



Real-time trigger-level analysis and calibration at LHC experiments



Rosen Matev (CERN)

2 March 2020

COMPASS Front-End, Trigger and DAQ Workshop

What is real-time analysis?

- Online we have finite time to decide what data to keep (forever)
- Here, RTA means to efficiently reduce data online
- If we are reducing, what do we keep?
 - plots for a paper is probably too extreme, but may be useful for a preliminary result!

- Briefly show the real-time analysis landscape at the LHC
- Delve a bit deeper into LHCb
- Focus on the software part

Motivation

- Triggering is expensive; must fit within computing constraints

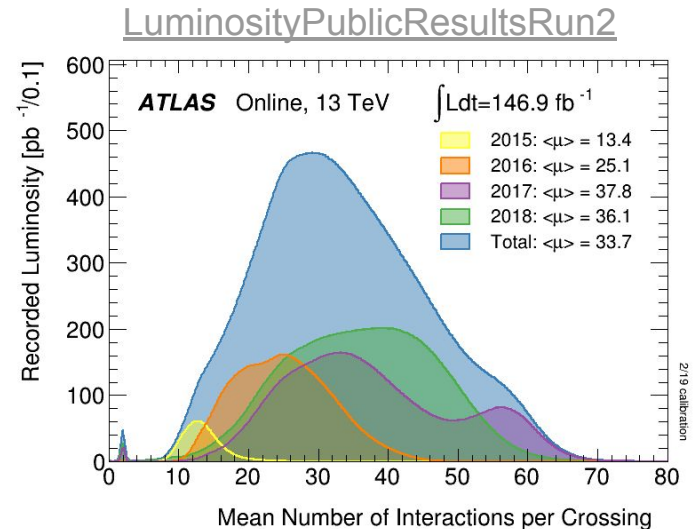
$$\text{Bandwidth [GB/s]} \propto \text{Accept rate [kHz]} \times \text{Event size [kB]}$$

- Want highest accept rate high to maximise $\epsilon_{\text{Sig.}}$ and reduce bias
 - Balanced against maximising $1 - \epsilon_{\text{Bkg.}}$
- Typically, can't do much to reduce the raw event size*; it's all or nothing!

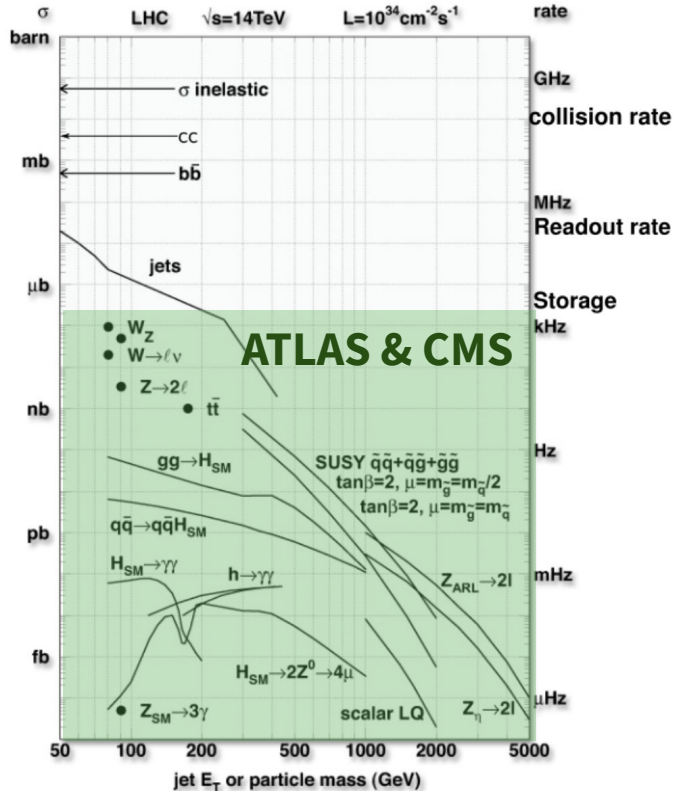
If event size is reduced, there's room for more physics!

Ever increasing pile-up

- Traditionally, we keep all raw data for events that contain signal
- Problem is **raw data bandwidth scales quadratically** with luminosity
 - more signal events, but much more bgr. data!
- The question is becoming less “Is this event/frame interesting?”
 - instead, “Which part of this event containing signal should we save?”
 - and how do we do it efficiently

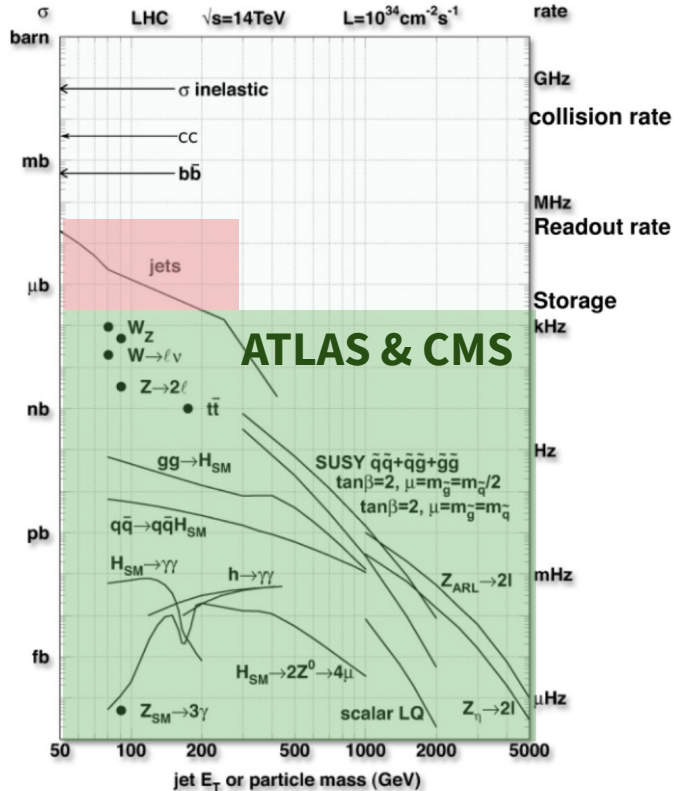


High mass physics



- 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 100 kHz is efficient with reasonably straightforward ET requirements

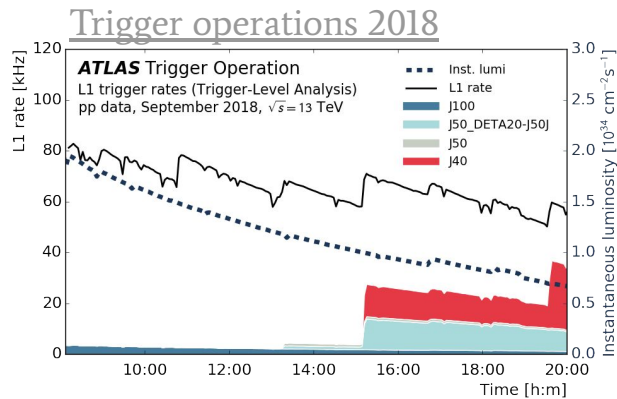
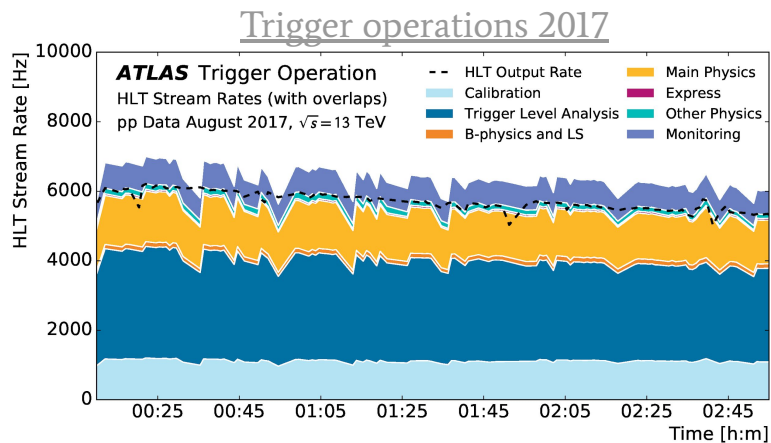
High mass physics



- 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 100 kHz is efficient with reasonably straightforward ET requirements
- What about that bit?

