Real-time computing solutions based on FPGAs

Federico Lazzari Università di Siena & INFN sezione di Pisa

Real Time Analysis with FPGAs

- With the slowdown of Moore's law, the need for more effective, specialized computing solutions in HEP is increasing
- One promising solution is the use of FPGAs, that are a standard solution where large computing power is needed
- In Run 3 the LHCb detector will be read out and events fully reconstructed at the full LHC crossing rate, while planning to runs at even higher luminosities in the future
- Intense R&D is being performed on alternativesolutions to traditional general-purpose computing[1]

Extremely parallel FPGA tracking architecture yields a better behaviour with detector occupancy than CPU based algorithms[2,3]

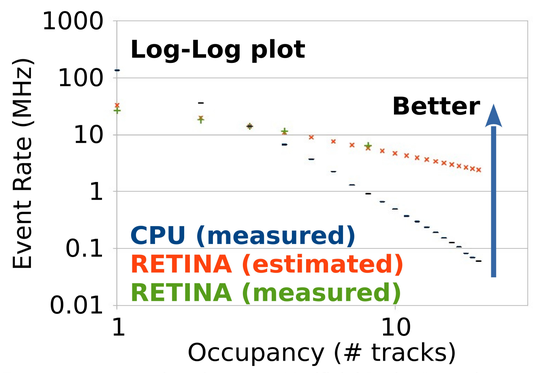

Event rate comparison between "Artificial Retina" running on FPGA and a CPU-based fitter for a generic 3 layers tracker

Clustering on FPGA[4,5]

- **Input**: position of active pixels inside the detector for a given event
- **Core idea**: clustering of separated areas of the detector **in parallel** by filling a set of **matrices** made of pixels. The matrix position inside the detector is set by the first pixel filling that matrix

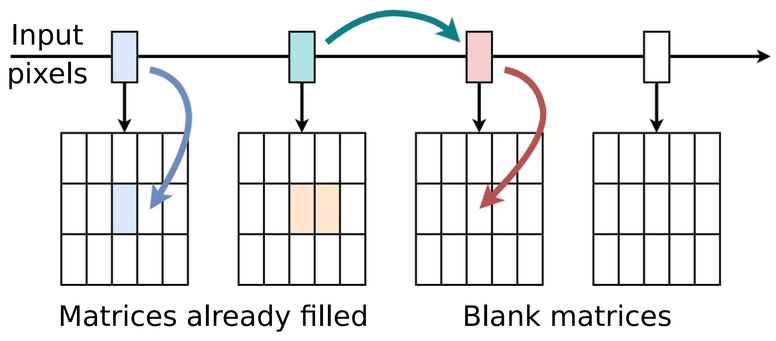

BLUE pixels belongs to the matrix, they fill the matrix

GREEN pixels do not belong to the matrix, they move forward

RED pixels do not belong to any previously filled matrix, they fill a blank matrix

- **Output**: centers of mass of the reconstructed clusters
- We designed and **fully tested** the FPGA clustering **firmware**, tailored for the LHCb **VELOPIX** detector, while comparing its output to an High-Level simulation
- The measured **event-rate** for the most occupied module is **38.9 MHz**
- The **tracking efficiencies** using FPGA clustering are basically **equivalent** to the CPU ones
(e.g. for long tracks CPU: 97.7 %, FPGA: 97.5 %)
- Clustering firmware was **integrated into VELO DAQ** firmware, using **~21 %** of available Logic Elements in the readout FPGAs

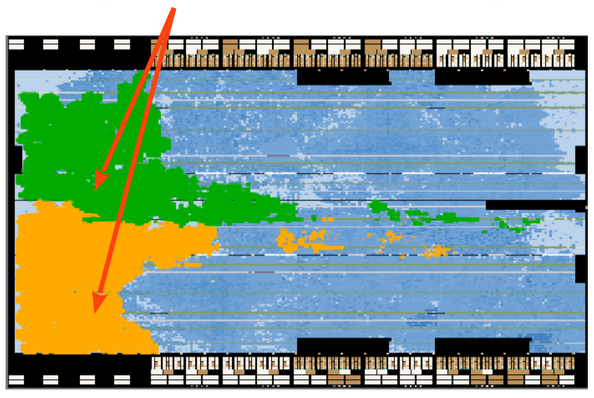

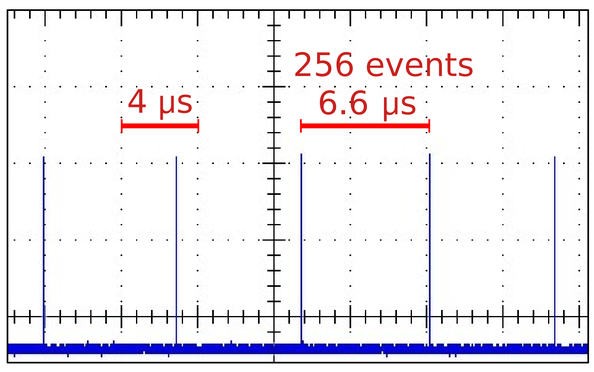

This R&D was performed within LHCb - Real Time Analysis project and INFN RETINA project whit fundings from INFN CSN1 and CSN5

