The LHCb Trigger in Run II

Rosen Matev, CERN
on behalf of the LHCb collaboration

20th Real Time Conference
2016, Padova

The LHCb experiment

The LHCb experiment

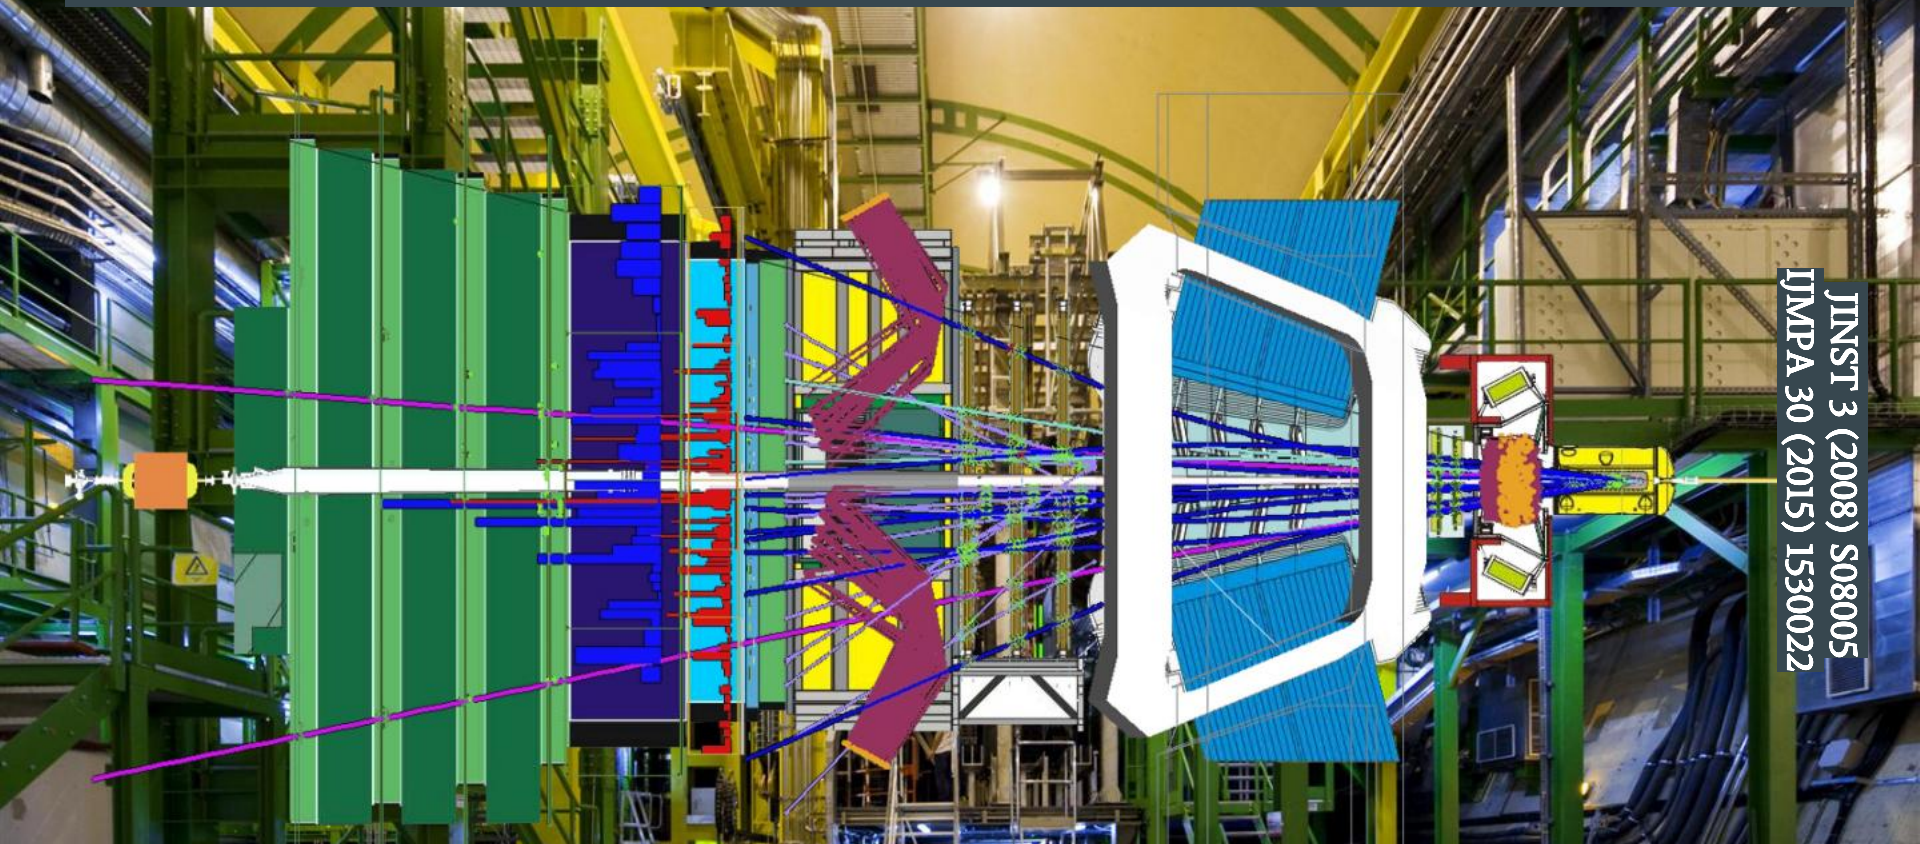

JINST 3 (2008) S08005
IJMPA 30 (2015) 1530022

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

