Online reconstruction and Trigger

Marianna Fontana, on behalf of the LHC collaborations

INFN bologna

LHCP 2023 - 23/05/2023
Belgrade - Serbia

Introduction

- The reconstruction of raw detector data and its processing in real time represents a **major challenge in HEP**
- **Demands for higher throughputs in upcoming years**

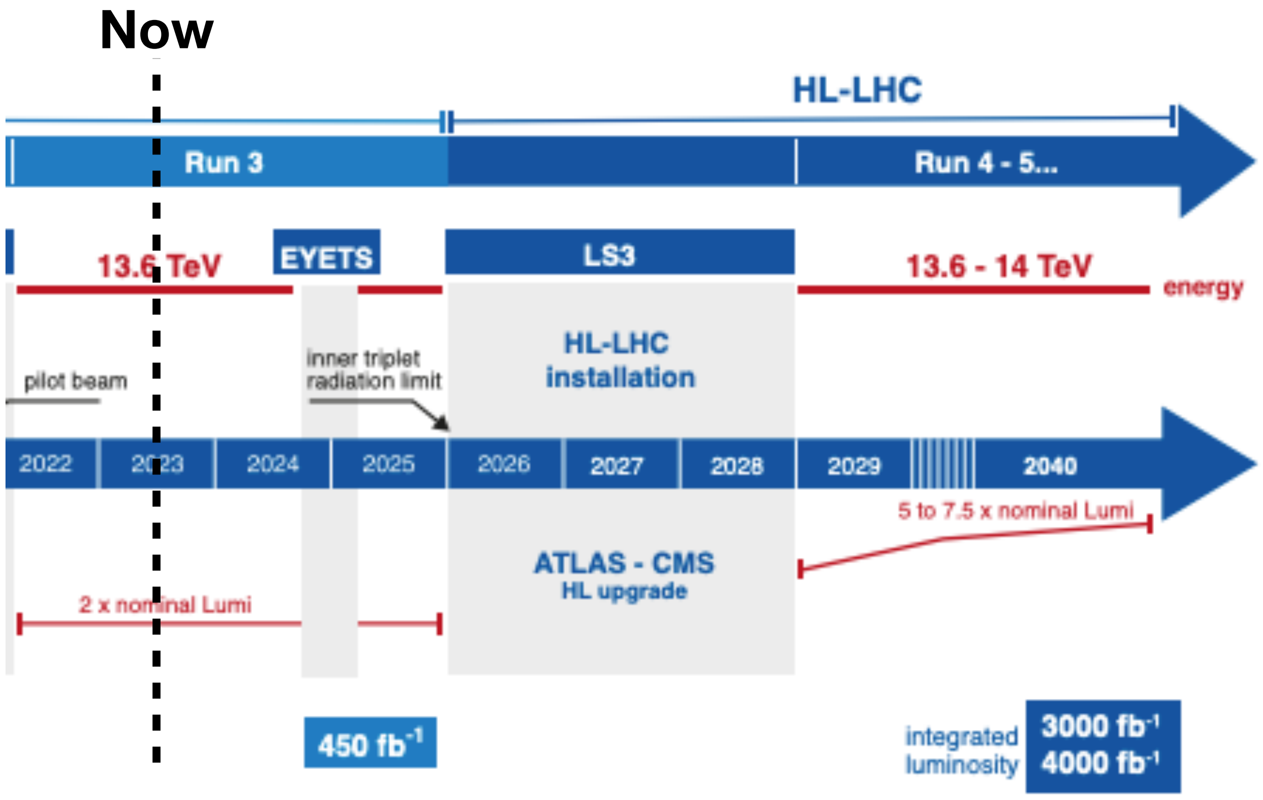

https://hilumilhc.web.cern.ch

- Two demands for Trigger:
    - Decrease throughput to backend DAQ
    - Keep trigger efficiency high
- Two trends:
    - Triggerless/continuous readout (ALICE-LHCb)
    - Higher-level reconstruction in hardware trigger (ATLAS-CMS)

Disclaimer

This presentation is in no way an exhaustive view of all LHC experiments trigger systems, rather a selection of some topic biased from my view

The ATLAS Run 3 trigger

Trigger system consists of Level-1 (hardware) and High Level Trigger (software):

- reduce data rate from ~ 40 MHz to O(1 kHz)

- **L1 improvements**: data from calorimeters and muon detectors processed by dedicated trigger systems with expected improved efficiency
  - L0Calo and L0Muon
- Highly selective triggering with strong channel dependence (~50% central Z —> ee, ~50% central HH->4b)
- High trigger thresholds for hadronic jets -> focus on new physics at high masses
- Online reconstructs partial or full event to offline-like quality
- CPU and readout limited at high beam intensities
- HLT software framework fully redesigned to be **multi-threaded compliant**
- **No accelerated computing for Run 3**, but lots of R&D for HL-LHC.

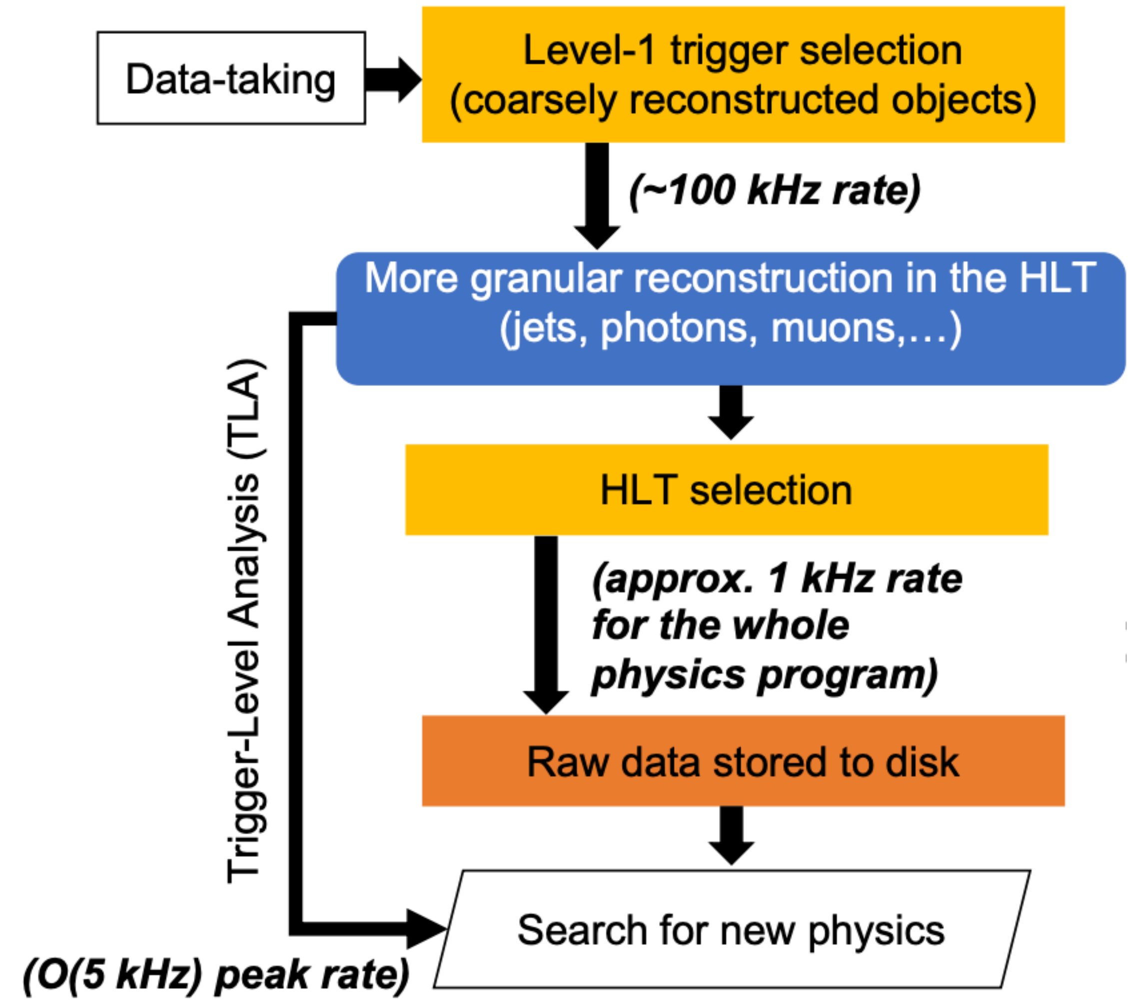

Courtesy of M. Amerl

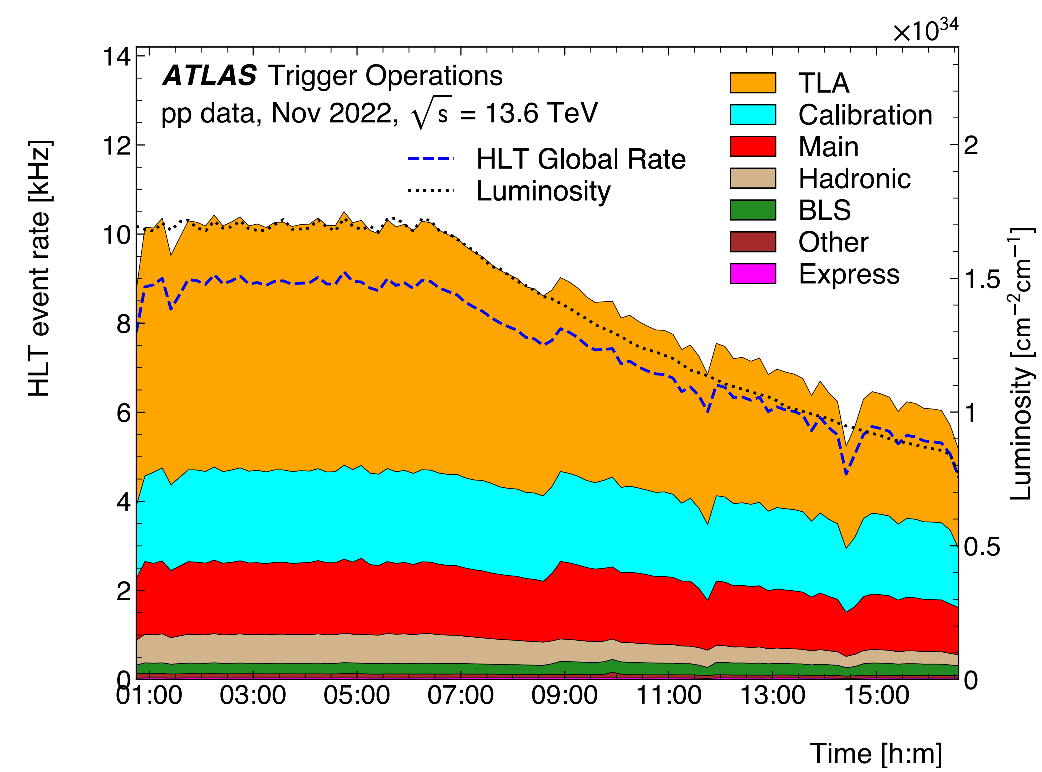

Trigger Operation Results

The ATLAS HLT improvements

ATLAS HLTTrackingPublicResults

- **Inner Detector Trigger MinBias tracking** in 900GeV collisions
  - The 4(6)GeV track-triggers have an unbiased tracking spectra below 4(6)GeV & very high efficiency
- Improvements to the offline large radius tracking (LRT) allow it to be used online in the trigger to target signatures with long lived particles
  - LRT outperforms standard tracking decision
- Improved jet and b-jet performance
  - b-jet efficiency of the new DIPS and DL1d algorithms outperforms the benchmark DL1r algorithm
  - Improved GN1 tagger [see poster by M. Chen]

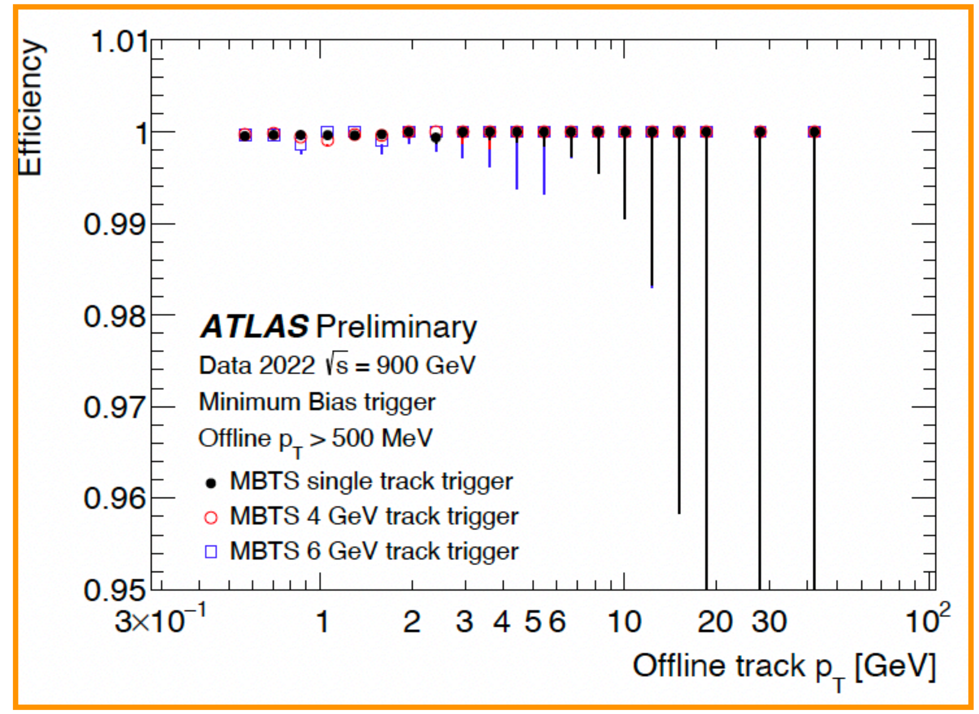

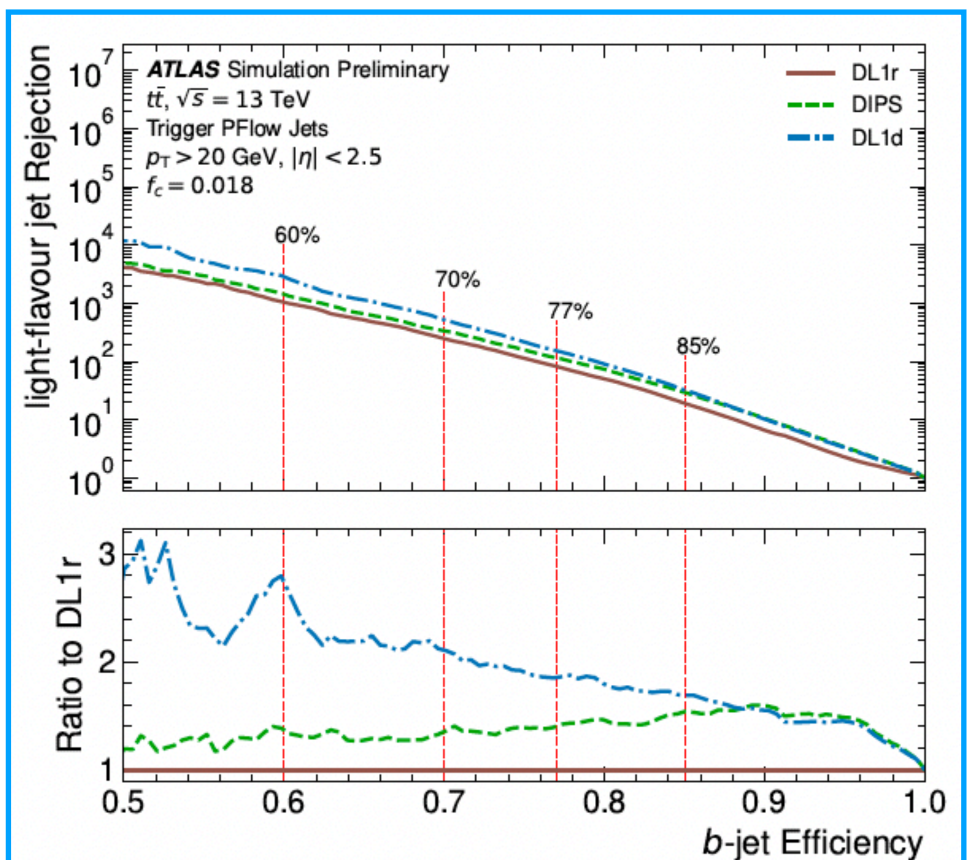

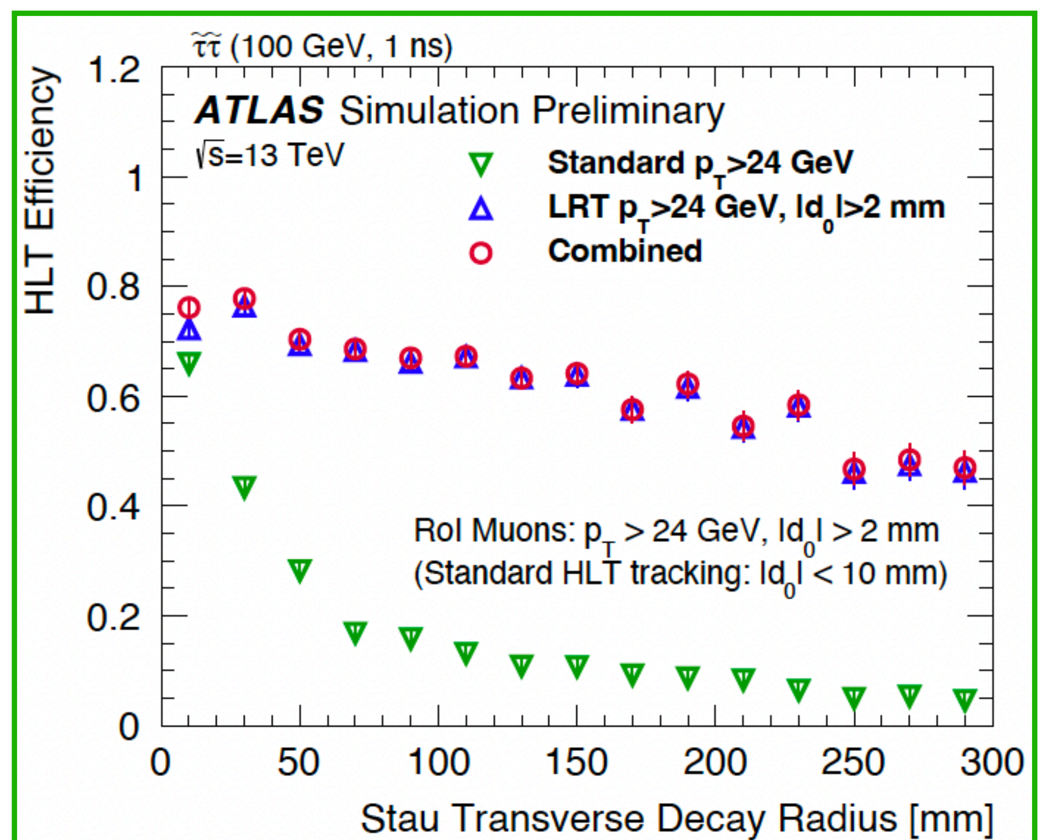

The CMS Run 3 trigger

All performance available at [CMS HLTTrackingResults]

Two level triggering system to reduce the rate to a manageable level

- L1 (< 100 kHz): large effort to develop new seed features, algorithms, and triggers to target LLP signatures and rare signals
- HLT (~ 1 kHz)
  - **New tracking** based on the optimized pixel track (= **Patatrack**) reconstruction
    - it allows to reduce the HLT tracking to a single-iteration approach
  - **Improved tagging** capabilities of low-level calorimeter objects + new L2 τ reconstruction using pixel tracks
  - New muon reconstruction and ML-based inside-out and outside-in seeding
  - **Improved performance of jets** reconstructed with Patatrack pixel tracks as inputs to the Particle Flow (PF) algorithm