[1] G. Punzi, F. Lazzari, G. Bassi, G. Tuci, M.J. Morello, R. Fantechi, S. Stracka, Real-time reconstruction of pixel vertex detectors with FPGAs, VERTEX 2019
[2] F. Lazzari, Development of a real-time tracking device for the LHCb Upgrade 1b, CERN-THESIS-2017-442
[3] R. Cenci, F. Lazzari, P. Marino, M.J. Morello, G. Punzi, L.F. Ristori, F. Spinella, S. Stracka, J. Walsh, Development of a High-Throughput Tracking Processor on FPGA Boards, PoS(TWEPP-17) 136
[4] F. Lazzari, G. Bassi, R. Cenci, M.J. Morello, G. Punzi, Real-time cluster finding for LHCb silicon pixelVELO detector using FPGA, ACAT 2019
[5] G. Bassi, R. Cenci, F. Lazzari, M.J. Morello, G. Punzi, A 2D FPGA-based clustering algorithm for the LHCb silicon pixel detector running at 30 Mhz, 2019 IEEE NSS-MIC
[6] G. Tuci, Reconstruction of track candidates at the LHC crossing rate using FPGAs, CHEP 2019
[7] Riccardo Cenci, Andrea Di Luca, Federico Lazzari, M.J. Morello, G. Punzi, Real-time reconstruction of long-lived particles atLHCb using FPGAs, ACAT 2019

Tracking on FPGA – The "Artificial Retina"

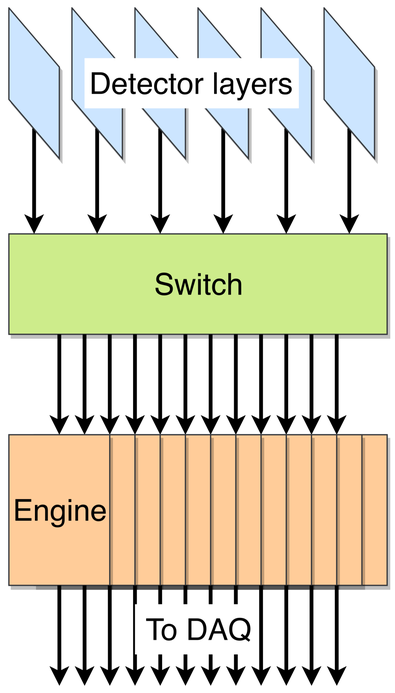

- Mathematically related to "Hough transform"
- Map tracks parameters space into a matrix of cells
- ← Tracks are forwarded to DAQ system
- ← Custom **Switch** delivers hits only to appropriate cells
- For each cell, the corresponding **Engine** performs a **weighted sum** of hits near the track trajectory
- Parameters are obtained interpolating responses of nearby cells
- Engine works in a **fully parallel** way
- ← Data flow from detector

- From a software simulation applied to the LHCb **VELO** pixel detector[6], the FPGA tracking **efficiencies** are **very close** to the efficiencies of LHCb baseline CPU algorithm

(e.g. for long tracks with P > 5 GeV/c and at least 5 hits on VELO CPU: 99.84 %, FPGA: 99.27 %)

