

The LHCb Experiment

L. de Paula^{1*}

¹ *LAPE, Instituto de Física, Universidade Federal do Rio de Janeiro,
C.P. 68528 - Rio de Janeiro, 21945-970, Brazil
E-mail: leandro@if.ufrj.br*

Received 7 January, 2000

An overview of the LHCb experiment, which has been approved by CERN in september 1998, is presented [1].

I Physics Motivation

I.1 Introduction

CP violation was first discovered in neutral kaon decays in 1964 [2]. Its origin is still one of the outstanding mysteries of elementary particle physics.

The Standard Model with three quark families can naturally generate CP violation in both weak and strong interactions although CP violation in the strong interaction has never been detected. CP violation in the weak interaction is generated by the complex three-by-three unitary matrix, known as the CKM matrix, introduced by Kobayashi and Maskawa [3]. Observed CP-violating phenomena in the neutral-kaon system are consistent with this mechanism. However, it cannot be excluded that physics beyond the Standard Model contributes, or even fully accounts for the observed phenomena.

CP violation also plays an important role in cosmology. It is one of the three ingredients required to explain the excess of matter over antimatter observed in our universe [4]. The level of CP violation that can be generated by the Standard Model weak interaction is insufficient to explain the dominance of matter in the universe [5]. This calls for new sources of CP violation beyond the Standard Model.

Since its discovery, CP violation has been detected only in the decay amplitude of K_L mesons. Experimental efforts in the kaon sector will continue for some time. In the B-meson system there are many more decay modes available, and the Standard Model makes precise predictions for CP violation in a number of these. The B-meson system is therefore a very attractive place to study CP violation, and to search for a hint of new physics.

CP violation

The elements of the CKM matrix

$$V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix},$$

V_{ij} , are related to the relative strengths of the transition of down-type quarks ($j = d, s, b$) to up-type quarks ($i = u, c, t$), normalised to $\sqrt{G_F}$, where G_F is the Fermi coupling constant. The matrix is uniquely defined by four parameters. Among various parameterisations, the most convenient one is that proposed by Wolfenstein [6]:

$$V_{\text{CKM}} \approx V_{\text{CKM}}^{(3)} + \delta V_{\text{CKM}}$$

where the expansion up to third order in λ is given by $V_{\text{CKM}}^{(3)}$

$$\begin{pmatrix} 1 - \lambda^2/2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \lambda^2/2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix}.$$

The parameter λ is given by the sine of the Cabibbo angle [7], measured to be 0.221 ± 0.002 [8] from decays involving s-quarks.

For a qualitative discussion of CP violation in B-meson systems, $V_{\text{CKM}}^{(3)}$ is sufficient and the second term δV_{CKM} , which is given by

$$\begin{pmatrix} 0 & 0 & 0 \\ -iA^2\lambda^5\eta & 0 & 0 \\ A(\rho + i\eta)\lambda^5/2 & (1/2 - \rho)A\lambda^4 - iA\lambda^4\eta & 0 \end{pmatrix},$$

is usually ignored. For CP violation in K^0 - \bar{K}^0 oscillations, the correction to V_{cd} is important. For B-meson

*On behalf of LHCb collaboration.

systems, the correction to V_{td} and V_{ts} becomes relevant once the sensitivity of experiments to measure CP-violation parameters becomes 10^{-2} or less. Note that $\eta \neq 0$ is required to generate CP violation.

Six of the nine unitarity conditions of the CKM matrix can be drawn as triangles in the complex plane. The two triangles relevant for the B-meson systems are shown in Fig. 1. The related unitarity conditions are given by

$$\begin{aligned} V_{ud}V_{ub}^* + V_{cd}V_{cb}^* + V_{td}V_{tb}^* &= 0 \\ V_{tb}V_{ub}^* + V_{ts}V_{us}^* + V_{td}V_{ud}^* &= 0. \end{aligned}$$

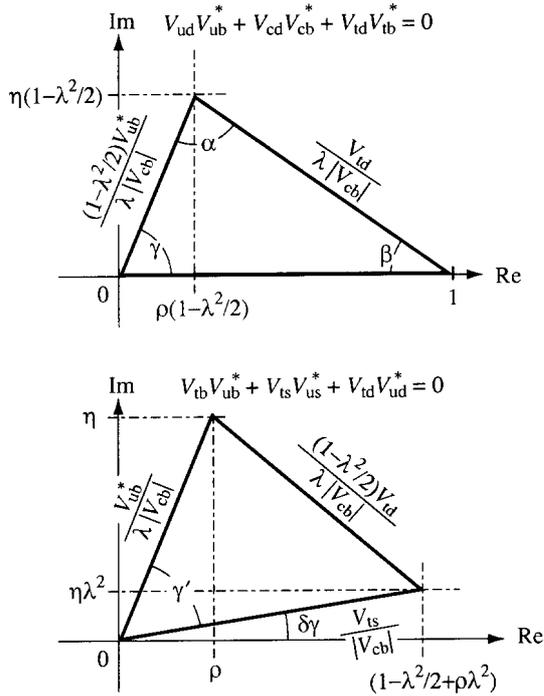


Figure 1. Two unitarity triangles in the Wolfenstein's parameterisation with an approximation valid up to $\mathcal{O}(\lambda^5)$.

The two triangles become identical if δV_{CKM} is ignored.

The angles of the triangles can be extracted either indirectly by measuring the lengths of the sides, or, within the Standard Model, directly from CP asymmetries. If the angles extracted by the two different methods disagree, this would indicate new physics.

Since λ is well known, the two triangles are completely determined by ρ and η , which can be derived from $|V_{cb}|$, $|V_{ub}|$ and $|V_{td}|$, as seen from Fig. 1. The parameter A is extracted from measurements of $|V_{cb}|$ and λ . Values of $|V_{cb}|$ and $|V_{ub}|$ are extracted from various B-meson decays and are currently known to be 0.041 ± 0.003 and 0.0033 ± 0.0009 [8], respectively. The large error on $|V_{ub}|$ is due to the limited experimental data available and theoretical uncertainties in the evaluation of strong interaction effects. Experiments at e^+e^- machines running at the $\Upsilon(4S)$, i.e. CLEO,

BABAR and BELLE, will reduce the errors on these elements. Their precisions will ultimately be limited by the theoretical uncertainties.

The value of $|V_{td}|$ is currently determined from the frequency of $B_d^0-\bar{B}_d^0$ oscillations. Due to difficulties in evaluating the effects of hadronic interactions, the extracted value has a large uncertainty, $|V_{td}| = 0.009 \pm 0.003$ [8]. This situation can be improved considerably once $|V_{ts}|$ is extracted from the frequency of $B_s^0-\bar{B}_s^0$ oscillations and $|V_{td}|/|V_{ts}|$ is used instead of $|V_{td}|$, since the Standard Model calculation of this ratio has a much reduced hadronic uncertainty. Experimentally, CDF, D0, HERA-B and SLD will try to measure the $B_s^0-\bar{B}_s^0$ oscillation frequency. However, this may not be possible before LHCb becomes operational, if the frequency is high.

Once ρ and η are derived from $|V_{cb}|$, $|V_{ub}|$ and $|V_{td}|$, the angles α , β , γ and $\delta\gamma$ can be calculated. At present, a non-zero value of η can only be obtained if CP violation in $K^0-\bar{K}^0$ oscillations is included in the analysis [9]. Future rare kaon decay experiments measuring $K \rightarrow \pi\nu\bar{\nu}$ will also provide information on ρ and η [10].

In the framework of the Standard Model, direct measurements can be made of the angles α , β , γ and $\delta\gamma$, or their combinations, from CP asymmetries in different final states of B-meson decays. Well known examples are [10]:

1. $\beta + \gamma$ from $B_d^0 \rightarrow \pi^+\pi^-$
2. β from $B_d^0 \rightarrow J/\psi K_S$
3. $\gamma - 2\delta\gamma$ from $B_s^0 \rightarrow D_s^\pm K^\mp$
4. $\delta\gamma$ from $B_s^0 \rightarrow J/\psi\phi$
5. γ from $B_d^0 \rightarrow \bar{D}^0 K^{*0}, D^0 K^{*0}, D_1 K^{*0}$,

where it is understood that the charge-conjugated decay processes are also measured, and D_1 is the CP = +1 state of the neutral D meson. Note that the angle α is not measured directly but can be determined only through the triangle relation $\alpha = \pi - \beta - \gamma$. Within the framework of the Standard Model, β , $\gamma - 2\delta\gamma$ and γ measured from the decay channels 2, 3 and 5 have very little theoretical uncertainty.

