

The LHCb Trigger

Eric van Herwijnen¹, on behalf of the LHCb collaboration

CERN

1211-CH Geneva-23, Switzerland

E-mail: eric.van.herwijnen@cern.ch

The Large Hadron Collider beauty experiment (LHCb) is a dedicated heavy flavour physics experiment at the LHC. The trigger system employs the finite lifetime and relative large mass of charm and beauty hadrons to distinguish heavy flavour and background from inelastic pp-scattering.

The LHCb trigger is a two level system. The first level is implemented in hardware, it reduces the visible interaction rate to a maximum of 1MHz, at which the whole detector can be readout. The second trigger level is a C++ application running on an Event Filter Farm composed of several thousand CPU nodes. The full trigger is operational in the experiment.

In this talk, an overview of the LHCb trigger system will be given. We put special emphasis on the experience obtained with the initial data taking at the LHC, and the commissioning and monitoring of the software trigger. The method to obtain the efficiency of the trigger from real data will be described, and first results will be presented.

35th International Conference of High Energy Physics (ICHEP2010)

Paris, France

July 22-28, 2011

¹ Speaker

1. Introduction

LHCb is designed to exploit the finite lifetime and large mass of charmed and beauty hadrons to distinguish heavy flavour particles from the background in inelastic pp scattering (see Figure 1). The aim of the trigger [1] is to reject not interesting events as soon as possible.

In this article, we assume that the luminosity equals the LHCb design luminosity of $2 \cdot 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.

The high production rate of b-pairs and the large inelastic pp scattering cross-section at LHC energies force the LHCb trigger to select more or less specific B meson decays of interest, rather than completely generic B signals.




Figure 1 A typical LHCb event

2. Structure of the trigger

Figure 2 shows the structure of the trigger and how it reduces the rate into LHCb. The trigger consists of two parts:




Figure 2 The trigger

1. The “level-0” (L0), which is implemented in hardware. It reduces the visible (e.g. at least 2 tracks inside the detector’s acceptance) interaction rate to a maximum of 1MHz. This presumes the “nominal” rate into LHCb.

2. The “high level trigger” (HLT), which is implemented as a C++ application running on an Event Filter Farm composed of several thousand CPU nodes.

At design luminosity we expect ~ 100 kHz of $b\bar{b}$ pairs inside the bunch crossings with visible pp interactions. However, only about 15% of these events will include at least one B with all its decay products contained in the spectrometer acceptance.

2.1 The L0 trigger

One of the main signatures of B decays is the presence of large p_T particles. The L0 trigger identifies this kind of events by using information from selected subdetectors that are read out at 40 MHz, thus reducing the rate from 10 MHz of bunch crossings with some activity in the detector to 1 MHz. Three types of L0 objects are built and transferred to the L0 Decision Unit to take the trigger decision:

1. L0 hadron: the highest transverse energy ($E_T \geq 3.5$ GeV) deposition found in the hadron calorimeter
2. L0 electromagnetic clusters: the highest transverse energy ($E_T \geq 2.5$ GeV) deposition measured in the electromagnetic calorimeter. L0 also distinguishes between neutral (γ) and charged showers (e)

3. L0 muons: the two highest transverse momenta values ($p_T^{\mu\mu} \geq 1 \text{ GeV}$) associated to muons ($\mu, \mu\mu$)

A positive decision is taken when a L0 object above the indicated threshold is found. The 1 MHz L0 output rate is limited by the readout speed of the various subdetectors.

2.2 The HLT trigger

To cope with a 1 MHz input rate, the HLT algorithms should run on about 2000 multi-core computing nodes. At the time of this conference, the LHCb Event Filter Farm has 20% of its final capacity. 21800 copies of the HLT program will run on the final farm.




Figure 3 The HLT trigger

The HLT has access to information from the whole detector but a complete event reconstruction is not possible at the L0 output rate. Figure 3 shows how the HLT trigger is divided in two sequential steps.

HLT first level (HLT1):

- Confirms the L0 decision using tracking¹
- Performs a reconstruction in the region of interest
- Cuts on simple signatures (p_T , impact parameter with respect to the primary vertex)

Thus the fraction of $c\bar{c}$ and $b\bar{b}$ is enhanced and the rate is low enough that full reconstruction can be performed at the next level, HLT second level (HLT2), which uses the full detector information to produce a mixture of inclusive and exclusive channels.

The selection cuts applied by the trigger are relaxed compared to the offline analysis. To monitor the efficiencies and systematic uncertainties of the L0 and HLT, the detailed trigger decision is included in the raw data.

