

CERN LHCC 2003-031
LHCb TDR 10
9 September 2003

LHCb

Trigger System Technical Design Report

Printed at CERN
Geneva, 2003
ISBN 92-9083-208-8

The LHCb Collaboration

Brasilian Center for Research in Physics, CBPF, Rio de Janeiro, Brasil

R. Antunes Nobrega, A. Franca Barbosa¹⁾, I. Bediaga, G. Cernicchiaro, E. Correa de Oliveira, J. Magnin, J. Marques de Miranda, A. Massafferri, E. Polycarpo, A. Reis

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

K. Akiba, S. Amato, T. da Silva, J.R.T. de Mello Neto, B. de Paula, L. de Paula, M. Gandelman, J.H. Lopes, B. Marechal, F. Marinho, D. Moraes²⁾, N. Pelloux¹⁾, C. Perreira Nunes

LAPP Annecy, IN2P3-CNRS, Annecy-Le-Vieux, France

J. Ballansat, D. Boget, P. Delebecque, I. De Bonis, D. Decamp, C. Drancourt, N. Dumont Dayot, C. Girard, B. Lieunard, M.-N. Minard, B. Pietrzyk, H. Terrier

LPC Clermont, IN2P3-CNRS and University Blaise Pascal, Clermont-Ferrand, France

Z. Ajaltouni, G. Bohner, C. Carloganu, R. Cornat, O. Deschamps, P. Henrard, J. Lecoq, R. Lefevre, S. Monteil, P. Perret, C. Rimbault, A. Robert

CPPM Marseille, IN2P3-CNRS and University of Aix-Marseille II, Marseille, France

E. Aslanides, J.-P. Cachemiche, B. Dinkespiler, P.-Y. Duval, R. Le Gac, V. Garonne, O. Leroy, P.-L. Liotard, M. Menouni, A. Tsaregorodtsev, B. Viaud

LAL Orsay, IN2P3-CNRS and University of Paris-Sud, Orsay, France

G. Barrand, C. Beigbeder-Beau, R. Beneyton, D. Breton, O. Callot¹⁾, D. Charlet, B. D'Almagne, B. Delcourt, O. Duarte, F. Fulda Quenzer, B. Jean-Marie, J. Lefrançois, F. Machefert, P. Robbe, M.-H. Schune, V. Tocut, K. Truong

Technical University of Dresden, Dresden, Germany

R. Schwierz, B. Spaan

Max-Planck-Institute for Nuclear Physics, Heidelberg, Germany

M. Agari, C. Bauer, D. Baumeister, J. Blouw, N. Bulian, H.P. Fuchs, W. Hofmann, K.T. Knöpfle, S. Löchner, A. Ludwig, M. Schmelling, B. Schwingenheuer

Physics Institute, University of Heidelberg, Heidelberg, Germany

S. Bachmann, P. Bock, H. Deppe, F. Eisele, S. Henneberger, P. Igo-Kemenes, R. Rusnyak, U. Stange, U. Trunk, M. Walter, D. Wiedner, U. Uwer

Kirchhoff Institute for Physics, University of Heidelberg, Heidelberg, Germany

I. Kisel, V. Lindenstruth, M.W. Schulz²⁾

Laboratori Nazionali dell' INFN, Frascati, Italy

G. Bencivenni, C. Bloise, F. Bossi, P. Campana, G. Capon, P. de Simone, C. Forti, G. Lanfranchi, F. Murtas, L. Passalacqua, V. Patera³⁾, M. Poli Lener, A. Sciubba³⁾

University of Bologna and INFN, Bologna, Italy

G. Avoni, G. Balbi, M. Bargiotti, A. Bertin, M. Bruschi, A. Carbone, S. de Castro, P. Faccioli, L. Fabbri, D. Galli, B. Giacobbe, F. Grimaldi, I. Lax, U. Marconi, I. Massa, M. Piccinini, N. Semprini-Cesari, R. Spighi, V. Vagnoni, S. Vecchi, M. Villa, A. Vitale, A. Zoccoli

University of Cagliari and INFN, Cagliari, Italy

W. Bonivento, S. Cadeddu, A. Cardini, V. de Leo, C. Deplano, A. Lai, D. Raspino, B. Saitta

University of Ferrara and INFN, Ferrara, Italy

W. Baldini, V. Carassiti, A. Cotta Ramusino, P. Dalpiaz, S. Germani, A. Gianoli, M. Martini, F. Petrucci, M. Savrié

University of Florence and INFN, Florence, Italy

A. Bizzeti, M. Lenti, M. Lenzi, G. Passaleva, P.G. Pelfer, M. Veltri

University of Genoa and INFN, Genoa, Italy

S. Cuneo, F. Fontanelli, V. Gracco, G. Mini, P. Musico, A. Petrolini, M. Sannino

University of Milano-Bicocca and INFN, Milano, Italy

T. Bellunato, M. Calvi, C. Matteuzzi, M. Musy, P. Negri, D. Perego, L. Trentadue⁴⁾

University of Rome, “La Sapienza” and INFN, Rome, Italy

G. Auriemma⁵⁾, V. Bocci, C. Bosio, E. Dane, D. Fidanza⁵⁾, A. Frenkel, G. Martellotti, G. Penso, S. Petrarca, D. Pinci, G. Pirozzi, W. Rinaldi, R. Santacesaria, C. Satriano⁵⁾, A. Satta

University of Rome, “Tor Vergata” and INFN, Rome, Italy

G. Carboni, S. de Capua, D. Domenici, R. Messi, G. Natali, L. Pacciani, E. Santovetti

NIKHEF, The Netherlands

G. van Apeldoorn(i,iii), N. van Bakel(i,ii), T.S. Bauer(i), M. van Beuzekom(i), J.F.J. van den Brand(i,ii), H.J. Bulten(i,ii), M. Doets(i), R. Hierck(i), L. Hommels(i), J. van Hunen(i), E. Jans(i), T. Ketel(i,ii), S. Klous(i,ii), M.J. Kraan(i), M. Merk(i), F. Mul(ii), J. Nardulli(i), A. Pellegrino(i), G. Raven(i,ii), H. Schuijlenburg(i), T. Sluijk(i), P. Vankov(i), J. van Tilburg(i), H. de Vries(i), L. Wiggers(i), M. Zupan(i)

(i) Foundation of Fundamental Research of Matter in the Netherlands

(ii) Free University Amsterdam

(iii) University of Amsterdam

Research Centre of High Energy Physics, Tsinghua University, Beijing, P.R.C.

M. Bisset, J.P. Cheng, Y.G. Cui, Y. Dai, Y. Gao, H.J. He, C. Huang, C. Jiang, Y.P. Kuang, Q.Li, Y.J. Li, Y. Liao, J.P. Ni, B.B. Shao, J.J. Su, Y.R. Tian, Q. Wang, Q.S. Yan

Institute of Nuclear Physics and University of Mining and Metallurgy, Krakow, Poland

K. Ciba, K. Galuszka, L. Hajduk, P. Kapusta, J. Michalowski, B. Muryn, Z. Natkaniec, A. Oblakowska-Mucha, G. Polok, M. Stodulski, M. Witek¹⁾, P. Zychowski

Soltan Institute for Nuclear Studies, Warsaw, Poland

M. Adamus, A. Chlopik, Z. Guzik, A. Nawrot, K. Syrczynski, M. Szczekowski

National Institute for Physics and Nuclear Engineering, IFIN-HH, Bucharest-Magurele, Romania

C. Coca, O. Dima, G. Giolu, C. Magureanu, M. Orlandea, S. Popescu¹⁾, A.M. Rosca⁶⁾, P.D. Tarta

Institute for Nuclear Research (INR), Moscow, Russia

S. Filippov, J. Gavrillov, E. Guschin, V. Kloubov, L. Kravchuk, S. Laptev, V. Laptev, V. Postoev, G. Rybkine, A. Sadovskii, I. Semeniouk, V. Strigin

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

S. Barsuk¹⁾, I. Belyaev¹⁾, B. Bobchenko, V. Dolgoshein, A. Golutvin, O. Gouchtchine, V. Kiritchenko, V. Kochetkov, I. Korolko, G. Pakhlova, E. Melnikov¹⁾, A. Morozov, P. Pakhlov, A. Petriaev, D. Roussinov, V. Rusinov, S. Semenov, S. Shuvalov, A. Soldatov, E. Tarkovski

Budker Institute for Nuclear Physics (INP), Novosibirsk, Russia

K. Beloborodov, A. Berdiouguine, A. Bondar, A. Bozhenok, A. Buzulutskov, S. Eidelman, V. Golubev, P. Krokovnyi, S. Oreshkin, A. Poluektov, S. Serednyakov, L. Shekhtman, B. Shwartz, Z. Silagadze, A. Sokolov, A. Vasiljev

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

K. Beloous, V. Brekhovskikh, R.I. Dzhelyadin, Yu.P. Gouz, I. Katchaev, V. Khmelnikov, V. Kisselev, A. Koblelev, A.K. Konoplyannikov, A.K. Likhoded, V.D. Matveev, V. Novikov, V.F. Obraztsov, A.P. Ostankov, V. Romanovski, V.I. Rykalin, M.M. Shapkin, A. Sokolov, M.M. Soldatov, V.V. Talanov, O.P. Yushchenko

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

G. Alkhazov, V. Andreev, B. Botchine, V. Ganja, V. Goloubev, S. Guetz, A. Kashchuk, V. Lazarev, E. Maev, O. Maev, G. Petrov, N. Saguidova, G. Sementchouk, V. Souvorov¹⁾, E. Spiridenkov, A. Vorobyov, An. Vorobyov, N. Voropaev

University of Barcelona, Barcelona, Spain

E. Aguilo, R. Ballabriga⁷⁾, M. Calvo, S. Ferragut, Ll. Garrido, D. Gascon, R. Graciani Diaz, E. Grauges Pous, S. Luengo⁷⁾, D. Peralta, M. Rosello⁷⁾, X. Vilasis⁷⁾

University of Santiago de Compostela, Santiago de Compostela, Spain

B. Adeva, P. Conde⁸⁾, C. Lois Gomez⁹⁾, A. Pazos, M. Plo, J.J. Saborido, M. Sanchez Garcia, P. Vazquez Regueiro

University of Lausanne, Lausanne, Switzerland

A. Bay, B. Carron, O. Dormond, L. Fernandez, R. Frei, G. Haefeli, J.-P. Hertig, C. Jacoby, P. Jalocha, S. Jimenez-Otero, F. Legger, L. Locatelli, N. Neufeld¹⁾, J.-P. Perroud, F. Ronga, T. Schietinger, O. Schneider, L. Studer, M.T. Tran, S. Villa, H. Voss

University of Zürich, Zürich, Switzerland

R. Bernet, R.P. Bernhard, Y. Ermoline, J. Gassner, St. Heule, F. Lehner, M. Needham, P. Sievers, St. Steiner, O. Steinkamp, U. Straumann, A. Vollhardt, D. Volyansky, M. Ziegler¹⁰⁾

Institute of Physics and Technologies, Kharkiv, Ukraine

A. Dovbnya, Yu. Ranyuk, I. Shapoval

Institute for Nuclear Research, National Academy of Sciences, Kiev, Ukraine

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

University of Bristol, Bristol, UK

N.H. Brook, R.D. Head, A. Muir, A. Phillips, A. Presland, F.F. Wilson

University of Cambridge, Cambridge, UK

A. Buckley, K. George, V. Gibson, K. Harrison, C.R. Jones, S.G. Katvars, J. Storey, C.P. Ward, S.A. Wotton

Rutherford Appleton Laboratory, Chilton, UK

C.J. Densham, S. Easo, B. Franek, J.G.V. Guy, R.N.J. Halsall, G. Kuznetsov, P. Loveridge, D. Morrow, J.V. Morris, A. Papanestis, G.N. Patrick, M.L. Woodward

University of Edinburgh, Edinburgh, UK

R. Chamonal, S. Eisenhardt, A. Khan, J. Lawrence, F. Muheim, S. Playfer, A. Walker

University of Glasgow, Glasgow, UK

A.G. Bates, A. MacGregor, V. O'Shea, C. Parkes, A. Pickford, M. Rahman, F.J.P. Soler¹¹⁾

University of Liverpool, Liverpool, UK

S. Biagi, T. Bowcock, G. Casse, R. Gamet, M. George, D. Hutchcroft, J. Palacios, G. Patel, I. Stavitskiy, M. Tobin, A. Washbrook

Imperial College, London, UK

L. Allebone, G.J. Barber, W. Cameron, D. Clark, P. Dornan, A. Duane, U. Egede, A. Howard, S. Jolly, R. Plackett, D.R. Price, T. Savidge, D. Websdale, R. White

University of Oxford, Oxford, UK

M. Adinolfi, J.H. Bibby, M.J. Charles¹²⁾, C. Cioffi, G. Damerell, N. Harnew, F. Harris, I.A. McArthur, C. Newby, J. Rademacker, L. Somerville, A. Soroko, N.J. Smale, S. Topp-Jorgensen, G. Wilkinson

CERN, Geneva, Switzerland

G. Anelli, F. Anghinolfi, N. Arnaud, F. Bal, A. Barczyk, J.C. Batista Lopes, M. Benayoun¹³⁾, V. Bobillier, A. Braem, J. Buytaert, M. Campbell, M. Cattaneo, Ph. Charpentier, J. Christiansen, J. Closier, P. Collins, G. Corti, C. D'Ambrosio, H. Dijkstra, J.-P. Dufey, D. Eckstein, M. Ferro-Luzzi, W. Flegel, F. Formenti, R. Forty, M. Frank, C. Frei, C. Gaspar, P. Gavillet, A. Guirao Elias, T. Gys, F. Hahn, S. Haider, J. Harvey, J.A. Hernando Morata, E. van Herwijnen, H.J. Hilke, R. Jacobsson, P. Jarron, C. Joram, B. Jost, S. Koestner, D. Lacarrère, M. Letheren, C. Lippmann¹⁴⁾, R. Lindner, M. Losasso, P. Mato Vila, M. Moritz, H. Müller, T. Nakada¹⁵⁾, C. Padilla, U. Parzefall, W. Pokorski, S. Ponce, F. Ranjard, W. Riegler, G. Aglieri Rinella, E.M. Rodrigues¹⁶⁾, D. Rodriguez de Llera Gonzalez, S. Roiser, T. Ruf, H. Ruiz Perez, B. Schmidt, T. Schneider, A. Schopper, A. Smith, F. Teubert, N. Tuning, O. Ullaland, P. Vannerem, W. Witzeling, K. Wylie, Y. Xie

¹⁾ also at CERN, Geneva

²⁾ now at CERN, Geneva

³⁾ also at Dipartimento di Energetica, University of Rome, "La Sapienza"

⁴⁾ also at Università degli Studi di Parma

⁵⁾ also at University of Basilicata, Potenza

⁶⁾ also at Humbolt University, Berlin

⁷⁾ also at departament d'Engineria Electronica La Salle, Universitat Ramon Llull, Barcelona

⁸⁾ now at DESY, Hamburg

⁹⁾ now at University of Zürich

¹⁰⁾ now at University of California, Santa Cruz

¹¹⁾ also at Rutherford Appleton Laboratory, Chilton

¹²⁾ now at University of Iowa, Iowa City

¹³⁾ now at Universités de Paris VI et VII (LPNHE), Paris

¹⁴⁾ now at GSI, Darmstadt

¹⁵⁾ also at Lausanne, on leave from PSI, Villigen

¹⁶⁾ supported by a Marie Curie Fellowship under contract number HPMF-CT-2002-01708

Technical Associates Institutes

Espoo-Vantaa Institute of Technology, Espoo, Finland

Ecole d'ingénieurs, Geneva, Switzerland

Acknowledgements

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.

“Trigger” on cover page:
The compact edition of the Oxford English Dictionary
Complete text reproduced micrographically
2 volumes
Oxford University Press
(23 edition, 1984)

Contents

1	Introduction	1
1.1	Physics requirements	1
1.2	Front-End Architecture Requirements	2
1.3	Implementation Overview	3
1.4	Organization of this Document	6
2	Level-0 Calorimeter Triggers	9
2.1	Concepts of the L0 Calorimeter Trigger	9
2.2	L0 Calorimeter Trigger performance	11
2.3	ECAL and HCAL FE card	11
2.4	PreShower FE card	12
2.5	Validation Card	13
2.5.1	ECAL candidates	13
2.5.2	HCAL candidates	14
2.6	SPD Multiplicity	15
2.7	Backplane and links	15
2.8	Selection Crate	15
2.8.1	Electromagnetic Candidates	15
2.8.2	SPD multiplicity	16
2.8.3	HCAL	16
2.8.4	Implementation	16
2.9	Latency	17
2.10	Debugging and Monitoring	17
3	Level-0 Muon Trigger	19
3.1	Overview of the muon system	19
3.2	Trigger implementation	21
3.3	Trigger performance	22
3.3.1	Low-energy background	24
3.3.2	Beam halo muons	24
3.3.3	Hardware parameters	24
3.4	Technical Design	25
3.4.1	ODE Trigger interface	25
3.4.2	Processing Board	26
3.4.3	Muon selection board	27
3.4.4	Controller board	28

3.4.5	Backplane	29
3.4.6	Latency	29
3.4.7	DAQ Event size	30
3.4.8	Debugging and monitoring tools	30
4	Level-0 Pile-Up System	31
4.1	Concept	31
4.2	Simulation Results	32
4.3	Technical Design	33
4.3.1	Beetle Chip	34
4.3.2	Prototype Vertex Finder Board	35
4.3.3	Trigger System Architecture	36
5	Level-0 Decision Unit	39
5.1	L0DU inputs	39
5.2	L0DU overview	40
5.3	L0DU Prototype	41
5.4	Studies for the final implementation	41
5.5	Debugging and monitoring	42
6	Level-1 and High Level Trigger	43
6.1	Level-1 and HLT hardware implementation	44
6.1.1	Architecture	45
6.1.2	Implementation	48
6.2	The Level-1 Algorithm	51
6.2.1	VELO 2D Track Reconstruction	52
6.2.2	Primary Vertex Search	52
6.2.3	Level-0 Object Matching to 2D Tracks	53
6.2.4	VELO 3D Track Reconstruction	53
6.2.5	Level-0 Object Matching to 3D Tracks	53
6.2.6	VELO TT Matching and Momentum Determination	53
6.2.7	L1 timing	54
6.3	The HLT Algorithm	55
6.3.1	VELO tracking	56
6.3.2	VELO TT tracking	56
6.3.3	Long tracking	56
6.3.4	Tracking summary	57
7	Performance Simulation	59
7.1	Performance of Level-0	59
7.1.1	Bandwidth Division	60
7.2	Performance of Level-1	62
7.2.1	Generic Algorithm	63
7.2.2	Specific Algorithm	63
7.2.3	Final decision	64
7.2.4	Efficiencies and bandwidth division	64
7.3	Performance of the High Level Trigger	64

7.3.1	Level-1 Confirmation	66
7.3.2	Exclusive Selection	66
7.4	Trigger Performance Robustness	67
7.4.1	Resolutions	68
7.4.2	Execution Time and Event Size	68
7.4.3	Performance	68
8	Project Organization	71
8.1	Cost and Funding	71
8.2	Schedule	71
8.3	Division of Responsibilities	71
A	Scalability of Level-1	75
	References	78

Chapter 1 Introduction

The LHCb experiment [1, 2] is designed to exploit the large number of $b\bar{b}$ -pairs produced in pp interactions at $\sqrt{s}=14$ TeV at the LHC, in order to make precise studies of CP asymmetries and rare decays in b-hadron systems. LHCb is a single arm spectrometer covering the range¹ $1.9 < \eta < 4.9$. The spectrometer is shown in Figure 1.1, and consists of the Vertex Locator (VELO), the Trigger Tracker (TT), the dipole magnet, two Ring Imaging Cherenkov detectors (RICH1&2), three tracking stations T1–T3, the Calorimeter system and the Muon system. LHCb uses a right handed coordinate system with the z -axis pointing from the interaction point towards the muon chambers along the beam-line, and the y -axis is pointing up-wards. All of the systems are described in detail in their respective Technical Design Reports [3-10], apart from TT and a new layout of RICH1 which are described in the LHCb Reoptimization TDR [2]. The LHCb experiment plans to operate at an average luminosity of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$, i.e. much lower than the maximum design luminosity of the LHC, which makes the radiation damage more manageable. A further advantage is that at this luminosity the number of interactions per crossing is dominated by single interactions, which facilitates the triggering and reconstruction by assuring low channel occupancy. Due to the LHC bunch structure and low luminosity the frequency of crossings with interactions visible² by the

spectrometer is about 10 MHz, which has to be reduced by the trigger to a few hundred Hz, at which rate the events are written to storage for further offline analysis. This reduction is achieved in three trigger levels: Level-0 (L0), Level-1 (L1) and the High Level Trigger (HLT). Level-0 is implemented in custom electronics, while Level-1 and the HLT are executed on a farm of commodity processors.

In the remainder of this introduction the requirements from both the physics and the Front-End implementation will be given, followed by an overview of the whole trigger architecture.

1.1 Physics requirements

At a luminosity of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ the 10 MHz of crossings with visible pp interaction are expected to contain a rate of about 100 kHz of $b\bar{b}$ -pairs. However, in only about 15% of the events will at least one B-meson have all its decay products contained in the acceptance of the spectrometer. Furthermore the branching ratios of B-mesons used to study CP violation are typically less than 10^{-3} . The offline selections exploit the relatively large b mass and lifetime to select those b-hadrons, and stringent cuts have to be applied to enhance signal over background and thus increase the CP sensitivity of the analysis. Hence the requirement for the trigger is to achieve the highest efficiency for these offline selected events. The

¹The pseudo-rapidity $\eta = -\ln(\tan(\theta/2))$ where θ is the angle between a track and the beam-line.

²An interaction is defined to be visible if it produces at

least two charged particles with sufficient hits in the VELO and T1–T3 to allow them to be reconstructible.




Figure 1.1: The layout of the LHCb spectrometer showing the VELO, the two RICH detectors, the four tracking stations TT and T1–T3, the 4 Tm dipole magnet, the Scintillating Pad detector (SPD), Preshower (PS), Electro-magnetic (ECAL) and Hadronic (HCAL) Calorimeters, and the five muon stations M1–M5.

trigger should be able to achieve high efficiency for a large variety of final states.

The trigger should allow overlapping and pre-scaled triggers. It should be possible to emulate the trigger from the data written to storage, which will give an additional handle on trigger efficiencies and possible systematics.

1.2 Front-End Architecture Requirements

A detailed description of the requirements to the Front-End (FE) electronics can be found in [11]. A schematic overview of the FE system is given in Figure 1.2. Here we limit ourselves to those requirements which have a particular influence on the design of the Level-0 and Level-1 triggers. The Level-0 and Level-1 decisions are transmitted to the FE electronics by the TTC system [12].

The latency of Level-0, which is the time elapsed between a pp interaction and the arrival of the Level-0 trigger decision at the

FE, is fixed to $4 \mu\text{s}$ [13]. This time includes the time-of-flight, cable length and all delays in the FE, leaving $2 \mu\text{s}$ for the actual processing of the data in the Level-0 trigger to derive a decision. The FE is required to be able to readout events in 900 ns, which combined with the 16-event-deep de-randomizer before the L1-buffer and a 1 MHz Level-0 accept rate gives less than 0.5% deadtime [14]. Level-0 will deliver its decision every 25 ns to the Readout Supervisor [15], which emulates the L0-buffer occupancy, and prevents buffer overflows by throttling the L0 accept rate.

The Level-1 trigger is a variable latency trigger and the data delivered to the trigger processors by the FE system must be delivered in chronological order and tagged with bunch and event identifiers. Despite the variable latency, Level-1 delivers a decision for each event in the same order to the Readout Supervisor. The maximum Level-1 output rate is fixed to 40 kHz, and overflow of the Level-1 de-randomizer buffers is prevented by a throttle system controlled by the Readout Supervisor. The




Figure 1.2: The FE architecture of LHCb.

depth of the L1-buffer³, combined with a 1 MHz Level-0 rate, and the requirement to deliver the decisions chronologically ordered allows a latency of up to 58 ms.

1.3 Implementation Overview

Given the physics requirement to achieve a maximum trigger efficiency for offline selected events, the main aim of the trigger

³The size of the Level-1 buffer is 2 M words. With a 36-word-long event fragment, which contains 32 channels and up to four data tags, this allows up to 58254 events to be buffered.

implementation is to enable trigger algorithms to have access to the same data as the offline analysis, and anticipates the selection algorithms as closely as possible at the highest possible rate. Figure 1.3 shows an overview of the sub-detectors participating in the three trigger levels.

