LHCb/2004–046 TRIG
LPHE/2004–05
July 13, 2004

B Tracks in Signal Events Rejected by the Level-1 Trigger

Federica Legger* and **Thomas Schietinger†**

Laboratory for High-Energy Physics (LPHE)
Swiss Federal Institute of Technology Lausanne (EPFL)

Abstract

We present a systematic analysis of signal B events rejected by the Level-1 trigger. Various causes for B signal tracks to be suppressed along the decision algorithm are identified and their respective frequencies are determined for some representative signal channels. The analysis is based on the trigger algorithms and data used for the TDR performance studies of 2003.

*Federica.Legger@epfl.ch
†Thomas.Schietinger@cern.ch

Contents

1 Introduction

The trigger strategy of Level-1 (L1) exploits the features of typical b decays, i.e. high transverse momentum and large impact parameter, due to the large mass and long lifetime of b hadrons. The offline selection exploits the same b signatures, but with more stringent requirements. The L1 decision algorithm was finalized and tuned on offline-selected signal events for the LHCb Trigger System [1] and Reoptimization [2] TDRs. The emphasis during the preparation of the TDRs was on providing a robust trigger algorithm with maximized signal efficiency, with little time left for a detailed understanding of why exactly some signal events are lost and why some minimum-bias events are accepted.

The large data samples generated under homogeneous conditions for the TDR studies represent an excellent opportunity for such a thorough evaluation of our trigger strategy. In this note, we present an in-debth study of inefficiencies in the L1 trigger, i.e. we try to find out why exactly a given "good" signal event, i.e. one that proved to be useful for offline analysis, did not set-off the L1 trigger. To this end, we identify the daughters of the signal B in the Monte Carlo (MC) truth banks and try to reconstruct the fate of each daughter in the L1 software chain in order to determine the reason for that daughter's not triggering. Possible reasons are the complete absence of a track (failed reconstruction), the elimination of the track by one of the selection cuts, or an insufficient contribution of the track to the over-all trigger variable.

In Sec. 2 we briefly recall the main features of the L1 trigger algorithm. The data samples and software version used for this study are described in Sec. 3. Sections 4 and 5 contain the main results of our analyses, i.e. the classification of tracks and events according to the types of trigger failures. In Sec. 6 we present a supplementary study on the origin of decision tracks in untriggered signal samples, before concluding in Sec. 7.

2 The Level-1 trigger

2.1 L1 data

L1 is a software trigger based on limited detector information, i.e. only data coming from the silicon vertex detector (Vertex Locator, VELO) and from the first silicon tracking stations (Trigger Tracker, TT) are used [1]. Level-0 (L0) data (muon track segments and calorimeter clusters with high transverse momentum) are also made available to L1 via the L0 decision unit. VELO data are used for tracking and vertex finding to estimate the track impact parameter [3], while extrapolation of VELO tracks to TT [4, 5] and L0 objects [6] are used to measure the momentum of the tracks.

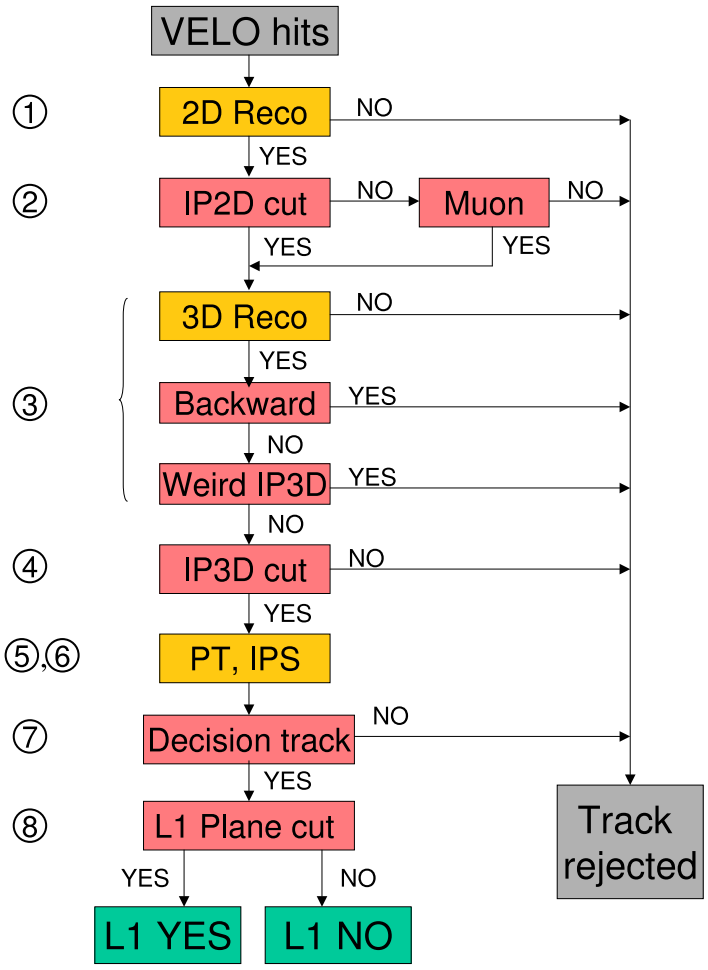

Figure 1: Schematic view of what can possibly happen to a track through the L1 generic algorithm flow (the numbers on the left represent the algorithm steps as presented in section 2.2).

