LHCb-PUB-2011-016
October 19, 2011

The HLT inclusive B triggers

V. V. Gligorov[1], C. Thomas[2], M. Williams[3].

[1] *CERN*
[2] *University of Oxford and STFC Rutherford Appleton Laboratory*
[3] *Imperial College*

Abstract

The inclusive HLT strategy relies on triggering any B decay based on two signatures: a single significantly displaced, high transverse momentum track, and a significantly displaced vertex containing this track and 1-3 other tracks, with high total transverse momentum. In order to provide optimal signal efficiency and background rejection the displaced vertex selection is implemented in a novel boosted decision tree algorithm incorporating information about the experimental resolution in the boosting procedure to protect against overtraining. The performance of these triggers has been commissioned using data taken during 2011 LHCb running and is evaluated here in a data-driven manner. The HLT inclusive triggers are found to have a rejection factor of around 1000 with respect to events selected by the L0 hardware trigger and a $b\bar{b}$ purity close to 100%.

LHCb-PUB-2011-016
20/10/2011

1 Introduction

This note describes the LHCb inclusive triggers for B physics, and their performance in 2011 data taking. Inclusive here means that the triggers are inherently designed to require only a part of the B decay to be reconstructed, which makes them efficient across the full range of B decay topologies which can be reconstructed inside the LHCb detector acceptance.

The LHCb trigger architecture [1] has two levels: Level-0 (L0) is a trigger implemented in hardware while the High Level Trigger (HLT) consists of a software application which runs on every CPU of the Event Filter Farm (EFF). The purpose of L0 is to reduce the rate of crossings with interactions to below the maximum rate dictated by the hardware, which is 1.1 MHz, and to a rate at which the HLT can process all events. L0 reconstructs the highest $E_{\rm T}$ hadron, electron and photon, and the two highest $p_{\rm T}$ muons. L0 is able to distinguish between electron and photon candidates by using the Scintillating Pad Detector (SPD) in front of the Ecal, and reduce their hadron contamination requiring Pre-Shower (PS) energy deposition.

The LHCb HLT trigger is split into two stages for reasons of timing: HLT1 and HLT2. The HLT1 stage performs a partial event reconstruction and selection, and aims to reduce the input rate of $\approx$1 MHz by a factor of around 20. In particular only very high ($>$1 GeV or so) transverse momentum tracks can be reconstructed at the HLT1 stage. This rate reduction then allows the HLT2 trigger stage to perform a more time intensive reconstruction, including all tracks above 500 MeV of transverse momentum, and hence to perform a more efficient final event selection. The HLT1 inclusive trigger is based around the concept of selecting one very good quality track [2], while the HLT2 inclusive trigger is based around a topological [3] selection of a 2-4 track displaced vertex. The design of the HLT1 trigger has not changed since [2], and hence only the performance on 2011 data will be described in this note. The topological trigger has, however, undergone a major change since [3], and is now based on a novel boosted decision tree technique which will be described here.

This note is organized as follows. The method for measuring trigger performance on data is described in Sec. 2. The datasets used for these measurements are described in Sec. 3. The HLT1 design and performance are described in Sec. 4 and 5, and the HLT2 design and performance are described in Sec. 6 and 7.

2 Method for determining efficiencies

The trigger efficiencies are computed using the TISTOS method, following [4] and [5]. In what follows, the term "signal" refers to the track, or combination of tracks, which are of interest in the offline analysis; typically this is the offline reconstructed B or D meson candidate. The term "trigger object" is used to refer to the collection of tracks which caused a particular trigger to return a positive decision. To recapitulate the definitions for the benefit of the reader, the TISTOS method defines three types of trigger decisions:

1. **TIS :** Events which are triggered independently of the presence of the signal. In order for an event to be TIS, there must exist at least one trigger object which does not have any overlap with the signal. The overlap between the signal and trigger objects is tested for by comparing the identifiers (LHCBIDs) of the detector elements which were hit by each track which is part of the signal or trigger object. For the purposes of TIS, two tracks are said to overlap if they share more than 1% of their hits; since a track in LHCb can have around 60 hits at most, this requirement means in practice that the tracks may not share a single hit. TIS events are trigger unbiased except for correlations between the signal B decay and the rest of the event, for example when triggering on the "other" B in the event and subsequently looking at the momentum distribution of the "signal" B.

2. **TOS :** Events which are triggered on the signal decay independently of the presence of the rest of the event. The TOS criterion is satisfied if there exists at least one trigger object all of whose tracks have overlap with the signal. In this case, two tracks are said to overlap if they share more than 70% of their hits (60% for muon segments).

3. **TOB :** Events which are neither TIS nor TOS. These events require both the signal and the rest of the event in order to be triggered, typically triggering because of a signal track combined with a ghost into a displaced vertex. In the case of a single track trigger it is possible to have a TOB event if, for example, the VELO segment of the signal track is combined with a T-station ghost, but such events occur only at the percent level. TOB events are problematic because their efficiency cannot be defined without constructing a model for the trigger efficiency on background events. Although trigger efficiencies are often quoted "globally", irrespective of the TISTOS classification, TOB events are of limited use to any analysis which need to know the trigger efficiency or acceptance.

The efficiency of TOS events is given by

$$\epsilon^{\mathrm{TOS}} = \frac{TOS \text{ and } TIS}{TIS}, \tag{1}$$

while the efficiency of TIS events is similarly given by

$$\epsilon^{\mathrm{TIS}} = \frac{TOS \text{ and } TIS}{TOS}. \tag{2}$$

All efficiencies quoted in this note will be for offline selected HLT TOS events, as these comprise the vast majority of the events available to the offline analyses (as they should, since a well designed trigger should trigger on the signal!).

The precision of the efficiencies computed by the TISTOS method depend on the amount of TIS events available, which are a small fraction of the total offline selected signal. For this reason the four highest yield B decay modes are chosen for this study:

$B^0 \to J/\psi K^{*0}$, $B^+ \to J/\psi K^+$, $B^0 \to D^+\pi^-$, and $B^+ \to D^0\pi^-$. Because of the correlation between the "signal" B and the other B in the event, mentioned above, efficiencies computed using the TISTOS method have to be shown as a function of the variables of interest. Three primary variables of interest are chosen: the B lifetime, the B momentum, and the B momentum transverse to the beamline. These are also the main discriminating variables used in the trigger selections. The following criterion is used for TIS events: any trigger at the L0 stage, HLT1 global TIS, and HLT2 global TIS. The L0 trigger is not TISTOSed since what is being measured here is simply the HLT efficiency.

Since all trigger efficiencies are computed relative to offline selected events, the offline selections used are listed in Tab. 1. In particular note that for the DiMuon channels the offline selection does not cut on any impact parameters or flight distances, while the hadronic selections cut reasonably hard on these quantities, which will make the trigger lifetime acceptance for these modes rather different. It is clear that one can always design an offline selection which will have a trigger efficiency of close to 100%, but the selections chosen here are representative of selections used in real LHCb analyses.

3 Datasets

The dataset used corresponds to 250pb^{-1} of data taken in 2011 with a single HLT configuration. The signal yields in the four modes of interest are shown in Fig. 1 for TIS events. All trigger efficiencies are computed in a 3 sigma mass window around the signal mean. They are background subtracted, and the background is taken from the upper mass sideband which is assumed to represent the background under the signal peak. This is of course an approximation since all the signal modes have some (small) quantity of peaking background, but this is negligible compared to the size of the dataset.