Level-0

The purpose of Level-0 is to reduce the LHC beam crossing rate of 40 MHz, which contains about 10 MHz of crossing with visible pp-interactions, to the rate at which in

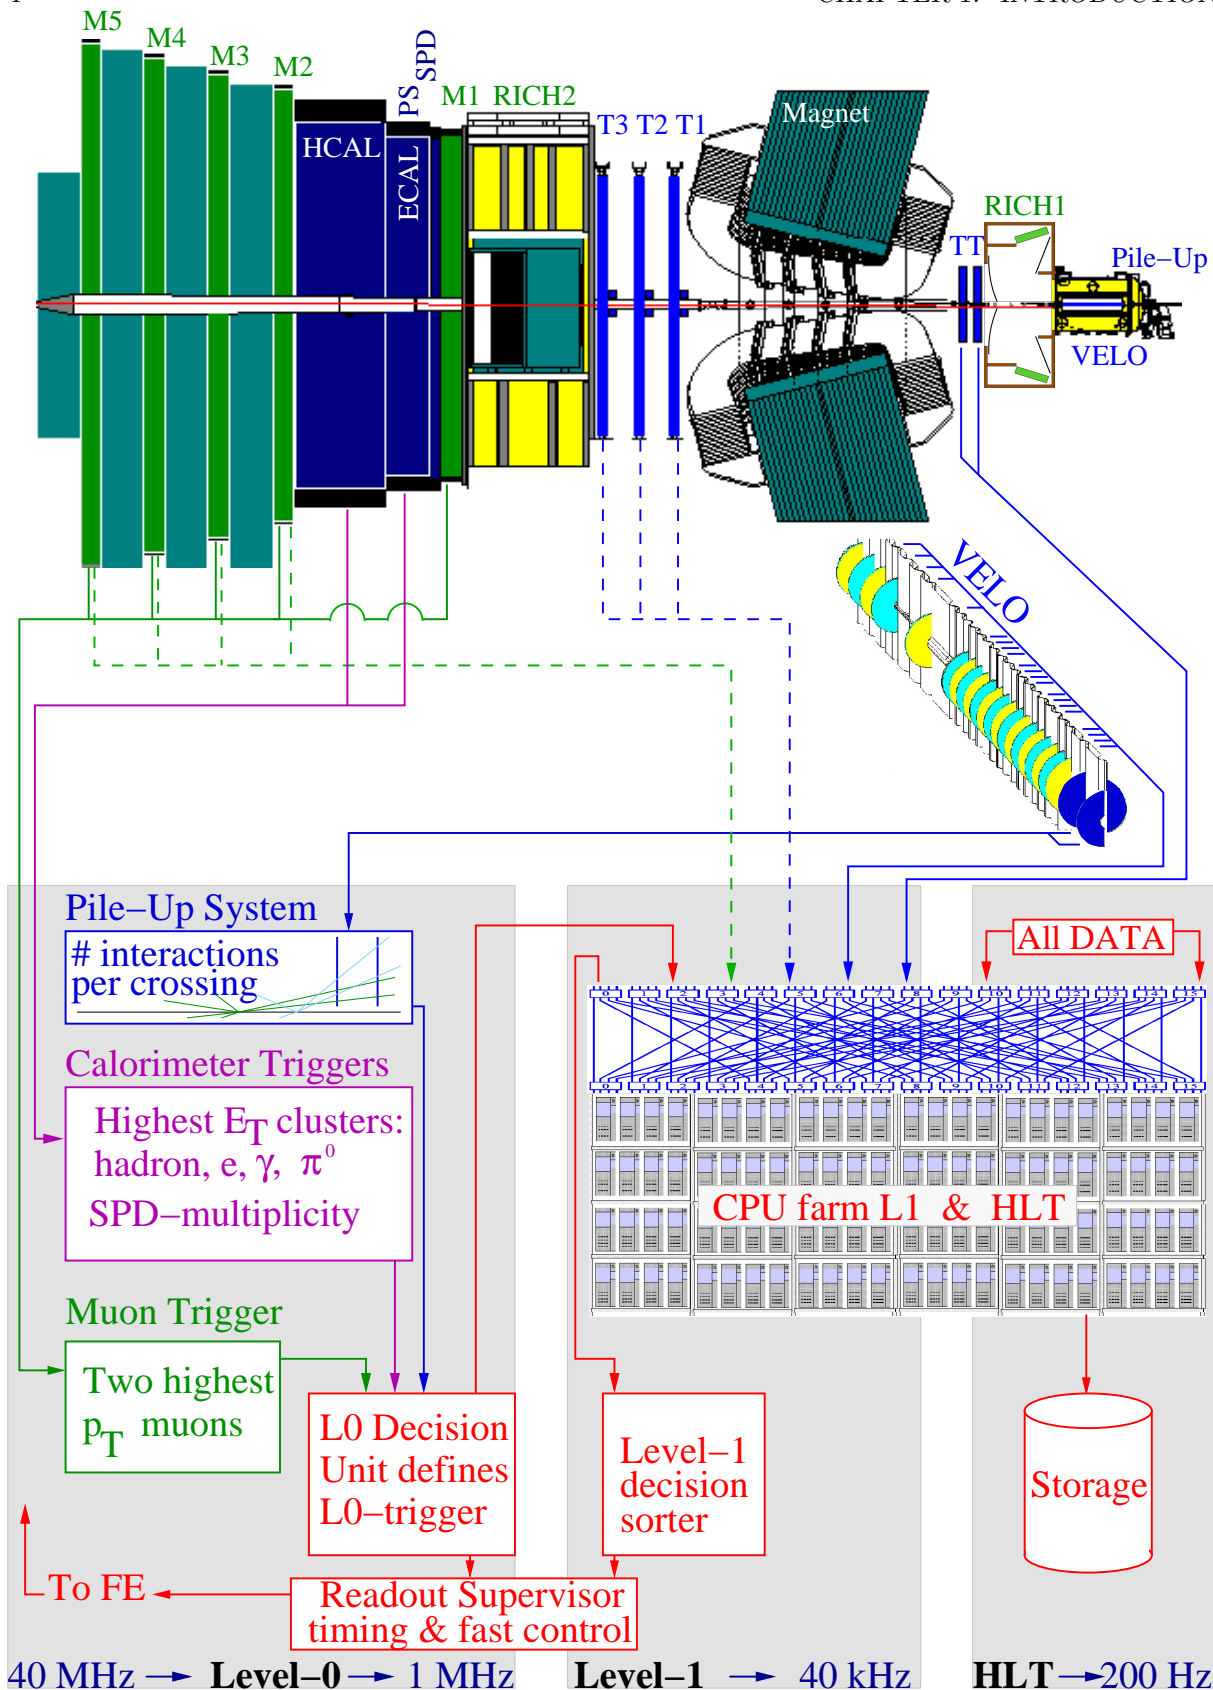


Figure 1.3: Overview of the three trigger levels. Stations M1–M5 are used to reconstruct two muons per quadrant. The SPD, PS, ECAL and HCAL are used to reconstruct the hadron, e , γ and π^0 with the largest transverse energy, the charged particle multiplicity, and the total energy. The Pile-Up detector is used to recognize multiple interactions per crossing. Level-1 uses the information from VELO, TT, and Level-0 to reduce the rate to 40 kHz. T1–T3 and M2–M4 could be included in Level-1. The HLT uses all data in the event apart from the RICH to reduce the rate to 200 Hz. Level-0 is executed in full custom electronics, while Level-1 and HLT are software triggers which share a commodity farm of 1800 CPUs.

principle all sub-systems could be used for deriving a trigger decision. Due to their large mass b-hadrons decay to give a large E_T lepton, hadron or photon, hence Level-0 reconstructs:

- the highest E_T hadron, electron and photon clusters in the Calorimeter,
- the two highest p_T muons in the Muon Chambers,

which information is collected by the Level-0 Decision Unit to select events. Events can be rejected based on global event variables such as charged track multiplicities and the number of interactions, as reconstructed by the Pile-Up system, to assure that the selection is based on b-signatures rather than large combinatorics, and that these events will not occupy a disproportional fraction of the data-flow bandwidth or available processing power in subsequent trigger levels.

All Level-0 triggers are fully synchronous, i.e. their latency does not depend upon occupancy nor on history. All Level-0 electronics is implemented in full custom boards.

The implementation of the calorimeter trigger is based on forming clusters by adding the E_T of 2×2 cells on the FE-boards, and selecting the clusters with the largest E_T . Clusters are identified as e, γ or hadron depending on the information from the Scintillating Pad Detector (SPD), Preshower (PS), Electromagnetic (ECAL) and Hadronic (HCAL) Calorimeter. The E_T of all HCAL cells is summed to reject crossings without visible interactions. The total number of SPD cells with a hit are counted to provide a measure of the charged track multiplicity in the crossing.

The muon chambers allow stand-alone muon reconstruction with a p_T resolution of 20%. Track finding is performed by processing units, which combine the strip and pad data from the five muon stations to form towers pointing towards the interaction re-

gion. One crate per quarter houses the trigger boards which reconstruct the two muons with the largest p_T .

The Pile-Up system aims at distinguishing between crossings with single and multiple visible interactions. It uses four silicon sensors of the same type as those used in the VELO to measure the radial position of tracks, covering $-4.2 < \eta < -2.9$. The Pile-Up system provides the position of the primary vertex candidates along the beam-line and a measure of the total backward charged track multiplicity. The Pile-Up information allows a relative luminosity measurement which is not affected by system downtime, and monitors the beam conditions.

The Level-0 Decision Unit (L0DU) collects all information from Level-0 components to form the Level-0 Trigger. The L0DU is able to perform simple arithmetic to combine all signatures into one decision per crossing. This decision is passed to the Readout Supervisor which transmits its Level-0 decision to the FE.

Level-1

At the 1 MHz output rate of Level-0 the remaining analogue data is digitized and all data is stored for the time needed to process the Level-1 algorithm. All sub-systems which deliver data to Level-1 make use of the same TELL1-board [16] to store the data in the L1-buffer, to perform zero-suppression and formatting, and to interface to Level-1. The Level-1 algorithm will be implemented on a commodity processors farm, which is shared between Level-1, HLT and offline reconstruction algorithms. The Level-1 algorithm uses the information from Level-0, the VELO and TT. The algorithm reconstructs tracks in the VELO, and matches these tracks to Level-0 muons or Calorimeter clusters to identify them and measure their momenta. The fringe field of the magnet between the VELO and TT is

used to determine the momenta of particles with a resolution of 20–40%. Events are selected based on tracks with a large p_T and significant impact parameter to the primary vertex.

The event building architecture is inspired by the one described in the Online System TDR [9], but adapted to profit from new technologies due to the delayed start-up of the LHC. The same event building network is used to collect the Level-1 decisions from all the processors, after which they are sorted according to their Level-0 event number and transmitted to the Read-out Supervisor, which transmits its Level-1 decision to the FE. The maximum Level-1 output rate has been fixed to 40 kHz to allow the FE to execute more elaborate zero-suppression algorithms than the ones used to prepare the information for the Level-1 trigger. The implementation is easily scalable to allow the inclusion of stations T1–T3 and M2–M5. This will improve the Level-1 performance, and this implementation is described in Appendix A, but all performance figures given in this TDR will assume that the Level-1 algorithm will not use T1–T3 or M2–M5.

High Level Trigger

The HLT will have access to all data. Since the design of the Online System [9] the LHCb spectrometer has been re-optimized [2], resulting in a considerably smaller data-size. Level-1 and HLT event building now share the same network, and this TDR supersedes the implementation as described in [9]. The HLT and Level-1 algorithms run concurrently on the same CPU nodes, with the Level-1 taking priority due to its limited latency budget. The HLT algorithm starts with reconstructing the VELO tracks and the primary vertex, rather than having this information transmitted from Level-1. A fast pattern recognition program links the VELO tracks to

the tracking stations T1–T3. The final selection of interesting events is a combination of confirming the Level-1 decision with better resolution, and selection cuts dedicated to specific final states. While the maximum output rates of the first two trigger levels are dictated by the implementations of the FE hardware, the output rate of the HLT is kept more flexible. Considering the channels currently under study one could envisage output rates of a few Hz. However, the RICH information is not currently used by the HLT, and selection cuts have to be relaxed compared to the final selection to study the sensitivity of the selections and profit from refinements to the calibration constants. These considerations lead to an output rate of 200 Hz of events accepted by the HLT. The total CPU farm will contain about 1800 nodes. It is derived from the the expected CPU power in 2007, and performance studies discussed below, that the L1 and HLT algorithms will use about 55% and 25% of the available computing resources respectively. The remaining resources are used to fully reconstruct events accepted by the HLT, including the particle identification, before being written to storage.

1.4 Organization of this Document

While the trigger is logically divided into three trigger levels, its implementation is divided into five hardware sub-systems, four Level-0 sub-systems and one sub-system for Level-1 and HLT combined. The Level-0 sub-systems are the Calorimeter Triggers, the Muon Trigger, the Pile-Up Trigger and the Level-0 Decision Unit. Level-1 and the HLT form one sub-system from the technical design point of view, since they share the same event building network and processor farm. In the following chapters each of the five sub-system

designs is described separately, including R&D and prototyping. Where applicable, the sensitivity of the performance of a subsystem to the LHC environment, so called robustness, will also be included, while the performance of the trigger as a whole will be described in Chapter 7. The last chapter deals with project organization.

Chapter 2 Level-0 Calorimeter Triggers

The purpose of the Calorimeter Triggers is to select and identify particles with high E_T deposit in the calorimeters. A schematic view of the calorimeter is shown on Figure 2.1, showing the four detectors involved:




Figure 2.1: Schematic side view of the calorimeter system.

- The SPD (Scintillator Pad Detector) identifies charged particles, and allows electrons to be separated from photons.
- The PreShower detector, after 2.5 radiation length of lead, identifies electromagnetic particles.
- The electro-magnetic calorimeter ECAL, of the shashlik type, measures the energy of electromagnetic showers.
- The hadronic calorimeter HCAL, made of iron with scintillator tiles, measures the energy of the hadrons.

The first three detectors have the same cell geometry, displayed in Figure 2.2. The cells are about $4 \times 4 \text{ cm}^2$ in the central region, $6 \times 6 \text{ cm}^2$ in the middle region and




Figure 2.2: Layout of the SPD, Preshower and ECAL cells. Each square represents 16 cells.

$12 \times 12 \text{ cm}^2$ in the outer region. The exact size of the cells is proportional to their distance from the vertex in order to obtain a pointing geometry, and the total number of cells in each detector is 5984. The HCAL contains 1468 cells, with only two sizes, $13 \times 13 \text{ cm}^2$ and $26 \times 26 \text{ cm}^2$, such that the HCAL cell boundaries project to ECAL cell boundaries. More details are given in the Calorimeter TDR [4].

2.1 Concepts of the L0 Calorimeter Trigger

The idea of the Calorimeter Triggers is to search for high E_T particles: electrons, photons, π^0 or hadrons. The way to identify each flavour is described in Section 2.5.

Showers are relatively narrow, with energy deposits in a small area. A zone of




Figure 2.3: Overall view of the Calorimeter Triggers.

2 by 2 cells is used, large enough to contain most of the energy, and small enough to avoid overlap between various particles. Only the particle with the highest E_T is looked at. Therefore at each stage only the highest E_T candidate is kept, to minimize the number of candidates to process.

These candidates are provided by a three step selection system:

- A first selection of high E_T deposits is performed on the Front-End (FE) card, which is the same for ECAL and HCAL. Each card handles 32 cells, and the highest E_T sum over the 32 sums of 2×2 cells is selected. To compute these 32 sums, access to cells in other cards is an important issue.
- The Validation Card merges the ECAL with the PreShower and SPD information, prepared by the PreShower FE card, to identify the type of electromagnetic candidate,

electron, photon and π^0 . Only one candidate per type is selected and sent to the next stage. The same card also adds the energy deposited in ECAL in front of the hadron candidates. A similar card computes the SPD multiplicity in the PreShower crates.

- The Selection Crate selects the candidate with the highest E_T for each type, it also produces a measure of the Total E_T in HCAL and the total SPD multiplicity.

An overall view of the Calorimeter Triggers is shown on Figure 2.3. Care has been taken to simplify the connections, and the system is fully synchronous, which will facilitate commissioning and debugging.

The first two steps are performed on the platform of the calorimeter, at a location with a radiation dose below 50 Gy during the whole lifetime of the experiment, and where Single Event Upsets (SEU) are

expected to occur. Each component has been tested for radiation and SEU robustness [17, 18]. Anti-fuse PGAs are used, and “triple voting” techniques for all permanent memories. In this technique, each bit is stored in three copies, and a majority vote is taken before using it, so that one single bit flip does not affect the value.

2.2 L0 Calorimeter Trigger performance

The L0 Calorimeter Trigger provides 7 inputs to the Decision Unit, as shown on Figure 2.3. Frequently, the same energy deposit produces several candidates, as the electron, photon and π^0 triggers are not exclusive. A hadron with a large deposit in ECAL may also produce electron, photon or π^0 candidates. This overlap has advantages in terms of robustness of the system and allows cross-checks of the trigger behaviour, but makes difficult an analysis of the exclusive performance of each type of candidate. The overall performance in term of trigger efficiency is discussed in Chapter 7.




Figure 2.4: Performance of the Hadron Trigger.

To illustrate how the signal events are different from the minimum-bias events, Figure 2.4 shows the fraction of events over a given hadron threshold, as function of this E_T threshold, both for minimum-bias events and for offline selected $B \rightarrow \pi\pi$ events. The nominal threshold is indicated, at which about 7% of the minimum-bias events are kept, with an efficiency around 70% on the signal.

2.3 ECAL and HCAL FE card

The processing on the FE card is described in [19]. It is divided into several steps, performed synchronously in pipeline mode:

- Preparation of the data: the Calorimeter Trigger uses as input the 12 bit ADC value of each cell. This digitisation is already pedestal corrected [20], but has to be converted to 8-bit E_T .
- Collection of the data: In order to compute all 32 sums over 2×2 cells, one has to access the neighbouring cells. A dedicated backplane connects neighbouring cards, while LVDS multiplexed links are used for the other connections.
- Computation of the 2×2 sums in parallel.
- Selection of the highest of the 32 sums, keeping track of its address.
- Computation of the total E_T on the card.
- Sending the result of the processing to the next stage.

The conversion from ADC value to E_T is performed by multiplying the ADC value by a 8 bit scale factor, and selecting the proper 8 bits. If the result would be larger than 8 bits, it is saturated to 255. This multiplication with saturation process is performed in one clock cycle [20]. The nominal scale factor is such that the value of E_T covers the

range 0 to 5.1 GeV. This choice is a compromise between the highest trigger threshold, around 4.5 GeV, and the loss of accuracy due to rounding effects. The gain is adjusted channel by channel to compensate for possible mis-calibration of the PM. The E_T range can be varied, if needed, from about 3 GeV up to 10 GeV full scale.

The next step is to get the E_T value for all the cells involved in the 32 sums. Each FE card covers an area of 8×4 cells. In order to build the 32 sums, 8 right neighbours, 4 top neighbours, and one top-right neighbour are needed, as shown on Figure 2.5.

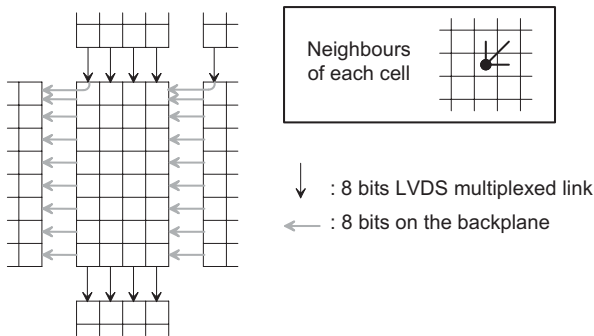


Figure 2.5: Connections to get the neighbouring cells.

Note that neighbouring cells of different size are not connected, as this would introduce unnecessary complications. The two halves of the detector are also not connected, as this would require to disconnect cables when opening the detector for maintenance. These two limitations introduce only a very small inefficiency, at the per mil level.

It is clear that local signals have to be delayed, to allow the remote information to arrive. The slowest signal is the corner one, which arrives in two steps. All other information is placed in pipe-lines, implemented using the internal memory of the PGA, with controllable length. The delay in these pipe-lines will be adjusted once the final layout is made, but it is estimated to be one clock cycle for every backplane link, and two clock cycles for an LVDS multiplexed link, since the cables will be around

2 m long. The corner signal waits 1 cycle, the LVDS signals wait 2 cycles, the backplane signals 3 cycles and the local signals 4 cycles. Note that these local signals are sent on 80 MHz multiplexed lines on the board, to reduce the number of I/O pins on the PGA. Details of this synchronisation, and of the backplane configuration, can be found in [19].

When the 45 signals are available, the 32 sums are computed in parallel. The sum is saturated when it overflows, around 5.1 GeV, which occurs for about one third of the $B \rightarrow \pi\pi$ events as can be seen on Figure 2.4. Saturation has no effect on the performance, as the trigger relies only on the presence of a high E_T cluster. The highest of the 32 sum is then selected, in a series of binary comparisons: First 16 comparisons of two sums, then 8 comparisons of the 16 previous results, and so on. Five steps are needed to get the highest E_T sum of the FE card, with its address on 5 bits. This is performed in pipeline mode, 3 cycles are needed for the sum and the comparisons.

The Total E_T of the card is also produced, by summing the appropriate 8 sums of 4 cells. This result is also saturated. This total sum is the input for the local π^0 trigger.

The card produces a 21 bit output on the backplane, for the Validation Card: 8-bit highest E_T sum, 5-bit address and 8-bit total E_T . It also sends on two multiplexed LVDS links, towards the PreShower FE card for ECAL and towards the Validation Card for HCAL, the 8-bit highest E_T sum and its 5-bit address, together with 8-bits Bunch Crossing Identifier (BCID) for synchronisation.

2.4 PreShower FE card

The PreShower FE card digitizes the PreShower analog signals, corrects them for pedestal, gain and for the spill-over

of earlier signals, and receives and synchronises the SPD information [21]. The PreShower trigger information is obtained by comparing the PreShower signal to a threshold, producing a Yes/No output. The SPD also provides binary information. The PreShower FE card handles 64 channels, and covers two ECAL cards. For each beam crossing, the PreShower FE card receives the address of two ECAL candidates, and for each candidate, sends to the corresponding Validation Card the 4 PreShower bits and the 4 SPD bits in front of the 4 cells of the candidate [21].

As for the Calorimeter card, access to the neighbouring information is needed. The same backplane is used, transporting this time only 2 bits per cell, but there are 8 vertical neighbours instead of 4, as the card covers 8×8 cells. The core of the processing is a 81×2 bit wide pipe-line (as the ECAL input arrives several clock cycles after the PreShower and SPD signals are ready) and the appropriate multiplexer to extract the 4×2 bits of the wanted cells for the proper beam crossing. A prototype to test this processing has been built and is shown on Figure 2.6.

The input is the 5-bit address produced by the ECAL card, with 8 bits of BCID to select the proper event. The output is 8 bits of data, plus the BCID for cross-checking. This is sent to the ECAL Validation Card. As already mentioned, there are two independent inputs and two outputs, each corresponding to one half of the card.

The card computes also the SPD multiplicity, by counting how many of the 64 bits have fired. This multiplicity is coded on 7 bits and is sent to the SPD Multiplicity card using the backplane lines that are used in the ECAL crate to connect to the Validation Card.

As the PreShower crates will be in the same racks as the ECAL crates, the cable length between ECAL and PreShower cards will be around 2 m.




Figure 2.6: Photograph of the prototype of the trigger part of the Preshower FE card.

2.5 Validation Card

The Validation Card [22] has two largely disconnected functions. First it handles the candidates from the 8 ECAL cards in the half-crate, doing a “validation” with the PreShower and SPD information, in order to produce electron, photon and π^0 candidates. From the 8 ECAL cards it is connected to, only one candidate of each type is selected, the one with the highest E_T . The second part handles the HCAL candidates, it adds to its energy the energy released at the same location in ECAL. Up to 4 HCAL candidates are connected to the Validation Card, and there is one output, with updated energy, for each input.

2.5.1 ECAL candidates

The 8-bit PreShower information is converted to a photon flag (PreShower and not SPD) and an electron flag (PreShower and

SPD). A Look Up Table (LUT) is used, with 3 bits output for each flag and triple voting, to be insensitive to SEU. Using a LUT allows the requirements to be modified, in particular the way the SPD is used, and a possible veto if too many of the 4 cells have fired. The ECAL inputs are delayed by about 5 cycles to wait for the PreShower information to arrive. Then the photons and the electrons candidates are sent to a selector 'highest in eight' similar to the one on the ECAL/HCAL FE board. The result is an 8-bit E_T candidate with a 8-bit address, the 3 new address bits keep track of which input was selected.

The local π^0 selection is quite similar, see [23]. For the so called "local π^0 ", where the two photons are expected on the same FE card, one uses the Sum E_T of the FE card as measure of the π^0 E_T . The highest in eight is selected. A similar validation by the PreShower and the SPD is foreseen.

The "global π^0 " candidate, when the two photons are on neighbouring cards, is obtained by adding the E_T candidates of two neighbouring cards. This is a simple add-and-saturate on 8-bit, followed by a selection "highest in eight". The address is somewhat arbitrary, it will be the address of the candidate of the first card.

These four outputs of the Validation Card are obviously quite similar. There is some flexibility in the validation by the PreShower and SPD, thanks to the use of a LUT to define which combinations are valid.

2.5.2 HCAL candidates

The motivation here is to add to the HCAL E_T the ECAL E_T in front, in order to improve the hadron energy estimate. Instead of bringing the ECAL information to the HCAL candidates, the HCAL candidates are sent to the ECAL crate. This reduces the number of connections between the two detectors by a factor 2.5, at the price of

some duplicate candidates [22]. As a side effect, the selection of the best version of the duplicated candidates has to be done in the Selection Crate.

The first processing is a time alignment, to handle the same event in ECAL and HCAL. Then the processing is in three steps:

- For each ECAL card, a single HCAL card can match. This is not the same pattern for each Validation Card, and is therefore performed by a programmable multiplexer.
- The ECAL and HCAL address are compared using a LUT, 5+5 bits input, 3 bits output for triple voting and SEU immunity. This indicates if the ECAL candidate is "in front" of the HCAL candidate. Eight such LUT are needed, one per ECAL card.
- As several ECAL candidates can be in front of an HCAL candidate, one selects for each HCAL card the matched ECAL candidate with the highest E_T and then adds its E_T to the HCAL E_T to obtain an updated HCAL candidate, which is sent to the selection crate.

This section of the Validation Card is shown in Figure 2.7. The internal memory of the PGA is used as a LUT.




Figure 2.7: HCAL validation logic.

All outputs of the Validation Card are optical links (32 bits at 40 MHz) towards the Selection Crate. The information on each link is similar: 8-bit E_T of the candidate, 8-bit address and 8-bits for BCID.

In the remaining 8-bits we intend to send a “cable number” field, allowing cabling checks.

2.6 SPD Multiplicity

In the PreShower crates, a card is located in the same slot as the Validation Card in the ECAL crate. This card receives via the backplane 8 SPD multiplicity values computed by 8 PreShower FE cards. It adds these 8 numbers and outputs the sum on an optical fiber, in a similar format as the ECAL Validation Card, using the two 8-bits address and E_T field described previously to transport the 10-bit multiplicity value. This will allow the computing the total SPD multiplicity in the Selection Crate.

2.7 Backplane and links

There is a large data flow between various boards, all at a frequency of 40 MHz. The problem has been simplified as much as possible by using a dedicated backplane [24] to implement most of the links. As shown in Figure 2.5, 9 of the 13 links between FE boards are via the backplane. The link between the FE board and the Validation Card is also via the backplane. They use multiplexed LVDS signals, where 4 pairs allow the transmission of 21 bits at 40 MHz. The same backplane is used for PreShower, ECAL and HCAL crates, the cost of any unused connection being overwhelmed by the simplification in debugging and maintenance.

