Triggering TB/s of data: The LHCb perspective

Marianna Fontana,
on behalf of LHCb
CHEP conference, 19-25 October 2025,
Krakow

The LHCb experiment

- Experiment dedicated to flavour physics
- Successfully took 9 fb^{-1} of data during Run 1-2
- **Major upgrade** of all subtectors **for Run 3**
- Factor 5 increase in instantaneous luminosity → pile-up of 5

CERN-LHCC-2012-007

- 100% of the readout electronics replaced
- **New data acquisition system** and data center

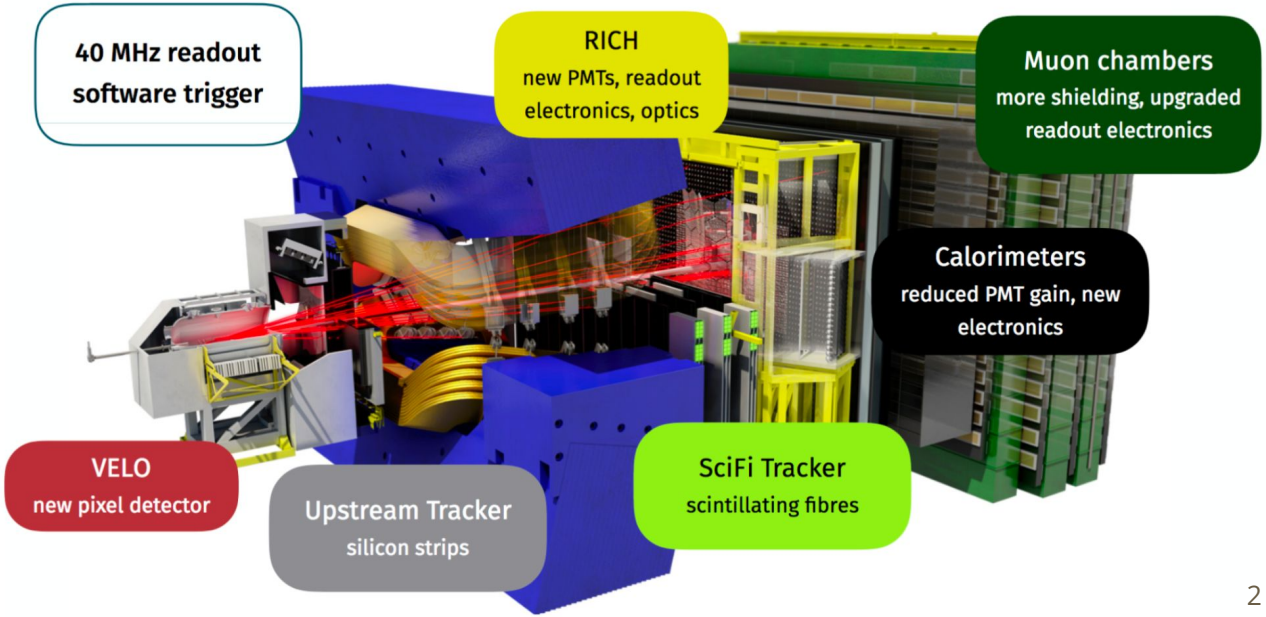

The trigger evolution: Run 1

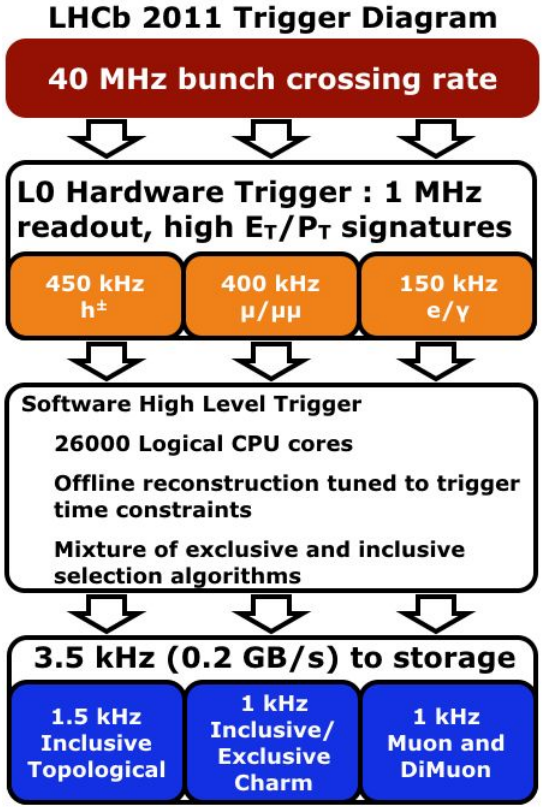

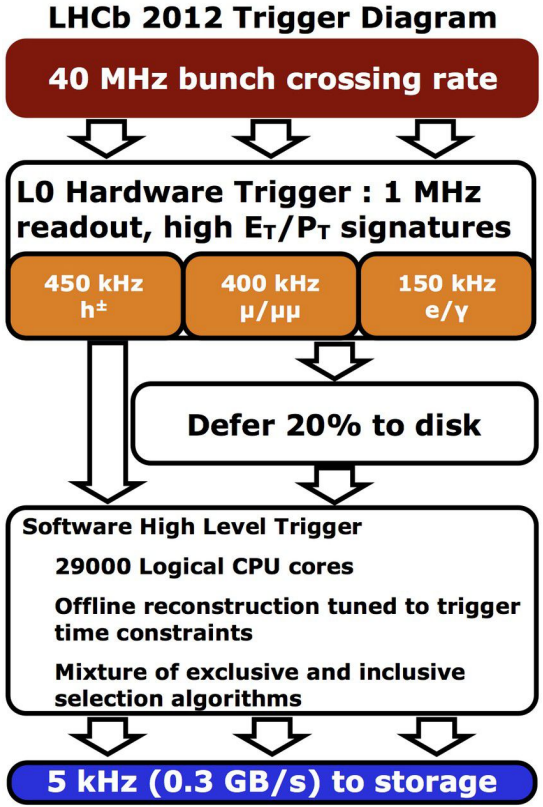

- L0 hardware level for high Et/pt signatures
- HLT1 running tracking (for high-pt) including Kalman filter
- HLT2 *almost* full event reconstruction
- Much bigger output rate than originally foreseen
  - Inclusive selections for full beauty programme
  - The charm programme initially not foreseen became a reality

