



Real-time tracking on FPGAs at LHCb

Giulia Tuci, on behalf of the LHCb-RTA project
giulia.tuci@cern.ch


Bari, 27/10/2022



中国科学院大学
University of Chinese Academy of Sciences

Introduction

- ❖ LHCb is a forward spectrometer designed to study b and c physics
- ❖ Instantaneous luminosity will reach $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ in Run 3
 - Tight p_T and E_T cuts saturate hadronic channels \rightarrow L0 hardware trigger has been removed
 - Software trigger will process events at the full LHC collision rate for an aggregate data flow of **40 Tb/s**



CERN-LHCC-2011-001

For further information on Run 3 reconstruction and trigger see talks by [A.Hennequin](#), [N.Schulte](#) and posters by [S.Akar](#) and [F.Reiss](#)


Heterogeneous computing at LHCb

- ❖ Heterogeneous computing allows event reconstruction in real time, from the the earliest stages of the processing, to the point of becoming embedded in the event building
- ❖ In Run 3, HLT1 trigger will run on GPUs
- ❖ Proposed upgrade for Run 5 (~2030) with another luminosity increase by a factor $\times(5\div 10)$
 - Exploring new solutions based on heterogeneous computing
 - LHCb established a coprocessor testbed to test them in realistic conditions



Real-time tracking on FPGAs


- ❖ Reconstruction of charged particle trajectories: large combinatorial problem that can be parallelized
- ❖ Reconstruction of downstream tracks is computationally expensive at HLT1
→ build a downstream tracking system on FPGAs



Downstream tracks at HLT1 needed to trigger on long-lived particles (K_S^0, Λ)

- ❖ Large and ambitious project for Run 4 → realize a smaller system to test the technology in Run 3
- ❖ **This talk:** demonstrator system for real-time tracking on FPGAs with the “artificial retina” architecture to reconstruct tracks in the Vertex Locator

The “artificial retina” architecture





Detector layers
Tracks

Parameter space divided into cells.

Each cell → pattern track



Receptors: interception of pattern tracks with detector layers.



Weights are assigned to hit-receptor distances.

For each cell all the weights are summed (excitation level).

