Demonstration of track reconstruction with FPGAs on live data at LHCb

Federico Lazzari on behalf of the LHCb-RTA project

26th International Conference on Computing in High-Energy and Nuclear Physics (CHEP2023) – May 8-12

Event reconstruction at LHCb

- LHCb is a forward spectrometer specialized to study c- and b-physics.
- Triggering on simple quantities is not possible due to high cross section of interesting events [1].
  → LHCb fully reconstructs the detector online,
  at the LHC average rate (~30 MHz).
- In Run 3 adopted a heterogeneous system:
  first level trigger (HLT1) on GPUs, second level on CPUs.
- In Run 5 (2035) luminosity will be increased by a factor 5÷10.
  → Computational power will increase by a factor 25 ($\propto n^2$).
- LHCb established a coprocessor testbed for:
  - Testing new heterogeneous computing solutions.
  - Providing realistic conditions.
- We developed a tracking system demonstrator for the LHCb Vertex Locator based on FPGAs.

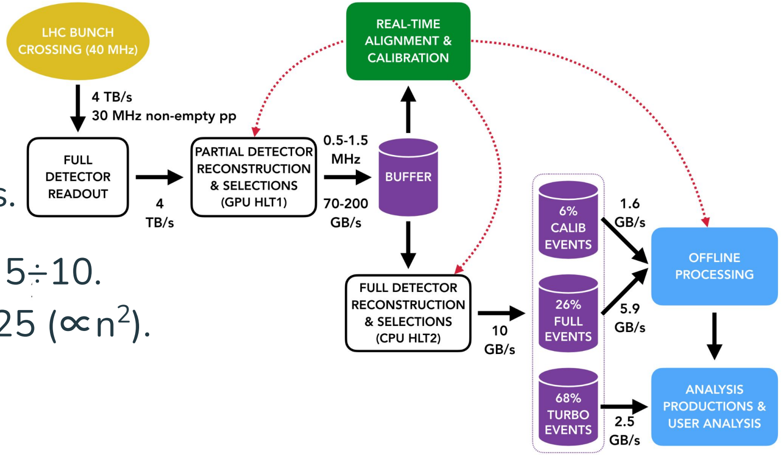

Why use FPGAs

- Modern FPGAs can perform highly parallel processing, with high throughputs and low latencies.
- **This allows to develop a tracking system that operates at the very first level of processing on unbuilt data.**
    - Can instrument just the desired sub-detectors.
    - Tracks could replace hits data in real time.
      → reduce data volume in Event Builder.
- Reconstructing tracks requires to combine data from several different layers, typically read out separately by the DAQ boards.
    - FPGAs have high-bandwidth transceivers (XCVRs) that allow to exchange information with great flexibility.
- The "Artificial Retina" is a highly-parallel architecture conceived for this scenario [4].