ATLAS “Trigger-Level Analysis”

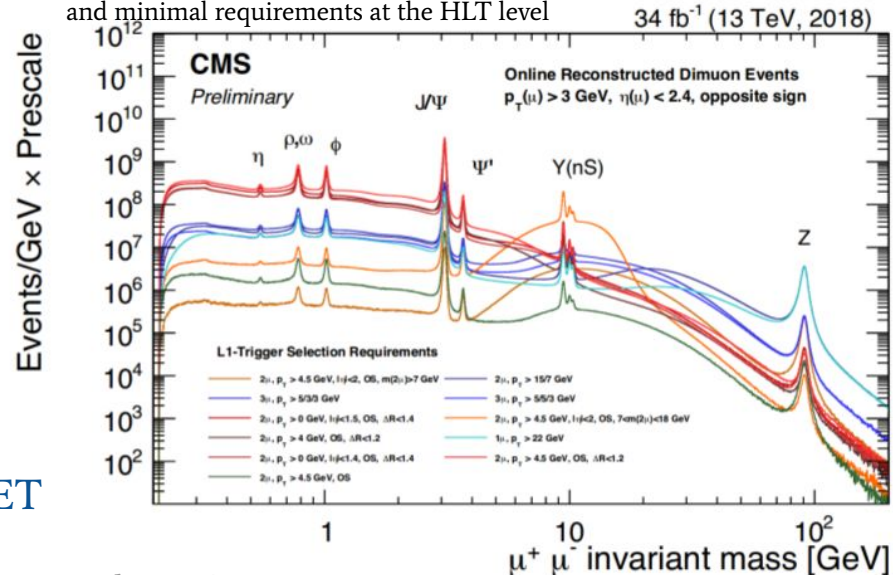
- Store only HLT jet 4-vectors and some summary info, e.g. N(constituents)
 - event is tiny, 0.5% of full size!
 - **all 3 kHz** of relevant triggered events saved
- Profit from available L1 rate during fill
 - save **up to 25 kHz** in 2018
- Limitations
 - Parts of the jet calibration “not quite real-time”
 - Coarse L1 algorithms \Rightarrow bad resolution
 - No tracking available
 - Ideas to improve for Run 3 and HL-LHC



CMS “Scouting”

- CaloScouting
 - save vertices, muons, calo jets, MET
 - L1-limited
- ParticleFlow Scouting
 - save vertices, PF muons, jets, cand, MET
 - CPU-limited
- Possible Run 3 extensions
 - PF scouting on all L1 events?
 - or restrict on L1 input to limit CPU
- HL-LHC: 40 MHz scouting
 - tracking in L1
 - streaming readout of detectors

dimuon events using a collection of L1 muon triggers, and minimal requirements at the HLT level



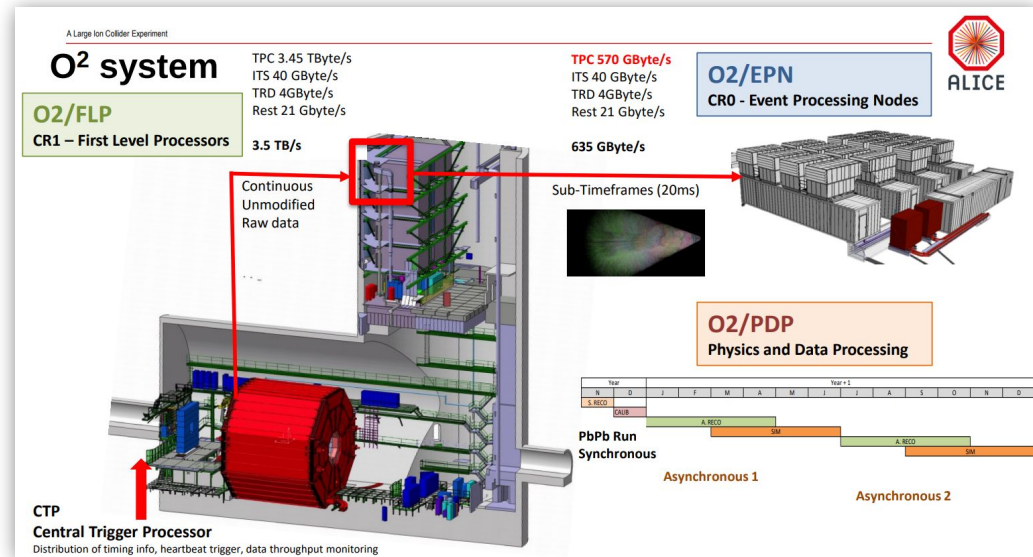
D.Sperka, HOW 2019

Stream	Rate (Hz)	Event Size	Bandwidth (MB/s)
PhysicsMuons	420	0.86 MB	360
PhysicsHadronsTaus	345	0.87 MB	300
ScoutingCaloMuon	4580	8.9 KB	40
ScoutingPF	1380	14.8 KB	20

Selected CMS stream rate, event size, and bandwidth at the beginning of LHC Fill 7334 (23 Oct. 2018, $L \approx 1.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$)

ALICE in Run 3

- Physics in ~all Pb-Pb events
 - LHC can deliver up to 50 kHz
- Zero suppression
 - non trivial, needs real-time calibration
- Compression with Huffman/ANS coding
 - save track parametrization + residuals
 - discard clusters not part of tracks
 - needs tracking
 - needs calibration! \Rightarrow feedback loop
- Big buffer that accumulates data
 - asynchronously processed 1-2 times in the following months of no beam period



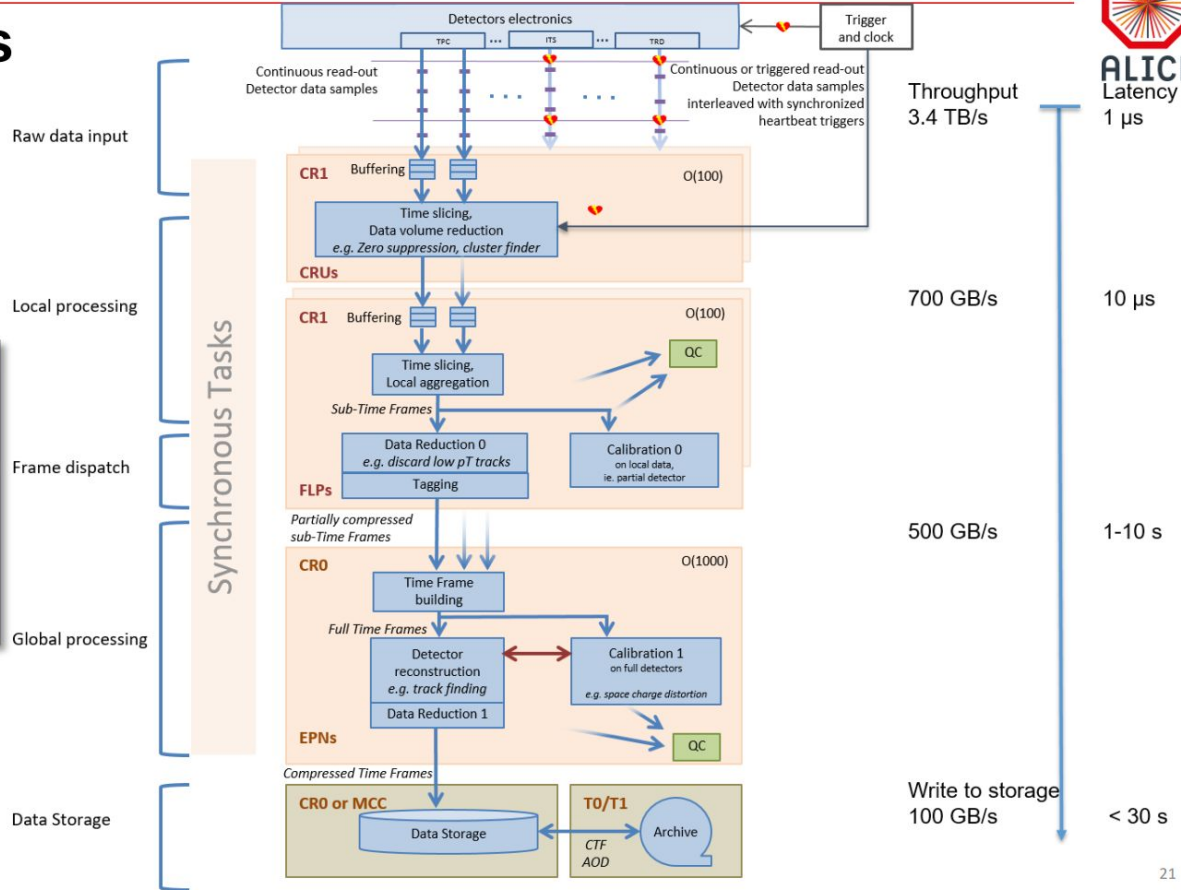
P. Vande Vyvre, EIC Streaming Readout, May 2019



