
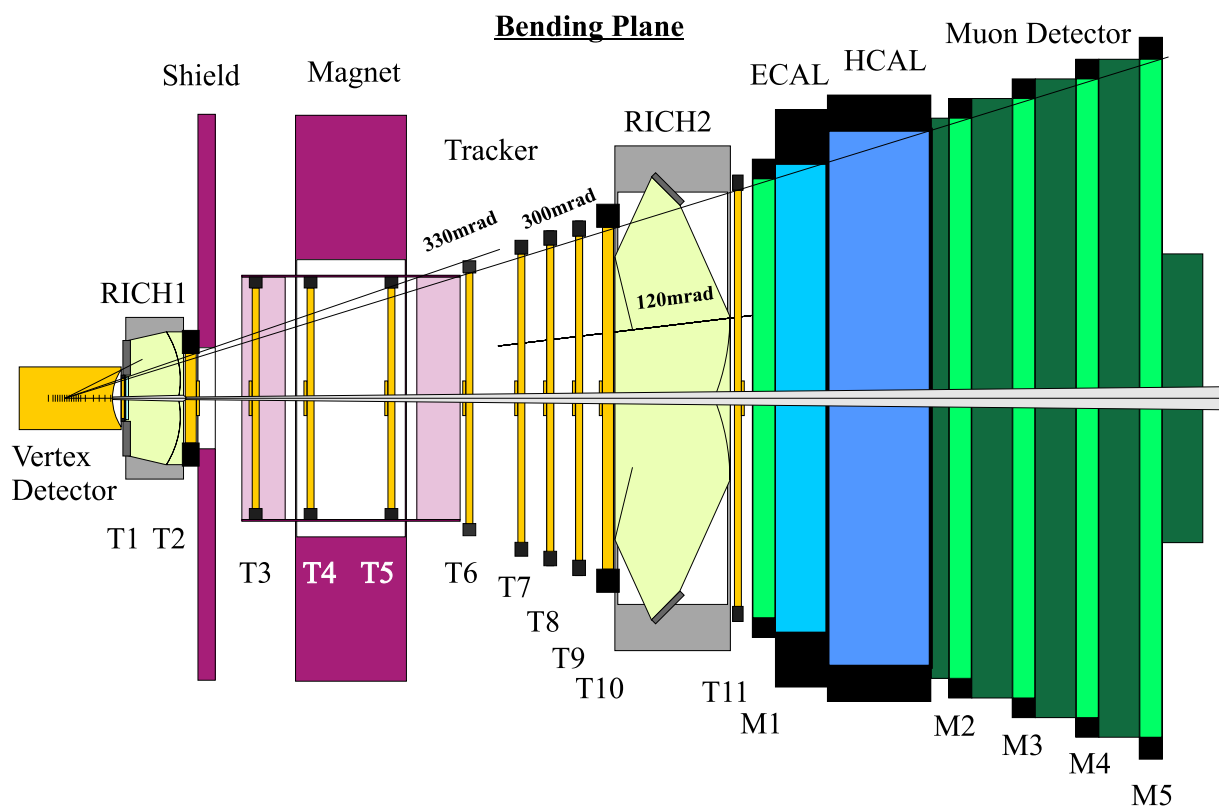


LHCb

Technical Proposal

A Large Hadron Collider Beauty Experiment for Precision Measurements of CP Violation and Rare Decays



Members of the LHCb Collaboration

The participation of Non-Member State groups is subject to the satisfactory conclusion of bilateral agreements between the funding agencies and CERN.

University of Rio de Janeiro, UFRJ, Rio de Janeiro, Brasil

S. Amato, J.R.T. de Mello, L. de Paula, M. Gandelman, J.H. Lopes, B. Marechal

University of Clermont-Ferrand II, Clermont-Ferrand, France

Z.Ajaltouni, A.Falvard, P.Perret

CPPM Marseille, Marseille, France

E. Aslanides, B. Dinkespiler *, R. Legac, C. Meessen *, M. Menouni, R. Potheau

University of Paris-Sud, LAL Orsay, Orsay, France

C.Beigbeder-Beau, D.Breton *, O.Callot, Ph.Cros *, B.D'Almagne, B.Delcourt, F.Fulda Quenzer, A.Hrisoho, A.Jacholkowska, B.Jean-Marie, J.Lefrancois, V.Tocut *, F.Vales *, I.Videau

College de France ¹, Paris, France

J. Seguinot ², T. Ypsilantis ³

¹ non-participating institute, ² emeritus, ³ retired

Humboldt University, Berlin, Germany

T.Lohse, R.Mankel

Technische Universität of Dresden, Dresden, Germany

R.Schwierz, B.Spaan

University of Freiburg, Freiburg, Germany

J.Braun, H.Fischer, J.Franz, F.H.Heinsius, K.Königsmann, J.Urban

Max-Planck-Institut für Kernphysik, Heidelberg, Germany

C. Bauer, N.Bulian, W.Fallot-Burghardt, H.P.Fuchs, W.Hofmann, K.T.Knoepfle, M.Schmelling, E.Sexauer, U. Trunk

Physikalisches Institut, University of Heidelberg, Heidelberg, Germany

P.Bock, F.Eisele, M.Feuerstack-Raible, J. von Krogh, P.Schleper, U.Straumann

University of Bologna and INFN, Bologna, Italy

A. Bertin, M. Bruschi, R. Dona, D. Galli, U. Marconi, M. Piccinini, N. Semprini-Cesari, R. Spighi, M. Villa, A. Vitale, A. Zoccoli

University of Genova and INFN, Genova, Italy

V.Gracco, A.Petrolini, G.Piana, M.Sannino

University of Milano and INFN, Milano, Italy

M.Alemi, M.Calvi, C.Matteuzzi, P.Negri, V.Pozdniakov

University of Rome, "La Sapienza" and INFN, Rome, Italy

C.Bosio, A.Frenkel, G.Martellotti, G.Penso

University of Rome, "Tor Vergata" and INFN, Rome, Italy

M.Adinolfi, G.Carboni, R.Messi, L.Paoluzi, E.Santovetti

NIKHEF, The Netherlands

G.v.Apeldoorn (3), N.v.Bakel(1,2), A.Buijs *(4), J.v.d.Brand(2), M.Doets *(1,2), R.v.d. Eijk(1), M.Ferro-Luzzi(1,2), G.Gracia-Abril(1,5), R.Hierck(1), W.Hoogland (1,3), B.Koene(1), P.Kooijman (3), M.Merk(4,6), O.Steinkamp(1), W.Ruckstuhl(1), L.Wiggers(1), H.Schuijlenburg *(1), T. Sluijk *(1), N.Zaitsev(3,7)

(1) Foundation of Fundamental Research of Matter in the Netherlands, (2) Free University Amsterdam, (3) University of Amsterdam, (4) University of Utrecht, (5) Partially supported by PN becas FPI en el extranjero SEUID PN-9603452864, (6) Supported by the Royal Netherlands Academy of Science KNAW, (7) On leave from Petersburg

IHPE Beijing, Beijing, P.R.C.

C.Jiang, Z.Ke, X.Li, F.Sun, H.Sun, Z.Zhu

University of Science and Technology of China, Hefei, P.R.C.

W. Mei, X. Y. Wu, Z.J. Yin, Y. M. Zhang

University of Nanjing, Nanjing, P.R.C.

D.Gao, Z.-X.Sha, X.Xi, N.-G.Yao, Z.-W.Zhang

University of Shandong, Shandong, P.R.C.

M.He, X.B.Ji, J.Y.Li, X.Y.Zhang, X.Zhao

Institute of Atomic Physics, Bucharest-Magurel, Romania

D.-V.Anghel, C.Coca, M.-L.Ion, C.Kusko, S.Popescu, T.Preda, A.Rosca, H.Petrascu, D.Dumitru, C.Magureanu, G.Giolu

Institute for Nuclear Research (INR), Moscow, Russia

Y.Gavrilov, E.Guschin *, L.Kravchuk, S.Laptev, V.Postoev, I.Semeniuk *

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

I.Belyaev, A.Golutvin, V.Kiritchenko, I.Korolko, G.Kostina, E.Novikov, P.Pakhlov, F.Ratnikov *, V.Rusinov *, S.Semenov, E.Tarkovski

Lebedev Physical Institute, Moscow, Russia

Yu Alexandrov, V. Baskov, L. Gorbov, B. Govorkov, V. Kim, P. Netchaeva, V. Polianski, L. Shtarkov, A. Verdi

Institute for High Energy Physics (IHEP-Serpukhov), Protvino, Russia

I.V.Ajinenko, G.V.Borissov, V.V.Brekhovskih, R.I Dzhelyadin, A.V.Dorokhov, V.I.Kochetkov, A.K.Konoplyannikov, A.K.Likhoded, V.D.Matveev, M.V.Medynski, V.F.Obratsov *, A.P.Ostankov *, V.A.Polyakov, A.V.Riazantsev, V.I.Rykalin, V.K.Semenov, M.M.Shapkin, M.M.Soldatov, V.V.Talanov, O.P.Yushchenko

Petersburg Nuclear Physics Institute, Gatchina, St. Petersburg, Russia

A.Atamanchuk, V.Borkowski, V.Golovtsov, A.Kashchuk, L.Kudin, V.Lazarev, V.Polyakov, B.Razmyslovich, N.Saguidova, E.Spiridenkov, A.Tsaregorodtsev †, A.Vorobyov, An.Vorobyov *

†presently at CERN

University of Santiago de Compostela, Santiago de Compostela, Spain

B.Adeva, A.Lopez-Aguera, P.Conde, F.Gomez, A.Pazos, M.Plo, M.J.Tobar, P.Vazquez

CERN, Geneva, Switzerland

J.T Alarcon, M. Alemi, A. Augustinus, D. Autiero, U. Becker, P. Binko, J. Buytaert, M. Campbell *, M. Cattaneo, J. Christiansen *, E. Collins, P. Colrain, O. Cooke, P. de Jong, H. Dijkstra, J.P. Droulez, J.P. Dufey, W. Flegel *, F. Formenti *, R. Forty, M. Frank, A. Go, T. Gys, F. Hahn *, S. Haider *, J. Harvey, H.J. Hilke, C. Joram, B. Jost, M. Koratzinos, D. Lacarrère, R. Lindner, P. Mato, R. McNulty, H. Müller, A. Pacheco, K. Potter *, F. Ranjard, F. Rohner, E. Rosso *, T. Ruf, J. Saborido, S. Saladino, B. Schmidt, O. Schneider, A. Schöning, A. Schopper, C. Shepherd-Themistocleous, W. Snoeys *, W. Tejessy, F. Teubert, I.R. Tomalin, O. Ullaland, E. van Herwijnen, R. Veness, G. von Holtey, P. Wertelaers *, P. Wicht, G. Wilkinson, K. Willey

University of Lausanne, Lausanne, Switzerland

P.Bartalini, A.Bay, C.Currat, O.Dormond, Y.Ermoline, R.Frei, J.-P.Hertig, P.Koppenburg, J.-R.Moser, A.Nippe, J.-P.Perroud, F.Ronga, D.Steele, M.Tareb, M.-T.Tran

Paul Scherrer Institute, Villigen, Switzerland

D.George, T.Nakada †, V.Vrankovic, J.Zichy

†presently at CERN

Institute of Physics and Techniques, Kharkiv, Ukraine

S.Maznichenko, O.Omelaenko, Yu.Ranyuk

Institute for Nuclear Research, Kiev, Ukraine

V.Aushev, V.Kiva, I.Kolomiets, Yu.Pavlenko, V.Pugatch, Yu.Vasiliev, V.Zerkin

University of Cambridge, Cambridge, U.K.

V.Gibson, C.M.Hawkes, S.G.Katvars, C.P.Ward, S.A.Wotton

University of Glasgow, Glasgow, U.K.

N.H.Brook, A.W.Halley, V.O'Shea *, I.Smith, P.Teixeira-Dias

University of Liverpool, Liverpool, U.K.

S.Biagi, T.Bowcock, P.Hayman, M.McCubbin, G.Patel

Imperial College, London, U.K

G.J.Barber *, D.Colling, A.Duane, J.Hassard *, M.J.John, R.McCarthy, J.G.McEwen, D.G.Miller, D.R.Price *, B.Simmons, D.Websdale

University of Oxford, Oxford, U.K

J.Bibby, N.Harnew, F.Harris †, T.Hessing, J.Holt, I.McArthur, J.Rademacker, N.Smale, S.Topp-Jorgensen

† presently at CERN

University of Virginia, Charlottesville, VA, USA

M. Arenton, G. Corti, S. Conetti, B. Cox, E.C. Dukes, C. Durandet, K.S. Nelson, B. Norum

Northwestern University, Evanston, IL, USA

K. Seth, A. Tomaradze

Rice University, Houston, TX, USA

D.Crosetto

* Also participating in another LHC experiment

Acknowledgments

The LHCb Collaboration is greatly indebted to all the technical and administrative staff for their important contributions to the design, testing and prototype activities. We are grateful for their dedicated work and are aware that the successful construction and commissioning of the LHCb experiment will also in future depend on their skills and commitment.

We would like to thank A. Mazzari, V. Brunner and in particular N. Grüb for their enthusiastic help in the preparation of this proposal and the related LHCb Notes.

Contents

I	Overview	1
1	Physics Motivation	3
1.1	Introduction	3
1.2	CP violation	3
1.3	LHCb performance	5
2	Detector	7
2.1	Introduction	7
2.2	General layout	7
2.3	Beam pipe	8
2.4	Magnet	8
2.5	Vertex detector system	8
2.6	Tracking system	10
2.7	RICH detectors	10
2.8	Calorimeters	11
2.9	Muon detector	12
2.10	Front-end electronics	12
3	Data Handling	13
3.1	Trigger/DAQ architecture	13
3.2	Trigger system	14
3.3	Data acquisition system	17
3.4	Computing	18
4	Project Planning	19
4.1	Organization of the Collaboration	19
4.2	Responsibilities for detector construction	19
4.3	Cost	19
4.4	Resources and schedule	21
II	Detector	23
5	Infrastructure	25
5.1	Experimental area	25
5.2	Interaction region	28
5.3	Dipole Magnet	30
6	Front-End Electronics	33
6.1	Introduction	33
6.2	Front-end model	33
6.3	Location of the electronics	35
6.4	Distribution of clock and triggers	35

7	Vertex Detector	37
7.1	Introduction	37
7.2	Geometry	37
7.3	Roman pots	39
7.4	Particle fluxes	39
7.5	Detector design	40
7.6	Front-end electronics	44
7.7	Performance	46
7.8	Future R&D	46
8	Main Tracking Detectors	49
8.1	Introduction	49
8.2	Inner tracking detectors	50
8.3	Outer tracking detector	56
8.4	Alignment	59
8.5	Performance	59
9	RICH Detectors	61
9.1	Introduction	61
9.2	Detector layout	62
9.3	Photodetectors	64
9.4	Pattern recognition	69
9.5	Prototype tests	72
9.6	Engineering design	74
10	Calorimeters	77
10.1	Introduction	77
10.2	Preshower detector	77
10.3	Electromagnetic calorimeter	79
10.4	Hadron calorimeter	83
10.5	Photodetectors and front-end electronics	85
11	Muon Detector	89
11.1	Introduction	89
11.2	Detector layout	89
11.3	Chamber technologies	91
11.4	Front-End electronics	92
11.5	Performance	93
11.6	Future R&D	94
III	Data Handling	95
12	Trigger	97
12.1	Introduction	97
12.2	Optimal luminosity	98
12.3	Level-0 Trigger	99
12.4	Level-1 Trigger	112
12.5	Level-2 Trigger	119
12.6	Level-3 Trigger	120

13 Data Acquisition	125
13.1 Introduction	125
13.2 Data traffic patterns	125
13.3 Design criteria	126
13.4 Readout protocol studies	127
13.5 Timing and trigger distribution	128
13.6 Error handling	129
13.7 Implementation studies	129
13.8 Detector Control System	132
14 Computing	135
14.1 Data-flow model	135
14.2 Computing requirements	135
14.3 Computing infrastructure	136
14.4 Software strategy	137
14.5 Project organisation	138
14.6 Project management	138
IV Physics Performance	141
15 Reconstruction	143
15.1 Simulation program	143
15.2 Flavour tag	144
15.3 $B_d^0 \rightarrow \pi^+\pi^-$	145
15.4 $B_d^0 \rightarrow \rho\pi$	147
15.5 $B_d^0 \rightarrow D^{*-}\pi^+$	148
15.6 $B_d^0 \rightarrow J/\psi K_S$	149
15.7 $B_s^0 \rightarrow D_s^- \pi^+$ and $B_s^0 \rightarrow D_s^\mp K^\pm$	151
15.8 $B_s^0 \rightarrow J/\psi \phi$	152
15.9 $B_d \rightarrow DK^*$	153
15.10 $B \rightarrow K^{*0}\gamma$	154
15.11 Summary of event yields	155
16 CP Sensitivities	159
16.1 Formalism	159
16.2 Rate normalisation	161
16.3 Systematic errors	162
16.4 Sensitivity to α	163
16.5 Sensitivity to β	165
16.6 Sensitivities to Δm_s and $\Delta\Gamma_s$	165
16.7 Sensitivity to $\gamma - 2\delta\gamma$	168
16.8 Sensitivity to γ	169
16.9 Sensitivity to $\delta\gamma$	169

Part I

Overview

Chapter 1 Physics Motivation

1.1 Introduction

CP violation was first discovered in neutral kaon decays in 1964 [1]. 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 interactions¹. 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 [2]. 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 [3]. 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 [4]. 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.

1.2 CP violation

The elements of the CKM matrix

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

¹CP violation in the strong interaction has never been detected.

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 [5]:

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

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 [6], measured to be 0.221 ± 0.002 [7] 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 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 Figure 1.1. The related unitarity conditions are given by

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

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




Figure 1.1: Two unitarity triangles in the Wolfenstein's parameterisation with an approximation valid up to $\mathcal{O}(\lambda^5)$.

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 Figure 1.1². 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 [7], 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, $|V_{td}| = 0.009 \pm 0.003$ [7]. 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 calculated³.

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 [9]:

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

³At 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 [8]. Future rare kaon decay experiments measuring $K \rightarrow \pi\nu\bar{\nu}$ will also provide information on ρ and η [9].

²The parameter A is extracted from measurements of $|V_{cb}|$ and λ .

Since only a small contribution from the new flavour-changing neutral current is expected in D^0 - \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 [10].

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 [11]. Similarly, a large effect could be seen in the energy asymmetry of the lepton pairs produced in $b \rightarrow s\ell^+\ell^-$ decays [12].

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,⁴ 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.

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

Table 1.1: 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 Reconstr.
$B_d^0 \rightarrow \pi^+\pi^- + \text{tag}$	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

⁴This notation refers to both particle and antiparticle states.

Table 1.2: 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

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 1.1 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 1.2 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 1.1 demonstrates that the LHCb experiment can study various other decay modes generated by the $b \rightarrow s\gamma$ process, such as $B_d^0 \rightarrow K^{**}\gamma^5$ 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.

References

- [1] J. Christenson et al., PRL **13** (1964) 138.
- [2] M. Kobayashi and K. Maskawa, Prog. Theor. Phys. **49** (1973) 652.
- [3] A. D. Sakharov, JETP Lett. **6** (1967) 21.
- [4] See for example, M.B. Gavela *et al.*, Modern Phys. Lett. **9A** (1994) 795.
- [5] L. Wolfenstein, PRL **51** (1983) 1945.
- [6] N. Cabibbo, Phys. Rev. Lett. **10** (1963) 531.
- [7] R.M. Barnett et al., Phys. Rev. **D54** (1996) 1, and 1997 off-year partial update for the 1998 edition (URL: <http://pdg.lbl.gov/>).
- [8] For a recent analysis, see for example P. Paganini et al., DELPHI-97-137 and LAL-97-79.
- [9] For a recent review, see A.J Buras and R. Fleischer, TUM-HEP-275/97 (1997).
- [10] For recent work, see M. Gronau and D. London, Phys. Rev. **D55** (1997) 2845; A.I. Sanda, Z.-Z. Xing, Phys.Rev. **D56** (1997) 6866; 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).
- [11] See for example L. Wolfenstein and Y.L. Wu, Phys. Rev. Lett. **73** (1994) 2809.
- [12] P. Cho, M. Misiak and D. Wyler, Phys. Rev. **D54** (1996) 3329.

⁵K** refers to orbitally-excited kaon states.

Chapter 2 Detector

2.1 Introduction

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. 2.1 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.

Figure 2.2 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.

To determine the momentum range required for the spectrometer, the $B_s^0 \rightarrow D_s^- \pi^+$ decay is stud-




Figure 2.1: Polar angles of the b- and \bar{b} -hadrons calculated by the PYTHIA event generator.




Figure 2.2: 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.

ied. The π^+ defines the high end of the momentum range, and the π^- from the D_s^- decay the low end. Figure 2.3 shows the momentum distributions for both pions, when they are within the spectrometer acceptance. Few tracks have momenta beyond 150 GeV/c.

The requirements for the vertex detector can be illustrated by the decay-length distribution of $B_d^0 \rightarrow \pi^+\pi^-$ shown in Fig. 2.4. The average decay length for the detected $B_d^0 \rightarrow \pi^+\pi^-$ is 1.0 cm.

2.2 General layout

The layout of the LHCb spectrometer is shown in Fig. 2.5. 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.

LHCb comprises a vertex detector system (in-

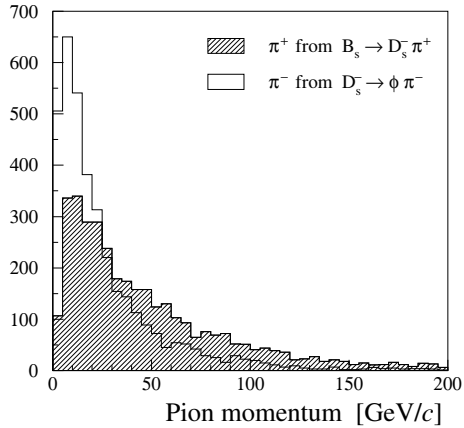


Figure 2.3: Momentum distributions for the π^+ from $B_s^0 \rightarrow D_s^- \pi^+$ decays and for the π^- from the subsequent $D_s^- \rightarrow \phi \pi^-$ decay, where the momenta are measured in the spectrometer.

cluding 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.




Figure 2.4: Decay-length distributions for $B_d^0 \rightarrow \pi^+ \pi^-$ decays in 4π , and those where both pions are detected in the spectrometer.

2.3 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.

2.4 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.

2.5 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




Figure 2.5: The LHCb detector seen from above (cut in the bending plane).

in a single bunch-crossing, by counting the number of primary vertices.

2.5.1 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.

2.5.2 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.

2.6 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. 2.5) 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 $B_d^0 \rightarrow \pi^+\pi^-$.

2.6.1 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 conduc-

tive 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%.

2.6.2 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.

2.7 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 $p > 70$ 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 m² and a detector granularity of 2.5 mm \times 2.5 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 .

2.8 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 con-

sists 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.

2.8.1 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.

2.8.2 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.

2.8.3 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 $9.0\text{ m} \times 7.0\text{ m}$ and the depth of the iron is 1.5 m . This provides $7.3 \lambda_1$.

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 $16\text{ cm} \times 16\text{ cm}$ cells have exceeded the expected performance.

2.9 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 $1\text{ cm} \times 2\text{ cm}$ to $8\text{ cm} \times 16\text{ cm}$. 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.

2.10 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 $3.2\text{ }\mu\text{s}$. 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 $256\text{ }\mu\text{s}$ 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.

Chapter 3 Data Handling

3.1 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. 3.1 and Fig. 3.2. 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.

Figure 3.3 shows the general architecture of the trigger and data acquisition (DAQ) system, including the main data and control flows involved in transporting the subdetector data from the front-end electronics to the event storage. The main parameters are summarised in Table 3.1.

The trigger strategy is based on four levels:

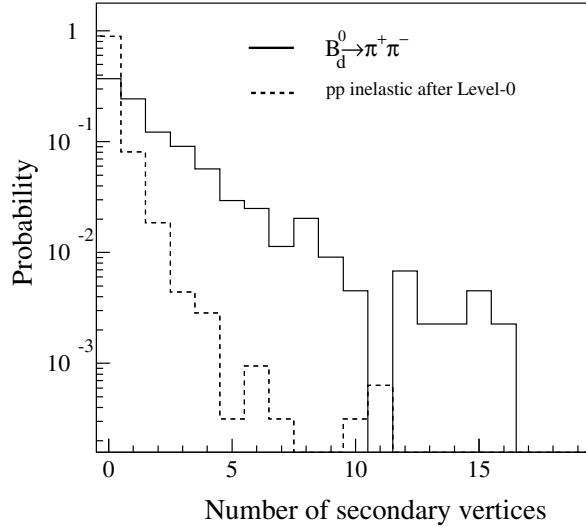


Figure 3.1: Number of secondary-vertex candidates reconstructed in the Level-1 trigger, for inelastic pp events (dashed) and $B_d^0 \rightarrow \pi^+ \pi^-$ events.

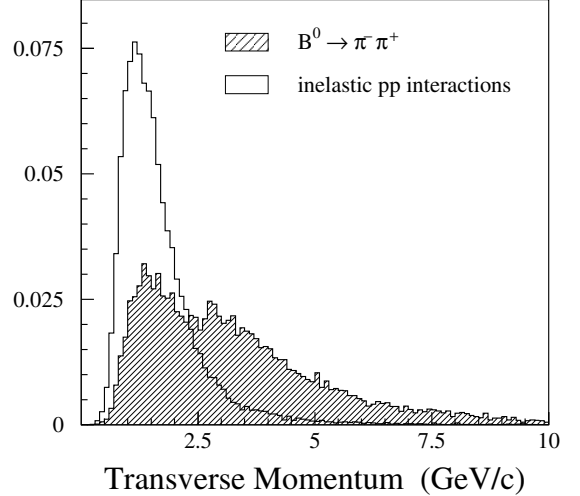


Figure 3.2: The p_T distributions of the charged hadrons with the highest p_T in the event, for pp inelastic events and $B_d^0 \rightarrow \pi^+ \pi^-$ events.

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 μs . 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 3.1: 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

3.2 Trigger system

3.2.1 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 $B_d^0 \rightarrow K^{*0} \gamma$ 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

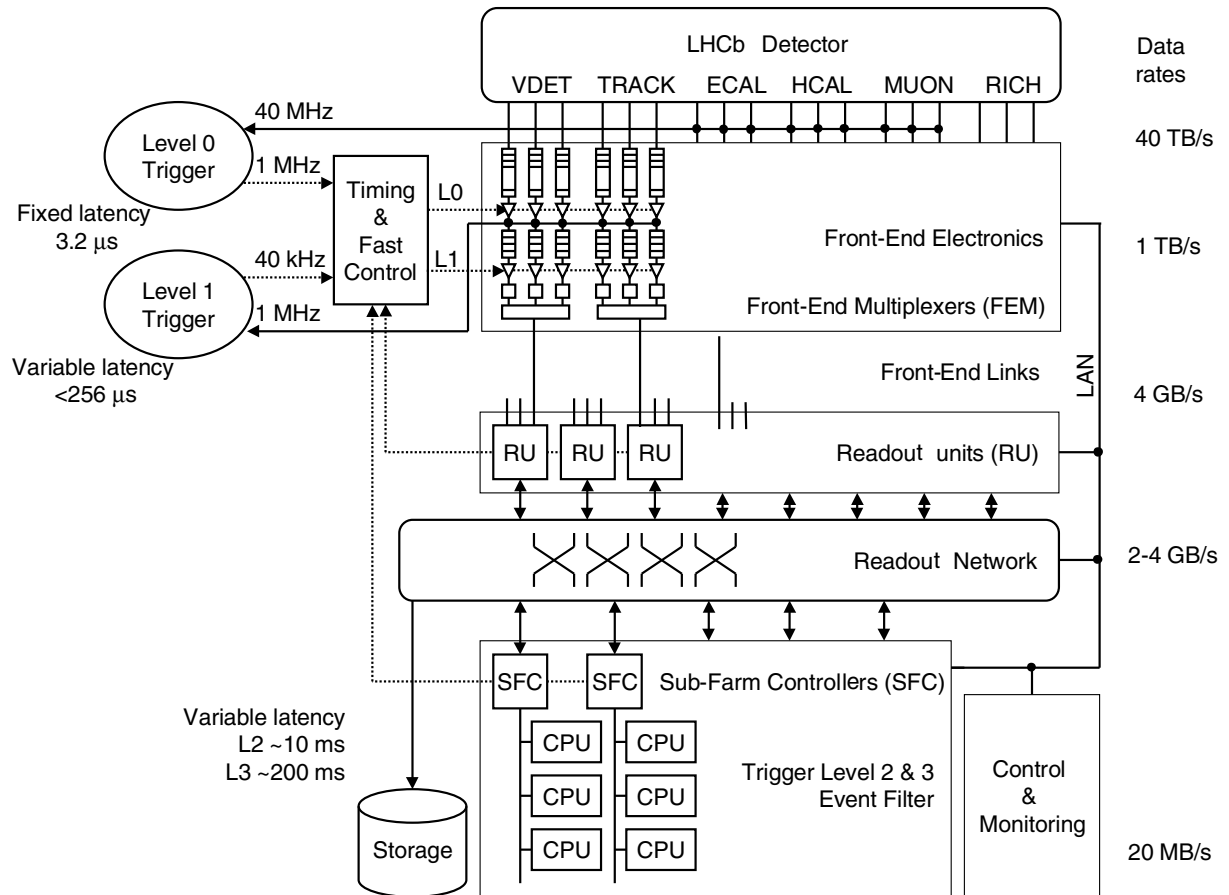


Figure 3.3: General architecture of the trigger and DAQ system.

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.

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 algo-

rithm, 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 obtained is comparable to that of the algorithm employed by the two previous implementations.

3.2.2 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.

3.2.3 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.

3.2.4 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.

3.2.5 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.

3.2.6 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.

3.2.7 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.

3.2.8 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 μm in x and y and 38 μm 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.

3.2.9 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.

3.3 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.

3.4 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.

Chapter 4 Project Planning

4.1 Organization of the Collaboration

Details of the LHCb organization are described in [1]. The **Collaboration Board** is the forum for decision and policy-making and represents the ultimate authority of LHCb. It has one member per institute (Table 4.1).

The Collaboration Board strives to reach consensus whenever possible. It elects its own Chairperson, as well as the Spokesperson (in consultation with the CERN Management) and the Technical Coordinator (in accordance with the CERN management). They are ex officio, non-voting members of the Collaboration Board. The present appointments are:

Spokesman:	T. Nakada
Technical Coordinator:	H.J. Hilke
Chairman of the Board:	N. Harnew

Major Collaboration decisions, such as for the appointment of the Spokesperson, the Collaboration Board Chairperson and the Technical Coordinator, for the admission of a new institute, and for any change to the Constitution, require a 2/3 majority, others a simple majority.

The executive responsibility rests with the Spokesperson and the Technical Coordinator. They are advised by the **Technical Board**, presently composed of the conveners of the major subsystems who have coordinated the work for the Technical Proposal ¹. The Collaboration has recently seen a substantial influx of new groups and expects others in the near future. The Collaboration Board will, therefore, decide on a new composition of the Technical Board after the approval of the experiment, once final responsibilities have been assigned. The Technical Board will then include the new project leaders of each subsystem.

The **Plenary Meeting**, chaired by the

¹W.Ruckstuhl (Tracking), R.Forty and T.Ypsilantis (RICH), G.Carboni and B.D'Almagne (Calorimetry), B.Cox (Muon), S.Conetti and U.Straumann (Trigger), J.Harvey (Computing), F.Harris (DAQ) and A.Tsaregorodtsev (Simulation).

Spokesperson, is open to everyone in the Collaboration. Physics and technical issues are discussed in the Plenary Meeting before any major decision is taken. Decisions taken by the Collaboration Board are reported at the Plenary Meeting.

4.2 Responsibilities for detector construction

Since the Collaboration is still experiencing substantial growth, the Collaboration Board has decided to postpone a detailed assignment of the responsibilities until after the approval of the experiment. The capabilities and resources of the institutes will constitute the major criteria, and are largely in line with the interests already expressed by the institutes and with their R&D activities in preparation for the Technical Proposal. These interests are summarized in Table 4.2.

4.3 Cost

The design of LHCb, described in this proposal, has evolved from the one presented in the Letter of Intent (LOI) [2]. A re-evaluation has produced some changes in the costs of the subsystems. However, the Collaboration has endeavoured to optimize the layout imposing the total cost given in the LoI as an overall ceiling. A summary of the investment cost for each component is given in Table 4.3. More details will be presented in a document which has been prepared for the Cost Review Committee (CORE).

The Collaboration intends to set up a Common Fund for about 28% of the total investment, i.e. 24 MCHF. This will include the cost of the magnet (16.5 MCHF), the computer farm and readout network (5.5 MCHF) and the long-distance cabling and piping (2.0 MCHF). Contributions to the Fund should be pro rata to the total investment contributions, with a fixed minimum amount.

Table 4.1: Present LHCb Collaboration Board members.

Institute	Contact Person	Institute	Contact Person
Bucharest	Coca C.	MPI, Heidelberg	Schmelling M.
Cambridge	Gibson V.	Nanjing	Yao N.-G.
CERN	Forty R.	NIKHEF - FOM	Ruckstuhl W.
Clermont-Ferrand	Perret P.	NIKHEF - Univ. of Amsterdam	Kooijman P.M.
CPPM Marseille	Aslanides E.	NIKHEF - Free Univ. of Amsterdam	v. d. Brand
Freiburg	Königsmann K.	NIKHEF - Univ. of Utrecht	Merk M.
Glasgow	Brook N.H.	Northwestern, Evanston	Rosen J.
Hefei	Yin Z-J.	Oxford	Harris F.
Humboldt, Berlin	Lohse T.	Paul Scherrer Institut (PSI)	Nakada T.
IHEP-Serpukhov, Protvino	Dzhelyadine R.I.	Heidelberg	Straumann U.
IHPE Beijing	Jiang C.	PNPI St. Petersburg	Vorobyov A.
INFN-Bologna	Piccinini M.	Rice University, Houston	Crosetto D.
INFN-Genova	Sannino M.	Rio De Janeiro/UFRJ	Marechal B.
INFN-Milano	Matteuzzi C.	Santiago De Compostela	Adeva B.
INFN-Roma I, La Sapienza	Martellotti G.	Shandong	He M.
INFN-Roma II, Tor Vergata	Carboni G.	Tech. Univ. Dresden	Spaan B.
INR Moscow	Kravchuk L.	Massachusetts, Amherst	Kreisler M.
ITEP Moscow	Golutvin A.	Virginia, Charlottesville	Cox B.
Kharkov	Ranyuk Y.		
Kiev	Pugatch V.	Ex-Officio Members	
LAL Orsay	D'Almagne B.		
Lausanne	Bay A.	Spokesman:	Nakada T.
Lebedev, Moscow	Govorkov B.	Technical Coordinator :	Hilke H.J.
Liverpool	Bowcock T.	Collaboration Board Chairman :	Harnew N.
Imperial College, London	Websdale D.		

Table 4.2: Interests expressed by the institutes in the construction of LHCb components.

Component	Institutes
Magnet	CERN, PSI
Vertex D./ Veto C.	CERN, Free Univ. of Amsterdam, Glasgow, Heidelberg, Imperial College, Kiev, Lausanne, Liverpool, MPI Heidelberg, NIKHEF-FOM, Roma I
Inner Tracker	Heidelberg, MPI Heidelberg, PNPI, Santiago de Compostela
Outer Tracker	Dresden, Free Univ. of Amsterdam, Freiburg, Humboldt, IHPE/Beijing, NIKHEF-FOM, Utrecht
RICH	Cambridge, CERN, Genova, Glasgow, Imperial College, Milano, Oxford
ECAL/Preshower	Bologna, Clermont-Ferrand, INR, ITEP, Lebedev, Milano, Orsay, Roma I, Roma II
HCAL	Bucharest, IHEP, Kharkov, Roma I
Muon Chambers	Beijing, CERN, Hefei, Nanjing, PNPI, Shandong, Rio, Virginia
DAQ/ Trigger/ Offline	Bologna, Cambridge, CERN, Clermont-Ferrand, Heidelberg, Lausanne, Lebedev, Marseille, NIKHEF-FOM, Orsay, Oxford, Rice, Virginia

4.4 Resources and schedule

Negotiations with the funding agencies are in very different stages of advancement: some firm commitments have been obtained, some negotiations have only just started. Preliminary estimates indicate, however, that the funding likely to be confirmed in the near term will fall short of the total cost estimate for the proposed detector. LHCb has therefore undertaken the following actions:

- The Collaboration is having discussions with potential new collaborators and will remain open in future.
- Possibilities for cost reductions not entailing serious deterioration of the performance are being investigated.
- Assuming that additional funding can be obtained in a few years' time, LHCb has started to elaborate "stretching scenarios", which permit late addition of substantial fractions of the various subsystems and are based on the latest dates for which funding is required to still guarantee completion and commissioning of the whole detector for first LHC beams. The Collaboration is confident that it will be able to convince the funding agencies by rapid construction progress (in particular on the LHCb-specific components like the vertex trigger and the RICH detectors), that additional funding for the experiment would be well placed.
- For the worst case, where no additional financing might be obtained in time to complete the detector as designed for "day 1", LHCb has started to look into "staging scenarios", which would minimize the loss in physics and still permit subsequent completion of the detector.

It is the firm intention of the Collaboration to be well prepared for efficient data taking on "day 1". LHCb has a unique chance to provide exciting physics immediately after LHC startup, since the experiment requires less than 5% of the LHC design luminosity for collecting about two orders of magnitude more B mesons than the earlier experiments, within the first year of operation.

The global schedule foresees:

1. about two more years of R&D and prototype studies,
2. the start of construction in early 2000,
3. the start of installation in autumn 2003.

Table 4.3: Estimates of investment costs.

Subsystem	LoI	Updated Cost(*)
	MCHF	MCHF
Magnet	15	16.5
Vertex D./Veto C.	5	5.1
Outer Tracker	9	8.5
Inner Tracker		4.0
RICH	9	7.0
ECAL	10	10.5
Preshower		1.9
HCAL	7	4.5
Muon Chambers	9	10.0
Data Handling	17	10.3
Trigger		3.7
Infrastructure	5	4.0
Total	86	86.0

(*) in 1998 prices

References

- [1] Organizational Structure of the LHCb Collaboration, LHCb 96-13.
- [2] LHC-B, Letter of Intent, CERN/LHCC 95-5 LHCC/I8.

Part II
Detector

Chapter 5 Infrastructure

5.1 Experimental area

5.1.1 General layout

The Interaction Point IP8, presently occupied by DELPHI, has been allocated to LHCb. To accommodate the single-arm spectrometer, the beam crossing point in the underground area UX85 has been displaced from the centre of the cavern by 11.25 m towards IP7, thus making full use of the existing hall without the need for substantial civil engineering work[1].

Figure 5.1 shows the general layout of the detector components in the cavern, viewed from the outside (the assembly area) towards the centre of the LHC ring. The latest version of the beam optics accommodates the detector length of 19.7m, measured from the IP to the end of the muon shield.

Figure 5.2 shows a top view and Fig. 5.3 an isometric view of the experimental zone. 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. This permits removal of detector modules, without breaking the LHC vacuum, and allows for in-situ bakeout of the beam pipe. The clearance currently available does not permit full retraction of the muon chambers on one side (see Fig. 5.2), and discussions are in progress on the layout of the LHC cryogenics in this area. The installation of the cryogenics, power supplies and discharge resistors for the LHCb magnet, is foreseen on an additional platform close to the dipole.

Most of the detector electronics and the on-line computer farm will be housed behind a shielding wall, at the other end of the UX hall, where they will be accessible during LHC operation. The selected event data will be sent via optical links to the data recording and control room in the surface building SX 8.

5.1.2 Safety

Radiation protection

The radiation shield presents the major modification to the UX cavern. Whilst the DELPHI detector itself provided sufficient radiation protection to allow people anywhere in the hall during LEP operation, LHC operation requires additional heavy shielding.

The radiation produced during normal operation and in a “worst-case” accident, and the protection provided by various geometries of the shielding wall have been calculated [2]. The design specifications are:

- For normal operation: a dose rate $< 10 \mu\text{Sv/h}$ everywhere behind the wall for a luminosity of $10^{33}\text{cm}^{-2}\text{s}^{-1}$, i.e. five times the nominal LHCb luminosity.
- For full beam loss: $< 50 \text{ mSv}$ at the back surface of the shield for the worst-case scenario, i.e. the loss of 4.7×10^{14} protons 5 m upstream from the centre of the cavern at the inner surface of the beam pipe.
- For beam loss at injection: $< 100 \text{ mSv/h}$ with loss of 4×10^{13} protons per 16.8 seconds.

The calculations show that the “worst-case scenario” beam loss defines the most stringent requirement. For standard concrete shielding, a thickness of 3 m for the lower 12.5 m and of 2 m for the upper part provide a good safety margin, even when a slanted cable shaft, a ‘chicane’ passage on the floor, and a straight opening for the ventilation ducts near the ceiling are added. The proposed layout of the shielding consists of two fixed pillars, erected on both sides up to crane level, with demountable central and top sections. This layout permits rapid assembly of the shield as well as the possibility to open a 13m x 11m access on the floor and/or a passage for the crane [3]. The shielding wall allows safe long-term access to the huts housing the data acquisition during LHC operation. Registration with film badge will be obligatory for any access to the




Figure 5.1: Layout of Detector components in the cavern.

cavern. During shutdowns, access to detector elements will be authorized by the Territorial Safety Officer. Detailed calculations on the build-up of induced radioactivity have yet to be carried out, but levels requiring special precautions are not expected except near the iron shield inside the muon detector.

Gases

Table 5.1 contains an inventory of the gases used in the detectors. All detectors operate close to atmospheric pressure. Due to the open geometry of the layout and the limited gas volumes and flow rates, there is no danger of asphyxiation, except possibly inside the dipole magnet and below the muon detector, where ventilation and leak detectors will be installed. The risk due to a possible helium leak in the cryogenics installation is being evaluated together with the cryogenics group and the Technical Inspection and Safety Commission (TIS).

The total amount of hydrogen-equivalent is anticipated to be below 5 kg, i.e. well below the limit of 40 kg (which defines a risk class III). With the flow rates envisaged, diffusion and convection will be sufficient to avoid flammability conditions except possibly inside the dipole and below the muon

Table 5.1: Gases for LHCb Detectors

Sub-Detectors	Gas	Vol. (m ³)	H ₂ -equiv. (kg)
RICH 1	C ₄ F ₁₀	4	0
RICH 2	CF ₄	50	0
Outer Tracker	CF ₄ /CH ₄ (80/20)%	8	0.5
Inner Tracker	Ar/DME (50/50)%	0.04	0.04
Muon MRPC	C ₂ F ₅ H/Ar/ i-C ₄ H ₁₀ (85/10/5)%	2	< 3
Muon CPC	CO ₂ /Ar/CF ₄ (60/30/10)%	< 1	0

detectors, where ventilation is foreseen. **Risk Class III is, therefore, excluded.**

Gas will be stored and mixed on the surface, as is now done for DELPHI, in the existing SG building. The available safety installations and the tested gas piping to the cavern through the PM85 shaft will be re-used.




Figure 5.2: Top view of the experimental zone.

Magnetic field

The large aperture dipole magnet produces a relatively strong stray field. Whenever the magnet is powered, a warning light will be actuated and a zone of restricted access - where standard CERN rules apply - will be fenced off.

Safety instruction

The GLIMOS will ensure that access to the cavern is permitted only for people having followed a course in safety procedure similar to the one established for the LEP experiments.

5.1.3 Services

Present estimates indicate that the installed electrical power of 1.5 MW for the detector electronics and the magnet and cooling should be adequate for LHCb. An effort will be made to re-use as much as possible from the existing cryogenics installation. However, the ventilation system will require substantial modifications, due to the increased induced radioactivity around the beam and the presence of

the shielding wall. It is foreseen to provide internal air filtering and circulation for the volume behind the shielding wall, where the electronic huts are sited, and a small overpressure with respect to the beam area.

The elevator for personnel in PZ85 and the access shaft PX84 for equipment are adequate for LHCb. The detector elements will be limited in weight and size so that handling is possible using the two existing 40 t cranes in the UX cavern.

5.1.4 Surface buildings

The surface buildings (Fig. 5.4) presently used by DELPHI are considered adequate for LHCb. These include building SX for storage space during assembly, a small workshop and the control room; SG for gas storage and mixing; the compressor area in SUH (20m × 17m) and the small Building 2889 with a meeting room and a few offices. It is important that the location of the new LHC compressor building SHM8 leaves space for the passage of large lorries behind SX 2885.




Figure 5.3: Isometric view of the experimental zone.

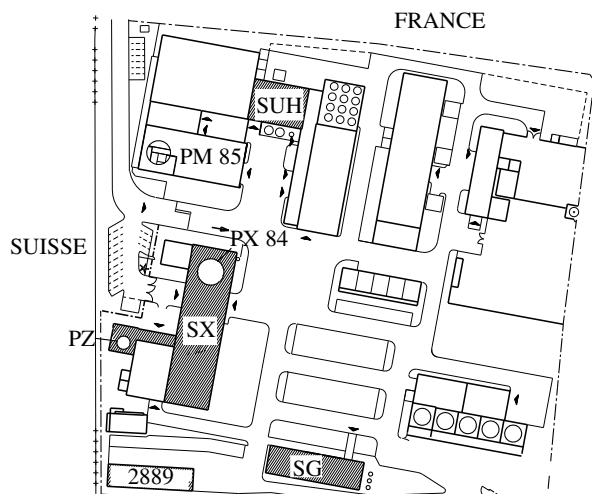


Figure 5.4: Surface buildings. The shaded buildings are to be used by LHCb.

5.1.5 Installation

Installation of the LHCb experiment will require 16 months. This includes the reception test of the superconducting dipole, which is foreseen to take

place underground. To reduce costs, the dipole will be tested at the factory, but not on the surface at CERN. It has been agreed with the accelerator group that an interruption (e.g. for injection tests) of the LHCb installation may be negotiated, as long as an integral of 16 months with full access for the experiment is guaranteed.

Since the movement of the 40 t cranes is limited, a platform has been designed for the assembly of the dipole yoke on the axis of PX84 [4]. Following assembly of the yoke and the mounting of the cryostats containing the coils, the magnet will be moved to the existing rails and along these into the beam position.

5.2 Interaction region

5.2.1 IP8 layout and optics

Circulating proton bunches in the LHC are separated by $\lambda/c = 25$ ns and counter-rotating bunches cross at points separated by a distance $\lambda/2$. The IP8 interaction point will be shifted by $3\lambda/2$ (11.25 m) in the direction of IP7, allowing the LHCb single-arm spectrometer to be accomo-

dated in the existing DELPHI hall, facing the clockwise rotating beam (LHC ring 1). To maintain the symmetry of the beam optics around this interaction point, approximately 160 m of machine elements on both sides, must be displaced. This includes the first focusing quadrupole triplet and the two separation dipoles D1 and D2. An LHC optics solution exists, which includes these changes, and provides additional space for the injection of protons into ring 2, downstream of IP8.

The luminosity is inversely proportional to β^* , the beta-value at the interaction point. The nominal LHC luminosity, $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ assumes $\beta^* = 0.5 \text{ m}$ in the two high luminosity interaction points IP1 and IP5. To provide an LHCb luminosity over the range $1 - 5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, the IP8 low-beta insertion must be tunable in the range $2.5 \text{ m} < \beta^* < 50 \text{ m}$. Fig. 5.5 shows how this LHCb luminosity range is obtained when the LHC runs at nominal luminosity and (during the early phase of operation) when it might provide only 10% of the nominal value.

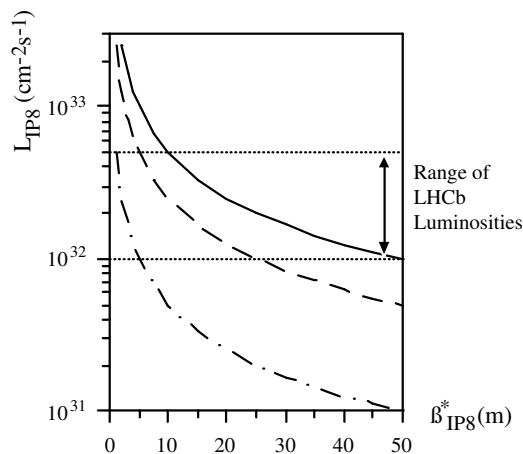


Figure 5.5: Luminosity as function of the beta-value in IP8 for three LHC luminosities: the nominal luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ (solid line), 50% (dashed line) and 10% (dash-dotted line) of the nominal luminosity.

The LHCb dipole magnet is expected to remain at its nominal field during the complete LHC cycle. It will give the circulating beams a horizontal angular kick, which must be compensated by adding magnets in the space between the vertex tank and the quadrupole triplet, upstream (towards IP7) of IP8. The spectrometer magnet and its compensation dipoles have been included into the LHC lattice. These optics also provide the required range for β -tuning in IP8 and allow reversal of the mag-

netic field. The aperture required around IP8 for the LHC beams is largest at injection energy and a minimum beam pipe diameter of 45 mm is required [5].

5.2.2 Beam pipe

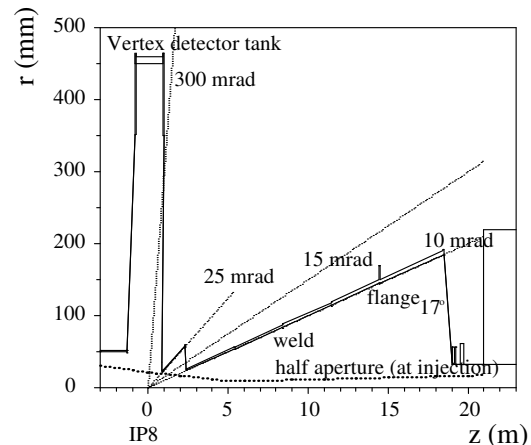


Figure 5.6: The proposed LHCb beam pipe comprises the vertex detector tank and two conical aluminium sections having 25 mrad and 10 mrad opening angle.

The proposed LHCb beam pipe is composed of three sections (see Fig. 5.6). The vertex detector tank of about 2 m total length is connected to a 25 mrad conical aluminium section of about 1.5 m length. This is followed by a 10 mrad conical Al section of about 16 m length. The different sections are interconnected by thin Al transition windows. The wall thickness of the conical Al sections increases from 1 mm to 6 mm, in proportion to the diameter. The conical sections are formed from several pieces welded together. Bellows and flanges separate the vertex detector tank from the fragile first part of the conical chamber. Pumping stations are located at both ends. To fulfil the stringent vacuum stability requirements, distributed non-evaporable getter pumping over most of the chamber length will also be needed and provision will be made for in-situ bakeout of the beam pipe.

A detailed design of the LHCb beam pipe is under discussion, taking into account particle background and vacuum requirements as well as the mechanical integration into the detector layout. Different chamber materials, in particular sections fabricated from beryllium, are being studied.




Figure 5.7: Conceptual Design of the Dipole Magnet.

5.3 Dipole Magnet

The dipole magnet has a window-frame design with a vertical gap of 3.6 m [6]. It provides a vertical magnetic field with good homogeneity across a 4.3 m horizontal aperture. The overall size is minimized by placing the magnet immediately downstream of RICH 1, where it provides the 300(h) x 250(v) mrad acceptance required for the downstream detectors. A schematic layout with the yoke opened is shown in Fig. 5.7.




Figure 5.8: Magnetic Field components as a function of z .

The choice of a superconducting magnet is determined by the high field required and the desire

to make maximum use of existing facilities from DELPHI. With the parameters given in Table 5.2, a maximum central field of 1.1 T is obtained, corresponding to a bending power along the beam axis of 4.0 Tm. Detailed 3-D field calculations have been carried out and some results are presented in Fig. 5.8.




Figure 5.9: Detail of coil winding and superconductor.

The field provides adequate momentum resolution with the proposed tracking system over the full acceptance. Fig. 5.9 shows details of the coil winding and the internally-cooled superconductor.




Figure 5.10: Support for a horizontal section of the coil.

Details of the support for a horizontal section of the coil are shown in Fig. 5.10.

It is foreseen to reverse the field direction regularly, in order to reduce systematic experimental errors. The time required for ramping down and up should be less than one hour. Throughout the LHC beam filling and acceleration, the dipole will be maintained at full field.

To reduce the stray field in the region of the Vertex Detector and RICH 1 (< 50 gauss around the photodetectors), an iron shielding plate (see Fig. 2.5) has been placed between the magnet and RICH 1.

Table 5.2: Dipole Specifications

Parameter:	Value:
Central Field	1.1 T
Conductor Size	$(14.4 \text{ mm})^2$
Number of Layers	2
Free Aperture (horizontal)	4.3 m
Free Aperture (vertical)	3.6 m
Coil Inductance	0.24 H
Ampère-turns	4 MA-turns
Operating Current	20 kA
Stored Energy	50 MJ
Cold Mass	4.0 t
Weight of Yoke	1200 t

References

- [1] D.Lacarrère, “LHCb Experimental Area”, CERN/EST-LEA/LHCB 97-11/dl.
- [2] H.Nakashima, M.Huhtinen, S.Roesler, G.R.Stevenson, “Radiation calculations for the main shielding wall for LHCb”, LHCB/98-023.
- [3] D.Lacarrère, R.Lindner, S.Roesler, G.R.Stevenson, “Main radiation shield for the LHCB Experiment at Point 8”, LHCB/98-024.
- [4] S.Karppinen, “Preliminary Study of the Mechanical Support Structure and Moving System for the LHCb Superconducting Dipole Magnet”, CERN/EST-LEA/LHCb 97-23/SK.
- [5] G. von Holtey, “Aperture requirements around interaction IP8”, LHCB/97-002 (EXPT).
- [6] D.C.George, R.K.Maix, V.Vrankovic and J.A.Zichy, “Design Study of a Large-Gap Superconducting Spectrometer Dipole”, LHCB/98-020.

Chapter 6 Front-End Electronics

6.1 Introduction

Features of the front-end electronics that are common to all LHCb detectors, and which are constrained by the requirements of the LHCb Trigger and DAQ system, are described here. Features specific to the requirements of a particular sub-detector are discussed in the appropriate chapter.

The front-end electronics receive analogue signals from the detector that are synchronised with the LHC bunch crossing at 40 MHz, and must deliver zero-suppressed and formatted data to the DAQ system over the read-out links. The detector data must be buffered in the front-end electronics for a time corresponding to the latency of the Level-0 and Level-1 triggers. Each trigger-level decision results in a reduction of the amount of data that must be transmitted downstream, and therefore after each stage, data from many buffers can be multiplexed onto the same output stream.

The impact of the LHCb trigger scheme on the buffering requirements in the front-end electronics is shown in Fig. 6.1. The Level-0 trigger makes the first event selection using signals from the calorimeters, the muon detector and the pile-up veto. It is implemented in specialised hardware with a latency of 128 clock cycles ($3.2 \mu\text{s}$). The Level-1 trigger uses tracking information from the vertex and the main tracking detectors. The implementation of the L1 algorithms uses a mixture of special purpose hardware and farms of high-performance processors. The latency of the Level-1 algorithms is variable, but limited to a maximum of $256 \mu\text{s}$. A special trigger decision unit is foreseen to deliver both Level-0 and Level-1 trigger decisions to the front-end electronics in chronological order, thereby simplifying the implementation of the L1 buffers. To ensure that the Level-1 front-end buffers do not overflow, they must be capable of storing up to 256 events accepted by Level-0.

The goal is to develop a model for the front-end electronics that fulfils the general LHCb requirements and that can be adapted to the specific needs of each sub-detector. In this way, common components can be more easily identified and developed for all appropriate sub-detectors.

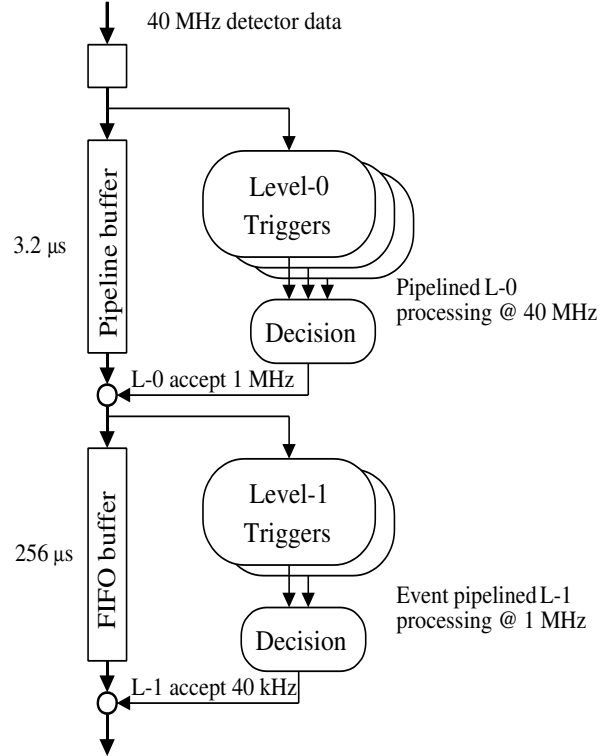


Figure 6.1: Trigger logic and front-end electronics.

6.2 Front-end model

A general view of the front-end electronics is shown in Fig. 6.2. Detector signals are amplified and shaped, when appropriate, and are stored in pipelined Level-0 buffers. Data in the Level-0 buffers can be analogue for detectors with limited dynamic range (such as the Vertex Detector, RICH, and inner tracker), digital for detectors with large dynamic range, binary or TDC data (such as the calorimeter, muon detector, and outer tracker). For detectors participating in the Level-0 decision, binary data (e.g. muon), or digital data with a reduced dynamic range data (e.g. 8 bit for calorimeters), are sent to the Level-0 trigger hardware.

Events are accepted at the output of the Level-0 pipeline at an average rate of 1 MHz. Following a Level-0 “accept” all data must be digitized, if not

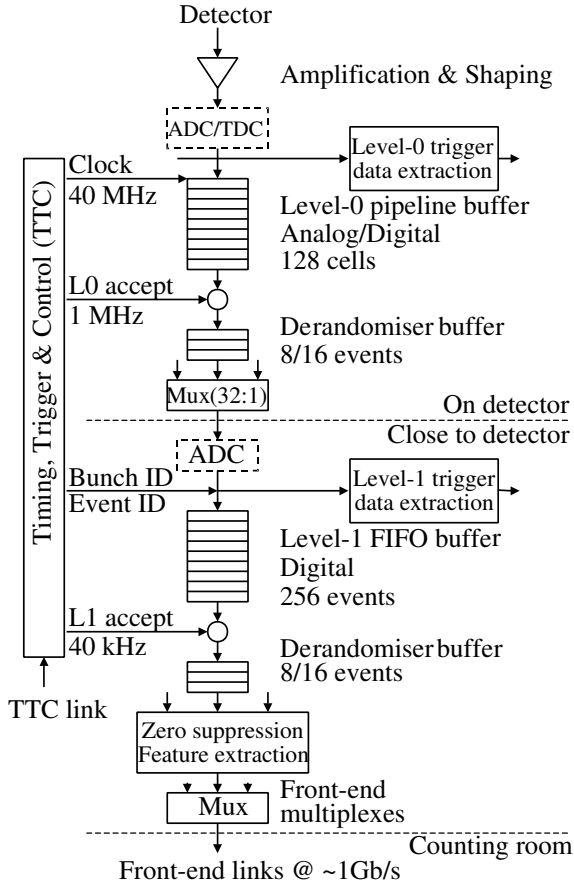


Figure 6.2: Front-end electronics architecture.

already done, before being stored in the Level-1 buffer. Level-0 “accepts” can be very close in time, in principle as close as 25 ns. A de-randomiser buffer (FIFO) is needed at the output of the pipeline and before the multiplexing stage, in order to match the data flow traffic with the speed capabilities of the output path. The optimal depth for the de-randomiser buffer is chosen to minimise the inefficiency that can arise should the buffer overflow, within constraints imposed by the cost of implementation. Table 6.1 shows the inefficiencies for different read-out time scenarios and buffer depths. The read-out time of 800 ns is of special significance, since it corresponds to the time needed to transmit multiplexed data from 32 channels using the 40 MHz clock. It can be seen that scenario C or D would both be acceptable as they result in a very small loss rate. Solution C uses the optimal readout time of 800 ns and an 8-deep de-randomiser, the depth implemented on many existing chips. The deadtime can be made negligible by increasing the depth of the de-randomiser to 16 (Solution D).

Table 6.1: Inefficiencies for different read-out scenarios.

Scenarios	Loss
A: De-randomiser depth 4, read out in 500ns	0.33%
B: De-randomiser depth 8, read out in 500ns	0.003%
C: De-randomiser depth 8, read out in 800ns	0.6%
D: De-randomiser depth 16, read out in 800ns	0.02%
E: De-randomiser depth 32, read out in 1000ns	1.5%
F: De-randomiser depth 64, read out in 1000ns	0.8%




Figure 6.3: Level-0 inefficiency as function of trigger rate.

Some detectors may require samples from more than one beam crossing to be extracted from the Level-0 pipeline for each accepted trigger. With front-end electronics designs developed by ATLAS and CMS [1], this can be implemented only by inhibiting the formation of triggers from consecutive bunch crossings. As LHCb has a high Level-0 trigger rate this would result in a relatively high deadtime of $\sim 5\%$ per MHz. This is clearly undesirable and therefore solutions will be sought to avoid this.