If a new flavour-changing neutral current is introduced by physics beyond the Standard Model, it can have a large effect on $B_d^0-\bar{B}_d^0$ and $B_s^0-\bar{B}_s^0$ oscillations, since the contribution of the weak interaction is second order. For such a case, the values of $|V_{td}|$ and $|V_{ts}|$ experimentally extracted from B- \bar{B} oscillations no longer correspond to their real values. The angles $\beta + \gamma$, β , $\gamma - 2\delta\gamma$ and $\delta\gamma$, extracted from the decay channels 1-4, are also affected. These angles, measured in the two ways explained above, will no longer agree.

The angle γ can be determined from channels 1 and 2 using B_d^0 decays, or from channels 3 and 4 using B_s^0

decays. Since $B_d^0\text{-}\bar{B}_d^0$ and $B_s^0\text{-}\bar{B}_s^0$ oscillations can be affected differently by the new flavour-changing neutral current, the two measurements of γ may disagree.

Since only a small contribution from the new flavour-changing neutral current is expected in $D^0\text{-}\bar{D}^0$ oscillations, γ extracted from the fifth decay mode would be very close to the Standard Model value. Therefore, γ calculated from the decay channels 2 and 3 and γ obtained from channel 5 will not agree.

In the Standard Model, $\delta\gamma$ is expected to be of the order of 10^{-2} , and the CP asymmetry in $B_s^0 \rightarrow J/\psi\phi$ decays should be very small. The new flavour-changing neutral current could, however, generate a large CP-violating effect in this decay channel.

This illustrates how new physics could be detected from precise measurements of CP violation in various B-meson decays, combined with ρ and η determined from other B-meson decays. Detailed discussion can be found elsewhere [11].

Another way to search for physics beyond the Standard Model is to study B-meson decays that are rare or even forbidden in the Standard Model. For example, B-meson decays generated by the penguin processes are first order in the weak interaction, and their branching fractions are less likely to be affected by new physics. However, these decay modes may exhibit a sizable CP-violating effect through interference, if new physics is present [12]. Similarly, a large effect could be seen in the energy asymmetry of the lepton pairs produced in $b \rightarrow sl^+\ell^-$ decays [13].

There are many ways to look for a sign of new physics. In all cases, large numbers of both B_s^0 and B_d^0 mesons are required, and many different decay modes have to be reconstructed.

1.3 LHCb performance

Compared to other accelerators that are in operation or under construction, the LHC will be by far the most copious source of B mesons, due to the high $b\bar{b}$ cross section and high luminosity. A variety of b-hadrons, such as B_u , B_d , B_s , B_c and b-baryons, are produced at high rate.

The LHCb experiment plans to operate with an average luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, which should be obtained from the beginning of LHC operation. Running at this luminosity has further advantages. The detector occupancy remains low, and radiation damage is reduced. Events are dominated by single pp interactions that are easy to analyse. The luminosity at the LHCb interaction point can be kept at its nominal value while the luminosities at the other interaction points are being progressively increased to their design values. This will allow the experiment to collect data for many years under constant conditions. About 10^{12} $b\bar{b}$ pairs are expected to be produced in one year of data taking.

The LHCb detector is designed to exploit the large number of b-hadrons produced at the LHC in order to make precision studies of CP asymmetries and of

rare decays in the B-meson systems. It has a high-performance trigger which is robust and optimised to collect B mesons efficiently, based on particles with large transverse momentum and displaced decay vertices.

Table 1. Expected numbers of events reconstructed offline in one year (10^7 s of data taking) with an average luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, for some channels.

Decay Modes	Visible Br. fraction	Offline Reconst.
$B_d^0 \rightarrow \pi^+\pi^- + \text{tag}$	0.7×10^{-5}	6.9 k
$B_d^0 \rightarrow K^+\pi^-$	1.5×10^{-5}	33 k
$B_d^0 \rightarrow \rho^+\pi^- + \text{tag}$	1.8×10^{-5}	551
$B_d^0 \rightarrow J/\psi K_S + \text{tag}$	3.6×10^{-5}	56 k
$B_d^0 \rightarrow \bar{D}^0 K^{*0}$	3.3×10^{-7}	337
$B_d^0 \rightarrow K^{*0}\gamma$	3.2×10^{-5}	26 k
$B_s^0 \rightarrow D_s^- \pi^+ + \text{tag}$	1.2×10^{-4}	35 k
$B_s^0 \rightarrow D_s^- K^+ + \text{tag}$	8.1×10^{-6}	2.1 k
$B_s^0 \rightarrow J/\psi\phi + \text{tag}$	5.4×10^{-5}	44 k

The detector can reconstruct a B-decay vertex with very good resolution and provide excellent particle identification for charged particles. Excellent vertex resolution is essential for studying the rapidly oscillating B_s mesons and in particular their CP asymmetries. It also helps to reduce combinatoric background when reconstructing rare decays.

Without separating kaons from pions, reconstructed $B_d \rightarrow \pi^+\pi^-$ decays are heavily contaminated by $B_d \rightarrow K^\pm\pi^\mp$, $B_s \rightarrow K^\mp\pi^\pm$ and $B_s \rightarrow K^\pm K^\mp$ decays. These introduce unknown systematic errors in the measured CP asymmetry in $B_d \rightarrow \pi^+\pi^-$ decays, since these decay modes may well have asymmetries too. The measurement of their asymmetries is also interesting. The ability to distinguish kaons from pions is also essential for $B_s \rightarrow D_s^\pm K^\mp$, where the main background comes from $B_s \rightarrow D_s^\pm\pi^\mp$ decays. The branching fraction of $B_s \rightarrow D_s^\pm\pi^\mp$ is estimated to be ten times larger than $B_s \rightarrow D_s^\pm K^\mp$, and no CP violation is expected. Therefore, without separating the two channels, CP asymmetries in $B_s \rightarrow D_s^\pm K^\mp$ decays would be seriously diluted. Particle identification is also needed for the reconstruction of $B_d \rightarrow DK^*$ decays, to reduce combinatoric background.

With the capabilities described above, LHCb is ideally suited to determine all the angles of the two unitarity triangles using high-statistics data. Table 1 shows the expected numbers of offline-reconstructed events for various B-meson final states in one year (10^7 s) of data taking. Simulation studies show that the LHCb detector is able to trigger and reconstruct, in addition to final states with only charged particles, also those including a photon or π^0 . This enhances the capability of the experiment to determine α without theoretical uncertainty from the penguin amplitude, and to allow the interesting radiative penguin decays to be studied.

Table 2. Expected precision on the angles of the unitarity triangles obtained by the LHCb experiment in one year of data taking. Special features of the detector, i.e. particle identification and excellent decay time resolution (σ_t), are indicated when they are important. Details can be found in Part V of this document.

Parameter	Decay Mode	σ [1 year]	Exploited features of LHCb
$\beta + \gamma$ (= $\pi - \alpha$)	B_d^0 and $\bar{B}_d^0 \rightarrow \pi^+ \pi^-$; no penguin	0.03	K/ π separation
	penguin/tree= 0.20 ± 0.02	0.03-0.16	K/ π separation
β	B_d^0 and $\bar{B}_d^0 \rightarrow J/\psi K_S$	0.01	-
$\gamma - 2\delta\gamma$	B_s^0 and $\bar{B}_s^0 \rightarrow D_s^\pm K^\mp$	0.05-0.28	K/ π separation and σ_t
γ	$B_d^0 \rightarrow \bar{D}^0 K^{*0}, D^0 K^{*0}, D_1 K^{*0}$ and $\bar{B}_d^0 \rightarrow \bar{D}^0 \bar{K}^{*0}, D^0 \bar{K}^{*0}, D_1 \bar{K}^{*0}$	0.07-0.31	K/ π separation
$\delta\gamma$	B_s^0 and $\bar{B}_s^0 \rightarrow J/\psi \phi$	0.01	σ_t
x_s	B_s^0 and $\bar{B}_s^0 \rightarrow D_s^\pm \pi^\mp$	up to 90 (95% CL)	σ_t

Table 2 summarises the expected precision on the angles of the unitarity triangles and the sensitivity to B_s^0 - \bar{B}_s^0 oscillations, obtained after one year of data taking. It also indicates the decay modes used and the important features of the LHCb detector discussed above.