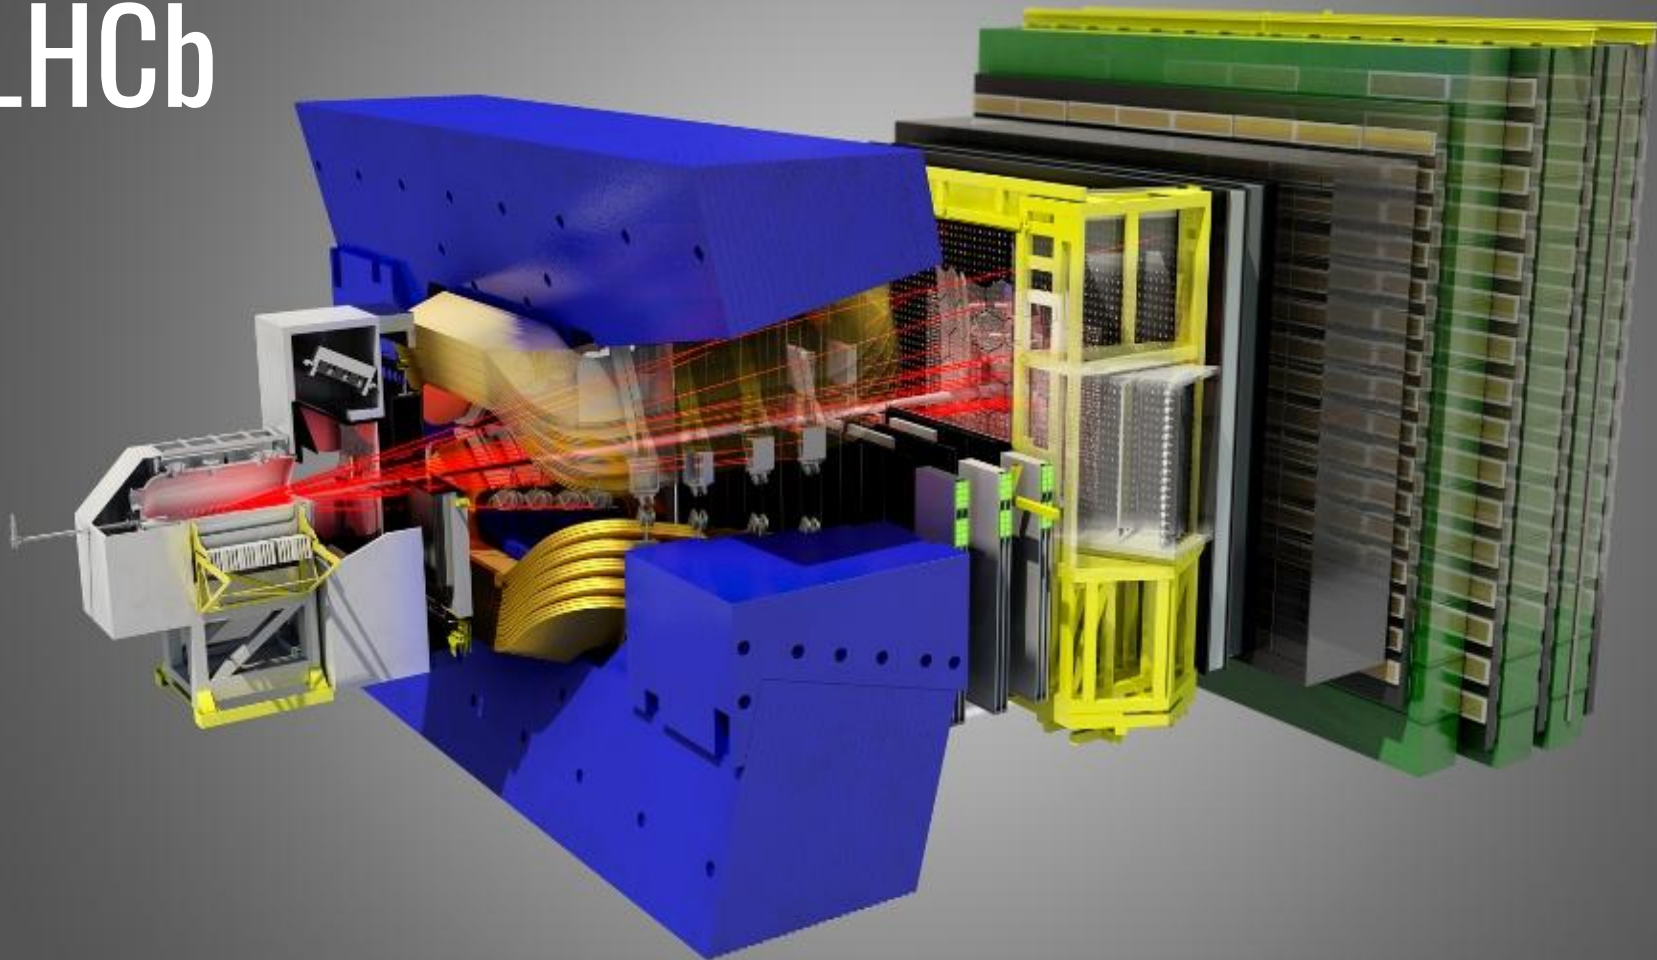
ALICE
Latency
1 μ s

Synchronous Processing

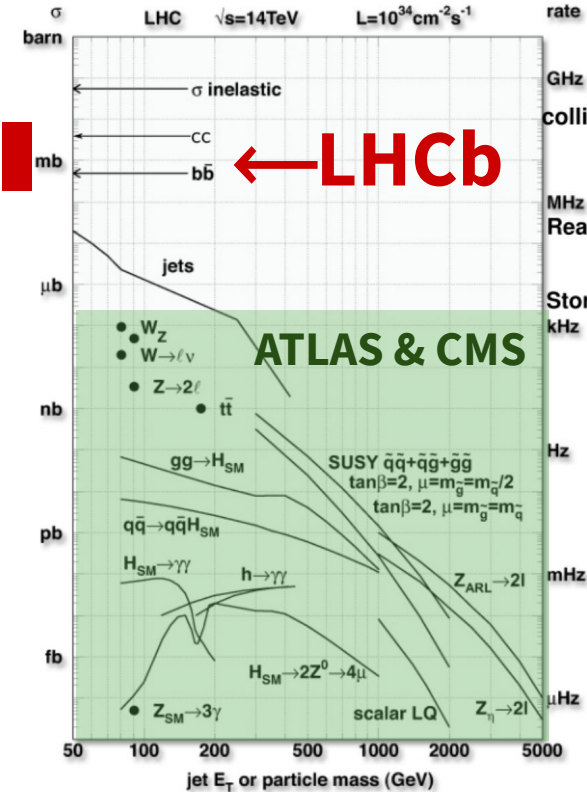
- Set of operations to be performed before data storage, i.e. all the operations to reduce the data volume
- Fast and simple calibration
- Reconstruction
- Production of the immutable Compressed Time Frames



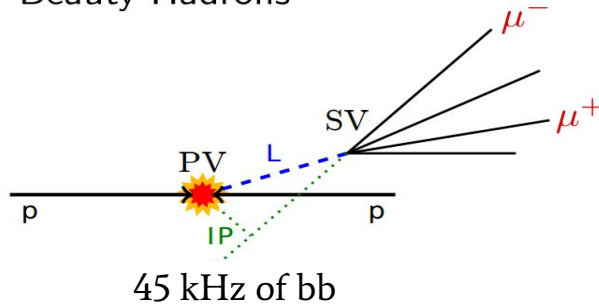
LHCb



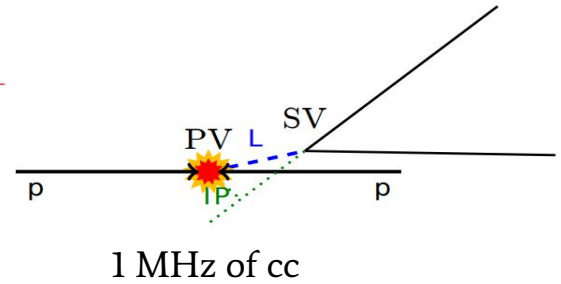
LHCb Trigger in Run 2



Beauty Hadrons

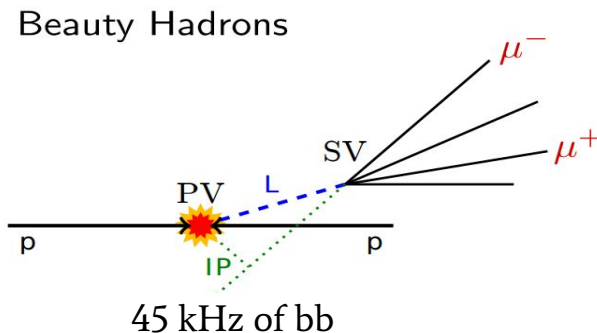
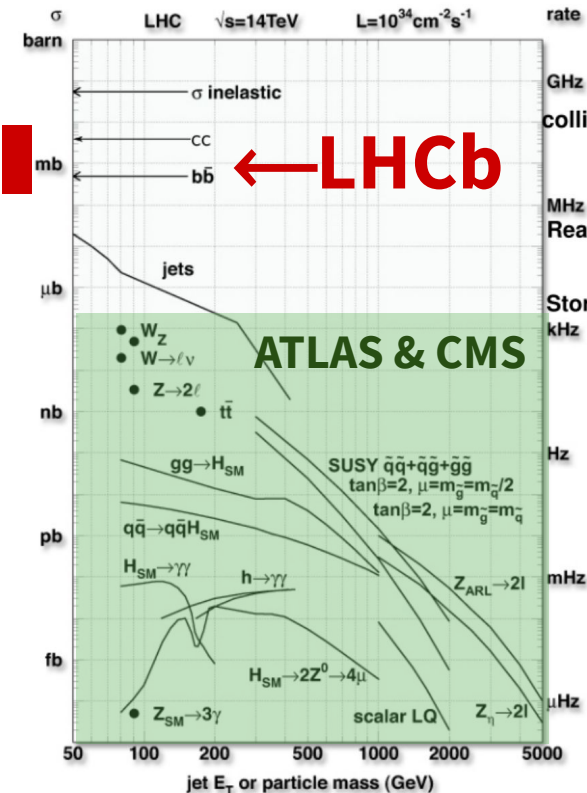


