FPGA acceleration for HLT1 at LHCb

Ao Xu (许傲)

CEPC Workshop 2024

October 24, 2024

Need for speed

- Progress of experiments goes together with increasing bandwidth
- Flavour physics at low p_{T} is most demanding ($\sigma_b \sim 10^{4/7}\sigma_{Z/H}$)
  - LHCb tops the chart despite smaller size and lower luminosity
  - Increasing $\mathcal{L}(t)$: $\mathcal{L}_{\mathrm{Run5}} = 7.5 \times \mathcal{L}_{\mathrm{Run3/4}} = 7.5 \times 5 \times \mathcal{L}_{\mathrm{Run1/2}}$
  - LHCb is effectively "processing-limited"

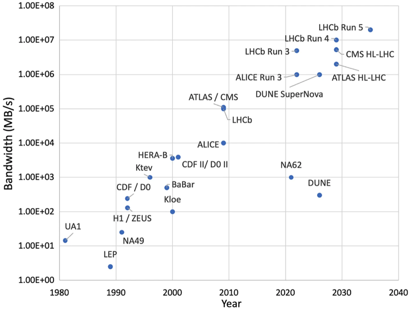

[EPJ Plus 138 1005 (2023)]

LHCb Data Processing model in Run 3

- Triggerless readout of whole detector + full event reconstruction
  - No inefficient hardware trigger on simple quantities (*e.g.* p_{T}, E_T)

[NIMA 1069 (2024) 169782]

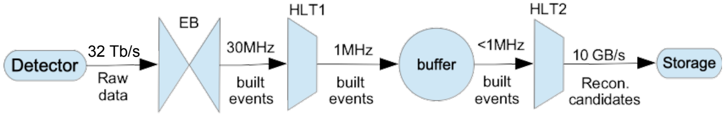

- Two-level trigger system
  - HLT1 (GPUs): full track reconstruction for trigger purpose
  - HLT2 (CPUs): full detector reconstruction and final physics selection
- Alignments performed between HLT1 and HLT2

Evolving towards primitive-based reconstruction

- Pre-EB: reconstruct intermediate data structures (*primitives*)

[NIMA 1069 (2024) 169782]

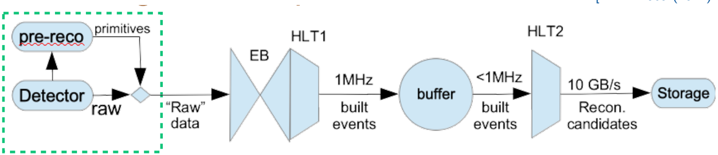

  - Clusters, track segments *etc.*
  - Locally embedded in the detector: looks like "Raw" data to DAQ
  - Accelerate HLT reconstruction and reduce data flow at the source

- Challenges: no time-multiplexing, little buffering
- Solutions
  - ASICs (on detector): *e.g.* CMS track vectors
  - **FPGAs (off detctor)**: for more complex primitives

LHCb Run 3 tracking system and sequence

- Current HLT1 reconstructions focus on Long tracks
  - Run 2 based on Forward tracking
  - Run 3 benefits also from Matching
- Add Downstream tracks to HLT1
  - Expand the LHCb physics program
- FPGA acceleration at pre-build level
  - **VELO clustering**: established as default method at LHCb
  - **T-track primitives**: approved enhancement for Run 4

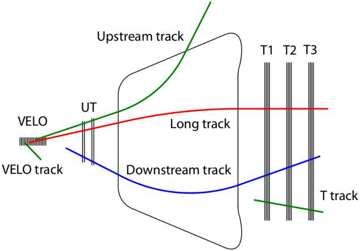

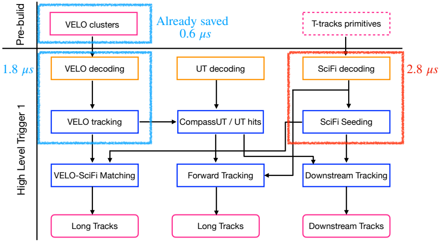

- Hits in VELO detector appear as 2D clusters of pixels
- Firmware deployed in Run 3 in FPGA readout boards (Arria 10) to make clusters on the fly
  - Pixels read out as 2×4 arrays (SuperPixels, SP)
  - Clusters found by unpacking them into active matrices, where each pixel actively checks if it belongs to a pattern
  - Centroid evaluated by LUTs
  - Dynamically allocate small matrixes where active pixels are found

