



Istituto Nazionale di Fisica Nucleare  
SEZIONE DI CAGLIARI

# Real time data processing with FPGAs at LHCb

**Andrea Contu**, Federico Lazzari, Francesco Terzuoli, Giovanni Bassi, Giovanni Punzi, Giulia Tuci, Michael J. Morello, Riccardo Fantechi, Sofia Kotriakhova, Wander Baldini **on behalf of the LHCb Collaboration**

## Topical Workshop on Electronics for Particle Physics

1-6 October 2023

# The LHCb detector in Run 1 and Run 2 (2011-2018)

- Excellent particle identification, IP and momentum resolution ( $\sim 13 \mu\text{m}$  on the transverse plane and  $\Delta p/p \sim 0.5\% - 0.8\%$ , respectively.)
- Huge beauty and charm production

$$\sigma(pp \rightarrow b\bar{b}X)_{2 < \eta < 5} = 144 \pm 1 \pm 21 \mu\text{b}$$

[PRL 119, 169901 (2017)]

$$\sigma(pp \rightarrow c\bar{c}X)_{p_T < 8 \text{ GeV}/c, 2.0 < y < 4.5} = 2369 \pm 3 \pm 152 \pm 118 \mu\text{b.}$$



[JHEP 05 (2017) 074]



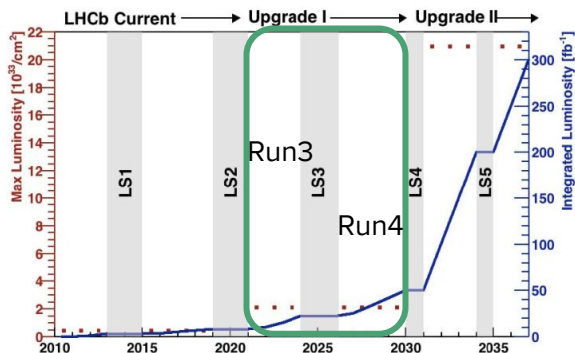
**LHCb Detector**

Weight: 5,600 tonnes  
Height: 10 m  
Length: 20 m


[JINST 3 S08005 \(2008\)](#)  
[JMP A 30, No. 07, 1530022 \(2015\)](#)