The first part of the backplane, containing the power lines and the distribution of the timing and control signals, is shown on Figure 2.8.

Links between crates are implemented with multiplexed LVDS signals. Using Cat-5 cables, safe transmission is possible for lengths up to 20 m [24]. Most of the connections are between crates in




Figure 2.8: Photograph of the power backplane.

the same rack, either inside ECAL, inside PreShower or between ECAL and PreShower. Longer connections exist between HCAL and ECAL. The crates are on two platforms on top of the calorimeters, which move independently when the detector is opened. The cable length should allow for opening without decabling, 10 m should be enough, which is safe for the quality of the link.

2.8 Selection Crate

As can be seen in Figure 2.3, the Selection Crate [25] handles a lot of inputs, 4 times 28 optical links for electromagnetic candidates, and 80 links for HCAL. It is made in two parts. One handles the electromagnetic candidates, essentially selecting the one with the highest E_T for each type, and the second part handles the HCAL candidates, in a slightly more complex way. One should note that the Selection Crate is in the barracks, and hence is not submitted to radiation or SEU problems, which allows the use of FPGAs.

2.8.1 Electromagnetic Candidates

Upon reception, the processing (after time alignment) is to select the highest E_T of the 28 inputs, by successive binary comparison. The address of the final candidate, 8-bit received and 5-bits from this selection, is converted to the official calorimeter cell identifier on 14 bits, using a LUT. The resulting candidate, 8-bit E_T and 14-bit address, plus 8-bit BCID and a 2-bit status, is

sent to the L0 Decision Unit (L0DU). The 4 types of candidates (electron, photon, local π^0 and global π^0) are handled exactly the same way.

2.8.2 SPD multiplicity

The functionality is similar: the 16 inputs are time aligned, then the 10-bit numbers are added without saturation, by a cascade of pair addition, and the result on 13 bits is sent to the L0 Decision Unit. The same hardware board can be used, with a different code for the processing FPGA. There is no address to send, and the 14-bit address field is used to send the result to the L0DU.

2.8.3 HCAL

The processing is similar, with two extra steps to eliminate duplicates and to obtain also the sum over the whole calorimeter. After time alignment, the duplicates are removed: 30 HCAL cards have their candidates sent to two Validation Cards, and thus to the Selection Crate. For each pair of inputs coming from the same HCAL card, only the one with the highest E_T is kept. Then, the HCAL card with the highest E_T is selected as in the ECAL case.

The sum of the 50 cards is performed, without saturating the result. This sum will be used to detect a minimal activity in the detector, with a threshold at a few GeV. It may also be used to detect dirty events, produced by piled-up interactions, and hence saturation at 5 GeV E_T is not allowed.

As 80 optical links cannot be received on a single board, the HCAL processing is performed on 3 boards, receiving respectively 28, 28 and 24 links. A simple connection allows one of the boards to perform the final selection for the highest E_T and total sum, based on the 3 intermediate results.

The output of the HCAL selection is then the highest E_T HCAL candidate, with

the same cell identifier processing and same final format as for the ECAL candidates, and the Total E_T in HCAL. As there is no address in this case, the 14-bit address field is used to send the value, as for the SPD multiplicity.

2.8.4 Implementation

Despite the diverse functionalities, the whole Selection Crate can be implemented with a single type of board, where some small part (inter-card connection, second output) is unused in the ECAL case. The boards will be equipped with three parallel 12-channel optical receivers connected to 28 low power consumption deserializers TLK2501 from Texas Instruments. After deserialization, the 28 16-bit words are demultiplexed 2:1 to 32 bits and synchronized to the TTC clock by 28 small FPGAs.

The synchronisation mechanism has been already successfully tested. It simply requires writing the data into a FIFO with the deserializer clock while reading it with the TTC clock.

The processing itself is reasonably simple, but requires a large number of connections. A single FPGA¹ with 812 I/O pins can do the job.

The 28 inputs of each board are saved and transmitted to the DAQ after the Level-0 decision, to enable a detailed monitoring of the correct behaviour of the system. Like for most sub-systems in LHCb the TELL1-board [16] is used for this purpose. A simple zero suppression, removing candidates with exactly zero E_T , gives an acceptable event size of around 500 bytes for final readout. The Selection Crate information is also made available to the Level-1 processors, as explained in Chapter 6, with a threshold at 1 GeV to reduce the average data size to 70 bytes.

A prototype of the processing part has been built and tested in 2002, and is shown

¹For example the FPGA XILINX XC2VP50-5FF 1148C

in Figure 2.9.




Figure 2.9: Photograph of the prototype of the Selection Board.

2.9 Latency

The latency can be analysed in terms of internal processing time, transport time and delays for synchronisation of the inputs.

- FE boards: Seven cycles. On the ECAL and HCAL FE boards, the processing is 1 cycle for converting the ADC to E_T and 3 cycles for the computation of the 2×2 sums and the selection of the highest. Time alignment of the inputs will require another 3 cycles.
- Validation Card: Eight cycles: The processing in the PreShower FE card adds to the latency: the ECAL candidate has to be received (2 cycles),

the answer extracted (1 cycle) and then transmitted (2 cycles) to the Validation Card. Five cycles are thus required for these operations. The ECAL input to the Validation Card have to wait during that time. Processing in the Validation Card is quite simple, and will take 3 cycles. The HCAL processing will take the same time, and the slower transmission (due to longer cables between HCAL and ECAL crates) is smaller than the latency due to the wait for the PreShower.

- Processing time in the Selection Crate takes 9 cycles for ECAL candidates, and 14 for HCAL candidates. Two cycles are requested to de-serialize and synchronize the data fluxes to the TTC clock. Data processing on the board takes 5 clock cycles. It takes two cycles to transfer data to the L0DU or to send them to the hadron master. Three more cycles have to be added to the previous 9 for the final hadron selection, and another two cycles to the hadron data transfer to the L0DU.

The total latency, not counting the optical transmission from the calorimeter platform to the barracks, is then below 30 cycles, or 750 ns, well within the budget, as discussed in Chapter 5.

2.10 Debugging and Monitoring

To monitor the correct behaviour of the system, the inputs are logged with the data: The 8-bits E_T of each ECAL and HCAL cell, and the PreShower and SPD bits of each cell, are read out. As mentioned earlier, the inputs of the Selection Crate are also logged with the event, allowing checking that they correspond to what is expected from the individual cell inputs. The result of the Selection Crate is logged by the

L0 Decision Unit, which permits to monitor the Selection Crate. Local tests of the FE cards and of the Validation Card are foreseen, with inputs from a memory writable by ECS and results logged in a FIFO readable by ECS. This will allow the debugging of the system and in-situ checks outside data taking periods.

Chapter 3 Level-0 Muon Trigger

The muon system has been designed to look for muons with a high transverse momentum: a typical signature of a b-hadron decay.

An overview of the muon system is given first followed by the description of the L0 muon trigger implementation, its performance as a function of various running conditions and its technical design.

3.1 Overview of the muon system

The muon detector [6] consists of five muon stations interleaved with muon filters (Figure 1.1). The filter is comprised of the electromagnetic and hadronic calorimeters and three iron absorbers. Stations M2-M3 are devoted to the muon track finding while stations M4-M5 confirm the muon identification. The first station M1 is placed in front of the calorimeter and plays an important role for the transverse-momentum measurement of the muon track. Station M1 improves the transverse momentum resolution by about 30%.

Each station has two detector layers with independent readout. A detector layer contains two gaps in station M2-M5. To achieve the high detection efficiency of 99% per station and to ensure redundancy, the signal of corresponding physical channels in the two gaps and two layers are logically OR-ed on the chamber to form a logical channel. The total number of physical channels in the system is about 120,000 while the number of logical channels is

25,920.

Each station is subdivided into four regions with different logical-pad dimensions, as shown in Figure 3.1. Region and pad sizes scale by a factor two from one region to the next. The logical layout in the five muon stations is projective in y to the interaction point. It is also projective in x when the bending in the horizontal plane introduced by the magnetic field is ignored.

The logical pad dimensions are summarized in Table 3.1. Compared to M1 the pad size along the x axis is twice smaller for M2-M3 and twice coarser for M4-M5.

Pads are obtained by the crossing of horizontal and vertical strips wherever possible. Strips are employed in stations M2-M5 while station M1 and region 1 (R1) of stations M4-M5 are equipped with pads.

Strips allow a reduction in the number of logical channels to be transferred to the muon trigger. The processor receives 25,920 bits every 25 ns forming 55,296 logical pads by crossing strips.

Each region is subdivided into *sectors* as shown in Figure 3.1. They are defined by the size of the horizontal and vertical strips and match the dimension of underlying chambers.

The L0 muon trigger looks for muon tracks with a large transverse momentum, p_T . The track finding is performed on the logical pad layout. It searches for hits defining a straight line through the five muon stations and pointing towards the interaction point (Figure 3.2). The position of a track in the first two stations allows the determination of its p_T .




Figure 3.1: Front view of one quadrant of muon station 2, showing the dimensions of the regions. Inside each region a sector is shown. It is defined by the size of the horizontal and vertical strips. The intersection of the horizontal and vertical strips, corresponding to the logical channels, are logical pads. The region and channel dimensions scale by a factor two from one region to the next.

To simplify the processing and to hide the complex layout of stations, we subdivide the muon detector into 192 *towers* pointing to the interaction point as shown in Figure 3.3. A tower contains logical pads with the same layout: 48 pads from M1, 2×96 pads from M2 and M3, 2×24 pads from M4 and M5. Therefore the same algorithm can be executed in each tower, the key element of the trigger processor. Each tower is connected to a Processing Unit (PU).

All logical channels belonging to a tower are sent to a PU using six high speed optical

links. The intersection between a tower and a station maps a sector. The corresponding logical channels are transported on a dedicated optical link to ease the connectivity between the muon detector and the trigger and the data distribution within a processor.

The data flow, however, is more complex for stations M2-M3 region R1 and R2. In region R1, a sector is shared by two towers while in region R2, a tower maps to two sectors (Figure 3.1 and Figure 3.3). The first case requires additional exchange of logical channels between PUs while the second one

Table 3.1: The logical pad size in the four regions of each station projected to M1, and the number of optical links per tower and their content in term of logical channels.

Station	Region	Pad size at M1 [cm ²]	links per tower	logical channels/link			
				pads	H-strips	V-strips	Total
M1	R1	1×2.5	2	24	–	–	24
	R2	2×5	2	24	–	–	24
	R3	4×10	2	24	–	–	24
	R4	8×20	2	24	–	–	24
M2 or M3	R1	0.5×2.5	1	–	16	12	28
	R2	1×5	2	–	4	12	16
	R3	2×10	1	–	4	24	28
	R4	4×20	1	–	4	24	28
M4 or M5	R1	2×2.5	1	24	–	–	24
	R2	4×5	1	–	8	6	14
	R3	8×10	1	–	4	6	10
	R4	16×20	1	–	4	6	10




Figure 3.2: Track finding by the muon trigger. For each logical-pad hit in M3, hits are sought in M2, M4 and M5, in a field of interest (highlighted) around a line projecting to the interaction region. When hits are found in the four stations, an extrapolation to M1 is made from the hits in M2 and M3, and the M1 hit closest the extrapolation point is selected. The track direction indicated by the hits in M1 and M2 is used in the p_T measurement for the trigger, assuming a particle from the interaction point and a single kick from the magnet. In the example shown, μ^+ and μ^- cross the same pad in M3.

requires eight optical links instead of six, as shown in Table 3.1. A unique *processing board* containing four PUs deals with all cases by programming differently the FPGAs and by grouping two interconnected PUs in region R2.

The L0 muon trigger is implemented with the four quadrants of the muon system treated independently.

3.2 Trigger implementation

The L0 muon trigger algorithm and its implementation are described in detail in

LHCb notes [26, 27].

The logical channels are transported from the Front-End electronics to the muon trigger through a total of 148 high speed optical ribbons of 12 fibres each.

Track finding in each region of a quadrant is performed by 12 PUs, arranged on processing boards in groups of four for regions R1, R3 and R4, and in pairs for region R2.

A PU collects data from the five muon stations for pads and strips forming a tower, and also receives information from neigh-




Figure 3.3: A quadrant of the muon system showing the tower layout. Thick lines delimit the fraction of the system analyzed by a processing board. In this view the interaction point is shifted to ∞ .

bouring towers, although they are in another region, to avoid inefficiency on boundaries. Logical channels are merged when they are transferred from region R_i to R_{i+1} . In the opposite direction, logical channels are transported as is and replicated in four channels to match the granularity of the receiving PU. Therefore all data collected in a tower have the same granularity.

Track finding in a PU starts from the 96 logical pads defined by the intersections of horizontal and vertical strips representing the unit's input from station M3. The track search is performed in parallel for all pads.

For each logical-pad hit in M3 (track seed), the straight line passing through the hit and the interaction point is extrapolated to M2, M4 and M5. Hits are looked for in these stations in search windows, termed Fields Of Interest (FOI), approximately centered on the straight-line extrapolation. FOIs are open along the x -axis for all stations and along the y -axis only for stations M4 and M5. The size of the FOI depends on the station considered, the level

of background, and the minimum-bias retention required. When at least one hit is found inside the FOI for each of the stations M2, M4 and M5, a muon track is flagged and the pad hit in M2 closest to the extrapolation from M3 is selected for subsequent use.

The track position in station M1 is determined by making a straight-line extrapolation from M3 and M2, and identifying in the M1 FOI the pad hit closest to the extrapolation point.

Since the logical layout is projective, there is a one-to-one mapping from pads in M3 to pads in M2, M4 and M5. There is also a one-to-one mapping from pairs of pads in M2 and M3 to pads in M1. This allows the track-finding algorithm to be implemented using only logical operations.

Once track finding is completed, an evaluation of p_T is performed for muon tracks. The p_T is determined from the track hits in M1 and M2, using look-up tables. The number of muon tracks per PU is limited to two. When more candidates are found, they are discarded and the PU gives an overflow.

The two muon tracks of highest p_T are selected first for each processor board, and then for each quadrant of the muon system. The information for up to eight selected tracks is transmitted to the L0 Decision Unit.

3.3 Trigger performance

The L0 muon trigger is designed for a minimum-bias output rate of around 200 kHz¹. This is obtained by optimizing the parameters of the algorithm given by the horizontal and vertical dimensions of the FOI and by the cut on p_T . Decreasing the dimension of the FOI and increasing the cut on p_T reduces the output rate. The size of the largest FOIs is an important

¹About 2/3 of the output rate is devoted to a single muon trigger and 1/3 to a di-muon trigger.

parameter for the processor since they define the number of data exchanged between PUs. This number determines the dimension of busses connecting PUs. The largest FOI² are given in Table 3.2.

Table 3.2: The maximum size of the FOI along the x and y coordinates. It is expressed in terms of pads with respect to the pad lying on the straight line passing through the hit in M3 and the interaction point. A FOI of ± 3 corresponds to a total width of 7 pads.

	M1	M2	M4	M5
x	± 3	± 5	± 3	± 3
y	-	-	± 1	± 1

Figure 3.4 shows the transverse momentum determined for L0 muon candidates found in minimum-bias and $B_s^0 \rightarrow J/\psi\phi$ samples when the FOI are optimized for an output rate of 125 kHz (single muon trigger). The corresponding trigger efficiency is shown in the bottom plot as a function of a cut on p_T . The origin of the muon candidates in the accepted minimum-bias events is given in Table 3.3. They mainly come from pion and kaon decays in flight. The resolution on the transverse momentum was measured to be 20% for muons coming from a b-quark.

Table 3.3: Origin of candidate triggering minimum-bias events when the rate for the single muon trigger is fixed to 125 kHz. The table includes hadron punch-through.

	[%]
b-hadron	2.2
c-hadron	3.3
Pion	63.2
Kaon	28.5
Other particles (p, n, τ ,...)	1.1
Ghost tracks	1.7

²They were obtained by optimizing the trigger efficiency while minimizing the size of the FOI. Studies were performed as a function of the output rate and running conditions [28, 29].

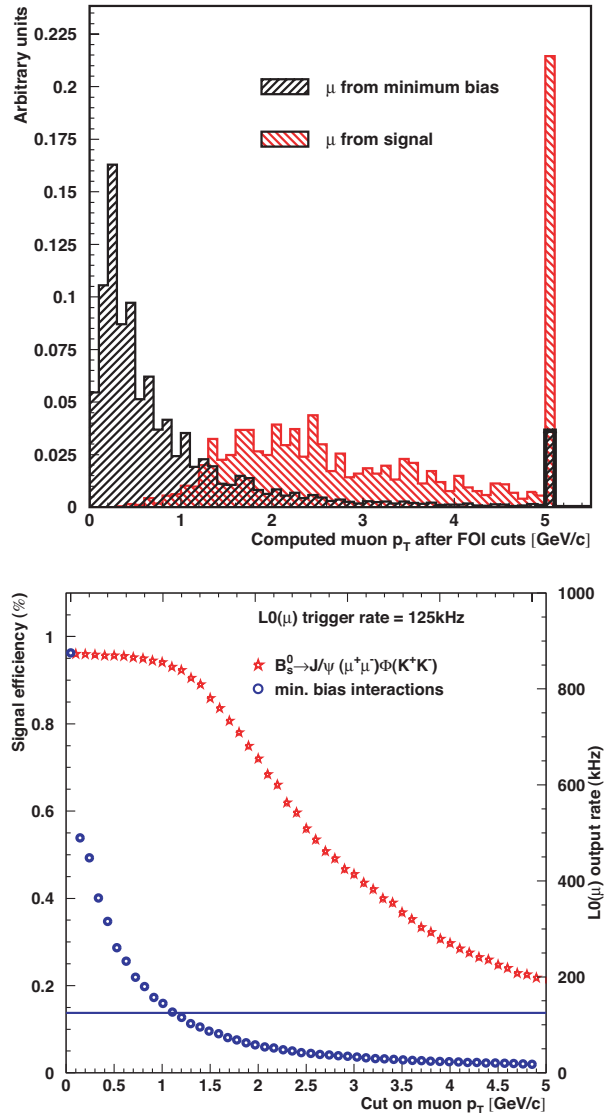


Figure 3.4: Top: reconstructed transverse momentum for minimum-bias and for $B_s^0 \rightarrow J/\psi\phi$ events. It is encoded on 8 bits and saturated to 5 GeV/c. Both samples are normalized to unity. Bottom: the trigger efficiencies for minimum-bias and for $B_s^0 \rightarrow J/\psi\phi$ events as a function of the cut on p_T . In both plots the dimensions of the FOI are optimized for an output rate of 125 kHz (single muon trigger). The $B_s^0 \rightarrow J/\psi\phi$ events are selected by offline reconstruction.

The robustness of the muon trigger implementation has been studied by varying the minimum-bias retention level, and the operating conditions defined by the level of low-energy background in the chambers, the level of beam halo muons, as well as

chamber parameters [28, 29]. The parameters of the trigger algorithm are optimized in each case. The performance on useful events selected by the reconstruction and tagging procedure are given in Chapter 7. Here, the relative loss in the efficiency when the running conditions are deteriorated is presented.

3.3.1 Low-energy background

The energy thresholds of the Geant simulation in the region behind the calorimeters are set to higher values than in the rest of the detector to save CPU time spent in tracking inside the iron filters. As a consequence the low energy component of the muon chambers hits rate [6] in stations M2-M5 is strongly suppressed. To restore the correct rate, background hits are added during digitization. They are extracted from a parametrization obtained with a different version of the simulation program which contains lower energy thresholds and a more detailed geometry of the detector and the beam optics. The low energy background is constituted by low energy particles, mainly electrons and charged hadrons and, for the large arrival time, to thermal neutrons. Since the simulation of these processes is affected by large uncertainties, in the robustness test, conservative safety factors from 2 to 5 have been applied to the total number of hits according to the relative importance of this component in the five muon stations. The loss induced by the level of low-energy background depends on the Level-0 output rate. It varies between 2% (300 kHz) and 8% (100 kHz).

3.3.2 Beam halo muons

The charged-particle flux associated with the beam halo in the accelerator tunnel contains muons of a rather wide energy spectrum with its largest flux at small radii [6]. In particular those halo muons traversing

the detector in the same direction as particles from the interaction point can cause a muon trigger. The average number of beam halo muons depends strongly on the level of residual gas in the beam pipe [30]. We define the *nominal condition* as the second year after 10 days of running and the *worst condition* by applying a safety factor of two on the expected level of residual gas and three on the beam current. In nominal condition, the average number of beam halo muons is equal to 0.015 per bunch for particle travelling from the entrance of the cavern toward the muon detector and 0.026 for particle coming from the other side. Studies, performed on minimum-bias samples with superimposed beam-halo particles [28, 29], show that the beam halo does not affect the trigger performance in nominal conditions. Increasing the level of residual gas in the beam pipe and the beam current, however, decreases the trigger efficiency by less than 8%. The magnitude of these losses is similar to the one induced by the maximum level of low-energy background and add linearly with it. These studies also show that the muon trigger is rather insensitive to the beam halo coming from the other side of the interaction point since particles are not in time in that case.

3.3.3 Hardware parameters

The chamber response depends on several parameters: cluster size, single gap efficiency and electronic noise. The most sensitive one is the cluster size [28]. An overall increase by 30% decreases the trigger efficiency by less than 5%.

The implementation of the muon trigger algorithm limits the number of muon candidates per PU to two, applies data compression algorithms when data are transferred between towers belonging to different regions, and encodes the p_T on 8 bits. These simplifications have no significant effect on the trigger performance.

3.4 Technical Design

The muon trigger is divided into four independent parts running on the quadrants of the muon detector. They are located behind the shielding wall in an environment without radiation. A processor is a 9U crate with 15 Processing Boards, 3 Muon Selection Boards and a Controller. All these boards are interconnected through a custom backplane as shown in Figure 3.5.

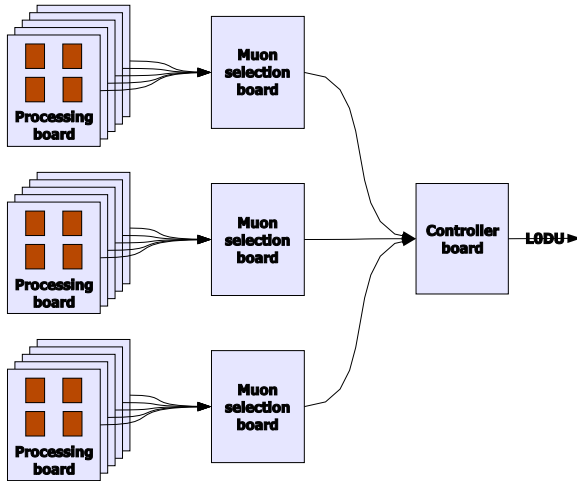


Figure 3.5: The data flow of a processor connected to a quadrant of the muon detector.

The architecture is fully synchronous, pipelined and massively parallel. The processing frequency is 40 MHz. Data exchange frequency between boards and between PUs is 80 MHz.

In this section, we present briefly the technical design of these components. A detailed description can be found in [27, 31].

3.4.1 ODE Trigger interface

The main task of the muon front-end electronics is to form the logical channels, to tag each logical signal with a bunch crossing identifier and to send time-aligned data to the trigger [6] as shown in Figure 3.6. The building of the logical channels from the physical ones is performed by 7632 Front-

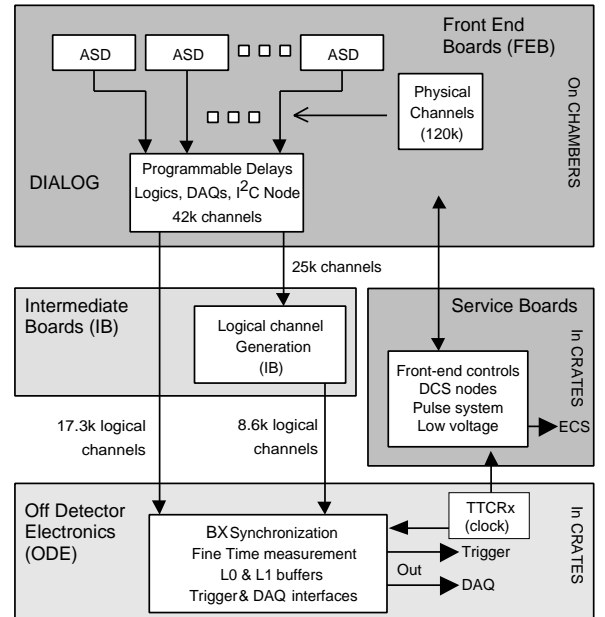


Figure 3.6: Simplified scheme of the muon front-end architecture.

End boards mounted on the detector and by 152 Intermediate Boards. The remaining tasks are handled by 148 Off Detector Boards (ODE) located on the left and right side of the muon detector.

A trigger interface located in an ODE board receives up to 192 logical channels and pushes every 25 ns twelve 32-bit words on 12 high-speed optical links grouped in one optical ribbon cable. The Gigabit Optical Link transmitter (GOL) [32], developed at CERN, encodes and serializes the 32-bit word with its clock using the 8B/10B protocol. The resulting 1.6 GHz electric signal is converted to an optical signal by a ribbon transmitter³.

The jitter of the input clock driving the GOL has to be lower than the jitter of the clock delivered by TTCrx components by a factor less than three to guarantee a bit error rate below 10^{-12} . A Filtering circuit is implemented in the interface. It will be either the *radiation hard jitter filter* ASIC from CERN, named QPLL [33],

³from Agilent™

or a discrete narrow bandwidth *phase lock loop* (PLL) controlling a *voltage crystal oscillator* [34].




Figure 3.7: The prototype of the ODE Board with its trigger interface.

We developed a prototype of a ribbon of high-speed optical links with a filter based on a narrow bandwidth PLL controlling a voltage crystal oscillator [34]. We obtained a bit error rate below 10^{-15} with the TTC clock. The effect of single event upsets has been estimated on the part of the ribbon optical link implemented in the ODE. We obtain an equivalent bit error rate of 3×10^{-11} in the radiation environment of the muon detector. However, the cross-section of the optical ribbon transmitter has still to be measured.