4 HLT1 Design

The HLT1 trigger stage is based around the Hlt1Track trigger, described in [2]. It selects B decays by looking for a single high transverse momentum track with a good track fit quality which is well displaced from all primary interactions. In order to increase efficiency for muon decays, if this track is identified as a muon (according to the same "IsMuon" criterion used offline [6]), the transverse momentum and track quality requirements are relaxed. Finally in events with an L0 Electron or Photon trigger the transverse momentum requirement on the track is slightly relaxed in order to increase efficiency for decays of the type Xγ – since the L0 requirement of a high transverse momentum photon makes it less likely that a high transverse momentum track can also be found in the same decay. These three triggers are called Hlt1TrackAllL0, Hlt1TrackMuon, and Hlt1TrackPhoton respectively. Only the performance of Hlt1TrackAllL0 and Hlt1TrackMuon will be evaluated here, but the same methods can eventually be applied to Hlt1TrackPhoton (radiative decays are rarer, however, so a TISTOS evaluation of their efficiency will need a larger datasample).

Table 1: Offline selection cuts applied for the modes $B^0 \to J/\psi K^{*0}$, $B^+ \to J/\psi K^+$, $B^0 \to D^+\pi^-$, and $B^+ \to D^0\pi^-$. The term "child" refers to all charged pions, kaons, and muons in the final state. The term "bachelor" refers to the child coming directly from the B decay (there is none in $B^0 \to J/\psi K^{*0}$). In addition, loose particle identification requirements are placed on all tracks, which do not significantly alter the momentum spectrum and are omitted for brevity. Mass windows are applied at approximately 3σ around the D,J/ψ,K^{*0} masses to aid purity.

Cut	$B^0 \to J/\psi K^{*0}$	$B^+ \to J/\psi K^+$	$B^0 \to D^+\pi^-$	$B^+ \to D^0\pi^-$
Child track χ^2	< 4			
Child P	> 2000 MeV			
Child P_T	> 300 MeV		> 250 MeV	
Child IP χ^2	-		> 9	
Bachelor (or K*) P_T	> 1000 MeV		> 500 MeV	
Bachelor (or K*) P	-		> 5000 MeV	
Bachelor (or K*) IP χ^2	-		> 16	
K* vertex χ^2	< 16	-		
J/ψ child P_T	> 500 MeV		-	
J/ψ (D) vertex χ^2	< 16		< 9	
J/ψ (D) DOCA χ^2	< 30		-	
J/ψ (D) P_T	-		> 2000	
J/ψ (D) flight distance χ^2	-		> 100	
B P_T	> 2000 MeV		> 1500 MeV	
B τ	> 0.3 ps		> 0.2 ps	
B IP χ^2	< 25		< 16	
B vertex χ^2	< 10		< 9	
B DIRA	-		> 0.9999	
B flight distance χ^2	-		> 16	

The HLT1 track reconstruction proceeds in stages : first a reconstruction of tracks in the VELO and the reconstruction of primary interaction vertices using these VELO tracks. Each primary vertex is required to have more than five VELO tracks. Then a selection of displaced VELO tracks (or VELO tracks matched to muon segments) is made, and for the selected displaced VELO tracks their track segment in the T-stations is sought to determine their momentum, a process called "forward tracking". The differences with respect to the offline reconstruction are

- In the trigger forward tracking is performed independently per VELO track, whereas offline a simulataneous forward reconstruction of all VELO tracks is performed in which the VELO tracks compete for hits with each other.
- Because only very high (transverse) momentum tracks are searched for, the trigger

introduces a momentum cut already at the reconstruction stage (by narrowing the search windows in the forward upgrade) to save time.

There have been no major changes to the trigger architecture since it was introduced. The list of cuts applied on the track of interest is listed in Tab. 2.

5 HLT1 Performance

The TOS efficiency of the 1Track triggers is shown in Fig. 2 (Hlt1TrackAllL0) and Fig. 3 (Hlt1TrackMuon). For the momentum and transverse momentum efficiencies a characteristic "turn-on" curve is observed in all modes: the trigger has reduced efficiency for low (transverse) momenta, this efficiency then rises and eventually plateaus. For the lifetime efficiency, however, the different offline selections lead to substantially different trigger efficiency curves. For the hadronic modes, where all tracks are required to be well displaced

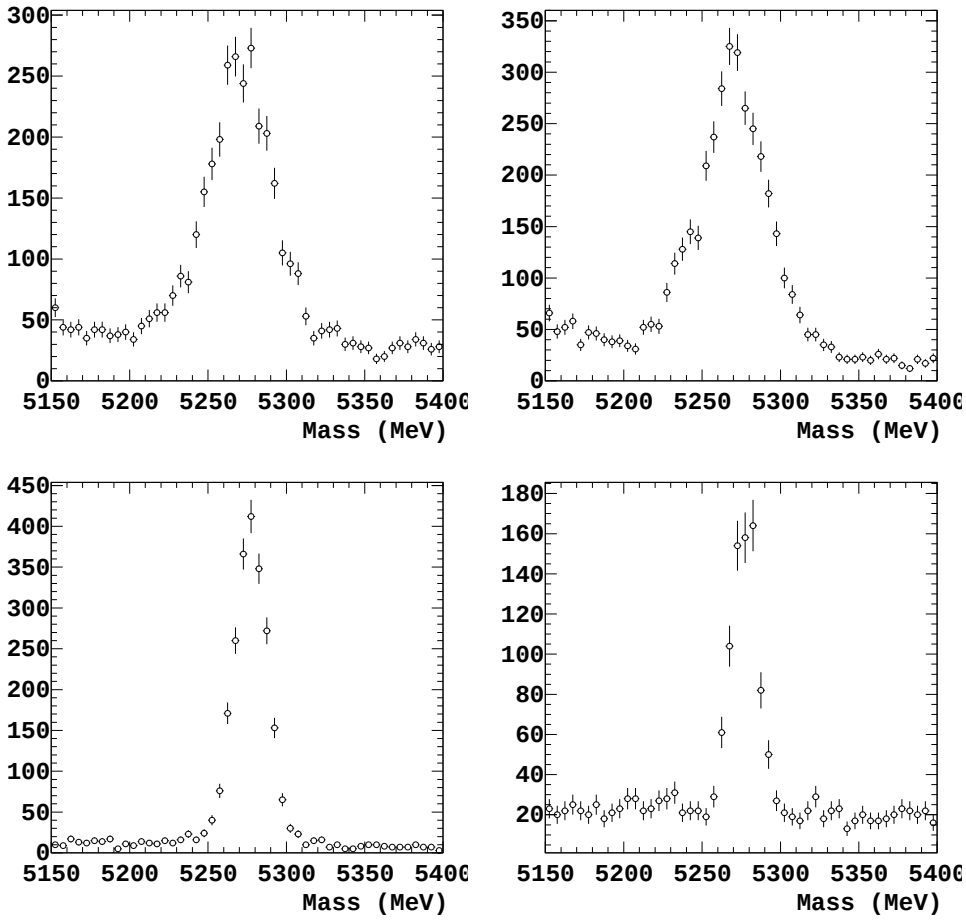

Figure 1: Mass distributions of the $B^0 \to D^+\pi^-$, $B^+ \to D^0\pi^-$, $B^0 \to J/\psi K^{*0}$, and $B^+ \to J/\psi K^+$ TIS candidates (clockwise from top left).

Table 2: Cuts applied in the 1Track trigger lines and HLT1 reconstruction. T-station hits means the total number of outer tracker (OT) and inner tracker (IT) hits, where each IT hit is given a weight of 2 to account for the smaller number of tracking layers in the IT region.