at 13 TeV and $L = 4\times10^{32}\ \mathrm{cm^{-2}s^{-1}}$

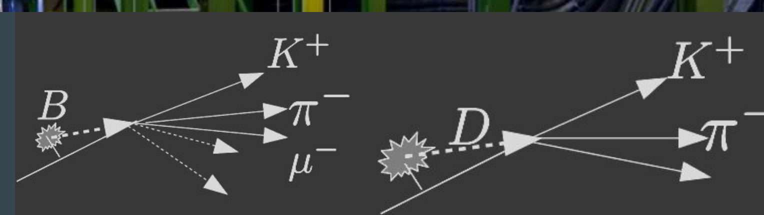

Online system architecture

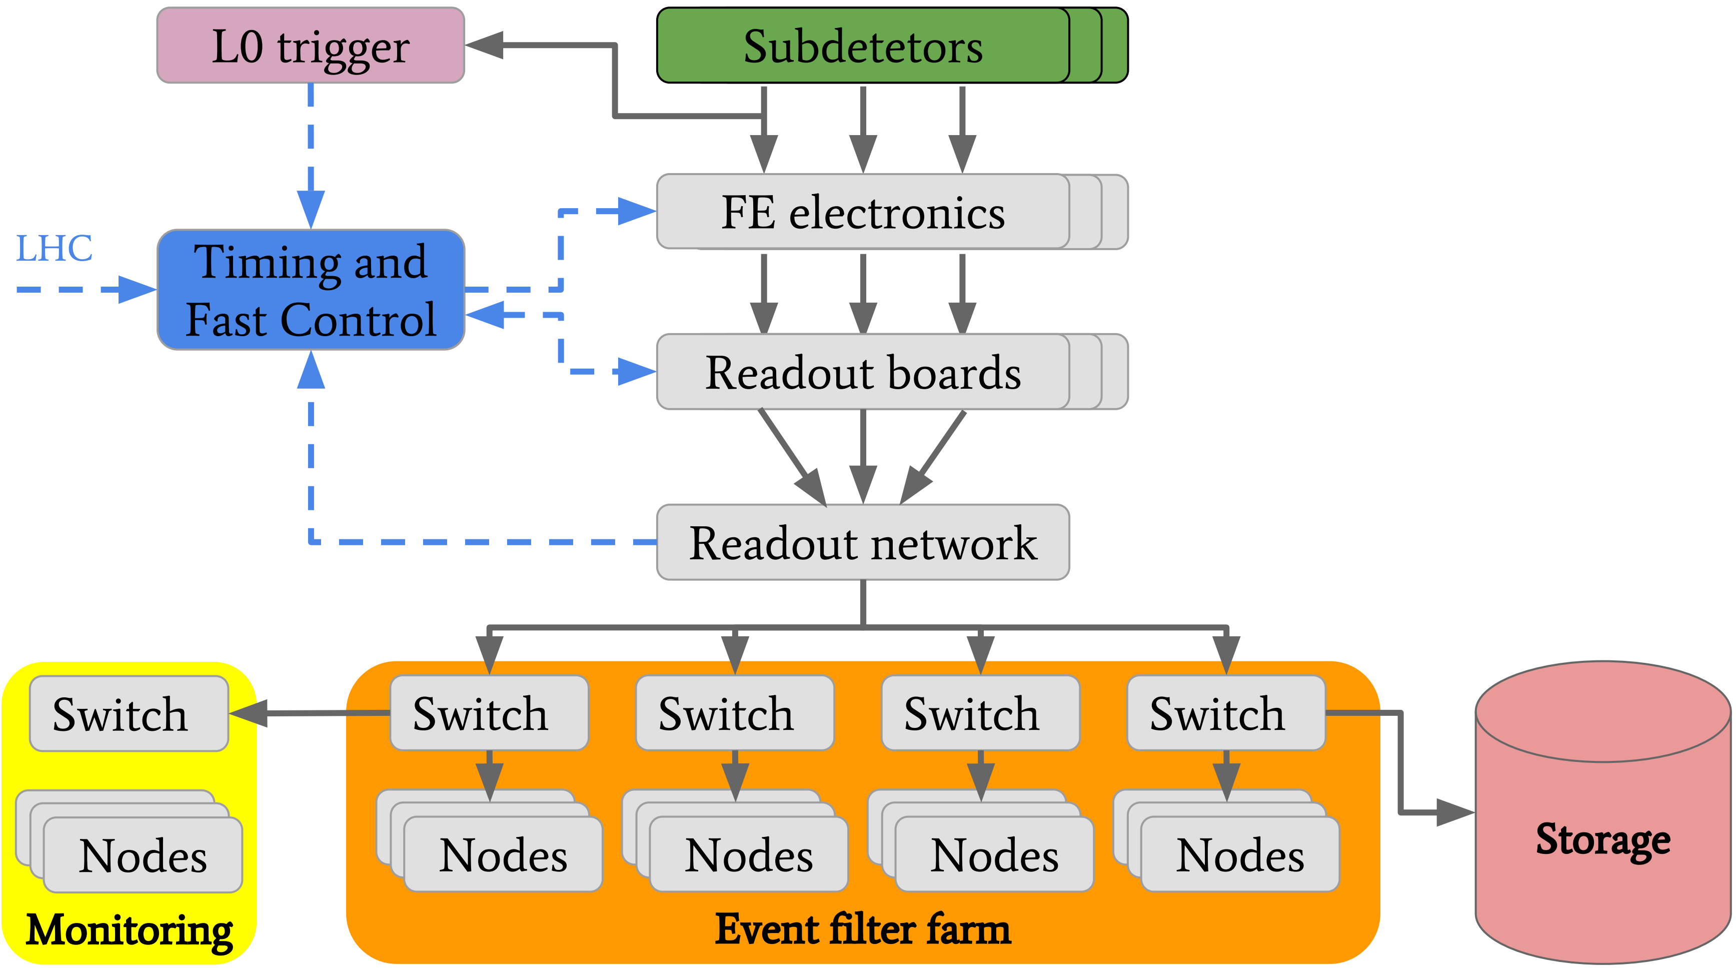

JINST 3 (2008) S08005

Event filter farm

- 62 homogeneous sub-farms
- 50880 logical cores
- 800 new nodes in Run II
  almost double performance
- Servers alone cost >5 MCHF