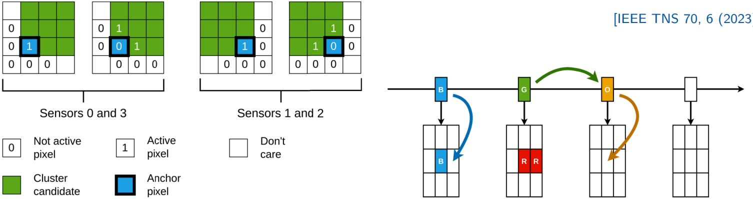

[IEEE TNS 70, 6 (2023)]

Performance of embedded cluster finding

- Quality of real-time clustering as good as CPU algorithm

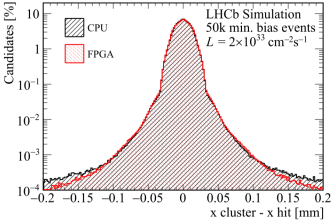

[IEEE TNS 70, 6 (2023)]

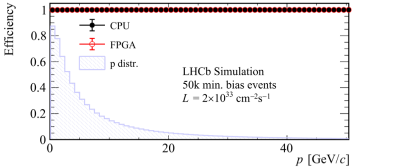

- 14% BW reduction: raw pixel information dropped and replaced by hit positions during readout
- 11% throughput increase and 1/50 power consumption
- Enable opportunities for further applications (*e.g.* precision monitoring of beamline): real-time availability of 10^11 hits/s in accessible way

DoWnstream Tracker

[CERN-LHCC-2024-001]

- **DoWnstream Tracker** will provide HLT1 with pre-formed T-track primitives in LHCb Run 4
  - Make room for Downstream tracking and other desirable enhancements
- Artificial retina architecture: highly-parallel architecture for pattern recognition
  - High throughput and low latency
  - Extreme parallelism and high connectivity
  - Computation similar to Hough transform
- Implemented in an array of FPGAs
  - Each board specialised to reconstruct a portion of parameter space
  - Each board processes each event
  - Connection between boards through a custom network

LHCb Scintillating Fibre Tracker

[CERN-LHCC-2014-001]

- **Three** tracking stations: T1, T2, T3
- Each consists of **four** detection planes: oriented $(0°, +5°, -5°, 0°)$
  - Modules have 2.5 m long scintillating fibres with a diameter of $250\,\mu m$ read out by SiPMs
  - Measurements of the co-ordinates (x, u, v, x)

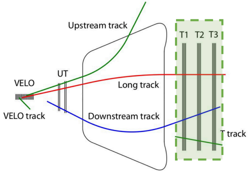

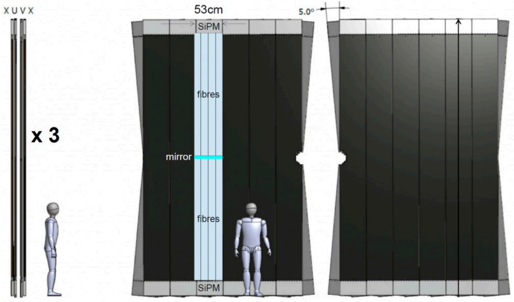

1. Axial (x-z plane) track parametrisation
   - $(\tilde{x}_0, \tilde{x}_{11})$: x-coordinates at the **first** and **last** SciFi layer
   - # of pattern cells for SciFi: **2×73k**

