The LHCb trigger in Run3

Alessandro Scarabotto

Jahrestreffen der deutschen LHCb-Gruppen (Theory and Experiment)

24/09/24, Bochum, Germany

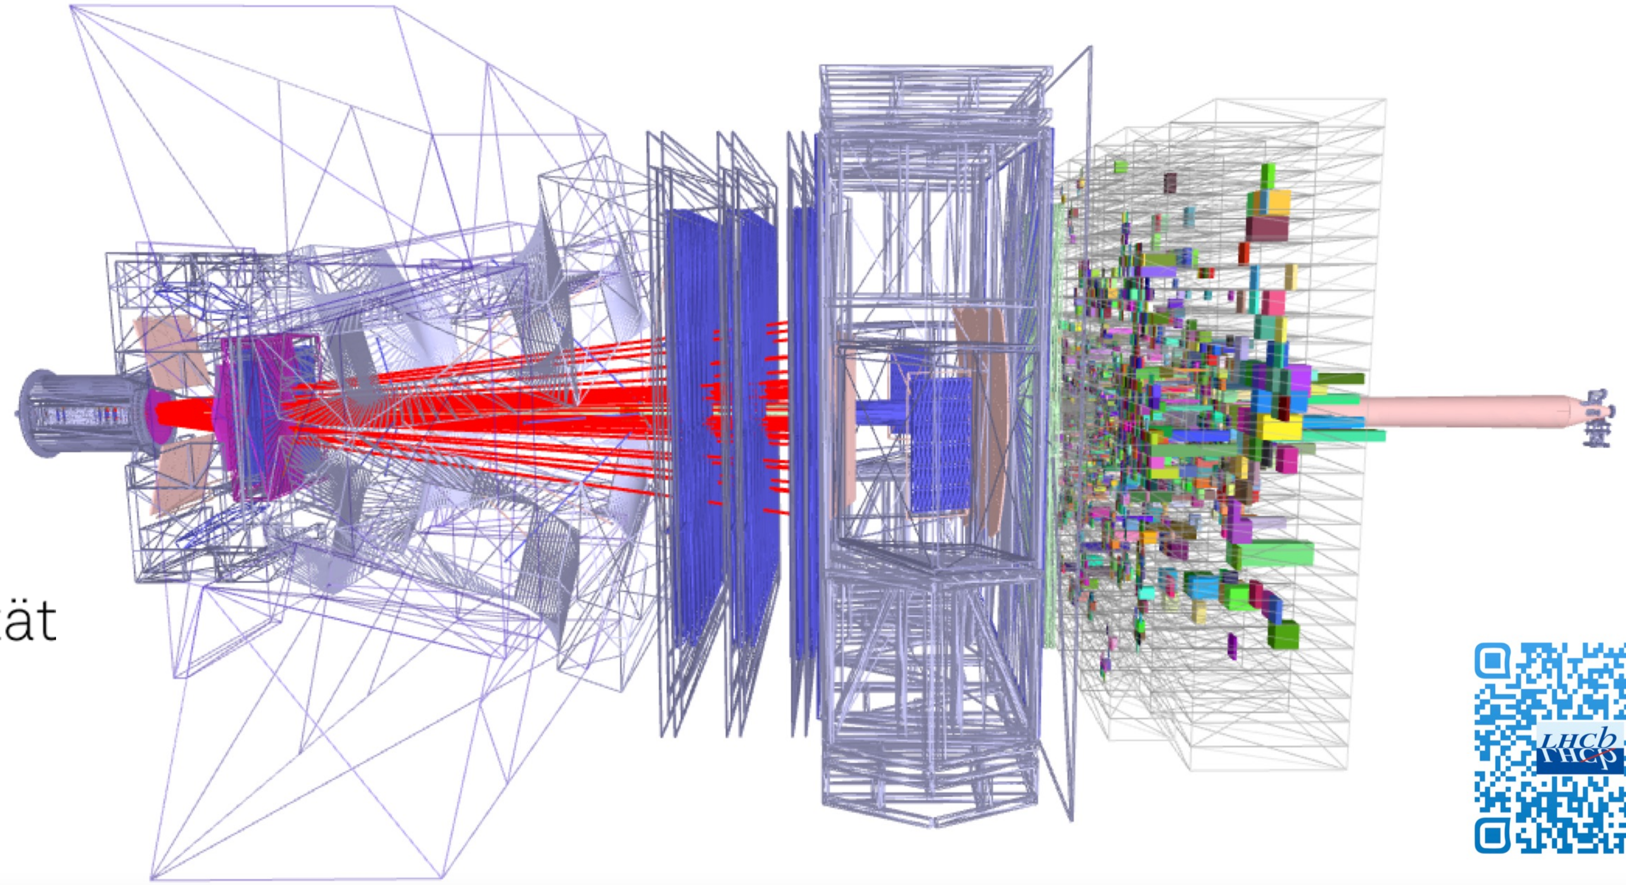

technische universität
dortmund

The LHCb Run3 upgrade

- Major upgrade to collect data with 5 times higher luminosity:

$$\mathcal{L} = 2 \cdot 10^{33}\, cm^{-2} s^{-1}$$

- Increasing pile-up $\langle \mu \rangle = 5.3$ (pp collisions per bunch crossing)
- Full replacement of tracking detectors needed to deal with higher occupancy environment
- What about triggering?

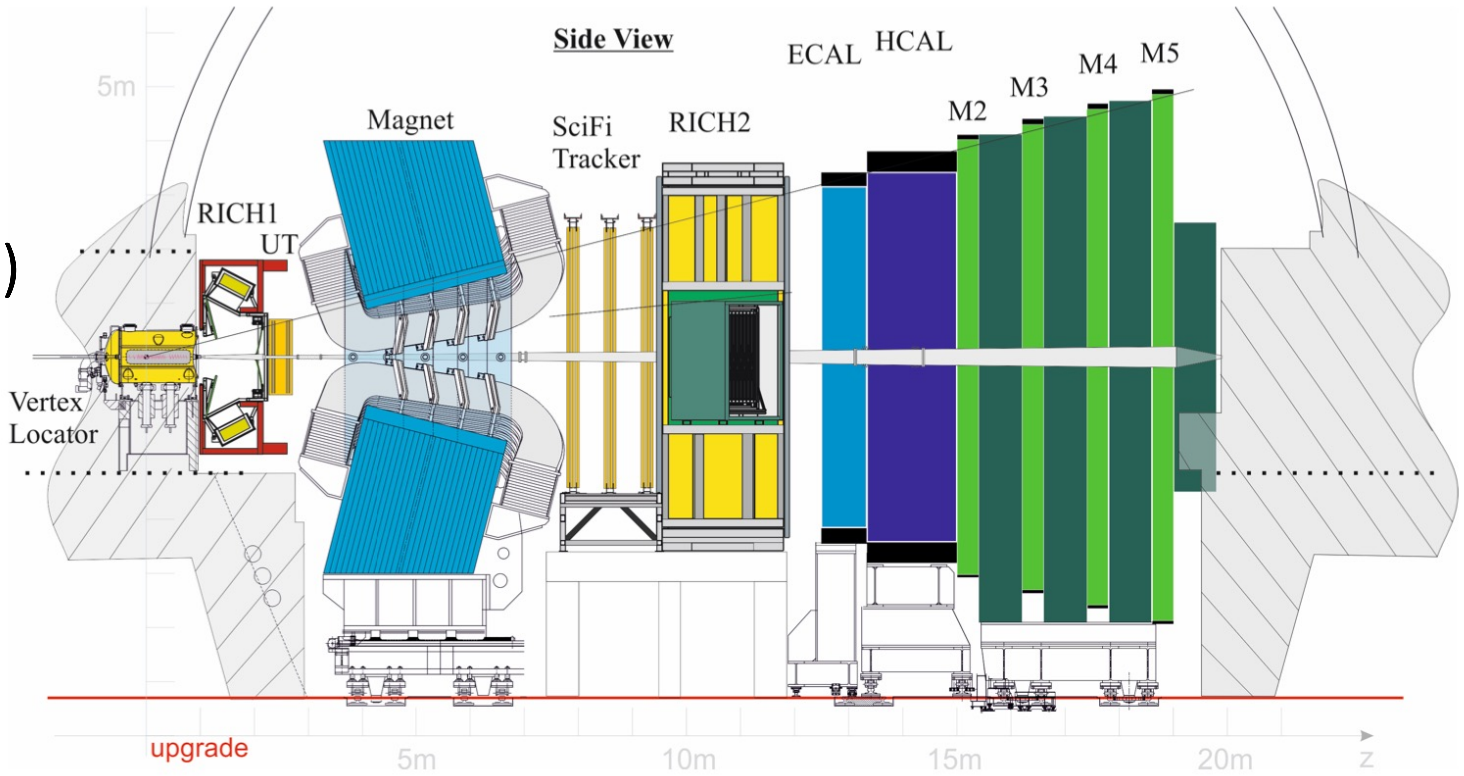

LHCb-TDR-12

The LHCb Run3 trigger

- Limitation of Run2 trigger is the first-level hardware stage (L0)
- Saturation of trigger yields by increasing luminosity
- Caused by tight momentum/energy requirements at L0 selections
- The Run3 LHCb:
    - Removal of L0! But …
    - Reconstruction at 30 MHz LHC pp collision rate for the High Level Trigger (HLT)