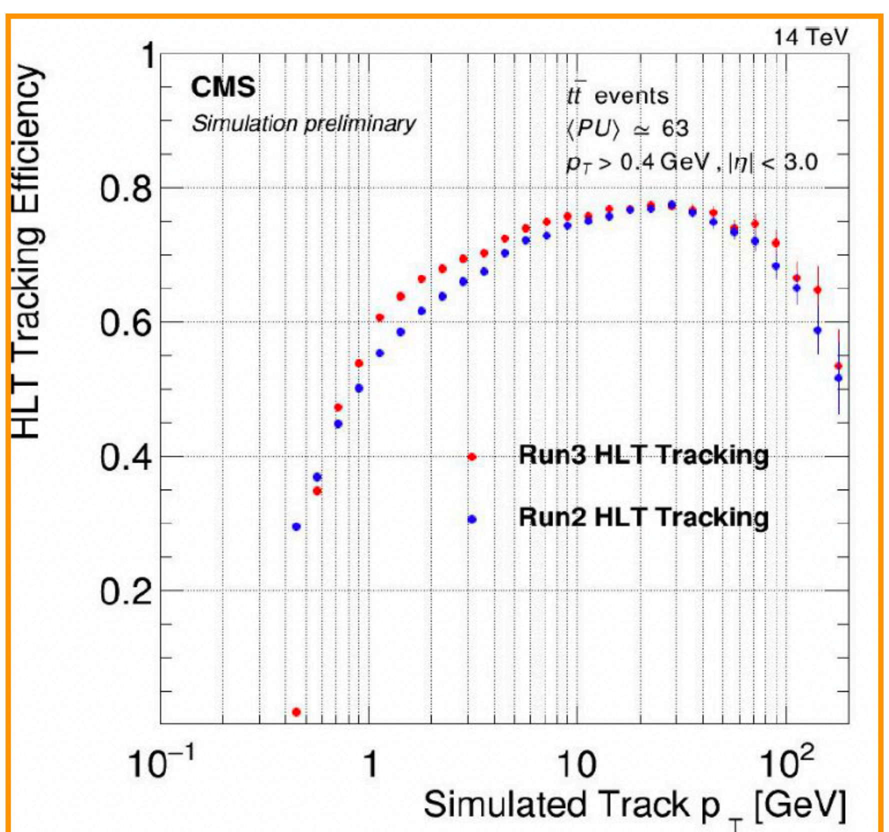

CMS-DP-2022-014

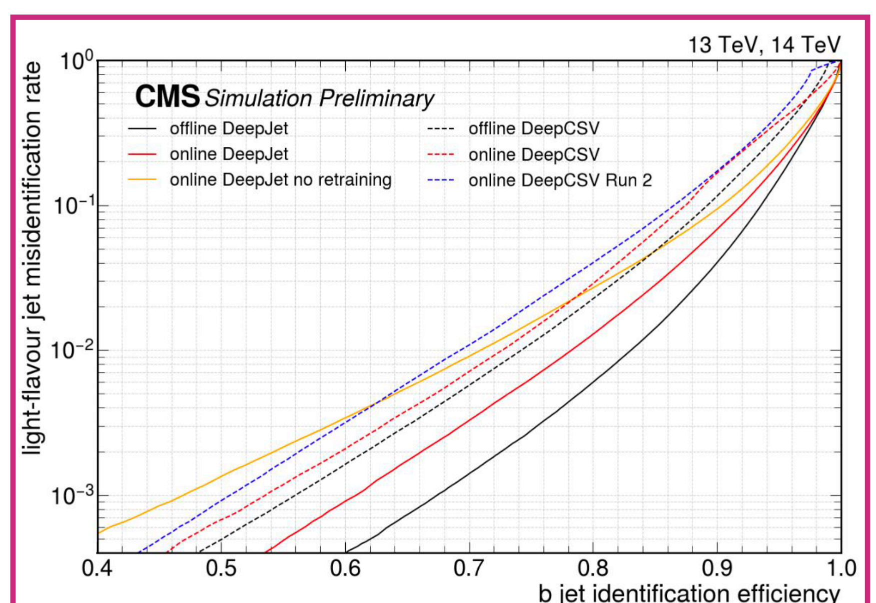

CMS-DP-2022-030

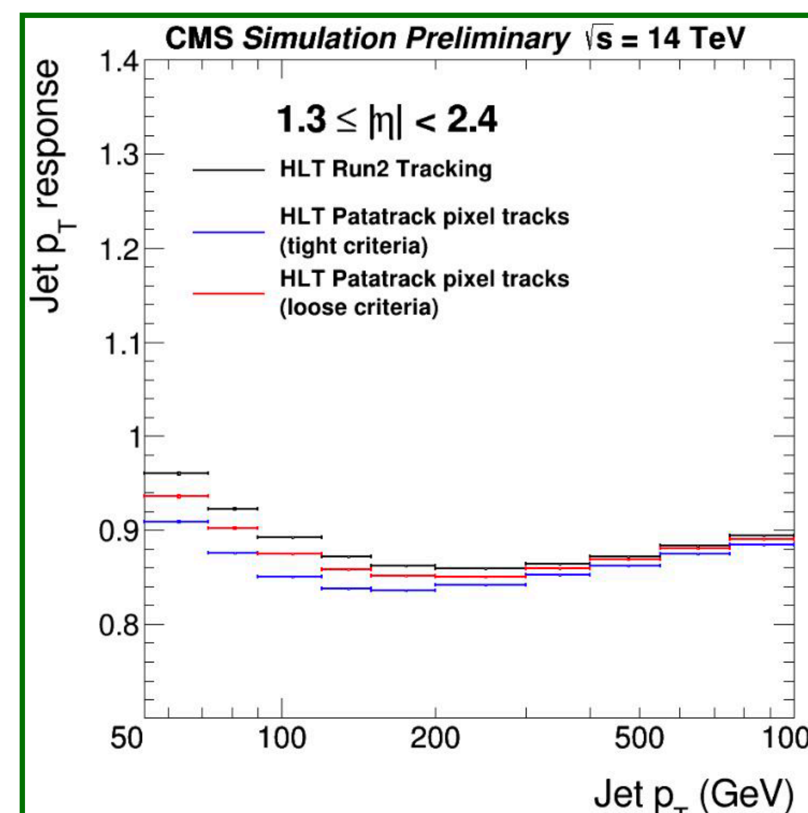

CMS-DP-2021-005

CMS HLT: a new CPU + GPU farm

CMS DP2023-024

Software based trigger with full event information available running on CPU + GPU based farm

- Run 3 HLT farm composed of **200 nodes (25600 CPU cores and 400 GPUs)**
  - Each node equipped with two AMD Milan 64-core CPUs and two NVIDIA Tesla T4 GPUs
- Increasing usage of GPUs at Run 3
  - Offloading **30% of the HLT reconstruction to GPU**
- GPU reconstruction implemented and fully commissioned
  - calorimeter and pixel local reconstruction + pixel tracking + vertex reconstruction
  - The **execution time per event** of was **reduced by ~40%**
- HLT throughput requirement ~500Hz:
  - Throughput increases by a factor of ~1.80
  - Power Consumption (per throughput) reduced by ~30%

CMS *Preliminary* 13.6 TeV

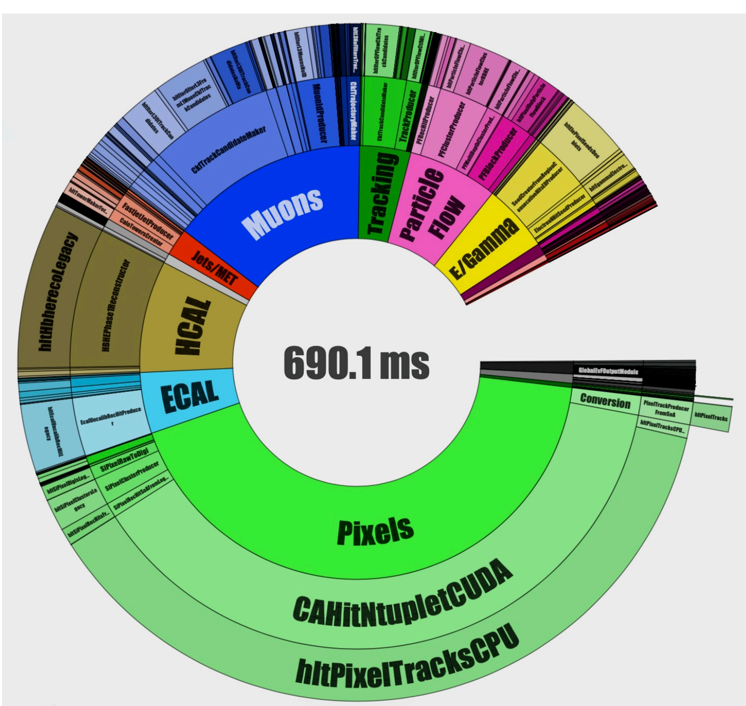

Link

Execution time CPU only

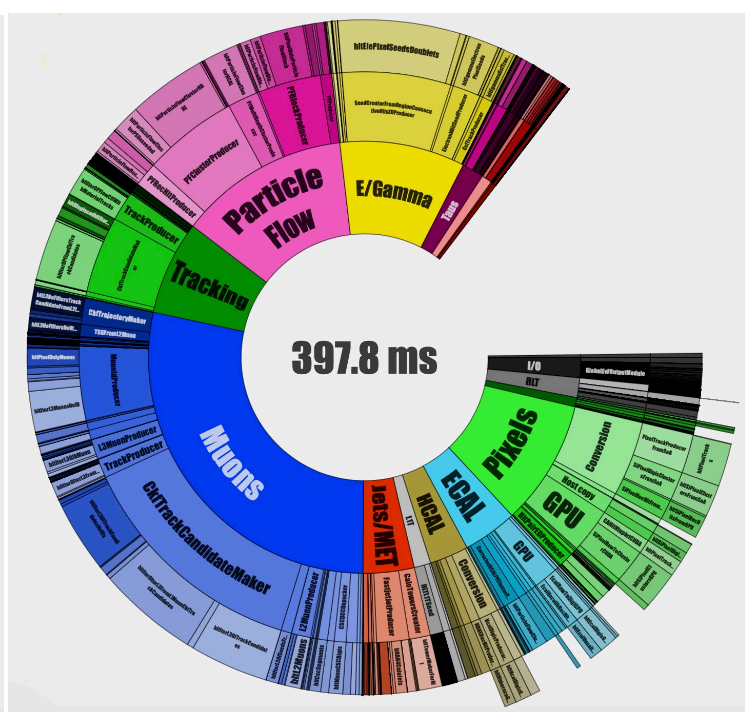

Link

Execution time CPU+GPU

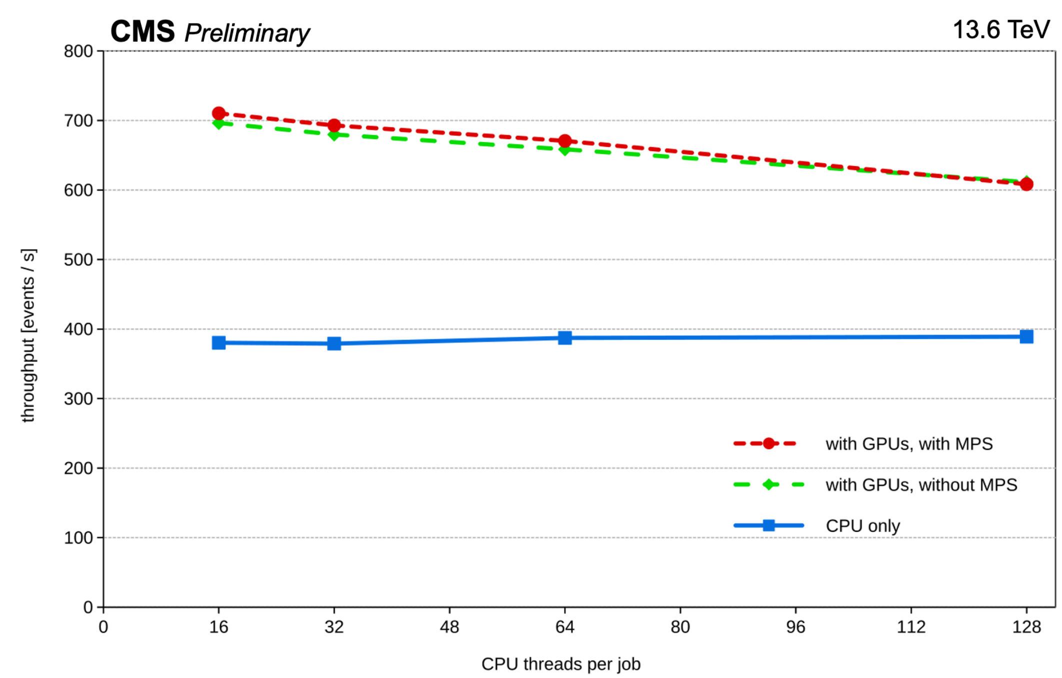

Link

HLT throughput

The ALICE Run 3

Challenges for Run 3:

- Completely **new detector readout** and substantial detector upgrades: new ITS, MFT, FIT. New GEM for TPC readout
- ~100 x more data than Run 2
- Many important physics signals have very small signal-to-noise ratio
- Triggering (selection) techniques very inefficient if not impossible
- Needs large statistics
    - **Read the data resulting from all interactions** (LHC will deliver min bias Pb-Pb collisions at 50 kHz)

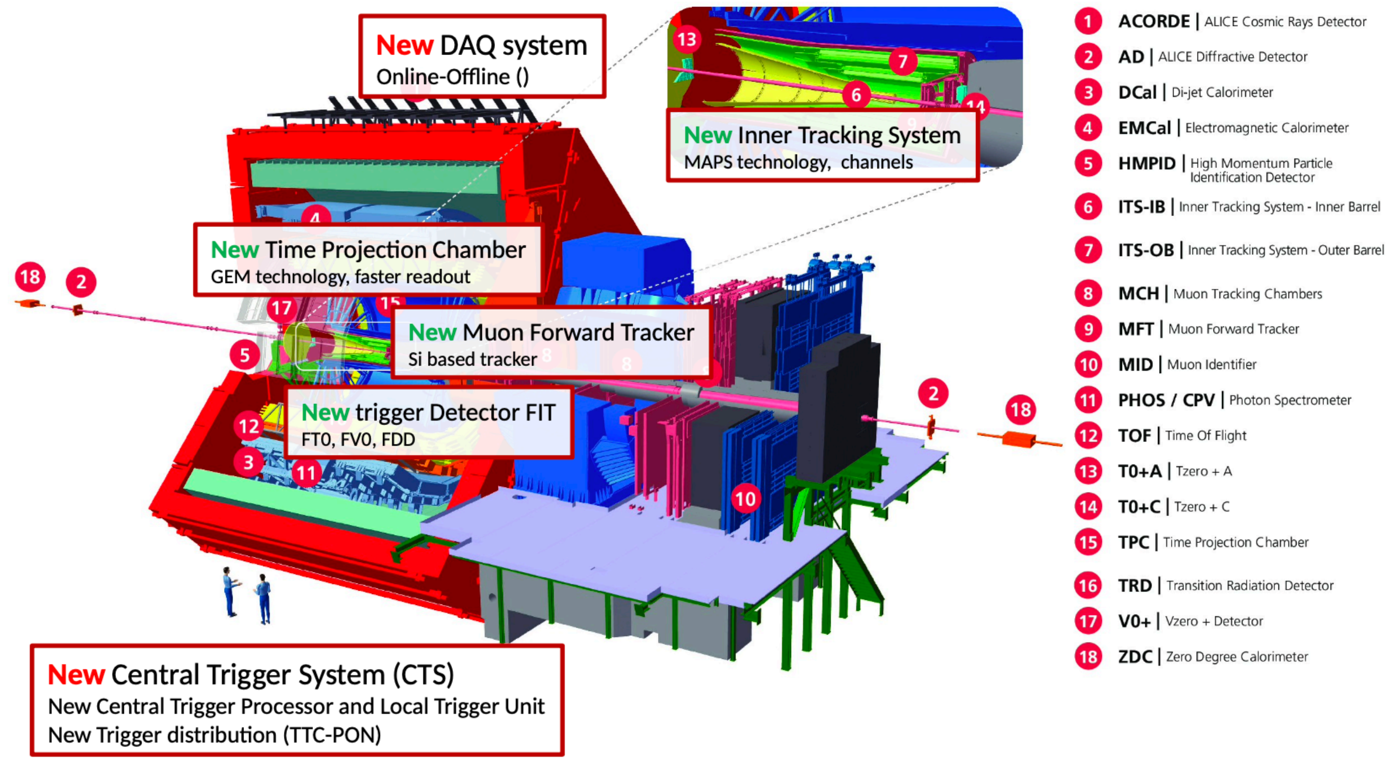

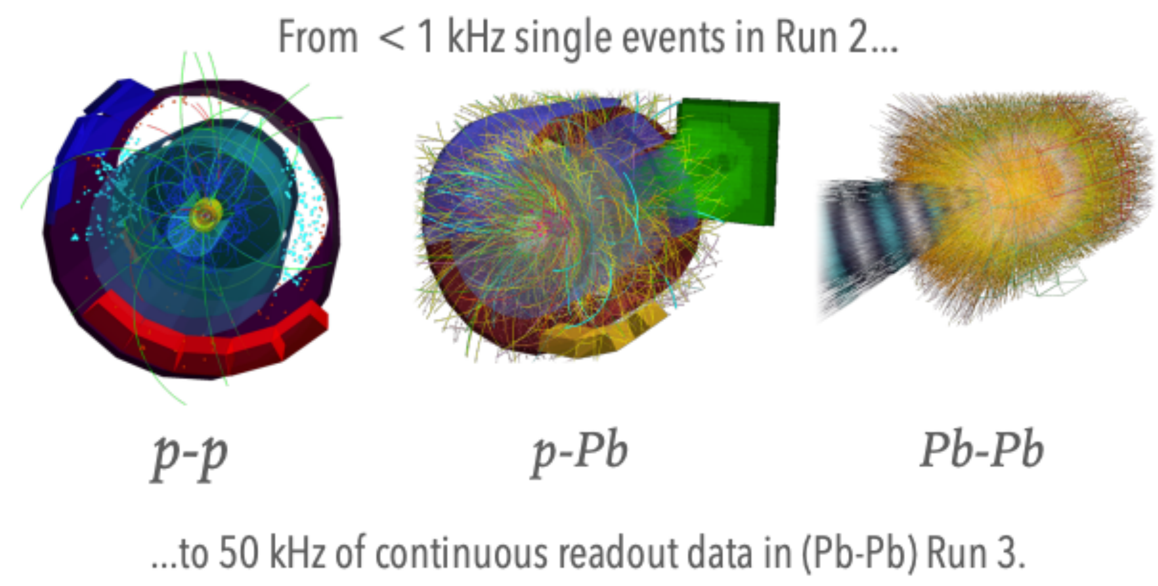

Courtesy of G. Eulisse

Reconstruction using GPUs in ALICE

- Trigger-less acquisition: **continuous detector read-out**
    - $L_{int} > 10\ nb^{-1}$ of PbPb data at 50kHz: 50x more than Run 2
- Reconstruction is two-stepped
    - **Synchronous phase** (beam circulating): for calibration and data compression
    - **Asynchronous phase** (no beam): full processing of data staged on a temporary buffer. One common online-offline system (O^2)

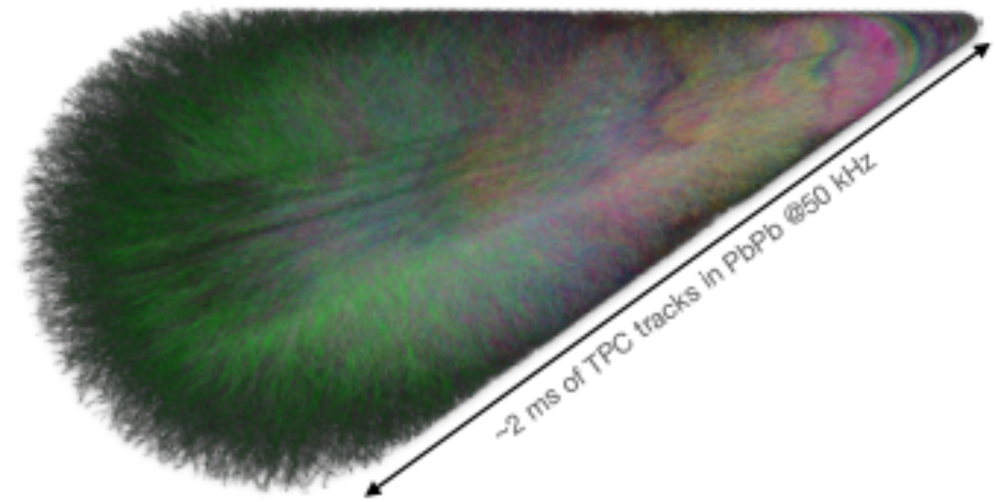

- ALICE uses **GPUs** to accelerate the process
    - TPC track reconstruction is the most time consuming during synchronous reconstruction and is therefore performed on GPUs (the most cost effective solution)
    - During the asynchronous reconstruction, the fraction of available GPU increases
    - Use those resources efficiently by offloading also ITS reconstruction there

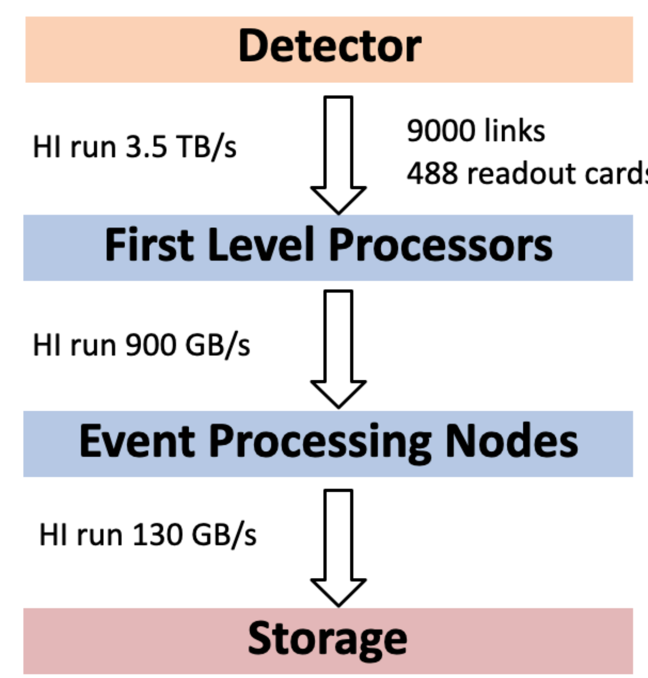

Courtesy of V. Barroso

The LHCb experiment

- Major upgrade of all subtectors:
  $\mathscr{L} = 2 \times 10^{33}\ cm^{-2}\ s^{-1}$
- 100% of the readout electronics replaced
- New data acquisition system and data center