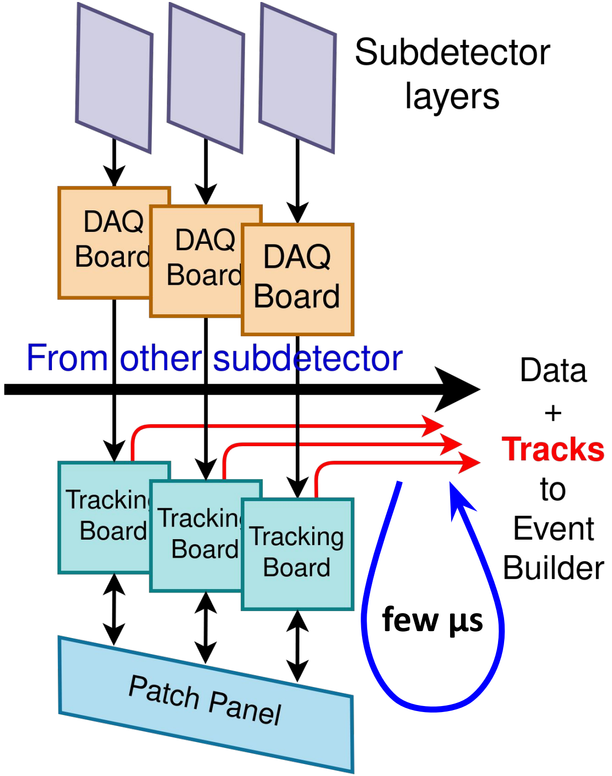

The “Artificial Retina” architecture

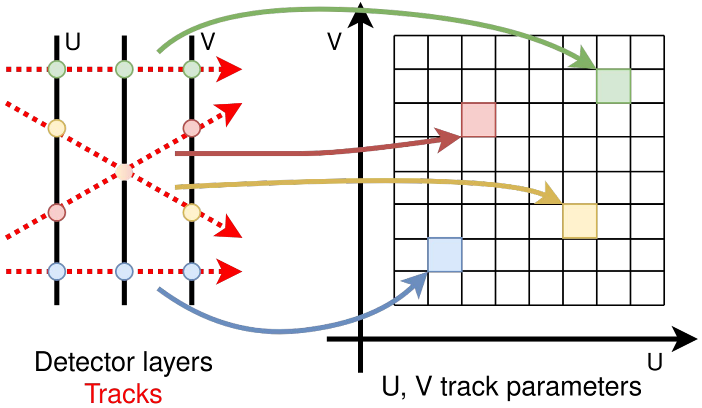

- Track parameter space represented in a matrix of processing units (cells).

The “Artificial Retina” architecture

Step 2: Accumulating weights (each cell)

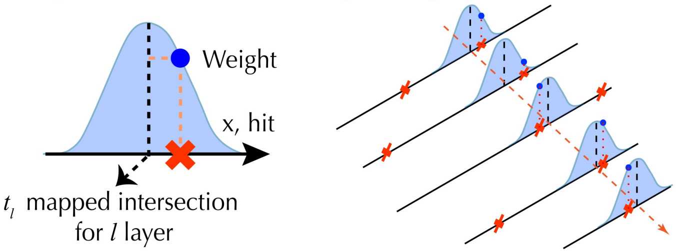

$$R = \sum^{\text{hits}} e^{-\frac{(x_l - t_l)^2}{2\sigma}}$$

R is close to N (# of layers) only if we have a set of hits near the mapped track

- Track parameter space represented in a matrix of processing units (cells).
- Each cell computes a weighted sum of hits near the reference track.

The “Artificial Retina” architecture

Step 3: Find the local maxima and compute centroid

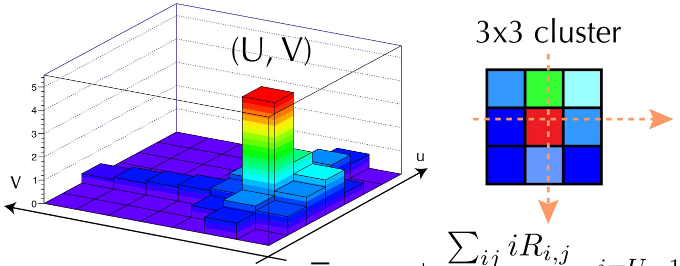

$$\bar{u} = u_0 + \frac{\sum_{ij} i R_{i,j}}{\sum_{ij} R_{i,j}} \qquad \bar{v} = v_0 + \frac{\sum_{ij} j R_{i,j}}{\sum_{ij} R_{i,j}}$$

$$i = U-1, U, U+1 \quad j = V-1, V, V+1$$

- Track parameter space represented in a matrix of processing units (cells).
- Each cell computes a weighted sum of hits near the reference track.
- Local maxima in the matrix of cells response correspond to reconstructed tracks.
- Real track parameters are obtained interpolating responses of nearby cells.

The “Artificial Retina” architecture

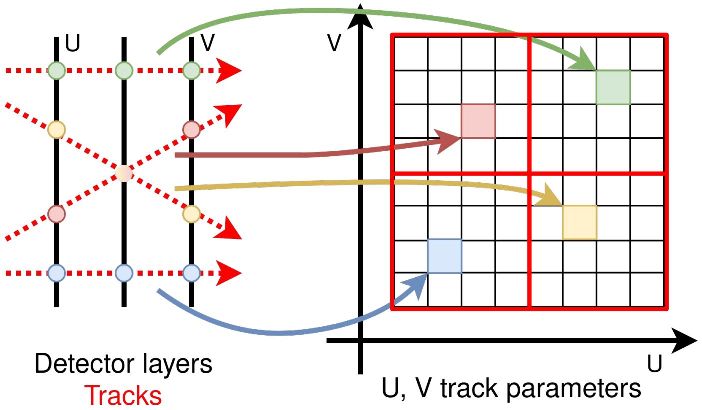

- Track parameter space represented in a matrix of processing units (cells).
- Each cell computes a weighted sum of hits near the reference track.
- Local maxima in the matrix of cells response correspond to reconstructed tracks.
- Real track parameters are obtained interpolating responses of nearby cells.
- To reach high-throughput and low-latencies, cells work in a fully parallel way.
- To overcome FPGA size limitations without increasing latency, cells are spread over several chips.

The system's implementation

- Input from detector.
- Data preparation (detector specific).
  - e.g. Hit clustering
- Distribution network:
  - Switch: routes hits only to specific cells.
  - Optical communication: exchanges hits between tracking boards.
- Cell:
  - Engine: computes and accumulates weights.
  - Max-finder: finds local maxima and interpolates responses.
- Output to Event Builder.