	Hlt1TrackAllL0	Hlt1TrackPhoton	Hlt1TrackMuon
Min. IP (μm)	100		
Min. num. of VELO hits	10		7
Max. num. of missed VELO hits	2		-
Min. P (MeV)	10000	6000	8000
Min. P_T (MeV)	1700	1200	1000
Min. number of T-station hits	17	16	-
Min. IP χ^2	16		
Max. track fit χ^2/nodf	2.5		

from all primary interactions, the 1Track trigger efficiency is flat as a function of the B propertime. This indicates that the lifetime biasing (in this case impact parameter) cuts in the 1Track trigger are fully efficient on the offline selected B mesons, and the trigger inefficiency in each B propertime bin is dominated by the (transverse) momentum cuts and residual tracking efficiency differences between the trigger and offline selections. For the DiMuon modes, however, the offline selection does not cut on the impact parameters of final state tracks, and hence for low B lifetimes these trigger cuts introduce an additional inefficiency.

For muon modes, allowing an "OR" of the Hlt1TrackAllL0 and Hlt1TrackMuon triggers gives an increased efficiency; this is shown in Fig. 4. The Hlt1TrackAllL0 allows triggering on any of the tracks, while the Hlt1TrackMuon allows triggering on one of the muons with reduced (transverse) momentum requirements, and relaxed track quality cuts. In particular, it is important to notice the high efficiency in the plateau region above $\tau(B) = 2$ ps (right most plot in each triplet), which is made possible by the near identical online and offline reconstructions, so that a signal track reconstructed offline is almost invariably also reconstructed in the trigger.

Two other HLT1 lines which are used later in this note are the SingleElectronNoIP line and the Hlt1DiMuon lines. The single electron line selects high transverse momentum tracks which are matched to ECAL clusters used in the L0 trigger decision, and is used as an input to the topological electron triggers described in 6. The DiMuon lines are described in detail in [7]: they trigger on either displaced dimuon vertices, in which case the invariant mass of the dimuon pair is required to be greater than 1 GeV, or on high mass (> 2.9 GeV) dimuon vertices, in which case no displacement from the primary vertex is required.

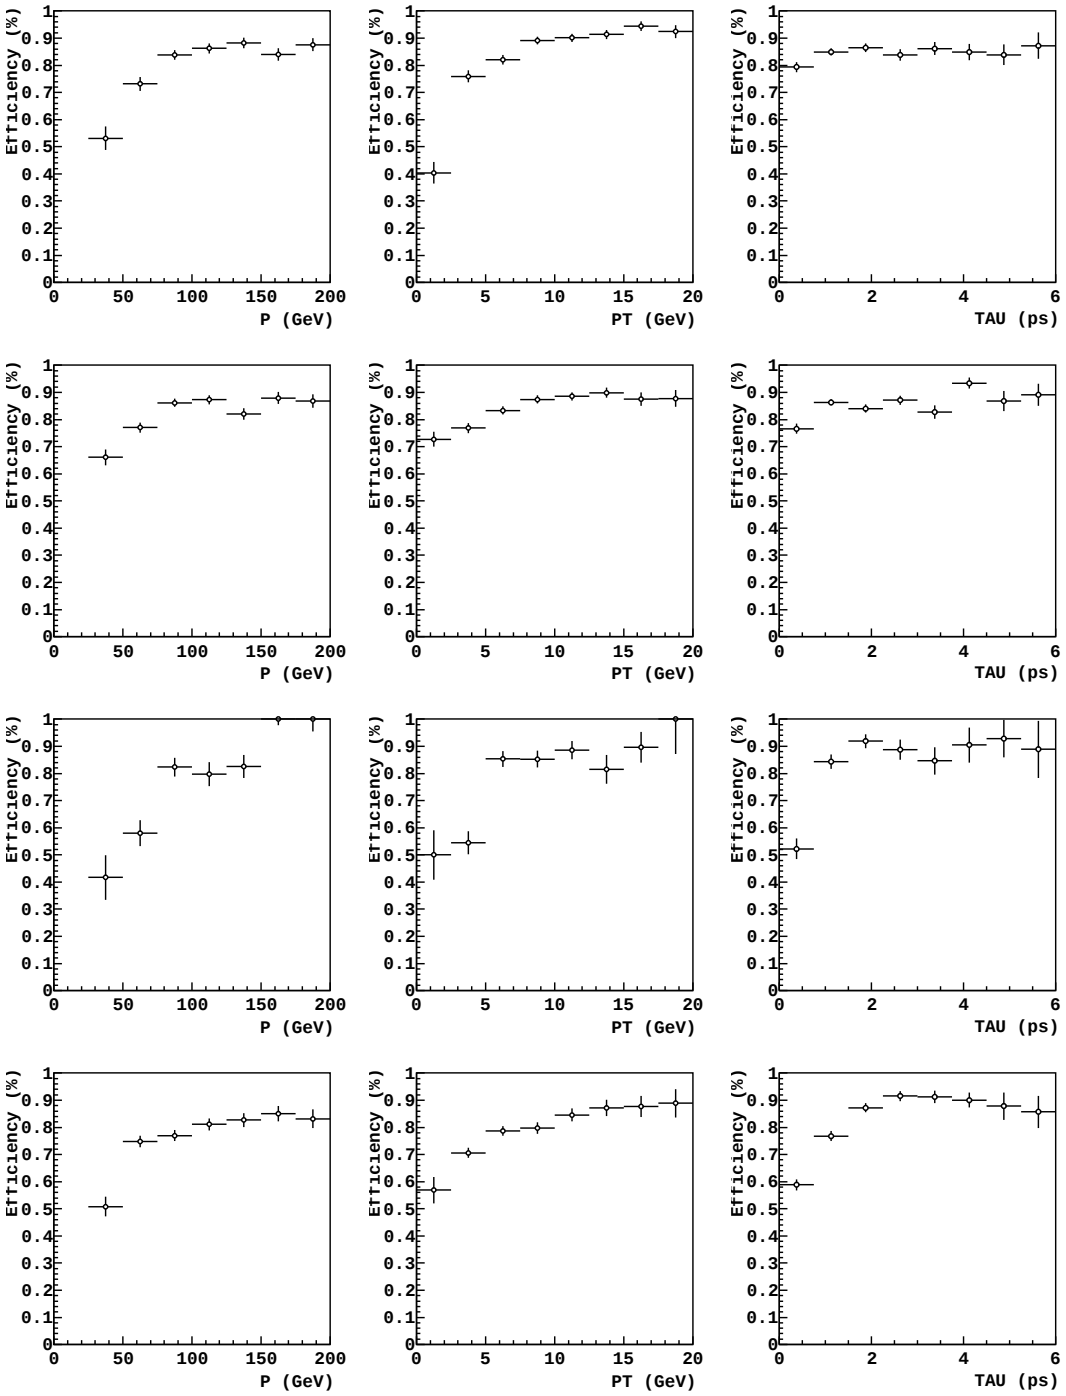

Figure 2: Efficiency of the Hlt1TrackAllL0 trigger for $B^0 \to D^+\pi^-$, $B^+ \to D^0\pi^-$, $B^0 \to J/\psi K^{*0}$, and $B^+ \to J/\psi K^+$ decays (top to bottom) as a function of B momentum, transverse momentum, and lifetime (left to right).