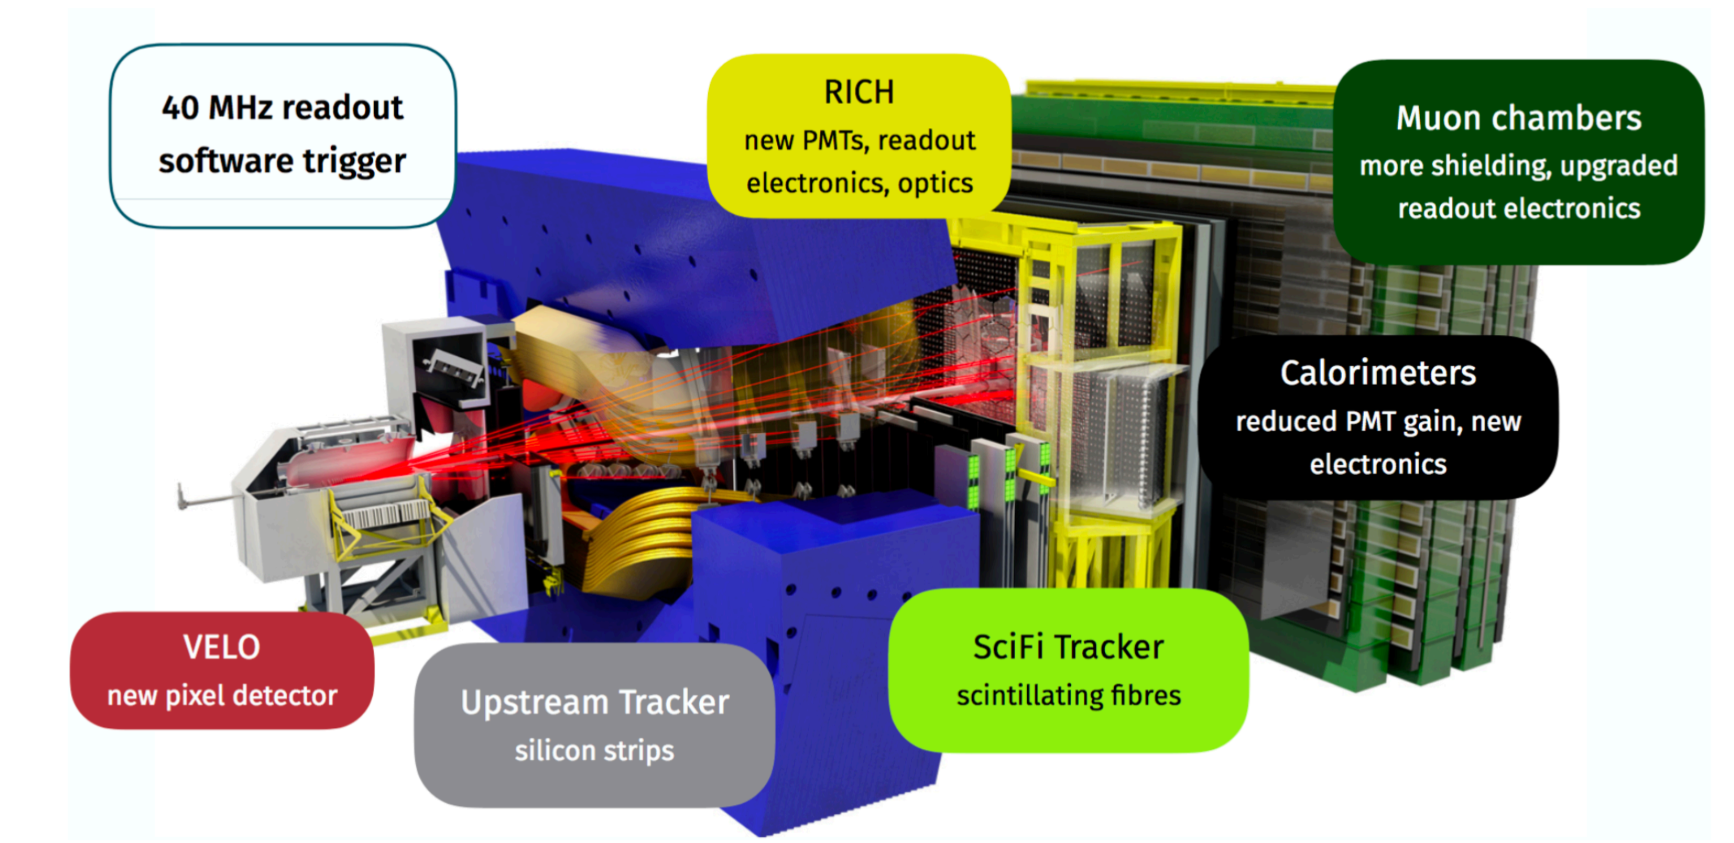

CERN-LHCC-2012-007

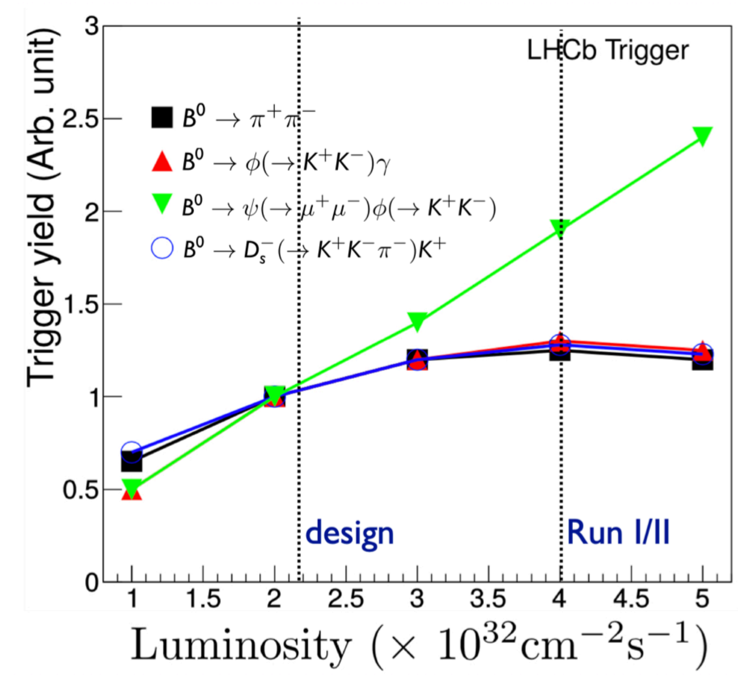

- Cannot effectively trigger on heavy flavour using hardware signatures
- Trigger for many hadronic channels saturated already at Run 1-2 luminosity
- **Solution**: fully software trigger

The LHCb data flow

LHCB-FIGURE-2020-016

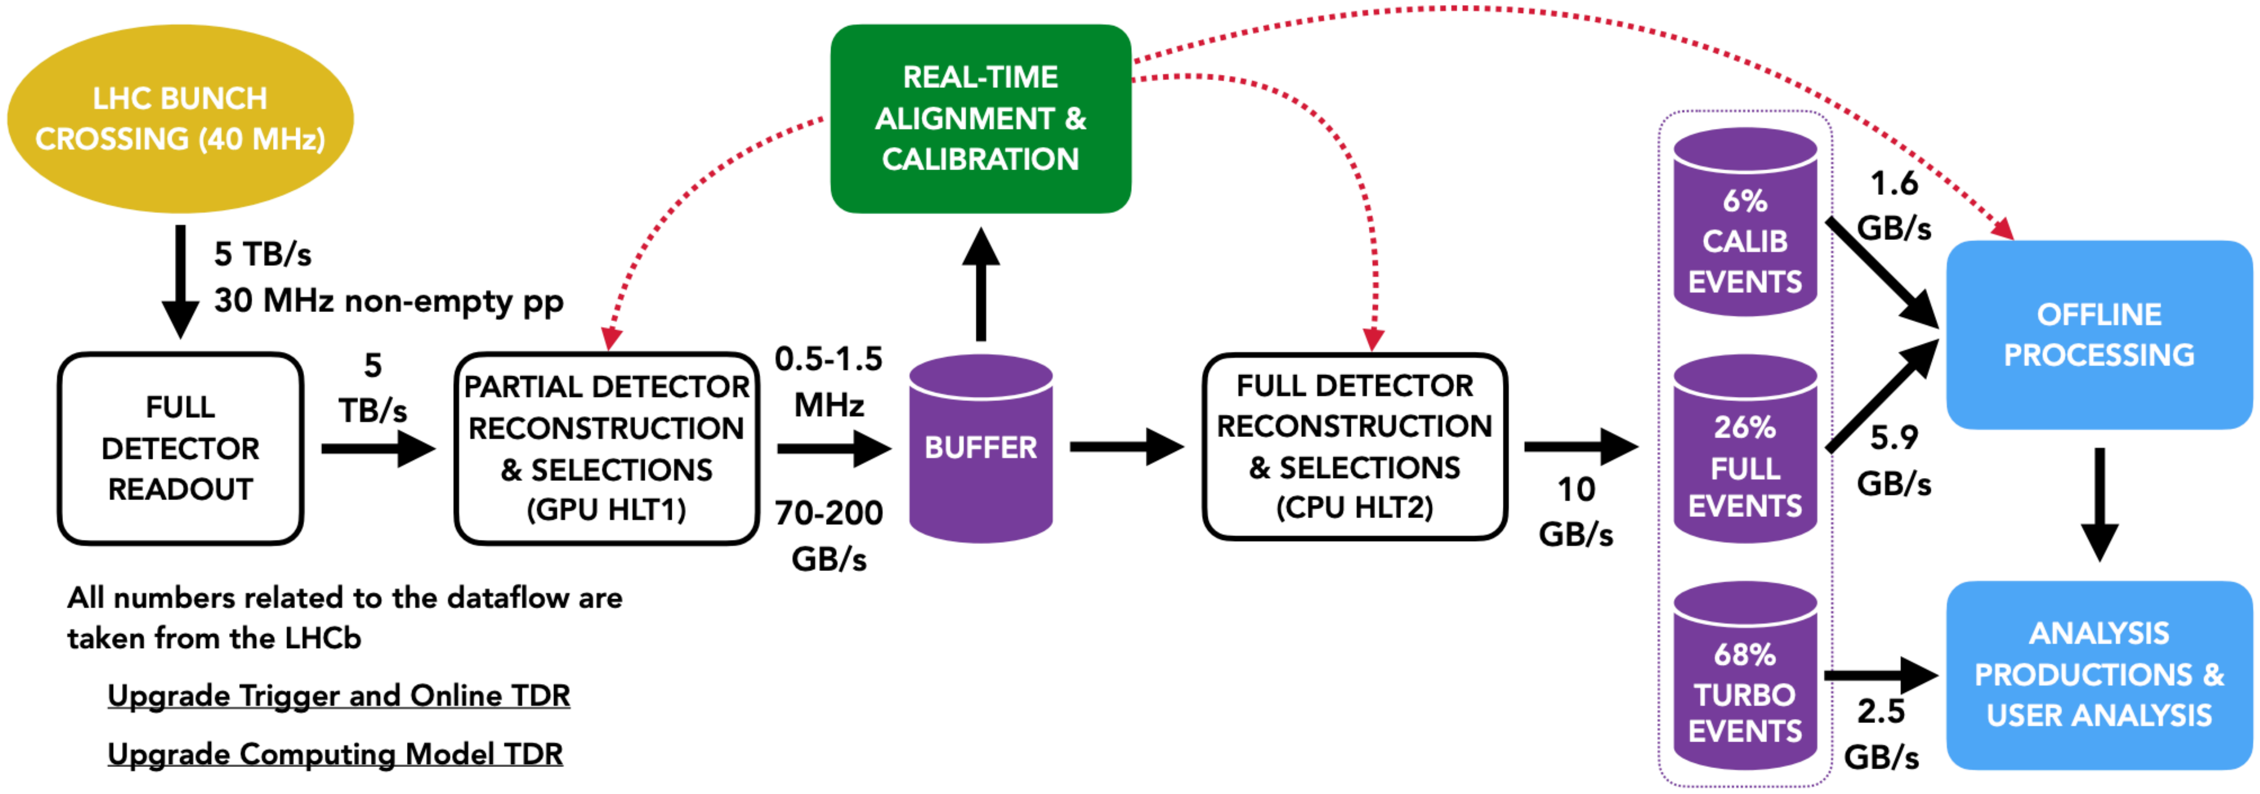

- Detector data @30 MHz received by O(500) FPGAs
- 2-stage software trigger, HLT1 & HLT2
- Real-time alignment & calibration
- After HLT2, 10 GB/s of data for offline processing

LHCb HLT1 trigger

The goal of HLT1

- Be able to intake the entirety of the LHCb raw data (5 TB/s) at 30 MHz
- Perform partial event reconstruction & coarse selection of broad LHCb physics cases
- Reduce the input rate by a factor of 30 (~1 MHz)
- 30 MHz benchmark can be achieved with **O(200) GPUs** (max number EB can host is 500)
- ~ 350 GPUs now installed -> Additional functionalities are being explored
- Throughput scales well with theoretical TFLOPs of GPU card

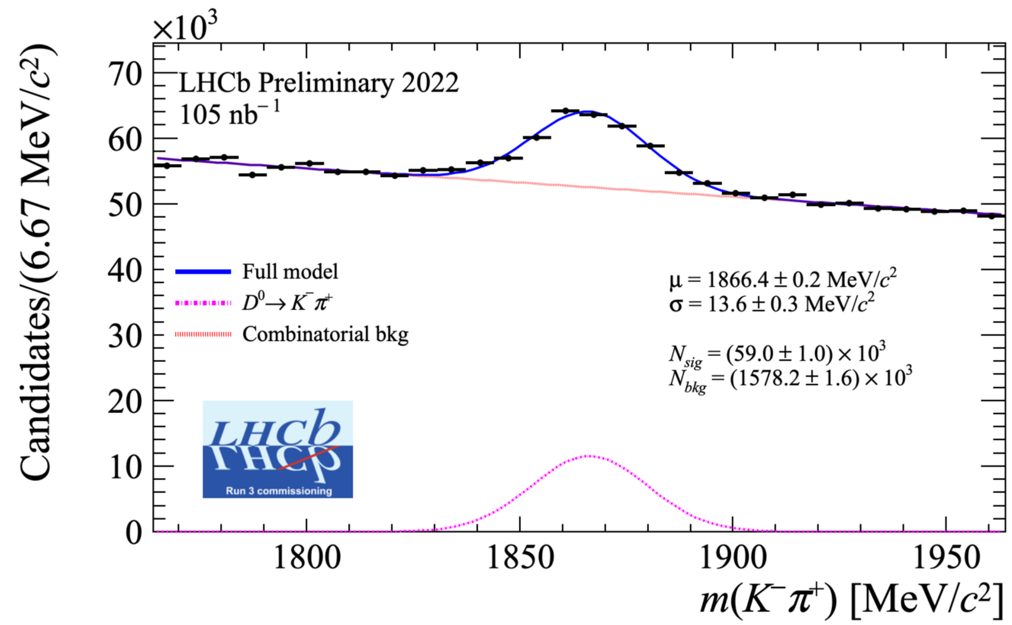

LHCb-FIGURE-2023-009

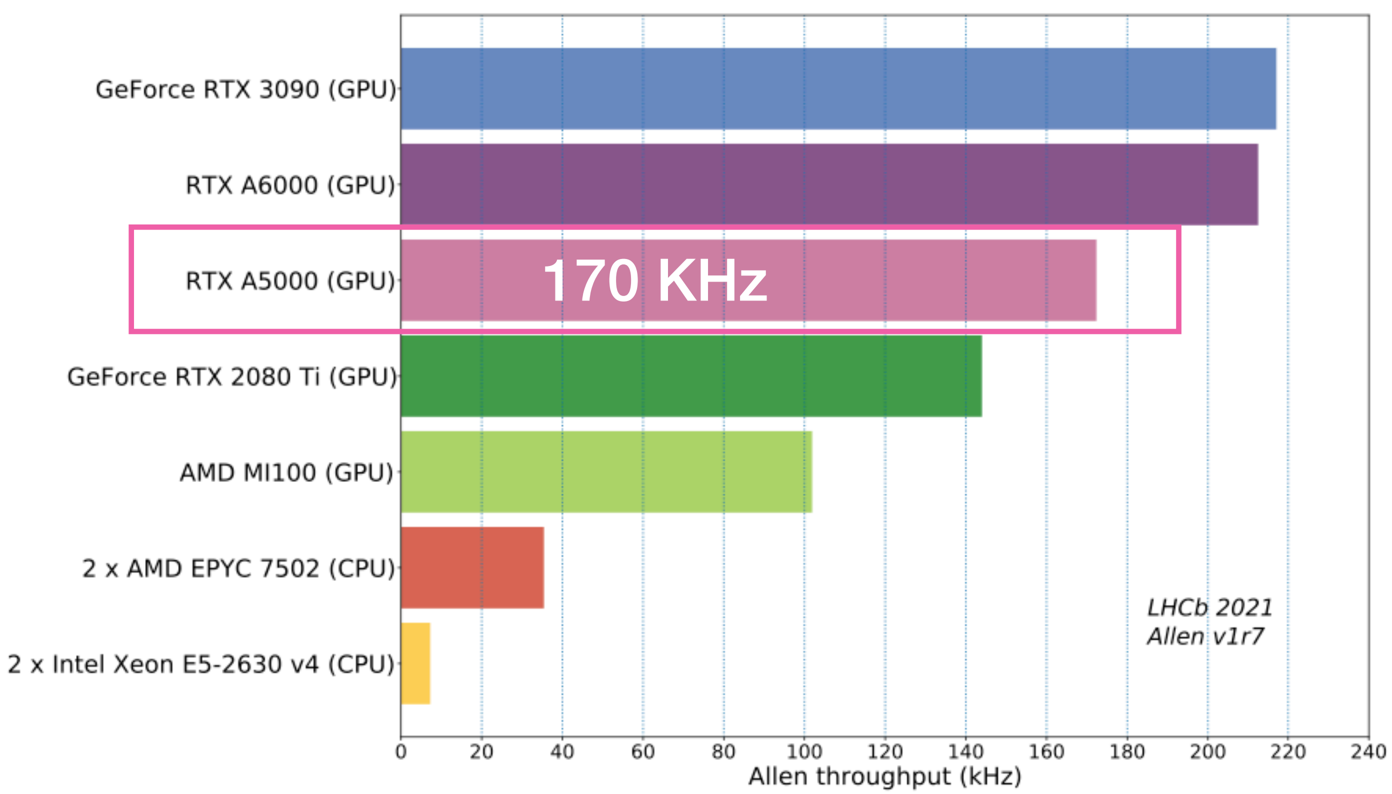

LHCb-FIGURE-2020-014

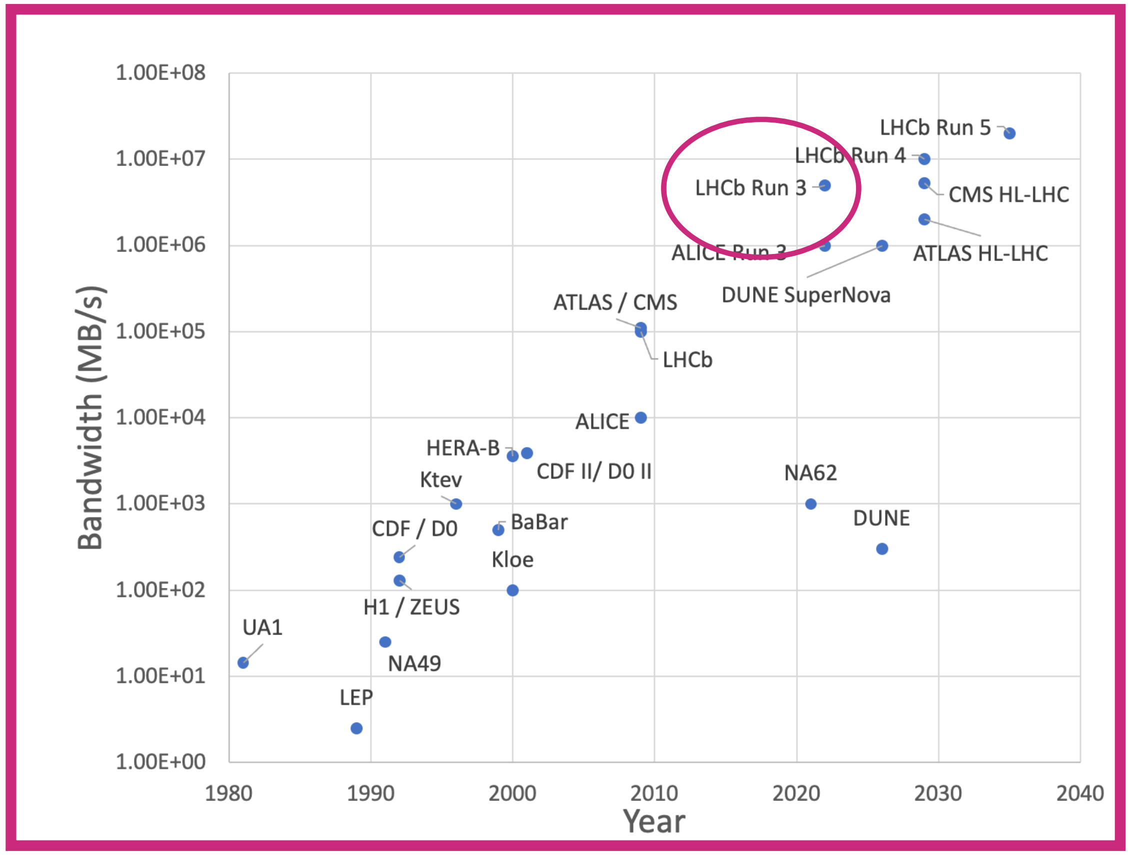

A. Cerri - University of Sussex

First complete high-throughput GPU trigger for a HEP experiment!

<u>Allen</u>: a GPU HLT1 trigger platform [Comput Softw Big Sci 4, 7 (2020)]

LHCb real-time alignment&calibration

- Store data selected in HLT1 in intermediate buffer for real-time alignment and calibration
- Fully aligned and calibrated data needed before running HLT2
- Online alignment and calibration pioneered in Run 2, crucial in Run 3
- Buffer capacity of O(10 PB) situated between HLT1 and HLT2
- Two types of processes
  - Alignment: VELO, RICH mirrors, UT, SciFi, Muon
  - Calibration: RICH, ECAL, HCAL

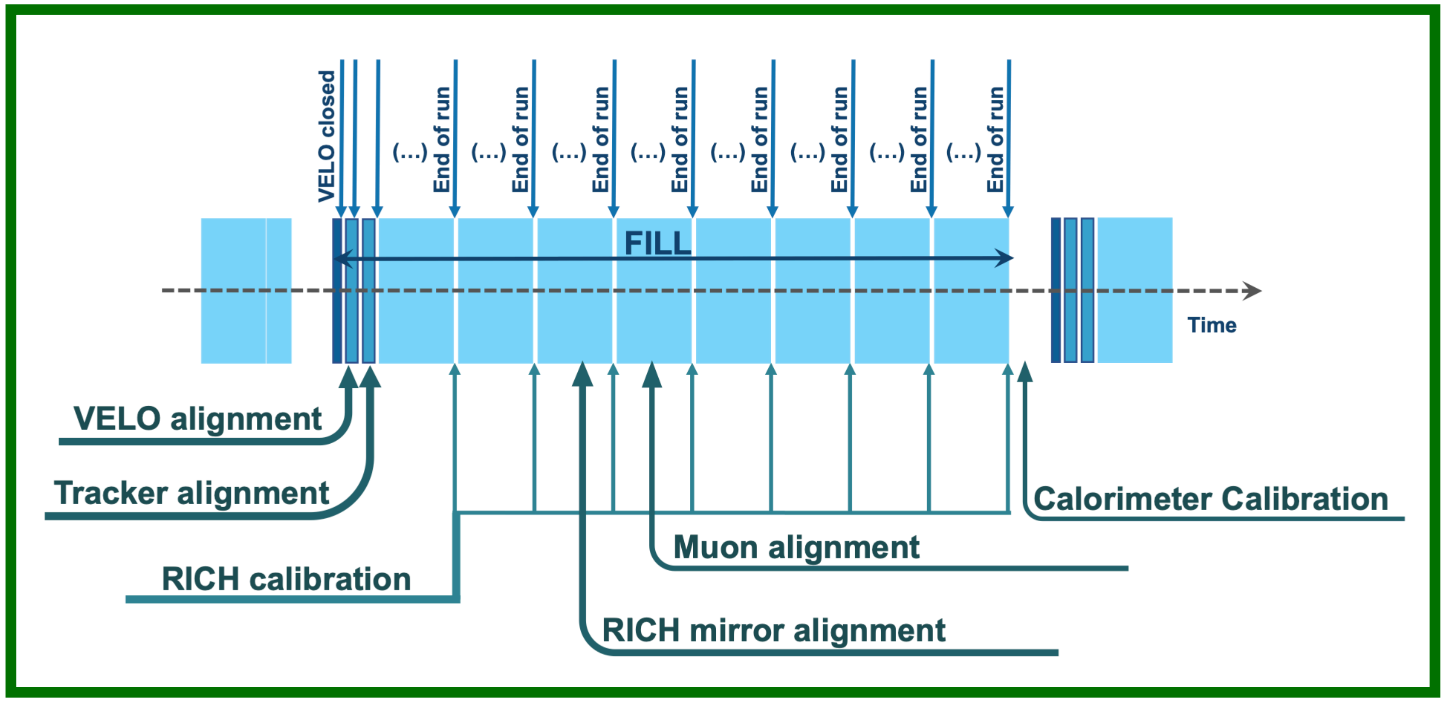

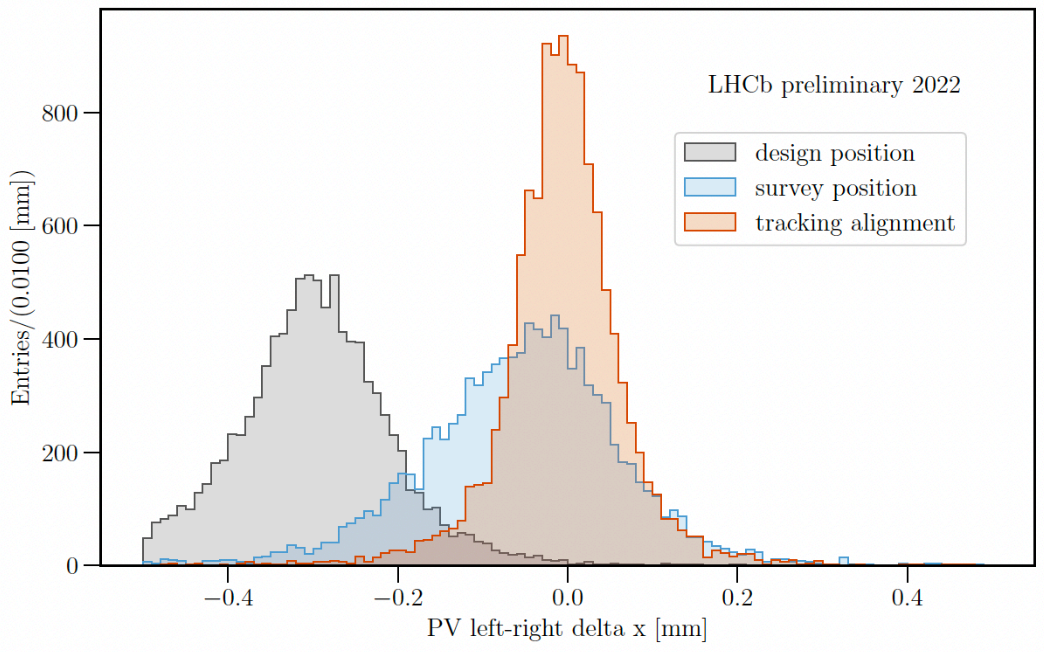

LHCb-FIGURE-2022-016

LHCb HLT2 trigger

- Given the hard limit on bandwidth and expected signal rate, event size is the only free parameter
- Instead of saving full event, only information needed for a physics analysis can be stored (**Turbo paradigm**)
- Successful hybrid strategy for Run 2 [J. Phys.: Conf. Ser. 664 082004]
- Significant reduction of data size ⇒ more events at same bandwidth
- Preliminary signals from 2022 are very encouraging!

