The Trigger and Real-time Reconstruction at LHCb

CPAD Instrumentation Frontier Workshop 2019

Daniel Craik
on behalf of the LHCb collaboration

Massachusetts Institute of Technology

9th December, 2019

The LHCb detector

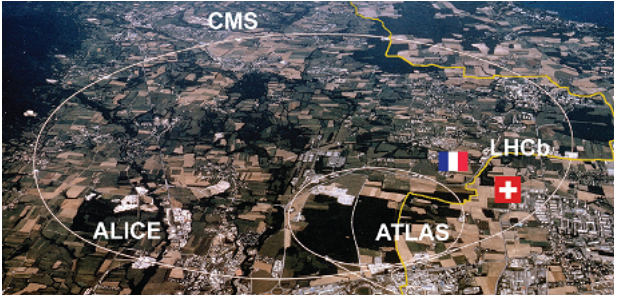

- Located at point 8 of the LHC
- General-purpose detector in the forward region
- Specialised in studying *b*- and *c*-decays

- Instrumented in the forward region to exploit forward-production of *c*- and *b*-hadrons

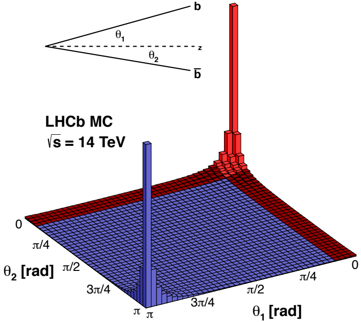

The LHCb detector

JINST **3** (2008) S08005

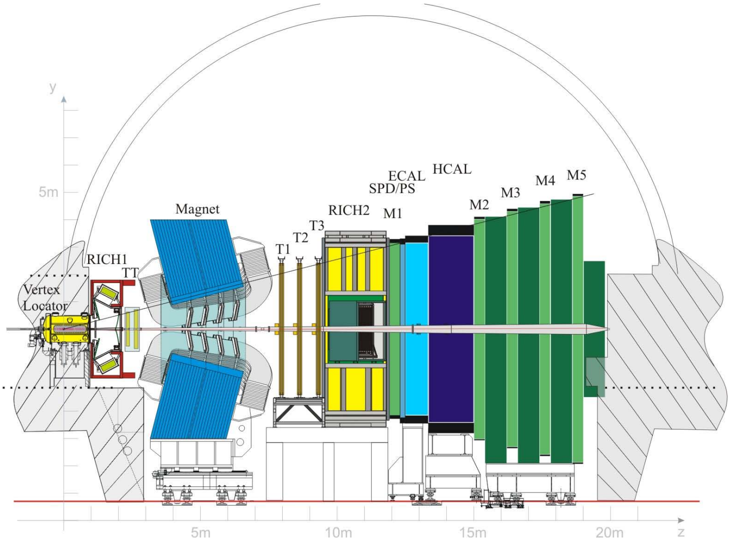

- Instrumentation in the forward region ($2 < \eta < 5$)
- Excellent secondary vertex reconstruction
- Precise tracking before and after magnet
- Good PID separation up to $\sim 100\,\mathrm{GeV}/c$