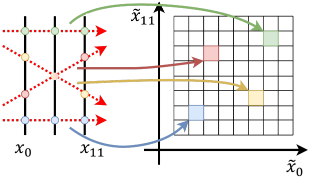

2. Weight accumulation
   - $w = \sum_{hits} \exp\left(-\frac{(x_l - t_l)^2}{2\sigma}\right)$ for $|x_l - t_l| < d_s$

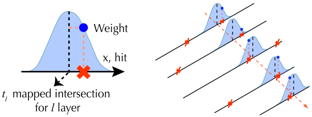

3. Identification of local maxima (axial track primitives)
   - **Maximum** above **threshold** in the centered 3×3 cluster

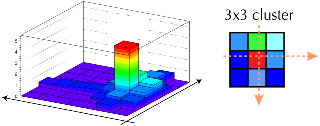

Ghost removal with axial track fit

- Linearised χ^2 fit for false maxima removal
- Parabolic model with cubic correction [1, 2]

$$x(z) = a_x + b_x \times z + c_x \times z^2 \times (1 + \mathrm{dRatio} \times z)$$

- For each local maximum determine the best fit over combinations of
  - 5 different axial layers out of 6
  - 1 out of 2 candidate hits on each layer

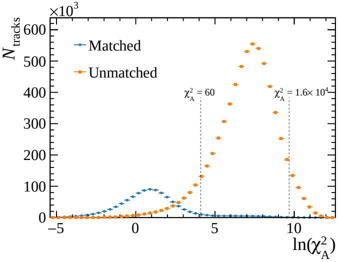

Before χ^2_A requirement

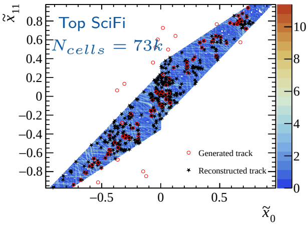

After $\chi^2_A < 60$

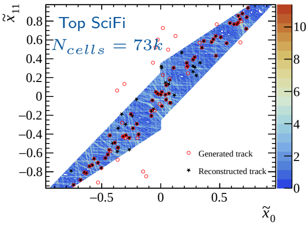

1. Stereo (y-z plane) track parametrisation
   - $\tilde{y}$: y-coordinate at the middle of SciFi
   - # of bins per axial track: 45
2. u/v hits distribution
   - Good axial track candidate $\longleftrightarrow$ Binned parametric space

$$x_{\text{pred},u/v} \xrightarrow{x_{\text{pred},u/v} - y \times \tan\alpha} x_{\text{meas},u/v}$$

3. Identification of local maxima (stereo track primitives)
   - **Maximum** above **threshold** in 1D histogram

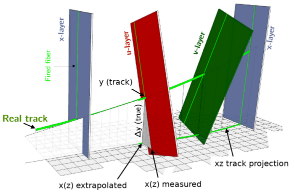

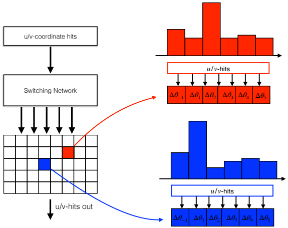

Ghost removal with stereo track fit

- Linearised χ^2 fit for false maxima removal
- Straight line: $y(z) = a_y + b_y \times z$
- For each local maximum determine best fit over combinations of
  - 5 different stereo layers out of 6
  - 1 out of all candidate hits on each layer
- 3D track primitives filtered with (χ^2_A, χ^2_S) requirement
  - Linear cut for illustration of performance