	CPU	# cores	RAM	disk
920 x	Intel x5650	12 (24)	24 GB	2 TB
100 x	AMD 6272	16 (32)	32 GB	2 TB
800 x	Intel E5-2630v3	16 (32)	32 GB	4 TB

Run I trigger

- Hardware trigger (L0)
  - Based on calorimeters and muon detectors
  - Fixed latency of 4 μs
  - Reduces rate to 1 MHz
- Software trigger (HLT)
  - Runs on event filter farm
  - HLT 1: partial reconstruction
  - HLT 2: full up front reconstruction
  - Events buffered to allow processing out of fill
  - Output rate 5 kHz
  - Total time budget ~35 ms/event

LHCb 2012 Trigger Diagram

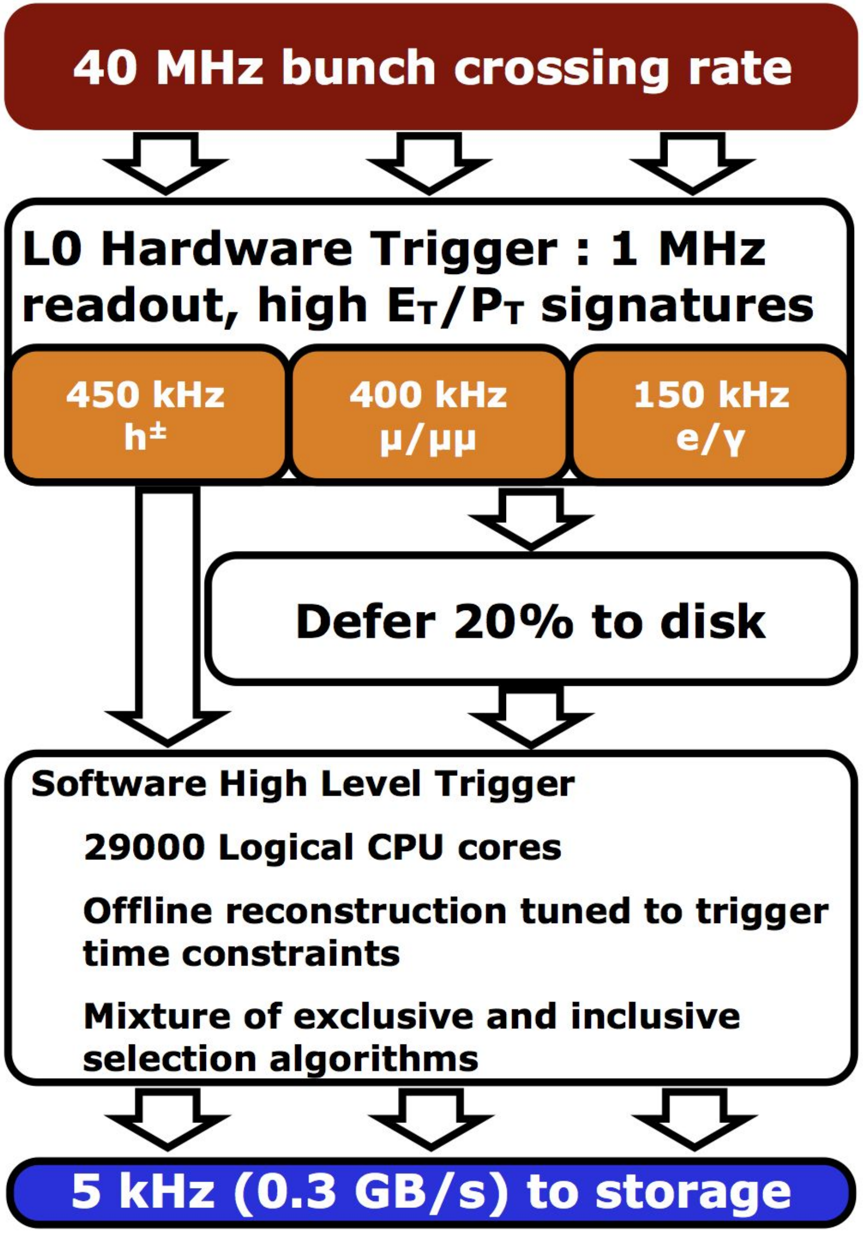

Run I: online and offline

- Online
  - Best possible within CPU budget
- Offline
  - Best available performance regardless of CPU
- Differences
  - Pattern recognition
  - Alignment
  - No hadron identification online
  - Particle selections
- Goal for Run II
  - Run offline reconstruction online
  - Offline quality alignment and calibrations needed online