Persistence method	Average event size [kB]
Turbo	O(10)
Turbo++/SP	O(10-100)
Raw event	O(100)

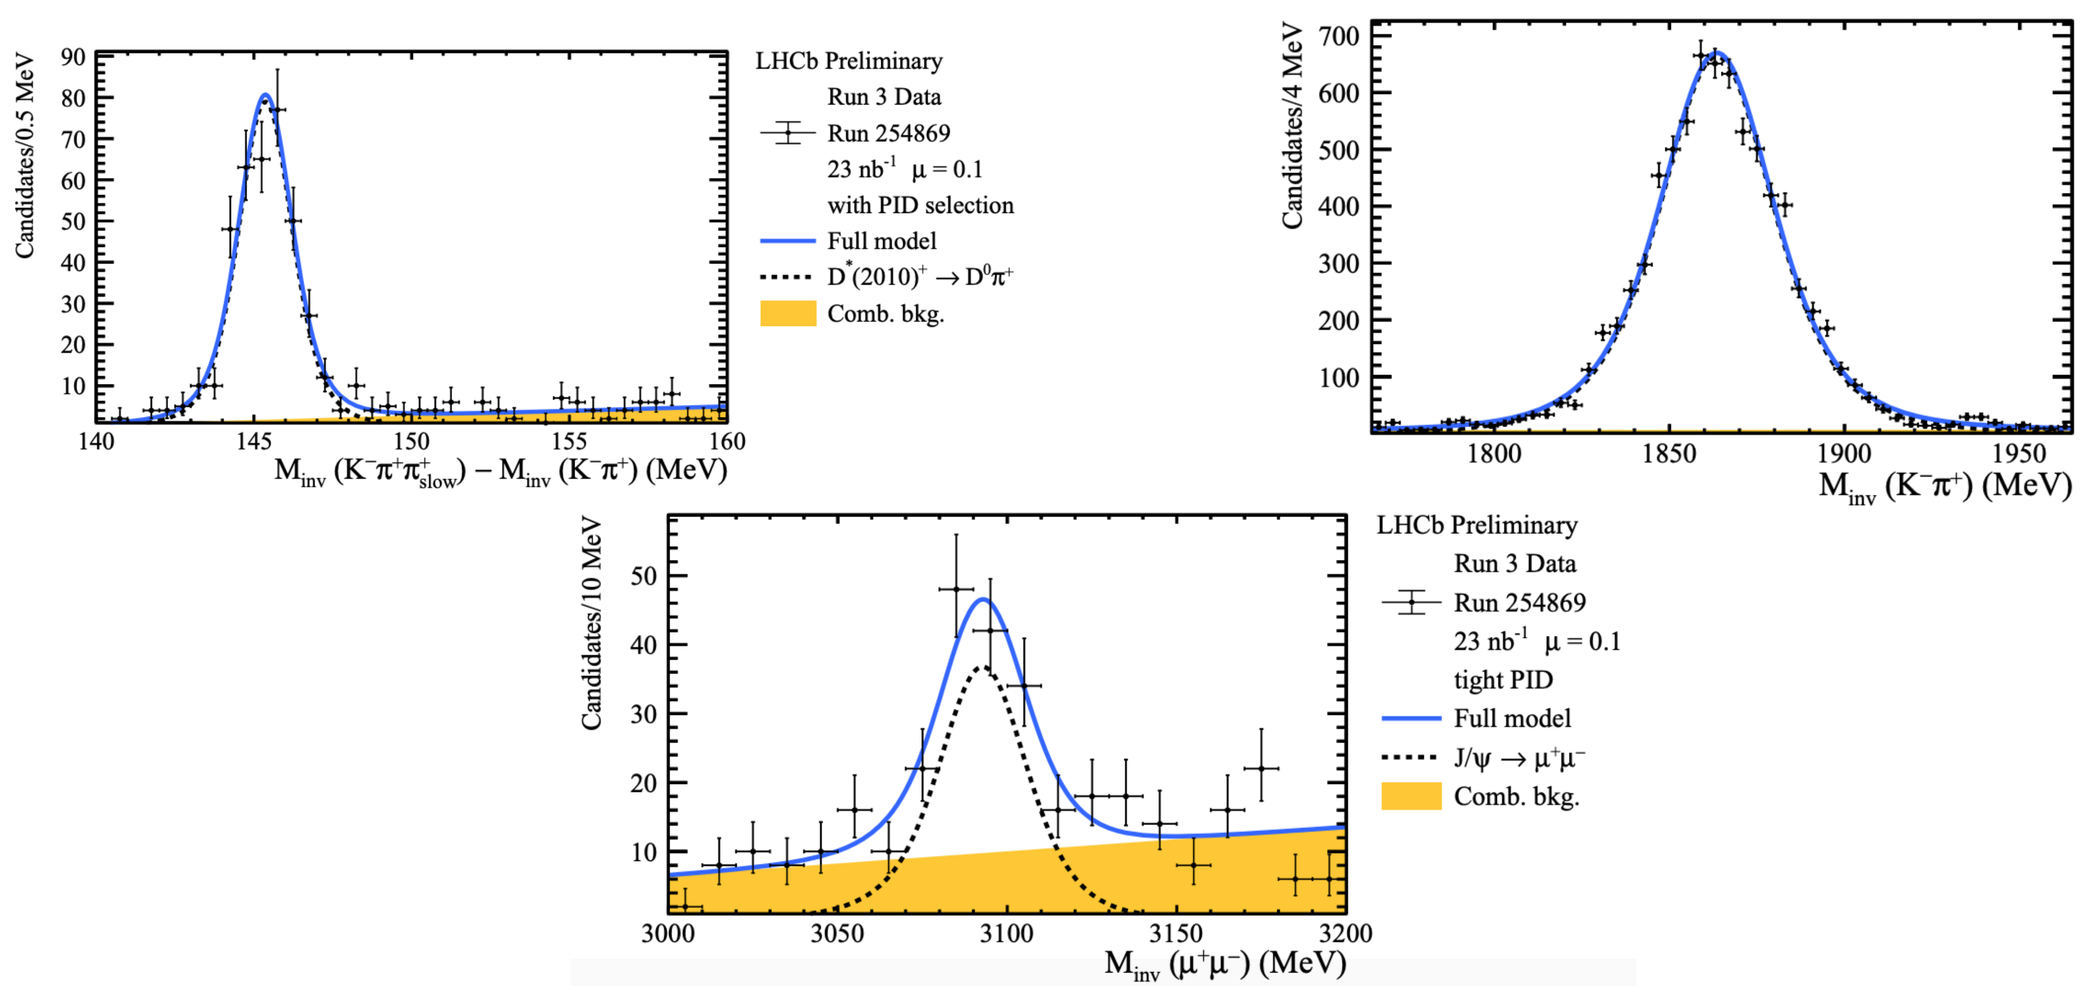

LHCb-FIGURE-2023-002

Towards the future

Trigger & data acquisition challenge for HL-LHC

- Luminosity: ATLAS/CMS $(2 \rightarrow 7.5) \times 10^{34}\ cm^{-2}\ s^{-1}$

  LHCb $(2 \times 10^{33} \rightarrow 1.5 \times 10^{34})\ cm^{-2}\ s^{-1}$

- Pileup: ATLAS/CMS $60 \rightarrow 200$

  LHCb $5 \rightarrow 40$

CMS phase II upgrade of Level 1 trigger

- The Particle Flow used for offline will be brought to L1 trigger
- Include tracking trigger in L1 trigger system
- Displaced vertex trigger for exotic events

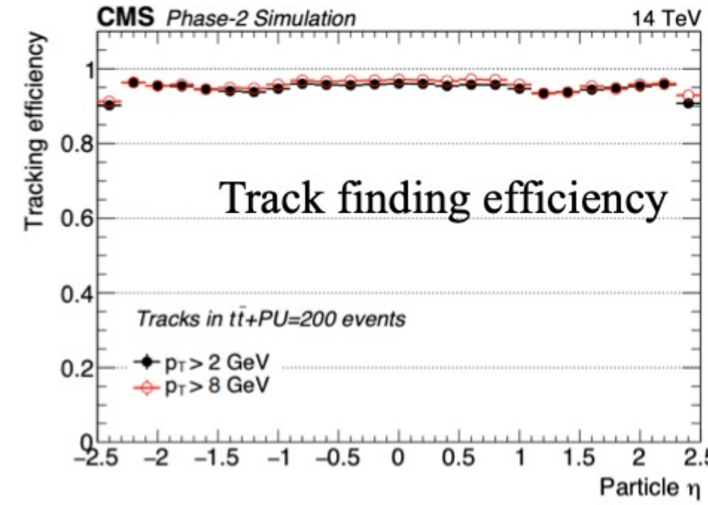

CMS-TDR-021

ATLAS L0 trigger for HL-LHC

- HLT-like object-level and event-level reconstruction and analysis at 40 MHz
- Collect all trigger data from a single event onto one FPGA

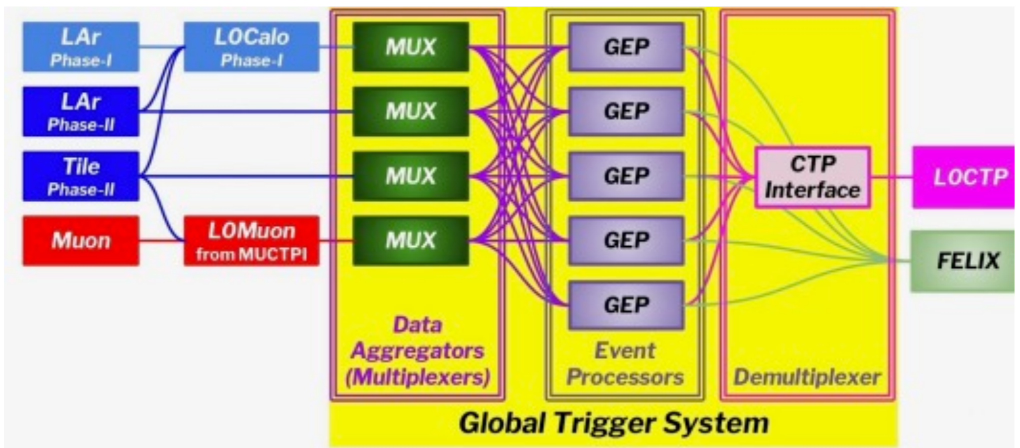

Courtesy of A. Negri

LHCb planning Upgrade II for LS4

- FTDR approved in March 22 [LHCB-TDR-023]
- Exciting challenges in trigger and DAQ
- 4D reconstruction: timing added to tracking to better isolate signals
- Actively considering which processors will work best

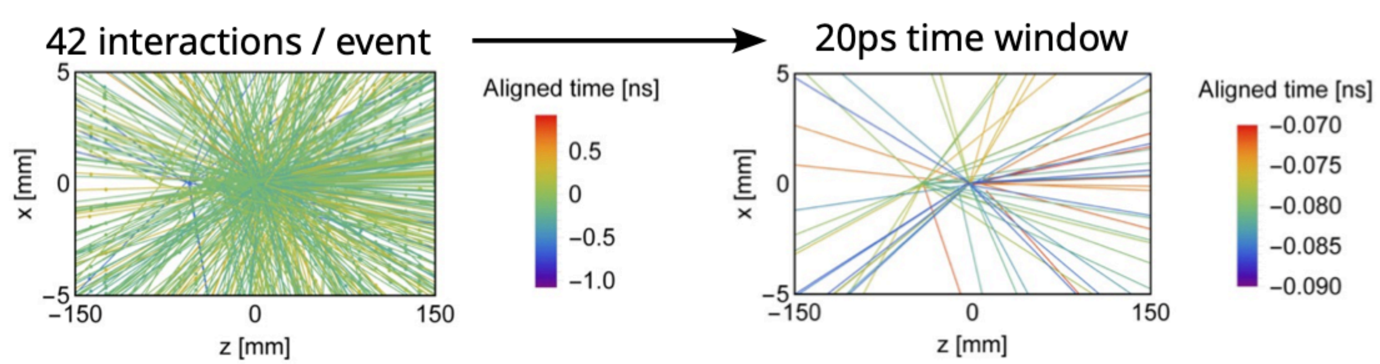

Thanks for you attention!

Backup

The ATLAS Trigger Level analysis

- Trigger bandwidth limitations:
  bandwidth = trigger rate × event size
- Trigger-Level Analysis (TLA): save only the HLT reconstructed objects (e.g. CMS Data Scouting (i), LHCb Turbo Stream (ii))
- Substantial event size reductions
  - (1.5MB (standard) vs. ~5kB (TLA))
  - Higher rate triggers
  - Small bandwidth footprint (right, (b))

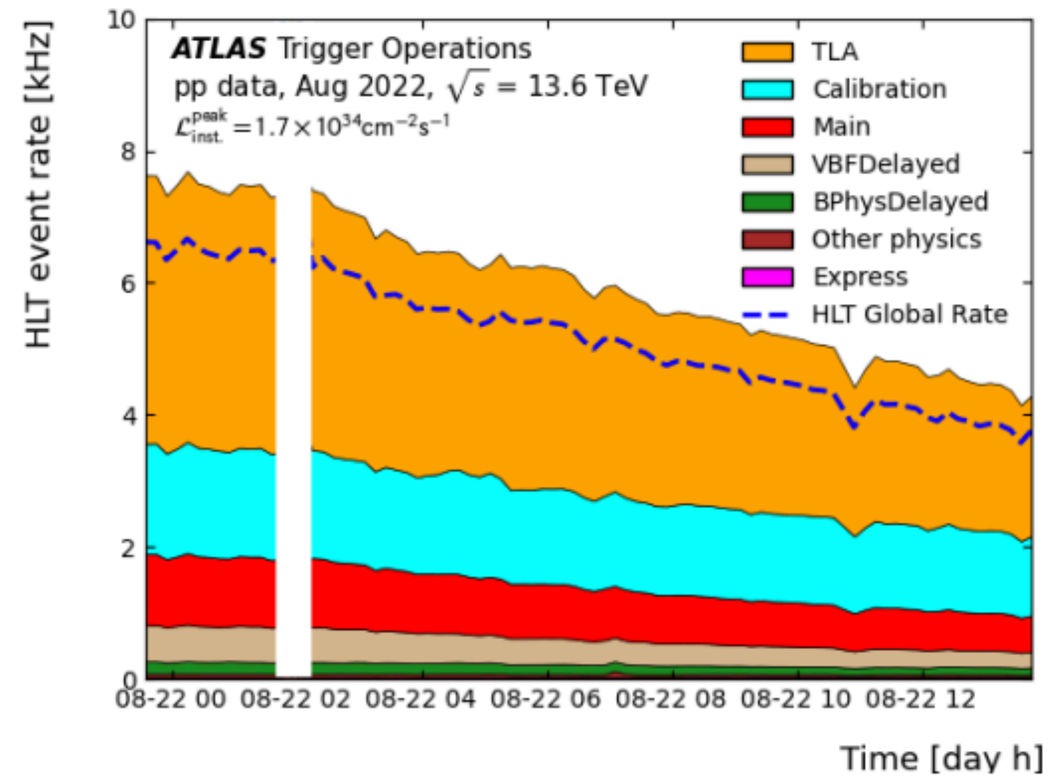

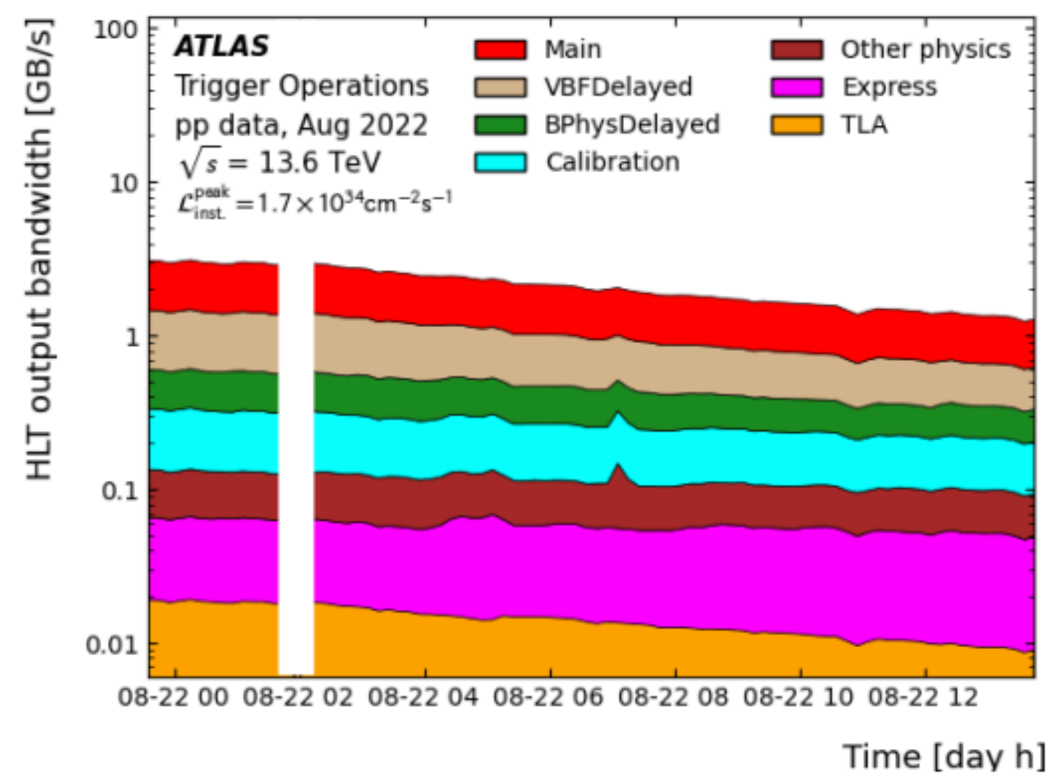

Trigger-Level Analyses require excellent performance of trigger physics objects

(i) pos.sissa.it/364/139, (ii) Sean Benson et al 2015 J. Phys.: Conf. Ser. 664 082004

The CMS Run 3 trigger

- Two level triggering system to reduce the rate to a manageable level
  - L1 (< 100 kHz): large effort to develop new seed features, algorithms, and triggers to target LLP signatures and rare signals

