Performance and Upgrade Plans of the LHCb Trigger System

V. V. Gligorov[a,*]

[a] *CERN, CH-1211, Switzerland*

Abstract

The trigger of the LHCb experiment consists of two stages : an initial hardware trigger, and a high-level trigger implemented in a farm of parallel-processing CPUs. It reduces the event rate from an input of 15 MHz to an output rate of around 4 kHz. In order to maximize efficiencies and minimize biases, the trigger is designed around inclusive selection algorithms, culminating in a novel boosted decision tree which enables the efficient selection of beauty hadron decays based on a robust partial reconstruction of their decay products. In order to improve performance, the LHCb upgrade aims to significantly increase the rate at which the detector will be read out, and hence shift more of the workload onto the high-level trigger. It is demonstrated that the current high-level trigger architecture will be able to meet this challenge, and the expected efficiencies in several key channels are discussed in context of the LHCb upgrade.

Key words: Trigger systems, Multivariate Selections, Flavour Physics, LHC
PACS: 07.05.Hd, 02.50.Sk, 13.20.He, 13.20.Fc

1. Introduction

The LHCb detector [1] is a single-arm forward spectrometer covering the pseudo-rapidity range $2 < \eta < 5$, designed for studing beauty and charm (heavy flavour) hadrons produced in proton-proton collisions at the Large Hadron Collider (LHC) at CERN. In what follows "transverse" means transverse to the LHC beamline. The detector includes a high-precision tracking system consisting of a silicon-strip vertex detector surrounding the pp interaction region, a large-area silicon-strip detector located upstream of a dipole magnet with a bending power of about 4 Tm, and three stations of silicon-strip detectors and straw drift tubes placed downstream. The combined tracking system has a momentum resolution $\Delta p/p$ that varies from 0.4% at 5 GeV/c to 0.6% at 100 GeV/c, an impact parameter[1] resolution of 20 μm for tracks with high transverse momentum, and a decay time resolution of 50 fs. Charged hadrons are identified using two ring-imaging Cherenkov detectors. Photon, electron and hadron candidates are identified by a calorimeter system consisting of scintillating-pad and pre-shower detectors, an electromagnetic calorimeter, and a hadronic calorimeter. Muons are identified by a muon system composed of alternating layers of iron and multiwire proportional chambers.

Although designed for a luminosity of $2{\cdot}10^{32}$ cm^{-2}s^{-1}, LHCb took data at an average luminosity of $2.8{\cdot}10^{32}$ cm^{-2}s^{-1} throughout 2011 and is currently taking data at $4 \cdot 10^{32}$ cm^{-2}s^{-1}. This improvement is made possible by the excellent performance of both the detector hardware and reconstruction software. Since the LHC currently runs with half its nominal number of bunches (50 ns bunch spacing), this implies an average of 1.7 proton-proton collisions per bunch crossing. The LHCb upgrade is planned to run at a luminosity of 10^{33} cm^{-2}s^{-1} but with a 25 ns LHC bunch spacing, which implies only a modest increase in the average number of proton-proton collisions per bunch crossing to around 2.1. Moreover although the centre-of-mass energy will increase from the current 8 TeV to 14 TeV, this is expected to have only a modest effect on the multiplicity of tracks produced in the proton-proton collisions. These facts mean that the current running conditions can be reliably extrapolated to predict the performance of an upgraded detector.

The characteristics of heavy flavour decays and the production environment of the LHC within which the LHCb trigger operates will be discussed first. The current performance of LHCb's trigger system will be characterized, and it will be demonstrated that the algorithms deployed in the present high-level trigger (HLT) not only perform well in current conditions but that they will scale to the conditions expected to face an upgraded LHCb detector. On the other hand, the LHCb upgrade strategy will be shown to rely on an ability to read out the detector at significantly higher rates than the current 1 MHz, allowing the HLT to do the vast majority of the event selection.

2. The anatomy of heavy flavour decays

The mass of decaying beauty and charm hadrons is transmuted into transverse momentum carried away by the child particles into which they decay, while the O(10^3) fs lifetimes of beauty and charm hadrons ensure that those child particles form vertices detached from the primary proton-proton interaction. A typical beauty or charm hadron produced at the LHC has a decay length of around 1 cm per mean lifetime, much

*On behalf of the LHCb collaboration
Email addresses: `vladimir.gligorov@cern.ch` (V. V. Gligorov)

[1] Impact parameter is the transverse distance of closest approach between a track and a vertex, most commonly the primary proton-proton interaction vertex.