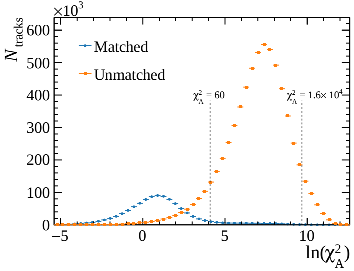

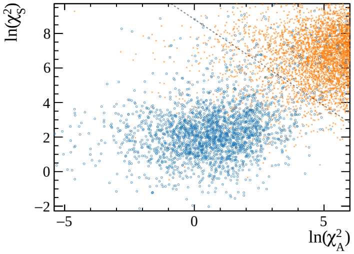

Physical implementation

- Input from detector over multiple lines
- Distribution network
  - Switch: routes hits only to appropriate cells
  - Optical communication: exchanges hits between boards
- Cell
  - Engine: computes and accumulates hit weights
  - Max-finder: finds tracks (local maxima)
- Primitive tracks are sent to the Event Builder

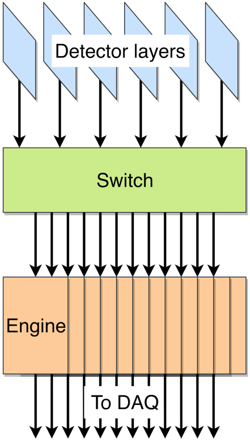

Modular design

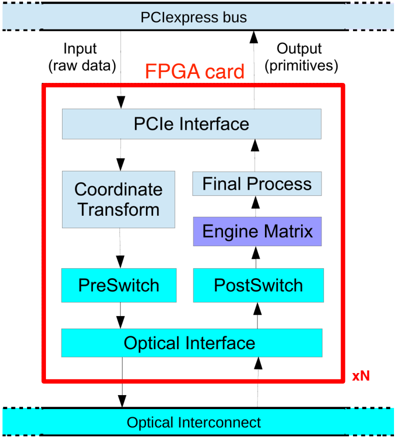

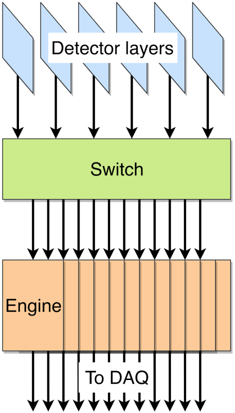

- Efficiency above 90% for high-momentum tracks
- Good efficiency for low-momentum ($p < 5\,\mathrm{GeV}$) tracks
  - Essential for downstream tracks (K^0_S and Λ)

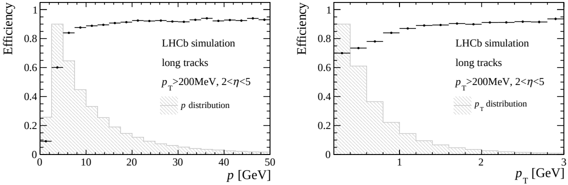

- Robust scaling with occupancy

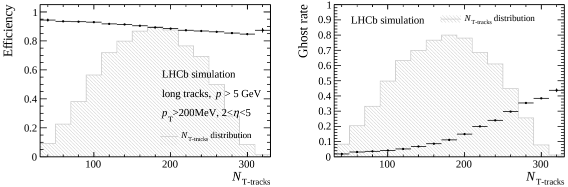

Throughput gain with DoWnstream Tracker

- HLT1 sequence `hlt1_pp_matching`

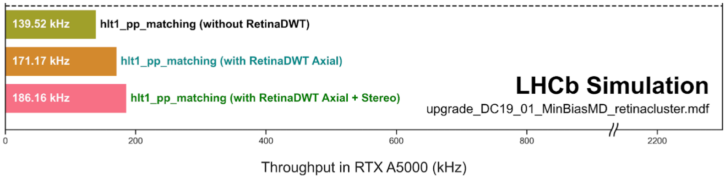

- Default sequence
  - Total (T-track reconstruction): $7.2\,\mu s$ ($1.5\,\mu s$)
- With T-track primitives from DWT
  - Total (Primitives decoding and refitting): $5.4\,\mu s$ ($0.06\,\mu s$)
- Throughput increased by a factor of **1.33**

Throughput scalability

- Event rate scales linearly with instantaneous luminosities
  - Luminosities obtained by merging events
- Event rate scales linearly with system size