Fig. 6.3 shows the inefficiency as a function of trigger rate for scenario D. The contribution to inefficiency from the limited buffer depth becomes significant only for trigger rates above 1.25 MHz. The figure also shows that adding the minimum bunch spacing requirement adds a loss of 5% per MHz, and that this dominates for trigger rates below 1.25 MHz.

Buffer overflows must be prevented, in order to simplify event assembly and preserve the integrity of the data. Clearly the buffer depths and read-out speeds must be well matched across all components of the front-end electronics, such that a global

scheme can be implemented to prevent triggers appearing at too fast a rate. The central readout supervisor can be programmed not to pass more than a given number of triggers in a given time.

The vertex detector and the tracking detectors must, after digitisation, supply binary data to the Level-1 trigger processing system. During the latency of Level-1, the event data will be stored in the Level-1 buffer, which comprises a fixed number of event slots. The data storage time of $256 \mu\text{s}$ requires that the Level-1 buffer is implemented as a digital buffer. Guaranteeing that Level-1 decisions arrive in correct sequence enables the Level-1 buffers to be implemented as simple FIFO's, with separators between blocks of event data. It is currently envisaged to implement the Level-1 buffer using discrete memory devices controlled by an ASIC, and this design could be common for all sub-detectors. An advantage of using discrete memory devices is that the Level-1 latency can be relatively easily extended, if required.

There is also a need for de-randomisation of the dataflow at the output of Level-1, following the same arguments as used for the Level-0 pipeline. The read-out and the subsequent data processing after the Level-1 buffer must be able to cope with an average event rate of 40 kHz. Since the positive Level-1 decisions can happen as close as $1 \mu\text{s}$, another FIFO is needed to store temporarily the Level-1 accepted events.

Full zero-suppression and data formatting must be performed before data are dispatched over the front-end links to the data acquisition system. The front-end multiplexers merge data from many modules or cards into structured blocks of data, and these are transmitted over standardised serial links.

6.3 Location of the electronics

The optimal physical location of the front-end electronics is still being investigated for each sub-detector. The LHCb detector, being an "open" structure, enables a significant part of the electronics to be located "off-detector". The general philosophy is that the analogue front-end and the Level-0 buffer can be located inside the detector. The Level-1 buffer should be off-detector, in a nearby location where radiation is sufficiently low that standard commercial components can be used. The 1 MHz Level-0 accept rate requires a significant bandwidth at this boundary.

Radiation levels inside the LHCb detector during normal operation are estimated to be of the order of 20 krad/year at a distance of 30 cm from

the beam pipe. Because of the approximate $1/r^2$ dependence, standard components may be used for radial distances from the beam greater than 4 metres. At this distance the integrated dose over 10 years is below 1 krad. Front-end electronics used inside the detector must be radiation-hard, or at least radiation tolerant.

The first level of off-detector electronics is placed in crates a few metres away from the detector. The possibility of damage to the electronics caused by thermal neutrons still has to be evaluated. It is planned to use standard power supplies and fans, so the effect of the magnetic field from the dipole must also be considered. These crates will only be accessible while LHC is stopped. The next level of electronics will be 60 metres away from the detector behind a thick concrete wall, allowing access during LHC operation.

6.4 Distribution of clock and triggers

In the current scheme, the distribution of the 40 MHz bunch-crossing clock and the Level-0 and Level-1 triggers to the front-end electronics is based on the same TTC (Trigger, Timing and Control) system as that used by the other LHC experiments [1]. The distribution of the Level-1 trigger, not provided in the present design of the TTC, is currently being studied and preliminary results indicate that the system is able to handle Level-1 decision rates up to 1 MHz.

References

- [1] Proceedings of "Workshop on Electronics for LHC Experiments", LHCC/97-60.

Chapter 7 Vertex Detector

7.1 Introduction

The Vertex Detector has to provide precise measurements of track coordinates close to the interaction region. These are used to reconstruct production and decay vertices of B-mesons and provide an accurate measurement of their flight time, and to measure the impact parameter of particles used to tag the B flavour. The Vertex Detector also provides information to the trigger, which enriches the B-decay content of the data.

Silicon detectors with a small strip-pitch and low multiple scattering have been chosen to achieve these requirements. For tracks within the spectrometer acceptance, the optimum arrangement of silicon detector planes is perpendicular to the beam. The extrapolation distance between the first measured point on a track and the primary vertex is reduced by placing the detectors at a radial distance from the beam which is smaller than the aperture required by the LHC during injection. This is achieved by mounting the detectors in Roman pots. A similar design is used for the HERA-B Vertex Detector [1], which experiences particle fluxes similar to those in the LHCb Vertex Detector.

The Vertex Detector is used in the Level-1 vertex trigger. The Level-1 algorithm requires stand-alone three-dimensional pattern recognition to distinguish B-events from those minimum bias events which are accepted by the Level-0 trigger. As a consequence the Vertex Detector has to provide measurements in three independent projections, and tracks in the acceptance of the downstream spectrometer are required to have at least three measured space points.

B-mesons that have all their decay products within the acceptance of the spectrometer are produced with a polar angle that is typically below 200 mrad. The projection of the impact parameter of the decay products to the primary vertex in the rz -plane is large, while in the plane perpendicular to the beam-axis ($r\phi$) it is similar to that of tracks originating from the primary vertex. The Level-1 trigger exploits this by first reconstructing all tracks in the rz projection, but reconstructing only those three-dimensional tracks which have a signifi-




Figure 7.1: Layout of the Vertex Detector. The detector has 17 stations, each consisting of 2 silicon discs (r and ϕ strips), divided into 6 sectors.

cant rz impact parameter. Hence the strip pattern on the detectors has strips with constant radius for the rz -track reconstruction, combined with radial-strip detectors having a $\pm 5^\circ$ stereo angle to allow the two projections to be combined.

The Level-0 trigger accepts events at a rate of 1 MHz, which requires the whole detector to be read out within $1 \mu\text{s}$. Analogue rather than binary readout has been chosen since it allows for better monitoring and control of the effects due to the very non-uniform radiation damage to the silicon detectors.

7.2 Geometry

The requirements outlined above are fulfilled using the geometry shown in Fig. 7.1. The length of the interaction region along the beam is $\sigma_z = 5.3$ cm. With a 4 cm spacing between the detectors and a radial coverage of 1 cm to 6 cm for all detector planes, it is assured that practically all particles

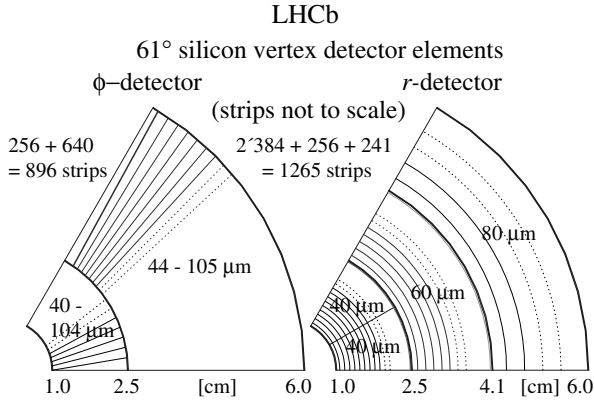


Figure 7.2: Schematic showing the strip layout of both the ϕ - and r -measuring detectors.




Figure 7.3: Channel occupancy for the r , ϕ detectors

measured in the downstream spectrometer traverse at least three stations. Twelve equidistant stations cover the interaction region, while five stations are placed further downstream to cover polar angles down to about 15 mrad. The detectors need to be protected against RF pickup from the beams, therefore they are placed in a secondary vacuum, separated from the LHC vacuum by a thin (100 μm) aluminium foil, which encapsulates each station.

The detector is split in two halves which are retracted by 3 cm during LHC injection. To allow an overlap between the top and bottom half of a station they are displaced by 2 cm along the beam axis. During injection the material at small radii extends inside the 1 cm radius since the detectors are equipped with complex guard-ring structures (1 mm width). In addition, the RF-shield needs to fold around the detectors with sufficient clearance (3 mm), and another 2 mm is reserved to accommodate the wake-field suppressor (see section 7.3), which is therefore situated 4 mm from the beam axis. The aperture required by the LHC machine during physics running conditions is only 1 mm.

Table 7.1: Dimensions of the silicon vertex detector

17 stations, 6 sectors/station, 102 r -sectors			
radius	strip pitch	coverage	channels
1.0-2.536 cm	40 μm	30.5°	2 \times 384
-4.072 cm	60 μm	61°	256
-6.000 cm	80 μm	61°	241
17 stations, 6 sectors/station, 102 ϕ -sectors			
radius	strip pitch	coverage	channels
1.0-2.536 cm	40-104 μm	61°	256
-6.000 cm	41-98 μm	61°	640

Each station consists of two discs, which measure r and ϕ coordinates respectively. Each disc is subdivided in six sectors. The sensitive area of each sector covers 61° in ϕ and 1 cm to 6 cm in radius. Fig. 7.2 shows the detector strip layout. Routing lines to connect the strips to the readout electronics, positioned at a radius of 7 cm, are not shown. The pitch and length of strips have been chosen such that the occupancy per channel, shown in Fig. 7.3, varies by less than a factor of three. The occupancy of the channels measuring the ϕ -coordinate could be further optimised using a more complicated routing of the signals to the electronics. Table 7.1 summarises the main parameters of the detectors. The chosen strip pitches result in a hit precision which varies from 6 μm (9 μm) for 40 μm strip pitch for double (single) channel clusters, to 10 μm (18 μm) for 80 μm strip pitch for double (single) channel clusters.

The Level-1—Level-3 triggers require the Vertex Detector to be aligned to a precision better than the reconstructed cluster position of 6 μm . The strategy is to align all the detector elements inside a half-station to better than 20 μm , and the stations relative to each other to better than 50 μm . A survey of the upper and lower halves of the detector separately should result in an alignment with a precision of 20 μm over the full length of the half-detector. The lower and upper halves need to be mechanically positioned with a precision better than 100 μm . The detectors have a 1° overlap in ϕ which allows for an internal alignment using reconstructed tracks, and will refine the survey results to the required precision of less than 6 μm .

Two silicon discs of similar geometry to those used for the Vertex Detector are placed at the upstream end (opposite to the spectrometer) of the Vertex Detector. They are equipped with fast readout and are used to implement the pile-up veto

in the Level-0 trigger. The detector sectors have 300 r -measuring strips with readout pitch of 120—240 μm , varying with radial distance from the beam. A total of 3600 strips are read out. The r coordinates of outgoing tracks, measured in the two planes, are used to reconstruct an approximate z -coordinate of the primary vertex. The pile-up veto algorithm suppresses bunch crossings containing multiple pp interactions. The detectors and the performance of the algorithm are described in further detail in Section 12.3.3 and in Ref. [2].

7.3 Roman pots

The design of the vacuum tank and support structures of the Vertex Detector needs to satisfy several requirements. It must have low mass in the acceptance region and provide for alignment and retractability of the Vertex Detector. It must take account of mechanical stresses induced by heat loading and maintain high-vacuum compatibility while providing signal feed-throughs, etc.

Figure 7.4 shows a draft design of the vacuum tank and support structures for the Vertex Detector. The tank is cylindrical with a length of 177 cm and a radius of 55 cm (60 cm) at the upstream (downstream) end. The heavy support of the exit window is rectangular and does not obstruct the 470 mrad corner-to-corner acceptance of the downstream tracking stations. Particles within the 330 mrad acceptance of the spectrometer traverse only silicon stations and the exit window. The Vertex Detector stations are mounted, at the top and the bottom, on two aluminium support frames, each attached to a pair of linear actuators via play-free hinges. The two silicon detector planes for the pile-up veto (not shown in Fig. 7.4) are accommodated upstream of the Vertex Detector.

This design allows retraction of the detectors by 3 cm from the beam axis during beam filling. The design parameters for the exit foil are determined by the mechanical stresses, the available space in the z -direction, the large angular acceptance and by multiple scattering. Numerical stress calculations have been performed to produce a realistic design of the exit foil. A self-supporting (spherical) shape has been chosen which minimises tensions on the connecting beam pipe. An aluminium foil of 2 mm thickness and 1 m radius of curvature satisfies all the requirements.

The detectors and front-end electronics will be located in a secondary vacuum. Compatibility issues have been discussed with the LHC group and do not *a-priori* present a problem. The feed-

through of 22 000 signal wires to the outside from the secondary vacuum appears feasible. Fifty twisted-pairs per hybrid will be transported outside the vacuum via “chimneys” and standard 50-pin connectors mounted on the top and bottom flanges.

Leakage currents in the detectors, the electronics and wake-field interactions lead to an expected heat load of about 25 W/station. The detectors will be cooled by heat conduction via the aluminium support frame, which will be maintained at a low temperature. A 100 μm Al foil separates the primary from the secondary vacuum and also provides partial RF-shielding of the detectors. Such a design is likely to exhibit strong coherent RF-coupling due to the ‘cavity-like’ structure of the detector encapsulation, possibly resulting in unacceptable RF pickup and heat load due to the circulating beams. To minimise such RF coupling a wake field suppressor will be mounted between the beam and the Vertex Detector. Different options are being simulated using the MAFIA[3] program. Possible solutions include the use of four metal strips ($\sim 5 \mu\text{m}$ thick stainless steel) or a number of conducting wires spanning the full length of the tank.

7.4 Particle fluxes

The detectors are positioned only 1 cm away from the colliding beams and will suffer radiation damage which may degrade the performance due to incomplete charge collection and increased shot noise resulting from increasing bulk leakage-current. The total fluence at the detector has been simulated [4] with the MARS [5] program. The simulation includes the following: the detectors with their hybrids and readout electronics, the vacuum tank, the concrete walls of the LHC-tunnel¹, the correction dipole (upstream, inside the tunnel) and the LHCb magnet downstream of the Vertex Detector. Fig. 7.5 shows the total charged particle and neutron flux, as a function of radius, for the station at the mean z -position of the interaction point. The contribution from particles originating from the primary interaction is also shown.

DTUJET93 has been used as event generator. At small radii the flux is completely dominated by the primary particles, while at a radius of 6 cm the thermal neutrons, which have a practically flat distribution as a function of the radius, make up one third of the total flux. PYTHIA (with the default parameters) gives fluxes from primary tracks

¹Due to the displaced vertex position the detector is placed in the tunnel rather than in the cavern.




Figure 7.4: Side view of the vacuum tank and support structures for the Vertex Detector. The bottom half has been omitted for clarity.

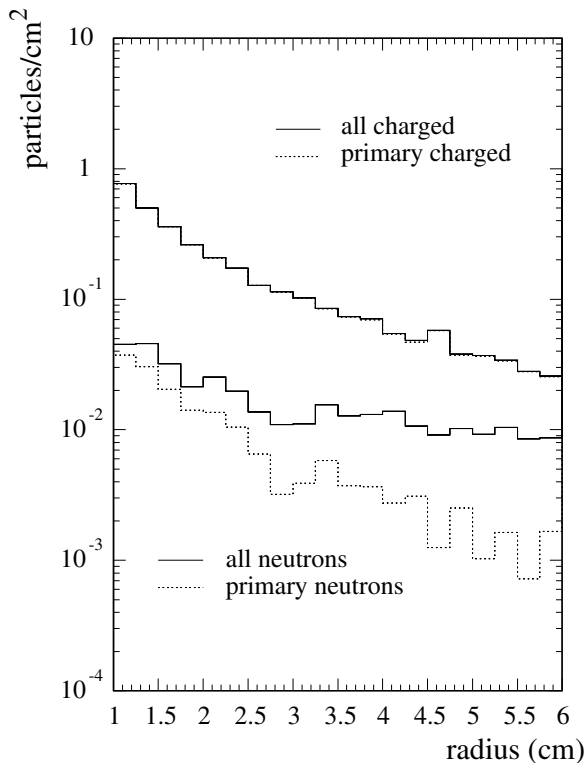


Figure 7.5: Charged particle (upper lines) and neutron (lower lines) fluxes, per interaction, as a function of radius, at station 6 of the Vertex Detector.

about 50% larger than those obtained with DTUJET93. The total charged particle flux is dominated by pions. The total flux varies little between the 17 stations, but the momentum distribution varies widely, from a most probable momentum of 0.8 GeV/c at $r = 1$ cm for station 6, to 20 GeV/c for station 17. As is shown in Fig. 7.6, this results in an equivalent dose per station when expressing the total fluence in terms of damage caused by 1 MeV/c neutrons using the Non-Ionizing Energy Loss (NIEL) tables [6]. The maximum 1 MeV/c neutron-equivalent fluence (over one year) is $0.84 \times 10^{14}/\text{cm}^2$ for the DTUJET93 generator, and $1.3 \times 10^{14}/\text{cm}^2$ for the PYTHIA generator with default parameters. These particle fluxes have consequences for the detector design, as described in the following section.

7.5 Detector design

7.5.1 Introduction

The impact-parameter resolution is most sensitive to the track hits at the smallest radii, and hence a conservative design is chosen for the silicon detectors which must run in the hostile environment close to the LHC beam. The most radiation tolerant silicon devices currently available are n-side readout single-sided detectors. Given the expected radiation doses and the present understanding of

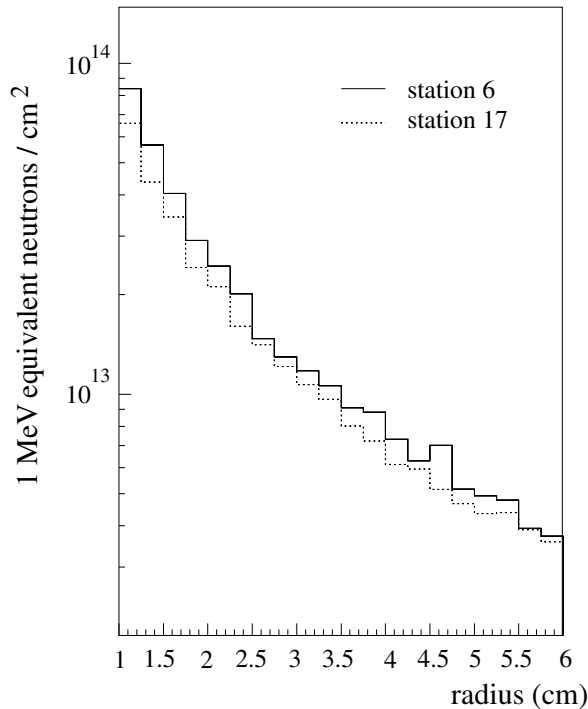


Figure 7.6: Total particle fluence at two stations of the Vertex Detector as a function of the radius. (For one year of running and using the DTUJET93 generator.)

the radiation resistance of these detectors they could last for several years. However, the baseline operating scenario is to replace the innermost detectors every year.

The bulk leakage current of the detectors increases linearly with the thickness and the full-depletion voltage is proportional to the square of the thickness. Hence a thin detector has low current and low total power dissipation, at the expense of a loss in signal and an increase in the capacitance to the backplane. Thin detectors have the additional advantage of lower multiple scattering and the low power dissipation reduces the amount of material needed to cool the inner edge of the detectors. These considerations have led to the adoption of $150\ \mu\text{m}$ thick detectors, which are the thinnest detectors that may be easily and economically manufactured.

7.5.2 Technology

It is foreseen that the detectors will employ polysilicon biasing, n-strip readout and individual p-stop isolation between the strips. Poly-silicon biasing avoids the problems of excess noise introduced by reach-through biasing as reported [7] by ATLAS

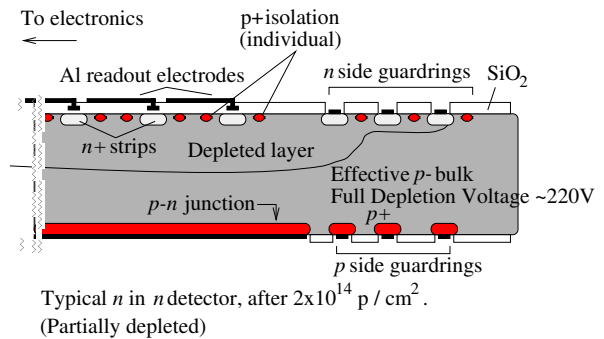


Figure 7.7: Schematic cross-section of the silicon detectors showing the depletion region after irradiation.

and CDF. Figure 7.7 is a schematic of these detectors showing the (partially) depleted region of the silicon, after type inversion due to irradiation. The n-implants are capacitively coupled to aluminium strips. The aluminium has a resistivity of about $20\ \Omega/\text{cm}$ and a capacitance to the implant of about $12\text{pF}/\text{cm}$.

The edge regions of these silicon detectors require careful design to avoid excess currents being drawn, severely compromising the performance of the detector. Indications from ATLAS tests are that multi-guard ring structures give adequate protection before and after irradiation. The detectors will have n-strips of width $10\ \mu\text{m}$ and p-stop (guards) of width $5\ \mu\text{m}$. The narrowest pitch foreseen for LHCb detectors of $40\ \mu\text{m}$ results in a separation of $6\ \mu\text{m}$ between adjacent guards and $7\ \mu\text{m}$ between the guard and its associated n-strip. To avoid micro-discharges at the ends of strips, all ends have been rounded.

Insufficient space is available (without introducing significant dead areas) to route out all the strips in a single metal plane. A double-metal technology as applied by DELPHI [8] has been chosen. This enables the signal from a strip close to the beam to be routed out over the top of other structures via a $1\ \mu\text{m}$ thick, $20\ \mu\text{m}$ wide aluminium strip (resistance $< 20\ \Omega/\text{cm}$) separated from the rest of the detector by a thin ($2\text{--}5\ \mu\text{m}$) SiO_2 or polyamide dielectric layer. The capacitance of the readout strip should be less than $1\text{pF}/\text{cm}$. The layout of the routing lines to the readout electronics is constrained by the need to minimise cross-talk and capacitance while preserving the strip order in the readout chain to facilitate the correction of local common base-line shifts. Figure 7.8 shows the routing lines of an LHCb prototype detector manufactured by Hamamatsu Photonics.




Figure 7.8: Schematic of the double-metal routing scheme of the 1997 prototypes detectors. The upper detector has ϕ -measuring strips and the lower has r -measuring strips.

7.5.3 Radiation resistance

The particle flux in the region of the Vertex Detector varies rapidly as a function of the radial distance from the beam, and after a few months of operation part of the detector will have undergone type-inversion. The response of this detector design, as a function of fluence, has been modelled using parameterisations from the ROSE (RD48) collaboration [9]. Figure 7.9 shows the expected full-depletion voltage, current and charge collected in 25 ns as a function of the particle fluence, the temperature and the radial distance from the beam line. An initial resistivity of 480 Ohm cm is assumed to move the inversion point of the silicon to higher fluences². This results in a full-depletion voltage for the detector of about 150V before irradiation. An exposure time of 8 months is assumed to acquire the total dose. Throughout this time the

²Not all detectors exhibit this same inversion point dependence as a function of the resistivity.




Figure 7.9: The expected total-depletion voltage, current and charge collected in 25 ns as a function of the particle fluence (Φ), temperature (T) and radius (R) to the beam line of $150\mu\text{m}$ thick detectors. For the current and charge collection curves an operating voltage of 200 V has been assumed. The quantities shown as a function of the temperature are given at $R=1$ cm.

temperature is kept constant.

The above parameterisations have been compared with detectors built by Hamamatsu Photonics for the ATLAS experiment. A major difference is that the LHCb detectors are equipped with a double metal readout layer and are $150\mu\text{m}$ thick compared to $300\mu\text{m}$ as in the ATLAS design. The ATLAS detectors were exposed to 2×10^{14} charged hadrons/cm² at -10°C and their properties studied after exposure as a function of time at 20°C . The model proposed by the ROSE collaboration agrees with the experimental results as obtained from the ATLAS detectors [10] on the essential properties, such as total leakage current and the charge collection efficiency.

The radiation-induced currents in the detectors show time-dependent effects which vary as a function of the temperature of the silicon. Beneficial (short term) annealing occurs at room temperatures as does a longer term “reverse annealing” which manifests itself as an increase in the full




Figure 7.10: Schematic of a silicon module. Two silicon detectors(1) are separated by a thermal spine(2). There is a separate hybrid(5) for each silicon detector. Inserts (3), (4) and (6) provide thermal insulation between the electronics and the detector.

depletion voltage. The most promising operating temperature appears to be about 5°C for the case of detectors which are replaced every year. For detectors which must operate for several years a lower operating temperature is preferable.

Following irradiation, the charge collection efficiency decreases due to trapping in the bulk; the capacitance of the strips and pickup lines will also change. Assuming that a perfect $150\mu\text{m}$ thick detector produces a signal of about 12500 electrons/MIP which is sampled over 25ns the performance after irradiation can be estimated. Although it is difficult to predict, and is one of the key measurements to be made on the prototype detectors, it is expected that such a detector operated at full depletion will give a signal larger than 11000 electrons for after exposure to a fluence of 2×10^{14} neutrons/cm².

7.5.4 Module design and cooling

Figure 7.10 shows the concept of a basic detector module. It consists of one r -measuring and




Figure 7.11: Thermal model of an LHCb module assuming uniform power dissipation in the silicon.

one ϕ -measuring sector, two hybrids carrying the front-end electronics, inserts to thermally insulate the heat-load of the front-end chips and to provide electrical screening from the silicon detector, and a heat-conducting spine to ensure a well-defined temperature at the inner edge of the detector. The thickness of one station, including the 0.1 mm Al RF-shields, amounts to 0.7% of a radiation length in the sensitive area of the detector.

Power dissipation in the silicon bulk heats the detector causing a local increase of the operating temperature which will then result in larger currents and full-depletion voltages as is shown in Fig. 7.9. This could cause thermal runaway and a cooling system is required to maintain a constant temperature, especially at the smallest radii. Figure 7.11 shows a thermal map of the module with the cooling system operating at -10°C ; 3 W for the front-end electronics and a uniform power dissipation over the module of $100\mu\text{W}/\text{mm}^2$ is assumed. Figure 7.12 shows the temperature at the tip of a module as a function of the power density. The model is time dependent and starts with the silicon at -10°C . It shows the temperature at the tip of the detector after uniform irradiation of 5×10^{14} charged hadrons/cm². The wake-field suppressors will ensure that the RF-shields do not heat up, so heat transfer from the shields to the detectors has not been considered in the model. For high power dissipation the temperature rises catas-




Figure 7.12: Operating temperature at the detector tip as a function of the power dissipation density.

trophically, and the onset of thermal runaway is at about $300\mu\text{W}/\text{mm}^2$. In reality the irradiation will fall with the square of the radius. A maximum power density of $150\mu\text{W}/\text{mm}^2$ will be obtained at the tip, falling rapidly with increasing radius. The modelling shows that the detectors can be safely operated well below thermal runaway.

7.6 Front-end electronics

7.6.1 Introduction

The detector signals from minimum ionising particles are about 12500 electrons. Analogue, rather than binary, readout has been chosen since it allows for a better monitoring and control of the radiation damage in the detector. Furthermore it provides better precision of the reconstructed hits and better discrimination between hits and clusters due to electronic noise [11]. It is also more robust against possible common baseline shifts. The general layout of the readout scheme is shown in Fig. 7.13. The 222 000 pre-amplified detector signals are stored in an analogue buffer for the period pending the Level-0 decision. Following a Level-0 Accept, 32 channels are multiplexed on one analogue readout line, and driven over about 10m to a FADC. Where possible the electronics is moved away from the detector to a radiation environment allowing standard components, but close enough to avoid the use of expensive optical-electronic analogue data transmission [12]. Following digitisation the information is split into two paths. In the first path, the data are pre-processed and sent to the Level-1 processor. In the second path the data are stored in the Level-1 buffer and sent, at Level-1

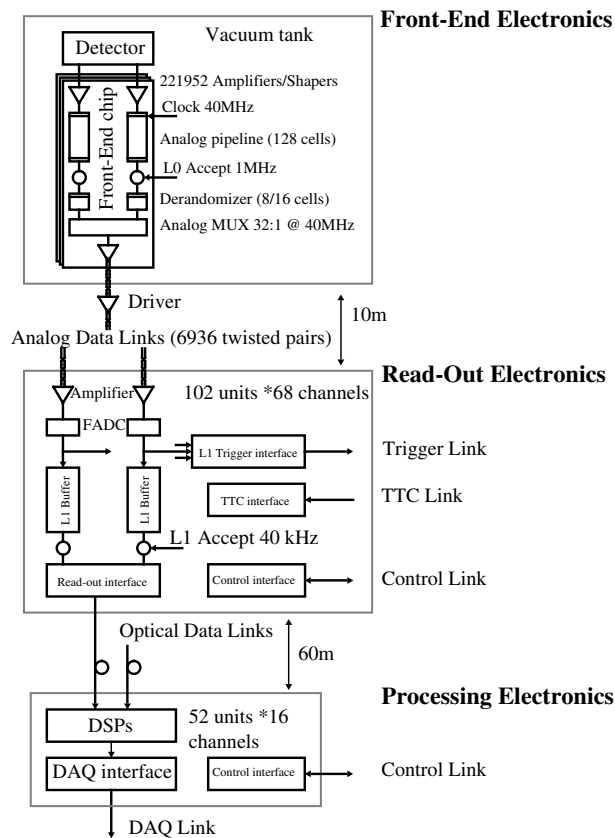


Figure 7.13: Overview of the readout scheme for the Vertex Detector.

Accept, to be processed using a more sophisticated zero-suppression algorithm.

7.6.2 On-detector electronics

After radiation damage a signal-to-noise (S/N) ratio better than eight is required to maintain a high efficiency with a tolerable level of clusters due to electronics noise [11]. The on-detector electronics are situated at about 7 cm from the beam-line, hence the resulting radiation tolerance of the chips must be better than a few hundred krad if they are to be replaced at most once per year. The SCT-A [13] is one chip that is close to meeting these requirements. The S/N performance expected from this chip, assuming the noise performance of the 1996 prototype ($\text{ENC}=720+33/\text{pF}$), as a function of the 1 MeV/c neutron equivalent flux in the detector, is shown in Fig. 7.14.

However, the requirements imposed by the Level-0 trigger rate mean that the SCT128A would need the following modifications [14]:

- The present length of the Level-0 pipeline




Figure 7.14: Expected S/N as a function of the 1 MeV/c neutron equivalent flux in the detector for a front-end chip with noise $ENC=720+33/\text{pF}$ for the strip at $r=1$ cm (full line) and $r=6$ cm (dashed line). The dotted curve is for a detector at $r=1$ cm with twice the leakage current of that described by the full line.

is 118 and another 8 for the derandomiser FIFO. To increase the pipeline length to 128, plus 16 (instead of 8) for the derandomiser FIFO, requires a major modification of the pipeline control logic.

- To read 32 multiplexed channels using only 40 clock cycles needs a re-design of the readout logic. The existing design uses 26 clocks for internal sampling, leaving only 14 clock cycles for the readout.
- In the 128-channel version of the chip the area occupied by the pipeline control logic is roughly the same as the area of the pipeline itself. A 32-channel version of the chip would be dominated by the control logic, hence a 128-channel version seems preferable, sharing the control logic between four multiplexers.

Furthermore, only every third position in the pipeline can be read out consecutively, leading to a 5% dead-time.

Given the number of modifications needed that are specific to the LHCb trigger scheme, it appears necessary to design a new front-end chip specifically for LHCb. One possibility is to port the experience from the AMS/HELIX-chip [15], developed for

HERA-B, to a DMILL version that satisfies LHCb requirements. This development could be undertaken together with that proposed (Section 8.2.1) for the readout of the MSGC detectors for the Inner Tracker.

The current design of the HELIX chip is not radiation hard (beyond 100 krad) and allows writing to the pipeline at the HERA bunch-crossing rate of 10 MHz. The maximum latency of the pipeline is 128 write-clock cycles and the derandomising buffer has a depth of 8. The major modifications to the HELIX design would thus affect the front-end and the multi-event buffer. To set up the readout of a specific pipeline cell an overhead of five write cycles is required and a 40 MHz readout of 32 multiplexed channels in $1 \mu\text{s}$ should be feasible without major redesign of the existing layout. The HELIX design allows consecutive readout of adjacent pipeline cells without deadtime loss. In addition it can supply the fast trigger output, which is required for the vertex pile-up trigger.

7.6.3 Off-detector electronics

With a multiplexing ratio of 32 the total number of analogue channels requiring digitisation is 6936. Figure 7.13 shows a possible solution using two modules. In the first, the data are digitised and stored in the Level-1 buffer. This module also performs the cluster preprocessing for the Level-1 trigger and executes the following steps.

- Pedestal subtraction: pedestals might depend upon the pipeline location and channel.
- Common baseline shift correction: due to the high voltages, micro-discharges might produce pick-up on the detector backplane.
- Cluster search: Clusters are found and encoded into addresses.

One detector sector contains 32×68 channels, hence a natural unit would pack 68 FADCs with corresponding Level-1 buffers and cluster finding hardware on one board, resulting in 102 boards. After a Level-1 Accept, the data are processed by DSP's. The output provides zero-suppressed data as well as pedestal and threshold information for the fast cluster search used in the Level-2 trigger. Assuming a conservative factor five in the optimisation of the DELPHI [8] algorithm and using 1000 MIP DSP's would require 816 DSP's, each DSP processing 272 channels. Given the high trigger rate of LHCb compared to DELPHI, it will not be necessary to update the channel-dependent information

Table 7.2: Precision of r -clusters as a function of the strip pitch and the number of channels in the cluster. The fraction of clusters with 1, 2 or more channels per cluster is listed for each pitch.

Pitch	Channels/clusters					
	1		2		≥ 3	
	σ	%	σ	%	σ	%
40 μm	9 μm	56	6 μm	33	20 μm	11
60 μm	13 μm	51	10 μm	37	30 μm	12
80 μm	18 μm	49	10 μm	39	40 μm	12

for every event, thus considerably reducing the processing time required to do the cluster search.

7.7 Performance

The Vertex Detector simulation is described in Ref. [16]. In the present version of the programme, charge is distributed in the detector according to the fraction of the path length in the silicon covered by that strip. Diffusion of the charge is not taken into account. Simulations show [11] that, with proper charge diffusion, the reconstructed cluster position will improve. The charge in a given channel is smeared by the expected electronic noise and hits are required to exceed three times the standard deviation of the noise.

Table 7.2 summarises the precision obtained for reconstructed clusters in the r -measuring detectors. No attempt has been made to decompose the clusters with three or more channels into the contributions from overlapping tracks, nuclear interactions, etc. The error for these clusters has been inflated to reduce their weight in the track fit.

The most probable number of stations traversed by tracks originating from $B_d^0 \rightarrow \pi\pi$ decays is five. 5.5% (2.7%) of these pions traverse less than three (two) stations. The precision in the rz -projected impact parameter of the decay pions is shown in Fig. 7.15 as a function of their momentum.

High radiation levels cause a decrease in the S/N ratio, as shown in Fig. 7.14. The online cluster-finding algorithm can either be adapted to keep a constant fraction of noise clusters, at the expense of a loss in efficiency, or by minimising the loss in efficiency, with a consequent increase in the number of noise clusters. The effect of this increase in noise on the Level-1 trigger is discussed in Section 12.4.1.

As an example, the efficiency in reconstructing the decay $B_d^0 \rightarrow \pi\pi$ has been studied [17]. The re-




Figure 7.15: The resolution of the rz -projected impact parameter of tracks as a function of their momentum. The impact parameter is calculated relative to the true position of the primary vertex of origin.

sults are summarised here for three different dose scenarios in terms of performance relative to the result obtained from a non-irradiated detector. Imposing a conservative requirement, that the number of noise clusters should be kept below 0.1%, reduces the cluster-finding efficiency to 92%, 61% and 0%. The S/N values at $r=1$ cm are 10, 8 and 5 respectively and correspond to 1 MeV/c neutron equivalent fluences of 5, 8 and more than $10 \times 10^{14}/\text{cm}^2$ (corresponding to 4, 6 and more than 8 years running at the nominal LHCb luminosity). For a fluence of $10 \times 10^{14}/\text{cm}^2$ the efficiency at $r=2$ cm is 92%.

The vertex of $B_d^0 \rightarrow \pi\pi$ events is reconstructed requiring the two pions to have a distance of closest approach smaller than $100 \mu\text{m}$. The number of events passing this cut, for the three scenarios is 100%, 99% and 93% respectively, compared to the result obtained with no irradiation. The corresponding vertex error increases by 2%, 10% and 60%.

7.8 Future R&D

R&D for the vertex detector is foreseen in the following areas:

- Silicon detectors. Prototype detectors (see




Figure 7.16: Three r -measuring and three ϕ -measuring prototype detectors are shown with a 19 mm diameter coin at the centre.

Fig. 7.7 and 7.16) have been designed and ordered from Hamamatsu Photonics. These detectors will be characterised in 1998, before and after irradiation, to enable the definition of a final detector design. A pre-production series of detectors can then be ordered to undergo long term performance tests and to construct a prototype half-station in 1999. Based on this half-station a final design can be made, allowing time for the detectors to be delivered before the end of 2002.

- Front-end electronics. LHCb is participating in a development of rad-hard BiCMOS analogue (25 ns integration time) readout chips (SCT128-A). Delivery is expected in spring 1998. Some of the prototype detectors will be equipped with these chips to evaluate their performance with a detector that is close to our final design. By summer 1998 a recommendation on the design of the front-end chip will be made. An LHCb front-end chip prototype should be ready by the end of 1999. The final chip should be produced and tested before the end of 2002.

The set-up used to evaluate transmission of analogue signals at 40 MHz [12], will be extended to include the full chain from detector, via vacuum feed-through of the signals, to the FADC.

The FADC and DSP boards and their link to the trigger processors will be reviewed in 1998, and a detailed design of the major com-

ponents prepared. A prototype chain will be made in 1999-2000. It will operate at LHCb speed from the detector to the Level-1 buffer. A review of the overall readout chain is foreseen in 2001, and production will start in 2002.

- Mechanics. Detailed design and simulation of the modules and their heat load will continue during 1998. Small scale prototypes to study cooling and vacuum feed-throughs will be produced by middle of 1999. The design of modules, vacuum tank and cooling will be completed by 2000.

References

- [1] M. Schmelling, "Design Issues of Vertex Detectors for B-physics at Proton Storage Rings, the case of the HERA-B Vertex Detector", LHCb/97-025.
R.St.Denis, et al., "First Experience and Results from the HERA-B Vertex Detector", Contribution to the Sixth International Workshop on Vertex Detectors, VERTEX 97, September, 1997, Rio de Janeiro, Brazil.
- [2] M.van Beuzekom, et al., "The use of Silicon Detectors for fast primary vertex reconstruction and pile-up rejection." LHCb/97-016 TRIG.
- [3] MAFIA: Users Guide. CST GmbH, Darmstadt, Germany.
- [4] V.Talanov, "Study of particle fluxes and absorbed radiation doses in the vertex detector" LHCb/98-019.
- [5] I.Azhgirey, I.Kurochkin and V.Talanov, "Development of MARS Code Package for Radiation Problems Solution of Electro-Nuclear Installation Design", Proc. of XV Conference of Charged Particle Accelerators, Protvino, October 22-24 (1996).
- [6] G. Lindstroem has kindly supplied us with the latest tables as compiled by the ROSE collaboration.
- [7] ATLAS Inner Detector TDR: CERN/LHCC/97-17, Vol II.
- [8] V. Chabaud et al., "The DELPHI silicon strip microvertex detector with double sided read-out", Nucl. Instr. and Meth. **A368** (1996) 314.

- [9] H. Feick, "Radiation Tolerance of Silicon Particle Detectors for High-Energy Physics Experiments" thesis, DESY F35D-97-08, August 1997.
Application of this model to the LHCb specific detectors is described in a note: LHCb/98-043.
- [10] Th. Bowcock, "Tests with prototype ATLAS silicon detectors." LHCb/98-048.
- [11] P. Koppenburg, "Comparison of analogue and binary read-out in the silicon strips vertex detector of LHCb". LHCb/97-020.
- [12] J. Buytaert et al., "Transmission of analogue information at 40 MHz." LHCb/98-032.
- [13] F. Anghinolfi et al, "SCTA - A Rad-Hard BiCMOS Analogue Readout ASIC for the ATLAS Semiconductor Tracker." IEEE, Trans. Nucl. Sci. Vol. 44, No. 3, June 1997.
- [14] Wladyslaw Dabrowski, private communication.
- [15] W. Fallot-Burghardt, A. Holscher, "HELIX-A Silicon Strip/MSGC Readout Amplifier" HD-ASIC-05-0595.
- [16] D. M. Steele, "The R- ϕ Silicon Vertex Detector Simulation as Implemented in SICB." LHCb 98-004.
- [17] O. Cooke, "The effect of radiation damage on the vertex detector efficiency." Note that the conversion between fluence, S/N and years uses old data. This is updated in the Technical Proposal. LHCb 97-023.

Chapter 8 Main Tracking Detectors

8.1 Introduction

The principal task of the tracking system is to provide efficient reconstruction of charged-particle tracks and precise measurements of their momenta. It also has to provide measurements of track directions in both x and y projections for the reconstruction of Cherenkov rings in the RICH detectors. Finally, tracking hits are used in the Level-1 and Level-2 triggers.

The main tracking system comprises eleven tracking stations, situated between the Vertex Detector vessel and the calorimeters. These are shown as T1—T11 in Fig. 2.5. The geometry of the detectors and the choice of technologies are motivated by the requirement of low occupancies. The required spatial resolution is then easily satisfied. Most of the detector acceptance is subject to a particle density which is low enough to use drift chambers with small cells. This Outer Tracking System consists of straw-tube-like chambers. The Inner Tracking System covers (in most stations) an area of $60 \times 40 \text{ cm}^2$ around the beam pipe, where a much finer granularity is necessary. For this inner region several alternative technologies are under study.

The tracking system provides precise particle trajectories in the bending plane of the spectrometer using near-vertical wires and strips. Three dimensional track reconstruction is achieved using small stereo angles of $\pm 5^\circ$. The stations situated up- and downstream of the RICH detectors are equipped with additional planes which provide precise measurements in the non-bending projection. These allow precise measurements of the track direction in both projections.

A track reconstruction program [1] has been developed, based on the technique of Kalman filters. For most tracks, the reconstruction is started by following tracks in the upstream direction. A track seed is found in station T10, which defines a search window in the next station, T9. The track is then followed, station by station into the Vertex Detector. In each station, the newly associated hits are used to update the track parameters. Using seeds found in the Vertex Detector, a second pass, in the downstream direction, reconstructs

those tracks which leave the acceptance of the spectrometer before they reach station T10 and extrapolates the remaining tracks through T11 and into the muon detectors.

The track reconstruction program is used to design the layout of the tracking system and to optimize its performance. Details are given in Ref. [1]. Stations T10 to T7 are used for track finding. In this region, there is almost no magnetic bending and redundancy is high, since the distance between the stations is small. The momentum measurement of the particles is provided principally using stations T6 to T3. The precision is dominated by multiple scattering, so the number of stations in this region is minimized. This also reduces the number of photon conversions. The distance between two stations is chosen so that less than 10% of the search windows contain a second track.

Chambers in the high magnetic field, close to the centre of the magnet, have a higher occupancy due to curling low-energy tracks. Access for cables and services is also limited in this region. Hence, the number of stations inside the magnet is kept as low as possible.

The stations T2 and T1 do not contribute significantly to the momentum measurement. They are required to link the tracks from the main tracking system to the track segments in the Vertex Detector. These stations are crucial for track following in the downstream direction, where the track segments found in the Vertex Detector are extrapolated through the magnet.

The Monte Carlo simulations of occupancies and the radiation levels presented in this chapter are based on the nominal LHCb luminosity of $\mathcal{L} = 2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$, which corresponds to an average of approximately one interaction every four bunch crossings. The LHCb experiment has to deal with particle densities which are similar to those experienced at HERA-B. The tracking detectors proposed for LHCb have a layout and employ technologies which are similar to those in HERA-B. The consequences related to short-term operation with higher luminosities (a maximum of $\mathcal{L} = 5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ is considered) differ for the

Table 8.1: Dimensions of the active area in the Inner Tracking System.

Station	z cm	Θ_{\min} mrad	x, y_{\min} cm	x_{\max} cm	y_{\max} cm
T1	100	50	5.0	35	35
T2	229	16	3.6	30	30
T3	321	14	4.5	30	20
T4	416	12	5.5	30	20
T5	556	12	6.9	30	20
T6	691	12	8.2	30	20
T7	781	12	9.1	30	20
T8	831	12	9.6	30	20
T9	881	11	10.1	30	20
T10	928	11	10.5	30	30
T11	1174	11	13.0	30	30

different detector technologies and are referred to in the following sections.

8.2 Inner tracking detectors

The Inner Tracking system covers the region around the beam pipe, where the particle flux is too high for conventional drift chambers. The positions and dimensions of the stations are summarized in Table 8.1. The outer boundaries are defined by the occupancy requirement in the outer tracker cells. Since the particle flux at a given distance from the beam varies only slowly with z , this boundary ($60 \times 40 \text{ cm}^2$) is the same in most stations. The stations immediately up- and downstream of the RICH counters include horizontal wires (or strips). To reduce the occupancy in the horizontal wires of the outer tracker, the inner tracker has dimensions $60 \times 60 \text{ cm}^2$ in these stations. The inner boundaries are defined by the conical beam-pipe.

Station T1 is special. The outer tracker would only cover a small region, which makes it more economical to construct T1 entirely using the inner detector technology. This station is situated close to the last disc of the Vertex Detector and does not have to cover the region close to the beam-pipe.

Several detector technologies meet the requirements of the inner tracker. Three options are described here. The MSGC/GEM detector has been chosen as the basis for cost estimates, although further R&D is required. The well-proven fall-back solution, silicon strip detectors, is estimated to be significantly more expensive. To reduce costs even further, a new type of detector, the MCSC,

is currently being investigated. The process and time schedule for making the technology choice is described at the end of the section.

8.2.1 MSGC/GEM detectors

The proposed microstrip gas chamber (MSGC) inner detector is based on the development and test of similar detectors for HERA-B. Particle rates and radiation doses expected at LHCb are about the same as those in HERA-B, which therefore provides an ideal test of the technology. There are some differences however. The bunch-crossing frequency at LHC is four times higher than at HERA. This requires a faster shaping of the signals, and different detector dimensions to limit the maximum drift times. In addition, the maximum Level-0 trigger rate is ten times higher at LHCb with corresponding requirements for the readout system. The detector for LHCb, described below, is based on presently available and tested technology with only minor ('safe') changes compared to the HERA-B detectors. Directions for further improvements are discussed below.

Severe aging and operation problems with MSGC detectors have been reported. The introduction of conductive 'diamond'-coating of the substrates, in combination with the use of appropriate materials and gases, has solved the aging problems and has led to detectors with good spatial homogeneity, good high-rate behaviour and a remarkable robustness of operation even in the presence of defects in the electrode or surface structure. A major problem was, however, the observation that MSGC's show gas discharges when operated at 'useful' gains of ~ 3000 in intense hadron beams. These discharges gradually destroy the electrode structures. The problem has been solved by the introduction of an additional amplification stage using a gaseous electron amplification (GEM) foil [2] on top of the MSGC. With a typical gain of 20 at the GEM, the gas amplification of the MSGC can be reduced below 200. The HERA-B inner tracker has shown that MSGC/GEM detectors, operated with Ar:DME 50:50, have excellent high-rate properties, show no gas discharges in intense hadron beams, and no aging up to integrated doses of at least 8 mC/cm strip length (equivalent to more than three years of LHCb operation at nominal luminosity). Similar results have been obtained by a CERN group using an Ar:CO₂ mixture [3]. Further tests at higher integrated doses are underway. It has also been shown that the detectors can be operated in a magnetic field (0.85 T) without loss of




Figure 8.1: Schematic cross section through a MSGC/GEM detector (not to scale).




Figure 8.2: Front view of one inner tracking layer consisting of four overlapping detectors. Dimensions are mm.

efficiency and resolution. The LHCb dipole magnet produces 1.1 T, so these results require confirmation at the higher field.

Detector configuration

A cross section through a typical detector is shown in Fig. 8.1. It has dimensions of $\sim 28 \times 25 \text{ cm}^2$ with a corner removed to accommodate the beam pipe. One layer is made of four overlapping detectors, having the same strip orientation, covering the sensitive region around the beam pipe. A view of the detector geometry for one layer is shown in Fig. 8.2.

The MSGC electrode structure for one chamber is manufactured on a single AF45 glass substrate with a conductive ‘diamond’ coating (\sim

$10^{15} \Omega/\text{square}$). The electrodes consist of gold (500 nm thick) with a pitch of $220 \mu\text{m}$, an anode width of $10 \mu\text{m}$ separated from the cathode by a gap of $60 \mu\text{m}$. The technology required to produce MSGC plates of this dimension has been established by the HERA-B production process¹. The GEM foil is kapton, $50 \mu\text{m}$ thick, and coated with $2 \times 5 \mu\text{m}$ of copper. It has $60 \mu\text{m}$ holes, staggered with a pitch of $140 \mu\text{m}$. These GEM foils are produced in the CERN workshop with good yield and quality. They allow safe operation at high rates up to gains of about 50 (operation at a gain of about 25 is preferred). GEM’s coated with a diamond layer show good stability and a much improved spatial homogeneity of the gain. The chamber construction is based on two hollow G10 frames, 8 mm wide and 2.0 mm high, which define the gas volume above and below the GEM foil. This frame is also used for the gas distribution in the detector. The chamber is closed by the drift electrode consisting of an aluminium-coated kapton foil ($100 \mu\text{m}$ thick). The detector components are glued together using UHV epoxy glue. This chamber construction is presently used for HERA-B except that the gap heights are 3.0 mm for HERA-B instead of 2.0 mm.

The effective detector thickness in the sensitive area, including support structures, is $0.55\% X_0$. The frame, which is 8 mm wide and covers 11% of the area, has a thickness of $1.0\% X_0$. The sensitive area starts at least 16 mm from the beam-pipe. The MSGC inner tracking system has a total of ten stations (if T1 is assumed to use silicon detectors), each with four detector layers having vertical and near-vertical strip orientations ($90^\circ, 95^\circ, 85^\circ, 90^\circ$). Including the horizontal-strip layers in T2, T10 and T11, this results in a total of 43 layers and 172 MSGC/GEM detectors each with 1280 strips. The total sensitive area covered is 13.5 m^2 , equipped with 221000 readout channels. This is similar to HERA-B where an area of 10.1 m^2 is read out by 140000 strips.

Operation characteristics

The gas gain achieved in an Ar:DME 50:50 mixture is shown as a function of cathode voltage in Fig. 8.3 for different GEM voltages. The dramatic reduction of cathode voltage which can be achieved for the same gas gain using the GEM technology is clearly visible.

¹Substrates of the necessary size (up to $30 \times 30 \text{ cm}^2$) are produced commercially. Diamond coating is presently produced by the Fraunhoferinstitut for Surface engineering and thin films, Braunschweig, Germany. Electrode structures are produced by IMT, Greifensee, Switzerland.

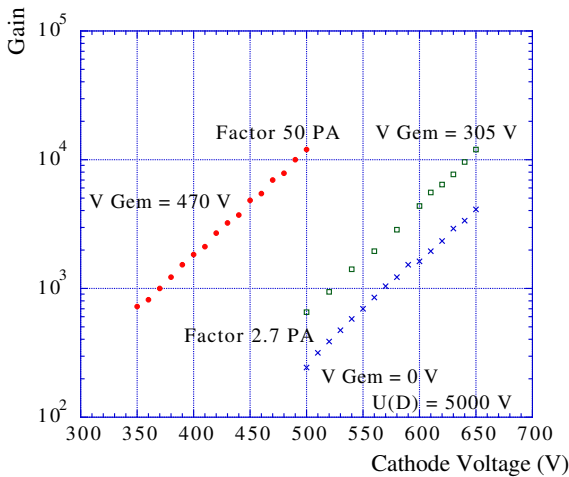


Figure 8.3: Gas gain for a MSGC/GEM detector with Ar-DME 50:50 mixture.

This gas mixture provides high primary ionization, reliable operation and good aging properties. It has been extensively tested and will be used for HERA-B. The chamber operates with a drift field of about 7.5 kV/cm resulting in a drift velocity of $52 \mu\text{m}/\text{ns}$ and a maximum drift time ≤ 40 ns. For operation in the magnet (1.1 T) the Lorentz angle at 7.5 kV/cm is about 9° , the expected strip multiplicity is 1.55 and the precision $\leq 65 \mu\text{m}$.

The mean primary ionization of $15 e^-$, a gas gain of 4000 and an effective integration time of 25 ns (50 ns pulse duration) lead to a mean detected signal of $22000 e^-/\text{MIP}$. This is adequate for a good signal-to-noise ratio and an efficiency $> 95\%$ (assuming noise of $\leq 1400 e^-$). The high gain which can be guaranteed by the MSGC/GEM technology is crucial. They show however a rate dependence of the gain due to charge-up and polarization effects which also lead to local inhomogeneities in the gas gain. Recently GEM's have been tested for HERA-B, where the GEM's are coated with a diamond layer. These show very good stability of gain and the homogeneity of the gain over the detector area is about 10% for diamond-coated GEM's, compared with 30% for the uncoated GEM's.

The expected maximum and average occupancies for the proposed LHCb detector layout are shown in Fig. 8.4. The occupancy never exceeds 4%. Correspondingly the maximum integrated charge per cm of strip-length is $\leq 2.5 \text{ mC}/\text{cm}$ per year at nominal luminosity. It is assumed that strips which fire in one bunch crossing are disregarded in the next bunch crossing. Pile-up will not, therefore, increase the occupancy but will decrease




Figure 8.4: Particle density per cm^2 for one interaction multiplied by r^2 [cm^2] in each station (left). Maximum and mean occupancy in each station at nominal luminosity (right).

the efficiency. These operating conditions provide adequate performance at the nominal LHCb luminosity. Faster gases with high primary ionisation are clearly desirable for high luminosity running and will form part of the LHCb-specific R&D programme.

Read-out electronics

The readout of the inner tracker has to provide the following:

- Fast shaping of signals to cope with the 25 ns bunch-crossing rate;
- Deadtimeless storage (or readout) of the strip information at the maximum Level-0 trigger rate of 1 MHz;
- Fast trigger signals for the Level-1 tracking triggers, for all events accepted by Level-0.

None of the readout chips developed for HERA-B or LHC is able to satisfy all requirements. The readout chip development for MSGC/GEM detectors could however benefit from the experience and partly from the design and layout of the HELIX128 chip developed at Heidelberg for the HERA-B strip detectors. The following developments are necessary:

- A new input stage with faster shaping;
- Radiation hard technology. The transfer to DMILL is already planned by HERA-B.
- A new data storage and readout scheme. Analogue data is stored in the Level-0 pipeline (see Fig. 8.5). Following the Level-0 trigger, two possible readout schemes are

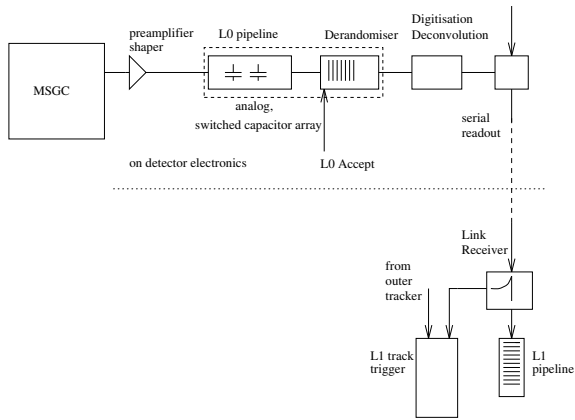


Figure 8.5: Schematic of the MSGC readout system.

considered, both requiring a substantial R&D effort.

- Multiplexing of 32 channels and analogue readout at 40 MHz, followed by analogue deconvolution and digitization in the link receiver.
- Digitization and read out with 3 bits per bunch crossing, using a GHz link for every 256 channels.

In both schemes the Level-1 trigger signals are derived from the Level-0 data buffers.

Supports and infrastructure

Detailed solutions for the chamber mechanics, cabling, gas and HV supplies have been worked out for HERA-B and can be used with minor modifications for LHCb. Each station is rigidly fixed and aligned to the outer tracking detectors.

Future R&D

The detector development and testing of the baseline MSGC/GEM solution will be done within the HERA-B project. In 1998 a fully equipped station will be operated in HERA-B. Mass production of detectors also starts early in 1998. This important field test will allow a realistic evaluation of the detector technology and performance by the end of this year.

LHCb-specific developments towards improved detector technology include:

- Optimization of the GEM technology.
- Investigation of faster gas mixtures with high primary ionization.

- Reduction of material, for example, through the development of 2-dimensional readout.
- Testing of advanced GEM technologies which might avoid the use of MSGC plates altogether.

8.2.2 Silicon Micro Strips

Silicon-strip detectors are well-proven and considered as a fall-back solution. They offer the following advantages:

- The occupancy is less than 2.5% per channel with a readout pitch of 250 μm .
- The resolution ($< 100 \mu\text{m}$) is more than sufficient.
- Silicon detectors can sustain high doses of radiation.
- No gases are involved.

The main argument against the choice of Si-strip detectors is the cost. However, if the number of readout channels is reduced, by using a large strip size or by combining strips, the cost can be comparable to that of other technologies.

Detector layout

Single-sided Si-detectors, similar to those used in the Vertex Detector are proposed. The advantages over double-sided detectors are:

- Simplicity of design, with readout only on one side of the detector.
- Single-sided p -on- n detectors remain operational after sustaining significant radiation doses.[4]

The disadvantage is that the detector will be thicker overall.

Four layers per station are proposed: Y , U , V , and X , with strip orientation 90° , 85° , 95° , 90° . All layers are identical and orientation is provided by rotating the whole layer. Overlap is provided by shifting layers slightly with respect to each other. The stations nearest the RICH detectors need a fifth layer in the X direction (0°).

Each layer is formed from ladders of daisy-chained Si-detector modules. A layer of $60 \times 40 \text{ cm}^2$ is built up of 12 short modules, ($6 \times 7.5 \text{ cm}^2$ of silicon) and 22 long ladders of two modules each, shown in Fig. 8.6. The hole for the beam pipe varies in size according to the location of the station. The




Figure 8.6: Layout of Si-detector ladders in a layer. Dimensions are mm.

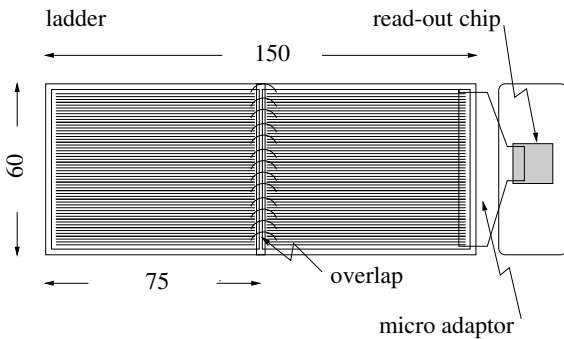


Figure 8.7: One ladder mounted on a carbon-fibre support.

smallest size is $x_{\min} = y_{\min} = 3.6$ cm. The layout of the ladders will differ from station to station, requiring a specific support structure for each station. The ladders themselves, however, are identical for all stations. They are mounted as “roof tiles”, in order to avoid dead space.

Detector design

The proposed Si-strip detectors are $300 \mu\text{m}$ thick, with a strip pitch of $125 \mu\text{m}$. Single-sided readout of every second strip is foreseen, so that the readout pitch is $250 \mu\text{m}$. The Si-modules will be glued onto carbon fibre supports and the strips bonded together in ladders as shown in Fig. 8.7. The readout chips are mounted on hybrids and glued onto the carbon-fibre support. They are bonded to the strips via pitch adaptors.

A stainless steel pipe with 2 mm diameter and $40 \mu\text{m}$ wall thickness is glued to the back of the assembly. Liquid at a temperature close to 0°C flows through this pipe and removes heat from the readout chips. It also helps to protect the Si-detectors from radiation damage. LHCb participates in a common evaluation of cooling techniques for LHC silicon detectors.

The amount of material a particle sees when it traverses a ladder corresponds to $0.51\% X_0$. Most of this ($0.28\% X_0$) is due to the silicon. At the location of the cooling pipe, the material amounts to $1\% X_0$. To avoid “hot spots”, the four layers which make up one station have their cooling pipes staggered. The overall material budget per station amounts to $2\% X_0$ on average.

Readout electronics

The readout electronics system used for the Vertex Detector will satisfy the requirements for the silicon detectors proposed here. At a readout pitch of $250 \mu\text{m}$, there will be 240 channels/module (2 readout chips), 34 modules/layer (= 8160 channels/layer) resulting in $\simeq 32640$ channels/station. Equipping all inner tracking stations with Si-detectors implies a total of about 350000 channels.

Signal to Noise ratio

The electronics noise in a silicon detector system depends on the properties of the pre-amplifier/shaper and of the silicon detector itself. Series noise is mainly determined by the capacity of the strip. The maximum strip length in the proposed detectors is 14 cm, corresponding to a capacity of 26 pF.

The noise is estimated for an amplifier-comparator chip with a rise time of 30 ns as used in Ref. [5]. The noise, expressed in electrons (ENC), as a function of the input capacitance parametrised [6] as

$$\text{ENC} = (730 + 31.5/\text{pF} \times C), \quad (8.1)$$

yields $\text{ENC} \simeq 1500$ electrons.

The MIP signal in a $300 \mu\text{m}$ thick silicon detector is about 25000 electrons, which implies a signal-to-noise ratio of 16 : 1. However, approximately 30% of the signal is lost to the capacitively coupled non-readout strip, which reduces the ratio to about 11:1. This is sufficient for good tracking.

