

CERN/LHCC 2006-001
CMS TDR 8.1
02 February 2006

CMS Physics

Technical Design Report

Volume I:

Detector Performance and Software

CMS Software and Physics, Reconstruction and Selection (PRS) Projects		
CMS Spokesperson	Michel Della Negra, CERN	Michel.Della.Negra@cern.ch
CMS Techical Coordinator	Alain Herve, CERN	Alain.Herve@cern.ch
CMS Collaboration Board Chair	Lorenzo Foa, Pisa	Lorenzo.Foa@cern.ch
CPT Project Manager	Paraskevas Sphicas, CERN and Athens	Paraskevas.Sphicas@cern.ch
CPT Detector PRS Coordinator	Darin Acosta, Florida	Darin.Acosta@cern.ch
CPT Analysis PRS Coordinator	Albert De Roeck, CERN	Albert.De.Roeck@cern.ch
CPT Software Coordinator	Lucia Silvestris, Bari	Lucia.Silvestris@cern.ch
CPT Software Coordinator	Avi Yagil, FNAL	Yagil@fnal.gov

Editor

D. Acosta

Chapter Editors

A. De Roeck, U. Gasparini, D. Marlow, N. Neumeister, F. Palla, J. Rohlf, C. Seez, L. Silvestris, Y. Sirois, I. Tomalin, C. Tully, T. S. Virdee, A. Yagil.

Cover Design

S. Cittolin

Acknowledgments

We wish to acknowledge the extraordinary effort from a large fraction of the CMS Collaboration in producing this Technical Design Report on the preparation for physics analysis at the LHC. Subprojects in all areas were involved (Detector, PRS, Software, and Computing) in order to produce the large Monte Carlo simulation samples needed, to develop the software to analyze those samples, to perform the studies reported in this Report, and to write and review our findings.

For their constructive comments and guidance, we would like to thank our CMS internal reviewers.

We also would like to thank our LHCC referees.

For their patience in meeting sometimes impossible demands, we wish to thank the CMS Secretariat: Kirsti Aspola, Madeleine Azeglio, Nadejda Bogolioubova, Dorothée Denise, Dawn Hudson, Guy Martin and Marie-Claude Pelloux.

We also would like to thank George Alverson and Lucas Taylor for their invaluable technical assistance in the preparation of this manuscript.

Finally, we wish to thank the CMS management for their strong support and encouragement.

ISBN 92-9083-268-1
ISBN 978-92-9083-268-3

Trademark notice: all trademarks appearing in this report are acknowledged as such.

Also available at: `http://cmsdoc.cern.ch/cms/cpt/tdr/`

CMS Collaboration

Yerevan Physics Institute, Yerevan, ARMENIA
G.L. Bayatian, S. Chatrchyan, A.M. Sirunyan

Institut für Hochenergiephysik der OeAW, Wien, AUSTRIA
W. Adam, T. Bergauer, M. Dragicevic, J. Erö, M. Friedl, R. Fruehwirth, V. Ghete, P. Glaser, J. Hrubec, M. Jeitler, M. Krammer, I. Magrans, I. Mikulec, W. Mitaroff, T. Noebauer, M. Pernicka, P. Porth, H. Rohringer, J. Strauss, A. Taurok, W. Waltenberger, G. Walzel, E. Widl, C.E. Wulz

Research Institute for Nuclear Problems, Minsk, BELARUS
A. Fedorov, M. Korzhik, O. Missevitch, R. Zuyeuski

National Centre for Particle and High Energy Physics, Minsk, BELARUS
V. Chekhovsky, O. Dvornikov, I. Emeliantchik, A. Litomin, V. Mossolov, N. Shumeiko, A. Solin, R. Stefanovitch, J. Suarez Gonzalez, A. Tikhonov

Byelorussian State University, Minsk, BELARUS
V. Petrov

Vrije Universiteit Brussel, Brussel, BELGIUM
J. D'Hondt, S. De Weirdt, R. Goorens, J. Heyninck, S. Lowette, S. Tavernier, W. Van Doninck[**1], L. Van Lancker

Université Libre de Bruxelles, Bruxelles, BELGIUM
B. Clerbaux, G. De Lentdecker, J.P. Dewulf, T. Mahmoud, P. Marage, L. Neukermans, V. Sundararajan, C. Vander Velde, P. Vanlaer, J. Wickens

Université Catholique de Louvain, Louvain-la-Neuve, BELGIUM
S. Assouak, J.L. Bonnet, G. Bruno, B. De Callatay, J. De Favereau De Jeneret, S. De Visscher, C. Delaere, P. Demin, D. Favart, E. Feltrin, E. Forton, G. Grégoire, T. Keutgen, G. Leibenguth, V. Lemaitre, Y. Liu, D. Michotte, O. Militaru, A. Ninane, S. Ovyn, T. Pierzchala, K. Piotrzkowski, V. Roberfroid, X. Rouby, D. Teyssier, O. Van der Aa, M. Vander Donckt

Université de Mons-Hainaut, Mons, BELGIUM
E. Daubie, P. Herquet, A. Mollet, A. Romeyer

Universiteit Antwerpen, Wilrijk, BELGIUM
W. Beaumont, E. De Langhe, E. De Wolf, L. Rurua, M. Tasevsky

Centro Brasileiro de Pesquisas Fisicas, Rio de Janeiro, RJ, BRAZIL
M. Henrique Gomes E Souza

Universidade do Estado do Rio de Janeiro, Rio de Janeiro, RJ, BRAZIL
A. Santoro

Instituto de Fisica - Universidade Federal do Rio de Janeiro, Rio de Janeiro, RJ, BRAZIL
J. Barreto, M. Vaz

Instituto de Fisica Teorica-Universidade Estadual Paulista, Sao Paulo, SP, BRAZIL
E.M. Gregores, S.M. Lietti, P.G. Mercadante, S.F. Novaes

Institute for Nuclear Research and Nuclear Energy, Sofia, BULGARIA
T. Anguelov, G. Antchev, I. Atanasov, J. Damgov, N. Darmenov[**1], L. Dimitrov, V. Genchev[**1], P. Iaydjiev, B. Panev, S. Piperov, S. Stoykova, G. Sultanov, I. Vankov

University of Sofia, Sofia, BULGARIA
A. Dimitrov, V. Kozhuharov, L. Litov, M. Makariev, A. Marinov, E. Marinova, S. Markov, M. Mateev, B. Pavlov, P. Petkov, C. Sabev, S. Stoynev, Z. Toteva[**1], V. Verguilov

Institute of High Energy Physics, Beijing, CHINA

J. Bai, J.G. Bian, G.M. Chen, H.S. Chen, Y.N. Guo, K. He, G. Huang, C.H. Jiang, Z.J. Ke, B. Li, J. Li, W.G. Li, H. Liu, G. Qin, J.F. Qiu, X. Shen, H. Sun, C. Teng, Y.Y. Wang, Z. Xue, M. Yang, X. Yue, S.Q. Zhang, Y. Zhang, W. Zhao, G.Y. Zhu, H. Zhuang

Peking University, Beijing, CHINA
Y. Ban, J. Cai, S.J. Qian, Z.C. yang, Y.L. Ye, J. Ying

University for Science and Technology of China, Hefei, Anhui, CHINA
J. Wu, Z.P. Zhang

Technical University of Split, Split, CROATIA
N. Godinovic, I. Puljak, I. Soric

University of Split, Split, CROATIA
Z. Antunovic, M. Dzelalija, K. Marasovic

Institute Rudjer Boskovic, Zagreb, CROATIA
V. Brigljevic, K. Kadija

University of Cyprus, Nicosia, CYPRUS
C. Nicolaou, A. Papadakis, P.A. Razis, D. Tsiakkouri

National Institute of Chemical Physics and Biophysics, Tallinn, ESTONIA
A. Hektor, M. Kadastik, K. Kannike, E. Lippmaa, M. Müntel, M. Raidal

Laboratory of Advanced Energy Systems, Helsinki University of Technology, Espoo, FINLAND
P.A. Aarnio

Helsinki Institute of Physics, Helsinki, FINLAND
S. Czellar[**1], E. Haeggstroem, M.A. Heikkinen, J. Härkönen, V. Karimäki[**1], R. Kinnunen, T. Lampén, K. Lassila-Perini, S. Lehti, T. Lindén, P.R. Luukka, S. Michal[**1], T. Mäenpää, J. Nysten, M. Stettler[**1], E. Tuominen, J. Tuominiemi, L. Wendland

Lappeenranta University of Technology, Lappeenranta, FINLAND
T. Tuuva

Laboratoire d'Annecy-le-Vieux de Physique des Particules, IN2P3-CNRS, Annecy-le-Vieux, FRANCE
J.P. Guillaud, P. Nedelec, D. Sillou

DSM/DAPNIA, CEA/Saclay, Gif-sur-Yvette, FRANCE
M. Anfreville, S. Beauceron, E. Bougamont, P. Bredy, R. Chipaux, M. Dejardin, D. Denegri, J. Descamps, B. Fabbro, J.L. Faure, S. Ganjour, F.X. Gentit, A. Givernaud, P. Gras, G. Hamel de Monchenault, P. Jarry, F. Kircher, M.C. Lemaire, B. Levesy[**1], E. Locci, J.P. Lottin, I. Mandjavidze, M. Mur, E. Pasquetto, A. Payn, J. Rander, J.M. Reymond, F. Rondeaux, A. Rosowsky, Z.H. Sun, P. Verrecchia

Laboratoire Leprince-Ringuet, Ecole Polytechnique, IN2P3-CNRS, Palaiseau, FRANCE
M. Anduze, S. Baffioni, F. Beaudette, M. Bercher, U. Berthon, S. Bimbot, J. Bourotte, P. Busson, M. Cerutti, D. Chamont, C. Charlot, C. Collard, D. Decotigny, E. Delmeire, L. Dobrzynski, A.M. Gaillac, Y. Geerebaert, J. Gilly, M. Haguenauer, A. Karar, A. Mathieu, G. Milleret, P. Miné, P. Paganini, P. Poilleux, T. Romanteau, I. Semeniouk, Y. Sirois

Institut de Recherches Subatomiques, IN2P3-CNRS - ULP, LEPSI Strasbourg, UHA Mulhouse, Strasbourg, FRANCE
J.D. Berst[**2], R. Blaes[**3], J.M. Brom, F. Didierjean, F. Drouhin[**1], J.C. Fontaine[**3], U. Goerlach, P. Graehling, L. Gross, L. Houchu, D. Huss, P. Juillot, A. Lounis, C. Maazouzi, D. Mangeol, C. Olivetto, T. Todorov[**1], P. Van Hove, D. Vintache

Institut de Physique Nucléaire de Lyon, IN2P3-CNRS, Université Claude Bernard Lyon I, Villeurbanne, FRANCE

M. Ageron, G. Baulieu, M. Bedjidian, A. Bonnevaux, G. Boudoul, E. Chabanat, C. Combaret, D. Contardo, R. Della Negra, P. Depasse, T. Dupasquier, H. El Mamouni, N. Estre, J. Fay, S. Gascon, N. Giraud, C. Girerd, R. Haroutounian, J.C. Ianigro, B. Ille, M. Lethuillier, N. Lumb, G. Maurelli, L. Mirabito[**1], S. Perries, O. Ravat, B. Trocme

High Energy Physics Institute, Tbilisi State University, Tbilisi, GEORGIA
R. Kvatadze

Institute of Physics Academy of Science, Tbilisi, GEORGIA
V. Roinishvili

RWTH, I. Physikalisches Institut, Aachen, GERMANY
R. Adolphi, R. Brauer, W. Braunschweig, H. Esser, L. Feld, A. Heister, W. Karpinski, K. Klein, C. Kukulies, J. Olzem, A. Ostapchuk, D. Pandoulas, G. Pierschel, F. Raupach, S. Schael, G. Schwering, M. Thomas, M. Weber, B. Wittmer, M. Wlochal

RWTH, III. Physikalisches Institut A, Aachen, GERMANY
A. Adolf, P. Biallass, M. Bontenackels, M. Erdmann, H. Fesefeldt, T. Hebbeker, S. Hermann, G. Hilgers, K. Hoepfner[**1], C. Hof, S. Kappler, M. Kirsch, D. Lanske, B. Philipps, H. Reithler, T. Rommerskirchen, M. Sowa, H. Szczesny, M. Tonutti, O. Tsigenov

RWTH, III. Physikalisches Institut B, Aachen, GERMANY
F. Beissel, M. Duda, G. Flügge, T. Franke, M. Giffels, T. Hermanns, D. Heydhausen, S. Kasselmann, G. Kaussen, T. Kress, A. Linn, A. Nowack, M. Poettgens, O. Pooth, A. Stahl, M. Weber

Deutsches Elektronen-Synchroton, Hamburg, GERMANY
J. Mnich

University of Hamburg, Hamburg, GERMANY
U. Holm, R. Klanner, U. Pein, P. Schleper, N. Schrim, G. Steinbrueck, M. Stoye, R. Van Staa, K. Wick

Institut für Experimentelle Kernphysik, Karlsruhe, GERMANY
P. Blüm, V. Buege, W. De Boer, G. Dirkes, M. Fahrer, M. Feindt, U. Felzmann, J. Fernandez Menendez, M. Frey, A. Furgeri, F. Hartmann[**1], S. Heier, C. Jung, T. Müller, T. Ortega Gomez, C. Piasecki, G. Quast, K. Rabbertz, A. Scheurer, D. Schieferdecker, A. Schmidt, H.J. Simonis, A. Theel, A. Vest, T. Weiler, C. Weiser, J. Weng[**1], V. Zhukov[**4]

University of Athens, Athens, GREECE
G. Karapostoli[**1], P. Katsas, P. Kreuzer, N. Marinelli, A. Panagiotou, C. Papadimitropoulos

Institute of Nuclear Physics "Demokritos", Attiki, GREECE
G. Anagnostou, M. Barone, T. Geralis, C. Kalfas, A. Koimas, A. Kyriakis, S. Kyriazopoulou, D. Loukas, A. Markou, C. Markou, C. Mavrommatis, K. Theofilatos, G. Vermisoglou, A. Zachariadou

University of Ioánnina, Ioánnina, GREECE
X. Aslanoglou, I. Evangelou, P. Kokkas, N. Manthos, I. Papadopoulos, G. Sidiropoulos, F.A. Triantis, P. Vichoudis[**1]

KFKI Research Institute for Particle and Nuclear Physics, Budapest, HUNGARY
G. Bencze[**1], L. Boldizsar, C. Hajdu[**1], P. Hidas, D. Horvath[**5], A. Laszlo, G. Odor, F. Sikler, N. Toth, G. Vesztergombi, P. Zalan

Institute of Nuclear Research ATOMKI, Debrecen, HUNGARY
J. Molnar

Kossuth Lajos University, Debrecen, HUNGARY
N. Beni, A. Kapusi, G. Marian, P. Raics, Z. Szabo, Z. Szillasi, G. Zilizi

Panjab University, Chandigarh, INDIA
K. Ashok, H.S. Bawa, S.B. Beri, V. Bhandari, V. Bhatnagar, M. Kaur, R. Kaur, J.M. Kohli, J.B. Singh

University of Delhi, Delhi, INDIA

A. Bhardwaj, S. Chatterji, B.C. Choudhary, P. Gupta, M. Jha, R.K. Shivpuri, A. Srivastava

Bhabha Atomic Research Centre, Mumbai, INDIA
S. Borkar, M. Dixit, M. Ghodgaonkar, S.K. Kataria, S.K. Lalwani, V. Mishra, A.K. Mohanty, A. Topkar

Tata Institute of Fundamental Research - EHEP, Mumbai, INDIA
T. Aziz, S. Banerjee[**1], N. Cheere, S. Chendvankar, P.V. Deshpande, M. Guchait[**6], A. Gurtu, M. Maity[**7], G. Majumder, K. Mazumdar, M.R. Patil, K. Sudhakar, S.C. Tonwar

Tata Institute of Fundamental Research - HECR, Mumbai, INDIA
B.S. Acharya, S. Banerjee, S. Bheesette, S. Dugad, S.D. Kalmani, V.R. Lakkireddi, N.K. Mondal, N. Panyam, P. Verma

Institute for Studies in Theoretical Physics & Mathematics (IPM), Tehran, IRAN
M. Arabgol, H. Arfaei, M. Hashemi, M. Mohammadi, M. Mohammadi Najafabadi, A. Moshaii, S. Paktinat Mehdiabadi

University College Dublin, Dublin, IRELAND
M. Grunewald

Università di Bari, Politecnico di Bari e Sezione dell' INFN, Bari, ITALY
M. Abbrescia, L. Barbone, A. Colaleo[**1], D. Creanza, N. De Filippis, M. De Palma, G. Donvito, L. Fiore, D. Giordano, G. Iaselli, F. Loddo, G. Maggi, M. Maggi, N. Manna, M.S. Mennea, S. My, S. Natali, S. Nuzzo, G. Pugliese, V. Radicci, F. Romano, G. Selvaggi, L. Silvestris, P. Tempesta, R. Trentadue, G. Zito

Università di Bologna e Sezione dell' INFN, Bologna, ITALY
G. Abbiendi, W. Bacchi, A. Benvenuti, D. Bonacorsi, S. Braibant-Giacomelli, P. Capiluppi, F. Cavallo, C. Ciocca, G. Codispoti, I. D'Antone, G.M. Dallavalle, F. Fabbri, A. Fanfani, P. Giacomelli[**8], C. Grandi, M. Guerzoni, L. Guiducci, S. Marcellini, G. Masetti, A. Montanari, C. Montanari, F. Navarria, F. Odorici, A. Perrotta, A. Rossi, T. Rovelli, G. Siroli, R. Travaglini

Università di Catania e Sezione dell' INFN, Catania, ITALY
S. Albergo, M. Chiorboli, S. Costa, M. Galanti, G. Gatto Rotondo, F. Noto, R. Potenza, G. Russo, C. Sutera, A. Tricomi, C. Tuve

Università di Firenze e Sezione dell' INFN, Firenze, ITALY
A. Bocci, G. Ciraolo, V. Ciulli, C. Civinini, R. D'Alessandro, E. Focardi, C. Genta, P. Lenzi, A. Macchiolo, N. Magini, F. Manolescu, C. Marchettini, L. Masetti, S. Mersi, M. Meschini, S. Paoletti, G. Parrini, R. Ranieri, M. Sani

Università di Genova e Sezione dell' INFN, Genova, ITALY
P. Fabbricatore, M. Fossa, R. Musenich, C. Pisoni

Istituto Nazionale di Fisica Nucleare e Universita Degli Studi Milano-Bicocca, Milano, ITALY
G. Cattaneo, A. De Min, M. Dominoni, F.M. Farina, F. Ferri, A. Ghezzi, P. Govoni, R. Leporini, S. Magni, M. Malberti, S. Malvezzi, S. Marelli, D. Menasce, L. Moroni, P. Musella, P. Negri, M. Paganoni, D. Pedrini, A. Pullia, S. Ragazzi, N. Redaelli, C. Rovelli, M. Rovere, L. Sala, S. Sala, R. Salerno, T. Tabarelli de Fatis, S. Vigano'

Istituto Nazionale di Fisica Nucleare de Napoli (INFN), Napoli, ITALY
G. Comunale, F. Fabozzi, D. Lomidze, S. Mele, P. Paolucci, D. Piccolo, C. Sciacca

Università di Padova e Sezione dell' INFN, Padova, ITALY
P. Azzi, N. Bacchetta, M. Bellato, M. Benettoni, D. Bisello, E. Borsato, A. Candelori, P. Checchia, E. Conti, M. De Mattia, T. Dorigo, U. Dosselli, V. Drollinger, F. Fanzago, F. Gasparini, U. Gasparini, M. Giarin, P. Giubilato, F. Gonella, A. Kaminskiy, S. Karaevskii, V. Khomenkov, S. Lacaprara, I. Lippi, M. Loreti, O. Lytovchenko, M. Mazzucato, A.T. Meneguzzo, M. Michelotto, F. Montecassiano[**1], M. Nigro, M. Passaseo, M. Pegoraro, P. Ronchese, E. Torassa, S. Vanini, S. Ventura, M. Zanetti, P. Zotto, G. Zumerle

Università di Pavia e Sezione dell' INFN, Pavia, ITALY
G. Belli, U. Berzano, C. De Vecchi, R. Guida, M.M. Necchi, S.P. Ratti, C. Riccardi, G. Sani, P. Torre, P. Vitulo

Università di Perugia e Sezione dell' INFN, Perugia, ITALY
F. Ambroglini, E. Babucci, D. Benedetti, M. Biasini, G.M. Bilei[**1], B. Checcucci, L. Fanò, P. Lariccia, G. Mantovani, D. Passeri, M. Pioppi, P. Placidi, V. Postolache, D. Ricci[**1], A. Santocchia, L. Servoli, D. Spiga

Università di Pisa, Scuola Normale Superiore e Sezione dell' INFN, Pisa, ITALY
P. Azzurri, G. Bagliesi, A. Basti, L. Benucci, J. Bernardini, T. Boccali, L. Borrello, F. Bosi, F. Calzolari, R. Castaldi, C. Cerri, A.S. Cucoanes, M. D'Alfonso, R. Dell'Orso, S. Dutta, L. Foà, S. Gennai, A. Giammanco, A. Giassi, D. Kartashov, F. Ligabue, S. Linari, T. Lomtadze, G.A. Lungu, B. Mangano, G. Martinelli, M. Massa, A. Messineo, A. Moggi, F. Palla, F. Palmonari, F. Raffaelli, A. Rizzi, G. Segneri, D. Sentenac, A.T. Serban, G. Sguazzoni, A. Slav, P. Spagnolo, R. Tenchini, G. Tonelli, A. Venturi, P.G. Verdini, M. Vos

Università di Roma I e Sezione dell' INFN, Roma, ITALY
S. Baccaro[**9], L. Barone, A. Bartoloni, F. Cavallari, S. Costantini, I. Dafinei, M. Diemoz, C. Gargiulo, E. Longo, P. Meridiani, G. Organtini, S. Rahatlou

Università di Torino e Sezione dell' INFN, Torino, ITALY
M. Arneodo[**10], R. Bellan, C. Biino, S. Bolognesi, N. Cartiglia, G. Cerminara, M. Cordero, M. Costa, G. Dellacasa, L. Demaria, C. Mariotti, S. Maselli, E. Menichetti, P. Mereu, E. Migliore, V. Monaco, M. Nervo, M.M. Obertino, N. Pastrone, G. Petrillo, A. Romero, M. Ruspa[**10], R. Sacchi, A. Staiano

Università di Trieste e Sezione dell' INFN, Trieste, ITALY
S. Belforte, F. Cossutti, G. Della Ricca

Kyungpook National University, Daegu, KOREA
K. Cho, S.W. Ham, D. Han, D.H. Kim, G.N. Kim, J.C. Kim, W.Y. Kim, K.H. Kwon, S.K. Oh, Y.D. Oh, S.R. Ro, D.C. Son, J.S. Suh

Chonnam National University, Kwangju, KOREA
J.Y. Kim

Konkuk University, Seoul, KOREA
S.Y. Jung, J.T. Rhee

Korea University, Seoul, KOREA
B.S. Hong, S.J. Hong, K.S. Lee, I.K. Park, S.K. Park, K.S. Sim, E. Won

Seoul National University, Seoul, KOREA
S.B. Kim

Universidad Iberoamericana, Mexico City, MEXICO
S. Carrillo Moreno

Centro de Investigacion y de Estudios Avanzados del IPN, Mexico City, MEXICO
H. Castilla Valdez, A. Sanchez Hernandez

Benemerita Universidad Autonoma de Puebla, Puebla, MEXICO
H.A. Salazar Ibarguen

Universidad Autonoma de San Luis Potosi, San Luis Potosi, MEXICO
A. Morelos Pineda

University of Auckland, Auckland, NEW ZEALAND
R.N.C. Gray, D. Krofcheck

University of Canterbury, Christchurch, NEW ZEALAND
M. Billinghurst, P.J. Bones, P.H. Butler, K.L. Kincade, J.C. Williams

National Centre for Physics, Quaid-I-Azam University, Islamabad, PAKISTAN
Z. Aftab, M. Ahmad, U. Ahmad, I. Ahmed, J. Alam Jan, M.I. Asghar, S. Asghar, M. Hafeez, H.R. Hoorani, M. Ibrahim, M. Iftikhar, M.S. Khan, N. Qaiser, I. Rehman, T. Solaija, S. Toor

Institute of Experimental Physics, Warsaw, POLAND
K. Bunkowski, M. Cwiok, H. Czyrkowski, R. Dabrowski, W. Dominik, K. Doroba, A. Kalinowski, M. Konecki, J. Krolikowski, I.M. Kudla, M. Pietrusinski, K. Pozniak[**11], W. Zabolotny[**11], P. Zych

Soltan Institute for Nuclear Studies, Warsaw, POLAND
M. Bluj, R. Gokieli, L. Goscilo, M. Górski, K. Nawrocki, P. Traczyk, G. Wrochna, P. Zalewski

Laboratório de Instrumentação e Física Experimental de Partículas, Lisboa, PORTUGAL
R. Alemany-Fernandez, C. Almeida, N. Almeida, P. Bordalo, R. Bugalho De Moura, J. Gomes, A. Jain, M. Kazana, N. Leonardo, S. Ramos, J. Rasteiro Da Silva, P.Q. Ribeiro, M. Santos, J. Semiao, I. Teixeira, J.P. Teixeira, J. Varela[**1], N. Vaz Cardoso

Joint Institute for Nuclear Research, Dubna, RUSSIA
S. Afanasiev, K. Babich, I. Belotelov, V. Elsha, Y. Ershov, I. Filozova, A. Golunov, I. Golutvin, N. Gorbounov, I. Gramenitski, V. Kalagin, A. Kamenev, V. Karjavin, S. Khabarov, V. Khabarov, Y. Kiryushin, V. Konoplyanikov, V. Korenkov, G. Kozlov, A. Kurenkov, A. Lanev, V. Lysiakov, A. Malakhov, I. Melnitchenko, V.V. Mitsyn, P. Moisenz, K. Moissenz, S. Movchan, E. Nikonov, D. Oleynik, V. Palichik, V. Perelygin, A. Petrosyan, E. Rogalev, V. Samsonov, M. Savina, R. Semenov, S. Shmatov, S. Shulha, V. Smirnov, D. Smolin, A. Tcheremoukhine, O. Teryaev, E. Tikhonenko, S. Vassiliev, A. Vishnevskiy, A. Volodko, N. Zamiatin, A. Zarubin, P. Zarubin, E. Zubarev

Petersburg Nuclear Physics Institute, Gatchina (St Petersburg), RUSSIA
N. Bondar, A. Goliach, V. Golovtsov, Y. Ivanov, V. Kim, V. Kozlov, V. Lebedev, G. Makarenkov, E. Orishchin, V. Sknar, I. Smirnov, V. Sulimov, V. Tarakanov, L. Uvarov, G. Velichko, S. Volkov, A. Vorobyev

Institute for Nuclear Research, Moscow, RUSSIA
I. Andreev, A. Anisimov, S. Gninenko, N. Golubev, D. Gorbunov, M. Kirsanov, A. Kovzelev, N. Krasnikov, V. Matveev, A. Pashenkov, V.E. Postoev, A. Sadovski, A. Solovey, A. Solovey, D. Soloviev, L. Stepanova, A. Toropin

Institute for Theoretical and Experimental Physics, Moscow, RUSSIA
V. Gavrilov, N. Ilina, V. Kaftanov[**1], I. Kiselevich, V. Kolosov, M. Kossov[**1], A. Krokhotin, S. Kuleshov, A. Oulianov, G. Safronov, S. Semenov, V. Stolin, V. Zaytsev

P.N. Lebedev Physical Institute, Moscow, RUSSIA
A.M. Fomenko, N. Konovalova, V. Kozlov, A.I. Lebedev, N. Lvova, S.V. Rusakov, A. Terkulov

Moscow State University, Moscow, RUSSIA
E. Boos, A. Ershov, A. Gribushin, V. Iylin, V. Klyukhin[**1], O.L. Kodolova[**1], I.P. Lokhtin, V. Mikhaylin, S. Petrushanko, L. Sarycheva, V. Savrin, A. Snigirev, K. Teplov, I. Vardanyan

State Research Center of Russian Federation - Institute for High Energy Physics, Protvino, RUSSIA
V. Abramov, I. Azhguirei, S. Bitioukov, P. Goncharov, V. Grishin, A. Inyakin, V. Kachanov, A. Khmelnikov, A. Korablev, V. Krychkine, A. Levine, I. Lobov, V. Petrov, V. Pikalov, S. Slabospitsky, A. Sourkov[**1], A. Sytine, L. Tourtchanovitch, S. Troshin, N. Tyurin, A. Uzunian, A. Volkov, S. Zelepoukine[**12]

Vinca Institute of Nuclear Sciences, Belgrade, SERBIA
P. Adzic, S. Drndarevic[**13], D. Maletic, P. Milenovic, J. Puzovic[**13], N. Smiljkovic[**1], M. Zupan

Centro de Investigaciones Energeticas Medioambientales y Tecnologicas, Madrid, SPAIN
M. Aguilar-Benitez, J. Alberdi, J. Alcaraz Maestre, M. Aldaya Martin, P. Arce[**1], C. Burgos Lazaro, J. Caballero Bejar, E. Calvo, M. Cerrada, M. Chamizo Llatas, N. Colino, M. Daniel, B. De La Cruz, C. Fernandez Bedoya, A. Ferrando, M.C. Fouz, P. Garcia-Abia, J. Hernandez, M.I. Josa, J.M. Luque,

J. Marin, G. Merino, A. Molinero, J.C. Oller, E. Perez Calle, L. Romero, J. Salicio,
C. Villanueva Munoz, C. Willmott

Universidad Autónoma de Madrid, Madrid, SPAIN
C. Albajar, J.F. de Trocóniz, M. Fernandez, I. Jimenez, R.F. Teixeira

Universidad de Oviedo, Oviedo, SPAIN
J. Cuevas, J.M. Lopez, H. Naves Sordo, J.M. Vizan Garcia

Instituto de Física de Cantabria (IFCA), CSIC-Universidad de Cantabria, Santander, SPAIN
A. Calderon, D. Cano Fernandez, I. Diaz Merino, L.A. Garcia Moral, G. Gomez, I. Gonzalez,
J. Gonzalez Sanchez, A. Lopez Virto, J. Marco, R. Marco, C. Martinez Rivero,
P. Martinez Ruiz del Arbol, F. Matorras, A. Patino Revuelta[**1], T. Rodrigo, D. Rodriguez Gonzalez,
A. Ruiz Jimeno, M. Sobron Sanudo, I. Vila, R. Vilar Cortabitarte

CERN, European Organization for Nuclear Research, Geneva, SWITZERLAND
D. Abbaneo, S.M. Abbas, L. Agostino, I. Ahmed, S. Akhtar, N. Amapane, B. Araujo Meleiro,
S. Argiro, S. Ashby, P. Aspell, E. Auffray, M. Axer, A. Ball, N. Bangert, D. Barney, C. Bernet, C. Bloch,
P. Bloch, S. Bonacini, M. Bosteels, V. Boyer, A. Branson, A.M. Brett, H. Breuker, R. Bruneliere,
O. Buchmuller, D. Campi, T. Camporesi, E. Cano, E. Carrone, A. Cattai, R. Chierici, T. Christiansen,
S. Cittolin, E. Corrin, M. Corvo, S. Cucciarelli, B. Curé, A. De Roeck, D. Delikaris, M. Della Negra,
A. Dierlamm, A. Elliott-Peisert, M. Eppard, H. Foeth, R. Folch, S. Fratianni[**14], W. Funk, A. Gaddi,
J.V. Galan Chiner[**15], M. Gastal, J.C. Gayde, H. Gerwig, K. Gill, A.S. Giolo-Nicollerat, F. Glege,
R. Gomez-Reino Garrido, R. Goudard, J. Gutleber, M. Hansen, A. Hervé, H.F. Hoffmann, A. Honma,
M. Huhtinen, G. Iles, V. Innocente, W. Jank, P. Janot, K. Kloukinas, C. Lasseur, M. Lebeau, P. Lecoq,
C. Leonidopoulos, M. Letheren, C. Ljuslin, R. Loos, G. Magazzu, L. Malgeri, M. Mannelli,
A. Marchioro, F. Meijers, E. Meschi, F. Moortgat, J. Nash, R.A. Ofierzynski, A. Oh, P. Olbrechts,
A. Onnela, M. Oriunno, L. Orsini, I. Pal, G. Papotti, R. Paramatti, G. Passardi, B. Perea Solano,
G. Perinic, P. Petagna, A. Petrilli, A. Pfeiffer, M. Pimiä, R. Pintus, H. Postema, R. Principe,
J. Puerta Pelayo, A. Racz, J. Rehn, S. Reynaud, M. Risoldi, P. Rodrigues Simoes Moreira, G. Rolandi,
P. Rosinsky, P. Rumerio, H. Sakulin, D. Samyn, F.P. Schilling, D. Schinzel, C. Schwick, C. Schäfer,
I. Segoni, A. Sharma, P. Siegrist, N. Sinanis, W. Snoeys, P. Sphicas[**16], M. Spiropulu, F. Szoncsó,
O. Teller, J. Troska, E. Tsesmelis, D. Tsirigkas, A. Tsirou, D. Ungaro, F. Vasey, L. Veillet, P. Wertelaers,
A. Wijnant, M. Wilhelmsson, I.M. Willers

Paul Scherrer Institut, Villigen, SWITZERLAND
W. Bertl, K. Deiters, W. Erdmann, K. Gabathuler, S. Heising, R. Horisberger, Q. Ingram, H.C. Kaestli,
D. Kotlinski, S. König, D. Renker, T. Rohe

Institut für Teilchenphysik, Eidgenössische Technische Hochschule (ETH), Zürich, SWITZERLAND
B. Betev, P. Cannarsa[**1], G. Davatz, G. Dissertori, M. Dittmar, L. Djambazov, J. Ehlers, R. Eichler,
G. Faber, K. Freudenreich, C. Grab, A. Holzner, P. Ingenito, U. Langenegger, P. Lecomte, A. Lister,
W. Lustermann, J.D. Maillefaud[**1], A. Nardulli, F. Nessi-Tedaldi, L. Pape, F. Pauss, U. Roser,
H. Rykaczewski, A. Starodumov[**17], F. Stoeckli, H. Suter, L. Tauscher, P. Trueb[**18], G. Viertel,
H.P. von Gunten

Universität Zürich, Zürich, SWITZERLAND
E. Alagoz, C. Amsler, V. Chiochia, C. Hoermann, K. Prokofiev, C. Regenfus, P. Robmann, T. Speer,
S. Steiner, L. Wilke

National Central University, Chung-Li, TAIWAN
S. Blyth, Y.H. Chang, E.A. Chen, A. Go, C.C. Hung, C.M. Kuo, W. Lin

National Taiwan University (NTU), Taipei, TAIWAN
P. Chang, Y. Chao, K.F. Chen, Z. Gao[**1], Y. Hsiung, J.G. Shiu, K. Ueno, Y. Velikzhanin, P. Yeh

Cukurova University, Adana, TURKEY
M. Bakirci, S. Cerci, I. Dumanoglu, S. Erturk, S. Esen, E. Eskut, A. Kayis Topaksu, H. Ozkurt,

A. Polatöz, K. Sogut, G. Önengüt

Middle East Technical University, Physics Department, Ankara, TURKEY
H. Gamsizkan, C. Ozkan, S. Sekmen, M. Serin-Zeyrek, R. Sever, E. Yazgan, M. Zeyrek

Bogaziçi University, Department of Physics, Istanbul, TURKEY
A. Cakir[**19], K. Cankocak[**20], M. Deliomeroglu, D. Demir[**19], E. Gülmez, E. Isiksal[**21], M. Kaya[**22], S. Ozkorucuklu[**23]

Institute of Single Crystals of National Academy of Science, Kharkov, UKRAINE
B. Grinev, V. Senchyshyn

National Scientific Center, Kharkov Institute of Physics and Technology, Kharkov, UKRAINE
L. Levchuk, P. Sorokin

University of Bristol, Bristol, UNITED KINGDOM
D.S. Bailey, T. Barrass, J.J. Brooke, R. Croft, D. Cussans, D. Evans, R. Frazier, N. Grant, M. Hansen, G.P. Heath, H.F. Heath, B. Huckvale, C. Lynch, C.K. Mackay, S. Metson, D.M. Newbold, V.J. Smith, R.J. Tapper

Rutherford Appleton Laboratory, Didcot, UNITED KINGDOM
S.A. Baird, K.W. Bell, R.M. Brown, D.J.A. Cockerill, J.A. Coughlan, P.S. Flower, V.B. Francis, M. French, J. Greenhalgh, R. Halsall, J. Hill, L. Jones, B.W. Kennedy, L. Lintern, A.B. Lodge, J. Maddox, Q. Morrissey, P. Murray, M. Pearson, S. Quinton, J. Salisbury, A. Shah, C. Shepherd-Themistocleous, B. Smith, M. Sproston, R. Stephenson, S. Taghavirad, I.R. Tomalin, J.H. Williams

Imperial College, University of London, London, UNITED KINGDOM
F. Arteche[**1], R. Bainbridge, G. Barber, P. Barrillon, R. Beuselinck, W. Bialas[**1], F. Blekman, D. Britton, D. Colling, G. Daskalakis, G. Dewhirst, S. Dris[**1], C. Foudas, J. Fulcher, S. Greder, G. Hall, J. Jones, J. Leaver, B.C. MacEvoy, O. Maroney, A. Nikitenko[**17], A. Papageorgiou, D.M. Raymond, M.J. Ryan, C. Seez, P. Sharp[**1], M. Takahashi, C. Timlin, T. Virdee[**1], S. Wakefield, M. Wingham, A. Zabi, Y. Zhang, O. Zorba

Brunel University, Uxbridge, UNITED KINGDOM
C. Da Via, I. Goitom, P.R. Hobson, P. Kyberd, C. Munro, J. Nebrensky, I. Reid, O. Sharif, R. Taylor, L. Teodorescu[**24], S.J. Watts, I. Yaselli

Boston University, Boston, Massachusetts, USA
E. Hazen, A.H. Heering, D. Lazic, E. Machado, D. Osborne, J. Rohlf, L. Sulak, S. Wu

Brown University, Providence, Rhode Island, USA
D. Cutts, R. Hooper, G. Landsberg, R. Partridge

University of California, Davis, Davis, California, USA
R. Breedon, M. Case, M. Chertok, J. Conway, P.T. Cox, R. Erbacher, J. Gunion, B. Holbrook, W. Ko, R. Lander, D. Pellett, J. Smith, A. Soha, M. Tripathi, R. Vogt

University of California, Los Angeles, Los Angeles, California, USA
V. Andreev, K. Arisaka, D. Cline, R. Cousins, S. Erhan, M. Felcini, J. Hauser, M. Ignatenko, B. Lisowski, C. Matthey, J. Mumford, S. Otwinowski, P. Schlein, Y. Shi, V. Valuev, R. Wallny, H.G. Wang, X. Yang, Y. Zheng

University of California, Riverside, Riverside, California, USA
R. Clare, D. Fortin, D. Futyan[**1], J.W. Gary, M. Giunta[**1], G. Hanson, G.Y. Jeng, S.C. Kao, H. Liu, G. Pasztor[**25], A. Satpathy, B.C. Shen, R. Stringer, V. Sytnik, R. Wilken, D. Zer-Zion

University of California, San Diego, La Jolla, California, USA
S. Bhattacharya, J.G. Branson, D. Del Re, J. Letts, T. Martin, M. Mojaver, H.P. Paar, M. Pieri, A. Rana, V. Sharma, A. White, F. Würthwein

University of California, Santa Barbara, Santa Barbara, California, USA
A. Affolder, C. Campagnari, C. Hill, J. Incandela, S. Kyre, J. Lamb, J. Richman, D. Stuart, D. White

California Institute of Technology, Pasadena, California, USA
J. Albert, T. Azim[**26], A. Bornheim, J. Bunn, J. Chen, G. Denis, P. Galvez, M. Gataullin, E. Hughes, Y. Kuznetsova, T. Lee, I. Legrand, V. Litvine, D. Nae, H.B. Newman, S. Ravot, S. Shevchenko, S. Singh, C. Steenberg, X. Su, M. Thomas, F. Van Lingen, B.R. Voicu[**1], A. Weinstein, R. Wilkinson, X. Yang, Y. Yang, L.Y. Zhang, K. Zhu, R.Y. Zhu

Carnegie Mellon University, Pittsburgh, Pennsylvania, USA
T. Ferguson, M. Paulini, J. Russ, N. Terentyev, H. Vogel, I. Vorobiev

Cornell Univesity, Ithaca, NY, USA
J. Alexander, D. Cassel, K. Ecklund, B. Heltsley, C. Jones, V. Kuznetsov, A. Ryd, J. Thom, P. Wittich

Fairfield University, Fairfield, Connecticut, USA
C.P. Beetz, G. Cirino, V. Podrasky, C. Sanzeni, D. Winn

Fermi National Accelerator Laboratory, Batavia, Illinois, USA
S. Abdullin[**17], M.A. Afaq[**1], M. Albrow, J. Amundson, G. Apollinari, M. Atac, J.A. Bakken, B. Baldin, L.A.T. Bauerdick, A. Baumbaugh, U. Baur, F. Borcherding, K. Burkett, J.N. Butler, H. Cheung, S. Cihangir, M. Demarteau, D.P. Eartly, J.E. Elias, V..D. Elvira, D. Evans, I. Fisk, J. Freeman, F.J..M. Geurts, D.A. Glenzinski, E. Gottschalk, G. Graham, D. Green, G.M. Guglielmo, Y. Guo, O. Gutsche, A. Hahn, J. Hanlon, S. Hansen, R.M. Harris, T. Hesselroth, S.L. Holm, B. Holzman, S. Iqbal, E. James, M. Johnson, U. Joshi, B. Klima, J. Kowalkowski, T. Kramer, S. Kwan, E. La Vallie, M. Larwill, S. Los, L. Lueking, G. Lukhanin, S. Lusin[**1], K. Maeshima, P. McBride, S.J. Murray, V. O'Dell, M. Paterno, D. Petravick, R. Pordes, O. Prokofyev, V. Rasmislovich, N. Ratnikova, A. Ronzhin, V. Sekhri, E. Sexton-Kennedy, T. Shaw, R.P. Smith, W.J. Spalding, L. Spiegel, M. Stavrianakou, G. Stiehr, I. Suzuki, P. Tan, W. Tanenbaum, S. Tkaczyk, S. Veseli, R. Vidal, H. Wenzel, J. Whitmore, W.M. Wu, Y. Wu, A. Yagil, J. Yarba, T. Yetkin[**27], J.C. Yun

University of Florida, Gainesville, Florida, USA
D. Acosta, P. Avery, V. Barashko, P. Bartalini, D. Bourilkov[**28], R. Cavanaugh, S. Dolinsky, A. Drozdetskiy, R.D. Field, Y. Fu, L. Gray, D. Holmes, B.J. Kim, S. Klimenko, J. Konigsberg, A. Korytov, K. Kotov, P. Levchenko, A. Madorsky, K. Matchev, G. Mitselmakher, Y. Pakhotin, H. Pi, C. Prescott, L. Ramond, P. Ramond, J.L. Rodriguez, M. Schmitt, B. Scurlock, H. Stoeck, S.M. Wang, J. Yelton

Florida International University, Miami, Florida, USA
W. Boeglin, V. Gaultney, L. Kramer, S. Linn, P. Markowitz, G. Martinez, B. Raue, J. Reinhold

Florida State University, Tallahassee, Florida, USA
A. Askew, M. Bertoldi, W. Dharmaratna, Y. Gershtein, S. Hagopian, V. Hagopian, M. Jenkins, K.F. Johnson, J. Mc Donald, H. Prosper, J. Thomaston, H. Wahl

Florida Institute of Technology, Melbourne, Florida, USA
M. Baarmand, L. Baksay[**29], M. Hohlmann, H. Mermerkaya, I. Vodopianov

University of Illinois at Chicago (UIC), Chicago, Illinois, USA
M.R. Adams, R.R. Betts, C. Gerber E., E. Shabalina, C. Smith, T. Ten

The University of Iowa, Iowa City, Iowa, USA
U. Akgun, A.S. Ayan, A. Cooper, P. Debbins, F. Duru, M. Fountain, N. George, E. McCliment, J.P. Merlo, A. Mestvirishvili, M.J. Miller, C. Newsom, J.E. Norbeck, Y. Onel, I. Schmidt, S. Wang

Iowa State University, Ames, Iowa, USA
E.W. Anderson, O. Atramentov, J.M. Hauptman, J. Lamsa

Johns Hopkins University, Baltimore, Maryland, USA
B.A. Barnett, B. Blumenfeld, C.Y. Chien, D.W. Kim, P. Maksimovic, S. Spangler, M. Swartz

The University of Kansas, Lawrence, Kansas, USA
P. Baringer, A. Bean, D. Coppage, M. Murray

Kansas State University, Manhattan, Kansas, USA
T. Bolton, A. Khanov[**17], Y. Maravin, F. Rizatdinova, R. Sidwell, N. Stanton, E. Von Toerne

University of Maryland, College Park, Maryland, USA
D. Baden, R. Bard, S.C. Eno, T. Grassi, N.J. Hadley, R.G. Kellogg, S. Kunori, A. Skuja

Massachusetts Institute of Technology, Cambridge, Massachusetts, USA
R. Arcidiacono, M. Ballintijn, G. Bauer, I. Kravchenko, C. Loizides, S. Nahn, C. Paus, S. Pavlon, C. Roland, G. Roland, K. Sumorok, G. Veres, B. Wyslouch

University of Minnesota, Minneapolis, Minnesota, USA
D. Bailleux, S. Corum, P. Cushman, A. De Benedetti, A. Dolgopolov, R. Egeland, G. Franzoni, W.J. Gilbert, J. Grahl, J. Haupt, Y. Kubota, J. Mans, N. Pearson, E.B. Petersen, R. Rusack, A. Singovsky

University of Mississippi, University, Mississippi, USA
L.M. Cremaldi, R. Godang, R. Kroeger, D. Sanders, D. Summers

University of Nebraska-Lincoln, Lincoln, Nebraska, USA
K. Bloom, D.R. Claes, A. Dominguez, M. Eads, D. Johnston, C. Lundstedt, S. Malik, G.R. Snow

Northeastern University, Boston, Massachusetts, USA
G. Alverson, E. Barberis, O. Boeriu, I. Britvitch, G. Eulisse, A. Kuznetsov, J. Moromisato, Y. Musienko[**30], S. Muzaffar, I. Osborne, S. Reucroft, J. Swain, L. Taylor, L. Tuura, D. Wood

Northwestern University, Evanston, Illinois, USA
B. Gobbi, M. Kubantsev, H. Schellman, E. Spencer, M. Velasco

University of Notre Dame, Notre Dame, Indiana, USA
B. Baumbaugh, N.M. Cason, M. Hildreth, D.J. Karmgard, A. Kharchilava, R. Ruchti, J. Warchol, M. Wayne

The Ohio State University, Columbus, Ohio, USA
B. Bylsma, L.S. Durkin, J. Gilmore, J. Gu, D. Herman, P. Killewald, K. Knobbe, T.Y. Ling

Princeton University, Princeton, New Jersey, USA
P. Elmer, D. Marlow, P. Piroué, D. Stickland, C. Tully, T. Wildish, S. Wynhoff, Z. Xie

Purdue University, West Lafayette, Indiana, USA
K. Arndt, K. Banicz, V.E. Barnes, G. Bolla, D. Bortoletto, A. Bujak, A.F. Garfinkel, O. Gonzalez Lopez, L. Gutay, N. Ippolito, Y. Kozhevnikov, A.T. Laasanen, V. Maroussov, D.H. Miller, J. Miyamoto, N. Neumeister, C. Rott, A. Roy, A. Sedov, I. Shipsey

Purdue University Calumet, Hammond, Indiana, USA
N. Parashar

Rice University, Houston, Texas, USA
G. Eppley, S.J. Lee, M. Matveev, T. Nussbaum, B.P. Padley, J. Roberts, A. Tumanov, P. Yepes

University of Rochester, Rochester, New York, USA
A. Bodek, H. Budd, Y.S. Chung, P. De Barbaro[**1], R. Demina, G. Ginther, Y. Gotra, E. Halkiadakis, S. Korjenevski, W. Sakumoto, P. Slattery, P. Tipton

Rutgers, the State University of New Jersey, Piscataway, New Jersey, USA
E. Bartz, J. Doroshenko, P.F. Jacques, M.S. Kalelkar, A. lath, A. Macpherson[**1], L. Perera, R. Plano, S. Schnetzer, S. Somalwar, R. Stone, G. Thomson, T.L. Watts, S. Worm

Texas Tech University, Lubbock, Texas, USA
N. Akchurin, K.W. Carrell, K. Gumus, H. Kim, V. Papadimitriou, A. Sill, M. Spezziga, E. Washington, R. Wigmans, L. Zhang

Vanderbilt University, Nashville, Tennessee, USA
T. Bapty, D. Engh, K. Givens, W. Johns, S. Neema, S. Nordstrom, S. Pathak, P. Sheldon, E. Vaandering, M. Webster

University of Wisconsin, Madison, Wisconsin, USA
Y.W. Baek, D. Bradley, D. Carlsmith, I. Crotty[**1], S. Dasu, F. Feyzi, T. Gorski, M. Grothe[**31], M. Jaworski, P. Klabbers, A. Lanaro, R. Loveless, M. Magrans de Abril, D. Reeder, W.H. Smith, D. Wenman

Yale University, New Haven, Connecticut, USA
G.S. Atoyan[**30], S. Dhawan, V. Issakov, H. Neal, A. Poblaguev, M.E. Zeller

Institute of Nuclear Physics of the Uzbekistan Academy of Sciences, Ulugbek, Tashkent, UZBEKISTAN
B.S. Yuldashev

**1: Also at CERN, European Organization for Nuclear Research, Geneva, Switzerland

**2: Also at Université Louis Pasteur, Strasbourg, France

**3: Also at Université de Haute-Alsace, Mulhouse, France

**4: Also at Moscow State University, Moscow, Russia

**5: Also at Institute of Nuclear Research ATOMKI, Debrecen, Hungary

**6: Also at Tata Institute of Fundamental Research - HECR, Mumbai, India

**7: Also at University of Visva-Bharati, Santiniketan, India

**8: Also at University of California, Riverside, Riverside, USA

**9: Also at ENEA - Casaccia Research Center, S. Maria di Galeria, Italy

**10: Also at Università del Piemonte Orientale, Novara, Italy

**11: Also at Warsaw University of Technology, Institute of Electronic Systems, Warsaw, Poland

**12: Also at Institut für Teilchenphysik, Eidgenössische Technische Hochschule (ETH), Zürich, Switzerland

**13: Also at Faculty of Physics of University of Belgrade, Belgrade, Serbia

**14: Also at Politecnico di Torino, Torino, Italy

**15: Also at Universidad Politecnica de Valencia, Valencia, Spain

**16: Also at University of Athens, Athens, Greece

**17: Also at Institute for Theoretical and Experimental Physics, Moscow, Russia

**18: Also at Paul Scherrer Institut, Villigen, Switzerland

**19: Also at Izmir Institute of Technology (IYTE), Izmir, Turkey

**20: Also at Mugla University, Mugla, Turkey

**21: Also at Marmara University, Istanbul, Turkey

**22: Also at Kafkas University, Kars, Turkey

**23: Also at Suleyman Demirel University, Isparta, Turkey

**24: Also at University of Bucharest, Bucuresti-Magurele, Romania

**25: Also at KFKI Research Institute for Particle and Nuclear Physics, Budapest, Hungary

**26: Also at National University of Sciences And Technology, Rawalpindi Cantt, Pakistan

**27: Also at Cukurova University, Adana, Turkey

**28: Also at Institute for Nuclear Research and Nuclear Energy, Sofia, Bulgaria

**29: Also at Kossuth Lajos University, Debrecen, Hungary
**30: Also at Institute for Nuclear Research, Moscow, Russia
**31: Also at Università di Torino e Sezione dell' INFN, Torino, Italy

Executive Summary

The Large Hadron Collider will provide extraordinary opportunities in particle physics based on its unprecedented collision energy and luminosity when it begins operation in 2007. The principal aim of this Technical Design Report is to present the strategy of CMS to explore the rich physics programme offered by the LHC: Volume 1 covering the operational procedures and reconstruction tools necessary to perform physics at the LHC, and Volume 2 demonstrating the physics capability of the CMS experiment based on this foundation.

In the first volume we highlight the final detector configuration as it will appear shortly after LHC start-up, including new detectors in the forward regions and for determining the beam luminosity. Results on the performance of the CMS detectors as obtained from detailed simulations are presented for realistic operating conditions, and validated where possible against test-beam or cosmic-ray data. Schemes to synchronize, calibrate, align, and monitor the detectors before, during and after LHC start-up are also given. Reconstruction algorithms developed to perform measurements of muons, electrons, photons, jets, taus, heavy-flavour quarks and the missing transverse energy using these detector subsystems are described. The performance of the reconstruction algorithms is determined from detailed simulations for realistic operating conditions, but techniques to measure the performance from LHC data are described as well. Parameterizations of the performance have been derived to facilitate faster simulations for some of the physics studies described in Volume 2. Included in this first volume are descriptions of the software components needed to implement all of the above, covering simulation, reconstruction, calibration and alignment, monitoring, and visualization.

The second volume covers the capability of the CMS experiment to address physics at the LHC. The prime goals of CMS are to explore physics at the TeV scale and to study the mechanism of electroweak symmetry breaking—through the discovery of the Higgs particle or otherwise. To carry out this task, CMS must be prepared to search for new particles, such as the Higgs boson or supersymmetric partners of the Standard Model particles, from the start-up of the LHC since new physics at the TeV scale may manifest itself with modest data samples of the order of 1 fb^{-1} or less. The experience of the Cosmic Data Challenge plays a crucial role in the preparation of CMS experiment, whereby calibration, alignment and reconstruction procedures are tested and made ready in advance of the LHC pilot and first physics runs. Lessons drawn from this test are described in this volume. Also described in Volume 2 is the online selection strategy of CMS as well as the analysis model and the data flow.

The tools that have been prepared in Volume 1 are applied in Volume 2 to study in great detail and with all the methodology of performing an analysis on CMS data specific benchmark processes upon which to gauge the performance of CMS. These processes cover several Higgs boson decay channels, the production and decay of new particles such as Z' and supersymmetric particles, B_s production and processes in heavy ion collisions. The simulation of these benchmark processes includes subtle effects such as possible detector miscalibration

and misalignment. Besides these benchmark processes, the physics reach of CMS is studied for a large number of signatures arising in the Standard Model and also in theories beyond the Standard Model for integrated luminosities of 10 fb^{-1} up to the asymptotic value of 300 fb^{-1}. The Standard Model processes include QCD, B-physics, diffraction, detailed studies of the top quark properties, and electroweak physics topics such as the W and Z^0 boson properties. The production and decay of the Higgs particle is studied for all observable channels, and the precision with which the Higgs boson properties can be derived is determined. About ten different supersymmetry benchmark points are analysed, and methods to extract e.g. the sparticle masses are tested. Furthermore, the discovery reach for a plethora of alternative models for new physics is explored, notably extra dimensions, new vector boson high mass states, little Higgs models, technicolour and others.

In summary, the content of these two volumes is meant to serve as a comprehensive reference for new CMS collaborators. It provides an entry point to the documentation of the standard simulation, reconstruction, and analysis tools and provides a measure of the expected detector performance and physics reach as we head into the LHC era.

Structure of Volume 1

Chapter 1, the Introduction, describes the context of this document.

Chapter 2 describes the software foundation being prepared for LHC data-taking, including the underlying framework, simulation tools, reconstruction, calibration, alignment, monitoring, and visualization of data.

Chapters 3–6 describe the design, operation and performance of the muon detectors, electromagnetic and hadron calorimeters, and inner tracking detectors.

Chapter 7 describes the design, operation and performance of detectors in the forward regions.

Chapter 8 describes techniques to measure the luminosity and the design, operation and performance of associated detectors.

Chapters 9–12 describe the reconstruction algorithms designed to measure muons, electrons, photons, jets, missing transverse energy, taus, and heavy-flavour quarks.

Contents

1 Introduction **1**
- 1.1 The LHC machine . . . 2
- 1.2 Coordinate conventions . . . 4
- 1.3 Physics and detector requirements . . . 4
  - 1.3.1 Summary of detector requirements . . . 6
- 1.4 Experimental challenge . . . 7
- 1.5 CMS: the overall concept . . . 7
  - 1.5.1 Magnet . . . 9
  - 1.5.2 Muon system . . . 11
  - 1.5.3 Electromagnetic calorimeter . . . 14
  - 1.5.4 Hadron calorimeter . . . 15
  - 1.5.5 Inner tracking system . . . 18
  - 1.5.6 Trigger and data acquisition . . . 21
  - 1.5.7 Software and computing . . . 24
  - 1.5.8 Performance: mass resolutions . . . 24
  - 1.5.9 Radiation levels in CMS . . . 24
- 1.6 The current state of the construction, installation and commissioning of CMS 25

2 Software Components **28**
- 2.1 CMS software architecture . . . 28
  - 2.1.1 Requirements . . . 29
  - 2.1.2 Architecture design . . . 29
  - 2.1.3 Software domain decomposition . . . 30
- 2.2 Framework . . . 31
  - 2.2.1 Event Data Model–data access . . . 31
  - 2.2.2 Framework-module types and communication . . . 32
  - 2.2.3 Framework services . . . 35
- 2.3 Event Filter . . . 38
  - 2.3.1 Filter Farm architecture . . . 40
  - 2.3.2 Data flow . . . 41
  - 2.3.3 Control flow . . . 44
- 2.4 Detector description . . . 45
  - 2.4.1 Geometry Service . . . 46
  - 2.4.2 Detector numbering scheme . . . 46

2.5 Simulation . . . 47
2.5.1 Generation of physics events . . . 48
2.5.2 Detailed simulation framework . . . 48
2.5.3 Pile-up treatment . . . 49
2.5.4 Digitization of simulated detector hits . . . 50
2.5.5 Simulation physics validation and performance . . . 52
2.6 Fast simulation . . . 55
2.6.1 Simulation of the tracker response . . . 56
2.6.2 Simulation of the calorimeter response to electrons and photons . . . 58
2.6.3 Simulation of the calorimeter response to hadrons . . . 60
2.6.4 Simulation of the detector response to muons . . . 61
2.7 Event selection and reconstruction . . . 65
2.7.1 Organization of the reconstruction code . . . 66
2.7.2 Local reconstruction . . . 67
2.7.3 Global reconstruction . . . 67
2.7.4 Combined reconstruction–high level objects . . . 68
2.8 Calibration and alignment . . . 70
2.8.1 Different source of technical data . . . 70
2.8.2 Calibration and alignment data . . . 72
2.8.3 Calibration and alignment model . . . 74
2.9 Data quality monitoring . . . 76
2.9.1 Architecture . . . 76
2.9.2 Components and data flow . . . 77
2.9.3 Control flow . . . 79
2.9.4 Subdetector monitoring . . . 80
2.10 Visualization . . . 80
2.10.1 Design of visualization software . . . 81
2.10.2 CMS visualization applications . . . 83
2.11 Software development environment, quality assurance, and testing . . . 85
2.11.1 Description and general concept . . . 85
2.11.2 Code management . . . 86
2.11.3 External packages . . . 87
2.11.4 Platforms and compilers . . . 87
2.11.5 Releases and release strategy . . . 87
2.11.6 Code distribution . . . 88
2.11.7 Documentation . . . 89
2.11.8 User support and mailing lists . . . 89
2.11.9 Software testing and validation . . . 89

3 Muon System **92**
3.1 System overview . . . 92
3.1.1 Detector layout . . . 92

3.1.2 Readout . . . 95
3.1.3 Data formats . . . 96
3.2 Operations . . . 97
3.2.1 Calibration . . . 97
3.2.2 Alignment . . . 101
3.2.3 Monitoring . . . 109
3.2.4 Databases . . . 111
3.3 Detector response . . . 114
3.3.1 Simulation . . . 114
3.3.2 Simulation and test-beam comparisons . . . 117
3.4 Level-1 trigger . . . 123
3.4.1 Overview of Level-1 muon trigger algorithms . . . 123
3.4.2 DT trigger . . . 124
3.4.3 CSC trigger . . . 124
3.4.4 RPC trigger . . . 125
3.4.5 Global Muon Trigger . . . 125
3.4.6 Test beam results and trigger emulator . . . 126
3.4.7 Muon trigger performance and trigger rates . . . 126
3.5 Detector Level Reconstruction . . . 130
3.5.1 Position reconstruction and track-segment reconstruction in the barrel drift tubes . . . 130
3.5.2 Clusterization and track-segment reconstruction in the endcap cathode strip chambers . . . 132
3.5.3 Clusterization in resistive plate chambers . . . 135
3.5.4 Comparisons with data . . . 136
3.5.5 RPC cosmic-ray tests . . . 137
3.6 Commissioning . . . 139
3.6.1 Detector start-up scenario . . . 139
3.6.2 Synchronization . . . 139
3.6.3 Muon triggers . . . 143
3.6.4 Momentum calibration . . . 144

4 Electromagnetic Calorimeter **146**
4.1 Description of the ECAL . . . 146
4.1.1 The ECAL layout and geometry . . . 146
4.1.2 Lead tungstate crystals . . . 148
4.1.3 Photodetectors . . . 148
4.1.4 Preshower detector . . . 150
4.2 Readout electronics . . . 151
4.2.1 Front end . . . 151
4.2.2 Off-detector electronics . . . 152
4.2.3 Trigger and readout . . . 154
4.2.4 Trigger primitive generation . . . 155

4.2.5 Selective readout . . . 156
4.2.6 Preshower electronics . . . 158
4.2.7 Synchronization . . . 160
4.3 Detector performance . . . 163
4.3.1 Amplitude reconstruction . . . 164
4.3.2 Energy resolution in test beam . . . 166
4.4 Calibration and alignment . . . 168
4.4.1 Intercalibration at start-up . . . 169
4.4.2 Phi independence . . . 172
4.4.3 Intercalibration using single electrons . . . 175
4.4.4 Use of $Z \rightarrow$ ee . . . 180
4.4.5 Intercalibration from π^0 and $\eta \rightarrow \gamma\gamma$ events . . . 184
4.4.6 Inner bremsstrahlung photons in **Z** boson decays to muons . . . 186
4.4.7 Preshower calibration . . . 187
4.4.8 Correction for crystal transparency changes . . . 189
4.4.9 Alignment . . . 194
4.4.10 Conditions/calibration database . . . 194
4.5 Monitoring . . . 195
4.5.1 Monitoring of operational conditions . . . 195
4.5.2 Monitoring using events . . . 197

5 Hadron Calorimeter **199**
5.1 HCAL design . . . 199
5.2 Detector operation . . . 204
5.2.1 Data path and event format . . . 204
5.2.2 Level-1 trigger . . . 206
5.2.3 Synchronization and timing . . . 206
5.3 HCAL Calibration . . . 209
5.3.1 Charge injector calibration for the ADC . . . 209
5.3.2 Radioactive source calibration . . . 210
5.3.3 Offline calibration and monitoring . . . 215
5.3.4 Database . . . 215
5.4 HCAL performance in test beams . . . 216
5.4.1 Longitunidal shower shape and magnetic field brightening . . . 217
5.4.2 Response and resolution . . . 219
5.4.3 Low energy e/π response . . . 220
5.4.4 Cluster-based response compensation . . . 222

6 Inner Tracking System **225**
6.1 Simulation . . . 225
6.1.1 Simulated geometry . . . 225
6.1.2 Simulation of the detector response . . . 226
6.2 Data handling . . . 233

6.2.1 Databases . . . 233
6.2.2 Data format and size . . . 235
6.3 Configuration of the tracker hardware and electronics . . . 237
6.3.1 Configuration of the Pixel tracker . . . 237
6.3.2 Configuration of the Silicon Strip Tracker . . . 239
6.4 Track reconstruction . . . 240
6.4.1 Tracker clusterization . . . 241
6.4.2 Seed generation . . . 244
6.4.3 Pattern recognition . . . 246
6.4.4 Ambiguity resolution . . . 251
6.4.5 Track fitting and smoothing . . . 251
6.4.6 Track reconstruction performance . . . 252
6.4.7 Fast track reconstruction with pixel detector only . . . 255
6.4.8 Track reconstruction in heavy-ion collisions . . . 255
6.5 Vertex reconstruction . . . 260
6.5.1 Vertex fitting . . . 261
6.5.2 The kinematic vertex fit . . . 264
6.5.3 Primary-vertex finding with Pixel Tracks . . . 264
6.5.4 Primary-vertex finding using fully reconstructed tracks . . . 266
6.5.5 Influence of tracker misalignment on vertex reconstruction . . . 268
6.5.6 Secondary-vertex reconstruction . . . 270
6.6 Alignment . . . 277
6.6.1 General alignment strategy . . . 277
6.6.2 Alignment data samples . . . 278
6.6.3 Laser Alignment System . . . 279
6.6.4 Simulation of misalignment . . . 281
6.6.5 Track-based alignment . . . 283
6.6.6 Alignment of a test-beam setup . . . 289
6.7 Data quality monitoring . . . 290
6.7.1 Quantities to monitor . . . 291
6.7.2 CPU estimate and requirements . . . 292
6.7.3 Event-display-based monitoring . . . 293

7 Forward Detectors **295**
7.1 Overview . . . 295
7.2 The Castor Calorimeter . . . 297
7.2.1 Description of the calorimeter . . . 297
7.2.2 Beam test results . . . 298
7.3 The Zero Degree Calorimeter . . . 306
7.3.1 Description of the calorimeter . . . 307
7.3.2 Calibration . . . 311
7.3.3 Trigger and DAQ . . . 311

7.3.4 Heat dissipation . . . 312
7.3.5 Performance . . . 312
7.4 Trigger for common CMS/TOTEM runs . . . 312

8 Luminosity Measurement **315**
8.1 Introduction and overview . . . 315
8.1.1 Relevant definitions . . . 315
8.1.2 Goals and requirements . . . 315
8.1.3 General strategy . . . 316
8.2 LHC beam parameters . . . 316
8.3 Online techniques . . . 317
8.3.1 HF . . . 317
8.3.2 Pixel Luminosity Telescope . . . 320
8.4 Offline and normalization techniques . . . 328
8.4.1 TOTEM . . . 328
8.4.2 W and Z rate measurements . . . 329
8.5 Sources of systematic effects . . . 329
8.6 Luminosity monitoring, reporting, and logging . . . 330
8.6.1 Luminosity accounting . . . 330

9 Muons **332**
9.1 Muon reconstruction . . . 332
9.1.1 Standalone muon reconstruction . . . 332
9.1.2 Global muon reconstruction . . . 333
9.1.3 Performance . . . 334
9.1.4 Cosmic-ray and beam-halo muon reconstruction . . . 338
9.1.5 Dimuon reconstruction in heavy-ion collisions . . . 343
9.2 Muon identification . . . 344
9.2.1 Algorithm . . . 345
9.2.2 Performance . . . 349
9.3 Muon isolation . . . 351
9.3.1 Principles of the muon-isolation algorithms . . . 351
9.3.2 Online event selection . . . 353
9.3.3 Example of the isolation procedure applied offline . . . 356
9.3.4 Sensitivity of the muon isolation cut efficiency to the underlying event uncertainties . . . 357
9.4 High-Level trigger . . . 360
9.4.1 Performance . . . 360
9.4.2 High-Level trigger for muons in heavy-ion collisions . . . 360

10 Electrons and Photons **365**
10.1 ECAL clustering and superclustering . . . 365
10.1.1 Superclustering algorithms . . . 365

10.1.2 Algorithmic energy corrections . . . 366
10.1.3 Position measurement . . . 367
10.1.4 Corrections for saturated channels . . . 367
10.2 Electron and photon trigger . . . 368
10.2.1 Level-1 trigger . . . 368
10.2.2 HLT selection of electron and photon candidates . . . 370
10.3 Photon reconstruction and selection . . . 375
10.3.1 Photon energy reconstruction . . . 375
10.3.2 Photon isolation . . . 376
10.3.3 Converted photons . . . 383
10.3.4 π^0 rejection . . . 388
10.3.5 Offpointing photons . . . 389
10.4 Electron reconstruction and selection . . . 390
10.4.1 Electron clustering . . . 391
10.4.2 Electron tracking . . . 392
10.4.3 Electron classes . . . 395
10.4.4 Electron energy scale corrections . . . 397
10.4.5 Track-cluster combination . . . 399
10.4.6 Electron isolation . . . 401
10.4.7 Electron identification . . . 401

11 Jets and Missing Transverse Energy **404**
11.1 Tower definition and thresholds . . . 404
11.2 Jet algorithms . . . 407
11.2.1 Iterative cone . . . 407
11.2.2 Midpoint cone . . . 408
11.2.3 Inclusive k_{T} algorithm . . . 408
11.3 Monte Carlo corrections . . . 409
11.4 Jet resolution . . . 409
11.5 Missing transverse energy . . . 411
11.5.1 Comparison with UA1 and CDF . . . 413
11.5.2 Resolution in minimum bias events . . . 414
11.5.3 Missing transverse energy resolution in QCD events . . . 414
11.5.4 Corrections to $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ . . . 416
11.5.5 Angular resolution . . . 419
11.6 Jet calibration . . . 420
11.6.1 Data-driven calibration strategy . . . 421
11.6.2 Dijet balancing . . . 422
11.6.3 γ+jet events . . . 423
11.6.4 Parton-level corrections . . . 427
11.6.5 Jet energy scale calibration using the W boson mass constraint in top quark events . . . 428

11.7 Association of jets with the signal vertex . . . 431
11.8 Jet energy correction using charged tracks . . . 432
11.9 Jet reconstruction in heavy-ion collisions . . . 432
11.9.1 Background subtraction algorithm . . . 434
11.9.2 Performance of jet reconstruction . . . 435

12 b and τ Tagging **437**
12.1 τ tagging tools . . . 437
12.1.1 Tau properties relevant to τ-jet reconstruction and identification . . . 437
12.1.2 Methods for τ tagging and performance . . . 441
12.1.3 High Level Trigger . . . 453
12.1.4 Calibration and tagging efficiency . . . 458
12.2 b-tagging tools . . . 461
12.2.1 General features . . . 461
12.2.2 Track impact parameter based tags . . . 463
12.2.3 Combined secondary vertex tag . . . 466
12.2.4 Soft lepton tags . . . 468
12.2.5 Performance of b-tagging algorithms . . . 471
12.2.6 HLT b tag . . . 473
12.2.7 Robustness of the performance . . . 476
12.2.8 Performance studies with data . . . 478

Glossary **483**

References **489**

Colour Plates **509**

Chapter 1

Introduction

Although the Standard Model (SM) of particle physics has so far been tested to exquisite precision, it is considered to be an effective theory up to some scale $\Lambda \approx$ TeV. The prime motivation of the Large Hadron Collider (LHC) is to elucidate the nature of electroweak symmetry breaking for which the Higgs mechanism is presumed to be responsible. The experimental study of the Higgs mechanism also can shed light on the mathematical consistency of the SM at energy scales above about 1 TeV. However, there are alternatives that invoke more symmetry such as supersymmetry or invoke new forces or constituents such as strongly-broken electroweak symmetry, technicolour, etc. An as yet unknown mechanism is also possible. Furthermore there are high hopes for discoveries that could pave the way towards a unified theory. These discoveries could take the form of supersymmetry or extra dimensions, the latter requiring modification of gravity at the TeV scale. Hence there are many compelling reasons to investigate the TeV energy scale. Hadron colliders are well suited to the task of exploring new energy domains, and the region of 1 TeV constituent centre-of-mass energy can be explored if the proton energy and the luminosity are high enough. The beam energy (7 TeV) and the design luminosity ($\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) of the LHC have been chosen in order to study physics at the TeV energy scale. Hence a wide range of physics is potentially possible with the seven-fold increase in energy and a hundred-fold increase in integrated luminosity over the current hadron collider experiments. These conditions also require a very careful design of the detectors.

The availability of high energy heavy-ion beams at energies over 30 times higher than at the present day accelerators will allow us to further extend the range of the heavy-ion physics programme to include studies of hot nuclear matter.

The focus of this volume is to present a concise yet complete overview of the strategy of the Compact Muon Solenoid (CMS) experiment [1] to exploit the physics opportunities presented by the LHC, from the operational procedures of the detectors and the handling of data to the tools needed to reconstruct high-level physics objects and perform the physics analyses. This foundation will form the basis not only for the detailed studies on the physics capability of CMS to be reported in Volume 2, but also for data-taking itself at LHC start-up.

The construction of the CMS subsystems [2]–[8] is nearly complete and installation and commissioning of some subsystems is well underway in the CMS surface assembly hall at Point 5 (Figs. CP 2–CP 12). Testing of the large superconducting 4 T solenoid, a key aspect of CMS, and the concurrent operation of a slice of the experiment in the so-called "Cosmic Challenge" is foreseen in the second quarter of 2006, and the initial results will be summarized in Volume 2 of this Report.

The software for the experiment is similarly well advanced, and has been used to perform

detailed simulations of the detector response and to implement sophisticated reconstruction algorithms already reported in [8], for example, as well as for most of the results reported here. Recently, however, the need was identified to restructure the software framework in preparation for CMS data-taking in order to implement calibration and alignment strategies, ensure tractable reconstruction results, simplify and standardize the reconstruction modules, and facilitate interactive analyses. A description of this new framework and event data model is described in Chapter 2.

1.1 The LHC machine

The machine parameters relevant for the operation of CMS are listed in Table 1.1. The LHC machine comprises 1232 dipole magnets, with r.f. cavities providing a "kick" that results in an increase in the proton energy of 0.5 MeV/turn. The luminosity is given by:

$$\mathcal{L} = \frac{\gamma f k_B N_p^2}{4\pi \varepsilon_n \beta^*} F, \tag{1.1}$$

where γ is the Lorentz factor, f is the revolution frequency, k_B is the number of bunches, N_p is the number of protons/bunch, ε_n is the normalized transverse emittance (with a design value of $3.75\ \mu\mathrm{m}$), β^* is the betatron function at the IP, and F is the reduction factor due to the crossing angle. The nominal energy of each proton beam is 7 TeV. The design luminosity of $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ leads to around 1 billion proton-proton interactions per second.

Table 1.1: The machine parameters relevant for the LHC detectors.

		pp	**HI**	
Energy per nucleon	E	7	2.76	TeV
Dipole field at 7 TeV	B	8.33	8.33	T
Design Luminosity*	$\mathcal{L}$	10^{34}	10^{27}	$\mathrm{cm}^{-2}\ \mathrm{s}^{-1}$
Bunch separation		25	100	ns
No. of bunches	k_B	2808	592	
No. particles per bunch	N_p	1.15×10^{11}	7.0×10^7	
Collisions				
β-value at IP	β^*	0.55	0.5	m
RMS beam radius at IP	σ^*	16.7	15.9	μm
Luminosity lifetime	τ_L	15	6	hr
Number of collisions/crossing	n_c	≈ 20	–	

* For heavy-ion (HI) operation the design luminosity for Pb-Pb collisions is given.

Another relevant feature is the bunch structure (Fig. 1.1) of the beam, since the various gaps are used for the purposes of synchronization, acquiring calibration data and providing resets to front-end electronics. The detailed beam structure is determined by the injection scheme and properties of the dump system. The bunches are formed in the 26 GeV Proton Synchrotron (PS) with the correct 25 ns spacing. The beam is subsequently accelerated to 450 GeV in the Super Proton Synchrotron (SPS) and transferred to the LHC. This operation is repeated 12 times for each counter-rotating beam. At each transfer, enough space has to

be reserved to accommodate the rise time of the injection kickers. Finally, a longer gap is reserved for the rise time of the dump kicker by eliminating 1 PS batch.

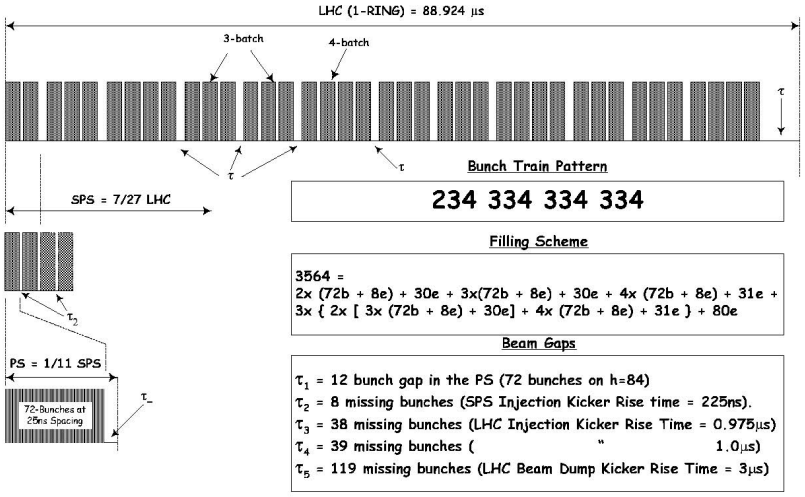


Figure 1.1: The bunch structure of the LHC beam, from [9].

The commissioning of the LHC machine with beams is expected to start in the second half of 2007 [10], first with a few bunches in single beams, followed by a low luminosity pilot physics run. This run will consist of machine development periods interleaved with data-taking runs. The number of bunches per beam is foreseen to be either 43 or 156, with a zero crossing angle.

The expected evolution of the machine performance is given in Table 1.2. During the pilot run the aim is to increase the bunch intensity from 10^{10} p/bunch to 4×10^{10} p/bunch and decrease β^* from 18 m to 2 m. With 156 bunches, a luminosity of 2×10^{31} $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ could be attained. Were it to be possible to approach full bunch intensity, a peak luminosity of 10^{32} $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ would be attainable. This would give on average 3 inelastic collisions per bunch crossing. During a month-long pilot physics run, an integrated luminosity of 10 pb^{-1} could be collected.

In 2008, following the pilot run, operation will start at 75 ns and subsequently 25 ns bunch spacing. The 75 ns operation is considered an important step in the commissioning of the LHC and the experiments. It will be especially useful for establishing synchronization quickly. In the 75 ns mode, each beam will comprise 936 bunches. A crossing angle will be needed, though it can be smaller than for the 25 ns mode.

It will be desirable to move quickly to a bunch spacing of 25 ns. The number of bunches will then be the nominal number (2808) with a crossing angle of 285 μrad. The beam current cannot exceed half the nominal value as part of the beam dump and collimation systems are staged. This will limit the initial luminosity to $\mathcal{L} = 2 \times 10^{33}$ $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ until the 2010 run.

During the first full year of physics running, the LHC should reach a peak luminosity of $\mathcal{L} = 2 \times 10^{33}$ $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. However, the integrated luminosity will most likely be limited by the time taken to master LHC operation. The integrated luminosity is likely to be about 5 fb^{-1} [10] in the first year. It may well be lower, as prolonged machine development periods may be required and higher than foreseen inefficiencies encountered. It should be noted that the way to maximize the integrated luminosity will be to run at a bunch spacing of 75 ns.

Therefore, we base the evaluation of the physics potential on several scenarios of integrated luminosities: $1\ \mathrm{fb}^{-1}$ during the initial operation, 10–30 fb^{-1} in the "low luminosity" phase ($\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) and 100–300 fb^{-1} over several years of operation at design or "high luminosity" ($\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$).

The first run with heavy-ion beams is expected in 2008, after the first pp physics run. Initially, 10 times less bunches will be filled and the beam will have a value of β^* 2 times larger, resulting in a factor ≈ 20 reduced maximum luminosity of $5.4 \times 10^{25}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$. In the following year, depending on the machine performance, the luminosity is expected to reach the nominal $10^{27}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$ that corresponds to a minimum-bias interaction rate of 8 kHz. The average luminosity and the lifetime of the beam in the machine will depend on the number of experiments collecting data, as well as the re-fill time of the accelerator, and the luminosity is expected to be between $4\text{–}6 \times 10^{26}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Table 1.2: Expected evolution of LHC performance parameters during 2007–2008 [11].

	Pilot run 2007	First Physics 2008
Number of bunches	43→156	936→2808
β^*	18 m→2 m	2 m→ 0.55 m
Protons per bunch	$10^{10} \rightarrow 4 \times 10^{10}$ (10^{11})	4×10^{10}
Luminosity	$3 \times 10^{29} \rightarrow 2 \times 10^{31}$ (10^{32})	$10^{32} \rightarrow 2 \times 10^{33}$
Integrated Luminosity	$10\ \mathrm{pb}^{-1}$	$< 5\ \mathrm{fb}^{-1}$

1.2 Coordinate conventions

The coordinate system adopted by CMS has the origin centered at the nominal collision point inside the experiment, the y-axis pointing vertically upward, and the x-axis pointing radially inward toward the center of the LHC. Thus, the z-axis points along the beam direction toward the Jura mountains from LHC Point 5. The azimuthal angle ϕ is measured from the x-axis in the x-y plane. The polar angle θ is measured from the z-axis. Pseudorapidity is defined as $\eta = -\ln\tan(\theta/2)$. Thus, the momentum and energy measured transverse to the beam direction, denoted by p_{T} and E_{T}, respectively, are computed from the x and y components. The imbalance of energy measured in the transverse plane is denoted by $E_{\mathrm{T}}^{\mathrm{miss}}$.

A one-quarter cross-sectional view of the CMS Experiment is shown in Figure CP 1 with some dimensions and lines of constant η superimposed. Some of the labels used to name the detector elements are also introduced.

1.3 Physics and detector requirements

Search for the Higgs Boson

In the design phase of CMS and ATLAS in the early 1990s, the detection of the SM Higgs boson was used as a benchmark to test the performance of the proposed designs. It is a particularly appropriate benchmark since there is a wide range of decay modes depending on the mass of the Higgs boson.

The current lower limit on the mass of the Higgs boson from LEP is $114.4\ \mathrm{GeV}/c^2$. In the vicinity of this limit, the branching fractions of the Higgs boson are dominated by hadronic decays, which are difficult to use to discover the Higgs boson at the LHC due to the large QCD backgrounds and the relatively poor mass resolution that is obtainable with jets. Hence, the search is preferentially conducted using final states that contain isolated leptons and photons, despite the smaller branching ratios.

The natural width of the Higgs boson in the *intermediate-mass* region ($114\ \mathrm{GeV}/c^2 < m_H < 2m_Z$) is only a few MeV, and the observed width of a potential signal will be dominated by the instrumental mass resolution. In the mass interval 114–130 GeV/c^2, the two-photon decay is one of the principal channels likely to yield a significant signal. Central exclusive production of the Higgs might offer the only way to access the $b\bar{b}$ decay mode. The Higgs boson should be detectable via its decay into 2 Z bosons if its mass is larger than about $130\ \mathrm{GeV}/c^2$ (one of the Z's is virtual when m_H is below the ZZ threshold). For $2m_Z < m_H < 600\ \mathrm{GeV}/c^2$ the ZZ decay, with its four-lepton final states, is the mode of choice.

In the region $600 < m_H < 1000\ \mathrm{GeV}/c^2$, the cross section decreases so that higher branching fraction modes involving jets or $E_{\mathrm{T}}^{\mathrm{miss}}$ from W or Z decays have to be used. The jets from W and Z decays will be boosted and may be close to each other in η-ϕ space.

The dominant Higgs-boson production mechanism, for masses up to about $700\ \mathrm{GeV}/c^2$, is gluon-gluon fusion via a t-quark loop. The WW or ZZ fusion mechanism becomes important for the production of higher-mass Higgs bosons. Here, the quarks that emit the Ws or Zs have transverse momenta of the order of W and Z masses. The detection of the resulting high-energy jets in the forward regions ($2 < |\eta| < 5$) can be used to tag the reaction, improving the signal-to-noise ratio and extending the range of masses over which the Higgs can be discovered. These jets are highly boosted and their transverse size is similar to that of a high-energy hadron shower.

More recently, the fusion mechanism has also been found to be useful for detecting an intermediate mass Higgs boson through channels such as $qq \to qqH$, followed by $H \to \tau\tau$.

Search for supersymmetric particles

The decays of supersymmetric particles, such as squarks and gluinos, involve cascades that, if R-parity is conserved, always contain the lightest SUSY particle (LSP). The latter is expected to interact very weakly, thus leading to significant $E_{\mathrm{T}}^{\mathrm{miss}}$ in the final state. The rest of the cascade results in an abundance of leptons and jets (particularly b-jets and/or τ-jets). In GMSB schemes with the LSP decaying into a photon and gravitino, an increased number of hard isolated photons is expected.

Search for new massive vector bosons

The detector requirements for high momenta can be determined by considering decays of high-mass objects such as $Z' \to e^+e^-$ and $\mu^+\mu^-$. The discovery of an object like a Z' boson will, very likely, be limited by the statistical significance of the signal. Ways of distinguishing between different models involve the measurement of the natural width and the forward-backward asymmetry, both of which require sufficiently good momentum resolution at high p_{T} ($\Delta p_{\mathrm{T}}/p_{\mathrm{T}} < 0.1$ at $p_{\mathrm{T}} \approx 1\ \mathrm{TeV}/c$) to determine the sign of the leptons and a pseudorapidity coverage up to $\eta = 2.4$.

Extra dimensions

The existence of extra dimensions can lead to a characteristic energy scale of quantum gravity, M_{D}, which is the analogue of the Planck mass in a D-dimensional theory, and which could lie just beyond the electroweak scale. In terms of experimental signatures, 3 regimes can be distinguished [12]: i) Cis-Planckian, where $E \ll M_{\mathrm{D}}$, leading to signals involving the emission of gravitons that escape into extra dimensions, e.g. $pp \rightarrow$ jet+graviton $\rightarrow$ jet+$E_{\mathrm{T}}^{\mathrm{miss}}$, ii) Planckian, $E \approx M_{\mathrm{D}}$, leading to model-dependent signatures. In string-theory motivated models there are Regge-like excitations that manifest themselves as Z-like resonances with $\approx$ TeV separations in mass, iii) Trans-Planckian, $E \gg M_{\mathrm{D}}$, leading to mini black hole production with spectacular decays involving equal and democratic production of fundamental particles such as leptons, photons, neutrinos, W, Z, jets, etc. The resulting production and kinematic distributions could allow the determination of the Hawking temperature, the mass of the black holes, the number of extra dimensions, etc.

Standard Model

The LHC will also allow studies of QCD, electroweak, and flavour physics. Precision studies can give indications for physics beyond the SM, providing complementary information with respect to direct searches. As an example, extensive tests of QCD will be possible through the measurement of the production of jets and direct photons with transverse energies up to 3–4 TeV and from cross-section measurements which fall by 11 orders of magnitude. Top quarks will be produced at the LHC with a rate measured in Hz, thus the opportunity to test the SM couplings and spin of the top quark is available provided good identification of b-jets in the decays is possible. Searches for flavour changing neutral currents, lepton flavour violation through $\tau \rightarrow 3\mu$ or $\tau \rightarrow \mu\gamma$, measurements of $B_s^0 \rightarrow \mu\mu$, measurements of triple- and quartic-gauge couplings, etc. can open a window onto new physics. Finally, in association with TOTEM, CMS will be able to cover the full range of diffractive physics as well.

Heavy-ion physics

The recent results from RHIC indicate that very strongly interacting nuclear matter is produced in high energy heavy-ion collisions. The most striking experimental signatures of the produced matter are the suppression of high p_T particles (jet quenching) and the strong elliptical flow approaching the hydrodynamic limit. The increase in collision energy from $\sqrt{s_{NN}} = 200$ GeV/c^2 at RHIC to 5500 GeV/c^2 at LHC will allow the extension of studies of jet quenching to much higher p_T and the identification of fully formed jets. The measurements of energy flow at LHC will stringently test the liquid-like behavior of the hot nuclear matter. The increased energy will also allow studies of presently inaccessible hard probes like Υ and Z^0. The studies of jet quenching, energy flow, and quarkonium production will require large-acceptance, high-resolution calorimeters and tracking devices, as well as a flexible trigger.

1.3.1 Summary of detector requirements

The detector requirements for CMS to meet the goals of the LHC physics programme can be summarized as follows:

- Good muon identification and momentum resolution over a wide range of momenta in the region $|\eta| < 2.5$, good dimuon mass resolution ($\approx 1\%$ at 100 GeV/c^2), and the ability to determine unambiguously the charge of muons with $p < 1$ TeV/c.
- Good charged particle momentum resolution and reconstruction efficiency in the inner tracker. Efficient triggering and offline tagging of τ's and b-jets, requiring

pixel detectors close to the interaction region.

- Good electromagnetic energy resolution, good diphoton and dielectron mass resolution ($\approx 1\%$ at $100\ \mathrm{GeV}/c^2$), wide geometric coverage ($|\eta| < 2.5$), measurement of the direction of photons and/or correct localization of the primary interaction vertex, π^0 rejection and efficient photon and lepton isolation at high luminosities.
- Good $E_{\mathrm{T}}^{\mathrm{miss}}$ and dijet mass resolution, requiring hadron calorimeters with a large hermetic geometric coverage ($|\eta| < 5$) and with fine lateral segmentation ($\Delta\eta \times \Delta\phi < 0.1 \times 0.1$).

