The LHCb Split HLT, How to Trigger More for Less

Roel Aaij

CERN, Geneva

April 12th 2016

DAQ@LHC

LHCb detector

2008 JINST 3 S08005

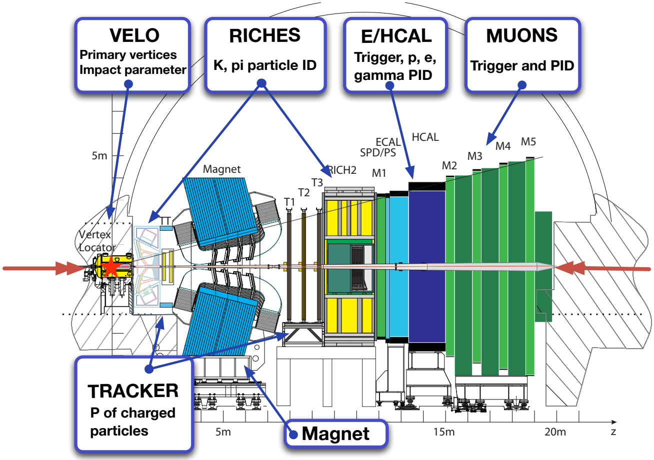

At 13 TeV and $\mathcal{L} = 4 \times 10^{32}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$:

$\sim$45 kHz $b\bar{b}$ pairs and $\sim$1 MHz $c\bar{c}$ pairs

Run I Trigger Overview

- LHCb detector read out at 1 MHz
- Hardware trigger (L0)
  - Based on multiplicity, calorimeters and muon detectors
  - Fixed latency of 4 μs
  - Reduces rate to 1 MHz
- Software trigger (HLT)
  - Runs on HLT farm
  - Split in two stages: HLT1 and HLT2
  - Events buffered to allow processing out of fill
  - Output rate 5 kHz
  - 65 kiB events
  - Total time budget O(35) ms/event

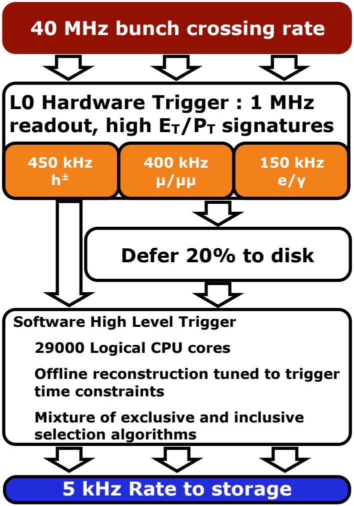

Deferred Trigger

- Stable beams 30% of the time.
- Disks are much cheaper than CPUs.
- Buffer events to local disks to allow for out of fill processing, effectively tripling the size of the HLT farm.
- Gives a large boost to trigger capabilities and discrimination power.
- $\sim 20\%$ of the events arriving from the L0 trigger written to disk.
- Implemented for the start of 2012.
- At the end of a run, HLT processes stopped and restarted with local disk input.
- Fully automatic procedure, driven by the control system.

From Deferred to Split Triggering

- Larger real-time reduction allows more efficient use of buffers
  $\rightarrow$ buffer after HLT1 $\rightarrow$ HLT split in 2 applications
- 5 PiB buffer on local disks (10 PiB in 2016)
- Space for 160 (320) hours of data with 150 kHz of 60 kiB events

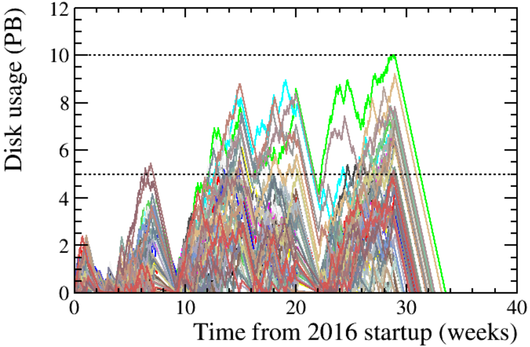

- Allows HLT1 output to be used for calibration and alignment
- Buffer utilisation tuned based on LHC schedule, including TS

Run II Trigger Overview

- LHCb detector read out at 1 MHz
- Hardware trigger (L0)
  - Based on multiplicity, calorimeters and muon detectors
  - Fixed latency of 4 μs
  - Reduces rate to 1 MHz
  - Higher thresholds in Run II
- Software trigger (HLT)
  - HLT farm nearly doubled.
  - HLT Split in two applications: HLT1 and HLT2
  - Events buffered after HLT1
  - Output rate 12.5 kHz
  - HLT software 40% faster
  - Same reconstruction online and offline!
  - Requires offline quality calibrations online.

