HPCs and ATLAS: vision and strategy

James Catmore (Oslo)
Davide Costanzo (Sheffield)

Two reasons why HPCs are important to ATLAS

Reason 1: we will need them to be able to exploit the HL-LHC

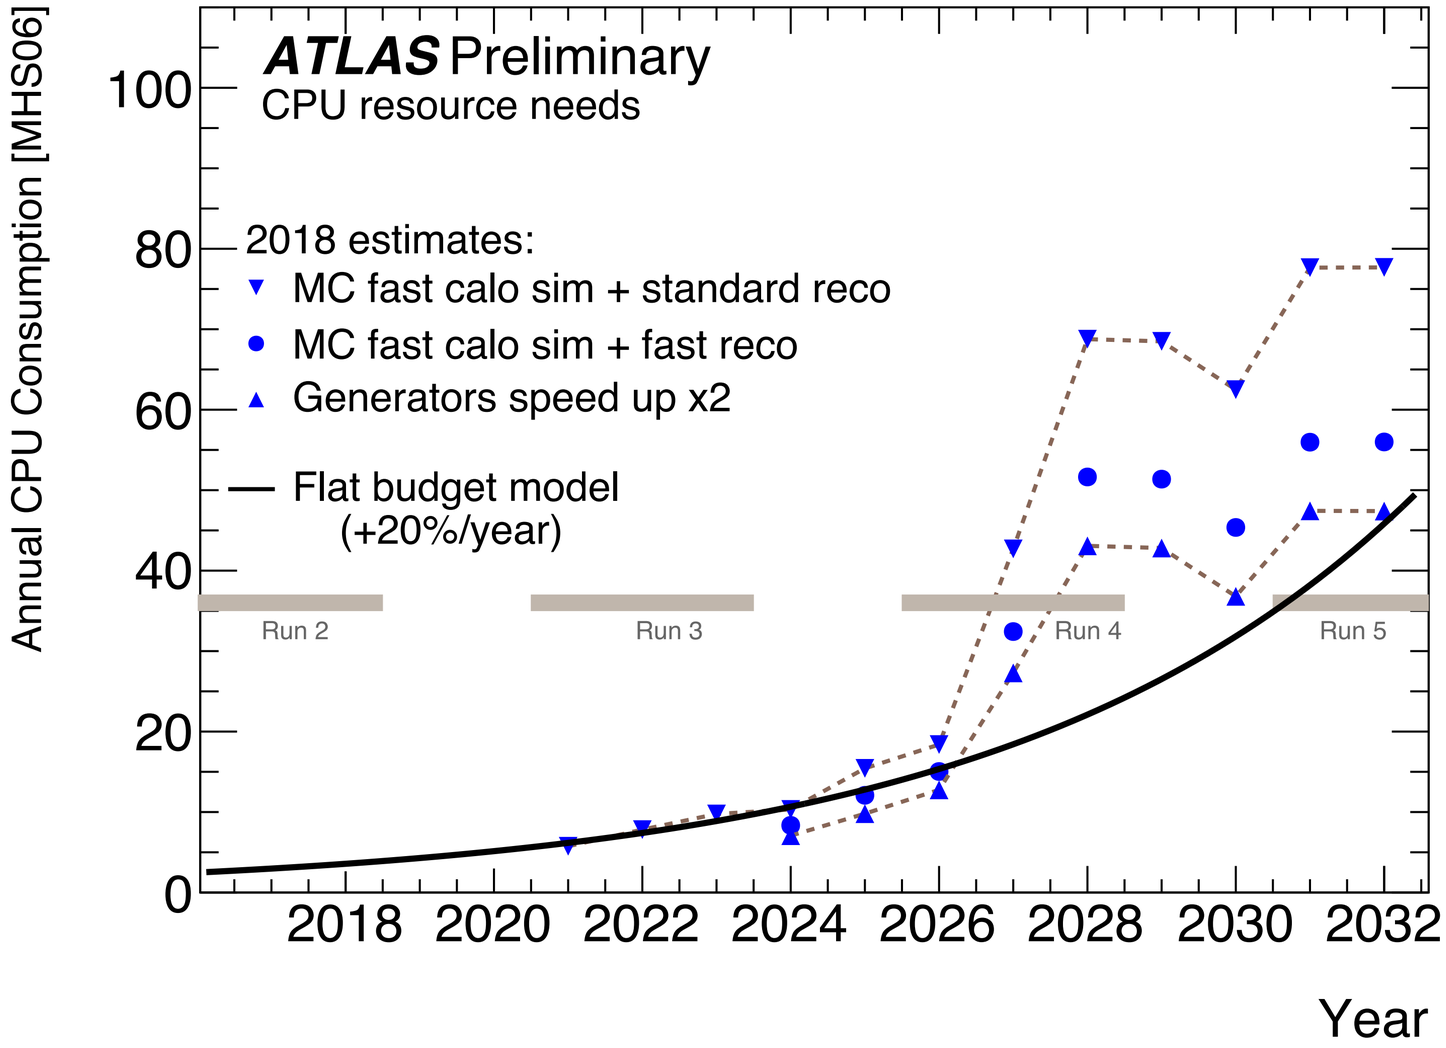