# The LHCb detector in Upgrade 1

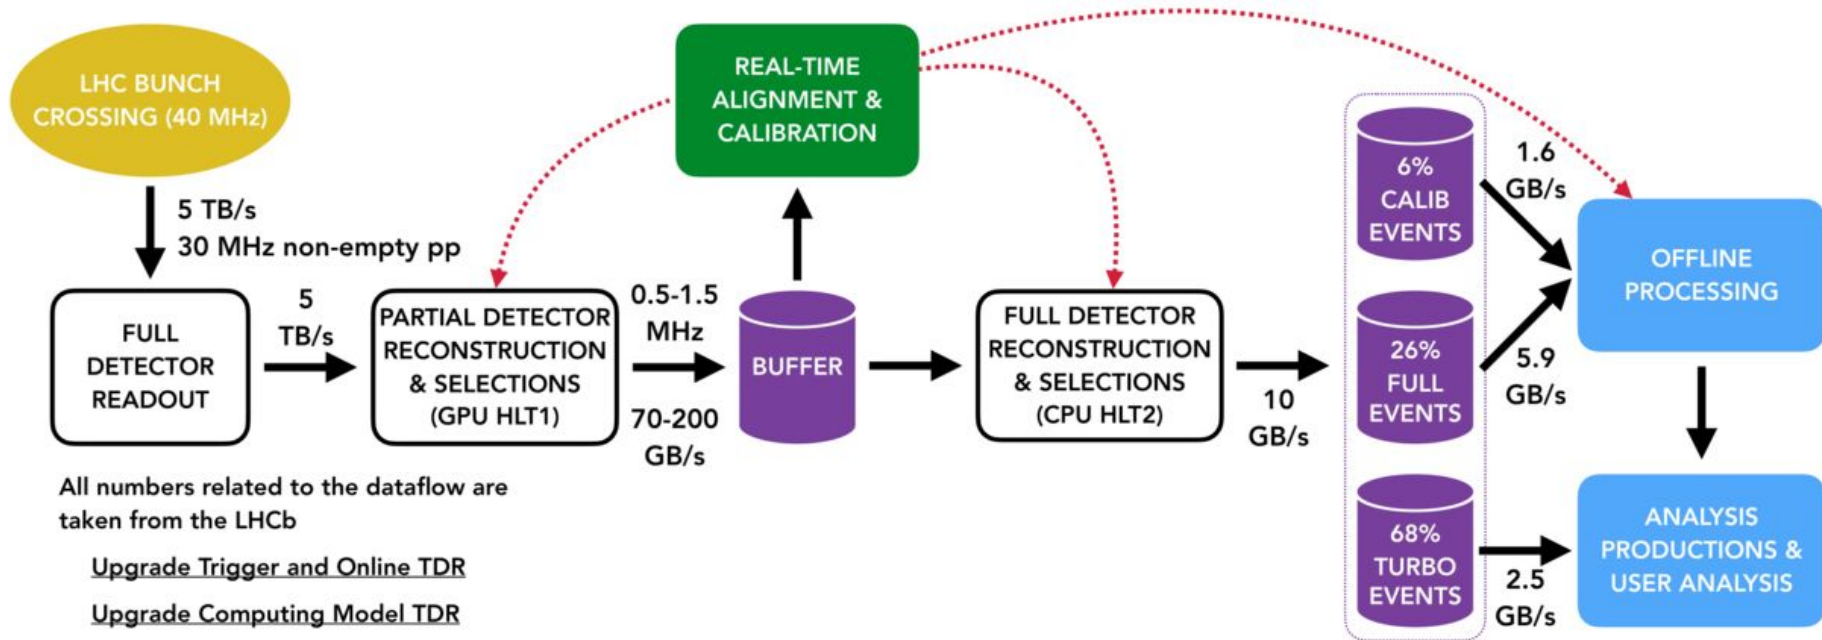


- Aim to collect  $\sim 50 \text{ fb}^{-1}$  at roughly  $\mathcal{L} = 2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$
- Keeping at least the same performance on Run 1&2



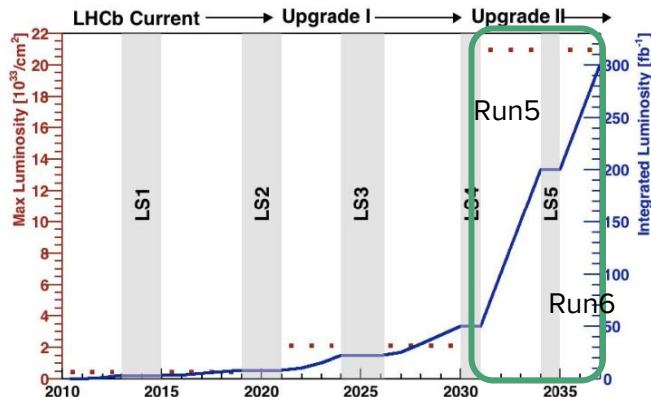
# The DAQ and trigger in Upgrade 1

**Fully software trigger running on GPU**, overcomes rate limitations in Run1&2 and builds on the successes of Run1 and Run2 (e.g. **real time alignment and calibration**)



# LHCb in Run 5&6 ?

- Target:  $\sim 300 \text{ fb}^{-1}$
- Pile-up:  $\sim 40$
- To keep the same performance in more difficult conditions, timing will be required in some sub-detectors
- **200 Tb/second data produced**
- **More processing has to be performed earlier in the DAQ Chain to reduce data offline**
- **Moving to a “heterogeneous-computing” paradigm**



[LHCb-TDR-023](#)