2.2 L1 trigger strategy

The L1 trigger strategy is fully described in [7]. We give here a shorter description necessary to understand the study of L1 inefficiencies. The L1 algorithm consists of two parts: the first one, which is a generic selection of all b decays, requires two tracks with high transverse momentum (p_T) and large impact parameter (IP), while the second one is tailored to enhance the efficiency for some specific decay channels (with dimuons, photons or electrons in the final state). This latter part, also called *bonus system*, is mainly based on L0 information. It is not taken into account in our analysis and therefore will not be discussed further. The flow of the generic L1 algorithm (see also Fig. 1) can be summarized in the following steps:

1. 2D tracks reconstruction (in the $r - z$ plane) and primary vertex finding,

starting from VELO hits. A first estimate of the impact parameter (IP2D) is given.

2. Tracks with an IP2D value between 0.15 and 3 mm (IP2D cut), or that have been matched to a L0 muon candidate, are selected and fully reconstructed in three dimensions.

3. The impact parameter is recalculated (IP3D). Removal of backward tracks and tracks with so-called 'weird IP' (larger than 5 mm or smaller than –2 mm).[1]

4. Reconfirmation of the L0 muon match (later use for p_T assignment and dimuon bonus) and impact parameter cut (0.15 mm < IP3D < 3 mm. Tracks that do not pass the IP3D cut are rejected.[2].

5. Extrapolation of VELO 3D tracks to TT to measure the p_T. For confirmed L0 muons the p_T is evaluated using the L0 muon candidate. If no p_T measurement is available, a default transverse momentum of 400 MeV/c is assigned to the track.

6. The error $\sigma_{\rm IP}$ on the IP measurement is evaluated by means of a simple parametrization, and each track is assigned an impact parameter significance IPS = IP/$\sigma_{\rm IP}$.

7. The two leading (highest p_T) tracks, labelled 1 and 2 in the following, are sorted out and used to take the L1 decision. Events with less than two 3D tracks at this level are rejected.

8. The trigger decision is based on a simple 2-dimensional cut (also referred to as 'L1 plane cut') in the plane spawned by the two variables log PT1 + log PT2 and log IPS1 + log IPS2 as explained in Fig. 2. The trigger decision is positive if $\Delta > -0.459658$, which corresponds to a 4% retention of L0-triggered minimum bias events, or an L1 output rate of 40 kHz.

As mentioned above, for this study we will not apply any of the specific algorithms, i.e. no bonus due to specific channels features (high photon and electron E_T, dimuon mass) is added to the trigger variable Δ.

3 Data samples and software

With the objective of a comprehensive analysis of L1 behaviour on signal events, some representative channels have been chosen for hadronic, muon and electromagnetic decays, as shown in Table 1. The analyzed data samples consist of events that pass L0 and the offline selection but do not pass the L1 trigger (L0&&!L1&&SEL).

[1]The 'weird IP' cut is only relevant for tracks matched to L0 muons; its purpose is to avoid muons that do not arise from the interaction region (beam halo).

[2]These tracks can still be used later, if the muon match is confirmed, for the dimuon mass bonus [7], which is beyond the scope of this study.

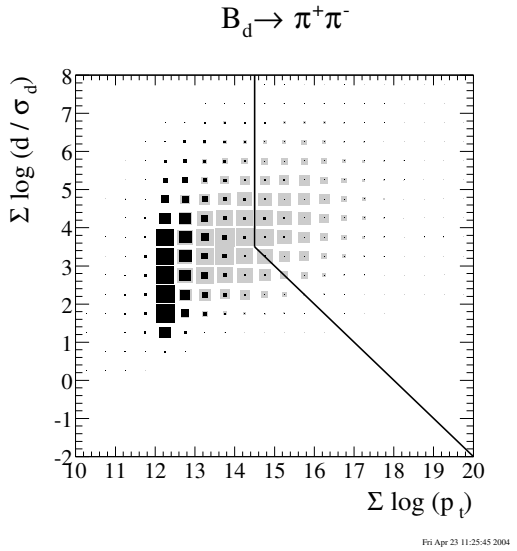

Figure 2: Distribution of minimum-bias events (in black) and $B_d \to \pi^+\pi^-$ signal events (in grey) in the plane spawned by log PT1 + log PT2 and log IPS1 + log IPS2. The signed normal two-dimensional distance of an event to the shown vertical-diagonal discriminant line is the trigger variable Δ. The sign of Δ is positive if the event lies to the right of the discriminating line.

