Lukas Calefice
On behalf of the
LHCb Collaboration

Experience and results from running the LHCb trigger system at 30MHz

19.07.2024,
ICHEP,
Prague

Title slide inspired by

Universitat de Barcelona

erc European Research Council
Supporting top researchers from anywhere in the world

LHCb RTA

LHCb

LHCb Run 3 detector (after Upgrade I)

[arXiv:2305.10515]

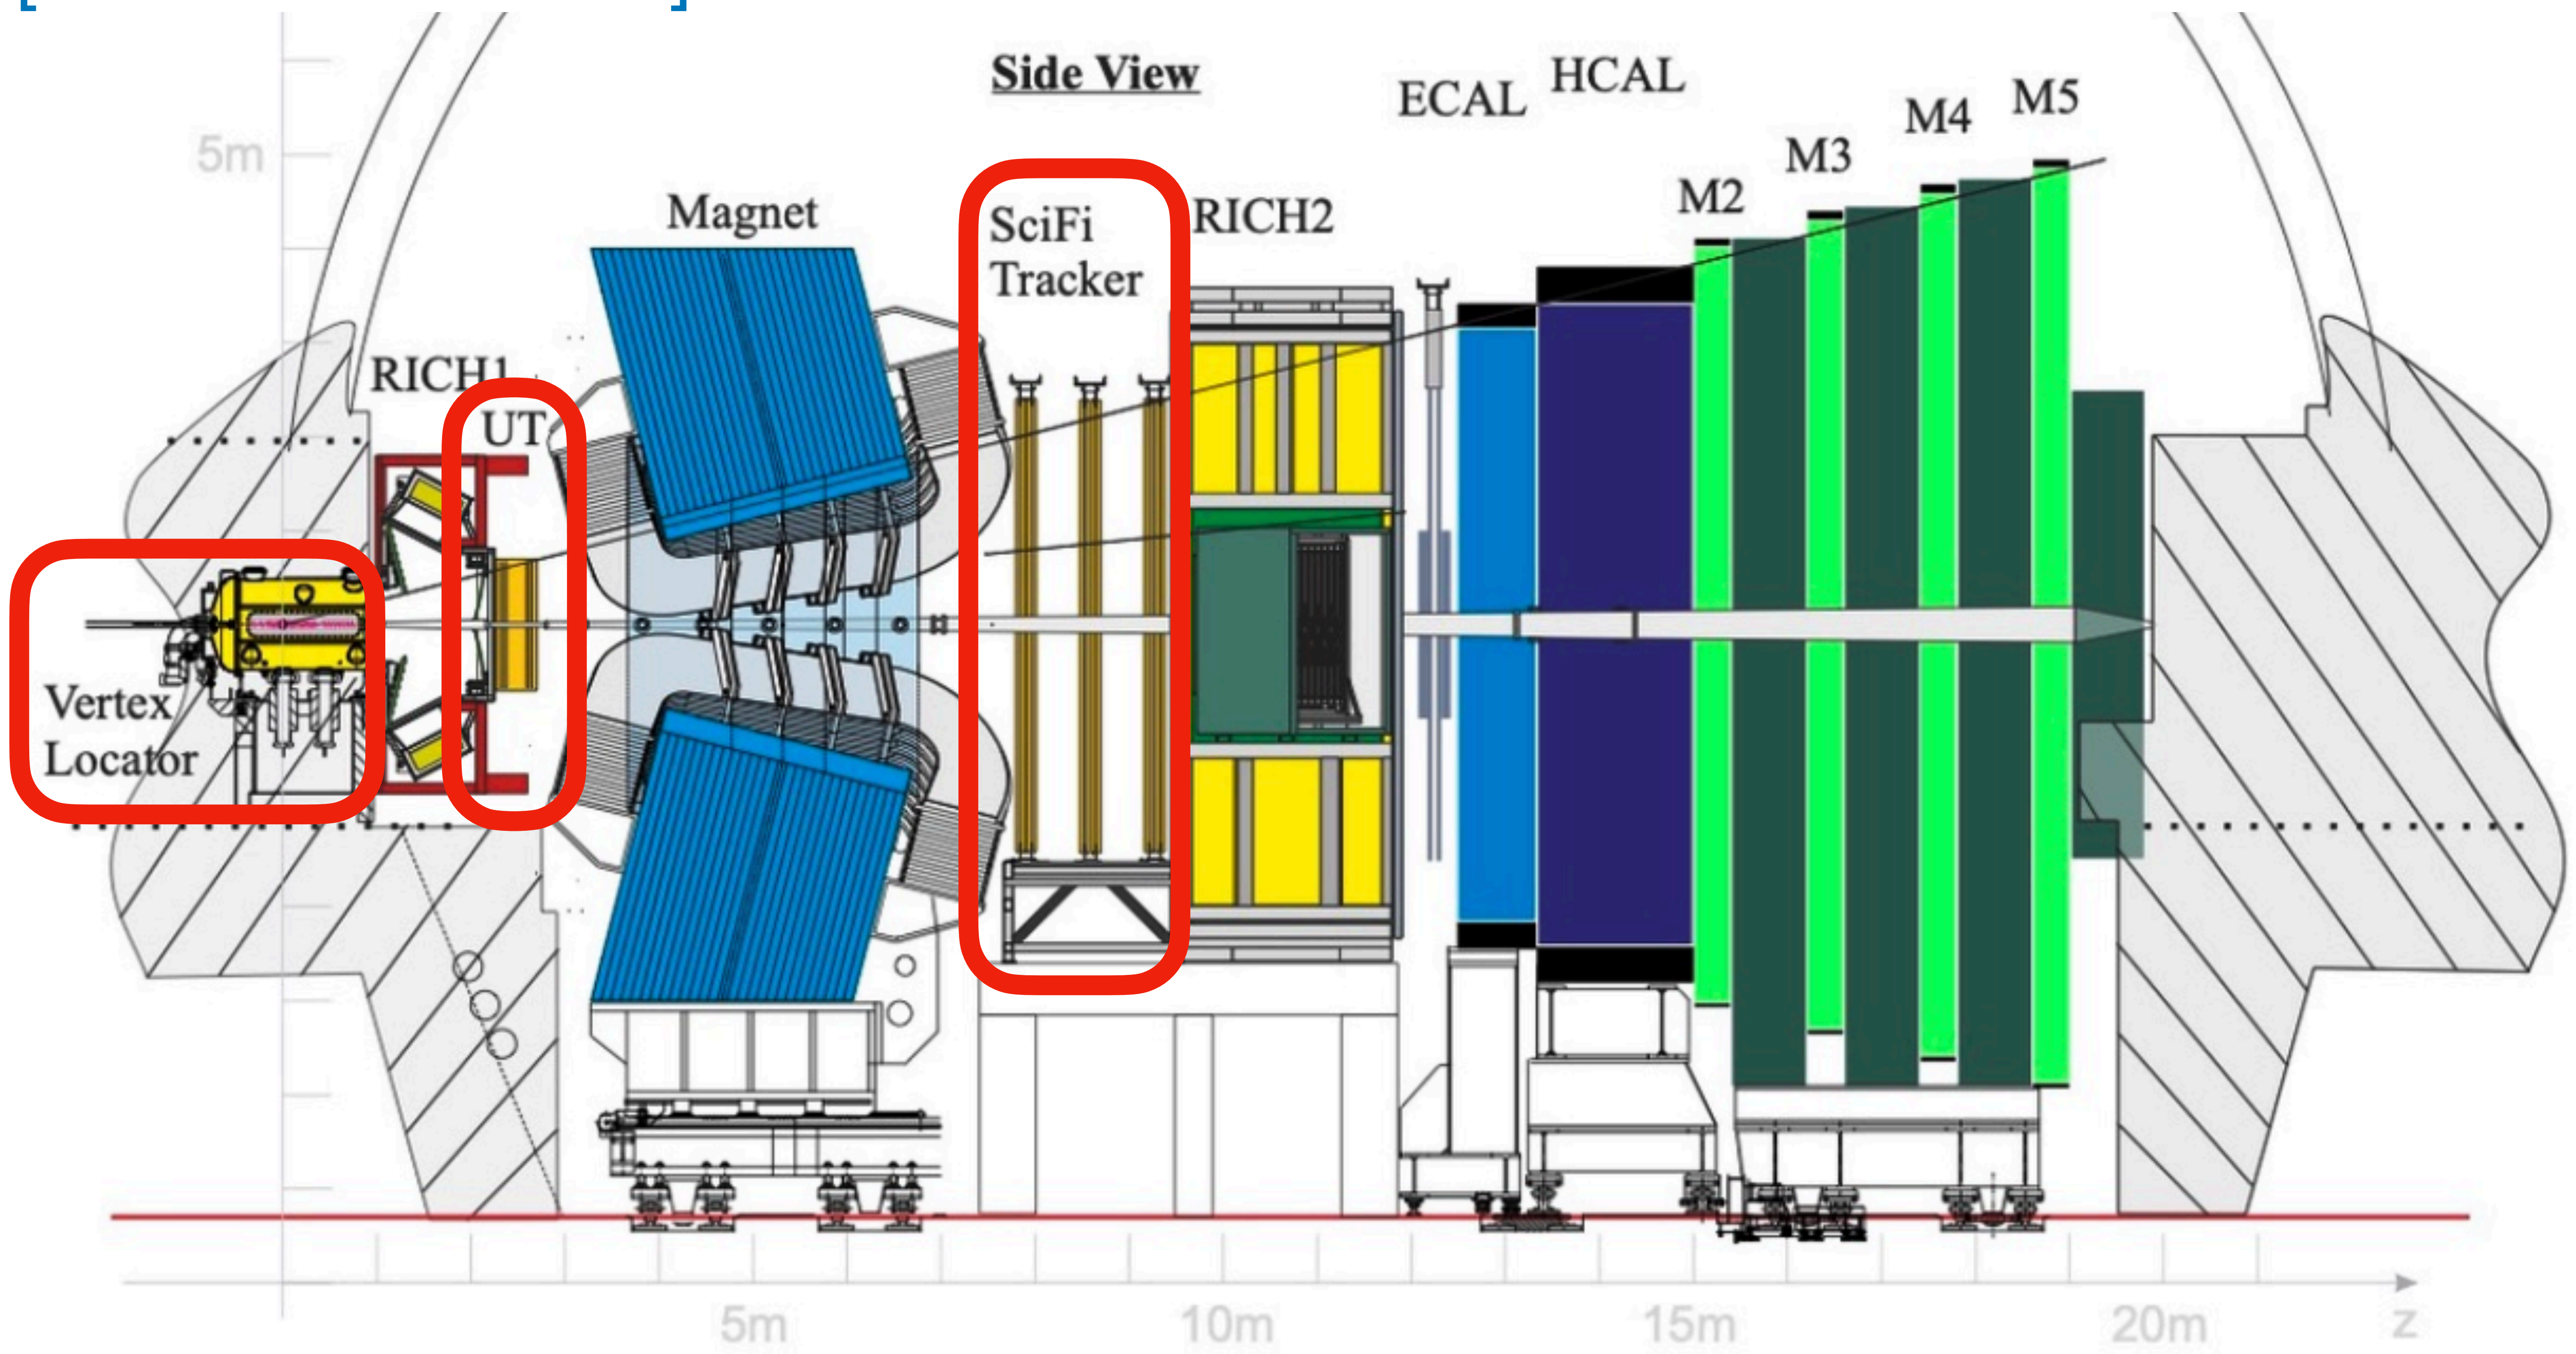

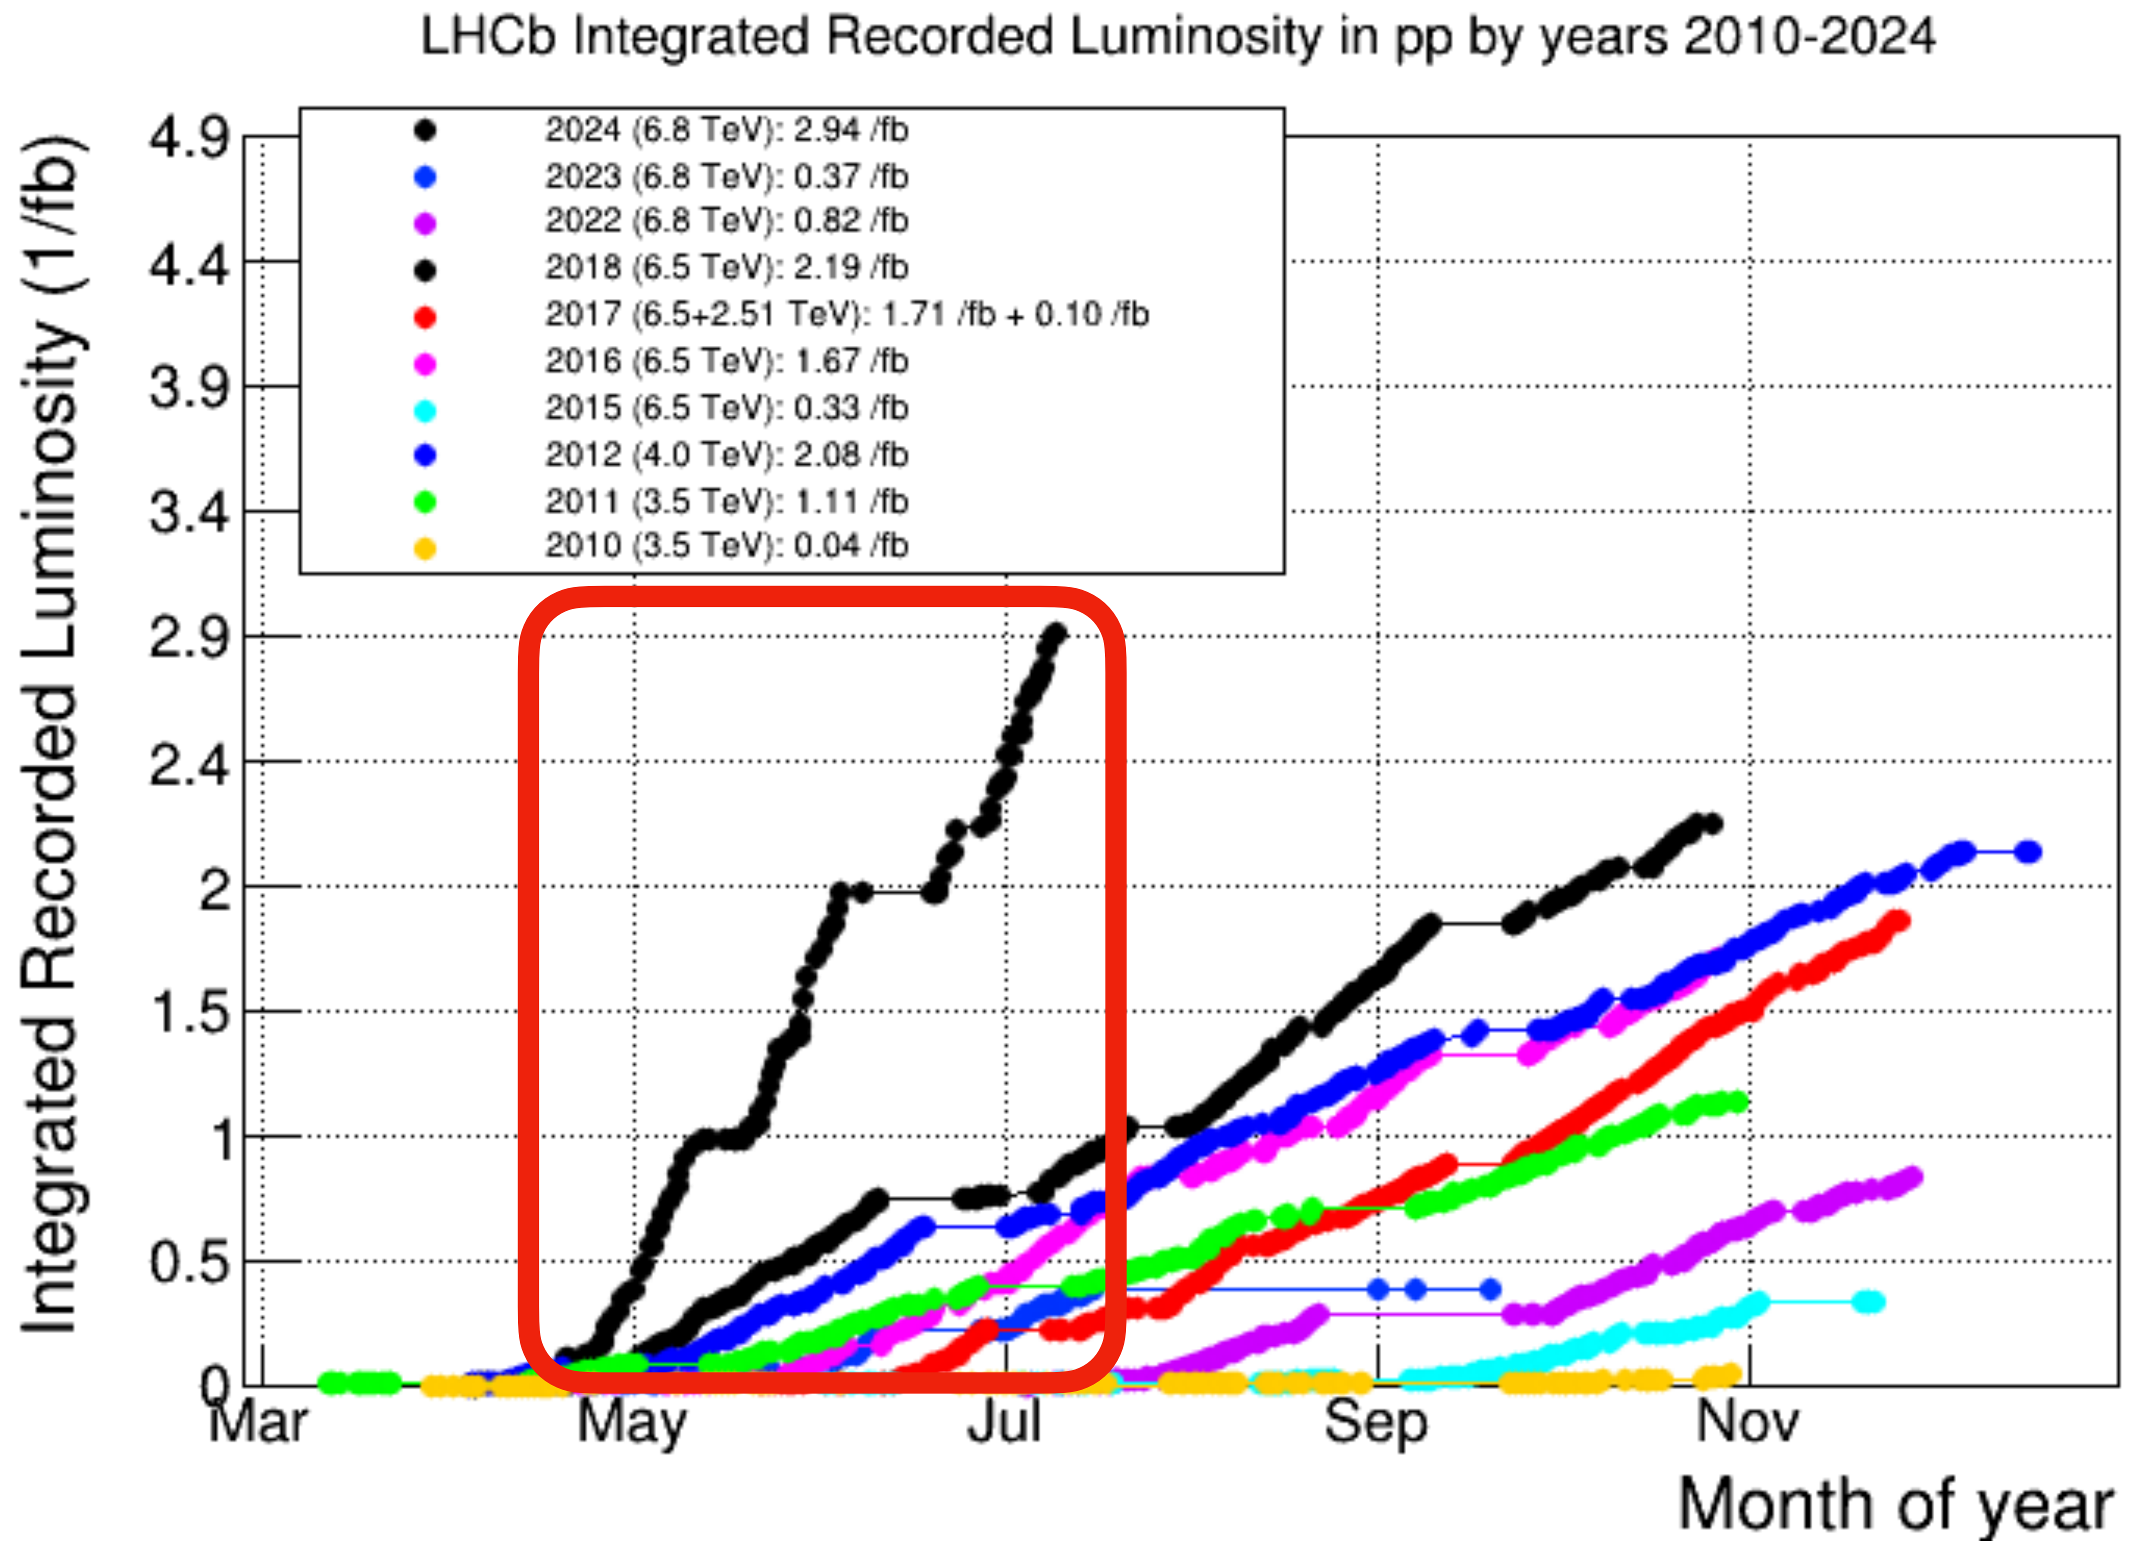

- Run at 5x higher instantaneous luminosity → pile-up of around 5
- New set of tracking detectors (VELO, UT, SciFi)
- Re-designed trigger system → Trigger-less sub-detector read-out at collision rate

Limitations of the Run 2 trigger system

- Run 2 trigger system:
  - hardware trigger (L0)
  - two-stage software trigger (HLT1 & HLT2)
- Tight p_T/E_T requirements by L0 $\rightarrow$ Trigger yields saturate with luminosity for fully hadronic modes
- Run 3 trigger system:
  - Removal of the hardware trigger
  - Run HLT1 directly at the effective collision rate (30MHz)

[arXiv:2305.10515]

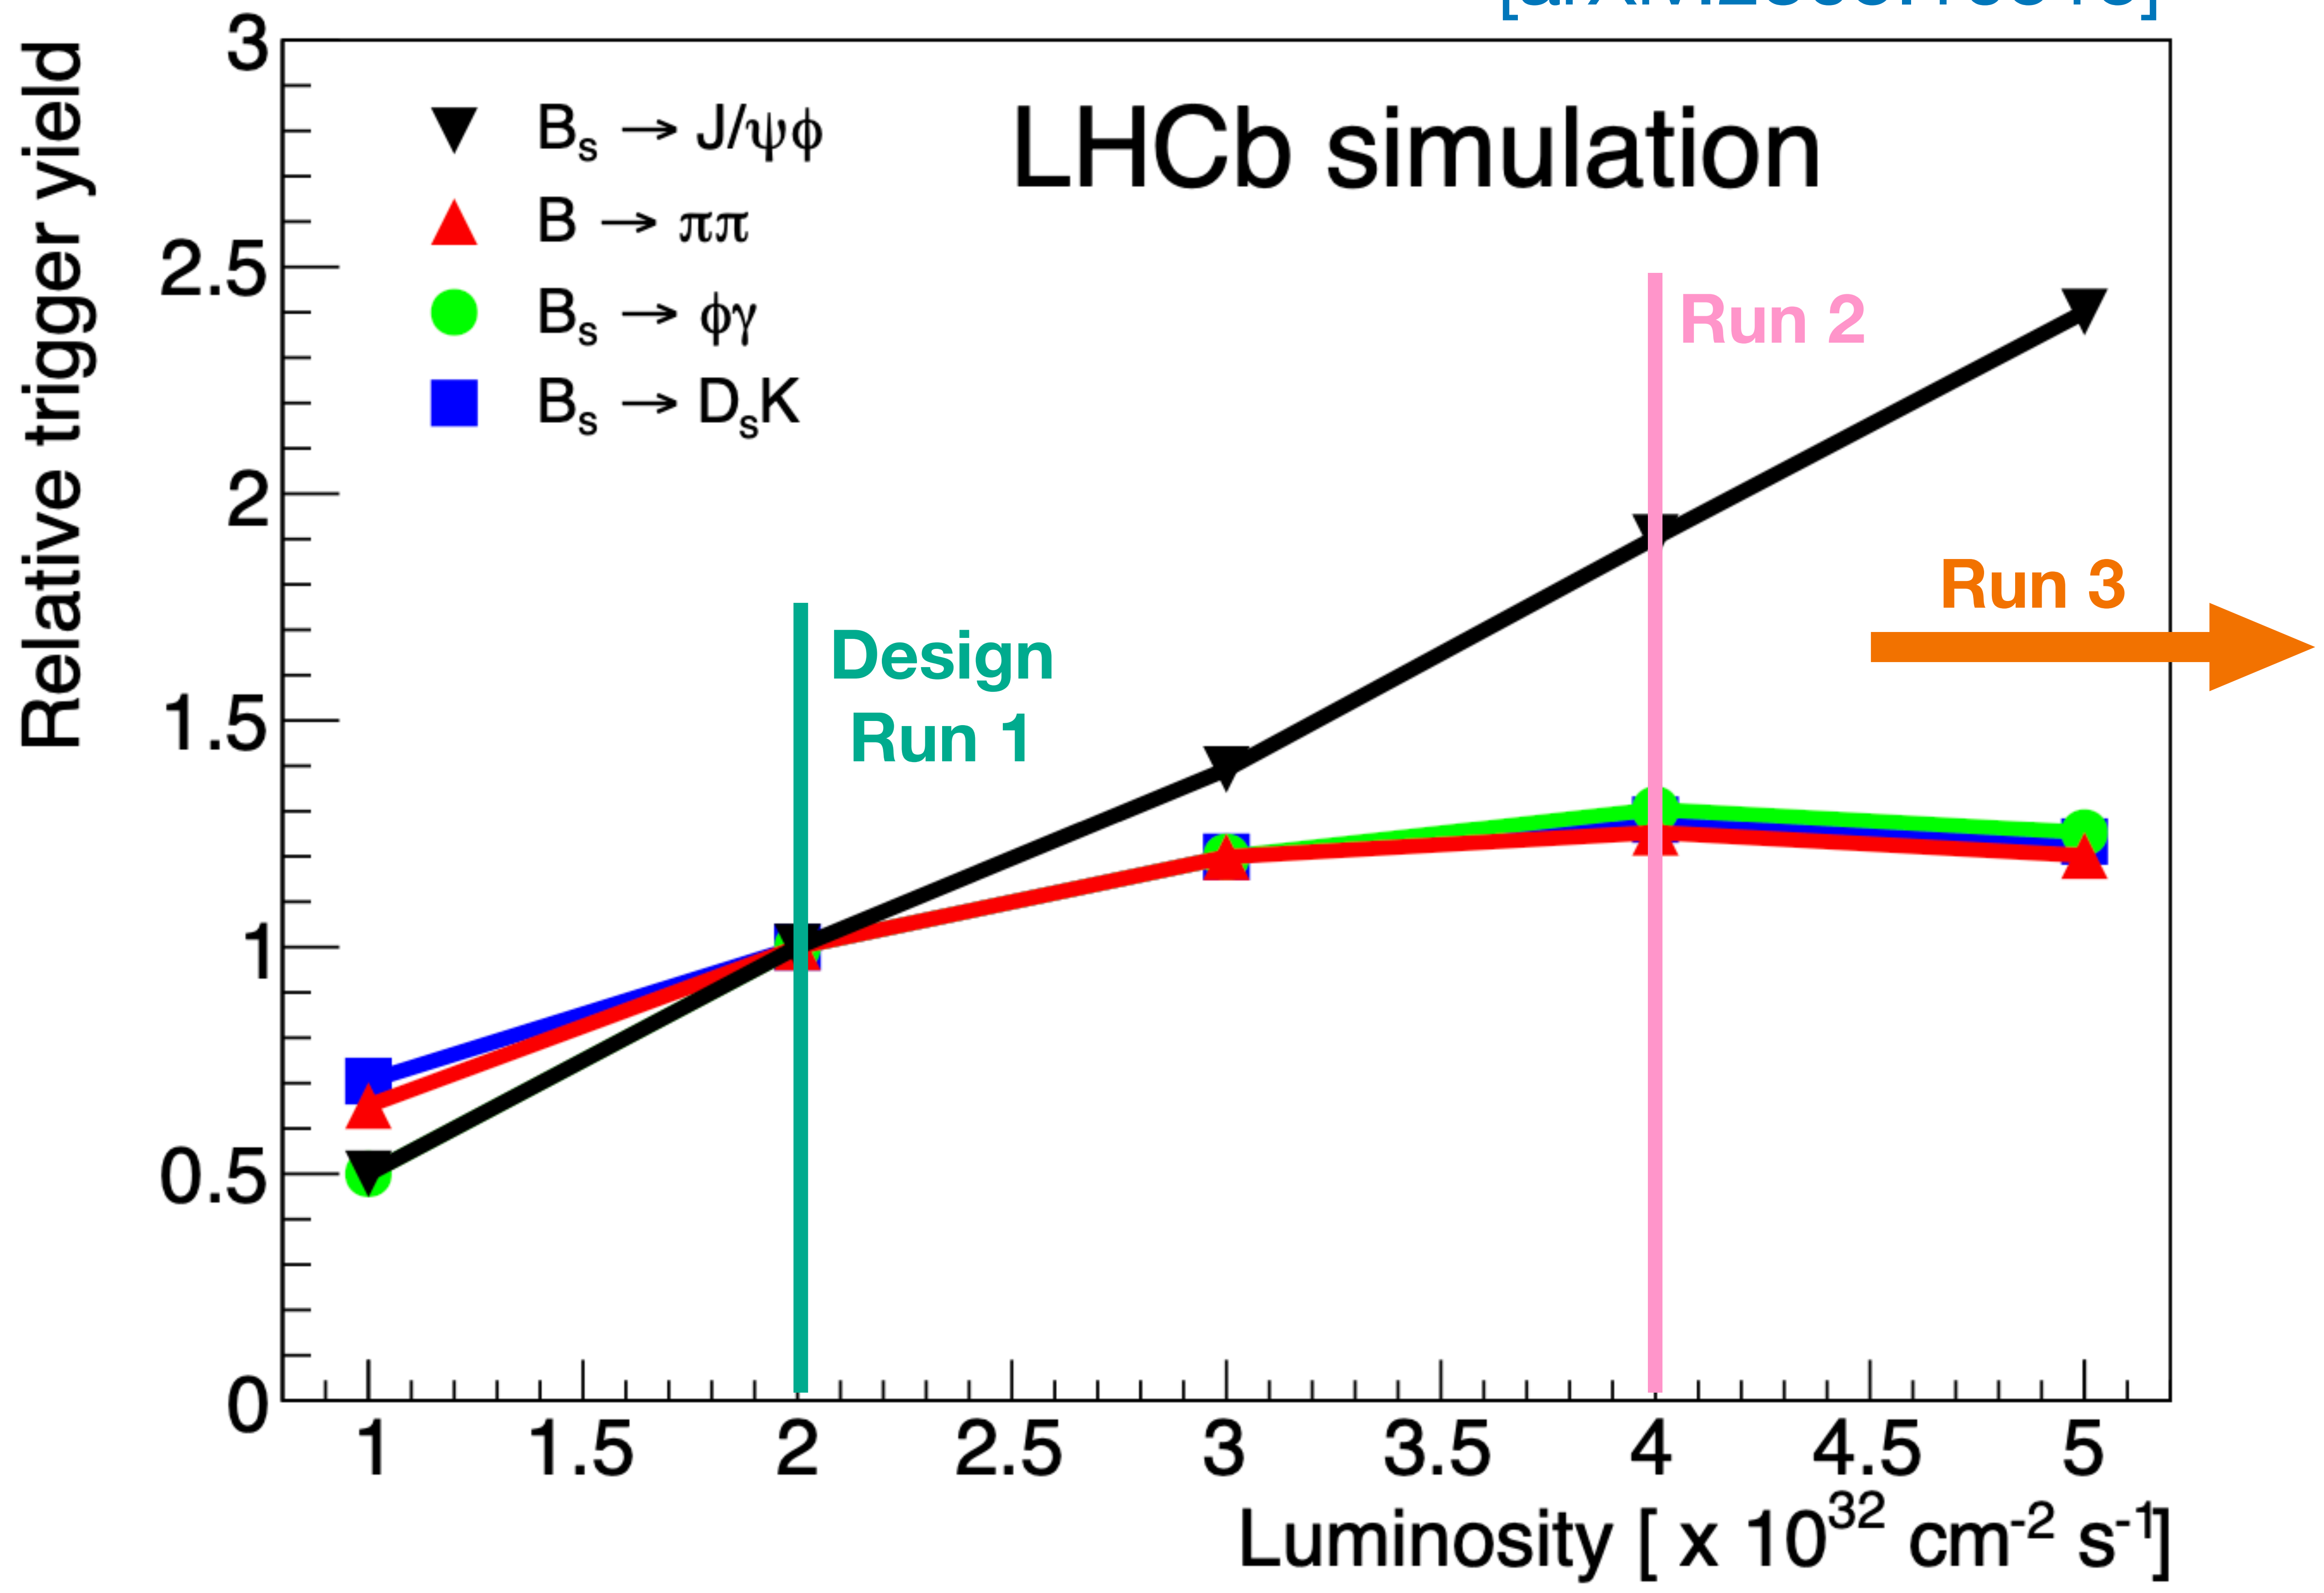

Online data flow in LHCb in Run 3

[LHCb-FIGURE-2020-016]

