Eur Phys J C **34**, s01, s419–s423 (2004)
Digital Object Identifier (DOI) 10.1140/epjcd/s2004-04-044-4

EPJ C direct
electronic only

The LHCb Level-1 Trigger

T. Schietinger

Institute of High-Energy Physics, University of Lausanne, CH-1015 Lausanne, Switzerland,
e-mail: Thomas.Schietinger@cern.ch

Received: 1 August 2003 / Accepted: 10 February 2004 /
Published Online: 13 July 2004 – © Springer-Verlag / Società Italiana di Fisica 2004

Abstract. The intermediate trigger stage of the LHCb experiment, operating after the initial hardware assortment but before the selection of specific decay channels, is described. Implemented in software, the trigger uses information on track impact parameters and transverse momenta to identify b hadron events. The main components, the decision algorithm and some performance figures are presented.

PACS. 2 5.70.Ef – 21.60.Gx – 27.30.+t

1 Introduction

1.1 LHCb in numbers

The aim of the LHCb experiment [1,2,3] is to harness the large $b\bar{b}$ cross section at the LHC to map the CKM sector of the Standard Model. Although the b cross section at the LHC energy is very large, estimated at around 0.5 mb, it is still overwhelmed by the total pp cross section at that energy, which is expected to reach 100 mb. The inelastic cross section amounts to about 80 mb, and some 60 mb are expected to give rise to interactions that are visible in the LHCb spectrometer. Factoring in the luminosity foreseen for the LHCb interaction point, 2×10^{32} cm^{-2}s^{-1} or 0.2 mb^{-1} per second, this results in a visible interaction rate of 12 MHz. The event rate seen by the detector is somewhat lower, 10 MHz, due to some interactions occurring in the same bunch crossing (pile-up). Comparison with the b cross section shows that the rate of events containing pairs of b hadrons is smaller by two orders of magnitude (100 kHz). While 200 000 b hadrons per second may sound like a lot, the rates of b decays that are both useful for the determination of CKM parameters and have the good taste of placing all their decay products inside the acceptance of the LHCb spectrometer – in such a way that they can be reconstructed completely – are much smaller: recent studies [4] indicate that on average we may reckon[1] with one decay of the type $B^0 \to J/\psi(\mu^+\mu^-)K^0_S(\pi^+\pi^-)$ per minute, one $B^0 \to \pi^+\pi^-$ every two minutes and about ten $B^0_s \to D_s(K^+K^-\pi^-)K^+$ decays in an hour, to name just a few representative channels. If the Standard Model perseveres, $B^0_s \to \mu^+\mu^-$ decays can be expected at a weekly rate.

[1] The event rates are derived from the annual yields given in [4], before trigger.

1.2 LHCb Trigger overview

It is the task of the trigger system to identify these rare useful events among the vast amount of collisions seen by the detector and to initiate ("trigger") the full chain of event read-out and reconstruction for such events. Several trigger levels are necessary to achieve the required large rejection factor. In LHCb, the trigger is based on three levels:

– **Level-0 Trigger:** In this first line of defense clusters in the calorimeters and track segments in the muon chambers are analyzed to look for evidence of objects with high transverse momentum (hadrons, photons, electrons, and muons), which would signal the presence of b decay products. At the same level information from a part of the silicon vertex detector is used to identify and reject events containing multiple interactions (pile-up veto). This trigger level, described in more detail in [5], is implemented in custom boards. It is designed to reduce the initial event rate of 10 MHz by a factor of ten, while increasing the b purity from a little less than 1% to 3%.
– **Level-1 Trigger:** The first of two software levels, this stage is the focus of this article. Straight-line tracking in the silicon vertex detector and a limited form of spectrometry using only one tracking station are used, in combination with the high-p_t objects found and passed on by Level-0, to bring the event rate further down from 1 MHz to 40 kHz. The b purity attains 9%.
– **High-Level Trigger:** The full event information is available at this level, with the exception of particle identification from the RICH detector. The algorithms are still being developed. It is currently foreseen to seek first a confirmation of the Level-1 decision, uti-

lizing the full tracking capabilities of the experiment. This is expected to cut the event rate in half. Further reduction to the storage rate of 100–200 Hz will be achieved by applying simplified versions of the offline physics analysis algorithms, which search for specific decay channels.

