LHCb Note 2005-047, TRIG & COMP
LPHE Note 2005-015
Revision 2: August 9, 2006

Exclusive HLT Performance

Luis Fernández[1]
LPHE, EPFL

Patrick Koppenburg[2]
CERN

Abstract

The present note describes the performance of the HLT selections, the exclusive b and the D^* streams, using the implementation of the HLT introduced in [1].

A comparison between the on-line and off-line reconstructions is made and the HLT efficiencies for the core physics channels are given, as well as the corresponding minimum-bias output rates and timing performance. Different sources of inefficiency are discussed and identified to motivate future improvements.

[1] Luis.Fernandez@epfl.ch
[2] Patrick.Koppenburg@cern.ch

Contents

1 Introduction

The LHCb High-Level-trigger (HLT) will have an output rate of 2 kHz and is divided in four streams:

Exclusive b (~ 200 Hz);

D^* (~ 300 Hz);

Di-muon (~ 600 Hz);

Inclusive b$\rightarrow\mu$ (~ 900 Hz).

Each of these streams will irrespectively of each other's decision send their selected events to storage. In this note we assume that the entire HLT chain, on-line reconstruction and selection, will be running on the on-line farm at a rate of 40 kHz and using about 400 CPUs in the initial strategy with a Level-1 trigger. The average HLT processing time should thus not exceed 10 ms per event. The change of scheme to the 1 MHz readout, yet to be implemented, will nevertheless not affect much the exclusive b and D^* streams as they will anyway be executed after pre-HLT algorithms.

The exclusive b HLT should efficiently select all final states relevant to the LHCb physics program. For the first implementation a list of ten core channels has been defined:

1. $B_s \rightarrow D_s h$
2. $B_d \rightarrow J/\psi K^0_S$
3. $B_s \rightarrow J/\psi\phi$
4. $B_q \rightarrow hh$
5. $B_d \rightarrow D^0 K^{*0}$
6. $B_s \rightarrow \phi\phi$
7. $B_d \rightarrow D^*\pi$
8. $B_d \rightarrow \mu^+\mu^- K^{*0}$
9. $B_s \rightarrow \mu^+\mu^-$
10. $B_s \rightarrow \phi\gamma$

The above list is meant to evolve and many more channels could be added. For instance new HLT selections involving electrons (e.g. $B \rightarrow K^{(*)}ee$) or π^0 ($B \rightarrow \rho\pi$) or K^0_S ($B \rightarrow \phi K^0_S$). Also the use of additional inclusive streams (ϕ, D_s, ...) is under study.

The present note describes the performance of the HLT exclusive b selections and the D^* streams, based on the design and implementation choices of [1], in the framework described in [2]. This note reports on the first complete and successful use of the on-line reconstruction to exclusively select signal channels, enhance the beauty and charm quark content and reject minimum-bias events to the design rates. The choice of selection criteria were chosen to prove the feasibility of the prototype and are suggested to revision by the different physics working groups.

We introduce the on-line environment and define the quantities used in Section (2), followed by a description of the on-line tracking performance in Section (3). In Sections (4) and (5) we extensively describe the HLT selections giving their selection criteria and the HLT performance in terms of signal efficiencies, output rates and timing. Finally we identify throughout this note some weaknesses of the current on-line reconstruction/selections.

2 On-line and HLT Environment

The design and the implementation choices of the HLT selections are described in [1]. This note is based on the HLT implementation for the Data Challenge 04 (DC04) compatible software. The picture for the HLT selections won't drastically change with the 1 MHz readout scheme and the additional data available at this higher rate. However, many improvements in terms of tracking strategy and performance are expected as well as a new design for the generic HLT, in particular its merge with the present Level-1 trigger, hence removing the current 40 kHz boundary. In the new design alleys are foreseen, i.e. the confirmation of previously found triggering objects, with as spin-off an even more "vertical" trigger between the Level-0 and the generic HLT lines. The 1 MHz scheme also avoids the repetition of identical tasks, such as buffer decoding.

The HLT selections use for both the exclusive b and the inclusive D^* streams the same generic algorithms and tools allowing the complete selection of events steered by options [2]. The use of common tools throughout the HLT selections ensures maximal correlations with a limited amount of code used. This implementation enables the quick addition or removal of streams (e.g. inclusive ϕ stream) from the configuration database.

The main steps of the on-line reconstruction are described in Section (2.1) and the dedicated on-line tools are discussed in Section (2.2). The program versions, data samples and a few technical definitions are also given.

2.1 From Tracks to Particles

The first stage in the HLT processing is the on-line reconstruction [3] in parallel with the generic HLT [4]. Note that there were some improvements with respect to the on-line pattern recognition described in [3] and they are summarized in the Appendix (B).

The main part of the pattern recognition is actually done in the generic HLT sequence. The full tracking is first run on a limited set of tracks without performing the rest of the reconstruction unless the event is accepted by the generic HLT trigger. The combined tracking and generic HLT procedure is briefly described hereafter. Each tracking step implicitly includes a data preparation and raw buffer decoding.

Velo tracking and primary vertex search

Velo-RZ tracking:
: Velo-RZ tracks are created from 2-d tracking in the $r-z$ projection using aligned r-cluster triplets in consecutive r-sensors.

Velo-space tracking:
: Velo-space tracks are reconstructed by extending to 3-d the Velo-RZ tracks using information from the Velo ϕ-sensors.

Primary vertex reconstruction:
: The primary vertices are looked for using Velo-space tracks as seeds with an iterative fit procedure.

Generic HLT and full tracking

Stand-alone muon pattern recognition:
: Information from the muon system (M2-M5) is used to reconstruct and identify muons, independently of the standard on-line tracking.

Velo-TT tracking:
The matching of Velo space tracks to TT clusters is done assuming that the effect of the magnetic field can be replaced by a kick of the track at a fixed z. All Velo tracks are then extrapolated to the z-positions of the measurements and the distance perpendicular to the strips between the extrapolation and the measurement is projected onto a reference plane. The solutions are looked for among the accumulations in the projection plane. When the track's extrapolation is outside TT's sensitive area (e.g. beam pipe, TT gaps) and under some special conditions, Velo tracks might be passed directly to the forward tracking without any TT clusters (or momentum).

Forward tracking:
The long tracks search is performed by matching Velo-space tracks to clusters in the T stations (inner and outer trackers), using the momentum estimate (if any) obtained from Velo-TT tracks to possibly restrict the search window for the T clusters, at a cost of some inefficiencies due to wrong TT clusters or acceptance. Using a seed in the T stations, the trajectory is parameterized depending on the track parameters, deflection and the field. All measurements compatible with this trajectory are collected and an iterative procedure removes the worse clusters along the trajectory.

Track errors parameterization:
The errors from the tracking are not used and the covariance matrix of all long tracks is replaced by a simplified and parameterized version as a function of the transverse momentum [5]. Note that the upstream tracks' errors are not overwritten.

HLT muon streams:
The single muon and di-muons are searched for the HLT decision of the inclusive b$\rightarrow\mu$ single muon ($p_{\mathrm{T}}(\mu) > 3\,\mathrm{GeV/c}$, $\mathrm{sIPS}(\mu) > 3$) and di-muon ($m_{\mu\mu} > 2.5\,\mathrm{GeV/c^2}$, $\chi^2_{\mu\mu} < 5$, $\mathrm{IP}_{\mu\mu} > 0.1$ mm) streams.

Generic HLT:
The generic HLT redoes the Level-1 generic line with better momentum determination. The decision is based on high transverse momentum and large impact parameter tracks with respect to all the primary vertices. The decision is also positive if either of the muon lines is triggered.

This first part of the HLT processing is followed by the creation of particles to be later used by the HLT selections. Muons are provided as di-muons (directly retrieved from the generic HLT) and all the upstream and forward tracks are made as both pions and kaons. Finally photons are made from on-line electromagnetic calorimeter (ECAL) clusters. Note that the possibility of creating protons (by just assigning the proton mass to tracks), electrons or resolved π^0 are available but so far not in use.

The possibility of using particle identification provided by the on-line RICH is at present not used. The latter allows to save CPU time by reducing combinatorics. However, the lower combinatorics do not compensate for the time lost in building the RICH information.

2.2 Dedicated On-line Software

The HLT will be running in the on-line farm with a limited processing time to reconstruct all the necessary information and to take the decision whether or not sending an event to storage. The same algorithms as in the off-line where an "unlimited" time is available can therefore not be run. In order to cope with this time restriction, we need fast dedicated

on-line software which will unavoidably give inefficiencies when comparing the on-line performance to the off-line one.

We list hereafter the main sources of inefficiency or differences that were identified for the HLT selections:

- On-line and off-line pattern recognitions are too different, yielding track finding inefficiencies for the tracks used in the off-line analyses.
- The differences in the primary vertex reconstruction strategy may give events with a different number of on-line and off-line reconstructed primaries. Moreover, a missed primary can fake the presence of a secondary vertex.
- The tracks parameters are not identical to the off-line ones (e.g. worse momentum resolution) and their covariance matrices are not comparable. In the HLT a simplified model for the track uncertainties is used [5], whereas in the off-line they are obtained from the track fitting.
- The HLT geometrical and vertexing tools assume straight-line trajectories (not really a source of inefficiency) for the tracks and use a constant parameterized covariance matrix [6]. The latter yields a few percent inefficiency in the vertexing for long-lived particles.
- The choice of selection criteria also degrades the HLT performance. For instance transverse momentum and impact parameter cuts are needed for timing reasons and applied to all long tracks in the HLT whereas several off-line selections do not apply any or very soft requirements on these variables.

All these differences between the on-line and off-line software used to perform physics analyses will inevitably lead to inefficiencies. Each of these issues will be addressed later in this note.

2.3 Selection Criteria Definitions

Ideally, in order to ensure a maximal correlation and efficiency with the off-line selections, the HLT selections should mimic what is done in the off-line with looser requirements whenever the output rate allows to. Since the computing time is limited and not to have too selection-specific kind of selection criteria ("cuts"), a limited set of simple cuts is used in the HLT. This allows for instance the introduction of shared composite particles with soft requirements to be later reused and refined to form the b-hadron candidates in the specific streams. In this way the reconstruction of the same particle several times per event can be avoided. The selection criteria are applied through a series of standard and reproducible filters, as described in [2]. Note that as most of the final selections apply dedicated criteria with different definitions (e.g. cuts related to the choice of the origin vertex of the b-hadron), the HLT cannot be fully correlated with the off-line selections.

The list of selection criteria used in the HLT is defined hereafter:

- Momentum p or transverse momentum p_{T} of a particle.
- Impact parameter IP or impact parameter significance IPS $\equiv \mathrm{IP}/\sigma_{\mathrm{IP}}$, where σ_{IP} denotes the error on IP. Impact parameter cuts are unsigned and always applied with respect to all reconstructed primary vertices, which is equivalent to cutting on the smallest impact parameter significance denoted by sIPS.

- Difference in mass δm between the reconstructed particle and the mass from the `ParticlePropertySvc` (generated Monte Carlo value).

- $\chi^2 \times$ ndf of the mass unconstrained vertex fit. As this quantity is not normalized to the number of degrees of freedom (ndf), it does not represent a χ^2 probability. Ignoring this, we will henceforth refer to $\chi^2 \times$ ndf simply as the χ^2 of the vertex fit.

- Flight distance FD or flight distance significance FS (ratio of FD to its computed error). The flight distance is defined as the distance between the decay vertex of a composite particle and the reconstructed primary vertices. The cut is on the minimal unsigned distance and is either applied with respect to all primaries (i.e. requiring the smallest FD or FS to be larger than some cut), or with respect to the primary vertex for which the composite particle has the smallest IPS.

- Pointing angle $\theta_{p,\mathrm{F}}$ defined as the angle between the reconstructed momentum $\vec{p}$ and the direction of flight $\vec{\mathrm{F}}$ for particles originating from a primary vertex (e.g. b-hadrons) chosen as the one for which the IPS of the composite is the smallest.

An identified source of background corresponds to events where the primary vertex reconstruction fails to find a true non-elastic primary interaction. These events are potentially dangerous since they offer an important collection of tracks of high momentum with presumably detached impact parameter with respect to all other reconstructed primary vertices. Due to the large combinatorics involved these missed primary vertices may be identified as secondaries and trigger the HLT since the standard filter criteria are in this case not discriminant enough.

A simple way for getting rid of theses cases is to apply vertex isolation cuts by requiring that only a limited number of tracks are compatible with the b-vertex, according to some criterion such as impact parameter significance. Isolation cuts enable to relax the standard selection criteria while being fully efficient on signal and keeping the output rate under control. This therefore improves the HLT efficiency on b-events significantly. Figure (1) shows the effect of missing primaries with the old primary vertex reconstruction. We see that a large fraction of the HLT triggered b-candidates in minimum-bias events have numerous high $p_{\mathrm{T}} > 300\,\mathrm{MeV/c}$ and sIPS > 2 particles pointing to it, within 4-sigmas in IP. Given the recent improvements in the primary vertex reconstruction, this effect becomes much less problematic and it was decided not to apply any isolation cut but to keep them as an extra handle to reduce the background [7].

2.4 Software and Data

All the results presented in this note were obtained with DaVinci `v12r15`, with `Hlt/HltSelections v4r1p1`, `Hlt/HltSelChecker v4r1` and `Hlt/HltInclusive v1r1`.

The Data Challenge 04 (DC04) data samples used are:

- Minimum-bias DC04-v1: 32,667 stripped events (with an obsolete DaVinci version) at 40 kHz. This sample was used to tune all the selection criteria, however all the results (e.g. rates, plots) presented in this note are based on the DC04-v2 sample.

- Minimum-bias DC04-v2: 131,320 Level-0 and Level-1 stripped events (with DaVinci v12r15) from an original sample of 48,254,400 events[3]. The stripped sample corresponds to ~ 3.28 s of LHCb running. This is the minimum-bias sample used to

[3]This is the result of our private stripping of the original sample. Many thanks to José-Angel Hernando for participating in this task.

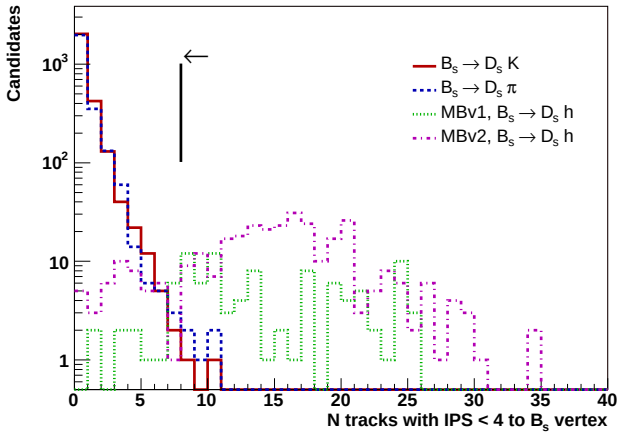

Figure 1: Old primary vertex reconstruction and number of b-candidates selected by the exclusive HLT $B_s \to D_s h$ selection as a function of the number of high p_T and high sIPS particles pointing back to the B_s vertex. The solid (red) and dotted (blue) lines are on $B_s \to D_s K$ and $B_s \to D_s \pi$ signals respectively whereas the dashed (green) and the dashed-dotted (magenta) lines are for minimum-bias. The histograms are not normalized.

assess the HLT performance presented in this note. Applying the generic HLT to this sample leaves 28,458 events which is equivalent to (8668.3 ± 45.5) Hz output rate.

- Signal DC04-v1 data: several signal types, after Level-0, Level-1 and corresponding specific off-line selection (here "TDR selections" were applied). The performance (efficiencies, resolutions, ...) on signal for the dedicated streams was tested with the following signal data, where the charged conjugate modes are implicit:
  - $B_s \to D_s h$ selection using $B_s \to D_s \pi$ and $B_s \to D_s K$ data, where $D_s \to K^+K^-\pi$;
  - $B_q \to J/\psi X$ selection using $B_s \to J/\psi\phi$ data with $J/\psi \to \mu^+\mu^-$ and $\phi \to K^+K^-$;
  - $B_q \to hh$ selection using $B_d \to \pi^+\pi^-$, $B_d \to K^+\pi^-$, $B_s \to K^+K^-$, and $B_s \to \pi^+K^-$ data;
  - $B_d \to D^0K^{*0}$ selection using $B_d \to D^0K^{*0}$ data where $D^0 \to K^+\pi^-$ and $K^{*0} \to K^+\pi^-$;
  - $B_s \to \phi\phi$ selection using $B_s \to \phi\phi$ data with $\phi \to K^+K^-$;
  - $B_d \to D^*\pi$ selection using $B_d \to D^*\pi$ where $D^* \to D^0\,\pi$ with $D^0 \to K^+\pi^-$;
  - $B_d \to \mu^+\mu^-K^{*0}$ selection using $K^{*0} \to K^+\pi^-$;
  - $B_s \to \mu^+\mu^-$ selection using $B_s \to \mu^+\mu^-$ data;
  - $B_s \to \phi\gamma$ selection using $B_s \to \phi\gamma$ data with $\phi \to K^+K^-$;
  - D^* stream selection using $B_d \to D^*\pi$ where $D^* \to D^0\,\pi$ with $D^0 \to K^+\pi^-$ for the efficiency on signal and some $B_d \to D^0K^{*0}$ data with $D^0 \to K^+K^-$ for mass windows checks.

  For $B_s \to D_s h$ and $B_q \to hh$ selections the samples used for the different flavors contain roughly the same number of events in order to equally assess the performance on all the decays.

3 On-line Reconstruction Performance

The performance of the HLT strongly relies on the quality and efficiency of the on-line tracking which outputs are used as ingredients for the HLT selections. The main tasks of the reconstruction are:

- Find all necessary tracks for the HLT decision, especially the signal tracks that will be used in an off-line analysis;
- Provide accurate track parameters and reliable errors;
- Find *all* possible primary vertices.

All these items will then be used to reconstruct the largest number of interesting events within the output rates specifications by reducing the minimum-bias.

We present in this Section a few aspects of the on-line reconstruction performance and discuss some of the tracking issues. In Section (3.1) we give on-line track finding performance on signal events for several distinct b-topologies and on minimum-bias. The track parameters and errors are shown in Section (3.2), and a comparison between on-line and off-line reconstruction quality is done in Section (3.3). Note that all the results shown are after the HLT-generic trigger.

3.1 On-line Track Finding

The track finding efficiencies are determined for three purely hadronic channels with a distinct topology: $B_d \to \pi^+\pi^-$, $B_s \to \phi\ (K^+\ K^-)\phi\ (K^+\ K^-)$ and $B_s \to D_s\ (K^+\ K^-\ \pi)\pi$. The timing and ghost rates are measured with the whole minimum-bias sample. For the timing we assume in this Section that we redo all the Velo-TT and forward tracking after the HLT-generic and ignore the fact that part of the forward tracking was already made. Moreover the quoted times were not determined on dedicated machines, hence they just represent an approximate order of magnitude. The detailed and accurate speed performance of each step will be given in Section (5.2.2).

A measure of the track finding performance is given by the ghosts rates, since these tracks have many wrong clusters and they will directly act on the later stages of the reconstruction. The fractions of ghosts in minimum-bias for each track type are given in Table (1), including the result for long tracks without any TT information. As expected Velo-TT tracks have the highest ghost rate since TT has only four planes and therefore there are not many constraints to assign the correct clusters to a Velo track. The use of TT actually allows to reduce the ghosts in the long tracks by almost a factor two.

The reduction of long track ghost rate when using TT reflects the fact that the momentum information obtained from Velo-TT tracks can be used in the forward tracking to narrow its search regions and considerably reduce its processing time. This can be seen in Table (2) where the timing for processing Velo-TT pattern recognition and the forward tracking are shown. The quoted values are for both the buffer decoding and the tracking parts, and the equivalent time when skipping TT is shown for comparison. We see that the use of TT gives a total time reduction of at least a factor two.

The use of TT offers a great speed advantage but will lead to inefficiencies first in Velo-TT tracks since the limited constraints to assign the correct clusters to a Velo track will lead to matching errors and moreover in the long track extrapolation since providing a wrong momentum estimate may cause the forward tracking to look for ST clusters in the wrong region. Before discussing the results, let us remind here two features of the association to the Monte Carlo truth for on-line tracks. First the quoted efficiencies for Velo-TT tracks

Table 1: Ghost rates determined from ~ 3.28 s minimum-bias events, for each type of on-line track. The samples contain over 1 M tracks each.

	Track types				
Minimum-bias	Velo-RZ	Velo-space	Velo-TT	Long	Long (no TT)
Ghost rate [%]	12.3	9.4	16.6	5.8	11.0

Table 2: Times for the Velo-TT and forward tracking parts as measured on minimum-bias with a 2.8 GHz Xeon processor, on a machine affected by the load. The quoted values include the times for buffer decoding and doing the full tracking for each event.

	Speed	
Algorithm	Standard on-line [ms]	On-line without TT [ms]
Velo-TT	1.9	-
Forward	8.4	22.3
Combined	10.3	22.3

include the effect of a track being later found by the forward tracking without any TT clusters. This is relevant if one wants to use upstream tracks in the HLT since the efficiency for finding the signal tracks with momentum information is a bit lower than the quoted ones. The second remark is somehow related to the first: in the association to the MC truth of the long tracks, the fraction of wrong TT hits is ignored for the decision.

The track finding efficiencies for the different signals are given in the Tables (3), (4) and (5) for each of the track types. Note that the on-line Velo tracking *should* be 100% efficient on off-line reconstructed tracks, as the same code is run. The inefficiencies are here due to different versions of the code used in the off-line reconstruction (stored on `DST`) and what is run in the HLT. The tracking results show what one could naively expect: the per-track efficiency decreases with the number of tracks involved in the b-decay. Low multiplicity decays to charged final states will have in average a higher momentum with a more straight-line like trajectory making them easier to reconstruct in Velo-TT. Moreover higher multiplicity decays will see their final states undergo more multiple scattering resulting in a worse Velo slope determination, strongly affecting the forward tracking performance.

For decays to purely hadronic charged final states involving only long tracks the fraction of events supposedly lost due to tracking inefficiencies and before applying any specific HLT selection is:

$\sim 2\%$ per track, ~ 3 - 4% for a 2-prong decay ($B_d \to \pi^+\pi^-$);

$\sim 4\%$ per track, $\sim 13\%$ for a 4-prong decay ($B_s \to \phi\phi$, $B_s \to D_s\pi$);

$\sim 5\%$ per track, $\sim 26\%$ for a 6-prong decay ($B_s \to \phi\ (K^+\ K^-)\ \eta_c\ (4h)$).

when checking that all tracks are reconstructed for the per-event efficiencies. A larger loss is expected for channels using upstream tracks, such as the slow π from D^*, since these tracks have in average lower momentum.

Table 3: Tracking efficiencies on signal $B_d \to \pi^+\pi^-$ events for each type of on-line track. Results are for $N^{evt}_{off} = 1225$ events after HLT-generic and off-line selection, for events with all signal final states ($N^{tr}_{fin} = 2$) off-line reconstructed. The uncertainties are statistical.