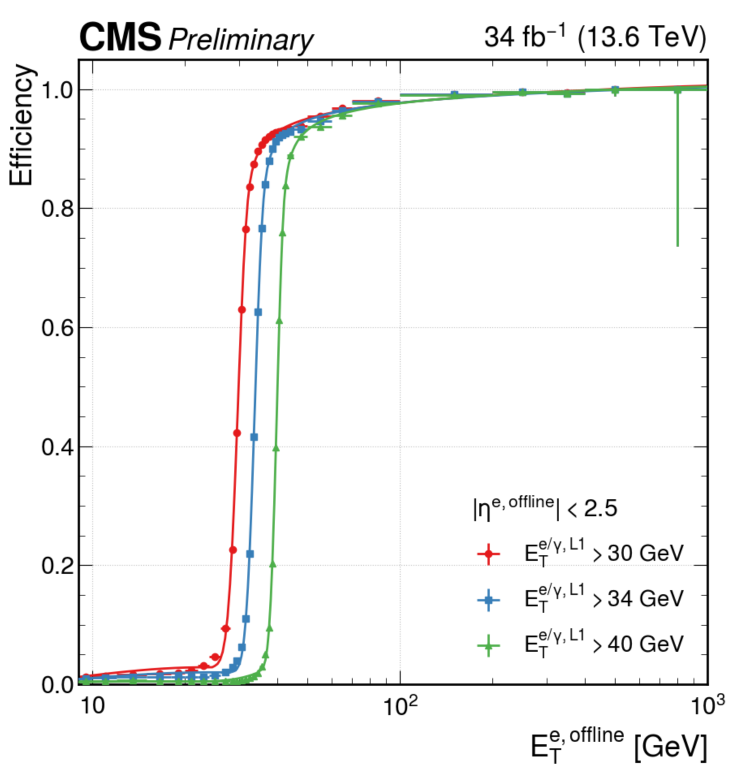

L1e/gamma

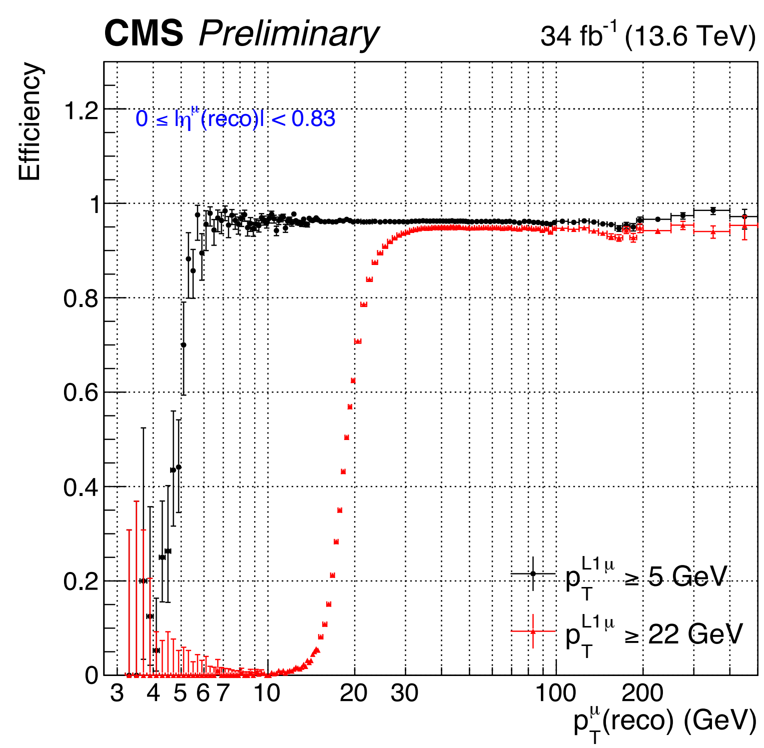

L1muons

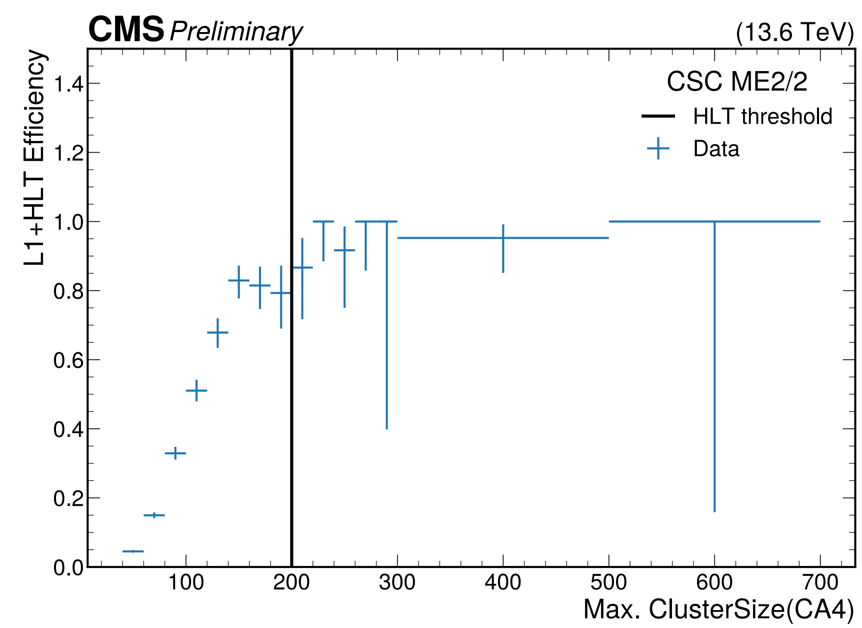

L1HMT

Upgraded LHCb online system

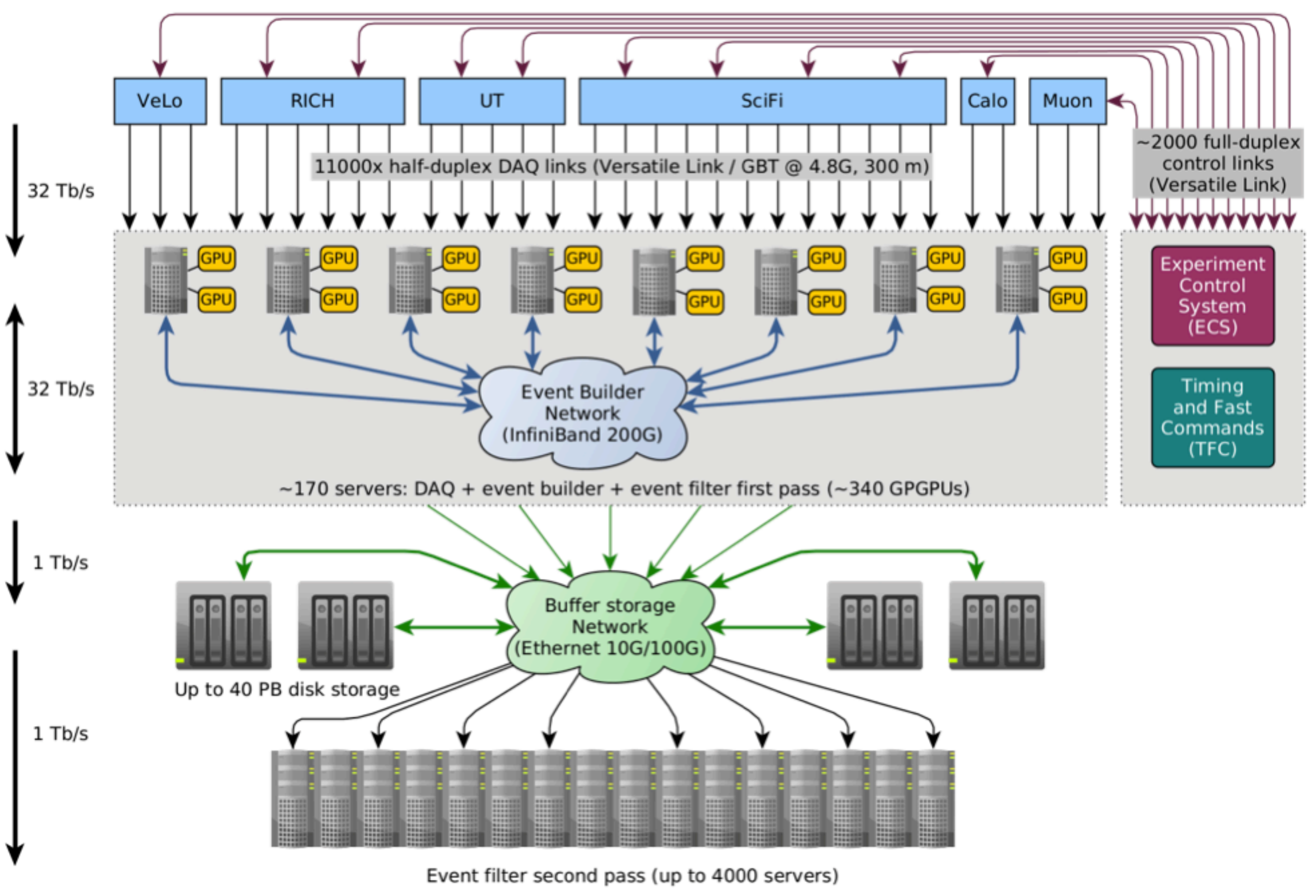

- Event builder farm equipped with 173 servers

- Each server has 3 free PCIe slots
  - can host GPUs
  - sufficient cooling and power
  - advantageous to have GPUs as self-contained processors
  - sending data to GPUs is like sending data to network card

- GPUs map well into LHCb DAQ architecture
- HLT1 tasks inherently parallelizable
- Smaller network between EB & CPU HLT
- Cheaper & more scalable than CPU alternative
- Implemented with 326 Nvidia RTX A5000 GPUs

LHCb HLT1 sequence

Tracking relies on

- **VELO**: clustering, tracking, vertex reconstruction
- **UT**: tracking, momentum estimate, fake rejection
- **SciFi**: tracking, momentum measurement

+

PID from **MUON** and **CALO** systems

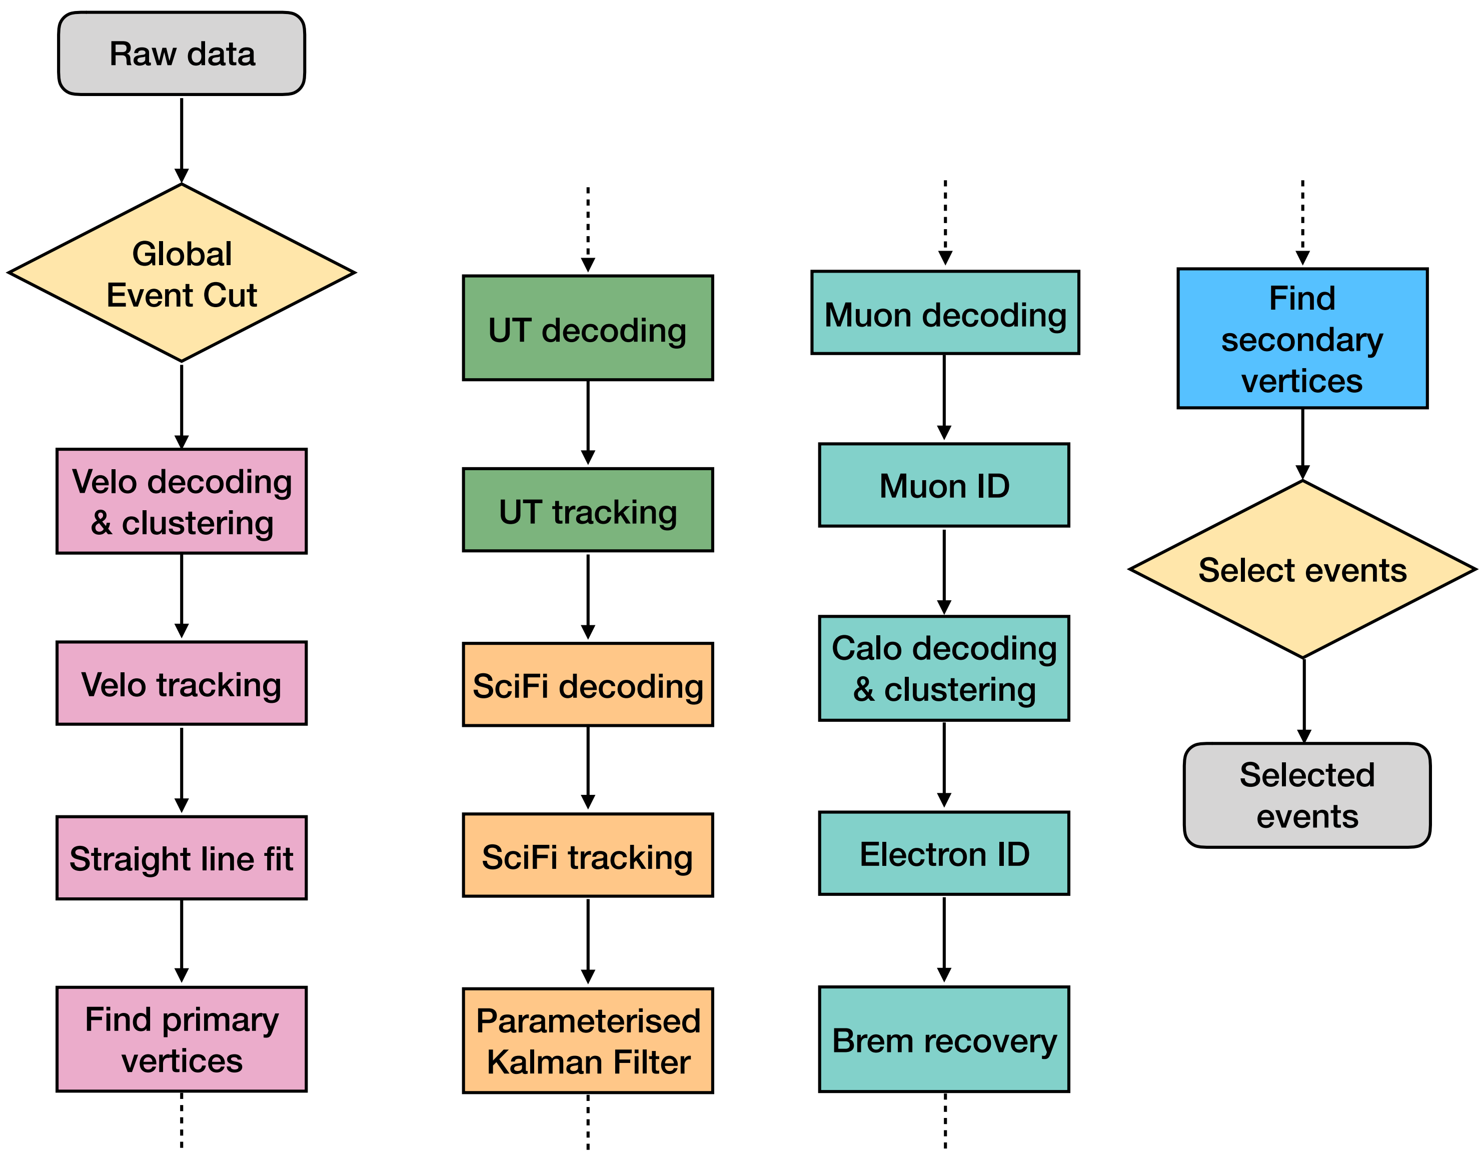

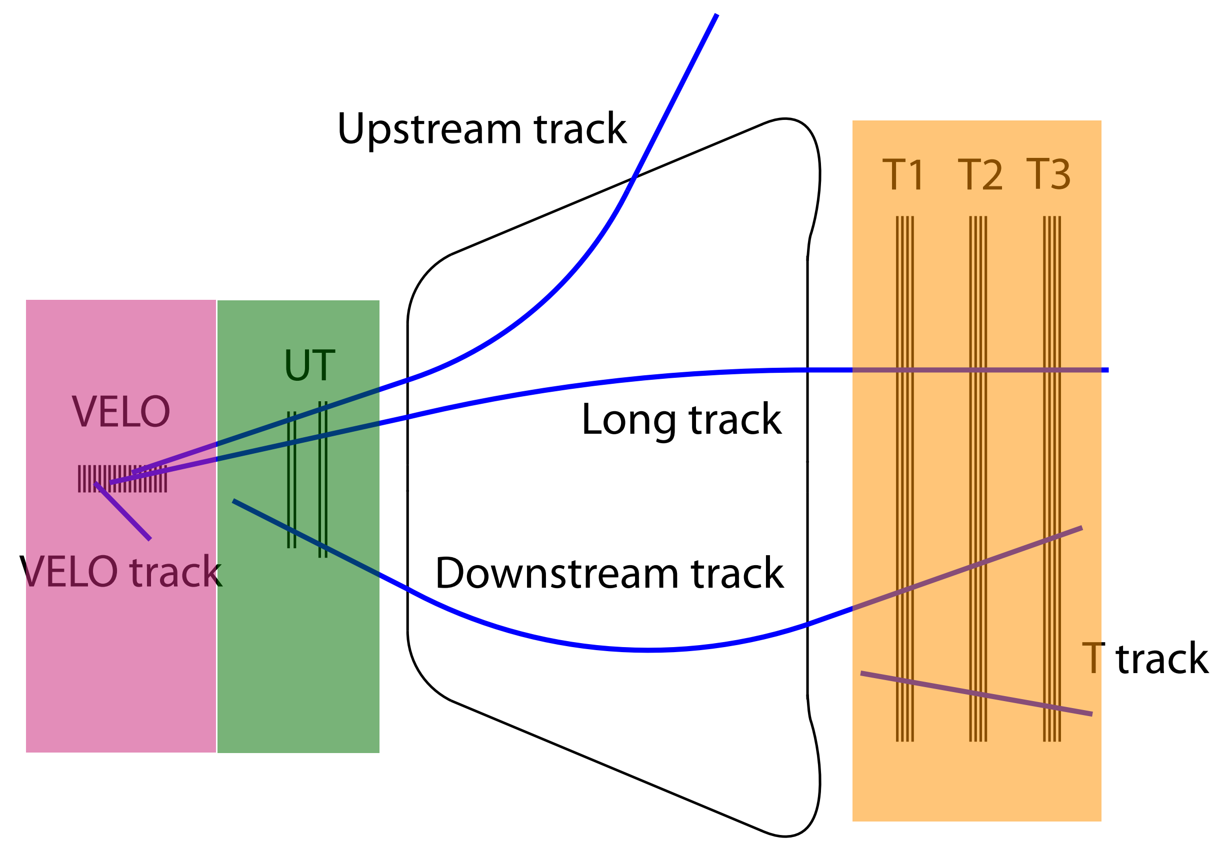

LHCb Track reconstruction

Journal of Computational Science, vol. 54, 2021

Velo tracking

- 26 silicon pixel modules with $\sigma_{x,y} \sim 5\ \mu m$
- Local paralleled clustering algorithm (Search by Triplet)
- Tracks fitted with simple Kalman filter assuming straight line model

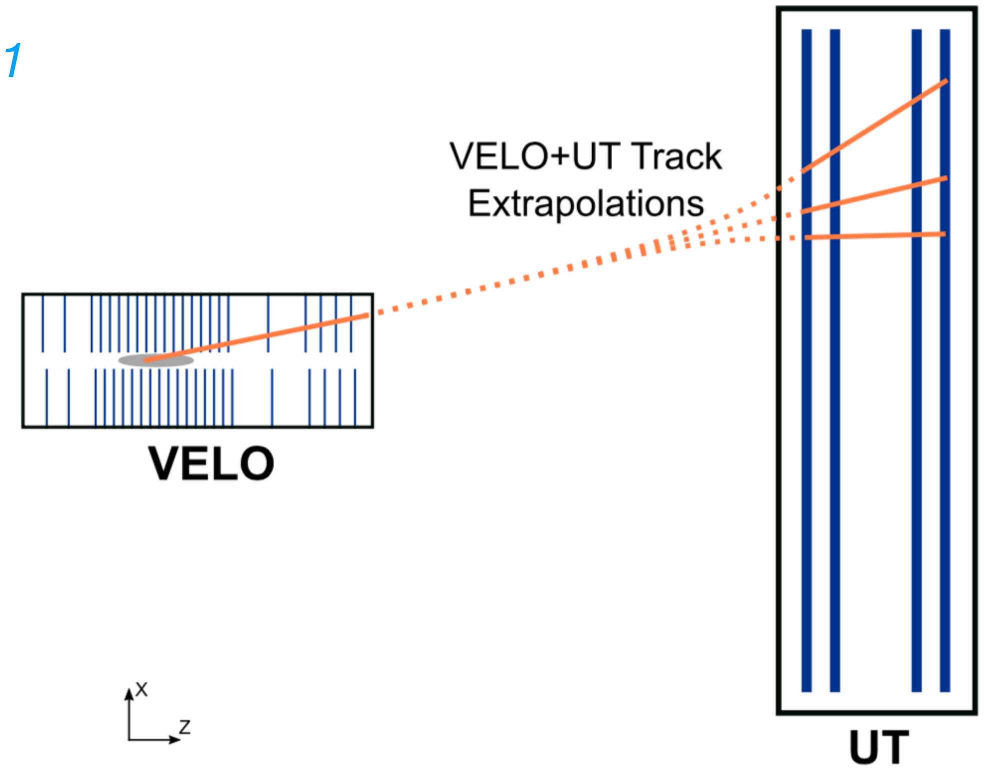

IEEE Access, vol. 7, pp. 91612-91626, 2019

UT tracking

- 4 layers of silicon strips
- Velo tracks extrapolated to UT taking into account B field
- Parallelized trackless finding inside search window requiring at least 3 hits

Comput. Softw. Big Sci 4, 7 (2020)

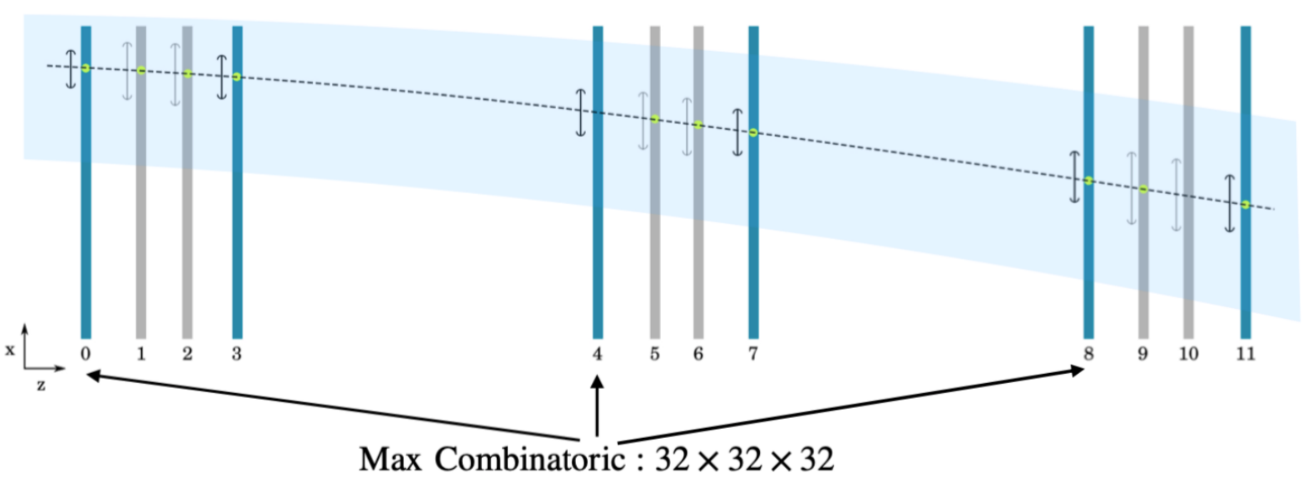

SciFi tracking

- 3 stations with 4 layers of Scintillating Fibres
- Velo-UT tracks extrapolated using parametrisation
- Parallelized Forward algorithm to reconstruct long tracks
  - Search windows from Velo-UT momentum estimate
  - From triplets and extend to remaining layers

LHCb Tracking performance

- **Run 2 performance maintained at x5 instantaneous luminosity**
    - Excellent track reconstruction efficiency (> 99% for VELO, 95% for high-p forward tracks)
    - Good momentum resolution and fake rejection

