


Intelligent embedded processing



Giulia Tuci
giulia.tuci@cern.ch

10/11/2021



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


Introduction

- ❖ Great challenges will be faced up by LHC experiments in next years:
 - LHCb and ALICE are doing their major upgrade now
 - HL-LHC will come next
- ❖ Data reconstruction and storage will be a tough issue
- ❖ The situation will be even worse in the future colliders scenario
 - Large data volumes, more complex events
 - Larger B/W to the trigger
 - More computing power needed
- ❖ A smaller, easily-extracted, portion of the event not sufficient anymore to reduce data for a later processing
- ❖ → need to move more complex event reconstruction at the earliest stage of the trigger

The LHCb case

see also F.Pisani talk

- ❖ LHCb is already facing these problems now, due to special requirements of flavor 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 will be 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

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 track reconstruction

- ❖ Reconstruction of charged particle trajectories: large combinatorial problem → calls for high parallelization
- ❖ **This talk:** real-time tracking on FPGAs with the “artificial retina” architecture
 - General idea
 - Demonstrator system: reconstruction of tracks in the LHCb vertex detector → will be tested in the coprocessor testbed
- ❖ CEPC environment will be of course different w.r.t LHC one
 - e^+e^- collisions → less busy events
 - But high luminosity, up to $3.2 \times 10^{35} \text{ cm}^{-2}\text{s}^{-1}$
 - → similar challenges
 - Cost of needed technology will be much lower at CEPC time-scale

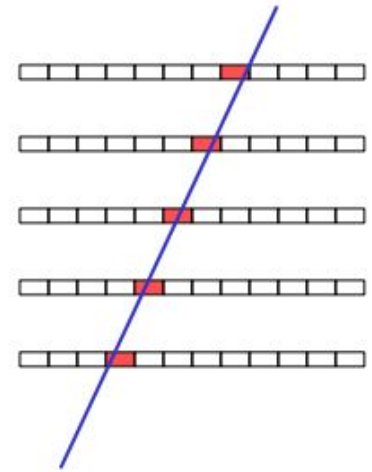
Past examples: real-time tracking by pattern-matching

- ❖ Fastest approach to tracking used up to now: direct matching to a bank of stored templates

- ❖ Successfully used by CDF (Tevatron)

- Real-time processor SVT capable of reconstructing high quality tracks in $\sim 10\mu\text{s}$ (input rate 0.03 MHz)
- Implemented in custom ASIC, Associative Memories

NIM A278 (1989), 436-440



- ❖ L0 track reconstruction in LHC environment more challenging


- Input rate is 40 MHz
- Very low number of clock cycles available per event (about 25 cycles)
- Latency must still be contained within few μs , to avoid the need of a large buffering space
- Is it a feasible task?

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




Natural vision features

- ❖ Only relevant data reaches a photoreceptor cell
 - Faster
 - Processing power can be spread over a network
- ❖ Continuous neuron response to light excitation, by interpolation of analog responses
 - Only yes/no response from Associative Memories
 - Save internal storage
- ❖ Can we implement a similar algorithmic in electronics?

The “artificial retina” algorithm

NIMA 453 (2000) 425-429






Parameter space divided into cells.

Each cell → pattern.

Receptors: interception of pattern tracks
with detector layers.

The “artificial retina” algorithm




Distance hit- receptor is calculated.
A weight is assigned to each distance.
For each cell all the weights are summed
(excitation level).

Parameter space divided into cells.
Each cell → pattern.
Receptors: interception of pattern tracks
with detector layers.

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



The “artificial retina” algorithm



Parameter space divided into cells.



Each cell → pattern.

Receptors: interception of pattern tracks with detector layers.



Distance hit- receptor is calculated.
A weight is assigned to each distance.
For each cell all the weights are summed (excitation level).

