The anatomy of an LHCb event in the upgrade era, and implications for the LHCb trigger

Public Note

Issue: 1
Revision: 0

Reference: LHCb-PUB-2014-027
Created: 2 May 2013
Last modified: May 21, 2014

Prepared by: Conor Fitzpatrick[a], Vladimir V. Gligorov[a]
[a]CERN,Switzerland

Abstract

This document studies the rates at which particles of interest are produced within the LHCb detector acceptance, in both Run 1 and Upgrade conditions. We present the event rates that could be selected by an idealised trigger system during the LHCb Upgrade, and compare these to the rates selected by the Run 1 trigger system. We discuss the implications of our findings for the proposed LHCb Upgrade trigger.

1 Introduction

After the second Long Shutdown, the LHC is expected to operate at the design energy of $14\ \mathrm{TeV}$, while the LHCb detector will be upgraded in order to function up to a luminosity of $2 \cdot 10^{33}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$. This running environment will present a number of challenges for LHCb, not least at the trigger level. In this note we attempt to describe the properties of a typical post-upgrade minimum bias event and compare these properties to those of a typical event seen by the current LHCb detector. Our aim is to establish whether the current LHCb trigger strategy can be continued in the upgrade; as such, it will be useful to give a brief overview of this strategy before proceeding further.

The Run 1 LHCb trigger strategy [1, 2, 3, 4, 5] consists of three stages. The first, hardware, stage, selects events based on high transverse energy calorimeter clusters or high transverse momentum muon track segments, and is mandated by the fact that the full LHCb detector can only be read out at $1\ \mathrm{MHz}$ while in the Run 1 LHC inelastic collision rate in LHCb is $15\ \mathrm{MHz}$. The second and third stages are implemented in a software High-Level-Trigger (HLT) farm of commodity processors. In the second stage, dimuon vertices or single high transverse momentum displaced tracks are searched for. In the third stage the dimuon candidates are refined, while inclusive B and D hadron candidates are built around the displaced track. An inclusive trigger is able to keep an event based only on a portion of the final state particles, while exclusive triggers require all final state particles to be reconstructed in order for the trigger to keep the event. While the inclusive B and D^* hadron, as well as dimuon, triggers take up the majority of the current LHCb trigger bandwidth, certain hadronic charm decay modes are also triggered using exclusive triggers in order to increase their efficiency.

The baseline trigger strategy for the LHCb upgrade consists of gradually removing the hardware trigger while increasing the size of the HLT farm in order to eventually process the entire LHC collision rate. The output rate to disk of such a trigger will be around 0.35 GByte/s which at an estimated event size of 100 kB implies a rate of 10-50 kHz. It has recently been confirmed [6] that such a strategy will be feasible from the point of view of the time taken to process these events. What remains to be demonstrated is whether the output rate constraint will allow an inclusive trigger strategy similar to that of the present detector to continue in the upgrade era, and that is the subject of our study.

2 Datasets

Two simulated LHCb minimum bias datasets are used in this note, corresponding to Run 1 conditions during 2011 and those of the nominal data taking conditions expected in the upgrade. The relevant conditions are: The average number of both inelastic and elastic proton-proton collisions per event, referred to as ν, the instantaneous luminosity, $\mathcal{L}$, and the collision energy, $\sqrt{s}$. During Run 1, spillover was not included in the simulation as the detector is robust to spillover at $50\ \mathrm{ns}$ bunch spacing. In the upgrade the LHC will deliver collisions at 25 ns, and so we simulate this in our upgrade sample. Table 1 describes the conditions and naming conventions of these samples. The Run 1 dataset includes the full detector digitisation and offline reconstructed physics objects. For the upgrade dataset the full 10M event sample consists only of generator-level information. This generator level only sample is complimented by a fully simulated sample of 97k events in which the upgraded Velo Pixel hardware has been simulated.

Property	Run 1, 2011	Upgrade, nominal luminosity
$\sqrt{s}$ [TeV]	7	14
ν	2	7.6
$\mathcal{L}$ [$\times 10^{33}\mathrm{cm}^{-2}\ \mathrm{s}^{-1}$]	0.4	2
Spillover	N	Y
Number of events	200k	10M (100k)

Table 1: Generator configuration for minimum bias datasets used in this report. Parameters are described in the text. The second number for upgrade events indicates the sample of fully simulated data.

b-hadrons		c-hadrons		Light long-lived	
Particle	|PDG ID|	Particle	|PDG ID|	Particle	|PDG ID|
B^0	511	D^+	411	K^0_S	310
B^+	521	D^0	421	Λ^0	3122
B^0_s	531	D^+_s	431	Σ^+	3112
B^+_c	541	Λ^+_c	4122	Ξ^0	3322
Λ_b	5122	Ξ^0_c	4132	Ξ^-	3312
Ξ^+_b	5312	Ξ^+_c	4232	Ω^-	3334
Ξ^0_b	5322	Ω^0_c	4332		
Ω^0_b	5332				

Table 2: Classification of particles used in this report. Charge conjugation is implied throughout.