3. Commissioning

Before the first pp collisions took place in November 2009, the L0 was commissioned using cosmics. It has been running smoothly from day one. The HLT has been prepared and tested offline using Monte Carlo and also with “unbiased” data. To test new online versions of the HLT a system was built (Full Experiment System Test) to inject real and MC data into the DAQ bypassing the L0. FEST has been used to detect and fix bugs, and to do the following benchmarks:

- The HLT configuration time. The executables and libraries are loaded from a distributed point on the network via NFS. As the EFF grows, the time from loading until running needs to be closely monitored.
- The time per event. This constrains the L0 input to the EFF.

¹ Since this conference, the trigger strategy has been changed due to the contamination of hadronic triggers by ghosts in the presence of pileup. Now, a single detached high momentum track is searched for in a region of interest defined by a straight line VELO track segment and the assumed track momentum, without L0 confirmation.

- HLT1 & 2 rejection rates. The 2kHz output rate is fixed by an upper limit of 500 MB/s for the maximum rate with which data can be written to the online storage and by the offline reconstruction capacity on the Grid.

4. Monitoring

To supervise the correct functioning of the trigger a number of counters and histograms displaying various rates are filled in real-time. The L0 rates are displayed in several places that are accessible from the run control. Figure 4 displays the HLT1 Physics accept rate and its composition. Histograms with the total trigger retention rate, the average time per event, and MB/s written to storage are also available.




Figure 4 A typical HLT monitoring histogram

Each HLT1&2 selection produces summary information which is written to storage for the accepted events. This information is used to check if an event would have triggered, even if the B decay of interest would not have participated in the trigger (see section 6).

5. Experience with early data taking

Up to July 2010, running conditions were a factor 100 lower than the design luminosity. Thus the cuts have been relaxed to achieve higher efficiencies, maximizing charm while exploiting the 2 kHz available for output. However, the increased luminosity per bunch




Figure 5 CPU/evt and HLT retention as a function of the number of pp interactions per visible event

due to a lower β^* means we have a higher average number of pp interactions per visible event² (1.5 instead of 1.2, but we even saw 2.3 and are expecting more). This has a dramatic effect on the CPU time/event and retention rate, as shown in Figure 5.

6. Trigger performance

The trigger performance can be calculated from data itself, without relying on Monte Carlo simulations [2]. For a particular channel (e.g. $J/\psi \rightarrow \mu\mu$) one looks at the number of candidates that triggered on the signal ($N_{\text{TOS}}(J/\psi \rightarrow \mu\mu)$), the set of candidates that are triggered independent of the signal ($N_{\text{TIS}}(\text{LO any})$) as well as the set of candidates in which both signal and non signal particles were necessary to trigger ($N_{\text{TOS}\&\text{TIS}}$). It can be shown that, for a given momentum bin, the efficiency to trigger on signal can be found by dividing the intersection between N_{TOS} and N_{TIS} by N_{TIS} .

Figure 6 and Figure 7 show the good performance of the LO muon and the HLT hadron trigger from the early data, determined using this method.

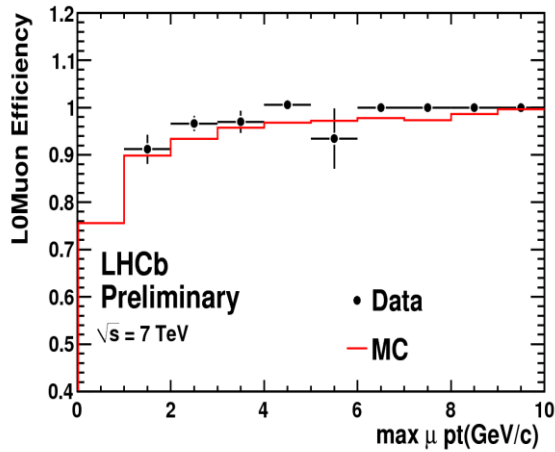


Figure 6 L0 Muon TOS efficiency to select a $J/\psi \rightarrow \mu\mu$ as a function of the p_T coming from the J/ψ 's μ




Figure 7 HLT efficiency for $D^* \rightarrow D(K\pi)\pi$ as a function of p_T . Given the large σ_{charm} a comparison could be made between a trigger which accepts all visible crossings (O) and the Hlt1 physics trigger (\square)

7. Conclusions

The full trigger is operational in the experiment. Efficiencies are as expected. At low luminosity we are running with relaxed thresholds and we are quickly adapting to more challenging than nominal conditions per crossing.

References

- [1] Chapter 7 of *The LHCb Detector at the LHC* (JINST), 2008 JINST 3 S08005.
- [2] E. Lopez Asamar et al., *Measurement of trigger efficiencies and biases*, LHCb-2008-073

² The L0 Calo rate measures visible pp interactions. Using Poissonian statistics, the number of visible pp interactions per bunch crossing is $-\ln(1 - (\text{L0 calo rate}/\text{nb of bunches} * \text{LHC frequency}))$. From this the number of pp interactions per visible event can be deduced.