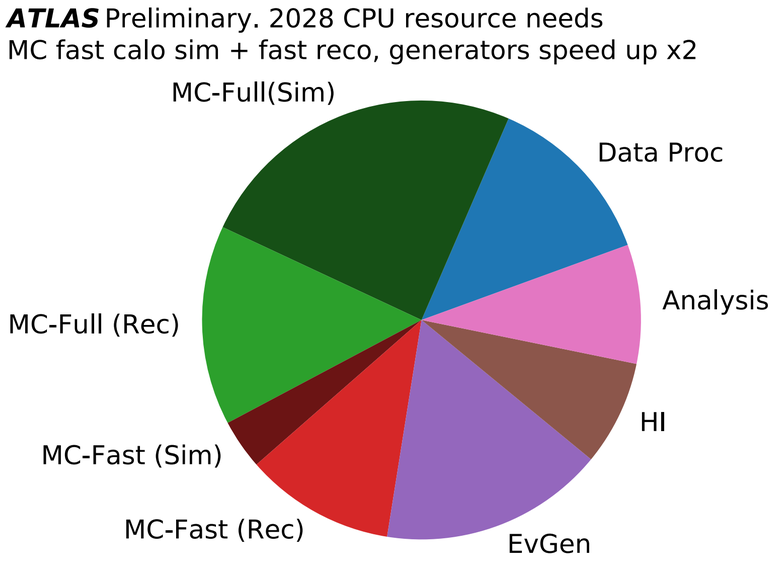

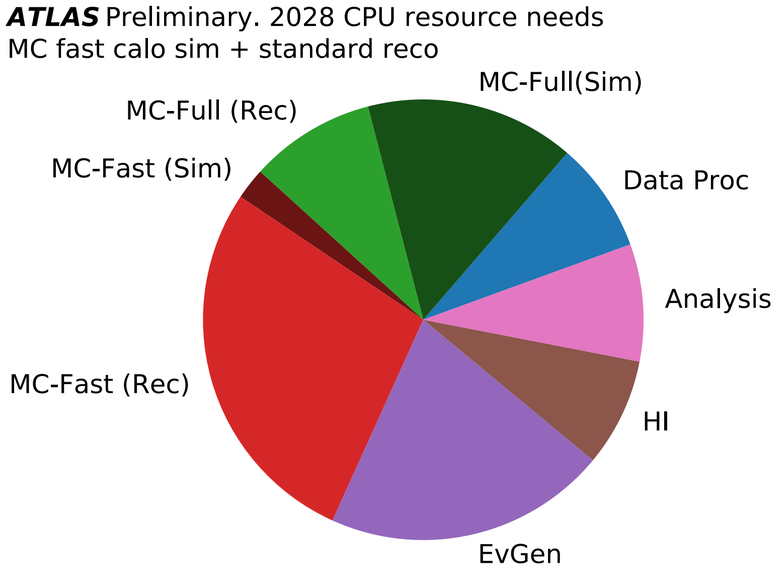

Two reasons why HPCs are important to ATLAS

Reason 2: they may form a part of our pledge in future years

(and they make a valuable beyond-pledge contribution already)