LHCb timeline

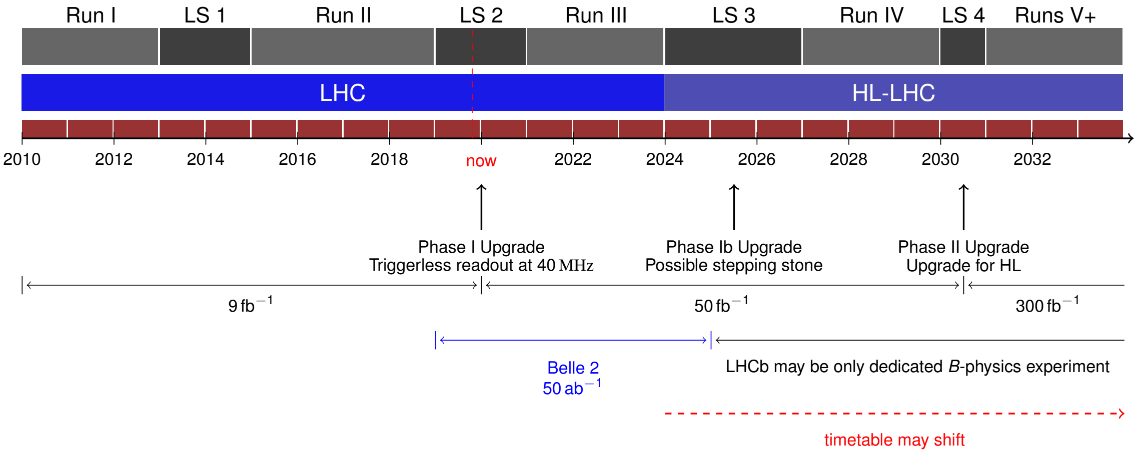

Real-time reconstruction in Run II

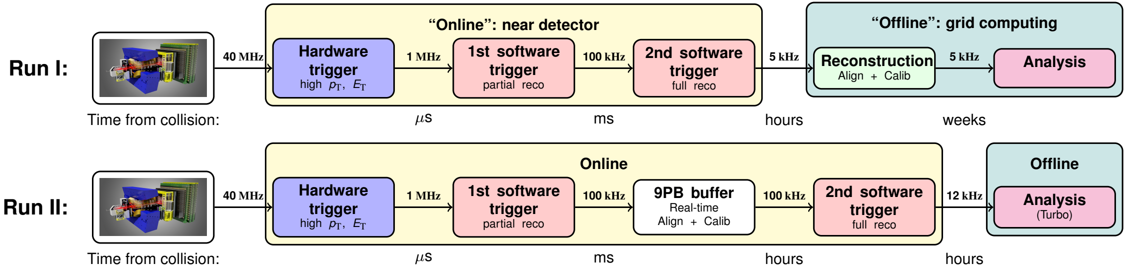

- Calibration and alignment of Run I data performed "offline" weeks after data taking
  - Trigger reconstruction different from offline
- In Run II, data buffered before final trigger stage
  - Allows for real-time alignment and calibration
  - Offline-like reconstruction within the trigger
  - Many analyses use "Turbo-stream" data – online reconstruction, full raw event not saved

Real-time reconstruction in Run II

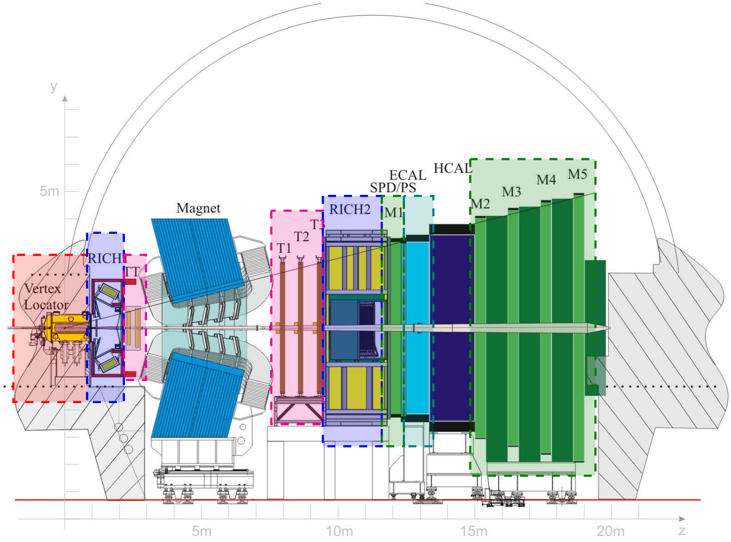

- Real-time alignment and calibration performed for vertex locator, RICH detectors, tracking stations, calorimeter and muon stations
- Will focus on velo and RICH
- Alignment particularly important for velo, which opens and closes between fills
- Gas-filled RICH detectors also require frequent calibration

Real-time reconstruction in Run II

- Each alignment task performed once per fill
- Alignment begins once a large enough dataset has been collected
- Calibration of RICH gas refractive index performed regularly to account for temperature/pressure changes within the radiator gas

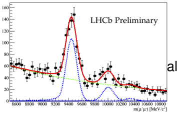

→ align + calib

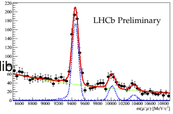

initial improved

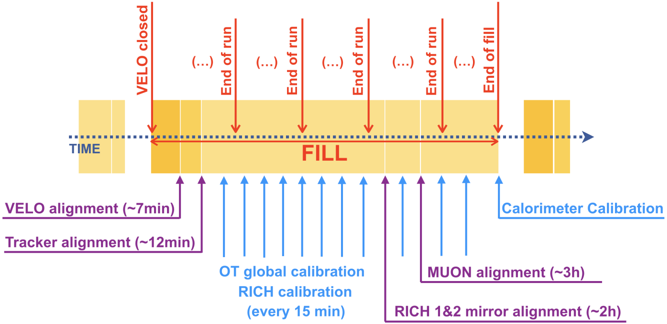

((~7min),(~12min),(~3h),(~2h)) - time needed for both data accumulation and running the task