$B_d \to \pi^+\pi^-$	Per track		Per event	
Track type	N^{tr}_{rec}	ε^{tr}_{rec} [%]	N^{evt}_{rec}	ε^{evt}_{rec} [%]
Velo-RZ	2447	99.9 ± 0.1	1222	99.8 ± 0.1
Velo-space	2438	99.5 ± 0.1	1213	99.0 ± 0.3
Velo-TT	2427	99.1 ± 0.2	1202	98.1 ± 0.4
Long	2411	98.4 ± 0.3	1186	96.8 ± 0.5
Long (no TT)	2415	98.6 ± 0.2	1191	97.2 ± 0.5

Table 4: Tracking efficiencies on signal $B_s \to \phi\phi$ events for each type of on-line track. Results are for $N^{evt}_{off} = 943$ events after HLT-generic and off-line selection, for events with all signal final states ($N^{tr}_{fin} = 4$) off-line reconstructed. The uncertainties are statistical.

$B_s \to \phi\phi$	Per track		Per event	
Track type	N^{tr}_{rec}	ε^{tr}_{rec} [%]	N^{evt}_{rec}	ε^{evt}_{rec} [%]
Velo-RZ	3754	99.5 ± 0.1	926	98.2 ± 0.4
Velo-space	3715	98.5 ± 0.2	895	94.9 ± 0.7
Velo-TT	3656	96.9 ± 0.3	845	89.6 ± 1.0
Long	3609	95.7 ± 0.3	809	85.8 ± 1.1
Long (no TT)	3638	96.4 ± 0.3	829	87.9 ± 1.1

Table 5: Tracking efficiencies on signal $B_s \to D_s\pi$ events for each type of on-line track. Results are for $N^{evt}_{off} = 2580$ events after HLT-generic and off-line selection, for events with all signal final states ($N^{tr}_{fin} = 4$) off-line reconstructed. The uncertainties are statistical.

$B_s \to D_s\pi$	Per track		Per event	
Track type	N^{tr}_{rec}	ε^{tr}_{rec} [%]	N^{evt}_{rec}	ε^{evt}_{rec} [%]
Velo-RZ	10306	99.9 ± 0.0	2568	99.5 ± 0.1
Velo-space	10237	99.2 ± 0.1	2502	97.0 ± 0.3
Velo-TT	10114	98.0 ± 0.1	2386	92.5 ± 0.5
Long	9955	96.5 ± 0.2	2241	86.9 ± 0.7
Long (no TT)	10042	97.3 ± 0.2	2317	89.8 ± 0.6

We see from Tables (3), (4), (5) that skipping TT improves the per track efficiencies by $\sim 0.2\%$ for a 2-prong and ~ 0.7 - 0.8% for a 4-prong, but for an additional time cost of a factor two, as shown in Table (2). The few percent lost are very hard to recover with the current implementation of the tracking, and a complete revision of the forward algorithm will have to be done for the next Data Challenge. For instance one could only use TT information to lower the ghost rate, but this implies serious improvements in speed for the forward tracking. The current implementation of Velo-TT is therefore the result of a compromise between speed and efficiency. Accepting more tracks at this stage will only slow the forward tracking, or even cause inefficiencies when providing too many ghosts. The Velo-TT or the forward tracking are not the only sources of inefficiency as we also see that there is ~ 0.5 - 1.0% loss in making Velo-space tracks. However, the latter loss is expected to disappear as the same Velo tracking code will be run on-line and off-line. Finally, we note that the ϕ decay products are more difficult to reconstruct at each step due to the small opening angle between the two kaons.

The investigation of TT inefficiencies shows that actually at least 60% of the missed tracks have less than 3 TT clusters, which is the minimum number of clusters required in Velo-TT. Moreover, the large majority of non-reconstructed tracks only have 2 TT clusters as a consequence of the vertical dead regions between the silicon sensors.

The use of TT in the on-line tracking improves the timing by at least a factor two, however it leads to inefficiencies that are difficult to recover and that would mean an additional time cost. If we want to have in the HLT all the tracks relevant to an off-line selection, then the on-line and off-line codes should have a better correlation and the use of TT should probably be avoided as an intermediate step before the extrapolation after the magnet. The use of common tools for the two pattern recognitions is under investigation for the next Data Challenge.

3.2 Track Parameters

The HLT selections apply geometrical selection criteria to reject minimum-bias background candidates originating from light quark flavor states. Criteria using the uncertainty offer a better signal–background discrimination while being more efficient compared to cutting on the bare quantity. We therefore need the knowledge of track uncertainties to be able to cut on the significance of variables.

The trajectory of a particle can be modeled as a collection of track segments which may be viewed as 'snapshots' of the particle's trajectory at a reference plane, and are referred to as track states. A track state is represented by a state vector depending on five track parameters at a constant z and its covariance matrix. The track parameters are given by the positions $x(z)$, $y(z)$, the slopes $t_x = \partial_z x$, $t_y = \partial_z y$ and the curvature q/p as a measure of the momentum p of a single charged particle ($q = \pm 1$) in the magnetic field.

The errors on the track parameters obtained from the on-line tracking not being reliable and over-estimated, they are recomputed from a parameterization[4] [5]. This method for the estimation of the track uncertainties is based on a fourth order polynomial parameterization as a function of p_{T}^{-1} of the x and y track parameters variances, for the first measured state of the track, i.e. the closest to the interaction point. This parameterization takes into account multiple scattering (p_{T}^{-1} dependence) and resolution effects (constant terms). The covariance matrix of each long track is then replaced with a simplified version and constant along the track ignoring all correlations. Only the x and y track parameter errors are non-null, neglecting the slopes and momentum elements of the track's covariance matrix as well. The track may be viewed as a straight line with a surrounding cylindrical

[4]We thank Hans Dijkstra and Hugo Ruiz for useful comments.

uncertainty envelope. Note that upstream tracks are not considered in the error parameterization.

The estimates for the track parameters together with their covariance matrices should be accurate enough to locate the primary vertices and to perform the reconstruction of decays in the HLT. The resolutions[5] and pulls[6] of the track parameters are important measures of the reconstruction quality since they will determine at which precision we can reconstruct b-decays, in particular secondary vertices and mass resolutions. The resolutions for the track parameters at the z coordinate of the first measured point are listed in Table (6), for the daughters of the $B_d \to \pi^+\pi^-$ decay. For the comparison, only fully associated B_d candidates for events with exactly one MC primary vertex, one reconstructed on-line and one off-line primaries are considered. The tracking errors are discussed in the next Subsection in terms of the impact parameter. Note that for the off-line, the errors are those obtained from the track fit whereas for the on-line the p_T parameterization is used.

Table 6: Resolutions (standard deviations of residuals) for the track parameters at the first measurement of the signal tracks in $B_d \to \pi^+\pi^-$. The values are the widths of single Gaussian fits.

Resolution	x [μm]	y [μm]	t_x [μrad]	t_y [μrad]	$\delta p/p$ [%]
On-line	16.1	16.3	0.25	0.11	0.79
Off-line	13.4	12.8	0.10	0.09	0.40

The on-line resolutions are a bit worse compared to the off-line, but in the acceptable range for the HLT needs. The on-line x-slope, though expected to have a worse resolution since in the bending plane, has a much worse resolution compared to the vertical slope. This is an artifact of the wrong correction of the Velo track x-slope when making long tracks to account for the residual magnetic field in the Velo.[7]

The long track momentum resolution is important since it directly translates to reconstructed mass resolutions. Moreover, momentum and mass cuts are primordial for the background rejection. We see that the track parameter momentum resolution, shown in Figure (2), is a factor two worse compared to the off-line tracks, but still quite accurate for the HLT purposes. The difference in performance is due to the purity of the tracks: in the on-line pattern recognition more wrong clusters are assigned to the tracks. Note that the momentum resolutions quoted here are for the daughters of a given B_d decay, after off-line selection and the generic HLT. The resolution vary with the origin of the tracks but the order of magnitude remains correct.

3.3 On-line Quality

The great advantage of the error parameterization approach is the introduction of fast geometrical and vertexing tools based on this error model [6], providing sufficiently good performance for the HLT selections. The transportation of the track states and matrix

[5]The resolution σ_x on a parameter x is the Gaussian width of the residual distribution $\delta x = x_{\rm rec} - x_{\rm true}$ between the reconstructed value $x_{\rm rec}$ and the true Monte Carlo value $x_{\rm true}$.

[6]The pull is defined as $\Sigma_x = \delta x/\sigma_x^{\rm err}$, where $\sigma_x^{\rm err}$ is the error on x from the covariance matrix. For an accurate estimate of x, Σ_x should follow a Normal Gaussian distribution.

[7]This effect is properly taken into account in `Trg/TrgForward v4r1`, but we do not use it this note. However, it has been checked that it only affects the HLT exclusive efficiencies at the per-mil level.

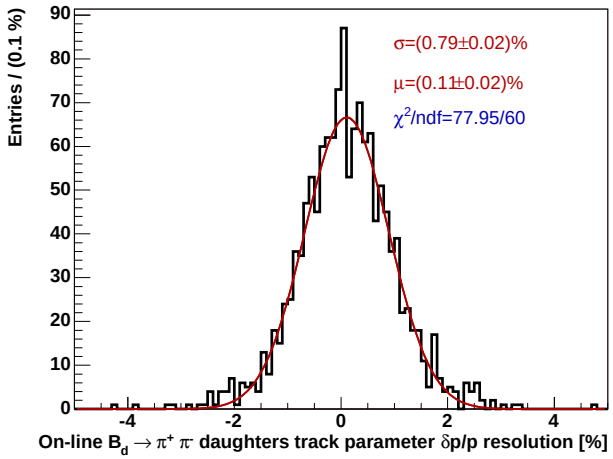

Figure 2: Momentum resolution [%] of $B_d \to \pi^+\pi^-$ on-line signal long tracks at the first measurement.

inversions become trivial given the simplified and constant covariance matrix, therefore improving considerably the timing.

We present in this Section a comparison between on-line and off-line of some of the important quantities related to a typical HLT selection. For the comparison we use again the $B_d \to \pi^+\pi^-$ decay as example and we only consider events with single primaries to avoid any ambiguity in the primary vertex assignment. All candidates are associated to the MC truth and are after off-line selection and the HLT-generic, and without applying any cut for the exclusive HLT selection. The error model of the previous Section is used for all on-line long tracks and tools.

An important geometrical quantity is the impact parameter defined as the shortest distance between a track and a vertex. Figure (3) shows the impact parameter and IP error distributions for the $B_d \to \pi^+\pi^-$ signal tracks. We see that the IP distributions for on-line and off-line are within the resolution very similar, meaning that the IP measurement is accurate in the HLT. The resolutions in impact parameter, computed as the width of the $(\mathrm{IP}_{\mathrm{rec}} - \mathrm{IP}_{\mathrm{true}})$ distribution, for the on-line and off-line are respectively $\sigma_{\mathrm{IP}}^{\mathrm{On}} = 20.0\,\mu$m and $\sigma_{\mathrm{IP}}^{\mathrm{Off}} = 16.5\,\mu$m. The errors on the impact parameter ($\sigma_{\mathrm{IP}}^{\mathrm{err}}$) are however very different, where the on-line has systematically its errors shifted towards higher values: $16.8\,\mu$m for the mean off-line error and $26.1\,\mu$m for the on-line. This is expected since we have a poorer resolution in the HLT, but the mean error should be compatible with the resolution for a correctly estimated uncertainty. We see that this is indeed the case for the off-line but not for the on-line where the mean error of $26.1\,\mu$m is deviated from the resolution by more than 20%.

The overestimation of the IP errors is confirmed when looking at the corresponding pull distributions in Figure (4). Whereas the off-line has correctly computed errors, the on-line pull shows a $\sim 22\%$ overestimation of the errors. This is a direct consequence of the simplified error which translates here into too large IP errors, for our b-tracks. The deviations of the pulls from unity for these b-decay tracks are a sign of an incomplete parameterization of the errors. This feature was already observed in Reference [5]. Looking at the residual dependences after the $1/p_{\mathrm{T}}$ correction, the main reasons for these too large errors are:

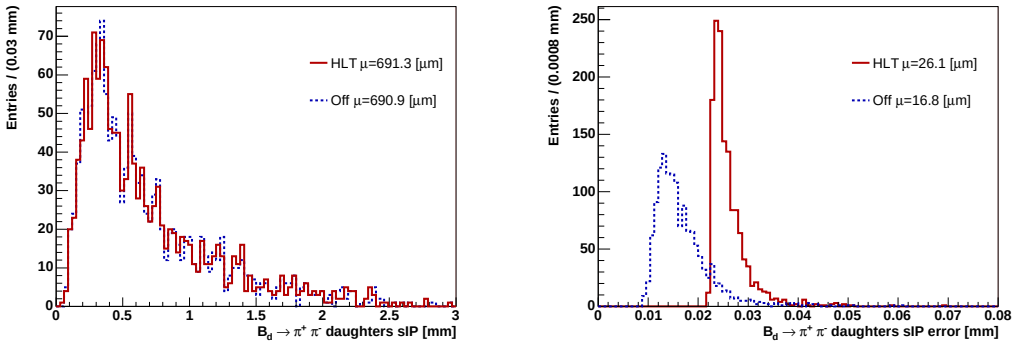

Figure 3: Left plot: sIP [mm] distributions for on-line tracks (red solid line) and off-line tracks (blue dotted line). Right plot: sIP error [mm] distributions for on-line tracks (red solid line) and off-line tracks (blue dotted line). Distributions are on $B_d \to \pi^+\pi^-$ events.

- The error parameterization ignores the primary vertex uncertainty. This effect increases with p_T and results in smaller pulls, which are here worsen by the fact that b-tracks are considered. This effect is of the order of $\sim 6\%$ for tracks with $p_T \gtrsim 4\,\mathrm{GeV/c}$, as it is the case of the $B_d \to \pi^+\pi^-$ final states. Thus the parameterization does not accurately reproduce the data when $1/p_T \lesssim 1\,\mathrm{GeV/c}^{-1}$.
- The parameterization has a dependence on the polar angle θ. Deviations of the IP pull from unity of the order of $\sim 20\%$ are observed for $\theta \gtrsim 100$ mrad. For $B_d \to \pi^+\pi^-$, the final states have in average $\theta \sim 130$ mrad.

We can however say that this procedure is satisfactory for HLT purposes where the main point is get to some estimate of tracking errors and not their precise measurement. Moreover, the large majority of tracks in minimum-bias have lower transverse momentum, compared to b-tracks, and in this range the parameterization works well. Given the residual dependences, a possible improvement would be the use of a look-up table rather than a 1-dimensional parameterization. A detailed study of the dependence on other parameters can be in [5].

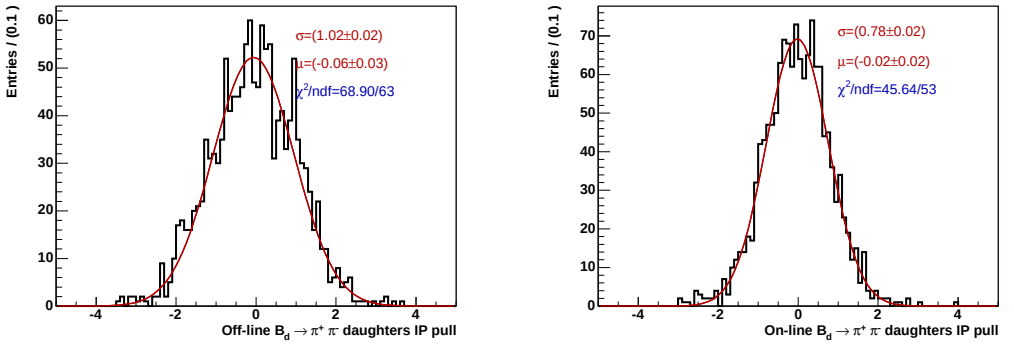

Figure 4: Impact parameter pull of the $B_d \to \pi^+\pi^-$ daughters. Left plot: off-line pull. Right plot: on-line pull. Single Gaussians are used to fit the standard deviations.

We must therefore be careful when directly comparing on-line and off-line significances

cuts. For instance the on-line errors are in average 70% larger than the off-line ones. Besides the fact that the on-line has too large errors, the primary vertex reconstruction efficiencies are not the same resulting in differences impact parameter distributions. In addition, the on-line has different errors when compared to the off-line due to a simplified error model and to a lesser extend worse resolution. We see from Figure (3) that actually the IP error distributions are very different. The off-line distribution looks much broader compared to the on-line one which is more peaked: only the x and y are considered in the error model and the correlations are ignored in the covariance matrix. This is reflected by the root-mean-square of the histograms which are given by $\mathrm{rms}^{\mathrm{Off}}_{\mathrm{IPerr}} = 5.8\,\mu$m and $\mathrm{rms}^{\mathrm{On}}_{\mathrm{IPerr}} = 4.0\,\mu$m. The differences in IP error determination, IP calculation and the ambiguity introduced by the different primary vertex reconstruction efficiencies represent sources of inefficiencies for the HLT selections.

The error on IP is used in the HLT to enable significance cuts which further discriminate the signal and background distributions, provided the errors are well estimated. The fact that a simplified error model is used has the important consequence of rending significance cuts less powerful in terms of signal–background rejection. Since the errors are spread over a smaller range, their impact on the separating power of significance cuts is reduced.

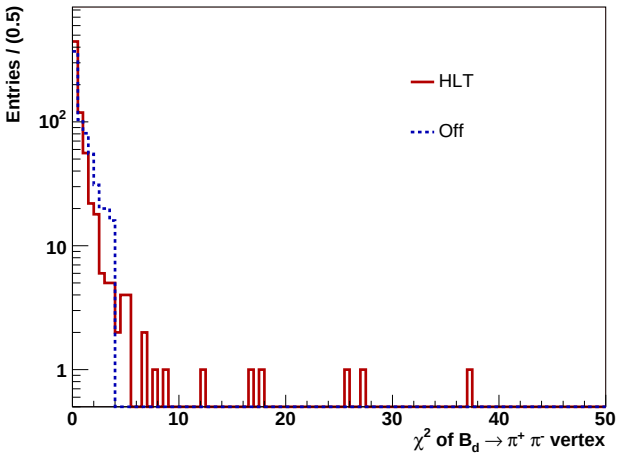

Figure 5: χ^2 of the vertex fit for $\mathrm{B_d} \to \pi^+\pi^-$ candidates. The red solid line is for the HLT and the blue dotted line for the off-line selection.

The reconstruction of decays in the HLT also requires a vertexing tool in order to create composite particles and their corresponding decay vertices when combining particles. The tool should be fast to be able to perform all combinatorics and should return an estimator of the quality, as an extra handle to reject background. The vertexing tool used in the HLT relies on the error parameterization and assumes straight trajectories [6]. In Figure (4) the distributions of the χ^2 of the fitted $\mathrm{B_d}$ vertex in $\mathrm{B_d} \to \pi^+\pi^-$ events are shown. Note that the candidates are shown on events passing the off-line selection and associated to the MC truth and where a $\chi^2 < 4.0$ cut is applied, therefore forming a good vertex. We see that some on-line vertices have large χ^2. The fraction of anomalous high χ^2 depends on the topology of the decay under study and this fraction is typically of order of a few percent, especially larger for very long-lived particles. For vertices significantly displaced with respect to the beam axis the vertex quality criterion can give bad results, yielding some inefficiencies in the HLT selections. In the above example, about 2% of the candidates are lost compared to the off-line when applying the same cut. In general the resulting χ^2

obtained from the on-line vertexing tool offers a satisfactory measure of the vertex quality to further reject the background. Depending on the decay channel *and* the off-line cut, the fraction of abnormal on-line χ^2 is a few percent.

From the study of the track parameters we know that on-line tracks have approximatively a two times worse momentum resolution. This effect will be propagated to the momentum of composite particles and to the reconstructed masses. The B_d momentum resolutions $\delta p/p$ are 0.64% and 0.32% for the on-line and off-line respectively. A positive bias of 0.13% is present in the on-line momentum resolution coming from the 3σ deviation of the mean in Figure (2) and resulting from the incorrect horizontal slope. This bias induces the observed mass shift of $5.4\,\mathrm{MeV/c^2}$ for the on-line reconstructed B_d mass, see Figure (6), when comparing the mean fitted on-line mass $m = 5284.8\,\mathrm{MeV/c^2}$ to the true value used in MC generation ($m^{\mathrm{MC}} = 5279.4\,\mathrm{MeV/c^2}$). The on-line mass resolution for $B_d \to \pi^+\pi^-$ is $32\,\mathrm{MeV/c^2}$, roughly two times the off-line resolution.

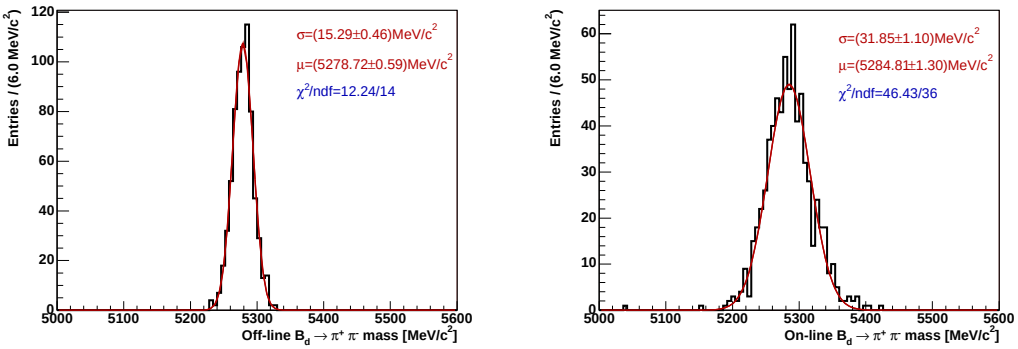

Figure 6: Invariant mass of the two pions combinations in $B_d \to \pi^+\pi^-$, for fully associated signal candidates. For the comparison purpose, simple Gaussian fits are applied to get the resolutions. Left plot: off-line mass [$\mathrm{MeV/c^2}$]. Right plot: on-line mass [$\mathrm{MeV/c^2}$].

Yet another standard cut applied in the HLT is the displacement between the primary and secondary vertices, which is equivalent to requiring the b-hadron to fly before decaying. We therefore need a good precision in the vertex positions to be able to cut on the flight distance. We already know from the study of the track parameters at the first measurement that the resolutions are somehow good enough for the HLT, but we explicitly check here for the $B_d \to \pi^+\pi^-$ decay. Table (7) shows the corresponding vertex resolutions obtained for events with exactly one reconstructed and generated primaries. For a simple on-line–off-line comparison, the resolutions are quoted as the standard deviations of single Gaussian fits. The v_x and v_y resolutions are comparable to the values obtained for the x and y track parameters and are as expected much better than the v_z resolution since the momentum is much lower in the transverse plane. The position measurement in the z-direction is critical since longitudinal to the beam and generally secondary vertices are produced under low angles in the forward direction. The primary vertex has a higher precision compared to the secondary since the multiplicity of tracks used is much larger and includes tracks with large opening angles between each other. Note that the primary vertex resolutions in Table (7) are obtained from single interaction signal events, which tend to be cleaner and have better resolutions.

The resolution of the primary vertex when measured on minimum bias and including multiple interaction events is shown in Figure (7) for the z-direction. The combined fitted resolution for v_z is about $66\,\mu$m, whereas for the transverse coordinates we get resolutions

of $\sigma(v_{x,y}) \sim 13\,\mu$m.

Table 7: Vertex resolutions in $B_d \to \pi^+\pi^-$ events with a single primary vertex on events passing the generic HLT and the off-line selection. The resolutions are the widths of single Gaussian fits.