Radiation hardness

The minimum inner radius of the Si-strip detectors in the Inner Tracker is 3.6 cm. At this radius, particle fluxes are about ten times lower than at the inner edge of the Vertex Detector, and decrease rapidly with distance from the beam. Since the high radiation dose is localised, the ladders could be recycled within one layer, i.e. by exchanging modules closest to the beam with those further away. We are confident that an inner tracking detector based on silicon microstrips can work. An R&D programme, together with the Vertex Detector group, will be carried out to determine an optimal solution taking into account physics performance, technology, cost, and the possible replacement and/or recycling of parts of the inner tracker after a few years of operation.

Future R&D

A prototype ladder will be produced, in order to test the mechanical aspects, the cooling aspects and the radiation hardness. Ideally, the prototype should be tested in an LHCb-like environment, e.g. near the HERA-B beam pipe. This development programme will be coordinated with that required for the silicon Vertex Detector modules.

8.2.3 Micro Cathode Strip Chambers

Micro Cathode Strip Chambers (MCSC's) measure analogue signals induced on narrow cathode strips, orthogonal to and separated by 0.6 mm from the anode wires. The geometry of the proposed device is shown in Fig. 8.8 and described in detail in Ref. [7]. The inner tracking stations are divided into several such MCSC modules.




Figure 8.8: Schematic of the MCSC chamber.

Precise coordinate measurement in a high particle flux requires a fast, low-noise front-end preamplifier/shaper. Prototype detectors have been tested in beams at CERN and with X-ray sources. The results [7] demonstrate that MCSC's can meet

the basic requirements for the inner tracker. A detailed simulation of the MCSC chamber, which includes the gas processes, noise and electronics response has shown good agreement [8] with the results obtained in the tests. An output signal with a peaking time of 20 ns and a full width of less than 50 ns was observed. The ADC must measure the collected charge with 8-bit precision. The design of the readout system [9] requires R&D effort to meet the requirements that are specific to LHCb.

An attractive feature of the MCSC option for the LHCb inner tracking system is that it could cover the area with less than half the number of readout channels, compared to the other options, however at the cost of higher occupancy.

8.2.4 Inner Tracker technology decision

All three inner tracking options have been implemented into the Monte Carlo simulation, which includes the appropriate material thickness, intrinsic spatial precision and occupancies. The results confirm that the measurement of the relevant physics parameters does not depend significantly on the choice of the technology.

The most important criteria for the choice will be reliability and cost. For each of the options, a programme of investigations will be defined with milestones. This programme will include studies of detector technologies, their implementation in the experiment, and also the readout electronics; in particular the development of chips compatible with LHCb requirements.

The inner tracker is a relatively small system which can be built and installed in less than two years. However, the time required to develop readout chips is much longer. The time scale for the inner tracker decision will need to be coordinated with the overall management of the front-end electronics.

In the cost estimate of the detector, silicon microstrip detectors have therefore been assumed for station T1, and MSGC's for the remainder. The particle flux is particularly high in the innermost part of the tracker in the stations situated in the magnet (see Fig. 8.4). Inner tracker stations of hybrid technology, i.e. MSGC's or MCSC's with a $10 \times 10 \text{ cm}^2$ silicon insert around the beam pipe are also considered for this region.

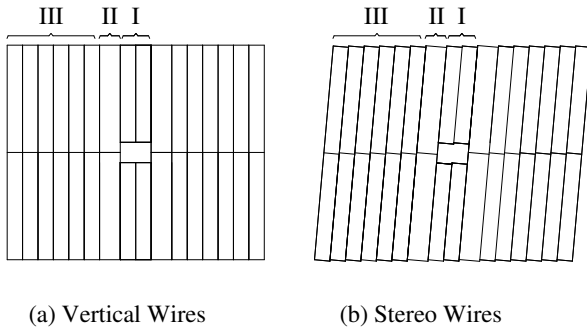


Figure 8.9: Modularity of detector planes. Module types I and II have 5 mm cells. Type III has 8 mm wide cells and is used in stations T7 – T11.

8.3 Outer tracking detector

8.3.1 Detector geometry

The outer detector consists of drift cells with straw-tube geometry. The configuration of each tracking station is summarized in Table 8.2. Since the emphasis is on tracking precision in the (x, z) bending plane, electron drift in most of the detector planes is in or near the x -direction. Stations T2 – T10 contain two planes with vertical sense wires, and two planes where the wire directions are tilted by $+5^\circ$ and -5° from the vertical (U and V in Table 8.2). In stations adjacent to a RICH detector, two additional planes with horizontal sense wires provide track directions with similar precision in the x - and y -projections. In station T11, just downstream of RICH 2, the $\pm 5^\circ$ stereo views have been omitted. They are not needed here because approximate space points provided by the nearby muon pad-chamber resolves stereo ambiguities.

Each plane listed in the table contains two staggered layers of drift cells. A “standard” station (T3 – T9) contains $4 \times 2 = 8$ cell layers in total. Figure 8.9 shows the schematic layout of a plane with (a) vertical or (b) stereo wires. A plane is assembled from modules with widths of 48 (type I), 64 (type II) or 32 (type III) cells per layer. Types I and II have 5 mm diameter cells, type III has 8 mm cells. Although a module of type II or III spans the full vertical extension of a station, the sense wires are split at $y = 0$. Thus, front-end electronics are mounted along the top and bottom edges of a station. Similarly, planes with horizontal sense wires carry front-end boards at the left and right edges. At a large distance from the beampipe, type III modules are split into two halves, not due to occupancy considerations but to reduce signal propagation time along the sense wire.

To keep the occupancy per cell below 10%, modules of type I are active only down to $|y| = 19$ cm, and modules of type II down to $|x| = 29$ cm. This leaves a rectangular hole around the beam-pipe, which is covered by the Inner Tracker. Simulation studies [10] show that the highest cell occupancy per interaction is 10%, occurring in the innermost cells of stations T2 and T11. One module of type II is necessary on either side of the beam pipe before switching to a larger cell size (type III, starting at $|x| = 68$ cm).

In addition to occupancy, signal latency is a major concern. The delay in signal arrival time at the preamplifier is the sum of drift time and propagation delay along the sense wire. The maximum propagation delay varies from about 4 ns in station T2 to as much as 17 ns in the horizontal wires of T11. For fast drift, a gas mixture containing CF_4 is needed. However, the possibility that detectors in or near the magnet, in stations T2 – T6, will be operated with a different mixture than stations T7 – T11 is being investigated. The drift times and trajectories have been studied with a detailed simulation program [11] which faithfully reproduces measurements [12] with honeycomb chambers. Tests of fast drift gases are in progress. We will also benefit from experience at HERA-B where a CF_4 -based mixture is used. Outside the magnetic field, the drift time in a 5 mm cell can be kept below 20 ns using $\text{CF}_4:\text{CH}_4$, $\text{Ar}:\text{CF}_4:\text{CO}_2$ or $\text{Ar}:\text{CF}_4:\text{CH}_4$ mixtures. In a field of 1.1 T, the maximum drift time in such a cell can be kept below 30 ns. However, the effective drift velocity $v_d \times \cos \alpha_L$, where α_L is the Lorentz angle, drops quite sharply at larger distances from the sense wire (the best result obtained with a 10 mm diameter cell is a maximum drift time of 85 ns). Therefore, stations T2 – T6 will have 5 mm diameter cells, even though occupancies would allow a larger cell size. In stations T7 – T11 the choice of cell size for the type III modules is 8 mm. With the HERA-B drift gas this will result in collection of all drift signals within 50 ns (two bunch-crossing intervals). The choice of a CF_4 -based drift gas entails some loss in precision due to electron attachment. The single-cell resolution is expected to be in the range 150–200 μm , which is adequate for momentum precision (see Section 8.5).

The fact that signal latency extends to two bunch-crossing intervals instead of one, has important consequences. The Level-1 track trigger must always consider two 25 ns time slices of pipelined data and track reconstruction will include signal overlap from interactions in consecutive beam crossings. At nominal LHCb luminosity

Table 8.2: Outer dimensions, arrangement of detection planes and numbers of readout channels for the Outer Tracking detector stations

station	z_{mean} (cm)	active area		planes (wire orient.)	channels	
		x_{max} (cm)	y_{max} (cm)		5 mm cells	8 mm cells
2	217	75	75	X X Y U V Y	6.0 k	
3	330	191	154	Y U V Y	10.2 k	
4	425	191	154	Y U V Y	10.2 k	
5	565	191	154	Y U V Y	10.2 k	
6	700	217	179	Y U V Y	11.6 k	
7	790	244	202	Y U V Y	3.6 k	6.3 k
8	840	260	214	Y U V Y	3.6 k	6.8 k
9	890	275	227	Y U V Y	3.6 k	7.4 k
10	940	291	240	Y U V Y X X	5.4 k	11.0 k
11	1165	360	297	Y Y X X	3.6 k	9.3 k
					68.0 k	40.8 k

there is about a 50 % chance that an accepted event contains hits from an interaction in either the preceding or the following bunch crossing.

The overall layout of the LHCb Outer Tracker is quite similar to that of the HERA-B experiment. The HERA-B Outer Tracker consists of 13 stations, at distances from about 2 m to 13 m from the target region. The largest station has a sensitive surface of $6\text{m} \times 4.5\text{m}$. Each station consists of the same three stereo views ($0, \pm 5^\circ$) as are used in LHCb, but there are no horizontal wires. The inner hole around the beampipe is $20\text{cm} \times 20\text{cm}$, and cell occupancies in HERA-B can be as high as 20%. However, at the 96 ns bunch interval of the HERA proton beam, overlap of signals from consecutive bunch crossings can be avoided.

8.3.2 Detector design

Detector modules will be constructed using a similar procedure as for HERA-B. A cell layer is produced from two preformed conductive foils. Instead of the true honeycomb structure employed by HERA-B, the LHCb design uses two staggered cell layers in a module. Vacuum forming of the cathode foils is currently being tested. However, should this prove unsatisfactory, the folding machine used for the HERA-B production is still available. Figure 8.10 shows cross sections of modules with 5 and 8 mm wide cells. A module provides its own rigidity, defines its own gas volume and has its own high voltage and front-end boards. The 5 mm cells have a circular cross section with two slightly flattened sides, the 8 mm wide cells are flattened to the same




Figure 8.10: Details of the end-pieces of module type I (top). Cross section of the border region between 5 mm and 8 mm cells (bottom): all dimensions are (mm)

depth as the 5 mm cells. The pitches are $5 + 1 = 6$ mm and $8 + 1 = 9$ mm. A flat conductive foil separates the two cell layers. This layout offers the following advantages with respect to a true honeycomb structure:

- It is easier and cheaper to build;
- There is no need for a hexagonal cell shape, which suffers from low drift fields near the sharp corners;
- In the proposed layout the fraction of active area per cell layer is $5/6$ (5 mm cells) or $8/9$

(8 mm cells) while it is only $2/3$ in a full honeycomb structure.

The choice of foil material is still open. Prototypes have used the same carbon-doped polycarbonate (trademark Pokalon) from which the HERA-B chambers have been built. The HERA-B chambers successfully withstood extensive aging tests with X-ray and beta sources, but stability problems have been observed in hadron beams. Independent of whether the problem is associated to the foil material, the same product will not be used for the LHCb outer detector, since it is no longer commercially available. Materials with volume conductivity (like Pokalon) or the use of a conductive surface layer are both under consideration for the cathode foil.

It follows from the cell geometry that the average hit multiplicity per track in a module is $5/3$ (5 mm cells) or $16/9$ (8 mm cells). For all realistic track angles, a particle passing through an insensitive region of one layer will hit a cell in the other layer in a region of almost 100 % efficiency. Efficiency near the cell edges depends not only on the cell shape, but also on the final choice of drift gas mixture (e.g. electron attachment to CF_4). Test results from prototype cells indicate that edge inefficiencies will be negligible compared to the geometrical loss factor, which is included in the LHCb detector simulation.

Module rigidity is provided by 10 mm-thick foam layers clad with $25 \mu\text{m}$ aluminium facings. Extensions of the inner facing are wrapped around the sides of the cell layers, thus providing the gas envelope. At 1 m intervals the sense wires are supported by isolating strips. The two sense wires of a cell, each covering half of its length, are terminated on a support strip placed halfway along the length of the module.

The geometrical thickness of a module is 30 mm (10 mm for the two cell layers and 10 mm foam on either side). This represents $0.4\% X_0$, assuming the same foil thickness ($75 \mu\text{m}$) as in prototype modules. There is no further material in the acceptance region, so a standard station of four detector planes represents 1.6 % of a radiation length. Consecutive planes will be mounted (almost) flush against each other, giving a station thickness in the acceptance region of slightly more than 120 mm (180 mm for stations T2 and T10). Outside the acceptance, the total extension in z for the mounting frame, front-end boards, gas distribution and cable paths is 180 mm (240 mm). The joint between adjacent modules in the same plane is 2 mm wide and has a thickness of $5\% X_0$. The terminating structure of

modules of type I extends 30 mm into the acceptance of the inner detector (see Fig. 8.10). Over 25 mm, it represents $1.5\% X_0$ per module, and about $8\% X_0$ over the remaining 5 mm closest to the beam pipe.

Stations will be mounted around the beam pipe by sliding left and right halves inwards, on a frame structure. The frame of the outer detector also serves as support for the inner detector, which is supported in front of the outer detector assembly on hollow carbon fibre rods.

8.3.3 Front-End electronics

The number of readout channels required for the Outer Tracking system is given in Table 8.2. It is intended to use the 8-channel ASD-BLR [13] preamplifier/discriminator chip developed by the ATLAS community. It provides ion-tail cancellation tailored to CF_4 -based drift gas mixtures, which tend to have very low ion mobility.

Given that the occupancy per interaction will be of order 5 - 10 % in an appreciable fraction of the outer detector channels, the TDC system will be driven by the bunch-crossing clock (rather than being data driven) and will have a Level-0 pipeline for each individual channel. Time measurement within the 25 ns bunch-crossing interval with 5 or 6-bit precision will be sufficient. An 8-channel TDC chip, currently under development at the University of Freiburg, could be adapted. Because of the 50 ns drift window, pipelined cells (following a Level-0 accept) must be fed pairwise to the Level-1 buffer and the track trigger electronics,

In stations inside the magnet coil (T3 - T5) only the preamplifiers will be mounted directly on the detector modules. In the other stations the front-end boards carry the TDC and pipeline chips.

8.3.4 R & D schedule

The design of the mechanics will be frozen by mid 1999; the detailed layout of the front-end electronics about one year later. The most important parameters — choice of drift gas, material and shaping method of the cathode foils — will be fixed before mid '99. A beam test of a telescope of two stations, one placed inside a 1 Tesla field and the other outside the field, is planned in 1998. The test stations will have smaller lateral extensions, but otherwise the geometry will be the same as presently foreseen for the final layout. In addition, a full scale model of a station will be built to test the design of the mounting structure, to verify the assembly procedure and to check the layout of

front-end boards, cables, gas lines, etc. Aging tests will be coordinated with the HERA-B outer tracker group. With a gas amplification of a few times 10^4 the accumulated charge in the hottest regions of the outer detector remains below 1 C/cm per year.

8.4 Alignment

It appears very difficult to assemble and to maintain the tracking stations, comprising the inner detectors and the large outer modules, with a precision below $100\ \mu\text{m}$. Instead, accurate displacement monitoring by means of an optical alignment system is foreseen. Two such systems are used in the ATLAS [14] muon spectrometer. Both systems are fully adequate for the alignment of the LHCb tracker and reach precisions below the $10\ \mu\text{m}$ level. One is developed by MPI and uses an infrared laser and semi-transparent amorphous silicon-strip sensors. The second system, RASNIK, consists of LED's with masks, lenses and CCD's. One device aligns three points with high precision in all three dimensions. More points can be aligned step-wise using several devices. The MPI system can align up to twelve points at a time, and is therefore the baseline option (as in HERA-B) for the alignment of the LHCb inner and outer tracking detectors.

8.5 Performance

The tracking system is designed to provide robust track-finding and track-following. This is obtained with redundant measurements in a relatively large number of layers and stations. The spatial precision is more than adequate and the momentum resolution is limited by multiple scattering. The relative precision, $\delta p/p = 0.3\%$, is approximately constant for tracks of momenta from 5 GeV to 200 GeV [1]. This propagates through to the precision in reconstruction of effective masses. Details are quantified (in Chapter 15) for each of the decay modes used in the physics analysis; examples are $B \rightarrow \pi\pi$ for which the r.m.s precision in mass is $15.2\ \text{MeV}/c^2$ and $D_s \rightarrow KK\pi$ (where the D_s is a B_s decay product), with a precision of $4.2\ \text{MeV}/c^2$. The contribution of the momentum measurement to the mass precision is comparable with that due to the error in the track angles as measured in the vertex detector. More details on the performance of the tracking system (including the Vertex Detector) are presented in Ref. [1].

References

- [1] R.v.d.Eijk , G.Gracia-Abril, M.Merk, O.Steinkamp and W.Ruckstuhl, "Track Reconstruction for LHCb", LHCb/98-045.
- [2] F. Sauli, Nucl. Instr. and Meth. A 386 (1996) 531-534.
- [3] F.Sauli, private communication.
- [4] E.Koffeman, private communication, (L. Andricek et al., in preparation, 1997)
- [5] E. Barberis et al., Nucl. Phys. **B** (Proc. suppl.) 32 (1993) 513; E. Barberis et al., IEEE Trans. Nucl. Sci. NS-40 (1993) 740.
- [6] E. Barberis et al., Nucl. Instr. and Meth. in Phys. Res. **A** 342 (1994) 90-95.
- [7] B.Bochin, et al, "Beam tests of the MCSC prototypes for the LHCb inner tracker", LHCb/98-008; O.Fedin, et al, "Study of MCSC with 36 KeV monochromatic X-rays", LHCb/98-010.
- [8] G.Velichko, "MCSC performance simulation", LHCb/98-016.
- [9] A.Kashchuk, et al, "Study of fast read-out electronics and charge interpolation for MCSC's", LHCb/98-011.
- [10] M.Merk, "Digitizations and Occupancies for the LHCb tracking detectors", LHCb/98-044.
- [11] R.Veenhof, GARFIELD, "A Drift Chamber Simulation Program", Version 5.05, CERN. S.Biagi, MAGBOLTZ, "A Gas Transport calculation Program", CERN.
- [12] RD5 collaboration, A.Bohrer et al, CERN-DRDC-91-53 (1992), CERN-DRDC-92-34 (1992).
- [13] M.Newcomer and A.Romaniouk, ATLAS Note INDET-122 (1996); B.Bevenue et al., ATLAS note INDET-80 (1996).
- [14] ATLAS Muon Spectrometer, Technical Design Report, CERN/LHCC/97-22.

Chapter 9 RICH Detectors

9.1 Introduction

The purpose of the RICH detectors is to identify the particle type of charged tracks.

Particle identification is crucial for the physics that LHCb is designed to study. It allows the decays of interest to be separated from other b-decay channels which would otherwise contaminate the signal. An example is $B_d^0 \rightarrow \pi^+\pi^-$ where the background from $B_d^0 \rightarrow K^+\pi^-$ gives an invariant-mass peak that significantly overlaps the signal peak, when the kaon is taken as a pion, and the peak corresponding to $B_s^0 \rightarrow K^+K^-$ lies right on top of the signal. Another example is $B_s^0 \rightarrow D_s^+K^-$, where the background from $B_s^0 \rightarrow D_s^-\pi^+$ would completely overwhelm the signal.

Even when the background does not peak underneath the signal, there are significant gains from particle identification in suppressing the combinatorial background, particularly for high-multiplicity decay channels.

The identification of kaons is crucial for one of the techniques used to determine the production state of a b hadron, by selecting a track from the accompanying \bar{b} decay (chosen as not being from the fully-reconstructed b hadron under study, but offset from the primary vertex of the event) [1]. If that track is a charged kaon produced by the cascade $\bar{b} \rightarrow \bar{c} \rightarrow \bar{s}$, then its charge will have the same sign as its parent \bar{b} quark.

Finally the particle-identification system can complement the calorimeters and muon system in the identification of electrons and muons. For high-mass particles it can also provide an improved momentum determination.

Acceptance requirements

The particle-identification system should cover the full angular acceptance of the LHCb spectrometer, from 10 mrad to 330 mrad.

The upper limit in momentum required for π -K separation is determined by tracks from $B_d^0 \rightarrow \pi^+\pi^-$ (or other two-body decays) as shown in Fig. 9.1 (a); 90% have $p < 150$ GeV/c. The identification of tagging kaons (and tracks from high mul-




Figure 9.1: Momentum distributions for (a) the highest momentum pion from $B_d^0 \rightarrow \pi^+\pi^-$ (unshaded) and $B_s^0 \rightarrow D_s^+\pi^+\pi^-\pi^-$ (shaded), (b) tagging kaons.

tiplicity decays) determines the lower momentum limit required. As shown in Fig. 9.1 (b), a limit of $p > 1$ GeV/c is sufficient.

Detector overview

The most suitable technique to cover the required momentum range of $1 < p < 150$ GeV/c is the detection of ring images of Cherenkov light produced by the passage of charged particles through various radiators in RICH detectors [2]. Three radiators have been selected: silica aerogel and two fluorocarbon gases, C_4F_{10} and CF_4 ; some details concerning these materials are given in Table 9.1. The radiators are incorporated in two RICH detectors, which are chosen to be in low magnetic field regions. This is so that the tracks do not curve appreciably whilst passing through the radiators, which would limit the resolution; low field is also important for the photodetectors.

The polar angle θ is plotted as a function of momentum for all tracks in fully-simulated b events in Fig. 9.2. As can be seen there is a clear correlation, with tracks at large angles having a softer momentum spectrum. The detector intended for the identification of low-momentum particles, RICH 1, is therefore designed to cover polar angles out to the full acceptance of the LHCb spectrometer, 330 mrad in both the horizontal and vertical projections. It is

Table 9.1: Some characteristics of the radiator materials used in the RICH system (for visible light at NTP); the lower part lists the contributions to the resolution (from emission-point, chromatic and pixel errors), the total resolution per photoelectron, the mean number of detected photoelectrons in the ring image, and the error on β per track.

Material		CF ₄	C ₄ F ₁₀	Aerogel
n		1.0005	1.0014	1.03
θ_c^{\max}	[mrad]	32	53	242
$p_{\text{thresh}}(\pi)$	[GeV/c]	4.4	2.6	0.6
$p_{\text{thresh}}(\text{K})$	[GeV/c]	15.6	9.3	2.0
$\sigma_\theta^{\text{emission}}$	[mrad]	0.21	0.54	0.36
$\sigma_\theta^{\text{chromatic}}$	[mrad]	0.22	0.54	1.21
$\sigma_\theta^{\text{pixel}}$	[mrad]	0.18	0.72	0.72
$\sigma_\theta^{\text{total}}$	[mrad]	0.35	1.10	1.45
N_{pe}		30	55	15
$\sigma_\beta^{\text{track}}$	/10 ⁻⁶	1.9	7.9	92.4

sited upstream of the magnet to catch low momentum tracks that will be swept out of the acceptance. It combines both aerogel and C₄F₁₀ radiators, with the gaseous radiator covering the acceptance down as close as possible to the beam pipe, which is conical with opening angle of 25 mrad in this region. The inner acceptance of the aerogel is limited to 50 mrad, which has been optimised taking into account the scattered light produced by high momentum, lower-angled tracks.

The detector for the identification of high momentum particles, RICH 2, is situated downstream of the magnet. Since the beam pipe has a smaller opening angle of 10 mrad in this region, the detector can have the best possible inner acceptance. The low-angle tracks are also better separated further from the interaction point, and the correspondingly longer focal length of the focussing mirror allows a higher-resolution device to be constructed, as required to maximise the momentum coverage. An outer acceptance of 120 mrad in the horizontal plane and 100 mrad in the vertical plane has been chosen, where the larger horizontal acceptance allows for the spreading of tracks by the dipole magnet. This accepts 90% of reconstructed pions from $B_d^0 \rightarrow \pi^+\pi^-$ decays with $p > 70$ GeV/c (i.e. beyond the limit for π -K separation in RICH 1).

9.2 Detector layout

The layout of the two RICH detectors can be seen in the overall view of the spectrometer, Fig. 2.5. An earlier design [3] had separate detectors for the




Figure 9.2: Polar angle versus momentum for all tracks in simulated $B_d^0 \rightarrow \pi^+\pi^-$ events.

aerogel and gaseous radiators, and used 45°-angled plane mirrors to bring the images out of the spectrometer acceptance, so that the material of the photodetectors did not compromise the performance of the other subdetectors. The same goal is now achieved by tilting the spherical focussing mirrors, but with a substantial increase in the photoelectron yield due to increased radiator lengths.

9.2.1 RICH 1

A schematic view of RICH 1 is shown in Fig. 9.3. A particle entering the detector from the interaction region first traverses a 5 cm thickness of silica aerogel, and then approximately 95 cm of C₄F₁₀ gas. This is the fluorocarbon with highest refractive index that remains gaseous at normal temperature and pressure, and is chosen for its low dispersion. The Cherenkov light produced by the radiators is focussed by two spherical mirrors, with 190 cm radius of curvature, one on each side of the beam pipe. They are tilted by 250 mrad to bring the image planes out of the acceptance of the spectrometer, with a rectangular array of photodetectors 60 cm × 100 cm on each side.

C₄F₁₀ resolution

Tilting the mirrors affects their focussing. When reconstructing the emission angle of a detected photon from a given track, an assumption must be made about the emission point: the mid-point along the track in the radiator is taken. As photons are emitted at other points in the radiator, the calculated Cherenkov emission angle is smeared, as shown in Fig. 9.4 (a).




Figure 9.3: Schematic layout of the RICH 1 detector (seen from above).

Another source of imperfect resolution is the dispersion in the refractive index of the radiator. For C_4F_{10} gas the index varies from 1.0013 at long wavelength, to 1.0015 at about 200 nm ($E_\gamma \sim 6$ eV). The wavelength spectrum of the detected photons is determined by the quantum efficiency of the photodetectors, and the assumed dependence on photon energy is shown in Fig. 9.5. The entrance window of the photodetector cuts off the energy spectrum; it is chosen to be UV glass giving $E_\gamma < 5.5$ eV, to optimise the resolution. The effect of this “chromatic error” is shown in Fig. 9.4(b).

Finally there is the effect of the finite granularity of the photodetector. This is chosen to be $2.5 \text{ mm} \times 2.5 \text{ mm}$ (to maintain high resolution whilst limiting the total number of channels) and gives the contribution shown in Fig. 9.4(c). Taking all three effects together, the RMS resolution per photoelectron for C_4F_{10} is 1.1 mrad.¹ This is for saturated rings, i.e. from particles with velocity $\beta \approx 1$. The Cherenkov angle θ_c is given by $\cos \theta_c = 1/\beta n$ for a radiator of refractive index n , and hence the uncertainty on n due to the dispersion leads to a term proportional to $1/\tan \theta_c$ in the resolution; this is accounted for in the pattern recognition studies.

The expected number of detected photoelectrons is given by [2]:

$$N_{\text{pe}} = \left(\frac{\alpha}{\hbar c} \right) L \epsilon_A \int Q \mathcal{R} \sin^2 \theta_c dE_\gamma, \quad (9.1)$$

where the first factor is a constant with value $370 \text{ eV}^{-1} \text{ cm}^{-1}$, $L \sim 95 \text{ cm}$ is the radiator length, $\epsilon_A = 0.73$ is the assumed coverage of the photodetector

¹The effect of imperfect knowledge of the track direction and momentum will be considered in Section 9.4.




Figure 9.4: Contributions to the resolution of the C_4F_{10} radiator: (a) emission-point uncertainty; (b) chromatic error; (c) pixel size; (d) overall resolution per detected photoelectron, with superimposed Gaussian fit.

active area, and $\mathcal{R} = 0.95$ is the assumed mirror reflectivity. Using the quantum efficiency Q from Fig. 9.5, Eq. (9.1) gives approximately 55 photoelectrons/track (for a saturated C_4F_{10} ring). Thus the resolution per track is about 0.15 mrad.

Aerogel

Silica aerogel is a colloidal form of quartz, solid but with extremely low density, that has a long-established use in threshold Cherenkov counters. Its refractive index can be chosen in the range 1.01–1.10, ideal for the identification of particles with momentum of a few GeV/c . The idea of using it in a ring-imaging detector [4, 3] has followed from the development of high quality, very clear, samples. The dominant cause of scattering of light within aerogel is Rayleigh scattering; this leads to the transmission of light with wavelength λ through a block of thickness L being proportional to e^{-CL/λ^4} , where C is a coefficient characterising the clarity of the sample. Recent samples have $C = 0.005\text{--}0.010 \mu\text{m}^4/\text{cm}$ [5].

The scattered photons can be at large angles to their initial direction,² and lead to background hits on the image plane. Since the Cherenkov light is produced uniformly along the path of the particle through the aerogel, the fraction of produced photons that emerge unscattered is $(1 - e^{-CL/\lambda^4})\lambda^4/CL$. The effect of scattering in the aerogel is shown by the shaded distribution in Fig. 9.5: it dominates at high energy, so a thin window (of Mylar) is placed after the aerogel to absorb the (mostly scattered)

²Single Rayleigh scattering gives a $\cos^2 \psi$ distribution for the scattering angle ψ , and this is modified somewhat by the possibility of subsequent scatterings.




Figure 9.5: Quantum efficiency as a function of incident photon energy assumed for the photodetectors in the simulation (unshaded), typical of a bialkali photocathode. When multiplied by the probability of the photon remaining unscattered by the aerogel, the shaded distribution results. The dashed lines indicate the assumed window cut-offs.

photons with $E_\gamma > 3.5$ eV. This also serves to reduce the chromatic aberration, and separate the aerogel from the gas. For a track passing through 5 cm of aerogel with $n = 1.03$ the resulting number of detected photoelectrons in a saturated ring image is expected to be approximately 15, from Eq. (9.1), with an additional 5 or so scattered over the detector plane.

The contributions to the resolution have been determined for the aerogel radiator in a similar manner to those for C_4F_{10} , and are listed in Table 9.1. The overall resolution per photoelectron is 1.4 mrad, which is reasonably well matched to the C_4F_{10} resolution, permitting the use of common photodetectors.

The radiation hardness of aerogel has been studied by a group from KEK [6]. No significant degradation of the properties were observed up to 10 Mrad equivalent dose, well beyond that expected for the aerogel in LHCb. Tests have shown that the properties will not degrade significantly over the timescale of the LHCb experiment if appropriate precautions are taken during its preparation, handling and storage [7]. The volume of aerogel required is anyway modest, ~ 30 l, so its replacement, if required, would be straightforward.

9.2.2 RICH 2

A schematic view of RICH 2 is given in Fig. 9.6. The radiator is approximately 180 cm of CF_4 gas, chosen for its suitable refractive index at normal temperature and pressure, and low dispersion. The spherical focussing mirrors have a radius of curva-




Figure 9.6: Schematic layout of the RICH 2 detector (seen from above).

ture of 820 cm, and are also tilted (by 370 mrad) to bring the image out of the acceptance of the spectrometer. In this case an additional flat mirror is required on each side, tilted by 240 mrad, to limit the extent of RICH 2 to 200 cm along the beam direction. The mirror positions and tilt angles have been optimised to give the best resolution for the detector, whilst keeping the flat mirror outside the RICH 2 acceptance to maximise the number of detected photoelectrons. The photodetectors are in a rectangular array of 72 cm \times 120 cm on each side, and the same granularity is assumed as for RICH 1.

The contributions to the resolution are listed in Table 9.1, with an overall resolution per photoelectron of 0.35 mrad. Approximately 30 detected photoelectrons are expected in a saturated ring.

9.3 Photodetectors

The RICH detectors each have two photodetector planes, with a total area of 2.9 m². Over this area it is necessary to detect single photons with the highest possible efficiency, with a detector granularity of 2.5 mm \times 2.5 mm. The photodetectors should be sensitive to light in the visible and ultraviolet: for the aerogel radiator visible-light sensitivity is essential as scattering dominates in the UV. The quantum efficiency that has been assumed for performance studies is shown in Fig. 9.5.

A fraction of 73% of the total area has been assumed to be active.³ With this assumption, the

³This corresponds to 90% in radius for a cylindrical device (i.e. 81% in area) with a further factor of 0.907 from hexagonal close-packing the tubes.

total number of channels is 340,000, so the cost per channel has to be low, $\mathcal{O}(10 \text{ CHF})$. The photodetectors must be fast enough for the time between bunch-crossings of 25 ns, and low noise from the readout electronics is important.

The fringe field from the spectrometer magnet does not exceed 40 gauss at the detector planes of RICH 1, but reaches 100 gauss at RICH 2, with a non-uniform field direction over the planes. The radiation dose predicted is a few krad/year at the RICH 1 photodetector planes, and less than 1 krad/year at RICH 2. The photodetectors should be able to survive for at least ten years in this environment.

Two detector technologies are candidates for meeting these requirements: the hybrid photodiode (HPD) and the multianode photomultiplier. The HPD [8] involves the use of a silicon pixel or pad detector inside a vacuum envelope (where “pixel” implies detector elements $< 1 \text{ mm}$ in size). A photocathode is deposited on a window in the envelope, and the photoelectron released by an incident photon is accelerated onto the silicon detector by an applied high voltage of $\sim 20 \text{ kV}$ (corresponding to about 5000 e^- released in the silicon).⁴ The multianode photomultiplier uses a dynode structure to amplify the photoelectrons, as in a traditional photomultiplier, but with the “metal-channel” technique many detection cells are incorporated in a single vacuum envelope.

Commercially-available examples exist of both technologies, but do not fully meet the requirements, particularly concerning the fraction of active area. A programme of R&D has therefore been pursued, focussing on the development of an HPD with high active area. This is most easily accomplished by increasing the size of the tube, since the inactive area is typically around the periphery, and is a lower fraction of the total as the tube size increases. For the fixed detector granularity this in turn leads to more channels per tube, $\mathcal{O}(10^3)$. This makes the provision of a feed-through per channel impractical, thus requiring electronics to be situated inside the vacuum envelope (also important for low noise). Two approaches have been pursued: the “Pixel HPD”, where a silicon pixel detector is bump-bonded to a readout electronics chip, and the “Pad HPD”, where a silicon pad detector is read out using routing lines on the silicon surface, wire-bonded to separate electronics chips.

⁴Radiation damage of the silicon due to the incident photoelectrons should not be a problem, as a 10% loss of gain has only been seen after an integrated charge of $\sim 0.1 \text{ C/mm}^2$ [8], corresponding to many years of operation in the RICH detectors of LHCb.




Figure 9.7: Photograph of a prototype 40 mm Pixel HPD tube.

9.3.1 Pixel HPD

The development of the Pixel HPD [9] is being carried out in close collaboration with industry.⁵ It is based on standard image-intensifier tube geometries that strongly focus photoelectrons onto a segmented silicon pixel array, bump-bonded to a binary readout chip with matching pixel electronics. The feasibility of this approach has already been demonstrated by the successful realization of the “ISPA-tube” in 1994 [10].

Cross-focussing geometries give good spatial resolution, and are robust with respect to external electric and magnetic field perturbations. Small pixels with bump-bond connections to the front-end electronics lead to small capacitances (giving low noise and high speed) and a compact anode structure. Binary electronics have low power consumption and are consequently well adapted to implementation in a vacuum tube; they are also compatible with the demanding bake-out cycles needed for high-quality photocathodes.

Current status

Two electrostatically-focussed prototype tubes with an active input diameter of 40 mm and an anode diameter of 11 mm have been constructed, as shown in Fig. 9.7. The LHC1 chip, which had been developed for tracking detectors in heavy ion experiments [11], is used for the anode. It permits measurements to be made with short peaking time (100 ns), close to the LHC requirement. However, the comparator threshold of 4000 e^- and its spread of 1000 e^- are too high for the HPD application. In addition the pixel size of $50 \mu\text{m} \times 500 \mu\text{m}$ is unnecessarily small, leading to charge sharing between pixels and difficulties with bump bonding.

⁵Delft Electronische Producten (DEP), the Netherlands.




Figure 9.8: Cross-section through the 40 mm prototype Pixel HPD tube.

The cross focussing has been simulated, as illustrated in Fig. 9.8; the spread of photoelectrons arriving at the silicon surface from a point source on the photocathode is around $40\ \mu\text{m}$. The expected demagnification and low distortion have been verified with the prototype tubes. No change has been seen in the performance of the pixel electronics after encapsulation. Preliminary measurements at an operating voltage of 18 kV show that 75% of the pixels are sensitive to photoelectrons. A further 8% exhibit too low a threshold and have been disabled, and the remaining 17% are insensitive due to high thresholds. Detailed studies are underway.

Future developments

The LHCb specifications require a timing resolution of 25 ns, and a threshold of $< 2000\ e^-$ with a spread of $< 200\ e^-$. Recent developments for other applications show that these can be reached [12]: a threshold of $1400\ e^-$ with an RMS of $90\ e^-$ (after 3-bit adjustment) has been achieved on a test chip with 25 ns shaping time. The chip has been demonstrated to be radiation tolerant up to 800 krad, well beyond the expected dose in the LHCb RICH detectors. An LHCb-specific chip will be designed, implementing these improved features, with $500\ \mu\text{m}$ square pixel dimensions and additional control logic related to the trigger requirements [13].

To increase the active-area fraction, larger tube prototypes are under development. Design studies indicate that the desired active-area fraction can be achieved. The present baseline dimensions of this tube are a 72 mm active photocathode diameter (80 mm total), and an 18 mm active silicon diameter (i.e. a linear demagnification by 4).

Magnetic field effects

The sensitivity of cross-focussed tube geometries to the expected stray field of the LHCb dipole magnet has been studied using a standard image inten-




Figure 9.9: Cross-section through the vacuum envelope of the Pad HPD.

sifier with an 18 mm active input diameter and a 7 mm image diameter. These tubes are found to be robust to low magnetic fields, because their electrode structure and material (Kovar) act as magnetic shields with high attenuation factors. A longitudinal magnetic field induces image rotation, with limited edge distortions ($< 4\%$ for 36 gauss). The effect of a transverse field of the same magnitude is a small image shift, with again limited edge distortions [14]. These effects can all be corrected offline.

9.3.2 Pad HPD

The Pad HPD [15] is housed in a cylindrical glass envelope of 127 mm diameter, capped with a spherical borosilicate UV-glass entrance window. A visible-light transmittive photocathode of K_2CsSb is vacuum-deposited on the inside surface of the window. The signal from single photoelectrons in the silicon sensor is detected by front-end electronic chips placed at the sensor circumference and wire-bonded to the routing lines from the pads. Focussing electrodes, fixed in the glass envelope, demagnify the “proximity focussed” image by a factor 2.3. Thus a $1\ \text{mm} \times 1\ \text{mm}$ silicon pad maps to a $2.3\ \text{mm}$ square pad on the photocathode surface, close to the RICH baseline requirement. The glass envelope is vacuum sealed to a stainless-steel base-plate with 40 vacuum feed-throughs which extract the multiplexed signals.

Vacuum envelope

The glass envelope with focussing electrodes is shown in Fig. 9.9. A Kovar metal skirt at the top of the cylinder is needed to decouple the thermal-expansion coefficients of normal and borosilicate glass, and in addition it facilitates the high-voltage connection to the photocathode. A Kovar metal skirt at the bottom is required for welding to the stainless-steel flange which incorporates a knife-edge, fitting the V-shaped groove in the base-plate.




Figure 9.10: Side view of the photocathode deposition apparatus.

This will be filled with degassed indium metal and closed (with up to 25 kN force) to make the vacuum seal. The envelope reduces in radius towards the bottom to prevent the flange from projecting beyond the radius of the window. The HV leads for the photocathode and focussing electrodes do protrude beyond this radius, but they are situated at 120° intervals and thus project into the gaps between the close-packed cylinders.

For the chosen electrode configuration, electrostatic simulations predict a unique mapping (onto the silicon sensor) of photoelectrons out to 60 mm radius. This corresponds to an overall active-area fraction of 81%. The UV-glass window transparency has been measured, and allows photon detection up to 5.2 eV (compared to 5.5 eV assumed in the baseline design). Two envelopes with fused quartz windows have been ordered, for coating with CsI transmittive photocathodes. They will allow the focussing to be studied, and will be used to produce the first vacuum-tight 2048-pad HPD.

A visible-light photocathode deposition apparatus has been constructed at CERN [16], as illustrated in Fig. 9.10, with a facility for making a high-vacuum seal of the base-plate onto the metal flange of the glass envelope. It is designed for a vacuum level of 10^{-10} mbar, and can be baked out at temperatures up to 300°C . Individual dispensers allow the deposition of Sb, Cs and K, and an optical system provides for online monitoring of the quantum efficiency.

Final assembly of the apparatus was performed at CERN in November 1997. After the first bake-out cycle a residual pressure of 10^{-9} mbar was




Figure 9.11: Layout of the pads of the 2048-pad silicon detector.

achieved. The indium sealing has been successfully tested [16], and testing of the evaporation system is now underway.

Silicon sensor

A 2048-pad silicon sensor has been developed, as shown in Fig. 9.11. It is divided into 16 sectors (each covering 22.5°). The wafer has a radius of 27.5 mm and thickness of $180\ \mu\text{m}$. The inner 25 mm region in radius contains the 1 mm pads and the outer 2.5 mm region is used to fan-in the 128 traces of a sector onto wire-bonding pads, for connection to the readout chip. The signals are AC-coupled, with a capacitor of 100 pF to the input of the readout amplifier and a resistor of $70\ \text{M}\Omega$ to ground.

This sensor has been successfully tested with photoelectrons using the VA3 readout chip [17], with $\sim 1\ \mu\text{s}$ peaking time. The measured pulse-height spectrum from a single pad is shown in Fig. 9.12, when the sensor is exposed to 20 keV photoelectrons in a pumped test setup [15]. The single-photoelectron signal is at $10\ \sigma$, where σ is the Gaussian width of the pedestal. This performance is worse than that measured for the previous 256-pad sensor [18], due to differences in the manufacturing of the silicon wafer [15]. A new production of 25 half-wafers and 25 full wafers is now in progress, aimed at recovering the better performance.

The silicon sensor is glued to the underside of a ceramic printed-circuit board, leaving the 25 mm sensitive radius of the sensor exposed to incident photoelectrons. The ceramic has four layers which contain all the routing lines for readout of the 16 chips through outputs connected to the feed-throughs of the base-plate.

The VA3 chip has been used in the first phase of R&D because it has low noise, is tested and reliable; however, it is too slow for operation at LHC. The SCT-128A is a fast (25 ns), low noise, radiation-hard 128-channel analogue chip, with a 128-element pipeline, developed for the ATLAS sil-




Figure 9.12: Pulse-height spectrum for a single pad of the 2048-pad silicon detector in the Pad HPD test setup. The peaks show the pedestal and resolved one and two photoelectron signals.

icon tracker [19]. Recently a noise level of $500 e^- + 50 e^-/\text{pF}$ has been measured in a prototype chip (with a low capacitance sensor). LHCb is participating in a new production of a slightly modified chip, aiming for a noise level of $\sim 600 e^-$ for the Pad HPD.

9.3.3 Readout electronics

To meet all of the LHCb requirements, further modifications will be required to the front-end electronics of both the Pixel and Pad HPD developments: in particular, the provision of a suitable derandomising buffer, and multiplexing in groups of 32 rather than 128 channels.

Data from the RICH detectors will not be used in the early levels of the LHCb trigger. The Level-0 trigger rate is 1 MHz, and the data will be stored in digital buffers situated a few metres from the photodetectors whilst the Level-1 trigger decision is being made. The data must therefore be read from the detectors within $1 \mu\text{s}$, which, at a rate of 40 MHz, requires multiplexing in groups of 32 channels. Given the number of channels, $\sim 10^4$ readout lines will be required, and the possibility of using twisted-pair copper cables is under investigation.

At the site of the Level-1 buffers, zero-suppression and feature extraction will also be performed, possibly using a DSP-based system. The requirements are very similar to those of the Vertex Detector (except that the data is not used in the Level-1 trigger decision) so a common solution is under study [13]. Further multiplexing will then reduce the number of 1 Gb/s readout links required to the DAQ system to 32. For the binary readout that is currently foreseen for the Pixel HPD, this number could be further reduced.

9.3.4 Multianode photomultipliers

Multianode photomultipliers have been adopted by the HERA-B collaboration to instrument their RICH detector. They use the 16-channel device R5900-M16 from Hamamatsu, with 4 mm cells. The recent development that makes these devices interesting for LHCb is a 64-channel device R5900-M64, with $2 \text{ mm} \times 2 \text{ mm}$ cells. The specified gain for these tubes is 10^6 at an operating voltage of 1000 V, which makes the signal processing straight-forward. There is a separate feed-through for each channel. Apart from the input stage (which requires much lower gain) the same electronics as for the Pad HPD could be used.

Assuming that the full photocathode area of $18.1 \text{ mm} \times 18.1 \text{ mm}$ is active, and staggering the tubes so that the effective outer dimension is 26.7 mm, this corresponds to an active fraction of 46%. About 4000 tubes would be required to cover the total RICH detector surface. To increase the active fraction, an optical lens system can be used in front of the tube, as adopted by HERA-B.⁶ A different lens arrangement is required for LHCb due to the large range of incident photon angles from the aerogel and gaseous radiators, but a possible solution has been found that would give an active-area coverage of close to 100% [20].

One of the 64-channel tubes was installed in the RICH prototype and tested in a beam in November 1997 [21]. The Cherenkov ring from air that filled the prototype vessel was arranged to fall within the active area of the tube, and was cleanly resolved. Using a Mylar filter to ensure the same spectrum, the photoelectron yield was measured to be 60% of that seen using a commercial HPD (described in Section 9.5). Most of this difference is ascribed to the lower collection efficiency of the multianode photomultiplier (the collection efficiency for the HPD is expected to be close to 100%).

9.3.5 Cost and schedule

The cost of most of the components of the Pad HPD are now established by solid commercial offers, with the major remaining uncertainty being the photocathode deposition. Allowing 5 kCHF for this, the total cost is 12 kCHF per HPD (6 CHF per channel). The cost of the Pixel HPD is expected to be ~ 10 CHF per channel.

The multianode photomultipliers are currently significantly more expensive than the estimates for the HPD's, but their price is expected to reduce

⁶An alternative approach using a multifaceted focussing mirror is also under study.




Figure 9.13: Event display of a simulated $B_d^0 \rightarrow \pi^+\pi^-$ event, with the photodetector planes of RICH 1 drawn side by side (scale in cm), and the Cherenkov rings superimposed.

to below 20 CHF per channel. At present they are considered to be a back-up solution.

The next step for the Pad HPD development is to test the focussing of the envelope with its electrodes in a pumped system, with a quartz window and CsI photocathode. Meanwhile the bialkali-photocathode deposition facility will be commissioned, and the parameters of the process tuned. Finally prototype tubes under vacuum will be produced with the bialkali photocathode, using the “in house” facility. The components for up to 25 tubes have been ordered, and some of these will be sent to commercial vendors to investigate the possibility of encapsulation by industry.

For the Pixel HPD, after the tests of the 40 mm tubes have been completed, the next step will be to produce an 80 mm prototype with a phosphor-screen anode, to study its photocathode and focussing properties.

Following the tests of prototypes of both HPD developments and the multianode photomultipliers in 1998, it is intended to make the final choice of photodetector in 1999, based on their demonstrated performance and price.

9.4 Pattern recognition

A simulated b-event in the two RICH detectors is shown in Fig. 9.13 and 9.14. The two detector planes of each RICH are drawn side by side, dots mark the positions of detected photoelectrons, and the expected ring images are superimposed. The event generator is PYTHIA, and a full GEANT treatment is used to provide track impact points on




Figure 9.14: Event display of the same event as Fig. 9.13, for RICH 2.

the entrance window of each RICH.⁷ A stand-alone routine is then used to simulate the generation of Cherenkov light, and trace it to the photodetectors.

The Cherenkov rings are not perfect circles, but are roughly elliptical in shape, with a degree of distortion that depends on the direction of the track within the acceptance. Instead of attempting to directly fit these rings, a substantial simplification is achieved by reconstructing the Cherenkov angles at emission (θ_c, ϕ_c) for each hit under the assumption that it originated from a given track [2]. That calculation accounts for the mirror geometry, and involves the solution of a quartic equation [22]. The hits which truly originate from that track will then all have the same value of polar Cherenkov angle θ_c (within the resolution), and have uniformly distributed azimuthal angle ϕ_c .

The task of the pattern recognition is to assign a particle type to each track, so as to best describe the observed hits. Two approaches have been developed: a “local” method which treats each track separately (and is therefore fast), and a “global” method that optimises the assignment of particle types for all tracks in the event simultaneously, to give the most accurate possible particle identification.

9.4.1 Local analysis

In the local method [23] each track is taken in turn, and the Cherenkov angle of each hit in the detector is calculated relative to that track; to save CPU time this is first calculated approximately to check that the hit is within the physically-accessible

⁷The effect of particles that are not reconstructed by the tracking system is under study, but has not been included in the results presented.

Table 9.2: Results from the local pattern recognition applied to 500 events. Each track gives one entry in the table, and X denotes tracks below threshold in all radiators; the rows give the reconstructed particle type, \mathcal{P} is the purity and ε the efficiency.

Rec	True particle type						\mathcal{P}
	e	μ	π	K	p	X	
e	6090	11	291	5	2	130	0.93
μ	47	200	512	7		136	0.22
π	16	18	12649	10		89	0.99
K	5	6	248	1028	11	60	0.76
p	2	3	104	11	381	2	0.76
X	90	12	432	51	33	3624	0.85
ε	0.97	0.80	0.89	0.92	0.89	0.90	

range of Cherenkov angles, before initiating the full calculation. For each track a log-likelihood function is calculated, proportional to:

$$\sum_i \ln \left(1 + \frac{1}{\sqrt{2\pi}\sigma_\theta\kappa} \exp \left[-\frac{(\theta_i - \theta_x)^2}{2\sigma_\theta^2} \right] \right), \quad (9.2)$$

where θ_i is the reconstructed emission angle of hit i , θ_x is the expected emission angle of the track under particle-type hypothesis $x = [e, \mu, \pi, K, p]$, and σ_θ is the angular resolution. κ is a hit selection parameter, which defines an effective range around the considered Cherenkov angle θ_i ; a value $\kappa = 1$ is found to give the best performance. The normalisation is chosen such that the value of the log-likelihood function corresponds to the number of hits expected at angle θ . The sum is performed over all hits, but excludes those with θ_i much greater than the saturated Cherenkov angle to save CPU time.

The log-likelihood function is used to calculate the number of hits which can be attributed to a given particle-type hypothesis, evaluated at the mean Cherenkov emission angle. A Poisson probability is then calculated from the comparison of the number of reconstructed hits with the number expected, from Eq. (9.1), and this is used to discriminate between the different particle hypotheses.

The performance of the algorithm has been tested using all of the tracks that pass through the RICH detectors in 500 simulated $B_d^0 \rightarrow \pi^+\pi^-$ events. The results are shown in Table 9.2. Each track gives a single entry: the column gives the true particle type (or X if the track is below threshold in all radiators) and the row gives the reconstructed particle type (or X if the track is reconstructed as being below threshold in all radiators). The performance can be quantified in terms of the efficiency ε (the fraction of true particles of a given type that

Table 9.3: Results from the global pattern recognition applied to 500 events (with the same notation as Table 9.2).

Rec	True particle type						\mathcal{P}
	e	μ	π	K	p	X	
e	6233	7	328				0.95
μ	8	224	554			31	0.27
π	5	10	13114	1		8	0.99
K		1	39	1083		11	0.96
p	1		4	1	427	1	0.98
X	3	8	197	27		3990	0.94
ε	0.99	0.90	0.92	0.97	1.00	0.99	

are identified correctly) and the purity \mathcal{P} (the fraction of tracks that have been identified as a given particle type that are truly that type). As can be seen, the efficiencies range between 80% and 97%, depending on particle type. The purities are also high, except for muons, which suffer from significant pion contamination due to the much larger number of pion tracks (close in mass to the muon).

The CPU time taken by this algorithm is 2.4 s/event on a machine of 100 MIPS processing power. Further optimisation of the speed is in progress, and the possibility of using it in the trigger, at Level-3, is under study.

9.4.2 Global analysis

Instead of treating each track separately, in the global method [24] the likelihood is constructed for the whole event. In this way the main “background” for a track in the local method, due to hits from other tracks, is correctly accounted for. For a given choice of particle type for each track, a likelihood is calculated that all the hits observed were produced by the tracks in the event. The particle-type assumptions are then changed and the likelihood recalculated; in this way the set of particle types that maximises the likelihood is searched for.

The event likelihood is calculated by comparing the number of photoelectrons detected in each pixel with the number expected in that pixel from all sources: signal (the Cherenkov rings from the various radiators), and backgrounds (from scattering in the aerogel, electronic noise, etc.). The number of detected photoelectrons is illustrated in Fig. 9.15 for the pixels in a zoomed region of the event in Fig. 9.13. A fitting function is calculated as the expected number of photoelectrons detected in each pixel, for a given choice of particle types for the tracks in the event. For the signal from a single track, that fitting function takes the form of a ring




Figure 9.15: Observed number of photoelectrons in each pixel, for a region of the RICH 1 detector.

with roughly Gaussian cross-section in radius.⁸ For all of the tracks in the event of Fig. 9.13, the fitting function over the same zoomed region has the form shown in Fig. 9.16 (for a given set of track hypotheses). The likelihood is then determined from comparison of the observed photoelectrons and the fitting function [24].

One advantage of this approach is that the detailed description of backgrounds is easily included. For example, the distribution of scattered photons from the aerogel, relative to the incident track direction, has been parametrized using the simulation. The resulting contribution has been included for each track in the fitting function.

The search for the maximum-likelihood solution is initiated with all tracks taken as pions (the most numerous particle type). The assumption for each track is then changed in turn to each of the other possible hypotheses, and the change which gives the largest increase in event likelihood is chosen. This procedure is then iterated until no further improvement in likelihood is seen. The CPU time taken is 11.5 s/event on a 100 MIPS machine.

Performance

The performance of the algorithm has been tested using all of the tracks that pass through the RICH detectors in 500 simulated $B_d^0 \rightarrow \pi^+\pi^-$ events. The results are shown in Table 9.3. The efficiencies and purities are now in the range 90–100%, except for the muon purity as before. However, instead of simply choosing the maximum-likelihood solution, the separation between different particle hypotheses can be varied. This is expressed in

⁸The parametrization is Gaussian as a function of the Cherenkov emission angle θ_c , and that is then converted to the detector plane using the RICH optics.




Figure 9.16: Expected number of photoelectrons in each pixel, for the same region as Fig. 9.15, under a given assumption of particle types for the tracks ($\times 10$ for the aerogel rings, for clarity).

terms of Gaussian sigma using the correspondence $N_\sigma = \sqrt{2 \Delta \ln \mathcal{L}}$, where $\Delta \ln \mathcal{L}$ is the difference in log-likelihood between the two hypotheses ($N_\sigma > 0$ for the nominal maximum-likelihood requirement). Adjusting the separation cut between pion and muon hypotheses to -2σ , almost all pions can be rejected whilst maintaining $\sim 85\%$ muon efficiency.

The separation between pion and kaon hypotheses is presented in Fig. 9.17 for true pions, as a function of their momentum. As can be seen, $> 3\sigma$ separation is achieved with the full pattern recognition, for almost all tracks. The separation is lowest at the two extremes in momentum, but the requirement of separation over the range $1 < p < 150$ GeV/c is satisfied.




Figure 9.17: π -K separation (in σ) as a function of momentum, for all true pions in 100 $B_d^0 \rightarrow \pi^+\pi^-$ events. The symbol indicates which detector(s) the pion passed through.




Figure 9.18: Layout of the $\frac{1}{4}$ -scale prototype of the RICH 1 detector.

The dependence of the particle-identification performance on the precision of the track parameters has been investigated [23, 24]. No significant degradation is seen as long as the momentum resolution satisfies $\Delta p/p < 0.01$, and the track angular resolution is better than 1 mrad in RICH 1 and 0.1 mrad in RICH 2. These requirements are comfortably satisfied by the tracking system of LHCb. The addition of noise in the photodetectors has been simulated, and little effect is seen until the probability of each pixel firing randomly is increased beyond 1%. This can be understood, because instead of searching for ring images in a stand-alone manner, the knowledge of the tracks has been used and the likelihood of different particle-type hypotheses compared; this comparison is robust to additional random hits.

9.5 Prototype tests

9.5.1 Aerogel beam test

A beam test of aerogel was performed at CERN in 1996 [25]. The apparatus used consisted of a light-tight box, flushed with nitrogen, containing an angled spherical mirror (of 90 cm radius of curvature) with a sample of aerogel supported in front; a one-inch photomultiplier mounted on a motorised stage was arranged so that it could be scanned horizontally across the focal plane of the mirror. This setup was exposed to 10 GeV pions from the T9 test beam of the CERN PS.

Scanning the photomultiplier across the image plane, clear peaks were observed in the count rate due to the Cherenkov light from the nitrogen gas




Figure 9.19: Photograph of the HPD array during assembly.

(at small radius) and from the aerogel (at a radius of 11 cm, corresponding to the nominal refractive index of 1.03). This was the first electronic detection of Cherenkov ring images from aerogel, and the observed photoelectron yield was consistent with expectation.

9.5.2 RICH prototype

A prototype of the RICH 1 detector has been constructed, and operated in the T9 test beam [26]. Focussed ring images from pions, kaons and protons have been detected on a plane of HPD's, using silica aerogel, air and C_4F_{10} gaseous radiators [27].

The aims of the prototype experimental programme were to investigate the performance of commercial HPD's, to demonstrate the simultaneous detection of Cherenkov rings from gaseous and aerogel radiators, and to measure the spatial precision of detected photons using a full-scale prototype of the RICH 1 device. In addition the scattering of light by the aerogel radiator was investigated using UV filters, and the effects due to ionising particles traversing the HPD's were studied.

A schematic diagram of the $\frac{1}{4}$ -scale prototype of the RICH 1 detector is shown in Fig. 9.18. A 240 mm focal-length mirror was used to simultaneously focus Cherenkov rings from aerogel and gaseous radiators onto an array of seven commercial HPD's at the photodetection plane. In this configuration the gaseous radiator (either air or C_4F_{10}) was contained within a volume of length 40 cm between a Mylar window and the mirror, and samples of aerogel were placed at the beam entrance window. Aerogel with nominal refractive index $n = 1.03$ was procured from KEK and Matsushita.⁹ The beam trajectory was measured using three planes of silicon detectors with 1.3 mm pads.

⁹Supplied by E. Nappi, Univ. of Bari, Italy.




Figure 9.20: Display of the hits in a run taken with a pion beam in the $\frac{1}{4}$ -scale prototype, showing the rings from aerogel and C_4F_{10} radiators.

Commercial HPD's

The 61-pixel HPD's were a new development contracted from DEP.¹⁰ They have a hexagonal array of 2 mm silicon pads, proximity focussed, with a fused-quartz window. Each HPD was connected to a front-end board, on which a 128-channel VA2 pre-amplifier chip (with 1.2 μ s shaping time) was mounted. The multiplexed analogue signals from each of the seven front-end boards were transmitted, via a distribution board, to a flash-ADC system in VME. The HPD array can be seen in Fig. 9.19.

The single-photoelectron signals from the HPD's were typically $3000 e^-$, at the 12 kV photocathode operating voltage. The complete readout and acquisition chain was initially tested using a weak source of light from a pulsed light-emitting diode. The signal/noise ratio was typically ~ 5.7 , corresponding to a noise of approximately $500 e^-$. The HPD's demonstrated very good pixel-to-pixel uniformity. They proved relatively trouble-free over three months of operation, between June and August 1997, during which 5 million triggers were recorded with over 96% of the channels fully operational.

Photon yields

Figure 9.20 shows data taken with a $10 \text{ GeV}/c \pi^-$ beam, with an 18 mm thickness of aerogel (KEK) and C_4F_{10} radiator. The C_4F_{10} ring is not quite contained within the central HPD; the radius is compatible with that expected from the refractive

¹⁰This followed from a development project supported by INFN Gruppo 2 and Univ. of Southampton under a European Union grant.

Table 9.4: Numbers of photoelectrons per event for the various radiator and detector filter combinations: N_{obs} is the observed number, N_{corr} is the number after background and acceptance correction and N_{pred} is the predicted number from the simulation. Aerogel sample 1 is from KEK and 2 from Matsushita.

Radiator	Filter	N_{obs}	N_{corr}	$N_{\text{corr}}/N_{\text{pred}}$
Air		4.9	4.8	0.99
C_4F_{10}		7.9	32.7	1.06
Aerogel 1		1.8	10.7	0.82
Aerogel 1	Glass	1.1	6.1	0.93
Aerogel 1	Mylar	1.3	3.5	0.75
Aerogel 2	Glass	0.9	4.3	1.12

index of 1.0014. The outer HPD's clearly exhibit the ring from the aerogel radiator. The number of photoelectrons per event has been measured for each of the radiators, and compared with expectations from a detailed simulation of the prototype geometry, incorporating measurements of optical transmission and reflection for each of the elements.

For each photodetector channel the width of the pedestal, the photoelectron peaks and the separation between peaks were determined by fitting multiple Gaussians to the spectra obtained from dedicated pedestal and LED runs. When analysing the data runs, a cut at 3σ of the pedestal width for each channel was used to select valid photon "hits". The number of photoelectrons per pixel per event was determined using the knowledge of the photoelectron-peak separation. Signal regions were defined as the area lying within three pixels of the observed ring arc for each photodetector.

