A 30 MHz software trigger for the LHCb Upgrade

Conor Fitzpatrick
on behalf of the LHCb collaboration

19th International Workshop on Advanced Computing
and Analysis Techniques in Physics Research
Saas-Fee, Switzerland

LHCb: The precision flavour experiment

- LHCb was built to study beauty and charm at the LHC:

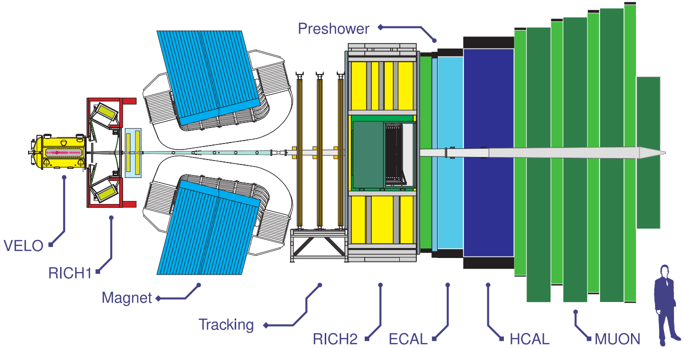

- Precise particle identification (RICH + MUON)
- Excellent decay time resolution: $\sim$45fs (VELO)
- High purity + Efficiency with flexible trigger

C. Fitzpatrick

March 12, 2019

The LHCb trigger

- A trigger is needed to reduce storage and readout costs
- A *good* trigger does so by keeping more signal than background

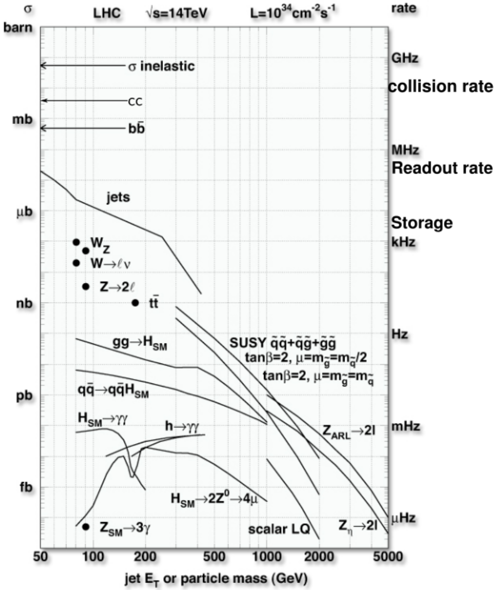

C. Fitzpatrick

March 12, 2019

The LHCb trigger

- A trigger is needed to reduce storage and readout costs
- A *good* trigger does so by keeping more signal than background
- General purpose LHC experiments are interested in signatures in the kHz region
  - Readout at 100kHz is efficient with reasonably straightforward E_T requirements

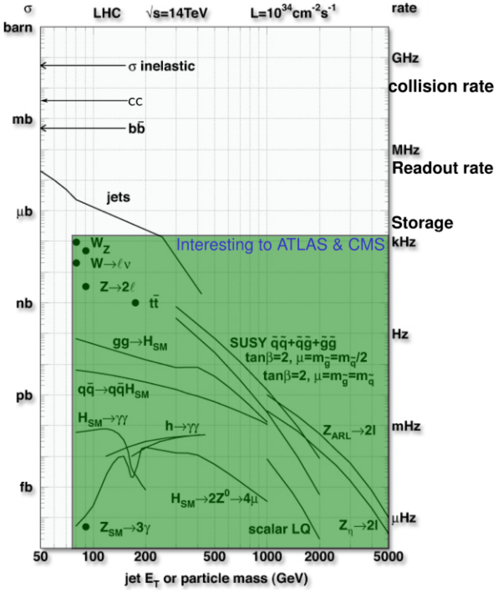

The LHCb trigger

- A trigger is needed to reduce storage and readout costs
- A *good* trigger does so by keeping more signal than background
- General purpose LHC experiments are interested in signatures in the kHz region
  - Readout at 100kHz is efficient with reasonably straightforward E_T requirements
- LHCb ($\mathfrak{L} = 4 \times 10^{32} \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) faces a unique challenge:
  - 45kHz of $b\bar{b}$, $\sim$1MHz of $c\bar{c}$
  - 1MHz readout is needed to stay efficient for beauty signals