Secondary vertex resolution	v_x [μm]	v_y [μm]	v_z [μm]
On-line	17.0	16.3	139.6
Off-line	13.1	12.5	127.9

Primary vertex resolution	v_x [μm]	v_y [μm]	v_z [μm]
On-line	9.3	9.5	48.4
Off-line	8.0	7.8	41.1

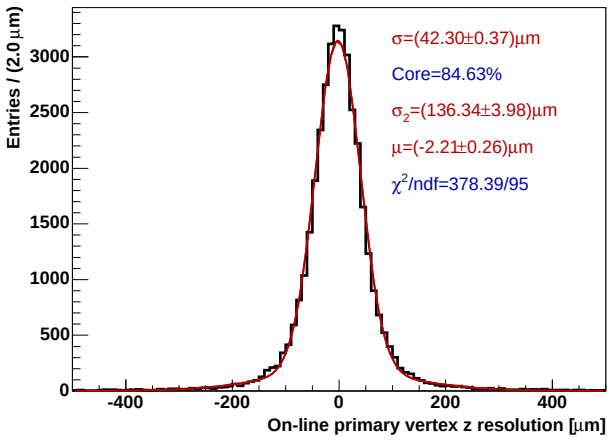

Figure 7: On-line primary vertex v_z residual [μm], on ~ 3.28 s minimum-bias events after the generic HLT. A double Gaussian with same mean is fitted to the distribution.

The primary vertex reconstruction plays a crucial role for the HLT retention rates. Missed primary vertices can fake a detached secondary vertex due to the large multiplicity of tracks in the vicinity. The HLT selections therefore strongly depend and rely on the primary vertex finding. The average number of HLT primary vertices found on minimum-bias after the generic HLT is 1.39 (39604 primaries in 28458 events). The on-line primary vertex reconstruction efficiencies[8] are given in Table (8), for reconstructible true Monte Carlo primary vertices (MCPV) matched to a reconstructed primary vertex. The following definitions are used to build the efficiency table:

- A MCPV is *reconstructible* if it has at least five tracks originating from it and reconstructed as Velo-space tracks;
- A MCPV is *matched* to a reconstructed primary vertex if more than 5% of the tracks used to form it are linked to true MC particles from the MCPV. Also to reject very

[8] Using `Phys/PVEff v1r0p1`.

badly reconstructed primaries, the matched PV should lay within 2 mm of the corresponding MCPV.

- A MCPV is said to be *isolated* (*close*) if the distance to the closest reconstructible MCPV is larger (smaller) than 10 mm.

Note that the total number of reconstructible MCPV matched to a PV ($N_{\text{ted}}^{\text{MCPV}} = 38401$) is smaller than the real number of reconstructed primaries (39604) since primaries not matched to any reconstructible MCPV are not considered. This fraction amounts to $\sim 3\%$ and generally represents fake or badly reconstructed primary vertices.

Table 8: On-line primary vertex reconstruction efficiencies $\varepsilon^{\text{MCPV}}$ on ~ 3.28 s minimum-bias events after the generic HLT. Only reconstructible MCPV are considered ($N_{\text{ible}}^{\text{MCPV}}$) and reconstructed primaries matched to a MCPV ($N_{\text{ted}}^{\text{MCPV}}$). The uncertainties are statistical.

MCPV type	$N_{\text{ted}}^{\text{MCPV}}$	$N_{\text{ible}}^{\text{MCPV}}$	$\varepsilon^{\text{MCPV}}$ [%]
Isolated	34250	35982	95.2 ± 0.1
Close	4151	5491	75.6 ± 0.6
Total	38401	41473	92.6 ± 0.2

The reconstruction efficiencies from Table (8) demonstrate the difficulties spanned by the presence of several close primary vertices. Though the primary finding efficiency for isolated MCPV is high, we should be careful since they may lead to fake b-vertices. For the close MCPV, the performance is much worse since $\sim 25\%$ of the reconstructible MCPV are lost. Note that the reconstruction algorithm requires at least six tracks to form a PV and we could expect a better efficiency when requiring less tracks. However, this has to be checked against the cost in the HLT efficiency on signal, especially for high multiplicity channels. For very close MCPV, we are at the border line for the distinction between a low multiplicity primary vertex and a high multiplicity decay. A possible improvement would be to use backward Velo tracks and large polar angles to separate primaries from secondaries. The total inefficiency for the reconstructible MCPV sets an upper limit on the HLT output rate of 642 Hz (from 8668.3 Hz after the generic HLT) that would be due missed primary vertices if we assume that this only happens in events without any b or c quark content and ignoring the non-matched primaries.

The new strategy for the primary vertex finding gives good results and enabled to considerably improve the HLT selections efficiencies by rejecting primary vertices identified as signal vertices. As it appears from Table (8) there is still room for improvement, especially for close primary vertices.

4 HLT Selections

We describe in this Section the selection procedure for the exclusive b and the inclusive D^* streams. We enter the HLT selections after a positive generic HLT decision and they are run on all events, even if the events are also triggered by the inclusive b$\rightarrow\mu$ and/or the di-muon streams. This is a demand from systematic studies (e.g. to determine whether an event was triggered because of the signal), that require the trigger to save all triggering objects. The current simplified data-flow of the HLT sequence is summarized in Figure (8). A description of the reconstruction, generic HLT and muon highways was briefly explained

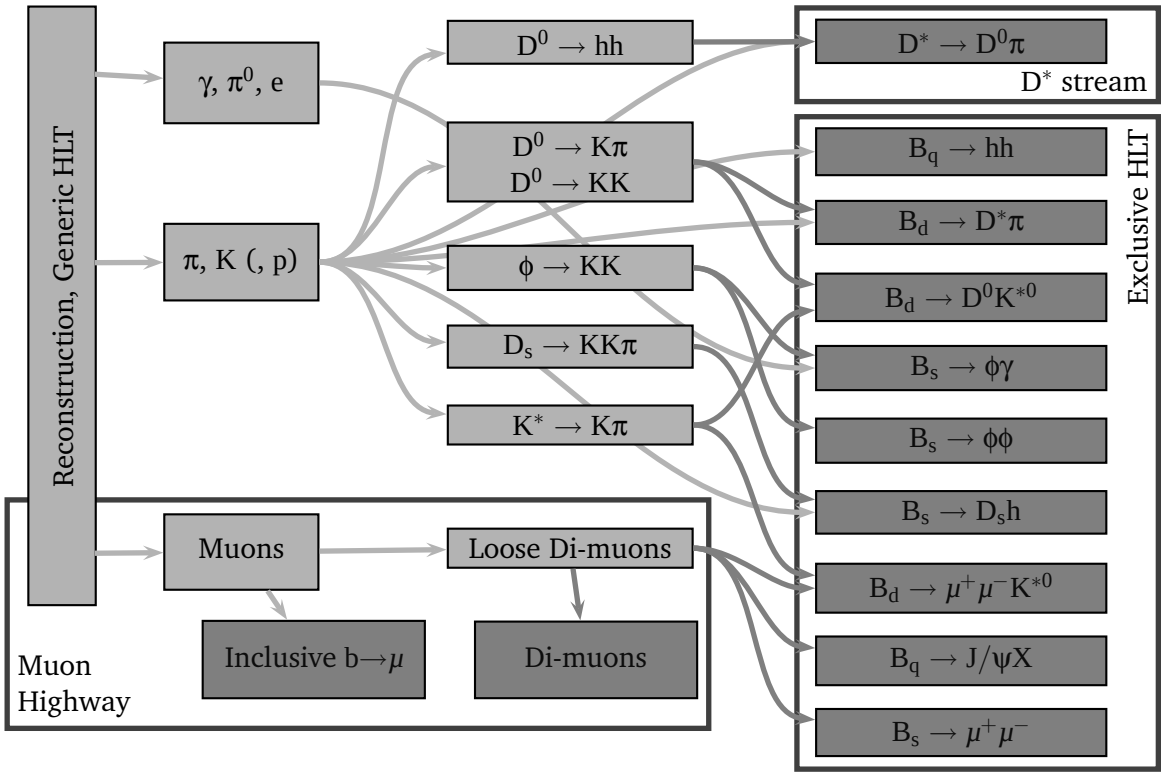

Figure 8: Simplified data flow in the HLT. Each gray box is an algorithm (or a set of algorithms).

in Section (2.1) and we present hereafter the steps leading to the exclusive b and the inclusive D^* streams:

- Preparation of the final state particles, in Sections (4.1, 4.5);
- Creation of the shared composite particles, in Section (4.2);
- HLT selections to fully reconstruct the different decay chains, in Sections (4.3, 4.4).

The first two steps are motivated by important requirements for the HLT, namely speed performance and output rate. Given the restricted available time budget to process the whole HLT, we cannot afford to perform several times identical operations unless it is absolutely necessary. Thus to avoid duplications the final states creation and the formation of composite particles are executed once per event, for each particle type. A few selection criteria are already applied at this stage, to limit the combinatorics and therefore save some additional processing time. Finally a set of common cuts is also applied to all final states made from long tracks to control the output rate.

These shared particles will then be reused and refined according to the needs of the dedicated HLT exclusive selections. This represents the part of the HLT selections whose purpose is to combine the shared composite particles and other final states with some refinement all the way to the signal b-hadron or D^* candidates. The modes used to reconstruct the different exclusive b and the D^* streams are presented in Section (4.3). Every exclusive selection is executed independently of each other on all events passing the generic HLT. The final decision for the exclusive selections b is then a logical *OR* of all the individual selection flags. Note that the D^* stream can be considered as an exclusive selection since it is *always* reconstructed with the same signature (slow pion and $D^0 \rightarrow$

hh) and with all tracks made as pions. From the decision point of view, the D^* selection nevertheless represents a separate stream with respect to the exclusive b bandwidth.

The exclusive selections are meant to mimic the off-line selections with if possible looser selection criteria. To ensure the stability and reproducibility of the cuts a limited list is applied, as defined in Section (2.3). The procedure for the determination of the selection criteria is explained in Section (4.4) and we give there all the values of the selection cuts for each of the HLT selections. The tuning of the selection criteria results from several iterations especially with the choice of preliminary cuts applied the final state particles, which is described in Section (4.5).

4.1 Final States

The HLT selections use final states particles as inputs and then combine them to form composite particles until the whole decay chain is reconstructed. The initial step is therefore to create the final states, using particle makers that assign particle identifications (PID) to the end products of the reconstruction and therefore define the particle's four-vector. We distinguish three different classes of particles used in the HLT:

Hadrons. Charged hadron particles (π, K, p) are made from Velo-TT or long tracks and reconstructed by the standard on-line pattern recognition. All tracks are made as both pion and kaon candidates. The pions should not be affected by PID requirements to ensure that the flavor-blind selections work for any particle hypothesis. The default HLT implementation does not use the RICH, as it represents an additional time cost. Note that protons are by default not produced.

Muons. Muon particles are made from muon segments and compatible with the μ hypothesis as identified by the stand-alone muon reconstruction. Muons are directly retrieved from the generic HLT sequence.

ECAL. ECAL particles are built from neutral on-line electromagnetic calorimeter clusters (photons γ, resolved π^0) or from long tracks matched to charged ECAL clusters that were not already identified as muons (electrons e, including bremsstrahlung correction). The ECAL reconstruction is only run after the generic HLT.

Once all the final state particles are created, they are distributed to different pools where a preliminary selection is applied in order to reduce combinatorics as early as possible. At this point electrons and π^0 are lost since they are not used in the current implementation of the HLT.

The choice of preliminary selection criteria for the final states is based on simple b-decays features (or c-decays to a lesser extend) and used as signatures for signal tracks:

- Some transverse momentum tracks with respect to the beam axis, as a result of the large mass difference between b-hadrons and their decay products.
- Some impact parameter tracks with respect to the interaction vertex, as a consequence of the presence of significantly displaced secondary vertices with respect to the production vertex due to the b-hadron long lifetime.

The preselection only affects hadron long tracks, and the list of cuts is given in Table (9). The choice of cuts will be justified in Section (4.5). Unless specified and for very rare cases, all the shared composite particles and HLT selections use these preselected particles as inputs.

Table 9: Preliminary selection criteria defining the different pools of particles used in the HLT selections.

π, K long tracks		
p [MeV/c]	$>$	2000.0
p_{T} [MeV/c]	$>$	300.0
IPS	$>$	2.0
Detached π (upstream, long)		
IPS	$>$	1.0
Photons γ		
p_{T} [MeV/c]	$>$	300.0

4.2 Shared Composite Particles

All the intermediate composite particles are built only once per event and using the filtered final states from the different pools as discussed in Section (4.1). This avoids the duplication of computing and the resulting composite particles are then used and refined by the different exclusive streams. Since these particles are meant to be shared among all possible exclusive selections, the selection criteria are loose enough to satisfy all use cases. The non-exhaustive existing list of shared composite particles is given by K^*, ϕ, D^0, $\mathrm{D_s}$ and J/ψ. Additional particles will be created as one includes new exclusive selections to the list of channels. The J/ψ are directly retrieved from the generic HLT without any mass, nor vertex cuts, nor any other cut at all and can thus be considered as di-muons.

The shared composite particles are reconstructed in the following modes together with their charged conjugated (*cc*) when indicated:

- $\mathrm{K}^{*0} \to \mathrm{K}^+\pi^-$ and *cc*;
- $\phi \to \mathrm{K}^+\mathrm{K}^-$;
- $\mathrm{D}^0 \to \pi^+\pi^-$, $\mathrm{K}^+\pi^-$, $\pi^+\mathrm{K}^-$, $\mathrm{K}^+\mathrm{K}^-$ (always labeled as "D^0" and never as "$\overline{\mathrm{D}}^0$");
- $\mathrm{D_s^-} \to \mathrm{K}^+\mathrm{K}^-\pi^-$ and *cc*.

All these shared particles are formed using particles made from long tracks, with the cuts given in Table (9). The list of preliminary cuts applied to these shared particles is given in Table (10) and will be justified later when discussing all the HLT selection criteria. Whenever possible, that is if all the exclusive selections using a given shared composite apply the same selection criteria, tight cuts are already applied at this stage to improve the timing performance. This is the case of the shared $\mathrm{D_s}$. If a new exclusive selection requires a softer shared particle then the criteria can just be reshuffled to the dedicated part. Besides the obvious time gain, using shared composite particles has the advantage of uniforming and standardizing the HLT selections.

The HLT mass cuts should be loose enough in order to keep sidebands and account for the worse on-line mass resolutions. The resolutions will justify the mass cuts applied when using the shared composite particles in the different exclusive streams. Running the HLT with cheated selections, that is we require all the signal tracks to be reconstructed and we apply no cuts at all, we are able to check the mass resolutions. To this end,

Table 10: Shared composite particles HLT selection criteria.

Shared K^{*0}		
δm [MeV/c^2]	$\pm$	160.0
p_T [MeV/c]	$>$	500.0
χ^2 vertex	$<$	100.0

Shared ϕ		
δm [MeV/c^2]	$\pm$	50.0
χ^2 vertex	$<$	100.0

Shared D_s		
δm [MeV/c^2]	$\pm$	45.0
p_T [MeV/c]	$>$	1750.0
χ^2 vertex	$<$	36.0
FS	$>$	4.0

Shared $D^0 \to K^+\pi^-$, π^+K^-, K^+K^-		
δm [MeV/c^2]	$\pm$	100.0
p_T [MeV/c]	$>$	500.0
χ^2 vertex	$<$	100.0

Shared $D^0 \to \pi^+\pi^-$		
Upper δm [MeV/c^2]	$+$	50.0
Lower δm [MeV/c^2]	$-$	700.0
p_T [MeV/c]	$>$	1250.0
χ^2 vertex	$<$	16.0
FS	$>$	6.0

we use off-line selected events passing the generic HLT and selected by the cheated HLT exclusive selection. Figures (9, 10, 11) show the corresponding reconstructed masses and resolutions for the shared particles. For the resonances with a width the reconstructed mass is not fitted with a Breit-Wigner but with a Gaussian just to quantity in a simple way the effects of the intrinsic width of the resonance and the detector resolution in terms of the standard deviation. The pure detector resolution component is obtained by fitting the difference between the reconstructed and the generated mass.

A comparison between on-line and off-line resolutions is given in Table (11). The resolutions are obtained from single Gaussian fits, on the same events as the previous plots. The HLT mass resolutions are roughly a factor two less good when compared to the off-line. As seen previously, this is a consequence of the worse track parameters momentum resolution.

Table 11: On-line (left box) and off-line (right box) shared composite particles standard deviation σ and pure detector mass resolution σ_{reso}. The quoted values are from single Gaussian fits.

On-line		
Shared	σ [MeV/c^2]	σ_{reso} [MeV/c^2]
K^{*0}	32.3	5.4
ϕ	3.7	2.0
D^0	12.9	12.9
D_s	10.5	10.5

Off-line		
Shared	σ [MeV/c^2]	σ_{reso} [MeV/c^2]
K^{*0}	31.7	3.0
ϕ	3.2	0.9
D^0	7	7
D_s	5	5

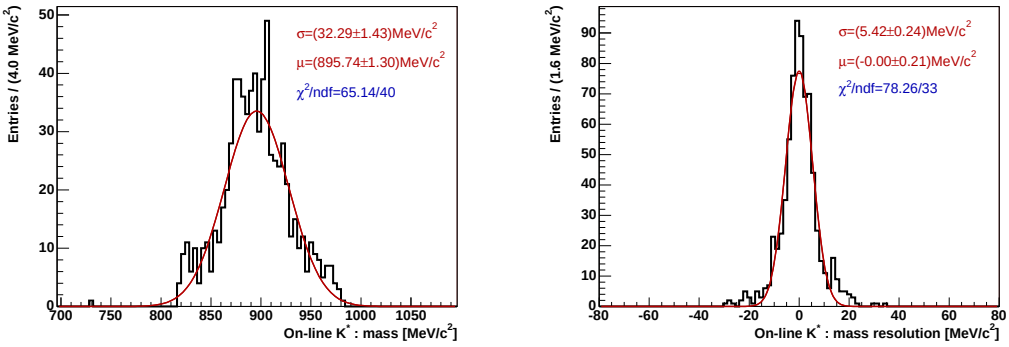

Figure 9: On-line reconstructed K^{*0} mass on $B_d \to D^0K^{*0}$ events passing the off-line selection, the generic HLT and with a cheated HLT selection. Left plot: invariant mass [MeV/c^2]. Right plot: mass residual [MeV/c^2] between reconstructed and true generated values. The curves are the results of a single Gaussian fit.

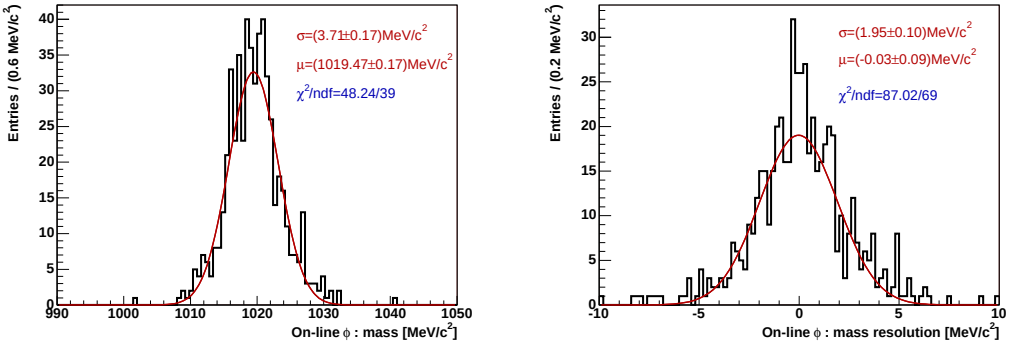

Figure 10: On-line reconstructed ϕ mass on $B_s \to \phi\gamma$ events passing the off-line selection, the generic HLT and with a cheated HLT selection. Left plot: invariant mass [MeV/c^2]. Right plot: mass residual [MeV/c^2] between reconstructed and true generated values. The curves are the results of a single Gaussian fit.

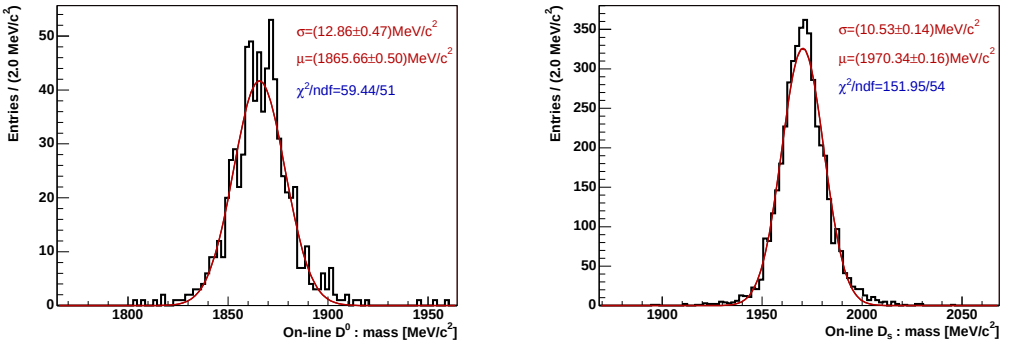

Figure 11: On-line reconstructed D^0 and D_s masses respectively on $B_d \to D^0K^{*0}$ and $B_s \to D_sh$ events passing the off-line selection, the generic HLT and with a cheated HLT selection. Left plot: D^0 mass [MeV/c^2]. Right plot: D_s mass [MeV/c^2]. The curves are the results of a single Gaussian fit.

4.3 Exclusive HLT Selections Channels

The exclusive b and the D^* streams are reconstructed in the modes described hereafter, using the different pools of particles as shown in the previous Sections and as illustrated in the data flow of Figure (8). The shared particles used to build the different channels are in general refined in the dedicated selections, though it is not explicitly mentioned here.

$B_s \to D_s h$ is always built as $B_s \to D_s \pi$ using the shared $D_s \to K^+K^-\pi$ and a bachelor π made from long tracks and with p_T and sIPS cuts. The mass window is further enlarged to accommodate $B_s \to D_s K$.

$B_q \to J/\psi X$ is reconstructed as di-muons from the generic HLT with large sidebands around the J/ψ nominal value. No extra tracks are used besides the μ such that this exclusive selection is a lifetime biased di-muon stream. As the rate of true J/ψ is much larger than the few hertz we would like to grant to this exclusive selection, tight cuts are applied to the J/ψ and the muons. As most of the di-muons will anyway be already triggered by the muon streams, we do not worry about the efficiency for this selection.

$B_q \to hh$ is always built as a PID-blind $B_d \to \pi\pi$ using opposite sign combinations with the π made from long tracks and with p_T and sIPS cuts. The mass window is widened enough to accommodate all $B_q \to hh$ modes, namely $B_d \to \pi^+\pi^-$, $B_d \to K^+\pi^-$, $B_s \to K^+K^-$, and $B_s \to \pi^+K^-$.

$B_d \to D^0K^{*0}$ is formed with the shared $D^0 \to K^+\pi^-$, π^+K^-, K^+K^-, and the shared $K^{*0} \to K^+\pi^-$, including wrong sign combinations;

$B_s \to \phi\phi$ is selected with the shared $\phi \to K^+K^-$.

$B_d \to D^*\pi$ is built using a shared $D^0 \to K^+\pi^-$, π^+K^-, K^+K^-, a slow pion made from upstream or long tracks and without any cuts for the $D^* \to D^0\,\pi$, and a fast pion made from long tracks and with p_T and sIPS cuts.

$B_d \to \mu^+\mu^-K^{*0}$ is built using di-muons from the generic HLT and the shared $K^{*0} \to K^+\pi^-$.

$B_s \to \mu^+\mu^-$ is built using di-muons from the generic HLT.