Real-time alignment of the Velo

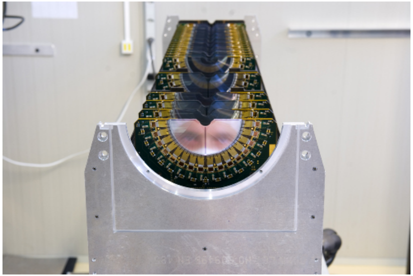

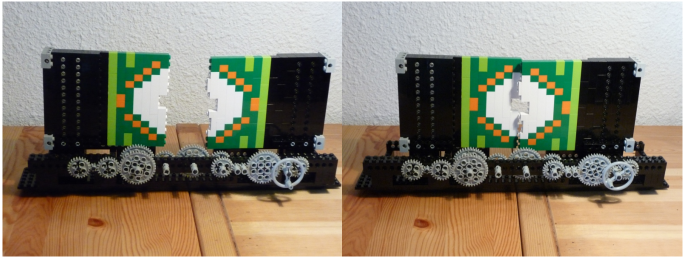

- Vertex locator modules sit 5 mm from the LHC beam
- Consists of two retractable halves (one shown)
- Modules formed of two sections – one on each velo half
- During beam injection, velo retracted to 35 mm for safety
- Closed once LHC beams are stable
- Moves every fill $\rightarrow$ align every fill

Real-time alignment of the Velo

- Alignment of velo based on minimising residuals between hits and reconstructed tracks
- Plot shows *x* and *y* translation between the two velo halves
- Tolerance of $\pm 2\ \mu m$ allowed without alignment update (empty markers)
- Updates may also be caused by other degrees of freedom
  - *e.g.* offsets or rotations within a velo half

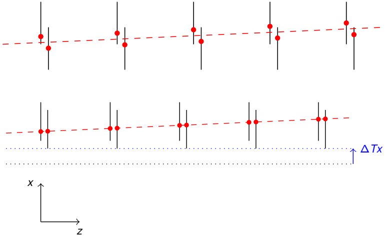

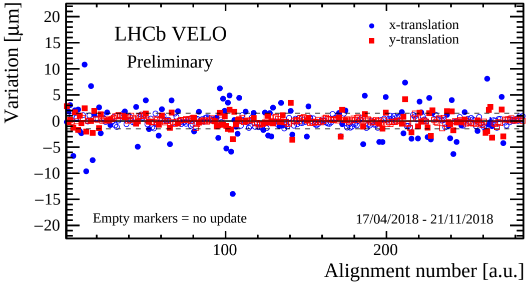

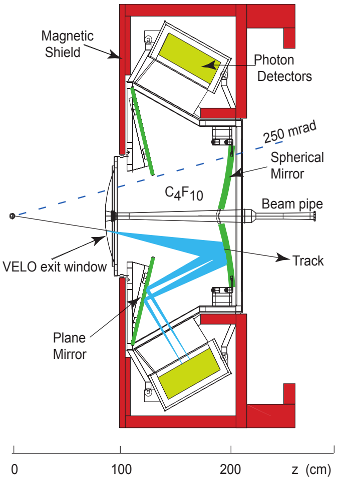

- RICH detectors provide particle ID information based on angle of Cherenkov radiation
- Index of refraction of the gas radiators sensitive to changes in temperature, pressure and composition
- These features are monitored but data-driven calibration also required
- Compare recorded and expected Cherenkov angles (bottom left)
- Alignment of mirrors also calibrated (bottom right)

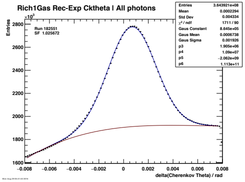

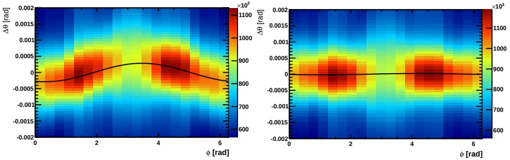

The turbo stream

- Save only parts of the event needed for offline analysis
- Multiple persistence levels
  - Only candidate ($\sim$ 7 kB)
  - Part of event ($\sim$ 16 kB)
  - Full event ($\sim$ 48 kB)
  - *cf.* Non-turbo ($\sim$ 69 kB)

