Digging computing numbers

A snapshot of the April 2018 ATLAS report to the LHC C-RSG

https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/ComputingModel#Interactions_with_the_Computing

Intro: the 2017 context

Pressure & challenges:

- Stellar LHC performance, tremendous data taking efficiency
- Pile up level of more than double the design value.
- High quality Run 2 analyses across the full physics spectrum -> hence simulation needs
- Support for 6 TDR for Phase-II detector upgrades, while work advances on the Phase-I projects
- Plans & developments for future HL LHC computing requirements

Progress on:

- Workflow efficiency (event server, multiple events loop & merge)
- Usage of the pledged resources (CPU, disk, tapes, network)
- Integration of opportunistic resources, and flexibility in sites allocation to workflows.
- Monitoring of data usage patterns & formats.

Comparisons for the public:

- CPU: 1 core = 10 HS06 ; a laptop = 2 cores → ATLAS pledge = 2 M HS06 = ~ 100 000 laptops
- Storage: a laptop = 16 GB → ATLAS disks pledges = 170 PB = ~ 10**7 times more
- Throughput: fibre @home ~ 500 bits/s, i.e, ~ 100 bytes/s → ATLAS raw data = 2 Gbytes/s

2017 raw data: 10 Petabytes

Overall data rate from SFO: > 2 GB/s (the maximum allowed rate for copy on tapes)
Backlog of ~ 10 days when the recording rate went up to 1.5 KHz

90% physics main stream − processed and delivered for analysis in a few days
The processing of “Non urgent streams” (10%) is delayed until technical stop.

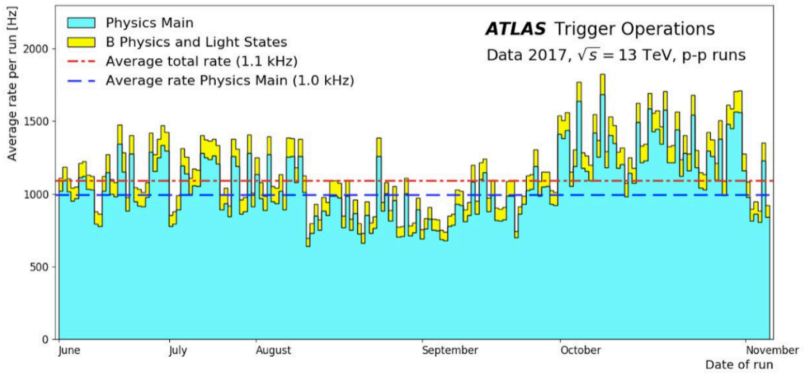

Figure 1: The ATLAS High Level Trigger output rate for physics streams per run in 2017 data taking, as a function of time. The physics events are classified in three streams, with different prompt reconstruction strategies for each stream.

Both collisions and simulated data are further processed several times, in “production campaigns” (calibration, filtering, production of ~100 physics analyses specific formats)

10 billion simulated events [10^{10}] were produced:

- Generators: pp collision physics process
- Simulation: emitted particles propagation through detector
- Reconstruction: same as for real data, to allow comparisons and understanding of experimental effects

Data processed by users analysis: ~ 10 to 15 petabytes per week

2017 ATLAS computing figures:

- Total CPU: ~ 2000 kHS06
- Disk: 170 PB
- Tape: 160 PB

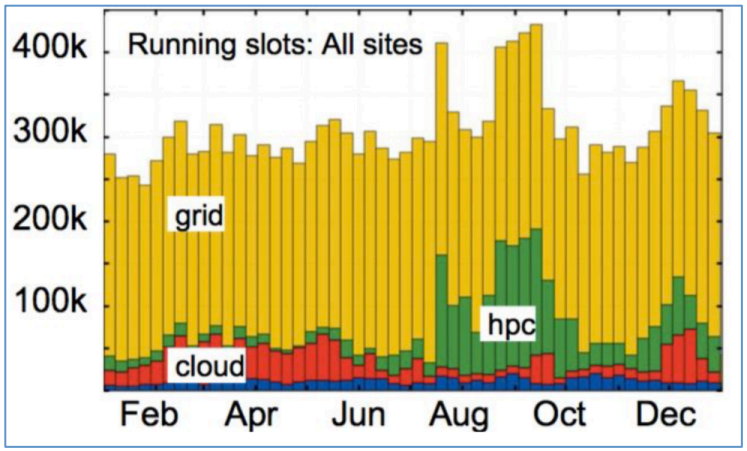

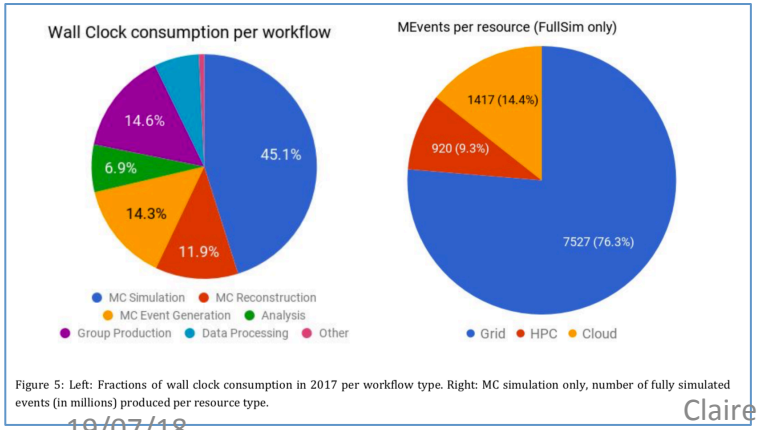

Figure 5: Left: Fractions of wall clock consumption in 2017 per workflow type. Right: MC simulation only, number of fully simulated events (in millions) produced per resource type.