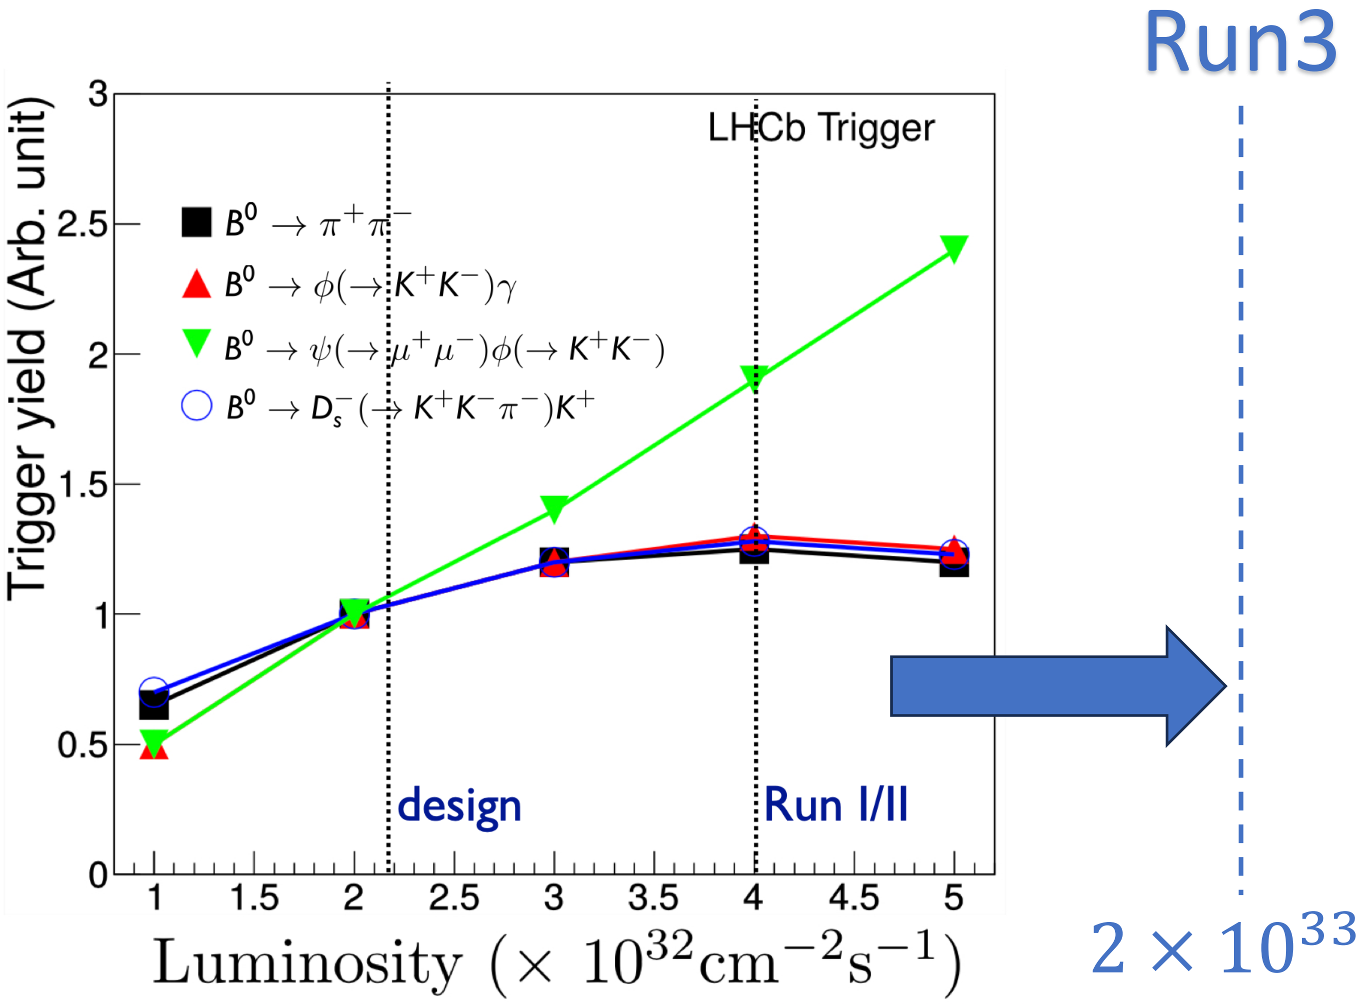

J. Phys.: Conf. Ser. **878** 012012

Trigger: Run2 vs Run3

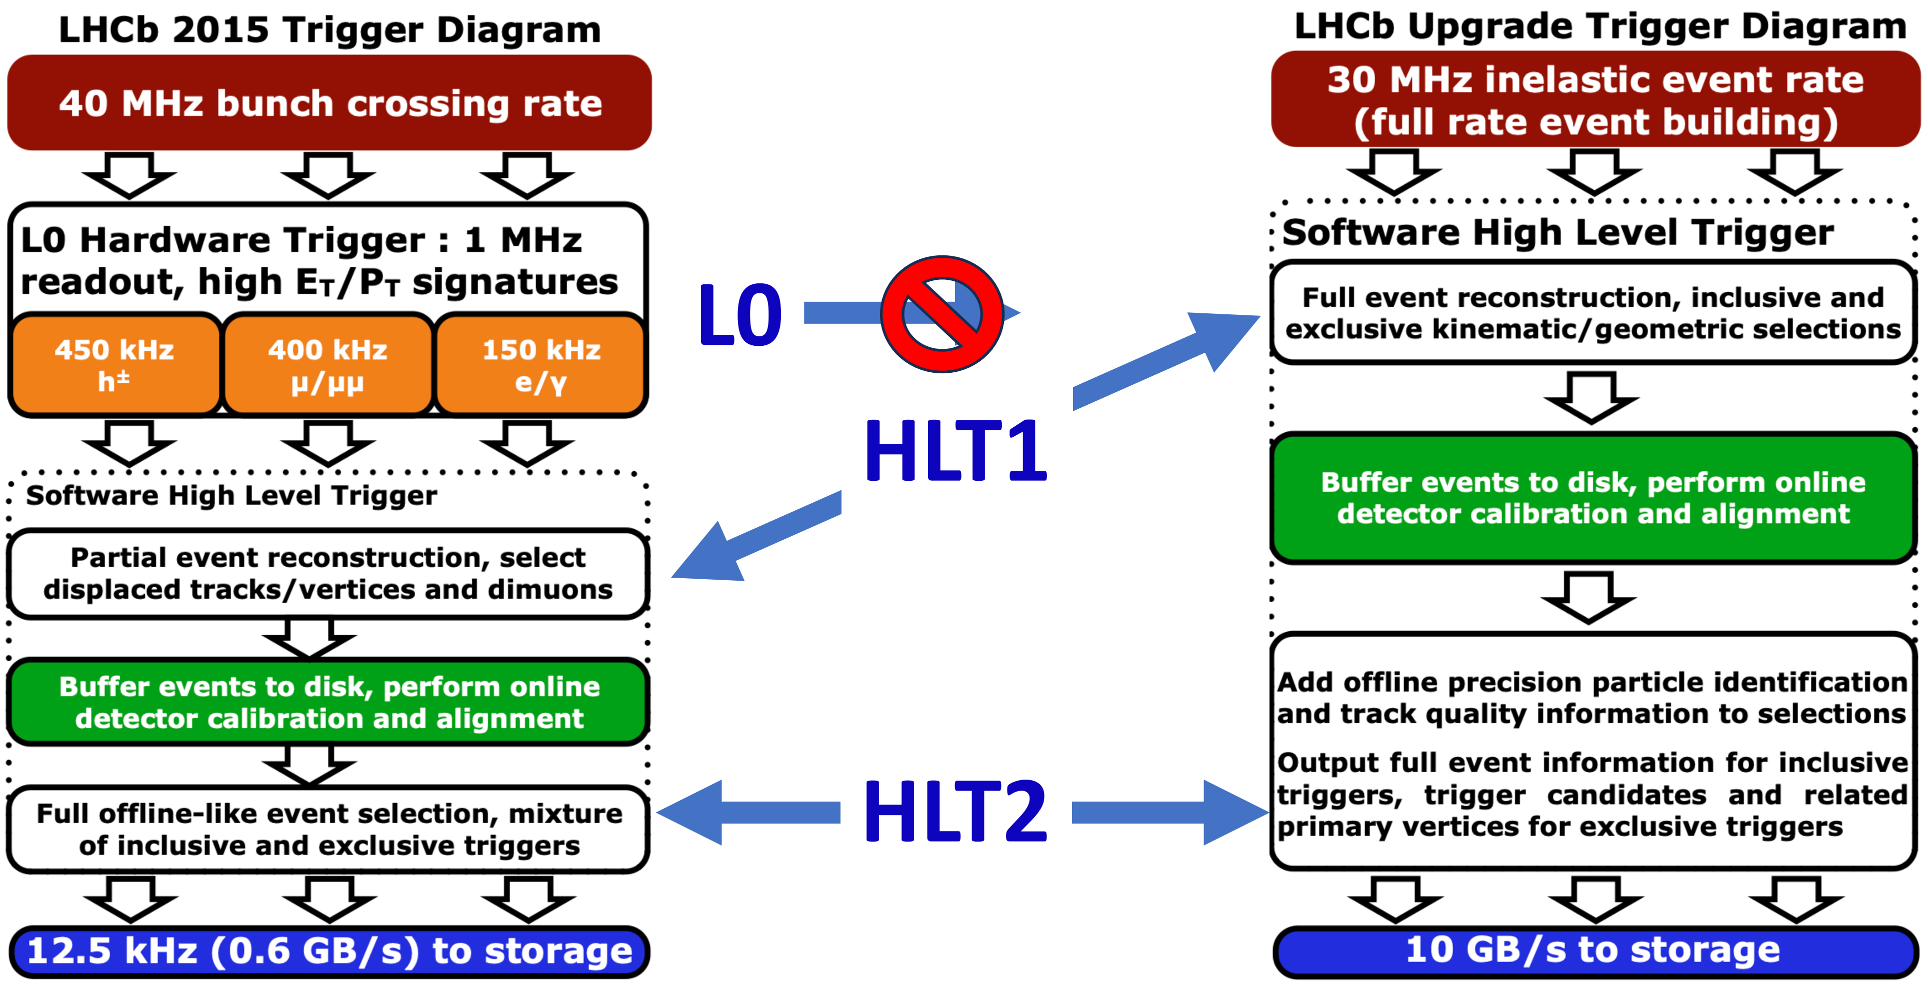

The Run3 LHCb dataflow

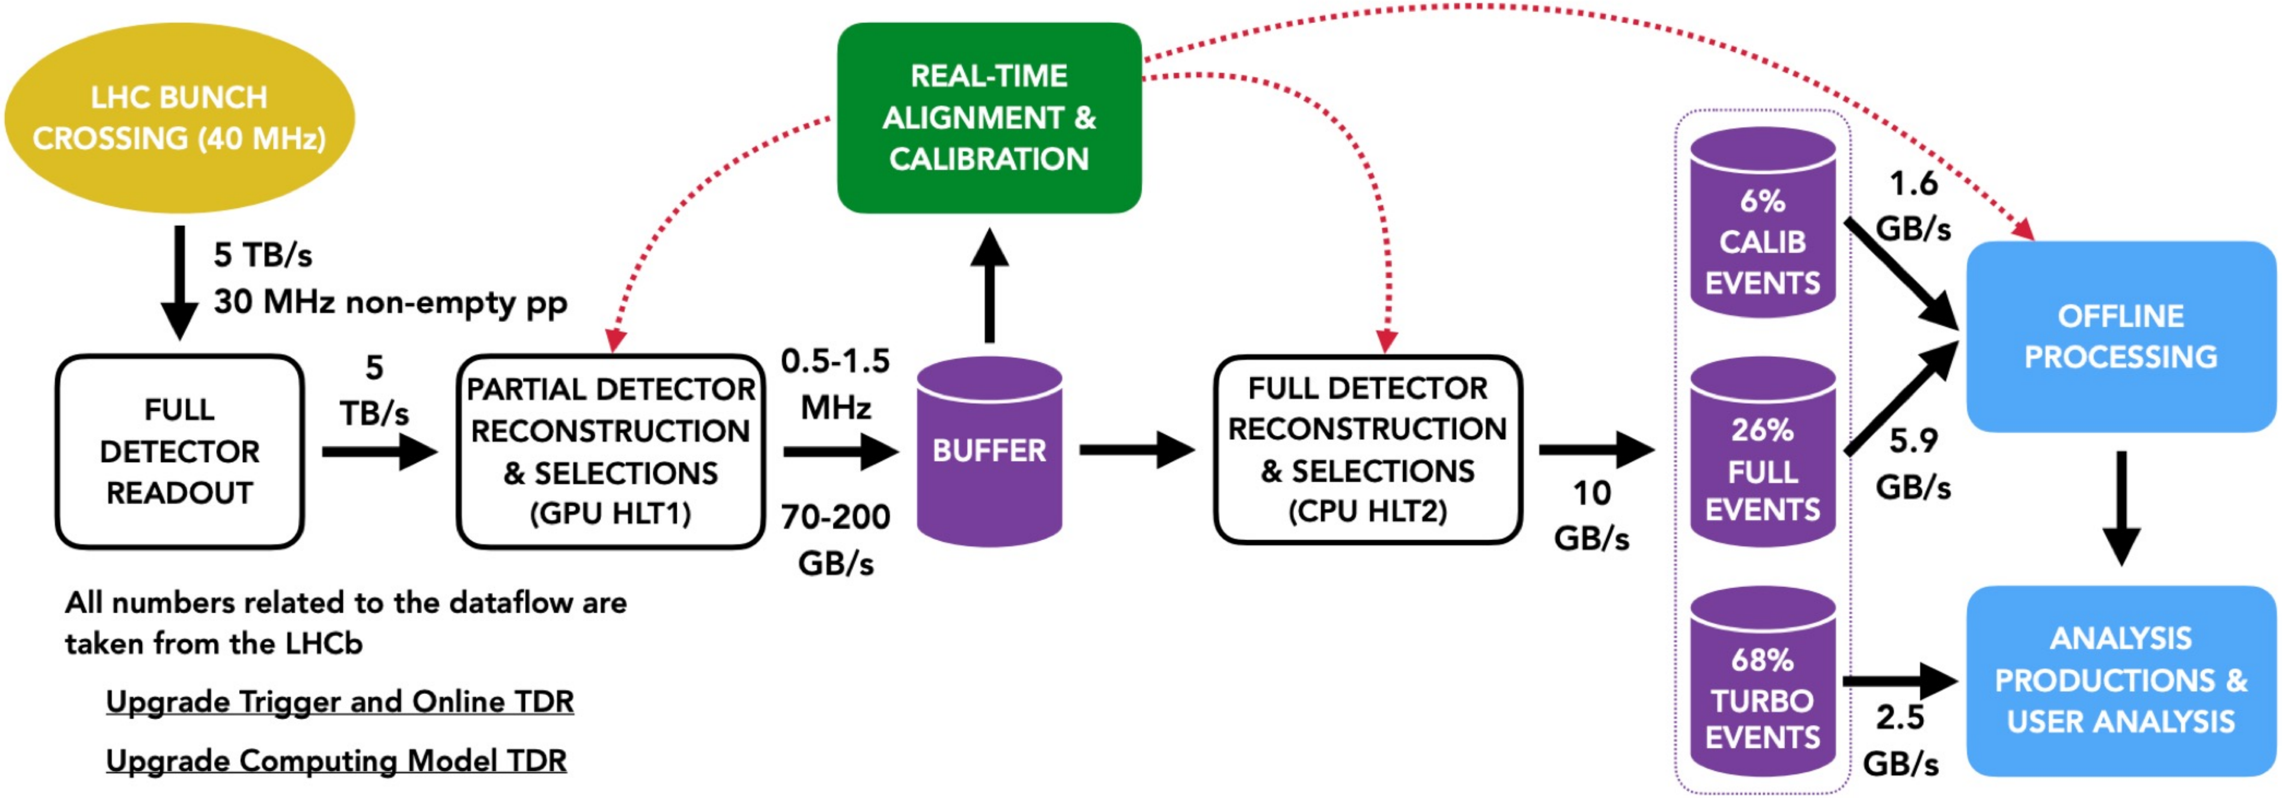

All numbers related to the dataflow are taken from the LHCb

- Upgrade Trigger and Online TDR
- Upgrade Computing Model TDR

LHCB-FIGURE-2020-016

Heterogeneous computing system

- Raw detector info is sent to the data processing center
- FPGA cards receive data at average 5 TB/s
- 163 Event Builder (EB) servers produce the packets of events
- Each EB server has 3 PCIe slots in which a GPU is installed
- Until August 2024: 2 GPUs per EB server = 326 Nvidia A5000 GPUs