In addition to investigating CP violation in B-meson decays, the physics programme of the LHCb experiment will include studies of rare B and τ decays, D- \bar{D} oscillations and B_c -meson decays. For example, the ability to reconstruct the large number of $B_d^0 \rightarrow K^{*0} \gamma$ decays given in Table 1 demonstrates that the LHCb experiment can study various other decay modes generated by the $b \rightarrow s \gamma$ process, such as $B_d^0 \rightarrow K^{**} \gamma$ and $B_s^0 \rightarrow \phi \gamma$. Reconstruction of $B_d^0 \rightarrow K^{*0} \mu^+ \mu^-$ should also be possible. The large numbers of reconstructed events expected allow searches to be made for surprising effects in these rare decay modes. Events needed to study these processes will pass the standard Level-0 to Level-2 trigger cuts. Only the Level-3 algorithm will need to be tuned accordingly.

II Detector

LHCb is a single-arm spectrometer with a forward angular coverage from approximately 10 mrad to 300 (250) mrad in the bending (non-bending) plane. The choice of the detector geometry is motivated by the fact that at high energies both the b- and \bar{b} -hadrons are predominantly produced in the same forward cone, a feature exploited in the flavour tag. This is demonstrated in Fig. 2 where the polar angles of the b- and \bar{b} -hadrons calculated by the PYTHIA event generator are plotted. The polar angle is defined with respect to the beam axis in the pp centre-of-mass system.

Fig. 3 shows the momentum distributions for $B_d^0 \rightarrow \pi^+ \pi^-$ decays into 4π , and for those where the momenta of both pions are measured in the spectrometer. The decrease of the detector acceptance for high momenta is due to the loss of particles below 10 mrad. In the low momentum region, the loss of acceptance is due to slow pions that do not hit enough tracking stations for their momenta to be measured.

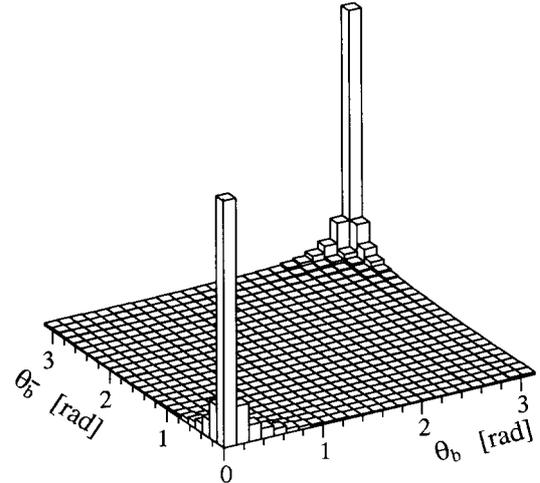


Figure 2. Polar angles of the b- and \bar{b} -hadrons calculated by the PYTHIA event generator.

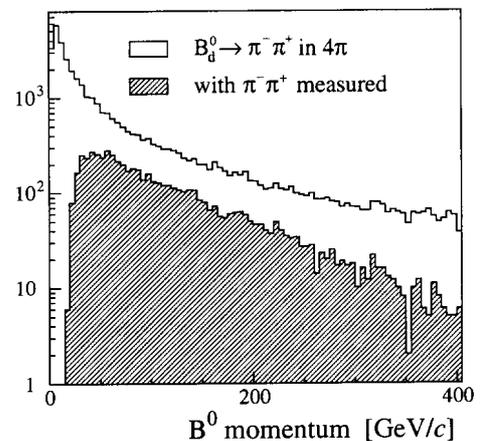


Figure 3. Momentum distributions for $B_d^0 \rightarrow \pi^+ \pi^-$ decays into 4π , and for those where the momenta of both pions are measured in the spectrometer.

To determine the momentum range required for the spectrometer, the $B_s^0 \rightarrow D_s^0 \pi^+$ decay is studied. The π^+ defines the high end of the momentum range, and the π^- from the D_s^- decay the low end. Fig. 2 shows the momentum distributions for both pions, when they are within the spectrometer acceptance. Few tracks have momenta beyond 150 GeV/c.

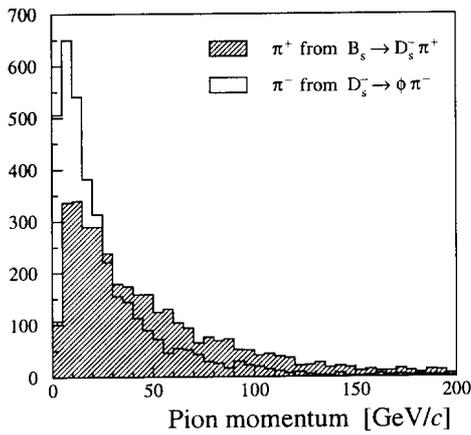


Figure 4. Momentum distributions for the π^+ from $B_s^0 \rightarrow D_s^- \pi^+$ decays and for the π^- from the subsequent $D_s^- \rightarrow \phi \pi^-$ decay, where the momenta are measured in the spectrometer.

The requirements for the vertex detector can be illustrated by the decay-length distribution of $B_d^0 \rightarrow \pi^+ \pi^-$ shown in Fig. 5. The average decay length for the detected $B_d^0 \rightarrow \pi^+ \pi^-$ is 1.0 cm.

II.1 General layout

The layout of the LHCb spectrometer is shown in Fig. 6. Intersection Point 8, currently used by DELPHI, has been allocated to the experiment. A modification to the LHC optics, displacing the interaction point by 11.25 m from the centre, has permitted maximum use to be made of the existing cavern by freeing 19.7 m for the LHCb detector components. A right-handed coordinate system is defined centred on the interaction point, with z along the beam axis and y pointing upwards.

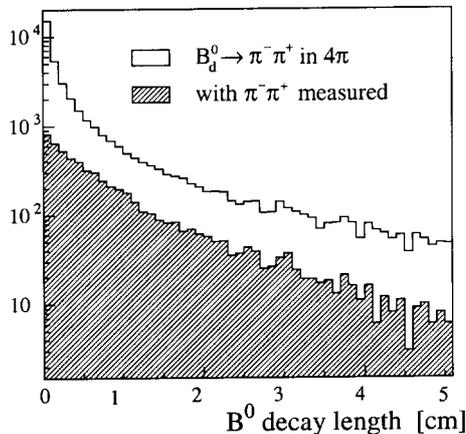


Figure 5. Decay-length distributions for $B_d^0 \rightarrow \pi^+ \pi^-$ decays in 4π , and those where both pions are detected in the spectrometer.

LHCb comprises a vertex detector system (including a pile-up veto counter), a tracking system (partially inside a dipole magnet), aerogel and gas RICH counters, an electromagnetic calorimeter with preshower detector, a hadron calorimeter and a muon detector. All

detector subsystems, except the vertex detector, are assembled in two halves, which can be separated horizontally for assembly and maintenance, as well as to provide access to the beam pipe.

II.2 Beam pipe

A 1.8 m-long section of the beam pipe around the interaction point has a large diameter of approximately 120 cm. This accommodates the vertex detector system with its retraction mechanics, and has a thin forward window over the full detector acceptance. This part is followed by two conical sections; the first is 1.5 m long with 25 mrad opening angle, and the second is 16 m long with 10 mrad opening angle. The current design is based on aluminium walls, but partial replacement by beryllium is under study.

II.3 Magnet

The spectrometer dipole is placed close to the interaction region, in order to keep its size small. Since tracks in the vertex detector are used in the trigger, it is desirable to have the vertex detector in a region of low magnetic field for fast track finding. The first RICH counter is designed to cover the momentum range down to 1 GeV/c. To maintain the necessary RICH 1 acceptance and to avoid tracks bending in the gas radiator, RICH 1 is also required to be in a region of low magnetic field. The magnet is therefore placed behind RICH 1, allowing an acceptance of 330 mrad in both projections upstream of the magnet.

A superconducting magnet is chosen to obtain a high field integral of 4 Tm with a short length. It benefits from the existing infrastructure of the DELPHI solenoid. The field is oriented vertically and has a maximum value of 1.1 T. The polarity of the field can be changed to reduce systematic errors in the CP-violation measurements that could result from a left-right asymmetry of the detector. The free aperture is 4.3 m horizontally and 3.6 m vertically. The coil is designed to maximise the field homogeneity. An iron shield upstream of the magnet reduces the stray field in the vicinity of the vertex detector and of RICH 1.

II.4 Vertex detector system

The vertex detector system comprises a silicon vertex detector and a pile-up veto counter. The vertex detector has to provide precise information on the production and decay vertices of b-hadrons both offline and for the Level-1 trigger. The latter requires all channels to be read out within 1 μ s. The pile-up veto counter is used in the Level-0 trigger to suppress events containing multiple pp interactions in a single bunch-crossing, by counting the number of primary vertices.