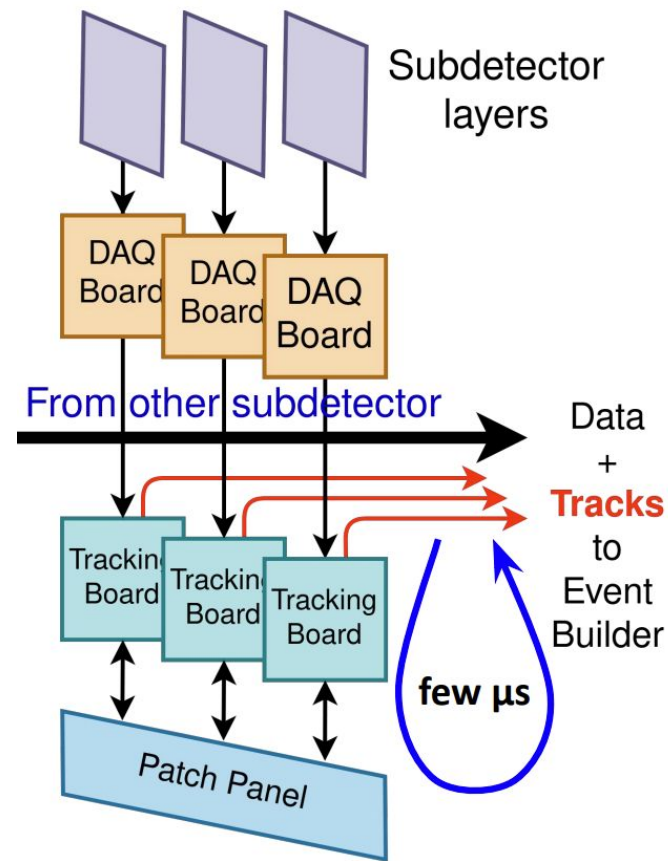


# Real time tracking with FPGAs

- Modern FPGAs can perform parallel data processing with high throughputs, low latencies and better energy efficiency than CPUs and GPUs (for certain tasks)
- **This talk:** demonstrator system for real-time tracking on FPGAs with the “**artificial retina**” architecture to reconstruct tracks in the Vertex Locator




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

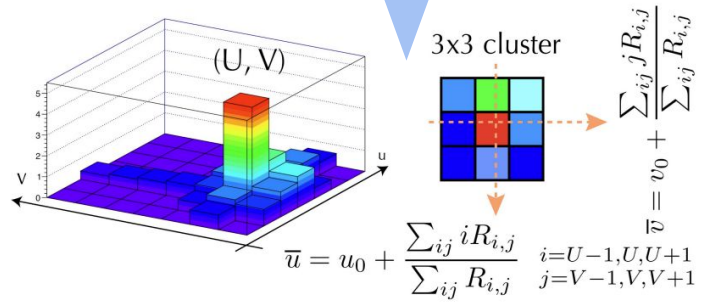



# The “artificial retina” architecture [\[NIMA 453 \(2000\) 425-429\]](#)

Track parameter space divided into cells (pattern tracks)