The L1 algorithm used for this study is the baseline version as described in the Trigger System TDR [1], without the bonus system. The offline selections also correspond to the TDR analyses [8]. Since the events under study have been selected by the offline analysis, the B tracks are known to be completely reconstructible, so in principle they have no reason not to trigger. The idea is then to follow the B daughters along the L1 algorithm flow, to see where and why they are either lost or found to contribute insufficiently to the trigger decision. In order to limit the proliferation of percentages and categories (all combinations of final particles and possible trigger failures), we do not show the results relative to all the B daughters, but restrain ourselves to one B daughter per channel, where we try to pick the one which on average is most likely to trigger (*hard* refers to high-p_T):

$B_d \to \pi\pi$**:** the harder of the two pions;

$B_s \to D_s K$**:** the bachelor kaon;

$B_d \to \phi K_S$**:** the harder of the two kaons from the ϕ;

$B_s \to J/\psi(\mu\mu)\phi$**:** the harder of the two muons;

$B_s \to J/\psi(ee)\phi$**:** the harder among the electron-positron pair;

$B_d \to K^*\gamma$**:** the harder among the two K^* daughters (pion or kaon).

Table 1: Data samples used for this study: the number of L0 accepted events (L0), L1 accepted (L1), offline selected (SEL), the L1 efficiency for the selected channel (ϵ_{L1}), and the number of offline selected events which are rejected by L1 (L0&&!L1&&SEL) are shown.

Channel	**L0**	**L1**	**SEL**	**L0&&L1 &&SEL** (ϵ_{L1})		**L0&&!L1 &&SEL**
$B_d \to \pi\pi$	106260	36303	7959	5609	(70.5%)	2350
$B_s \to D_s K$	99057	34719	2447	1745	(71.3%)	702
$B_d \to \phi K_S$	96103	22540	2291	1159	(50.6%)	1132
$B_s \to J/\psi(\mu\mu)\phi$	181864	75465	19511	12999	(66.6%)	6512
$B_s \to J/\psi(ee)\phi$	84786	26067	3953	2116	(53.5%)	1837
$B_d \to K^*\gamma$	59137	12857	1300	587	(45.1%)	713

4 Non-triggering B tracks in Level-1

4.1 Main types of trigger failures

Based on the structure of the L1 generic algorithm, we distinguish three main stages where a signal track can be lost or make an insufficient contribution to the trigger variable:

1. **lost at 2D:**
   - the 2D reconstruction fails, or
   - the 2D track does neither pass the IP2D cut nor is it matched to an L0 muon candidate.
2. **lost at 3D:**
   - the 3D reconstruction fails, or
   - the 3D track does not pass the IP3D cut.
3. **low p_T:**
   the 3D track does not become a decision track because its measured p_T is too small.
4. **small Δ:**
   the 3D track is a decision track, but the calculated Δ variable of the event is too small to pass the L1 plane cut.

Our method consists in identifying the signal decay tree for every event using MC truth. Each MC particle identified as a signal B daughter is then followed through the L1 algorithm flow via association to 2D and 3D tracks, and finally classified according to where it is lost or found to make an insufficient contribution to the trigger variable. The results of this coarse classification into the main categories are shown in Table 2. For each channel, the B daughter which is a priori more likely to trigger is analyzed (as explained in Sec. 3).

Table 2: Classification of B tracks according to the main types of L1 trigger failures.

Channel	**lost at 2D**		**lost at 3D**		**low p_T**		**small Δ**	
$B_d \to \pi\pi$	1347	(57.3%)	317	(13.5%)	374	(15.9%)	312	(13.3%)
$B_s \to D_s K$	354	(50.4%)	71	(10.1%)	119	(16.9%)	158	(22.5%)
$B_d \to \phi K_S$	686	(60.6%)	144	(12.7%)	115	(10.2%)	187	(16.5%)
$B_s \to J/\psi(\mu\mu)\phi$	818	(12.6%)	4366	(67.0%)	173	(2.7%)	1155	(17.7%)
$B_s \to J/\psi(ee)\phi$	1262	(68.7%)	214	(11.7%)	144	(7.8%)	217	(11.8%)
$B_d \to K^*\gamma$	432	(60.6%)	58	(8.1%)	58	(8.1%)	165	(23.2%)

4.2 Detailed breakdown

4.2.1 Tracks lost at 2D

As stated above, a track can be lost at 2D either because of failed 2D reconstruction, or because it does not pass the IP2D cut (and is not matched to L0 muon candidate at the same time). The latter category may be split in two according to whether the rejection of the track based on its IP is justified or not: By using MC truth, we determine whether the true IP is inside or outside the IP selection window to further classify the track. Results are shown in Table 3. The fraction of (leading) tracks that pass this level since they are matched to an L0 muon candidate is of the order of 1% for all channels but $B_s \to J/\psi(\mu\mu)\phi$ where this fractions amounts to about 63%.

4.2.2 Tracks lost at 3D

Similarly, a track can be lost at 3D either because of failed 3D reconstruction, or because it does not pass the IP3D cut.Again, this last category can be split in two using the information from MC truth, i.e. a track is lost either correctly (true IP outside the selection window) or incorrectly because of bad reconstruction (true IP inside the selection window). The results are shown in Table 4. The fraction of (leading) tracks that arrive at this level since they are matched to an L0 muon candidate is of the order of 1% for all channels but $B_s \to J/\psi(\mu\mu)\phi$ where this fractions amounts to about 60%.