- Now: 3rd GPU installed → 489 GPUs

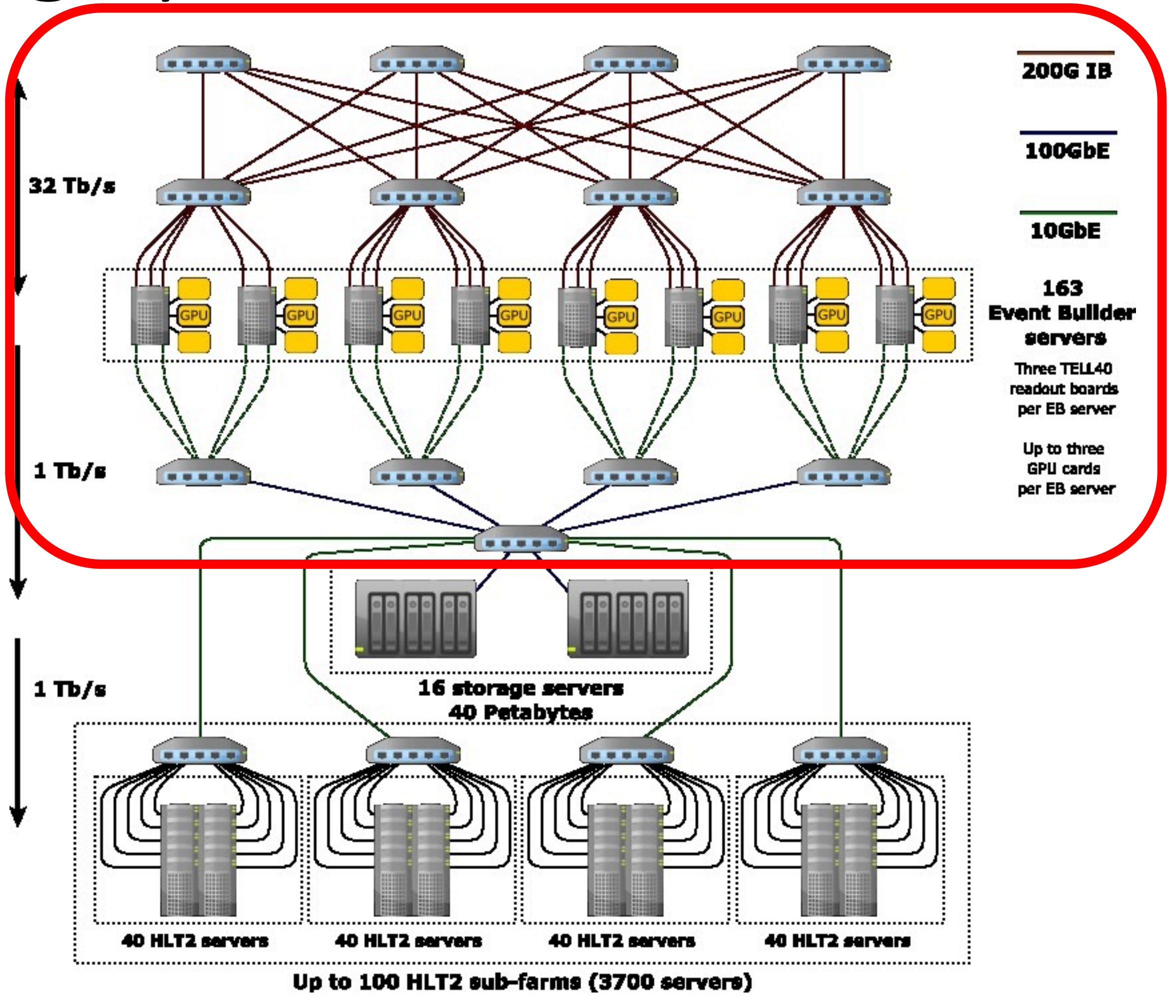

[Comput.Softw.Big Sci. 6 (2022) 1, 1]

Why GPUs?

- Zero overhead costs to install GPUs in EB servers
- Most of HLT1 tasks naturally lend themselves to very high degree of parallelism
- [Allen project]
- HLT1 throughput scales linearly with theoretical 32-bit TFLOPs performance → higher luminosity can be handled by more performant GPUs (no saturation expected)

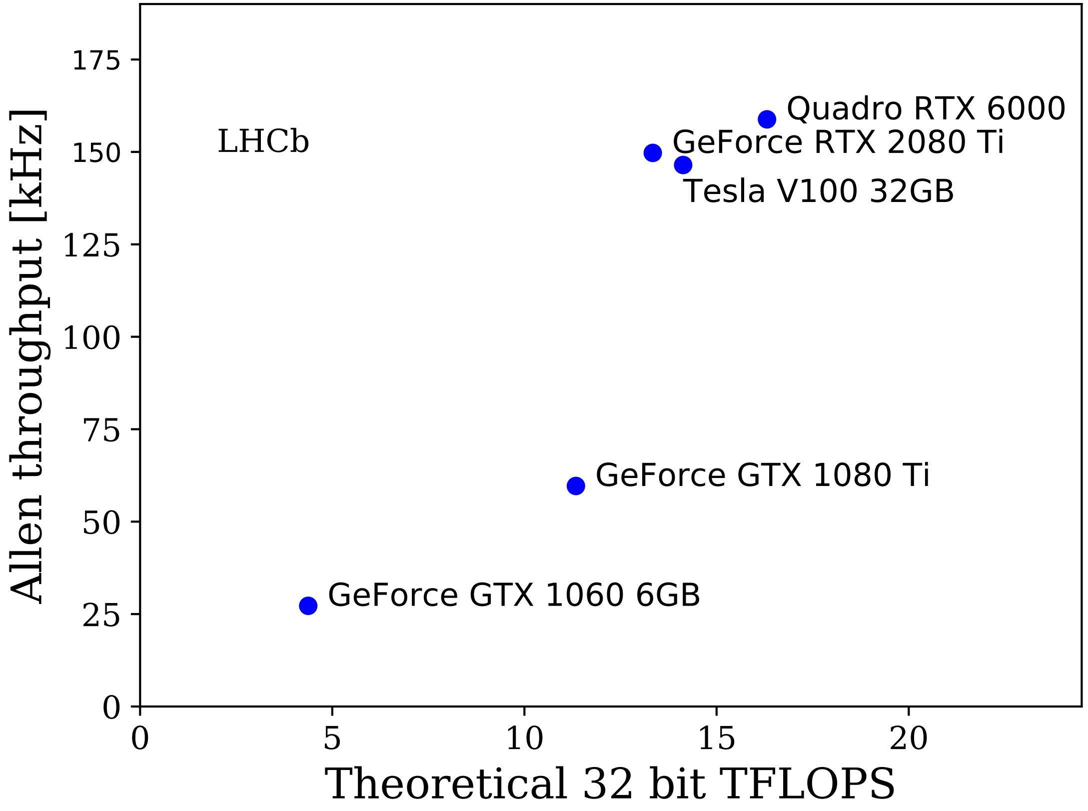

LHCB-FIGURE-2020-014

First-level trigger on GPUs

- Make full-event reconstruction respecting throughput constraints (30 MHz)
- Reconstruct tracks with good momentum resolution (< 1%) and vertexing