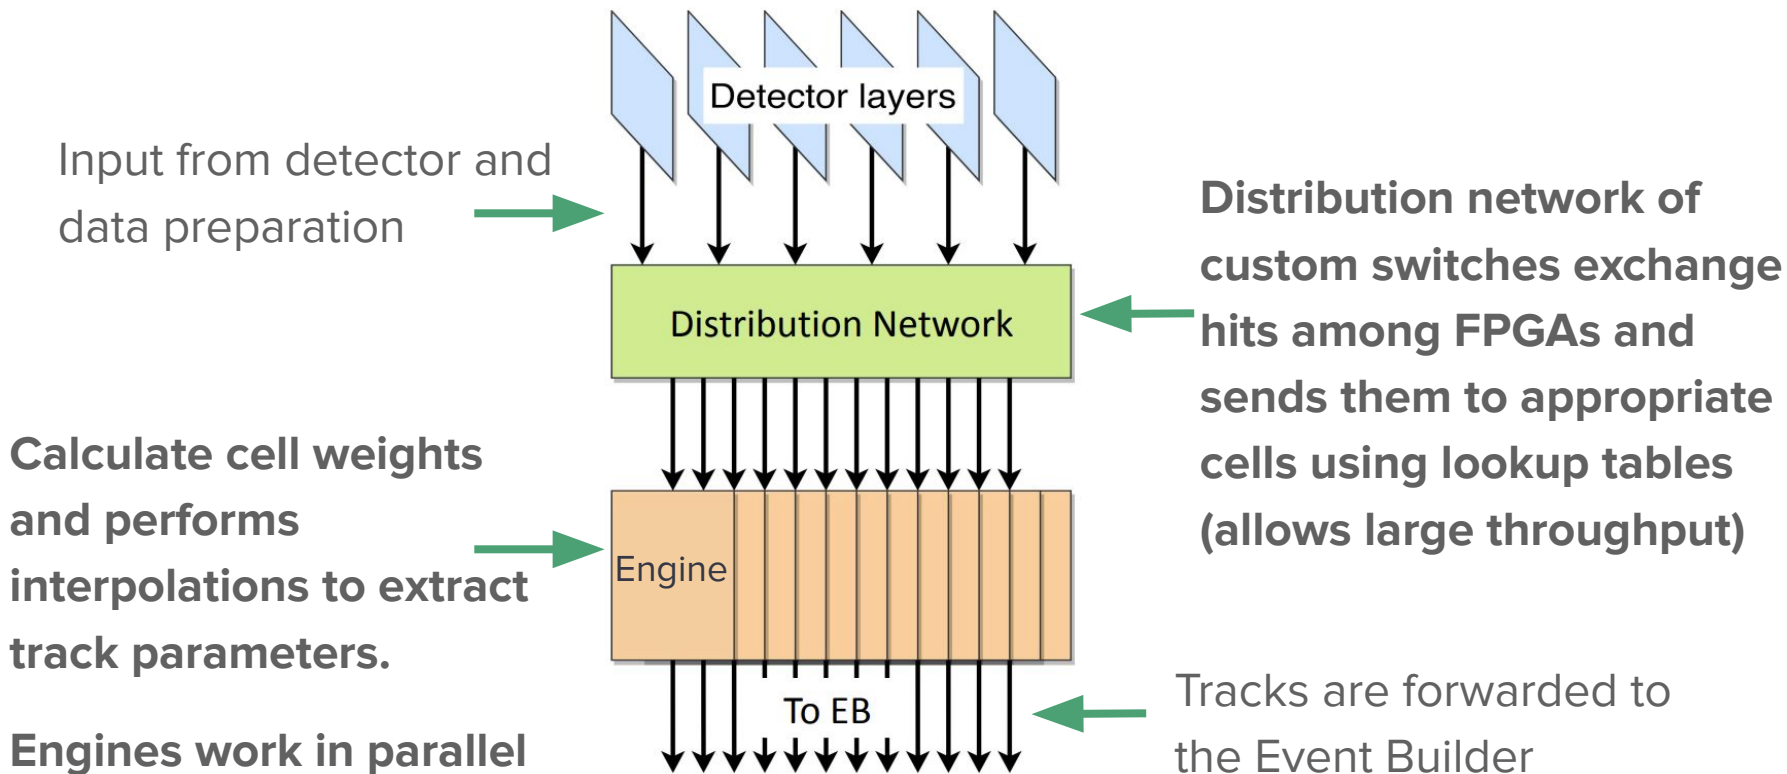
Each cell computes a weighted sum of hits near the reference track



- Reconstructed tracks correspond to local maxima in the matrix of cells response
- Final track parameters from interpolating responses of nearby cells.

- **Cells work in parallel:** high-throughput and low-latencies
- **FPGA size limitations overcome by spreading cells over several chips** (without increasing latency).


# The “artificial retina” architecture [\[NIMA 453 \(2000\) 425-429\]](#)






# A realistic test: track reconstruction in the VELO

- The VELO is a crucial subdetector for LHCb [see [A.F. Prieto's talk](#)]