The number of photoelectrons per event were counted for the gaseous radiators and the aerogel (with and without borosilicate glass or Mylar filters). The background to the raw photoelectron count was estimated by looking in those pixels not in the signal region. For the air and C_4F_{10} radiators the background correction is small ($\sim 5\%$), whereas for the aerogel samples it is $\sim 25\%$. The background consists of photoelectrons from scattered photons in the aerogel, a contribution from backscattered photoelectrons in the HPD's, and electronic noise. Efficiency corrections include losses due to the 3σ cut and the geometrical acceptance of the photodetectors. The results are shown in Table 9.4. The observed values are in excellent agreement with expectations for the gaseous radiators. For the aerogel the agreement is good except for the KEK sample with Mylar filter, where 25% less light is observed than expected; however, in this case the systematic error from the background

Table 9.5: Predicted Cherenkov-angle resolution contributions (in mrad) and comparison with the observed value (with and without Mylar filter).

Source	No Mylar	Mylar
Finite pixel size	0.56	0.56
Emission point uncertainty	0.58	0.58
Chromatic aberration	1.03	0.20
Gas pressure variations	0.02	0.02
Particle trajectory error	0.28	0.58
Total (predicted)	1.34	1.01
Total (observed)	1.40	1.10

subtraction is large.

From Eq. (9.1) a figure of merit can be defined as $N_0 = N_{pe}/L \sin^2 \theta_c$. The observed photon yield from the air and C_4F_{10} radiators correspond to a figure of merit $N_0 \sim 250 \text{ cm}^{-1}$. The photon yield from the aerogel samples (with Mylar filter) indicate a corresponding value for $N_0 \sim 50 \text{ cm}^{-1}$.

Cherenkov angle resolution

The performance of the RICH particle identification depends upon the precision with which the Cherenkov angle can be reconstructed using the photodetectors. This is limited by their finite pixel size, variations in gas pressure, knowledge of the particle's trajectory, the photon emission point and the effects of chromatic aberration. These contributions to the resolution have been studied using the simulation, and compared with the observed values determined from reconstruction of Cherenkov rings in a modified configuration of the prototype.

A full-scale RICH 1 prototype was constructed by adding extension tubes to increase the gaseous-radiator length to 100 cm and to accommodate a 1143 mm focal-length mirror. With the longer focal-length mirror, the rings from C_4F_{10} now span the outer six HPD's. Using a π^-/K^- beam of 15.5 GeV/c momentum, the pion and kaon rings were clearly resolved.

The rings are reconstructed in the detector plane by fitting an ellipse to the observed hits. The centre of the ellipse is measured to an event-by-event precision of 0.58 mrad and 0.28 mrad for the detectors with and without a Mylar filter placed in front of them, respectively. This is then used to give an estimate of the beam direction. The predicted resolution contributions are listed in Table 9.5, and their total is in good agreement with the observed value. When the silicon telescope is used to determine the beam trajectory, a resolution of 1.06 mrad is measured (with Mylar filter), close to the expectation of 0.96 mrad from simulation.




Figure 9.21: Front view of the mechanical support for RICH 2, with the two halves slightly separated; the region filled with radiator gas is shaded.

Future prototyping

The next step is to study RICH 2, for which a modification to the prototype is being prepared to extend the radiator length and use a longer focal-length mirror. Beam time has been reserved at the SPS test beam T1-X7, from April 1998, with pions of momentum up to 200 GeV/c available. It is also planned to use the RICH prototype system to test the various photodetectors as they become available.

9.6 Engineering design

Of the two RICH detectors, the engineering design of RICH 2 is most challenging, as it is much larger but requires greater angular precision than RICH 1 (and has a higher magnetic fringe field). The design studies have therefore started with RICH 2.¹¹

As far as possible, dense material such as the supporting structure is placed outside the acceptance of the spectrometer. The total volume of the gas enclosure is approximately 90 m^3 . It is proposed to reduce this volume to about 50 m^3 with the use of low-mass insertions (air bags) at the top and bottom of the gas enclosure. The proposed gas is CF_4 at atmospheric pressure and room temperature, with a total mass of about 200 kg. The pressure difference between the top and bottom of the gas enclosure is about 3 mbar and is limited to less than 1 mbar across the active volume.

¹¹Although the design of RICH 1 is expected to be simpler, it does have one complication: as seen in Fig. 9.13 a significant fraction of the photons produced by aerogel in one half of the detector are detected in the other half. For these photons to be detected the separating wall would need to be transparent.




Figure 9.22: Top view of the engineering design for one half of RICH 2.

Mechanical structure

The supporting mechanical structure is built around two vertical C-structures, as illustrated in Fig. 9.21. The main building block is a high I-profile stainless-steel beam, and further stabilisation of the structure is obtained by interconnecting stainless-steel sheets. Finite element calculations show that this structure is stable against sagging and twisting to better than $100\ \mu\text{m}$.

The two C-structures are not self-supporting. Before disconnecting the two halves, a precision stainless-steel beam is introduced into both, thereby closing the open end of the C and ensuring the mechanical integrity of the two halves. These beams are split in the horizontal plane and will be lodged above and below the supporting structure during data taking.

Low mass corrugated sandwich panels (G10-PMI) are currently being considered for the walls that define the gas volume. Inside the active area their thickness will be limited to about 0.5% of a radiation length. The walls between the two halves are proposed to be made from pre-formed 1 mm-thick G10 plates, and will incorporate half-conical shells that surround the beam pipe.

Photodetector housing

The photodetectors are placed at the two extremes of RICH 2, as shown in Fig. 9.22. They are housed in a gas-tight box which is separated from the Cherenkov gas by a thin quartz plate placed parallel to the photocathodes. This housing is surrounded by a shell of ferromagnetic material to provide shielding from the fringe field of the spectrometer magnet. Simulation of different arrangements for this shielding material indicates that the field seen by the photodetectors can be reduced below 10 gauss [28].

The base of the photodetector tubes are fixed to a high-precision base-plate, fixed in turn to a trolley which allows a positioning of the plate to within $100\ \mu\text{m}$. The positioning mechanism of the plate allows free access to each individual tube and its cabling. The design also makes it possible to change one tube without disturbing the alignment of the base-plate or any other tube. Cooling lines can be incorporated in the walls of the housing box.

Mirror support

The flat mirrors will be glued onto a supporting structure with low mass and high mechanical stability, such as a sandwich structure. These panels will span the full height of the detector, and will be suspended from the top with a freedom of rotation around the vertical axis of the panel. The alignment between the top and the bottom of the panels will be ensured with an optical measuring system. The precision should be better than 0.1 mrad.

The spherical mirrors will be glued onto similar panels which have been pre-formed to give a near-perfect spherical shape. The panels will be suspended from the top. These mirrors will have a fixed position. The aging effects on composite materials with CF_4 will be studied: the avoidance of swelling or shrinking will have an impact on the choice of the support-panel material.

The currently favoured substrate for the mirrors is a classic 6 mm thickness of glass. The use of composite materials or beryllium to give lower mass is a possibility, and a study is underway to verify if sufficient resolution can be achieved. Ultraviolet-enhanced aluminium coating with an overcoating of MgF_2 is proposed for the reflecting surface.

Alignment and calibration

A matrix of light spots can be projected from the centre of the gas volume onto the photodetector plane, via the mirrors, to monitor the alignment of

the system. The light source would be housed in the lower support structure, and a mirror placed near the beam axis would reflect the light onto the spherical mirror. The possible use of RASNIK or other alignment techniques [29] is under evaluation.

The pressure and temperature of the radiator gas will be recorded with a precision of respectively 0.1 mbar and 0.1 °C. Continuous monitoring of the refractive index will be performed using sonar for relative changes and a Fabry-Perot interferometer for the absolute measurement [30].

References

- [1] M. Tobar, “B flavour tagging with kaons”, LHCb/98-018.
- [2] T. Ypsilantis and J. Seguinot, Nucl. Instr. Meth. **A 343** (1994) 30.
- [3] LHCb Letter of Intent, “A dedicated LHC collider beauty experiment for precision measurements of CP violation”, CERN/LHCC 95-5.
- [4] D.E. Fields *et al.*, Nucl. Instr. Meth. **A 349** (1994) 431.
- [5] I. Adachi *et al.*, Nucl. Instr. Meth. **A 355** (1995) 390;
- [6] S.K. Sahu *et al.*, Nucl. Instr. Meth. **A 382** (1996) 441.
- [7] A. K. Gougas, D. Ilie, S. Ilie and V. Pojidaev, “Behavior of hydrophobic aerogel used as a Cherenkov medium”, EST/SM/98-1, submitted to Nucl. Instr. Meth. **A**.
- [8] R. De Salvo *et al.*, Nucl. Instr. Meth. **A 315** (1992) 375.
- [9] M. Campbell *et al.*, “Development of Pixel Hybrid Photon Detectors for the RICH counters of LHCb”, LHCb/98-035.
- [10] T. Gys *et al.*, Nucl. Instr. Meth. **A 355** (1995) 386.
- [11] E. Heijne *et al.*, Nucl. Instr. Meth. **A 383** (1996) 55.
- [12] W. Snoeys *et al.*, contributed paper to the IEEE Nucl. and Space Radiation Effects Conf., Newport Beach, July 1998.
- [13] J. Bibby and T. Gys, “Readout electronics for the RICH detectors”, LHCb/98-036.
- [14] T. Gys and D. Piedigrossi, “Performance of electrostatically-focussed image intensifier tubes in low magnetic fields”, LHCb/97-026.
- [15] E. Chesi *et al.*, “The 5-inch 2048-pad HPD development project”, LHCb/98-037.
- [16] A. Braem *et al.*, “An apparatus for the fabrication of large area hybrid photodiodes”, LHCb/98-007.
- [17] P. Weilhammer *et al.*, Nucl. Instr. Meth. **A 383** (1996) 89.
- [18] E. Chesi *et al.*, Nucl. Instr. Meth. **A 387** (1997) 122.
- [19] F. Anghinolfi *et al.*, IEEE Trans. Nucl. Sci. **44** (1997) 298.
- [20] R. Forty, “Use of lenses to increase the RICH photodetector coverage”, LHCb/98-038.
- [21] M. John *et al.*, “Cherenkov rings detected with a Multianode PMT”, LHCb/98-039.
- [22] R. Forty, Nucl. Instr. Meth. **A 384** (1996) 167; the term $4a(b^2 + c^2)R$ in Eq. (1) should read $4a^2(b^2 + c^2)$.
- [23] A. Schöning, “Particle identification in the RICH detectors and study of impact parameters”, LHCb/97-018.
- [24] R. Forty and O. Schneider, “RICH pattern recognition”, LHCb/98-040.
- [25] R. de Leo *et al.*, “Electronic detection of focussed Cherenkov rings from aerogel”, Nucl. Instr. Meth. **A 401** (1997) 187.
- [26] E. Albrecht *et al.*, “First observation of Cherenkov ring images using hybrid photodiodes”, CERN-EP/98-09, submitted to Nucl. Instr. Meth. **A**.
- [27] G. Wilkinson *et al.*, “Recent results from the LHCb ring-imaging Cherenkov test beam using hybrid photodiodes”, LHCb/98-041.
- [28] M. Alemi, “Passive magnetic shielding calculation for the photodetectors of RICH 2”, LHCb/98-017.
- [29] J.A. Paradiso, Nucl. Instr. Meth. **A 386** (1997) 409, and references therein.
- [30] M. L. Andrieux *et al.*, Nucl. Instr. Meth. **A 371** (1996) 259; A. Filippas *et al.*, Nucl. Instr. Meth. **A 371** (1996) 255.

Chapter 10 Calorimeters

10.1 Introduction

The main purpose of the LHCb calorimeters is to identify electrons and hadrons and to provide measurements of their energy and position. These measurements are required for the various trigger algorithms and for the off-line event analysis. Furthermore, the electromagnetic calorimeter should enable the reconstruction of photons and π^0 mesons. The selectivity demanded at the first level of triggering implies longitudinal segmentation of the electromagnetic shower detection. The structure chosen consists of three elements: a single-layer Preshower detector followed by a Shashlik electromagnetic calorimeter (ECAL), and a scintillating tile hadron calorimeter (HCAL).

Geometric matching of acceptance and of lateral detector segmentation facilitates the combination of the information from corresponding cells of the three detectors in the trigger processors.

The polar-angle acceptance starts at 30 mrad from the beam axis. The acceptance improvement which could be gained at smaller angles does not justify the additional cost and complications of the radiation-hard detection techniques that would be required. The outer limits of polar-angle acceptance are 300 mrad horizontally and 250 mrad vertically. The dynamic range of the readout extends up to 200 GeV in the electromagnetic calorimeter, and 300 GeV in the hadron calorimeter.

The three elements of calorimetry employ similar technologies, i.e. scintillators coupled to wavelength-shifting fibres read out by fast photodetectors. The front-end electronics will be largely unified for the whole calorimeter system.

10.2 Preshower detector

10.2.1 Performance requirements

Identification of electrons and photons at the trigger level is extremely important in LHCb. By selecting clusters exceeding an E_T threshold the EM single-particle trigger should provide a Level-0 rejection factor of at least 30 for minimum bias

events. Good photon identification is necessary for detection of radiative B decays and for reconstructing π^0 mesons.

It was found [1] that pion rejection based on leakage into the HCAL has inherent limitations in regions of high track density: in particular the energy lost by hadrons in the HCAL near to an electron can increase the probability of mistaking the electron for a pion. As a consequence, pion rejection using only the information given by the electromagnetic and hadron calorimeters is not sufficient, and a Preshower detector is needed to provide better identification of electromagnetic particles. The transverse segmentation of the Preshower detector will match that of the ECAL. A finer subdivision will not substantially improve the detector performance, and would add considerably to its complication and cost.

A Preshower structure made of lead and scintillator, respectively 14 mm ($2.5 X_0$) and 10 mm thick, has been chosen. The optimum pulse-height cut in the scintillator corresponds to 9 MIPs, yielding the performance characteristics listed in Table 10.1 [2]

Table 10.1: Single-particle Preshower performance for a cut of 9 MIPs. ε_e (ε_π) is the fraction of true electrons (pions) which pass this cut.

Energy	ε_e	ε_π	Enrichment ($\varepsilon_e/\varepsilon_\pi$)
5 GeV	0.904	0.099	9.13
10 GeV	0.975	0.073	13.36
50 GeV	0.998	0.081	12.32

10.2.2 Detector layout

The Preshower detector is located immediately upstream of the ECAL, with one-to-one correspondence between ECAL towers and Preshower cells. The Preshower is mechanically supported by the ECAL frame, the detectors being arranged in a series of drawers which can slide horizontally on rails

on the left and right sides of the frame, allowing easy access to the scintillator and readout elements.

Each cell is made of a 14 mm-thick lead plate followed by a square scintillator, 10 mm thick. The tile has the same transverse dimensions as the corresponding ECAL tower. i.e. a square having 4, 8 or 16 cm sides. The scintillator will be chosen to have a high light yield and fast response. Bicron BC408 is adequate, but cheaper injection-moulded styrenic scintillators or acrylic scintillators could also be suitable. The light will be extracted from the scintillator using wavelength-shifting (WLS) fibres. This method allows the photodetectors and the readout electronics to be sited away from the beam, thus reducing the radiation dose, if necessary. Several methods are available to couple the fibres to the scintillator with good efficiency, and two possibilities are shown in Fig. 10.1:

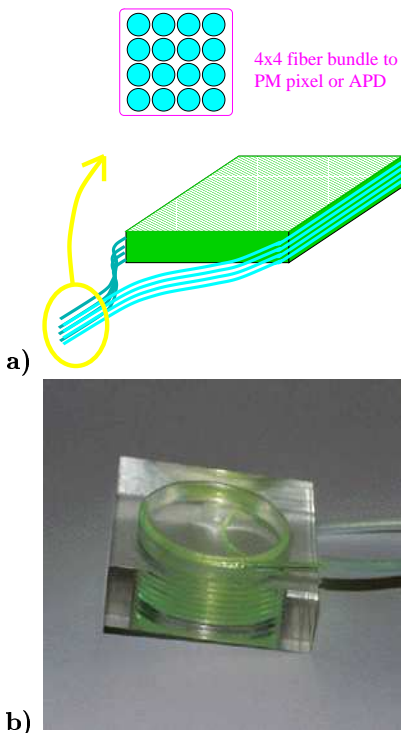


Figure 10.1: Two possible methods of coupling the scintillator tiles to the WLS fibres (see text).

Scheme (a) is similar to that employed in the hadron calorimeter and has given good results. Tests show [3] that a MIP produces more than 3 photoelectrons (p.e.) on the photomultiplier (using 1 mm diameter Kuraray WLS fibres coupled to a 1 mm BC408 scintillator plate). Scaling to 10 mm scintillator thickness, 30 p.e. will be available (the

value can be almost doubled by making one end of the fibres reflecting.) In this case a bundle of 16 fibres will be directly coupled to the photodetector.

Scheme (b) is known as “spiral” coupling. A deep groove is machined inside the tile, and the WLS fibre is coiled inside it. The amount of light is similar to that measured with the first setup. The advantage is that only one fibre is used, thus simplifying the coupling to a small-area photodetector. The method is well suited to small tiles, where the total length of the coiled fibre is short and problems of transit time and attenuation are reduced.

Detailed tests will be conducted in order to study the uniformity of light response over the tile surface, and the effects on the overall detector performance.

10.2.3 Readout

The baseline solution for reading out the fibres uses multi-anode photomultiplier tubes (PMT’s). The 16-pixel Hamamatsu H6568 phototube (pixel size $4 \times 4 \text{ mm}^2$, pitch 4.5 mm) is well suited for coupling the 16-fibre bundle of a detector channel to a single pixel. It is compact, has low cross-talk and sufficient gain, uniformity and linearity. The readout electronics will be the same as for the calorimeters, but an ADC with only 8 bits can be used resulting in a slight cost reduction of the chain.

The electronics can be placed close to the PMT. To reduce the radiation dose the readout system (including the PMT) will be mounted on the upstream side of the radiator. If necessary, the PMT’s and the electronics could be placed as far from the beam axis as necessary by making the fibres longer, with a modest reduction in light output (see Fig. 10.2).

10.2.4 Calibration

For calibration it is foreseen to use (i) a radioactive source, positioned in front of the tiles, (ii) an optical system for light injection into the tiles and (iii) a system for light injection in the PMT pixels for measuring the PMT gain. In addition the electronics chain will have (iv) a test input to be used for gain monitoring. Parts of the calibration system for the ECAL and the Preshower detector will be common.

10.2.5 APD readout option

Avalanche Photo Diodes (APD’s) might become an interesting alternative to the PMT readout.




Figure 10.2: Light yield measured with different WLS fibres as a function of fibre length (1 mm BC408 scintillator).

These would have the advantage of compactness and higher quantum efficiency (Q.E.) than PMT's, and potentially could be more economical.

Manufacturers such as EG&G (Canada) and Hamamatsu commercially produce large area APD's (up to 100 mm²), with typical gains of 50 and more. A new manufacturer (ITE, Poland) is now producing good quality APD's that have been tested successfully in a beam test.

CMS [4, 5] is executing an intensive R&D programme on APD's in collaboration with Hamamatsu to improve characteristics and to reduce costs below the \$60/unit barrier, and LHCb can benefit from their experience.

For use in the Preshower, the main problem with APD's is the small light output from the tile. Taking the figures quoted above, and assuming the diode has a 60 % Q.E. in the blue region, 150 p.e. per MIP could be obtained. With a gain of 50 the signal charge would be 7500 electrons, with a detector capacitance of typically 30 – 300 pF, thus posing a serious challenge to the front-end electronics (speed vs. noise.) Therefore a larger light yield would be desirable. A collaboration with a scintillator firm (Pol.Hi.Tech) is being setup, in order to develop scintillators and/or WLS fibres with an improved matching to the APD spectral response which is peaked in the red.

10.2.6 Beam tests

First tests of a Preshower prototype using APD readout have been conducted using the X7 Test Beam in November 1997. A Preshower cell was built using a 10 mm lead radiator followed by a 10

mm thick scintillator plate (4 cm side) with “spiral” readout (see Fig. 10.1b) The 1 mm WLS fibre (Kuraray Y11) was coupled to an ITE APD. This APD had a sensitive area of 7 mm², a Q.E. \approx 70% at 500 nm, a capacitance of 40 pF and operated at a gain of 200-400. The diode was connected to a preamplifier [6] with 11 ns peaking time. The Preshower prototype was mounted upstream of a Shashlik calorimeter module. Pulse-height spectra for 20 GeV pions and electrons are shown in Fig. 10.3 and demonstrate good separation. The efficiency for electrons over the tile surface was measured to be uniform, even in the region of the groove. The results are encouraging and more detailed tests will be conducted.




Figure 10.3: Signals from the APD Preshower prototype for 20 GeV electrons and pions.

10.3 Electromagnetic calorimeter

10.3.1 Performance requirements

The main aims of the electromagnetic calorimeter are:

- To provide electron identification at different trigger levels and at the offline reconstruction stage,
- To measure the energy of electrons and photons,
- To provide for π^0 reconstruction.

These aims define the following set of requirements for the performance of the ECAL:

Acceptance:

The z -position of the upstream face of the ECAL is 12.4 m from the interaction point. The outer dimensions are matched projectively to those of the tracking system, thus providing the polar-angle coverage of $\theta_x < 300$ mrad and $\theta_y < 250$ mrad. Due to severe radiation levels close to the beam pipe the active calorimeter volume has a central square hole of 64×64 cm².

Response time:

To simplify the implementation of the Level-0 trigger the signal collection time should be below 25 ns.

Transverse granularity:

The rates of photons and charged particles at the front face of the ECAL as a function of the distance from the beam pipe are shown in Fig. 10.4. The hit density varies over the calorimeter surface by two orders of magnitude. To minimize the number of readout channels the calorimeter should have a variable granularity. It is composed of inner, middle and outer regions with transverse segmentation adjusted to ensure at most a 5% charged-particle occupancy per calorimeter cell. The fine transverse segmentation is essential to improve the efficiency of π^0 reconstruction and the discrimination between electrons and charged hadrons overlapping with photons from π^0 decays.

Energy resolution:

The requirement on the energy resolution of the ECAL is rather modest. A resolution of

$$\frac{\sigma(E)}{E} = \frac{10\%}{\sqrt{E}} \oplus 1.5\%$$

(E in GeV) is sufficient to provide electron/hadron separation at a level of about 100. This energy resolution is also adequate for π^0 reconstruction. Figure. 10.5 shows the signal-to-background ratio for the π^0 signal reconstructed in minimum bias events, which is close to unity.

Radiation resistance:

The expected radiation dose strongly depends on the distance from the beam pipe, as shown in Fig. 10.6. For the innermost part of the ECAL peak doses of up to 0.4 Mrad can be reached after 1 year of LHCb operation. The longitudinal dose distribution is shown in Fig. 10.7. Since a relatively small part of the calorimeter in the vicinity of the




Figure 10.4: The number of particles entering the calorimeter per cm² per bunch crossing; (A) vs. x and (B) vs. y . Electrons and photons have energies above 1 MeV, hadrons have energies above 10 MeV.

shower maximum position suffers the maximum radiation, no significant deterioration of the performance is expected during 10 years of LHCb operation. Modules of “shashlik” calorimeter, made of injection moulded polystyrene based scintillator and Y-11 wavelength shifter (WLS) fibres, have shown only minor degradation of the energy resolution after irradiation up to 5 Mrad [7, 8, 9].

10.3.2 Design considerations

The electromagnetic calorimeter employs the “shashlik” technology of a sampling scintillator/lead structure read out by plastic WLS fibres. This technology has been successfully developed by the HERA-B collaboration [10, 11] at DESY, the PHENIX collaboration at BNL [12] and the RD36 collaboration [13] at CERN. The combination of good energy resolution, fast response time and relatively low cost-to-performance ratio fulfil the requirements of the LHCb electromagnetic calorimeter.

The transverse view of the ECAL is shown in Fig. 10.8. The calorimeter is constructed from supermodules, shown in the figure as rectangular




Figure 10.5: The invariant mass of two photons in minimum bias events. The minimum photon energy is 2 GeV.

blocks with transverse dimensions of $64 \times 32 \text{ cm}^2$. Each supermodule is a 4×2 matrix of ECAL modules which constitute the basic construction units of the calorimeter. To provide a variable transverse granularity the calorimeter is subdivided into three regions. The inner, middle and outer modules comprise 16, 4 and 1 calorimeter cell respectively. The cell parameters, as well as the number of read-out channels, are listed in Table 10.2 separately for each region of the calorimeter. The super-modular geometry of the electromagnetic calorimeter facilitates its integration into the LHCb detector. The supermodules are connected mechanically at the front and backward faces and form a self-supporting structure. The structure is rigid enough to allow the left and right calorimeter halves to move sideways on rails in order to provide access for servicing. In total, there are 258 supermodules.

10.3.3 Construction features

The structure of a calorimeter module is shown in Fig. 10.9. Inner, middle and outer modules have a similar design differing by the number of calorimeter cells per module and by the WLS




Figure 10.6: The expected annual radiation dose. The vertical line is the inner edge of the ECAL.

fibre density per scintillator tile. The module is constructed from alternating layers of 2 mm-thick lead, white reflecting paper (TYVEK) and 4 mm-thick injection moulded polystyrene-based scintillator tiles. The complete stack is compressed using stainless steel side covers, pre-tensioned and welded to front and rear steel lids. In total, there are 70 layers resulting in a total depth of $25 X_0$. The transverse granularity is achieved by assembling each module layer using 16, 4 or 1 identical scintillator tiles with aluminized edges to improve light collection and prevent tile-to-tile light crosstalk. The light is read out via 1 mm diameter WLS fibres which penetrate the entire length of the module. The fibres belonging to each calorimeter cell are bundled at the end of the module and polished. A 1 mm air gap separates the surface of the WLS bundle and the phototube. The density of the fibres is a subject for optimization between light yield, uniformity of response and cost.

Readout of the scintillation light:

The operational features required from the phototubes are similar for the hadronic and electromagnetic calorimeters. The criteria governing the choice of phototube are discussed in




Figure 10.7: The longitudinal dose distribution for the inner region of the ECAL.




Figure 10.9: Plan view of a middle-type ECAL supermodule with its Preshower elements.




Figure 10.8: Transverse view of the ECAL. Dimensions are mm.

section 10.5.

10.3.4 Calibration and monitoring

In-situ calibration constants for ECAL will be obtained from experimental data using isolated particles with momenta precisely measured in the tracking system. A sufficiently large data sample ($\sim 10^3$ electrons and $\sim 10^5$ hadrons) can be collected for the calibration of each ECAL cell within a few hours of LHCb running. A complementary calibration can be performed using a sample of π^0 mesons where one of the photons converts upstream of the magnet.

The stability-monitoring system for the large number of phototubes must correct for time-dependent gain variations. Two monitoring techniques, based on the use of either multiple LED assemblies or a high intensity nitrogen laser, are under study. Monitoring systems using pulsed light sources have demonstrated better than a few percent stability. The choice will be made following further R&D.

Table 10.2: Parameters of the LHCb ECAL.

Type	Shashlik		
Absorber	Lead (2 mm plates)		
Sampling	Scintillator (4 mm plates)		
No. of layers	70		
Moliere radius	3.6 cm		
Radiation length	1.64 cm		
Depth	43.8 cm ($25 X_0$)		
Section	Inner	Middle	Outer
Cell size (cm)	4×4	8×8	16×16
No. of channels	3072	1344	1536
Height (cm)	320	512	832
Width (cm)	192	384	640
Weight (tonnes)	9.2	16.2	55.7
Scint. weight (kg)	1420	2486	13,442
WLS fibre type	Kuraray	BCF	BCF
	Y-11	91a	91a
Fibre length (km)	48.4	54.2	139.4
Photomultiplier	R5600	R5600	R5900

10.4 Hadron calorimeter

10.4.1 Performance requirements

The purpose of the hadron calorimeter (HCAL) is to provide data for the trigger and to assist in background suppression when reconstructing B decays.

At the trigger level the HCAL must:

- provide a single-particle transverse-energy measurement of high p_T hadrons,
- improve electron and hadron separation especially for high-energy particles.

In addition, at the reconstruction stage, the HCAL will improve electron and muon identification.

The performance parameters expected from the HCAL design are summarized in Table 10.3.

10.4.2 HCAL design

A schematic view of the HCAL is shown in Fig. 10.10.

Like the ECAL it consists of three regions with different granularity. The lateral cell segmentation has to satisfy the following criteria:

- it must measure the shower position, with a small contribution to the p_T measurement error of a single hadron,
- it must allow good two-hadron spatial separation,

Table 10.3: HCAL design parameters.

Calorimeter feature	Value
Inner acceptance, ($x \times y$)	32×27 mrad
Outer acceptance, ($x \times y$)	300×250 mrad
Energy range	$0.1 \div 300$ GeV
Time resolution	< 25 ns
Shower separation	15 mrad
Radiation resistance	3 kGy/yr
Depth	$7.3 \lambda_I$
Energy resolution, $\sigma(E)/E$	$0.8/\sqrt{E} \oplus 0.05$




Figure 10.10: Front view of the Hadron Calorimeter.

- The number of channels should be kept low to minimize the cost.

A four to one correspondence in the segmentation between the HCAL and ECAL will be maintained throughout the detector. Hence the cell sizes range from 8×8 cm² at the centre, through 16×16 , to 32×32 cm² at the periphery.

The technology consists of a scintillator/iron sampling structure with the scintillating tiles parallel to the beam axis and with passive radiator made of steel plates. The overall dimensions of the active surface are 8.96 m \times 7.04 m. The beam pipe passes through a 0.64 m \times 0.64 m square hole at the centre. Shielding material is placed between the beam pipe and the inner edge of the HCAL.

The scintillation light is transferred by the WLS fibres to the rear of the calorimeter where it is detected by photo-multipliers. This technology, already tested in the ATLAS TileCal [14] design, has several advantages:

- the calorimeter has the ‘natural’ depth of passive material between tiles, of the order λ_I along the direction of the hadron shower development,
- in the transverse direction, the spacing be-

tween tiles is of the order of one radiation length, X_0 , suitable for electromagnetic cascade development,

- the fibres have a large contact surface with the tiles providing efficient light collection,
- the scintillator is produced by injection-moulding techniques and is cheap,
- the detector can be segmented both transversally and longitudinally,
- the detector is easily assembled and has good mechanical properties.

The design parameters of the Hadron Calorimeter are summarized in Table 10.4.

Table 10.4: Hadron Calorimeter: summary of features.

Active area m^2	8.96×7.04
Average thickness	$68 X_0, \quad 7.3 \lambda_I$
Cell size, cm^3	$8 \times 8 \times 153$ $16 \times 16 \times 153$ $32 \times 32 \times 137$
Sampling structure	6 mm steel plate 4 mm steel spacer 3 mm scintillator
Total weight, tonnes	597
Segmentation	2-fold in depth
Number of cells:	
— with $8 \times 8 cm^2$	512 cells
— with $16 \times 16 cm^2$	336 cells
— with $32 \times 32 cm^2$	736 cells
Total number of channels	3144
Photomultiplier	R5900 Hamamatsu

10.4.3 Calibration and monitoring

The hadron calorimeter is required to supply data for the Level-0 trigger. Its stability and performance must therefore be continuously checked and maintained. Three different types of calibration are foreseen:

- a ^{137}Cs radioactive source is driven hydraulically inside the calorimeter through a 8 mm diameter stainless steel tube,
- LED/laser pulses monitor the stability of the PMT-electronics chain,

- test-charge injection calibrates the ADC's and checks the stability of the trigger logic.

The calibration system has to maintain the overall stability of the calorimeter response within a tolerance of a few percent, compatible with the resolution of the detector at the high end of the energy range. The principles of design and performance of the calibration system were tested during the prototype construction and in beam tests. They are described in the following section.

10.4.4 Prototype beam test

During 1997, three prototype modules of dimensions $16(H) \times 96(W) \times 153(D) cm^3$ were constructed and tested. The scintillating tiles, made of PSM 115 with 1.5% PTP and 0.015% POPOP were injection moulded. The Pol.Hi.Tech. type S250 WLS fibres, aluminised at the upstream end, were glued in bundles and polished to an optical finish on the rear side of the module.

All tiles had dimensions $157 \times 156 \times 3 mm^3$ and were wrapped in a TYVEK envelope, the WLS fibres running between scintillator and TYVEK on both sides of the tiles. Measurements of tile samples using a β -source show uniformity of the light response within 5% (Fig. 10.11).




Figure 10.11: The typical tile response (pulse height, in arbitrary units) to a β -source scanned between the two readout fibres.

The prototype detector contained a 8 mm diameter stainless steel pipe through which a ^{137}Cs radioactive source could be hydraulically driven (Fig. 10.12).




Figure 10.12: The schematic view of the HCAL calibration system. The zoom illustrates the scintillators in the tile array with the vertical scale magnified 10x compared to the horizontal.

Using FEU84 PMT's, with a quantum efficiency of 12-15%, a yield of 30-50 photoelectrons/GeV was obtained. Tests of the prototype response at different energies show good agreement compared to Monte-Carlo predictions using a GEANT simulation of hadronic showers. Data were taken with a Shashlik ECAL module placed upstream of the HCAL. The scatter plot shown in Fig. 10.13 illustrates the combined ECAL and HCAL performance in separating electrons hadrons and muons.

10.5 Photodetectors and front-end electronics

The Preshower, ECAL and HCAL detectors generate fast secondary light signals by the conversion in wavelength-shifting fibres of the primary scintillation of plastic plates. The signals from these detector elements will be digitised and channeled to the data acquisition system according to the general LHCb front-end electronics standards. They will also be transmitted to the Level-0 trigger hardware. The main difference between the three systems is

ECAL and HCAL BEAM TEST




Figure 10.13: Combined beam test of ECAL and HCAL. Muons, electrons and pions are clearly separated.

their light yield. Common solutions for detector readout will therefore be promoted, hopefully leading to cost-effective construction and simpler operation.

10.5.1 Phototubes:

The list of requirements for the photomultipliers is given in Table 10.5. The light yield in the ECAL cells will be in excess of 1000 photoelectrons per GeV. The mean anode current has been chosen small enough ($5 \mu\text{A}$) to ensure a gain degradation of less than 10% per year of operation. WLS fibres incorporating the K27 dye emit a light spectrum peaked at 490 nm. Phototubes with green-extended photocathodes are therefore preferred.

The same requirements apply to the HCAL, but the light yield will be of the order of 50 photoelectrons per GeV. The 0.1-300 GeV dynamic range will be covered by operating the phototubes at a higher gain.

The Hamamatsu R5600 and 5900 series meet these requirements. They are compact, have good timing properties, high limits on mean anode current and good linearity.

Uniformity of the calorimeter response requires the photocathode to be uniform over the area coupled to the fibres or, if this is inadequate, the use of a light mixer in front of the PMT window.

The high voltage supply has to provide for stable PMT operation at high particle rates up to 10

Table 10.5: Requirements for the ECAL phototubes.

	Inner section	Middle section	Outer section
E_e max	200 GeV	100 GeV	50 GeV
$\langle E \rangle$ max/bx	0.12 GeV	0.015 GeV	0.024 GeV
Light yld./GeV	1200 p.e.	1200 p.e.	4500 p.e.
$\langle I_A \rangle$	$\leq 5 \mu A$	$\leq 5 \mu A$	$\leq 5 \mu A$
I_A max	10 mA	27 mA	10 mA
gain	$\leq 5 \cdot 10^3$	$\leq 4 \cdot 10^4$	$\leq 7 \cdot 10^3$

MHz. A Cockcroft-Walton base can provide an average output current as large as 0.1 mA and current pulse amplitude up to 200 mA. Several thousand channels of such a system have been built and tested for the HERA-B calorimeter.

10.5.2 Front-end electronics

Design requirements:

Taking into account the performance requirements for the energy resolution of the calorimeters, a single design with a limited dynamic range is foreseen. The following conditions should be satisfied :

- Shaping within 25 ns to avoid pile up between consecutive bunch crossings. The noise should be smaller than the digitisation error,
- Stable pedestal,
- 12-bit dynamic range. In the ECAL, the energy range covered will vary with the polar angle:
 - 50 MeV \rightarrow 200 GeV for the inner region
 - 25 MeV \rightarrow 100 GeV for the middle region
 - 12.5 MeV \rightarrow 50 GeV for the outer region

No B-physics is lost at the high-energy scale (there are almost no B events with e^-, γ, π^0 above an E_T of 10 GeV/c). At the low-energy end the binning error of ± 1.3 digits, which is obtained after pedestal subtraction and clusterisation, will remain negligible compared to the detector energy resolution at 4 GeV, 2 GeV and 1 GeV, respectively in the 3 calorimeter regions,

- Digitisation of the events at 40 MHz and their storage in a 128-deep pipeline buffer providing 3.2 μs latency,

- An 8-bit transverse-energy measurement must be supplied to the Level-0 trigger logic.

Analogue signal shaping:

Various solutions are being studied for the amplifier shaper: a 5-fold 2.6 nsec integration (Gaussian shaping), or delay-line-clipped integrator. Both unipolar and bipolar versions are under study. The final choice will result from a compromise between critical timing (better for the unipolar delay-line-clipped integrator) and baseline shift at high rate (better for the bipolar Gaussian shaper).

A delay-line-clipped shaping electronics schematic, which incorporates a slow integrator for calibration and monitoring purposes is displayed in Fig. 10.14. This circuit generates a flat-topped output pulse which is robust against jitter.

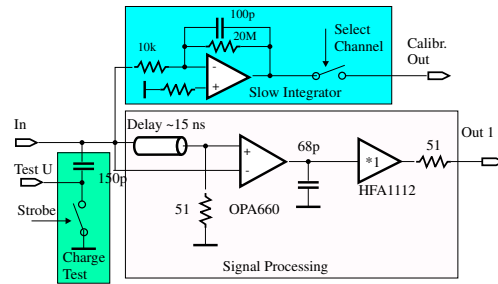


Figure 10.14: Block diagram of delay-line-clipped shaping electronics.

A custom designed implementation of front-end electronics:

The system consists of front-end cards serving 16 channels (Fig. 10.15) and controller cards used to read out 8 front-end cards. The front-end cards contain, for each channel, an amplifier-shaper, a 12-bit 40 MHz ADC with adjustable strobe timing, a register to resynchronize the ADC's output, a 4 K memory used as a look-up table to correct the pedestal, the gain, and possible nonlinearity. A second memory is used to convert the ADC signal to E_T values, as required for the Level-0 trigger. The memory is followed by a FIFO pipeline, 128 deep, which is used to store the signal for a latency of 3.2 μs . After the Level-0 strobe, data are buffered and sent by a 48 line bus to a controller card.

The amplifier-shaper would be built as an ASIC for 4 channels. The remaining electronics use commercial components. For example, the memories are presently available as 8-bit 8K 40 MHz memories, the pipelines as 18-bit 256-cell FIFO's shared

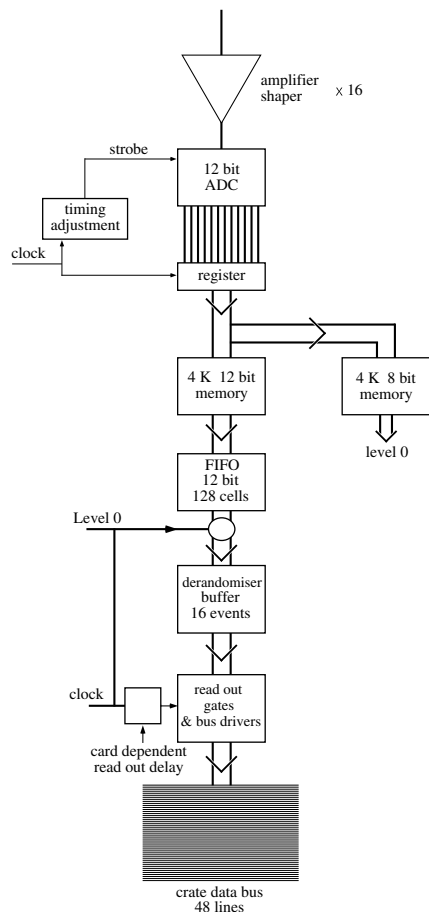


Figure 10.15: Schematics of a custom designed front-end electronics chain.

between channels. The buffer and gate readout would be implemented as one FPGA for 4 channels.

The readout card receives the data from 8 front-end cards, transferred by multiplexing in 32 steps after a Level-0 trigger. Channel address, bunch ID and event ID information are added at this stage (24 bits). It is followed by a 128-event buffer (which could be implemented as 4×4 K 18 bit commercial FIFO). Following a Level-1 trigger, data would be zero suppressed and sent by a high speed link to the DAQ. A single card in a crate can host two controllers, each serving 8 readout cards (see Fig. 10.16).

An implementation based on the FERMI system:

A solution considered for signal digitization and processing is the digital FERMI system [15], suitably adapted to LHCb requirements. Initially designed for high precision calorimetry it is im-

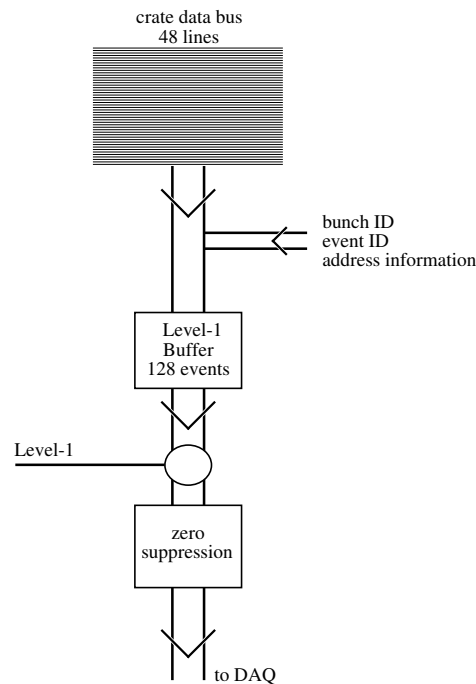


Figure 10.16: Readout card; functional scheme.

plemented as a silicon-on-silicon multichip microsystem with a variety of functional components located on large Si substrates ($\sim 5 \times 5$ cm²) with 4 metal layers.

It could incorporate ADC, pipeline delay (Level-0 and Level-1 buffers), trigger data extraction and service ASIC's.

A number of FERMI microsystems can be mounted on a board, with external circuitry providing the following functions:

- combine data from several FERMI's into one trigger-cell value which is transmitted to the global Level-0 trigger processor,
- generate addresses at each clock signal, pointing at free storage locations for the sampled input data and to generate the pointers required for readout,
- transfer data to the external front-end driver units,
- initialize, control and supervise the FERMI's.

Figure. 10.17 displays the general organisation of this electronics chain.

The choice between the various possible electronics designs will be made after thorough assessment of their performance and costs.

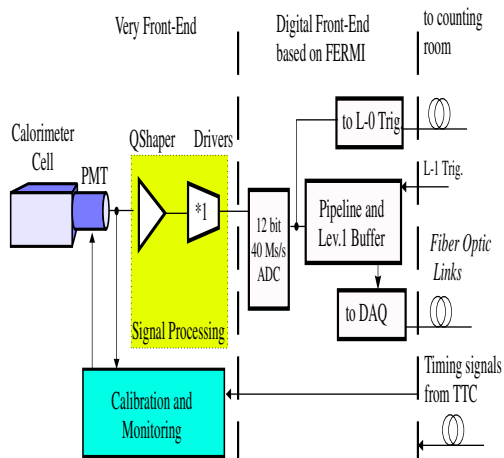


Figure 10.17: Schematic of the calorimeter electronics chain using the FERMI system.

Location of electronics:

All front-end electronics associated to the calorimeters will be housed in crates, close to the detectors. Cable lengths between PMT and electronics will be 6-8 m and the crates can be located where the annual radiation dose should not exceed a few hundred rads, allowing the use of standard components.

References

- [1] S. Conetti, I. Korolko and G. Wilkinson, "Performance of the LHCb Calorimeter Triggers", LHCb/98-014.
- [2] M. Adinolfi, "Optimization of the LHCb Preshower", LHCb/98-015.
- [3] M. Adinolfi et al., "Readout optimization for the KLOE QCAL tile calorimeters" - Proc. 7th Pisa Meeting on Advanced Detectors (1997). La Biodola (1997).
- [4] CMS Collaboration, "CMS Technical Proposal", CERN/LHCC 94-38 (1994).
- [5] A. Karar et al., "Investigation of Avalanche Photodiodes for EM Calorimeter at LHC", CMS-TN/95-095 (1995).
- [6] P. Aspell et al., (RD36 Collaboration), Nucl. Instr. and Meth. in Phys. Res. A 376, 17 (1996).
- [7] SDC Collaboration, "SDC Technical Design Report", SDC-92-201, SSCL-SR-1215 (1992).
- [8] V.G. Vasil'chenko, "New results on plastic scintillators radiation damage", CMS TN 94-220 (1994).
- [9] E. Tarkovsky, "The HERA-B electromagnetic calorimeter", NIM A379 (1996) 515-517.
- [10] HERA-B Collaboration, HERA-B Proposal, DESY-PRC 94/02 May 1994.
- [11] HERA-B Collaboration, HERA-B Design Report, DESY-PRC 95/01 January 1995
- [12] Bazilevsky A. et al., IEEE Transactions on Nuclear Science v.43, No. 3 (1996).
- [13] J. Badier et al., "Shashlik Calorimeter: Beam Test Results", NIM A348 (1994) 74-86.
- [14] ATLAS - Tile Calorimeter TDR, CERN/LHCC 96-42.
- [15] B. Lofstedt et al., "FERMI: A digital Front End and Readout Microsystem for High Resolution Calorimetry", Proc. Second workshop on electronics for LHC experiments, Balaton, 23-27 Sep. 1996, (1996) 84.

Chapter 11 Muon Detector

11.1 Introduction

The LHCb Muon Detector performs two functions, muon identification and the Level-0 muon trigger formation. The components of the Muon Detector are shown in Fig. 11.1.

The muon detector contains a longitudinally segmented shield to attenuate hadrons, photons and electrons. The components of the shield comprise the electromagnetic and hadronic calorimeter and four layers of steel. Five muon tracking stations, M1–M5, are situated as shown in Fig. 11.1. To penetrate this shield and reach station M5 a muon must have an energy of at least 6 GeV. The technologies chosen for the muon stations are Multigap Resistive Plate Chambers (MRPC) for most of the coverage and Cathode Pad Chambers (CPC) for regions which are subjected to the highest rates. Present data on MRPC's indicate that above 5 kHz/cm² the efficiency begins to decrease significantly [1]. Based on a 5 kHz/cm² limit and simulations of 14 TeV interactions in the LHCb spectrometer, CPC's will be used for the unshielded M1 station and for instrumentation of the region close to the beam in stations M2–M5.




Figure 11.1: Muon Detector Components.

Table 11.1: Muon Detector Components

Detector Element	Composition	Depth		
		m	z/\bar{X}_0	z/λ_0
M1	CPC	0.2		
ECAL+Presh.	Pb/Scint.	0.5	27.7	1.3
HCAL	Fe/Scint.	1.5	68.0	7.3
Shield-1	Fe	0.3	16.8	1.8
M2-5	MRPC/CPC	0.4		
Shield-2,3,4	Fe	0.7	39.3	4.2
Total Thickness			230	23

11.2 Detector layout

11.2.1 The Muon shield

Table 11.1 gives the thicknesses of the shield and other components of the Muon Detector. The overall thickness of the shield is constrained by the size of the experimental hall and the position of the LHC quadrupole close to IP8. The total weight of the steel shield is approximately 1800 tonnes.

The muon shield and the muon stations are designed so that the left and right halves of the muon detector can be separated to allow access to the beam pipe. In addition to the shield components in Table 11.1, M2–M5 are protected from secondary particles which emerge from the beam pipe. This shielding extends from the beam pipe at 10 mrad out to 12 mrad in both x - and y -projections. Additional shielding immediately behind M5 protects it

from backscplash from nearby LHC beam elements and from hadrons and muons emerging from the collider tunnel.

11.2.2 The Muon stations

Muon station M1 is positioned at 12.15 metres from the interaction point, immediately in front of the electromagnetic detector. The remaining four stations, M2–M5, are embedded in the muon shield at average positions of 15.5, 16.6, 17.7 and 18.8 metres from the interaction point. They fit into the 40 cm space available between the steel shielding

walls.

The angular coverage of the muon detector involves a compromise between cost, geometry of the required beam shielding, detector dimensions, sustainable detector rates, trigger rates from $\pi/K \rightarrow \mu$ decays, and $B \rightarrow \mu$ acceptance. The inner and outer acceptance (horizontal \times vertical) have been set to 25 mrad \times 15 mrad and 294 mrad \times 245 mrad respectively.

The chamber technology used in the various regions of the Muon Detector is determined by the rates experienced by the detectors. The highest rates are experienced in the M1 station which is not protected by the shield. The MARS [2] program was used to calculate the particle rates, including hits due to neutrons, expected in the chambers. The rates per cm^2 per interaction as a function of r , the distance from the beam, are shown in Fig. 11.2 for stations M1 and M2. The rates in stations M3–M5 are approximately one half those in M2. At a rate of $5 \text{ kHz}/\text{cm}^2$, the present MRPC technology begins to lose efficiency and CPC's, which can withstand higher rates, are needed. The horizontal line in Fig. 11.2 indicates the $5 \text{ kHz}/\text{cm}^2$ level at a luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. In stations M2–M5 it is only the region near the beam where this limit is exceeded. In M1 however, a significant area exceeds the limit. In these regions CPC's will be used.

In spite of the relatively high rates in the small angle region of the detector planes, the occupancies are acceptable because each module has a 2D pad structure. Even for M1, the occupancy is less than 10% beyond 60 cm (assuming a signal pulse width of 100 ns). The occupancies in M2–M5 are below 10% everywhere.

Station M1 comprises a mosaic of four-layer CPC modules. The four layers are required in order to achieve single beam-crossing timing resolution and good efficiency for muon detection. Stations M2–M5 are a combination of four-layer CPC chambers covering the angular region ≤ 60 mrad and two layers of 4×0.8 mm gap MRPC's covering the region beyond 60 mrad. Two layers of MRPC per station are required for good efficiency. The Level-0 muon trigger algorithm requires a hit in all five stations. Since the Level-0 trigger efficiency depends on the fifth power of the individual muon station efficiency good efficiency is essential in each station.

The size of the individual CPC chamber modules in the M1 mosaic is $x \times y = 2.4 \text{ m} \times 1.3 \text{ m}$. Fourteen such modules are required. The central region near the beam pipe is covered by two smaller




Figure 11.2: Rates per cm^2 per interaction: in muon stations M1 and in M2.

modules which have maximum dimensions of 1.45 m \times 1.3 m with a corner of each module removed to allow for the beam pipe.

The MRPC module sizes scale with distance from the interaction region. They vary from $\approx 2.2 \text{ m} \times 1.2 \text{ m}$ in M2 to $2.7 \text{ m} \times 1.5 \text{ m}$ in M5. Twenty four modules are required for each of the four stations. The CPC modules covering the beam region will vary in size from $\approx 1.6 \text{ m} \times 1.1 \text{ m}$ at M2 to $2.0 \text{ m} \times 1.2 \text{ m}$ at M5. Two modules of this type, with corners removed, will be required per station.

The muon station modules will be attached to steel support plates and hung from a rail system. The modules on the left and right sides of each station can be positioned precisely relative to one another and withdrawn as units for servicing and repair without requiring movement of the shield walls.

11.2.3 The chamber pad structure

The necessity of forming the Level-0 muon trigger within $3.2 \mu\text{s}$ requires a 2D pad structure in M1–M5 so that muon 3D hit information is available at the earliest possible time. The pad structure required for the Level-0 muon trigger is referred to below as the “logical” pad structure. The largest logical pad sizes would generate unacceptable noise and time slewing due to their capacitance. In these cases, the logical pads will be divided into smaller “physical” pads which will be ORed to form the logical pad input signal to the Level-0 trigger.

In order to minimise the number of logical pads, each of the five muon stations is divided into four regions with different pad sizes. The pad sizes double from region to region as one proceeds from the beam to the outer edge of the detector, thereby ap-

proximately equalising the occupancies and rates experienced by pads in each region. Furthermore, the x and y pad dimensions in each muon station are scaled so that the pad configuration of M1–M5 is projective to the interaction point.

The x dimensions of the logical pads in stations M1 and M2 are determined by the precision required to obtain good resolution of muon p_T for the Level-0 trigger. The y dimensions are determined by the required rejection of machine associated backgrounds [3] by pointing at the interaction region in the y projection. Since more precision is required in the bending plane (x) to determine the muon p_T than in the non-bending plane (y), a y/x aspect ratio of two to one has been adopted for the pads in M1–M2.

The logical pad sizes in M3–M5 are set by the requirements of pattern recognition of penetrating muons and the formation of triple coincidences. The pad configuration for M3–M5 follows a similar pattern except that the x pad dimension has been doubled relative to that in M1 and M2, reflecting the lesser spatial resolution required to find the triple coincidences. Hence, M3–M5 pads are square but still project to the interaction point.

Following studies of the effects of pad sizes on pattern recognition and Level-0 trigger rates and efficiency, the pad structure shown in Fig. 11.3 has been adopted for station M1. With this choice of pad structure, stations M1 and M2 each have 12904 logical pads. Stations M3–M5 each have 6452 logical pads. This results in a total of 45164 logical pads to be transported to the Level-0 trigger.

For the reasons stated above, the large logical pads in the outer regions will be divided into small physical pads, ORed together locally at the individual modules. Moreover, the pads from the different layers in one station must be combined. In the case of the two-layer MRPC's, an OR of corresponding pads from each layer will be formed. For the four-layer CPC's, a 3 of 4 majority logic will be formed using the corresponding pads from the four layers. The OR and majority logic circuitry will be sited at the periphery of the Muon Detector. The total number of physical pads and, hence, discriminator/amplifiers required for the Muon Detector system is approximately 236000.

11.3 Chamber technologies

Two technologies have been chosen for the muon station detectors: Cathode Pad Chambers and Multigap Resistive Plate Chambers. Both CPC's and MRPC's are under study as a part of the LHCb

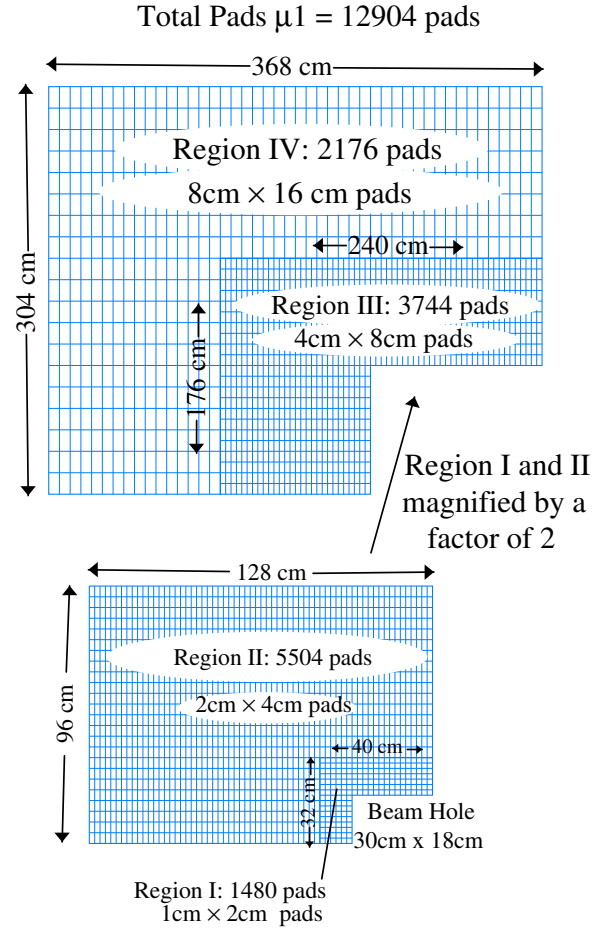


Figure 11.3: M1 logical pad configuration.

R&D programme.

11.3.1 Multigap Resistive Plate Chambers (MRPC)

The interest of the LHCb collaboration in multigap resistive pad chambers has been generated by the work of Williams et al. [4]. The MRPC resembles the conventional RPC technology except that the MRPC has more than one gap, and the gaps are much smaller. The use of multiple narrow gaps results in better timing and less streamer activity in comparison to single-gap RPC's. It is foreseen to use MRPC chambers with four gaps of 0.8 mm (see Fig. 11.4) in the LHCb muon stations M2–M5.

The resistive plates of the MRPC's are 0.9 mm-thick melamine-phenolic sheets with the melamine forming a $\approx 250 \mu\text{m}$ layer on either side. The chamber operates at approximately 16.4 kV. It has been shown [1] that this type of chamber, operated with




Figure 11.4: Schematic section of one end of a Multigap Resistive Plate Chamber (not to scale).

a gas mixture of 5% iso- C_4H_{10} , 10% Argon and 85% C_2F_5H bubbled through water at $10^\circ C$, has a $\geq 98\%$ efficiency at a “flood” illumination of 5 kHz/cm^2 . The “time walk” of the leading edge of the signals (for varying rates up to 5 kHz) is approximately 2 ns and the leading-edge time resolution (at a fixed rate) is a few ns. This allows single beam-crossing resolution in stations M2–M5. A 5 kHz/cm^2 rate limit for M2–M5 has been adopted, although it is expected that with “spot” illumination rather than “flood” illumination, the MRPC’s can withstand higher rates.

11.3.2 Cathode Pad Chambers (CPC)

Cathode Pad Chambers are similar to the Cathode Strip Chambers [5] which were proposed for the muon chambers in the LHCb Letter of Intent [6]. In Fig. 11.5 the cross section of the CPC under study for LHCb is shown. Each CPC chamber will be composed of two such submodules, each containing two layers. The anode wires, as shown in Fig. 11.5, are placed asymmetrically in the 5 mm gas gap at 1 mm from the pads and 4 mm from the cathode. The wires are stretched in the y (vertical) direction. The wire diameter is $30 \mu\text{m}$ and the wire spacing is 2 mm .

The cathode sheets are Cu-clad G10 glued to a 2 cm -thick honeycomb plate. Anode bars are fixed on both sides of the cathode plane, and the wires are wound such that both anode-wire planes are formed on either side of the cathode plane at the same time. The pad-plane boards are G10 clad on both sides with $20 \mu\text{m}$ of Cu. The pad geom-




Figure 11.5: Cross Section for Cathode Pad Chamber (units in mm).

etry is chemically etched on one side of the pad plane and connected to electrical leads on the other side of the board via electrochemically metallized holes. The pad boards are glued to a honeycomb plate with grooves which accommodate the zero-insertion-force connectors, fixed on the pad boards, and the flexible cables with front-end readout chips.

A possible gas mixture is 50-60% CO_2 , 30% Ar and 20-10% CF_4 . This provides a high drift velocity ($80 \mu\text{m/ns}$) and results in good aging properties, allowing operation up to $4C/cm$ (equivalent to a rate of $50\text{kHz cm}^{-2}\text{s}^{-1}$ for ten years) with a gain of 5×10^4 . With the CPC geometry as described, the signal will be induced with high probability on a single pad.

11.4 Front-End electronics

No special preamplifier design features are anticipated for the muon detectors. Bi-polar (BJT) amplifier/discriminators will be used since they are relatively radiation hard and have lower voltage noise than CMOS devices. Two possible candidate chips are the Penn ASD8 and the Rutherford RAL110. The Penn ASD8 preamplifier can operate at a threshold of 1 fC with a 50 pF input capacitance. It has previously been used with MRPC’s. In the present chamber design the maximum physical pad size is approximately 50 cm^2 . This corresponds to a capacitance of $\leq 75 \text{ pF}$ in both MRPC’s and CPC’s. No special pulse shaping of the preamplifier output is anticipated. Even for the largest

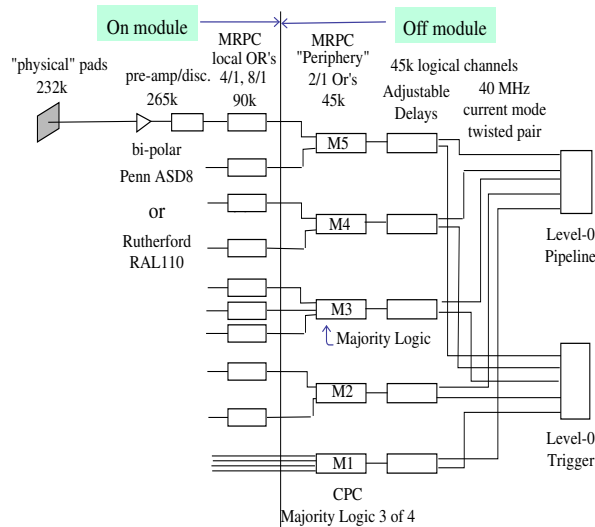


Figure 11.6: Schematic of pad signal flow from LHCb Muon chambers.

physical pads of the MRPC's, peaking times of less than 10 ns are expected. The peaking times for the CPC's will be somewhat longer. The MRPC and CPC pads will be connected to the preamplifier chips using 100 Ω leads.

Taking into account the number of layers in each muon station, the number of physical pads is about 78000 in M1, 47000 in M2 and 37000 in each of M3–M5, giving a total count of 236000 preamplifier and discriminator channels. These will be mounted at the edge of each module. Appropriate combinations (as described in Section 11.2.3) of the physical pad signals will be made locally, on the chamber modules, to form the 45164 logical pad signals used for the Level-0 muon trigger.

The radiation levels in stations M2–M5 are less than 1 krad per year at the positions of the preamplifiers and discriminators and no special radiation-resistant electronics are required. For station M1 the situation is more demanding, with an expected radiation level of 10-30 krad/year, depending on where the preamplifiers can be positioned. Radiation tolerant electronics may be required for this station.