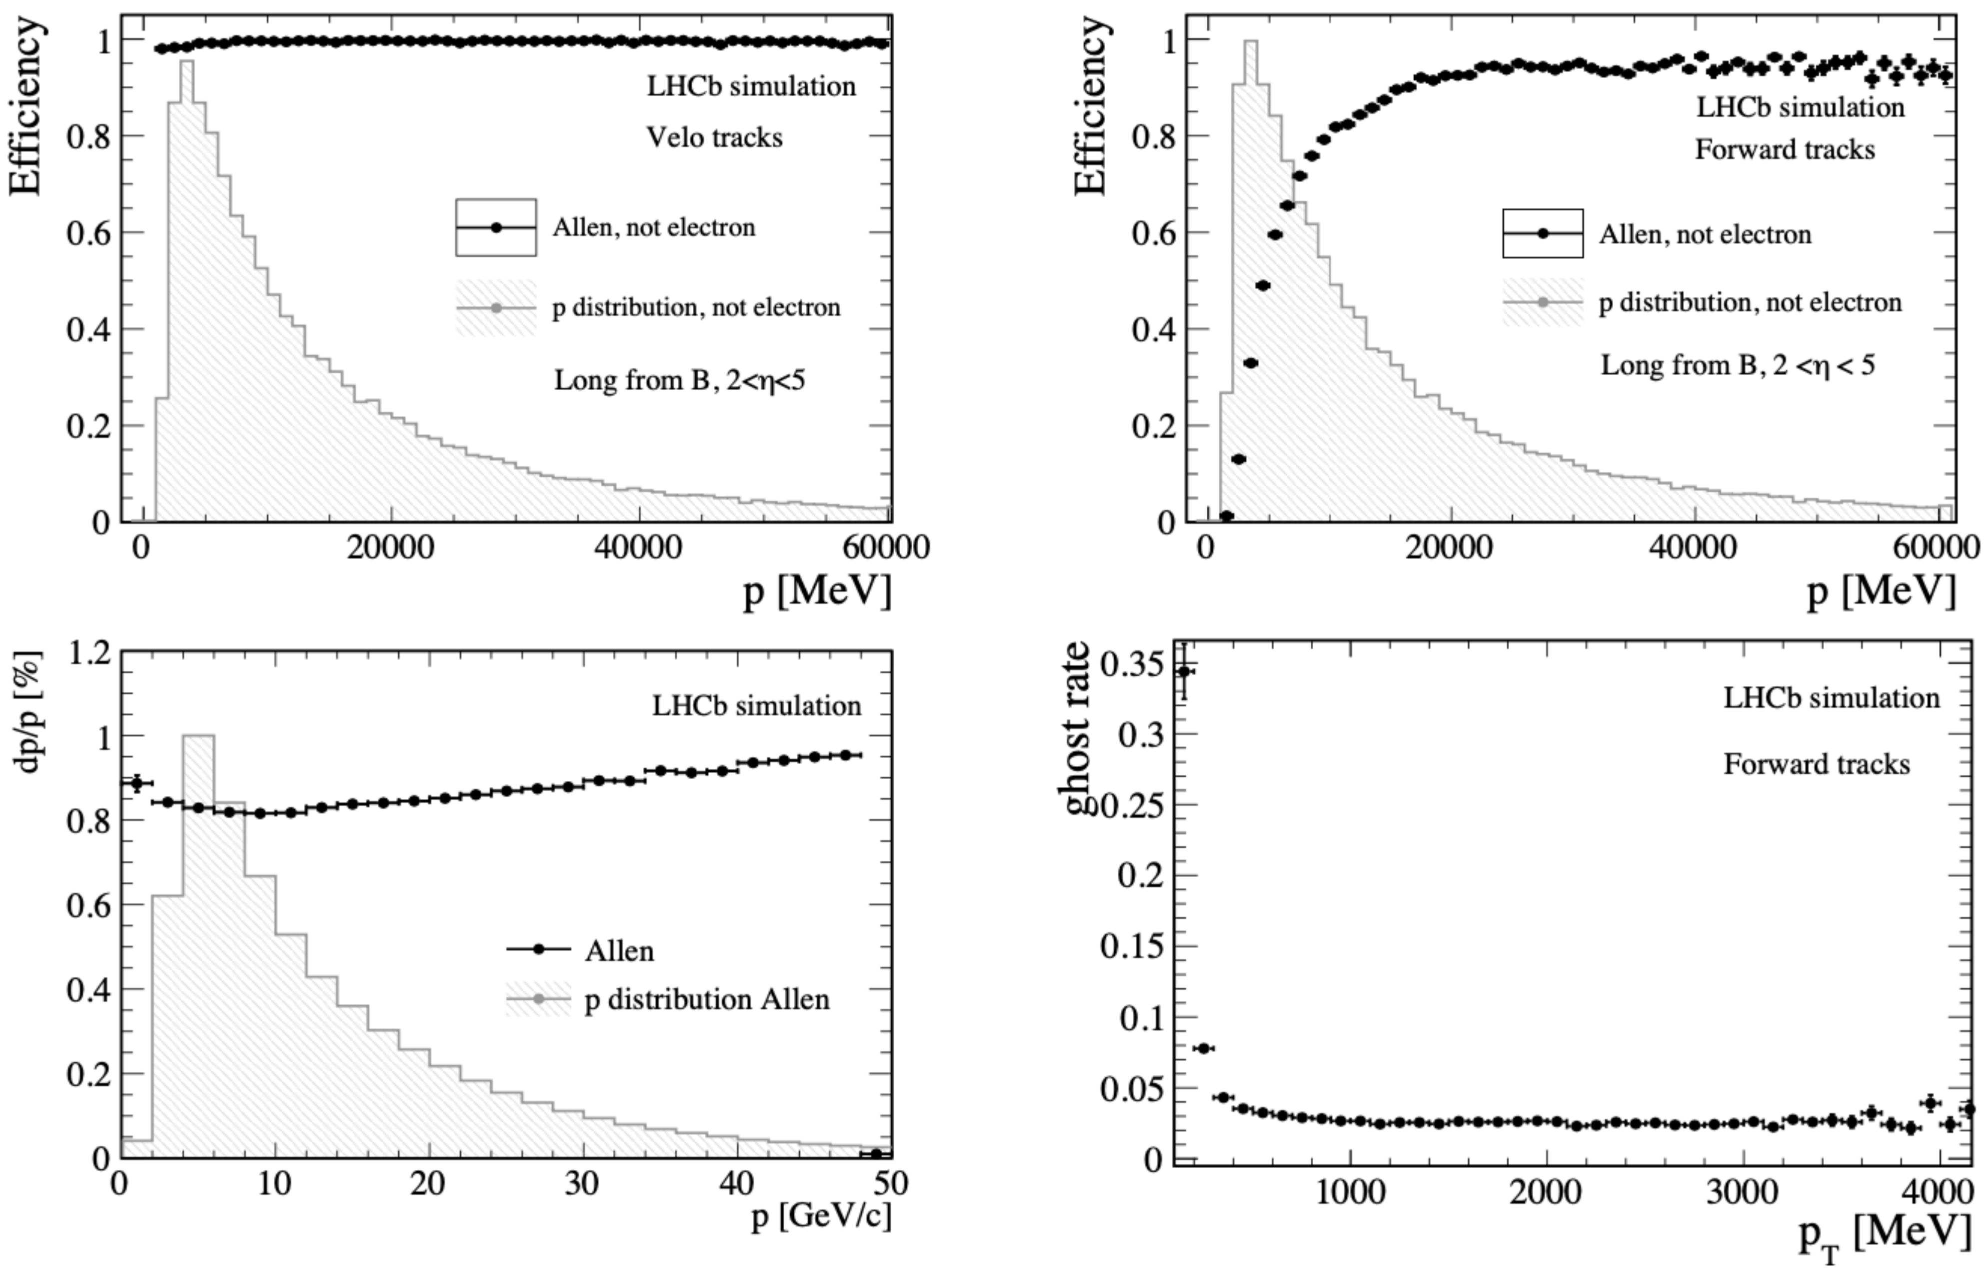

LHCb-FIGURE-2020-014

LHCb Tracking without the UT

- In 2022, the UT detector is unfortunately not be available for data-taking
- Tracking performance and throughput maintained, at the cost of larger fake rate
- Commissioning two options, which **both maintain the current throughput**

1. Forward without UT

- Extrapolate VELO track as a straight line, make two windows — assuming positive/negative charge
- Assume $p > 5$ GeV, $p_T > 1$ GeV (low-p tracks get bent out of the SciFi acceptance anyway)

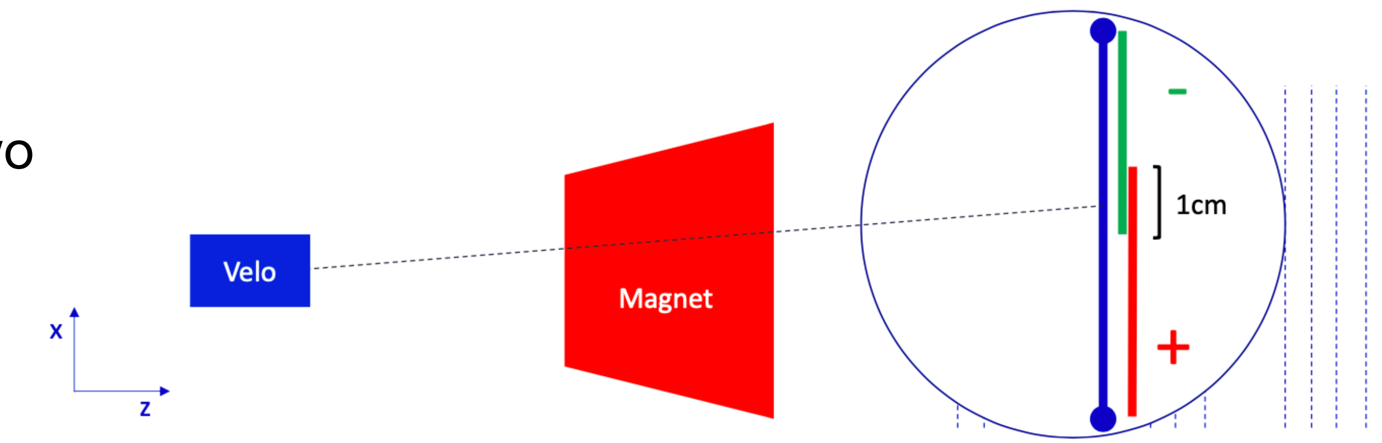

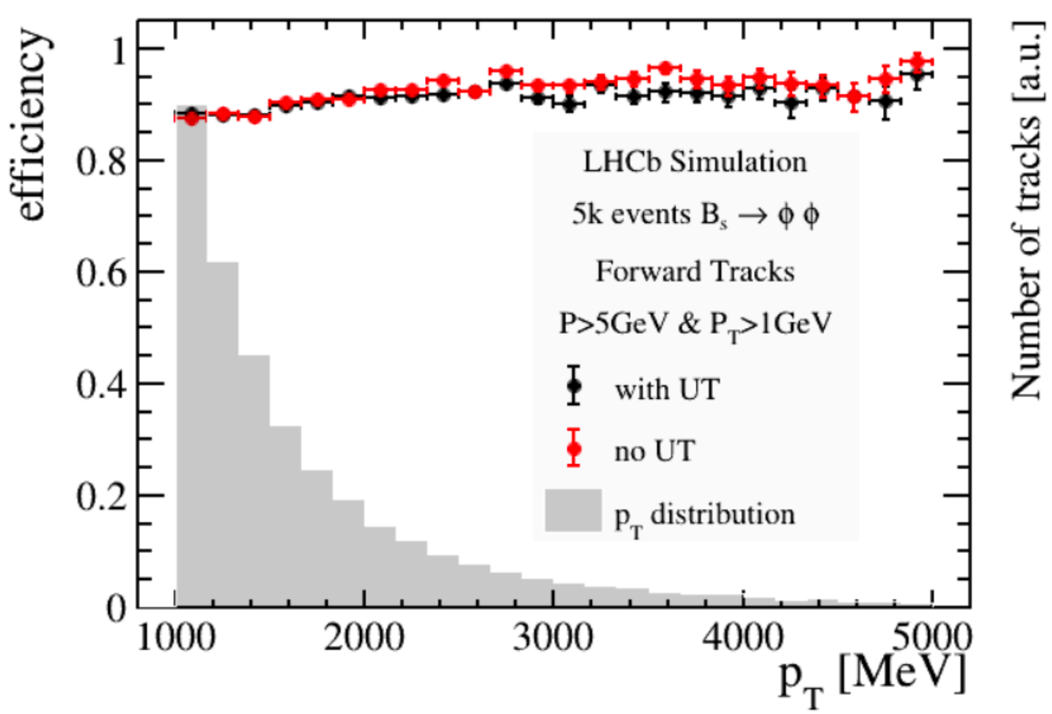

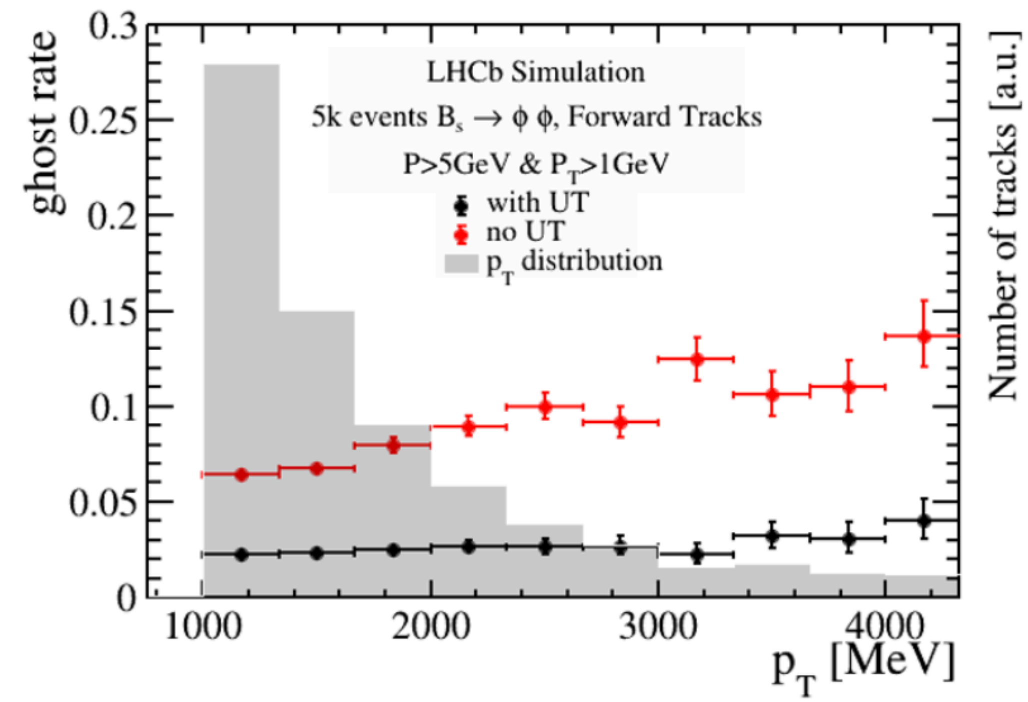

LHCB-FIGURE-2022-007

LHCb Tracking without the UT

- In 2022, the UT detector is unfortunately not be available for data-taking
- Tracking performance and throughput maintained, at the cost of larger fake rate
- Commissioning two options, which **both maintain the current throughput**

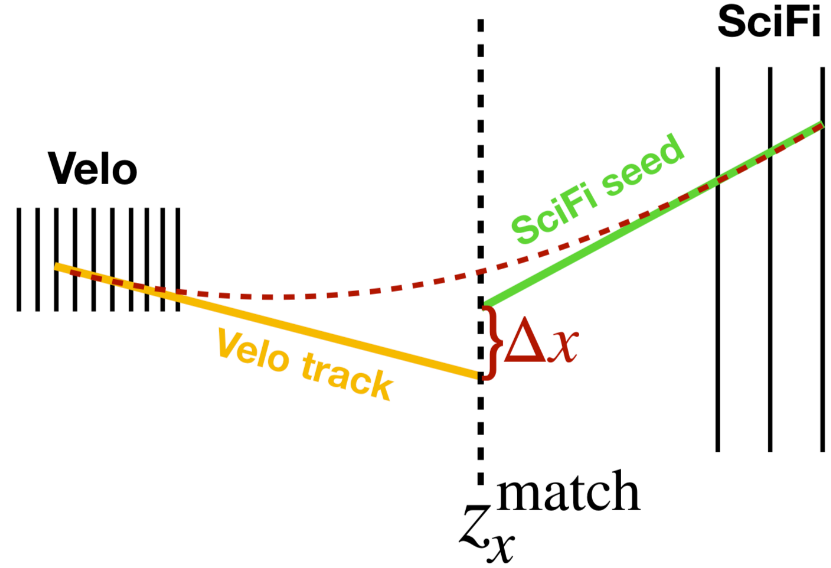

2. Seeding+matching

- Standalone SciFi reconstruction & matching to VELO seeds
- Highly efficient for low momenta
- Opens the door to additional physics cases in HLT1 (downstream and SciFi tracks)

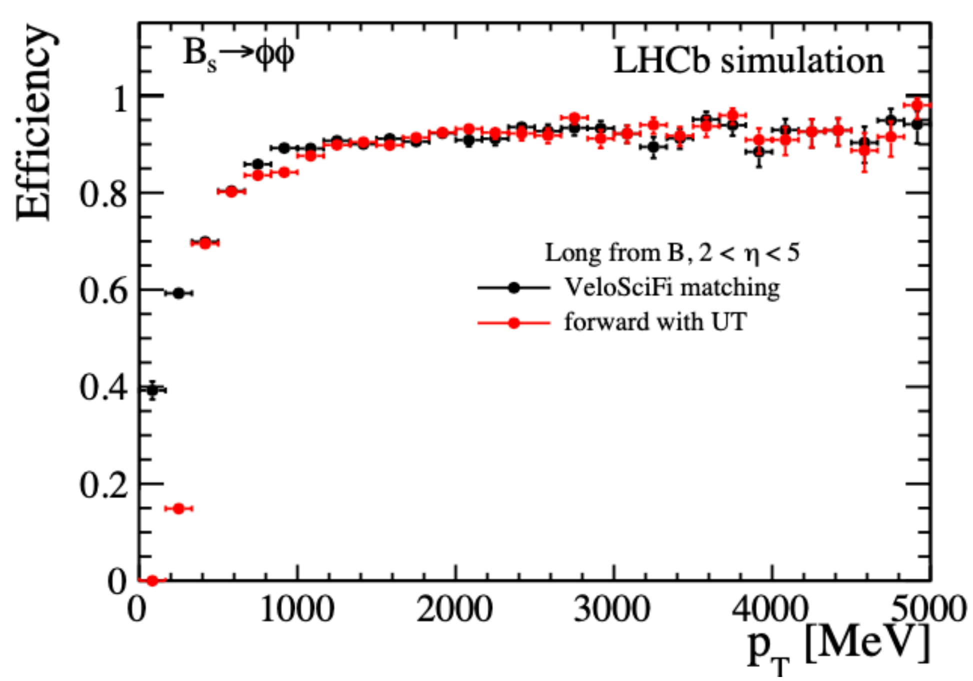

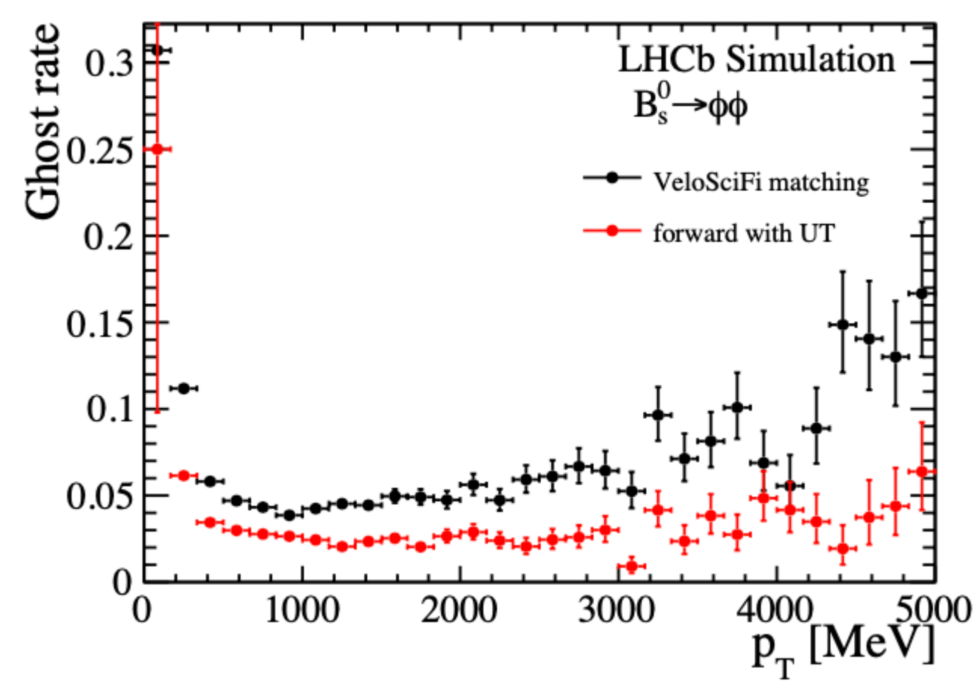

LHCB-FIGURE-2022-010

LHCb Tracking without the UT

- In 2022, the UT detector will unfortunately not be available for data-taking
- Tracking performance and throughput maintained, at the cost of larger fake rate
- Opportunity to commission 2 options, which **both maintain the current throughput**

