# The upgrade of the LHCb trigger for Run III

•••

Rosen Matev, CERN
on behalf of the LHCb collaboration

ACAT, 21-25 August 2017
University of Washington, Seattle

# The LHCb experiment

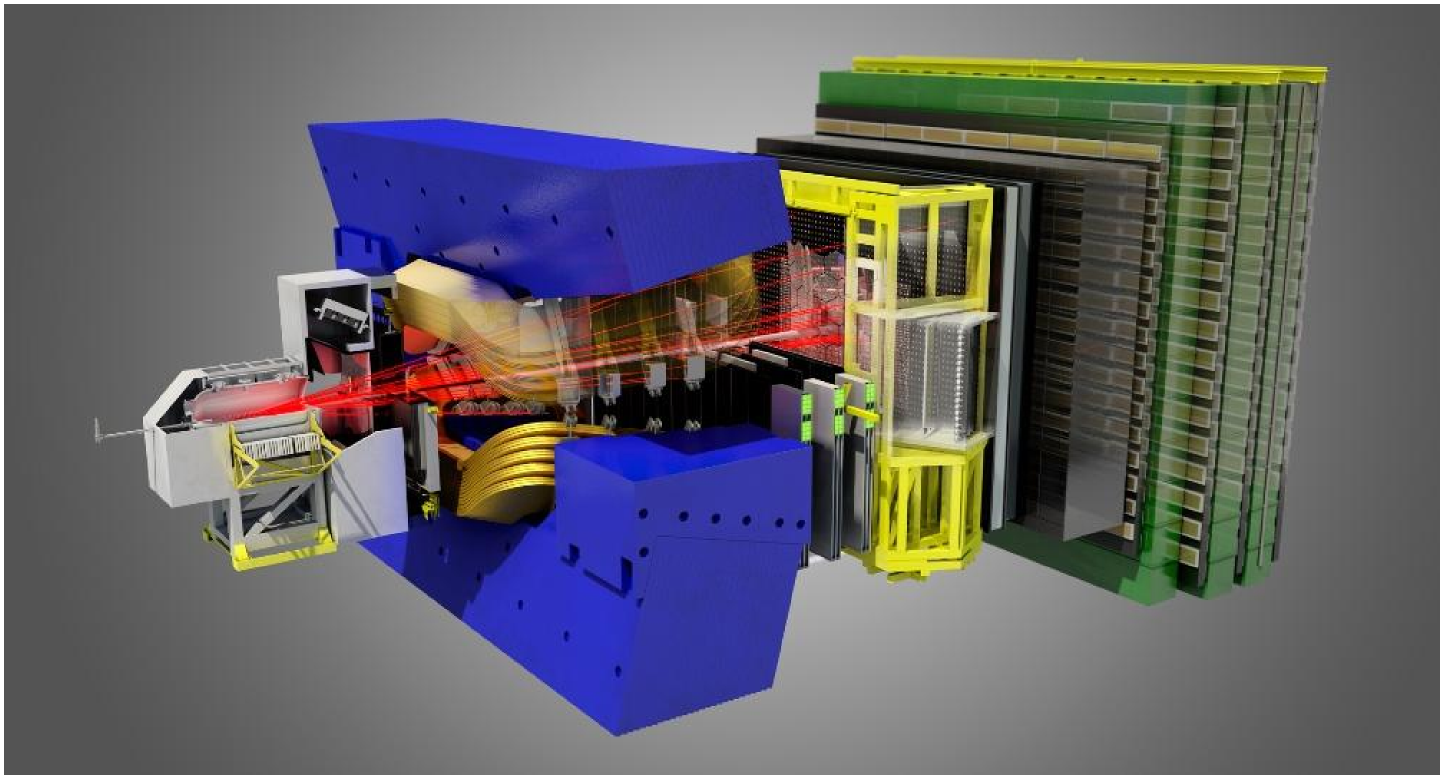

JINST 3 (2008) S08005
IJMPA 30 (2015) 1530022

**~45 kHz** bb pairs and **~1 MHz** cc pairs

at 13 TeV and L = $4 \times 10^{32}$ cm$^{-2}$s$^{-1}$

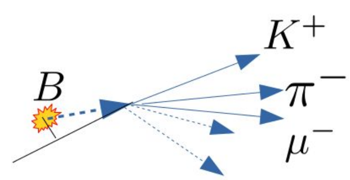


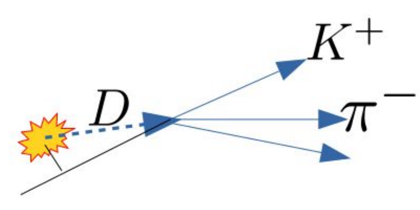

# LHCb upgrade

- Precision of many physics measurements at LHCb will be statistically limited at the end of Run II
- Upgrade to cope with 5× more luminosity ($L = 2\times10^{33}\ cm^{-2}s^{-1}$)
  - Sub-detectors
  - Triggerless read-out
  - Software trigger

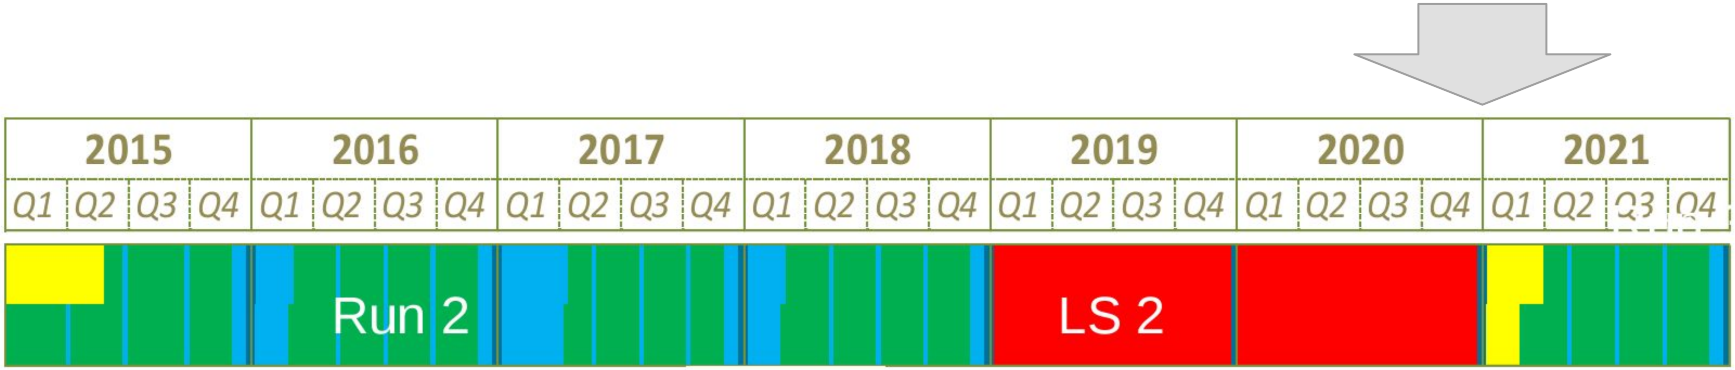

# LHCb upgrade

See more on read-out in Tommaso Colombo's talk

- Move from 1 to 40 MHz read-out
  - Requires upgrade of many detectors and front-end electronics
- Fully-software trigger with no further offline processing
  - Best possible quality reconstruction
  - Real-time alignment and calibration
  - Persist only the needed high level objects
- All of this already in place today!
  - although in much less challenging conditions

