

Real-time analysis from the trigger candidates and novel calibration strategy at the LHCb experiment

Giulio Gazzoni^{*†}

Université Clermont Auvergne, LPC Clermont, France

E-mail: giulio.gazzoni@cern.ch

The availability of computing resources is a limiting factor in data collection at the LHCb experiment, due to the high production rate of beauty and charm hadrons. For Run 2, LHCb has implemented a novel approach to make optimal use of these resources: the output of the first software trigger stage is buffered to disk and the second stage is executed asynchronously, using 100% of the available trigger farm even between LHC fills. As an integral part of the new strategy, the detector is aligned and calibrated, and the data are fully reconstructed, in real-time, permitting offline-quality signal selections. Without the need to reconstruct offline, it is possible to save only the information needed for analysis directly from the trigger. This concept, called the “LHCb Turbo stream”, maximizes the signal rate saved to disk. The analysis of the data collected also required appropriate calibration samples to determine the tracking and PID performance. A novel strategy has been introduced in Run 2, where the selection of calibration samples is implemented as a Turbo stream. A further processing of the data is required in order to provide background subtracted samples for the determination of performance, which is achieved through a centralized production that makes highly efficient use of computing resources. These data are also used in the development of new algorithms to evaluate the detector performance in LHC upgrade scenarios. The aim of this document is twofold. It covers the major steps of the implementation, and details the use of the calibration samples to determine the PID performance and tracking efficiency. It also presents how LHCb dynamically adapts the output rate of the first trigger stage to the LHC efficiency, and how the Turbo paradigm was extended in 2017 to dynamically persist any information required for analysis. This can range from only the signal candidate object to the complete event, with user-selectable granularity.

The European Physical Society Conference on High Energy Physics

5-12 July, 2017

Venice

^{*}Speaker.

[†]On behalf of the LHCb Collaboration.

1. Introduction

The LHCb detector [1] at the Large Hadron Collider (LHC) is a single-arm spectrometer with a forward geometry covering a pseudorapidity range of about $2 < \eta < 5$. Designed to perform precision measurements in decays of b and c hadrons, it is composed by several sub-detectors that form a tracking system and a particle identification system.

The tracking system is designed to identify the primary and secondary vertices, reconstruct the particles trajectories and measure their momentum. It consists of a silicon-strip Vertex Locator (VELO) [2] surrounding the pp interaction region, a large-area silicon strip detector (TT), a dipole magnet and three Tracking Stations (T1-T3) placed downstream the magnet and composed of an Inner [3] and an Outer [4] tracker. The particle identification system, that provides the identification of photons, leptons and hadrons, is composed by two Ring-Imaging Cherenkov detectors (RICH1 and RICH2) [5] placed upstream and downstream the magnet, the electromagnetic (ECAL) and hadronic (HCAL) calorimeters [6] and by five muon stations [7] (MUONS).

The online event selection is performed by a trigger system [8], that consists of a hardware stage (L0), based on information from the calorimeter and muon systems and which reduces the event rate to below 1 MHz, followed by a two software stages (HLT1 and HLT2) that perform a full event reconstruction and its preliminary selection.

The physics goals of LHCb are achieved thanks to the excellent performances [9] of the tracking and particle identification systems, which in turn crucially rely on the spatial alignment and calibration of the various sub-detectors. Thanks to the implementation of major changes in the LHCb trigger for Run 2, such alignment and calibration procedures are now performed automatically and in real-time [15]. As discussed hereafter, with this new approach it was possible to redesign and improve the LHCb data taking strategy as well as the computing strategy for producing the calibration samples.

2. Evolution of the LHCb trigger

During the first long shutdown the LHCb Event Filter Farm (EFF) was almost doubled (~ 52000 logical CPU cores compared to ~ 29000 in Run 1) and the High Level Trigger software was improved, allowing a deep revisit of the trigger strategy for Run 2.

After the L0 hardware trigger stage events are sent to the HLT1, which performs a partial event reconstruction using information from the VELO and tracking stations (TT, T1-T3). If at least one track satisfies strict quality, separation from any primary vertex and transverse momentum criteria, the event is passed to the second level of the software trigger.

At the HLT2 stage the full event reconstruction is performed, and events passing this step are sent to storage at a rate of 5kHz in Run 1 and 12.5 kHz in Run 2. However, in order to satisfy the time constraints, the HLT event reconstruction in Run 1 was simpler and quicker than used offline. Additionally it did not use the latest alignment and calibration constants which were calculated later and used in the offline data reconstruction [10].

In order to better exploit the computing resources during the LHC downtime periods and to reduce the fraction of time that the farm spent in idle, in 2012 $\sim 25\%$ of the events passing the L0 stage were buffered to disk and sent later to the HLT software to be processed.