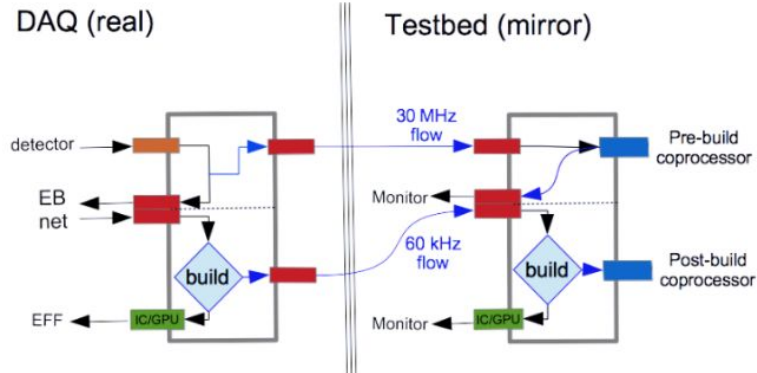
VELO top view



- **52 modules (38 in forward region), 10% of LHCb data size. 25% HLT1 time**
- Need compact FPGA system, studies already exist [[PoS Vertex2019 \(2020\) 047](#), [EPJ Web. Conf. 245, 10001 \(2020\)](#)]
- **A good test-case for future and larger-scale applications**
- **VELO pixel clustering originated from retina algorithm is already running on FPGA in Run3, integrated in the DAQ boards firmware!**

# The LHCb Co-processor Testbed

- A testbed dedicated to study new technologies has been set-up at LHCb site.
- It can host various projects with the one discussed here being the most advanced
- It allows to run parasitically (not at full rate yet) with the “normal” DAQ to be able to test in realistic scenarios



Front

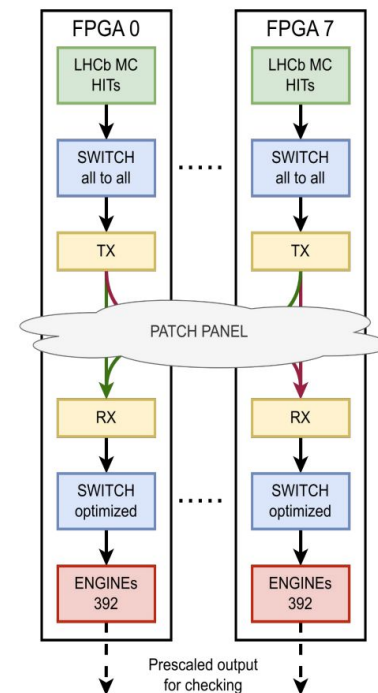


Back



# The RETINA demonstrator at the testbed

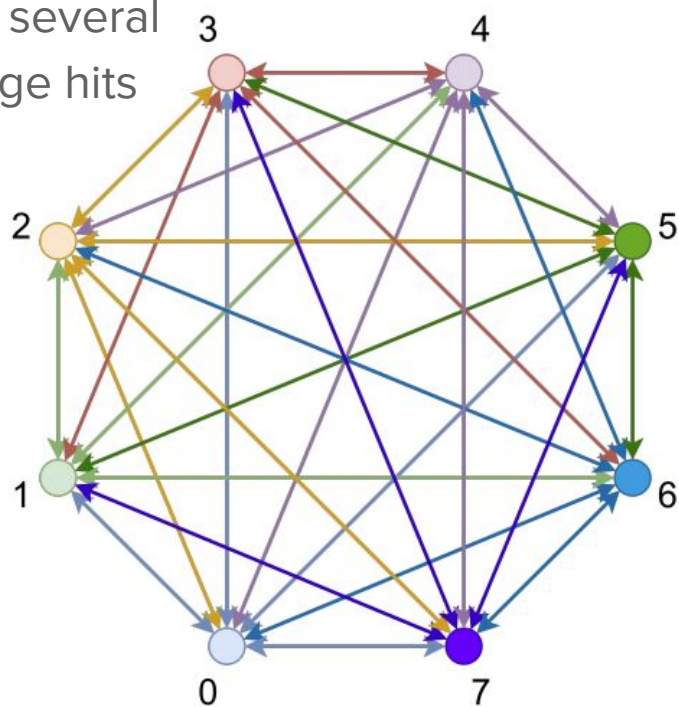

- Simulated data used for high rate tests
- **Now live data from the LHCb monitoring farm**
- Demonstrates that a RETINA based tracking on FPGAs possible in HEP experiments:
- Current setup:
  - Reconstructs tracks of a VELO quarter
  - Spread over multiple PCIe-hosted FPGA cards. **8 cards** are sufficient
  - Scalable to cover the whole detector with additional FPGA cards.