4.2.3 Tracks with low p_T

In this category we find B tracks that do not become decision tracks because their p_T is not high enough, i.e. there are at least two tracks in the event with higher p_T. We may apply the same classification as in the case of the IP selections above, i.e. failed reconstruction, correct rejection and wrong rejection: In the first case, the reconstruction of the track in TT fails and no p_T estimate is available. The track is assigned the default p_T of 400 MeV/c. If the track has a measured p_T, we check using MC truth whether or not the track would have become a decision track with

Table 3: Detailed breakdown of signal B tracks lost at 2D. The percentages given in the first and second column after the track numbers represent absolute fractions (normalized to all leading signal B tracks for that channel) and relative fractions (normalized to the tracks that enter this table), respectively. The label 'correct' ('wrong') indicates that the algorithm would (not) have rejected the track based on MC truth instead of reconstructed information.

Channel	**Failed 2D Reco**			**Bad IP2D, no μ (correct)**			**Bad IP2D, no μ (wrong)**		
$B_d \to \pi\pi$	79	3.4%	5.7%	573	24.4%	42.5%	695	29.6%	51.6%
$B_s \to D_s K$	14	2.0%	4.0%	216	30.8%	61.0%	124	17.7%	35.0%
$B_d \to \phi K_S$	98	8.7%	14.3%	382	33.7%	55.7%	206	18.2%	30.0%
$B_s \to J/\psi(\mu\mu)\phi$	169	2.6%	20.7%	437	6.7%	53.4%	212	3.3%	25.9%
$B_s \to J/\psi(ee)\phi$	62	3.4%	4.9%	936	50.9%	74.2%	264	14.4%	20.9%
$B_d \to K^*\gamma$	19	2.7%	4.4%	269	37.7%	62.3%	144	20.2%	33.3%

Table 4: Detailed breakdown of signal B tracks lost at 3D. The percentages given in the first and second column after the track numbers represent absolute fractions (normalized to all leading signal B tracks for that channel) and relative fractions (normalized to the tracks that enter this table), respectively. The label 'correct' ('wrong') indicates that the algorithm would (not) have rejected the track based on MC truth instead of reconstructed information.

Channel	**Failed 3D Reco**			**Bad IP3D (correct)**			**Bad IP3D (wrong)**		
$B_d \to \pi\pi$	290	12.3%	91.4%	16	0.7%	5.0%	11	0.5%	3.5%
$B_s \to D_s K$	55	7.8%	77.5%	14	2.0%	19.7%	2	0.3%	2.8%
$B_d \to \phi K_S$	114	10.1%	79.2%	16	1.4%	11.1%	14	1.3%	9.7%
$B_s \to J/\psi(\mu\mu)\phi$	426	6.5%	9.7%	3123	48.0%	71.5%	817	12.5%	18.7%
$B_s \to J/\psi(ee)\phi$	178	9.7%	83.2%	27	1.5%	12.6%	9	0.5%	4.2%
$B_d \to K^*\gamma$	38	5.3%	65.5%	17	2.4%	29.3%	3	0.4%	5.2%

perfect momentum measurement and classify it accordingly. Results are shown in Table 5.

4.2.4 Decision tracks

The 3D track originating from the signal B is one of the two decision tracks (PT1 or PT2), i.e. it is used to compute the L1 decision variable Δ, but the event is rejected anyway because Δ turns out to be too small. In this case, we can apply MC truth to check whether the event would have triggered if the track in question had been measured with infinite precision, in other words if or not the rejection of the event is justified.[3] Results are shown in Table 6.

4.3 Discussion

Figure 3 gives a graphical summary of the results averaged over the signal channels except $B_s \to J/\psi(\mu\mu)\phi$, which is special because of the treatment of muons in the L1 algorithm. We will discuss this channel separately at the end of this subsection. For all other channels we can draw the following conclusions:

- The first IP selection in 2D (IP2D) turns out to be the most significant killer of signal B tracks: roughly 50% of the tracks under investigation are lost at this level. The loss is even more dramatic for channels containing a J/ψ (around 65%), since the signal selections for these channels do not apply a strong IP cut because of the strong signature given by the J/ψ. The MC truth check shows that the tracks are correctly rejected in about two thirds of the cases on average, except for $B_d \to \pi\pi$, where the rejection is more often wrong than correct (29.6% vs. 24.4%).

- The fraction of signal B tracks lost due to track reconstruction failure at 2D is not significant. The worst channel in this respect is $B_d \to \phi K_S$ with an absolute loss of 8.7% of all tracks under consideration. The losses due to track reconstruction failure at 3D are somewhat more significant, reaching up to 12.3% of the tracks in the case of $B_d \to \pi\pi$. It is the dominant reason for trigger failure at the 3D level.

- The effect of the IP selection in 3D is small, since for most tracks it only confirms the selection already applied at 2D. An exception are those tracks that have been matched to L0 muons in 2D and therefore evaded the IP cut in 2D. About half of the signal B tracks cut by the IP3D requirement are of that type.

- As for signal B tracks that are not used in the L1 decision because their p_T is too low, we see that this category (the low-p_T category) is completely dominated by tracks carrying the default momentum of 400 MeV/c, i.e. tracks

[3] Strictly speaking, one should also verify the correctness of all the other tracks involved in the decision before deciding whether the rejection of an event was justified or not.

for which the extrapolation to TT failed. These tracks account for 7–16% of all non-triggering signal tracks, depending on the channel.

