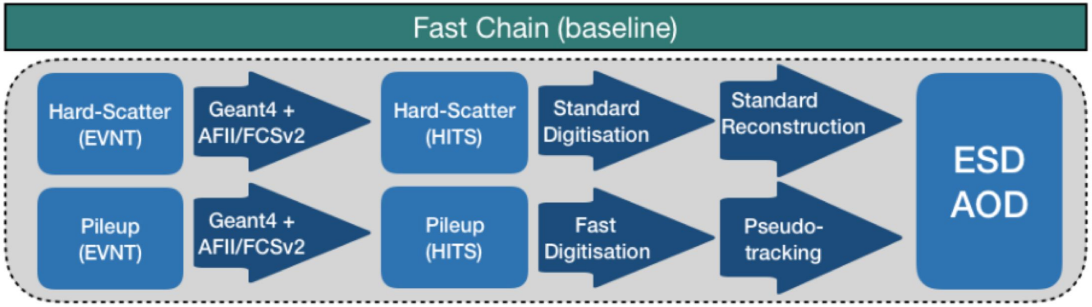

Single photon shower shapes

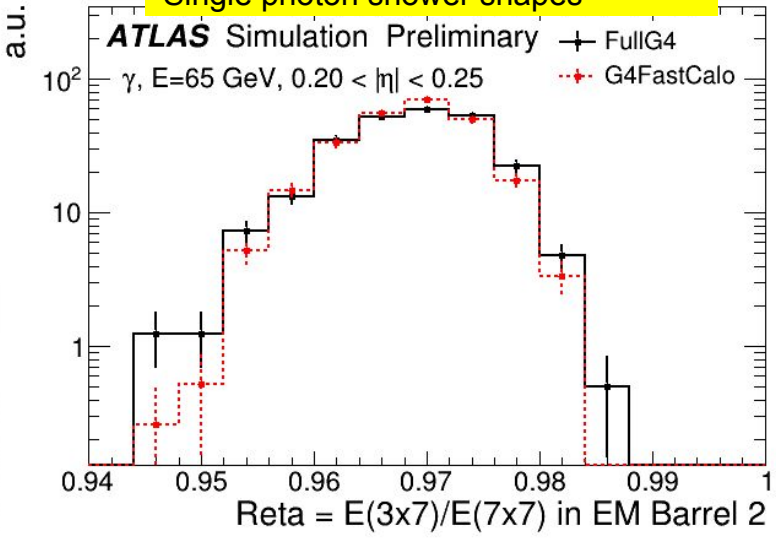

Using accelerators and HPCs

- Supercomputers are evolving away from the “usual” hardware we have on WLCG resources (it’s all in the flops)
  - Eg Summit Power9 + Nvidia V100
  - Other architectures becoming popular - ARM
  - Challenge of portability
- How do we use of accelerators
  - GPU, FPGA, TPU, …
  - GPGPUs demonstrator in ATLAS in 2016
- More in Charles’ talk

