A GPU High Level Trigger 1 for the upgraded LHCb detector

Marianna Fontana on behalf of the LHCb collaboration

LPNHE (Paris)

EPS conference 2021

Virtual, 26 - 30 July 2021

The upgraded LHCb detector

JINST 3 (2008) S080005

CERN-LHCC-2012-007

- General-purpose single arm forward spectrometer
- Optimized for heavy-flavour physics
- $0.5 - 1$ % momentum resolution
- $15 + 29/p_T$ [GeV] $\mathrm{m}u$ impact parameter resolution
- $(1 + 10/\sqrt{E}$ [GeV]$)$% ECAL resolution
- Excellent particle identification over the $2 - 100$ GeV p_T range

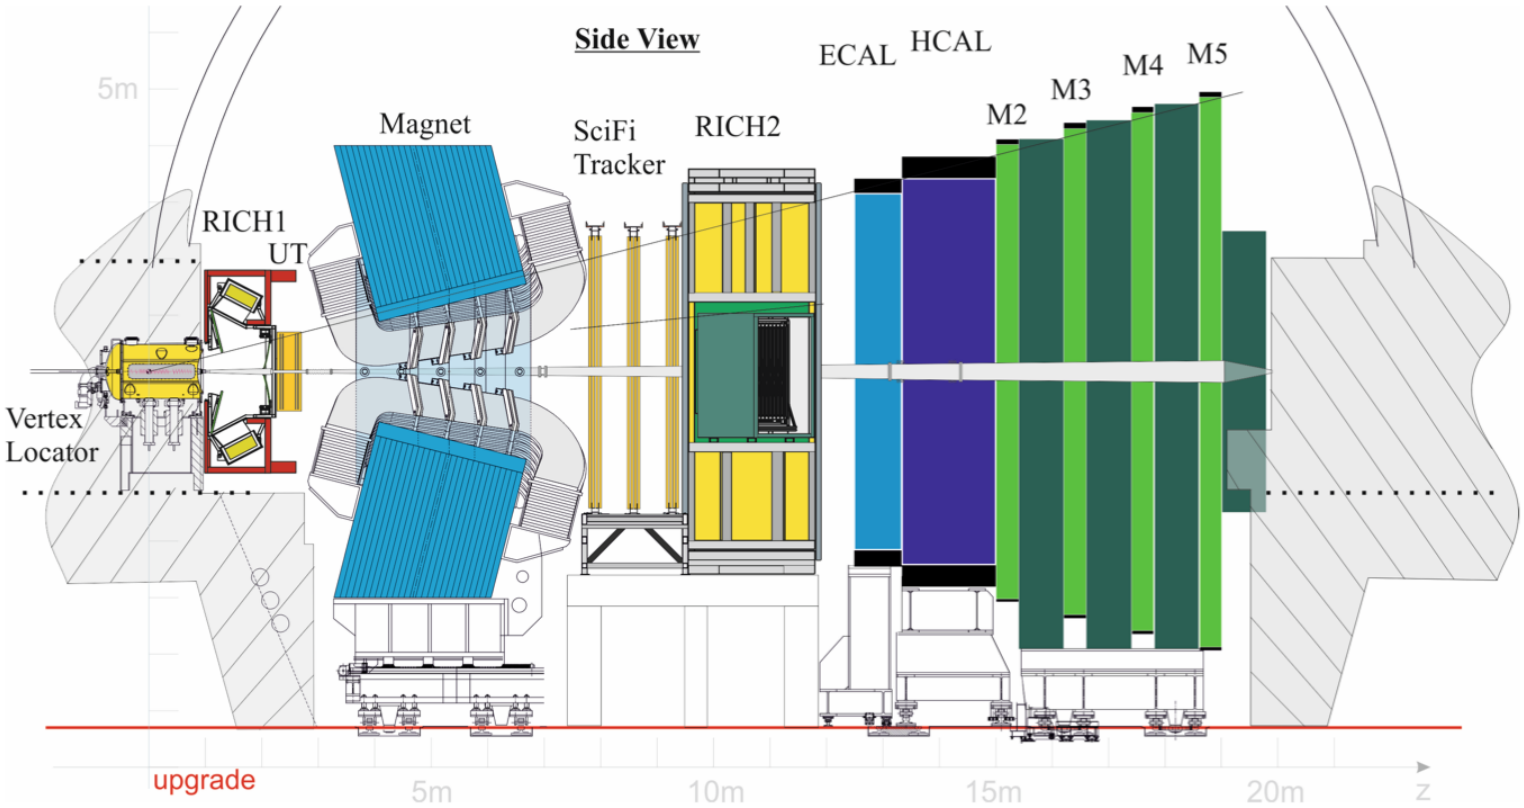

The LHCb trigger challenge

- Instantaneous luminosity:
  4×10^{32} (Run 2) $\rightarrow 2 \times 10^{33}$ cm^{-2}s^{-1}
- L0 rate limit of 1 MHz saturates fully hadronic modes
- **Solution**: read full event at bunch-crossing rate and apply track reconstruction selections