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

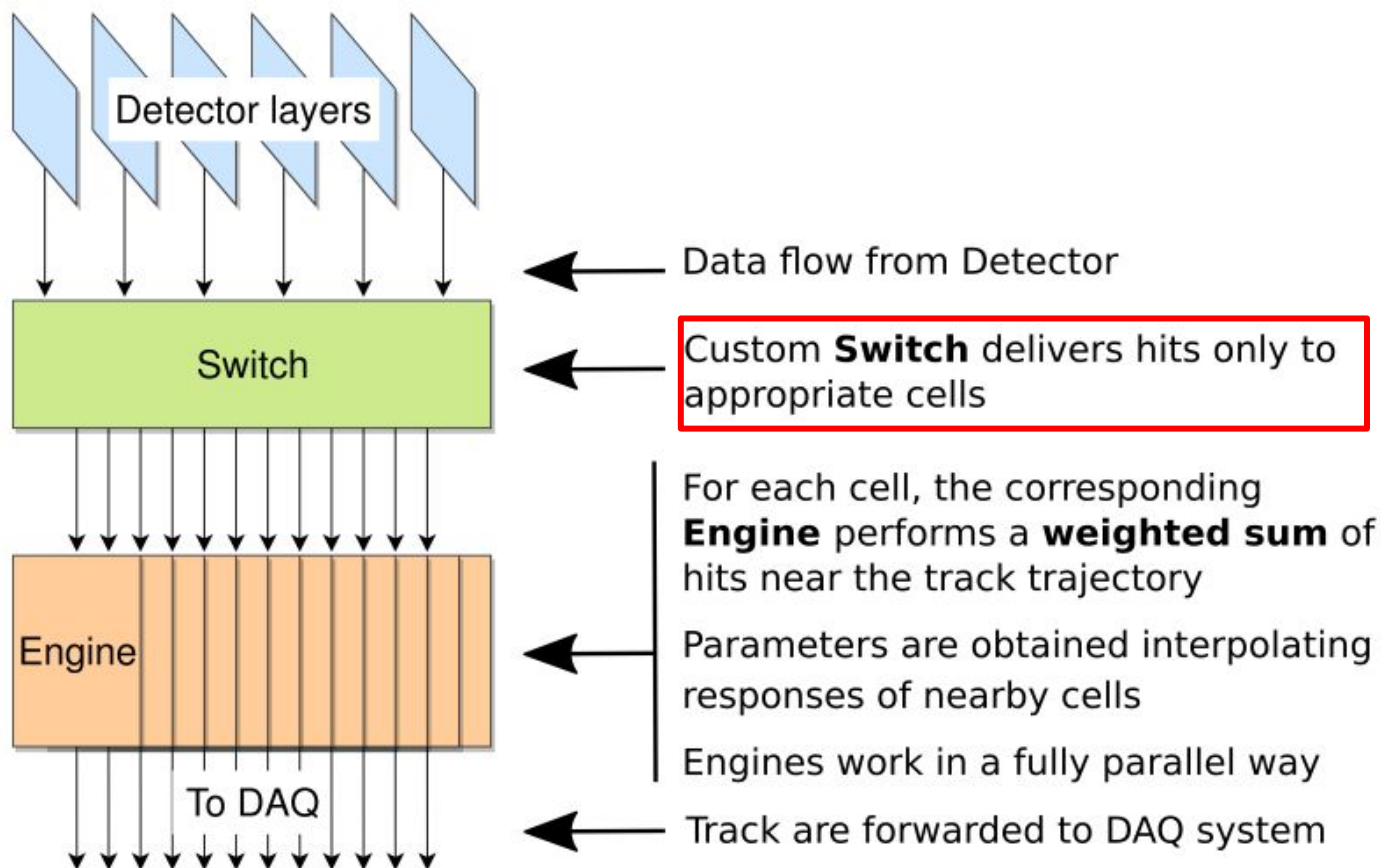


Search of local maxima over a threshold → reconstructed tracks

Real track parameters are obtained interpolating responses of nearby cells

Everything is executed in parallel!