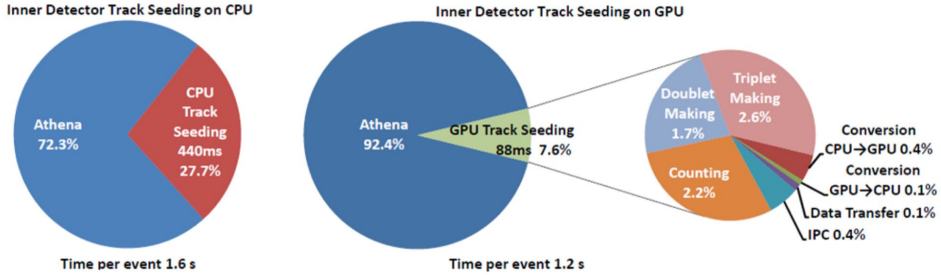

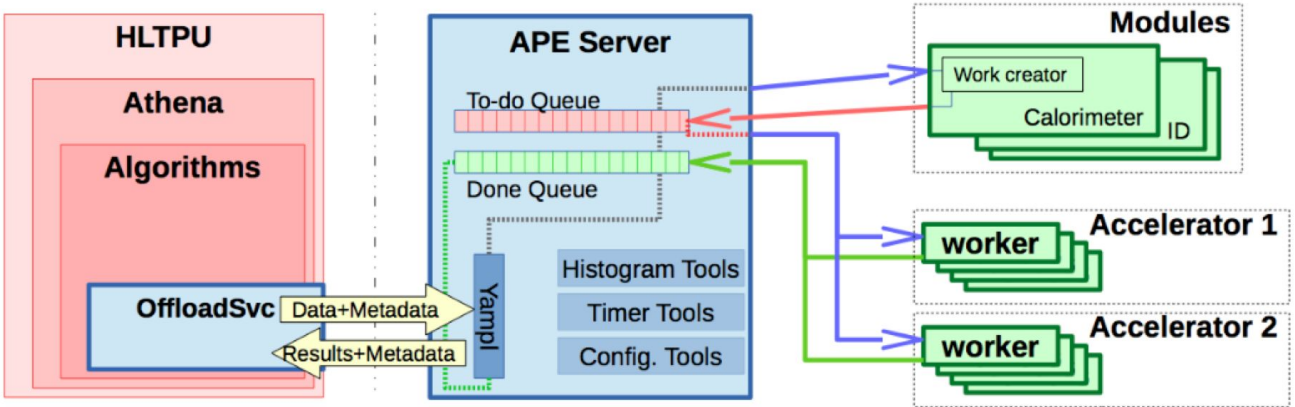

Challenge:
Use accelerators, avoid a lock-in to a specific technology

Disk storage projections for HL-LHC

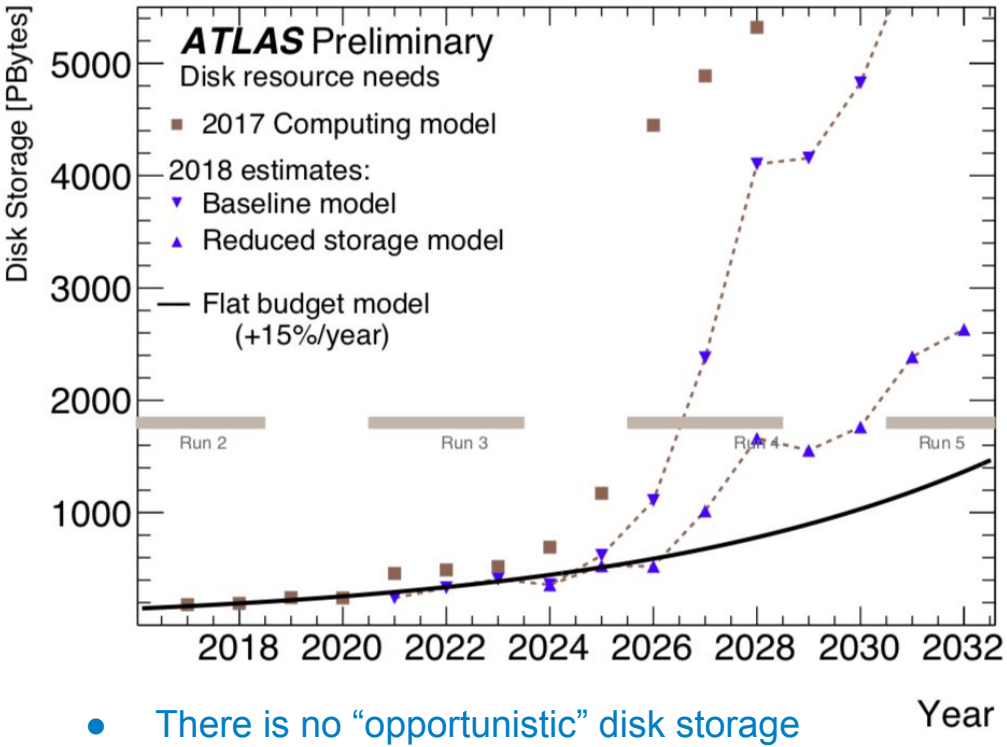

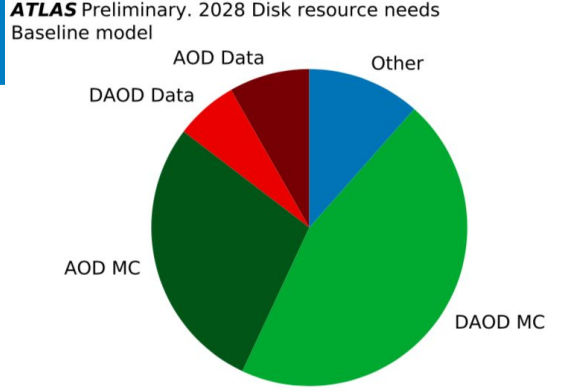

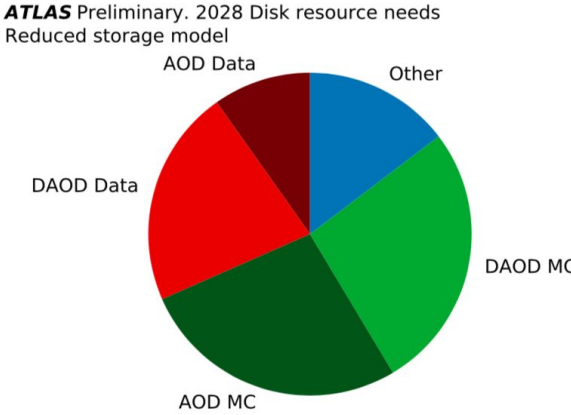

- There is no “opportunistic” disk storage
  - We are on our own…
- Solutions from physics (analysis model)
- Technical solutions
  - Cross-experiment: DOMA -- Session on Wed

Analysis model evolution

- Most disk space is used by analysis formats
  - Access is often time critical due to conferences, paper deadlines, competition with CMS, …
  - Situation currently requires active management from ATLAS
- Ongoing Analysis Model Study Group for Run-3 (AMSG-R3)
  - Further evolution expected for Run-4
  - Analysis session on Wed