so we **have** to use them efficiently

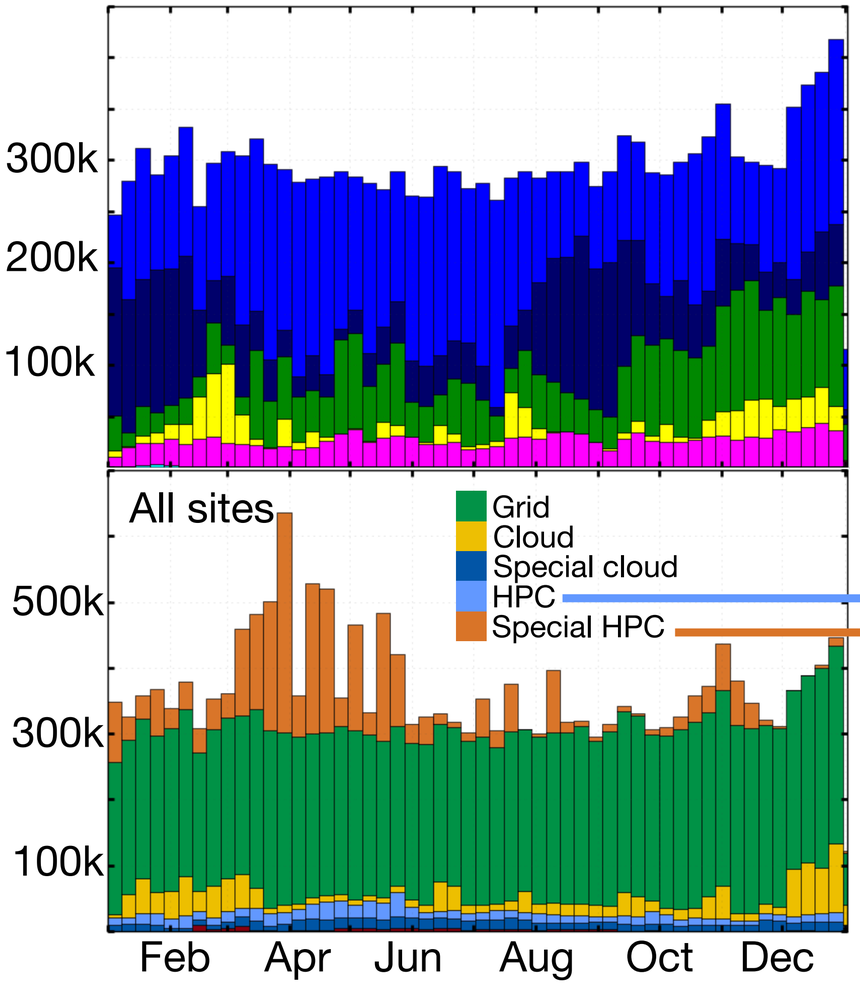

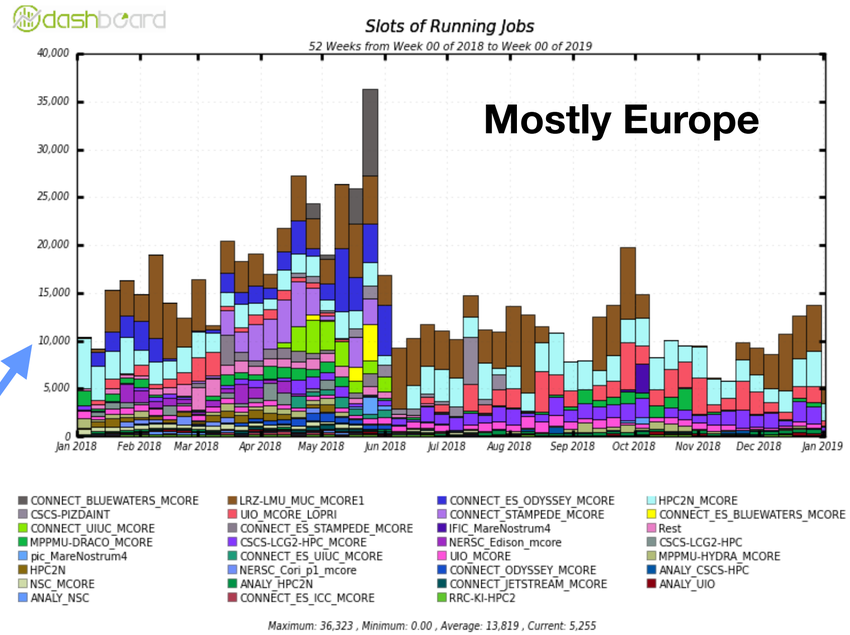

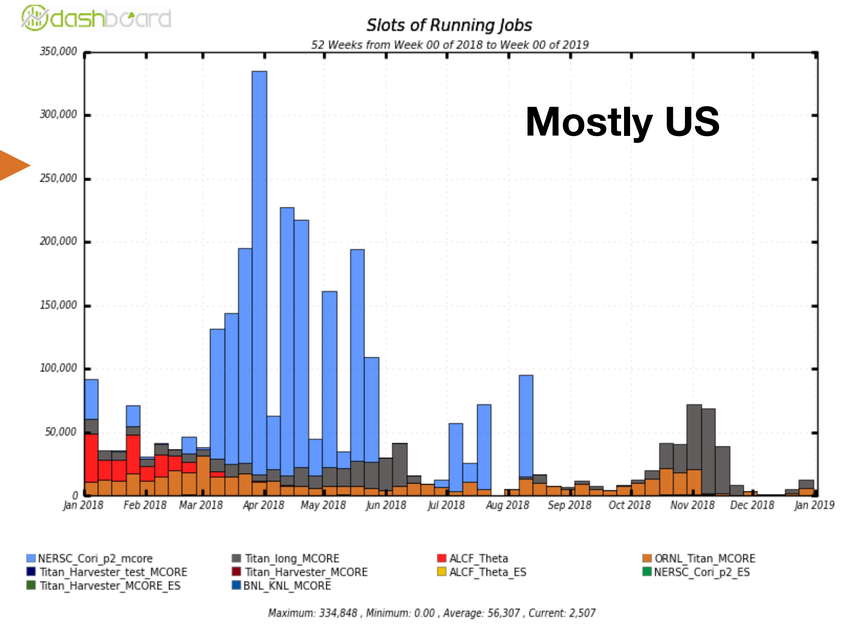

Side comment: what are the plans for Asian HPCs?

Success story: ease of use

- Ideally we should be able to seamlessly access HPC resources using the normal distributed computing machinery
  - Harvester/Pilot2/event service are making this a reality provided the hardware works with our software
- Completion of the development work in LS2 on Harvester, Pilot2 and the event (streaming) service will be crucial

Distributed computing deliverables	Run 3 physics	Run 4 physics
Grid and HPC jobs submission with Harvester	CRITICAL	BLOCKER
Pilot2 deployment and comissioning on Grid and HPC	CRITICAL	BLOCKER
Support for containers, event based workflows and GPU usage	MAJOR	BLOCKER