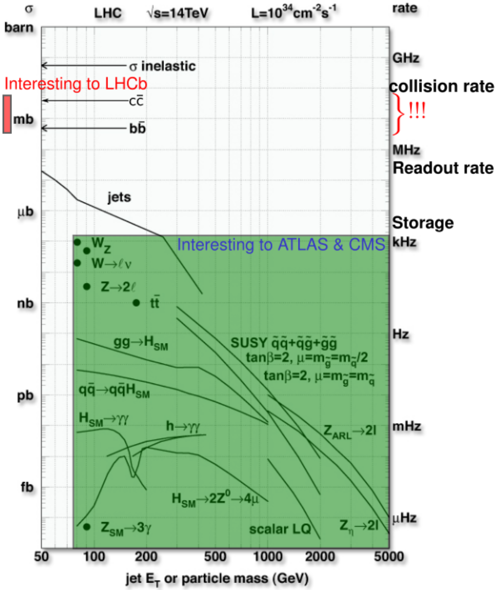

C. Fitzpatrick

March 12, 2019

The LHCb Run 2 trigger in two plots

- The LHCb trigger has to cover extremes of data taking:

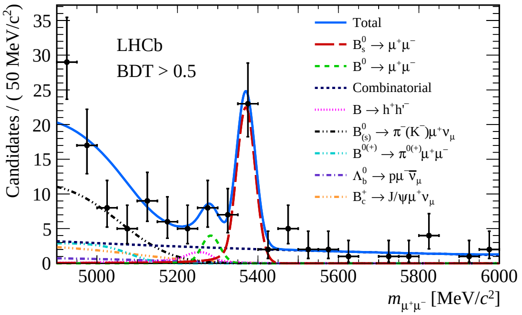

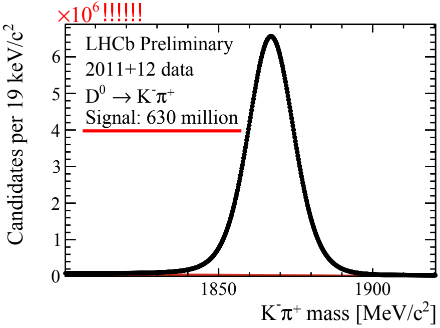

- High efficiency to collect rare decays like $B_s^0 \to \mu\mu$[1]
- High purity for enormous charm signals like $D^0 \to K\pi$[2]
- Must be flexible to operate in both extremes simultaneously: After readout, HLT has access to 100% of event in software

[1] Phys. Rev. Lett. 118, 191801 (2017)
[2] LHCb-CONF-2016-005

The Run 2 LHCb Trigger

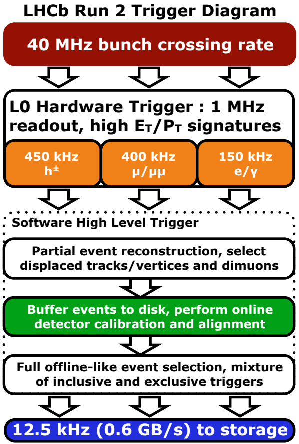

- The LHCb Run 2 trigger (2015-2019)
- Three trigger levels, with a hardware L0 stage:
  - Level-0 trigger buys time to readout the detector with Calo, Muon p_T thresholds: $40 \rightarrow 1$MHz
  - Events built at 1MHz, sent to HLT farm (~ 27000 physical cores)
  - HLT1 has $40\times$ more time, fast tracking followed by inclusive selections 1MHz $\rightarrow$ 100kHz
  - HLT2 has $400\times$ more time than L0: Full event reconstruction, inclusive + exclusive selections using whole detector
- Flexibility comes from software-centric HLT design[3]

[3] arXiv:1812.10790 [hep-ex], submitted to JINST

HLT1

- Beauty and charm hadron typical decay topologies:

