ATLAS in GridPP5 (and after)

Roger Jones

ATLAS Upgrade Timeline

- The major upgrade to ATLAS will come in Phase II (2023-24) and be installed for *High Luminosity LHC* (2025)
  - Also known as Long Shutdown 3, followed by Run4
  - Significant upgrades to the tracker in place now for Run 2
  - In Run 2 we have already the Insertable B-Layer
  - The output rate of the trigger will more than double to 1kHz
  - The pile-up rate will increase through Run 2
  - Even before Run 4, the compute requirement will grow significantly
- Reconstruction event complexity is naively m! (factorial)

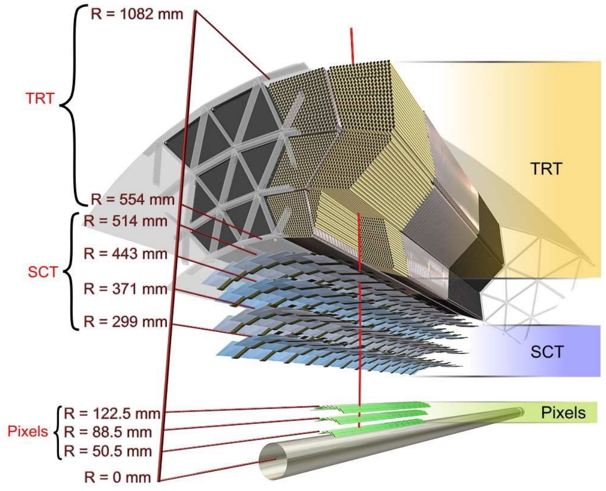

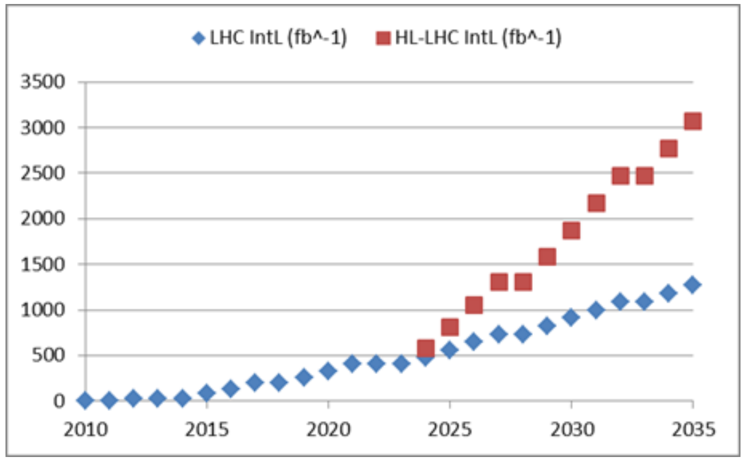

Scale Numbers

	HLT Output	Events per year	RAW per Event	RAW data per year
Run1	600Hz	3.6B	0.7MB	2.5PB
Run2	1kHz	5B	1.0MB	6PB
Run3	1kHz	5B	1.2MB	7.2PB
Run4	5kHz	25B	2.5MB	75PB

How hard will it get?

- Event Generation
  - Not intrinsically harder at high luminosities, however better generators and studying rare processes will mean using more cycles; volume increase to scale with events
  - Good fit for novel architectures and HPC facilities Optimisations possible
  - Relatively easy code to run massively parallel (MPI)
- Simulation
  - Main scaling of simulation per event is with energy (so ~constant in Runs 2, 3, 4, …); however, more data needs more simulation to accompany it, so volume increases
  - Many concurrency opportunities from independence of particles
  - Geant V targets vectoriseable parts of the process — improvements should feed back into G4
  - **More and better fast simulation!**

How hard will it get?

- Reconstruction
  - Definitely *very hard* at high pile up; scaling is naively m! (factorial) for tracking; certainly the biggest challenge faced in software; combines with volume increases
  - Highly *serial* implementation to reject poor track candidates early and minimise wasted cycles
  - Great improvements in Run2 already
    - x4 improvement in overall reconstruction speed, mainly from tracking
    - Greatly improved track seeding strategy has also improved physics quality!
  - But need new ideas for high lumi - machine learning?