$$R = \sum_{\text{all hits}} e^{-\frac{s_i^2}{2\sigma^2}}$$

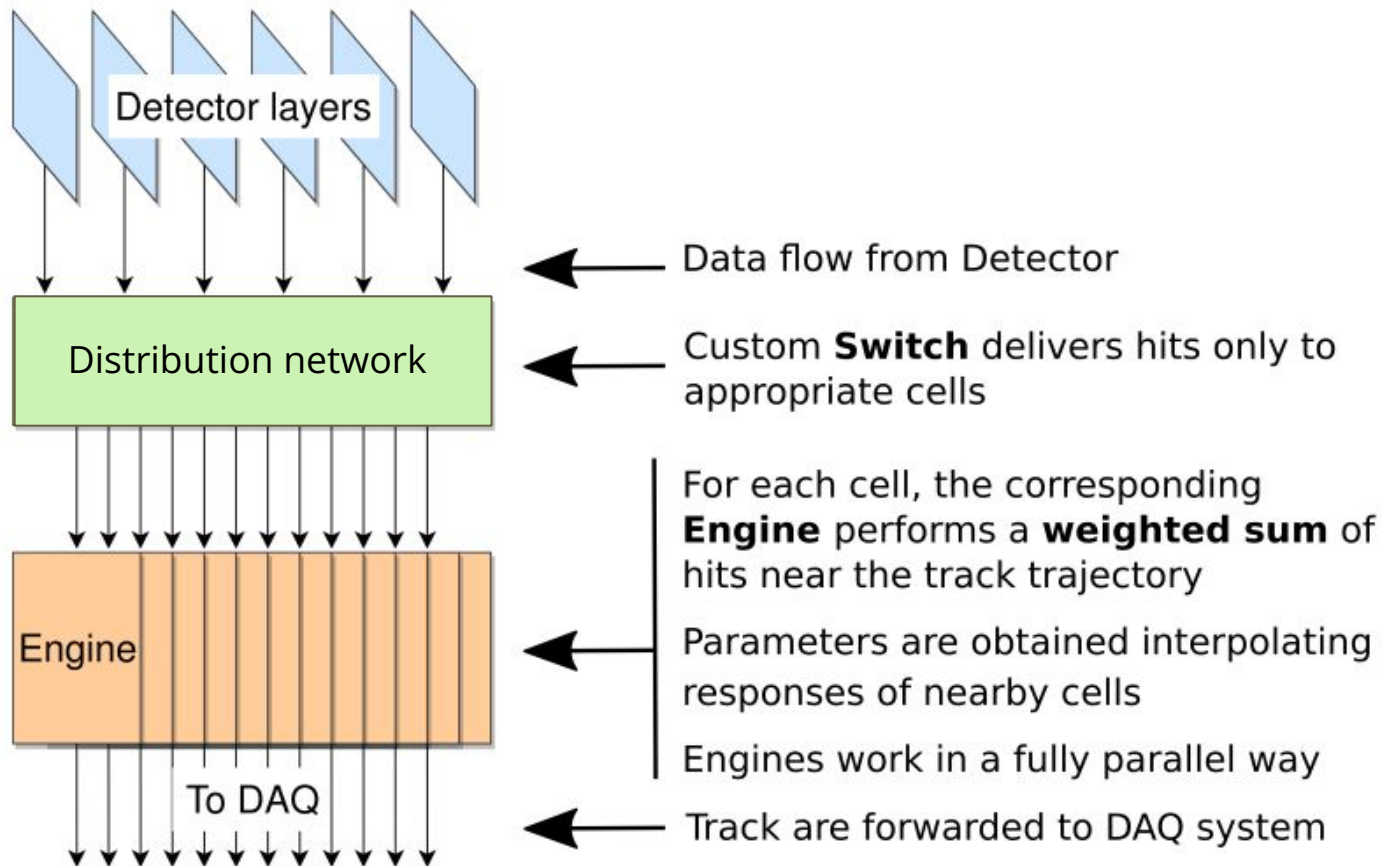


Search of local maxima over a threshold → reconstructed tracks

Real track parameters are obtained interpolating responses of nearby cells

Everything is executed in parallel!



The “artificial retina” architecture (2)



Using Lookup Tables, the distribution network delivers to each cell only hits close to the parametrized track, enabling large system throughput

Hits distribution

- ❖ To overcome FPGA size limitations without increasing latency, cells are spread over several chips
- ❖ Hits must be delivered only to the cells that need them → there can be more than one!
- ❖ Bandwidth profile increases in the distribution network, but output value lower than input after tracks are found





Prototype: track reconstruction in the LHCb VELO

- ❖ The Vertex Locator (VELO) is a crucial subdetector for LHCb physics program
- ❖ Composed of 52 silicon pixel modules, 38 in the forward region (< 10 % of LHCb data size)


PoS Vertex2019 (2020) 047
- ❖ Relatively compact FPGA system

EPJ Web. Conf. 245, 10001 (2020)

