


Triggering with GPUs

DPF Early Career Instrumentation Award 2024

Roel Aaij

November 22nd, 2024

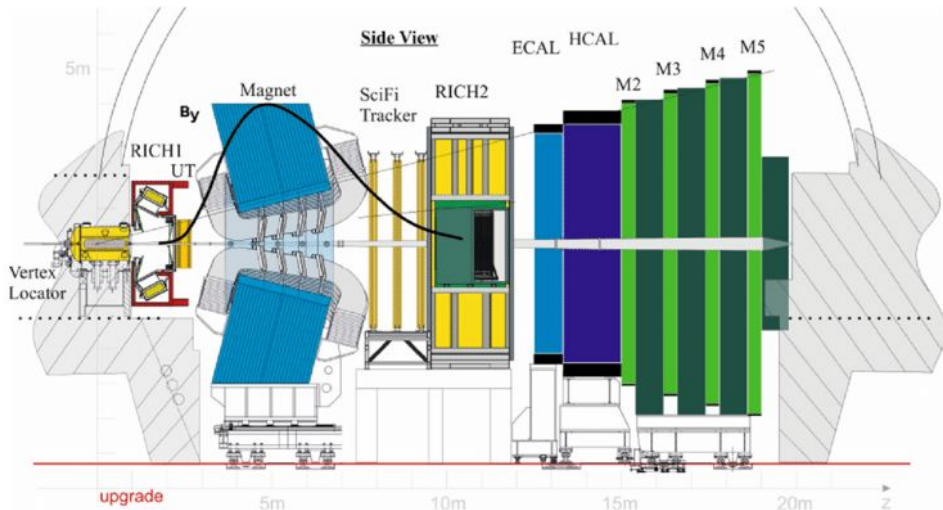
Nikhef



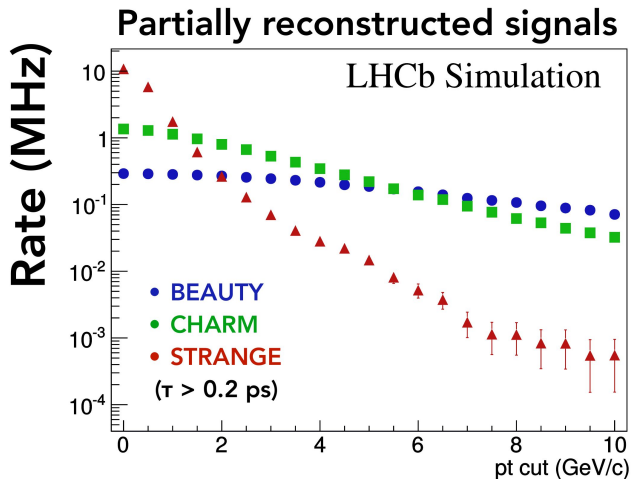
A Bit of (my) History

- Joined LHCb (R. Aaij et. al.) as an undergraduate student in 2008 (UvA)
- Started my Ph.D. in 2010 (VU Amsterdam);
 - Analysis of time-dependent CP violation in decays of $B_s^0 \rightarrow J/\psi\phi$
 - The LHCb software trigger (HLT)
- In 2010/2011, the HLT strategy was updated several times, because new ideas, additional resources and better software allowed more physics
- Applied Fellow at CERN starting in 2015
 - “Split” HLT1
 - HLT2 framework and monitoring
 - Responsible for HLT operations
 - Development and commissioning of online calibration and alignment tasks
- Staff physicist at Nikhef since 2018
 - GPU HLT1 at LHCb

The Run 3 LHCb Detector




LHCb Upgrade Trigger Motivation



**30 MHz (4 TB/s) of input contains a ~MHz of signal,
while we can only store 10 GB/s long-term**

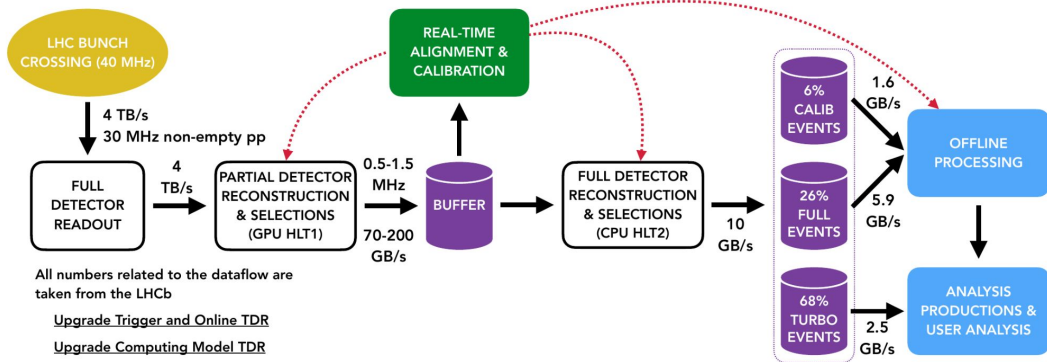
LHCb Upgrade Trigger Motivation

[J.Phys.: Conf. Ser. 878 012012](#)