$B_s \to \phi\gamma$ is formed using shared $\phi \to K^+K^-$ and a high transverse energy E_T photon γ. Note that photons do not participate in the vertex determination and are just used to compute the resulting 4-momentum and its covariance matrix.

D^* **stream** is selected using a shared $D^0 \to \pi^+\pi^-$ and a detached slow pion made from upstream or long tracks. This stream is PID-blind and use asymmetric mass windows to in-globe the other 2-body D^0 decays.

4.4 Exclusive HLT Selections Criteria

The HLT selections have a bandwidth of ~ 200 Hz for the exclusive b and of ~ 300 Hz for the inclusive D^* stream. In order to lower the minimum-bias rate down to these specifications a series of filters is applied to all the channels. The goal is of course to have the lowest possible rate while losing as few as possible interesting events. To this end we imposed the constraints for the cuts determination listed hereafter:

- Only the selection criteria from the list given in Section (2.3) are used, with simple values to avoid any excessive and too specific tuning of the cuts.

- Unless specified, the preliminary requirements introduced in Section (4.1) are applied to all selections and the standard composite particles of Section (4.2) are used. According to the needs of the different selections, these shared particles may be further refined.

- The mass cuts though extremely useful in terms of background rejections are loose enough to include sidebands and generic PID-blind selections. The mass cuts for the intermediate particles correspond to at least three root-mean-square[9] and the b-hadron mass windows are 500 MeV/c^2 or more. Finally for PID-blind selections the mass windows must be enlarged sufficiently to accommodate the different particle hypotheses.

- The selection cuts were determined on the minimum-bias DC04-v1 sample and then blindly applied to the minimum-bias DC04-v2 sample to get an unbiased estimate of the output rates.

- The cuts are tuned such that the output rate of each of the dedicated exclusive b streams is at most of the order of 10 Hz and ignoring any correlation between each of these selections. All the available bandwidth is intentionally *not* filled thus leaving room for additional new exclusive selections. For the inclusive D^* stream, all the bandwidth is used.

The procedure to determine the selection criteria can be summarized as follows:

1. We use the signal data after off-line selection and generic HLT to study the effect of the cuts. A set of loose preselection cuts is then defined for each stream using a cheated HLT selection and applying no cuts whatsoever in order to see the effect of the on-line environment. The cheated selection ignores the PID when required, namely for the PID-blind selections such as $B_q \to hh$. The values of the preselection cuts are obvious choices when looking at the cut variable distribution of the off-line versus the on-line selections.

2. The preselection cuts are then applied when running on minimum-bias to limit the combinatorics. Based on the distributions of the cuts for minimum-bias versus the HLT preselected and associated signal candidates, the final cuts are chosen such that the output rate is lowered to the required level while cutting the tails of the signal distributions. There is no fine tuning of the selection cuts and we impose the cuts such that we get rid of signal candidates with abnormal parameters due to on-line environment or simply not in the core of the variable distributions. As the HLT selections can tolerate to have some background within the output rate, there is no optimization of the significance of the signal.

The final selection criteria and the preselection cuts used for their determination are given in the Tables (12-21).[10]

4.5 Final States Filtering

The preliminary filtering of the final states introduced in Section (4.1) and in particular the IP cuts do not affect muons, as these do not have any significant impact on combinatorics.

[9] Actually a single Gaussian is fitted to the reconstructed mass. The resulting standard deviation includes both the detector resolution and the intrinsic width if the particle is a resonance.

[10] All the distributions used in the determination of the selection criteria for each of the HLT selections can be found on the `http://lhcb-trig.web.cern.ch/lhcb-trig/HLT/Default.htm` web page.

Table 12: $B_s \to D_s h$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

D_s products			
p [MeV/c]	$>$	2000.0	(1500.0)
p_T [MeV/c]	$>$	300.0	(200.0)
IPS	$>$	2.0	(1.0)

D_s			
δm [MeV/c^2]	$\pm$	45.0	(80.0)
p_T [MeV/c]	$>$	1750.0	(1000.0)
χ^2 vertex	$<$	36.0	(100.0)
FS	$>$	4.0	$(-)$

Bachelor h			
p [MeV/c]	$>$	2000.0	(1000.0)
p_T [MeV/c]	$>$	700.0	(500.0)
IPS	$>$	3.0	(1.0)

B_s			
δm [MeV/c^2]	$\pm$	600.0	(600.0)
χ^2 vertex	$<$	9.0	(100.0)
IPS	$<$	3.0	(9.0)
FS	$>$	1.0	(1.0)
FD [mm]	$>$	1.0	$(-)$
$\cos\theta_{p,F}$	$>$	0.9998	(0.99)

Table 13: $B_q \to J/\psi X$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

J/ψ products			
p_T [MeV/c]	$>$	500.0	$(-)$
IPS	$>$	2.0	$(-)$

J/ψ			
δm [MeV/c^2]	$\pm$	120.0	$(-)$
FS	$>$	2.0	$(-)$
FD [mm]	$>$	2.5	$(-)$
IPS	$>$	3.0	$(-)$
χ^2 vertex	$<$	25.0	$(-)$

Table 14: $B_q \to hh$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

B_q products			
p [MeV/c]	$>$	3000.0	(2000.0)
p_T [MeV/c]	$>$	800.0	(600.0)
IPS	$>$	4.0	(2.0)

B_q			
δm [MeV/c^2]	$\pm$	600.0	(600.0)
p_T [MeV/c]	$>$	1000.0	(800.0)
χ^2 vertex	$<$	9.0	(50.0)
FS	$>$	8.0	(4.0)
FD [mm]	$>$	1.0	$(-)$
IPS	$<$	3.0	(9.0)
$\cos\theta_{p,F}$	$>$	0.9995	(0.999)

Table 15: $B_d \to D^0 K^{*0}$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

D^0 and K^{*0} products			
p [MeV/c]	$>$	2000.0	(1000.0)
p_T [MeV/c]	$>$	300.0	(200.0)
IPS	$>$	2.0	(1.0)

D^0			
δm [MeV/c^2]	$\pm$	50.0	(100.0)
p_T [MeV/c]	$>$	1000.0	(500.0)
χ^2 vertex	$<$	36.0	(100.0)
IPS	$>$	4.0	$(-)$

K^{*0}			
δm [MeV/c^2]	$\pm$	100.0	(160.0)
p_T [MeV/c]	$>$	1000.0	(500.0)
χ^2 vertex	$<$	36.0	(100.0)
IPS	$>$	4.0	$(-)$

B_d			
δm [MeV/c^2]	$\pm$	500.0	(500.0)
χ^2 vertex	$<$	16.0	(100.0)
FS	$>$	4.0	$(-)$
FD [mm]	$>$	1.0	$(-)$
IPS	$<$	4.0	(9.0)
$\cos\theta_{p,F}$	$>$	0.9999	(0.999)

Table 16: $B_s \to \phi\phi$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

ϕ products			
p [MeV/c]	$>$	2000.0	(1000.0)
p_T [MeV/c]	$>$	300.0	(200.0)
IPS	$>$	2.0	(1.0)

ϕ			
δm [MeV/c^2]	$\pm$	20.0	(50.0)
χ^2 vertex	$<$	49.0	(50.0)

B_s			
δm [MeV/c^2]	$\pm$	500.0	(500.0)
χ^2 vertex	$<$	49.0	(100.0)
FS	$>$	6.0	(4.0)
FD [mm]	$>$	1.0	(0.0)
IPS	$<$	6.0	(9.0)
$\cos\theta_{p,F}$	$>$	0.9995	(0.999)

Table 17: $B_d \to D^*\pi$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

D^0 products			
p [MeV/c]	$>$	2000.0	(1000.0)
p_T [MeV/c]	$>$	300.0	(200.0)
IPS	$>$	2.0	(1.0)

D^0			
δm [MeV/c^2]	$\pm$	50.0	(100.0)
p_T [MeV/c]	$>$	1000.0	(500.0)
χ^2 vertex	$<$	36.0	(100.0)

Slow π (upstream and long tracks)
No cuts

D^*			
δm [MeV/c^2]	$\pm$	50.0	(100.0)
$m(D^* - D^0)$ [MeV/c^2]	$<$	6.0	(60.0)
p_T [MeV/c]	$>$	1000.0	(500.0)
χ^2 vertex	$<$	36.0	(100.0)

Fast π			
p [MeV/c]	$>$	2000.0	(1000.0)
p_T [MeV/c]	$>$	1300.0	(1000.0)
IPS	$>$	2.0	(1.0)

B_d			
δm [MeV/c^2]	$\pm$	500.0	(500.0)
χ^2 vertex	$<$	25.0	(100.0)
FS	$>$	3.5	(2.0)
FD [mm]	$>$	1.0	(0.0)
IPS	$<$	4.0	(9.0)
$\cos\theta_{p,F}$	$>$	0.9995	(0.995)

Table 18: $B_d \to \mu^+\mu^- K^{*0}$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

K^{*0} products			
p [MeV/c]	$>$	2000.0	(1000.0)
p_T [MeV/c]	$>$	300.0	(200.0)
IPS	$>$	2.0	(1.0)

K^{*0}			
δm [MeV/c^2]	$\pm$	120.0	(160.0)
p_T [MeV/c]	$>$	850.0	(500.0)
χ^2 vertex	$<$	36.0	(100.0)

μ			
p_T [MeV/c]	$>$	900.0	(500.0)

B_d			
δm [MeV/c^2]	$\pm$	500.0	(500.0)
χ^2 vertex	$<$	100.0	(150.0)
FS	$>$	5.0	$(-)$
FD [mm]	$>$	1.0	$(-)$
IPS	$<$	5.0	(9.0)
$\cos\theta_{p,F}$	$>$	0.9995	(0.998)

Table 19: $B_s \to \mu^+\mu^-$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

μ		
p_T [MeV/c]	> 1200.0	(500.0)

B_s		
δm [MeV/c^2]	$\pm$ 500.0	(500.0)
p_T [MeV/c]	> 2900.0	(1000.0)
χ^2 vertex	$<$ 49.0	(100.0)
FS	$>$ 12.0	$(-)$
IPS	$<$ 5.0	(9.0)
$\cos\theta_{p,F}$	$>$ 0.999	$(-)$

Table 20: $B_s \to \phi\gamma$ final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

ϕ products		
p [MeV/c]	> 2000.0	(1000.0)
p_T [MeV/c]	$>$ 300.0	(200.0)
IPS	$>$ 2.0	(1.0)

ϕ		
δm [MeV/c^2]	$\pm$ 20.0	(50.0)
p_T [MeV/c]	> 1000.0	$(-)$
χ^2 vertex	$<$ 49.0	(100.0)

γ		
E_T [MeV]	> 2800.0	(2000.0)

B_s		
δm [MeV/c^2]	$\pm$ 500.0	(500.0)
FD [mm]	$>$ 1.0	(0.0)
IPS	$<$ 3.5	(9.0)
$\cos\theta_{p,F}$	$>$ 0.999	(0.995)

Table 21: D^* final HLT selection criteria. The rightmost values in parenthesis indicate the preselection used in the cuts determination.

D^0 products		
p [MeV/c]	> 2000.0	(1000.0)
p_T [MeV/c]	$>$ 300.0	(200.0)
IPS	$>$ 2.0	(1.0)

D^0		
δm [MeV/c^2]	$+$ 50.0	(100.0)
δm [MeV/c^2]	$-$ 700.0	(1300.0)
p_T [MeV/c]	> 1250.0	(500.0)
χ^2 vertex	$<$ 16.0	(100.0)
FS	$>$ 6.0	$(-)$

Slow π (upstream and long tracks)		
IPS	$>$ 1.0	$(-)$

D^*		
δm [MeV/c^2]	$+$ 50.0	(100.0)
δm [MeV/c^2]	$-$ 600.0	(1200.0)
$m(D^* - D^0)$ [MeV/c^2]	$<$ 10.0	(60.0)
p_T [MeV/c]	> 1250.0	(500.0)
χ^2 vertex	$<$ 16.0	(100.0)

We therefore only consider pions and kaons for the determination of the preselection cuts. The effect of cutting on the transverse momentum (p_{T}) and on the (unsigned) impact parameter significance with respect to *all* reconstructed primary vertices (sIPS) has been investigated in terms of:

1. Processing time on minimum-bias for all the exclusive b and D^* streams operations.
2. Output rate on minimum-bias of the exclusive b stream.
3. Efficiency of the dedicated selections for each of the core signal channels involving hadrons.

A scan for different values of the preliminary selection criteria p_{T} and sIPS applied to all π and K made from long tracks is done for each of the above points. Besides the obvious gain in timing and retention rate achieved when applying such a selection at the particle creation level, the choices for the cuts values were motivated by the fact that in the longer term many more exclusive channels, $\sim \mathcal{O}(100)$, will be considered besides the core signal channels. The final preliminary selection cut values thus represent a conservative choice in terms of output rate, and are in some cases in detriment of signal efficiencies. This is actually one place where we are forced to cut harder compared to several official off-line selections. However, the performance could be improved in a next iteration and tuning of off-line selections that presently apply very loose or no p_{T} and sIPS[11] requirements.

For the p_{T} and sIPS scan, the final selection criteria as presented in Section (4.4) are applied, except for the p_{T} and sIPS of all hadrons. For exclusive HLT selections requiring stronger requirements of the final states, such as $\mathrm{B_q} \to \mathrm{hh}$ or the bachelor in $\mathrm{B_s} \to \mathrm{D_s h}$, the preliminary selection criteria on the final states have no effect. Moreover, only long hadron tracks are affected, i.e. particles made from upstream tracks are not considered in the scan.

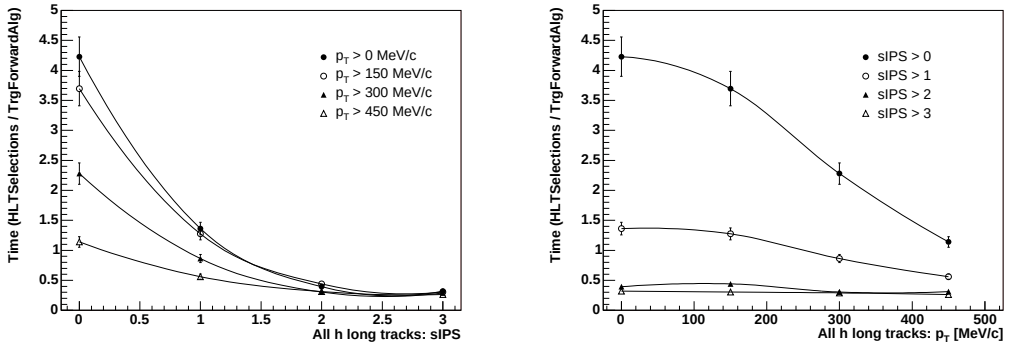

Figure 12: HLT selections (`HLTSelections`) execution time normalized to the forward tracking algorithm (`TrgForwardAlg`) CPU time on minimum-bias events. Left plot: time as a function of the sIPS cut on all long hadron tracks. Right plot: time as a function of the p_{T} [MeV/c] cut on all long hadron tracks. The errors $\sim 7.8\%$ represent the spread of the normalized time distribution weighted by the mean in a series of identical jobs.

The impact on the processing time of the different p_{T} and sIPS cuts in minimum-bias events is shown in Figure (12). We consider the HLT selections CPU time normalized to a

[11] Except when the off-line analyses can afford, in terms of background, a lifetime unbiased selection and when they are triggered by inclusive (lifetime unbiased) streams.

reference algorithm chosen to be the forward tracking since it is the most time consuming part of the HLT and to avoid side effects due to the machine load. In the HLT selections contributions are counted the times for the creation and refinement of the final state and shared composite particles as well as all the exclusive b selections and the D^* stream. The normalized time is not representative of the relative times between the HLT selections and the forward tracking algorithm since the efficiency of the generic HLT ($\sim 21.7\%$) is not taken into account in the forward tracking time and moreover we consider the time to make the forward tracking on all tracks. The relevant observable is therefore the relative normalized times between different p_{T} and sIPS configurations. A detailed comparison of the timing of all the HLT parts will be given in Section (5.2.2).

We can see in Figure (12) how the timing drops with increasing sIPS cuts and converges at sIPS $= 2-3$ for the different p_{T} cuts. At this point cutting harder has almost no effect. The impact parameter thus represents a powerful handle to reduce the combinatorics that directly act on the timing performance. For instance applying no p_{T} cut and sIPS > 2 reduces the timing by a factor ~ 11 compared to when no cuts are imposed. The effect of cutting on the p_{T} is less drastic but still significant. Applying a selection of $p_{\mathrm{T}} > 300\,\mathrm{MeV/c}$ without sIPS cut improves the speed by a factor ~ 2 relatively to the no requirements case. Applying very strong p_{T} cuts is not sufficient and cuts in sIPS are required to achieve an acceptable timing performance.

The most time consuming part in the HLT selections is the creation of shared composite particles. As an example, if we apply a mellow selection on the final states with $p_{\mathrm{T}} > 150\,\mathrm{MeV/c}$ and sIPS > 1, the time to form all the shared composite particles represents $\sim 68\%$ of the total time taken by the HLT selections. From this fraction the largest contribution, $\sim 65\%$, comes from $D_s \to K^+K^-\pi$ since it is the highest multiplicity intermediate decay considered in the HLT.

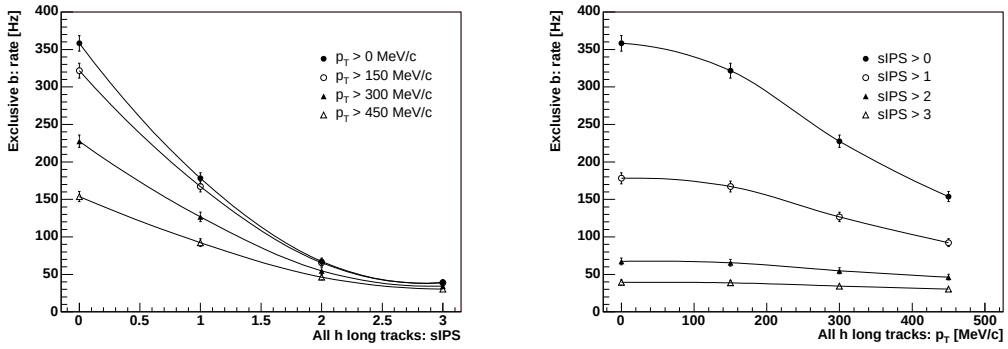

Figure 13: Overall HLT exclusive b selections retention rate on minimum-bias events. Left plot: rate as a function of the sIPS cut on all long hadron tracks. Right plot: rate as a function of the p_{T} [MeV/c] cut on all long hadron tracks. The uncertainties are statistical.

The HLT exclusive b output rates versus the different p_{T} and sIPS cuts are shown in Figure (13). The different shapes are similar to what is obtained for the timing, which is a sign that we select a large fraction of light quarks combinatorial background. The fall-off of the rate versus sIPS is however slower and the rate curves become flat for sIPS > 3, independently of the p_{T} cuts. This limit corresponds to the components of events with b or c content, and b-candidates identified as primaries. As a result, these components are difficult to suppress without significant losses in signal efficiencies.

We see from Figure (13) that without any sIPS cut the output rate is close to the available bandwidth of 200 Hz for all the exclusive b HLT selections, even when applying strong p_T cuts. In order not to fill all the bandwidth with only the few core signal channels, sIPS cuts are again required. Moreover, it is more desirable to cut here rather than on the mass of composite particles since we would like to have large enough sidebands for background studies. Let us consider a few example of cuts. Without any requirements on the final states, the output rate is of ~ 358 Hz. As expected, the exclusive $B_s \to D_s h$ selects a large part of the bandwidth when applying no requirements, $\sim 41\%$, including the overlaps with the other exclusive selections. As we will see later this is not dramatic since the efficiency on signal when applying sIPS and p_T does not suffer much for this channel. A different noisy selection when not applying any preliminary cuts is $B_d \to D^* \pi$, which retains $\sim 24\%$ of the ~ 358 Hz. In this case the effect of the requirements on the long tracks will be much more important. As a last example if we require $p_T > 300\,\mathrm{MeV/c}$ alone, the rate is lowered down to ~ 228 Hz, which is still above specifications. Then if in addition we cut on sIPS > 2, the rate is brought ~ 55 Hz.

We have demonstrated the benefits of the sIPS and p_T requirements for timing and output rate on minimum-bias events and we now consider their impact against the different signal core channels involving hadrons, see Figures (14, 15, 16, 17, 18, 19). These plots show the selection efficiencies as a function of the p_T and sIPS cuts on all the hadrons made from long tracks. The efficiencies are computed on off-line selected and generic HLT triggered signal events and only the events selected by the dedicated HLT exclusive selection of the channel under study are considered. As for the minimum-bias, all the final HLT selection criteria are applied, except for the p_T and sIPS of all long hadrons. The efficiencies include both the inefficiencies due to the selection cuts and to the track finding. We could have combinatorial background in the signal selection and select events where not all the signal tracks are on-line reconstructed. The purity of an exclusive selection does not matter but it should be efficient on signal and within the assigned bandwidth.

As we can see from the efficiency curves, the behavior when applying p_T and sIPS criteria results in general in large losses in efficiency for requirements above $150\,\mathrm{MeV/c}$ in p_T and 1 in sIPS. When applied separately and up to these values, the efficiencies remains almost flat with respect to p_T and with a small loss of at most $\sim 1\%$ for the sIPS requirement. However, stronger cuts are needed to improve the speed by reducing the combinatorics. Using Figure (12), the choice of sIPS > 2 is the best choice as far as timing is concerned. The effect of an additional p_T cut on the combinatorics is negligible, but it can further slightly lower the retention rate as it can be seen in Figure (13). Having an extra handle to directly control the output rate and the processing time at an early stage of the HLT makes it more robust and prepared to unexpected situations. The on-line errors may not be available from the start of data taking since we need to calibrate them on data and moreover the errors might be wrongly estimated, an additional p_T requirement is therefore adequate to safely control timing and rate, given the good momentum resolution of long tracks below 1%. Finally the list of exclusive HLT channels is meant to evolve and accommodate many more selections, including additional and new shared composite particles (e.g. $\eta_c \to 4h$) rending the p_T selection criterion purposeful. Based on these observations and considering Figure (14) as reference since $B_s \to D_s h$ is the most critical selection, the choice of preliminary selection criteria are given in Table (9) where additional soft momentum cuts are applied to the long tracks and in which the loose requirements for the other classes particles are also listed.

The performance of selection criteria of Table (9) on the signal efficiencies strongly depends on the channel and on whether or not this kind of cuts are applied off-line. If we first consider the $p_T > 300\,\mathrm{MeV/c}$ alone we see that there is a no loss for $B_s \to \phi\gamma$ or a tiny

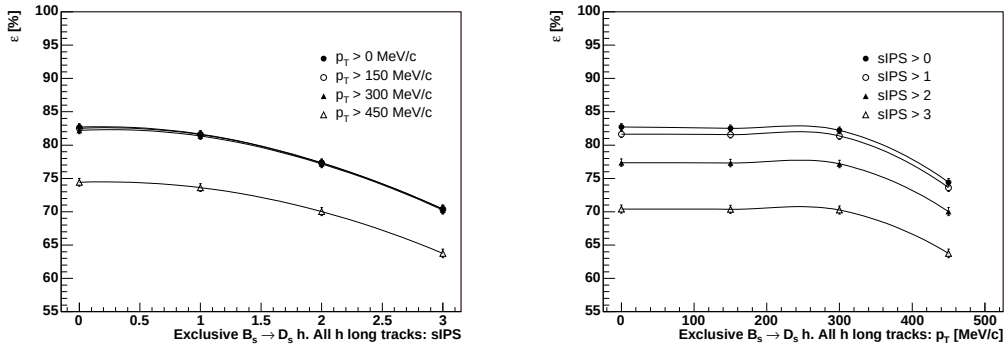

Figure 14: Dedicated exclusive $B_s \to D_s h$ efficiency ε after the generic HLT on signal off-line selected events and applying all HLT cuts except for the p_T and sIPS on all long hadrons tracks. Left plot: efficiency versus the sIPS cut. Right plot: efficiency versus the p_T cut [MeV/c]. The uncertainties are statistical.