 - Good first test-case for future and larger-scale applications



Data coming from detector



Track-reconstruction on FPGAs



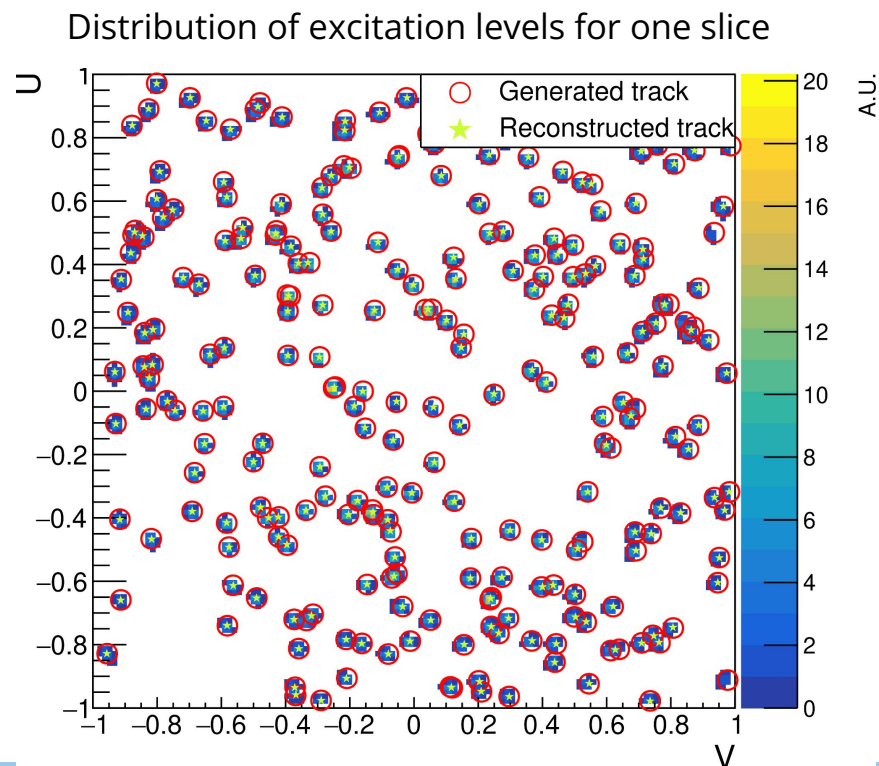

Track candidates + list of hits

Remaining HLT sequence on GPUs/CPU (e.g. fitting, selections)

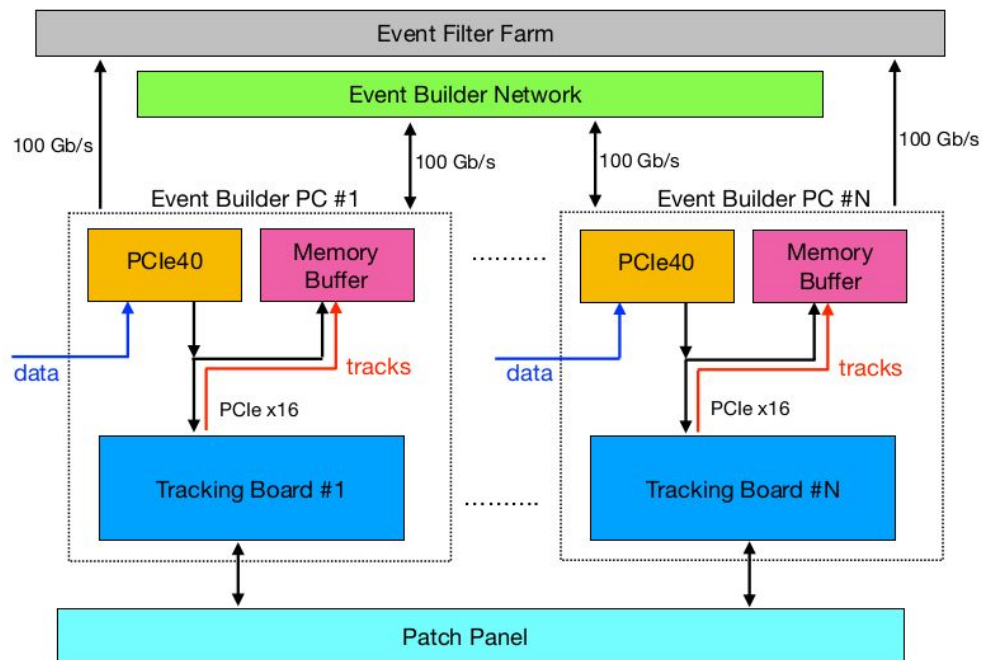
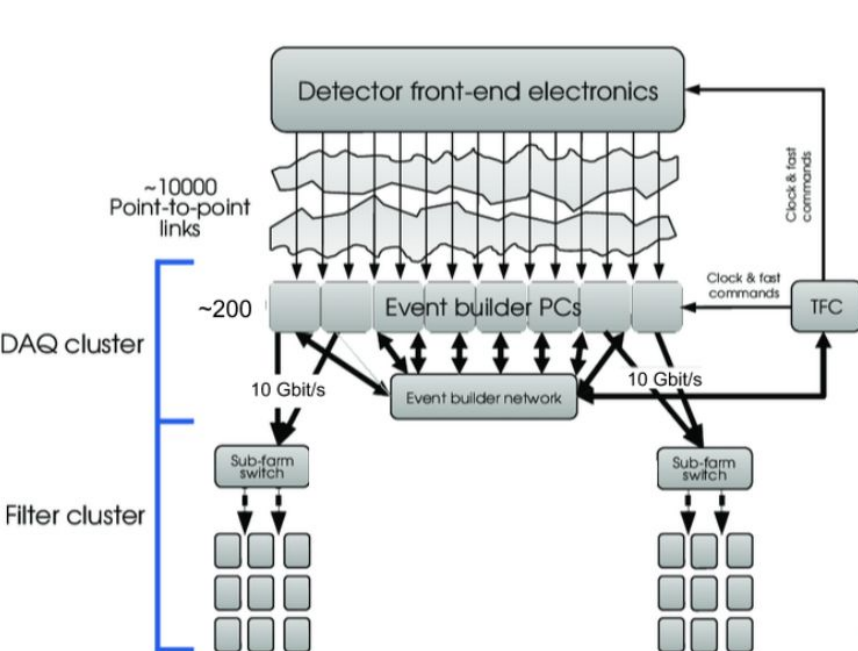


