

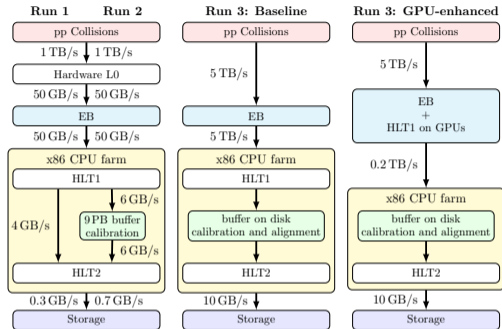
The Allen Project: a GPU trigger for LHCb

Daniel Craik, Mike Williams

2020-05-28



The Allen Project



- LHCb upgrading to streaming trigger for upcoming Run III
- Move to a GPU-based HLT1 with GPUs installed on the Event Builder servers
- Free up full CPU farm for HLT2 and save on networking between event builders and CPU farm
- Selected as baseline trigger implementation

MIT Personnel:

- Mike Williams
- Dan Craik (IRIS-HEP funded)
- Tom Boettcher
- Sean Condon (IRIS-HEP summer fellow)

Collaborators:

- Wider LHCb Allen team
- LPNHE, Paris
- NIKHEF/Maastricht
- ...

Milestones: years 1 & 2

Algorithm development:

- Introduce the ability to configure parameters of GPU algorithms at run time ✓
- Develop functionality to perform asynchronous monitoring ✓

Scale-up Testing:

- Implement realistic frequency of run and alignment changes
Work ongoing to test run and alignment changes
- Monitor throughput, decision fraction, hardware performance *etc.* during extended running time ✓

Integration Testing:

- Run on two servers similar to LHCb event builder nodes ✓
- Test monitoring by mimicking failures ✓
- Implementing run and alignment changes
Functionality developed to interface Allen with the LHCb event control system and further developments ongoing to track run numbers *etc* internally.

Deployment:

- Merged and deployed in the official (publicly available) Allen software stack ✓

Year 3

- Working on defining milestones now in consultation with the Allen team
- Likely to focus on monitoring and data quality

Longer term

- Allen gives us access to all experimental info from every single event at LHCb
 - Access to low- and high-level objects in single stage
- Implementing automatic (experimental) anomaly detection in Allen would be an interesting – and fun! – (grand?) challenge

Expected outcomes

- Develop and deploy algorithms to configure and monitor fully GPU-based HLT application – and have it all running at LHCb
- Possibly deploy automatic anomaly detection