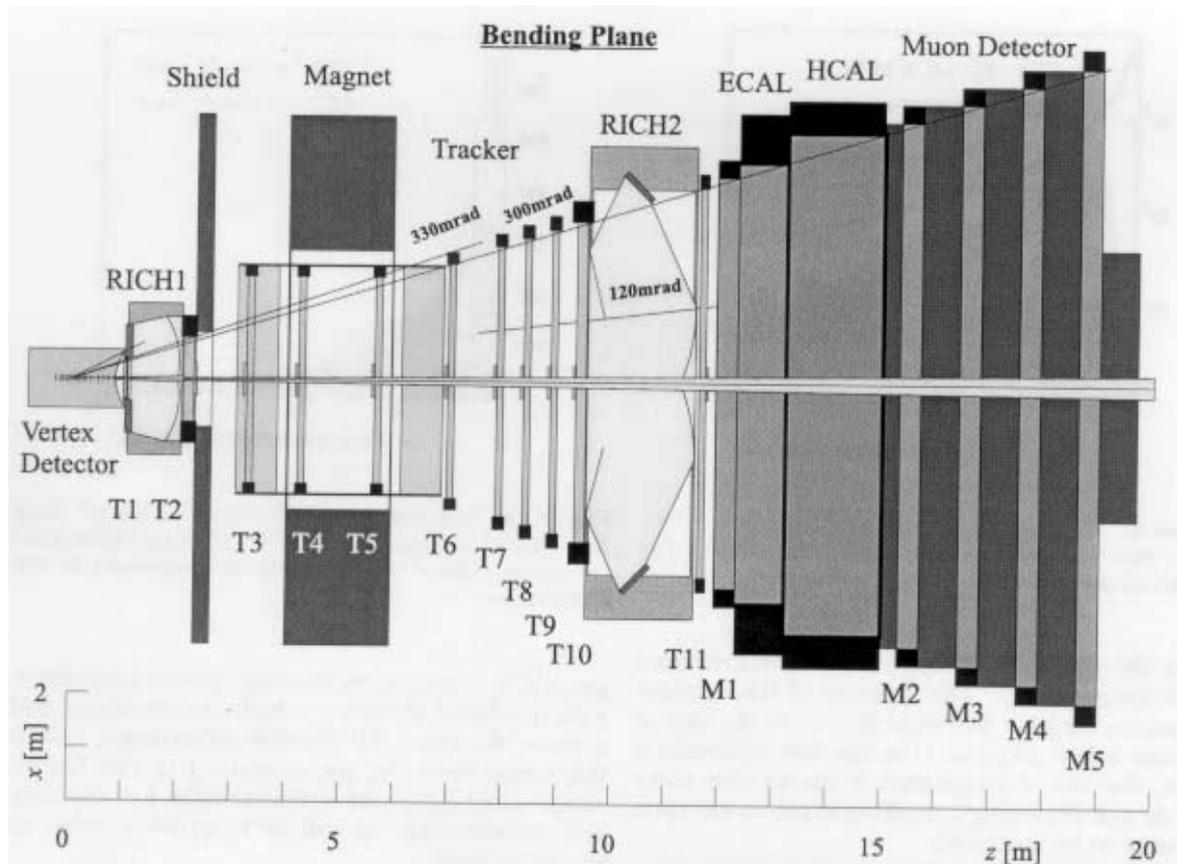


Figure 6. The LHCb detector seen from above (cut in the bending plane).

II.4.1 Vertex Detector

The Vertex Detector layout consists of 17 stations between $z = -18$ cm and $+80$ cm, each containing two discs of silicon detectors, with circular (r) and radial (ϕ) strips, respectively. The discs are positioned perpendicular to the beam. Planes with radial strips alternate between a $+5^\circ$ and -5° stereo angle. Tracking coverage extends radially from 1 to 6 cm, and provides at least three space points on tracks with polar angles down to 15 mrad. A silicon thickness of $150 \mu\text{m}$ is used to reduce multiple scattering. The front-end electronics up to the Level-0 buffers are mounted at approximately 7 cm from the beam axis. Analogue information from 220,000 amplifiers is transmitted on 7000 twisted-pair cables through the vacuum tank to the readout electronics at a distance of about 10 m from the detector. At nominal luminosity, the innermost part of the detector is expected to give acceptable performance (i.e. a minimum signal-to-noise ratio of 8) for at least one year, when operated at 5°C .

The readout pitch varies from 40 to $80 \mu\text{m}$ for the r -strips and from 40 to $104 \mu\text{m}$ for the ϕ -strips. This provides hit resolutions between 6 and $10 \mu\text{m}$ for double-channel clusters and between 9 and $18 \mu\text{m}$ for single-channel hits. A resolution of approximately $40 \mu\text{m}$ on

the impact parameter of high-momentum tracks is obtained. Each detector element covers 61° in azimuth, the innermost r strips being read out in two segments. This ensures that the strip occupancy stays below 0.8%.

Material from guard rings on the detectors, from an RF-shield and a wake-field suppressor reaches to within 4 mm of the beam axis. This is acceptable during normal LHC operation for physics data taking. However, the material has to be retracted by 3 cm during injection. This is possible because the complete assembly is constructed in two halves, which can be moved apart vertically. With a longitudinal offset of the two halves by 2 cm, an overlap of the sensitive areas can be achieved. The design foresees a secondary vacuum for the detectors inside a $100 \mu\text{m}$ aluminium RF-shield.

II.4.2 Pile-up Veto Counter

Two dedicated planes of silicon detectors act as a pile-up veto which is available in time for the Level-0 trigger. These detectors have circular strips and are placed upstream of the main Vertex Detector, opposite to the spectrometer arm. Each plane is subdivided into 60° sectors containing 300 strips with a pitch between 120 and $240 \mu\text{m}$. The on-detector electronics include a discriminator, and the digital information of the 3600 channels is fed through the vacuum tank to nearby

vertex-finding processors. Simulations show that a primary vertex is reconstructed with a resolution of 1 mm along the beam direction. This rejects 80% of double interactions while retaining 95% of single interactions.

II.5 Tracking system

The tracking system, consisting of Inner and Outer Tracker, provides efficient reconstruction and precise momentum measurement of charged tracks, track directions for ring reconstruction in the RICH, and information for the Level-1 and higher-level triggers.

The system comprises 11 stations (T1–T11 in Fig. 6) between the vertex detector and the calorimeters. Precise coordinates in the bending plane are obtained from wires or strips at 0° and $\pm 5^\circ$ with respect to the vertical (y, u, v). Stations immediately up- and downstream of the RICH counters contain additional planes, providing precise measurements in the non-bending plane (wires/strips along x). The choice of technology is determined by the requirement of low occupancy. For most of the acceptance, covered by the Outer Tracker, particle fluxes are below $1.4 \times 10^5 \text{ cm}^{-2}\text{s}^{-1}$ and permit the use of honeycomb-like drift chambers. For the Inner Tracker, which lies inside a boundary of $60\text{cm} \times 40\text{cm}$, a different technology handling fluxes up to $3.5 \times 10^6 \text{ cm}^{-2}\text{s}^{-1}$ is required. Because of its reduced dimensions, station T1 contains only Inner Tracker modules.

The expected momentum resolution for the chosen design is approximately 0.3% for momenta from 5 to 200 GeV/ c , limited mainly by multiple scattering. Mass resolutions are good, e.g. $17 \text{ MeV}/c^2$ for $B_d^0 \rightarrow \pi^+\pi^-$.

II.5.1 Outer Tracker

The detector proposed is based on honeycomb-chamber technology: a multi-cell layer with mostly 5 mm cells is built from two pre-formed conductive foils. Two staggered layers, separated by a conductive foil and sandwiched between two foam layers, make up a self-supporting module for one coordinate measurement. Such a module is $0.4\% X_0$ thick. Most stations are assembled from four modules with y, u, v, y orientations. Stations 2, 10 and 11 contain two additional x modules. The total number of electronic channels is 110,000.

With a fast CF_4 -based drift gas, signal latency will be two bunch-crossing intervals. Single-cell resolution is expected to be $< 200 \mu\text{m}$. The inner boundary of the Outer Tracker stations is determined by demanding a maximum cell occupancy of less than 10%.

II.5.2 Inner Tracker

For coverage inside $|x| = 30 \text{ cm}$ and $|y| = 20\text{--}30 \text{ cm}$, chambers using a Microstrip Gas Chamber (MSGC) with a Gaseous Electron Multiplier (GEM) are taken for the cost evaluation. Until all problems with this technology are solved, LHCb is keeping silicon strip detectors as a fall-back solution, since they are estimated to be more expensive. Microstrip Cathode

Chambers, which might turn out to be cheaper than the MSGC/GEM solution, are also under study.