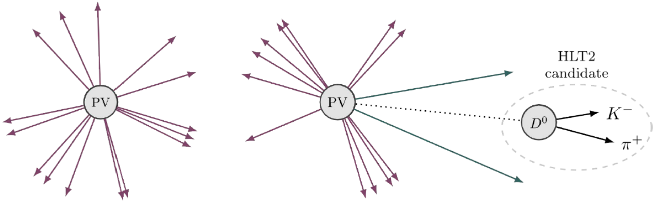

Used by many analyses, *e.g.*

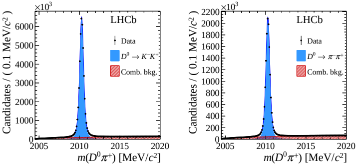

CP violation in charm decays

PRL **122** (2019) 211803

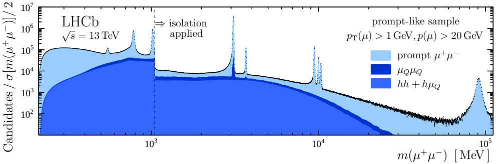

Search for dark photons decaying to dimuons

PRL **120** (2018) 061801

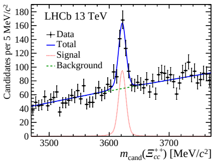

Observation of Ξ_{cc}^{++}

PRL **119** (2017) 112001

The LHCb detector: Run III upgrade

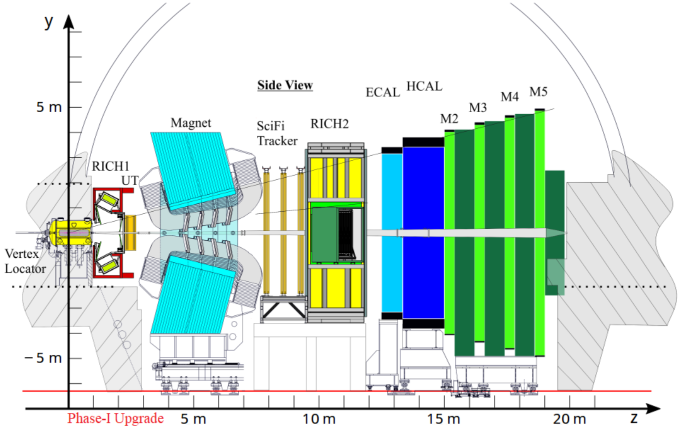

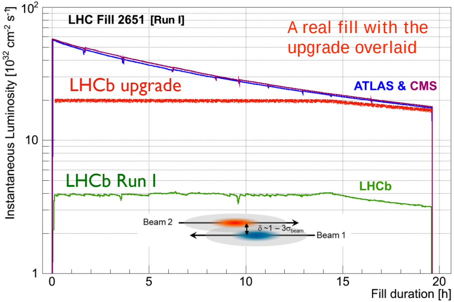

- Run at 5× higher luminosity
- Triggerless readout at 40 MHz
- New vertex locator
- New tracking (UT, SciFi)

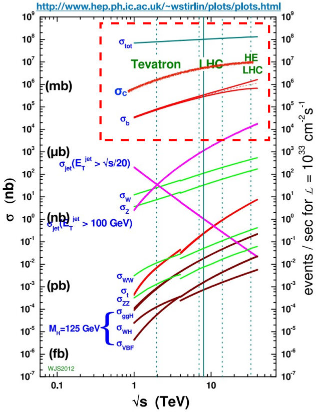

- At increased luminosity, charm (beauty) in 24 % (2 %) of bunch crossings
  - Cannot write out charm at 7 MHz
- Trigger must distinguish signal from less-interesting signal as well as from background
- No longer feasible to have first trigger based on calorimeters and muon detectors alone
- Need as much information about an event as soon as possible