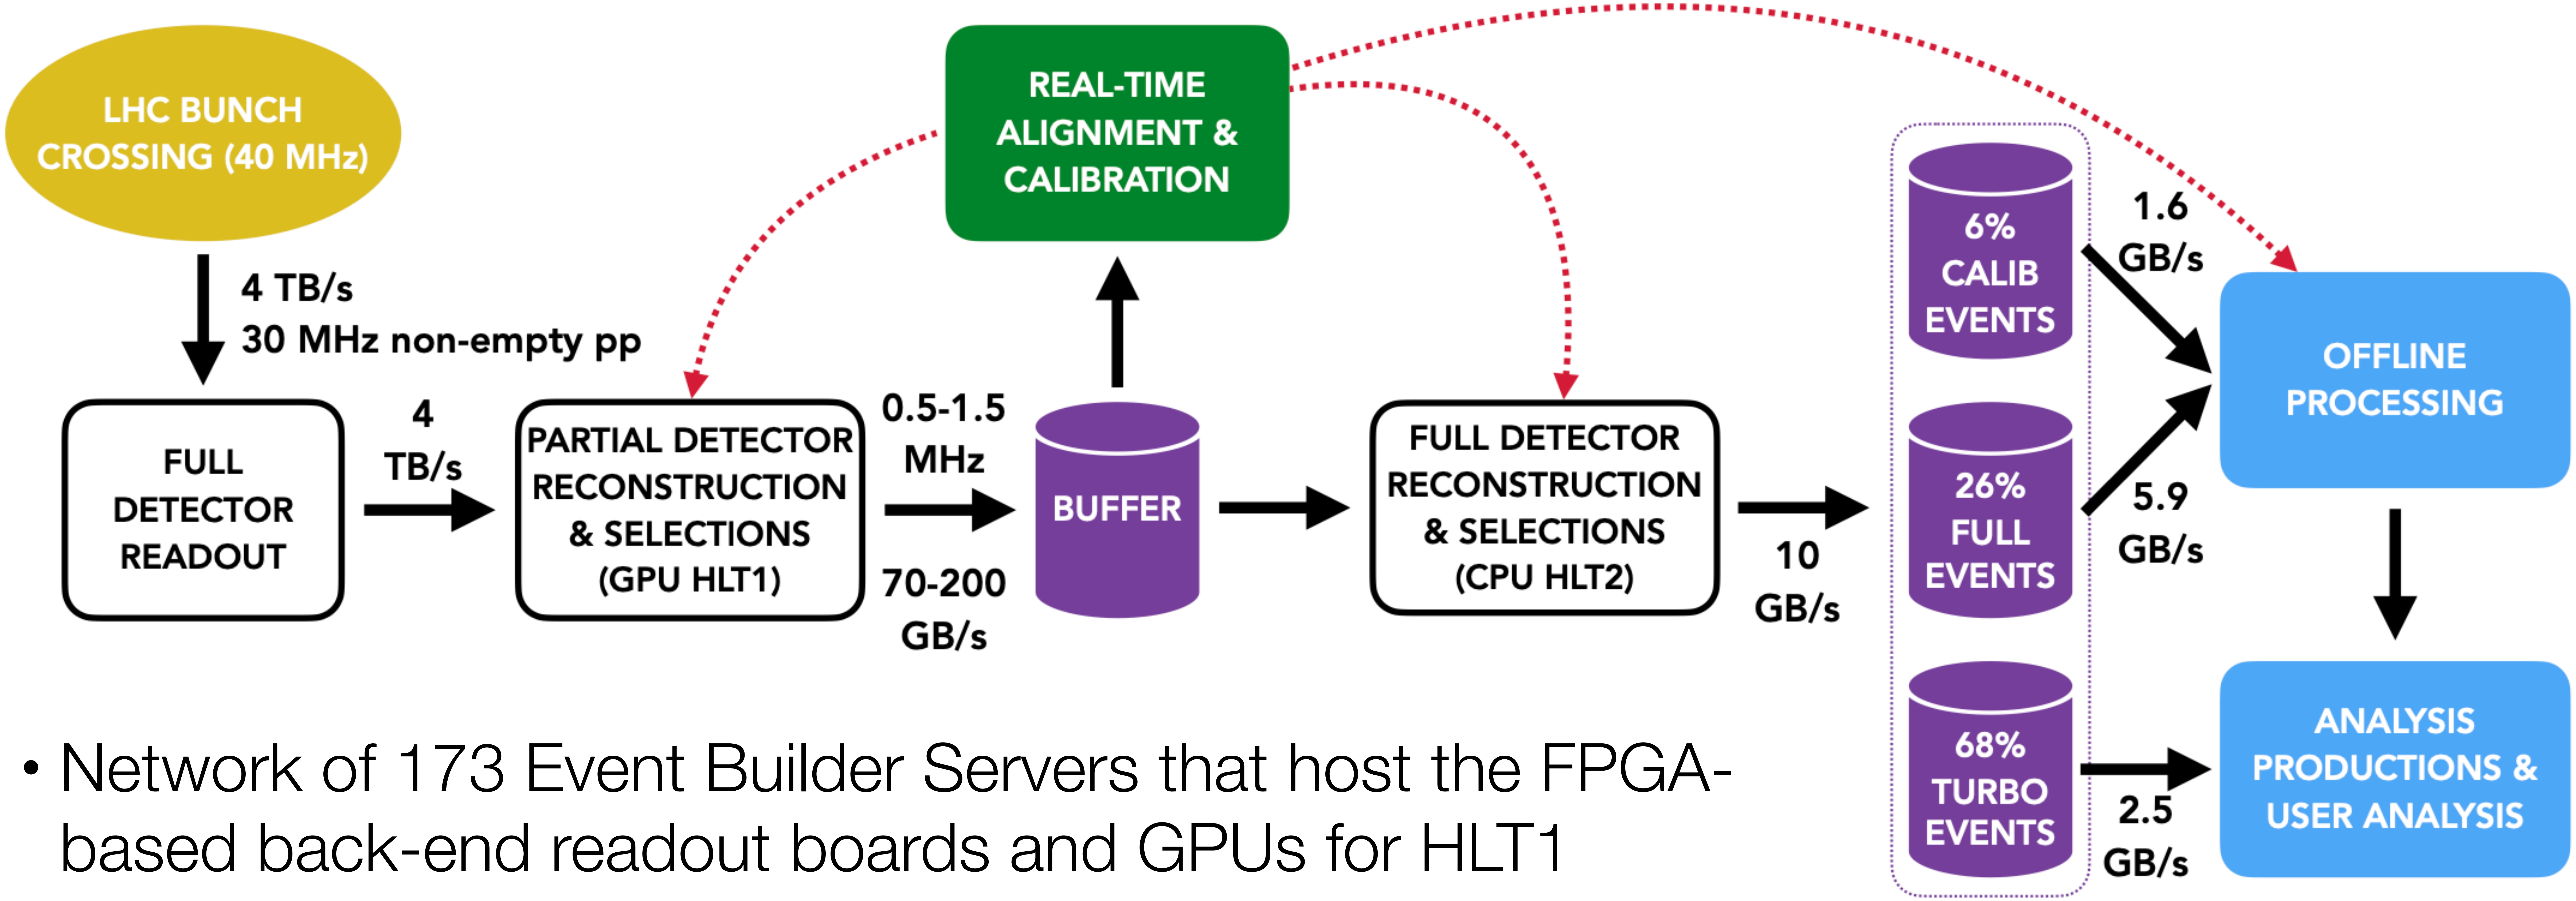

- Network of 173 Event Builder Servers that host the FPGA-based back-end readout boards and GPUs for HLT1
- Buffer system between HLT1 & HLT2 with 40PB of storage

HLT1 tasks

Partial event reconstruction at throughput of 30MHz

- Track reconstruction
  (Pattern recognition and track fitting)
- Vertex reconstruction
  (Primary and decay vertices)
- Electron clustering and bremsstrahlung recovery
- Muon identification

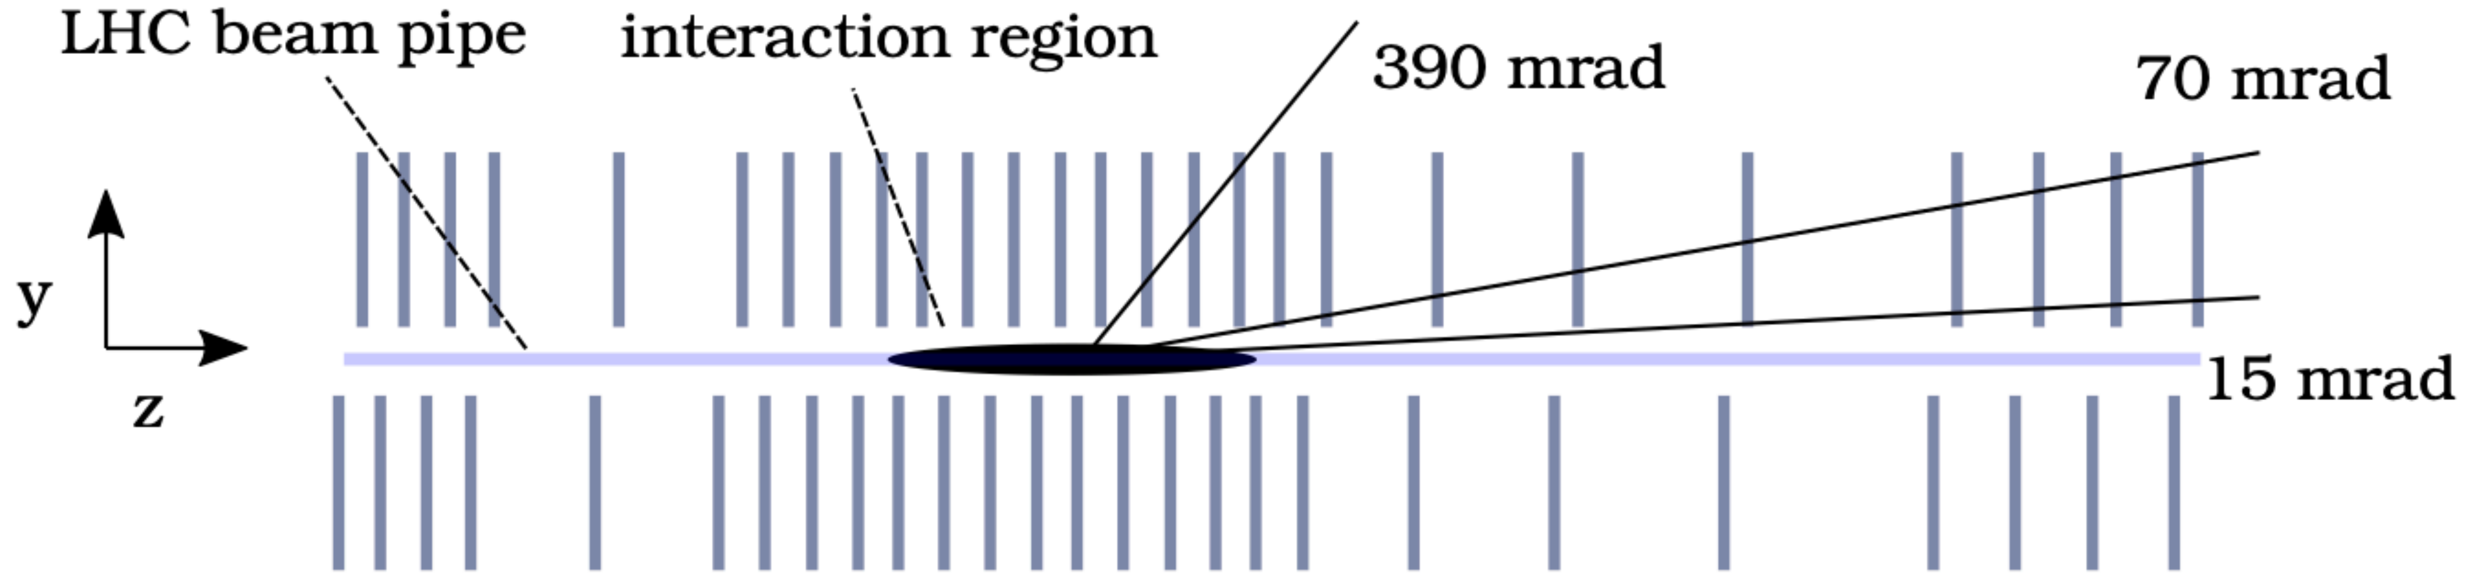

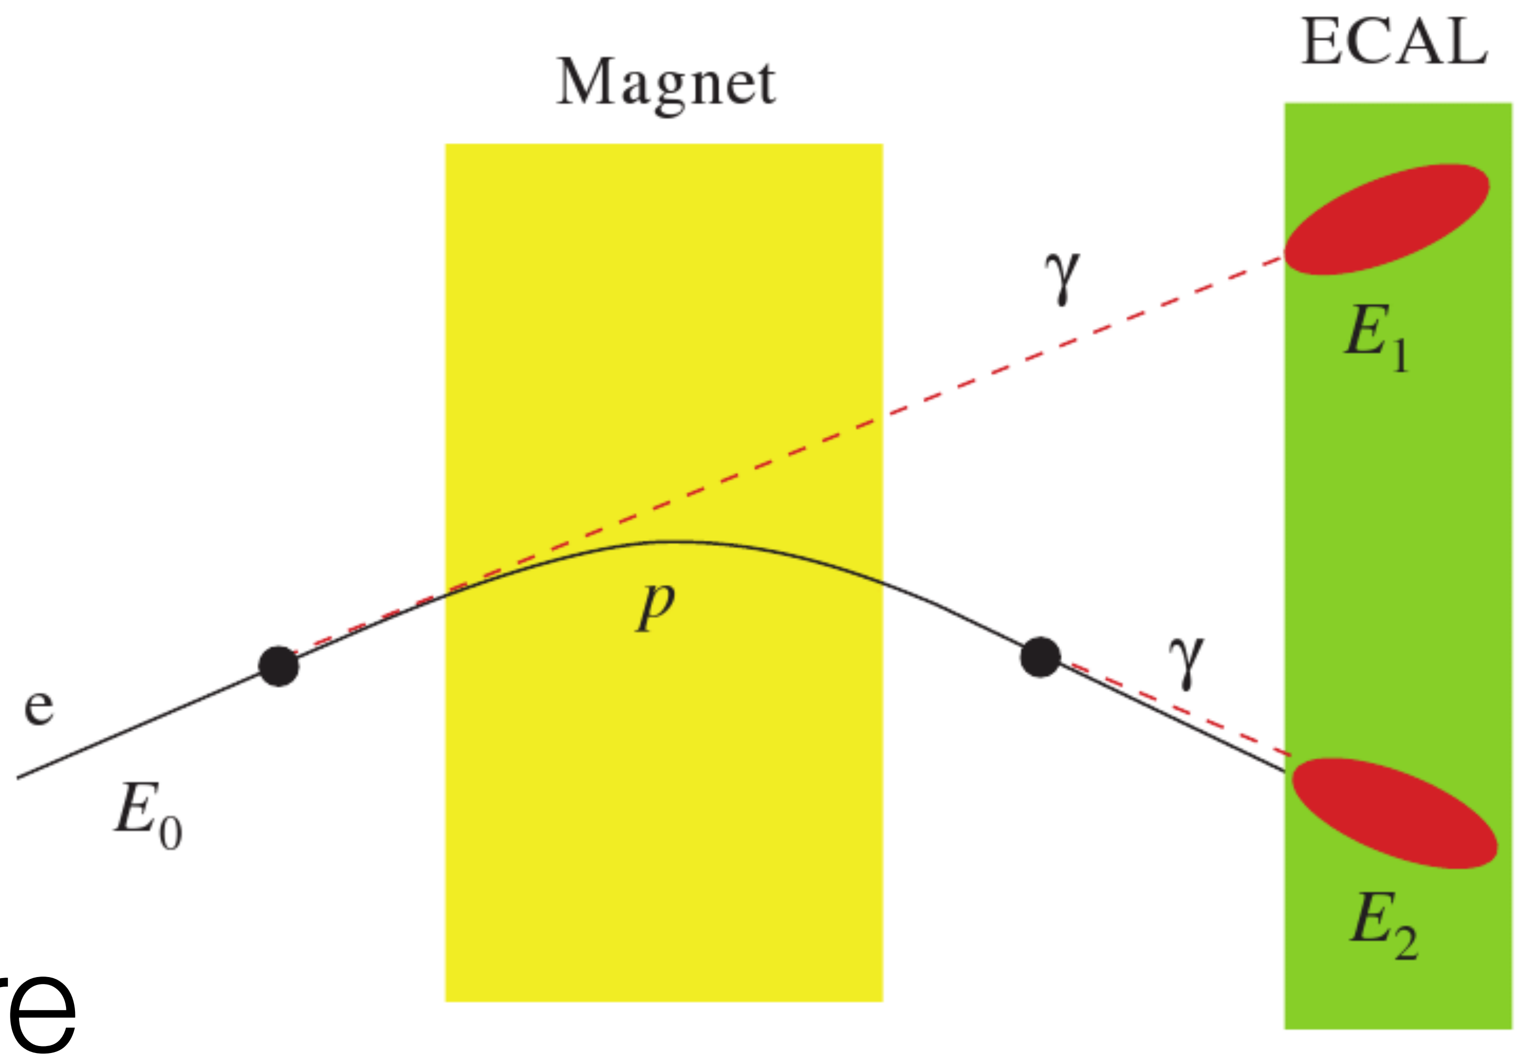

- Event selection to **reduce data rate by factor 30**
  - Cover broadly the LHCb physics programme
  - Topological lines (displaced, high-pT tracks) and more dedicated physics lines

Why GPUs?

Matches the DAQ architecture of LHCb:

- can be hosted by the Event Builder Nodes via PCIe slots
- reduced costs due to shared powering and cooling and smaller network

HLT1 tasks are suited for parallelisation:

- Events can be treated independently
- Objects of reconstruction (Tracks, vertices, …) are independent

Implemented on **323 NVIDIA RTX A5000 GPUs**

Does it work? - Challenges 2024

- Major challenges of this year:
  - Going up to pile-up of ~5
  - Include the UT in the reconstruction
- HLT1 runs successfully with a measured throughput of 30.5MHz at pile-up of ~5
- UT is still being commissioned on the detector side, so far only being used in the HLT2 reconstruction
- Throughput is still a major challenge due to the combinatorics in the pattern recognition of the tracking algorithms
- New features e.g. downstream track reconstruction [see talk by Izaac Sanderswood] could be included if throughput can be increased

Does it work?: Trigger efficiencies at HLT1

[LHCB-FIGURE-2024-014] [LHCB-FIGURE-2024-006] [LHCB-FIGURE-2024-007]

- Significant improvements in trigger efficiencies at HLT1 level
  - Huge gain at low-p_T

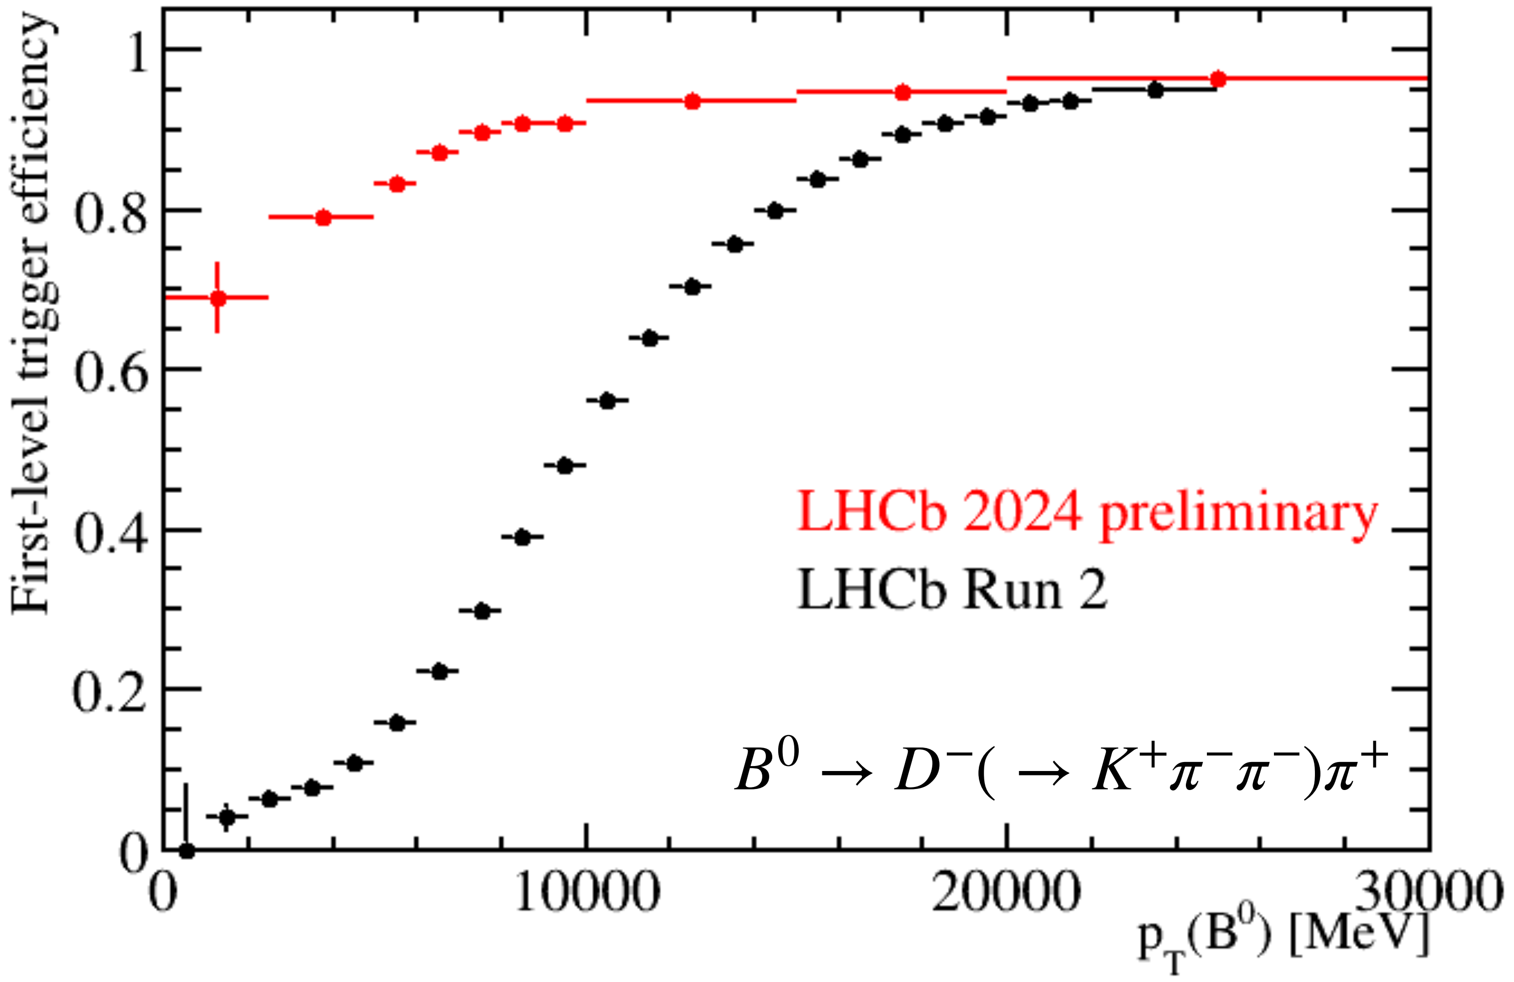

Does it work?: Trigger efficiencies at HLT1

[LHCB-FIGURE-2024-014] [LHCB-FIGURE-2024-006] [LHCB-FIGURE-2024-007]

- Significant improvements in trigger efficiencies at HLT1 level
  - Huge gain at low-p_T
    → Very beneficial for our charm physics programme