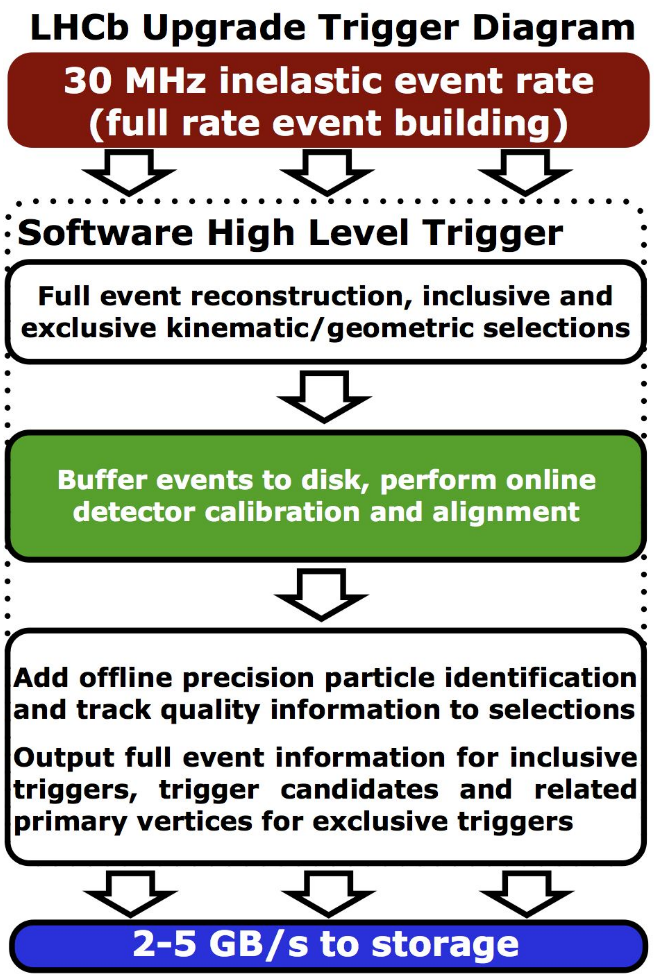

# The MHz signal era

- A paradigm shift from Run II
  - 24% (2%) of events contain a reconstructible charm (beauty) hadron
  - 80 (27) GB/s worth of events usable for analysis
  - We can only afford storing 2-10 GB/s offline
- Not only separate signal and background decay topologies
  - but effectively separate signal decays from other signal decays


**Triggers today**

**Triggers in the future**

# The MHz signal era

- A paradigm shift from Run II
  - 24% (2%) of events contain a reconstructible charm (beauty) hadron
  - 80 (27) GB/s worth of events usable for analysis
  - We can only afford storing 2-10 GB/s offline
- Not only separate signal and background decay topologies
  - but effectively separate signal decays from other signal decays
- Exclusive selections will be the standard
  - Some high rate channels cannot be saved fully even with 100% purity
  - Retain some inclusive triggers for breadth of the physics programme
  - Should be almost the offline selections - aim for high purity and efficiency
  - More sensitivity to detector performance effects (e.g. asymmetries)
  - Real-time alignment and calibration will be crucial

CERN-LHCC-2014-016 LHCb-PUB-2017-005

# Reconstruction

See more on Run II in
Mike Sokoloff's poster

- Take advantage of the Run II trigger strategy
  - Perform a fast reconstruction and selection (HLT1)
    - mainly tracking, vertex finding and inclusive selections
    - reduce bandwidth to a manageable level
  - Buffer events on disk and perform detector alignment and calibration
  - Perform the full reconstruction and selection (HLT2)
    - ultimate track quality and particle identification

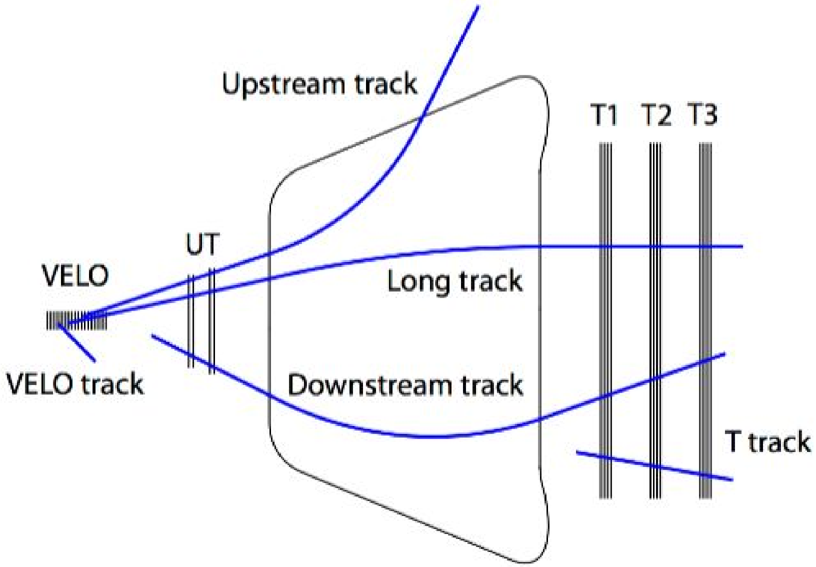


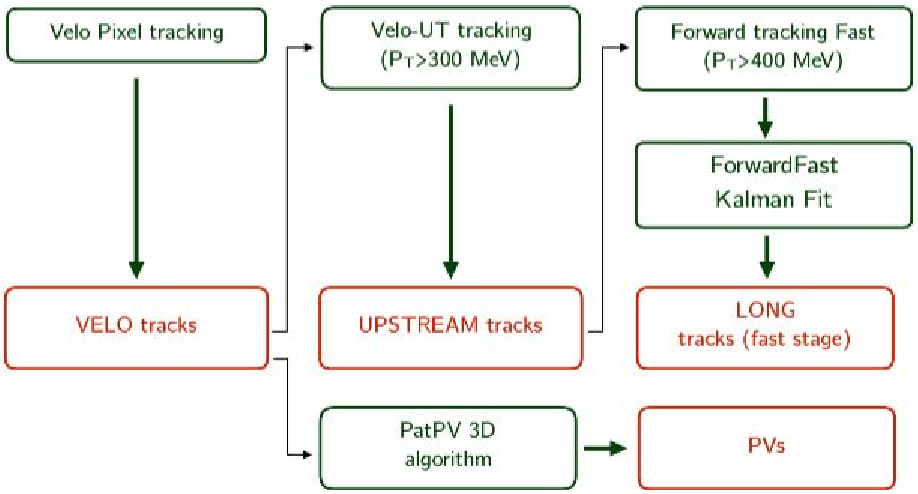


LHCb-PUB-2017-005

# Fast reconstruction stage

- Several algorithms to reconstruct tracks fully or partially traversing the detector, and find primary vertices
- The event topology is more complex at the upgrade conditions
  - 3-4 times more primary vertices and 2-3 times higher track multiplicity
  - Challenging to keep good physics performance and to lower processing time