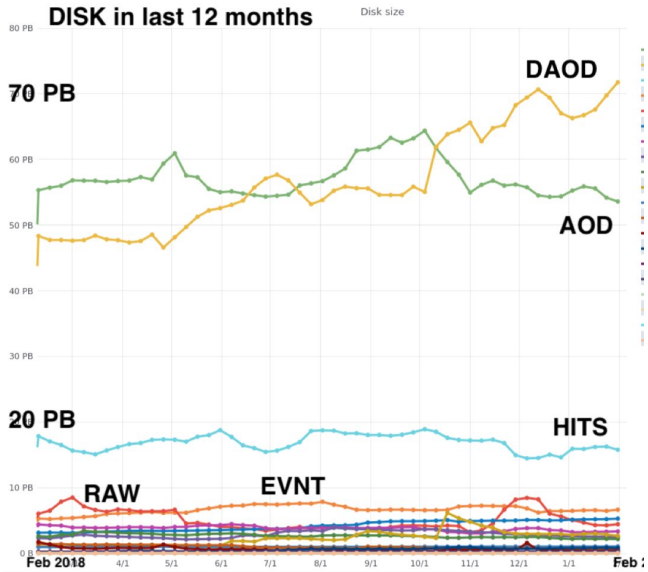

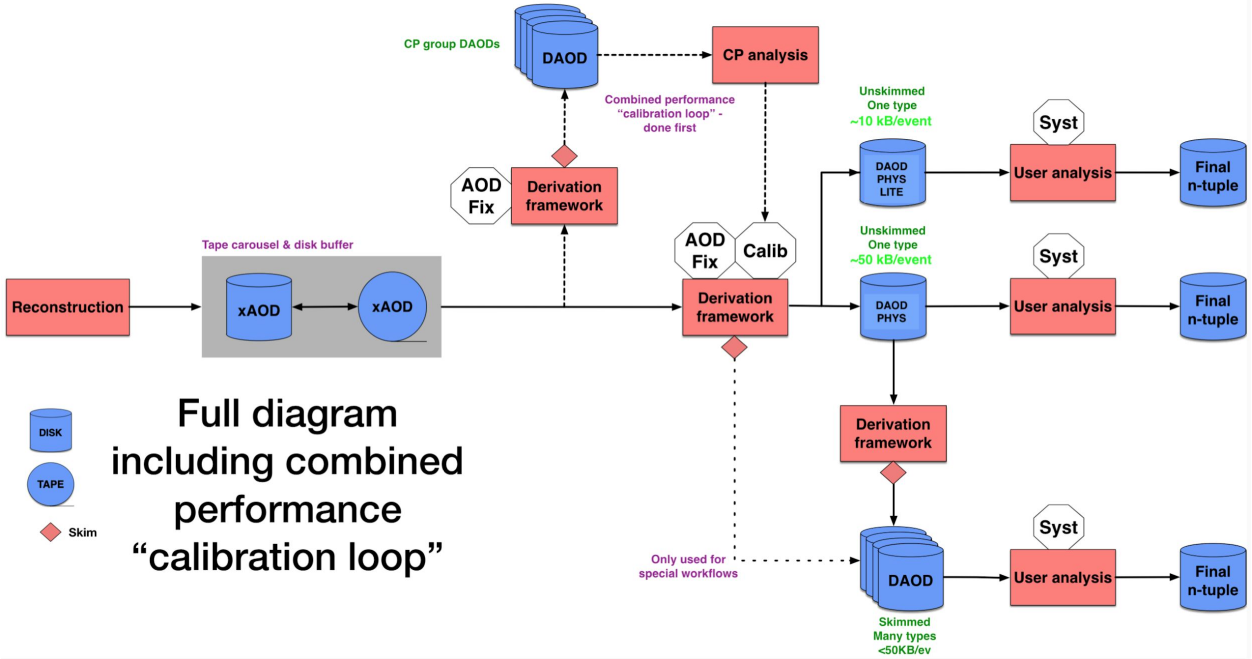

Tape vs Disk at HL-LHC

- I call tape a “cold storage”, cheaper but slow access
- Large needs just for two custodial copies of the raw data
- Keep on disk data needed quickly
  - Analysis formats AOD and DAOD
- Move to tape data that is processed less often with set schedules (data carousel)
  - Currently capable to stage ~0.8 PB/day bulk staging

- Need to evaluate the relative cost of storage vs reprocessing
  - Can we avoid storing intermediate formats?

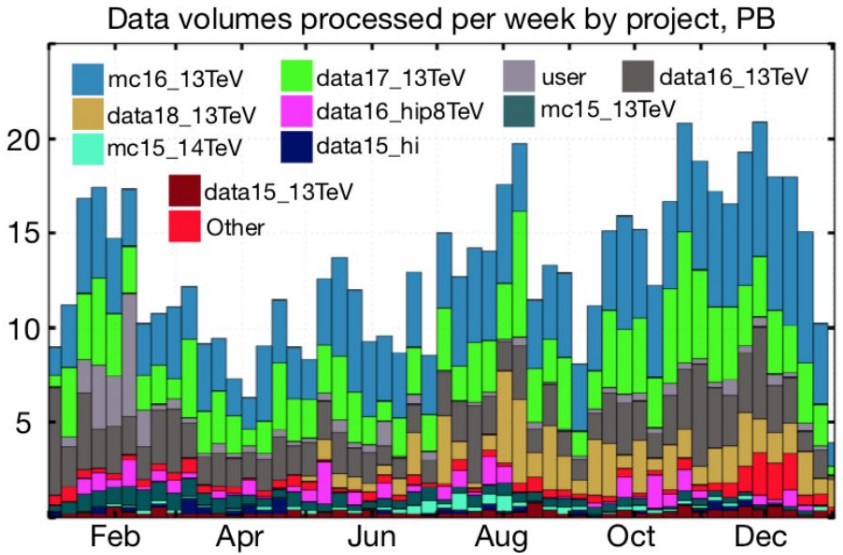

Comments on resources evolution and costs

- Very efficient at using all resources available
- We will certainly use up all the flat budget increases