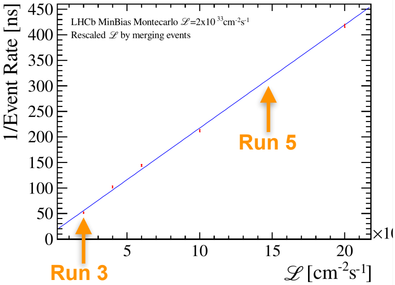

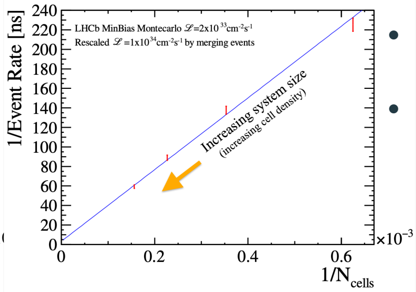

Resources and integration in LHCb Run 4 DAQ

- Number of FPGAs: 64 (axial) + 32 (stereo)
- DWT Boxes (up to 6 FPGA each) connected to SciFi EB nodes
- Modular, scalable, and minimal disturbance to current DAQ

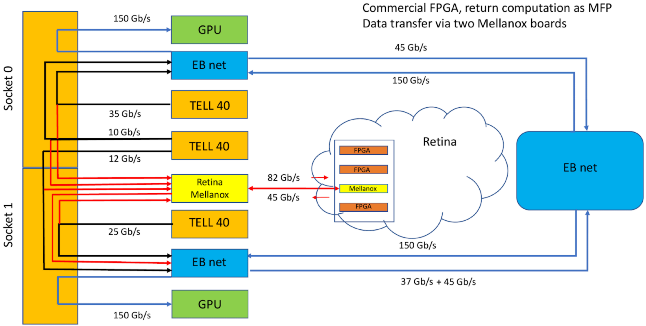

Hardware demonstrator with live LHCb Run 3 data

- A complete demonstrator installed and currently Running parasitically at the LHCb TestBed facility (Point 8)
- Smooth long-term operation without errors
- Tracks from demonstrator show very similar distribution to HLT2 reconstruction output

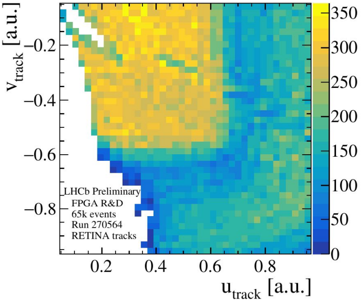

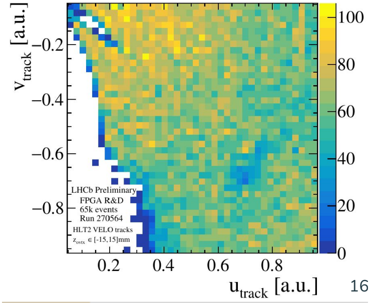

Summary and outlook

- Two major efforts of FPGA acceleration at LHCb
  - VELO cluster finding: established as the default method in Run 3
  - DoWnstream Tracker: approved as part of LHCb DAQ Enhancement in Run 4
- Increase throughput and decrease power consumption
  - Performance as good as software algorithms
- R&D ongoing for LHCb Run 5

Backup

LHCb RETINA team

Wander Baldini[1,2], Giovanni Bassi[3,4], Andrea Contu[5], Riccardo Fantechi[3], Jibo He[6,7], Brij Kishor Jashal[8], Sofia Kotriakhova[1,9], Federico Lazzari[3,10], Maurizio Martinelli[11,12], Diego Mendoza[8], Michael J. Morello[3,4], Arantza De Oyanguren Campos[8], Lorenzo Pica[3,4], Giovanni Punzi[3,10], Qi Shi[6], Francesco Terzuoli[3,13], Giulia Tuci[14], Ao Xu[3], Jiahui Zhuo[8]