Charm Hadrons

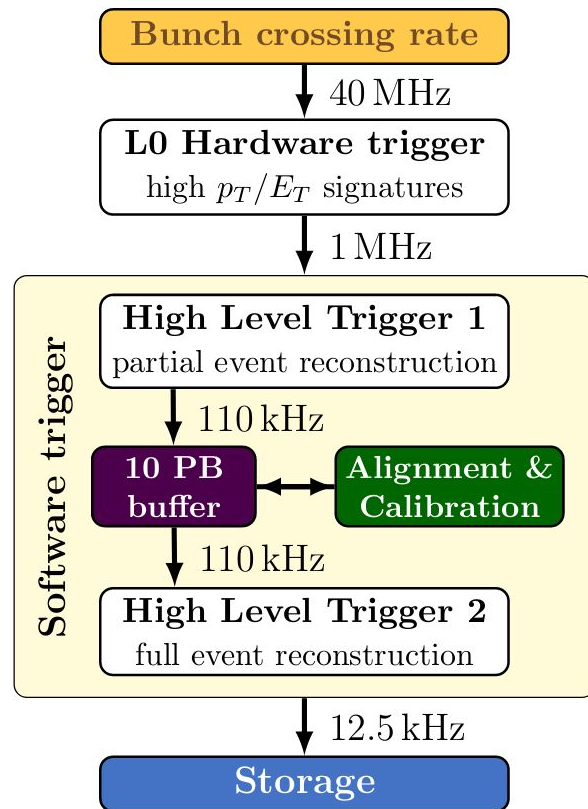


1 MHz readout is needed to stay efficient for beauty signals

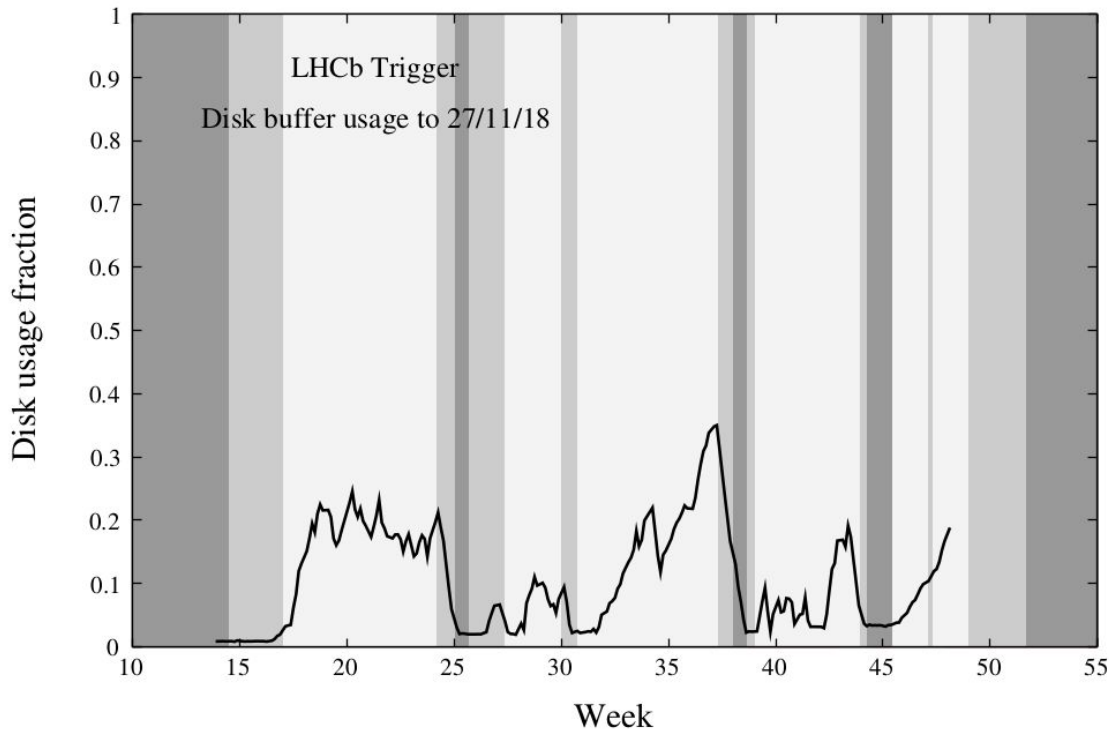
LHCb Trigger in Run 2



1 MHz readout is needed to stay efficient for beauty signals



Disk buffer



11 PB of disk capacity

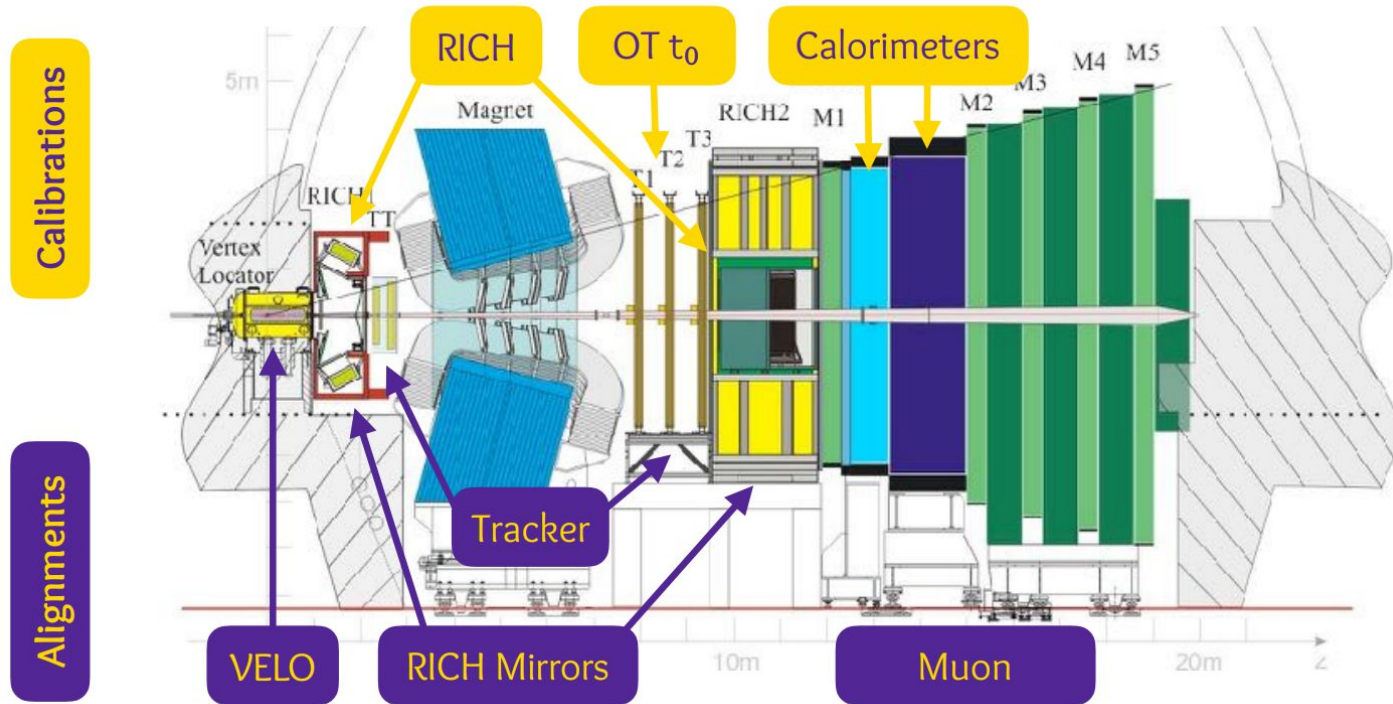
HLT1 writes at 110 kHz in fill

HLT2 processes at 30/90 kHz
in/out-fill

Effectively **doubles the
trigger CPU capacity.**

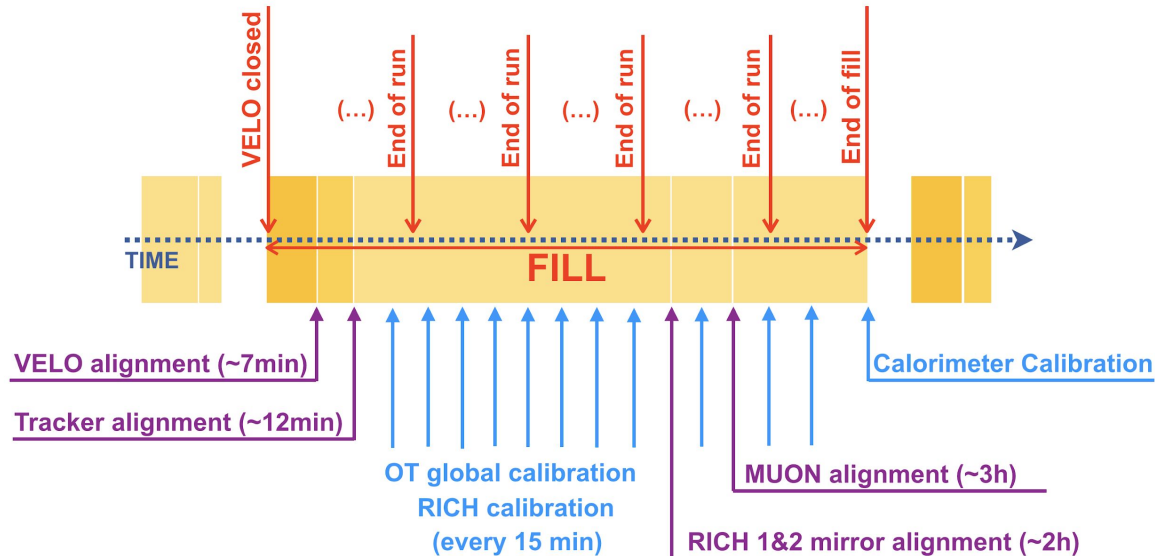
Full event reconstruction
becomes feasible.