In the following, we outline the strategy employed by the Level-1 trigger, describe how the input quantities are obtained from the sub-detectors, and explain the decision logic. Some details concerning the implementation as well as a few performance figures based on simulation data are given in the second half of the article.

2 Trigger strategy

In Level-1, we take advantage of the fact that the b hadrons are the "elephants" in the particle zoo: they are both heavy and long-lived. The direct measurement of masses and lifetimes, however, would require a rather complete reconstruction of specific b decays, for which there is not enough time at an input rate of 1 MHz. We therefore resort to simple approximations of these properties at the track level: the Level-1 algorithm looks for tracks that have at the same time a large impact parameter with respect to the primary vertex – indicating the decay of a long-lived particle, and high transverse momentum (p_t) – implying the disintegration of a massive object. In addition, the requirement for high p_t ensures that the large measured impact parameter is not feigned by multiple scattering in the vertex detector. It was found that requiring the presence of at least two such tracks is a very effective means to suppress background. Both quantities, impact parameter and transverse momentum, can be extracted with sufficient precision by reading out only a part of the detector.

3 Input quantities

3.1 Impact parameter measurement

The identification of tracks with large impact parameter with respect to the primary vertex relies on the silicon vertex detector, called Vertex Locator (VELO) in LHCb [6]. The VELO is a silicon microstrip detector consisting of 21 stations, each equipped with 2 r- and 2 ϕ-sensors, surrounding the beam pipe in the range -17.5 cm $< z <$ 75 cm around the interaction point, where z is measured in the forward direction along the beam pipe. The sensors are 220 μm thick silicon half-disks with sensitive areas between 0.8 and 4.2 cm radius. The strips of the ϕ-sensors are arranged at stereo angles of $\pm 10°$ and $\mp 20°$, with a pitch that varies between 37 and 98 μm. The r-sensors have strip pitches ranging from 40 to 103 μm and are divided into 45° sectors, thus providing some limited information in ϕ by themselves. VELO data are read out in 170k channels. An average event will send about 1000 clusters to Level-1.

Tracks are first reconstructed in the rz projection only, using the information from the r-sensors. Figure 1 shows an event display of the result of the 2D track search in a 45° sector of the VELO. These 2D tracks form the input of a primary vertex finding algorithm, which achieves a precision of 60 μm in z (rms). The impact parameter of each 2D track can now be determined. In addition, it is checked whether there is a calorimeter cluster or muon track segment found by Level-0 which is consistent with stemming from the same particle as the 2D track. In the next step, only those tracks that are found to have a large impact parameter (between 0.2 and 3.0 mm) or to be consistent with matching a high-p_t object found by Level-0 are subject to the more time-consuming 3D reconstruction using the information from the ϕ-sensors. The upper end of the impact-parameter window (3.0 mm) serves to protect against wrongly reconstructed tracks and tracks from K_S and Λ decays.

3.2 Transverse momentum measurement

An estimate of the transverse momentum of a track necessitates the extrapolation of the track into the detector to a region where the magnetic dipole field of the experiment has had a chance to bend the track. In the standard offline analysis, momenta are measured using the three main tracking stations of LHCb, T1–T3, located just after the magnet. Information from these tracking stations is, however, not available to the Level-1 trigger.[2] Two alternative methods have been developed to assess the transverse momentum of tracks measured in the VELO.

3.2.1 Trigger Tracker (TT)

In the first of these methods, VELO tracks are extrapolated to a smaller full-silicon tracking station, appropriately called the Trigger Tracker (TT), which is installed just in front of the magnet. The Trigger Tracker is a key feature of the re-optimized LHCb design. It consists of four layers of silicon microstrip detectors (500 μm thick, 200 μm strip pitch) with a total area of 7 m^2. The first pair of detector planes is separated from the second one by a gap of 30 cm, which allows not only the determination of the deflection of a particle with respect to the straight-line extrapolation from the VELO, but also the measurement of its direction at TT, improving considerably the momentum resolution. The fringe magnetic field between VELO and TT adds up to $\int B\,\mathrm{d}l \approx 115$ mT m in the re-optimized LHCb design, enough to deflect a 10-GeV track by 3.4 mm at the position of the TT station. The p_t resolution achieved with TT at the trigger level is between 20% and 40%, depending on the particle momentum.