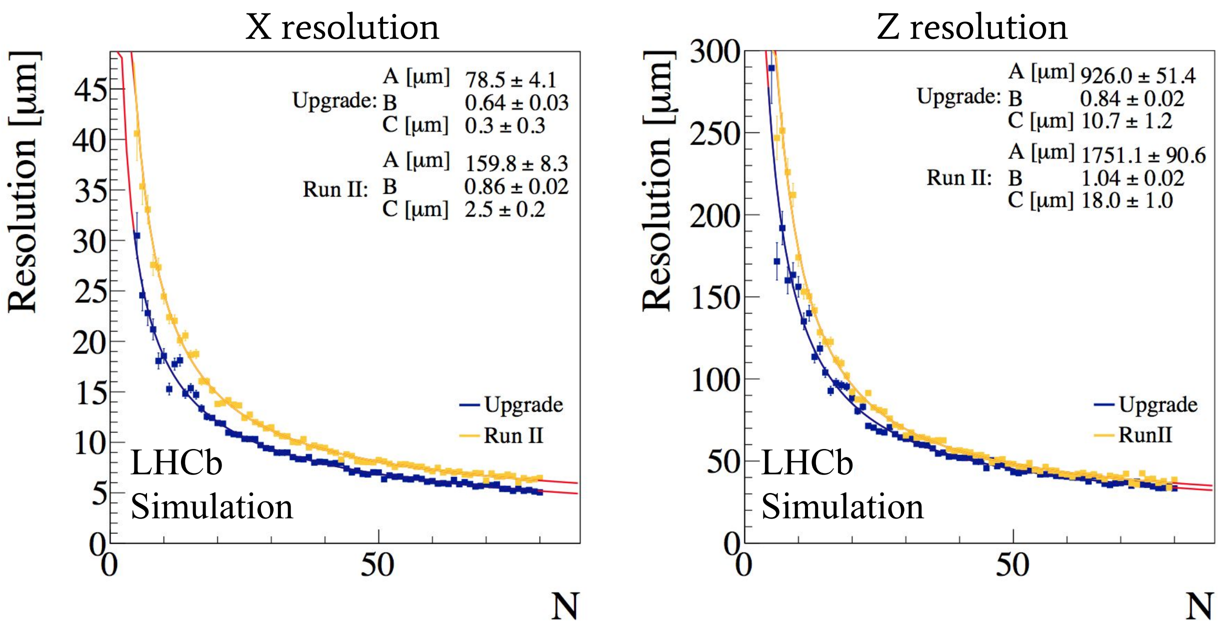


Predicted PV position resolution as function of number of tracks using MC simulation

LHCb-PUB-2017-005

# Fast reconstruction stage

- Several algorithms to reconstruct tracks fully or partially traversing the detector, and find primary vertices
- The event topology is more complex at the upgrade conditions
  - 3-4 times more primary vertices and 2-3 times higher track multiplicity
  - Challenging to keep good physics performance and to lower processing time

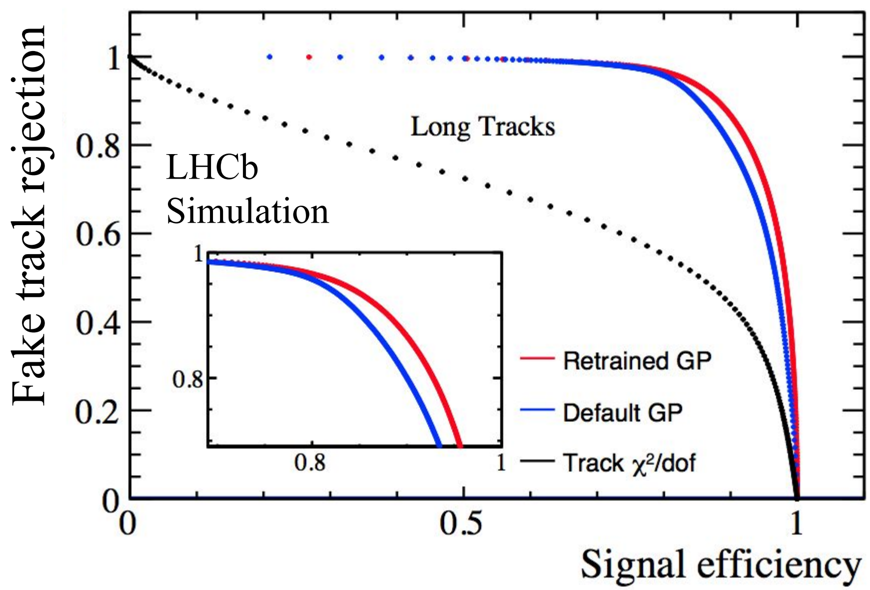


Predicted ROC curves of fake track classifiers using MC simulation.

Optimisation ongoing.

# Performance and timing

- Tracking efficiencies look promising
  - Equal to or better than those in the trigger TDR

| Fake probability | Trigger TDR | Fast stage |
|---|---|---|
| Fake rate | 10.9% | 5.6% |
| long | 42.7% | 42.9% |
| long, from B | 72.5% | 72.7% |
| long, from B, $p_T > 0.5GeV$ | 92.3% | 92.5% |

# Performance and timing

- Tracking efficiencies look promising
  - Equal to or better than those in the trigger TDR
- Timing inline with trigger TDR
  - Using the same framework and single threaded paradigms as today's software

| Timing [ms] | Trigger TDR | Fast stage |
|---|---|---|
| VELO tracking | 2.0 | 2.0 |
| VELO-UT tracking | 1.3 | 0.5 |
| Forward tracking | 1.9 | 2.3 |
| PV finding | 0.4 | 1.1 |
| total | 5.6 | 6.0 |

# Performance and timing

- Tracking efficiencies look promising
  - Equal to or better than those in the trigger TDR
- Timing inline with trigger TDR
  - Using the same framework and single threaded paradigms as today's software
- Throughput performance targets challenging to meet
  - Hardware performance growth at equal cost is slowing dramatically
  - Timing above measured using the software designed a decade ago
    - Concurrency was not a concern in HEP at the time
- A lot of work on new software underway
  - Core framework (Gaudi): built-in thread safety, flexible scheduling, etc.
  - Experiment software: major redesign of algorithms
  - Computing TDR expected at the end of the year
  - See more in Stefan Roiser's and Niko Neufeld's talks