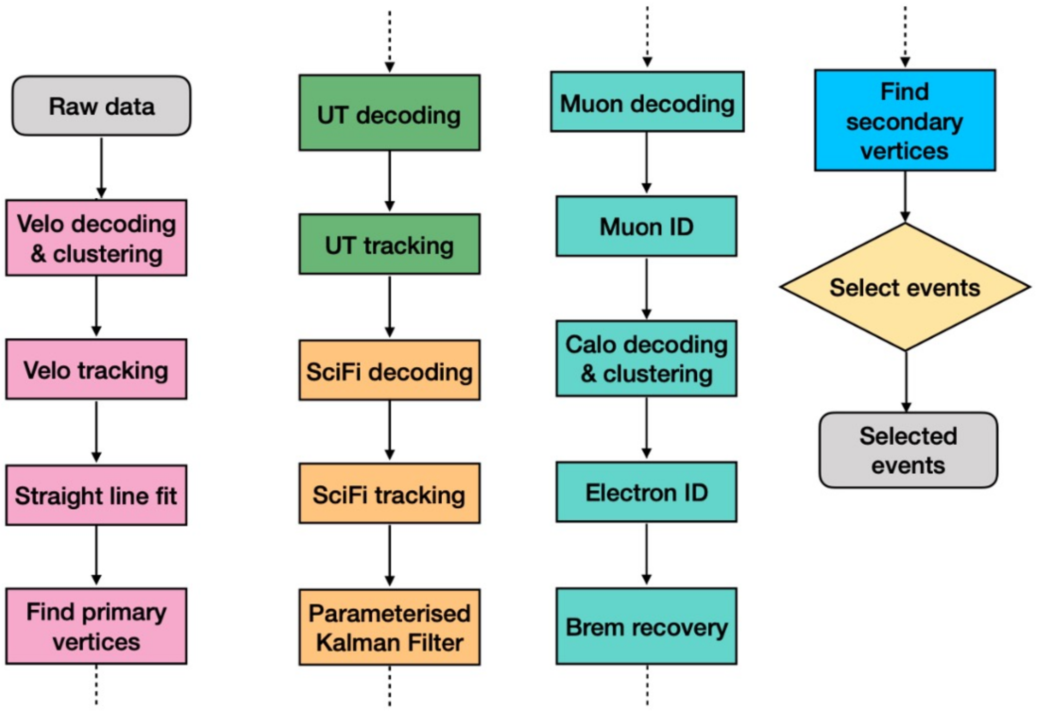

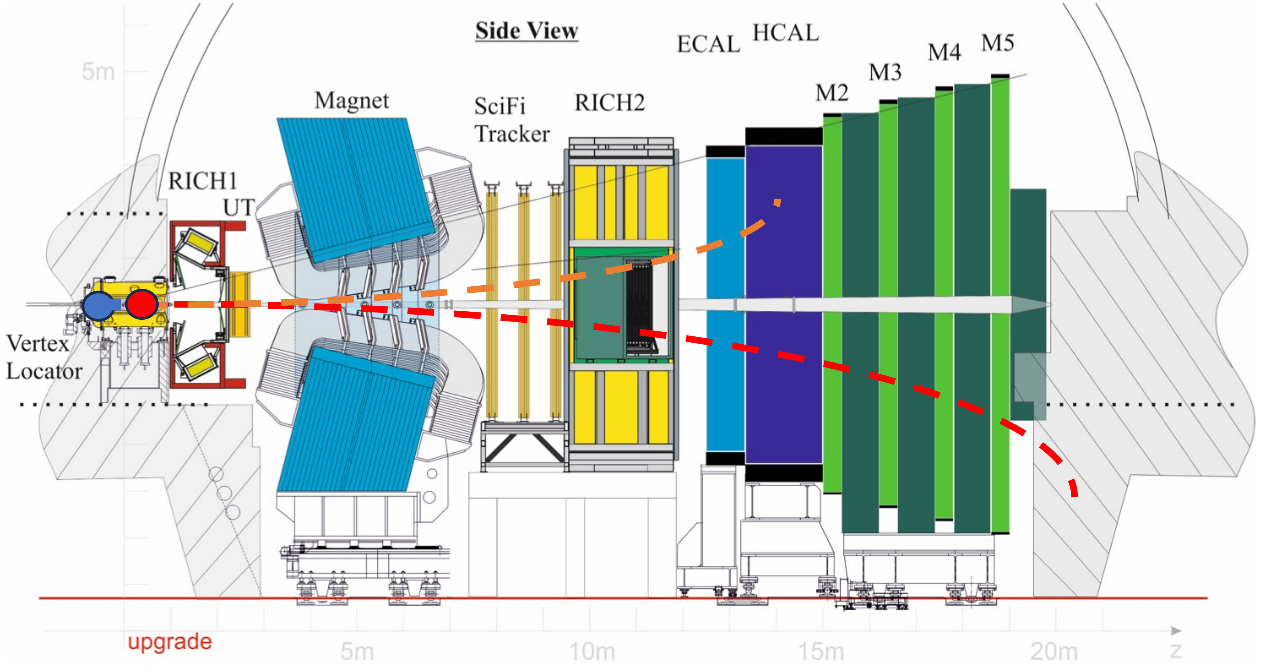

Reconstruction in GPU

- How to fully exploit parallelization power of GPUs?
- Parallelization levels when reconstructing tracks traversing the whole LHCb detector:
  1. Over events, independent p-p collisions
  2. Over input tracks, extrapolate straight tracks in VELO+UT into the magnetic field reaching the SciFi
  3. Over hits in SciFi, meaning possible extrapolations segments

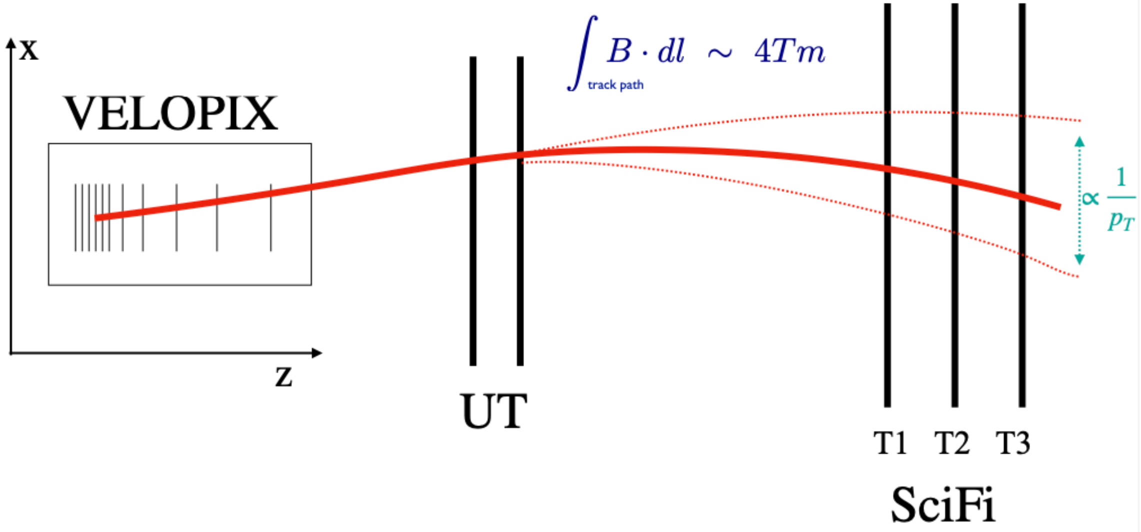

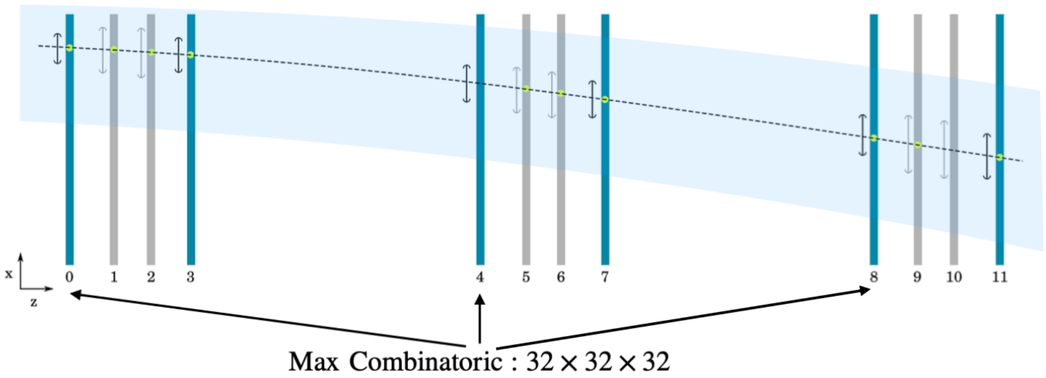

IEEE Access 12 (2024) 114198-114211

HLT1 performance on 2024 data

- Exploiting reconstructed objects to select decays of interest
- Looking for displaced signatures with high transverse momentum
- Comparison with Run2 trigger efficiencies, limited by L0 selections
- Clear gain at low momentum for hadronic and electronic B-mesons modes