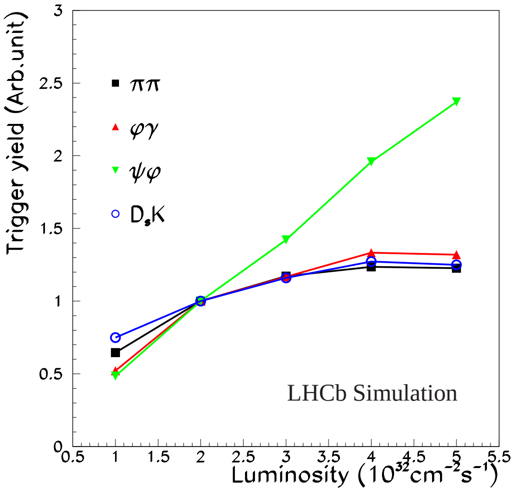

Figure 1: Efficiencies of the L0 at a given luminosity for selected channels measured on simulated events. Reproduced from [2].

greater than the impact parameter resolution of 20 µm. Because the magnitude of the transverse momentum of the child particles is directly linked to the mass of the parent particle, the scalar sum of the transverse momenta of these child particles is a better discriminating variable than the vector sum (which is proportional to the transverse momentum which the parent is produced with). Furthermore the produced particles are generally asymmetric : there is usually a single child which takes a large proportion of the total (transverse) momentum. These are the key trigger signatures of heavy flavour.

3. The LHC environment

As the job of a trigger is to accept interesting events, and reject others, it is important to understand the production cross-sections of various kinds of processes at the LHC. In the context of the LHCb trigger, the most important consideration is the ratio of prompt charm to prompt beauty production, which is measured to be around 20 at 7 TeV. This is critical for two reasons. Firstly, charm hadrons are lighter than and don't live as long as beauty hadrons, but they nevertheless do form decay vertices detached from the primary proton-proton interaction and produce child particles with substantial (> 1 GeV) transverse momentum. They are therefore the most dangerous backgrounds for beauty hadron triggers. Just as importantly, charm hadrons are very interesting to study in their own right, and should be triggered as efficiently as possible within the available output bandwidth. It follows that a good separation of charm and beauty hadron decays is mandatory for the LHCb trigger. In addition, the trigger must maintain high efficiency on dimuon final states, as well as being able to generically trigger on heavy long lived objects in order to allow the discovery of potential new particles produced at the LHC.

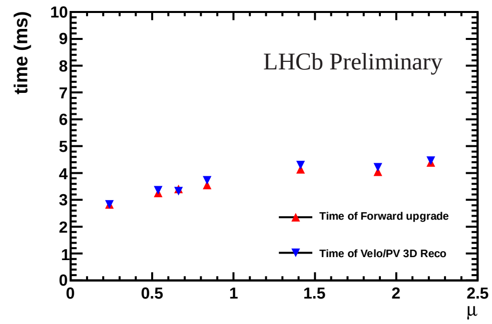

Figure 2: Timing of the VELO reconstruction and forward upgrade as a function of the mean number of interactions per event (μ). Reproduced from [3].

4. Performance of the current LHCb trigger

The LHC proton-proton collision rate of 15 MHz must be reduced by the data acquisition system (trigger) to a few kHz which are saved for further analysis. The LHCb trigger consists of a hardware and a software stage. The hardware stage, also called the Level-0 trigger (L0), reduces the interaction rate to 1 MHz, which is the rate at which all of LHCb's detectors can presently be read out. It has a latency of 4 μs and makes decisions based on energy deposits in the calorimeter system and track stubs in the muon detectors. The HLT is implemented in a software farm of around 30000 logical cores, giving it around 30 ms in which to process each event. It reuses LHCb's offline reconstruction and selection software, adjusted for the online speed requirements through specific options, for example tightening tracking search windows, or fitting vertices using a simplified description of the detector geometry and material. In general the HLT has significantly higher signal efficiency than the L0, since it has far more information available with which to make its decisions.