It is foreseen to transmit all logical binary signals in parallel to the Level 0 trigger processor. A delay trim device that compensates for different lead and cable lengths is needed to equalise the arrival times of these bits at the Level-0 trigger and the Level-0 data pipeline. The Level-0 trigger and pipeline for the muon detector will be located close to the detector and transmission to these de-

vices will be at 40 MHz in differential current mode via ≈ 10 m length twisted-pair cables. The flow of the pad information from the detector to the Level-0 trigger and the Level-0 pipeline is shown in Fig. 11.6.

The movement of data after the Level-0 pipeline follows the standard philosophy common to the various LHCb detector components. The 45K bits will pass into a derandomising buffer, the output of which will then be multiplexed by 32:1 into a zero suppression/encoding stage. The time estimated for the encoding process is 10 ns per hit which leads to a negligible deadtime. The encoded events are then passed to the Level-1 trigger and FIFO buffer. On receipt of a Level 1 trigger, the appropriate event is extracted from the FIFO buffer and routed to Level 2 in the counting room. For the muon detector, the pipeline, derandomising buffer, 32:1 multiplexer and encoding electronics are close to but not on detector. These issues are discussed more fully in Ref. [7].

11.5 Performance

11.5.1 The Level-0 muon trigger

If the muon p_T threshold is set to a relatively small value ($p_T \geq 0.7$ GeV/ c) then about 3.0% of the minimum bias interactions will be retained while achieving an efficiency for $B \rightarrow \mu + X$ within the muon detector acceptance of 50%. This minimum bias retention corresponds to a Level-0 muon trigger rate of 150 kHz at an LHCb luminosity of $2.0 \times 10^{32} \text{cm}^{-2} \text{s}^{-1}$ (assuming the LHCb multiple-interaction crossing veto selects the 31% of beam crossings which have only one interaction). Alternatively, the minimum bias trigger rate can be reduced below 1%, corresponding to a $B \rightarrow \mu$ efficiency of 28%, by raising the Level-0 muon p_T threshold to 1.6 GeV/ c .

11.5.2 Muon identification

For the purposes of evaluating the muon identification efficiency and the $B \rightarrow \mu$ misidentification probability, a “good muon” is defined in the Muon Detector by a triple coincidence associated with a combination of M1–M2 hits in a given search window. A track segment downstream of the magnet is then identified as a muon if that track segment is the one that points closest to the centre of the pad in station M1 that is hit by this good muon.

There are hit configurations that satisfy the “good muon” criteria given above but do not re-

sult from true muons. Hadron punch-through contributes to such a false muon signal (but causes no loss of efficiency for B events). Preliminary estimates of this background indicate that it is negligible.

A more serious problem is the case in which a true muon produces a good muon signature in the Muon Detector but is not connected to the correct upstream track due to multiple scattering or to pattern recognition problems. This can contribute to the wrongly-tagged class of B decays. This problem was investigated for two different classes of muons from $B \rightarrow \mu$ decays: those that cause a trigger and those that do not (because they are below the trigger p_T cut, which is set to 1.0 GeV/c). The second class of muon may be used as a tag in the event, even though it does not cause a trigger. The percentage of $B \rightarrow \mu$ that satisfy the Level-0 trigger but have a wrong association is $\sim 2\%$. The percentage of $B \rightarrow \mu$ that do not cause a trigger, that can be used as a tag, but have a wrong association is 9%.

11.6 Future R&D

Further R&D is planned for the muon chambers in the following areas:

- high rate performance of both CPC's and MRPC's
- charge sharing in both the CPC's and MRPC's for the 1cm \times 2cm pads
- cross talk, noise and time slewing for the larger pads to verify that single-bucket resolution can be achieved
- efficiencies and time resolutions for the two chamber technologies.

It will also be important to determine the dependence of chamber performance on construction techniques and materials. A two year programme of beam tests with a series of prototypes will be carried out by the LHCb collaboration, beginning in summer 1998.

References

[1] E. Cerron Zeballos *et al.*, "Micro-streamers and the Micro-gap Resistive Plate Chamber", CERN/LAA-MC 97-04, (Oct, 1997), submitted to Nucl. Instr. Meth. A.

- [2] I. Azhgirey, *et al.*: "MARS: A code package for Radiation Problems Solution of Electro-Nuclear Installations Design", Proc. of XV Conference on Charged Particle Accelerators, Protvino, October 1996.
- [3] A.I. Droshdin, M. Huhtinen and N.V. Mokhov, "Accelerator Related Background at the CMS Detector at the LHC", CERN CMS TN/96-056, (April, 1996).
- [4] E. Cerron Zeballos *et al.*, "A New Type of Resistive Plate Chamber: The Multigap RPC", Nucl.Instr.Meth. A 374 (1996),132.
E. Cerron Zeballos *et al.*, "Avalanche Fluctuations Within the Multigap Resistive Plate Chamber", Nucl.Instr.Meth. A 381 (1996),569.
- [5] H. Faissler *et al.*, CERN/DRC/93-49, RD5 Status Report, (1994)25; GEM Technical Design Report, GEM-TN-93-262, SSCL CR, 1219; CMS Technical Proposal, CERN/LHCC/94-38.
- [6] LHCb Letter of Intent, CERN/LHCC/95-5, LHCC/I8,65.
- [7] G. Corti, B. Cox, K. Nelson and A. Tsaregorodtsev, "LHCb Muon Detector Electronics", LHCb/98-003.

Part III

Data Handling

Chapter 12 Trigger

12.1 Introduction

Level-0 and Level-1 algorithms use localised detector components for fast decisions, and this is shown schematically in Figure 12.1. In most cases the detector information is delivered with a reduced resolution.

Transverse momentum can be quickly measured by the calorimeter system for hadrons and electrons, and by the muon system for muons. In total, the number of channels to be read for the p_T algorithm is about 60k. The high- p_T trigger is therefore applied first (Level-0). Events with multiple pp interactions, which are difficult to analyze, are also rejected at Level-0 by a pile-up veto. This is based on counting the number of primary vertices and measuring the total energy. The overall rate reduction factor of Level-0 is about 40.

Identification of detached vertices is more complicated and requires more time. It is therefore applied at the next trigger level, Level-1. Events are selected based on the impact parameters of tracks and reconstructed secondary vertices, as measured with the vertex detector. The high- p_T tracks found by the Level-0 trigger are also re-examined using the tracking system. The overall rate reduction factor of Level-1 is ~ 25 .

The DAQ system reads event data at the Level-1 output rate of 40 kHz. The goal is to write events to storage at a rate of 200 Hz. The high level triggers will reduce the rate from 40 kHz to 200 Hz, the data recording rate, in 2 stages. In the first stage (Level-2), fast algorithms acting on partial event data are used, based on combining vertex and tracking detector data at full resolution. The goal is to reduce the rate to 5 kHz, by improved vertex reconstruction. In the second stage (Level-3), complete event data are used and full B-decay analysis is performed to select interesting B physics channels.

The trigger system has been designed to be flexible, robust and efficient. Events are selected by various criteria that can be easily adjusted according to the experimental conditions. There is no strong dependence on a particular trigger component, and the rate reduction is fairly evenly spread

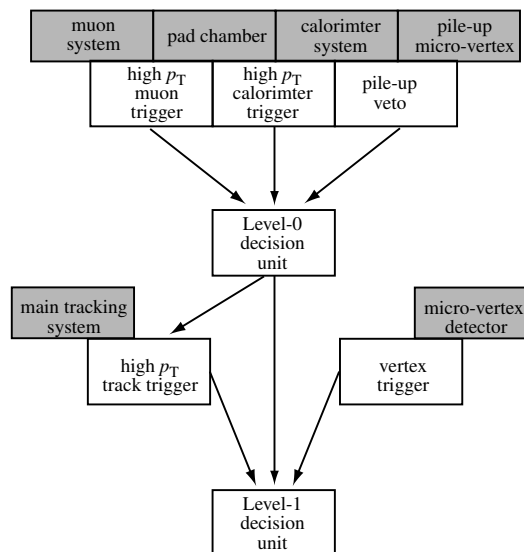


Figure 12.1: A schematic view for the LHCb Level-0 and Level-1 trigger with the detector components used in the algorithms.

over the four trigger levels. Trigger conditions are optimized to select interesting B events with a high probability of being reconstructed offline.

12.1.1 Level-0

The Level-0 trigger has an input rate of 40 MHz and an output rate of ~ 1 MHz. It is based on the identification of leptons, hadrons and photons with high- p_T in calorimeters and muon chambers, combined with a pile-up veto. The pile-up veto unit identifies bunch crossings with more than one pp interaction using a dedicated piece of the vertex detector system and the calorimeters.

The Level-0 decision unit decides whether an event is kept for the next level. A bunch crossing with a high probability of containing more than one pp interaction is rejected. About 90% of the output rate is given to events with a single high- p_T muon, electron or hadron. The p_T thresholds can be adjusted according to physics needs. At present it is assumed that they share 20%, 10% and 60% of the total trigger rate, respectively. The remaining

10% of the output rate is reserved for events with a single high- p_T photon, multi-leptons and some other combinations. Examples of decay channels that can be selected by the single high- p_T triggers are as follows:

$$e: b \rightarrow e + X \text{ (flavour tag}^1\text{)}$$

$$B \rightarrow J/\psi(e^+e^-) + X$$

$$\mu: b \rightarrow \mu + X \text{ (flavour tag)}$$

$$B \rightarrow J/\psi(\mu^+\mu^-) + X$$

$$h: B_d^0 \rightarrow \pi^+\pi^-, B_s^0 \rightarrow D_s^+K^-$$

$$\gamma: B_s^0 \rightarrow K^{*0}\gamma$$

In the case of a positive decision, the Level-0 decision unit passes the information on the high- p_T particle to the Level-1 track trigger to be used as a seed in the tracking algorithm. The same information is also given to the Level-1 decision unit. The Level-0 decision unit is programmable so that the selection criteria can be easily changed to reflect physics interests.

12.1.2 Level-1

The Level-1 trigger has an input rate of ~ 1 MHz and output rate of ~ 40 kHz. It is based on two independent systems, the vertex trigger and the track trigger. The maximum latency of Level-1 is $256 \mu\text{s}$, which allows sufficient time to run sophisticated algorithms providing a high selection efficiency.

A key feature of b quark events is the existence of displaced secondary vertices, and this is exploited by the vertex trigger using the vertex detector. After reconstruction of the primary vertex, tracks with a significant impact parameter to this vertex are selected. A secondary vertex search looks for combinations of pairs of those tracks that are close in space. The output of the vertex trigger is an event probability based on the number of secondary vertices found and on their separation from the primary vertex.

A significant fraction of the events passing the high- p_T triggers is due to muons produced by particles decaying in flight, electrons from γ -conversion and overlapping showers simulating highly energetic single particles. The tracking trigger aims to refine the Level-0 trigger by adding information from the tracking system. Starting from the Level-0 seeds in the calorimeter and the muon systems, corresponding tracks are searched for using a simplified Kalman filter algorithm. Events are rejected

if no corresponding track is reconstructed. Additional rejection power can be obtained by applying a cut on the reconstructed transverse momentum.

The Level-1 decision unit decides whether an event is to be read out by combining the results from the vertex trigger and tracking trigger. The vertex trigger alone can deliver the required rate reduction with a high efficiency for b events triggered by the Level-0 single high- p_T triggers. However, the additional reduction provided by the track trigger gives an extra safety factor. For the rate calculations in this document, it is assumed that only the vertex trigger is used.

12.1.3 Trigger Levels 2 and 3

Algorithms for both Level-2 and Level-3 triggers run in farms of commercial processors. Level-2 reduces the number of events without b-quarks by a factor of eight to ~ 5 kHz by refining the vertex trigger. Fake displaced secondary vertices are generated by low momentum tracks due to significant multiple scattering. By adding momentum information for the tracks forming the secondary vertices, rejection of non-b events is greatly improved.

The Level-3 trigger further reduces the data rate by a factor of 25 to ~ 200 Hz. The events reaching the Level-3 stage consist of $b\bar{b}$, $c\bar{c}$ and light quark events in a ratio of about 1:2:1. The inclusive $b\bar{b}$ rate is thus 1.3 kHz and already exceeds the Level-3 output rate by an order of magnitude. At Level-3, it is therefore necessary to reconstruct specific b-hadron decay modes using all the detector components, including the RICH.

12.1.4 LHCb trigger performance

The performance of the trigger is best demonstrated by quoting the trigger efficiencies for events where the decay final states are fully reconstructed in the offline analysis. If the final states are not self-tagging, a correct flavour tag is required in addition. Table 12.1 shows the trigger efficiencies for those “useful” b-quark events with five selected decay modes, which are sensitive to different angles of the unitarity triangles. The Level-0, -1 and -2 triggers are applied consecutively. The Level-3 trigger does not introduce any additional loss for the signal events.

12.2 Optimal luminosity

The LHC operates with a nominal bunch crossing frequency of 40 MHz. The Level-0 pipeline of the

¹This is defined in Chapter 15.2

Table 12.1: Trigger efficiencies of Level-0, -1 and -2 for events where the final state is reconstructed offline and its flavour is correctly tagged, for non-specific flavour final states. Level-0, -1 and -2 triggers are applied consecutively.

Decay Mode	Physics	Level-0	Level-1	Level-2
$B_d^0 \rightarrow \pi^+\pi^- + \text{right tag}$	$\beta + \gamma$	0.76	0.48	0.83
$B_d^0 \rightarrow J/\psi(e^+e^-)K_S + \text{right tag}$	β	0.72	0.42	0.81
$B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S + \text{right tag}$	β	0.88	0.50	0.81
$B_s^0 \rightarrow D_s^+K^- + \text{right tag}$	$\gamma - 2\delta\gamma$	0.54	0.56	0.92
$B_d^0 \rightarrow \bar{D}^0K^{*0}$	γ	0.37	0.59	0.95




Figure 12.2: Probabilities for having 0,1,2,3 and 4 pp interactions per bunch crossing as a function of the machine luminosity at LHCb.

front-end electronics is clocked at this frequency. With the displaced interaction point of LHCb, only 2649 of the 3564 bunches, i.e. 74.3%, will collide [1]. The average bunch-bunch crossing rate at IP8 is, therefore, 30 MHz. Figure 12.2 [2] shows the probabilities for having zero, one, two, three and four pp interactions per bunch crossing as a function of the machine luminosity, where the inelastic pp cross section is assumed to be 80 mb. The number of bunch crossings with one pp interaction is maximum at $\sim 4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$.

Compared to events with one pp interaction, those with multiple pp interactions are more difficult to reconstruct due to the increased particle density. In the forward region, the decay length of a B-meson is mainly given by the momentum

component along the beam direction, z . Since the spread of the primary vertex along the z direction is large, $\sigma_z \approx 5 \text{ cm}$, the z coordinate of the B-meson production point becomes ambiguous if there is more than one primary vertex per bunch crossing. This makes it difficult to measure the decay length of a B-meson. The flavour tag with the decay products from the accompanying b-hadron also becomes difficult. Moreover, the high- p_T triggers retain a higher proportion of multiple interaction bunch crossings than single interaction bunch crossings, the retention factor being almost proportional to the number of interactions. Therefore, given that the trigger accept rate is limited, the rejection of pile-up events is essential in order to maximize the number of triggered $b\bar{b}$ events.

Other important factors influencing the choice of the running luminosity are the radiation damage that may result, and the occupancy of the detectors.

Taking all these points into account, the average running luminosity of the LHCb experiment has been chosen to be $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. Bunch crossing rates with only one, and with more than one, pp interaction are 9.3 MHz and 3.0 MHz, respectively. With the pile-up veto, these rates are reduced to 8.92 MHz and 0.6 MHz, respectively. In order to achieve the Level-0 output rate of 1 MHz, a rejection factor of ~ 10 is required for the triggers from the non pile-up bunch crossings. The detector is designed to withstand a luminosity of $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ for short periods.

12.3 Level-0 Trigger

12.3.1 Calorimeter triggers

Electron Trigger: The primary purpose of the electron trigger is to identify prompt electrons from B decays by searching for clean isolated electron

showers with high transverse energy (E_T). The main backgrounds come from conversion electrons, photons and hadrons. The detectors used are the electromagnetic calorimeter (ECAL), the preshower detector and the pad chamber in front of the ECAL (M1). What follows is a description of the algorithm that uses energy clusters defined by 3×3 cells. An alternative algorithm using 2×2 cells is described in the implementation section.

1. ECAL 3×3 clusters are found using a procedure in which the central cell is required to have more energy than any of its eight neighbours. Clusters are required to be above some minimum energy, typically 3 GeV, depending on the minimum E_T cut finally applied.
2. The centre-of-gravity of the cluster is calculated in both the x and y coordinates. For the remainder of the algorithm this defines the cluster position, and the candidate's trajectory is assumed to be a straight line between the interaction region and this point.
3. Hits are required in the pad chamber (M1) in positions corresponding to the candidate's trajectory. The granularity of the pad chamber matches that of the ECAL in the non-bending plane, but is twice as fine in the bending plane. A window of one pad is used for the former, and two for the latter.
4. An energy deposition above that produced by three minimum ionising particles is demanded in the preshower cell along the candidate's trajectory.
5. The energy (E_c) of the central cell is required to be at least 40% of the total energy of the cluster (E_9).
6. The E_T of each cluster is calculated. For this calculation, the energy has been calibrated according to the region of the calorimeter in which the cluster lies. The single particle electron trigger is defined by requiring that the E_T of the leading cluster exceeds a given threshold.

The effect of these successive cuts can be seen in Table 12.2, for events containing a semi-leptonic B decay to an electron, where the electron lies within the acceptance of the electromagnetic calorimeter, and for inelastic pp interactions. It should be noted that the geometrical acceptance for the electron to hit the calorimeter, in these events, is 16%. The




Figure 12.3: Variation in electron trigger efficiency as a function of E_T threshold for signal and inelastic pp interactions. Curves are shown for both $B \rightarrow eX$ events, with the electron within the electromagnetic calorimeter acceptance, and $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$ events reconstructed offline.

E_T threshold has been chosen to achieve a suppression of ~ 100 for inelastic pp interactions. To illustrate the full spectrum of performance available, Figure 12.3 shows the variation in signal efficiency against retention of inelastic pp interactions for such events, and for $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$ events where $J/\psi K_S^0$ is reconstructed offline.

Note that this algorithm uses no information from the hadron calorimeter. The preshower alone provides powerful hadron rejection and no significant further suppression is achieved by using the hadron calorimeter in addition [3].

It can be seen from Table 12.2 that the most important contribution for non-b events passing the trigger is electrons from photon conversions. Figure 12.4 shows the distribution of the z coordinate of these conversion points and reflects the position of the detector material. The description of the material in the simulation program is based on conservative estimates for all detector components.

Photon Trigger: High- E_T photons can be selected with a small modification to the electron trigger algorithm [4]. Firstly the pad chamber should have no hits (step 3). Secondly a significantly higher E_T threshold must be applied as compared to the electron trigger (step 6). For example, inelastic

Table 12.2: Number of clusters per event passing successive stages of the electron trigger, for signal and for background events. All the signal events have an electron within the acceptance of the electromagnetic calorimeter, whereas the inelastic pp interaction events are generated over 4π . Both samples are made up of 10k events.

Contribution from	B \rightarrow eX			Non-b		
	γ	e	h	γ	e	h
All ECAL Clusters with $E > 3$ GeV	6.4	4.2	9.8	5.3	2.9	9.0
Pad chamber hit found	0.75	2.7	3.2	0.51	1.80	2.9
Preshower > 3 MIP	0.57	2.4	0.67	0.37	1.50	0.59
ECAL $E_c/E_0 > 0.4$	0.52	2.20	0.41	0.34	1.40	0.36
ECAL $E_T > 2.34$ GeV	0.028	0.160	0.005	0.002	0.005	0.003




Figure 12.4: Origin in z of gamma conversions passing the electron trigger.

pp interaction events are suppressed by a factor of ~ 150 with a cut at 4 GeV. Therefore, the effect of the photon trigger on the overall trigger rate is negligible. The physics potential of this trigger is discussed in Chapter 15.10.

Hadron Trigger: Accepting events with high E_T clusters in the hadron calorimeter (HCAL) efficiently selects B decays to hadronic final states, such as $B_d^0 \rightarrow \pi^+\pi^-$ or $B_s^0 \rightarrow D_s^-\pi^+$ where at least one stable hadron can have high- p_T . The algorithm is as follows:

1. A search is made for HCAL 3×3 clusters with an identical algorithm to that used in the ECAL for the electron trigger. The higher E_T thresholds of the hadron trigger, compared to that of the electron trigger, mean that only clusters of > 5 GeV need be considered.
2. The cluster's position and trajectory is calculated as for the electron trigger.
3. ECAL 3×3 clusters lying on the candidate's trajectory are added to obtain the total trigger cluster energy.
4. The E_T of each cluster is calculated. The single hadron trigger is defined by requiring that the E_T of the leading cluster exceeds a given threshold.

This simple algorithm gives a good separation between signal and background. Figure 12.5 shows the trigger efficiency as a function of the cut on the reconstructed E_T for b-quark events with $B_d^0 \rightarrow \pi^+\pi^-$ and $B_s^0 \rightarrow D_s^-\pi^+$ decays, where the final states are reconstructed offline and correctly flavour-tagged. As expected, the lower multiplicity decay has higher efficiency. Also indicated is the suppression the cut gives on inelastic pp interaction events. It can be seen that a cut at $E_T \sim 2.4$ GeV retains 60% of the $B_d^0 \rightarrow \pi^+\pi^-$ events while suppressing background by a factor of ~ 17 . This performance is significantly better than is achievable by the lepton triggers in this channel.

It is possible to simplify the algorithm still further without suffering from a dramatic loss in performance. Neglecting the ECAL contribution to the trigger cluster, an E_T cut providing a factor 30 suppression of inelastic pp interactions gives a trigger efficiency of 28% for b-events with $B_d^0 \rightarrow \pi^+\pi^-$ decays.




Figure 12.5: Performance of the hadron trigger as a function of E_T cut. Signal efficiency is shown for $B_d^0 \rightarrow \pi^+\pi^-$ (solid circles), and $B_s^0 \rightarrow D_s^-\pi^+$ (solid triangle). Suppression for inelastic pp interactions is also indicated (open squares, right-hand scale).

Rather than performing a centre-of-gravity calculation, the cluster coordinates can be approximated by the position of the central HCAL cell with negligible change in E_T discrimination. However, the more accurate procedure is retained for uniformity with the electron trigger, and to provide a more precise seed for the Level-1 track trigger.

A modest energy resolution of $\approx 80\%/\sqrt{E}$ is sufficient [3]. The hard p_T spectrum of hadrons in the decays of interest allows E_T thresholds to be set high above the steeply falling p_T spectrum of particles in the non-b events.

Figure 12.6 shows the true p_T spectra of the leading hadron in b-events with $B_d^0 \rightarrow \pi^+\pi^-$. Also shown are the equivalent spectra for those inelastic pp interaction events passing a cut on the reconstructed E_T of 2.4 GeV. A significant fraction of those clusters in the non-b events passing the trigger are of genuinely high- p_T . Occasional overlaps lead to a few fake triggers from low- p_T tracks.

Hardware implementation using 3×3 clusters

Three different approaches are being studied in order to demonstrate the feasibility of implementing the Level-0 high- p_T trigger. Two approaches attempt to reuse existing hardware architectures for which expertise is contained in the collaboration. The first of these uses a general purpose programmable device, called 3D-Flow, which is based on an assembly of programmable ASICs. The sec-




Figure 12.6: True p_T spectra of the leading hadron in $B_d^0 \rightarrow \pi^+\pi^-$ events, and inelastic pp interaction events after a cut on the reconstructed E_T of 2.4 GeV.

ond uses processing units based on those designed for the HERA-B pretrigger. An alternative method has been studied, based on a slightly different algorithm using energy clusters defined by 2×2 cells.

3D-Flow: Although not a commercial product, “3D-Flow” [5] is a general-purpose system based on a microprocessor ASIC, specifically designed for High Energy Physics experiments. In a typical application, a 3D-Flow system consists of a set of processors assembled into planar layers (Figure 12.7), each processor being directly associated with one or more detector elements. Given that the algorithm’s execution time in the processors is in general larger than the time separation between successive events, 25 nsec, several layers of processors are stacked together, with successive events being routed to successive layers. The total number of layers required for a given application depends on the total number of clock cycles required to execute the desired algorithm, the clock rate of the processor and the event input rate.

A 3D-Flow processor operating at 80 MHz is envisaged for our application [6]. A 40 MHz input rate and an algorithm requiring 10 instructions, for example, would then imply a stack of five separate layers. Input-output (I/O) of the 3D-Flow operates with 16 bit words. Running at 80 MHz, a maximum of 2 words (32 bits) can be brought in or taken out of the processors when working at the 40 MHz




Figure 12.7: A schematic view of the 3D-Flow trigger processor system.

event input rate.

After processing data within one processor layer, further processing takes place only for the fraction of events where potentially good trigger candidates are found. Information from the specific processors that have found data of interest must be routed down to a single final processor. To achieve this, a stack of processor planes is followed by a sequence of processor layers, each layer containing four times less processors than the previous layer, in a pyramid-like structure. The relevant data trickles through this cascade of processor planes, until it reaches the apex of the pyramid, where it is made available to the final trigger decision logic.

Four processors are integrated into one ASIC structure. Some important characteristics of the processor are as follows:

- A 128 bits-wide instruction word allows concurrent operation of the processor’s internal units (Arithmetic Logic Units, Look Up Tables, I/O busses, etc.) in a single clock cycle.
- A strong I/O capability is built around four bi-directional ports (North, East, South and West) and two mono-directional ports (Top and Bottom). All of the ports can be accessed simultaneously within the same clock cycle. N, E, S, and W ports are used to exchange data between processors associated with neighbouring detector elements within

the same layer. T and B ports receive input data and transmit results of calculations along successive layers.

- A built-in pipelining capability is realized using a “bypass mode”. In bypass mode, a processor will ignore data at its Top port, and automatically transmit it to the Top port of the processor in the next layer. This feature therefore provides an automatic procedure to route the incoming events to the correct layer.

Another important property of the processors is that they can be individually accessed via an RS-232 line by a supervising host. This gives the possibility to perform continuous monitoring of the system performance, to download test programs and verify their results. This, together with the intrinsic properties of programmability, modularity and expandability, constitute the appeal of such a system.

In the case of the LHCb calorimeter trigger, the amount of data to be examined at the input rate of 40 MHz is rather large, namely the ~ 6000 ADC channels from the ECAL and the corresponding preshower data and the pad chamber, amounting to ~ 12000 bits. However, data from most cells can be discarded with simple checks. These must be applied on the full detector granularity, each cell requiring a dedicated processor working in parallel with its neighbours. A small subset of processors, working on reduced event samples, are required to execute the final steps of the algorithm.

In the first stage, the input to the processor is a 16 bit word containing the encoded information generated by the front-end electronics, i.e. 7 bits of calibrated ECAL energy data, 1 bit of preshower information and an 8 bit pattern of the relevant elements from the pad chamber. Since up to 32 bits of input data can be handled by the 80 MHz processor, HCAL information can be easily included if required.

Upon reception of the input word, all of the ~ 6000 processors execute in parallel the following steps:

1. Gather the 8 energy values from the neighbouring cells. This is performed concurrently with the broadcast of the processor’s own energy value to its 8 neighbours.
2. Compute the total energy sum of the 3×3 array.
3. Calculate the energy asymmetries, for subsequent determination of the point of impact; $\Delta_x = E_W - E_E$ and $\Delta_y = E_N - E_S$.

4. Test whether the energy of the central cell is larger than any of its neighbours.
5. Test whether the 3×3 energy sum is above a given E_T threshold.

This set of operations represents the minimum number of steps required to compute all of the quantities needed in the second stage. The second stage is performed by a greatly reduced number of processors, given that the number of clusters is reduced to less than one per event. Simulations show that, by applying an E_T cut of 1.5 GeV, the number of cluster candidates is about 0.3 per single inelastic pp interaction.

These steps can be accomplished in 6 instructions with the 3D-Flow processor. At this point, only the processors that found useful clusters need transmit their data along the pyramid for further processing. The output data is packed into two 16 bit words carrying the following information: 3×3 energy sum, energy of the central cell, Δ_x , Δ_y , the pre-shower and pad chamber bit maps. Preparing the processor for the next cycle and branching requires two more cycles and so the total cycle count is therefore 8. This requires the depth of the first stage of the 3D-Flow system to be 4 layers, each containing 5952 processors.

Before entering the pyramid, a dedicated additional layer of processors is needed to add the cell ID and the bunch crossing ID (time stamp). Given that the data is now sparsified, the 32-bit per bunch crossing constraint on I/O no longer applies. A smaller set of 64 processors located near the pyramid's apex can then perform the final steps of electron and photon recognition:

1. Test whether the central cell contains a sizeable fraction of the total energy.
2. Compute Δ_x/E_9 and Δ_y/E_9 and use results to get bit masks to compare with the pad chamber pattern.
3. If a match is found and the preshower bit is set then a good electron candidate has been found.
4. If there is no match with the pad chamber, a good high- E_T photon candidate can form a trigger, providing a higher threshold condition is satisfied.

The total execution time is the sum of the actual computation times (stage 1 + stage 2) and the worse case transmission time along the pyramid. Our simulations show the total to be less

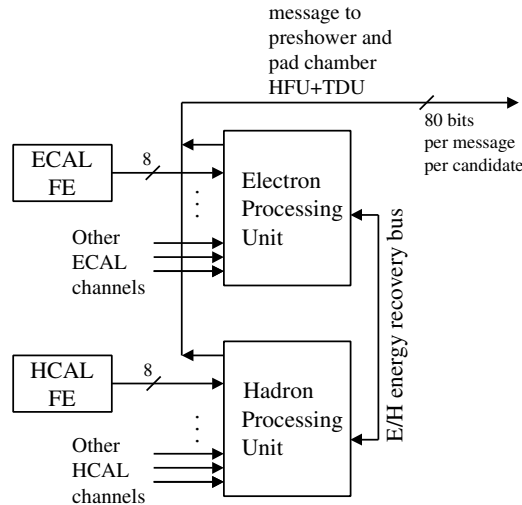


Figure 12.8: Hardware implementation for the ECAL and HCAL Level-0 trigger using HERA-B electron pretrigger processing units.

than $1.5 \mu\text{s}$, which falls within the allocated time of $2 \mu\text{s}$ for the execution of the algorithm.

The operation of the hadron trigger is much more straightforward, since the algorithm requires only the 3×3 energy sums of ECAL and HCAL and threshold checks. Moreover, the total number of HCAL trigger channels is 1/4 of the ECAL. The most economical solution is to have a separate set of processors dedicated to the hadron trigger. Simulation of the hadron trigger shows that the algorithm can be executed in 8 processor instructions requiring a stack of 4 layers with ~ 1500 processors each.

Application-specific hardware: The electromagnetic calorimeter and electron trigger algorithm of the HERA-B experiment are very similar to those of the LHCb experiment. Therefore an implementation of the Level-0 electron trigger is being studied based on processing units similar to those developed for the HERA-B electron pretrigger [7]. The architecture is shown schematically in Figure 12.8.

The digitised energy information already calibrated from ECAL (HCAL) front-end electronics is sent to electron (hadron) processing units [8]. These units produce streams of bits (“messages”) containing the shower energy, calculated on a 3×3 cell matrix, and the centre of gravity of the shower. The message also includes the addresses of Regions Of Interest (ROI), which indicate where to look for




Figure 12.9: Electromagnetic and hadronic Level-0 trigger concept developed on the basis of the HERA-B trigger system. The shaded areas on preshower and pad chambers show schematically ROI addressed by the electron and hadron Level-0 trigger units.

the preshower and pad-chamber hits (Figure 12.9).

When an electromagnetic (hadronic) shower candidate is found, the corresponding cluster energy in HCAL (ECAL) is added to the message. This can be implemented on the HERA-B trigger boards, which extract the total energy information of the 3×3 cell matrix from other boards [8].

In the next step, the messages formed in the ECAL (HCAL) processing units are used to address ROI's in the preshower and pad chambers. This is done by the processing units shown schematically in Figure 12.9. These units [7] are a straightforward extension of the HERA-B Track Finding and Track Decision Units [9]. They are formed by two distinct logic blocks:

- **Hit Finding Unit (HFU):** The hits of the preshower and the pad-chamber are stored in a Dual Ported RAM (DPRAM) using the bunch crossing number as an address. The string of hits addressed by the ROI's are extracted from these memories and sent to a look-up table that gives as output a preshower flag, a pad-chamber flag and a correction to the particle direction evaluated by the ECAL (HCAL) trigger processing unit.
- **Trigger Decision Unit (TDU):** This evaluates a particle ID (electron, photon, hadron) and applies a loose cut on E_T . Information is then sent to the Level-0 decision unit.

In the trigger architecture studied, the look-up table based processors allow great flexibility in the trigger algorithm and in system debugging. A change of the granularity of the detectors can be

handled by increasing (or reducing) the number of processing units. The electronics modules already constructed for HERA-B can be taken as a starting point for the implementation of the system. Implementing an ASIC version of the main logic functions seems to be technically feasible and affordable.

Hardware implementation using 2×2 clusters

An approach to the electron trigger has been studied that places emphasis on ease of implementation, both from the point of view of integrating trigger and front-end electronics and also the connectivity required between processors implementing the trigger algorithms. The basic idea is to reduce the large multiplicity of channels by performing first a local search for hot spots, and then collecting these hot spots to check if the trigger conditions are fulfilled. The “local” processing will be done on the cards housing the readout components (ADC, FIFO, and Level-1 buffer) to avoid extra cabling. A similar approach can be applied to the HCAL trigger.

Algorithm The procedure is somewhat different from the one described in the previous section. Firstly, a cluster consists of 2×2 cells instead of 3×3 . This decreases significantly the number of connections between calorimeter cells that have to be made in the electronics, and gives only a small loss of energy. The fraction of energy in the 2×2 area is about 0.93 for the 4×4 cm² cells and 0.97 for the 8×8 cm² cells, while the shower is fully contained in the 16×16 cm² cells. This leakage effect can be corrected for.

The energy of each cell is converted to transverse energy using the position of the cell's centre. The transverse energies are then added in each 2×2 area, and only these sums are used in the following steps. The local maximum in the transverse energy sums is searched in a small area of the calorimeter, small enough to avoid a large E_T photon or π^0 overlapping the good electron, but large enough to have a manageable number of such areas in total. The local area is chosen to be a square of 8×8 calorimeter cells of the same size, called a “block”. Due to the complex geometry, some blocks are incomplete. The 2×2 sums are also performed across block boundaries, except at the boundaries where the cell size is changing. This seems too difficult to implement, and affects only a small fraction of the calorimeter. The total number of blocks is 106 in the current version.




Figure 12.10: Schematic view of the electron trigger based on 2×2 cluster.

The preshower and the pad chamber information are accessed, but only for the local maximum of each block. The same 2×2 area in the preshower is summed, and compared to a threshold. The presence of a hit in the pad chamber is searched in the corresponding area, or a smaller area corresponding to the vertical half with the highest E_T . This does not require center-of-gravity computation, and is therefore also less sensitive to the relative alignment of the various detectors. The output of each of the 106 blocks consists of a large word with the maximum's address, its E_T and the preshower and pad chamber validation. The outputs of all the blocks are then collected, and the most relevant ones are selected, extracted and sent to a decision unit, where the presence of a high energy photon or electron, or of several medium energy electrons, can be identified. The selection of the most relevant blocks could be performed by applying a threshold, but we present here a method to select the 8 highest blocks, since it can give more flexibility for the final decision.

The hardware implementation and the algorithm were designed simultaneously, with careful attention to the connectivity and to the time constraints of the pipe-line mode of operation. A schematic view of the various elements and their connections is depicted in Figure 12.10.

Local Selection The calorimeter cells are processed in blocks of 8×8 cells. All the cells in a block have the same size. For this study, the assumption is that each group is built of 8 local-cards, placed in a local-crate, each local-card handling 8 channels. A physical-card will most probably contain more than one local-card, and a physical-crate will con-

tain 2 or 4 local-crates, giving a more manageable volume of electronics.

For each channel, the 12 bit ADC output used for the readout is converted to E_T using a fast look-up table. There is one table ($4k \times 8$ bits) per channel, which performs pedestal and gain corrections, energy to transverse energy conversion, and 2×2 leakage correction as described earlier.

The E_T of the cells above, to the right, and above and right, are summed to form the 2×2 transverse energy. This implies that neighbours above, and to the right, are available for every cell. The neighbour above is usually on the same card. Each local-card is connected to its right neighbour within the crate by the backplane. There is a link using a Serial Transceiver from each card to the corresponding card in the upper crate to receive the missing neighbour, and an eight channel link to receive the 8 channels from the crate to the right. This requires the use of two additional cards to receive the information from the crate at the right and to send to the crate at the left.

The sum of the 2×2 8-bit inputs is performed in two steps, using fast adders, first producing the 8 sums of each channel with its vertical neighbours on the card, then adding these partial sums between cards. This can be done in a single ASIC for 8 channels of a card.

The selection of the highest of the 64 clusters in the block is performed in two steps. Firstly the highest of the 8 channels in the card is found, then the highest of these is selected in a dedicated "summary card", using dedicated lines on the backplane. This selection of the "highest in eight" can be implemented on the ASIC used to perform the sums, and the same ASIC re-used for the summary card. The summary card also fetches the information for the selected region from the preshower detector and from the pad chamber. The summary card has then one word of output, built of the address of the highest E_T cluster (5 bits local plus 7 bits for the crate number), the E_T of the cluster (8 bits) and the preshower and pad chamber information (1 bit each), giving a total of 22 bits.

All these cards work in pipeline mode, and 8 to 10 cycles are needed as follows: energy conversion to E_T using look-up table (1 cycle), data collection from neighbouring cards (1 cycle), sum 2×2 cells (2 cycles), select highest on the card (1 cycle), select highest in the crate (1 cycle), fetch the relevant information from the other detectors (2 to 4 cycles).

Main Selection Crate The information from the 106 blocks is sent to the Main Selection Crate,

where a selection of the 8 clusters with the highest E_T takes place. This value of 8 is large enough to always contain all the interesting clusters, and small enough in terms of bandwidth and number of connections. The data are sent by optical link using Fast Serial Transceivers, and stored in a block-receiver.

The fast selection of the highest clusters can be implemented by finding the value of E_T for which at most 8 block-receivers have a higher value. This involves a binary search of this E_T value, using an analogue method to measure the number of selected blocks. A binary search requires as many iterations as bits to be set in E_T plus one more to allow for the blocks passing the last threshold to know if they have to output their information. The 8 pieces of information are then collected, and sent to the next stage of the trigger. Thus the Main Selection Crate requires one cycle for reception, 9 cycles for decision and up to 8 cycles for collection.

Trigger Processor Crate The block of information produced by the Main Selection Crate is sent via optical fibers to one of the processors in the Trigger Processor Crate. Each processor deals with one event to test if there are high E_T electrons or photons above given threshold values. The information is then sent to the Level-0 decision unit. The simplest implementation is to allocate a fixed amount of time for this operation and to distribute in turn each event to one of a fixed number of processors. A latency of $1.0 \mu s$ would require an array of 40 processors.

Performance The performance of this trigger approach is being studied with the standard LHCb simulation program. It reproduces the proposed sequence of operations, with the proposed crate structure, a pad chamber matching the projection of the calorimeter geometry, and the preshower with the same cell size as ECAL.

The only implemented trigger is an electron trigger, requiring at least 5 MeV in the 2×2 area of the preshower, and the presence of a hit in the 1×2 area (vertical half with highest E_T) of the pad chamber. A retention of 3% for inelastic pp interactions is obtained with an E_T cut at 1.80 GeV, giving an efficiency of 54% for pp interactions with $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$ where the final state is reconstructed offline and its flavour correctly tagged.

12.3.2 Muon Trigger

Algorithm and performance

The Level-0 muon trigger must suppress π/K decays while maintaining good efficiency for detecting $B \rightarrow \mu X$. The results of detailed studies made to optimize the muon system and the trigger algorithm are found in [10]. The muon trigger uses all five pad chambers, M1 to M5, of the muon detector.

The Level-0 muon trigger algorithm is executed in the following steps:

1. For a given M3 pad hit, a specified search region in M2, M4 and M5, is opened.
2. If at least one hit is found in each of the four muon stations, the pad hit on M2 closest to the centre of the search window is taken.
3. The x intercept at M1 is calculated using the M2 and M3 combination.
4. The closest pad hit to the x intercept in M1 is found.
5. The y slope is calculated using the M1 and M2 combination, and the y intercept at the plane $z = 0$ is obtained.
6. The y intercept is required to be within a certain distance from the interaction point.
7. The p_T of the muon candidate is calculated from the x slope determined from the M1 and M2 pad hits and from the y slope determined from the M1 pad hit and the nominal interaction point. Finally, a minimum p_T requirement is imposed.

When a positive decision is made by the Level-0 decision unit, slopes in x and y , the M1 pad hit and the trigger p_T are sent to the Level-1 track trigger.

The trigger rates have been calculated using the standard LHCb simulation program based on Fluka [11], i.e. full shower generation in the muon shield with charged hadrons, electrons, photons, and neutrons tracked down to 1 MeV. Figure 12.11 shows the muon trigger p_T resolution achieved with the pad structure described in [10] and the algorithm described above.

Figure 12.12 shows the performance of the muon trigger using two signal channels, pp interactions with $B \rightarrow \mu X$ where the muon hits the muon system and $B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$ where $J/\psi K_S^0$ is reconstructed offline. Signal efficiencies are shown as a function of the retention for inelastic pp interactions. The results are obtained by varying the threshold p_T value.




Figure 12.11: The difference between the p_T measured by the muon trigger and that generated for muons from the semileptonic B-meson decays.




Figure 12.12: Efficiencies for inelastic pp interactions as a function of those for pp interactions with $B \rightarrow \mu + X$ and $B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$ where the muon hit the muon system for the first decay mode and $J/\psi K_S^0$ is reconstructed offline for the second case.

The muon trigger rate due to hadrons, produced by punch-through or overlap with background muons is found to be only 0.2% of the total trigger rate at a p_T cut of 1.0 GeV/c.




Figure 12.13: LHCb muon pad information necessary for a 3D-Flow processor.

Implementation

Two approaches have been studied, one based on the same 3D-Flow system used for the high- p_T trigger study and the other a specific implementation for the LHCb muon algorithm.

3D-Flow: An implementation of the LHCb Level-0 muon trigger has been studied [12] using 3D-Flow ASICs [13, 14]. There are two stacks of processor layers. A layer in the first stack receives all the 45k channels from the muon pad chambers in parallel. Each processor receives 31 pad channels, 5 pads each from M2 to M5 and 11 pads from M1, as 31 bit information from the top port. It exchanges pad information with the neighbouring 8 processors and builds up a group of pads consisting of 9×3 pads each for M5 and M4, 5 pads for M3, 11×3 pads for M2 and 21×3 pads for M1.

A seed hit in M3 is then searched for. If there is a seed, hits are searched for in corresponding 5×3 and 3×3 pads in M5 and M4 respectively in order to form a triple coincidence. The average number of such coincidences is only 26% per bunch crossing.

Once a triple coincidence is found, the M1 to M5 pad information around the M3 seed shown in Figure 12.13 is routed to a layer in the second stack of 3D-flow where only a small number of processors is needed per layer. In the second stack, the rest of the algorithm to select a clean muon and compute its p_T is executed. Results are transmitted to the Level-0 decision unit.

Simulation studies show that 3 layers with ~ 1300 processors each are needed for the first stack. The number of processors in the second stack is very small.

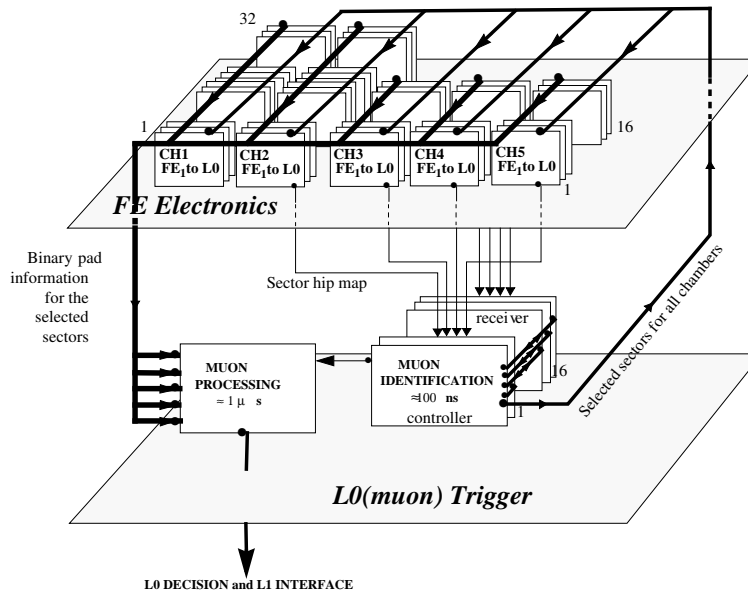


Figure 12.14: A schematic view of the muon trigger implementation with application specific hardware.

Application-specific hardware: In the muon system, the average number of muon tracks in one bunch crossing is ~ 0.12 for inelastic pp interactions at a luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. An alternative solution for the implementation of the Level-0 muon trigger shown in Figure 12.14 [15] has been studied based on this feature.

Data are organized within the muon chamber front-end electronics boards, so that an effective zero suppression is performed. This is done in order to minimize the data flow from the front-end electronics to the trigger processors. This is achieved by defining pad-sectors dimensioned adequately, where an “OR” of all the discriminated pad signals is made. A fast muon track identification algorithm in chambers M2 to M5 is then executed based on the above pad sector structure. A simulation study shows an identification efficiency for muons with $p_T \geq 1 \text{ GeV}/c$ of $99.0 \pm 0.5\%$.

Pad hits from all muon chambers, including M1, are then extracted but only from within those sectors in the vicinity of the identified muon track. This results in an effective data suppression of about 20.

Full muon track processing, as discussed in the algorithm section, is then performed on the extracted pad hits in order to determine the transverse momentum of the muon candidates, and to transfer all relevant information to the Level-0 decision unit. First estimates of the total processing time, from the beginning of the data transfer to the

broadcast of the Level-0 muon decision, are 2.4 to $2.7 \mu\text{s}$, well within the Level-0 fixed latency of $3.2 \mu\text{s}$.

This hardware implementation is simple, rather compact, with a small number of connections. The total number of boards (18) and optical links (89) is reasonably small. It can be built with current technologies, such as FPGA’s, DSP’s and 1 Gbs^{-1} optical links. Debugging, auto-test, calibration and monitoring functions are feasible with the proposed architecture [15].

12.3.3 Pile-up Veto

Algorithm

At the LHC interaction region, the primary vertex is well defined in the transverse plane; $\sigma_x = \sigma_y = \mathcal{O}(100) \mu\text{m}$. Along the beam, it has a spread of $\sigma_z = \sim 5 \text{ cm}$. Therefore, primary vertices created due to multiple pp interactions are spread along the z axis and can be distinguished with a modest vertex resolution.

The z coordinate of a primary vertex (z_{PV}) can be obtained using only the radial information of the tracks. By assuming that the primary vertex is at $x = y = 0$, the radial coordinates of a track from a primary vertex (R_A and R_B) measured at two planes (A and B) perpendicular to the beam and placed at z_A and z_B (Figure 12.15), have the




Figure 12.15: Basic principle of detecting the primary vertex PV .

following relation;

$$\frac{R_B}{R_A} = \frac{z_B - z_{PV}}{z_A - z_{PV}}. \quad (12.1)$$

Points obtained by combining the two coordinates from the same track lie on a line through the origin in the R_A - R_B plane, if the tracks originate from the same primary vertex. The slope of the line is determined by the z position of the primary vertex; z_{PV} . Points made by the two coordinates from different tracks are randomly distributed over the plane.

If points with all the two-coordinate combinations in one event are plotted in the R_A - R_B plane, points from the real tracks line up and the number of lines corresponds to the number of primary vertices. Such lines can be found by scanning the R_A - R_B plane with a line through the origin by varying z_{PV} and counting the number of points on that line for a given z_{PV} . By histogramming the number of points as a function of z_{PV} , a primary vertex appears as a peak in the histogram (Figure 12.15).

To implement the algorithm described above, two planes of silicon detector equipped with circular strips are used. They are placed behind the main vertex detector system at the opposite side of the LHCb interaction region. Each plane is split into six elements covering 60° in azimuth. Each element has 300 circular strips with an effective pitch of 120 to 240 μm depending on the radius. The front-end electronics is placed very close to the detector in the Roman Pot and reads all the 3504 strips within 25 ns. After discrimination, only binary information is sent to off-detector electron-

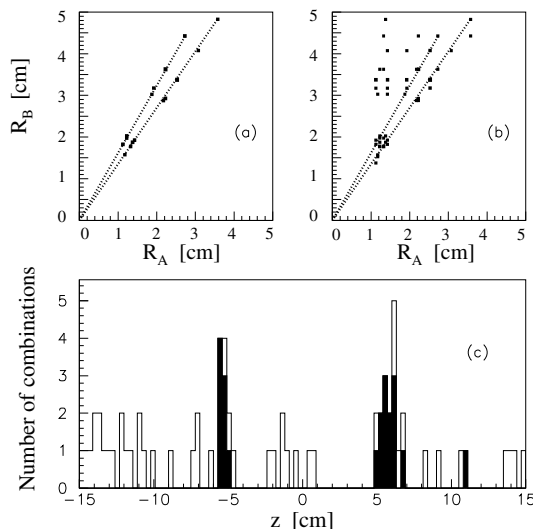


Figure 12.16: Correlation of the measurements in the two planes. A typical event with two vertices is displayed. (a) True combinations, i.e. both hits from the same track, (b) all combinations are shown separately. The slope of a dotted line corresponds to the z -position of a vertex. The z histogram is shown in histogram (c); the shaded bins correspond to true combinations.

ics outside the Roman Pot over 3600 10 m long twisted-pair cables.

Simulated R_A - R_B plots for a double pp interaction are shown in Figure 12.16-a for combinations corresponding to real tracks and Figure 12.16-b for all combinations. The transverse spread of the primary vertex and detector resolution are taken into account in the simulation program. Figure 12.16-c shows the z_{PV} histogram within a defined search window of ± 15 cm. A primary vertex resolution of $\sigma_z \approx 1$ mm is obtained. The two peaks seen in the figure indicate the two pp interactions. The peak height of the second highest peak is sent to the Level-0 decision unit.

Figure 12.17 shows the retention factor for bunch crossings with two pp interactions as a function of that for the single pp interaction bunch crossings. A working point for the LHCb trigger is set to reject 80% of two pp interactions which gives 95% retention for single pp interaction events. More details can be found in [16].

Another parameter that may be used to reject multiple pp interactions is the total energy seen in the calorimeters (E_{tot}). As seen from Figure 12.17, this method is not effective for rejecting double pp interactions due to the large overlap in E_{tot} distributions for single and double pp interactions [2].




Figure 12.17: Retention of double pp interaction events as a function of retention for single pp interaction events obtained by varying the height of the second peak in the pile-up veto trigger. The performance using the total energy in the calorimeters is also shown.

However, it is useful for rejecting bunch crossings with three or more pp interactions. At a luminosity of $5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ they become significant as seen from Figure 12.2.

Implementation

The binary signals of the radial strips are sent out from the Roman Pot to the off-detector electronics (Figure 12.18). The Event Demultiplexer distributes the events to the six vertex finders, each dedicated to one event.

The vertex finder correlates points in the R_A and R_B plane and counts the number of points on the line. The system can handle 30000 points expected within the search window. The procedure is repeated for all 6 pairs of azimuthal detector elements. The resulting histograms are summed and processed for further analysis. A histogram has about 100 bins, which can be represented by a 6-bit number. The two largest peaks are searched for. The Comparator/Encoder extracts the two vertex candidate positions and the height of the second peak. They are sent to the Level-0 decision unit through the Result Multiplexer.

The amount of logic needed for the Vertex Finder is about 350k gates. If it is built in a single integrated circuit, $2 \times 300 = 600$ inputs and about 50 outputs and control lines are required. In the near

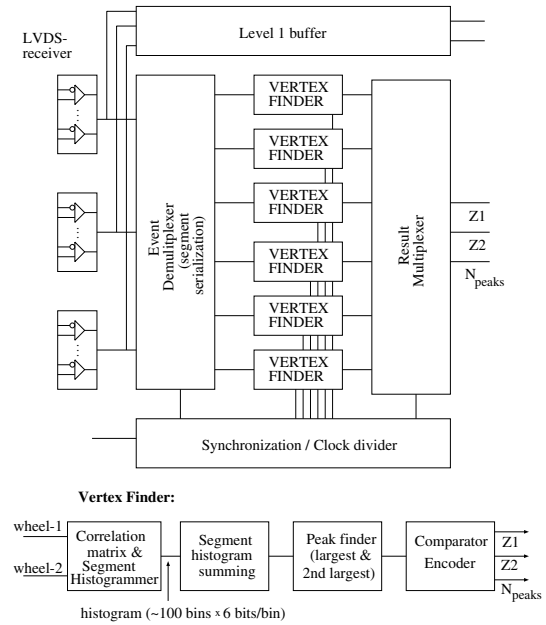


Figure 12.18: Block diagram of the Pile-up Veto.

future commercial FPGAs will be able to satisfy these requirements. For example, devices such as the Altera FLEX 10K250 with 250k gates and 640 input-outputs (I/O) will be available soon. The size of such devices is expected to grow to 2M gates in the year 2000, e.g. for Altera Raphael architectures.

12.3.4 Level-0 Decision Unit

A global decision on whether an event must be kept for the next trigger level is taken by the Level-0 decision unit. The muon high- p_T trigger units send information on the muon candidates: their p_T values, slopes and hit coordinates at the pad chamber. The calorimeter high- E_T trigger units send information on the electron, hadron and photon candidates: their E_T values, slopes and hit coordinates at the pad chamber for hadrons and electrons and at the preshower for photons. In addition, a probability that two pp interactions occurred in the event is sent by the pile-up veto unit.

The event is accepted if there is no pile-up signal and if one of the following conditions is fulfilled.

- $p_T(\mu) > 1.0 \text{ GeV}/c$
- $E_T(e) > 2.34 \text{ GeV}/c$
- $E_T(h) > 2.4 \text{ GeV}/c$
- $E_T(\gamma) > 4.0 \text{ GeV}/c$

- 2 μ 's
- 2 e's
- 1 μ AND 1 e

Implementation of the unit is based on a look-up-table, in which modification and extension of the list can be easily made to match the running conditions and physics requirements.

If the event is kept for the next level, the Level-0 decision unit informs Level-1 of which trigger condition is fulfilled. When the triggering particles are muons, electrons and hadrons, track information is sent to the Level-1 track trigger.

12.3.5 Allocation of trigger rates

Allocation of the trigger rates for the single high- p_T triggers is obtained by optimising the physics performance of the LHCb experiment with the five pilot decay modes introduced in Table 12.1.

$$\begin{aligned} B_d^0 &\rightarrow \pi^+\pi^- + \text{correct tag} \\ B_d^0 &\rightarrow J/\psi(e^+e^-)K_S + \text{correct tag} \\ B_d^0 &\rightarrow J/\psi(\mu^+\mu^-)K_S + \text{correct tag} \\ B_s^0 &\rightarrow D_s^+K^- + \text{correct tag} \\ B_d^0 &\rightarrow \bar{D}^0K^{*0} \end{aligned}$$

The statistical errors on the angles of the unitarity triangles are inversely proportional to the ‘‘trigger power’’, which is defined for a given decay final state f , trigger mode i ($i=e, \mu$ and h) and tagging mode j ($j=e, \mu$ and K) as

$$P_i^j(f) = \sqrt{\epsilon_{\text{trig}}^i(f)\epsilon_{\text{tag}}^{i,j}(f)} \left[1 - 2\omega_{\text{tag}}^{i,j}(f) \right] \quad (12.2)$$

where $\epsilon_{\text{trig}}^i(f)$ is the trigger efficiency for the event with the final state f where f is reconstructed off-line, $\epsilon_{\text{tag}}^{i,j}$ is the tagging efficiency for the triggered event and $\omega_{\text{tag}}^{i,j}$ is the wrong-tag fraction. For the final state f and trigger mode i , the power is then given by

$$P_i(f) = \sqrt{\sum_{j=e,\mu,K} P_i^{j^2}(f)} \quad (12.3)$$

Note that for a self-tagging decay such as $B_d^0 \rightarrow \bar{D}^0K^{*0}$; $P_i(f) = \sqrt{\epsilon_{\text{trig}}^i(f)}$.

Allocation of the rates to different trigger components for a given channel f is then optimized by maximizing the total power

$$P(f) = \sqrt{\sum_{i=e,\mu,h} P_i^2(f)}. \quad (12.4)$$

with a total retention factor of 0.09 for a single inelastic pp interaction.

Figure 12.19 shows the powers $P_i(f)$ for the five pilot channels as a function of retention factor for an ordinary single pp interaction obtained by single electron, muon and hadron high- p_T triggers. The rate allocation for the three triggers is made by optimizing the errors on α , $\gamma - 2\delta\gamma$ and γ without sacrificing the β measurement.

12.4 Level-1 Trigger

12.4.1 Vertex Trigger

Algorithm

A detailed description of the vertex trigger algorithm can be found in [17], only a summary is given here. The vertex detector simulation is described in [18].

The implementation of the track and vertex finding algorithm is simplified if the silicon detectors have an r, ϕ geometry for the following reasons:

- Seeds for the track finding are located in the 12 inner r -measuring sectors. The average charged track multiplicity in each of these sectors is about 1, which limits the number of hit combinations to be tried.
- To first order the r -coordinates define the z -position of the vertices.

The algorithm consists of five consecutive steps. Note that the momentum of the tracks is unknown and the error estimation of the track parameter is derived from purely geometric considerations, assuming straight tracks.

1. Track finding in the r - z projection: The track search starts in the inner r -subsectors. Three collinear hits in successive stations define a track. Several segments of the same track may be found in different triplets. A rejection process removes duplicate segments, leaving the one with a cluster closest to the beam. Those tracks are referred to as 2D-tracks.

2. Calculation of the primary vertex position: An estimate of the longitudinal position of the primary vertex is made by making a histogram for the z -coordinates of the two track vertices. A resolution of $80 \mu\text{m}$ is obtained. Intersecting the tracks with a plane perpendicular to the beam at the position of the primary vertex determines the x and y position of the primary vertex, with an error of $20 \mu\text{m}$.




Figure 12.19: The statistical power as a function of retention factor for inelastic pp interactions with a single electron, muon and hadron high- p_T trigger for $B_d^0 \rightarrow \pi^+\pi^-$, $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$, $B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$, $B_s^0 \rightarrow D_s^+K^-$, and $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ decays. Trigger-power is defined in the text.

3. Impact parameter determination: For each 2D-track in the forward direction, the impact parameter with respect to the primary vertex is calculated. A probability that this track originates from the primary vertex is defined as the fraction of tracks in inelastic pp interactions having an impact parameter exceeding the calculated value. Look-up tables containing these probabilities are determined as a function of the track length, the extrapolation distance to the primary vertex and the polar angle. Bunch crossings with multiple pp interactions passing the Level-0 trigger are rejected by requiring an upper limit on the number of tracks with a large impact parameter with respect to one of the primary vertices. The dominant sources for tracks with a large impact parameter not coming from b decays are low momentum tracks suffering from multiple scattering, the decay products of strange hadrons, and γ -conversions.

4. Reconstruction of tracks with large impact parameter: For 2D-tracks having a probability of less than 20% of originating from the primary vertex, and an impact parameter less than 3 mm, the ϕ cluster information is now added to reconstruct the track in three-dimensions (3D-track). The fact that the ϕ strips are rotated by a small angle helps to resolve hit ambiguities. Remaining

ambiguities are resolved by accepting only the combination with the smallest x - y impact parameter which is defined as the distance between the primary vertex and the crossing point of the candidate on the x - y plane placed at the primary vertex. The reconstruction efficiencies obtained are about 98% for the 2D-tracks and 95% for the 3D-tracks.

5. Search for two track vertices: The algorithm finds secondary vertices which are separated significantly from the primary vertex.

All vertices with a distance between two 3D-tracks smaller than $200\ \mu\text{m}$ are accepted. For each accepted secondary vertex, the smallest probability that any of the two tracks originates from the primary vertex is assigned. The product of all the assigned probabilities define the total event probability to be non-b events. This value is sent to the Level-1 decision unit.

Performance

The Level-1 vertex trigger algorithm is tailored to meet the selection criteria proposed for offline use in the reconstruction of B-mesons. The decay mode $B_d^0 \rightarrow \pi^+\pi^-$ is used to demonstrate the performance of the algorithm, since this is one of the most demanding modes.




Figure 12.20: The retention of inelastic pp interactions against the trigger efficiency of $B_d^0 \rightarrow \pi^+\pi^-$ events.

In order to optimize the vertex trigger performance for B events that can be reconstructed offline, only events passing the standard offline cuts are used with Level-0 applied. By varying the cut on the distance between primary and secondary vertices, and on the number of secondary vertices, the trigger efficiency can be determined as a function of the retention of non-b events (Figure 12.20). A 4% retention for inelastic pp interactions is achieved with a 45% efficiency for $B_d^0 \rightarrow \pi^+\pi^-$ events.

Noise study

Because of the high radiation level, we expect a decrease with time in the signal to noise ratio in the silicon detectors. In order to keep the single hit efficiency constant, the thresholds have to be lowered and this results in an increase in the number of ghost hits. The performance of the Level-1 trigger has been studied under different noise conditions i.e. from no noise up to a noise level which is three times higher than the one expected after one year of operation ($\sim 1 \times 10^{14}$ particles/cm² at a radius of 1 cm). For a 4.4% retention of inelastic pp interactions, an efficiency of 44% is obtained for $B_d^0 \rightarrow \pi^+\pi^-$ events [17], indicating no significant change in the trigger performance.

