LHCb detector performance

The LHCb Collaboration*

Received 19 December 2014

Accepted 22 December 2014

Published 5 March 2015

The LHCb detector is a forward spectrometer at the Large Hadron Collider (LHC) at CERN. The experiment is designed for precision measurements of CP violation and rare decays of beauty and charm hadrons. In this paper the performance of the various LHCb sub-detectors and the trigger system are described, using data taken from 2010 to 2012. It is shown that the design criteria of the experiment have been met. The excellent performance of the detector has allowed the LHCb collaboration to publish a wide range of physics results, demonstrating LHCb's unique role, both as a heavy flavour experiment and as a general purpose detector in the forward region.

Keywords: Large detector systems for particle and astroparticle physics; particle tracking detectors; gaseous detectors; calorimeters; Cherenkov detectors; particle identification methods; detector alignment and calibration methods; trigger; LHC.

PACS numbers: 29.40.Cs, 29.40.Gx, 29.40.Ka, 29.40.Vj, 29.40.Wk, 29.85.Ca

Contents

1. Introduction	2
1.1. Physics goals of the LHCb experiment	2
1.2. Overview of the experimental setup	4
1.3. Data taking periods and operating conditions	7
2. Charged Particle Reconstruction	11
2.1. Hit efficiencies and hit resolutions of the tracking detectors	11
2.1.1. Vertex Locator	11
2.1.2. Silicon Tracker	14
2.1.3. Outer Tracker	16
2.1.4. Muon system	18
2.2. Track reconstruction	21
2.2.1. Track finding efficiency	23
2.2.2. Mass and momentum resolution	25

This is an Open Access article published by World Scientific Publishing Company. It is distributed under the terms of the Creative Commons Attribution 4.0 (CC-BY) License. Further distribution of this work is permitted, provided the original work is properly cited.

*See Appendix A for full author list.

2.3.	Spatial alignment of the tracking detectors	27
2.3.1.	Vertex locator alignment	28
2.3.2.	Alignment of the silicon tracker and outer tracker	30
2.3.3.	Muon system alignment	32
2.4.	Vertexing and decay time resolution	34
2.4.1.	Primary vertex reconstruction	34
2.4.2.	Impact parameter resolution	35
2.4.3.	Decay time resolution	36
2.4.4.	V^0 reconstruction	37
3.	Neutral Particle Reconstruction	38
3.1.	Calibration of the calorimeter system	38
3.2.	Selection of neutral energy deposits in ECAL	40
3.3.	Photon reconstruction	41
3.4.	Neutral pion reconstruction	41
4.	Particle Identification	43
4.1.	Calorimeter system based particle identification	43
4.1.1.	Photon and merged π^0 identification	43
4.1.2.	Electron identification	45
4.2.	RICH system based particle identification	47
4.2.1.	Cherenkov angle resolution	48
4.2.2.	Photoelectron yield	49
4.2.3.	Particle identification performance	50
4.3.	Muon system based particle identification	50
4.4.	Combined particle identification performance	53
5.	Trigger	55
5.1.	Data driven trigger performance determination	55
5.2.	Level-0 hardware trigger	56
5.3.	High level trigger	58
5.3.1.	First level	58
5.3.2.	Second level	60
5.4.	Deferred trigger	63
5.5.	Trigger performance summary	64
6.	Conclusion and Outlook	65
	Acknowledgments	65
	References	66
	Appendix A. LHCb Collaboration	68

1. Introduction

1.1. *Physics goals of the LHCb experiment*

LHCb is a dedicated heavy flavour physics experiment at the LHC. Its main goal is to search for indirect evidence of new physics in CP violation and rare decays of beauty and charm hadrons, by looking for the effects of new particles in processes that are precisely predicted in the Standard Model (SM) and by utilising the distinctive flavour structure of the SM with no tree-level flavour-changing neutral currents. Quark mixing in the SM is described by the Cabibbo–Kobayashi–Maskawa (CKM) matrix,^{1,2} which has a single source of CP violation. Since the level of CP violation in weak interactions cannot explain the matter–antimatter asymmetry in

the universe,³ new sources of CP violation beyond the SM are needed. The effect of such new sources might be seen in heavy flavour physics, where many models of new physics produce contributions that change the expectation values of the CP violating phases or the branching fractions of rare decays. Some models even predict decay modes that are forbidden in the SM. To examine such possibilities, CP violation and rare decays of hadrons containing b and c quarks must be studied with large data samples, using many different decay modes.

Thanks to the large beauty and charm production cross-section at the LHC,^{4,5} the LHCb experiment collected $\sim 10^{12}$ heavy flavour decays during 2011 and 2012. Despite these large yields, at the LHC centre-of-mass energies of $\sqrt{s} = 7\text{--}8$ TeV the charm and beauty cross-sections are approximately a factor 10 and 200 smaller than the total cross-section, respectively. To separate the decays of interest from the background, both displaced vertex and high transverse momentum signatures are exploited. Excellent vertex resolution is required to measure impact parameters and to achieve a good decay time resolution, which is essential to resolve B_s^0 flavour oscillations and to reject various sources of background. Good momentum and invariant mass resolution are important to minimise combinatorial background and resolve heavy-flavour decays with kinematically similar topologies. Charged particle identification is essential in any flavour physics programme, for instance to isolate suppressed decays and for b -quark flavour tagging. Detection of photons, in addition to charged particles, allows the reconstruction of rare radiative decays and more common decays with a π^0 or an η meson in the final state. Finally, to benefit from the high event rate at the LHC, a high-bandwidth data acquisition system and a robust and selective trigger system are required.

LHCb has various advantages over the $e^+e^- B$ factories, including a higher cross-section, a larger boost and the fact that all species of b hadrons are produced. Less attractive characteristics of the LHC environment are the generally increased background levels encountered, inherent to hadronic collisions, which result in a number of experimental compromises, such as reduced b flavour tagging efficiency and the difficulty in reconstructing final states with missing or neutral particles. Despite these challenges, the results⁶ obtained from data taken between 2010 and 2013 (LHC Run I) have clearly established LHCb as the next generation flavour physics experiment. Thanks to efficient charged particle tracking and dedicated triggers for lepton, hadron and photon signatures, LHCb has the world's largest sample of exclusively reconstructed charm and beauty decays. With these samples, LHCb has already made many key results, such as the first evidence for the rare decay $B_s^0 \rightarrow \mu^+\mu^-$ (Refs. 7 and 8) and measurements of angular distributions in the $B^0 \rightarrow K^{*0}\mu^+\mu^-$ decay,^{9,10} which are particularly sensitive to deviations from the SM. Another example is the measurement of the CP violating phase (ϕ_s) in the interference between mixing and decay of B_s^0 mesons, where the value predicted within the SM is small, but much larger values are possible in new physics models. LHCb has measured this phase with results that are at present consistent with the SM within the uncertainties.^{11,12} The measurement of the angle γ of the Unitarity

Triangle from $B \rightarrow DK$ decays is a crucial component in the determination of the parameters of the CKM quark mixing matrix. The γ results from LHCb^{13,14} already dominate the global averages. In the charm sector, one of the most interesting observables (A_Γ) is the difference in the inverse effective lifetimes between D^0 and $\bar{D}^0$ decays. The most precise measurement of A_Γ to date has been presented by LHCb.¹⁵ These are just some of the results from LHCb that have made a significant impact on the flavour physics landscape.

The physics output of LHCb extends well beyond this core programme. Examples of other topics include: measurements of the production of electroweak gauge bosons in the forward kinematic region, uniquely covered by the LHCb acceptance;^{16,17} measurements of the properties of newly discovered exotic hadrons;^{18,19} searches for lepton number and lepton flavour violation,^{20,21} and measurements of heavy quarkonia in proton–lead collisions.^{22,23} These illustrate the wide variety in electroweak and QCD topics covered by the LHCb experiment and establish LHCb as a general purpose detector in the forward region at a hadron collider.

In the remainder of this introduction an overview of the LHCb detector is given, together with a summary of the data-taking periods and the operating conditions. Thereafter, the paper discusses charged particle reconstruction, vertexing and decay-time resolution in Sec. 2, neutral particle reconstruction in Sec. 3 and particle identification in Sec. 4. The performance results shown are indicative, and depend on the specific requirements set by a physics analysis, for example to achieve high efficiency or high purity. Section 5 discusses the trigger and the paper concludes with a short summary in Sec. 6.

1.2. Overview of the experimental setup

LHCb is a single-arm spectrometer with a forward angular coverage from approximately 15 mrad to 300 (250) mrad in the bending (non-bending) plane.²⁵ The choice of the detector geometry is driven by the fact that at high energies production of the b - and $\bar{b}$ -hadrons is highly correlated, such that they are predominantly produced in the same forward or backward cone. The layout of the LHCb spectrometer is shown in Fig. 1. Most detector subsystems are assembled in two halves, which can be moved out horizontally for assembly and maintenance purposes, as well as to provide access to the beam-pipe. They are referred to as the detector A- and C-sides. A right-handed coordinate system is defined with z along the beam axis into the detector, y vertical and x horizontal. Cylindrical polar coordinates (r, ϕ, z) are also used, as appropriate.

The spectrometer magnet, required for the momentum measurement of charged particles, is a warm dipole magnet providing an integrated field of about 4 Tm, which deflects charged particles in the horizontal plane. The field of the spectrometer magnet also has an impact on the trajectory of the LHC beams. Three dipole magnets are used to compensate for this effect and to ensure a closed orbit for the beams.²⁶

Fig. 1. View of the LHCb detector.²⁴

The tracking system consists of the Vertex Locator (VELO), situated around the interaction region inside a vacuum tank, and four planar tracking stations: the Tracker Turicensis (TT) upstream of the dipole magnet, and tracking stations T1–T3 downstream of the magnet. Silicon microstrips are used in TT and the region close to the beam-pipe (Inner Tracker, IT) of stations T1–T3, whereas straw tubes are employed for the outer parts (Outer Tracker, OT). Charged particles require a minimum momentum of $1.5 \text{ GeV}/c$ to reach the tracking stations, T1–T3.

The VELO contains 42 silicon modules arranged along the beam, each providing a measurement of the r (R sensors) and ϕ (Φ sensors) coordinates. The pitch within a module varies from $38 \mu\text{m}$ at the inner radius of 8.2 mm , increasing linearly to $102 \mu\text{m}$ at the outer radius of 42 mm . For detector safety, the VELO modules are retracted by 29 mm in the horizontal direction during injection of the LHC beams and are subsequently moved back, using a fully automated procedure once stable conditions have been declared. From the declaration of stable beams the VELO takes, on average, 210 seconds to close. During LHC Run I approximately 750 closing procedures were performed.

The TT and IT detectors use silicon microstrip sensors with a strip pitch of $183 \mu\text{m}$ and $198 \mu\text{m}$, respectively. The TT is about 150 cm wide and 130 cm high, with a total active area of around 8 m^2 . The IT covers a 120 cm wide and 40 cm high cross-shaped region in the centre of the three tracking stations T1–T3. The total active area of the IT is approximately 4 m^2 . Each of the tracking stations has four detection layers in an x – u – v – x arrangement with vertical strips in each of the

two x layers, and strips rotated by a stereo angle of -5° and $+5^\circ$ in the u and v layers, respectively.

The Outer Tracker is a drift-tube gas detector consisting of approximately 200 gas-tight straw-tube modules with drift-time read-out. Each module contains two staggered layers of drift-tubes with an inner diameter of 4.9 mm. As a counting gas, a mixture of Argon (70%), CO_2 (28.5%) and O_2 (1.5%) is chosen to guarantee a drift time below 50 ns and a spatial resolution of $200\ \mu\text{m}$. As for the IT part of T1–T3, the OT has four layers arranged in an x – u – v – x geometry. The total active area of a station is $597\text{ cm} \times 485\text{ cm}$.

Charged hadron identification in the momentum range from 2 to 100 GeV/ c is achieved by two Ring Imaging Cherenkov detectors (RICH1 and RICH2) read out by Hybrid Photon Detectors (HPDs). The upstream detector, RICH1, covers the low momentum charged particle range from about 2 to 60 GeV/ c and uses Aerogel and C_4F_{10} as radiators, while the downstream detector, RICH2, covers the high momentum range from about 15 GeV/ c to 100 GeV/ c , using a CF_4 radiator. RICH1 has a wide acceptance, covering the LHCb acceptance from $\pm 25\text{ mrad}$ to $\pm 300\text{ mrad}$ (horizontal) and $\pm 250\text{ mrad}$ (vertical), while RICH2 has a limited angular acceptance of $\pm 15\text{ mrad}$ to $\pm 120\text{ mrad}$ (horizontal) and $\pm 100\text{ mrad}$ (vertical).

The calorimeter system is composed of a Scintillating Pad Detector (SPD), a Preshower (PS), a shashlik type electromagnetic calorimeter (ECAL) and a hadronic calorimeter (HCAL). It provides the identification of electrons, photons and hadrons as well as the measurement of their energies and positions, and selects candidates with high transverse energy for the first trigger level (L0). The SPD improves the separation of electrons and photons. A 15 mm lead converter with a thickness of 2.5 radiation lengths (X_0) is placed between the planes of rectangular scintillating pads of the SPD and the PS. The background from charged pions is reduced by a measurement of the longitudinal partitioning of the electromagnetic shower in the PS detector and the main section of ECAL. The ECAL is made of a sampling scintillator/lead structure with a total thickness of $25\ X_0$. The calorimeter system has a variable lateral segmentation which takes into account the variation in hit density of two orders of magnitude over the calorimeter surface. A segmentation into three different sections has been chosen for the ECAL with a corresponding projective geometry for the SPD and PS detectors, meaning that all of their transverse dimensions scale with the distance from the interaction point. The outer dimensions match projectively those of the tracking system, while the square hole around the beam-pipe approximately limits the inner acceptance to projective polar angles $\theta_{x,y} > 25\text{ mrad}$. The hadron calorimeter (HCAL) is a sampling device made from iron and scintillating tiles, as absorber and active material, respectively. The special feature of this sampling structure is the orientation of the scintillating tiles which run parallel to the beam-axis. Given the dimensions of the hadronic showers, the HCAL is segmented into two zones with different lateral dimensions. The thickness of the HCAL is limited to 5.6 nuclear interaction lengths (λ_i) due to space constraints.

The muon detection system provides muon identification and contributes to the L0 trigger of the experiment. It is composed of five stations (M1–M5) of rectangular shape equipped predominantly with Multi Wire Proportional Chambers (MWPC), except in the highest rate region of M1, where triple Gas Electron Multiplier (GEM) detectors are used. The full system comprises 1380 chambers and covers a total area of 435 m². Stations M2 to M5 are placed downstream of the calorimeters and are interleaved with 80 cm thick iron absorbers to select penetrating muons. The minimum momentum that a muon must have to traverse the five stations is approximately 6 GeV/c. The total absorber thickness, including the calorimeters, is approximately 20 λ_i . Station M1 is placed in front of the calorimeters and is used to improve the p_T measurement in the trigger. The geometry of the five stations is projective, with each station divided into four regions, R1 to R4, with increasing distance from the beam axis. The linear dimensions of the regions R1, R2, R3, R4, and their segmentation scale in the ratio 1:2:4:8. With this geometry, the channel occupancies are comparable in each of the four regions of a given station.

The LHCb trigger system consists of two levels. The first level is implemented in hardware and is designed to reduce the event rate from the nominal LHC bunch crossing rate of 40 MHz to a maximum of 1.1 MHz. The complete detector is then read out and the data is sent to the High Level Trigger (HLT) implemented on the Event Filter Farm (EFF), which had about 30,000 processing cores in 2012. The HLT is a software trigger, running a simplified version of the offline event reconstruction to accommodate the more stringent CPU time requirements.

1.3. Data taking periods and operating conditions

At the end of 2009, LHCb recorded its first pp collisions at the injection energy of the LHC, $\sqrt{s} = 0.9$ TeV. These data have been used to finalise the commissioning of the sub-detector systems and the reconstruction software, and to perform a first alignment and calibration of the tracking, calorimeter and particle identification (PID) systems. In this period, the VELO was left in the open position, due to the larger aperture required at lower beam energies.

During 2010 the operating conditions changed rapidly due to the ramp-up of the LHC luminosity. A critical parameter for LHCb performance is the pile-up μ_{vis} , defined as the average number of visible interactions per beam–beam crossing.²⁷ The evolution of the LHCb operating conditions during LHC Run I is shown in Fig. 2. Starting with luminosities $\sim 10^{28} \text{ cm}^{-2} \text{ s}^{-1}$ and almost no pile-up, the luminosity reached $10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ with $\mu_{\text{vis}} \approx 2.5$.

While the highest luminosity in 2010 was already 75% of the LHCb design luminosity, the pile-up was much larger than the design value due to the low number of bunches in the machine. It was demonstrated that the trigger and reconstruction work efficiently under such harsh conditions with increased detector occupancy due to pile-up, and that the physics output was not compromised.

Fig. 2. Average number of visible interactions per bunch crossing ('pile-up', top) and instantaneous luminosity (bottom) at the LHCb interaction point in the period 2010-2012. The dotted lines show the design values.

The LHC beam energy was 3.5 TeV during 2010 and 2011. In the first part of the 2011 data taking the number of bunches in the machine increased in several steps to about 1300, the maximum possible with 50 ns bunch spacing. Due to the larger number of bunches the pile-up over the year could be reduced, while LHCb took the majority of the data at a luminosity of $3.5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. This was 1.75 times more than the design luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, as shown in Fig. 2. In 2011 a luminosity levelling procedure was introduced at the LHCb interaction point. By adjusting the transverse overlap of the beams at LHCb, the instantaneous luminosity could be kept stable to within about 5% during a fill, as illustrated in Fig. 3. For this particularly long fill, a maximum overlap with head-on beams was reached only after 15 hours. The luminosity levelling procedure minimises the effects of luminosity decay, allowing to maintain the same trigger configuration during a fill and to reduce systematic uncertainties due to changes in the detector occupancy.

In 2012 the LHC beam energy was increased to 4 TeV. LHCb took data at a luminosity of $4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, twice the LHCb design luminosity. The LHC delivered stable beams for about 30% of the operational year. An effort was made in 2012 to use more efficiently the processing power available in the Event-Filter-Farm (EFF), which otherwise would have been idle during 70% of the time. The mechanism put in operation defers a fraction of the HLT processing to the inter-fill time, typically several hours, between the LHC collision periods. In this approach about 20% of the L0 accepted events during data-taking are temporarily saved on the local disks of the EFF nodes and are processed only after the end of stable

Fig. 3. Development of the instantaneous luminosity for ATLAS, CMS and LHCb during LHC fill 2651. After ramping to the desired value of $4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ for LHCb, the luminosity is kept stable in a range of 5% for about 15 hours by adjusting the transversal beam overlap. The difference in luminosity towards the end of the fill between ATLAS, CMS and LHCb is due to the difference in the final focusing at the collision points, commonly referred to as the beta function, β^* .

beams. This deferred triggering method allowed LHCb to increase the data sample available for physics analysis.

The integrated luminosity recorded by LHCb was 38 pb^{-1} in 2010, 1.11 fb^{-1} in 2011 and 2.08 fb^{-1} in 2012. The evolution of the integrated luminosity for the years 2010 to 2012 is shown in Fig. 4.

Luminosity calibrations were carried out with the LHCb detector for the various centre-of-mass energy $\sqrt{s}$ at which data has been taken. Both the “van der Meer scan” and “beam-gas imaging” luminosity calibration methods were employed.²⁸

Fig. 4. Integrated luminosity in LHCb during the three years of LHC Run I. The figure shows the curves for the delivered (dark coloured lines) and recorded (light coloured lines) integrated luminosities.

For proton–proton interactions at $\sqrt{s} = 8$ TeV a relative precision of the luminosity calibration of 1.47% was obtained using van der Meer scans and 1.43% using beam-gas imaging, resulting in a combined precision of 1.12%. Applying the calibration to the full data set determines the luminosity with a precision of 1.16%. This represents the most precise luminosity measurement achieved so far at a bunched-beam hadron collider.

The average operational efficiency, defined as the ratio of recorded over delivered luminosity, was 93% during LHC Run I, reaching 95% on average in 2012. The inefficiency contains two irreducible sources. The first one is the detector-safety procedure for the VELO closing, amounting to 0.9%, which is in line with expectations. The second originates from non-conformities in the implementation of the read-out protocol of some sub-detector front-end systems and introduces 2.4% of dead-time at 1 MHz read-out frequency. The remaining 3.6% is related to short technical problems with the sub-detector electronics or the central read-out system. About 99% of the recorded data is used for physics analyses.

After a short pilot run in 2012, the LHC delivered for the first time proton–lead collisions in January and February 2013. The beam energy of the proton beam was 4 TeV, while the corresponding nucleon energy of the lead beam was 1.58 TeV, corresponding to a centre-of-mass energy of 5 TeV. The LHC delivered collisions with both protons and lead nuclei as the clockwise, and anti-clockwise beams, which made it possible for LHCb to collect data in the forward and backward direction of proton–lead collisions. The integrated recorded luminosity during the proton–lead run was 1.6 nb^{-1} .

Since the LHCb magnet deflects positive and negative particles in opposite directions in the x – z plane, a difference in performance of the left and right sides of the detector leads to charge detection asymmetries. To reach its design sensitivity in CP violation measurements, LHCb aims to control such detection asymmetries to a precision of 10^{-3} or better. This is achieved by changing the direction of the magnetic field regularly and then combining data sets with different polarity to cancel left–right asymmetries. In Run I the polarity of the magnet was inverted about two times per month, such that smoothly varying changes in data-taking conditions or detector performance would not jeopardise the cancellation.

The LHCb operation with both field polarities leads to different effective crossing angles between the two beams, in particular when the beam crossing is performed in the horizontal plane, as it was the case in 2010 and 2011. The effective total crossing angles varied between about $40 \mu\text{rad}$ and $1040 \mu\text{rad}$ for the two spectrometer polarities. During 2012 the beam crossing was performed in the vertical plane. Together with the deflection caused by the LHCb spectrometer magnet this led to more similar total effective crossing angles of about $\pm 470 \mu\text{rad}$ in the horizontal plane for the two spectrometer polarities, respectively, and of $\pm 200 \mu\text{rad}$ in the vertical plane. However, the physics performance of the experiment has not been affected by the various beam crossing scenarios mentioned here.