Real-time alignment and calibration



Real-time alignment and calibration

- Data collection & analysis fully automated
- New constants automatically applied
- Shift crew verifies updates

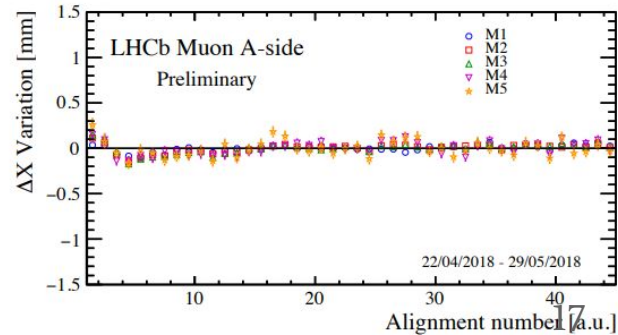
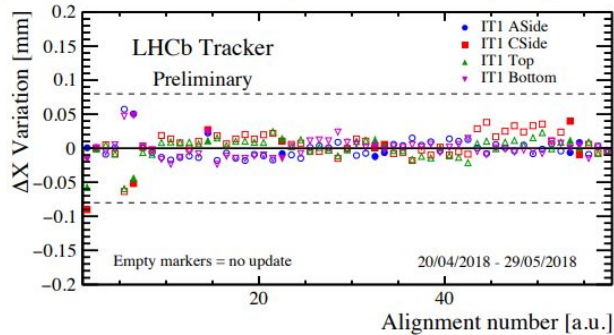
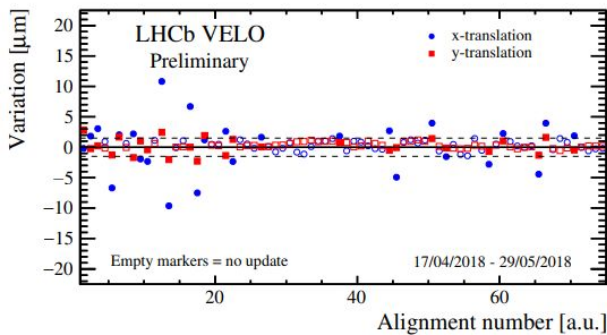


((~7min),(~12min),(~3h),(~2h)) - time needed for both data accumulation and running the task

Tracking alignments: minimise the global χ^2

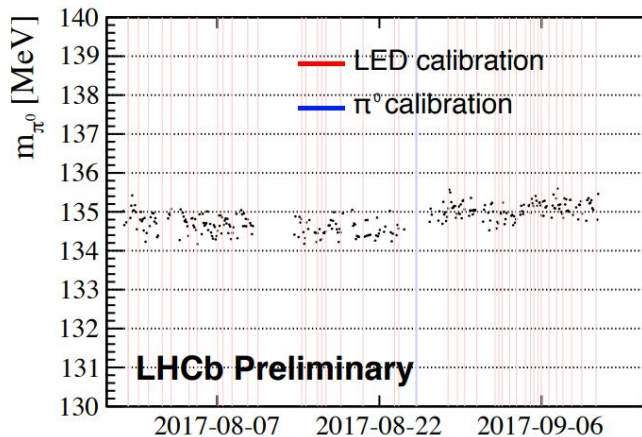
- Velo:
 - Sample collected ~immediately, alignment takes ~2 minutes
 - Frequent updates due to movement at the beginning of each fill
- Tracker:
 - Sample collected in ~immediately, alignment takes ~7 minutes
 - Updates mostly expected after magnet polarity changes
- Muon:
 - Sample collected in ~3 hours, alignment takes ~7 minutes
 - No movement expected except after physical intervention

Changes of alignment constants each time the alignment is ran with solid markers represent the alignments that triggered and update. The horizontal dashed lines represent the minimum change required to trigger an update.

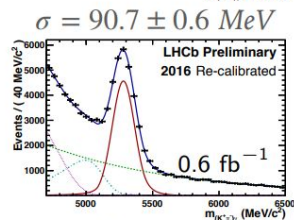
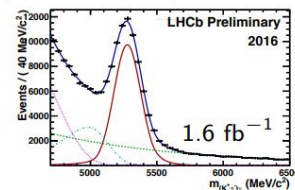
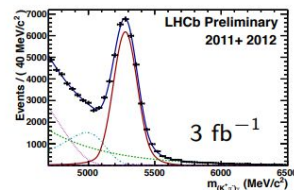


Calorimeter calibration

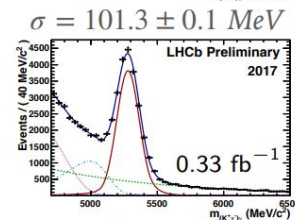
- Required to counteract changes and ageing of the detector material
- **Relative calibration:** end of every fill
 - Compare LED monitoring system to a reference and update HV
 - Reference updated after each absolute calibration
- **Absolute calibration:** ~once a month
 - HCAL: Caesium scan performed during technical stops
 - ECAL: Use 300M randomly selected events to fit mass in each cell



Effect of calibration in $B_d^0 \rightarrow (K^* \rightarrow K^+\pi^-) \gamma$ decays



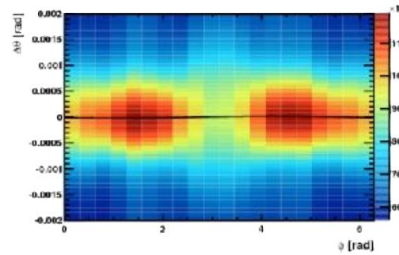
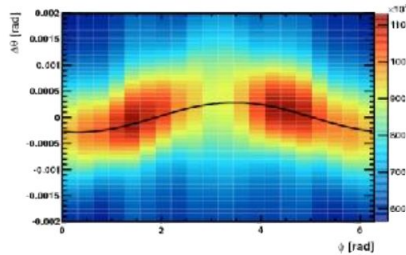
$\sigma = 90.7 \pm 0.6$ MeV



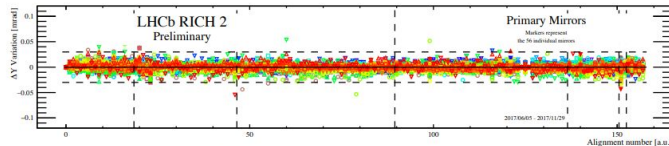
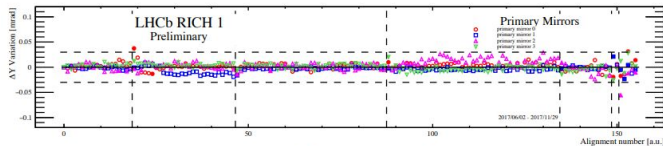
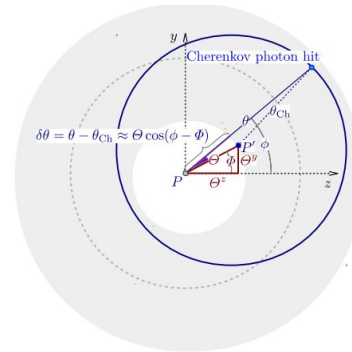
$\sigma = 101.3 \pm 0.1$ MeV

Alignment of RICH detectors

- Primary and secondary mirrors need to be aligned (110 mirror pairs)
- Fit the variation of Cherenkov angle $\Delta\theta$ as a function of polar angle
- Ran every fill, parameters typically change with magnet polarity flips
- Takes ~ 2 hours to collect data and ~ 20 minutes to run procedure



Cherenkov angle vs phi for misaligned (left) and correctly aligned (right) mirrors