- Signal B tracks which are actually used by L1 as decision tracks but still do not trigger cover about 15% of the cases. Most of them ($\approx$ 90%) are leading tracks (PT1), and in most cases, they would not have triggered the event even with perfectly measured p_T. These tracks therefore represent an irreducible source of trigger inefficiency (unless the over-all trigger threshold can be lowered thanks to improvements elsewhere in the trigger algorithm).

  It is interesting to note that a perfect p_T measurement for the sub-leading track (PT2) would save more events from rejection than it would for the leading track (PT1). This is because the majority of the PT2 tracks have unmeasured p_T and have therefore been assigned the default 400 MeV/c. In principle, PT1 tracks may also be of this type, but this case is extremely rare among signal events.

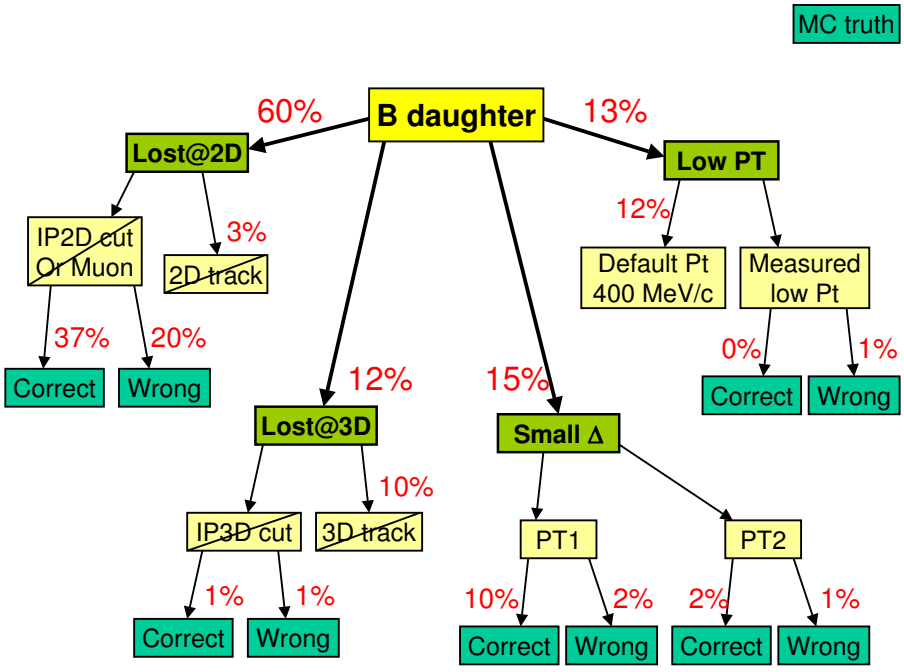

Figure 3: Summary of trigger inefficiencies (see text). The numbers are averages over all considered channels except $B_s \to J/\psi(\mu\mu)\phi$.

Adding up all categories, we may determine the percentage of signal tracks that do not trigger the event because of trigger design (the subcategories marked 'correct'), and compare them to the tracks that do not trigger the event because of resolution problems, bad measurements and reconstruction failures (all the others), see Table 7. For hadronic channels the two classes of tracks are comparable in size, whereas for electromagnetic and muon channels the trigger inefficiencies are mainly

the result of the trigger design. Clearly the design of the generic L1 algorithm is not optimal for these channel and we need the specific algorithms (bonus system) to boost the L1 efficiencies for these decays. The channel $B_d \to \pi\pi$ suffers the most from resolution effects. This may be traced back to the substantial loss of 2D tracks due to bad IP measurements for this channel.

4.3.1 The muon channel

The main difference between the muon and the other channels is the treatment it undergoes after the 2D reconstruction. In fact, 2D tracks which have been matched to an L0 muon (63%) are reconstructed in 3D even if they do not pass the IP2D cut (73%). Therefore only 10% of the tracks are lost at the 2D level. After the 3D reconstruction, however, most of these tracks do not pass the IP3D cut (61%). Some 99% of the tracks cut by the IP3D requirement are muon candidates. The MC truth check reveals that 48% of the tracks lost at the 3D level are correctly rejected, and only about 13% are lost in error. If we add to these percentages the ones from the 2D level, we arrive at failure rates that are comparable to those obtained for the other channels at the 2D level only.

We emphasize that the muons rejected by the IP3D selection cut are still considered fot the evaluation of the dimuon mass and may therefore still trigger the event via the bonus system (which is beyond the scope of this study).

Table 5: Detailed breakdown of signal B tracks with too low p_T for being decision tracks. The percentages given in the first and second column after the track numbers represent absolute fractions (normalized to all leading signal B tracks for that channel) and relative fractions (normalized to the tracks that enter this table), respectively. The label 'correct' ('wrong') indicates that the algorithm would (not) have rejected the track based on MC truth instead of reconstructed information.