2. Charged Particle Reconstruction

The trajectories of charged particles inside the LHCb detector are reconstructed using dedicated tracking detectors. The VELO detector encompassing the interaction region, the TT stations before the spectrometer magnet and the T1–T3 stations further downstream. By determining the deflection of the charged particles after traversing the magnetic field, their momentum can be determined. The high spatial resolution of the VELO enables a precise determination of the particle's flight direction close to the primary interaction point, resulting in a good vertex resolution.

2.1. Hit efficiencies and hit resolutions of the tracking detectors

The hit efficiencies and hit resolutions of the different tracking detectors are discussed in the following sections. Hit efficiencies in general exceeding 99% were achieved, more than sufficient for an efficient track reconstruction. The hit resolutions of all tracking detectors are as expected from test-beam measurements. Hit occupancies for the 2011 data taking conditions, although running at much higher luminosity and pile-up than originally planned, are well within acceptable levels, only mildly affecting the track finding efficiency and rate of wrongly reconstructed trajectories.

2.1.1. Vertex Locator

The overall performance of the VELO is described in detail in Ref. 29. A summary of the hit efficiency, hit resolution, occupancy and radiation damage given below.

The VELO hit efficiency is evaluated by two methods. The cluster finding efficiency³⁰ is determined by removing each sensor in turn from the track reconstruction, extrapolating the tracks to this sensor and searching for a hit around the intercept point. Alternatively, the channel occupancy spectra is analysed to identify strips with a substantially lower or higher number of hits than the average. The two methods are in agreement. At the end of LHC Run I, the occupancy method identified 0.6% inefficient strips and 0.02% noisy strips in the detector, these numbers are effectively identical to those at the start of operations in 2010.

The hit resolution in silicon devices depends on the inter-strip read-out pitch and the charge sharing between strips. The charge sharing varies with the operational bias voltage and the projected angle of the track. The bias voltage was 150 V throughout the physics data taking in 2010 to 2013. The projected angle³⁰ provides information on the number of strips that the particle crosses while it traverses the thickness of the silicon sensor. Initially the resolution improves with increasing angle due to the charge sharing between strips, allowing more accurate interpolation of the hit position. The optimal resolution is obtained when the tracks cross the width of one strip when traversing the 300 μm thickness of the sensor. For the VELO the optimal projected angle varies between about 7° at the lowest inter-strip pitch of 40 μm to about 18° for the largest 100 μm pitch strips, as shown in Fig. 5 (right).

Fig. 5. The VELO hit resolution as a function of the inter-strip pitch (left) evaluated with 2010 data for the R sensors. Results²⁹ are shown for two projected angle ranges and the expected resolution of a single-hit binary system is indicated for comparison. Resolution divided by pitch as a function of the track projected angle for four different strip pitches (right).

Above the optimal angle the resolution begins to deteriorate due to the fluctuations in the charge on the strips and because the signal to noise ratio on individual strips may drop below the clustering threshold.

The VELO reads out analogue pulse-height information from the strips, and this information is used offline to calculate the cluster position³⁰ using the weighted average of the strip ADC values. The resolution of the sensors is determined from the residual between the extrapolated position of the fitted track and the measured cluster position. The use of the evaluated cluster position in the track fit gives rise to a bias in the residual, for which a correction is applied.

The resolution is determined as a function of the strip pitch and of the projected angle. For each bin, the resolution is determined from the width of the fit of a Gaussian function to the distribution of the corrected residuals. The resolution is evaluated using tracks that have hits in the tracking stations behind the magnet and hence for which the momentum measurement is available. The tracks are required to have a momentum greater than 10 GeV/ c to reduce the dependence of the estimation on the multiple scattering effect, and a number of other track quality criteria are applied to reject fake tracks. The results are presented here for the R sensor. The Φ sensor results are compatible with those of the R sensor but the almost radial geometry of the Φ sensor strips means that tracks primarily have small projected angles.

The measured hit resolution has a linear dependence on the strip pitch in projected angle bins, as shown in Fig. 5 (left). The best hit precision measured is around 4 μm for an optimal projected angle of 8° and the minimum pitch of approximately 40 μm .

The detector occupancy is a key parameter in the performance of the pattern recognition and tracking algorithms. The cluster occupancy was measured to be

Fig. 6. The effective depletion voltage versus fluence for all VELO sensors up to the end of LHC Run 1 at 3.4 fb^{-1} delivered integrated luminosity.²⁹

around 0.5% in randomly triggered events and 1% in events passing the high level trigger, for data with a $\mu_{\text{vis}} = 1.7$. The pitch of the strips on the sensors increases with radius, keeping the local occupancy values to within 25% of these typical values. The occupancy from noise is negligible compared with that from particles; in the absence of circulating beams the occupancy is below 0.01%.

The proximity of the VELO sensors to the LHC pp collisions results in the sensors receiving a significant radiation dose. A study of the observed effects is available in Ref. 31. During LHC Run I the sensors have been exposed to a range of fluences up to a maximum value of 1.8×10^{14} 1 MeV neutron equivalents/ cm^2 (n_{eq}) at the radius of the inner strip of 8.2 mm.

The current drawn from a silicon sensor increases linearly with fluence. The sensor current is composed of two dominant sources, referred to as *bulk* and *surface* currents. Studying the current as a function of the temperature allows the two sources to be separated, and dedicated data is taken to allow this study to be performed. At the operational sensor temperature of approximately -7°C , the average rate of sensor current increase is $18 \mu\text{A}$ per fb^{-1} , in agreement with predictions.³¹

Dedicated data are taken around three times a year to study the charge collection of the VELO as a function of the bias voltage of the sensors. The bias voltage required to extract a fixed fraction of the maximum charge can then be determined. From this measurement the ‘effective depletion voltage’ can be determined,³¹ and this is shown as a function of fluence in Fig. 6. Each sensor contributes multiple points to this plot in each data sample as the sensors are divided in the analysis into radial regions that have received similar fluences, as denoted by the different colours in the figure. The n -bulk sensors undergo space-charge sign inversion under irradiation, and hence their depletion voltage initially reduces with irradiation. This continues until type inversion occurs, after which it increases with further irradiation. The first observation of n^+ -on- n sensor space-charge-sign-inversion at the LHC was made during 2011,³¹ occurring at a fluence of a round 15×10^{12} 1 MeV n_{eq} . The effective depletion voltage at the maximal fluence at the end of LHC Run I

was approximately 100 V, and followed the expectation. The current detector is predicted to deliver an acceptable physics performance until the end of LHC Run II with an operating voltage below 500 V.

2.1.2. Silicon Tracker

The Tracker Turicensis (TT) and the Inner Tracker (IT) are constructed from p^+ -on- n silicon microstrip detectors. The TT sensors are 9.64 cm wide, 9.44 cm long and 500 μm thick. The TT modules have read-out sectors with one, two, three or four sensors bonded together, and are arranged such that the single-sensor sectors are closest to the beam-pipe in the region with the highest flux of particles. The sensors in the IT are 7.6 cm wide, 11 cm long and are either 320 μm or 410 μm thick. Two 410 μm thick sensors are bonded together for the IT modules on either side of the beam-pipe while the modules above and below the beam-pipe use one 320 μm sensor. In total, there are 280 (336) read-out sectors with 512 (384) strips in the TT (IT).

The cluster finding efficiency of the detectors depends on the fraction of working channels and the intrinsic hit efficiency of the silicon sensor. The number of working channels is affected by problems with the read-out, and the masking of dead or noisy strips found during the calibration of the detector. The fraction of working channels varied during data taking. The luminosity-weighted average of the fraction of working channels during Run I is calculated to be 99.7% and 98.6% for the TT and IT, respectively. Repairs can be made to the TT read-out during short technical stops whereas problems with the IT read-out can only be fixed during the LHC shutdowns at the end of each year. Two read-out sectors were disabled in the IT as they could not be properly configured.

The intrinsic hit efficiency of the silicon sensors can be measured using reconstructed tracks to probe whether or not the expected hits on a track are found. The efficiency is defined as the ratio between the number of hits found and the number of hits expected for a given sector. The measurement uses daughter tracks from clean samples of $J/\psi \rightarrow \mu^+ \mu^-$ decays. The method looks for hits in a window around the intersection point between a track and each sensor on the track where a hit is expected. The tracks are required to have momentum greater than 10 GeV/ c to reduce the effect of multiple scattering. Additional cuts are placed on the track quality to minimise the effect of fake tracks on the efficiency measurement. The efficiency is calculated relative to the number of working channels, i.e. hits are not expected to be found when a channel or group of channels is disabled. The overall hit efficiency is determined to be greater than 99.7% and 99.8% for TT and IT, respectively.

The hit resolution is determined from the residuals between the measured hit position and the extrapolated track position. The unbiased residual is calculated by removing the hit from the track fit and calculating the distance between the hit and the extrapolated track position. The resolution is given by the spread of the unbiased residual distribution after correcting for the uncertainty in the track

Fig. 7. Hit resolution measured for all modules in the TT. The sector number corresponds approximately to the x -direction. The resolution improves in the outer regions of the “A-side” and “C-side” regions where there is more charge sharing due to the larger track angle. It is almost constant in the sectors in the “Central” region where the occupancy is highest. The labels X1, U, V and X2 correspond to the four detection layers arranged with an $(x - u - v - x)$ geometry in the TT box.

Fig. 8. Hit resolution measured for modules in IT1 (bottom), IT2 (middle) and IT3 (top). The sector number corresponds approximately to the x -direction. The resolution in the 1-sensor sectors in the boxes above (Top) and below (Bottom) the beam-pipe are constant. The resolution improves for the 2-sensor sectors in the A- and C-side boxes with increasing distance from beam-pipe where the track angle is typically larger. The labels X1, U, V and X2 correspond to the four detection layers arranged with an $(x - u - v - x)$ geometry in each box.

parameters. The hit resolution measured using the 2011 data is $52.6 \mu\text{m}$ for the TT and $50.3 \mu\text{m}$ for the IT. The resolution measured using the 2012 data is shown as a function of the sector number in Fig. 7 for the four TT layers and in Fig. 8 for the IT. The resolution is worse in the central regions closest to the beam-pipe where

Table 1. Summary of the hit efficiency and resolution measurements made using 2011 and 2012 data. Results are also shown for simulated events.

Detector	Measurement	2011 Data	2012 Data	2011 MC	2012 MC
TT	Hit efficiency	99.7%	99.8%	99.9%	99.9%
	Hit resolution	52.6 μm	53.4 μm	47.8 μm	48.0 μm
IT	Hit efficiency	99.8%	99.9%	99.9%	99.9%
	Hit resolution	50.3 μm	54.9 μm	53.8 μm	53.9 μm

the track angles are smallest and, consequently, where there is the least amount of charge sharing between strips.

The measurements of the hit efficiency and the resolution are summarised in Table 1 for the 2011 and 2012 data taking periods. The results are compared with the expectation from simulations for 2011 and 2012 data taking conditions, respectively. The measured hit resolutions are in agreement with those expected from simulation. The small differences observed can be partially explained by the remaining misalignment of the modules. The measured hit efficiency is well above 99% in all cases.

The particle density falls significantly as the distance from the beam-pipe is increased. The occupancy in each of the read-out sectors was estimated using a data sample containing events randomly selected after the Level-0 trigger with $\mu = 1.7$. The average occupancy in the TT varies between 1.9% for the sectors closest to be beam-pipe compared to 0.2% for the outermost modules. Similarly, the average occupancy was found to vary between 1.9% and 0.2% for the IT sectors.

2.1.3. Outer Tracker

The outer parts of the tracking stations T1–T3 are equipped with a straw-tube detector (OT).³² Charged particles traversing the straw-tubes will ionise the gas along their trajectory. The drift-times of the ionisation electrons to the wire located at the centre of the straw are measured with respect to the beam crossing signal. The distribution of the recorded drift-time, which is proportional to the distance of the particle trajectory to the wire, is shown in Fig. 9 (right). The calibration of the drift-time to distance relation³² has been done on data.

The maximum drift time in the straw tubes is about 35 ns, but to account for variations in the time-of-flight of the particles, the signal propagation time through the wire, and variations in time offset constants in the electronics, three bunch crossings are read out upon a positive L0 trigger on the first bunch crossing, corresponding to a time window of 75 ns.

During Run I, the LHC was operating predominantly in either 75 ns or 50 ns bunch spacing schemes. A short running period in 2012 with 25 ns bunch spacing was also performed, allowing a study of the detector performance under these conditions to be undertaken. The contribution from earlier and later bunch crossings

Fig. 9. Drift time distribution (left) for the modules located closest to the beam ("M8"). Drift time versus distance relation (right) where the red-dotted lines indicate the centre and the edge of the straw, corresponding to drift times of 0 and 36 ns, respectively.³²

is visible in the drift time spectrum, see Fig. 9. These additional hits from different bunch crossings increase the occupancy as shown in Fig. 10. The occupancy for the most central modules is reduced with respect to the neighbouring modules, as these modules are located further away from the beam, in the vertical direction. The straws with highest average occupancy for typical running conditions in 2011 (i.e. 50 ns bunch spacing conditions and about 1.4 visible overlapping events) amounts to about 17%. This increases to about 25% for 25 ns bunch spacing conditions with on average 1.2 overlapping events. The average pile-up conditions in 2012 were slightly different, corresponding to about 1.8 visible overlapping events, resulting in a higher multiplicity compared to 2011.

A scan of the hit efficiency as a function of the predicted distance between the expected hit and the centre of the considered straw is performed on 2011 and 2012

Fig. 10. Straw occupancy for (red) 75 ns, (black) 50 ns and (blue) 25 ns bunch-crossing spacing, for comparable pile-up conditions.³² The modules are indicated by 'M', and contain 256 straws each. The width of the module is 340 mm.

Fig. 11. Example of the OT efficiency profile as a function of the distance between the extrapolated track position and the centre of the straw for hits in the detector modules on either side of the beam-pipe (type M7).³² The vertical bars represent the edges of the straw cell.

data. An efficiency profile of the detector single cell is thus obtained, an example of which is shown in Fig. 11. The efficiency drops close to the cell edges due to two effects. First, the probability for ionisation to occur decreases for shorter path lengths inside the straw. Secondly, a fraction of the hits are positioned outside the straw volume due to the uncertainty on the track extrapolation. The average single cell efficiency for tracks in the central half of the straw, closer than 1.25 mm to the wire, amounts to 99.2%. Radiation damage could in principle lead to a decrease in signal amplitude. This was monitored during the 2011 and 2012 running periods and no degradation is observed.³³

The single hit resolution is determined by comparing the predicted hit position from the track with the hit position obtained from the drift-time. The hit under study is not used in the reconstruction of the extrapolated track, in order not to bias the resolution determination. The resulting single hit resolution is $205\ \mu\text{m}$, close to the design value of $200\ \mu\text{m}$. Only tracks with a momentum larger than $10\ \text{GeV}/c$ are used, and the residual is corrected for the uncertainty in the track parameters, caused by effects such as multiple scattering.

2.1.4. Muon system

To discriminate muons against the abundant hadronic background, muon candidates are formed from aligned hits in each of the five stations. Since LHCb aims at a trigger efficiency for muons larger than 95%, the average efficiency of each muon station must exceed 99%. To meet this stringent requirement, a redundant design was chosen for the muon chambers, consisting of four active layers per chamber for the M2–M5 stations, and two for the M1 station.²⁵ Chambers are operated with a gas gain providing a signal detection efficiency, for the logical OR of the different layers, well above the required 99%. A conflicting demand, also dictated by the L0

Fig. 12. Average cluster size in each detector region for data and simulation.³⁵ The labels refer to the stations (M1 to M5) and to the four regions with different granularity used in each station (from the innermost R1 to the outermost R4). Only isolated muon tracks are used, and angular effects are corrected for.

trigger algorithm, is to minimise cross-talk between channels. The cluster size of muon track hits, defined as the average number of adjacent pads fired by an isolated muon track, is measured using 2010 data. The result depends on the station (M1 to M5) and region (from the innermost R1 to the outermost R4), since twenty chamber types of different size and granularity are used. As shown in Fig. 12, the cluster size values observed in the data are in reasonable agreement with the simulation and are sufficient to meet the L0 trigger requirements.³⁴

Due to the redundant design, all of the 1380 muon chambers were continuously operating during the whole data taking. The few cases of a broken MWPC or GEM detector layer only caused a locally limited reduction of efficiency. Dead detector channels were only due to faulty components in the read-out chain, and never affected more than 0.2% of the total detector surface. Their effect on the muon trigger efficiency is estimated to be less than 1%. The other main sources of inefficiencies are incorrect time reconstruction and dead-time of the read-out electronics.

Since the signals must be detected within the 25 ns LHC time gate around a bunch crossing, the detector time resolution is required to be smaller than 4.5 ns. The 122,112 physical channels were aligned in time with an accuracy of 1 ns using samples of cosmic rays³⁶ and tracks from the first pp collisions in the detector.³⁵ The timing performance is measured from special calibration runs where events triggered by the calorimeters were acquired in a 125 ns wide gate around the triggered collision. A high-purity sample of muon candidates is obtained by reconstructing track segments from aligned muon detector hits in all of the five stations, and matching such segments with high-quality tracks reconstructed by the tracking detectors. The time resolution of muon detector hits associated to these tracks is measured to be between 2.5 and 4 ns, depending on the detector region.³⁵ The inefficiency

Fig. 13. Average measured hit efficiency, in percent, for the different regions of the muon detector. Statistic and systematic uncertainties are added in quadrature. The effect of the few known dead channels is not included. Measurement in the 2010 and 2011/2012 data taking periods are shown separately due to different pile-up conditions.

due to tails in the time measurement is estimated by counting the fraction of muon tracks having one or more hits outside the 25 ns time gate. In each of the time alignment runs acquired during the data taking, such inefficiency is found to never exceed 1.2%.

The total hit efficiency of the muon chambers is measured using muon candidates in events triggered independently of the muon detector during the normal data taking. The efficiency for each station is estimated by searching hits around the position predicted by the segment reconstructed using only the other four stations, which must have a good matching with a high-quality track. For station M1, which is located upstream of the calorimeter system, candidate muons are also required to originate from a J/ψ decay. Tracks close to the known dead channels are removed from the sample. The contribution of background hits accidentally matching the candidate track is subtracted using a statistical model.

The resulting efficiencies, measured separately for the twenty chamber types, and for the 2010 and 2011–2012 data taking conditions, are shown in Fig. 13. The 2010 values are compatible with the inefficiencies due to incorrect time reconstruction.³⁵ The larger inefficiency observed in 2011 and 2012 is due to the different beam conditions, with 50 ns bunch spacing and higher luminosity, causing a non-negligible dead-time of the read-out chain.

The dead-time of the front-end read-out chips varies from 50 to 100 ns, depending on the region and on the signal amplitude. This affects in particular the inner regions having the highest channel occupancy, reaching average values of 2.5% in M1R1 and 0.6% in M2R1 for the 2012 data taking. A second source of dead time is

the finite length of the digital output signals, 18 to 25 ns, depending on the region and the data taking period. In order to reduce the number of off-detector read-out channels, these signals are formed from the logical OR of several contiguous physical channels. The occupancy of these logical channels is thus larger than the occupancy of physical channels, and can lead to measurable dead-time effects, even in the outer detector regions. This happens in particular for station M5, which is affected by spurious hits due to back-scattering from the beam-line elements located behind the detector. Since the detector was operated at twice the nominal luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, the dead-time effect is larger than originally expected. Nonetheless, most regions meet the 99% efficiency requirement. Taking into account the combined response of the five stations, the detector is found to provide muon identification for trigger and offline reconstruction with an efficiency larger than 95%.

2.2. Track reconstruction

The trajectories of the charged particles traversing the tracking system are reconstructed from hits in the VELO, TT, IT and OT detectors. Depending on their paths through the spectrometer, the following track types are defined, as illustrated in Fig. 14:

- **Long tracks** traverse the full tracking system. They have hits in both the VELO and the T stations, and optionally in TT. As they traverse the full magnetic field they have the most precise momentum estimate and therefore are the most important set of tracks for physics analyses.
- **Upstream tracks** pass only through the VELO and TT stations. In general their momentum is too low to traverse the magnet and reach the T stations. However, they pass through the RICH1 detector and may generate Cherenkov photons if they have $p > 1 \text{ GeV}/c$. They are therefore also used to understand backgrounds in the particle identification algorithm of the RICH.
- **Downstream tracks** pass only through the TT and T stations. They are important for the reconstruction of long lived particles, such as K_S^0 and Λ , that decay outside the VELO acceptance.
- **VELO tracks** pass only through the VELO and are typically large-angle or backward tracks, which are useful for the primary vertex reconstruction.
- **T tracks** pass only through the T stations. They are typically produced in secondary interactions, but are still useful during the treatment of RICH2 data for particle identification.

The long track reconstruction starts with a search in the VELO for straight line trajectories.^{37,38} To be reconstructed as VELO tracks, traversing particles must provide at least three hits in the R sensors and three hits in the Φ sensors. Then, there are two complementary algorithms to add information from the downstream tracking stations to these VELO tracks. In the first algorithm, the forward tracking,³⁹

Fig. 14. A schematic illustration of the various track types:²⁵ long, upstream, downstream, VELO and T tracks. For reference the main B -field component (B_y) is plotted above as a function of the z coordinate.

the VELO tracks are combined with information from the T stations. The momentum of a particle and its trajectory through the detector are fully determined from the information provided by the VELO and a single T station hit. Further hits in the T stations are then searched along this trajectory to find the best possible combination of hits defining the long track. In the second algorithm, called track matching,^{40,41} the VELO tracks are combined with track segments found after the magnet in the T stations, using a standalone track finding algorithm.⁴² In order to form such a track segment, particles traversing the T stations need to provide at least one hit in the x layers and one in the stereo layers in each of the three stations. The candidate tracks found by each algorithm are then combined, removing duplicates, to form the final set of long tracks used for analysis. Finally, hits in the TT consistent with the extrapolated trajectories of each track are added to improve their momentum determination.