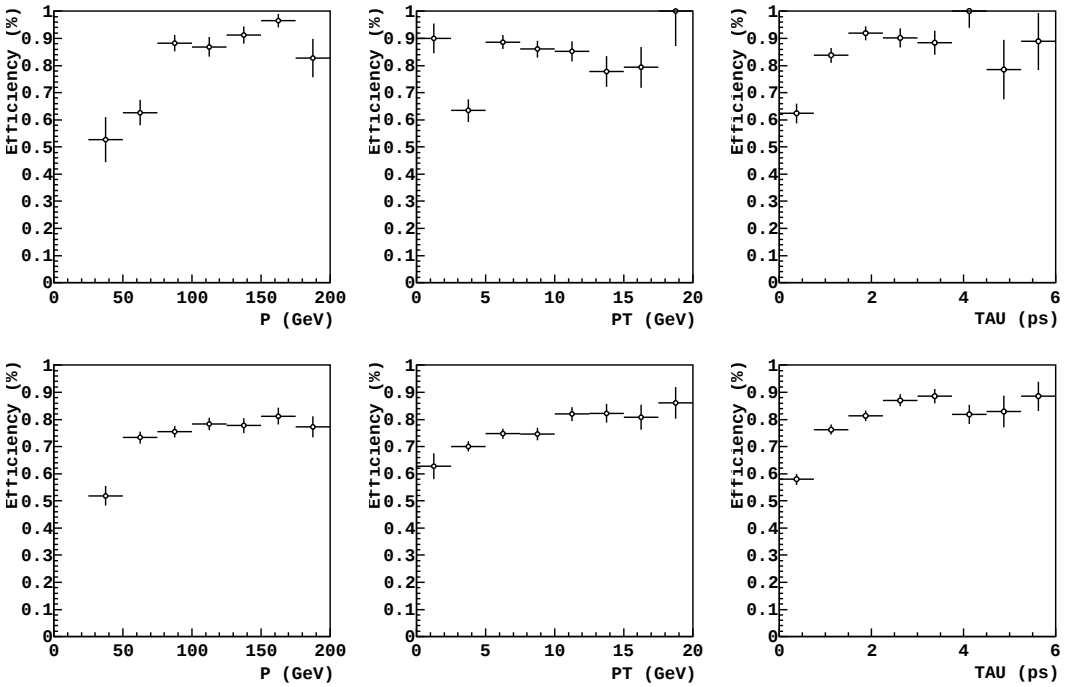

Figure 3: Efficiency of the Hlt1TrackMuon trigger for $B^0 \to J/\psi K^{*0}$ and $B^+ \to J/\psi K^+$ decays (top to bottom) as a function of B momentum, transverse momentum, and lifetime (left to right).

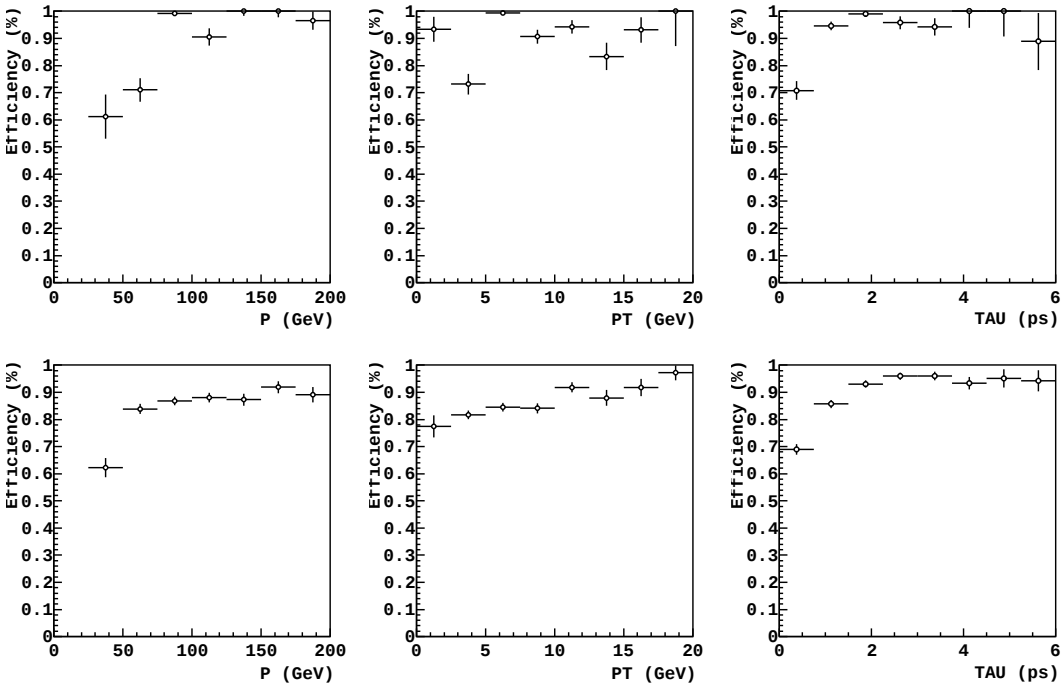

Figure 4: Efficiency of the logical OR of the Hlt1TrackMuon and Hlt1TrackAllL0 triggers for $B^0 \to J/\psi K^{*0}$ and $B^+ \to J/\psi K^+$ decays (top to bottom) as a function of B momentum, transverse momentum, and lifetime (left to right).

6 HLT2 Design

6.1 Methodology

The HLT2 topological lines are designed to trigger efficiently on any B decay with at least two charged daughters. HLT2 starts with the forward tracking of all VELO tracks to determine their momenta. Requiring $P_T > 500$ MeV/c and $P > 5$ GeV/c in the algorithm narrows the search windows and, thus, saves CPU time. Because of this, the topological trigger must take an inclusive approach to obtain the desired efficiency. *I.e.*, the topological trigger is designed to trigger on partially reconstructed b-hadron decays. For this reason tight cuts cannot be applied to exclusive quantities like the mass of the candidate or its impact parameter to a PV. Decays with long-lived resonances (*e.g.*, D mesons) must also be accommodated for. This section will describe the basic HLT2 topological strategy, while the next one will describe the multivariate selection criteria introduced in 2011.

All tracks are required to have IP$\chi^2 > 4$. Tracks identified as muons (IsMuon is true) must have a track $\chi^2 < 4$, while tracks identified as electrons (PIDe > -2) must have a track $\chi^2 < 5$. All other tracks are required to have a track $\chi^2 < 3$. Tracks passing these cuts will be referred to as *input particles*.

Two-body proto-candidates are made from input particles that have the same *best* PV (a particle's best PV is the one to which it has the smallest IP) and a distance of closest approach DOCA < 0.2 mm. A candidate must have an invariant mass less than 7 GeV. Its vertex is required to have a flight-distance $\chi^2 > 100$ and to be downstream of its best PV. No cut is placed on the vertex χ^2 at this stage to allow for the possibility of having one track originating directly from the B decay and the other from a subsequent D decay. These 2-body proto-candidates will be used as input for the 2, 3 and 4 body topological lines.

The 2 body topological lines are constructed by filtering the 2-body proto-candidates by requiring $\sum |p_T| > 3$ GeV and that the candidate is TOS with respect to at least one HLT1 track trigger (this reduces the output rate by 1/3). These 2-body filtered candidates are then passed on to the multivariate selection (described in the next section) to create the 2 body topological lines. The 2-body proto-candidates are filtered in a different way to be used as input for the 3 and 4 body topological lines: their invariant mass is required to be less than 6 GeV and their vertex χ^2 is required to be less than 10. These *2-for-n* proto-candidates are then combined with input particles to make the 3 and 4 body filtered candidates.

The 3-body proto-candidates are made by combining a 2-for-n proto-candidate with an input particle that shares the same best PV. The same invariant mass, DOCA (where DOCA is taken between the 2-for-n proto-candidate and the input particle) and flight distance cuts that are applied to produce the 2-body proto-candidates are applied to make these candidates as well. The 3-body filtered candidates are made by requiring $\sum |p_T| > 4$ GeV and HLT1 track TOS and then passed on to the multivariate selection to make the 3-body topological lines. The 3-for-4 proto-candidates are made by requiring

the 3-body proto-candidates to have an invariant mass less than 6 GeV. These proto-candidates are then combined with input particles that share the same best PV and subjected to the same cuts as the 2 and 3 body proto-candidates to create 4-body proto-candidates. The 4-body filtered candidates are made by requiring $\sum |p_T| > 4$ GeV and HLT1 track TOS and then passed on to the multivariate selection to make the 4-body topological lines.