Ex.: $B_s^0 \to J/\psi\ \varphi$ lifetime measurement

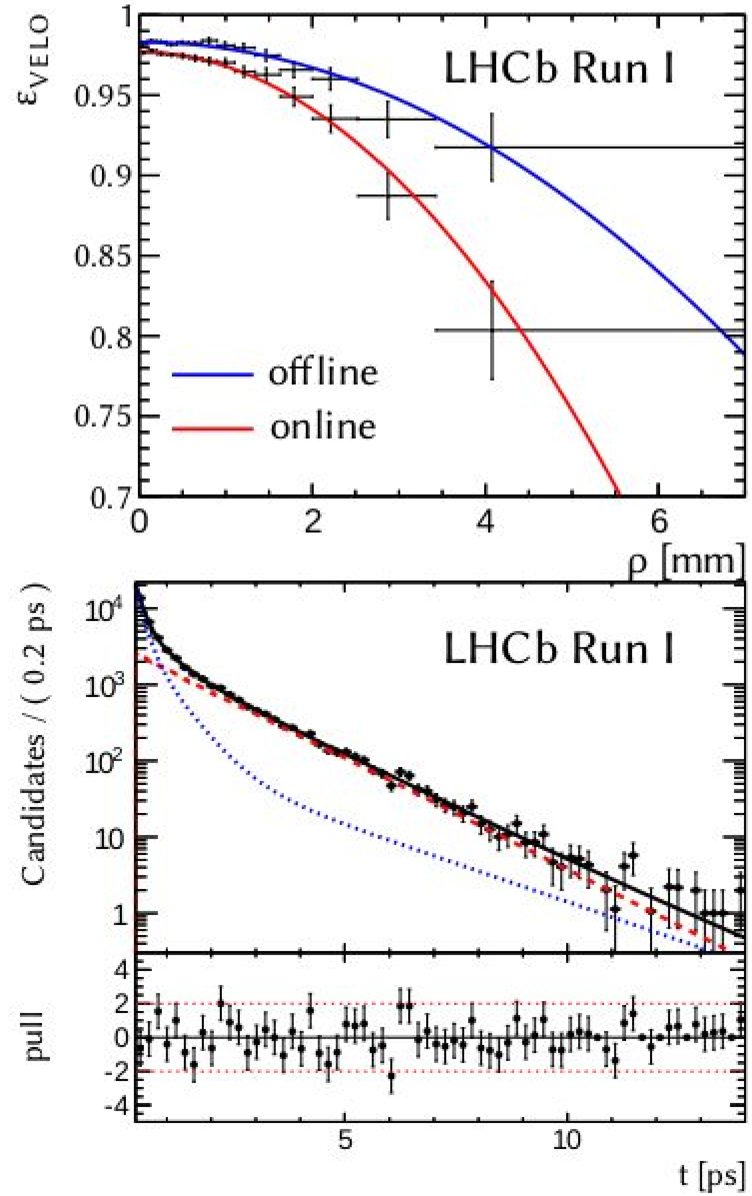

JHEP 04 (2014)114

Deferred triggering

- Stable beams ~35% 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)
    - Sufficient buffer given LHC's performance comparable to 2015

Disk usage models

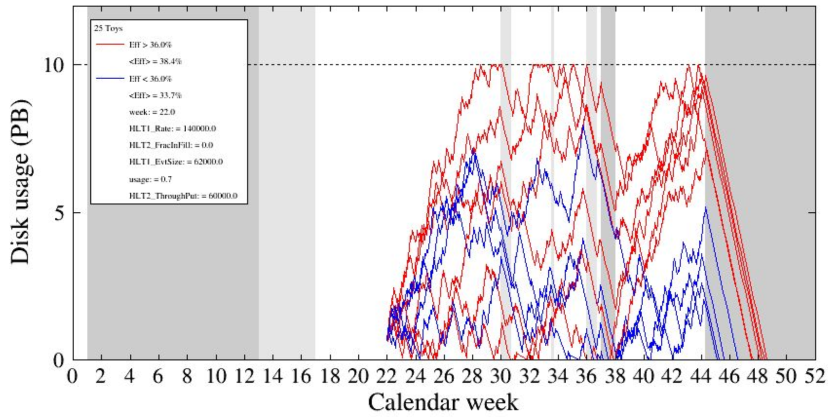

Run II trigger

- 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
    - HLT 1 and HLT 2
  - Events buffered after HLT 1
  - Output rate 12.5 kHz
  - HLT software 40% faster