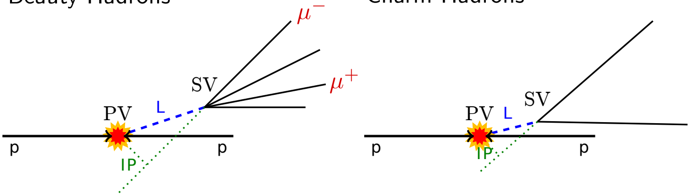

- $B^{\pm}$ mass ~ 5.28 GeV, daughter p_T $\mathcal{O}(1$ GeV$)$
- $\tau \sim 1.6$ ps, Flight distance ~ 1 cm
- Important signature: Detached muons from $B \to J/\psi X$, $J/\psi \to \mu\mu$

- D^0 mass ~ 1.86 GeV, appreciable daughter p_T
- $\tau \sim 0.4$ ps, Flight distance ~ 4 mm
- Also produced as 'secondary' charm from B decays.

Underlying HLT1 strategy:

- Fast reconstruction: Primary Vertices, High p_T tracks, optional Muon ID
- Inclusive triggering using MVAs on 1&2-track signatures

Disk Buffer

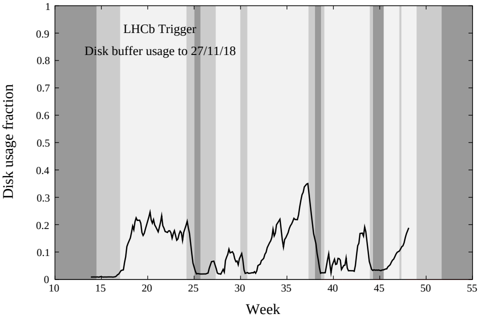

- HLT Farm is off-the shelf servers: Considerable (11PB) disk capacity
- HLT1 accepted events written to the disk in-fill at 100kHz: 2 week contingency
- HLT2 throughput in-fill is 30kHz, out of fill 90kHz when HLT1 isn't running
- Effectively doubles trigger CPU capacity, Farm is used twice for HLT, excess used for simulation
- Asynchronous HLT has another big advantage though...

Real-time Alignment + Calibration

- With Run 2 signal rates, efficient & pure output requires full reconstruction at HLT2
  - Online selections $\rightarrow$ offline selections
  - Reduces systematic uncertainties and workload for analysts
- Alignment and calibration of full detector in the trigger needed
- While HLT1 is written to disk, alignment & calibration tasks run