The trigger evolution: Run 2

Comput. Phys. Commun. 208 (2016) 35-42
JINST 14 P04006

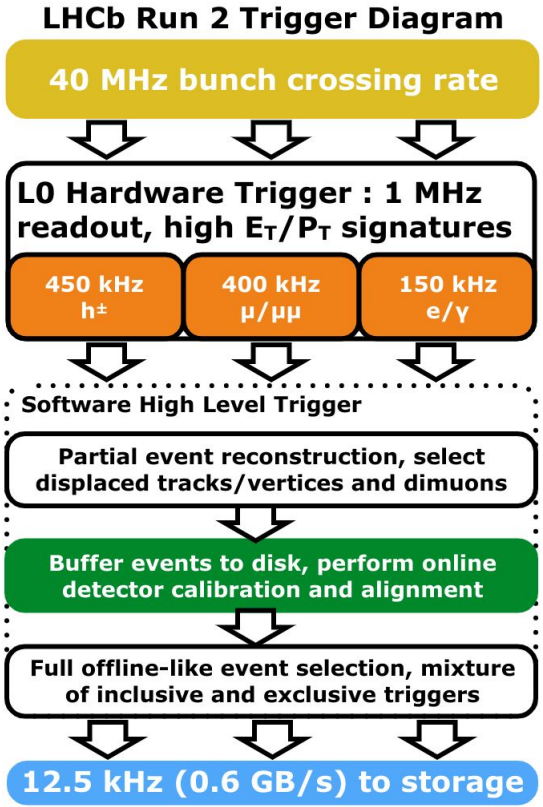

Disk buffer moved between HLT1 and HLT2 → increased number of CPUs and enabled

- Real-time alignment and calibration
- Real-time reconstruction with analysis quality reconstruction
- Ability to use trigger output for analysis and discard raw detector information in trigger (**Turbo stream**) [J. Phys.: Conf. Ser. 664 082004]
  - System fully commissioned already in 2015 with physics publications. It became the baseline for a good fraction of the Run 2 physics programme
- Adopted as the baseline approach for Run 3