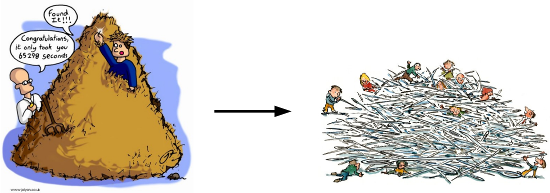

LHCb trigger in Run III

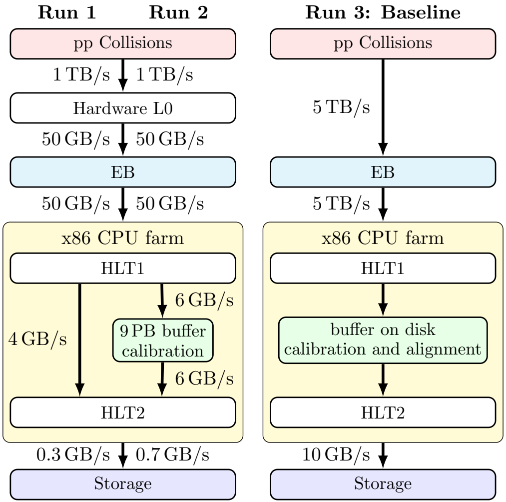

- Hardware trigger to be removed from Run III
- HLT1 software trigger must perform at 30× higher rate with 5× the pileup
- Buffer will reduce from $\mathcal{O}$(weeks)$\rightarrow$ $\mathcal{O}$(days)
- Significant increase in data transfer rates
- New trigger setup offers up to $\sim 10\times$ efficiency improvement for some physics channels

Run III baseline HLT1 performance

LHCB-TDR-017 LHCB-FIGURE-2019-002

- Significant progress made to optimise tracking algorithms
- $\sim 4\times$ improvement in throughput from vectorising, improvements to event model and optimisation

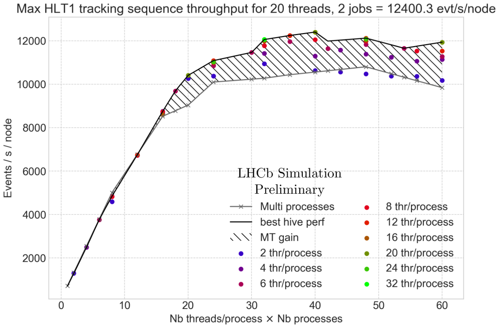

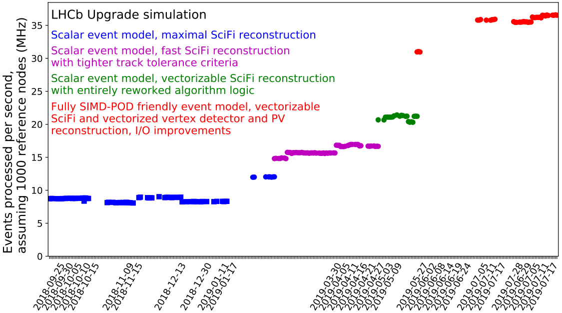

- multi-threaded processes offer gains over running more processes in parallel
- Optimal CPU architecture under investigation – new AMD architecture offers significant cost/benefit improvements

- Allows for a flexible and configurable sequence
- Physics performance of single-track and two-track selections studied
- Loose (L) and tight (T) versions of algorithms simulated with different p_T thresholds ($500 - 1000\,\mathrm{MeV}/c$)
- (Top) $\sim 1\,\mathrm{MHz}$ output rate achievable based on “minimum bias” simulation
- Two-track selection remains efficient
- Single-track selection still requires work