The TT station also plays a key role in the offline reconstruction of events, where it is used to measure tracks originating downstream of the VELO (K_S) or tracks that

[2] Upgrade scenarios for the Level-1 trigger include the use of all tracking stations.

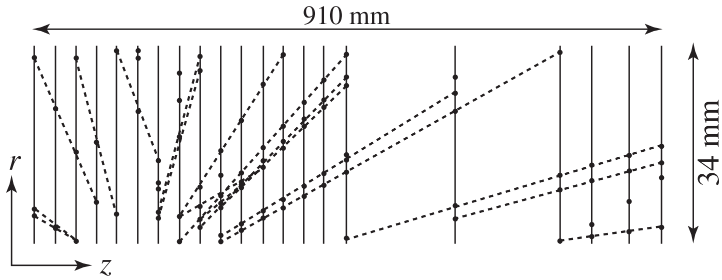

Fig. 1. Event display showing the result of the 2D tracking in one 45° sector of the VELO. Hits in the r-sensors are indicated as dots. The complete event contains 72 forward tracks as reconstructed in 2D

have too low momentum to reach the main tracking stations, such as slow pions from D^* decays or kaons used for flavour tagging.

3.2.2 Matching to Level-0 high-p_t objects

In the second method, the extrapolation of the $\mathrm{d}x/\mathrm{d}z$ and $\mathrm{d}y/\mathrm{d}z$ slopes of a given VELO track is compared with the position of the high-p_t objects found by the Level-0 trigger (calorimeter clusters and muon track segments). If a good match is found, the VELO track and the Level-0 object are likely to stem from the same particle, in which case the trajectory and hence the (transverse) momentum can be determined. Momentum resolutions of 6%, 12% and 15% are achieved for Level-0 muons, electrons and hadrons, respectively. While the matching efficiencies are well above 90% for cases in which both the VELO track and the Level-0 object are present, it is obvious that the overall efficiency for measuring b decay products in this way is rather low, as by far not every one of those gives rise to a calorimeter cluster or muon track that catches the attention of Level-0. The method also suffers from a rather high rate of wrong matches leading to low purities of 52%, 32% and 27% for samples of matched Level-0 muons, electrons and hadrons, respectively. Nevertheless, although the method cannot compete in efficiency with the p_t measurement using TT, its superior momentum resolution and in particular its "poor-man's" particle identification capability open up interesting opportunities. As an example we mention the dimuon invariant mass, an extremely powerful tool for the selection of decays involving $J/\psi \to \mu^+\mu^-$ decays (see Fig. 2). Similar signatures involving electrons and hadrons are currently being studied.

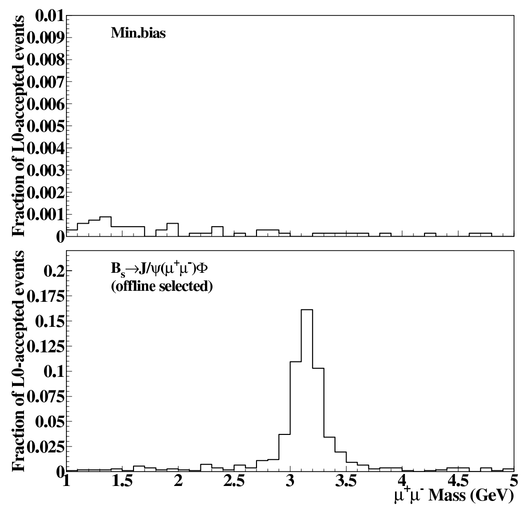

Fig. 2. Invariant mass distribution of dimuons consisting of two VELO tracks that have been matched to Level-0 muon track segments, for minimum-bias events (*top*) and signal events of the decay $B_s^0 \to J/\psi(\mu^+\mu^-)\phi(K^+K^-)$ (*bottom*)