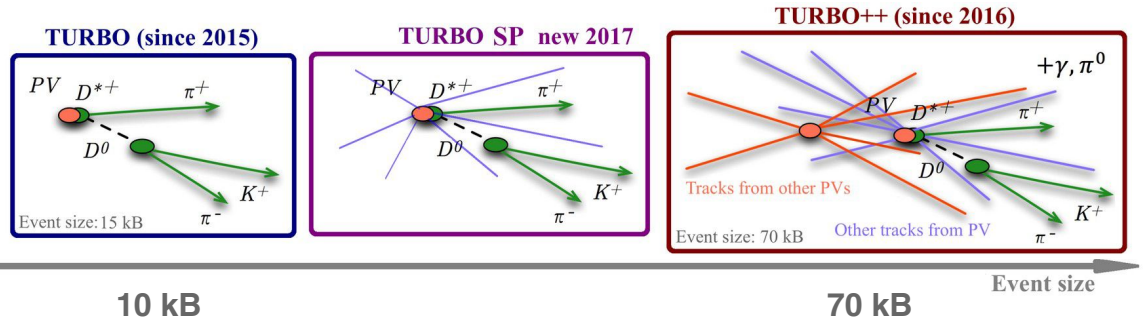

The trigger (R)evolution: Run 3

J. Phys.: Conf. Ser. **878** 012012

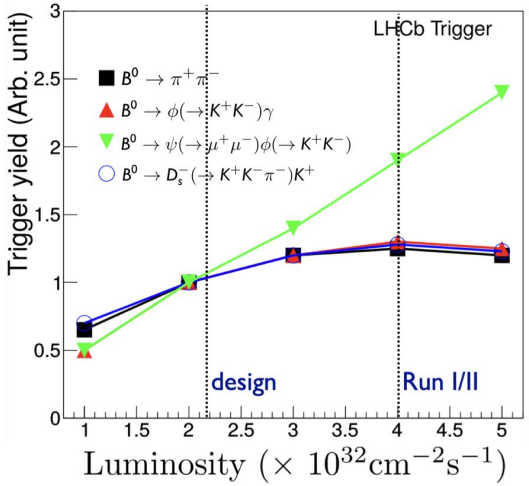

https://home.web.cern.ch/news

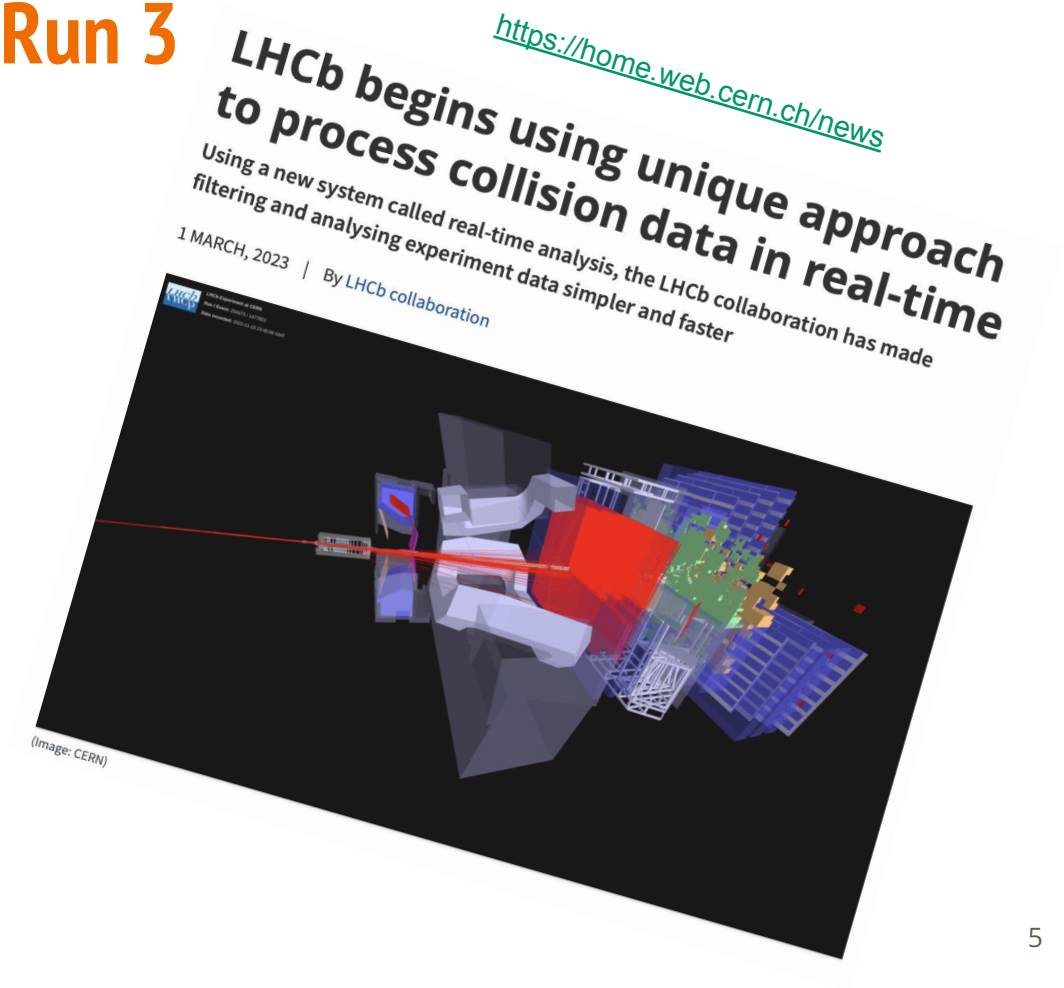

- In Run 1-2 couldn't efficiently trigger on heavy flavour using hardware signatures
- Trigger for many hadronic channels saturated
- Solution: **fully software trigger**

The Run 3 data flow

LHCb-FIGURE-2020-016

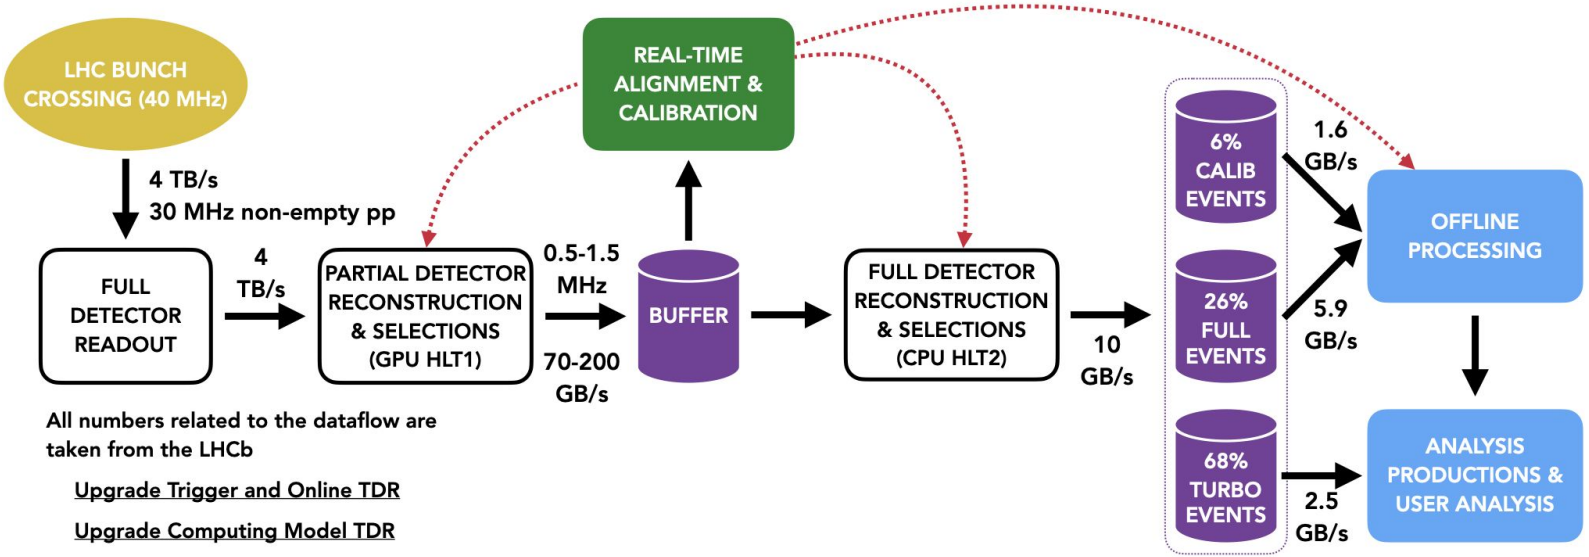

- Detector data @30 MHz received by O(500) FPGAs
- 2-stage software trigger, HLT1 & HLT2
- Real-time alignment & calibration
- After HLT2, 10 GB/s of data for offline processing