[1] *INFN Sezione di Ferrara, Ferrara, Italy*
[2] *European Organization for Nuclear Research (CERN), Geneva, Switzerland*
[3] *INFN Sezione di Pisa, Pisa, Italy*
[4] *Scuola Normale Superiore, Pisa, Italy*
[5] *INFN Sezione di Cagliari, Monserrato, Italy*
[6] *University of Chinese Academy of Sciences, Beijing, China*
[7] *Hangzhou Institute for Advanced Study, UCAS, Hangzhou, China*
[8] *Instituto de Fisica Corpuscular, Centro Mixto Universidad de Valencia - CSIC, Valencia, Spain*
[9] *Università di Ferrara, Ferrara, Italy*
[10] *Università di Pisa, Pisa, Italy*
[11] *INFN Sezione di Milano-Bicocca, Milano, Italy*
[12] *Università di Milano Bicocca, Milano, Italy*
[13] *Università di Siena, Siena, Italy*
[14] *Physikalisches Institut, Ruprecht-Karls-Universitat Heidelberg, Heidelberg, Germany*

Performance with simulation

- LHCb simulation with Run 3-4 condition
  - $\sqrt{s} = 14\,\text{TeV}$, bunch $25\,\text{ns}$, $\nu = 7.6$
  - Samples: Minimum Bias, $D^0 \to K^0_S \pi^+ \pi^-$, $B^0_s \to \phi\phi$

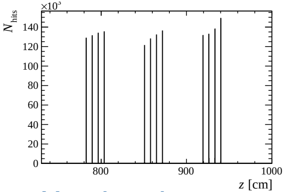

Hit distribution z

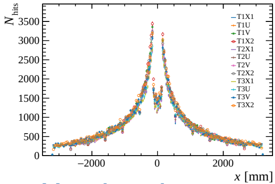

Hit distribution x

- DoWnstream Tracker emulator
  - C++ software emulator of an FPGA-based system for reconstruction of T-track primitives
  - Use integers to emulate the firmware implementation at **bit-level**
- Indicators: efficiency and ghost rate

Performance of T-track at primitive level

- Fiducial requirements: $p_{\mathrm{T}} > 200\,\mathrm{MeV}$ and $2 < \eta < 5$
- Efficiencies **comparable** with GPU-HLT1 and CPU-HLT2 Seeding
  - Higher efficiencies could be achieved with looser χ^2 requirements
- Ghost rate is **under control**
  - As a reference: below 15% (6%) for GPU-HLT1 tracking

Track type	MinBias	$D^0 \to K^0_{\mathrm{S}}\pi^+\pi^-$	$B^0_s \to \phi\phi$
Long, $p > 3\,\mathrm{GeV}/c$	85 (86)	83 (84)	84 (85)
Long, $p > 5\,\mathrm{GeV}/c$	**90 (91)**	**89 (90)**	**89 (89)**
Long from B not $e^\pm$, $p > 3\,\mathrm{GeV}/c$	-	-	88 (87)
Long from B not $e^\pm$, $p > 5\,\mathrm{GeV}/c$	-	-	90 (90)
Down, $p > 3\,\mathrm{GeV}/c$	84 (85)	83 (84)	83 (84)
Down, $p > 5\,\mathrm{GeV}/c$	**89 (91)**	**88 (89)**	**88 (89)**
Down from strange not $e^\pm$, $p > 3\,\mathrm{GeV}/c$	-	83 (83)	-
Down from strange not $e^\pm$, $p > 5\,\mathrm{GeV}/c$	-	88 (88)	-
Down from strange not long not $e^\pm$, $p > 3\,\mathrm{GeV}/c$	-	83 (83)	-
Down from strange not long not $e^\pm$, $p > 5\,\mathrm{GeV}/c$	-	88 (89)	-
ghost rate	**16 (10)**	**17 (12)**	**17 (13)**
ghost per real track	0.2 (0.1)	0.2 (0.1)	0.2 (0.1)