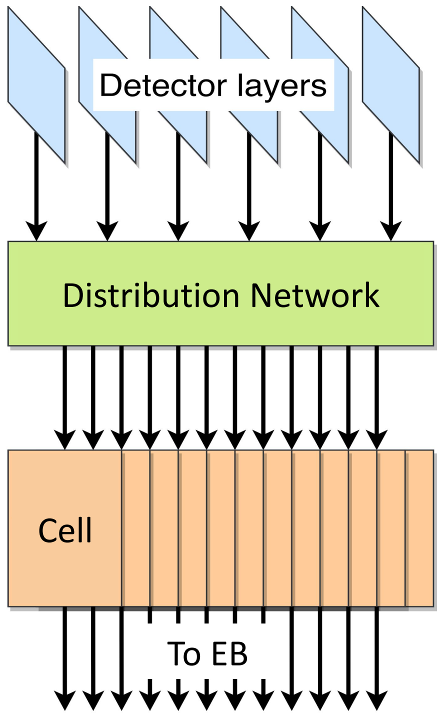

The VELO detector

- Crucial subdetector for LHCb physics program.

- 38 modules in the forward region (< 10% of LHCb data size).
- Time consuming (25% of HLT1 time budget [2]).
  → A relatively compact FPGA system can save
  a large fraction of HLT1 time.

- Physics performances study of FPGA tracker system available [3].

- A good test-case for future and larger-scale applications.

- Hit clustering (data preparation) currently adopted in Run 3:
  - Architecture originated from “Artificial Retina”.
  - Integrated in DAQ boards.
  - See G. Bassi talk May 11, 12:00 PM, track 11.

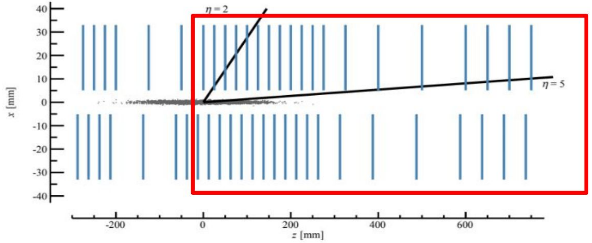

VELO top view

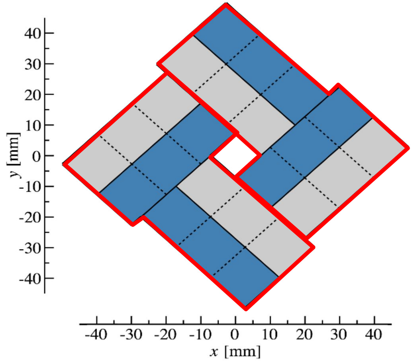

VELO front view

The Demonstrator

- “Artificial Retina” demonstrator installed at the LHCb site.
- Receiving live data from the LHCb monitoring farm.
- Demonstration that this architecture is ready to be adopted in HEP experiments:
    - Reconstructing an actual detector (VELO quadrant).
    - Working in real-time.
    - At nominal LHCb luminosity ($2x10^{33}$ $cm^{-2}s^{-1}$).
    - Integrated in the LHCb DAQ environment.
- Spread over multiple PCIe-hosted FPGA cards.
- Scalable to cover the whole detector adding further FPGA cards.

Design overview

- 8 FPGA boards.
- Clock speed: 250 MHz.
- Data buffered in LoopRAMs (2 VELO modules per board).
    - Data generated with LHCb simulation at LHCb luminosity.
    - During Run 3 operations data will be feed on-line from the detector monitoring farm.
- Data exchanged by distribution network (Switch + optical communication).
- Engines reconstruct tracks.
    - Engines implemented in different boards cover different regions on track parameters space.
- Output is prescaled and compared bitwise to C++ simulation.

