The LHCb Trigger System: Present and Future

2015 J. Phys.: Conf. Ser. 623 012003

(http://iopscience.iop.org/1742-6596/623/1/012003)

The LHCb Trigger System: Present and Future

Johannes Albrecht
TU Dortmund University
On behalf of the LHCb Collaboration

E-mail: `johannes.albrecht@tu-dortmund.de`

Abstract. The current LHCb trigger system consists of a hardware level, which reduces the LHC inelastic collision rate of 30 MHz to 1 MHz, at which the entire detector is read out. In a second level, implemented in a CPU farm, the event rate is reduced to about 5 kHz. The major bottleneck in LHCb's trigger efficiencies for hadronic heavy flavour decays is the hardware trigger. The LHCb experiment plans a major upgrade of the detector and DAQ system in the LHC shutdown of 2018. In this upgrade, a purely software based trigger system is being developed, which will have to process the full 30 MHz of inelastic collisions delivered by the LHC. Both the current trigger system and its planned upgrade are discussed in these proceedings.

1. Introduction

The LHCb detector at the LHC is a precision experiment dedicated to beauty and charm physics, covering a rapidity range of $2 < \eta < 5$ [1]. Ref. [2] summarizes the performance of the LHCb detector in Run 1 (2010 – 2012). The trigger system during Run 1 [3, 4] consisted of a fixed latency hardware level which reduced the visible bunch crossing rate to 1 MHz. At this rate, the whole detector was read out and a flexible software High-Level Trigger (HLT) was run to further reduce the rate to about 5 kHz, which was written to offline storage. This configuration allowed LHCb to record the largest beauty and charm hadron samples, at a very high signal purity.

For the second data taking period (Run 2, 2015 – 2018), the detector will be mostly unchanged with respect to Run 1, while the trigger software will undergo an incremental update. The software trigger will be split in two separate instances, the first one running synchronous with the LHC collisions while the second stage will be run asynchronously. This split allows to to perform a real-time calibration and alignment of the detector. The software trigger is thus able to select events based on a reconstruction with a quality almost identical to the offline processing.

The LHCb experiment plans major upgrades in preparation for the third LHC data taking period (Run 3, 2019++). As part of the detector upgrade, all of the front-end devices will be replaced in order to read the detector out at the full bunch crossing rate of 40 MHz [5]. Removing the readout bottleneck allows the complete event reconstruction at the full collision rate and hence an unprecedented trigger performance, especially for hadronic modes.

This note discusses the LHCb trigger system and its performance in Run 1 as well as the prospects and preparations for Run 2 and 3.

2. Run 1: Trigger Performance in 2009 – 2012

The trigger system of the LHCb experiment in Run 1 consists of two levels: the first level, implemented in hardware (L0), and the High Level Trigger (HLT), implemented in a CPU farm of about 29 000 logical cores. The LHCb trigger system and its performance during 2011 data taking has been described in detail in Ref. [4]. Here, the adjustments of the system for 2012 and its performance are discused. The LHCb detector was originally designed for an instantaneous luminosity of $\mathcal{L} = 2 \cdot 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. During the running period of 2012, this design luminosity was doubled to $\mathcal{L} = 4 \cdot 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, which was possible thanks to the excellent performance of both the detector hardware and the trigger system. In these conditions, the average number of visible interactions per bunch crossing is $\mu = 1.6$. One novel feature of the Run 1 trigger was that of event deferral: 20% of all L0 accepted events were stored on disk to be processed by the HLT during the LHC inter-fill time, which made an effective 25% of extra CPU available.

The trigger efficiency is evaluated using events that are reconstructed using the full offline software and selected with the final analysis selection for the respective channel. Thus, the trigger efficiency contains only the additional inefficiency due to simplifications used in the trigger, possible alignment inaccuracies, worse resolution than the offline reconstruction or harder cuts imposed by rate and processing time limitations. The method to evaluate the trigger performance and its implementation are discussed in Refs. [4, 6].

The rate of visible collisions in LHCb is about 13 MHz, which then is reduced by the L0 trigger to 1 MHz, at which the full detector is read out. The L0 trigger is implemented in custom hardware and has a latency of $4\,\mu$s. It triggers on high transverse momentum (p_T) muons and on large transverse energy (E_T) deposition in the calorimeter.