Muonic

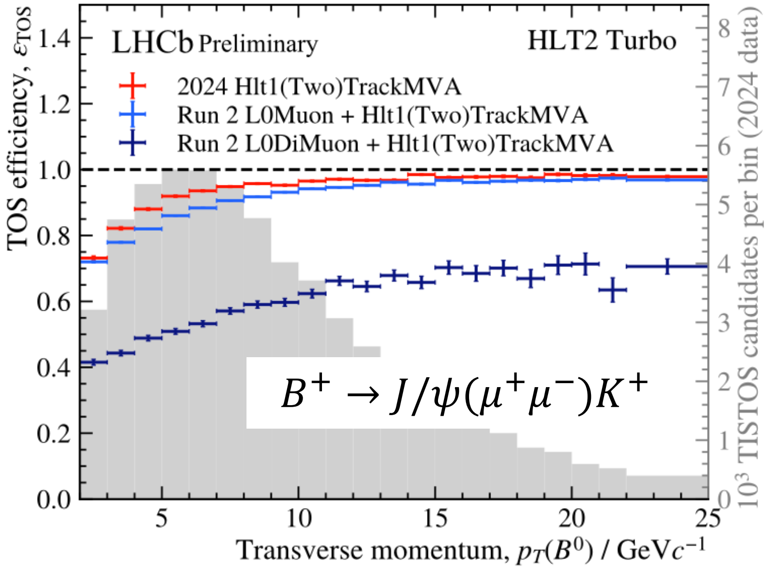

Hadronic

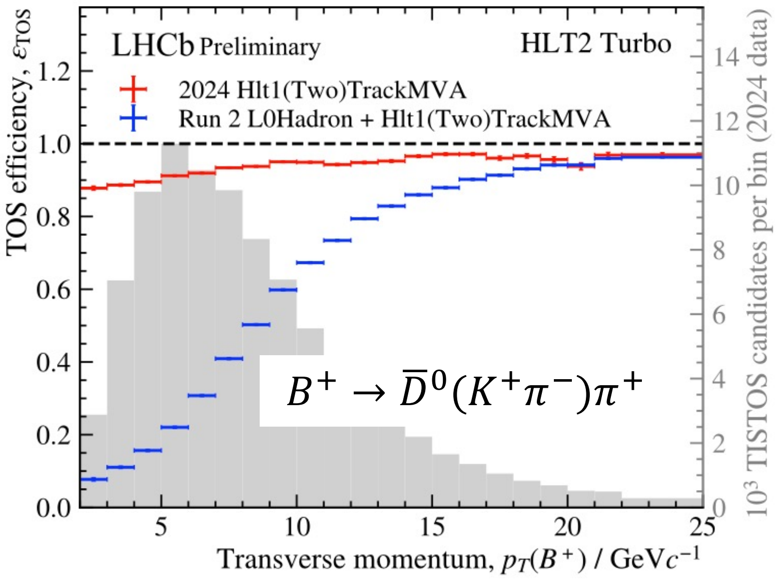

Electronic

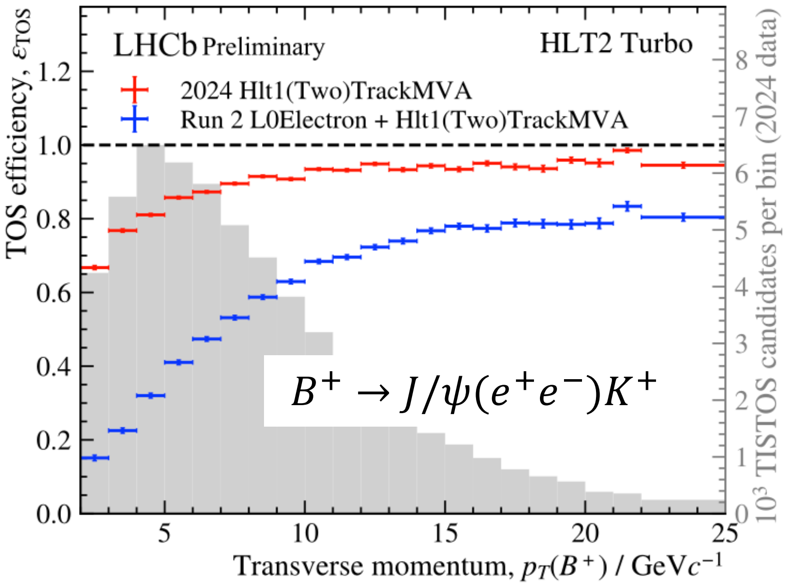

Alignment & Calibration

- HLT1 sends data at a rate of 1 Tb/s to a 40 Petabytes storage
- While data is stored, perform full detector alignment and calibration

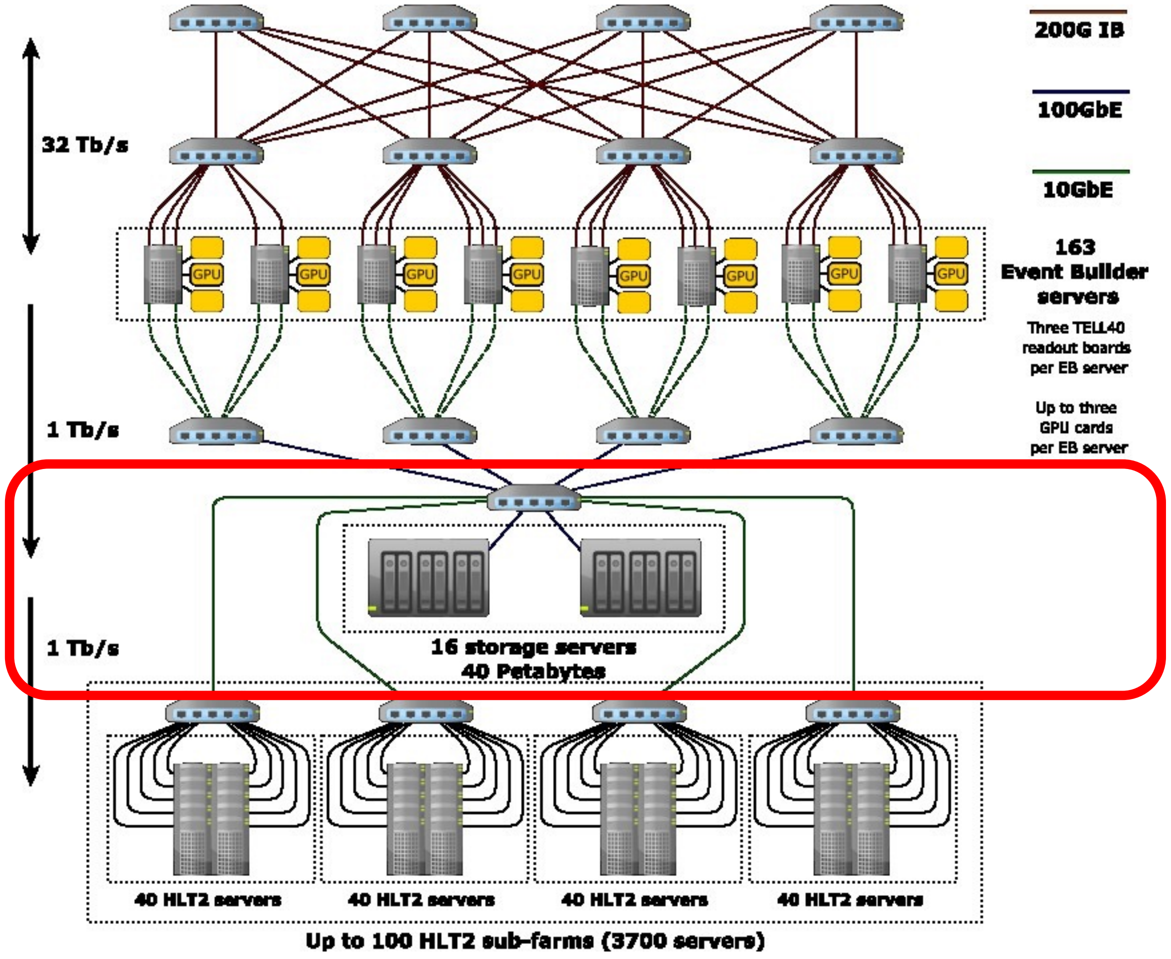

[Comput.Softw.Big Sci. 6 (2022) 1, 1]

Alignment & Calibration

- Critical to ensure offline-like quality of HLT2 reconstruction
- Alignement of the tracking detectors, muon chambers and RICH mirrors
- Calibration of RICH detectors and calorimeters
- Performed at each LHC fill or more frequently

More details in Miguel's talk

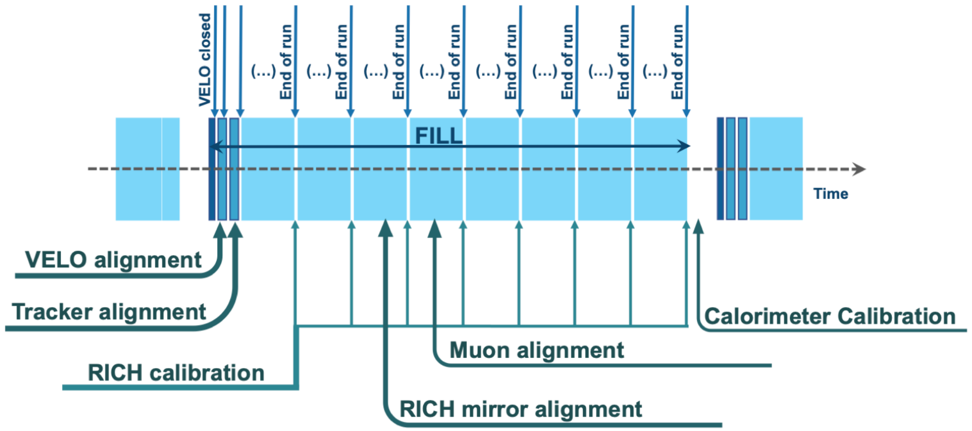

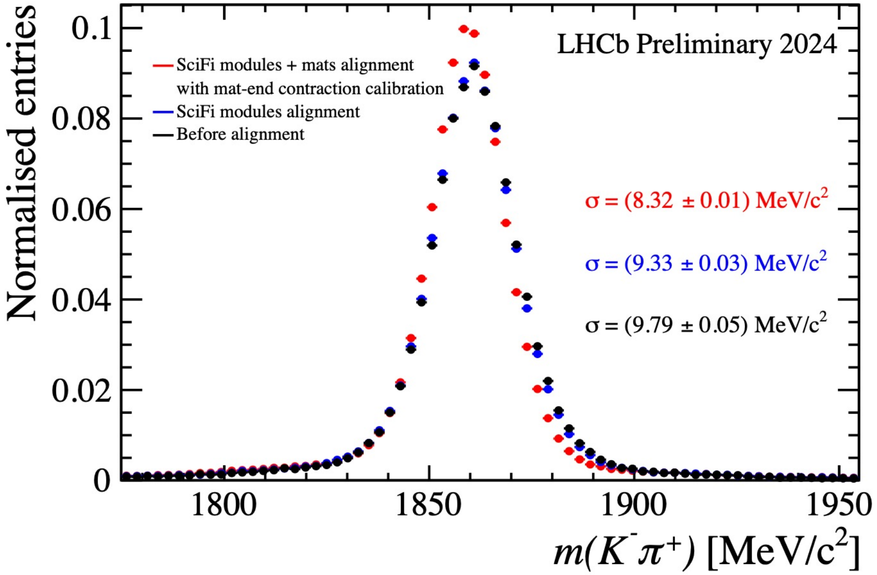

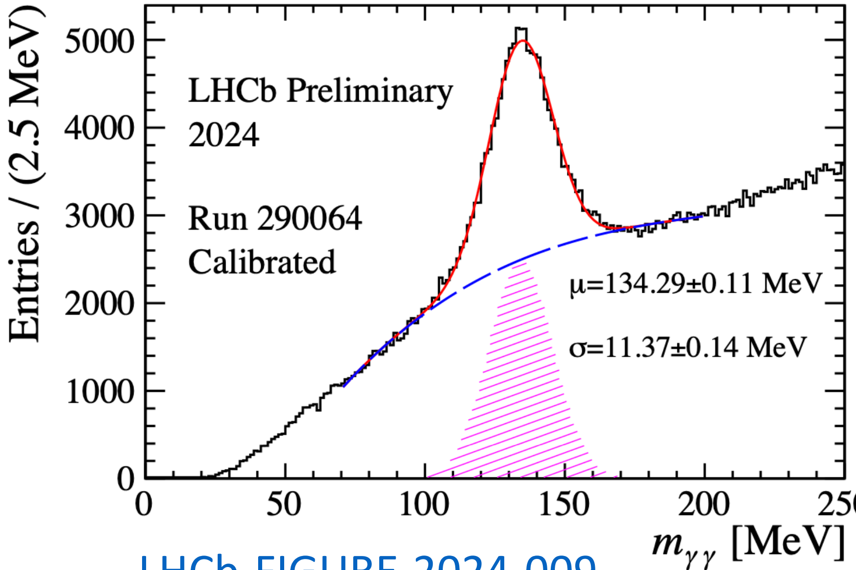

LHCb-FIGURE-2024-009

The second high-level trigger: HLT2

- HLT2 can be run asynchronously to HLT1 once the full alignment&calibration is performed
- HLT2 needs to process data at rate greater than half of the HLT1 output (1 MHz): minimum 500 kHz