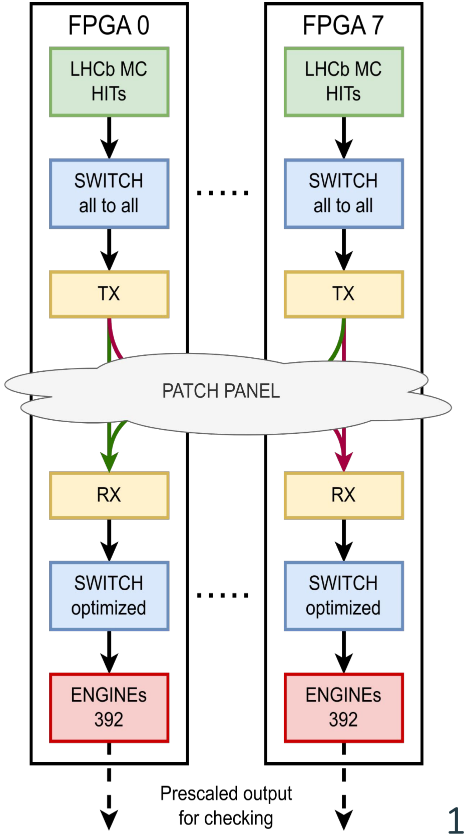

Distribution network

- Implements a portion of the **whole VELO distribution network**.
- Implements one stage of the optical communication:
  - 8 nodes full-mesh network.
  - 28 full-duplex links at 25.8 Gbps.
  - Total bandwidth 1.41 Tbps.

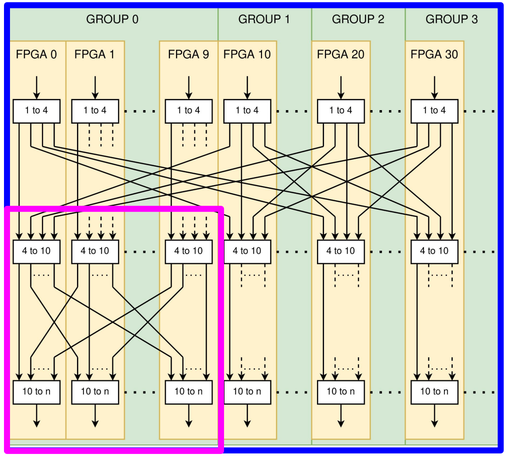

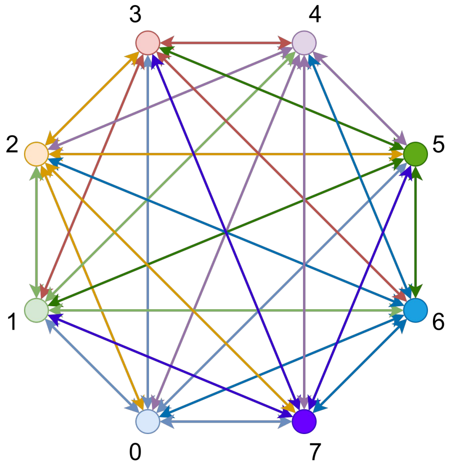

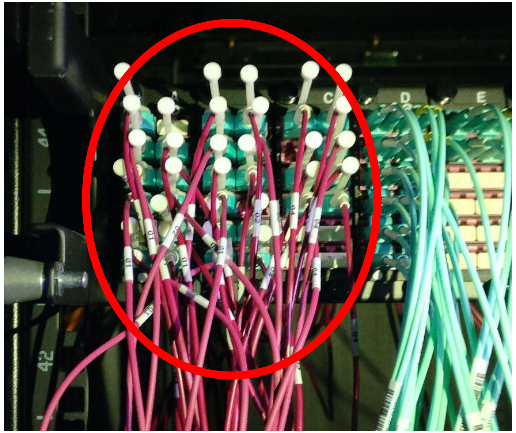

Engines

- Accepts 2D hits (VELO hits).
- Multiple inputs (N_{in} = 4) for accepting up to 4 hits per clock cycle.
- Cover a **quadrant** of the **track parameters space**.

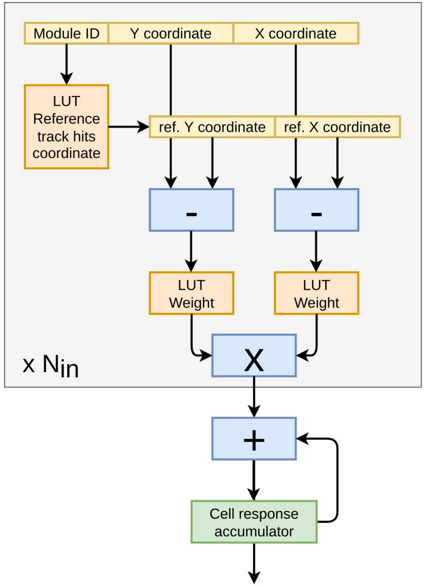

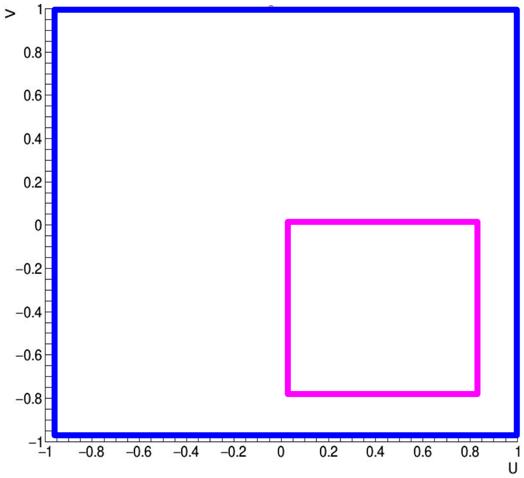

Results

- Runned several days (~10) without hiccup (much higher than typical bunch life).
- Every tracks correspond bitwise to the tracks reconstructed by the C++ simulation.
- Every tracks reconstructed by the C++ simulation is produced by the Demonstrator.
- **Currently running at 16.2 MHz.**