Table 12.3: Results of benchmarking the vertex trigger algorithm

Algorithm step	Time [μ s]
Triplet finding	417
Triplet removal	117
Primary vertex	139
Impact parameter	10
3D-tracks	108
Secondary vertex	52
TOTAL	843

Effect of detector misalignments

The effect of misalignment of the vertex detector has been studied [19]. The vertex trigger is insensitive to the beam position within a few 100 μ m. If individual detector elements are misaligned by more than 10 μ m perpendicular to the beam or by more than 50 μ m parallel to the beam, then new alignment constants are required.

12.4.2 Implementation of the Vertex Trigger

Analysis of the requirements

The detector simulation [18] shows that the mean number of clusters per event is about 600. For estimates on data bandwidth we have assumed 400 hits per event due to noise, giving a total of 1000 clusters per event on average. Since a 2-byte number can represent each cluster within a station, the average size of an event is 2000 bytes and the average input data rate to the vertex trigger system is 2 GB/s. The maximum event size has been taken to be 4000 bytes which corresponds to 2000 clusters per event since the distribution of the number of clusters does not have long tails.

All the steps of the vertex algorithm have been benchmarked. Table 12.3 shows results for the average processing time of the different steps using an Intel 180 MHz Pentium-based PC running Windows NT 4.0, and compiled with Visual C++ version 5.0. Benchmarks give a CPU power of 140 MIPS for this processor.

The average processing time defines the number of processors required to implement the algorithm. To implement the complete algorithm at an input rate of 1 MHz, 117 processors would be required, each having 1000 MIPS processing power, providing the system is designed to ensure that the

processors are constantly kept busy.

The maximum processing time, which is of importance for the depth in the Level-1 buffers, is determined by the tails of the processing time distribution. If one implements the complete algorithm in software, then the maximum processing time for the vertex trigger is of the order of 300 μ s using 1000 MIPS processors with the current version of the software. By optimizing the code and algorithm, we envisage reducing this time by a sizable factor with no loss in performance.

Design considerations

Since the 3D track reconstruction is complicated, it must be executed in software on a general-purpose processor. However we have studied in detail the possibility of using preprocessors to execute the initial steps of the algorithm before the data are sent to the CPUs. This would reduce the bandwidth requirement for data collection and the total processing time of the Level-1 trigger, but at the expense of a less flexible algorithm and a reduced precision.

In this implementation study [20], the r -cluster information is preprocessed in parallel by an array of special purpose processors (FPGA's) and the r data volume is reduced by a factor 3 to 4 by sending the list of 2D tracks instead of the initial r -clusters. Finding 2D tracks in hardware reduces the processing time by about 40%.

However, the benefit of such a design seems not to balance the loss in flexibility and the reduced precision. The performed simulations of minimum bias rejections and signal efficiencies show that it is necessary to use the best possible precision, and to have a flexible algorithm that can deal with misalignments and inefficiencies. Therefore, a solution without preprocessing has been devised.

Two possible protocols have been considered for controlling the transfer of data to the processors. The first uses a data on demand scheme, where a processor requests any data that it requires for a specific event during the execution of the trigger algorithm. For example, the processor would request the ϕ data for the sectors where 2D-tracks have been found and not for the others. This approach minimizes data flow but has a complicated control sequence. Our calculations have shown that the gain in event bandwidth does not compensate for the resulting complexity of the design. An alternative scheme in which all data for every event is pushed through to the processors is therefore proposed.

Proposed solution

All the trigger data belonging to the same event are sent to a single processor in a farm of processors. Details of the design can be found in [21]. The trigger algorithm is implemented as a software program running in the processor and has the complete data available in its memory. The main problem to be faced is an event building rate of 1 MHz and an aggregate sustained throughput of 2 GB/s. A diagram of the implementation architecture can be seen in Figure 12.21 where the event builder function has been implemented using a special switch made of a grid of dual-port memories.

Trigger data blocks containing the r and ϕ coordinates of the strip clusters are received by the trigger system from the front-end electronics sitting near the detector. We have assumed that the number of data sources is 17, corresponding to the number of stations. A header containing the event identifier, the trigger number, and the station number precede each data block. The average size of a data block is 100 bytes.

All the data blocks are pushed from the 17 sources into the memory of the processor where the algorithm will be executed. The protocol is kept to a minimum by having enough data buffers in the final processor and in the intermediate stages to store sufficient events. A mechanism for handling buffer overflow conditions is foreseen.

The implementation of the switch we are considering is based on a series of independent horizontal and vertical buses, with one interface per bus and a series of dual port memories (DPMs) at every bus intersection. The interface to the horizontal or vertical bus is the only master on the bus, thus making the design of the bus simple. Both buses have similar requirements in terms of throughput. If one assumes 32 bit wide synchronous buses, the required frequency is 25 MHz for a maximum data throughput of 100 MByte/s. Designing the buses for an average load of 50% requires 50 MHz, which is easily feasible with current technology.

The dual-port memories situated at the crossings for each vertical and horizontal bus will be used to store temporarily the event fragments during the event building. The interface to the vertical bus (Front-End Receiver or FER) pushes the event fragments into addresses, which map different dual-port memories for each event in a round-robin fashion. Each FER acts asynchronously with respect to the others.

The role of the horizontal interface (Sub-Farm Controller or SFC), is to transfer all data fragments of the same event into a CPU memory. The hori-

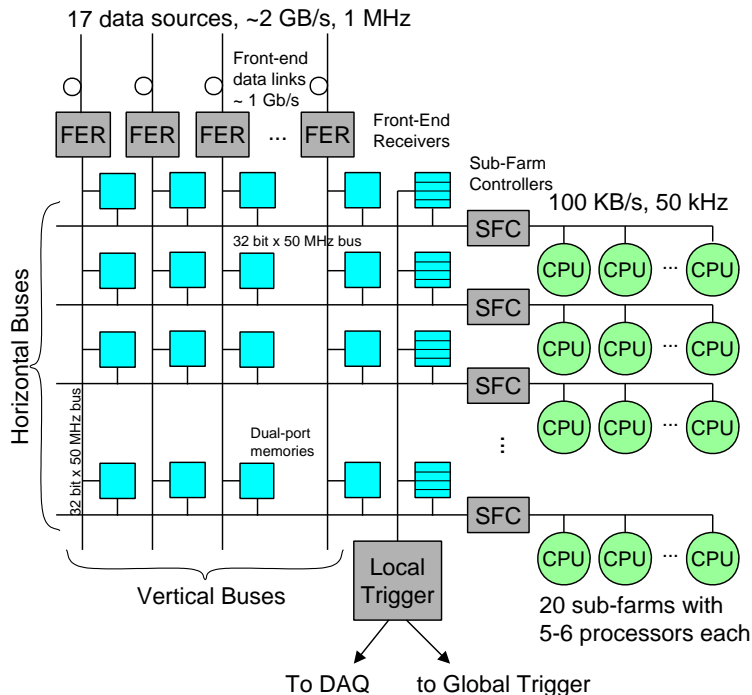


Figure 12.21: A schematic view of the vertex trigger implementation.

zontal transfer can be initiated as soon as the first data fragment arrives or after all of them are there. The latter requires waiting for the longest fragment (maximum 3 times the average). Mechanisms for checking if the complete event is present can be implemented with a single bus cycle. A simple local algorithm using the number of events queued on each processor determines the final processor to which the data are pushed.

The processors are organized into 20 sub-farms. Each sub-farm runs independently of the others. The rate of input events to each sub-farm is 50 kHz. Each processor in the sub-farm will be a commercially available processor (PC-like) and will be interconnected with a high performance link to the horizontal bus interface. The minimum required aggregate throughput for such an interconnect is 100 MByte/s. Today several technologies, such as SCI and Myrinet, are able to provide the required throughput.

The sub-farm controller acts as a bridge between a horizontal bus in the switch and the network connecting all the processors on the sub-farm. Its main function is to assemble event fragments from the dual-port memories and send complete events to the processors. It can also be used to collect the results from the processors and put them into a FIFO to be picked up by the local trigger unit

using a vertical bus. In this case, since the controller knows which processors have finished processing, the assignment of a processor can take into account the number of events scheduled and completed. This gives a simple way of balancing the load on each sub-farm.

The local trigger unit is a specialized hardware processor that receives all the trigger decisions from all the processors and sorts them into chronological order. It also provides an interface to the Level-1 decision unit and to the DAQ system, thus providing a means of monitoring the system performance. The DAQ will read out the summary of the results of the trigger processing from the local trigger unit. This processor must work at 1 MHz input and output rate, though the amount of data is relatively small. The minimum data required will be the event identifier and a decision bit.

Since the overall latency must be kept manageable, an upper limit will be set on the total time available for taking the Level-1 decision. In other words, events in the tail of the latency distribution must be accepted or rejected as soon as this time limit expires. However first indications are that the rate corresponding to the number of events in the tail is small compared to the retention rate for inelastic pp interactions (4%). This can therefore be resolved by simply accepting all these events.

Table 12.4: Estimates of the average latencies of the different phases

	Min. [μ s]	Typical [μ s]	Max. [μ s]
Coordinate encoding	3	3	3
Write into memories	0.2	1	3
Wait complete event	0.5	2	3
Read from memories	4	20	50
Algorithm execution	30	120	250
Collect results	5	5	5
TOTAL	43	151	314

Performance

Since algorithms are executed in software on standard processors the performance can be evaluated off-line. Table 12.4 shows our first estimate for the latencies of the different steps in the trigger processing. More accurate numbers will be obtained using the Foresight modelling tool to simulate the system behaviour.

One of the strengths of this design is its flexibility. There are no constraints on the algorithm, since all data are available to the processor for every event. The architecture is scalable. It can easily incorporate more data sources or more processing units without any redesign. The system can be easily upgraded by replacing CPUs with faster ones.

12.4.3 Track Trigger

A significant fraction of the Level-0 high- p_T triggers in inelastic pp interactions do not originate from single particles coming from the primary vertex region, but are due to secondary interactions in the detector material, to decays, and to overlapping showers. In all of these cases, the addition of information from the tracking chambers gives a means of reducing this background at Level-1. Since the Level-0 trigger information is used, full pattern recognition is not required, and only a few track candidates are processed in each event.

The approach taken here follows similar ideas adopted by the HERA-B collaboration, however with greatly reduced requirements on trigger latencies and occupancies.

Algorithm and performance

A graphical view of the track trigger algorithm is shown in Figure 12.22. The algorithm can be seen

as a 3-step process:

1. Track seed selection.
2. Pattern recognition and momentum measurement.
3. Trigger decision.

These steps are briefly described below. A detailed account, as well as the corresponding performance studies, are given in [22].

In the first step of the algorithm, track candidates are selected from the patterns of high- p_T particles found by Level-0. Using these “seeds” to define initial track parameters, hits are searched for in the tracking station closest to the calorimeters. Hit coordinates in 3 dimensions are used to improve on the track parameters by applying a simplified Kalman filter technique, which uses only diagonal parts of the covariance matrices of the track parameters. The track is extrapolated to the next tracking station towards the vertex, where the next hit is searched for. This sequence of backward extrapolation, hit finding and track parameter updating is repeated station by station towards the vertex. The size of the search window is influenced by the accuracy of track parameters, the distance between adjacent stations, multiple scattering and the magnetic field. Chamber inefficiencies are allowed for by accepting tracks with up to two tracking stations without hits in the search window. A track is also kept for the final trigger decision if no hit is to be expected because of the limited inner acceptance of the tracking chambers. Both inner and outer trackers are considered in the track search.

Sufficient rejection power for inelastic pp interactions is obtained by extrapolation to station 6, where the track momentum is measured with a resolution of $\sigma_p/p \approx 3\%$ independent of the momentum. The final trigger decision is made by applying a p_T threshold.

Simulation results for the retention rates for inelastic pp interaction events are shown in Figure 12.23 as functions of b-events with specific B -meson decays, $B^0 \rightarrow J/\psi \rightarrow \mu^+\mu^-$, $B^0 \rightarrow J/\psi \rightarrow e^+e^-$ and $B_d^0 \rightarrow \pi^+\pi^-$ triggered by the muon, electron and hadron Level-0 trigger respectively. Plots start from points which are obtained by the Level-0 trigger with particular values of p_T thresholds.

The track trigger is able to reduce the number of inelastic pp interactions by as much as a factor of ~ 10 with a loss for signal events of $\sim 30\%$.




Figure 12.22: Schematic view of the Level-1 track trigger.

Implementation

The tracking trigger algorithm described in the previous section can be implemented in the following manner. Each tracking station represents an independent processing unit, which consists of an input unit, several geographical segment units and an output unit.

The input unit receives track candidates as seeds, either from the previous station or from the Level-0 decision unit. Processing of each candidate includes the calculation of the expected hit position and search window size for this tracking station. Then the track candidate is sent to all segments that overlap with its search window.

Each segment unit has two tasks. First it has to prepare the hit information from the tracking chamber to which it is connected. Inner tracking readout channels are added together to make a granularity of 1 mm. In the outer tracking, only wire hit signals are used without TDC information, while each pair of staggered planes is combined giving an effective granularity of twice the number of wires. In both inner and outer tracking, hit information of two consecutive bunch crossings has to be stored for each of the three orientations 0, +5 and -5 degrees for later use in the track finding algorithm together with its Level-0 trigger number.

The second task of the segment processing is to receive a track candidate, and to search for hits that lie within the search window and originate from the correct Level-0 trigger number that were stored earlier. These hits are then combined to two

dimensional (x and y) coordinates using the stereo information of the tracking systems and this information is added to the track candidate, which is finally sent to the output processing unit.

The output processing unit recombines the track candidates of the same seed origin sent to different segment units, and selects the hit coordinate with the smallest distance to the track. If a hit has been found, the track parameters need to be improved, otherwise the track quality and the geometrical acceptance has to be checked and the track is possibly dropped.

The 256 μs maximum latency of the Level-1 trigger means that the input, output and most of the segment processing can be implemented by digital signal processors (DSP's). The data preprocessing in the segment unit can be implemented in a more application-specific hardware. Typically a few μs processing time is allocated for each of the steps of input, output and segment operations. Assuming 10 seeds, a sustained throughput of 10 MHz is required at the nominal Level-0 trigger rate of 1 MHz.

The proposed upgrade program of the H1 detector at HERA [23] contains a new tracking vertex trigger, which has similar requirements for throughput, latency and hit combining algorithm in the segment unit. The LHCb tracking trigger will profit a great deal from the H1 development work.




Figure 12.23: Efficiencies of the Level-1 track trigger for inelastic pp events passed by Level-0, as functions of those for events with specific B -decay modes: $B^0 \rightarrow J/\psi \rightarrow \mu^+\mu^-$ (top), $B^0 \rightarrow J/\psi \rightarrow e^+e^-$ (middle), $B_d^0 \rightarrow \pi^+\pi^-$ (bottom). Each data point corresponds to a different p_T cut, as described in the text. The open circles indicate the efficiencies when only the Level-0 trigger is applied. The data points show the efficiencies if the track trigger is applied using p_T thresholds after station 6; they are (in GeV) 0 (i.e. only pattern recognition), 0.8, 1.0, 1.2, 1.4, 1.6, and 1.8 for muon seeds, 0, 0.8, 1.0, 1.2, 1.3, 1.4, and 1.8 for electron seeds, and 0., 2.0, 2.2, 2.4, 3.0, 4.0 and 5.0 for hadron trigger seeds.

12.4.4 Level-1 Decision Unit

Event selection is made by the Level-1 decision unit based on information from the vertex and track triggers. It also receives the trigger condition fulfilled in Level-0. This allows optimal combinations of the vertex and track triggers to be made for different Level-0 trigger conditions. For example, an event with two high- p_T muons of $p_T > 2$ GeV/ c each would be kept irrespective of the vertex trigger information if the two muons are confirmed by the track trigger.

As is the case for the Level-0 decision unit, the implementation is based on a look-up table where modification and extension of the trigger requirements can be easily made to match the running conditions and physics requirements.

12.5 Level-2 Trigger

The main Level-2 algorithm is designed to match the vertex information provided by the silicon detector with the momentum information provided by the tracking stations of the spectrometer [24]. Fake displaced secondary vertices can be generated in non- b events by low momentum tracks with significant multiple scattering. Knowing the momentum of each track, the error of the corresponding impact parameter can be computed properly and the rate of fake displaced secondary vertices can be substantially reduced [24].

Since Level-2 is implemented by a farm of commercial processors, the algorithm can be continuously improved. However it can already be shown that the desired reduction of the rate, a factor of 8, is possible without a substantial loss in the B -meson signals which are used in offline analyses.

12.5.1 Event reconstruction

The Level-2 algorithm begins by reconstructing three-dimensional tracks in the vertex detector, in both forward and backward hemispheres. This is done with the same simple algorithm used in Level-1. On average 45 tracks are found per event, of which 18% have at least one incorrectly assigned coordinate. This percentage is reduced to 11% by rejecting tracks which share common coordinates or which have fewer than the expected number of vertex detector hits. For each track, the momentum is initially taken to be proportional to the measured value of $\sin^{-1}\theta$, where θ is the polar angle of the track and the constant of proportionality is determined from simulation studies. This approximation relies on the rather narrow transverse momentum distribution of the tracks. A resolution σ_d on the impact parameter d_0 is then estimated taking into account contributions from both detector resolution and multiple scattering. Finally the primary vertex is reconstructed by minimizing the impact parameter significances d_0/σ_d of all the tracks. This yields a primary vertex resolution of $\sigma_x = \sigma_y = 9$ μm and $\sigma_z = 38$ μm .

Tracks in the forward hemisphere with an impact parameter greater than twice the estimated resolution (about 6 tracks per event), are extrapolated as a straight line up to the first tracking station (in front of RICH 1). A Kalman filter technique with the full tracking detector information is then applied to extrapolate these tracks down to station 5 (well inside the magnetic field). The hit closest to the predicted point is used to update the track parameters at each tracking station.

With the requirements that at least four hits be associated to the track in the first five tracking stations (each tracking station can have up to 12 hits), and at least one hit inside the magnetic field (stations 3 to 5), about 90% of the high momentum ($p > 2$ GeV) tracks are extended. The efficiency to extend low momentum tracks ($p < 2$ GeV) is significantly smaller (76%), which, in fact, helps in the selection discussed later. The impact parameter resolution σ_d is then recalculated for these tracks using their measured momenta.

The relative resolution on $1/p$ is ~ 0.01 , with small tails due to the tendency to overestimate the momentum for low momentum tracks. Since the overall Level-2 trigger performance does not deteriorate even with a resolution ten times worse, there is some room to simplify the tracking algorithm. The estimated errors on the fitted parameters are correct, except for very low momentum tracks where there is a tendency to underestimate the errors.

12.5.2 Event selection

Events are kept if there are at least three tracks with an impact parameter significance of $d_0/\sigma_d > 3$. Background from K_S^0 and Λ decays is reduced by also requiring $d_0 < 2$ mm.

Table 12.5 shows the Level-2 trigger performance for events which have already passed both the Level-0 and Level-1 triggers. The efficiencies quoted the exclusive b decay channels are obtained using events where the final states are selected by the offline physics analyses, and flavour tagged if necessary. The efficiencies given for inclusive $udsc$ or b production are obtained using events generated over the full 4π solid angle. Level-2 achieves the required rejection factor for $udsc$ events. It also tends to reject inclusive $b\bar{b}$ events where the b-hadron decay products are only partially inside the detector acceptance. This explains the lower efficiency quoted for these events. Few signal B-meson events are lost, although it would be desirable to boost the efficiency for low multiplicity decays such as $B_d^0 \rightarrow \pi^+\pi^-$. This can probably be achieved by also keeping events with only two large impact parameter tracks, providing that one of those tracks has a large transverse momentum.

12.5.3 CPU requirements

With a 1000 MIPS CPU, the time required to reconstruct 45 tracks per event in three-dimensions using the Level-1 algorithm is 0.3 ms. The primary vertex reconstruction takes negligible time relative

Table 12.5: Efficiency of the Level-2 trigger for events previously selected by the Level-0 and Level-1 triggers and, in the case of the exclusive decays, also found by the offline analyses.

	Efficiency (%)
$udsc \rightarrow X$	6.5
$b \rightarrow X$	66
$B_d^0 \rightarrow \pi^+\pi^-$	83
$B_d^0 \rightarrow J/\psi(e^+e^-)K_S$	81
$B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$	81
$B_s^0 \rightarrow D_s^-\pi^+$	92
$B_d^0 \rightarrow \bar{D}^0K^{*0}$	95

to this. Most of the Level-2 CPU requirement comes from the momentum measurement, which takes about 6 ms per track. This algorithm is not yet optimized for trigger use and should easily be improved. Once momenta are measured, the time needed for the decision is negligible. Therefore, our goal of achieving a total latency of ~ 10 ms/event appears to be feasible.

12.6 Level-3 Trigger

12.6.1 Classification of signal channels

The exclusive b-hadron decay modes are divided into four topological classes with distinctive signatures suitable for the identification by Level-3. A fifth class is reserved to encompass non b-physics.

1. Charged two-body decays

A selection of two-track vertices compatible with the B-meson mass isolates decays such as $B_{d,s}^0 \rightarrow \pi^+\pi^-$, $B_{d,s}^0 \rightarrow K^+\pi^-$, $B_{d,s}^0 \rightarrow K^+K^-$ and $B_{d,s}^0 \rightarrow \mu^+\mu^-$.

2. Dilepton decays

Modes such as $B_d^0 \rightarrow J/\psi K_S^0$, $B_d^0 \rightarrow J/\psi K^{*0}$, $B_s^0 \rightarrow J/\psi\phi$ are all characterized by a pair of leptons from a displaced decay vertex compatible with the J/ψ mass.

3. Low multiplicity decays with neutrals

Channels including $B_d^0 \rightarrow \rho^+\pi^-$, $\rho^0\pi^0$ and $B_{d,s}^0 \rightarrow K^{*0}\gamma$ are seen as two-track decay vertices which, with the addition of a π^0 or photon, give the B-meson mass.

4. Decays with D mesons

To select $B_d^0 \rightarrow \bar{D}^0 K^{*0}$, $B_d^0 \rightarrow D^{*+} \pi^-$, $B_s^0 \rightarrow D_s \pi / D_s K$ requires the reconstruction of D/D_s mesons. If available, particle identification is a valuable tool for many of these decay modes.

5. Non b-physics channels

Other physics goals, for instance measurements in the τ and charm systems requiring specific algorithms.

Independently of the topology, all events can be classified as either those which require a flavour tag to be useful in the physics analysis, or those which do not.

12.6.2 Algorithms

Independent sets of algorithms are used to select each topology. Class 4 decays need a further subdivision as they contains the most diverse event structures. Under the assumption that different algorithms reject different background events, each algorithm should have a suppression factor of 100 to 200 in order to achieve the required output rate. The exact factors vary from channel to channel, adjusted by their physics weight.

The filter algorithm is largely based on the offline analysis, but with looser cuts. Due to time limitation, reconstructed tracks have a slightly worse resolution than that available for the offline. The fact that offline calibration and alignment constants are not available also has to be considered. Therefore, the selection windows are chosen as broad as possible for the highest efficiency and sufficient safety for understanding systematic effects. Particle identification can be available for selected tracks only [25].

Two different scenarios have been considered, depending on whether or not particle identification from the RICH information can be used at Level-3.

Figure 12.24 shows schematically a data flow diagram for Level-3 exploiting RICH particle identification. In step-1, the full pattern recognition and the track parameter fit are performed using the vertex detector and the tracking stations. A simplified Kalman filtering technique with an approximation for transporting tracks is used for an increased speed with a reduced resolution. A momentum resolution, σ_p/p , of 1% meets the Level-3 requirements. Simultaneously, lepton candidates are searched for in the muon chambers and the electromagnetic calorimeter.

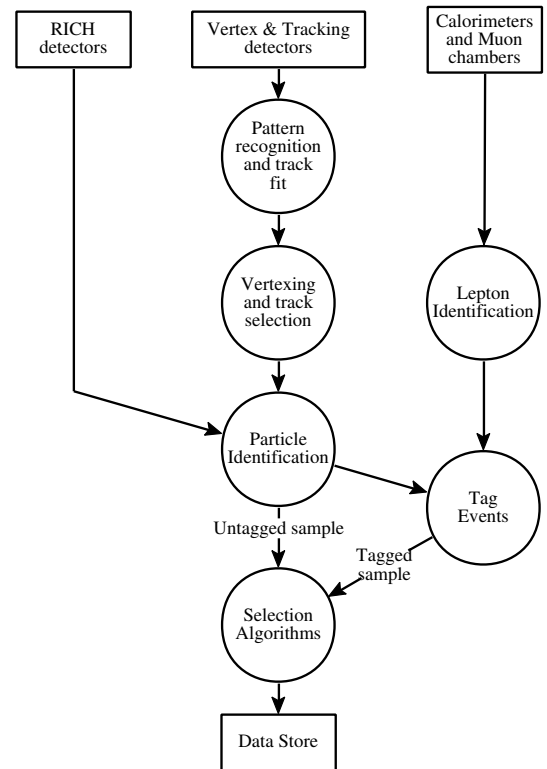


Figure 12.24: Data-flow diagram for Level-3 trigger

The fitted tracks are then subject to a precise vertex reconstruction and to a common preselection, based on the impact parameters for example, which provides a reduced set of tracks as input for the RICH particle identification. A fast RICH algorithm picks up single tracks and searches for the kaon pattern with a knowledge of the track momentum.

Finally, the algorithm identifies a flavour tag (lepton or kaon), and then reconstructs the invariant mass and submasses of the b-hadron decay chain of interest. The self-tagging decay modes together with the non-b channels are selected independently of any tag.

The alternative option follows the same scheme, but misses the RICH particle identification and the tagging step.

Detailed studies made for $B_d^0 \rightarrow \pi^+ \pi^-$, $B_s^0 \rightarrow D_s^- \pi^+$ and $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ show that a global rejection factor of 25 can be reached with or without the RICH. However, the advantage of using kaon identification is twofold. Firstly, events are selected by demanding a good flavour tag. Secondly, combinatorial background is suppressed for the many-body decay modes involving kaons, such as $B_d^0 \rightarrow \bar{D}^0 K^{*0}$

and $B_s^0 \rightarrow D_s \pi$. This leads to an additional safety factor of about 5 to 10. However, it has to be proven that the RICH algorithm, which works well in simulation, provides sufficient robustness to be applied to the real data.

The execution time is dominated by the pattern recognition and track parameter fit. Using the simplified Kalman filter, and with further optimisations of the code, the goal of a ~ 200 ms execution time (normalised to a 1000 MIPS CPU), should be achieved. This gives a requirement of 10^6 MIPS for the CPU farm, as far as Level-3 is concerned.

References

- [1] The LHC Study Group, "THE LARGE HADRON COLLIDER Conceptual Design Report", CERN/AC/95-05(LHC), 1995.
- [2] "Choice of running Luminosity for LHC-B, and Performance of Pile-up Tag" by G. Wilkinson and N. Zaitsev, LHCb 97-014 TRIG.
- [3] S.Conetti, I.Korolko and G.Wilkinson, "Performance of the LHCb Calorimeter Triggers", LHCb 98-14
- [4] G. Kostina and P. Pakhlov, "High E_T Photon Trigger", LHCb/97-015
- [5] D.Crosetto, "Massively Parallel Processing System with 3D-Flow Processor, Proc. of First International Conference on Massively Parallel Computing Systems", Ischia, Italy, May 2-6, 1994, pp. 355-374, IEEE Computer Society Press 0-8186-6322-7/94
- [6] S.Conetti and D.Crosetto, "The LHCb Calorimeter Trigger and its implementation with the 3D-Flow system", LHCb 98-13
- [7] V. Alberico et al., "The HERA-B electromagnetic pre-trigger and its possible adaptations to the LHCb Level-0 calorimeter trigger", LHCb NOTE 98-034 TRIG
- [8] C.Baldanza et al., "HERA-B Ecal Pretrigger Board Description", HERA-B NOTE 97-178, Trigger 97-013
- [9] E. Hartouni et al., "HERA-B An Experiment to Study CP Violation in the B System Using an Internal Target at the HERA Proton Ring", HERA-B Technical Design Report, DESY-PRC 95/01, January 1995
- [10] G. Corti and B. Cox, "The LHC-B Muon Trigger", LHCb 97-01 TRIG 97-01, (Jan, 1997); M. Borkovsky, A. Tsaregorodtsev, and A. Vorobyov, "Study of the LHC-B Muon Trigger", LHCb 97-07 TRIG
- [11] A. Fasso et al., "FLUKA92" Proc. of the Workshop on Simulating Accelerator Radiation Environments, Santa Fe, USA, 1993
- [12] G. Corti, B. Cox, and D. Crosetto, "An Implementation of the L0 Muon Trigger Using the 3D Flow", LHCb 97- TRIG.
- [13] D. Crosetto and S. Conetti, "Implementing the Level-0 Trigger", Proceedings of the IEEE Real Time 95 Conference, IEEE Trans. Nucl. Sc. **43**. 170(1996).
- [14] D. Crosetto, "High Speed Parallel, Pipelined Processor Architecture for Front End Electronics and Method of Use Thereof", LHC-B-96-02
- [15] E. Aslanides, B. Dinkespiler, R. Le Gac, M. Menouni, R. Potheau, "An alternative architecture of the Level-0 μ processor", LHC-B/97-024
- [16] "The use of Silicon detectors for fast primary vertex reconstruction and pile-up rejection.", M. van Beuzekom et.al., LHC-B 97-016 TRIG.
- [17] T. Ruf, "The L1 Vertex Trigger Algorithm and its Performance", LHCb 98-006, TRIG
- [18] D. Steele, "The $r \phi$ Silicon Vertex Detector Simulation as Implemented in SICB", LHCb 98-004, TRAC
- [19] M. Koratzinos and T. Ruf, "Vertex Detector trigger alignment requirements", LHCb 98-012 TRIG
- [20] H. Müller, "Vertex Trigger Implementation using shared memory technology", LHCb 98-33, TRIG
- [21] M. Koratzinos and P. Mato, "Vertex Trigger Implementation using Commercial Processors", LHCb 98-022, TRIG
- [22] M.Cuje et al., "The LHCb Track Trigger", LHCb 98-021
- [23] M. Cuje et al, "H1 Luminosity Upgrade 2000, CIP and Level-1 Vertex Trigger", proposal submitted to the DESY PRC-98/02

- [24] J.Holt, F.Teubert, I.Tomalin, "The LHCb Level-2 Trigger", LHCb 98-47
- [25] A.Schoning, "Particle identification in the RICH detectors and study of impact parameters", LHCb 97-018

Chapter 13 Data Acquisition

13.1 Introduction

The role of the DAQ system is to read and buffer data from the front-end electronics following a Level-1 trigger, to assemble complete events, and to provide storage facilities for event data and for calibration and monitoring information. The architecture of the LHCb Trigger and Data acquisition system is shown schematically in Figure 3.3, and the main parameters are summarised in Table 3.1.

The main functional components of the DAQ system are as follows:

- The Front-End Multiplexers (FEM) are responsible for multiplexing the data from many detector channels onto the front-end links. The multiplexing factor will be determined by the capacity of the links, and by the physical characteristics of the front-end electronics, i.e. the convenience with which channels can be grouped together. For the design of the readout protocols, we assume that the output of the FEMs carry zero-suppressed data in a detector independent format.
- The number of Front-End Links (FEL) has been chosen to match the expected data rate with the capabilities of the link technology, such that at least one link is allocated for each detector segment. This boundary condition is needed to take into account the segmentation and physical layout of the detectors.
- The Readout Units (RU) receive event fragments from several front-end links and assemble them into larger sub-events. Once a sub-event is assembled, the RU transfers it to the next stage for further event building. The multiplexing factor in the RU is chosen to match the bandwidths on its inputs with the output bandwidth towards the readout network.
- The Readout Network has the task of providing the connectivity and the required data transfer bandwidth such that all sub-event data belonging to a given trigger, distributed

amongst all the RU's, arrive at a single destination. This destination will vary from trigger to trigger. It will thus support building of full events from the sub-events assembled previously in the RU's.

- The Sub-Farm Controller (SFC) has to assemble the sub-events arriving via the readout network into complete events. Once this is complete, it sends the event to one of the CPUs it manages. After processing by the high level triggers, the SFC has the role of dispatching accepted events, via the Readout Network, to the storage sub-system. It will also monitor its buffer usage and raise a request to slow down the trigger if the free space becomes low.

The detailed requirements [1] have been used as the basic input for design and implementation studies.

13.2 Data traffic patterns

The sizes, distributions and rates of data coming into the DAQ system on the Front-End links have been analysed using simulation data [1]. The results are summarised in Table 13.1 in terms of the average event sizes for each sub-detector. The channel occupancies shown assume detector dependent processing for clustering and zero suppression. In the case of the tracking detectors an extra 30% has been added to the estimate for the average event size, in order to allow for the fact that data from two beam crossings are read for these sub-detectors. The number of segments shows the natural physical segmentation of a sub-detector. This is the number of stations in the case of tracking devices, and the lateral segmentation (inner, middle, outer) in the case of calorimeters.

The allocation of front-end links to a sub-detector will require optimisation according to the natural segmentation of the sub-detector, and the distribution of data in the segments. Table 13.1 shows a first assignment of front-end links, which have been calculated under the assumption of a 1

Table 13.1: Average event data sizes for sub-detectors and their possible segmentation

Sub-detector	Channels (10 ³)	Occ (%)	Data (kB)	Segs	FE-links
Vertex	220	0.5	4	17	17
Inner Tracker	220	1.0	10	11	44
Outer Tracker	110	4.5	24	10	40
RICH 1	140	2.0	11	2	16
RICH 2	200	0.5	4	2	16
Preshower	6	2.5	1	3	3
ECAL	6	25.0	6	3	12
HCAL	3	10.0	1	2	2
Muon	45	0.5	1	5	5
Trigger			5	1	10
Total	950		67	56	165

Gbit/s link capacity. The allocation shown is generous with respect to bandwidth, since average link occupancies range from 10 to 30%.

Table 13.1 shows an average size of 67 kByte. There is a large variation from event to event as shown in Figure 13.1. In the readout protocol and implementation studies, an average event size of 100 kByte is assumed in order to allow for already foreseen factors, such as noise in detector electronics. Although the event building rate for LHCb is comparable to that of ATLAS and CMS (i.e. 40kHz), the average event size is a factor of ten smaller, which reduces the scale of the event building requirements by the same factor.

An average of 4 GB/sec comes from the front-end links into the DAQ. All event data goes to a single processor, but the event building requirement depends on whether all the data are immediately sent, or whether they are buffered and dispatched in several phases according to the trigger strategy. In the ‘phased’ approach, the event building requirement can be less than 2 GB/s.

The target Level-3 accept rate is 200 Hz. Thus the basic requirement for writing raw data is ~ 20 MB/s.

13.3 Design criteria

The principal criteria for the design are simplicity and ease of maintenance, scalability to meet new requirements, and the ability to follow developments in technology. A scalable architecture for the DAQ system is essential in order to allow for changes in the running conditions of the experiment that will affect trigger rates, event sizes and trigger algorithms and result in increasing demands on data




Figure 13.1: Total Event Size distribution.

bandwidth and processing requirements.

Commercial components will be used for front-end links, the readout network and the processing farm, and interfaces suitably designed to prevent dependence on a particular technology. It is important that each component can be tested and validated individually, since this greatly simplifies the integration and commissioning of the complete DAQ system.

Data transmission errors must be detected at all stages where communication links are used. Automatic detection and recovery procedures are needed to ensure the integrity of the data and the event assembly mechanism.

It is important that tests can be run in parallel on different sub-detectors, particularly during calibration and commissioning periods. The concept of partitioning will be used to permit this, and this has important consequences for the design of the DAQ system. A partition is defined to be a subset of the DAQ system that has been configured to function independently of the rest of the system. More than one partition may exist at any time thus permitting parallel data streams. Each stream will have its own set of Readout Units (RU’s), Sub-Farm Controllers (SFC’s) and its own independent trigger source. The facilities for distribution of trigger and timing information, event building and event storage, must support the partitioning concept.

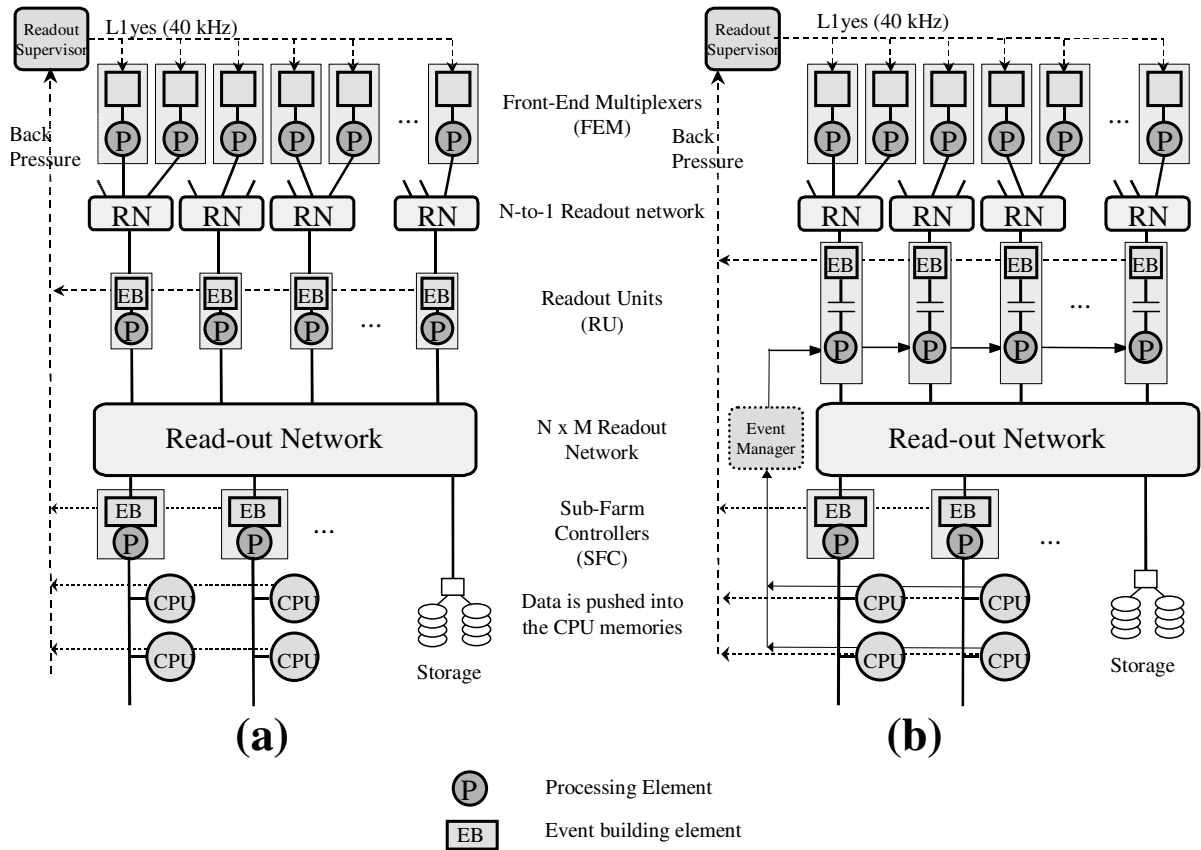


Figure 13.2: Functional elements of (a) the full-readout protocol and (b) the phased-readout protocol. Note the presence of an event manager and the explicit buffering in the readout units to bridge the latency of the Level-2 decision making process in the phased-readout option.

13.4 Readout protocol studies

Within the overall framework of the DAQ architecture, two different protocols have been studied for managing the collection of event data[2]. The first, called the ‘full-readout protocol’, requires that data are immediately transmitted by all RU’s through the readout network, such that a complete event is directly assembled in a destination processor. In the second approach, called the ‘phased-readout protocol’, data from subsets of the RU’s are sent in several phases corresponding to the demands of the high level trigger algorithms. The full-readout approach uses a simpler protocol but places increased bandwidth requirements on the readout network. The functional components of the two protocols are illustrated in Figure 13.2a and Figure 13.2b.

In both approaches the Front-End Multiplexers (FEM’s) push data to their appropriate RU after a Level-1 accept decision. There is no further protocol between the input stage of the RU and the FEM. After the RU’s the two protocols have im-

portant differences, and these are elaborated in the following sections.

13.4.1 Full-readout Protocol

Data from the FEM’s are assembled in the RU and are immediately dispatched through the readout network to an SFC. The algorithm for determining the SFC address can be extremely simple, for example using the event number modulo the number of SFC’s. The buffering capacities of the SFC’s will be designed such that in normal operation there is always sufficient space available to receive data from the RU’s. Obviously the system must also be designed to cope with bursts of data. For this a throttle signal is sent to the Readout Supervisor anticipating the buffer overflow situation, and is used to reduce the trigger rate such that the buffer occupancies can be stabilised. The latency of the throttle signal translates into buffering requirements at the different levels of the readout. First simulation results show that tens of millisec-

onds are available for the throttle to be acted upon before the buffers overflow. Hence this feed-back does not have to be a fast signal, and can even be implemented using the control network.

The SFC collects all event fragments, and once the event is complete passes it to a processor running the high level trigger algorithms. Since all data belonging to a trigger is immediately available to the high level trigger algorithms, there is no real need to distinguish between Level-2 and Level-3. Thus the high level trigger algorithms can evolve with complete freedom.

The bandwidth through the readout network required to implement this protocol is of the order of 4 GByte/s, assuming a 40 kHz Level-1 rate and 100 kByte event size.

13.4.2 Phased-readout Protocol

Using this protocol, the transfer of data from the RU's to the high level trigger processors takes place in two or more phases¹. In the first phase, the subset of the event data that is needed by the Level-2 algorithm is transferred from the appropriate RU's to a processor. The data from the remaining RU's must be buffered for the duration of the Level-2 latency, i.e. ~ 10 ms. The Level-2 decision must be transferred to all of these RU's. On reception of a Level-2 "No" decision the data are discarded. On reception of a Level-2 "Yes" decision the data must be sent to the processor that ran the Level-2 algorithm to execute further filtering algorithms (Level-3) on the complete event.

As in the case of the full-readout protocol the RU's will push the data to the SFC's in each phase, and again a throttle mechanism is required to prevent buffer overflows.

The reduction factor of the bandwidth required of the readout network depends on the rejection power of the Level-2 algorithm and the fraction of the complete event needed to execute the algorithm. At a retention of 1 event in 8, and assuming only 30% of the event data is needed by the Level-2 trigger, the required bandwidth is of order of 35% of the value for the full-readout option, i.e. ~ 1.4 GByte/s.

The phased-readout protocol requires a new functional element, called an 'Event Manager'. Its function is to collect the Level-2 decisions and distribute them to the appropriate RU's. The presence of an event manager would permit the dy-

namic load balancing across the processors of the high-level trigger farm.

13.4.3 Comparison of the two protocols

The full-readout protocol requires a larger scale readout network, both in terms of bandwidth and in number of RU's and SFC modules. However the complexity of these modules is less than in the phased-readout case, since the protocol is much simpler. It also has to be borne in mind that the smaller the number of RU's, the more ambitious are the requirements on the multiplexing of the data from the FEM's into the RU's.

From the point of view of simplicity of the protocols, and the flexibility it allows for the high-level trigger algorithms, the full-readout protocol is to be preferred. It is the availability of affordable high bandwidth network technologies that will determine which protocol will be implemented. The decision can be delayed for several years, during which time the high level trigger strategy will evolve, and the technology used for the readout network can be evaluated in terms of cost and market share.

In the short term, effort will be concentrated on making simulations ("virtual prototypes") of the entire readout system. Many design issues will be studied, such as the impact of transmission errors on the performance of the system and the impact of latency of the throttle on the buffering capabilities at the different levels of the readout.

13.5 Timing and trigger distribution

LHCb is the only experiment at LHC that has to distribute more than one level of trigger decision to the front-end electronics. The partitioning scheme also places specific requirements on the distribution of different trigger signals to different parts of the DAQ that may be running concurrently, asynchronously and under control of different operators.

In the R&D phase for the LHC experiments considerable work was done in the RD-12 project [3] to develop a Trigger, Timing and Control distribution system (TTC) for the LHC experiments. The TTC system comprises a transmitter that distributes timing and trigger signals over a passive optical fibre network, and a special receiver chip is under development for decoding this information and distributing it to the front-end electronics. One basic feature of this system is that it offers two

¹More than two phases could be implemented. However with the current performance of the Level-2 algorithms this would result in little gain. A large number of phases would eventually result in a "data on demand" scheme.

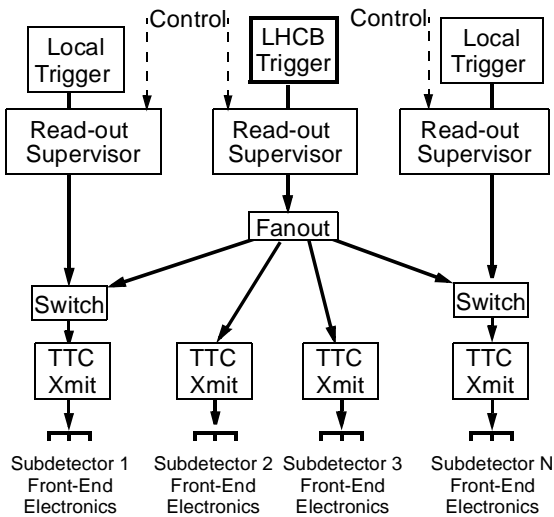


Figure 13.3: Architecture of the LHCb Timing and Trigger Distribution system. The figure shows two subdetectors with their own local trigger and two which only run with the LHCb trigger.

transmission channels (A and B) of 40 MHz bandwidth each. Channel A is reserved to transmit a clock synchronous to the LHC clock and one level of trigger decision, whereas channel B allows commands to be sent to individual receiver chips or groups of chips (broadcasts).

Careful study has shown that this TTC system can be used in the LHCb environment to distribute both Level-0 and Level-1 trigger decisions to the front-end electronics[4]. Channel A will be used to transmit the LHC clock and the Level-0 trigger decision, and channel B to transmit the Level-1 trigger decision. The scheme depicted in Figure 13.3 will be used to support partitioning. This would enable a sub-detector to run either within LHCb, or in stand-alone mode. Both the TTC transmitter and receiver chip could be used without modification. All other components, i.e. programmable switch, fan-out and readout supervisor, will be built specifically for LHCb. At least one transmitter will be deployed for each sub-detector, and a total of a few thousand receivers will be required.

13.6 Error handling

There is a non-negligible probability that data transmission errors will occur, both in the collection of the data and in the distribution of control and sta-

tus information[2]. The error rate depends strongly on characteristics of the technology used to transfer the data, e.g. the frame sizes, transmission media etc. One essential requirement on the data transmission technology is the ability to detect that an error has occurred.

Transmission errors in the acquisition of data can have two types of consequence:

- Errors leading to a violation of the readout protocol through corruption of protocol information, e.g. the event number. These will result in incomplete events being assembled. The protocol has to be designed to handle these cases, usually manifesting themselves as timeouts.
- Errors leading to corruption of the contents of the event data. This will result in missing or wrong information in small pieces of data.

Since the two readout protocols are very similar, the effect of data transmission errors on their performance is also very similar. Neither of the high-level protocols described above implements a re-send mechanism in case the receiver detects an error². If the expected error rate of the chosen technology is unacceptable, error correction codes can be implemented to overcome the problem. Error correction codes have a minimal impact on the total data volume ($\sim 2\%$), and allow correction of single bit errors and detection of double bit errors [2].

The most serious transmission errors are those in the TTC system, since they cannot be detected immediately but only at a later stage in the readout of the data. For example, a missing clock signal or a wrong Level-0 decision can cause a desynchronisation of the pipelines in the front-end electronics. An appropriate scheme will be devised to recover from these errors once their expected rate is measured. These measurements are currently being performed in the context of the RD-12 project.

13.7 Implementation studies

The DAQ implementation studies aim at investigating the feasibility of the proposed DAQ architecture and protocols, identifying the critical components and suggesting strategies for their implementation. Different technologies have been con-

²There are technologies, e.g. SCI, that implement an automatic re-transmission of data at the physical level when an error is detected.

sidered : ATM and SCI which are well advanced standards, the Ethernet family for which Gigabit Ethernet is an emerging standard, and Myrinet a single company solution. A detailed account of the implementation studies and of an ATM implementation is given in [5]. A scheme based on SCI is proposed in [6].

13.7.1 Requirements

The following assumptions on data traffic patterns have been used in simulation studies of the implementation. The data distribution by detector and front-end link is obtained by normalizing the values from Table 13.1 to a total event size of 100 KByte. The Level-1 trigger frequency is taken to be 40 kHz. The inter trigger delay is variable, but is not of great importance, since the DAQ system provides several levels of de-randomizing buffers. The latency of Level-2 is assumed to be 10 msec, and for Level-3 200 msec. For phased event building, the hypothesis is $\sim 40\%$ of data required for the Level-2 algorithm and a rejection factor of 8.

Safety Factor: The figures above characterize the working point of the DAQ system. To provide a safety factor, it is required that the load at the working point should represent 50% of the maximum load that the DAQ system is able to sustain. A load of 100% would be realised by increasing the trigger frequency by k_F and the event size by k_E with $k_E \times k_F = 2$.

13.7.2 Implementation model

The event building system (Figure 13.4) consists of sources (i.e. RU's) and destinations (i.e. SFC's) connected by a single network. RU's are all identical but allow for a variable number of Front-end Link receivers. The destination modules are all identical as well. The functionality of each component is summarised below.

Readout Units (RU)

A Readout Unit collects data from one or more front-end links, trying to achieve a load of its network link of $\sim 50\%$. The determination of the multiplexing factor depends on the front-end link data rate, on the bandwidth connecting the RU to the network, and on the requirement of partitioning that forbids links from different sub-detectors to be mixed. These constraints result in an average load somewhat lower than 50%.

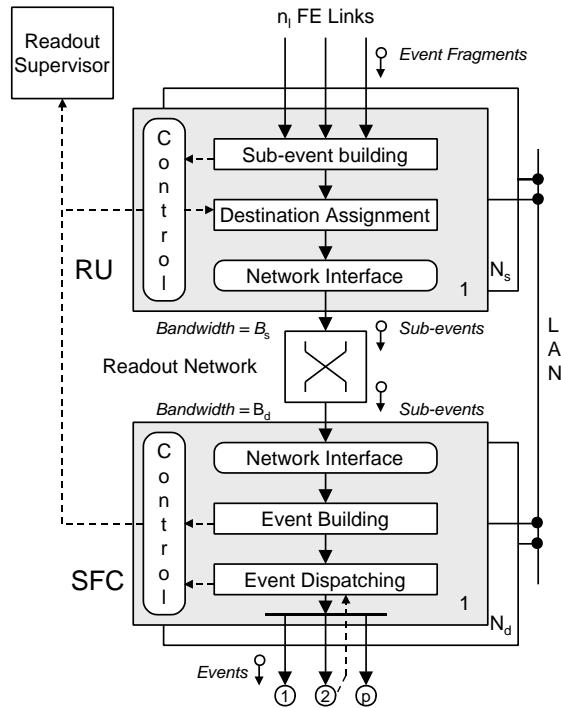


Figure 13.4: Implementation and functional model of the event building system.

The RU's have to cope with a total frequency of event fragments several times higher than the Level-1 rate. For example, a network providing a bandwidth of 1 Gb^{-1} can require a multiplexing factor of 4 in some sources, thus leading to a rate of event fragments at 160 kHz. Higher bandwidths lead to higher multiplexing factors. The implementation of RU's must rely on hardware solutions, or at least on dedicated processors, to sustain such rates. A detailed table giving the assignment of Front-end Links by detector can be found in [5].

The assignment of a sub-event to a destination can be determined locally on the basis of a look-up table or by a dedicated controller that collects information on the status of the SFC's and event processing units. Although the first solution is distributed, it requires a "slow" backpressure control from the destination to the readout supervisor in order to control the Level-1 trigger rate thus avoiding buffer overflow in a destination. Another solution based on a central controller allows the load to be balanced while keeping an optimal overall throughput. It can easily be implemented if the control messages are carried by the same data network as the main DAQ stream. Possible implementations of the event building control are discussed in detail in [5].

Table 13.2: Number of SFC modules ($\sim 50\%$ load on link)

Network Technology	Full Readout		Phased Readout	
	# of SFC's	CPU/SFC	# of SFC's	CPU/SFC
Ethernet 100 Mbs ⁻¹	800	2 - 3	400	5
ATM 155 Mbs ⁻¹	400	5	250	8
ATM 622 Mbs ⁻¹	120	17	60	34
Ethernet 1000 Mbs ⁻¹	80	25	35	57
Myrinet 1280 Mbs ⁻¹	56	36	27	73
SCI 4096 Mbs ⁻¹	132	15 - 16	67	30

If phased event building is implemented, the source modules must implement a buffering scheme where Level-3 data are stored until a Level-2 decision is made. The additional buffer space required is modest: 1-2 MByte per RU (assuming event building latency of a few msec, processing time of 10 ms and the highest rate of 80 kHz); it represents a shift of buffer capacity from the sub-farms to the RU's.

Sub-Farm Controllers (SFC)

The SFC has the task of assembling sub-events to full events and implements the event building protocol that copes with sub-event losses. It must also implement the partial event building scheme required by the partitioning requirements.

In a simple model, the SFC dispatches full events to one of the processors in the sub-farm that it controls, the individual processors not being "visible" from the RU's. This gives the possibility of implementing a local load balancing that issues an overflow warning only when all the local resources become scarce.

The number of SFC's is not constrained by bandwidth considerations. However the number of processors required to process the flux of incoming events in an SFC is proportional to the readout network link bandwidth and is determined from the average processing time per event and the number of events per second arriving in the destination. Table 13.2 gives the number of destinations for various technologies, assuming a bandwidth occupancy of $\sim 50\%$. The table has an entry for both protocols.

Readout Network

The main function of the network is to route all the parts of a given event to one destination. As such it participates in the event building process. The concentration of data towards one destination results in contention for shared resources. Depending on the mechanism used by a particular technology to resolve contention, this may lead to reducing the usage of the available bandwidth, thus limiting the maximum load that the network can tolerate. This is an important issue that can be studied by simulation.

The network can also be used for routing control messages. The problem of contention is not important since most messages travel in the opposite direction to the data transfer.

The main advantage of phased event building is to reduce the network size by a factor 2-3, depending on the data requirements for Level-2 and on the rejection factor. The cost is a higher multiplexing of Front-end Links for Level-3 data, and a more complex control of the Level-3 RU's, the RU having to implement an event retrieval or discard logic.

13.7.3 An ATM implementation

The RU's are connected to the readout network by 622 Mbs⁻¹ links. The Level-2 RU's need to multiplex at most 2 Front-end Links. In the phased event building case, a multiplexing factor of up to 4 is required for the RU's providing data for the Level-3 trigger. For the destinations, we have selected links at 155 Mbs⁻¹ to have a small number of processors in the sub-farms, possibly using 'commodity' symmetric multiprocessors (note however that phased event building requires 8×1000 MIPS processors per sub-farm). The average load on a destination link is $\sim 50\%$, at the normal working point conditions. The network consists of interconnected switches offering 10 Gbs⁻¹ aggregate bandwidth each, or the equivalent of 64 ports at 155 Mbs⁻¹. Table 13.3 gives the network characteristics for both event building protocols.

Simulations of roughly equivalent networks, with ports at 155 Mbs⁻¹ only, have been run: a 1024 ports switching network for the full event building and 512 ports for the phased event building, equally distributed in both cases between sources and destinations. Results show that no congestion occurs for events of 100 KByte at frequencies up to at least 70 kHz (Figure 13.5). The buffer occupancy inside the network due to contention is much lower than the capacities usually provided by commercial

Table 13.3: Number of RU's and SFC's for the ATM scenario

	Full-readout	Phased-readout
# of L2 RU's (622 Mbs ⁻¹)	125	50
# of L3 RU's (155 Mbs ⁻¹)	–	60
# of SFC's (155 Mbs ⁻¹)	494	244




Figure 13.5: Tail distribution of Queue occupancy in switching nodes.

switches.

13.7.4 Outlook

This example of an implementation enhances our confidence that the proposed DAQ architecture is feasible. Alternative or emerging networking technologies might provide equally valuable solutions and will deserve further studies. There are 2 important subjects for further investigation. The first is a detailed study of the feasibility of RU's, mainly from the point of view of sub-event building. Reference [5] suggests a layout that could easily be studied without hardware developments. The second field of investigation is network simulation adapted in particular to technologies other than ATM.

13.8 Detector Control System

The Detector Control System (DCS) will be used to monitor and control the operational state of the LHCb detector and the associated experimental equipment such as gas systems, high voltages,

and readout electronics. The system should be able to operate the experiment as a whole from the control room during physics data-taking periods, but it should also allow the operation of the different sub-systems in a stand-alone manner at other times. The DCS will also acquire slowly changing data from the detector (for instance environmental parameters such as temperatures, positions, etc.), and store them permanently to be accessed by the reconstruction and analysis programs.

The number of control and monitoring channels will be of the order of many tens of thousands. To handle this complexity, these channels will need to be grouped into logical devices and devices will be grouped into sub-systems. Each sub-system should offer only summary data and high level commands to other sub-systems, and to the operator.

Robustness and reliability are important requirements for the DCS. The DCS must be fully operational at all times, regardless of the state of the experiment, of the DAQ system or of the LHC accelerator. The DCS should be well integrated with the DAQ system, since both need to exchange information and commands, and to ensure that a coherent interface is provided for control and status display in the control room.

13.8.1 The DCS architecture

The architecture of the DCS implementation is shown schematically in Figure 13.6. It is a distributed control system consisting of control stations (I/O servers) and general-purpose workstations or servers, interconnected by a local area network. The equipment is connected to the control stations by means of a field bus or by separated analog or digital lines. The control stations will constantly read all I/O channels and report to clients running on the workstations and servers only when changes are detected. Specialized process control stations will be implemented using commercial programmable logical controllers (PLC). They will also be connected to control stations.

A key feature of the software architecture is the central database. This database will contain all the configuration parameters for the complete system including the configuration of the data points, relationships between equipment, data presentation parameters, graphics, etc.

13.8.2 Organization and planning

The DCS will be part of the computing project of LHCb. As for the DAQ, the central DCS will provide the frameworks and the infrastructure to the

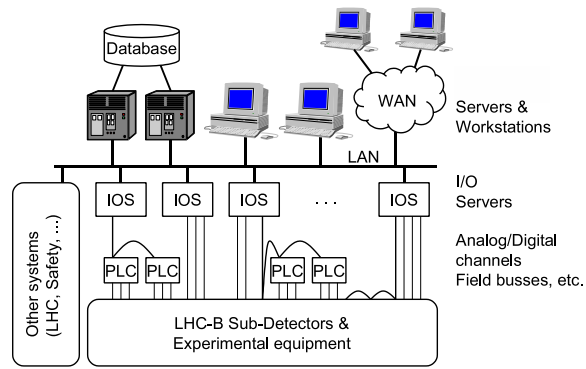


Figure 13.6: Hardware architecture for the Detector Control System.

sub-detectors groups which will need to configure and develop the specific control functions for their sub-detector.

A common project for the DCS for all LHC experiments has recently started. The objective of this project is to explore ways of providing a common DCS for all 4 experiments in a resource effective manner and recommend and support which control system to use. LHCb is participating in the project, and is currently engaged in writing the user requirement document for the DCS kernel [7].

References

- [1] F.Harris, M.Frank, "LHCb Data Flow Requirements", LHCb 98-027 DAQ.
- [2] B.Jost, "DAQ Architecture and Readout Protocols", LHCb 98-028 DAQ.
- [3] O. Villalobos Baille et al., "Timing Trigger and Control Systems for LHC Detectors - Status Report on the RD12 Project", CERN/LHCC 97-29, 1997.
- [4] B.Jost, "Trigger and Timing Distribution in LHCb", LHCb 98-031 DAQ.
- [5] J.P.Dufey, I. Mandjavitzze, "DAQ Implementation Studies", LHCb 98-029 DAQ.
- [6] H.Müller, "SCI Implementation Study for LHCb Data Acquisition", LHCb 98-30 DAQ.
- [7] P.Mato, "Detector Control System for an LHC experiment", LHCb 98-005 DAQ.

Chapter 14 Computing

14.1 Data-flow model

The data produced by the detector and by Monte Carlo simulation undergo a number of transformations before reaching the desktops of the physicists performing the analysis. Each transformation places specific requirements and constraints on the system. For example there are requirements on processing power, data storage and network bandwidth, and constraints from the availability of calibration and alignment data, or on the quality and reliability of the algorithms.

The data flow model of the LHCb experiment is shown in Figure 14.1. The model is similar to that of other HEP experiments. There are however certain features which are particularly relevant to LHCb:

- The Level-2 and Level-3 software triggers must achieve rate reduction factors of 8 and 25 respectively. This will require sophisticated algorithms similar to those employed in the full reconstruction, but which must be fast enough to cope with the expected input rate and highly reliable since only the data passing these triggers will be written to permanent storage. Calibration and alignment constants will need to be of sufficient quality to allow the triggers to work efficiently, implying that at least some part of the calibration and alignment must be performed in real time. The full reconstruction will make use of the results of the trigger algorithms.
- The raw data is written to permanent storage at a rate of 20 MB/s. The reconstruction program adds a similar amount of data to the data store, at a similar rate. Assuming a data-taking year of 10^7 seconds, the permanent data store will have to be capable of accepting in excess of 500 TB of data per year at a rate of 40 MB/s. This is 2-3 orders of magnitude greater than current experiments, and similar to the other LHC experiments [1].
- Nearly all processes in the diagram access at least some fraction of the data store. This is

a database of enormous size and complexity, which will need to be designed, maintained and populated in such a way that all processes will have transparent access to the necessary information.