HLT1 trigger

See talks from
- A. Scarabotto
- J. Horsvill

- Take as input LHCb raw data (**4 TB/s**) at 30 MHz
- Perform partial event reconstruction & coarse selection to cover the full breadth of LHCb physics
- Reduce the input rate by a factor of 30 (~1 MHz)
- **~ 500 GPUs NVIDIA RTX A5000 GPUs** installed
  - The baseline TDR design could be achieved with 300 GPUs
  - Extra GPU power used to extend the improvements beyond-TDR

The GPU choice matches the DAQ architecture of LHCb

- GPUs can be hosted by the Event Builder Nodes via PCIe slots
- reduced costs due to shared powering and cooling and smaller network

HLT1 tasks are suited for parallelisation:

- Events can be treated independently
- Objects of reconstruction (tracks, vertices, ...) are independent

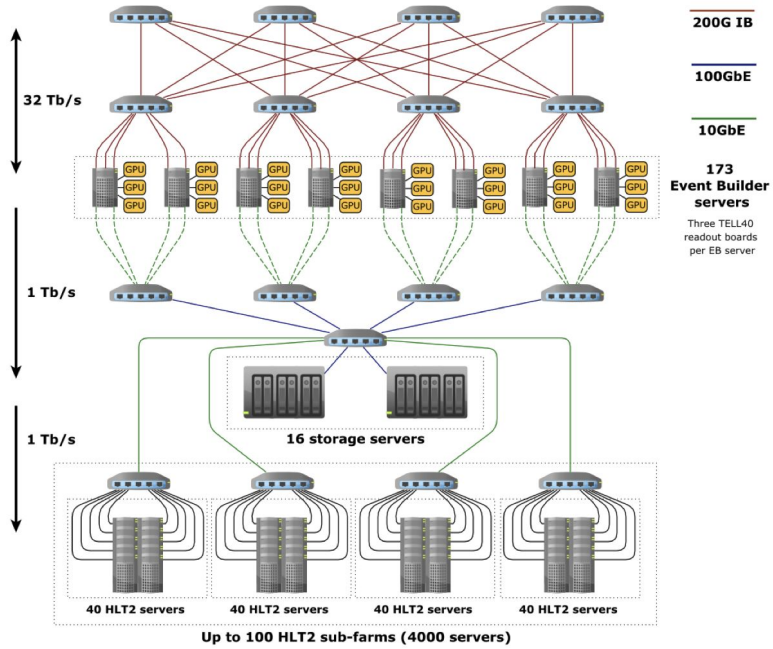

Comput.Softw.Big Sci. 6 (2022) 1, 1

Allen: LHCb HLT1 trigger

Comput. Softw. Big. Sci. **4**, 7 (2020)

Partial event reconstruction through

- Track reconstruction for all the track types used in physics analysis (Long and Downstream* tracks) [See talk by J. Zhuo]
- Vertex reconstruction
- Electron clustering* and bremsstrahlung recovery*
- Muon identification