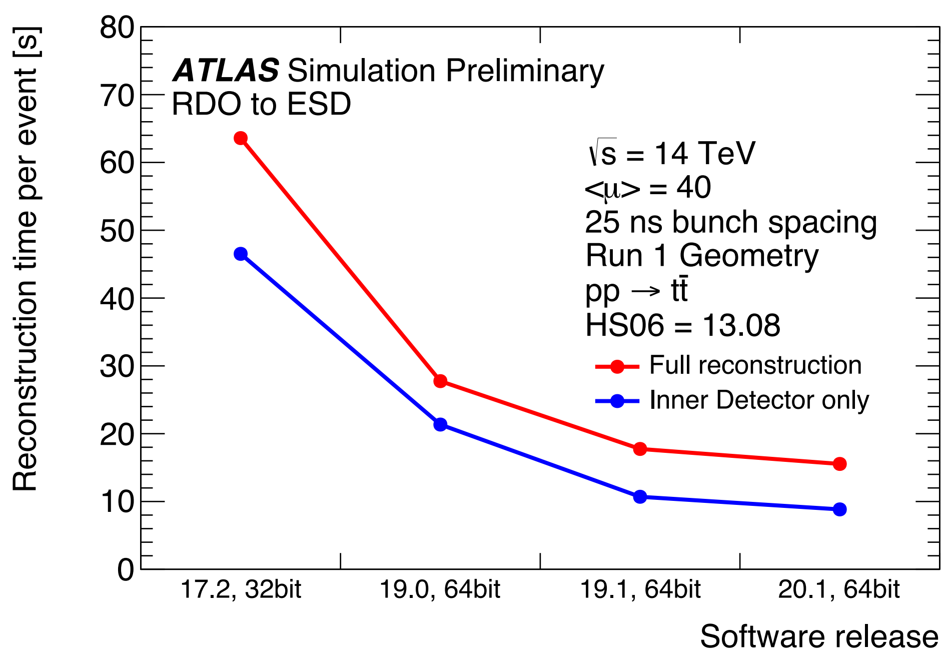

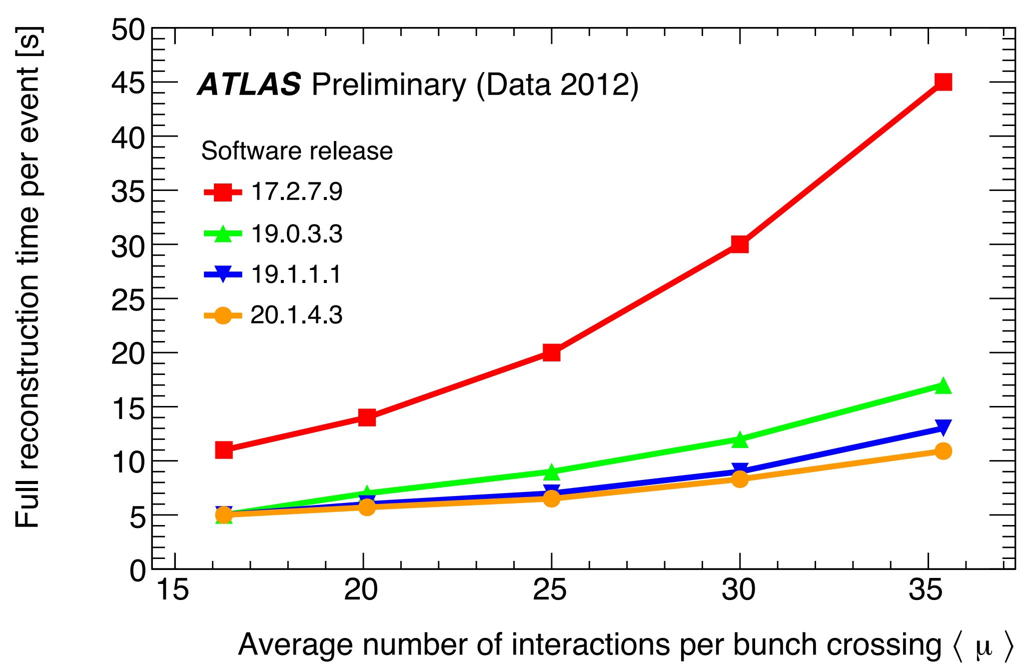

New Framework: GaudiHive

- Memory constraints, especially on non-Xeon server architectures make reducing memory footprint imperative
  - High luminosity and hard tracking conditions only increase this pressure
- Need to move to a multi-threading framework (beyond AthenaMP)
  - Memory savings can be huge as all heap memory is shared
  - However, a more difficult programming model as threads can interfere with each other: data races and deadlocks
- Development to introduce parallelism into the Gaudi framework used by ATLAS and LHCb
- Take advantage of parallelism between algorithms and across multiple events

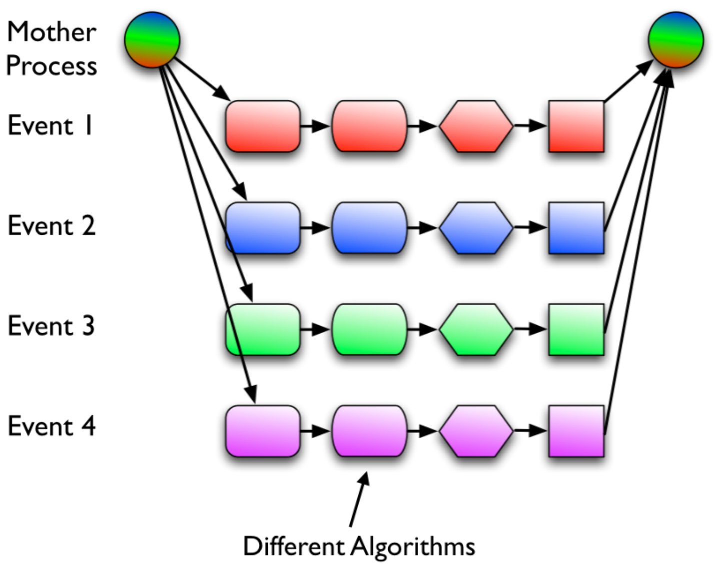

Run2 AthenaMP multi-processing: Each worker uses a separate process, but read-only memory pages are shared

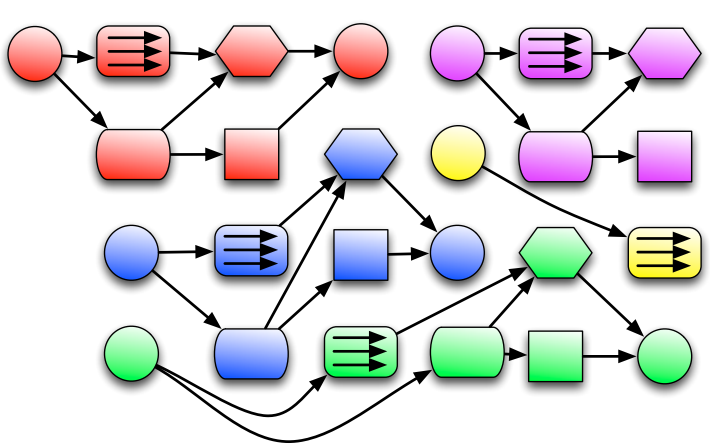

Run3 multi-threaded reconstruction: Colours represent different events, shapes different algorithms; all one process running multiple threads

Analysis

- Most likely linear with data volumes, but analysis can already be i/o bound.
  - i/o becomes a serious problem; need to optimise across huge range of workloads
- Smart slimming and skimming frameworks used to bring data volumes under control
  - At the expense of some data duplication (though also augmentation used)
  - Must keep data volume and cpu costs under control
- Limited i/o capacity pushes us towards ***train models***
  - One job reads an input AOD file, writes multiple analysis output formats
  - Maximise use of staged data
    - Staged can mean staged from **tape**, moved from **mass storage disk** to local **SSD**, data that has undergone persistent to **transient conversion**, data that has been moved from **main memory** into the **CPU cache** hierarchy…
- Internally analysis may use multi-threading to have multiple events in flight
  - Remains to be seen how useful/possible this model is re. programming difficulties (sandbox each event?)
- Object stores as a disruptive technology here…?