6.2 Bonsai Boosted Decision Tree

The n-body topological lines are constructed by applying a multivariate selection to the n-body filtered candidates described in the previous section. The bulk of the rejection power of the topological trigger lines is achieved at this stage. In fact, the multivariate selection described below only accepts a few percent of the filtered candidates that pass the simple cut-based criteria described in the previous section. A multivariate classifier known as a boosted decision tree (BDT) was chosen.

A decision tree is a multivariate classifier that is built by performing repetitive one-dimensional splits of the data. The criteria for where to split is based on some figure of merit (in this case signal significance). To limit the effects of *overtraining* (*i.e.*, inaccurate fine tuning due to limited sample sizes) one can *boost* the DT. The type of boosting used in the topological trigger is known as *bagging*. This technique involves creating a large number, in this case 1000, bootstrap copy training samples produced by sampling with replacement from the original one. A separate DT is trained on each sample and the response for any event passed to the BDT is then simply the fraction of these DT's in which the event is in a signal leaf (as opposed to a background one). This procedure greatly enhances the classifying power of the DT.

All multivariate classifiers select n-dimensional regions of a multivariate space to keep by learning from the training samples provided to them. One difficulty that needs to be overcome is that the selected regions could be small relative to the resolution or stability of the detector. This could cause the signal to oscillate in and out of the keep regions resulting in, at best, a less efficient trigger or, at worst, a trigger whose systematics are very difficult to understand. Furthermore, the topological trigger is supposed to be inclusive; however, it is simply not possible to use every known B decay in the training. It is vitally important to ensure that the multivariate classifier is learning common B-decay traits and not a large sum of specific ones. Another concern is that a bagged DT is an extremely large set of if/else statements and it can take a long time to evaluate the response for each event.

The simplest way to avoid these issues is to discretize all of the variables. This then limits where the splits of the data can be made and, in effect, allows the grower of the tree to control and shape its growth; thus, we are calling it a *bonsai boosted decision tree* (BBDT).

This technique works because it enforces that the smallest interval that can be used satisfies $\Delta x_{\text{min}} > \delta_x$ for all x values and on all leaves, where $\delta_x = \text{MIN}\{|x_i - x_j| : x_i, x_j \in x_{\text{discrete}}\}$. The constraints governing the choice of $\{x_{\text{discrete}}\}$ are then as follows:

Table 3: Signal samples used to train the BBDT.

parent	daughters
$B^{\pm}$	$K\pi\pi, D_{[K\pi]}\pi, D_{[hhhh]}K, D_{[K_S\pi\pi]}K, D_{[K\pi\pi]}K\pi$
B^0	$K^*_{[K\pi]}\mu\mu, K^*_{[K\pi]}ee, D_{[K\pi\pi]}\pi, K\pi, D_{[K\pi]}K\pi, D^*_{[D(K\pi)\pi]}\mu\nu, D_{[K\pi\pi]}K\pi\pi$
B_s	$D_{s[KK\pi]}\pi, D_{s[KK\pi]}K\pi\pi, K^*_{[K\pi]}K^*_{[K\pi]}$
Λ_b	$\Lambda_{c[pK\pi]}\pi, \Lambda_{c[pK\pi]}K\pi\pi$

(1) δ_x should be greater than the resolution on x in LHCb and be large with respect to the expected online variations in x and (2) the discretization should reflect common B-decay properties (this is discussed more below). The discretization also allows us to convert the extremely large number of if/else statements into a one-dimensional array of response values; one-dimensional array look-up speeds are extremely fast, much faster than vertexing particles. Thus, by construction the BBDT is fit for purpose for use in the HLT2 topological trigger.

A large number of variables were tested in the BBDT but in the end it was found that the following seven variables were all that is needed: $\sum |p_T|$, $p_T^{\min}$, mass, *corrected mass*, DOCA, candidate IPχ^2 and flight distance χ^2. The corrected mass is defined as follows:

$$m_{\rm cor} = \sqrt{m^2 + |p_T^{\rm miss}|^2} + |p_T^{\rm miss}|, \tag{3}$$

where m is the mass and $p_T^{\rm miss}$ is the missing momentum transverse to the direction of flight of the candidate assuming it originates from its best PV. It is the minimum correction to the mass if any daughters are missing. The optimal discretization scheme for each variable was determined by first training a BBDT with a very large number of discretization values and then gradually decreasing this number as low as possible without losing much in performance. Table 3 gives a list of all of the decays used in the training. The signal samples were each offline reconstructible Monte Carlo, while the background sample was taken from minimum-bias 2010 data.

Table 4 shows the discretization scheme for each of the variables used in the BBDT. Most variables' schemes are completely determined by physics. *E.g.*, IPχ^2 only has one allowed split point: IP$\chi^2 = 20$; thus, the BBDT can only split a node into *points* and *doesn't point* (or it can ignore the pointing). This ensures that differences in the tails of the IPχ^2 distributions of various $B \to X$ decays will not affect the performance of the topological trigger.

6.3 μ and e Lines

The 2011 topological trigger also has muon and electron lines. The μ lines require that at least one of the particles in the candidate has IsMuon equals true. This substantially reduces the number of candidates and allows us to cut looser on the BBDT response. The electron topological lines require that at least one particle in the candidate has PID$e > -2$

Table 4: Allowed split points in the bonsai boosted decision tree.

variable	cuts(2,3,4-body)	allowed splits
$\sum \lvert p_T \rvert$	$>$ 3,4,4 GeV	3.5,4,4.5,5,6,7,8,9,10,15,20 (GeV)
mass	$<$ 7 GeV	2.5,4.75 (GeV)
DOCA	$<$ 0.2 mm	0.05,0.1,0.15 (mm)
IPχ^2		20
corrected mass		2,3,4,5,6,7,8,9,10,15 (GeV)
$p_T^{\min}$	$>$ 0.5 GeV	0.6.0.7.0.8,0.9,1,1.25,1.5,1.75,2,2.5,3,4,5,10 (GeV)
FDχ^2	$>$ 100	2,3,4,5,6,7,8,9,10,25,50,100 $\times$100

and that the L0 electron and that the HLT1 track or Hlt1SingleElectron triggers passed the event (a global event filter with no TOS requirement). This also permits a looser cut on the BBDT response to be made. By using the same BBDT for the standard, muon and electron lines the overlap of the backgrounds is maximized which reduces the total output rate of the topological trigger. The muon and electron lines enhance the efficiency of the HLT2 topological trigger on $B \to \mu X$ and $B \to eX$ decays; the performance of muon triggers is described in Section 7, while the real data performance of electron triggers will be described in a future publication as the currently existing signals are not large enough to allow the efficiency to be extracted from data alone. In the case of $B \to K^* e^+ e^-$ simulated events, the electron topological gains 7% in efficiency compared to having only the regular topological lines in place.

6.4 K_S's as Input Particles

Another new feature in the 2011 topological trigger is the use of K_S's as input particles. These are constructed by taking pions with a track $\chi^2 < 3$ and an IP$\chi^2 > 16$ and vertexing them. If the resulting invariant mass is within 30 MeV of the K_S mass and the vertex $\chi^2 < 10$, then the candidate is accepted as a proto-K_S. The proto-K_S candidates are then filtered by requiring $p_T^{K_S} > 500$ MeV, $p^{K_S} > 5$ GeV, IP$^{K_S}\chi^2 > 4$, FD$^{K_S}\chi^2 > 1000$ and that the K_S vertex is downstream of its best PV. The K_S's are then treated just like any other input particle when forming the n-body candidates discussed above. Adding these particles to the list of inputs increases the efficiency on $B \to K_S X$ (any channel with a K_S in the decay chain) by around 5%. The output rate of the topological trigger lines is virtually unaffected.