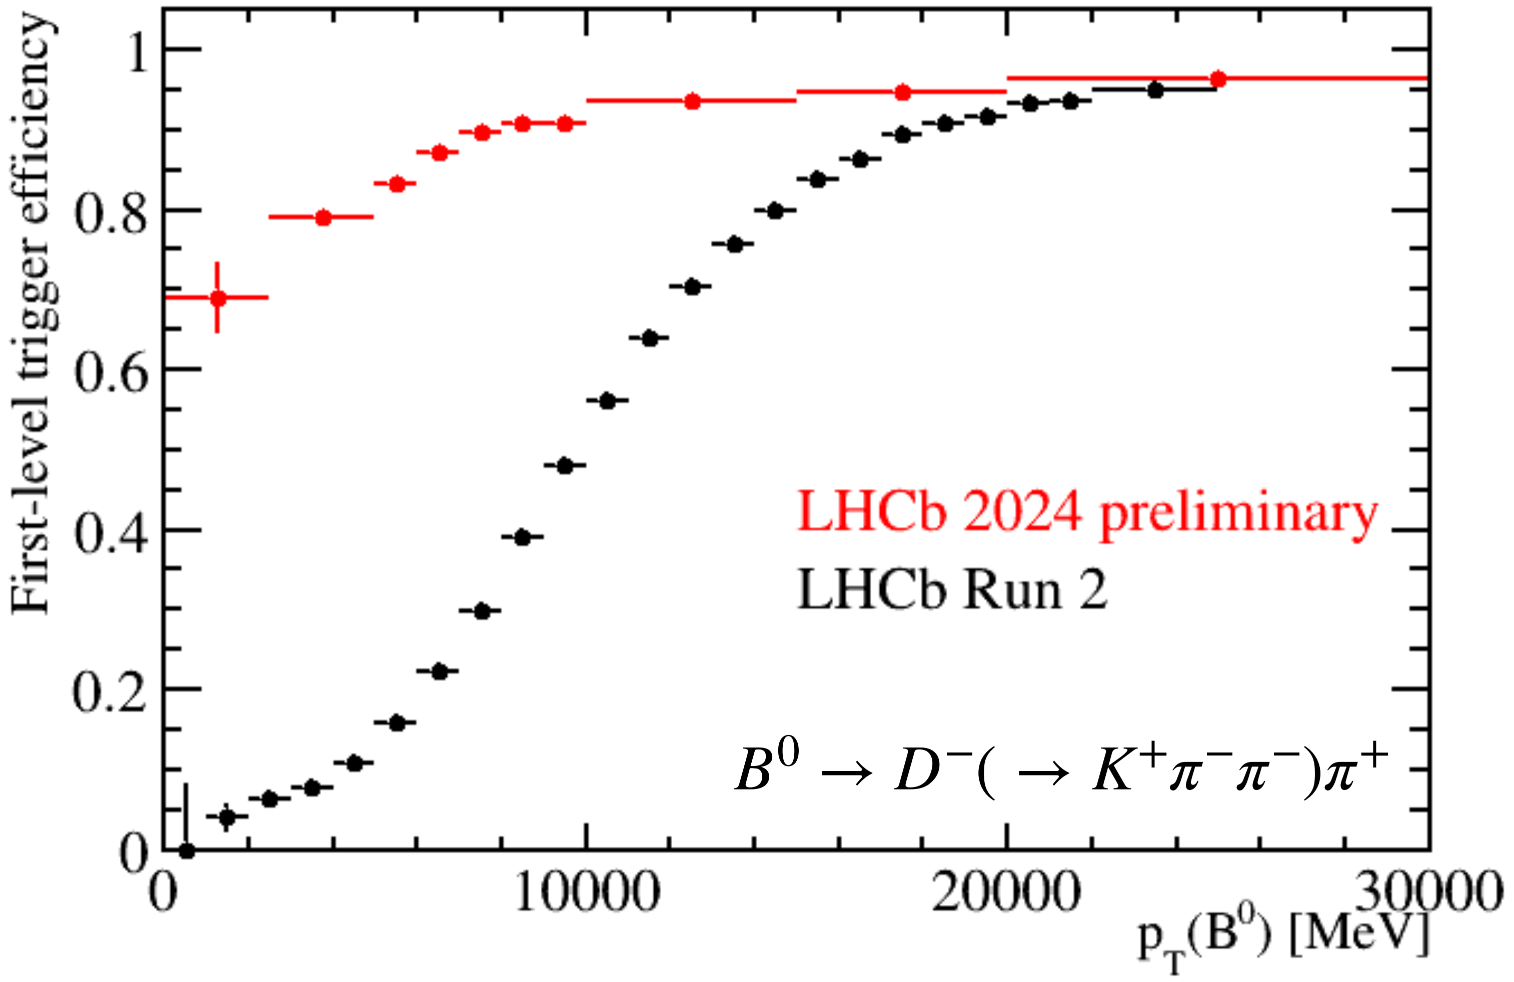

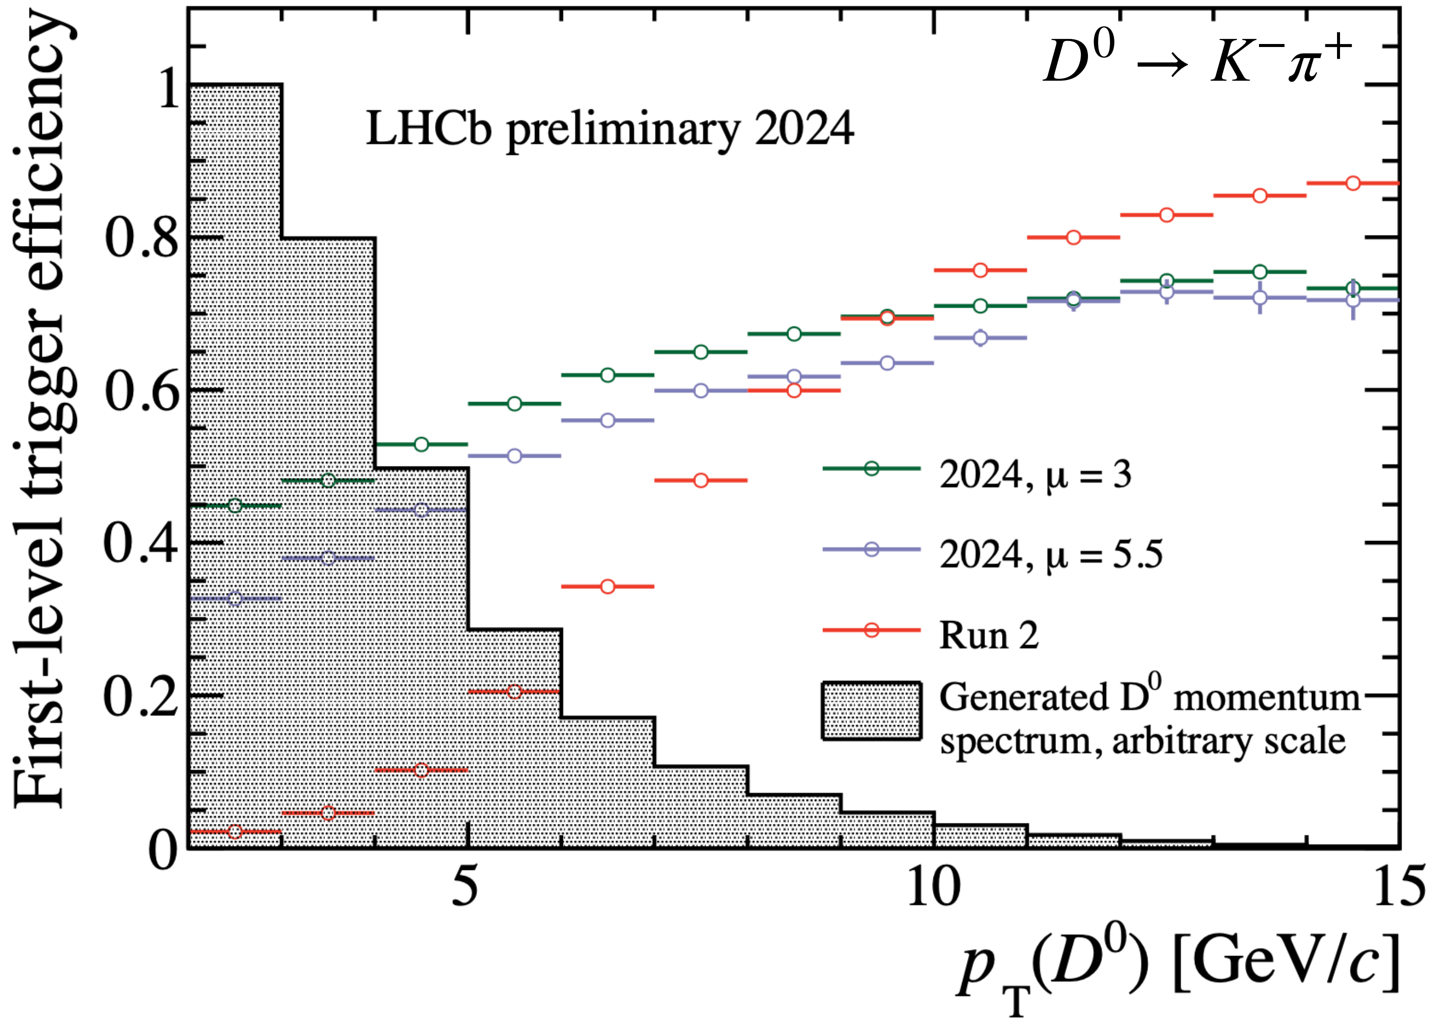

Does it work?: Trigger efficiencies at HLT1

[LHCB-FIGURE-2024-014] [LHCB-FIGURE-2024-006] [LHCB-FIGURE-2024-007]

- Significant improvements in trigger efficiencies at HLT1 level
  - Huge gain at low-p_T → Very beneficial for our charm physics programme
  - Muon channels at similar performance as in Run 2

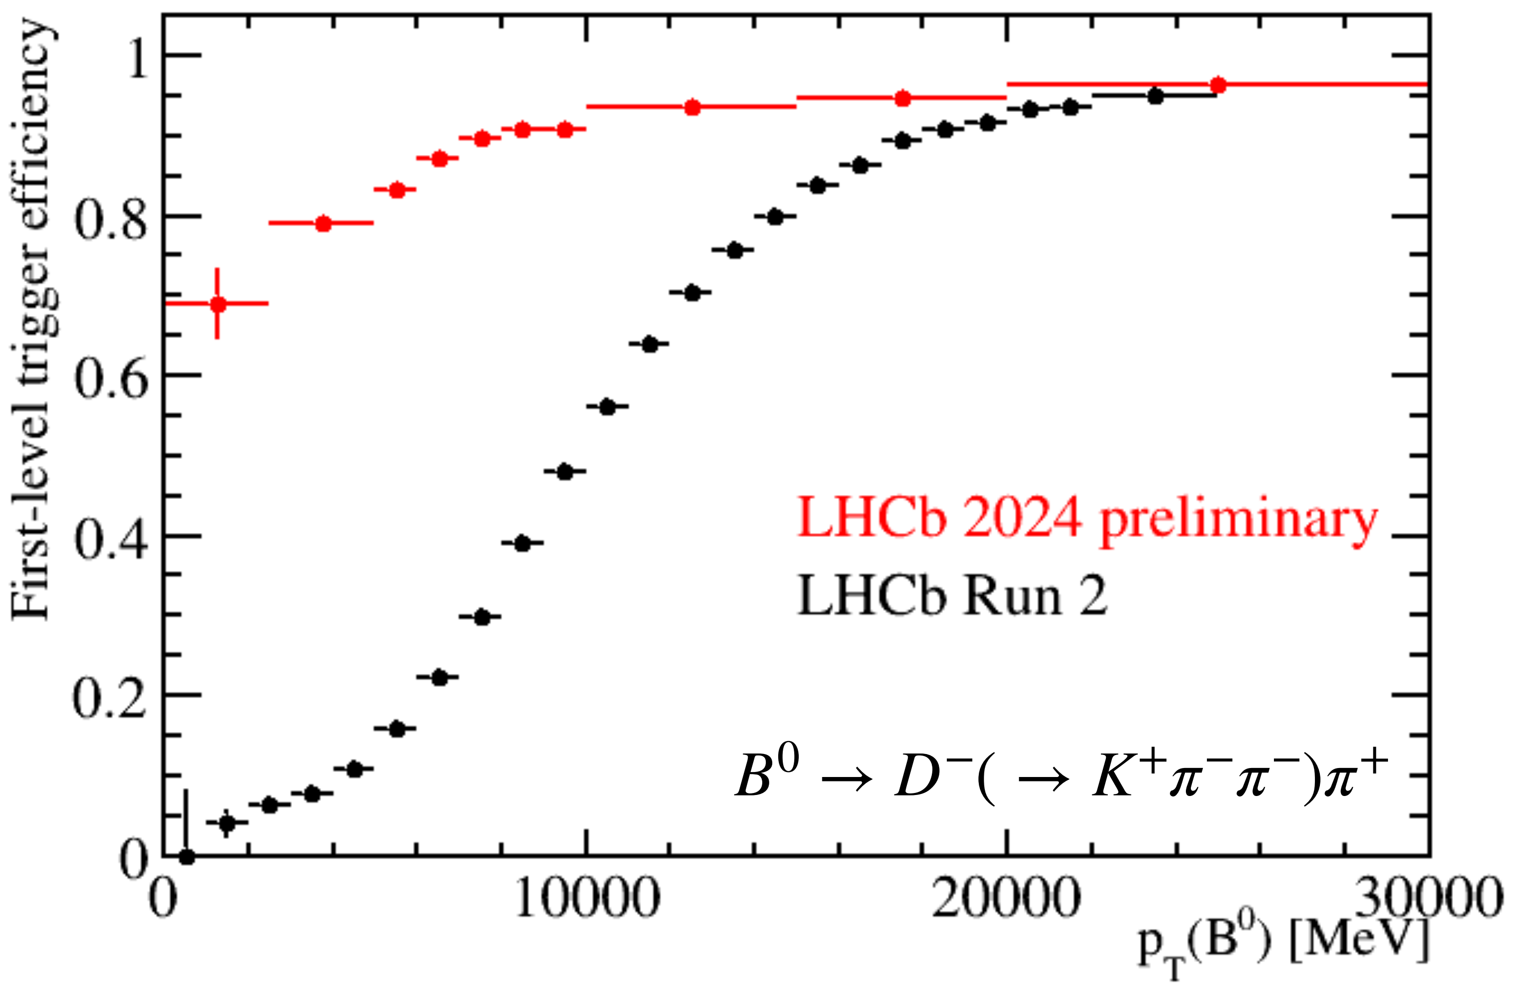

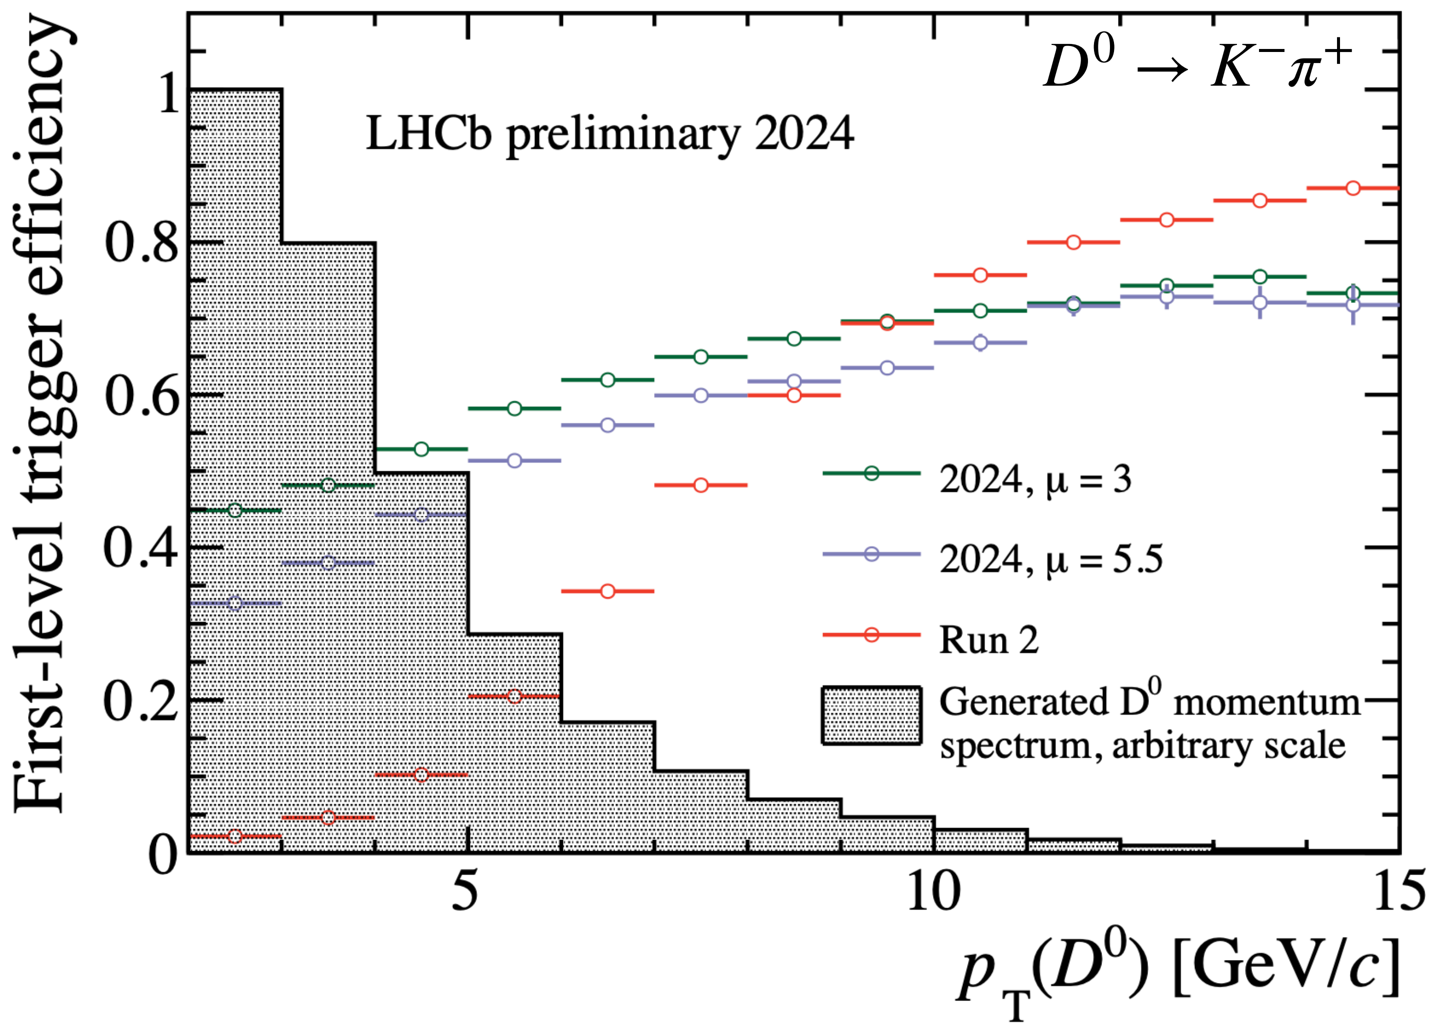

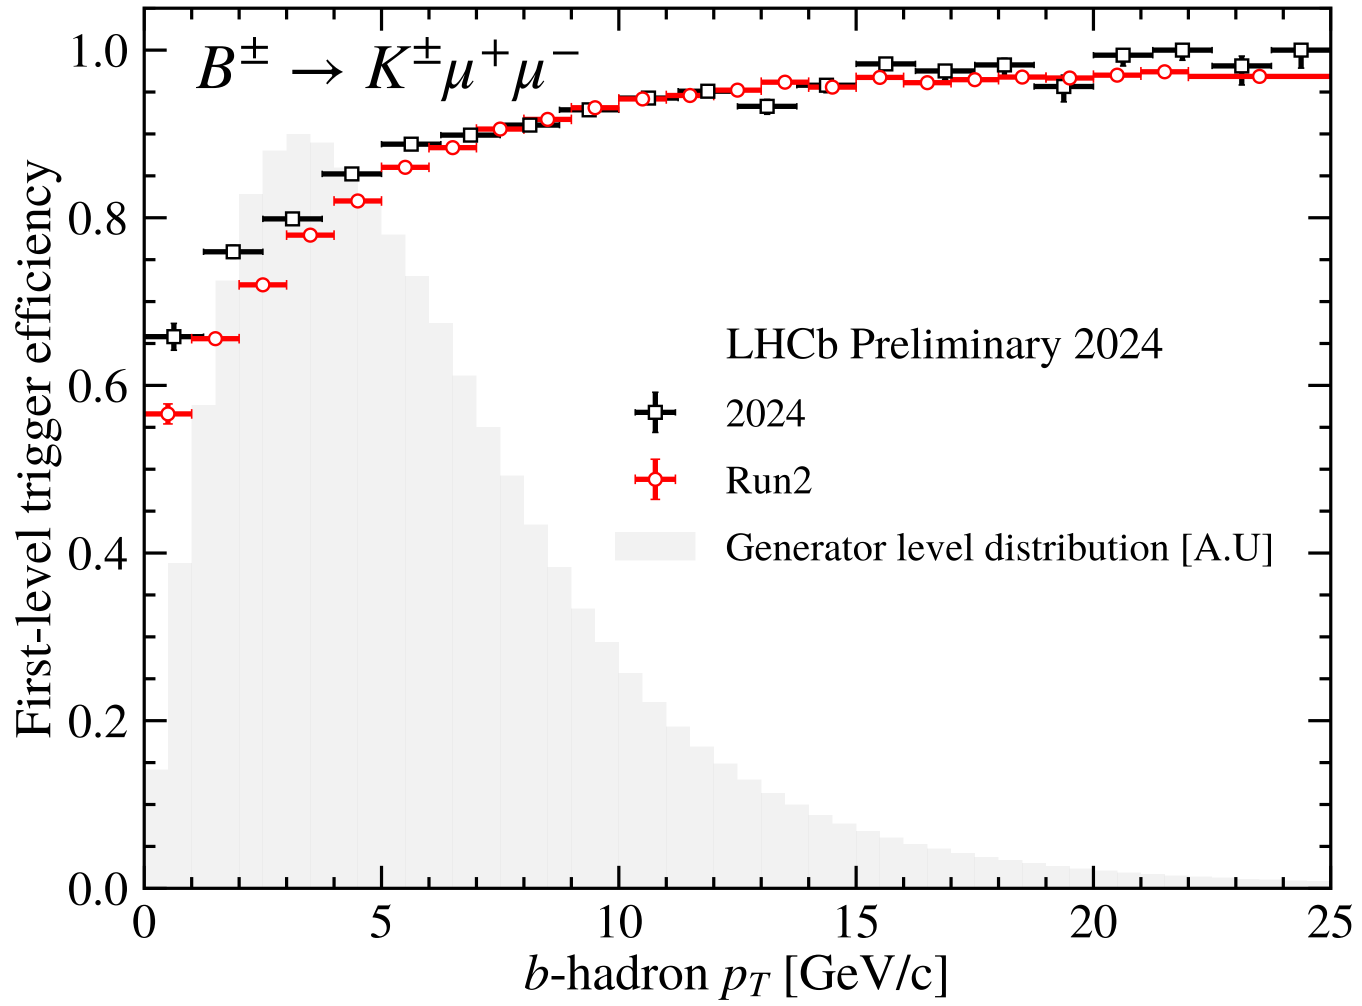

Does it work?: Trigger efficiencies at HLT1

[LHCB-FIGURE-2024-014] [LHCB-FIGURE-2024-006] [LHCB-FIGURE-2024-007]

- Significant improvements in trigger efficiencies at HLT1 level
  - Huge gain at low-p_T
    → Very beneficial for our charm physics programme
  - Muon channels at similar performance as in Run 2
  - Large impact for electron channels