# Distribution network

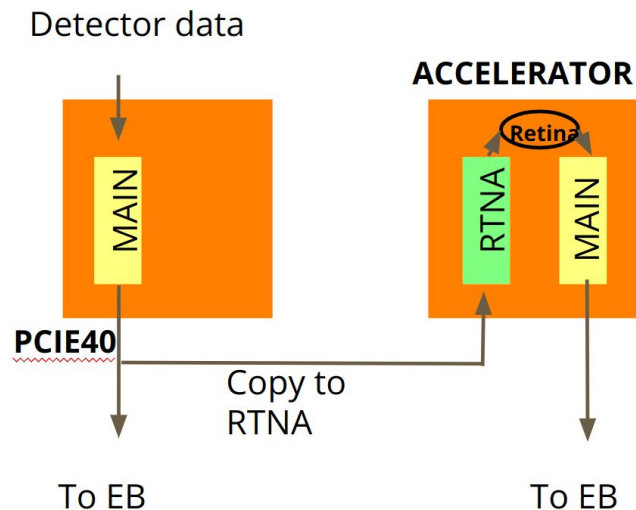
- As the RETINA algorithm is spread among several boards, a distribution is needed to exchange hits among boards:

- 8 nodes full-mesh network
- 28 full-duplex links at 25.8 Gbps
- Total bandwidth 1.41 Tbps



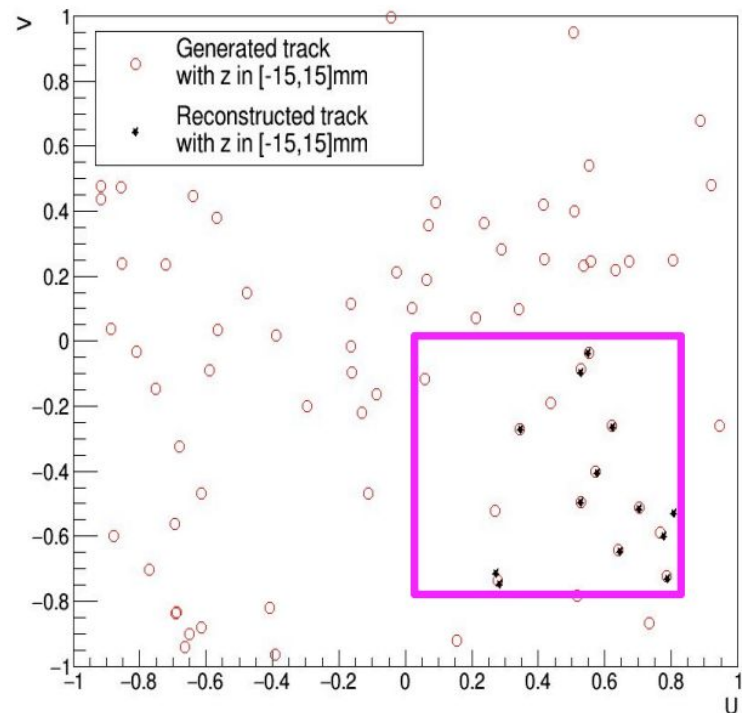
# First attempt of LHCb DAQ integration at the testbed

- **Pair one accelerator board running retina to a readout board (PCIE40), they can be the same board but with different firmware**
- A modified driver is implemented to perform a copy of the detector data to the accelerator (via PCIE) **before sending it to the Event Builder**
- The accelerator sends data (=tracks) to the EB as if it was another subdetector
- **Transparent to the EB, no significant drop in throughput is noticed**



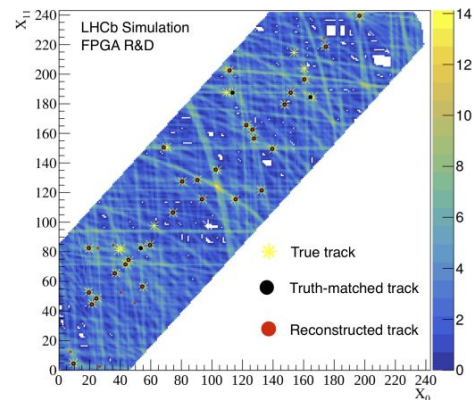

# Demonstrator Results

- No issues in several days of running
- Tracks match the tracks reconstructed by the C++ simulation.
- **Currently running at 16.2 MHz**
- Working on design optimisation, expect to reach the desired 30 MHz throughput
- **Data processing on FPGA is possible!**
- **Now moving towards a concrete proposal for Run4 & 5**