Downstream tracks are found starting with T tracks, extrapolating them through the magnetic field and searching for corresponding hits in the TT.^{43,44} Upstream tracks are found by extrapolating VELO tracks to the TT where matching hits are then added in a procedure similar to that used by the downstream tracking. At least three TT hits are required to be present by these algorithms.⁴⁵

Fig. 15. Display of the reconstructed tracks and assigned hits in an event in the x - z plane.²⁵ The insert shows a zoom into the VELO region in the x - y plane.

In a final step, the tracks are fitted using a Kalman filter.^{46,47} The fit takes into account multiple scattering and corrects for energy loss due to ionisation. The χ^2 per degree of freedom of the fit is used to determine the quality of the reconstructed track. After the fit, the reconstructed track is represented by state vectors $(x, y, dx/dz, dy/dz, q/p)$ which are specified at given z -positions in the experiment. If two or more tracks have many hits in common, only the one with most hits is kept. Figure 15 shows the tracks reconstructed in a typical event.

Mis-reconstructed (fake) tracks are those that do not correspond to the trajectory of a real charged particle. Due to the large extrapolation distance in traversing the magnet, most of these fake tracks originate from wrong associations between VELO tracks and tracks in the T stations. The fraction of fake tracks in minimum bias events is typically around 6.5%, increasing to about 20% for large multiplicity events.⁴⁸ This fake rate is significantly reduced, at the cost of a small drop in efficiency, with a neural network classifier which uses as input the result of the track fit, the track kinematics and the number of measured hits in the tracking stations versus the number of expected hits.

2.2.1. Track finding efficiency

The tracking efficiency is defined here as the probability that the trajectory of a charged particle that has passed through the full tracking system is reconstructed. In particular it does not account for interactions with the material, decays in flight and particles that fly outside of the detector acceptance.

Fig. 16. Tracking efficiency as function of the momentum, p , the pseudorapidity, η , the total number of tracks in the event, N_{track} , and the number of reconstructed primary vertices, N_{PV} .⁵⁰ The error bars indicate the statistical uncertainty.

The efficiency is measured using a tag-and-probe technique with $J/\psi \rightarrow \mu^+ \mu^-$ decays. In this method one of the daughter particles, the “tag” leg, is fully reconstructed, while the other particle, the “probe” leg, is only partially reconstructed. The probe leg should carry enough momentum information such that the J/ψ invariant mass can be reconstructed with a sufficiently high resolution. The tracking efficiency is then obtained by matching the partially reconstructed probe leg to a fully reconstructed long track. If a match is found the probe leg is defined as efficient. In the trigger and offline selection of the J/ψ candidates, no requirements are set on the particle used for the probe leg to avoid biases on the measured efficiency.

Two different tag-and-probe methods^{49,50} are used to measure the efficiency for long tracks. The overall efficiency depends on the momentum spectrum of the tracks and the track multiplicity of the event. The tracking efficiency is shown in Fig. 16 as a function of the absolute momentum, p , of the pseudorapidity, η , of the total number of tracks in the event, N_{track} , and of the number of reconstructed primary vertices, N_{PV} . The performance in the 2012 data is slightly worse, which is partially due to the higher hit multiplicity at the higher centre-of-mass energy. As can be seen, the average efficiency is above 96% in the momentum range $5 \text{ GeV}/c < p < 200 \text{ GeV}/c$ and in the pseudorapidity range $2 < \eta < 5$, which covers the phase space of LHCb. Only in high multiplicity events ($N_{\text{track}} > 200$) it is slightly less than 96%. The track reconstruction efficiency has been shown to be well reproduced in simulated events.⁵⁰

Fig. 17. Relative momentum resolution versus momentum for long tracks in data obtained using J/ψ decays.

2.2.2. Mass and momentum resolution

The momentum resolution for long tracks in data is extracted using $J/\psi \rightarrow \mu^+\mu^-$ decays. The mass resolution of the J/ψ is primarily defined by the momentum resolution of the two muons. Neglecting the muon masses and considering decays where the two muons have a similar momentum, the momentum resolution, δp , can be approximated as:

$$\left(\frac{\delta p}{p}\right)^2 = 2\left(\frac{\sigma_m}{m}\right)^2 - 2\left(\frac{p\sigma_\theta}{mc\theta}\right)^2, \quad (1)$$

where m is the invariant mass of the J/ψ candidate and σ_m is the Gaussian width obtained from a fit to the mass distribution. The second term is a correction for the opening angle, θ , between the two muons, where σ_θ is the per-event error on θ which is obtained from the track fits of the two muons. Figure 17 shows the relative momentum resolution, $\delta p/p$, as a function of the momentum, p . The momentum resolution is about 5 per mille for particles below 20 GeV/c, rising to about 8 per mille for particles around 100 GeV/c.

The mass resolution is compared for six different dimuon resonances: the J/ψ , $\psi(2S)$, $\Upsilon(1S)$, $\Upsilon(2S)$ and $\Upsilon(3S)$ mesons, and the Z^0 boson. These resonances are chosen as they share the same topology and exhibit a clean mass peak. A loose selection is applied to obtain the invariant mass distributions, as shown in Fig. 18.

The momentum scale is calibrated using large samples of $J/\psi \rightarrow \mu^+\mu^-$ and $B^+ \rightarrow J/\psi K^+$ decays, as is done for the precision measurements of b -hadron and D meson masses.^{51–54} By comparing the measured masses of known resonances with the world average values,⁵⁵ a systematic uncertainty of 0.03% on the momentum scale is obtained. As shown in Fig. 17 the momentum resolution depends on the momentum of the final-state particles, and therefore the mass resolution is not expected to behave as a pure single Gaussian. Nevertheless, a double Gaussian function is sufficient to describe the observed mass distributions. Final-state radiation creates a low-mass-tail to the left side of the mass distribution, which is modelled by an additional power-law tail. To describe the Z^0 mass distribution, a single

Fig. 18. Mass distributions for (top left) J/ψ , (top right) $\psi(2S)$, (bottom left) $\Upsilon(1S)$, $\Upsilon(2S)$ and $\Upsilon(3S)$, and (bottom right) Z^0 candidates. The shapes from the mass fits are superimposed, indicating the signal component (dotted line), the background component (dashed line) and the total yield (solid line).

Table 2. Mass resolution for the six different dimuon resonances.

Resonance	Mass resolution (MeV/c^2)
J/ψ	14.3 ± 0.1
$\psi(2S)$	16.5 ± 0.4
$\Upsilon(1S)$	42.8 ± 0.1
$\Upsilon(2S)$	44.8 ± 0.1
$\Upsilon(3S)$	48.8 ± 0.2
Z^0	1727 ± 64

Gaussian function with power-law tail is convolved with a Breit–Wigner function, where the natural width is fixed to $2495.2 \text{ MeV}/c^2$.⁵⁵ In all cases, an exponential shape models the background. The results from the fits are overlaid in Fig. 18. The overall mass resolution is calculated as the root mean square of the double Gaussian function. The mass resolution obtained from the fits are shown in Table 2. The uncertainties are statistical only. Figure 19 shows the mass resolution and relative mass resolution versus the mass of the resonance. It can be seen that the relative mass resolution, σ_m/m , is about 5 per mille up to the Υ masses.

Fig. 19. Mass resolution (σ_m) (left) and relative mass resolution (right) as a function of the mass (m) of the dimuon resonance. The mass of the muons can be neglected in the invariant mass calculation of these resonances. The mass resolution is obtained from a fit to the mass distributions. The superimposed curve is obtained from an empirical power-law fit through the data points.

2.3. Spatial alignment of the tracking detectors

The alignment of the LHCb tracking detector uses information from optical and mechanical surveys and from reconstructed charged particle trajectories. To ensure adequate tracking performance, the position and orientation of detector elements in the global reference frame must be known with an accuracy significantly better than the single hit resolution. Since LHCb is a forward spectrometer, the requirements in terms of absolute units of distance are different for the different coordinate axes: tracks are less sensitive to displacements of elements in the z direction compared to equally sized displacements in x and y . Similarly, rotations around the z axis are more important than those around the x and y axis.

Although the final alignment precision is obtained with reconstructed tracks, a precise survey is indispensable both as a starting point for the track-based alignment and to constrain degrees of freedom to which fitted track trajectories are insensitive. For example, the knowledge of the z scale of the vertex detector originates solely from the pre-installation survey. Ultimately this is what limits, for example, certain measurements such as the B_s^0 oscillation frequency.

Several methods have been deployed for track-based alignment in LHCb. One technique used for the VELO divides the alignment in three stages, corresponding to different detector granularity.^{56,57} The relative alignment of each Φ sensor with respect to the R sensor in the same module is performed by fitting an analytical form to the residuals as a function of ϕ . The relative alignment of the modules within each VELO half are obtained with a χ^2 minimisation based on an implementation of Millepede method.⁵⁸ The relative alignment of one VELO half with respect to the other half is also based on the Millepede method. It is performed using a track sample crossing the overlap region between the two halves and with a χ^2 minimisation that exploits the difference in the position between primary vertices reconstructed in both halves. Similar approaches based on Millepede have been considered for the OT⁵⁹ and IT.

The implementations of the Millepede algorithm in LHCb use a simplified model of the track, ignoring the effects of the magnetic field, multiple scattering and energy loss. These effects are accounted for in the default LHCb track fit, which is based on a Kalman filter. Therefore, another global χ^2 minimisation that uses the default track fit has been implemented.^{60,61} The algorithm can align all tracking detectors simultaneously. The correct treatment of magnetic field and material effects facilitates the use of relatively low-momentum tracks in the alignment, which helps to constrain the z scale of the spectrometer. Another novel aspect is that tracks can be combined in vertices, allowing for the use of primary vertex and mass constraints.⁶²

All methods were used during the commissioning of the detectors and in the initial pp collisions and found to be in good agreement.⁶³ The method using the Kalman track fit is used routinely for the tracking alignment updates.

2.3.1. Vertex locator alignment

The most stringent alignment requirements apply to the vertex detector. In order not to degrade impact parameter or decay time resolutions, the VELO sensors need to be aligned with a precision of a few microns in x and y and a few tens of microns in z . Components of the detector have been surveyed at various stages of the assembly at ambient temperature. The relative position of the Φ sensor with respect to the R sensor in each module has been measured with an accuracy of about $3\ \mu\text{m}$ for the x and y translation and with an accuracy of about $20\ \mu\text{rad}$ for rotations around the x and y axis. The relative module position within each half of the detector has been measured with a precision of about $10\ \mu\text{m}$ for the translations along x , y and z . The position of the two VELO halves has been determined with an accuracy of $100\ \mu\text{m}$ for the translations and $100\ \mu\text{rad}$ for the rotations.

The main degrees of freedom in the track alignment of the VELO sensors and modules are the x and y translation and the rotation around the z axis. The alignment for the x and y translation can be evaluated at the sensor level, while the one for the rotation around the z axis can be determined only at the module level, as only the Φ sensors are sensitive to this degree of freedom. The misalignment due to the other three degrees of freedom (the z translation and the rotations around the x and y axis) causes a second-order effect. To obtain the desired sensitivity, a track sample with a wide distribution of the angle between the track and the strips in the sensor plane is required. Consequently, the alignment for these degrees of freedom can be evaluated only for the R sensors.

The track-based alignment is insensitive to the overall z scale, xz and yz shearing and to the global position and orientation of the VELO.⁵⁶ To constrain these degrees of freedom the position of two modules in each half are fixed to their nominal survey position in the VELO half frame. The average position and rotation of the two halves is also fixed. After correcting for differences in temperature, the position of the modules and sensors evaluated by the alignment with tracks is found to be in good agreement with the metrology.

Fig. 20. Run dependence of the relative misalignment of the two VELO halves along the x axis evaluated with primary vertices.²⁹

The alignment of the two VELO halves relies on two constraints. The first one is determined by tracks that cross both halves of the detector, in particular those that traverse the region where the sensors in the two halves overlap. This gives sensitivity to misalignment due to x and y translations and to rotation around the z axis. In addition, reconstructed primary vertices are used which adds sensitivity to relative x , y and z translations and to rotations around the x and y axis.

The operating temperature was found to have an effect on the alignment and hence is kept sufficiently stable such that variations can be ignored. A more important issue is the fact that the VELO halves are moved every fill in order to put them at a safe distance from the beam during LHC injection. This movement corresponds to about 29 mm in x . The VELO is closed only once stable beam conditions are declared. The position of the VELO stepper motors is measured using resolvers mounted on the motor axes and is reproducible with a precision better than $10\ \mu\text{m}$. This measurement is then used as an alignment correction. Figure 20 shows the distribution of the difference between the x position of primary vertices that are separately reconstructed in the left and right detector halves as a function of time. The variation illustrates that the resolver position measurement is accurate to about $5\ \mu\text{m}$.

The uncertainty in the z scale of the VELO is important for precision measurements of b -hadron lifetimes and $B_{(s)}^0$ mixing frequencies. At the time of assembly the length of the VELO base plate was measured with an accuracy of approximately $100\ \mu\text{m}$ over the full length of the VELO.²⁹ This translates into a length scale uncertainty of about 0.01%. To verify the understanding of the survey, the measurements are compared to the track-based alignment. In the latter, the length scale is fixed by constraining two modules in each half to their nominal position. The RMS of the differences in the z positions of unconstrained modules is $20\ \mu\text{m}$, in agreement with the estimated survey uncertainty. To interpret this as a length scale,

Fig. 21. IP_x resolution as a function of $1/p_T$, comparing different qualities of alignment, measured on 2010 data.

the RMS of the distribution of the z positions of the first hits on typical VELO track segments is conservatively used. In combination with the number above this leads to a total systematic uncertainty on the length scale of 0.022%.

To illustrate the effects of misalignment on the VELO performance, the impact parameter (IP) resolution is examined (see Subsec. 2.4.2). Figure 21 shows the IP_x resolution versus $1/p_T$ obtained at different stages of the alignment, namely by using the alignment from the commissioning phase, after a track-based alignment that only corrects for the relative alignment of the two halves, and after the full alignment of the sensors. The refinement of the alignment improves the IP resolution by about 25% at high transverse momentum. As the remaining alignment uncertainties are smaller than the corrections obtained in the last stage, the residual misalignment has no significant effect on the IP resolution.

2.3.2. Alignment of the silicon tracker and outer tracker

The rest of the spectrometer is aligned relative to the VELO using long tracks. The alignment is performed at different levels of granularity, exploiting differences in the precision of survey between ‘small’ and ‘large’ structures. Typical alignment degrees of freedom are displacements in x and rotations around the z axis for the smallest structures (modules in the OT and ladders in IT and ST) and displacements in z for the layers.

Global deformations are a concern in a forward spectrometer with parallel detector planes, in particular x scaling, z scaling, xz shearing and curvature (q/p) bias. An x scaling corresponds to a displacement along x of detector modules assembled in a single layer proportional to the x coordinate. In the TT and IT detectors such a scaling is constrained by tracks that traverse neighbouring ladders in the same layer. To profit from this constraint the sample of tracks is enriched by preferentially selecting such ‘overlap’ tracks.

In principle, the z scaling of the spectrometer is fixed by the z scale of the vertex detector, which comes from a survey. In practice, this leads to a relatively poor constraint on the tracking layers downstream of the magnet. It has been verified that the last OT detector plane downstream of the magnet can be fixed to its survey position without introducing a momentum bias.

A global xz shearing can also be fixed using information from the VELO survey. Whilst this leads to a relatively weak constraint on the tracking layers downstream of the magnet, any remaining shearing between the VELO and the rest of the tracking system is absorbed in the curvature bias, which is a global deformation that is typical for alignment with tracks in a non-zero field.^{62,64} The curvature bias is constrained by including mass constraints from cleanly selected $D^0 \rightarrow K^-\pi^+$ or $J/\psi \rightarrow \mu^+\mu^-$ candidates.⁶²

Another concern are observed displacements in the z coordinate, in particular in TT and IT1, the detectors closest to the magnet. In the presence of a magnetic field, tracks are sensitive to the position of the tracking detectors relative to the dipole field. An alignment performed with early 2010 data indicated a displacement of approximately 1 cm of the entire spectrometer along the z axis. In winter 2011 an *in-situ* measurement of the magnetic field map on a finite number of points along the x , y and z axes in the centre of the magnet confirmed this displacement.

The survey of the tracking detectors was performed with the dipole magnet switched off. After anomalously large differences between survey data and track-based alignment were observed in IT1, the position of all IT boxes was monitored before and after ramping the field, revealing movements of up to 5 mm. This illustrates that data collected in the absence of magnetic field are only of limited value in the alignment.

As for the VELO a crucial aspect of the alignment is stability over time. Detectors may be moved, for example, for maintenance during accelerator technical stops. These occurred at least once every two months in the first years of LHC running. The dipole field is reversed about twice per month, which also affects alignment. Consequently, the detector is realigned after every technical stop and every magnetic field reversal. Remaining misalignments in the relative position of neighbouring detector modules are estimated from hit residual distributions to be approximately $10\ \mu\text{m}$ in IT and $30\ \mu\text{m}$ in TT.

Figure 22 shows the position of the peak of the $J/\psi \rightarrow \mu^+\mu^-$ invariant mass distribution as a function of time in a period in which the operating temperature of the TT modules was varied by 15°C in order to study detector performance. The temperature change causes the support structure on which the modules are mounted to contract by an amount that is large enough to affect the curvature measurement, as shown by the bias in the J/ψ mass. After a separate track-based alignment is performed for each period with constant temperature, the bias in the mass disappears.⁶⁵ This illustrates the importance of operating the detector under stable conditions.

Fig. 22. Fitted position of the peak of the $J/\psi \rightarrow \mu^+\mu^-$ invariant mass distribution as a function of run number in a two-week period in which the operating temperature of TT modules was varied. The mass is evaluated using the same alignment for the full period on the top and using a dedicated track alignment for each period with constant temperature on the bottom.

Although the average curvature bias is constrained with $D^0 \rightarrow K^-\pi^+$ decays, other misalignment effects, uncertainties in modelling of the magnetic field (evaluated to be about 0.1%) and detector material still affect the reconstructed invariant mass. In particular, small variations in the invariant mass as function of particle momentum are observed. To obtain precise mass measurements corrections on the momentum scale are tabulated and calibrated using samples of $J/\psi \rightarrow \mu^+\mu^-$ and $B^+ \rightarrow J/\psi K^+$ decays (see Subsec. 2.2.2).

2.3.3. Muon system alignment

The read-out pads of the muon detector are less fine-grained than the read-out channels of the other tracking systems, leading to a coarser spatial resolution. Consequently the alignment of the muon chambers is in general less critical. However, misalignments larger than a few mm in the first two muon stations can affect the efficiency of the L0 muon trigger and introduce a charge asymmetry. In the L0 trigger the muon momentum is estimated by the x coordinate of hits in stations M1 and M2. Studies on simulated events have shown that an alignment precision of 1 mm is enough to guarantee charge symmetry of the trigger efficiency and momentum measurement to the 0.1% level. The alignment is even less critical for the other stations, which are not used for momentum estimate and have lower spatial resolution.

The muon chambers are mounted in support structures called ‘half stations’. The alignment accuracy of the chambers within a half station is about 1 mm in the x and y direction. Each half station can be independently moved on rails in the

Fig. 23. Alignment of the ten muon half stations for the 2012 run.³⁵ The values of the inner edge in the x position (left) and of the median y position are shown as a function of the station position along z . The empty and full dots represent the results from the survey measurements and the software alignment respectively. The error bars, when visible, show the sum in quadrature of statistic and systematic uncertainties. The dashed lines represent the ideal projective alignment of the detector in the closed position.

x direction. Due to mechanical limitations observed during installation, the half stations could not be put exactly in their nominal position. Moreover, in order to keep a safe distance from the beam-pipe, a small separation between the two sides was maintained, preserving as much as possible the symmetry and projectivity with respect to the interaction point. The position of the closed half stations was surveyed using four reference points on each half station. The result of the survey is shown in Fig. 23. Displacements with respect to the ideal projective position are found to be within 2 mm.

The alignment obtained from the survey is refined offline using reconstructed tracks after the alignment of the tracking system. Muon hits are attached to reconstructed tracks matching standalone muon segments with a good χ^2 . The global χ^2 of an ensemble of tracks is minimised using the same method as used for the rest of the tracking system. As shown in Fig. 23 for 2012 data, small, but significant, differences are found for the translational degrees of freedom, while no significant rotations are found. The result of this procedure is confirmed, independently of the alignment of the other tracking detectors, using standalone muon segment reconstruction from events selected without the muon trigger. The resulting differences with respect to the survey positions are within 1.5 mm in x and y . The accuracy of the alignment based on tracks is 1 mm or less, sufficient to avoid any detector efficiency effects that could introduce charge asymmetries in the L0 trigger. The alignment results are thus used by the L0 muon trigger for the computation of transverse momenta and are accounted for in the subsequent offline reconstruction of muons.

Fig. 24. Decay time distribution for $B_s^0 \rightarrow D_s^- \pi^+$ candidates tagged as mixed (different flavour at decay and production; red, continuous line) or unmixed (same flavour at decay and production; blue, dotted line).⁶⁶

2.4. Vertexing and decay time resolution

The study of CP violation and rare decays in the heavy flavour sector requires the accurate measurement of production and decay vertices and track impact parameters, both for flavour tagging and for background rejection. The most stringent demands on the vertex reconstruction arise from the decay time resolution requirements to resolve the fast flavour oscillations induced by $B_s^0 - \bar{B}_s^0$ mixing. LHCb has made the world's most precise measurement of the B_s^0 oscillation frequency using the decay $B_s^0 \rightarrow D_s^- \pi^+$.⁶⁶ The decay time resolution in LHCb is sufficient to observe the oscillations in the flavour tagged decay time distribution, as illustrated in Fig. 24.

2.4.1. Primary vertex reconstruction

The primary vertex (PV) resolution is measured by comparing two independent measurements of the vertex position in the same event. This is achieved by randomly splitting the set of tracks in an event into two and reconstructing the PVs in both sets. The width of the distribution of the difference of the vertex positions is corrected for a factor $\sqrt{2}$ to extract the vertex resolution. The number of tracks making a vertex ranges from 5 (the minimum required by the PV reconstruction) to around 150, and this technique allows the resolution to be measured using up to around 65 tracks. The PV resolution is strongly correlated to the number of tracks in the vertex (the track multiplicity). To determine the vertex resolution as a function of the track multiplicity, only vertex pairs with exactly the same number of tracks are compared. The result for the resolution in the x and y direction is shown in Fig. 25. A PV with 25 tracks has a resolution of $13 \mu\text{m}$ in the x and y coordinates and $71 \mu\text{m}$ in z .