- Now working on optimizing a few points of the actual design.
- Expect to achieve a further factor 2x in throughput, yielding the desidered speed of 30 MHz.

- Proposal for construction of a real tracking system for Run 4 is currently under review by LHCb collaboration.

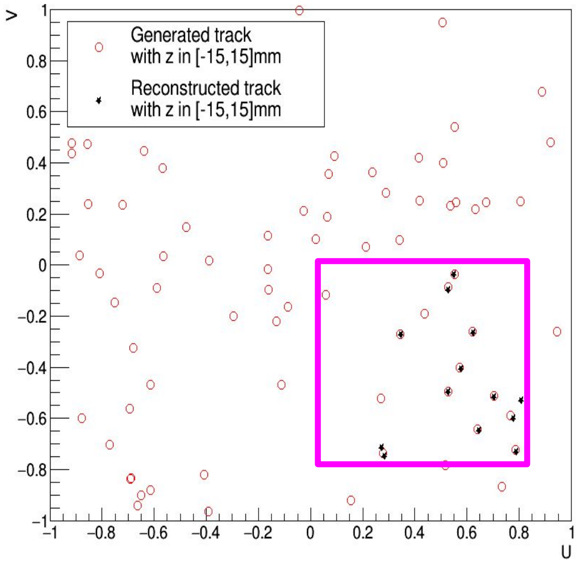

Implications

- The “Artificial Retina” is a highly-parallel tracking system with good physics performances.
- We developed a demonstrator for the VELO detector at LHCb.
  - Currently tuning final detail to reach the design throughput.
- This demonstrator is the cornerstone for the realization of a tracking system based on FPGA.
- Construction proposal of the Downstream Tracker (DWT) for Run 4 under review [5 6].
  - Tracks generated outside the VELO (mostly daughters of long lived particles).
  - Not reconstructed by HLT1 due to computational time limits.
  - With DWT, LHCb can improve its physics program.

More details in Xavier Vilasís Cardona's talk.

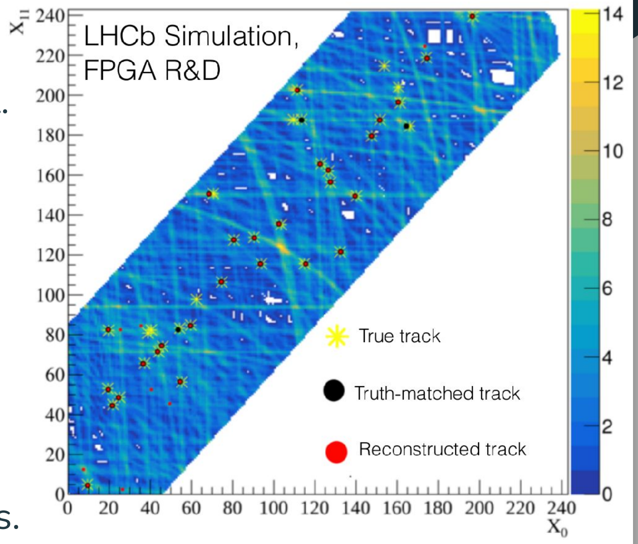

Backup

Bibliography

1. LHCb Collaboration, **LHCb Trigger and Online Upgrade Technical Design Report**, CERN-LHCC-2014-016
2. LHCb Collaboration, ***LHCb Upgrade GPU High Level Trigger Technical Design Report***, CERN-LHCC-2020-006
3. G. Tuci, ***Reconstruction of track candidates at the LHC crossing rate using FPGAs***, CHEP 2019
4. G. Punzi et al. on behalf of the LHCb Real-Time Analysis project, ***Real-time reconstruction of pixel vertex detectors with FPGAs***, PoS(Vertex2019) - Tracking and vertexing
5. LHCb Collaboration, ***Expression of Interest for a Phase-II LHCb Upgrade: Opportunities in flavour physics,*** and beyond, in the HL-LHC era, CERN-LHCC-2017-003
6. M.J. Morello et al., ***Real-time reconstruction of long-lived particles at LHCb using FPGAs***, ACAT 2019