The addition of a GEM to the MSGC reduces drastically the danger of sparking in the MSGC, as the gas amplification is divided into two stages. Detector modules of linear dimensions up to 30 cm are proposed, with a 2 mm drift gap, a $50 \mu\text{m}$ GEM foil and a 2 mm amplification gap. With a gain of 20 in the GEM, safe operation of the MSGC at a gain of 200 can be obtained. Glass substrates with diamond-like coating are used, with a readout pitch of $220 \mu\text{m}$. In a 1.1 T field, the strip hit multiplicity is around 1.5, and the hit resolution is better than $65 \mu\text{m}$. The occupancy is on average below 1%, and at most 4% at a minimum distance of 36 mm from the beams. A total sensitive area of 14 m^2 has to be equipped with about 220,000 readout channels. For the Level-1 trigger, four channels are combined to provide a granularity of 0.88 mm.

II.6 RICH detectors

The RICH system has the task of identifying charged particles over the momentum range 1–150 GeV/ c , within an angular acceptance of 10–330 mrad. Particle identification is crucial to reduce background in selected final states and to provide an efficient kaon tag. To achieve this goal, the system consists of an upstream detector (RICH 1) with both a silica aerogel and a C_4F_{10} gas radiator, and a downstream detector (RICH 2) with a CF_4 gas radiator. RICH 1 has 25–330 mrad acceptance in both x and y projections and is situated upstream of the magnet to detect low-momentum tracks that are swept out of the acceptance. Although RICH 2 has a reduced acceptance, 10–120 mrad in x and 10–100 mrad in y , it catches a large fraction of the high momentum tracks, e.g. 90% of reconstructed pions from $B_d^0 \rightarrow \pi^+\pi^-$ with $p > 70 \text{ GeV}/c$ (i.e. beyond the limit for π -K separation in RICH 1).

The pion thresholds are 0.6, 2.6, and 4.4 GeV/ c for the aerogel, C_4F_{10} and CF_4 radiators, respectively. Kaon thresholds are 2.0, 9.3 and 15.6 GeV/ c . The Cherenkov photons are focussed by mirrors to detector planes positioned outside the spectrometer acceptance. Hybrid photodiodes (HPD's) with pixel or pad readout are foreseen. A total image surface of about 2.9 m^2 and a detector granularity of $2.5 \text{ mm} \times 2.5 \text{ mm}$ is required. A major effort goes into the development of HPD's, which must provide a large fraction of active area at an acceptable cost. Multianode photomultipliers are considered as a back-up solution. Assuming an active-area fraction of 73%, approximately 340,000 electronic channels are required.

RICH 1 has a 5 cm-thick aerogel radiator and a 95 cm-long C_4F_{10} radiator. The expected numbers of detected photoelectrons are 15 and 55, respectively, for tracks with $\beta = 1$. Preliminary prototype results are in agreement with expectations, including Cherenkov-angle resolution and photon yields.

RICH 2 is filled with a CF_4 radiator with an approximate length of 180 cm. To shorten the overall length of the detector, the image from the spherical mirror is reflected by a second flat mirror to the detector planes. Approximately 30 detected photoelectrons are expected for tracks with $\beta = 1$.

Using a maximum likelihood analysis, all reconstructed tracks from simulated b-events are identified with efficiencies and purities above 90%. A 3σ separation of pions from kaons is achieved over the momentum range 1–150 GeV/c.

II.7 Calorimeters

The main purpose of the calorimeters is to provide identification of electrons and hadrons for trigger and offline analysis, with measurements of position and energy. The required Level-0 trigger selectivity demands a longitudinal segmentation of the electromagnetic calorimeter (ECAL). The structure consists of a single-layer preshower detector followed by a Shashlik ECAL. A scintillating-tile geometry is used for the hadron calorimeter (HCAL).

Acceptance and lateral detector segmentation of the three subsystems are geometrically matched to facilitate trigger formation. Polar acceptance starts at 30 mrad as a compromise between performance, cost and radiation dose. The outer transverse dimensions are matched projectively to the tracker acceptance. Front-end electronics can be largely unified due to the choice of similar scintillator and wavelength-shifting (WLS) fibre technology. It is expected that signal collection times less than 25 ns can be achieved.

II.7.1 Preshower detector

The cells of the Preshower detector are made up from 14 mm-thick lead plates followed by square scintillators, 10 mm thick. Transverse dimensions of 4, 8 or 16 cm correspond to the segmentation of the ECAL. Wavelength-shifting fibres, tightly coupled to the scintillators, will be read out by multi-anode phototubes, or possibly avalanche photodiodes.

II.7.2 Electromagnetic calorimeter

A fine transverse segmentation is necessary for efficient π^0 reconstruction, and for discrimination between electrons and charged hadrons with overlapping photons. To limit cost and complexity, the lateral segmentation is adjusted radially in three steps, in such a way that hit occupancy does not exceed 5%. A modest energy resolution with a 10% statistical and 1.5% constant term is sufficient for a hadron rejection factor of about 100, as well as for π^0 reconstruction.

The ECAL submodule is constructed from 70 layers, consisting of 2 mm-thick lead plates and 4 mm-thick polystyrene-based scintillator plates. The length corresponds to $25 X_0$. Light is collected by 1 mm-diameter WLS fibres traversing the whole length of the stack. The transverse granularity is varied by merging the readout of either 1, 4 or 16 identical submodules.

Tests with injection-moulded polystyrene-based scintillators and Y-11 fibres indicate that operation should be possible for more than 10 years, assuming a maximum annual radiation dose of 0.4 Mrad.

II.7.3 Hadron calorimeter

The HCAL module is constructed from scintillator tiles embedded in an iron structure. The tiles are placed parallel to the beam direction, in a staggered arrangement. The overall transverse dimensions of the sensitive volume are 9.0 m \times 7.0 m and the depth of the iron is 1.5 m. This provides $7.3 \lambda_I$.

The lateral segmentation of the submodules has a 4:1 correspondence between the ECAL and HCAL. Hence, cells have square cross-sections with sides of 8, 16, and 32 cm.

Approximately 50 photoelectrons/GeV are expected to be detected with phototubes, resulting in an energy resolution with an 80% statistical and a 5% constant term. Tests with a full-scale prototype using 16 cm \times 16 cm cells have exceeded the expected performance.

II.8 Muon detector

The Muon Detector provides muon identification and Level-0 trigger information. It consists of four stations M2–M5 embedded in an iron filter and a special station M1 in front of the calorimeter. All stations have pad readout to achieve fast trigger response. The sizes of the logical pads (used for triggering and reconstruction) vary from 1 cm \times 2 cm to 8 cm \times 16 cm. To reduce capacitive noise, the largest logical pads have to be formed by combining the information from four physical pads, each connected to a separate amplifier. Multigap Resistive Plate Chambers (MRPC's) are proposed for most of the coverage of M2–M5, where particle fluxes are below $5 \times 10^3 \text{ cm}^{-2} \text{ s}^{-1}$. Station M1 and the inner regions of stations M2–M5 experience the highest fluxes and are therefore constructed from Cathode Pad Chambers (CPC's). These chambers extend down to 25 mrad in x and 15 mrad in y . The complete Muon Detector has 45,000 readout channels formed from 230,000 physical pads.

An MRPC has four 0.8 mm gaps, read out by a single pad plane and sandwiched between two honeycomb plates. To achieve high-rate capability, the chambers have to be operated in the avalanche mode with the lowest possible gain. A multigap structure is proposed to improve response time and to decrease the probability of streamer formation. Peaking times of 10 ns are expected. Two MRPC layers per station are required for good efficiency. Occupancy is below 10% everywhere.

A CPC module contains two multiwire proportional chambers with vertical wires placed asymmetrically and forming 1 and 4 mm gaps, each with cathode-pad readout from the 1 mm gap. The chambers are sandwiched

between, and separated by, honeycomb plates. A muon station is formed by two such modules.

II.9 Front-end electronics

The subdetectors will use a common architecture for the front-end electronics, which has to accommodate the specific trigger requirements of LHCb, making maximum use of existing components. All analogue and digital signals arriving at 40 MHz will be stored in Level-0 pipelined buffers, 128 cells deep, to await the Level-0 trigger decision taken after a fixed delay of $3.2 \mu\text{s}$. Events accepted at an average rate of 1 MHz are transmitted to short derandomising buffers to avoid overflow due to limited output speed. The data are then multiplexed and digitised, if they were still analogue, and sent to Level-1 buffers, 256 events deep, to allow up to $256 \mu\text{s}$ for the next trigger selection. The average rate of events accepted by Level-1 is 40 kHz. Accepted events pass zero suppression and data formatting, are multiplexed and sent via the “front-end links” to the data acquisition system, located approximately 60 m from the detector.