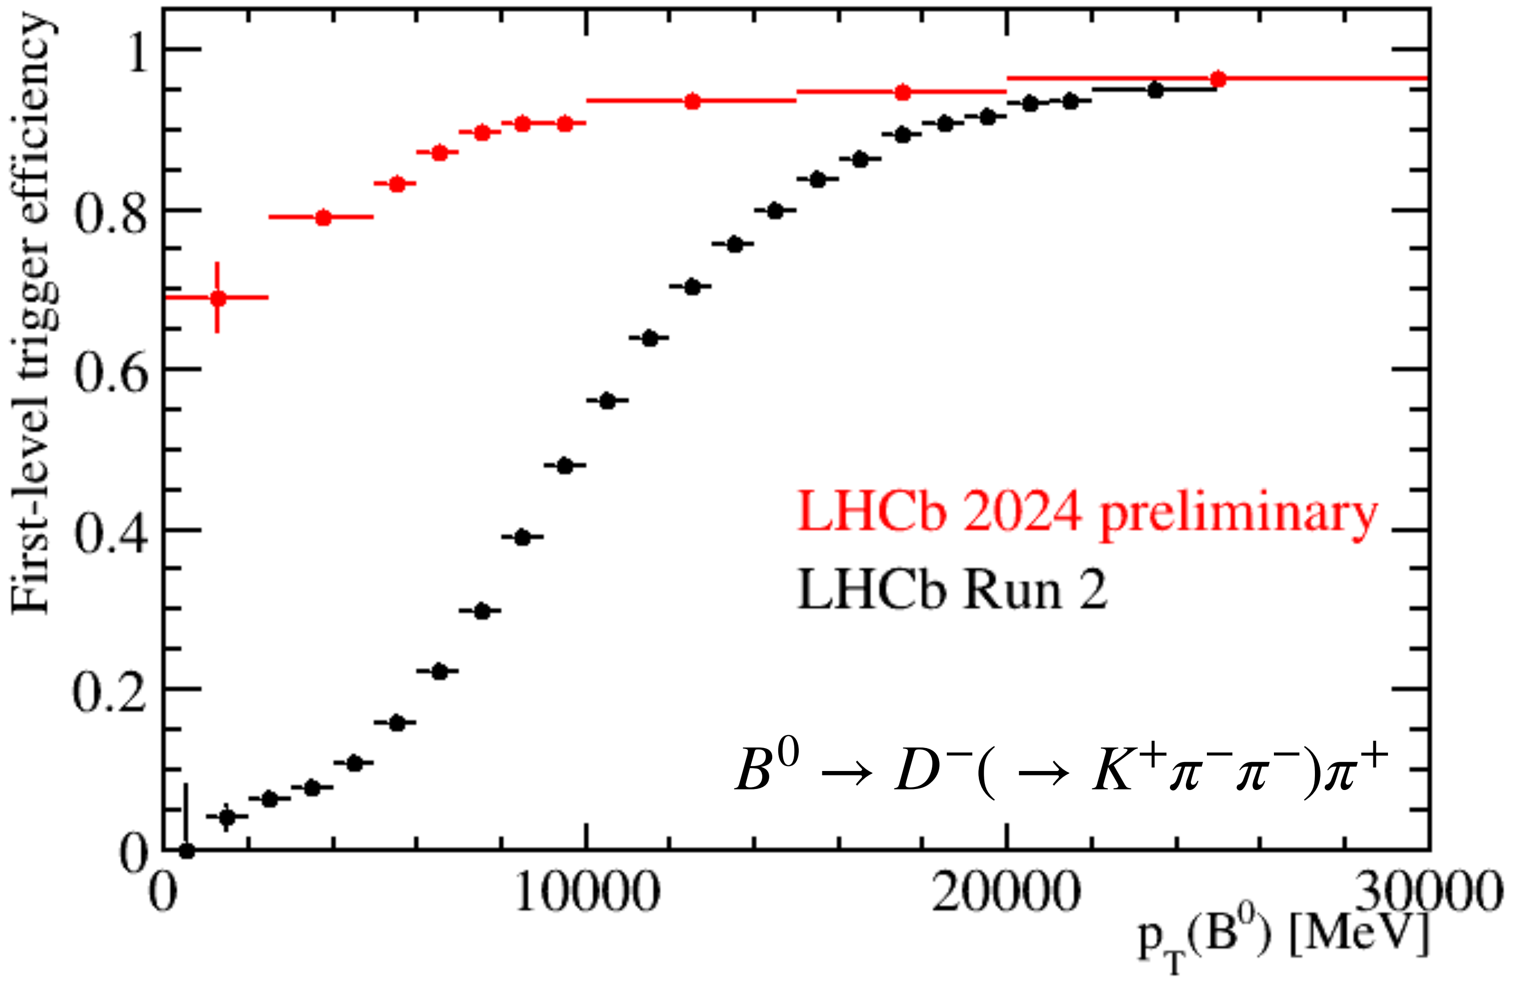

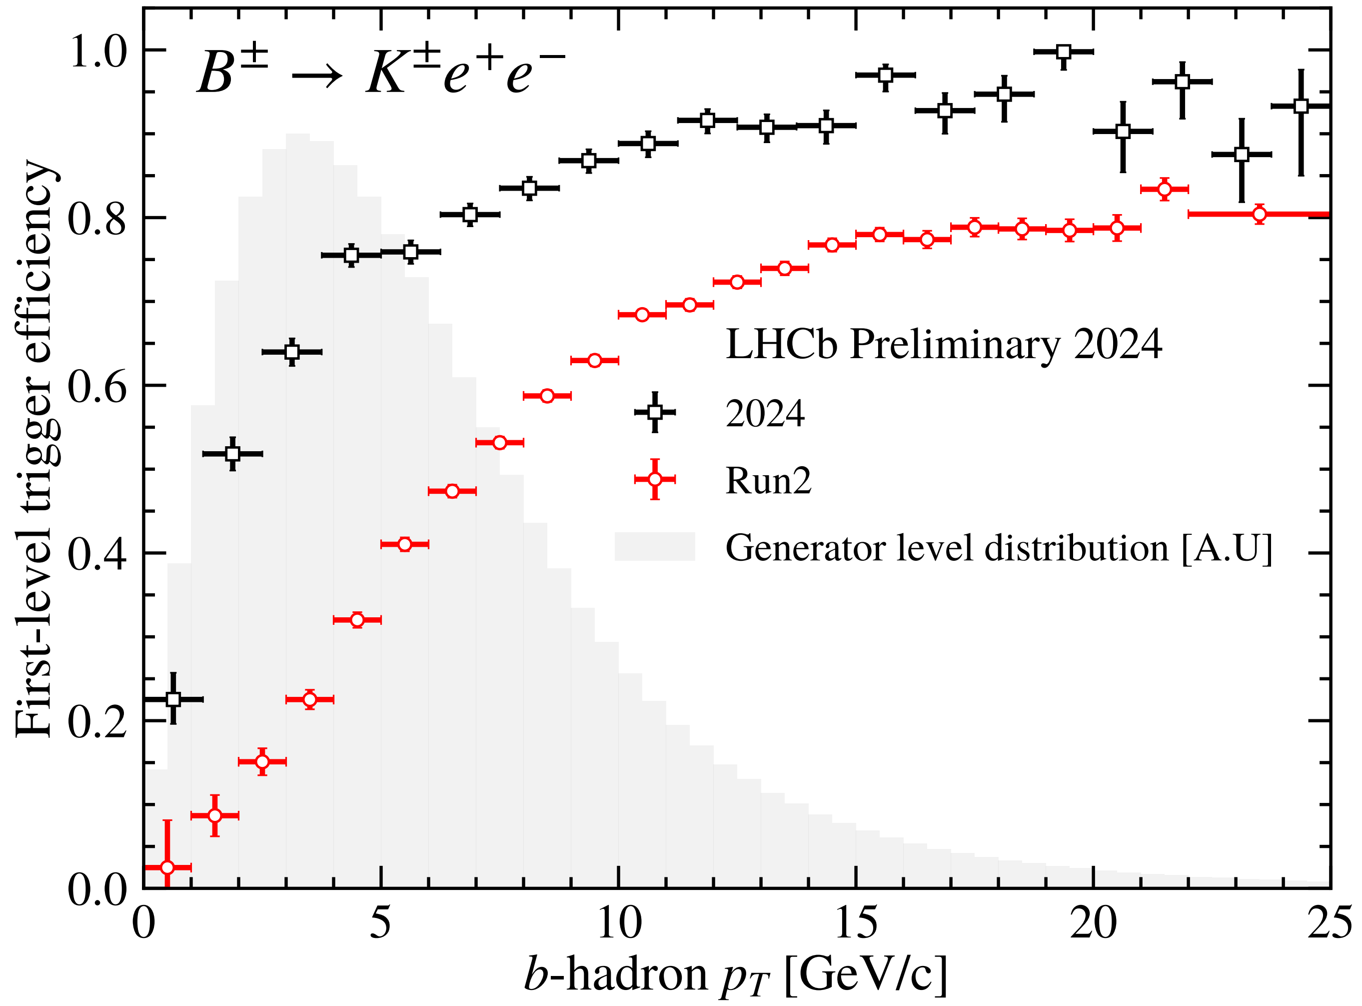

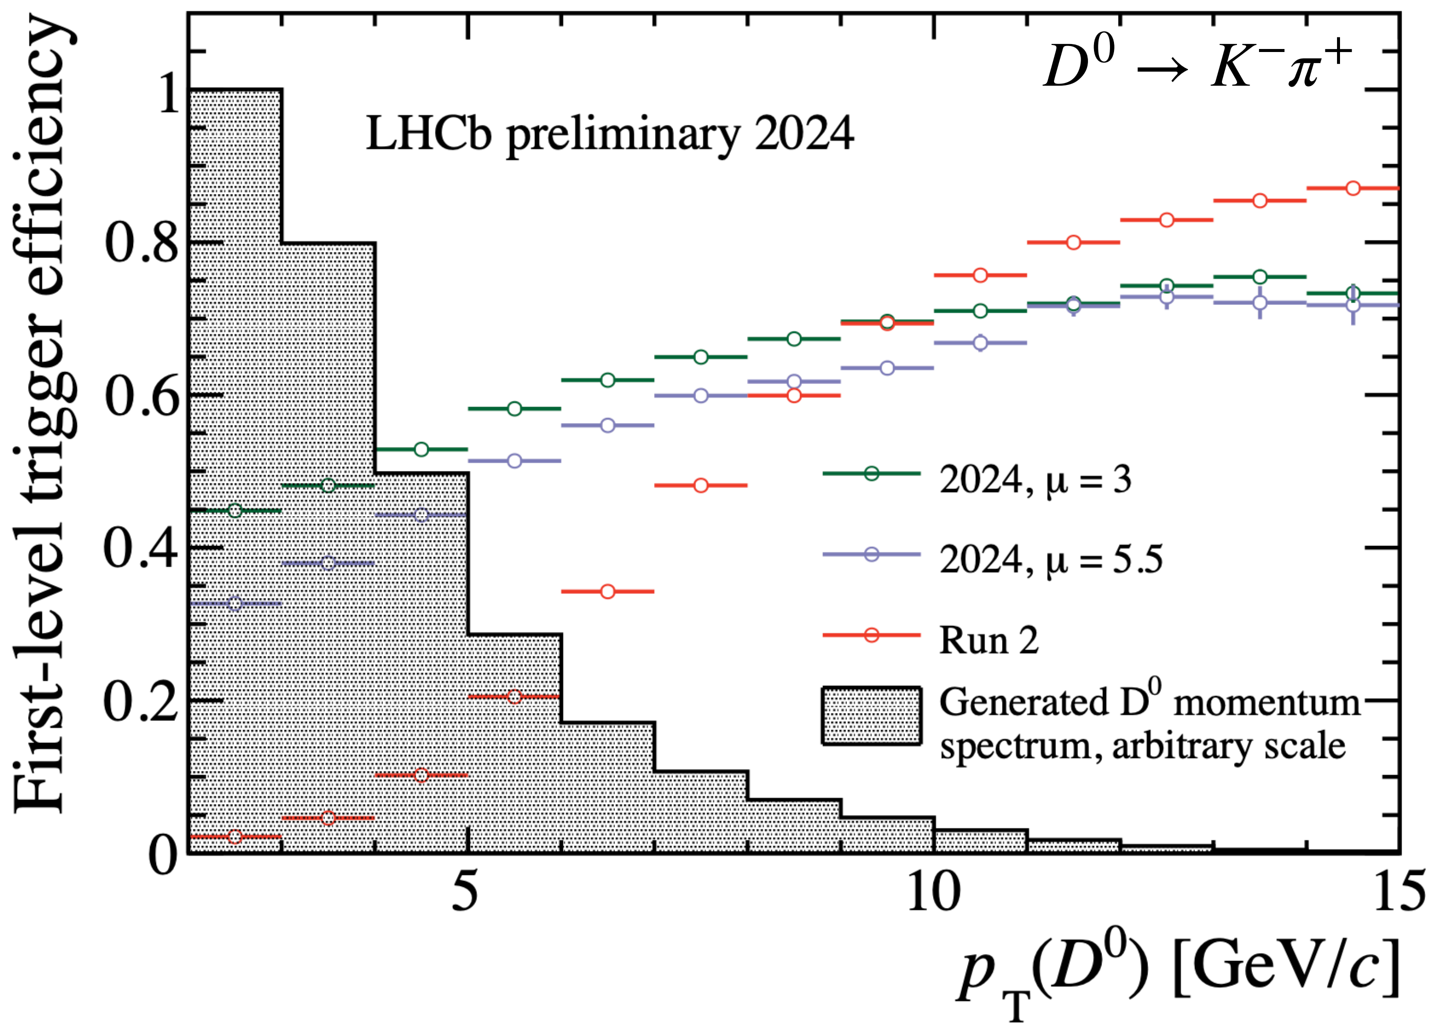

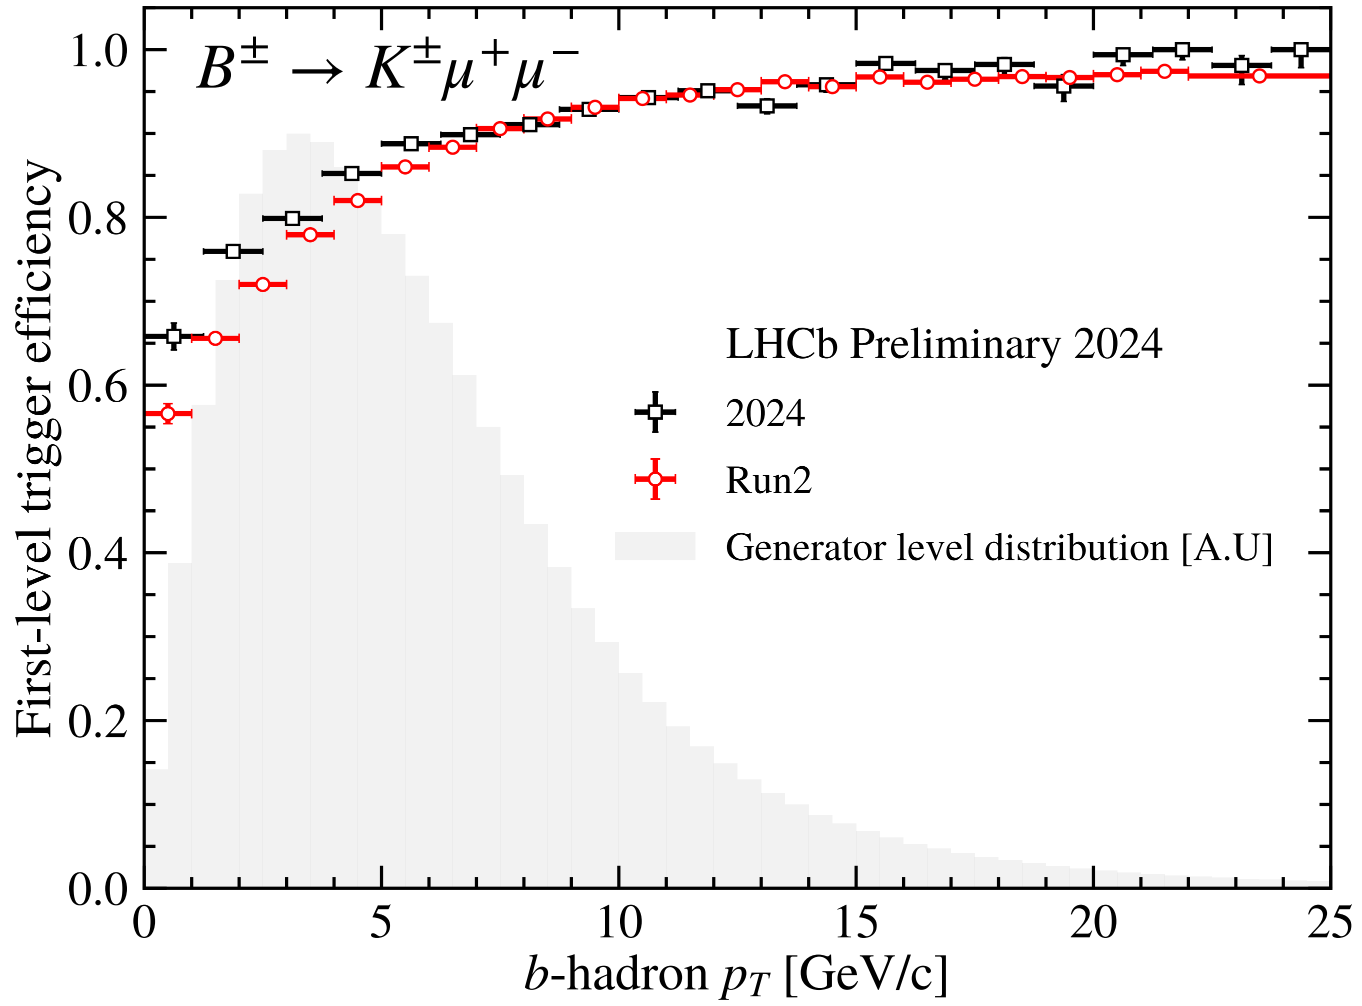

Real-time alignment and calibration

- Essential part of the upgraded trigger system to ensure offline-level quality of the HLT2 reconstruction

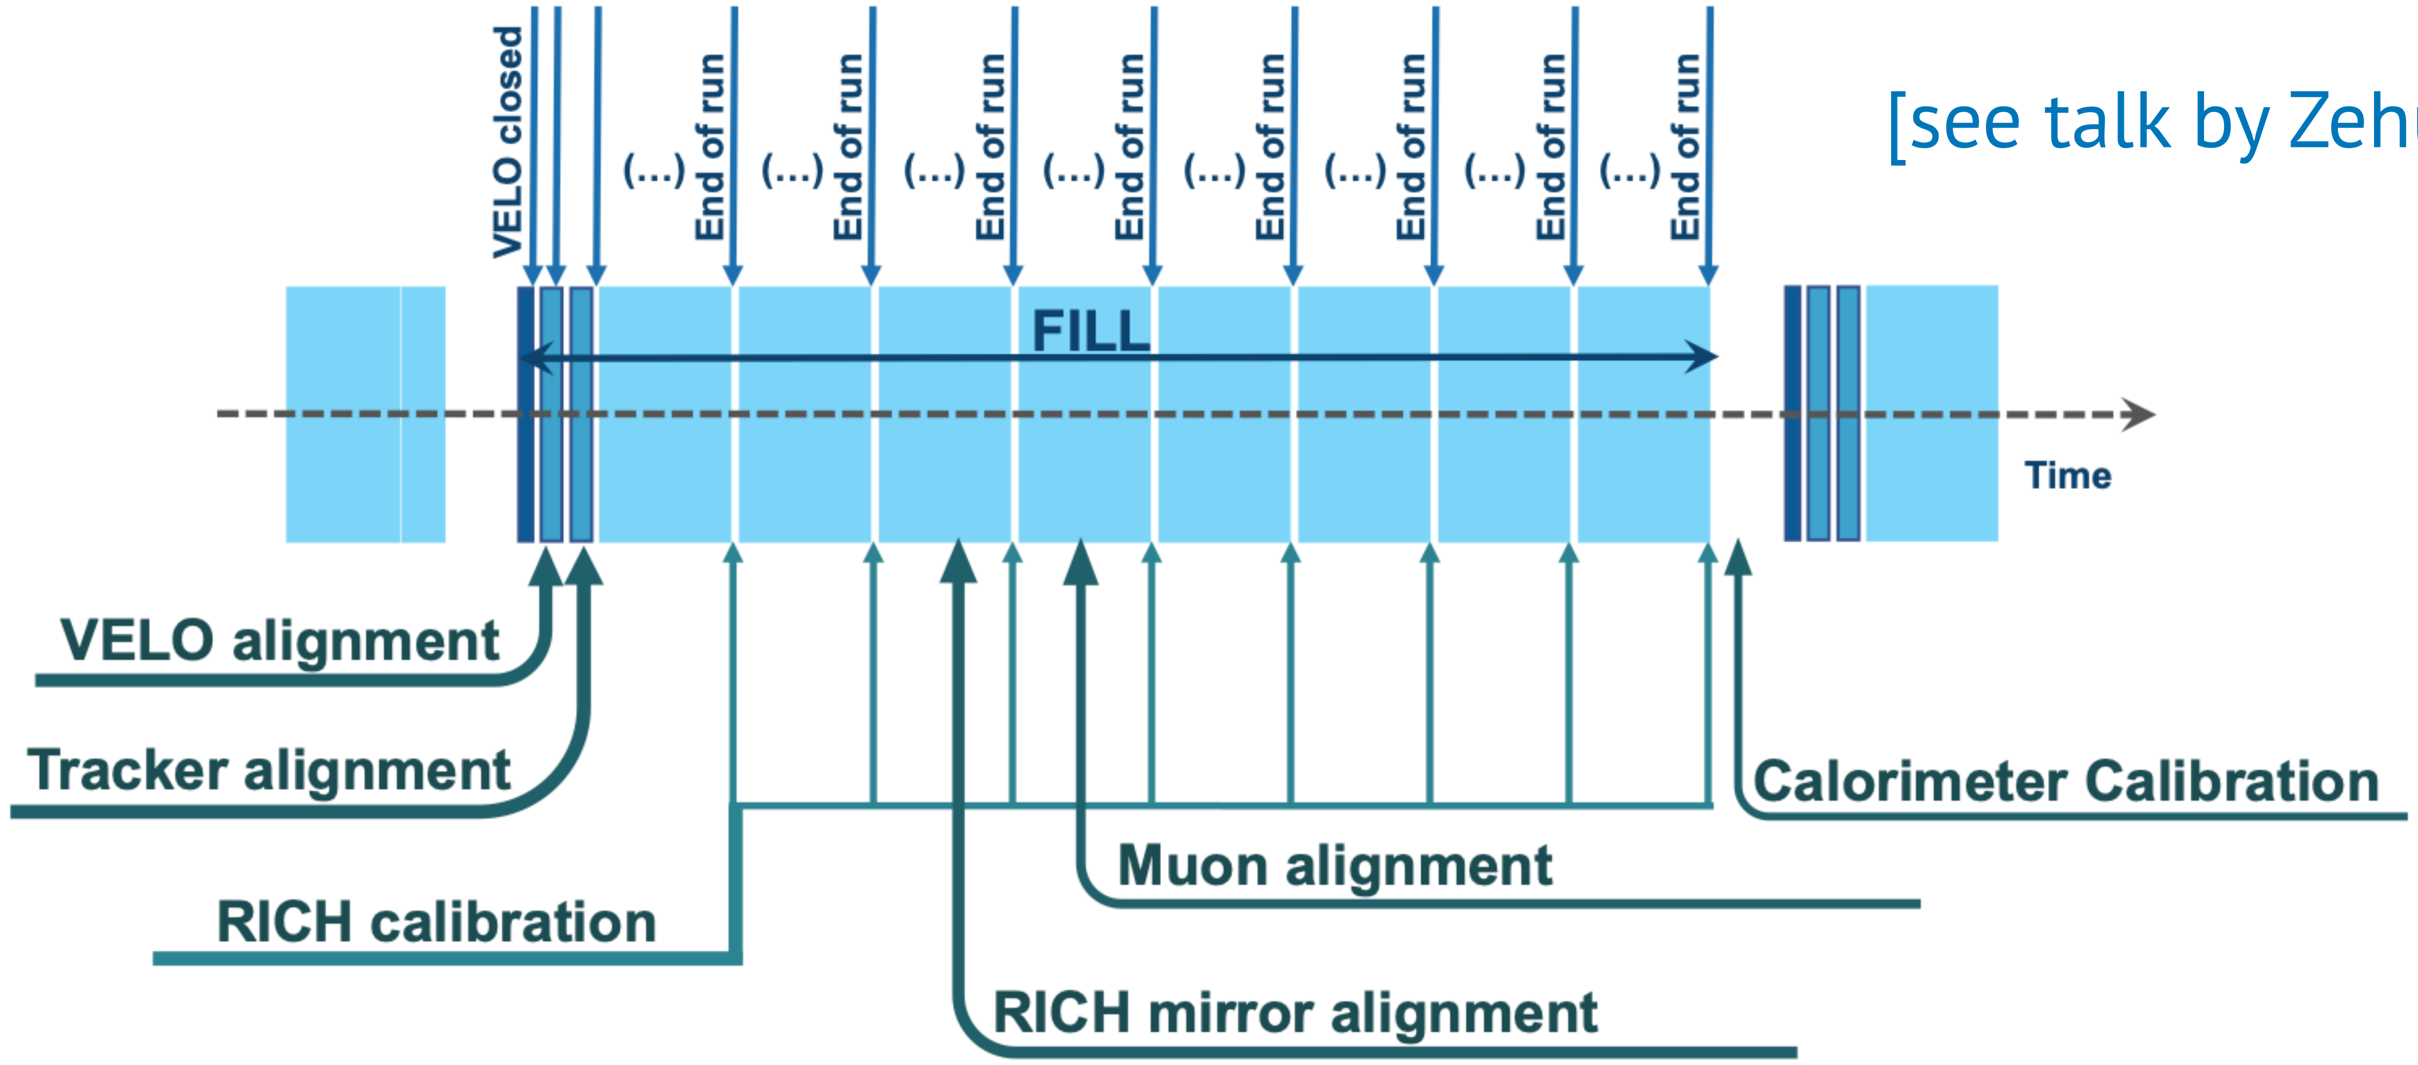

[see talk by Zehua Xu]

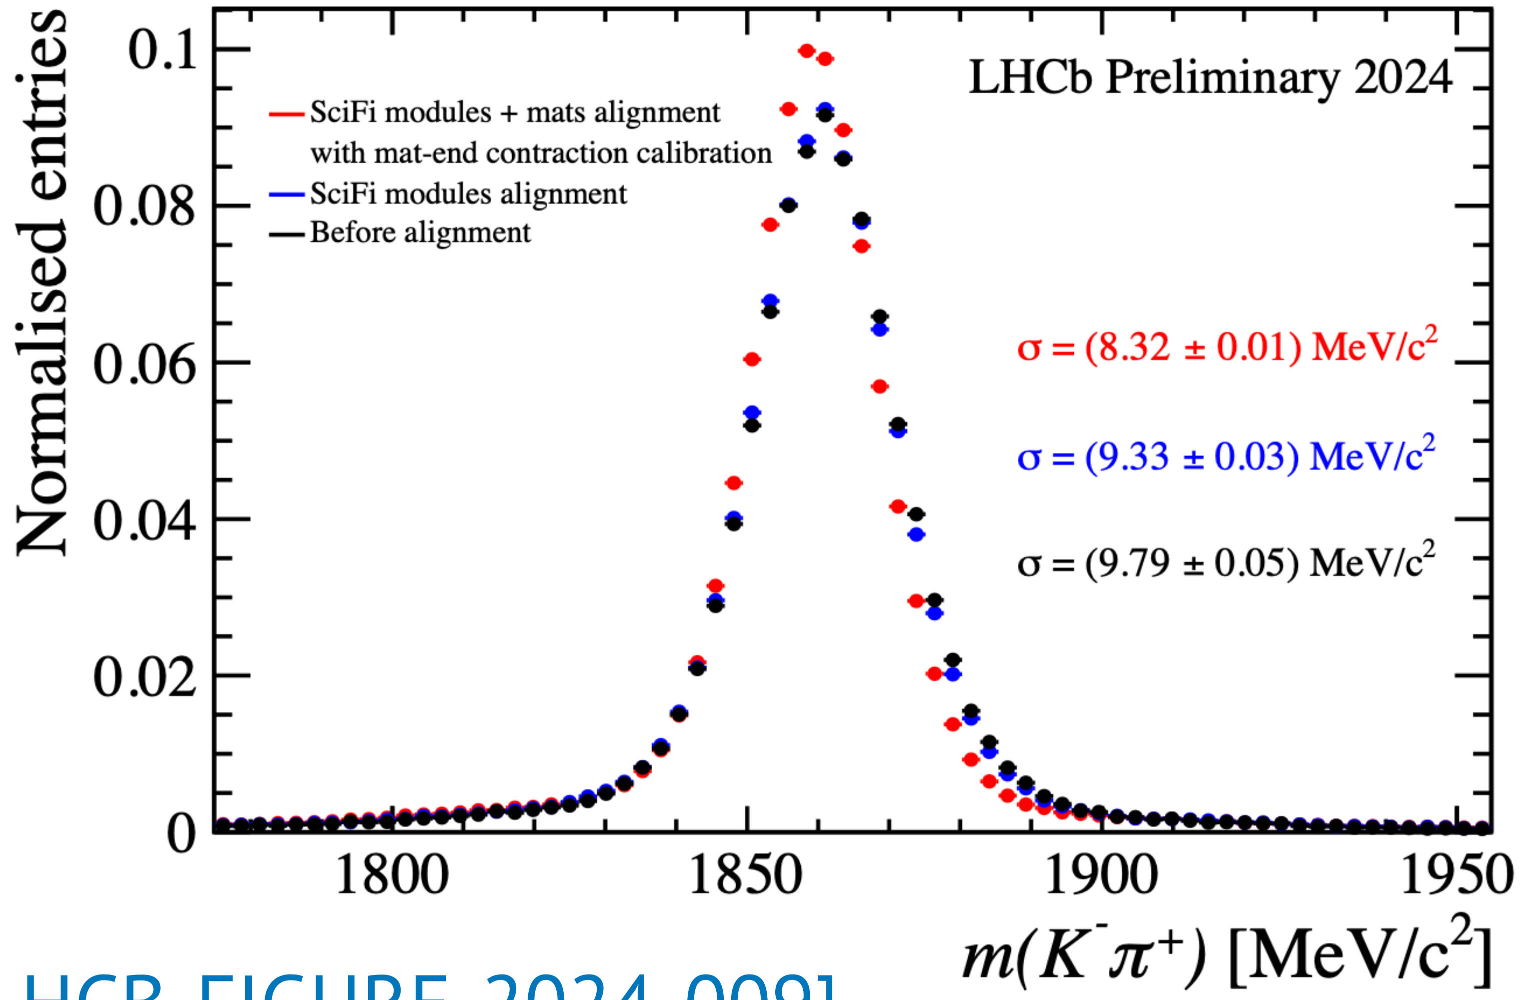

[LHCB-FIGURE-2024-009]

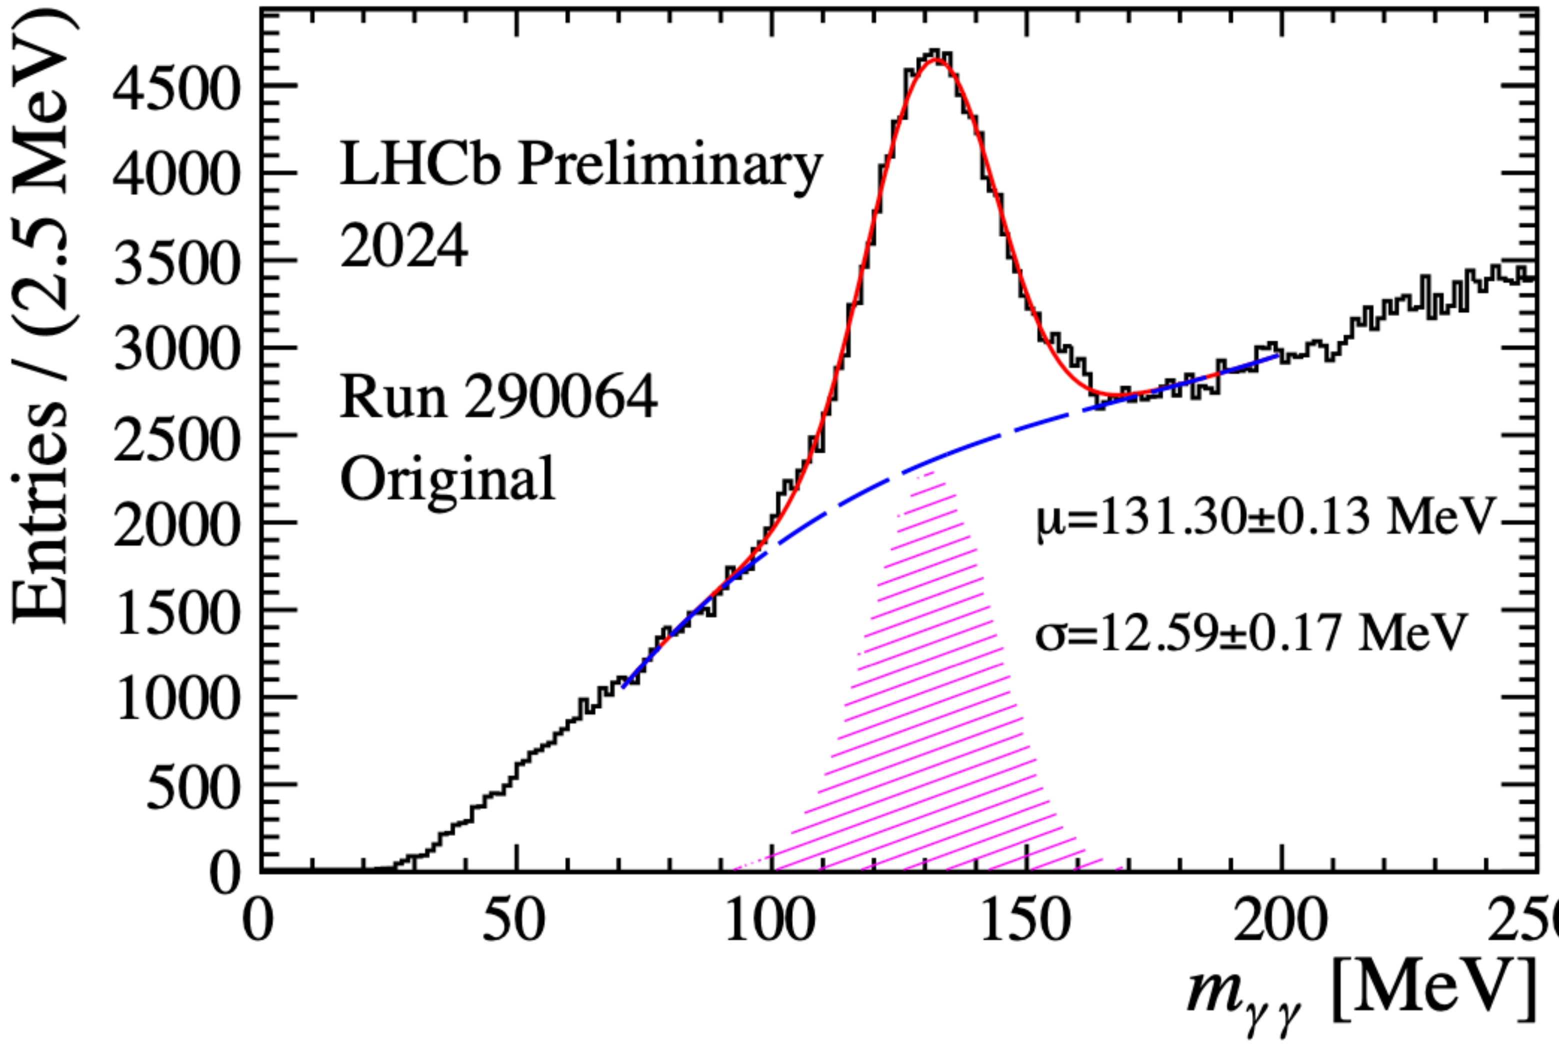

- Alignment of the tracking detectors, muon chambers and RICH mirrors
- Calibration for RICH and calorimeters
- Performed each LHC fill or more frequent

HLT2

- **Full event reconstruction** (including PID) at ~0.5MHz
- Dedicated trigger selections representing the broad LHCb physics programme

  → O(2700) selections developed by physics analysts

  - Bandwidth limit of 10GB/s
- Developments to meet throughput requirement:

  Event model as Structure of Arrays (SoA) collections

  → fewer memory accesses, vectorise reconstruction algorithms

  → matches SIMD (single instruction, multiple data) approach

  Throughput **Or**iented **(ThOr)** selections

  → Cached, no just-in-time compilation

  → Agnostic to I/O type

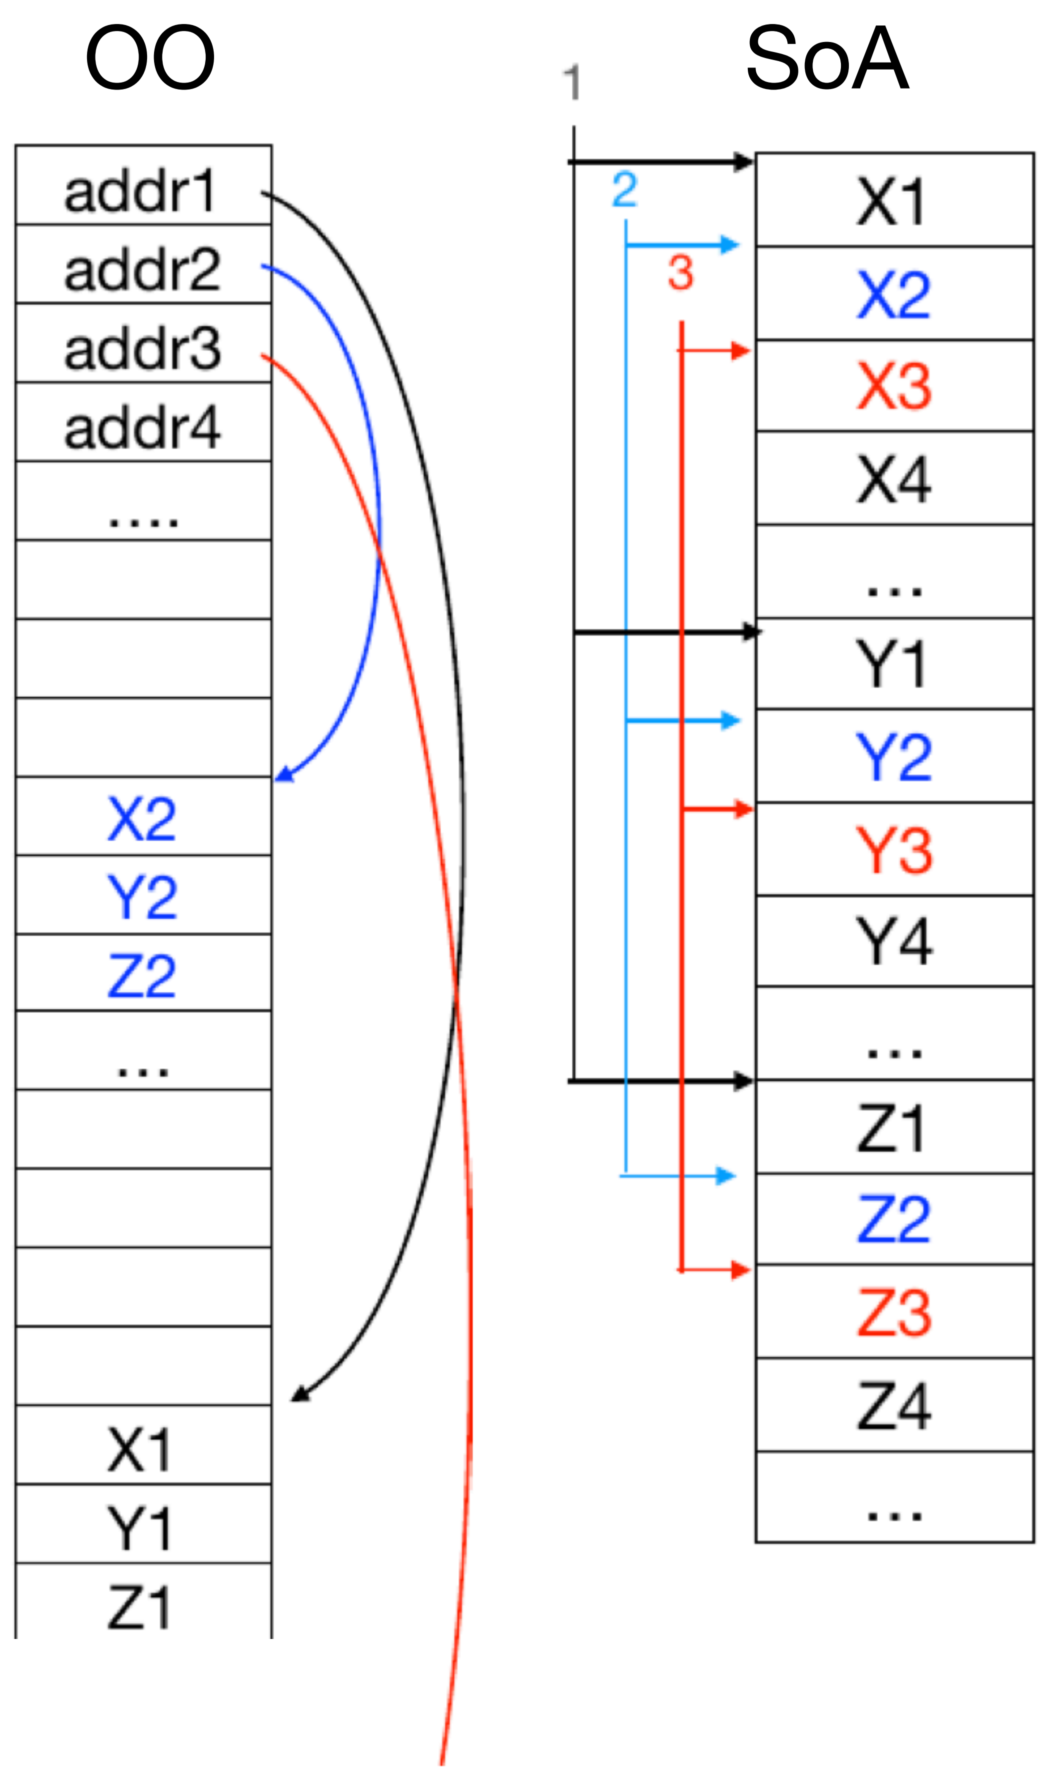

Selective persistency

- Baseline persistency approach for Run 3: **Turbo model**
- Reduce event size from O(100kB) to O(10kB)
- Requires offline-quality HLT2 reconstruction
- Hybrid solution: **Selective persistency**: save additional objects for e.g. flavour tagging, offline calibrations, isolation variables

→ ~70% of the events via Turbo model

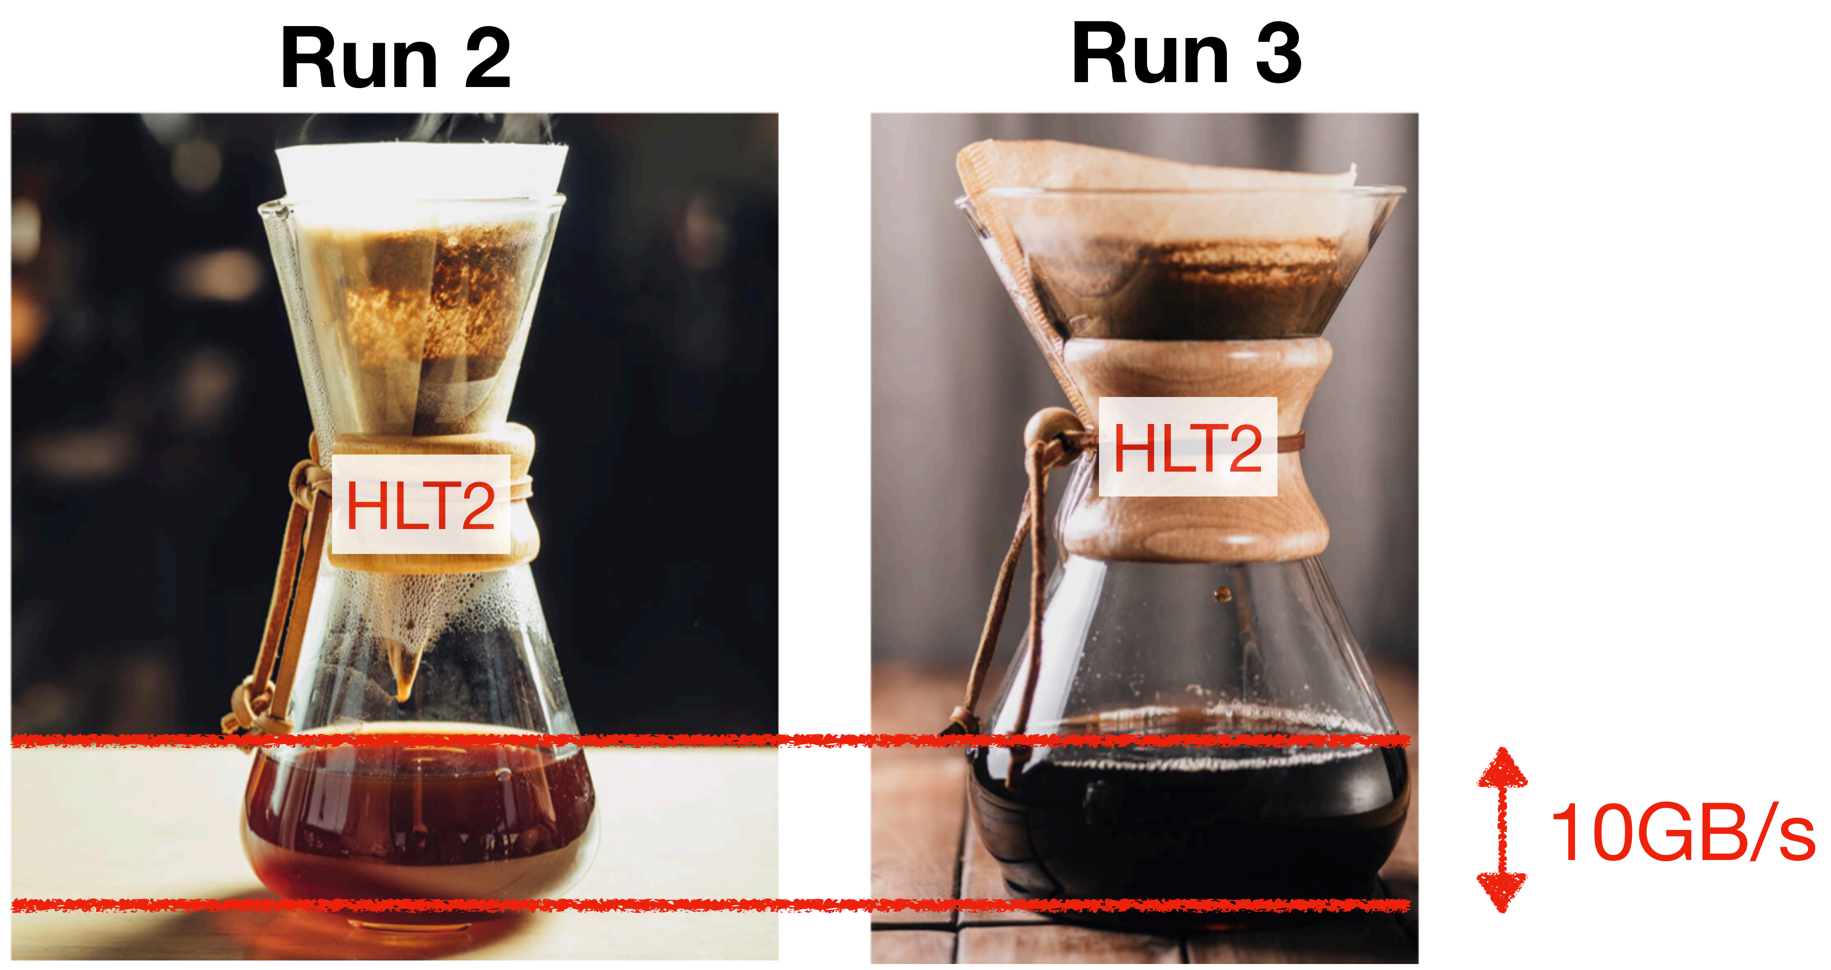

[arXiv:1903.01360]

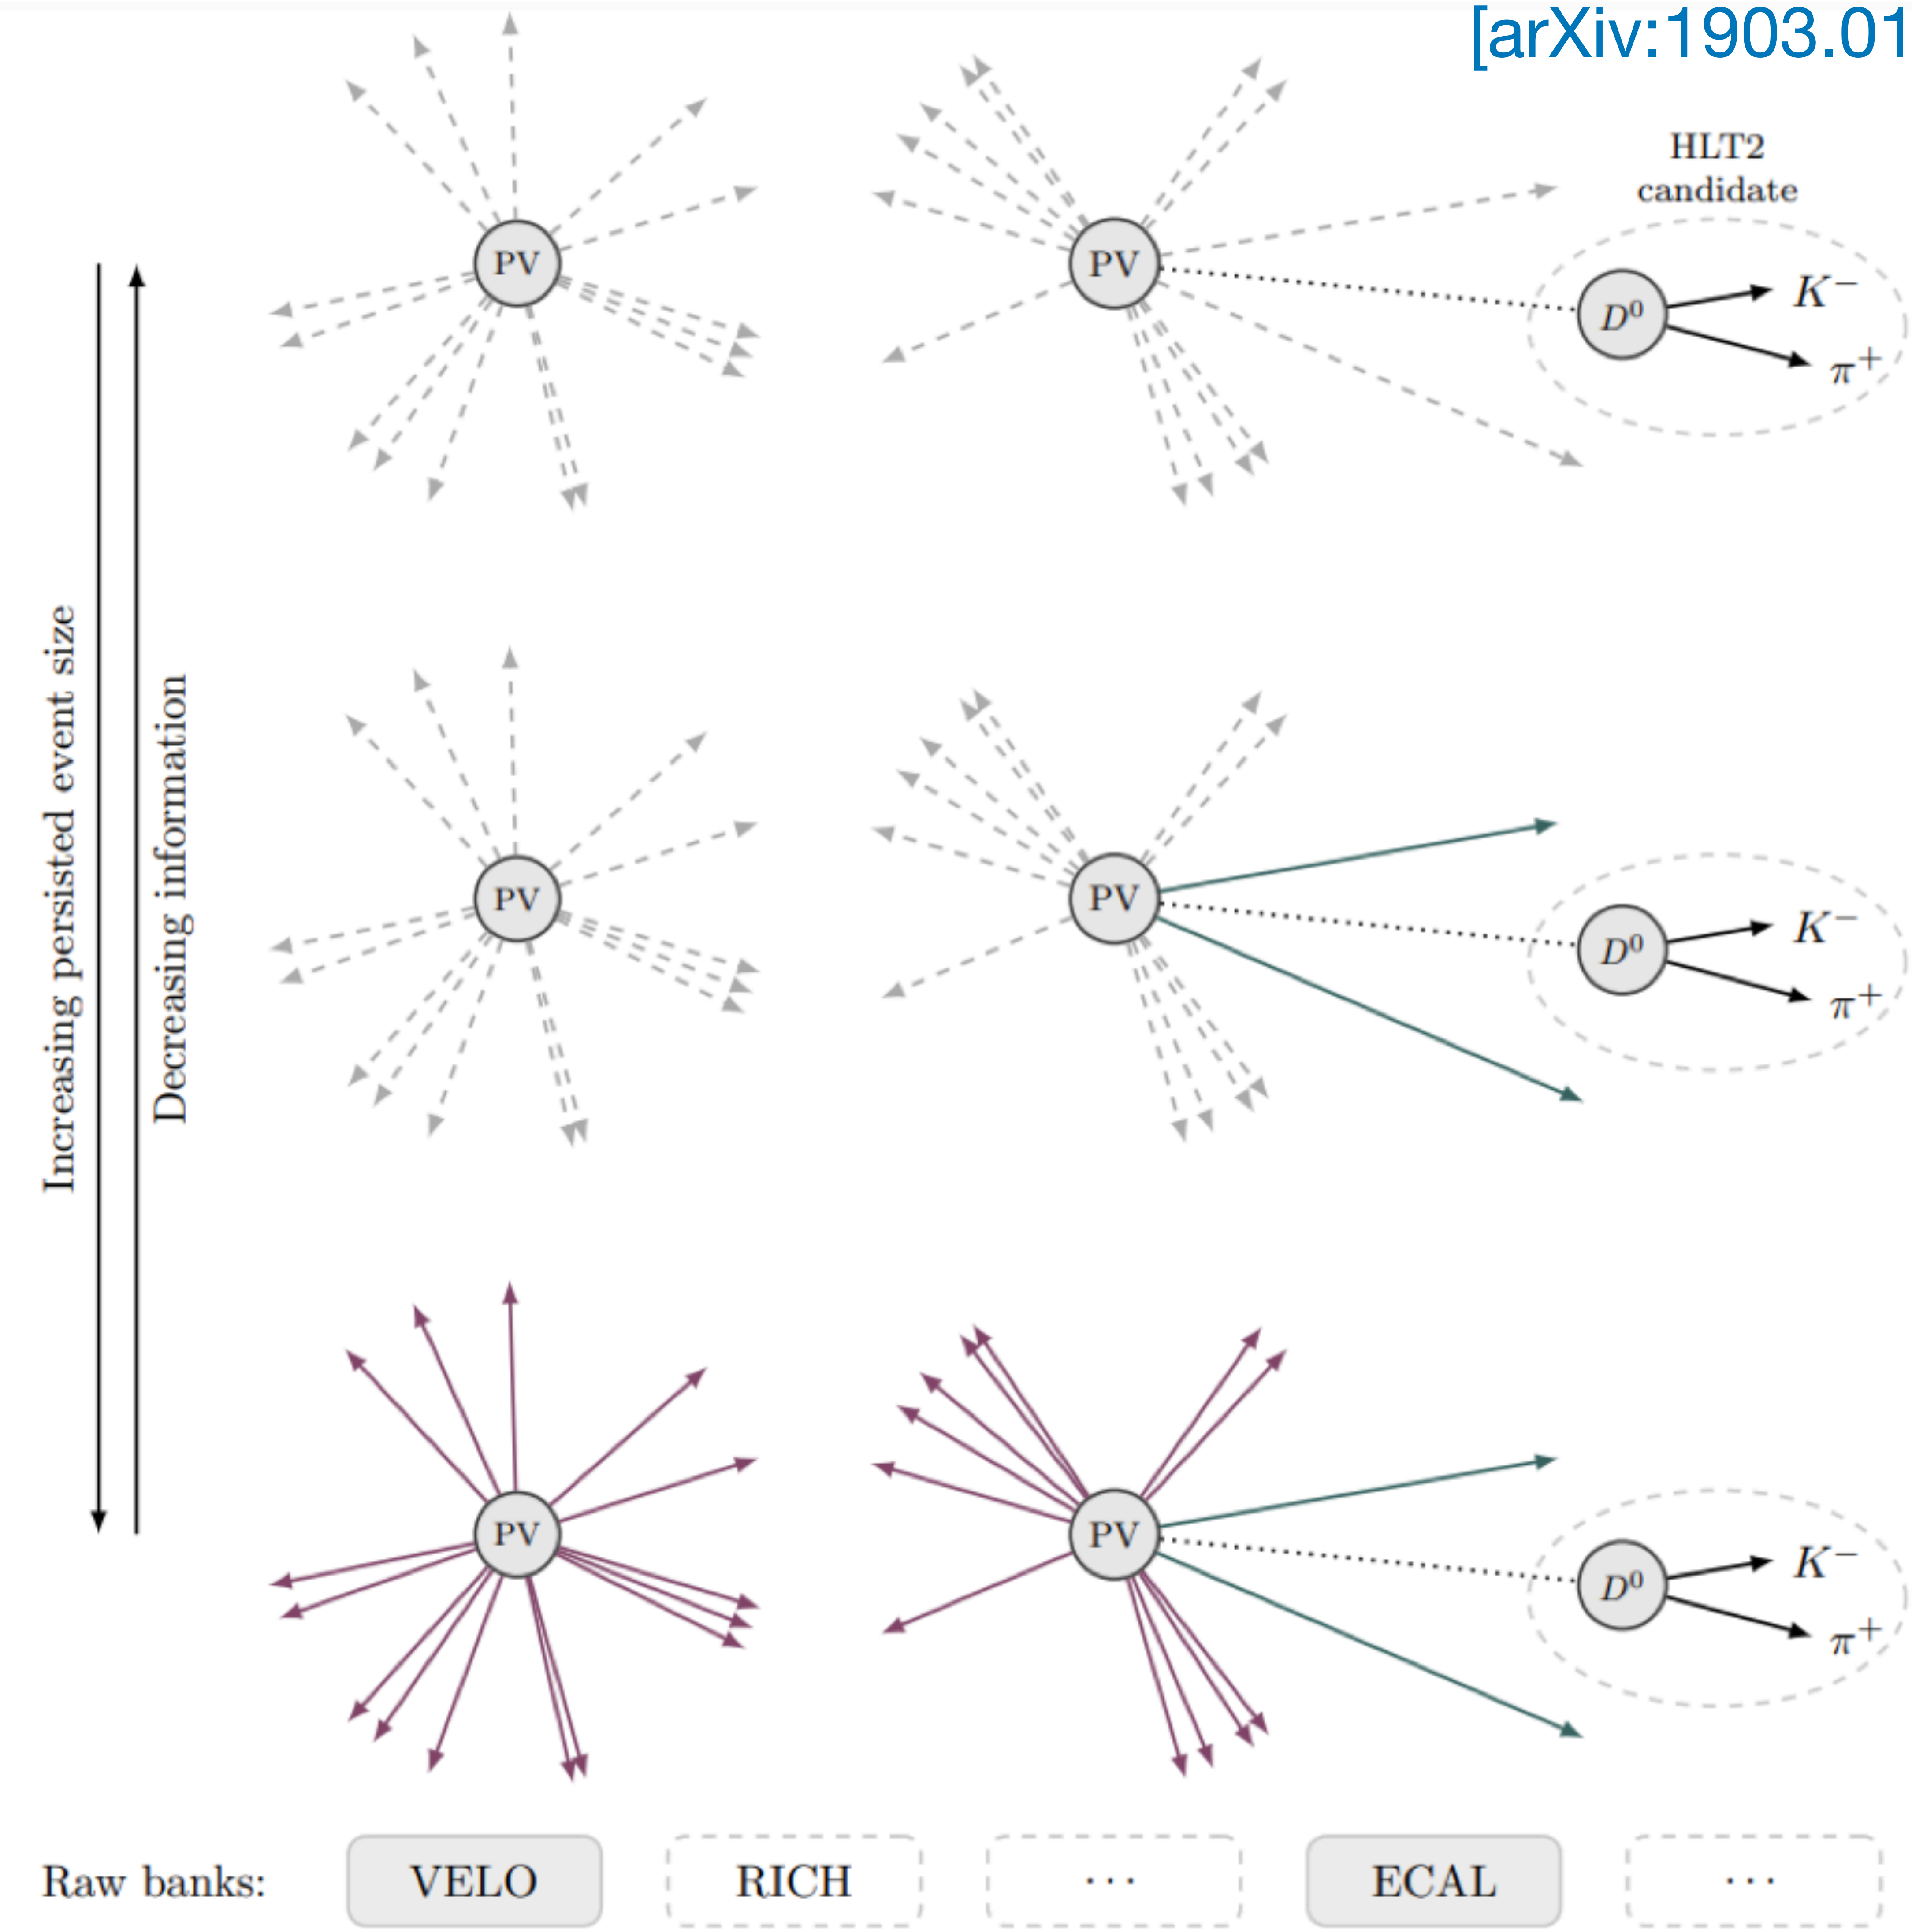

HLT2 performance

- Excellent vertex resolutions, good track reconstruction (performance with UT is being re-evaluated once UT)
- Stable PID performance for hadrons, electrons and muons