Hardware

- Prototyping board,
  2 Intel Stratix V FPGAs,
  96 optical links

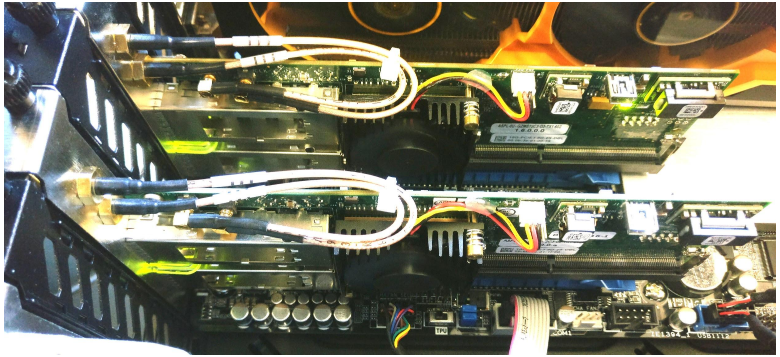

- PCIe 8x board, 1 Intel Arria V GX FPGA, 8 optical links

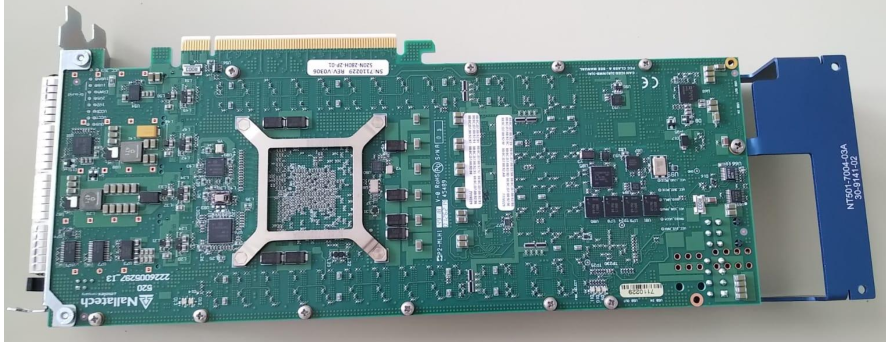

- PCIe 16x board, 1 Intel Stratix 10 FPGA, 16 optical links

Physics performances

- Studies done with C++ simulation of the “Artificial Retina” architecture.
- *Bs*$\rightarrow$*ΦΦ* sample from official LHCb full simulation.
- Inject result into LHCb’s track performance benchmark.
- Comparison with standard CPU algorithm shows very close efficiency performance on fiducial tracks (-200mm < z < 200mm) [3].

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

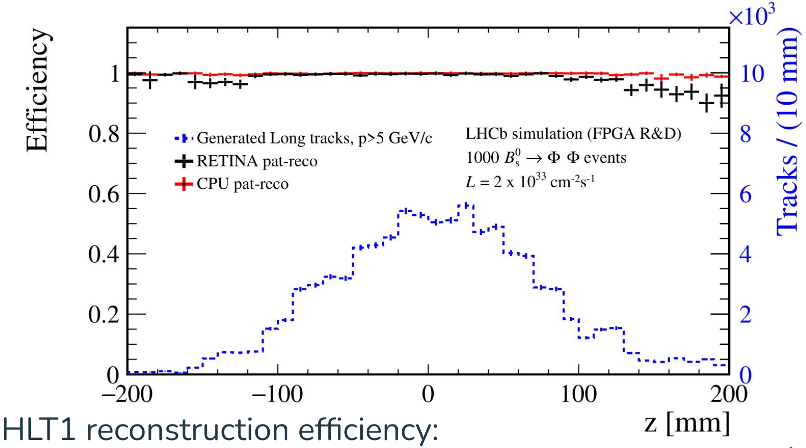

HLT1 reconstruction efficiency:
comparison between CPU and “Artificial Retina”.

Integration in DAQ system

- The "Artificial Retina" could find a place in the Event Builder nodes using PCIe boards.
- The Event Builder collects the tracks and performs the building, treating the "Artificial Retina" like a virtual sub-detector.