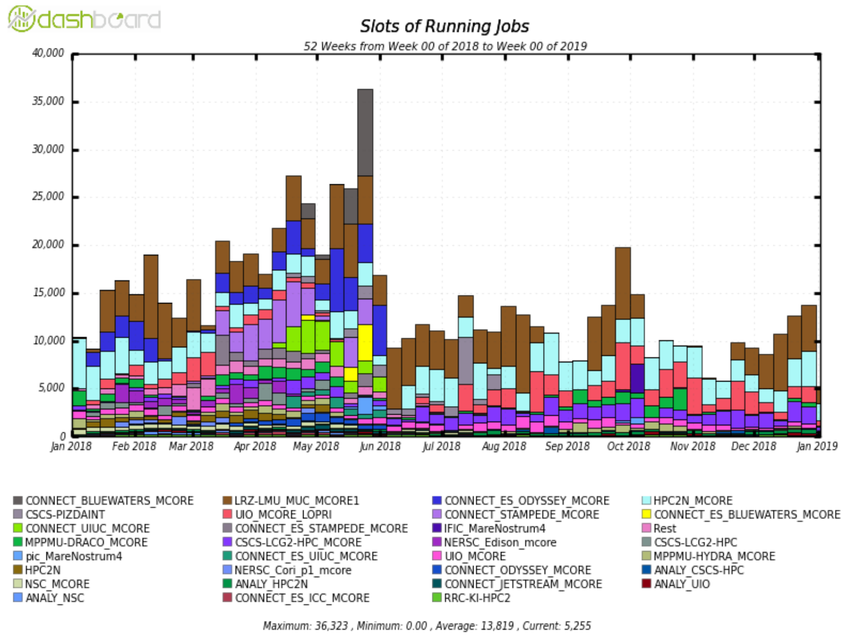

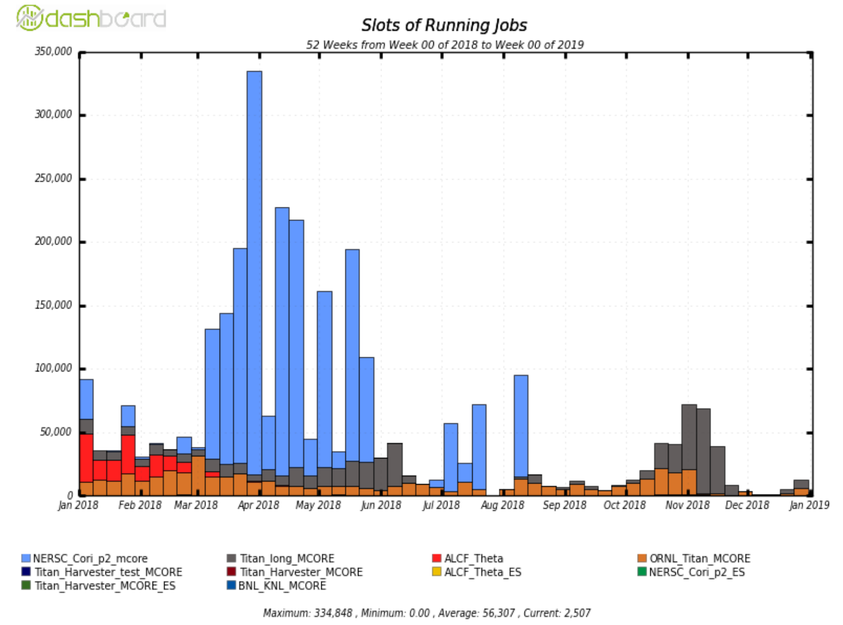

Suitable workflows

- What should we use HPCs for?
- What features of the workflows will enable us to extract the best value from HPC resources?
  - Low I/O, high CPU?
  - Minimal requirement for data transfer from the wider grid?
  - Most amenable to back-filling?
- This would point to the following:
  - **Definitely**: Event generation and simulation (sim is the overwhelming user at the moment)
  - **Probably not**: user analysis
  - **Maybe**: reconstruction and analysis data format production?
- Would increased use of fast simulation put a different complexion on this?