The High Level Trigger consists of two stages, HLT1 and HLT2. The first stage, HLT1, performs a partial event reconstruction and an inclusive selection of signal candidates. At the reduced rate of 80 kHz, HLT2 performs a full event reconstruction with only minor adjustments as compared to the offline reconstruction sequence. After this reconstruction, a set of inclusive and exclusive selections reduces the trigger rate to 5 kHz, which are saved for later offline analysis. The rates discussed above are average rates from the 2012 run of the LHC, in 2011 the HLT1 output rate was approximately 40 kHz and the HLT2 output rate was 3 kHz.

2.1. Hadron trigger performance

The L0 hardware trigger performance for hadronic modes is shown in Fig. 1. It is most efficient on two prong beauty decays (∼40%), and least for four prong charm decays (∼20%). The other modes lie in between.

Hadronic signals are generally selected in the HLT most efficiently by the inclusive beauty and charm trigger lines for both levels of the HLT. The HLT1 line selects good quality track candidates based on their p_T and displacement from the primary vertex. This trigger line gets the dominant part of the HLT1 bandwidth allocated. It is the dominant trigger line for most physics channels that do not contain leptons in the final state. The performance of HLT1 for hadronic signatures is shown in Fig. 2 as a function the p_T of the decaying charm or beauty hadron.

A multivariate selection is used to trigger B decays into charged hadrons in an inclusive selection based on two- three- and four-prong vertices. These trigger lines are based on a BDT classifier that uses discretized input variables [7] which ensures a fast and robust implementation. A crucial input to the BDT is the corrected mass, which takes the missing momentum transverse to the direction of flight into account. This allows the trigger to select heavy flavour decays even when some final state particles are not reconstructed. Fig. 3 shows the efficiency for the inclusive HLT2 trigger lines on a set of hadronic beauty decays.

2.2. Lepton trigger performance

The L0 hardware trigger selects muons based on their track segment reconstructed in the Muon chambers. The corresponding efficiency is shown in Fig. 4 for the decay $B^- \to J/\psi K^-$ with J/ψ ($\mu^+\mu^-$). The single muon trigger contributes the dominant part to the efficiency. The largest

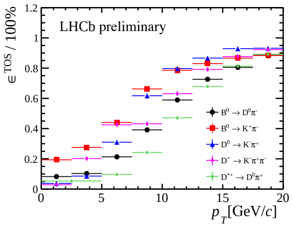

Figure 1. L0 hadron trigger performance: Trigger efficiency for beauty and charm decay modes, listed in the legend.

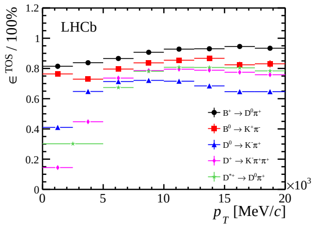

Figure 2. HLT1 inclusive trigger performance: efficiency for various channels as a function of p_{T} of the decaying hadron.

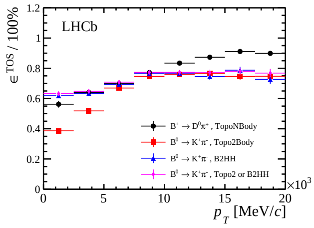

Figure 3. HLT2 inclusive beauty trigger performance as a function of p_{T} of the decaying hadron.

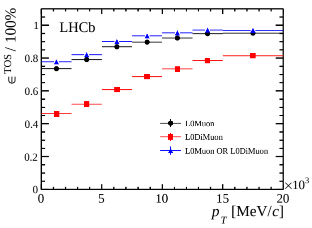

Figure 4. L0 muon trigger performance: Trigger efficiency for selected $B^- \to J/\psi K^-$ candidates.

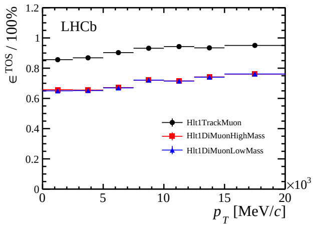

Figure 5. HLT1 muon trigger performance: Efficiency for $B^- \to J/\psi K^-$ candidates as function of B^+ p_{T}.

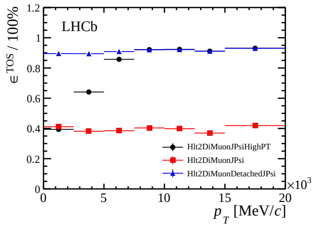

Figure 6. HLT2 muon trigger performance.

inefficiency originates in the tight muon identification requirements inside the L0 reconstruction algorithm. The L0 dimuon trigger selects a small fraction of additional candidates at lower transverse momenta. The combined efficiency for both L0 muon triggers, integrated over p_{T}, is 89%.