The front-end electronics mounted inside the detector must be radiation hard or tolerant, the dose integrated over 10 years amounting to 0.2 Mrad at 30 cm. Part of the electronics, probably from the Level-1 buffer onwards, will be mounted at least 4 m from the beam to permit standard components to be used.

III Data Handling

III.1 Trigger/DAQ architecture

Events with B mesons can be distinguished from other inelastic pp interactions by the presence of secondary vertices, and particles with high transverse momentum (p_T). This is illustrated in Fig. 7 and Fig. 8. However, events with fully-reconstructed interesting $b\bar{b}$ final states represent only a small fraction of the total $b\bar{b}$ sample, due to the small branching ratios and the limited detector acceptance. The LHCb trigger system must therefore be selective and efficient in extracting the small fraction of interesting events from the large number of $b\bar{b}$ and other pp inelastic events.

Fig. 9 shows the general architecture of the trigger and data acquisition (DAQ) system, including the main data and control flows involved in transporting the sub-detector data from the front-end electronics to the event storage. The main parameters are summarised in Table 3.

The trigger strategy is based on four levels:

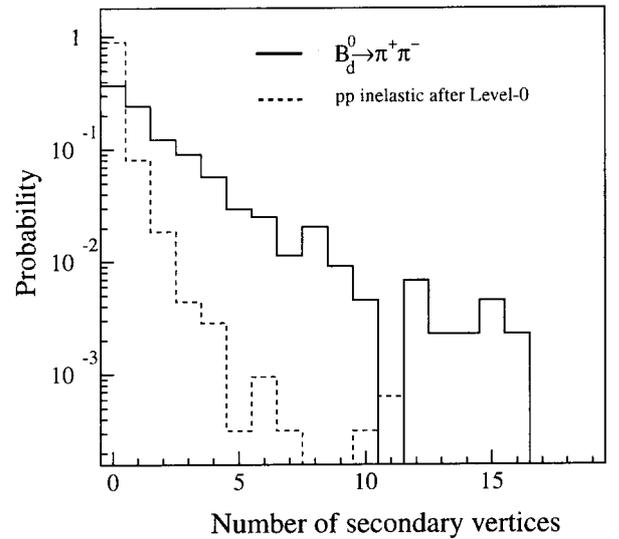


Figure 7. Number of secondary-vertex candidates reconstructed in the Level-1 trigger, for inelastic pp events (dashed) and $B_d^0 \rightarrow \pi^+ \pi^-$ events.

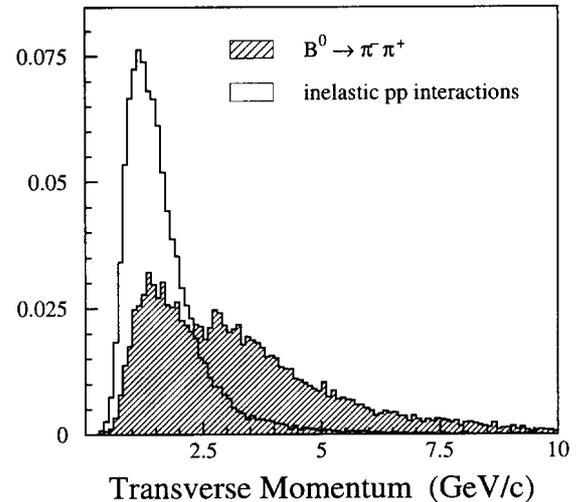


Figure 8. The p_T distributions of the charged hadrons with the highest p_T in the event, for pp inelastic events and $B_d^0 \rightarrow \pi^+ \pi^-$ events.

Level-0 comprises three high p_T triggers, operating on muons, electrons and hadrons, together with a pile-up veto that suppresses bunch-crossings with more than one pp interaction. Level-0 operates at the bunch-crossing frequency of 40 MHz, and is designed to achieve a total suppression factor of 40. This includes a factor of about 10 suppression of single inelastic pp interactions by the high p_T trigger.

The Level-0 trigger has a fixed latency of $3.2 \mu\text{s}$. This is sufficient to cover the execution time of the Level-0 algorithm ($\sim 2 \mu\text{s}$), and the time to collect trigger input data and deliver the Level-0 decision to

the front-end electronics ($\sim 1 \mu\text{s}$). The choice of latency is compatible with the standard LHC design of pipelines (128 steps) in the front-end electronics.

Level-1 has two components. The first, the vertex trigger, selects events containing one or more secondary vertices. The second, the track trigger, seeks to confirm the high p_T trigger by searching for high p_T tracks in the tracking system. Level-1 operates at the Level-0 accept rate, nominally 1 MHz, and has a suppression factor of 25.

The transfer of data from the front-end electronics to the DAQ system is initiated by a positive Level-1 decision. The average event length is ~ 100 kB and thus the DAQ must be able to cope with data rates of order 4 GB/s.

The Level-1 latency is variable and is largely determined by the execution time of the Level-1 algorithm. The processing-time distribution has an average of $120 \mu\text{s}$, assuming a CPU power of 1000 MIPS, and has long tails. The time for transferring data and broadcasting the Level-1 decision is $\sim 30 \mu\text{s}$. The maximum latency of Level-2 is set to $256 \mu\text{s}$. The buffer space required in the front-end electronics to bridge the Level-1 latency is not a serious issue, since it can be implemented as a digital FIFO and can be of arbitrary length.

Level-2 eliminates events with fake secondary vertices, by using momentum information. Fake vertices are typically caused by multiply-scattered low-momentum tracks. Level-2 operates at the Level-1 accept rate, nominally 40 kHz, and achieves a suppression factor of 8. It correlates information from the vertex detector and tracking stations.

The Level-2 latency is determined by the execution time of the algorithm. There is clear interest to make this as efficient as possible, in order to minimise the required installed CPU power, as the trigger operates at 40 kHz. An average latency of 10 ms is required, assuming a CPU power of 1000 MIPS. The buffering requirements in the readout units are estimated to be ~ 50 events for every ms of latency, and are therefore not a critical issue.

Level-3 uses full and partial reconstruction of final states to select events associated with specific b-hadron decay modes. It makes use of information from all detector components, including the RICH. The suppression factor of Level-3 is ~ 25 and the data recording rate is ~ 200 Hz.

The Level-3 latency is estimated from existing reconstruction algorithms to be ~ 200 ms. At an input

rate of ~ 5 kHz, this gives a deployed CPU power requirement of 10^6 MIPS. Again, buffer capacity is not a critical issue, as each CPU receives only ~ 30 Hz of events.

Table 3. Trigger/DAQ parameters.

Parameter	Value
Readout channels	950,000
Average event size	100 kB
Average Level-0 accept rate	1 MHz
Level-0 latency (fixed)	$3.2 \mu\text{s}$
Average Level-1 accept rate	40 kHz
Level-1 latency (variable)	$< 256 \mu\text{s}$
Front-end Links	165
Readout Units (RUs)	~ 120
Event Building bandwidth	2-4 GB/s
Sub-farm Controllers (SFCs)	~ 120
Average Level-2 processing time (on a 1000 MIPS CPU)	10 ms
Average Level-2 accept rate	5 kHz
Average Level-3 processing time (on a 1000 MIPS CPU)	200 ms
Average Level-3 accept rate	200 Hz
Data storage rate	20 MB/s

Trigger system

III.2.1 Level-0 calorimeter triggers

The electron trigger requires isolated showers in the ECAL, together with correlated hits in the pad chamber (M1) in front of the calorimeter, and energy deposition in the Preshower. Simulation shows that a suppression factor of 100 for inelastic pp interactions can be achieved with an E_T threshold of 2.34 GeV, whilst maintaining an efficiency of more than 19% for $B \rightarrow e + X$ events, for which the electron hits the calorimeter. A photon trigger is implemented by requiring a higher E_T threshold and no correlated pad chamber hits. For example, a suppression factor of ~ 150 can be achieved for inelastic pp events with a cut of 4 GeV, whilst retaining 22 % of $B_d^0 \rightarrow K^* \gamma$ where the decay products are within the detector acceptance. Events with high E_T in the HCAL are accepted in order to select B decays to hadronic final states. A cut of $E_T > 2.4$ GeV retains 60% of the $B_d^0 \rightarrow \pi^+ \pi^-$ events that can be reconstructed offline, whilst suppressing the background by a factor of 17.

Three different approaches to the implementation of these triggers are under investigation.