Fig. 25. The primary vertex resolution (left), for events with one reconstructed primary vertex, as a function of track multiplicity. The x (red) and y (blue) resolutions are separately shown and the superimposed histogram shows the distribution of number of tracks per reconstructed primary vertex for all events that pass the high level trigger. The impact parameter in x resolution as a function of $1/p_T$ (right). Both plots are made using data collected in 2012.²⁹

2.4.2. Impact parameter resolution

The impact parameter (IP) of a track is defined as its distance from the primary vertex at its point of closest approach to the primary vertex. Particles resulting from the decay of long lived B or D mesons tend to have larger IP than those of particles produced at the primary vertex. Selections on IP and $IP \chi^2$ are extensively used in LHCb analyses to reduce the contamination from prompt backgrounds. Consequently, an optimal IP resolution and a good understanding of the effects contributing to the IP resolution are of prime importance to LHCb performance.

The IP resolution is governed by three main factors: multiple scattering of particles by the detector material; the resolution on the position of hits in the detector from which tracks are reconstructed; and the distance of extrapolation of a track between its first hit in the detector and the interaction point. The minimisation of these factors is achieved in the design of the VELO. The sensors are positioned close to the beams, separated from them by only a thin aluminium foil. The first active strips are only 8 mm away from the beams during physics collisions. The detector provides high-precision hit position measurements as shown in Subsec. 2.1.1.

As the IP is defined as the distance between a point and a line, it is not a Gaussian distributed quantity. It is therefore customary to divide the IP in two quantities that follow a normal distribution by projecting out two independent components. In LHCb these are the components of the IP vector in the transverse plane,

$$IP_x = x - x_{PV} - (z - z_{PV})t_x \quad (2)$$

and similarly for y , where $(x, y, z)_{PV}$ is the position of the primary vertex, (x, y, z) is the point on the track of closest approach to the primary vertex and $(t_x, t_y, 1)$ is the direction vector of the track. Figure 25 shows the IP_x resolution as a function $1/p_T$. The IP_y resolution is similar. The linear dependence on $1/p_T$ is a consequence

of multiple scattering and the geometry of the vertex detector. At asymptotically high p_T the IP_x resolution is about $13\ \mu\text{m}$.²⁹

2.4.3. Decay time resolution

The distance between the production and secondary decay vertices of long lived mesons is used to reconstruct the particle's decay time. This is required for lifetime measurements and for resolving flavour oscillations in time-dependent CP violation measurements. Consequently, the performance of the VELO is illustrated here with an analysis of the decay time resolution of $B_s^0 \rightarrow J/\psi \phi$ decays.⁶⁷

Time dependent CP violation effects are measured as the amplitude of an oscillation in the B decay time distribution. The size of the observed amplitude is damped by a dilution factor from the finite decay time resolution.⁶⁸ Hence, achieving optimal decay time resolution is important and any bias in the estimated decay time resolution leads to a bias in the measurement of the CP violating effect.

The reconstructed decay time in the rest frame of the decaying particle can be expressed in terms of the reconstructed decay length l , momentum p and mass m of the particle in the LHCb frame as $t = ml/p$. The decay time is computed with a vertex fit that constrains the decaying particle to originate from the primary vertex. The uncertainty on the decay length l and on the momentum p are essentially uncorrelated in LHCb. Consequently, the decay time uncertainty can be expressed in terms of the decay length uncertainty σ_l and the momentum uncertainty σ_p as

$$\sigma_t^2 = \left(\frac{m}{p}\right)^2 \sigma_l^2 + \left(\frac{t}{p}\right)^2 \sigma_p^2. \quad (3)$$

This expression shows an explicit dependence on the decay time. However, for decay times up to a few times the B meson lifetime, the uncertainty is dominated by the σ_l term, motivating the use of a 'prompt' control channel to calibrate the decay time uncertainty. The decay time resolution depends on the topology of the decay and is calibrated for each final state on data. For $B_s^0 \rightarrow J/\psi \phi$ decays, the calibration method uses prompt combinations that fake signal candidates. Subtracting the small contribution from signal candidates and long-lived background using the *sPlot* technique,⁶⁹ the shape of the decay time distribution is determined only by the resolution function.

Figure 26 shows the resolution as a function of the (fake) B candidate momentum. It should be noted that the decay time resolution is essentially independent of the B momentum, illustrating that $\sigma_l \propto p$. This is a consequence of the fact that the larger the momentum is, the smaller the opening angle, and hence the larger the uncertainty on the position of the vertex in the direction of the boost. The resolution is also shown as a function of the per-event estimated uncertainty in the decay time, which is obtained from the vertex fit. As expected, the resolution is a linear function of the estimated uncertainty. A decay time resolution of $\sim 50\ \text{fs}$ is obtained in LHCb. For a mixing frequency of $17.7\ \text{ps}^{-1}$, such as for B_s^0 oscillations, this decay time resolution leads to a dilution of the CP asymmetry by a factor ~ 0.7 .

Fig. 26. Decay time resolution as a function of momentum (left) and as a function of the estimated decay time uncertainty (right) of fake, prompt $B_s^0 \rightarrow J/\psi \phi \rightarrow \mu^+ \mu^- K^+ K^-$ candidates in 2011 and 2012 data.²⁹ Only events with a single reconstructed primary vertex are used. The superimposed histogram shows the distribution of momentum (left) and estimated decay time uncertainty (right) on an arbitrary scale.

Fig. 27. Distribution of the invariant mass of $K_S^0 \rightarrow \pi^+ \pi^-$ candidates with a decay vertex at a significant distance to the PV, for long tracks (left) and downstream tracks (right). A mass resolution of $3.5 \text{ MeV}/c^2$ is achieved for the candidates reconstructed from long tracks and $7 \text{ MeV}/c^2$ for those using downstream tracks.

2.4.4. V^0 reconstruction

Reconstructed V^0 decays ($K_S^0 \rightarrow \pi^+ \pi^-$ and $\Lambda \rightarrow p \pi^-$) are an essential ingredient of many LHCb analyses. If the decay time is sufficiently small, the daughter particles are reconstructed as long tracks, and for these decays the invariant mass resolution is as good as for short-lived resonances (see Subsec. 2.2.2 and Fig. 27). For V^0 's that decay outside the VELO acceptance, but before the magnet, the daughter particles are reconstructed as downstream tracks from hits in the TT and T stations. As the resolution on the track direction reconstructed in the layers of the TT is not as good as in the VELO, the invariant mass resolution for the downstream category is worse than for the short-lived category, as shown in Fig. 27. For K_S^0 momenta typical

of B decay products, about two thirds of the reconstructed K_S^0 decays are found using downstream tracks, illustrating the importance of the downstream tracking for physics performance.

3. Neutral Particle Reconstruction

Neutral particle reconstruction is based on information provided by the four systems (SPD, PS, ECAL and HCAL), which together form the calorimeter. The SPD and the PS both consist of a plane of scintillator tiles, separated from each other by a thin lead layer, while the ECAL and HCAL have shashlik and sampling constructions, respectively. In all four cases, the light produced in the organic scintillators is transmitted to photomultiplier tubes (PMT) by optical fibres.^{25,70} In general, the detected signal pulses are longer than the nominal read-out window of 25 ns, and this must be taken into account to minimise spill-over effects. In the ECAL and HCAL detectors, this is performed by first clipping the signal to fit within the read-out window. In the PS and SPD detectors, the effects of spill-over are removed by subtracting a fraction of the signal integrated in the previous clock cycle.

The SPD uses a single bit for each cell to indicate whether or not it was traversed by a charged particle, with a discriminator comparing the energy deposited in the given cell to half of that expected from a minimum ionising particle (MIP). The signal from the PS detector is digitised using a 10-bit ADC with a dynamic range of 0.1–100 times the corresponding MIP energy deposit. The ECAL and the HCAL have the same read-out electronics, which digitises signals with a 12-bit precision and a dynamic range that results in a maximum detectable transverse energy of 10 GeV, optimised for the typical energy deposits that occur in LHCb events. The operational status and stability of the ECAL and the HCAL detectors are examined using dedicated systems based on light emitting diodes (LEDs). The average PMT responses when illuminated by pulses of known intensities are used to monitor the behaviour of the corresponding read-out channels.

Reliable neutral particle reconstruction requires calibration^{71,72} of the calorimeter system, which is therefore outlined briefly before describing the methods used to reconstruct neutral particles

3.1. Calibration of the calorimeter system

The SPD calibration is performed by adjusting the discriminator threshold for each channel to half of the expected MIP energy deposit. These values are established at the beginning of each data taking period, and the fraction of tracks pointing to SPD cells that have an associated SPD hit is used to monitor performance during data-taking.⁷³ Thresholds are adjusted to ensure high and uniform efficiency.

The PS calibration consists of equalising the ADC response of all channels, based on the most probable values of MIP energy deposits. The MIP sample is composed of reconstructed tracks of particles with momentum greater than 2 GeV/ c to ensure that they reach the PS.⁷⁴ As for the SPD, the calibration is established at the

Fig. 28. Invariant mass distribution for $\pi^0 \rightarrow \gamma\gamma$ candidates upon which the fine calibration algorithm is applied. The red curve corresponds to the distribution before applying the method, while the blue curve is the final one. Values in the red (blue) box are the mean and sigma of the signal peak distribution in MeV/c^2 before (after) applying the fine calibration method.

beginning of each data-taking period, checked regularly and corrected as necessary during the run.

The calibration of the ECAL requires two main steps: the first to set an initial calibration, and the second to refine the values through an iterative procedure. The first step was performed at the beginning of LHCb commissioning by using test-beam measurements to reproduce the design energy range. With these settings, a $\pi^0 \rightarrow \gamma\gamma$ mass resolution at the level of 10% was achieved for the first collisions in 2009. In subsequent years, initial settings were obtained using the LED system, achieving a similar accuracy of 8–10%. In the second step, initial calibration parameters are refined by studying the energy deposited over many events and requiring continuity across cell boundaries.⁷⁵ The calibration constants are improved by performing fits to the invariant mass distribution of $\pi^0 \rightarrow \gamma\gamma$ decays,⁷⁶ combining a photon hitting the cell to be calibrated with another reconstructed photon. The procedure is repeated until all coefficients are stable. Figure 28 shows the change in the fitted π^0 invariant mass distribution, before and after the calibration procedure. By applying the method to a sample of miscalibrated simulated events, the final precision of the cell-to-cell intercalibration is estimated to be approximately 2%. In order to apply the π^0 mass fit to every cell in the ECAL, several hundred thousand events are required.

The calibration of the HCAL uses two ^{137}Cs sources of $\sim 10 \text{ mCi}$, one per detector side. This procedure takes about an hour during which they are transported through all of the scintillator cells by a hydraulic system. The response of the PMTs is measured by a dedicated system of current integrators. The relationship between the integrated anode current and the particle energy was measured in test-beam

Fig. 29. Ratio between the rate of events triggered by the L0-hadron trigger, based on the HCAL, and Muon-based triggers. Distinct increases in rate, e.g. as at a recorded luminosity of around 50 pb^{-1} , correspond to the application of a new set of PMT gains.

and is used to set the values for the HCAL parameters, obtaining a cell-to-cell intercalibration at the level of 5%. The use of the *in-situ* source limits the calibration procedure of the HCAL to technical stops, which occur bi-monthly.

The performance of the ECAL and the HCAL is monitored during the data-taking periods using the built-in LED system: the response to the LED flashes is found to agree well with the response to actual particles. In addition, the distribution of E/pc , for electrons in the case of the ECAL and hadrons in the case of the HCAL, can be compared to simulations and is used for monitoring purposes.

While no significant degradation of the SPD and the PS performance has been seen, ageing has been observed in both the ECAL and the HCAL. The main cause of this is a decrease of the PMT gains due to the degradation of the dynode system at high integrated anode current. The gain losses are compensated by increasing the voltage between the PMT dynodes. The time period over which the gradual gain reduction takes place is shorter than the interval between which absolute calibrations can be performed for the ECAL (using π^0 mass reconstruction) and the HCAL (using the ^{137}Cs sources). Therefore, the relative corrections required to maintain performance are estimated using dedicated procedures. For the ECAL, the E/pc ratio is used whilst the HCAL is monitored using the LED pulse system. The change in performance with time is illustrated in Fig. 29, which shows the ratio between the L0-hadron trigger rate, based on the HCAL hardware, and a combination of muon-based triggers. As a result of the ageing, a gradual reduction of the rate is observed, with intermittent increases, corresponding to the application of a new set of PMT gains.

3.2. Selection of neutral energy deposits in ECAL

ECAL cells with energy deposits are grouped together to form clusters by applying a 3×3 cell pattern around local energy deposition maxima. Consequently, the centres of the reconstructed clusters are always separated by at least one cell. If a

given cell is present in more than one cluster, its energy is redistributed between the clusters under consideration according to the total energies of the clusters. Although this process is iterative, it converges rapidly because the effective Molière radius (3.5 cm) of the ECAL is smaller than the size of the cells, which have lengths of side of 4.04 cm, 6.06 cm and 12.12 cm for the inner, middle and outer regions, respectively. Each cluster is characterised by its energy-weighted moments up to second order, namely the total energy, the energy-weighted position and the two-dimensional energy spread matrix.

Clusters corresponding to energy deposits of neutral particles are identified as those without an associated charged track. This is done using the procedure summarised below. First, all reconstructed tracks in the event are extrapolated to the calorimeter. Next, all pairwise combinations of extrapolated tracks and reconstructed clusters are formed. The matching between tracks and clusters is evaluated using the χ^2_{2D} metric,

$$\chi^2_{2D} = (\vec{r}_{tr} - \vec{r}_{cl})^T (\mathcal{C}_{tr} + \mathcal{S}_{cl})^{-1} (\vec{r}_{tr} - \vec{r}_{cl}), \quad (4)$$

where $\vec{r}_{tr}$ and $\vec{r}_{cl}$ represent the local coordinates of tracks and clusters, respectively, at the z barycentres of clusters, $\mathcal{C}_{tr}$ is the covariance matrix of the $\vec{r}_{tr}$, and $\mathcal{S}_{cl}$ is the cluster energy spread matrix. The z barycentre of a cluster is the average energy-weighted position of clusters in z , corrected assuming logarithmic energy dependence. A cluster generated by a neutral particle is considered to be isolated, and hence a photon candidate, if it has a minimum value of χ^2_{2D} with respect to any extrapolated track of at least 4. This cut significantly suppresses the clusters due to other charged particles while keeping high efficiency for photons.⁷⁷

3.3. Photon reconstruction

The photon energy is determined from the total cluster energy in the ECAL and the reconstructed energy deposit in the PS.⁷⁷ The photon direction is derived from an assumed origin for the photon and the energy-weighted position of the photon candidate: the transverse profile is corrected for the spread of the cluster, and the z barycentre calculated as for the χ^2_{2D} matching.

The performance of high-energy photon reconstruction is illustrated by the reconstructed $B^0 \rightarrow K^{*0} \gamma$ mass distribution shown in Fig. 30. The mass resolution obtained for this radiative decay is dominated by the ECAL energy resolution and is found to be 93 MeV/ c^2 .⁷⁸ A comparison of the data with simulated samples shows that this corresponds to an accuracy of the cell-to-cell intercalibration of around 2%.

3.4. Neutral pion reconstruction

Neutral pions with low transverse momenta are mostly reconstructed as pairs of well-separated photons (resolved π^0 candidates). A mass resolution of 8 MeV/ c^2

Fig. 30. Mass distribution of reconstructed $B^0 \rightarrow K^{*0}(K^+\pi^-)\gamma$ candidates obtained in the 2011 data sample. The blue curve corresponds to the mass shape fit. The $K^{*0}\gamma$ signal (green dotted line) and the various background contaminations are shown.⁷⁸

is obtained for such neutral pions. However, due to the finite ECAL granularity, photon pairs from the decay of sufficiently high momentum π^0 cannot be resolved as individual clusters. This essentially holds for all π^0 meson decays with transverse momentum above 2 GeV/c. To reconstruct such ‘merged’ π^0 candidates, a procedure has been designed to identify overlapping clusters. The algorithm consists of splitting each single ECAL cluster into two 3×3 subclusters built around the two highest energy deposits of the original cluster. The energy of the common cells is then distributed between the two assumed subclusters by fitting the energy distribution with that of two photons, using the expected transverse profile obtained from simulations. Since the position of the two subcluster barycentres is a function of the energy distribution, this procedure requires an iterative process.

The performance of neutral pion reconstruction is illustrated in Fig. 31, which shows the invariant mass distribution for $D^0 \rightarrow K^-\pi^+\pi^0$ candidates for resolved

Fig. 31. Mass distributions of reconstructed $D^0 \rightarrow K^-\pi^+\pi^0$ candidates with resolved π^0 (left) and merged π^0 (right). Both are obtained from the 2011 data sample. The overall mass fit⁷⁹ is represented by the blue curve, with the signal (red dashed line) and background (green dash-dotted line and purple dotted lines) contributions also shown.

and merged π^0 candidates.⁷⁹ In this example, the estimated invariant mass resolution is 30 MeV/ c for the merged π^0 candidates and 20 MeV/ c for the resolved ones.

4. Particle Identification

Particle identification in LHCb is provided by four different detectors: the calorimeter system, the two RICH detectors and the muon stations. In the following sections the performance of the individual sub-systems is presented first, followed by a description of the methods used to combine the information for charged particles in a single set of variables that provide optimal particle identification performance.

4.1. Calorimeter system based particle identification

The main role of the calorimeters in terms of particle identification is to provide for the recognition of photons, electrons and π^0 candidates. Distinguishing charged from neutral particles is performed by studying the presence or absence of tracks in front of the energy deposits using the techniques described in Subsec. 3.2. For energy deposits related to neutral particles, the shape of the cluster is used to distinguish between photons and π^0 candidates. The photon hypothesis is established by taking into account the possibility that photons convert when interacting with the detector material upstream of the calorimeter. When an energy deposit corresponds to a charged particle, the electron hypothesis is constructed to distinguish electrons from hadrons. Outline descriptions of how the photon and electron hypotheses are built are given below.

4.1.1. Photon and merged π^0 identification

Two independent estimators are built to establish the photon hypothesis, one each for the converted and non-converted candidates. Non-converted photons are identified by computing a photon hypothesis likelihood from the signal and background probability density functions of several variables, namely the PS energy deposited in front of the ECAL cluster cells, the matching estimator χ^2_{2D} between the cluster and any track defined for a charged particle, and the ratio between the energy of the central cell of the ECAL cluster and the total ECAL energy. Because the non-converted photon estimator depends on the energy and the calorimeter zone, several probability density functions are constructed from simulations, for both signal and background, corresponding to each of the zones. The difference in log-likelihood between the photon and the background hypotheses ($\Delta \log \mathcal{L}$) is calculated and used to identify photons. Figure 32 shows the performance of the photon identification in terms of the efficiency and purity obtained for candidates with $p_T > 200$ MeV/ c .

To avoid the misidentification of photons with high- E_T merged π^0 candidates, the difference between the distribution of the expected energy deposit of a photon with respect to that of a π^0 is used. This difference is evaluated by a neural network

Fig. 32. Performance of the photon identification. Purity as a function of efficiency for (green) the full photon candidate sample, (blue) converted candidates according to the SPD information and (red) non-converted candidates (left). Photon identification efficiency as a function of π^0 rejection efficiency for the $\gamma - \pi^0$ separation tool for simulation, the red curve, and data, the blue curve (right).

classifier (specifically, a multi-layer perceptron) trained with photons from a simulated $B^0 \rightarrow K^{*0}\gamma$ sample as signal and π^0 mesons from a mixture of B decays as background. These π^0 mesons are reconstructed and selected as photons using the same $B^0 \rightarrow K^{*0}\gamma$ preselection used for the signal sample. A photon identification efficiency of 95% can be obtained while rejecting 45% of the merged π^0 meson background that are reconstructed as photons. Figure 32 shows the photon identification efficiency with respect to the π^0 rejection efficiency for simulation and data.

Photons converted before the magnet are reconstructed from electron–positron tracks. The electron (positron) is selected on the basis of its electron PID variables and electron confidence level, requiring a minimum p_T value and an E/pc value within a selected range. The algorithm only combines electron–positron pairs for which the associated clusters have energy-weighted positions that are closer than 3σ of cluster extent (and 200 mm) in the vertical plane, at the average z barycentre of clusters. Pairs are selected on the basis of their transverse momenta, their di-electron masses and their reconstructed vertex positions. The electron energy is corrected by including any bremsstrahlung photons measured by the calorimeters that are compatible with the electron–positron pair.

Figure 33 shows the ratio of photon detection efficiencies between converted and non-converted photons coming from the decay of π^0 mesons for both simulation and data. The simulation provides a good description of the photon reconstruction efficiency implying that the detector material where the conversions occur is modelled well, and that the reconstruction algorithms work equally well in data and simulation. The level of performance is illustrated by analyses that benefit from the good resolution obtained using converted photons, such as $\chi_c \rightarrow J/\psi\gamma$ (Ref. 80) or $\chi_b \rightarrow \Upsilon\gamma$ (Ref. 81). In the case of the χ_c , for instance, the resolution on the mass difference $\Delta M = M(\mu^+\mu^-\gamma) - M(\mu^+\mu^-)$ is about 5 MeV/ c^2 . With this resolution, the χ_{c0} , χ_{c1} and χ_{c2} states can be disentangled from one another.⁸⁰

Fig. 33. Ratio of photon detection efficiencies $\epsilon(\gamma \rightarrow e^+e^-)/\epsilon(\gamma_{\text{CALO}})$ from the decay of π^0 mesons in data (red) and simulations (blue).