3 Particle classification

For the purpose of classification, we look for 21 parent particle types and their charge conjugates in each event. We make no immediate assumptions about their final-state topology or their suitability for physics analysis purposes, and we do not require that they come directly from a primary proton-proton interaction. This approach is justified by our experience of LHCb's first running period (2010-2012) during which we have already seen many analyses which were once considered impossible or uninteresting become very interesting indeed; therefore, whatever the efficiencies of any specific trigger configuration, we maintain that the trigger must retain the flexibility to select any decay topology efficiently relative to the offline analysis. We loosely categorise these 21 particle types based upon their quark content: b-hadrons, c-hadrons, and *light*, *long-lived* candidates. The categories are presented in Table 2. This list is not, of course, complete, but it should be representative of almost all possible topologies of interesting events.

4 Software and analysis procedure

The analysis is first performed on the HepMC [7] generator-level information. Each event contains several HepMC GenEvents corresponding to a single PV. Each of these GenEvents is traversed and when a candidate corresponding to those listed in Table 2 is found, its properties are retained for further analysis. These properties consist of:

- The position of the PV to which the particle is associated
- The particle's production vertex
- The particle's end vertex
- The particle's momentum four-vector
- The particle's decay products, their production- and end-vertices, and momenum 4-vectors

Derived quantities are also retained, including the particle's decay time, flight distance to its associated PV and impact parameter with respect to the PV. Later, we study fully simulated events, where the same properties are available and include detector resolution effects.

4.1 Definitions

In this document we refer to the particles listed in Table 2 as *candidates*. Each candidate can have several possible decay products which are defined as follows:

- Particles with an end vertex are defined as *composite*.
- We define a particle as being a *track* if it is charged and has no end vertex.

The ability to fully reconstruct a specific candidate is an important metric. For Monte-Carlo productions LHCb often chooses to apply a "Decay Products in acceptance" cut when generating Monte-Carlo samples in order to preserve disk space. This cut rejects candidates whose detectable daughter particles are not fully contained within the $10 - 400$ mrad spectrometer acceptance. For the purpose of this paper we investigate the efficiency of candidates passing a similar, but more strict definition as follows:

- Photons are not required to be within the acceptance unless they come from η, π^0 decays.
- All daughter charged tracks must be within 10-400 mrad.
- All daughter neutral tracks must be within 5-400 mrad.
- The standard acceptance cut does not require K^0_S, Λ^0 daughters to be within the acceptance. For this study we require that they are.

Additionally we study the number of candidates that have a reconstructible vertex within the LHCb Vertex Locator, the VELO. We define a track as being within the VELO acceptance if it has positive momentum in the forward direction and traverses at least three VELO stations. For this we do not take into account detection efficiency (which is 99% [8]), we only require that the track intersects a with a VELO station. A candidate is considered to have decayed in the VELO acceptance if at least two daughter tracks traverse at least three VELO stations each. The reason for this definition is that inclusive triggers are generally based on finding a displaced vertex, and two is the minimum number of tracks needed to form a vertex. For pre-upgrade events we determine whether or not a candidate meets these requirements using the current (pre-upgrade) VELO. For the upgrade events we use the nominal VELO pixel geometry. This allows us to compare directly with fully simulated, reconstructed events of both types.

5 Expected Yields

In this section we explore the expected number of candidates of each category and the effect of applying some basic requirements.

5.1 Generator-level yields

We provide here an estimate of the per-event yields of each candidate category in both Run 1 and Upgrade datasets. These yields are determined from a subset of 1M upgrade events and the full 210k of Run 1. The yields are normalised by the subset size in order to provide the per-event yield. Table 3 presents the per-event yields in Run 1 conditions for b, c, and light, long-lived hadrons respectively. Table 4 presents the same information for events generated assuming upgrade conditions. In both tables we also include the percentage of these candidates that leave two tracks in the VELO, and the percentage that meet both this VELO requirement and that have all daughter tracks fully contained within the LHCb acceptance. These tables give some indication of the relative complexity of proton-proton collisions pre- and post-upgrade. Of particular interest is the number of light, long-lived candidates: Prior to the upgrade a little over half of all events contain one such candidate with at least two VELO tracks and are fully contained within the LHCb acceptance. Post upgrade this increases to *more than two per event*. We may infer from this that any trigger solely dependent upon VELO track information would be capable of rejecting half of all events prior to the upgrade, but would have no selectivity at $2 \times 10^{33}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$. After the upgrade we can expect roughly a factor of five increase in charm and a factor of six increase in beauty hadrons within the LHCb acceptance and which leave tracks in the pixel velo with respect to pre-LS1 running conditions.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
b-hadrons	0.0258 ± 0.0004	$30.5 \pm 0.6\%$	$11.1 \pm 0.4\%$
c-hadrons	0.297 ± 0.001	$21.9 \pm 0.2\%$	$14.2 \pm 0.1\%$
light, long-lived hadrons	8.04 ± 0.01	$6.67 \pm 0.02\%$	$6.35 \pm 0.02\%$