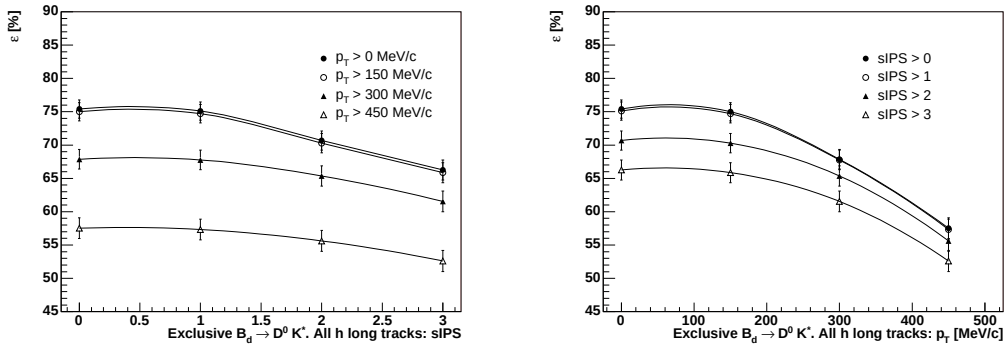

Figure 15: Dedicated exclusive $B_d \to D^0 K^{*0}$ efficiency ε after the generic HLT on signal off-line selected events and applying all HLT cuts except for the p_T and sIPS on all long hadrons tracks. Left plot: efficiency versus the sIPS cut. Right plot: efficiency versus the p_T cut [MeV/c]. The uncertainties are statistical.

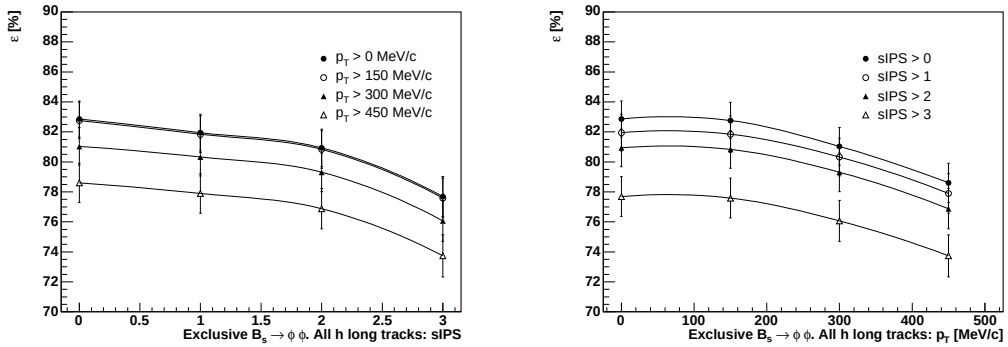

Figure 16: Dedicated exclusive $B_s \to \phi\phi$ efficiency ε after generic HLT on signal off-line selected events and applying all HLT cuts except for the p_T and sIPS on all long hadrons tracks. Left plot: efficiency versus the sIPS cut. Right plot: efficiency versus the p_T cut [MeV/c]. The uncertainties are statistical.

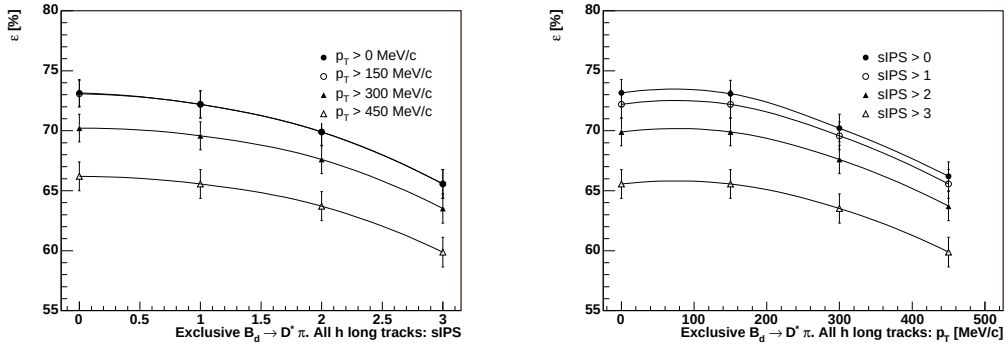

Figure 17: Dedicated exclusive $B_d \to D^*\pi$ efficiency ε after the generic HLT on signal off-line selected events and applying all HLT cuts except for the p_T and sIPS on all long hadrons tracks. Left plot: efficiency versus the sIPS cut. Right plot: efficiency versus the p_T cut [MeV/c]. The uncertainties are statistical.

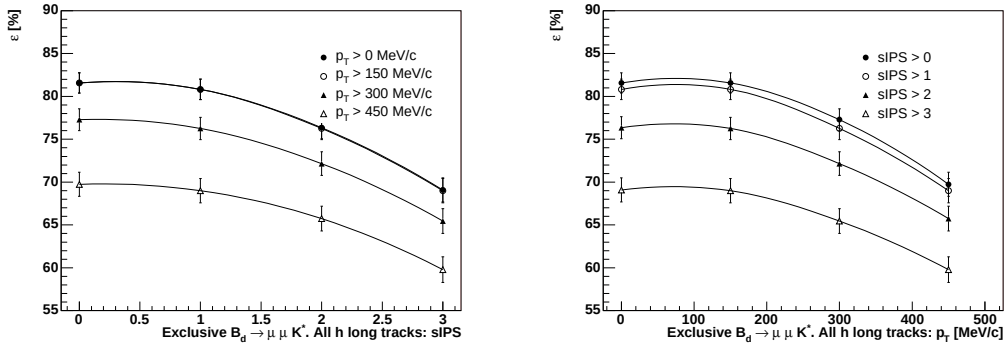

Figure 18: Dedicated exclusive $B_d \to \mu^+\mu^-K^*$ efficiency ε after the generic HLT on signal off-line selected events and applying all HLT cuts except for the p_T and sIPS on all long hadrons tracks. Left plot: efficiency versus the sIPS cut. Right plot: efficiency versus the p_T cut [MeV/c]. The uncertainties are statistical.

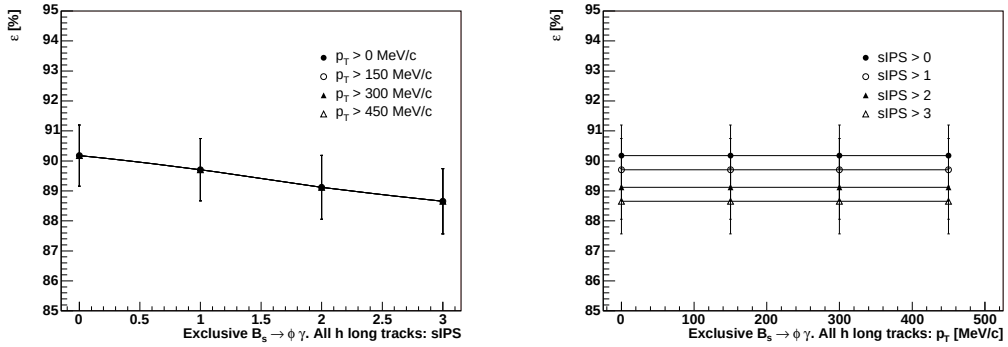

Figure 19: Dedicated exclusive $B_s \to \phi\gamma$ efficiency ε after the generic HLT on signal off-line selected events and applying all HLT cuts except for the p_T and sIPS on all long hadrons tracks. Left plot: efficiency versus the sIPS cut. Right plot: efficiency versus the p_T cut [MeV/c]. The uncertainties are statistical.

effect for $B_s \to D_s h$ which is due to on-line momentum resolution effects.[12] Cutting on p_T can be really inefficient, as for $B_d \to D^0 K^{*0}$ where we lose an additional $\sim 7.5\%$. Looking at sIPS solely, the inefficiencies are much larger, even for off-line selections applying sIPS cuts. For instance for $B_s \to \phi\phi$ we lose $\sim 1.9\%$, and for $B_s \to D_s h$ over 5% compared to the no preliminary cuts configuration.

The loss in exclusive dedicated efficiency for the HLT channels involving hadrons when applying all HLT selection criteria except for the p_T and sIPS on all long tracks hadrons is given in Table (22). The efficiencies are with respect to off-line selected events and after the generic HLT. The large loss for some channels is due to the off-line selection choosing regions of phase-space not accessible in the trigger, typically low momentum tracks.

Table 22: Exclusive HLT efficiency loss when adding to the HLT cuts the $p_T > 300\,\mathrm{MeV/c}$ and sIPS > 2 preliminary selection criteria. The efficiencies are from the dedicated exlclusive selection only and on off-line selected events passing the generic HLT.

Channel	Loss [%]
$B_s \to D_s h$	5.5 ± 0.3
$B_d \to D^0 K^{*0}$	10.0 ± 1.0
$B_s \to \phi\phi$	3.6 ± 0.6
$B_d \to D^* \pi$	5.5 ± 0.6
$B_d \to \mu^+\mu^- K^{*0}$	9.4 ± 0.9
$B_s \to \phi\gamma$	1.1 ± 0.4

The losses are very different among the channels. The various sources of inefficiencies will be presented when discussing the overall effect of the HLT selections in Section (5). Note that the different losses include the tracking inefficiencies. Finally, it is worth mentioning that the p_T knob is probably the safest parameter to tune to control the output rate, given the good momentum resolution and the fact that the on-line errors are not perfectly estimated.

5 HLT Selections Performance

The exclusive selections should reproduce as much as possible the off-line selections, with softer requirements whenever doable. The choice of standard and limited cuts to ensure the stability of the HLT selections, as described in Section (4), will inevitably lead to some inefficiencies as the off-line selections use very different type of cuts. Moreover as we have seen in Section (4) the preliminary filtering of the final states required for timing and output rate is an important source of inefficiency.

The HLT efficiencies for all the core channels are shown in Section (5.1) and the different sources of inefficiency are further described for each HLT selection. In Section (5.2) we present the results obtained on minimum-bias, such as retention rates, the quark contents and the timing performance.

[12] The off-line selection applies the same p_T cuts.

5.1 HLT Selections Signal Performance

The selection efficiencies are determined on off-line selected events by one or more off-line selections relevant to the HLT selection under study and we always apply the generic HLT, as explained in Section (A.2). The signal data samples used are listed in Section (2.4). We give hereafter for each of the HLT selections and the core channels the corresponding HLT efficiencies and we consider the dedicated specific efficiency and all the other HLT streams. Note that the quoted HLT efficiencies in-globe the on-line track finding efficiencies. The mass resolutions for the particles not considered so far are shown as well as the effect of the flavor blind selections on the mass spectra. For the sake of clarity of the on-line–off-line comparison, single Gaussians are fitted to the mass distributions, ignoring several contributions (upstream and long tracks) or peculiar shapes (e.g. Crystal-Ball for radiative energy losses).

Exclusive $B_s \to D_s h$

As we always reconstruct $B_s \to D_s h$ with $h = \pi$ the performance is checked on both $B_s \to D_s \pi$ and $B_s \to D_s K$ events with samples of approximatively the same size not to have any flavor bias. As given in Table (12) the mass window compared to other channels is enlarged to $\pm 600\,\mathrm{MeV/c^2}$ around the nominal B_s mass to account for the mass shift induced by the wrong PID hypothesis. This effect can be seen in Figure (20) where the mass distribution on $B_s \to D_s K$ is moved to lower values due to the underestimation of the bachelor mass with the apparition of a long negative tail. For $B_s \to D_s \pi$ data the on-line mass resolution is $\sim 29.4\,\mathrm{MeV/c^2}$ to be compared to the off-line combined $B_s \to D_s h$ resolution $\sim 13.2\,\mathrm{MeV/c^2}$. Given the π-K mass difference and this resolution, the mass window is appropriate.

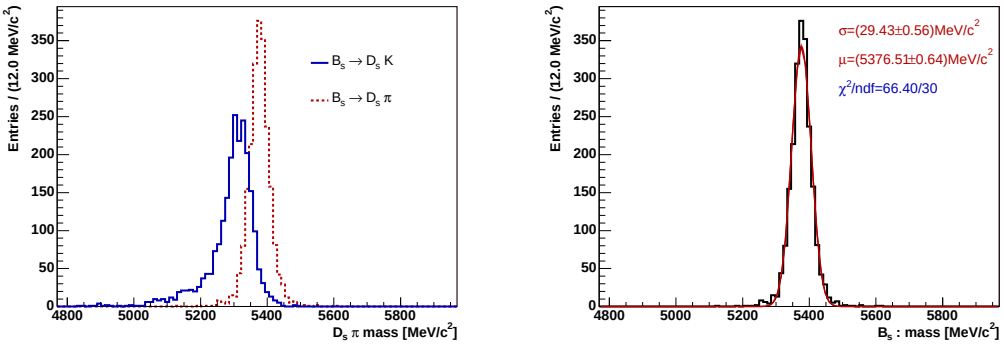

Figure 20: The reconstructed mass of the selected associated B_s signal candidates after the dedicated $B_s \to D_s h$ exclusive HLT selection and on off-line selected events passing the generic HLT. Left plot: reconstructed mass [$\mathrm{MeV/c^2}$] on $B_s \to D_s \pi$ (dashed red line) and on $B_s \to D_s K$ events (solid blue line). The histograms are not normalized to the relative branching fractions. Right plot: on-line reconstructed mass [$\mathrm{MeV/c^2}$] on $B_s \to D_s \pi$ events with a Gaussian fit.

The efficiency of the generic HLT for $B_s \to D_s h$ is $\varepsilon_{\mathrm{HLTGen/L1}} \sim 93.3\%$. This loss with respect to the Level-1 can be mainly explained by the better momentum determination therefore discarding Level-1 p_T mistakes. A fraction of this loss is also due to the 'off-diagonal' part of the trigger since the version of the generic HLT used in this note is only based on the momentum and the muon streams. For instance $\sim 2\%$ of the Level-1

triggered events for $B_s \to D_s h$ were not triggered neither by the generic nor by any of the muon lines. This fraction could be recuperated in the new trigger scheme with the introduction of ECAL lines in the generic HLT.

The efficiency of the cheated selection after the generic HLT is $\varepsilon_{HLT}^{cheated} \sim 86\%$ giving a measure of the efficiency of having all final states on-line reconstructed, which is a bit smaller than the value given in Section (3.1). This is due to the fact that for the computation of $\varepsilon_{HLT}^{cheated}$ we do not require all the final states to be off-line reconstructed. Moreover the off-line selection always select a few percent of non-associated candidates on signal, i.e. its purity is not 100%.

The number of events selected by the dedicated $B_s \to D_s h$ HLT selection and the corresponding specific efficiency on N_{off}^{evt} signal $B_s \to D_s h$ events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	N_{off}^{evt}	$N_{HLT}^{specific}$	$\varepsilon_{HLT}^{specific}$ [%]	N_{HLT}^{exb}	$N_{HLT}^{b\to\mu}$	$N_{HLT}^{\mu\mu}$	$N_{HLT}^{D^*}$	N_{HLT}^{tot}
$B_s \to D_s h$	5701	4400	77.2 ± 0.6	4407	206	26	196	4496

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	ε_{HLT}^{exb} [%]	$\varepsilon_{HLT}^{b\to\mu}$ [%]	$\varepsilon_{HLT}^{\mu\mu}$ [%]	$\varepsilon_{HLT}^{D^*}$ [%]	ε_{HLT}^{tot} [%]
$B_s \to D_s h$	77.3 ± 0.6	3.6 ± 0.3	0.5 ± 0.1	3.4 ± 0.2	78.9 ± 0.5

We notice that the other exclusive selections do not contribute significantly to the overall ε_{HLT}^{exb} efficiency. This is not surprising as $B_s \to D_s h$ is the core channel with a displaced tertiary vertex with 3 tracks. There is nevertheless a small correlation with the exclusive $B_d \to D^0 K^{*0}$ stream. The later selects $\sim 3.6\%$ of the events triggered by the exclusive $B_s \to D_s h$, partly due to the non-negligible D^0 lifetime.

As expected, $B_s \to D_s h$ is primarily triggered by its dedicated exclusive selection with little help from the other streams. Given the loss due to the preliminary selection of the final state, see Table (9), and the tracking inefficiencies we see that a small fraction of the events is still lost due to reconstruction problems and to difference with respect to the off-line selection.

Exclusive $B_q \to J/\psi X$

For this channel we do not actually look for the X but directly rely on the presence of a di-muon pair, as provided by the generic HLT. The efficiency is expected to be far from 100% since we are limited by the rate of true di-muons and the available bandwidth. Strict cuts are thus applied keeping however very large sidebands with a mass cut of $\pm 120\,MeV/c^2$ around the J/ψ central mass. The reconstructed masses for both the on-line and off-line J/ψ are shown in Figure (21). The corresponding on-line mass resolution is $\sim 19.2\,MeV/c^2$ whereas the off-line is $\sim 10.8\,MeV/c^2$.

The efficiency of the generic HLT for $B_s \to J/\psi\phi$ is $\varepsilon_{HLTGen/L1} \sim 95.7\%$, higher that for purely hadronic channels thanks to the muon streams. The cheated selection efficiency is lower than for a 2-prong hadronic decay, $\varepsilon_{HLT}^{cheated} \sim 86.8\%$.[13] The number of events selected by the dedicated $B_q \to J/\psi X$ HLT selection and the corresponding specific efficiency on N_{off}^{evt} signal $B_s \to J/\psi\phi$ events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

[13] Note that, as explained in Appendix (A.2), we use some arbitrary association for on-line muons.

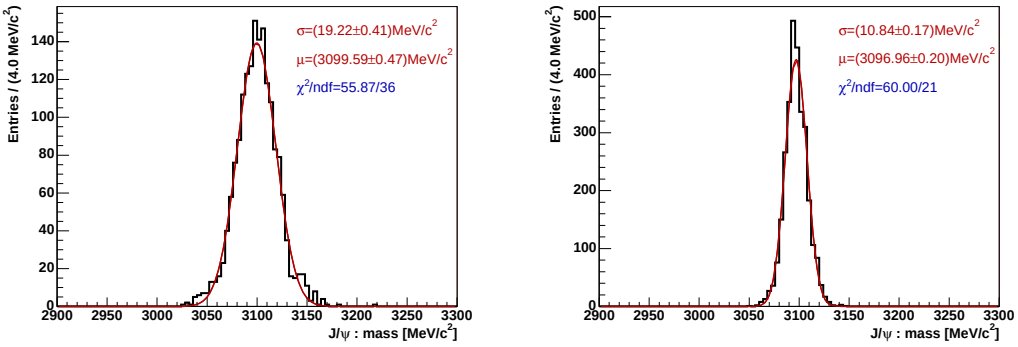

Figure 21: The reconstructed mass of the selected associated J/ψ signal candidates after the dedicated exclusive HLT selection and on off-line selected $B_s \to J/\psi\phi$ events passing the generic HLT. Right plot: on-line reconstructed mass [MeV/c^2]. Left plot: off-line reconstructed mass [MeV/c^2]. The masses are fitted with a single Gaussian.

Channel	N_{off}^{evt}	$N_{HLT}^{specific}$	$\varepsilon_{HLT}^{specific}$ [%]	N_{HLT}^{exb}	$N_{HLT}^{b\to\mu}$	$N_{HLT}^{\mu\mu}$	$N_{HLT}^{D^*}$	N_{HLT}^{tot}
$B_q \to J/\psi X$	3008	1853	61.6 ± 0.9	2197	1389	2449	84	2796

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	ε_{HLT}^{exb} [%]	$\varepsilon_{HLT}^{b\to\mu}$ [%]	$\varepsilon_{HLT}^{\mu\mu}$ [%]	$\varepsilon_{HLT}^{D^*}$ [%]	ε_{HLT}^{tot} [%]
$B_q \to J/\psi X$	73.0 ± 0.8	46.2 ± 0.9	81.4 ± 0.7	2.8 ± 0.3	93.0 ± 0.5

We see that the other exclusive selections contribute significantly to the $B_q \to J/\psi X$ efficiency, essentially due to $B_d \to \mu^+\mu^- K^{*0}$. Note that the di-muon efficiency is close to the cheated selection efficiency showing that we trigger on most of the true di-muons. $B_q \to J/\psi X$ is much less affected by the tracking issues since it can benefit from the single muon line, that is not necessarily requiring the two muons to be reconstructed (and identified). The total HLT efficiency is therefore better compared to other large multiplicity channels.

Exclusive $B_q \to hh$

The exclusive $B_q \to hh$ is meant to select all $B_d \to \pi^+\pi^-$, $B_d \to K^+\pi^-$, $B_s \to K^+K^-$, and $B_s \to \pi^+K^-$ in one by only looking at $h = \pi$ and always reconstructing a B_d ($q = d$). We must therefore be careful in the choice of the mass window since we must consider the effect of the π-K *and* the B_d-B_s mass differences. The choice of the mass window is $\pm 600\,MeV/c^2$ around the nominal B_d mass based on the plot of Figure (22). We can see how is the reconstructed mass affected when using the false flavor hypothesis for each type of $B_q \to hh$ data. For $B_d \to \pi^+\pi^-$ there is no effect and the mass resolution is $\sim 33.4\,MeV/c^2$, compared to the off-line of $\sim 16.5 MeV/c^2$. The lowest invariant masses are for $B_d \to K^+\pi^-$ as we underestimate the kaon mass. The largest invariant masses correspond to the $B_s \to \pi^+K^-$ when we underestimate the kaon mass compensated by the B_d-B_s mass difference, whereas for $B_s \to K^+K^-$ the two misidentifications and the wrong B_q flavor results in a value close to the B_d mass.