Channel	No PT (default 400 Mev/c)			Low PT (correct)			Low PT (wrong)		
$B_d \to \pi\pi$	373	15.9%	99.7%	0	0.%	0.%	1	0.%	0.2%
$B_s \to D_sK$	115	16.4%	96.6%	0	0.%	0.%	4	0.6%	3.4%
$B_d \to \phi K_S$	104	9.2%	90.4%	3	0.3%	2.6%	8	0.7%	7.0%
$B_s \to J/\psi(\mu\mu)\phi$	170	2.6%	98.3%	0	0.%	0.%	3	0.%	1.7%
$B_s \to J/\psi(ee)\phi$	133	7.2%	92.4%	0	0.%	0.%	11	0.6%	7.6%
$B_d \to K^*\gamma$	54	7.6%	93.1%	1	0.1%	1.7%	3	0.4%	5.2%

Table 6: Detailed breakdown of signal B tracks that are one of the two decision tracks (PT1 or PT2) of the event, but where the event is rejected anyway. The percentages given in the first and second column after the track numbers represent absolute fractions (normalized to all leading signal B tracks for that channel) and relative fractions (normalized to the tracks that enter this table), respectively. The label 'correct' ('wrong') indicates that the algorithm would (not) have rejected the track based on MC truth instead of reconstructed information.

Ch.	PT1 correct			PT1 wrong			PT2 correct			PT2 wrong		
$\pi\pi$	146	6.2%	46.8%	115	4.9%	36.9%	16	0.7%	5.1%	35	1.5%	11.2%
D_sK	118	16.8%	74.7%	21	3.0%	13.3%	10	1.4%	6.3%	9	1.3%	5.7%
ϕK_S	136	12.0%	72.7%	13	1.1%	6.8%	23	2.0%	12.3%	15	1.3%	8.0%
$(\mu\mu)\phi$	1049	16.1%	90.8%	72	1.1%	6.2%	19	0.3%	1.6%	15	0.2%	1.3%
$(ee)\phi$	126	6.9%	58.1%	65	3.5%	29.9%	10	0.5%	4.6%	16	0.9%	7.4%
$K^*\gamma$	135	18.9%	81.8%	12	1.7%	7.3%	11	1.5%	6.7%	7	1.0%	4.2%

Table 7: Number of events lost for bad resolution or for trigger design.

Channel	L1 Design		Resolution	
$B_d \to \pi\pi$	751	32.0%	1599	68.0%
$B_s \to D_sK$	358	51.0%	344	49.0%
$B_d \to \phi K_S$	560	49.5%	572	50.5%
$B_s \to J/\psi(\mu\mu)\phi$	4628	71.1%	1884	28.9%
$B_s \to J/\psi(ee)\phi$	1099	59.8%	738	40.2%
$B_d \to K^*\gamma$	433	60.7%	280	39.3%

5 Trigger inefficiency at the event level: the example $B_d \to \pi\pi$

For simplicity, we so far have analyzed trigger inefficiencies at the track level, i.e. for each signal B track, we have determined the reason why that particular track has not been able to trigger the event. It is straight-forward to extend our analysis to the event level by combining the n track failure categories into n^j event failure categories where j is the number of signal tracks in the decay tree that are supposed to trigger. Since such a procedure results in a proliferation of numbers and percentages with questionable benefit in terms of interpretation, we give a simplified breakdown for one example only, the decay $B_d \to \pi\pi$.

We regroup the possible trigger failures in the following categories:

- **NO2D**: Failed 2D track reconstruction.
- **NO3D**: Failed 3D track reconstruction.
- **IP2D**: The track does not pass the IP2D cut and is not a muon.
- **PT1**: The track is the leading decision track but Δ is too small.
- **PT2**: The track is the sub-leading decision track but Δ is too small.
- **LOWPT**: The track is neither of the two decision tracks because its p_T is too small.
- **IP3D**: The track does not pass the IP3D cut.

The result of this classification is shown in Fig. 4. It is confirmed that the IP2D cut is the main reason for trigger failure, even for events where the other B daughter provides the leading decision track.