Table 3: Candidates per event and efficiencies for generator-level events in MC11a. A breakdown of each category is available in Tables 7-9. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the current VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
b-hadrons	0.1572 ± 0.0004	$34.9 \pm 0.1\%$	$11.9 \pm 0.1\%$
c-hadrons	1.422 ± 0.001	$24.73 \pm 0.04\%$	$15.12 \pm 0.03\%$
light, long-lived hadrons	33.291 ± 0.006	$7.022 \pm 0.004\%$	$6.257 \pm 0.004\%$

Table 4: Candidates per event and efficiencies of generator-level events after the upgrade. A breakdown of each category is available in Tables 10-12. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the upgrade VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

6 Reconstructed yields

This section presents the per-event yields and efficiencies for candidates that have been fully simulated and reconstructed within the LHCb simulation framework. Candidates are partially reconstructed by forming a vertex from two charged tracks which are determined to have come from genuine $\pi, \mu, \mathrm{K}, \mathrm{e}, \mathrm{p}$ particles. The vertex is then matched to a composite particle which is either one of those listed in Table 2 or a daughter of one. No additional cuts are applied. Table 5 indicates the number of these candidates per-event for each type in the present detector configuration. Current analysis experience shows that, broadly speaking, candidates with a p_{T} above 2 GeV/c and a decay-time above 0.2 ps have a purity acceptable for offline analyses. For this reason, we show the efficiency of these two cuts when applied on the patrially reconstructed but fully reconstructible final state offline. By this we mean that while we only reconstruct a two track vertex we ensure that all other daughter tracks are reconstructible. Any fully reconstructible candidates passing these cuts are considered to be potentially interesting for further offline analysis and hence define the minimum which an ideal trigger ought to be selecting. Table 6 provides the same information for the available 100k of fully simulated post-upgrade minimum bias events.

The last row in these tables presents the output rate of an ideal trigger, which selects signal with 100% efficiency and purity in these scenarios. Figures 1 shows how the rates of fully reconstructible signal candidates varies as a function of p_{T} and decay time cuts in the upgrade scenario.

In Run 1 the rate of b hadrons with substantial p_{T} and lifetime which could be reconstructed inside the detector was already very substantial, although this issue was partly mitigated by the comparatively

	b-hadrons	c-hadrons	light, long-lived hadrons
Reconstructed yield	$(4.0 \pm 0.1) \times 10^{-3}$	0.0196 ± 0.0003	0.0792 ± 0.0006
$\epsilon(p_{\mathrm{T}} > 2\mathrm{GeV/c})$	$83 \pm 1\%$	$47.2 \pm 0.8\%$	$2.0 \pm 0.1\%$
$\epsilon(\tau > 0.2\ \mathrm{ps})$	$89 \pm 1\%$	$64.2 \pm 0.7\%$	$99.53 \pm 0.05\%$
$\epsilon(p_{\mathrm{T}}) \times \epsilon(\tau)$	$73 \pm 2\%$	$30.2 \pm 0.7\%$	$1.9 \pm 0.1\%$
$\epsilon(p_{\mathrm{T}}) \times \epsilon(\tau) \times \epsilon(\mathrm{LHCb})$	$29 \pm 1\%$	$22.3 \pm 0.6\%$	$1.9 \pm 0.1\%$
Output rate	17.3 kHz	66.9 kHz	22.8 kHz

Table 5: Per-event yields determined from 0.21M of Run 1 minimum-bias events after partial offline reconstruction. The first row indicates the number of candidates which had at least two tracks from which a vertex could be produced. The last row shows the output rate of a trigger selecting such events with perfect efficiency, assuming an input rate of 15 MHz from the LHC, as during 2012 running. A breakdown of each category is available in Table 13.

	b-hadrons	c-hadrons	light, long-lived hadrons
Reconstructed yield	0.0317 ± 0.0006	0.118 ± 0.001	0.406 ± 0.002
$\epsilon(p_{\mathrm{T}} > 2\mathrm{GeV/c})$	$85.6 \pm 0.6\%$	$51.8 \pm 0.5\%$	$2.34 \pm 0.08\%$
$\epsilon(\tau > 0.2\ \mathrm{ps})$	$88.1 \pm 0.6\%$	$63.1 \pm 0.5\%$	$99.46 \pm 0.03\%$
$\epsilon(p_{\mathrm{T}}) \times \epsilon(\tau)$	$75.9 \pm 0.8\%$	$32.6 \pm 0.4\%$	$2.30 \pm 0.08\%$
$\epsilon(p_{\mathrm{T}}) \times \epsilon(\tau) \times \epsilon(\mathrm{LHCb})$	$27.9 \pm 0.3\%$	$22.6 \pm 0.3\%$	$2.17 \pm 0.07\%$
Output rate	270 kHz	800 kHz	264 kHz

Table 6: Per-event yields determined from 100k of upgrade minimum-bias events after partial offline reconstruction. The first row indicates the number of candidates which had at least two tracks from which a vertex could be produced. The last row shows the output rate of a trigger selecting such events with perfect efficiency, assuming an input rate of 30 MHz from the LHC, as expected during upgrade running. A breakdown of each category is available in Table 14.