The efficiency of the generic HLT for $B_q \to hh$ is $\varepsilon_{HLTGen/L1} \sim 94.7\%$. The cheated selection efficiency is high $\varepsilon_{HLT}^{cheated} \sim 96.5\%$, as we have only two final states. The number

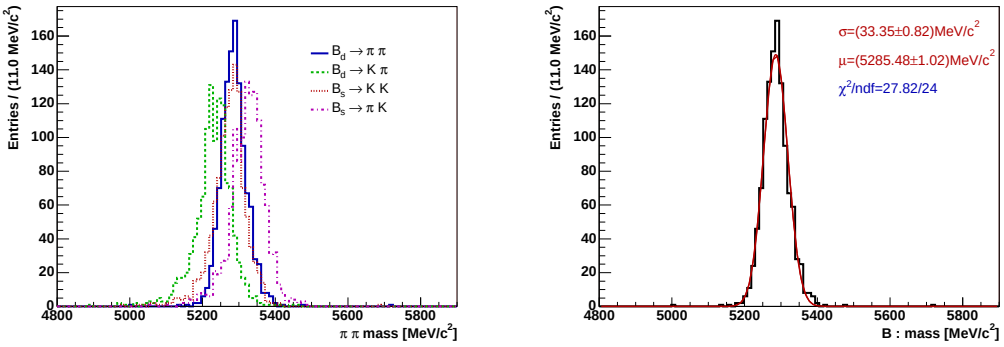

Figure 22: The reconstructed mass of the selected associated B_q signal candidates after the dedicated $B_q \to hh$ exclusive HLT selection and on off-line selected events passing the generic HLT. Left plot: contributions to the reconstructed mass [MeV/c^2] on the $B_d \to \pi^+\pi^-$ (blue solid line), $B_d \to K^+\pi^-$ (green dotted line), $B_s \to K^+K^-$ (dashed red line), and $B_s \to \pi^+K^-$ (dashed-dotted magenta line) decay events. The histograms are not normalized to the relative branching fractions. Right plot: on-line reconstructed mass [MeV/c^2] on $B_d \to \pi^+\pi^-$ events with a Gaussian fit.

of events selected by the dedicated $B_q \to hh$ HLT selection and the corresponding specific efficiency on N^{evt}_{off} signal $B_d \to \pi^+\pi^-$, $B_d \to K^+\pi^-$, $B_s \to K^+K^-$, and $B_s \to \pi^+K^-$ events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	N^{evt}_{off}	$N^{specific}_{HLT}$	$\varepsilon^{specific}_{HLT}$ [%]	N^{exb}_{HLT}	$N^{b\to\mu}_{HLT}$	$N^{\mu\mu}_{HLT}$	$N^{D^*}_{HLT}$	N^{tot}_{HLT}
$B_q \to hh$	4957	4637	93.5 ± 0.4	4638	134	11	59	4652

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	ε^{exb}_{HLT} [%]	$\varepsilon^{b\to\mu}_{HLT}$ [%]	$\varepsilon^{\mu\mu}_{HLT}$ [%]	$\varepsilon^{D^*}_{HLT}$ [%]	ε^{tot}_{HLT} [%]
$B_q \to hh$	93.6 ± 0.4	2.7 ± 0.2	0.2 ± 0.1	1.2 ± 0.2	93.9 ± 0.3

For $B_q \to hh$ practically all the events are triggered by the exclusive selection. The HLT efficiency is high, as these low multiplicity channels are much less affected by the tracking inefficiencies and as the off-line requires high IP and p_T tracks. Only a few percent are lost compared to the cheated selection, meaning that the HLT selection is quite optimal.

Exclusive $B_d \to D^0K^{*0}$

The performance for $B_d \to D^0K^{*0}$ is checked against off-line selected $B_d \to D^0$ ($\to K^+\pi^-$) K^{*0} ($\to K^+\pi^-$) events. The on-line B_d mass resolution is $\sim 24.8\,MeV/c^2$ and the off-line one is $\sim 11.3\,MeV/c^2$.

The efficiency of the generic HLT for $B_d \to D^0K^{*0}$ is $\varepsilon_{HLTGen/L1} \sim 95.0\%$ and the cheated selection efficiency is $\varepsilon^{cheated}_{HLT} \sim 83.6\%$, significantly worse than $B_s \to D_sh$. The number of events selected by the dedicated $B_d \to D^0K^{*0}$ HLT selection and the corresponding specific efficiency on N^{evt}_{off} signal events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	N^{evt}_{off}	$N^{specific}_{HLT}$	$\varepsilon^{specific}_{HLT}$ [%]	N^{exb}_{HLT}	$N^{b\to\mu}_{HLT}$	$N^{\mu\mu}_{HLT}$	$N^{D^*}_{HLT}$	N^{tot}_{HLT}
$B_d \to D^0 K^{*0}$	996	651	65.4 ± 1.5	653	21	0	50	675

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	ε^{exb}_{HLT} [%]	$\varepsilon^{b\to\mu}_{HLT}$ [%]	$\varepsilon^{\mu\mu}_{HLT}$ [%]	$\varepsilon^{D^*}_{HLT}$ [%]	ε^{tot}_{HLT} [%]
$B_d \to D^0 K^{*0}$	65.6 ± 1.5	2.1 ± 0.5	0.0 ± 0.0	5.0 ± 0.7	67.8 ± 1.5

This channel has the worse exclusive selection among all the core channels. The difference compared to $B_s \to D_s h$ is more than 10% and it is mostly due to the loss caused by the preliminary selection on the final states, see Table (9). This is closely related to the choice of final off-line selection cuts making use of low momentum tracks. As we shall see later, even for a 6-prong the efficiency is significantly better. We therefore advocate a revision of the off-line selection cuts in order to have a better alignment with the HLT. We can note that that the D^* stream helps here since there is a D^0 in the exclusive selection.

Exclusive $B_s \to \phi\phi$

The mass resolutions in the $B_s \to \phi\phi$ selection are $\sim 21.1\,\mathrm{MeV/c^2}$ and $\sim 11.5\,\mathrm{MeV/c^2}$ respectively for the on-line and off-line cases.

The efficiency of the generic HLT for $B_s \to \phi\phi$ is $\varepsilon_{HLTGen/L1} \sim 89.4\%$, which is the lowest of all the purely hadronic core signal channels. This is possibly due to the complete absence of momentum cuts in the off-line selection, resulting in smaller efficiency of the generic p_T line. The cheated selection efficiency is $\varepsilon^{cheated}_{HLT} \sim 83.1\%$, smaller than the tracking efficiency quoted previously due to the purity of the off-line events, of the order of $\sim 95.6\%$. Moreover the ambiguity induced by the small opening angle of the two kaons leads to association issues. The number of events selected by the dedicated $B_s \to \phi\phi$ HLT selection and the corresponding specific efficiency on N^{evt}_{off} signal events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	N^{evt}_{off}	$N^{specific}_{HLT}$	$\varepsilon^{specific}_{HLT}$ [%]	N^{exb}_{HLT}	$N^{b\to\mu}_{HLT}$	$N^{\mu\mu}_{HLT}$	$N^{D^*}_{HLT}$	N^{tot}_{HLT}
$B_s \to \phi\phi$	986	782	79.3 ± 1.3	787	37	3	26	794

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	ε^{exb}_{HLT} [%]	$\varepsilon^{b\to\mu}_{HLT}$ [%]	$\varepsilon^{\mu\mu}_{HLT}$ [%]	$\varepsilon^{D^*}_{HLT}$ [%]	ε^{tot}_{HLT} [%]
$B_s \to \phi\phi$	79.8 ± 1.3	3.8 ± 0.6	0.3 ± 0.2	2.6 ± 0.5	80.5 ± 1.3

The exclusive HLT efficiency for this channel can hardly be improved by a smarter choice of selection cuts, since we can see that only a few percent of the events are lost because of the selection cuts when compared to the cheated selection efficiency. The main loss is due to the p_T cuts on the ϕ daughters.

Exclusive $B_d \to D^*\pi$

The selection of $B_d \to D^*\pi$ makes use of both upstream and long tracks for the slow pion from the D^*. As upstream tracks get their rough momentum estimate from the Velo-TT matching, the resulting momentum resolution is at least an order of magnitude worse compared to what is achieved for long tracks. The contribution of Velo-TT tracks translates to non-Gaussian tails in the invariant D^* mass distribution shown in Figure (23). The resulting mass including the two type of tracks is $\sim 12.5\,\mathrm{MeV/c^2}$ and $\sim 5.0\,\mathrm{MeV/c^2}$ respectively for the on-line and off-line cases. Note that very strict mass cuts are applied in the off-line. The contribution of upstream and long tracks can be seen on Figure (24) where the reconstructed B_d mass is plotted. Though it is small compared to the long track case, we see a non-negligible fraction of the candidates formed with upstream tracks. These tracks clearly contribute to the mass with a larger resolution compared to the forward matched tracks. The core resolution assuming a single contribution is for the on-line B_d mass $\sim 29.8\,\mathrm{MeV/c^2}$ and $\sim 14.7\,\mathrm{MeV/c^2}$ for the off-line. These resolutions are obtained on $B_d \to D^*\ (\to D^0\ (\to K^+\pi^-)\pi)\pi$ events.

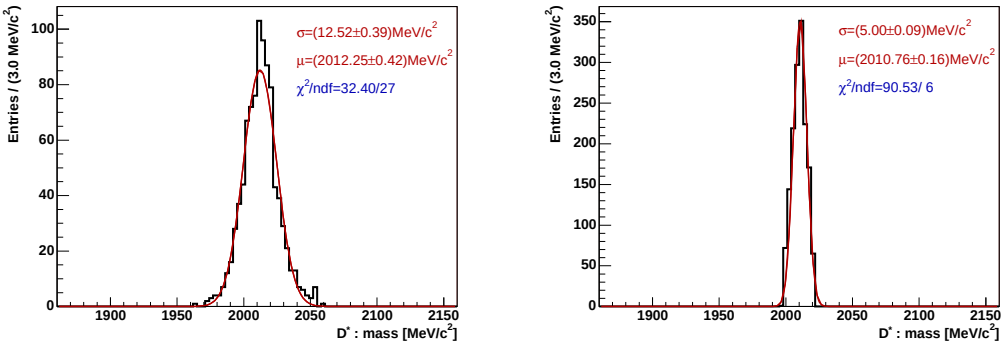

Figure 23: The reconstructed mass of the selected associated D^* signal candidates after the dedicated exclusive HLT selection and on off-line selected $B_d \to D^*\pi$ events passing the generic HLT. Right plot: on-line reconstructed mass [MeV/c^2]. Left plot: off-line reconstructed mass [MeV/c^2]. The masses are fitted with a single Gaussian.

The efficiency of the generic HLT for $B_d \to D^*\pi$ is $\varepsilon_{\mathrm{HLTGen/L1}} \sim 95.3\%$ and the cheated selection efficiency is $\varepsilon^{\mathrm{cheated}}_{\mathrm{HLT}} \sim 78.3\%$, essentially due to the slow pion. The number of events selected by the dedicated $B_d \to D^*\pi$ HLT selection and the corresponding specific efficiency on $N^{\mathrm{evt}}_{\mathrm{off}}$ signal events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	$N^{\mathrm{evt}}_{\mathrm{off}}$	$N^{\mathrm{specific}}_{\mathrm{HLT}}$	$\varepsilon^{\mathrm{specific}}_{\mathrm{HLT}}$ [%]	$N^{\mathrm{exb}}_{\mathrm{HLT}}$	$N^{\mathrm{b}\to\mu}_{\mathrm{HLT}}$	$N^{\mu\mu}_{\mathrm{HLT}}$	$N^{\mathrm{D}^*}_{\mathrm{HLT}}$	$N^{\mathrm{tot}}_{\mathrm{HLT}}$
$B_d \to D^*\pi$	1568	1060	67.6 ± 1.2	1063	59	7	996	1161

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	$\varepsilon^{\mathrm{exb}}_{\mathrm{HLT}}$ [%]	$\varepsilon^{\mathrm{b}\to\mu}_{\mathrm{HLT}}$ [%]	$\varepsilon^{\mu\mu}_{\mathrm{HLT}}$ [%]	$\varepsilon^{\mathrm{D}^*}_{\mathrm{HLT}}$ [%]	$\varepsilon^{\mathrm{tot}}_{\mathrm{HLT}}$ [%]
$B_d \to D^*\pi$	67.8 ± 1.2	3.8 ± 0.5	0.5 ± 0.2	63.5 ± 1.2	74.0 ± 1.1

This channel clearly benefits from the D^* stream. Besides the tracking inefficiency, this channels suffers from large combinatorics introduced by the slow pion on which no

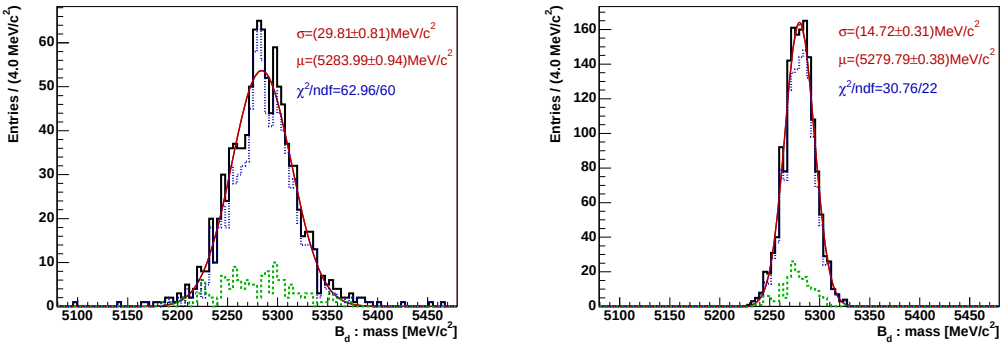

Figure 24: The reconstructed mass of the selected associated B_d signal candidates after the dedicated exclusive HLT selection and on off-line selected $B_d \to D^*\pi$ events passing the generic HLT. The contributions from the slow pions made from long (dashed blue line) and from upstream (dotted green line) tracks are superimposed. Right plot: on-line reconstructed mass [MeV/c^2]. Left plot: off-line reconstructed mass [MeV/c^2]. The masses are fitted with a single Gaussian.

cuts are applied. Indeed cutting on this particle reduces significantly the signal. Then the off-line selection strongly relies on one key selection criterion which is the D^*-D^0 mass difference and by applying an extremely strict cut. We cannot do the same in the HLT as we wish to keep large enough sidebands. The mass window for the D^*-D^0 mass difference was chosen for the exclusive HLT selection to approximatively correspond to the difference with respect to the $D^* \to D^0\pi$ threshold energy.

Exclusive $B_d \to \mu^+\mu^- K^{*0}$

The performance for $B_d \to \mu^+\mu^- K^{*0}$ is checked against $B_d \to \mu^+\mu^-$ K^{*0} ($\to K^+\pi^-$) events. The on-line B_d mass resolution is $\sim 25.2\,\mathrm{MeV/c^2}$ and the off-line one is $\sim 14.4\,\mathrm{MeV/c^2}$.

The efficiency of the generic HLT for $B_d \to \mu^+\mu^- K^{*0}$ is $\varepsilon_{\mathrm{HLTGen/L1}} \sim 98.2\%$ and the cheated selection efficiency is $\varepsilon_{\mathrm{HLT}}^{\mathrm{cheated}} \sim 79.2\%$.[14] The number of events selected by the dedicated $B_d \to \mu^+\mu^- K^{*0}$ HLT selection and the corresponding specific efficiency on $N_{\mathrm{off}}^{\mathrm{evt}}$ signal events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	$N_{\mathrm{off}}^{\mathrm{evt}}$	$N_{\mathrm{HLT}}^{\mathrm{specific}}$	$\varepsilon_{\mathrm{HLT}}^{\mathrm{specific}}$ [%]	$N_{\mathrm{HLT}}^{\mathrm{exb}}$	$N_{\mathrm{HLT}}^{\mathrm{b}\to\mu}$	$N_{\mathrm{HLT}}^{\mu\mu}$	$N_{\mathrm{HLT}}^{\mathrm{D}^*}$	$N_{\mathrm{HLT}}^{\mathrm{tot}}$
$B_d \to \mu^+\mu^- K^{*0}$	1074	775	72.2 ± 1.4	784	671	619	33	1015

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	$\varepsilon_{\mathrm{HLT}}^{\mathrm{exb}}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mathrm{b}\to\mu}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mu\mu}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mathrm{D}^*}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mathrm{tot}}$ [%]
$B_d \to \mu^+\mu^- K^{*0}$	73.0 ± 1.4	62.5 ± 1.5	57.6 ± 1.5	3.1 ± 0.5	94.5 ± 0.7

For this channel the preliminary cuts on the hadrons appear to be important sources of inefficiency. The off-line selection criteria should thus be subject to revision. Even though the HLT streams are far from being optimal, a pretty high total HLT efficiency can

[14] Note that, as explained in Appendix (A.2), we use some arbitrary association for on-line muons.

be achieved. This indirectly shows that the HLT streams are not that correlated. This is also apparent for instance by comparing the efficiencies for the generic HLT and the muon streams.

Exclusive $B_s \to \mu^+\mu^-$

This channel is expected to be the easiest one regarding the trigger selection, with no combinatorics at all. Indeed, requiring for two high p_T whose invariant mass is close to the B_s one and well above the J/ψ mass will easily kill all the background. The mass resolutions for the on-line and off-line B_s are respectively $\sim 29.1\,\mathrm{MeV/c^2}$ and $\sim 18.4\,\mathrm{MeV/c^2}$ as shown in Figure (25).

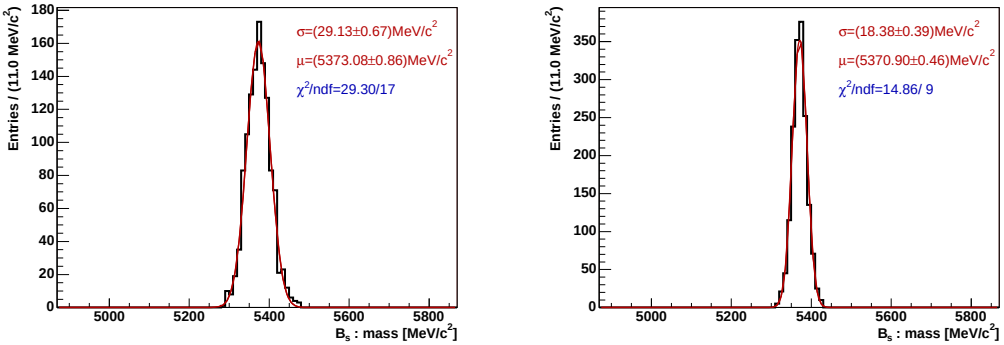

Figure 25: The reconstructed mass of the selected associated B_s signal candidates after the dedicated exclusive HLT selection and on off-line selected $B_s \to \mu^+\mu^-$ events passing the generic HLT. Right plot: on-line reconstructed mass [MeV/c^2]. Left plot: off-line reconstructed mass [MeV/c^2]. The masses are fitted with a single Gaussian.

The efficiency of the generic HLT for $B_s \to \mu^+\mu^-$ is $\varepsilon_{\mathrm{HLTGen/L1}} \sim 98.8\%$ and the cheated selection efficiency is $\varepsilon_{\mathrm{HLT}}^{\mathrm{cheated}} \sim 89.5\%$.[15] The number of events selected by the dedicated $B_s \to \mu^+\mu^-$ HLT selection and the corresponding specific efficiency on $N_{\mathrm{off}}^{\mathrm{evt}}$ signal events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	$N_{\mathrm{off}}^{\mathrm{evt}}$	$N_{\mathrm{HLT}}^{\mathrm{specific}}$	$\varepsilon_{\mathrm{HLT}}^{\mathrm{specific}}$ [%]	$N_{\mathrm{HLT}}^{\mathrm{exb}}$	$N_{\mathrm{HLT}}^{\mathrm{b}\to\mu}$	$N_{\mathrm{HLT}}^{\mu\mu}$	$N_{\mathrm{HLT}}^{\mathrm{D}^*}$	$N_{\mathrm{HLT}}^{\mathrm{tot}}$
$B_s \to \mu^+\mu^-$	1645	1486	90.3 ± 0.7	1573	1567	1459	26	1639

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

Channel	$\varepsilon_{\mathrm{HLT}}^{\mathrm{exb}}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mathrm{b}\to\mu}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mu\mu}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mathrm{D}^*}$ [%]	$\varepsilon_{\mathrm{HLT}}^{\mathrm{tot}}$ [%]
$B_s \to \mu^+\mu^-$	95.6 ± 0.5	95.3 ± 0.5	88.7 ± 0.8	1.6 ± 0.3	99.6 ± 0.2

When combining all the streams we see that the total HLT efficiency is $\approx 100\%$. Looking at the dedicated exclusive selection $\varepsilon_{\mathrm{HLT}}^{\mathrm{specific}}$ we notice that we actually lose a large fraction of the events. This is not due to the selection cuts but to the missing muons. Besides the tracking issues, there are two additional sources of inefficiency: the online momentum cut

[15] Note that, as explained in Appendix (A.2), we use some arbitrary association for on-line muons.

in the muon identification, and the inefficiency of the on-line muon identification.[16] The total exclusive efficiency $\varepsilon^{\text{exb}}_{\text{HLT}}$ is much higher than the dedicated one as we add the result of the $B_q \to hh$ exclusive selection that triggers most of the signal events, approximatively $\sim 88.5\%$.

Exclusive $B_s \to \phi\gamma$

The quality of the photon reconstruction in the HLT is within errors compatible with the off-line one as it can be seen from Figure (26). Factorizing in the stochastic and calibration contributions the photon energy resolutions are 2.58% and 2.50% for the on-line and off-line. We note the better calibration in the on-line parameters. This photon resolution translate to the observed invariant B_s mass resolutions shown in Figure (27). For the on-line we have $\sim 69.3\,\text{MeV}/c^2$ and for the off-line it is slightly better, $\sim 64.5\,\text{MeV}/c^2$ due to the charged tracks, but we are in both cases clearly dominated by the photon contribution.

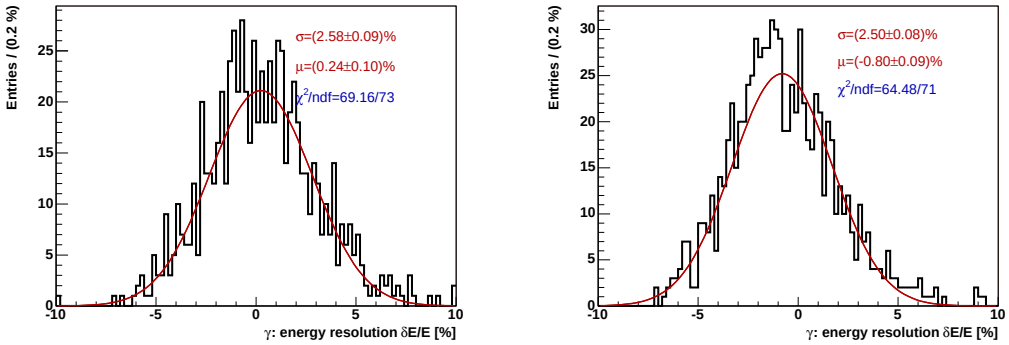

Figure 26: The residual of the photon energy from the associated B_s signal candidates after the dedicated exclusive HLT selection and on off-line selected $B_s \to \phi\gamma$ events passing the generic HLT. Right plot: on-line γ energy resolution [%]. Left plot: off-line γ energy resolution [%]. The residuals are fitted with a single Gaussian.

The efficiency of the generic HLT for $B_s \to \phi\gamma$ is only $\varepsilon_{\text{HLTGen/L1}} \sim 73.4\%$. This is merely due to the absence of ECAL trigger in the generic HLT and hence we lose all the bandwidth from the Level-1 photon line. The cheated selection efficiency is $\varepsilon^{\text{cheated}}_{\text{HLT}} \sim 95.1\%$, including the photon reconstruction. Basically all the photons found in the off-line are found in the on-line as well, hence leaving us with the usual tracking inefficiencies. Note that the presence of a ϕ further worsen the track finding efficiency, as opposed to the $B_q \to hh$ channels. The number of events selected by the dedicated $B_s \to \phi\gamma$ HLT selection and the corresponding specific efficiency on $N^{\text{evt}}_{\text{off}}$ signal events after the generic HLT are given below, together with the detailed number of events of all the HLT streams. The results are:

Channel	$N^{\text{evt}}_{\text{off}}$	$N^{\text{specific}}_{\text{HLT}}$	$\varepsilon^{\text{specific}}_{\text{HLT}}$ [%]	$N^{\text{exb}}_{\text{HLT}}$	$N^{\text{b}\to\mu}_{\text{HLT}}$	$N^{\mu\mu}_{\text{HLT}}$	$N^{\text{D}^*}_{\text{HLT}}$	$N^{\text{tot}}_{\text{HLT}}$
$B_s \to \phi\gamma$	855	762	89.1 ± 1.1	762	29	0	15	771

The corresponding full breakdown of the HLT efficiencies for the different streams is listed hereafter:

[16] Private communication from Alessia Satta.