The L0 makes its decisions based on information from the calorimeters and muon stations. In the case of the muon triggers the momentum resolution is approximately 25% in the region of interest ($P > 6\,\text{GeV}/c$ and $P_{\text{T}} > 1\,\text{GeV}/c$). The resolution for the electromagnetic and hadronic calorimeters is approximately $10\%/\sqrt{(E)}$ and $80\%/\sqrt{(E)}$ respectively, and the requirement is placed on the transverse energy of the calorimeter cluster computed under the assumption that the decays originate from z = 0 in the LHCb coordinate system[2]. The L0 performance varies according to the nature of the event

- For decays involving multiple muons in the final state, whether of beauty or charm hadrons, the L0 is > 90% efficient with respect to offline selected events. This is firstly because muons are not produced directly in proton-proton collisions (their most abundant source being decays of $c\overline{c}$ states), so identifying a muon is already a powerful background discriminant. Decays involving multiple electrons

[2] In which the z axis is along the beamline, and the zero point lies roughly in the middle of the vertex detector.

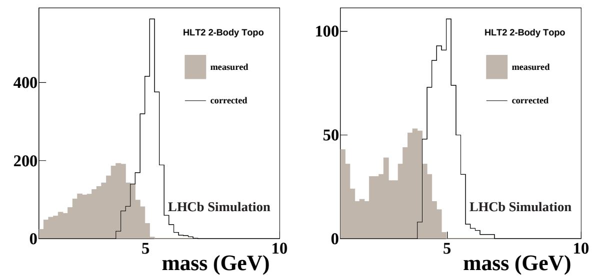

Figure 3: The measured and corrected masses of two-track vertices formed from $B^+ \to h^+h^+h^-$ (left) and $B^0_d \to K^{*0}\mu^+\mu^-$ (right) signal decays. Reproduced from [5].

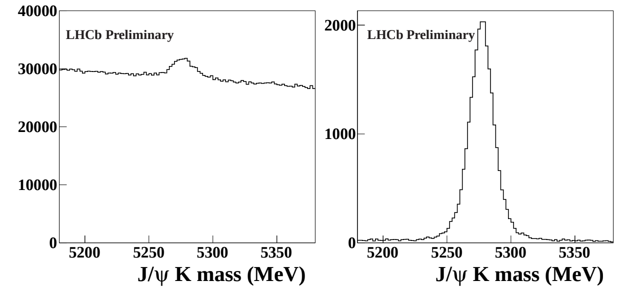

Figure 4: Candidate $B^+ \to J/\psi K^+$ events from 7 TeV collision data collected in 2011 by LHCb. On the left, selected using dimuon triggers only. On the right, the subset of the same events which fire the topological trigger. Reproduced from [6].

behave in a similar way, but have lower efficiencies because of greater backgrounds which impose tighter transverse energy requirements.

- For beauty decays involving a single muon the L0 maintains $> 70\%$ efficiency, a fraction of which comes from triggering on the hadronic decay products.

- For decays involving photons, for example $B^0_s \to \phi\gamma$, the offline analysis requires a high transverse energy photon, $E_{\mathrm{T}} > 2.8\,\mathrm{GeV}/c$, which can then also be imposed within the L0 while maintaining $> 90\%$ efficiency.

- For purely hadronic decays, the L0 has efficiencies of O(10%) for charm decays and O(20%) for beauty meson decays. This is partly because of the worse energy resolution, but largely because of the much higher backgrounds which impose tighter transverse energy requirements of $3.5\,\mathrm{GeV}/c$ or greater.

The performance of the L0 depends strongly on the instantaneous luminosity, as seen in Fig. 1. The y-axis represents the signal yield in each mode in some arbitrary units. This rises linearly with instantaneous luminosity for muons, where the L0 is best able to discriminate between signal and background, but it plateaus rapidly for other kinds of decays, limiting the maximum instantaneous luminosity at which the detector can be usefully operated.