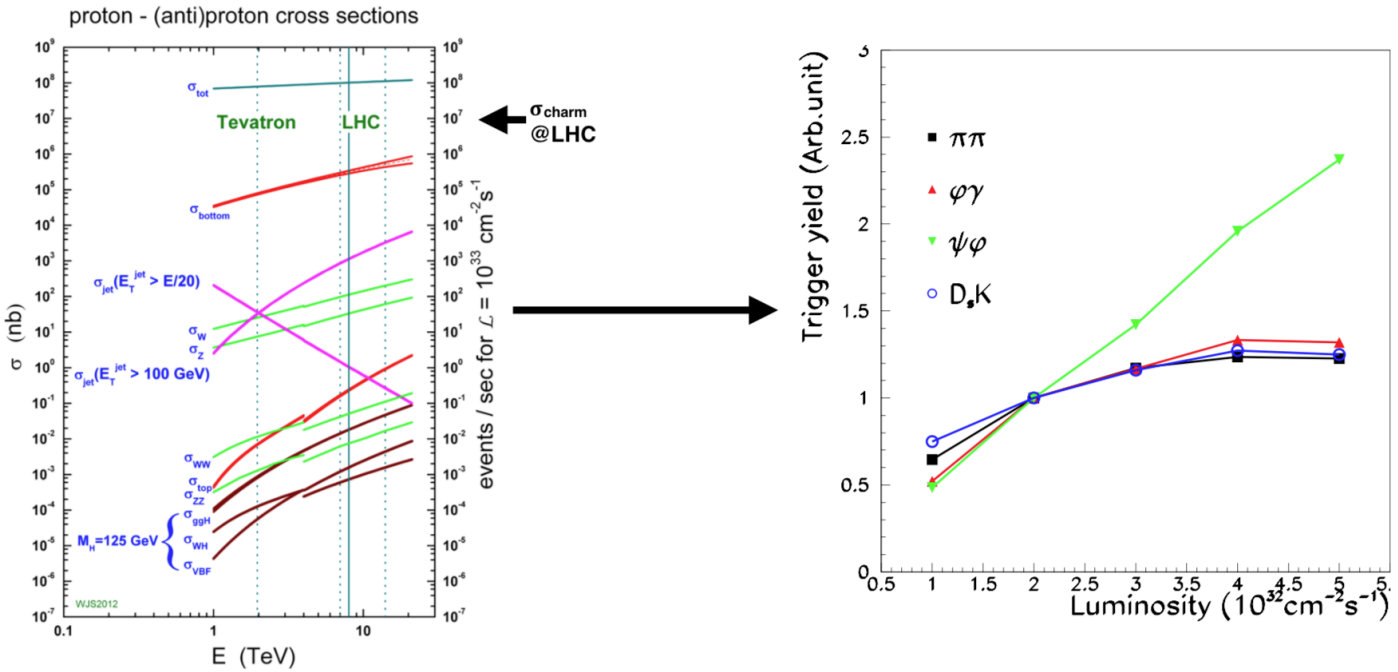

LHCb data flow in the upgrade [LHCb-FIGURE-2020-016]

- A partial reconstruction and selection is needed before running a full reconstruction