7 HLT2 Performance

The timing of the HLT2 topological trigger lines is excellent; they take up a few percent of the total HLT2 timing. The BBDT response cut values that maximize the $pp \to b\bar{b}X$ efficiency while delivering a rate of just over 1 kHz for all of the topological lines combined

Table 5: Cuts applied on the BBDT response for various HLT2 topological trigger lines.

Type	2-Body Cut	3-body Cut	4-body Cut
standard	0.4	0.4	0.3
μ	0.1	0.1	0.1
e	0.1	0.1	0.1

are given in Tab. 5. The output rate is stable under any running conditions feasible prior to an LHCb upgrade; this trigger is very robust.

As with the HLT1 efficiencies, the topological efficiencies can be plotted as a function of some variables of interest. These plots are shown for the standard topological lines in Figs. 5, and for the OR of the standard and muon topological lines in Fig. 6. As in HLT1, efficiencies for modes involving muons are improved by the addition of the muon lines, and the proper time acceptances show clearly the difference between the muon offline selections (which are not lifetime biasing) and the hadron offline selections (which are).

The purity of the output of the HLT2 topological lines can be studied in data. For example, a dedicated J/ψ trigger[1] produces an extremely clean J/ψ sample. We can add a track with $P_T > 500$ MeV, PIDK > 2 and track $\chi^2 < 4$ and require the vertex $\chi^2 < 10$ and produce the $B \to J/\psi K$ candidate sample show in Fig. 7(left). The background level is very high; however, if we simply then require the candidate to be TOS in the HLT2 topological lines then we get the sample shown in Fig. 7(right). All that is left is a very pure B mass peak. Figure 8 shows the results following the same procedure to add an additional track and cut around the K^* mass. Again, the full sample is dominated by background but the topological TOS sample is very clean.

Triggering on $B \to DX$ decays is more difficult due to the presence of the long-lived D. A very pure $D^* \to D\pi$ sample is obtained from the exclusive 2-body charm trigger lines. We then add a pion (with IP$\chi^2 > 4$) and require that the B points and has a decent vertex χ^2. Figure 9 shows full and topological TOS samples. The high-mass side-band is very small after requiring topological trigger TOS. The lower-mass side-band certainly contains a lot of partially reconstructed $B \to D^*X$ decays. The efficiency is also clearly very good. From these studies we conclude that while rigorously determining the $pp \to b\bar{b}X$ purity coming out of the HLT2 topological lines would be difficult, we can certainly say that the purity is very high.

From Monte Carlo we can estimate the $pp \to b\bar{b}X$ purity at each stage of the trigger which yields the following:

$pp \to b\bar{b}$X purity (from MC)

$$\boxed{1.0 \pm 0.2\%} \xrightarrow[\text{L0}]{} \boxed{2.9 \pm 0.6\%} \xrightarrow[\text{1-Track}]{} \boxed{8.9 \pm 1.9\%} \xrightarrow[\text{Topo}]{} \boxed{100 \pm 24\%}$$

These numbers use the visible [8] and $b\bar{b}$ [9] cross sections measured by LHCb and PYTHIA

[1] DiMuonJPsi is a dedicated J/ψ trigger which selects high transverse momentum DiMuon pairs near the J/ψ mass, described more fully in [7].

to generate the required samples of each type of event. Figure 10 shows the BBDT response for minimum bias (MB) 2010 LHCb data along with $pp \to c\bar{c}X, b\bar{b}X$ and MB Monte Carlo. The normalization again uses cross section measurements from LHCb; *i.e.*, the Monte Carlo is not normalized to the data. The agreement of the size and shape of the MB data above 0.2 with the $pp \to b\bar{b}X$ Monte Carlo suggests a large $b\bar{b}$ purity.