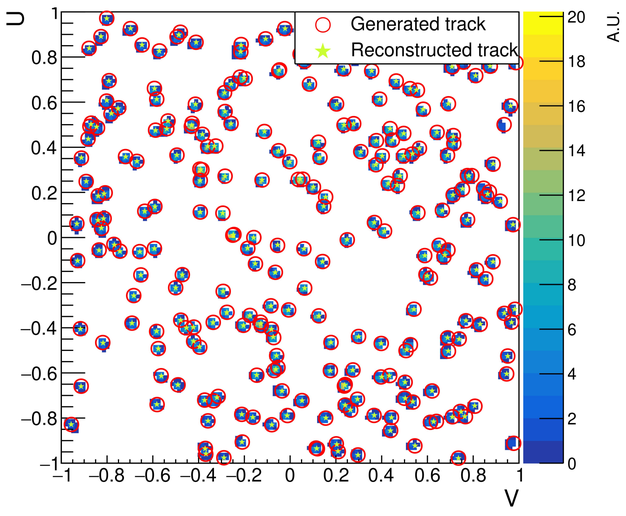

- We designed and tested the FPGA **firmware** for a 6-layers single-coordinate tracker like the axial layers of LHCb **SciFi** detector[3]
- We implemented an actual hardware prototype, and measured **event-rate > 30 Mhz** and **latencies < 500 ns**
- From a software simulation applied to the LHCb **SciFi**[4] detector, we estimate that **O(10²) current FPGAs** are sufficient for reconstructing tracks in the whole SciFi. The same task would need **O(10⁴) CPU** Xeon E5-2630 v4

The dedicated network

- Thanks to a dedicated switching network, the RETINA architecture can be distributed on several FPGAs O(40) that process a portion of the same event
- Modern FPGA boards have several high-speed transceivers with a total bandwidth of hundreds of Gbps (~400 Gbps full-duplex)
- At LHCb Vertical Slice (VST) we are performing extensive tests

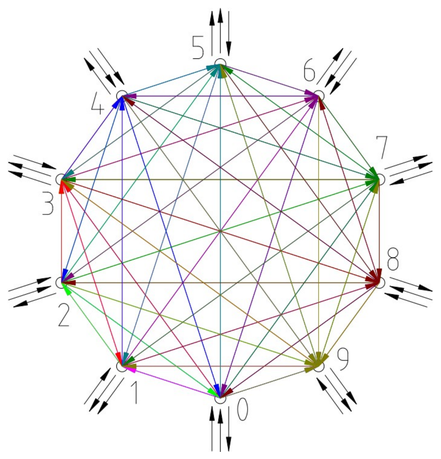

Topology for a 40 FPGAs network, black arrow are links towards other 3 similar networks. The total bandwidth is 1,6 Tbps

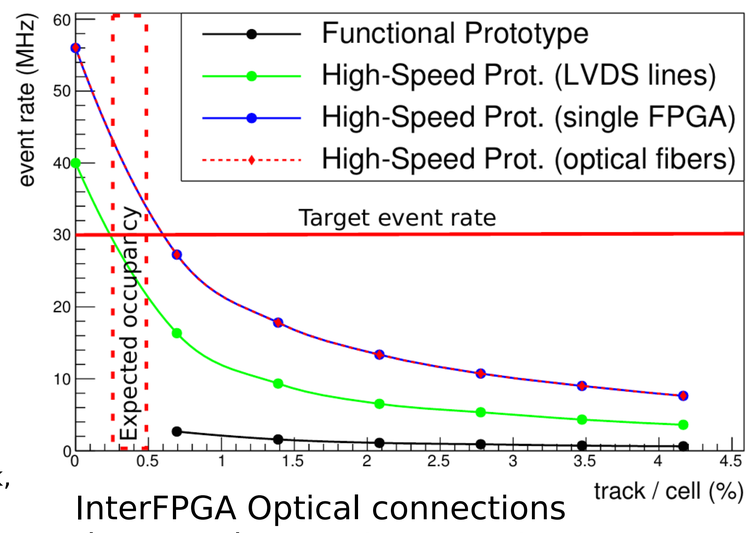

InterFPGA Optical connections do not reduce system event-rate

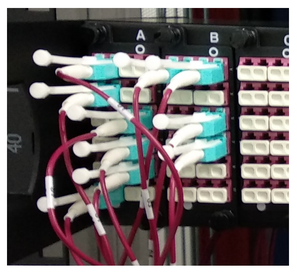

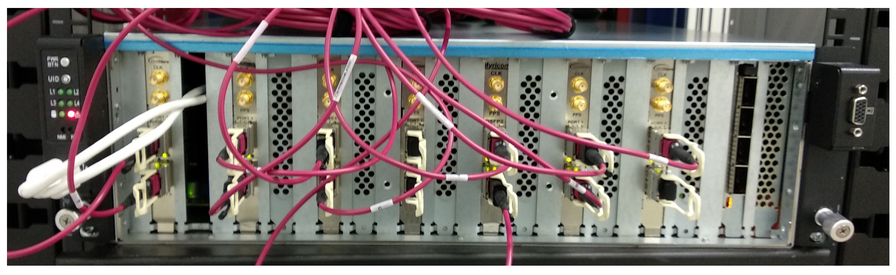

Setup of a 5-FPGA full-mesh network successfully tested for 24 days continuous operation