The design of CMS, detailed in Section 1.5, meets these requirements. The main distinguishing features of CMS are a high-field solenoid, a full silicon-based inner tracking system, and a fully active scintillating crystals-based electromagnetic calorimeter.

1.4 Experimental challenge

The total proton-proton cross-section at $\sqrt{s} = 14$ TeV is roughly 100 mb. At design luminosity the general-purpose detectors will therefore observe an event rate of approximately 10^9 inelastic events/s. This leads to a number of formidable experimental challenges [13].

The online event selection process ("trigger") must reduce the approximately 1 billion interactions/s to no more than about 100 events/s for storage and subsequent analysis. The short time between bunch crossings, 25 ns, has major implications for the design of the readout and trigger systems.

At the design luminosity, a mean of about 20 inelastic (hard-core scattering) collisions will be superimposed on the event of interest. This implies that around 1000 charged particles will emerge from the interaction region every 25 ns. The products of an interaction under study may be confused with those from other interactions in the same bunch crossing. This problem clearly becomes more severe when the response time of a detector element and its electronic signal is longer than 25 ns. The effect of this pile-up can be reduced by using high-granularity detectors with good time resolution, resulting in low occupancy. This requires a large number of detector channels. The resulting millions of detector electronic channels require very good synchronization.

The large flux of particles coming from the interaction region leads to high radiation levels, requiring radiation-hard detectors and front-end electronics.

1.5 CMS: the overall concept

The construction, installation and commissioning of CMS is progressing well, though not without challenges, towards the goal of being ready for collisions in the second half of 2007.

An important aspect driving the detector design and layout is the choice of the magnetic field configuration for the measurement of the momentum of muons. Large bending power is needed to measure precisely the momentum of charged particles. This forces a choice of superconducting technology for the magnets. The design configuration chosen by CMS [1] is discussed below.

The overall layout of CMS is shown in Figure 1.2. At the heart of CMS sits a 13-m-long, 5.9 m inner diameter, 4 T superconducting solenoid. In order to achieve good momentum resolution within a compact spectrometer without making stringent demands on muon-chamber resolution and alignment, a high magnetic field was chosen. The return field is large enough to saturate 1.5 m of iron, allowing 4 muon "stations" to be integrated to ensure robustness and full geometric coverage. Each muon station consists of several layers of aluminium drift tubes (DT) in the barrel region and cathode strip chambers (CSCs) in the endcap region, complemented by resistive plate chambers (RPCs).

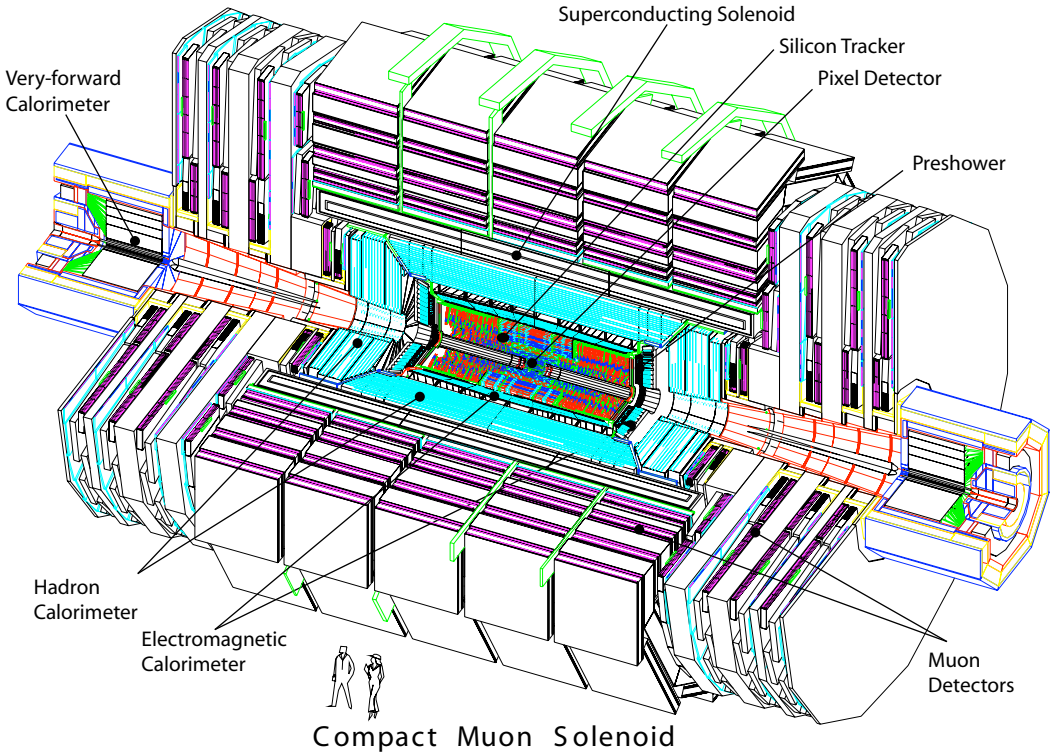


Figure 1.2: An exploded view of the CMS detector.

The bore of the magnet coil is also large enough to accommodate the inner tracker and the calorimetry inside. The tracking volume is given by a cylinder of length 5.8 m and diameter 2.6 m. In order to deal with high track multiplicities, CMS employs 10 layers of silicon microstrip detectors, which provide the required granularity and precision. In addition, 3 layers of silicon pixel detectors are placed close to the interaction region to improve the measurement of the impact parameter of charged-particle tracks, as well as the position of secondary vertices. The EM calorimeter (ECAL) uses lead tungstate ($PbWO_4$) crystals with coverage in pseudorapidity up to $|\eta| < 3.0$. The scintillation light is detected by silicon avalanche photodiodes (APDs) in the barrel region and vacuum phototriodes (VPTs) in the endcap region. A preshower system is installed in front of the endcap ECAL for π^0 rejection. The ECAL is surrounded by a brass/scintillator sampling hadron calorimeter with coverage up to $|\eta| < 3.0$. The scintillation light is converted by wavelength-shifting (WLS) fibres embedded in the scintillator tiles and channeled to photodetectors via clear fibres. This light is detected by novel photodetectors (hybrid photodiodes, or HPDs) that can provide gain and operate in high axial magnetic fields. This central calorimetery is complemented by a

"tail-catcher" in the barrel region—ensuring that hadronic showers are sampled with nearly 11 hadronic interaction lengths. Coverage up to a pseudorapidity of 5.0 is provided by an iron/quartz-fibre calorimeter. The Cerenkov light emitted in the quartz fibres is detected by photomultipliers. The forward calorimeters ensure full geometric coverage for the measurement of the transverse energy in the event.

The overall dimensions of the CMS detector are a length of 21.6 m, a diameter of 14.6 m and a total weight of $12\,500$ tons. The thickness of the detector in radiation lengths (Fig. 1.3) is greater than 25 X_0 for the ECAL, and the thickness in interaction lengths (Fig. 1.4) varies from 7–11λ_{I} for HCAL depending on η. Also shown in both figures is the material depth at each muon station. Details on the material budget of the inner tracker are given in Section 6.1.1.

1.5.1 Magnet

The required performance of the muon system, and hence the bending power, is defined by the narrow states decaying into muons and by the unambiguous determination of the sign for muons with a momentum of $\approx 1\ \mathrm{TeV}/c$. This requires a momentum resolution of $\Delta p/p \approx 10\%$ at $p = 1\ \mathrm{TeV}/c$.

CMS chose a large superconducting solenoid, the parameters of which are given in Table 1.3 [2]. A large bending power can be obtained for a modestly-sized solenoid, albeit a high-field superconducting one, as the bending starts at the primary vertex. A favourable length/radius ratio is necessary to ensure good momentum resolution in the forward region as well. The CMS magnet has been assembled in the surface experiment hall (Figure CP 2). The solenoid will be taken up to full current in the second quarter of 2006.

Table 1.3: Parameters of the CMS superconducting solenoid.

Field	4 T
Inner Bore	5.9 m
Length	12.9 m
Number of Turns	2168
Current	19.5 kA
Stored energy	2.7 GJ
Hoop stress	64 atm

The main features of the CMS solenoid are the use of a high-purity aluminium-stabilised conductor and indirect cooling (by thermosyphon), together with full epoxy impregnation. This technique was successfully used previously in the construction of the large solenoids for ALEPH and DELPHI at LEP and for H1 at HERA. However, the large increase in some parameters such as magnetic field, Ampere-turns, forces and stored energy (2.7 GJ) necessitated changes. In particular, a four-layer winding has been adopted using a novel conductor with a larger cross-section that can withstand an outward pressure (hoop stress) of 64 atmospheres. The conductor carries a current of 20 kA and has a compound structure. The Rutherford-type cable is co-extruded with pure aluminium, which acts as a thermal stabiliser. This "insert" is then electron-beam-welded to 2 "plates", made of a high-strength aluminium alloy, for the mechanical reinforcement. The overall conductor cross section is 64×22 mm^2. The conductor was manufactured in twenty continuous lengths, each with a length of 2.65 km. Four lengths were wound to make each of the 5 coil modules. These modules were assembled

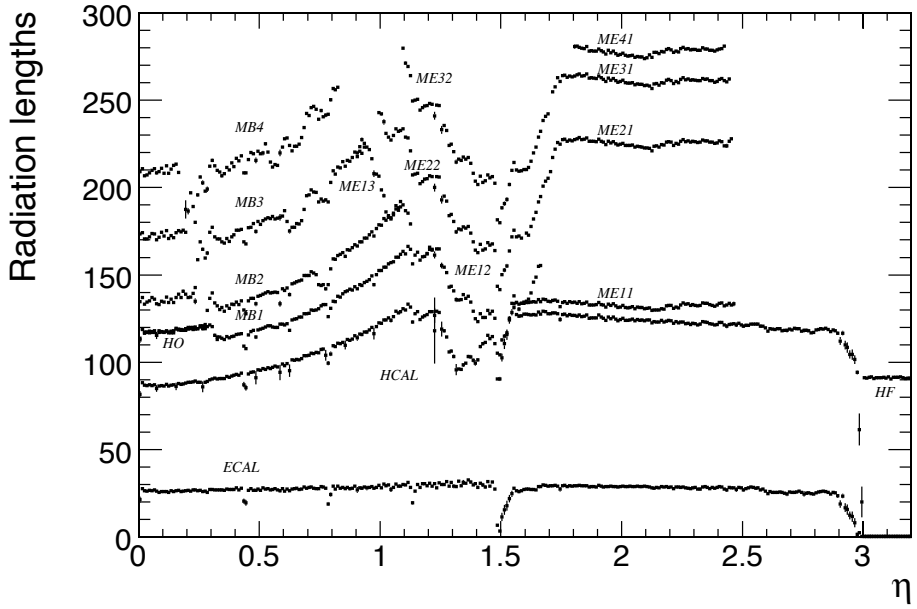


Figure 1.3: Material thickness in radiation lengths after the ECAL, HCAL, and at the depth of each muon station as a function of pseudorapidity. The thickness of the forward calorimeter (HF) remains approximately constant over the range $3 < |\eta| < 5$ (not shown).

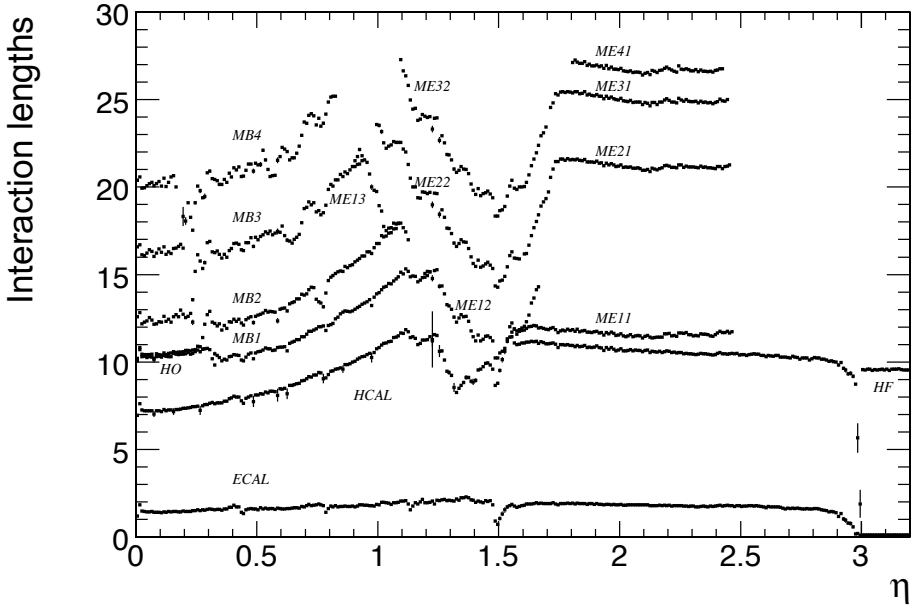


Figure 1.4: Material thickness in interaction lengths after the ECAL, HCAL, and at the depth of each muon station as a function of pseudorapidity. The thickness of the forward calorimeter (HF) remains approximately constant over the range $3 < |\eta| < 5$ (not shown).

and connected together in SX5 at Point 5.

1.5.2 Muon system

Centrally produced muons are measured 3 times: in the inner tracker, after the coil, and in the return flux. Measurement of the momentum of muons using only the muon system is essentially determined by the muon bending angle at the exit of the 4 T coil, taking the interaction point (which will be known to $\approx 20\ \mu\mathrm{m}$) as the origin of the muon. The resolution of this measurement (labelled "muon system only" in Figure 1.5) is dominated by multiple scattering in the material before the first muon station up to p_{T} values of 200 GeV/c, when the chamber spatial resolution starts to dominate. For low-momentum muons, the best momentum resolution (by an order of magnitude) is given by the resolution obtained in the silicon tracker ("inner tracker only" in Figure 1.5). However, the muon trajectory beyond the return yoke extrapolates back to the beam-line due to the compensation of the bend before and after the coil when multiple scattering and energy loss can be neglected. This fact can be used to improve the muon momentum resolution at high momentum when combining the inner tracker and muon detector measurements ("full system" in Figure 1.5).

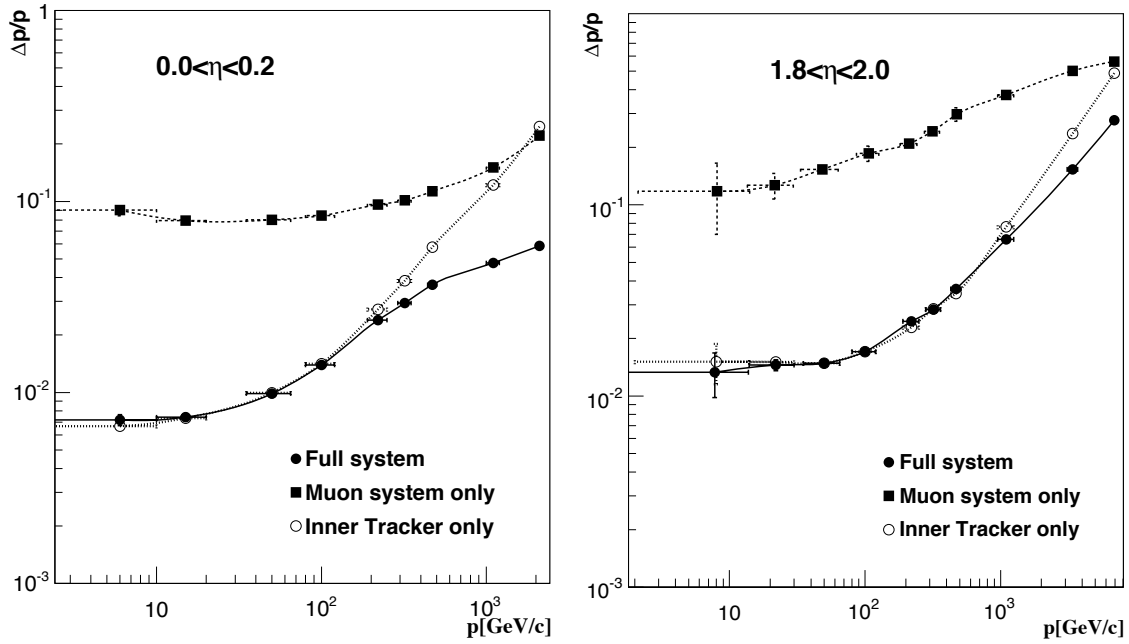


Figure 1.5: The muon momentum resolution versus p using the muon system only, the inner tracker only, or both ("full system"). a) barrel, $|\eta| < 0.2$; b) endcap, $1.8 < |\eta| < 2.0$.

Three types of gaseous detectors are used to identify and measure muons [4]. The choice of the detector technologies has been driven by the very large surface to be covered and by the different radiation environments. In the barrel region ($|\eta| < 1.2$), where the neutron induced background is small, the muon rate is low and the residual magnetic field in the chambers is low, drift tube (DT) chambers are used. In the 2 endcaps, where the muon rate as well as the neutron induced background rate is high, and the magnetic field is also high, cathode strip chambers (CSC) are deployed and cover the region up to $|\eta| < 2.4$. In addition to this, resistive plate chambers (RPC) are used in both the barrel and the endcap

regions. These RPCs are operated in avalanche mode to ensure good operation at high rates (up to 10 kHz/cm^2) and have double gaps with a gas gap of 2 mm. A change from the Muon TDR [4] has been the coating of the inner bakelite surfaces of the RPC with linseed oil for good noise performance. RPCs provide a fast response with good time resolution but with a coarser position resolution than the DTs or CSCs. RPCs can therefore identify unambiguously the correct bunch crossing.

The DTs or CSCs and the RPCs operate within the first level trigger system, providing 2 independent and complementary sources of information. The complete system results in a robust, precise and flexible trigger device. In the initial stages of the experiment, the RPC system will cover the region $|\eta| < 1.6$. The coverage will be extended to $|\eta| < 2.1$ later.

The layout of one quarter of the CMS muon system for initial low luminosity running is shown in Figure 1.6. In the Muon Barrel (MB) region, 4 stations of detectors are arranged in cylinders interleaved with the iron yoke. The segmentation along the beam direction follows the 5 wheels of the yoke (labeled YB-2 for the farthest wheel in $-z$, and YB$+2$ for the farthest is $+z$). In each of the endcaps, the CSCs and RPCs are arranged in 4 disks perpendicular to the beam, and in concentric rings, 3 rings in the innermost station, and 2 in the others. In total, the muon system contains of order 25 000 m^2 of active detection planes, and nearly 1 million electronic channels.

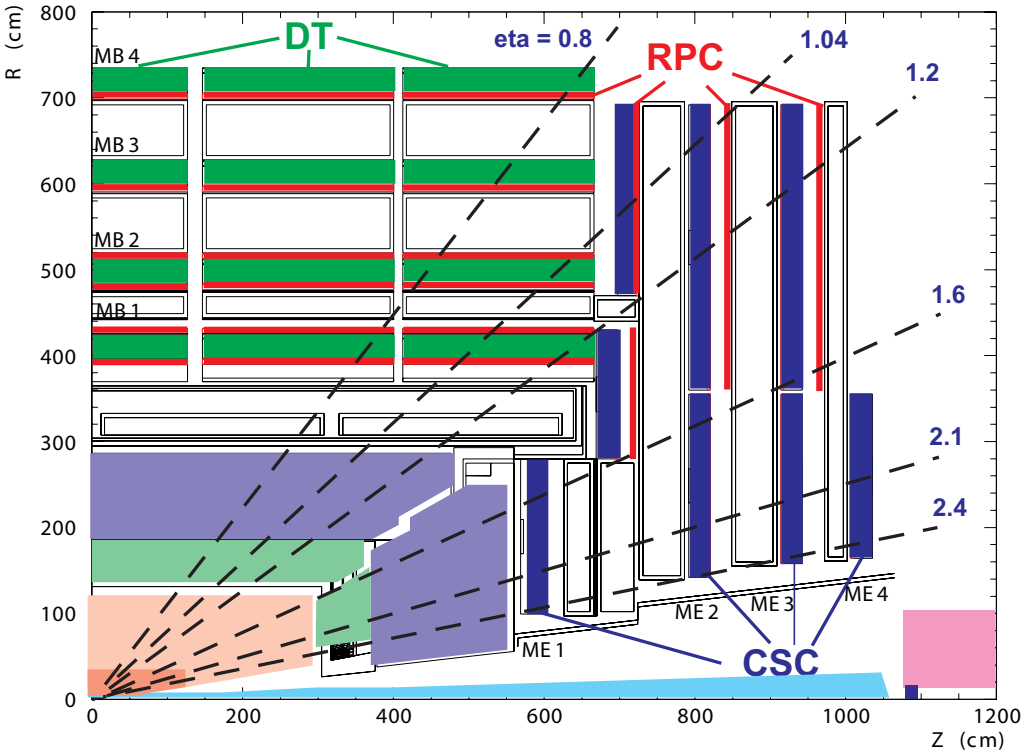


Figure 1.6: Layout of one quarter of the CMS muon system for initial low luminosity running. The RPC system is limited to $|\eta| < 1.6$ in the endcap, and for the CSC system only the inner ring of the ME4 chambers have been deployed.

1.5.2.1 Drift tube chambers

The Barrel Detector, shown partially installed into the YB+1 yoke in Figure CP 3, consists of 250 chambers organized in 4 layers (stations labeled MB1, MB2, MB3 and MB4 with the last being the outermost) inside the magnet return yoke, at radii of approximately 4.0, 4.9, 5.9 and 7.0 m from the beam axis. Each of the 5 wheels of the Barrel Detector is divided into 12 sectors, with each covering a 30° azimuthal angle. Chambers in different stations are staggered so that a high-p_{T} muon produced near a sector boundary crosses at least 3 out of the 4 stations. There are 12 chambers in each of the 3 inner layers. In the 4th layer, the top and bottom sectors host 2 chambers each, thus leading to a total of 14 chambers per wheel in this outermost layer. The MB1, 2 and 3 chambers consist of 12 planes of aluminium drift tubes; 4 r-ϕ measuring planes in each of the 2 outermost "superlayers," separated by about 20 cm and sandwiching a z-superlayer comprising 4 z-measuring planes. The MB4 station does not contain the z-measuring planes. The maximum drift length is 2.0 cm and the single-point resolution is $\approx 200\ \mu\mathrm{m}$. Each station is designed to give a muon vector in space, with a ϕ precision better than $100\ \mu\mathrm{m}$ in position and approximately 1 mrad in direction.

Each DT chamber has 1 or 2 RPCs coupled to it before installation, depending on the station. In stations MB1 and MB2, each package consists of 1 DT chamber sandwiched between 2 RPCs. In stations MB3 and MB4, each package comprises 1 DT chamber and 1 RPC, which is placed on the innermost side of the station. A high-p_{T} muon thus crosses up to 6 RPCs and 4 DT chambers, producing up to 44 measured points in the DT system from which a muon-track candidate can be built.

1.5.2.2 Cathode strip chambers

The Muon Endcap (ME) system comprises 468 CSCs in the 2 endcaps. The complete installation of 1 disk is shown in Figure CP 4. Each CSC is trapezoidal in shape and consists of 6 gas gaps, each gap having a plane of radial cathode strips and a plane of anode wires running almost perpendicularly to the strips. All CSCs except those in the third ring of the first endcap disk (ME1/3) are overlapped in phi to avoid gaps in the muon acceptance. There are 36 chambers in each ring of a muon station, except for the innermost ring of the second through fourth disks (ME2/1, ME3/1, and ME4/1) where there are 18 chambers. The gas ionization and subsequent electron avalanche caused by a charged particle traversing each plane of a chamber produces a charge on the anode wire and an image charge on a group of cathode strips. The signal on the wires is fast and is used in the Level-1 Trigger. However, it leads to a coarser position resolution. A precise position measurement is made by determining the centre-of-gravity of the charge distribution induced on the cathode strips. Each CSC measures up to 6 space coordinates (r, ϕ, z). The spatial resolution provided by each chamber from the strips is typically about $200\ \mu\mathrm{m}$ ($100\ \mu\mathrm{m}$ for ME1/1). The angular resolution in ϕ is of order 10 mrad.

1.5.2.3 Forward resistive plate chambers

The system described in the Muon TDR [4] comprises 4 stations covering the pseudorapidity region up to $|\eta| < 2.1$. However, a shortfall of funds has led to the staging of the chambers sitting beyond $|\eta| > 1.6$. RPCs in the first endcap station are also used to help resolve ambiguities in the CSCs. There are 36 chambers mounted in each of 2 rings in each of the endcap stations.

1.5.3 Electromagnetic calorimeter

The Electromagnetic Calorimeter (ECAL) is a hermetic, homogeneous calorimeter comprising 61200 lead tungstate ($PbWO_4$) crystals mounted in the central barrel part, closed by 7324 crystals in each of the 2 endcaps.

CMS has chosen lead tungstate scintillating crystals for its ECAL. These crystals have short radiation (X_0 = 0.89 cm) and Moliere (2.2 cm) lengths, are fast (80% of the light is emitted within 25 ns) and radiation hard (up to 10 Mrad). However, the relatively low light yield (30 γ/MeV) requires use of photodetectors with intrinsic gain that can operate in a magnetic field. Silicon avalanche photodiodes (APDs) are used as photodetectors in the barrel and vacuum phototriodes (VPTs) in the endcaps. In addition, the sensitivity of both the crystals and the APD response to temperature changes requires a temperature stability (the goal is $0.1°$C). The use of $PbWO_4$ crystals has thus allowed the design of a compact calorimeter inside the solenoid that is fast, has fine granularity, and is radiation resistant.

The barrel section (EB) has an inner radius of 129 cm. It is structured as 36 identical "supermodules," each covering half the barrel length and corresponding to a pseudorapidity interval of $0 < |\eta| < 1.479$. The crystals are quasi-projective (the axes are tilted at $3°$ with respect to the line from the nominal vertex position) and cover 0.0174 (i.e. $1°$) in $\Delta\phi$ and $\Delta\eta$. The crystals have a front face cross-section of $\approx$ 22×22 mm^2 and a length of 230 mm, corresponding to 25.8 X_0.

The endcaps (EE), at a distance of 314 cm from the vertex and covering a pseudorapidity range of $1.479 < |\eta| < 3.0$, are each structured as 2 "Dees", consisting of semi-circular aluminium plates from which are cantilevered structural units of 5×5 crystals, known as "supercrystals." In the ECAL TDR [5] the basic mechanical unit was envisaged to hold 6×6 crystals. The change was accommodated by a corresponding increase in the lateral size of the crystals. The endcap crystals, like the barrel crystals, off-point from the nominal vertex position, but are arranged in an x-y grid (i.e. not an η-ϕ grid). They are all identical and have a front face cross section of 28.6×28.6 mm^2 and a length of 220 mm (24.7 X_0). A preshower device is placed in front of the crystal calorimeter over much of the endcap pseudorapidity range. The active elements of this device are 2 planes of silicon strip detectors, with a pitch of 1.9 mm, which lie behind disks of lead absorber at depths of 2 X_0 and 3 X_0.

1.5.3.1 Electronics readout

After amplification by a multi-gain preamplifier, the signal, shaped to peak after about 50 ns, is sampled and digitized at 40 MHz in 1 of 3 selected 12-bit ADCs used for each channel. A dynamic range of over 15 bits is attained. For each trigger, consecutive digitizations within a defined time frame (250 ns) are read out. In order to obtain the amplitude of a digitized pulse, the samples within the time frame are weighted and summed. The noise performance has been measured in several supermodules and found to be close to the original specification of approximately 40 MeV/channel.

1.5.3.2 Performance of the electromagnetic calorimeter

The performance of a supermodule was measured in a test beam. Representative results on the energy resolution as a function of beam energy are shown in Figure 1.7. The energy

resolution, measured by fitting a Gaussian function to the reconstructed energy distributions, has been parameterized as a function of energy:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2, \tag{1.2}$$

where S is the stochastic term, N the noise and C the constant term. The values of these parameters are listed in the figure.

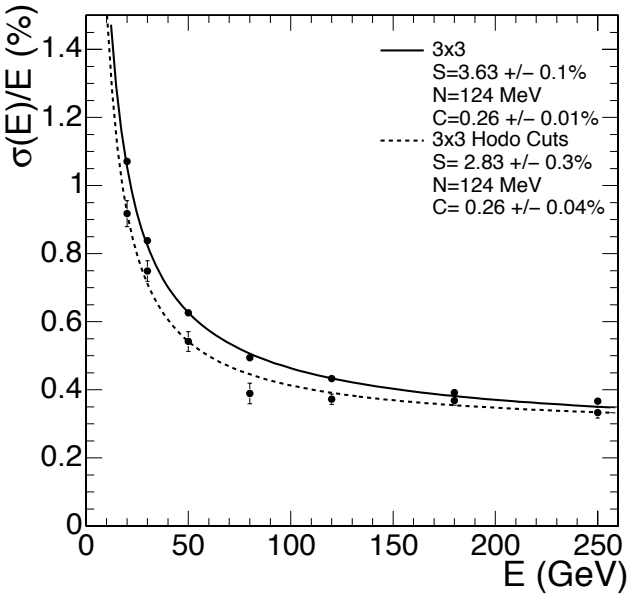


Figure 1.7: ECAL supermodule energy resolution, σ_E/E, as a function of electron energy as measured from a beam test. The upper series of points correspond to events taken with a 20×20 mm^2 trigger and reconstructed using a containment correction described in Section 4.3.2.2. The lower series of points correspond to events selected to fall within a 4×4 mm^2 region. The energy was measured in an array of 3×3 crystals with electrons impacting the central crystal.

1.5.4 Hadron calorimeter

The design of the hadron calorimeter (HCAL) [3] is strongly influenced by the choice of magnet parameters since most of the CMS calorimetry is located inside the magnet coil (Fig. CP 1) and surrounds the ECAL system. An important requirement of HCAL is to minimize the non-Gaussian tails in the energy resolution and to provide good containment and hermeticity for the $E_{\mathrm{T}}^{\mathrm{miss}}$ measurement. Hence, the HCAL design maximizes material inside the magnet coil in terms of interaction lengths. This is complemented by an additional layer of scintillators, referred to as the hadron outer (HO) detector, lining the outside of the coil. Brass has been chosen as absorber material as it has a reasonably short interaction length, is easy to machine and is non-magnetic. Maximizing the amount of absorber before the magnet requires keeping to a minimum the amount of space devoted to the active medium. The tile/fibre technology makes for an ideal choice. It consists of plastic scintillator tiles read out with embedded wavelength-shifting (WLS) fibres. The WLS fibres are spliced to high-

attenuation-length clear fibres outside the scintillator that carry the light to the readout system. This technology was first developed by the UA1 collaboration [14] and at Protvino [15] and has been used in the upgrade of the CDF endcap calorimeter [16]. The photodetection readout is based on multi-channel hybrid photodiodes (HPDs). The absorber structure is assembled by bolting together precisely machined and overlapping brass plates so as to leave space to insert the scintillator plates, which have a thickness of 3.7 mm. The overall assembly enables the HCAL to be built with essentially no uninstrumented cracks or dead areas in ϕ. The gap between the barrel and the endcap HCAL, through which the services of the ECAL and the inner tracker pass, is inclined at 53° and points away from the centre of the detector.

1.5.4.1 Hadron barrel

The hadron barrel (HB) part of HCAL consists of 32 towers covering the pseudorapidity region $-1.4 < \eta < 1.4$, resulting in 2304 towers with a segmentation $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$. The HB is constructed in 2 half barrels (see Fig. CP 8). Details of the HB design, together with the performance of production modules measured in CERN test beams, may be found in [17]. The HB is read out as a single longitudinal sampling. There are 15 brass plates, each with a thickness of about 5 cm, plus 2 external stainless steel plates for mechanical strength. Particles leaving the ECAL volume first see a scintillator plate with a thickness of 9 mm rather than 3.7 mm for the other plates. The light collected by the first layer is optimized to be a factor of about 1.5 higher than the other scintillator plates.

1.5.4.2 Hadron outer

The hadron outer (HO) detector contains scintillators with a thickness of 10 mm, which line the outside of the outer vacuum tank of the coil and cover the region $-1.26 < \eta < 1.26$. The tiles are grouped in 30°-sectors, matching the ϕ segmentation of the DT chambers. They sample the energy from penetrating hadron showers leaking through the rear of the calorimeters and so serve as a "tail-catcher" after the magnet coil. They increase the effective thickness of the hadron calorimetry to over 10 interaction lengths, thus reducing the tails in the energy resolution function. The HO also improves the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution of the calorimeter.

HO is physically located inside the barrel muon system and is hence constrained by the geometry and construction of that system. It is divided into 5 sections along η, called "rings" -2, -1, 0, 1, and 2. The fixed ring-0 has 2 scintillator layers on either side of an iron absorber with a thickness of about 18 cm, at radial distances of 3.850 m and 4.097 m, respectively. The other mobile rings have single layers at a radial distance of 4.097 m. Each ring covers 2.5 m in z. HO scintillators follow the HCAL barrel tower geometry in η and ϕ.

1.5.4.3 Hadron endcap

Each hadron endcap (HE) of HCAL consists of 14 η towers with 5° ϕ segmentation, covering the pseudorapidity region $1.3 < |\eta| < 3.0$. For the 5 outermost towers (at smaller η) the ϕ segmentation is 5° and the η segmentation is 0.087. For the 8 innermost towers the ϕ segmentation is 10°, whilst the η segmentation varies from 0.09 to 0.35 at the highest η. The total number of HE towers is 2304.

Details of the HE design, together with the performance of production modules measured in CERN test beams, may be found in [18].

1.5.4.4 Hadron forward

Coverage between pseudorapidities of 3.0 and 5.0 is provided by the steel/quartz fibre Hadron Forward (HF) calorimeter. Because the neutral component of the hadron shower is preferentially sampled in the HF technology, this design leads to narrower and shorter hadronic showers and hence is ideally suited for the congested environment in the forward region. The front face is located at 11.2 m from the interaction point. The depth of the absorber is 1.65 m. The signal originates from Cerenkov light emitted in the quartz fibres, which is then channeled by the fibres to photomultipliers. The absorber structure is created by machining 1 mm square grooves into steel plates, which are then diffusion welded. The diameter of the quartz fibres is 0.6 mm and they are placed 5 mm apart in a square grid. The quartz fibres, which run parallel to the beam line, have two different lengths (namely 1.43 m and 1.65 m) which are inserted into grooves, creating 2 effective longitudinal samplings. There are 13 towers in η, all with a size given by $\Delta\eta \approx 0.175$, except for the lowest-η tower with $\Delta\eta \approx 0.1$ and the highest-η tower with $\Delta\eta \approx 0.3$. The ϕ segmentation of all towers is 10°, except for the highest-η one which has $\Delta\phi$ = 20°. This leads to 900 towers and 1800 channels in the 2 HF modules.

1.5.4.5 Performance of the hadron calorimeter

In Chapter 11 the simulated single particle energy response and resolution is compared with test beam data from all 3 geographic parts of the HCAL. However, for gauging the performance of the HCAL, it is usual to look at the jet energy resolution and the missing transverse energy resolution. The granularity of the sampling in the 3 parts of the HCAL has been chosen such that the jet energy resolution, as a function of E_{T}, is similar in all 3 parts. This is illustrated in Figure 1.8. The missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$) resolution is given by $\sigma(E_{\mathrm{T}}^{\mathrm{miss}}) \approx 1.25\sqrt{\Sigma E_{\mathrm{T}}}$, if energy clustering corrections are not made.

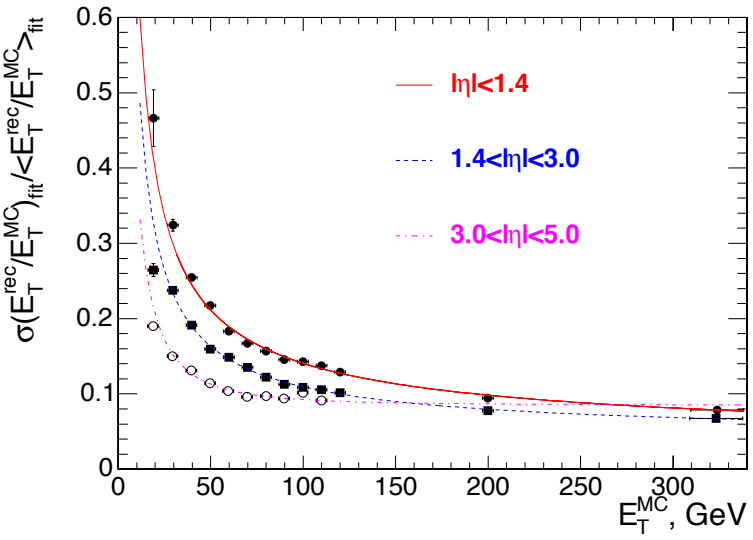


Figure 1.8: The jet transverse energy resolution as a function of the simulated jet transverse energy for barrel jets ($|\eta|$ <1.4), endcap jets (1.4< $|\eta|$ <3.0) and very forward jets (3.0< $|\eta|$ <5.0). The jets are reconstructed with the interative cone $R = 0.5$ algorithm. See Section 11.4 for further details.

1.5.5 Inner tracking system

By considering the charged particle flux at various radii at high luminosity (Table 1.6), 3 regions can be delineated:

- Closest to the interaction vertex where the particle flux is the highest ($\approx 10^7$/s at $r \approx 10$ cm), pixel detectors are placed. The size of a pixel is $\approx 100 \times 150\ \mu\mathrm{m}^2$, giving an occupancy of about 10^{-4} per pixel per LHC crossing.
- In the intermediate region ($20 < r < 55$ cm), the particle flux is low enough to enable use of silicon microstrip detectors with a minimum cell size of 10 cm $\times 80\ \mu\mathrm{m}$, leading to an occupancy of $\approx$ 2–3%/LHC crossing.
- In the outermost region ($r > 55$ cm) of the inner tracker, the particle flux has dropped sufficiently to allow use of larger-pitch silicon microstrips with a maximum cell size of 25 cm $\times$ 180 μm, whilst keeping the occupancy to $\approx 1\%$.

Even in heavy-ion (Pb-Pb) running, the occupancy is expected to be at the level of 1% in the pixel detectors and less than 20% in the outer silicon strip detectors, permitting track reconstruction in the high density environment.

The layout of the CMS tracker is shown in Figure 1.9. The outer radius of the CMS tracker extends to nearly 110 cm, and its total length is approximately 540 cm.

Close to the interaction vertex, in the barrel region, are 3 layers of hybrid pixel detectors at a radii of $4, 7$, and 11 cm. The size of the pixels is $100\times150\ \mu\mathrm{m}^2$. In the barrel part, the silicon microstrip detectors are placed at r between 20 and 110 cm. The forward region has 2 pixel and 9 microstrip layers in each of the 2 Endcaps. The barrel part is separated into an Inner and an Outer Barrel. In order to avoid excessively shallow track crossing angles, the Inner Barrel is shorter than the Outer Barrel, and there are an additional 3 Inner Disks in the transition region between the barrel and endcap parts, on each side of the Inner Barrel. The total area of the pixel detector is ≈ 1 m^2, whilst that of the silicon strip detectors is 200 m^2, providing coverage up to $|\eta| < 2.4$. The inner tracker comprises 66 million pixels and 9.6 million silicon strips [19].

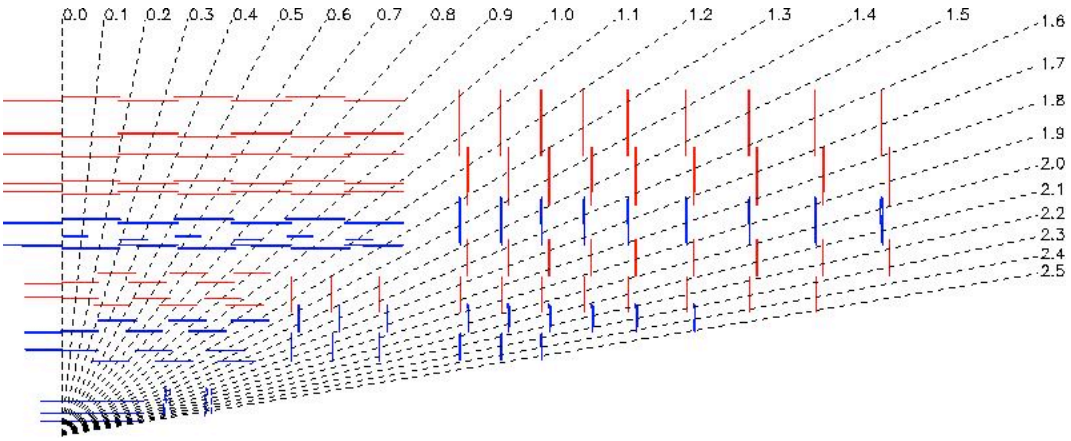


Figure 1.9: The tracker layout (1/4 of the z view).

1.5.5.1 Strip tracker

The barrel tracker region is divided into 2 parts: a TIB (Tracker Inner Barrel) and a TOB (Tracker Outer Barrel). The TIB is made of 4 layers and covers up to $|z| < 65$ cm, using silicon sensors with a thickness of 320 μm and a strip pitch which varies from 80 to 120 μm. The first 2 layers are made with "stereo" modules in order to provide a measurement in both r-ϕ and r-z coordinates. A stereo angle of 100 mrad has been chosen. This leads to a single-point resolution of between 23–34 μm in the r-ϕ direction and 23 μm in z. The TOB comprises 6 layers with a half-length of $|z| < 110$ cm. As the radiation levels are smaller in this region, thicker silicon sensors (500 μm) can be used to maintain a good S/N ratio for longer strip length and wider pitch. The strip pitch varies from 120 to 180 μm. Also for the TOB the first 2 layers provide a "stereo" measurement in both r-ϕ and r-z coordinates. The stereo angle is again 100 mrad and the single-point resolution varies from 35–52 μm in the r-ϕ direction and 52 μm in z.

The endcaps are divided into the TEC (Tracker End Cap) and TID (Tracker Inner Disks). Each TEC comprises 9 disks that extend into the region 120 cm$< |z| < 280$ cm, and each TID comprises 3 small disks that fill the gap between the TIB and the TEC. The TEC and TID modules are arranged in rings, centred on the beam line, and have strips that point towards the beam line, therefore having a variable pitch. The first 2 rings of the TID and the innermost 2 rings and the fifth ring of the TEC have "stereo" modules. The thickness of the sensors is 320 μm for the TID and the 3 innermost rings of the TEC and 500 μm for the rest of the TEC.

The entire silicon strip detector consists of almost $15\,400$ modules, which will be mounted on carbon-fibre structures and housed inside a temperature controlled outer support tube. The operating temperature will be around -20°C.

Table 1.4: Detector types in the silicon tracker

part	No. detectors	thickness (μm)	mean pitch (μm)
TIB	2724	320	81/118
TOB	5208	500	81/183
TID	816	320	97/128/143
TEC	2512	320	96/126/128/143
TEC(2)	3888	500	143/158/183

1.5.5.2 Pixel tracker

The pixel detector consists of 3 barrel layers with 2 endcap disks on each side on them (Fig. 1.10). The 3 barrel layers are located at mean radii of 4.4 cm, 7.3 cm and 10.2 cm, and have a length of 53 cm. The 2 end disks, extending from 6 to 15 cm in radius, are placed on each side at $|z| = 34.5$ cm and 46.5 cm.

In order to achieve the optimal vertex position resolution, a design with an "almost" square pixel shape of 100×150 μm^2 in both the (r,ϕ) and the z coordinates has been adopted. The barrel comprises 768 pixel modules arranged into half-ladders of 4 identical modules each. The large Lorentz effect (Lorentz angle is $23°$) improves the r-ϕ resolution through charge sharing.

The endcap disks are assembled in a turbine-like geometry with blades rotated by $20°$ to also benefit from the Lorentz effect. The endcap disks comprise 672 pixel modules with 7 different modules in each blade.

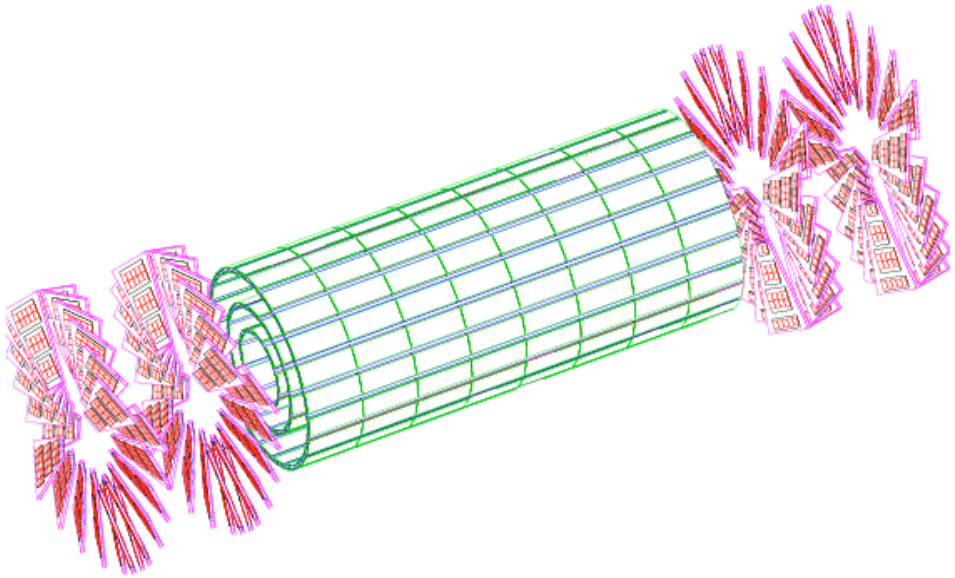


Figure 1.10: Layout of pixel detectors in the CMS tracker.

The spatial resolution is measured to be about $10\ \mu\mathrm{m}$ for the r-ϕ measurement and about $20\ \mu\mathrm{m}$ for the z measurement. The detector is readout using approximately $16\,000$ readout chips, which are bump-bonded to the detector modules.

1.5.5.3 Tracker control and readout scheme

The Silicon Strip Tracker (SST) readout system is based on a front-end APV25 readout chip [20], analogue optical links [21] and an off-detector Front-End Driver (FED) processing board [22]. The APV25 chip samples, amplifies, buffers and processes signals from 128 channels of a silicon strip sensor. Each microstrip is readout by a charge sensitive amplifier with τ = 50 ns. The output voltage is sampled at the beam crossing rate of 40 MHz. Samples are stored in an analogue pipeline for up to the Level-1 latency of $3.2\ \mu\mathrm{s}$. Following a trigger, a weighted sum of 3 samples is formed in an analogue circuit. This confines the signal to a single bunch crossing and gives the pulse height. The buffered pulse height data from pairs of APV25 chips are multiplexed onto a single line and the analogue data are converted to optical signals before being transmitted via optical fibres to the off-detector FED boards. The output of the transmitting laser is modulated by the pulse height for each strip. The FEDs digitize, process and format the pulse height data from up to 96 pairs of APV25 chips, before forwarding zero-suppressed data to the DAQ online farm. The electronics noise/channel of the tracking system is about 1000 to 1500 electrons before and after irradiation, respectively. The SST control system comprises ≈300 control rings that start and end at the off-detector Front-End Controller (FEC) boards [23]. Slow-control commands, clock and Level-1 triggers are distributed via digital optical links to Digital Opto-Hybrids (DOH) [24], which perform optical-to-electrical conversion before the control signals are distributed to the front-end electronics.

The Pixel Tracker readout system is described in detail in [6]. A single pixel barrel module is readout by 16 Read-Out Chips (ROCs). In the endcaps, the number of ROCs per module varies from 2 to 10. Each ROC reads an array of 52×80 pixels. Analogue signals and corresponding pixel addresses are stored in a data buffer, waiting for the Level-1 trigger decision.

Following a Level-1 trigger accept, data are transmitted on optical links to FED boards. In the barrel, groups of 8 or 16 ROCs are connected to 1 link, whereas in the endcaps there are 21 or 24 ROCs per link. The 40 Pixel FEDs perform digitization and data formatting.

1.5.5.4 Performance of the tracker

The performance of the tracker is illustrated in Figure 1.11, which shows the transverse momentum and impact parameter resolutions in the r-ϕ and z planes for single muons with a p_{T} of 1, 10 and 100 GeV/c, as a function of pseudorapidity. Track reconstruction efficiency as a function of pseudorapidity for single muons and pions is shown in Figure 1.12. The material inside the active volume of the tracker increases from $\approx 0.4X_0$ at $\eta = 0$ to around $1X_0$ at $|\eta| \approx 1.6$, before decreasing to $\approx 0.6X_0$ at $|\eta| = 2.5$.

1.5.6 Trigger and data acquisition

The LHC bunch crossing rate of 40 MHz leads to $\approx 10^9$ interactions/sec at design luminosity. Data from only about 10^2 crossings/sec can be written to archival media; hence, the trigger system has to achieve a rejection factor of nearly 10^6.

The CMS trigger and data acquisition system [7, 8] consists of 4 parts: the detector electronics, the Level-1 trigger processors (calorimeter, muon, and global), the readout network, and an online event filter system (processor farm) that executes the software for the High-Level Triggers (HLT).

1.5.6.1 Level-1 trigger

The size of the LHC detectors and the underground caverns imposes a minimum transit time for signals from the front-end electronics to reach the services cavern housing the Level-1 trigger logic and return back to the detector front-end electronics. The total time allocated for the transit and for reaching a decision to keep or discard data from a particular beam crossing is $3.2\ \mu\mathrm{s}$. During this time, the detector data must be held in buffers while trigger data is collected from the front-end electronics and decisions reached that discard a large fraction of events while retaining the small fraction of interactions of interest (nearly 1 crossing in 1000). Of the total latency, the time allocated to Level-1 trigger calculations is less than $1\ \mu\mathrm{s}$.

Custom hardware processors form the Level-1 decision. The Level-1 triggers involve the calorimetry and muon systems, as well as some correlation of information between these systems. The Level-1 decision is based on the presence of "trigger primitive" objects such as photons, electrons, muons, and jets above set E_{T} or p_{T} thresholds. It also employs global sums of E_{T} and $E_{\mathrm{T}}^{\mathrm{miss}}$. Reduced-granularity and reduced-resolution data are used to form trigger objects. At startup the Level-1 rate will be limited to 50 kHz (the design value is 100 kHz). Taking a safety margin of a factor of 3 into account for simulation uncertainties, as well as beam and detector conditions not included in the simulation programs, leads to an estimated rate of 16 kHz. The design value of 100 kHz is set by the average time to transfer full detector information through the readout system.

Much of the logic in the trigger system is contained in custom Application Specific Integrated Circuits (ASICs), semi-custom and gate-array ASICs, Field Programmable Gate Arrays (FPGAs), Programmable Logic Devices (PLDs), and discrete logic such as Random Access Memories that are used for memory Look-Up Tables (LUTs). Where possible and where the added

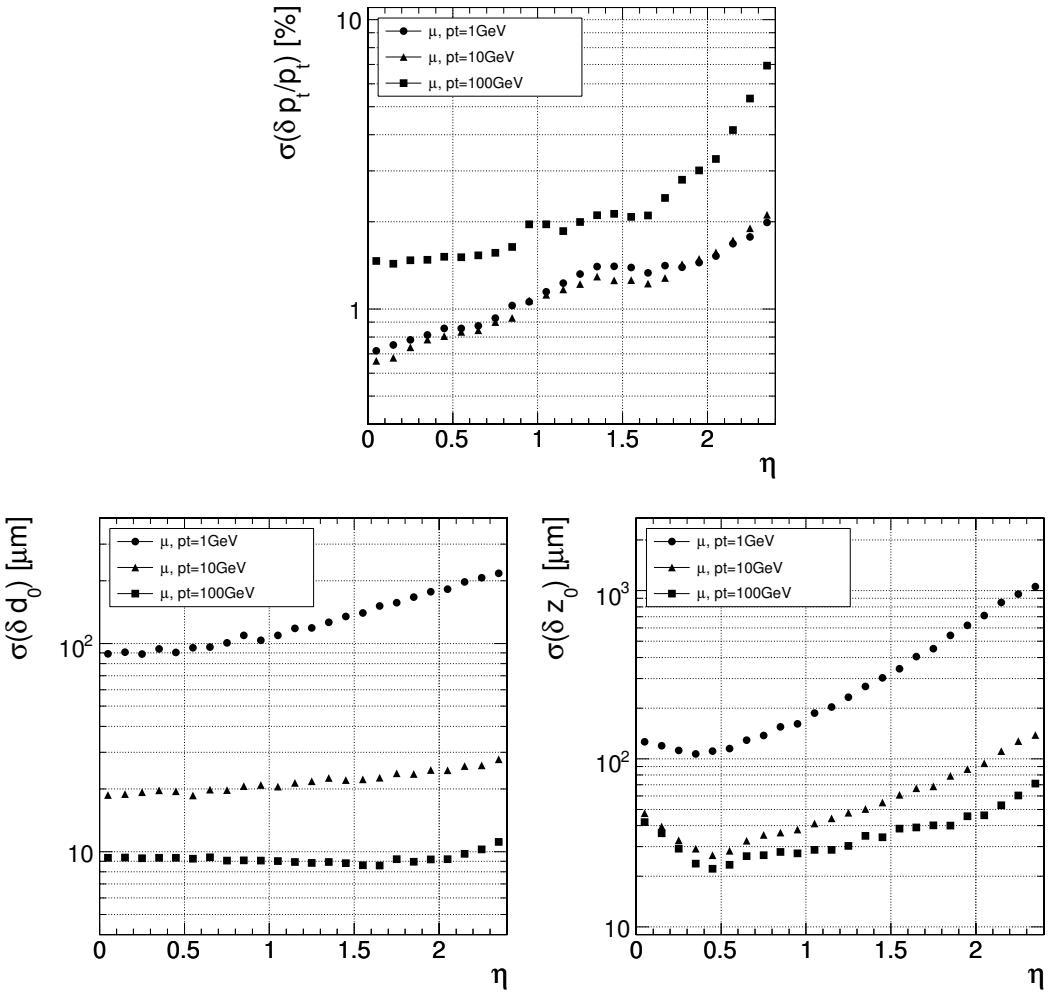


Figure 1.11: Resolution of several track parameters for single muons with transverse momenta of 1, 10 and 100 GeV/c: (upper) transverse momentum, (lower left) transverse impact parameter, and (lower right) longitudinal impact parameter.

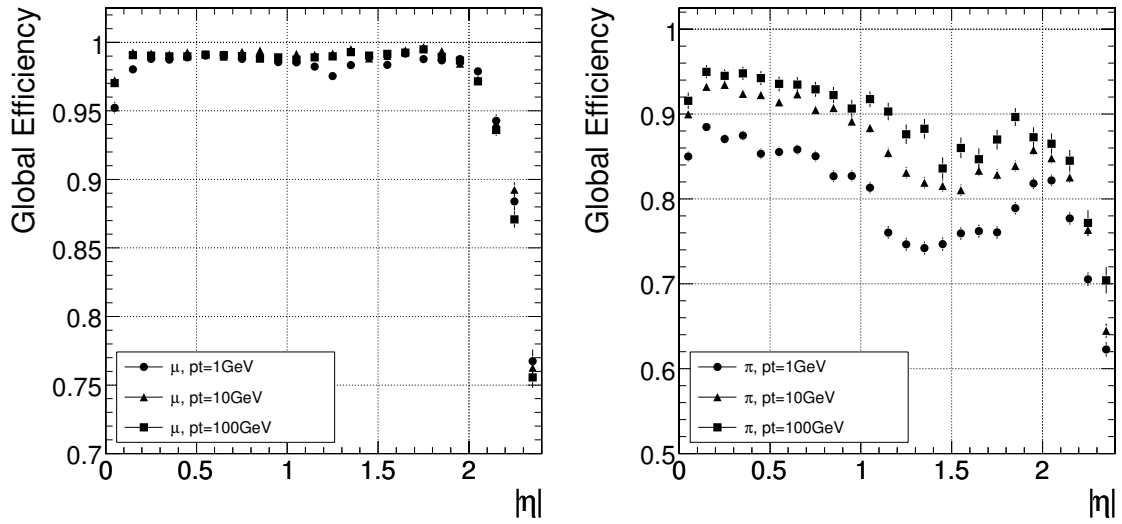


Figure 1.12: Global track reconstruction efficiency for muons (left) and pions (right) of transverse momenta of 1, 10 and 100 GeV/c.

flexibility offers an advantage and is cost effective, designs incorporate new FPGA technology.

During the Level-1 decision-making period, all the high-resolution data is held in pipelined memories. Commodity computer processors make subsequent decisions using more detailed information from all of the detectors in more and more sophisticated algorithms that approach the quality of final reconstruction.

1.5.6.2 High-Level triggers

Upon receipt of a Level-1 trigger, after a fixed time interval of about 3.2 μs, the data from the pipelines are transferred to front-end readout buffers. After further signal processing, zero-suppression and/or data-compression, the data are placed in dual-port memories for access by the DAQ system. Each event, with a size of about 1.5 MB (pp interactions), is contained in several hundred front-end readout buffers. Through the event building "switch," data from a given event are transferred to a processor. Each processor runs the same high-level trigger (HLT) software code to reduce the Level-1 output rate of 100 kHz to 100 Hz for mass storage.

The use of a processor farm for all selections beyond Level-1 allows maximal benefit to be taken from the evolution of computing technology. Flexibility is maximized since there is complete freedom in the selection of the data to access, as well as in the sophistication of the algorithms.

Various strategies guide the development of the HLT code. Rather than reconstruct all possible objects in an event, whenever possible only those objects and regions of the detector that are actually needed are reconstructed. Events are to be discarded as soon as possible. This leads to the idea of partial reconstruction and to the notion of many virtual trigger levels, e.g., calorimeter and muon information are used, followed by use of the tracker pixel data and finally the use of the full event information (including full tracking).

1.5.7 Software and computing

The CMS software and computing systems need to cover a broad range of activities including the design, evaluation, construction, and calibration of the detector; the storage, access, reconstruction and analysis of data; and the support of a distributed computing infrastructure for physicists engaged in these tasks.

The storage, networking and processing power needed to analyse these data is well in excess of today's facilities and exceed any reasonably projected capabilities of CERN's central computing systems. The CMS computing model [25] is therefore highly distributed, with a primary "Tier-0" centre at CERN being supplemented by Tier-1 and Tier-2 computing centres at national laboratories and universities worldwide.

New computing grid technologies will be used to facilitate the seamless exploitation of these distributed centres. Close collaboration is maintained with running HEP experiments to learn from their experience and adopt and extend appropriate computing technologies they have developed. The "LHC Computing Grid," a joint project of the experiments and laboratories, is proceeding towards integration and deployment of grid technologies for LHC.

Much software has to be developed and verified for simulation and physics analysis, as well as for common libraries, tools and frameworks. All of this has to be ready and tested by the start of data taking in 2007. To this end, "Data Challenges" of increasing size and complexity are being performed by CMS in order to subject the steadily developing computing environment and software to progressively more realistic tests. Of particular note is the Computing, Software and Analysis challenge foreseen in the second half of 2006.

1.5.8 Performance: mass resolutions

An indication of the performance of CMS can be obtained by looking at the mass resolution for various states (Table 1.5).

Table 1.5: Mass resolution for various states at a luminosity of $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

	CMS GeV/c^2
$B \to \pi\pi$	0.031
$B \to J/\Psi K^0_S$	0.016
$\Upsilon \to \mu^+\mu^-$	0.050
$H(130\ \mathrm{GeV}/c^2) \to \gamma\gamma$	0.90
$H(150\ \mathrm{GeV}/c^2) \to ZZ^* \to 4\mu$	1.3
$A(500\ \mathrm{GeV}/c^2) \to \tau\tau$	75.0
$W \to$ jet-jet	10.0
$Z'(1\ \mathrm{TeV}/c^2) \to \mu^+\mu^-$	45.0
$Z'(1\ \mathrm{TeV}/c^2) \to e^+e^-$	5.0

1.5.9 Radiation levels in CMS

The high particle fluxes emanating from the interaction region lead to high radiation levels requiring radiation hard detectors and front-end electronics. The expected radiation levels in CMS for the various detectors are given in Tables 1.6 and 1.7 for an integrated luminosity of 500 fb^{-1}, corresponding to the first 10 years of running.

Table 1.6: Hadron fluence and radiation dose in different radial layers of the CMS Tracker (barrel part) for an integrated luminosity of 500 fb^{-1} (≈ 10 years).

Radius (cm)	Fluence of fast hadrons (10^{14} cm^{-2})	Dose (kGy)	Charged Particle Flux (cm^{-2}s^{-1})
4	32	840	10^8
11	4.6	190	
22	1.6	70	6×10^6
75	0.3	7	
115	0.2	1.8	3×10^5

Table 1.7: Radiation dose in CMS Calorimeters for an integrated luminosity of 500 fb^{-1} (≈ 10 years).

Pseudorapidity (η)	ECAL Dose (kGy)	HCAL Dose (kGy)	ECAL Dose Rate (Gy/h)
0–1.5	3	0.2	0.25
2.0	20	4	1.4
2.9	200	40	14
3.5	–	100	–
5	–	1000	–

1.6 The current state of the construction, installation and commissioning of CMS

The CMS detector is being assembled in the surface assembly hall, SX5, at Point 5 in Cessy.

For ease of assembly, installation and maintenance, CMS has been designed along the following lines:

- All the subdetectors should be maintainable by opening CMS in large sections,
- The movements of these sections should be possible without un-cabling the attached subdetectors and without breaking the chain of services (cooling, LV, etc.) allowing fast re-commissioning after closing,
- access must be possible to the Tracker flange for at least 1 day during a 10-day shut-down.

The main design consequences are as follows:

- The Barrel yoke is sectioned into 5 ring-sections and each Endcap yoke into 3 disk-sections (4 with YE4) to allow maintenance of the Muon stations.
- The Hadron Forward (HF) calorimeters have been pushed outside of the yoke to allow easy sliding of the Endcaps along the beam-pipe.
- This allows the beam-pipe to flare out, which reduces the background from secondary interactions, and improves pumping without having to put a vacuum pump too near the interaction point.

Civil Engineering (CE) work at Point 5 (located at Cessy, France) has finished. Installation of the counting room and service infrastructure in cavern USC55 is advancing well. The counting room is expected to be ready for subdetector readout crate installation in the second quarter of 2006. The experiment cavern will be ready to receive detector elements in May 2006.

The magnet yoke has been assembled at Point 5 for some time. The assembly of the solenoid (coil integrated into the vacuum tank) is expected to be finished in January 2006 and cool-down started. The coil will be powered to full current in the second quarter of 2006.

There are 3 main phases in the assembly of the Tracker: module production, assembly of modules into TOB-rods, TIB-shells, and TEC-petals, and finally the integration of TOB rods, complete TIB and complete TEC structures into the Tracker Support Tube. The first phase is almost complete, the second is progressing and the third is expected to be completed by the summer of 2006. The commissioning of the full tracker, 25% at a time, is foreseen in the autumn of 2006, before transport to Point 5 at the end of 2006.

The delivery of crystals defines the critical path for the ECAL. At the time of writing, some $46\,000$ out of $61\,200$ barrel crystals have been delivered. The serial integration of electronics into bare supermodules (SMs) has started. These SMs comprise 1700 crystals, which are tested after assembly in the laboratory for a period of 1 week and with cosmic rays for another 1 week. The first half-barrel should be lowered into UX5 by mid-2006. The whole barrel ECAL should be ready for the pilot physics run in 2007.

All HCAL module types [HB (barrel), HE (endcap), HO (outer) and HF (forward)], including absorber and optics, are completed. Photodetectors and electronics have been installed or are being installed prior to a comprehensive calibration of HCAL using Co-60 sources. HF will be the first subdetector to be lowered in May of 2006.

All of the CSC chambers (a total of 496, including 6% spares) are at CERN. At the time of writing, some 342 (75%) have been installed on the magnet yoke disks; almost all of the installed chambers have been fully commissioned using cosmic rays.

The assembly of DT chambers should be completed by the beginning of 2006, except for the MB4 chambers which will be finished by April 2006. The installation of DT chambers in yoke wheels YB+2 and YB+1 (84 DT and RPC packages) is complete. The installation will be completed in the other wheels before they are lowered. The commissioning of YB+2 chambers using cosmic rays has been finished and the cabling operation has just started.

The installation of RE1 chambers should be completed by the end of this year.

The muon-alignment hardware for the magnet test is in production. All MABs (carbon-fibre position reference structures) are at CERN. Installation and cabling of components has started on YE+.

The trigger system is mostly in production. Production of some trigger components has already finished. There is much work underway on software and firmware. Integration tests of detector primitive generators, trigger system and DAQ are underway. Components of the trigger system are being thoroughly exercised and integrated with other trigger and detector electronics systems in the Electronics Integration Centre (Building 904 on the Prevessin site) in preparation for installation in USC55. Some components of the final system will be used in data taking during the magnet test at SX5 in early 2006.

The production of the Data-to-Surface (D2S) custom components has been completed and the electrical and functional tests have been passed. The final system tests have started in Building 904. The D2S FED builder system (1024 Myrinet 2XP link boards, 10 Myrinet crates and switches) has been delivered, and part of the modules have been installed in the pre-series system at Point 5. The D2S components (except PC and fibres) are ready for installation and the start of the readout commissioning in USC. A 16×16 readout builder with 16 FRL-

FED columns is permanently running in Point 5 and used to test the integration of the central DAQ with the detector systems participating in the Magnet Test.

In the current CMS Master Schedule, v34.2, the initial detector is foreseen to be installed and closed for beam on 30 June 2007 and be ready for first collisions in late summer of 2007. Installation of the pixel tracker, although ready in summer 2007, and the ECAL endcaps is foreseen during the 2007/2008 winter shutdown, in time for the first physics run in spring 2008. The staged items include part of the fourth endcap muon station ME4/2, RPC chambers at low angles ($|\eta| > 1.6$), 60% of the DAQ online farm and the third pair of forward pixel disks.

Chapter 2

Software Components

The reconstruction algorithms described in this Volume (detector chapters) as well as the simulation tools used to study the performance of these algorithms have been implemented in the object-oriented framework COBRA [26]. The collection of detector simulation and reconstruction code developed in this framework is known as ORCA [27], and with COBRA, has served the simulation and design needs of the Collaboration since development began in 1998, including studies reported in previous Technical Design Reports [7, 8].

In 2003 we completed a transition from a detector simulation based on GEANT3 [28], a package known as CMSIM [29], to an object-oriented version based on GEANT4 [30], a package called OSCAR [31]. The simulated samples of LHC collisions and detector effects using OSCAR make up most of the samples used for the studies described in this Report.

Recently, a comprehensive set of changes have been made in the underlying software framework, the services it provides, and in the model for data storage to prepare the experiment for LHC data-taking. These changes address additional requirements on the software to implement calibration and alignment strategies, ensure tracking and reproducibility of the reconstruction results, simplify and standardize the way physicists develop reconstruction algorithms, and facilitate interactive analysis (where we aim to eliminate the need for post-reconstruction "ntuples").

The overall collection of software, now referred to as CMSSW [32], is built around a Framework, an Event Data Model, and Services needed by the simulation, calibration and alignment, and reconstruction modules that process event data so that physicists can perform analysis. Already several instances of the software and analysis reported in the two Volumes of this Technical Design Report have migrated to this newer software foundation.

This chapter thus describes the software components in development or migration to prepare CMS for LHC collisions in 2007, and lays the foundation for the detector operation and reconstruction methodologies described in later chapters.

2.1 CMS software architecture

The high-level goals of the CMS software are to process and select events inside the High Level Trigger Farm [8], to deliver the processed results to experimenters within the CMS Collaboration, and to provide tools for them to analyze the processed information in order to produce physics results.

Many technical requirements should be considered such as the memory consumption and

processing time per event necessary to meet financial constraints, and the physics performance requirements encompassing the ability to reproduce faithfully details of the underlying physics processes based on the detector data.

2.1.1 Requirements

CMS has identified [33] the following underlying principles which motivate the overall design of the CMS software architecture.

- **Multiple Environments**: Various software modules must be able to run in a variety of environments as different computing tasks are performed. Examples of these environments are High Level triggering, production reconstruction, program development, and individual analysis;
- **Migration between environments**: a particular software module may start out being developed for one environment, then later used in other unforeseen environments as well;
- **Migration to new technologies**: hardware and software technologies will change during the lifetime of the experiment: a migration to a new technology should require a finite effort localized in a small portion of the software system, ideally without involving changes to physics software modules;
- **Dispersed code development**: the software will be developed by organizationally and geographically dispersed groups of experimenters;
- **Flexibility**: not all software requirements will be fully known in advance, therefore the software systems must be adaptable without requiring total rewrites;
- **Ease of use**: the software systems must be easily usable by collaboration physicists who are not computing experts and cannot devote large amounts of time to learning computing techniques.

These requirements imply that software should be developed keeping in mind not only performance but also modularity, flexibility, maintainability, quality assurance and documentation. CMS has adopted an object-oriented development methodology, based primarily on the C++ programming language.

2.1.2 Architecture design

The requirements described above (see Section 2.1.1) on the software architecture result in the following overall structure for the CMS software:

- an application framework customizable for each of the computing environments;
- physics software modules with clearly defined interfaces that can be plugged into the framework;
- services and utility toolkits that can be used by any of the physics modules.

The framework defines the top level abstractions, their behavior and collaboration patterns. It comprises two components: a set of classes that capture CMS specific concepts like detector components and event features and a control policy that orchestrates the instances of those classes taking care of the flow of control, module scheduling, input/output, etc. This control policy is tailored to the task at hand and to the computing environment.

The physics and utility modules are written by detector groups. The modules can be plugged into the application framework at run time, independently of the computing environment. One can easily choose between different versions of various modules. The physics modules do not communicate with each other directly but only through the data access protocols that are part of the framework itself.

The service and utility toolkit consists of two major categories of services: physics type services (histogrammers, fitters, mathematical algorithms, geometry and physics calculation routines) and computer services (data access, inter module communication, user interface, etc.). This toolkit is based on LCG Application Area components [34]. Both the application framework and the service and utility toolkit shield the physics software modules from the underlying technologies which will be used for the computer services. This will ensure a smooth transition to new technologies with changes localized in the framework and in specific components of the service toolkit.

We are starting to put in place a more robust and reliable development process with emphasis on guidance and use of automated tools and extensive testing and validation releases.

2.1.3 Software domain decomposition

Several orthogonal domain decompositions have been identified.

The first spans between the primary data-processing activities:

- The Event Farm and High-Level triggering;
- Simulation including event generation, pile-up and digitization;
- Local and Global Reconstruction;
- Calibration and Alignment processing;
- High Level Objects (muons, electrons, jets, etc.);
- Physics tools and visualization;
- Physics and Data Quality Monitoring.

The second domain decomposition spans the CMS subdetectors. These are described in the different detector chapters of this volume, and include:

- Tracker detector, which itself is comprised of the pixel detector, the silicon strip detector;
- Electromagnetic Calorimeter;
- Hadron Calorimeter;
- Muon Detectors (DT, CSC, RPC);
- Forward detectors.

Further domain decompositions cover the infrastructure needed to support the software development activity, and components that derive from the overall architectural vision.

The overall structure of the CMS Software Project is the following:

- Framework and Services. A common event processing framework has been adopted based on plug-compatible components and abstract interfaces. A variety of services such as scripting, magnetic field, etc. have been developed in support of this

framework. The framework is described in Section 2.2. Services are described in detail in Sections 2.2.3, 2.4, and 2.10.

- Event generators, simulation, pile-up and digitization. These are described in Section 2.5.
- Event selection and reconstruction. These are described in Section 2.7.
- Calibration and alignment. The high-level requirements for subdetector calibration and alignment as well as the DB infrastructures are described in Section 2.8.
- Event Filter and Data Quality Monitoring. These are described in Sections 2.3 and 2.9
- Software Development tools and validation. This task provides the services that support the software development process including the source code repository, tools for building coherent releases of the software, and for testing and validating those releases, and documentation tools and services. These are described in detail in Section 2.11.

2.2 Framework

The primary goal of the CMS Framework and Event Data Model (EDM) is to facilitate the development and deployment of reconstruction and analysis software. Ease-of-use is very high in the design priorities; for example, developers are encouraged to make their stored data objects as simple as possible in order to remove the need for specialized *private* analysis data formats and the code used to make them. All reconstruction results should be made persistent in a format that can be directly used by an analysis without any additional layer.

Another way to ensure consistency and ease of use of the software is automation: the Framework provides ways to guarantee reproducibility by automatically maintaining and recording sufficient *provenance* information for all application results. This avoids developers having to record such information separately on how a particular reconstruction result was obtained.

This section is organized as follows: Section 2.2.1 will focus on the Event Data Model and Section 2.2.2 will describe the main features of the framework. Finally in Section 2.2.3 we will describe the various services associated with the framework (i.e. geometry, non-event data, magnetic field, etc.)

2.2.1 Event Data Model–data access

The event data model is centered around the "Event." The Event holds all data that was taken during a triggered *physics event* as well as all data derived from the taken data. Auxiliary information needed to process an Event is accessed via the "EventSetup."

Events are processed by passing the Event through a sequence of modules. The exact sequence of modules is specified by the user. When an Event is passed to a module, that module can get data from the Event and put data back into the Event. When data is put into the Event, the provenance information about the module that created the data will be stored with the data into the Event.

The Event class[1] represents the observed and inferred products of a single triggered readout of the CMS detector. The Event is responsible for managing the lifetime of, and relationships between, it contents. The contents of the Event can include objects representing the raw detector output, reconstruction products, simulation products, and analysis objects relating to a single beam crossing or simulation thereof. The Event also contains "metadata" describing the configuration of the software used for the reconstruction of each contained data object and the conditions and calibration data used for such reconstruction.

2.2.2 Framework-module types and communication

The purpose of a "module" is to allow independent development and verification of distinct elements of triggering, simulation, reconstruction, and analysis. The concept of an event-processing module, each of which encapsulates a unit of clearly defined event-processing functionality, is introduced to support this goal. Such modules are not allowed to communicate directly with each other, which allows them to be independently tested and reused. Instead, modules communicate only through the Event. Figure 2.1 shows the collaborating components.

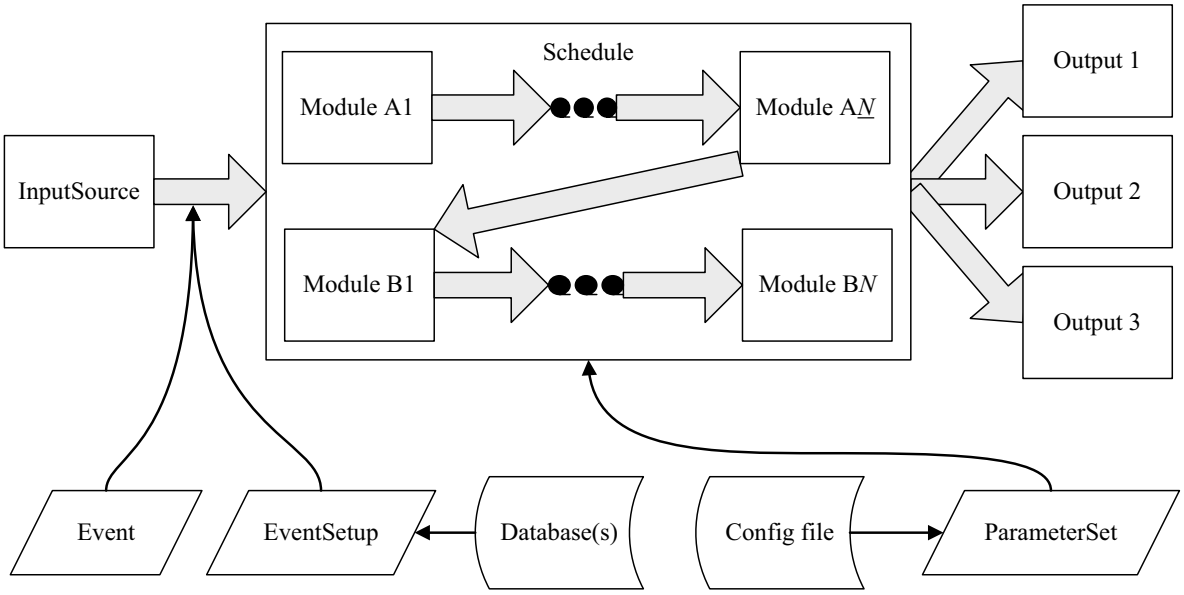


Figure 2.1: Components of the CMS Framework and EDM.

2.2.2.1 Framework configuration

The CMS framework executable (**cmsRun**) is configured, by "ParameterSets," which are collections of named parameter/value pairs. These ParameterSets are created from a user-written configuration file. The ParameterSets used to configure a job will be stored in the same file as the event data written by that job.

[1] In C++, a *class* bundles together some amount of data with the set of functions relevant for manipulating those data.

2.2.2.2 General characteristics of modules

Here are some characteristics of modules: a "module" is the generic term for a "worker" in the Framework. Not all modules have the same interface. Modules are scheduled by the "ScheduleBuilder," and invoked by the "ScheduleExecutor." Each module instance is configured with a ParameterSet. Modules must not interact directly with (i.e. call) other modules.

Only modules are *configurable*. An internal algorithm is configured by percolating ParameterSets to the algorithm, through the module that contains the algorithm. In order to provide for modular testing, which is important for quality assurance of the physics results, we require that modules communicate *only* through the Event, by putting "EDProducts" into the Event. Furthermore, we require that one may "cut" the event-processing chain between any two modules, and save the state of the event at that instant. This requires that all EDProducts have the ability to be made persistent. While each EDProduct can be made persistent, this does not imply each one *must be* for every event. The event output mechanism must be capable of selective writing of EDProduct instances to several output streams.

Here is a (possibly non-exhaustive) list of framework module types:

- event data producers ("EDProducers")—used in triggering, reconstruction, and simulation, these are the modules that put data products into the Event.
- output—each of these modules write the event data to one of several persistent forms.
- filter ("EDFilters")—used in triggering, these modules control the flow of processing for the trigger lists.
- analyzers ("EDAnalyzers")—these modules do not modify the event data, but can use it to create histograms or other event summaries.
- input—while not in the Framework's technical sense a module, input modules are used to read event data (e.g., from persistent storage, or from the DAQ system) and to deliver the data to the Framework.

2.2.2.3 Scheduler and paths

We will support two different "styles" of event-processing application in the same software framework. One style of application supports "reconstruction on demand," in the style of the previous COBRA framework [26]. The other style is more similar to the style of the CDF [35] and DØ [36] trigger and reconstruction frameworks. We call these two styles *unscheduled* and *scheduled*, respectively.

For both the unscheduled and the scheduled applications, EDProducer instances are the objects that actually perform the task of reconstruction. An author of an EDProducer does not need to choose to support one or the other style of use; any EDProducer is able to be used in either mode. For both styles of application, the same EDProduct classes are used, and the same EDProduct instances will be produced from identically-configured EDProducers. For both styles of application, the same parameter set system is used to configure the EDProducers. For both styles of application, the same input and output formats are supported.

2.2.2.3.1 Scheduled application A scheduled application is configured by specifying a module instance path through which the Event will flow. More derived or calculated products will be added to the Event as it moves through the path. This is our baseline mode.

The responsibility of getting the proper dependency ordering within an explicitly specified path lies with the person configuring the job. However each path can be considered independent of the others. It is the job of the framework to optimize the schedule given a set of fully self-contained reconstruction paths.

2.2.2.3.2 Unscheduled application An unscheduled application is configured by specifying: a selection of independent top-level EDProducts to be written out, or a selection of independent high-level triggers to be run, or an analysis module to be run, or some combination of the above. Calculated products will be added to the event as a result of requesting them from the Event.

The responsibility of getting the proper dependency ordering is handled automatically by the code. Each top-level EDProduct request can be considered independent of the others.

2.2.2.4 Provenance

It is critical for users to be able to unambiguously identify how each reconstruction result was produced. There are several varieties of information that constitute this identification.

Collectively, we refer to all this information as the *provenance* of the EDProduct. Each EDProduct is associated with a "Provenance" object that records this information. Where appropriate, Provenance objects are shared between EDProduct instances.

1. Module configuration

   (a) The unique identifier representing all (the names and values) of the run-time configuration parameters given to the module.

   (b) A string giving the fully-qualified class name of the module.

2. Parentage

   A vector of the unique identifiers of the EDProducts used as inputs for this bit of reconstruction. The identifiers are unique to the Event.

   Although a module can make use of more than 1 input to create its output, we make no attempt to specify the *type* of the EDProduct to which each of the entries in this vector refer. If such identification is needed in a particular EDProduct, that product can store the information in its own member data. However, a *mapping* of class name to "EDP_id" is deemed unnecessarily complex for the simple use to which the *parentage* information in this general form is put.

3. Executable configuration

   (a) A human-friendly string called a "module label," which is a unique identifier (within a job) used for EDProducts created by the module configured by this label. This label comes from a module configuration parameter with a fixed name. Each module has exactly 1 of these.

The label configuration parameter is special. Changing the label in the configuration will cause a new module to come into existence because a unique ParameterSet determines module instances. However, the label is not part of the permanently generated ID.

(b) A single version number that defines the code for the entire executable. The user can obtain specific library version numbers by querying a central database, using this version number.

The value is only meaningful for tagged releases. This number specifies which libraries were *available* when building the application; it does not indicate that *all* such libraries were used.

4. Conditions Data

   An identifier representing the calibration and alignment set that was used in the construction of this EDProduct.

   We assume here that calibration and alignment are handled in the same way and that this single, high-level identifier refers to all the calibration information used for this event. It is possible that individual calibrations (e.g., inner tracker, calorimeter, muon detector) will also have IDs associated with them and that each of these will need to be recorded instead of the "set" ID.

   Other conditions data IDs may also be needed here, such as geometry version or hardware configuration.

5. Job configuration

   A physical process name. A job starts up in a particular context such as HLT or Reconstruction. This name identifies the process under which the job was started and is likely to be a run-time property.

All of this *provenance data* is distinguished from the *event data* because their home is in an ancillary database, although a copy may be readily accessible from the event data (e.g., within the POOL file [37] that contains events).

A "Provenance" serves to collect the relevant information describing *how* a given EDProduct was created. Each EDProduct is associated (in an Event) with *one* Provenance.

2.2.3 Framework services

The Framework identifies 2 categories of services based on whether or not the service can affect physics results. The 2 categories are handled by different system: the "ServiceRegistry" system and the "EventSetup" system.

2.2.3.1 ServiceRegistry System

The ServiceRegistry system provides services which are application extensions and, as such, their use should have no effect on physics results. Therefore, the ServiceRegistry is used to deliver services such as the error logger or a debugging service which provides feedback about the state of the Framework (e.g., what module is presently running). What services

are to be used in a job and the exact configuration of those services is handled through the standard configuration mechanism, i.e., via a ParameterSet.

Because ServiceRegistry services are meant to be application extensions that have no effect on physics results, a service can be accessed globally (e.g., the error logging service can be called from any piece of code). All services are accessed using a consistent interface to make the system easier to learn and maintain. If an algorithm asks for a service which is not available, the error is handled using the same method used when data is not available in the Event.

Services are informed about the present state of the Framework, e.g., the start of a new Event or the completion of a certain module. Such information is useful for producing meaningful error messages from the error logger or for debugging.

2.2.3.2 EventSetup system

To be able to fully process an Event, for example in an analysis, requires additional information outside of the Event itself (e.g., magnetic field measurements). These non-Event data are data whose "interval of validity" (IOV) is longer than 1 Event. We have 2 types of IOVs which are distinguished by whether or not the DAQ system initiated the interval of validity transition. IOVs initiated by DAQ (such as the Event or a Run transition) are to be handled by the Event system. All other IOVs are handled by the EventSetup system. The EventSetup system provides a unified access model for all services that deliver non-Event data.

The EventSetup provides a uniform access mechanism to all data/services constrained by an IOV. This will include all calibrations, alignments, geometry descriptions, magnetic field and run conditions recorded during data acquisition. Figure 2.2 shows the main concepts for the EventSetup, which are:

1. Record: holds data and services which have identical IOVs.
2. EventSetup: holds all Records that have an IOV which overlap with the "time" of the Event being studied.

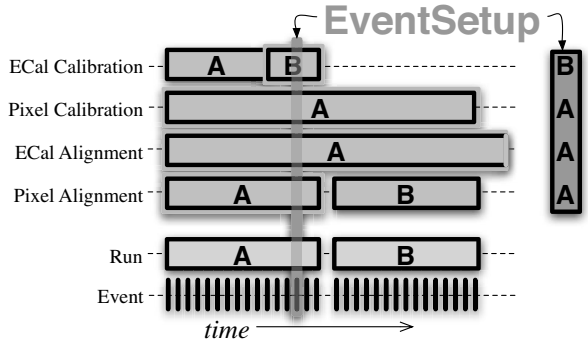


Figure 2.2: The EventSetup is formed from the Records that have an IOV that overlaps with the moment in time that is being studied.

2.2.3.2.1 EventSetup As shown in Figure 2.2, the EventSetup provides access to the various Records it contains. If the requested Record is not available, the Framework will take appropriate action (e.g., stop the job or skip that Event). In this way physicists can safely assume that the Record they want exists. In addition to access to Records, the EventSetup provides access to information about the "instance in time" for which the EventSetup is describing (e.g., the run and event number).

2.2.3.2.2 Records Records provide safe, read-only access to the objects they contain. This is done in a method analogous to data access from the Event. A Record also provides access to its interval of validity (IOV).

The EventSetup system sets no requirements on the data/services which may be placed in a Record. The only restriction is the lifetime within the program of the data/services within a Record is only guaranteed to be as long as the IOV for which the Record is appropriate (e.g., if the IOV of the data is only good for 1 run, then the object holding the data is only guaranteed to be available for that 1 run). This does not mean that an object that holds a datum/service within a Record cannot be reused across an IOV transition (e.g., one run to the next). It only means code that reads the object from a Record should not assume that it will be reused and therefore should be re-read after every IOV transition. The ability to reuse an object is meant to allow for speed optimizations within the program.

In the case where the data or service is only meant to come from one Record (e.g., the ECAL pedestal data only come from the ECAL pedestal Record), then the "default" Record can be declared. If a default Record has been declared for some data, then users can access that data directly from the EventSetup without having to specify the Record. For example, a physicist can ask for the ECAL pedestal data directly from the EventSetup rather than first asking the EventSetup for the ECAL pedestal Record and then asking that Record for the data. This is intended to decrease the workload on experimenters since they have to know less information (e.g., what data is in what Records) and can write fewer lines of code.

Sometimes an algorithm in the EventSetup is dependent on data coming from more than 1 Record. For example, the aligned tracking geometry is dependent on the "ideal geometry" and on the tracking alignment values. In such a case, the Record used by that algorithm needs to be declared "dependent" on the other Records. The IOV of a dependent Record is the intersection of the IOV of all the Records to which it depends. For example, the IOV of the aligned tracking geometry must change when either the IOV of the ideal geometry or of the alignment values changes. The EventSetup system guarantees that the proper relationships between the IOVs is preserved. Dependent Records allow access to only those Records to which they are dependent. In this way, the IOV dependencies between Records can be enforced (if you can not read data from another Record then you can not be dependent on that Record).

2.2.3.2.3 EventSetup system components The EventSetup system design uses 2 categories of components to do its work: "ESSource" and "ESProducer." These components are configured using the same configuration mechanism as their Event counterparts, i.e., via a ParameterSet.

ESSource: an ESSource is responsible for determining the IOV of a Record (or a set of Records). The ESSource may also deliver data/services. An ESSource normally reads its information

from a "persistent store" (e.g., a database) although it is not required to do so. For example, the ECAL pedestals will be delivered via an ESSource that reads the appropriate values from a database.

ESProducer : an ESProducer is, conceptually, an algorithm whose inputs are dependent on data with IOVs. The ESProducer's algorithm is run whenever there is an IOV change for the Record to which the ESProducer is bound. For example, an ESProducer is used to read the ideal geometry of the tracker as well as the alignment corrections and then create the aligned tracker geometry from those 2 pieces of information. This ESProducer is told by the EventSetup system to create a new aligned tracker geometry whenever the alignment changes.

2.3 Event Filter

The CMS Trigger and Data Acquisition System (TriDAS) is designed to inspect the detector information at the full crossing frequency and to select events at a maximum rate of $\mathcal{O}(10^2)$ Hz for archiving and later offline analysis. The required rejection power of $\mathcal{O}(10^5)$ is too large to be achieved in a single processing step, if a high efficiency is to be maintained for the physics phenomena CMS plans to study. For this reason, the full selection task is split into 2 steps. The first step (Level-1 Trigger) is designed to reduce the rate of events accepted for further processing to less than 100 kHz. The second step (High-Level Trigger, or "HLT") is designed to reduce this maximum Level-1 Accept (L1A) rate of 100 kHz to a final output rate of approximately 100 Hz. The design of the Level-1 Trigger has already been extensively documented in Volume 1 of the TriDAS Technical Design Report [7]. The design of the Data Acquisition System is documented in Volume 2 [8].