- HL-LHC requirements are above flat-budget
  - Priority will to store and process all the data
  - Reduced MC → larger uncertainty from MC stats and larger systematic uncertainties
  - Cost model WG on Wed
  - Analysis processing time also an important consideration

CPU usage 2018

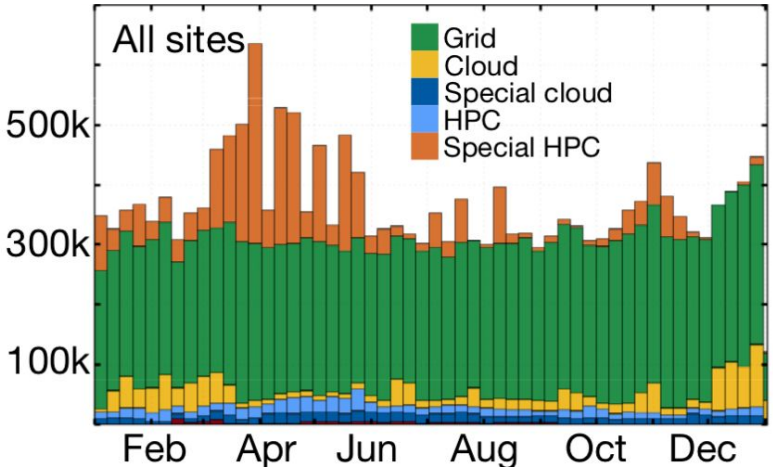

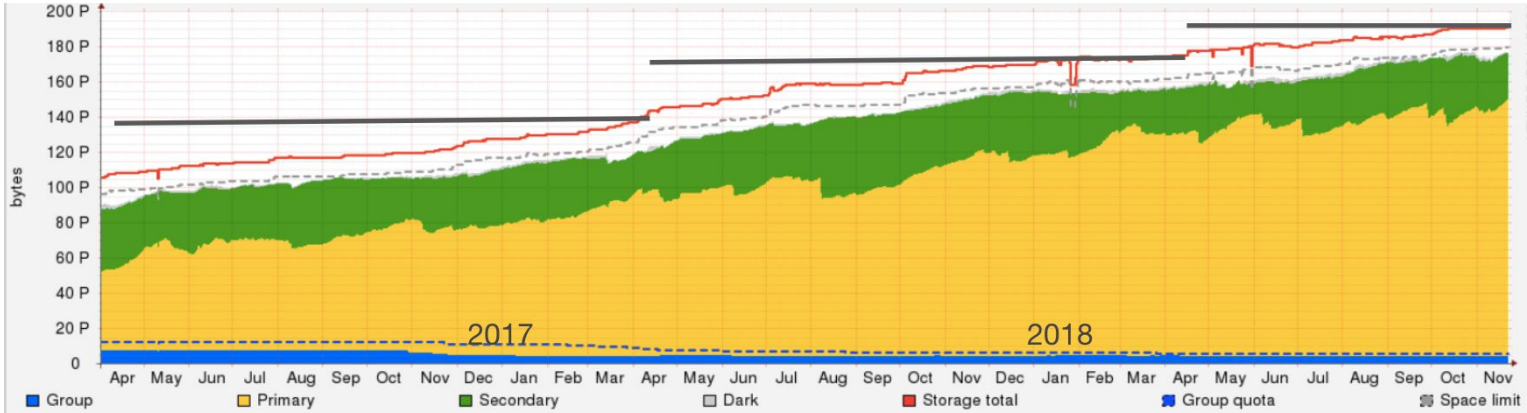

The role of machine learning

- Deep learning is known to be well suited for accelerators
  - Especially GPU, driving new hardware development (more in Charles' talk)
- Artificial intelligence techniques are developed in the HEP area
  - This is moving really fast
  - It could radically speed up simulation and reconstruction at the inference step
- Current usage still at the ML driven reconstruction/analysis
  - Some tasks are particularly well suited (eg flavour tagging)
  - Role on track reconstruction and simulation (big CPU consumers) still to be fully proven
- Training of neural networks becoming a major user of GPU resources
  - Hyperparameter scanning on the grid
  - Training phases will pause once networks are deployed
  - Neural networks should not be considered as the only way to fill Supercomputers
  - Deep learning of complex patterns requires lots of simulation → Fast simulation

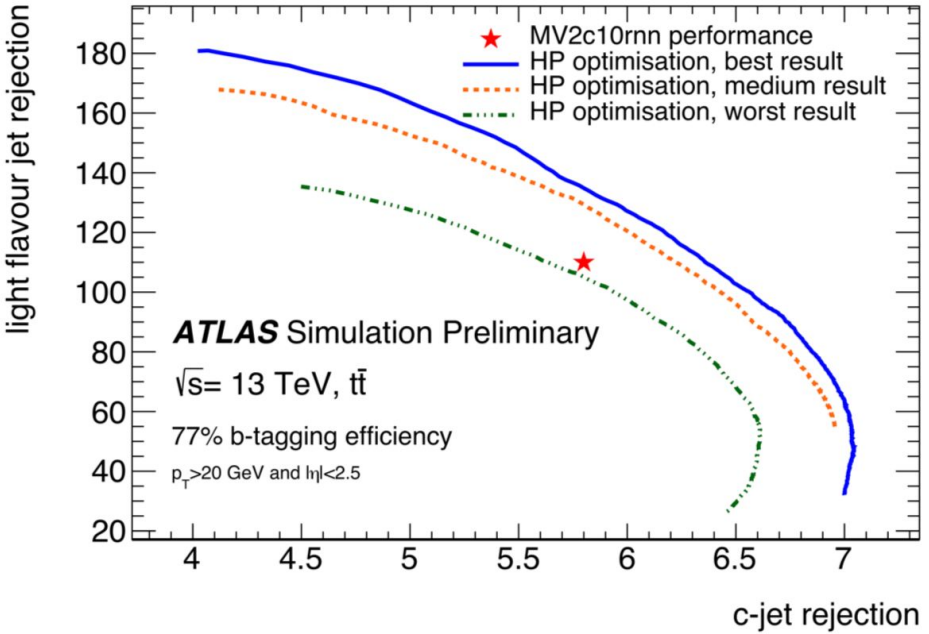

Synergies and collaboration with other collaborations

- Luckily we are not alone
  - The challenges of CMS are very similar to ours! LHCb and ALICE also in the same boat
  - We need to collaborate with other HEP and non-HEP experiments (spirit of this workshops!)
- Data management tools should be shared - DOMA project
  - Rucio adopted by many communities
- Grid WLCG infrastructure is already shared across LHC experiments
  - Good practice, delivered over many years
  - Session on Thurs (and throughout the workshop)

- Some software tools are shared, eg ROOT, Geant4, …
- Can we really share experimental software?
  - Probably no, but we can share experience and ideas!

- Training and careers
  - This is really something we can and should do together as a community
  - Younger colleagues may not be trained to use C++
  - Ensure there is a professional development for people who work on software (and computing) for large physics experiments

Conclusions

- Many challenges ahead of us
  - But also a strong community to support these challenges
  - And people are attracted by interesting challenges
- Training and ensuring a healthy turnover is critical
  - We need young people and fresh ideas
  - This is very valuable as people move on to great new jobs
  - But we would like to retain some of these people in our community!
- It's great to work together
  - We should keep the exchange of ideas and tools flowing
- LHCC review and TDR are the next milestones