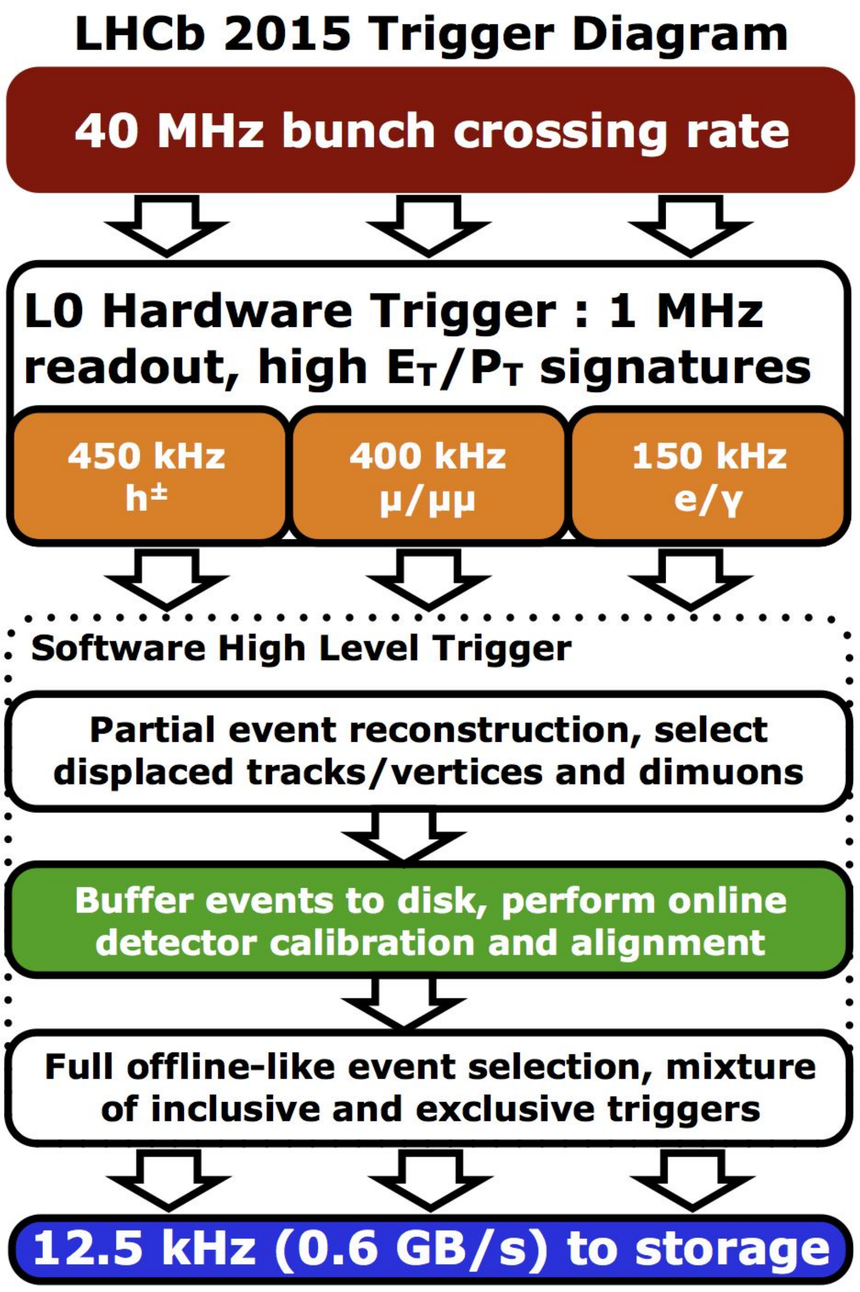

HLT 1 overview

- 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

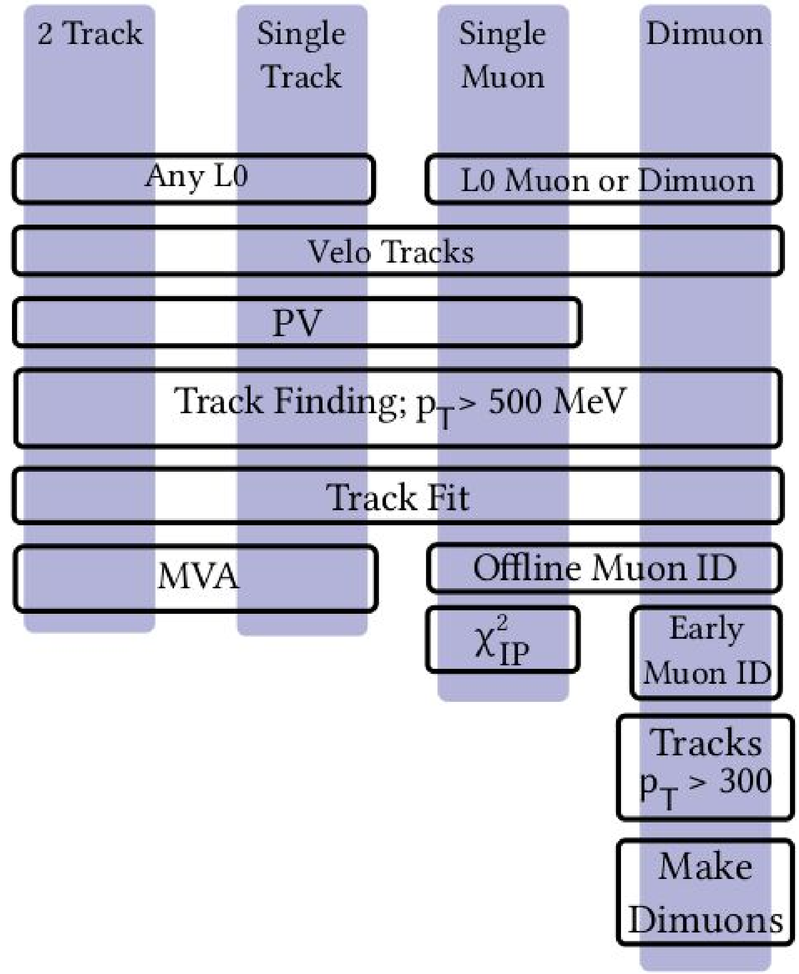

Real-time alignment and calibration

- 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

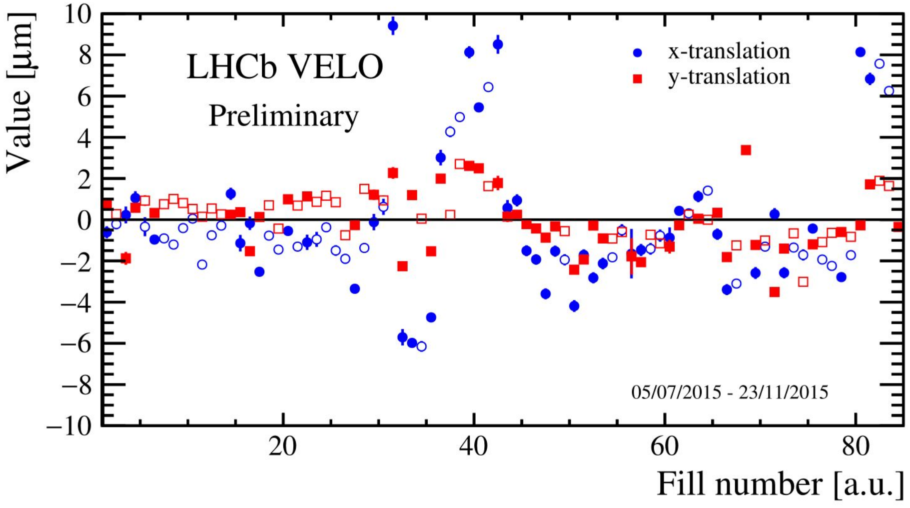

HLT 2 reconstruction

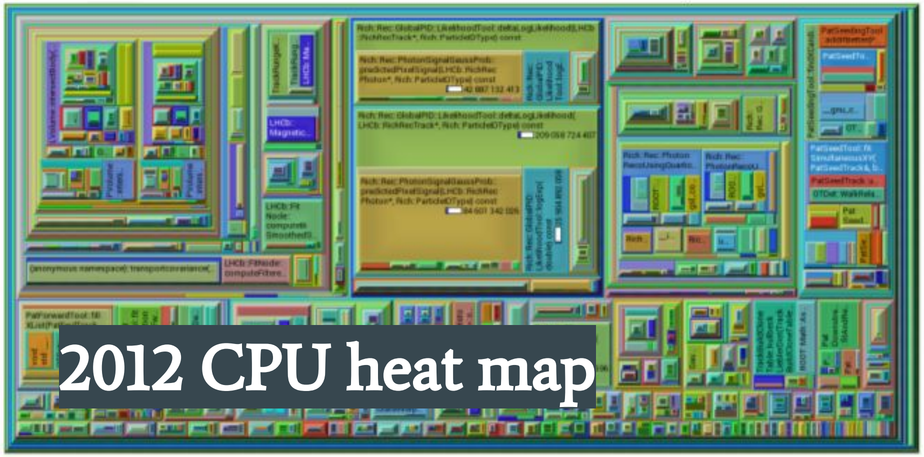