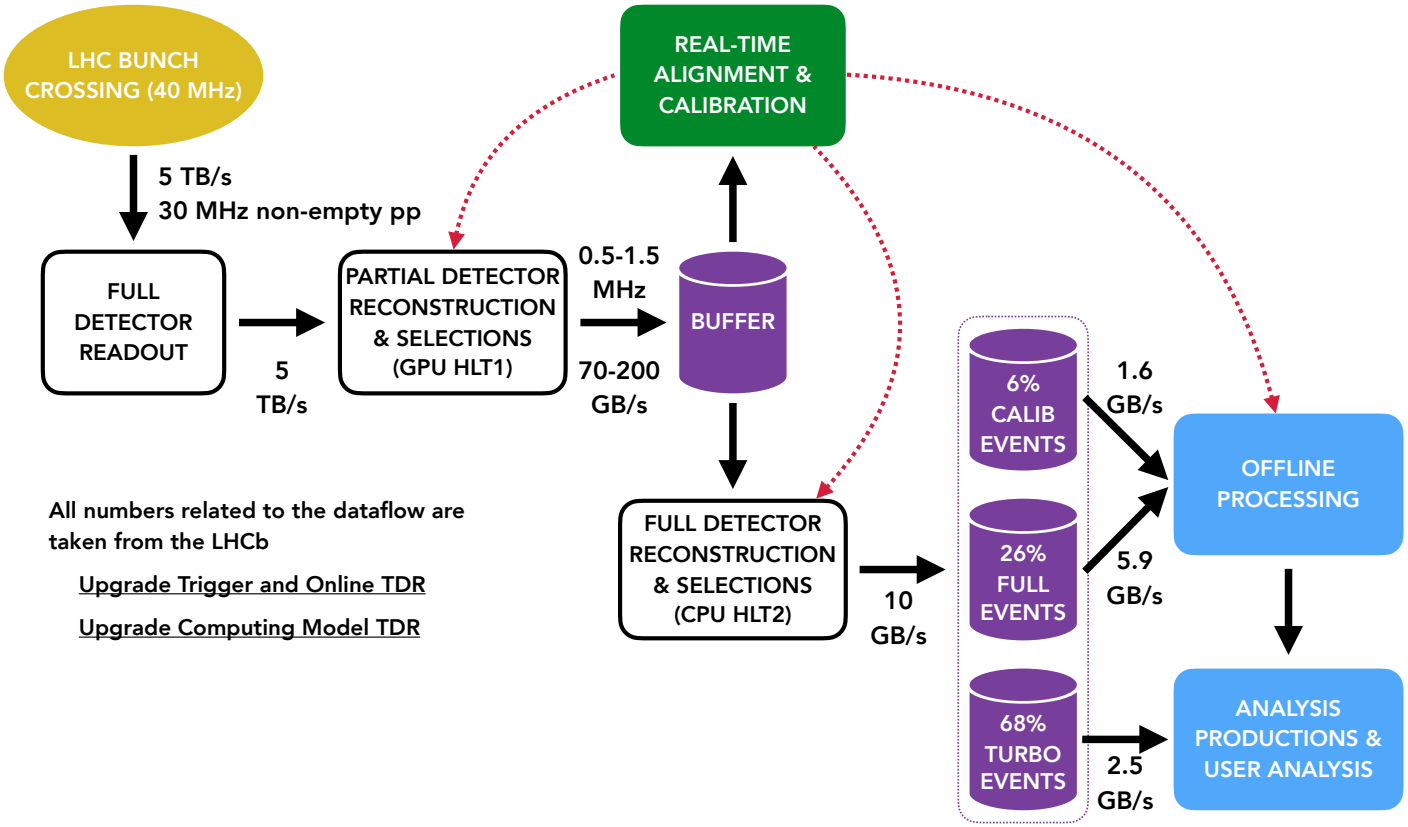

Allen project: HLT1 on GPU [Comput Softw Big Sci 4, 7 (2020)]

- Full track reconstruction at 30 MHz $\rightarrow$ computing challenge
- Track reconstruction is an inherently parallel problem $\rightarrow$ tracking algorithms can be designed to map well to the many-core architecture of GPUs
- `Allen` is a framework for GPU-based execution of an algorithm sequence [Gitlab repo]
- 40 Tbit/s input $\rightarrow$ 500 PCIe GPU cards to feed all data
- Each GPU process a slice of $\mathcal{O}(1000)$ events in parallel

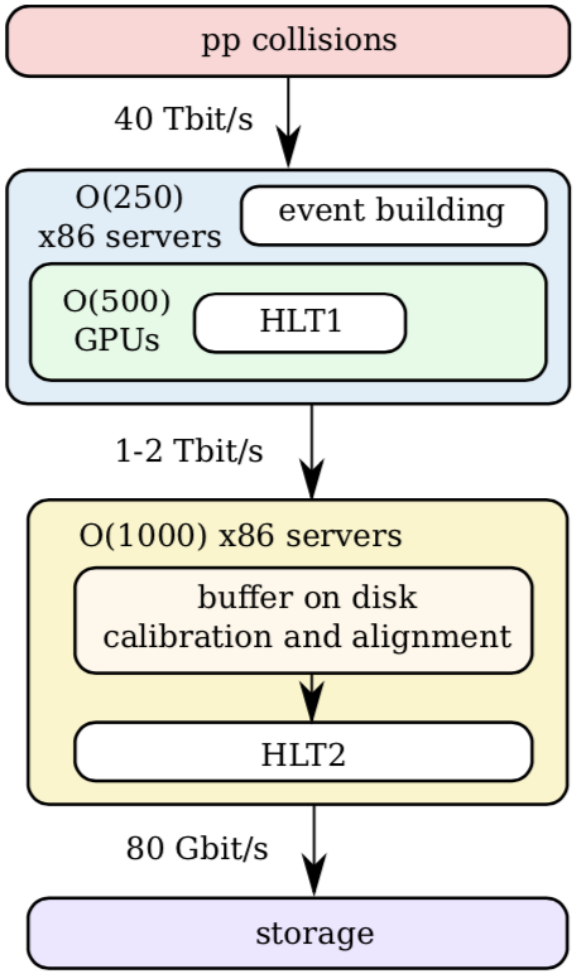

Integration with the Online system [LHCb Upgrade GPU TDR]