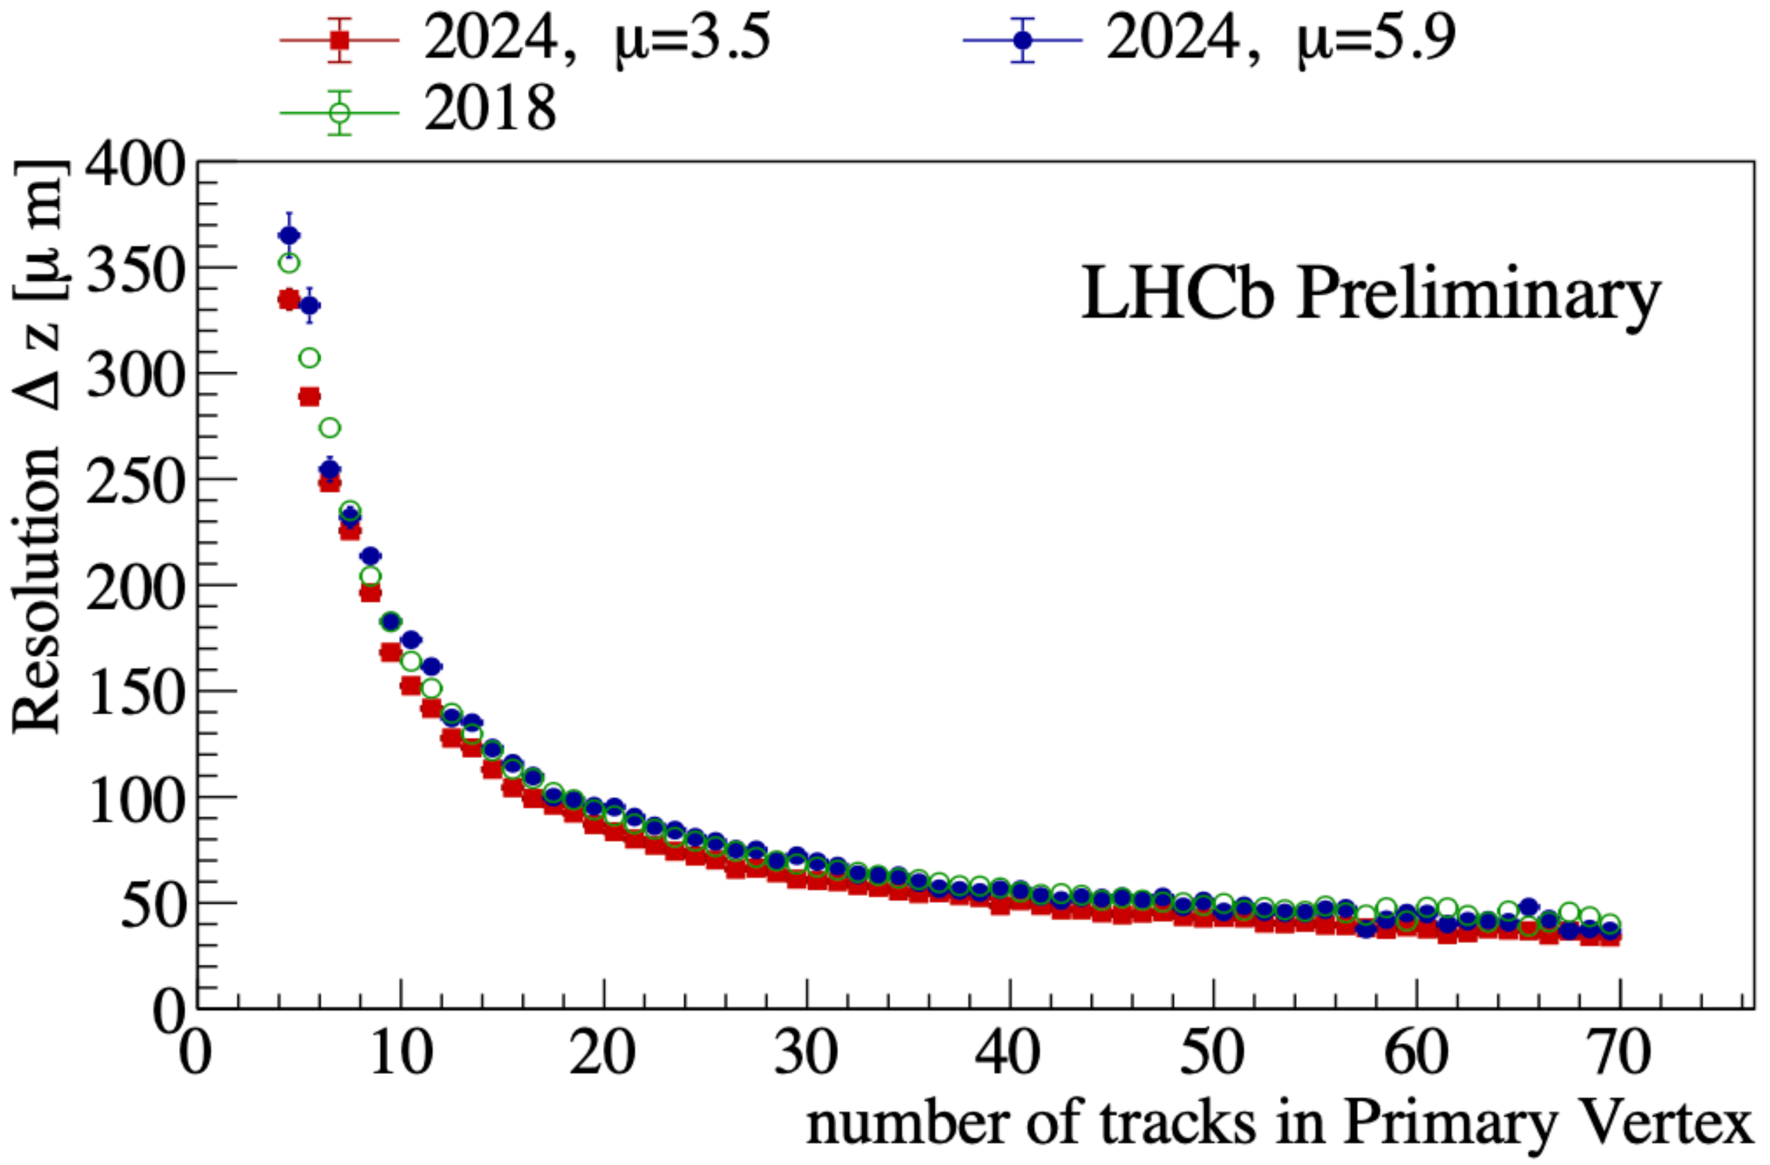

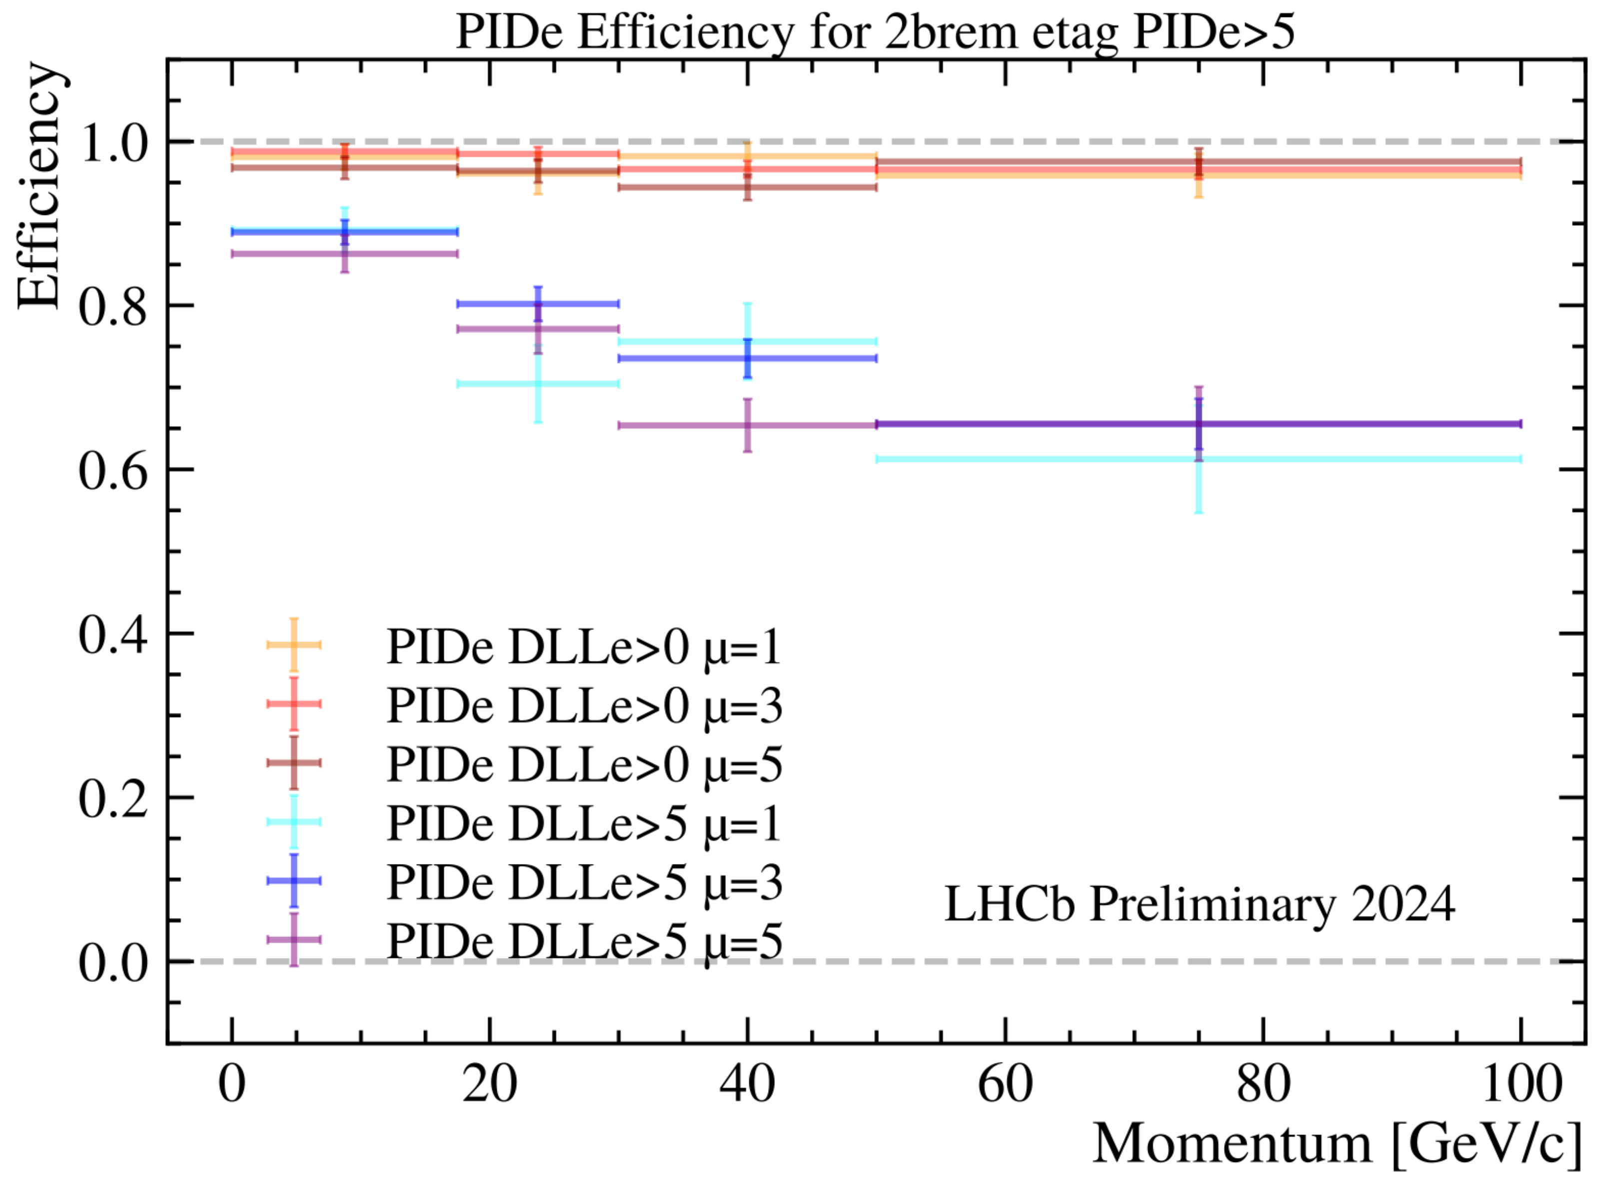

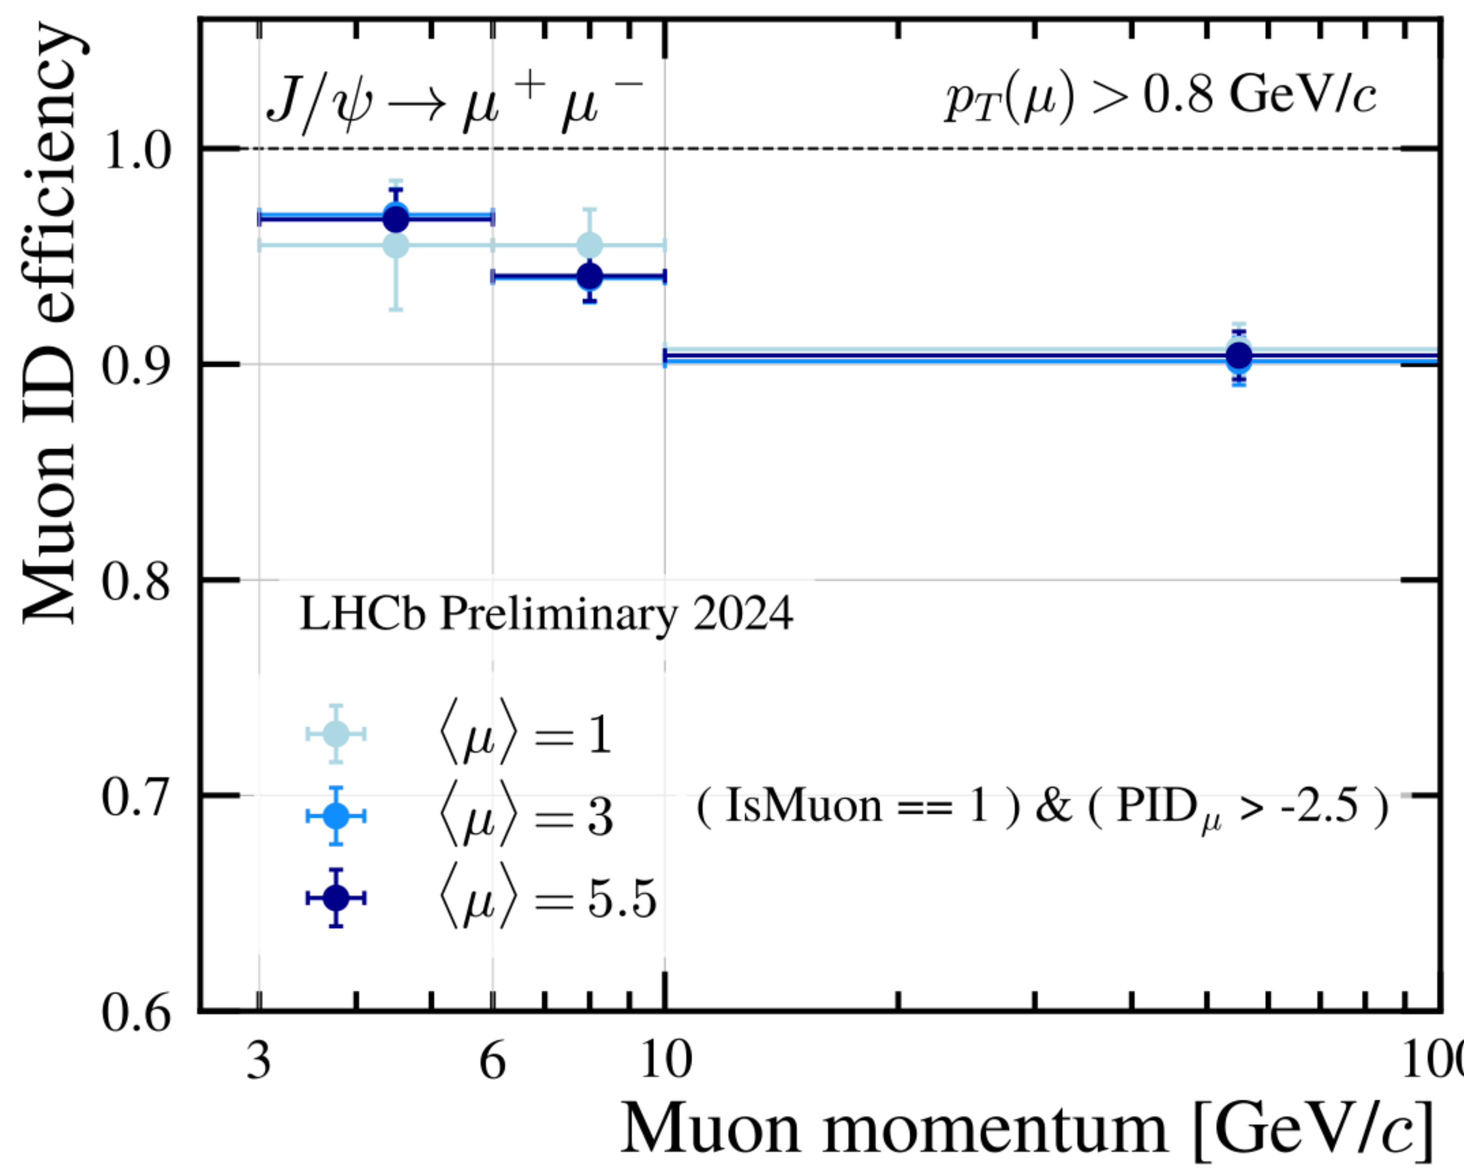

[LHCB-FIGURE-2024-010] [LHCB-FIGURE-2024-011]

Online monitoring

- Key ingredient for spotting problems quickly
  - Identify problems with hardware/data taking conditions
  - Check trigger selections
  - Flag data quality for physics analysis
- HLT1 monitors reconstruction/selection output at 30MHz: track qualities, mass peaks, trigger rates, …

  → **extension to high-level quantities**: tracking efficiencies, trigger efficiencies, muonID
- HLT2 has a real-time monitoring of the reconstruction + data-quality monitoring once it is fully run

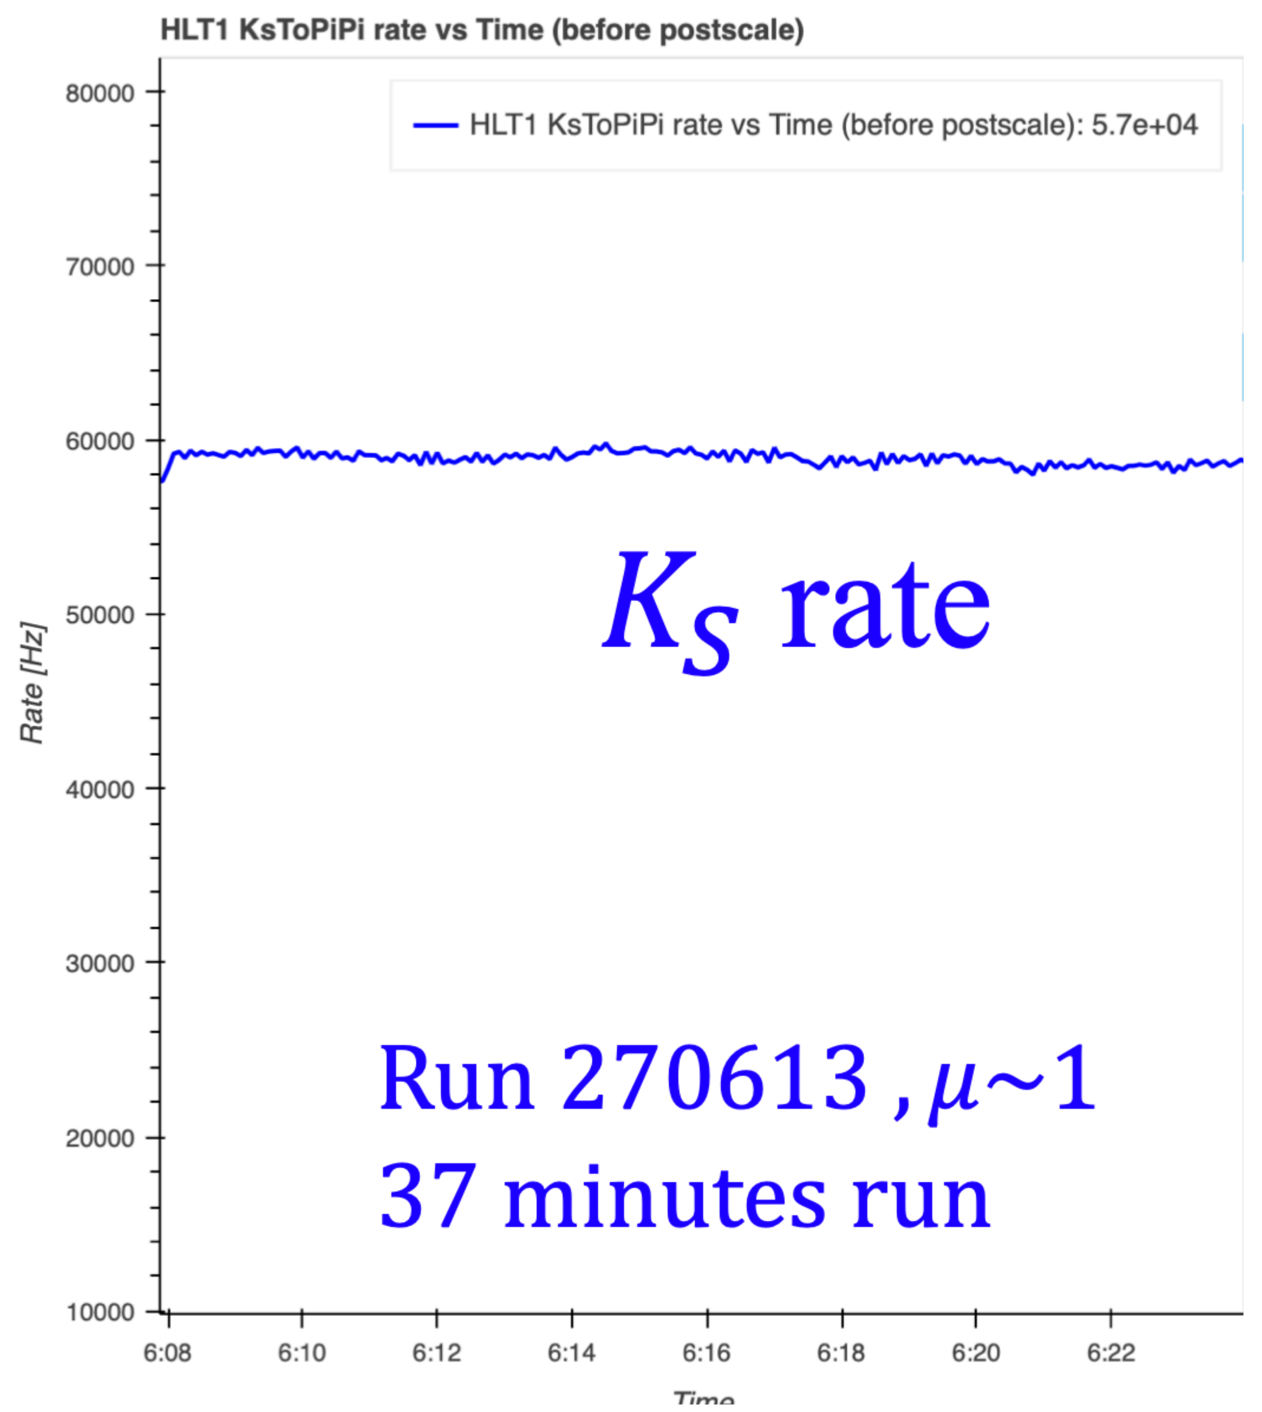

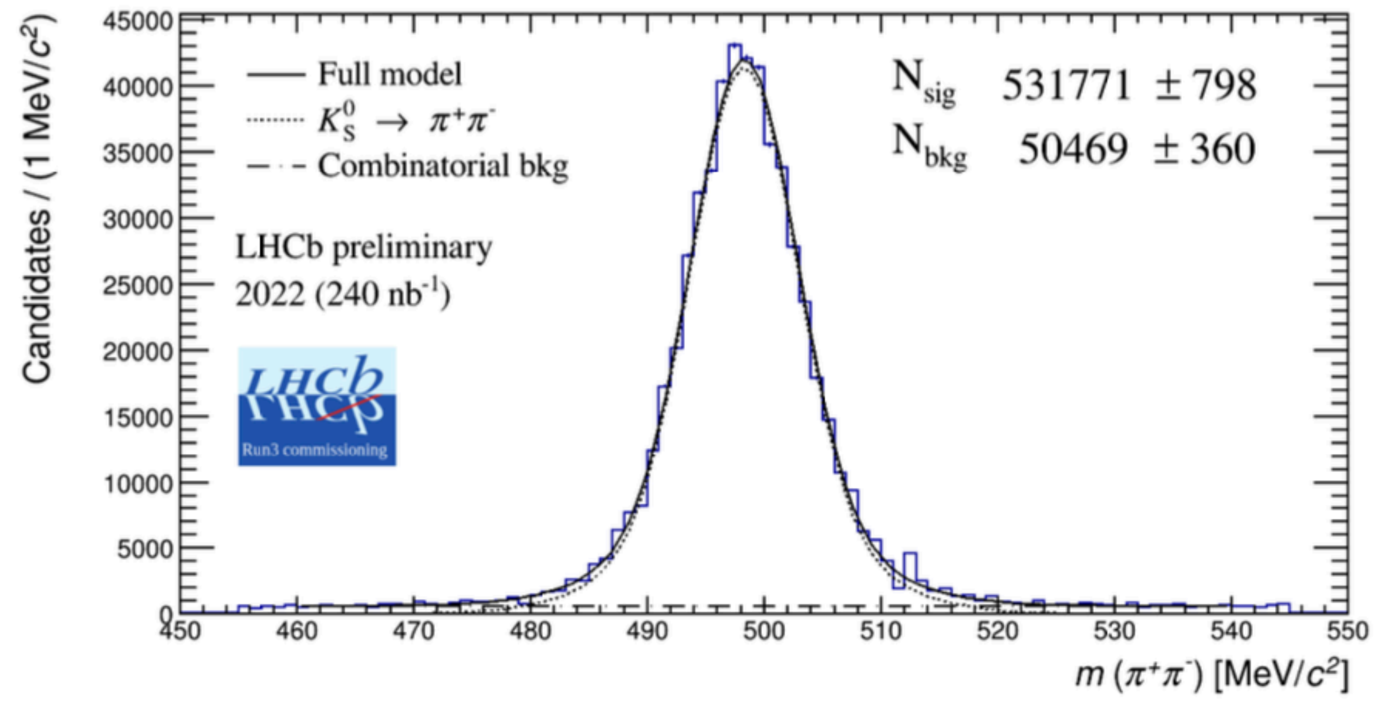

[LHCB-FIGURE-2023-005]

Conclusions

- L0 Hardware trigger has been removed for LHCb Run 3
- **Fully software-based trigger system runs successfully at 30MHz**
- HLT1 implemented on GPUs
- Real-time alignment and calibration enables an offline quality for the HLT2 reconstruction
- HLT2 developments meet throughput and rate requirements and allows for **higher physics output through Turbo persistency model**
- Challenges in this years data taking: going up in pile-up, inclusion of the UT
- $3\ \mathrm{fb}^{-1}$ of physics data recorded, $9\ \mathrm{fb}^{-1}$ planned for this year

  → Stay tuned for a new era of LHCb results!

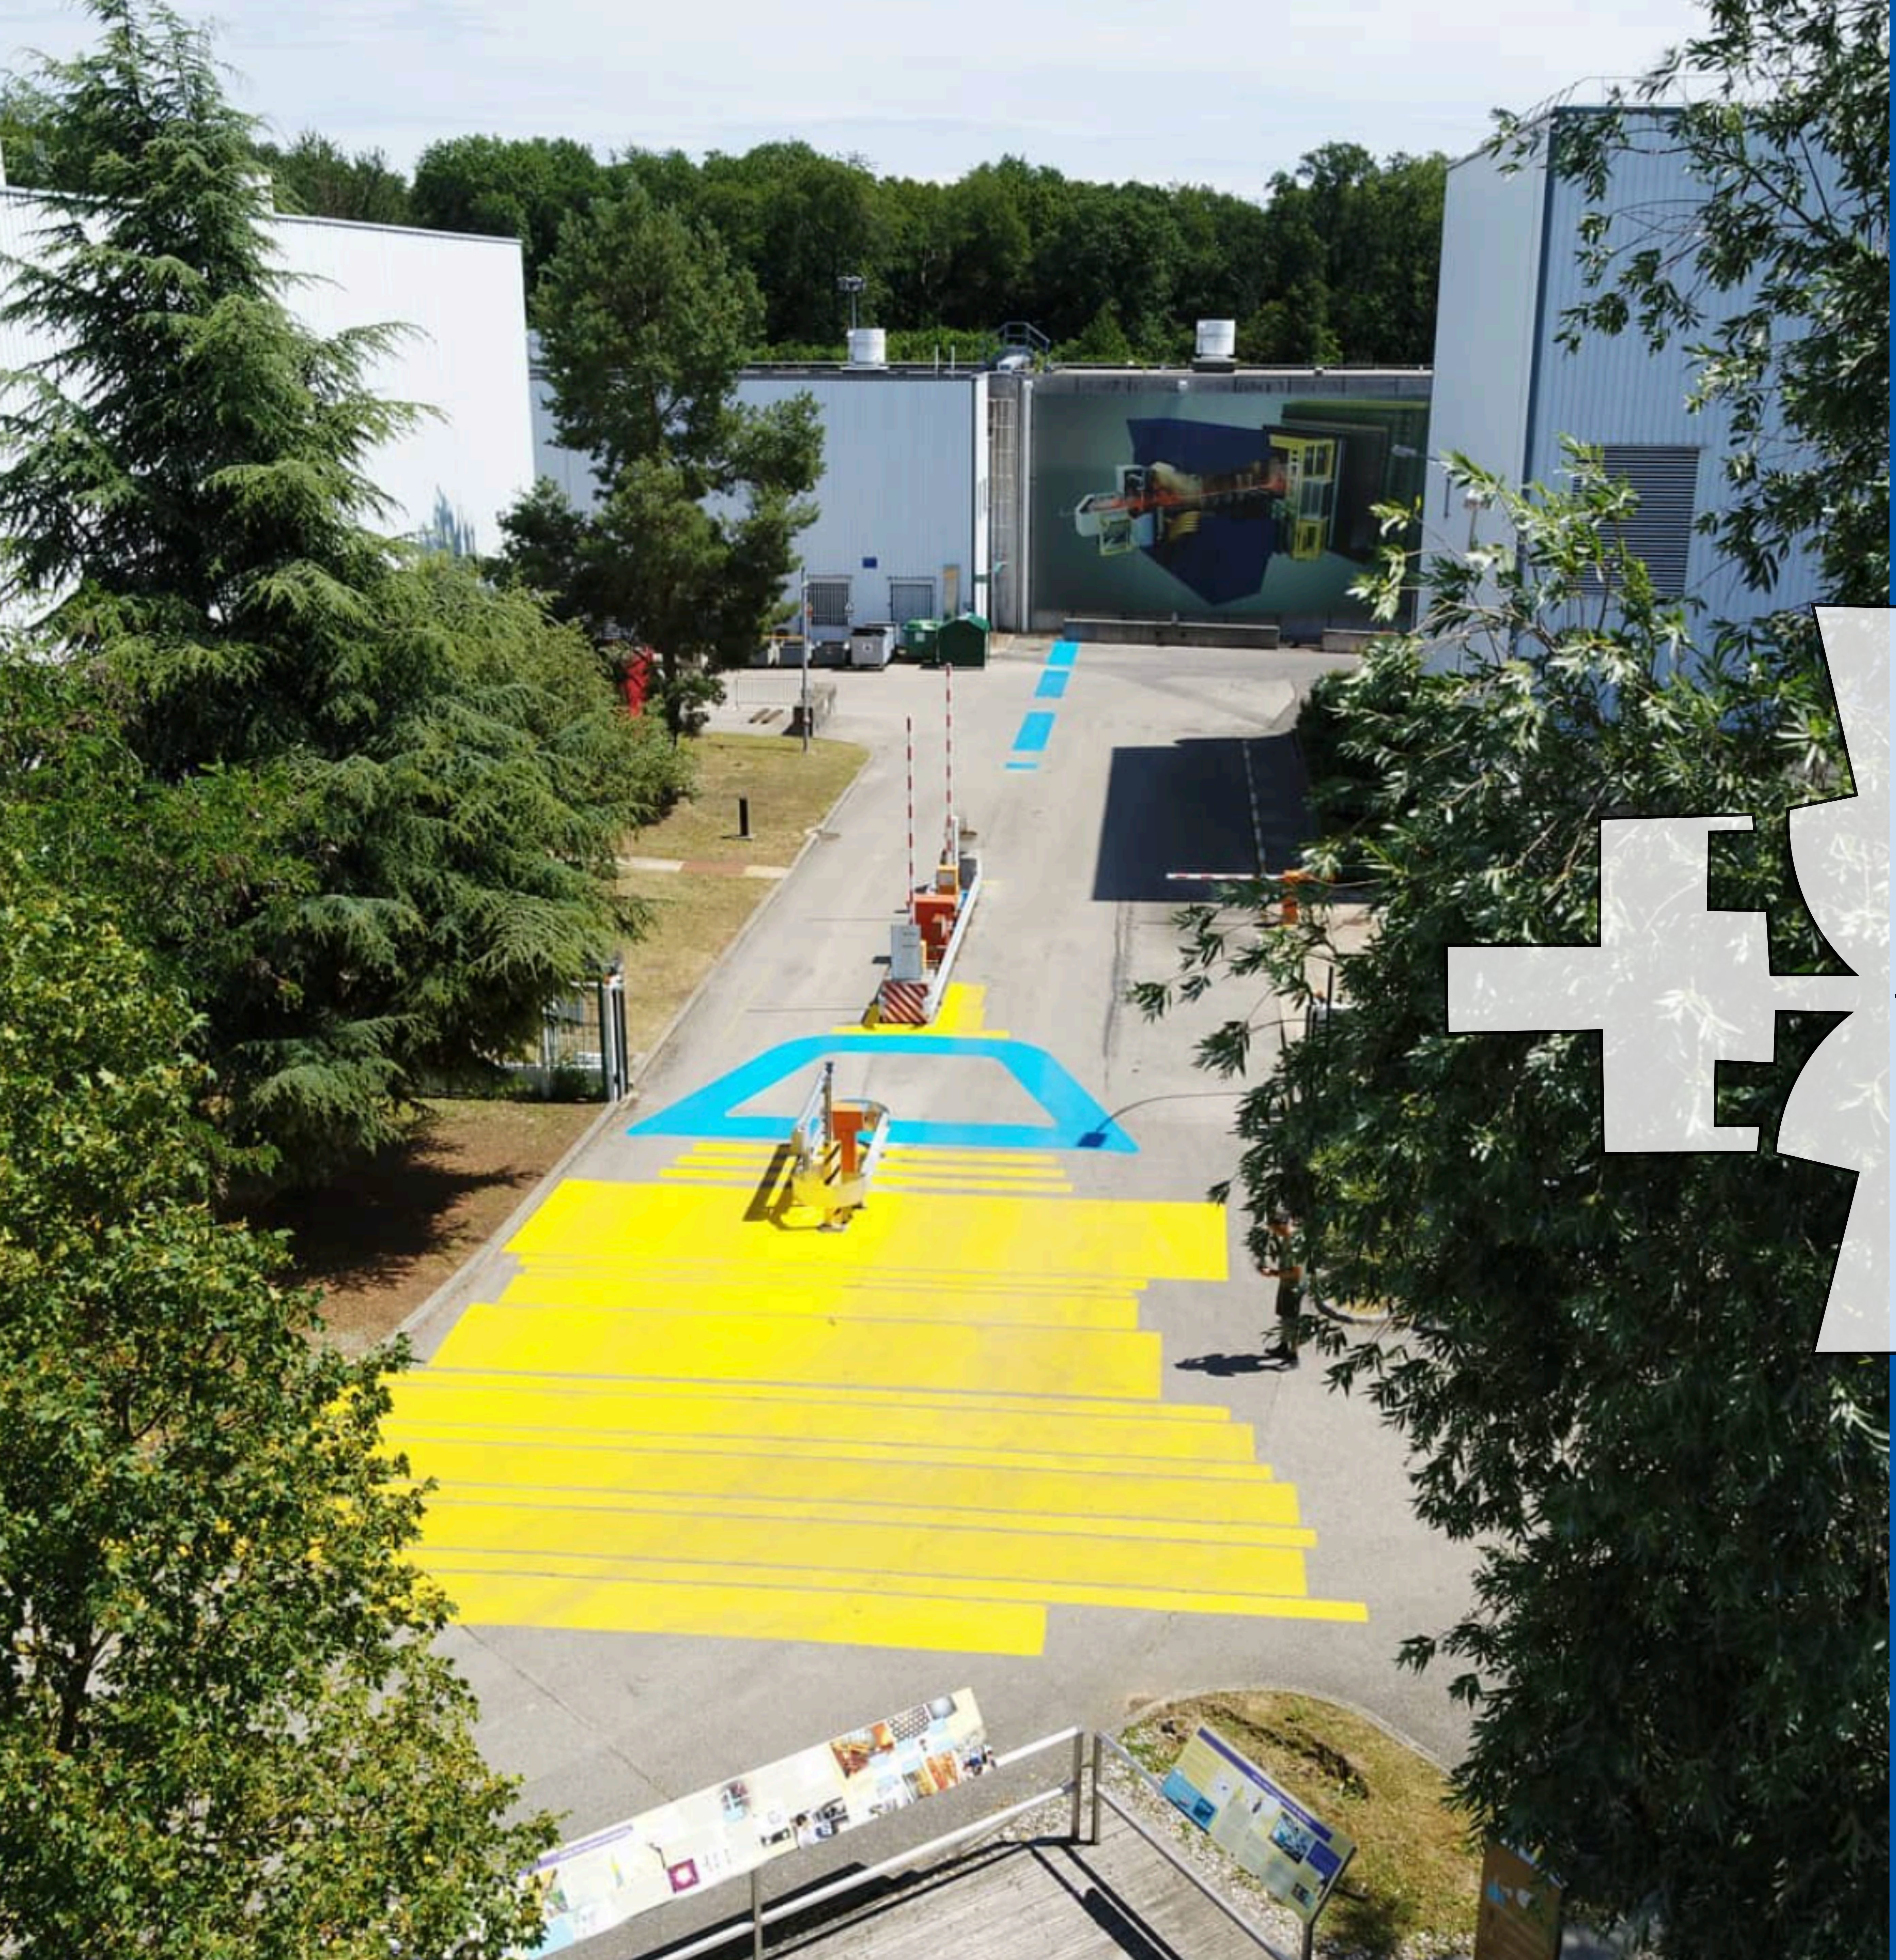

BACKUP SLIDES

Heterogeneous computing architecture

Sub-detector front-end electronics:
- Read-out at 40MHz

173 **Event Building (EB) Servers**:
- 3 Back-end readout boards (PCIe40) based on FPGAs per EB node
- Packing data fragments of the sub-detectors together (**event building**)
- Packing ~1000 events together to multi-event packages
- 2 GPU cards for **partial event reconstruction and selection (HLT1)**