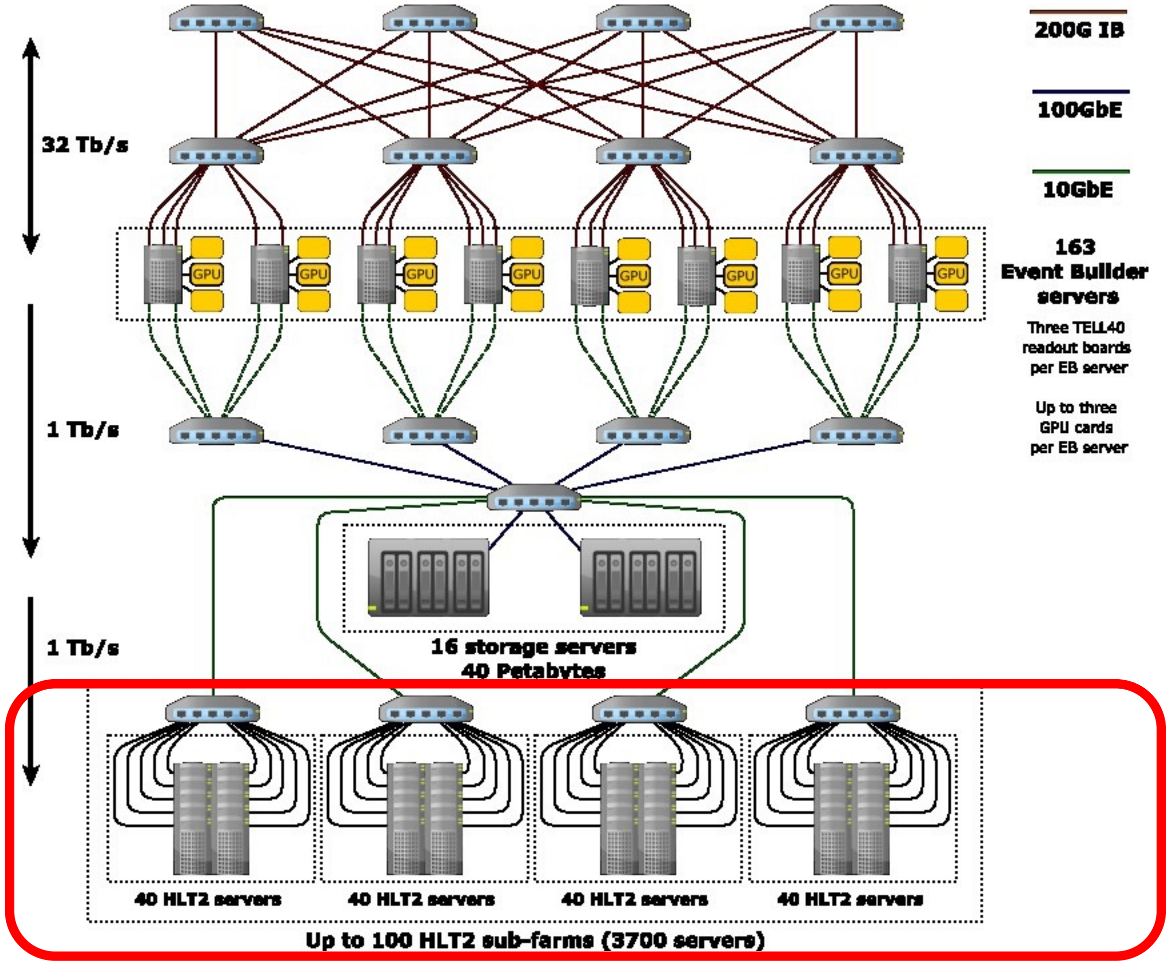

[Comput.Softw.Big Sci. 6 (2022) 1, 1]

The second high-level trigger

- Full event reconstruction exploiting best alignement&calibration and PID info
- Dedicated trigger selections O(2700) covering the broad LHCb physics program
- Inclusive (focused on heavy-flavour decay signatures) and exclusive (specific decays) selections
- Running at a minimum of 500 kHz throughput with limited bandwidth of 10 GB/s
- How do we do it?

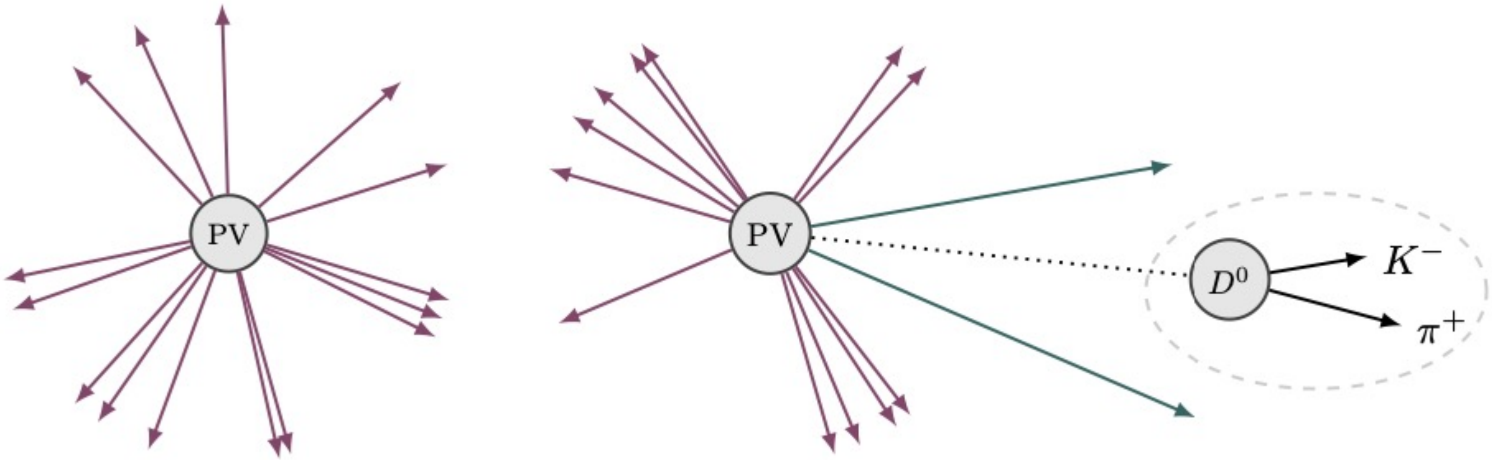

HLT2 throughput improvements

- Exploiting Structure-of-Arrays collections for faster data access using vectorisation and multi-threading: access in loops same information in different data, Single Instruction Multiple Data (SIMD)
- Throughput-Oriented selections: Thor functors
  - Functors built at compile time into cache memory
  - Agnostic on I/O type such that they can be combined e.g. vertex().position().x()

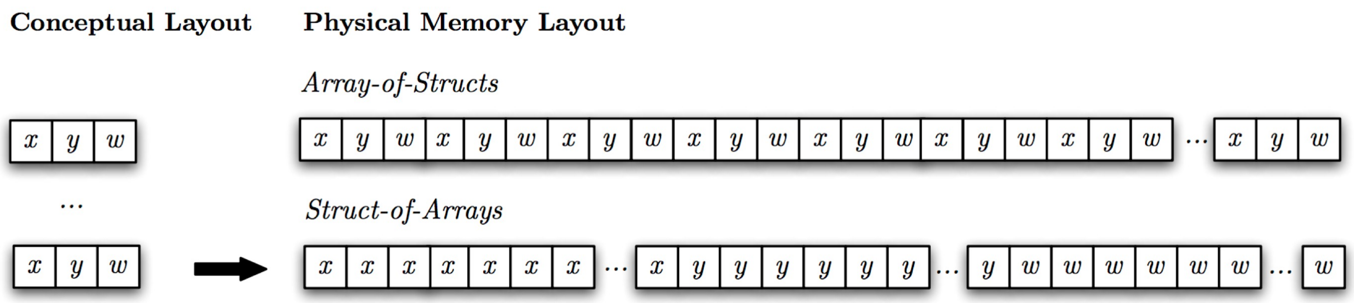

Implementation	$D^+ \to K^+\pi^+\pi^-$ execution time
CombineParticles	256 µs
NBodyDecays	77.1 µs
ThorParticleCombiner	38.8 µs
ThOrCombiner Scalar	10.2 µs
ThOrCombiner SSE	7.5 µs
ThOrCombiner AVX2	6.9 µs

CERN-THESIS-2020-331

HLT2 bandwidth

JINST 14 (2019) P04006

- Limited bandwidth of 10 GB/s
- Do we need to save the whole event?
- Different options:
  - Turbo: saving only info related to the signal candidate
  - Selective persistency: save additional objects relative to the signal
  - Full persistency
- About 70% of the events are saved in the Turbo stream reducing the event size by a factor 10