4.1.2. Electron identification

The identification of electrons in the calorimeter system uses information derived from the ECAL, the PS and the HCAL. The procedure to combine these different sources of information is based on signal and background likelihood distributions constructed for each sub-detector. In each case, reference histograms correlating the energy measurement with the particle momentum are produced. For example, Fig. 34 shows the E/p_c distribution in the ECAL for electrons and hadrons, produced using the first 340 pb^{-1} recorded in 2011. The electron distribution has been produced using reconstructed electrons from photon conversions and the hadron distribution using pions and kaons from D^0 meson decays. From these distributions, the log-likelihood difference between electrons and hadrons is derived.

Fig. 34. Distribution for the ECAL of E/p_c for electrons (red) and hadrons (blue), as obtained from the first 340 pb^{-1} recorded in 2011.

Fig. 35. Electron identification efficiency versus misidentification rate.

For the ECAL, the log-likelihood difference for electron and hadron hypotheses $\Delta \log \mathcal{L}^{\text{ECAL}}(e - h)$ is computed based on both E/p_c and the χ^2_{2D} estimator defined in Subsec. 3.2. The electron hypothesis likelihoods for the PS, $\Delta \log \mathcal{L}^{\text{PS}}(e - h)$ and the HCAL $\Delta \log \mathcal{L}^{\text{HCAL}}(e - h)$ are built using the energy deposits in each sub-detector. A combined estimator is then formed for the calorimeter system by taking the sum of the individual estimators from the PS, the ECAL and the HCAL,

$$\begin{aligned} \Delta \log \mathcal{L}^{\text{CALO}}(e - h) &= \Delta \log \mathcal{L}^{\text{ECAL}}(e - h) + \Delta \log \mathcal{L}^{\text{HCAL}}(e - h) \\ &+ \Delta \log \mathcal{L}^{\text{PS}}(e - h). \end{aligned} \quad (5)$$

Figure 35 shows the combined electron identification efficiency defined above versus the misidentification rate obtained by varying the selection criteria applied to the likelihood difference.

The electron identification performance is evaluated using the data recorded in 2011, which are sufficient for it to be measured using a tag-and-probe method. This is applied to $B^\pm \rightarrow J/\psi K^\pm$ candidates with $J/\psi \rightarrow e^+e^-$, where one of the electrons is required to be identified by its electron ID (e_{tag}) while the second electron is selected without using any information from the calorimeter system (e_{probe}). This second electron is then used to estimate the efficiency of the electron ID.

The efficiency and the misidentification rate as a function of the e_{probe} momentum are presented in Fig. 36 for several cuts on $\Delta \log \mathcal{L}^{\text{CALO}}(e - h)$. The electron identification efficiency is observed to be lower for $p < 10$ GeV/c. As expected, the higher momenta particles have higher misidentification rates as illustrated in Fig. 36. To quantify the typical identification performance of the entire calorimeter system, the average identification efficiency of electrons from the $J/\psi \rightarrow e^+e^-$ decay in $B^\pm \rightarrow J/\psi K^\pm$ events is $(91.9 \pm 1.3)\%$ for a misidentification rate of $(4.54 \pm 0.02)\%$ after requiring $\Delta \log \mathcal{L}^{\text{CALO}}(e - h) > 2$.

Fig. 36. Electron identification performances for various $\Delta \log \mathcal{L}^{\text{CALO}}(e-h)$ cuts: electron efficiency (left) and misidentification rate (right) as functions of the track momentum.

Fig. 37. $\Delta\theta_C$ distributions for the RICH 1 gas (top left), RICH 2 gas (top right) and Aerogel (bottom).⁸²

4.2. RICH system based particle identification

The primary role of the RICH system is the identification of charged hadrons (π , K , p). The information provided is used both at the final analysis level, and as part of the software trigger (see Sec. 5). In addition, the RICH system can contribute to the identification of charged leptons (e , μ), complementing information from the calorimeter and muon systems, respectively.

4.2.1. Cherenkov angle resolution

One of the primary measures of the RICH performance is $\sigma(\theta_C)$, the resolution of the Cherenkov angle with which the photons, radiated from the particles as they traverse the various radiator volumes, can be reconstructed. The distributions for $\Delta\theta_C$, the difference between the reconstructed and expected photon Cherenkov angles, are shown in Fig. 37 for 2011 data, after all detector alignment and calibration procedures have been performed.⁸² The expected Cherenkov angles for each track are calculated using reconstructed momenta and radiator refractive index information. Only high-momentum tracks are selected, to ensure that the Cherenkov angle is close to saturation.

The values of $\sigma(\theta_C)$, extracted from a simple fit to the $\Delta\theta_C$ distributions, are determined to be 1.618 ± 0.002 mrad for RICH1 gas (C_4F_{10}) and 0.68 ± 0.02 mrad for RICH2 (CF_4), comparable with the expectations from simulation of 1.52 ± 0.02 mrad and 0.68 ± 0.01 mrad respectively. The disagreement seen between data and simulation for C_4F_{10} are largely attributed to imperfect corrections for distortions in the RICH photon detector images caused by the residual magnetic field in the vicinity of the RICH1 detector. Enhancements to the procedures used to compute these corrections are foreseen for Run II, thus improving the resolutions achieved in data.

For the RICH1 aerogel radiator, where the distribution is not symmetric, the standard deviation is estimated to be 5.6 mrad. This value is about a factor of 1.8 larger than the expectation from simulation. This discrepancy is, at least partially, explained by the unmodelled absorption of C_4F_{10} gas by the very porous aerogel radiator, with which it is in contact.

Due to the high average track multiplicity in LHCb events, a reconstructed Cherenkov ring will generally overlap with several neighbouring rings. Solitary rings from isolated tracks, where no overlap is found, provide a useful test of the RICH performance, since isolated rings can be cleanly and unambiguously associated with a single track. Figure 38 shows the Cherenkov angle as a function of particle momen-

Fig. 38. Reconstructed Cherenkov angle for *isolated* tracks, as a function of track momentum in the C_4F_{10} radiator.⁸² The Cherenkov bands for muons, pions, kaons and protons are clearly visible.

Table 3. Comparison of photoelectron yields (N_{pe}) determined from $D^{*+} \rightarrow D^0\pi^+$ decays in simulation and data, and $pp \rightarrow pp\mu^+\mu^-$ events in data.

Radiator	N_{pe} from data		N_{pe} from simulation $D^0 \rightarrow K^-\pi^+$
	$D^0 \rightarrow K^-\pi^+$	$pp \rightarrow pp\mu^+\mu^-$	
Aerogel	5.0 ± 3.0	4.3 ± 0.9	8.0 ± 0.6
C ₄ F ₁₀	20.4 ± 0.1	24.5 ± 0.3	28.3 ± 0.6
CF ₄	15.8 ± 0.1	17.6 ± 0.2	22.7 ± 0.6

tum using information from the C₄F₁₀ radiator for isolated tracks selected in data ($\sim 2\%$ of all tracks). As expected, the events populate distinct bands according to their mass.

4.2.2. Photoelectron yield

The average number of detected photons for each track traversing the Cherenkov radiator media, called the photoelectron yield (N_{pe}), is another important measure of the performance of a RICH detector. The yields for the three radiators used in LHCb are measured in data using two different samples of events.⁸² The first sample is representative of normal LHCb data taking conditions, and consists of the kaons and pions originating from the decay $D^0 \rightarrow K^-\pi^+$, where the D^0 is selected from $D^{*+} \rightarrow D^0\pi^+$ decays. The second sample consists of low detector occupancy $pp \rightarrow pp\mu^+\mu^-$ events, which provide a clean track sample with very low background levels. In both samples, only high-momentum tracks are selected, to ensure that the Cherenkov angle is close to saturation.

Table 3 shows the results of the photoelectron yield extraction, performed on both real and simulated data. In data, the $D^{*+} \rightarrow D^0\pi^+$ events have values of N_{pe} that are less than those for $pp \rightarrow pp\mu^+\mu^-$ events. This is mainly due to the higher charged track multiplicities of the $D^{*+} \rightarrow D^0\pi^+$ events, reducing the effective N_{pe} , and the track geometry cut that is applied to the $pp \rightarrow pp\mu^+\mu^-$ events increasing their N_{pe} yield. The aerogel N_{pe} data values have a large uncertainty due to the significant background levels in the $\Delta\theta_C$ distributions and the additional uncertainty in the shape of the signal peak.⁸²

The photoelectron yields for data are lower than those predicted by the simulation. One reason for this is a small detector read-out inefficiency, which was identified during high trigger rate data taking in Run I. The results presented include a retuning of the read-out settings, applied during data taking to minimise the impact of the inefficiency. A further optimisation will be performed for LHC Run II to reduce the effect to the negligible level, and is expected to improve the yields further by a few percent. The remaining discrepancy is accounted for by an over-estimate of the yield in the simulation, which will be addressed by improved simulation tunings. It must be stressed however, that the smaller yield measured in data does not have

a significant impact on the final particle identification performance, as described in Subsec. 4.2.3.

4.2.3. Particle identification performance

To determine the RICH particle identification performance on data, large samples of genuine π , K and p tracks are required. Such control samples must be selected independent of RICH information which would otherwise bias the results. The strategy employed is to reconstruct exclusive decays purely from kinematic selections. Only decay modes with large branching fractions, for which large samples can be easily collected, are used to allow for precise calibration over a range of track kinematics.

The following decays, and their charge conjugates, are identified: $K_S^0 \rightarrow \pi^+\pi^-$, $\Lambda \rightarrow p\pi^-$ and $D^{*+} \rightarrow D^0(K^-\pi^+)\pi^+$. This ensemble of final states provides a complete set of charged particle types needed to assess comprehensively the hadron PID performance. Utilising the track samples obtained from these exclusive control decay modes, Fig. 39 demonstrates the kaon efficiency (kaons identified as kaons) and pion misidentification (pions misidentified as kaons) fraction achieved in LHCb data, as a function of momentum. For illustration the data is shown with two different PID requirements, one optimising the efficiency, the other minimising the misidentification rate.

For each track the likelihood that it is an electron, muon, pion, kaon or proton is computed. In the first approach it is required that, for each track, the likelihood for the kaon mass hypothesis is larger than that for the pion hypothesis, i.e. $\Delta \log \mathcal{L}(K - \pi) > 0$. When averaging over the momentum range 2–100 GeV/ c one finds the kaon efficiency to be $\sim 95\%$ with a pion misidentification rate of $\sim 10\%$. A stricter PID requirement, $\Delta \log \mathcal{L}(K - \pi) > 5$, reduces the pion misidentification rate to $\sim 3\%$ at a modest loss in kaon efficiency of $\sim 10\%$ on average. Figure 39 also shows the performance in simulation, for the same exclusive control channels and PID requirements as above for data. Good agreement with data is observed for both sets of PID requirements.

The Run I conditions, with multiple interactions per bunch crossing and the resulting high particle multiplicities, provide an insight into the RICH performance at possible future higher luminosity running. Figure 40 shows the pion misidentification fraction versus the kaon identification efficiency as a function of track multiplicity and the number of reconstructed primary vertices, as the requirement on the likelihood difference $\Delta \log \mathcal{L}(K - \pi)$ is varied. The results demonstrate some degradation in PID performance with increased interaction multiplicity. However, the performance is still excellent and gives confidence that the RICH system will continue to perform well during LHC Run II.

4.3. Muon system based particle identification

The identification of a track reconstructed in the tracking system as a muon is based on the association of hits around its extrapolated trajectory in the muon system.⁸³

Fig. 39. Kaon identification efficiency and pion misidentification rate as measured using data (left) and from simulation (right) as a function of track momentum.⁸² Two different $\Delta \log \mathcal{L}(K-\pi)$ requirements have been imposed on the samples, resulting in the open and filled marker distributions, respectively.

Fig. 40. Pion misidentification fraction versus kaon identification efficiency as measured in 7 TeV LHCb collisions: (left) as a function of track multiplicity, and (right) as a function of the number of reconstructed primary vertices.⁸² The efficiencies are averaged over all particle momenta.

A search is performed for hits within rectangular windows around the extrapolation points where the x and y dimensions of the windows are parameterised as a function of momentum at each station and separately for each muon system region. The parameters are optimised to maximise the efficiency and at the same time provide low misidentification probabilities of pions as muons. The same criterion is used to define the number of stations required to have hits within a window as a function of momentum. A minimum momentum of 3 GeV/ c is necessary for a muon to traverse the calorimeters and reach the M2 and M3 stations, while above 6 GeV/ c they traverse all five of the stations. For each muon candidate, likelihoods for the muon and non-muon hypotheses are computed, based on the average squared distance of the hits that are closest to the extrapolation points.

The performance of the muon identification is obtained from data using muons from $J/\psi \rightarrow \mu^+ \mu^-$ decays, protons from $\Lambda \rightarrow p \pi^-$ decays and kaons and pions from $D^0 \rightarrow K^- \pi^+$, where the D^0 is selected from $D^{*+} \rightarrow D^0 \pi^+$ decays. These

Fig. 41. Top left: efficiency of the muon candidate selection based on the matching of hits in the muon system to track extrapolation, as a function of momentum for different p_T ranges. Other panels: misidentification probability of protons (top right), pions (bottom left), and kaons (bottom right) as muon candidates as a function of momentum, for different p_T ranges.

samples can be selected without using PID information and are characterised by relatively high statistics and low background. The latter is subtracted by fitting the appropriate invariant mass distribution. Figure 41 shows, as a function of the track momentum and for different ranges of transverse momentum, the efficiency of the muon candidate selection, and the probabilities of incorrect identification of protons, pions and kaons as muons.

The incorrect assignment of the muon identity to a proton occurs either due to a combination of spurious hits in the different muon stations that are aligned with the proton direction, or due to the existence of a true muon in the event that points in the same direction as the proton in the muon system. This muon can be produced close to the interaction point or in the calorimeter shower. Since the window dimension decreases with momentum and an increasing number of hits is required for tracks above 10 GeV/c, a strong reduction of the proton misidentification rate is seen in the interval 3–30 GeV/c and for $p_T < 1.7$ GeV/c. For higher p_T values, the protons have a high polar angle and therefore fall outside of the high-occupancy part of the detector. Decays in flight are the main cause of misidentification of

pions and kaons as muons. To a good approximation, the misidentification rate is the sum of the contribution of decays in flight and the proton misidentification probability.

The background rejection power can be improved by the computation of a likelihood for the muon and non-muon hypotheses, based on the pattern of hits around the extrapolation to the different muon stations of the charged particles trajectories reconstructed with high precision in the tracking system. The logarithm of the ratio between the muon and non-muon hypotheses, $\Delta \log \mathcal{L}(\mu)$, is used as a discriminating variable. The likelihood for the non-muon hypothesis is calibrated using proton data, since the other charged hadrons (pions or kaons) selected as muons will have a component identical to the protons and a component very similar to the true muons, due to decays in flight before the calorimeter. The muon likelihood has been calibrated with muons from $J/\psi \rightarrow \mu^+ \mu^-$ decays selected from data, while the non-muon likelihood has been calibrated with a simulated sample of $\Lambda \rightarrow p\pi^-$ decays.

4.4. Combined particle identification performance

The PID information obtained separately from the muon, RICH, and calorimeter systems is combined to provide a single set of more powerful variables. Two different approaches are used. In the first method the likelihood information produced by each sub-system is simply added linearly, to form a set of combined likelihoods, $\Delta \log \mathcal{L}_{\text{comb}}(X - \pi)$, where X represents either the electron, muon, kaon or proton mass hypothesis. These variables give a measure of how likely the mass hypothesis under consideration is, for any given track, relative to the pion hypothesis. A second approach has been subsequently developed to improve upon the simple log likelihood variables both by taking into account correlations between the detector systems and also by including additional information. This is carried out using multivariate techniques,⁸⁴ combining PID information from each sub-system into a single probability value for each particle hypothesis.

To illustrate the improvement made by combining information from different sub-detectors, the performance of the variable $\Delta \log \mathcal{L}_{\text{comb}}(e - \pi)$ is first considered, using a similar tag and probe technique to that of Subsec. 4.1.2 in which the calorimeter-only performance is presented. Figure 42 shows the pion misidentification versus electron identification probability, for various cuts on $\Delta \log \mathcal{L}_{\text{comb}}(e - \pi)$. Compared to Fig. 35 the improvement in the misidentification rate can clearly be seen, e.g. at $\sim 90\%$ electron efficiency the pion misidentification rate drops from $\sim 6\%$ to $\sim 0.6\%$.

An improvement in performance can also be seen for the muon identification, as illustrated by one of the most prominent LHCb results, the measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ branching fraction and the search for $B^0 \rightarrow \mu^+ \mu^-$ decays.⁸ The $B_{(s)}^0 \rightarrow h^+ h^-$ decay modes, where $h = (K, \pi)$, can fake a signal if both hadrons are misidentified as muons. Therefore the minimisation of these backgrounds is of

Fig. 42. Electron identification performance using the $\Delta \log \mathcal{L}_{\text{comb}}(e - \pi)$ variable, as measured in 8 TeV collision data, using a tag and probe technique with electrons from the decay $B^\pm \rightarrow (J/\psi \rightarrow e^+e^-)K^\pm$. Left, pion misidentification rate versus electron identification probability when the cut value is varied. Right, electron identification efficiency and pion misidentification rate as a function of track momentum, for two different cuts on $\Delta \log \mathcal{L}_{\text{comb}}(e - \pi)$.

Fig. 43. Background misidentification rates versus muon (left) and proton (right) identification efficiency, as measured in the $\Sigma^+ \rightarrow p\mu^+\mu^-$ decay study. The variables $\Delta \log \mathcal{L}(X - \pi)$ (black) and ProbNN (red), the probability value for each particle hypothesis, are compared for 5–10 GeV/c muons and 5–50 GeV/c protons, using data sidebands for backgrounds and Monte Carlo simulation for the signal.

paramount importance for this analysis. This double misidentification probability has been evaluated by folding the $K \rightarrow \mu$ and $\pi \rightarrow \mu$ fake rates extracted from a $D^0 \rightarrow K\pi$ sample from $D^{*+} \rightarrow D^0\pi^+$ decays, in bins of p and p_T , into the spectrum for selected simulated $B_{(s)}^0 \rightarrow h^+h^-$ events. If the tracks identified as muons are also required to satisfy a selection using the combined PID information ($\Delta \log \mathcal{L}_{\text{comb}}(K - \pi) < 10$ and $\Delta \log \mathcal{L}_{\text{comb}}(\mu - \pi) > -5$), the $B_{(s)}^0 \rightarrow h^+h^-$ misidentification probability is reduced by a factor of ~ 6 , whilst only $\sim 3\%$ of the $B_s \rightarrow \mu^+\mu^-$ signal is lost.

The possible improvement of the multivariate approach with respect to the simple log likelihood may also be illustrated by the ongoing search for the flavour-changing neutral current decay $\Sigma^+ \rightarrow p\mu^+\mu^-$. In Fig. 43 the misidentification rates versus efficiency curves for the $\Delta \log \mathcal{L}(X - \pi)$ and the probability value for each

particle hypothesis variables are shown. The improvement is clearly visible for both muons and protons. These multivariate variables will be further developed and utilised more extensively during Run II.

5. Trigger

The LHCb trigger consists of two levels; Level-0 (L0) and the High Level Trigger (HLT). In Run I the trigger reduced the rate of events to be saved for physics analysis to 2–5 kHz.

The L0 trigger is implemented in hardware and makes decisions based on information from the calorimeter and muon systems in order to reduce the event rate to below 1 MHz, at which point the whole detector can be read out. The HLT is a software application running on the event filter farm (EFF). A fraction of L0 accepted events are deferred to disk for processing by the HLT during the inter-fill time, optimising the use of available EFF resources.

After the HLT, events are stored and later processed with a more accurate alignment and calibration of the sub-detectors, and with the reconstruction software described in Subsec. 2.2. This part of the reconstruction and subsequent selection of interesting events is referred to as the offline reconstruction and selection for the remainder of this discussion.

Subsection 5.1 describes a data-driven method to determine the efficiency and purity of the signals that are selected by the LHCb trigger. The implementation of the L0 trigger²⁵ is briefly summarised in Subsec. 5.2 and the performance of the HLT⁸⁵ is discussed in Subsec. 5.3. The selection criteria used in the two trigger levels during 2011 and 2012 are described in detail in Ref. 86. The deferral system is discussed in Subsec. 5.4 and a short summary of the LHCb trigger is provided in Subsec. 5.5.

The results presented here are based on the configuration and performance of the LHCb trigger during 2012 when the deferral system was first used and the majority of the Run I dataset was collected.

5.1. Data driven trigger performance determination

The trigger efficiencies are evaluated using events reconstructed with the full offline software, and are calculated with respect to candidates selected by the full offline analysis of the respective channel, excluding the trigger. They therefore quantify the inefficiencies due to the simplified reconstruction, possible misalignments and reduced resolution, as compared to the offline reconstruction. They also account for any tighter selection requirements that are needed to satisfy the rate and processing time limitations. The following decay channels are chosen to highlight the performance of the trigger:^a

^aHere, and in the following, charge conjugated decays are implicitly included.

- $B^+ \rightarrow J/\psi K^+$ decays, with $J/\psi \rightarrow \mu^+ \mu^-$. This decay evaluates the muon trigger efficiency and serves as a proxy for several key physics decay channels like $B_s^0 \rightarrow J/\psi \phi$, $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ or $B_s^0 \rightarrow \mu^+ \mu^-$.
- $B^0 \rightarrow K^+ \pi^-$ as a typical two-body hadronic beauty decay.
- $B^0 \rightarrow D^+ \pi^-$ with $D^+ \rightarrow K^- \pi^+ \pi^+$, as a typical four-body beauty decay.
- $D^0 \rightarrow K^- \pi^+$ as a two-body charm decay.
- $D^+ \rightarrow K^- \pi^+ \pi^+$ represents a three-body charm decay.
- $D^{*+} \rightarrow D^0 \pi^+$, followed by the four-body charm decay $D^0 \rightarrow K^- \pi^+ \pi^- \pi^+$.

These decay channels and their selections are representative of the trigger needs of the core physics analyses of the LHCb experiment. The selected charm modes cover the topologies that are most sensitive to CP violating effects. All samples used in this study carry a large signal to background ratio. Nevertheless, the yields are determined by fitting the signal peaks in the invariant mass distributions in order to subtract the residual background.