- GPUs map well onto LHCb's data acquisition architecture

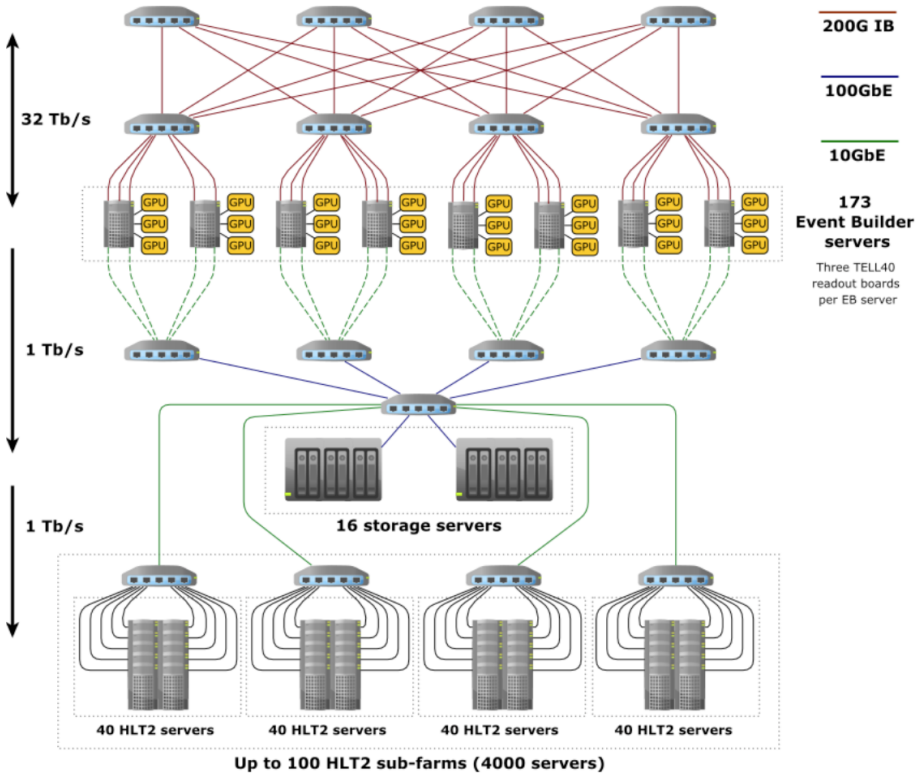

- $\mathcal{O}(500)$ FPGA readout boards receive data from detector and built the events
- The three PCIs slots can be filled with GPU cards and reduce the data rate to 1 Tb/s

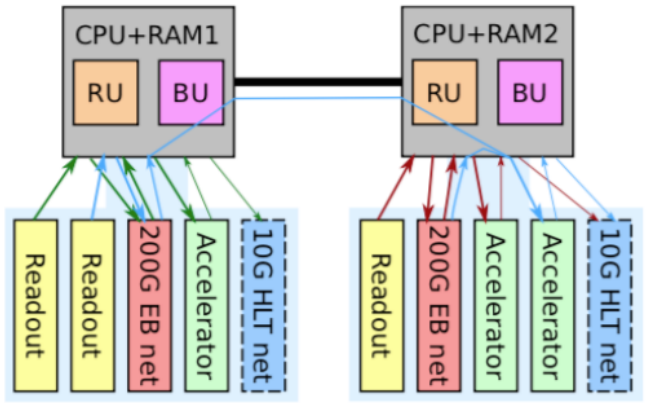

GPU-equipped event builder PC, with traffic of all three readout cards.