# Turbo stream in Run II

- Turbo: analysis with the trigger output
  - Save offline storage by removing raw and uninteresting data
  - Crucial for analyses needing large samples
  - Real-time data reduction ⇒ be flexible in monitoring quality and updating
- Originally, only exclusive decays selected
- Since 2016 the full reconstruction can be persisted
- Since 2017 we can selectively persist anything in between

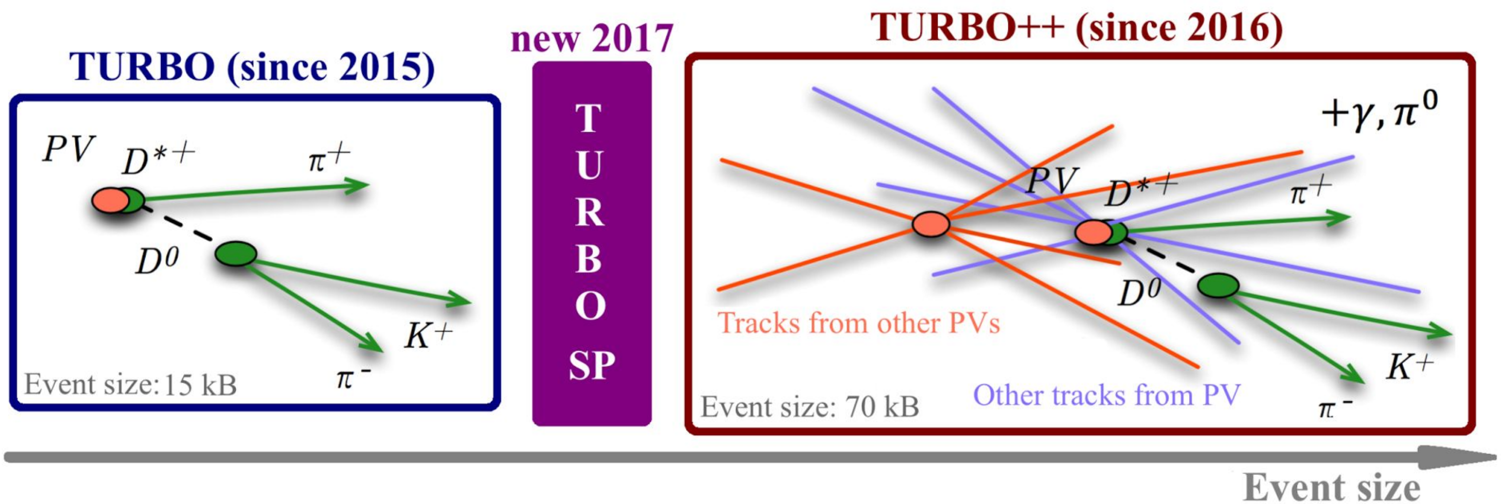


Comput. Phys. Commun. 208 (2016) 35-42

# Output bandwidth division

- How do we divide up the trigger output bandwidth?
  - This is the output to offline storage
  - Finite disk space limits the output BW - not the network or trigger
  - TURBO stream: reduced event size - more signal events for the same amount of disk space
- Use an automated method to divide between channels
  - BW per channel defined by number of channels and physics priority
  - Need a way to tune the output BW consumed per channel
    - Here we study it using a multivariate classifier approach
  - Proof of principle study using four charm decay modes

| Channel | Event size |
|---|---|
| $D^+ \to K^+ K^- \pi^+$ | 14 kB |
| $D^0 \to K^+ K^-$ | 12 kB |
| $D^0 \to K^+ K^- \pi^+ \pi^-$ | 14 kB |
| $D^0 \to K^0_S \pi^+ \pi^-$ | 70 kB (14kB) |

* Event sizes taken from Run II data

LHCb-PUB-2017-006

# Output bandwidth division

- Minimise the $\chi^2$ by varying the MVA response for each decay
  - Channel weight (unity here, but can be used to prioritise)
  - Channel efficiency
  - Maximum channel efficiency (when given the full output BW)

$$\chi^2 = \sum_{i}^{\text{channels}} \omega_i \times \left(1 - \frac{\varepsilon_i}{\varepsilon_i^{\max}}\right)^2$$

- Assign these channels a 60 MB/s bandwidth limit between them and use the algorithm to divide it up
  - Efficiency calculated from signal MC samples
  - Bandwidth calculated from minimum bias MC sample

$$\text{BW[GB/s]} = \text{retention} \times \textbf{rate} \times \text{event size[kB/evt.]}$$

# Output bandwidth division

- Minimise the $\chi^2$ by varying the MVA response for each decay
  - Channel weight (unity here, but can be used to prioritise)
  - Channel efficiency
  - Maximum channel efficiency (when given the full output BW)

$$\chi^2 = \sum_{i}^{\text{channels}} \omega_i \times \left(1 - \frac{\varepsilon_i}{\varepsilon_i^{\max}}\right)^2$$

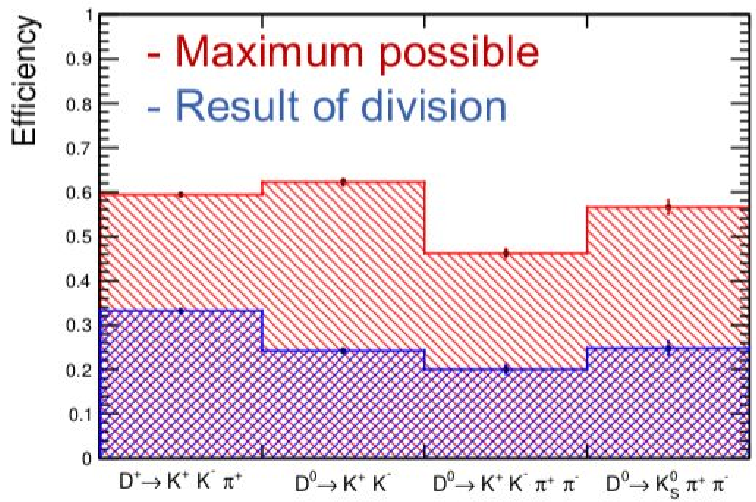


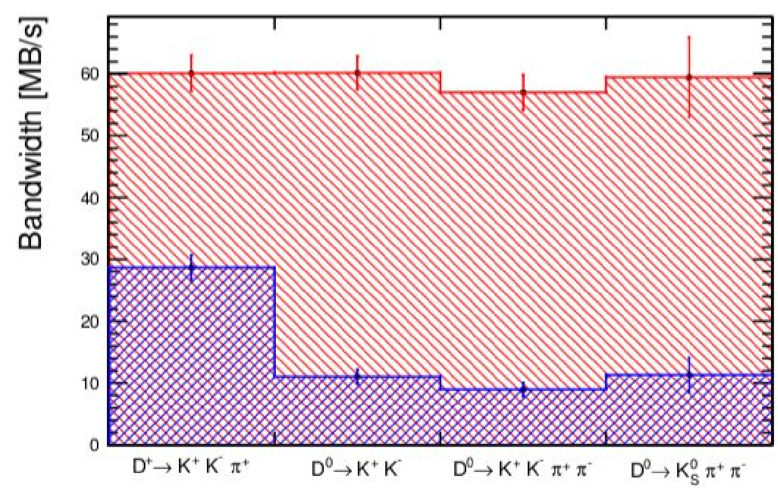

# Output bandwidth division

- Minimise the $\chi^2$ by varying the MVA response for each decay
  - Channel weight (unity here, but can be used to prioritise)
  - Channel efficiency
  - Maximum channel efficiency (when given the full output BW)

$$\chi^2 = \sum_{i}^{\text{channels}} \omega_i \times \left(1 - \frac{\varepsilon_i}{\varepsilon_i^{\text{max}}}\right)^2$$