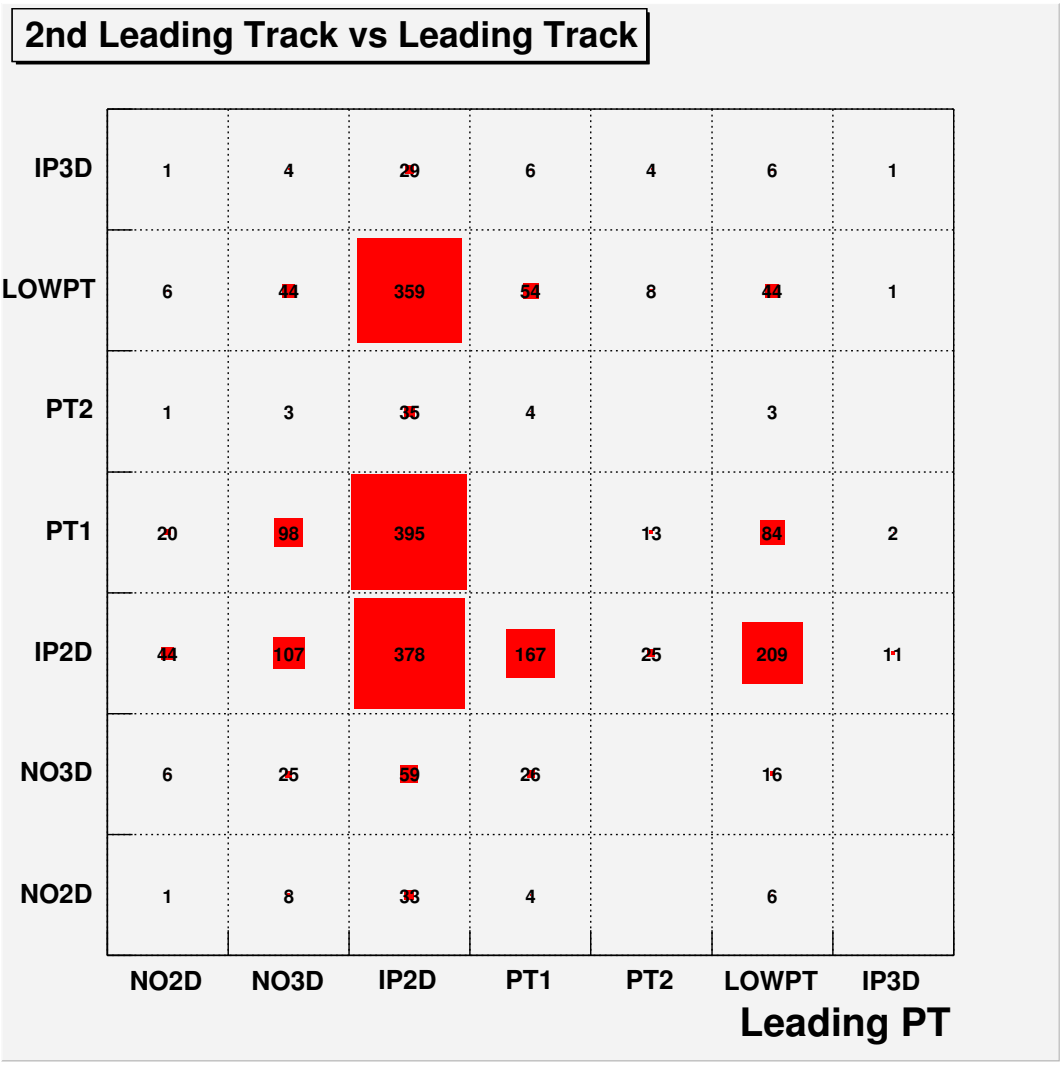

Figure 4: Distribution of untriggered signal events according to trigger failures for the two pions in $B_d \to \pi^+\pi^-$. A description of the categories is given in the text. The categories itemized horizontally refer to the leading pion (higher p_T), the categories on the vertical axis refer to the other pion. For example, there are 359 events in the untriggered signal event sample where the leading pion did not pass the IP selection in two dimensions (IP2D) and the second pion did not become a decision track because of too low or unmeasured p_T (LOWPT).

6 Origin of the decision tracks in the untriggered sample

In a supplementary study we have investigated the origin of the decision tracks in the untriggered event samples. In particular, we are interested in the fraction of B tracks and signal B tracks among the decision tracks in these samples. Table 8 summarizes the B content among the decision tracks for each of the samples under study. We see that even in the untriggered samples, in between 65% ($B_d \to \phi K_S$) and 84% ($B_s \to D_s K$) of the events at least one of the decision tracks is a B track.

In Tables 9 and 10) we split the leading and sub-leading decision tracks, respectively, into samples according to their arising from the signal B, the other B or from somewhere else. We observe that the leading decision track has a large probability of being a signal B track (even in the untriggered sample). In events of the types $B_d \to \phi K_S$ and $B_d \to K^* \gamma$ the decision tracks are more likely to come from the other B than in the other channels, reflecting the relatively weak signatures of these two channels (relatively low p_T of the ϕ and K^* daughters, decay of K_S outside of the VELO).

Table 8: Breakdown of untriggered signal events where: the two decision tracks come both from a B (both B), one comes from a B and the other does not (one B), none of them is a B child (no B), and at least one of the two leading tracks has no association to a MC particle (No MC asso).

channel	**both B**		**one B**		**no B**		**no MC**	
$B_d \to \pi\pi$	520	(18.0%)	1471	(50.8%)	849	(29.3%)	56	(1.9%)
$B_s \to D_s K$	292	(33.4%)	442	(50.6%)	136	(15.6%)	3	(0.4%)
$B_d \to \phi K_S$	381	(29.4%)	461	(35.6%)	426	(32.9%)	27	(2.1%)
$B_s \to J/\psi(\mu\mu)\phi$	2755	(34.8%)	3011	(38.0%)	1999	(25.2%)	153	(1.7%)
$B_s \to J/\psi(ee)\phi$	675	(31.7%)	860	(40.4%)	559	(26.3%)	34	(1.6%)
$B_d \to K^*\gamma$	237	(29.1%)	303	(37.2%)	250	(30.7%)	24	(3.0%)

Table 9: Origin of the leading track in untriggered signal events. It may come either from the signal B, the other B, is not a B child, or no association to the MC particle has been found.