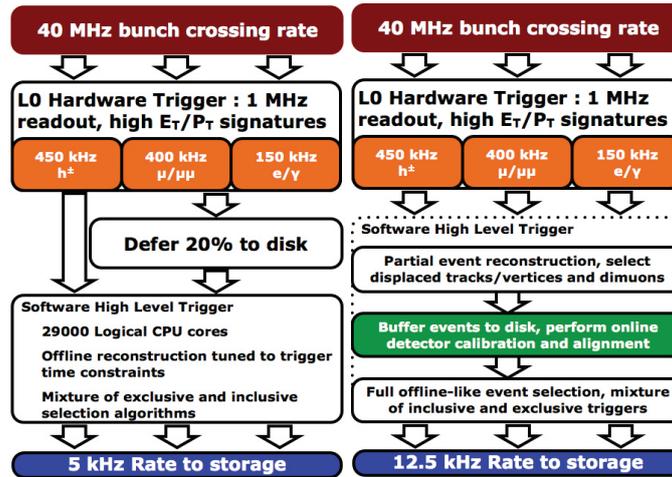


Figure 1: Schematic diagram of the LHCb trigger data-flow in Run 1 data taking (left) compared to the data-flow in Run 2 (right). Software High Level Trigger indicates HLT1 and HLT2 stages.

In Run 2 the L0 trigger remains the same, while HLT1 and HLT2 become two independent softwares executed asynchronously. This is achieved by sending all the events passing the HLT1 stage to a 10 PB buffer and deferring the HLT2 step, that now can be executed later. The comparison between Run 1 and Run 2 trigger layouts is shown in Fig. 1. With this new configuration part of the events out of the HLT1 are processed within few minutes after the LHC fills and taken as input to update the constants needed to align and calibrate the various sub-detectors. This real-time alignment and calibration procedure is fully automated and performed run-by-run at regular intervals depending on needs. The high-quality calibrations are then implemented before sending events to the HLT2, that now shares the same constants used in the offline reconstruction. This is very important, especially for the hadron identification, as it allows to achieve the best performances directly in the trigger.

With the total resources allocated in the computing farm, HLT1 events typically require ~ 35 ms to be processed, while HLT2 events need ~ 650 ms.

3. The Turbo Stream

Thanks to the changes in the trigger configuration, the LHCb data processing model significantly evolved and new concepts were implemented. The scheme of the LHCb data workflow for Run 2 is shown in Fig. 2.

In the traditional data flow, raw event data undergoes a complete reconstruction using the sub-detector raw data banks, which contain the detector information. This additional reconstruction was designed for a data processing model in which final calibrations were significantly improved compared to the calibrations initially available. After the offline reconstruction, selection criteria based on typical b -hadron and c -hadron topologies are applied to identify the decay candidates for user analysis. After a final merging step, the datasets are ready for physics analyses [11].

In Run 2 candidates passing the HLT2 stage have an offline quality level and do not need a further processing step to be analyzed: such events can then be directly saved removing the sub-

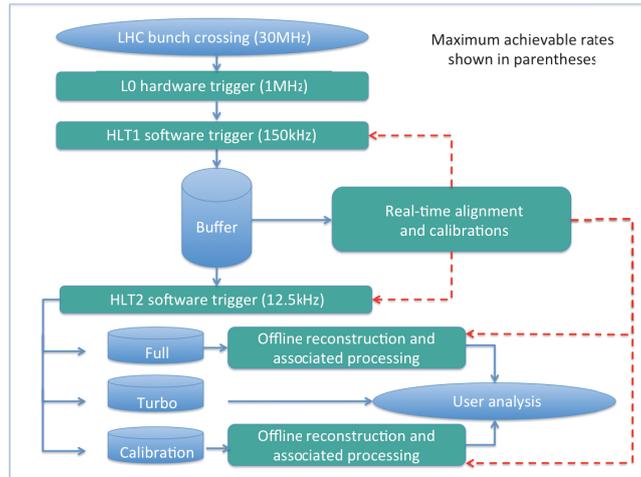


Figure 2: Schematic diagram showing the overall data processing model in Run 2, where the blue solid line represents data flow, and the red dashed line the propagation of calibrations.

detector information. This new implementation is called Turbo Stream [12] and exploits 20% of the HLT2 bandwidth, namely 2.5 kHz. The other 10 kHz (Full stream and Calibration stream) of the HLT2 output undergo the traditional data processing flow. The Turbo Stream has some considerable advantages; first of all it requires less storage space because a part of the raw event is discarded and secondly it provides a much faster output (< 1h against ~ 30 h for events following the traditional Full stream). This way data can be accessed earlier to perform physics analyses.