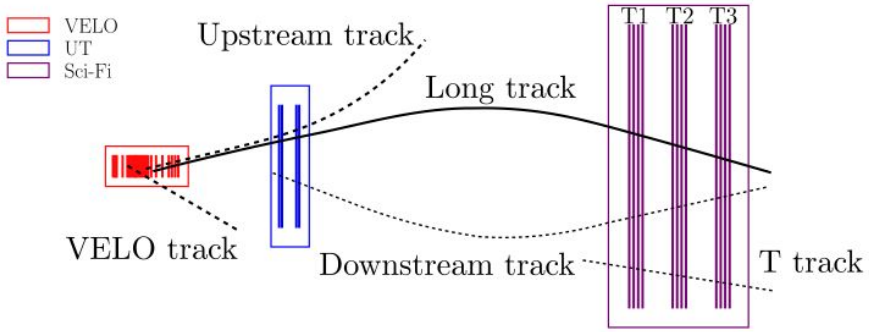

HLT1 Throughput per GPU

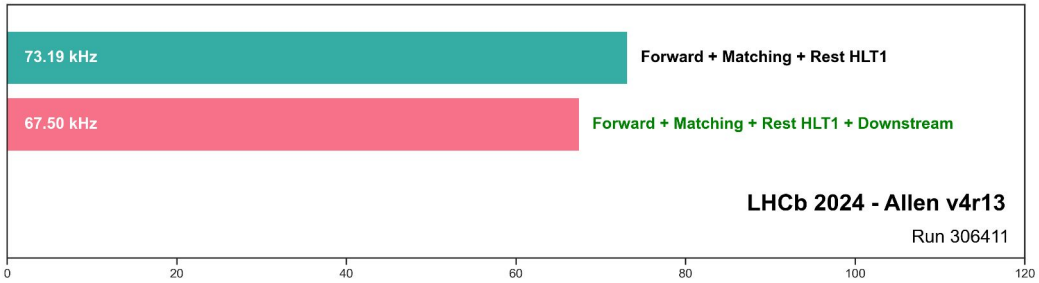

LHCb-FIGURE-2024-035

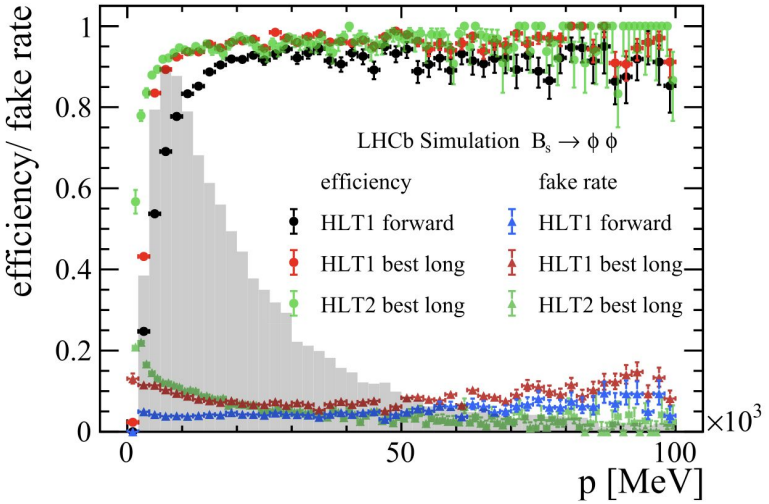

IEEE Access, vol. 12, 2024

* beyond TDR

HLT1 performance

- **The real-time analysis philosophy proved to be valid**
- Significant improvements in trigger efficiencies
- Huge gain a low-pT
  - Beneficial for the charm and strange physics programme
- Large impact for electron channels
- Muon channels gained from the removal of the global event cuts

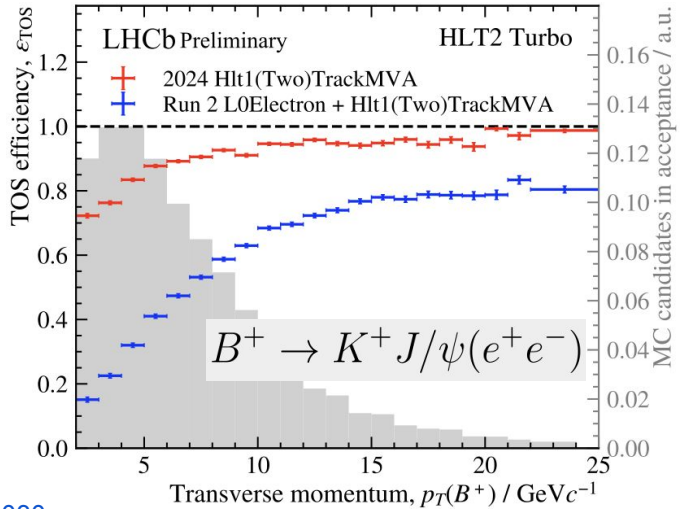

LHCb-FIGURE-2024-030

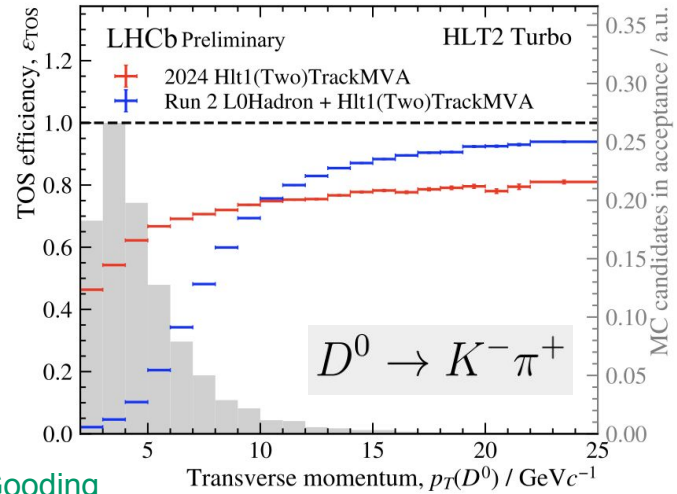

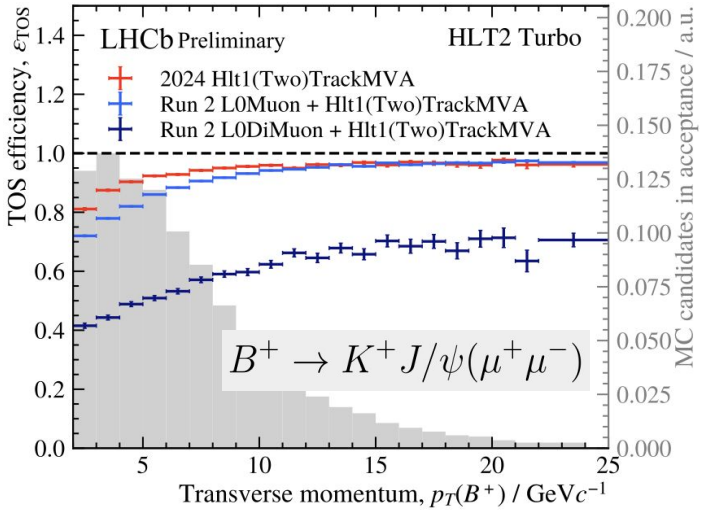

See talk by J. Gooding

Alignment and calibration

- Store data selected in HLT1 in intermediate buffer of O(30 PB) for real-time alignment and calibration
- Fully aligned and calibrated detector needed to have offline-quality reconstruction in HLT2
- Online alignment and calibration pioneered in Run 2, crucial in Run 3
- Two types of processes
  - Alignment: VELO, RICH mirrors, UT, SciFi, Muon
  - Calibration: RICH, ECAL, HCAL