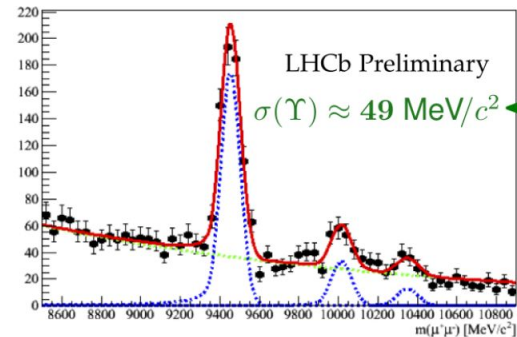
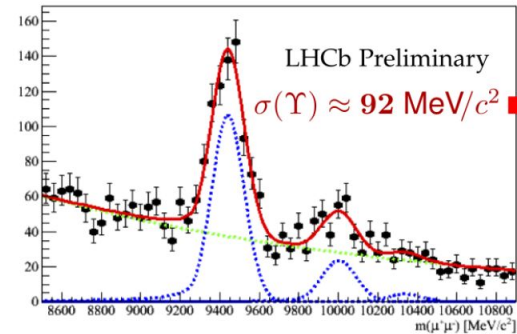


Difference in alignment constants of the primary mirrors of RICH 1 and RICH 2

What this buys us

- Offline-equivalent, fully aligned and calibrated physics objects in HLT2
- Can include offline selections in the trigger with no associated systematic effects
- Offline reprocessing of the raw data is not necessary to recover information

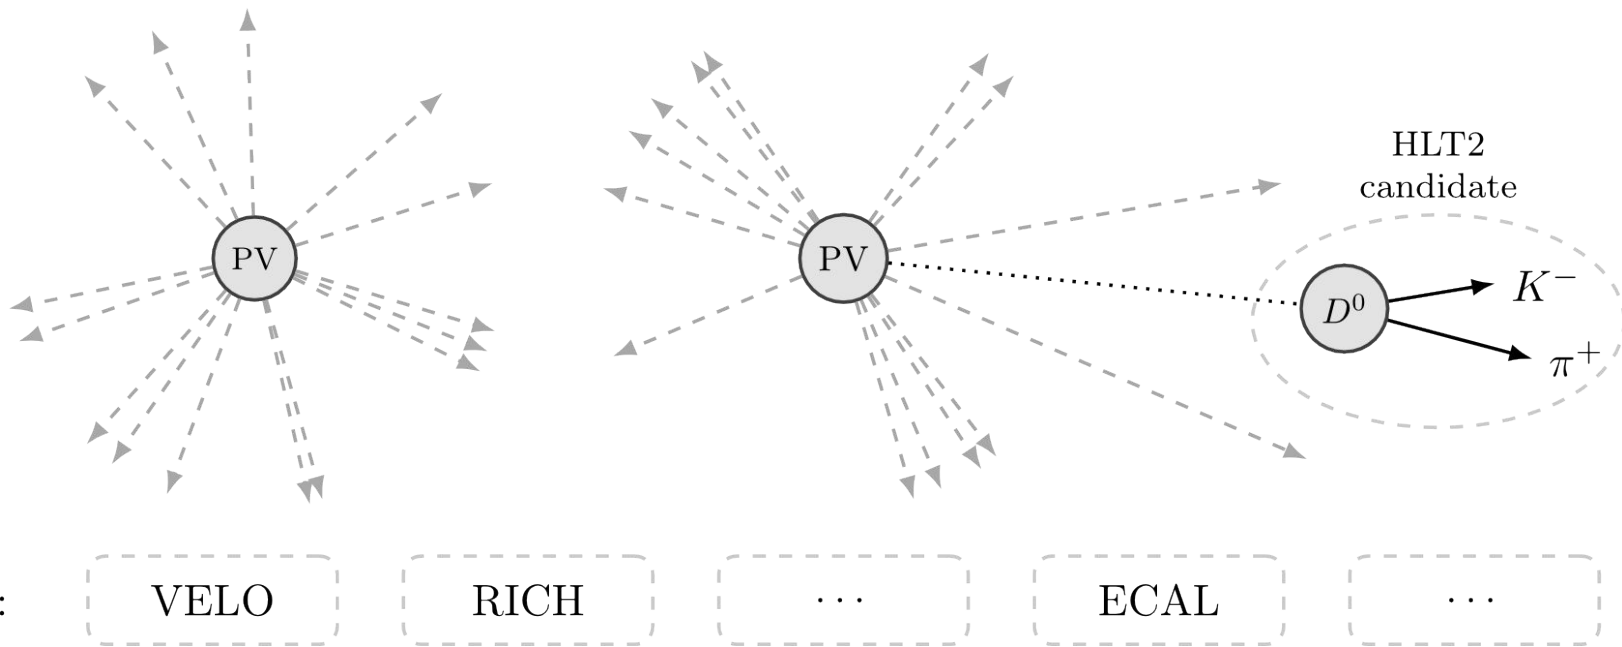
Real-time analysis with offline-quality physics objects



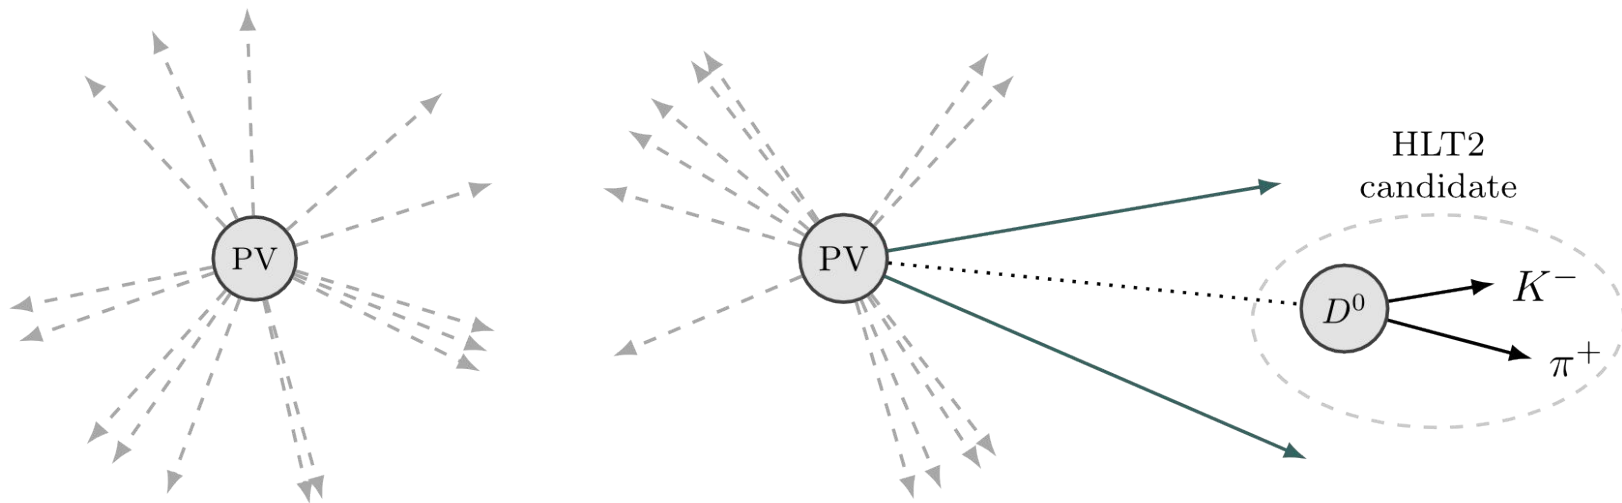
“Turbo” persistence model

- Persist objects from HLT2 directly, analyse only these offline
- Each trigger selection has complete control over what objects are saved
- Evolved over time to meet increasing needs

Persistence granularity



Persistence granularity



Raw banks:

VELO

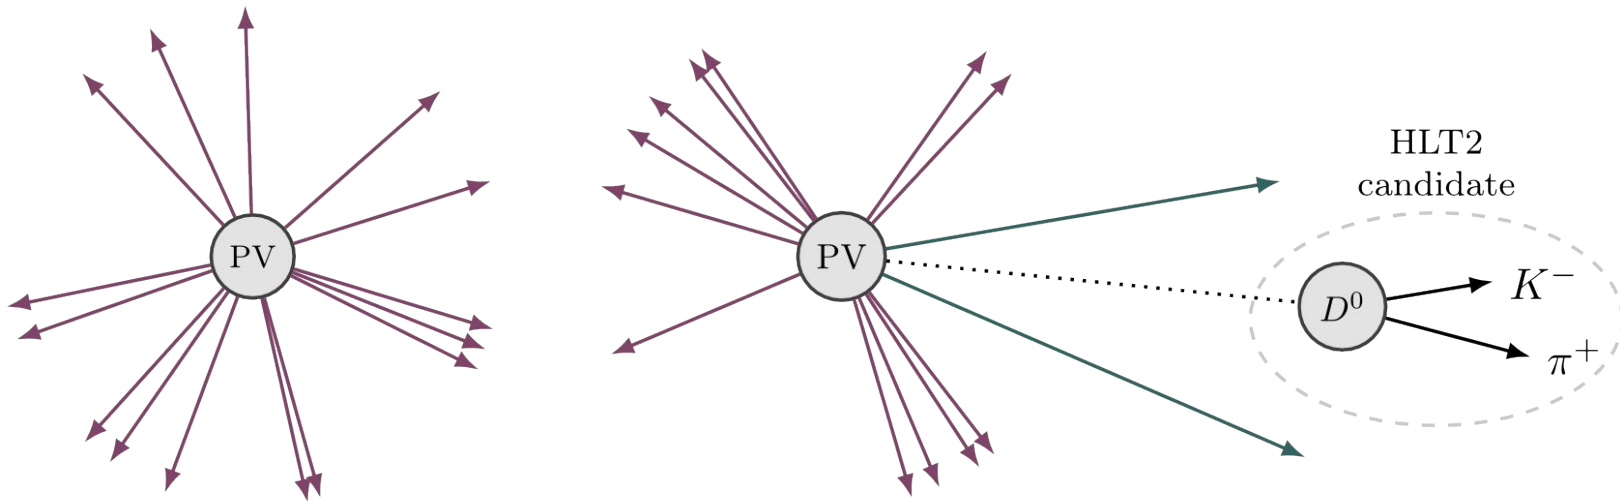
RICH

...