Figure 3.7 shows a photograph of the prototype of the ODE board with its trigger interface.

The 148 ribbon cables coming from ODE boards are connected to a passive patch

panel. It merges fibres related to a tower coming from different stations into a single output ribbon. The input cable is about 80 m long. Output ribbons are connected to processing boards.

3.4.2 Processing Board

The diagram of the processing board [31] is shown in Figure 3.8. A board contains:

- two receivers for two ribbon optical links corresponding to 24 single optical channels;
- six FPGAs: one for each PU, one for the BCSU (Best Candidates Selection Unit) and one for the L1MU (L1 Management Unit);
- eight look-up tables;
- one ECS interface based on a *credit card PC* [9];
- one interface to the custom backplane.

The hardware of the processing board is unique but the programming of the 60 PUs housed in a processor depends on their location in the system.

Six optical channels coming from a tower are connected to a PU. The corresponding FPGA receives six 16-bit words, their 80 MHz clocks and 6×2 control bits. The input data are time aligned using the bunch crossing identifier and control bits. Data are exchanged with neighbouring PUs and the muon finding algorithm is executed. Table 3.4 shows the maximum information exchange between PUs.

Each PU outputs a 38-bit word with addresses of hits in station M1, M2 and M3 for the two candidates, the bunch crossing identifier and a status.

For each candidate, the p_T is computed using a look-up table and encoded on an 8-bit word. The look-up table is implemented in a $32k \times 8$ static RAM.

The next step is performed by the BCSU which selects the two candidates with the

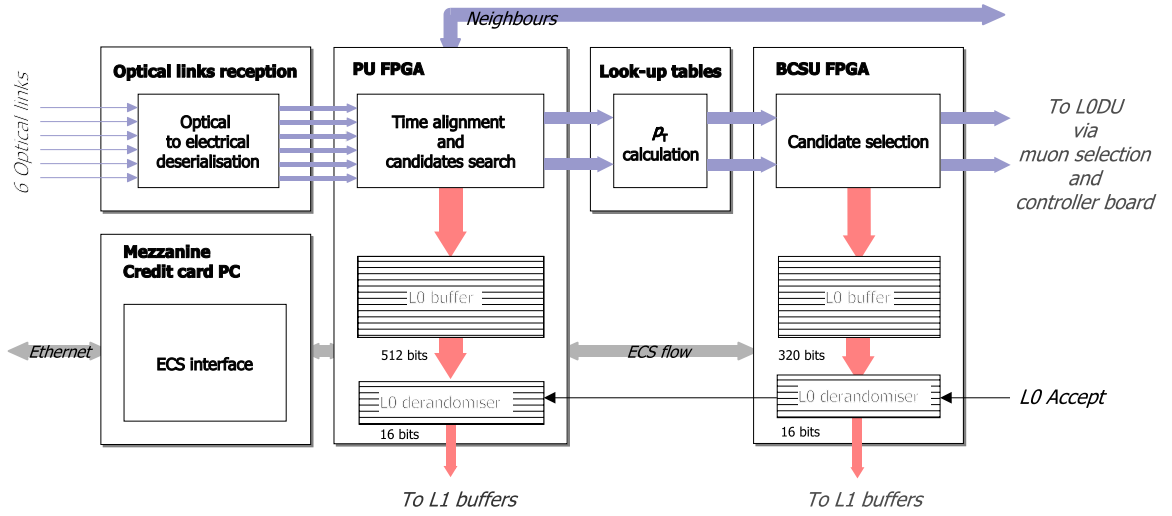


Figure 3.8: Scheme of the processing board where the data flow is only shown for one PU.

Table 3.4: The number of logical channels exchanged between PUs. Busses named “Horizontal”, “Vertical” and “Crossing” link PUs located in the same board while the bus named “Backplane” connects PUs spread over several boards.

	Top		Bottom	
	Left	Right	Left	Right
Backplane From	94	90	88	86
Backplane To	110	82	96	96
Vertical From	42	42	42	42
Vertical To	42	42	42	42
Horizontal From	82	72	72	82
Horizontal To	72	82	82	72
Crossing From	2	12	12	2
Crossing To	2	12	12	2

highest p_T among the eight proposed by the PUs. Results are stored in a 60-bit word which is sent to the muon selection board via an 80 MHz point-to-point connection. It contains the addresses of two candidates in station M1, M2 and M3, their p_T , the bunch crossing identifier and a status.

Each PU and BCSU houses a L0 buffer and its derandomizer buffer. It receives the input and output words of the components. Their width is 512 bits for a PU and 320 bits for a BCSU.

Each L0 buffer is connected to an L1 buffer through a bus, 16-bits wide, running at 40 MHz. The L1MU activates transfers between L0 and L1 buffers and between the L1 buffer and the controller board. It also houses the L1 derandomizer buffer. A L1 buffer is implemented with a synchronous dynamic RAM, 1M×16 wide. The data are transferred to the controller via a serial point-to-point 4-bit wide link.

Figure 3.9 shows a photograph of the prototype of the Processing Board. This prototype is very close to the final design except for the size of the L1 buffer, which is too small. Each FPGA exchanges data with its neighbours at 80 MHz and is connected to the ECS interface.

3.4.3 Muon selection board

The Muon Selection Board contains only one FPGA with a functionality very similar to the best-candidate selection unit. It is connected to five processing boards and receives their best candidates. The chip selects the two candidates with the highest p_T among the 10 proposed. The 60-bit output word is sent to the controller board through a 80 MHz point-to-point connection.




Figure 3.10: Scheme of the controller board.

3.4.4 Controller board

The diagram of the controller board is shown in Figure 3.10.

The controller board is connected to the three muon selection boards and receives their best candidates. An FPGA with programming very similar to the BCSU selects the two candidates with the highest p_T among the six proposed. The result is encoded into two 32-bit words, one for each candidate, containing the address of the candidate in station M1, its p_T values and a status. These two words are sent to the L0 Decision Unit. Inputs and output of this component are stored in a 320-bit wide L0 buffer.

The controller board is also linked to the

15 processing boards of the card to receive the output of their L1 derandomizer buffers. A FPGA builds the events, strips duplicated information such as bunch crossing and event identifiers and applies a zero suppression algorithm. The output is sent to the data acquisition system via two Gigabit Ethernet links.

The controller board is the interface with the TTC system [12]. It receives TTC signals through the TTC receiver chip and distributes them to processing and controller boards via dedicated links running on the backplane.




Figure 3.9: The prototype of the Processing Board where a PU is implemented in one FPGA with 600,000 gates and 652 pins.

3.4.5 Backplane

The backplane distributes power supplies (+5 V, +3.3 V, +48 V, GND), the main 40 MHz clock and service signals coming from the TTC system. Clocks are sent individually from the controller to each processing boards by point-to-point links while service signals are broadcast on a common bus at 40 MHz.

The backplane connects:

- processing boards together to exchange neighbouring information;
- a processing board to a muon selection board and muon selection boards to the controller;
- each processing board to the controller to transfer the content of L1 buffers.

All these connections rely on point-to-point links running at 80 MHz.

Analog simulations have been made to find the most appropriated impedance

matching scheme. All point-to-point links are terminated either on the processing boards or on the controller board side while bussed signals are adapted on the backplane. Table 3.5 summarizes the number of pins required to connect a board to the backplane. We implement a compact PCI connector with shield and guide lugs for centering.

Table 3.5: Number of pins required to connect a processing, muon selection and controller boards to the backplane.

	Processing board	Muon selection board	Controller
Signal	443	163	219
Ground	198	198	198
Power	106	106	106
Free	43	323	267
Total	790	790	790

3.4.6 Latency

The latency for the muon trigger is 1050 ns as shown in Table 3.6, well within specification given in Chapter 5. It starts with the first data arriving on the fastest optical link and it ends when the results are serialized on the link connected to the L0 Decision Unit.

Table 3.6: Breakdown of the latency for the muon trigger expressed in terms of LHC clock cycles.

	Clock cycle
Optical link deserialization	2
Optical link synchronization	4
Neighbouring exchange	5
Muon tracking	1
M1 pad finding	1
Candidates selection within a PU	5
p_T computation	1
p_T selection within a board	4
Final selection	17
Serialization to L0DU	2
Total	42

3.4.7 DAQ Event size

We log the input and output words of each processing element in the L0 buffers to monitor the trigger during data taking and to trace any bias which might be introduced. The target event size is about 1 kBytes after zero suppression.

3.4.8 Debugging and monitoring tools

The strategy to debug a processing board or a processor relies on the ECS interface, L0 buffers and a simulation of the hardware. The ECS interface can fill specific buffers located in the PUs with test patterns. They emulate data transport by the optical links for 16 consecutive events. The test buffers inject data in place of the buffers receiving the output of the optical transceiver stage. The trigger then runs on 16 consecutive events and stops. Input values provided by the test buffers, neighbouring data from adjacent PUs and p_T computation are logged to L0 buffers and systematically transferred to L0 derandomizer buffers. The ECS interface can read the content of L0 derandomizer buffers and we can compare the results with the expected values provided by a simulation of PUs and BCSUs.

Chapter 4 Level-0 Pile-Up System

Upstream of the VELO system [7], a set of two planes of silicon strip detectors is used to determine the number of primary interactions within one bunch crossing. The silicon detectors of this Pile-Up system are equipped with special fast readout electronics to allow their data to be made available at Level-0. LHCb aims to run with an average luminosity of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$, however to achieve this average all sub-systems are able to cope with luminosities up to $5 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$. Figure 4.1 shows the rate of crossings with their expected number of visible interactions¹ in the luminosity range for which the spectrometer has been optimised.




Figure 4.1: Rate of crossings with their number of pp-interactions assuming $\sigma_{\text{visible}} = 63 \text{ mb}$, as a function of luminosity.

Crossings with multiple interactions trigger at Level-0 and subsequent trigger levels more based on combinatorics rather than on genuine b-decay candidates, and in

¹Chapter 7 describes the physics simulation, and defines visible interactions, which are expected to have a cross-section $\sigma_{\text{visible}} = 63 \text{ mb}$.

addition tend to occupy a disproportional large share of the event building bandwidth and the available processing power. Removing these crossings can even give a gain in the number of signal events collected, since other trigger cuts can be relaxed to saturate the allowed bandwidth. Note that the Pile-Up system detects only tracks in the backward direction, and hence it cannot mistake B-decays in the acceptance of LHCb for pile-up interactions.

The Pile-Up system also provides a relative measurement of the luminosity, since in the luminosity range of LHCb the rate of crossings with zero, one and multiple interactions allows its determination using Poisson statistics.

4.1 Concept

The Pile-Up system [35] - [37] consists of two planes (*A* and *B*) perpendicular to the beam-line and located upstream of the VELO, as shown in Figure 4.2. Each plane consists of two overlapping VELO R-sensors [38], which have strips at constant radii, and each strip covers 45° . In both




Figure 4.2: Top view of the layout of VELO planes and the Pile-Up detector planes *A* and *B* at $-22.0/23.5$ and $-30.0/31.5 \text{ cm}$ respectively. The interaction region containing 95% of the luminosity is expected to be 16 cm wide along the beam-line, and is indicated as well.

planes the radii of track hits, r_a and r_b ,

are recorded. The hits belonging to tracks from the same origin have the simple relation $k = r_b/r_a$, giving:

$$z_v = \frac{kz_a - z_b}{k - 1} \quad (4.1)$$

where z_b , z_a are the detector positions and z_v is the position of the track origin on the beam axis, *i.e.* the vertex. The equation is exact for tracks originating from the beam-line. All hits in the same octant of both planes are combined according to equation 4.1 and the resulting values of z_v are entered into an appropriately binned histogram, in which a peak search is performed, as shown in Figure 4.3. The resolu-

tion of z_v is limited to around 3 mm by multiple scattering and the hit resolution of the radial measurements. To limit the number of channels which have to be processed, four neighbouring strips of the sensors are OR-ed on the FE-chip, and hence the latter effect dominates. All hits contributing to the highest peak in this histogram are masked, after which a second peak is searched for. The height of this second peak is a measure of the number of tracks coming from a second vertex, and a cut is applied on this number to select crossings.

4.2 Simulation Results

The effect on the number of signal events is shown in Figure 4.4, where for different signal channels the combined Level-0 and Level-1 efficiency is plotted as a function of a cut on the number of track candidates from the second largest multiplicity vertex as detected by the Pile-Up at a luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. For every Pile-Up cut, the thresholds for the other trigger variables are modified to fill the allowed bandwidth of the two trigger levels. The Level-0 algorithm described in Chapter 7 is found to give better results for channels with muons if the Pile-Up cut is not applied if the sum of the transverse momenta of the two largest p_T muons is above its threshold, and these channels show a different sensitivity compared to hadronic decays. Figure 4.4 also shows that the system reduces the average event size, which allows a smaller event building network in Level-1, and reduces the necessary processing time in subsequent trigger levels. The gain in event yield due to pile-up detection increases with luminosity, which is shown in Figure 4.5, where the expected yield for offline reconstructed and triggered $B_s \rightarrow D_s K$ events is given for the nominal Pile-Up cut of 3, and without the Pile-Up System.

The number of hits detected in the Pile-

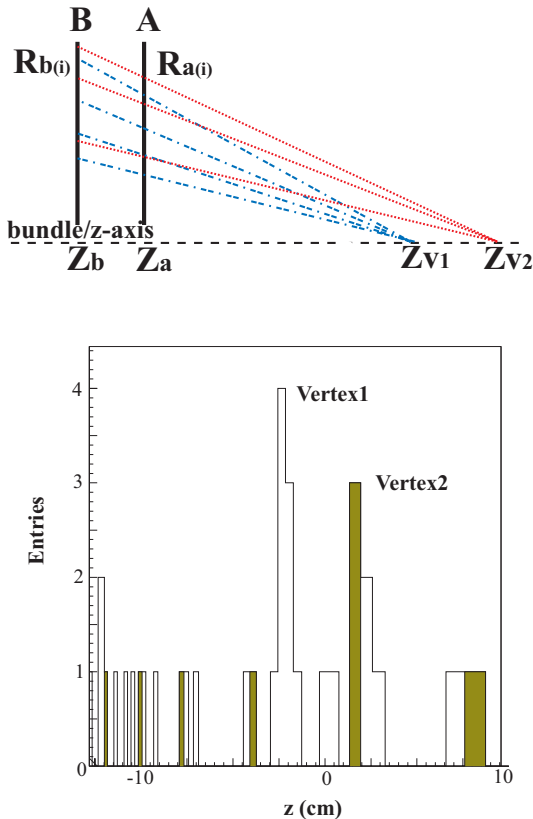


Figure 4.3: Basic principle of detecting vertexes in a crossing. The readout hits of plane A and B are combined in a coincidence matrix. All combinations are projected onto a z_v -histogram. The peaks indicated correspond to the two interaction vertexes in this particular Monte-Carlo event. After the first vertex finding the hits corresponding to the two highest bins are masked, resulting in the hatched histogram.




Figure 4.4: Combined L0xL1 efficiency for three physics channels as function of a cut on the number of tracks detected in the second vertex. Also shown is the efficiency without the Pile-Up system (No Cut). The bottom plot shows the corresponding number of VELO+TT clusters and Level-1 execution time, normalized to their “No Cut” values, in minimum-bias events after Level-0. The nominal Pile-Up Cut is indicated by the dashed line.




Figure 4.5: Expected yield for $B_s \rightarrow D_s K$ per year after Level-1 as a function of the luminosity with and without the Pile-Up System. The cut on the number of tracks in the second vertex is 2 and 3 respectively.

Up system gives a measure of the charged track multiplicity in the event close to the primary vertex, and is used in combination with the SPD-multiplicity in addition to the number of interactions to flag ‘complicated’ events [39].

4.3 Technical Design

The Pile-Up system is an integral part of the VELO [7] as far as the sensors, their mechanical mounting and the readout of the analog pipeline after Level-0 are concerned. Also the control system and power supplies are identical to the VELO. In addition, however, the Pile-Up system uses the signals of the integrated comparators of the Beetle [40] chips on its four hybrids. The output of the four neighbouring comparators is OR-ed, resulting in 256 LVDS links running at 80 Mbit/s per hybrid, which send the L0-signals via the Repeater Station on the vertex tank to an Optical Transmission Station. From where the data of the 1024 signal pairs will be transferred via optical links to the trigger logic which is located in the radiation-free electronics barracks. Figure 4.6 shows an overview of the system.

The radiation levels at the hybrid, Repeater Station and Optical Transmission Station are given in Table 4.1. The sensors are located in a radiation area necessitating their replacement every few years, like for the VELO. The hybrids, although far less sensitive, have then to be replaced as well. For the Pile-Up System no active elements will be placed at the Repeater Station. The radiation level at the Optical Station is tolerable for using commercial optical links at that location .

The design of the optical links for the muon system (Chapter 3) will be followed. Timing information is lost when using serial optical transmission links. Therefore a time stamp consisting of part of the BCID

Table 4.1: Yearly radiation levels at several locations of the Pile-Up System electronics.

Location	Dose
Hybrid	2 kGy
Repeater Station	200 Gy
Optical Station	<1 Gy




Figure 4.6: Overview of the Pile-Up System.

is included in the data. Hence a TTCrx chip is included in the Optical Transmission Station, the receiving side is passive. The time stamp data occupies part of the optical links. The number of needed connections is then 2 ribbons per hybrid, giving 96 optical fibers in total.

4.3.1 Beetle Chip

The Beetle [40] chip is used in LHCb in the VELO, TT and IT stations, in the RICH and in the Pile-Up. It is designed in a com-




Figure 4.7: Prototype VELO hybrid with a 182° Si-detector mounted. The detector strips are circular arcs, the pitch increases with the distance to the beam axis.

mercial $0.25 \mu\text{m}$ CMOS technology and has a die size of $6.1 \times 5.4 \text{ mm}^2$. In case of the VELO and Pile-Up systems, the readout chip will be positioned only 5 cm from the LHC beam, and the Beetle has been designed to be radiation hard, and to avoid risk of Single Event Latch-up. Radiation hardness of the Beetle has been demonstrated for up to 300 kGy [41].

Beetle 1.1 was used to check the full analog operation [42], including a 16-chip hybrid, with a prototype R-sensor, as shown in Figure 4.7.

In the Beetle chip four detector input

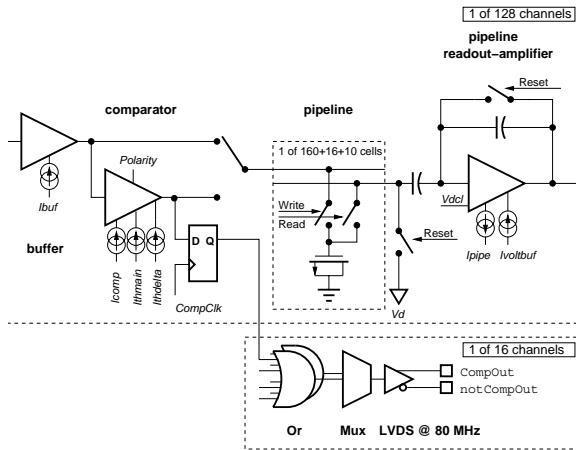


Figure 4.8: Comparator, pipeline and output part of the Beetle chip.

channels are combined, at the cost of a decrease in vertex accuracy, by a Logic-OR at the comparator stage for providing fast signals that immediately can be used in the Level-0 trigger system (Figure 4.8). Two groups are multiplexed on one output line, giving 16 LVDS outputs at 80 Mbit/s per chip. There are two output modes: a) tracking mode with a time over threshold pulse or b) pulse mode for one clock period. The latter where one signal will not produce spillover in the next bunch crossing, is used. The comparator part has been tested for single chips [43]. Typical threshold curves for the Beetle 1.1 are shown in Figure 4.9. The sigma on the threshold is about 0.07 MIP. Despite a satisfactory performance in noise and efficiency for single channels, not all channels could be operated in discriminator mode due to a too large offset spread combined with a too small range in the DACs to set the individual thresholds. These deficiencies have been rectified in the Beetle 1.3 design.

Since the Pile-Up sensors and hybrids are not inside the acceptance of LHCb, thicker sensors than the VELO, i.e. $300 \mu\text{m}$, will be used to have more signal, and the hybrid design is not limited by radiation length consideration, hence allowing




Figure 4.9: Threshold scans for three Beetle 1.1 comparator channels with different offsets.

easier suppression of common mode. The hybrid shown in Figure 4.7 has been tested in the CERN testbeam [42]. Based on this VELO design an eight-layer Pile-Up hybrid has been designed for the Beetle 1.2/1.3, now including the LVDS outputs of the discriminators.

4.3.2 Prototype Vertex Finder Board

Since the Vertex Finder Board (VFB) is the most complicated board of the processing system, it has been prototyped first as a 6U VME board (Figure 4.10). The trigger algorithm has been partitioned into code for two FPGA's. The XCV3200E is with 4 M system gates the largest FPGA of the Xilinx Virtex-E 1.8 family. With the present design the device utilisation is about 70%, close to the maximum possible. Timing analysis and tests on the prototype board show that the required 40 MHz operation is feasible. The present implementation of the trigger algorithm in the FPGAs takes 50 steps of 25 ns, close to initial predictions. Additional steps are required for data alignment and serialisation. Studied is whether




Figure 4.10: Prototype Vertex Finder Board.




Figure 4.11: Prototype Vertex Finder Board diagram.

the algorithm can be optimized further to regain some processing steps to fulfill the overall latency requirements.

A Test Board with a smaller FPGA provides the logic to supply test patterns to the Vertex Finder Board. The test patterns are loaded via VME into the FPGA, which stores these patterns in memory.

In Figure 4.11 the schematics of the pro-

tototype of a Vertex Finder Board is shown. In the left FPGA all hit patterns of the Si-detectors are first stored in a correlation matrix. Hits from tracks with the same origin have equal ratio $k = r_b/r_a$. All channel combinations are stored in the Coincidence Matrix. A z -histogram is formed by summing all entries of wedges between lines of constant ratio k in the matrix. The number of processing steps is always the same, does not depend on whether detector strips are hit or not. A linear search for the highest peak in the histogram is performed. Then the input bits related to that peak (hence having the same k value) are removed from the data stream that is passed on further to the next FPGA. There the second highest peak is searched for. All processing is pipelined, with 100 ns intervals. Results are output via LVDS lines.

4.3.3 Trigger System Architecture

Input data are routed further by the Multiplexer Boards. The processing of the vertex finding algorithm (Figure 4.12) is performed in the Vertex Finder Boards. The Multiplexer Boards distribute the events by round-robin scheduling. Each Vertex Finder Board processes one event as indicated. Processing results are demultiplexed by the Output Board. The Output Board interfaces the processor system to the central Level-0 trigger. Just one 9U/40 cm crate will be needed for the whole processor system. The crate layout and the internal connections are shown in Figure 4.13.

Multiplexer Boards

The number of input signals is 1024. Eight optical ribbons will be connected to four Multiplexer Boards. The optical to electrical transition will be directly at the Multiplexer Board level. Each Multiplexer Board is connected to all Vertex Finder Boards




Figure 4.12: Pile-Up processing plus the data multiplexing and serialising scheme.

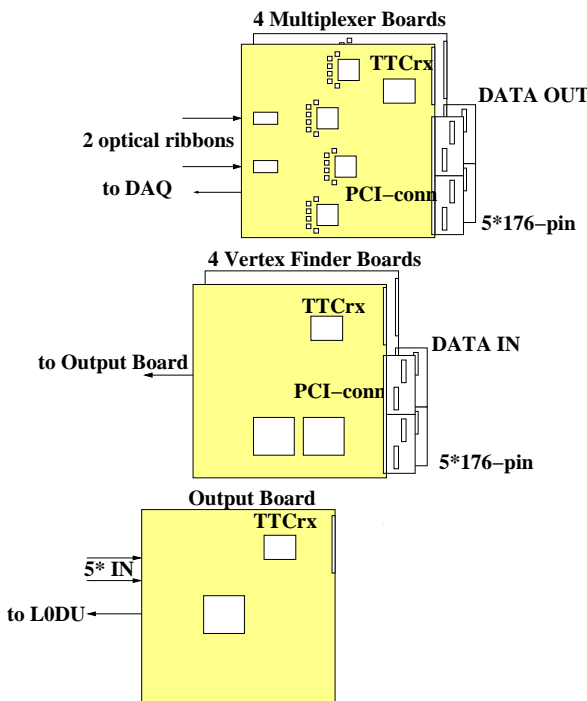


Figure 4.13: System crate layout and data connections.

via point-to-point connections running at 80 Mbit/s, using PCI-connectors/cables at the backplane. In the Multiplexer Board the data will be round-robin routed by a FPGA to the Vertex Finder Boards. The input data is also copied directly into memory (L0 buffer) for inclusion in the DAQ chain. A TELL1-board [16] will be used for that purpose. Both at the Beetle level as at the Multiplexer Board level noisy channels

can be masked.

Vertex Finder Boards

In total four Vertex Finder Boards are planned to be used, where each board handles every fourth event. Minor configuration parameters as threshold levels should easily be adaptable. Algorithms for different beam or geometrical conditions will be pre-programmed and loaded on demand. Binning in the vertex histogram and the masking width can be adapted.

Future FPGAs are expected to be even larger than the XCV3200E, providing the possibility to combine all tasks in just one FPGA. The specific elements for the luminosity processing still have to be defined in detail and require also extra FPGA resources.

Output Board

The Output Board is a simple board combining the inputs of the Vertex Finder Boards and outputting the data to the LODU. The trigger information (0, 1, 2 interactions) is histogrammed at the LODU level. These histograms will be the basis for determining the luminosity with the Pile-Up system.

Latency

The breakdown of the latency of the Pile-Up system is given in Table 4.2.

Debugging and Monitoring

The following items have to be monitored regularly:

- Noisy channels: channels that give spurious hits could flood the processing system with uncorrelated entries. Automatically these channels should be looked for and removed from the input of the processing system.

Table 4.2: Breakdown of the latency for the Pile-Up system.