The functionality of the CMS DAQ/HLT system can be summarized in 4 points:

- perform the readout of the front-end electronics after a Level-1 Trigger accept and assemble data from a given bunch-crossing in a single location (the memory of a computer);
- execute physics selection algorithms on the events read out, in order to accept the ones with the most interesting physics content;
- forward accepted events, as well as a small sample of the rejected ones, to online services monitoring the performance of the CMS detector
- provide the means of archiving accepted events in mass storage

2.3.0.3 DAQ architecture overview

The following is a schematic summary of the main functional elements as depicted in Figure 2.3:

- Detector Front-ends: the modules that store the data from the detector front-end electronics upon the reception of a Level-1 Trigger accept signal.
- Readout Systems: the modules that read the data from the detector Front-End and store them, until they are sent to the processor which will analyze the event.
- Builder Network: the collection of networks that provide the interconnections between the Readout and the Filter Systems. It is a large switching fabric, capable

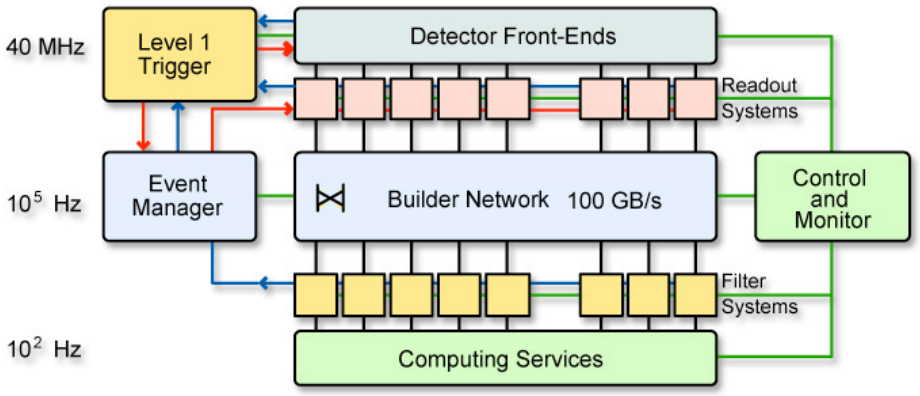


Figure 2.3: General architecture of the CMS DAQ System.

of supplying 800 Gb/s sustained throughput to the Filter Systems.

- Event Manager: the entity responsible for controlling the Builder Network data flow.
- Controls: all the entities responsible for the user interface and the configuration and system monitoring of the DAQ.
- HLT Systems: the ensemble of the components providing control, input data, monitoring, error detection services, and the processors executing High-Level Trigger algorithms. Approximately 500 Builder Units receive the incoming data fragments corresponding to an individual event and build them into full event buffers. An appropriate number of Filter Units are connected to each Builder Unit, to provide the necessary processing power to carry out the High-Level Trigger selection.
- Services: all the processors and networks which receive or route complete or partial events, or online monitoring information, from the Filter Farm.

2.3.0.4 HLT system requirements

The High-Level Trigger must reduce the event rate output by the Level-1 Trigger by a factor ≈ 1000 for a total output to storage of ≈ 100 Hz. At the design luminosity of the LHC, this total expected to be output to mass storage, corresponds to a cross section of 10 nb. Given that the $W \rightarrow e\nu_e$ production cross section alone is of this order, a significant degree of physics selection has to take place online. It is this aspect of the HLT system that places the most stringent requirements on the system. The main requirements on the system are thus:

- The system has to provide enough bandwidth and computing resources to minimize the dead-time at any luminosity, while maintaining the maximum possible efficiency for the discovery signals. The current goal is to have a total dead-time of less than 2%. Half of the dead-time is currently allocated to the Level-1 Trigger system.
- The HLT system should tag selected events with the specific trigger selection paths that were satisfied. This information can then be used by the offline system for a quick sorting of the events into primary datasets.

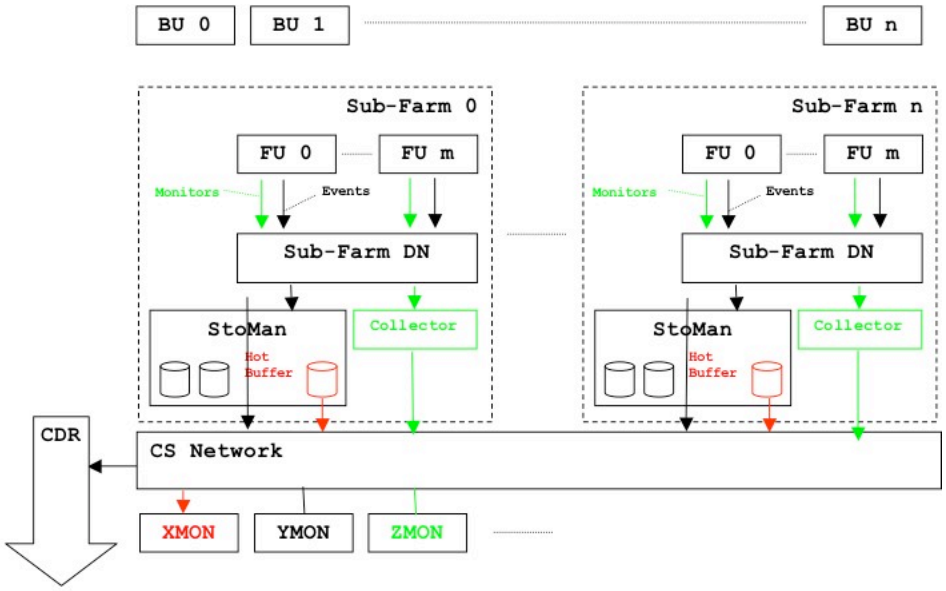


Figure 2.4: Architecture and Data Flow of the Filter Farm.

- The system should allow for the readout, processing and storage of events that will be needed for calibration purposes.
- The HLT efficiency should not depend strongly on changes of the calibration and detector alignment constants. It must in any case be possible to validate trigger selections and compute their overall efficiency using only the data itself, with as little reference to simulation as possible.
- The system has to be flexible enough to adapt to changing run and/or store conditions. As an example, the instantaneous luminosity is expected to drop in the course of a store, and therefore an optimal allocation of resources might require changing trigger conditions, for instance by lowering trigger thresholds or decreasing prescale factors for select channels. All such changes, along with any other changes in the running conditions, must be logged.
- The system should provide enough resources to monitor the status of the CMS detector, and to provide enough information to the experimenters in case of problems.
- To maximize the efficiency of the filtering process, an event should be rejected as soon as possible. Furthermore, the system should not rely on the presence of all the information from the CMS detectors.

Given the unprecedented rate of online rejection, a most important task of the HLT is, after achieving this rejection, to provide enough information on what is rejected. It is for this reason that a major aspect of the HLT System concerns the control and monitoring of the algorithm running in the Filter Farm.

2.3.1 Filter Farm architecture

The architecture of the Filter Farm is described in detail in Volume 2 of the TriDAS Technical Design Report [8] and is schematically illustrated in Figure 2.4. Based on the analysis of the

CPU needs of existing HLT algorithms, it is estimated that the HLT will run on a farm of PCs consisting of ≈ 1000 dual-CPU nodes (2007 units). To match the organization of the DAQ system, improve the scaling behavior of the control system, and allow a staged deployment, the farm is organized in a number of smaller subfarms, each with ≈ 100 processors (Filter Units, FU).

The Filter Units connect to a subset of Builder Units (BUs) via a switched network (Filter Data Network) capable of sustaining the incoming data traffic from the Event Builder running at the nominal L1A rate of 100 kHz. The BU-FU communication is realized with a simple request-response protocol: a FU can ask for 1 or more event(s) which are delivered asynchronously by the BU. When a decision has been reached for a given event in the FU, the event is discarded by issuing the appropriate message to the BU. Each FU in a subfarm is assigned at configuration time to a given BU, to which it places requests.

All FUs in a subfarm forward accepted events to a Storage Manager (SM). The main task of the SM is to write accepted events in the local disk buffer. Each SM has enough local disk buffer to store event data for at least 24 hours of steady data taking.

Accepted events contain the original raw-data plus any additional by-product of the reconstruction carried out to reach the HLT decision. In any case, they contain enough information to enable checking of the HLT decision offline. The FU-SM communication and the architecture of the Storage Manager are discussed in more detail in a following section.

Each subfarm is controlled as a whole by a Subfarm Manager (SFM), relaying commands from Run Control and keeping track of the states of individual nodes.

All software applications running in the Filter Farm are based on the CMS Data Acquisition framework, XDAQ [38, 39, 40]. XDAQ components are used to configure, control, and monitor the applications, as well as to support intercommunication amongst them.

2.3.2 Data flow

For each L1A, a total of ≈ 75 million electronics channels are read out for the various CMS subdetectors. Digitized data are collected by ≈ 600 Front End Driver boards (FEDs) operating synchronously with the LHC clock. Event fragments collected by individual FEDs are encapsulated in a common data structure by adding a header and trailer that mark the beginning and the end of the event fragment. Information in the header and trailer includes bunch-crossing and Level-1 Accept identifiers, as well as fragment size and Cyclic Redundancy Check (CRC) information used by the DAQ to check for data-transfer errors. Each FED is identified by a unique source-id which enables the Filter Unit to unambiguously associate a FED fragment to a certain group of subdetector channels. Event fragments are asynchronously transferred to the DAQ via a 64-bit serial link (S-LINK64 [41, 42]).

The DAQ Event Builder uses a complex switch fabric to assemble all fragments from a given event in the memory of 1 Builder Unit. This operation is carried out at pace with the L1A rate, up to a nominal maximum average accept rate of 100 kHz. For efficiency reasons, the event fragments are organized in the BU memory using blocks of fixed size.

The Filter Farm data flow is schematically illustrated in Figure 2.4. Built events are served to Filter Units connected to a given BU upon request. Events assigned to a Filter Unit remain in the memory of the BU until a discard message is received from the FU.

The Filter Unit operation, from the DAQ standpoint, consists of collecting data messages from the BU to form a complete event, and rearranging the event data in memory as individual FED fragments, using the information in the common header and trailer. Data are checked in the process to identify synchronization problems (fragments from different L1As) or other inconsistencies which might have been introduced in the event building. Subsenquently, event ownership is transferred to the HLT Event Processor that carries out reconstruction and selection. When the HLT processing is complete, the event ownership is returned to the DAQ segment that frees memory and issues a discard message to the BU.

The Filter Unit Event Processor uses software modules based on the CMS Framework(2.2) to run reconstruction and selection algorithms necessary to reach the final HLT decision. The sequence of operations that lead to acceptance or rejection of one event is guaranteed to be unique and reproducible. Each reconstruction or selection algorithm runs in a predefined sequence, starting from raw data unpacking down to sophisticated topological cuts on high level physics objects.

Raw data unpacking modules deal with subdetector specific raw-data formats. Unpacked data objects ("digis") are associated with subdetector channels through a channel map using the FED block identifiers. Consistency checks of individual subdetector raw information are carried out at this stage. All subsequent reconstruction and selection modules only deal with digi objects, so that no other part of the code is exposed to the original FED-formatted data. The latter can be optionally stored with the event for debugging or diagnostic purposes. To address potential performance issues related to the CPU requirements of the unpacking phase, the framework will support selective unpacking of fractions of the subdetector data corresponding to regions of interest identified by means of some fast criterion (e.g. Level-1 candidate location). All intermediate results of any reconstruction module, and information that enables tracing of their provenance, can be associated with the event. The operation of the HLT algorithms can be monitored making use of the Physics and Data Quality Monitoring infrastructure (DQM). Information collected by the DQM on the FU is periodically forwarded to a collector process, which is not part of the subfarm. The FU is able to continue its operation undisturbed even if the DQM is malfunctioning. For more details on the DQM infrastructure, see Section 2.9.

Events accepted by the HLT are serialized and wrapped with preamble information, before being forwarded to the Storage Manager (SM) process Figure 2.5. The preamble contains enough information to enable the Storage Manager to route the accepted event to the appropriate output stream without looking into its content. It also contains information that can be cross-checked with the local configuration information available to the SM (e.g. version number, run number etc.). In the serialization, certain portions of the event data can be dropped according to a general policy or to a specific policy for certain types of events.

The Storage Manager is responsible of routing accepted events to various predefined online streams. These can be disk streams for physics data, or network streams for the usage of consumer processes carrying out calibration or monitoring tasks. The machine hosting the SM is also directly or indirectly managing the subfarm disk pool. The SM is capable of routing multiple copies of 1 event to different streams. For disk streams, the SM will use file formats supported by POOL [37]. It can optionally directly store serialized events in binary format. This second option requires later post-processing to recreate the POOL files to be transferred to the tier-0 for permanent storage.

Since the Filter Units process events in parallel, there is no guarantee that accepted events

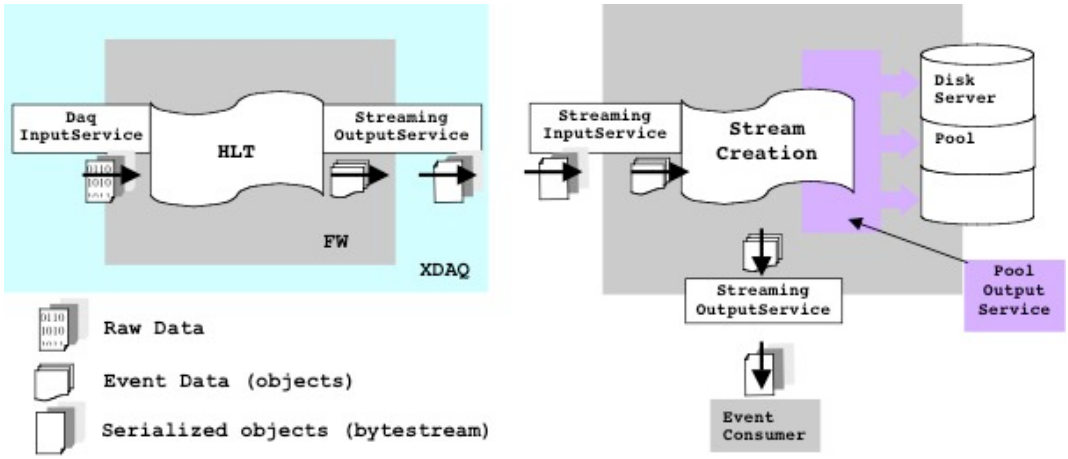


Figure 2.5: Storage Manager Data Flow and Components.

forwarded to the SM for storage will arrive in the correct time order. However, events contain sufficient information to enable reordering according to their L1A time in the subsequent post-processing stage. The post-processing phase has the following responsibilities:

- Collect datafiles for a given online stream from all subfarms
- If events are stored in serialized binary format, deserialize them
- Reorder events from all datafiles according to their Level-1 accept time
- Split the reordered stream in sections of fixed duration (run sections)
- Create time-ordered primary datasets according to L1+HLT trigger path
- Store them in POOL

POOL files containing time-ordered sections of primary datasets are the input to the first pass production in the tier-0.

2.3.2.1 Raw data types

As mentioned above, FED blocks of digitized data received by the FU contain information from the various subdetectors in a format designed to minimize data traffic in the Event Builder. FED blocks are wrapped with a standard header and trailer containing, among other things, a unique source identifier that enables routing of the block to the appropriate unpacking module. Ranges of source identifiers are pre-assigned to the various subdetectors by the DAQ group. Upon reception in the FU, FED blocks are arranged in a data structure indexed by the source identifier, preserving their internal format, including the standard header and trailer. This data structure (FEDRawData) can be stored along with the rest of the event, enabling later checking of the unpacking stage. Each subdetector specific unpacking module extracts from the event the group of FEDRawData corresponding to its input. It then proceeds to extract the relevant information for each detector channel and creating the corresponding digi object. The correspondance between a certain detector element and its electronics channel is obtained from a Channel Map in the conditions database. The digi objects contain enough information to uniquely identify the corresponding channel in the context of the reconstruction framework (DetId), and the low-level, uncalibrated content of that channel (e.g. ADC counts), and are arranged in collections enabling fast retrieval by a

reconstruction module. Collections of digi objects (one for each individual subdetector, or logical component of a subdetector) are stored along with the event. They are the standard input for subsequent reconstruction, both in the HLT and in the following stages.

2.3.2.2 Output service and Storage Manager

An HLT process running in the Filter Unit uses a special Output Module to serialize accepted events to be sent to the Storage Manager. Serialized events are handed over to a Sender thread running in the Filter Unit, which is responsible of maintaing the necessary handshake between the FU and the Collector running in the SM. The protocol used for communication consists of control and data messages. Control messages are used for

- Establishing the connection
- Checking the consistency of initialization information between the 2 applications (code version, trigger table, etc.)
- Flow control

Data messages are used to send event data from 1 FU to the SM. The SM Collector transfers collected events to 1 or more of the storage queues according to their trigger configuration, as it appears in the event preamble. The SM must be capable of handling a number of connection in excess of the maximum total number of FUs in a subfarm. An SM Event Processor is responsible of creating a physical disk or network stream from the events in 1 queue. Multiple SM Event Processors may run in parallel on multiple threads or processes for performance reasons.

2.3.3 Control flow

The Filter Farm is operated as a DAQ component, under the supervision of the global DAQ Run Control system. Run Control commands are relayed to individual subfarm PCs by a Subfarm Manager (SFM). The behavior of the Filter Units, and the Storage Manager, is controlled by standardized state machines. Each state transition is monitored by the SFM and reported to run control upon completion. At the request of the SFM, a daemon process on each PC of the subfarm spawns the necessary XDAQ processes that will host individual components. Initialized processes of all the components are subsequently configured. When all processes are ready, all DAQ components are enabled. Run Control can at this point start a run. The run start command is accompanied by a run identifier (Run Number) which can be subsequently used to uniquely identify the corresponding event data. Only at the start of a run the HLT is configured. The HLT configuration consists of sequences of reconstruction and trigger modules, and their parameters, and of all the calibration constants and other condition data necessary to carry out the selection, including the Level-1 trigger configuration, forming the HLT trigger table. The HLT configuration is also made available to the SM at the run start in order to allow proper streaming of accepted events according to the trigger accept path. A run can only be ended as a consequence of a run end command issued by the global Run Control. When a run is ended, the Level-1 trigger is stopped, and the Event Builder is flushed. When the flush is complete, each individual subfarm in the Filter Farm is instructed to close the run. All Filter Units must report having successfully closed the run before the SM can finally close all output streams. This sequence is carried out under control of the Subfarm Manager.

2.3.3.1 Fault tolerance, exception, logging, and error handling

The complex of the Event Builder and the Filter Farm system are designed for fault tolerant behavior. In particular, failure of one or more Filter Units in a subfarm only leads to "graceful" performance degradation of the Event Builder. Since events are not removed from the BU until a discard message, indicating that the event has been completely processed, is received, no data can be lost due to failure of a Filter Unit. In the Filter Unit, all exceptional conditions generated by the execution of reconstruction or selection algorithms, are either recovered locally or caught by the top level HLT Event Processor. Recovered exceptions generate a log message which is sent to the central DAQ logging system. Unrecoverable exceptions caught by the Event Processor also generate an error message, and result in a state transition of the HLT Event Processor to the error state. Failure of the HLT Event Processor does not result in a data loss, unless if it causes corruption of the DAQ segment memory. The Subfarm Manager can reset and restart a failed HLT or DAQ process in the Filter Unit.

In case of failure of the Storage Manager, the FUs can switch to local storage mode, whereby accepted events are stored on each FU's local disk, so that data taking can continue. This can also happen in case the SM local disk buffer becomes full due to malfunctions in the Data Management system. Failure of the SM can result in data losses if one or more of the output files become corrupted, or in case of failure of the file systems hosting the local buffer. Various redundant storage systems are being considered to reduce the risk of losses due to disk failures.

Log messages generated by Filter Farm components are collected by the central DAQ logging system. They become part of the run conditions database and can be analyzed offline to detect and understand problems in the data streams. Error messages generated by the Filter Unit and the Storage Manager are forwarded to the Subfarm Manager. The Sub-farm Manager is capable of analyzing and filtering error messages to avoid network storms. Errors which can not be handled locally by the Sub-farm Manager are forwarded to central DAQ Error Handler. Unlike the logging system, the error handling system guarantees delivery of all error messages.

2.4 Detector description

The CMS detector description includes geometrical shapes and dimensions, material information, and relative placements of each part of the detector in an ideal world. The software service providing this information uses a master source which can have derivatives. Sources derived from the master source can be databases, flat-file caches or other more optimized storage.

The Detector Description Database (DDD) [43] provides access to the ideal geometry (geometry defined for the nominal placement of all geometrical objects) for CMS. Historically the DDD has referred to both a C++ interface and ideal geometry data sources. The master data source is a collection of Extensible Markup Language (XML) files which can be validated using the XML-Schema Detector Description Language.

The online relational database model and the offline objects made persistent using the POOL [37] Object Relational Access mechanism [44] are both derived from the master source files. The offline tables can be in either SQLite, MySQL or Oracle relational databases.

The DDD internal model is a directed acyclic graph which optimizes the definitions of materials, solids, rotations and parts so that the in-memory model is small. This means that each part as it is in the real detector does not exist in memory unless one iterates over the graph through all nodes and edges this producing a tree of nodes. This tree represents the global transformation (translation and rotation matrix) in the CMS global reference frame. Filtering is provided on the graph so that the tree view can be constrained to show only certain detector parts.

Developers of basic software components for simulation, reconstruction and visualization use the DDD C++ interface to retrieve the ideal detector description. The graph is provided to the client software (simulation, reconstruction and visualization) via the "DDCompactView". One can apply filters on this view using the "DDFilteredView" or access the collections of parts and positioning instructions.

2.4.1 Geometry Service

The Geometry Service provides the necessary geometrical information for an event to be processed using the real geometry of detector elements as actually measured at the time of the event. Framework components are used to provide the ideal geometry from the Detector Description as well as Alignment corrections from the Conditions Service. Other framework components ensure the synchronization of the conditions to the event time. Subdetector specific software applies the alignments to the ideal geometry.

The CMS Geometry Service relies on the underlying software framework to provide the ideal geometry and alignment corrections necessary to deliver a user-level geometry for physics analysis. The EventSetup mechanisms, see Section 2.2.3.2, is used for this service. *ESSources* (EventSetup Sources) provide the DDCompactView from the XML source or a database. Another *ESSource* provides Alignment via the he Condition Service. Subdetector specific software modules provide the geometry to the user-level, combining the 2 sources transparently for the user.

2.4.2 Detector numbering scheme

The detector Numbering Scheme (NS) provides the identification of the in the reconstruction active CMS detector elements. A 32-bit integer is sufficient to define the NS. The general detector numbering scheme structure is specified in Table 2.1.

Table 2.1: General structure of CMS Detector Numbering Scheme.

content	DetectorID	SubDetectorID	specific SubDetector
size	4 bits	3 bits	25 bits

The NS is arranged in such a way that it uniquely reflects the hierarchical structure and positions of active elements inside each subdetector. As an example, the specific numbering for the Tracker Outer Barrel modules is constructed from the Layer number, the Rod number and the modules position inside the Rod.

The detector numbering scheme is used in the simulation and reconstruction software in order to navigate through the active elements (e.g. "DetUnits") and, therefore, represents a

very convenient way to access relevant reconstruction or simulation information via a simple 32-bit number.

Besides its use in simulation and offline reconstruction, the NS is also crucial for querying information from the Databases (DB). For that reason, the numbering scheme is written in both the offline and online DBs (see Section 2.8). In the offline DB the numbering scheme is associated with the volume name while in the online DB it is linked with the hardware identifier of the detector. This concept insures that all relevant information for given (active) detector element can be accessed via this unique identifier.

2.5 Simulation

The detailed CMS detector and physics simulation, a package referred to as OSCAR [31], is currently based on the GEANT4 [30] simulation toolkit and the CMS object-oriented framework and event model. The simulation is implemented for all CMS detectors in both the central region (Tracker, Calorimeters and Muon Systems) and in the forward regions (CASTOR calorimeter, TOTEM telescopes, Roman Pot detectors and the Zero Degree Calorimeter, ZDC), including the field map from the 4 T solenoid. In addition, several test-beam prototypes and layouts have been simulated. The full simulation program implements the sensitive detector behavior, track selection mechanisms, hit collection and digitization (i.e. detector response).

The simulation has been deployed since the 2004 CMS Data Challenge and has been extensively validated by detailed comparison with test-beam data as well as results from its predecessor, the GEANT3-based program CMSIM [29]. To date, a total of more than 100 million fully simulated events have been produced with the GEANT4-based simulation program; these include various LHC physics signals, backgrounds, and validation samples.

The detailed simulation workflow is as follows:

- A physics group configures an appropriate Monte Carlo event generator (several are used) to produce the data samples of interest;
- The production team/system runs the generator software to produce generator event data files in HepMC format, see Section 2.5.1;
- The physics group validates the generator data samples and selects a configuration for the GEANT4 simulation (detector configuration, physics cuts, etc.);
- The production team/system runs the GEANT4-based simulation of CMS, with generator events as input, to produce (using the standard CMS framework) persistent hits in the sensitive detectors;
- The physics group validates this hit data which are then used as input to the subsequent digitization step, allowing for pile-up. This step converts hits into digitizations (also known as "digis") which correspond to the output of the CMS electronics.

These generator and simulation steps above, culminating in digitizations, are described in the remainder of this section. The subsequent reconstruction step is described in Section 2.7.

2.5.1 Generation of physics events

Most of the event generators that provide the collision events as input for the detector simulation are still written in FORTRAN, notably PYTHIA [45] and HERWIG [46]. Several projects to write event generators in C++ are ongoing (PYTHIA8, HERWIG++, SHERPA [47]), and mature code is expected by 2007. The Generator Production Framework therefore has to be able to handle both FORTRAN and C++ generators.

The design of the Generator Interface in the CMS software logically comprises the following steps:

- An input file containing all parameters needed to steer the event generator is read by the production framework, which then passes the parameters to the generator itself. The framework parameter and provenance handling mechanisms are used.
- The generator (FORTRAN or C++, as appropriate) is initialized and events are produced. The generic format of Monte Carlo events throughout the CMS software is HepMC. GenEvents in HepMC are based on the natural structure of an event, i.e., vertices with ingoing and outgoing particles. If needed, the event output format of the generator (e.g. HEPEVT) is translated to HepMC objects using specific convertors.
- Optionally, events may be selected/rejected based on generator event information to enrich the physics content.
- In the final step, the events are made persistent by writing the HepMC objects as EDProducts to a tuple using POOL [37].

ASCII file readers for externally produced HepMC events are implemented, as well as particle guns.

The physics of the generators is covered in more detail in the references cited above and in Volume II of this TDR.

2.5.2 Detailed simulation framework

The full simulation relies on the GEANT4 toolkit [30]. GEANT4 provides a rich set of physics processes describing electromagnetic and hadronic interactions in detail. It also provides tools for modelling the full CMS detector geometry and the interfaces required for retrieving information from particle tracking through these detectors and the magnetic field. The physics and other aspects of GEANT4 itself are covered in detail elsewhere [30], while the validation of GEANT4 in the context of CMS is described in detail in Section 2.5.5.

The CMS GEANT4-based simulation program uses the standard CMS software framework and utilities, as are used by the reconstruction programs. The GEANT4 functionality is interfaced to the core CMS Framework (Section 2.2) to yield the following elements of the simulation software system:

- "EventSetup" to create the application context in terms of the detector and magnetic field description;
- "Input Services" to pass the generated events to the simulation application (using the IOMC Input Output MonteCarlo service);
- "Module factories" to select and load "physics lists" (models and parameters describing particle interactions in matter) and miscellaneous utilities;

- "ParameterSets" for run-time configuration;
- "Signaling mechanism" for user monitoring actions; and
- "Event Producer" for loading (saving) the hits (digitizations) from the sensitive detectors and the Monte Carlo simulation generated quantities such as particle tracks and vertices.

A fully interactive (and non-CMS specific) GEANT4 visualization system has been established to debug the GEANT4 simulation, in particular the geometry modelling and magnetic field description. It is based on the IGUANA toolkit. Figure 2.6 shows the IGUANA display of the full CMS detector as used for the GEANT4 simulation.

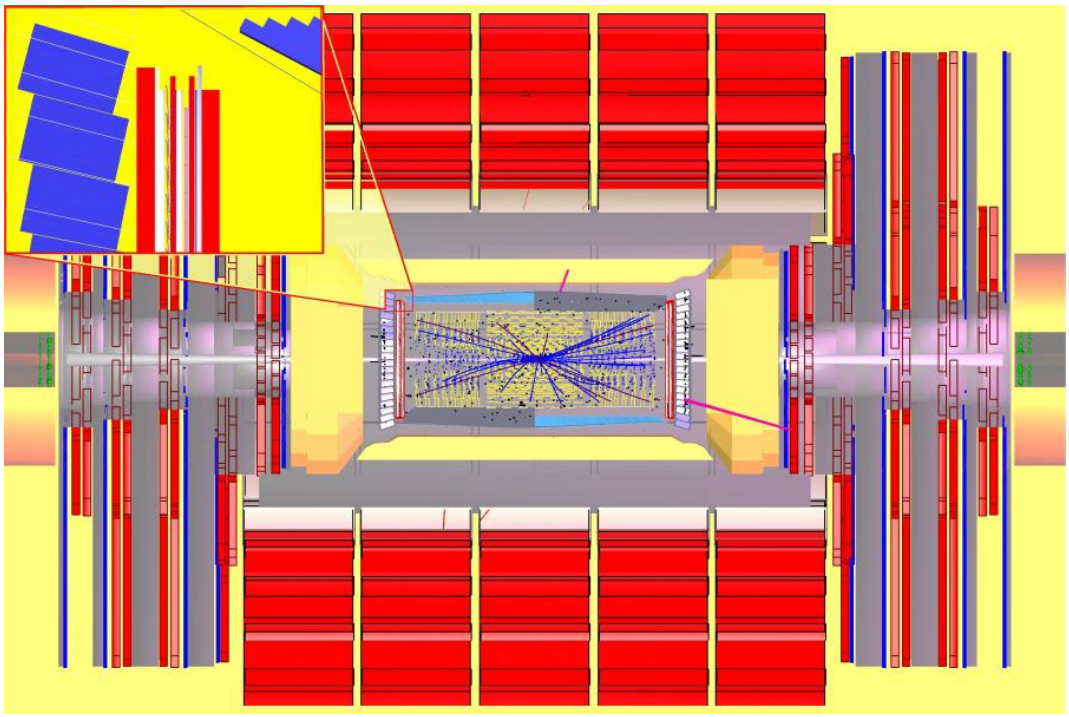


Figure 2.6: CMS GEANT4-based simulation of the $H \rightarrow \gamma\gamma$ process showing the full CMS detector with the IGUANA-based interactive display. An inset zoom of the detector detail is also shown.

2.5.3 Pile-up treatment

During the low luminosity ($\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) and high luminosity ($\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) phases of its operation, the LHC accelerator will produce, respectively, an average of about 3.5 and 17.5 inelastic (hard-core) pp collisions per bunch crossing that will "pile-up" on top of the signal collision firing the trigger. For the Monte Carlo simulations used in this report, diffractive collisions are also considered, which increases the pile-up total to 5.0 and 25 collisions, respectively, for low and high luminosity operation.

Moreover, in addition to the in-time pile-up, it is necessary to account in the simulation for out-of-time pile-up coming from bunch crossings before and after the triggered event. The number of crossings to consider before and after the nominal one depends on the front-end time response of the different subdetectors. Special cases such as bunch crossings with no

pile-up either before or after the nominal one are also considered by the framework.

Since the addition of pile-up collisions occurs much faster than the detector simulation, and since it depends on the LHC luminosity and run conditions, pile-up collisions are simulated separately from the signal collisions. Both outputs are merged in a second step, using a luminosity dependent pile-up contribution. Generated signal collisions are therefore reused for producing samples corresponding to different luminosities.

The pile-up simulation is performed using the same GEANT4 toolkit as for the signal simulation, producing simulated hits in exactly the same format. The pile-up collisions to be merged are randomly chosen from a pregenerated sample. In order to avoid correlations between event subsets overlapped to the same pile-up event sequence, a pile-up sample is never reused in the same order. Moreover, simulated pile-up collisions that pass trigger requirements are filtered out to avoid a low statistics bias in the many bunch crossings that would use such an event.

The "mixing module" is the software framework that merges events from a primary stream with a number of events from a secondary stream. The product of this module is a "CrossingFrame," which in turn is the input to the digitization module, the next step in the simulation chain. The CrossingFrame handles the sequence of bunch crossings and provides access to the information the digitizers need. It also merges the true information from the different pile-up and signal events that constitute the CrossingFrame, and provides access to such true information. To produce the CrossingFrame, the mixing module collaborates closely with a specialized secondary input service. The average number of events, the minimum and maximum number of bunch crossings after the nominal, as well as the bunch spacing are configurable parameters. The secondary input service is also configurable, to allow for different options to randomly select pile-up events from the pregenerated minimum bias sample.

2.5.4 Digitization of simulated detector hits

The digitization step, following the hit creation step, constitutes the simulation of the electronic readout used to acquire data by the detector and DAQ systems. It starts from the hit positions and simulated energy losses in the sensitive detectors, and produces an output that needs to be as close as possible to real data coming from CMS. Information from the generation stage (e.g. particle type and momentum) is preserved in the digitization step.

2.5.4.1 Inner tracking detector digitization

The energy loss of a charged particle crossing a layer of one of the inner tracking detectors (strips and pixels) is distributed along a path between the entry and exit points within the detector module. Landau fluctuations are taken into account. The simulated drift of the charges to the detector surface takes into account the Lorentz drift and diffusion in the perpendicular plane.

On the detector surface, charges corresponding to each pixel or strip are integrated, and Gaussian-distributed noise is added. Noise is also added to other channels if it exceeds a given threshold. Couplings between channels are taken into account, and the conversion to digital counts is applied using the gain of the detector and the time with respect to the

signal bunch crossing. If zero suppression is required, only channels with a signal exceeding a suppression threshold are saved. Further details can be found in Section 6.1.2.

2.5.4.2 Calorimeter digitization

In ECAL digitization, the sensitive volumes of the ECAL are the crystals and the silicon strips of the preshower detector. For these volumes, the deposited energy and arrival time of hits is recorded.

For the crystals the variation of the light collection efficiency along the length of the crystal is simulated by multiplying the energy, as it is deposited, by a nominal longitudinal light collection curve. This is a simple function of the distance from the crystal front-face, and is flat in the front half of the crystal, with a 5% linear rise over the rearmost 100 mm. Deviations from this nominal curve, as measured on production crystals, deteriorate the energy resolution and are accounted for by smearing the energy.

The digitization simulates the nominal signal pulse as a function of time for each hit, so that out of time pile-up results in realistic out of time signal pulses which may be superimposed upon a signal. The signal pulses are then sampled at 25 ns intervals to produce a time frame of 10 consecutive digitizations. Additional information can be found in Section 4.3.1.

In HCAL digitization, the simulation of the electronics for the HB, HE, and HO starts with the information provided from GEANT on the energies deposited in the scintillator and their corresponding timing. The energies are converted to numbers of photoelectrons, and fluctuations are applied assuming a Poisson distribution. Noise is then added, using a Gaussian distribution for the energy, equal to 1.5 photoelectrons per time sample per read out depth segment. This is uncorrelated between time buckets and corresponds to about 240 MeV after the corrections described below.

The pulses from the different energy depositions (both from the current crossing and from up to 5 previous and 3 subsequent crossings) are then added. The HF electronics uses a conventional photomultiplier tube instead of a hybrid photodiode, but otherwise the same electronics as for the HB, HE, and HO. The HF pulse shape is short enough to be entirely contained in one bunch crossing and is thus not affected by pile-up from previous or later bunch crossings. The magnitude of the noise used in the HF simulation is 0.125 photoelectrons and the ADC count size is set at 0.43 photoelectrons.

2.5.4.3 Muon detector digitization

The digitization step of the Muon Drift Tube system involves simulating the TDC responses. Particular care is taken in simulating the behaviour of the drift cells as a function of the muon direction and impact position with respect to the sense wire, and of the residual magnetic field in the air gaps where the detectors are placed within the iron yokes. The resulting drift time is smeared so as to obtain a 4 ns resolution, corresponding to an intrinsic cell resolution of about $220\ \mu\mathrm{m}$, as measured in test-beam data. The TDC output signal for the hit reconstruction is obtained from this drift time by adding the muon time-of-flight from the collision vertex and the propagation time of the signal along the cell wire. The average muon time-of-flight from the collision vertex to a given chamber is assumed to have been subtracted at TDC level.

The digitization step of the Cathode Strip Chamber system involves simulating the responses of the ADCs and discriminators connected to the strips and wires. To create the analogue signals seen by the CSC wire and strip electronics, parameterizations of the amplifier and shaper response are convoluted with the ion drift collection time. The signal may contain contributions from drifting electrons due to background hits from other beam crossings. Cross-talk, both capacitive and resistive, is included in the strip signal. Finally, the storage of the strip signals in Switched Capacitor Arrays (SCA) is simulated. The signal shape is sampled and stored at 8 times, each 50 ns apart.

The response of a resistive plate chamber is assumed to take place within 20 ns of the passage of a charged particle through the detector with a 3 ns Gaussian-distributed time jitter, which also accounts for the contribution from the front-end electronics and the cables to the link board. The 20 ns wide time gates were adjusted in order to accommodate triggering signals. The RPC cluster size is set to 1.5 strips.

Additional and more specific information about the digitization of muon detector signals can be found in Section 3.3.1.

2.5.5 Simulation physics validation and performance

High energy physics experiments depend critically on the accuracy of physics generators and detector simulations. Simulated events are used for detector design optimization, calibration, object identification, and physics analysis. The size of systematic uncertainties associated with particle discoveries, mass, or cross section measurements is tightly associated with how accurately the simulations describe the actual performance of the detector in measuring electrons, photons, and hadrons. It is essential for the success of a HEP experiment to understand and tune the physics of the simulation tool to agree with measurements. In this section we mention some of the comparisons made between data and GEANT4 (and in some cases GEANT3), and the discrimination between some of the models (physics lists).

2.5.5.1 Tracker validation

The inner tracker simulation has played a key role in the development and optimization of the simulation infrastructure and the validation process. The tracker material budget, which can be correctly estimated only with a very detailed description of all active and passive detector components, directly affects the electromagnetic calorimeter physics performance and places stringent requirements on the accuracy of the detector description and geometry construction. A navigable listing of the particle information from the Monte Carlo simulation is necessary for the correct reconstruction of decay chains as well as for the proper treatment of hard electron bremsstrahlung, both of vital importance in b-τ studies.

With the above requirements satisfied, the tracker performance has been extensively validated in terms of the tracking and hit distributions for single particles, minimum bias interactions, and other physics events. For details see Section 6.1.2.

2.5.5.2 Electromagnetic calorimeter validation

Initial studies based on a comparison between a GEANT4-based simulation and test-beam data provide evidence that GEANT4 gives an excellent representation of electromagnetic showers.

The overall ECAL performance, in terms of energy and position resolution, is dominated by effects that are not part of the shower simulation, such as electronics noise, photostatistics, longitudinal uniformity of light collection, and crystal intercalibration. These effects are taken into account in the digitization step of the simulation process. For this reason, only gross errors are identified by a comparison of energy and position resolution taken from GEANT shower information.

On the other hand, the largest sensitivity is to changes or errors in the radiation and showering in the tracker material. Unfortunately the accurate simulation of this effect cannot be validated in the test beam. The shower lateral distribution, and its fluctuations from shower-to-shower is an important quantity which can be validated comparing the Monte Carlo simulations with test beam measurements. In particular, parameters sensitive to the lateral shower shape, which effects the fraction of incident energy contained in ECAL clusters, are measured in the test beam. For details see Section 4.3.

2.5.5.3 Hadron calorimeter validation

Studies on the HCAL energy resolution, linearity, e/π ratio, and shower profile are instrumental in validating the GEANT4 hadronic physics simulation in the context of the LCG physics simulation validation project. They are based on comparisons between single particle measurements in test beam experiments and GEANT4 based simulations of the associated detector setup.

In 2002–2004, several HCAL test beam experiments exposed different HCAL modules, preceded by an electromagnetic calorimeter prototype, to beams of pions, electrons and muons over a large energy range. The data were compared with GEANT4 simulations using the hadronic physics parametric (LHEP) and microscopic (QGSP) models. The pion energy resolution and response linearity as a function of incident energy derived from the simulations are in good agreement with the data measurement within the large systematic uncertainties in the latter. Transverse and longitudinal shower profiles were studied in the 1996 and 2004 test beam experiments. Pion showers predicted by GEANT4 are narrower than those predicted by GEANT3. Showers predicted by the QGSP physics list (version 2.7) are shorter than those predicted by the LHEP (version 3.6) list, with LHEP predictions being closer to those from the GEANT4-Geisha model. More precise test beam measurements during 2006 will hopefully allow us to tune the default QGSP physics list to the data. Details and results from these studies are described in Section 5.4.

A combined HCAL-CSC test beam experiment in 2004 also provided important data to differentiate between the GEANT4 hadronic physics models in the form of HCAL shower leakage studies. As described in Section 3.3.2.2, a CSC placed downstream of the calorimeter was used to measure the fraction of incident pions that led to a reconstructed hit in the chamber. The LHEP model appears to describe the punch-through probability best over the entire tested momentum range.

2.5.5.4 Muon system validation

Single muons with momenta in the 10 GeV/c–10 TeV/c range have been simulated in the CMS detector using the GEANT3 and GEANT4 packages. While ionization is modelled very similarly in both packages, muon bremsstrahlung, e^+e^- production, and in particular, muon-nuclear interaction are significantly different, due to newer theoretical developments in-

cluded in GEANT4. Multiple scattering is significantly smaller in GEANT4, in agreement with experimental results [48]. GEANT4 results also show an improvement with respect to GEANT3 in the precision of the propagation of the muons along the detector.

The production thresholds on secondary particles in the different regions of the detector were set to a large value to increase performance, but not so large as to kill the particles which originated in a passive region with enough energy to reach the sensitive detectors.

The production of hits in the simulation of the muon system was tested by comparing the Monte Carlo simulation predictions with DT test beam data. The test beam experiment consisted of 2 muon chambers with and without an iron slab in between them, to investigate the effect of the muon showers in the passive material. The analysis, based on muons in the 50–300 GeV/c momentum range, shows that GEANT4 slightly underestimates soft delta ray production in cell volumes, while hard delta rays and electromagnetic showers are correctly modelled. In spite of this discrepancy, local track reconstruction efficiency and resolution is well reproduced by the simulation. For details see Section 3.3.2.1.

2.5.5.5 Forward detectors

The forward detectors, such as the CASTOR and Zero Degree Calorimeter (ZDC), and the Totem telescopes are also incorporated into the simulation framework. They are essential tools for the diffractive and heavy-ion programs. For example, the ZDC is a Cerenkov detector designed to collect any remaining neutral fragments of the colliding nuclei and may be used as a measure of the collision centrality. In pp collisions, the ZDC may be incorporated in the study of forward physics and photon production.

Simulation studies are underway to study issues such as energy resolution and energy leakage. New test beam data will allow for a more systematic validation of the simulation results. Current validation efforts involve the implementation of a RHIC-design ZDC in the GEANT4 simulation to take advantage of existing test beam data. For details see Chapter 7).

2.5.5.6 Parameterized showers

The detailed simulation of electromagnetic showers is computationally intensive. A parameterization of the spatial energy distribution of an electromagnetic shower, based on probability density functions, allows us to speed up the process without compromising the simulation accuracy. A shower parameterization model called GFLASH, based on 3 probability density functions, was developed and used by the H1 experiment [49].

As part of the GEANT4 software distribution, GFLASH is available in OSCAR, the CMS detector simulation package. In CMS, GFLASH is used to parameterize electrons and positrons in the barrel and endcap electromagnetic calorimeter. Comparisons between the GFLASH-based and the full simulation of the energy depositions in the central crystal, and the 3×3 and 5×5 crystal matrices show good agreement at the 1% level, as illustrated in Figure 2.7.

Figure 2.8 shows that the transverse and longitudinal shower profiles are also well modelled by GFLASH to within 1-3$\%$. The GFLASH shower parameterizations allow a time performance gain of a factor of 3–10 in the simulation. The gain in speed depends on the event type, the particle energy and the detector η region. For instance, a single electron or photon with an energy of 100 GeV in the ECAL barrel is simulated 10 times faster using GFLASH.

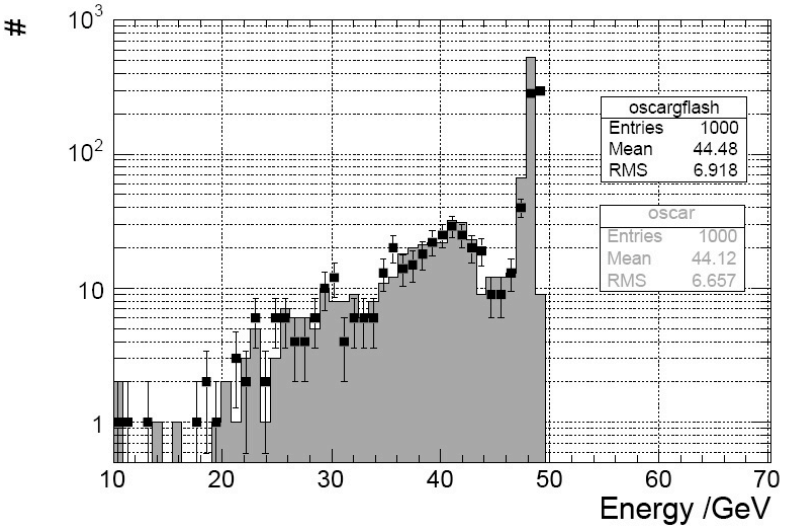


Figure 2.7: Energy depositions in a 5×5 crystal matrix for 50 GeV electrons. The histogram corresponds to the full GEANT4 simulation and the markers to the shower parameterization.

For $pp \rightarrow \gamma + G$ events, with a single photon above 1 TeV, as predicted by models based on extra-dimensions, the gain in speed is a factor of 4.

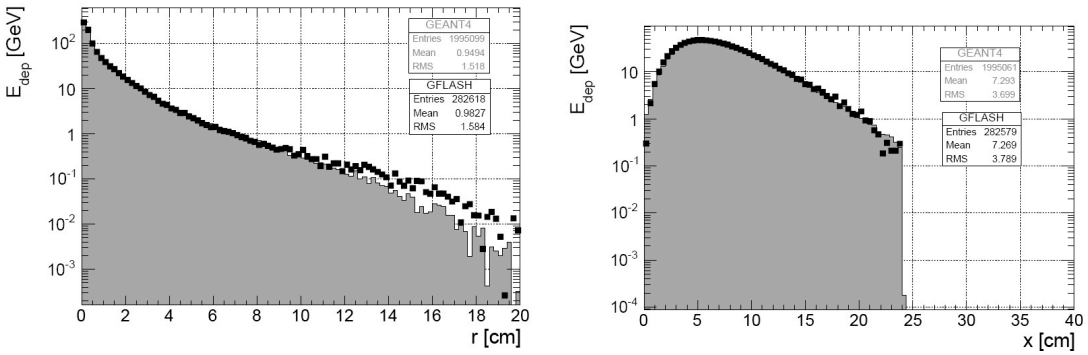


Figure 2.8: Transverse (left) and longitudinal (right) shower profiles for 50 GeV photons.

2.6 Fast simulation

A framework for fast simulation of particle interactions in the CMS detector, called FAMOS, has been recently developed, and is intended to be used for most physics analysis, in view of the Volume 2 of this Physics TDR and beyond. It is an object-oriented system for which C++ has been chosen as programming language. The acronym FAMOS stands for FAst MOnte-Carlo Simulation. As it is a work in progress, only the current status and performance are described in this section.

The input of FAMOS is a list of particles (originating from an event generator or a simple particle gun) characterized by their momentum and origin vertex, with mother and daughter

relationships to follow the various decay chains in the event. Upon user request, each of the (quasi)-stable particles in this list is then propagated in the CMS magnetic field to the different layers of the various CMS subdetectors, which it may interact with. While propagating, these quasi-stable particles are also allowed to decay according to their known branching fractions and decay kinematics. The particles resulting from the interactions with the detector layers or from the decays in flight are added to the original list, and propagated/decayed in the same way. Events from pile-up interactions in the same bunch crossing as the original event are read from pre-generated files, added to the list according to a Poisson distribution with a user-defined average, and follow the same treatment.

The interactions simulated in FAMOS are 1) electron Bremsstrahlung; 2) photon conversion; 3) charged particle energy loss by ionization; 4) charged particle multiple scattering; and 5) electron, photon, and hadron showering. The first 4 are applied to particles traversing the thin layers of the tracker (Section 2.6.1), while the latter is parameterized in the electromagnetic (Section 2.6.2) and hadron (Section 2.6.3) calorimeters. The plans are to simulate all 5 interactions together with synchrotron radiation for the muon propagation through the tracker, calorimeters and muon chambers. The current muon simulation is, however, mostly based on a parameterization of resolutions and efficiencies, as is explained in Section 2.6.4.

As output, FAMOS delivers a series of "high-level objects", such as reconstructed hits for charged particles in the tracker layers, energy deposits in calorimeter cells, which can then immediately be used as inputs of the same "higher-level algorithms" (track fitting, calorimeter clustering, b tagging, electron identification, jet reconstruction and calibration, trigger algorithms, etc.) as in the full reconstruction and analysis package. This parallelism between the fast simulation and the complete reconstruction allows comparisons between fast and full simulations and subsequent tuning in a straightforward manner, with the use of identical analysis programs. It also allows the development of new reconstruction algorithms with the fast simulation, for later use in the complete reconstruction package.

Last but not least, the computer time needed to simulate an event in FAMOS is about 3 orders of magnitude smaller than that needed in the full chain, for a level of agreement aimed at the percent level or below. At the time of writing, the simulation of 1 LHC event in FAMOS in its highest level of sophistication takes a second or thereabout on a 1.8 GHz computer, with expected optimizations still to come.

2.6.1 Simulation of the tracker response

A simplified version of the tracker geometry, but deemed adequate for the required level of accuracy, is used. Is it made of over 30 thin nested cylinders representing the sensitive layers of, from inside to outside, the pixel detector (3 barrel layers and 2 endcap disks), the 4 tracker inner barrel layers, the 3 tracker inner disks, the 6 tracker outer barrel layers and the 9 tracker endcap disks, interleaved with non-instrumented cylinders with dead material (cables, support, etc.). The material, assumed to be pure Silicon, is also assumed to be uniformly distributed over each cylinder barrel (respectively, endcap). A transverse view of this simplified geometry is shown in Figure 2.9a, in which the vertices of converted photons collected in a large number of simulated events are displayed. For comparison, the higher level of details present in the full simulation is visible in Figure 2.9b.

The thickness of each layer was tuned in the fast simulation to reproduce the fully simulated number of Bremsstrahlung photons above a certain energy threshold radiated by energetic

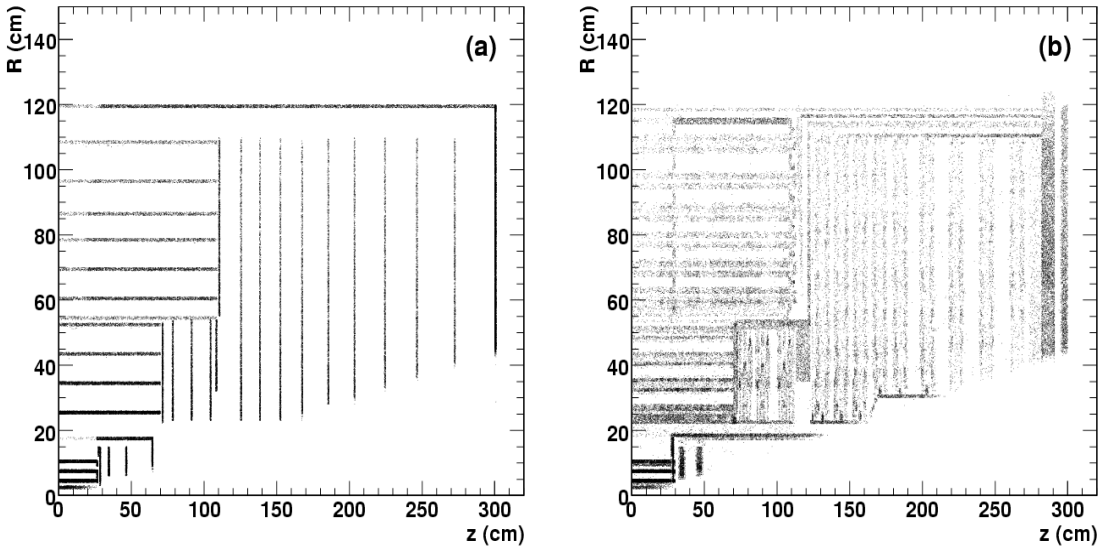


Figure 2.9: A radiography of a quarter of the simulated tracker geometry in the (a) fast and (b) full simulation. The higher level of details in the full simulation is clearly visible.

electrons traversing this layer. The distribution of this number is shown for 35 GeV/c electrons as a function of their pseudorapidity η in Figure 2.10, for the complete tracker as well as for some of its subsets (pixel detector, inner and outer tracker). The same level of agreement between the fast and full simulations is obtained for each single layer. After tuning, the total number of radiation lengths traversed in the tracker reaches $1.42X_0$ at a pseudorapidity $\eta = 1.65$, in agreement with the full geometry. This agreement demonstrates in addition that the Bremsstrahlung model implemented in FAMOS reproduces that of the full GEANT simulation.

While being propagated in the magnetic field through the tracker layers, charged particles experience multiple scattering and energy loss by ionization. The intersections between the modified trajectories and each tracker layer define the position of "simulated hits". Each simulated hit is turned with a certain efficiency to a "reconstructed hit", the position of which is obtained from a Gaussian smearing of the simulated hit position. In the Silicon tracker, the Gaussian resolution in each of the 2 directions (longitudinal and transverse to the beam direction), obtained from a fit of the residuals with respect to the reconstructed charged particle tracks in the full simulation, is essentially a constant for each layer. In the pixel detector, the Gaussian resolution in each of the 2 directions is parameterized according to the pixel cluster size (itself generated according to its fully simulated η-dependent distribution) and on the incident angle of the particles with respect to the layer. This detailed procedure was developed in view of reproducing the b-tagging performance with the requested level of accuracy. The accuracy of the parameterization is illustrated in Figure 2.11, in which the distributions of the reconstructed transverse impact parameter and the corresponding uncertainty for 10 GeV/c single muons generated with a uniform η distribution, are displayed for fast and full simulations.

To save execution time, no pattern recognition is performed to reconstruct charged particle tracks. The reconstructed hits belonging to a given simulated charged particle are, instead, fit to form a reconstructed track, with the same fitting algorithms as in the complete recon-

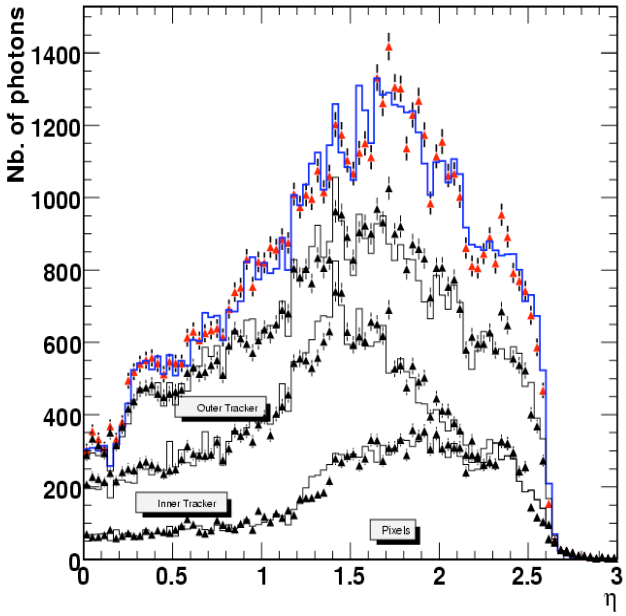


Figure 2.10: The number of photons radiated with an energy above 500 MeV by 35 GeV electrons when traversing the various tracker layers as a function of the pseudorapidity η, for the fast (triangles with error bars) and the full (histogram) simulation. Shown are the distributions after the pixel detector, the inner tracker, the outer tracker and the whole tracker and cables.

struction. This procedure, together with the fact that 2 hits simulated at the same position are never merged into a single reconstructed hit, may have to be revised in a high luminosity environment to reproduce the rate of fake tracks and the reconstruction efficiency.

2.6.2 Simulation of the calorimeter response to electrons and photons

The showers of electrons and photons which impinge on the ECAL are simulated in 2 steps. These electrons and photons may come from the primary vertex, from particle decays, or from interaction in the tracker material.

In a first step, the shower is developed following the well-tuned Grindhammer parameterization [49], as if the ECAL were an homogeneous medium. This approximation is realistic because the CMS calorimeter is made of contiguous crystals. In this parameterization, an electron shower consists of several thousands energy spots, longitudinally distributed according to a Γ function, the parameters of which fluctuate from one shower to the other. The deposited energy is integrated over $2X_0$-thick longitudinal slices, including uncertainties due to the limited photo-statistics and the longitudinal non-uniformity in the crystals.

In each slice, and as a second step, the energy spots are distributed in space according to the radial profile of [49] and placed into the actual crystal geometry, under the realistic assumption that no energy is lost in the small inter-crystal gaps. The time used in this step is kept to a reasonable value by limiting the two-dimensional spot-crystal assignment to a small 7×7 crystal grid (and even smaller for low energy electrons) in a plane perpendicular to the shower longitudinal development. The energy collection simulation is then refined

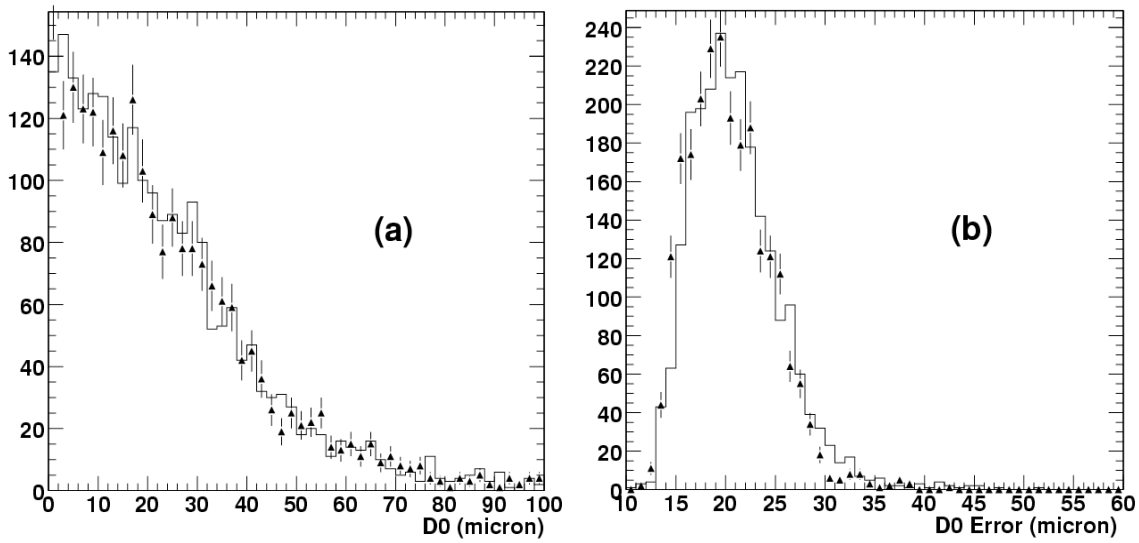


Figure 2.11: Distributions of (a) the reconstructed transverse impact parameter and (b) the corresponding uncertainty, for 10 GeV/c single muons generated with a uniform η distribution, in the fast (triangles) and full (histogram) simulations.

by simulating a number of effects such as rear and front leakage, energy losses in the gaps between ECAL modules, and shower enlargement due to the magnetic field.

In front of the ECAL endcaps, electrons may first cross the preshower. In this case, the corresponding showers are developed from the preshower entrance, and the energy deposit in each layer is converted into a number of MIPs (with related statistical uncertainties), assigned to the relevant strips according to the shower radial profile. Very energetic electrons (above several hundred GeV) can extend their shower substantially beyond the ECAL. A $2X_0$-thick gap between the rear side of the crystals and the entrance of the HCAL is assumed, in which all the energy integrated from the Γ-distribution tail is lost. The rest is assigned to the HCAL towers according to the shower radial profile at this depth, and the energy of each spot is accommodated for an e/π factor.

The Grindhammer parameterization only applies to electrons. Photons are first converted in the ECAL (or preshower) material at a varying depth (according to the number of radiation lengths traversed). Each of the resulting e^+e^- pairs gives rise to 2 separate showers generated as explained above along the same longitudinal direction and, therefore, with the same transverse crystal grids.

Finally, at rapidity values not covered by the electromagnetic calorimeter ($|\eta| > 3$), electrons and photons are propagated directly to the forward hadron calorimeter. Here, the detector response is evaluated from the full simulation of electrons with energies of 30, 100, 300, 1000 and 3000 GeV as a function of pseudorapidity, in a way similar to that explained in Section 2.6.3.

When all electrons and photons are processed, the electronic noise is simulated and the zero suppression applied. At last (when the hadron showers are simulated as well, as is explained in Section 2.6.3), a list of reconstructed hits is built and stored in a format readable by the clustering algorithms, electron reconstruction, etc. Altogether, the various cluster reconstructed energies in the fast and the full simulations agree at the level of the permil in the ECAL barrel (Fig. 2.12a), and at the percent level in the endcaps, for energies ranging from 1 GeV to 1 TeV.

The transverse shower profiles are also adequately reproduced, as is shown in Figure 2.12b, where the distribution of the ratio of the energy in the most energetic crystal to that in the surrounding 3×3 crystal window, S_1/S_9, is displayed.

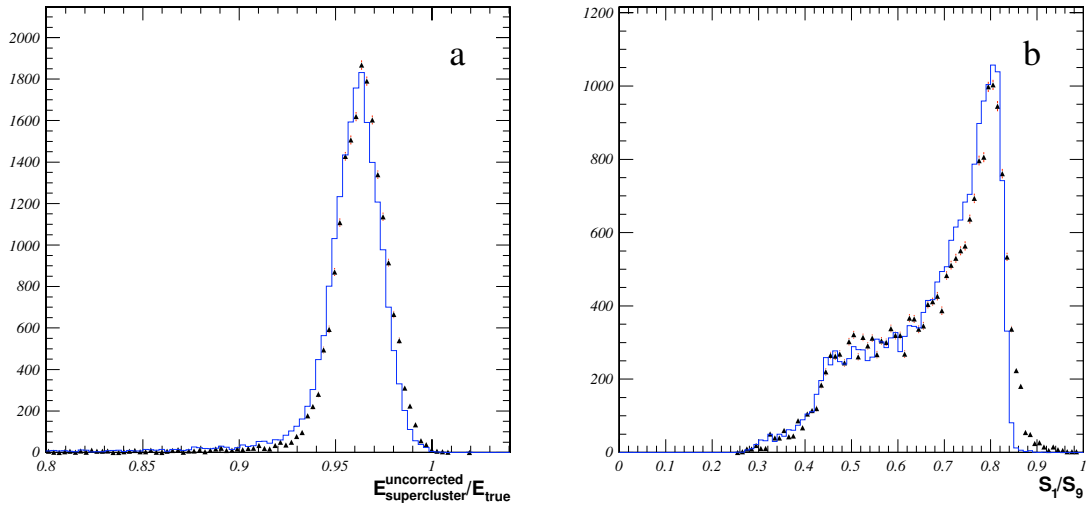


Figure 2.12: Reconstructed supercluster energy over true energy (a) and S1/S9, the ratio of the energy in the most energetic crystal to that in the surrounding 3×3 crystal window (b) for 25 GeV unconverted photons in the barrel, in the fast (triangles with error bars) and full (histogram) simulation. The vertical scale is arbitrary.

2.6.3 Simulation of the calorimeter response to hadrons

Charged and neutral hadrons are also propagated to the ECAL, HCAL and HF entrance after their interactions with the tracker layers. Their energy response is derived from the full simulation in the following way. Single charged pions are fully simulated for p_{T} values of 2, 5, 10, 30, 50, 100 and 300 GeV/c, uniformly distributed in pseudorapidity between -5.0 and $+5.0$. The reconstructed energy is collected in 5×5 HCAL tower matrices and in the corresponding 25×25 ECAL crystals. The energy distributions are then sliced into η bins of 0.1. These distributions are fitted to a Gaussian, the mean value and the sigma of which are tabulated as a function of the energy and pseudorapidity, used in turn to smear the hadron energy response in the fast simulation. Linear interpolation is used for p_{T} in the range from 2 to 300 GeV/c, and extrapolation for particles with transverse momentum outside this range. Figures 2.13a and b show the distribution of the single pion transverse energy as a function of pseudorapidity reconstructed with fast and full simulations, for transverse momenta of 5 and 100 GeV/c.

This smeared energy is then distributed in the calorimeters using parameterized longitudinal and shower profiles, with shower-to-shower variations, following an approach similar to that of GFLASH [50]. The actual implementation in FAMOS is, however, slightly different from that of the original GFLASH, so as to speed up the code and to better adjust the shower shape and the ECAL energy fraction to the full simulation. Distributions of quantities related to the shower transverse shapes are displayed in Figure 2.14 for 100 GeV/c pions in the fast and full simulations.

Other hadrons are simulated as if they were charged pions with the same transverse energy.

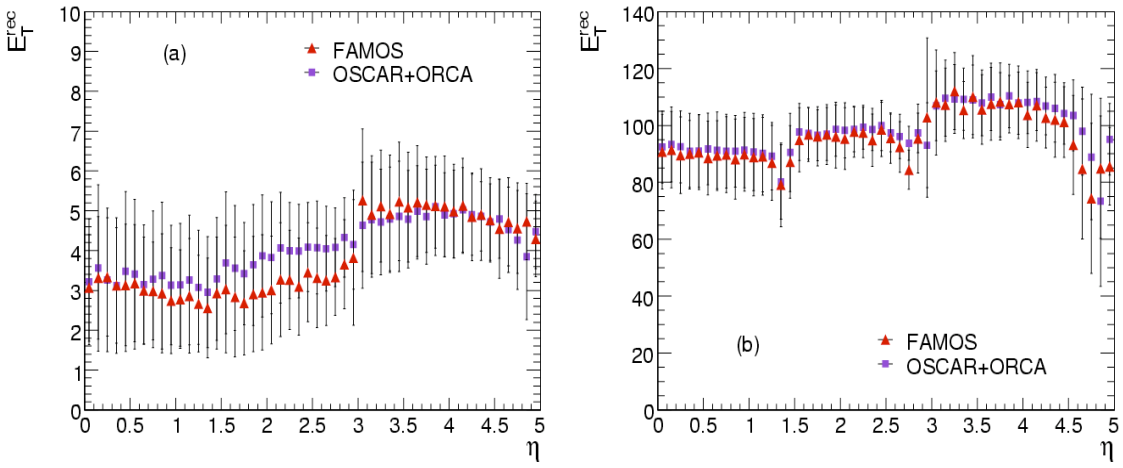


Figure 2.13: Reconstructed transverse energy of pion with (a) $p_{\mathrm{T}} = 5\,\mathrm{GeV}/c$ and (b) $100\,\mathrm{GeV}/c$ as a function of pseudorapidity, in the fast simulation (triangles) and the full simulation (squares). The error bars show the standard deviation of the distribution in each pseudorapidity bin.

Future plans include specific treatments for 1) neutral hadrons, which do no leave, unlike charged particles, any signal in the calorimeters prior to the first nuclear interaction; and 2) protons/neutrons and anti-protons/anti-neutrons, which are expected to behave differently due to particles stopping and anti-particles annihilating in the detector material.

The performance of the fast and full simulations is best compared with hadronic jet reconstruction. Jets are reconstructed from the energy deposits (reconstructed hits) in each HCAL tower and the corresponding 5×5 ECAL crystal window, with the Iterative Cone algorithm and a cone size $R = 0.5$ in the (η, ϕ) plane. No jet energy corrections are applied at this level. Figure 2.15 shows the ratio (mean value and Gaussian sigma) of the reconstructed jet transverse energy to that of the generated jet as a function of jet pseudorapidity, for 2 generated E_{T} values. The agreement between the fast and the full simulations is satisfactory.

The CASTOR detector extends the calorimetric coverage package in the pseudorapidity range $5.2 < |\eta| < 6.5$. In the fast simulation, the energies of the generated particles undergo a Gaussian smearing according to the relative energy resolution curve of the CASTOR prototype obtained after beam tests at different energies (Chapter 7) and from GEANT4 simulations. This Gaussian smearing describes the resolution well. At the time of writing, the CASTOR simulation package considers electromagnetic response while propagation of hadrons is foreseen for upcoming resolution energy dependence to be obtained from further geometry development and ongoing beam tests.

2.6.4 Simulation of the detector response to muons

At the time of writing, the detailed particle propagation of muons in FAMOS stops at the entrance of the calorimeters. The implementation of the calorimeter response to muons is done in a way very much similar to that for pions (Section 2.6.3), with the only difference that the response (mean value and Gaussian sigma) of fully simulated muons is tabulated for transverse momenta of 10, 30, 100, 300 GeV/c, and in η regions corresponding to barrel,

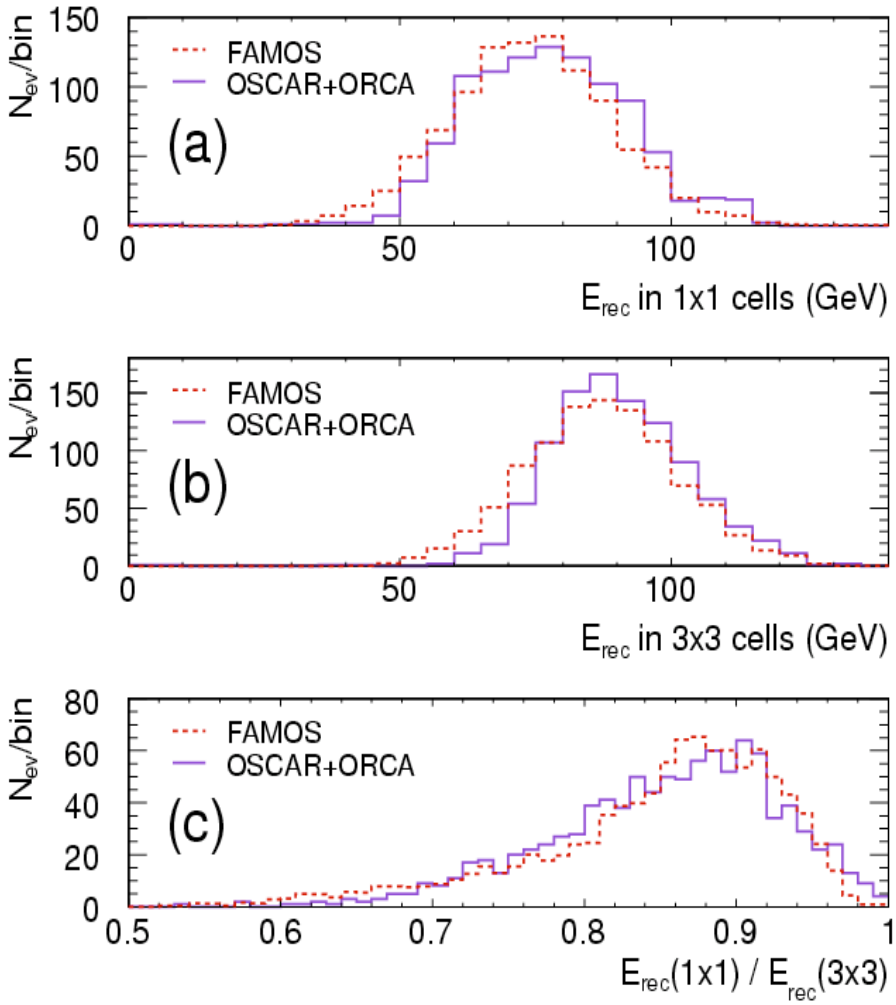


Figure 2.14: Transverse-shower-shape-related distributions for 100 GeV/c pions: (a) Energy in the central tower; (b) energy in the 3×3 surrounding tower window; and (c) ratio of the two.

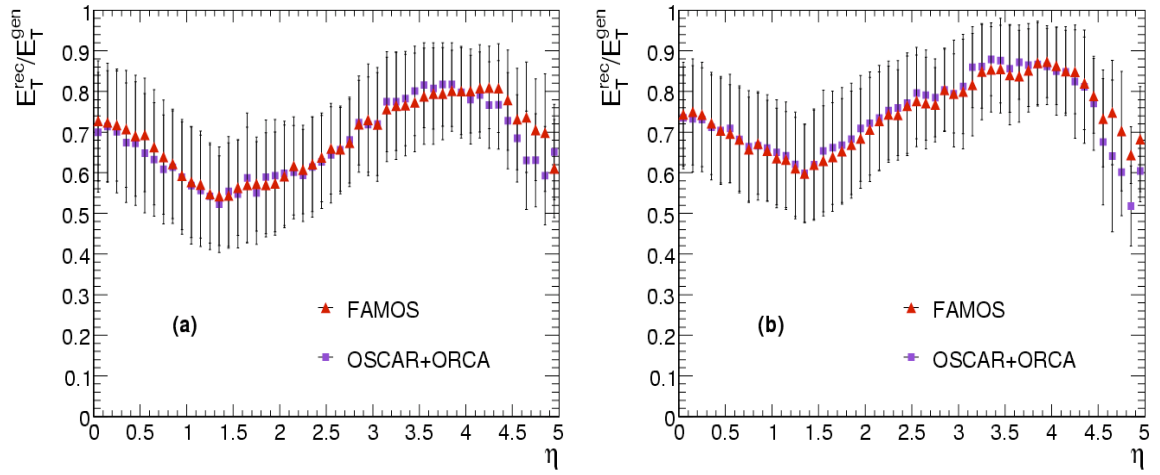


Figure 2.15: Ratio of the reconstructed to the generated jet E_T as a function of jet pseudorapidity for generated E_T values between 20 and 30 GeV (a) and between 80 and 120 GeV (b), in the fast (triangles) and the full (squares) simulations.

endcap and forward hadron calorimeters.

The response of the muon chambers is simply parameterized to reproduce the efficiencies and resolutions of the full simulation. There are plans, however, to extrapolate the tracks through the calorimeters, the coil and the muon detectors, with all pertaining interactions, so that the final tracks can be also reconstructed by fitting all hits, including those in the muon chambers, as is done in the reconstruction of raw data or data from the full simulation.

The muon detectors enter both at the trigger level and in the final reconstruction. In FAMOS, muon parameterizations are available which emulate the output of the Level-1 Trigger, that of the High-Level Trigger (HLT), and the fully reconstructed final offline muon (global muon).

The fast Level-1 Muon trigger relies on 3 quantities, the transverse momentum (with coarse granularity), the pseudorapidity η, and the azimuthal angle ϕ. The Level-1 muon p_T is derived from the "true" particle p_T via a look-up table tuned on the results of the full simulation. The probability of a wrong charge assignment as a function of p_T is also taken into account. It was found adequate to keep the "true" quantities of η and ϕ. Figure 2.16 shows the turn-on efficiency curves in the barrel and in the endcaps for the single muon trigger when the p_T trigger threshold is set at 20 GeV/c, determined from a sample of single muons with uniform distribution in η and ϕ: the curves for the fast simulation superimpose very well with those obtained with the full simulation.

The High-Level Muon Trigger provides quickly reconstructed muons whenever a correlation is found between patterns obtained in the muon chambers and a fast reconstructed track segment in the inner tracker. The parameterization of the algorithmic efficiency and of the p_T resolution is tuned on the full simulation, while for η and ϕ, the resolutions given by the tracker alone are kept. Figure 2.17a shows the corresponding algorithmic selection efficiency as a function of η for single muons in the fast and in the full simulation.

Global muons are identified and reconstructed offline with full tracker and muon information. In FAMOS, the parameterized global muon keeps the correlation with the muon track in

the inner tracker and with the muon information recorded by the muon chambers (i.e., with the Level-1 and HLT information as described above). A parameterization is obtained from a sample of fully simulated global muons. It accounts for the selection efficiency as a function of η, ϕ and p_{T}, and for the Gaussian spread of the reconstructed muon transverse momentum. The pseudorapidity and azimuthal angles determined with the tracker are retained.

To check the overall performance of the above algorithms, events with SUSY Higgs bosons that decay into 2 muons (with $m_{\mathrm{H,A}} = 200\ \mathrm{GeV}/c^2$) were processed through the full and the fast simulations. The dimuon mass spectra thus obtained are displayed in Figure 2.17b. The dimuon (algorithmic) selection efficiencies agree within $\pm 1.5\%$. The reconstructed mean values and widths of the dimuon mass peaks are 200.0 ± 0.1 (200.1 ± 0.1) GeV/c^2 and 3.5 ± 0.1 (3.2 ± 0.1) GeV/c^2 for the full (fast) simulations.

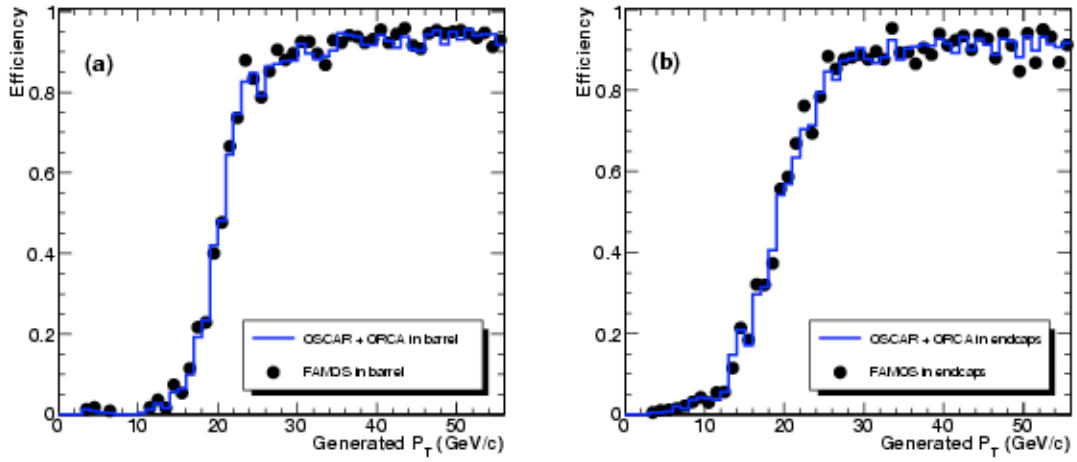


Figure 2.16: Level-1 single-muon trigger efficiency in the barrel (a) and in the endcaps (b) as a function of the generated muon p_{T}, when a threshold of 20 GeV/c is chosen for the Level 1 single muon trigger output, in fast (dots) and full (curve) simulations.

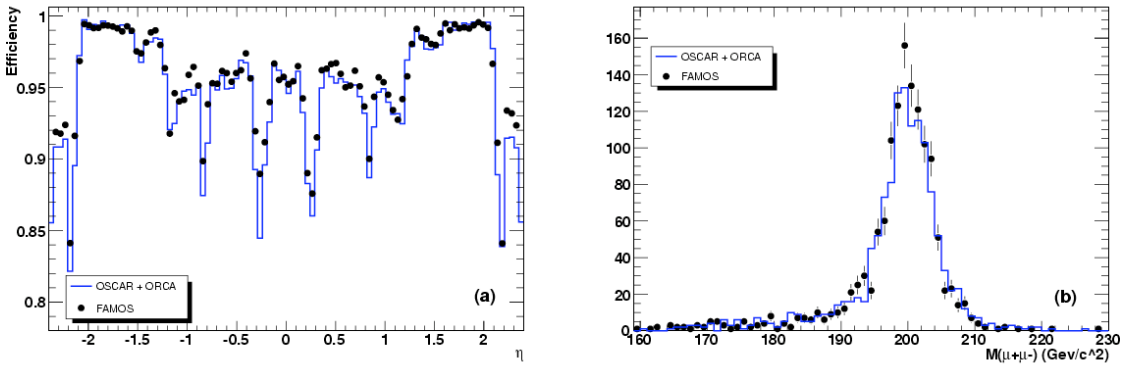


Figure 2.17: (a) Overall Level-1 and High-level trigger efficiency for single muons, as a function of the generated η, and (b) reconstructed dimuon mass spectra in events where a Higgs particle decays into $\mu^+\mu^-$, in fast (dots) and full (curve) simulations.

2.7 Event selection and reconstruction

Reconstruction is the operation of constructing physics quantities from the raw data collected in the experiment. As a software process, reconstruction is therefore the procedure of data reduction whose main client is the data analysis.

The reconstruction process is seen as a collection of independent units, each one providing a set of corresponding reconstructed objects as output. Each reconstruction unit is implemented in the Framework (Section 2.2.2) as an "EDProducer," which retrieves information, in the form of "EDProducts," from the input event, and produces and adds to the event new EDProducts. The following section describes the underlying software design of the reconstruction, whereas the specific algorithms and performance of the reconstruction are described in the detector chapters.

Reconstruction units may use as input

- EDProducts provided by the DAQ system (real data) or by a simulation program (simulated data);
- EDProducts produced by other reconstruction units;
- "EventSetup" information, i.e. event-independent environmental data, e.g. detector description, detector status, calibrations, alignment;
- "ParameterSet" data, i.e. parameters to steer the reconstruction algorithms.

The reconstruction process can be divided into 3 steps, corresponding to local reconstruction within an individual detector module, global reconstruction within a whole detector, and combination of these reconstructed objects to produce higher-level objects.

The reconstruction units providing local reconstruction in a detector module use as input real data from the DAQ system or simulated data representing the real data. These data in either case are called "digis". The output from the reconstruction units are "RecHits," reconstructed hits which are typically position measurements (from times or clusters of strips or pixels) in tracking-type detectors (Muon and Tracker systems) and calorimetric clusters in Calorimeter systems. The RecHits are added to the event as EDProducts, and used as the input to the global reconstruction.

In the global reconstruction step information from the different modules of a subdetector are combined, although information from different subdetectors is not. For example, Tracker RecHits are used to produce reconstructed charged particle tracks and Muon RecHits are used to produce candidate muon tracks. Once again, the objects produced are added to the event as EDProducts.

The final reconstruction step combines reconstructed objects from individual subdetectors to produce higher-level reconstructed objects suitable for high-level triggering or for physics analysis. For example, tracks in the Tracker system and tracks in the Muon system are combined to provide final muon candidates, and electron candidates from the Calorimeter system are matched to tracks in the Tracker system. Figure 2.7 shows how different units/products are combined to form a high level product, and represents an implementation of the flow shown in Figure 2.1.

Reconstruction is performed as a Framework application (Section 2.2.2), in which all steps are explicitly scheduled in advance, and care is taken to ensure that any EDProduct required by a given reconstruction unit will have already been prepared by a prior reconstruction

unit.

Different reconstruction units may produce the same EDProduct, to allow convenient substitution of different algorithms which produce the same objects. The "Provenance" (Section 2.2.2.4) mechanism in each case provides a full record of the generation path of each product.

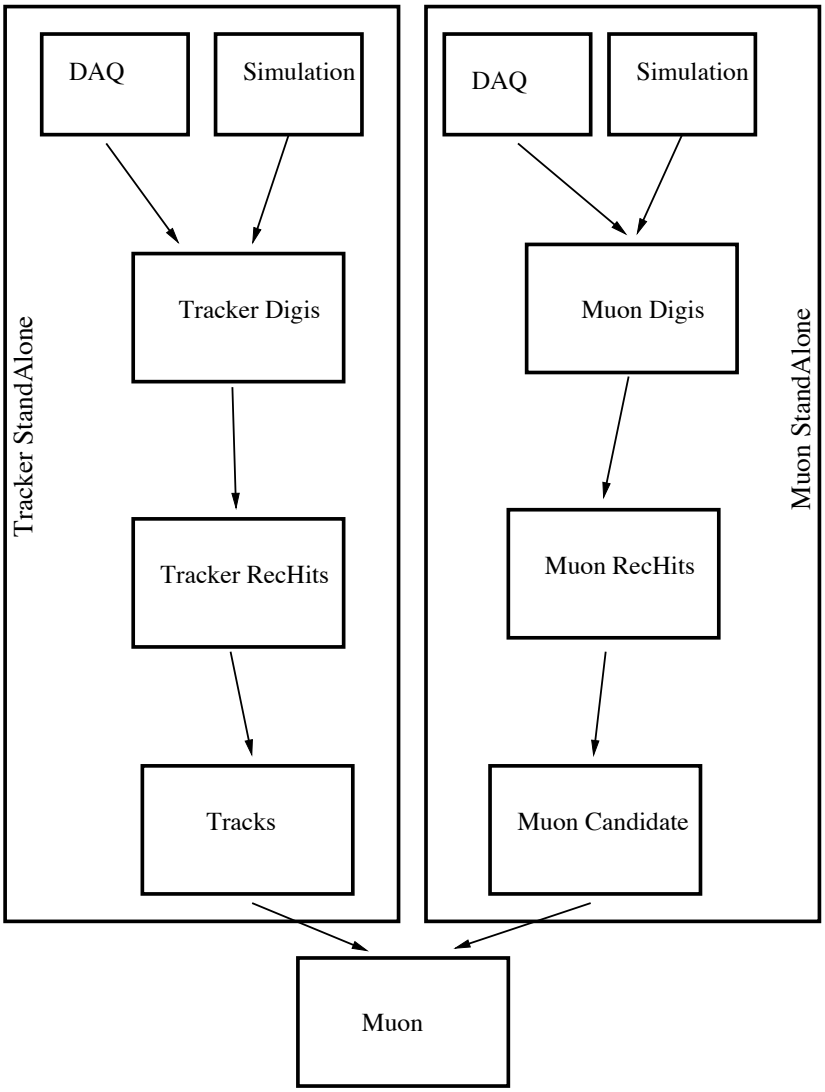


Figure 2.18: Path to produce a Muon using information from Tracker and Muon Systems.

2.7.1 Organization of the reconstruction code

The reconstruction code is packaged in separate modules, providing the granularity necessary for incremental development and release.

The algorithms for local and global reconstruction within a subdetector reside in separate modules. For example, for the Tracker these are RecoLocalTracker and RecoTracker. The product definitions EDProduct are grouped in a module called DataFormats, which is sufficient to interpret data files. The global reconstruction code, as in the third step described

above, resides again in different modules (i.e. RecoTracking, RecoJets, etc.)

The same code is designed to run in the environment of offline reconstruction and the High-Level Trigger (HLT, Section 2.3). In general, the algorithms are the same, while the parameters used for the 2 use cases are different. For offline reconstruction, a global approach is used, in which all the detector parts are analyzed and reconstructed with the best available set of parameters, without constraints on the processing time (at least to first approximation). When the same code is used in the HLT, the set of parameters is optimized to give the best reconstruction within reasonable processing time, and the reconstruction is regional, i.e. only the regions of the detector from which a Level-1 trigger was received are considered.

2.7.2 Local reconstruction

Local reconstruction in individual detector modules leads to RecHits, which contain information about the energy deposition and positions of the particles interacting in the detectors.

In the Tracker detectors (strips and pixels), local reconstruction algorithms search for strips / pixels with a signal exceeding a threshold, and use these as seeds for clusters. Clusters are built by adding neighboring strips/pixels.

In the Muon Drift Chambers (DTs), local reconstruction provides the position of a muon hit in a drift cell, determined from the drift time measurement and the effective drift velocity. Three-dimensional track segments within a superlayer are built from hits in each component layer.

In the Muon Cathode Strip Chambers (CSCs), local reconstruction provides position and time of arrival of a muon hit from the distribution of charge induced on the cathode strips. Two-dimensional hits are obtained in each layer, and these can be combined to create three-dimensional track segments within each chamber (of 6 layers).

In the Muon Resistive Plate Chambers (RPCs), local reconstruction gives the position of a muon hit from the position of clusters of hit strips.

In the Electromagnetic Calorimeter (ECAL), local reconstruction identifies the position, time of arrival, and energy of localized electromagnetic energy depositions.

In the Hadron Calorimeter (HCAL), local reconstruction likewise identifies the position, time, and energy of localized hadronic energy depositions.

2.7.3 Global reconstruction

The global reconstruction algorithms use the objects created in the local reconstruction within a single detector module, combining them with the objects arising from other modules of the same subdetector to produce further objects which represent the best measurement from that subdetector. At this stage, no attempt is made to link the information from different subdetectors; this is part of the later Combined Reconstruction step (Section 2.7.4).