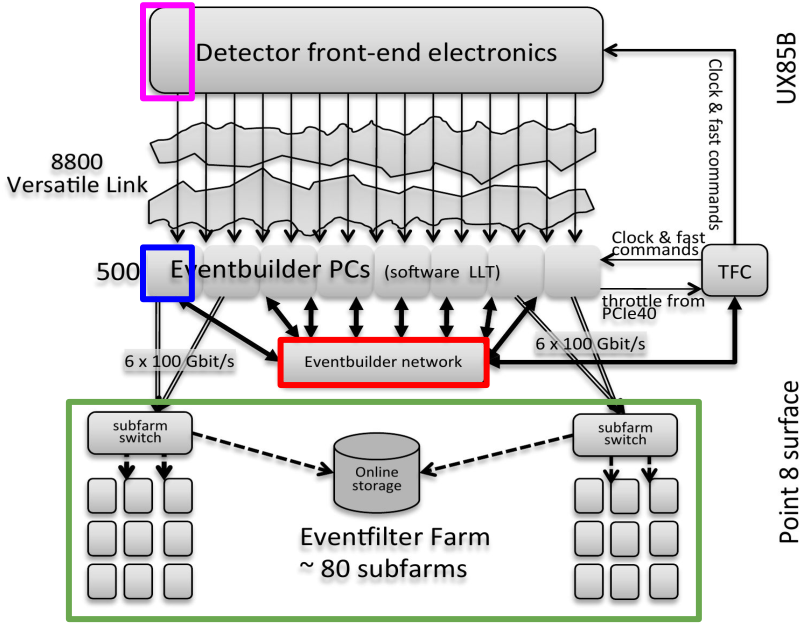

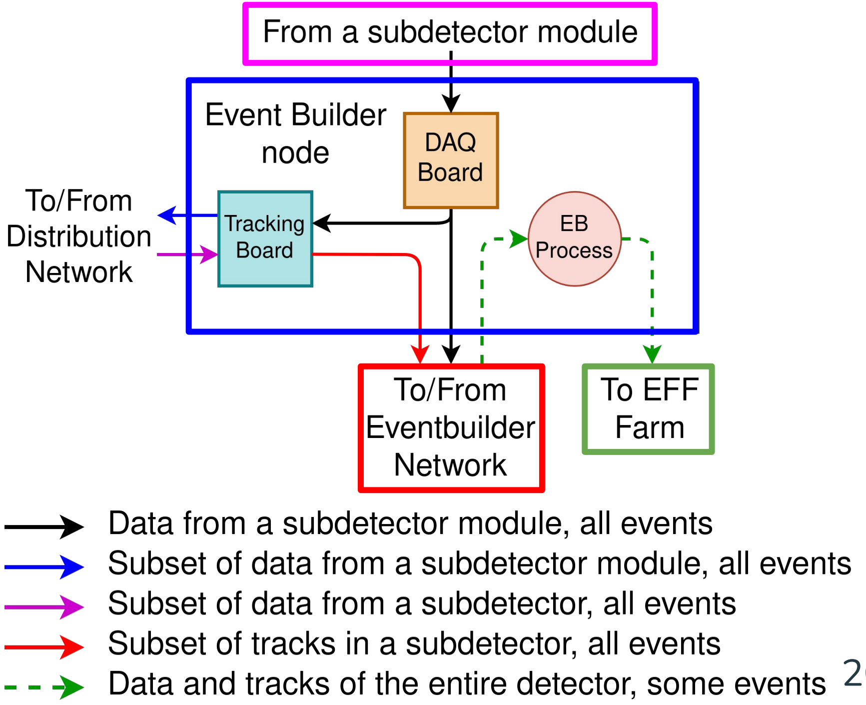

The firmware paradigms

Pipeline:

- Like an assembly line, an event is processed as soon as possible, without waiting for the previous one to go through all the steps.
- This paradigm is extended to the hit level → 1 hit/clk cycle.

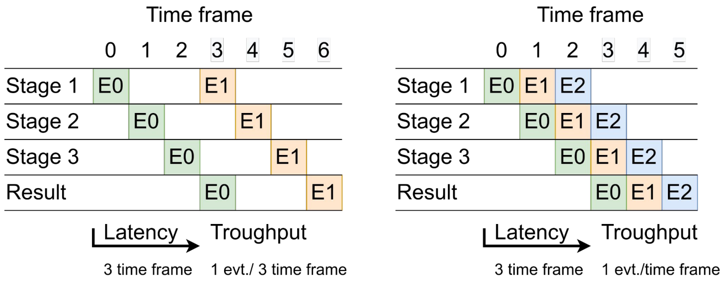

Parallel computing:

- Hits flow through the distribution network via parallel lines.
- Cells work in a fully parallel way (both weight accumulation and maxima finding).
- Cells have also parallel inputs to process more hit per clock cycle.
- A bigger system has more parallel processor, so its throughput is similar to the one of a small system.

Modularity:

- Each component (switch, matrix of cell, ecc.) is a repetition of basic blocks.
- A bigger system is implemented instantiating more copies of the same modules.
- Modules can be freely spread over multiple devices overcoming FPGA size limitation.

This is different from other systems that rely to time multiplexing.