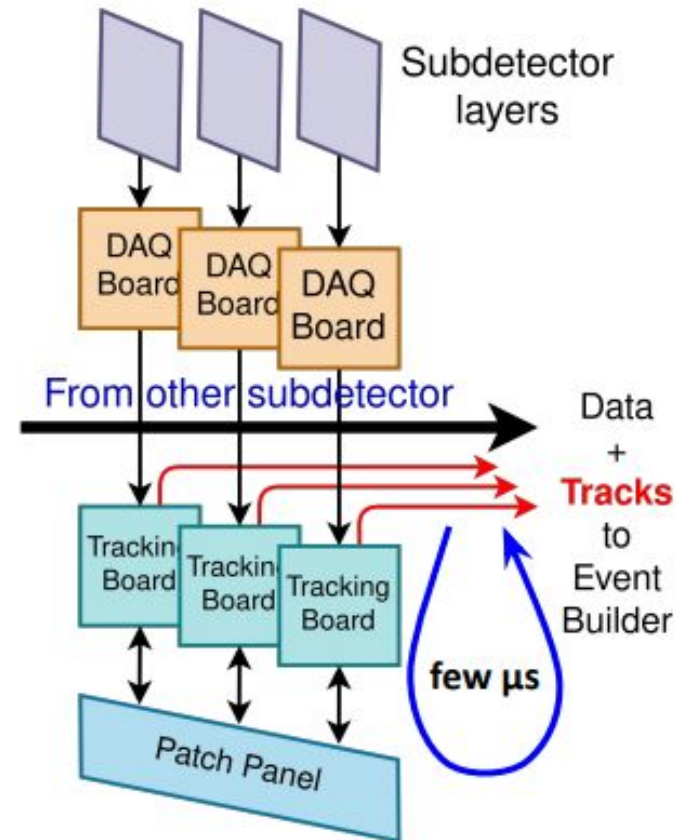
System architecture



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

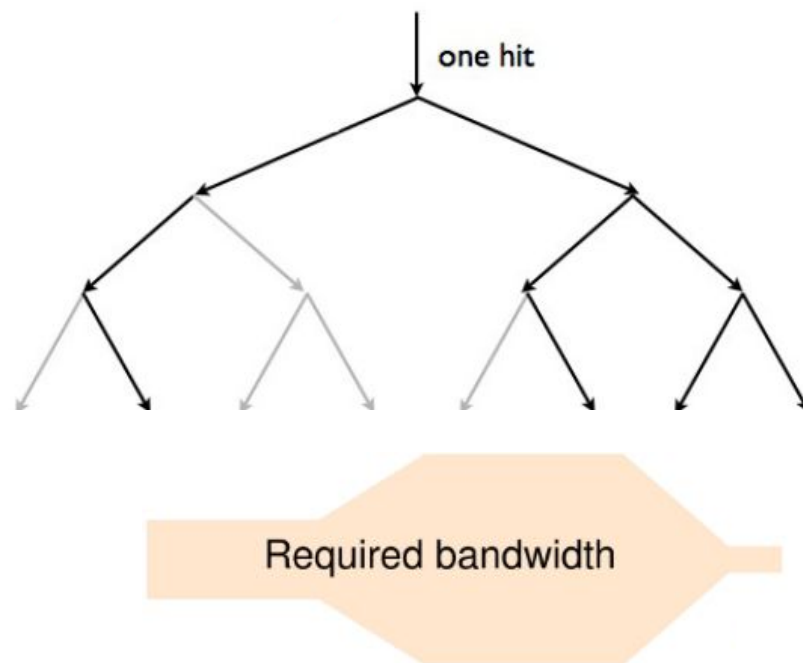

Technology choice: FPGAs

- ❖ **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
 - Tracks could replace hits data in real time, reducing data volume before even performing Event Building
- ❖ 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**





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: tracks 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)