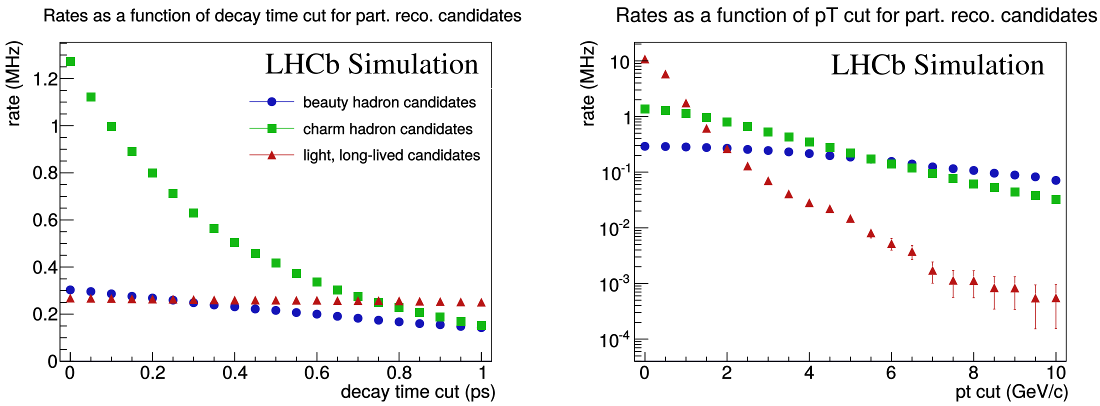

Figure 1: HLT partially reconstructed (but fully reconstructible) signal rates as a function of decay time for candidates with $p_{\mathrm{T}} > 2\ \mathrm{GeV/c}$ (left) and transverse momentum cuts for candidates with $\tau > 0.2\ \mathrm{ps}$(right). The rate is for two-track combinations that form a vertex only for candidates that can be fully reconstructed offline, ie: All additional tracks are also within the LHCb acceptance.

low efficiency of the L0 hardware trigger. The enormous rate of c hadrons explains why LHCb mainly relies on exclusive triggers for c hadrons already during current data taking.

These signal rates are already high, but the signal rates facing the upgraded detector are much higher still, particularly compared to the currently budgeted output rate of $10 - 50$ kHz. We could select 270 kHz of b hadron **signal** alone using an inclusive trigger, and three times as many c events. Within these constraints it is clear that any inclusive trigger strategy will have to favour some kinds of signals with respect to others. In other words, signals not explicitly used in training the inclusive triggers may get treated as part of the background, weakening one of the primary strengths of an inclusive trigger strategy : the ability to invent new analyses after the data has already been taken. Of course this problem always exists to an extent in any trigger, but it will become critical for the LHCb upgrade.

To see this even more clearly, consider the rate of the $D^0 \to K^+K^-$ decay with p_T above 2 GeV/c and a decay-time above 0.2 ps. This decay is one of the most important signals for the upgraded detector as it is sensitive to the indirect $\mathcal{CP}$ violation parameter A_Γ. From Table 11 we can see that 14% of all upgrade events will contain a D^0 meson whose products are in the detector acceptance. The branching fraction of this decay is 0.4% and Table 14 tells us that the efficiency of the listed p_T and τ cuts is 31.4%, resulting in a total output rate of approximately 5.4 kHz for an LHC interaction rate of 30 MHz. This rate can be significantly decreased by requiring that the D^0 came from a $D^{*+} \to D^0\pi^+$ decay chain, so that the slow pion tags its flavour, a requirement which would anyhow be applied by the offline analysis. From current experience this can be used to decrease the event yield by a factor of around 5, still leaving us with around 1 kHz of signal, or about 5% of the baseline trigger rate! A similar calculation yields a signal rate of around 2.4 kHz for the $D_s^+ \to \pi^+\pi^-\pi^+$ decay mode which is also very important as it is highly sensitive to direct $\mathcal{CP}$ violation : we have now used 17% of the trigger bandwidth simply to select two specific charm hadron signatures. It is clear that the output rate of signal is the single most serious difficulty facing the upgraded trigger system.

7 Conclusion and outlook

The upgrade of the LHCb experiment will demand fundamentally new approaches to triggering and online event selection. Past triggers at comparable detectors were either trivial (LEP, B-factories) or aimed largely at preselecting a small signal out of a huge background which was itself of no physics interest (CDF, ATLAS, CMS). In either case inclusive trigger strategies were ideally suited to the task, as they minimize selection biases and sensitivity to the running conditions while achieving the required reduction rate. The fact that 2% of upgraded LHCb events will contain a b hadron which is reconstructible within the detector while 24% will contain a reconstructible c hadron, however, changes the qualitative purpose of a trigger. If we are only able to write around 0.05% of events to disk then it is clearly impossible to maintain high efficiencies for all potential signal decay topologies, as an inclusive strategy aims to do, even in the ideal case where all background is neglected. Instead the trigger strategy must shift towards a greater reliance on exclusive selections which select specific, "golden", signal decay topologies with high efficiencies while suppressing other signal decay topologies. In particular, we know that most of the sensitivity to interesting observables comes from Cabibbo-suppressed decays; the trigger must be able to separate these from topologically identical Cabibbo-favoured decay topologies in a minimally biasing manner, which implies ever greater use of hadronic particle identification in the trigger.