The former Turbo stream implementation allows only the objects selected by the trigger to be saved, and the output is limited to a reduced set of variables. In 2016 this paradigm was extended in order to save not only the trigger candidates but also any reconstructed objects or any additional variables such as isolation or hits in a cone region. The goal of this implementation called Turbo++ (see Fig. 3) is to choose how many, and indeed which, variables of the events should be saved according to the physics channel and desired measurement. The Turbo Stream can then be used by more physics analyses.

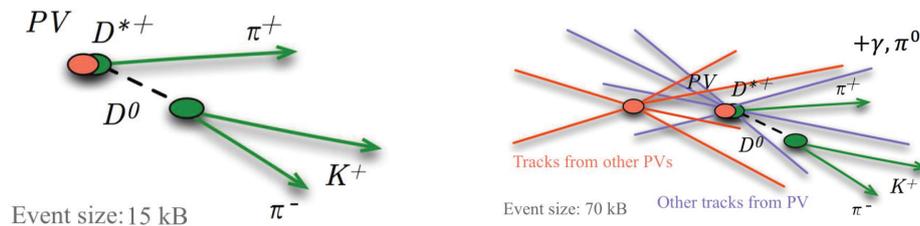


Figure 3: Illustrative sketch of the 2015 Turbo (left) and 2016 Turbo++ (right) implementations.

For 2017 data taking a new intermediate solution for Turbo has been developed. This implementation, called Turbo SP (Selective Persistence), allows to save candidates and a subset of the reconstruction, like for example additional particles near the primary vertex. Thanks to the possibility of choosing what to persist, this approach is very flexible and the variety of physics analyses that can be performed on the Turbo stream increases. In addition, compared to Turbo++, this is on average less demanding in terms of storage space. In the Turbo SP implementation a utility to help select the interesting particles is also foreseen.

4. Calibration samples

In order to perform a complete physics analysis it is necessary to evaluate the PID and Tracking efficiencies with a data-driven approach using dedicated high statistics calibration samples. This step is needed because the simulation does not reproduce variables with sufficient accuracy. During Run 1 the calibration samples were produced with offline selections, giving rise in some cases to lack of statistics and poor coverage in some phase-space regions.

In Run 2 the production strategy of the calibration samples radically changed exploiting the Turbo Stream implementation. The samples are now selected out of the HLT2 and saved in a dedicated stream called TurboCalib [13], allowing the calculation of PID related variables for both online (Turbo) and offline (Full) streams. The stream is then processed centrally to perform the matching of online and offline candidates. With a further processing step the matched candidates are used to produce background subtracted samples that can be exploited by the final user in order to extract the real efficiency for a given PID cut.

Run 2 charged PID calibration samples contain hundreds of millions of events of high-purity candidates and, thanks to the new production strategy, have a significantly improved phase-space coverage. As an example, a comparison of the distributions in the (p, η) plane for protons from Run 1 and Run 2 is shown in Fig. [14]. In 2017 the TurboCalib implementation has been extended to provide analogous calibration samples also for neutral objects (see Tab. 1), *i.e.* π^0 and γ .

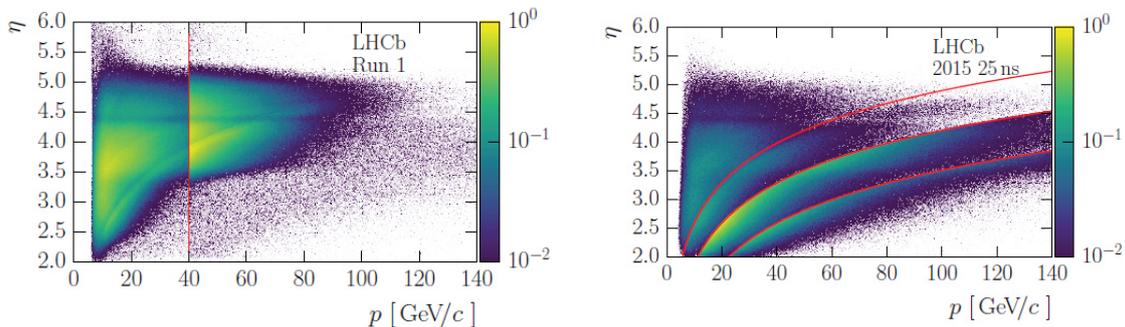


Figure 4: Comparison of distributions between protons from (left) Run 1 and (right) Run 2 calibration samples in the (p, η) plane. The plots are normalised and the red lines separate samples selected by different trigger lines.