- Signal efficiencies will ultimately depend on analyst's ability to define powerful selections
  - Already using machine learning in the trigger and will be more common
  - Reduction of the event size, more signal for the same BW usage

# Output bandwidth division

Event size is very important: leads to huge efficiency gains

Change from 70 to 14 kB for $D^0 \rightarrow K^0_S \pi^+ \pi^-$

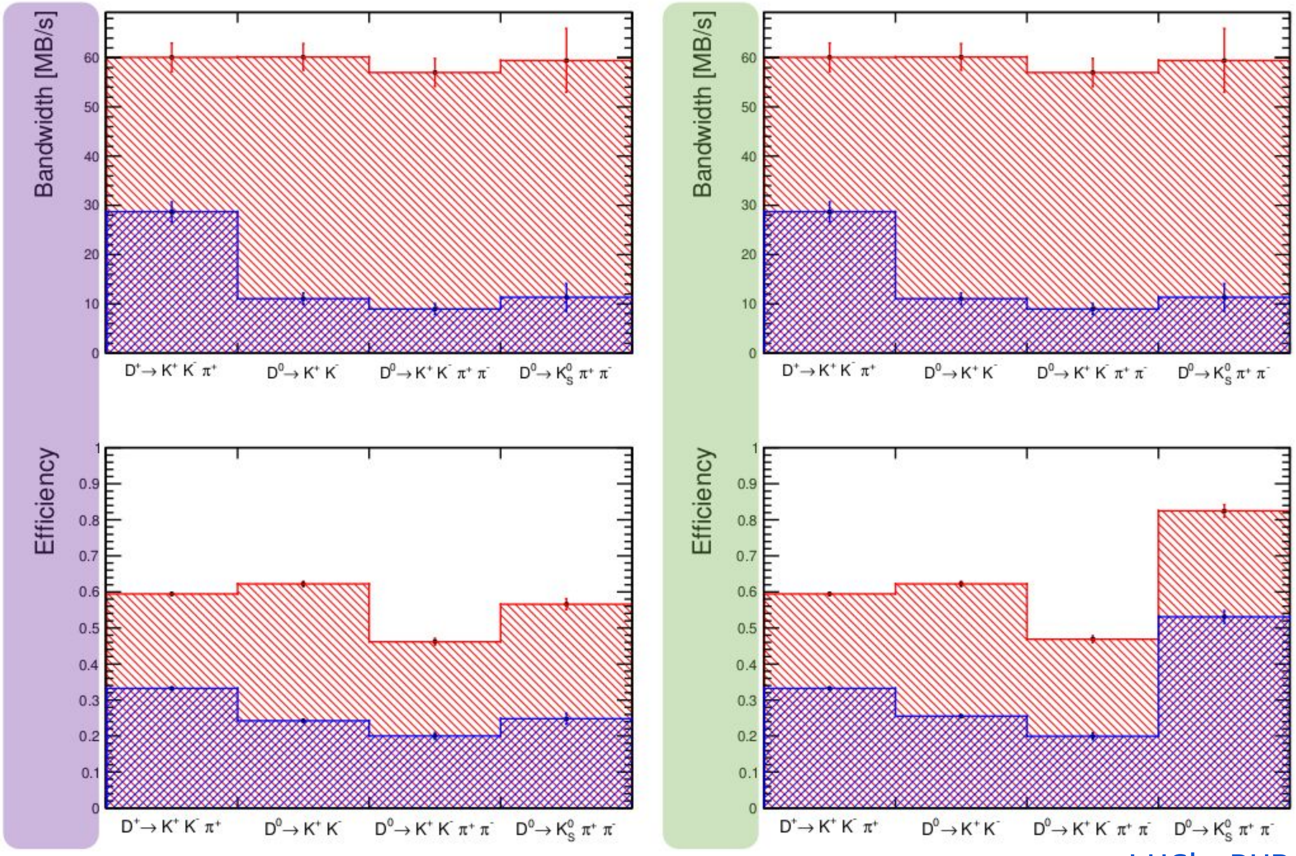

# Summary

- In Run II we can do offline quality selections in the trigger thanks to
  - real-time alignment and calibration
  - fully using the online CPU resources (to fit the full offline reconstruction)
- Selectively persisting subset of the event data can already be done
- LHCb upgrade trigger studies well underway
  - Promising performance on simulated data
  - Throughput will improve with adaptation to multithreaded running, employing vectorization and redesigning reconstruction algorithms.
  - Proof of principle for bandwidth division
- Lots more to come in the next couple of years