In the following, the term ‘signal’ refers to a combination of tracks forming the offline reconstructed and selected b - or c -hadron candidates. To determine the trigger efficiency, trigger objects are associated with signal tracks. The criteria used to associate a trigger object with a signal track are as follows. An event is classified as Trigger on Signal (TOS) if the trigger objects that are associated with the signal candidate are sufficient to trigger the event. An event is classified as Trigger Independent of Signal (TIS) if it has been triggered by trigger objects that are not associated with the signal. Some events can be classified as TIS and TOS simultaneously (TIS & TOS), which allows the determination of the trigger efficiency relative to the offline reconstructed events from data alone. The efficiency to trigger an event independently of the signal, ϵ^{TIS} , is given by

$$\epsilon^{\text{TIS}} = N^{\text{TIS \& TOS}} / N^{\text{TOS}},$$

where N^{TOS} is the number of events classified as TOS. The efficiency to trigger an event on the signal alone, ϵ^{TOS} , is given by

$$\epsilon^{\text{TOS}} = N^{\text{TIS \& TOS}} / N^{\text{TIS}},$$

where N^{TIS} is the number of events classified as TIS. This method of measuring trigger efficiencies is discussed in detail in Ref. 87.

5.2. Level-0 hardware trigger

The L0 trigger is divided into three independent units; the L0-Calorimeter trigger, the L0-Muon trigger and the L0-PileUp trigger, the latter being used only for the determination of the luminosity.²⁷ The L0 trigger system is fully synchronous with the 40 MHz bunch crossing rate of the LHC. The latencies are fixed and are independent of the occupancy or the bunch crossing history.

The L0-Calorimeter system uses information from the SPD, PS, ECAL and HCAL detectors, as described in Sec. 3. It computes the transverse energy, E_T ,

Table 4. Typical L0 thresholds used in Run I.⁸⁶

	p_T or E_T		SPD 2011 and 2012
	2011	2012	
Single muon	1.48 GeV/ c	1.76 GeV/ c	600
Dimuon $p_{T1} \times p_{T2}$	(1.30 GeV/ c) ²	(1.60 GeV/ c) ²	900
Hadron	3.50 GeV	3.70 GeV	600
Electron	2.50 GeV	3.00 GeV	600
Photon	2.50 GeV	3.00 GeV	600

deposited by incident particles in clusters of 2×2 cells. From these clusters, the following three types of candidates are built. **L0Hadron** is the highest E_T HCAL cluster, which also contains the energy of the matching ECAL cluster. **L0Photon** is the highest E_T ECAL cluster with 1 or 2 PS hits in front of the ECAL cluster and no hit in the SPD cells corresponding to the PS cells. **L0Electron** has the same requirements as **L0Photon**, with the additional condition of at least one SPD cell hit in front of the PS cells. The E_T of each candidate is compared to a fixed threshold and events containing at least one candidate above threshold fire the L0 trigger. The total number of hits in the SPD is also determined, and is used to veto events that would take a disproportionately large fraction of the available processing time in the HLT. The SPD hit multiplicity requirements are listed in Table 4.

The L0 muon processors look for the two highest p_T muon tracks in each quadrant. The position of a track in the first two stations allows the determination of its p_T with a measured momentum resolution of roughly 25%. The trigger considers the eight candidates and sets a single threshold on either the largest transverse momentum, p_{T}^{largest} , or on the product, $p_{T}^{\text{largest}} \times p_{T}^{\text{2nd largest}}$. These thresholds are listed in Table 4. The tightening of L0 thresholds in the 2012 configuration is a consequence of the increased luminosity and beam energy.

The trigger efficiencies are measured on offline selected events, using the techniques described in Subsec. 5.1. The efficiencies of the L0 muon triggers evaluated on $B^+ \rightarrow J/\psi K^+$ events are shown in Fig. 44. The majority of events are accepted by the single muon trigger. The largest inefficiency originates from the tight muon identification requirements inside the L0 reconstruction algorithm. The L0 dimuon trigger selects a small fraction of additional candidates at lower transverse momenta. The combined efficiency for both L0 muon triggers is evaluated to be $89 \pm 0.5\%$.⁸⁷

The L0 hadron efficiency is shown in Fig. 44 for the two- and three-prong beauty decays $B^0 \rightarrow K^+\pi^-$ and $B^0 \rightarrow D^+\pi^-$ and the two-, three- and four-prong charm decays $D^0 \rightarrow K^-\pi^+$, $D^+ \rightarrow K^-\pi^+\pi^+$ and $D^{*+} \rightarrow D^0\pi^+$, as a function of the B or D meson p_T . The two-prong beauty decay is triggered with highest efficiency by the L0 hadron E_T criterion ($\epsilon^{\text{TOS}} = 40\%$) while the four-prong charm decay $D^{*+} \rightarrow D^0\pi^+$ is selected with the lowest efficiency ($\epsilon^{\text{TOS}} = 22\%$). The other modes lie in between, as shown in Fig. 44. With the inclusion of TIS triggers, the total efficiencies increase significantly, e.g. from 40% to 53% for $B^0 \rightarrow K^+\pi^-$.

Fig. 44. (left) L0 muon trigger performance: TOS trigger efficiency for selected $B^+ \rightarrow J/\psi K^+$ candidates. (right) L0 hadron trigger performance: TOS trigger efficiency for different beauty and charm decay modes.

The total output rate of the L0 trigger is limited to 1 MHz, which is the maximum rate at which the LHCb detector can be read out. This output rate consists of approximately 400 kHz of muon triggers, 500 kHz of hadron triggers and 150 kHz of electron and photon triggers, where the individual triggers have an overlap of about 10%.

5.3. High level trigger

Events accepted by L0 are transported by the data acquisition network to one of the processors of the EFF. The HLT is a software application, of which 29,500 instances run on the EFF, and each instance consists of independently operating trigger “lines,” each of which consists of selection parameters for a specific class of events.

The HLT is based on the same software framework used throughout LHCb. Given the available resources in the EFF, the time per event is around fifty times smaller than in the offline processing. The HLT is divided in two levels. In the first level (HLT1), a partial event reconstruction is performed. In the second level (HLT2), the complete event is reconstructed. Where time allows, the HLT uses the same reconstruction algorithms as employed offline, with some simplifications that are needed to satisfy the time constraints.

5.3.1. First level

The offline VELO reconstruction algorithm which performs a full 3D pattern recognition is sufficiently fast to be run on all events entering the HLT. However, the offline algorithm makes a second pass on unused hits to enhance the efficiency for tracks pointing far away from the beam-line, while in the HLT this second pass is not used due to CPU constraints. Vertices are reconstructed from a minimum of five intersecting VELO tracks. Vertices within a radius of 300 μm of the mean position of the pp -interaction envelope are considered to be primary vertices.

Fig. 45. HLT1 inclusive track trigger performance: TOS efficiency for various channels as a function of B or D p_T (left). HLT1 muon trigger performance: TOS efficiency for $B^+ \rightarrow J/\psi K^+$ candidates as function of B^+ p_T (right).

HLT1 limits the number of VELO tracks that are passed through the forward tracking algorithm that searches for matching hits in the tracking stations. VELO tracks must have a significant IP with respect to all PVs, or be matched to muon chamber hits by a fast muon identification algorithm. This algorithm is only run in events that were triggered by a muon line in the hardware step. To further limit the processing time, the forward track search has a minimum momentum requirement that varied between 3 and 6 GeV/ c during Run I. VELO tracks without matching muon hits are also subject to a minimum p_T threshold that varied between 0.5 and 1.25 GeV/ c . The reconstructed forward tracks are fitted using a Kalman filter with a simplified detector geometry description and fewer iterations than in the offline configuration. The invariant mass resolution of $J/\psi \rightarrow \mu^+\mu^-$ candidates reconstructed with this procedure is found to be 3% worse than the 15.1 MeV/ c^2 obtained offline. For tracks identified as muon candidates, the basic offline muon identification algorithm⁸³ is applied to increase the purity of the muon sample.

The inclusive beauty and charm trigger is based on the properties of one good quality track candidate. The selection is based on track p_T (typical value $p_T > 1.6$ – 1.7 GeV/ c) and displacement from the primary vertex (typical value IP > 0.1 mm). This trigger line produces around 58 kHz of output, which is the largest contribution to the allocated HLT1 bandwidth. It is the most efficient line for physics channels that do not contain leptons in the final state. The performance of HLT1 for hadronic signatures is shown in Fig. 45 as a function of the p_T of the resonance considered. The inclusive one-track based trigger also exists in a version for electrons or photons identified by L0, with reduced thresholds relative to the inclusive version. The output rate of these lines is around 7 kHz.

A similar line exists for tracks that are matched to hits in the muon chambers.⁸⁸ This single muon trigger line selects good quality muon candidates that are displaced from the primary vertex and satisfy $p_T > 1$ GeV/ c . Single muon candidates that satisfy $p_T > 4.8$ GeV/ c are selected by a special trigger line without any vertex separation requirements.

Dimuon candidates are either selected based on their mass ($m_{\mu\mu} > 2.5 \text{ GeV}/c^2$) without any displacement requirement, or based on their displacement without the mass restriction. The dominant inefficiency for these lines originates from the online muon identification algorithms. The performance of HLT1 on muonic signatures as a function of p_T of the B^+ parent is shown in Fig. 45. The single muon line has an efficiency of around 90% to select $B^+ \rightarrow J/\psi K^+$ decays, while the dimuon lines have an efficiency of around 70%. The combination of muon trigger lines produces an output rate of around 14 kHz.

In addition to the trigger lines discussed above, several dedicated lines are implemented to enhance the trigger performance for events containing candidates for high p_T electrons, di-protons, displaced vertices or high E_T jets. A set of technical lines including selections for luminosity and beam-gas measurements completes the list of HLT1 triggers.

5.3.2. Second level

HLT1 reduces the event rate to about 80 kHz, which is sufficiently low to allow the forward tracking of all VELO tracks in HLT2. As described in Subsec. 2.2, the offline reconstruction uses two complementary tracking algorithms. Due to the CPU constraints, HLT2 only searches for long tracks based on VELO seeds. This simplification leads to a lower tracking efficiency compared to the offline reconstruction of 1–2% per track.

The processing time is further reduced by restricting the search to tracks with $p > 3 \text{ GeV}/c$ and $p_T > 0.3 \text{ GeV}/c$. Muon identification in HLT2 is performed using the offline muon identification algorithm. Tracks are also associated to ECAL clusters to identify electrons. Photons and neutral pions are built starting from the energy clusters reconstructed by the L0-Calorimeter system.

Generic beauty trigger

A significant portion of the output rate of HLT2 is selected by the ‘topological’ lines, which are designed to trigger on partially reconstructed b -hadron decays. These topological lines cover all b -hadron decays with at least two charged particles in the final state and a displaced decay vertex. The inclusive nature of these lines makes them less susceptible to the 1–2% loss in efficiency per reconstructed track in HLT2. Tracks are selected based on their track fit χ^2/ndf , IP and muon or electron identification. Two-, three- or four-body vertices are constructed from the selected tracks with a requirement on their distance of closest approach (DOCA).

Candidate n -body combinations are selected based on the following variables: $\sum |p_T|$, $p_T^{\min}$, n -body invariant mass (m), DOCA, $\text{IP}\chi^2$ and flight distance (FD) χ^2 . In addition, the corrected mass is defined as $m_{\text{corr}} = \sqrt{m^2 + |p'_{T\text{miss}}|^2} + |p'_{T\text{miss}}|$, where $p'_{T\text{miss}}$ is the missing momentum transverse to the line of flight between the n -body vertex and the PV to which it has the smallest IP.⁸⁹ Figure 46 shows the

Fig. 46. Simulated $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ events. The reconstructed 2-body mass is shown in red and the corrected mass (m_{corr} , see text for definition) is shown in black.

Fig. 47. HLT2 inclusive beauty trigger performance as a function of B p_T . The efficiency for the exclusive $B^0 \rightarrow K^+ \pi^-$ trigger line is also given.

reconstructed 2-body and corrected mass distributions for $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ events. These variables are combined using a boosted decision tree trained on simulated signal events and data taken in 2010.⁹⁰ An explicit veto on prompt charm is also applied to reduce its output rate.

Figure 47 shows the efficiency for the topological trigger lines for $B^0 \rightarrow K^+ \pi^-$ and $B^0 \rightarrow D^+ \pi^-$ events as well as the additional efficiency that can be gained by an exclusive selection for $B^0 \rightarrow K^+ \pi^-$ in the low p_T regime. The output rate of the topological trigger is 2 kHz, in which it efficiently selects beauty decays to charged tracks. For example, the efficiencies for $B^0 \rightarrow K^+ \pi^-$ and $B^0 \rightarrow D^+ \pi^-$ decays are approximately 78% and 76%, respectively.

If one of the tracks forming the generic beauty trigger decision is identified as a muon, the selection on the boosted decision tree classifier is loosened which enhances the efficiency for muonic beauty decays like $B^0 \rightarrow K^{*0} \mu^+ \mu^-$.

Fig. 48. HLT2 muon trigger performance for the J/ψ trigger lines (left). The triggers Hlt2DiMuonJPsi and Hlt2DiMuonJPsiHighPT are the two prompt J/ψ triggers and Hlt2DiMuonDetachedJPsi is the trigger that selects J/ψ candidates that are inconsistent with coming from the primary vertex. HLT2 charm trigger performance for inclusive and exclusive selections (right). The decay $D^{*+} \rightarrow D^0 \pi^+$ is followed by $D^0 \rightarrow K^- \pi^+ \pi^+ \pi^-$.

Muon triggers

Several trigger lines select events that contain one or two muons. The muon identification procedure in HLT2 is identical to that which is used offline.⁸³ Single muon candidate events are selected if the muon passes a tight p_T requirement ($p_T > 10$ GeV/ c) or if candidates are inconsistent with originating from the PV and they satisfy moderate p_T ($p_T > 1.3$ GeV/ c) and tight track quality requirements. Candidates selected with the latter criteria are prescaled by factor of two.

Dimuon candidate events are selected without a mass requirement if the dimuon vertex is separated from the primary vertex. If the mass of the muon pair is within ± 100 MeV/ c^2 ($\sim 8\sigma$) of the J/ψ mass, three trigger selections are considered. Decay-time unbiased $J/\psi \rightarrow \mu^+ \mu^-$ candidates are extensively used for calibration of the LHCb decay time acceptance. Therefore, all muon pairs with $p_T > 2$ GeV/ c are considered, as well as a fraction of those without any p_T requirement. These two prompt selections are complemented by an additional detached J/ψ trigger. The separation requirement of this trigger is looser than that of the generic detached dimuon trigger described above.

This set of lines is optimised to fully exploit the large physics potential of both prompt J/ψ and $B \rightarrow J/\psi X$ decays. Relative to the offline selection, the trigger efficiencies are typically above 90%. Figure 48 shows the performance of the J/ψ triggers, where the effective prescale of about a factor of two on the prompt J/ψ line is visible, as well as the p_T acceptance of the high p_T line. The total output rate of all single and dimuon trigger lines is around 1 kHz.

Charm triggers

In the 2012 running conditions, roughly 600 kHz of $c\bar{c}$ -events are produced within the acceptance of LHCb. This large rate of charm production means that tight

exclusive selections are required in the trigger. The exception is the decay chain $D^{*+} \rightarrow D^0 \pi^+$, which can be selected inclusively, i.e. only reconstructing two charged tracks from the D^0 decay matched to a slow pion from the D^{*+} decay. The mass difference between the D^{*+} and D^0 candidates remains a good discriminating variable, enabling the rate to be sufficiently reduced. The D^0 is partially reconstructed in all combinations containing $\pi^\pm$, $K^\pm$, p , $\mu^\pm$, K_S^0 or Λ^0 enabling both rare decay and CP violation measurements.

The dominant exclusive selections for prompt charm are the hadronic two body and three body lines. The efficiency of these triggers is summarised in Fig. 48. Additional selections for hadronic, leptonic and semi-leptonic D and Λ_c^+ decays are implemented. The total output rate of all charm selections is ~ 2 kHz, but the trigger efficiencies strongly depend on the offline selection: using a pure sample of $D^0 \rightarrow K^- \pi^+$ events from D^{*+} decays, the HLT2 efficiency is 90%, while the trigger efficiency for the multibody decay $D^{*+} \rightarrow D^0 \pi^+$ is 26%.

Exclusive and technical lines

The HLT incorporates a large number of exclusive and technical trigger lines to complement the signal selection by the inclusive lines discussed above. For example, 100 Hz of random events are recorded throughout the full data taking period. These events can be used to understand the trigger system, for fast monitoring of the data and for luminosity determinations. A number of trigger lines also exist that maximise the performance for decays with electrons or photons in the final state.⁹¹ A single muon line with hard p_T requirements is designed to select heavy particles decaying promptly to one or more muons, like $W^\pm$ or Z^0 for electroweak measurements.

The remaining lines are required for luminosity measurements, prescaled physics lines with looser cuts, lines which trigger on low multiplicity events and lines which look for large transverse momentum jets. The trigger also contain lines that pass events for monitoring to allow fast feedback on the quality of the data.

5.4. Deferred trigger

The LHC delivers stable beams $\sim 30\%$ of the available time, and thus unless otherwise occupied the EFF would be idle the remaining $\sim 70\%$ of the time. Therefore to maximise the use of available resources, LHCb developed a deferred triggering system for data taking in 2012.^{92,93} Around 20% of the L0 accepted events are temporarily saved on the EFF node local disks and these events are subsequently processed by the HLT during the inter-fill periods. The effective increase in CPU resources this provides was used to reduce the p_T requirement in the forward tracking algorithm and to include additional tracking algorithms that allow the reconstruction of tracks from particles that decay beyond the VELO acceptance. These modifications significantly enhance the efficiency for charm decays.

Table 5. Efficiencies of selected channels for the whole trigger chain using the 2012 trigger configuration. The efficiencies are normalised to the number of events that are offline selected.

Channel	L0	HLT1	HLT2
$B^+ \rightarrow J/\psi K^+, J/\psi \rightarrow \mu^+ \mu^-$	89%	92%	87%
$B^0 \rightarrow K^+ \pi^-$	53%	97%	80%
$B^0 \rightarrow D^+ \pi^-, D^+ \rightarrow K^- \pi^+ \pi^+$	59%	98%	77%
$D^+ \rightarrow K^- \pi^+ \pi^+$	44%	89%	91%
$D^{*+} \rightarrow D^0 \pi^+, D^0 \rightarrow K^- \pi^+ \pi^- \pi^+$	49%	93%	30%

An additional benefit of this system is that it provides redundancy against problems downstream in the dataflow. Prior to the adoption of this scheme, such problems would quickly lead to back-pressure all the way up into the farm and lead to dead-time. In the deferred trigger scheme, although the fraction of deferred events would increase, no dead-time would be incurred. The deferred trigger requires a substantial amount of local disk space, which was added to the farm before the 2012 run. The requirements on the performance of the disks are modest,^b such that inexpensive desktop hard drives could be used.

5.5. Trigger performance summary

The LHCb trigger system has delivered a range of physics modes with high efficiency in the first running phase of the LHC. It is primarily designed to select charm and beauty hadrons over a large range of momentum and decay time, and its efficiency can be determined directly from the data.

The flexible design of the HLT, which is fully deployed in software, allows for the rapid adaptation to changes in running conditions and physics goals. Several innovative concepts have been developed that enable inclusive selections to be utilised in the full trigger chain and thus provide an efficient trigger for nearly any beauty decay to charged particles. The deferred triggering permits optimisation of computing resources for mean instead of peak usage, leading to an effective increase in farm size in 2012 of 20–30%. Multivariate selections allow the inclusive triggering of beauty decays to charged tracks with high efficiency.

Typical trigger efficiencies for selected signals are summarised in Table 5. The trigger efficiency is high for muonic b -decays. For hadronic b -decays, the hardware trigger stage causes a significant loss of efficiency, but the software component remains efficient. The trigger efficiency for multibody charm decays is lower than for b -hadron decays, due to the softer momentum spectra of the final state particles.

The LHCb trigger system also efficiently covers physics beyond the core beauty and charm programme. This includes $W^\pm$ or Z^0 production, inclusive particle

^bThe input data rate per farm-server is about 60 MB/s, which is significantly lower than the typical disk performance of about 120 MB/s.

production, and exotic phenomena, for instance displaced vertices from heavy long-lived particles.

6. Conclusion and Outlook

During the first running period of the LHC between 2009 and 2013, the LHCb experiment recorded a total of about 3.2fb^{-1} of integrated luminosity with pp collisions at centre-of-mass energies of 0.9, 7 and 8 TeV, and 1.6nb^{-1} of integrated luminosity with proton–lead collisions. The majority of the data were recorded under conditions corresponding to a luminosity of $4 \times 10^{32}\text{cm}^{-2}\text{s}^{-1}$, a bunch spacing of 50 ns and a pile-up of 1.8. Despite the fact that these are significantly more challenging than the conditions originally foreseen for the experiment, it has been demonstrated that the performance of each sub-system and the global performance of the detector are in good agreement with the original expectations presented in the LHCb detector paper.²⁵ For some physics analyses the original expectations are even exceeded, thanks to new ideas and well understood backgrounds.

During Run II of the LHC, the LHCb experiment expects to collect an additional 5fb^{-1} of data, which improves the prospects for observing physics beyond the Standard Model using heavy quark flavours as a tool. However, the read-out and trigger scheme of the current LHCb detector limit the luminosity that can be recorded. To overcome these limitations, the LHCb experiment will be upgraded^{94–99} to allow the detector to be read out at the maximum LHC bunch crossing rate of 40 MHz with a flexible software-based trigger. This will not only allow the data rate to be increased substantially, but will also increase the trigger efficiency, especially for channels currently triggered at the hardware level by energy deposits in the calorimeters. In addition to the significant increase in sensitivity for flavour physics, the experiment will continue to explore other interesting signatures and thus act as a general purpose detector in the forward region.