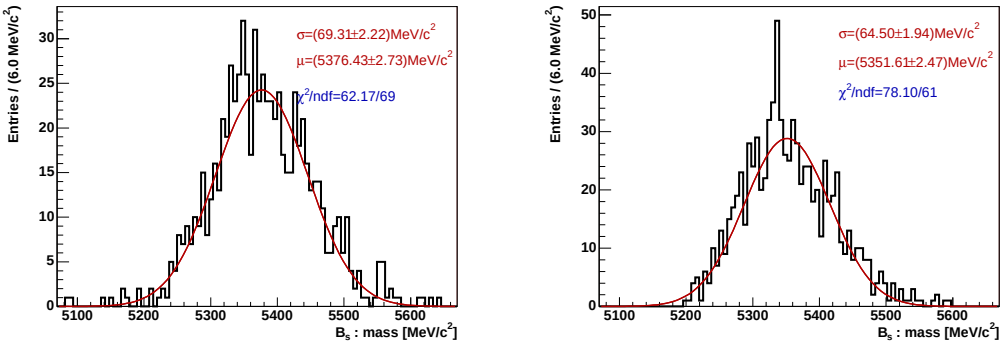

Figure 27: The reconstructed mass of the selected associated B_s signal candidates after the dedicated exclusive HLT selection and on off-line selected $B_s \to \phi\gamma$ events passing the generic HLT. Right plot: on-line reconstructed mass [MeV/c^2]. Left plot: off-line reconstructed mass [MeV/c^2]. The masses are fitted with a single Gaussian.

Channel	ε^{exb}_{HLT} [%]	$\varepsilon^{b\to\mu}_{HLT}$ [%]	$\varepsilon^{\mu\mu}_{HLT}$ [%]	$\varepsilon^{D^*}_{HLT}$ [%]	ε^{tot}_{HLT} [%]
$B_s \to \phi\gamma$	89.1 ± 1.1	3.4 ± 0.6	0.0 ± 0.0	1.8 ± 0.5	90.2 ± 1.0

The $B_s \to \phi\gamma$ HLT selection is challenging as we have a loosely determined ϕ vertex, we are overwhelmed by photons that do not contribute to the position determination, and we cannot apply strict mass cuts to reduce combinatorics. Requiring large transverse momentum and energy for the ϕ and γ enables to reduce the bandwidth while keeping a high efficiency, a few percent smaller than the $B_q \to hh$ one.

D^* stream

We use the $B_d \to D^*\pi$ decay channel, with $D^0 \to K^+\pi^-$ to check the performance on signal of the D^* stream. The D^* stream is always reconstructed assuming the pion hypothesis for all the stable decay products, using the $D^0 \to \pi^+\pi^-$ mode. There are loose mass asymmetric cuts on both the D^* and the D^0 as given in Table (21). Given their mass resolutions, the upper bound corresponds to a cut of more than four standard deviations. The lower limit is set to remove the long negative tails, though we have to be careful since we underestimate the mass for the $D^0 \to K^+\pi^-$ and $D^0 \to K^+K^-$ cases as we only consider the D^0 to two pion decays in the D^* stream. The reconstructed D^* stream mass distribution on signal with $D^0 \to K^+\pi^-$ is shown in Figure (28), where the contributions from long and upstream tracks are also drawn. We see that almost all signal candidates fall in the asymmetric mass window, which justifies our choice. The D^0 mass cuts have also been tested for the $D^0 \to K^+K^-$ case, using $B_d \to D^0\, K^{*0}$ data. The signal candidates falling in the negative tails have been found to be badly reconstructed. Moreover, these candidates also fall in the upper tails of the distribution of the D^*-D^0 mass difference because of the strong correlation induced by the small available phase space for the slow pion.

The efficiency of the D^* stream is 63.5% as it can be read off the table given for the $B_d \to D^*\pi$ selection, which is a few percent smaller than the dedicated counterpart. Though we have a much larger available bandwidth for the D^* stream, we need to cut much harder if we want to keep loose mass windows. Moreover, we do not have the extra constraints of the fast pion and of forming a B_d to suppress the background. In particular,

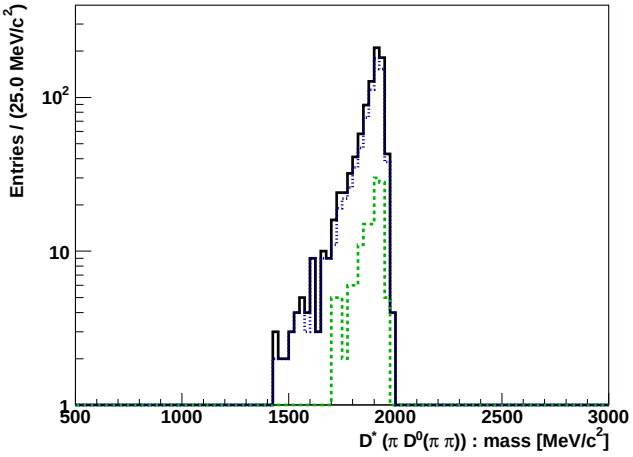

Figure 28: The reconstructed mass [MeV/c^2] of the selected associated D^* signal candidates after the dedicated D^* stream HLT selection and on off-line selected $B_d \to D^*\pi$ events passing the generic HLT. The contributions from the slow pions made from long (dashed blue line) and from upstream (dotted green line) tracks are superimposed.

we have to impose a sIPS cut on the slow pion to reject the large combinatorics introduced by the pions. Finally the exclusive and D^* streams are not completely correlated as the combination of the two results in a significant gain for the total efficiency on the $B_d \to D^*\pi$ signal channel.

Table 23: The number of events selected by the dedicated HLT selections and the corresponding specific efficiency on N^{evt}_{off} signal events after the generic HLT. The uncertainties are statistical.

Channel	N^{evt}_{off}	$N^{specific}_{HLT}$	$\varepsilon^{specific}_{HLT}$ [%]	N^{exb}_{HLT}	$N^{b\to\mu}_{HLT}$	$N^{\mu\mu}_{HLT}$	$N^{D^*}_{HLT}$	N^{tot}_{HLT}
$B_s \to D_s h$	5701	4400	77.2 ± 0.6	4407	206	26	196	4496
$B_q \to J/\psi X$	3008	1853	61.6 ± 0.9	2197	1389	2449	84	2796
$B_q \to hh$	4957	4637	93.5 ± 0.4	4638	134	11	59	4652
$B_d \to D^0 K^{*0}$	996	651	65.4 ± 1.5	653	21	0	50	675
$B_s \to \phi\phi$	986	782	79.3 ± 1.3	787	37	3	26	794
$B_d \to D^*\pi$	1568	1060	67.6 ± 1.2	1063	59	7	996	1161
$B_d \to \mu^+\mu^- K^{*0}$	1074	775	72.2 ± 1.4	784	671	619	33	1015
$B_s \to \mu^+\mu^-$	1645	1486	90.3 ± 0.7	1573	1567	1459	26	1639
$B_s \to \phi\gamma$	855	762	89.1 ± 1.1	762	29	0	15	771

The number of events considered and all the HLT efficiencies on signal are summarized in Table (23) and Table (24). As we can see from all the efficiencies on signal, the performance of the HLT can be very different depending on the decay channel we consider. The five dominant sources of performance differences identified are:

1. The on-line track finding. A large fraction of the signal events is lost because the

Table 24: HLT efficiencies on generic HLT triggered events and after the dedicated off-line selections. The uncertainties are statistical.

Channel	ε^{exb}_{HLT} [%]	$\varepsilon^{b\to\mu}_{HLT}$ [%]	$\varepsilon^{\mu\mu}_{HLT}$ [%]	$\varepsilon^{D^*}_{HLT}$ [%]	ε^{tot}_{HLT} [%]
$B_s \to D_s h$	77.3 ± 0.6	3.6 ± 0.3	0.5 ± 0.1	3.4 ± 0.2	78.9 ± 0.5
$B_q \to J/\psi X$	73.0 ± 0.8	46.2 ± 0.9	81.4 ± 0.7	2.8 ± 0.3	93.0 ± 0.5
$B_q \to hh$	93.6 ± 0.4	2.7 ± 0.2	0.2 ± 0.1	1.2 ± 0.2	93.9 ± 0.3
$B_d \to D^0 K^{*0}$	65.6 ± 1.5	2.1 ± 0.5	0.0 ± 0.0	5.0 ± 0.7	67.8 ± 1.5
$B_s \to \phi\phi$	79.8 ± 1.3	3.8 ± 0.6	0.3 ± 0.2	2.6 ± 0.5	80.5 ± 1.3
$B_d \to D^*\pi$	67.8 ± 1.2	3.8 ± 0.5	0.5 ± 0.2	63.5 ± 1.2	74.0 ± 1.1
$B_d \to \mu^+\mu^- K^{*0}$	73.0 ± 1.4	62.5 ± 1.5	57.6 ± 1.5	3.1 ± 0.5	94.5 ± 0.7
$B_s \to \mu^+\mu^-$	95.6 ± 0.5	95.3 ± 0.5	88.7 ± 0.8	1.6 ± 0.3	99.6 ± 0.2
$B_s \to \phi\gamma$	89.1 ± 1.1	3.4 ± 0.6	0.0 ± 0.0	1.8 ± 0.5	90.2 ± 1.0

tracks are not reconstructed. Large multiplicity decays are the most affected by this as they have more tracks and with smaller momenta. We also observe inefficiencies in the reconstruction and identification of muons.

2. The type of decay. Having at least a muon in the decay chain helps to cover for the tracking inefficiencies as we benefit from the generous bandwidth granted to muons. Channels with large combinatorics, such as the presence of a D_s, will undergo larger losses as we need to cut harder to suppress the background.

3. The on-line–off-line reconstruction quality differences. The momentum and IP resolutions, the vertexing and more importantly the difference in errors lead to inefficiencies that increase with the decay multiplicity.

4. The primary vertex reconstruction. A missed PV gives rise to large combinatorics possibly faking the presence of a secondary vertex.

5. The on-line–off-line selection differences. The preliminary cuts imposed to all long tracks result in a loss of events due to the quality difference and to the choice of off-line selection cuts. Selections applying very loose such cuts, in particular the hadronic channels, will inevitably lose events since we anyways apply these cuts in the HLT. The use of common standard composite particles with momentum cuts on the decay products will certainly indirectly improve the HLT selection efficiencies.

As an illustration of the fact that a larger correlation and alignment of the off-line selections with the HLT results in better efficiencies we can consider the $B_s \to \phi\,(K^+\,K^-)\eta_c$ (4h) decay channel, which involves six hadron final states. Even though this channel tremendously suffers from the missing on-line tracks, $\sim 26\%$ are lost because of it, the overall HLT efficiency is $\sim 73\%$, for a negligible minimum-bias output rate: < 1 Hz. The loss due to the HLT selection criteria is very small thanks to the use of the same code with looser cuts for the exclusive HLT selection and the efficiency is for instance much better than for $B_d \to D^0\,K^{*0}$ and even comparable to that of $B_d \to D^*\pi$.

For completeness we give in Table (25) a comparison between the on-line and off-line of all the mass resolutions involved in the HLT selections and that were not already given in Table (11).

Table 25: Mass resolutions σ_m of the on-line and off-line composite particles and b-hadron. The quoted values are from single Gaussian fits.

Channel	On-line σ_m [MeV/c^2]	Off-line σ_m [MeV/c^2]
$B_s \to D_s h$	29.4	13.2
J/ψ ($B_q \to J/\psi X$)	19.2	10.8
$B_q \to hh$	33.4	16.5
$B_d \to D^0 K^{*0}$	24.8	11.3
$B_s \to \phi\phi$	21.1	11.5
D^*	12.5	5.0
$B_d \to D^*\pi$	29.8	14.7
$B_d \to \mu^+\mu^- K^{*0}$	25.2	14.4
$B_s \to \mu^+\mu^-$	29.1	18.4
$B_s \to \phi\gamma$	69.3	64.5

5.2 HLT Selections Performance on Minimum-bias

The HLT performance on minimum-bias can now be assessed using the whole trigger chain. For that we consider 131,320 minimum-bias events (DC04-v2) at 40 kHz, which is equivalent to ~ 3.28 s of data taking. The HLT selection criteria have been tuned on a different sample (DC04-v1) and are thus blindly applied. The generic HLT already removes an important part of the minimum-bias events, as we are left with about 9 kHz after it, including the 1.5 kHz from the muons streams. Hence we see that we still need to reduce the rate by at least a factor 15 to bring the rate for the HLT selections down to 500 Hz. This is actually the tough part given the natural charm and beauty components in the minimum-bias. The task is even more arduous for the exclusive HLT given the numerous channels that it will contain and since only a small fraction of the total bandwidth is granted, 200 Hz. We will give in the next Sections the HLT output rates on minimum-bias and the timing performance of the whole HLT chain.

5.2.1 HLT Output Rates

Each of the HLT selections is ran on all events, after a positive decision of the generic HLT and the exclusive b stream will have a positive decision if one of the dedicated selections finds a candidate. As mentioned previously, the selection criteria were tuned to have at most $O(10\,\text{Hz})$ for the exclusive selections individually, whereas all the bandwidth for the D^* stream is used. The number of selected events in the minimum-bias sample and the corresponding output rates are given in Table (26) and Table (27). The rates of the other HLT streams ($b \to \mu$, $\mu\,\mu$ and D^*) for events selected by the dedicated HLT selection are given to show the correlations. We see from these tables that the maximal exclusive b output rate is (57.0 ± 4.2) Hz, i.e. when adding all exclusive b selected events and ignoring the correlations between them. The exclusive b rate is thus well within specification. For the D^* stream we have an output rate just below the nominal 300 Hz.

The exclusive selections involving muons obviously have large correlations with the muon streams. For instance all the events passing the $B_d \to \mu^+\mu^- K^{*0}$ selection are also

Table 26: Number of selected events $N^{specific}$ by each of the specific HLT selections on 131,320 L0-L1 stripped minimum bias events, with the result of the other streams $N^{specific}_{b\to\mu}$, $N^{specific}_{\mu\mu}$ and $N^{specific}_{D^*}$ on these $N^{specific}$ events.

Channel	$N^{specific}$	$N^{specific}_{b\to\mu}$	$N^{specific}_{\mu\mu}$	$N^{specific}_{D^*}$
$B_s \to D_s h$	36	2	2	4
$B_q \to J/\psi X$	55	18	48	1
$B_q \to hh$	12	2	0	1
$B_d \to D^0 K^{*0}$	15	0	0	3
$B_s \to \phi\phi$	0	0	0	0
$B_d \to D^*\pi$	34	1	1	6
$B_d \to \mu^+\mu^- K^{*0}$	7	2	7	0
$B_s \to \mu^+\mu^-$	0	0	0	0
$B_s \to \phi\gamma$	28	1	2	3
D^* stream	964	18	71	964

Table 27: Output rates $R^{specific}$ of each of the specific HLT selections on 131,320 L0-L1 stripped minimum bias events, with the shared rates of the other streams $R^{specific}_{b\to\mu}$, $R^{specific}_{\mu\mu}$ and $R^{specific}_{D^*}$ on the $R^{specific}$. The uncertainties are statistical.

Channel	$R^{specific}$ [Hz]	$R^{specific}_{b\to\mu}$ [Hz]	$R^{specific}_{\mu\mu}$ [Hz]	$R^{specific}_{D^*}$ [Hz]
$B_s \to D_s h$	11.0 ± 1.8	0.6 ± 0.4	0.6 ± 0.4	1.2 ± 0.6
$B_q \to J/\psi X$	16.8 ± 2.3	5.5 ± 1.3	14.6 ± 2.1	0.3 ± 0.3
$B_q \to hh$	3.7 ± 1.1	0.6 ± 0.4	0.0 ± 0.0	0.3 ± 0.3
$B_d \to D^0 K^{*0}$	4.6 ± 1.2	0.0 ± 0.0	0.0 ± 0.0	0.9 ± 0.5
$B_s \to \phi\phi$	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
$B_d \to D^*\pi$	10.4 ± 1.8	0.3 ± 0.3	0.3 ± 0.3	1.8 ± 0.8
$B_d \to \mu^+\mu^- K^{*0}$	2.1 ± 0.8	0.6 ± 0.4	2.1 ± 0.8	0.0 ± 0.0
$B_s \to \mu^+\mu^-$	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
$B_s \to \phi\gamma$	8.5 ± 1.6	0.3 ± 0.3	0.6 ± 0.4	0.9 ± 0.5
D^* stream	293.6 ± 9.4	21.6 ± 2.6	9.3 ± 1.7	293.6 ± 9.4

triggered by the di-muon streams. The overlap is also very large for the $B_q \to J/\psi X$ due to the high rate of true J/ψ. This selection is the one taking the largest fraction of the exclusive bandwidth, but can easily be reduced for instance by reducing the sidebands. Most of these events are anyways already triggered by the $b \to \mu$ and/or the $\mu\,\mu$ streams.

The output rates for a few channels can easily be under control, such as the 2-prong decays or $B_s \to \phi\phi$, due to their limited combinatorics or the presence of muons and charged tracks. Interestingly a true $B_d \to \pi^+\pi^-$ was found by the $B_q \to hh$ selection. The largest bandwidth among the purely hadronic channels is taken by $B_s \to D_s h$, because of the numerous D_s combinations. In general there is a very small correlation between the exclusive b selections and the muons streams. Even though the statistics are limited, we observe that the D^* stream shares a larger bandwidth with high multiplicity hadronic channels, which is a sign of the tendency to trigger on dirty events or events with a non-reconstructed primary vertex. Finally we can note that $\sim 7\%$ of the D^* bandwidth was also triggered by the inclusive muon stream.

We give in Table (28) the total output rates for the whole HLT, assuming 40 kHz Level-1 output rate. We see that the di-muon stream is well above the specification as a result of the latest Level-1 bandwidth share giving more bandwidth to muons. The overall HLT retention rate on minimum bias is below the design 2 kHz. The exclusive b output rate is obtained by taking the logical *OR* of the different output rates given in Table (27). All the selection efficiencies on minimum-bias for each of the HLT streams, channels and their correlations are given in the Table (30) of the Appendix (C).

Table 28: Minimum-bias trigger output rates. The uncertainties are statistical.

Trigger	N^{evt}	Rate [Hz]
Level-1	131320	40 kHz reference
Generic HLT	28458	8668.3 ± 45.5
HLT	6355	1935.7 ± 23.7
Exclusive b stream	180	54.8 ± 4.1
$b \to \mu$ stream	2547	775.8 ± 15.2
Di-muon stream	2981	908.0 ± 16.4
D^* stream	964	293.6 ± 9.4

An important criterion of the trigger performance is the fraction of interesting events in the selected bandwidth. Each step of the trigger is meant to successively reduce the minimum-bias output rate while enriching its content of beauty b and c charm quarks as these are the events we wish to reconstruct at LHCb. The quark content of the selected events after each of the HLT streams is given in Table (29), for events containing in the following order: a b-hadron and a b- or c-hadron. The charm component includes all charmed hadrons and not necessarily only the open charm particles. We first notice that the di-muon streams has the lowest 'purity', which is intended as we want all $J/\psi \to \mu^+\mu^-$. The HLT selections, i.e. the exclusive b and the D^*, have comparable fractions of events with a b and a b or c quarks. A large fraction of charm component is given by the $B_q \to J/\psi X$ selecting true J/ψ. The detail of the quark content for each of exclusive selection can be read off Table (30) of the Appendix (C). As we can see from the quark content table, there is still a significantly large fraction of events selected by the HLT selections that do not contain any b or c quarks, of the order of $\mathcal{O}(30\%)$. These events need further

investigation in order to understand the nature of our minimum-bias background.

Table 29: Quark contents, b and b or c, of the HLT triggered minimum-bias events.

Trigger	b [%]	b or c [%]
Generic HLT	34.1	56.2
HLT	39.9	66.6
Exclusive b stream	52.2	71.7
$b \rightarrow \mu$ stream	70.6	90.5
Di-muon stream	11.9	46.0
D^* stream	52.8	73.4

The reconstructed mass of all the selected candidates in the minimum-bias is shown in Figure (29) which illustrates the mass windows used in the HLT selections. The peak centered at the J/ψ mass correspond to the $B_q \rightarrow J/\psi X$ selected candidates, which are all true J/ψ. The exclusive b mass spectrum shows that the selected candidates have an evenly distributed mass due to the combinatorics. Nevertheless there is a distinctive peaking background at low mass. These candidates are presumably due to partially reconstructed b decays, as we do select more than 50% of the events with a b quark.

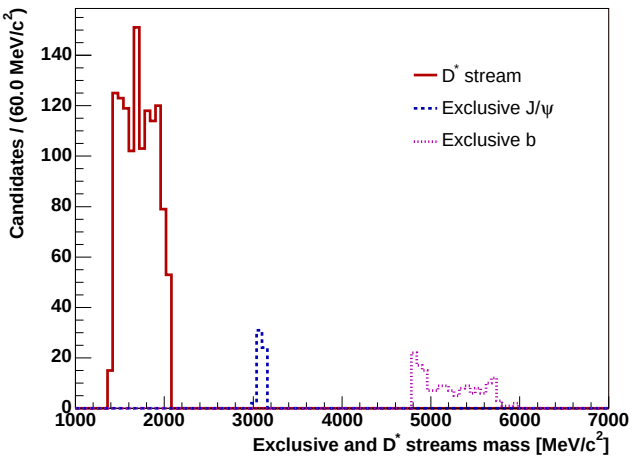

Figure 29: Reconstructed mass [MeV/c^2] of the HLT selected candidates with the exclusive b and D^* streams.

In order to understand the nature of the selected events without any b nor b quarks, we now introduce a classification of the selected candidates according to the primary vertex reconstruction, as this is the main suspect for our combinatorial events.

The first step of the classification is to match the true primary vertices (MCPV) to the reconstructed (PV) ones in order to spot if we have missed primaries. The matching of a MCPV and PV is done by first assigning to a MCPV the closest PV when computing the 3-d distance. Then to take into account the resolution effects, we arbitrarily require that for each directions $i = x,y,z$ the distance between the MCPV and its closest PV should not differ by more than 5 σ: $|\mathrm{PV}_i - \mathrm{MCPV}_i| < 5\sigma_i$. If those conditions are satisfied then the

MCPV and the closest PV are matched. The σ_i are chosen to approximatively correspond to the primary vertex resolutions, namely $\sigma_{x,y} = 10\,\mu$m and $\sigma_z = 50\,\mu$m.

The second step of the classification is to compute the distance in 3-d between the primaries MCPV and the secondary vertices of the reconstructed HLT selected candidates to see if we can identify it as a missed primary vertex. We use the same procedure as before and we match a secondary reconstructed vertex V^{sec} to the closest $MCPV^{sec}$ if for each $i = x,y,z$ we have $|V_i^{sec} - MCPV_i^{sec}| < 10\sigma_i$, with a larger tolerance due to the worse secondary vertex determination.