14.2 Computing requirements

The tasks identified in the data flow model are analysed by breaking each of them down into subsystems, such as individual subdetectors for the simulation, subdetectors plus tracking, calorimetry and particle identification for the reconstruction, and individual physics channels for the analysis. For each subsystem, studies are made of existing simulation, reconstruction and analysis algorithms to measure processing power requirements, to understand the types and volume of data required as input and produced as output, and to evaluate the dependence on calibration and alignment data.

The measurements are extrapolated to the final system by making reasonable assumptions on the evolution of the algorithms and on the expected luminosity and trigger efficiency. Where suitable algorithms do not yet exist, such as for pattern recognition in the tracking system or for all but the simplest analyses, estimates are based on extrapolations from existing experiments. This approach should ensure that the estimates of data storage, processing and networking needs are realistic. It should give a detailed understanding of access patterns to the data and of the data flow logistics, allowing reasonable constraints to be placed on the computing model.

The current processing power and data storage requirements of LHCb are summarized in Table 14.1 and Table 14.2 respectively, and compared with estimates of the annual requirements during data taking. The detailed breakdown of the estimates by computing task can be found in [2]. The figures for reconstruction assume two passes through the real data in order to extract calibration constants etc. One complete re-processing of all the real data per year is also assumed. It should be noted that the total CPU requirements do not

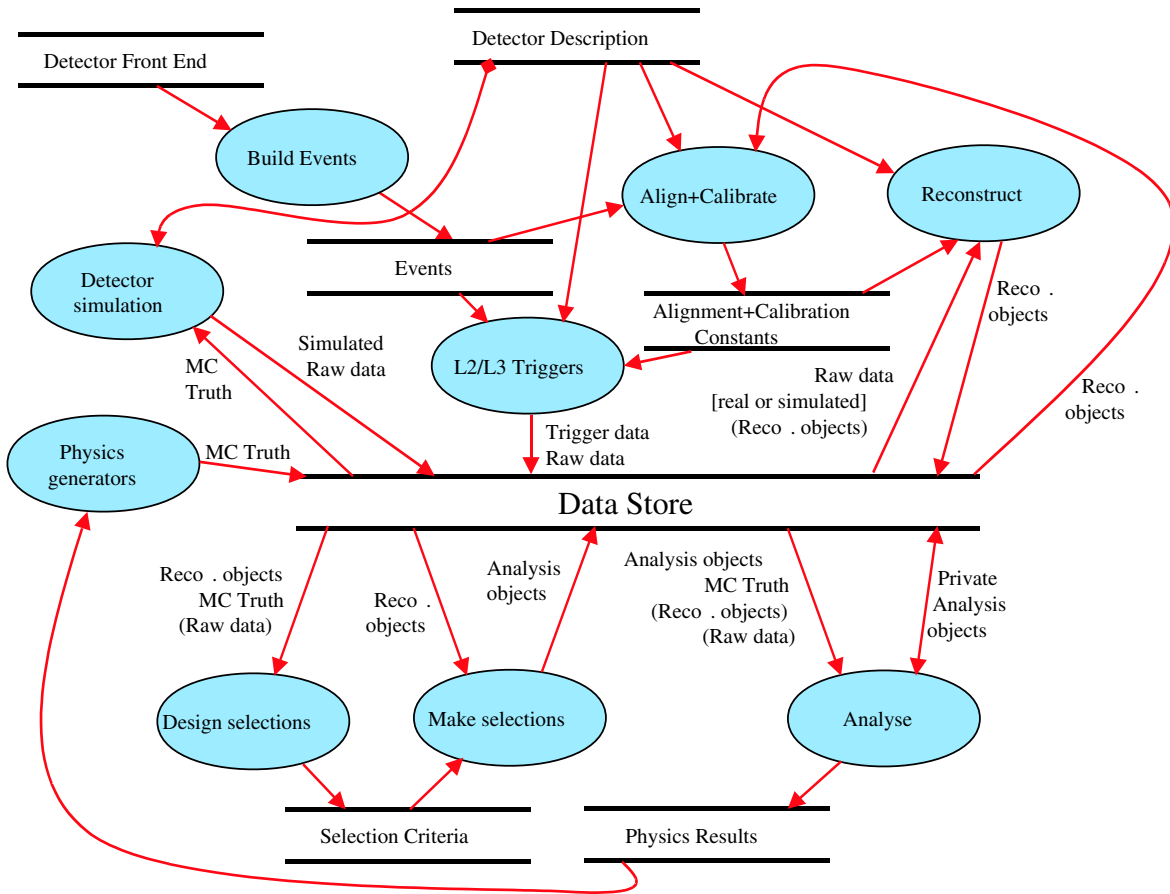


Figure 14.1: LHCb data flow model.

necessarily map to processors installed. Only a detailed computing model will tell how exactly the mapping of requirements to processors will be done. For example, some of the CPU power installed for triggering could be used for re-processing outside data-taking periods.

14.3 Computing infrastructure

Until very recently LHCb relied almost entirely on public Unix services at CERN for its computing needs. In the last few months access has been gained to regional facilities outside CERN and to the public PC simulation farm (PCSF) at CERN for large Monte Carlo productions. Although processing requirements are expected to grow by $\sim 50\%$ per year until the end of 2000, it should be possible to satisfy these needs by use of these public facilities during this period. The simulation program will produce sequential data until the end of 1999, and therefore investment in object oriented

Table 14.1: Processing power requirements

Processing type	CPU used (MIPS, 1997)	CPU needed (MIPS, 2006)
Triggering		
Level-2	-	4×10^5
Level-3	-	1×10^6
Reconstruction		
Real Data	-	1.6×10^6
Sim. Data	200	2×10^5
Simulation		
B-meson	600	1.2×10^6
Background	-	1.2×10^3
Analysis		
Real Data	-	5×10^4
Sim. Data	150	5×10^5
Total	950	$\sim 5 \times 10^6$

Table 14.2: Data storage requirements

Storage type	Storage used (TB, 1997)	Storage needed (TB/year, 2006)
Raw	0.01	200
Monte Carlo Raw	1.0	100
Calibration & Alignment	-	< 1
Reconstructed (Raw+MC)	1.0	200
Analysis	0.05	~2
Total	~2.0	~500

databases is not expected before 2000.

As test beam activity increases and Technical Design Reports are prepared, a more significant increase in CPU needs is anticipated. This will be the appropriate time to begin investing in private facilities. Investment in full scale facilities will be left as late as possible, to benefit fully from continuous improvements in the price/performance ratio of computing equipment.

The main issues relating to infrastructure that must be addressed during the preparation phase are:

- The ability to build and operate large scale computer farms containing several hundred CPUs.
- The handling of petabytes of data stored in a central database.
- The access to the data from HEP institutes outside CERN.

These issues are of common interest to all LHC experiments, and are discussed in detail elsewhere [3]. Current thinking is that experimental data in the LHC era will be logically stored in a single, experiment-wide Object database, with the physical storage managed by a mass storage system, as currently being investigated by the RD45 [4] project. The distribution of storage and computing resources at CERN and amongst regional centres will be optimized, in the light of networking, political and sociological constraints.

The mass consumer market will have a big impact on the provision of the necessary computing resources at an affordable cost. PC technology is evolving rapidly, and industry standards must be followed closely to avoid making inappropriate

choices that might result in a premature commitment to an obsolescent platform, language or operating system. None of the requirements can be considered outside the limits of current technology trends.

14.4 Software strategy

The LHCb approach to managing the production of software is outlined below. More detailed information can be found in [5].

14.4.1 Objectives

The strategy adopted for ensuring good software quality will have important consequences for the success of the experiment. Intensive use is made of software in the trigger, and the event reconstruction and selection algorithms must be implemented correctly and be fully tested by the time first colliding beams are delivered. Bulk processing of TByte data samples require large CPU facilities, and re-processings will be costly and must be minimised.

Software quality can best be ensured by using modern engineering practices, modern technologies and by putting in place sound management practices for controlling the development process. Effort will be put into building high quality software components, and the extra investment in manpower that this represents will be recuperated by extracting as much use as possible out of these components in all application domains. Commercial products will be used whenever appropriate technically, and providing they can be licensed to all collaboration institutes.

14.4.2 Technology

The software industry has invested heavily in the development of Object Technologies. A good level of maturity has been reached in the definition of standards and commercial products, such as databases and development tools, are becoming widely available. Supported features, such as patterns, components and frameworks, will provide opportunities for sharing concepts, designs and code across projects and will therefore help to achieve the re-useability goals already mentioned.

The most widely used OO language today is C++. However C++ is known to have deficiencies and other languages could be more appropriate providing they become widely used and attract industrial support. Java is an example today of a new

language with this potential. The two most popular design methodologies, OMT and Booch, are converging to define a single method, called UML, incorporating the best features of both. UML is supported, at least partially for the moment, by the major CASE tools, such as Rose and StP.

At this moment a combination of C++ and UML offers a robust and well supported software development environment. However it is important that the architecture and design of the software is independent of the language and design tool, since changes in the choice of methods and tools may become necessary.

14.4.3 Migration policy

The process of understanding an object-oriented language is typically quite time-consuming and it can take several months to become an effective OO developer. The transition to OO programming is one of the most important challenges we face in the immediate future and its introduction requires careful planning.

At present the LHCb development environment is based on FORTRAN. The simulation program is the basic tool for making detector optimisation and physics performance studies, and this uses GEANT3 as a framework. Tools for reconstruction and analysis written in FORTRAN also form part of the maintained suite of software. It is our intention to re-engineer these codes using Object Technologies. There is an incentive to make this transition as soon as possible in order to minimise the amount of “legacy code” that has to be maintained.

In the first instance, the language, tools and programming environment will be selected and a suitable infrastructure (i.e. user interface, data management tools etc.) assembled such that application development can begin. Much of the impetus for this work will come from the release of the LHC++ library[6], which implements some of the key HEP specific toolkits for data analysis, a new detector simulation framework (GEANT4) and object-persistency services based on a commercial ODBMS (Objectivity). Using these components, new simulation, reconstruction and analysis frameworks will be developed. Once the development tools and frameworks are in place, a rigorous training programme will be organised for collaboration members. At this time, further developments of the existing FORTRAN code will be frozen, and it will be discarded altogether as soon as the new OO based programs are put into production. A detailed plan exists showing details of the tasks and

their timescales[7].

14.5 Project organisation

All computing related activities will be managed within one project framework (Figure 14.2). The scope includes on-line and off-line systems and covers both hardware and software aspects. The aim of this approach is to minimise barriers between different activities and encourage the sharing of concepts, tools and software. Teams dedicated to support software engineering and computing facilities will provide the environments used to develop the software and the computing infrastructure. The development of software of common interest to all applications will also be done by a dedicated team. The Computing Project will be managed by a Steering Group, comprising the Computing Coordinator and the leaders of the various subprojects.

14.6 Project management

The lifetime of the LHCb experiment, from its conception to its completion, spans more than 15 years. The main phases and activities during this time are summarised below. More detailed information on current planning, including task schedules and manpower estimates can be found in [7].

Preparation phase (now – 2000) Functional specifications of all components will be developed. Knowledge and experience of relevant technologies must be accumulated, and in some cases detailed evaluations made through prototyping and simulation. This will be an intense period of learning and will involve formal training. Public computing facilities will be used for the processing of simulation and testbeam data during this period.

Implementation phase (2001 – 2003) Final choices will be made on software and hardware technologies, and all sub-systems developed. Reviews will be made of the architecture and design of software in the light of experience obtained during the preparation phase. Hardware components will be built and tested, and installed in the experimental area. During this phase, data processing needs will increase sharply, due to increased test beam activity, and investment in private CPU capacity will be needed.

Commissioning Phase (2004) A significant period of time must be allocated for the assembly

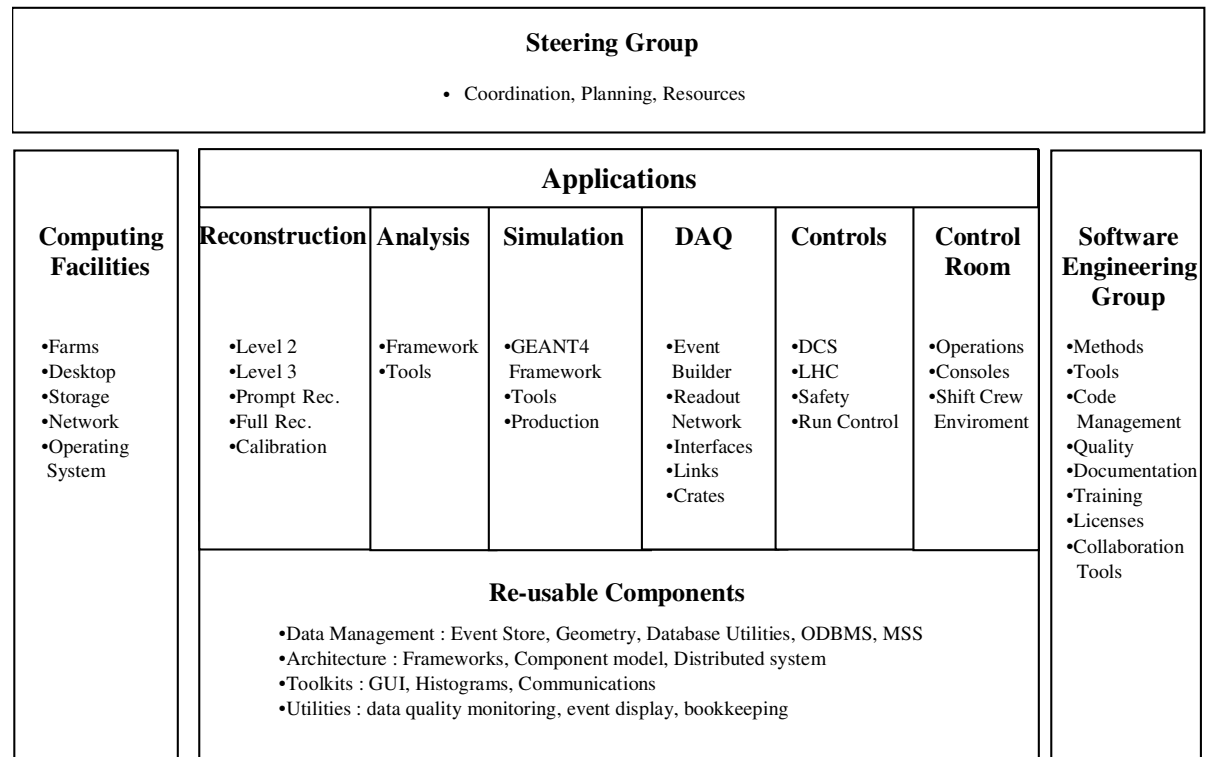


Figure 14.2: Breakdown of Computing activities into subprojects.

of all the various computing subsystems, and for thoroughly testing them under realistic data taking conditions, since the detection and elimination of faults can take a considerable time. The software must also be tested in advance to ensure that the event selection mechanisms and reconstruction algorithms are working correctly, and that the expected data volumes and processing loads can be sustained. Operational experience must be obtained with large scale CPU farms. The commissioning phase must therefore start towards the beginning of 2004.

Operation Phase (from 2005) This starts when beams are delivered and involves maintenance of the computing environment to match evolving requirements and to keep up to date with technology changes.

References

- [1] ALICE Technical Proposal, CERN/LHCC/95-71; ATLAS Computing Technical Proposal, CERN/LHCC/96-43; CMS Computing Technical Proposal, CERN/LHCC/96-45
- [2] P. Binko, "LHCb Computing Tasks", LHCb/98-042 COMP
- [3] M. Frank, A. Pacheco, "LHCb Computing Model", LHCb/97-046 COMP
- [4] RD45 Collaboration, "Object Databases and their Impact on Storage-Related Aspects of HEP Computing", CERN/LHCC 97-7
- [5] E. van Herwijnen, "LHCb Software Strategy", LHCb/98-025 COMP
- [6] LHC++ , Version 1.0 Users Guide, 3 Nov 1997
- [7] J. Harvey, LHCb Project Plan for Computing Technical Note, LHCb/98-026 COMP

Part IV

Physics Performance

Chapter 15 Reconstruction

15.1 Simulation program

For physics performance studies, events are produced using the PYTHIA 5.7/JETSET 7.4 generator [1]. Default values are used for the physics parameters of PYTHIA. The LHCb detector response is simulated using a program (SICb) based on the GEANT 3.21 package [2]. A description of the detector elements is given in Part II of this Technical Proposal.

SICb contains a simplified model of the LHC beam elements, the LHCb cavern and infrastructure, and a detailed field map of the dipole magnet that includes inhomogeneities and stray field. The beam pipe and the vacuum tank, containing the Vertex Detector, are fully reproduced in the program. They have a large effect on the particle fluxes in the spectrometer generated by secondary interactions.

Simulation of the Vertex Detector takes into account details such as energy deposition in the silicon wafers and charge sharing between the strips. The support structure and the R.F. shield are included in the simulation.

The RICH simulation includes generation and tracking of Cherenkov photons up to the surface of the photodetectors. A detailed digitisation package is implemented for the HPD photodetectors. However, only parameterised photodetector efficiencies and resolutions have been used for physics studies.

In the Outer Tracker, the drift time is simulated from the position of the particle trajectory in the drift cell. The Inner Tracker hits are generated with a single-hit r.m.s. of $60 \mu\text{m}$.

Energy deposition in the ECAL is calculated in the GEANT framework with energy thresholds of 1 MeV for electrons and photons, and 10 MeV for hadrons. With these threshold values, the shower development in the media of the calorimeter, including energy leakage at the boundaries, can be accurately simulated.

Energy deposition in the HCAL is determined with an energy threshold of 500 MeV for all particles. A limited amount of data produced with thresholds of 1 MeV for electrons and photons and 10 MeV for hadrons shows that the results obtained

with the higher threshold values are valid.

For both calorimeters the digitised energy is then calculated by smearing the deposited energy assuming Gaussian resolution.

The Muon Detector includes the iron absorbers and the muon stations. Low-energy particles produced in secondary interactions are taken into account when generating the pad hits.

In the event reconstruction, a full pattern recognition is made for tracks in the Vertex Detector, including those outside of the spectrometer acceptance. Pattern recognition in the main tracking system is not yet fully implemented. However, a study made for the Level-2 trigger shows that the efficiency for finding a track element in the main tracking system, using a seed found in the Vertex Detector, is very close to 100%.

Track parameters are obtained by a fit procedure based on the Kalman-filter algorithm, where continuous energy losses and multiple scattering are taken into account.

The primary vertex is determined using all the tracks found in the Vertex Detector, including those outside of the spectrometer acceptance. For tracks with no momentum measurement, momenta are estimated from the polar angles of the tracks. The resolution of the reconstructed primary vertex is $\sigma_x = \sigma_y = 9 \mu\text{m}$ and $\sigma_z = 38 \mu\text{m}$ for pp inelastic events.

For particle identification with the RICH counters, parameterised angular resolutions of the measured Cherenkov angles in the three radiators are used to calculate probabilities for the e, μ , π , K and p hypotheses. Studies of full pattern recognition have been used to verify the efficiency of this method. Information from the ECAL or Muon Detector is added, if available, for the electron candidates and muon candidates, respectively.

If an electromagnetic cluster has no corresponding hits in the pad chamber (M1), it is considered to be a photon.

15.2 Flavour tag

Identification of the initial flavour is required in order to study neutral B-meson decays involving flavour oscillations. This is achieved by measuring the flavour of the accompanying b-hadron from the charge of decay leptons and kaons. If the tagging lepton is a direct product of the decay of the b-hadron (i.e. $b \rightarrow \ell$), or the tagging kaon contains the s-quark from the decay chain $b \rightarrow c \rightarrow s$, then the flavour of the parent b-hadron at decay will be correctly identified.

Leptons are required to be identified using information from the Muon Detector and the calorimeters, combined with information from the RICH. To reduce the contamination by electrons from conversions, tracks are required to have at least one space point measured in the Vertex Detector. The p_T of all lepton-tag candidates is required to be above 1.5 GeV/c. As seen in Fig. 15.1, this cut reduces leptons from $b \rightarrow c \rightarrow \ell$ decays, that produce a wrong tag as they have opposite charge to those produced by $b \rightarrow \ell$. If more than one muon- or one electron-tag candidate exists, the candidate with the highest p_T in each tag class is chosen.

Kaon-tag candidates are selected using RICH particle identification. As seen from Fig. 15.2, kaons from b-hadron decays have a large impact parameter, d_0 . Therefore, the candidate tracks are required to have an impact parameter with respect to the primary vertex with a significance, d_0/σ_d , greater than 3.0, where σ_d is the error on d_0 . In addition, tracks must have an impact parameter less than 30 mm and a p_T greater than 0.4 GeV/c. It is also required that the tracks have at least one space




Figure 15.1: The measured p_T distribution for muons from $b \rightarrow \mu$, $b \rightarrow c \rightarrow \mu$ and other sources, in events where one B meson decayed into $\pi^+\pi^-$.




Figure 15.2: Measured distribution of the impact parameter significance for charged kaons with $p_T > 0.8$ GeV/c from b-hadron decays and from other sources, in events where one B meson decayed into $\pi^+\pi^-$. The kaons are required to have at least one associated Vertex Detector hit.

point measured in the Vertex Detector to obtain a good impact-parameter measurement. If more than one kaon-tag candidate survives these cuts, then the candidate with the highest p_T is chosen.

The statistical error of an asymmetry measurement is inversely proportional to $\sqrt{\varepsilon_{\text{tag}}(1 - 2\omega_{\text{tag}})}$ where ε_{tag} is the efficiency for finding a tag and ω_{tag} is the wrong-tag fraction. Oscillations of neutral B mesons introduce an irreducible source of wrong tags. For B_d mesons, the wrong-tag fraction is enhanced if cuts are applied which reject decays at early proper time. The B_s mesons oscillate rapidly and their decay products do not provide useful tags. Another source of wrong tags is misidentification of particles. The cuts discussed above are chosen to maximise $\sqrt{\varepsilon_{\text{tag}}(1 - 2\omega_{\text{tag}})}$.

Up to three simultaneous tags could be present in one event. If more than one tag exists, the best tag is chosen according to a hierarchy: muon, then electron, then kaon.

Tagging efficiencies and wrong-tag fractions could differ for different triggers and different tag categories. By combining the separately measured CP asymmetries for each trigger and tag category, the smallest statistical error can be obtained. However, for the physics performance studies the overall tagging efficiency and wrong-tag fraction have been used. Simulation studies show that the tagging efficiency and wrong-tag fraction do not vary significantly between different decay modes (within

a statistical error of a few percent). They are

$$\varepsilon_{\text{tag}} = 0.40, \quad \omega_{\text{tag}} = 0.30.$$

15.3 $B_D^0 \rightarrow \pi^+\pi^-$

To isolate well-measured tracks that are potentially from B decays, the following requirements are imposed:

- At least one three-dimensional space point in the Vertex Detector;
- An impact parameter with respect to the primary vertex of $d_0 < 5$ mm, an associated error of $\sigma_d < 500$ μm , and a normalised impact parameter of $d_0/\sigma_d > 1$;
- χ^2 per degree of freedom of the track fit less than 50.

The $B_D^0 \rightarrow \pi^+\pi^-$ reconstruction then proceeds by attempting to form a vertex of all opposite-sign combinations from such tracks, and calculating the invariant mass under the pion hypothesis. Demanding that the combination falls in a initial mass window of 5.0–5.5 GeV/c^2 , that the vertex χ^2 be below 5, and that the flight distance be at least 0.5 mm, results in a signal efficiency of 62% with respect to triggered events within the acceptance, and suppresses the triggered background by a factor of 100.

After this preselection, the following vertex cuts are applied to further reduce the combinatoric background:

- The significance of impact parameter, d_0/σ_d , for both candidate pions to be greater than 3;
- The separation in z between the secondary and primary vertices normalised by the assigned error, $\delta_z/\sigma_{\delta_z}$, to exceed 3;
- The momentum vector of the reconstructed B to be consistent with the flight path from the primary to the secondary vertex, such that $\cos\theta_B > 0.95$, where θ_B is the opening angle between these two vectors;
- No additional track can be added to the candidate vertex which results in a vertex-fit χ^2 per degree of freedom below 500;
- No alternative vertex of 3 tracks or more with a flight distance compatible with a B decay,

and χ^2 per degree of freedom below 300 can be made using either of the two tracks from the candidate $B_D^0 \rightarrow \pi^+\pi^-$ vertex.

The final two cuts ensure that the pions are uniquely associated to the candidate $B_D^0 \rightarrow \pi^+\pi^-$ vertex, and that this vertex is unambiguously made from the two tracks.

In addition, kinematic cuts are imposed on the candidate pions and on the reconstructed B meson:

- The polar angle of the pions, θ_π , in the B rest frame to satisfy $|\cos\theta_\pi| < 0.9$;
- The transverse momentum, p_T^π , of both pions to exceed 1 GeV/c , and at least one pion to have $p_T^\pi > 3.5$ GeV/c ;
- The transverse momentum of the reconstructed B, p_T^B to exceed 3 GeV/c ;

The event selection is summarised in Table 15.1, and the distributions of signal and background in the most important variables are shown in Fig. 15.3.

The net efficiency of these cuts on the unbiased $B_D^0 \rightarrow \pi^+\pi^-$ event sample is 18.3%, but after the events have passed the Level-0 and Level-1 triggers the efficiency becomes 28.6%. This indicates the high correlation between the trigger and the reconstruction cuts. Those $B_D^0 \rightarrow \pi^+\pi^-$ events surviving have a mass resolution of 17 MeV/c^2 . To further suppress the combinatoric background, a tight




Figure 15.3: Distributions of signal (points) and background (broken line) for the most important vertex and kinematic cuts in the $B_D^0 \rightarrow \pi^+\pi^-$ selection.

Table 15.1: Summary of event selection in the $B_d^0 \rightarrow \pi^+\pi^-$ analysis. Efficiencies are given with respect to the preceding step in the selection. “LHCb acceptance” refers to that fraction of events where both pions fall within the acceptance and are momentum-analysed by the spectrometer.

Requirement	$B_d^0 \rightarrow \pi^+\pi^-$ efficiency
LHCb acceptance	0.136
Trigger	0.168
Preselection	0.658
Vertex	0.603
Kinematic	0.724
Tight mass window	0.898
Particle id	0.941
Total efficiency	5.5×10^{-3}

mass window of $\pm 30 \text{ MeV}/c^2$ is imposed, together with the condition that both tracks are found to be pions by the RICH system. The criteria for the pion identification are the “ π best” criteria discussed below. The total reconstruction efficiency for $B_d^0 \rightarrow \pi^+\pi^-$ events passing the trigger is 24.2%. In one year of running at nominal luminosity, 6.9 k reconstructed and tagged events are expected.

The proper-time resolution is found to be 0.04 ps. Figure 15.4 shows the acceptance curve of surviving events as a function of proper time. This has been fitted to a function of the form:

$$A(t) \propto \frac{(at)^3}{1 + (at)^3} + b, \quad (15.1)$$

giving $a = 0.96 \pm 0.06 \text{ ps}^{-1}$ and $b = -0.09 \pm 0.03$.




Figure 15.4: Acceptance of signal events as a function of proper time for the $B_d^0 \rightarrow \pi^+\pi^-$ analysis. The line is a fit to Equation 15.1.

No events from a generic $b\bar{b}$ sample of 500k events pass this selection. By loosening certain cuts and studying how many events are selected, it is seen that the signal to combinatoric background ratio is consistent with 1. The background contribution from specific two-body decay modes is considered separately below.

Rejection of two-body backgrounds

A crucial requirement of the $B_d^0 \rightarrow \pi^+\pi^-$ analysis is the ability to suppress other B decays with the same two-track topology. Most important of these are $B_d^0 \rightarrow K^+\pi^-$, $B_s^0 \rightarrow K^+K^-$ and $B_s^0 \rightarrow K^-\pi^+$. These channels have unknown CP asymmetries and therefore, if present in the selected sample, will confuse the interpretation of the observed CP asymmetry. LHCb is well equipped to suppress these backgrounds by virtue of its good mass resolution and the π/K separation provided by the RICH system.

The $\pm 30 \text{ MeV}/c^2$ mass window alone is sufficient to suppress these backgrounds by a factor of 3. To achieve the necessary rejection, however, the RICH information must be used. Three criteria of “pion identification” have been investigated:

- **π consistent** – that the pion probability returned by the algorithm be 5% or greater for both tracks;
- **π best** – that the best hypothesis returned by the RICH algorithm be a pion or lighter particle, for both tracks;
- **K inconsistent** – that the kaon probability returned by the algorithm be less than 5% for both tracks.

The efficiencies of these criteria relative to the “no particle identification” case on the channels of interest are shown in Table 15.2. Also shown is the signal-to-background ratio, calculated on the assumption that the branching fractions of $B_d^0 \rightarrow \pi^+\pi^-$ and $B_s^0 \rightarrow K^-\pi^+$ are identical, and that the branching fractions of $B_d^0 \rightarrow K^+\pi^-$ and $B_s^0 \rightarrow K^+K^-$ are 2.1 times higher. The “ π best” condition gives the best statistical precision on the CP asymmetry, and is consequently adopted for the yield estimates. In combination with the tight mass cut it gives a signal-to-background ratio of almost 20, thereby reducing the effect of systematic biases in the background to a negligible level. Figure 15.5 shows the reconstructed mass peak before and after the pion identification.

To ensure that the results of the fast RICH parameterisation used for these studies are reliable, a

Table 15.2: Summary of the effect of RICH criteria and mass cuts on the signal-to-background ratio (S/B) in the $B_D^0 \rightarrow \pi^+\pi^-$ analysis, with efficiencies shown in the two-body decay channels with respect to the case where no RICH information is used and a wide mass window is imposed. The wide mass window is $5.0 - 5.5 \text{ GeV}/c^2$ whereas the tight mass window is $\pm 30 \text{ MeV}/c^2$ around the B_D^0 mass. The RICH criteria and assumed relative branching fractions are stated in the text.

RICH criterion	Mass cut	$B_D^0 \rightarrow \pi^+\pi^-$	$B_D^0 \rightarrow K^+\pi^-$	$B_s^0 \rightarrow K^-\pi^+$	$B_s^0 \rightarrow K^+K^-$	S/B
None	Wide	1.00	1.00	1.00	1.00	0.4
None	Tight	0.90	0.28	0.16	0.55	0.9
π best	Wide	0.92	0.08	0.07	0.01	4.9
π best (full simulation)	Wide	0.95	0.12	0.09	0.02	3.4
π consistent	Wide	0.99	0.17	0.12	0.05	2.4
K inconsistent	Wide	0.79	0.03	0.01	< 0.01	12.5
π best	Tight	0.85	0.02	< 0.01	< 0.01	18.0

limited-statistics analysis has been repeated using full simulation of the photodetector response, and complete pattern recognition as described in Chap-

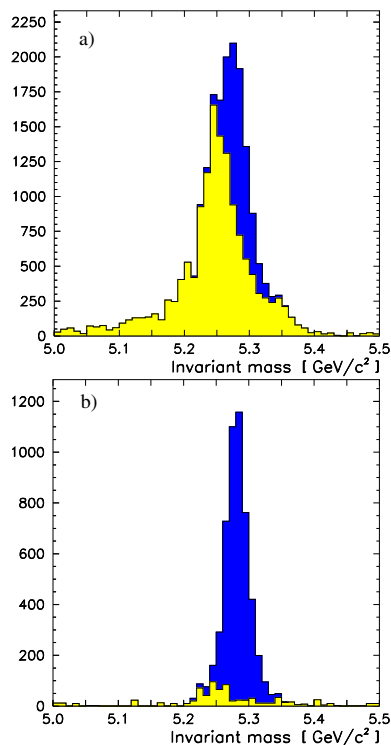


Figure 15.5: The reconstructed mass peak of two-body B decays under the $B_D^0 \rightarrow \pi^+\pi^-$ hypothesis (a) without particle identification, (b) with particle identification. The dark shading indicates the signal and the light shading the two-body background. The particle-identification criteria are given in the text.

ter 9. Despite the fact that the pattern recognition has not yet been tuned in this environment, a good performance is observed.

The background surviving the particle identification is dominated by events with tracks of too high a momentum to be reliably analysed in RICH 2. About 70% of this background has a track of momentum greater than $150 \text{ GeV}/c$, whereas only 20% of the signal is of such a nature. By imposing a momentum cut and tight RICH criteria an extremely pure sample of signal events could be isolated.

Alternatively, the RICH information can be used to identify kaons and thereby select clean samples of the two-body decays other than $B_D^0 \rightarrow \pi^+\pi^-$. For instance, by requiring that one of the tracks be “ π best”, the other be “K best” (that is with a best hypothesis of a kaon or heavier particle), and calculating the invariant mass of events accordingly, enables $B_D^0 \rightarrow K^+\pi^-$ events to be isolated with a purity of 99%. The total efficiency of this selection is 4.9×10^{-3} .

15.4 $B_D^0 \rightarrow \rho\pi$

In the $B_D^0 \rightarrow \rho\pi$ analysis kinematic cuts are applied to the charged tracks and photons in order to suppress combinatoric background prior to the ρ and π^0 reconstruction. These cuts depend on the polar angle θ and are given in Table 15.3. After these cuts there are on average ~ 4 photon candidates per event. Asymmetric π^0 decays are eliminated by demanding both photons to be above a common energy threshold.

The π^0 reconstruction only considers decays with the two ECAL clusters clearly distinguishable. With such a procedure, a π^0 -mass peak is reconstructed

Table 15.3: Initial kinematic cuts applied to charged pion and photon candidates in the $B_d^0 \rightarrow \rho\pi$ analysis.

θ (mrad)	p_π (GeV/c)	E_γ (GeV)
< 70	> 2.0	> 6.0
70 – 140	> 1.5	> 4.0
> 140	> 1.0	> 2.0

as shown in Fig. 15.6 (a). The width varies between 5 and 10 MeV/c^2 depending on the region of the ECAL used and the energy of the π^0 . In the selection a window of 15 MeV/c^2 is imposed.

The ρ peak found in the $\rho^\pm\pi^\mp$ channel has a width of 100 MeV/c^2 . For both the $\rho^\pm\pi^\mp$ and $\rho^0\pi^0$ channels a mass window of $\pm 250 \text{ MeV}/c^2$ is imposed.

Finally vertex cuts are applied to reject the in-




Figure 15.6: Reconstructed mass for (a) π^0 candidates and (b) B_d^0 in the $B_d^0 \rightarrow \rho\pi$ events.

clusive $b\bar{b}$ background. These cuts are similar to those of the $B_d^0 \rightarrow \pi^+\pi^-$ analysis, although at present not all of the possible variables have been studied.

The final B-mass peak is shown in Fig. 15.6 (b), with a width of 44 MeV/c^2 . A mass window of $\pm 100 \text{ MeV}/c^2$ gives total efficiencies of 24% in the $\rho^\pm\pi^\mp$ channel and 23% in the $\rho^0\pi^0$ channel, for triggered events. The proper-time resolution and acceptance is found to be compatible with the $B_d^0 \rightarrow \pi^+\pi^-$ case. The efficiency of the event selection is summarised in Table 15.4. With this performance the numbers of reconstructed and tagged events in one year of running are expected to be 551, 214 and 491 for $B_d^0 \rightarrow \rho^+\pi^-$, $B_d^0 \rightarrow \rho^-\pi^+$ and $B_d^0 \rightarrow \rho^0\pi^0$ respectively (assuming the branching fractions given in Table 15.10). This yield can be improved by refining the π^0 -finding algorithm to select high-momentum decays.

Although the mass resolution is more than two times worse than for $B_d^0 \rightarrow \pi^+\pi^-$ decays, the sub-mass constraints provide additional tools for background suppression. At present there are inadequate statistics in simulated generic $b\bar{b}$ events to measure the background-to-signal ratio, but an upper limit of 40 has been determined in the $B_d^0 \rightarrow \rho^+\pi^-$ channel.

15.5 $B_d^0 \rightarrow D^{*-}\pi^+$

The selection of $B_d^0 \rightarrow D^{*-}\pi^+$ events has been studied using two methods. In the first the decay

Table 15.4: Summary of event selection in the $B_d^0 \rightarrow \rho\pi$ analysis. Efficiencies are given with respect to the preceding step in the selection. ‘‘LHCb acceptance’’ refers to that fraction of events where both the charged pions are momentum-analysed by the spectrometer, and both photons are seen as isolated clusters above 2 GeV in the ECAL; hence it includes a large part of the π^0 -finding efficiency.

Requirement	$\rho^\pm\pi^\mp$	$\rho^0\pi^0$
LHCb acceptance	0.009	0.013
Trigger	0.10	0.13
Kinematic cuts	0.65	0.80
π^0 identified	0.93	0.95
ρ identified	0.93	0.92
Vertex/kinematic	0.35	0.35
Tight mass window	0.95	0.94
Total efficiency	1.7×10^{-4}	3.9×10^{-4}

chain $\bar{D}^0 \rightarrow K^+\pi^-$ is reconstructed; the second is an inclusive reconstruction of $\bar{D}^0 \rightarrow X^+Y^-$ (+ neutrals), a topology which dominates \bar{D}^0 decays.

In the $\bar{D}^0 \rightarrow K^+\pi^-$ reconstruction, loose RICH criteria are used to select the candidate \bar{D}^0 decay products. The $D^{*-} - \bar{D}^0$ mass peak has a resolution of $1 \text{ MeV}/c^2$ and is shown in Fig. 15.7 (a). A window of $\pm 1.7 \text{ MeV}/c^2$ is imposed on this mass difference. A requirement of $500 \mu\text{m}$ separation in z between the B and primary vertices is made, together with kinematic demands on the transverse momentum of the decay products. The final mass peak has a width of $13 \text{ MeV}/c^2$. Selecting candidates within a window of $\pm 30 \text{ MeV}/c^2$ about the nominal B mass results in a reconstruction efficiency of 45% for events surviving the trigger. The steps in the reconstruction are summarised in Table 15.5. This performance implies 85.9 k reconstructed and tagged events in one year's operation at nominal luminosity.

In the inclusive analysis the B mass is recon-

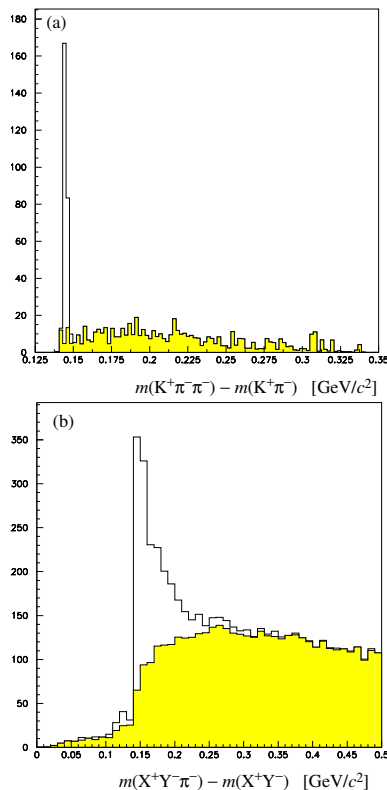


Figure 15.7: (a) The $D^{*-} - \bar{D}^0$ mass peak in the exclusive $B_D^0 \rightarrow D^{*-}\pi^+$ analysis, (b) the $m(\pi^- X^+ Y^-) - m(X^+ Y^-)$ mass peak in the inclusive analysis. The shaded area is the background, superimposed with arbitrary normalisation.

structed solely from the candidate pions of the B and D^{*-} decay. These tracks form a vertex with a resolution of $\sim 35 \mu\text{m}$ in x and y and $\sim 400 \mu\text{m}$ in z . Assuming that the direction of the slow pion from the D^{*-} is that of the D^{*-} itself allows the absolute value of the momentum to be determined, and in turn the B mass. After quality cuts, this mass resolution is $250 \text{ MeV}/c^2$.

Candidate pseudo- \bar{D}^0 s are formed from oppositely-charged tracks with significant impact parameter, and combined with the slow pion. This enables a cut to be made on the mass difference $m(\pi^- X^+ Y^-) - m(X^+ Y^-)$. This distribution is shown in Fig. 15.7(b). By cutting at $< 200 \text{ MeV}/c^2$ a substantial fraction of the background is removed. Additional background rejection comes from requirements on the B and \bar{D}^0 vertices, and on the transverse momentum of the pion from the B decay. A total reconstruction efficiency of 7% is obtained for those events passing the trigger. A total of 446 k events are expected to be reconstructed and tagged for one year's running at nominal luminosity.

With present inclusive $b\bar{b}$ statistics the signal-to-background ratio is estimated to be better than two in both analyses.

15.6 $B_D^0 \rightarrow J/\psi K_S$

The reconstruction of the $B_D^0 \rightarrow J/\psi K_S$ channel begins by selecting lepton pairs with opposite charges

Table 15.5: Summary of event selection in the $B_D^0 \rightarrow D^{*-}\pi^+$ analysis. Efficiencies are given with respect to the preceding step in the selection. In the exclusive case ‘‘LHCb acceptance’’ refers to that fraction of events where all decay products are momentum-analysed in the spectrometer. In the inclusive case the definition is loosened to the requirement that at least two oppositely-signed charged tracks from the \bar{D}^0 decay chain are found.

Requirement	$D^0 \rightarrow K^+\pi^-$	$D^0 \rightarrow X^+Y^-$
LHCb acc.	0.079	0.098
Trigger	0.21	0.21
\bar{D}^0 found	0.78	0.99
$D^{*-} - \bar{D}^0$ cut	0.96	0.86
Kinematic cuts	0.76	0.30
Vertex cuts	0.79	0.28
Tight mass cut	0.99	0.92
Total efficiency	7.1×10^{-3}	1.4×10^{-3}

which make good J/ψ candidates, followed by searching for K_S candidates. Finally they are combined to make B_d^0 mesons.

The tracks of the lepton pair selected for the J/ψ reconstruction are required to be measured in the Vertex Detector to ensure a good vertex resolution. Furthermore, the χ^2 per degree of freedom of the vertex fit has to be less than 5. The J/ψ produced at the primary vertex are rejected by requiring the longitudinal distance between the primary and J/ψ vertices to be $> 400 \mu\text{m}$.

The mass resolution of the $J/\psi \rightarrow \mu^+\mu^-$ peak is $9 \text{ MeV}/c^2$, and a mass window of $\pm 27 \text{ MeV}/c^2$ around the nominal J/ψ mass is used to select the signal. In the case of $J/\psi \rightarrow e^+e^-$, a long tail to the lower-mass region is present due to energy losses, therefore a wide mass window of $\pm 1 \text{ GeV}/c^2$ is applied. For those combinations within the mass windows, momenta and vertices are fitted under J/ψ -mass and vertex constraints, requiring a χ^2 per degree of freedom < 5 . The vertex resolution obtained is $153 \mu\text{m}$. Figure 15.8 shows the invariant-mass distribution of $\mu^+\mu^-$ pairs and e^+e^- pairs for signal events that pass all analysis cuts.

For the reconstruction of K_S , two oppositely charged particles identified as compatible with pions with a probability $> 5\%$ are combined. In order to reject pions from the primary vertex, a cut on the impact parameters (d_0) of the two pions,

$$\frac{|d_0^1 - d_0^2|}{d_0^1 + d_0^2} < 0.85,$$

is applied. The χ^2 per degree of freedom of the vertex fit is required to be < 5 and the distance between the primary and the K_S -decay vertices is




Figure 15.8: Invariant mass distribution for (a) $\mu^+\mu^-$ and (b) e^+e^- pairs, for $B_d^0 \rightarrow J/\psi K_S$ events after the cuts described in the text.




Figure 15.9: Invariant-mass distribution of $\pi^+\pi^-$, for $B_d^0 \rightarrow J/\psi K_S$ events after the cuts described in the text.

required to be $> 5 \text{ cm}$. Figure 15.9 shows the invariant mass plot for selected $\pi^+\pi^-$ combinations for triggered $B_d^0 \rightarrow J/\psi K_S$ events. The K_S mass resolution is found to be $3 \text{ MeV}/c^2$ and a mass window of $\pm 10 \text{ MeV}/c^2$ is used to select K_S candidates. For those candidates, a fit with vertex and K_S -mass constraints is applied, requiring a χ^2 per degree of freedom < 5 . The vertex resolution obtained is $450 \mu\text{m}$ along the direction of flight.

Finally, the J/ψ and the K_S are combined to make the B_d^0 . The reconstructed J/ψ vertex is taken for the B_d^0 vertex. The impact parameter of the reconstructed B_d^0 is required to be $< 60 \mu\text{m}$. Figure 15.10 shows the mass distributions for




Figure 15.10: Reconstructed $J/\psi K_S$ invariant mass for (a) $B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$ and (b) for $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$, for signal events after the cuts described in the text.

Table 15.6: Summary of efficiencies to reconstruct $B^0 \rightarrow J/\psi K_S$ decays.

Requirement	$J/\psi(e^+e^-)K_S$	$J/\psi(\mu^+\mu^-)K_S$
LHCb acc.	0.09	0.10
Trigger	0.13	0.25
B_d reconstr.	0.19	0.46
Total	2.3×10^{-3}	11.5×10^{-3}

the reconstructed $B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$ and $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$. The mass resolutions are $7 \text{ MeV}/c^2$ and $20 \text{ MeV}/c^2$ respectively. The mass windows for the accepted B_d^0 candidates are ± 21 and $\pm 30 \text{ MeV}/c^2$, respectively.

The resolution of the B-meson proper time is found to be 0.045 ps for $B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S$ and 0.055 ps for $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$, independent of proper time. The reconstruction efficiency as a function of proper time is fitted to Equation 15.1 with $b = 0$, giving $a = 2.19 \pm 0.12 \text{ ps}^{-1}$.

The geometrical acceptance and trigger and reconstruction efficiencies are summarised in Table 15.6. Since the major source of background is the accidental combination of a real J/ψ produced away from the primary vertex with a real K_S , the J/ψ 's in the signal sample were combined with other possible K_S 's in the event to study the combinatorial background. Those J/ψ 's produced at the primary vertex are largely rejected by the cut in the longitudinal distance between the primary and J/ψ vertices, and by the condition that they should form a good B_d^0 mass when combined with the K_S . The number of background events within the mass window is multiplied by the ratio of the branching fraction for $B_d^0 \rightarrow J/\psi+X$ to that for $B_d^0 \rightarrow J/\psi K_S$. This gives a background-to-signal ratio of 0.4.

For $B_d^0 \rightarrow J/\psi(e^+e^-)K_S$ events, the decay mode $B_d^0 \rightarrow J/\psi K^{*0}$ followed by $K^{*0} \rightarrow K_S \pi^0$ was studied as a source of background, and was found to give a negligible contribution.

15.7 $B_s^0 \rightarrow D_s^- \pi^+$ and $B_s^0 \rightarrow D_s^\mp K^\pm$

The $B_s^0 \rightarrow D_s^- \pi^+$ and $B_s^0 \rightarrow D_s^\mp K^\pm$ channels are chosen to measure Δm_s and the angle γ respectively. The selection and reconstruction of B_s^0 candidates in these channels need to optimise the proper-time resolution which is crucial for resolving the fast B_s^0 oscillations.

All analysis cuts were optimised without any trigger requirement on a sample of 28 k B_s^0 mesons generated within 600 mrad of the beam axis,

and decaying to $D_s^- \pi^+$ followed by $D_s^- \rightarrow \phi \pi^-$ ($\phi \rightarrow K^+ K^-$). The $B_s^0 \rightarrow D_s^- \pi^+$ candidates are then reconstructed from all combinations of two oppositely-charged tracks assumed to be kaons and two oppositely-charged tracks assumed to be pions. Each track must be well measured, i.e. have at least 26 hits in the tracking system including at least two space points in the Vertex Detector, a measured momentum greater than $2 \text{ GeV}/c$ with at least 1.5% precision, and a RICH particle identification compatible with the assumed mass hypothesis at the 5% level.

The D_s^- candidate vertex is obtained from a mass-constrained fit of the two kaons and one of the pions, and a χ^2 probability cut is applied at 2%. No D_s^- mass cut is necessary after the above requirements, but the reconstructed $K^+ K^-$ mass is required to be within $\pm 30 \text{ MeV}/c^2$ of the nominal ϕ mass. The second pion is then added to form a B_s^0 candidate. The reconstructed $D_s^- \pi^+$ mass is required to be within $\pm 35 \text{ MeV}/c^2$ of the nominal B_s^0 mass and the χ^2 probability of the B_s^0 vertex fit to be greater than 1%. The B_s^0 proper time is estimated as

$$t = L m_{B_s} / p_{B_s},$$

where m_{B_s} and p_{B_s} are the nominal B_s^0 mass and the reconstructed B_s^0 momentum. The B_s^0 decay length L is estimated as

$$L = L_S - L_P, \quad L_{S,P} = \vec{x}_{S,P} \cdot \vec{p}_{B_s} / p_{B_s},$$

where \vec{x}_S is the measured B_s^0 decay point and \vec{x}_P is the primary-vertex position determined from the remaining tracks in the event.

In order to reduce the background, candidates where the B_s^0 vertex is reconstructed downstream of the D_s^- vertex or with $L/\sigma_L < 5$ are rejected (σ_L being the event-by-event estimate of the error on L obtained from the covariance matrices of \vec{x}_S and \vec{x}_P). Furthermore, the transverse momentum of the D_s^- with respect to the beam axis is required to be greater than $2 \text{ GeV}/c$. Cuts are finally applied on the impact parameters d_0 with respect to \vec{x}_P : the B_s^0 must point back to the primary vertex ($d_0^{B_s} < 50 \mu\text{m}$) while its decay products must not ($d_0^{D_s} > 50 \mu\text{m}$ and $d_0^\pi > 75 \mu\text{m}$).

The fraction of signal events that are in the detector acceptance, triggered and reconstructed is $(0.55 \pm 0.03)\%$ (see also Table 15.11). The time dependence of this combined efficiency is parameterised using the function given in Equation 15.1 with $b=0$, giving $a = 1.7 \pm 0.2 \text{ ps}^{-1}$. After all trigger and reconstruction requirements, the RMS resolutions on the D_s^- mass, the B_s^0 mass and the B_s^0 momentum are found to be $4.0 \pm 0.1 \text{ MeV}/c^2$,




Figure 15.11: Proper-time resolution for triggered and reconstructed $B_s^0 \rightarrow D_s^- \pi^+$ events.

$11.0 \pm 0.4 \text{ MeV}/c^2$ and $(0.25 \pm 0.01)\%$, respectively. The vertex positions L_P and L_S have core resolutions of $38 \pm 2 \mu\text{m}$ and $162 \pm 9 \mu\text{m}$. This results in an RMS resolution on the proper time of $\sigma_t = 43 \pm 2 \text{ fs}$ (see Figure 15.11).

Negligible background from charm events is expected after trigger and selection. Backgrounds from b-events consists of real D_s^- associated with a random track (“ D_s^- background”) and of fake D_s^- formed from random tracks (“combinatorial background”). Their rate has been estimated from the number of background candidates found in the signal events. In order to increase the statistics, no trigger requirement is made and the B_s^0 mass window is enlarged to $\pm 2 \text{ GeV}/c^2$. The number of observed events (4 for the D_s^- background and 3 for the combinatorial background) are scaled by appropriate factors obtained from the efficiency of the tight B_s^0 mass cut and under the assumption that $\text{BR}(\text{b-hadron} \rightarrow D_s^\pm X)/\text{BR}(\bar{b} \rightarrow B_s^0 \rightarrow D_s^- \pi^+) = 500$. The fractions of D_s^- background and combinatorial background are estimated to be $(1.0 \pm 0.5)\%$ ($< 2\%$ at 90% CL) and $(4.0 \pm 3.7)\%$ ($< 12\%$ at 90% CL) respectively.

The $B_s^0 \rightarrow D_s^\mp K^\pm$ reconstruction is similar to the $B_s^0 \rightarrow D_s^- \pi^+$ reconstruction, except that the K^\pm from the B_s^0 is subject to tighter requirements in order to remove the $B_s^0 \rightarrow D_s^- \pi^+$ background where the pion could be misidentified as a kaon. The K^\pm candidate is required to have a momentum greater than $3 \text{ GeV}/c$ and to be positively identified by the RICH system. Monte Carlo studies

show that a pion-misidentification rate of 0.1% can be achieved in this momentum range, while maintaining a 90% efficiency to identify kaons. This reduces the background due to $B_s^0 \rightarrow D_s^- \pi^+$ decays to less than 2% of the selected candidates. The combined geometrical, trigger and selection efficiency for $B_s^0 \rightarrow D_s^\mp K^\pm$ events is 0.49%.

Other D_s^- decay modes contributing to the $K^+ K^- \pi^-$ final state, in particular $D_s^- \rightarrow K^{*0} K^- \rightarrow (K^+ \pi^-) K^-$, will also be reconstructed and used for both the $B_s^0 \rightarrow D_s^- \pi^+$ and $B_s^0 \rightarrow D_s^\mp K^\pm$ channels. For the events yields presented in Table 15.12, an effective $D_s^- \rightarrow K^+ K^- \pi^-$ branching ratio of 4.0% is assumed, with the same efficiency and purity as for $D_s^- \rightarrow \phi \pi^-$. This effective branching ratio corresponds to the sum of the $\phi \pi^-$ and $K^{*0} K^-$ modes.

15.8 $B_s^0 \rightarrow J/\psi \phi$

In the reconstruction, only those tracks where the kaon hypothesis has the highest probability are considered as kaons. A similar condition is applied for leptons. With those selected tracks, all possible $K^+ K^- \ell^+ \ell^-$ combinations are formed. Only those combinations where at least two tracks are measured with the Vertex Detector (with no less than 3 space points) are kept. Since the opening angle between the kaons from $\phi \rightarrow K^+ K^-$ is small, at least one of the tracks measured in the Vertex Detector is required to be a lepton.

A vertex is fitted with the four tracks and accepted if the χ^2 per degree of freedom is less than 3. The ϕ candidates within a mass window of $400 \text{ MeV}/c^2$ are accepted.

The reconstructed J/ψ 's from muon pairs have a mass resolution of $9.3 \text{ MeV}/c^2$. Those within 3σ are accepted, and B_s mesons are reconstructed with a mass resolution of $11.9 \text{ MeV}/c^2$. Only those B_s candidates with reconstructed mass within $\pm 35.7 \text{ MeV}/c^2$ from the nominal B_s mass are kept. Reconstructed B_s decay vertices have resolutions of $\sigma_x = \sigma_y = 12 \mu\text{m}$ and $\sigma_z = 135 \mu\text{m}$, resulting in a decay-length resolution of $\sigma_l = 137 \mu\text{m}$. The proper-time resolution of the reconstructed B_s meson is $\sigma_t = 0.032 \text{ ps}$. A separation of the primary and secondary vertices, $l/\sigma_l > 3$, is applied.

For the J/ψ candidates from $e^+ e^-$ pairs, those with invariant masses above 2 GeV are selected and combined with a ϕ . Only those B_s candidates with an invariant mass above 4 GeV are considered. A kinematic fit with the B_s and J/ψ masses as constraints is then applied, and only events with a χ^2 per degree of freedom less than 2.7 are kept. The

Table 15.7: Total reconstruction efficiencies for $B_s \rightarrow J/\psi\phi$ decays.

	$J/\psi(ee)\phi$	$J/\psi(\mu\mu)\phi$
LHCb acceptance	0.10	0.11
Level-0 triggered	0.50	0.82
Level-1 and -2	0.33	0.36
B_d^0 reconstructed	0.53	0.68
Total	.009	.022

resolution of the proper time for the reconstructed B_s meson is the same within statistical errors as that obtained from data with muons.

The total reconstruction efficiencies are summarised in Table 15.7. As for the case of $B_d^0 \rightarrow J/\psi K_S$, combinatorial background due to the J/ψ from B-meson decays combined with fake or real ϕ 's is studied using $B_s^0 \rightarrow J/\psi\phi$ events. The background-to-signal ratios are found to be 3% and 30% for the muon and electron samples, respectively.

The time dependence of the reconstruction efficiency is described by the acceptance function defined by Equation 15.1 with $b = 0$, giving $a = 1.5 \text{ ps}^{-1}$ and 2.0 ps^{-1} , for $J/\psi(ee)\phi$ and $J/\psi(\mu\mu)\phi$, respectively.

15.9 $B_d \rightarrow DK^*$

For these decay modes, the decay fractions integrated over proper time are measured. The decay $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ is a self-tagging mode, i.e. no additional flavour tag is required.

Given the small visible branching fractions of 10^{-8} to 10^{-7} , suppressing background while keeping a high efficiency for the signal is a challenge. The following requirements are made for reconstructing $B_d^0 \rightarrow \bar{D}^0 K^{*0}$:

- Charged tracks are required to have at least one hit in the Vertex Detector and the particle identification probability (from the RICH) has to be greater than 0.2, for the selected particle-type hypothesis.
- The \bar{D}^0 and K^{*0} candidates made from $K^+\pi^-$ combinations must form good vertices, i.e. χ^2 per degree of freedom < 10 , and their masses must be within $14 \text{ MeV}/c^2$ or $40 \text{ MeV}/c^2$ of the nominal masses, respectively.
- A ratio $|(d_0^K - d_0^\pi)/(d_0^K + d_0^\pi)|$ of less than 0.9 is required for the impact parameters of the K^* and D-meson decay products.

- For a B_d^0 candidate, the $\bar{D}^0 K^{*0}$ combination is required to have a mass within $25 \text{ MeV}/c^2$ of the B_d^0 mass.
- The B- and D-meson vertices are required to satisfy $0 < z_D - z_B < 3 \text{ cm}$ and the decay length of the B_d^0 is required to be more than $500 \mu\text{m}$ and less than 3 cm .
- The normalised decay distance in z between the \bar{D}^0 and K^{*0} , $(z_D - z_{K^*})/(z_D + z_{K^*})$, is required to be in the range -0.1 to 0.7 .
- The cosine of the angle between the momentum vector and the flight direction of the B_d^0 candidate has to be greater than 0.9999 .
- The absolute value of the cosine of the angle between the momentum vector of the kaon in the D rest frame and the D momentum must be less than 0.9 .
- The separation between the primary vertex and the B vertex is required to exceed 4σ .

$D_1 \rightarrow K^+K^-$ and $\pi^+\pi^-$ decay channels are reconstructed in a similar way.

Table 15.8 summarises the signal efficiencies after various reconstruction steps. Figure 15.12 shows the invariant-mass distribution of the $\bar{D}^0 K^{*0}$ candidates for triggered $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ events in which two good vertices are reconstructed. The mass resolutions of the D^0 and the B_d^0 mesons are 7.0 and




Figure 15.12: Invariant-mass distribution of $K^-\pi^+K^+\pi^-$ for $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ events after a loose selection (shaded area). The distribution is also shown without any particle identification (open area).

Table 15.8: Signal efficiencies after different reconstruction steps for $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ events.

	$\bar{D}^0 K^{*0}$
LHCb acceptance	0.083
Triggered	0.13
B^0 reconstructed	0.50
Topological cuts	0.42
Total	2.3×10^{-3}

13.0 MeV/ c^2 , respectively. The mass distribution obtained after applying particle-identification criteria is superimposed, demonstrating their importance. In inclusive $b\bar{b}$ events the latter result in a suppression factor of more than 100.

In order to reach the required sensitivity with a restricted sample of simulated background events, selection criteria are loosened for those cuts where the individual efficiencies factorise. This has demonstrated that a final background-to-signal ratio of less than 1 can be achieved for the decay $B_d^0 \rightarrow \bar{D}^0 K^{*0}$. With the present statistics of simulated events, the signal-to-background ratio is less than 5 for the $D^0 K^{*0}$ and $D_1 K^{*0}$ modes. The branching fraction of the latter depends in addition on the angle γ and the strong phase Δ .

15.10 $B \rightarrow K^{*0} \gamma$

The photons from $B_d^0 \rightarrow K^{*0} \gamma$ decays have a relatively-large transverse energy compared to the photons from other sources, as shown in Figure 15.13. Therefore, a large suppression of background can be achieved by requiring high- E_T photon candidates. The K^{*0} vertex gives another means to suppress background.

In the reconstruction, the following selection criteria are applied:

- The transverse energy of the photon candidate to exceed 4 GeV;
- Electromagnetic cluster isolation: the ratio of energy deposition in 3×3 cell clusters and 5×5 cell clusters, $E_{3 \times 3} / E_{5 \times 5} > 0.94$ for the inner part of ECAL and > 0.97 for the middle and outer parts;
- Tracks to be consistent with the K^- and π^+ hypotheses with a probability $> 5\%$;
- The mass of $K^- \pi^+$ to lie in the interval $|m(K^- \pi^+) - m(K^{*0})| < 55$ MeV/ c^2 ;




Figure 15.13: The E_T spectrum of the photons from $B_d^0 \rightarrow K^{*0} \gamma$ (points) and $b\bar{b}$ background (histogram).