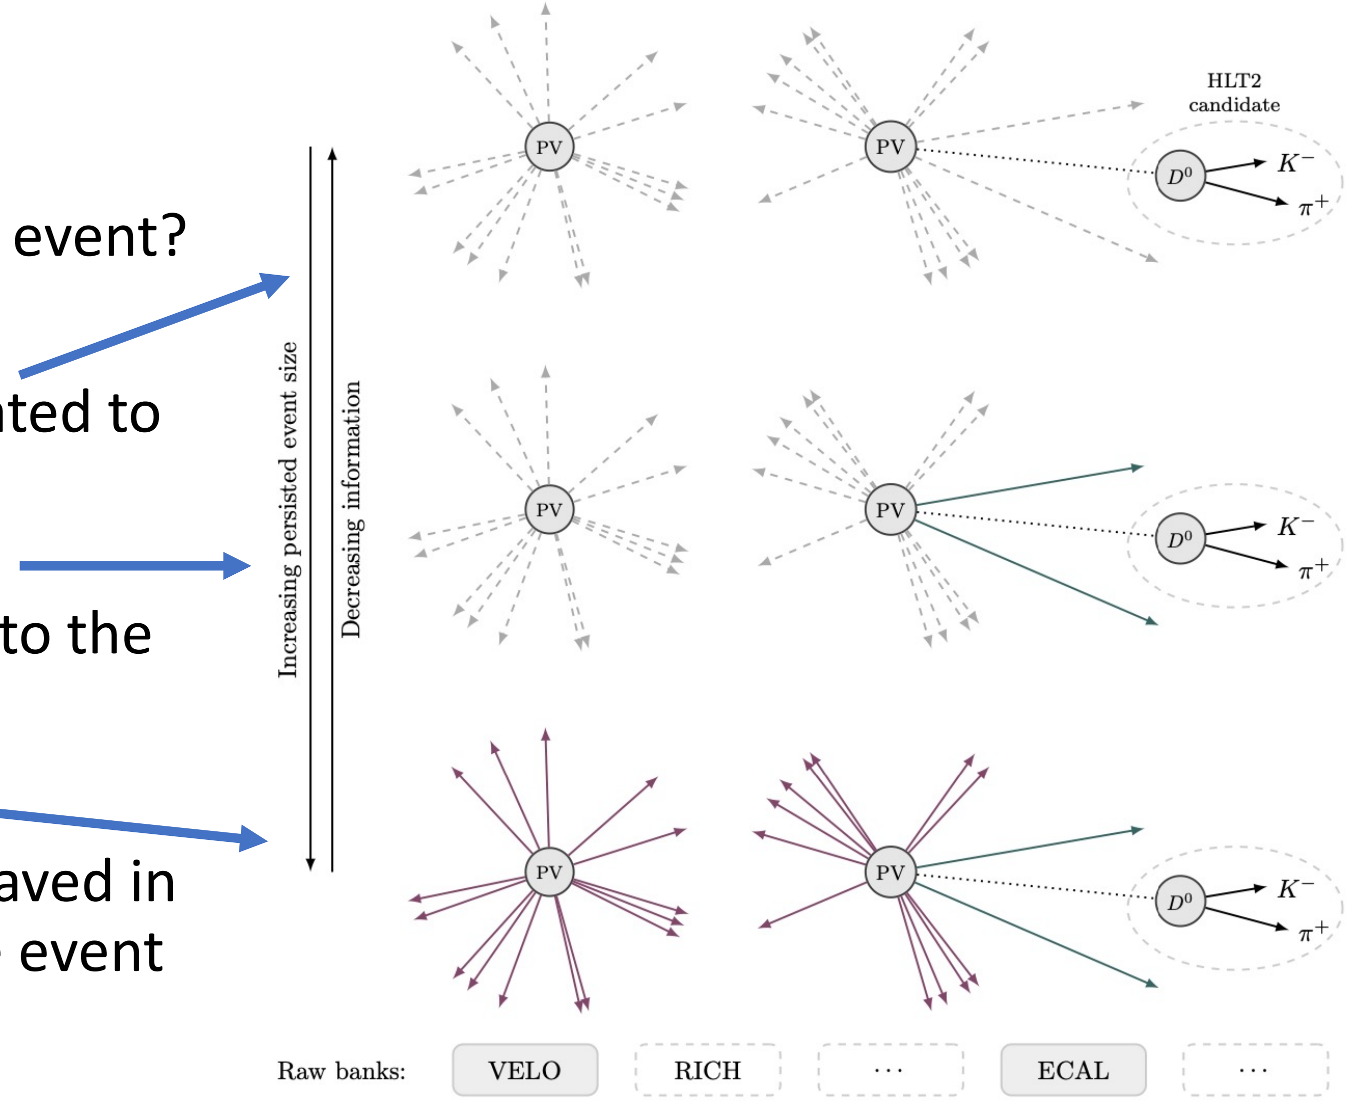

The second high-level trigger: HLT2

- During summer 2024:
  - 4400 CPUs replaced going from 8 to 14 cores
  - 204 HLT2 new servers installed with each 128 cores
- Minimum of 500 kHz HLT2 throughput … we reached 900kHz !
- Keeping the bandwidth around 9-10 GB/s!

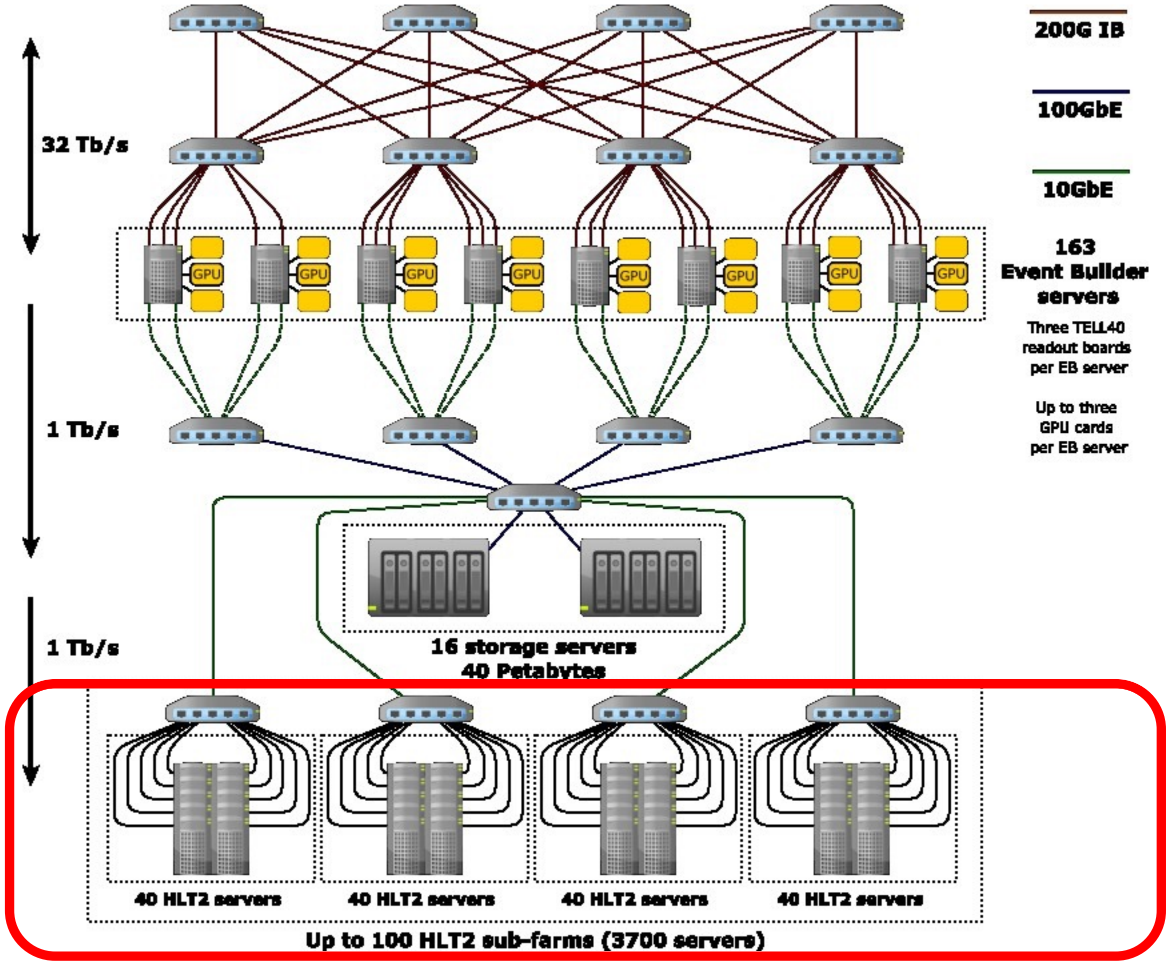

[Comput.Softw.Big Sci. 6 (2022) 1, 1]

Performance of HLT2 on 2024 data

- Achieving excellent vertex resolutions, good track reconstruction and stable PID performance for muons, hadrons and electrons

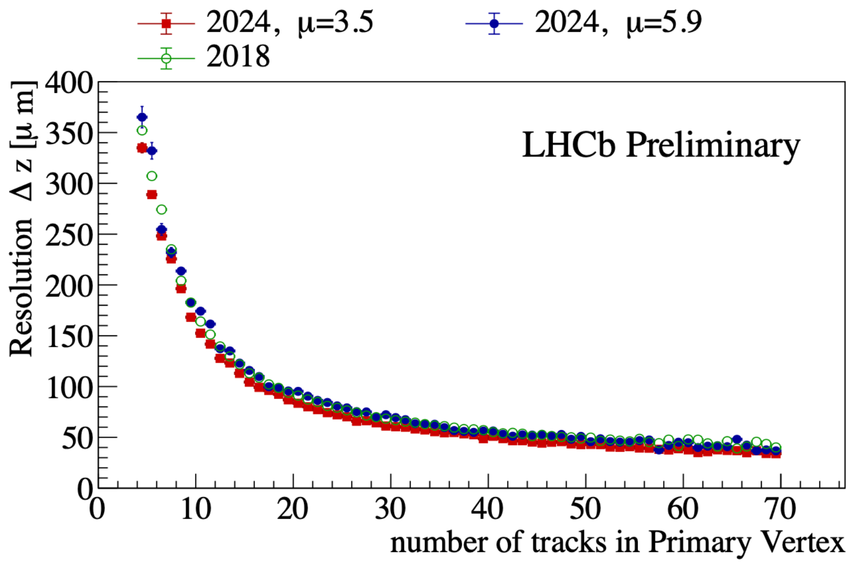

LHCB-FIGURE-2024-011

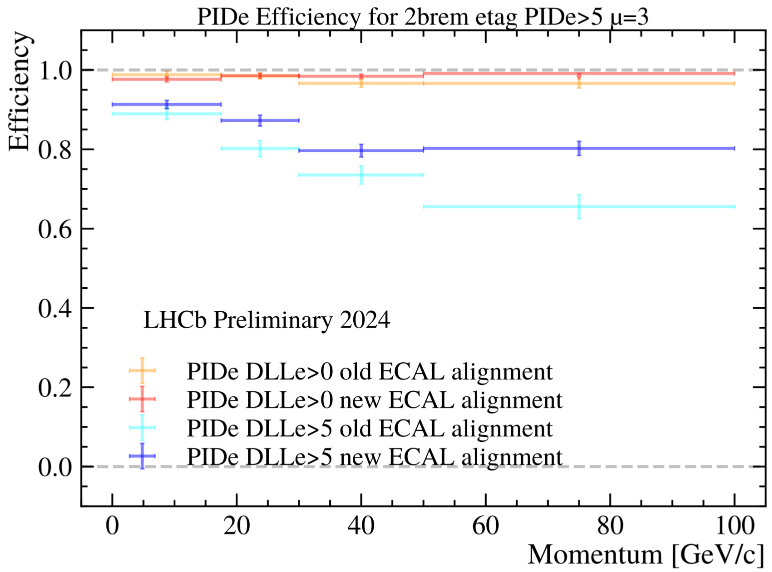

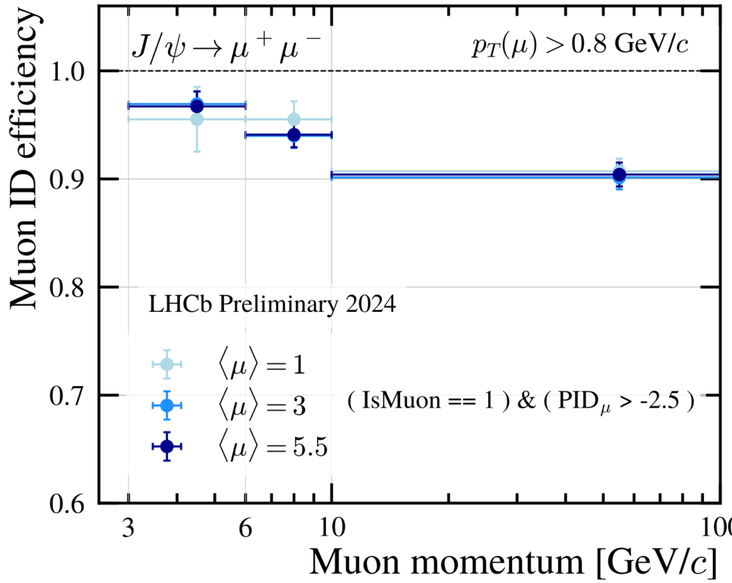

LHCB-FIGURE-2024-010

Performance of HLT2 on 2024 data

- Achieving excellent vertex resolutions, good track reconstruction and stable PID performance for muons, hadrons and electrons
- Leading to reconstruction and selection of decays of interest