# Proposal for Run4: the Downstream Tracker [[ACAT2019](#)]

- Proposal for a **downstream tracker (DT)** RETINA-like tracking in Run4 under scrutiny at LHCb
- Not (yet) included in HLT1 as it is computationally expensive
- The DT could significantly extend LHCb's physics reach for long lived particles
- **DAQ integration is crucial**, strong constraints from available servers/PCIE slots/bandwidth, several options under discussion to minimise impact on operations and maximise performance




# Conclusions


- LHCb has put in place a testbed for heterogeneous computing tests
- FPGA usage for data processing is feasible and could lower the computational burden (and cost) down the DAQ chain
- Proceeding by steps:
  - a. FPGA-based VELO clustering already in production
  - b. VELO tracking demonstrator close to desired performance
  - c. Now proposing a Downstream Tracking system for Run4
  - d. Computing cost is a huge factor in Run5, designing a system that includes co-processors from the start seems the way to go
- Stay tuned for more news in the coming years!



# Backups

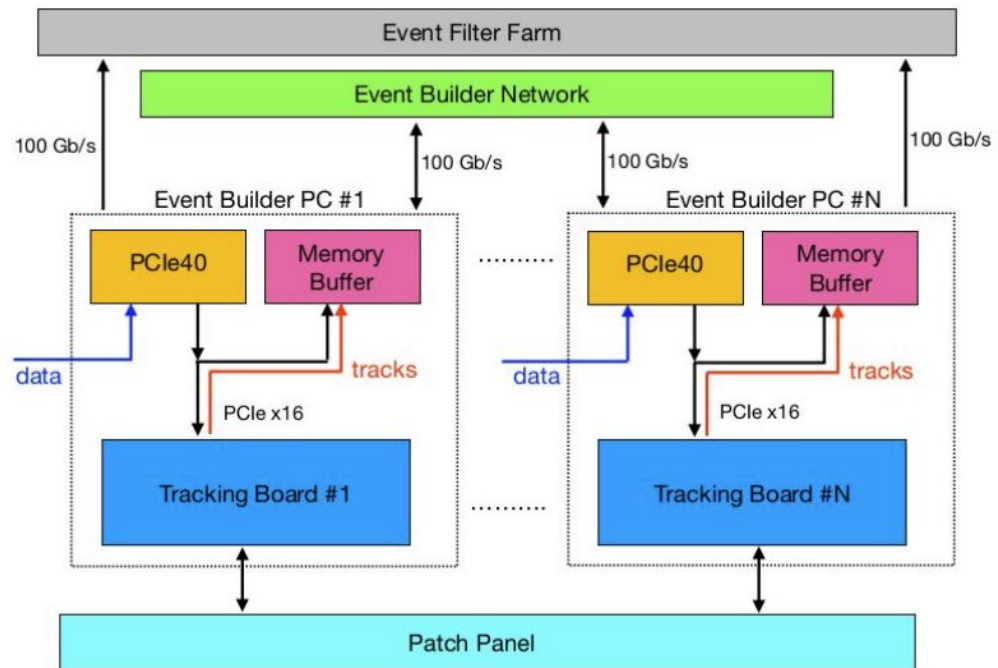
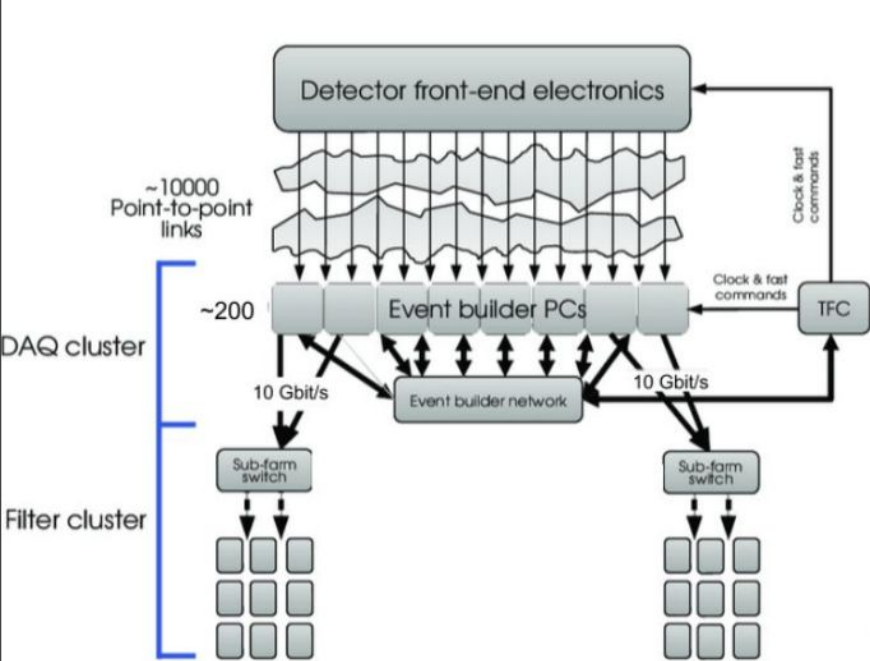