Of course inclusive triggers will always be required in order to enable the widest possible physics reach, but equally achieving the maximal possible efficiency will require exclusive triggers for many signals. We should be clear about what this means from an operational point of view : the number of exclusive triggers will undoubtedly grow over the lifetime of the experiment, which means that the proportion of trigger time spent making high-level objects such as composite particles will grow. A fast combinatorics engine for combining tracks into such high-level objects is therefore a necessary precondition for successful triggering in the LHCb upgrade. In addition, the trigger will necessarily become more sensitive to operational fluctuations in the detector : monitoring the trigger performance and being able to rapidly fix operational problems will therefore be of much greater importance to the upgraded experiment than it is today. Finally, as the exclusive triggers will essentially implement a slightly looser version of the offline selection algorithms, they will be much more sensitive to low-level performance parameters such as track-finding efficiencies and asymmetries, momentum and lifetime resolutions, and so on. A continuous, real time, alignment and calibration of the detector will therefore also be crucial to the success of the LHCb upgrade.

For all these challenges, exclusive triggers also offer certain important advantages. In particular, we emphasize that exclusive triggers which select fully-reconstructed decay topologies allow us to select candidates using the physics quantities of interest, especially mass and lifetime, thus avoiding complex acceptance effects which would later have to be modelled by the offline analysts. The elimination of lifetime acceptances for the most important hadronic decay modes could prove crucial to the upgrade, particularly in light of LHCb's unique ability to make precise measurements of time-dependent $\mathcal{CP}$ violation in B_s^0 mesons.

Acknowledgements

We would like to thank Hans Dijkstra, Gerhard Raven, and Johannes Albrecht for useful discussions.

8 References

[1] Vladimir V Gligorov. A single track HLT1 trigger. Technical Report LHCb-PUB-2011-003.

[2] Vladimir V Gligorov, C. Thomas, and M. Williams. The HLT inclusive B triggers. Technical Report LHCb-PUB-2011-016.

[3] Roel Aaij and Johannes Albrecht. Muon triggers in the High Level Trigger of LHCb. Technical Report LHCb-PUB-2011-017.

[4] R. Aaij et al. The LHCb trigger and its performance in 2011. *JINST*, 8:P04022, 2013.

[5] J. Albrecht, V.V. Gligorov, G. Raven, and S. Tolk. Performance of the LHCb High Level Trigger in 2012. 2013.

[6] J. Albrecht et al. Impact of hardware choices on the LHCb upgrade trigger strategy. Technical Report LHCb-INT-2013-035. CERN-LHCb-INT-2013-035, CERN, Geneva, Jun 2013.

[7] Matt Dobbs and Jrgen Beck Hansen. The hepmc c++ monte carlo event record for high energy physics. *Computer Physics Communications*, 134(1):41 – 46, 2001.

[8] Jr. Alves, A. Augusto et al. The LHCb Detector at the LHC. *JINST*, 3:S08005, 2008.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
B^0	0.0111 ± 0.0002	$31 \pm 1\%$	$11.3 \pm 0.7\%$
B^+	$(9.4 \pm 0.2) \times 10^{-3}$	$30 \pm 1\%$	$10.8 \pm 0.7\%$
B_s^0	$(3.4 \pm 0.1) \times 10^{-3}$	$31 \pm 2\%$	$11 \pm 1\%$
B_c^+	0 ± 0	-	-
Λ_b	$(1.7 \pm 0.1) \times 10^{-3}$	$32 \pm 2\%$	$12 \pm 2\%$
Ξ_b^+	$(10 \pm 6) \times 10^{-6}$	0 ± 0	0 ± 0
Ξ_b^0	$(2 \pm 1) \times 10^{-5}$	$25 \pm 22\%$	$25 \pm 22\%$
Ω_b^0	0 ± 0	-	-
Total	0.0258 ± 0.0003	$30.5 \pm 0.6\%$	$11.1 \pm 0.4\%$

Table 7: Per-event yields and efficiencies of b-hadron candidates in Run 1 conditions. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the current VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
D^0	0.177 ± 0.001	$23.6 \pm 0.2\%$	$16.1 \pm 0.2\%$
D^+	0.0608 ± 0.0005	$18.7 \pm 0.3\%$	$11.7 \pm 0.3\%$
D_s^+	0.0330 ± 0.0004	$19.3 \pm 0.5\%$	$11.2 \pm 0.4\%$
Λ_c^+	0.0230 ± 0.0003	$21.3 \pm 0.589\%$	$12.2 \pm 0.5\%$
Ξ_c^0	$(1.74 \pm 0.09) \times 10^{-3}$	$16 \pm 2\%$	$9 \pm 1\%$
Ξ_c^+	$(1.57 \pm 0.09) \times 10^{-3}$	$17 \pm 2\%$	$10 \pm 2\%$
Ω_c^0	$(6 \pm 2) \times 10^{-5}$	$17 \pm 11\%$	$0 \pm 0\%$
Total	0.297 ± 0.001	$21.9 \pm 0.2\%$	$14.2 \pm 0.1\%$