HLT1 reconstruction sequence [LHCb Upgrade GPU TDR]

- Subdetector reconstruction:
  - VELO: clustering, tracking, vertex reconstruction
  - UT, SciFi: tracking
- Global event reconstruction:
  - Track fit (Kalman filter)
  - Reconstruction of secondary vertices
- Perform trigger selections

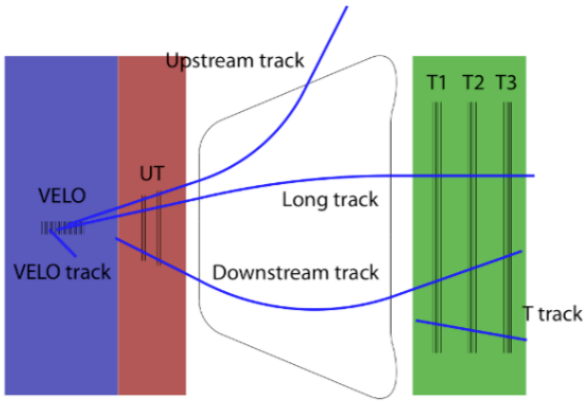

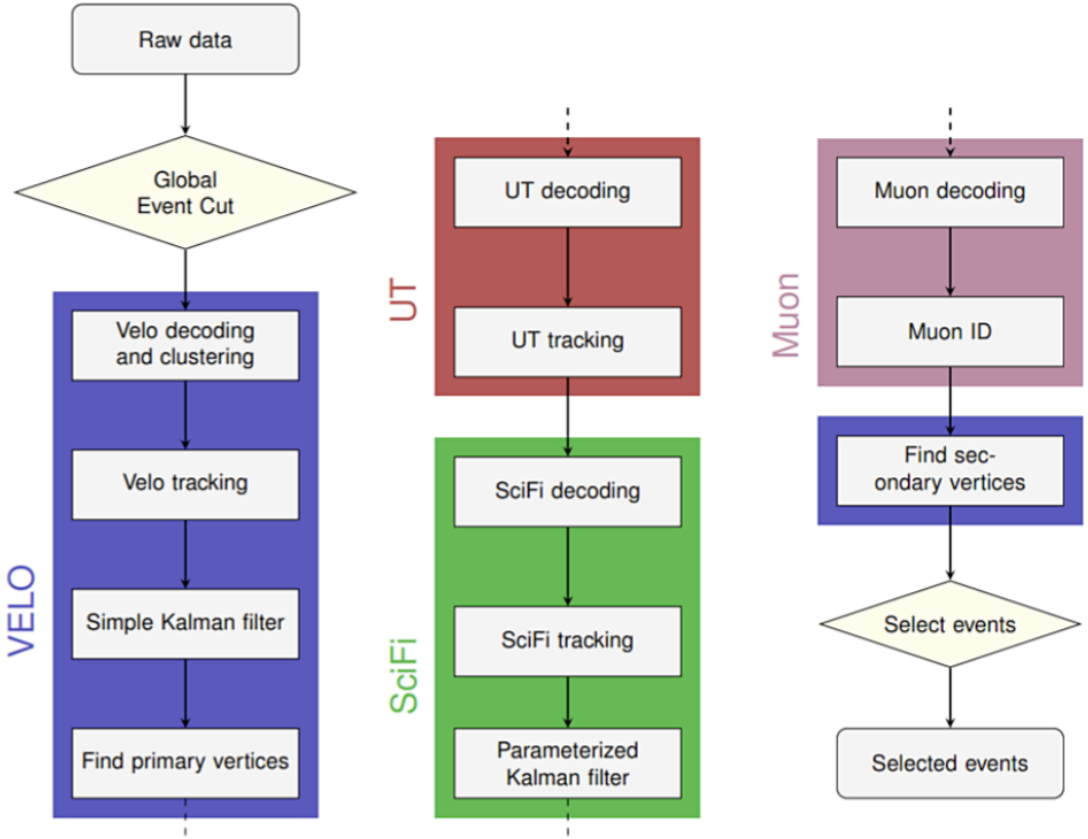

HLT1 throughput performance [LHCb-FIGURE-2020-016]

- The minimum rate per GPU necessary for processing the 30 MHz input rate with 500 GPUs is 60 kHz
- Throughput scaling very well with theoretical TFLOPS of GPU card