The L0 outputs 1 MHz of events, of which around 100 kHz actually contain a charm or beauty hadron. These events are subsequently processed by the HLT, which begins with the reconstruction of tracks and primary vertices in the vertex detector. Two types of tracks are then identified : those which are detached from the primary interaction, and those which can be matched to track segements in the muon detectors. These requirements select an average of 12 tracks per event for further processing. The selected tracks are subsequently extrapolated to the main tracking system, in search windows constructed by assuming the desired p_{T} of the track, roughly $> 1.5\,\mathrm{GeV}/c$ for detached tracks and $> 0.5\,\mathrm{GeV}/c$ for tracks identified as muons. Successfully extrapolated tracks undergo a further track quality selection, including a kalman filter based track fit. An event is selected if it contains either a high p_{T} detached track, a single detached muon track, a detached dimuon candidate, or a dimuon candidate near or above the J/ψ meson mass. Certain special triggers for W, Z^0, and exotic particles are also executed. This stage of the HLT reduces the trigger rate to around 50 kHz, and takes around half the total time available to the HLT. The timing of this trigger stage does not depend greatly on the average number of interactions per event, as shown in Fig. 2. It is $> 80\%$ efficient on any beauty hadron decay involving multiple charged tracks, and around 50% efficient on certain key charm hadron decay modes such as $D^0 \to K^+K^-$. Only a few percent of events which do not contain a charm or beauty hadron survive this trigger stage, further details of which can be found in [3, 4].

The second HLT stage performs a global pattern recognition similar to the offline one, and is able to reconstruct tracks down to 300 MeV of p_{T} and 3 GeV of p. It carries out three major kinds of selections : firstly, it searches for charmed hadrons produced in the primary interaction decaying into a variety of exclusive modes; secondly, it searches for inclusive di- and multi-muon signals both with and without requiring that the muons are detached from the primary interaction; thirdly, it performs an inclusive search for two-, three-, and four-track decay vertices which are detached from the primary interaction and are consistent with the decay of a beauty hadron [5, 6]. The inclusive selection of beauty-like detached vertices is based around a seven-variable bagged boosted decision tree [7] algorithm known as the "topological trigger", whose input variables are discretized according to their online resolution as measured from the data. This discretization protects against overtraining, since the algorithm cannot select regions of parameter space smaller than the resolution, and makes the algorithm sufficiently fast for use in a trigger by limiting the total number of "keep" regions to a manageable number. A key variable is the "corrected mass" of the vertex, which is defined as

$$m_{\mathrm{corr}} = \sqrt{m^2 + |p_{\mathrm{T}}^{\mathrm{miss}}|^2} + |p_{\mathrm{T}}^{\mathrm{miss}}|, \tag{1}$$

where $p_{\mathrm{T}}^{\mathrm{miss}}$ is the missing momentum transverse to the flight direction of the candidate vertex with mass m [8]. This variable allows the trigger to recover, and subsequently cut on, the mass of the decaying beauty hadron even when using only a subset of its tracks to reconstruct the displaced vertex, as seen in Fig. 3. The topological trigger is $> 75\%$ efficient on generic B

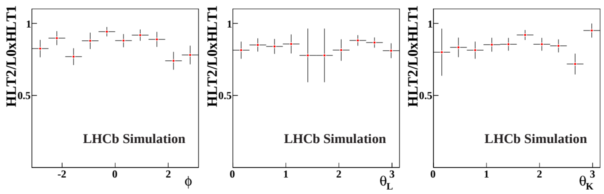

Figure 5: The efficiency of the topological trigger as a function of the decay angles in the $B \to K^{*0}\mu\mu$ decay. Reproduced from [5].

decays into charged tracks, while suppressing almost all events which do not contain a B hadron. This is illustrated in Fig. 4 for $B^+ \to J/\psi K^+$ decays, with the candidates as reconstructed from the dimuon triggers on the left and those candidates passing the topological trigger on the right. Moreover, its inclusive nature means that it does not introduce signifiant biases in physical observables, as shown in Fig. 5 by its flat efficiency as a function of the angular observables in the decay $B_d^0 \to K^{*0}\mu\mu$.