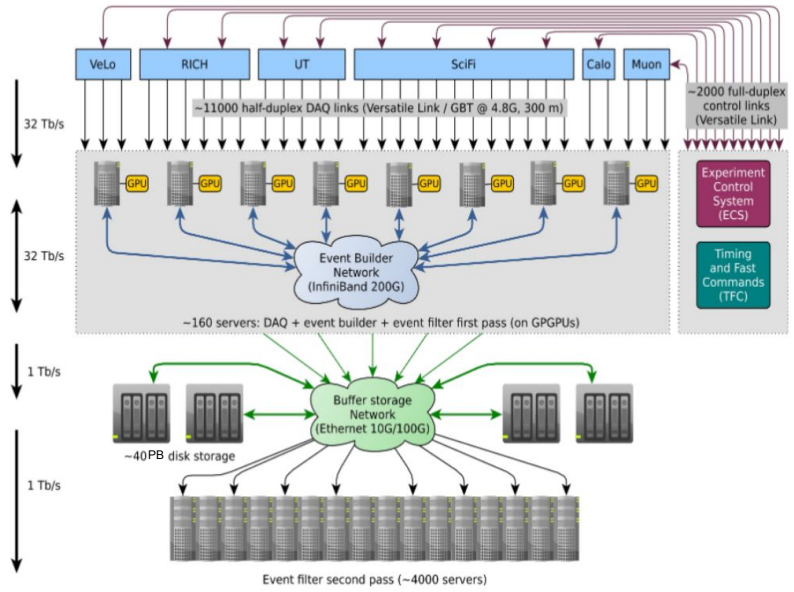
Run I/II hardware trigger doesn't scale

LHCb Upgrade Dataflow



**HLT1 challenge: reduce ~4 TB/s to 100 GB/s
in real-time with high physics efficiency**

LHCb Upgrade Trigger and DAQ



HLT1/Allen Philosophy

- Must not interfere with event building
- Do everything on the GPU: raw data in, decisions and candidates out
- Maximise (GPU) algorithm performance
- Start with barebones framework and write kernels in CUDA
- Implement performant reconstruction algorithms, i.e. significantly faster/\$ than on CPU
- Batches of ~1000 events with per-event logical control flow
- Data model will evolve so keep it simple
 - little to no dynamic memory allocation
 - SOA containing (small structs of) PODs
 - Count first, write later
- Minimise serialization of event data
- Opportunistic use of the CPU
 - Prefix sums
 - Monitoring
 - Low IPC algorithms that require little data



[Frances E. Allen](#)

1-Slide HLT1/Allen Framework

- gitlab: gitlab.cern.ch/lhcb/Allen
- C++17 (20 in progress), CUDA (12.X), HIP (6.X re-work in progress)
- Built with CMake and runs on CPU (for simulation) and GPU (in production)
- Standalone builds and “stack” builds

- Single precision throughout
- Batches of ~1000 events (~100 kb/event)
- GPUs have their own memory; framework provides functions to copy data
- ~10 batches in parallel using GPU streams (1 CPU thread per stream)
- No dynamic allocations
- Configurable (in Python) sequence of algorithms
- It's own asynchronous event loop
- All algorithms written from scratch for good performance on GPUs
- documentation: <https://allen-docs.docs.cern.ch/>


Integration with LHCb stack

- Need geometry and conditions data
- LHCb conditions change slowly (a 10 minute run is considered short)
- All required geometry and conditions data are converted to blobs that can be memcopy'd to the GPU; content optimized for parallel reads
- Use parameterisations when possible, e.g. currently no magnetic field map on the device

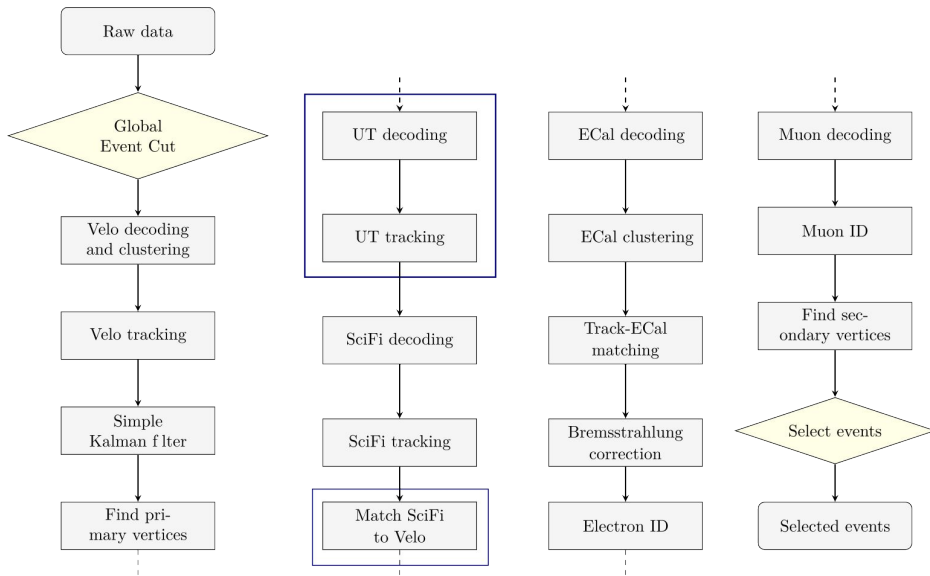
- Setup an LHCb-like fake event loop “on the side”
- When data with a new run number arrives:
 - Finish processing data of previous run
 - Trigger a single event in the fake event loop to update blobs
 - Copy new blobs to the device
 - Restart processing
- LHCb Online code used to allow steering by the control system
- Input and output to the DAQ in dedicated classes