# Trigger yield vs lumi in Run 1&2



**Table 1.** Axial only ( $\varepsilon_A$ ) and three-dimensional ( $\varepsilon_{AS}$ ) averaged reconstruction efficiencies for different simulated samples and different track categories. The ghost rate is also shown. The downstream strange tracks are mainly pions from  $K_S^0 \rightarrow \pi^+\pi^-$  decay.

| Track type                              | Minimum Bias    |                    | $D^0 \rightarrow K_S^0 \pi^+ \pi^-$ |                    | $B_s^0 \rightarrow \phi\phi$ |                    |
|---|-----------------|--------------------|-------------------------------------|--------------------|------------------------------|--------------------|
|   | $\varepsilon_A$ | $\varepsilon_{AS}$ | $\varepsilon_A$                     | $\varepsilon_{AS}$ | $\varepsilon_A$              | $\varepsilon_{AS}$ |
| T-track                                 | 75.0            | 71.4               | 74.4                                | 70.0               | 73.9                         | 67.4               |
| T-track, $p > 3\text{GeV}/c$            | 87.0            | 83.0               | 85.9                                | 80.8               | 85.1                         | 77.2               |
| T-track, $p > 5\text{GeV}/c$            | 90.3            | 85.7               | 88.2                                | 82.7               | 86.6                         | 77.4               |
| Long                                    | 81.7            | 78.8               | 84.1                                | 79.5               | 84.2                         | 77.2               |
| Long, $p > 3\text{GeV}/c$               | 87.3            | 84.2               | 87.1                                | 82.3               | 87.3                         | 79.8               |
| Long, $p > 5\text{GeV}/c$               | 90.6            | 86.9               | 88.1                                | 83.1               | 88.1                         | 79.9               |
| Downstream                              | 80.1            | 77.7               | 83.0                                | 78.6               | 82.6                         | 76.2               |
| Downstream, $p > 3\text{GeV}/c$         | 87.0            | 84.4               | 87.1                                | 82.5               | 86.5                         | 79.3               |
| Downstream, $p > 5\text{GeV}/c$         | 90.5            | 87.5               | 88.8                                | 83.6               | 87.9                         | 80.2               |
| Downstream strange                      | -               | -                  | 84.7                                | 82.8               | -                            | -                  |
| Downstream strange, $p > 3\text{GeV}/c$ | -               | -                  | 89.4                                | 86.7               | -                            | -                  |
| Downstream strange, $p > 5\text{GeV}/c$ | -               | -                  | 93.0                                | 87.2               | -                            | -                  |
| ghost rate                              | 12.1            | 15.7               | 16.3                                | 20.2               | 18.4                         | 24.7               |



# Downstream Tracker: DAQ integration options