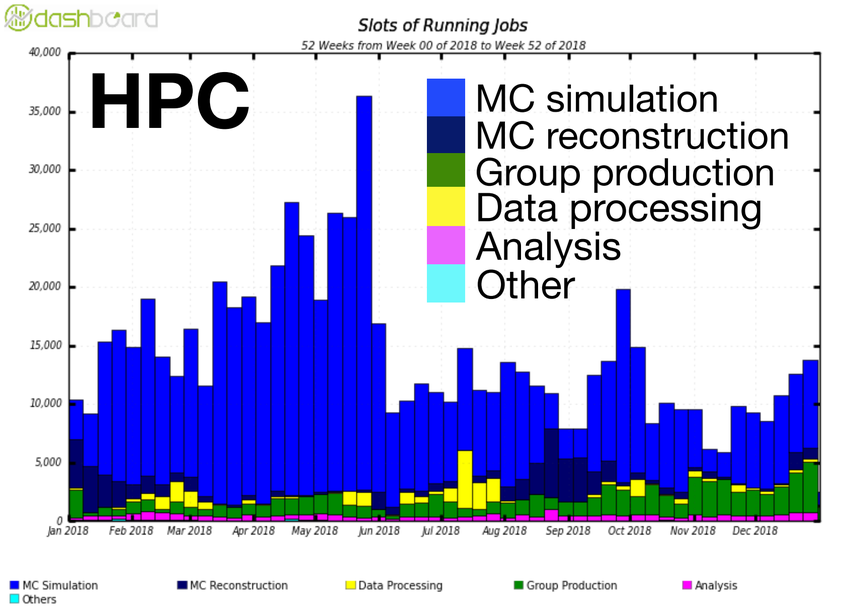

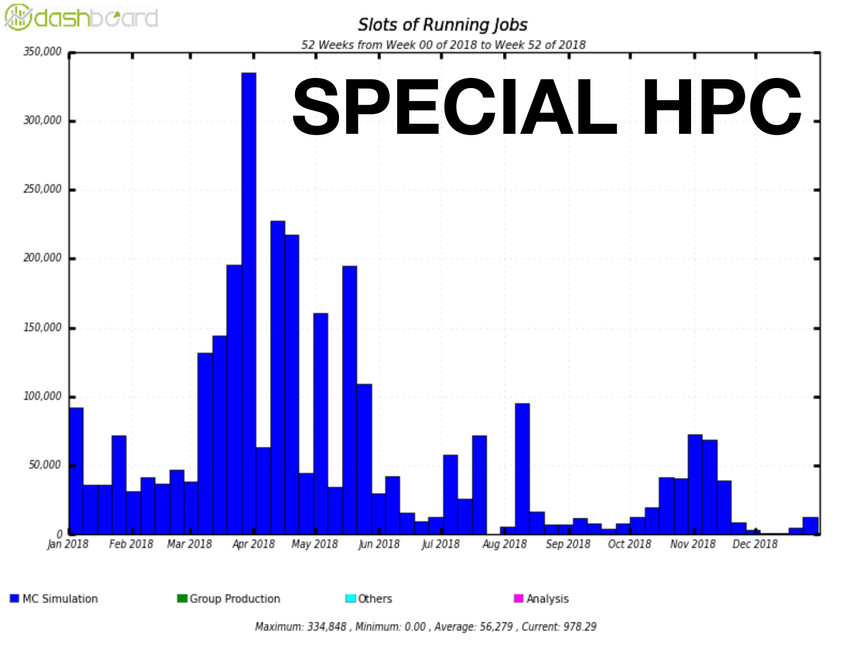

New architectures and accelerators

- We assume that accelerators will be a major part of the new "exascale" machines in the US
- Probably essential to use accelerators in Run 4
- Does "accelerators" really mean GPUs? From Nvidia? Or FPGAs? Some new CPU-accelerator hybrid?
- Non-x86 CPUs? Power9? ARM?
- We couldn't be expected to re-implement large parts of the code base every time a new architecture appeared - some kind of stability would be needed
- Note that CMS claim to have this largely sorted
- But note that ATLAS tried something similar with the HLT and it wasn't clear whether the human effort required to make it work was commensurate with the performance gains

Application	Advantages	Drawbacks
Deep learning training for physics analysis	Known to work and be useful; let Tensorflow do the work	Might not keep an exascale machine busy for long enough
Deep learning for (parts of) reconstruction/ simulation	Potential to integrate accelerators into the production; let Tensorflow do the work	Not proven to work effectively; might not keep an exascale machine busy for long enough
Port relevant parts of the existing reconstruction/ simulation	Uses existing workflows; not relying on unproven techniques; would keep machines busy	Would require a huge amount of work; would the gain be worth it?
Run event generation (phase space integration step) on GPUs	Big physics impact; would keep machines busy; small code base and low I/O	Who would write the code?
Limit setting and other statistical operations at the end of the analysis	Substantial progress in porting statistical tools to GPU (GooFit)	Might not keep an exascale machine busy for long enough; maybe more human effort than gain?

ATLAS strategy for HPCs: distributed computing

- Intense work within ATLAS to develop tools to enable HPCs of various types to be used more transparently and seamlessly for more types of workflow
- **Harvester:** seamless access to resources, be they pilots, VMs, MPI workers, batch workers etc. Can therefore be used to submit work to grid, cloud, HPC
  - All US HPCs are accessed with Harvester; GPU payloads under investigation
  - Status report: Thursday @ 11am (Fernando Megino, Tadashi Maeno)
- **Pilot2**: integration with Harvester; also includes container support
  - For HPC: operational at Titan - 15M events produced at the end of 2018
  - Commissioning for other US HPCs under way now
  - Status report: Wednesday @ 11:15 (Paul Nilsson)
- **Event (streaming) service:** event-by-event data access - essential for exploiting opportunistic resources
- **Containers**: essential for getting software distributions to non-networked HPCs and for deep learning on GPUs
  - Reports on Wednesday @ 9:20, 9:35 and 16:35 (Wei Yang and Alessandra Forti)

ATLAS strategy for HPCs: software

- Ensure release 22 can run multi-threaded by 2021
- Continue the intensive work to understand how deep machine learning could be integrated into the ATLAS bulk data processing (reconstruction and simulation)
- Begin work to understand how GPUs could be used in the ATLAS software
  - Work starting in the core software area and SPOT
    - Understanding how GPUs perform for our kind of work
    - Try to identify parts of the code which could in principle be factored out and farmed out to GPUs
    - Try to understand how to make this independent of architecture (e.g. not assuming Nvidia)
  - Interact with HSF to encourage common development efforts
    - Event generation looks like it might be low-hanging fruit
- Understand how we could run ATLAS software on non-x86 architectures without changing the code base