- **Reconstruction in the Tracker system**: in the high multiplicity charged-particle environment of pp collisions at LHC, a global tracking approach is unlikely to be an efficient use of computing resources. Instead, specialized approaches can serve different use cases, e.g. low/high p_{T} tracks, searches for displaced vertices, etc.

In order to satisfy multiple use cases, different reconstruction units must be permitted to run in parallel, each producing a set of tracks applicable to a specific use case. CMS has implemented and tested several different tracking algorithms, each implemented in a common framework in which the various components, from seed finding to propagation and final fit, are cleanly separated and modularized. More than one algorithm is usually available for each component, allowing an easy adaptation to different use cases. Additional information can be found in 6.4.

- **Reconstruction in the Calorimeter system**: a Calorimetric Tower (CaloTower) links matching clusters in ECAL and HCAL to produce a projective tower in the calorimetry system. The towers have a definite position in the (η, ϕ) plane, and hence can be used as the basis for Jet reconstruction as described in Section 2.7.4.3.
- **Reconstruction in the Muon system**: global reconstruction in the Muon system is often called "Standalone muon" since it does not make use of Tracker hits; and it is also used in the Level 2 trigger algorithms. The reconstruction makes use of the track hits and track segments from the local reconstruction step in the individual muon subdetector modules of the CSC, DT, and RPC detectors. The algorithm starts from the locally-reconstructed muon track segments. A segment in one of the innermost detector stations (those closest to the interaction point) is used as a seed for a Kalman filter, which builds trajectories in the radially-increasing direction. A χ^2 cut is applied to reject hits unlikely to be associated with the track, which can arise from showering, delta rays, and pair production. The trajectory is propagated using a detailed map of the magnetic field and taking account of energy loss in the detector material (mainly the steel of the magnet return yoke), until the outermost detector layer of the Muon system is reached. A backward Kalman filter is then applied, working from outside in, and the track parameters are defined at the innermost muon station. Finally, the track is extrapolated to the nominal interaction point and a vertex-constrained fit to the track parameters is performed. In this fit, since the magnetic field is inhomogeneous and nonuniform in the 2 endcap regions, the two-dimensional hits in CSC layers are used instead of the track segments which were used for seeding. Additional information can be found in Section 9.1.1.

2.7.4 Combined reconstruction–high level objects

The final stage of reconstruction combines input objects created in the global reconstruction within each subdetector detector, creating objects based on the complete CMS detector. For example, a standalone muon candidate can be extrapolated into the Tracker detector, thus improving the measured muon track parameters using the high precision of the Tracker measurements. Another common example is matching of ECAL and HCAL clusters and their combination into jet candidates. These examples and others are discussed in the following subsections.

2.7.4.1 Photon and electron identification

The global selection of electrons and photons proceeds in 3 steps. The first step uses the Calorimeter information only. The second step requires hits in the pixel detectors, consistent with an electron candidate. The success of the matching of an ECAL "supercluster"

(Section 10.1) to hits in the pixel detector flags the candidate as an electron; otherwise, the candidate is flagged as a photon. In the final step, the selection of electrons uses full track reconstruction, seeded from the pixel hits obtained by the matching step. The selection of photons can instead use isolation cuts and rejection of π^0s based on lateral shower shape and the reconstruction of converted photons. See Sections 10.3 and 10.4 for additional information.

2.7.4.2 Muon identification

Global muon identification starts from the Standalone muon, adding associated silicon tracker hits and performing a final fit to the track.

Isolation criteria can be applied to the muon candidates to provide additional rejection, using the number of pixel tracks in a region around the projected muon trajectory. This suppresses nonprompt muons from b, c, and K decays. Additional information is available in Section 9.1.

2.7.4.3 Jet reconstruction

Jet reconstruction in pp collisions aims to reconstruct and identify jets arising from the hadronization of a scattered parton, in order to reconstruct its direction and energy. Many reconstruction algorithms exist in the literature, and vary in speed, efficiency, and resolution.

Most algorithms use a clustering technique, in which calorimetric towers close in (η, ϕ) to a high E_{T} tower are summed together, subject to some constraints. For example, in the cone algorithm, a seed tower is selected (typically according to high E_{T}) and then all objects sufficiently close in (η, ϕ) are used to form a proto-jet. The process of association is iterated until the parameters of the proto-jet have stabilized, and then the associated towers are considered to comprise a jet candidate. The procedure is repeated with the remaining unassociated towers, until no seeding tower with sufficiently high E_{T} remains.

2.7.4.4 Missing E_{T} reconstruction

Many channels of discovery at the LHC present as a clear signature for new physics a large missing transverse energy (for example, SUSY decays with a LSP escaping detection by CMS). A large effort has been placed on the design of calorimeters to have as complete η coverage as possible to allow for the needed measurement accuracy.

2.7.4.5 Data tiers in CMS

Within the current understanding, the CMS Event Data Model distinguishes between

- FEVT (Full EVenT): all the data collections from all the producers, plus the RAW data; useful for debugging and probably only at the startup. It is expected to require 1–2 MB/event.
- RECO (RECOnstructed Data): contains selected objects from reconstruction modules; it is still quite large, at the level of 500 kB/event.

- AOD (Analysis Object Data): a subset of the previous, containing only high-level objects; should be sufficient for most analyses and substantially smaller than RECO (50 kB/event).

2.7.4.6 RECO/AOD objects

Objects contained in the RECO data tier are a subset of Event data specifically for analysis tasks. At the current status of development of CMSSW, all the basic objects up to the level of tracks and jets have been defined, and are listed in Table 2.2. The objects are designed in such a way as to provide useful information for analyses without requiring external information, like the magnetic field or geometry. Care is also taken in separating *core* quantities (the ones which must survive in each information tier) from *extra* quantities, which can be stripped out if smaller data sizes are needed. Some higher level objects must still be introduced to fulfil physics analyses; most notably, in Table 2.2, objects for b-tagged and τ-tagged jets are currently missing. These will be implemented soon as an extension to the objects described here (for example, extending the jet object with tagging quantities). A possible solution (one of several possible) is shown in Table 2.3.

2.8 Calibration and alignment

Calibration and alignment information are part of a larger class of data referred to as "non-event data." These data are gathered during the construction and operational phases for the experiment and are needed to fully understand the physics data collected from the detector.

CMS has identified 4 logical categories of nonevent data: 1) Construction, 2) Equipment management, 3) Configuration and 4) Conditions. The data have been classified according to their needs for metadata (data to describe the data) and are described in the following sections.

2.8.1 Different source of technical data

2.8.1.1 Construction data

During production of the detector, data are gathered about the construction process and the produced items. Some of the construction data also belong in the data categories described below and will therefore be moved to the common data storage at the end of the construction phase. The CMS subdetectors will keep their construction data available for the lifetime of the experiment enabling them to trace possible production errors. The construction data and its storage will not be described further in this document.

2.8.1.2 Equipment management data

Detector items should be tracked to log their history of placement and repairs. The classification of CERN as INB (Installation Nucleaire de Base) requires additionally to keep a continuous trace of the location of irradiated items. Equipment management data therefore include these location history for all items being installed at the experiment. They contain detector parts as well as off-detector electronics. The required metadata are the interval of validity (IOV).

Table 2.2: RECO Objects in CMSSW.

System	Object	
TrackReco	Track	Contains *core* track info, like helix parameters, covariance matrix, fit χ^2, number of hits.
	TrackExtra	contains, for a given track, outermost extrapolation and references to reconstructed hits used.
VertexReco	Vertex	Contains position and covariance matrix, number of degrees of freedom, references to used tracks.
MuonReco	Muon	Extends a Track adding Muon System specific objects.
EGammaReco	Cell	Stores energy, η, ϕ, position, layer code.
	BasicCluster	A cluster of Cells; stores energy, position, uncorrected energy, χ^2 and references to Cells.
	Cluster	Contains energy, position, η-ϕ covariance, π^0 discriminator, radius.
	SuperCluster	Contains E_{T}, η, ϕ, position, uncorrected energy, reference to the seeding cluster, reference to the clusters.
	EGammaCandidate	Contains reference to the SuperCluster, plus trigger information.
	Gamma	Contains E_{T}, η, ϕ, primary vertex along z, isolation, reference to the EGammaCandidate.
	Electron	Contains E_{T}, η, ϕ from both clusters and track, track charge, E/p, hadronic fraction, $\Delta\eta$, isolation, references to EGammaCluster and Track.
JetReco	CaloJet	Stores 4-momentum as a TLorentzVector, references to the constituent CaloTowers, energy in ECAL and HCAL.
METReco	$E_{\mathrm{T}}^{\mathrm{miss}}$	Stores missing E_{T} in x and y, corrected and uncorrected.
HLTReco	HLTResult	Contains an array of bits, packed into 16-bit words.

Table 2.3: Additional objects needed for physiscs analyses; one possible implementation.

System	Object	
b-Tagging	BTaggedJet	Contains b-tagging information (and reference to) for a jet.
	JetWithTracks	Contains b-tagging information for a jet, and gives access to the tracks.
Tau Tagging	TauJetWithTracks	Contains a tau jet giving access to the individual Tracks.

2.8.1.3 Configuration data

The data needed to bring the detector to a running mode are classified as configuration data. These comprise voltage settings of power supplies as well as programmable parameters for front-end electronics. The configuration data will have versions and tags of consistent sets of versioned configurations as metadata.

2.8.1.4 Conditions data

Conditions data include several types of information and the use in the online and offline environments are somewhat different. In the online system conditions data are used for post-mortem analysis of detector errors. In the offline the conditions are needed for event reconstruction as well as data quality monitoring. The conditions needed offline are a subset of the online conditions. The conditions metadata are interval of validity, version and tag. Unfortunately a unique classification of data items into the above classes is not possible. A configuration for instance becomes a condition, once it has been used in the detector. The general data flow can be described as follows: configuration data are prepared using the equipment management information and loaded into all of the subdetector systems.

2.8.2 Calibration and alignment data

Calibration and alignment are special cases of conditions data needed offline to characterize the detector running conditions for particular periods of time. The offline conditions subset is extracted from the online nonevent data and sent to the offline sites. Prompt or first time reconstruction of Event data is performed using the offline conditions data. A better understanding of the detector with time will require more careful calibration and alignment by creation of new versions of conditions data.

2.8.2.1 Conditions database architecture

The initial nonevent data repository on the online side is called Online Master Data Storage (OMDS). It holds the configuration and online conditions data of all systems. In addition common data sets like the equipment management data, the detector geometry and the electronics logbook are stored here. The offline conditions subset is cached at the experiment site in a database referred to as ORCON (Offline ReConstruction ON line subset). The conditions data needed by the HLT farm will be accessed from ORCON. Data will then be transferred to ORCOF (Off line ReConstruction OFfline subset) which is the main conditions database for the CERN Tier 0 and Tier 1 computing centres. From ORCOF, data will be distributed to the other tiered computing centres. Calibrated data will be written back to ORCOF and transferred to ORCON if required by the HLT. The underlying database technology for all online and offline databases will be Oracle [51].

2.8.2.2 Condition database online model

Online conditions data comprise data describing the detector configuration, status and logging information. The online conditions data are stored in a relational database with schemas customized for the needs of each subdetector. This enables searching and comparing values

needed to perform error tracking and other diagnostic and operational activities. Each subsystem has specific database tables holding their own conditions data. Only a small subset of the online conditions will be transferred to the offline system.

Because Oracle is used for online and offline databases, intrinsic mechanisms for the transfer and synchronization between on and offline can be used. The only meta-information needed for conditions data is the time of validity, and therefore all data items must have a start and end time associated with them. The end time might be "open ended" showing that this data item is currently valid. To take the ambiguities between configuration and conditions data into account, conditions data might just be a link to a specific configuration, where its IOV indicates the period for which these configuration parameters are valid. The online to offline transfer mechanism will resolve those links, ensure that at least one copy of the corresponding configuration data exists in the offline conditions DB, and create an offline IOV referencing those parameters.

2.8.2.3 Condition database HLT/offline model

Applications using the conditions data, e.g. reconstruction software, get the data from the offline database through the "EventSetup" system in the framework (Section 2.2.3.2). The database software is responsible for providing a data source implementation to the interface defined by the EventSetup. The connection between the data source and the EventSetup guarantees that the data being delivered to the user are valid for the current event.

The Object-Relational (OR) database in POOL (POOL-ORA) [44] is the baseline technology for the CMS offline conditions database software. The POOL-ORA implementation sits on top of a generic relational database layer that hides the underlying relational database technology from the user.

The POOL-ORA interface used for handling nonevent data is identical to that of POOL used for handling event data. The essence of the object-relational approach of POOL is that the transient object model drives the database data model. The designers of the offline data model are unaware of the tabular representation of the data in the database. The offline database schema is automatically created from the definitions of the persistent-capable objects by following the Object Relational Mapping (ORM) rules. The data are retrieved from the underlying relational database then stored as C++ objects in memory by following the SEAL [52] dictionary information and then finding the corresponding entries in the OR mapping files.

2.8.2.4 Tools for transfer from online to HLT/offline databases

A mechanism is being developed for the Online to Offline Database transfer and transformation. It will move and translate the online data to a form that will meet the offline POOL-ORA requirements. This tool provides a generic data movement for all the subdetectors to the databases used for HLT and offline. Data can be transferred from online to offline using a script or compiled application. Since Oracle is employed for both instances of views (a view is a customized slice of a table or set of tables), other Oracle specific approaches might be used. In order to trigger the data transfer process either a "cron" job or some other driving mechanism will be employed. This trigger process will be configurable and might be different for different types of data or various subdetector groups.

Some subdetectors (Tracker) store their data in BLOBs (Binary Large OBjects), whereas subdetectors like HCAL store their data in numeric fields, and data will be selected from the Online database in the format that is needed for the Offline side. This tool also provides a mechanism to monitor the data which is transferred from the Online to Offline database. Moreover a fail safe mechanism is required to check that the transferred data is correctly synchronized. Once data is transformed and transferred to the Offline side to match the POOL-ORA API, these tables will be used to generate POOL-ORA objects for the event reconstruction as well as other physics analysis purposes.

2.8.2.5 Offline access and deployment

Offline access to the detector nonevent data is through a system of servers that connect to the offline database. The approach is that of FroNTier [53] that uses a Tomcat server [54] to extract information from the database through a network of proxy caching servers, specifically Squids [55], to CMS client applications. This approach has many advantages over a system of direct client connections to the database server as it significantly reduces the activity to the central database. All of the servers in the system are stateless, which means they can "failover" (automatically switch) to other redundant servers providing a reliable and "performant" system. Connections to the Oracle server originate from the Tomcat server and are pooled together and carefully managed.

The system is made robust by providing several parallel Tomcat servers at the Tier-0 computing centre that provide load balancing and failover capabilities. The deployment of the system is straightforward with an instance of the Squid proxy-caching server needed at each Tier-1 and Tier-n centre. The clients connect through the Squid proxy-caching servers and data returned to the client are cached locally in the server. For the client, this entire mechanism is hidden behind the POOL API and only a connection string is configured in the offline software to enable it.

The FroNTier approach is considered the baseline solution for distributing calibration and alignment data. An alternative also being considered employs a distributed database system that replicates the needed portions of the database to remote sites. This is done by the replication through Oracle streams, or by MySQL or other database technologies.

2.8.3 Calibration and alignment model

2.8.3.1 Calibration

Calibration is based on detector signal measurements taken under controlled conditions. These data are characterized by several parameters per electronic channel, e.g. pedestal, gain, time offset, drift velocity, dead and hot channels, and others. The data are often collected for sets of channels within a given subdetector, and an algorithm is used to produce a calibration set that is considered valid for a certain interval of time called Interval Of Validity (IOV). The calibration information for any given IOV can have one or more versions of the calculation of the calibration values made with different algorithms. This version information might be made of both an algorithm name and version to allow for different algorithms, and slight changes to a particular algorithm. The procedure varies considerably from detector to detector.

2.8.3.1.1 Calorimeter calibration HCAL uses a radioactive source fed through a small tube to each tower as the starting point of its gain calibration. The signals produced by this source are recorded as a function of the position of the source in the module, and a gain calculated based on an algorithm. Pedestals are measured when the module is completely quiet, with no sources of energy. Electronic noise causes the pedestals to fluctuate and an algorithm is used to determine the best value for each channel. It is anticipated that the HCAL gain and pedestals will not vary significantly over many days or weeks of operation. The ECAL calibration is much more challenging due to the very high precison that is being sought. The target intercalibration precision is $\approx 0.5\%$. Many small effects which are negligible at lower precision need to be treated with care as this level of precision is approached. The nature and technology of the ECAL provides no convenient or *a priori* way of intercalibrating the channels and the target precision can only by achieved using physics events. A further complication is the changing transparency of the lead-tungstate crystals with radiation exposure, due to the formation of colour centres and subsequent annealing. The transparency is measured for each crystal approximately every 20 minutes by injecting laser pulses and reading out the response.

2.8.3.2 Alignment

A precision alignment of the CMS tracking devices—pixel, microstrip, and muon detectors—is one of the most challenging offline calibration tasks CMS faces. A detailed description of the general alignment concepts and the procedures involved for the muon and Tracker are presented in Sections 3.2.2 and 6.6, respectively. This section is devoted to an overview of the application and storage of alignment data in the context of the offline condition service.

A set of alignment corrections for an active detector element used in the reconstruction, such as a silicon sensor or a muon chamber, is defined to be a vector containing 6 terms: δx, δy, δz, and, $\delta\alpha$, $\delta\beta$, $\delta\gamma$ where the first 3 represent the displacements of the 3 translation degrees of freedom while the last 3 define Euler angle corrections for the 3 possible rotations. All corrections are calculated with respect to the ideal geometry as delivered by the geometry service (Section 2.4.1). Therefore, the ideal geometry together with a set of alignment corrections defines a new reconstruction geometry. Similar to all other calibrations that need to be made available in the offline software framework, the alignment corrections are stored as POOL-ORA objects and queried via the standard framework tools.

2.8.3.2.1 Microstrip and pixel detector alignment While track-based alignment is the main source of alignment corrections, also the Laser Alignment System (LAS) of the Microstrip detector will provide important information about the alignment of the high level support structures for the Microstrip detector. Although the alignment corrections stemming from track-based and LAS alignment are usually defined at different levels of detector geometry granularity (i.e. the LAS provides corrections for support structures, while the track based alignment can align individual sensors), the framework provides a consistent treatment of the different alignment corrections to insure a well-defined reconstruction geometry.

Currently it is believed that a full alignment of the Microstrip and Pixel Detector down to the lowest level of granularity will be carried out every 4 to 5 days. This time estimate is mainly based on the number of tracks originating from $Z \rightarrow \mu\mu$ and $W \rightarrow \mu\nu$ decays that can be recorded at low luminosity. In the future it is foreseen to also study other tracks from

central interactions for the alignment. If this is successful, the alignment of the roughly $20\,000$ active elements in the tracker reconstruction can be carried out within a single day.

2.8.3.2.2 Muon chamber alignment Similar to the Microstrip detector, the muon system also possesses a hardware alignment system. Therefore the alignment corrections from both the hardware alignment and the track-based alignment needs to be properly merged in order to obtain the final reconstruction geometry. However, in contrast to the Tracker, the overall number of elements to be aligned is significantly smaller. With roughly 1000 chambers only 6000 alignment parameters are needed for the Muon chambers. It is also expected that the frequency of geometry updates due to changes in the alignment is significantly less for the Muon chambers than for the Microstrip and Pixel Detectors.

2.9 Data quality monitoring

The Physics and Data Quality Monitoring system (DQM) aims at providing a homogeneous monitoring environment across various applications related to data taking at CMS. Its infrastructure must be fairly flexible and easily customizable so as to be usable by different groups across the experiment. Applications that can benefit from a unified approach to monitoring range from the high-level trigger algorithms in the Filter Farm to local DAQ supervision by a subdetector group up to off-line reconstruction jobs carrying out "production validation". The primary goal of the DQM system is however to guarantee the quality of physics data collected by the general data acquisition.

The DQM infrastructure provides a generic interface independent of the specific technology implementation for the creation and update of monitoring objects (e.g. histograms), allowing direct insertion of monitoring statements in the reconstruction code. Producers publish a list of available information to be delivered to consumers upon connection. They accept subscription requests for delivery of regular updates of a given piece of information. The DQM infrastructure provides functionality to collect and organize information received from a number of producers, and redirect it to consumers according to their subscription requests. This interface can be accessed from standalone programs, or can be used from within reconstruction applications and modules. On the client side, tools are provided for evaluating the consistency of received information to reference information retrieved from a database, update these references, set thresholds, raise alarms, and create error messages for use by the central error logging facility.

2.9.1 Architecture

The DQM framework is designed to deal with sets of objects ("monitoring elements", or ME) from the creation in monitoring producers ("sources"), to the organization and redistribution, on a periodical basis, in the "collectors," to their final use by clients. Sources are defined as individual nodes that have either direct access to or can process and produce information we are interested in. The creation and update of MEs at the source can be the result of processing input event data (event consumers) or input monitor elements (monitor consumers).

At the other end of this architecture are the monitoring information consumers ("clients"). Clients are notified of available monitoring information ("monitorables") from all sources

combined. They can subscribe to and receive periodic updates of any desired subset of the monitorables, in a classic implementation of the "publish-subscribe" service. A hierarchical system of collector nodes is responsible for the communication between sources and clients (e.g. subscription requests) and the actual monitoring transfer. These nodes serve as collectors for the sources and as monitoring servers for the clients.

In order to minimize the interference with the "main" application running in the source process (e.g. analysis, calibration/alignment, trigger algorithm, etc), DQM operations at the source are reduced to a minimum. All CPU-intensive tasks (e.g. comparison to reference monitoring element, display, database storage, etc.) are to be carried out at the client's side.

The above design aims at

- shielding the sources from connecting clients that could slow down the main application or threaten the stability of the source
- facilitating the quick transfer of the monitoring information from the sources to the collectors.

To this end, sources are connected to only 1 collector each (but a collector can connect to multiple sources). Clients do not have direct access to the sources. All source-client communication is carried out through the collector (or collectors).

In this design, the production of the monitoring information is clearly separated from the collection and the processing. The collectors act as the "middle man": they are responsible for advertising the monitorables to different clients and serve monitoring requests. The nature of the collection and processing of the monitoring information is statistical by construction. In particular, the DQM

- is meant to help the experts identify problems that occur (and are monitored) *over a period of time* and is not expected to be capable of spotting punctual problems
- does *not* give access to particular events
- does *not* guarantee that 2 clients will receive identical monitoring information

2.9.2 Components and data flow

The DQM infrastructure supports various monitorable types. 1D-, 2D- and 3D-histograms, 1D- and 2D-profiles, scalars (integer and real numbers) and string messages can be booked and filled/updated anywhere in the context of reconstruction and analysis code. The infrastructure takes care of publishing, tracking updates, and transporting these updates to subscriber processes. The DQM infrastructure does *not* provide support for publishing/transport of individual event data. Distribution of data to "event consumers" is provided by a separate system (see Section 2.3.2.2). Access to booking, filling and modifying MEs is provided via abstract interfaces in every component. MEs are organized in directory structures with virtually unlimited depth, from which monitor consumers can "pick and choose". In every component, it is possible at any point in time to create root-tuples with "snapshots" of the monitoring structure for debugging and reference.

2.9.2.1 Sources

Data Quality Monitoring services available to the source not only keep track of updates to existing monitor elements, they also enable dynamic modifications to the monitoring struc-

ture. The list of available monitoring information can be modified at run-time by booking or deleting MEs via the public DQM service interface. Updated information from an individual source is distributed to all consumers (through the collector) by an update task (MonitorDaemon) running periodically is a separate thread of the source process. The interval between 2 MonitorDaemon updates, which defines a monitoring cycle, can be configured for each individual source process. A "reset" switch can be used to specify for each ME whether monitoring contents should be reset at the end of a cycle[2].

The MonitorDaemon maintains the connection with the collector and uses the DQM service to collect updates to be transmitted to it. The main application (which could be a critical one, like a HLT process), is not affected by the failure of a downstream component in the DQM system As an example, the source can continue to run even if the connection to the collector is lost (e.g. the collector has crashed).

2.9.2.2 Collectors

Collectors serve as dispatch points between sources and clients. Unlike sources and clients, collectors are completely standardized and do not need any customization. A collector accepts network connections from sources and clients. A source can post messages to the collector advertising available monitor contents. All connected clients are dispatched with the entire published content available at the collector. The collector receives subscription messages from clients that are relayed to the appropriate sources. When a source sends an update message containing new data, this is relayed to all subscribed clients. Individual sources and clients can be added or removed at run-time. The collector is responsible of keeping track of active connections.

2.9.2.3 Clients

A generic DQM client application is distinct from a source in that it normally only deals with monitor data, and not with event data. Client input comes in the form of updates of subscribed information from one or more collector instances the client is connected to. As mentioned above, connections can be dropped without affecting the overall functionality of the DQM system and its sources. Standard components are provided that allow the client to:

- start and configure itself, making connections to the relevant collector(s)
- load an initial subscription list at configuration; this list can be later edited and saved
- be notified of the data taking configuration (DAQ configuration, trigger tables) and of run start and stop
- subscribe to selected subsets of data available from the connected sources
- add or remove available items from the subscription list at runtime
- receive and keep track of periodic updates of monitoring information from multiple sources
- collate information from different sources

[2]This option should be turned on for MEs that describe dynamic content (e.g. hit occupancy of a subdetector) and off for MEs that describe accumulating quantities (e.g. number of events processed, number of errors, or counters of rare events).

- maintain a local list of available data; this includes all input data, results of collation, and all other monitor data created in the client itself
- create groups of monitor elements for analysis or display
- create and attach rules to monitor elements: these rules are evaluated for each update and used for diagnostic reports

Client applications must be customized for the use by an individual subsystem.

2.9.2.4 Client customization

As discussed in Section 2.9.1, the majority of the operations involving MEs takes place on the client side. Here we list a set of tools used to customize these tasks, accessible only by clients.

- Analysis tools for monitoring element operations: "reset", "accumulate", "collate", "compare to reference".
- Status flags to summarize with a single discrete parameter the status of hierarchical components of a subsystem. This is convenient for summary pages on a GUI that can give the overall status of components subdetector e.g. through a colour-coded system. Problem flags can be set according to rules, alarms raised or masked.
- Archival facility to store (and retrieve) custom sets of MEs to be "played back", used as reference, or for historical analysis
- Graphical User Interface to give interactive access to the custom operations discussed in this section via a standard set of graphical interface widgets, and to provide visual feedback to the user (overall and hierarchical status display, representative plots)
- Display of arbitrary sets of monitor elements

Generic graphic clients are also provided that can be connected to a given application, allowing the standard DQM interface to be visualized as complex live displays of monitor information.

2.9.3 Control flow

DQM applications that are part of the online systems are controlled as a whole by a "DQM Supervisor" which runs under the general Run Control. The DQM Supervisor is responsible for the central initiation of all relevant processes and the transmission of common configuration information to them. All DQM clients implement a standardized state machine and states are constantly reported to the supervisor to give feedback to the shift crew in case problems are detected. Run Control also uses the Supervisor to notify all clients when a run is started or ended.

During the lifetime of a DQM client, interaction with the user happens through individual application control interfaces based on the web protocol.

2.9.4 Subdetector monitoring

Implementation of DQM tools is under active development by the different subdetector groups. The current focus has been on addressing the short-term operational needs of the Cosmic Challenge. For example, a Drift Tubes DQM histogram browser with a GUI based on IGUANA (see following section) has been used to monitor the drift tubes as part of the commissioning effort at SX5 (Fig. 2.19). This development will continue on to the development of monitoring programs for the long-term needs of the subdetectors. A discussion of the amount of information and resources which are necessary to monitor the operation and performance of each subsystem is given in the subsequent subdetector chapters.

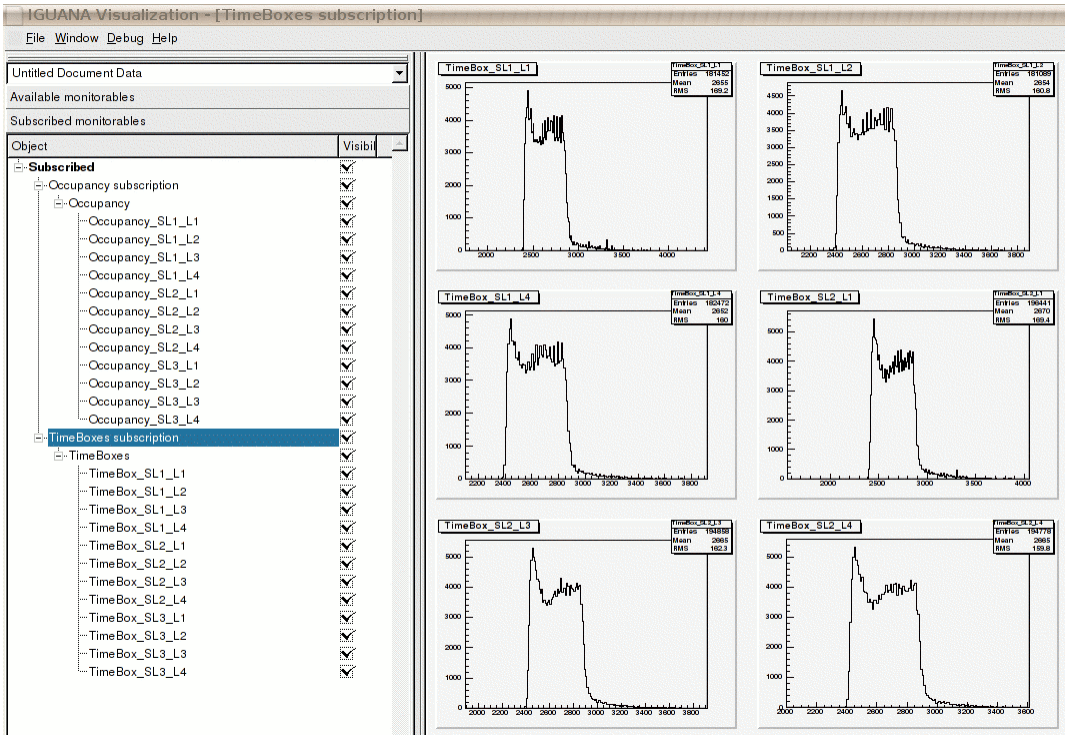


Figure 2.19: Screen shot of the Drift Tubes DQM browser of online data sources: occupancy and time boxes subscription, using the interactive IGUANA GUI and tree controller to display several embedded ROOT canvas components. The data sources are the cosmics muons taken at SX5 during commissioning of the installed DT.

2.10 Visualization

Visualization is an essential part of physics analysis and is an invaluable debugging tool in a wide variety of areas including: offline simulation and reconstruction (event display), data analysis, test beams, and detector monitoring in the commissioning and running phases.

An effective visualization system should act as a highly intuitive tool that allows the user to rapidly navigate through the complex event data structures and their associated contextual data, such as detector configuration. Users should be able to explore the features of the data that are of specific interest to themselves, without being overwhelmed by the other data. Depending on the context, the user should be offered different views of the data, such as: 2D

(slices and projections of events and the detector); 3D (perspective and orthographic views of events and the detector); 1D and 2D histograms (e.g. cell maps or lego plots of calorimeter energies); or text views (e.g. tables of track parameters).

A coherent graphical user interface (GUI) enables the user to control the behavior of the application supporting actions such as: display (selected parts of) event; display (selected parts of) detector; define magnification and orientation of view; and select items in the views (e.g. by mouse) and get feedback (e.g. textual data about selected objects). In order to see only features of interest to individual users, the GUI controls should allow users to control the visibility, colors, and transparency of individual objects and groups of objects and select them either individually or using filters based on, for example, their properties such as name, composition, density.

The CMS visualization philosophy is based on the following tenets:

1. Exploit existing public-domain graphics tools as much as possible.
   This provides better functionality at lower cost (than in-house equivalents).
2. Build a generic toolkit of graphics representations and viewers to support event data and detector visualization needs, based on the afore-mentioned public-domain tools.
   This maximizes re-use of components across applications (simulation, reconstruction, etc.) and environments (online, offline, or even other experiments).
3. Keep experiment-specific software decoupled from graphics software through the use of relatively modest pieces of "adapter" code.
   This enables the CMS systems to evolve and be maintained independently of the bulk of the graphics systems, to facilitate the integration of new graphics toolkits and versions, and to enable new applications to be rapidly developed as new needs arise.

2.10.1 Design of visualization software

The tenets above are reflected in the structure of the CMS visualization systems, which is shown schematically in Figure 2.20. The core of the visualization is the Interactive Graphics for User Analysis (IGUANA) system [56, 57, 58]. This is a generic HEP visualization toolkit, mostly written in C++, which has an emphasis on detector and event display. Although it was primarily developed in the context of CMS, it is deliberately designed to be completely independent of CMS software. Therefore it may be freely used by other experiments (for example DØ has produced an event display using IGUANA). IGUANA relies on a few generally available external graphics packages, notably: OpenInventor and OpenGL for performant 2D/3D graphics [59] and the Qt Graphical User Interface toolkit [60].

IGUANA extends these basic toolkits to provide HEP extensions for the following:

- graphics representations;
- viewers / controller components;
- desktop and web-based graphical user interface clients.

The **graphics representations** correspond to common physics objects (e.g. lines represent tracks, ellipsoids represent vertices, cones represent jets, etc.), as well as detector elements (e.g. trapezoids, conic sections, etc.). The full range of GEANT4 shapes is supported.

The **viewers / controllers** are graphical window elements (widgets) and GUI control elements (e.g. a 2D event display canvas or a tree widget for browsing a hierarchical data

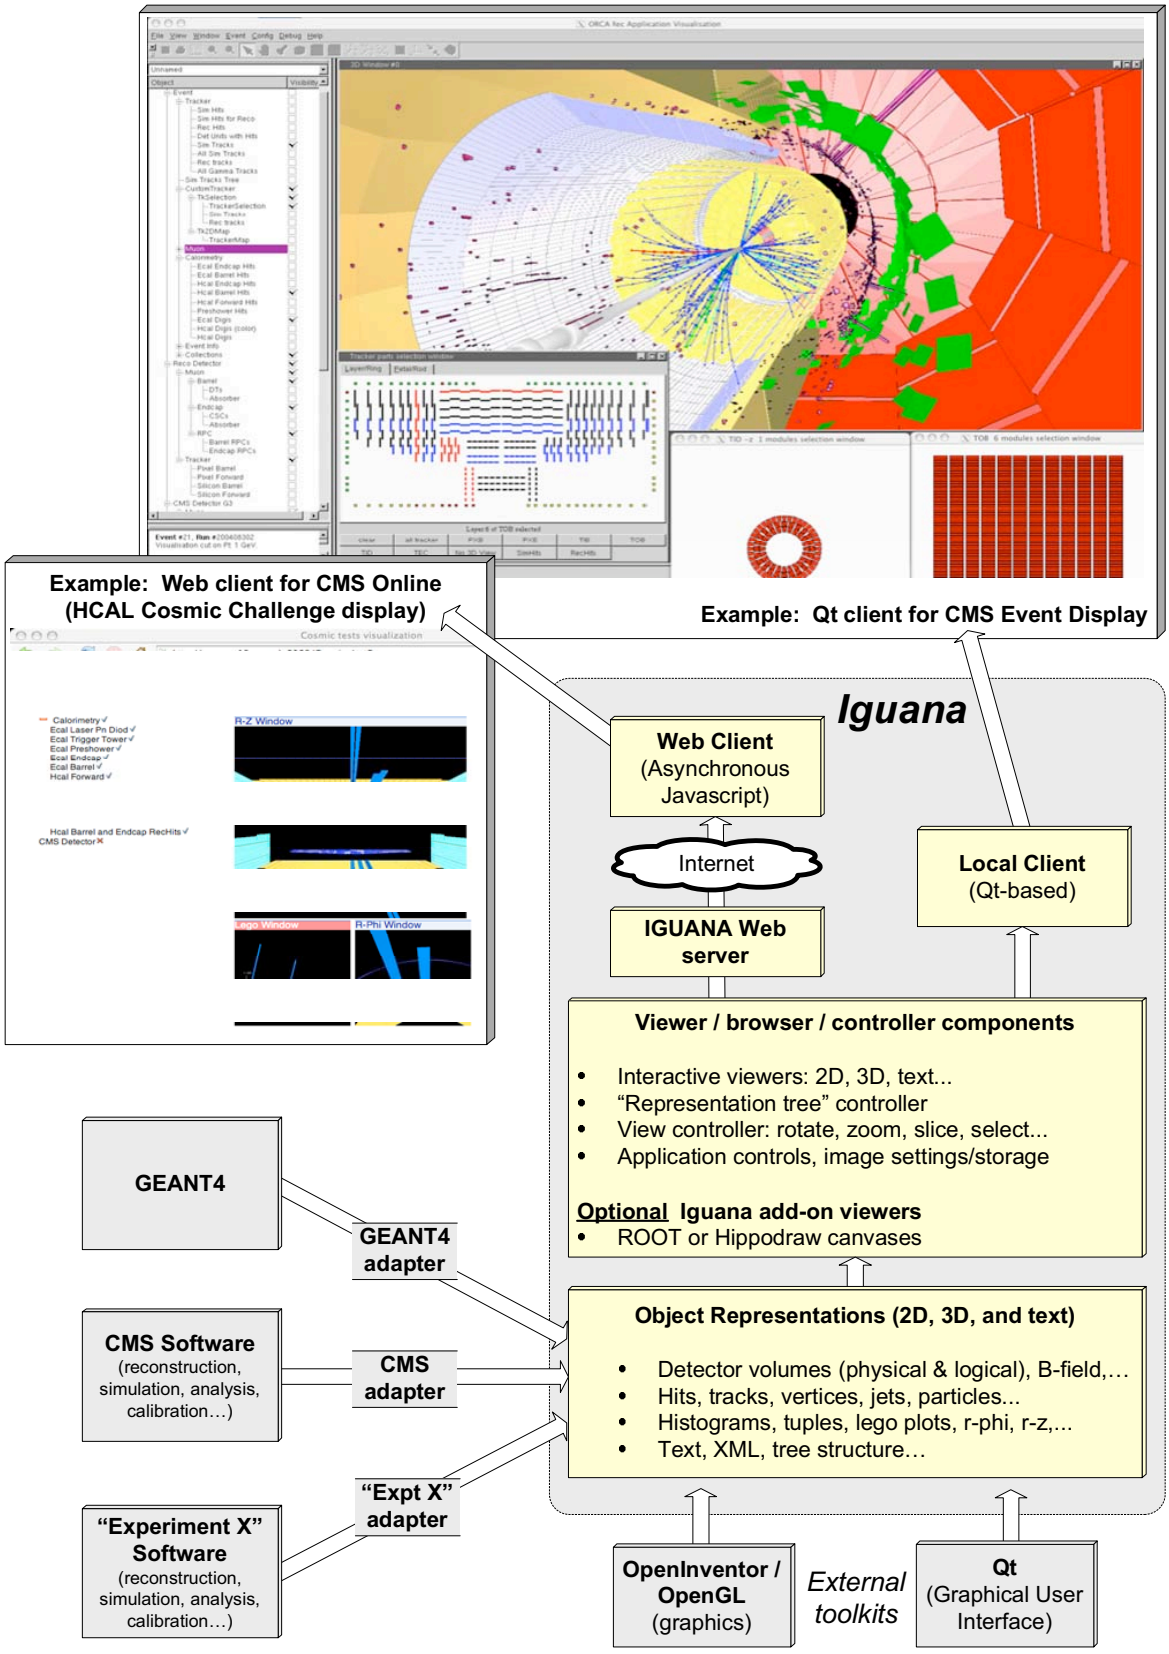


Figure 2.20: Schematic structure of the CMS visualization systems, showing the CMS-specific software, the external graphics systems, and the IGUANA toolkit.

structure). These viewers display and allow the manipulation of arbitrary user-defined collections of graphics representations in 2D or 3D. The user may translate, zoom, rotate, and select objects either by moving the mouse in the graphics window or by using specific widgets.

A broad range of configurable components are provided by IGUANA including:

- a startup wizard to configure an application;
- toolbars for common actions (save, print, select pointer, axes, grids, etc.);
- a multi-document interface for viewers (2D, 3D, lego, Python shell, etc.);
- customizable pull-down menus, tear-offs, and keyboard shortcuts;
- a control centre for animators, clip planes, lights, slices, view properties (perspective, orthographic, wire-frame, solid,...), and camera viewpoints;
- object representation manipulators for colors, transparency, etc.; and
- a configuration parameter controller (e.g. for changing cut settings).

The **desktop and web-based graphical user interface clients** are provided to enable the viewer and controllers to be integrated into a coherent end-user application. The desktop client is a Qt-based system which runs on many flavors of the UNIX operating system as well as the Windows and MacOS operating systems.

The web IGUANA is based on a client-server model. An IGUANA back-end server runs the full application, which creates, manages, and renders the graphics representations. This server produces rasterized images (e.g. jpeg or gif) which are then shipped to the client. The client is a JavaScript component which can run in an arbitrary web browser, just by opening the URL. The web client is particularly useful for occasional and remote users as it requires just a URL and no local installation of CMS, IGUANA, or other external graphics software packages. The use of the asynchronous JavaScript technique enables the client to be highly performant such that the image in the web browser can refresh fast enough to, for example, smoothly rotate the CMS detector in 3D.

2.10.2 CMS visualization applications

CMS visualization applications are constructed by loading a CMS-specific visualization subframework which provides communication between the CMS framework and IGUANA. The application specific "adapters" connect the CMS data structures (provided through the CMS framework) to the graphics representations (provided by IGUANA). The representations are then handled by the IGUANA viewers, controllers, and desktop or web clients. The applications manifest themselves as session types of IGUANA through the mechanism of plugins. The user may run IGUANA either through a GUI or through a command line interface for interactive or batch use; Python scripting is supported but is not yet heavily used by CMS. These and other aspects of creating IGUANA-based applications are described in reference [61] and references therein.

The main CMS-specific visualization applications are described below, including: the interactive simulation visualization system; the detector and event display for reconstruction; and the online event display and data quality monitoring visualization applications.

2.10.2.1 Visualization for simulation

The CMS simulation is based on the GEANT4 toolkit as described in section 2.5. Although GEANT4 provides some basic visualization tools, these use a pipeline model for graphics rendering. This is unsuitable for a highly interactive application with incremental user control and rendering, which requires feedback to the original objects after, for example, their selection with the mouse. As a result of these limitations, the integration of GEANT4 into the CMS interactive graphical environment required some customization of a portion of the GEANT4 visualization framework, manager and command line shell, for example to synchronize the GUI and GEANT4 threads. These features are packaged as a non-CMS-specific optional extension to IGUANA and may be easily used to visualize any GEANT4 detector using the full IGUANA capabilities.

The system allows users to choose which parts of the detector and event are to be visualized. The detector structures may be viewed as a physical or logical tree, to which filters may be applied to select volumes based on name, material or other properties. Specialized visualization modules have been written to help debug errors in the definition of the detector, such as the display of erroneously overlapping volumes or, for example, to calculate the mass of volumes made of a particular material. Figure 2.21 shows the CMS detector visualisation; this particular view shows some details of the endcap muon system.

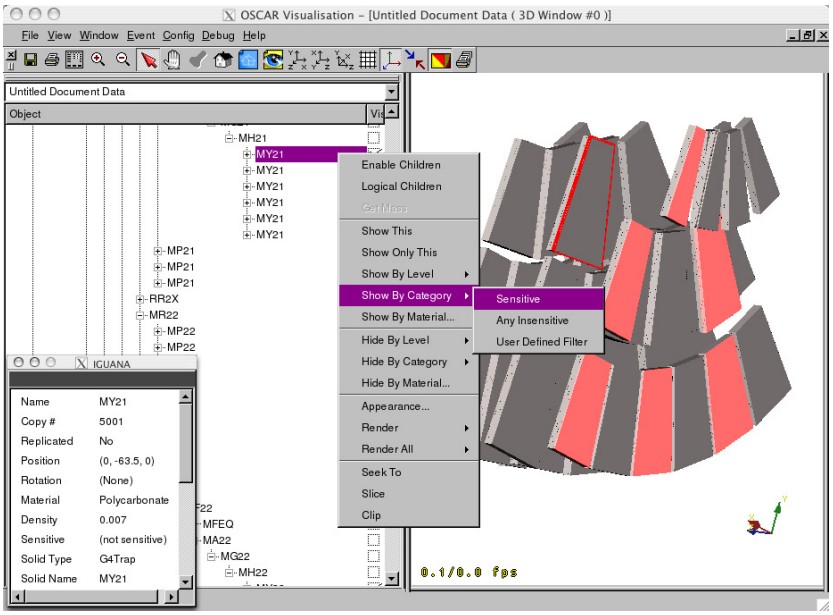


Figure 2.21: Visualization of the CSC "slice test" geometry, using IGUANA, showing the hierarchical detector tree browser, the material filter/selectors, and a 3D viewer of the detector; 2D and magnetic field views are possible but not shown here.

2.10.2.2 Event display for reconstruction

The CMS visualization for reconstruction is used for debugging detector configurations, digitizations, and the event objects created in the reconstruction such as tracks, vertices, jets, etc. It supports different startup configurations according to specific needs of particular detector or physics groups. Applications may or may not load a selective list of plug-ins and

different views according to the task they are to perform. The visualization application for RECO/AOD events is similar to the full reconstruction visualization application, with the difference that the input data collection contains reconstructed objects but no digitizations. Figure 2.20 and figure CP 13 show typical user sessions for the CMS visualization program, with the Qt client, used to study CMS reconstructed data.

2.10.2.3 Visualization for online / data quality monitoring

Visualization is an invaluable tool for giving rapid and intuitive feedback on the performance of the detector. Online displays include generic event displays, similar to the offline programs, as well as more specific Data Quality Monitoring (DQM) displays which show, for example, cell maps for specific subdetectors or histogram displays of selected quantities.

CMS is already starting to develop such displays for the Cosmic Challenge. Figure CP 14 shows one such example for the HCAL online visualization application based on recent real cosmic muons used for debugging the CMS HCAL. It shows examples of the 2D (r-ϕ and r-z) views, a 3D view, and the lego histogram view. It displays a simplified version of the CMS detector as the tracker and ECAL have not yet been installed inside the HCAL at this time.

Displays of various 1D and 2D histogram views are well-supported by tools such as ROOT [62] or HippoDraw [63]. Therefore, rather than duplicate such functionality, component plug-ins to IGUANA were created to wrap the native ROOT and HippoDraw histograms and drawing area (canvas) such that they may be used as standard IGUANA components. In the IGUANA Qt GUI the integration of ROOT with IGUANA is provided through the QtRoot component [64]. For the web client, IGUANA provides a means for creating a rasterized image of the ROOT model that can then be served to the remote web client application.

As discussed in Section 2.9.4, Figure 2.19 shows the Drift Tubes DQM histogram browser with an IGUANA-based GUI. This has a tree browser of available data sources and a set of embedded ROOT canvas components.

2.11 Software development environment, quality assurance, and testing

2.11.1 Description and general concept

The community of contributors to the CMS Software project is a vast and dispersed one. In order to obtain a software cycle that converges rapidly, a basic infrastructure must be provided in order to allow the single developer to work at his/her best, and to integrate seamlessly the work of several contributors. Such infrastructure is the Software Development Environment. Its aim is to coordinate the work of collaborators in the following areas:

- Development: provide a common development environment, a central location for source code and documentation, coding rules and guidelines.
- Integration: provide a schema to facilitate the integration of the work of the various collaborators
- Release: devise a clear release strategy and setup convenient release procedures.

- Distribution: provide the sources and binaries needed to install software in any if the development and/or production centers worldwide.
- Quality assurance: keep the quality of the code under control, prevent regression, track bugs.
- Documentation: ease the process of documenting software units and assure every software unit is properly documented.

In the following sections we will describe in more detail how we intend to achieve these goals.

2.11.2 Code management

In order to enable several developers to contribute to the same project, a code management system must be available to keep old versions of source code files, keep a log of who, when, and why changes occurred, and lastly to allow the consistent retrieval of the code from various authors. In order to make proficient use of the code management system, a set of conventions and guidelines must be established. Once the source code is conveniently managed, the developers must be provided with a convenient build system and package management. Additional development tool are often useful to provide high quality code: memory usage checks, profiling, dependency analysis.

2.11.2.1 Source configuration management

The CVS system has been chosen as the source configuration management tool [65]. Despite the fact that more modern tools exist, CVS is still highly configurable and familiar to most developers.

The source code is structured in a series of **subsystems**, each consisting of one or several **packages**. In order to be able to commit source code for a particular package, a developer must be registered in the developer list for that package. Developer status for a package is granted by the administrator of the associated subsystem. Subsystem administrators have the right to create new packages within that subsystem. Project coordinators create subsystems and grant administrator rights.

2.11.2.2 Coding conventions

Adherence to coding conventions eases the long-term maintenance of the code base. They improve the readability of the code and ease its use and navigation. Adherence of the code to these conventions is guaranteed constantly by the nightly build system (see Section 2.11.2.4), which runs the code checking tools described later in Section 2.11.2.5.

2.11.2.3 Code management tool (SCRAM)

The code management tool `scram` [66] is a configuration management tool, a distribution system, a build system, and a resource manager that manages local resources and applications in a transparent way. In addition it provides a common development environment. `scram` has been used successfully by the CMS collaboration for several years.

2.11.2.4 Nightly builds

In order to allow developers to keep pace with each other without investing time in compiling large portions of code, a nightly build system is under deployment. The choice fell on a tool already successfully used by other LHC experiments [67]. We intend to use this system also as a quality assurance and checking tool, as it will be integrated with the tools described in Section 2.11.2.5 and the testing and validation system described in Section 2.11.9.

2.11.2.5 Code analysis tools

When managing a large project such as the CMS Software, code analysis tools are mandatory. They ease the detection of circular dependencies, help in controlling the quality of the code and its adherence to certain rules or standards, and provide static (configuration) and dynamic (runtime) analysis of the project. The Ignominy suite [68] collects a number of such tools, and is integrated with the nightly build system to keep constant surveillance of the consistency of the code base.

2.11.2.6 Issue tracking

An issue tracking system is essential in a large, distributed project. It is a helpful tool for users, developers and coordinators. For the user, it is a way to communicate to developers and ask for additions, improvements, and bug fixes. For developers, it is a way to keep pace in the development cycle. For coordinators, it is useful to see how far the project has gone and what is missing. The tool of choice for this kind of task is the `Savannah` portal [69].

2.11.3 External packages

The CMS software relies on a number of external packages. Most of them are used and distributed by the LCG Application Area project. We plan however to compile and repackage all the externals in RPM packages to be distributed with the `xcmsi` system (see Section 2.11.6). This will guarantee the reproducibility of the distribution package and the installation procedure. The coherent configuration of the external tools is guaranteed by the configuration management and build system discussed in Section 2.11.2.3.

2.11.4 Platforms and compilers

Currently, the CMS software supports only one platform. Every effort will be made to support multiple platforms, the next candidates being the Linux operating system on the AMD64/IA64 platform and the MacOS operating system. The nightly build machinery will help in building on several platforms simultaneously.

2.11.5 Releases and release strategy

The CMSSW release strategy is being optimized to allow the developers to easily follow the fast development track and cope with the tight release schedule [70]. This strategy is described in detail in [71].

When in production, we plan to distinguish between development releases and production releases. The first will be issued every two weeks, and it will be required to pass a reduced set of testing and validation checks. The latter will be issued in view of production, analysis or data taking/monitoring needs, and will undergo an extended set of validation checks.

Development and production releases will both be associated with a clear workplan in the sense that, for every release, we will plan in advance the features that are expected to be added to the system. The `Savannah` [69] issue tracking system will help the developer to focus on the items required for the next release.

2.11.5.1 Development and release cycle

During the normal development cycle, developers take care of the issues assigned to them in the `savannah` system. Near the time of a scheduled release, each package administrator reviews the goals for that release. At that time, package administrators are requested to publish a tag for the package they are responsible for. The release manager will collect the tags for a given release and provide a deadline. Tags published after that deadline will not be considered for the current release. The release manager will use the tag collector (Section 2.11.5.2) to decide which tags will actually be used for the open release. To ensure consistency of the code, the nightly build system (Section 2.11.2.4) will periodically compile and test the latest published tags, therefore preventing problems from arising at a delicate time such as the build of a release.

2.11.5.2 Tag collector

The tag collector is essentially a web interface to the release management. Package administrators use it to submit tags for a particular release or for the main development stream (nightly build). Release coordinators use it to configure the release by choosing the package versions and checking out a particular version of the code base. We plan to use a modified version of the Web-SRT system previously adopted by the BaBar experiment [72].

2.11.6 Code distribution

For data analysis and user code development in the CMS collaboration, it is important to have an efficient procedure to distribute, install, and update the centrally maintained software. In our case, the aim is to be able to install the CMS software on systems ranging from physicists' notebook computers up to local computing clusters and Grid-enabled clusters. The solution to this problem is a set of two tools: the first one is RpmGen, a Perl script for the automated generation of RPM packages[3], and the second one is xcmsi[73], a tool written in Perl and Perl/Tk, to facilitate the installation, update, and verification of the software included in the RPM packages.

The basic development, the compilation, and the final testing of the CMS software is done centrally at CERN. The CMS software installation comprises experiment-specific programs of the CMS collaboration, common LCG projects[34], and CMS-specific installations of other external programs. Because of the external components which are usually not installed everywhere and since the compilation of the whole source code is very time consuming, the best way is to prepare packages containing readily compiled binaries.

[3] Named after the RedHat package manager.

Concerning the generation of RPM packages with RpmGen, two classes of packages are differentiated: software projects and tools. Software projects are managed by `scram` and have to be registered in a local `scram` database during the installation. Software projects are split into three kinds of packages: source code, documentation, and platform-specific binaries. The tools do not need any special treatment during the installation, their packages usually contain platform-specific binaries only. All packages are built from the central installation at CERN and contain all necessary dependency information and pre- and post-installation scripts, if necessary.

Xcmsi eases the installation of the entire CMS software or selected parts of it on a given computing resource. It provides a graphical user interface for a user-friendly configuration, installation, update, and verification of the software as well as command line tools for batch mode installation, update, and verification. Furthermore the creation and submission of an installation job via LCG2 is supported. Xcmsi was designed to allow the installation of the software on different platforms in an arbitrary directory without root privileges. The configuration file for xcmsi is also used to set the CMS environment for users.

RpmGen/xcmsi has proven to be a reliable and flexible solution to the problem of software distribution and installation. Up to now, RpmGen was used to build around 1000 RPM packages for nearly 240 different CMS projects and versions. Xcmsi is frequently and widely used for around 40 network-based installations of parts of the CMS software per day on nearly 200 sites in 36 countries.

2.11.7 Documentation

Documentation is essential in any software project. We intend to pursue extensive documentation of the system in the following ways:

- User Manuals.
- Reference Manuals (automatically generated with the doxygen tool). These are generated from the source code and require the discipline from the developers to document the code as they are writing it.
- Web pages and "howtos."

2.11.8 User support and mailing lists

We envisage to provide user support through two main channels:

- Issue tracking system, for bug reports and requests for enhancements.
- Mailing lists, for a more human-like support, as in the case of advice, rationales, etc.

2.11.9 Software testing and validation

2.11.9.1 Description and general concept

In a large project such as the CMS software system, testing and validation are essential, indispensable steps in the development cycle. With this in mind, we anticipate that our code

can be divided into 2 main categories: *infrastructural code* and *physics* code. Infrastractural code is suited for **unit testing**. Physics code is suited for **validation**.

A unit test is a method of testing the correctness of a particular module of source code. The idea is to write test cases for every non-trivial function or method in the module, so that each test case is separate from the others if possible. In this scenario, a bug is due to the absence of the test case that checks the correctness of the corresponding code. When a bug is found, the corresponding test should be added in order to prevent regression and check that the correct fix was provided. Among the benefits of unit testing are the following [74]:

- Allows painless refactoring and prevents regression. A unit test provides a written contract of what a particular code unit will do. Programmers can refactor, optimize and simplify the code, and test that it still does what it was expected to do.
- Simplifies integration. Unit testing helps eliminate uncertainty in the pieces themselves and can be used in a bottom-up testing style approach. Testing the parts of a program first and then testing the sum of its parts will make integration testing easier.
- Simplifies collaboration. Each developer can make sure his changes do not break the code from another developer by running unit tests.
- Documents the code. Unit testing provides a sort of ”living document” for the class being tested. Clients wanting to learn how to use the class can look at the unit tests to determine how to use the class to fit their needs.

Not every portion of code in the CMS software project is suited for unit testing. We can take the case of a tracking algorithm: there is no single correct answer when reconstructing a track. The goodness of a tracking algorithm when compared to another can be estimated, for instance, from their resolution. In this case, we don't need to test that the functionality of a code unit is the expected one, but rather we need to check the quality of the result. This is where validation comes in place: the developer can compare the result of his algorithm with a reference one, and make sure the resolution is maintained above a certain standard.

Use of unit testing and validation simplifies the release procedure. The release manager can detect early if a release is consistent by running unit and acceptance tests.

Unit and acceptance testing are an essential component of the Quality Assurance program.

2.11.9.2 Testing framework

It is convenient for developers to write unit tests in a uniform way and have an error reporting system. CppUnit [75] is the testing framework of choice for LCG. For every failing test case, the CppUnit framework reports the line of code where the failure occured and the code that failed, making it easy to detect problems.

The CMS policy is to have a CppUnit test for every class that deserves one. In particular, code units related to framework and services should have a unit test. A dedicated build target allows us to run all the unit tests and inspect their output.

2.11.9.3 Test coverage

We plan to use test coverage tools to find out which parts of the code are covered by unit tests. This will also help in assessing some basic performance statistics, such as:

- how often each line of code executes
- what lines of code are actually executed
- how much computing time each section of code uses

For this purpose, we will use the `lcov` testing suite, which is supported by LCG [76].

2.11.9.4 Physics validation

As mentioned earlier, validation of physics related code requires a different approach. Monte Carlo simulations which include the detector simulation, digitization, and the physics object reconstruction are a critical component of physics analysis. The validation of the simulation sofware is therefore essential to guarantee the quality and accuracy of the Monte Carlo samples. CMS is developing a simulation validation suite consisting of a set of packages associated with the different subdetector systems: tracker, electromagnetic calorimeter, hadron calorimeter, and muon detector. The suite also contains packages to verify detector geometry parameters and the magnetic field. The simulation validation suite uses the OVAL [77] software to control the execution of the packages as well as the comparison tests and "difference" analysis. Each package consists of one or more tests running on single-particle or collider samples, producing distributions of validation quantities which are checked against reference values, using either χ^2 or Kolmogorov-Smirnov tests. The tests are performed at different levels or modes, verifying different components from basic objects such as "hits" to more complex physics quantities such as resolutions and shower profiles, depending on whether the validation is performed on pre-release, regular release, or production software. "digis" and reconstructed physics objects such as electrons, photons, muons, and jets will be incorporated into the validation process in the near future.

Chapter 3

Muon System

3.1 System overview

The 3 types of gaseous detectors used to identify muons and measure their momenta, and the reasons these detector systems were selected, were described in Chapter 1. In this section, after a short description of the detector layout, we discuss the changes made since the Muon TDR [4] and describe the main components of the readout system, giving information on the data format relevant to the reconstruction task.

3.1.1 Detector layout

The Muon system, hosted in the magnet return yokes of CMS, is divided into a central part (Barrel Detector, $|\eta| < 1.2$) and forward parts (Endcap Detector, $|\eta| < 2.4$)(Fig. 1.6). Each Endcap Detector consists of 4 disks that enclose both ends of the barrel cylinder.

3.1.1.1 Barrel detector

The Barrel Detector (Fig. CP 3) consists of 4 concentric "stations" of 250 chambers inside the magnet return yoke of CMS, which is in turn divided into 5 wheels. Each wheel is divided into 12 sectors, each covering a 30° azimuthal angle. Wheels are labeled consecutively from YB-2 for the furthest wheel in $-z$ to YB$+2$ for the furthest is $+z$, while sectors are labeled in order of increasing ϕ beginning with the sector centered at $\phi = 0$. Sectors 3 and 4 in wheels YB-1 and YB$+1$, respectively, host the chimneys for the magnet cryogenic lines: all the chambers in these sectors are shorter by 40 cm along the beam direction than the chambers in the other sectors.

The 2 innermost stations, named MB1 and MB2, consist of "sandwiches" made of a DT chamber placed between 2 RPCs. The 2 outermost stations, MB3 and MB4, consist of packages of a DT chamber coupled to a layer made of 1, 2, or 4 RPCs, depending on the sector and station, placed on the innermost side of the station.

Each DT chamber in the 3 innermost stations, MB1–MB3, consists of 12 layers of drift tubes divided into 3 groups of 4 consecutive layers, hereafter called SuperLayers (SL) (Fig. 3.1). The tubes inside each SL are staggered by half a tube. Two SLs measure the r-ϕ coordinate in the bending plane (they have wires parallel to the beam line), and the third SL measures the z-coordinate running parallel to the beam. A honeycomb structure separates an r-ϕ SL from the other 2 SLs. This gives a lever arm length of about 28 cm for the measurement of the track direction inside each chamber in the bending plane. In the outermost station, MB4, each DT chamber has only the 2 SLs that measure the r-ϕ coordinate.

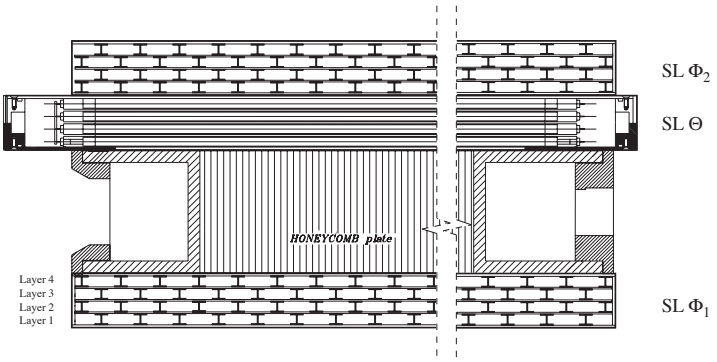


Figure 3.1: The layout of a DT chamber inside a muon barrel station.

In the DT subdetector, an important modification of the basic element of detection, the drift-tube cell, led to a slightly wider drift cell and a new design of the cathode I-beams that separate the drift cells, resulting in a mechanically more robust chamber [78]. In addition, the wire pitch and hence the cell size was increased from 4.0 to 4.2 cm to optimize the electronic segmentation and acceptance, leading to a reduction in the total number of channels from $192\,000$ to $172\,000$. Analysis of test-beam and cosmic-ray muon data collected on both prototype and full-size final chambers has shown no degradation of the performance of the detector in terms of linearity of response, time resolution, noise level, or efficiency of hit collection [78, 79, 80].

Each RPC detector consists of a double-gap bakelite chamber (Fig. 3.2), operating in avalanche mode [81]. The gaps have a 2 mm width. Trigger requirements demand that in each station the strips, which run along the beam direction, be segmented into 2 parts for stations MB1, MB3, and MB4; in the MB2 station, either the innermost or the outermost layer is segmented into 3 parts, depending on the position in the wheels and sectors. The strip length is thus 130 cm, except in the MB2 station where it is either 85 cm or 130 cm. There are 480 RPCs in the Barrel, for a total of 1020 double-gap modules.

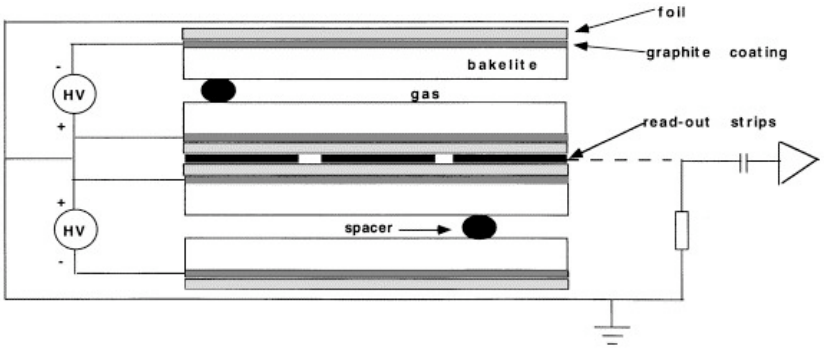


Figure 3.2: Schematic view of the RPC double-gap structure. The read-out strips in the Barrel chambers run along the beam direction.

3.1.1.2 Endcap detector

There are 468 CSCs in the 2 Muon Endcaps. Each Endcap consists of 4 "stations" of chambers, labeled ME1 to ME4 in order of increasing distance from the collision point, which are mounted on the disks enclosing the CMS magnet, perpendicular to the beam direction. In each disk the chambers are divided into 2 concentric rings around the beam axis (3 for ME1 chambers) (Fig. CP 4).

Each CSC (Fig. 3.3) is trapezoidal in shape and consists of 6 gas gaps, each gap having a plane of radial cathode strips and a plane of anode wires running almost perpendicularly to the strips. All CSCs, except those in ME1/3, are overlapped in ϕ to avoid gaps in the muon acceptance. There are 36 chambers in each ring of a muon station, except for the innermost (highest η) rings of ME2–ME4 which have 18 chambers. The gas ionization and subsequent electron avalanche caused by a charged particle traversing each plane of a chamber produces a charge on the anode wire and an image charge on a group of cathode strips. Thus, each CSC measures the space coordinates (r, ϕ, z) in each of the 6 layers.

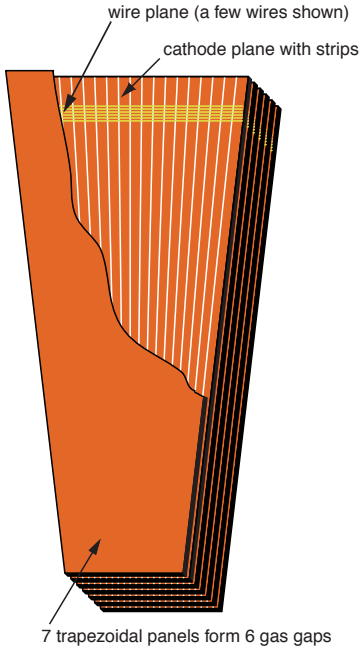


Figure 3.3: Schematic view of a CSC chamber.

There have been several important changes in the CSC system from the set-up described in [4]. The 10°-chambers in the fourth station (ME4/2) have been staged, reducing the total number of chambers from 540 to 468. The number of wire channels per plane for ME1/2 chambers was increased to 64, while for ME1/3 chambers it was decreased to 32. The ME1/1 chambers, which have split cathode strips, now have 48 strips for the higher-rapidity part ($|\eta| > 2.1$) and 64 strips for the lower-rapidity part. The Anode Local-Charged-Track (ALCT) trigger boards are placed on the top face of the chambers, while the corresponding Cathode Local-Charged-Track (CLCT) trigger boards have been combined with the Trigger MotherBoards (TMB) and moved into the peripheral crates placed around the edge of the steel disks. The higher-rapidity cathode strips of the ME1/1 chambers have been removed from the muon trigger and ganged together in groups of 3 (every 16th strip is combined) before

being connected to the cathode front-end electronics. The ME1/1 chambers now also use exactly the same anode and cathode electronics as the other CSCs.

Like in the Barrel, there are layers of double-gap RPCs in the Endcaps, however, for the initial low-luminosity run there are RPCs only in the outer rings of each station, while they are staged in the internal rings. The RPC Endcap system is thus limited to $\eta < 1.6$ for the first period of data taking.

3.1.2 Readout

3.1.2.1 DT chambers

The DT-chamber local trigger and readout electronics reside in a Mini-Crate (MC) mounted on the side opposite to the HV supply of each chamber, inside one of the "C" profiles surrounding the honeycomb structure shown in Figure 3.1. The MC hosts all the boards dedicated to the first level of trigger electronics: the Bunch and Track Identifier (BTI), the Track COrrelator (TRACO), and the Trigger Server (TS) devices [82]), and the Readout Boards (ROB), each equipped with four 32-channel High Performance Time to Digital Converters (HPTDC) [83]. There are 5 to 7 ROBs in each MC depending on the chamber size (the number of readout channels per chamber varies from 608 (MB1) to 800 (MB3)).

The data from the ROBs are sent via LVDS signals (using 2 cables per chamber) to a Readout Server Board (ROS) hosted in VME crates situated on the balcony of CMS. The ROS collects all data from one sector (up to 3268 TDC channels) and sends them to a Detector Dependent Unit (DDU), which is the detector front-end to the CMS Data Acquisition System. There are 5 DDUs (one for each Barrel wheel) for the DT system, each one collecting inputs from 12 ROSs. Trigger data are also sent via 2 LVDS cables from the MC Trigger Board to the Sector Collector Board (SCB) hosted in the same VME crate, which packs the data from 4 chambers (5 in sectors 4 and 10) and sends them via optical link to the Regional Muon Trigger.

3.1.2.2 CSC

There are 2 data paths for the CSCs: one for the cathode strip data and one for the anode wire data. The signals from the anode front-end boards (FEB) go to the Anode Local-Charged-Track (ALCT) trigger board located on the face of each chamber. The ALCT looks for tracks from the wire hits in the 6 chamber layers that point back toward the vertex. It sends its results to one of the Trigger Mother Boards (TMB), which sit in 60 peripheral crates located around the edge of the disks of the magnetic flux return yoke.

The cathode strip pulse heights from the cathode FEBs are sent to the Cathode Local-Charged-Track (CLCT) trigger logic, which is located on the TMB. In a similar way to the ALCT, the CLCT looks for tracks in the hit patterns from the 6 cathode strip layers in a chamber. The TMB tries to match the ALCT and CLCT tracks that were found in a chamber, and sends the results to a Muon Port Card (MPC). The TMB also passes all the anode and cathode raw data on to the Data-acquisition MotherBoard (DMB). Both the MPC and DMB are located in the same peripheral crate as the TMB. The DMB is responsible for initiating all data digitization and readout, as well as providing slow-control signals to all the front-end electronics. The DMB buffers the raw data and sends them via optical fibres to a CSC Detector Dependent Unit (DDU) in USC55. Each DDU collects, combines, and checks the data from 13 DMBs. A

Data Concentrator Card (DCC) then merges the information from 9 DDU boards and sends it via 2 S-LINKs [41, 42] to the overall CMS data acquisition system.

3.1.2.3 RPC

Analog RPC signals are discriminated at the Front-End Boards (FEBs) attached to the chambers. They are then sent as LVDS-standard signals through copper cables to the Link Boards (LB) located on the CMS balconies. One LB receives up to 96 signals from the strips of one chamber. The Link Board electronics synchronizes the signals to the 40 MHz clock and compresses the data (zero suppression). Then the data from up to 3 LBs are multiplexed, converted to optical signals, and transmitted through optical fibers to the Trigger Boards (TB) located in the CMS counting room. One TB receives up to 18 optical links; signals from most of the links are split into 2 or 4 TBs. On the TB, the data from the links are deserialized and transmitted by 18 parallel buses to the RPC trigger Pattern Comparator (PAC) mezzanine boards and the Readout Mezzanine Boards (RMB) [84]. On the RMB, the data are demultiplexed (selected channels can be masked to avoid duplication of data) and stored in a FIFO memory, awaiting a trigger signal. The data originating from the bunch crossing of a trigger are transmitted via optical links to the Data Concentrator Cards (DCC). Three DCC boards concentrate the optical links from all TBs, each DCC taking data from 36 RMBs. The Event Merger (EM) on the DCC merges data from 36 channels into one packet, and the Event Builder (EB) builds events for each trigger, conforming to the Common Data Format (CDF). The packet is then sent to the DAQ by a Slink64 LVDS interface.

3.1.3 Data formats

For the DT system, at each L1A signal the TDC FIFO buffers containing the measured drift time for each recorded hit in the drift tube cells, within a programmable time window of the order of 20-30 bunch crossings, are transferred to the ROS. The drift times are given in TDC units corresponding to a programmable time resolution (in the test-beam DAQ this value was set to 25/32 = 0.781 ns). Multiple hits in a cell are efficiently recorded if their time separation is larger than a preset value (typically 150 ns). The trigger quality information from the Sector Collector [7] is included in the DAQ buffer. Complete data from the chamber trigger system, which are input to the Level-1 DTTrackFinder, are available in the DTTrackFinder DDU buffer.

The ROS/ROB/TDC numbering is not convenient for direct use in the hit reconstruction. It is thus transformed in the DAQ Filter Farm into the "digi" format given in Table 3.1. This is done for each hit in a given chamber, using the mapping stored in the DT Configuration Database (see Section 3.2.4.1).

Table 3.1: The DT "digi" data format used by the reconstruction software.

name	range	number of bits
wire	0–95	7
layer	0–3	2
superlayer	0–2	2
TDC data	$0–2^{20}$	20

The unpacked data from the CSC byte stream are stored as several types of "digi" objects, which are convenient for passing into the reconstruction software. So far, we have used 2 types, a "strip digi" containing information from the CSC cathode data, and a "wire digi" containing information from the CSC anode data. There is one strip digi per strip channel and one wire digi per wire(-group) channel. It is likely we will extend the number to encompass up to 4 additional types of trigger-related CSC information: from the CSC comparator logic, from the Anode LCTs, from the Cathode LCTs, and from the combined Anode/Cathode LCTs. A strip digi contains a 7-bit identifier for that particular strip, and the 8 (or 16, depending on the configuration) time samples of the Switched Capacitor Array (SCA) pulse heights (each with 12-bit precision). The wire digi contains just a 7-bit identifier, denoting that this wire-group fired. In the current simulation. The strip digi also contains information from the CSC comparator logic; it is more realistic to split this off into a separate "comparator digi" and it is intended to do this.

The RPC part of the CDF raw data is unpacked and stored as RPC digis, where the detector identification, the fired strip number, and the bunch crossing number with respect to the L1A bunch crossing are packed. Only digis having a bunch crossing index equal to 0 are used for reconstruction, while the others are useful for detector and electronic debugging purposes.

3.2 Operations

This section gives a general overview of the procedures needed to bring the muon detectors into operating condition, specifically regarding detector calibration and alignment. In addition, the main features of the monitoring and database systems needed to control and run the online/offline reconstruction are reviewed.

3.2.1 Calibration

Two types of calibrations are usually needed for a detector. First, there are calibrations that provide the values of parameters needed to set up the detector for data taking: determining hardware thresholds, masks for noisy or dead channels, etc. These types of calibrations are available to all the muon detectors and were used prior to detector operation during test-beam activities, in bench tests, and for chamber commissioning. The constants they produce are either fixed in hardware or stored in the configuration database and are continuously monitored during data taking. Second, there are calibration data needed to determine the parameters for optimal offline reconstruction and HLT event processing. These data can be obtained by analysing special calibration runs or from special analyses of standard data. Only detectors involved in muon reconstruction (DTs and CSCs) use this kind of data. The RPC system, used purely for triggering, does not need such calibrations.

3.2.1.1 Barrel muon chambers

The barrel chambers measure drift times and, therefore, need special calibration runs to compute the relative t_0 of each wire due to the different signal path lengths to the readout electronics. The "calibration" command generates a signal on the input to the front-end chip for 4 wires at once, simulating a muon track crossing a Superlayer. These simultaneous signals generate a trigger in the DT system and cause the start of the local readout. A set of masks

allows the pulsing of only 4 wires at a time, and a special "sequence advance" command allows scanning of an entire chamber.

Calibration data will be taken before beam injection in the LHC. These data will provide a verification of the status of the entire electronics chain. The time needed for the execution of one calibration cycle is such that it can be issued during the abort gap as a private trigger. A complete check of the DT electronics can be done every few seconds by working on all chambers in parallel.

Once each relative t_0 is determined, the drift times are synchronized and a global t_0 is computed. Figure 3.4a shows the arrival time recorded by the TDC in all the cells of a chamber Superlayer, as obtained in test-beam data after relative t_0 equalization. A similar distribution can be obtained from a minimum-bias muon sample at the LHC. Figure 3.4b shows an enlargement of the short-time region; the global t_0 is obtained by fitting the derivative of this distribution (Fig. 3.4c). The maximum of the derivative distribution minus 2σ, where σ is the Gaussian width, gives the value of the global t_0.

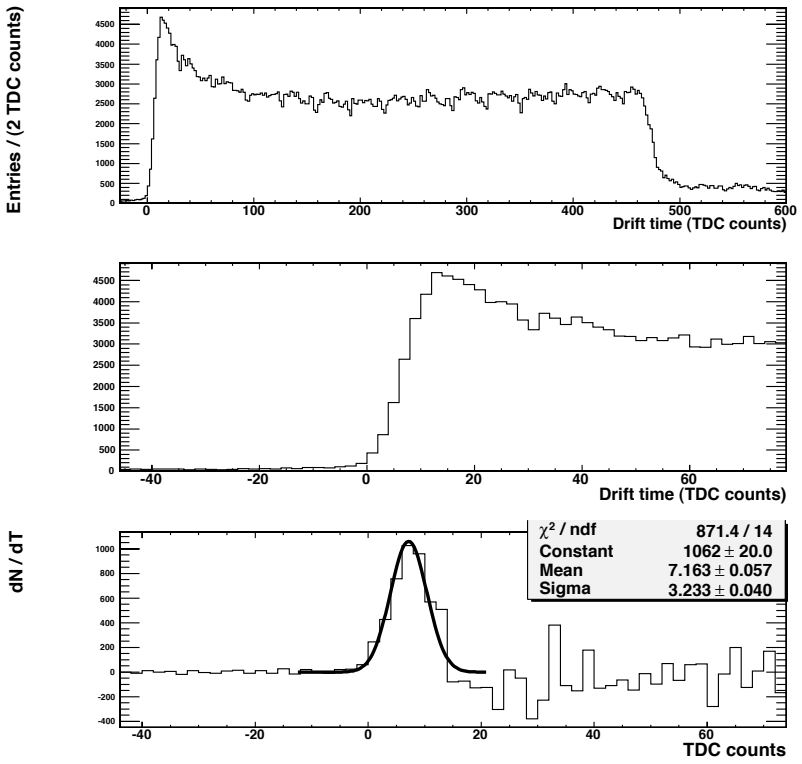


Figure 3.4: Determination of the global t_0 from the time derivative of the drift time distribution.

The same muon sample can be used to determine the drift velocity, which is evaluated from the mean-time (MT) distribution. The quantity $MT = (t_1 + 2t_2 + t_3)/2$ for any 3 consecutive layers is equal to the drift time, T_{max}, for one-half of a cell length, independent of the muon incidence angle. Hence, the drift velocity is the ratio of the average of the MT distribution (Fig. 3.5) and one-half the wire pitch. Only a few hundred events are needed to compute the average drift velocity in a chamber to better than 0.2%.

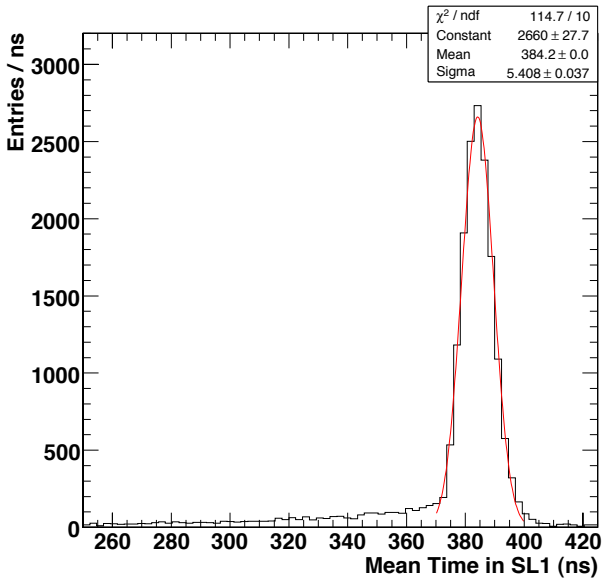


Figure 3.5: An example of the mean-time distribution for a muon chamber Superlayer.

3.2.1.2 Endcap muon chambers

The Endcap muon detectors are multiwire proportional cathode strip chambers (CSC), which have the high coordinate precision and fast response time needed for good muon tracking and effective bunch crossing number identification. The calibration measurements described here were performed on all CSCs during their testing at the Final Assembly and System Testing (FAST) sites. The results were stored in a database to serve as a reference for future calibrations and monitoring.

The signals from the cathode strips are read out by the cathode front-end boards (CFEB), providing a precise measurement of the azimuthal coordinate of the hit for offline analysis. For this measurement, we need to know the gain and cross-talk of the cathode ADC channels.

Additional calibration constants are measured that specify the amplifier linearity, as well as the individual cell offsets ($< 1\%$) of the switched capacitor array (SCA) [85]. Since the amplifier/shapers are capacitively coupled, the signal baseline values are rate dependent. Therefore, baseline subtraction is done using the first 2 time bins sampled before the pulse height rise of each cathode strip signal. Calibration constants are measured using the precision ($< 1\%$) test-pulse generator on the DAQ motherboard (DMB). A 12-bit DAC on the DMB sets the amplitude of the test pulse, while the arrival time of the pulse can be delayed in 6.25 ns steps.

The gain and cross-talk are measured by sending a test pulse of fixed amplitude ($Q_{in} = 100$ fC) through high precision capacitors, one strip at a time, for each CFEB and CSC layer (Fig. 3.6). A scan over the test-pulse delays provides an accurate measurement of the output pulse shape, yielding the gain and cross-talk. A scan over the test-pulse amplitudes allows a calibration of the slope and intercept of the preamplifier's output signal versus the DAC code of the input test pulse and quantifies the non-linearity of the preamplifier response. The SCA offsets are measured by randomly sampling the quiescent outputs of the amplifiers.

The calibration of the trigger path for the CFEBs consists of measurements of the comparators' thresholds, noise, and output signal timing. The DMB test pulse is used to find the

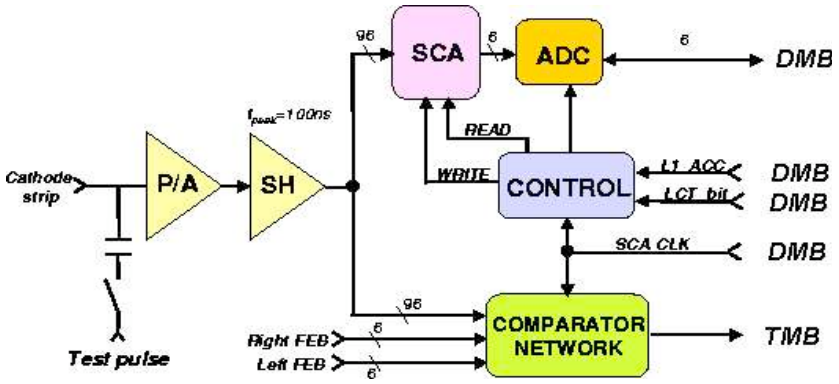


Figure 3.6: Block diagram of the CSC cathode front-end electronics board.

comparator thresholds at 2 test-pulse amplitudes of $Q_{in} \approx 15$ fC and ≈ 40 fC. The results yield the threshold slope, offset, and noise of each comparator. The comparator output timing is checked by means of a test pulse of fixed amplitude, which is scanned over a time delay of 15 steps of 6.25 ns each. The resulting comparator time offsets are calculated with respect to the average value per CFEB.

The operational scenario for the calibration of CFEBs assumes that it will be done with both spy data and electronic pulsing offline between beam fills. The cathode calibration constants have been found from previous tests to be quite stable, so the frequency of the calibration remains to be determined. With a data taking rate of about 27 Mb/pulse and a data flow rate of 600 Mb/s through 2 Gb fibres to each PC, the calibration is fast. The calibration analysis will be performed locally on a farm of 18 PCs. The files containing calibration constants for $\approx 200\,000$ cathode strips are small and in total do not exceed 50–60 MB per run.

The anode front-end electronics [86] (Fig. 3.7) supplies a precise timing measurement from a hit anode wire group. The thresholds and delays of the anode electronics are controlled by the anode local charged track (ALCT) boards. To calibrate the anode thresholds, an amplitude-controlled test pulse from the ALCT's test-pulse generator is used to fire simultaneously all 16 channels of each anode front-end board (AFEB) through their own internal test capacitors. A set of threshold measurements is done using 2 different test-pulse amplitude settings, corresponding to input charges of ≈ 30 fC and ≈ 50 fC. The threshold settings are varied around these 2 input charges to get the threshold turn-on curves. These data are then used to measure the corresponding threshold settings and the noise, as well as the threshold slope and offset.

The propagation time for signals from the AFEB to reach the ALCT has a spread among the front-end boards, and also has variations due to differences in the AFEB–ALCT cable lengths and the muon flight times. To equalize the arrival times of the anode raw hits at the ALCT within one CSC chamber, a set of control delay chips is used as input circuits on the ALCT. The delays in these chips can be set between 0 and 30 ns in 2 ns steps. At each delay setting, an asynchronous (with respect to the 40 MHz clock) ALCT test pulse is sent to all AFEBs, and the anode hits are read back by the ALCT. This allows a measurement of the delay slope and offset for each delay chip and the equalization of all the delay settings. This measurement is an important first step in the synchronization procedure for all the Endcap muon CSC chambers.

Data from CSC tests done at FAST sites, ISR, and SX5 show that the anode calibration con-

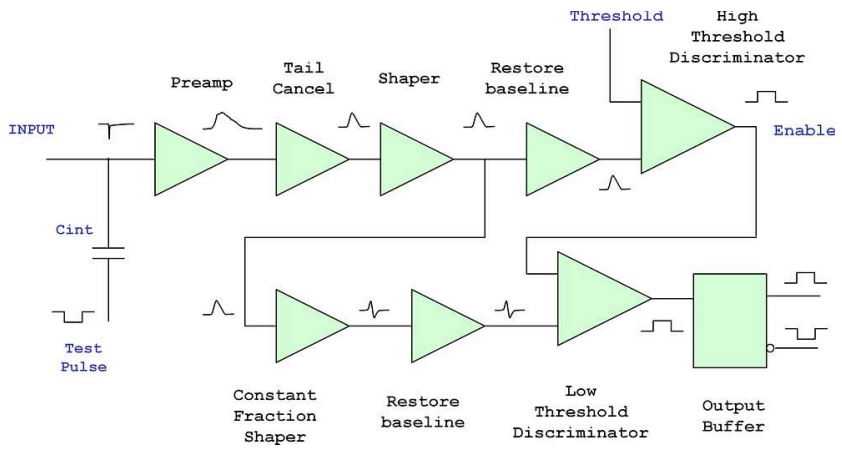


Figure 3.7: Circuit block diagram of one anode amplifier-shaper-discriminator channel.

stants are stable. It is reasonable to assume that the anode calibration will be performed at the same time as the cathode calibration but in separate runs. The volume of calibration data for the anode front-end is much less than for the cathode electronics. In addition, the anode calibration constants are defined per board, not per channel. An analysis of the anode calibration data will be done locally on the 18-PC farm. Calibration constants for $\approx 10\,000$ anode boards occupy a space of ≈ 1 MB per run.

Most of the CSCs have been tested at the default nominal high voltage of 3.6 kV. In one of the standard FAST site tests, the absolute gas gains were measured with cosmic rays using the positions of the Landau peak in the cathode strip signal-amplitude distributions. The voltage offsets (relative to 3.6 kV) for each CSC high voltage segment were found to equalize the gas gain in all layers of the CSC. A high-voltage scan was done in another test where ALCT and CLCT cosmic-ray rate plateaus were measured across the entire CSC. It is planned to calibrate the gas gains and voltage offsets periodically during normal CMS operation to monitor the stability of the CSCs.

At least one time at CMS start-up we will likely perform a high-voltage scan with cosmic rays and measure the actual nominal high-voltage settings individually for each high-voltage segment of each CSC with the working gas mixture and nominal thresholds on the anode discriminators and cathode comparators. Care should be taken in the analysis to ensure that selected tracks are going through the scanned segment and have inclinations similar to tracks that originate in the centre of the CMS detector.

3.2.2 Alignment

There are several potential sources of misalignment in the muon spectrometer, from chamber construction tolerances to final detector operation conditions, including

- Chamber construction tolerances: these are unavoidable geometrical tolerances in the production of the chamber parts (e.g., mis-positioning of wires within a layer, relative shifts in the Superlayer assembly).
- Detector assembly, detector opening and closing, and solenoid effects: assembly tolerances (per wheel and when closing the detector), and gravitational and magnetic field distortions on the return yoke that can lead to static deformations of the steel support. The overall effect can result in displacements of the chambers

with respect to their nominal positions of up to a few cm. The reproducibility after magnet power off/on cycles or detector opening/closure will be in general on the order of 1 mm.

- During operation: thermal instabilities can cause dynamic misalignments at the sub-millimetre level.

For optimal performance of the muon spectrometer over the entire momentum range up to $1\ \mathrm{TeV}/c$, the different muon chambers must be aligned with respect to each other and to the central tracking system to within 100–500 μm. To control misalignment during commissioning and to monitor further displacements during operation, CMS will combine measurements from an optical-mechanical system with the results of algorithms based on muon tracks crossing the spectrometer.