5. The LHCb trigger upgrade

As discussed in Sec. 1 , the current LHCb running conditions closesly resemble the nominal upgrade running conditions in terms of the detector occupancy. Since the trigger timing is anyway relatively insensitive to the average number of interactions, from Fig. 2, the current HLT design will be able to perform its job in the upgrade era so long as the available computing power scales to the higher input rate. As the LHC upgrade is scheduled to begin taking data in 2018, and Moore's law implies [9] a doubling of computing power every two years for the same money, an input rate ten times higher than the current one should be achievable. This output rate of 10 MHz is therefore taken as the nominal running point of the upgraded detector. The hardware trigger formerly known as the L0 is renamed to the low-level trigger (LLT) in the upgrade, to reflect its reduced role in the data selection.

Performance studies for the upgrade have been performed with simulated events generated in the nominal LHCb simulation but assuming 14 TeV centre-of-mass energy and 2.1 average proton-proton interactions per bunch crossing. Table 1 shows the expected rate and signal efficiency of the upgraded detector in some key modes, both for the nominal configuration and for a more pessimistic scenario in which the LLT output rate is 5 MHz. The pessimistic scenario implies LLT cuts which are only slightly looser than those used in present datataking, but even though the signal efficiency is unchanged the upgrade is able to collect almost ten times more signal events per second of LHC running because of the twice higher cross section for producing B hadrons and 4 times higher instantaneous luminosity compared to 2011 datataking. In the more optimistic scenario, hadronic modes gain almost another factor two, while small gains can also be made for other channels. In practice it is of course likely that LHCb will adiabatically increase the size of the HLT computing farm as experience is gained taking data with the upgraded detector.

Table 1: HLT performance for some key channels in two upgrade scenarios.

Running conditions	5 MHz LLT	10 MHz LLT
HLT output rate	16 kHz	26 kHz
Signal Efficiency		
$B_s^0 \to \phi\phi$	29%	50%
$B_s^0 \to \phi\gamma$	43%	53%
$B_d^0 \to K^{*0}\mu\mu$	75%	85%

6. Conclusions

The LHCb experiment is currently taking data at twice its nominal instantaneous luminosity at the LHC, thanks in large part to a flexible and robust software trigger whose performance has been shown to be largely insensitive to the LHC running conditions. The maximum instantanous luminosity at which the detector can be usefully operated is largely limited by the performance of the hardware trigger, which rapidly becomes inefficient for hadronic decay modes at high luminosities owing to the limited range of background discriminants available to it. For this reason the LHCb upgrade aims to increase the detector readout rate from 1 MHz to 40 MHz, allowing more of the work to be performed within the software trigger.

It has been demonstrated that the current software trigger performs well. It is $> 60\%$ efficient on typical B decays into charged tracks selected offline and its inclusive nature makes it inherently robust against operational problems. It has also been demonstrated that the software trigger will scale to the conditions expected to face an upgraded detector, allowing LHCb to collect between ten and twenty times more signal events per second of datataking than collected with the current detector in 2011.

Acknowledgments

VVG is supported by a Marie Curie Action: "Cofunding of the CERN Fellowship Programme (COFUND-CERN)" of the European Community's Seventh Framework Programme under contract number (PCOFUND-GA-2008-229600).

References

[1] LHCb Collaboration (A. Alves et al.), The LHCb Detector at the LHC, JINST, 3:S08005 (2008).
[2] LHCb Collaboration (R. Aaij et al.), Letter of Intent for the LHCb Upgrade, CERN-LHCC-2011-001 (2011).
[3] V. V. Gligorov, A Single Track HLT1 Trigger, LHCb-PUB-2011-003 (2011).
[4] J. Albrecht and R. Aaij, Muon triggers in the High Level Trigger of LHCb, LHCb-PUB-2011-017 (2011).
[5] M. Williams et al., The HLT2 Topological Lines, LHCb-PUB-2011-002 (2011).
[6] V. V. Gligorov, C. Thomas, and M. Williams, The HLT inclusive B triggers, LHCb-PUB-2011-016 (2011).
[7] L. Breiman, Bagging Predictors, Machine Learning 24, 123-140 (1996); 553-568 (1997).
[8] SLD Collaboration (K. Abe et al.), Measurement of R_b using a Vertex Mass Tag, Phys. Rev. Lett., 80:660-665 (1998).
[9] G. E. Moore, Cramming more components onto integrated circuits, Electronics Magazine, 4 (1965).