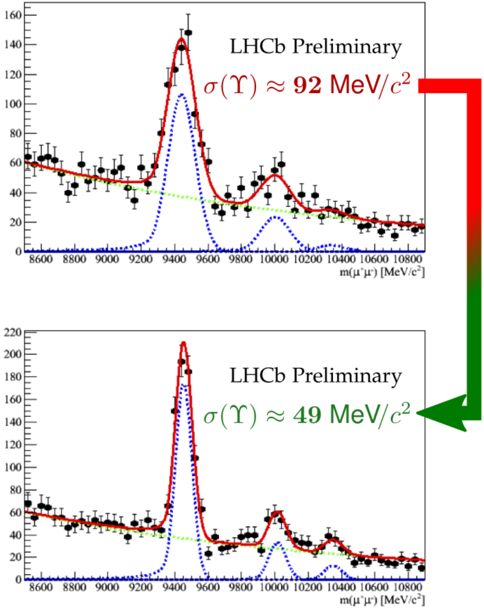

C. Fitzpatrick

March 12, 2019

A fully aligned detector

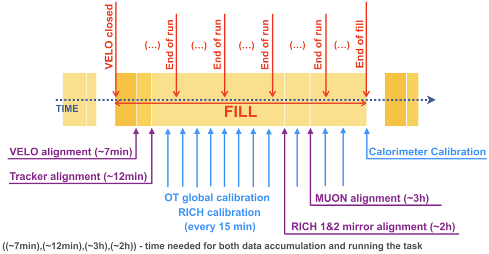

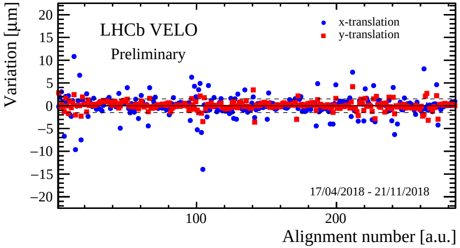

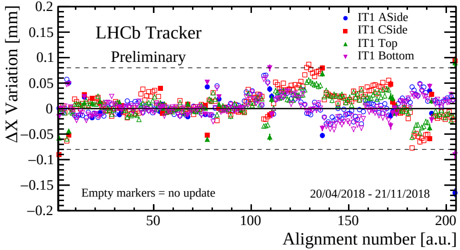

- All detectors are aligned & calibrated in-situ using the full HLT1 output rate
- Updates applied automatically if needed prior to HLT2 starting

HLT2: Reduced event formats

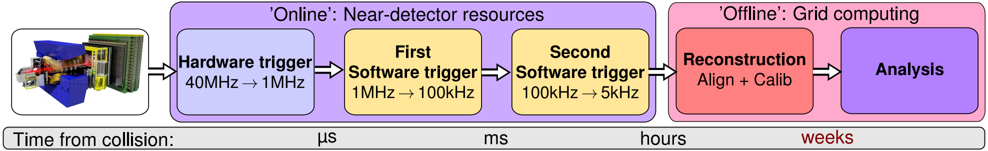

- Trigger *rates* aren't important, output *bandwidth* is
- Offline reprocessing previously needed to recover best quality

C. Fitzpatrick

March 12, 2019

HLT2: Reduced event formats

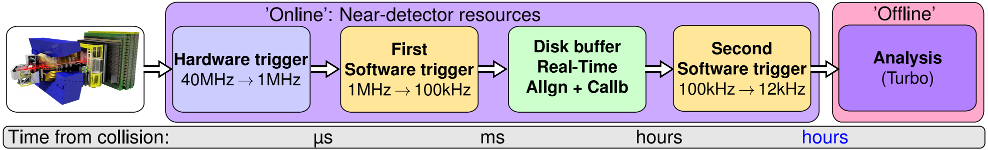

- Trigger *rates* aren't important, output *bandwidth* is
- Offline reprocessing previously needed to recover best quality
- After alignment: online == offline, why reprocess? Do analysis on trigger objects at HLT2, write only the relevant objects offline
- Significant reduction in event size → higher rates for the same bandwidth

HLT2: Reduced event formats

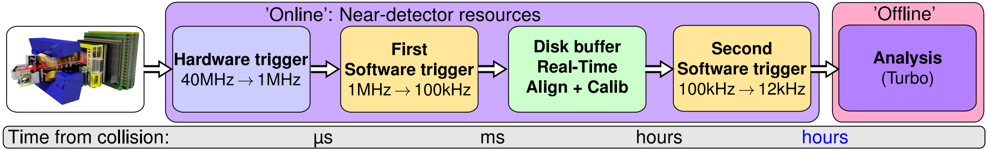

- Trigger *rates* aren't important, output *bandwidth* is
- Offline reprocessing previously needed to recover best quality
- After alignment: online == offline, why reprocess? Do analysis on trigger objects at HLT2, write only the relevant objects offline
- Significant reduction in event size → higher rates for the same bandwidth
- Added bonus: offline CPU freed up for simulation.

Turbo

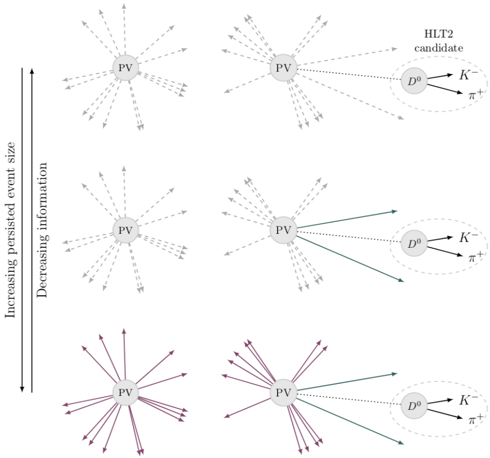

- Turbo is the LHCb paradigm for reduced event format data[4]
- High degree of flexibility: Save only as much of the event as is needed for analysis
  - Keep all reconstructed objects, drop the raw event: 70kB
  - Keep only objects used to trigger: 15kB
  - 'Selective Persistence' objects used to trigger + user-defined selection: $15 \rightarrow 70$kB