	Time [ns]
Beetle	50
Copper cable	90
Optical Transmission Station	125
Optical fibre	270
Multiplexer Board	250
Algorithm	1175
Output Board	150
To L0DU	90
Total	2200

- **Optical Station:** The ECS interface can fill specific buffers located at the Optical Station with test patterns. They emulate data transport by the optical links for consecutive events.
- **Readout:** the input signals for the Pile-Up system are transferred to the DAQ. Regularly offline checks have to be performed to see whether the results of online and offline processing agree.
- **Vertex Checks:** the number of entries in the histograms and the location of the vertices should be followed closely to check the behaviour of the system and the machine background conditions.
- **Processing:** test patterns can be fed into the Vertex Finder Boards to check the overall processing of the system.

Although the requirements on alignment are not very stringent ($\Delta r < 100 \mu\text{m}$), a check on the correct position of the detectors is necessary. This check is part of the overall VELO geometry alignment.

Chapter 5 Level-0 Decision Unit

The Level-0 Decision Unit (L0DU) receives information from the Calorimeter, Muon and Pile-Up sub-triggers at 40 MHz, which arrive at different fixed times. The L0DU latency budget is 500 ns, counted from the latest arrival of the sub-system data. Table 5.1 lists the breakdown of latency budget for the Level-0 sub-systems. The computation of the decision can start

Table 5.1: Breakdown of the latency of Level-0 in ns.

	Muon	Calo	Pile-Up
TOF+Cables	975	850	1000
Processing	1200		
Subtotal	2175	2050	2200
L0DU	500		
RS+TTC→FE	800		
Total=max	3500		
Contingency	500		
Total latency	4000		

with a sub-set of information coming from a Level-0 sub-trigger, after which the sub-trigger information is time aligned. An algorithm is executed to determine the trigger decision, and a summary bank (L0Block) is constructed. The L0Block is made available to Level-1 and the HLT. The decision is sent to the Readout Supervisor [15], which has the ultimate decision about whether to accept an event or not. The Readout Supervisor is able to generate and time-in all types of self-triggers (random triggers, calibration, etc) and to control the trigger rate by taking into account the status of the different components in order to prevent buffer overflows and to enable/disable the

triggers at appropriate times during resets etc.

The L0DU performs simple arithmetic combining the signatures into one decision per crossing. It can set several thresholds per candidate, and allows the downscaling of triggers. It can also base its decision on some information of the two preceding and two subsequent crossings, and this information is also included in the L0Block. It will monitor the Level-0 performance with counters which are made available via the ECS, and allows quick interrogation of the trigger source via an explanation word included in the L0Block. Despite the available flexibility, the results presented in Chapter 7 are based on a simple algorithm, which sets thresholds on the E_T of all the candidates. If the SPD and Pile-Up multiplicities, or the number of tracks from a secondary vertex are above a given value, the event is tagged as a Pile-Up-Event and rejected. In addition, events are accepted if Σp_T^k is larger than a threshold, irrespective of the Pile-Up-Event tag.

The L0DU will be installed in the barracks. It is a custom-built board [44], implemented using a 40 MHz pipeline logic in an FPGA.

5.1 L0DU inputs

Table 5.2 summarizes the L0DU input/output ports.

Each L0 trigger processor sends the corresponding data synchronously with its own latency. The trigger processor data fits into

Table 5.2: L0DU Input/Output summary

External system	I/O	# bits
CALO	I	224@40MHz
MUON	I	256@40MHz
Pile-Up	I	64@40MHz
Reserved	I	96@40MHz
Readout Supervisor	O	16@40MHz
L1	O	704@1MHz
HLT	O	1024@40kHz
ECS	IO	-

normalized 32 bits words corresponding to a candidate. It includes a bunch identification number allowing the synchronization between the data sources.

- The Calorimeter trigger (see Chapter 2) sends seven words to the L0DU. They correspond to the highest E_T electron, photon, local π^0 [23], global π^0 and hadron trigger candidates, the sum of the total transverse energies measured in HCAL (Total E_T) and the total multiplicity of the event N_{SPD} measured with SPD detector [45].
- The Muon trigger processor (see Chapter 3) sends eight words to the L0DU corresponding to eight muon candidates, two per quadrant of the muon system.
- The Pile-Up System (see Chapter 4) sends two words, indicating the number of tracks per interaction to the L0DU.

A total of 20×32 bits @ 40 MHz is expected as input of the L0DU. Only one electrical standard on the L0DU input interface will be used in order to simplify test and maintenance.

5.2 L0DU overview

For each data source, a Partial Data Processing system performs a specific part of

the algorithm and the synchronisation between the various data sources. Then a trigger definition unit combines the information from the above systems to form a set of trigger conditions based on multi-source information (Figure 5.1).

The trigger conditions are logically ORed to obtain the L0DU decision after they have been individually downscaled if necessary.

Then a decision word (16 bits) is sent to the Readout Supervisor [15]. This word includes the decision itself (1 bit) and 12 bits for the bunch number. One additional bit is reserved for a forced trigger. The L0Block is built for each event and stored in pipeline memories waiting for the L0 accept signal. Then it is transferred to the Level-1 buffer and a subset is sent to the Level-1 trigger.

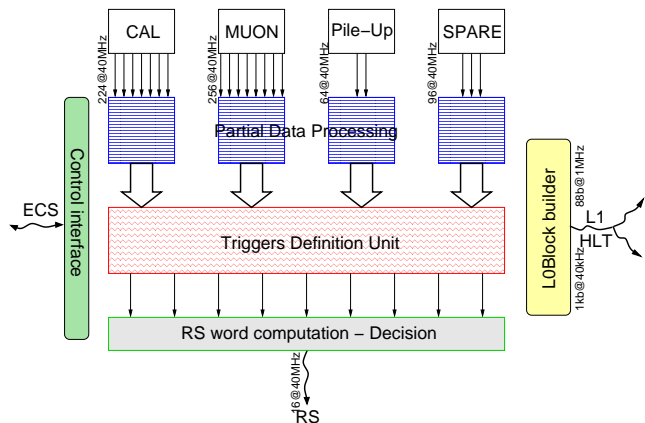


Figure 5.1: L0DU logical architecture

The ECS control interface manages the algorithm parameters, the algorithm behaviour, reset scenarios, slow control, online debugging and monitoring and many other tasks like FPGA programming.

Like a detector Front-End electronics board, the decision unit is able to send data to L1 and the HLT.

5.3 L0DU Prototype

An L0DU prototype was assembled at the beginning of 2002. It is a simplified version of what is foreseen for the final L0DU. It has neither ECS nor TTC connection and has only a reduced number of inputs (96 bits) and outputs (32 bits). Inputs and outputs are transmitted in LVDS format at 40 MHz via standard Ethernet cables. Figure 5.2 shows the prototype of the decision unit being tested. This prototype offers a maximum flexibility and adaptability to test a large part of the final L0DU functionalities including L0Block building operations.



Figure 5.2: L0DU prototype and its test set-up

To avoid connections between different modules the final L0DU will be implemented in a single board.

In the prototype, the use of a single FPGA would have been sufficient in terms of number of input/output and internal memory resources, but to be more realistic in relation to the final version, the prototype was made up of five interconnected FPGA¹.

Several simple algorithms were implemented and tested successfully [46] when the latencies of the various sources were emulated.

After a first step of time alignment of the different sources, thresholds are applied on the data. Each intermediate condition

is individually downscaled, rate divided or masked under a given set of parameters. Finally the decision is taken if a combination of conditions is realised.

Tests of the L0DU require the design of a specific test bench. A first version was built to test the L0DU prototype. It is made up of several “memory” boards synchronized by a “clock generator” board. Each board allows 64 bi-directional input/outputs to be driven or received into standard Ethernet cable with LVDS levels at a 40 MHz frequency. The memory boards are both used to store the stimuli and the outputs of the tested system. The maximum number of data words is 2^{16} in the current design. Three memory boards and one clock generator were used to test the first L0DU prototype (Figure 5.2).

The user defined stimuli and the data from the system under tests are downloaded or read out through a VME bus.

In addition, the clock generator board delivers a synchronization signal as a reference for the whole test bench timing. Up to 16 memory boards can be synchronised.

A C++ acquisition program running on a Linux PC controls the memory boards through VME. It makes a bit-to-bit comparison of the L0DU prototype outputs with the results of a C++ simulation performed with the same stimuli.

The prototype was fully functional. Due to the synchronous pipe-lined architecture it should be easily scaled to the final version of the unit. Recent FPGA technologies providing more input/output ports in many formats and more internal memory will be used.

5.4 Studies for the final implementation

In order to reduce the number of incoming cables and connectors, the data will be serialized. Optical links, allowing the transfer

¹ACEX1K100 were used

of 16 bits at 80 MHz, are a good candidate to exchange data with the L0 sub-triggers.

Most sub-detectors of LHCb use the TELL1-board [16] for buffering their data and interfacing to L1 and the HLT. This board includes many components needed by the L0DU; it is described in more detail in Section 6.1.2. Its use to implement directly L0DU is appealing and under study.

The TELL1-board has some modular input mezzanines, where is implemented the receiver part, linked to the mother board through 4 connectors. The L0DU itself could be a specific mezzanine (9U height and about 16.5 cm wide) including :

- input interfaces from trigger processors (re-use of the design of the optical mezzanines);
- the FPGA implemented L0DU algorithms;
- output interfaces to the Readout Supervisor with LVDS signals;
- Level-0 pipeline.

The mother board connectors would be used to receive TTC and ECS signals and send data to L1 and the HLT. The L0DU still remains a specific board but strongly linked to the TELL1-board.

5.5 Debugging and monitoring

A simplified version of the external test bench, described in Section 5.3, is foreseen to be integrated on the L0DU. It will ensure that the L0DU is still working correctly by injecting data patterns designed to make an easy and fast diagnostic of possible problems. Meanwhile the full debugging and the maintenance will be performed with the external test bench. In that way, a copy of the L0DU will be maintained permanently to ensure the full availability of the unit.

Online monitoring functions will be implemented through ECS. Counters will pro-

vide statistics on taken decisions and intermediate results allowing a measurement of the L0 trigger performances.

Chapter 6 Level-1 and High Level Trigger

In this chapter we describe the technical design of Level-1 and the HLT, addressing both its hardware implementation and the algorithms used to take the trigger decisions.

Essential input parameters for the design are the average values of (a) the amount of event data sent from the FE electronics to the CPU farm and (b) the CPU processing time of the trigger algorithms. The former sets the scale for the network and, for a given readout network technology, it defines the number of data sources per subdetector. The latter sets the minimum number of CPUs needed in the farm. The distributions of these input parameters must also be known, to some extent, in order to verify that the system is well-behaved even in the presence of tail events with large data sizes and/or large processing times. All event data sizes and processing times shown in this document have been obtained from the standard LHCb simulation framework, which is described in Chapter 7. Minimum-bias events were processed to produce L0-accepted and L1-accepted event samples with realistic data sizes and processing times. These samples have then been used as input to the L1/HLT network and CPU farm simulations.

The requirement for the implementation is to be flexible in the assignment of processing nodes to either Level-1 or HLT, and to be easily scalable as the need arises.

L1 makes use of data from the VELO, TT, L0DU and Calorimeter Selection Crate, whereas the HLT has access to all data. The VELO and TT provide the min-




Figure 6.1: The number of clusters per event in the VELO and TT for events passing Level-0.

imum information required to obtain precise impact parameter measurements and a rough estimation of the particle momenta by using angles and deflections of tracks in the upstream fringe field of the spectrometer magnet. Events are selected by requiring at least two tracks with large p_T and significant impact parameter to the primary vertex. The muons from the L0DU and clusters from the Calorimeter Selection Crate allow further enhancement of the signal purity by matching VELO tracks to these L0 high- E_T candidates.

For VELO and TT, the L1 cluster information can be encoded in 2 Bytes with sufficient spatial resolution. Hence, the data size per event is roughly given by $N_{cl} \times 2$ Bytes, plus some L1 board header information (4 Bytes per board). Here, N_{cl} is the number of L1 clusters in an event. The N_{cl} distributions after L0 for VELO and TT are shown in Figure 6.1. The number of data sources and event fragment sizes

are summarized in Table 6.1.

HLT has access to the full event data, and is executed on the same commodity CPU farm as L1. The algorithm first confirms the L0 and L1 triggers with better precision, and then mimics the off-line selection algorithms for the various channels to reduce the rate to 200 Hz, at which rate events will be written to storage. The total raw event size is approximately 31 kBytes.

Table 6.1: Number of L1 data sources and average event fragment size per source, which does not include the transport overhead.

Subsystem	Number of sources	Data/source [Bytes]
VELO	76	36
TT	48	24
L0DU	1	86
Calo Crate	1	70

The total size of the readout network for Level-1 and the HLT has been based on the simulation results shown above. Rather than overdesigning the system to be able to cope with the unforeseen, the system is required to be scalable to be able to adapt quickly to actual needs. In addition, the 40 kHz L1 output rate is dominated by events which trigger because a track is wrongly assigned to have large transverse momentum. Hence, the L1 algorithm would benefit significantly from a more precise momentum determination. This can be achieved by providing L1 with Tracker (T1–T3) data. In addition, the use of Muon (M2–M5) data would increase significantly the L1 efficiency for channels with muons in the final state. The scalability of the L1/HLT system is presented in Appendix A.

Next, we give a detailed description of the L1/HLT technical design, discussing first the hardware architecture and implementation, and subsequently the L1 and HLT algorithms.

6.1 Level-1 and HLT hardware implementation

The Level-1 trigger and HLT algorithms operate on general-purpose CPUs. The input data come from the front-end electronics of the detectors included in the system, which are the VELO and TT, together with data from the L0-trigger. In this section the Data Acquisition (DAQ) system for the Level-1 and HLT is described, which will collect the event fragments from the Front-End electronics boards, assemble them into complete events and deliver them to a CPU in a computer farm.

In the case of the Level-1 trigger the event data are buffered in the front-end electronics until a decision has been taken. It is therefore important to keep the latency for the whole process of data transport and event assembly as short as possible, to allow maximal time for the execution of the algorithm. The system provides an environment for the physics algorithms, in which they can run unchanged from the “offline” environment. Some adaptations of low-level services of the software framework (“Gaudi” [53]) are however required.

The technological challenge in the system consists of handling the high rate of data using commercial and (to a large extent) commodity equipment, while transporting and assembling the data as quickly as possible.

The High Level Trigger uses the full event data and operates at the accept rate of the Level-1 trigger. From a data acquisition point of view the problem is very similar, the main difference being that there are many more data sources sending larger fragments, but at a much reduced rate. However the aggregated traffic is significantly smaller than that from the Level-1 trigger. The HLT algorithm also needs completely assembled events and runs on a general purpose CPU. In this case there is no latency

limit due to limited front-end buffers since the buffering of the events is done in the CPUs.

A system for performing the DAQ for the HLT has been described in the Online System TDR [9]. The system described here is an evolution of the architecture described there, which does the data acquisition and event assembly for both trigger levels using the same infrastructure. The system presented here supersedes what has been written on the data-flow in [9]. The other parts of [9], which deal with the TFC, ECS and general infrastructure remain unchanged, except that their scale is adapted accordingly. The key characteristics of the data-flow system are:

- Copper Gigabit Ethernet is used as a link technology. The connectivity between the sources and the destinations is provided by a large switching network.
- Data are pushed through, every source sends when it is ready to do so. Flow control is exercised centrally by disabling the trigger at the level preceding the one at which the problem is detected via the TFC system.
- HLT and Level-1 data share the infrastructure and the HLT and Level-1 algorithms run concurrently in the CPUs.
- Event fragments are packed in the data sources to reduce the packet rate in the system.

6.1.1 Architecture

The architecture is most easily explained by following the data-flow from the sources, the front-end electronics boards, to the ultimate destinations, the CPU nodes, as shown in Figure 6.2.

The front-end electronics is required to be able to store 58254 events [11]. The minimal time between events entering the sys-

tem is 900 ns. This means that in order to avoid buffer overflow the maximum time from the entry of the event into the system until the trigger decision is transmitted to the front-end electronics is 52.4 ms. The decisions are distributed via the TFC system, described in [9]. All front-end boards use the same standard Gigabit Ethernet plugin card to send the data [47]. This card has two output links, one of which is used for HLT data and one for Level-1 data, if applicable¹.

In order to reduce the packet rate in the system, the front-end electronics is required to pack several event fragments in a multi-event packet (MEP). Due to fluctuations in the fragment size, the resulting MEP can become bigger than the Maximum Transmission Unit (MTU) of Ethernet, which is 1500 bytes of payload [48]. The front-end electronics must then split the MEP into several Ethernet frames. Because we want to use standard protocols wherever possible, the front-end electronics is required to format the MEP as an Internet Protocol (IP) packet. The IPv4 standard defines also the way to split up a packet into Ethernet frames [49].

The packing factor, i.e. how many event fragments are put into one MEP, is an adjustable parameter of the system. For Level-1 the maximum packing factor is defined to be 32 and for HLT 16. Packing factors which frequently require a packet to be split into several Ethernet frames are not useful, because it is the Ethernet frame rate which is the problem for the receiving end.

From the current knowledge of the data fragment sizes per board, good working points for the packing factors are 25 for the Level-1 and 10 for the HLT.

To reduce the size of the central readout network it is foreseen to do some multiplexing using Ethernet switches before going into the main readout network switch.

¹Technically it is also possible to use both links for L1 and HLT traffic.

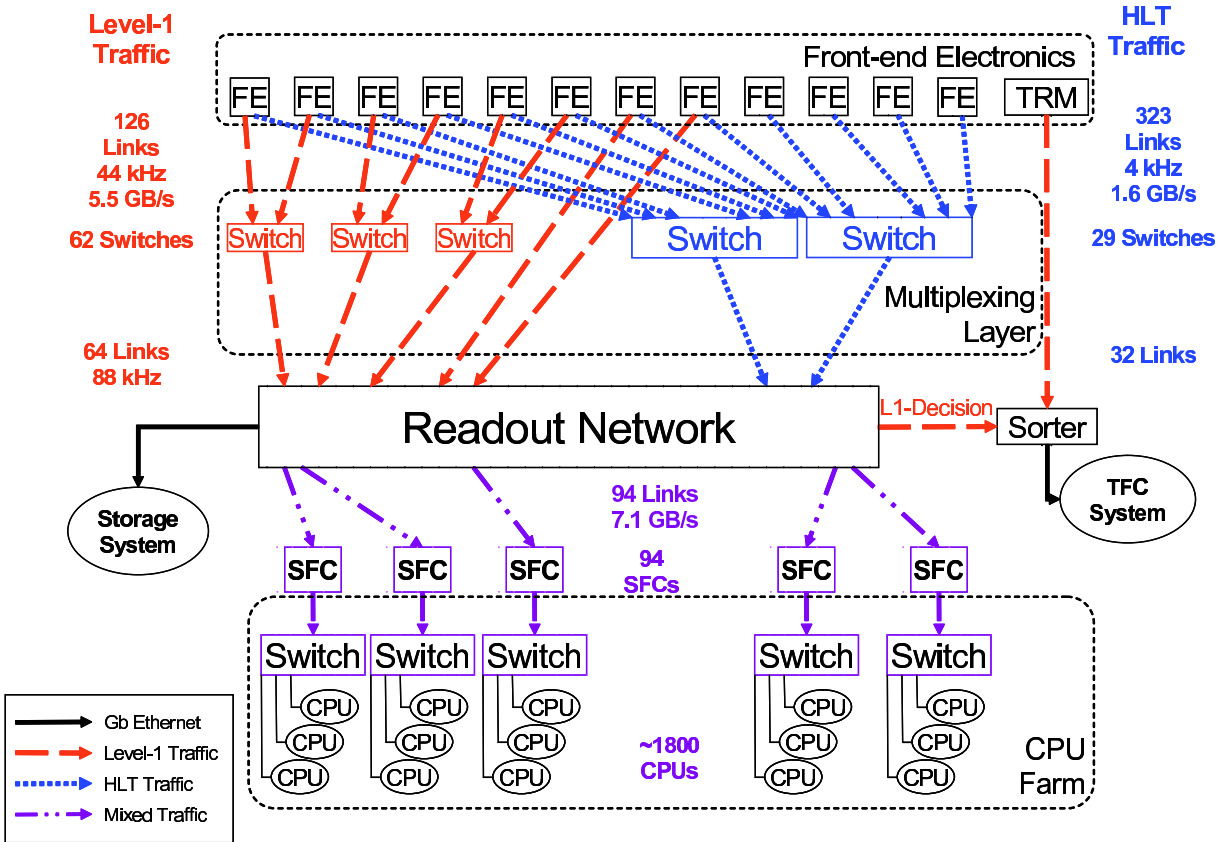


Figure 6.2: The architecture of the DAQ system. The ECS and the throttle signals are not shown.

The multiplexing increases the packet-rate on the outgoing link towards the readout network². Data are then pushed through a large high-performance Ethernet switch to a sub-farm controller.

The sub-farm controllers (SFC) sit at the down-stream end of the readout network. They perform the event-building, where individual event fragments from the MEPs are assembled in correct order into events. They distribute the events to the compute nodes connected to them via another Gigabit Ethernet switch. The SFC exercises dynamic load balancing among the nodes. Each node is only processing one Level-1 event at any given time. This means that Level-1 events will have to queue in the SFC, when there is no available node, and that time-out mechanisms must be implemented in the nodes.

²The total rate is of course not increased.

Simulation has been used to investigate the additional latency suffered by events due to the queueing in the SFCs and due to the packing of events into MEPs [50]. Figure 6.3 shows the number of events in time-out as a function of the maximum processing time allowed. The underlying simulation includes a complete model of all features of the queueing and load balancing in a sub-farm and a model of the expected performance of the Level-1 trigger algorithm. The model for the processing time of the Level-1 algorithm has a cut-off at 50 ms, which explains the steep drop around 50 ms in the insert. No significant increase of the latency has been found.

Since the mean time for reaching a Level-1 decision is much shorter than the allowed maximum, the nodes will not always be busy. Optimal usage of the total available CPU power is achieved by running the




Figure 6.3: Fraction of events exceeding the maximum processing time in a sub-farm as a function of the time cut-off. The solid and dotted lines show the simulation result for a scenario with and without event-packing, respectively.

HLT as a background task, which is interrupted whenever a Level-1 event needs to be processed. Switching between the two tasks is done by the operating system, and has been measured to take less than $10 \mu\text{s}$ [50].

After a trigger algorithm has finished, a result is sent back to the SFC. For a Level-1 event the decision contains only a short summary block, which is forwarded to the *Level-1 decision sorter* described in the next paragraph. In the case of the HLT data, accepted events will undergo full reconstruction and the reconstructed data will be sent together with the raw data to permanent storage. To this end the SFCs will be connected to the storage either via the event-building network itself or via a small dedicated network.

Level-1 decisions

The Level-1 decisions are small Ethernet packets, which are sent to the TFC system. In particular they are received by the Readout Supervisor, which makes the ultimate decision about whether to accept an event or not, because it knows about the state of the throttle signals. Since it also sends the trigger decisions to the front-end it has an easy way to measure the actual time an event has spent in the system, and can react to a time-out for Level-1 events³.

The Readout Supervisor requires the Level-1 trigger decisions to arrive in the order they entered the system, i.e. in ascending order of Level-0 event numbers. This is the task of the Level-1 decision sorter, shown on the right side of Figure 6.2. The sorter must be informed when an event enters the system. This is done by the Trigger Receiver Module (TRM), shown in the upper right part of Figure 6.2.

Destination Assignment

The destination assignment is central and static. The Readout Supervisor broadcasts the destination for each Level-1 and HLT MEP. The broadcast contains among other information the 10 least significant bits of the IP address of the SFC to which this MEP should be sent. The Readout Supervisor keeps a table from which it selects the destinations. By making SFCs appear more often than others in this table a coarse static load balancing can be achieved. The advantage for the front-end electronics is to avoid keeping a relatively big table of addresses. Details can be found in [54].

Control and Monitoring

All the software and hardware components described here will be configured, controlled

³Time-outs for HLT events are not so critical, because there are no buffers in the front-end electronics which could overflow, so time-outs are decided locally in the farm, just to avoid wasting time on an excessively complicated event.

and monitored using the Experiment Control System (ECS), which is described in [9, 51]. The ECS interfaces to the equipment via Ethernet. Following a general design principle of the online system to separate everywhere data and control paths, a separate Ethernet Local Area Network (LAN) is used to connect the equipment. This is done either directly by using a second network interface card in the farm-nodes or SFCs, or indirectly by means of a controls PC, which in turn accesses the hardware by one of the agreed interfaces to electronics described in [9]. The ECS also collects data from all farm-nodes for online monitoring and quality checking.

6.1.2 Implementation

The first stage of the system is part of the front-end electronics of the sub-detectors. All detectors are required to use the same custom-made Gigabit Ethernet interface card. Its design is finished and a prototype is expected soon. In many aspects it corresponds to a standard network interface controller (NIC), found in PCs [47]. The current implementation has two independent Gigabit Ethernet ports and thus can support a theoretical maximum data output rate of 250 MB/s. A future version is planned which will have 4 ports, with a maximum data output rate of 500 MB/s.

FE interface to the DAQ

The Gigabit Ethernet interface is limited to the Ethernet protocol. The formatting of the MEP data into a IPv4 packet, which may span several Ethernet frames, is the responsibility of the motherboard. Since IP is a very simple protocol, this poses no problem for the powerful FPGAs used in the FE. All sub-detectors which send data to Level-1 use the TELL1-board [16] to receive the data from the Level-0 electronics, process them for Level-1 and send them to the

event builder. The data are buffered until a Level-1 decision has been reached, and, if accepted, are sent to the event builder for HLT processing. In the following a short description of the key features of the TELL1-board is given.

The data from the various sub-detectors can be received either via digital optical or analogue electrical links. For this purpose mezzanine cards are used, which are shown in the top part of Figure 6.4. The data are digitised for the electrical links and deserialised for the optical links. To cope with the quite diverse processing requirements of the sub-detectors the data are then passed through several large FPGAs. An example of the processing to be done is given in [55], focusing on the VELO, but applicable also to TT and IT. The data are then stored in a buffer implemented using standard DDR-SDRAM memory. The Level-1 buffer holds 58254 events. For detectors, which send data to the Level-1 trigger the data are then forwarded to the SyncLink FPGA (shown in the lower part of Figure 6.4), which performs the following tasks:

- link the fragments from all processing FPGAs into one fragment;
- perform any processing required for completely assembled fragments;
- buffer the assembled fragments until the number of events in a MEP has been reached;
- pack the MEP into an IPv4 packet;
- if necessary segment the IPv4 packet into several Ethernet frames;
- send the frames via the Gigabit Ethernet card (RO-Tx in the figure) to the event builder.