# Thank you

# Trigger schematic in Run II and Run III

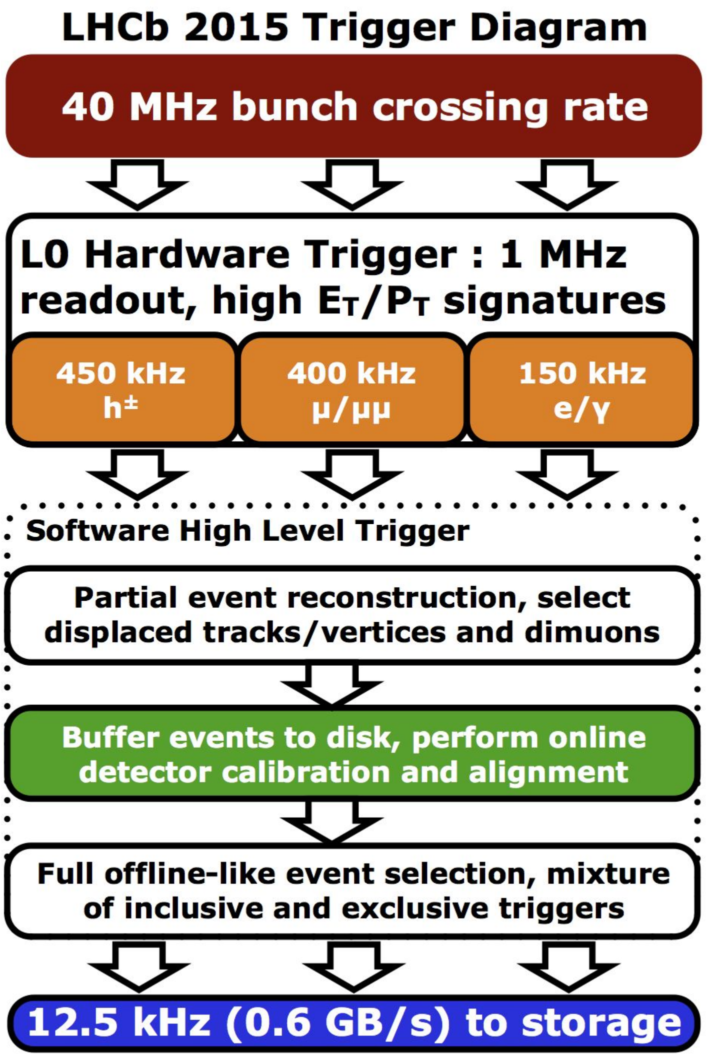


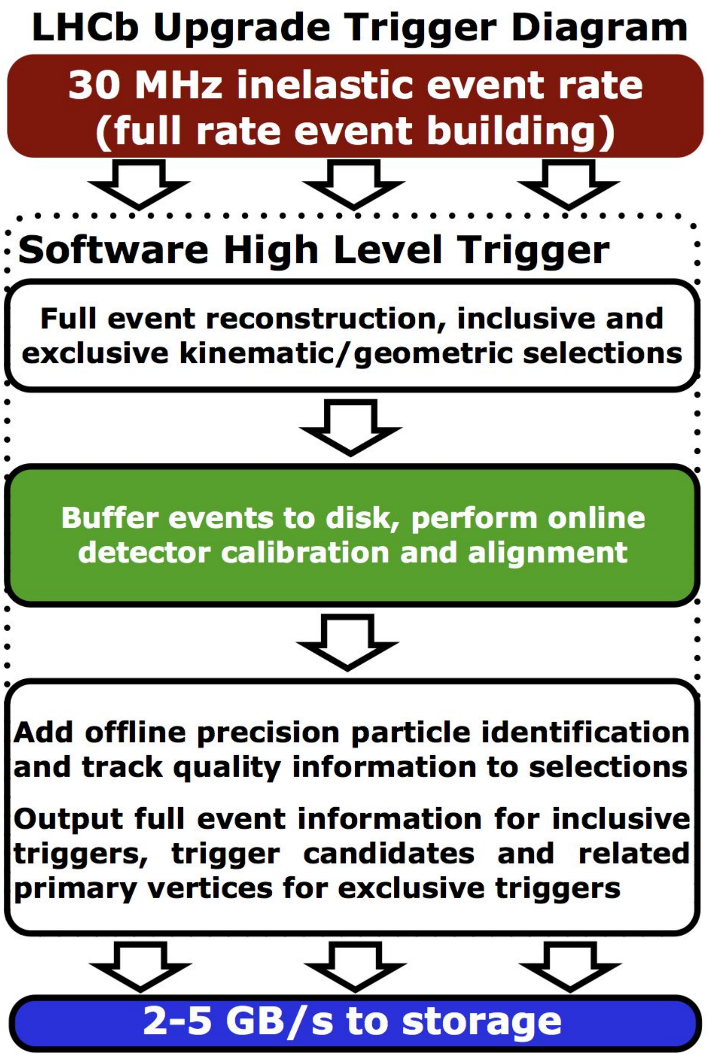

# Run II online system architecture

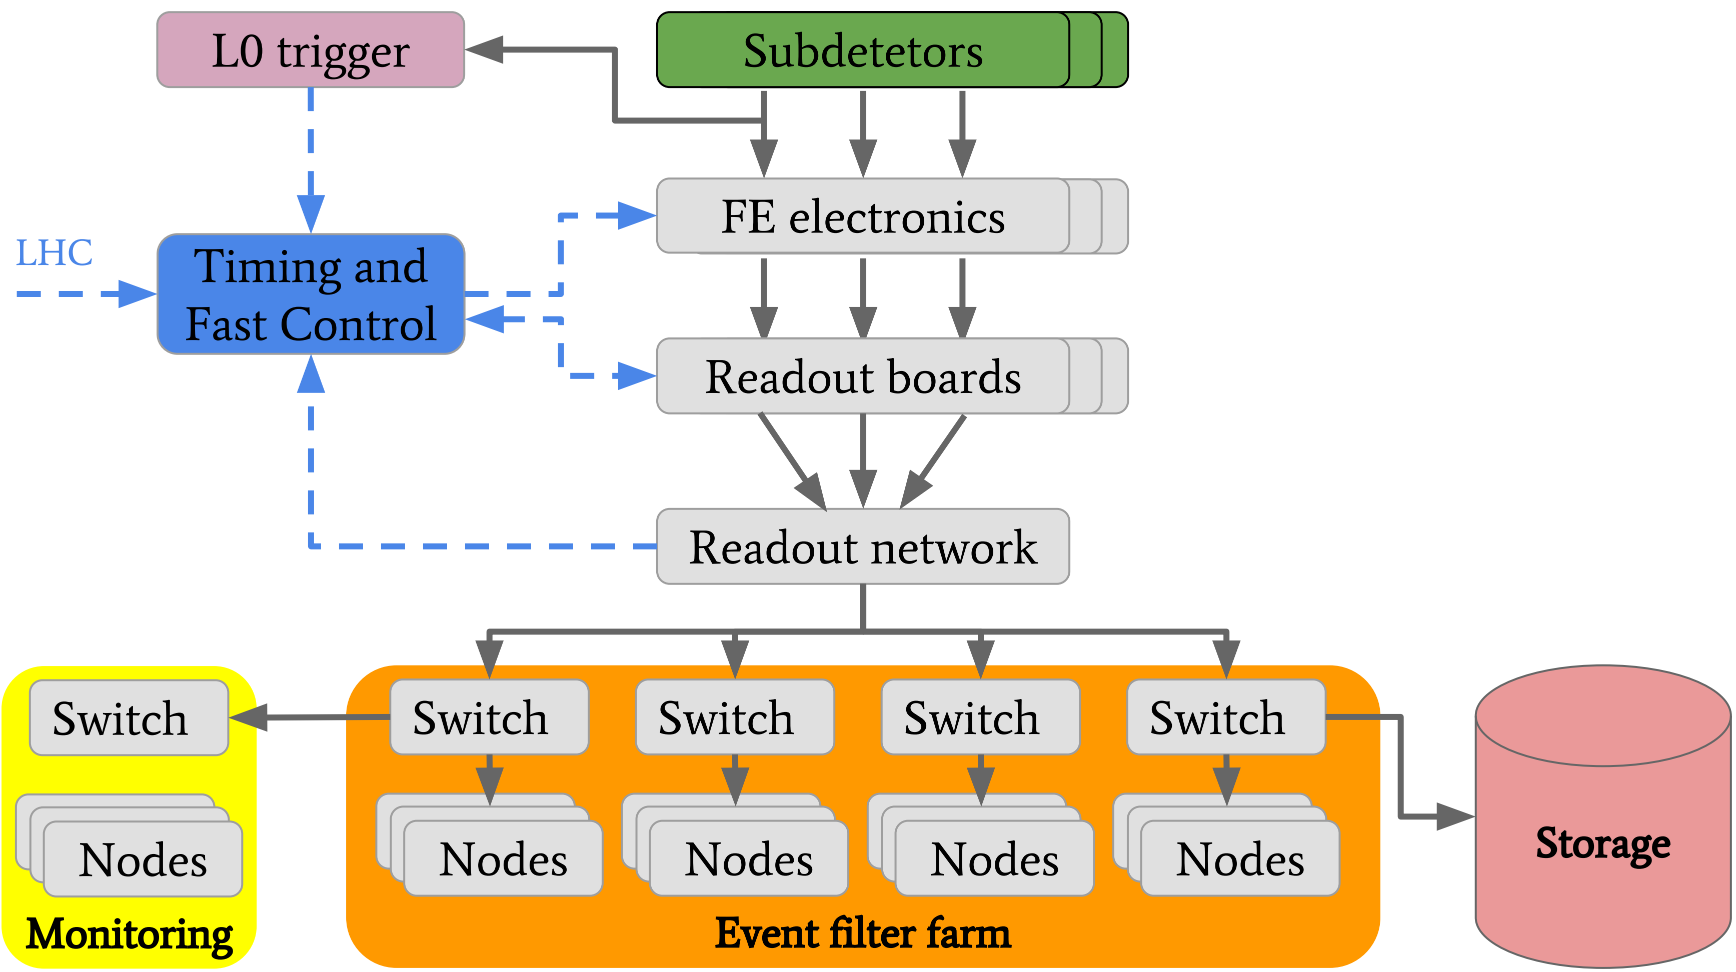


JINST 3 (2008) S08005

# Deferred triggering

- Stable beams ~50% of the time
- Buffer events to disk and process between fills
- Run I
  - Defer 20% of L0 accepted events
  - Effectively 25% more CPU
- Run II
  - Defer 100% of HLT 1 accepted events
  - More efficient use of buffers due to larger real-time reduction
  - Save 100% of events at 150 kHz instead of 20% at 1 MHz
  - Use HLT 1 output for calibration and alignment
  - 10 PiB in farm (half in 2015)

Disk usage models

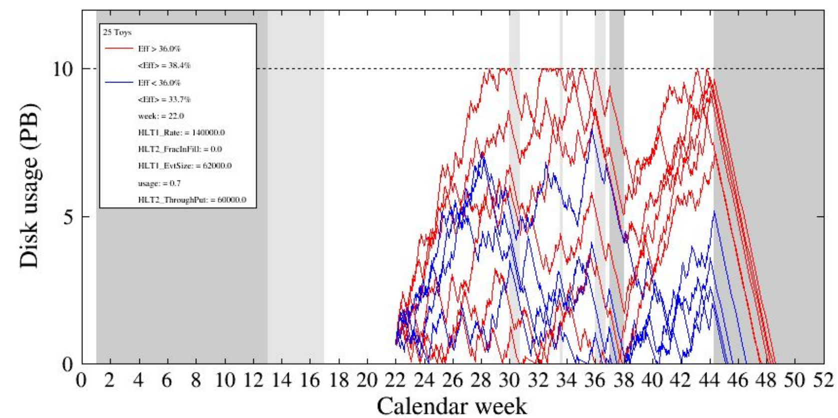