Table 8: Per-event yields and efficiencies of c-hadron candidates in Run 1 conditions. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the current VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
K_S^0	5.000 ± 0.005	$7.34 \pm 0.03\%$	$7.01 \pm 0.02\%$
Λ^0	2.332 ± 0.003	$4.83 \pm 0.03\%$	$4.66 \pm 0.03\%$
Σ^+	0.413 ± 0.001	$11.0 \pm 0.1\%$	$10.6 \pm 0.1\%$
Ξ^0	0.146 ± 0.001	$2.00 \pm 0.08\%$	$1.00 \pm 0.06\%$
Ξ^-	0.144 ± 0.001	$5.7 \pm 0.1\%$	$3.9 \pm 0.1\%$
Ω^-	$(3.8 \pm 0.1) \times 10^{-3}$	$5.6 \pm 0.8\%$	$4.1 \pm 0.7\%$
Total	8.04 ± 0.01	$6.67 \pm 0.02\%$	$6.35 \pm 0.02\%$

Table 9: Per-event yields and efficiencies of light, long-lived candidates in Run 1 conditions. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the current VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
B^0	0.06776 ± 0.0003	$35.1 \pm 0.2\%$	$11.9 \pm 0.1\%$
B^+	0.0566 ± 0.0002	$34.8 \pm 0.2\%$	$12.1 \pm 0.1\%$
B_s^0	0.0211 ± 0.0002	$34.7 \pm 0.3\%$	$11.7 \pm 0.2\%$
B_c^+	$(5.1 \pm 0.7) \times 10^{-5}$	$41 \pm 7\%$	$14 \pm 5\%$
Λ_b	0.0111 ± 0.0001	$34.6 \pm 0.5\%$	$12 \pm 0.3\%$
Ξ_b^+	$(6.10 \pm 0.8) \times 10^{-5}$	$33 \pm 6\%$	$2 \pm 2\%$
Ξ_b^0	$(7.90 \pm 0.9) \times 10^{-5}$	$44 \pm 6\%$	$15 \pm 4\%$
Ω_b^0	$(2.10 \pm 0.5 \times 10^{-5}$	$33 \pm 10\%$	$14 \pm 8\%$
Total	0.1572 ± 0.0004	$34.9 \pm 0.1\%$	$11.9 \pm 0.1\%$

Table 10: Per-event yields and efficiencies of b-hadron candidates after the upgrade. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the upgrade VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
D^0	0.8372 ± 0.0009	$26.89 \pm 0.05\%$	$17.02 \pm 0.04\%$
D^+	0.2971 ± 0.0006	$21.20 \pm 0.07\%$	$12.39 \pm 0.06\%$
D_s^+	0.1635 ± 0.0004	$21.8 \pm 0.1\%$	$11.68 \pm 0.08\%$
Λ_c^+	0.111 ± 0.0003	$22.8 \pm 0.1\%$	$12.8 \pm 0.1\%$
Ξ_c^0	$(7.77 \pm 0.09) \times 10^{-3}$	$16.8 \pm 0.4\%$	$10.2 \pm 0.3\%$
Ξ_c^+	$(7.47 \pm 0.09) \times 10^{-3}$	$16 \pm 0.4\%$	$8.3 \pm 0.3\%$
Ω_c^0	$(3.9 \pm 0.2) \times 10^{-4}$	$17 \pm 2\%$	$8 \pm 1\%$
Total	1.422 ± 0.001	$24.73 \pm 0.04\%$	$15.12 \pm 0.03\%$

Table 11: Per-event yields and efficiencies of c-hadron candidates after the upgrade. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the upgrade VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Category	In 4π	ϵ(VELO)	ϵ(VELO) $\times$ ϵ(LHCb)
K_S^0	20.813 ± 0.005	$7.782 \pm 0.006\%$	$6.941 \pm 0.006\%$
Λ^0	9.547 ± 0.003	$4.943 \pm 0.007\%$	$4.506 \pm 0.007\%$
Σ^+	1.723 ± 0.001	$11.84 \pm 0.03\%$	$10.82 \pm 0.02\%$
Ξ^0	0.6196 ± 0.0008	$1.81 \pm 0.02\%$	$0.86 \pm 0.01\%$
Ξ^-	0.6021 ± 0.0008	$5.01 \pm 0.03\%$	$3.29 \pm 0.02\%$
Ω^-	0.0161 ± 0.0001	$6.8 \pm 0.2\%$	$4.3 \pm 0.2\%$
Total	33.291 ± 0.006	$7.022 \pm 0.004\%$	$6.257 \pm 0.004\%$

Table 12: Per-event yields and efficiencies of light, long-lived candidates after the upgrade. ϵ(VELO) is the efficiency for candidates having at least two tracks traversing at least three modules in the upgrade VELO. ϵ(LHCb) is the efficiency for candidates having all daughter tracks contained in the LHCb acceptance.