ATLAS Data Management

- Today in Rucio
  - 1,000 ATLAS users
  - 140PB
  - 700M files
  - 130 storage sites
- Run4
  - 1,000 ATLAS users
  - 500PB/year with 100 PB of RAW
  - 4B files/year
  - 130 storage sites (+ volatile storage: Cloud ?)

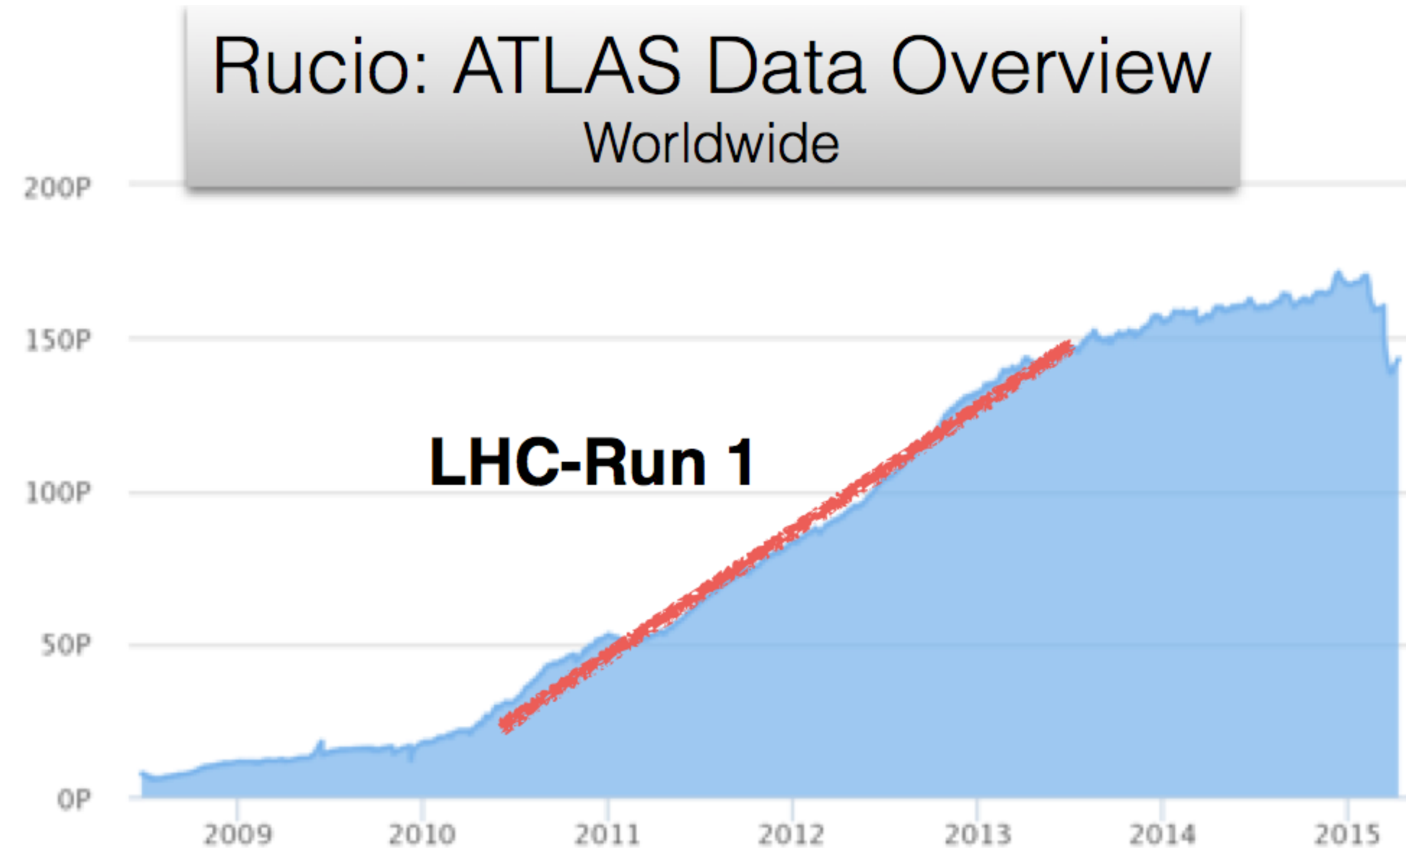

- DDM will have to scale with the (cumulative) number of files, bytes and data operations, e.g., 2.5M transferred files/day in 2015
- Most of DDM implementations for the LHC experiments are based on dataset/file catalogs, e.g., Oracle, and will follow the advances in data bases, middleware and open & standard technologies
  - Flexible design with no dependence on particular middleware
  - Horizontal scalability as a strong requirement

DDM Future

- Integration of new storage types might be needed
  - Object stores: Opportunity to handle data at the object level, e.g., event, and to store (physics) metadata with data
  - Volatile storage resources: cloud storage, HPC to increase the total storage capacity
- The biggest predictable gain will come from network (x100?) and will influence the experiments computing models
- Custodial data management with 2 copies on tape would stay unchanged but not necessary for the derived and secondary data, e.g., intelligent and content delivery (CDN) networks
  - Store vs. cache vs. recompute
  - Intelligent and maximal use of data when requested

Workload Management

- Major restructuring already for Run2: DEFT and JEDI to handle tasks and job splitting
  - Dynamic job definition to fit available resources
    - HPC, Volunteer, Grid, Cloud, …
- Fine grained event service jobs for different workloads, beyond simulation
  - Relies on smart caches for storage and fine grained access to events
    - Retrieve, buffer locally and deliver events to workers
  - Meshes very nicely with new framework design
- Easier installation of PanDA as a local resource manager
- Utilise network knowledge to smartly marry data sources and sinks to CPU resources
  - Initially passively, but eventually actively using smart networks?

Evolving Roles

- Some discussion of different topology for WLCG in Okinawa