# HLT 1 overview in Run II

- Inclusive selections, ~100 kHz
  - Single and two track MVA selections
- Inclusive muon selections, ~40 kHz
  - Single and dimuon selections
  - Additional low pT track reconstruction
- Exclusive selections
  - Lifetime unbiased beauty and charm selections
  - Selections for alignment
- Low multiplicity trigger for central exclusive production analyses

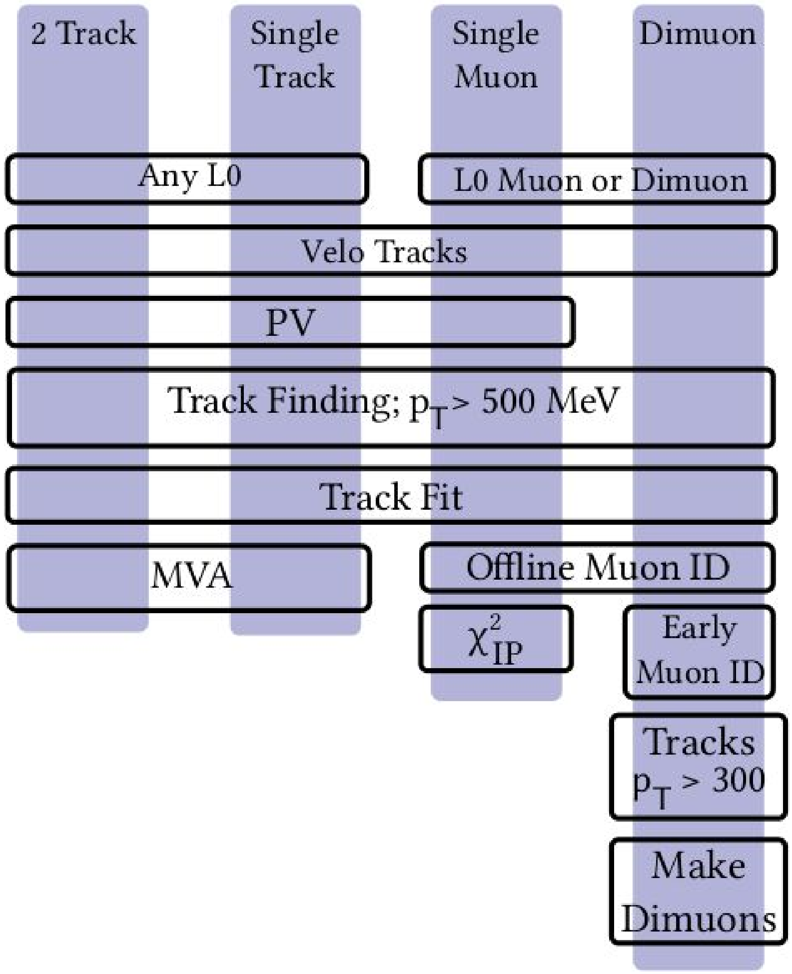

# Real-time alignment and calibration

- Alignment and calibration crucial for optimal physics performance
- Alignment per fill
  - Collect suitable data with dedicated HLT 1 selections
  - Run alignment workers on the HLT farm (1 per node)
  - Controller iterates until converged, ~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 ~1 min after collection of data