Event-averaged values shown in brakets

Physics performance of axial T-track primitives

- Working point set for $\varepsilon = 90\%$ **of long tracks with** $p > 5\,\mathrm{GeV}$
  - Number of pattern cells for SciFi: 2×73k
  - Efficiencies comparable with CPU-HLT2 Hybrid Seeding and GPU-HLT1 Seeding
  - Ghost rate about 35% (25%) $\Longrightarrow$ 0.5 (0.4) fake track for each real track
    - For reference 22% of (axial-only) GPU-HLT1

Track type	$\varepsilon(\mathrm{MinBias})$	$\varepsilon(D^0 \to K^0_\mathrm{S}\pi^+\pi^-)$	$\varepsilon(B^0_s \to \phi\phi)$ [%]
T-track, $p > 3\,\mathrm{GeV}$	83 (85)	82 (83)	83 (84)
T-track, $p > 5\,\mathrm{GeV}$	90 (91)	89 (90)	88 (89)
Long, $p > 3\,\mathrm{GeV}$	86 (87)	84 (85)	85 (86)
Long, $p > 5\,\mathrm{GeV}$	**91 (92)**	**90 (91)**	**89 (90)**
Long from B not $e^\pm$, $p > 3\,\mathrm{GeV}$	-	-	89 (88)
Long from B not $e^\pm$, $p > 5\,\mathrm{GeV}$	-	-	92 (91)
Down, $p > 3\,\mathrm{GeV}$	85 (86)	83 (84)	84 (85)
Down, $p > 5\,\mathrm{GeV}$	90 (91)	89 (90)	89 (90)
Down from strange not $e^\pm$, $p > 3\,\mathrm{GeV}$	-	83 (83)	-
Down from strange not $e^\pm$, $p > 5\,\mathrm{GeV}$	-	89 (89)	-
ghost rate [%]	32 (22)	35 (28)	35 (27)
ghost per real track	0.5 (0.3)	0.5 (0.4)	0.5 (0.4)

Event-averaged values are shown in parenthesis

Definition of efficiency and ghost rate

- Event-integrated quantity

$$\varepsilon \equiv \frac{\sum_i n^i_{\texttt{tracks,matched}}}{\sum_i n^i_{\texttt{tracks,reconstructible}}}$$

$$\texttt{ghost rate} \equiv \frac{\sum_i n^i_{\texttt{tracks,unmatched}}}{\sum_i n^i_{\texttt{tracks,reconstructed}}} = \sum_i \frac{n^i_{\texttt{tracks,reconstructed}}}{\sum_i n^i_{\texttt{tracks,reconstructed}}} \times \frac{n^i_{\texttt{tracks,unmatched}}}{n^i_{\texttt{tracks,reconstructed}}}$$

- Event-averaged quantity

$$\varepsilon \equiv \sum_i \frac{1}{N_{\texttt{evt}}} \times \frac{n^i_{\texttt{tracks,matched}}}{n^i_{\texttt{tracks,reconstructible}}}$$

$$\texttt{ghost rate} \equiv \sum_i \frac{1}{N_{\texttt{evt}}} \times \frac{n^i_{\texttt{tracks,unmatched}}}{n^i_{\texttt{tracks,reconstructed}}}$$

Physics performance (axial): effciency VS momentum

- Working point set for $\varepsilon = 90\%$ of long tracks with $p > 5\,\mathrm{GeV}$

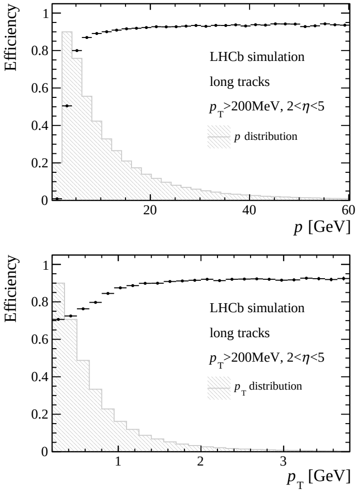

Physics performance (axial): efficiency VS η and ϕ

- Working point set for $\varepsilon = 90\%$ of long tracks with $p > 5\,\mathrm{GeV}$