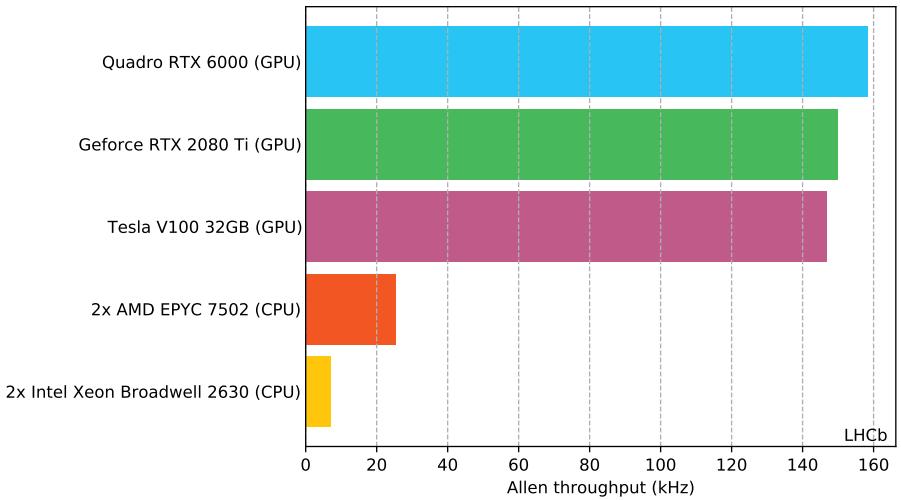

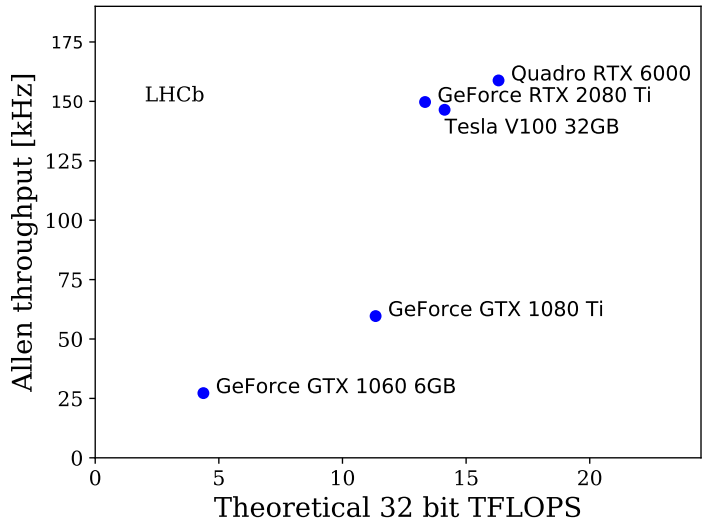

HLT1 tracking performance [LHCb-FIGURE-2020-016]

- The overall VELO reconstruction efficiency is around 99%
- The implemented reconstruction is significantly more performant than Run 1-2 HLT1 reconstruction
- The Forward reconstruction close to 95% efficiency at high p for tracks with $p_T > 500$ MeV