channel	**signal B**		**other B**		**no B**		**no MC**	
$B_d \to \pi\pi$	1409	(48.7%)	412	(14.2%)	1052	(36.3%)	23	(0.8%)
$B_s \to D_s K$	562	(64.4%)	111	(12.7%)	199	(22.8%)	1	(0.1%)
$B_d \to \phi K_S$	479	(37.0%)	274	(21.2%)	531	(41.0%)	11	(0.8%)
$B_s \to J/\psi(\mu\mu)\phi$	4430	(55.9%)	922	(11.6%)	2509	(31.7%)	57	(0.7%)
$B_s \to J/\psi(ee)\phi$	1007	(47.3%)	379	(17.8%)	734	(34.5%)	8	(0.4%)
$B_d \to K^*\gamma$	338	(41.5%)	158	(19.4%)	306	(37.6%)	12	(1.4%)

Table 10: Origin of the second leading track in untriggered signal events. It may come either from the signal B, the other B, is not a B child, or no association to the MC particle has been found.

channel	**signal B**		**other B**		**no B**		**no MC**	
$B_d \to \pi\pi$	120	(4.1%)	570	(19.7%)	2099	(72.5%)	107	(3.7%)
$B_s \to D_s K$	191	(21.9%)	162	(18.6%)	511	(58.5%)	9	(1.0%)
$B_d \to \phi K_S$	223	(17.2%)	247	(19.1%)	792	(61.2%)	33	(2.5%)
$B_s \to J/\psi(\mu\mu)\phi$	1878	(23.7%)	1291	(16.3%)	4524	(57.1%)	225	(2.8%)
$B_s \to J/\psi(ee)\phi$	431	(20.3%)	393	(18.5%)	1247	(58.6%)	57	(2.7%)
$B_d \to K^*\gamma$	129	(15.8%)	152	(18.7%)	497	(61.1%)	36	(4.4%)

7 Conclusion and outlook

We summarize the main conclusions of our systematic analysis of trigger inefficiencies at L1 for offline-selected signal events (see also Fig. 3):

- There are no trivial losses of signal tracks due to coding errors or obvious shortcomings of the trigger strategy. All types of inefficiencies encountered in this study were expected beforehand (but their quantitative contributions unknown).

- By far the largest suppression of B signal tracks is due to the IP cut applied after the track and primary vertex reconstruction in two dimensions. More than half of all signal tracks are lost at this stage. In the majority of cases, the rejection is justified based on the IP selection window chosen for L1.

- The second largest source of inefficiency is due to insufficient p_T: the signal track either does not become one of the two decision tracks or it does not contribute enough to the trigger variable to pass the trigger threshold. In the former case the problem is almost always caused by reconstruction failure in TT. The latter case is in fact dominated by tracks with truly low p_T, such that it must be considered an irreducible source of inefficiency, accounting for about 10% of all non-triggering signal tracks.

Given that the main reason for trigger inefficiency turns out to be the IP selection, we conclude that the most crucial tunable parameter is the required minimum IP at the 2D reconstruction stage. A concurrent study on minimum-bias triggers at L1 has come to a similar conclusion [9], i.e. most of the minimum-bias triggers at L1 are the result of a bad IP determination. Perhaps a more flexible cut on IP could further improve the trigger performance.

More effort should also go into investigating the reasons for failed reconstruction in TT, and into finding the optimal way of handling such tracks (assigning a default p_T of 400 MeV/c may not be optimal).

Many of these issues are already being addressed within the new trigger software framework developed for the 2004 data challenge (DC'04). Within the same framework, we plan to automate the classification of signal tracks according to their behaviour in the trigger chain as presented here and make it available in the form of a software tool. With this tool it will be easy to monitor trigger inefficiencies simply by switching on the corresponding option in the L1 trigger code.

On a longer timescale, we should think about developing methods to study the trigger performance without relying on Monte-Carlo truth, i.e. in a situation as we will face it when real data taking begins in 2007 or 2008.

References

[1] R. Antunes Nobrega *et al.* (LHCb collaboration), *LHCb Trigger System Technical Design Report*, CERN/LHCC 2003–031, LHCb TDR 10, September 2003.

[2] R. Antunes Nobrega *et al.* (LHCb collaboration), *LHCb Reoptimized Detector (Design and Performance) Technical Design Report*, CERN/LHCC 2003–030, LHCb TDR 9, September 2003.

[3] V. Lindenstruth *et al.*, *The LHCb Level-1 Trigger: Architecture, Prototype, Simulation and Algorithm*, LHCb/2003–064, TRIG (Chapter 5).

[4] H. Dijkstra *et al.*, *The use of the TT1 tracking station in the Level-1 trigger*, LHCb/2002–045 TRIG.

[5] M. Witek, *VELO–TT matching and momentum determination at Level-1 trigger*, LHCb/2003–060 TRIG.

[6] N. Tuning, *Matching VELO tracks to L0 objects for L1*, LHCb/2003–039, TRIG.

[7] C. Jacoby and T. Schietinger, *Level-1 decision algorithm and bandwidth division*, LHCb/2003–111 TRIG.

[8] `http://lhcb-phys.web.cern.ch/lhcb-phys/task_force/TDR_studies_2002-2003/event_lists_TDR_Dv8r2/`

[9] L. Fernández and T. Schietinger, *Decision Tracks in Minimum-Bias Events Selected by the Level-1 Trigger*, LHCb/2004–047 TRIG.