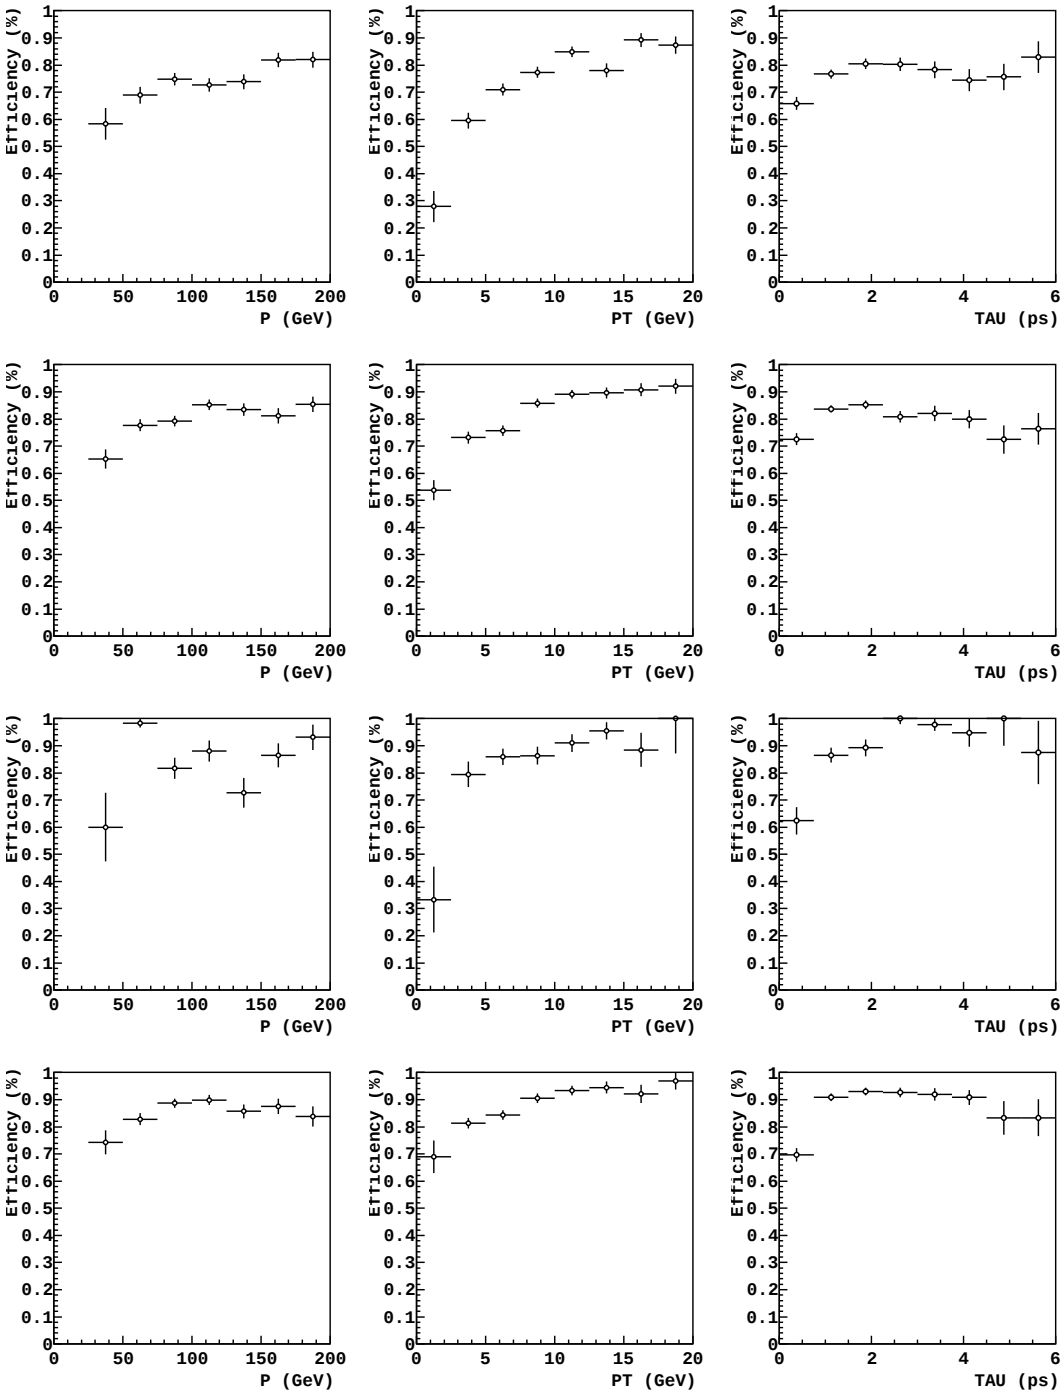

Figure 5: Efficiency of the HLT2 topological trigger (standard lines) for $B^0 \to D^+\pi^-$, $B^+ \to D^0\pi^-$, $B^0 \to J/\psi K^{*0}$, and $B^+ \to J/\psi K^+$ decays (top to bottom) as a function of B momentum, transverse momentum, and lifetime (left to right). This efficiency is measured relative to offline selected events which are TOS in the Hlt1TrackAllL0 trigger line.

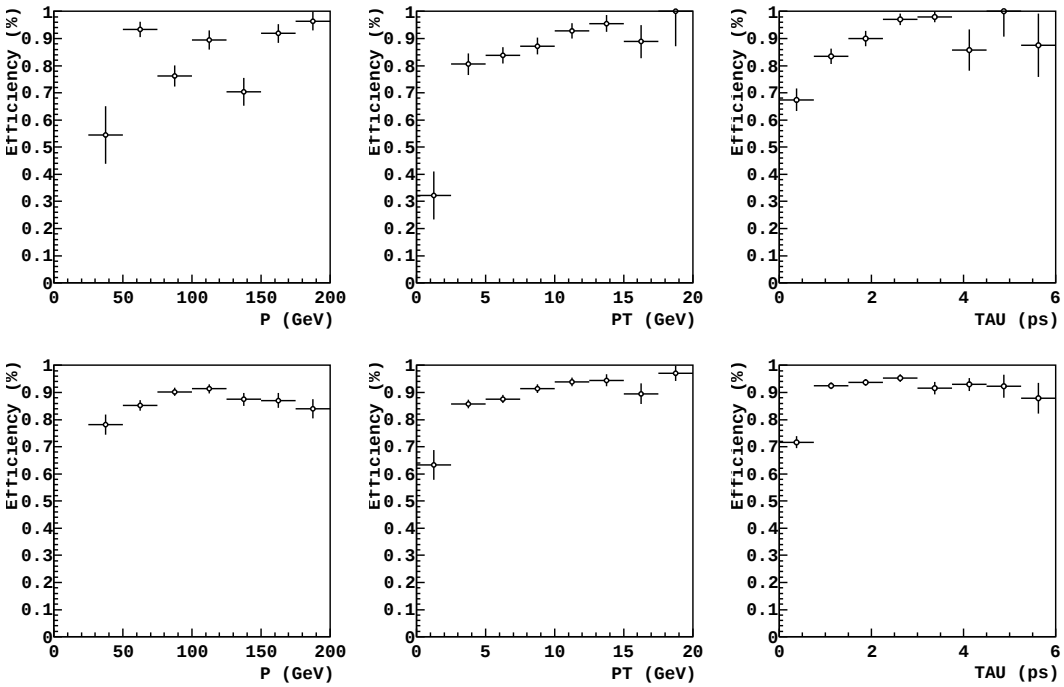

Figure 6: Efficiency of the logical OR of the HLT2 standard and μ-topological triggers for $B^0 \to J/\psi K^{*0}$ and $B^+ \to J/\psi K^+$ decays (top to bottom) as a function of B momentum, transverse momentum, and lifetime (left to right). This efficiency is measured relative to offline selected events which are TOS in either the Hlt1TrackAllL0 or the Hlt1TrackMuonL0 trigger line.

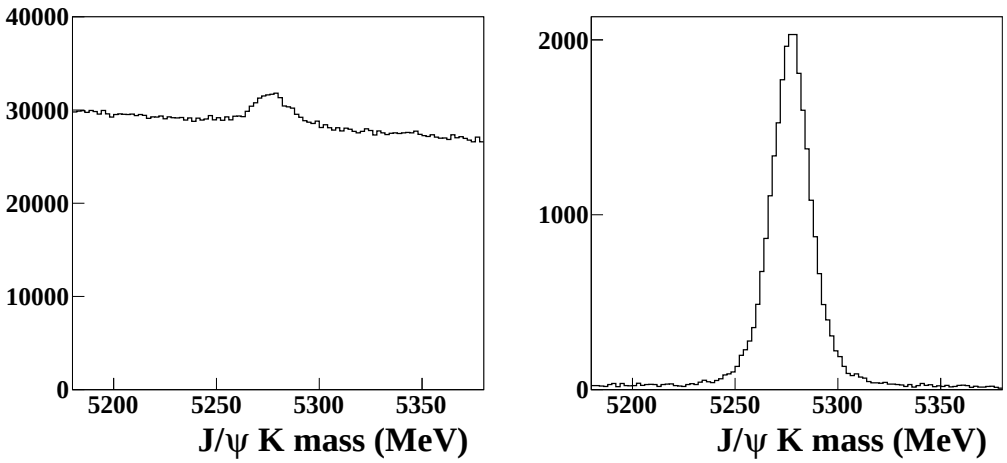

Figure 7: $B \to J/\psi K$ candidates using J/ψ's that are TOS in the DiMuonJPsi line combined with a Kaon with some very loose cuts (see text for details): (left) all candidates; (right) candidates that are TOS in the HLT2 topological trigger.

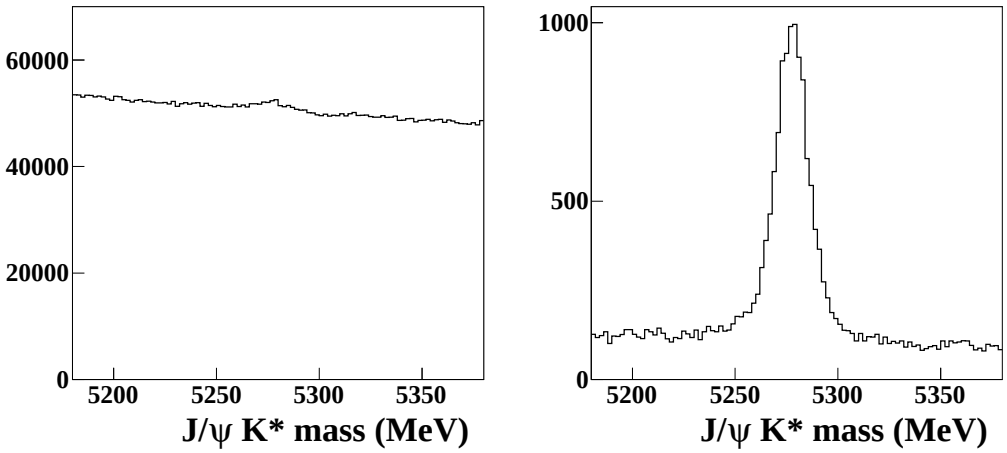

Figure 8: $B \to J/\psi K^*$ candidates using J/ψ's that are TOS in the DiMuonJPsi line combined with a Kaon with some very loose cuts (see text for details): (left) all candidates; (right) candidates that are TOS in the HLT2 topological trigger.