4 Decision algorithm

A decision algorithm is needed to turn the track information into an event probability for containing an interesting b hadron decay. The following algorithm has been used for recent evaluations of the trigger performance:

- Among the tracks with large impact parameter, we select the two with the highest p_t. These are the Level-1 trigger tracks. Here, the p_t of a given track is determined from a match to a Level-0 muon if present, or else, much more likely, from the extrapolation to the TT station. If neither method has been successful, the average p_t of 400 MeV/c is assigned to the track.
- Using the assigned p_t values, the significances of the impact parameters d of the two trigger tracks, i.e. d/σ_d, are estimated from a parameterization.
- A two-dimensional cut is applied in the plane spawned by the two variables $\Sigma \log(p_t)$ and $\Sigma \log(d/\sigma_d)$, where the sums extend just over the two trigger tracks. The cut is illustrated in Fig. 3.
- The cut is *relaxed* in the presence of particularly promising signatures, such as a high dimuon invariant mass, or very-high-p_t electrons and photons seen by the Level-0 trigger.

The algorithm should be regarded as an example of how to make use of the track information available at Level-1. It may still change according to the evolving priorities of the experiment.

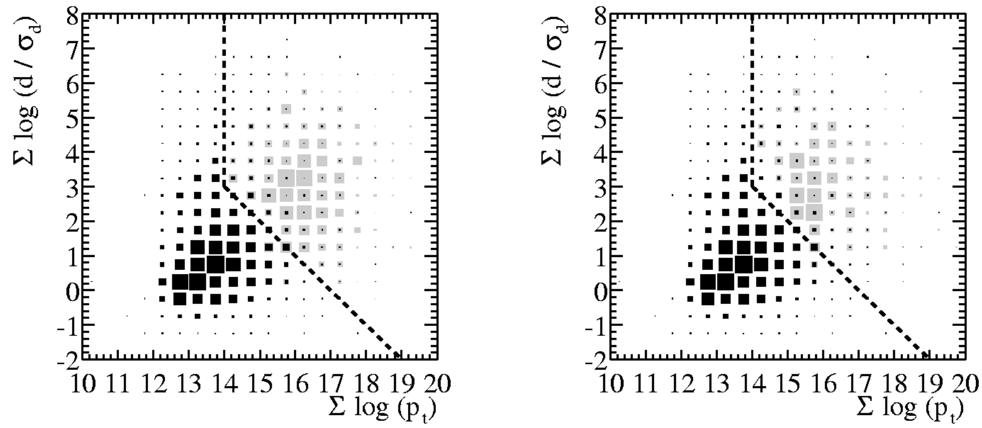

Fig. 3. Distribution of minimum-bias events (black) and signal events passing the offline physics selection (grey) in the plane of the two variables $\Sigma\log(p_t)$ and $\Sigma\log(d/\sigma_d)$ (see text): *left* for $B_d^0 \to \pi^+\pi^-$, *right* for $B_s^0 \to D_s^-(K^+K^-\pi^-)K^+$. The *dashed line* shows the vertical-diagonal cut applied for the Level-1 trigger selection

5 Performance

Table 1 summarizes the trigger efficiencies[3] of a few representative decay channels, as they were obtained for a status report [4] of the experiment in early 2003. The reconstruction and decision algorithms have been improved since, and updated figures will be published in the forthcoming Trigger System TDR [7].

Figure 4 illustrates the relation between the retention of minimum-bias events and the achieved signal efficiencies for three physics channels representing different decay topologies. The curves are obtained by successively tightening the two-dimensional cut described in the previous section.