For all detectors upon reception of a Level-1 yes via the TTC system, the event fragments are pulled into the SyncLink FPGA. It then goes through the same steps as for Level-1 events, except that the




Figure 6.4: Building blocks of the TELL1-board. In the first row the input stage is shown with optical and electrical input mezzanines. The second row shows the processing FPGAs and the Level-1 buffers. The third row shows the SyncLink FPGA, which assembles the sub-fragments and pushes the formatted events to the network. The fourth and last row shows the interfaces to the external systems like the ECS, the TFC and the network (RO-Tx).

fragment processing for HLT is different⁴ and an additional link is used to send the frames.

PCs and switches

Most of the other components are commercially available. The aggregation and sub-farm switches are relatively cheap Gigabit Ethernet switches, typically found in high performance LAN (Local Area Network) installations. Full connectivity at maximum speed is not required, because most of the links are not fully loaded. On the other hand, the core switch of the readout network must provide full performance and generous buffering to cope with the traffic. Such devices are typically found in the

⁴At 40 kHz more sophisticated processing can be performed.

backbone of large campus networks. Key parameters for this switch can be extracted from simulation [50]. If monolithic switches with a sufficient number of ports cannot be found, or are very expensive, then the switch can be built from smaller components. Various interconnection topologies are possible. While detailed studies can be found in [50], in general it is found that a system built from several layers has advantages because of the distributed buffering. The forwarding latency, however, is increased slightly. From an operational point of view, a single unit is easier to manage. In the end this question will be decided based on the cost of the solutions.

The sub-farm controller, whose architecture is described in [9], is a high-performance PC. The emphasis on performance is mainly on the I/O capabilities, because it is required to handle at least two Gigabit/s of data. Such PCs are already available. To achieve maximum throughput care must be taken in selecting high-performance network interfaces, which support advanced DMA (Direct Memory Access) features and buffering. The required packet rate of 80 kHz can be sustained with today's hardware without problems; better hardware and custom software will allow raising this limit even further.

The farm-nodes will be chosen according to the best obtainable price/performance ratio. They will be operated disk-less and require, apart from CPU power, sufficient memory to run without a page-file and two network interfaces in order to maintain the separation of data and control paths.

The detailed implementation of the CPU farm is the subject of ongoing studies. These comprise the physical realisation of the farm as a system of water-cooled racks with rack mounted PCs, as well as the software to control and monitor the CPUs and the event distribution from the SFCs to the worker CPUs.

Packet loss and error handling

Ethernet does not guarantee reliable frame delivery. For performance reasons, no reliable higher-level protocol like TCP is used in the system, except for sending events to permanent storage. The system is however error detecting at all stages, using sequence numbers in the transport protocols and time-outs. There are two major reasons for packet loss: bit errors and dropped packets in the switches due to congestions. Bit errors can be detected reliably by the Ethernet checksum. Measurements with modern Ethernet hardware on unshielded twisted pair cables considerably longer than anything foreseen in the experiment have shown bit error rates better than 10^{-14} . Packet dropping in the switch can only be avoided by selecting a suitable switch. The parameters for such a switch will be obtained from simulation. From current simulations it is known that large output buffer memories are required. Candidate switches will be verified in a test-bed. Experience with recent high-end routers (such as used for the CERN campus network) shows that packet-loss in the switches is extremely rare.

TRM and Level-1 decision sorter

The TRM is implemented using a TELL1-board [16], which simply sends a packet to the Level-1 decision sorter upon reception of the Level-0 accept decision from the TFC system.

The decision sorter itself is implemented using a commercially available PCI card, which allows fast packet processing using a Network Processor (NP). The code for sorting and performance results are described in [50]. An alternative implemented entirely in the TFC system is described in [54].

The implementation of the remaining infrastructure and hardware is as described

in [9].

Size of the system

To determine the size of the system, the number of front-end boards is the first basic input. This defines the number of links to be read out. For the Level-1 trigger, where VELO, TT, L0DU and the Calorimeter Selection Crate are read-out, this results in 126 before aggregation and 64 links into the event building network; for the HLT there are 323 before aggregation and 32 after. Assuming a packing factor of 25 (i.e. 25 events are sent in one packet) for the Level-1 and 10 for the HLT, multiplexing factors can be calculated to get to an average link load of 80%.

Table 6.2: Base target parameters for the L1/HLT implementation. The first four are given externally, the others are chosen. The overheads describe the size of the header, which is needed to describe the data contents for the event-builder [52].

L0 accept rate	1 MHz
L1 accept rate	40 kHz
L1 transport overhead / MEP	48 bytes
HLT transport overhead / MEP	48 bytes
L1 packing factor	25
HLT packing factor	10
Input link rate	< 100 MB/s
Output link rate	< 100 MB/s
Frame rate at output	< 80 kHz

The target parameters are summarised in Table 6.2. They either stand for available resources like CPU nodes and switch ports, or load factors like the link rates and frame rates. The other input into the system is given by the expected average data size per fragment per front-end board. These numbers are taken from the full detector simulation described in Chapter 7. At the level of the system design only the overheads due to the data transport format (in particular the IP header) are added. A relatively straightforward minimisation exercise yields the fi-

nal required numbers of switch-ports in the event-building switch. Extra ports need to be added to connect the storage-system and the L1-decision sorter.

Table 6.3: Key performance figures of the system.

Event Building	
Total Frame Rate at RN input [kHz]	7248
RN output links	94
RN output link rate [MB/s]	47.9
Frame rate (L1) per link [kHz]	59.9
Frame rate (HLT) per link [kHz]	20.0
Total frame rate at RN output [kHz]	79.6
MEP rate (L1) per link [kHz]	0.47
MEP rate (HLT) per link [kHz]	0.05
Total MEP rate	0.53
Trigger farms	
Sub-farms	94
Event rate/sub-farm (L1) [kHz]	11.7
Event rate/sub-farm (HLT) [kHz]	0.4
Processors/sub-farm	21
Event rate per processor (L1) [kHz]	0.56
Event rate per processor (HLT) [kHz]	0.02

In Table 6.3 the key performance figures of the system are summarised including numbers for an extended network to include the tracking detectors. The system is dominated by the Level-1 traffic. It should be noted that the output links from the event building network are not very heavily loaded. In fact the number is determined by the frame rate limit of 80 kHz, which is chosen to protect the SFC from too high an interrupt rate. It is likely that with better hardware and custom software significantly better results can be achieved, which would allow reducing the size of the system.

6.2 The Level-1 Algorithm

Level-1 exploits the finite lifetime of the B-mesons in addition to the large B-meson mass as a further signature to improve the purity of the selected events. All results assume the following information is used by

Level-1:

1. The L0DU summary information and data from the Calorimeter Selection Boards.
2. The VELO measurements of the radial and angular position of the tracks, in silicon planes perpendicular to the beam-line between radii of 8 mm and 42 mm. The strip layout of the sensors is shown in Figure 6.5.




Figure 6.5: The strip layout of the VELO sensors as viewed along the beam line. Two R-sensors (top) are shown, with their strips subdivided in octants. In ϕ -sensors (bottom) the strips make an angle between $10-20^\circ$ with the radial, and strips are subdivided in two regions. The dotted lines indicate the two ϕ -detectors downstream, which are rotated around the y -axis. The lines to route the signals to the electronics located at the periphery of the sensors are not shown.

3. The Trigger Tracker (TT) measurements from its four silicon planes, two with vertical strips and two with a $\pm 5^\circ$ stereo angle.




Figure 6.6: Event display of the result of the 2D tracking in the VELO detector, showing all hits and reconstructed tracks in a slice of 45° of the VELO R-sensors in an event where 72 forward 2D tracks were reconstructed.

B-mesons with their decay products in the LHCb acceptance move predominantly forward along the beam-line, which implies that the projection of the impact parameter in the plane defined by the beam-line and the track is large, while in the plane perpendicular to the beam it is almost indistinguishable with respect to primary tracks. The L1-algorithm exploits this by reconstructing so-called 2D tracks using only the VELO R-sensors. The 2D tracks are sufficient to measure the position of the primary vertex since the strips at constant radius are segmented in 45° ϕ -slices. Muon tracks are identified by matching 2D tracks to Level-0 muon candidates. A fraction of the 2D tracks is selected based on their impact parameter and their match to Level-0 muons, and these 2D tracks are combined with the ϕ -sensor clusters to form 3D tracks. By combining 3D tracks with hits in TT the momentum of these tracks can be measured using the fringe field of the magnet between VELO and TT. In the following sections the reconstruction algorithm and its performance will be described in more detail, while the combination of B-signatures to form a trigger decision will be given in Chapter 7.

6.2.1 VELO 2D Track Reconstruction

Using the information from the R-sensors, tracks are reconstructed in the rz view, in which view tracks originating from the beam-line form straight lines. The method is based on triplet search, a triplet defined as three points in consecutive VELO stations, and in the same octant, compatible with a straight line. The triplets are combined into longer segments using a dedicated fast algorithm based on a manipulation of integer indexes [56]. Figure 6.6 shows an event display of the result of the 2D track search in a 45° slice of the VELO.

6.2.2 Primary Vertex Search

In the next step a primary vertex search is performed [56]. 2D tracks are projected onto the central axis of their octant, and these projections are combined with tracks from perpendicular octants to estimate the position of the primary vertex. The initial estimate of the primary vertex is based on histogramming and peak finding, followed by iteratively rejecting outliers and re-fitting the vertex. After three iterations, a precision of $\text{RMS}_{x,y}^{\text{PV}} = 25 \mu\text{m}$ and $\text{RMS}_z^{\text{PV}} = 60 \mu\text{m}$ is reached.

6.2.3 Level-0 Object Matching to 2D Tracks

The electron and hadron candidates above a given E_T threshold, and the largest p_T muons, are matched to 2D tracks to identify candidates for the 3D tracking described below. The details of the procedure can be found in [57]. The matching is performed by comparing dr/dz and azimuthal angle ϕ between VELO tracks and Level-0 objects. A χ^2 of the match is formed based on the p_T kick and the energy resolution of the Level-0 objects and $\delta(\phi) = 45^\circ/\sqrt{12}$ of the 2D track. For the moment only muon candidates with relatively small χ^2 are accepted for 3D confirmation.

6.2.4 VELO 3D Track Reconstruction

In addition to the selection as explained in the previous section, 2D tracks with an impact parameter in the range 0.15 to 3 mm to the primary vertex are also selected. Only these 2D tracks are reconstructed in 3D to reduce the execution time of the algorithm. The ϕ -sensor information is linked to the 2D tracks taking into account the estimated position of the primary vertex [56]. The combined 2D and 3D reconstruction efficiency for reconstructible B-decay products⁵ is 94% and the ghost rate is 5.9%. The momentum information is not available at the level of VELO reconstruction and thus the covariance matrix of a track cannot include the contribution from multiple scattering in the traversed material. It was shown [58] that assuming multiple scattering contributions corresponding to a 3 GeV particle gives optimal parameters for extrapolation to TT, and for the measurement of the impact parameter.

⁵A particle is considered reconstructible if it has at least three hits in VELO R-sensors, three hits in ϕ -sensors, and at least one x and one stereo hit in each station of T1–T3. A track is considered to be found if 70% of its reconstructed hits originate from a single Monte Carlo particle.

6.2.5 Level-0 Object Matching to 3D Tracks

In the next step of the reconstruction the matching between Level-0 objects and 2D tracks is confirmed for the corresponding 3D tracks to improve the ghost rejection thanks to the increased precision in ϕ . The uncertainties of the slopes of Level-0 objects dominate the error, and hence the contribution from the VELO tracks uncertainty to the matching χ^2 is ignored. The purities and efficiencies obtained for typical χ^2 cuts are presented in Table 6.4. The purity is defined as a fraction of correct matches normalized to all matches in a signal sample, while the efficiency is given for B-decay tracks and shows the rate of tracks matched to L0 candidate, provided that the track and L0 candidate are both reconstructed. Only the matching to muons is used for the moment in the performance given in Chapter 7.

Table 6.4: The performance of Level-0 object matching to a 3D VELO track

3D tracks	χ_{\max}^2	purity	efficiency	$\sigma(p)/p$
muons	16	51.2%	94.7%	6%
electrons	4	32.9%	95.8%	12%
hadrons	4	26.9%	92.8%	15%

6.2.6 VELO TT Matching and Momentum Determination

The reconstructed 3D VELO tracks are combined with hits in TT, which is located about 250 cm downstream from the interaction point. A perspective view of the TT station as it is modelled in the simulation is shown in Figure 6.7. The 3D track parameters at the farthest downstream point observed in the VELO are used to search for the track segment passing the four layers of TT. A 3D track is extrapolated to the z -position of each layer, and its distance




Figure 6.7: A front view of the TT station. Adjacent silicon sensors of the same colour are daisy chained in the readout.

to all hits inside a ± 20 mm window is calculated. These distances are then rescaled to a reference plane perpendicular to the beam line at the middle position of TT, taking into account the expected slope of the track, and filled into a histogram. In this histogram accumulations of at least three clusters in five consecutive bins are searched for. They are analyzed in an increasing order of their mean distances to the straight-line extrapolation. For each accumulation a least-squares fit of a straight line is made, and based on their χ^2 , the worst points are rejected iteratively. The procedure stops when the χ^2 becomes acceptable for at least three TT hits. Otherwise the accumulation is rejected and the next one from the list is considered until the list is exhausted. For each accepted accumulation the momentum is determined by a fit to the slopes of the track at VELO and at TT, and the integral of magnetic field in-between. The vertical component of the magnetic field and the corresponding $\int Bdl$ between the VELO and TT are shown in Figure 6.8, and is close to optimum [59] for the Level-1 performance. The VELO-TT reconstruction was tuned to optimize the Level-1 performance, which requires that for high- p_T purity prevails over efficiency. Figure 6.9 shows how the efficiency and resolution vary as a func-




Figure 6.8: The vertical component of the magnetic field (a) and the corresponding $\int Bdl$ (b) of the fringe field. The vertical lines indicate the position of the TT station.

tion of p_T at that optimized working point.




Figure 6.9: The efficiency (a) and p_T resolution (b) of VELO-TT matching at the optimal working point as a function of p_T .

6.2.7 L1 timing

A crucial aspect of the Level-1 algorithm is its execution time. A full study of Level-1 execution-time optimization of the

code which is used for all performance numbers in this TDR and the Reoptimization TDR can be found in [60]. All critical parts were identified and the logic of the algorithm was tuned to meet the timing constraints. The execution flow was set-up to minimize a number of calculations needed to make the negative decision that is expected for the majority of Level-0 accepted events. On average 8.5 tracks need to be reconstructed in 3D from 58.4 forward 2D tracks in a minimum-bias event that passes Level-0. The final phase of the algorithm, in which all B-signatures are combined to obtain the Level-1 decision, is described in Chapter 7, but its execution time is accounted for here. The timing of the most important phases of the Level-1 algorithm are detailed in Table 6.5. On average

Table 6.5: The timing of various phases of the Level-1 algorithm as measured on a 1 GHz Pentium III Linux processor.

Level-1 phase	time [ms]
initialization	1.1
2D tracks	2.4
PV fit	0.5
2D selection	1.3
3D tracks	1.4
refit of 3D tracks	0.2
TT matching	0.9
L0 3D matching	0.4
decision	0.1

8.3 ms is spent in Level-1 with the algorithm described above on a 1 GHz Pentium III processor. The time was measured as the real time elapsed between start and stop of the Level-1 algorithm with a granularity of 1 μ s. An overhead of about 20% due to communication between different parts based on creation of dynamic containers has been subtracted. The distribution of total execution time is presented in Figure 6.10. The code described above has not been further optimized for execution time performance to allow stable code to be used for




Figure 6.10: The distribution of the Level-1 execution time on a 1 GHz Pentium III processor for minimum bias events accepted by Level-0.

the large scale production for the physics studies. In the mean time faster code has been developed [62] with a similar track finding efficiency. With this new code we expect to have an average execution time per event of less than 1 ms in 2007⁶.

6.3 The HLT Algorithm

The purpose of the HLT algorithm is to reject as fast as possible events not compatible with an interesting b decay, using the final quality information from the tracking system. Basically, the tracking is performed first in the VELO, then in the whole tracking system to get the track's momentum with almost the full accuracy. Particle identification is limited to muon and electron identification, as processing the full RICH information would be too slow. With an L1 processing time of less than 1 ms, the processing power remaining for the HLT corresponds to around 25% of the farm, which translates to about 10 ms per event, 60 ms on today's 1 GHz Pentium III processor.

⁶A 1 GHz PIII processor delivers 425 CERN Units (CU), while it is expected [61] that in 2007 the performance will be 2500 CU per node.

6.3.1 VELO tracking

An implementation of the HLT VELO tracking has been performed [62] with both efficiency and speed as main concerns. Starting from the HLT event, it uses the full VELO information to get a more accurate cluster position than in L1. The 2D tracking in the rz -plane is performed, using only the R-sensors, and then the full space tracking, collecting clusters in the ϕ -sensors compatible with the 2D track.

The efficiency is measured on signal datasets, where the main concern is to reconstruct the tracks of the interesting b decay. This is not exactly the same as to get the best efficiency on all tracks, as the interesting b decay products tend to have larger transverse momentum.

Speed is measured on the expected input events, this means minimum-bias events passing the L0 and L1 triggers, on a 1.0 GHz Pentium III Linux CPU. Those events tend to be busier than average signal events, and thus the processing speed is a bit slower. The reconstruction efficiency for tracks passing the magnet is $95.5 \pm 0.2\%$, and for tracks from a b-decay $98.3 \pm 0.3\%$. The ghost rate is 5.7%, and the time per minimum-bias event is 5.9 ms.

6.3.2 VELO TT tracking

The first step to measure the momentum of a track is performed by finding the relevant hit in the TT station. Even if the momentum estimate has a low accuracy, at the 30% level for $\sigma(p)/p$, this reduces significantly the search region in the T stations, hence speeding up by a factor close to two the second step of the track search.

For pattern recognition purposes, the field between the VELO and TT has the same effect as an angular kick, proportional to $1/p$, of the track at a fixed z position z_{bend} . But multiple scattering is important and dilutes this simple relation. The VELO

TT tracking method is to loop on VELO tracks, and to project all TT hits onto a fixed plane, joining them to the track impact in the z_{bend} plane. Only a small subset of the hits needs to be handled, as a minimal momentum for the track implies a maximum distance for the hit. The window is typically ± 21 mm in the bending direction, corresponding to a minimal momentum of 1.5 GeV. In fact, the window is reduced at low angle, corresponding to a minimal transverse momentum of 100 MeV. The pattern recognition is simply a search for an accumulation of projected hits, with at least 3 of the 4 planes fired. Priority is given to 4-plane candidates, and to the highest momentum solution.

As the TT stations do not cover completely the acceptance, mainly around the beam pipe, tracks with less than 4 expected hits in TT are accepted, without momentum estimate. Tracks for which no sufficient accumulation was found are rejected, they are low momentum tracks. About 30% of the VELO tracks are discarded this way. The efficiency for interesting b decay tracks is over 99% compared to found VELO tracks. It takes less than 5 ms on a 1 GHz Pentium III to perform this first momentum estimate for about 40 tracks per event.

6.3.3 Long tracking

The search for hits in the T stations is based on the so-called ‘‘Forward tracking’’ [63], which uses the fact that as soon as one point is known after the magnet, the momentum is known and thus the complete trajectory. A fast method, using polynomial parametrizations, allows all hits to be projected onto a common plane. With the momentum estimate from the VELO-TT track or a minimal momentum requirement of 2 GeV/c, the range of hits can be reduced. An accumulation is searched for in the projection, and an iterative procedure,

fitting and removing the worst point, allows the best extrapolation of the VELO track to be found.

The key ingredient is a parametrization of the trajectory in the field, to obtain a fast projection. As the field is reasonably regular, and the minimum momentum $2 \text{ GeV}/c$, only a few polynomial terms need to be computed, as described in [63]. The efficiency to find the correct hits is over 98% for the pions of $B \rightarrow \pi\pi$ and a bit lower, 94% , for larger multiplicity decays. The speed of the algorithm is about 40 ms per event on a 1 GHz PIII CPU, with a momentum resolution $\sigma(p)/p$ around 0.6%.

6.3.4 Tracking summary

The total tracking efficiency for b decay tracks in the HLT varies between 92% and 96.5% depending on the decay channel. This takes around 50 ms per minimum-bias event accepted by L0 and L1. This figure is already in the allowed range, but can be improved by technical changes such as using a faster compiler, by improving the handling of the geometry information, and by tuning the cuts. For example, the “Long tracking” speed depends quite strongly on the minimum (transverse) momentum. The main concern is to be able to perform high quality tracking in the available time budget, and hence the tracking package was developed keeping the execution time in mind. The algorithm to efficiently select the signal events is under development, and hence emphasizes flexibility rather than optimization of execution time. This selection will be described in Section 7.3.

Chapter 7 Performance Simulation

Minimum-bias proton-proton interactions at $\sqrt{s} = 14$ TeV are generated using the PYTHIA 6.2 program [66], including hard QCD processes, single and double diffraction, and elastic scattering. Some of the PYTHIA parameters have been tuned to reproduce measured charged particle distributions for \sqrt{s} less than 1.8 TeV, and a detailed description of this tuning can be found in [2] and references therein. PYTHIA predicts a total inelastic cross-section of $\sigma_{\text{inel}} = 79.2$ mb, of which visible collisions¹ correspond to $(79.1 \pm 0.2)\%$ of σ_{inel} . Several inelastic proton-proton collisions may occur in the same bunch crossing, which is simulated assuming $\sigma_{\text{inel}} = 80$ mb and a non-empty bunch crossing frequency at the LHCb interaction point of 30 MHz. The luminosity L is assumed to decrease exponentially with a 10-hour luminosity lifetime in the course of 7-hour fills, with an average value of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$, which implies $\sim 2.8(1.4) \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$ at start(end) of fill. Studies have shown [67] that over this range of luminosities the efficiency for selecting signal events while re-tuning the trigger settings does not vary significantly. In this chapter all threshold settings will be given for a luminosity of $2 \times 10^{32} \text{ cm}^{-2}\text{s}^{-1}$.

Generated particles are traced through a detailed description of the spectrometer and its surroundings using the GEANT3 package [68], and the detector response

¹A collision is defined to be visible if it produces at least two charged particles with sufficient hits in the VELO and T1-3 to allow them to be reconstructible. Elastic scattering never results in tracks observed in the spectrometer.

is simulated taking the effect of the two preceding and one following bunch crossings into account, but ignoring the LHC bunch structure. All expected performances are based on these simulated detector responses, which include efficiencies, noise and cross-talk as appropriate for each sub-system.

Minimum-bias data corresponding to about 2 s of LHCb running have been generated, together with generating signal events by forcing a generated B-meson to decay into a specific final state. In the following sections the trigger settings and corresponding signal efficiencies will be presented which allow the minimum-bias event rate to be reduced to 1 MHz, 40 kHz and 200 Hz by Level-0, Level-1 and the HLT, respectively. In the final section of this chapter the sensitivity of the trigger to alternate settings of the PYTHIA parameters and a less than expected spectrometer performance will be presented.

7.1 Performance of Level-0

The performance of Level-0 is expressed in terms of efficiency $\varepsilon_{\text{L0}}^{\text{channel}}$, the fraction of offline selected events that pass Level-0 for a given signal channel.

The maximization of the Level-0 efficiencies is done by optimizing the bandwidth given to each sub-trigger taking into account correlations between them. Ideally, Level-0 Level-1 and the HLT should be optimized simultaneously. This scenario, however, was simplified by first determin-

ing cuts on the global event variables, hence SPD and Pile-Up multiplicity and the number of interactions [69], for a chosen set of thresholds on the other variables. Figure 7.1 shows $\varepsilon_{L0} \times \varepsilon_{L1}$ for a few channels as a function of the cuts applied on the SPD multiplicity. For each value, all thresholds




Figure 7.1: The top plot shows $\varepsilon_{L0} \times \varepsilon_{L1}$ as a function of the SPD multiplicity cut for events which have been accepted by the Pile-Up cut, and have a Pile-Up hit multiplicity below 112. The bottom plot shows the data size after Level-0 and the Level-1 execution time for the same SPD multiplicity cuts, normalized to no SPD cut. The dashed line indicates the chosen working point.

on the transverse energy of B-decay candidates were scaled by a common factor to re-adjust the total output of Level-0 to 1 MHz, and Level-1 to 40 kHz on minimum-bias events. As is explained in the next section, events are accepted if the sum of the transverse momenta of the two largest p_T muons is above its threshold irrespective of the global cuts. Hence, channels with muons in their final state show an increased effi-

ciency while tightening the global cuts. The values of the cuts on the global event variables are given in Table 7.1, and these cuts are used throughout the evaluation of the trigger performance described in the next sections. Cuts on the global event variables allow to tune the size of the events which pass Level-0, and their corresponding execution time in Level-1, with only small variations in signal efficiency. The chosen cuts are rather conservative, just to show the principle.

Events are only accepted if “Total E_T ”, which is a measure of the total energy deposited in the HCAL, is above 5 GeV, to reduce the possibility of triggers on halo-muons in crossings without interactions.

Table 7.1: List of L0-cuts on the global event variables.

Global Cuts	Value
Tracks in 2 nd vertex	3
Pile-Up Multiplicity	112 hits
SPD Multiplicity	280 hits
Total E_T	5.0 GeV

7.1.1 Bandwidth Division

We determine the L0 thresholds using a set of decay channels given in Table 7.4, which are representative both in giving access to the CKM parameters, and in the way they rely on the different trigger components [70].