~40PB buffer storage system

Event Filter Farm:
- CPU servers for **full event reconstruction and selection (HLT2)**

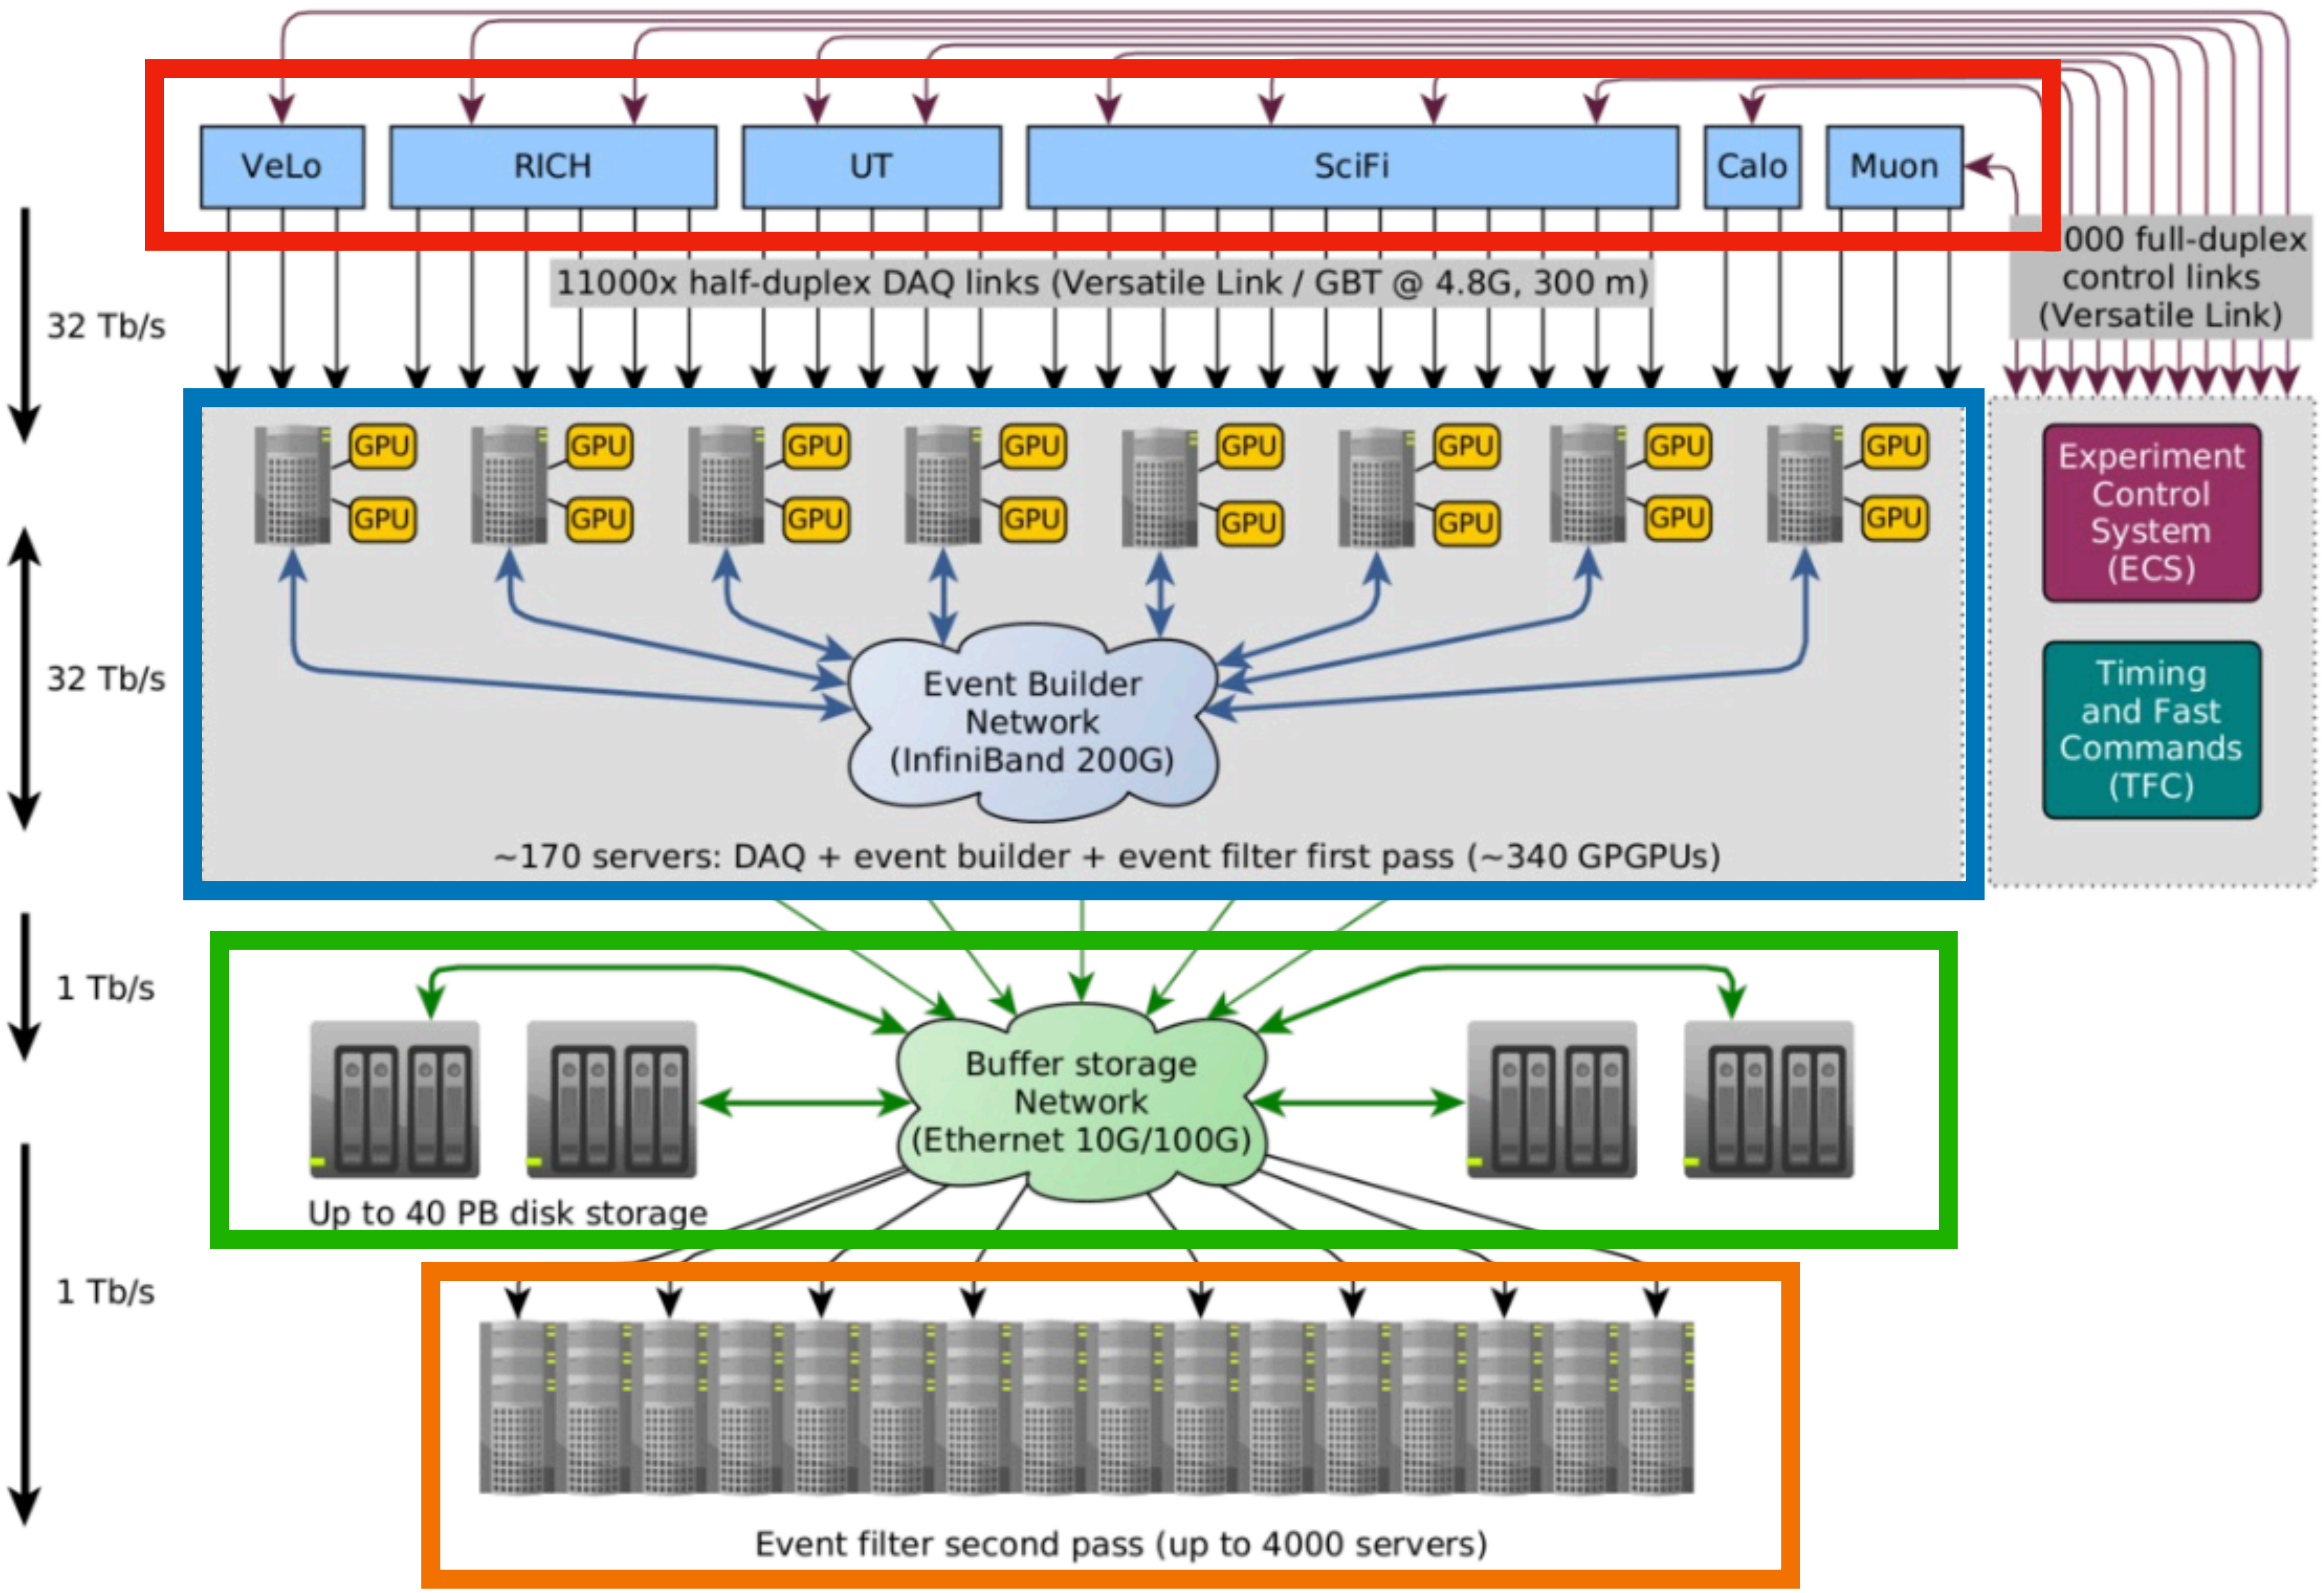

MuonID performance Run 2 & Run 3

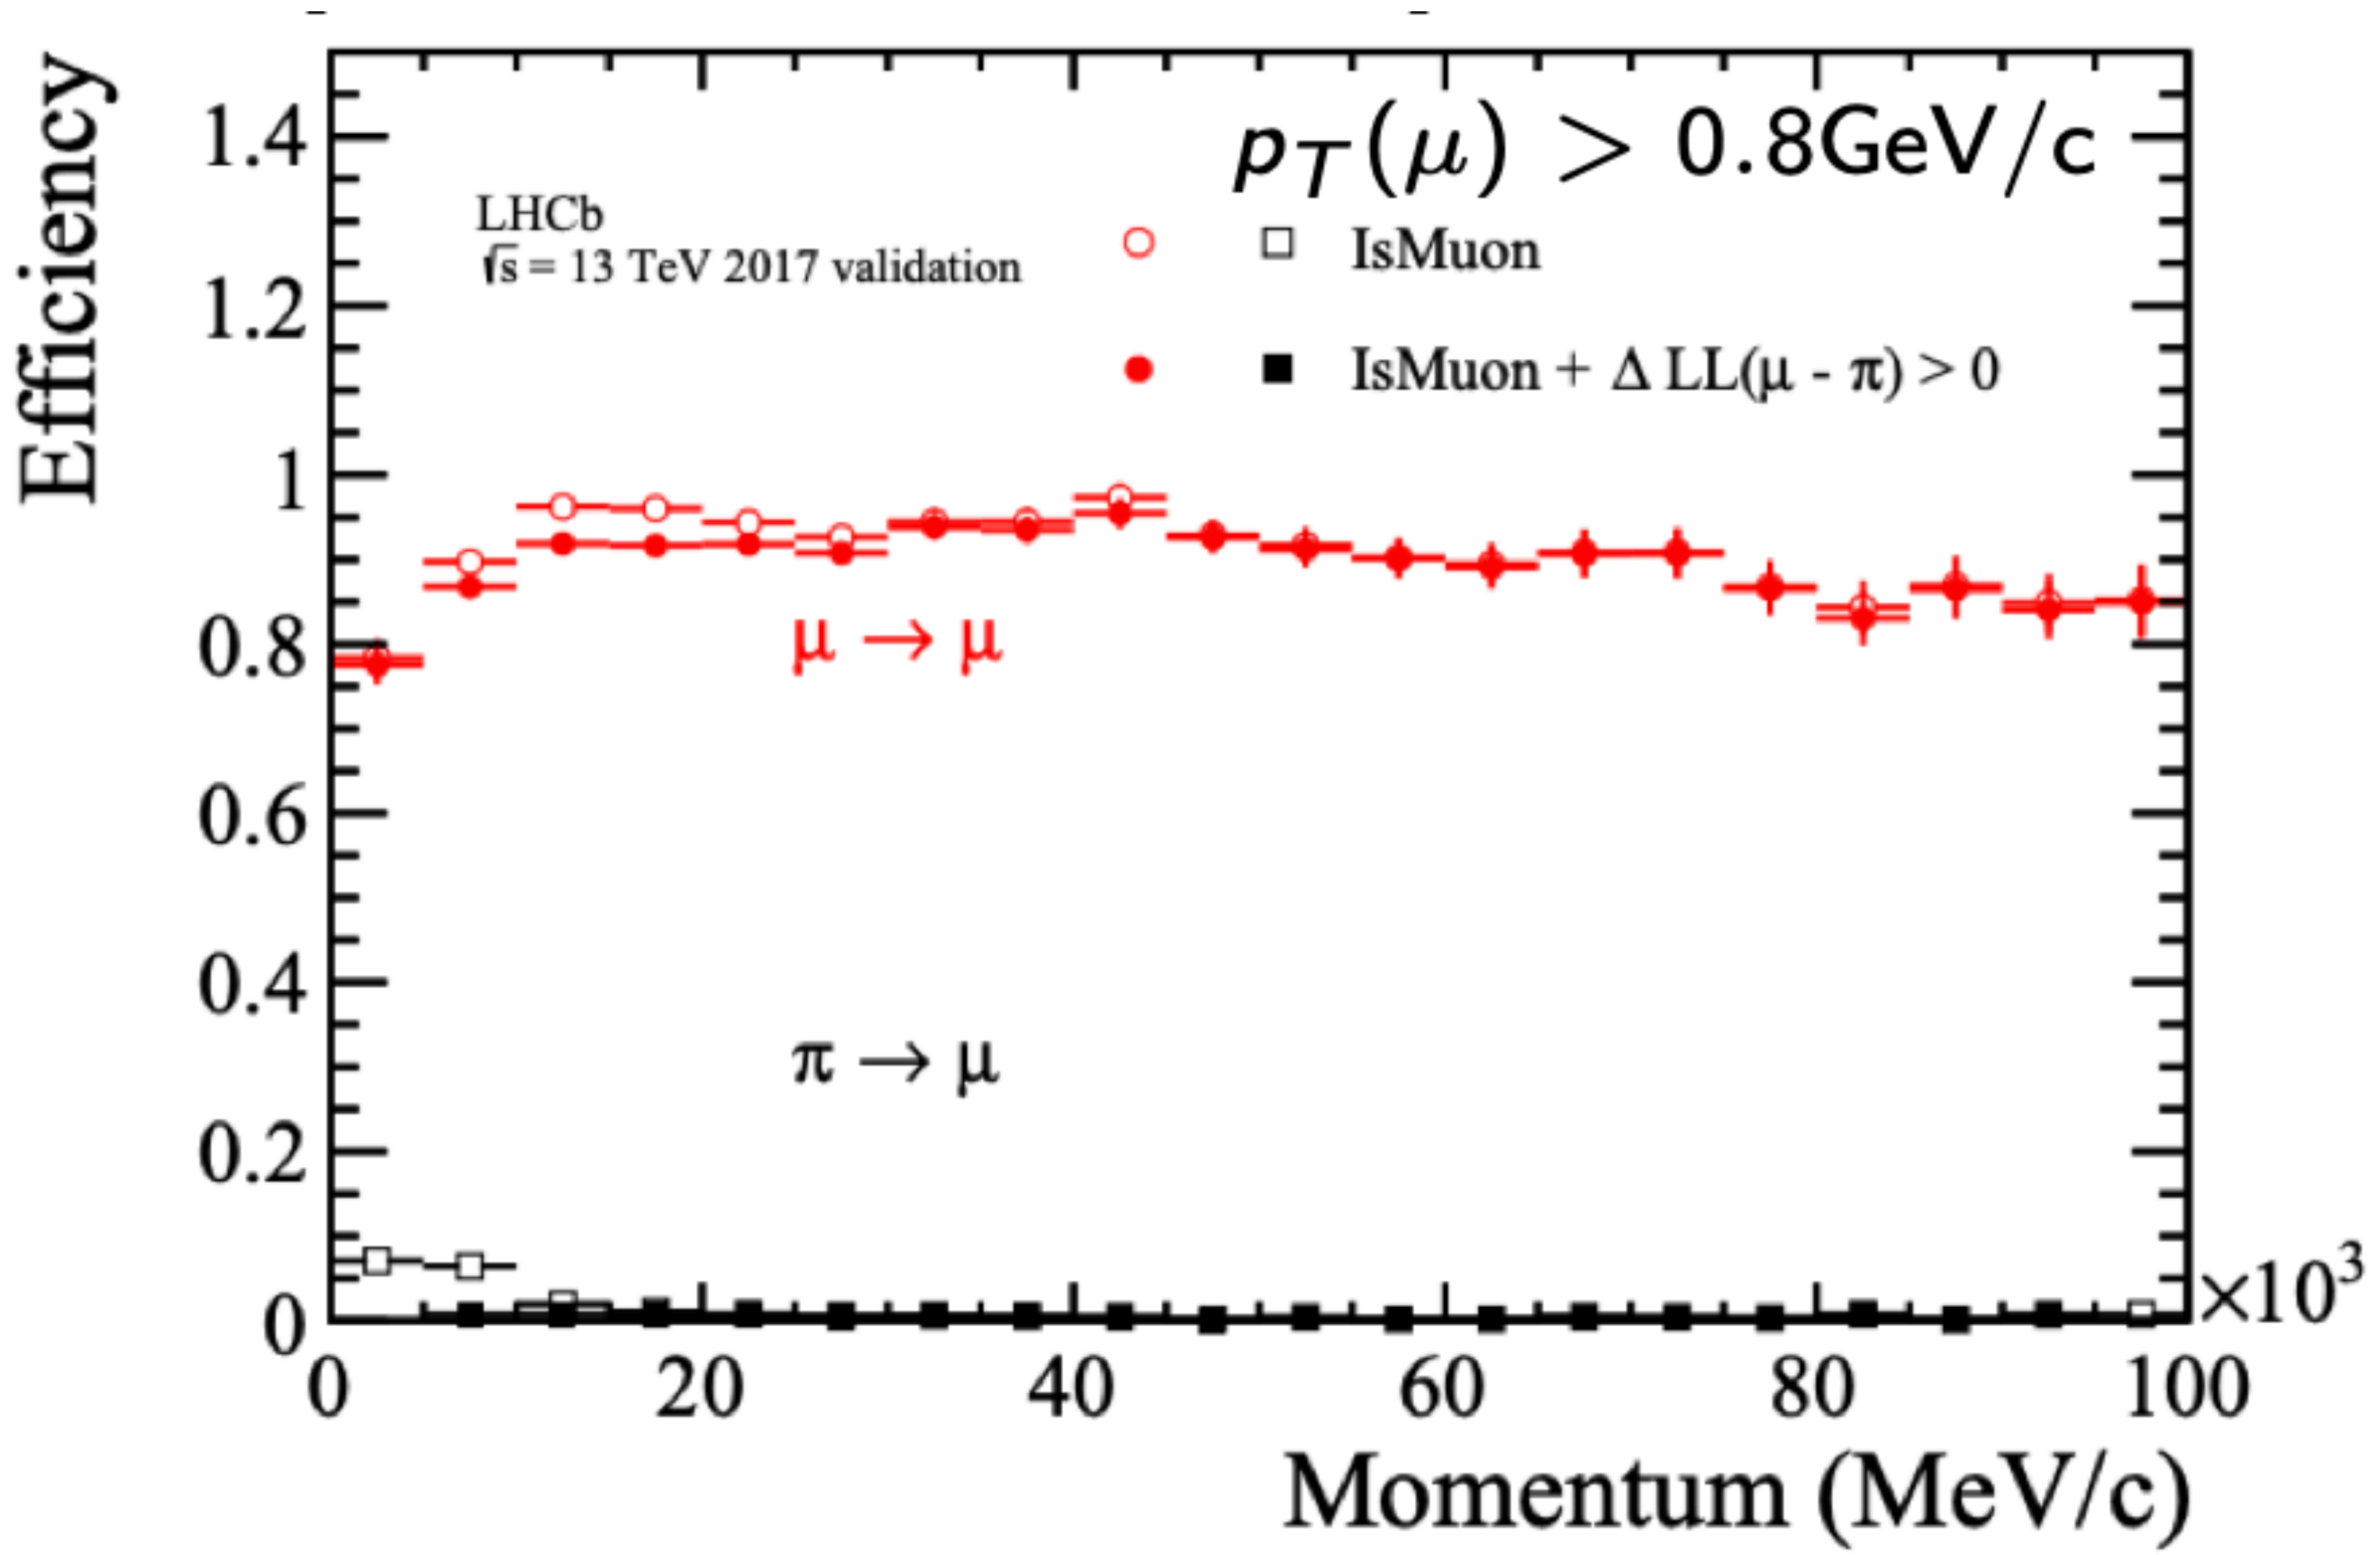