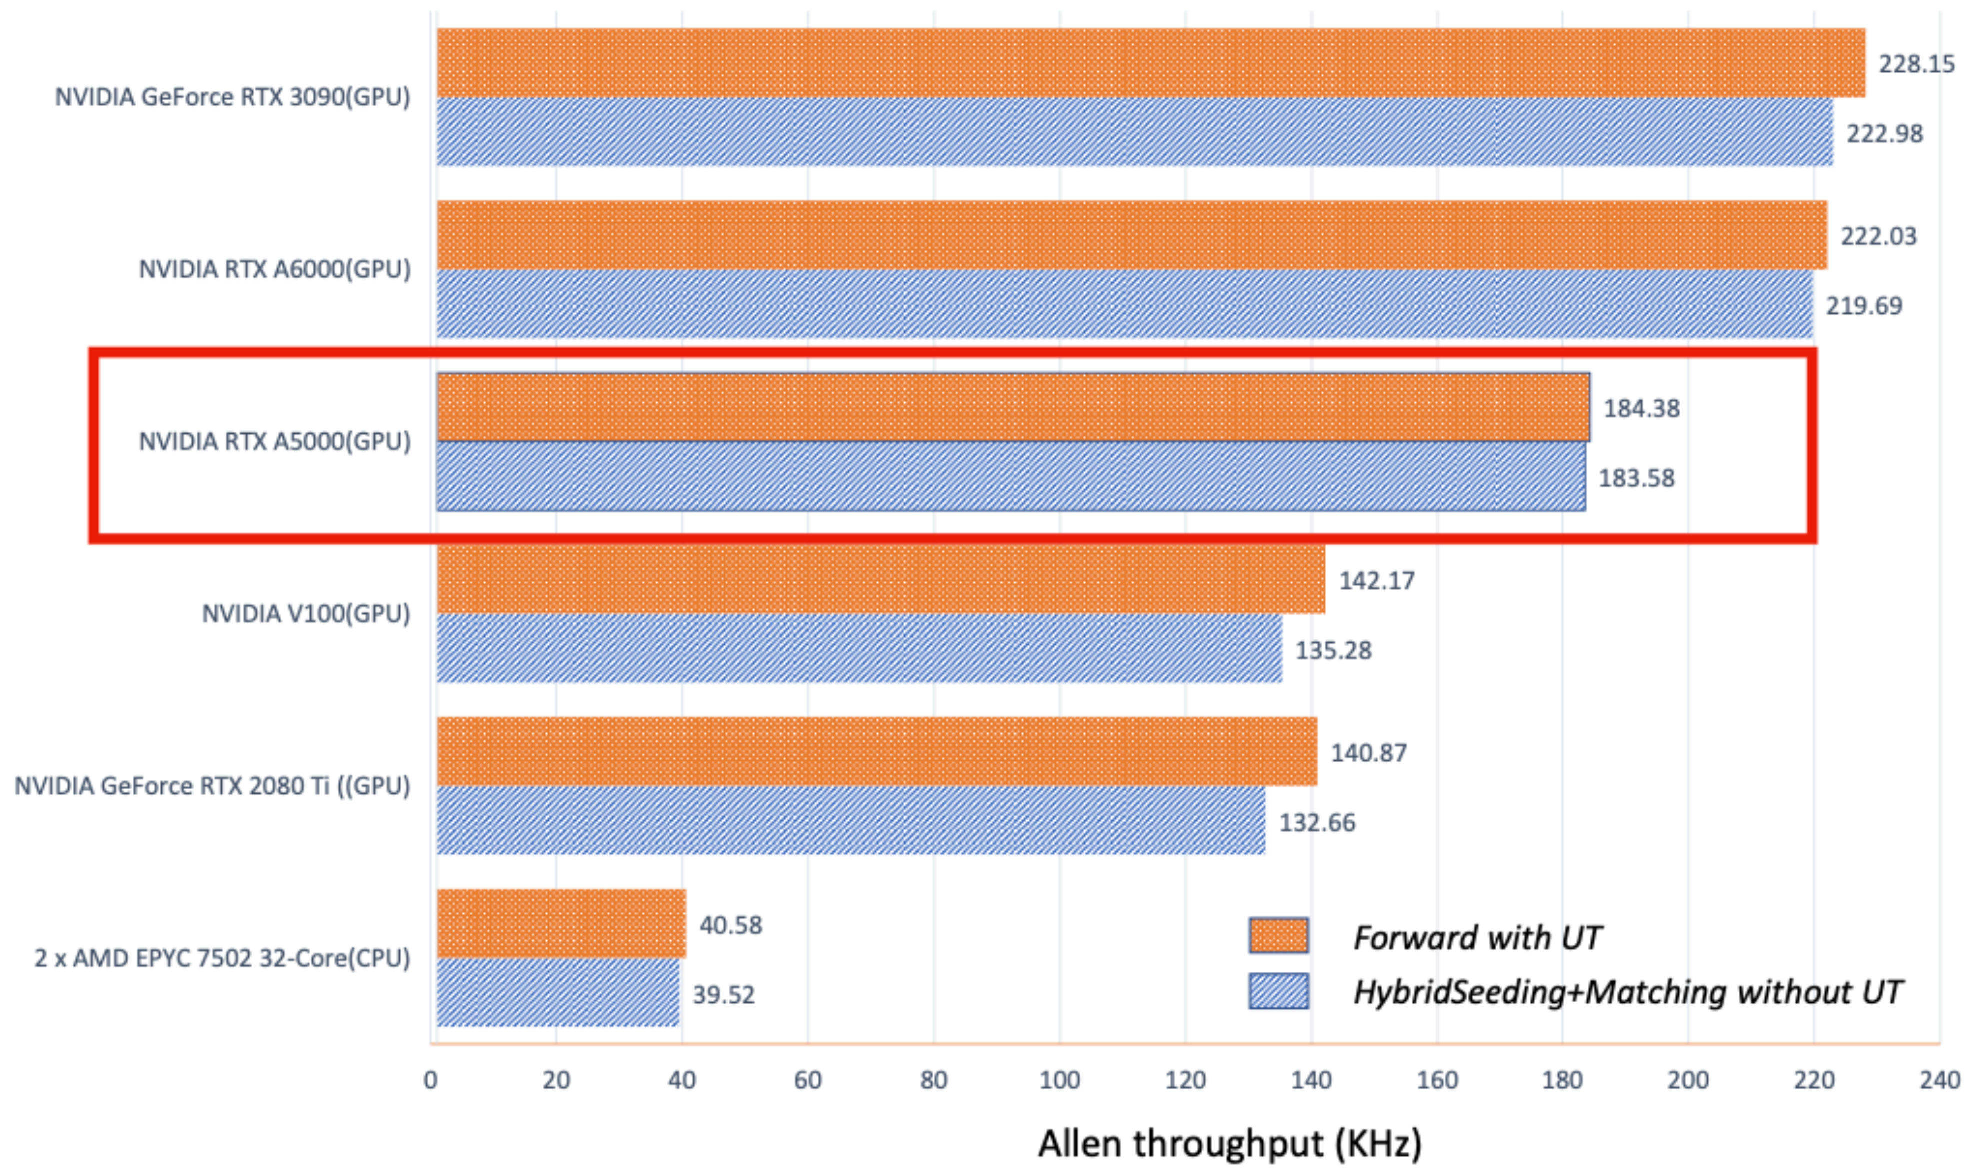

LHCB-FIGURE-2022-010

LHCb Vertex reconstruction

- Primary vertices found from **clusters** in the closest approach of tracks to the beamline
- 1-1 mapping between tracks and vertices requires **serialization**
  - Instead, every track assigned to every vertex based on **weight**
- **Efficiency > 90%** for vertices with number tracks > 10

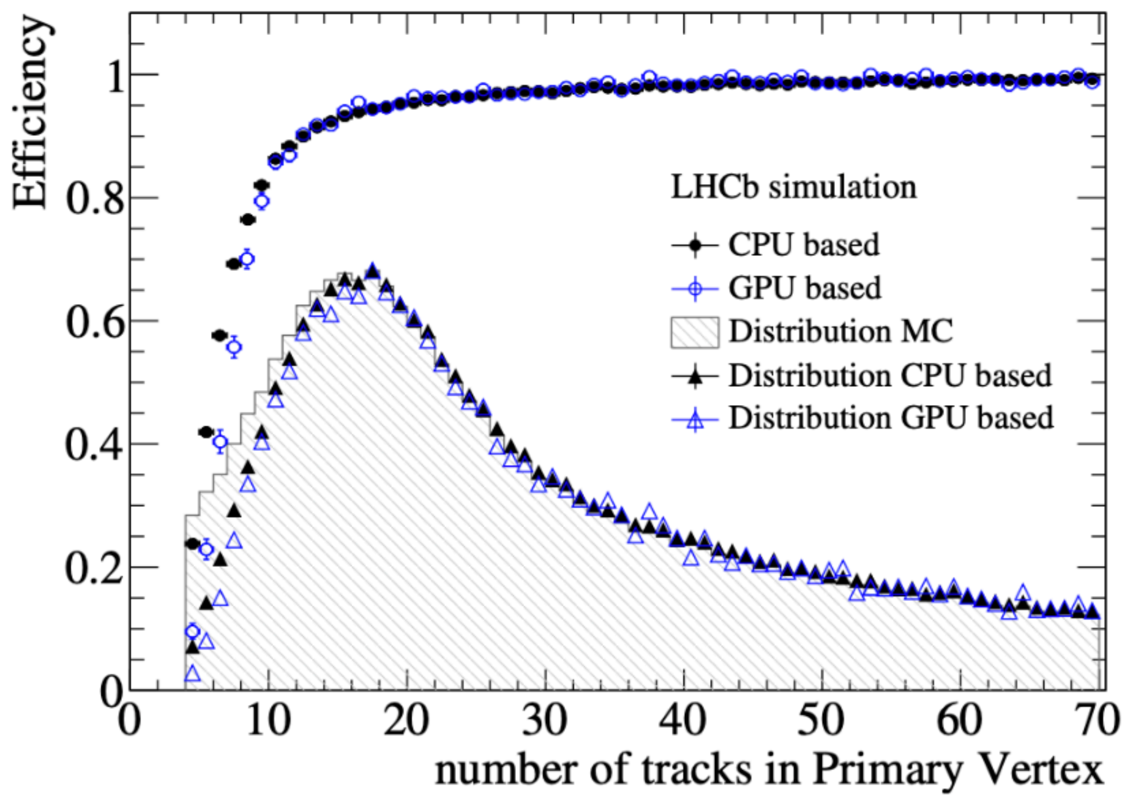

Comput. Softw. Big Sci. 6 (2022) no.1, 1

LHCb Calorimeter and muon PID

- **CALO reconstruction**
  - Loop over calorimeter cells and look for energetic clusters
  - Originally not foreseen within the baseline TDR, but outcome of ambition and good design (and lots of optimisation)
  - The very first algorithm that was tested with real Run 3 data!

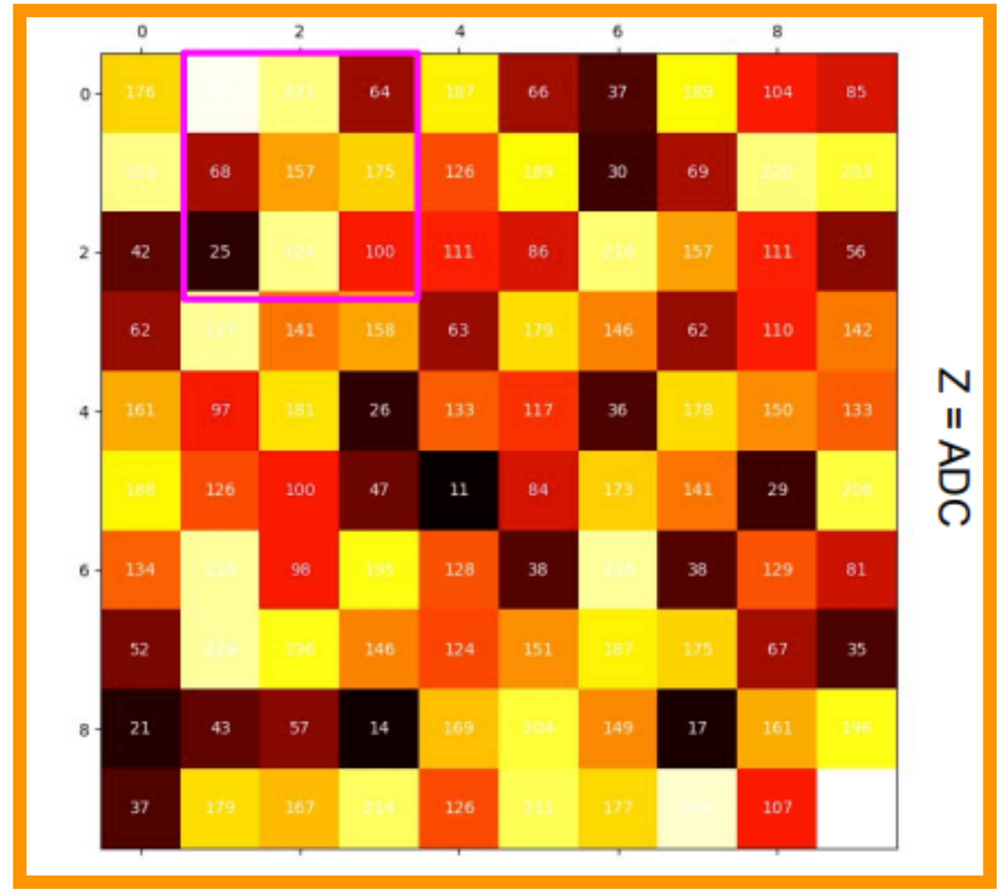

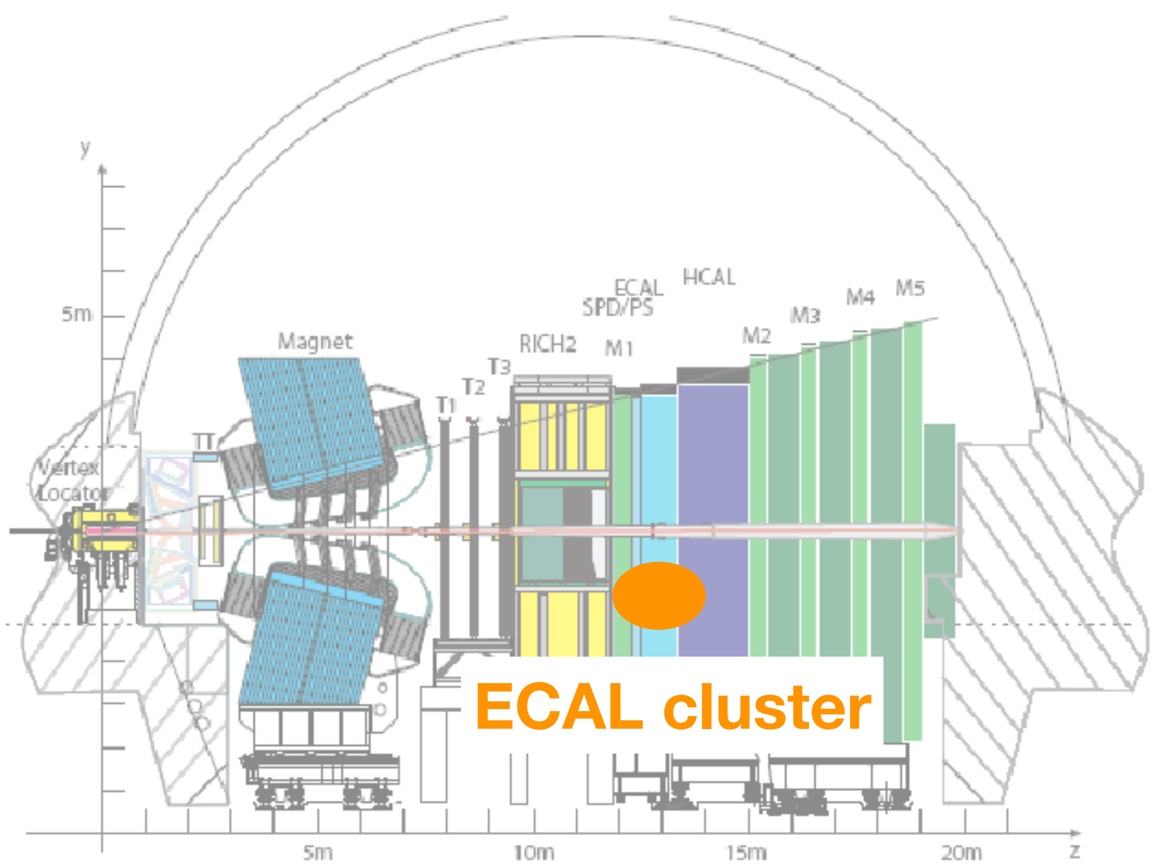

LHCb Calorimeter reconstruction

- **CALO reconstruction**
  - Loop over calorimeter cells and look for energetic clusters
  - Originally not foreseen within the baseline TDR, but outcome of ambition and good design (and lots of optimisation)
  - The very first algorithm that was tested with real Run 3 data!

LHCB-FIGURE-2021-003

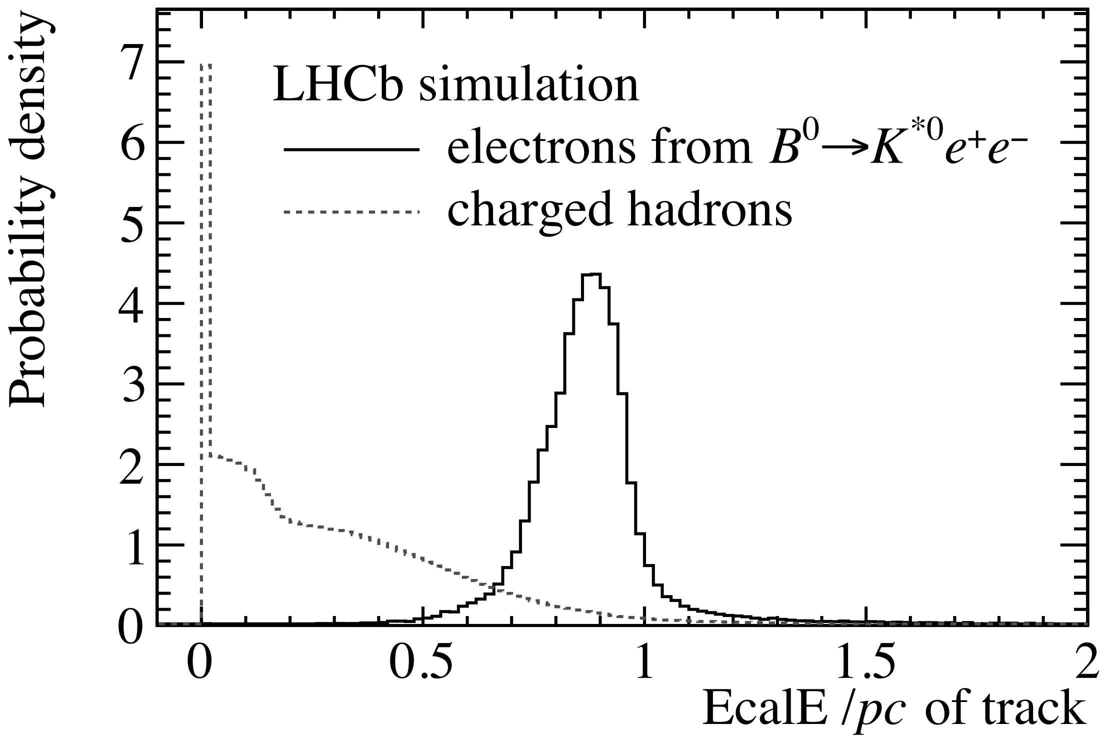

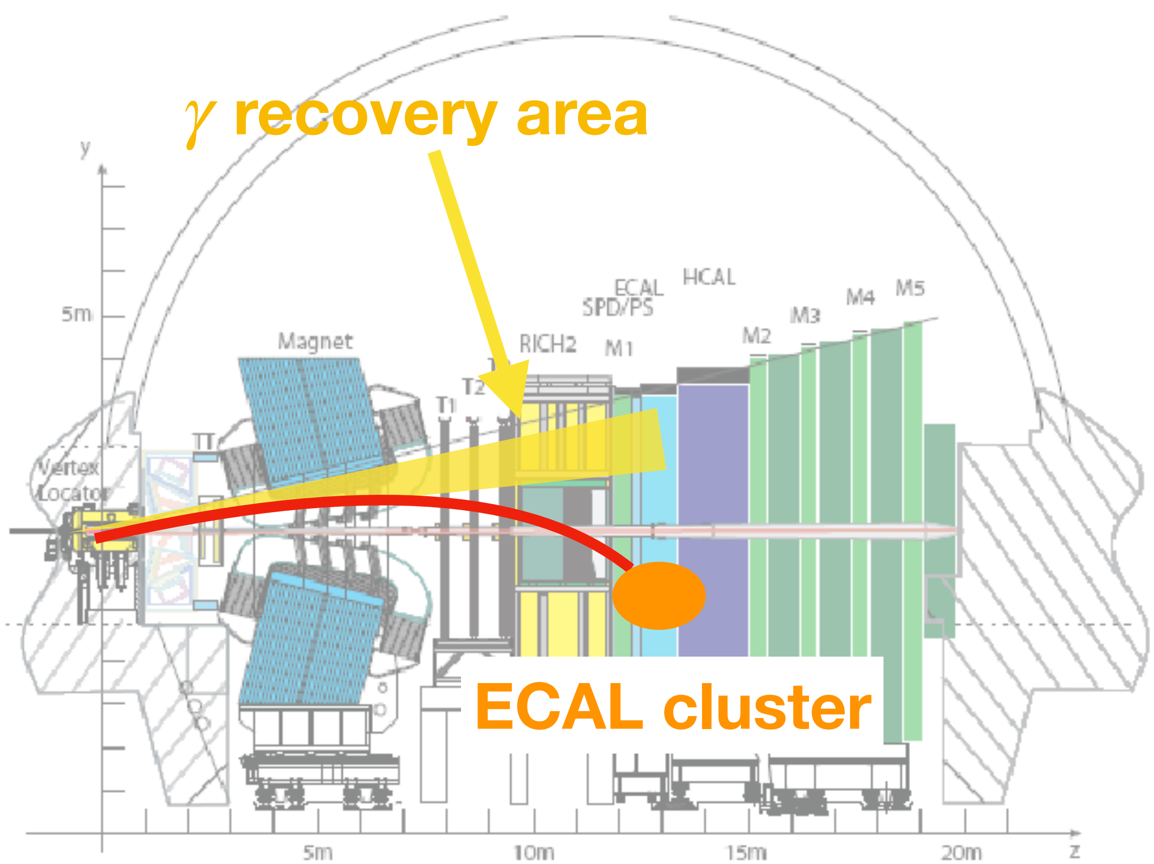

- Calo digits attached to long tracks for electrons
- Momentum is corrected if clusters are found in the Bremsstrahlung recovery area

LHCb Muon PID

- **Muon particle identification**
    - Extrapolate tracks from SciFi to Muon stations
    - Match hits to tracks in a field of interest
    - Excellent muon identification and misID background rejection