The Distribution Network

- Hits are provided to different Tracking boards arranged by sub-detector DAQ board.
- A custom distribution network rearranges the hits by track parameters coordinates (similar to a “change of reference system”).
- Using Lookup Tables (LUTs), the Distribution Network delivers to each cell only hits close to the parametrized track, enabling large system throughput.
- The Distribution Network is a single entity transversal to all the Tracking boards.
- We designed a modular Distribution Network spread over the same array of FPGAs performing the tracking.

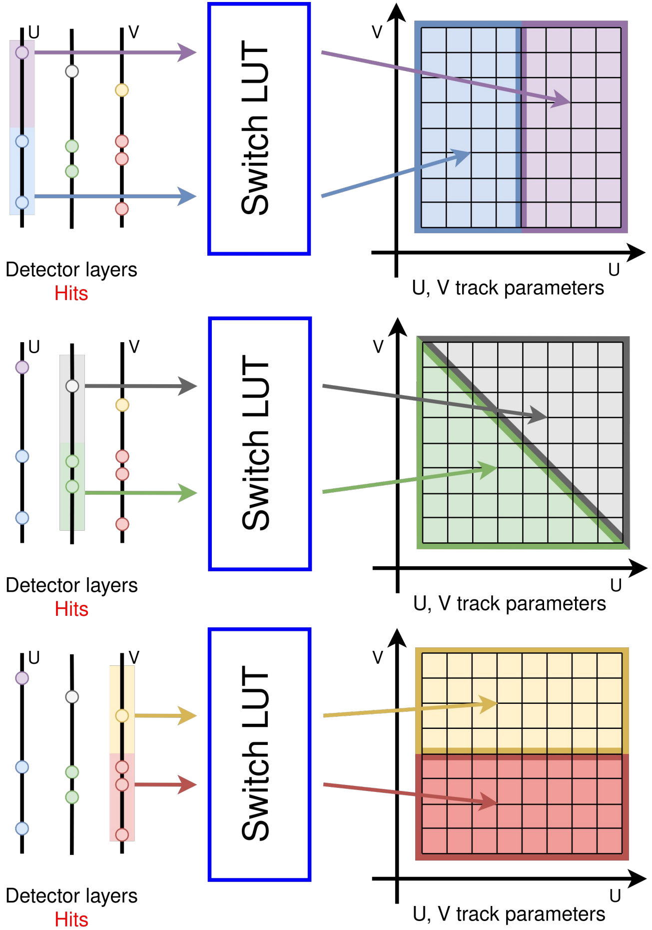

Switch

- 2-way dispatcher (2d): 2 splitters (1 input - 2 outputs) and 2 mergers (2 inputs - 1 output).
- Combining 2-way dispatchers is possible to build a switch with the desired number of lanes:
  - Switch with $N = 2^n$ lanes requires M 2-way dispatchers: $\begin{cases} M(0) = 0 \\ M(n) = 2M(n-1) + 2^{n-1} \end{cases}$
- We can implement any 2^n lanes switch changing a single parameter.

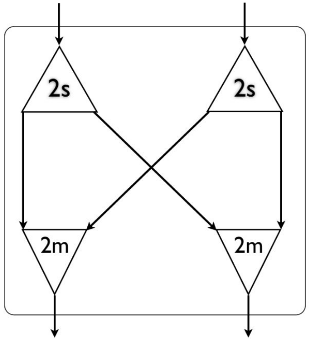

2-way dispatcher (2d)

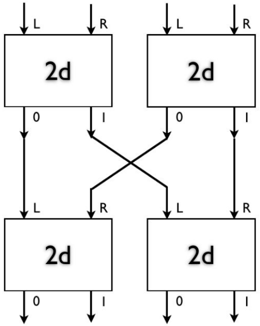

4-way dispatcher (4d)

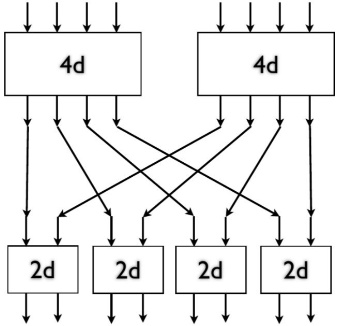

8-way dispatcher (8d)

Optical communication

- Uses Intel SuperLite II v4 communication protocol.
  - Fully free and available in source code.
  - Supports flow control.
  - Can be used to connect various FPGA families (already available on A10, S10, Agilex).
- Design adapted to implement the desired number of independent links.
- Extensively tested:
  - Long run: up to 2 months.
  - High-speed: up to 26 Gbps.
  - Multiple boards: up to 5 boards.
  - Large patch-panel: up to 64 links.