Acknowledgments

We express our gratitude to our colleagues in the CERN accelerator departments for the excellent performance of the LHC. We thank the technical and administrative staff at the LHCb institutes. We acknowledge support from CERN and from the national agencies: CAPES, CNPq, FAPERJ and FINEP (Brazil); NSFC (China); CNRS/IN2P3 (France); BMBF, DFG, HGF and MPG (Germany); SFI (Ireland); INFN (Italy); FOM and NWO (The Netherlands); MNiSW and NCN (Poland); MEN/IFA (Romania); MinES and FANO (Russia); MinECo (Spain); SNSF and SER (Switzerland); NASU (Ukraine); STFC (United Kingdom); NSF (USA). The Tier1 computing centres are supported by IN2P3 (France), KIT and BMBF (Germany), INFN (Italy), NWO and SURF (The Netherlands), PIC (Spain), GridPP (United Kingdom). We are indebted to the communities behind the multiple open source software packages on which we depend. We are also thankful for the computing resources and the access to software R&D tools provided by

Yandex LLC (Russia). Individual groups or members have received support from EPLANET, Marie Skłodowska-Curie Actions and ERC (European Union), Conseil général de Haute-Savoie, Labex ENIGMASS and OCEVU, Région Auvergne (France), RFBR (Russia), XuntaGal and GENCAT (Spain), Royal Society and Royal Commission for the Exhibition of 1851 (United Kingdom).

References

1. N. Cabibbo, *Phys. Rev. Lett.* **10**, 531 (1963).
2. M. Kobayashi and T. Maskawa, *Prog. Theor. Phys.* **49**, 652 (1973).
3. A. Sakharov, *Pisma Zh. Eksp. Teor. Fiz.* **5**, 32 (1967).
4. LHCb Collab. (R. Aaij *et al.*), *Phys. Lett. B* **694**, 209 (2010), arXiv:1009.2731.
5. LHCb Collab. (R. Aaij *et al.*), *Nucl. Phys. B* **871**, 1 (2013), arXiv:1302.2864.
6. LHCb Collab. (R. Aaij *et al.*), *Eur. Phys. J. C* **73**, 2373 (2013), arXiv:1208.3355.
7. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **110**, 021801 (2013), arXiv:1211.2674.
8. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **111**, 101805 (2013), arXiv:1307.5024.
9. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **08**, 131 (2013), arXiv:1304.6325.
10. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **111**, 191801 (2013), arXiv:1308.1707.
11. LHCb Collab. (R. Aaij *et al.*), *Phys. Lett. B* **736**, 186 (2014), arXiv:1405.4140.
12. LHCb Collab. (R. Aaij *et al.*), Precision measurement of CP violation in $B_s^0 \rightarrow J/\psi K^+ K^-$ decays, submitted to *Phys. Rev. Lett.*, arXiv:1411.3104.
13. LHCb Collab. (R. Aaij *et al.*), *Phys. Lett. B* **726**, 151 (2013), arXiv:1305.2050.
14. LHCb Collab., Improved constraints on γ : CKM2014 update, LHCb-CONF-2014-004.
15. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **112**, 041801 (2014), arXiv:1310.7201.
16. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **06**, 058 (2012), arXiv:1204.1620.
17. LHCb Collab. (R. Aaij *et al.*), Measurement of the forward W boson production cross-section in pp collisions at $\sqrt{s} = 7$ TeV, to appear in *J. High Energy Phys.*, arXiv:1408.4354.
18. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **110**, 222001 (2013), arXiv:1302.6269.
19. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **112**, 222002 (2014), arXiv:1404.1903.
20. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. D* **85**, 112004 (2012), arXiv:1201.5600.
21. LHCb Collab. (R. Aaij *et al.*), *Phys. Lett. B* **724**, 36 (2013), arXiv:1304.4518.
22. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **02**, 072 (2014), arXiv:1308.6729.
23. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **07**, 094 (2014), arXiv:1405.5152.
24. LHCb Collab., LHCb reoptimized detector design and performance: Technical Design Report, CERN-LHCC-2003-030, LHCb-TDR-009.
25. LHCb Collab. (A. Alves *et al.*), *JINST* **3**, S08005 (2008).
26. L. Evans and P. Bryant, *JINST* **3**, S08001 (2008).
27. LHCb Collab. (R. Aaij *et al.*), *JINST* **7**, P01010 (2012), arXiv:1110.2866.
28. LHCb Collab. (R. Aaij *et al.*), *JINST* **9**, P12005 (2014), arXiv:1410.0149.
29. R. Aaij *et al.*, *JINST* **9**, P09007 (2014), arXiv:1405.7808.
30. C. Parkes, T. Ruf and T. Szumlak, Reconstruction of cluster positions in the LHCb VELO, Tech. Rep. LHCb-2007-151, CERN-LHCb-2007-151, CERN, Geneva, December 2007.
31. A. Affolder *et al.*, *JINST* **8**, P08002 (2013), arXiv:1302.5259.
32. R. Arink *et al.*, *JINST* **9**, P01002 (2014), arXiv:1311.3893.
33. D. van Eijk *et al.*, *Nucl. Instrum. Methods A* **685**, 62 (2012).
34. E. Aslanides *et al.*, Performance of the muon trigger with a realistic simulation, CERN-LHCb-2002-041.
35. A. A. Alves Jr. *et al.*, *JINST* **8**, P02022 (2013), arXiv:1211.1346.

36. M. Anelli *et al.*, *JINST* **5**, P10003 (2010), arXiv:1009.1963.
37. O. Callot, FastVelo, a fast and efficient pattern recognition package for the VELO, LHCb-PUB-2011-001.
38. D. Hutchcroft, VELO pattern recognition, LHCb-2007-013.
39. O. Callot and S. Hansmann-Menzemer, The forward tracking: Algorithm and performance studies, LHCb-2007-015.
40. M. Needham and J. van Tilburg, Performance of the track matching, LHCb-2007-020.
41. M. Needham, Performance of the track matching, LHCb-2007-129.
42. O. Callot and M. Schiller, PatSeeding: A standalone track reconstruction algorithm, LHCb-2008-042.
43. O. Callot, Downstream pattern recognition, LHCb-2007-026.
44. S. Stahl, Reconstruction of displaced tracks and measurement of K_S^0 production in proton-proton collisions at $\sqrt{s} = 900$ GeV at the LHCb experiment, Master's thesis, Heidelberg University, 2010.
45. O. Callot, M. Kucharczyk and M. Witek, VELO-TT track reconstruction, LHCb-2007-010.
46. R. Fruhwirth, *Nucl. Instrum. Methods A* **262**, 444 (1987).
47. J. A. N. van Tilburg, Track simulation and reconstruction in LHCb, Ph.D. thesis, Vrije Universiteit, Amsterdam, The Netherlands, 2005, CERN-THESIS-2005-040.
48. LHCb Collab. (R. Aaij *et al.*), *Eur. Phys. J. C* **74**, 2888 (2014), arXiv:1402.4430.
49. A. Jaeger *et al.*, Measurement of the track finding efficiency, LHCb-PUB-2011-025.
50. LHCb Collab. (R. Aaij *et al.*), Measurement of the track reconstruction efficiency at LHCb, submitted to *JINST*, arXiv:1408.1251.
51. LHCb Collab. (R. Aaij *et al.*), *Phys. Lett. B* **708**, 241 (2012), arXiv:1112.4896.
52. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **109**, 232001 (2012), arXiv:1209.5634.
53. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **110**, 182001 (2013), arXiv:1302.1072.
54. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **06**, 065 (2013), arXiv:1304.6865.
55. Particle Data Group (K. A. Olive *et al.*), *Chin. Phys. C* **38**, 090001 (2014).
56. S. Viret, C. Parkes and M. Gersabeck, *Nucl. Instrum. Methods A* **596**, 157 (2008), arXiv:0807.5067.
57. M. Gersabeck, S. Viret and C. Parkes, *Nucl. Instrum. Methods A* **596**, 164 (2008), arXiv:0807.5069.
58. V. Blobel and C. Kleinwort, A new method for the high precision alignment of track detectors, arXiv:hep-ex/0208021.
59. M. Deissenroth, Software alignment of the LHCb outer tracker chambers, Ph.D. thesis, Heidelberg University, 2010, CERN-THESIS-2010-063.
60. W. Hulsbergen, *Nucl. Instrum. Methods A* **600**, 471 (2009), arXiv:0810.2241.
61. J. M. Amoraal, Alignment with Kalman filter fitted tracks and reconstruction of $B_s^0 \rightarrow J/\psi\phi$ decays, Ph.D. thesis, Vrije Universiteit, Amsterdam, The Netherlands, 2011, CERN-THESIS-2011-011.
62. J. Amoraal *et al.*, *Nucl. Instrum. Methods A* **712**, 48 (2013), arXiv:1207.4756.
63. S. Borghi, *Nucl. Instrum. Methods Phys. Res. Sec. A* **623**, 156 (2010) [*1st International Conference on Technology and Instrumentation in Particle Physics*].
64. E. E. Simioni, New physics from rare beauty, Ph.D. thesis, Vrije Universiteit, Amsterdam, The Netherlands, 2010, CERN-THESIS-2010-031.
65. R. S. Märki, Measurements of b -baryon properties at LHCb, Ph.D. thesis, École Polytechnique Fédérale de Lausanne, Switzerland, January 2014, CERN-THESIS-2014-029.
66. LHCb Collab. (R. Aaij *et al.*), *New J. Phys.* **15**, 053021 (2013), arXiv:1304.4741.
67. LHCb Collab. (R. Aaij *et al.*), *Phys. Rev. Lett.* **108**, 101803 (2012), arXiv:1112.3183.

68. H. Moser and A. Roussarie, *Nucl. Instrum. Methods A* **384**, 491 (1997).
69. M. Pivk and F. R. Le Diberder, *Nucl. Instrum. Methods A* **555**, 356 (2005), arXiv:physics/0402083.
70. LHCb Collab., LHCb calorimeters: Technical Design Report, CERN-LHCC-2000-036, LHCb-TDR-002.
71. C. Abellan Beteta *et al.*, *JINST* **7**, P08020 (2012).
72. LHCb Calorimeter Group (I. Machikhiliyan), *J. Phys. Conf. Ser.* **293**, 012052 (2011).
73. R. V. Gomez, *J. Phys. Conf. Ser.* **293**, 012059 (2011).
74. V. Niess, *J. Phys. Conf. Ser.* **293**, 012060 (2011).
75. A. Martens, Towards a measurement of the angle γ of the Unitarity Triangle with the LHCb detector at the LHC (CERN): Calibration of the calorimeters using an energy flow technique and first observation of the $B_s^0 \rightarrow \bar{D}^0 \bar{K}^{*0}$ decay, Ph.D. thesis, Orsay, Orsay, 2011, CERN-THESIS-2011-128, LAL-11-172.
76. I. Belyaev *et al.*, *J. Phys. Conf. Ser.* **331**, 032050 (2011).
77. O. Deschamps *et al.*, Photon and neutral pion reconstruction, LHCb-2003-091.
78. LHCb Collab. (R. Aaij *et al.*), *Nucl. Phys. B* **867**, 1 (2013), arXiv:1209.0313.
79. D. A. R. Romero, Study of charmless $B_{d,s}^0 \rightarrow h^+ h'^- \pi^0$ decays in LHCb, Ph.D. thesis, Université Blaise Pascal, Clermont-Ferrand, France, 2013, CERN-THESIS-2013-051.
80. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **10**, 115 (2013), arXiv:1307.4285.
81. LHCb Collab. (R. Aaij *et al.*), *J. High Energy Phys.* **11**, 031 (2012), arXiv:1209.0282.
82. M. Adinolfi *et al.*, *Eur. Phys. J. C* **73**, 2431 (2013), arXiv:1211.6759.
83. F. Archilli *et al.*, *JINST* **8**, P10020 (2013), arXiv:1306.0249.
84. A. Hoecker *et al.*, *PoS (ACAT2007) 040* (2007), arXiv:physics/0703039.
85. LHCb HLT Group (J. Albrecht, V. Gligorov, G. Raven and S. Tolk), *J. Phys. Conf. Ser.* **513**, 012001 (2014), arXiv:1310.8544.
86. A. Puig, The LHCb trigger in 2011 and 2012, LHCb-PUB-2014-046, CERN-LHCb-PUB-2014-046.
87. S. Tolk, J. Albrecht, F. Dettori and A. Pellegrino, Data driven trigger efficiency determination at LHCb, LHCb-PUB-2014-039.
88. R. Aaij and J. Albrecht, Muon triggers in the high level trigger of LHCb, LHCb-PUB-2011-017.
89. SLD Collab. (K. Abe *et al.*), *Phys. Rev. Lett.* **80**, 660 (1998), arXiv:hep-ex/9708015.
90. V. V. Gligorov and M. Williams, *JINST* **8**, P02013 (2013), arXiv:1210.6861.
91. R. Aaij *et al.*, *JINST* **8**, P04022 (2013), arXiv:1211.3055.
92. D. Campora, N. Neufeld and R. Schwemmer, Improvements in the LHCb DAQ, in *Proceedings of the 19th IEEE Real Time Conference*, Nara, Japan, 2014, Technical Report.
93. M. Frank *et al.*, *J. Phys. Conf. Ser.* **513**, 012006 (2014).
94. LHCb Collab., Letter of intent for the LHCb upgrade, CERN-LHCC-2011-001. LHC-C-I-018.
95. LHCb Collab., Framework TDR for the LHCb Upgrade: Technical Design Report, CERN-LHCC-2012-007, LHCb-TDR-012.
96. LHCb Collab., LHCb VELO Upgrade Technical Design Report, CERN-LHCC-2013-021, LHCb-TDR-013.
97. LHCb Collab., LHCb PID Upgrade Technical Design Report, CERN-LHCC-2013-022, LHCb-TDR-014.
98. LHCb Collab., LHCb Tracker Upgrade Technical Design Report, CERN-LHCC-2014-001. LHCb-TDR-015.
99. LHCb Collab., LHCb Trigger and Online Upgrade Technical Design Report, CERN-LHCC-2014-016. LHCb-TDR-016.

Appendix A. LHCb Collaboration

R. Aaij,⁴¹ B. Adeva,³⁷ M. Adinolfi,⁴⁶ A. Affolder,⁵² Z. Ajaltouni,⁵ S. Akar,⁶
 J. Albrecht,⁹ F. Alessio,³⁸ M. Alexander,⁵¹ S. Ali,⁴¹ G. Alkhazov,³⁰
 P. Alvarez Cartelle,³⁷ A. A. Alves, Jr.,^{25,38} S. Amato,² S. Amerio,²² Y. Amhis,⁷ L. An,³
 L. Anderlini,^{17,g} J. Anderson,⁴⁰ R. Andreassen,⁵⁷ M. Andreotti,^{16,f} J. E. Andrews,⁵⁸
 R. B. Appleby,⁵⁴ O. Aquines Gutierrez,¹⁰ F. Archilli,³⁸ A. Artamonov,³⁵ M. Artuso,⁵⁹
 E. Aslanides,⁶ G. Auriemma,^{25,n} M. Baalouch,⁵ S. Bachmann,¹¹ J. J. Back,⁴⁸
 A. Badalov,³⁶ C. Baesso,⁶⁰ W. Baldini,¹⁶ R. J. Barlow,⁵⁴ C. Barschel,³⁸ S. Barsuk,⁷
 W. Barter,⁴⁷ V. Batozskaya,²⁸ V. Battista,³⁹ A. Bay,³⁹ L. Beaucourt,⁴ J. Beddow,⁵¹
 F. Bedeschi,²³ I. Bediaga,¹ S. Belogurov,³¹ K. Belous,³⁵ I. Belyaev,³¹ E. Ben-Haim,⁸
 G. Bencivenni,¹⁸ S. Benson,³⁸ J. Benton,⁴⁶ A. Berezhnoy,³² R. Bernet,⁴⁰
 M.-O. Bettler,⁴⁷ M. van Beuzekom,⁴¹ A. Bien,¹¹ S. Bifani,⁴⁵ T. Bird,⁵⁴ A. Bizzeti,^{17,i}
 P. M. Bjørnstad,⁵⁴ T. Blake,⁴⁸ F. Blanc,³⁹ J. Blouw,¹⁰ S. Blusk,⁵⁹ V. Bocci,²⁵
 A. Bondar,³⁴ N. Bondar,^{30,38} W. Bonivento,^{15,38} S. Borghi,⁵⁴ A. Borgia,⁵⁹ M. Borsato,⁷
 T. J. V. Bowcock,⁵² E. Bowen,⁴⁰ C. Bozzi,¹⁶ T. Brambach,⁹ J. Bressieux,³⁹ D. Brett,⁵⁴
 M. Britsch,¹⁰ T. Britton,⁵⁹ J. Brodzicka,⁵⁴ N. H. Brook,⁴⁶ H. Brown,⁵² A. Bursche,⁴⁰
 J. Buytaert,³⁸ S. Cadeddu,¹⁵ R. Calabrese,^{16,f} M. Calvi,^{20,k} M. Calvo Gomez,^{36,p}
 P. Campana,^{18,38} D. Campora Perez,³⁸ A. Carbone,^{14,d} G. Carboni,^{24,l}
 R. Cardinale,^{19,38,j} A. Cardini,¹⁵ L. Carson,⁵⁰ K. Carvalho Akiba,² G. Casse,⁵²
 L. Cassina,^{20,k} L. Castillo Garcia,³⁸ M. Cattaneo,³⁸ Ch. Cauet,⁹ R. Cenci,⁵⁸
 M. Charles,⁸ Ph. Charpentier,³⁸ M. Chefdeville,⁴ S. Chen,⁵⁴ S.-F. Cheung,⁵⁵
 N. Chiapolini,⁴⁰ M. Chrzasczcz,^{40,26} K. Ciba,³⁸ X. Cid Vidal,³⁸ G. Ciezarek,⁵³
 P. E. L. Clarke,⁵⁰ M. Clemencic,³⁸ H. V. Cliff,⁴⁷ J. Closier,³⁸ V. Coco,³⁸ J. Cogan,⁶
 E. Cogneras,⁵ V. Cogoni,^{15,e} L. Cojocariu,²⁹ G. Collazuol,²² P. Collins,³⁸
 A. Comerma-Montells,¹¹ A. Contu,^{15,38} A. Cook,⁴⁶ M. Coombes,⁴⁶ S. Coquereau,⁸
 G. Corti,³⁸ M. Corvo,^{16,f} I. Counts,⁵⁶ B. Couturier,³⁸ G. A. Cowan,⁵⁰ D. C. Craik,⁴⁸
 A. C. Crocombe,⁴⁸ M. Cruz Torres,⁶⁰ S. Cunliffe,⁵³ R. Currie,⁵³ C. D'Ambrosio,³⁸
 J. Dalseno,⁴⁶ P. David,⁸ P. N. Y. David,⁴¹ A. Davis,⁵⁷ K. De Bruyn,⁴¹ S. De Capua,⁵⁴
 M. De Cian,¹¹ J. M. De Miranda,¹ L. De Paula,² W. De Silva,⁵⁷ P. De Simone,¹⁸
 D. Decamp,⁴ M. Deckenhoff,⁹ L. Del Buono,⁸ N. Deléage,⁴ D. Derkach,⁵⁵
 O. Deschamps,⁵ F. Dettori,³⁸ A. Di Canto,³⁸ H. Dijkstra,³⁸ S. Donleavy,⁵² F. Dordei,¹¹
 M. Dorigo,³⁹ A. Dosil Suárez,³⁷ D. Dossett,⁴⁸ A. Dovbnya,⁴³ K. Dreimanis,⁵²
 G. Dujany,⁵⁴ F. Dupertuis,³⁹ P. Durante,³⁸ R. Dzhelyadin,³⁵ A. Dziurda,²⁶
 A. Dzyuba,³⁰ S. Easo,^{49,38} U. Egede,⁵³ V. Egorychev,³¹ S. Eidelman,³⁴ S. Eisenhardt,⁵⁰
 U. Eitschberger,⁹ R. Ekelhof,⁹ L. Eklund,⁵¹ I. El Rifai,⁵ Ch. Elsasser,⁴⁰ S. Ely,⁵⁹
 S. Esen,¹¹ H. M. Evans,⁴⁷ T. Evans,⁵⁵ A. Falabella,¹⁴ C. Färber,¹¹ C. Farinelli,⁴¹
 N. Farley,⁴⁵ S. Farry,⁵² R. Fay,⁵² D. Ferguson,⁵⁰ V. Fernandez Albor,³⁷
 F. Ferreira Rodrigues,¹ M. Ferro-Luzzi,³⁸ S. Filippov,³³ M. Fiore,^{16,f} M. Fiorini,^{16,f}
 M. Firlej,²⁷ C. Fitzpatrick,³⁹ T. Fiutowski,²⁷ P. Fol,⁵³ M. Fontana,¹⁰ F. Fontanelli,^{19,j}
 R. Forty,³⁸ O. Francisco,² M. Frank,³⁸ C. Frei,³⁸ M. Frosini,¹⁷ J. Fu,^{21,38} E. Furfaro,^{24,l}
 A. Gallas Torreira,³⁷ D. Galli,^{14,d} S. Gallorini,^{22,38} S. Gambetta,^{19,j} M. Gandelman,²
 P. Gandini,⁵⁹ Y. Gao,³ J. García Pardiñas,³⁷ J. Garofoli,⁵⁹ J. Garra Tico,⁴⁷