3.2.2.1 Optical alignment system

The Muon Alignment (MA) system is described in detail in [4]; we give here only a brief summary of its main characteristics. The basic geometrical segmentation consists of 3 r-z alignment planes with 60° staggering in ϕ. Within each plane the 3 tracking subdetectors of CMS (Central Tracker and Barrel and Endcap muon detectors) are linked together. Figure 3.8a shows a schematic longitudinal view of CMS, with the light paths indicated. This segmentation is based on the 12-fold geometry of the Barrel muon detector (Fig. 3.8b). Furthermore, the Barrel and Endcap monitoring systems can work in standalone mode, in which they provide reconstruction of the full geometry of each independent subdetector. The system uses 2 types of light sources (LEDs and laser beams), together with precise distance- and angle-measuring devices. The layout of the optical paths allows the monitoring of each of the 250 DT chambers (with 3D position measurements at the 4 chamber corners), while only $\approx 23\%$ of selected CSCs (12 chambers per Endcap layer) in the 4 Endcap stations are directly monitored. The links among the Tracker, Barrel, and Endcap are made at the external Barrel Wheels (YB$\pm$2).

The MA system is designed to provide continuous monitoring of the detector geometry, with or without collisions in the accelerator. The dynamic range of the system allows it to work at solenoid magnetic fields between 0 and 4 T. Its goal is to provide independent monitoring of the CMS tracking detector geometry with respect to an internal light-based reference system. This will help disentangle geometrical errors from sources of uncertainty present in a track-based alignment approach; for example, knowledge of the magnetic field, material-budget description, drift velocity, etc.

A real-size validation of one half of an alignment plane was performed in the I4 ISR hall [87]. It proved the validity of the system with precision close to the nominal one: $\approx 200\ \mu$m in $r\phi$ displacement, $400\ \mu$m in z displacement, and from 40 to 100 μrad in orientation.

3.2.2.2 Operation and validation

The relative positioning of the different internal mechanical components of a chamber (e.g., the wires and layers positioning inside a DT chamber SuperLayer) was measured during construction to be within the required tolerance ($100\ \mu$m). At the chamber production sites, systematic variations on the final assembled chambers were checked with cosmic muon data, showing good correlation between those measurements and the results of the muon track fit

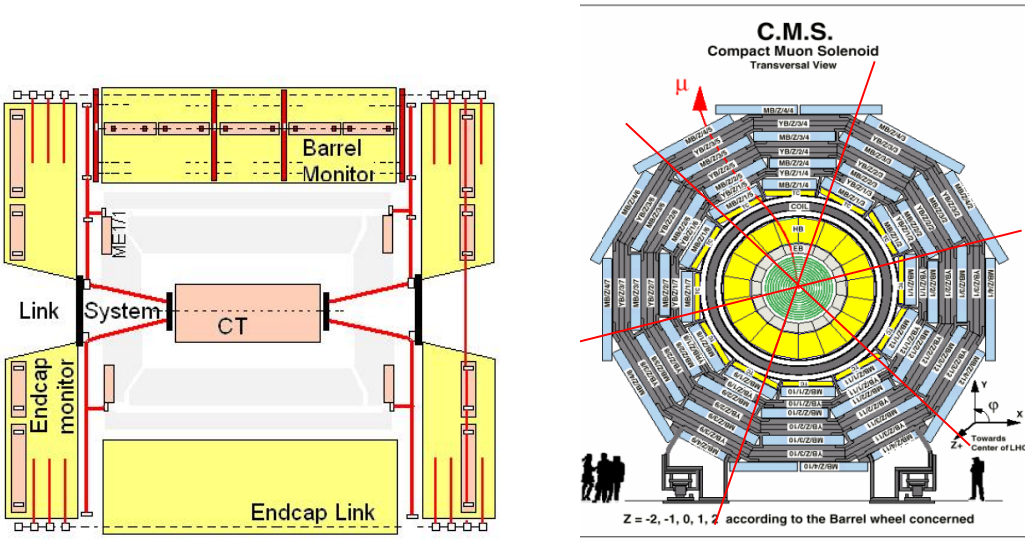


Figure 3.8: Schematic of the longitudinal view of CMS showing the optical alignment light paths (left), and transverse view of the Barrel Muon detector (right). The crossing lines indicate the r-z alignment planes with 60° staggering in ϕ.

in the chamber. For the DT chambers, the displacement of the SuperLayers with respect to their nominal position is less than a few hundred microns for most chambers [88] and can be corrected using the optical measurements performed on all the chambers at the ISR storage area before their insertion in the CMS steel yokes.

The internal chamber geometry is cross-checked at the CERN site using optical and survey measurements. These measurements are compared with the construction drawings and results from cosmic data to provide corrections to the design geometry when necessary. The individual chamber data are stored in a web-accessible DB [89]. An example is shown in Figure 3.9, with the distribution of the difference between ideal and real x coordinate and angular orientations (Fig. 3.10) for SuperLayer 3.

After the full detector assembly, photogrammetry measurements will be performed for each wheel, absorbing installation tolerances plus static (mostly gravitational) steel deformations, monitoring the different chamber positions and orientations, and providing an initial current geometry of the whole system. The measurements already made on the complete YB+2 wheel [90] show that, in general, the differences from the ideal positions are within 2–5 mm, the maximum deviation being about 8 mm. The sectors seem to move in one block so the displacement should be dominated by the wheel deformation. After detector closure and solenoid power-on, this geometry is expected to vary significantly in several coordinates and hence has to be monitored and corrected with the MA system and with cosmic-ray tracks. To validate the system, the 3 sets of measurements—survey, MA, and Tracks—will be compared before the solenoid is powered on. MA and tracks will also be compared at the working conditions of 4 T.

The entire scheme will be validated during the solenoid magnet test, where about 1/3 of the alignment system will be implemented. Barrel sectors 10 and 11 will be instrumented with DT chambers and all the corresponding alignment components; likewise one complete

Endcap (ME+) will be instrumented, which will allow cross-checks of the MA system against cosmic tracks.

The MA system is designed to take data continuously, completing a full cycle of measurements several times per hour, monitoring any movements on that same time scale, and with high redundancy for slower structural distortions.

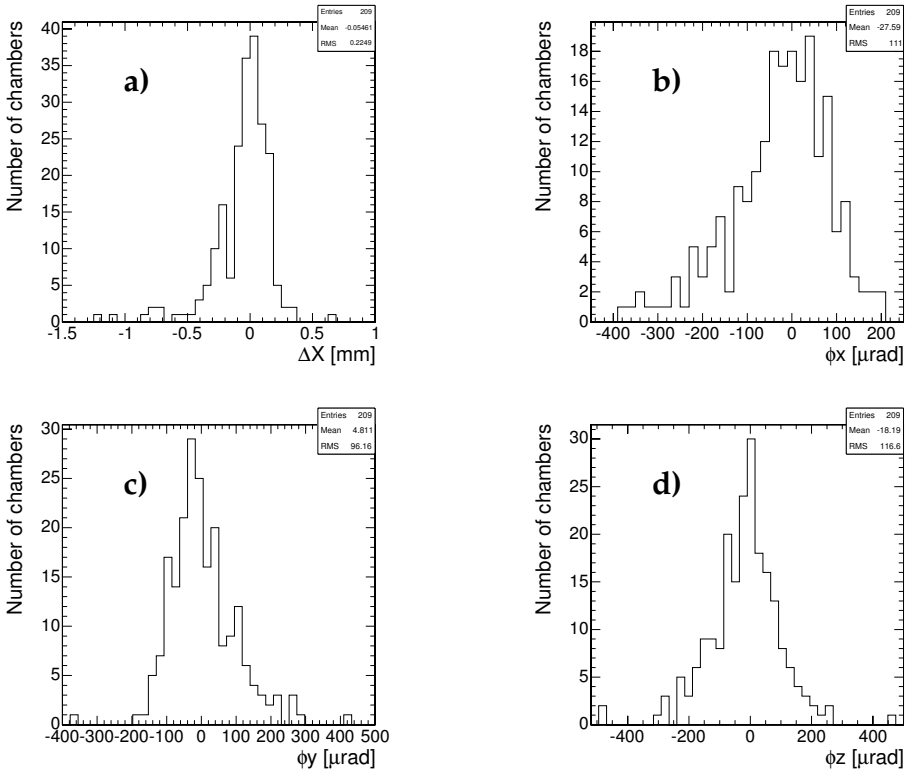


Figure 3.9: Survey results of the DT Geometry for SuperLayer 3: difference between ideal and real positions in a) the x coordinate and b) the orientation around the x, c) y, and d) z axis.

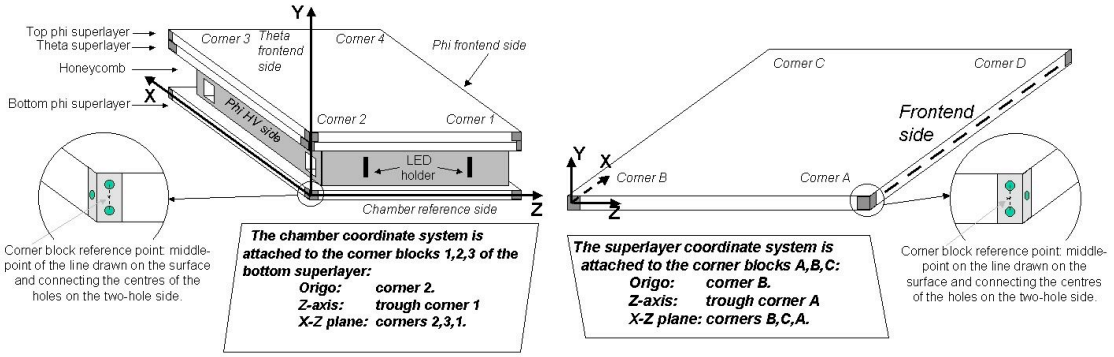


Figure 3.10: DT chamber scheme and coordinate system for survey.

3.2.2.3 Reconstruction software

The data taken by the MA system will be written into an Oracle database using the CMS Detector Control System (DCS, see [8]) system. The transferred data will amount to at most a few MB per second. The alignment reconstruction software, COCOA (CMS Object-oriented Code for Optical Alignment) [91], will read each set of optical alignment data and reconstruct the positions and rotations of the muon chambers. To do this, COCOA needs the positions of the different hardware elements of the MA system, as well as their optical properties and internal calibrations. The nominal positions of the chambers and of a few large elements in the system will be read from the same CMS Geometry Database that is used by OSCAR and ORCA. The internal alignment sensor calibration data will be read from the Alignment Database, which is part of the CMS Construction Database. With the description of the system, COCOA will read a set of MA system data and then make a non-linear least-squares fit to determine the positions and angles of the muon chambers (together with other unknown parameters of the system that are used as intermediate steps but that are not needed by the ORCA software). The uncertainties in the calibrations will also be propagated to obtain the errors in the chamber positions, as well as the correlation matrix among the different parameters. All these results will be written to the Alignment Database. The chamber positions and orientations will be written to the CMS Geometry Database so that they can be used by the ORCA software.

3.2.2.4 Track-based alignment

To complement the MA system, tracks from physics-run collisions and from cosmic-ray data will be used to measure the chamber positions and orientations to a precision similar to the MA in a totally independent measurement (with independent systematic uncertainties), which will contribute to the robustness of the system. This process will be performed offline in a continuous mode, updating the geometry whenever movements significantly larger than the achieved precisions are observed.

By construction, the various chambers can be considered as rigid bodies within the desired precision. Any mis-positioning in the construction will be already accounted for by the survey, optical, cosmic-ray, and beam-test measurements. Hence, the remaining task during operation is limited to obtaining the 6 degrees of freedom for each of the 790 chambers. Several simplifications can be performed in this procedure, but even in the more complex case, the problem consists of calculating about 5000 parameters and all their correlations. Despite being large, this is a manageable number given current mathematical and computational tools.

The alignment of the muon system has some unique features in comparison with the alignment of other sub-detectors, for example, that of the CMS tracker. The muon alignment involves a relatively small number of parameters, but with the different tracking systems (Tracker-muon or Barrel muon-Endcap muon) being separated by a large amount of material and with the tracks, which cross the 2 systems, travelling through a nonuniform magnetic field. Another difference is that once the chambers are assumed as rigid, they provide not only a spatial point but also a vector direction for the track, which can be used in the alignment. Two alternative approaches are foreseen [92] for the muon alignment: an alignment with respect to the central tracker and a standalone muon-system alignment.

The baseline alignment algorithm takes advantage of the redundancy in CMS for reconstruct-

ing muons. It uses isolated muons, which are detected both in the central tracker and in the muon system. Once such a muon is identified, the extrapolation of the tracker-only track to the muon system is compared with the local reconstruction in each muon station. This allows one to decouple the alignment of each chamber, factorizing the 5000-parameter fit into 790 simpler 6-parameter fits. This simplifies the calculations, avoids numerical problems, and allows a more direct estimation of the systematic errors. The method only needs to store the differences between the reconstructed and extrapolated 4-D coordinates, together with their errors, which can be done during the first-pass reconstruction. After rejecting events in the tails of the distributions, the 6 degrees of freedom for each chamber can be obtained from a χ^2 fit to the data. Examples of such distributions for a $W^{\pm}$ sample equivalent to 50 h at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ are shown in Fig. 3.11. They show the estimator for the displacement in the $r\phi$ coordinate for a single MB1 chamber for an ideally aligned detector and for a hypothetical case when the chamber is rotated with respect to the beam axis by 0.25 mrad.

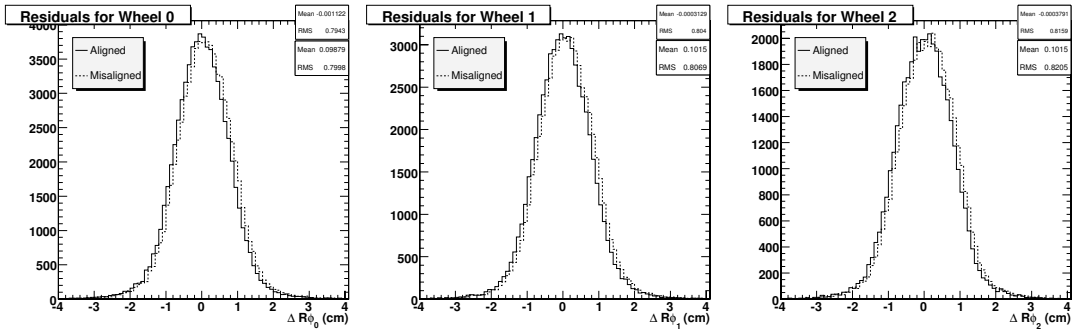


Figure 3.11: $r\phi$ estimator in a single MB1 chamber in different wheels for aligned (solid line) and 0.25 mrad rotated (dashed line) samples of $W^{\pm}$ equivalent to 50 h of data taking at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Although in principle a single pass should be enough to obtain a precise solution, this is only strictly true if the errors are purely Gaussian and the displacements and rotations are small. At least one additional pass will be needed to reach the desired precision whenever the detector has moved significantly. Obviously, this procedure relies on a good tracker alignment, but it has been shown that the results are not significantly degraded when using the tracker alignment errors expected at the beginning of data taking (the so-called "First Data Taking Scenario" that will be discussed in Section 6.6.4).

Alternatively, the muon system can be aligned without relying on the central tracker using tracks reconstructed with the muon system. This procedure will be used if the tracker misalignment is too large to efficiently reconstruct full tracks. The technique will also provide initial information on the muon system's internal alignment, without having to wait for the final tracker alignment. The relative position of the entire muon system with respect to the tracker can then be calculated at a later stage. In this case, the Blobel approach [93], as used for the tracker (Section 6.6.5), will be followed.

Besides solving the full problem of 5000 parameters, which provides the best estimation of the parameters, their errors, and correlations, the problem can also be factorized with good approximation into several almost independent fits with 20–50 parameters each. These fits will be used to study the consistency of the results, to obtain an initial estimate of the systematic errors, and for a more direct comparison with the optical alignment system. One

such example is shown in Fig. 3.12 for the same $W^{\pm}$ sample, corresponding to the alignment of a Barrel Muon sector when all the chambers are displaced alternatively by ± 1 mm. The variables shown correspond to the alignment estimators in each chamber derived from the simultaneous minimization of track and alignment parameters following the Blobel formulation. These optimal estimators are obtained as linear combinations of the track residuals, defined as the difference between the hit and fitted (with ideal alignment) track r-ϕ spatial coordinates. The coefficients of this combination are defined by the method as a function of quantities such as the geometry, the field integrals, and the measurement errors. In general the problem is nonlinear and is solved by iteratively linearizing it. In this particular example, we show the result after the first iteration. This formulation permits the simultaneous acquisition of the alignment parameters and the corrections to the track's fit parameters, taking as input the difference between measured and fitted (with ideal alignment) track parameters. All of these procedures will be tested in detail during the Cosmic Challenge.

For both methods of alignment, the basic data sample that will be used during physics runs is from inclusive high-p_{T} single muons. This will be complemented by dimuon samples, mainly from Z decays, obtaining additional information from the known invariant mass. Cosmic-ray samples will also be used, particularly during commissioning and before data taking has begun. The precise muon lower-momentum cut-off used for the alignment depends on a balance between reduction of systematic errors, computing resources, and the frequency of updating, but under reasonable assumptions it can be set to $40\ \mathrm{GeV}/c$ when running at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. This cut-off can be lowered or even removed at start-up, when significantly lower luminosity is expected but the inclusion of lower momentum tracks will not significantly improve the results. The large rates expected during physics runs with this threshold will allow a statistical precision of about $100\ \mu\mathrm{m}$ within a few hours of running at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Even at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, the rates are enough to reach this precision within a few days of running time. However, different systematic errors are expected to affect this measurement. They are, in general, reduced with an increase in the momentum cut-off, and hence can be overcome by accumulating more data. The following systematic errors are for a minimum muon momentum of $40\ \mathrm{GeV}/c$ and correspond to initial estimations based on simulations:

- bias in the alignment estimators due to tails and non-Gaussian behaviour (40 μm)
- uncertainty in the magnetic field outside the tracker volume (200 μm)
- uncertainty in the material budget estimation (100 μm)

The Cosmic Challenge is expected to provide much more precise information on the achievable precision.

3.2.2.5 Using alignment information in reconstruction and HLT

The outcome of the different alignment measurements will be combined and the current displacements and rotations with respect to the chamber nominal positions will be stored in the CMS Geometry Database, if significantly different from 0.

These data can be read by the ORCA reconstruction program, which can change the position of the reconstructed hits when their coordinates are translated from the local reference system of the chamber to the global CMS system. Consequently, within the global reconstruction, hits are placed at their true position with respect to the CMS global reference system.

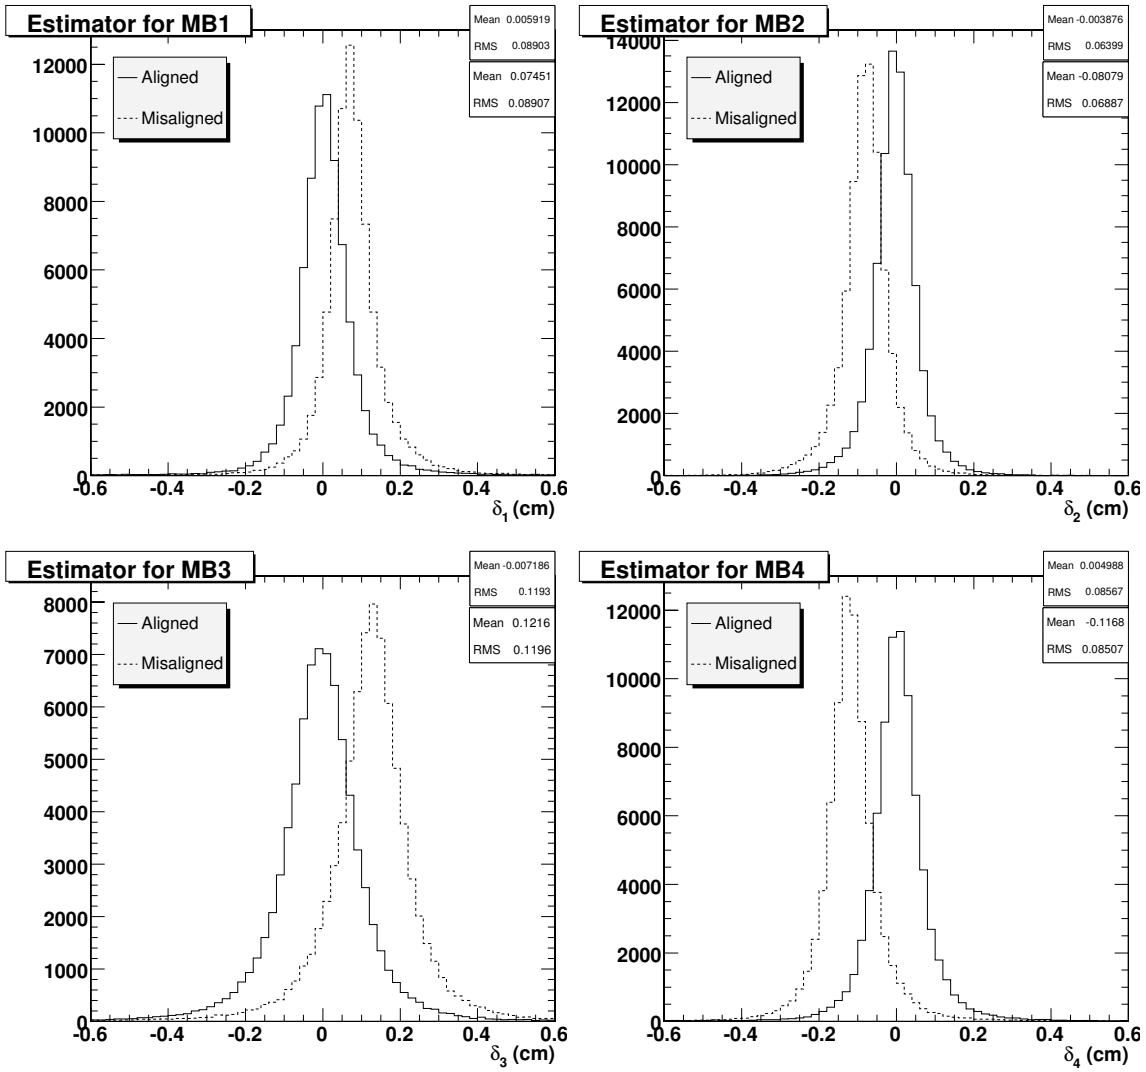


Figure 3.12: $r\phi$ muon standalone estimator for chambers in different wheels for aligned (solid line) and ± 1 mm displaced (dashed line) samples of $W^{\pm}$ equivalent to 50 h of data taking at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

This code can also be used to simulate the effect of misalignments in the physics reconstruction. To do this, hits with alignment constants different from 0 are reconstructed using samples simulated with an ideal geometry. In this way, misalignment is introduced at the reconstruction level, saving the large amount of CPU time that would be needed to re-simulate the samples. Similar to what is described for the tracker alignment simulation, 2 default *misalignment scenarios* [94] have been implemented. These scenarios approximately simulate the situation expected for the first CMS data taking (First Data Taking Scenario) and when the nominal precision is asymptotically reached (Long Term Scenario). Each chamber is moved and rotated according to the expected precisions in each case (summarized in Table 3.2). Additionally, a correlated error for the entire muon system, which simulates the tracker-to-muon relative positioning and orientation, is put into the simulation, using the values also shown in the table.

Table 3.2: Alignment precisions used in the 2 misalignment scenarios. Displacements are in mm and rotations are in mrad.

	MB	CSC	Tracker-Muon
First Data Taking Scenario			
Displacement (x-y)	1.0	1.0	1.0
Displacement (z)	1.0	2.0	1.0
Rotations	0.25	0.5	0.2
Long Term Scenario			
Displacement (x-y)	0.2	0.2	0.2
Displacement (z)	0.2	0.4	0.2
Rotations	0.05	0.1	0.04

The effect of the misalignments [95] considered in the 2 scenarios is found to be negligible as far as the muon reconstruction efficiency is concerned, while it might have an important impact on the momentum resolution (Fig. 3.13).

3.2.3 Monitoring

Online monitoring of the muon system performance requires a coherent suite of tools to quickly detect hardware malfunctions and to provide clear feedback in the form of alarms or warnings. With approximately one million readout channels of very diverse natures, this is a complex task. There are 2 distinct ways of monitoring the system operation.

First, the operational status of various hardware elements can be monitored by accessing dedicated on-board sensors/registers and using communication channels that bypass the physics data stream, i.e., via the Detector Control System (DCS, see [8]) and the Fast Merging Module (FMM, see [96]) paths. Monitoring through DCS deals with high voltage for muon chambers, low voltage for electronic boards, flow/pressure of chamber gas and cooling water for electronic boards, environmental and radiation sensors, and alignment. FMM is designed to report critical errors in electronic board operation, such as out-of-synch states, timeouts, format corruptions, and FIFO overflows. Detection of these errors can be programmed into on-board FPGAs. DCS and FMM are described elsewhere.

Second, detector operation performance can be evaluated by spying on the actual physics data flowing through either the DAQ or Trigger paths, and performing data analyses of varying complexity. Such online data analyses, referred to as Data Quality Monitoring (DQM), can be subdivided into 4 sub-groups:

1. DAQ-path raw-data analyses: e.g., general data-format integrity; inefficient, dead, and noisy electronics elements; cross-talk and after-pulsing; delays, gains, and pedestals; synchronization checks (match of L1A numbers, BX numbers); and current FIFO status.

2. Level-1 Trigger path raw-data analyses: e.g., trigger rates (to be corrected for luminosity and presented in terms of cross sections) and various distributions related to trigger objects. These analyses are best for identifying Level-1-path electronics problems.

3. Level-1-Trigger simulations based on the DAQ raw data to check that the results match the Level-1-Trigger data. Note that DAQ data contain trigger primitives sent to the Level-1 Trigger. These cross-correlation analyses, based on both DAQ and Trigger data,

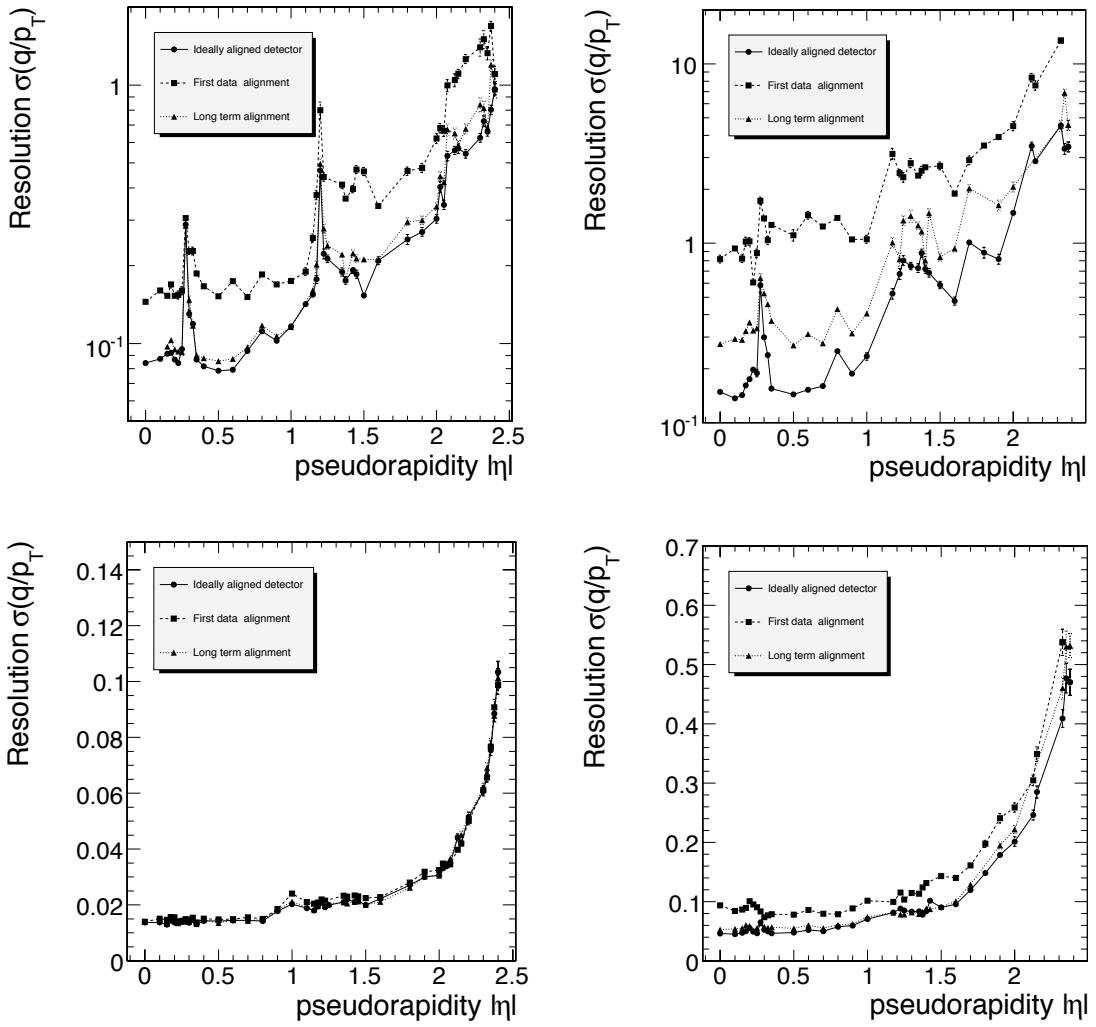


Figure 3.13: Effect of the different alignment scenarios on the muon momentum resolution as a function of η for $p_{\mathrm{T}} = 100\ \mathrm{GeV}/c$ (left) and $p_{\mathrm{T}} = 1000\ \mathrm{GeV}/c$ (right), for standalone muon reconstruction (top) and global muon reconstruction (bottom).

will allow one to sort out whether any problems seen by the Level-1-Trigger DQM are related to the trigger hardware (in the case of a mismatch) or to the front-end electronics (in the case of matching, e.g., a dead trigger sector is aligned with chambers that never report data).

4. Physics-object-oriented analyses: e.g., multiplicity of reconstructed muons, muon-track quality, muon momentum and angular distributions, and opposite-sign and same-sign di-muon mass distributions. These analyses are based on fully reconstructed physics objects and are likely to involve more than one CMS sub-system. Results of such tests are hard to relate to particular hardware element problems, but, nevertheless, they can be viewed as the ultimate judge of whether the overall system performance satisfies the requirements of physics analyses (e.g., a single dead channel, as seen at the lower level hardware-oriented DQM analyses, does not necessarily disturb the ability of the system to reconstruct physics objects).

The first 2 categories of DQM tests can be performed both on the CMS Event Filter Farm and by spying on sub-system data via the corresponding sub-system Local DAQs (CSC, DT, and RPC). The last 2 categories of analyses, requiring fully-built events, can be done only at the level of the CMS Filter Farm.

The scope and diversity of the front-end electronics of the muon system naturally imply that the quantitative validation of data that adequately covers the full range of possible failure modes will require a large number of tests. The total number of histograms to be filled and analyzed in the muon system is estimated to be $200\,000$ (about 10^7 bins). Also, because the CMS muon system has a large number of channels that are read out only when trigger primitives are present, the average occupancy per channel per recorded event is very low. Therefore, to produce a meaningful statistical analysis per channel in a reasonable time, DQM must be able to access and digest data at a substantial rate, currently estimated to be at least 1 kHz per sub-system. This requires intercepting data right after the Level-1 Trigger. Events rejected by HLT may be of special value for DQM as well.

The first muon DQM prototypes were exercised in the beam tests of 2004. The current ongoing tests at SX5 (slice test) are also indispensable for developing DQM tools and getting early experience with possible problems. Many test analyses and corresponding histograms have been implemented and used to evaluate DQM performance. For example, the EMU beam tests in 2004 had up to 4-5 chambers in the readout, which is just the right scope—it is estimated that a typical CMS event will have 2-3 CSCs in the readout. Therefore, we were able to time the data unpacking and DQM analyses, which gave us a benchmark number of about 100 events per second per CPU node. This defines the necessary scope of a farm with about 10 CPUs per sub-system to perform DQM analyses at the desired rate of 1 kHz per subsystem. Work on developing tools for automatic detection and user-friendly representation of problems among $200\,000$ histograms is in progress.

3.2.4 Databases

There are several kinds of muon-system databases created in the CMS database framework: Detector Geometry Database (DGD), Construction Database, Equipment Management Database (EMDB), Configuration Database, and Conditions Database [97].

Calibration data, needed for the reconstruction software, are included in the online Conditions Database. They are transferred to the offline database by standard methods (see Section 2.8) and are accessible to the offline reconstruction as calibration objects.

Oracle database technology is mainly used for muon-system databases, while in addition, MySQL is used for RPC Construction databases.

3.2.4.1 DT databases

Construction databases exist at each DT production site and are accessible locally. The Central Production database can be seen worldwide at the web address [98]. It gives information on the DT chambers, superlayers, layers, and cells (parameters such as wire position, tension, noise, efficiency, etc.), and data about production. Information on racks, crates, and cables is entered into the EMDB by the integration group at the time of installation.

The Configuration Database contains configuration data such as lookup tables, thresholds, and masks. Lookup tables (LUTs) are needed by the trigger system to compute track pa-

rameters (radial angle and bending angle). The delay between the bunch-crossing time and the signal collection, due to time-of-flight and propagation along the anode wires, is needed for a correct bunch-crossing assignment when triggering. Front-end boards are to be programmed with the signal thresholds and channel masks. Most of these quantities are not expected to vary significantly during detector operation, so the total amount of data for all configuration parameters is foreseen to be quite small (≈ 35 MB).

The DT Conditions/Calibration DB will contain data necessary to perform reconstruction in offline and HLT analyses, and will monitor the data collected during detector operation.

The hit position of a track inside a DT chamber is reconstructed using the electron drift time from the track's crossing point in the cell to the anode wire. So 2 parameters are needed in the operation:

- an offset time to be subtracted from the TDC response to get the actual drift time, with a single-cell granularity;
- the electron drift velocity to compute the distance once the time is given, with a "superlayer" granularity.

These parameters are determined by the methods explained in previous sections, and are stored in DB tables. The data size is estimated to be about 750 kbytes for the time offsets and 8 kbytes for the drift velocities. In addition to the above "calibration" quantities, information about the mapping from readout channel to physical drift tube is to be stored and retrieved.

Much other data, not used in reconstruction, will also be read from the DT system and written into the database. Data such as temperatures, voltages, and currents are written through the DCS system (PVSS). Information on the electronics status, including single-card power, thresholds, clock delays, and error flags will be read from each minicrate every minute. They will be written into the database with a frequency depending on the data themselves, e.g., when errors occur or when values are out of range.

3.2.4.2 CSC databases

Access to the CSC databases is from the web page [99]. The Construction Database contains information on the production and testing of the Endcap Muon system components: tracking of electronics boards and chambers, tracking of HV components, information on the production of the ALCT and Mezzanine cards, and tests of the AFEB boards. Information on cables, racks, crates, and installed CSC chambers is included into the EMDB. An interface to retrieve information on cables and installed CSC chambers is provided.

The Configuration Database contains settings for all the front-end and peripheral crate electronics boards, FED boards, TTC boards, Track-Finder boards, the XDAQ [38, 39, 40] environment, and the local DAQ parameters. All of these parameters are needed to start a run. The total amount of data is about a few hundred kilobytes. This information must be put into a tree of relational tables and read by a standard XDAQ application. The table structure should also maintain the versioning of the configuration settings and their tagging.

The first prototype of the CSC Conditions Database is the Test-Beam Database [100]. It saves all the parameters characterizing a test-beam run and the configuration parameters for the run. In some cases, the entire XML files are uploaded, which initially contained the configuration information. This DB is also used as a logging system.

The CSC Conditions Database contains both information on the electronics performance and the calibration values needed for reconstruction. In principal, the Endcap Muon system is designed in such a way that it can successfully work without calibration. The baseline can be determined from the first time slices of any signal. Channel gains are stable to within 3%, even between different chips (and to 1.2% inside a chip). Cross-talk corrections can be applied universally for all channels. However, for more precise measurements, the following calibration constants will be measured and saved in the database: gains (400 kB), cross-talk (1.6 MB), and probably linearity (2 MB). For refined measurements further parameters can be saved: pedestals and their RMS values (1.6 MB), and predefined parameters of the signal shape in time and in space across the strips (done per layer, 70 kB).

There are about 30 tests for checking the electronics performance, corresponding to about 24k numbers per chamber (50 MB in total). Measurements of the anode thresholds (350 kB) will go to the Configuration Database for downloading onto the AFEB boards. In addition, there is the measurement of all the pedestals of the SCA (Switch Capacitor Array, 96 capacitors per cathode channel, 20 M numbers, 40 MB), which will be performed in the initial period of detector operation and then only occasionally.

All values from the calibration measurements and performance tests are expected to be stable and measured rarely. The rate of obtaining the conditions data will be found from experience.

3.2.4.3 RPC databases

Since the RPC Barrel and Endcap production lines have proceeded in parallel and independently, 2 Construction Databases were installed where all the information related to chamber structure, and quality control and assurance reside. The Barrel RPC Construction Database has existed since March 1999, and is accessible world-wide from the web page [101]. The database contains information related to the different steps of production and Quality Control. It will also contain information related to the commissioning procedures, as they develop. It can create interactive online plots that allow the time dependence of parameters and their comparison in different periods of time to be shown. The Endcap RPC Construction Database was developed [102] later. It has a similar capability to create various plots interactively, reflecting the content of the database.

Information on racks, crates, cables, etc. is being included in the EMDB. Information on chamber installation will be included during the massive installation. Other important information (e.g., the electronics chain necessary to generate the first-level trigger signal for muons, a comprehensive description of the relative hardware (Link, Distribution, and Control Boards), and relations among the different Components) will also be stored.

The Configuration Database (as well as the Conditions Database and EMDB) is unique for the Barrel, Endcap, and RPC Trigger groups. It includes all the information related to the front-end electronics, in particular the location of the FEBs on the chambers (and the location of the chambers themselves), the thresholds, the relative I2C addresses needed by XDAQ to perform this operation, and the parameters needed to configure the Link Board chain. Integration between the Configuration DB and the High Voltage and Low Voltage systems, which will be handled by PVSS/JCOP, remains to be done. The Configuration Database will have 3 different operational modes: 1) Off Mode (during shutdown), 2) Standby Mode (during accelerator filling), and 3) Running Mode (during data acquisition).

A first set of parameters to be stored in the Conditions Database will be the configuration settings, for instance, the threshold of the FEBs, and the voltage set on the chambers. In addition, all parameters characterizing the chamber behaviour will have to be stored. Parameters coming from the HV, LV, and gas systems will be handled through the PVSS system. In particular, the currents and voltages on all the chambers of the system will be read out every minute and stored, together with the environmental conditions, to give a first glimpse of the chamber performance. Finally, the gas system and its quality will be constantly monitored, and the relevant information stored. The Conditions DB will therefore be constantly uploaded with different classes of data related to the different parts of the system being monitored.

3.3 Detector response

3.3.1 Simulation

The physics response of all CMS subdetectors is simulated with GEANT4, which takes into account the geometry and material of the detector as a whole and simulates the physics processes that occur as particles travel through each subdetector. As discussed in detail in Section 2.5, different models of the particle interactions in matter ("physics lists" in GEANT4) have been considered in the simulation for the modelling of the hadronic processes. The output of this stage of the simulation is GEANT "hits" in "sensitive detector material." For the case of tracking detectors, like those in the Muon subsystems, each such "SimHit" contains the entry and exit points of a simulated particle as it traverses a detector gas layer. GEANT is not used to model the interaction of this primary particle in the gas, but just to estimate its energy loss as it passes through.

The interaction in the gas itself is considered as part of the next step of the simulation, which we refer to as detector "digitization." This involves simulation of the electronic signal processing and signal readout performed by hardware either mounted locally (on the detector) or remotely, and ideally leads to a representation of what real data will look like as they stream into the DAQ system. The output of this procedure is a set of "digis" that approximate the raw data expected in reality. In practice, these digis for the muon subsystems contain values representing the essential measurements of the respective detectors rather than a bit stream approximation of the actual raw data, although it is intended later to extend the modelling to this level.

Digitization is subdetector specific since it depends on the particular detector technology, and thus is completely different for the DT, CSC, and RPC components of the CMS Muon detector system. We summarize below the simulation of the digitization in each muon subsystem.

3.3.1.1 Digitization of the DTs

A DT digi consists of a TDC time that includes several contributions: the drift time of the electrons in the gas, the propagation time of the signal along the cell wire, the time-of-flight of the muon from the interaction point to the wire, and a t_0-like offset.

The current digitization in ORCA is based on a parameterization [103] of the cell response obtained with a GARFIELD [104] simulation; the drift time is parameterized as a function of

the distance of the SimHit from the wire, the incidence angle of the muon, and the magnetic field components parallel and orthogonal to the wire. This time is smeared according to a double-half-Gaussian distribution, where the 2 sigmas are also parameterized [103].

For the case of several SimHits in the same cell, the drift times are computed independently, but separate digis are created only if the difference in drift times is larger than a configurable dead time (the default value is 150 ns). Hits from energetic delta rays and any other secondaries (e.g., e^+e^- from pair production) crossing the whole cell are handled the same way as muons. Soft electrons stopping in the gas or entering and exiting on the same face of the cell due to bending in the magnetic field cannot be handled by the parameterization and are currently ignored. Delta rays produced in the gas are also ignored, since they are already generated by GARFIELD and their effect is included (statistically) in the parameterization. Preliminary simulation studies using CMSIM/GEANT3 [29, 28] show that these cases account for less than 4% of the SimHits.

The signal propagation time from the hit to the front-end electronics is computed assuming a propagation velocity of 0.244 m/ns [80]. The particle's time-of-flight from the collision vertex is computed by GEANT and includes a proper offset for out-of-time pile-up. The t_0-like offset represents all the delays due to cables, phases, and readout; in the simulation there is no need for a different offset for each channel, therefore the default value is 500 ns for all the cells. Moreover, since in the current implementation the Level-1 trigger emulator requires digis synchronized by chamber, the time-of-flight of an infinite-momentum particle travelling from the nominal interaction point to the centre of the chamber is subtracted.

Currently, the effects of cell inefficiency are not included in the parameterization of the cell response. These are particularly important in the I-beam and wire regions. Furthermore, in the current implementation there is no simulation of the after-pulses which are due to photo-electric extraction of electrons from the cathodes and strips in the DT cell.

3.3.1.2 Digitization of the CSCs

The description given here summarizes and updates the content of [105]. In the current digi implementation for the CSCs, we build independent strip and wire digis, representing the data measured on the cathode strips and the anode wires (which are grouped into bundles). The strips give a precise measurement (a few hundred μm) of the position of passage of a charged particle in the global ϕ coordinate by sampling the induced charge distribution which spreads over several strips. The wire groups give a coarse radial measurement (several cm) but a precise time at which the electron charge arrives. The digis are obtained through a detailed simulation of the CSC gas ionization and collection from each SimHit, and the electronic processing of the resulting, possibly overlapping, signals. Each CSC contains 6 independent gas layers, each with a layer of strips and a layer of wires on which signals are collected. The strip signals are measured by a switched capacitor array (SCA) [85], which works as a continuously refilled buffer of the pulse heights on the strips, sampled every 50 ns.

The ionization of the CSC gas is modelled using an adaptation of GEANT's thin-layer approximation for modelling energy loss. Production of secondary delta electrons is included. The electrons and ions produced are transported to the anode wires and cathode strips using parameterizations of more complete drift simulations from GEANT. Near the wire, the gas gain and avalanche fluctuations are modelled, as well as the electron attraction in the gas

and the anode collection efficiency. The result is a charge collected on an anode wire. The induced charge on the cathode planes is distributed according to the Gatti parameterization for signal shapes on the strips in the CSCs [106].

The electronics used to detect the signals on the CSC strips and wires are modelled using parameterized distributions of the measured amplifier and pulse-shaper responses. The resulting signals are overlapped where necessary, and wires are ganged as in the hardware. Signal propagation times and delays are modelled.

The wire information is written into the "wire digi." This information includes the wire-group label, the time at which the signal on the wire reached 50% of its maximum value (modelling a constant-fraction discriminator), and the bunch-crossing tag (defined as the constant-fraction-discriminator time with the expected time-of-flight of an ultra-relativistic particle travelling from the IP to the centre of the chamber subtracted, normalized by the bunch spacing, typically 25 ns). Noise is simulated by jittering the threshold of the constant-fraction discriminator.

The strip signals are given random baseline shifts dependent on the chamber length. Noise in the SCA is modelled by superimposing noise samples from test-beam pedestal runs, in order to obtain realistic time correlations. We also apply a 3% Gaussian smearing to account for variation in amplifier gain, and a 3-ns Gaussian smearing to account for variation in the shaper peaking time.

Strip-to-strip cross-talk is modelled with a capacitive and a resistive component. The capacitive component is proportional to the slope of the signal, and the resistive component is proportional to the signal itself. The constants of proportionality are chosen to obtain cross-talk signals similar to those seen in test beams. The capacitive coefficient varies from chamber to chamber in proportion to strip length. The resistive component is set at 2% of the original signal. The cross-talk is applied in order of decreasing signal size: charge is added to neighbouring strips and subtracted from the original signal. Wire-to-wire and wire-to-strip cross-talk are expected to be smaller and are as yet unmodelled.

The CSC information entering the Level-1 Trigger does not require the full precision of the SCA readout and needs only to localize muon hits within about half a strip width per chamber plane. This is achieved using hardware comparator logic, which is also modelled in the software. There is one comparator circuit for each pair of strips, so only the comparator output from one strip of each pair, the even-numbered strip, is required to specify the overall output. For each comparator, every 25 ns we check whether a strip signal is above the comparator threshold (8 times the RMS pedestal noise), and whether that signal is larger than either of its neighbours 50 ns after it exceeds the threshold. If so, the comparator returns a code (one of 4 integer values) corresponding to the "half-strip" in which the hit is expected to be, according to whether the left or the right neighbouring strip pulse height was larger. The comparator is then suppressed from further readout until the strip which triggered it drops below the comparator threshold.

In the real DAQ, the CSC readout is initiated by the recognition of a Local Charged Track (LCT) in a chamber, in coincidence with a Level-1 Trigger Accept. This triggers the readout from the SCA of a group of 16 strips in all 6 layers of the chamber, which gives a 400 ns wide "snapshot" of the pulse shape in those strips, in 8 time samples, each separated by 50 ns. In the simulation, we must supply CSC digi information to the trigger simulation package to allow verification of the trigger algorithms. So instead, we start the CSC readout when a

comparator fires, since this is in any case the primary input to the Level-1 CSC trigger. Thus, when a comparator fires, the readout is triggered for a group of 16 strips in that layer. Of these 16 strips, those strips with actual charge deposition are read out. Noise is simulated on those empty strips that are neighbours of a strip with a signal and, in order to reduce the volume of data, the other empty strips are suppressed. If the strip that fires the comparator is within 2 strips of a boundary of these 16-strip groups, the neighbouring 16 strips are also read out. Once a group of strips has been read out, these strips cannot be read out again for 200 ns, which is the effective dead-time of the readout electronics. The output from the strips is written to the "strip digi." This includes the strip number, the comparator output and its time, and the 8 SCA pulse-height samples.

3.3.1.3 Digitization of the RPCs

The RPC detector response is simply the digital information on the strips that are fired by a ionizing particle. The time resolution of about 3 ns allows the RPC performance to be described just in terms of efficiency and the number of adjacent strips, called the cluster size, involved in the detection process. The efficiency is assumed to be uniform over the chamber surface, and the cluster size is determined by a simple exponential shape. The intrinsic noise of the RPC is also included as a parameter in the simulation.

3.3.2 Simulation and test-beam comparisons

3.3.2.1 DT results

In Figure 3.14 the drift-time distribution for test-beam data is compared with Monte Carlo simulation. Data were collected with a beam of 300 GeV muons with normal incidence with respect to the chamber plane. The excess observed in real data for small arrival times is likely due to soft δ-rays stopping in the volume of the cell, which are not simulated in the current version of the digitization process. (Only the δ-rays produced in the aluminium plates and in the material of the chamber structure crossing a cell are fully simulated.) This difference can be seen in the distribution shown in Figure 3.15 of the quantity $T_{\max}$, computed from the arrival times, T_i, of 3 consecutive layers in a SL using

$T_{\max}^{123} = (T_1 + T_3)/2 + T_2.$

If all 3 signals detected in the layers come from a crossing muon (i.e., they are not masked by the occurrence of a δ-ray), this quantity measures the drift time (approximately 380 ns) corresponding to one half of the cell size (i.e., 2.1 mm, including the I-beam), independent of the track position and inclination. As will be shown later, the observed differences have a negligible impact on the correct simulation of the local-segment reconstruction in a chamber. More details can be found in [107].

3.3.2.2 CSC results

How well the simulation of the CSCs agrees with reality can only be seen by comparing results from real operating CSCs with those from the simulation in equivalent situations. Prototype and actual CSCs for CMS have operated in standalone situations, with and without the final readout electronics, over a period of several years. Cosmic-ray data have been accumulated, and chambers have operated in a number of test beams.

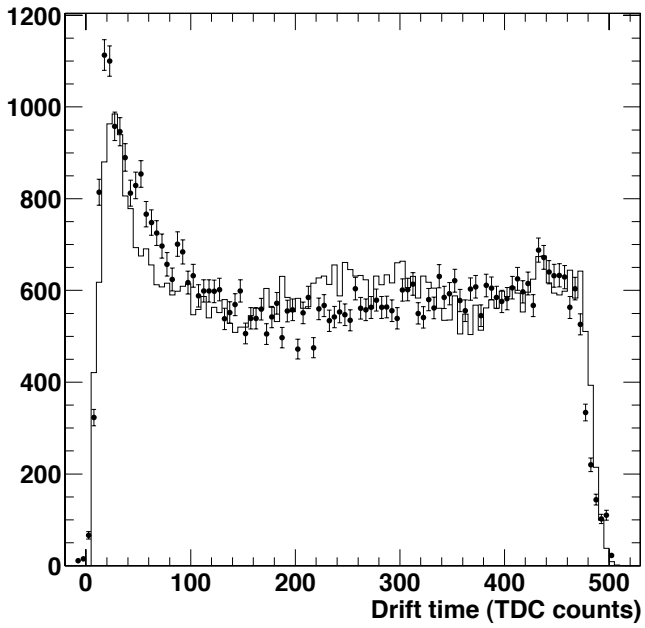


Figure 3.14: Drift-time distribution for test-beam data (dots) and from simulation (solid line). One TDC count = 0.781 ns.

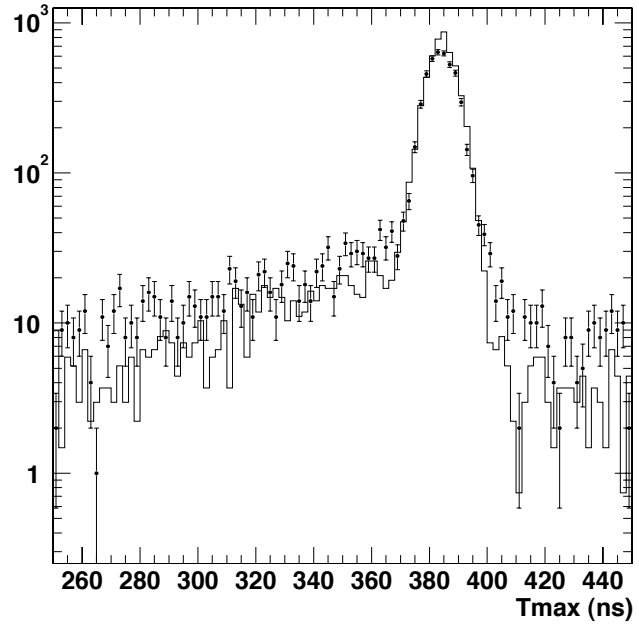


Figure 3.15: Mean-time distribution for test-beam data (dots) and from simulation (solid line).

We mention here 4 particular aspects of the CSCs and the overall CMS detector simulation that were checked in beam tests: the strip pulse shapes sampled by the SCA, the cross-talk between strips, the induced secondary clusters accompanying a muon, and the punchthrough (charged hadrons escaping the calorimeters). A detailed description of the studies focusing on the first 2 aspects can be found in [108], while [109] deals with the hadron punchthrough analysis.

The SCA [85] samples the pulses on hit strips every 50 ns. These values are digitized by 12-bit ADCs during the readout of the strips in chamber regions where a Level-1 CSC Local Charged Track occurs in coincidence with a Level-1 Trigger accept. It is from these values that the highest-precision measurement of a muon hit in each CSC layer is reconstructed, based on the height and position of the peak of the pulse. The quality of the simulation can therefore be examined by comparing the simulated SCA pulse shape with that obtained from the real CSC hardware in a muon test beam.

Cross-talk between strips can lead to a pulse of several percent of the signal from one strip appearing on the 2 neighbouring strips. This cross-talk is the combined result of several effects: strip-to-strip and strip-to-ground capacitances, preamplifier input impedance, and shaper response. This can deteriorate the position resolution, although in principle it can be calibrated out of the measurement.

Measurements on a real CSC mounted in the SX5 assembly hall using a pulser gave the strip response to a delta-function charge input. If these resulting pulse shapes are convoluted with functions representing the ion drift time and the electron collection time, the resulting distributions can be compared with pulses arising from test-beam muons in the equivalent chamber type.

The SCA shape around the maximum can be well fit by a semi-Gaussian function involving the charge on the strip (Q), the arrival time of the pulse (T_s), and the preamplifier peaking time ($4 \times T_0$). The cross-talk between neighbouring strips is taken into account using a functional form determined from the analysis of data collected while pulsing the chambers. From these fits, values for T_0, T_s, and a cross-talk coefficient, C_t, can be extracted [108]. The average value of T_0 is found to be about the same in all chambers with the same gas gap for the 3 types of data used: from 150 GeV muons in a test beam, from a pulser, and from the simulation. The value C_t is about the same for all chambers with the same strip geometry for the test-beam and pulser data. The value is about 15% lower in the simulated data, which is not surprising considering the rather simplistic model of cross-talk currently implemented in the simulation. The T_s value can be subtracted from the time of each SCA time sample for a given strip in each event, leading to a restored average pulse shape measured with much finer time resolution than the original 50 ns sampling [108]. Figure 3.16 compares the restored average pulse shape (6.25 ns binning) from the 3 types of data: the real data are from a 150 GeV$/c$ muon test beam, and the simulated data from single muons with p_{T} = 100 GeV$/c$ and $1.3 < \eta < 1.6$ (to approximately match the test-beam kinematics.) It can be seen that the simulated shape agrees well with reality, demonstrating that the full response of the CSC detectors and electronics is well modelled in the simulation.

Figure 3.17 demonstrates the quality of the modelling of cross-talk in the simulation. We wish to examine the contribution to the SCA pulse height on a strip which arises purely from the cross-talk coupling to the signal on a neighbouring strip. These contributions are extracted from a similar analysis to that described above for the extraction of the restored SCA pulse shape on a strip, but now corrected to eliminate the electrostatic component of the induced charge on the neighbours to a primary signal strip. (This permits test-beam and simulation data to be compared with pulser data in which only the central strip is pulsed.) From this analysis, the cross-talk signal for strips on either side of a primary signal strip can be obtained [108]. The figure presents the ratio of the extracted signal on either neighbouring strip to the sum of the signals on the central strip and its 2 neighbours, as a function of time. The conclusion is that the cross-talk in the test-beam and pulser data are in close agreement,

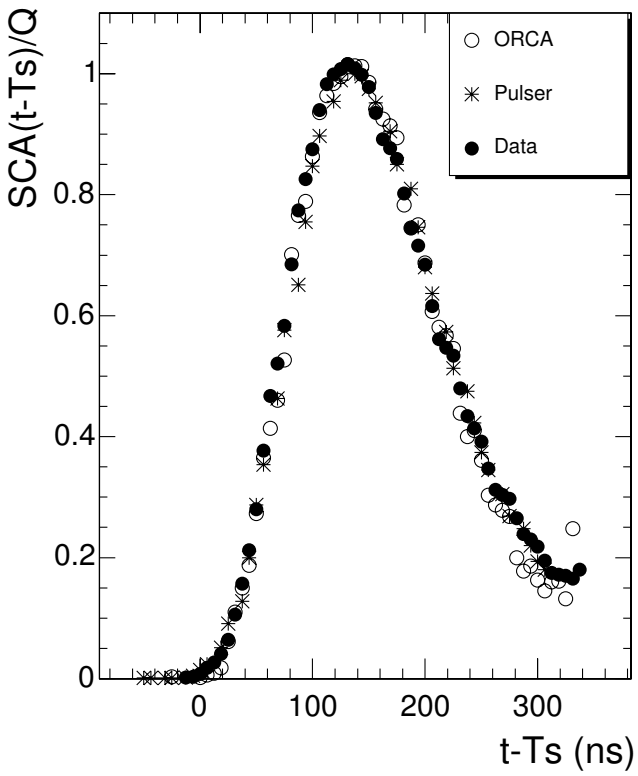


Figure 3.16: Comparison of the SCA-restored pulse shape on strips versus time for test-beam data in an ME2/2 CSC (closed circles), pulser data (crosses), and simulated data (open circles).

while the cross-talk in the simulation is within about 20% of the test-beam measurement. These results will be used to improve the modelling of the cross-talk in the simulation.

An estimation of the calorimeter shower leakage in the ME1/1 region of the CMS detector was done by analysing experimental data taken during the combined test beam of ME1/1 and HE detectors at the CERN H2 beam facility in 2004 [109]. To study the punchthrough effect and the accuracy of the CMS detector simulation, 2 experimental configurations were considered: one with and one without an ECAL module in front of the HCAL module.

In Figure 3.18 the probability of electromagnetic secondaries for muons (knock-on δ-rays, bremsstrahlung, or pair production) in ME1/1 is compared with a GEANT4-based simulation of the test-beam setup. In this case, the secondaries are defined as additional separate clusters (clusters with a distance larger than 2 cm from the muon beam position) in any ME1/1 layer. One can see that the simulation predicts a distribution close to that measured in the test beam.

To estimate the punchthrough probability of charged hadrons, test-beam data taken with various incident pion beam energies were analysed [109], and the fraction of events in which at least one reconstructed hit was found in ME1/1 was considered. For this study, the corresponding simulation results depend significantly on the GEANT4 physics list used. As shown in Figures 3.19 and 3.20, with the ECAL module inserted and removed, respectively, the simulation curves obtained using the LHEP showering model are in good agreement with the measured data.

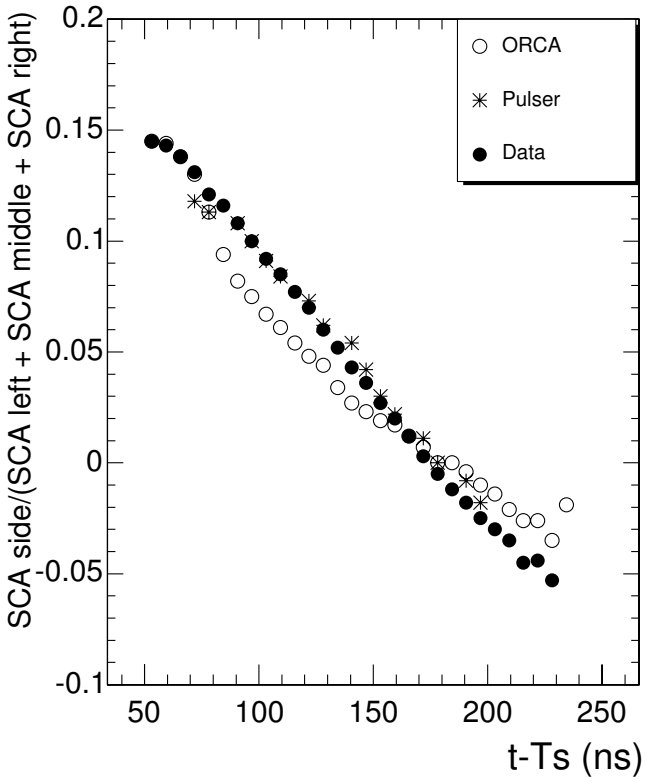


Figure 3.17: The ratio of the SCA-restored pulse height due to cross-talk on a strip to the sum of the pulse heights on the (neighbouring) signal strip and its 2 immediate neighbours, as a function of time for test-beam data in an ME2/2 CSC (closed circles), pulser data (crosses), and simulated data (open circles).

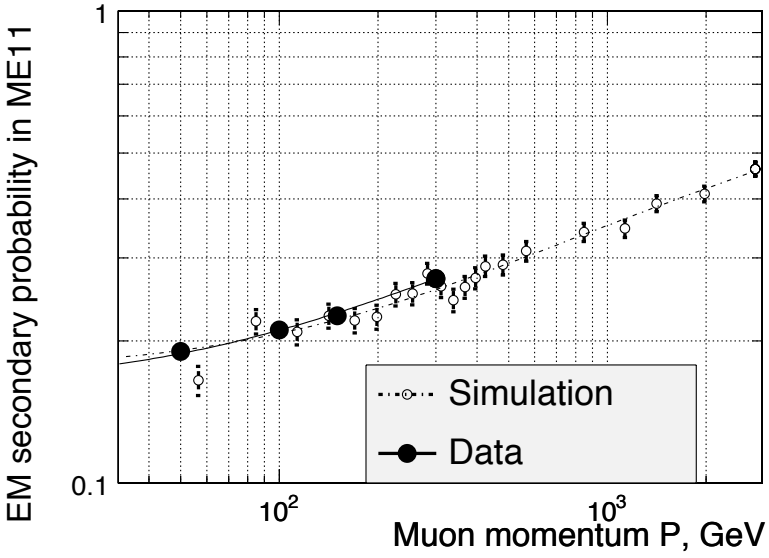


Figure 3.18: Electromagnetic secondaries probability in ME1/1 vs muon beam energy.

The combined analysis of both secondaries and punchthrough effects shows that the simu-

lation based on the LHEP physics list should preferably be used for physics simulation.

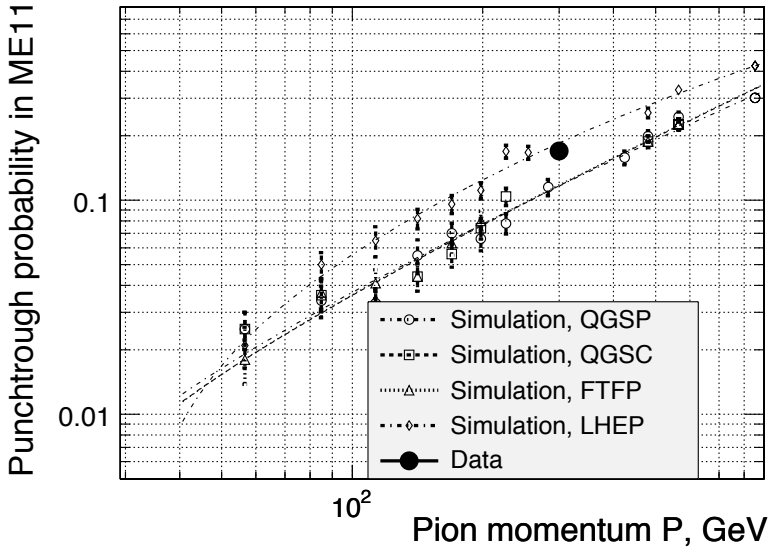


Figure 3.19: Punchthrough probability vs. pion beam energy for the EMU+ECAL+HCAL configuration. Simulations using 4 GEANT4 physics lists are also shown.

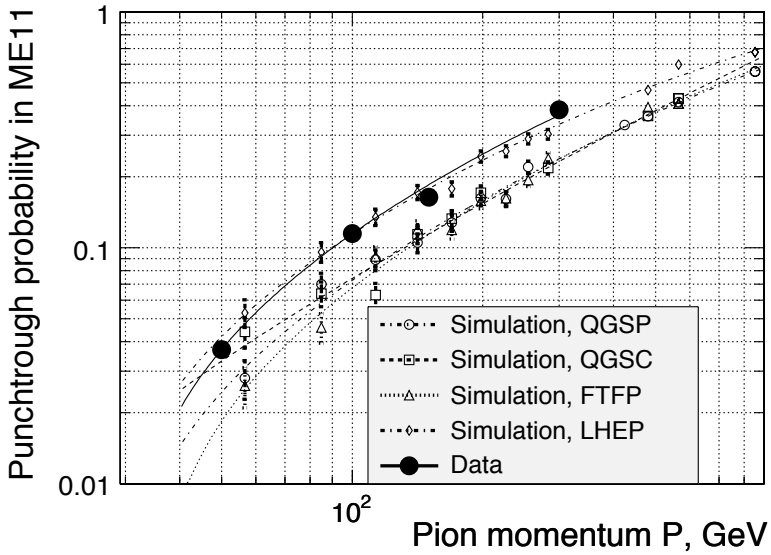


Figure 3.20: Punchthrough probability vs. pion beam energy for the EMU+HCAL configuration. Simulations using 4 GEANT4 physics lists are also shown.

3.3.2.3 RPC Results

Figure 3.21 shows a comparison of the cluster size between RPC data and simulated events. Data were taken in a cosmic-muon test facility [110] using a chamber operating with a detection efficiency of 98%. A simulation sample of single-W events decaying into isolated prompt muons was used for the comparison with the cosmic-muon data. The average cluster size parameter in the simulation was set to 1.6 strip units, as extracted from the data. The distri-

bution of the incidence angle of the muons with respect to the chamber surface was similar in the selected cosmic-muon data and in the simulated sample. The difference seen in the figure is due to the simple functional form used in the simulation, which is better suited for an RPC operating at lower voltage.

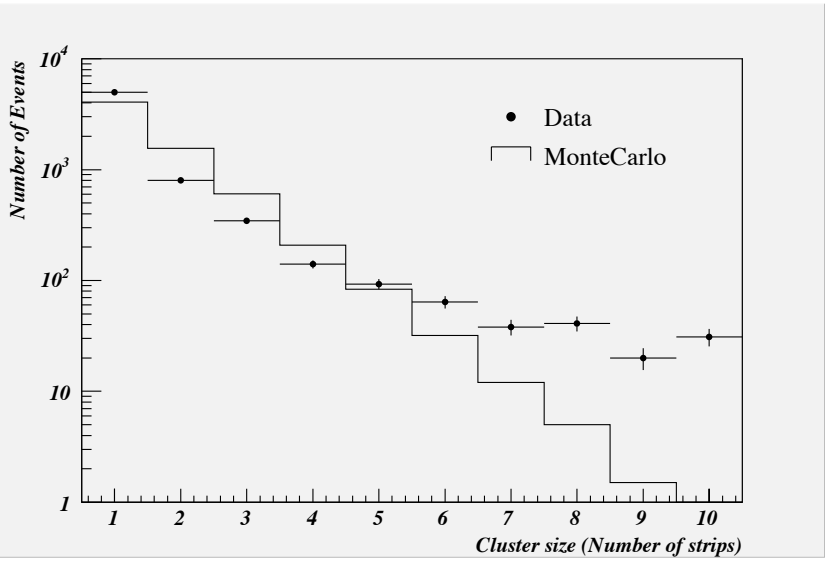


Figure 3.21: Cluster size distribution in an RPC. The dots with error bars correspond to cosmic-muon data; the solid line is obtained from isolated muons in a simulated sample of $W \to \mu\nu$ decays.

3.4 Level-1 trigger

3.4.1 Overview of Level-1 muon trigger algorithms

The purpose of the Level-1 muon trigger of the CMS experiment is to identify muons, assign them to a particular beam crossing, and determine their transverse momenta and location. The compilation of this information and the information from the Level-1 calorimeter trigger (discussed in Sections 4.2 and 5.2.2) is used to decide whether to keep the data from a particular beam crossing for further processing.

The Level-1 muon trigger is based on custom electronics and is organized into subsystems representing the 3 different muon detectors: the DT trigger in the barrel, the CSC trigger in the endcap and the RPC trigger covering both barrel and endcap. The Level-1 muon trigger also has the Global Muon Trigger (GMT) that combines the trigger information from the DT, CSC, and RPC muon subsystems, as well as from the calorimeter subsystem, and sends it to the Level-1 Global Trigger.

Each of the Level-1 muon trigger subsystems has its own trigger logic. The DT and CSC electronics first process the information from each chamber locally, delivering 1 vector (position, direction, bunch crossing, and quality) per muon per station. Such vectors produced by local triggers are often referred to as trigger primitives. Trigger primitives from different stations are collected by the Track Finders (TF), which build them into tracks and assign a transverse momentum value to each. Therefore, the TF plays the role of a regional trigger.

The DT and CSC Track Finders exchange track segment information in the pseudorapidity region where these systems overlap. Up to 4 best (highest p_{T} and quality) muon candidates from each subsystem are selected and sent to the GMT.

In the case of the RPC, there is no local processing apart from synchronization and cluster reduction. Hits from all stations are collected by the Pattern Comparator Trigger logic. If they are aligned along a possible muon track, a candidate is formed and a p_{T} value is assigned. Found muon candidates are ranked based on their quality and p_{T}; up to 4 best candidates from the barrel and 4 from the endcaps are sent to the GMT.

The Global Muon Trigger attempts to correlate the DT and CSC muon candidates with the RPC candidates. Bits delivered by the calorimeter trigger are used to determine if these muons are isolated. The final ensemble of muons is sorted based on their quality, correlation, and p_{T}, and the 4 best muons are then transmitted to the Global Trigger. Finally, transverse momentum thresholds are applied by the Global Trigger for all trigger conditions.

A detailed description of the CMS Level-1 muon trigger can be found elsewhere [7, 8]. In the following, we give only a brief overview, mentioning in more detail recent changes in the algorithm and updating estimates of its performance.

3.4.2 DT trigger

The DT front-end trigger electronics, called the Bunch and Track Identifier (BTI), forms track segments from coincidences of at least 3 aligned hits in 4 layers of 1 drift tube Superlayer. The BTI also performs the bunch-crossing assignment for every found segment. The positions and directions of the segments are sent to the Track Correlator (TRACO), which attempts to combine the segments from 2 Superlayers and produce a ϕ coordinate. The Trigger Server selects, in each chamber and at every bunch crossing, 2 segments with the highest quality and the smallest angular distances with respect to the radial direction to the vertex (i.e., with the highest p_{T}), and transmits them to the Sector Collector, where the trigger information is coded and sent via optical links to the DT Track Finder (DTTF).

The DTTF system is divided into sectors, each of them covering 30° in ϕ angle. Each Sector Processor is in turn logically divided into 3 functional units: the Extrapolator Unit, the Track Assembler, and the Assignment Unit. The Extrapolator Unit tries to match pairs of track segments from distinct stations, using the spatial coordinate ϕ and the direction of the segments. The 2 best-matched segments per each source segment are forwarded to the Track Assembler, which attempts to find the 2 tracks in a sector with the highest rank, i.e., exhibiting the highest number of matching track segments and the highest extrapolation quality. Once the track segment data are available to the Assignment Unit, memory-based look-up tables are used to determine the transverse momentum of the candidate track, its ϕ and η coordinates, and the track quality. Finally, the Wedge and Barrel Sorters apply cancellation schemes to reduce the number of tracks found in neighbouring sectors and wedges, and forward the 4 highest-p_{T} tracks to the GMT.

3.4.3 CSC trigger

The CSC local trigger uses the 6-layer redundancy of the CSCs to measure precisely the ϕ coordinate and to find the correct bunch crossing for a track. Muon segments, also known as Local Charged Tracks (LCT), are found separately in the nearly orthogonal cathode and

anode projections by somewhat different algorithms and by different electronics boards; the cathode electronics design is optimized to measure the ϕ value and the direction with high precision, while the anode electronics design is optimized to determine the muon bunch crossing and the η value. Up to 2 cathode and 2 anode LCTs can be found in each chamber during any bunch crossing. The 2 projections are then combined into three-dimensional LCTs by a timing coincidence in the Trigger MotherBoard (TMB). The Muon Port Card receives the LCTs from all of the TMB cards in one 60° azimuthal sector of 1 endcap muon station, selects the 3 best LCTs, and sends them over optical fibre links to the CSC Track Finder (CSCTF).

Each CSCTF unit finds muon tracks in a 60° sector. A single extrapolation unit forms the core of the CSCTF logic. It takes the three-dimensional spatial information from 2 LCTs in different stations, and tests if those 2 segments are compatible with a muon originating from the nominal collision vertex with a curvature consistent with the bending in the magnetic field in that region. Each Sector Processor can find up to 3 muon candidates and measure their parameters using look-up tables. The Muon-Sorter module selects the 4 best muon candidates in the entire CSC system, and sends them to the Global Muon Trigger.

An important change in the CSC trigger logic was the extension of the pseudorapidity coverage of the dimuon trigger from $|\eta| < 2.1$ to $|\eta| < 2.4$. Since the electronics of the CSC chambers covering the $|\eta|$ region from 2.1 to 2.4 in the first muon station was descoped (see Section 3.1), muon tracks in this pseudorapidity region can only be reconstructed from segments in the other 3 stations and, because of the reduced magnetic field in these stations, with poor p_{T} resolution. Detailed studies showed, however, that improvements in the Sector Processor firmware allow one to use these low-quality tracks in the dimuon trigger without increasing the rate of fake dimuons. This extension of the η coverage of the Level-1 dimuon trigger resulted in a higher trigger efficiency for a number of interesting physics channels [111].

3.4.4 RPC trigger

The RPC Pattern Comparator Trigger (PACT) is based on the spatial and time coincidence of hits in the RPC muon stations. The PACT electronics performs 2 functions: it requires a time coincidence of hits in patterns along a certain road and assigns a p_{T} value. The time coincidence gives the bunch-crossing assignment for a candidate track formed by a pattern of hits that matches with 1 of many pre-defined patterns obtained from simulation. The patterns are divided into classes with a p_{T} value assigned to each of them. PACT is a threshold trigger: it gives a momentum code if an actual hit pattern is straighter than any of the pre-defined patterns with a lower momentum code.

The PACT muon candidates are then filtered through the ghost-cancellation logic and sorted according to their quality and p_{T}; up to 8 best muon candidates (4 from the barrel and 4 from the endcaps) are sent to the GMT.

3.4.5 Global Muon Trigger

The GMT receives the best 4 barrel DT and the best 4 endcap CSC muon candidates and combines them with the best 4+4 candidates sent by the RPC. It performs a matching based on the proximity of the candidates in (η, ϕ) space. If 2 muons are matched, their parameters are combined to give optimum precision. If a muon candidate cannot be confirmed by the

complementary system, quality criteria are applied to decide whether to forward it or not. The GMT also contains logic to cancel ghost tracks that arise when a single muon is found by more than 1 muon system and is not otherwise matched, such as at the boundary between the DT and CSC muon systems.

The selected muon candidates are ranked based on their detector type, p_{T}, quality, and η. The GMT also extrapolates the muon tracks back to the calorimeter trigger towers and appends the corresponding ISO (isolation) and MIP (compatibility with a minimum ionizing particle in $\Delta\eta \times \Delta\phi = 0.35 \times 0.35$ regions) bits to the track data consisting of p_{T}, sign of the charge, η, ϕ, and quality. The 4 best muon candidates in the entire CMS detector are sent to the Global Trigger.

3.4.6 Test beam results and trigger emulator

Numerous tests of the trigger electronics prototypes, performed prior to 2002 in preparation for their mass production, are described in [7], [112], and [113], and the references therein. In recent years, the performance of the production on-chamber electronics and near-final prototypes of the off-chamber electronics were tested in CERN test beams with a 25 ns bunch structure similar to that of the LHC. The results obtained [82] further validated the trigger design and showed that the trigger system meets all requirements.

One important goal of the tests was the validation of various parts of the Level-1 muon trigger emulator. This emulator was developed as a tool to be used in the framework of the official CMS reconstruction package, ORCA [27]. The emulator must exactly reproduce the workings of all the trigger electronics boards and their firmware, including all the features and algorithms of the trigger subsystems outlined in the previous sections. In many cases, a bit-by-bit comparison between the output trigger data and the response of the emulator of a given board to the input trigger bits was performed, with the emulator response found to be in perfect agreement with the data.

Several examples of the results obtained in the 25 ns-structured test beams in 2003 are given below. Figure 3.22 [82] shows the efficiency to reconstruct a DT trigger track segment and to assign it to the correct bunch crossing as a function of the track angle; the test-beam measurements (dark circles) are compared with the predictions of the detector simulation and trigger emulator (open circles). The agreement is better than 0.1% at angles of incidence less than 25°; the efficiency drop at the largest angle is not a problem since it corresponds to very low-p_{T} muons.

Figure 3.23 shows typical distributions of the difference in the muon arrival time (in 25 ns clock ticks) between the anode trigger circuitry of the CSCs and the timing from the scintillator hodoscope. As one can see, the efficiency to identify the correct bunch crossing for a muon segment in the CSCs approaches 99%. The predictions of the emulator agree with the measurements to better than 0.1%.

The analysis of the test-beam data is still in progress. More results are expected to be published soon, and more refined versions of the muon trigger emulator will become available.

3.4.7 Muon trigger performance and trigger rates

The Level-1 muon trigger emulator discussed in the previous section has been used to evaluate the performance of the muon trigger in CMS and to estimate the expected single-muon

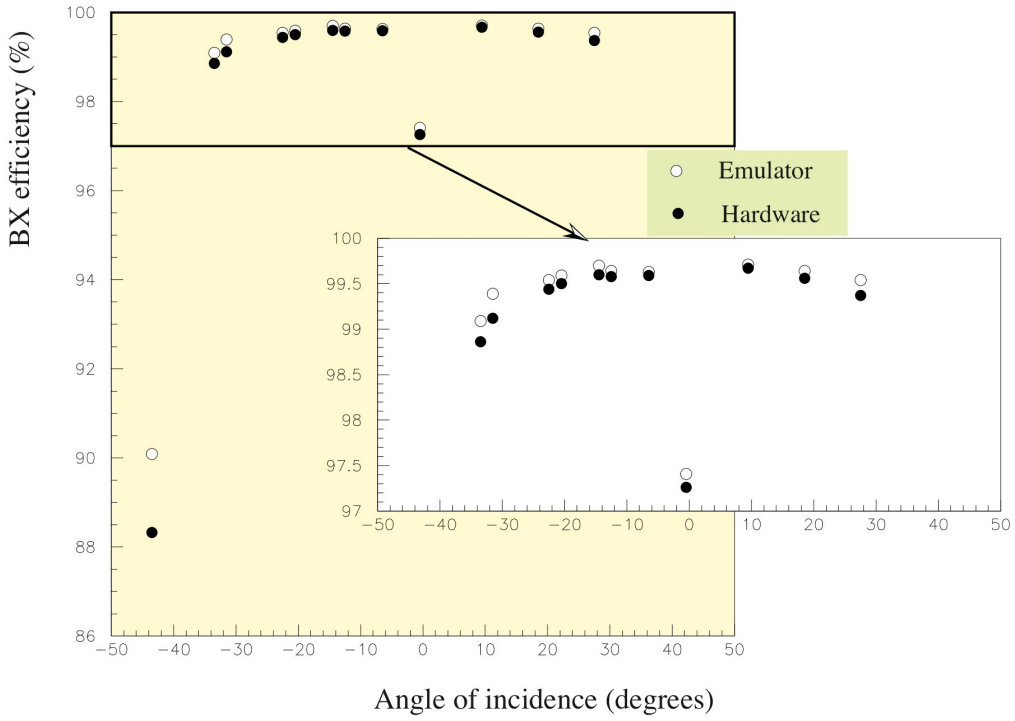


Figure 3.22: Bunch-crossing efficiency of the DT local trigger as a function of the angle of incidence [82]. The efficiency measured in the test-beam data (dark circles) is compared with that computed by the trigger emulator (open circles).

and dimuon trigger rates.

Figure 3.24 shows the efficiency to identify single muons in various Level-1 muon trigger subsystems as a function of the muon pseudorapidity. The muons were generated flat in the intervals $5 < p_{\mathrm{T}} < 100$ GeV/c and $|\eta| < 2.4$. The average identification efficiency of the Global Muon Trigger is 98.3%; the losses of efficiency in some $|\eta|$ regions are due to the gaps between the muon chambers. The corresponding probability to reconstruct extra fake muons ("ghosts") is 0.3%. More results on the performance of the Level-1 muon trigger, notably the "turn-on" curves and the p_{T} resolutions, are reported in [8] and [7]. (These and other results were used to tune the parameterizations of the Level-1 muon trigger in the fast simulation of CMS, FAMOS, discussed in Section 2.6.4.)

The Level-1 single-muon and dimuon trigger rates for low luminosity as a function of the p_{T} threshold are shown in Figure 3.25; also shown are the generated rates. As can be seen from the single-muon rate plot, a large increase in the trigger rates would occur if the RPC system or the combined DT/CSC system operated in a standalone mode, without the optimization of the GMT. The dimuon rate plot shows 2 main contributions to the total dimuon rate: 2 muons produced in the same pp interaction, and 2 muons from different pp interactions in the same bunch crossing; at low lumunosity, the former dominates at all p_{T} values. All rates were calculated according to the first of the 2 methods described in [114], whereby the effect of pile-up is taken into account analytically. We have checked that the expected rates do not change significantly when the effect of the misalignment of the muon chambers (simulated with the help of the misalignment scenarios described in Section 3.2.2) is taken into account [95].

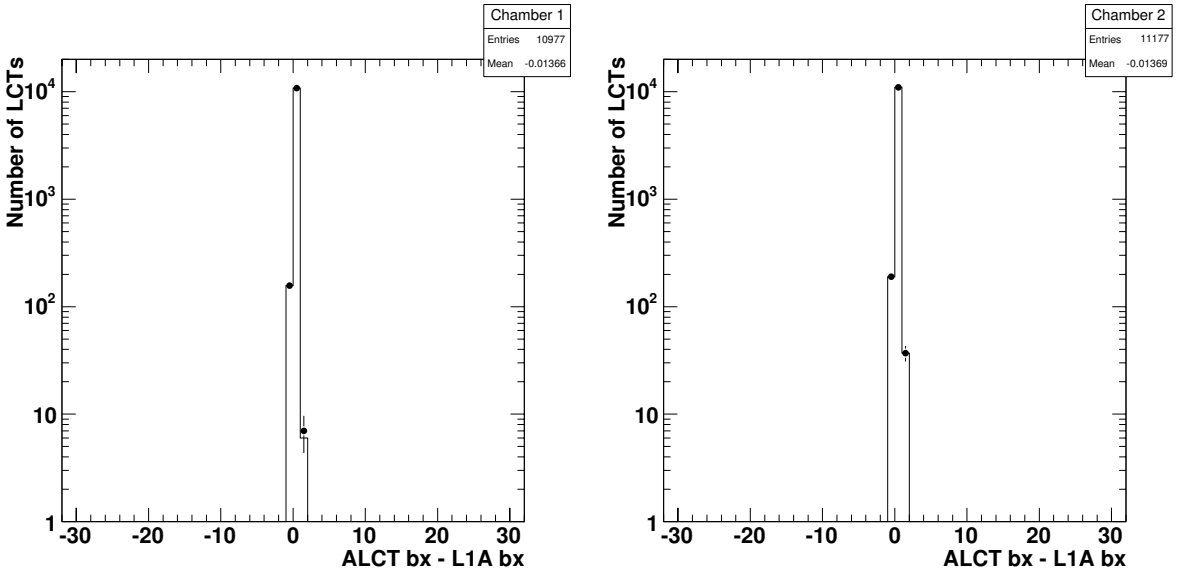


Figure 3.23: Difference in the muon arrival time (in 25 ns clock ticks) between the anode trigger circuitry and the timing from the scintillator hodoscope for 2 CSCs in a typical run at the X5 test beam at CERN in 2003 (points with error bars) and for the trigger emulator using the anode raw hits recorded in this run (histogram).

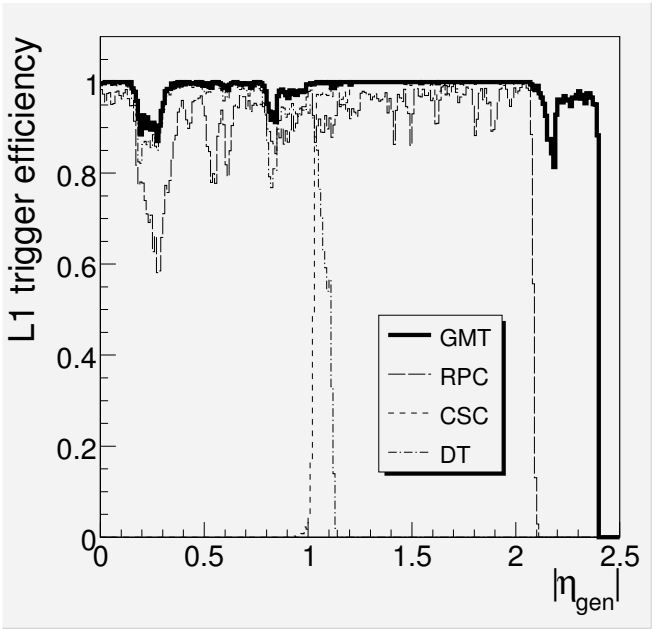


Figure 3.24: Efficiency for identifying single muons as a function of the muon pseudorapidity for the Global Muon Trigger (solid line) and for the DT, CSC, and RPC trigger subsystems (shown in dash-dotted, dotted, and dashed lines, respectively).

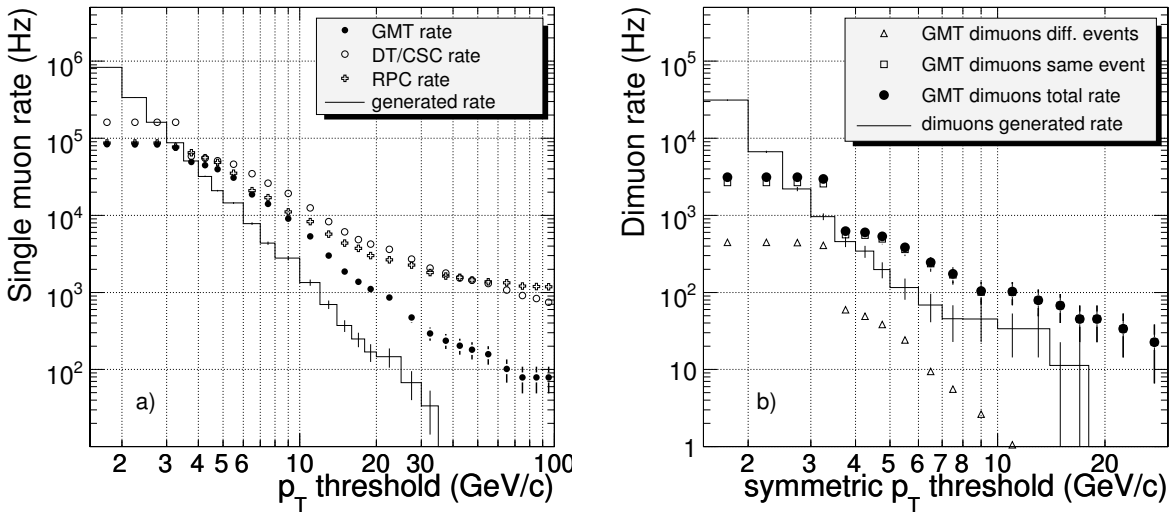


Figure 3.25: Level-1 trigger rate at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ as a function of p_{T} threshold for a) single-muon trigger and b) dimuon trigger, at the generator level (histogram) and from the Global Muon Trigger (dark circles with error bars). The single-muon rate plot also shows the trigger rates that would occur if the RPC system or the combined DT/CSC system operated standalone (crosses and open circles). The dimuon rate plot shows separately the contributions from the same (squares) and different (triangles) pp collisions.

Figure 3.26 shows contours of equal total rate (in kHz) from the Level-1 single-muon and dimuon triggers, in the plane of the applied p_{T} thresholds, for low and high luminosity; the threshold for the dimuon trigger applies to both muons. Assuming the allocation of Level-1 triggers proposed in [8], a bandwidth of about 4 kHz is available for muon triggers at low luminosity, and of about 8 kHz at high luminosity. In [8], the operating thresholds were chosen to be 14 GeV/c and 3 GeV/c for the single muon and dimuon triggers, respectively, at low luminosity, and 20 GeV/c and 5 GeV/c at high luminosity. This choice remains adequate to this day, corresponding to a total muon trigger rate of 4.9 ± 0.2 kHz at low luminosity, and of 6.6 ± 0.6 kHz at high luminosity, where the uncertainties are statistical only.