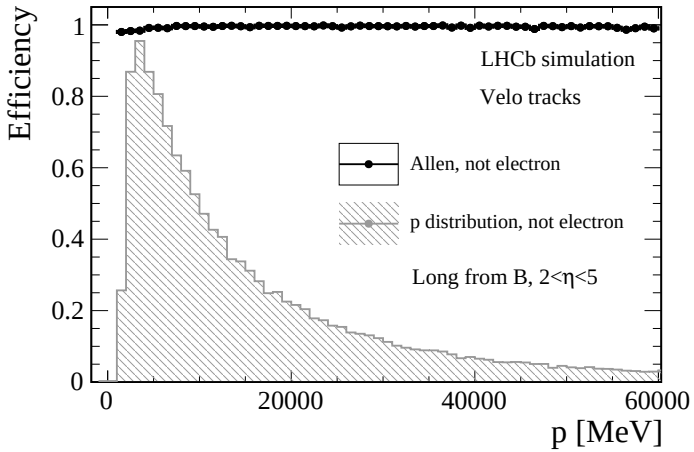

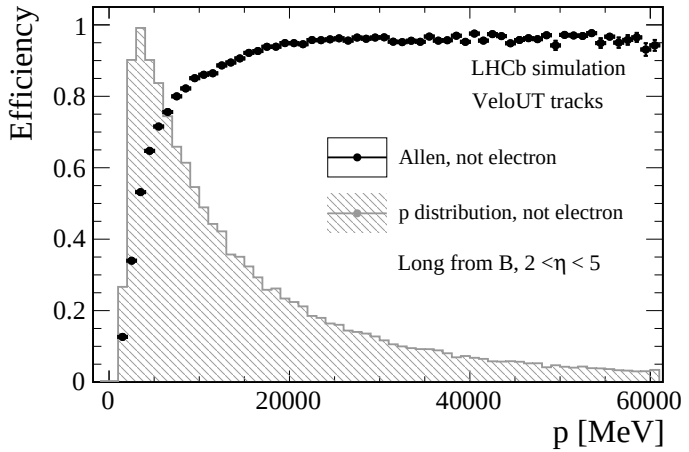

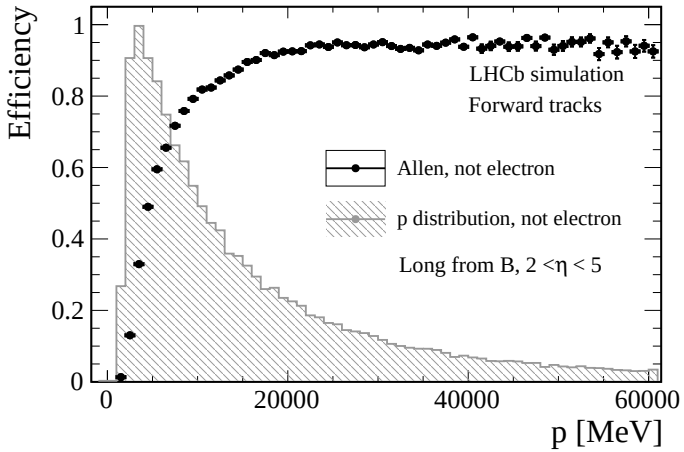

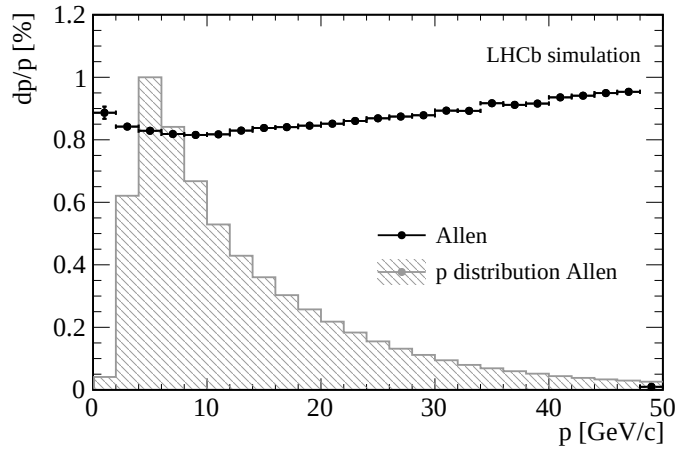

HLT1 muonID performance [LHCb-FIGURE-2020-016]