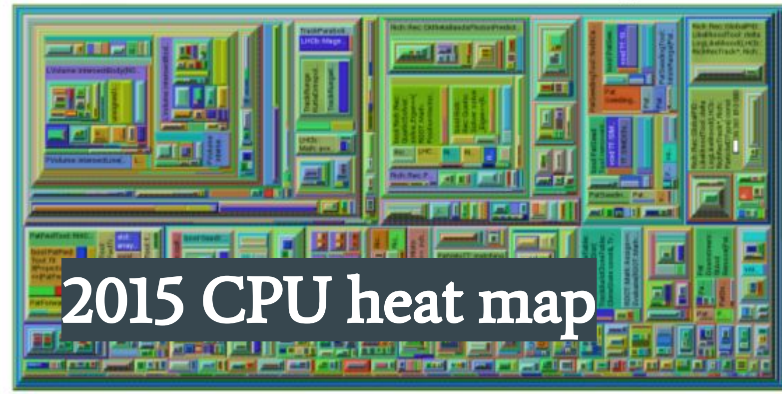

	Run I	Run II
HLT 1 time	~20 ms	~35 ms
HLT 1 rate	~80 kHz	~150 kHz
HLT 2 time	~150 ms	~650 ms
HLT 2 rate	~5 kHz	~12.5 kHz

- Full event reconstruction
  - HLT 1 vertices and tracks
  - All charged tracks
  - Neutral particles
  - RICH, Muon and Calo PID
  - Same reconstruction online and offline
  - 30% speedup achieved