Cand.	Reco. yield	$\epsilon(p_T > 2$ GeV/c)	$\epsilon(\tau > 0.2$ ps)	$\epsilon(p_T) \times \epsilon(\tau)$	$\epsilon(p_T) \times \epsilon(\tau) \times \epsilon$(LHCb)
B^0	$(1.77 \pm 0.092) \times 10^{-3}$	$85.4 \pm 1.8\%$	$86.5 \pm 1.77\%$	$72.8 \pm 2.31\%$	$28.2 \pm 1.8\%$
B^+	$(1.50 \pm 0.085) \times 10^{-3}$	$81.8 \pm 2.3\%$	$91.7 \pm 1.56\%$	$74.8 \pm 2.45\%$	$29.8 \pm 2.1\%$
B_s^0	$(3.91 \pm 0.43) \times 10^{-4}$	$73.2 \pm 4.9\%$	$87.8 \pm 3.61\%$	$65.9 \pm 5.24\%$	$27.2 \pm 3.5\%$
B_c^+	0	-	-	-	-
Λ_b	$(3.53 \pm 0.41) \times 10^{-4}$	$85.1 \pm 4.1\%$	$89.2 \pm 3.61\%$	$77 \pm 4.89\%$	$29.3 \pm 4.6\%$
Ξ_b^+	0	-	-	-	-
Ξ_b^0	0	-	-	-	-
Ω_b^0	0	-	-	-	-
Total	$(4.01 \pm 0.14) \times 10^{-3}$	$82.9 \pm 1.3\%$	$88.8 \pm 1.09\%$	$73.2 \pm 1.53\%$	$28.8 \pm 1.2\%$
D^0	0.014 ± 0.00026	$46.9 \pm 0.92\%$	$61.8 \pm 0.897\%$	$28.7 \pm 0.834\%$	$21.7 \pm 0.68\%$
D^+	0.00308 ± 0.00012	$47.8 \pm 2.0\%$	$83.6 \pm 1.46\%$	$40.6 \pm 1.93\%$	$30.1 \pm 1.6\%$
D_s^+	$(1.80 \pm 0.093) \times 10^{-5}$	$49.3 \pm 2.6\%$	$69.5 \pm 2.37\%$	$33.2 \pm 2.42\%$	$21.8 \pm 1.8\%$
Λ_c^+	$(1.19 \pm 0.075) \times 10^{-5}$	$47.6 \pm 3.2\%$	$37.2 \pm 3.06\%$	$18.8 \pm 2.47\%$	$13.1 \pm 1.8\%$
Ξ_c^0	$(3.34 \pm 1.3) \times 10^{-5}$	$14.3 \pm 13\%$	$0 \pm 0\%$	$0 \pm 0\%$	$0 \pm 0\%$
Ξ_c^+	$(8.1 \pm 2.0) \times 10^{-5}$	$41.2 \pm 12\%$	$58.8 \pm 11.9\%$	$17.6 \pm 9.25\%$	$9.5 \pm 5.5\%$
Ω_c^0	0	-	-	-	-
Total	0.020 ± 0.00031	$47.2 \pm 0.767\%$	$64.2 \pm 0.737\%$	$30.2 \pm 0.706\%$	$22.3 \pm 0.6\%$
K_S^0	0.068 ± 0.00057	$1.69 \pm 0.11\%$	$99.5 \pm 0.0608\%$	$1.66 \pm 0.107\%$	$1.63 \pm 0.11\%$
Λ^0	0.010 ± 0.000223	$3.63 \pm 0.40\%$	$99.8 \pm 0.0918\%$	$3.63 \pm 0.401\%$	$3.58 \pm 0.40\%$
Σ^+	0	-	-	-	-
$\overline{\Xi}^0$	$(1.43 \pm 0.83) \times 10^{-5}$	0 ± 0	100 ± 0	-	-
Ξ^-	$(3.05 \pm 0.38) \times 10^{-4}$	$3.1 \pm 2.2\%$	$100 \pm 0\%$	$3.12 \pm 2.17\%$	$3.12 \pm 3.80\%$
Ω^-	$(2.4 \pm 1.1) \times 10^{-5}$	$20 \pm 18\%$	$100 \pm 0\%$	$20 \pm 17.9\%$	$20 \pm 21\%$
Total	0.079 ± 0.00061	$1.96 \pm 0.11\%$	$99.5 \pm 0.0539\%$	$1.93 \pm 0.107\%$	$1.90 \pm 0.11\%$

Table 13: Per-event yields determined from 0.21M of Run 1 minimum-bias events after partial offline reconstruction. The first column indicates the number of candidates which had at least two tracks from which a vertex could be produced. The second column indicates the efficiency of a $p_T > 2$GeV/c cut applied to the fully reconstructed candidate.