LHCb 2015 Trigger Diagram

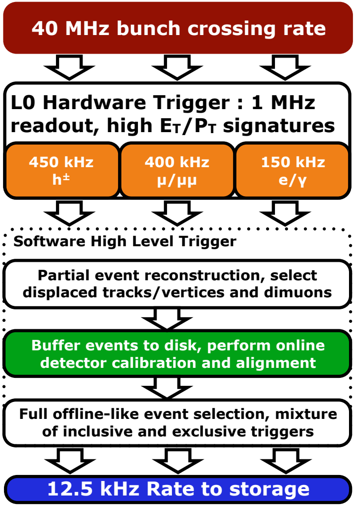

Control System

- As a result of the split into HLT1 and HLT2, HLT2 runs fully asynchronously.
- Different nodes potentially processes different runs with HLT2.
- Should be possible to run HLT1 and HLT2 in parallel to fully utilise farm.
- We have 2 shifters → running of HLT2 fully automated.
- HLT processes are checkpointed for fast start up and forked to maximise memory sharing.
- Processes can be started and stopped dynamically.

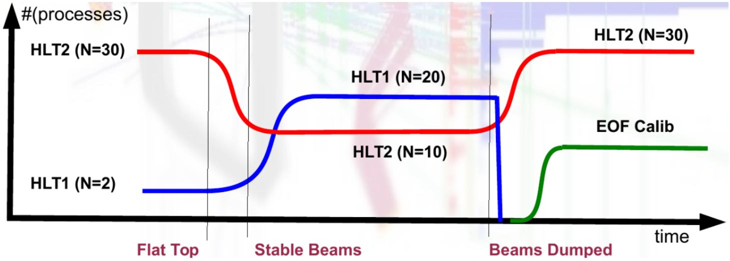

HLT1

- Inclusive selections:
  - Single and two track MVA selections
  - $\sim$ 100 kHz
- Inclusive muon selections
  - Single and dimuon selections
  - Additional low p_T track reconstruction
  - $\sim$ 40 kHz
- Exclusive selections
  - Lifetime unbiased beauty and charm selections
  - Selections for alignment
- Low multiplicity trigger for central exclusive production analyses
- Throughput $\sim$ 1.2 MHz

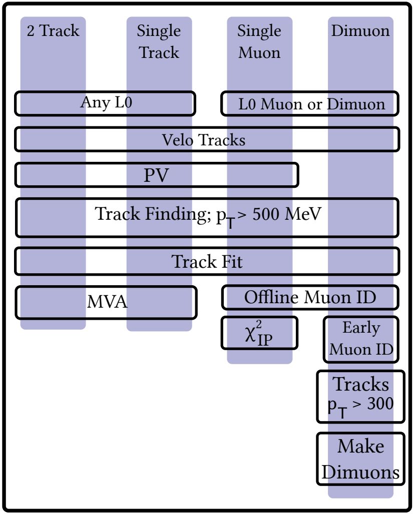

Real-Time Calibration and Alignment

- Same online and offline reconstructions requires prompt alignment and calibration
- Alignment per fill:
  - Collect suitable data with dedicated HLT1 selections, e.g. $D^0 \to K^+\pi^-$ and $J/\psi \to \mu^+\mu^-$
  - Run alignment workers on the HLT farm (1 per node)
  - Controller iterates until converged, O(5) min
  - Apply updates of Velo and/or tracker alignment if needed
  - RICH mirror alignment and muon alignment for monitoring
  - ECAL gain calibration
- Calibration per 1 h run:
  - RICH and Outer Tracker t_0
  - Available O(1) minute after collection of data

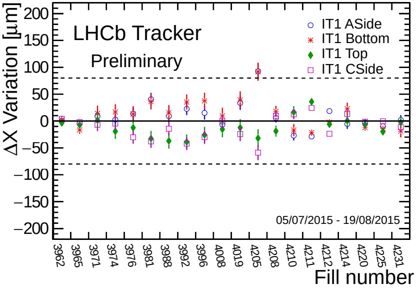

HLT2 Reconstruction

2012 CPU Heat Map

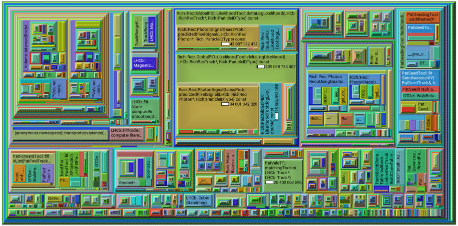

2015 CPU Heat Map

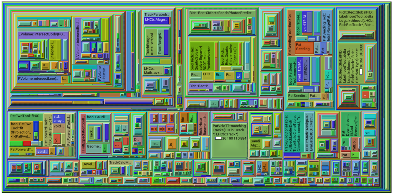

- Full event reconstruction
- Starts from HLT1 objects
- All charged tracks
- Neutral particles
- RICH, Muon and Calo PID
- Same reconstruction online and offline
- 30% speedup achieved
- Throughput $\sim$ 60 kHz