Real-time analysis – TURBO stream

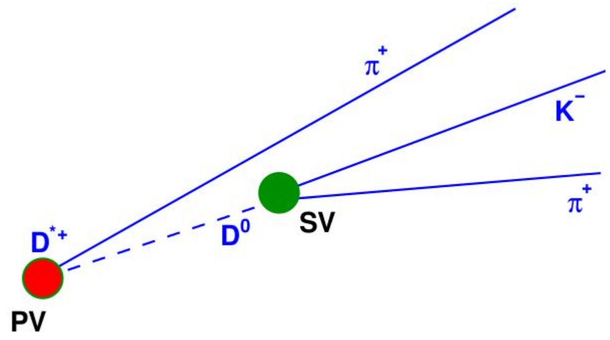

arXiv:1604.05596

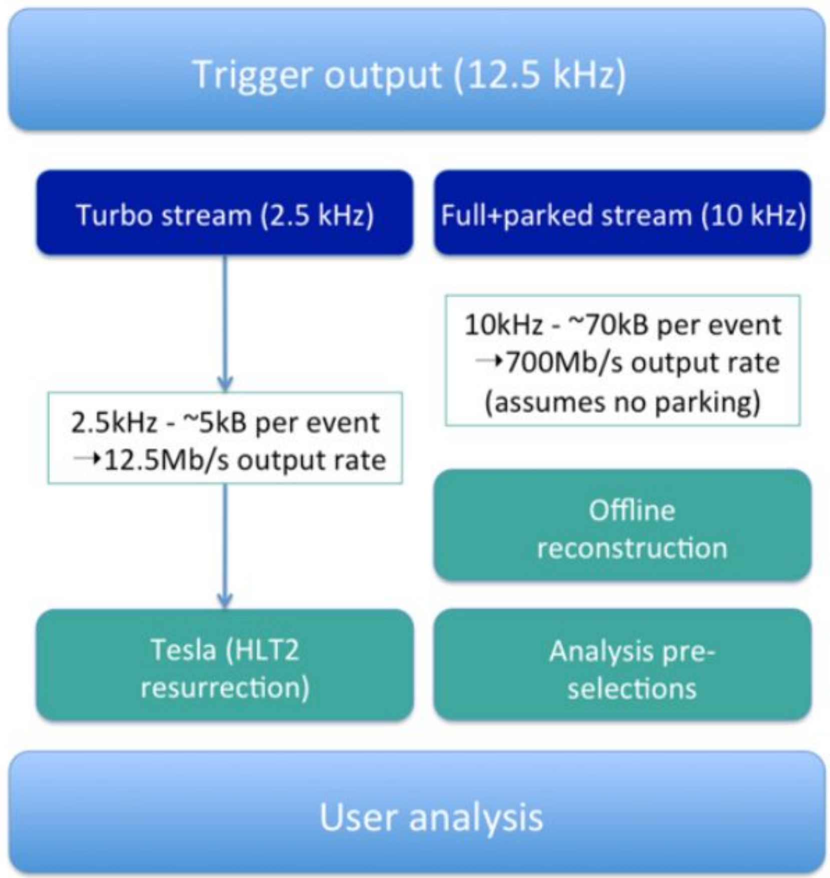

- Online = offline
  - Do physics analysis with HLT reconstructed objects
- TURBO stream
  - Store HLT candidate information
  - Remove most of detector raw data
  - Space required reduced by > 90 %
- Ideal for high-yield analyses
- Very fast turn around (~24 h)

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

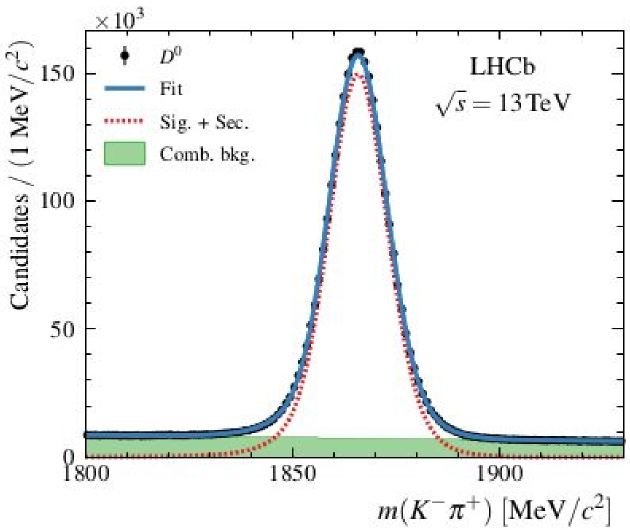

JHEP 03 (2016)159

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

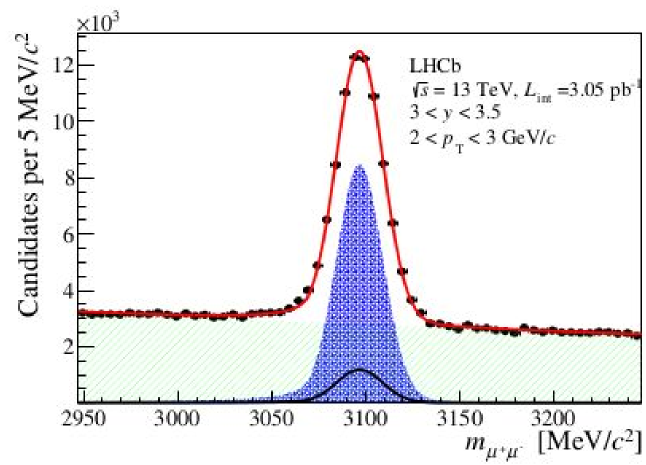

JHEP 10 (2015)172

TURBO++

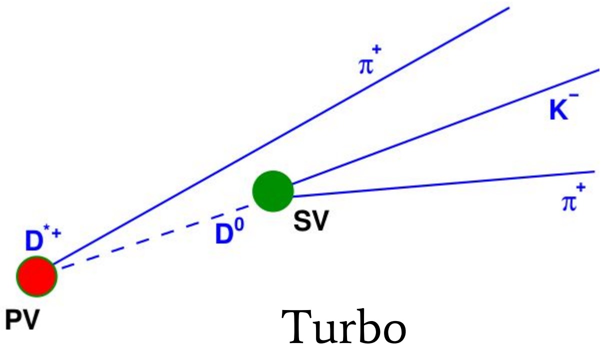

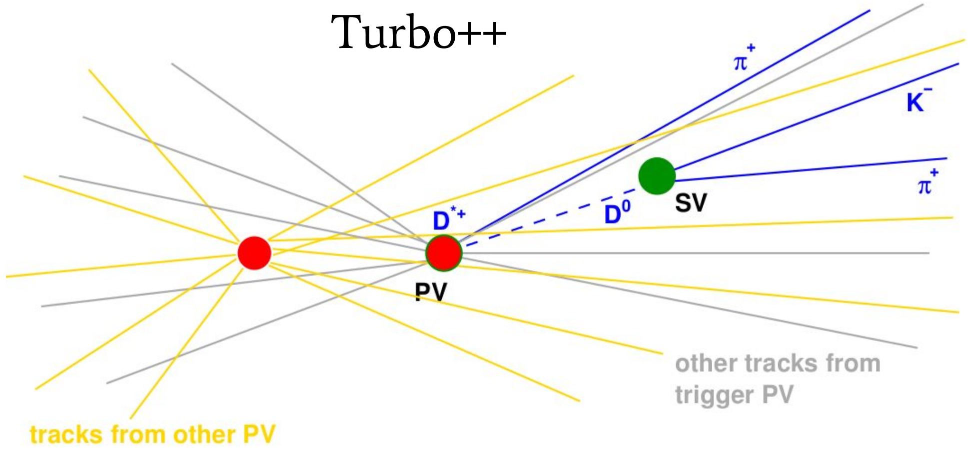

- New in 2016
- Persist arbitrary variables like isolation with HLT candidate
- Can save HLT candidate + any reconstructed objects
  - Custom binary serialization in SOA format, LZMA compression per event
  - Event size of 50 kB, including a minimal subset of the raw data
- Can do qualitatively new things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