L. Garrido,³⁶ D. Gascon,³⁶ C. Gaspar,³⁸ R. Gauld,⁵⁵ L. Gavardi,⁹ A. Geraci,^{21,v}
 E. Gersabeck,¹¹ M. Gersabeck,⁵⁴ T. Gershon,⁴⁸ Ph. Ghez,⁴ A. Gianelle,²² S. Gianì,³⁹
 V. Gibson,⁴⁷ L. Giubega,²⁹ V. V. Gligorov,³⁸ C. Göbel,⁶⁰ D. Golubkov,³¹
 A. Golutvin,^{53,31,38} A. Gomes,^{1,a} C. Gotti,^{20,k} M. Grabalosa Gándara,⁵
 R. Graciani Diaz,³⁶ L. A. Granado Cardoso,³⁸ E. Graugés,³⁶ E. Graverini,⁴⁰
 G. Graziani,¹⁷ A. Grecu,²⁹ E. Greening,⁵⁵ S. Gregson,⁴⁷ P. Griffith,⁴⁵ L. Grillo,¹¹
 O. Grünberg,⁶² B. Gui,⁵⁹ E. Gushchin,³³ Yu. Guz,^{35,38} T. Gys,³⁸ C. Hadjivasiliou,⁵⁹
 G. Haefeli,³⁹ C. Haen,³⁸ S. C. Haines,⁴⁷ S. Hall,⁵³ B. Hamilton,⁵⁸ T. Hampson,⁴⁶
 X. Han,¹¹ S. Hansmann-Menzemer,¹¹ N. Harnew,⁵⁵ S. T. Harnew,⁴⁶ J. Harrison,⁵⁴
 J. He,³⁸ T. Head,³⁸ V. Heijne,⁴¹ K. Hennessy,⁵² P. Henrard,⁵ L. Henry,⁸
 J. A. Hernando Morata,³⁷ E. van Herwijnen,³⁸ M. Heß,⁶² A. Hicheur,² D. Hill,⁵⁵
 M. Hoballah,⁵ C. Hombach,⁵⁴ W. Hulsbergen,⁴¹ P. Hunt,⁵⁵ N. Hussain,⁵⁵
 D. Hutchcroft,⁵² D. Hynds,⁵¹ M. Idzik,²⁷ P. Ilten,⁵⁶ R. Jacobsson,³⁸ A. Jaeger,¹¹
 J. Jalocha,⁵⁵ E. Jans,⁴¹ P. Jaton,³⁹ A. Jawahery,⁵⁸ F. Jing,³ M. John,⁵⁵ D. Johnson,³⁸
 C. R. Jones,⁴⁷ C. Joram,³⁸ B. Jost,³⁸ N. Jurik,⁵⁹ S. Kandybei,⁴³ W. Kanso,⁶
 M. Karacson,³⁸ T. M. Karbach,³⁸ S. Karodia,⁵¹ M. Kelsey,⁵⁹ I. R. Kenyon,⁴⁵
 T. Ketel,⁴² B. Khanji,^{20,k} C. Khurewathanakul,³⁹ S. Klaver,⁵⁴ K. Klimaszewski,²⁸
 O. Kochebina,⁷ M. Kolpin,¹¹ I. Komarov,³⁹ R. F. Koopman,⁴² P. Koppenburg,^{41,38}
 M. Korolev,³² A. Kozlinskiy,⁴¹ L. Kravchuk,³³ K. Kreplin,¹¹ M. Kreps,⁴⁸ G. Krocker,¹¹
 P. Krokovny,³⁴ F. Kruse,⁹ W. Kucewicz,^{26,o} M. Kucharczyk,^{20,26,k} V. Kudryavtsev,³⁴
 K. Kurek,²⁸ T. Kvaratskheliya,³¹ V. N. La Thi,³⁹ D. Lacarrere,³⁸ G. Lafferty,⁵⁴
 A. Lai,¹⁵ D. Lambert,⁵⁰ R. W. Lambert,⁴² G. Lanfranchi,¹⁸ C. Langenbruch,⁴⁸
 B. Langhans,³⁸ T. Latham,⁴⁸ C. Lazzeroni,⁴⁵ R. Le Gac,⁶ J. van Leerdam,⁴¹
 J.-P. Lees,⁴ R. Lefèvre,⁵ A. Leflat,³² J. Lefrançois,⁷ S. Leo,²³ O. Leroy,⁶ T. Lesiak,²⁶
 B. Leverington,¹¹ Y. Li,³ T. Likhomanenko,⁶³ M. Liles,⁵² R. Lindner,³⁸ C. Linn,³⁸
 F. Lionetto,⁴⁰ B. Liu,¹⁵ S. Lohn,³⁸ I. Longstaff,⁵¹ J. H. Lopes,² N. Lopez-March,³⁹
 P. Lowdon,⁴⁰ D. Lucchesi,^{22,r} H. Luo,⁵⁰ A. Lupato,²² E. Luppi,^{16,f} O. Lupton,⁵⁵
 F. Machefert,⁷ I. V. Machikhiliyan,³¹ F. Maciuc,²⁹ O. Maev,³⁰ S. Malde,⁵⁵
 A. Malinin,⁶³ G. Manca,^{15,e} G. Mancinelli,⁶ A. Mapelli,³⁸ J. Maratas,⁵
 J. F. Marchand,⁴ U. Marconi,¹⁴ C. Marin Benito,³⁶ P. Marino,^{23,t} R. Märki,³⁹
 J. Marks,¹¹ G. Martellotti,²⁵ A. Martens,⁸ A. Martín Sánchez,⁷ M. Martinelli,³⁹
 D. Martinez Santos,^{42,38} F. Martinez Vidal,⁶⁴ D. Martins Tostes,² A. Massafferri,¹
 R. Matev,³⁸ Z. Mathe,³⁸ C. Matteuzzi,²⁰ A. Mazurov,⁴⁵ M. McCann,⁵³ J. McCarthy,⁴⁵
 A. McNab,⁵⁴ R. McNulty,¹² B. McSkelly,⁵² B. Meadows,⁵⁷ F. Meier,⁹ M. Meissner,¹¹
 M. Merk,⁴¹ D. A. Milanes,⁸ M.-N. Minard,⁴ N. Moggi,¹⁴ J. Molina Rodriguez,⁶⁰
 S. Monteil,⁵ M. Morandin,²² P. Morawski,²⁷ A. Mordà,⁶ M. J. Morello,^{23,t} J. Moron,²⁷
 A.-B. Morris,⁵⁰ R. Mountain,⁵⁹ F. Muheim,⁵⁰ K. Müller,⁴⁰ M. Mussini,¹⁴ B. Muster,³⁹
 P. Naik,⁴⁶ T. Nakada,³⁹ R. Nandakumar,⁴⁹ I. Nasteva,² M. Needham,⁵⁰ N. Neri,²¹
 S. Neubert,³⁸ N. Neufeld,³⁸ M. Neuner,¹¹ A. D. Nguyen,³⁹ T. D. Nguyen,³⁹
 C. Nguyen-Mau,^{39,q} M. Nicol,⁷ V. Niess,⁵ R. Niet,⁹ N. Nikitin,³² T. Nikodem,¹¹
 A. Novoselov,³⁵ D. P. O'Hanlon,⁴⁸ A. Oblakowska-Mucha,^{27,38} V. Obraztsov,³⁵
 S. Oggero,⁴¹ S. Ogilvy,⁵¹ O. Okhrimenko,⁴⁴ R. Oldeman,^{15,e} C. J. G. Onderwater,⁶⁵
 M. Orlandea,²⁹ J. M. Otalora Goicochea,² P. Owen,⁵³ A. Oyanguren,⁶⁴ B. K. Pal,⁵⁹

A. Palano,^{13,c} F. Palombo,^{21,u} M. Palutan,¹⁸ J. Panman,³⁸ A. Papanestis,^{49,38}
 M. Pappagallo,⁵¹ L. L. Pappalardo,^{16,f} C. Parkes,⁵⁴ C. J. Parkinson,^{9,45} G. Passaleva,¹⁷
 G. D. Patel,⁵² M. Patel,⁵³ C. Patrignani,^{19,j} A. Pearce,^{54,49} A. Pellegrino,⁴¹
 G. Penso,^{25,m} M. Pepe Altarelli,³⁸ S. Perazzini,^{14,d} P. Perret,⁵ M. Perrin-Terrin,⁶
 L. Pescatore,⁴⁵ E. Pesen,⁶⁶ G. Pessina,²⁰ K. Petridis,⁵³ A. Petrolini,^{19,j}
 E. Picatoste Olloqui,³⁶ B. Pietrzyk,⁴ T. Pilař,⁴⁸ D. Pinci,²⁵ A. Pistone,¹⁹ S. Playfer,⁵⁰
 M. Plo Casasus,³⁷ F. Polci,⁸ A. Poluektov,^{48,34} I. Polyakov,³¹ E. Polycarpo,²
 A. Popov,³⁵ D. Popov,¹⁰ B. Popovici,²⁹ C. Potterat,² E. Price,⁴⁶ J. D. Price,⁵²
 J. Prisciandaro,³⁹ A. Pritchard,⁵² C. Prouve,⁴⁶ V. Pugatch,⁴⁴ A. Puig Navarro,³⁹
 G. Punzi,^{23,s} W. Qian,⁴ B. Rachwal,²⁶ J. H. Rademacker,⁴⁶ B. Rakotomiamanana,³⁹
 M. Rama,¹⁸ M. S. Rangel,² I. Raniuk,⁴³ N. Rauschmayr,³⁸ G. Raven,⁴² F. Redi,⁵³
 S. Reichert,⁵⁴ M. M. Reid,⁴⁸ A. C. dos Reis,¹ S. Ricciardi,⁴⁹ S. Richards,⁴⁶ M. Rihl,³⁸
 K. Rinnert,⁵² V. Rives Molina,³⁶ P. Robbe,⁷ A. B. Rodrigues,¹ E. Rodrigues,⁵⁴
 P. Rodriguez Perez,⁵⁴ S. Roiser,³⁸ V. Romanovsky,³⁵ A. Romero Vidal,³⁷
 M. Rotondo,²² J. Rouvinet,³⁹ T. Ruf,³⁸ H. Ruiz,³⁶ P. Ruiz Valls,⁶⁴
 J. J. Saborido Silva,³⁷ N. Sagidova,³⁰ P. Sail,⁵¹ B. Saitta,^{15,e} V. Salustino Guimaraes,²
 C. Sanchez Mayordomo,⁶⁴ B. Sanmartin Sedes,³⁷ R. Santacesaria,²⁵
 C. Santamarina Rios,³⁷ E. Santovetti,^{24,l} A. Sarti,^{18,m} C. Satriano,^{25,n} A. Satta,²⁴
 D. M. Saunders,⁴⁶ D. Savrina,^{31,32} M. Schiller,⁴² H. Schindler,³⁸ M. Schlupp,⁹
 M. Schmelling,¹⁰ B. Schmidt,³⁸ O. Schneider,³⁹ A. Schopper,³⁸ M.-H. Schune,⁷
 R. Schwemmer,³⁸ B. Sciascia,¹⁸ A. Sciubba,^{25,m} A. Semennikov,³¹ I. Sepp,⁵³ N. Serra,⁴⁰
 J. Serrano,⁶ L. Sestini,²² P. Seyfert,¹¹ M. Shapkin,³⁵ I. Shapoval,^{16,43,f} Y. Shcheglov,³⁰
 T. Shears,⁵² L. Shekhtman,³⁴ V. Shevchenko,⁶³ A. Shires,⁹ R. Silva Coutinho,⁴⁸
 G. Simi,²² M. Sirendi,⁴⁷ N. Skidmore,⁴⁶ I. Skillicorn,⁵¹ T. Skwarnicki,⁵⁹ N. A. Smith,⁵²
 E. Smith,^{55,49} E. Smith,⁵³ J. Smith,⁴⁷ M. Smith,⁵⁴ H. Snoek,⁴¹ M. D. Sokoloff,⁵⁷
 F. J. P. Soler,⁵¹ F. Soomro,³⁹ D. Souza,⁴⁶ B. Souza De Paula,² B. Spaan,⁹ A. Sparkes,⁵⁰
 P. Spradlin,⁵¹ S. Sridharan,³⁸ F. Stagni,³⁸ M. Stahl,¹¹ S. Stahl,¹¹ O. Steinkamp,⁴⁰
 O. Stenyakin,³⁵ S. Stevenson,⁵⁵ S. Stoica,²⁹ S. Stone,⁵⁹ B. Storaci,⁴⁰ S. Stracka,^{23,t}
 M. Straticiu,²⁹ U. Straumann,⁴⁰ R. Stroili,²² V. K. Subbiah,³⁸ L. Sun,⁵⁷
 W. Sutcliffe,⁵³ K. Swientek,²⁷ S. Swientek,⁹ V. Syropoulos,⁴² M. Szczekowski,²⁸
 P. Szczypka,^{39,38} T. Szumlak,²⁷ S. T'Jampens,⁴ M. Teklishyn,⁷ G. Tellarini,^{16,f}
 F. Teubert,³⁸ C. Thomas,⁵⁵ E. Thomas,³⁸ J. van Tilburg,⁴¹ V. Tisserand,⁴ M. Tobin,³⁹
 S. Tolk,⁴² L. Tomassetti,^{16,f} D. Tonelli,³⁸ S. Topp-Joergensen,⁵⁵ N. Torr,⁵⁵
 E. Tournefier,⁴ S. Tourneur,³⁹ M. T. Tran,³⁹ M. Tresch,⁴⁰ A. Tsaregorodtsev,⁶
 P. Tsopelas,⁴¹ N. Tuning,⁴¹ M. Ubeda Garcia,³⁸ A. Ukleja,²⁸ A. Ustyuzhanin,⁶³
 U. Uwer,¹¹ C. Vacca,^{15,e} V. Vagnoni,¹⁴ G. Valenti,¹⁴ A. Vallier,⁷ R. Vazquez Gomez,¹⁸
 P. Vazquez Regueiro,³⁷ C. Vázquez Sierra,³⁷ S. Vecchi,¹⁶ J. J. Velthuis,⁴⁶ M. Veltri,^{17,h}
 G. Veneziano,³⁹ M. Vesterinen,¹¹ B. Viaud,⁷ D. Vieira,² M. Vieites Diaz,³⁷
 X. Vilasis-Cardona,^{36,p} A. Vollhardt,⁴⁰ D. Volyanskyy,¹⁰ D. Voong,⁴⁶ A. Vorobyev,³⁰
 V. Vorobyev,³⁴ C. Voß,⁶² J. A. de Vries,⁴¹ R. Waldi,⁶² C. Wallace,⁴⁸ R. Wallace,¹²
 J. Walsh,²³ S. Wandernoth,¹¹ J. Wang,⁵⁹ D. R. Ward,⁴⁷ N. K. Watson,⁴⁵
 D. Websdale,⁵³ M. Whitehead,⁴⁸ J. Wicht,³⁸ D. Wiedner,¹¹ G. Wilkinson,^{55,38}
 M. P. Williams,^{48,49} M. Williams,⁵⁶ H. W. Wilschut,⁶⁵ F. F. Wilson,⁴⁹ J. Wimberley,⁵⁸

J. Wishahi,⁹ W. Wislicki,²⁸ M. Witek,²⁶ G. Wormser,⁷ S. A. Wotton,⁴⁷ S. Wright,⁴⁷
K. Wyllie,³⁸ Y. Xie,⁶¹ Z. Xing,⁵⁹ Z. Xu,³⁹ Z. Yang,³ X. Yuan,³ O. Yushchenko,³⁵
M. Zangoli,¹⁴ M. Zavertyaev,^{10,b} L. Zhang,⁵⁹ W. C. Zhang,¹² Y. Zhang,³ A. Zhelezov,¹¹
A. Zhokhov,³¹ L. Zhong,³ A. Zvyagin³⁸

¹ Centro Brasileiro de Pesquisas Físicas (CBPF), Rio de Janeiro, Brazil

² Universidade Federal do Rio de Janeiro (UFRJ), Rio de Janeiro, Brazil

³ Center for High Energy Physics, Tsinghua University, Beijing, China

⁴ LAPP, Université de Savoie, CNRS/IN2P3, Annecy-Le-Vieux, France

⁵ Clermont Université, Université Blaise Pascal, CNRS/IN2P3, LPC, Clermont-Ferrand, France

⁶ CPPM, Aix-Marseille Université, CNRS/IN2P3, Marseille, France

⁷ LAL, Université Paris-Sud, CNRS/IN2P3, Orsay, France

⁸ LPNHE, Université Pierre et Marie Curie, Université Paris Diderot, CNRS/IN2P3, Paris, France

⁹ Fakultät Physik, Technische Universität Dortmund, Dortmund, Germany

¹⁰ Max-Planck-Institut für Kernphysik (MPIK), Heidelberg, Germany

¹¹ Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany

¹² School of Physics, University College Dublin, Dublin, Ireland

¹³ Sezione INFN di Bari, Bari, Italy

¹⁴ Sezione INFN di Bologna, Bologna, Italy

¹⁵ Sezione INFN di Cagliari, Cagliari, Italy

¹⁶ Sezione INFN di Ferrara, Ferrara, Italy

¹⁷ Sezione INFN di Firenze, Firenze, Italy

¹⁸ Laboratori Nazionali dell'INFN di Frascati, Frascati, Italy

¹⁹ Sezione INFN di Genova, Genova, Italy

²⁰ Sezione INFN di Milano Bicocca, Milano, Italy

²¹ Sezione INFN di Milano, Milano, Italy

²² Sezione INFN di Padova, Padova, Italy

²³ Sezione INFN di Pisa, Pisa, Italy

²⁴ Sezione INFN di Roma Tor Vergata, Roma, Italy

²⁵ Sezione INFN di Roma La Sapienza, Roma, Italy

²⁶ Henryk Niewodniczanski Institute of Nuclear Physics Polish Academy of Sciences, Kraków, Poland

²⁷ AGH – University of Science and Technology, Faculty of Physics and Applied Computer Science, Kraków, Poland

²⁸ National Center for Nuclear Research (NCBJ), Warsaw, Poland

²⁹ Horia Hulubei National Institute of Physics and Nuclear Engineering, Bucharest-Magurele, Romania

³⁰ Petersburg Nuclear Physics Institute (PNPI), Gatchina, Russia

³¹ Institute of Theoretical and Experimental Physics (ITEP), Moscow, Russia

³² Institute of Nuclear Physics, Moscow State University (SINP MSU), Moscow, Russia

³³ Institute for Nuclear Research of the Russian Academy of Sciences (INR RAN), Moscow, Russia

³⁴ Budker Institute of Nuclear Physics (SB RAS) and Novosibirsk State University, Novosibirsk, Russia

³⁵ Institute for High Energy Physics (IHEP), Protvino, Russia

³⁶ Universitat de Barcelona, Barcelona, Spain

³⁷ Universidad de Santiago de Compostela, Santiago de Compostela, Spain

³⁸ European Organization for Nuclear Research (CERN), Geneva, Switzerland

³⁹ Ecole Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

⁴⁰ Physik-Institut, Universität Zürich, Zürich, Switzerland

⁴¹ Nikhef National Institute for Subatomic Physics, Amsterdam, The Netherlands

⁴² Nikhef National Institute for Subatomic Physics and VU University Amsterdam, Amsterdam, The Netherlands

⁴³ NSC Kharkiv Institute of Physics and Technology (NSC KIPT), Kharkiv, Ukraine

- ⁴⁴*Institute for Nuclear Research of the National Academy of Sciences (KINR), Kyiv, Ukraine*
- ⁴⁵*University of Birmingham, Birmingham, United Kingdom*
- ⁴⁶*H.H. Wills Physics Laboratory, University of Bristol, Bristol, United Kingdom*
- ⁴⁷*Cavendish Laboratory, University of Cambridge, Cambridge, United Kingdom*
- ⁴⁸*Department of Physics, University of Warwick, Coventry, United Kingdom*
- ⁴⁹*STFC Rutherford Appleton Laboratory, Didcot, United Kingdom*
- ⁵⁰*School of Physics and Astronomy, University of Edinburgh, Edinburgh, United Kingdom*
- ⁵¹*School of Physics and Astronomy, University of Glasgow, Glasgow, United Kingdom*
- ⁵²*Oliver Lodge Laboratory, University of Liverpool, Liverpool, United Kingdom*
- ⁵³*Imperial College London, London, United Kingdom*
- ⁵⁴*School of Physics and Astronomy, University of Manchester, Manchester, United Kingdom*
- ⁵⁵*Department of Physics, University of Oxford, Oxford, United Kingdom*
- ⁵⁶*Massachusetts Institute of Technology, Cambridge, MA, United States*
- ⁵⁷*University of Cincinnati, Cincinnati, OH, United States*
- ⁵⁸*University of Maryland, College Park, MD, United States*
- ⁵⁹*Syracuse University, Syracuse, NY, United States*
- ⁶⁰*Pontifícia Universidade Católica do Rio de Janeiro (PUC-Rio), Rio de Janeiro, Brazil, associated to ²*
- ⁶¹*Institute of Particle Physics, Central China Normal University, Wuhan, Hubei, China, associated to ³*
- ⁶²*Institut für Physik, Universität Rostock, Rostock, Germany, associated to ¹¹*
- ⁶³*National Research Centre Kurchatov Institute, Moscow, Russia, associated to ³¹*
- ⁶⁴*Instituto de Física Corpuscular (IFIC), Universitat de Valencia-CSIC, Valencia, Spain, associated to ³⁶*
- ⁶⁵*Van Swinderen Institute, University of Groningen, Groningen, The Netherlands, associated to ⁴¹*
- ⁶⁶*Celal Bayar University, Manisa, Turkey, associated to ³⁸*
- ^a*Universidade Federal do Triângulo Mineiro (UFMT), Uberaba-MG, Brazil*
- ^b*P.N. Lebedev Physical Institute, Russian Academy of Science (LPI RAS), Moscow, Russia*
- ^c*Università di Bari, Bari, Italy*
- ^d*Università di Bologna, Bologna, Italy*
- ^e*Università di Cagliari, Cagliari, Italy*
- ^f*Università di Ferrara, Ferrara, Italy*
- ^g*Università di Firenze, Firenze, Italy*
- ^h*Università di Urbino, Urbino, Italy*
- ⁱ*Università di Modena e Reggio Emilia, Modena, Italy*
- ^j*Università di Genova, Genova, Italy*
- ^k*Università di Milano Bicocca, Milano, Italy*
- ^l*Università di Roma Tor Vergata, Roma, Italy*
- ^m*Università di Roma La Sapienza, Roma, Italy*
- ⁿ*Università della Basilicata, Potenza, Italy*
- ^o*AGH – University of Science and Technology, Faculty of Computer Science, Electronics and Telecommunications, Kraków, Poland*
- ^p*LIFAEELS, La Salle, Universitat Ramon Llull, Barcelona, Spain*
- ^q*Hanoi University of Science, Hanoi, Vietnam*
- ^r*Università di Padova, Padova, Italy*
- ^s*Università di Pisa, Pisa, Italy*
- ^t*Scuola Normale Superiore, Pisa, Italy*
- ^u*Università degli Studi di Milano, Milano, Italy*
- ^v*Politecnico di Milano, Milano, Italy*