Segmentation of parameter space

- ❖ Algorithm optimized via C++ bit-by-bit emulator for FPGA pattern recognition
- ❖ Best configuration:
 - u, v are the x, y coordinates of intercept with a fixed plane
 - Algorithm based on a 2D space \rightarrow segment track space in 10 z -slices
 - 10000 cells for each slice (100k total cells)
 - Fits well FPGA logic available in current COTS PCIe cards.




Integration in the DAQ



- ❖ Event Builder collects tracks and performs building, treating the output of “retina system” like a virtual subdetector

- 1000 $B_s \rightarrow \Phi\Phi$ simulated events (Run 3 luminosity)




$$\epsilon = \frac{\# \text{ reconstructed tracks}}{\# \text{ reconstructible tracks}}$$

Reconstructible track: at least 3 hits in different VELO modules


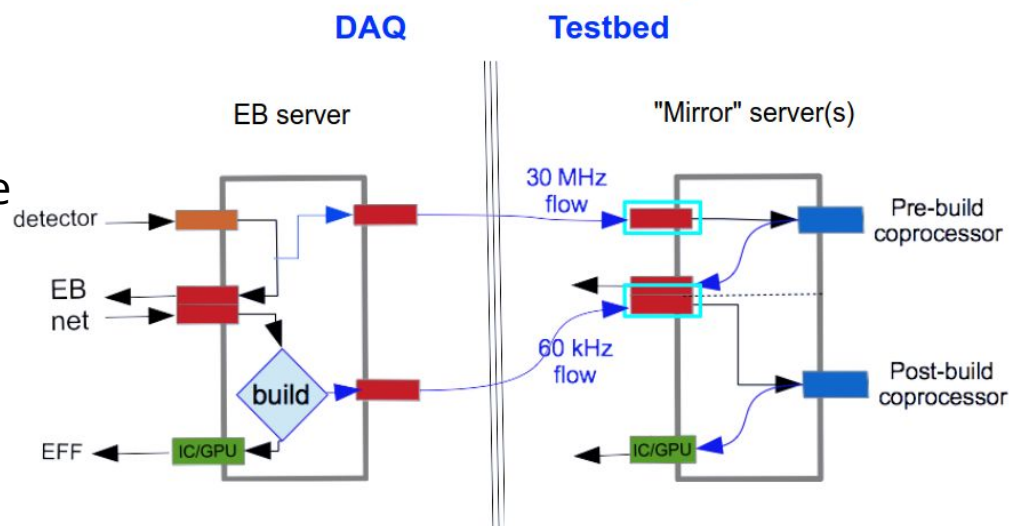
Long tracks \rightarrow crossing all tracking detectors

- Main limitation: number of cells should fit in available FPGA logic
 - Future advancements in technology \rightarrow price/performance drops

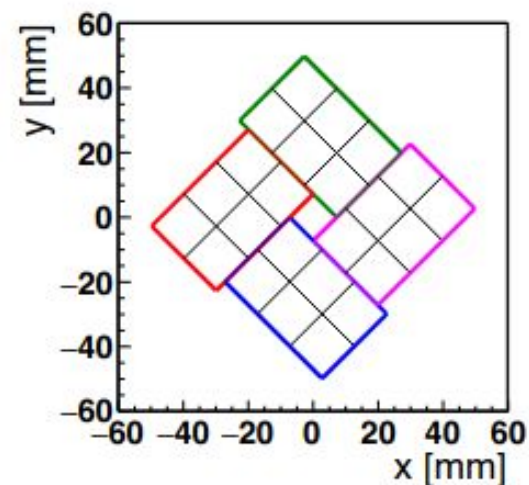


Implementation status @ coprocessor testbed

- ❖ Physically located in the main building of LHCb site
- ❖ Currently receiving live data from the monitoring farm (lower rate, post HLT1 processing)
- ❖ Unpack and buffer them to test our system with VELO hits
- ❖ Demonstrator system: 10 FPGA cards with Stratix 10 to reconstruct a VELO quadrant