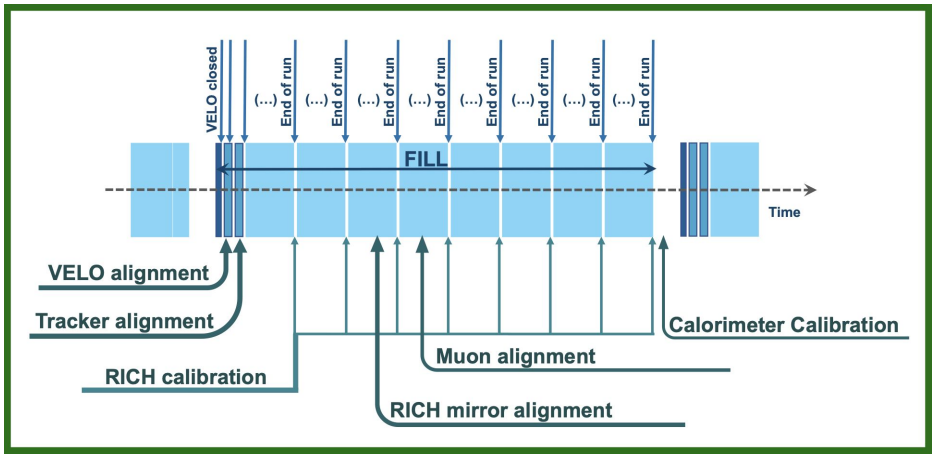

LHCb-FIGURE-2024-025

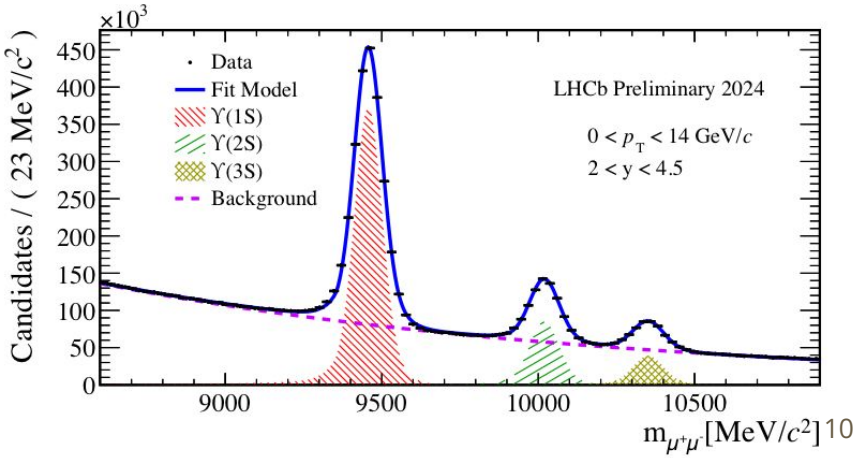

See talk by R. Hunter

LHCb HLT2 trigger

- HLT2 runs a full reconstruction and all the necessary selections (inclusive but mostly exclusive) for the wide LHCb physics programme (~3000 lines)
- Given the hard limit on bandwidth (10 GB/s to tape and 3.5 GB/s on disk) and expected signal rate, event size is the only free parameter
- Need to "persist" all the reconstructed objects for offline analysis
- The successful strategy of the Turbo paradigm used at full speed also in Run 3

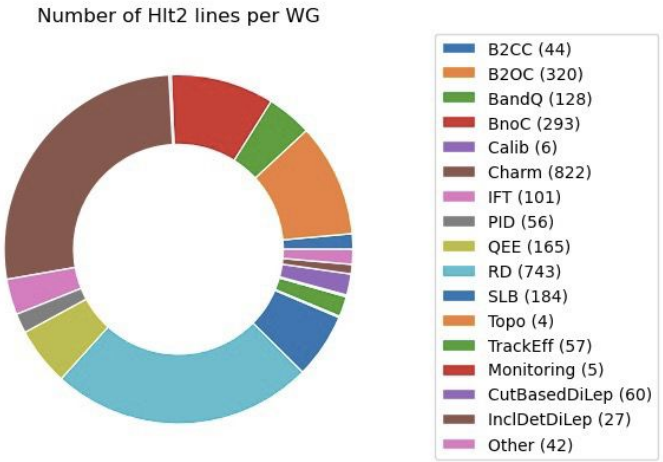

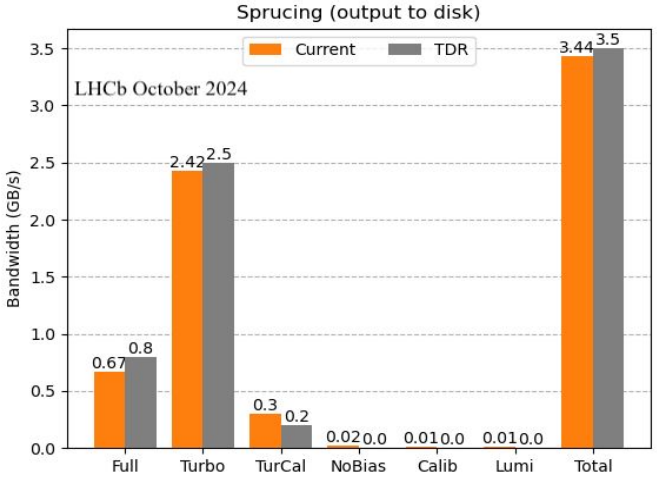

LHCb-FIGURE-2024-034

HLT2 performance

Achieving TDR performance for vertex resolutions, track reconstruction and PID performance