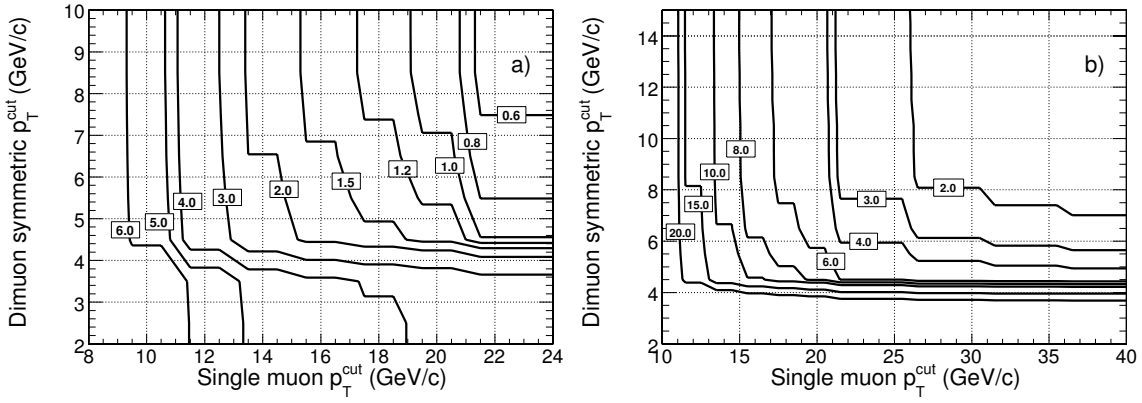


Figure 3.26: Contours of equal rate (in kHz) in the plane of p_{T} thresholds for Level-1 single-muon and dimuon triggers at (a) $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and (b) $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Excellent trigger performance and low operating p_{T} thresholds ensure high efficiency for

many physics channels involving muons. For example, at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, the combined efficiency of single-muon and dimuon Level-1 triggers for the Standard Model Higgs boson with a mass of $120\ \mathrm{GeV}/c^2$ and decaying via the channel $H \to WW^{(*)} \to \mu\mu\nu\nu$ is 97.6 $\pm$ 0.4 (stat.)%; the corresponding efficiency for $H(115\ \mathrm{GeV}/c^2) \to ZZ^{(*)} \to 4\mu$ is 100%. The trigger efficiencies for these channels remain high at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, 91.3 $\pm$ 0.7 (stat.)% for $H(120\ \mathrm{GeV}/c^2) \to \mu\mu\nu\nu$ and 99.6 $\pm$ 0.1 (stat.)% for $H(115\ \mathrm{GeV}/c^2) \to 4\mu$. For both channels, the efficiencies increase slightly with the Higgs boson mass, and remain practically unchanged when the effect of the misalignment of the muon chambers is taken into account. The trigger efficiency stays high up to very large values of muon transverse momentum; for example, the efficiency expected for dimuon decays of a new heavy resonance, such as a Z′ or graviton, with a mass of $3\ \mathrm{TeV}/c^2$ is about 99%. All the efficiencies listed here were calculated relative to the number of events with all final-state muons (2 or 4, depending on the channel) inside the full acceptance of the muon system, $|\eta| < 2.4$.

The lowest possible p_T thresholds and large acceptance of the muon triggers are crucial for the detection of heavy quarkonia resonances, which are an important tool to study the hot, dense matter produced in heavy ion collisions. The detectable muons from the decays of J/ψ and Υ produced in these interactions are expected to be at relatively high pseudorapidity. For instance, the simulation predicts that almost 40% of the J/ψ decaying into 2 muons capable of reaching the muon detectors will give at least 1 muon at $|\eta| > 2.1$. With the extension of the η coverage of the CSC dimuon trigger described above, the Level-1 trigger logic foreseen for pp collisions is also suitable for heavy ion interactions. Lower p_T thresholds will be used for a number of ion species; e.g., the default Level-1 trigger configuration for Pb-Pb collisions at the luminosity $\mathcal{L} = 10^{27}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ is expected to include the single muon trigger with the minimal p_T thresholds (ranging from about 3.5–4 GeV in the barrel to 1.5–2 GeV in the endcap) [7].

In addition to the single-muon and dimuon triggers, muons reconstructed by the GMT are used in a number of cross-channel triggers; these triggers are discussed in Vol. 2 of this report.

3.5 Detector Level Reconstruction

3.5.1 Position reconstruction and track-segment reconstruction in the barrel drift tubes

3.5.1.1 Input to the local reconstruction

The information stored in a DT digi is a TDC measurement. In order to extract the drift time, a synchronization procedure is needed. In fact, the raw time also contains contributions from the time-of-flight of the muon from the interaction point to the cell, from the propagation time of the signal along the wire, and an offset due to the electronics and cables. In the simulated data, various delays are added to the drift time during digitization in order to fake the real case. (See Section 3.3.1.1 for details).

The synchronization procedure for data is described in Section 3.2.1.1. Further corrections for the TOF and the signal propagation along the wire with respect to the average values, computed calculating the t_0 offset with the SL granularity, can be applied during reconstruction, as soon as information about the hit position within the cell volume is available.

3.5.1.2 Reconstruction of hit positions within the cell

The primary objects that result from the DT local reconstruction are points in the cell volume: their distances with respect to the wire are computed by converting drift times to drift distances.

Two reconstruction algorithms are currently available in ORCA, one that uses a constant drift velocity over the entire cell, and another, used as default, that is based on a time-to-distance parameterization obtained using a GARFIELD [104] simulation of the cell behaviour.

This parameterization is based on the same GARFIELD simulation of the drift cell that is used in the digitization step (Section 3.3.1.1). The measured coordinate $x_{\rm drift}$ is computed using a function of the drift time; $B_{\parallel}$ and $B_{\perp}$, the components of the B field parallel and perpendicular, respectively, to the wire in the radial direction; and α, the incidence angle with respect to the normal direction to the chamber:

$$x_{\rm drift} = x_{\rm drift}(t_{\rm drift}, B_{\parallel}, B_{\perp}, \alpha) \ . \tag{3.1}$$

The component of the magnetic field parallel to the drift lines can be neglected since it has no measurable effect on the drift time.

Since $B_{\parallel}$, $B_{\perp}$, and α are not known at the level of the individual hit, a 3-step reconstruction procedure is implemented. The first step assumes a crude estimate of the impact angle and the hit position along the wire. The hit is then updated twice: after it has been used to build a 2D r-ϕ or r-z segment (second step) and after it has been used in the 3D segment fit.

By default, the errors on the reconstructed hit positions are obtained from a Gaussian fit of the residual distributions (the difference between simulated and reconstructed distances from the wire).

The reconstruction algorithm, which uses a constant drift velocity to compute the hit distance from the wire using a constant drift velocity, neglects any dependency on the magnetic field or the track angle. It is possible to use different drift velocities for different channels (with the desired granularity) estimating them with the calibration procedures described in Section 3.2.1.1.

3.5.1.3 Segment reconstruction

The segment reconstruction works on the r-ϕ and r-z projections independently: only at the end of the procedure are the 2 projections combined and a three-dimensional segment is built. The reconstruction in each projection is an exercise in pattern recognition and linear fitting, with the complication of left-right ambiguity. The reconstruction is performed in 3 steps:

- a segment candidate is built from a set of aligned hits;
- the best segments among those sharing hits are selected, solving conflicts and suppressing ghosts
- the hit reconstruction is then updated using the information from the segment and the segments themselves are eventually refitted.

The first step begins by selecting 2 hits in different layers, starting from the most separated. Both hypotheses, left and right, are considered. The pair is kept if the angle of this proto-segment is compatible with a track pointing to the nominal interaction point. For each pair,

additional compatible hits are searched for in all layers. It is possible that the left and right hypotheses are both compatible with the segment. In this case, the ambiguity is solved later. At this stage, a linear fit is performed using the hit positions and errors for each collection of hits. The segment with the maximum number of hits and the minimum χ^2 is retained; all the others are rejected. Finally, a quality criterion is applied, requiring the number of hits ≥ 3 and the $\chi^2/\text{ndf} < 20$. The same algorithm is applied for a single superlayer r-z or, in the case of the r-ϕ projection, for 2 SLs, considering all 8 layers together. In all cases, a protection against the detector being too noisy is applied, so the reconstruction is not attempted if the number of hits is greater than a predefined number.

The pattern recognition described above produces a set of segment candidates. A consistency check is performed in order to test whether 2 segments share the same hit. Conflicts are solved on the basis of the number of hits and the χ^2 of the segment.

The hits from the leftover candidates are updated, taking into account the incidence angle as reconstructed by the segment. The segment linear fit is then redone using the updated hits.

Up to this point, the r-ϕ and r-z projections are treated independently. As the 2 projections are orthogonal, a segment in 1 projection cannot be used to validate or invalidate a segment in the other. By construction, all combinations of segments from the 2 projections are kept. In practice, more than 1 segment is found for either projection in a given chamber in less than 1% of the cases.

Knowledge of the position along the wire, as well as the angle with respect to the drift cell, are provided to the hit reconstruction algorithm before combining the 2 projections performing the final fit of the segment. The result is a segment with parameters in space, suitable for use in the global reconstruction.

3.5.1.4 Performance

The plots in Figure 3.27 show the resolution on the reconstructed hit distance from the wire in the DT cells of r-ϕ superlayers. They have been obtained comparing the simulated hits with the results of the 3-step reconstruction procedure based on the time-to-distance parameterization for muons generated with a flat p_{T} spectrum in the range 5-100 GeV/c. The long tail in the negative values is due to the contribution of secondaries. Combining the hits of r-ϕ SLs in a 2D segment the achieved resolution on the position is about 70 μm while the resolution of the direction is about 2 mrad.

3.5.2 Clusterization and track-segment reconstruction in the endcap cathode strip chambers

3.5.2.1 Input to the local reconstruction

The inputs to the local reconstruction are the detected signals from the cathode strips and anode wires. The information about these signals are contained in the strip and wire digis, respectively. In the simulation software, the digis are the output of the preceding simulation of the electronics response of the detector to the passage of generated particles. In the real world, the digis will be derived from the actual DAQ output stream.

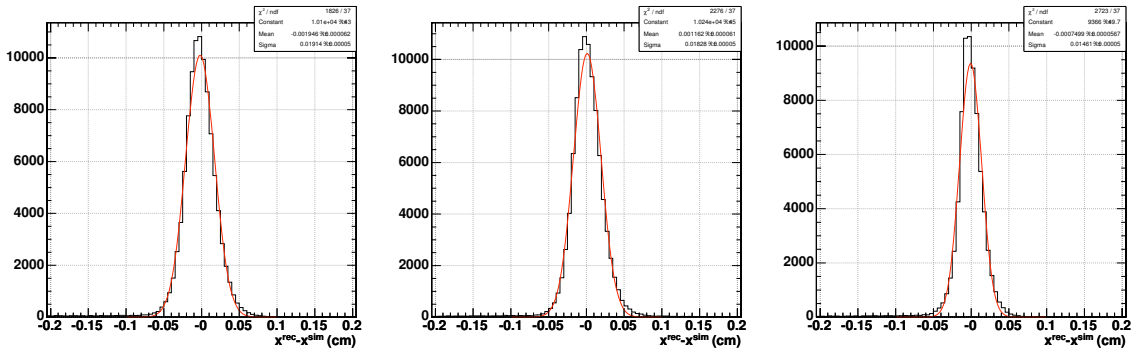


Figure 3.27: Residuals between the reconstructed and the simulated hit distances from the wire. The 3 plots refer to the first (left) second (center) and third (right) reconstruction steps of the drift time parameterization in superlayers r-ϕ.

The charge distribution due to the passage of a single charged particle through a chamber layer is typically distributed over 3 to 5 strips, with a shape described by the "Gatti" parameterization [106]. This charge distribution is reconstructed from the strip-charge samples measured by the SCA, and the position at which a muon passed through the chamber is identified with the position of the peak of the Gatti distribution.

3.5.2.2 Local reconstruction

The basic procedures for local reconstruction in the endcap muon CSCs have already been described in section 8 of an earlier CMS Note [115], and expanded in another [116].

The simple approach is to obtain the pulse height in each strip from the SCA information (in a strip digi), and then to cluster neighbouring strips to determine the probable position of incidence of the incident muon.

Each of the 6 layers of a chamber is considered independently. A two-dimensional "RecHit" is created at each intersection of a 3-strip cluster and a wire group. At a later stage of local reconstruction, the RecHits in the 6 layers of a chamber are fit to form a "track segment" for use in the first stage of the muon track reconstruction.

3.5.2.2.1 Finding the pulse height on a strip We determine the pulse height on each strip from the SCA information. We find the peak of the distribution in time of the 8 samples stored by the SCA, 50 ns apart. Simply selecting the largest of the 8 samples is as satisfactory as attempting a fit over several time bins, since doing likewise with neighbouring strips, sampled at the same time, will give a good representation of the charge distribution across the strips. We use the average of the pulse heights from the first 2 time samples (before the pulse has developed) to correct for baseline shifts by subtracting this average from the maximum pulse height. The estimated error on this pulse height is obtained by adding in quadrature a 3% calibration uncertainty and the expected noise. The noise on a strip channel is dominated by the electronics noise of the input preamplifier. Therefore low-noise amplifiers have been used, and the simulated noise is equivalent to about 1.5% of the pulse height at the Landau peak.

3.5.2.2.2 Fitting clusters of strips To obtain optimal position resolution from the strip measurements, we fit a Gatti distribution to the pulse heights on 3 neighbouring strips. We loop over all strips with a signal (strip digi), starting with the strip with the highest charge. The pulse height on each strip is corrected by a small amount to compensate for the expected cross-talk between channels. We find the local y coordinate of the intersection of each strip with each fired wire group (wire digi). We restrict the association to those strip and wire digis which occur within 2 bunch crossings of each other. The starting value for the fit to the Gatti distribution is the charge centroid; if the fit fails or has a poor χ^2, we take the centroid as the position.

3.5.2.2.3 The two-dimensional RecHit We build the CSC two-dimensional RecHit from the local x and y values, calculating their uncertainties from the wire measurement resolution ($w/\sqrt{12}$, where w is the width of the wire group) and the error on the strip position from the fit.

3.5.2.3 Segments reconstruction

The final stage of the local reconstruction in the CSCs is to build track segments within each CSC from the RecHits reconstructed in each of the 6 layers. The algorithm starts with the first and last RecHit in a chamber (the hit with the most negative local x in layer 1 and the hit with the most positive local x in layer 6) and constructs a straight line between them. The starting and ending points are required to have an r-ϕ separation of at most $1 \ \mathrm{cm}$. Then for each intermediate layer, we successively attempt to add 1 hit and update the linear fit accordingly. Only hits reasonably close (within 2.5 mm in r-ϕ) to the line are considered, and the χ^2/d.o.f. of the updated fit must be reasonable. The hit giving the best fit from all compatible hits on a layer is kept. A track segment is defined only if there are at least 4 hits on such a line. The associated hits are then flagged as 'used', and the procedure is iterated starting with another pair of unused hits for the end points of a line. The fit is performed simultaneously in the local z projections for x and y. If no segment is built using the default hit-association windows, a second pass is allowed in which the windows are broadened.

There has been as yet little investigation of possible biases introduced by this algorithm. For example, it is conceivable that a bias toward low-p_{T} tracks might arise from the choice of starting and ending points. However, studies using a version of the algorithm in which all possible combinations of starting and ending hits are tried, the resulting segments sorted in order of decreasing quality (χ^2/d.o.f.), and segments successively selected so that no hits are shared on the final segments, show that such biases do not seem to be significant. The increased combinatorics of this procedure, however, quickly render it much slower than the original one in the presence of many hits in the same chamber.

3.5.2.4 Performance

Figures 3.28 and 3.29 [116] show the r-ϕ residuals on the azimuthal position, for various subsets of chambers, obtained comparing simulated and reconstructed 2D hits for a sample of muons with flat p_{T} in the range $5 < p_{\mathrm{T}} < 100 \ \mathrm{GeV}/c$. All expected sources of electronic noise are simulated. The resolution is dominated by electronics noise and cross-talk between neighbouring strips. The typical resolutions obtained per layer for single-muon tracks, using simulated cross-talk and noise, are on order the of 100 μm for ME1/1 chambers and 200 μm

for the other chambers. If cross-talk and electronics noise are neglected, these resolutions are reduced to a range of about 35 to 55 μm. When a track passes close to the centre of a strip, the charge distribution tends to be symmetrically distributed over each neighbouring strip. The difference between the smaller signals on the neighbours is then controlled by electronic noise, which subverts the accuracy of the reconstructed hit position. The position resolution thus depends sensitively on the distance of the hit position from the centre of the strip. It also depends on the strip width and the gas gap.

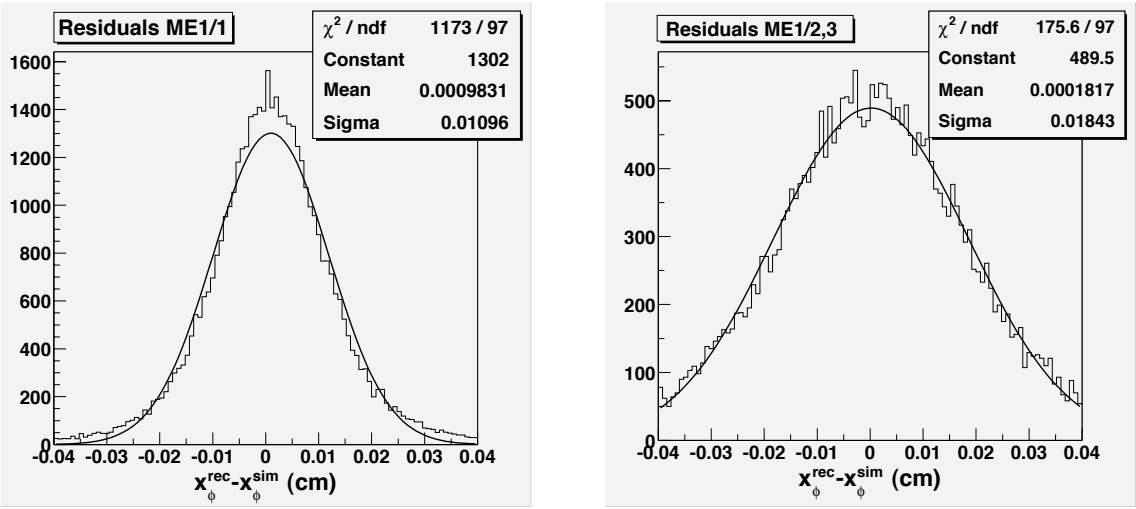


Figure 3.28: Resolution of the reconstructed hit position in the azimuthal direction (in cm) for endcap CSCs in station 1. The left plot is for ring 1 and the right plot for rings 2 and 3.

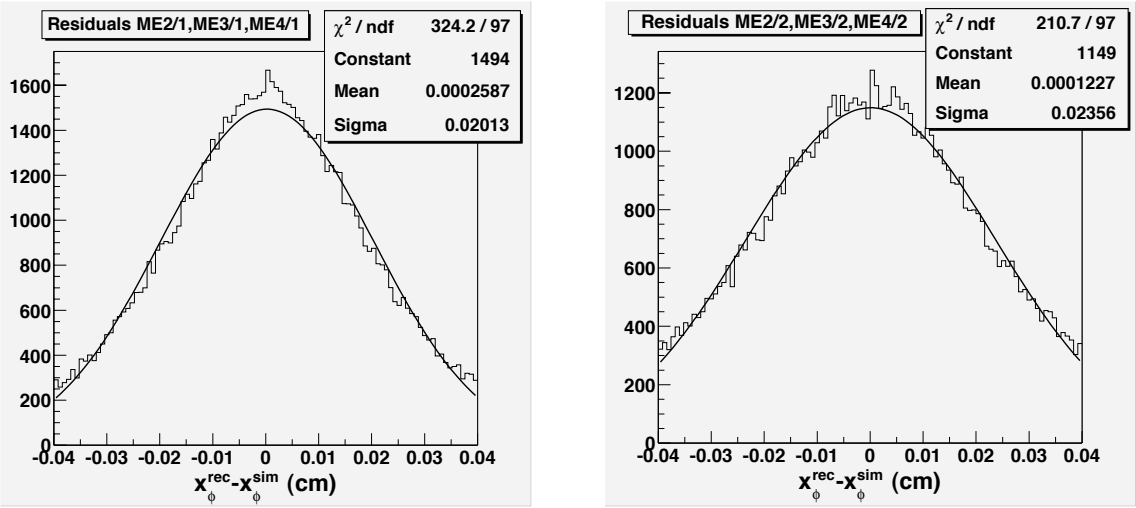


Figure 3.29: Resolution of the reconstructed hit position in the azimuthal direction (in cm) for endcap CSCs in stations 2, 3, and 4. The left plot is for ring 1 and the right plot for ring 2.

3.5.3 Clusterization in resistive plate chambers

In the RPCs, the results of the local reconstruction [116] are points in the plane of the detector. A clustering procedure is first performed, using all the strips that have a signal. The procedure combines all adjacent fired strips into a cluster. Once all the clusters are formed, the

reconstructed point for each is defined as the "center of gravity" of the area covered by that cluster of strips. In the barrel, where the strips are rectangular, this point is simply the center of a rectangle. In the endcap, the computation is more complicated since the area covered by the clusters are trapezoids of variable shape. The assumption is that each group of strips is fired as the result of a single particle crossing the chamber plane, and that this crossing can have taken place anywhere with flat probability over the area covered by the strips of that cluster. Errors are computed under the same assumption of flat probability. Therefore, they are simply $\sigma_i = L_i/\sqrt{12}$ $(i = 1, 2)$, where L_i is the length of the i^{th} side of the rectangle.

3.5.3.1 Performance

The derived resolution function along the direction orthogonal to the strips is shown in Figure 3.30. Barrel and endcap detectors have been considered separately, using a sample of single muons distributed uniformly in η, ϕ, and in the momentum range $5 < p_{\mathrm{T}} < 100$ GeV/c. The shape of the distribution for the barrel detectors is essentially determined by the fact that the strips are rectangular with a width ranging from 2.1 cm to 4 cm. The resolution function for the endcap detectors is instead more peaked because of the trapezoidal shape of the strips having a minimal width comparable with overall space resolution of RPC.

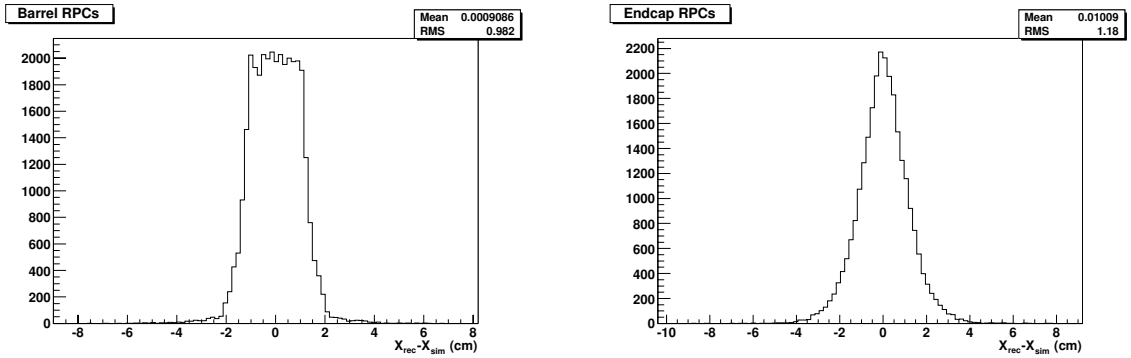


Figure 3.30: Distribution of residuals for the RPC measurement along the direction orthogonal to the strips. Barrel (left) and endcap (right) detectors have been considered separately using a sample of single muons flatly distributed in η, ϕ, and in the $5 < p_{\mathrm{T}} < 100$ GeV/c range [116].

3.5.4 Comparisons with data

3.5.4.1 DT test-beam results

In the 2004 DT test-beam set-up [107], 2 Muon Barrel chambers (one MB1 and one MB3) were placed at a distance of about 60 cm along the beam direction; two 5 cm thick steel slabs were placed between them, the first immediately after MB1 and the second just in front of MB3. This configuration reproduces well the situation in the CMS steel yoke, as far as the study of muon showering effects is concerned.

The plot on the left in Figure 3.31 displays the multiplicity of reconstructed hits in the MB3 chamber for a 300 GeV incident muon beam; the distributions for real and simulated data are normalized to the number of selected events having only 1 clean track and no shower activity

in MB1. It can be seen that the reconstructed-hit multiplicity is slightly underestimated in the Monte Carlo simulation, because of the soft δ-rays neglected in the digitization step of the simulation, as previously discussed in Section 3.3.2. On the other hand, the high-multiplicity component of the distribution due to muon showering is well reproduced by the simulation. Despite the small observed differences, the capability to reconstruct a muon track segment is well simulated, as can be seen in the plot on the right in Figure 3.31, which shows the number of reconstructed tracks in the r-ϕ view of the MB3 chamber having more than 5 reconstructed hits included in the fit. The plot is referred to an incident beam direction normal to the chamber. The 3% inefficiency of the track reconstruction is due to the aluminium I-beam separating the cells and is well reproduced by the simulation. It can be also seen that the Monte Carlo simulation slightly overestimates the efficiency to reconstruct additional track segments, coming from hard δ-rays passing through both ϕ SLs of the chamber.

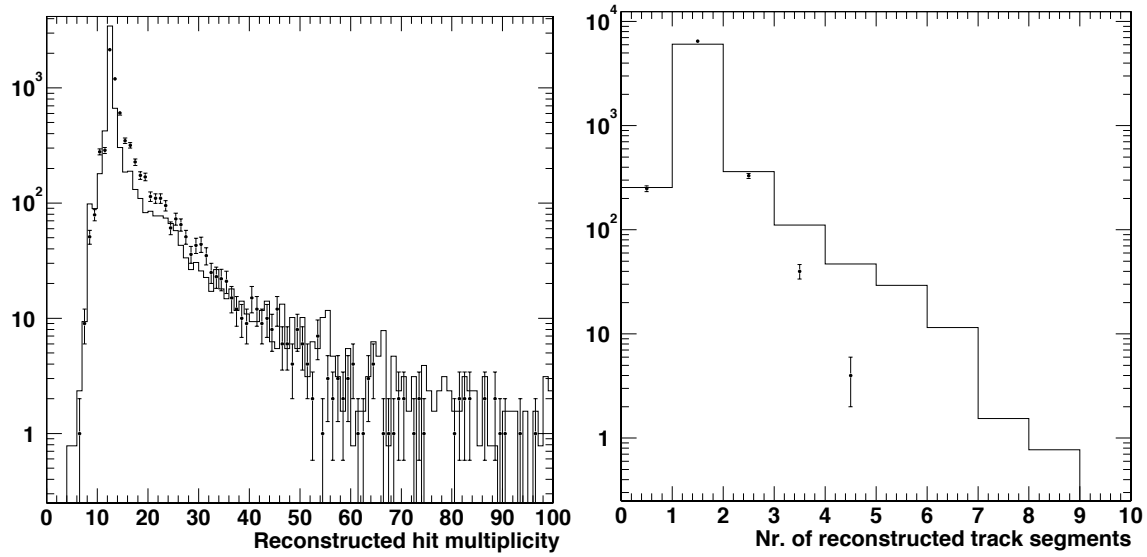


Figure 3.31: The plot on the left shows the reconstructed-hit multiplicity in MB3 chamber for testbeam data (dots) and simulation (full line). The plot on the right shows the number of reconstructed tracks in the r-ϕ view of an MB3 chamber for test-beam data (dots) and simulation (full line).

Figure 3.32 displays the reconstruction efficiency as a function of the track impact point in the chamber, for track segments in the r-ϕ view of the MB3 chamber placed in the test beam after the 2 steel slabs. The left (right) plot refers to a beam incident angle of 0° (10°) with respect to the direction normal to the chamber plane. The efficiency predicted by the simulation is also shown in the figure. In the left plot, the efficiency drop due to the I-beams (1.3 mm thick) separating the DT cells is clearly visible. The distribution of the angular difference between track segments measured in the 2 DT chambers is shown in Figure 3.33. The angular resolution is σ = 1.40 mrad and σ = 1.25 mrad in the real and simulated data, respectively.

3.5.5 RPC cosmic-ray tests

Experimental data were taken in the RPC cosmic-ray test areas located in Bari, Pavia, and Sofia [110]. They consist of towers where chambers are placed horizontally and read out in coincidence with the passage of muons triggered by a set of scintillator counters.

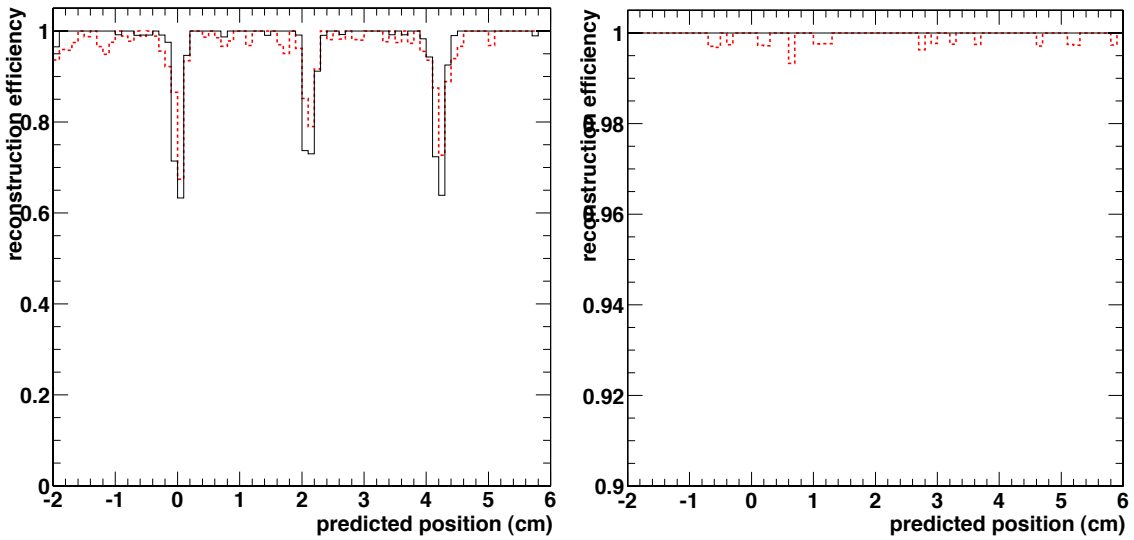


Figure 3.32: Track-segment reconstruction efficiency vs. predicted track position in MB3. Left: muon beam at normal incidence; right: beam direction inclined by 10° with respect to the normal direction: test-beam data (dotted line), simulation (full line).

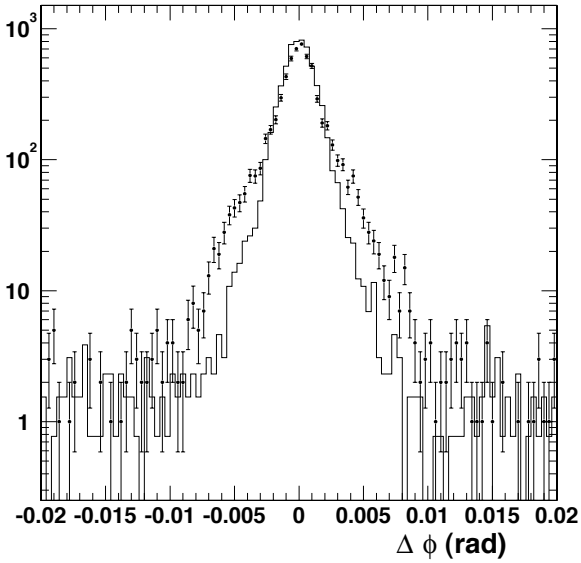


Figure 3.33: Angular difference between track segments measured in MB1 and MB3 chambers: test-beam data (dots), simulation (full line).

The clusterization is followed by muon reconstruction in a two-dimensional view with coordinates defined by the strip's position along the chamber and the chamber position in the tower. One chamber is chosen as a test chamber and not used in the reconstruction. Successfully reconstructed trajectories are required to have at least 1 used hit in 4 chambers. A typical reconstructed muon event is shown in Figure 3.34(a). After the reconstruction, the track impact point on the test chamber is determined. The residuals are computed as the difference between the impact point and the nearest cluster center. The distribution obtained for an RPC with 28 mm strip width is show in Figure 3.34(b). The root mean square of the

distribution represents the spatial resolution of the chamber, being the error on the impact point negligible. The spatial resolution value of 8.16 ± 0.07 mm is in agreement with the expected $28/\sqrt{12}$ mm.

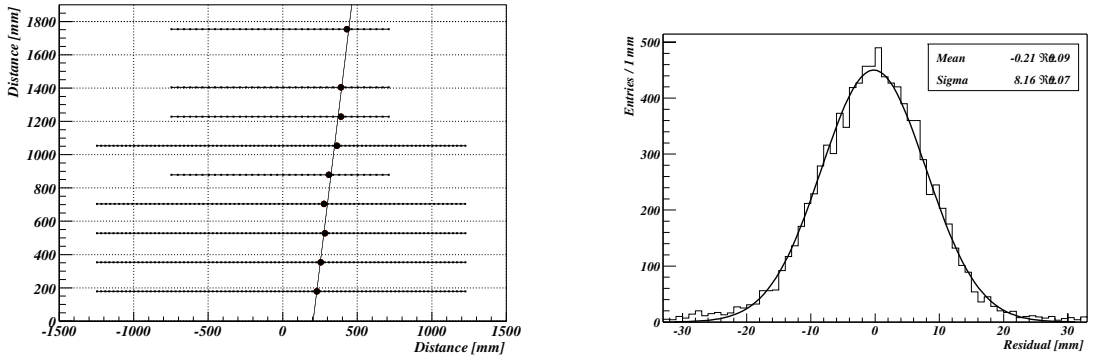


Figure 3.34: Left: Picture of a typical reconstructed muon track. Right: Distribution of residuals for RPC hits, with a fit to a Gaussian giving the average spatial resolution.

3.6 Commissioning

3.6.1 Detector start-up scenario

At startup of LHC, the full muon detector is expected to be operating, except for the innermost rings of the RPC endcap trigger, which will cover only up to $\eta = 1.6$.

During the detector commissioning stage and the CMS Magnet Test, extensive tests of the system will be performed using cosmic muons to check both the functionality of the subdetectors at different levels (individual chambers, sectors, barrel-endcap overlapping regions) and the alignment system. Knowledge of relative alignment of the sub-detectors is expected to be as described in the "first data" scenario quoted in Section 3.2.2.5 and it will be precise enough for safe Level-1 and HLT trigger operations.

3.6.2 Synchronization

The LHC bunch-crossing frequency of 40.08 MHz puts important requirements on the timing of signals. The CMS subdetectors produce data only after particles pass through them. In the case of the muon chambers, the times-of-flight are different enough between modules to cause a delay larger than 1 bunch crossing. Furthermore, most of the data communication will be done using long cables or optical fibres, which will produce delays among different parts of the detector. It is, therefore, necessary to develop electronics, methods, and procedures that permit the association of signals recorded in each channel of the detectors to the interaction that produced the detected particle to assure correct triggering and readout of the full CMS event.

Both trigger and readout channels need to be synchronized. The readout is generally asynchronous, except in the front-end electronics. There, to read out the correct data, the front-end electronics needs to be pre-configured with the precise number of bunch crossing between the storage of data in its internal data buffer and the arrival of the Level-1 Accept

signal. Additionally, the association of each recorded signal to a particular event is done synchronously by Level-1 Accept number. The Level-1 Accept number is crucial to the building of an event by the data acquisition system. After readout, in the high-level trigger and offline, the signals are additionally associated and verified by comparing the trigger bunch-crossing number. For readout purposes, it is enough to make sure that the front-end electronics records the Level-1 Accept and bunch crossing signals at the correct times. Because the trigger, on the other hand, is completely synchronous, it must be timed correctly at every step in the hardware, a more challenging task.

Trigger synchronization must be performed in several steps at different levels in all the subdetectors. The first step is to achieve reliable data transmission within the readout and trigger data systems. The second step is a time-equalization of all detector channels within a detector-defined basic block (normally a muon station). The third step is a relative synchronization between these basic blocks. Finally, an absolute synchronization to the LHC machine is required.

It should be mentioned that some synchronization procedures require iteration. For instance, a change in synchronization will, in general, affect the timing of the Level-1 trigger. That will in turn affect the time at which the Level-1 Accept signal returns to the front-end cards, whose timing must then be adjusted in order to read out the correct signals. This may then affect or even invalidate data used in previous synchronization steps. It is clear, therefore, that synchronization procedures have to be done very carefully to avoid these problems.

3.6.2.1 Synchronization for reliable data transmission

The absolute first step in synchronizing the detector is to achieve reliable readout and trigger data transmission between electronics boards. Most electronics boards use a clock at the LHC bunch-crossing frequency or a multiple of it that is derived from the TTC system. Some sophisticated systems such as high-speed optical links auto-synchronize and produce error bits in the rare cases when synchronization fails. In these cases, no action is generally needed except for monitoring the error bits. Some types of links synchronize automatically but with variable latency: such links are not used in the trigger system, which requires fixed latency. Most inter-board data connections use simpler methods to achieve synchronization, such as by adjusting the phase of a clock used to latch data in the receiving board until the data are known to be received correctly. In these cases, a mechanism for generating known data, such as pulsing of the front-end, is useful for verifying the correct synchronization phases. Cosmic-ray muon particles can also provide data for this.

3.6.2.2 Time-equalization of channels

Signals produced by the same muon will arrive at different front-end electronics with a jitter due to propagation time and muon time-of-flight. A time-equalization of the channels must be performed in order to guarantee that the signals are all assigned to the same clock period by the boards collecting the information from a basic detector electronics block (e.g. 1 chamber). Although different, all the developed methods rely on the signals broadcast by the TTC system through the on-board TTCrx chips.

The time-equalization of the RPC channels is done inside the Link Board by a dedicated Synchronization Unit module. The synchronization method is shown in Figure 3.35. A so-called WINDOW signal is derived from the TTCrx . The width of this signal is adjustable

in 1 ns steps and will be set large enough to latch all the signals associated with the same muon coming from the chambers connected to the Link Board. All the delayed signals are then output simultaneously synchronized with the TTCrx clock.

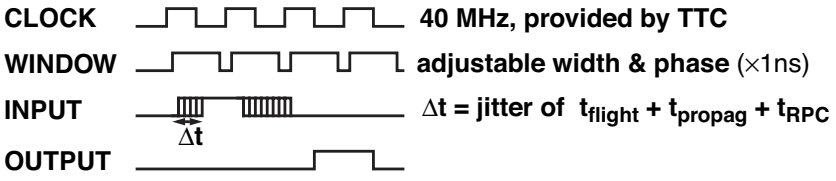


Figure 3.35: Illustration of the method to be used for time-equalization of the RPC channels in the Link Board.

In the CSCs, there are signals coming from charge collected on the anodes and other signals from charge induced on the cathodes. The anode electronics are used to achieve precise identification of the muon bunch crossing, while the cathode electronics are allowed to vary by ± 1 bunch crossing. Hence, precise time-equalization is not required for the cathode channels, while it is compulsory for the anode channels. The equalization is achieved inside the ALCT card, which performs the 4/6 coincidence of anode signals for trigger purposes, by time-delaying each set of 16 anode channels (2 layers of 8 wiregroups) by 0 to 30 ns adjustable in 2 ns steps. The relative timing between anodes and cathodes is adjusted by integer bunch crossings in the TMB board. Additionally, the bunch crossing numbers for the 3 boards used to fully read out a CSC are adjusted by offsets to make them equal, by offline examination of readout data from cosmic rays or LHC muons.

In both the RPC and CSC cases, since the time-of-flight depends only on geometry, and the signal propagation depends only on cable routing, the delays can be measured before data taking begins using test-pulse signals.

The case of the DT chambers is complicated by the very long drift times before the charge is collected. Furthermore, the efficiency of the trigger algorithm depends on the correct phase relation between the sampling clock and the machine clock. A single DT station is hardware timed-in by means of a precise cable distribution scheme to guarantee equal propagation delays for every channel. However, the trigger ASICs will sample these signals with an unknown phase relation with respect to the clock distributed by the TTCrx card serving each chamber. This phase depends on the drift velocity (i.e., magnetic field strength, gas composition, and HV setting), and its best value must be found in situ using real muons. The optimal value for the phase (different for each muon barrel station) must be determined by finding the TTCrx fine delay setting that gives the highest trigger efficiency. Several algorithms have been proposed to do this [117]. The results from one of them, optimizing the HL/HH ratio, are shown in Figure 3.36. An automated procedure to perform this algorithm was developed and tested during the 2004 test beam.

3.6.2.3 Relative synchronization

Once each basic detector block is time-equalized, the signals coming from different blocks will be synchronous with respect to the TTCrx clock. However, they may still not be aligned with each other, although the differences will be in integer numbers of clock periods. All the muon detectors provide a block-to-block synchronization by means of synchronization pipelines. The general approach to synchronize the complete system is to align the arrival of signals from chambers crossed by the same muon.

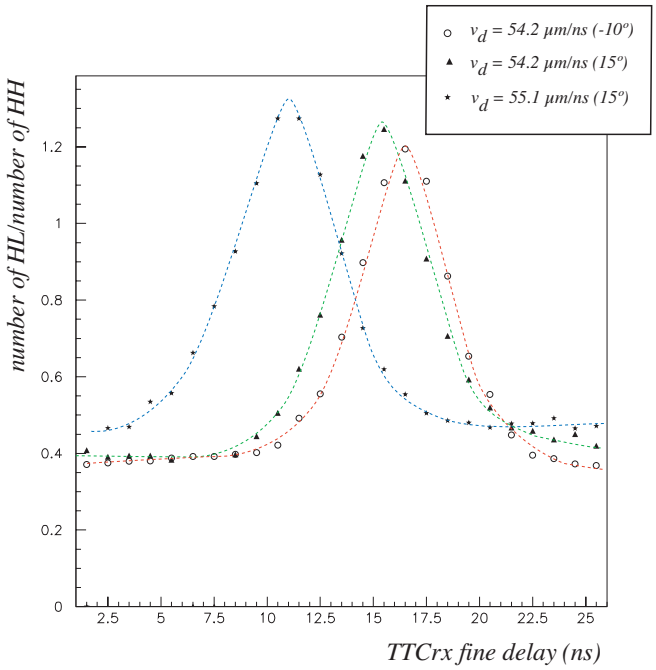


Figure 3.36: Dependence of the HH/HL ratio on the TTCrx fine time delay for 3 difference drift velocities and entrance angles. The optimal phase corresponds to the maximum of these distributions.

The RPC electronics provide a 3-clock-deep Synchronization Buffer on the Link Board. Therefore, the trigger data will be aligned at the PACT input. The same strategy is used in the DT chambers using the synchronization delays set at the Sector Collector input and, therefore, providing aligned signals to the DTTF. For the CSC system, there are various types of chambers having different cable delays sending trigger data through a single Muon Port Card in each Peripheral Crate (covering a trigger sector, i.e., 30° or 60° slice in ϕ in a single muon station). Their synchronization is done by delaying trigger signals using a pipeline in the Trigger Motherboards before transmission to the Muon Port Card. Moreover, there will be relative timing between CSC trigger sectors of signals arriving at the CSC Track Finder that can be adjusted by uniform additional delays of all the Trigger Motherboard pipelines in a Peripheral Crate. Cosmic rays have proven very useful for relative time alignment of the CSC trigger, although care has to be taken that the cosmic ray paths are similar to the expected paths of LHC collision muons.

3.6.2.4 Absolute synchronization

After the relative synchronization, the muon detectors will be self-aligned in time, but will not share the same time origin. Therefore, an absolute synchronization is needed. The absolute synchronization can only be performed by comparing the known LHC bunch structure (bunch numbers 0–3563, with certain empty bunches) with the structure of bunch-crossing assignments from each trigger branch, accumulated as a histogram during data taking. Example of histograms corresponding to correct, late, and early timings are shown in Figure 3.37. The comparison will provide to each detector a value for the coarse TTCrx delay that correctly sets the absolute trigger time.

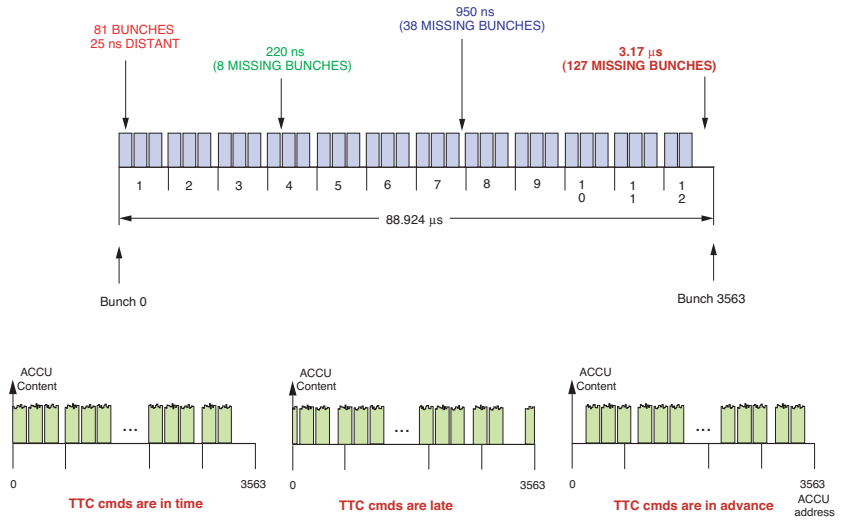


Figure 3.37: The LHC bunch structure (top). Examples of bunch-crossing histograms accumulated by a detector with correct, late, and early timing (bottom).

3.6.3 Muon triggers

Muon triggers are needed for cosmic-ray running, single-beam (halo muon), and LHC collisions. The first 2 kinds of trigger are very important for early detector commissioning, in particular for alignment studies. Beam-halo muon triggers from CSCs will be provided both at the commissioning stage and during normal collision running. An extensive study of the trigger capability and reconstruction performance using cosmic-muon and LHC-beam-halo simulations has been reported in [118], as will be further discussed in Section 9.1.4.

3.6.3.1 Cosmic-ray triggers

The CMS muon trigger system is not optimized as a cosmic-ray trigger due to the strict requirement of vertex pointing. However, such a trigger is needed in the early set-up phase of the CMS detector to study the detector performance. A cosmic-ray trigger can be obtained only as a technical trigger with low and varying efficiency around the detector. A partial exception to this situation is in the lower part of the barrel chambers where a pointing trigger can in fact be set up.

The cosmic-ray trigger will be accomplished by a time coincidence of 2 or more detector chambers, skipping the reconstruction of the track done in the PACT processor for the RPCs or the Regional Trigger Processors in the case of the DTs and CSCs. The low cosmic-ray rate (especially in the cavern) does not require any further algorithm, due to the power of the CMS data-processing farm.

A pointing cosmic-ray trigger can be set up in the lower part of the CMS muon barrel detector by making use of the fact that muon tracks are already fully reconstructed in both the bending and non-bending detector views (even if with much lower resolution and without transverse momentum determination with respect to the DTTF processor). The tracks in the non-bending view can be individually selected as pointing to the vertex with about a 3° precision. Tracks in the bending view will then be forwarded to the DTTF only if they fulfil this requirement. No special set-up will be needed in the DTTF, although the transverse momentum assignment will not be accurate.

3.6.3.2 Halo muons

Halo-muon triggers can be set up only in the forward detectors. The barrel detectors will be blind to halo muons, since they will pass through the entire sensitive zone of the RPC and DT chambers. This will create a large amount of charge to be deposited and, therefore, cause the electric field to decrease, with a resulting drop in chamber efficiency.

3.6.4 Momentum calibration

Reconstruction of mass resonances in their dimuon decay mode will be a fundamental tool to check the scale of the reconstructed muon momenta, which relies on knowledge of the magnetic field map.

Figure 3.38 shows the invariant mass distribution for reconstructed muon pairs as expected from 1 day of data taking at a LHC luminosity $L = 2 \times 10^{33} cm^{-2} s^{-1}$. The mass spectrum is obtained using the Global Muon Reconstruction in ORCA, as described in Section 9.1, after HLT and offline selection requiring 2 isolated, opposite-sign muons with $p_{\mathrm{T}} > 20 and 10$ GeV$/c$, respectively. The Z-mass peak is clearly visible over the background predicted from QCD events, which were included in the simulation (dashed line in the top plot).

As seen by the error of the fits shown in the plots, statistical errors of the order or smaller than 100 MeV on the mass scale can be achieved in the different regions of the detector in a relatively short time. Taking into account the Z natural width, the measured resolution is compatible with the momentum resolution obtained from the simulation by comparing the reconstructed and the generated muon tracks (see Chapter 9).

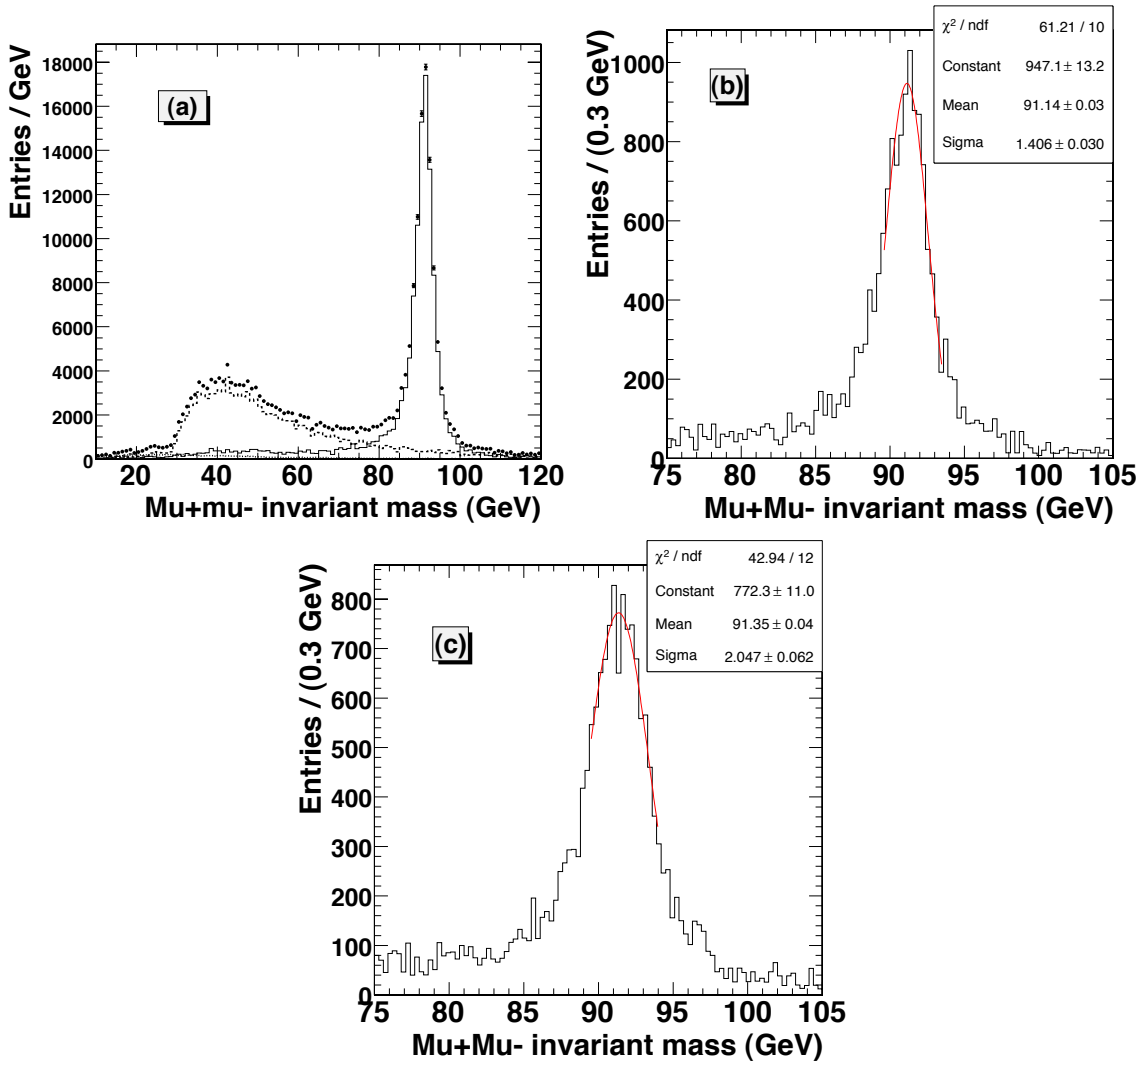


Figure 3.38: The reconstructed mass spectrum for 2 isolated muon tracks with $p_{\mathrm{T}} > 20 and 10$ GeV$/c$, respectively, as expected from 1 day of data taking at $\mathcal{L} = 2 \times 10^{33}$ cm^{-2} s^{-1}: a) all reconstructed dimuon events (solid line: Z signal, dashed line: QCD background); b) both muons in the "barrel region" ($\eta < 0.8$); c) 1 muon in the "overlap region" ($0.8 < \eta < 1.2$) and 1 muon in the "endcap region" ($\eta > 1.2$).

Chapter 4

Electromagnetic Calorimeter

4.1 Description of the ECAL

In this section, the layout, the crystals and the photodetectors of the Electromagnetic Calorimeter (ECAL) are described. The section ends with a description of the preshower detector which sits in front of the endcap crystals. Two important changes have occurred to the geometry and configuration since the ECAL TDR [5]. In the endcap the basic mechanical unit, the "supercrystal," which was originally envisaged to hold 6×6 crystals, is now a 5×5 unit. The lateral dimensions of the endcap crystals have been increased such that the supercrystal remains little changed in size. This choice took advantage of the crystal producer's ability to produce larger crystals, to reduce the channel count. Secondly, the option of a barrel preshower detector, envisaged for high-luminosity running only, has been dropped. This simplification allows more space to the tracker, but requires that the longitudinal vertices of $H \to \gamma\gamma$ events be found with the reconstructed charged particle tracks in the event.

4.1.1 The ECAL layout and geometry

The nominal geometry of the ECAL (the engineering specification) is simulated in detail in the GEANT4/OSCAR model. There are 36 identical supermodules, 18 in each half barrel, each covering 20° in ϕ. The barrel is closed at each end by an endcap. In front of most of the fiducial region of each endcap is a preshower device. Figure 4.1 shows a transverse section through ECAL.

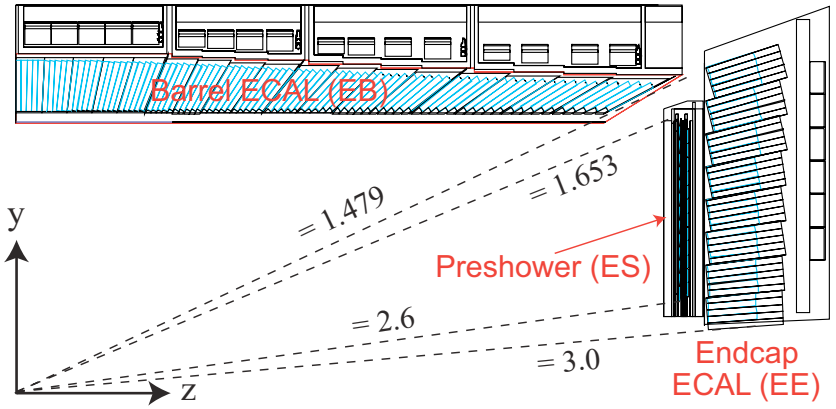


Figure 4.1: Transverse section through the ECAL, showing geometrical configuration.

The barrel part of the ECAL covers the pseudorapidity range $|\eta| < 1.479$. The barrel granularity is 360-fold in ϕ and (2×85)-fold in η, resulting in a total of 61 200 crystals.The truncated-pyramid shaped crystals are mounted in a quasi-projective geometry so that their axes make a small angle (3°) with the respect to the vector from the nominal interaction vertex, in both the ϕ and η projections. The crystal cross-section corresponds to approximately $0.0174 \times 0.0174°$ in η-ϕ or 22×22 mm^2 at the front face of crystal, and 26×26 mm^2 at the rear face. The crystal length is 230 mm corresponding to 25.8 X_0.

The centres of the front faces of the crystals in the supermodules are at a radius 1.29 m. The crystals are contained in a thin-walled glass-fibre alveola structures ("submodules," as shown in Fig. CP 5) with 5 pairs of crystals (left and right reflections of a single shape) per submodule. The η extent of the submodule corresponds to a trigger tower. To reduce the number of different type of crystals, the crystals in each submodule have the same shape. There are 17 pairs of shapes. The submodules are assembled into modules and there are 4 modules in each supermodule separated by aluminium webs. The arrangement of the 4 modules in a supermodule can be seen in the photograph shown in Fig. 4.2.

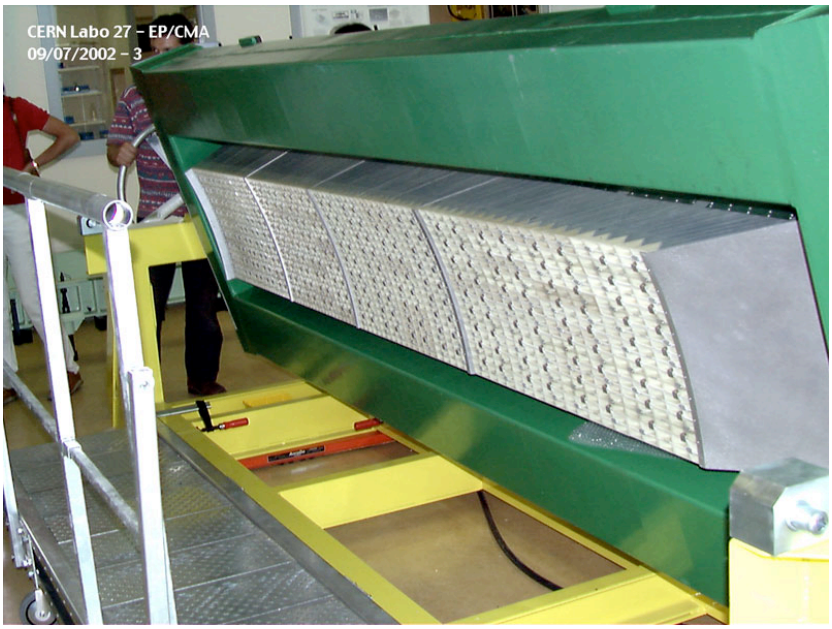


Figure 4.2: Photograph of supermodule, showing modules.

The thermal screen and neutron moderator in front of the crystals are described in the model, as well as an approximate modelling of the electronics, thermal regulation system and mechanical structure behind the crystals.

The endcaps cover the rapidity range $1.479 < |\eta| < 3.0$. The longitudinal distance between the interaction point and the endcap envelop is 3144 mm in the simulation. This location takes account of the estimated shift toward the interaction point by 2.6 cm when the 4 T magnetic field is switched on. The endcap consists of identically shaped crystals grouped in mechanical units of 5×5 crystals (supercrystals, or SCs) consisting of a carbon-fibre alveola structure. Each endcap is divided into 2 halves, or "Dees" (Fig. CP 6). Each Dee comprises 3662 crystals. These are contained in 138 standard SCs and 18 special partial supercrystals on the inner and outer circumference. The crystals and SCs are arranged in a rectangular

x-y grid, with the crystals pointing at a focus 1300 mm beyond the interaction point, so that the off-pointing angle varies with η. The crystals have a rear face cross section $30\times30\,\text{mm}^2$, a front face cross section $28.62\times28.62\,\text{mm}^2$ and length of 220 mm (24.7 X_0). The simulated endcap geometry includes the aluminum backplate and approximate modelling of the support structure and electronics behind the crystals.

4.1.2 Lead tungstate crystals

Lead Tungstate crystals ($PbWO_4$) are produced for CMS by the Bogoroditsk Techno-Chemical Plant in Russia and by the Shanghai Institute of Ceramics in China. The characteristics [119] of these production crystals make them an appropriate choice for operation at LHC. The high density ($8.3\,\text{g/cm}^3$), short radiation length (0.89 cm) and small Molière radius (2.2 cm) results in a fine granularity and a compact calorimeter. The scintillation decay time is of the same order of magnitude as the LHC bunch crossing time: about 80% of the light is emitted in 25 ns. The light output is relatively low: about 4.5 photoelectrons per MeV are collected in both the avalanche photodiodes (APDs) and the vacuum phototriodes (VPTs), where the higher APD quantum efficiency is balanced by their smaller surface coverage on the back face of the crystal. The crystals emit blue-green scintillation light with a broad maximum at 420 nm [120, 121]. The light output variation with temperature, -1.9% per °C at 18°C, requires an ECAL cooling system capable of extracting the heat dissipated by the readout electronics and of keeping the crystal temperature stable within ±0.05°C to preserve energy resolution. To exploit the total internal reflection for optimum light collection on the photodetector, the crystals are polished after machining. This is done on all but one side for EB crystals. For fully polished crystals, the truncated pyramidal shape makes the light collection non-uniform along the crystal length, and the needed uniformity [122] is achieved by depolishing one lateral face. In the EE, the light collection is naturally more uniform because the crystal geometry is nearly parallelepipedic, and just a mild tuning is being considered.

The crystals have to withstand the radiation levels and particle fluxes [5] anticipated throughout the duration of the experiment. Ionizing radiation produces absorption bands through the formation of colour centres due to oxygen vacancies and impurities in the lattice. The practical consequence is a wavelength-dependent loss of light transmission without changes to the scintillation mechanism, a damage which can be tracked and corrected for by monitoring the optical transparency with injected laser light. The damage reaches a dose-rate dependent equilibrium level which results from a balance between damage and recovery at 18°C [123, 120]. To ensure an adequate performance throughout LHC operation, the crystals are required to exhibit radiation hardness properties quantified as an induced light attenuation length always greater than 3 times the crystal length even when the damage is saturated. Hadrons have been measured to induce a specific, cumulative reduction of light transmission, but the extrapolation to LHC indicates that the damage will remain within limits required for good ECAL performance [124].

4.1.3 Photodetectors

4.1.3.1 Barrel: avalanche photodiodes

In the barrel, the photodetectors are Hamamatsu type S8148 reverse structure (i.e., with the bulk n-type silicon behind the p-n junction) avalanche photodiodes (APDs) specially developed for the CMS ECAL. Each APD has an active area of $5\times5\,\text{mm}^2$ and 2 are glued to the

back of each crystal. The APDs are sorted according to their operating voltage into bins 5 V wide, and then paired such that each pair has a mean gain of 50. The main properties of the APDs at gain 50 and 18°C are listed in Table 4.1.

Table 4.1: Properties of the APDs at gain 50 and 18°C.

Sensitive area	$5\times5\,\mathrm{mm}^2$
Operating voltage	340–430 V
Breakdown voltage - operating voltage	45 ± 5 V
Quantum efficiency (430nm)	$75 \pm 2\%$
Capacitance	80 ± 2 pF
Excess noise factor	2.1 ± 0.2
Effective thickness	6 ± 0.5 μm
Series resistance	< 10 Ohm
Voltage sensitivity of the gain (1/M.dM/dV)	$3.1 \pm 0.1\%$/V
Temperature sensitivity of the gain (1/M.dM/dT)	$-2.4 \pm 0.2\%/°$C
Rise time	< 2 ns
Dark current	< 50 nA
Typical dark current	3 nA
Dark current after 2×10^{12} n/cm^2	5 μA

The sensitivity to ionizing radiation traversing the APD (nuclear counter effect) is given by the effective thickness of 6 μm, which translates into a signal from a minimum ionizing particle traversing an APD equivalent to about 100 MeV deposited in the $PbWO_4$.

For acceptance for the ECAL each APD was required to be fully depleted and to pass through a screening procedure involving 5 kGy of ^{60}Co irradiation and 1 month of operation at 80°C. Each APD was tested to breakdown and required to show no significant noise increase up to a gain of 300. The screening and testing aimed to ensure reliable operation for 10 years under high luminosity LHC conditions for over 99% of the APDs installed in the ECAL [125]. Based on tests with hadron irradiations it is expected that the dark current after such operation will have risen to about 5 μA, but that no other properties will have changed. Small samples of APDs were irradiated with a ^{251}Cf source to monitor the effectiveness of the screening procedure selection of radiation resistant APDs.

4.1.3.2 Endcap: vacuum phototriodes

In the endcaps, the photodetectors are vacuum phototriodes (VPTs) (type PMT188 from National Research Institute Electron in St. Petersburg). Vacuum phototriodes are photomultipliers having a single gain stage. These particular devices were developed specially for CMS and have an anode of very fine copper mesh (10 μm pitch) allowing them to operate in the 4 T magnetic field. Each VPT is 25 mm in diameter; one VPT is glued to the back of each crystal. The VPTs delivered to date have a mean quantum efficiency of the bialkali photocathode (SbKCs) of 22% at 430 nm, and a mean gain of 10.2 at 0 T.

When placed in a strong axial magnetic field, the response is slightly reduced and there is a modest variation of response with the angle of the VPT axis with respect to the field over the range of angles relevant to the CMS endcaps (6° to 26°). The mean response in a magnetic field of 4 T, with the VPT axis at 15° to the field direction, is 94.5% of that in zero field.

All VPTs are tested by the manufacturer before delivery, without an applied magnetic field. All VPTs are also tested on receipt by CMS to determine their response as a function of magnetic field up to 1.8 T. Each device is measured at a set of angles with respect to the applied field, spanning the range of angles covered by the endcaps. In addition, at least 10% of the tubes, selected at random, are also tested in a 4 T superconducting magnet, at a fixed angle of 15°, to verify satisfactory operation at the full field of CMS.

The estimated doses and particle fluences for 10 years of LHC operation are 0.5 kGy and 5×10^{13} n/cm^2 at the outer circumference of the endcaps and 20.0 kGy and 7×10^{14} n/cm^2 at $|\eta|$ = 2.6. The VPTs are expected to be insensitive to such neutron fluences. The VPTs are required to show a loss of anode response of no more than 10% after a dose of 20.0 kGy.

4.1.4 Preshower detector

The principal aim of the CMS Preshower detector (ES) is to identify neutral pions in the endcaps within a fiducial region $1.653 < |\eta| < 2.6$. It also helps the identification of electrons against minimum ionizing particles, and improves the position determination of electrons and photons with it superior granularity. The ES is a sampling calorimeter with 2 layers: lead radiators initiate electromagnetic showers from incoming photons/electrons whilst silicon strip sensors placed after each radiator measure the energy deposited and the transverse shower profiles.

The material thickness of the ES traversed at $\eta = 1.653$ before reaching the first sensor plane is 2 X_0, followed by a further 1 X_0 before reaching the second plane. Thus about 95% of single incident photons start showering before the second sensor plane. The orientation of the strips in the 2 planes is orthogonal. A major design consideration is that all lead is covered by silicon sensors, including the effects of shower spread, primary vertex spread etc. For optimum Level-1 trigger performance the profile of the outer edge of the lead should follow the shape of the ECAL crystals behind it. For the inner radius the effect of the exact profiling of the lead is far less critical, and thus a circular shape has been chosen. The lead planes are formed from 2 Dees that join close to the vertical axis.

Each silicon sensor measures 63×63 mm^2, with an active area of 61×61 mm^2 divided into 32 strips (1.9 mm pitch). The nominal thickness of the silicon is 320 μm; a minimum ionizing particle (MIP) will deposit around 3.6 fC of charge in this thickness (normal incidence). The sensors are precisely glued to ceramic supports, which also support the front-end electronics assembly, and this is in turn glued to an aluminium tile that allows a 2 mm overlap of the active part of the sensors in the direction parallel to the strips.

The micromodules are placed on baseplates in groups of 7, 8 or 10 that, when coupled with an electronics motherboard placed above the micromodules, form a ladder. The spacing between silicon strips (at the edges) in adjacent micromodules within a ladder is 2.4 mm, whilst the spacing between strips in adjacent ladders is normally 2.5 mm; for the region where the 2 Dees join this spacing is increased to 3.0 mm.

The ladders are attached to the radiators in an x-y configuration. Around 500 ladders are required, corresponding to a total of around 4300 micromodules and 137 000 individual readout channels. Further details of the layout can be found in [126].

4.2 Readout electronics

This section describes the readout of the ECAL crystals and preshower detector. It also describes the generation and readout of the trigger data, the trigger primitives, passed to the Level-1 trigger; and the selection, by the Selective Readout Processor, of the channels for which the full precision data will be sent to the DAQ event builder. Finally the issue of the synchronization of the complete system is discussed.

The CMS-ECAL electronics can be divided into 2 subsystems. The on-detector electronics, composed of radiation-resistant circuits located just behind the crystals, and the off-detector electronics housed in underground counting rooms close to the experimental area. Both systems communicate through 90-m-long high-speed optical links, operated at 800 Mb/s.

4.2.1 Front end

The front end electronics of the ECAL must first amplify and shape the signal from the sensors, digitize the signal at 40 MHz, buffer the data until receipt of a Level-1 trigger, and then transmit the data to the off-detector electronics for insertion in the CMS data stream. In addition, the front end electronics uses the digitized data to calculate trigger primitives which are transmitted at 40 MHz to be used in the Level-1 Trigger decision.

The basic building block of the front end electronics is a group of 25 crystals (grouped in a 5×5 geometry) a trigger tower in EB or supercrystal in EE. The trigger towers are composed of 4 different electronics boards. Each trigger tower contains a motherboard (MB), a Low Voltage Regulator Board (LVRB), 5 Very Front End (VFE) boards, and a Front End (FE) card.

The motherboards are located beneath the cooling system for the electronics and are used to route the signals from the photodetectors, APDs in EB or VPTs in EE, to the VFE cards, to distribute high voltage to the photodetectors, and to distribute Low Voltage to the VFE cards. Each supermodule contains 68 motherboards which each connect to the photodetectors of 25 crystals via kapton flexible-print cables. In addition, signals for temperature monitoring thermistors are routed from the sensor capsule to the VFE cards.

The LVRBs are connected directly to the external Low Voltage power supplies which sit in the CMS racks attached to the outside of the CMS iron yoke, approximately 20 metres from the supermodule. Each LVRB contains radiation-hard voltage regulators which provide the 2.5 V needed by the front end electronics. This regulated 2.5 V is distributed to the FE card by a small connector on the LVRB, and to the 5 VFE cards in a trigger tower via the motherboard.

Each VFE card contains amplification and digitization for the signals from 5 crystals. In order to achieve the low noise and high dynamic range requirements for the ECAL, 2 new radiation-hard ASICs were developed. The Multi Gain Pre-Amplifier (MGPA) contains 3 parallel gain stages which process the sensor signals. The signals from these 3 stages are routed to the AD41240, a custom designed Analog to Digital Converter which contains 4 channels each with 12 bits of information and an effective number of bits equal to 11. Three of the 4 ADCs on the AD41240 are used for each crystal. The ASIC digitizes these 3 inputs in parallel, determines whether each channel has saturated, and then outputs the data from the channel which has the highest gain, and was not saturated. The 3 MGPA gains are arranged so that the highest gain range has an amplification of a factor of 12 and a least significant bit of $\simeq$35 MeV. This range has a noise of around 40 MeV for the barrel and saturates at approximately 160 GeV. The other 2 gains are a factor of 6 and unity. The unity gain determines the

upper end of the dynamic range and saturates at approximately 1.7 TeV for the barrel, and 3.0 TeV for the endcap. The ADC is also designed so that once a range saturates, the ADC returns the next 5 samples without changing the gain. This prevents a second (return) gain change within the waveform of the digitized sample. The same chip is used for both the APDs and the VPTs with only a small change of external components on the VFE.

The signals from the 5 VFEs are collected on the FE card. Here they are buffered in a custom ASIC designed for the ECAL front end – the FENIX ASIC. Each FE card contains 7 FENIX ASICs. These are multipurpose and are used to

1. buffer the sampled data from the ADCs until a Level-1 trigger is received, and then transmit the data the ECAL off detector electronics. The amount of data which is transferred is configurable, but typically consists of the 10 ADC samples surrounding the beam crossing,
2. sum the samples from a group of 5 channels (called a strip) at 40 MHz,
3. sum the samples from the 5 strips (all 25 channels) for transmission to the calorimeter trigger.

The data are transmitted to the Level-1 trigger and the DAQ system using 2 opto-hybrids on each FE card. These hybrids contain radiation-tolerant laser diodes for electrical to optical conversion, and the CERN-developed radiation-hard GOL ASIC which provides parallel to serial conversion. The FE card also contains the clock distribution, and the control unit for allowing local configuration of all ASICs via I2C protocol.

Clock signals and configuration are distributed to the 68 FE cards using 8 independent electrical (i.e. the interconnect is not optical) token rings. These rings communicate optically to the off-detector electronics using a digital opto-hybrid (DOH) developed originally for the CMS tracker project. The DOHs are mounted on token ring link boards which then connect to the token rings. There are 2 rings which run in parallel and are designed to allow the recovery of the ring, in the event that one of the FE cards fails.

A fast test pulse, clocked by the ADC clock, can be injected into the input of each preamplifier by the MGPA. A simple 8 bit DAC allows injection of a sufficient range of charges to verify the functionality of the 3 gain ranges.

Fuller details of the front end electronics can be found in Refs. [127, 128, 129, 130].

4.2.2 Off-detector electronics

The ECAL off-detector readout and trigger architecture [131, 132] is illustrated schematically in Fig. 4.3. The system is composed of different electronic boards sitting in 18 VME-9U crates (the CCS, TCC and DCC modules) and in 1 VME-6U crate (the selective readout processor, SRP, system). The system serves both the DAQ and the trigger paths. In the DAQ path, the DCC performs data readout and data reduction based on the selective readout flags computed by the SRP system. In the trigger path, at each bunch crossing, trigger primitive generation started in the FE boards are finalized and synchronized in the TCC before transmission to the regional calorimeter trigger.

The clock and control system (CCS) board distributes the system clock, trigger and broadcast commands, configures the FE electronics and provides an interface to the trigger throttling

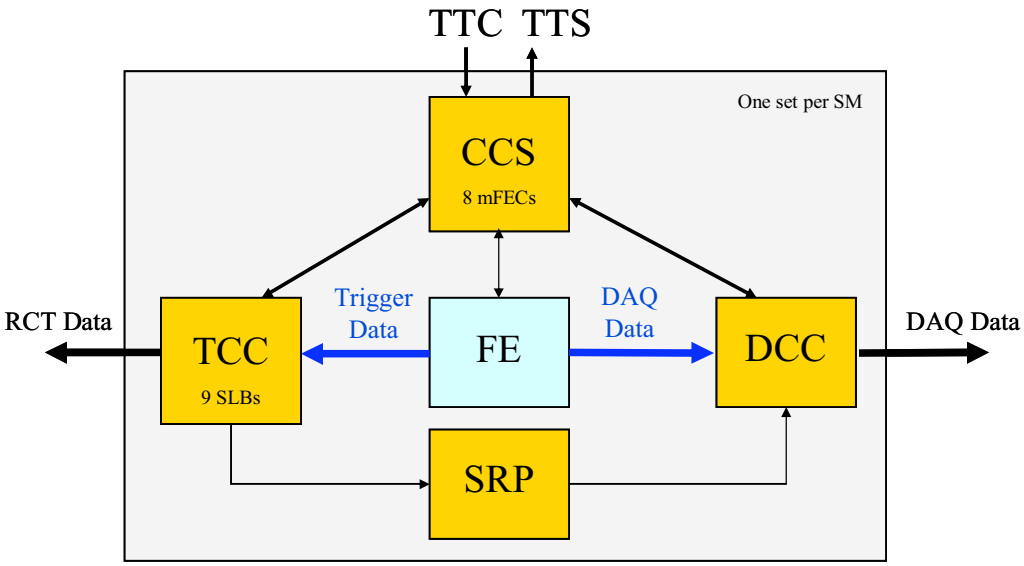


Figure 4.3: Schematic view of ECAL off-detector electronics.

system. The TTC signals are translated and encoded by suppression of clock edges and sent to the mFEC mezzanine cards. The mFEC interfaces optically with a FE token ring. The 8 mFECs of the CCS board control a supermodule. The TCC and the DCC cards in the off-detector crate receive the encoded TTC signals from the CCS card through a dedicated backplane.

The trigger concentration card (TCC) [133] main functionalities include the completion of the trigger primitive generation and their transmission to the synchronization and link board (SLB) mezzanines [134] at each bunch crossing, the classification of each trigger tower and its transmission to the Selective Readout Processor at each Level-1 trigger accept signal, and the storage of the trigger primitives during the Level-1 latency for subsequent reading by the DCC.

Each TCC collects trigger data from 68 FE boards in the barrel, corresponding to a supermodule, and from 48 FE boards in the endcaps corresponding to the inner or outer part of a 20° sector. In the endcaps, trigger primitive computation is completed in the TCCs, which must perform a mapping between the collected pseudo-strips trigger data from the different supercrystals and the associated trigger towers. The encoded trigger primitives (8 bits for the nonlinear representation of the trigger tower E_{T} plus 1 bit for the fine grain veto) are time aligned and sent to the regional trigger processors by the SLB. The trigger primitives are stored in the TCC during the Level-1 latency for subsequent reading by the DCC. In the barrel region a single TCC is interfaced with 1 DCC. In the endcap region, a DCC serves 4 TCCs covering a 40° sector

The data concentration card (DCC) [135, 136] is responsible for collecting crystal data from up to 68 FE boards. Two extra FE links are dedicated to the readout of laser monitoring data (pin diodes). The DCC also collects trigger data transmitted from the TCC modules and the selective readout flags transmitted from the SRP system. A data suppression factor near 20 is attained using a programmable selective readout algorithms. When operating in the selective readout mode the SRP flags indicate the level of suppression that must be

applied to the crystal data of a given FE readout. For the application of zero suppression, time samples pass through a finite impulse response filter with 6 consecutive positions and the result is compared to a threshold. If any time sample of the 6 has been digitized at a gain other than the maximum, then zero suppression is not applied to the channel.

Data integrity is checked, including verification of the event-fragment header, in particular the data synchronization check, verification of the event-fragment word count and verification of the event-fragment parity bits. Identified error conditions, triggered by input event-fragment checks, link errors, data timeouts or buffer memory overflows are flagged in the DCC error registers and incremented in associated error counters. Error conditions are flagged in the DCC event header.

Input and output memory occupancy is monitored to prevent buffer overflows. If a first occupancy level is reached, the Trigger Throttling System (TTS) signal "Warning Overflow" is issued, requesting a reduction of the trigger rate. In a second level a TTS signal "Busy" inhibits new triggers and empty events (events with just the header words and trailer) are stored. DCC events are transmitted to the central CMS DAQ using the S-LINK64 [41, 42] at a maximum data rate of 528 MB/s, while an average transmission data flow of 200 MB/s is expected after ECAL data reduction. Laser triggers (for crystal transparency monitoring) will occur with a programmable frequency and synchronously with the LHC gap. No data reduction is applied for these events, which are readout following a TTC test enable command. A VME memory is used for local DAQ, allowing VME access to physics events and laser events in spy mode.

The selective readout processor (SRP) [137] is responsible for the implementation of the selective readout algorithm. The system is composed by a single 6U-VME crate with twelve identical algorithm boards (AB). The AB computes the selective readout flags in different calorimeter partitions. The flags are composed by 3 bits, indicating the suppression level that must be applied to the associated readout units.

4.2.3 Trigger and readout

The ECAL data, in the form of trigger primitives, are sent to the Level-1 calorimeter trigger processor, for each bunch crossing. The trigger primitives each refer to a single trigger tower and consist of the summed transverse energy deposited in the tower, and a compactness bit, which characterizes the lateral extension of the electromagnetic shower. The accept signal, for accepted events, is returned from the global trigger in about 3 μs. The selected events are read out through the data acquisition system to the Filter Farm where further rate reduction is performed using the full detector data.

The readout system is structured into sets of 5×5 crystals. The FE card stores the data, in 256-clock cycles deep memory banks, awaiting a Level-1 trigger decision during at most 128 bunch crossings after the collision occurred. It implements most of the Trigger Primitives Generation (TPG) pipeline (Section 4.2.4). In the barrel, each FE is served by 3 optical links: 2 dedicated fibres for sending the data and trigger primitives respectively, and a third link which transmits the clock, control and Level-1 trigger signals.

In the barrel, these 5×5 crystal sets correspond to the trigger towers. Each trigger tower is divided into 5 ϕ-oriented strips, whose energy deposits are summed by the FE board trigger pipeline to give the total transverse energy of the tower, called the main trigger primitive.

In the endcaps, the readout modularity maps onto the 5×5 mechanical units (supercrystals). However the sizes of the trigger towers vary in order to approximately follow the η, ϕ geometry of the HCAL and Level-1 trigger processor. The supercrystals are divided into groups of 5 contiguous crystal. These groups are of variable shape and referred to as pseudo-strips. The trigger towers are composed of several pseudo-strips and may extend over more than one supercrystal. Since the readout structure does not match the trigger structure, only the pseudo-strip summations are performed on the detector. The total transverse energy of the trigger tower is computed by the off-detector electronics. Hence, each endcap FE board is served by 7 optical links, 5 of them being used to transmit the trigger primitives.

After time alignment (Section 4.2.7) the ECAL trigger primitives are sent to the regional calorimeter trigger, via 10-m-long 1.2 Gb/s electrical cables, where together with HCAL trigger primitives, the electron/photon and jets candidates are computed as well as the total transverse energy.

4.2.4 Trigger primitive generation

The TPG logic implemented on the FE boards combines the digitized samples delivered by the VFE boards to determine the trigger primitives and the bunch crossing to which they should be assigned. The logic must reconstruct the signal amplitude to be assigned to each bunch-crossing from the continuous stream of successive digitizations.

The TPG logic is implemented as a pipeline, operated at the LHC bunch crossing frequency. The trigger primitives are delivered to the regional calorimeter trigger after a constant latency of 52 clock cycles, of which 22 are used for the transmission over the optical fibres and cables. The signal processing performed in the VFE and FE barrel electronics has a total duration of only 17 clock cycles. The remaining part of the latency is mainly due to formatting and time alignment of the digital signals. Ideally, the output of this processing should be a stream of zeroes, unless there is a signal in the tower resulting from a bunch crossing exactly 17 clock cycles before. In this case the output is a word encoding the summed transverse energy in the tower together with the compactness bit. The endcap pipeline is split between the on-detector and off-detector electronics and implements very similar algorithms. The trigger primitives are expected to be delivered to the regional calorimeter trigger in 50 clock cycles in the endcap case.

The trigger primitive generation was studied using data taken in a test beam in November 2004. Electron beams with energies ranging from 15 to 120 GeV were used. In the test-beam experiment, electrons hit the detector at random times with respect to the 40 MHz clock used for sampling the analogue signals. Moreover, the parameters loaded in the chips (intercalibration constants and amplitude filter coefficients) were not optimized and only a limited amount of data were recorded. In order to assess the quality of the trigger primitives in the full range of energy and in LHC-like conditions, a functional model of the trigger primitives electronics and a Monte-Carlo simulation, based on GEANT4, has been developed to reproduce the online processing of electromagnetic showers. The actual noise characteristics of the VFE electronics were introduced by superimposing onto the simulated digital responses the recorded electronics responses of channels when the beam was off. The full simulation chain was tuned on the recorded data events in order to obtain a perfect matching of the simulation results with respect to the trigger primitives recorded online.

The response of the trigger primitives generation was first studied in 2 different geomet-

rical configurations, namely with the electron impact in the centre and in the corner of a trigger tower. For both configurations, the TPG shows good linearity versus the energy of the impinging electrons. Using the full simulation, the resolution expected for the ECAL electronics chain processing the trigger primitives was estimated. Figure 4.4 shows the relative resolution of the total transverse energy of an electromagnetic shower at the level of the ECAL trigger primitive generation. Two sets of points corresponding to the 2 geometrical configurations are presented in this figure. There is very little difference in the ECAL trigger primitives generated in these 2 extreme configurations.

Figure 4.5 shows the bunch crossing identification efficiency (BCID) versus the total transverse energy of the electromagnetic shower. The BCID for electromagnetic showers with energy $\geq 1\,\text{GeV}$ is almost 100%. These characteristics will enable the time alignment of the trigger primitives (Section 4.2.7) to be performed without difficulty.

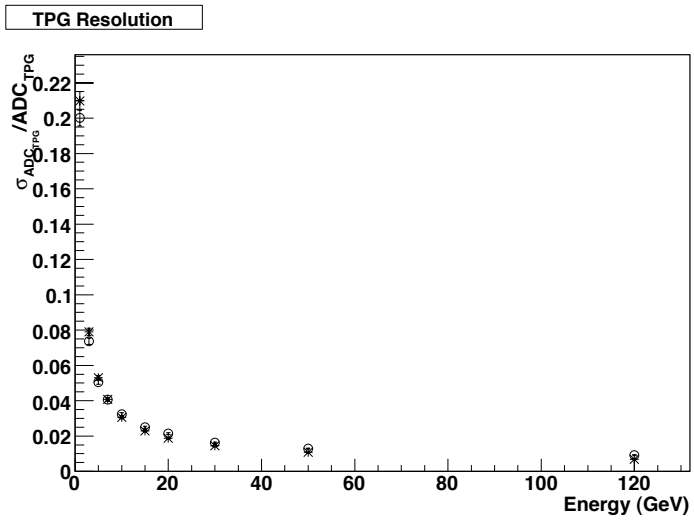


Figure 4.4: Trigger primitive resolution function (Monte Carlo simulation tuned using experimental data). Star markers refer to electrons impacting the centre of a trigger tower, while the circle markers refer to impacts at the corner of a trigger tower.

4.2.5 Selective readout

About $100\,\text{kB}$ per event has been allocated for ECAL data. The full ECAL data for an event, if all channels are read out, exceeds this target by a factor of nearly 20. Reduction of the data volume, "selective readout", can be performed by the Selective Readout Processor [132, 137] so that the suppression applied to a channel takes account of energy deposits in the vicinity. For the measure of the energy in a region, the trigger tower sums are used. In the barrel the readout modularity corresponds exactly to the 5×5-crystal trigger towers. In the endcap, the situation is more complex. The simplified and illustrative description below is given for the barrel case.

The selective readout algorithm classifies the trigger towers of the ECAL into 3 classes using the Level-1 trigger primitives. The energy deposited in each trigger tower is compared to 2

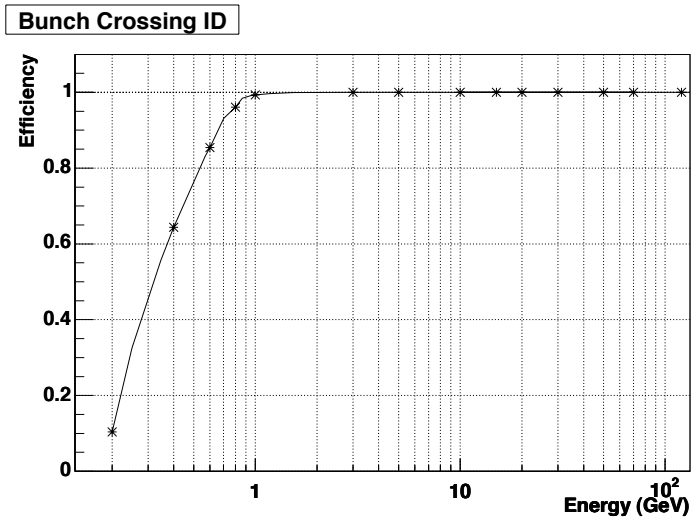


Figure 4.5: Bunch crossing assignment efficiency (Monte Carlo simulation tuned using experimental data).

thresholds. Trigger towers with an energy above the higher threshold are classified as high interest trigger towers, those with an energy between the 2 thresholds as medium interest, and those with an energy below the lower threshold as low interest trigger towers.

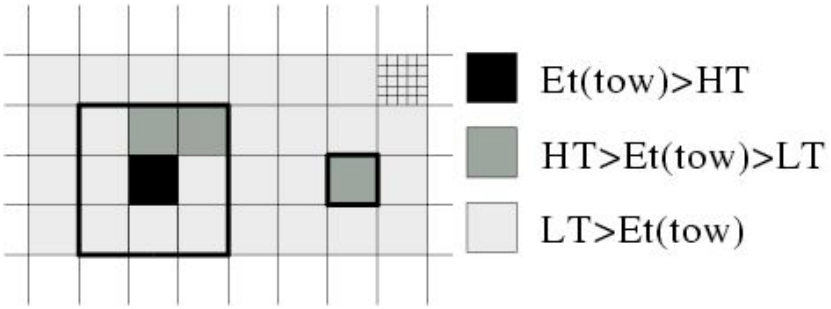


Figure 4.6: Selective readout regions. The figure illustrates the case of one trigger tower with a high transverse energy deposit (above the higher threshold HT, in black): the crystals of the 3×3 trigger-tower matrix around this trigger tower are read without zero suppression threshold. The crystals of the trigger tower with a medium transverse energy deposit (between the higher and lower thresholds, HT and LT, in dark grey) are also read without zero suppression. The remaining towers, shown in grey, are read out with zero suppression

These classifications, illustrated in Fig. 4.6, can be used flexibly to implement a range of algorithms by using different thresholds to define the classes, and different suppression levels for the readout of the channels within each class. The algorithm currently used in the simulation provides adequate data reduction even at high luminosity. The algorithm functions as follows: if a trigger tower belongs to the high interest class ($E_{\mathrm{T}} > 5$ GeV) then the crystals of this trigger tower and of its neighbour trigger towers (225 crystals in the barrel case) are read with no zero suppression. If a trigger tower belongs to the medium interest

class ($E_{\mathrm{T}} > 2.5$ GeV), then the crystals of this trigger tower (25 crystals in the barrel case) are read with no suppression. If a trigger tower belongs to the low interest class and it is not the neighbour of a high interest trigger tower, then the crystals in it are read with zero suppression at about $3\sigma_{\mathrm{noise}}$.