ECAL

...

Persistence granularity



Raw banks:

VELO

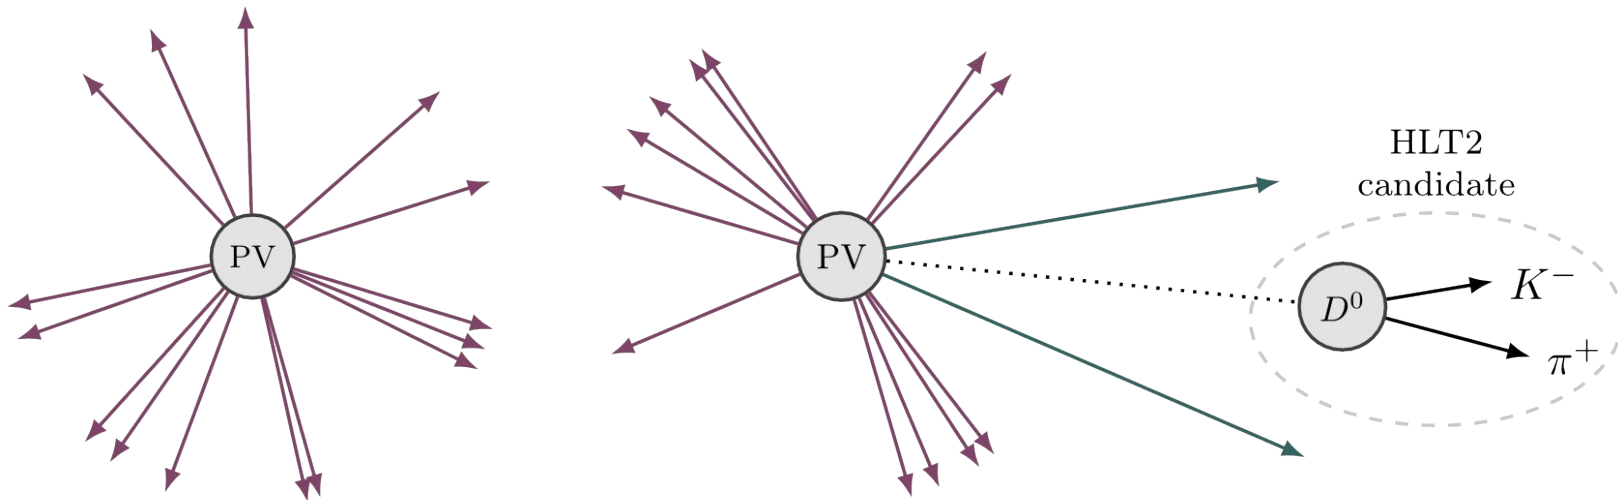
RICH

...

ECAL

...

Persistence granularity



Raw banks:

VELO

RICH

...

ECAL

...

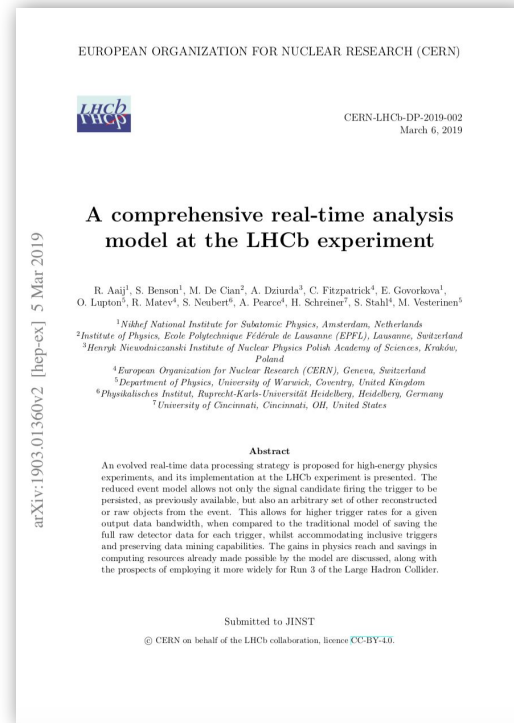
Rewards

Much smaller average event size

⇒ more physics within our resources

Persistence method	Average event size (kB)
Turbo	7
Selective persistence	16
Complete persistence	48
Raw event	69

Accounted for around 25% of the trigger rate in Run 2.
For 10% of the bandwidth!

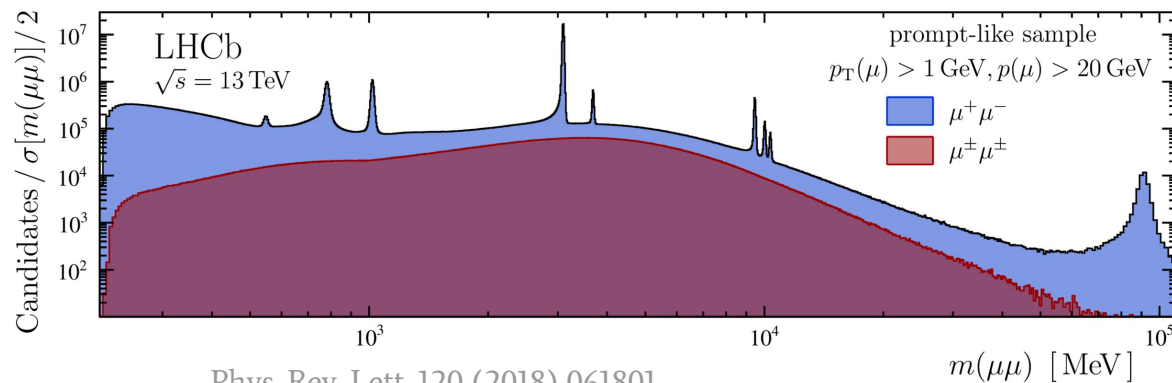


Looking back

- Must overcome fear of losing information
- There's always room for improvement
 - Selective persistence allowed us to reduce Turbo bandwidth, then added new inclusive charm baryon lines
- Must support users in transitioning to any new features

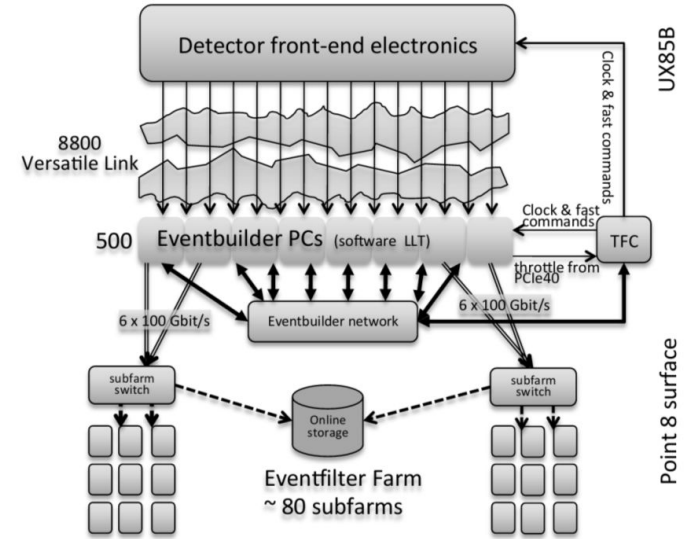
Looking back

- Must overcome fear of losing information
- There's always room for improvement
 - Selective persistence allowed us to reduce Turbo bandwidth, then added new inclusive charm baryon lines
- Must support users in transitioning to any new features



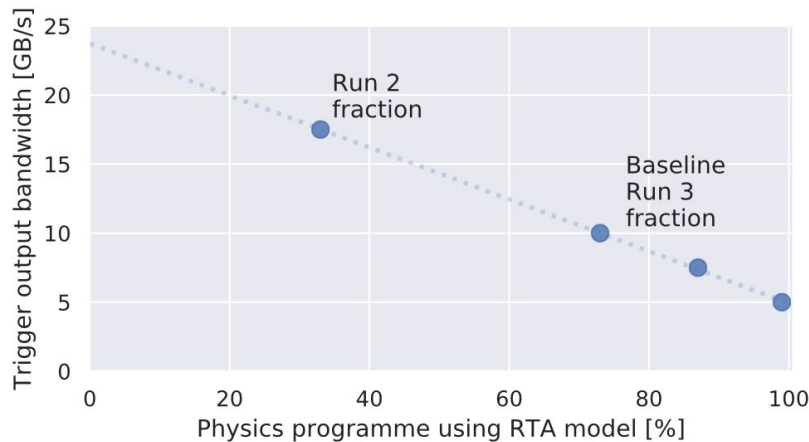
Looking forward

- Run 3 luminosity increases 5x
- Triggerless readout, full software trigger
 - Removal of hardware trigger increases efficiency of hadronic signals >2x
 - but 4 TB/s into HLT1
- Huge increase in signal rate!