[4] arXiv:1604.05596, NEW arXiv:1903.01360

Turbo usage in Run 2

- 528 trigger lines at HLT2. 50% are Turbo
- 25% of the trigger rate is Turbo but it counts for only 10% of the bandwidth
- Many analyses would not be possible without Turbo[5]

CERN-EP-2017-248
LHCb-PAPER-2017-038
October 5, 2017

Search for dark photons produced in 13 TeV pp collisions

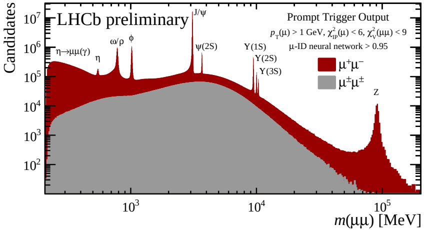

CERN-EP-2018-172
LHCb-PAPER-2018-026
October 18, 2018

First observation of the doubly charmed baryon decay $\Xi_{cc}^{++} \rightarrow \Xi_c^+ \pi^+$

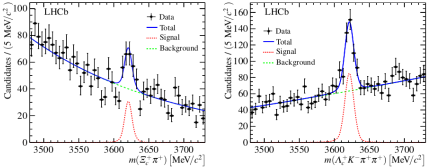

[5] Phys. Rev. Lett. 120, 061801 (2018), Phys. Rev. Lett. 121, 162002 (2018)

The MHz signal era

- Starting in 2021, LHCb will run at $\mathcal{L} = 2 \times 10^{33}\ \text{cm}^{-2}\,\text{s}^{-1}$: $5 \times$ more collisions per second

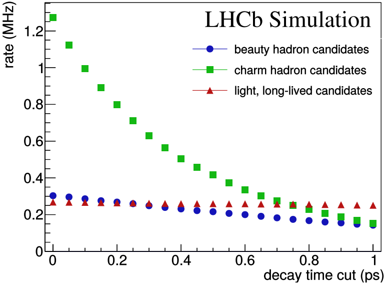

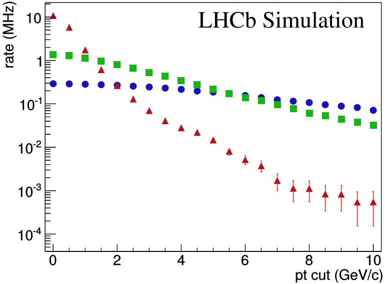

- Readout becomes a bottleneck as signal rates $\rightarrow$ MHz even after simple trigger criteria [6]

[6] LHCb-PUB-2014-027

So what 'stuff' can we throw away?

- The problem is no longer one of rejecting (trivial) background
- Fundamentally changes what it means to trigger

- Instead, we need to categorise different 'signals'
  - Requires access to as much of the event as possible, as early as possible

Reading out at 30MHz

- Solution: Readout and reconstruct 30 MHz of collisions in software

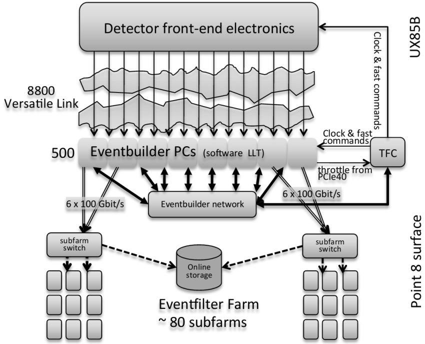

- Detector readout at the LHC bunch crossing frequency:
- Event builder, trigger farm & disk buffer in modular containers at the LHCb experiment area