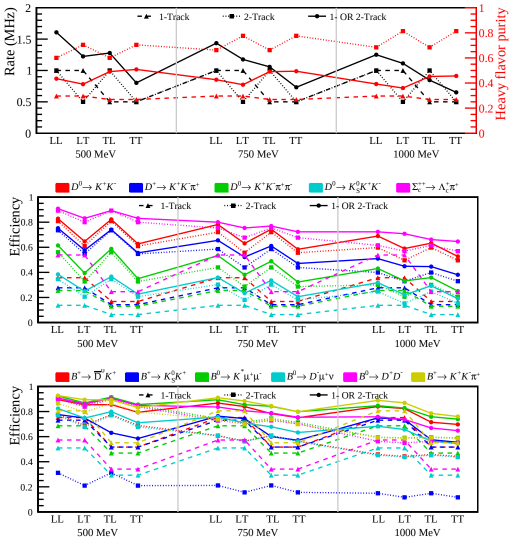

LHCb trigger in Run III

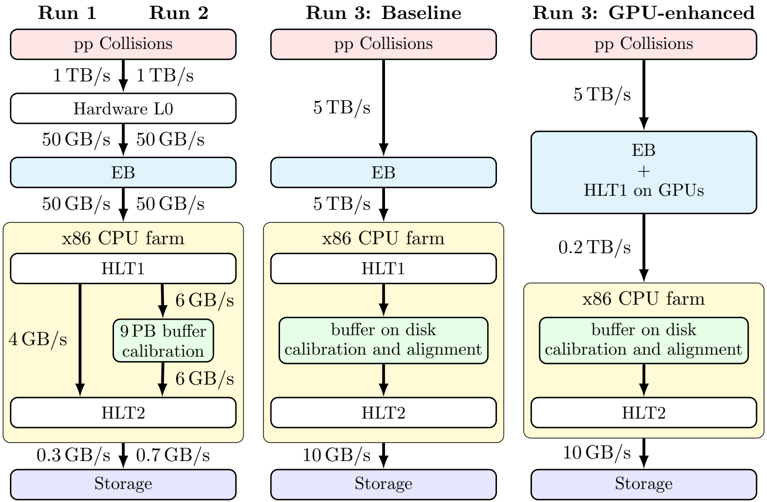

- Option to move to a GPU-based HLT1 with GPUs installed on the Event Builder servers
- Free up full CPU farm for HLT2 and save on networking between event builders and CPU farm
- Demonstrated technical feasibility
- Decision due in early 2020

Why GPUs?

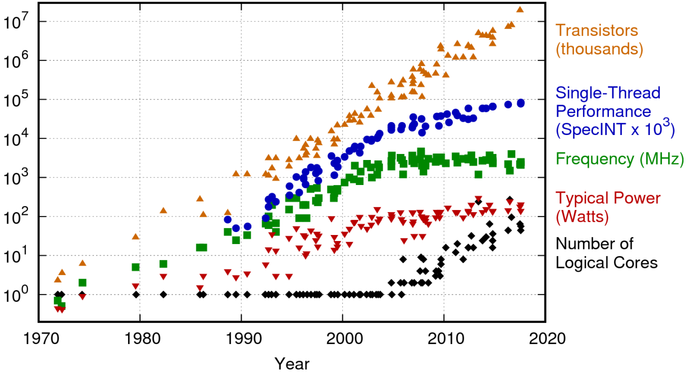

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

- Moore's law still holds but single-thread performance has levelled off
- Gains now to be made through parallelisation
- GPUs specialised for massively parallel operations (100s–1000s of cores)

- HLT1 involves decoding, clustering and track reconstruction for all tracking subdetectors
- Algorithms also perform Kalman filter and trigger selection
- All stages of the process may be parallelised

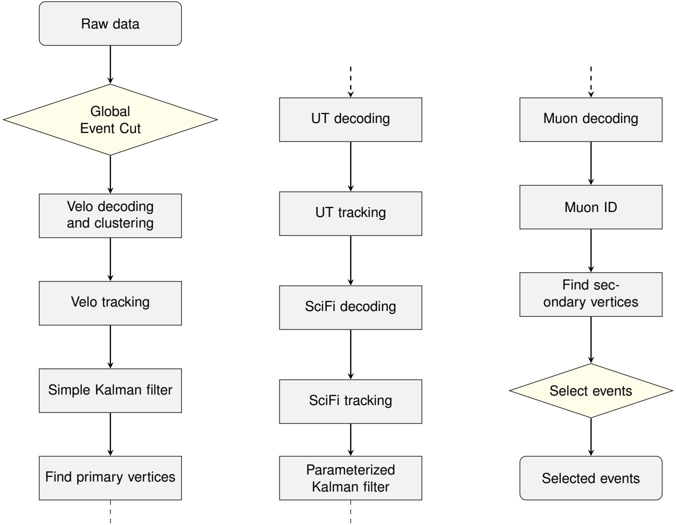

The Allen project

- Generic configurable framework for GPU-based execution of an algorithm sequence
  - Data passed to GPU device
  - All algorithms executed in order
  - Results passed back to the host
- Process thousands of events in a single sequence
  - Opportunity for massive parallelisation
- Configurable sequences at compile time
- Configurable algorithms at run time
- Custom GPU memory management – no dynamic allocation
- Built in validation and monitoring
- Cross-platform compatibility with CPU architectures
- Named for Frances E. Allen
- Implement HLT1 on GPUs