LHCb HLT1

- Inclusive trigger
- Reduce rate from 30 MHz to 1 MHz
- TDR Requirements:
 - Velo tracks
 - Primary Vertices
 - “Long” tracks (Velo->UT->SciFi)
 - Muon ID
- Beyond-TDR ingredients:
 - ECal reconstruction
 - Electron ID
 - Photons
 - “T-Tracks” (SciFi)
 - “Downstream” tracks (UT-SciFi)
 - SMOG: LHCb as a fixed target experiment
- Avoid global event cuts




HLT1/Allen Kernels



In the Datacenter



In the Datacenter




In the Datacenter



DAQ with GPUs

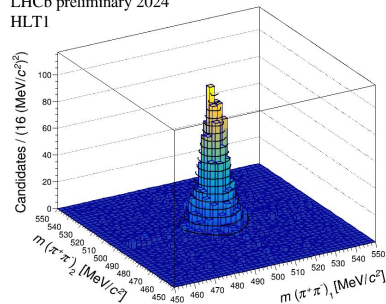
- 500 GPUs installed in Event Builder servers
- Input data copied to GPUs in EB format:
~1000 multi-fragment-packets in
multi-event-packets of 30000 events
- 20-25 GB/s per server
- Event data memory layout “transposed” with
respect to event-by-event
- Input data directly from shared memory
- Output in custom binary format to DAQ
- Experiment Control System steers HLT1
- Obtain geometry and conditions from LHCb
software on-the-fly



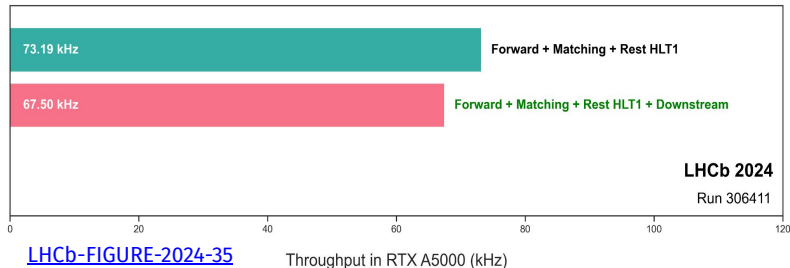
Performance

- Real Time Analysis proven to work
- Required HLT1 throughput achieved
- Significant improvements in trigger efficiencies
- Huge gain at low p_T
 - Particularly important for charm and strange physics
- Large impact on channels with electrons

LHCb preliminary 2024
HLT1



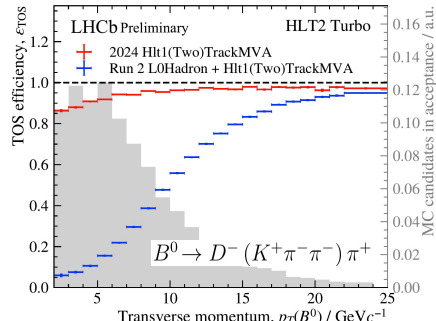


[LHCb-FIGURE-2024-13](#)



Performance

- Real Time Analysis proven to work
- Required HLT1 throughput achieved
- Significant improvements in trigger efficiencies
- Huge gain at low p_T
 - Particularly important for charm and strange physics
- Large impact on channels with electrons

[LHCb-FIGURE-2024-30](#)




People

- HLT1/Allen project grew quite fast from the start
- Part of an excellent core team with complementary skills
 - GPU performance and parallel programming
 - Knowledge of the experiment and existing stack
 - Project management
 - Modern C++/CUDA
 - Mixed CS/HEP background
- Support from DAQ/Online

- Most contracts are short, so bringing new people and skills on board has to happen all the time
- Knowledge transfer is hard
- Skilled people determine the success of a project such as HLT1/Allen
 - They are highly motivated, but career prospects matter
 - HLT1/Allen has already contributed to at least 4 people getting a big grant or a permanent position

Conclusion

- HLT1/Allen has been a big success
- It has been a pleasure to work on and lead the project
- Success of the project depends on many other people
 - LHCb subdetectors
 - Calibration and Alignment
 - LHCb DAQ/Online
 - The broader RTA project
 - The entire LHCb commissioning and operations team



BACKUP

Reconstruction Performance