Cand.	Reco. yield	$\epsilon(p_{\mathrm{T}} > 2\ \mathrm{GeV/c})$	$\epsilon(\tau > 0.2\ \mathrm{ps})$	$\epsilon(p_{\mathrm{T}}) \times \epsilon(\tau)$	$\epsilon(p_{\mathrm{T}}) \times \epsilon(\tau) \times \epsilon(\mathrm{LHCb})$
B^0	0.014 ± 0.00038	$86.0 \pm 0.95\%$	$87 \pm 0.917\%$	$75.4 \pm 1.17\%$	$27.7 \pm 0.45\%$
B^+	0.013 ± 0.00036	$86.2 \pm 0.99\%$	$88.3 \pm 0.925\%$	$76.2 \pm 1.22\%$	$28.1 \pm 0.47\%$
B_s^0	$(3.42 \pm 0.19) \times 10^{-3}$	$85.8 \pm 1.9\%$	$91 \pm 1.57\%$	$78 \pm 2.27\%$	$28.3 \pm 0.86\%$
B_c^+	$(4.1 \pm 2.1) \times 10^{-5}$	$75 \pm 22\%$	$50 \pm 25\%$	$25 \pm 21.7\%$	$8.3 \pm 7.4\%$
Λ_b	$(2.36 \pm 0.16) \times 10^{-3}$	$80.3 \pm 2.6\%$	$90.8 \pm 1.91\%$	$74.7 \pm 2.87\%$	$27.2 \pm 1.1\%$
Ξ_b^+	0	-	-	-	-
Ξ_b^0	0	-	-	-	-
Ω_b^0	0	-	-	-	-
Total	0.032 ± 0.00058	$85.6 \pm 0.63\%$	$88.1 \pm 0.58\%$	$75.9 \pm 0.77\%$	$27.9 \pm 0.29\%$
D^0	0.080 ± 0.00091	$51.3 \pm 0.57\%$	$60.7 \pm 0.554\%$	$31.4 \pm 0.526\%$	$22.4 \pm 0.38\%$
D^+	0.0182 ± 0.00043	$51.2 \pm 1.2\%$	$82.2 \pm 0.908\%$	$40.9 \pm 1.17\%$	$27.6 \pm 0.79\%$
D_s^+	0.0118 ± 0.00035	$54.4 \pm 1.5\%$	$67.8 \pm 1.38\%$	$37.3 \pm 1.43\%$	$23.5 \pm 0.90\%$
Λ_c^+	$(7.50 \pm 0.28) \times 10^{-3}$	$54.6 \pm 1.8\%$	$35.4 \pm 1.77\%$	$18.2 \pm 1.43\%$	$12.1 \pm 0.95\%$
Ξ_c^0	$(2.16 \pm 0.47) \times 10^{-4}$	$38.1 \pm 11\%$	$4.76 \pm 4.65\%$	$4.76 \pm 4.65\%$	$3.2 \pm 3.2\%$
Ξ_c^+	$(2.26 \pm 0.48) \times 10^{-4}$	$40.9 \pm 11\%$	$59.1 \pm 10.5\%$	$36.4 \pm 10.3\%$	$23.6 \pm 6.7\%$
Ω_c^0	$(1.0 \pm 1.0) \times 10^{-5}$	$100 \pm 0\%$	$0 \pm 0\%$	$0 \pm 0\%$	$0 \pm 0\%$
Total	0.118 ± 0.0011	$51.8 \pm 0.47\%$	$63.1 \pm 0.451\%$	$32.6 \pm 0.438\%$	$22.6 \pm 0.30\%$
K_S^0	0.35 ± 0.0019	$1.99 \pm 0.076\%$	$99.5 \pm 0.039\%$	$1.95 \pm 0.0751\%$	$1.84 \pm 0.071\%$
Λ^0	0.055 ± 0.00075	$4.47 \pm 0.28\%$	$99.8 \pm 0.0647\%$	$4.43 \pm 0.281\%$	$4.18 \pm 0.27\%$
Σ^+	$(2.1 \pm 1.5) \times 10^{-5}$	$50 \pm 35\% 100 \pm 0\%$	$50 \pm 35.4\%$	$47 \pm 33\%$	
$\overline{\Xi}^0$	$(1.0 \pm 1.0) \times 10^{-5}$	$0 \pm 0\%$	$100 \pm 0\%$	$0 \pm 0\%$	$0 \pm 0\%$
Ξ^-	$(1.52 \pm 0.13) \times 10^{-3}$	$4.7 \pm 1.7\%$ $100 \pm 0\%$	$4.73 \pm 1.74\%$	$4.3 \pm 1.6\%$	
Ω^-	$(1.65 \pm 0.41) \times 10^{-5}$	$6.3 \pm 6.1\%$	$100 \pm 0\%$	$6.25 \pm 6.05\%$	$4.8 \pm 4.7\%$
Total	0.406 ± 0.0020	$2.34 \pm 0.076\%$	$99.5 \pm 0.0347\%$	$2.30 \pm 0.076\%$	2.17 ± 0.071

Table 14: Per-event yields determined from 100k of upgrade minimum-bias events after partial offline reconstruction. The first column indicates the number of candidates which had at least two tracks from which a vertex could be produced. The second column indicates the efficiency of a $p_{\mathrm{T}} > 2\mathrm{GeV/c}$ cut applied to the fully reconstructed candidate.