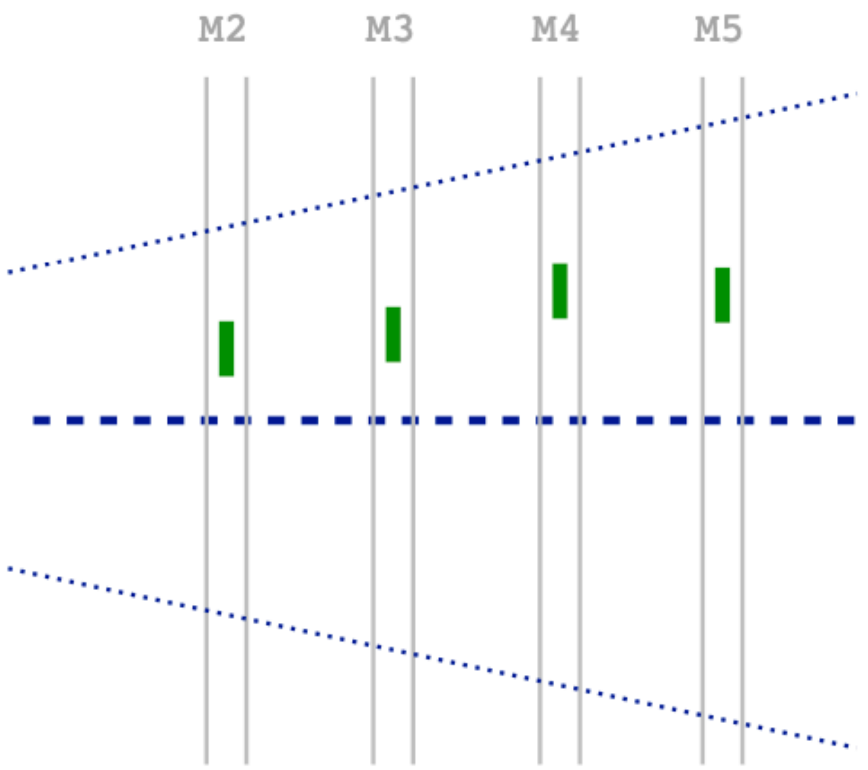

LHCB-FIGURE-2020-014

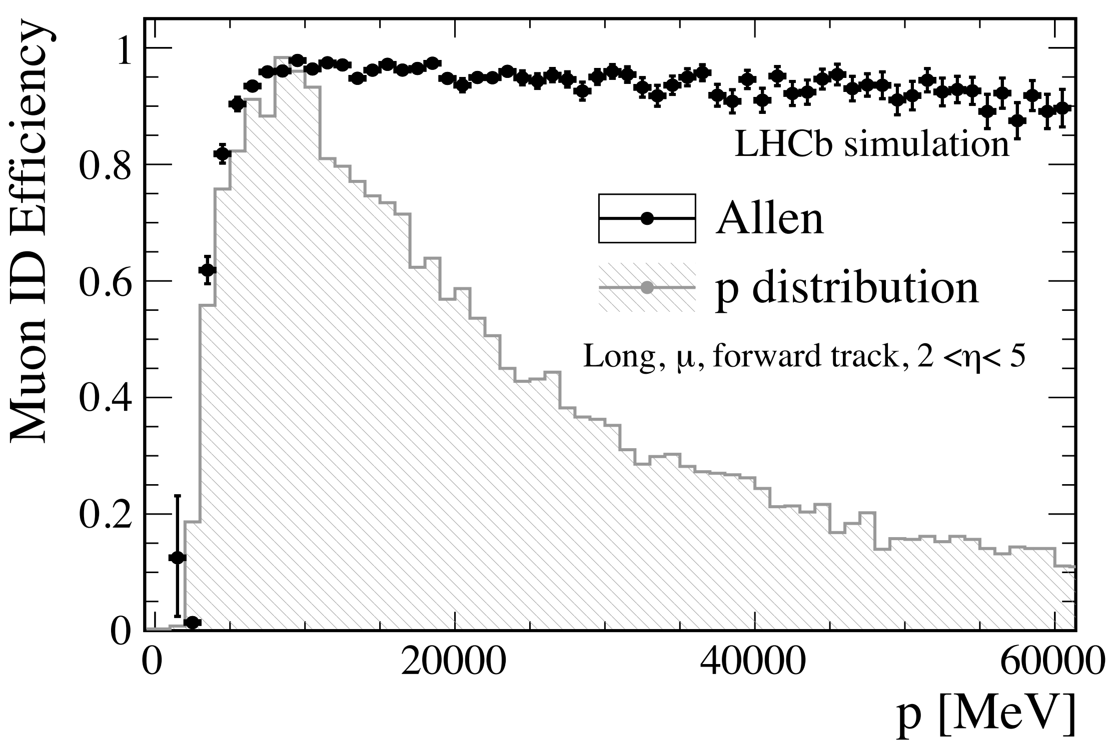

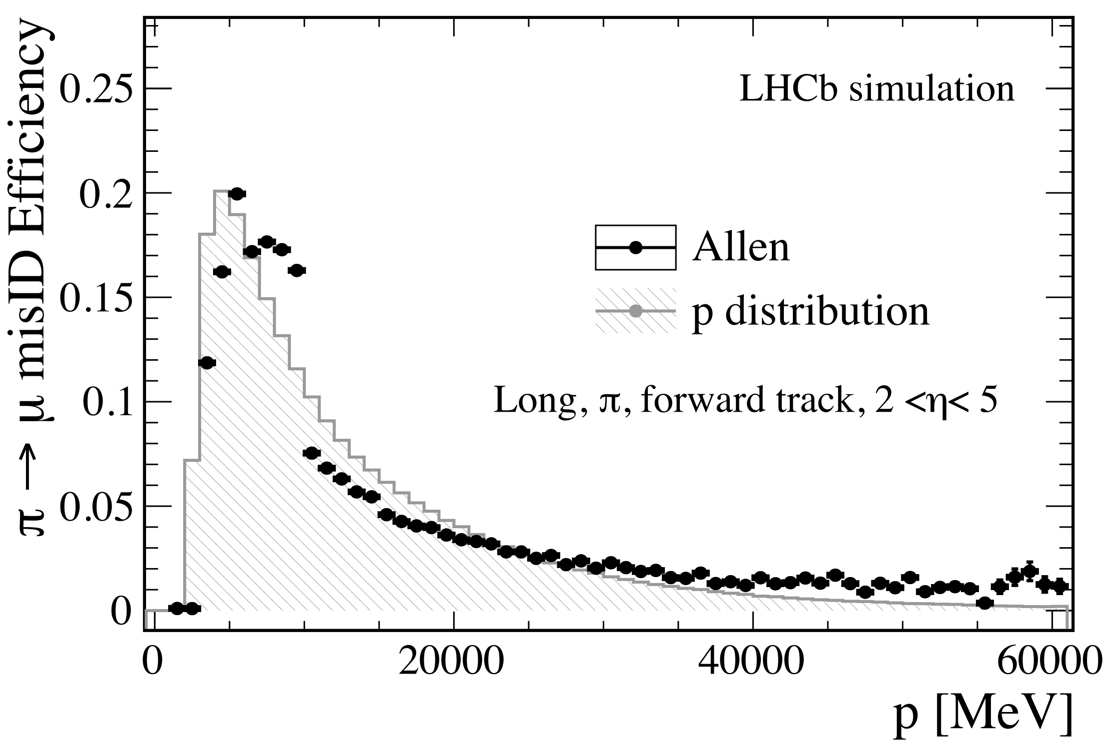

LHCb HLT1 selection performance

- Inclusive rate for the main HLT1 lines ~ 1 MHz
- O(30) lines implemented so far:
    - Cover majority of LHCb physics programme (B, D decays, semileptonic, EW physics)
    - Special lines for monitoring, alignment and calibration

TwoTrackMVA line for $B_s^0 \to \phi\phi$

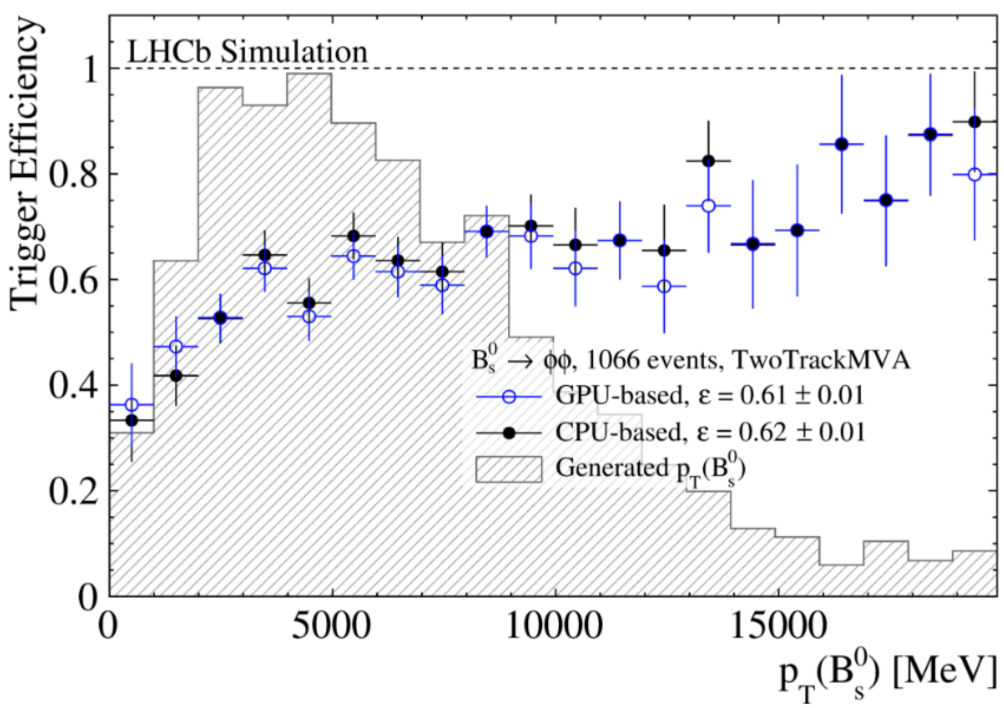

SingleHighPtMuon line for $Z \to \mu^+\mu^-$

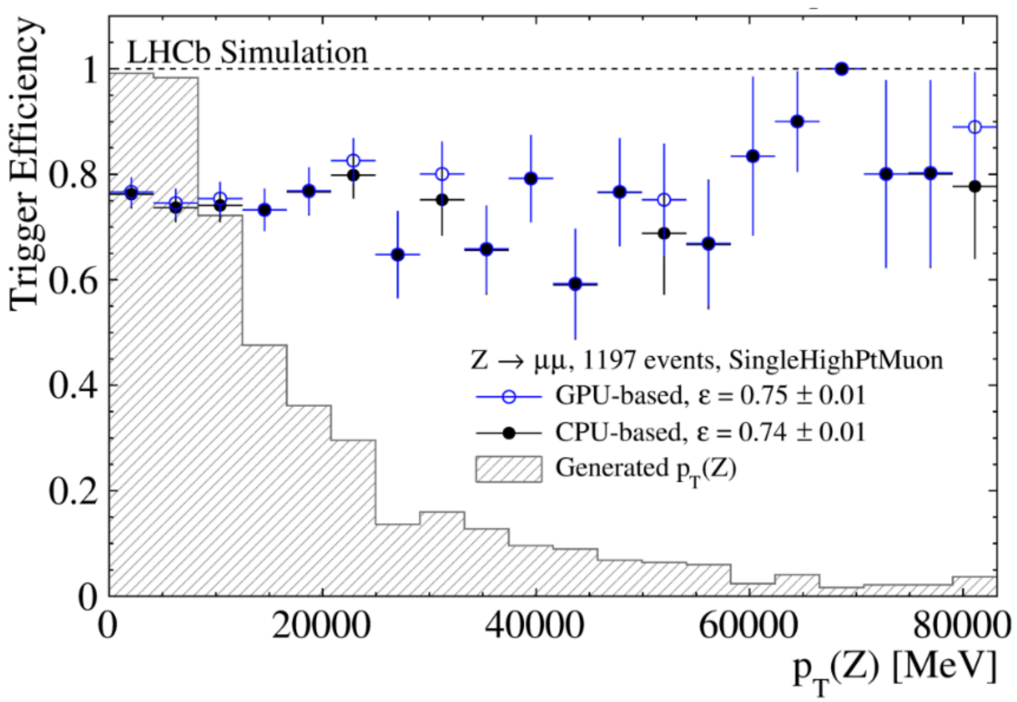

Comput. Softw. Big Sci. 6 (2022) no.1, 1

Allen: a GPU HLT1 LHCb trigger platform

- Public software project: [gitlab repo]
- Supports three modes:
  - Standalone
  - Compiling within the LHCb framework for dat acquisition
  - Compiling within the LHCb framework for simulation and offline studies
- Runs on CPU, Nvidia GPU (CUDA, CUDACLANG), AMD GPUs (HIP)
- GPU code written in CUDA
- Cross-architecture compatibility (HIP, CPU) via macros

Allen

pipeline passed

Welcome to Allen, a project providing a full HLT1 realization on GPU.

Documentation can be found here.

Mattermost discussion channels

- Allen developers - Channel for any Allen algorithm development discussion.
- Allen core - Discussion of Allen core features.
- AllenPR throughput - Throughput reports from nightlies and MRs.

Performance monitoring

- Allen throughput evolution over time in grafana
- Allen dashboard with physics performance over time

» Welcome to Allen's documentation! Edit on GitLab

Welcome to Allen's documentation!

Allen is the LHCb high-level trigger 1 (HLT1) application on graphics processing units (GPUs). It is responsible for filtering an input rate of 30 million collisions per second down to an output rate of around 1-2 MHz. It does this by performing fast track reconstruction and selecting pp collision events based on one- and two-track objects entirely on GPUs.

This site documents various aspects of Allen.

LHCb HLT1 CPU/GPU comparison

Compatible performance between CPU and GPU!

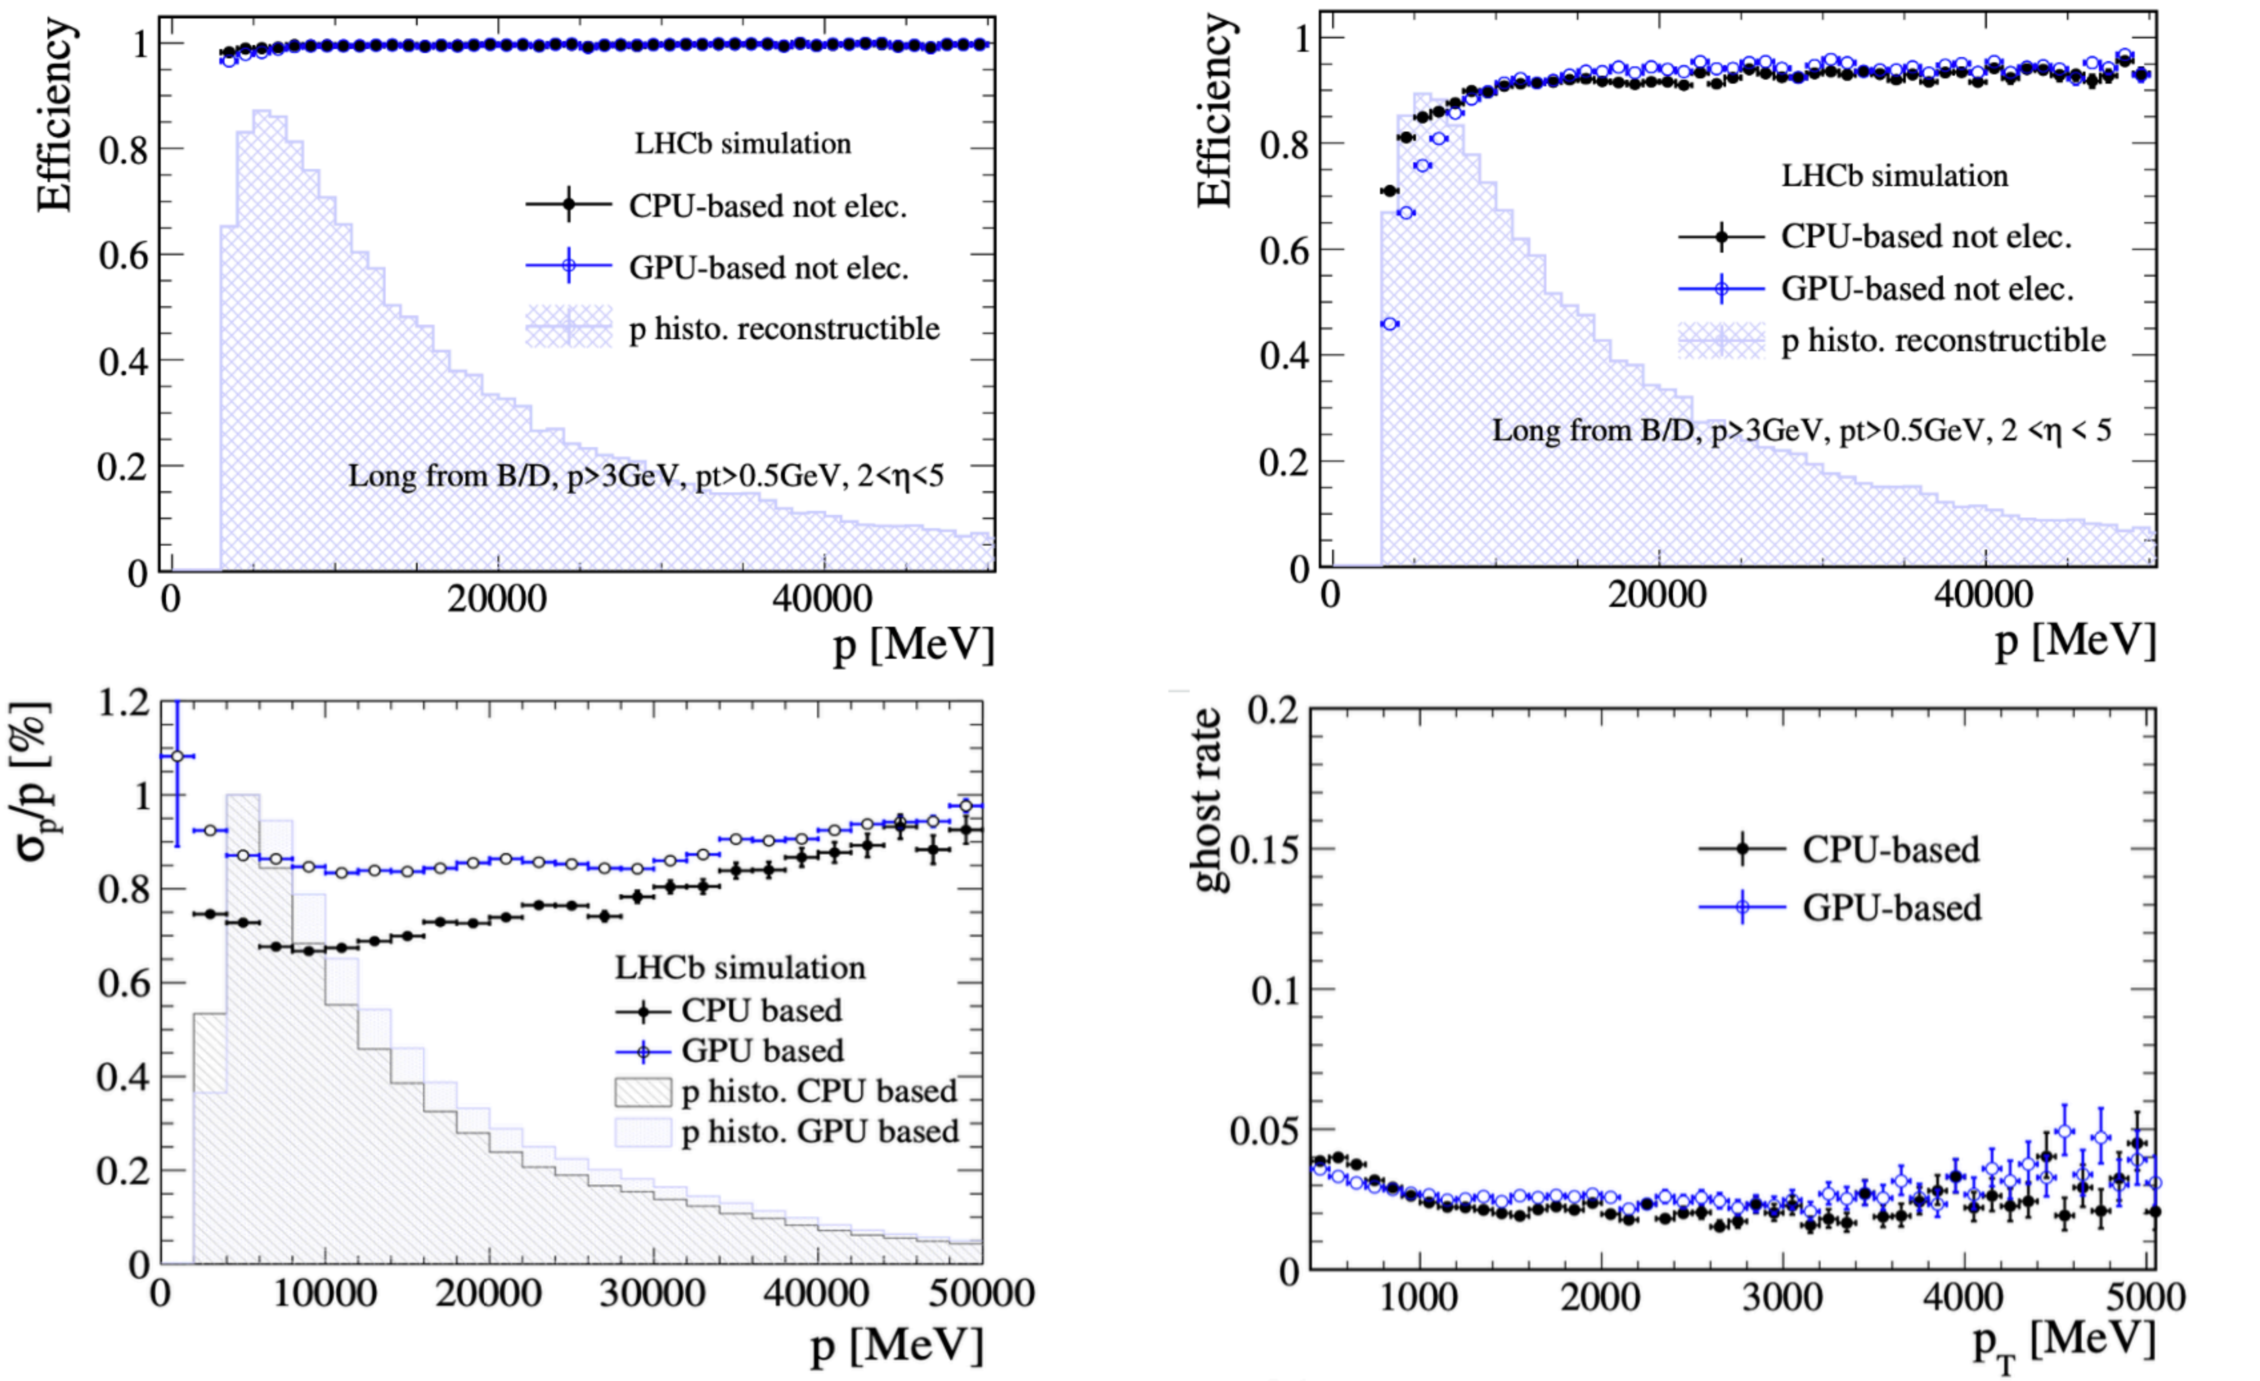

Comput. Softw. Big Sci. 6 (2022) no.1, 1

LHCb HLT1 CPU/GPU comparison

- Compatible performance between CPU and GPU [*Comput. Softw. Big Sci. 6 (2022) no.1, 1*]()