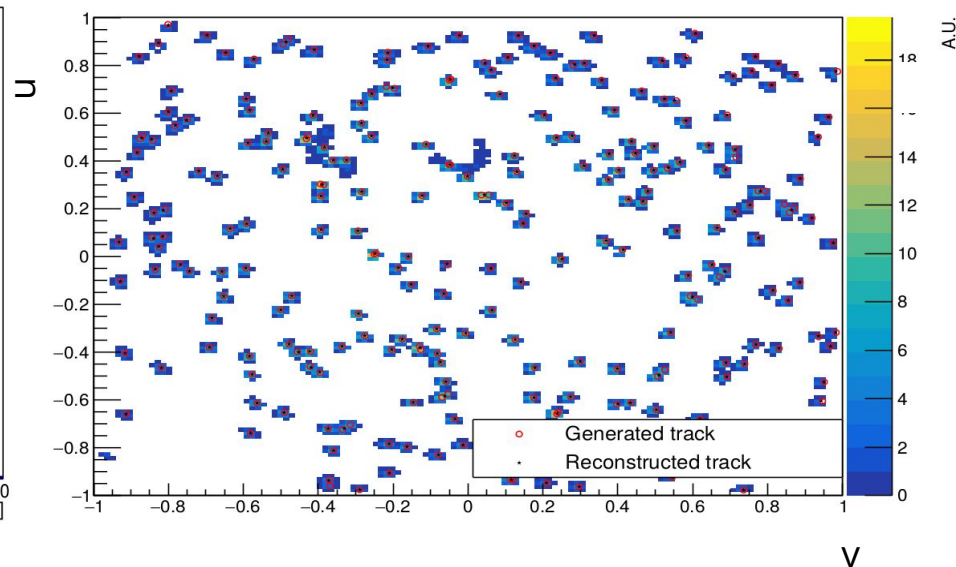


Segmentation of parameter space

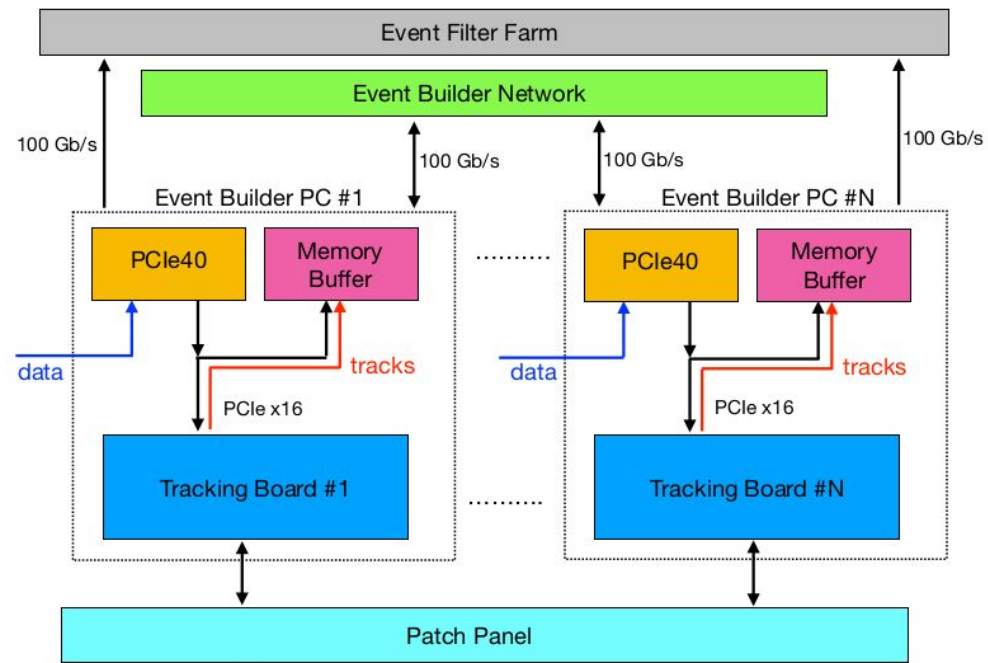
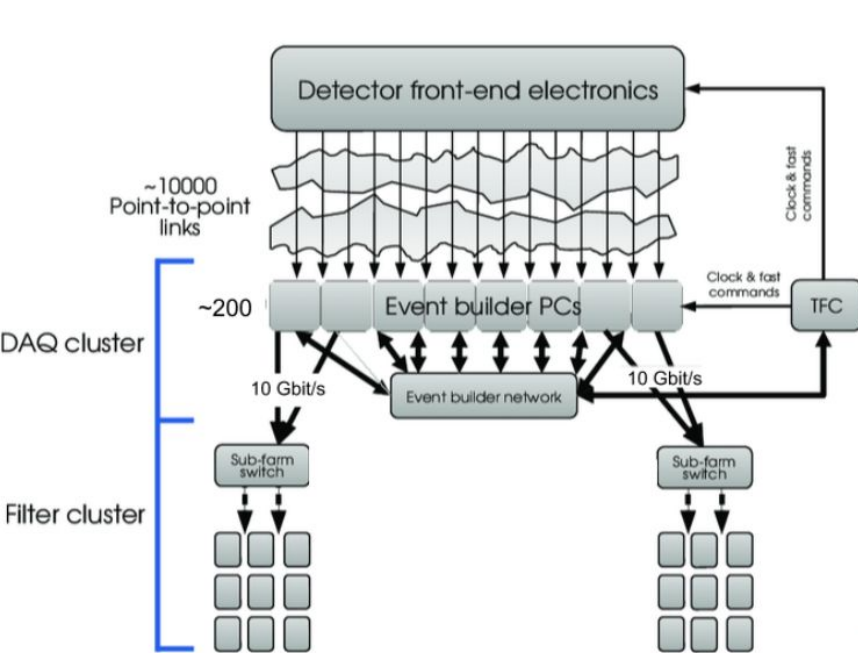
- ❖ Algorithm optimized via C++ bit-by-bit emulator for FPGA pattern recognition
- ❖ Best configuration:
 - u, v are the x, y coordinates of intercept with a fixed plane
 - Algorithm based on a 2D space → segment track space in 10 z-slices
 - 10000 cells for each slice (100k total cells)
 - Fits well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



Distribution of excitation levels for one slice



Integration in the DAQ




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

Performance

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


LHCb-FIGURE-2019-011



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

Reconstructible track: at least 3 hits in different VELO modules

Long tracks \rightarrow crossing all tracking detectors



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

Implementation status

PoS(TWEPP-17) 136

- ❖ Old hardware prototype of generic 6-layer tracker already demonstrated throughput exceeding 30 MHz, but need to specialize for the VELO (ongoing)
- ❖ Logic resources: estimate to fit comfortably in the 38 current COTS FPGA cards, needed to match the readout boxes of the VELO

Commercial board equipped with Altera Stratix 10 FPGA is under testing



- ❖ Successfully tested optical network with 8 links on Stratix 10 boards (26 Gbps each) using SuperLite II v4 communication protocol

PROC-CTD2020-23

TWEPP-2021

Final remarks

- ❖ Future 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
- ❖ Detailed simulation studies show fairly good tracking performance for fast VELO tracking and practical feasibility of implementation in current COTS FPGA systems!
 - Ongoing implementation of a demonstrator system which will operate at coprocessor testbed at LHCb
- ❖ Limiting factors related to available logic in current chips, not to the architecture itself
- ❖ Huge cost/performance drops expected in CEPC time-scale