- Excellent muon identification and misID background rejection

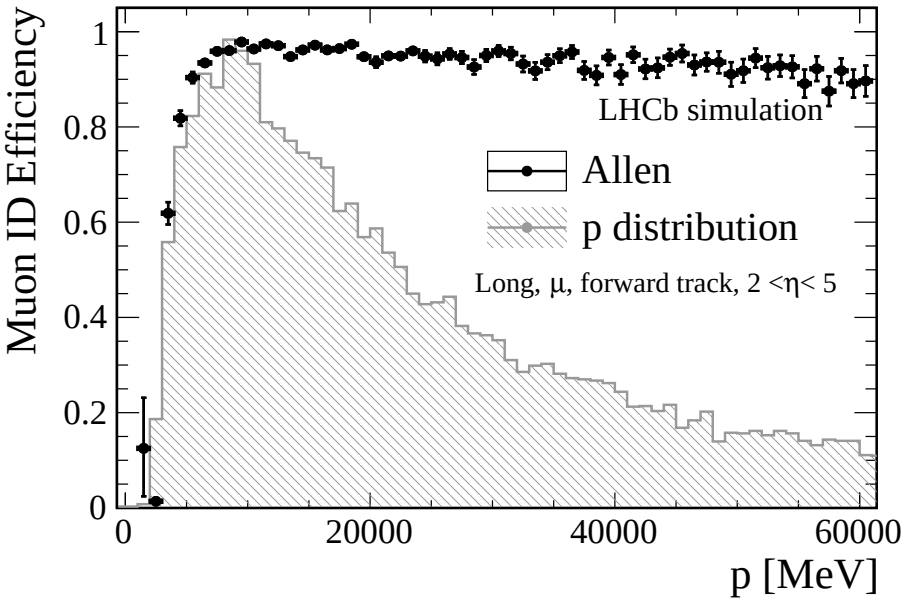

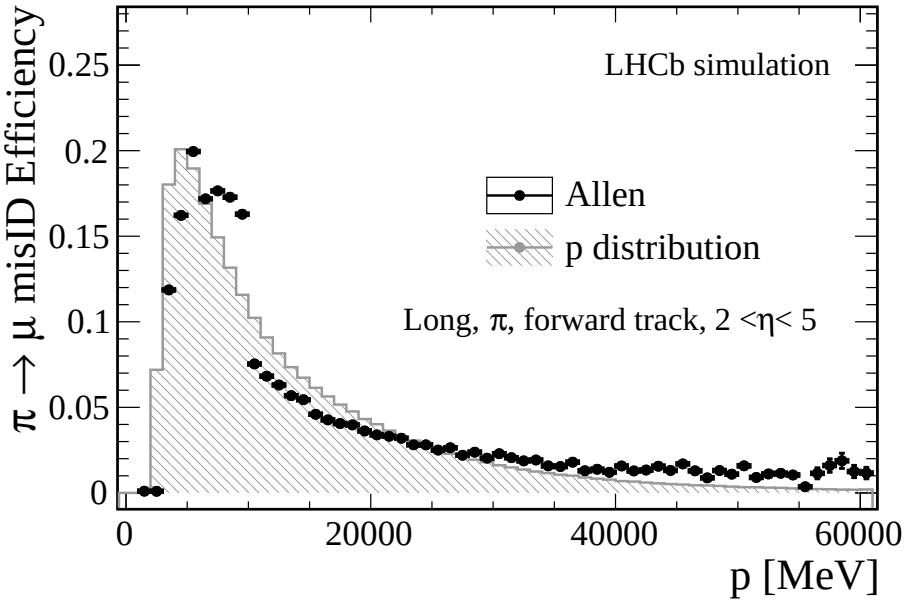

HLT1 selection performance [LHCb Upgrade GPU High Level Trigger TDR]

- Some of the selections still have to be tuned to reach a total output rate of about 1 MHz
- The main inclusive selections (TrackMVA and TwoTrackMVA) have reasonable output rates

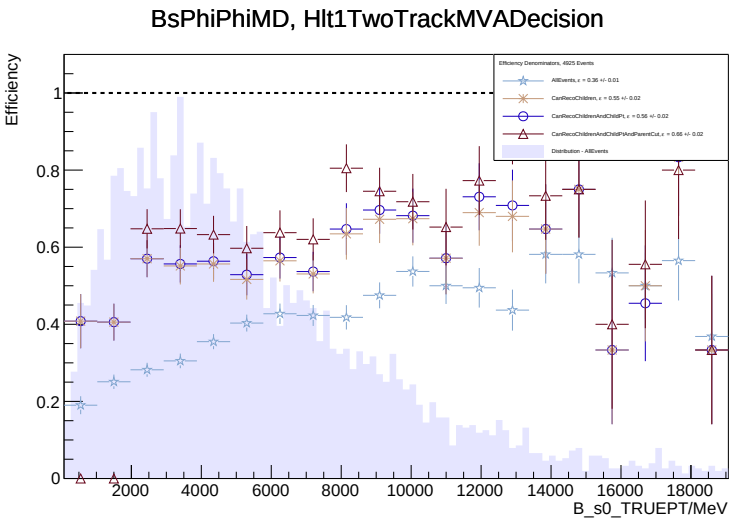

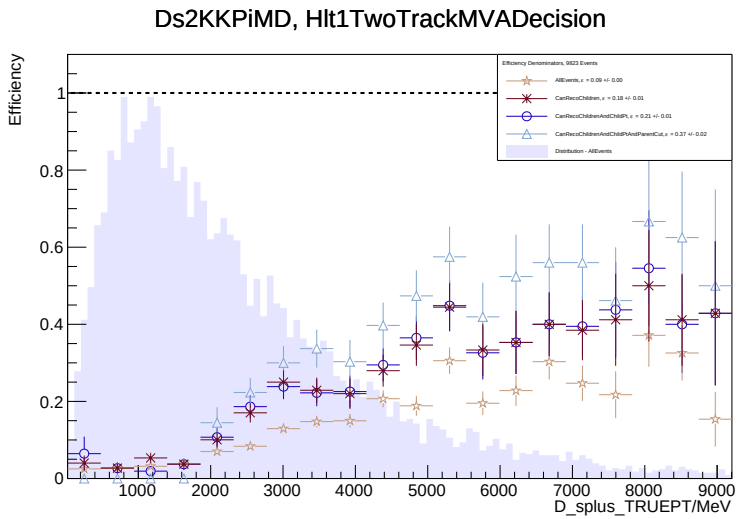

Conclusion

- LHCb upgrade ongoing, Run 3 to start in 2022
- Aim to increase instantaneous luminosity to $2 \times 10^{33} \text{cm}^{-2}\text{s}^{-1}$
- Major redesign of readout and trigger in LHCb upgrade
- Remove hardware L0 stage, read out full detector at 30 MHz non-empty bunch crossing rate
- LHCb is ready to face the challenge, also thanks to the HLT1 fully implemented on GPUs
- A lot of work ongoing to implement further algorithms and possibly extend the reconstruction to neutral and long-lived particles
- Stay tuned for the commissioning!!

Thanks for your attention!

Backup