- The ‘few big sites’ model suggested by some will not fit the ATLAS requirements

- We need:
  - Fewer sites (to ease the management burden, increase responsiveness)
  - But we need lots of capacity
  - People matter! We need the help from the staff at the many sites

- Sites need to be seen to be doing analysis etc, not just MC, to lever extra resources

An ATLAS Vision

- The ATLAS view of the world:
  - Tier 1s remain important as data sources
    - Reprocessing important, but can be done elsewhere
  - Tier 2s become more important as networking improves
    - More role in reprocessing etc
    - But fewer, larger, well supported sites sites
    - The disk/cpu balance remains large but the disk becomes largely dynamic cache - more PD2P
    - Event service, FAX etc become more important
  - The current UK networking arrangements are seem perfectly adequate (assuming a sensible uplift programme)
  - LHCOne connection as a virtual network if and when required

An ATLAS Vision

- Other resources:
  - HPC sites good for some workflows
    - Generation can be done easily - ~10% of workload
    - Other workflows can be done, but disruptive and inefficient
  - Commercial and other private clouds for bursting
    - Again, particularly cpu-bound tasks
  - BOINC etc
    - Again, good for cpu-bound workflows, non-critical data
    - Not a good general solution
    - Demonstrated as a proof of principle, but *not* yet near production quality

Conclusions

- We know the directions that we should be moving in
  - Multi-threading, data oriented designs, parallel algorithms, architectural flexibility, simplicity
  - Smart use of i/o, from tape through to online and local buffers
    - Dynamic popularity and smart workload management
  - Make use of network capacity in matching data capacity to processing capacity
- Many ideas will be advanced in an evolutionary way for Run3 (2020), starting noe
  - Testbed for Run4 challenges
- Longer term
  - We will have to live within our means (tape, disk, cpu and network) — adaptive process between physics goals and practical affordable computing
- Through GridPP5
  - Mainly need the capacity
  - But we also need the people
  - Fewer, well managed sites