| Species | Soft p_T | Hard p_T |
|-----------|---|---|
| e^\pm | – | $B^+ \rightarrow [e^+e^-]_{J/\psi} K^+$ |
| μ^\pm | $D_s^+ \rightarrow [\mu^+\mu^-]_\phi \pi^+$ | $J/\psi \rightarrow \mu^+\mu^-, B^+ \rightarrow [\mu^+\mu^-]_{J/\psi} K^+$ |
| π^\pm | $K_S^0 \rightarrow \pi^+\pi^-$ | $D^{*+} \rightarrow [K^-\pi^+]_{D^0} \pi^+$ |
| K^\pm | $D_s^+ \rightarrow [K^+K^-]_\phi \pi^+$ | $D^{*+} \rightarrow [K^-\pi^+]_{D^0} \pi^+$ |
| p^\pm | $\Lambda^0 \rightarrow p^+\pi^-$ | $\Lambda^0 \rightarrow p^+\pi^-, \Lambda_c^+ \rightarrow p^+K^-\pi^+, \Lambda_b^0 \rightarrow [p^+K^-\pi^+]_{\Lambda_c^+} \mu^-\bar{\nu}_\mu$ |
| π^0 | – | $D^{*+} \rightarrow [K^-\pi^+\pi^0]_{D^0} \pi^+$ |
| γ | $D_s^{*+} \rightarrow [K^+K^-\pi^+]_{D_s^+} \gamma, D_{(s)}^+ \rightarrow [\rho^0\gamma]_{\eta'} \pi^+$ | $B^0 \rightarrow [K^+\pi^-]_{K^*} \gamma, B_s^0 \rightarrow [K^+K^-]_\phi \gamma$ |

Table 1: Summary of the modes collected in the calibration samples, divided by species and p_T regimes.

The TurboCalib stream developed for the PID calibration samples has also been used to provide the samples needed to extract the real tracking reconstruction efficiency. More details about the LHCb tracking reconstruction in Run 2 can be found in [15].

5. Conclusions

The LHCb computing model significantly evolved during the first long shutdown and the high level trigger underwent important updates. This allowed a deep revisit of the LHCb trigger strategy for Run 2 and to implement a fully automated real-time alignment and calibration procedure based on the new trigger features.

A novel approach exploiting such a procedure was implemented. This new concept, called Turbo stream, allows to save candidates directly after the HLT2 stage without any further reconstruction and with an offline quality level. A more flexible variation of this paradigm has also been developed for 2017 data taking. This extension, the Turbo SP (Selective Persistence), allows to save candidates out of the HLT2 together with a subset of the reconstruction with user-selectable granularity.

The Turbo stream concept was also exploited to revisit the production strategy of PID and tracking calibration samples thanks to the TurboCalib stream implementation. This allowed to speed up the process and to guarantee samples with higher statistics and better kinematic coverage.

References

- [1] A. A. Alves Jr. *et al.*, The LHCb Detector at the LHC, *JINST* **3** S08005 (2008).
- [2] LHCb Collaboration, LHCb VELO (VVertex LOcator) : Technical Design Report, *CERN-LHCC-2001-0011* (2001).
- [3] LHCb Collaboration, LHCb inner tracker : Technical Design Report, *CERN-LHCC-2002-029* (2002).
- [4] R. Arink *et al.*, Performance of the LHCb Outer Tracker, *JINST* **9** P01002 (2014).
- [5] M. Adinolfi *et al.*, Performance of the LHCb RICH detector at the LHC, *Eur. Phys. J.* **C73** 2431 (2013).
- [6] LHCb Collaboration, LHCb calorimeters: Technical Design Report, *CERN-LHCC-2000-036* (2000).
- [7] A. A. Alves Jr. *et al.*, Performance of the LHCb muon system, *JINST* **8** P02022 (2013).

- [8] R. Aaij *et. al.*, The LHCb trigger and its performance in 2011, *JINST* **8** P04022 (2013).
- [9] R. Aaij, *et. al.*, LHCb detector performance, *Int. J. Mod. Phys. A* **30** 1530022 (2015).
- [10] B. Sciascia, LHCb Run 2 trigger performance, *LHCb-PROC-2016-020* (2016).
- [11] R. Aaij *et. al.*, Tesla : an application for real-time data analysis in High Energy Physics, *Computer Physics Communications* **208** 35-42 (2016).
- [12] S. Benson *et. al.*, The LHCb Turbo Stream, *J. Phys. Conf. Ser.*, **664** 082004 (2015).
- [13] L. Anderlini *et. al.*, Computing strategy for PID calibration samples for LHCb Run 2, *LHCb-PUB-2016-020* (2016).
- [14] O. Lupton *et. al.*, Calibration samples for particle identification at LHCb in Run 2, *LHCb-PUB-2016-005* (2016).
- [15] A. Dziurda, Performance and recent developments of the real-time track reconstruction and alignment of the LHCb detector, in these proceedings.