For debugging purpose, the selective readout can be deactivated and either a global zero suppression (same threshold for every channel) or no zero suppression applied. Even when the selective readout is not applied the selective readout flags are inserted into the data stream and can be used offline for debugging purpose.

4.2.6 Preshower electronics

The preshower electronics can best be described in 2 parts: on-detector and off-detector, both of which are described in more detail in [138]. Figure 4.7 shows a schematic view of the complete electronics chain with this division marked.

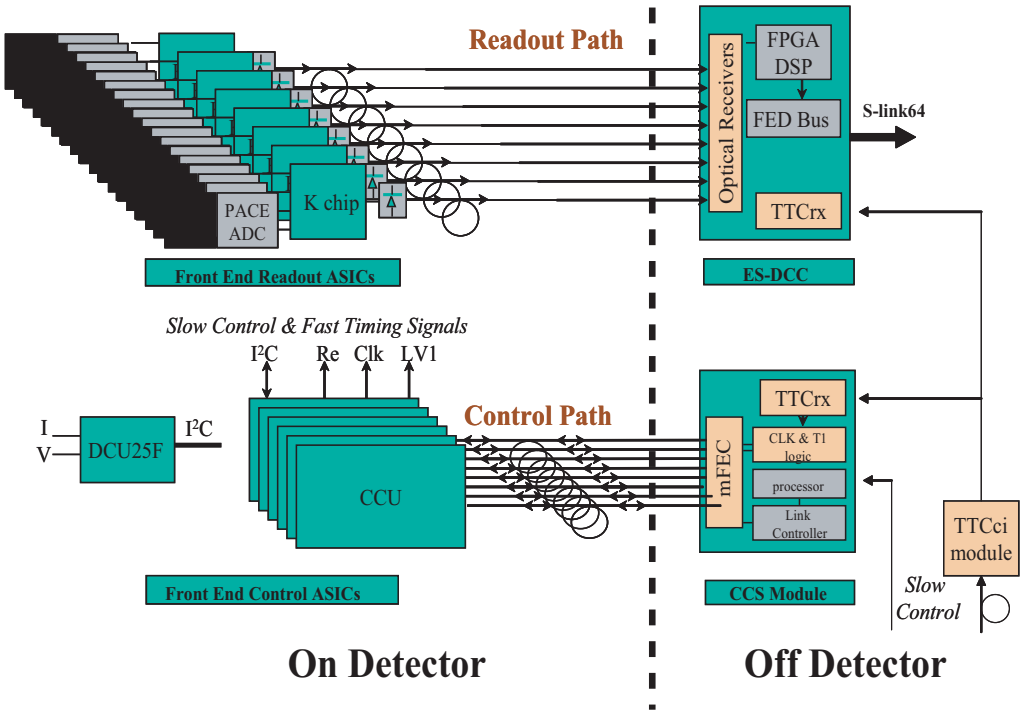


Figure 4.7: Preshower electronics chain.

4.2.6.1 On-detector electronics

Each of the approximately 4300 32-channel silicon sensors in the Preshower detector is DC-coupled to a front-end ASIC (PACE3 [139]) that performs preamplification, signal shaping and voltage sampling into a high dynamic range 192-cell deep analogue memory at 40 MHz. For each Level-1 trigger received, 3 consecutive cells of the memory, corresponding to samples on the baseline, near the peak and after the peak, are read-out for all 32 channels through a 20 MHz multiplexer. The PACE3 has a switchable gain:

- Low gain: For normal physics running with a high dynamic range (0–1600 fC[1]) with a S/N of around 3 for a single MIP,
- High gain: For calibration purposes, with a reduced dynamic range (0–200 fC) but with a S/N approaching 10 for a single MIP.

The PACE3 has many programmable features, including an internal injection pulse that is used to intercalibrate the 2 gains, programmable biases and currents through 8-bit DACs, variable latency, etc. The DACs can be calibrated via multiplexed analogue outputs to DCU25F chips mounted on the same front-end hybrids as the PACE3. The sensor and hybrid are mounted on a ceramic support and an aluminium "tile" that allows overlapping of the sensors in one dimension, to form a "micromodule". The micromodules are then mounted on the absorber plates. To avoid DC-coupling (and associated noise problems) of the sensor/electronics to the absorber planes, the aluminium tiles are constructed in 2 parts, with a glass-fibre insulation layer between.

Groups of 7, 8 or 10 micromodules are connected, via polyimide cables embedded in the front-end hybrids, to "system motherboards" (SMBs). The SMBs contain AD41240 12-bit ADCs for digitizing the analogue data from the PACE (1 ADC is used for 1 or 2 PACE3). The digital data from up to 4 PACE3 are then formatted and packaged by a second Preshower ASIC called the K-chip [140]. The K-chip also performs synchronization checks on the data, adds bunch/event counter information to the data packets, and transmits the data to the off-detector VME electronics via gigabit optical hybrids (GOH). The SMB also contains an implementation of the CMS Tracker control system, based on the CCU25 chip and its associated PLL25, QPLL, LVDSmux4p and LVDSbuf. The CCU25 provides a method to communicate to the K-chips and PACE3s via the I^2C protocol, to program registers, etc. The combination of micromodules and an SMB, along with heatsinks and an aluminium baseplate, forms a "ladder".

Groups of up to 12 ladders are connected via polyimide cables to form "control rings". Each of these control rings communicates to the off-detector board "Clock and Control System" (CCS) modules via digital opto-hybrids (DOH) mounted on 2 of the SMBs in the control ring (for redundancy purposes). Each Preshower plane contains 12 control rings. There are thus 48 control rings for the complete Preshower system.

4.2.6.2 Off-Detector electronics

The Preshower does not contribute to the Level-1 trigger. Consequently there are only 2 main parts to the off-detector electronics: the CCS module and the preshower data concentrator card - ES-DCC.

The CCS module is identical to that used by the EB and EE, except that for the ES only 3 or 4 of the 8 FEC mezzanines are mounted. One CCS module communicates with 3 or 4 control rings.

The ES-DCC reads the raw data from multiple GOH and performs bunch-crossing assignment, pedestal subtraction, common-mode rejection, charge reconstruction and zero suppression (threshold application), before formatting and sending sparsified data to the central DAQ via the S-LINK64 interface.

[1] A single minimum ionizing particle deposits around 3.6 fC in the 320 μm silicon sensors.

A "spy mode" will also be implemented, allowing full (non-zero-suppressed) data to be transferred to a local DAQ via the VME backplane or via the S-LINK64 (special condition, not during normal physics running). This will be used for startup conditions, pedestal runs, etc.

The design of the ES-DCC is ongoing, but will be largely based upon the ECAL DCC. The number of optical inputs may be reduced for the ES due to the larger amount of internal processing required in comparison to the EB and EE. The number of Slink64 conections reserved for the ES is 56, setting an upper limit to the number of ES-DCC.

4.2.7 Synchronization

In this section the procedures of synchronization of individual readout channels, as well as of the trigger primitives are described. A further issue related to timing, a procedure to verify the constancy of the phase of the sampling clock with respect to the signals is also discussed.

4.2.7.1 Readout synchronization

After digitization, data are held in a buffer pipeline awaiting a possible Level-1 Accept decision. If such a decision is given, the data are read out. The Level-1 Accept corresponding to a given bunch crossing has to match the data from the same bunch crossing. For the ECAL this means, in particular, that the sample corresponding to the maximum has to be propagated to the DAQ at a fixed position in the 10-sample timeframe which is read out. This can be achieved by setting the readout pointers in the pipeline memories inside the FE boards. The most important source of variation between the pointer values among the channels of a supermodule is due to the Level-1 signal distribution and is in the range of 2 to 3 units of the 40 MHz clock. The values of the readout pointers will be obtained during the commissioning period of the supermodules prior to their installation in the CMS detector. Laser pulses are sent to every channel and the corresponding 10-sample frames are analysed offline. The readout pointer values can be easily deduced from the position of the maximum of the digitized pulses inside the 10-sample streams. Correction of these settings will be calculated using the *a priori* knowledge of the different mean time of flight of the particles reaching the different parts of ECAL.

Because ECAL data is also sent to the Level-1 trigger, 2 possibilities exist for monitoring of the synchronization of the readout pipelines. They both are based on the generation of the Trigger Primitives. In the first method, a full readout is made of all trigger towers above a threshold of 1-2 GeV and a stream of 16 samples is extracted from the pipeline memory. An offline comparison of the recorded trigger primitives with the results of an emulation of the TPG (based on the recorded 16-sample stream) allows the readout synchronization to be verified. This is achieved by requesting that the results of the emulation match perfectly the online values of the TPG in the same trigger tower. A second method based on the same technique developed for the trigger primitives synchronization is described below.

4.2.7.2 Trigger Primitives synchronization

The CMS trigger system is a pipelined and synchronous system working at the LHC reference clock frequency. At each processing stage trigger data must be synchronized. The trigger primitives generated by the front-end electronics are collected by the Trigger Concentra-

tor Cards (TCC) which are part of the off-detector electronics. They are sent to the Regional Calorimeter Trigger which combines them in order to extract the total transverse energy in a pair of trigger towers. Non-negligible differences on the trigger primitives arrival time to the processors are introduced by different particle flight paths, different optical transmission fibre lengths and different phase lock delays in the electronic serializers. Thus a synchronization procedure is necessary. The synchronization is performed in the Synchronization and Link Board (SLB), a mezzanine board of the TCC. Trigger channel synchronization relies on the TTC Bunch Crossing Zero (BC0) broadcast command that can be adjusted relative to the LHC orbit signal, on a synchronization FIFO and on an accumulator histogram that reflects the LHC bunch crossing structure. A common hardwired control signal distribution guarantees aligned trigger data transmission through all high speed (1.2 Gb/s) transmission links. The SLB trigger channel accumulator histograms are analysed by an online readout program which is responsible for setting-up the trigger channel synchronization before data taking.

Trigger channel occupancy at each bunch crossing is dominated by minimum bias events. Simulations with a minimum-bias event rate corresponding to low luminosity ($\mathcal{L} = 10^{32}$ $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) give a trigger channel occupancy of 0.8×10^{-4} at transverse energies higher than 1.0 GeV. Although lower transverse energies imply higher channel occupancy and faster accumulation histogramming, energy deposits near the tower noise have lower BCID efficiency. For energies higher than 1.0 GeV the BCID filter can be considered fully (100%) efficient. A Poisson generator was used to simulate particle collisions along the LHC orbit for this level of occupancy.

Channel synchronization relies on a correct identification of the LHC bunch structure. The channel alignment can be investigated using the discrete correlation function between the ideal bunch structure and the structure acquired in the SLB trigger channel accumulator. In Fig. 4.8 are shown the correlation curves after 2 minutes and 25 minutes of a LHC run for the barrel case. As expected the maximum of the correlation function is obtained for a correct alignment (ΔBX=0). The errors associated with the correlation points decrease with the acquisition time since higher statistics are acquired, therefore the correlation peak become more significant for long acquisition periods. The plot in Fig. 4.9 represents the correlation significance as a function of the accumulation time.

Correct alignments can be considered to be obtained if the significance is greater than 5. This is achieved after 5 minutes of LHC run. For high luminosity ($\mathcal{L} = 10^{34}$ $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$), the same performance would be obtained a hundred times faster, i.e. in 3 seconds. Using the same time alignment procedure for the endcap part would require an acquisition time longer because in this case the procedure is applied to the pseudo-strip signals. Calorimeter trigger channel synchronization using identification of the LHC bunch structure can be considered an efficient and fast process.

4.2.7.3 Phase adjustment monitoring

As described in Section 4.2.1 each individual signal is sampled with a 40 MHz clock and digitized by an ADC after amplification. From the recorded set of sample values the amplitude of the corresponding signal can be reconstructed. The amplitude reconstruction procedure presented in Section 4.3 assumes the phase of the sampling clock with respect to the signal to be fixed and constant. A verification procedure to ensure this will be necessary.

The clock distribution system allows adjustment of the phase of the sampling clock by steps

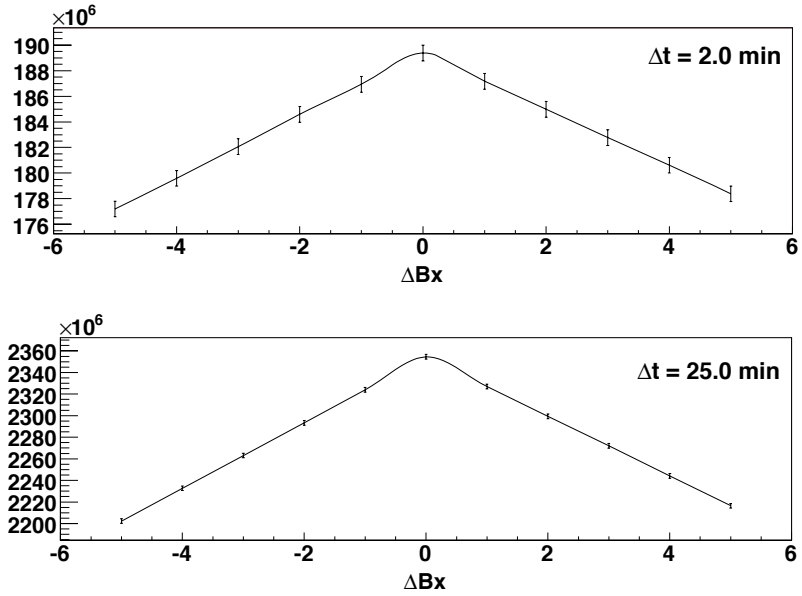


Figure 4.8: Correlation curves between ideal LHC bunch structure and the bunch structure acquired in the SLB trigger channel accumulator (after 2 and 25 minutes) as a function of the bunch crossing phase between the 2 structures.

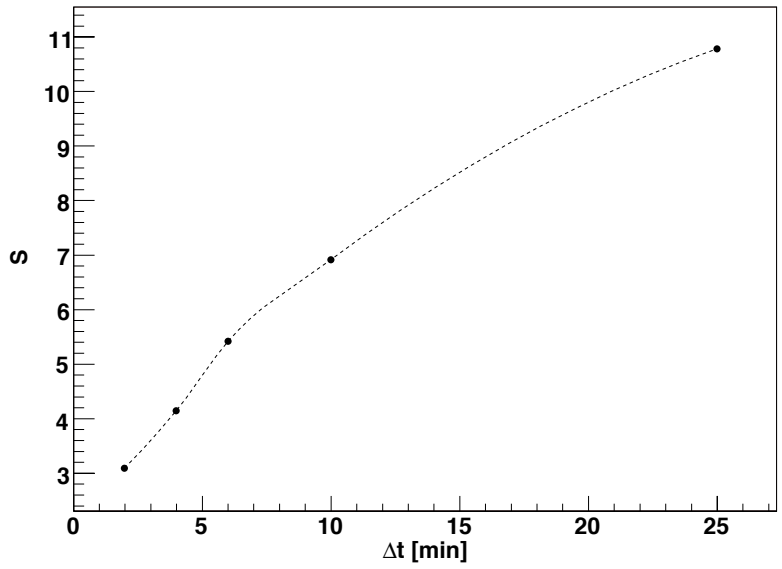


Figure 4.9: Peak significance for the correlation function after a given accumulated period (barrel case).

of 1 ns in a range of 25 ns for groups of 25 channels (corresponding to a trigger tower in the barrel and a supercrystal in the endcap).

The phase of signals in the installed detector, before any adjustment, will not be constant across the detector because of the geometry and architecture of the clock distribution, and because of the different time of flight from the interaction point to the detector. Laboratory measurements, performed on a large fraction of the final VFE electronics channels, show that the total spread of the peaking time distribution is less than 1 ns.The time of flight ranges. from 3.9 ns for $\eta = 0$ to 10.2 ns at $\eta = 1.5$.

Event-to-event phase variation caused by fluctuations on the time of flight because the longitudinal spread in the position of the primary vertex is less than ± 0.2 ns. Such fluctuations should not be of any consequence.

In order to adjust the sampling phase and subsequently verify that it remains constant, it is proposed to measure the phase using a digital filtering technique. This technique has been tested on beam data recorded in 2004. Figure 4.10 shows the distribution of the resolution of the difference between the time extracted using the digital filtering technique and the test beam trigger time (given by TDC) as a function of signal amplitude. The signal amplitude is expressed in units of the noise (i.e. around 40 MeV). It can be seen that a resolution less than 1 ns can be achieved for signals with an amplitude greater than 2 GeV. It has been shown that this time resolution is dominated by the test beam trigger timing measurement, which amounts to 0.6–0.7 ns.

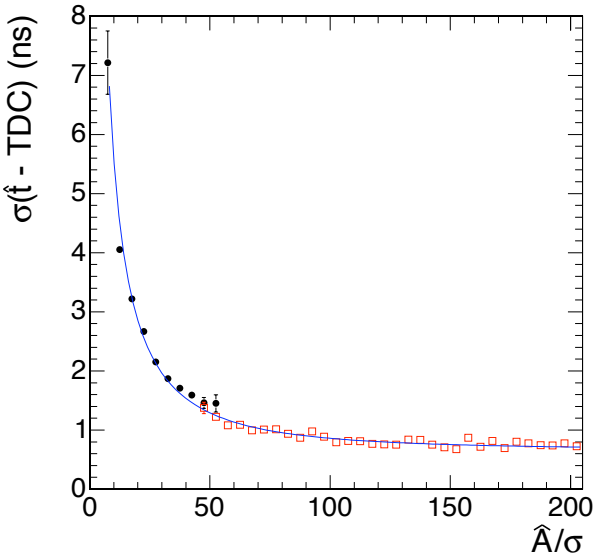


Figure 4.10: Resolution of the digital filtering technique versus the amplitude of the pulse expressed in units of noise.

4.3 Detector performance

In November 2004 a complete barrel supermodule (SM10) fully equipped with near final electronics was tested with high energy electrons ($20 \leq E \leq 250$ GeV) in the CERN H4 beam line. The data taken have allowed verification of the performance of a completed element of

the final detector. Detailed studies have been made of amplitude reconstruction, electronics noise and energy resolution. The results of these studies is summarized in this section.

The ECAL performance for unconverted photons or for electrons in the test beam, is entirely dominated by contributions from effects which have to be added to the shower simulation. This is done when the electronics response is simulated—the "digitization" simulation. The results presented below make it clear that the values used for two of the most important of these contributions, noise and photo-statistics, are consistent with what is measured in data. The results also verify that the algorithm used to reconstruct the signal magnitude from a time frame of consecutive digitizations does not itself add a contribution to the energy resolution. The most important contribution to the intrinsic ECAL performance, intercalibration, is the subject of Section 4.4.

4.3.1 Amplitude reconstruction

The raw data for a single channel consists of a series of consecutive digitizations of the signal making up a time frame. The number of samples is adjustable (2+4n) with a default of 10. The digitizations are made at the bunch crossing frequency of 40 MHz using an ADC clock that is locked to the LHC bunch structure. The timing of the signal will be adjusted in LHC running so that the signal pulse maximum corresponds to one of the samplings. The channel to channel variation, after this adjustment, is expected to be less than 1 ns.

The simplest method of reconstructing the amplitude is to take the sampling on the maximum as the measurement of the signal. However, one of the reasons for reading out a larger number of samples is to allow more sophisticated digital processing of the signal to reduce the noise contribution. The other reason is to enable identification of out-of-time (other bunch-crossing) pile-up. The signal amplitude is computed as a linear combination of discrete time samples as shown in equation 4.1.

$$\hat{\mathcal{A}} = \sum_{i=0}^{i=N} w_i \times S_i \tag{4.1}$$

where w_i are the weights, S_i the time sample values in ADC counts and N is the number of samples used in the filtering. The weights are determined to minimize the noise contribution. Details of the mathematical formalism of the optimization procedure can be found in [141].

4.3.1.1 Test beam studies

The shape of the signal pulse as a function of time is shown in Fig. 4.11. If there are no correlations of noise between time samples, the total noise level decreases with the number of samples used in the reconstruction. With the shape shown in Fig. 4.11, the use of 5 samples to reconstruct the amplitude should give noise level about 40% lower than the single sample noise. Considering more than 5 samples does not significantly improve the reduction of the noise.

Test beam data have been used to study the choice of amplitude reconstruction algorithm [142]. It is observed that the correlations between the samplings prevent the 40% noise reduction factor being achieved. The noise level obtained using 5 samples is very similar to that obtained from a single sample. There is also some noise correlation between nearby channels,

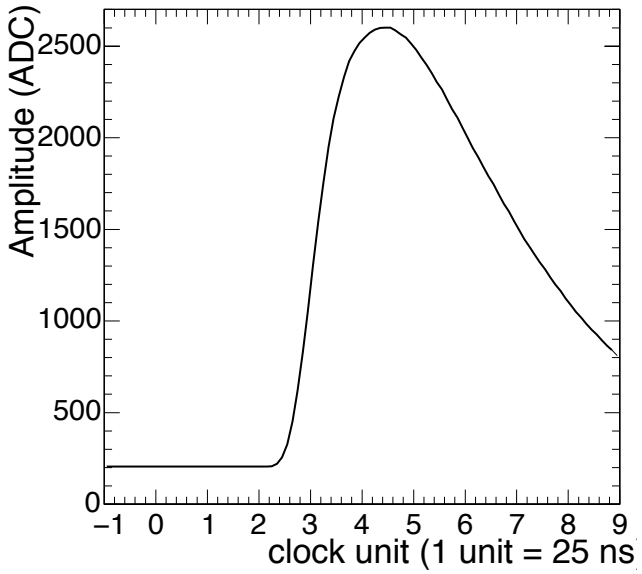


Figure 4.11: Profile of the signal pulse from a crystal of the SM10 supermodule using an electron beam of 120 GeV.

so that the noise measured in a sum of 25 channels (a typical cluster size used for ECAL shower reconstruction) is, in fact, slightly worse when the reconstruction uses 5 samples rather than a single sample.

This correlation is due to the presence of a small level of low frequency noise (i.e., noise at a much lower frequency than the 40 MHz digitization) that remains despite the grounding and screening. Because of this, the signal pulse sits on a baseline which varies from event to event. The weights can be determined so as to remove this varying pedestal. This is best done by additionally using samplings taken before the signal pulse together with the signal samplings. Using 3 samplings taken before the signal pulse, and 5 samplings on the signal pulse reduces the total noise seen in a sum of 25 channels by 20% as compared to what is measured by reconstruction with a single sampling followed by pedestal subtraction using an average pedestal value. With this reconstruction the total noise seen in a sum of 25 channels is almost exactly 5 times the noise seen in a single channel showing that the coherent noise has been very effectively removed. The average value of the noise seen in 1000 channels in supermodule tested in the H4 test beam in October 2004 is slightly less than 40 MeV. The implementation of a noise covariance matrix, and its use in the derivation of the weights does not improve the noise performance.

The optimization of the weights to minimize the noise contribution is insensitive to the small variation of signal pulse shapes present in the supermodule tested. Thus it is possible to use a single set of weights, which can be determined before LHC startup, for the amplitude reconstruction of all channels.

Small variations of the pulse maximum timing from channel to channel generate a small bias in the reconstructed amplitude which, in the absence of any precise overall absolute calibration, are simply swallowed by the intercalibration of the calorimeter. Systematic variation of the timing of maximum might result in an unacceptable variation of the channel response with time. A shift of 1 ns corresponds to a response variation of about 0.02%, but a shift of

3 ns corresponds to 0.2%. To prevent degradation of the resolution induced by such variation, the time of maximum must be carefully monitored.

The reconstruction of the signal is affected by jitter on the timing. The expected jitter in LHC running is about 0.2 ns, dominated by the longitudinal spread of the interaction vertex. The resolution is not degraded by such a small uncertainty on the signal timing.

4.3.1.2 Use of other MGPA gain ranges

For signal pulses resulting from very large energy deposits in a single crystal ($E > 160\,\text{GeV}$ in the barrel, and $E > 300\,\text{GeV}$ in the endcap) the data-frame contains samples in different MGPA gain ranges. Reconstruction using the peak sample alone greatly simplifies the treatment of such data-frames, and the noise is negligible at these energies. The relative calibration of the different ranges must be determined precisely to avoid any degradation of the resolution in gain transition and introducing a nonlinearity of the response.

Further details concerning the amplitude reconstruction, the determination of gain ratios, the effect of pile-up, as well as the complications of reconstructing test beam data where the ADC clock is asynchronous with respect to the signal pulse, are given in [142].

4.3.2 Energy resolution in test beam

4.3.2.1 Resolution for central impact

Showers in the ECAL are reconstructed by building clusters of crystals. In the test beam the best performance is obtained using a simple sliding window centred on the crystal having the maximum energy, summing 3×3 or 5×5 crystals. The energy contained in such a cluster varies with the shower position. The simplest verification of the performance of the ECAL is provided by a measurement of the energy resolution when the incident electrons are restricted, using the position measured in a set of beam hodoscope detectors, to a small region so that this variation is negligible. Figure 4.12 shows examples of the distribution of the reconstructed energy, fitted by a Gaussian to obtain the energy resolution σ_E/E.

4.3.2.2 Resolution for uniform impact

The test beam data was taken using a trigger covering a $20\times20\,\text{mm}^2$ area, roughly matching the beam size ($\sigma \approx 10\,\text{mm}$) and only slightly smaller than the crystal granularity. When reconstructing showers from incident electrons spread over this area the energy resolution can be improved by making a correction for the varying containment as a function of the shower position as measured by the ECAL. In the analysis reported here, the position was measured independently in the 2 lateral coordinates using the logarithm of the ratio of crystal energies on either side of a reference line defined by the crystal edge. A single correction function, parameterized from the data, was used for all energies, and for all regions in the supermodule. Excellent performance is obtained even, for example, for runs where the supermodule was positioned such that the beam (and the $20\times20\,\text{mm}^2$ trigger) was centred on the corner of a crystal. Figure 4.13 shows an example of the energy reconstructed in such a case, together with the energy distribution before applying the correction.

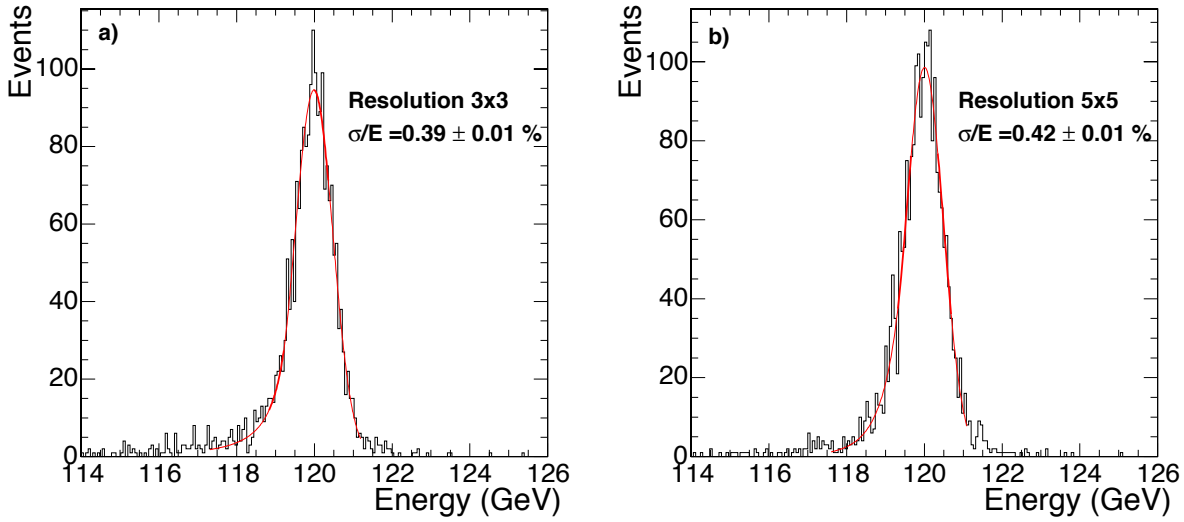


Figure 4.12: Distribution of the energy reconstructed in (a) a 3×3 cluster and (b) a 5×5 cluster, when 120 GeV electrons are incident in a 4×4 mm^2 region.

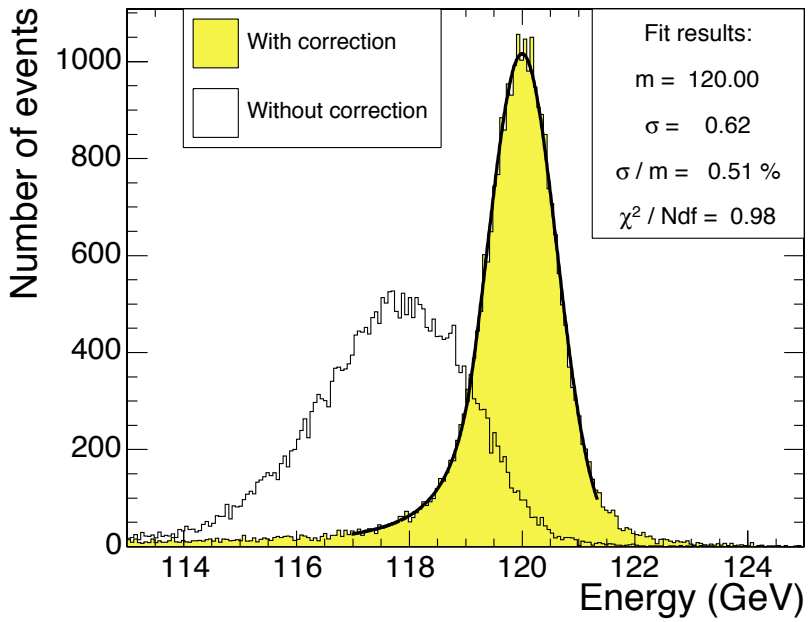


Figure 4.13: Distribution of the energy reconstructed in a 3×3 cluster when the beam of 120 GeV electrons is centred on a crystal corner. The outline histograms show the energy reconstructed before applying the energy containment function.

4.3.2.3 Energy resolution as a function of energy

The energy resolution, measured by fitting a Gaussian function to the reconstructed energy distributions, has been parameterized as a function of energy,

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2, \tag{4.2}$$

where S is the stochastic term, N the noise and C the constant term. As presented earlier, Fig. 1.7 shows an example of the energy resolution as function of energy, together with the fitted function. Similar results are obtained for sets of data (energy scans) taken throughout the supermodule. The upper series of points are the resolution values obtained from fits to the energy distribution obtained for unselected events taken with a $20\times20\,\mathrm{mm}^2$ trigger and reconstructed using the containment correction described in the previous section. The lower series of points are obtained from events selected to fall within a $4\times4\,\mathrm{mm}^2$ region. Due to the tight selection, the statistical error on the measured energy resolution in the lower set is rather large for some data points.

4.4 Calibration and alignment

Calibration is a severe technical challenge for the operation of the CMS ECAL. It is naturally seen as composed of a global component, giving the absolute energy scale, and a channel-to-channel relative component, which is thereafter referred to as intercalibration.

Many small effects which are negligible at low precision need to be treated with care as the level of precision of few per mil is approached. The essential issue is stability in both time and space, so that showers in different locations in the ECAL in data recorded at different times are accurately related to each other.

The main source of channel-to-channel response variation in the barrel is the crystal-to-crystal variation of scintillation light yield which has an RMS of $\approx 8\%$. In the endcap the VPT signal yield, the product of the gain, quantum efficiency and photocathode area, has an RMS variation of almost 25%. The nature and technology of the ECAL provides no convenient or *a priori* way of intercalibrating the channels and the target precision can only be achieved using physics events. Over the period of time in which the physics events used to provide an intercalibration are taken the response must remain stable and constant to high precision. Where there is a source of significant variation, as in the case of the changes in crystal transparency caused by irradiation and subsequent annealing, the variation must be precisely tracked by an independent measurement. The changes in crystal transparency are tracked and corrected using the laser monitoring system.

The final goal of the calibration strategy is to achieve the most accurate energy measurement for electron and photons. Schematically the reconstructed energy might be decomposed to 3 factors,

$$E_{\mathrm{e},\gamma} = G \times \mathcal{F} \times \sum_i c_i \times A_i, \tag{4.3}$$

where G is a global absolute scale. The function $\mathcal{F}$ is a correction function depending on the type of particle, its position, its momentum and of the clustering algorithm used. The

c_i factors are the intercalibration coefficients while the A_i are the signal amplitudes, in ADC counts, which are summed over the clustered crystals.

The correction function $\mathcal{F}$, discussed in more detail in Chapter 10, corrects for energy loss due to bremsstrahlung and various containment variations. In the current working model, $\mathcal{F}$ =1 is chosen for the reference 5×5 crystal shower reconstruction algorithm, used for unconverted photons, or for electrons in the test beam. Different reconstruction algorithms are used to estimate the energy of different electromagnetic objects, i.e., unconverted photons, electrons and converted photons, each of them having their own correction functions. At present these "algorithmic" corrections are obtained from the simulated data by accessing the generated parameters of the Monte Carlo simulation. For some of the corrections, for example the containment corrections, this is an acceptable procedure provided that test beam data is used to verify the simulation, so that, in effect, the simulation is being used only as a means of interpolating and extrapolating from data taken in the test beam. In other cases, where the test beam provides no useful information, for example in issues related to conversions and bremsstrahlung radiation in the tracker material, it is necessary to use information that can be obtained from data taken *in situ* with the running detector. Two particularly useful channels which can be used to obtain such information are under investigation: $Z \rightarrow ee$, and $Z \rightarrow \mu\mu\gamma$ (the photon coming from inner bremsstrahlung). The latter, in the case of unconverted photons, is also able to set the global scale G.

Preliminary estimates of the intercalibration coefficients c_i are obtained from laboratory measurements of crystal light yield, test beam precalibration of some supermodules, and the commissioning of further supermodules with cosmic rays. Imposing the ϕ-independence of energy deposited in the calorimeter can be used to rapidly improve on this start-up intercalibration for fixed η regions. The method of intercalibration with physics events that has been investigated in the most detail, uses the momentum of the abundant electrons as measured in the tracker, mainly from $W \rightarrow e\nu$, which have a similar p_T to the photons of the benchmark channel $H \rightarrow \gamma\gamma$. A complementary method, not relying on the momentum measurement, is based on $\pi^0 \rightarrow \gamma\gamma$ and $\eta \rightarrow \gamma\gamma$ mass reconstruction.

The achievable precision, the methods used, and the path to full precision will be a function of time and available luminosity.

4.4.1 Intercalibration at start-up

4.4.1.1 Laboratory measurements

The calorimeter is being assembled in 2 regional centres: at CERN and at INFN-ENEA Casaccia near Rome. During the assembly phase, all the detector components are characterized [143, 144] and the data are saved in the construction database. It is thus possible to predict the calibration c_i of each channel i using the laboratory measurements as

$$\frac{1}{c_i} \propto LY \cdot \varepsilon_Q \cdot c_{ele} \cdot M, \tag{4.4}$$

where LY is the Light Yield of the crystals, M and ε_Q are respectively the gain and quantum efficiency of the photo-detectors and c_{ele} is the calibration of the electronics chain. The crystal LY is measured in the laboratory with a photo-multiplier tube, exciting the crystal with a ^{60}Co source. This gives an average $\mathrm{LY_{PMT}}$ for the $PbWO_4$ crystals of 10 pe/MeV at 18°C.

It is a delicate and difficult measurement because the $PbWO_4$ crystals have a rather low LY and the energy of the ^{60}Co γ is only 1.2 MeV. Due to the different percentage of the crystal rear face covered by the PMT with respect to the VPT or APD and to the quantum efficiency, different values are measured in equipped ECAL units: ≈ 4.5 pe/MeV for the barrel channels read out with APDs and a similar number in the endcap channels read out by VPTs (the larger surface area of the VPT photocathodes being balanced by the higher quantum efficiency of the APDs).

The measurements span about 7 years of crystal production, so the stability of the LY bench calibration is crucial and is constantly controlled using reference crystals. The LY measurement of each crystal is rescaled using the daily reference crystal measurements [145].

The laboratory LY measurement can be improved exploiting the correlation between the crystal LY and the Longitudinal Transmission at 360 nm (LTO_{360}). The latter can be measured with better precision and stability [146, 147].

Test-beam intercalibration constants, as described in Section 4.4.1.2, can be used to measure the precision achieved by laboratory estimations of intercalibration. Figure 4.14a shows the intercalibration coefficients calculated from laboratory measurements plotted against the test-beam measurements. Figure 4.14b shows that a resolution of about 4% is obtained with this method.

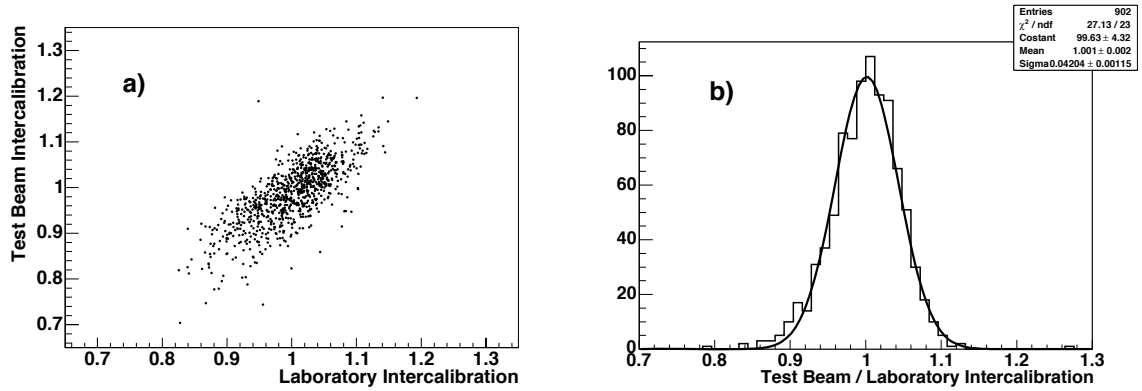


Figure 4.14: (a) Intercalibration coefficients obtained with supermodule 10 at the 2004 test-beam with high energy electrons versus intercalibration coefficients calculated from laboratory measurements and (b) distribution of the fractional difference between the laboratory estimated coefficients and coefficients measured in the test-beam.

4.4.1.2 Testbeam precalibration

The intercalibration method presented here was developed using test beam data taken at CERN between year 2000 and 2004. In the test beam, supermodules are mounted on a rotating table which allows rotation in both the η and ϕ coordinates and a full scan of the supermodules with high energy electron beams. The electron position is measured with a set of hodoscopes.

The response of a single crystal to electrons depends on the electron impact position. The dependence in the 2 lateral coordinates can be factorized and fitted with a 4^{th} order polynomial separately. The corrected response of the single crystal S_{corr} can be obtained from the

measured amplitude $S_{\rm meas}$ as:

$$S_{\rm corr} = S_{\rm meas} \frac{P_x^{\rm max} P_y^{\rm max}}{P_x(x) P_y(y)}, \tag{4.5}$$

where x and y are the measured positions of the incident electron in the two lateral coordinates, and $P_{x,y}^{max}$ is the maximum of the polynomial. Figure 4.15 shows the fit and the crystal response after correction. Only the events impinging in a central $7\times7\,\mathrm{mm}^2$ window are used, (about 25% of events taken with a $20\times20\,\mathrm{mm}^2$ trigger). The intercalibration coefficients c_i for crystal i are defined as the ratio of the mean value of the corrected response with respect to a reference value.

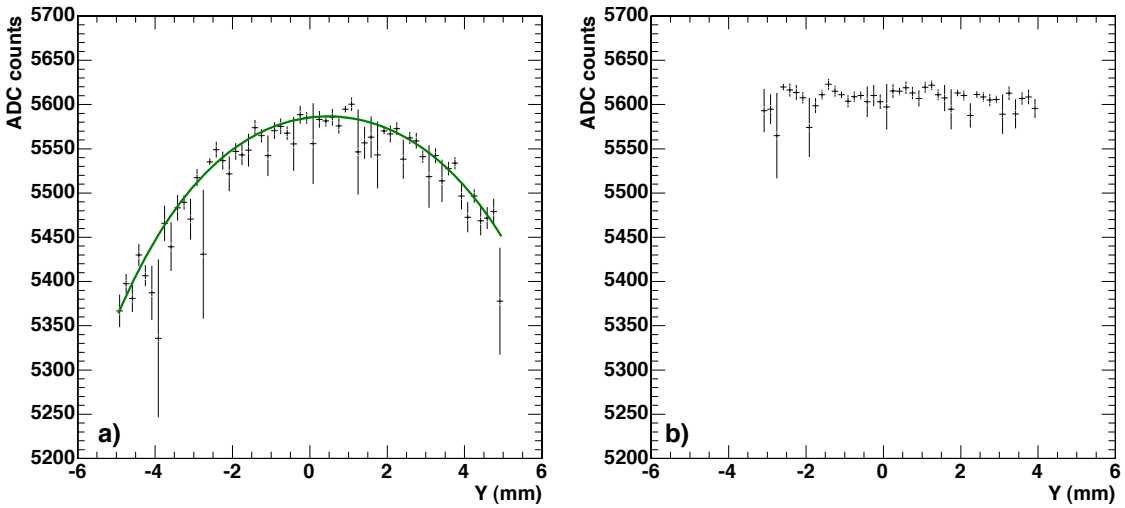


Figure 4.15: a) Polynomial fit to energy in a single crystal as a function of a lateral coordinate; b) energy in a single crystal as a function of a lateral coordinate after application of polynomial corrections for both lateral coordinates.

The statistical uncertainty remains negligible (less than 0.1%) provided that at least 1000 events are taken per crystal. Several tests have been done to evaluate the robustness of this method: using a different polynomial for each crystal or the same polynomial for all crystals. The intercalibration precision, when these constants are used *in situ*, is expected to be limited by variations occurring in the time between their determination in the test beam and their utilization in the installed detector.

At least 5 supermodules will be precalibrated in the H4 beam during the summer of 2006. The same supermodules will also be calibrated with cosmic rays, as described in Section 4.4.1.3 in order to compare the response of the 2 methods. It is foreseen to repeat the precalibration of at least one supermodule in order to quantify the reproducibility of the procedure, and the transportation of the constants over a period of a couple of months.

4.4.1.3 Measurements with cosmic ray muons

Intercalibration coefficients for barrel supermodules are also obtained using cosmic muons which are well aligned with the crystal axes [148]. For this measurement the APD gain is

increased by a factor 4 with respect to the gain to be used during normal data taking by increasing the bias voltage.

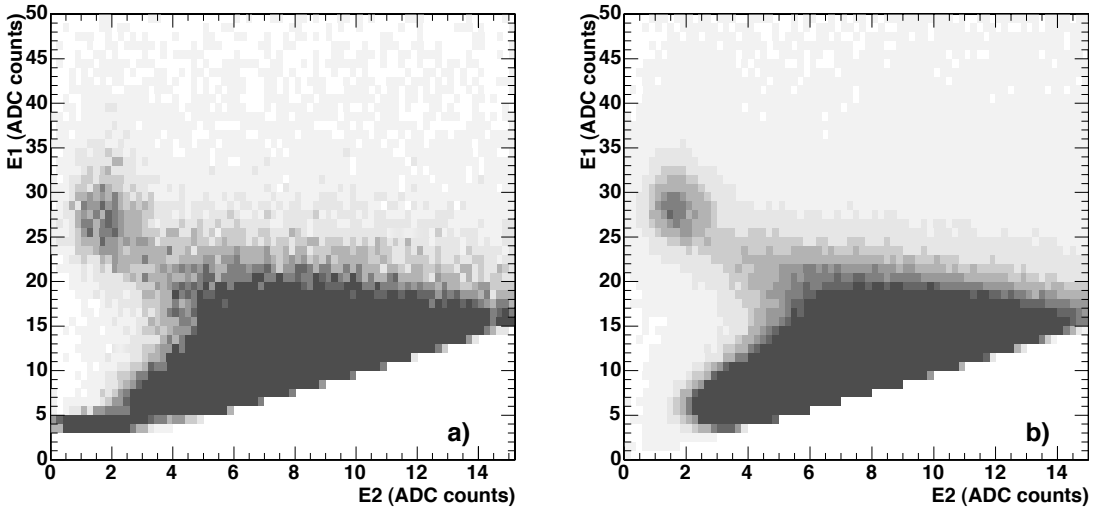


Figure 4.16: Scatter plot of E_1 vs E_2 in real cosmic muon data (a) and in simulation (b). The number of events in the plots corresponds to 41 hours for the data and to 1 week for the simulation. Edge crystals have been excluded.

The feasibility of this method was verified by measurements on a supermodule which was exposed to cosmic ray muons for 41 hours in November 2004, and with a detailed GEANT4 simulation [149]. Good agreement was found between data and simulation. Well aligned cosmic rays, giving a large signal in the crystal they pass through (E_1), are selected by vetoing on signals above a rather low threshold in the adjacent crystal with highest signal (E_2), obviating the need for external tracking (Fig. 4.16).

In the region covered by the trigger, which roughly corresponds to Module 1, an agreement of about 3% was achieved with respect to the testbeam calibration as shown in Fig. 4.17. The statistical contribution to the overall uncertainty was estimated to be 2%.

Some difficulties emerge when calibrating edge crystals (for which the veto based on neighbouring channels is inefficient) and the extreme crystals of Module 4 (due to the unfavorable angle with respect to the dominant cosmic ray direction). However, by using a more selective trigger setup in the cosmic telescope and by inclining the supermodule by $\approx 10^\circ$, an overall precision of 3% should be achievable in 1 week of data taking for channels in Modules 1–3, and 3.5% for channels in Module 4.

4.4.2 Phi independence

A method taking advantage of the ϕ-symmetry of deposited energy to intercalibrate crystals within rings at constant η has been studied. Intercalibration is performed by comparing the total energy deposited in each crystal with the mean of the distribution of total energies for all crystals at that pseudorapidity. Two choices of event trigger have been investigated: random bunch crossings [150], and Level-1 jet triggers [151]. The results shown in the following are based on a Level-1 jet-triggers simulated sample.

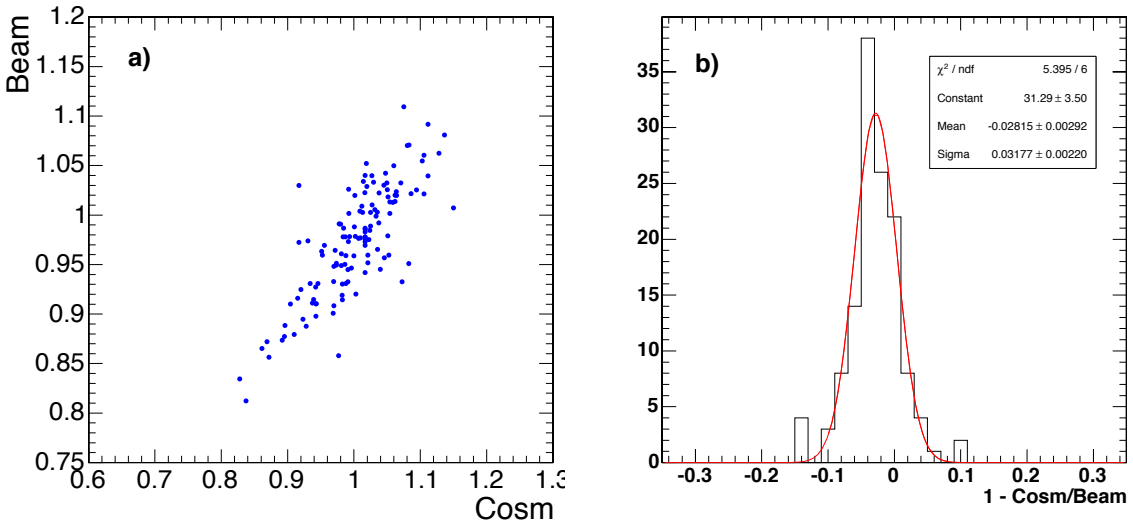


Figure 4.17: Correlation between (a) the testbeam calibration coefficients and cosmic ray muon coefficients, and (b) relative precision of the cosmic muon calibration constants. Only those channels with more than 50 cosmic muon events are considered.

The value of the total transverse energy, ΣE_{T}, deposited in each crystal from all selected events is determined. In the case of jet triggers, only crystals with transverse energy deposits in the range $1 < E_{\mathrm{T}} < 6\,\mathrm{GeV}$ (barrel), and $1 < E_{\mathrm{T}} < 4\,\mathrm{GeV}$ (endcaps) contribute to the energy sums The lower threshold excludes noise, and the upper threshold improves the stability of the sums. In addition, in order to avoid trigger bias, crystals associated with the jet which has triggered the event (the highest-E_{T} jet in the event) are excluded by requiring that crystals are separated from the position of the triggering jet by more than 1 rad. The intercalibration precision for a given η is obtained from the Gaussian width of the distribution of ΣE_{T} for the pair of rings of crystals at that absolute value of η. The precision is determined separately for each of 85 pairs of rings in the barrel and 39 pairs of rings in the endcaps.

With perfect ϕ-symmetry the intercalibration precision is expected to vary as $1/\sqrt{N}$, where N is the number of events. A limit on the precision arises due to non-uniformities in ϕ, primarily from the inhomogeneity of tracker material, but also from geometrical asymmetries such as the varying off-pointing angle of endcap crystals, and the boundaries between barrel supermodules. These non-uniformities result in a spread in the ΣE_{T} values which cannot be reduced by increasing the statistics of the event sample. The limit s on the precision is determined by calculating the precision for different of numbers of events, N, and fitting a two parameter function

$$f(N) = \sqrt{s^2 + (m/\sqrt{N})^2}.$$

to the measured points. The parameter, s, measures the limit on the attainable precision. The technique has been tested directly by applying a set of miscalibration factors, chosen randomly from a Gaussian distribution of width 4.5%, to crystals in a pair of rings, and determining the residual miscalibration after correction. Figure 4.18 shows the distribution of residual miscalibrations ϵ_R after 2 iterations of the correction procedure, obtained with 1.1×10^7 Level-1 jet trigger events. The Gaussian width of the distribution is (2.7±0.1)%.

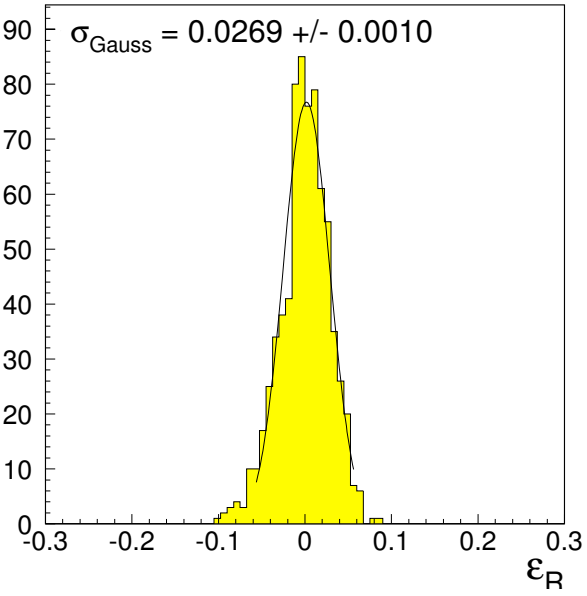


Figure 4.18: Distribution of residuals ϵ_R, for the 720 crystals in the pair of rings at $|\eta| = 0.23$, after 2 iterations of the correction procedure, for 1.1×10^7 Level-1 jet trigger events.

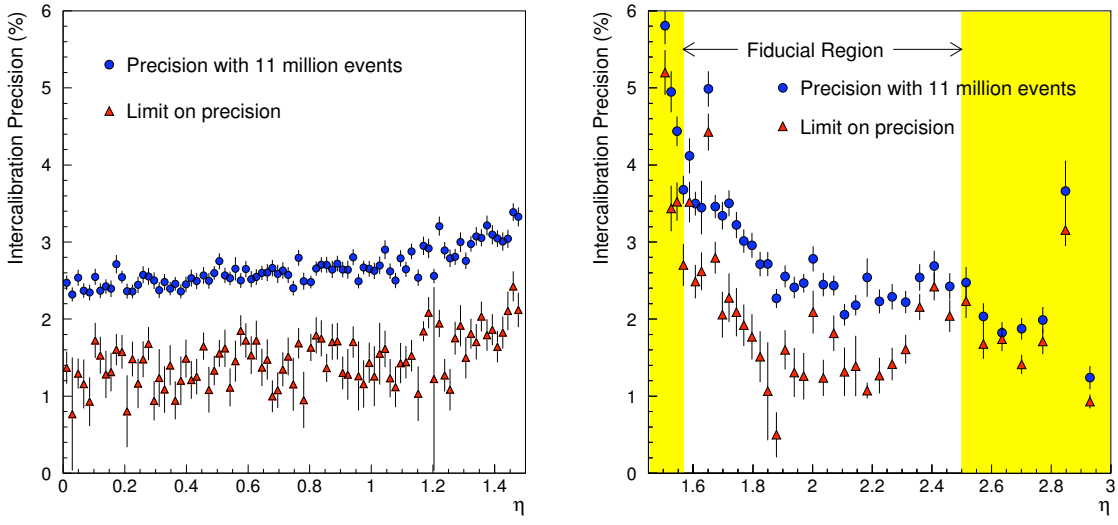


Figure 4.19: Intercalibration precision which can be obtained with 11 million Level-1 jet trigger events and the limit on the intercalibration precision due to tracker material inhomogeneity as a function of η.

It can be seen in Fig. 4.19 that without using any knowledge about the material distribution in the tracker, the limit on the precision is close to 1.5% throughout the barrel and between 3.0% and 1.0% for the fiducial region of endcaps. It can be expected that the limit on the precision will be closely approached with a few tens of millions of events.

This is equivalent to about 10 hours of data taking, under the assumption that 1 kHz of Level-1 bandwidth is allocated to single jet triggers, and that the calibration software has

access to this rate, either running on the Filter Farm, or more probably, running offline on a highly compacted data stream (a few tens of channels stored per event). With increasing knowledge of the material distribution in the tracker, after the start of LHC running, the attainable precision of the method is expected to increase, with the potential of providing rapid and repeated calibration of the ECAL. Intercalibration of different η rings can be done with one of the following methods.

4.4.3 Intercalibration using single electrons

Once the Tracker is fully operational and well aligned, intercalibration of different crystals within a single module can be performed using the momentum measurement of isolated electrons. The main difficulty in using electrons for intercalibration is that they radiate in the tracker material in front of the ECAL, and both the energy and the momentum measurement are affected. Moreover the average amount of bremsstrahlung varies with tracker material thickness.

Detailed Monte Carlo simulation studies of calibration using electrons have been made using fully simulated $W \rightarrow e\nu$ events, digitized with pile-up corresponding to low luminosity ($\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$) [152]. While the methodology and techniques used are applicable to the whole ECAL, the characteristics of the events, the material budget and the geometry of the detector differs between the barrel and the endcap regions, and the 2 regions have been studied separately.

For these studies the ECAL energy was measured by summing the 5×5 array of crystals around the crystal with the maximum signal (S25). In the endcap the energy measured in the preshower and associated with the electron cluster is added to the energy summed in the crystals. The energy in a cluster was used for intercalibration, rather than the single crystal calibration used in the ECAL test beam. The cluster energy does not require the complexity of a single crystal containment correction, which would always be a potential source of error. Moreover, it makes better use of the available information by using also the energy deposited in crystals other than the one containing most of the energy. The choice to use S_{25} rather than the "superclustering" algorithms (Section 10.1), which are generally used for electron reconstruction, was motivated by the wish to cleanly separate the intercalibration from the algorithmic corrections required for the superclustering algorithms.

Each energy measurement contains the contribution of many crystals, each with its own calibration constant. In order to extract those constants the individual crystal contributions must be unfolded, while minimizing the difference between the energy and momentum measurements. Two algorithms to achieve this minimization have been considered: an iterative technique which was used for the in-situ calibration of the BGO crystals in the L3/LEP experiment [153] and a matrix inversion algorithm. The results, both in terms of precision and in terms of speed of algorithm, are similar, and show no dependence on the technique used.

4.4.3.1 The barrel case

The event selection was based on variables which are sensitive to the amount of bremsstrahlung emission, and consequently measure the quality of the energy and momentum reconstruction. These variables are the number of valid hits and the χ^2/n.d.f. of the electron track, S_9/S_{25} measured by the ECAL (where S_9 is the energy contained in a 3×3 array of crys-

tals around the crystal with the maximum signal), and the ratio S_{25}/P, where P is the track momentum.

The selection cuts were chosen by scanning the four-dimensional phase space of the selection variables and choosing the point giving the best precision (global minimum) in the determination of the calibration constants. The optimal point is dependent on both the number and the quality of the input electron measurements, both of which depend on the average amount of bremsstrahlung in the region considered. The sensitivity of the results to the selection cuts was extensively investigated. The calibration precision attained shows a rather wide global minimum with respect to variations of the selection cuts, which guarantees the stability of the selection process and allows for potential small differences between simulated events, used to derive the selection cuts, and data.

Due to the variation of the average value of S_{25}/P with pseudorapidity, caused by the variation of the amount of material in front of the ECAL, the calibration task was divided into 2 steps. In the first step crystals in small regions in η, over which the average value of the S_{25}/P is rather constant, were intercalibrated. In the second step the small regions are intercalibrated with each other.

To test the calibration algorithms, a Gaussian miscalibration has been applied crystal by crystal, with a spread of 4% around 1. This represents a conservative estimate of the residual miscalibration expected from the precalibration, and early intercalibration of the crystals at start-up (Section 4.4.1). For each test many Monte Carlo experiments ($\geq$50) were performed with different randomly chosen sets of miscalibration constants. The results given are the averages of the results of the experiments.

The calibration precision versus η achievable for a fixed integrated luminosity follows the tracker material budget distribution (Fig. 4.20). There are fewer electrons that have radiated only a little at higher η. Fewer events are selected by the High Level Trigger (HLT) here and those selected are still, on average, less well measured electrons. The Monte Carlo simulation data used to obtain these results correspond to about 5 fb^{-1}. This estimation uses the PYTHIA cross section for W-production with no k-factor.

The calibration precision was also extensively studied in different ϕ regions keeping the same η interval. There is no evidence of any ϕ dependence.

For electrons impacting crystals on the module boundaries, there is a loss in the collected energy. The calibration procedure artificially increases the coefficients for these crystals to compensate for energy loss. A suitable method must be used to deal with this problem. This may involve applying a correction for clusters spanning intermodule boundaries. Such a correction can be derived from test beam data.

It is found that extending the intercalibration area in ϕ direction, keeping the same η region, does not affect the intercalibration accuracy once a module boundary correction is applied.

After crystals within regions are intercalibrated, the regions have to be calibrated among themselves. This task is accomplished by selecting electrons with minimum energy loss due to bremsstrahlung. After this selection, the resulting values of the peaks of the S_{25}/P distributions that are found are consistent with the pseudorapidity dependence of shower containment.

Dependence on integrated luminosity

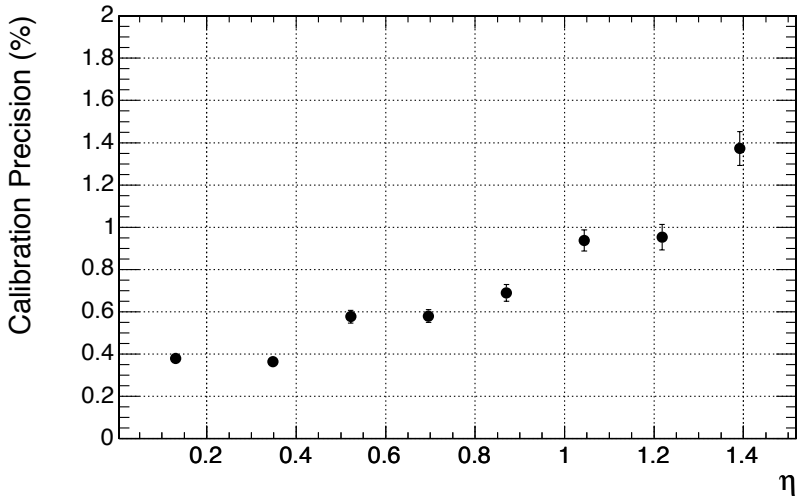


Figure 4.20: Calibration precision versus η for 5 fb^{-1} of integrated luminosity.

The calibration precision achievable is strongly dependent on the available statistics, as shown in Fig. 4.21. The dependence of the precision on the number of electrons per crystal output by the HLT, i.e. the number of electrons available to be used for calibration, can be fitted by the function:

$$\text{resolution} \;=\; \frac{A}{\sqrt{\text{HLT events per Crystal}}} \oplus C \tag{4.6}$$

The parameters A and C depend on η and have greater values as η increases. The constant term, C, sets the limit on the precision that can be reached with infinite statistics. In Table 4.2 the fitted values of A and C are listed for the different η regions. The corresponding fitted curves are shown in Fig. 4.21.

Table 4.2: A and C parameters for different η regions in the barrel.

η	A (%)	C (%)
0.000–0.261	6.19	0.12
0.783–0.957	10.7	0.27
1.305–1.479	15.0	0.42

4.4.3.2 The endcap case

For the study of the endcap GSF tracking was used for the electron track reconstruction (Section 10.4.2), and the event selection parameters for tracks are consequently different. As in the barrel case, the event selection uses variables which are sensitive to the amount of bremsstrahlung emission, and thus measure the quality of the energy and momentum reconstruction. The variables used were: S_9/S_{25} and the ratio $P_{\text{out}}/P_{\text{in}}$ between the momentum measured at the end of the electron track and that measured at the vertex. As in the barrel case a cut was also made on the ratio S_{25}/P itself. The selection efficiency varied between 10% and 30%. The sensitivity of the calibration precision shows only a small dependence on the selection cuts.

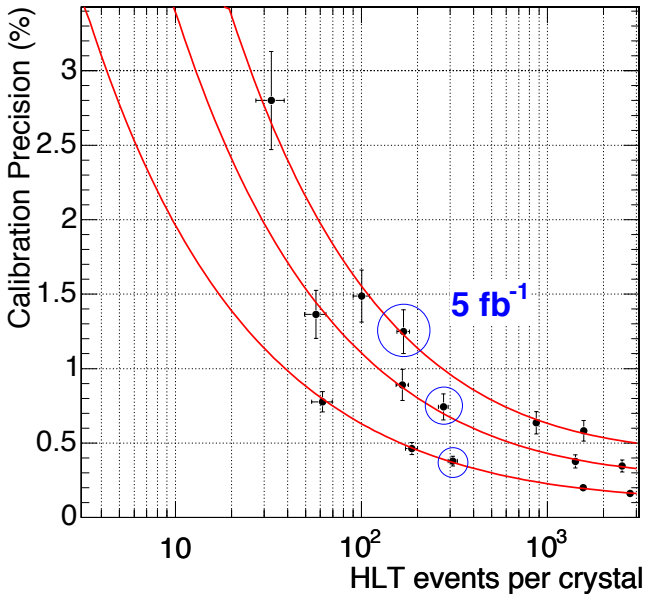


Figure 4.21: Calibration precision versus HLT events per crystal for different η regions. Upper curve: the last 10 crystals in the ECAL Barrel (1.305$<\eta<$1.479), Middle curve: 10 crystals in the middle of ECAL Barrel ($0.783 < |\eta| < 0.957$), Lower curve: the first 15 crystals in the ECAL Barrel ($0.0 < |\eta| < 0.261$). The third point along each line gives the precision for 5 fb^{-1} of integrated luminosity.

The calibration precision versus η achievable for a fixed integrated luminosity is shown in Fig. 4.22. No evidence of any ϕ dependence was found. The precision is limited by the momentum resolution, which is worse than in the barrel. No significant effect on the precision achieved was seen for crystals at the border of the supercrystal (Section 4.1) structure.

The calibration precision achievable is strongly dependent on the available statistics. The dependence of the precision on the number of electrons per crystal output by the HLT was fitted using the parametrization in Equation 4.6 for the region at $\eta \approx 2.0$. The data points and fitted curve is shown in Fig. 4.23.

4.4.3.3 Background contamination

The rate of the dijet background in the single electron trigger stream (HLT output) is estimated to be 27 Hz at low luminosity (Section 10.2.2) out of which 16 Hz are expected in the barrel. The residual background has been investigated for the barrel case. After the selection described above is applied the surviving background corresponds to a rate of 2.3 Hz. One third of this rate comes from $b/c \rightarrow e$ semileptonic decays. Such decays might be useful in the calibration process, increasing the overall calibration statistics. In fact their S_{25}/P distribution is very similar to the $W \rightarrow e\nu$ signal electrons. If required, the background can easily be further reduced by a factor 10 using isolation cuts, with only a small effect on the signal efficiency.

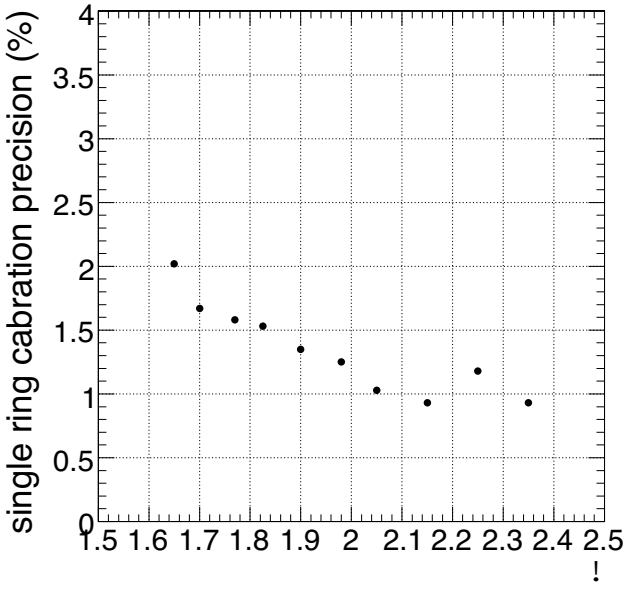


Figure 4.22: Calibration precision versus η for 7 fb^{-1} of integrated luminosity.

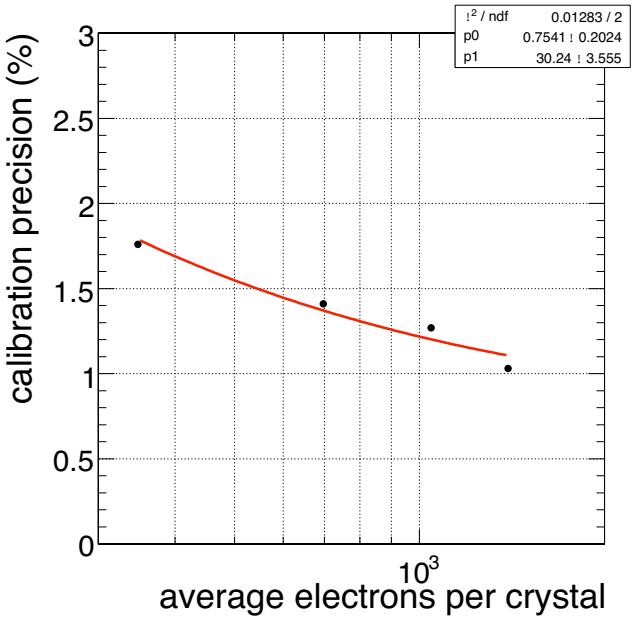


Figure 4.23: Calibration precision versus HLT events per crystal for a region at $\eta \approx 2.0$. The last point gives the precision for 7 fb^{-1} of integrated luminosity.

4.4.3.4 Systematic Uncertainties

It is assumed that the tracker can be well aligned with $5\,\text{fb}^{-1}$ of data. Nevertheless, in order to study the effect of a misaligned/miscalibrated tracker, the inverse of the track momentum was smeared by 2%. In that case, the calibration precision achievable using $5\,\text{fb}^{-1}$ of data worsens by a relative 20%.

If the initial miscalibration error is increased to 8%, the calibration precision achieved, for a given number of events, is only very slightly worsened. Such an initial high value corresponds to the calibration uncertainty for completely unmeasured crystals in the barrel.

It is inevitable that there will be a few non-functioning channels. As in the case of crystals at the module boundaries, the calibration procedure artificially increases the coefficients of the crystals in a 3×3 window around the dead one, in order to compensate for the energy loss caused the missing energy. A comprehensive correction procedure for dead channels might be to parameterize the energy content of the crystals in the 5×5 matrix as function of energy, η, ϕ, and particle impact point. This might be achieved by looking in areas without dead crystals and making a one-to-one correspondence of the energy deposition pattern using a technique similar to the one described in Section 10.1.4.

4.4.3.5 Summary of intercalibration using single electrons

Although further study remains to be done, the preliminary results from studies of intercalibration using single electrons are quite encouraging. The intercalibration precision attainable with $10\,\text{fb}^{-1}$, averaged over the barrel, is about 0.6%. The single electron rates have been calculated rather conservatively using the PYTHIA cross sections and no k factors, and not including additional contributions from electrons in b decays. Further optimization may involve more complete separation of the task of intercalibration of local regions where the amount of tracker material is fairly constant, from the task of intercalibrating those local regions to each other. This would allow the use of clustering algorithms which are less sensitive to bremsstrahlung than the 5×5 clusters, hence allow the use of a significantly larger fraction of the available electrons.

4.4.4 Use of $Z \rightarrow$ ee

The Z mass constraint in $Z \rightarrow ee$ decays is a powerful tool for ECAL calibration. A number of different uses are envisaged, from the tuning of the algorithmic corrections for electron reconstruction, to the intercalibration of regions of the ECAL, for example as a complement to the ϕ symmetry method at the start-up.

A study of the use of the $Z \rightarrow ee$ has been performed with a fully simulated data sample digitized with low luminosity pile-up, corresponding to an integrated luminosity of $2.4\,\text{fb}^{-1}$ [154]. This study was restricted to Z decays with both electrons in the barrel, but the method can be extended to the whole ECAL.

An iterative method has been developed to tune the algorithmic corrections and to extract intercalibration constants of regions or individual crystals from samples of $Z \rightarrow ee$ events. Starting from the 2 electron invariant mass reconstruction, the average quantity ϵ in event i is defined as

$$\bar{\epsilon}^i = \frac{1}{2} \cdot \left[\left(\frac{M^i_{inv}}{M_Z} \right)^2 - 1 \right], \tag{4.7}$$

The calibration constants are obtained from the peak of the $\bar{\epsilon}^i$ distribution, so as to reduce the sensitivity to tails. Convergence is reached after 10–15 iterations. The intercalibration precision reached is insensitive to the magnitude of the initial miscalibration.

4.4.4.1 Tuning of algorithmic corrections

The $Z \rightarrow ee$ sample also allows algorithmic corrections ($\mathcal{F}$ in Equation 4.3) to be determined or cross-checked. This was tested for the case of the functions used to correct the supercluster η dependence for different electron classes. It was assumed that the crystals were intercalibrated with a precision of only 2%. Miscalibration constants, with an RMS spread of 2% were applied to the crystals.

The electron energy was reconstructed with the superclustering algorithms described in Section 10.1, and the electron classification described in Section 10.4 was used. Events are selected where either both electrons belong to the golden class, or to the showering class (the "best" and "worst" classes, respectively). The energy is reconstructed in superclusters but without applying the algorithmic correction for variation in η. Figure 4.24 compares the correction factors obtained from $Z \rightarrow ee$ events used as real events can be used, with those obtained from Monte Carlo studies where the true values of the generated kinematics were used. The agreement in the shape is excellent, but small overall rescaling values are needed which are different for different classes of electrons and which have been used in the figure (1.005 for the golden electrons, and 1.008 for the showering electrons). The origin of these rescaling factors is under investigation, one suggestion is that the difference is due to the presence of pile-up in the $Z \rightarrow ee$ sample, and the absence of pile-up in the electron sample used to extract original correction factors.

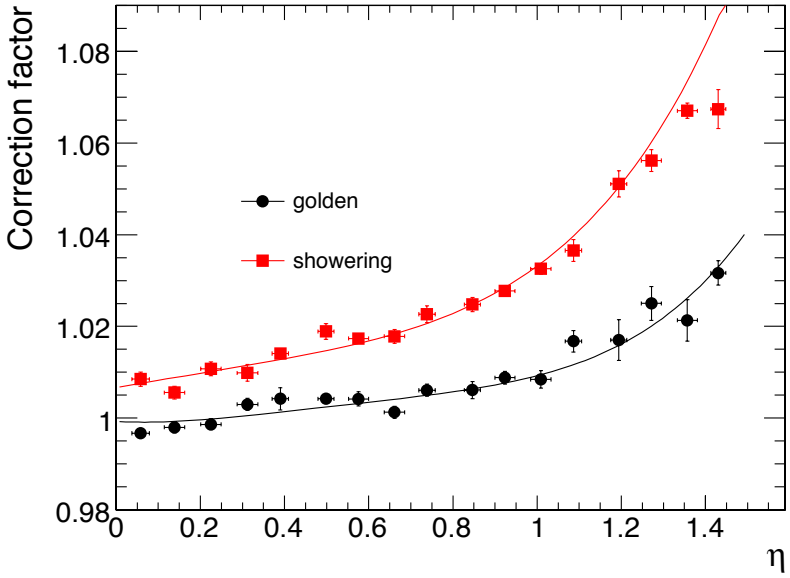


Figure 4.24: Correction $f(\eta)$ obtained using the Monte Carlo simulation compared to that obtained from 2 fb^{-1} of $Z \rightarrow ee$ with intercalibration precision at 2%.

4.4.4.2 Intercalibration of crystal rings

In a start-up scenario, where the algorithmic correction factors are taken from Monte Carlo simulation, $Z \rightarrow ee$ events might be used to obtain a preliminary estimate of the intercalibration factors between rings. The golden electron class from the classification described in Section 10.4 is chosen since these electrons radiate little, and their reconstructed energy shows the least dependence on the tracker material, and hence η. A 4% algorithmic correction is still needed to equalize the supercluster energy measurement over the length of the barrel. Only events with two golden electrons are selected.

The method has been tested taking the calorimeter regions as rings of crystals (at fixed η) in the ECAL barrel. There are 170 such rings of 360 crystals each. The miscalibration constants can be set up to have a different RMS variation between crystals within a ring to the variation of the average miscalibration value between rings. To reduce the sensitivity to tails, the peak of the mass distribution is used to obtain the intercalibration coefficients. The peak is obtained from an iterative Gaussian fit in a region $[-2.5\sigma, 2.5\sigma]$ around the peak. Crystals on the module borders are excluded.

The performance of the method was measured by comparing the applied miscalibration factors with the calibration coefficient determined by the method. The results given here are obtained when starting from a 5% miscalibration between rings and a 2% miscalibration between crystals within a ring. Using events corresponding to an integrated luminosity of 2.0 fb^{-1} the distribution of the residual miscalibration is shown in Fig. 4.25. 2The RMS spread of this distribution, corresponding to 0.6%, gives the achieved ring intercalibration precision.

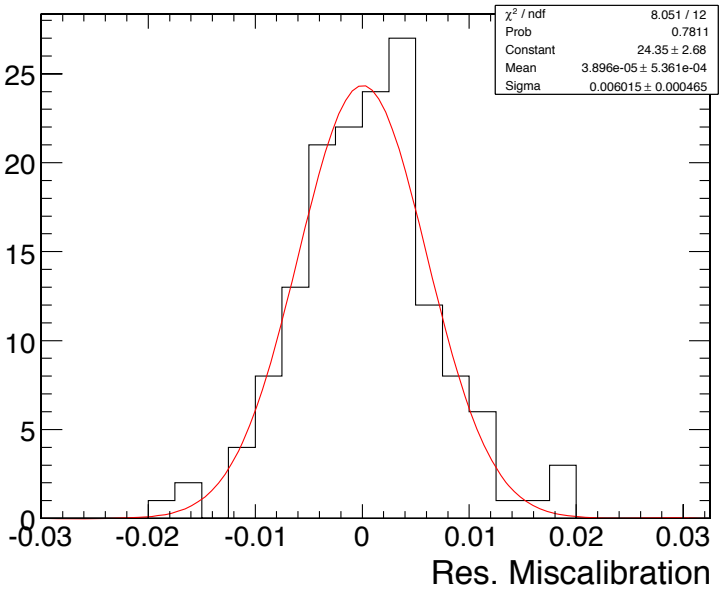


Figure 4.25: Distribution of the residual miscalibration using 2.0 fb^{-1}.

The ring intercalibration precision as a function of the number of events per ring is shown in Fig. 4.26. Table 4.3 shows the accuracy reached for each module type for an integrated luminosity of 2.0 fb^{-1}.

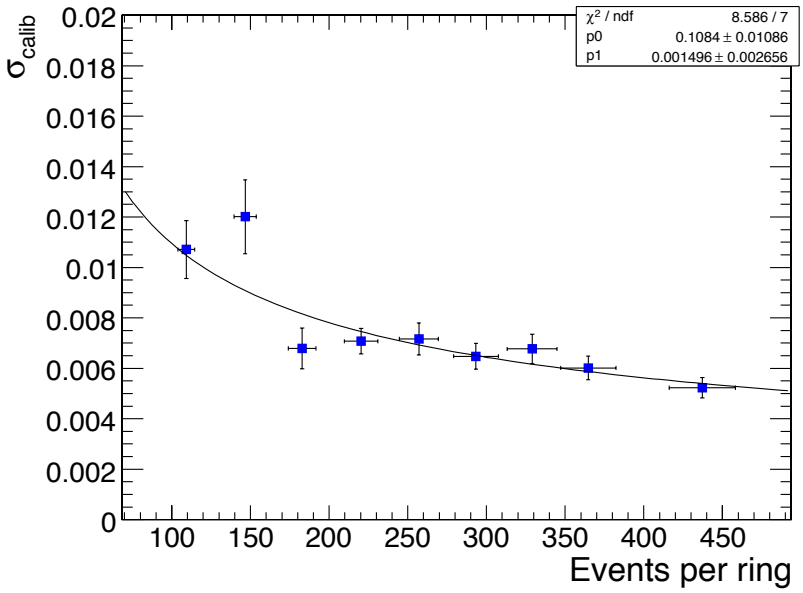


Figure 4.26: Ring-to-ring calibration precision achieved versus average number of events per ring. An average of 370 events per ring corresponds to an integrated luminosity of 2.0 fb^{-1}. The full curve is a fit to the sum of a statistical term and a constant term.

Table 4.3: Intercalibration precision between rings using a data sample equivalent to 2.0 fb^{-1}.

Module 1	$0.32\% \pm 0.05\%$
Module 2	$0.5\% \pm 0.1\%$
Module 3	$0.6\% \pm 0.1\%$
Module 4	$1.3\% \pm 0.4\%$

4.4.5 Intercalibration from π^0 and $\eta \to \gamma\gamma$ events

The possibility of intercalibrating the ECAL using mass reconstruction of π^0 and $\eta \to \gamma\gamma$ is being investigated. Experiments such as CLEO and L3 have successfully used low energy signals to calibrate crystal calorimeters to high accuracy. In the case of CMS, these low mass particles could provide an important additional calibration tool which is useful for

- relatively rapid intercalibration of all crystals;
- study of the effects of crystal transparency corrections from the laser monitor;
- and rapid check-out and monitoring of detector performance.

The intercalibration obtained from low-energy π^0s is not sensitive to tracker material as much as the intercalibration obtained from electrons is. First, π^0 decays into photons that do not convert are unaffected. Converted low-energy photons give rise to low-energy electrons, which reach the ECAL far from the expected photon impact point because of the magnetic field. As a consequence, a selection of π^0s based on the selection of pairs of closeby ECAL clusters retains mostly either unconverted photons or photons which convert just in front of the calorimeter. For this reason, the energy resolution does not deteriorate and no energy bias is introduced at high η. The only effect of the tracker material is a rate loss at larger η values.

The minimum separation for 2 photons in the barrel from a π^0 with E_{T} of 5 GeV is about 65 mm, corresponding to nearly 3 times the crystal granularity. Such π^0s occur in almost any kind of event. It has been shown that π^0s useful for calibration can be located within events using the ECAL Level-1 trigger information. They can then be reconstructed in a small region identified by the trigger. This operation would be similar to a normal first step of the High Level Trigger, requiring very little processing time to extract the small amount of information relevant for calibration.

To obtain an estimate of how many π^0 can be reconstructed with little background using data read out with all Level-1 triggers, a large data set of high E_{T} jet events was studied. It was found that about 1.4% of them have a usable π^0 in the barrel and that almost all of these are tagged by the isolated electron Level-1 trigger information. The selection applied to get this rate is rather simple. First, some fairly stringent shower shape cuts were applied to the individual photon candidates: $0.5 < R_1 < 0.9$, $R_4 > 0.2$ and $0.9 < R_9 < 1$, where R_n is the fraction of the cluster energy in a square array of n crystals centered on the cluster. Energy cuts were also made on the photon candidates. These cuts, $1.5 < E_\gamma < 5\,\mathrm{GeV}$, are slightly tightened at high η. To reduce the combinatorial background, it is important to only consider π^0 candidates with small opening angles. The cuts used correspond to photon separations at the ECAL front face between 60 and 90 mm.

The reconstructed mass of the candidates selected in this way is shown in Fig. 4.27 in 3 bins in η in the ECAL barrel. There is a π^0 mass peak with relatively little background in all 3 regions, and the mass resolution is about 8% in each case. Reconstructed π^0s are found coming from both the jet event and from pile-up events in the same bunch-crossing.

This result suggests that it will be possible to intercalibrate the ECAL using the reconstructed mass of π^0. With 1000 events per crystal, a statistical precision of 0.5% can be estimated for the intercalibration constants. This needs to be demonstrated, and sources of systematic error must be investigated. Such a demonstration is in progress but extremely large numbers of Monte Carlo simulated events are required to probe a precision of 0.5%.

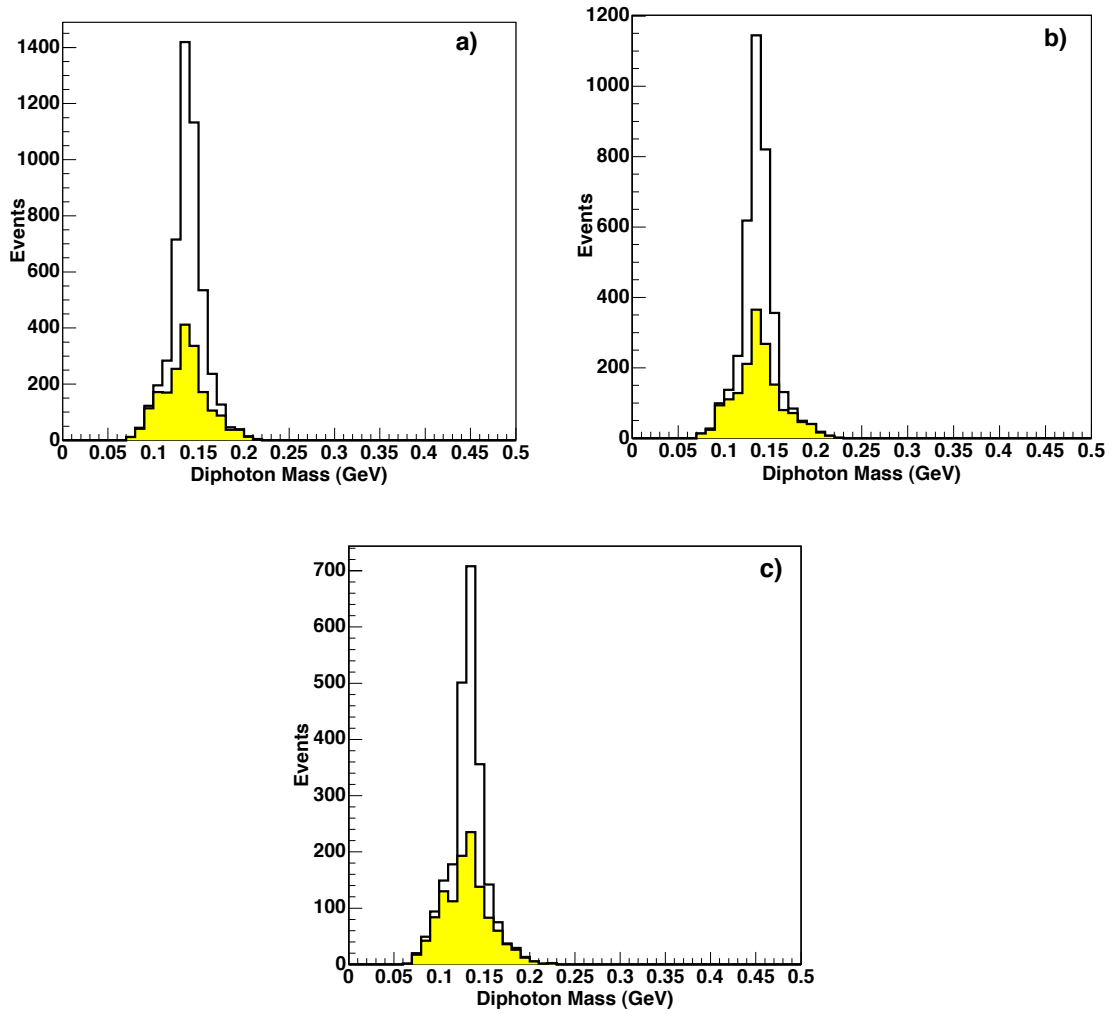


Figure 4.27: Reconstructed mass of π^0 candidates in 3 regions of the barrel: a) $|\eta| < 0.5$, b) $0.5 < |\eta| < 1.0$, and c) $|\eta| > 1.0$. The unshaded histograms represent the signal while the shaded ones the background.

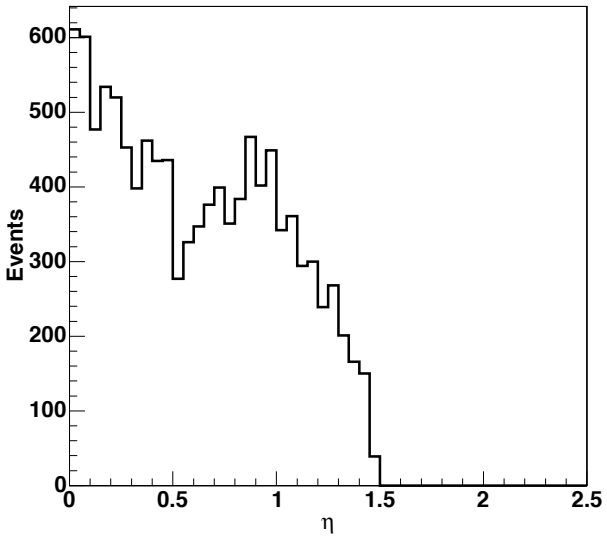


Figure 4.28: The η distribution of π^0 candidates in the barrel.

As can be seen from the number of π^0s in the 3 histograms of Fig. 4.27, the rate decreases with η, but there is still large rate in the last bin at the edge of the barrel. The η distribution of the π^0 candidates is shown in Fig. 4.28. With an assumed a Level-1 global trigger rate of 25 kHz, and with 1.4% of these triggers assumed to have usable π^0s, about 100 π^0s per crystal can be obtained in a running period of less than 5 hours.

Events from $\eta \rightarrow \gamma\gamma$ are also being studied. The signal has a much lower rate once the background is reduced sufficiently, but the mass resolution is about 3%. The η should be a useful calibration tool at higher energy, although it will take longer, and may prove very useful in the endcap.

4.4.6 Inner bremsstrahlung photons in Z boson decays to muons

A significant rate of high-p_{T} photons with very little background and an energy which can be known independently of the ECAL is available in radiative decays of $Z \rightarrow \mu\mu$. These photons are being investigated as a valuable tool for various calibration related tasks, as well as a probe for measuring photon reconstruction efficiency. They can be used, for example, to intercalibrate different regions of the ECAL (coefficient c_i of Equation 4.3), and to tune the various cluster correction algorithms (coefficient $\mathcal{F}$) and the overall energy scale (coefficient G). They can also be used to relate the energy scale of unconverted photons to that of electrons (from converted photons).