PCIe 16x board, 1 Intel Stratix 10 FPGA, 16 optical links



Distribution network

- ❖ The track reconstruction is performed at pre-build stage
 - Receptors are spread over all the boards → hits need to be distributed
- ❖ Optical network connecting the boards is a key component of the retina system
 - It allows to exchange hits between FPGA boards



Test of the distribution network


- ❖ Test done using three Stratix 10 boards connected using two triangular full mesh networks by means of the optical patch panel

- ❖ Optical links at full speed:

- Link bandwidth: 25.8 Gbps
- Total bandwidth: 620 Gbps

- ❖ Test running without error for 1417 h 55 m 40 s

- $7.13 \cdot 10^{16}$ bit tested
(~10 PB)

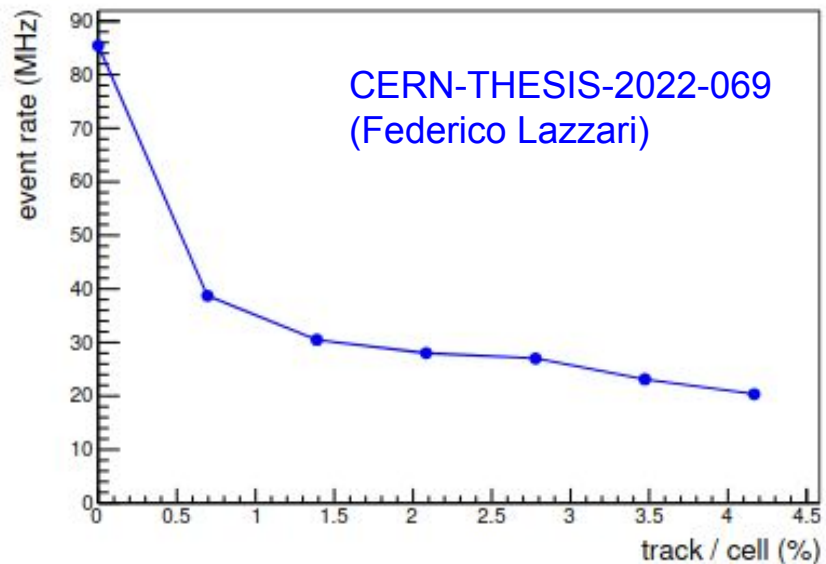
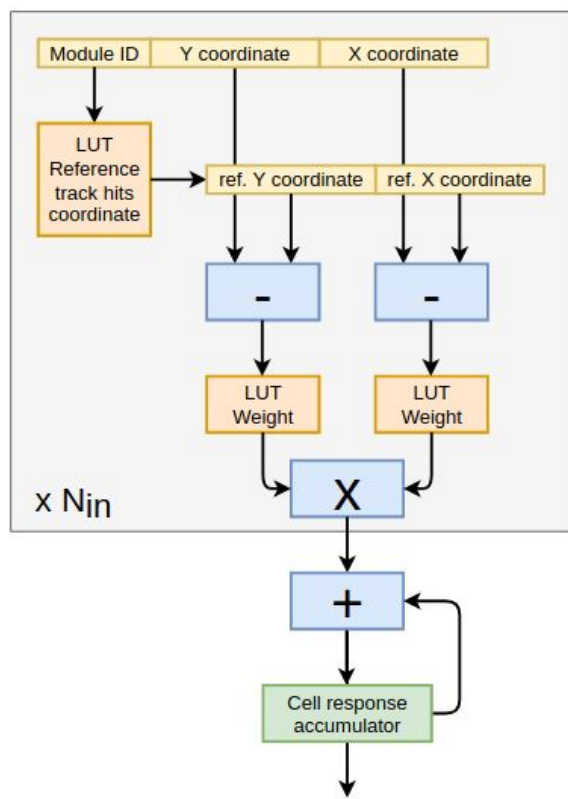


- ❖ Bit Error Ratio $< 4.21 \cdot 10^{-17}$ (CL = 95%)

- ❖ → Switch logic and optical communication successfully working!