Using the above classification, we can look at the distribution of the unsigned z-distance Δ_z between the secondary vertex V^{sec} and the corresponding closest $MCPV^{sec}$, as shown in Figure (30), for all the reconstructed candidates passing the exclusive HLT. The case where the $MCPV^{sec}$ actually corresponds to the same MCPV as the one matched to the reconstruct PV means that the reconstruction is fine. From the plot on the left hand side we see that most of the HLT candidates where the $MCPV^{sec}$ is not matched to the reconstructed PV have a very small distance difference, which is a clear sign that we may be reconstructing a primary vertex instead of a true secondary. The right hand side shows the Δ_z distribution only for candidates whose closest $MCPV^{sec}$ is not matched to any close PV, and hence potentially missed primaries. We can see from this plot the contribution of the candidates whose $MCPV^{sec}$ corresponds to the V^{sec}, in order words candidates that are in fact a missed primary vertex.

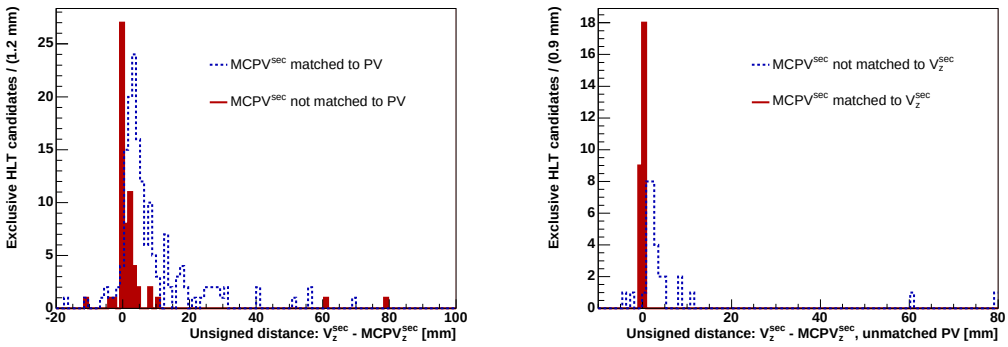

Figure 30: Unsigned distance in z [mm] between the secondary vertex V^{sec} of the selected exclusive HLT candidate and the corresponding closest $MCPV^{sec}$. Left plot: candidates for which $MCPV^{sec}$ is matched to the reconstructed PV (dashed blue line) or not (solid red line). Right plot: candidates for which $MCPV^{sec}$ is matched to the reconstructed PV and in turn matched to the V^{sec} (solid red line) or not (dashed blue line).

From the above simple classification we therefore see that the primary vertex reconstruction is crucial as a missed primary vertex represents a large source of combinatorics, especially for the hadronic modes.

5.2.2 HLT Timing Performance

We give in this Section the timing performance for the entire HLT sequence as measured on minimum-bias events. We assume an input rate after the Level-1 of 40 Hz, with 400 CPUs of the on-line farm dedicated to the on-line reconstruction and the HLT. The average computing time per event should thus not exceed 10 ms. Note that the new 1 MHz readout scheme will imply many modifications and shuffling of the resources for the parts prior

to the HLT selections and that we ignore here as they have not been implemented yet. The timing measurement is defined in Section (A.1) and all the times quoted hereafter correspond to clock times. These times were determined on a dedicated machine only running the HLT code and are therefore unaffected by the machine load.[17]

In the timing measurement given below, we have taken into account the efficiency of the generic HLT and the fact that the full tracking is not executed for each event after the generic HLT. To be conservative we ignore if some of the tracks were already reconstructed in the generic HLT and we count the time to reconstruct all the tracks after a positive generic HLT decision. Also the empty sequencers contributions were subtracted from the total time.

The per-event execution time for the full HLT is measured to be 30.04 ms. The error on this time is approximatively 0.05 ms, determined in a series of identical jobs. The contributions from the main components are given below, always normalized to the total time such that the time equivalent value can directly be obtained.

HLT Velo tracking $\sim 20.2\%$. It includes the times for the raw Velo buffer decoding, the Velo $r-z$ and 3-d tracking, and the primary vertex reconstruction. The main contribution to this part are:

- Velo-space tracking: $\sim 14.5\%$.

Generic HLT sequence $\sim 52.3\%$. It includes the times for the Velo-TT decoding and tracking, the partial forward tracking, the error parameterization, the muon reconstruction, the generic HLT and the muon streams decisions. The main contributions to this part are:

- TT decoding: $\sim 3.2\%$;
- Velo-TT tracking: $\sim 5.5\%$;
- Forward decoding: $\sim 5.3\%$;
- Forward tracking: $\sim 34.0\%$.

The following parts are only executed on generic HLT triggered events.

Full tracking $\sim 18.0\%$. It corresponds to the time for the rest of the forward tracking.

HLT particles $\sim 3.4\%$. It represents the times for the final state particles making. The main contributions to this part are:

- ECAL reconstruction: $\sim 2.1\%$;
- Particle making: $\sim 1.3\%$.

HLT selections $\sim 6.1\%$. It includes the times for the final states filtering, the shared composite particles making, the exclusive b and the D^* stream selections. The relevant contributions to this part are:

- Final states copying and filtering: $\sim 2.1\%$;
- Shared particles: $\sim 2.0\%$;
- Inclusive D^* selection: $\sim 0.3\%$;
- Exclusive b selections: $\sim 1.6\%$.

[17] We thank Loic Brarda and Niko Neufeld for setting up the pclbonsrv05 machine.

The time taken by the HLT selections including particle making amounts to less than 10% of the total time. Adding extra exclusive selections is not an issue as far as timing is concerned as most of the time spent in the HLT selections is in the creation and filtering of the shared particles. These particles will be in most cases already available to the new channels. The most consuming part is by far the forward tracking which represents more than 50% of the total.

The overall time of 30.04 ms corresponds to the performance for a single process. In the real experiment we will have at least two CPUs running on the *same* machine. i.e. 2 out of 4 cores for each box in the on-line farm dedicated to the HLT. If we compare this time with what will be required during data taking we therefore see that the overall timing is practically in the good range. The performance must however be re-evaluated once the new 1 MHz scheme will be implemented, especially with the numerous trigger alleys and hopefully the new pattern recognition for the forward tracking will be faster.

6 Outlook and Conclusion

The exclusive b and the D^* HLT selections are up and running. The selection criteria have been tuned to achieve the design output rates. These criteria will be revisited by the different physics working groups as new selections are added to the HLT, to optimize the physics goals. The purity of the selected bandwidth in terms of beauty and charm quarks is above 70%, where the remainder is due to reconstruction problems, such as the primary vertex search. The timing performance is almost in the required range, though further improvements are needed especially for the forward tracking and moreover with the 1 MHz readout scheme.

The total HLT efficiencies on the core signal channels vary in a large range, 70-100%, and strongly depend on the type and multiplicity of the decay. The specific efficiencies are essentially limited by the on-line–off-line differences, rather by than by a fine tuning of the selection requirements. The restricted processing time demands a fast dedicated on-line software that inevitably leads to HLT inefficiencies. The observed and identified main sources of inefficiency are summarized hereafter:

- On-line and off-line pattern recognitions are different yielding track finding inefficiencies for the tracks used in the off-line analyses, where the loss becomes important for high multiplicity decays. The use of TT is fundamental in the current implementation for the speed performance, but yields inefficiencies due to the lack of constraints and its geometry. Without TT and regardless of the timing, the forward tracking still misses a few percent of the off-line reconstructed tracks. The forward tracking is currently under revision for speed and efficiency. Tracking inefficiencies in the muon reconstruction are also observed, though these are covered by the high muon rate streams.

- There was a significant improvement in the primary vertex reconstruction strategy allowing a better discrimination of the minimum-bias background. However, part of the HLT bandwidth is due to the reconstructed signal candidate in place of a primary vertex.

- The use of a simplified covariance matrix for the tracks parameter give a few percent inefficiency for the HLT selections as the errors are different and systematically overestimated compared to the off-line. A more complete parameterization could be implemented (e.g. using a look-up table) but the performance is satisfactory for

the HLT needs, where anyway some inefficiency should be tolerated. The geometrical and vertexing tools using this constant parameterized covariance matrix also give reliable results, with a few limitations for vertices significantly displaced with respect to the beam axis (few percent effect).

- The HLT uses preliminary sIPS and p_{T} cuts on all the final states long tracks, introduced for speed and rate reasons. Thus off-line selection applying no or very loose requirements on these parameters will inevitably lose events as the HLT applies them anyway.

- The choice of final selection criteria indirectly degrades the performance as the HLT applies a series of standardized cuts in order to form shared composite particles. Thus the use of these kind of standard particles by the off-line selections would make them more uniform and in turn improve the HLT selection efficiencies. All efficiencies are with respect to off-line selected samples and in some cases the off-line criteria are difficult to justify or understand. The physics working groups should have a critical look at them since the overall HLT–off-line efficiency could only profit.

The first complete HLT prototype is fully implemented yielding reasonable results in terms of efficiency, timing and output rates. The development and study of the HLT selections performance has emphasized a few weaknesses of the on-line software. These issues should be considered and taken into account in the next iteration of the HLT related software, especially now that the Level-1 boundary has been removed.

Appendix

A Technicalities

In the study of the HLT performance the key points are:

Speed. The time performance should be measured on minimum-bias having passed the previous trigger levels, as these represent the kind of events the HLT will process.

Efficiency. The efficiency should be measured on signal events. We call here signal whatever decays are relevant for our physics goals but especially b-decays.

We give hereafter various definitions used in this note to analyze the HLT performance.

A.1 Timing Measurement

The time measurements are performed with the equivalent of a 2.8 GHz Xeon processor by means of the `SequencerTimerTool` used by the `GaudiSequencer`. Two different measures of the time are used in this note:

CPU time. It gives the CPU cycles used by the process with a millisecond sampling. This is well suited for studying the time taken by the HLT selections as a function of the p_T and sIPS of all long tracks when normalizing to some reference algorithm and independently of the machine load, see Section (4.5).

Clock time. It measures the elapsed time between the beginning and the end of an algorithm. This time is affected by the load of the machine. In order to have a reliable time measurement of the full HLT sequence, a 'clean' replica of a lxbatch node with a 2.8 GHz Xeon processor was used (pclbonsrv05), only running the HLT code.

All the timing estimates are obtained on minimum-bias events and applying the triggers prior to the HLT.

A.2 Efficiency Measurement

Throughout the development of the HLT selections the Monte Carlo (MC) truth was used to understand and tune the different algorithms. The access to the MC truth depends on what we wish to measure and determine. We also need to specify how are efficiencies defined, namely with respect to what sample or quantity we normalize them. For this reason we introduce hereafter various definitions related to the MC truth.

The association to the MC truth for on-line `TrgTrack` tracks is done using `Linker` objects associating the `ChannelID` of a measurement to `MCParticles` [8]. As each reconstructed track carries a collection of measurements from the tracking detectors, we can determine how many clusters are associated to a given `MCParticle`. The criteria to associate an on-line track actually depend on the part of the track we consider, which translate to the track types as:

- A requirement of 70% of the Velo measurements to be associated to the same MC particle is asked for a track to be reconstructed as a Velo track.
- A Velo-TT track is associated if its Velo part is and if the TT part has at most one wrong measurement. Tracks accepted by the Velo-TT algorithm without any momentum are counted as reconstructed even if there are no TT clusters attached to it, provided the Velo part is correctly reconstructed.

- For long tracks, 70% of the measurements for both and independently the Velo and the T stations (inner and outer tracker) need to be associated to the same MC Particle and the TT measurements are ignored. In particular a track could be reconstructed as long even if the TT part has more than one wrong cluster.

- A track is a ghost if it is not associated to any MC particle. Ghost are reconstructed tracks made of clusters from different MC particles, and/or from noise.

The same procedure extends to `Particles` which are just tracks with an assigned particle identification, hence a well-defined 4-vector. The association between particle and MC particle is done using the `Particle2MCLinks` (or the `CompositeParticle2MCLinks`) associator algorithm, with a work-around for particles made from on-line tracks for which the keys are needed to access the association tables. When running flavor-blind selections (e.g. reconstructing all $B_q \to hh$ modes with $h = \pi$), the particle identification is ignored in the association. There are yet two particular cases for the on-line particle to MC particle association. For HLT muons produced by the stand-alone muon reconstruction the matching to the MC truth is done by comparing the transverse momentum components with some arbitrary tolerance of 350 MeV/c, approximatively corresponding to a $\sim 5\sigma$ resolution cut. For HLT photons the association is done by matching on-line clusters to the off-line electromagnetic calorimeter clusters and by asking the momentum components to be matched within 2 GeV/c.

For the on-line tracking performance we are interested in reconstructing all the tracks used off-line in a final selection. As most final selections do not have a 100% purity because of combinatorial background or ghosts in the signal sample, the track finding efficiencies are computed on events where all the charged signal final states are off-line reconstructed as long tracks[18] in order to have better estimates. Moreover the off-line selection and the generic HLT are applied to these events. The different tracking efficiencies are then computed as:

- Per-track efficiency: $\varepsilon^{\rm tr}_{\rm rec} = N^{\rm tr}_{\rm rec}/(N^{\rm tr}_{\rm fin} * N^{\rm evt}_{\rm off})$;

- Per-event efficiency: $\varepsilon^{\rm evt}_{\rm rec} = N^{\rm rec}_{\rm evt}/N^{\rm evt}_{\rm off}$;

with the following definitions:

- $N^{\rm tr}_{\rm fin}$: number of charged final states in the decay;

- $N^{\rm evt}_{\rm off}$: number of events after generic HLT and off-line selection, for events with all signal final states off-line reconstructed;

- $N^{\rm tr}_{\rm rec}$: number of on-line reconstructed tracks;

- $N^{\rm rec}_{\rm evt}$: number of events with *all* final states on-line reconstructed.

Note that for the per-event efficiency $\varepsilon^{\rm evt}_{\rm rec}$ we require all the final states to be on-line reconstructed and hence it is not determined as $(\varepsilon^{\rm tr}_{\rm rec})^{N^{\rm tr}_{\rm fin}}$. The ghost rates are determined on minimum-bias events, after the generic HLT selection.

The HLT selection efficiencies on off-line selected signal events are computed after the generic HLT, independently if the off-line selected event contains a correctly reconstructed candidate. The selection efficiencies include the track finding efficiencies. In particular we will consider the following definitions:

[18] A charged MC particle is reconstructed as an off-line long track if it has an associated `ProtoParticle` originating from a long `TrStoredTrack`.

- $\varepsilon_{\rm HLT}^{\rm specific}$: the efficiency for a given decay with the dedicated specific HLT selection;
- $\varepsilon_{\rm HLT}^{\rm exb}$: the HLT efficiency of the exclusive b stream;
- $\varepsilon_{\rm HLT}^{\rm b\to\mu}$: the HLT efficiency of the inclusive b$\to\mu$ single muon stream;
- $\varepsilon_{\rm HLT}^{\mu\mu}$: the HLT efficiency of the di-muon stream;
- $\varepsilon_{\rm HLT}^{\rm D^*}$: the HLT efficiency of the D^* stream;
- $\varepsilon_{\rm HLT}^{\rm tot}$: the total HLT efficiency of all HLT streams.

A measure of the total track reconstruction efficiency is given by the cheated HLT selection efficiency $\varepsilon_{\rm HLT}^{\rm cheated}$, which is equivalent to running the HLT selections without any cut and only with the tracks associated to the signal. This efficiency corresponds to the per-event tracking efficiency *provided* we reconstruct the mother of the decay with all the associated final states since the same MC particle could be associated more than once, and only if the vertex fit converges.

In all the signal efficiencies presented in this note, we always require the events to have previously triggered the generic HLT. To give a hint of the generic HLT performance on signal we will also give its efficiency $\varepsilon_{\rm HLTGen/L1}$ with respect to the Level-1 triggered events, on a of subsample off-line selected signal events.

B Recent On-line Reconstruction Improvements

We describe hereafter the most significant improvements to the on-line reconstruction and used in our performance study.

Primary vertex (PV) reconstruction [9]

The primary vertex search has largely been improved and consequently the fraction of non-reconstructed primaries identified as secondary b-vertices is better under control. The new implementation is a two-step procedure using Velo-3d tracks where first the vertex seeds are looked for and second the iterative vertex fit is performed. The fit now can merge (split) close (large) primary vertices to avoid the misidentification of primaries and secondaries. The borderline between a primary and a secondary vertex is currently set by requiring a minimum of six tracks to form a primary vertex.

Velo-TT tracking [10]

The investigation of the HLT inefficiencies has shown that they were dominated by the track finding problems. The Velo-TT tracking has large inefficiencies because of the lack of constraints one can impose on the TT segments, since there are only four plane. Moreover, many tracks are not in the TT sensitive area: they are either not in the acceptance (in central hole or outside planes) or traverse the dead regions (gaps in y between consecutive bonded sensors). Whereas TT clusters are *not* required in the off-line tracking, in the HLT the Velo-TT tracks are the "seeds" of the forward tracking: the momentum estimate obtained from TT is primordial for the speed performance of the long tracking. Based on these observations, and with a compromise between speed and performance, a new strategy for the Velo-TT tracking has been developed for both the L1 and the HLT. [19] The Velo-TT reconstruction for the HLT

[19] The strategy is however a bit different for the Level-1 for timing reasons: a single pass is done, and all created tracks have a defined momentum estimate.

has two passes: the first with strong matching tolerances and global clean-up of candidates and the second with looser tolerances for Velo tracks not matched in the first step and only with unused TT clusters. As a result, a significant improvement in the tracking efficiencies was achieved for a relatively small extra time penalty. This improvement translates to $\sim 1\%$ increase in efficiency per-long track. The improvements have a larger impact for higher multiplicity decays. For instance the track finding efficiency for having all final states in $B_s \to D_s\pi$ goes from $\sim 82\%$ to $\sim 87\%$.

Muon tracking and identification [11]

The muon identification has a dedicated pattern recognition in the HLT. Muon track segments are first looked for using information from the muon stations M5 to M2 and then matched to Velo tracks for which the standard on-line pattern recognition is run. Finally, the differences in x and y at M2 between the track extrapolation and the muon segments are compared to identify muons. Using this approach the muon reconstruction for the HLT does not directly depend on the standard on-line pattern recognition, and the muons can be reconstructed very fast.

C HLT Correlations

The correlations for each of the HLT streams and selections as obtained on the 131,320 minimum-bias events are given in Table (30). The channels not selecting any events are not included and all the entries are normalized to the generic HLT trigger. The absolute efficiency of each algorithm from the "Alg." column is given in the "Eff." column. The algorithms numbering is:

1. Generic HLT;
2. Full HLT;
3. $b \to \mu$ stream;
4. Exclusive b stream;
5. $\mu\ \mu$ stream;
6. D^* stream;
7. $B_q \to hh$;
8. $B_d \to D^0 K^{*0}$;
9. $B_d \to D^*\pi$;
10. $B_s \to D_s h$;
12. $B_s \to \phi\gamma$;
14. $B_d \to \mu^+\mu^- K^{*0}$;
15. $B_q \to J/\psi X$;
16. Beauty b event;
17. Beauty b or charm c event.

In the rest of the table, each entry in line indexed by l and column indexed by c gives the efficiency for algorithm l provided that algorithm c found something. For example the efficiency of $B_d \to D^*\pi$ ($l = 9$) is 0.62% when a D^* was found by the D^* stream ($c = 6$).

Table 30: HLT correlations on 131,320 minimum-bias events.

Alg.	Eff.	1	2	3	4	5	6	7	8	9	10	12	14	15	16	17
1	100.00%	######	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%
2	22.33%	22.33%	######	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	26.12%	26.46%
3	8.95%	8.95%	40.07%	######	14.44%	4.56%	7.36%	16.66%	0.00%	2.94%	5.55%	3.57%	28.57%	32.72%	18.53%	14.42%
4	0.63%	0.63%	2.83%	1.02%	######	1.91%	1.86%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	100.00%	0.96%	0.80%
5	10.47%	10.47%	46.90%	5.34%	31.66%	######	3.21%	0.00%	0.00%	2.94%	5.55%	7.14%	100.00%	87.27%	3.67%	8.57%
6	3.38%	3.38%	15.16%	2.78%	10.00%	1.04%	######	8.33%	20.00%	17.64%	11.11%	10.71%	0.00%	1.81%	5.24%	4.42%
7	0.04%	0.04%	0.18%	0.07%	6.66%	0.00%	0.10%	######	0.00%	0.00%	2.77%	0.00%	0.00%	0.00%	0.05%	0.05%
8	0.05%	0.05%	0.23%	0.00%	8.33%	0.00%	0.31%	0.00%	######	0.00%	5.55%	0.00%	0.00%	0.00%	0.09%	0.06%
9	0.11%	0.11%	0.53%	0.03%	18.88%	0.03%	0.62%	0.00%	0.00%	######	2.77%	0.00%	0.00%	0.00%	0.10%	0.11%
10	0.12%	0.12%	0.56%	0.07%	20.00%	0.06%	0.41%	8.33%	13.33%	2.94%	######	0.00%	14.28%	1.81%	0.23%	0.16%
12	0.09%	0.09%	0.44%	0.03%	15.55%	0.06%	0.31%	0.00%	0.00%	0.00%	0.00%	######	0.00%	1.81%	0.04%	0.06%
14	0.02%	0.02%	0.11%	0.07%	3.88%	0.23%	0.00%	0.00%	0.00%	0.00%	2.77%	0.00%	######	1.81%	0.06%	0.03%
15	0.19%	0.19%	0.86%	0.70%	30.55%	1.61%	0.10%	0.00%	0.00%	0.00%	2.77%	3.57%	14.28%	######	0.42%	0.34%
16	34.08%	34.08%	39.87%	70.59%	52.22%	11.94%	52.80%	41.66%	60.00%	29.41%	63.88%	14.28%	85.71%	74.54%	######	60.67%
17	56.17%	56.17%	66.56%	90.53%	71.66%	45.95%	73.44%	66.66%	66.66%	52.94%	75.00%	35.71%	85.71%	100.00%	100.00%	######

References

[1] PATRICK KOPPENBURG AND LUIS FERNÁNDEZ. HLT Exclusive Selections Design and Implementation. LHCb-2005-015, LPHE-2005-011, 2005.

[2] PATRICK KOPPENBURG AND LUIS FERNÁNDEZ. DAVINCI for Busy People. LHCb-2005-016, LPHE-2005-012, 2005.

[3] OLIVIER CALLOT. Online Pattern Recognition. LHCb-2004-094, 2004.

[4] JOSÉ-ANGEL HERNANDO ET AL. HLT-generic package: `Hlt/HltGeneric v2r8`. LHCb Software, 2006.

[5] HUGO RUIZ. Parameterization of track uncertainties for the HLT. LHCb-2005-012, 2005.

[6] HUGO RUIZ. Fast tools for vertexing and geometry calculations for the HLT. LHCb-2005-013, 2005.

[7] LUIS FERNÁNDEZ. HLT rates: old HLT PV vs new HLT PV. T-Rec Meeting, CERN, December 19, 2005.

[8] OLIVIER CALLOT. A new implementation of the Relations: The Linker objects. LHCb-2004-007, 2005.

[9] MARIUSZ WITEK. Improved vertex finding for HLT. T-Rec Meeting, CERN, December 12, 2005.

[10] LUIS FERNÁNDEZ. New Velo-TT tracking implementation. T-Rec Meeting, CERN, June 13, 2005.

[11] ALESSIA SATTA. Muon identification in the LHCb High Level Trigger. LHCb-2005-071, 2005.