The Level-0 bandwidth division minimizes the overall loss in efficiency by maximizing:

$$\sum_{\text{channels}} \frac{\varepsilon_{L0}^{\text{channel}}}{\varepsilon_{L0-\text{max}}^{\text{channel}}} \quad , \quad (7.1)$$

where $\varepsilon_{L0-\text{max}}^{\text{channel}}$ is the trigger efficiency for a channel when the full bandwidth is available, and $\varepsilon_{L0}^{\text{channel}}$ is obtained using one fixed set of thresholds for all channels simultaneously. Thresholds are given in Table 7.2, with their corresponding inclusive rates on

Table 7.2: List of cuts obtained after the combined Level-0 optimization. The last column gives the inclusive L0 output rate on minimum-bias events after the global event cuts.

E_T thresholds	Value (GeV)	M. B. rate (kHz)
hadron	3.6	705
electron	2.8	103
photon	2.6	126
π^0 local	4.5	110
π^0 global	4.0	145
muon	1.1	110
$\sum p_T^\mu$	1.3	145

minimum-bias events after the global event cuts described above. The thresholds of the ECAL triggers are highly correlated since these triggers are to a large extent redundant.

An event triggers Level-0 if (1) it passes the global selection and if at least one of the candidates' E_T exceeds its threshold, or (2) if the sum of the transverse momenta of the two largest p_T muons ($\sum p_T^\mu$) is above its threshold, irrespective of the global cuts. The FOIs of the Muon Trigger (see Chapter 3) have been optimized correspondingly, and their values are listed in Table 7.3.

Table 7.3: The size of the FOI along the x and y coordinates as used in the Muon Trigger for the thresholds listed in Table 7.2.

	M1	M2	M4	M5
x	± 2	± 5	± 1	± 1
y	0	0	± 1	± 1

Figure 7.2 shows the sensitivity of ε_{L0} to the Level-0 output rate. The values of ε_{L0} from 0.5–1.0 MHz were obtained after a combined optimization of Level-0 as described above for each output rate. The values of ε_{L0-max} are shown in the right-hand part of the figure, indicating how much is lost in efficiency per channel while sharing the bandwidth over all channels in a democratic way. This bandwidth division results

in small losses, apart from channels with the decay $J/\psi \rightarrow ee$, since in the bandwidth division optimization it is combined with $J/\psi \rightarrow \mu\mu$ for calculating $\varepsilon_{L0}^{\text{channel}}$.




Figure 7.2: Level-0 efficiencies as a function of the Level-0 output rate. The rightmost set of data points refers to the efficiency obtained after individual optimization of each channel.

The Level-0 efficiencies for the set of channels used to tune the thresholds are given in Table 7.4, while Table 7.5 gives the efficiencies for channels which did not participate in the optimization. We also include the inclusive hadronic, electromagnetic (electron, photon and π^0 s) and muon trigger, to show the correlation between them. The contributions from the ECAL triggers have been grouped together, since they are to a large extent redundant and their relative contributions depend on the choice of their highly correlated thresholds.

PYTHIA predicts a $b\bar{b}$ cross section of $633 \mu\text{b}$, which results in 1.1% of the crossings with at least one inelastic pp-interaction containing at least one $b\bar{b}$ -pair. A $c\bar{c}$ -pair is produced in 5.6% of the crossings, where $c\bar{c}$ -pairs are only counted if no $b\bar{b}$ -pair is present. The beauty enrichment of the data after the various triggers is listed

Table 7.4: L0 efficiencies at 1 MHz for several offline selected signal channels which have been used to determine the thresholds. The last three columns show the inclusive trigger efficiencies for the hadronic, electromagnetic (electron, photon, π^0 's) and muon triggers.

Decay Channel	$\varepsilon_{L0}(\%)$	Inclusive efficiencies (%)		
		had. trig.	elec. trig.	muon trig.
$B_d^0 \rightarrow \pi^+\pi^-$	53.6 ± 0.4	47.6 ± 0.5	14.1 ± 0.3	6.8 ± 0.2
$B_s^0 \rightarrow D_s^-(K^+K^-\pi^-)\pi^+$	49.4 ± 0.6	42.2 ± 0.6	13.1 ± 0.4	8.3 ± 0.4
$B_s^0 \rightarrow D_s^-(K^+K^-\pi^-)K^+$	47.2 ± 0.3	39.4 ± 0.3	11.7 ± 0.2	8.2 ± 0.2
$B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K_S^0(\pi^+\pi^-)$	89.3 ± 0.5	18.6 ± 0.7	8.3 ± 0.5	87.2 ± 0.6
$B_d^0 \rightarrow J/\psi(e^+e^-)K_S^0(\pi^+\pi^-)$	48.3 ± 1.0	21.5 ± 0.8	37.4 ± 0.9	7.0 ± 0.5
$B_s^0 \rightarrow J/\psi(\mu^+\mu^-)\phi(K^+K^-)$	89.7 ± 0.1	20.0 ± 0.2	8.4 ± 0.1	87.4 ± 0.1
$B_d^0 \rightarrow K^{*0}(K^+\pi^-)\gamma$	72.9 ± 1.0	32.7 ± 1.1	68.1 ± 1.1	7.8 ± 0.6

Table 7.5: L0 efficiencies at 1 MHz for offline selected signal channels which have not been used to tune the thresholds. The last three columns show the inclusive trigger efficiencies for the hadronic, electromagnetic (electron, photon, π^0 's) and muon triggers.

Decay Channel	$\varepsilon_{L0}(\%)$	Inclusive efficiencies (%)		
		had. trig.	elec. trig.	muon trig.
$B_d^0 \rightarrow K^+\pi^-$	54.1 ± 0.8	48.3 ± 0.8	12.3 ± 0.5	7.2 ± 0.4
$B_s^0 \rightarrow K^-\pi^+$	56.5 ± 1.1	51.2 ± 1.1	13.2 ± 0.7	6.7 ± 0.5
$B_s^0 \rightarrow K^+K^-$	51.8 ± 0.3	46.0 ± 0.3	11.6 ± 0.2	6.5 ± 0.2
$B_d^0 \rightarrow \pi^+\pi^-\pi^0$	77.2 ± 1.6	39.4 ± 1.9	66.2 ± 1.8	7.9 ± 1.1
$B_d^0 \rightarrow D^{*-}(\bar{D}^0\pi^-)\pi^+$	49.0 ± 1.1	41.7 ± 1.1	14.0 ± 0.8	8.4 ± 0.6
$B_d^0 \rightarrow \bar{D}^0(K^+\pi^-)K^{*0}(K^+\pi^-)$	53.0 ± 1.4	45.3 ± 1.4	13.9 ± 0.9	8.1 ± 0.7
$B_d^0 \rightarrow \bar{D}^0(K^+K^-)K^{*0}(K^+\pi^-)$	50.7 ± 1.2	43.4 ± 1.1	13.6 ± 0.8	8.4 ± 0.6
$B_d^0 \rightarrow J/\psi(\mu^+\mu^-)K^{*0}(K^+\pi^-)$	91.2 ± 0.3	23.1 ± 0.4	9.3 ± 0.3	88.6 ± 0.3
$B_u^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$	90.3 ± 0.4	26.2 ± 0.5	9.1 ± 0.4	87.1 ± 0.4
$B_s^0 \rightarrow J/\psi(e^+e^-)\phi(K^+K^-)$	49.0 ± 0.6	22.9 ± 0.5	38.3 ± 0.5	7.0 ± 0.3
$B_s^0 \rightarrow J/\psi(\mu^+\mu^-)\eta(\gamma\gamma)$	92.1 ± 0.8	19.2 ± 1.2	37.2 ± 1.5	88.4 ± 1.0
$B_s^0 \rightarrow \eta_c(4\pi, 2K2\pi)\phi(K^+K^-)$	47.0 ± 3.0	41.5 ± 2.9	12.5 ± 1.9	8.0 ± 1.6
$B_s^0 \rightarrow \phi(K^+K^-)\phi(K^+K^-)$	41.8 ± 0.9	28.7 ± 0.9	9.7 ± 0.6	8.6 ± 0.5
$B_d^0 \rightarrow \mu^+\mu^-K^{*0}(K^+\pi^-)$	93.6 ± 0.7	24.9 ± 1.2	10.3 ± 0.8	91.8 ± 0.7
$B_s^0 \rightarrow \phi(K^+K^-)\gamma$	69.6 ± 1.6	33.1 ± 1.6	65.8 ± 1.7	7.7 ± 0.9
$B_c^+ \rightarrow J/\psi(\mu^+\mu^-)\pi^+$	92.6 ± 0.5	29.4 ± 0.8	9.9 ± 0.5	89.5 ± 0.6

in Table 7.6.

The bandwidth division described above should be regarded as an example only, another possibility is to optimize on effective tagging efficiency [70] rather than just considering Level-0 efficiency as used above.

7.2 Performance of Level-1

The principal idea of Level-1 is to combine the two most characteristic properties of b tracks available at this early stage, impact parameter and transverse momentum, to form an efficient selection of events

Table 7.6: The fraction of crossings with at least one inelastic pp-interaction containing at least one $b\bar{b}$ -pair, or $c\bar{c}$ -pair. Where $c\bar{c}$ -pairs are only counted if no $b\bar{b}$ -pairs are present.

	bb %	c \bar{c} %
Generated	1.1	5.6
Level-0	3.0	10.6
Level-1	9.7	14.2
HLT (L1-confirmation)	14.0	14.7

containing b-hadrons. In addition, use is made from the signatures already found by Level-0 and passed on to Level-1.

The Level-1 decision algorithm [71] con-




Figure 7.3: Distribution of offline-selected signal events and minimum-bias events in the plane of the two variables $\ln(\text{PT1})+\ln(\text{PT2})$ versus $\ln(\text{IPS1})+\ln(\text{IPS2})$. The solid line is an example of the vertical-diagonal discriminant applied to determine the Level-1 trigger variable.

sists of two parts: in the first, generic algorithm, a trigger variable is computed based on the properties of the two tracks with highest transverse momentum p_T . This part is sensitive to very generic b-hadron signatures. In the second, specific algorithm, the trigger variable is weighted according to signatures involving L0 objects, such as dimuons or high- E_T electrons and photons, that are present in the event. This means that good L0 signatures have the effect of relaxing the generic requirement.

7.2.1 Generic Algorithm

An average (most probable) number of 8.5 (4) tracks per minimum-bias event are reconstructed in 3D (see Section 6.2). Requiring the 3D impact parameter to be between 0.15 and 3 mm reduces this number to 6.5 (4). Using this set of tracks, two event variables are defined as B-signatures, $\ln(\text{PT1})+\ln(\text{PT2})$ and $\ln(\text{IPS1})+\ln(\text{IPS2})$, where PT1(2) is the p_T of the 3D track with the highest (second-highest) p_T , and IPS1(2) are their respective impact parameter significances with respect to the primary vertex, defined as d/σ_d [72]. The error σ_d is estimated based on the p_T of the track. Figure 7.3 shows the distribution of minimum-bias events and two types of signal events

in these two variables. Also shown is the vertical-diagonal discriminant line that is used to determine the trigger variable Δ : the distance of an event to this line, with negative sign if the event lies to its left. The choice of the diagonal discriminant line was originally motivated by an earlier version of the L1-algorithm which had poorer resolution in impact parameter and therefore a surplus of high- p_T tracks originating from the primary vertex, i.e. at low impact parameters. We decided to keep this feature in order to guard against possible degradation of the tracking performance in the robustness tests.

7.2.2 Specific Algorithm

The following event variables are considered for relaxing the generic trigger condition:

- $m_{\mu\mu}^{\max}$ – The highest invariant dimuon mass, where dimuons consist of two VELO tracks that have been matched in 3D to L0 muon tracks, without requirements on the charges nor the vertex of the two tracks. This variable is sensitive to channels involving $J/\psi \rightarrow \mu^+\mu^-$ or $B \rightarrow \mu^+\mu^-(X)$ (see Figure 7.4).
- $E_T^{\gamma, \max}$ – The highest photon transverse energy found by Level-0, if above

3 GeV. Sensitive to channels such as $B \rightarrow K^* \gamma$.

- $E_T^{e,\max}$ – The highest electron transverse energy found by Level-0, if above 3 GeV. Sensitive to channels involving $J/\psi \rightarrow e^+e^-$.

In each case, a “bonus” β is calculated depending on the variable and added to the generic trigger variable.




Figure 7.4: Dimuon invariant mass distribution for several signal channels in comparison with minimum-bias events.

For dimuons, $\beta_{\mu\mu}$ is set to a very high value (i.e. completely overwriting the decision of the generic trigger) if $m_{\mu\mu}^{\max}$ lies either within ± 500 MeV/ c windows around the J/ψ and the B mass, or above the latter. For other values $\beta_{\mu\mu}$ increases linearly with $m_{\mu\mu}^{\max}$, where the linear coefficient is chosen as a compromise between enhancing the efficiency for $B \rightarrow \mu^+ \mu^- K^*$ and minimizing the used bandwidth.

In a similar way, β_γ and β_e are computed as functions of $E_T^{\gamma,\max}$ and $E_T^{e,\max}$, respectively. A minimum of 3 GeV is required,

and β increases linearly starting from this value. In the case of photons, both $E_T^{\gamma,\max}$ and $E_T^{e,\max}$ are required to be above 3 GeV for $\beta_\gamma > 0$, whereas only $E_T^{e,\max}$ is used to compute the value of β_e . Again the linear coefficients have been chosen to minimize bandwidth use while giving significant improvements in the efficiencies for channels containing photons and electrons.

7.2.3 Final decision

An event passes Level-1 if

$$\Delta + \beta_{\mu\mu} + \beta_\gamma + \beta_e > \Delta_{0.04}, \quad (7.2)$$

where $\Delta_{0.04}$ is determined empirically by requiring a minimum-bias retention rate of 4% of all L0-triggered events, corresponding to an output rate of 40 kHz.

7.2.4 Efficiencies and bandwidth division

Figure 7.5 shows the trigger efficiencies of a few selected signal channels as a function of the Level-1 output rate. Table 7.7 summarizes signal efficiencies for the design output rate of 40 kHz. Figure 7.6 illustrates the bandwidth division between generic and specific sub-triggers. The specific improvements use up 24.7% of the bandwidth. The current trigger balance should be regarded as an example only – by adjusting the various parameters the division among the signatures can easily be adapted to the needs of the experiment.

The $b\bar{b}$ enrichment after Level-1 is listed in Table 7.6.

7.3 Performance of the High Level Trigger

The HLT can access the same data as used for the offline selection, but at its input rate of 40 kHz. The reconstruction and selection algorithms which will be employed are




Figure 7.5: L1 efficiencies as a function of the L1 output rate. The last bin refers to the maximum efficiency obtained after individual optimization of each channel. The efficiencies are normalized to L0-triggered events that have been selected by the offline analysis. Indicated errors are statistical.




Figure 7.6: Bandwidth division among the various trigger components as explained in the text.

constrained by the available computing resources of a few hundred CPU nodes. The HLT algorithm is under development, and the following strategy guarantees a high selection efficiency, and is estimated to give an affordable execution time:

- A Confirm the Level-1 algorithm, but using T1–T3 to improve on the momentum resolution compared to the VELO-TT tracks of Level-1.
- B Full pattern recognition of long tracks, and lepton identification.
- C Exclusive selection of channels.

Table 7.7: L1 efficiencies at 40 kHz output rate for several signal channels. The efficiencies are normalized to L0-triggered events that are used for offline analysis.

Decay channel	$\epsilon_{L1}(\%)$
$B_d^0 \rightarrow \pi^+ \pi^-$	62.7 ± 0.5
$B_d^0 \rightarrow K^+ \pi^-$	61.5 ± 1.0
$B_d^0 \rightarrow K^- \pi^+$	65.0 ± 1.4
$B_s^0 \rightarrow K^+ K^-$	60.0 ± 0.4
$B_d^0 \rightarrow \pi^+ \pi^- \pi^0$	46.6 ± 2.2
$B_d^0 \rightarrow D^{*-}(\bar{D}^0 \pi^-) \pi^+$	56.0 ± 1.6
$B_d^0 \rightarrow \bar{D}^0(K^+ \pi^-) K^{*0}(K^+ \pi^-)$	66.7 ± 1.8
$B_d^0 \rightarrow \bar{D}^0(K^+ K^-) K^{*0}(K^+ \pi^-)$	61.6 ± 1.6
$B_s^0 \rightarrow D_s^-(K^+ K^- \pi^-) \pi^+$	63.0 ± 0.9
$B_s^0 \rightarrow D_s^-(K^+ K^- \pi^-) K^+$	62.6 ± 0.4
$B_d^0 \rightarrow J/\psi(\mu^+ \mu^-) K_S^0(\pi^+ \pi^-)$	67.7 ± 0.9
$B_d^0 \rightarrow J/\psi(e^+ e^-) K_S^0(\pi^+ \pi^-)$	54.9 ± 1.4
$B_d^0 \rightarrow J/\psi(\mu^+ \mu^-) K^{*0}(K^+ \pi^-)$	76.8 ± 0.3
$B_u^+ \rightarrow J/\psi(\mu^+ \mu^-) K^+$	76.0 ± 0.5
$B_s^0 \rightarrow J/\psi(\mu^+ \mu^-) \phi(K^+ K^-)$	71.4 ± 0.2
$B_s^0 \rightarrow J/\psi(e^+ e^-) \phi(K^+ K^-)$	57.2 ± 0.8
$B_s^0 \rightarrow J/\psi(\mu^+ \mu^-) \eta(\gamma\gamma)$	70.3 ± 1.5
$B_s^0 \rightarrow \eta_c(4\pi, 2K2\pi) \phi(K^+ K^-)$	59.0 ± 4.0
$B_s^0 \rightarrow \phi(K^+ K^-) \phi(K^+ K^-)$	60.3 ± 1.5
$B_d^0 \rightarrow \mu^+ \mu^- K^{*0}(K^+ \pi^-)$	78.5 ± 1.1
$B_d^0 \rightarrow K^{*0}(K^+ \pi^-) \gamma$	51.9 ± 1.4
$B_s^0 \rightarrow \phi(K^+ K^-) \gamma$	49.3 ± 2.0
$B_c^+ \rightarrow J/\psi(\mu^+ \mu^-) \pi^+$	65.6 ± 0.9

D Inclusive selection of channels.

The purpose of step A is to reject events as soon as possible with an algorithm which is cheap in execution time. With this algorithm the rate is reduced to 20 kHz, and its expected performance will be described in the next section. For the remaining events after A, full pattern recognition is performed, including lepton identification, but not using the RICH information. The tracking part of step B will require most of the execution time, and is described already in Section 6.3. The average number of tracks reconstructed per event with both VELO and T information is 32.

Selection C aims at getting the highest possible efficiency for those channels which are considered as the most important for the physics goals. Below it will be shown for a few selected channels that low

enough output rates can be achieved, even without the particle identification using the RICHes.

The $b\bar{b}$ -pair content of the 20 kHz of events which have to be analysed after step A by the HLT is given in Table 7.6, and shows that the sample is still dominated by light quark events. In step D commonalities in the offline selection of B-decay channels with similar kinematics are used to define a series of selections which should result in high efficiencies for all interesting channels, without having to resort to an exclusive reconstruction per channel, and hence should guarantee that LHCb will write with high efficiency a large spectrum of B-like events to storage for subsequent analysis. This last selection is least developed, and hence no results are presented on its expected output rate here, but it will be reported as part of the Computing TDR.

7.3.1 Level-1 Confirmation

Section 6.3 describes the pattern recognition algorithm and its performance. While the VELO-TT tracks of Level-1 have a momentum resolution between 20–30%, the HLT reconstructs tracks with a $\sigma(p)/p$ around 0.6%. With the HLT tracks the L1-algorithm is repeated as described in Section 7.2, i.e among the about eight tracks with an impact parameter between 0.15–3 mm the momentum is measured as described in Section 6.3, the two tracks with the largest p_T are selected, the trigger variable Δ is computed, and a bonus is added depending on $m_{\mu\mu}^{\max}$, $E_T^{\gamma,\max}$ or $E_T^{e,\max}$. Figure 7.7 shows the trigger efficiencies of a few selected signal channels as a function of trigger rate. This shows that the L1-confirmation reduces the rate from 40 kHz to 20 kHz with only a few % loss in signal efficiency. To achieve this reduction not all long tracks have to be reconstructed, hence reducing the necessary execution time. Unlike the algorithm deployed in Level-1, the




Figure 7.7: L1-confirmation efficiencies for events accepted L0, L1 and the offline analysis for a few selected signal channels as a function of minimum-bias rate.

HLT algorithm has not yet been packaged to allow an execution time measurement of the L1-confirmation step separately, however based on the total track reconstruction it is estimated that this step will take less than 4 ms in 2007, hence leaving around 14 ms per event for the further analysis of the 20 kHz of accepted events. Out of this about 6 ms will be needed to reconstruct all remaining tracks.

7.3.2 Exclusive Selection

The combined output rate of all channels under study at the moment in LHCb [2], including the expected background, is less than 1 Hz of data taking rate. However, the exclusive selection in the HLT will have a significantly larger output rate due to the need to relax the final selection cuts to be able to study the sensitivity and systematics. In particular side bands in invariant mass distributions are necessary to fit the contribution of the background. In addition, the global RICH reconstruction benefits from reconstructing all track segments around the RICH detectors, and con-

Table 7.8: Expected rate in the HLT in minimum-bias events due to the exclusive selection algorithms aimed at selecting the indicated channels. All errors are statistical only.

Algorithm	HLT rate (Hz)
$B \rightarrow \psi(\mu^+\mu^-)X$	21 ± 4
$B \rightarrow h^+h^-$	12 ± 3
$B \rightarrow D_s h$	12 ± 3
$B_d^0 \rightarrow K^{*0}(K^+\pi^-)\gamma$	13 ± 3
$B_d^0 \rightarrow \phi(K^+K^-)\gamma$	14 ± 3

sequently its execution time only allows it to be executed at a rate of a few hundred Hz. The exclusive selection power has been checked by passing generated minimum-bias events through the first two trigger levels, and then applying relaxed offline selection cuts [73, 74, 75, 76]. Table 7.8 lists the expected rates for the channels considered. The $\psi(\mu^+\mu^-)$ rate contains 5 Hz of genuine $B \rightarrow \psi X$ decays. This shows that the exclusive selection can reduce the rate to well below a few hundred Hz without having to resort to RICH information.

7.4 Trigger Performance Robustness

The expected trigger performance depends on the imperfections of our simulation program: the underlying physics process may be slightly different in real data, the description of the detector and its performance may not be perfect, the LHC machine conditions may be different than expected, and so on.

The LHCb trigger system is designed to be flexible to adapt to such unexpected situations. In order to illustrate this point, several robustness scenarios have been fully simulated, and the trigger execution time and performance have been measured on these samples. The scenarios considered are described in detail in [64] and they correspond to:

- **PYTHIA Test:** PYTHIA settings (extrapolated to the LHC energy) from a recent tuning on CDF data [65].
- **Global Test:** General LHCb degraded detector and worse PYTHIA settings (as in [2]).
- **VELO Test:** Degraded VELO detector, with increased material, worse cluster resolution and worse signal-over-noise ratio (as in [64]).
- **Beam Spot Test:** LHC beam position uncertainty increased by a factor three in the perpendicular plane, from $70 \mu\text{m}$ to $210 \mu\text{m}$.
- **LHC Background Test:** LHC machine backgrounds worse by an order of magnitude (as in [30]).

The uncertainty on the simulation of the underlying physics process is taken into account by the changes in the PYTHIA settings in the ‘‘PYTHIA Test’’ and ‘‘Global Test’’. The most visible effect is the change in the track multiplicity, which turns out to be the most relevant variable in the evaluation of the trigger performance. The worse VELO detector resolutions intends to simulate the effect of possible misalignments. The ‘‘Beam Spot Test’’ checks the dependence of the trigger performance with the position of the beam.

The effect of the LHC machine background is found to be negligible. In nominal conditions, the average number of beam halo muons traversing the detector in any direction is 0.04 per bunch. Even with an order of magnitude worse conditions, the event rate increases after Level-0 and Level-1 by only 12 ± 3 kHz. The most significant effect is seen in the Level-0 muon system and has been described in Chapter 3.

In the next two sections, the robustness of the algorithm used in Level-1 is determined through the effect in the resolution of the track parameters, and through the changes in the execution time of the algorithms and event size. The overall effect




Figure 7.8: Impact parameter resolution (in mm) as measured by the Level-1 algorithm as a function of the transverse momentum of the track, for several MC simulations. The impact parameter is computed with respect to the reconstructed primary vertex.

on the trigger decision will be discussed in Section 7.4.3.

7.4.1 Resolutions

The impact parameter and the determination of the p_T of the tracks used in the trigger algorithms depend on the performance of the VELO detector and the ability to reconstruct the primary vertex. The “Global Test” and the “VELO Test” both degrade the measurement of the track parameters in the VELO detector, while the “Beam Spot Test” degrades the ability to reconstruct the primary vertex. The resolution of the impact parameter with respect to the reconstructed primary vertex, as measured by the Level-1 algorithm, is shown in Figure 7.8 as a function of p_T . The resolution is worse by 30% when the material of the VELO detector is increased, in particular the RF foil. The resolution of p_T as measured by the Level-1 algorithm is shown in Figure 7.9 as a function of p_T . In the case




Figure 7.9: $\sigma(p_T)/p_T$ as measured by the L1 algorithm as a function of the transverse momentum of the track, for several MC simulations.

of Level-1 the small changes in the p_T resolution have a negligible effect.

7.4.2 Execution Time and Event Size

The execution time of the Level-1 algorithm is dominated by the time it takes to perform the track reconstruction. The only significant changes are observed when the expected multiplicity of the event is different, i.e. in the “Global Test” and the “PYTHIA Test”. The event multiplicity is increased by 30% in the “Global Test” while it is reduced by 20% using the CDF tuning. The measured execution time and event size for the Level-1 follow the same pattern.

7.4.3 Performance

The output of Level-0 is fixed to 1.0 MHz, while the output of the Level-1 trigger is fixed to 40 kHz. Hence, the cuts used in the trigger algorithms need to be adapted for each robustness scenario. One convenient way to do this for Level-0 is just to change

Table 7.9: Level-0/Level-1 efficiencies for different robustness scenarios divided by the efficiencies using the default settings.