- The χ^2 per degree of freedom of the secondary-vertex fit to be < 9 ;
- The separation in z between primary and secondary vertices, $\delta z > 1.5$ mm;
- The impact parameter (d_0) to the primary vertex of both tracks to exceed $400 \mu\text{m}$;
- The transverse momentum of the K^{*0} , $p_T(K^{*0})$, to exceed 1 GeV/ c ;
- K^{*0} polarisation cut: the angle between the B_d^0 and K^- in the K^{*0} rest frame to satisfy

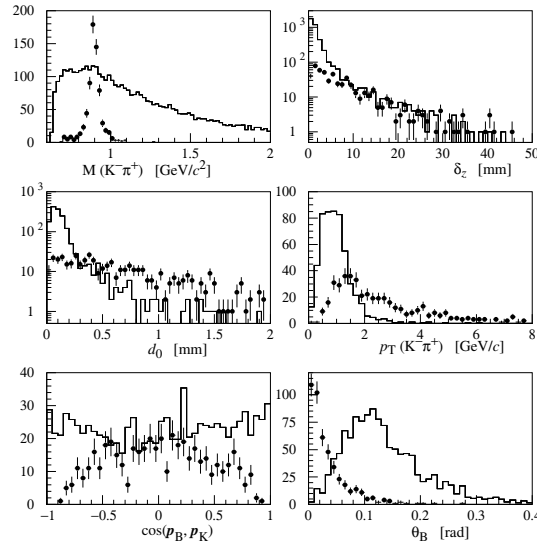


Figure 15.14: The distributions of the main selection variables for $B_d^0 \rightarrow K^{*0} \gamma$ (points) and inclusive $b\bar{b}$ background (histogram).

Table 15.9: Summary of the efficiencies for $B_d^0 \rightarrow K^{*0}\gamma$ decays, and for $b\bar{b}$ background in 4 π . The numbers are given without the trigger selection applied.

Cut	$K^{*0}\gamma$	$b\bar{b}$
LHCb acceptance	0.16	-
L0 Triggered	0.22	0.0059
L1+L2 Triggered	0.30	0.14
K^- and π^+ reconstruction	0.90	-
K^- and π^+ identification	0.90	0.21
K^{*0} mass cut	0.73	0.14
χ^2 p.d.f. secondary vertex	0.98	0.75
$z_{prim} - z_{sec} > 1.5$ mm	0.70	0.11
$d_0 > 400$ μm	0.62	0.064
K^{*0} polarisation	0.82	0.57
$p_T(K^{*0}) > 1$ GeV/c	0.87	0.51
$\theta_B < 0.1$ rad	0.96	0.16
Total efficiency	0.0018	5×10^{-9}

$$\left| \cos(\vec{P}_B, \vec{P}_K) \right| < 0.6;$$

- The opening angle between the momentum vector of the reconstructed B and the flight path from the primary to the secondary vertex, $\theta_B < 0.1$ rad.

Distributions of the discriminating quantities used, for signal and inclusive $b\bar{b}$ events, are shown in Fig. 15.14.

The B-meson invariant mass distribution before and after these cuts is shown on Fig. 15.15 (a) for triggered events. The width of the reconstructed B mass is $67 \text{ MeV}/c^2$. The signal window is defined to be $\pm 200 \text{ MeV}/c^2$ around the nominal B mass.

Background is estimated by scaling the number of accepted events in an enlarged mass window




Figure 15.15: The $K^- \pi^+ \gamma$ mass distribution before (open histogram) and after (hatched histogram) the selection is applied for (a) $B_d^0 \rightarrow K^{*0}\gamma$; (b) $b\bar{b}$ background.

from 4 to 6 GeV/c^2 . The $b\bar{b}$ background mass distribution before and after the selection is applied is shown in Fig. 15.15 (b) for triggered events. Only three events out of 2 million generated $b\bar{b}$ events pass all reconstruction requirements.

Efficiencies for the signal and inclusive $b\bar{b}$ events are summarised in Table 15.9. With $\text{BR}(B_d^0 \rightarrow K^{*0}\gamma) = 4.9 \times 10^{-5}$ [4] the signal-to-background ratio is calculated to be ~ 1 .

15.11 Summary of event yields

In this section, all the necessary quantities for calculating event rates for the pilot channels are summarised. With the following assumptions:

- a 30 MHz bunch-crossing rate for two filled proton bunches colliding at IP8 (see Section 12.2),
- a nominal luminosity of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$,
- an inelastic pp cross-section of 80 mb and a $b\bar{b}$ cross-section of 500 μb ,

only 59% of $b\bar{b}$ pairs are produced in bunch-crossings with a single pp interaction, which are used in the offline analysis. Furthermore, 5% of single pp-interaction bunch-crossings are rejected by the Level-0 pile-up veto. Therefore, only 56% of produced $b\bar{b}$ pairs are used in the offline analysis; i.e. 5.6×10^{11} useful $b\bar{b}$ in one year (10^7 s) of LHCb operation.

The B^+ , B_d^0 , B_s^0 and $\bar{\Lambda}_b$ production rates from \bar{b} quarks are taken to be 40%, 40%, 12% and 8%, respectively. The same fractions are assumed for the charge-conjugated states. Numbers of B mesons available for the offline analysis in one year are then given by:

$$\begin{aligned} B^+ X_b + \text{c.c.} &= 4.5 \times 10^{11} \\ B_d^0 X_b + \text{c.c.} &= 4.5 \times 10^{11} \\ B_s^0 X_b + \text{c.c.} &= 1.3 \times 10^{11} \end{aligned}$$

where X_b and c.c. stand for hadrons with a b-quark, and charge-conjugated states, respectively.

The branching fractions relevant to the physics performance studies for B_d^0 and B_s^0 decays are shown in Table 15.10.

Table 15.11 shows visible branching fractions, geometrical acceptances of the LHCb spectrometer for all the decay products of a given final state, trigger efficiencies for the accepted events, and offline-reconstruction efficiencies for those geometrically accepted and triggered. It also shows the tagging efficiencies with e, μ and kaon tags.

Table 15.10: Assumed branching ratios for B_d^0 and B_s^0 decays.

Mode	B.R.	Comments
$B_d^0 \rightarrow$		
$\pi^+\pi^-$	0.7×10^{-5}	calculated from [3]
$\pi^0\pi^0$	0.3×10^{-5}	calculated from [3]
$K^+\pi^-$	1.5×10^{-5}	Ref. [3]
$\rho^+\pi^-$	1.8×10^{-5}	scaling $B_d^0 \rightarrow \pi^+\pi^-$
$\rho^-\pi^+$	0.7×10^{-5}	scaling $B_d^0 \rightarrow \pi^+\pi^-$
$\rho^0\rho^0$	0.7×10^{-5}	scaling $B_d^0 \rightarrow \pi^0\pi^0$
$D^{*-}\pi^+$	2.6×10^{-3}	Ref. [4]
$J/\psi K_S$	4.4×10^{-4}	half of $B_d^0 \rightarrow J/\psi K^0$ [4]
$J/\psi K^{*0}$	1.5×10^{-3}	Ref. [4]
$\bar{D}^0 K^{*0}$	1.3×10^{-5}	estimate
$D^0 K^{*0}$	1.7×10^{-6}	scaling $B_d^0 \rightarrow \bar{D}^0 K^{*0}$
$D_1 K^{*0}$	$\propto f(\gamma, \Delta)$	depends on γ and Δ
$K^{*0}\gamma$	4.9×10^{-5}	Ref. [4]
$B_s^0 \rightarrow$		
π^+K^-	0.7×10^{-5}	same as $B_d^0 \rightarrow \pi^+\pi^-$
K^+K^-	1.5×10^{-5}	same as $B_d^0 \rightarrow K^+\pi^-$
$D_s^-\pi^+$	3.0×10^{-3}	same as $B_d^0 \rightarrow D^-\pi^+$ [4]
$D_s^-\pi^+$	3.0×10^{-3}	same as $B_d^0 \rightarrow D^-\pi^+$ [4]
$D_s^{*-}\pi^+$	3.0×10^{-3}	same as $D_s^-\pi^+$
$D_s^-K^+$	2.0×10^{-4}	scaling $B_s^0 \rightarrow D_s^-\pi^+$
$D_s^+K^-$	3.1×10^{-5}	scaling $B_d^0 \rightarrow D_s^+D^-$ [4]
$J/\psi\phi$	9.3×10^{-4}	Ref. [4]

Based on Table 15.11 and the numbers of B-mesons produced, the number of events for each final state reconstructed (and tagged, in the case of flavour non-specific decays) are summarised in Table 15.12 for one year (10^7 s) of LHCb operation. The wrong-tag fractions and the background-to-signal ratios are also shown.

References

- [1] T.Sjöstrand, CERN-TH 7112/93 (1993).
- [2] CERN Program Library Long Writeup W5013 (1993).
- [3] R. Godang et al. (CLEO), CLNS 97-1522 (1997).
- [4] R.M. Barnett et al., Physical Review **D54**, (1996) 1; and 1997 off-year partial update for the 1998 edition available on the PDG WWW pages (URL: <http://pdg.lbl.gov/>).

Table 15.11: Summary of visible branching fractions, geometrical acceptance, trigger, offline reconstruction and tagging efficiencies, calculated for B_d^0 and B_s^0 . The function $f(\gamma, \Delta)$ is defined in Equation 16.17.

Mode	Visible Br.	Acc. ($\times 10^{-2}$)	Trigger ($\times 10^{-2}$)	Offline ($\times 10^{-2}$)	Tag (e+ μ +K) ($\times 10^{-2}$)
$B_d^0 \rightarrow$					
$\pi^+\pi^-$	0.7×10^{-5}	14	17	23	40
$K^+\pi^-$	1.5×10^{-5}	14	17	21	-
$\rho^+(\pi^+\pi^0)\pi^-$	1.8×10^{-5}	0.9	10	19	40
$\rho^-(\pi^-\pi^0)\pi^+$	0.7×10^{-5}	0.9	10	19	40
$\rho^0(\pi^+\pi^-\pi^0)$	0.7×10^{-5}	1.3	13	23	40
$D^{*-}[\overline{D}^0(K^+\pi^-\pi^-)]\pi^+$	6.7×10^{-5}	7.9	21	43	40
$D^{*-}[\overline{D}^0(\text{incl.})\pi^-]\pi^+$	1.8×10^{-3}	9.8	21	6.8	40
$J/\psi(e^+e^-)K_S(\pi^+\pi^-)$	1.8×10^{-5}	9.2	13	19	40
$J/\psi(\mu^+\mu^-)K_S(\pi^+\pi^-)$	1.8×10^{-5}	10	25	46	40
$\overline{D}^0(K^+\pi^-)K^{*0}(K^+\pi^-)$	3.3×10^{-7}	8.3	13	21	-
$D^0(K^-\pi^+)K^{*0}(K^+\pi^-)$	4.3×10^{-8}	8.3	13	21	-
$D_1^0(K^+K^-)K^{*0}(K^+\pi^-)$	$\propto f(\gamma, \Delta)$	8.3	13	21	-
$D_1^0(\pi^+\pi^-)K^{*0}(K^+\pi^-)$	$\propto f(\gamma, \Delta)$	8.3	13	21	-
$K^{*0}(K^+\pi^-)\gamma$	3.2×10^{-5}	16	6.6	17	-
$B_s^0 \rightarrow$					
$D_s^-(K^+K^-\pi^-)\pi^+$	1.2×10^{-4}	8.9	16	40	40
$D_s^-(K^+K^-\pi^-)K^+$	8.1×10^{-6}	8.9	16	36	40
$D_s^+(K^+K^-\pi^+)K^-$	1.3×10^{-6}	8.9	16	36	40
$J/\psi(e^+e^-)\phi(K^+K^-)$	2.7×10^{-5}	10	17	53	40
$J/\psi(\mu^+\mu^-)\phi(K^+K^-)$	2.7×10^{-5}	11	30	68	40

Table 15.12: Number of final states which are fully reconstructed, and tagged if required, together with the background-to-signal ratio (B/S) and the wrong-tag fraction (ω_{tag}), for B_d^0 and B_s^0 .

Mode	N_{offline}	B/S	ω_{tag}
$B_d^0 \rightarrow$			
$\pi^+\pi^- + \text{c.c.}$	6900	1	0.3
$K^+\pi^- + \text{c.c.}$	33,000	0.5	-
$\rho^+\pi^- + \text{c.c.}$	551	< 40	0.3
$\rho^-\pi^+ + \text{c.c.}$	214	n.a.	0.3
$\rho^0\pi^0 + \text{c.c.}$	491	n.a.	0.3
$D^{*-}(\text{D excl.})\pi^+ + \text{c.c.}$	85,900	< 0.5	0.3
$D^{*-}(\text{D incl.})\pi^+ + \text{c.c.}$	446,000	< 0.5	0.3
$J/\psi K_S + \text{c.c.}$	55,600	0.4	0.3
$\overline{D}^0 K^{*0} + \text{c.c.}$	337	< 1	-
$D^0 K^{*0} + \text{c.c.}$	45	~ 1	-
$K^{*0}\gamma + \text{c.c.}$	26,000	~ 1	-
$B_s^0 \rightarrow$			
$D_s^-\pi^+ \text{ c.c.}$	86,000	0.05	0.3
$D_s^-\pi^+ + \text{c.c.}$	34,500	0.05	0.3
$D_s^-K^+ + \text{c.c.}$	2100	0.08	0.3
$D_s^+K^- + \text{c.c.}$	320	0.08	0.3
$J/\psi(\mu\mu)\phi + \text{c.c.}$	31,000	0.03	0.3
$J/\psi(ee)\phi + \text{c.c.}$	13,000	0.30	0.3

Chapter 16 CP Sensitivities

16.1 Formalism

Many excellent articles describing CP violation in B-meson decays can be found elsewhere [1]. In this section, the formalism is presented in a way which can be easily applied for both B_d^0 and B_s^0 systems.

The symbols B and \bar{B} are used to denote the particle and antiparticle states for both B_d^0 and B_s^0 systems. The most general expressions of the time-dependent rates for initial B and \bar{B} decaying into a final state f at a given time t can be expressed as:

$$\begin{aligned} R_f(t) &\propto \frac{|A_f|^2}{2} e^{-\Gamma t} [I_+(t) + I_-(t)] \\ \bar{R}_f(t) &\propto \frac{|A_f|^2}{2} \left| \frac{p}{q} \right|^2 e^{-\Gamma t} [I_+(t) - I_-(t)] , \end{aligned} \quad (16.1)$$

where A_f is the instantaneous decay amplitude for $B \rightarrow f$ and Γ is the average decay width for the two mass eigenstates $|B_l\rangle$ and $|B_h\rangle$, which are given by

$$|B_{l(h)}\rangle = \frac{1}{\sqrt{p^2 + q^2}} [p|B\rangle + (-)q|\bar{B}\rangle] , \quad (16.2)$$

i.e. $\Gamma = (\Gamma_l + \Gamma_h)/2$. Two time-dependent functions, $I_+(t)$ and $I_-(t)$ are given by

$$\begin{aligned} I_+(t) &= (1 + |\eta|^2) \cosh \frac{\Delta\Gamma}{2} t - 2\Re \eta \sinh \frac{\Delta\Gamma}{2} t \\ I_-(t) &= (1 - |\eta|^2) \cos \Delta m t - 2\Im \eta \sin \Delta m t \end{aligned} \quad (16.3)$$

where the decay width- and mass-difference are given by $\Delta\Gamma = \Gamma_l - \Gamma_h$ and $\Delta m = m_h - m_l$. The parameter η is given by

$$\eta = \frac{q \bar{A}_f}{p A_f} , \quad (16.4)$$

where \bar{A}_f is the instantaneous decay amplitude for $\bar{B} \rightarrow f$, and p and q are defined in Equation 16.2.

The rates of CP-conjugated final states \bar{f} are

then given by

$$\begin{aligned} \bar{R}_{\bar{f}}(t) &\propto \frac{|\bar{A}_{\bar{f}}|^2}{2} e^{-\Gamma t} [\bar{I}_+(t) + \bar{I}_-(t)] \\ R_{\bar{f}}(t) &\propto \frac{|\bar{A}_{\bar{f}}|^2}{2} \left| \frac{q}{p} \right|^2 e^{-\Gamma t} [\bar{I}_+(t) - \bar{I}_-(t)] , \end{aligned} \quad (16.5)$$

where $\bar{A}_{\bar{f}}$ is the instantaneous decay amplitude for $\bar{B} \rightarrow \bar{f}$. Two time-dependent functions, $\bar{I}_+(t)$ and $\bar{I}_-(t)$ are given by

$$\begin{aligned} \bar{I}_+(t) &= (1 + |\bar{\eta}|^2) \cosh \frac{\Delta\Gamma}{2} t - 2\Re \bar{\eta} \sinh \frac{\Delta\Gamma}{2} t \\ \bar{I}_-(t) &= (1 - |\bar{\eta}|^2) \cos \Delta m t - 2\Im \bar{\eta} \sin \Delta m t \end{aligned} \quad (16.6)$$

where the parameter $\bar{\eta}$ is given by

$$\bar{\eta} = \frac{p \bar{A}_{\bar{f}}}{q A_{\bar{f}}} \quad (16.7)$$

and $A_{\bar{f}}$ is the instantaneous decay amplitude for $B \rightarrow \bar{f}$.

In the time-dependent decay-rate asymmetries, defined as

$$\begin{aligned} \mathcal{A}_f(t) &= \frac{R_f(t) - \bar{R}_f(t)}{R_f(t) + \bar{R}_f(t)} \\ \mathcal{A}_{\bar{f}}(t) &= \frac{R_{\bar{f}}(t) - \bar{R}_{\bar{f}}(t)}{R_{\bar{f}}(t) + \bar{R}_{\bar{f}}(t)} , \end{aligned} \quad (16.8)$$

any acceptance effect of the detector is cancelled, since the asymmetries are built for identical final states. They are therefore better quantities to be studied experimentally than the individual decay rates.

For a CP eigenstate, f and \bar{f} are identical and only the decay rates given in Equation 16.1 need to be considered. The rate asymmetries $\mathcal{A}_f(t) = \mathcal{A}_{\bar{f}}(t)$ become a CP asymmetry. If f is a flavour-specific final state, i.e. $A_{\bar{f}} = 0$ and $\bar{A}_f = 0$, it follows that $\eta = \bar{\eta} = 0$ and

$$\begin{aligned} I_+(t) = \bar{I}_+(t) &= \cosh \frac{\Delta\Gamma}{2} t \\ I_-(t) = \bar{I}_-(t) &= \cos \Delta m t . \end{aligned}$$




Figure 16.1: The tree diagram and a penguin diagram (gluonic) generating $B_d^0 \rightarrow \pi^+ \pi^-$ decays.

Even in the presence of physics beyond the Standard Model, the approximation $|q/p| \approx 1$ is valid up to a few times 10^{-3} . Therefore, it is generally assumed that $|q/p| = 1$. The decay-rate asymmetries (Equation 16.8) are then given by

$$\mathcal{A}_f(t) = \frac{I_-(t)}{I_+(t)}, \quad \mathcal{A}_{\bar{f}}(t) = -\frac{\bar{I}_-(t)}{\bar{I}_+(t)} \quad (16.9)$$

For the B_d^0 meson, $\Delta\Gamma \approx 0$ can be safely assumed, even in the presence of new physics. For the B_s^0 meson, $\Delta\Gamma/\Gamma \approx 0.1$, and cannot be neglected.

It should be noted that, within the Standard Model,

$$B_d^0 : \frac{q}{p} = e^{-2i\beta}, \quad B_s^0 : \frac{q}{p} = e^{2i\delta\gamma},$$

where β and $\delta\gamma$ are angles of the unitarity triangles discussed in Chapter 1.

$B_d^0 \rightarrow J/\psi K_S$: The final state $J/\psi K_S$ is a CP eigenstate and is generated dominantly by the $b \rightarrow c$ tree diagram. In addition, the leading penguin contribution has the same weak phase as the tree diagram; i.e. $A_f/\bar{A}_f = -1$. It follows that the CP asymmetry is given by

$$\mathcal{A}_{J/\psi K_S}(t) = -\sin 2\beta \sin \Delta m t, \quad (16.10)$$

which allows the extraction of β .

$B_d^0 \rightarrow \pi^+ \pi^-$: The final state $\pi^+ \pi^-$ is a CP eigenstate and generated by the $b \rightarrow u$ tree diagram, T , and by the small (but not negligible) $b \rightarrow d$ penguin diagram, P (Fig. 16.1). It follows that

$$\mathcal{A}_{\pi\pi}(t) = a \cos \Delta m t + b \sin \Delta m t, \quad (16.11)$$




Figure 16.2: Two tree diagrams generating B_s^0 and $\bar{B}_s^0 \rightarrow D_s^- K^+$ decays.

where

$$a = 2 \left| \frac{P}{T} \right| \sin \Delta_{P/T} \sin \alpha$$

and

$$b = -\sin 2\alpha - 2 \left| \frac{P}{T} \right| \cos \Delta_{P/T} \cos 2\alpha \sin \alpha,$$

where $\Delta_{P/T}$ is the difference between the strong phases of the penguin and tree diagrams. Note that the triangle relation $\alpha = \pi - \beta - \gamma$ is used to obtain α in the expressions. Once the fraction of the penguin contribution P/T becomes known, α and $\Delta_{P/T}$ can be extracted.

$B_s^0 \rightarrow D_s^+ K^-$: Both B_s^0 and \bar{B}_s^0 can instantaneously decay into $D_s^- K^+$ with tree diagrams T_1 , $b \rightarrow c + W(\rightarrow \bar{u} + s)$, and T_2 , $b \rightarrow u + W(\rightarrow \bar{c} + s)$, respectively (Fig. 16.2). The two branching fractions are of the same order, i.e. $|\eta| = |\bar{\eta}| = |T_2/T_1| \approx \mathcal{O}(1)$. Therefore the rate asymmetries, $\mathcal{A}_{D_s^- K^+}(t)$ and $\mathcal{A}_{D_s^+ K^-}(t)$, can be large. It follows that

$$\begin{aligned} \arg \eta &= -\gamma + 2\delta\gamma + \Delta_{T_1/T_2} \\ \arg \bar{\eta} &= \gamma - 2\delta\gamma + \Delta_{T_1/T_2} \end{aligned} \quad (16.12)$$

where Δ_{T_1/T_2} is the strong-phase difference between T_1 and T_2 . From the two decay asymmetries, $-\gamma + 2\delta\gamma$ and Δ_{T_1/T_2} can be obtained.

$B_s^0 \rightarrow D_s^- \pi^+$: Only the B_s^0 can instantaneously decay into $D_s^- \pi^+$. It follows that

$$\mathcal{A}_{D_s^- \pi^+}(t) = \frac{\cos \Delta m t}{\cosh \frac{\Delta\Gamma}{2} t}, \quad (16.13)$$

which can be used to measure $\Delta m(B_s)$.




Figure 16.3: Two colour-suppressed tree diagrams generating $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ and $D^0 K^{*0}$ decays.

$B_s^0 \rightarrow J/\psi\phi$: The $J/\psi\phi$ is a mixture of $CP=+1$ and $CP=-1$ states, depending on the orbital angular momentum between the J/ψ and ϕ . It follows that

$$\eta = \frac{1-r}{1+r} e^{2i\delta\gamma}, \quad (16.14)$$

where r is the ratio of the B_s^0 decay amplitudes into the $CP=-1$ and $CP=+1$ states. If the two amplitudes have the same strong-interaction phase, then r is real. In this case, r and $2\delta\gamma$ can be obtained from the decay asymmetry

$$\begin{aligned} \mathcal{A}_{J/\psi\phi}(t) &= \frac{(1-|\eta|^2) \cos \Delta m t - 2\Im \eta \sin \Delta m t}{(1+|\eta|^2) \cosh \frac{\Delta\Gamma}{2} t - 2\Re \eta \sinh \frac{\Delta\Gamma}{2} t}. \end{aligned} \quad (16.15)$$

$B_d^0 \rightarrow \bar{D}^0 K^{*0}$: The final states $\bar{D}^0 K^{*0}$, $D^0 K^{*0}$ and $D_1 K^{*0}$, where D_1 is the $CP=+1$ eigenstate, are all flavour specific, i.e. the K^{*0} indicates that they are the decay products of a B_d^0 at the time of decay. Therefore, no flavour tag is required for measuring the branching fractions. The decays $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ and $B_d^0 \rightarrow D^0 K^{*0}$ are produced by the tree diagrams T_1 , $b \rightarrow c + W(\rightarrow \bar{u} + s)$, and T_2 , $b \rightarrow u + W(\rightarrow \bar{c} + s)$, respectively, which are both colour suppressed (Fig. 16.3). The two branching fractions are expected to be of the same order. Therefore, $D^0 \rightarrow K^-\pi^+$ and $\bar{D}^0 \rightarrow K^+\pi^-$ decays are used to identify the D^0 and \bar{D}^0 , and the doubly Cabibbo-suppressed decays $D^0 \rightarrow K^+\pi^-$ and $\bar{D}^0 \rightarrow K^-\pi^+$ can be safely neglected. The decay amplitudes corresponding to the T_1 and T_2 tree diagrams have the

following relations:

$$\begin{aligned} |A_{\bar{D}^0 K^{*0}}| &= |\bar{A}_{\bar{D}^0 \bar{K}^{*0}}| \equiv a_1 \\ |A_{D^0 K^{*0}}| &= |\bar{A}_{D^0 \bar{K}^{*0}}| \equiv a_2. \end{aligned} \quad (16.16)$$

The decay rates to the CP eigenstates are then given by

$$\begin{aligned} a_-^2 &\equiv |A_{D_1 K^{*0}}|^2 \\ &= \frac{a_1^2 + a_2^2 + 2a_1 a_2 \cos(\gamma - \Delta_{T_1/T_2})}{2} \\ &\equiv f(\gamma, \Delta), \\ a_+^2 &\equiv |\bar{A}_{D_1 \bar{K}^{*0}}|^2 \\ &= \frac{a_1^2 + a_2^2 + 2a_1 a_2 \cos(\gamma + \Delta_{T_1/T_2})}{2}, \end{aligned} \quad (16.17)$$

where Δ_{T_1/T_2} is the strong-interaction phase difference between the T_1 and T_2 diagrams. By measuring six branching fractions, for $B_d^0 \rightarrow \bar{D}^0 K^{*0}$, $D^0 K^{*0}$ and $D_1 K^{*0}$, and $\bar{B}_d^0 \rightarrow D^0 \bar{K}^{*0}$, $\bar{D}^0 \bar{K}^{*0}$ and $D_1 \bar{K}^{*0}$, consistency of the measurements can be used to test Equation 16.16, and γ and Δ_{T_1/T_2} can be extracted using Equation 16.17.

16.2 Rate normalisation

Before extracting the CP -violating parameters from the measured decay-rate asymmetries given in Equation 16.8, the effect of other sources producing a B and \bar{B} decay-rate difference have to be corrected for:

- 1) The production rates η_B and $\eta_{\bar{B}}$ of B and \bar{B} mesons in pp interactions are different, and depend on the B momentum vector.
- 2) The tagging efficiencies ε_B and $\varepsilon_{\bar{B}}$ for B and \bar{B} mesons are different because of charge asymmetries of the detector. This depends in general on the momentum of the tagging particle.

The corrected decay-rate asymmetry is given by:

$$\mathcal{A}_f(t) = \frac{N_f(t) - r_B(t) \times \bar{N}_f(t)}{N_f(t) + r_B(t) \times \bar{N}_f(t)},$$

where $N_f(t)$ ($\bar{N}_f(t)$) is the total number of reconstructed B (\bar{B}) decays at decay eigentime t , and $r_B(t)$ is an average normalisation factor close to unity given by

$$r_B(t) = \frac{\int \frac{\eta_B[\Omega(t)] \varepsilon_B[\Omega(t)]}{\eta_{\bar{B}}[\Omega(t)] \varepsilon_{\bar{B}}[\Omega(t)]} n[\Omega(t)] d\Omega(t)}{\int n(\Omega, t) d\Omega},$$




Figure 16.4: The ratios of the number of produced B-hadrons and \bar{B} -hadrons as a function of the hadron momentum, for B_u (top left), B_d (top right) and B_s (bottom centre) hadrons. Straight-line fits to the data are superimposed.

where $\Omega(t)$ represents the phase space for given eigentime t and $n[\Omega(t)]$ is the distribution of the reconstructed events in this phase space.

A large contribution to the rate normalisation factor r_B can come from the charge asymmetry of the detector, introducing different tagging efficiencies for B and \bar{B} . Such an asymmetry would be caused by the difference in detector efficiencies at the right and left sides of the detector, together with the dipole field, which separates particles according to their charge to the right or left sides. At the position of the calorimeters, this separation is about 80%, and further downstream at the Muon Detector it is almost 100%. Differences in the efficiencies of the detector from one side to the other introduce a rate difference of B and \bar{B} decays. This problem is considerably reduced if the magnet polarity is switched frequently compared to the typical time constant of changes in the detector efficiency.

Once the charge symmetry of the detector is restored, the average normalisation factor $r_B(t)$ will depend only on the difference of the production rates of B and \bar{B} hadrons in pp interactions. This difference can be determined as a function of the B-meson momentum using final states for which no CP violation is expected, e.g. $B_d^0 \rightarrow J/\psi K^{0*}$. The ratio of \bar{B}^0/B^0 as a function of the momentum, predicted by PYTHIA, is shown in Fig. 16.4.

Although these differences are small, they need to be determined experimentally.

Figure 16.4 also shows that the normalisation of the initial B and \bar{B} rate is different depending on the B-meson type, i.e. B_u , B_d or B_s .

If the momentum dependence of the rate normalisation is small (or small within a given B-momentum bin), the rate normalisation simply introduces an offset in the theoretical asymmetry in the following way:

$$\mathcal{A}_{\text{measured}}(t) = \mathcal{A}_{\text{theory}}(t) - \frac{1}{2}(r_B - 1). \quad (16.18)$$

Since the theoretical asymmetry has a time dependence, the rate normalisation could in principle be determined simultaneously with the CP-violation parameters. However, this leads to increased statistical uncertainties for the CP parameters (up to a factor of 2) due to their correlation with the initial rate normalisation.

16.3 Systematic errors

As discussed in the previous section, the initial rate difference of B and \bar{B} mesons in pp interactions will be determined using control channels, where no CP violation is expected, e.g. $B_d^0 \rightarrow J/\psi K^{0*}$. The total statistics from control channels will exceed the statistics from the CP-violating decay channels, and therefore the systematic error due to the initial rate normalisation will always be smaller than the statistical error of the CP parameters. In a difference of CP-violating asymmetries, e.g. $\mathcal{A}_{\pi\pi} - \mathcal{A}_{J/\psi K_S}$, the contribution of the rate normalisation cancels to first order.

The background contamination of the reconstructed B-decay mode will be determined based on the reconstructed invariant-mass distribution. The background can be studied as a function of the decay eigentime, and for \bar{B} and B separately, taking possible CP asymmetries of the background into account.

Difference in the tagging efficiencies between B and \bar{B} mesons can be reduced by regularly switching the magnetic-field polarity, as discussed in the previous section.

Wrong tags dilute the magnitude of the decay-rate asymmetry, since

$$\mathcal{A}_{\text{measured}}(t) = (1 - 2\omega_{\text{tag}}) \mathcal{A}_{\text{theory}}(t),$$

where ω_{tag} , the wrong-tag fraction, has to be known for extracting CP-violation parameters. This is best studied by measuring ω_{tag} with fully-reconstructed charged B decays. This procedure

assumes no correlation between the B used for tagging and the other B. This must eventually be verified with detailed simulation studies.

Another source of systematic error is the proper-time resolution. In order to correct for the proper-time resolution, the resolution function has to be extracted from simulation. This can be checked by comparing the measured lifetime in one decay channel with other decay channels, or by comparing with other experiments.

Since there will be significantly more events available for control channels than signal events, systematic errors are neglected in the sensitivity studies of CP parameters.

Simulations show that a total of 600k $J/\psi K^{*0}$ and $J/\psi \bar{K}^{*0}$ decays and 600k $J/\psi K^\pm$ decays can be reconstructed in one year of data taking. From the $J/\psi K^{*0}$ and $J/\psi \bar{K}^{*0}$ decays, the normalisation factor r_B can be obtained with a precision of 2.0×10^{-3} . Similarly, an error on ω_{tag} of 1.1×10^{-3} can be obtained from the $J/\psi \bar{K}^{*0}$ decays, assuming $\varepsilon_{\text{tag}} = 0.4$ and $\omega_{\text{tag}} = 0.3$.

16.4 Sensitivity to α

16.4.1 Determination from $B_d^0 \rightarrow \pi^+ \pi^-$

The sensitivity to α in the channel $B_d^0 \rightarrow \pi^+ \pi^-$ has been estimated by performing least-squared fits to events generated with a stand-alone simulation program. Samples have been generated containing 6900 $B_d^0 \rightarrow \pi^+ \pi^-$ events, corresponding to the expected yield in one year's running, as given in Table 15.12. The proper-time acceptance and resolution have been applied, and a B/\bar{B} normalisation factor of $r_B = 1$ assumed. The signal has been diluted by the introduction of wrong tags and background events. The statistical fluctuations in these contributions increase the error by a factor

$$\frac{\sqrt{1 + B/S}}{1 - 2\omega_{\text{tag}}}$$

where ω_{tag} is the wrong-tag fraction and B/S is the background-to-signal ratio. The values of these parameters for the $B_d^0 \rightarrow \pi^+ \pi^-$ analysis are summarised in Table 15.12.

The background was generated, and fitted, without CP asymmetry. This is unlikely to be the case in reality, as the residual background from two-body B decays is expected to exhibit some CP violation. However, the isolation of clean samples of these different two-body decay modes, relying on particle identification and the good mass resolution, will enable the background asymmetries to

be measured sufficiently well, such that they will not induce a significant systematic error in the α measurement.

To span the full range of possible values of α , samples were generated at six separate values, between 0° and 180° . In addition to this, two scenarios for the penguin diagram were considered. In the first, the penguin contribution to the CP asymmetry was taken to be zero. In the second, it was assumed to be present with a fraction $|P/T| = 0.20$. For the latter the strong-phase difference between penguin and tree, $\Delta_{P/T}$, was in turn set to 0° , 30° and 60° .

The samples were fitted to Equation 16.11. For the ‘‘no penguin’’ samples only the $\sin \Delta mt$ term was considered. For samples including the penguin contribution, α , $\Delta_{P/T}$ and $|P/T|$ were fitted, but with $|P/T|$ constrained to its actual value with a 10% Gaussian uncertainty, that is $|P/T| = 0.20 \pm 0.02$. The wrong-tag fraction and the B/\bar{B} normalisation factor can be determined from control channels, as explained in Section 16.2, and these parameters were set to their central values.

The results of the fit to the ‘‘no penguin’’ cases are given in Table 16.1. Because of the form of the asymmetry in this scenario, there are only two distinct results for the six considered values of α . The error includes contributions from the uncertainty in the wrong-tag fraction and the normalisation factor. These components are largest for $\alpha = 30^\circ$ where they contribute 0.35° and 0.38° , respectively; they are therefore small compared to the statistical uncertainty. The error of 0.056 on $\sin 2\alpha$ may be compared favourably with the error of 0.039 quoted in [2], where a three-times higher branching ratio was assumed.

The results of the fits to the cases with a penguin contribution are given in Table 16.2. The errors are given both for α and $\Delta_{P/T}$, with the correlation between the two variables in the fit. The form of the expression for the asymmetry leads to a complicated and significant dependence of the precision on α on the value of α and of $\Delta_{P/T}$. Note

Table 16.1: The fitted errors on α and $\sin 2\alpha$ in $B_d^0 \rightarrow \pi^+ \pi^-$ for different values of α with no penguin contribution present. These are for one year of data taking.

α	0°	30°
σ_α	1.60°	3.35°
$\sigma_{\sin 2\alpha}$	0.056	0.052

Table 16.2: The fitted errors on α , $\Delta_{P/T}$ and the correlation between the two, $\rho_{\alpha \Delta_{P/T}}$, in $B_d^0 \rightarrow \pi^+ \pi^-$ for different values of α and $\Delta_{P/T}$. These are for one year of data taking.

		α	30°	60°	90°	120°	150°
$\Delta_{P/T}$	0°	σ_α	5.3°	2.3°	2.2°	6.5°	3.7°
		$\sigma_{\Delta_{P/T}}$	22°	14°	11°	12°	33°
		$\rho_{\alpha \Delta_{P/T}}$	0.2	-0.4	-0.6	-0.2	0.5
	30°	σ_α	6.2°	2.4°	2.5°	9.7°	4.8°
		$\sigma_{\Delta_{P/T}}$	26°	19°	16°	19°	37°
		$\rho_{\alpha \Delta_{P/T}}$	0.3	0.0	0.4	0.4	0.8
	60°	σ_α	9.1°	4.9°	5.0°	8.6°	4.8°
		$\sigma_{\Delta_{P/T}}$	35°	29°	25°	28°	37°
		$\rho_{\alpha \Delta_{P/T}}$	0.7	0.5	0.9	0.7	0.8




Figure 16.5: A fit to the CP asymmetry in $B_d^0 \rightarrow \pi^+ \pi^-$ of the angle α and the strong-phase difference $\Delta_{P/T}$. The sample was generated with $\alpha = 90^\circ$, $\Delta_{P/T} = 0^\circ$.

that when α is zero the asymmetry vanishes, and there is no sensitivity to the penguin contribution. An example fit to a typical data set from one year's operation is shown in Fig. 16.5.

To assess what fraction of the error on α is attributable to the lack of knowledge on the penguin amplitude, the fits have been repeated at $\alpha = 90^\circ$, $\Delta_{P/T} = 30^\circ$, with $|P/T| = 0.20 \pm 0.00$ and $|P/T| = 0.20 \pm 0.10$. When $|P/T|$ is exactly known the error on α decreases by a factor 0.8. When the penguin contribution is only known to 50% the

error increases by a factor of 3. Thus a reasonable knowledge of $|P/T|$ is necessary for a worthwhile measurement of α to be made in the $B_d^0 \rightarrow \pi^+ \pi^-$ channel. Theoretical understanding of penguin processes can be improved, by measuring branching fractions of decays that are dominated by them, such as $B_d^0 \rightarrow K^+ \pi^-$. Many of those decay modes can be studied by LHCb.

16.4.2 Determination from $B_d^0 \rightarrow \rho\pi$

The angle α can also be determined by analysing the sample of selected $B_d^0 \rightarrow \rho\pi$ events. A $B_d^0 \rightarrow \pi^+ \pi^- \pi^0$ Dalitz plot exhibits bands corresponding to the three separate channels, with interference between the Breit-Wigners of the ρ^+ , ρ^- and ρ^0 occurring at the intersection of these bands. Note that the density of events in these interference regions is enhanced by the helicity amplitudes in the ρ decay. With initial-state flavour tagging, a time-dependent analysis can be made within this space, which provides simultaneous information on the tree and penguin amplitudes, and the strong-phase difference between the two. This enables a clean determination of α to be made, as demonstrated in [3] where an unbinned log-likelihood fit to a sample of 1000 signal events gave a precision of $\sigma_\alpha \sim 5^\circ$.

The $B_d^0 \rightarrow \pi^+ \pi^- \pi^0$ Dalitz plot is shown in Fig. 16.6, for reconstructed signal events. The kinematic cuts necessary to suppress the combinatoric background from inclusive $b\bar{b}$ events eliminate the $\rho^+ \rho^-$ interference region; however, the other two corners remain well populated. It is still possible to measure α in this reduced space, but with lower sensitivity.

This analysis can provide both a pure measurement of α and insight into the role of penguin graphs. Over several years significant statistics will be col-




Figure 16.6: The Dalitz plot for $B_d^0 \rightarrow \pi^+\pi^-\pi^0$ decays. Only signal events are shown.

lected in the $B_d^0 \rightarrow \rho\pi$ channels. Studies are in progress to quantify the expected precision on α . It should also be noted that by including $B^\pm \rightarrow \rho\pi$ decays additional information may be obtained on the weak phase of the penguin contributions. Furthermore, the analysis of $B_d^0 \rightarrow \rho K$ modes is experimentally possible at LHCb and of interest in penguin studies [4, 5].

Determination from $B_d^0 \rightarrow D^{*-}\pi^+$

Flavour-tagged samples in the channel $B_d^0 \rightarrow D^{*-}\pi^+$ can be used to form two separate asymmetries from the decay distributions of the initial states B_d^0 and \bar{B}_d^0 into $D^+\pi^-$ and $D^-\pi^+$. These asymmetries depend on the quantities $\gamma+2\beta$ ($\equiv \pi-\alpha+\beta$), a strong-phase difference, and $|\eta|$ given by Equation 16.4. Constraining the value of β from alternative measurements allows the other three variables to be fitted.

This is an attractive method of measuring α , as it gives a result which is theoretically clean. It has been demonstrated in Section 15.5 that LHCb can accumulate a tagged sample of order 500k such events through an inclusive selection of D^{*-} decays. However, because the sensitivity to α is low, the statistical error from the fit is several times larger than in the $B_d^0 \rightarrow \pi^+\pi^-$ analysis. Nevertheless, the channel remains promising and studies are continuing.

Conclusion

LHCb is able to make complementary measurements of the angle α in a variety of channels. Simulation has demonstrated the feasibility of reconstructing the modes $B_d^0 \rightarrow \rho\pi$ and $B_d^0 \rightarrow D^{*-}\pi^+$, and work is ongoing to quantify the α sensitivity of these channels. The measurement of α with the best statistical precision comes from $B_d^0 \rightarrow \pi^+\pi^-$. Here LHCb estimates an error of $\sigma_{\sin 2\alpha} = 0.056$ in one year of operation, for the case where there is no penguin contamination. Because of the kaon rejection provided by the RICH system, the systematic uncertainty induced by two-body backgrounds is negligible. Reasonable knowledge of the size of the penguin amplitude is necessary to retain good precision when this effect is included.

16.5 Sensitivity to β

The performance studies of LHCb have not focussed heavily on the measurement of β , as it is expected that using the “golden” channel $B_d^0 \rightarrow J/\psi K_S$ other experiments, including the general purpose LHC detectors, will be able to perform this measurement.

Equation 16.10 is used to fit the measured CP asymmetries constructed from initial B_d^0 and \bar{B}_d^0 to extract the angle β . From a fit similar to the one used for the $B_d^0 \rightarrow \pi^+\pi^-$ study, the following precision is obtained in the determination of β , for several values of this angle:

$$\begin{aligned} \sigma_{\sin 2\beta} &= 0.017 & (\beta = 0^\circ) \\ \sigma_{\sin 2\beta} &= 0.014 & (\beta = 20^\circ) \\ \sigma_{\sin 2\beta} &= 0.011 & (\beta = 30^\circ) . \end{aligned}$$

16.6 Sensitivities to Δm_s and $\Delta\Gamma_s$

The channel $B_s^0 \rightarrow D_s^-\pi^+$ is chosen to measure Δm_s because the flavour of the B_s^0 at its decay is tagged unambiguously by the sign of the D_s^- . Together with the initial flavour information, this allows a separation between the events that have oscillated and those that have not, and hence the extraction of the oscillation frequency. The $B_s^0 \rightarrow D_s^-\pi^+$ channel also has the largest visible branching ratio of all B_s^0 channels under consideration, and is therefore the most suitable to perform a direct measurement of $\Delta\Gamma_s$, by fitting the sum of two exponential functions to the proper-time distribution of untagged events.

Table 16.3: Statistical uncertainty on $\tau_{B_s} = 1/\Gamma_s$ and $\Delta\Gamma_s/\Gamma_s$ obtained from fits of untagged samples of $B_s^0 \rightarrow D_s^- \pi^+$ candidates, equivalent to one year of data taking.

True $\Delta\Gamma_s/\Gamma_s$	$\sigma(1/\Gamma_s)$	$\sigma(\Delta\Gamma_s/\Gamma_s)$
0.00	± 0.005 ps	± 0.07
0.02	± 0.010 ps	± 0.07
0.04	± 0.009 ps	± 0.06
0.08	± 0.014 ps	± 0.06
0.16	± 0.015 ps	± 0.05
0.40	± 0.017 ps	± 0.02
0.80	± 0.032 ps	± 0.02

In order to study the sensitivity to $\Delta\Gamma_s$ and Δm_s , maximum-likelihood fits are performed on samples of fast Monte Carlo events. The fast generator assumes the parameters given in Table 16.5. Each sample is generated with a statistics corresponding to one year of data taking, i.e. including 86,000 untagged signal events for the $\Delta\Gamma_s$ measurement, or 34,500 tagged signal events for the Δm_s measurement. The background distribution is taken as being the same as the signal distribution for $\Delta\Gamma_s = \Delta m_s = 0$. This is conservative as the average reconstructed proper time will probably be lower for the background than for the signal. The fraction of background tagged as having oscillated is assumed to be 50%. The same average proper-time resolution and average wrong-tag probability are used for each event.

Untagged samples of events are fitted simultaneously for the mean B_s^0 lifetime $\tau_{B_s} = 1/\Gamma_s$ and the relative B_s^0 decay-width difference $\Delta\Gamma_s/\Gamma_s$. The statistical precision obtained on these parameters is given in Table 16.3 and indicates that direct evidence for a non-zero value of $\Delta\Gamma_s/\Gamma_s$ can already be obtained after one year, if $\Delta\Gamma_s/\Gamma_s$ is at least 20%. After ten years of running, the statistical power should be sufficient to distinguish $\Delta\Gamma_s/\Gamma_s = 5\%$ from zero.

Table 16.4 shows the results of fits for Δm_s using tagged events generated with a B_s^0 width difference, chosen to be the central value of a recent theoretical prediction of $\Delta\Gamma_s/\Gamma_s = 0.16^{+0.11}_{-0.09}$ [7]. In these fits both Γ_s and $\Delta\Gamma_s$ are fixed to their true values. Examples of proper-time distributions of the candidates tagged as having oscillated are shown in Fig. 16.7. The damping of the B_s^0 oscillation due to $\Delta\Gamma_s/\Gamma_s$ is not significant at the chosen value of 16%, but could be significant for large values of $\Delta\Gamma_s/\Gamma_s$, as demonstrated in Figure 16.8.




Figure 16.7: Proper-time distributions of $B_s^0 \rightarrow D_s^- \pi^+$ candidates generated with two different values of Δm_s . The curves display the result of the maximum-likelihood fit to the total samples.




Figure 16.8: Fraction of events tagged as having oscillated as a function of proper time, for two different values of $\Delta\Gamma_s/\Gamma_s$ (and $\Delta m_s = 10 \text{ ps}^{-1}$). The curves display the result of the maximum-likelihood fit to the total samples.

Table 16.4: Statistical uncertainty on Δm_s obtained from fits of tagged samples of $B_s^0 \rightarrow D_s^- \pi^+$ candidates, equivalent to one year of data taking.

True Δm_s	$\sigma(\Delta m_s)$
10 ps ⁻¹	± 0.008 ps ⁻¹
20 ps ⁻¹	± 0.011 ps ⁻¹
30 ps ⁻¹	± 0.013 ps ⁻¹
40 ps ⁻¹	± 0.016 ps ⁻¹
50 ps ⁻¹	± 0.018 ps ⁻¹

If the correct minimum of the negative log-likelihood function is found (which should always be the case if the value of Δm_s is not too large), then the statistical uncertainty on the measured value of Δm_s is very small, typically less than 0.1%. However, the probability to find the wrong minimum increases with Δm_s . For very large values of Δm_s a measurement is no longer possible, but individual values of Δm_s can be excluded in the region where the analysis has some sensitivity. The statistical significance \mathcal{S} of a B_s^0 oscillation signal can approximately be written as [6]

$$\mathcal{S} \approx \sqrt{N_{\text{tag}}/2} f_{B_s} (1 - 2\omega_{\text{tag}}) e^{-(\Delta m_s \sigma_t)^2/2}, \quad (16.19)$$

where N_{tag} is the total number of reconstructed and tagged candidates in the sample (signal and background), f_{B_s} the fraction of B_s^0 signal in the sample, ω_{tag} the wrong-tag probability, and σ_t the proper-time resolution (see Table 16.5). The quantity \mathcal{S} decreases very quickly as Δm_s increases and is very sensitive to σ_t , which is therefore a crucial parameter.

A more reliable estimate of \mathcal{S} can be obtained directly from the samples of fast Monte Carlo events using the ‘‘amplitude fit method’’ [6]. In this method a new parameter, the B_s^0 amplitude \mathcal{A} , multiplies the term $\cos(\Delta m_s t)$ in the proper-time distribution functions of the B_s^0 signal. The amplitude \mathcal{A} is measured at fixed values of Δm_s , using a standard maximum-likelihood method. Its statistical uncertainty is an estimate of $1/\mathcal{S}$, which takes into account all the features of the fast Monte Carlo sample, including the proper-time dependence of the efficiency, the distribution of the background, and the oscillation damping due to $\Delta \Gamma_s$. Figure 16.9 shows the significance $\mathcal{S} = 1/\sigma_{\mathcal{A}}$ as a function of Δm_s ; it indicates that one expects to be able to make an observation of B_s^0 oscillations with a statistical significance of at least 5σ (i.e. $\mathcal{S} > 5$) if the true value of Δm_s is less than 48 ps⁻¹, or exclude

Table 16.5: Parameters used in the fast Monte Carlo generator for Δm_s sensitivity studies, using fully-reconstructed $B_s^0 \rightarrow D_s^- \pi^+$ candidates from data taken in one year, and the resulting reach in Δm_s .

Efficiency parameter	$a = 1.7$ ps ⁻¹
Reconstructed events	$N = 90,500$
Rec. and tagged events	$N_{\text{tag}} = 36,300$
B_s^0 purity	$f_{B_s} = 0.95$
Wrong tag probability	$\omega_{\text{tag}} = 0.30$
Proper time resolution	$\sigma_t = 43$ fs
Mean B_s^0 lifetime	$1/\Gamma_s = 1.57$ ps
B_s^0 width difference	$\Delta \Gamma_s = 16\% \cdot \Gamma_s$
5σ measurement of Δm_s	up to 48 ps ⁻¹
95% CL excl. of Δm_s values	up to 58 ps ⁻¹




Figure 16.9: Statistical significance of the B_s^0 oscillation signal as a function of Δm_s . The significance is obtained as $1/\sigma_{\mathcal{A}}$, where $\sigma_{\mathcal{A}}$ is the statistical uncertainty on the oscillation amplitude, measured in a fast Monte Carlo sample generated with $\Delta m_s = \infty$ and $\Delta \Gamma_s/\Gamma_s = 16\%$. The band delimited by the two curves reflects the 1σ statistical uncertainty on the proper-time resolution of $\sigma_t = 43 \pm 2$ fs.

values of Δm_s at 95% CL up to 58 ps⁻¹ (where $\mathcal{S} = 1.645$). These values of Δm_s correspond to $x_s = 75$ and 91, respectively.

Within the Standard Model, one has [7]

$$\Delta \Gamma_s/\Delta m_s = (5.6 \pm 2.6) \cdot 10^{-3}, \quad (16.20)$$

with an uncertainty independent of CKM matrix

elements and mainly due to the hadronic uncertainty, which is expected to decrease in the future. Therefore, if Δm_s is larger than 48 ps^{-1} , a measurement of $\Delta\Gamma_s$ will be possible after one year of data taking (since in that case $\Delta\Gamma_s/\Gamma_s > 0.42 \pm 0.20$ according to Equation 16.20) and will provide an indirect measurement of Δm_s within the Standard Model, although with a large uncertainty.

16.7 Sensitivity to $\gamma - 2\delta\gamma$

The angle $\gamma - 2\delta\gamma$ is measured from B_s decays into $D_s^\pm K^\mp$. In this study, separate maximum-likelihood fits are used for decay asymmetries of the $D_s^+ K^-$ and the $D_s^- K^+$ final states (Equation 16.9). Table 15.12 summarises the expected event yields for one year of data taking. Although the decay-amplitude ratio $|A(B_s \rightarrow D_s^+ K^-)/A(B_s \rightarrow D_s^- K^+)|$ is predicted to be ~ 0.39 , the decay rates for the initial B_s into $D_s^+ K^-$ and into $D_s^- K^+$ are almost identical due to the rapid oscillations.

Events are generated with $\Delta\Gamma_s/\Gamma_s = 0.16$ for various values of $x_s = \Delta m_s/\Gamma_s$ and strong-phase differences Δ_{T_1/T_2} . The free parameters are $|\eta|$ and $\arg \eta$ ($|\bar{\eta}|$ and $\arg \bar{\eta}$) in the fit to $D_s^+ K^-$ ($D_s^- K^+$) final states, respectively (cf Equation 16.12). Since x_s , $\Delta\Gamma_s/\Gamma_s$ and the wrong-tag fraction ω_{tag} are well measured using $B_s^0 \rightarrow D_s^- \pi^+$ decays, they are fixed to the values used for the event generation.

Using Equation 16.12 both the angle $\gamma - 2\delta\gamma$ and the strong-interaction phase difference Δ_{T_1/T_2} can be determined. However, the sensitivity to $\gamma - 2\delta\gamma$ depends on Δ_{T_1/T_2} and x_s , as shown in Table 16.6 for different values of x_s and $\gamma - 2\delta\gamma$ and Δ .

The effect of two contributions to the background have been studied in detail. An uncertainty of 4% in the amount of $D_s \pi$ background and of 10% in the amount of background from other B decays has a negligible effect compared to the statistical error on $\gamma - 2\delta\gamma$, even after 10 years of running. As shown in Section 15.7 the expected background level is much lower than this. Also, an error in the decay-time resolution of 20% does not change the fit results significantly.

It is interesting to note that $\arg \eta$ has a twofold ambiguity when $\Delta\Gamma_s/\Gamma_s = 0$. In this case, $\Re \eta$ is not determined, as can be seen from Equation 16.3. With the statistics expected in one year of running, this ambiguity can be resolved if $\Delta\Gamma_s/\Gamma_s \geq 4\%$. Figure 16.10 shows the negative log-likelihood distribution of the fit as a function of $\arg \eta$, obtained with the events expected in one year of data taking for $x_s = 30$ and $\Delta\Gamma/\Gamma = 0.05$. Although there are two minima, one solution (the correct one) is

Table 16.6: Expected error on $\gamma - 2\delta\gamma$ for one year of data taking.

$\gamma - 2\delta\gamma$	$x_s = 10$	$x_s = 20$	$x_s = 30$
$\Delta = 0^\circ$			
0°	$\pm 3.1^\circ$	$\pm 6.1^\circ$	$\pm 7.2^\circ$
30°	$\pm 5.8^\circ$	$\pm 7.6^\circ$	$\pm 9.0^\circ$
90°	$\pm 10.9^\circ$	$\pm 13.0^\circ$	$\pm 14.6^\circ$
120°	$\pm 9.0^\circ$	$\pm 10.5^\circ$	$\pm 14.2^\circ$
150°	$\pm 5.9^\circ$	$\pm 7.3^\circ$	$\pm 14.0^\circ$
$\Delta = 30^\circ$			
0°	$\pm 6.3^\circ$	$\pm 7.1^\circ$	$\pm 10.3^\circ$
30°	$\pm 8.0^\circ$	$\pm 8.8^\circ$	$\pm 10.4^\circ$
90°	$\pm 9.3^\circ$	$\pm 12.2^\circ$	$\pm 13.7^\circ$
120°	$\pm 8.3^\circ$	$\pm 11.1^\circ$	$\pm 16.5^\circ$
150°	$\pm 6.5^\circ$	$\pm 8.3^\circ$	$\pm 11.9^\circ$
$\Delta = 60^\circ$			
0°	$\pm 9.5^\circ$	$\pm 10.3^\circ$	$\pm 12.2^\circ$
30°	$\pm 9.6^\circ$	$\pm 9.5^\circ$	$\pm 12.5^\circ$
90°	$\pm 5.9^\circ$	$\pm 7.2^\circ$	$\pm 13.8^\circ$
120°	$\pm 6.8^\circ$	$\pm 10.4^\circ$	$\pm 11.3^\circ$
150°	$\pm 8.4^\circ$	$\pm 11.9^\circ$	$\pm 13.9^\circ$




Figure 16.10: Negative log-likelihood distribution of the fit as a function of $\arg \eta$ obtained with the events expected in one year of data taking. Simulated data are generated with $x_s = 30$ and $\Delta\Gamma/\Gamma = 0.05$.

clearly favoured.

Table 16.7: Summary of reconstructed event yields for one year of data taking.

	$\Delta = 0^\circ$	$\Delta = 30^\circ$	$\Delta = 60^\circ$
$B_d^0 \rightarrow D_1 K^{*0}$			
$\gamma = 15^\circ$	22	22	20
$\gamma = 30^\circ$	21	22	21
$\gamma = 45^\circ$	20	22	22
$\gamma = 60^\circ$	18	21	22
$\gamma = 75^\circ$	16	20	22
$B_d^0 \rightarrow D_1 \bar{K}^{*0}$			
$\gamma = 15^\circ$	22	20	16
$\gamma = 30^\circ$	21	18	14
$\gamma = 45^\circ$	20	16	11
$\gamma = 60^\circ$	18	14	9
$\gamma = 75^\circ$	16	11	8

16.8 Sensitivity to γ

Extracting γ requires the measurement of the decay rates of the following channels:

$$\begin{aligned}
 B_d^0 &\rightarrow \bar{D}^0 X \rightarrow K^+ \pi^- X \\
 B_d^0 &\rightarrow D^0 X \rightarrow K^- \pi^+ X \\
 B_d^0 &\rightarrow D_1 X \rightarrow (K^+ K^- \text{ or } \pi^+ \pi^-) X
 \end{aligned}$$

$$\begin{aligned}
 \bar{B}_d^0 &\rightarrow D^0 \bar{X} \rightarrow K^- \pi^+ \bar{X} \\
 \bar{B}_d^0 &\rightarrow \bar{D}^0 \bar{X} \rightarrow K^+ \pi^- \bar{X} \\
 \bar{B}_d^0 &\rightarrow D_1 \bar{X} \rightarrow (K^+ K^- \text{ or } \pi^+ \pi^-) \bar{X}
 \end{aligned}$$

with X being the final state $K^{*0} \rightarrow K^+ \pi^-$.

The expected yields for one year of data taking for reconstructed $B_d^0 \rightarrow \bar{D}^0 K^{*0}$ and $D^0 K^{*0}$, and their charge-conjugated decays, are summarised in Table 15.12. The branching fractions for $B_d^0 \rightarrow D_1 K^{*0}$ and $\bar{B}_d^0 \rightarrow D_1 \bar{K}^{*0}$ depend on the angle γ , as well as on the strong-interaction phase difference Δ , and can be obtained from Equation 16.17. Table 16.7 lists the expected yields for one year of data taking for reconstructed $B_d^0 \rightarrow D_1 K^{*0}$ and $\bar{B}_d^0 \rightarrow D_1 \bar{K}^{*0}$, for various combinations of γ and Δ . The total reconstruction efficiencies of the two decay modes are assumed to be the same as that for $B_d^0 \rightarrow \bar{D}^0 K^{*0}$.

In order to determine γ from the number of observed signal events, the individual channels have to be corrected for relative efficiency differences only. One way that this could be achieved is to compare the numbers of reconstructed $\bar{D}^0 K^{*0}$ and $D^0 \bar{K}^{*0}$, which should be the same within statistical fluctuations. The number of reconstructed $D^0 K^{*0}$

and $\bar{D}^0 \bar{K}^{*0}$ must also be the same. The final statistical error on γ is given in Table 16.8 for different values of γ and Δ . A signal-to-background ratio of 1 has been assumed.

Table 16.8: The expected error on γ for different values of γ and Δ .

	$\Delta = 0^\circ$	$\Delta = 30^\circ$	$\Delta = 60^\circ$
$\gamma = 15^\circ$	4.0	7.9	15.5
$\gamma = 30^\circ$	8.2	7.8	14.5
$\gamma = 45^\circ$	12.0	11.5	13.0
$\gamma = 60^\circ$	15.9	14.5	11.7
$\gamma = 75^\circ$	18.8	16.6	14.6

16.9 Sensitivity to $\delta\gamma$

The expression given in Equation 16.15 is fitted to the observed time-dependent asymmetry between tagged B_s^0 and \bar{B}_s^0 decaying into $J/\psi\phi$, with $\delta\gamma$ and the amplitude ratio r (discussed in Section 16.1) as free parameters. In the fit, perfect knowledge of $\Delta\Gamma$ and Δm_s are assumed. The event yield, wrong-tag fractions and background-to-signal ratios are listed in Table 15.12.

Table 16.9 summarises the uncertainties on r and $\delta\gamma$ obtained with one year of data taking, for various combinations of r and x_s . The value of $\delta\gamma$ is set to 0.

References

- [1] For a recent review, see A.J. Buras and R. Fleischer, TUM-HEP-275/97 (1997).
- [2] LHC-B, Letter of Intent, CERN/LHCC 95-5 LHCC/I8 (1995).

Table 16.9: Sensitivities on $\delta\gamma$ and r for one year of data taking, obtained from B_s^0 and $\bar{B}_s^0 \rightarrow J/\psi\phi$ decays.

$x_s = \Delta m_s / \Gamma_s$		10	20	30
$r = 0$	$\sigma_r [10^{-2}]$	0.82	0.98	1.17
	$\sigma_{\delta\gamma} [10^{-2}]$	0.85	0.97	1.17
$r = 0.15$	$\sigma_r [10^{-2}]$	0.94	1.09	1.27
	$\sigma_{\delta\gamma} [10^{-2}]$	0.93	1.06	1.25
$r = 0.30$	$\sigma_r [10^{-2}]$	1.16	1.24	1.49
	$\sigma_{\delta\gamma} [10^{-2}]$	1.02	1.14	1.39

- [3] A. Snyder and H. Quinn, *Phys. Rev. D* **48** (1993) 2139.
- [4] R. Aleksan et al., *Phys. Lett. B* **356** (1995) 95.
- [5] M. Ciuchini et al., CERN-TH/97-188, July 1997.
- [6] H.-G. Moser and A. Roussarie, *Nucl. Instrum. Methods A* **384** (1997) 491.
- [7] M. Beneke, G. Buchalla and I. Dunietz, *Phys. Rev. D* **54** (1996) 4419.