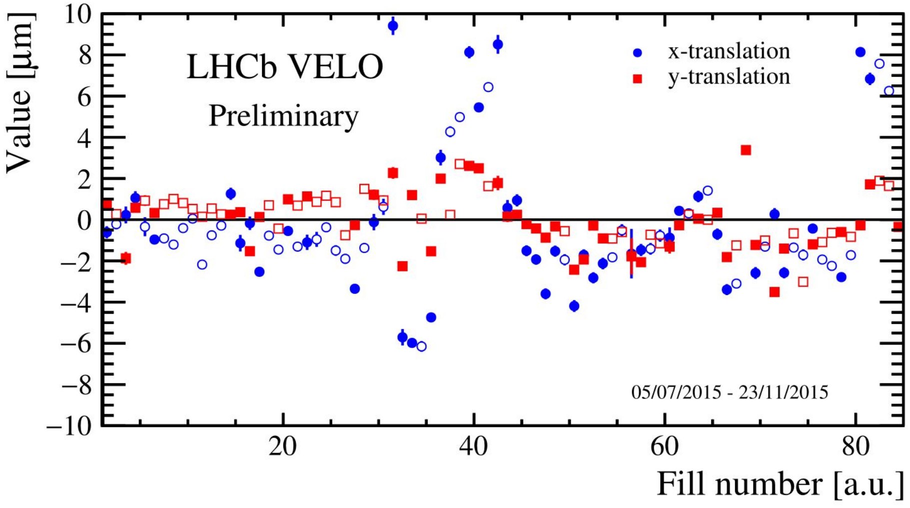

# Analyses with TURBO

- Presented preliminary results **one** week after data taking!
- Published Run II measurements performed exclusively with Turbo

**Measurements of prompt charm production cross-sections in $pp$ collisions at $\sqrt{s} = 13$ TeV**

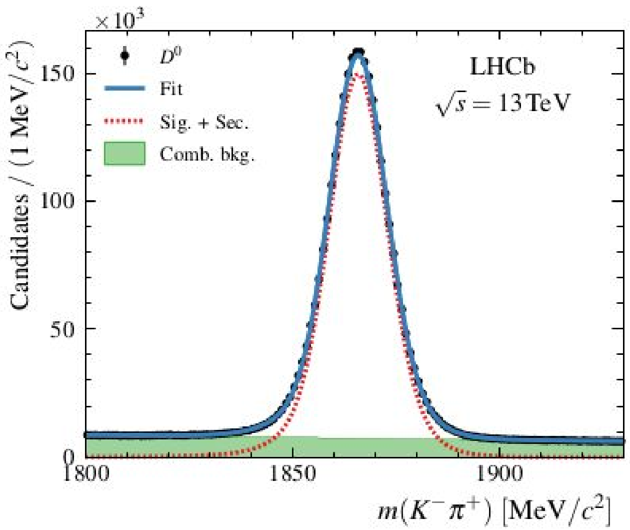


JHEP 03 (2016)159

**Measurement of forward $J/\psi$ production cross-sections in $pp$ collisions at $\sqrt{s} = 13$ TeV**

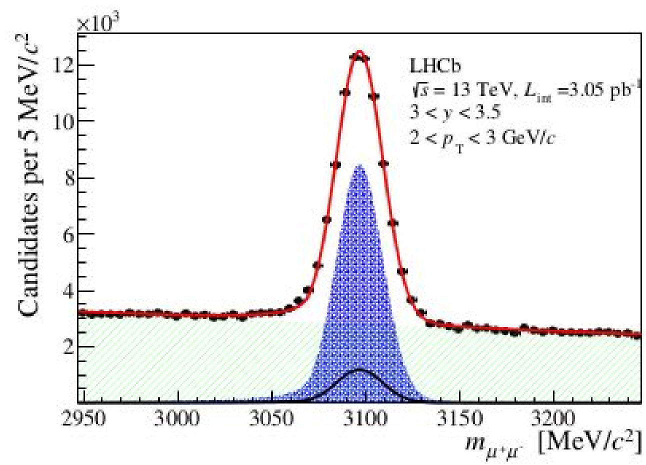


JHEP 10 (2015)172

# TURBO++

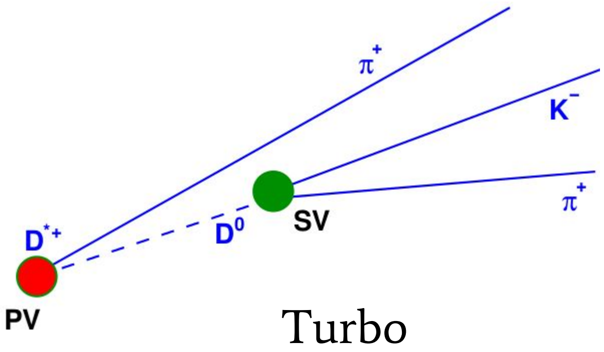


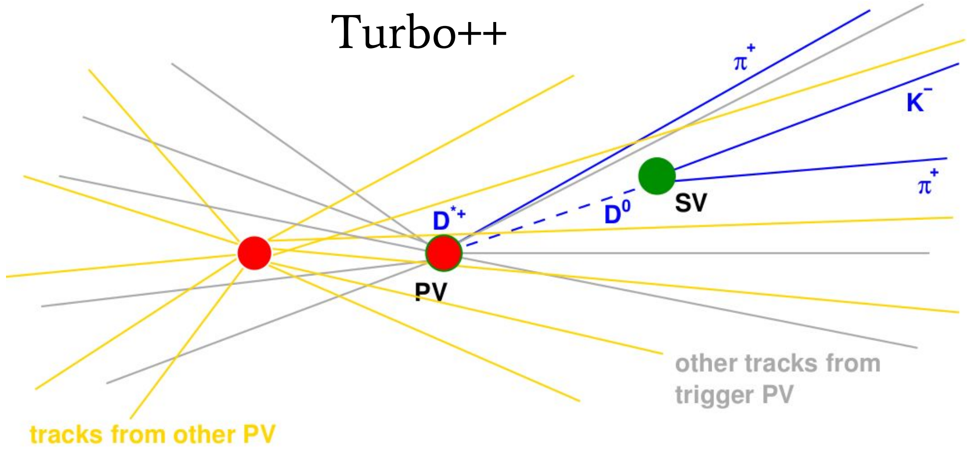


- New in 2016
- Persist arbitrary variables like isolation with HLT candidate
- Can save HLT candidate + any reconstructed objects
  - Custom binary serialization in SOA format, LZMA compression per event
  - Event size of 50 kB, including a minimal subset of the raw data
- Can do qualitatively new things on HLT output
  - Entire analysis can be done on trigger output, incl. flavour tagging
  - e.g. in charm spectroscopy: $D^* \rightarrow D^0( K^-\pi^+)\pi^+$