Allen selection ingredients

Primary vertices

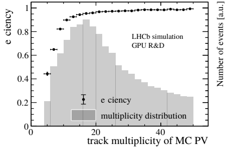

Tracks

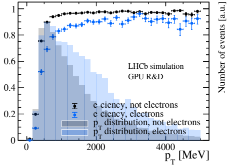

Selections

- One track
- Two tracks
- Single muon
- Two muons (displaced)
- Two muons (high-mass)

Secondary vertices

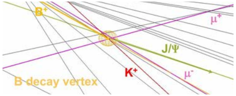

Momentum

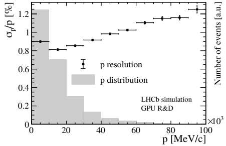

Impact parameter

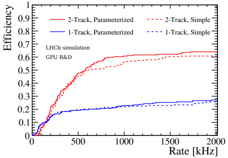

Muon ID

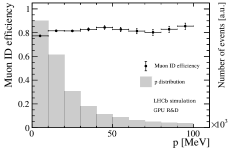

Allen performance

Trigger	Rate [kHz]
1-Track	249 ± 18
2-Track	663 ± 30
High-p_T muon	1 ± 1
Displaced dimuon	50 ± 8
High-mass dimuon	101 ± 12
Total	971 ± 36

- Total rate reduced 30 $\rightarrow$ 1 MHz
- Physics performance consistent with x86 baseline

Signal	GEC	TIS -OR- TOS	TOS	$\mathrm{GEC} \times \mathrm{TOS}$
$B^0 \rightarrow K^{*0}\mu^+\mu^-$	89 ± 2	85 ± 2	78 ± 3	69 ± 3
$B^0 \rightarrow K^{*0}e^+e^-$	84 ± 3	69 ± 4	62 ± 4	53 ± 3
$B^0_s \rightarrow \phi\phi$	83 ± 3	70 ± 3	65 ± 4	54 ± 3
$D^+_s \rightarrow K^+K^-\pi^+$	82 ± 4	62 ± 5	38 ± 5	32 ± 4
$Z \rightarrow \mu^+\mu^-$	78 ± 1	97 ± 1	97 ± 1	75 ± 1

GEC = global event cut, TIS = trigger independent of signal, TOS= trigger on signal

Allen throughput

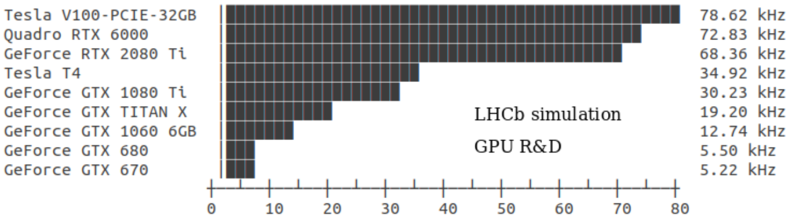

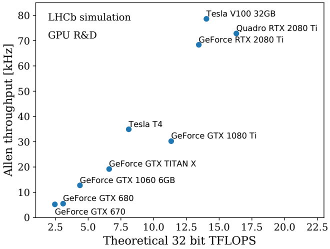

- Full HLT1 algorithm can be run on $\sim$ 500 current GPUs
- Buy GPUs instead of networking
- Performance scales with GPU so can expect more from 2021 GPUs
  - Not yet limited by Amdahl's law
  - Potential to perform more tasks within HLT1

Summary

- Real-time reconstruction and calibration a success story for LHCb in Run II
- Offline-quality reconstruction allowed for many trigger selections to be moved to the turbo stream
  - Selections can be based on offline-quality features
  - Smaller event size $\rightarrow$ higher event rate for same disk rate
  - Tradeoff – full raw event not saved $\rightarrow$ cannot rerun reconstruction offline
  - Already crucial for charm decays in Run II
- LHCb detector and DAQ upgrades for Run III
  - No hardware trigger
  - First-stage software trigger must perform track reconstruction at bunch-crossing rate
- Baseline x86 implementation of first-stage trigger significantly optimised to deal with higher throughput
- Allen project offers a GPU-implementation
  - Generic framework allows for configurable algorithm sequence
  - Feasibility for possible use in LHCb Run III already demonstrated