The first approach (3D-Flow) uses programmable ASIC's assembled into planar layers, each processor being directly associated with one or more detector elements. Several layers are needed to achieve a pipeline mode of operation, where the number of layers required

depends on the latency of the algorithm. Each processor exchanges information with the 4 nearest neighbours in its layer in order to implement the cluster algorithm. Studies based on a 3×3 clustering show that the first stage of the electron and photon algorithms can be implemented with a 3D-Flow system containing 4 layers, each layer comprising 6000 processors. The

hadron algorithm requires 4 layers of 1500 processors. The number of additional processors required to collect the results and apply cuts is small by comparison. Simulation studies show that the total execution time of the Level-0 algorithms is less than $1.5 \mu\text{s}$, which falls within the allocated time.

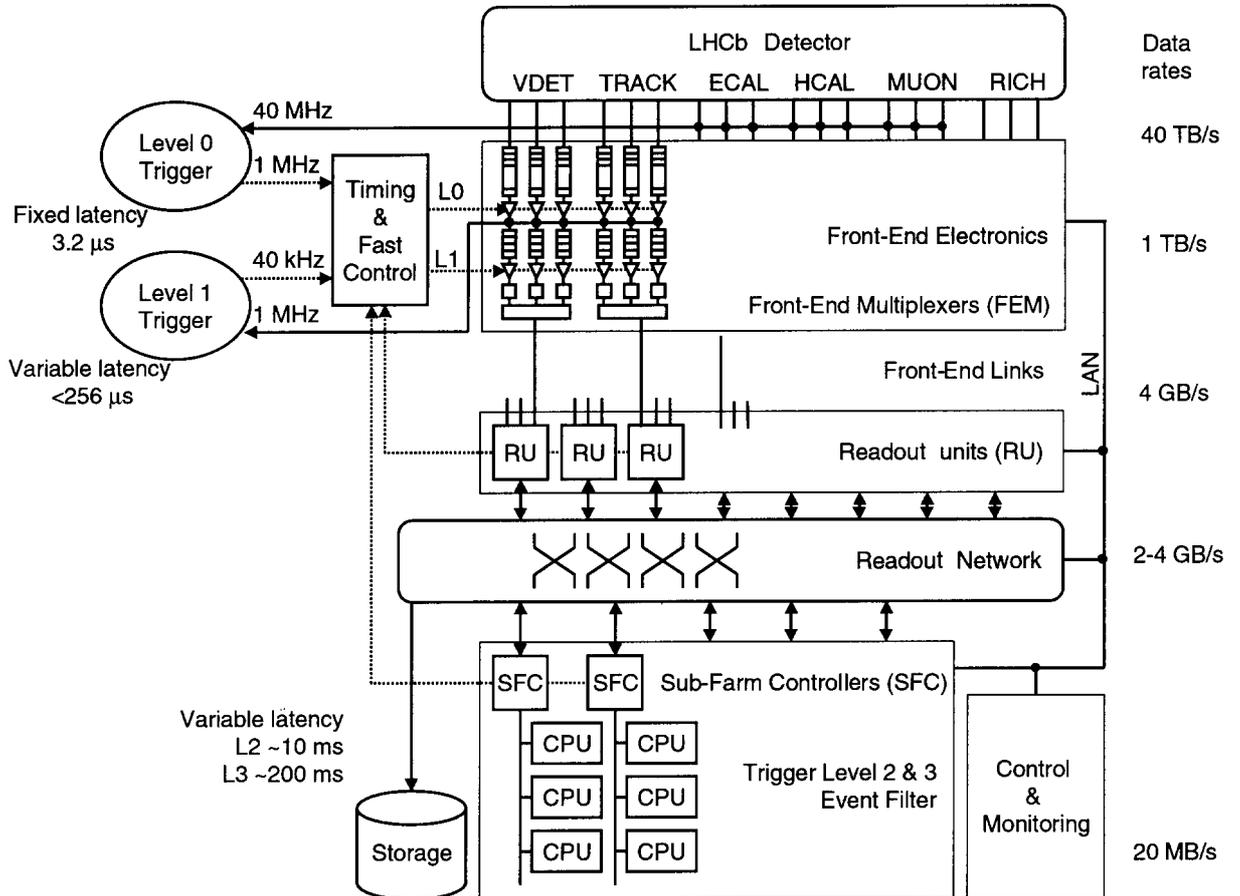


Figure 9. General architecture of the trigger and DAQ system.

The second implementation study, also based on 3×3 clustering, seeks to exploit the similarity of the LHCb and HERA-B electron triggers. Energy information that has been digitised and calibrated is sent from the front-end electronics into processing units, which make extensive use of look-up tables. Specialised units look for correlated hits in the pad chamber and Preshower, and apply the E_T cut.

The third approach uses an alternative algorithm, based on 2×2 clustering and searches for energy clusters in the front-end electronics. The aim is to reduce the large number of connections. The performance ob-

tained is comparable to that of the algorithm employed by the two previous implementations.

III.2.2 Level-0 muon trigger

The muon trigger must maintain good efficiency for detecting $B \rightarrow \mu + X$, whilst suppressing muons from π and K decays. The algorithm searches for tracks in all five stations of the muon detector and imposes a minimum p_T requirement. Simulations show that a p_T resolution of $230 \text{ MeV}/c$ can be achieved, which gives signal efficiencies of 30% for muons within the acceptance of the spectrometer, with a retention of $< 2\%$ for inelastic pp interactions.

Two implementation studies have been conducted for the muon trigger. The first uses the same 3D-Flow architecture as that used in the calorimeter trigger. Data from all 45,000 channels are fed to a stack of processors comprising three layers for making the track search, each layer containing 1300 processors. A second, much smaller, stack of processors is needed to apply the final steps of the algorithm.

The second approach aims to exploit the low track multiplicity in the muon chambers by already identifying muon-track candidates with a coarse granularity in the front-end electronics. Detailed hit information is sent to the trigger units only from those regions in the vicinity of the candidate muon tracks. This results in a data suppression factor of about 20. Full muon track processing is then performed on these pad hits.

III.2.3 Level-0 pile-up veto

Since primary vertices are spread along the beam axis with $\sigma_z = 5.3$ cm, multiple primary vertices, in a bunch crossing with more than one pp interaction, can be identified by reconstructing them with a modest resolution. An algorithm has been developed to exploit the fact that the z coordinate of the primary vertex can be determined from only the radial position of the emerging tracks. Simulation studies show that a vertex resolution of 1 mm along z can be obtained using the Pile-up Veto Counter. The performance of this veto is sufficient to give an 80% rejection of bunch-crossings with two or more pp interactions, whilst retaining 95% of single pp interactions. The vertex-based veto can be complemented by a calorimeter total energy measurement.

III.2.4 Level-0 decision unit

The results from all Level-0 trigger components are fed to a single unit where they are combined and a final decision is made. The acceptable combinations can be programmed to give flexibility and allow the trigger to be adapted to running conditions.

III.2.5 Level-1 vertex trigger

The r - ϕ geometry of the Vertex Detector is chosen to facilitate the implementation of the vertex trigger algorithm. After tracks are found in the r - z projection within the 61 degree sectors ("2D tracks"), they are combined with other 2D tracks in the opposite sector to build two-track vertices. The position of the primary vertex is determined by combining all two-track vertices, giving a resolution of 80 μm in z and 20 μm in x and y . The impact parameter with respect to the primary vertex is then calculated for each 2D track, such that tracks coming from the primary vertex can be eliminated. The ϕ -cluster information is subsequently

added for all remaining 2D tracks and the track is reconstructed in three dimensions. Finally, the algorithm finds secondary vertices that are significantly separated from the primary vertex. Simulation studies show that a factor of 25 suppression of inelastic pp interactions can be achieved with a 45% efficiency for $B_d^0 \rightarrow \pi^+\pi^-$ events that can be reconstructed offline.

The following implementation is proposed for the vertex trigger. All Vertex Detector data for a specific event, at a reduced resolution, are collected and dispatched to a single processor in a processor farm, where the trigger algorithm is executed entirely in software. Each event contains about 1000 clusters on average, which results in a data collection rate of 2 GB/s. Processing requirements are estimated from benchmarks made by running the algorithm on a 180 MHz Pentium CPU. Assuming the Level-1 trigger operates at 1 MHz, ~ 120 processors of 1000 MIPS CPU capacity would be needed to implement the vertex-trigger algorithm. The maximum latency of the algorithm would be of the order of 300 μs on a 1000 MIPS processor, with the current version of the software. Optimisation of the code has still to be done and the goal of 256 μs should be easily achievable. A solution for the readout network has been devised and elements of the computer farm can be assembled from commercial components.