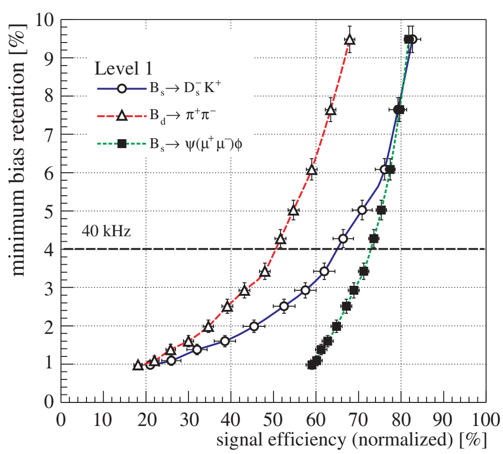

Fig. 4. Retention of minimum-bias events versus signal efficiency for some physics channels. The Level-1 nominal output rate of 40 kHz, corresponding to 4% minimum-bias retention, is indicated by the *dashed horizontal line*

[3] The algorithm used for that evaluation did not yet include the use of the Level-0 photons, hence the rather poor efficiency for $B_d^0 \to K^*\gamma$ (now above 50% for Level-1).

6 Implementation and speed

The Level-1 software will be run on commodity computer processors, endowing us with maximal flexibility in both implementation and physics strategy. While a dedicated farm for Level-1 was foreseen originally, recent changes in detector and read-out design (smaller over-all event size but larger Level-1 event size due to the inclusion of TT data) led us to a scheme in which the Level-1 processors are a part of the LHCb online farm.

The number of processors allocated to the trigger must match the product of the 1-MHz event rate and the average processing time per event. Preliminary studies have shown an average Level-1 processing time of around 5 ms per event, clocked on 1-GHz Pentium-III processors. About 60% of the time is spent on the two-dimensional tracking, some 10% for finding the primary vertex, another 10% for the three-dimensional reconstruction of selected tracks, and the remaining time for the determination of p_t through extrapolations to the TT station and Level-0 objects. Taking into account a factor 6 expected for the increase in computing power per processor between 2002 and the start of the experiment in 2007 [8], we would require some 850 processors to run the Level-1 trigger code in its current form. Many parts of the code have not yet been optimized for speed, so further improvements are to be expected.

Another 200–300 processors will be added to the farm for the High-Level Trigger. The allocation between Level-1 and High-Level will be flexible, permitting the optimal use of resources depending on the experimental conditions.

The latency of Level-1 is given by the buffer length of 58k events, i.e. more than 50 ms.

7 Summary

We have presented the main ingredients of the LHCb Level-1 software trigger, which achieves efficiencies between 50% and 80% for the majority of the decay channels

Table 1. Level-1 and combined Level-0 × Level-1 efficiencies for some physics channels. The efficiencies are normalized to signal events which pass the offline physics selection (not taking into account flavour tagging)

channel	$\varepsilon_{\mathrm{L1}}$	$\varepsilon_{\mathrm{L0\times L1}}$
$B_d^0 \to \pi^+\pi^-$	50.5%	30.9%
$B_s^0 \to D_s^-(K^+K^-\pi^-)K^+$	65.4%	28.6%
$B_d^0 \to J/\psi(\mu^+\mu^-)K_S(\pi^+\pi^-)$	71.1%	64.8%
$B_s^0 \to J/\psi(\mu^+\mu^-)\phi(K^+K^-)$	73.1%	67.9%
$B_d^0 \to K^*(K^+\pi^-)\gamma$	32.7%	26.7%

on the LHCb physics menu. For more details and updated performance numbers we refer to the forthcoming Technical Design Report of the LHCb Trigger System [7], to be published in the fall of 2003.

References

1. LHCb Collaboration: *LHCb Technical Proposal*, CERN/LHCC 98-4
2. LHCb Collaboration: *Reoptimized LHCb Detector Design and Performance Technical Design Report*, CERN/LHCC 2003-030
3. T. Nakada: these proceedings
4. LHCb Collaboration: *Status of the LHCb Detector Reoptimization*, CERN/LHCC 2003-003
5. A. Satta: these proceedings
6. LHCb Collaboration: *LHCb VELO Technical Design Report*, CERN/LHCC 2001-011; see also [2]
7. LHCb Collaboration: *LHCb Trigger System Technical Design Report*, CERN/LHCC 2003-031
8. The LHC Technology Tracking Team for **P**rocessors, Memory, **A**rchitectures, **ST**orage and **TA**pes (PASTA): Final Report (2003), available at `http://lcg.web.cern.ch/LCG/peb/PASTAIII/pasta2002Report.htm`