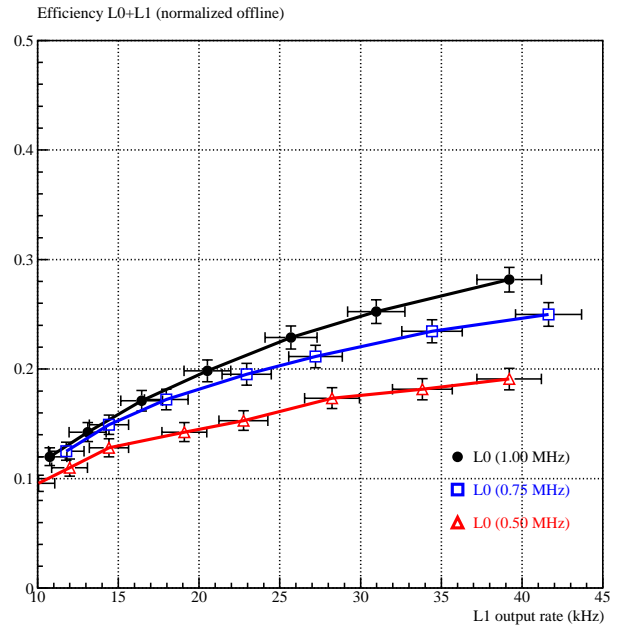
	$B_s^0 \rightarrow D_s^+ K^\pm$	$B_s^0 \rightarrow J/\psi(\mu\mu)\phi$	$B_s^0 \rightarrow K^+ K^-$	$B^0 \rightarrow K^{*0} \gamma$
L0 Global Test	0.93 ± 0.02	1.01 ± 0.01	0.95 ± 0.01	0.93 ± 0.02
L0 PYTHIA Test	1.06 ± 0.02	1.01 ± 0.01	1.11 ± 0.01	1.03 ± 0.02
L0 Beam Spot Test	0.98 ± 0.02	1.01 ± 0.01	1.00 ± 0.01	0.96 ± 0.02
L0 VELO Test	1.00 ± 0.02	1.01 ± 0.01	1.01 ± 0.01	1.00 ± 0.02
L1 Global Test	0.84 ± 0.03	0.92 ± 0.01	0.87 ± 0.01	0.82 ± 0.03
L1 PYTHIA Test	0.97 ± 0.03	0.99 ± 0.01	0.97 ± 0.01	0.99 ± 0.03
L1 Beam Spot Test	0.90 ± 0.03	1.00 ± 0.01	0.94 ± 0.01	0.97 ± 0.03
L1 VELO Test	0.92 ± 0.03	0.96 ± 0.01	0.89 ± 0.01	0.89 ± 0.03

the SPD/PU multiplicity cuts quoted in Table 7.1 while keeping the rest of the cuts unchanged. The cut on the Level-1 variable, $\Delta_{0.04}$, described in Section 7.2 also needs to be modified to keep the 40 kHz output rate.

We have not tried to modify the algorithms to adapt to each robustness scenario, but rather we quote the efficiencies as they come out from the same algorithms, which may be regarded as an estimate of the order of magnitude of the uncertainties in the trigger performance. The results are quoted in Table 7.9. In general, the Level-0 performance is stable within 10%, while the Level-1 performance is stable within 20%.

Independently of the uncertainties in the simulation, we have also considered the scenario in which on “day one” we do not have the full CPU power, and we start running the experiment with part of the event building network, and part of the nominal number of CPUs assigned to L1/HLT. The trigger efficiency for $B_s^0 \rightarrow D_s^+ K^\pm$ events, normalized to offline selected events, is shown as a function of different Level-0 and Level-1 output rates in Figure 7.10.

In conclusion, the trigger performance satisfies the physics requirements of the experiment. It is also a robust system. The uncertainty on the expected efficiency is evaluated to be not larger than 25% even for the most pessimistic scenarios considered.

Figure 7.10: Trigger efficiency for $B_s^0 \rightarrow D_s^+ K^\pm$ events, normalized to offline selected events, as a function of different Level-0/Level-1 output rates.

Chapter 8 Project Organization

This chapter deals with the managerial aspects of the trigger systems. Information is presented on the current cost estimates of the systems, the planning schedules and the distribution of the responsibilities.

8.1 Cost and Funding

The breakdown of the cost of the trigger is shown in Table 8.1.

For both the Calorimeter Triggers and the Pile-Up System infrastructure is shared between the trigger components and their respective detectors. The cost listed excludes those items already accounted for in their detector TDRs.

For the prices of commercial hardware like PCs and switches, used in the implementation of Level-1 and HLT, the numbers projected by the technology tracking group at CERN have been used [61, 77].

In the Online System TDR [9] the cost of the part of the Online System which excludes TFC, ECS and general infrastructure was 4,244 kCHF. This value is now superseded by 5,711 kCHF of the combined Level-1&HLT system as indicated in Table 8.1.

8.2 Schedule

The detailed project schedules of the sub-systems are shown in Figures 8.1. All schedules assume that the commissioning of individual sub-systems has to be finished by September 2006, after which date

LHCb starts the overall commissioning of the spectrometer to be ready for beam in the beginning of 2007.

Some sub-systems contain several different boards, in which case the milestones refer to the total number of boards, irrespective of their type.

The L1&HLT System will have enough functionality in 2006 to allow a test of the whole detector, but not at the design bandwidth. The bulk of the CPUs will be acquired and installed as late as possible. Table 8.2 contains a set of milestones extracted from the project schedules.

8.3 Division of Responsibilities

Institutes currently involved in the LHCb trigger are Annecy, Barcelona, Bologna, CERN, Clermont-Ferrand, Krakow, Lausanne, NIKHEF, Marseille, Orsay and Rio de Janeiro. The sharing of responsibilities is listed in Table 8.3.

Table 8.1: Components (needed + spares) and cost of the trigger.

Item	Quantity	Total cost [kCHF]
Calorimeter Triggers		
CaloFE Card, 9U (trigger part)	238+31	
PrsFE Card, 9U (trigger part)	100+10	
Validation Card, 9U	28+3	
SPD Multiplicity Card (part), 9U	16+2	
Crates (part) and Backplane	26+4	
Optical links	208+30	
Selection Cards, 9U	8+2	
Selection Crate	1+1	
Readout Card, 9U	1+1	
Total		950
Muon Trigger		
PU board, 9U	60+6	
Muon Selection board 9U	12+2	
Controller board, 9U	4+2	
Backplane	4+2	
Crates	4+2	
Short distance OL ribbon	120+12	
Total		1000
Pile-Up System		
Hybrids	4+1	
Cables		
Repeater station	1	
Optical ribbon connection	8+1	
Processing system	1	
Sensors	4+2	
SC/HV/LV		
Crates	1	
TELL1 board	5	
Total		420
L0 Decision Unit		
TELL1 board	1+2	
Mezzanine Cards	1+2	
Total		60
Level-1 and the HLT DAQ		
Mux switches L1	62 + 7	
Mux switches HLT	29 + 3	
Sub-farm switches	94 + 10	
Readout network ports	190 + 19	
SFCs	94 + 10	
CPUs	1786 + 188 ^a	
TRM	1 + 1	
Decision sorter	1 + 1	
Total		5711
Total		8141

^aAll CPUs available will be active in the system anytime. The spares can be seen as “hot spares”




Figure 8.1: Project schedule of the trigger.

Table 8.2: List of major milestones.

Milestone	Date
Calorimeter Triggers	
Start of board production	9/2004
50% of boards produced	10/2005
100% of boards produced and tested	6/2006
Muon Trigger	
Start of board production	8/2005
25% of boards produced	2/2006
100% of boards produced and tested	8/2006
Pile-Up System	
Detector produced	1/2006
Trigger boards produced and tested	5/2006
Level-0 Decision Unit	
Boards produced and tested	6/2006
Level-1 & HLT	
Sorter Implementation Decision	05/2004
Event Builder Switch specification	06/2004
Subfarm Controller specification	01/2005
10(5)% of network(farm) ready	07/2005
100% of network ready	10/2006
Full farm installed	3/2007
Trigger installed	9/2006

Table 8.3: Breakdown of the responsibilities of the Trigger Systems.

System	Institute
Level-0 Calorimeter Triggers	
ECAL/HCAL FE + coordination	Orsay
PreShower FE	Clermont-Ferrand
SPD multiplicity board	Barcelona
Validation Card	Annecy
Selection Crate	Bologna
Level-0 Muon Trigger	Marseille
Level-0 Pile-Up System	NIKHEF
Level-0 Decision Unit	Clermont-Ferrand
Level-1 & HLT	
Data movement upstream of farm and event-building	CERN/Lausanne
Gigabit Ethernet Mezzanine Card	CERN
Installation, configuration and procurement of hardware	CERN
Implementation of sub-farm	Bologna
Simulation of network and switch specification	Bologna/CERN
Online adaption of offline software framework	Rio de Janeiro
L1&HLT algorithms	CERN/Krakow/Lausanne

Appendix A Scalability of Level-1

The total size of the readout network for Level-1 and the HLT has been based on simulation. As has been shown in Section 7.4, with different conditions one would alter the data size and the performance of Level-1. Hence, the DAQ is required to be flexible enough to be able to adapt quickly to actual needs.

In this section an extension of the data made available to Level-1 is discussed as an example of what this scalability encompasses. The additional data comprise Level-1 data from the tracking detectors T1 to T3 and the Muon detector stations M2 to M5, and its size is given in Table A.1.

Table A.1: Number of L1 data sources and average event fragment size per source, which does not include any transport overheads. IT, OT and M2–M5 are the extra sources added to L1.

Subsystem	Number of sources	Data/source [Bytes]
VELO	76	36
TT	48	24
L0DU	1	86
Calo Crate	1	70
IT	42	23
OT	48	60
M2–M5	8	54

The numbers given here for the additional components required for the DAQ system illustrate the scalability of the system. The numbers for the system presented in the main text are repeated for comparison.

The system is perfectly scalable to include only part of this maximum extension.

The input parameters remain as shown in Table 6.2.

The following remarks should illustrate how the numbers in Table A.2 are derived:

- The target number of worker nodes remains the same, $\sim 1800^1$.
- The number of SFCs must be increased, to cope with the increased bandwidth at the output of the network. Each sub-farm is still limited to 100 MB/s (c.f. Table 6.2).
- The number of worker nodes is required to be identical for each sub-farm. Thus the actual total number of worker-nodes and consequently the event-rate per CPU node varies slightly compared to Table 6.3.

The figures for the extended system are indicated in the architecture shown in Figure A.1. Table A.3 lists the number of components required, if data for IT, OT and Muon stations M2 to M5 would be included in Level-1. It illustrates that the system scales quite nicely, taking into account that the event size almost doubles as shown in A.2.

The cost estimate is based on prices for a system ready in 2007. The additional cost for this enlarged system using the cost figures from [61, 77] is ~ 2100 kCHF. An “upgrade” at a later time will be cheaper.

¹It might seem odd at first, that here with the same number of CPUs a larger Level-1 event is processed for the Level-1-trigger. However, the increased information available to the Level-1 algorithm will improve its rejection power, thus allowing to reduce the Level-1 output rate and consequently the number of CPUs required for HLT and reconstruction.

Table A.2: Key performance figures of the standard and the extended Level-1 and HLT DAQ system

	VELO TT	+ IT OT M2-M5
Event Building		
Total Frame Rate at RN input [kHz]	7248	12990
RN output links	94	175
RN output link rate [MB/s]	47.9	52.8
Frame rate (L1) per link [kHz]	59.9	68.9
Frame rate (HLT) per link [kHz]	20.0	10.7
Total frame rate at RN output [kHz]	79.6	79.9
MEP rate (L1) per link [kHz]	0.47	0.25
MEP rate (HLT) per link [kHz]	0.05	0.03
Total MEP rate	0.53	0.28
Trigger farms		
Sub-farms	94	175
Event rate/sub-farm (L1) [kHz]	11.7	6.3
Event rate/sub-farm (HLT) [kHz]	0.4	0.2
Processors/sub-farm	21	12
Event rate per processor (L1) [kHz]	0.56	0.53
Event rate per processor (HLT) [kHz]	0.02	0.02

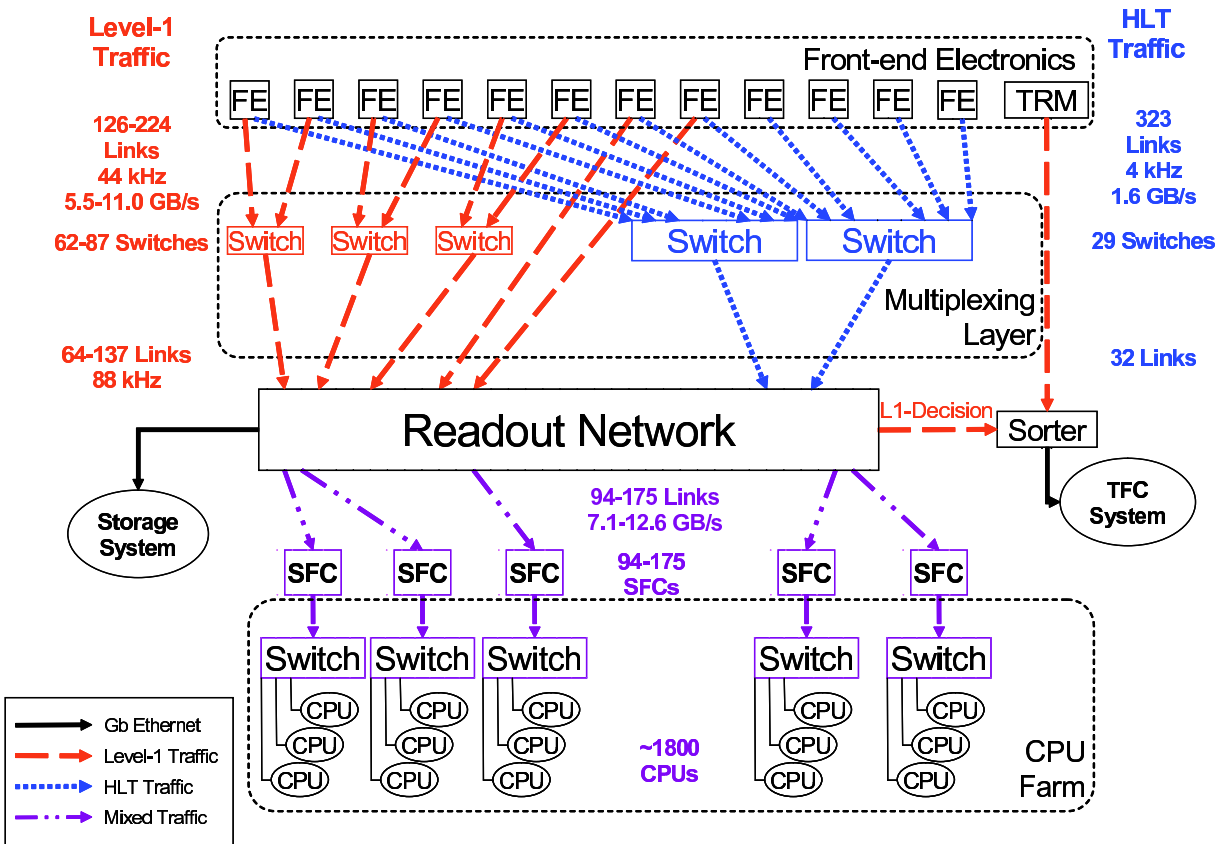


Figure A.1: The architecture of the Level-1 and HLT DAQ system. The numbers on the left side show the increase in scale required by the additional data in Level-1.

Table A.3: Number of items for the hardware implementation of the Level-1 and HLT DAQ system. The number of spares is listed separately.

Item	VELO TT Quantity	+ IT OT M2–M5 Quantity
Mux switches L1	62 + 7	87 + 9
Mux switches HLT	29 + 3	29 + 3
Sub-farm switches	94 + 10	175 + 18
Readout network ports	190 + 19	344 + 35
SFCs	94 + 10	175 + 18
CPUs	1786 + 188	1925 + 175
TRM	1 + 1	1 + 1
Decision sorter	1 + 1	1 + 1

References

- [1] *LHCb Technical Proposal*, LHCb, CERN/LHCC 98–4.
- [2] *LHCb Reoptimized Detector Design and Performance Technical Design Report*, LHCb, CERN LHCC 2003–030.
- [3] *LHCb Magnet Technical Design Report*, LHCb, CERN LHCC 2000–007.
- [4] *LHCb Calorimeter Technical Design Report*, LHCb, CERN LHCC 2000–036.
- [5] *LHCb RICH Technical Design Report*, LHCb, CERN LHCC 2000–037.
- [6] *LHCb Muon System Technical Design Report*, LHCb, CERN LHCC 2001–010.
- [7] *LHCb VELO Technical Design Report*, LHCb, CERN LHCC 2001–011.
- [8] *LHCb Outer Tracker Technical Design Report*, LHCb, CERN LHCC 2001–024.
- [9] *LHCb Online System Technical Design Report*, LHCb, CERN LHCC 2001–040.
- [10] *LHCb Inner Tracker Technical Design Report*, LHCb, CERN LHCC 2002–029.
- [11] *Requirements to the L0 front-end electronics*, J. Christiansen, LHCb 2001–014.
Requirements to the L1 front-end electronics, J. Christiansen, LHCb 2003–078.
- [12] *Timing, Trigger and Control (TTC) Systems for the LHC*, <http://ttc.web.cern.ch/TTC/intro.html>
- [13] *The latency of the Level-0 trigger*, J. Christiansen *et al.*, LHCb 99–015.
- [14] *Simulation of the LHCb front-end*, J. Christiansen and I. Garcia Alfonso, LHCb 99–047.
- [15] *Readout supervisor design specifications*, R. Jacobsson, B. Jost and Z. Guzik, LHCb 2001–012.
- [16] *Common L1 read out board for LHCb specification*, A. Bay *et al.*, LHCb 2003–007.
- [17] *Single Event Effects, Actel AX PGA*, F. Machefert, LHCb 2002–072.
- [18] *LHCb Calorimeter Front-End Electronics, Radiation Dose and Single Event Effects*, C. Beigbeder *et al.*, LHCb 2002–021.
- [19] *The Trigger Part of the Calorimeter Front-End Card*, C. Beigbeder *et al.*, LHCb 2003–037.
- [20] *Functional specification of the PGAs for the ECAL/HCAL Front-End Card*, C. Beigbeder *et al.*, LHCb 2003–036.
- [21] *Front-End Electronics for LHCb Preshower Trigger Part*, G. Bohner *et al.*, LHCb 2003–068.
- [22] *The Validation Card for the Calorimeter Triggers*, C. Drancourt *et al.*, LHCb 2003–120.
- [23] *Implementation and performance of Level-0 trigger system for neutral pions*, O. Deschamps and P. Perret, LHCb 2003–067.
- [24] *The Backplane of the Calorimeter Front-End crate*, C. Beigbeder *et al.*, LHCb 2003–038.
- [25] *The Selection Crate for the L0 Calorimeter Trigger*, G. Balbi *et al.*, LHCb 2003–095.
- [26] *A realistic algorithm for the L0 muon trigger*, E. Aslanides *et al.*, LHCb 2002–042.
- [27] *A synchronous architecture for the L0 muon trigger*, E. Aslanides *et al.*, LHCb 2001–010.
- [28] *Performance of the muon trigger with a realistic simulation*, E. Aslanides *et al.*, LHCb 2002–041.

- [29] *Muon trigger performance with the re-optimized LHCb detector*, E. Aslanides *et al.*, LHCb 2003–074.
- [30] *Machine halo in LHCb for various vacuum conditions*, G. Corti and G. von Holtey, LHCb 2003–086.
- [31] *Specification of the muon trigger processing board*, E. Aslanides *et al.*, LHCb 2002–003.
- [32] *Gigabit Optical Link (GOL) Transmitter* Cern Microelectronics Group, <http://proj-gol.web.cern.ch/proj-gol>
- [33] *Quartz Crystal Phase-Locked Loop (QPLL)*, Cern Microelectronics Group, <http://proj-qpll.web.cern.ch/proj-qpll>
- [34] *High speed ribbon optical link for the L0 muon trigger*, E. Aslanides *et al.*, LHCb 2003–008.
- [35] *Study of the LHCb pile-up trigger and the $B_s \rightarrow J/\psi \phi$ decay*. N. Zaitsev, Thesis, University of Amsterdam, 2000.
- [36] *The LHCb vertex triggers*, N. Zaitsev and L. Wiggers, Nucl.Instrum.Meth. A447, 235–243, 2000.
- [37] *The LHCb Vertex Locator and Level-1 Trigger*, H. Dijkstra, Nucl.Instrum.Meth. A453, 126–130, 2000.
- [38] *R-Sensor sectors and Strip Pitch*, L. Wiggers *et al.*, LHCb 2003–012.
- [39] *Pile-Up System Simulations*, M. Zupan and M. Ferro-Luzzi, LHCb 2003–070.
- [40] *The Beetle Reference Manual*, N. van Bakel *et al.*, LHCb 2001–046.
- [41] *Performance of the Beetle readout chip for LHCb*, N. van Bakel *et al.*, Proc. 8th Workshop on Electronics for LHC Experiments, Colmar, 121–124, 2002.
- [42] *Investigation of the Beetle.1.1 chip in the X7 testbeam*, N. van Bakel *et al.*, LHCb 2002–053.
- [43] *Beetle Comparator Implementation* M van Beuzekom and H. Verkooijen, LHCb 2003–071.
- [44] *The Level 0 Decision Unit for LHCb*, R. Cornat, J. Lecoq and P. Perret, LHCb 2003–065.
- [45] *Using the SPD multiplicity in the Level-0 trigger*, O. Callot, M. Ferro-Luzzi and P. Perret, LHCb 2003–022.
- [46] *Test of the first LODU prototype*, L. Bernard, R. Cornat, J. Lecoq, R. Lefèvre and P. Perret, LHCb 2003–066.
- [47] *GiGabit Ethernet mezzanines for DAQ and Trigger links of LHCb*, H. Muller *et al.*, LHCb 2003–021.
- [48] *Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*, IEEE, IEEE Std 802.3, 2000 ed.
- [49] *Internet Protocol*, IETF, RFC 791 and *Requirements for Internet Hosts*, IETF, RFC 1122.
- [50] *A common implementation of the Level-1 and High Level Trigger Data Acquisition*, A. Barczyk *et al.*, LHCb 2003–079.
- [51] *An integrated experiment control system: the LHCb approach*, C. Gaspar *et al.*, in Proc. IEEE NPSS Real Time Conference 2003.
- [52] *Raw data transport format*, B. Jost and N. Neufeld, LHCb 2003–063.
- [53] *Gaudi Project*, <http://proj-gaudi.web.cern.ch/proj-gaudi>
- [54] *Implementing the L1 trigger path*, R. Jacobsson, LHCb 2003–080.
- [55] *LHCb VELO off detector electronics preprocessor and interface to the Level-1 trigger*, A. Bay *et al.*, LHCb-2001–043.
- L1-type Clustering in the VELO on Test-beam Data and Simulation*, N. Tuning, LHCb-2003–073.

- [56] *The LHCb Level-1 Trigger: Architecture, Prototype, Simulation and Algorithm*, V. Lindenstruth *et al.*, LHCb 2003-064, Chapter 5.
- [57] *Matching VELO tracks to L0 objects*, N. Tuning, LHCb 2003-039.
- [58] *VELO-TT matching and momentum determination at Level-1 trigger*, M. Witek, LHCb 2003-060.
- [59] *The relevance of the magnetic field in the Level-1 trigger*, H. Dijkstra *et al.*, LHCb 2003-110.
- [60] *Execution time optimization of Level-1 algorithm*, M. Witek, LHCb 2003-061.
- [61] *Processors, Memory and Basic Systems*, Working Group A, PASTA 2002 ed. <http://lcg.web.cern.ch/LCG/PEB/PASTAIII/pasta2002Report.htm>
- [62] *VeloTracking for the High Level Trigger*, O. Callot, LHCb 2003-027.
- [63] *The Forward tracking, an optical model method*, M. Benayoun and O. Callot, LHCb 2002-008.
- [64] *Study of the LHCb Trigger Performance Robustness*, F. Teubert, LHCb 2003-059.
- [65] R. Field, private communication, and talks available at http://www.phys.ufl.edu/~rfield/cdf/rdf_talks.html; we have considered the PYTHIA 6.2 settings referred to as “tune A” on http://www.phys.ufl.edu/~rfield/cdf/tunes/rdf_tunes.html.
- [66] *High-Energy-Physics Event Generation with PYTHIA 6.1*, T. Sjöstrand *et al.*, Computer Physics Commun. **135** (2001) 238.
- [67] *Some comments on the running scenario for LHCb*, O. Callot, LHCb 2000-095.
- [68] *GEANT Detector Description and Simulation Tool*, CERN Program Library Long Write-up W5013 (1994).
- [69] *Effect of Multiplicity Cuts on the L0 and L1 Triggers*, M. Ferro-Luzzi *et al.*, LHCb 2003-047.
- [70] *Level-0 Trigger Bandwidth Division*, E. Rodrigues, LHCb 2003-048.
- [71] *Level-1 decision algorithm and bandwidth division*, C. Jacoby and T. Schietinger, LHCb 2003-111.
- [72] *The use of the TT1 tracking station in the level-1 trigger*, H. Dijkstra, *et al.*, LHCb 2002-045.
- [73] *The control channel $B^0 \rightarrow J/\psi(\mu\mu)K^{*0}$* , L. de Paula and E.C. de Oliveira, LHCb 2003-108.
- [74] *Selection of $B_{(s)}^0 \rightarrow h^+h^-$ decays at LHCb*, V. Vagnoni *et al.*, LHCb 2003-123.
- [75] *$B_s^0 \rightarrow D_s^\mp K^\pm$ and $B_s^0 \rightarrow D_s^- \pi^+$ event selection*, A. Golutvin, R. Hierck, J. van Hunen, M. Prodkudin and R. White, LHCb 2003-127.
- [76] *Radiative B decays with LHCb*, I. Belyaev and G. Pakhlova, LHCb 2003-090.
- [77] *Networking Technology*, Working Group D, PASTA 2002 ed. <http://lcg.web.cern.ch/LCG/PEB/PASTAIII/pasta2002Report.htm>