$B^0 \to D^-(K^+\pi^-\pi^-)\pi^+$

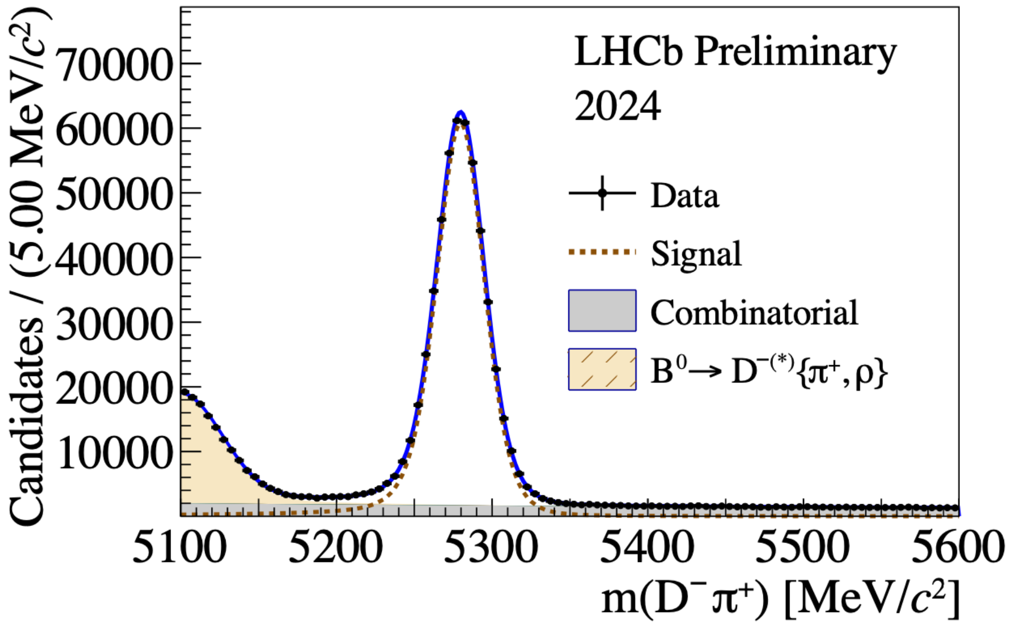

LHCB-FIGURE-2024-021

$B^+ \to K^+\mu^+\mu^-$

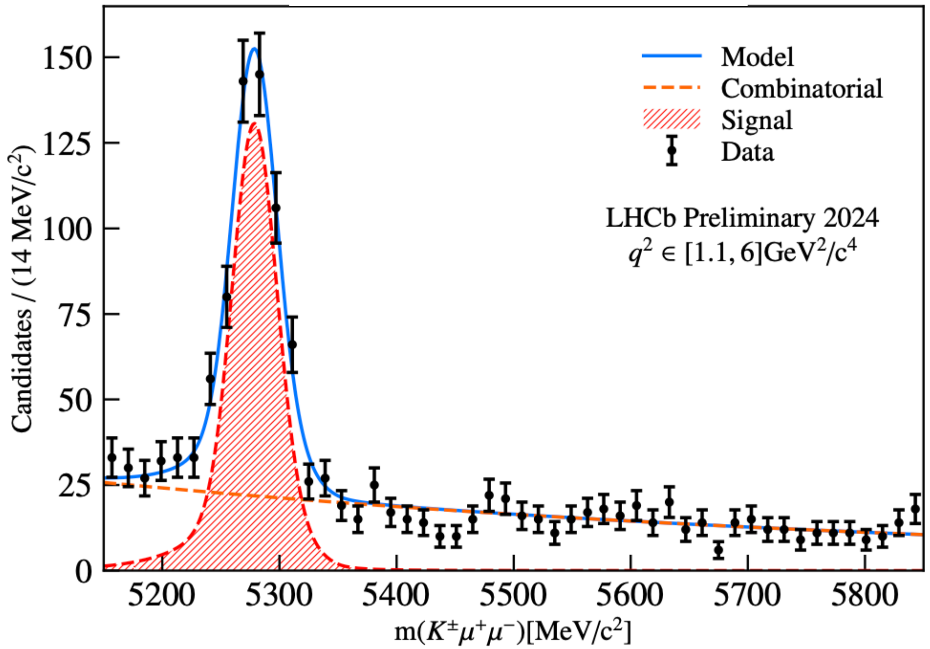

LHCB-FIGURE-2024-021

Very successful 2024 data taking year

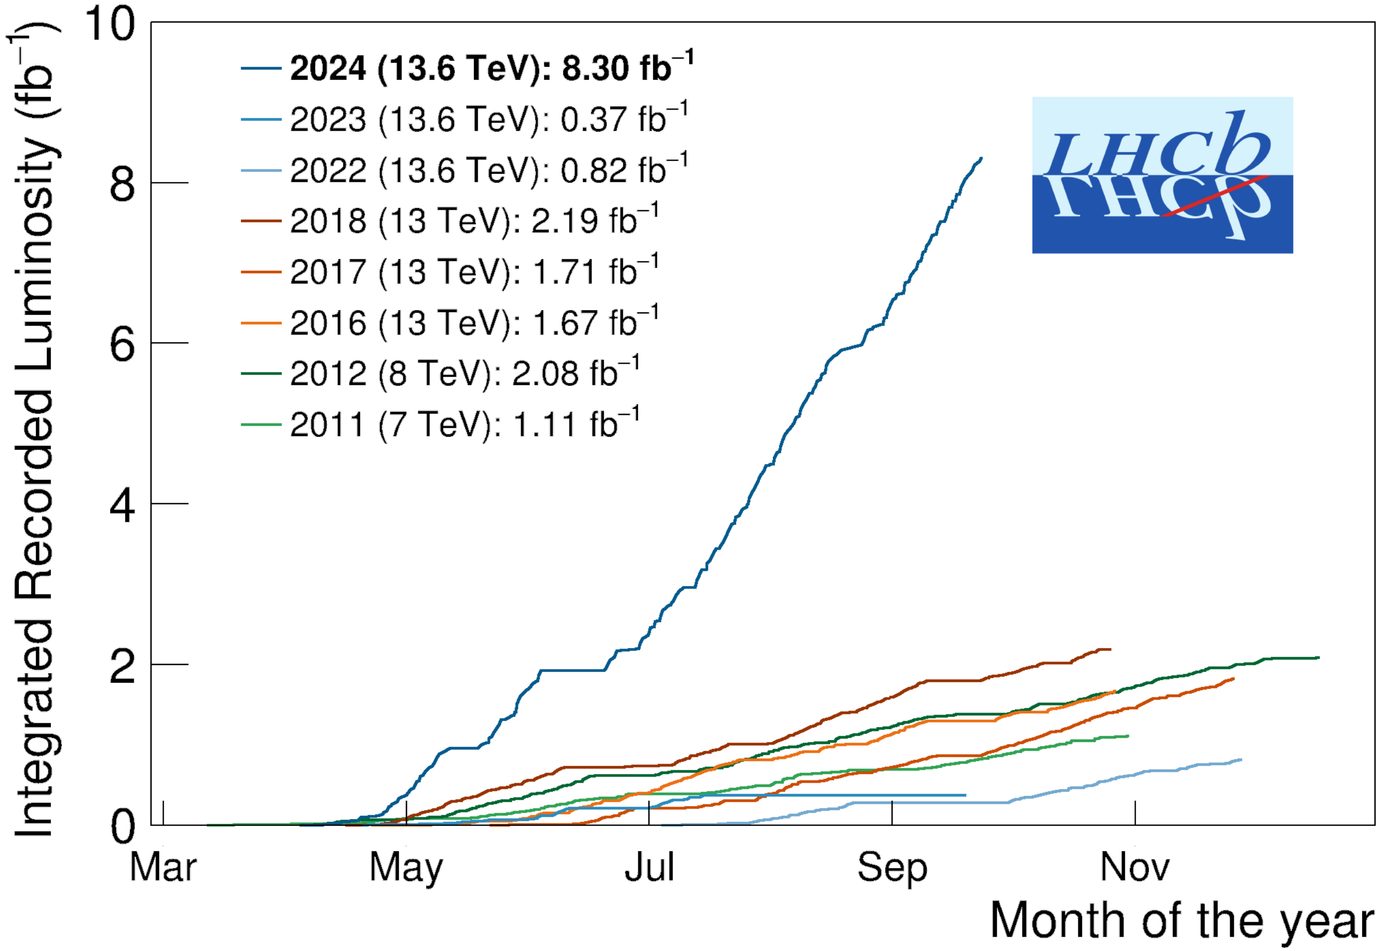

Plans for the future

- Now that the 3rd GPU is installed, we can exploit the extra throughput to include new features in HLT1:
  - Near future: downstream tracking ready to be included in the next weeks with dedicated selections for long-lived particles
  - Work in progress for next years:
    - Include full Kalman filter to improve momentum resolution and ghost rejection
    - Include RICH reconstruction → distinguish pion, kaon, protons already at HLT1
- Far future (probably Upgrade2 from 2032):
  - Implement HLT2 in GPUs : porting HLT2 algorithms from CPU to GPU
  - Use the same reconstruction algorithms in HLT1 and HLT2 with different requirements based on throughput and selections

Conclusions

- LHCb taking data with a fully-software trigger: successfully in Run3!
- First trigger stage (HLT1) optimised on GPU and dealing with 30 MHz LHC input rate
- Performing alignment&calibration before running second stage (HLT2)
- HLT2 performs offline-like reconstruction including PID information
- HLT1, alignement and HLT2 achieve expected performance during 2024 data taking (more than $8\ fb^{-1}$ collected up to now)
- Many ideas for the future to improve our trigger system, also exploiting our knowledge on heterogenous architectures

Backup