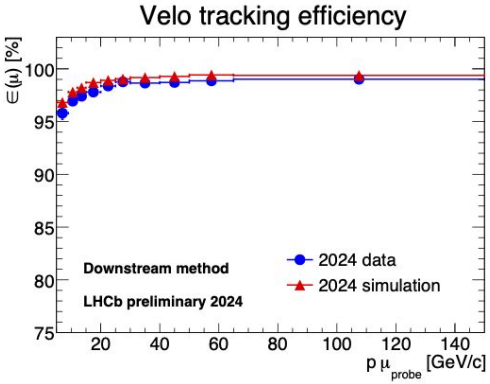

LHCb-FIGURE-2024-032

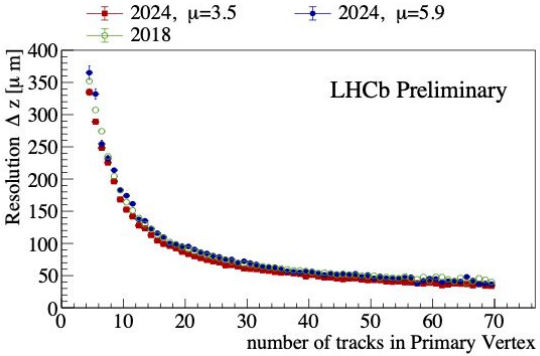

LHCb-FIGURE-2024-011

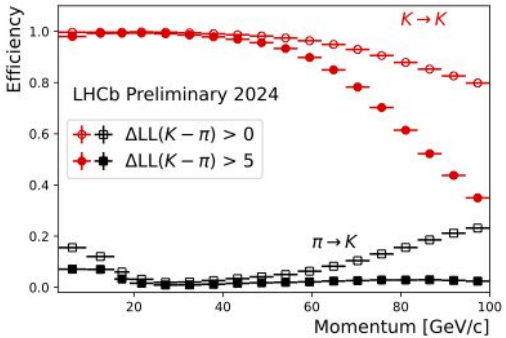

LHCb-FIGURE-2024-031

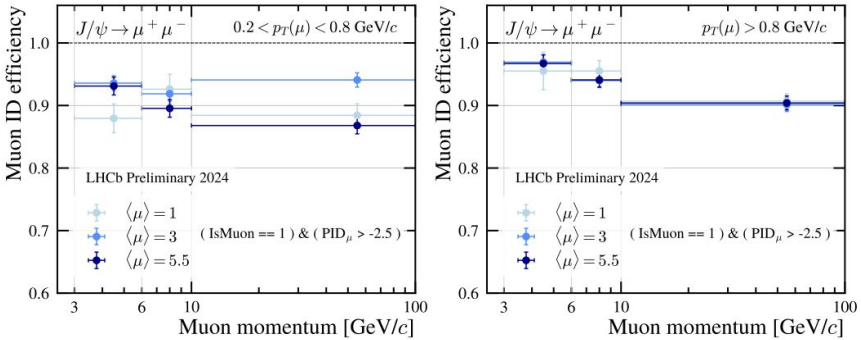

LHCb-FIGURE-2024-010

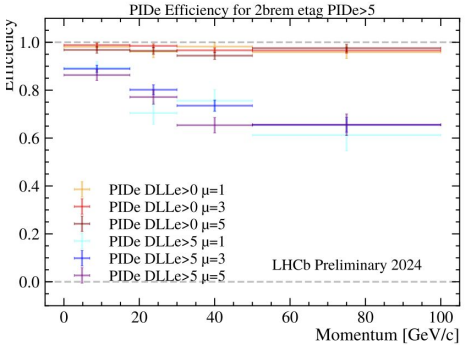

LHCb-FIGURE-2024-010

Towards the future

LHCb planning **Upgrade II for LS4**

- FTDR approved in March ‘22 and Scoping document in preparation
- Luminosity: $(2 \times 10^{33} \rightarrow 1.5 \times 10^{34})\ cm^{-2}\ s^{-1}$
- Pile-up: $5 \rightarrow 40$
- Exciting challenges in trigger and DAQ
  - 200 TB/s of data, to be processed in real time and reduced by ~4 orders of magnitude before sending to permanent storage
  - data processing will be based around pile-up suppression
  - 4D reconstruction: timing added to tracking and ECAL detectors to better isolate signals

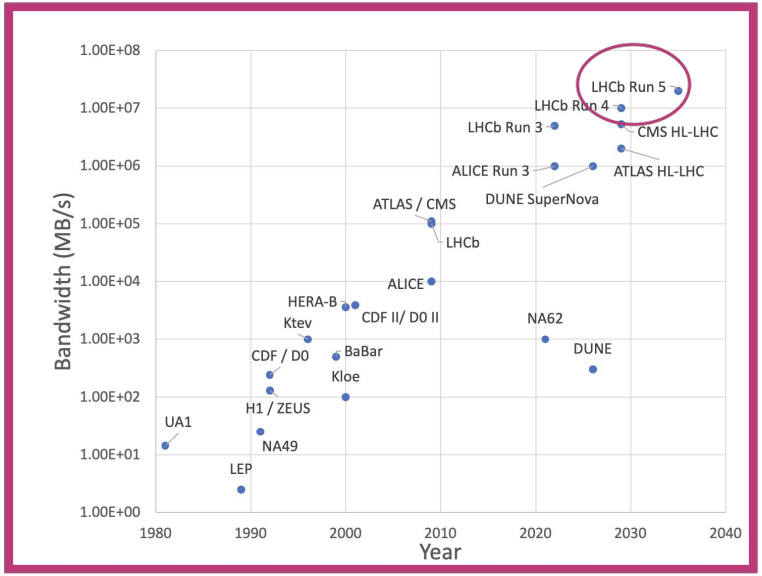

Courtesy of A. Cerri

The biggest data challenge in HEP!