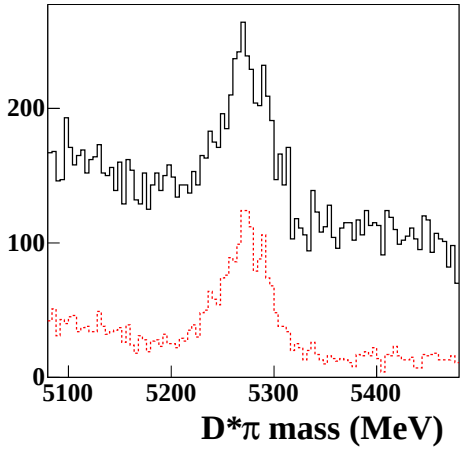

Figure 9: $B \to D^*\pi$ candidates using D^*'s that are TOS in the Charm lines combined with a pion with some very loose cuts (see text for details): (solid black) all candidates; (dashed red) candidates that are TOS in the HLT2 topological trigger.

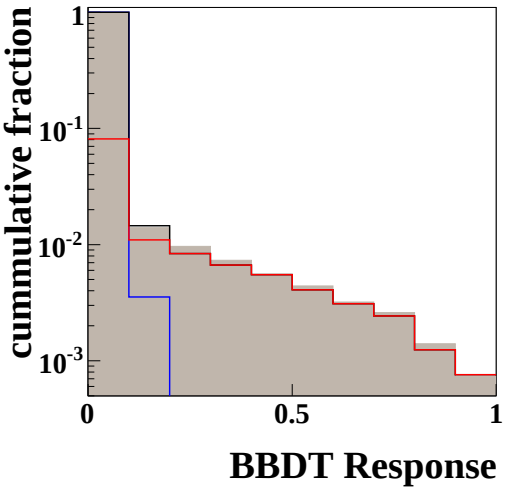

Figure 10: Response from the BBDT for minimum bias LHC*b* 2010 data (shaded grey), $pp \to c\bar{c}X$ Monte Carlo (blue), $pp \to b\bar{b}X$ Monte Carlo (red) and all minimum bias Monte Carlo (black). The Monte Carlo is not normalized to the data (see text for details). *N.b.*, no muon or electron requirements were used when making this plot.

Table 6: The TOS efficiency of each trigger stage for the decay modes listed in Sec. 3, relative to the offline selections defined in Tab. 1. HLT2 efficiencies are listed relative to HLT1 efficiencies. In the case of the topological trigger, the efficiency is the OR of the (2,3,4) body topological triggers. The total HLT efficiency is listed for the combination of (Hlt1TrackAllL0 and Hlt2TopoBBDT) TOS in the case of the hadron modes, and the combination (Hlt1Track(Muon or AllL0) and Hlt2TopoMuBBDT) TOS in the case of the muon modes. The final line gives the TIS efficiency for each channel, measured with respect to events TOS in the 1Track and Topological triggers. All uncertainties are purely statistical.

Efficiency	$B^0 \to J/\psi K^{*0}$	$B^+ \to J/\psi K^+$	$B^0 \to D^+\pi^-$	$B^+ \to D^0\pi^-$
Hlt1TrackAllL0	$(78 \pm 5)\%$	$(79 \pm 2.5)\%$	$(83 \pm 2)\%$	$(84 \pm 2)\%$
Hlt1TrackMuon	$(81 \pm 5)\%$	$(76 \pm 2.5)\%$	N/A	N/A
Hlt1Track(Muon or AllL0)	$(88 \pm 5)\%$	$(86 \pm 2.5)\%$	N/A	N/A
Hlt2TopoBBDT	$(86 \pm 5)\%$	$(87 \pm 2.5)\%$	$(75 \pm 2)\%$	$(81 \pm 2)\%$
Hlt2Topo(Mu or BBDT)	$(87 \pm 5)\%$	$(90 \pm 2.5)\%$	N/A	N/A
Total HLT	$(78 \pm 5)\%$	$(77 \pm 2.5)\%$	$(63 \pm 2)\%$	$(68 \pm 2)\%$
TIS Efficiency	$(3.2 \pm 0.2)\%$	$(2.9 \pm 0.1)\%$	$(4.0 \pm 0.1)\%$	$(3.6 \pm 0.1)\%$

8 Combined HLT Performance

In order to summarize the HLT performance, Table 6 lists the TOS efficiency of each trigger stage for the decay modes listed in Sec. 3, relative to the offline selections defined in Tab. 1. Unlike the earlier plots, which binned the efficiency in variables of interests, the efficiencies quoted here are integrated over the full range of these variables.

9 Conclusion

The current design of the LHCb trigger, and in particular the reliance on a few highly inclusive trigger lines, is a direct result of the experience gained commissioning the detector during 2010. It is a fundamentally different trigger to anything deployed in previous experiments, particularly in its reliance on tight track quality cuts and the use of multivariate selection criteria. The HLT inclusive triggers have been demonstrated to reduce the L0 trigger output rate of 1 MHz to a rate of 1 kHz for writing to tape, while increasing the proportion of events containing a b-quark from around 3% at the L0 stage to almost 100% at the output of the HLT. Furthermore the efficiency of the triggers has been measured on data for a representative subsample of B decay modes as a function of the B momentum, transverse momentum, and lifetime. The total HLT efficiency is found to be 60-80% for a broad range of topologies.

Acknowledgments

We would like to thank Hans Dijkstra for many useful discussions on the design and implementation of these triggers.

References

[1] **LHCb** Collaboration, A. Alves et al., *The LHCb Detector at the LHC*, *JINST* **3** (2008) S08005.

[2] V. V. Gligorov, *A single track HLT1 trigger*, Tech. Rep. LHCb-PUB-2011-003, CERN, Geneva, Jan, 2011.

[3] M. Williams et al., *The HLT2 Topological Lines*, Tech. Rep. LHCb-PUB-2011-002, CERN, Geneva, Jan, 2011.

[4] H. Dijkstra, N. Tuning, and N. Brook, *Some remarks on systematic effects of the trigger and event generator studies*, Tech. Rep. LHCb-2003-157, CERN, Geneva, Dec, 2003.

[5] B. Souza de Paula, *Studies on Systematic Effects of the Trigger on Flavour Tagging at the Generator Level*, Tech. Rep. LHCb-PUB-2009-014. CERN-LHCb-PUB-2009-014. LHCb-2006-046, CERN, Geneva, Aug, 2009.

[6] G. Lanfranchi, X. Cid Vidal, S. Furcas, M. Gandelman, J. A. Hernando, J. H. Lopez, E. Polycarpo, and A. Sarti, *The muon identification procedure of the lhcb experiment for the first data*, Tech. Rep. LHCb-PUB-2009-013. CERN-LHCb-PUB-2009-013, CERN, Geneva, Aug, 2009.

[7] **LHCb** Collaboration, R. Aaij and J. Albrecht, *Muon triggers in the High Level Trigger of LHCb*, *LHCb-PUB-2011-017* (2011).

[8] **LHCb** Collaboration, R. Aaij et al., *Absolute luminosity measurements with the LHCb detector at the LHC*, *To be Published* (2011).

[9] **LHCb** Collaboration, R. Aaij et al., *Measurement of J/psi production in pp collisions at sqrt(s)=7 TeV*, *Eur.Phys.J.* **C71** (2011) 1645, [arXiv:1103.0423].