In the first part of the HLT, muons are selected both by an inclusive single track line (analogous to the hadron line discussed above) and by dimuon triggers. The dominant inefficiency for these lines originates in the online muon identification algorithms. The performance of HLT1 at selecting muonic signatures is shown in Fig. 5 as a function of p_{T} of the B^+ candidate.

In HLt2, several trigger lines select events with one or two identified muons. The muon identification procedure is identical to the one used in offline analysis. Single muon candidate events are selected with very tight p_{T} and vertex separation requirements. Dimuon candidate events are selected with either mass or vertex separation requirements, or with a combination of both. The efficiency of the HLT2 muon selections is shown in Fig. 6.

3. Run 2: Trigger preparation for 2015 – 2018

In 2015, the collision energy will be increased compared to 2012 (from 8 TeV to 13 TeV), and the bunch spacing will reach the design value of 25 ns. In these conditions the 2012 instantaneous luminosity of $4 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ can be achieved at a lower pile-up (1 interaction/crossing instead of 1.6). The yield will further benefit from increased cross-sections for b- and c-hadron production. The requirements on the High Level Trigger will be partly relaxed by increasing the output rate for storing data to tape (12.5 kHz). A part of this data will be saved in a reduced data format for analysis without offline processing. On the other hand, the maximum read-out rate will be still limited to 1 MHz, and the new conditions will be challenging for the L0 trigger.

The computing resources will be adjusted to the 2015 conditions by doubling the Event Filter Farm CPU power and increasing the disk space for deferral. Changes will also be introduced in the architecture: both HLT levels will be split in two separate physical processes, allowing the second stage to run asynchronously. This allows a more efficient use of disk space and of the deferral mechanism, ultimately leading to an increase in the resources available for the High Level Trigger processing. A real time alignement and calibration can be performed on the output of HLT1 and used to improve the quality of the reconstruction in HLT2.

The track reconstruction sequence will be improved such that it allows to adopt multi-track trigger selections already in the first stage of the High Level Trigger.

4. Run 3: Trigger strategy from 2018 onwards

The LHC will provide LHCb with an instantaneous luminosity of $2 \times 10^{33}\text{cm}^{-2}\text{s}^{-1}$. This implies signal rates of about 300 kHz of beauty hadrons and 800 kHz for charm hadrons that are reconstructible inside the detector [8]. This represents a fundamental change in trigger requirements of LHCb. It is no longer the case that the trigger must reject background from signal, it must now categorise signal according to physics requirements. Such a strategy requires significantly more information in the trigger than can be provided by low-latency, hardware based solutions.

With the removal of the L0 hardware trigger, LHCb will be the first hadron collider experiment to deploy a trigger exclusively in software, using off-the-shelf hardware. The upgraded LHCb trigger will select and categorise events at the full collision rate. The advantages of such a system are clear. Software is easily modified, allowing an unprecedented flexibility as and when the LHCb physics programme changes. It also has the advantage that computing power is readily upgradeable and benefits from the ability to purchase more CPU power for the same price at a later stage. The deferral technique used in Run 1 will be leveraged in Run 2 and Run 3. Subdetector alignment and calibration will be performed while events are buffered to