III.2.6 Level-1 track trigger

The tracking detectors are used to reject fake high- p_T tracks that passed Level-0, which originate from background due to secondary interactions in detector material, particle decays and overlapping showers. The tracking algorithm uses high- E_T clusters and tracks found by Level-0 as seeds, thereby eliminating the need for full pattern recognition. Tracks are identified by a series of backward extrapolations, starting from the seed. After each extrapolation to the next tracking station, the track parameters are recalculated using a simplified Kalman filter technique. It is sufficient to stop the extrapolation at station T6 where the track momentum is already measured with a resolution of approximately 3%. Simulation shows that the track trigger reduces the number of inelastic collisions by as much as a factor of 10 whilst retaining between 50% and 80% of signal events, depending on the B decay mode.

Each tracking station will be equipped with a processing unit for collecting data and executing the algorithm. Application-specific hardware would be used for the pre-processing steps and, since the maximum latency is 256 μsec , commercial Digital Signal Processors can be used for executing the algorithm.

III.2.7 Level-1 decision unit

The results of the vertex and track triggers are sent

to the Level-1 decision unit. This unit also receives information from the Level-0 decision unit. Optimal combinations can then be made based on all information available before the final Level-1 decision is formed. This unit will be implemented using look-up tables, so that the trigger can be adapted to running conditions and physics requirements.

III.2.8 Level-2 trigger

Level-2 is designed to match the vertex information provided by the Vertex Detector with the momentum information provided by the tracking detectors. Firstly, all tracks are reconstructed in 3D in the Vertex Detector, and an initial estimate of the momentum is taken from the polar angle of the track. The primary vertex position is found with a resolution of $9 \mu\text{m}$ in x and y and $38 \mu\text{m}$ in z . Tracks in the forward hemisphere having an impact parameter larger than twice the resolution (~ 6 tracks per event) are followed to station T5. The impact parameter and its error are then recalculated using the measured momenta. Events are kept if there are at least three tracks with an impact parameter greater than 3 times the resolution. Background from K_S and Λ decays is reduced by requiring that the impact parameter is < 2 mm. Simulations show that Level-2 loses very few B-meson events, while it gives good rejection of light quark events. Another positive feature is that it also suppresses $b\bar{b}$ events where the b-hadron decay products are only partially contained inside the detector acceptance.

The Level-2 algorithms are executed on a farm of general-purpose processors. The goal of a Level-2 latency of ~ 10 ms/event should be easily achievable, and the total CPU requirement for Level-2 is 4×10^5 MIPS.

III.2.9 Level-3 trigger

A set of independent algorithms is used to select exclusive b-hadron decays according to their topology. The filter algorithms are based on reconstruction and analysis codes, but since they are applied in real-time, final calibration and alignment constants are not necessarily available and cuts are deliberately loosened. Fitted tracks are subjected to a precise vertex reconstruction, and, for those tracks having a large impact parameter, a fast RICH pattern-recognition algorithm is executed to search for kaons. Finally all information is used to make a flavour tag and the invariant masses are calculated for the b-hadron decay of interest. It has been found that a rejection factor of 25 can be achieved. Benchmarks of the Level-3 algorithms indicate a total CPU requirement of 10^6 MIPS.

III.3 Data acquisition system

The role of the DAQ system is to read zero-suppressed data from the front-end electronics, to assemble complete events and to provide sufficient CPU power for execution of the Level-2 and Level-3 algorithms. The flow of data through the DAQ system is being studied using simulation data. The input rates are determined by the average event size, which is ~ 100 kB, and the Level-1 accept rate, nominally 40 kHz. This results in a total data-acquisition rate of 4 GB/s. The output rate is governed by the nominal Level-3 accept rate (~ 200 Hz), leading to a data storage requirement of ~ 20 MB/s.

The current design contains the following functional components. Readout Units (RU's) receive data from one or more front-end links and assemble the fragments belonging to each event. Full event building is achieved by having all RU's dispatch their data into a readout network such that fragments belonging to the same event arrive at the same destination. Complete events are assembled at the destination by a unit called the Sub-Farm Controller. This unit also has the role of allocating each event to one of the free CPU's it manages, and the Level-2 and Level-3 algorithms are executed on this CPU. Accepted events are written to storage devices.

Two protocols are being studied for managing the flow of data and for ensuring events are assembled correctly. In the first of these (full readout), data are immediately routed through the readout network to the destination as soon as they appear at the RU. Thus complete events are immediately made available, and there is complete flexibility in defining and applying the high-level trigger algorithms. In the second approach (phased readout), data from subsets of the RU's are sent in several phases corresponding to the requirements of the high-level triggers. In this approach use is made of the fact that high-level triggering is performed in a sequence of steps, each step requiring a subset of the data. Thus the bandwidth required of the readout network can be reduced (by $\sim 50\%$), since data need not be transmitted from all detectors if the event can be rejected at an early stage. The full readout protocol is superior from the point of view of simplicity of protocol, and the flexibility in application of the high-level trigger algorithms. However, the future availability of high-bandwidth network technologies will determine whether this solution can be realised at an affordable cost.

The clock, Level-0 and Level-1 decisions must all be transmitted to the front-end electronics. The Trigger, Timing and Control distribution system (TTC), developed in the context of the RD12 project and used by the other LHC experiments, can also be used for LHCb, despite the added requirement of distributing the Level-1 decision.

The Detector Control System (DCS) will be used to monitor and control the operational state of the LHCb detector, and to acquire slowly-changing data from the detector to keep a permanent record of environmental parameters. The DCS will be fully integrated with the DAQ system.

III.4 Computing

A data-flow analysis has been made of the computing system, in order to identify the key tasks and data flows. Estimates of computing requirements are made by analysing existing LHCb simulation, reconstruction and analysis programs, and, in cases where they do not yet exist, by extrapolations from similar codes developed by other experiments (such as HERA-B). Results show that the total CPU requirements needed for simulation, for Level-2 and Level-3 triggers, and for reconstruction and analysis of simulated and real data, amounts to $\sim 5 \times 10^6$ MIPS. The data-storage requirements for both real and simulated data amount to ~ 500 TB per year.

At present, LHCb depends almost entirely on central computing facilities for satisfying its computing needs. A modest increase in these needs ($\sim 50\%$ per year) is expected up to the end of 2000, and the use of central services, inside and outside of CERN, will continue for this period. There will then be a significant increase in activity in test beams, and investment in private CPU capacity is therefore envisaged. Investment in full scale facilities will be left as late as possible, to benefit from improvements in the price/performance ratio of computing equipment. It is considered that the computing requirements are within the limits of current technology trends, and that the equipment will be

affordable.

Attention will be given to the quality of software, particularly for the high-level trigger algorithms. The extra effort that this implies can be recovered by extracting as much use as possible out of each software component by ensuring its reuse in all application domains. This implies investment in sound engineering and management practices. The LHCb simulation program is written in FORTRAN, but will be re-engineered using Object Technologies.

References

- [1] LHCb Technical Proposal. LHCb Collaboration, CERN/LHCC 98-4.
- [2] J. Christenson *et al.*, Phys. Rev. Lett. **13**, 138 (1964).
- [3] M. Kobayashi and K. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).
- [4] A. D. Sakharov, JETP Lett. **6**, 21 (1967).
- [5] See for example, M.B. Gavela *et al.*, Modern Phys. Lett. A **9**, 795 (1994).
- [6] L. Wolfenstein, Phys. Rev. Lett. **51**, 1945 (1983).
- [7] N. Cabibbo, Phys. Rev. Lett. **10**, 531 (1963).
- [8] R.M. Barnett *et al.*, Phys. Rev. D **54**, 1 (1996), and 1997 off-year partial update for the 1998 edition (URL: <http://pdg.lbl.gov/>).
- [9] For a recent analysis, see for example P. Paganini *et al.*, DELPHI-97-137 and LAL-97-79.
- [10] For a recent review, see A.J. Buras and R. Fleischer, TUM-HEP-275/97 (1997).
- [11] For recent work, see M. Gronau and D. London, Phys. Rev. D **55**, 2845 (1997); A.I. Sanda, Z.-Z. Xing, Phys. Rev. D **56**, 6866 (1997); Y. Grossman, Y. Nir, and R. Rattazzi, SLAC-PUB-7379 (1997); R. Fleischer, 7th Int. Symp. on Heavy Flavour Physics, Santa Barbara, CERN-TH-97-241 (1997).
- [12] See for example L. Wolfenstein and Y.L. Wu, Phys. Rev. Lett. **73**, 2809 (1994).
- [13] P. Cho, M. Misiak and D. Wyler, Phys. Rev. D **54**, 3329 (1996).