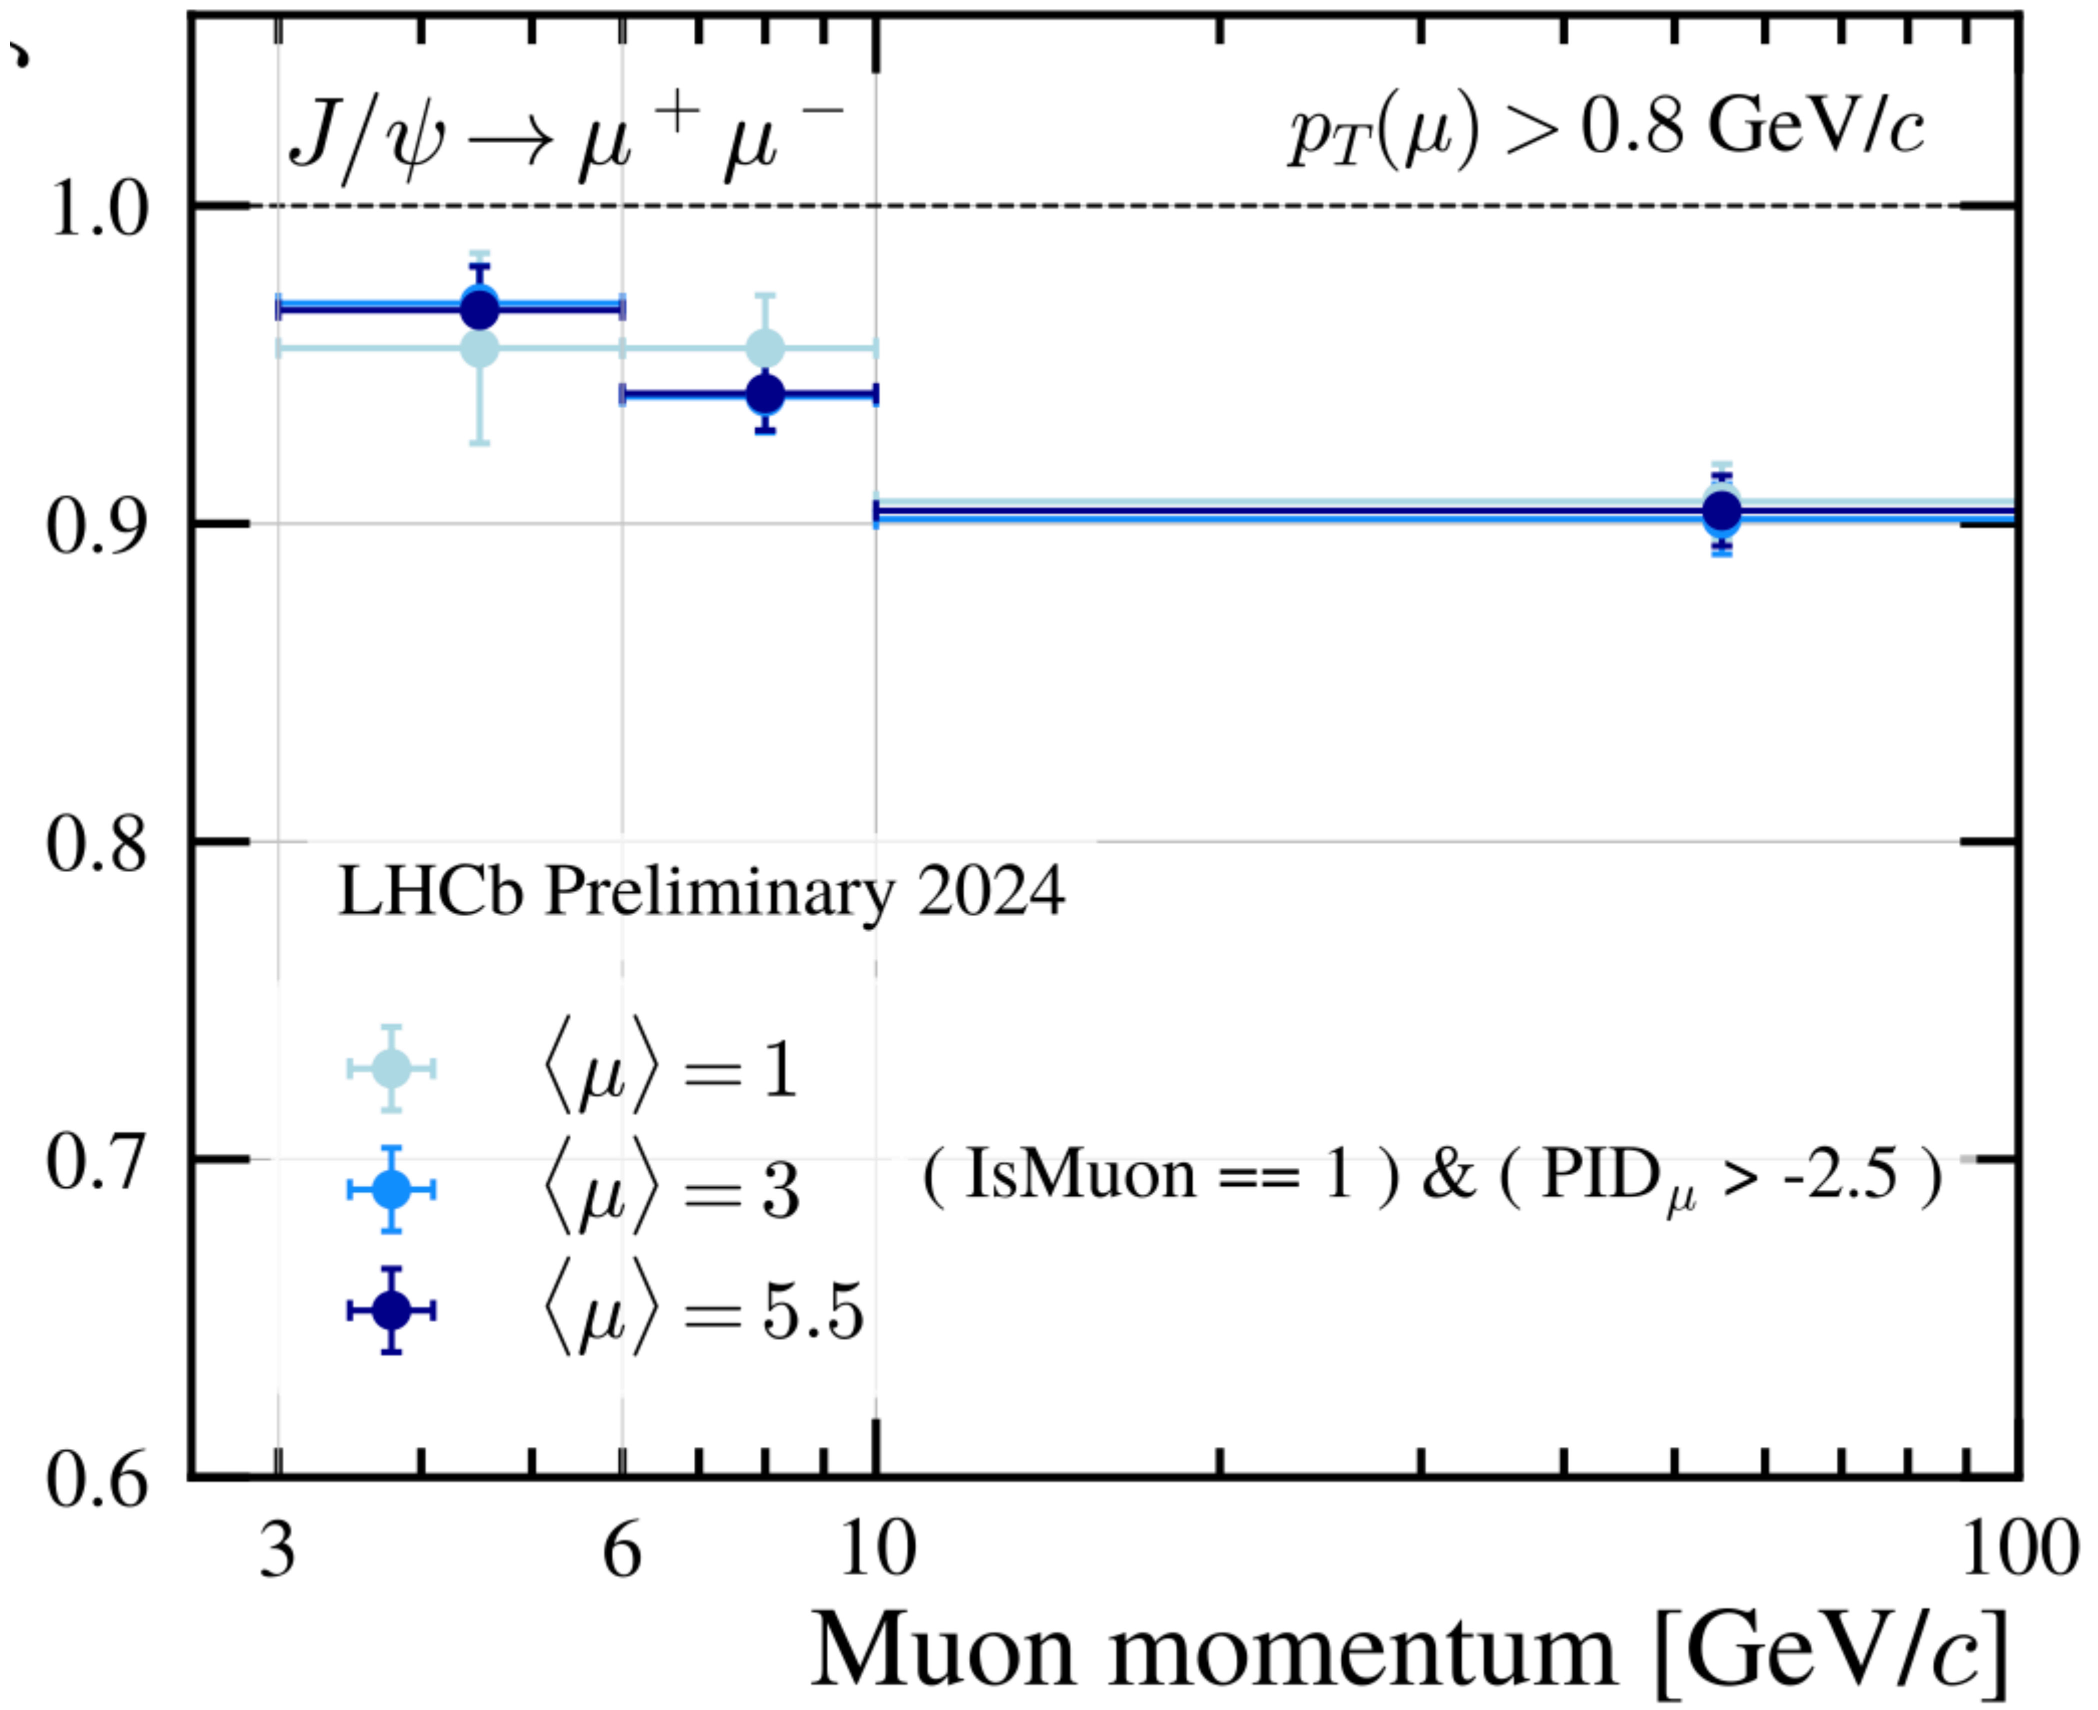

ElectronID performance Run 2 & Run 3

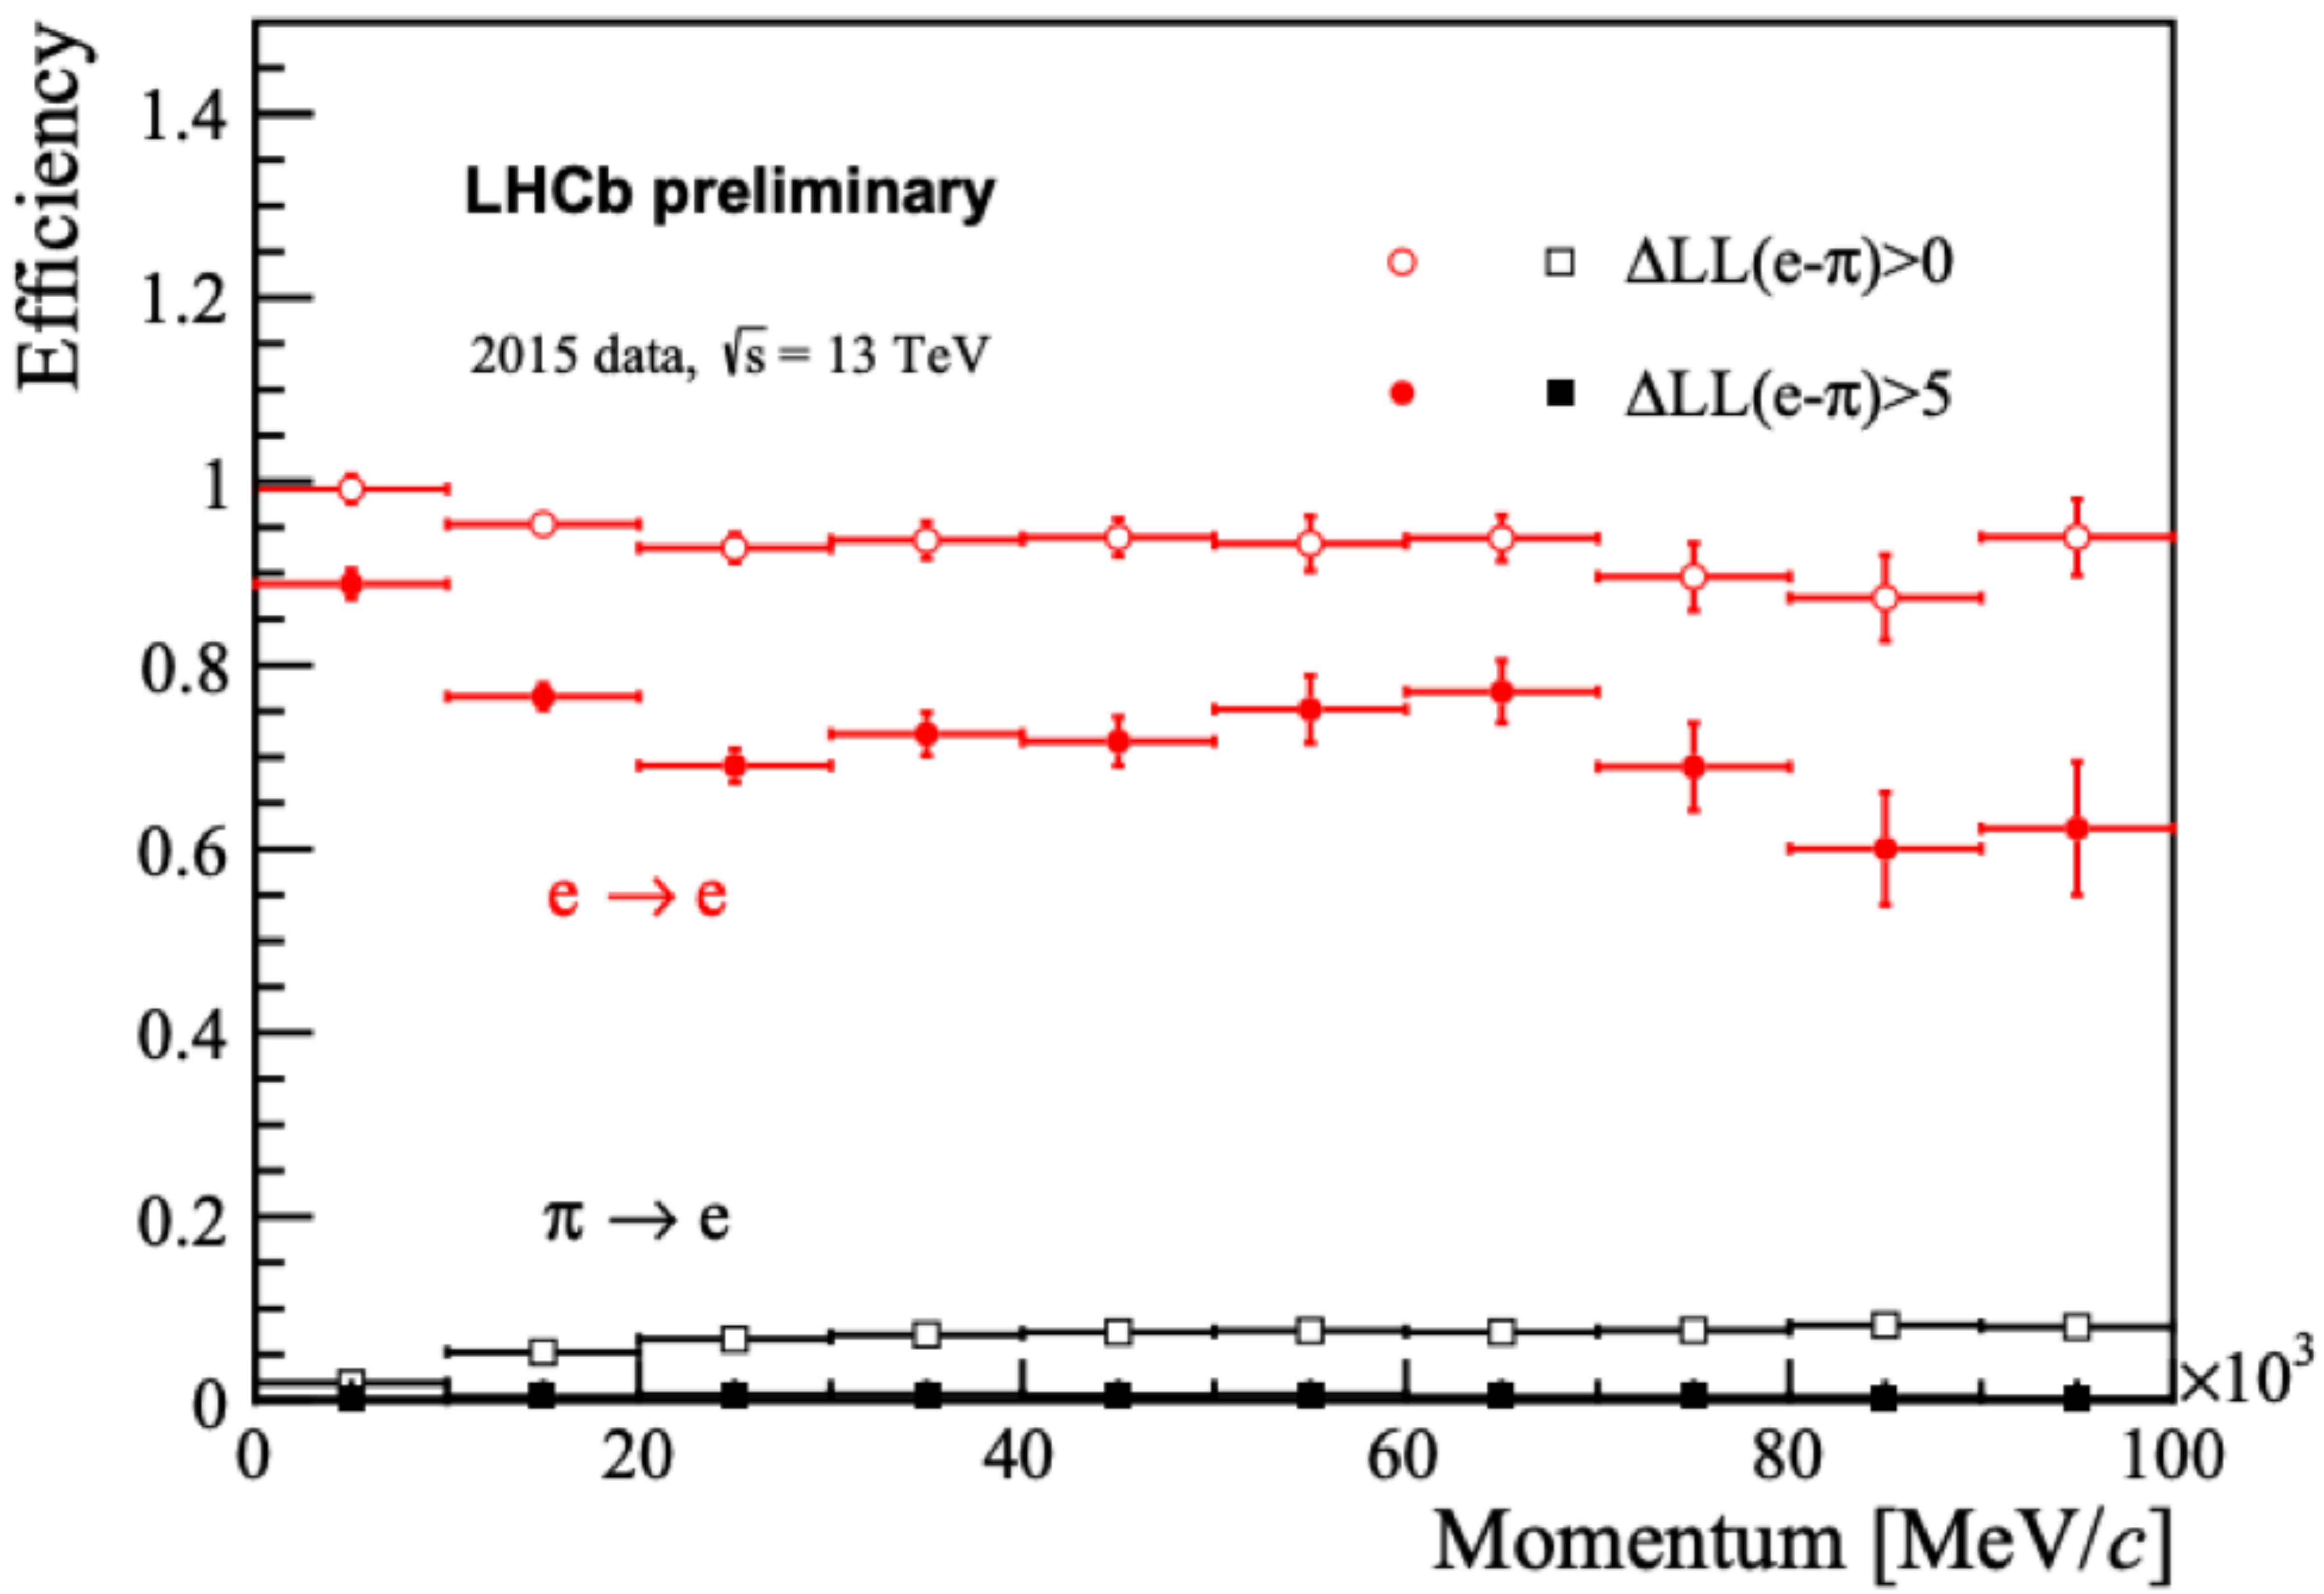

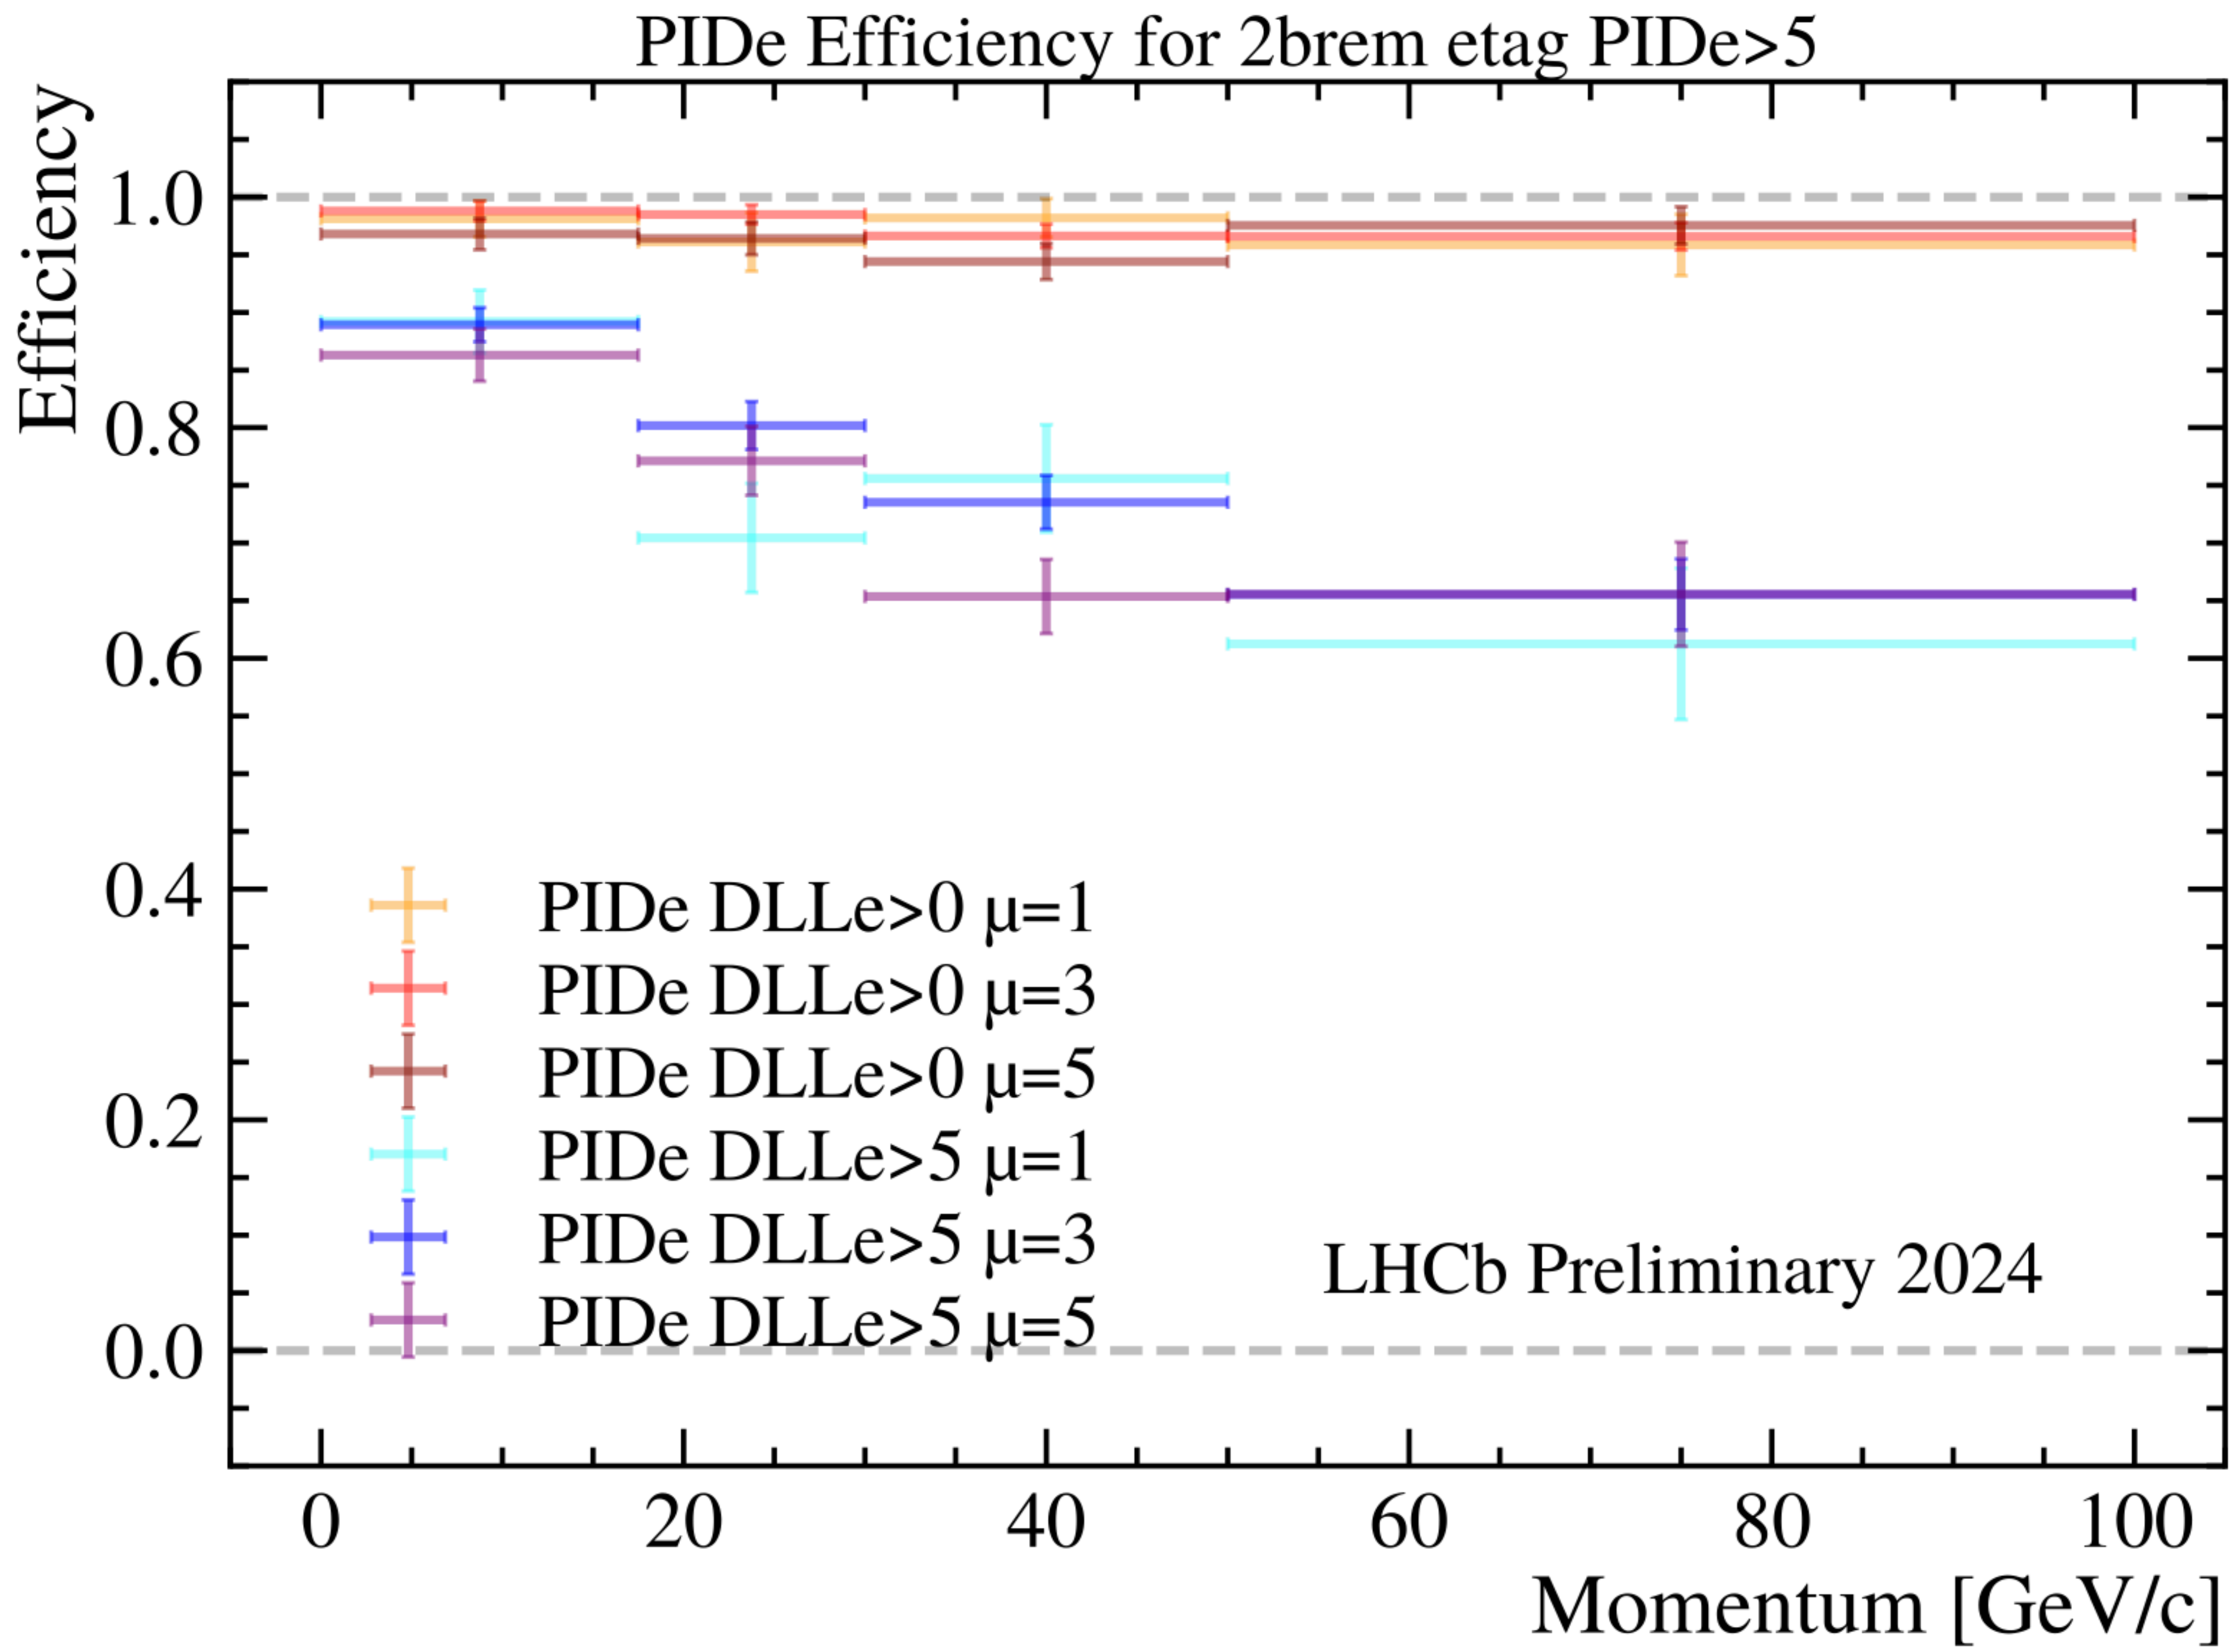

Hadron PID performance Run 2 & Run 3