Other ongoing tests

- ❖ PCIe communication between boards and the server
- ❖ Engine logic developed and tested with Questa Advanced Simulator



Final remarks


- ❖ HEP experiments will increasingly depend on large computing power
 - → a key to progress will be the capability of real-time, embedded reconstruction of tracks by specialized processors
- ❖ The artificial retina architecture could represent a viable solution
 - LHCb is investigating the possibility to realize a downstream tracker based on this in Run 4
- ❖ Ongoing implementation of a demonstrator system operating at coprocessor testbed at LHCb that will reconstruct $\frac{1}{4}$ of the pixel Vertex Locator

Backup slides

Inspiration from vision processing in the brain

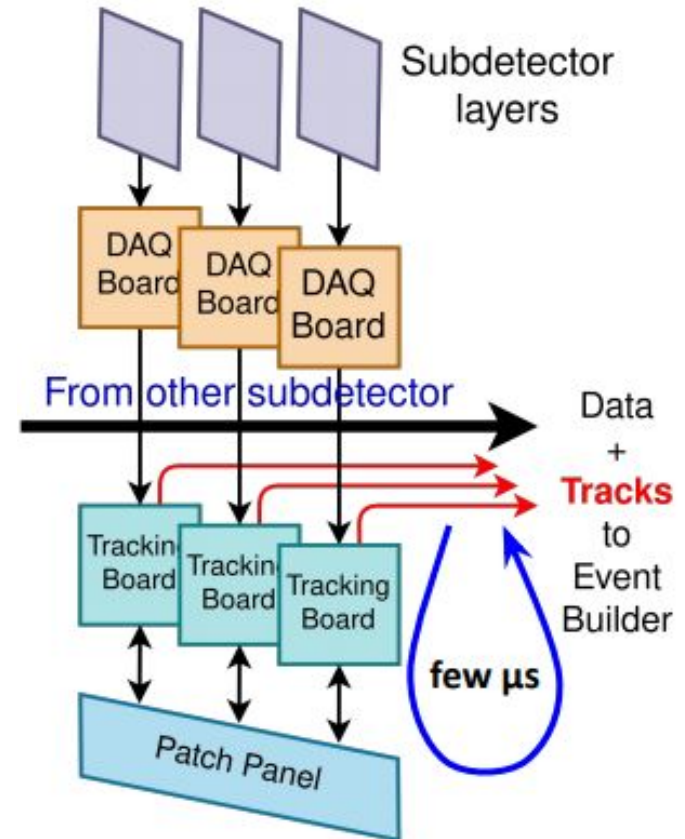
PloS one vol. 8,7 e69154

- ❖ Unconventional example: human brain
- ❖ Its early visual areas produce a recognizable sketch of the image in about 30 ms with a maximum neuron firing frequency of about 1 kHz → 30 clock cycles for image
- ❖ Similarities with HEP:
 - Lots of complex data
 - Same number of clock cycles available to process a LHC event with current electronics
 - Constrained computing resources



Why FPGAs?

- ❖ FPGA is the appropriate technology: aim for **high bandwidth and low latency**, comparable with that of other elements in the detector DAQ
 - Programmable and flexible devices
 - Low power consumption
- ❖ But tracks reconstruction requires to combine data from several different layers, typically read out separately by the DAQ
 - **Quick exchange of data between FPGA modules via fast optical network**



Performance (2)

LHCb-FIGURE-2019-011

- ❖ 1000 $B_s \rightarrow \Phi\Phi$ simulated events (Run 3 luminosity)

Track type	ϵ CPU pat-reco (%)	ϵ FPGA pat-reco (%)	
		all z	fiducial z-region
Long tracks with $p > 5$ GeV/c and hits in VELO > 5	99.84 ± 0.02	99.27 ± 0.06	99.45 ± 0.05
Long tracks from b with $p > 5$ GeV/c and hits in VELO > 5	99.61 ± 0.13	99.24 ± 0.21	99.41 ± 0.18
Long tracks from c with $p > 5$ GeV/c and hits in VELO > 5	99.89 ± 0.12	98.50 ± 0.53	98.62 ± 0.53

fiducial z-region: $-200\text{mm} < z < 200\text{mm}$

- ❖ **Comparison with standard CPU algorithm shows very close efficiency performance**