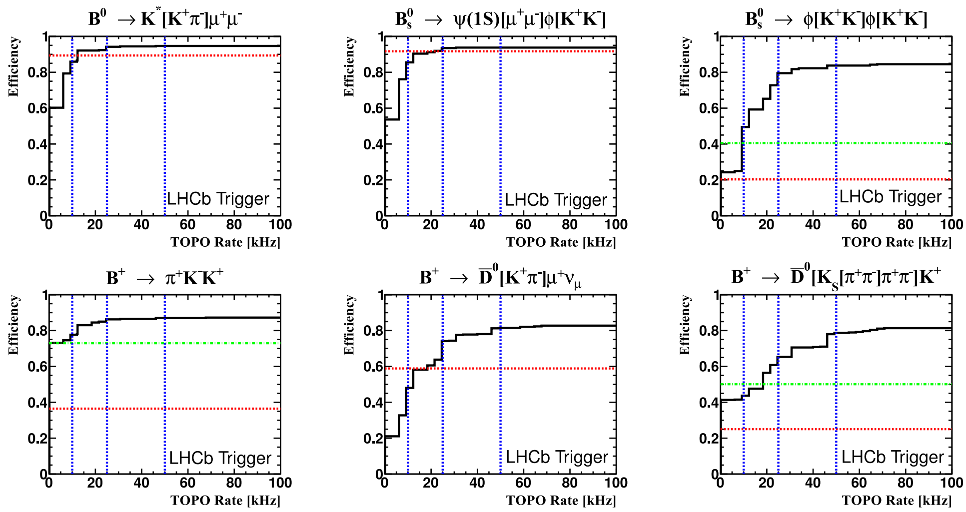

Figure 7. Efficiencies as a function of the output rate for several benchmark decay modes. Red lines indicate the efficiency during Run 1. Green lines indicate twice the Run 1 efficiency. Blue lines indicate three prospective output rates at 10, 25 and 50 kHz respectively.

disk, permitting the use of offline-quality reconstruction, and reducing the need for reprocessing of data at a later stage.

In the Run 3 trigger, all tracks will be made available at the earliest stages of the trigger and the full tracking sequence will be performed upfront [9]. In the trigger tracking sequence the same algorithms as in the offline sequence are used, but the sequencing and configuration are modified to reconstruct the most valuable tracks first and perform some basic candidate selection using these. Slower more specialised reconstruction algorithms only need to be executed later in the HLT algorithm flow.

After the tracking sequence, the remaining timing budget is available to apply offline-quality trigger selections. An inclusive, multivariate beauty signal selection is then performed directly on the 30 MHz of inelastic collision data. The performance of the inclusive trigger selections is summarised in Fig. 7. This inclusive trigger selection is complemented by a set of exclusive trigger selections, that are designed to maximize the efficiencies for specific core modes or to minimize systematic uncertainties. One example are trigger selections that do not bias the decay time of the selected mesons.

5. Summary

The LHCb trigger system has performed extraordinarily well in the first running phase of the LHC. It is designed to select charm and beauty hadrons in a large range of decay modes. The flexible design of the HLT, fully deployed in software, allows to quickly adjust to changes in running conditions and physics goals. Inclusive selections in the full trigger chain allow an efficient trigger for basically any beauty decay to charged tracks. Several innovative concepts have enabled this performance: the deferred triggering allows to optimise the trigger usage for mean instead of peak usage of the available computing resources. Multivariate selections allow the inclusive selection of beauty decays into charged tracks with high efficiency.

For the upcoming Run 2, several improvements of the trigger system are planned: the two software trigger levels will be decoupled which allows calibrations between them and thus a performance much closer to the one achieved offline. The use of event buffering for alignment and calibration permits selections that would otherwise only be available in the offline analysis environment.

The upgraded LHCb trigger represents a turning point in the design of hadron collider trigger systems. The approach to trigger purely in software running on commercially available hardware allows an inexpensive, flexible and scalable design. The upgraded trigger will perform offline-quality tracking at the full collision rate, which allows a number of very efficient inclusive and exclusive trigger selections. Subject to output bandwidth requirements, gains of up to a factor of four in efficiency are possible for several LHCb benchmark modes.

6. References

[1] A. A. Alves Jr. et al., The LHCb Detector at the LHC, JINST 3, S08005 (2008).
[2] R. Aaij *et al.* [LHCb Collaboration], LHCb Detector Performance, Int.J.Mod.Phys. A30 (2015) 1530022
[3] R. Aaij et al., The LHCb trigger and its performance in 2011, JINST 8, P04022 (2013).
[4] J. Albrecht et al., Performance of the LHCb High Level Trigger in 2012, J. Phys.: Conf. Ser. 513, 012001 (2014).
[5] I. Bediaga et al., LHCb Trigger and Online Upgrade Technical Design Report, CERN-LHCC-2014-016 (2014).
[6] S. Tolk, J. Albrecht, F. Dettori and A. Pellegrino, Data driven trigger efficiency determination at LHCb, LHCb-PUB-2014-039
[7] V. V. Gligorov and M. Williams, Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree, JINST 8, P02013 (2013).
[8] C. Fitzpatrick and V. Gligorov, Anatomy of an upgrade event in the upgrade era, and implications for the LHCb trigger, LHCb-PUB-2014-027
[9] T. Head, J. Albrecht, and K. Dungs, Track finding scenarios for the software trigger of the upgraded LHCb experiment, LHCb-PUB-2014-028