Challenges

Run 3 physics programme is bandwidth-constrained like charm was in Run 2



- Turbo fraction must increase: baseline is 70%
- Must migrate some inclusive triggers to the RTA model
- What if we cannot achieve online/offline parity in HLT2?

Takeaway

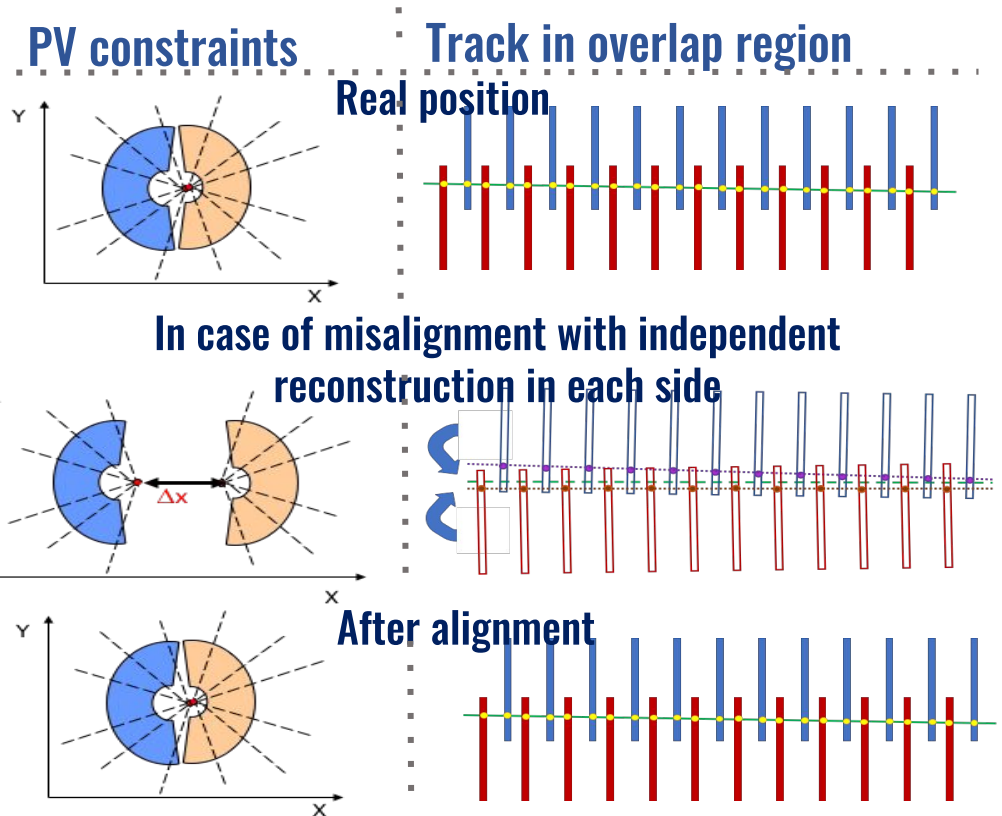
- Going “triggerless” helps if you have the processing power
- Align and calibrate your detector online
 - helps with improving efficiency and reducing background
- Squeeze the offline A&C and reconstruction online
 - you are sure to have the best physics objects for analysis
 - you can be much tighter on selections
- After that, it’s “easy”
 - just throw away what is not necessary from the events
 - still, make sure you’ve convinced yourself first it’s ok
 - still, make sure your QA/QC is solid as there is no going back

References

- D. Rorh, Real-time analysis model in ALICE, HOW 2019
- W. Kalderon, Real-time analysis model in ATLAS, HOW 2019
- D. Sperka, Real-time analysis model in CMS, HOW 2019
- ATLAS collaboration, Trigger-object Level Analysis with the ATLAS detector at the Large Hadron Collider: summary and perspectives, ATL-DAQ-PUB-2017-003 (2017)
- CMS collaboration, Data Parking and Data Scouting at the CMS Experiment, CMS-DP-2012-022 (2012)
- R. Aaij et al., Performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC, arXiv:1812.10790
- R. Aaij et al., A comprehensive real-time analysis model at the LHCb experiment, arXiv:1903.01360
- LHCb Collaboration, Computing Model of the Upgrade LHCb experiment, LHCb-TDR-018

Example: VELO alignment

- VELO centred around the beam for each fill
 - Resolver X, Y position accuracy of $10\ \mu\text{m}$
- Kalman filter based method, minimizing the track hit residuals with PV constraints
- Automatic alignment of VELO halves in less than 5 minutes

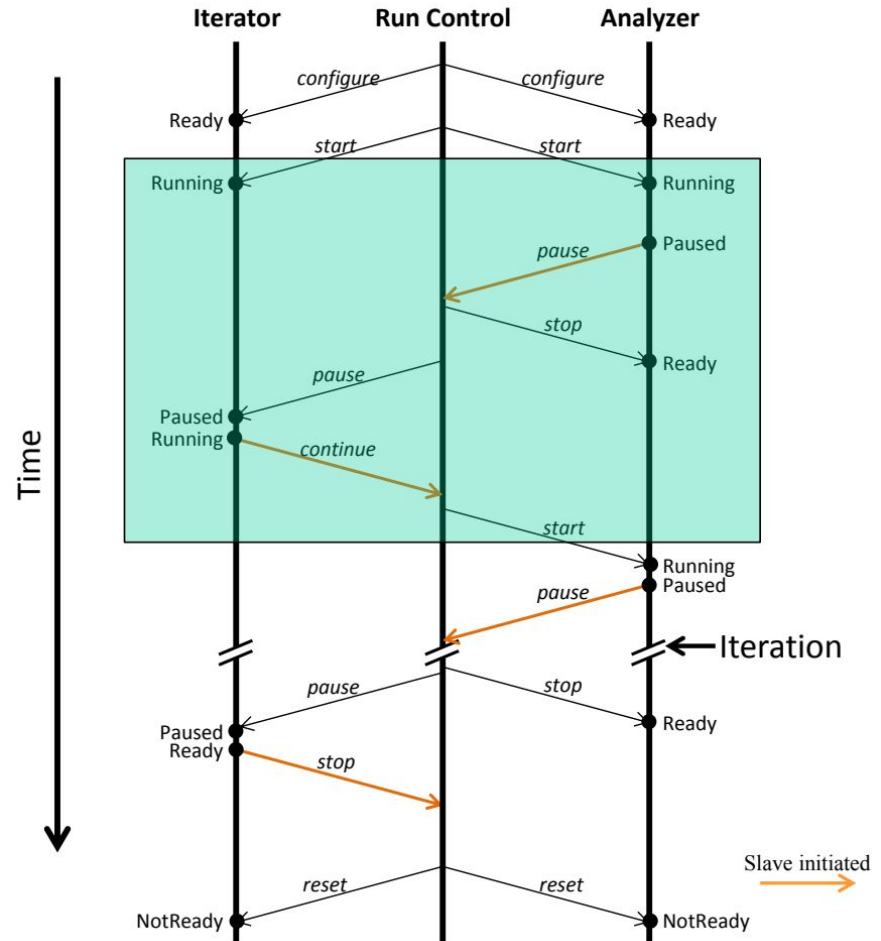


Automated tasks

Task	Update	Sample	Data collection	Duration	When?
Velo alignment	Automatic	50k minbias + beamgas	< 1 min	2 min	Every fill
Tracker alignment	Automatic	100k $D^0 \rightarrow K \pi$	< 1 min	7 min	Every fill
RICH mirror alignment	Automatic	3M good tracks	2 h	20 min	Every fill
Muon alignment	Expert	250k $J/\psi \rightarrow \mu^+ \mu^-$	3 h	7 min	Every fill
OT t_0 calibration	Automatic	Some minbias	15 min	O(min)	Every run
RICH Calibration	Automatic	Good tracks	15 min	O(min)	Every run
Relative CALO calibration	Automatic	LED monitoring system	N/A	2 min	Between fills
Absolute HCAL calibration	Expert	Caesium scan	N/A	2 hours	Technical stops
Absolute ECAL calibration	Automatic	300M minbias	O(4 weeks)	2 hours	When sample ready

Alignment FSM

- Each online alignment and calibration task is controlled by the same finite state machine
- One process of the analyser task runs on each of the ~1600 nodes in the trigger farm
- Overview of sequence:
 1. Iterator writes conditions in XML
 2. Each analyser reads these conditions and reconstructs events to produce a binary file “alignsummarydata” (ASD)
 3. Iterator combines the ASDs to compute the new conditions constants and writes these to XML
 4. Steps 2 & 3 repeat until the procedure converges. The new constants are then copied to the trigger area.



Turbo internals