Data processing and Data access : Run 1 model

- **CERN: T0** computing centre's primary task is to save, process and distribute collision data.
- **Grid: T1** national computing centres hold copies of the data, handle reprocessing campaigns and simulation.
- **Grid: T2** university or regional centres are the primary access point for scientists analysing the data. Their CPU is also used for simulation

Run 2 facts:

T2 provide ~50% of the simulated data, twice more than expected, and many are as reliable as T1.
Network / monitoring / deployment easier & mature
The contribution of opportunistic resources (HPCs and cloud) is growing.

→ Moved for Run 2 to a "network oriented" and dynamic model of "Nucleus" and "Satellites", where data and tasks are automatically redistributed to match resources.

~ 30 nucleus sites (including T1) in 2017

Data transfer activities:

- During data taking, the T0 exports ~ 1PB per week
- Rest = simulation & “train” processing ~ 5 PB per week total
  - Transfer input data to processing site
  - consolidate output to final destination
- See in orange at the bottom “user subscription” = chaotic / random access by scientists

Number of events stored on disk:

- 20 billions collision events
- 16 billions simulated events

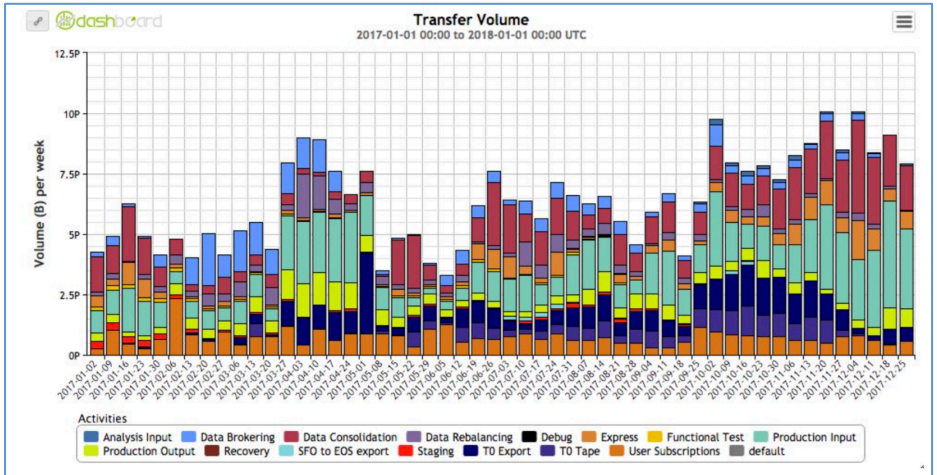

Figure 12: The volume of data in PB transferred by the ATLAS data management system (Rucio) between storage endpoints at Grid sites, in 2017.

Data access by scientists:

- “Derived data” (AOD, xAOD) have to be stored on disk for efficient analysis
- “Raw” (collisions data) and “Hits” (simulated data) are on disk while needed for reprocessing or specific studies

Reconstruction time per event at the T0, as a function of “pile up”

High lumi: expect ~ 200

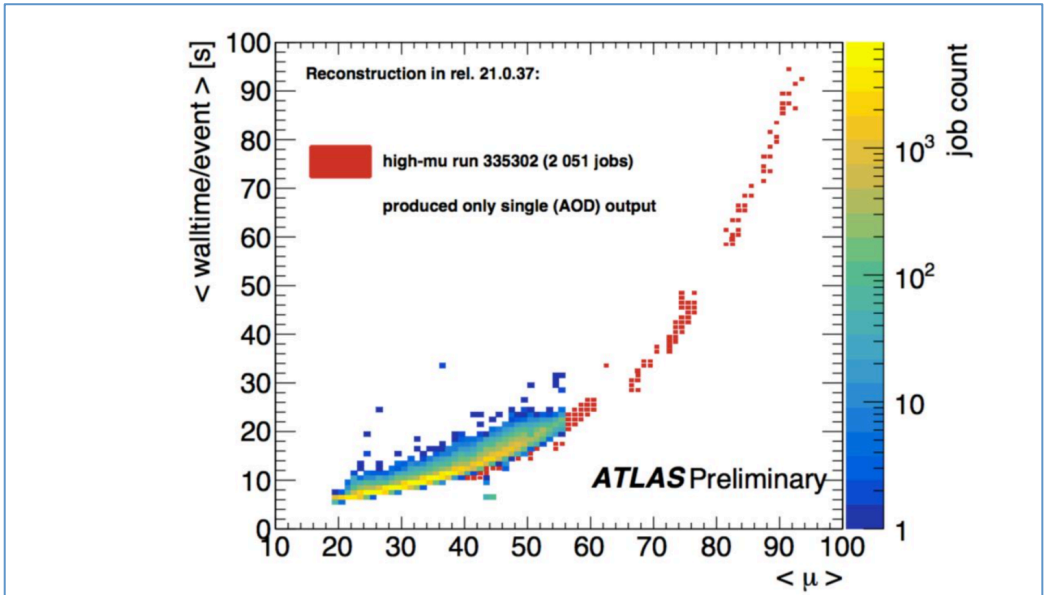

Figure 7: The wall-clock time per event measured at the ATLAS T0 reconstructing the 2017 data as a function of the event pile

Levelling: ATLAS chose a maximum of <mu> = 60 [this is a 2018 plot]

Some comparisons with previous runs