Reconstruction	Run II		Run I
HLT1 rate	$\sim$150	kHz	$\sim$ 80 kHz
HLT1 time	$\sim$ 35	ms	$\sim$ 20 ms
Track finding	$\sim$240	ms	
Track fit	$\sim$100	ms	
Calorimeter reco	$\sim$ 50	ms	
RICH PID	$\sim$130	ms	
Muon ID	$\sim$ 2	ms	
Total HLT2	$\sim$630	ms	$\sim$150 ms
HLT2 rate	$\sim$ 12.5	kHz	$\sim$ 5 kHz

HLT2 Selections

- Inclusive beauty selections:
  - MVA based 2, 3, and 4 body detached vertices
  - Dimuon selections
- Exclusive beauty selections:
  - E.g. $B \to \varphi\varphi$, $B \to \gamma\gamma$
- Charm selections
  - Inclusive selection of $D^* \to (D^0 \to X)\,\pi^+$
  - Charmed baryons
  - Final states with K^0_S
  - 2,3,4,5—body final states
- Electroweak bosons
- ...
- More than 400 selections in total
- 12.5 kHz to tape

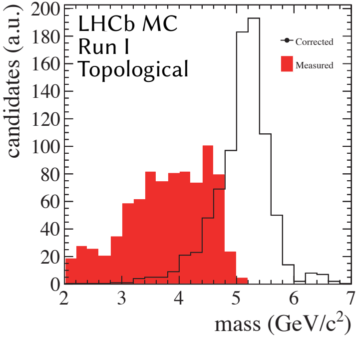

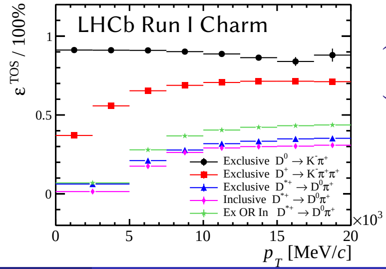

Int. J. Mod. Phys. A 30, 1530022 (2015)

Turbo Stream

- Offline reconstruction available online
  $\rightarrow$ do physics analysis with HLT candidates
- Turbo stream:
  - Store HLT candidate information.
  - Remove most of detector raw data.
  - Space required reduced by > 90 %.
- Ideal for high-yield analyses.
- O(24) h turn-around.
- This year, subset of HLT selections also store reconstructed objects.

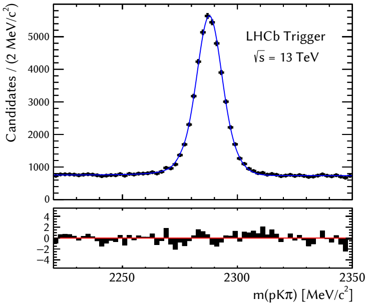

Turbo Stream

- Offline reconstruction available online $\rightarrow$ do physics analysis with HLT candidates
- Turbo stream:
  - Store HLT candidate information.
  - Remove most of detector raw data.
  - Space required reduced by > 90 %.
- Ideal for high-yield analyses.
- O(24) h turn-around.
- This year, subset of HLT selections also store reconstructed objects.

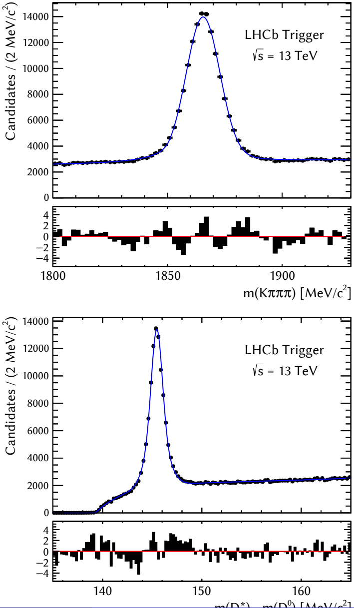

Summary

- Split HLT with intermediate buffering to local disks.
- Additional HLT farm purchased, now effectively 2 times larger.
    - 1800 servers
    - 27000 physical cores
    - 10 PiB disk space
- Full offline-quality reconstruction available online.
- Calibration and alignment running online.
- Software optimised to fit reconstruction in time budget.
- Turbo stream implemented, including possibility to store all reconstructed objects.