The Run 3 Trigger

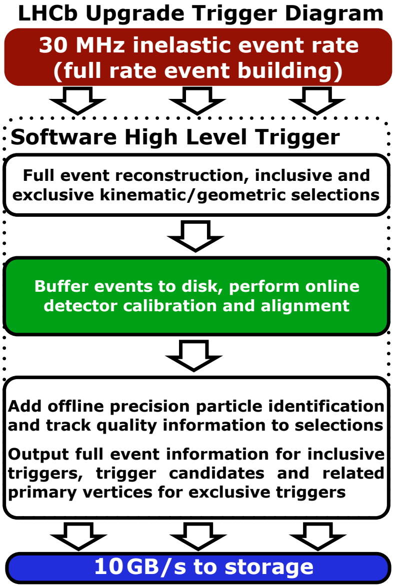

- Run 2: has proven the strategy at 1 MHz at a pileup of ~ 1
- Run 3: must now process full 30 MHz at $5\times$ the pileup
- Overall strategy similar, but:
  - HLT1 $\rightarrow$ first level trigger. Output 100kHz $\rightarrow$ $\sim$1MHz
  - Disk buffer has contingency of O(days) instead of weeks
  - HLT2 $\rightarrow$ second level trigger. 10GB/s mostly turbo output

Run 3 first level trigger

- 1- and 2- track performance under study[7]
  - MVA parameters for Loose and Tight configurations
  - Several tracking thresholds 500 → 1000MeV
- Results with minimal changes from Run 2:
  - 1-track needs more work
  - 2-track remains efficient

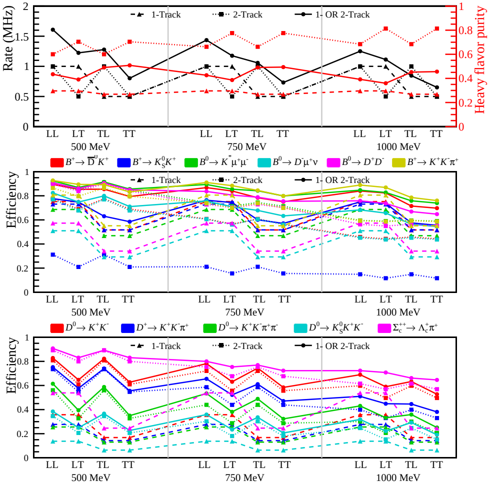

[7] LHCb-PUB-2017-006

Run 3 second level trigger

- Fully embrace the turbo paradigm: More exclusive selections than in Run 2, with wide adoption of MVAs
- Recent work to develop multivariate selections to select tracks generically coming from B and D decays[8]
- With many (> 500) trigger lines, sharing output bandwidth equitably is a challenge
- Genetic algorithm based procedure makes this easier, analysts decide between event size and output rate[9]:

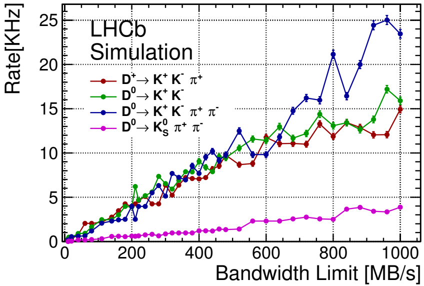

[8] NEW arXiv:1903.01360
[9] LHCb-PUB-2017-006

There's no turning back. . .

- Throwing away most of the event means care must be taken
- Turbo relies on never needing to reprocess:
  - Online monitoring & data quality are even more important
  - In Run 2 the disk buffer allows up to 2 weeks of safety margin
  - Not so in Run 3, where buffer will have O(days)
- Integration testing, real-time monitoring & robust procedures are critical components of the trigger
- During Run 2, we never needed to reprocess thanks to these procedures

Conclusions

- LHCb signal rates in the Upgrade change the definition of a trigger:
  - 'Rejects background' $\rightarrow$ 'categorises signal'
  - 'Reduces rate' $\rightarrow$ 'Reduces bandwidth'
- In order to efficiently categorise MHz signals, LHCb will use a triggerless readout into a software trigger
- Offline quality selections mean only subset of the event has to be saved for analysis
  - Requires fully aligned & calibrated detector in the trigger
- Not without its challenges: Extensive upgrades to the software as well as the detector
  - See talks in this session from N. Nolte, M. Cattaneo