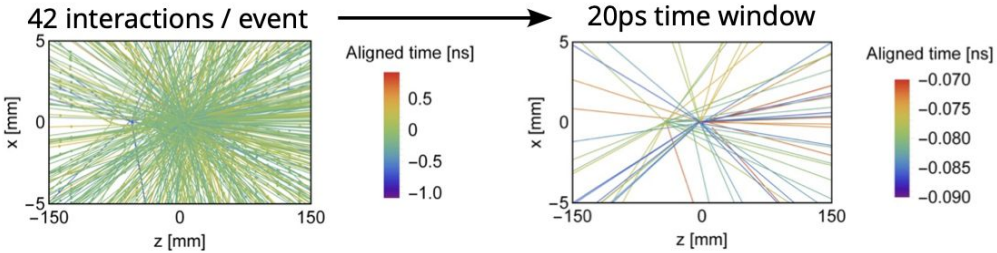

The trigger evolution: Run 5

See talk by F. Lazzari

- Triggerless design philosophy will remain correct and scalable
- Partial and full detector reconstruction (and selections?) both on GPUs
- Complementary R&D activities focusing on two main areas
  - Building subdetector primitives, for example tracks or calorimeter clusters, on FPGAs [LHCb-PUB-2024-001]
  - Exploiting other architectures such as the IPU or even more exotic hardware

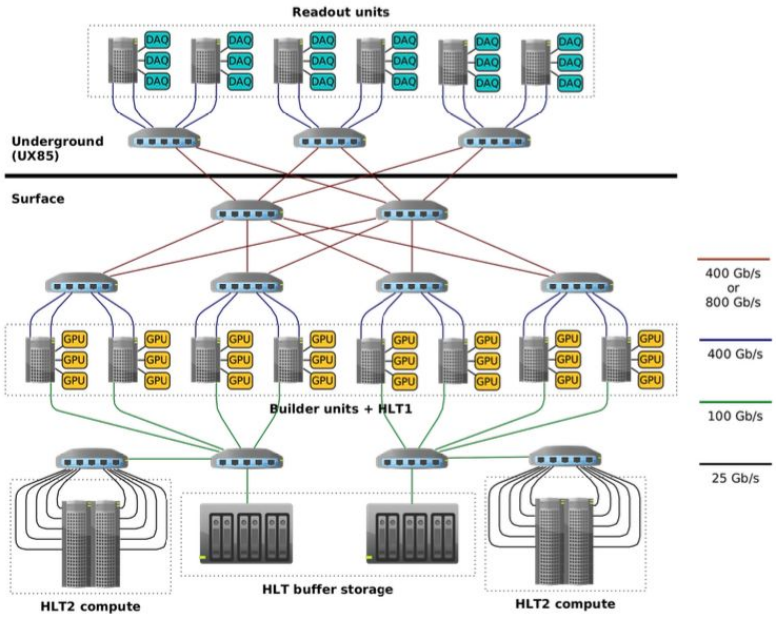

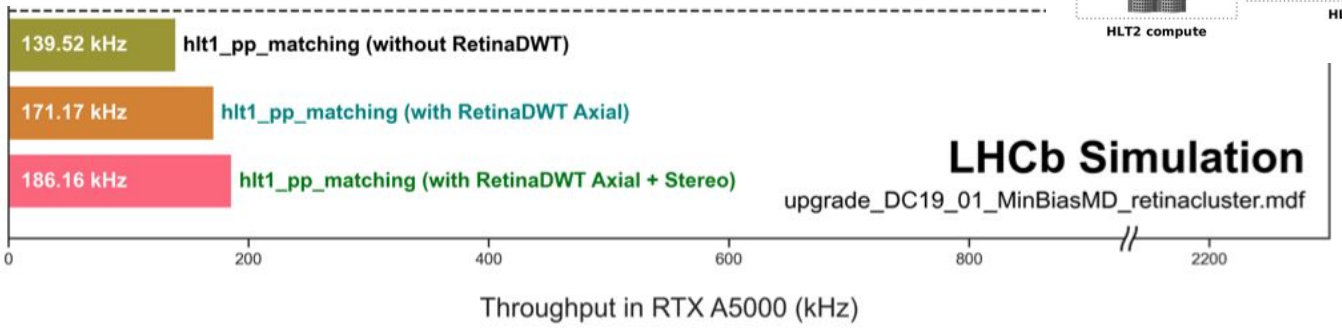

Conclusion

- LHCb underwent its first major upgrade in order to increase its instantaneous luminosity by x5
- Major changes in the trigger strategy:
  - Remove L0 hardware trigger, read-out full detector at 30 MHz
  - First level trigger run on GPUs
- The new trigger system has been successfully commissioned at nominal luminosity, even going beyond-expectations
- About 9.5 fb^{-1} of data have been taken and currently being analysed for a great physics outcome
- The LHCb Upgrade II is becoming a reality and this will pose very interesting challenges

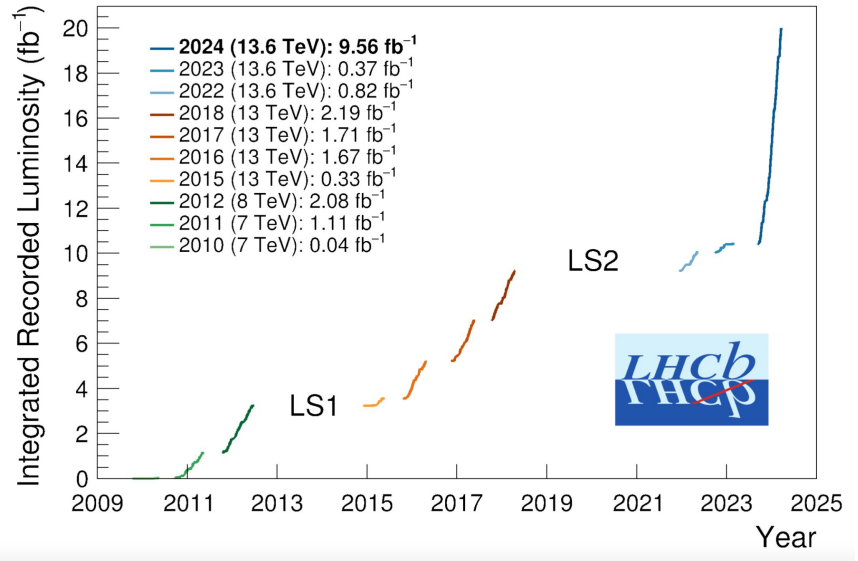

Thanks a lot for your attention!