TURBO++

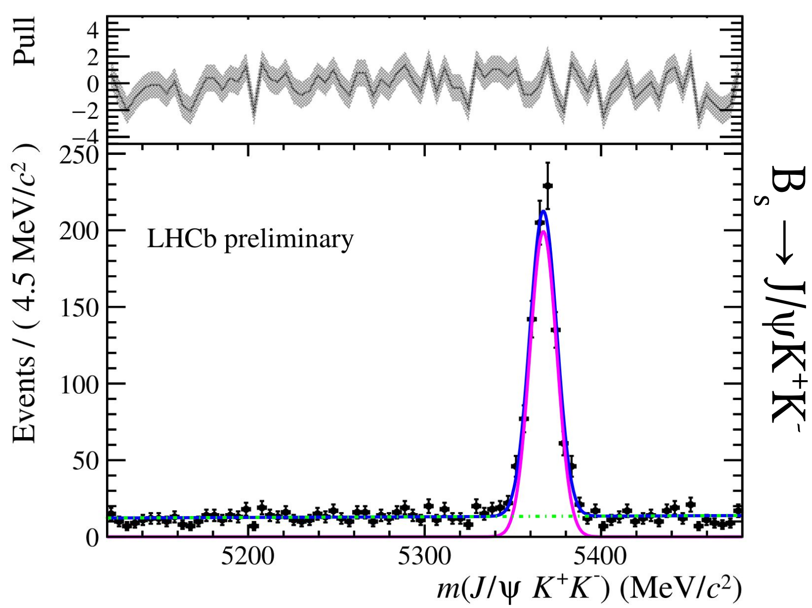

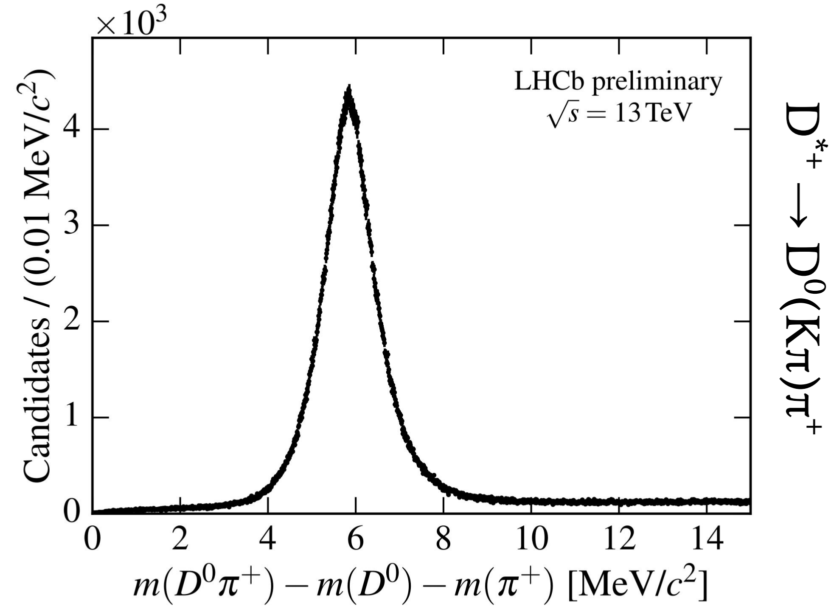

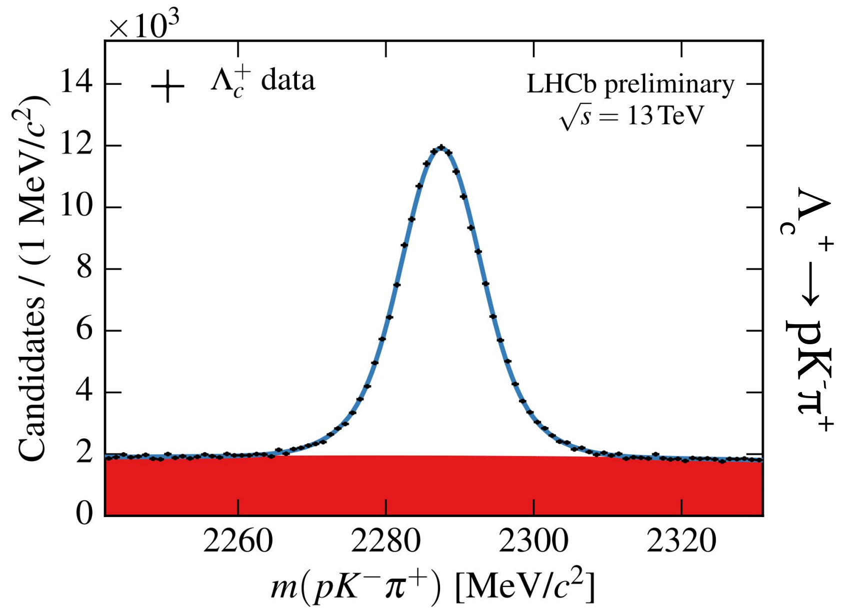

Successful use of the TURBO++ with the new 2016 data!

Conclusions

- Two times more CPU power in HLT farm
- Full offline-quality reconstruction available online
- Calibration and alignment running online
- TURBO stream – do physics with trigger objects
  - More physics per byte
  - 2015 – save exclusive decays
  - 2016 – can save all reconstructed objects
  - Results already published and more to come
- Critical for the upgrade (Run III)
  - Real-time alignment and calibration
  - Selective persistence of raw and reconstructed data


Thank you