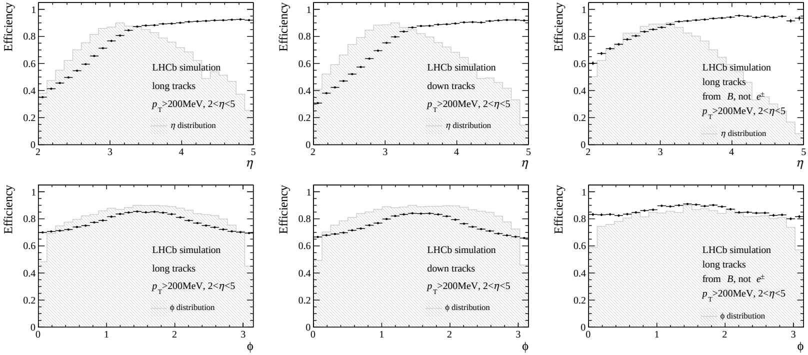

FPGA architecture

- Programmable logic blocks (CLB)
  linked by programmable interconnections (SM)
- Programmable Input and Output (I/O)
- More than a million of CLB and a thousand of I/O

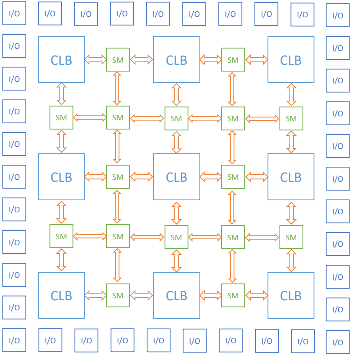

Track reconstruction: general concepts

- Track reconstruction of charged particles

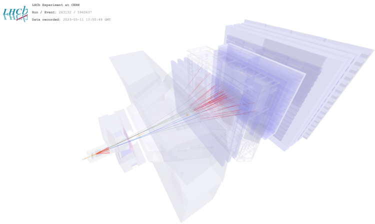

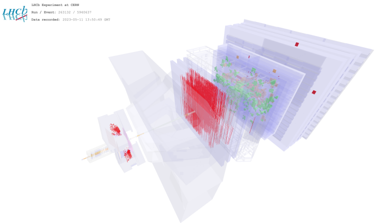

- Pattern recognition: partition of signals (detector measurements) into disjoint sets (track candidates)
  - Local method: select one track candidate at a time
  - **Global method**: all hits enter the algorithm in the same way
- Track fitting: determination of track parameters considering multiple scattering and energy loss

Hardware demonstrator

- All individual components tested in 10 years R&D
- A complete demonstrator installed and currently Running parasitically at the LHCb TestBed facility (Point 8)
  - Implemented in 8 PCIe-hosted FPGA cards
  - Reconstruct a VELO quadrant
  - Monitoring Farm provides to the TestBed facility events @ 1 kHz
  - Online LHCb alignment constant applied on the fly

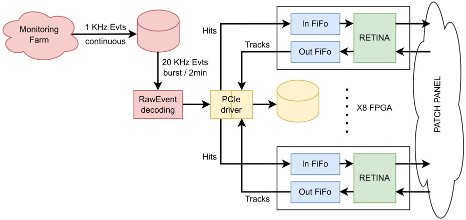

Milestones

- RTA proposal in the Framework TDR for the LHCb Upgrade II
- Dedicated workshop on 27/02/2023
- LHCb public note (LHCb-PUB-2024-001) submitted to U2PG and RTA on 24/10/2023
- LHCb internal U2PG review from 24/10/2023
- RTA endorsement on 30/11/2023
- LHCb TDR for LS2 enhancement submitter to LHCC on 27/02/2024