The available rates have been investigated using a sample of events generated using a full matrix element calculation of radiative decays. The background from Z bosons produced with additional jets has been studied also. Events are selected by placing a loose cut on the dimuon mass ($40 < M_{\mu\mu} < 80$ GeV/c^2) and searching for a photon with $p_{\mathrm{T}} > 15\,\mathrm{GeV}/c$ within a radius of $\Delta R < 0.8$ of either muon, where ΔR is defined as $\sqrt{\Delta\phi^2 + \Delta\eta^2}$. Figure 4.29 shows the signal and background rates as a function of photon p_{T}. The rate is calculated after a mass window of $87.2 < M < 95.2$ GeV/c^2 has been placed on the three-

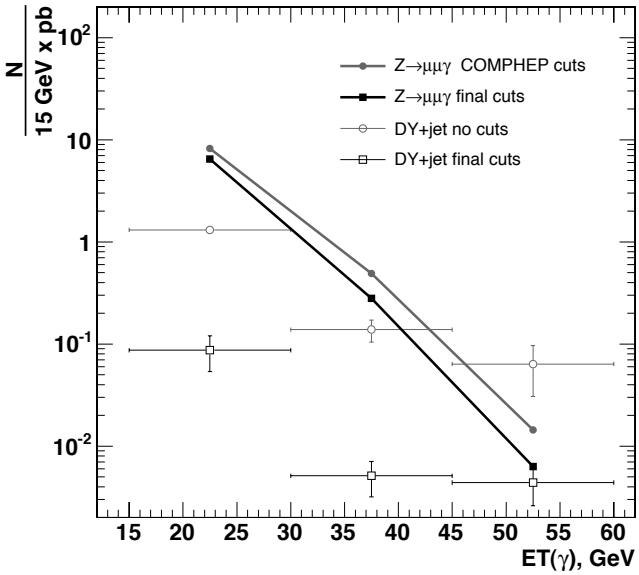


Figure 4.29: Signal and background yields, as a function of photon E_{T}, for three 15 GeV bins. The rates are shown before and after selection cuts.

body mass. With the selection cuts described, the signal-to-background ratio is ≈ 80 in the bin $15 < E_{\mathrm{T}}^{\gamma} < 30$ GeV, with a rate of 65 pb^{-1} over the region $|\eta| < 2.5$. For an integrated luminosity of only 1 fb^{-1}, an average of nearly 1 such photon per crystal will be collected. Therefore, a statistical precision better than 0.1% would be achieved with unconverted photons for the intercalibration of 10-crystal-wide rings of crystals (i.e. groups of 3600 crystals) in the barrel.

4.4.7 Preshower calibration

4.4.7.1 Overview

The energy deposited in the lead of the preshower will be estimated using the charge measurements made by 2 layers of silicon sensors. To preserve the overall photon energy resolution, a 5% charge measurement is necessary over a very large dynamic range (1–400 MIPs equivalent), resulting in a required strip calibration accuracy of a few per cent. There are 2 principle sources of response variation (sensor-to-sensor and channel-to-channel) at start-up: sensor thickness (RMS of $\approx 1\%$) and gain uniformity of the electronics (RMS $\approx 5\%$). Radiation damage to the sensors (decrease in charge collection efficiency by $\approx 17\%$ over 10 years) and the electronics (decrease in gain by $\approx 2\%$ over 10 years) will need to be assessed by periodic in-situ calibrations.

The principle difficulty in calibrating the preshower is that the only physics reference points are "zero" (zero output for zero input) and one MIP. The full dynamic range can only be explored systematically using an internal calibration circuit (ICC). Single MIPs with a S/N ratio suitable for calibration can only be measured when the front-end electronics is in high gain mode (HG). Consequently, the ICC must be used to transfer this information from HG to low gain (LG)—i.e. intercalibrate the 2 gains.

There are thus 2 parts to the full preshower calibration procedure:

- Absolute calibration using real MIPs
- Relative calibration between gains and between strips using an internal electronic charge generation system

Each of these parts is described in the sections below, followed by a summary of the procedures to execute before preshower installation in CMS, at start-up of CMS and subsequent periodic calibrations.

4.4.7.2 Determination of the MIP scale

The absolute calibration scale can be found by detecting single MIPs in HG mode. Prior to installation of the preshower in CMS, stacks of up to 12 ladders will be built, sandwiched between scintillators to allow triggering and detection of cosmic rays. The cosmic rate is approximately 1 per minute per cm^2 and are mostly muons with a mean energy of around 4 GeV [119] equating to 1 MIP per strip per minute. Running the ladders for 8–10 hours will thus give sufficient statistics per strip for an accurate absolute measurement of the MIP scale for each strip.

In CMS MIPs can be approximated with high energy muons and/or charged pions. The occupancy in the preshower is rather low, around 0.2% on average at low luminosity. Noise is thus the dominant "signal" in the strips, necessitating the use of the tracker and/or muon chambers to point to strips that have been traversed.

Studies of full jet and muon events at low luminosity have shown [155] that indeed both muons and charged pions can be used for calibration. Muon events are cleaner, but the muon chamber coverage only extends to $|\eta| < 2.4$. Pions are plentiful in both jet and minimum bias (i.e. pile-up) events and the tracker extends to $|\eta| < 2.5$.

The overall efficiency for muons and high p_{T} pions is around 15%, taking into account geometrical factors and tracking efficiencies. Using a combination of both types of event, a sensor-by-sensor calibration to the required few percent accuracy can be made in the order of a week, with the inter-strip calibration being achieved with the ICC. This takes into account the HLT rates and the LHC duty cycle. As the in-situ calibrations will be mostly following the drop in the charge collection efficiency (which in turn is a linear function of the accumulated radiation damage [156]), geometrical factors (forward-backward endcap symmetry; use of "rings" of sensors at similar η etc.) can be used to decrease the calibration time substantially.

These MIP calibrations must be taken with the front-end in HG mode. The dynamic range of the preshower will thus be reduced during the calibration periods.

4.4.7.3 The internal calibration circuit

The ICC inside the front-end electronics provides precise charge injections into the preamplifier, mimicking real signals. The amplitude of the injected signal is controlled by an on-board 8-bit DAC coupled with a switchable *precision*: high precision produces injection signals between -16 and 12 MIPs equivalent, in ≈ 0.1 MIP steps, whilst low precision allows signals between -120 and 500 MIPs in ≈ 2.4 MIP steps.

The ICC uses a voltage step applied to a fixed "injection" capacitor. Each of the 32 channels in a front-end board (corresponding to the 32 strips of the silicon sensor) has an individual

injection capacitor, with a variation in value significantly less than 1%.

The ICC can be used to find an intercalibration between the 2 gain modes and serves to intercalibrate the electronics channels within a single micromodule if necessary (e.g. to check the uniformity of the decrease in preamplifier gain with radiation damage). It is also used to explore the linearity of response of the electronics over the full dynamic range.

Before it can be used for these purposes, the ICC itself must be calibrated. There are several steps to calibrate the ICC.

- With the ICC set to high precision and the front-end in HG, the internal DAC is set to provide an output comparable to the real MIP signal. This, coupled with the zero point, gives a "MIPs per DAC step in high precision" value.
- The DAC is then set to a high value (e.g. 220) and the corresponding number of MIPs, M_1, measured.
- The ICC is switched to low precision and the DAC adjusted to provide again M_1 on the output. The "MIPs per DAC step in low precision" can then be estimated

Switching the front-end to LG mode then allows the intercalibration of the 2 gains.

The time required for the calibration of the ICC (including the intercalibration of the 2 gains) and the intercalibration of the strips should be of the order of a few hours maximum for the complete preshower and can be done between LHC fills.

4.4.7.4 The preshower calibration schedule

The absolute MIP calibration with cosmic rays, together with a first calibration between the 2 gains and the examination of the full dynamic range, will be done before preshower installation in 2007/2008.

A first in-situ calibration (with MIPs) will be performed during the preshower commissioning phase in 2008 at low luminosity. As mentioned above, a sensor-by-sensor calibration can be achieved in the order of a week. It is likely that this commissioning time will be coupled with that of the endcap, resulting in a longer period being available, and thus a strip-by-strip calibration of the preshower can be performed. In either scenario, this first in-situ calibration will be a check of the calibration with cosmic rays.

Subsequent in-situ calibrations should be performed at the beginning and end of each year, and are essentially to follow the effects of radiation damage. A few days maximum will allow a sufficiently accurate calibration to be achieved.

4.4.8 Correction for crystal transparency changes

4.4.8.1 Crystal behaviour under irradiation

Although radiation resistant, ECAL crystals show rapid loss of optical transmission under irradiation due to the production of colour centres within $PbWO_4$ which absorb a fraction of the transmitted light. This process is accompanied by transmission recovery due to self annealing processes, leading to cyclic transparency behavior between LHC collision runs and machine refills (Fig.4.30). The magnitude of the changes is dose-rate dependent, and is expected to range between 1 or 2 per cent at low luminosity in the barrel, to tens of per cent

in the high η regions of the endcap at high luminosity. The inter calibration of the crystals would be unacceptably degraded by these radiation induced transparency changes were they not measured and corrected for.

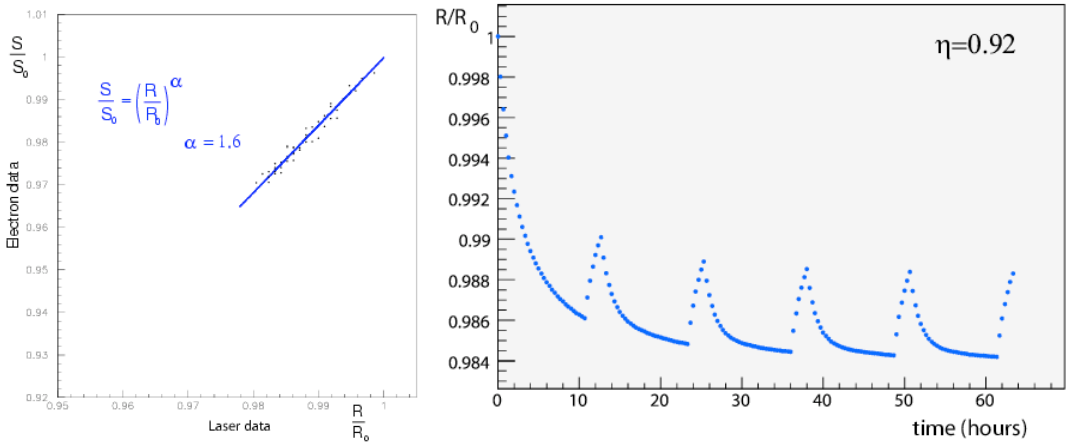


Figure 4.30: (a) Relation between the transmission losses for scintillation light and for laser light; (b) simulation of crystal transparency evolution at LHC based on test-beam results. For this illustrative example a luminosity of $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ was assumed, together with a machine cycle consisting of a 10 hour coast followed by 2 hours filling time. The crystal behaviour under irradiation was modeled on data taken during a crystal irradiation in the test beam.

The evolution of the crystal transparency is measured using laser pulses injected into the crystals via optical fibres. The response is normalized by the laser pulse magnitude measured using silicon PN diodes. Thus $R(t) = APD(t)/PN(t)$ is used as the measure of the crystal transparency. The laser monitoring system [5] performing this task is briefly outlined in the next section. Because of the different optical paths and spectra of the injected laser pulses and the scintillation light, the change in response to the laser light is not equal to the change in response to scintillation light. For attenuations $< 10\%$ the relationship between the changes can be expressed by a power law,

$$S(t)/S(t_0) = [R(t)/R(t_0)]^{\alpha}, \tag{4.8}$$

where $S(t)$ represents the response to scintillation light and α is characteristic of the crystal ($\alpha \approx 1.53$ for BCTP crystals, and $\alpha \approx 1.0$ for SIC crystals). This power law describes well the behaviour of all the crystals that have been evaluated in the test beam, and this formula is expected to be valid in the barrel for both low and high luminosity at LHC.

4.4.8.2 Laser-monitoring system overview

Figure 4.31 shows the basic components of the laser-monitoring system: Two laser wavelengths are used for the basic source. One, blue, at λ=440 nm, very close to the scintillation emission peak, which is used to follow the changes in transparency due to radiation, and the other, infra-red, at λ=796 nm far from the emission peak, hence very little affected by changes in transparency, which can be used to verify the stability of other elements in the system. The spectral contamination is less than 10^{-3}. The lasers are operated such that the full width at half maximum of the pulses is ≈ 30 ns. The lasers can be pulsed at a rate of ≈ 80 Hz, and the

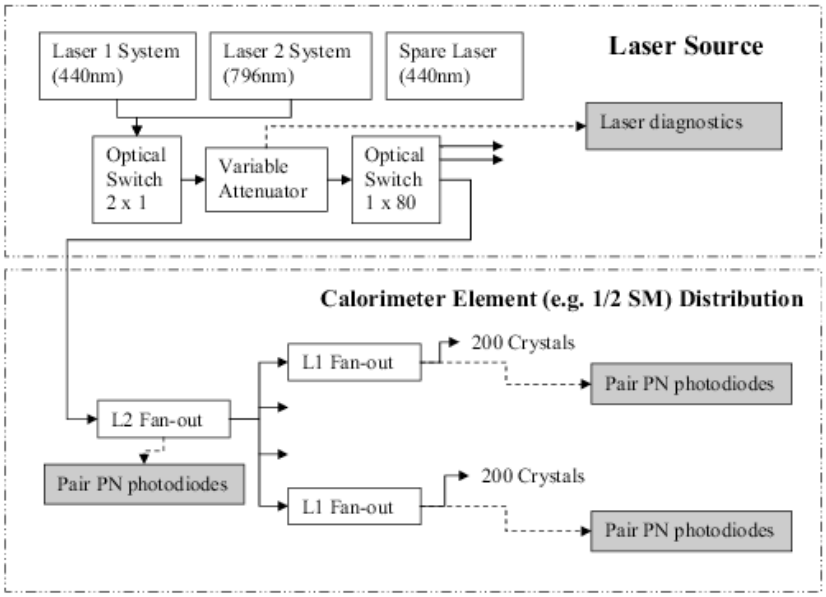


Figure 4.31: The components of the laser monitoring system.

pulse timing jitter is less than 3 ns to allow adequate trigger synchronization with the LHC bunch train and ECAL ADC clock.

The pulse energy of 1 mJ/pulse at the principal monitoring wavelength corresponds to ≈ 1.3 TeV, and a linear attenuator allows 1% steps down to 13 GeV. The pulse intensity instability is less than 10% which guarantees a monitoring precision of 0.1% by using the PN silicon photodiode normalization.

There are 3 light sources, 2 blue and 1 infrared. The duplication of the blue source provides fault tolerance and allows maintenance of one while the other is in use, ensuring that a source at the wavelength used to track changes in transparency is always available. Each source consists of an Nd:YLF pump laser, its power supply and cooler unit and corresponding transformer, a Ti:Sapphire laser and its controller, and a NESLAB cooler for an LBO crystal in the Ti:S laser. Each pair of the YLF and Ti:S lasers and their corresponding optics are mounted on an optical table. Each source has its own diagnostics, 2 fibre-optic switches, internal monitors and corresponding PC based controllers. Further details can be found in [157].

The monitoring light pulses are distributed via a system of optical fibres. A fibre optic switch at the laser directs the laser pulses to 1 of 88 calorimeter regions (72 half supermodules in the barrel and 8 regions in each endcap). A two-stage distribution system mounted on each calorimeter region delivers the light to each crystal.

To provide continuous monitoring, about 1% of the 3.17 μs beam gap in every 88.924 μs LHC beam cycle will be used to inject monitoring light pulses into crystals. The time needed to scan the entire ECAL is expected to be about 30 minutes.

4.4.8.3 Operational procedures

Each measurement of the crystal transparecy on 1 of the 80 calorimeter elements is envisaged to consist of 600 laser pulses (7.5 s at 80 Hz). The data will consist of the time frames read

out from the illuminated channels ($\approx$ 1000 channels), the time frames read out from the illuminated PN diodes, and the data read from the 2 Gs/s flash ADC. The latter samples the laser pulse viewed by a fast PIN diode located at the source to allow correction for variation of the laser pulse width.

Conceptually the monitoring and correction of changes in crystal transparency involves 3 steps:

- detailed data quality monitoring of the laser calibration data,
- the extraction of the $R(t) = APD(t)/PN(t)$ values from the laser data,
- the setting and tuning of the reference starting point $R(t_0)$ and the value(s) of α used for the crystals

It is intended that the first 2 steps will be performed online. The data will be rapidly processed in a small farm of processors to extract the mean value $R(t) = APD(t)/PN(t)$ for each measurement time, t. This processing is expected to involve correction for the non-linearity of the PN diode system, and may involve correction for variation of the laser pulse width.

The output of this processing will be a value of $R(t)$ for each crystal, measuring the crystal transparency, which will be stored in the online database. Since the scan over the entire ECAL is expected to take about 30 minutes, a new value for each crystal will be produced every 30 minutes. It is intended that this initial processing will be made sufficiently reliable and robust so that it will not need to be repeated. However, the raw data will also be stored to allow reprocessing if necessary.

4.4.8.4 Crystal transparency correction and systematic uncertainties

The transmission corrections may have systematic uncertainties, which result in intercalibration errors, which in turn contribute to the ECAL energy resolution. The magnitude of these contributions can be estimated based on experience from beam tests [158, 159]. The 4 main sources of error are 1) the dispersion of α between crystals, 2) the stability and reproducibility of the monitoring system itself, 3) laser pulsewidth changes, and 4) photodetector gain changes.

1. Dispersion of α
   The value of α is determined by the optical properties of the crystals, and hence by the details of the crystal production. Using data from the test beam, α has been measured to be 1.53 (1.0) with a relative uncertainty 5% (10%) for crystals produced by BCTP (SIC); Although the relative dispersion is larger for the SIC crystals, the absolute contribution to the systematic uncertainty is similar to the BCTP crystals. The magnitude of this uncertainty is proportional to the magnitude of the changes being corrected for, and is larger for a group of crystals under a higher or a wider range of radiation exposure rates. Unfortunately it will be not possible to measure the parameter α for all crystals before LHC operation. There is good reason to believe, however, that a direct evaluation for a group of crystals can be obtained using physics calibration events (for example, $W \rightarrow e\nu$) taken *in situ*.

2. Monitoring stability
   The light pulses injected into the crystals are simultaneously monitored by a local pair of PN photodiodes mounted close to the point of injection. This monitoring is done

for groups of 200 crystals. These reference photodiodes are stable, radiation hard and insensitive to the magnetic field. The laser monitoring stability is thus controlled by the stability of the relative calibration of injected laser light on each group of crystals. The redundancy of normalization (there are 10 PN photodiodes per supermodule) is used to detect relative miscalibration of PN reference between groups of crystals sharing the same normalization.

3. Laser pulsewidth
When illuminated with laser pulses from the monitoring system, the pulse shape at the

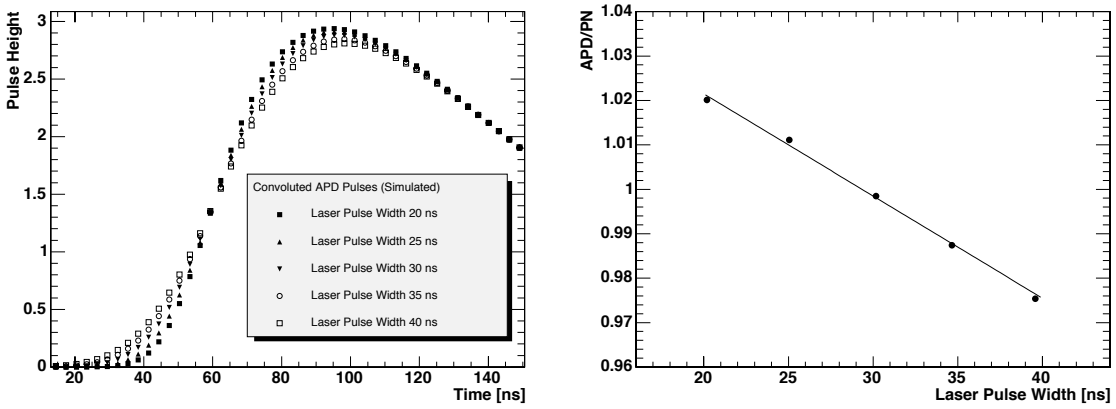


Figure 4.32: (a) Convolutions of the preamplifier shaping and Gaussian functions simulating the laser pulse widths of $20\,\mathrm{ns}$, $25\,\mathrm{ns}$, $30\,\mathrm{ns}$, $35\,\mathrm{ns}$ and $40\,\mathrm{ns}$ respectively. The PN diode response is stable due to its long shaping time. (b) Variation of the ratio APD/PN, which is used to monitor the change in the crystal transparency, as a function of the laser pulse width.

output of the APD and the PN readout electronics can be described as a convolution of the laser pulse shape and the shaping function of the respective readout systems. The laser pulse shape, measured in the test beam, is Gaussian and the electronics shaping is well described by the formula

$$A(t) = A_0 \left(\frac{x - t_0}{t_{\mathrm{rise}}} \right)^a e^{-a\left(\frac{x - t_0 + t_{\mathrm{rise}}}{t_{\mathrm{rise}}} \right)}, \tag{4.9}$$

where a and t_{rise} are measured on experimental pulse shapes. For the APD electronics $t_{\mathrm{rise}} \approx 50\,\mathrm{ns}$ and for the PN diode electronics $t_{\mathrm{rise}} \approx 800\,\mathrm{ns}$, so that the sensitivity to the laser pulse width is very different between the two.

To investigate the impact of the variation of the laser pulse width this convolution has been simulated for different laser pulse widths. As can be seen in Fig. 4.32a, the convolution results in a significant variation of the maximum height of the signal in the APD. The effect in the PN is negligible due to the long shaping time. Thus the crystal transparency measurement is very sensitive to changes of the laser pulse width as shown in Fig.4.32b. With a laser pulse width of about 30 ns, the sensitivity of the ratio APD/PN is 0.23%/ns.

This effect can be corrected by measuring the laser pulse width event by event with a precision better than 1 ns using the flash ADC at 2 Gs/sec.

4. APD gain
A change in the gain of any part of the electronics chain will affect equally the scintil-

lation response and the response to laser pulses. Since the constant of proportionality, α, between the magnitude of response variation of the laser signal and the scintillation signal is not 1, any such change of gain will lead to systematic errors in the transparency correction if not independently monitored and corrected for. The most significant change of gain expected is for the APD, where the gain is directly related to the high voltage applied. This depends on the HV power supply and the dark current of the APD. Such effects can be controlled by the monitoring of the high voltage and the dark current of each APD (the LSB for this measurement is 322 nA which is sufficient to control the APD gain with a precision better than 0.1%). Such systematic contributions can also be identified by comparing the crystal response to blue and infrared laser light.

4.4.9 Alignment

The ECAL is capable of providing quite precise measurements of shower position, and a resolution of about 0.5 mm is achievable for very high energy showers [160]. Because of the 3° off-pointing angle of the crystal axes the shower-to-shower variation of shower depth results in an uncertainty in lateral shower position approaching 0.5 mm. This sets a limit to the useful precision of ECAL alignment with respect to the tracker.

The key step of the electron HLT selection matches shower position to pixel detector hits. The full width of the selected region is $\Delta\phi \approx 0.04$ rad, corresponding to about 60 mm in the ECAL, and thus places only very mild demands on the alignment precision required.

Precision measurements made on a number of supermodules during construction show that the variations of the crystal lateral positions with respect to their nominal positions within the supermodule are only a few hundred microns. It thus seems likely that the nominal construction geometry can be taken for all relative positions within a supermodule. Similarly, it is expected that nominal positions will be adequately precise within endcap Dees. Alignment for the ECAL then consists of measuring the position of 36 supermodules and 4 Dees with respect to the tracker. The initial placement errors are not expected to be larger than a few mm. These errors will be rapidly reduced using track/cluster matching of electrons.

4.4.10 Conditions/calibration database

The conditions database stores data which relates to and describes the condition of the detector. For the ECAL, it will contain data from many sources. Much effort has been put into listing and summarizing the type and amount of ECAL data in the conditions database, and it has been estimated that about 1 GB per day of such data will be produced. The Detector Control System (DCS [8]) will store the relevant values of the High Voltages, Low Voltages, Temperatures and other environment sensors and the ECAL Safety System. The temperature sensors which are located on the crystal and the near the photodetectors will be read via the Token Ring, as well as all other temperature sensors on the Very Front End cards (VFE) and Low Voltage Regulator Boards (LVRB) and the APD dark currents. These data will go in the condition database to track changes and produce history plots. The ECAL OD and DAQ will monitor the status of the read-out and will periodically produce summary information that will also be stored to track problems in the read-out and trigger system.

The standard reconstruction requires a few parameters per channel (of the order of 10–20). Many of these parameters have to do with the electronics chain and are unlikely to change

often. Some parameters (like the pedestals) may need to be monitored on a daily basis. The changes in crystal transparency, monitored using the laser system, will require a new value for each crystal every 30 minutes, and interpolation may be needed between 2 consecutive values. Of all values used in the reconstruction, the latter are expected to have the shortest interval of validity.

The detector control system monitors the stability of environment variables (like temperature, high voltage, low voltage, etc.) to verify that they remain constant to a precision which guarantees a stable response. From beam-test experience, it can be expected that these variables will not change significantly over a short time-scale. Thus it is expected that corrections will not need to be applied. If it is necessary the environment variables data will be copied from the online to the offline database, otherwise a simple channel status indicator can be computed to summarize the stable status of each channel.

4.5 Monitoring

In this section the different elements which contribute to the ECAL monitoring are described: the detector monitoring and the data quality monitoring. The task of the former is to continuously inspect the operational conditions of the hardware, and that of the latter is to check the detector performances on different categories of collected events. They will allow the detection of either hardware problems or software configuration errors. Any unexpected behaviour needs to be identified and investigated rapidly.

4.5.1 Monitoring of operational conditions

4.5.1.1 ECAL detector and safety control system

The ECAL detector control system (DCS) monitors the detector status, in particular various environmental parameters, as well as the detector safety system, which will generate alarms and hardwired interlocks in situations which could lead to damage of the detector hardware. It consists of the following subsystems: Precision Temperature Monitoring (PTM), Humidity Monitoring (HM), ECAL Safety System (ESS), High Voltage (HV), Low Voltage (LV) and monitoring of the laser operation, the cooling system and of the parameters (temperatures in capsules, temperatures on the printed circuit boards, APD leakage currents) read out by the DCUs on the VFE and LVRBs.

A major aspect for the ECAL detector control is the monitoring of the system temperature and the verification that the required temperature stability of the crystal volume and the APDs, expected to be $(18 \pm 0.05)^{\circ}$C, is achieved. The PTM is designed to read out thermistors, placed on both the front and back of the crystals, with a relative precision better than 0.01°C. In total there are ten sensors per supermodule. Two immersion probes measure the temperature of the incoming and outgoing cooling water, whereas two sensors per module, one on the grid and one on the thermal screen side of the crystal volume, monitor the crystal temperature. The readout is based on the Embedded Local Monitoring Board (ELMB) developed by ATLAS and completely independent of the DAQ and control links. However, the latter are used to read out further thermistors, mounted in the capsules, via the DCUs placed on the VFE boards.

The purpose of the ESS is to monitor 1) the air temperature of the VFE and FE environment (expected to be around 25–30°C) and 2) the water leakage detection cable, which is routed inside the electronics compartment, 3) to control the proper functioning of the cooling system and 4) to automatically perform pre-defined safety actions and generate interlocks in case of any alarm situation. One pair of temperature sensors is placed at the centre of each module. The completely redundant readout system is independent of the DAQ and control links and based on a Programmable Logic Controller (PLC) situated in the underground service cavern (USC55). It evaluates the data from the temperature sensors, from the water leakage detection system and the signals received from the cooling system, which has its own PLC. If any critical value is detected hardwired interlock signals will be routed to the relevant crates in order to switch off the LV and/or to the cooling PLC in order to stop the water flow on a certain cooling line. The proper functioning of the ESS PLC itself is monitored by the general CMS detector safety system. In addition to this very robust, independent and continuously running system, the temperature of the VFE and LVRBs can be monitored with the DCUs placed on them. These data are read out via the control link and the local DAQ. They will not be used to trigger automatic actions based on hardwired links, but warnings and alarms might be triggered through software for notification of the operators. This is also true for the HM system, which monitors the humidity level in the electronics compartment and consists of one humidity sensor per module and an ELMB-based readout.

The whole DCS software is based on the commercial SCADA package PVSS II. A distributed system is built out of several applications dedicated to the PTM, HM, ESS, the monitoring of the status of the cooling system and the laser operation, as well as the monitoring and control of the HV and LV systems. The DCU data read out via the local DAQ will be written directly to the conditions database, but can be retrieved with user interfaces based on PVSS and/or ROOT. Every application is implemented as a Finite State Machine (FSM) and linked to a supervisory level, which summarizes the overall ECAL DCS status and itself incorporates a FSM. Finally, this ECAL DCS supervisor is linked to the general CMS DCS supervisory node, in order to communicate the status and alarms and to receive commands which are propagated down to the relevant subsystems. In general it is the CMS DCS supervisor which will communicate with the CMS Run Control, but it is possible that the Run Control and/or the local ECAL DAQ system communicate with the ECAL DCS supervisor via the exchange of SOAP (simple object access protocol) messages.

4.5.1.2 Temperature

The number of scintillation photons emitted by the crystals and the amplification of the APD are both temperature dependent. Both variations are negative when the temperature increases and in both cases the magnitude is measured to be about $-2\%/°\mathrm{C}$ [161]. Much effort has been put into ensuring the thermal stability of the ECAL, and measurements indicate that no significant or measurable variation of response due to temperature variation is to be expected during ECAL operation [161]. During operation numerous temperature probes will monitor the temperature of the ECAL, and, in particular, precision measurements from sensors fixed to the back surface of every tenth crystal in the barrel, and one in 25 crystals in the endcap, will be made frequently and regularly.

4.5.1.3 Dark current

The APD dark current will increase during CMS operation due to bulk damage of the silicon structure by neutrons. Part of this damage anneals, but the overall effect will be an increase of electronics noise, due to an increasing dark current, over the lifetime of the detector. The dark current of all APD channels will be continuously monitored.

4.5.1.4 HV and LV

The APD gain dependence on high voltage is $\approx 3\%/\mathrm{V}$ at the chosen operating gain (50). The High Voltage system has been designed and qualified to maintain the voltage stable over about 30 days with a maximum deviation of 66 mV. This limits the effect on the calibration to below 0.2%. Such stability is obtained using sense wires that measure the voltage at the load, correcting for instability or varying voltage drop over the cables from the HV crates to the experiment. The high voltage boards will require periodic calibration in order to insure the stability over longer periods. The DCS system will monitor the voltage and the current delivered by each channel.

The ECAL amplification and digitization electronics located on the VFE electronics cards require a very stable low voltage to maintain constant signal amplification. The system uses Low Voltage Regulators that guarantee the required stability of the signal amplification. The Low Voltage Regulators Boards are equipped with DCUs that measure the voltages and these measurements are read via the Token Ring.

4.5.2 Monitoring using events

The purpose of the ECAL Data Quality Monitor is to follow the time evolution of relevant quantities describing the running conditions of the detector and its performance. The aim is twofold, first, to spot possible problems or system instabilities, and promptly correct for them during th data-taking, and second to monitor the high-level quantities that are needed by the offline reconstruction. The data to be monitored consist not only of standard physics trigger data, but also of special ECAL-only triggers acquired during the normal physics data-taking by triggering during the LHC machine orbit gap. Some of these special triggers will need to be taken outside of normal physics data-taking periods, and the stand-alone VME readout option may be needed for this purpose. These special ECAL-only triggers are

- laser triggers;
- test-pulse triggers;
- empty bunch crossing triggers.

Test-pulse triggers, and empty bunch triggers used to measure the pedestal in the gain ranges other than the highest, require special settings to be downloaded to the FE. So also do laser triggers with the MGPA gain range forced, which will be needed in order to allow the study of the relative gains of the different gain ranges. In normal circumstances special settings will not be downloaded to the FE during physics data taking periods.

The histograms of the "monitorable quantities" will all be made available through a common DQM histogram framework using the Physics and Data Quality Monitoring package (Section 2.9), and the histograms and the results of the monitoring will be available also offline from the ECAL Conditions Database (Section 4.4). A prototype system has already been

commissioned and will be tested at the end of 2005. It will include a *source* and a first set of *clients* as defined in Section 2.9.2.3. They will allow to visualize the histograms, to save them to ordinary ROOT files, and to save the relevant informations to the conditions database for future needs.

From the experience gained to date, they can be classified in different categories:

- Raw data quantities: they include error flags from the DAQ, bad synchronization of the different FE streams, problems in the decoding of the raw data, data integrity, correct read-out order of towers and active channels, reasonable read-out values.
- Single channel raw quantities: the number and location of dead, noisy and in general malfunctioning channels will need to be monitored. This category includes the monitoring of
  - signal pulse phase: the time of signal pulse maximum ($t_{\max}$) needs to be monitored for both physics and laser events.
  - pedestal events: the mean values and widths of the pedestals will be monitored, for example using trace plots. Given the large number of channels, it will be useful to monitor the average deviation of pedestals and widths of all channels, comparing them to the last high statistics pedestal run taken before the current fill. For some channels, correlations between samples and between channels will be computed and monitored.
  - laser events: the response of each channel to the injection of a laser light (normalized to the laser light detected by reference PN diodes) will be monitored as a function of time.
- Global raw quantities: the operation of selective readout will need to be monitored and validated. The monitoring of the zero suppression algorithms will be possible using a map of active channels, averaged over several events.
- Reconstructed quantities: after pedestal subtraction, gain correction, and energy reconstruction, it will be possible to identify problematic channels using a map of the average energy deposition.
- Physics quantities: several useful higher level quantities will need to be monitored in standard physics triggers. A short and incomplete list includes:
  - average number of crystals above a given threshold per event
  - average number of energy clusters per supermodule
  - average number of crystals per energy cluster
  - average number of energy clusters per event
  - raw/calibrated energy per trigger tower
  - average energy per event in (phi, eta) bins
  - average energy per event
  - pedestal values (only pedestals in the lowest gain range are available in normal running conditions)

The final set of quantities to be monitored is expected to be fully verified, tested and possibly further extended during the ECAL precalibration with cosmic muons, the magnet test, and the test beam data taking in 2006.

Chapter 5

Hadron Calorimeter

The hadron calorimeters in conjuction with the ECAL subdetectors form a complete calorimetry system for the measurement of jets and missing transverse energy [3]. The central barrel and endcap HCAL subdetectors completely surround the ECAL and are fully immersed within the high magnetic field of the solenoid (Fig. CP 1). The barrel (HB) and endcap (HE) are joined hermetically with the barrel extending out to $|\eta| = 1.4$ and the endcap covering the overlapping range $1.3 < |\eta| < 3.0$. The forward calorimeters are located 11.2 m from the interaction point and extend the pseudorapidity coverage overalapping with the endcap from $|\eta| = 2.9$ down to $|\eta| = 5$. The forward calorimeters (HF) are specifically designed to measure energetic forward jets optimized to discriminate the narrow lateral shower profile and to increase the hermeticity of the missing transverse energy measurement. Central shower containment in the region $|\eta| < 1.26$ is improved with an array of scintillators located outside the magnet in the outer barrel hadronic calorimeter (HO).

5.1 HCAL design

The barrel hadron calorimeter is an assembly of two half barrels, each composed of 18 identical 20° wedges in ϕ (Fig. CP 8). The wedge is composed of flat brass alloy [1] absorber plates parallel to the beam axis. The innermost and outermost absorber layers are made of stainless steel for structural strength. There are 17 active plastic scintillator tiles interspersed between the stainless steel and brass absorber plates. The first active layer is situated directly behind the ECAL. This layer has roughly double the scintillator thickness to actively sample low energy showering particles from support material between the ECAL and HCAL. The longitudinal profile in the barrel going from an inner radius of 1777 mm to an outer radius of 2876.5 mm is given by

- (Layer 0) 9 mm Scint/61 mm Stainless Steel
- (Layers 1-8) 3.7 mm Scint/50.5 mm Brass
- (Layers 9-14) 3.7 mm Scint/56.5 mm Brass
- (Layers 15+16) 3.7 mm Scint/75 mm Stainless Steel/9 mm Scint

where the layer number refers to the active scintillator layer. The individual tiles of scintillator are machined to a size of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ and instrumented with a single wavelength shifting fiber (WLS). The WLS fibers are spliced to clear fibers, and the clear fibers

[1] Cartidge brass # 260 with 70% Copper and 30% Zinc

are run down the length of the half-barrel where they are optically added to corresponding projective tiles from each of the 17 active layers, thus forming 32 barrel HCAL "towers" in η (Tab. 5.1). The exceptions are towers 15 and 16 located at the edge of the HB half-barrel where multiple optical readouts are present, as shown in Figure 5.1. The optical signal from the HCAL towers is detected with a pixelated hybrid photodiode (HPD) mounted at the ends of the barrel mechanical structure. An additional layer of scintillators, the outer hadron calorimeter (HO), is placed outside of the solenoid and has a matching $\Delta\eta \times \Delta\phi$ projective geometry with a separate optical readout. Specific details of the geometry can be found in [162].

The endcap hadron calorimeter is tapered to interlock with the barrel calorimeter and to overlap with tower 16, as shown in Figure 5.1. The HE is composed entirely of brass absorber plates in an 18-fold ϕ-geometry matching that of the barrel calorimeter. The thickness of the plates is 78 mm while the scintillator thickness is 3.7 mm, hence reducing the sampling fraction. There are 19 active plastic scintillator layers. In the high η-region, above $|\eta| = 1.74$, the ϕ-granularity of the tiles is reduced to 10° to accommodate the bending radius of the WLS fiber readout, as shown in Figure 5.2. For the purpose of uniform segmentation in the Level-1 calorimeter trigger, the energies measured in the 10° ϕ-wedges are artificially divided into equal shares and sent separately to the trigger. The $\Delta\eta \times \Delta\phi$ tower size matches that of the barrel in the range $1.3 < |\eta| < 1.74$. For $|\eta| > 1.74$, the η size increases as shown in Table 5.1. The number of depth segments in the HE includes a pseudo-EM compartment starting with tower 18, the first tower beyond the η coverage of the ECAL barrel. During startup, the first depth segment of the HE will be used to feed the regional calorimeter trigger (RCT) in place of the ECAL endcap signals. The rear compartments will form the hadronic energy inputs for the RCT. From Table 5.1, the size of rear compartment of tower 28 is unusually large with $\Delta\eta = 0.35$. In the precision readout the two front compartments of this tower are split in η and readout separately to provide finer granularity, whereas the front compartments are combined in the trigger readout.

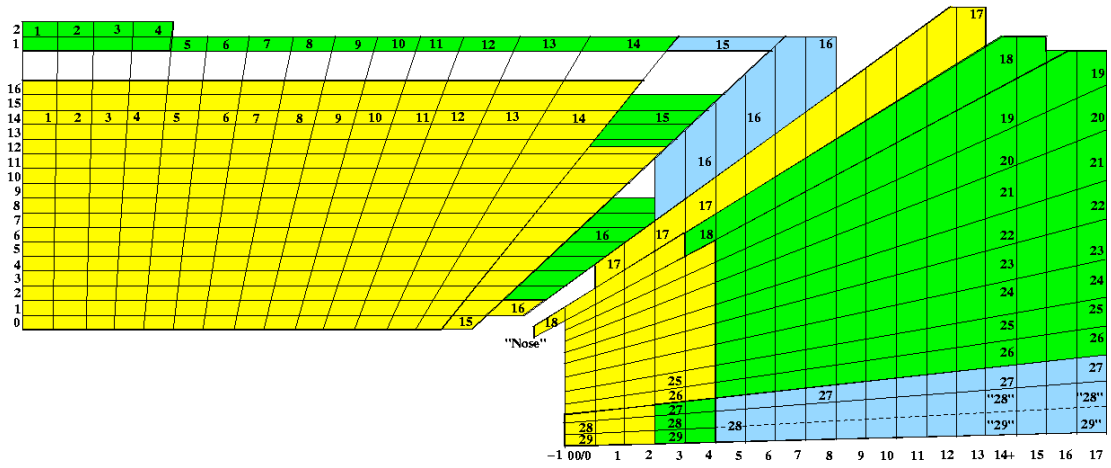


Figure 5.1: A schematic view of the tower mapping in r-z of the HCAL barrel and endcap regions.

The outer barrel hadron calorimeter consists of layers of scintillator located outside of the magnet coil. Since these are located within the return yoke along with the barrel muon detector, the segmentation of these detectors closely follows that of the barrel muon system.

Tower index	η range Low	η range High	Detector	Size η	Size ϕ	Depth segments
1	0.000	0.087	HB, HO	0.087	5°	HB=1, HO=1
2	0.087	0.174	HB, HO	0.087	5°	HB=1, HO=1
3	0.174	0.261	HB, HO	0.087	5°	HB=1, HO=1
4	0.261	0.348	HB, HO	0.087	5°	HB=1, HO=1
5	0.348	0.435	HB, HO	0.087	5°	HB=1, HO=1
6	0.435	0.522	HB, HO	0.087	5°	HB=1, HO=1
7	0.522	0.609	HB, HO	0.087	5°	HB=1, HO=1
8	0.609	0.696	HB, HO	0.087	5°	HB=1, HO=1
9	0.696	0.783	HB, HO	0.087	5°	HB=1, HO=1
10	0.783	0.870	HB, HO	0.087	5°	HB=1, HO=1
11	0.879	0.957	HB, HO	0.087	5°	HB=1, HO=1
12	0.957	1.044	HB, HO	0.087	5°	HB=1, HO=1
13	1.044	1.131	HB, HO	0.087	5°	HB=1, HO=1
14	1.131	1.218	HB, HO	0.087	5°	HB=1, HO=1
15	1.218	1.305	HB, HO	0.087	5°	HB=2, HO=1
16	1.305	1.392	HB, HE	0.087	5°	HB=2, HE=1
17	1.392	1.479	HE	0.087	5°	HE=1
18	1.479	1.566	HE	0.087	5°	HE=2
19	1.566	1.653	HE	0.087	5°	HE=2
20	1.653	1.740	HE	0.087	5°	HE=2
21	1.740	1.830	HE	0.090	10°	HE=2
22	1.830	1.930	HE	0.100	10°	HE=2
23	1.930	2.043	HE	0.113	10°	HE=2
24	2.043	2.172	HE	0.129	10°	HE=2
25	2.172	2.322	HE	0.150	10°	HE=2
26	2.322	2.500	HE	0.178	10°	HE=2
27	2.500	2.650	HE	0.150	10°	HE=3
*28	2.650	3.000	HE	0.350	10°	HE=3
29	2.853	2.964	HF	0.111	10°	HF=2
30	2.964	3.139	HF	0.175	10°	HF=2
31	3.139	3.314	HF	0.175	10°	HF=2
32	3.314	3.489	HF	0.175	10°	HF=2
33	3.489	3.664	HF	0.175	10°	HF=2
34	3.664	3.839	HF	0.175	10°	HF=2
35	3.839	4.013	HF	0.174	10°	HF=2
36	4.013	4.191	HF	0.178	10°	HF=2
37	4.191	4.363	HF	0.172	10°	HF=2
38	4.363	4.538	HF	0.175	10°	HF=2
39	4.538	4.716	HF	0.178	10°	HF=2
40	4.716	4.889	HF	0.173	20°	HF=2
41	4.889	5.191	HF	0.302	20°	HF=2

Table 5.1: Sizes of the HCAL readout towers in η and ϕ as well as the segmentation in depth. The HF has a non-pointing geometry, and therefore the tower η ranges provided here correspond to $|z| = 11.2$ m. *The first two depth segments of tower 28 have a finer η segmentation, divided at $|\eta| = 2.868$.

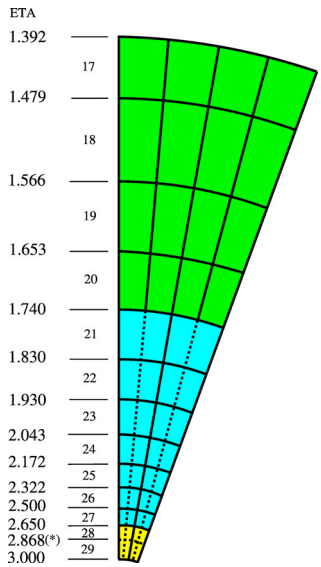


Figure 5.2: The η-ϕ view of a 20° HE endcap section showing the 5° regions and "split" 10° regions above $\eta = 1.740$ in detector pseudo-rapidity. The tower 28/29 split in η is also shown.

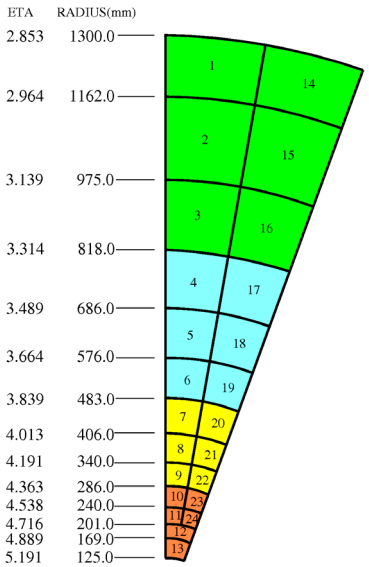


Figure 5.3: The r-ϕ view of an HF wedge (η at $z = 11.2$ m). The shaded regions correspond to the level-1 trigger sums.

The entire assembly is divided into 5 rings (2.54 m wide along the z-axis) each having 12 sectors. Figure 5.4 shows one such panel being mounted on a ring. The central ring (ring 0) has two layers of 10 mm thick scintillators on either side of the 'tail catcher' iron (18 cm thick) at radial distances of 3850 mm and 4097 mm, respectively. All other rings have a single layer at a radial distance of 4097 mm. The panels in the 12 sectors are identical except those in rings ± 1. This is due to the chimney structure in the magnet. To accommodate this structure, special panels were built with a single row of scintillator tiles removed. The HO covers $|\eta| < 1.26$ with the exception of the space between successive muon rings in the η direction, the space occupied by 75 mm stainless steel support beams separating the 12 layers in ϕ and the chimney structure. The inclusion of the HO layers extends the total depth of the calorimeter system to a minimum of $11\lambda_{int}$ for $|\eta| < 1.26$, as shown in Figure 1.3.

The forward calorimeters are located 11.2 m from the interaction point. They are made of steel absorbers and embedded radiation hard quartz fibers, which provide a fast collection of Cherenkov light. Each HF module is constructed of 18 wedges in a non-projective geometry with the quartz fibers running parallel to the beam axis along the length of the iron absorbers. Long (1.65 m) and short (1.43 m) quartz fibers are placed alternately with a separation of 5 mm. These fibers are bundled at the back of the detector and are readout separately with phototubes. The η ranges of the cells are listed in Table 5.1 at a distance of 11.2 m from the interaction point. The r-ϕ view of an HF wedge is shown in Figure 5.3.

The overall segmentation of HCAL is illustrated in a lego plot of a simulated multi-jet event (Fig. 5.5) which gives a sense of the jet structure available with this segmentation.

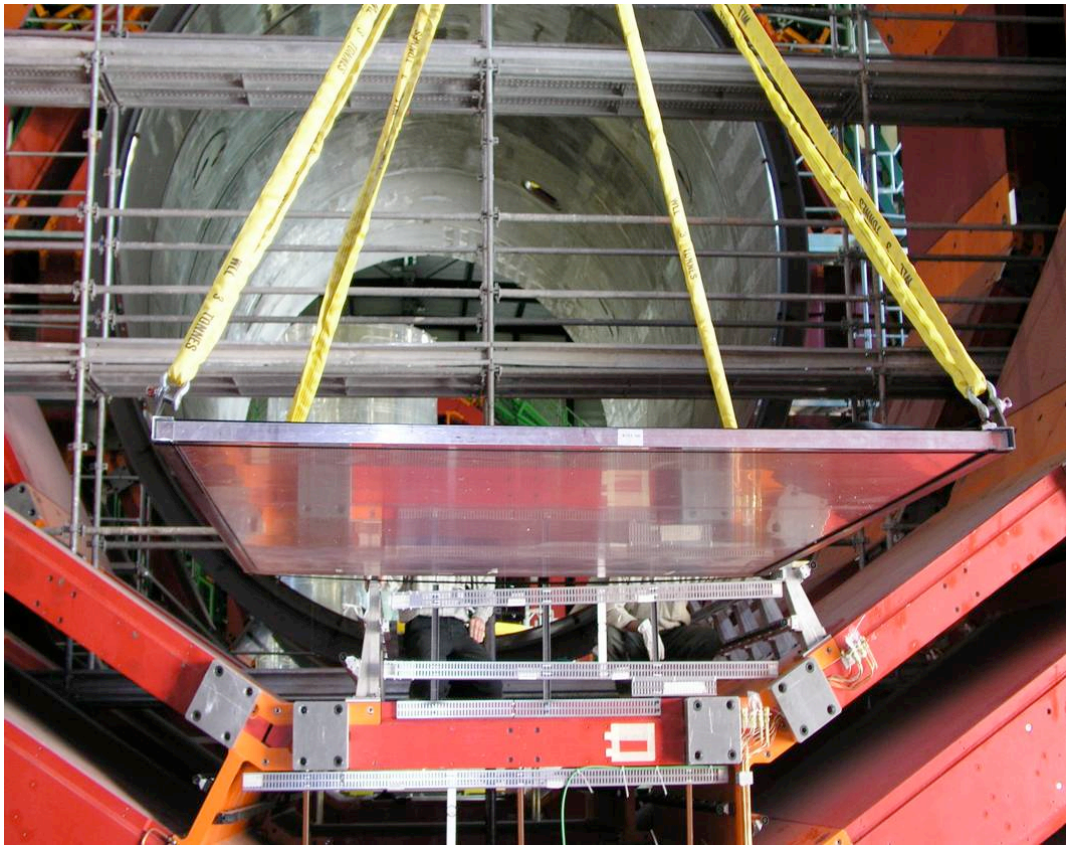


Figure 5.4: A fully assembled panel of the outer hadron calorimeter being mounted on one of the outer rings(ring +2) of the magnet yoke.

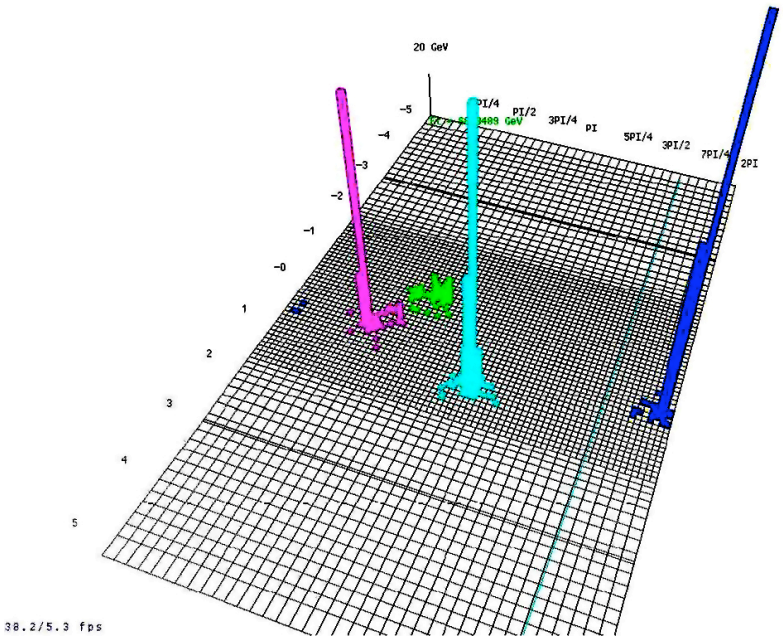


Figure 5.5: Lego plot of a multi-jet event showing the η, ϕ segmentation of HCAL.

5.2 Detector operation

There is extensive experience in operating HCAL, beginning in 2002 [17] with the testing of production modules in a CERN test beam and continuing with additional test beams in 2003 and 2004, and continuing operations at the surface hall SX5 beginning in 2004. The test beams in 2004 included simulated LHC operation with a 25 ns structured beam [163].

The overview of the HCAL controls is given in Fig. 5.6. Several PCs in the CMS control room operated through PVSS are used to control high and low voltages. The control system also downloads pedestal DAC and timing parameters to front-ends and controls many of the calibration and monitoring systems including the source calibration drivers, the LED pulsers, and the laser system each of which is separately in Section 5.3.

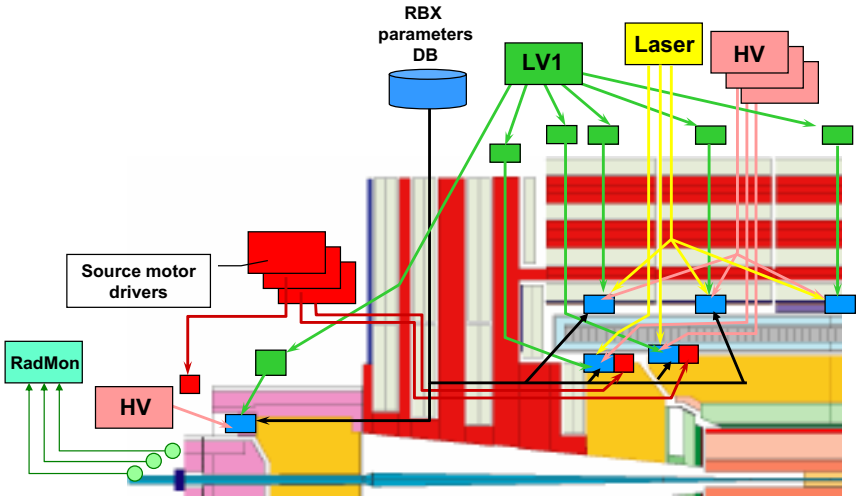


Figure 5.6: Overview of HCAL detector controls.

5.2.1 Data path and event format

The overview of the HCAL readout is shown in Fig. 5.7. The readout consists of an optical to electrical transducer followed by a fast charge-integrating ADC. The digital output of the ADC is transmitted for every bunch over a gigabit digital optical fiber to the counting house. In the counting house, the signal is deserialized and used to construct trigger primitives which are sent to the calorimeter trigger. The data and trigger primitives are also pipelined for transmission to the DAQ upon a Level-1 Accept (L1A) decision.

The optical signals from the scintillator-based detectors (HB, HE, and HO) are converted to electrical signals using multichannel hybrid photodiodes (HPDs) which provide a gain of ≈ 2000. In the forward calorimeter, where the magnetic fields are much smaller than in the central detector, conventional photomultiplier tubes (Hamamatsu R7525HA) are used.

The analogue signal from the HPD or photomultiplier is converted to a digital signal by a charge-integrating ADC ASIC called the QIE (Charge-Integrator and Encoder). The QIE internally contains four capacitors which are connected in turn to the input, one during each 25 ns period. The integrated charge from the capacitors is converted to a seven-bit non-linear scale to cover the large dynamic range of the detector. The ADC is designed so its contribution to the detector energy resolution over its multi-range operation is negligible, as

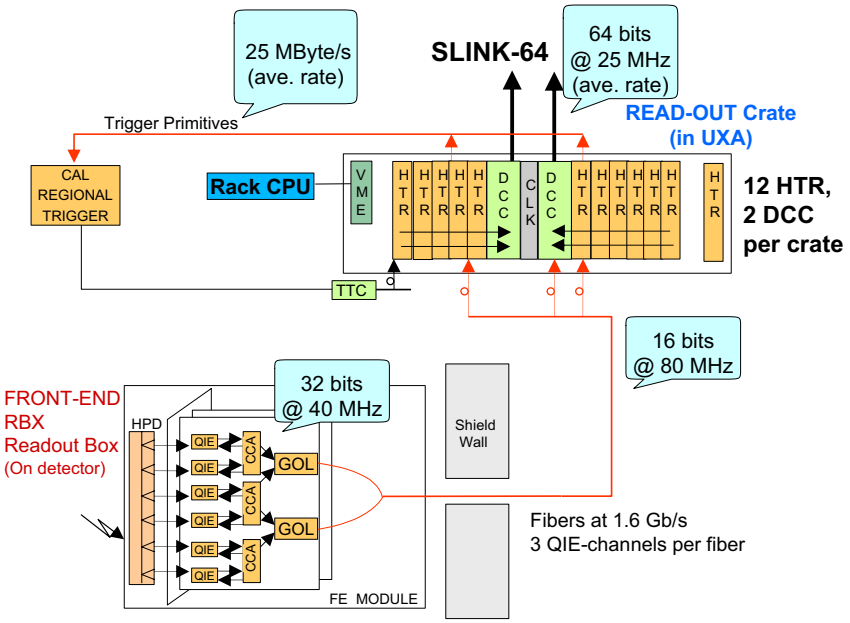


Figure 5.7: Overview of HCAL readout electronics.

shown in Fig. 5.8. The QIE input characteristics were chosen from test beam data to optimize speed and noise performance. This resulted in a per channel RMS noise of 4600 electrons (0.7 fC) corresponding to about 180 MeV.

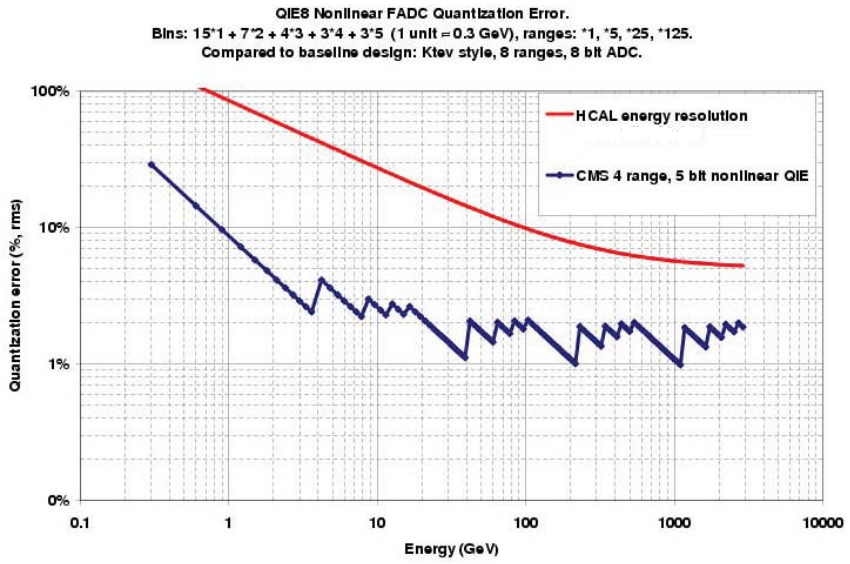


Figure 5.8: Contribution of the FADC quantization error to the resolution, compared with a representative HCAL resolution curve.

The digital outputs of three QIE channels are combined with some monitoring information to create a 32-bit data word. The 32-bit data, at a rate of 40 MHz, is fed into the Gigabit Optical Link (GOL) chip and transmitted using 8b/10b encoding off the detector to the counting house. In the counting house, the data is received by the HCAL Trigger/Readout (HTR) board. The HTR board contains the Level-1 pipeline and also constructs the trigger prim-

itives for HCAL. These trigger primitives are sent to the Regional Calorimeter trigger via Serial Link Board mezzanine cards. The HTR board receives data for 48 channels (16 data fibers) and may host up to six SLBs.

When a L1A is received by the HTR through the TTC system, it prepares a packet of data for the DAQ including a programmable number of precision readout values and trigger primitives around the triggered bunch crossing. For normal operations, the HTR will transmit 7 time samples for each non-zero channel and a single trigger primitive for every trigger tower. These packets of data, each covering 24 channels, are transmitted by LVDS to the HCAL Data Concentrator Card (DCC). The DCC is the HCAL Front-End Driver (FED) and concentrates the data from up to 360 channels for transmission into the DAQ.

5.2.2 Level-1 trigger

The Level-1 trigger primitives (TPG) are calculated in the HTR modules. The QIE data are linearized and converted to transverse energy with a single look up table. Two or more consecutive time samples are summed. A sum over depth is made for those towers having longitudinal segmentation. A final look up table is used to make a data compression before sending the data across the trigger link to the regional calorimeter trigger. Table 5.2 summarizes the geometry of the trigger towers. The HF towers are summed in η and ϕ before being sent to the trigger (Fig. 5.3).

Tower index	$\lvert\eta_{max}\rvert$	Detector	Size	
			η	ϕ
1-15	0.087×η	HB	0.087	5°
16	1.392	HB, HE	0.087	5°
17-20	0.087×η	HE	0.087	5°
21	1.830	HE	0.090	5°
22	1.930	HE	0.100	5°
23	2.043	HE	0.113	5°
24	2.172	HE	0.129	5°
25	2.322	HE	0.150	5°
26	2.500	HE	0.178	5°
27	2.650	HE	0.150	5°
28	3.000	HE	0.350	5°
29	3.314	HF	0.461	20°
30	3.839	HF	0.525	20°
31	4.363	HF	0.524	20°
32	5.191	HF	0.828	20°

Table 5.2: Sizes of the HCAL trigger towers in η and ϕ.

5.2.3 Synchronization and timing

The QIE integration clock is controlled by the Channel Control ASIC (CCA) which allows for fine-skewing of the integration phase of each tower relative to the machine clock. This allows each channel's integration phase to correct for differences in the time-of-flight from the interaction region as well as differences in the optical pathlength within the detector.

Additionally, the detector can be operated with an integration phase tuned to minimize the effect of out-of-time pile-up. The selection of a phase from the full available set depends on the HCAL pulse shape.

The pulse shape of the hadron calorimeter was measured using a photomultipler during the 2002 test beam period. This was accomplished by constructing a special optical unit that coupled the light from one 5° ϕ slice of a wedge (16 η towers) into a single 10-stage, 2-inch photomultiplier (RCA 6342A). The output of the photomultipler was fed into a LeCroy digital scope which recorded voltage vs. time in 0.4 ns steps. Events were recorded in each of 4 categories: a) 300 GeV pion showers, b) 20 GeV pion showers, c) 100 GeV electron showers, and d) minimum ionizing muons. Sample events are shown in Figure 5.9. In general, the pulses show that 68% of the pulse is contained within a 25 ns window leaving 32% of the pulse subject to out-of-time pile-up.

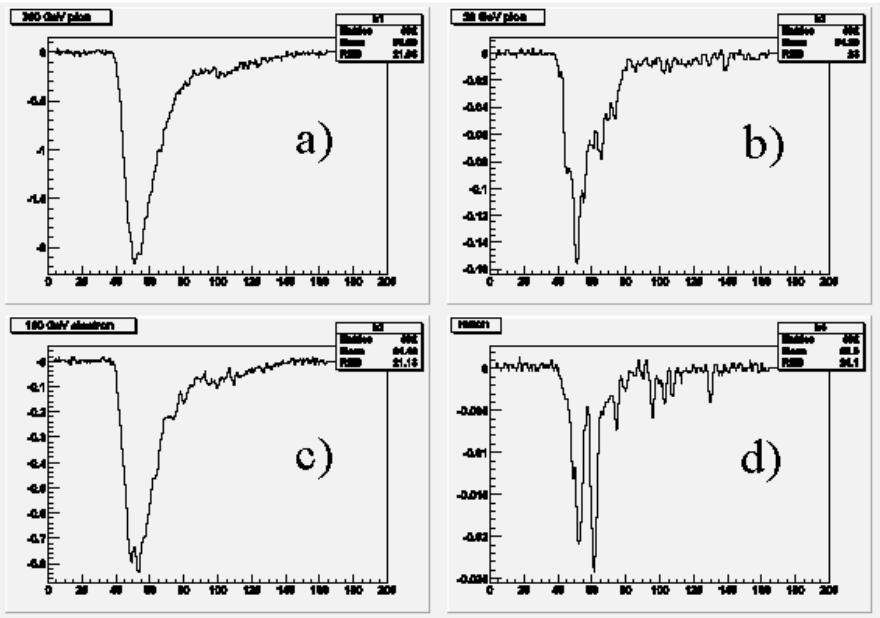


Figure 5.9: Measured single event pulse shape using a photomultiplier for a) 300 GeV pion, b) 20 GeV pion, c) 100 GeV electron, d) muon.

Figure 5.10 shows the pulse shape for the forward calorimeter. The Cerenkov process and phototubes used in the forward calorimeter are extremely fast, so the pulse in HF is only 10 ns wide. The HF is thus subject only to in-time pile-up which is important in the highly active forward region of CMS.

An additional important effect on the HCAL pulse timing comes from the input stage of the QIE. The QIE has an amplitude-dependent impedence which implies a faster pulse shape for large signals than for small ones. This effect was noticed during the test beam 2003 and confirmed in bench measurements as seen in Fig. 5.11. The amount of time slewing is dependent on the noise characteristics of the QIE, so the final QIE ASICs for the barrel and endcap were chosen to limit the timeslew to the "medium" case in exchange for somewhat increased noise. In the outer calorimeter, the noise level is a critical factor for muon identification and pile-up is much less important so the quieter "slow" characteristics were chosen for the HO QIEs.

The in-situ synchronization of HCAL will be performed using the HCAL laser system. The

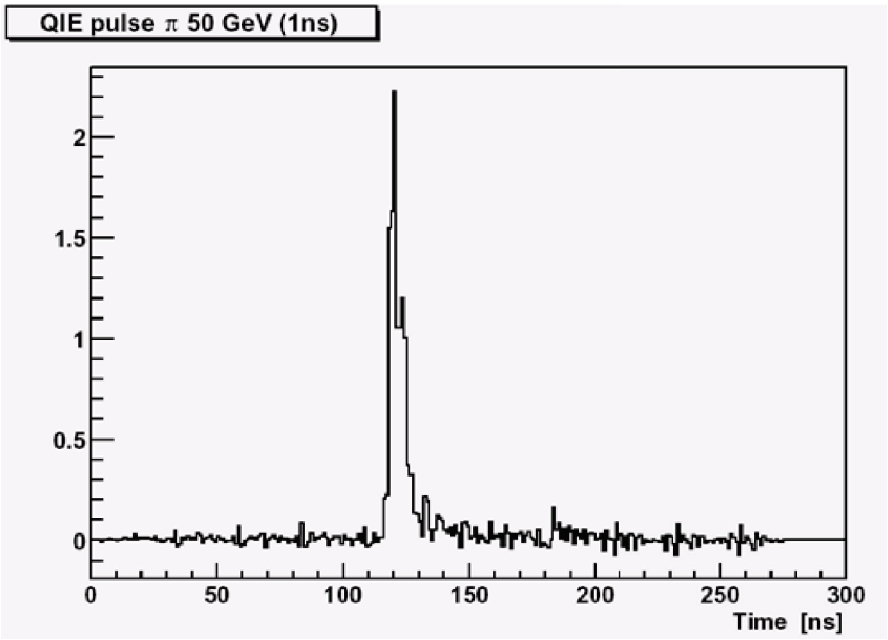


Figure 5.10: Measured pulse shape, energy collected vs. time, for HF.

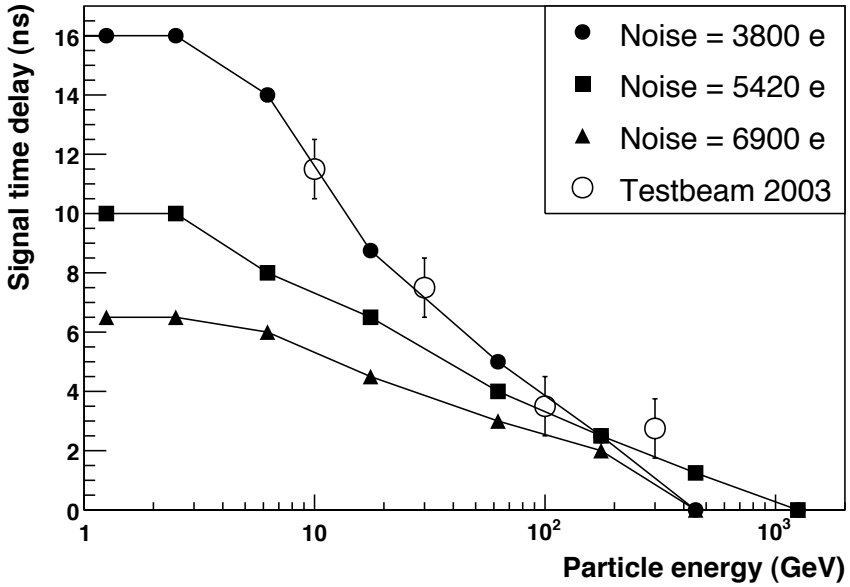


Figure 5.11: Pulse time variation as a function of signal amplitude as measured on the bench (solid points) for several input amplifier configurations compared with testbeam measurements from 2003.

laser system consists of a single UV laser which can illuminate an entire half-barrel of HB or a single endcap at once through a series of optical splitters. The quartz fibers which lead from the laser to the detector have been carefully controlled to equalize the optical path length to each wedge. The laser can be directed either directly onto a scintillator block connected to the HPD or into the wedge. Within layer 9 of each wedge is an arrangement of optical fibers which mimic the time-of-flight from the interaction region. This arrangement allows the timing of HCAL to be flattened and monitored, as has been demonstrated in testbeam

data taking which verified the timing determined by the laser using the synchronized beam. In the HO and HF detectors, only the photodetector can be illuminstated so the alignment will be based on construction and testbeam data.

The channel-by-channel bunch synchronization of HCAL will be determined using a histogramming procedure in the serial link boards (SLBs) which determine the bunch synchronization using the beam structure of LHC. The event and bunch synchronization will be monitored using fast control signals originination from the TTC system which are transmitted in the data stream between the front-ends and the HTR. On a global scale, the bunch and event synchronization between the HCAL and CSC Muon systems was demonstrated during a testbeam run in 2004. These data were also used to measure the punch-through rate into the forward muon chambers, ME1, and are reported in Section 3.3.2.2.

5.3 HCAL Calibration

The HCAL calibration system will be used to set the initial absolute energy scale, understand the detector response and uniformity, and will monitor the time stability during physics data taking.

The initial calibration information came from quality control tests performed with: 1) a collimated radiation source which guaranteed that the scintillating tiles matched specifications, and 2) a charge injector used to validate the digital converters.

The energy scale constants are obtained by combining test beam data taken with $e^{\pm}$, $\pi^{\pm}$, and muon beams with radiation source data taken for a limited number of modules. The detector is not assumed to be uniform. Instead, the test beam energy scale information is translated to the full system by comparing the radiation source data taken *in situ* for every channel.

Updates to the constants will come from the analysis of physics events.

The full response of each channel (scintillator-photodetector-electronics) will be monitored using radiation sources and the UV-laser. Additional monitoring of the electronics will also be done with the UV-laser system and a blue LED pulser.

As described below, a combination of all of the above systems will allow us to have up-to-date values of the basic calibration constants for all $\approx 10\,000$ readout channels.

5.3.1 Charge injector calibration for the ADC

The ADC-to-charge conversion factors for each range of a QIE channel (see section 5.2.1) were obtained using a charge injector. These calibration constants are given by the ratio of the measured response (fC) over the input response (fC). Figures 5.12a) and 5.12b) show the measured values for the QIE's installed in HB/HE and HF. For both types of QIE's we have a 2-3% spread, while the average values are 0.91 ADC/fC and 0.36 ADC/fC, respectively. The lower values for HF are required by the difference in gain with respect to HB/HE/HO.

Figures 5.12c) and 5.12d) show the conversion factors needed for the source calibration data. In this case QIE's are initialized with a range that is about three times smaller, to cope with the small signals.

The pedestal value for a QIE's depends on which of the four internal capacitors (CAP-ID) is being used. As a consequence, all the conversion factors are obtained for each of the capacitors separately as well.

All the ADC-to-charge conversion factors will be kept in the calibration database.

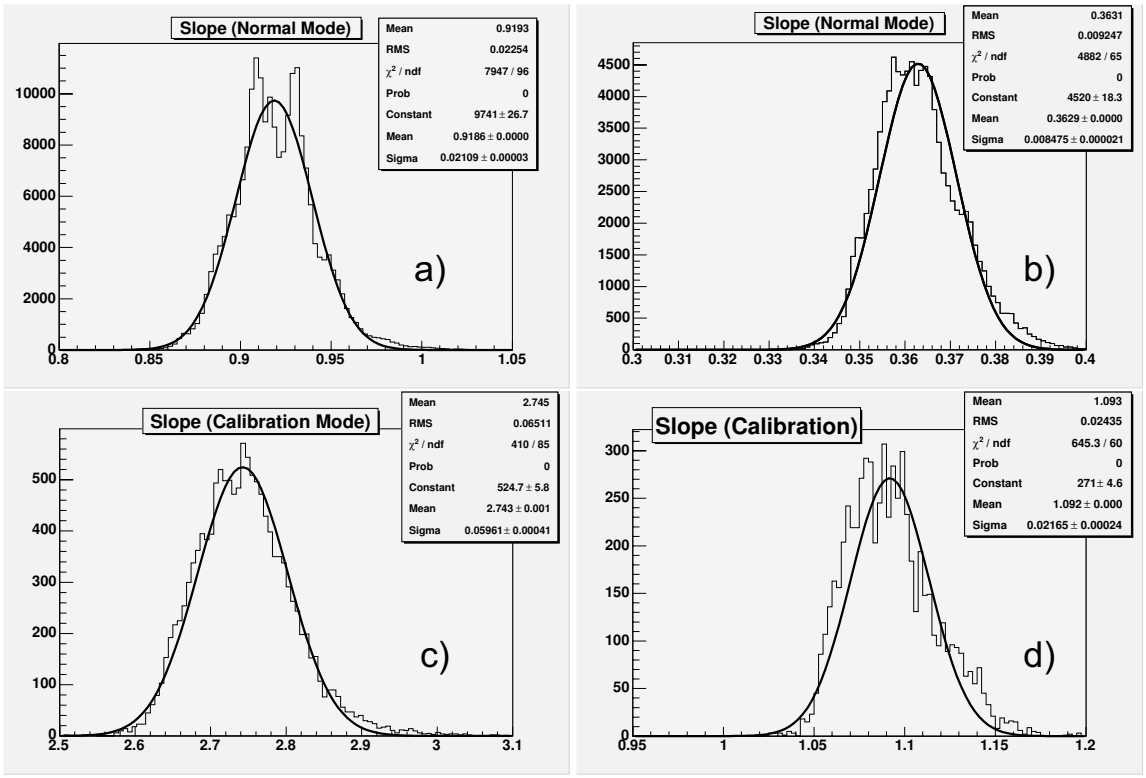


Figure 5.12: Slope distribution of QIE cards operated in "normal" mode for HB/HE (a) and HF (b). The bottom distributions are the QIE slope in "calibration" mode for HB/HE (c) and HF (d).

5.3.2 Radioactive source calibration

Source calibration has to be made for each scintillator tile in every layer of the calorimeter. For each calorimeter tile the source calibration coefficient has to be derived. The ratio between source radioactive signal to the response of the different particles in the tower provides initial calibration coefficients for the HCAL towers that were not exposed in test-beams.

Figure 5.13 shows the responses of a HB scintillating tile as the radioactive source passes by it. The distribution of scintillating tile response is presented in Fig.5.14. The width of the distribution is 8%. The difference in peak response taken when the radiative source is extending versus when it is retracting shows a reproducibility of the calibration signal to better than 1%, as demonstrated in Fig.5.15. The dependence of the ratio of radioactive signal to electron signal on the η-number of tower is shown in Fig.5.16 and the same dependence for the ratio of radioactive signal to the muon signal is presented in Fig.5.17. The comparison between the radioactive signal and the electron test-beam data gives 0.22–0.25 GeV/ADC count or 0.2–0.23 GeV/fC. Similarly, the test-beam HE analysis yields 0.163 GeV/ADC count.

Figure 5.18 shows the measured signal from the *in situ* source calibration for a series of η-towers in HB+ (half-barrel on positive η-side), averaged over all ϕ and layers of a given tower. (Only the layers with 4mm-thick scintillators were used.) The signal increases steadily with η, as expected due to decreasing attenuation of the signal from the shorter length of fiber between the scintillator tile and the HPD photodetector. This behavior must be established for each value of ϕ. Figure 5.19 shows the uniformity of response at a fixed η stepping through values of ϕ. As examples, η tower numbers 4 and 12 are shown, as they should have distinctly different responses as shown in Fig. 5.18. The plots show good uniformity as a function of ϕ. The values are fitted to a level function and the result is reported in the box as "A0." The scatter about the fitted value reflects the variation of the performance of individual tiles, and the gains of the photo-detectors. It is this variation which the source calibration will quantify so that it can be taken into account for the sake of obtaining the best possible energy resolution.

These constants needed to be corrected for the fact that the source data were taken in the absence of a magnetic field. The magnetic field improves the HPD's behavior by suppressing cross-talk between pixels. In addition, a 10% increase in the scintillator brightness is expected in the nominal 4 T field. These corrections will be calculated using data to be taken during the magnet test.

HB and HE have the similar strategy for source precalibration and a 2% measurement can be achieved for a single channel from a measurement of an individual layer. The main difference between the HB and HE analysis is that for HE is more sensitive to the fact that the source is not collimated due to the smaller size of its tiles.

The source will also be use to quantify the radiation damage of the tiles during the lifetime of the experiment. For this purpose, the individual layer information for a given tower will be tracked in a database.

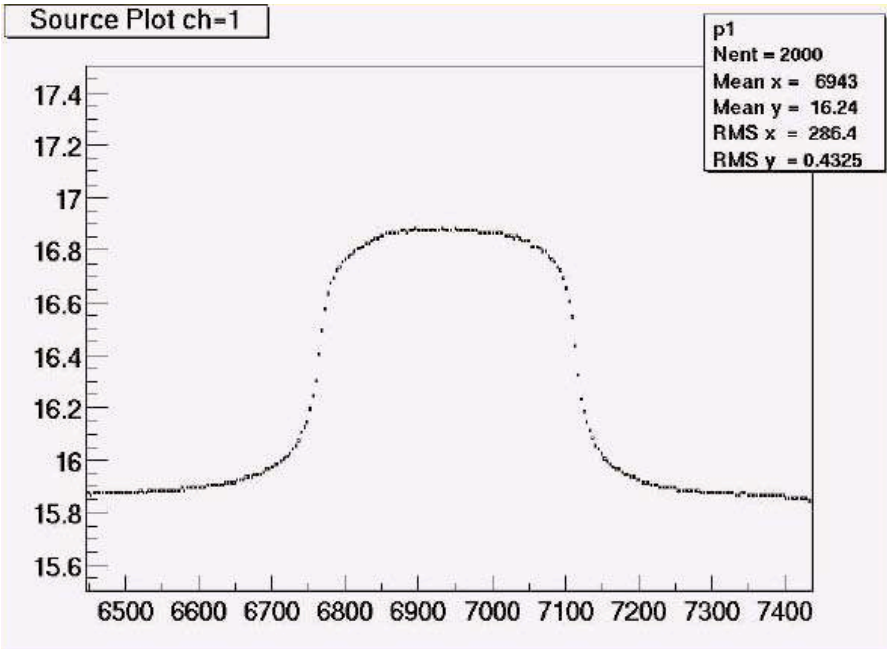


Figure 5.13: Response of a scintillating tile as the radioactive source passes by it.

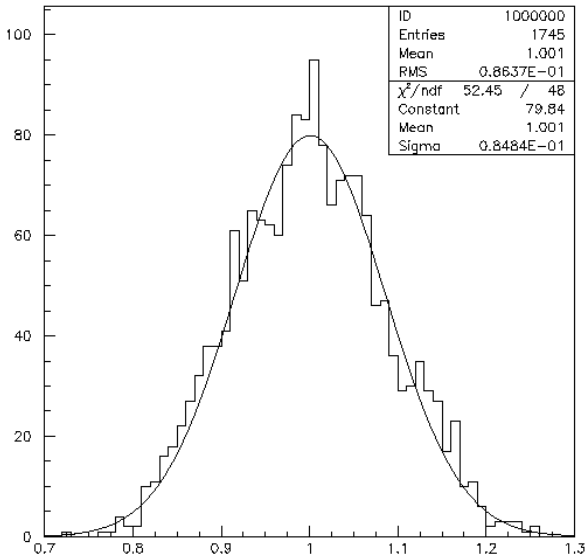


Figure 5.14: Distribution of scintillating tile response to the source calibration. The observed RMS spread of about 8% is consistent with that observed by source measurements taken during tile assembly.

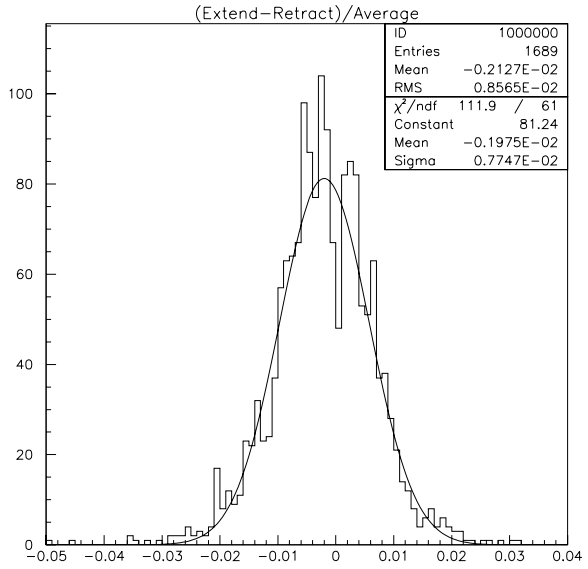


Figure 5.15: Difference in peak response when the radiative source is extending or retracting.

Initial precalibration of HF cells energy response will be based on the measurement of the position of the single photoelectron peak in each cell. It was measured in test beam that the single photoelectron peak corresponds to 4 GeV to within 10%. Some improvement of the precalibration accuracy will be achieved basing the cell response on the data taken with a ^{60}Co 5 mCi radioactive source. Comparison of such source calibration with the test beam results shows this method to be accurate to 5%.

All calibration coefficients will be kept in the calibration database, as discussed in Section 5.3.4.

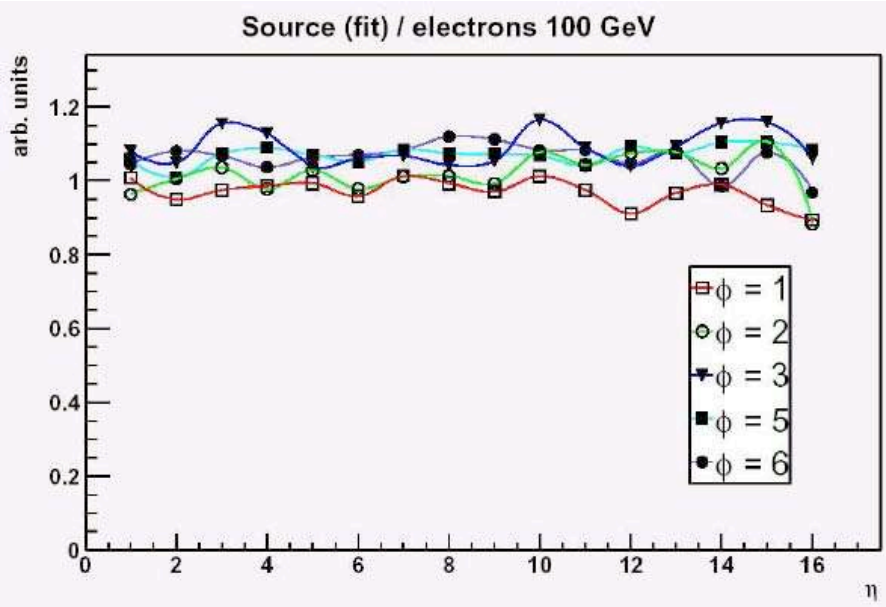


Figure 5.16: Ratio of radioactive source signal to electron signal.

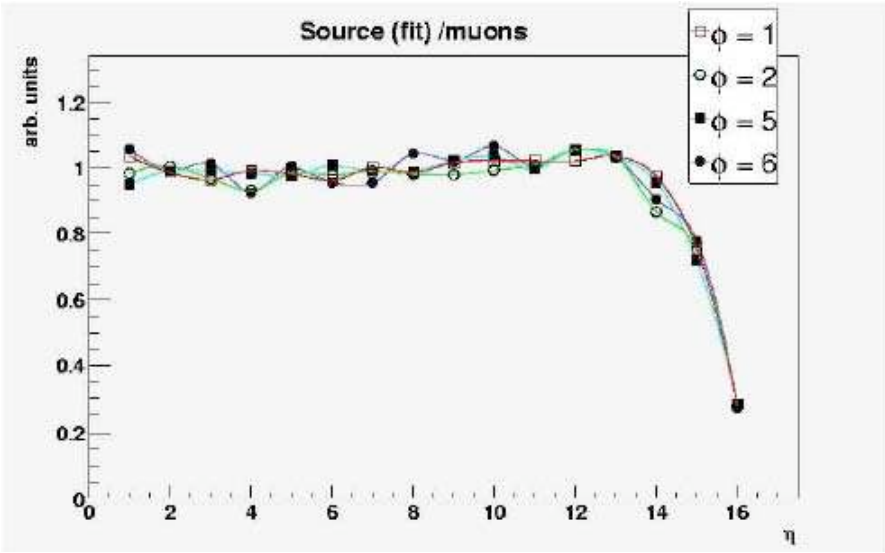


Figure 5.17: Ratio of radioactive source signal to muon signal.

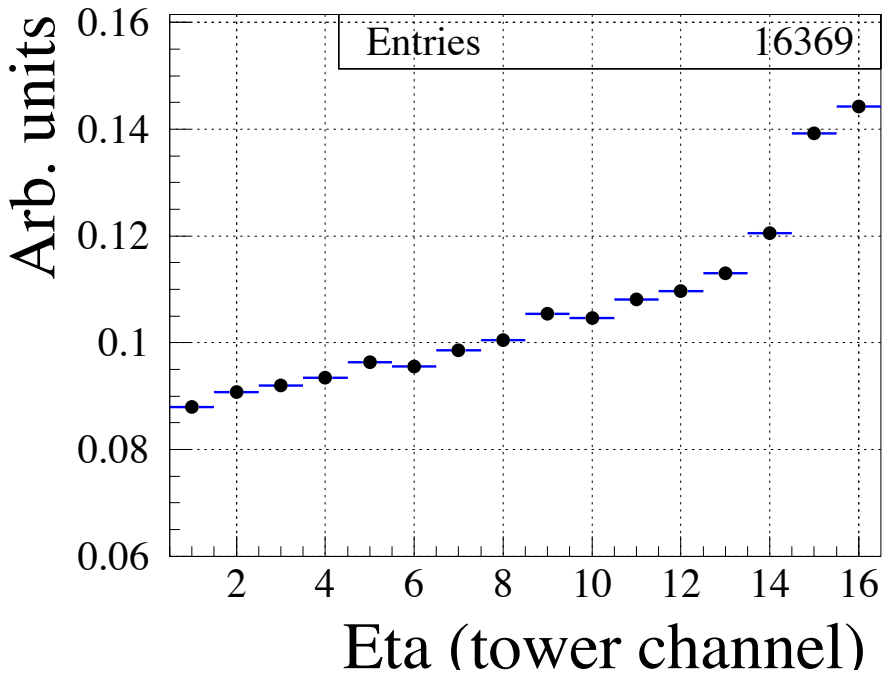


Figure 5.18: Average radioactive source signal for HB+ with a 1.5 mCi Co-60 source.

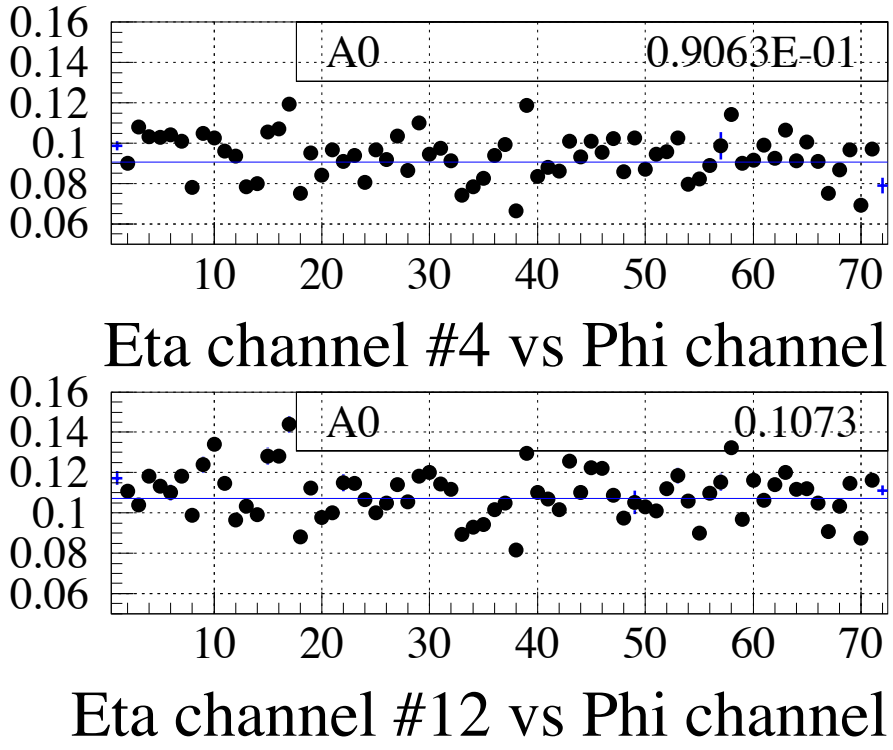


Figure 5.19: Radioactive source signal in HB+ for tower numbers 4 and 12 as a function of ϕ channel.

5.3.3 Offline calibration and monitoring

During run time the HCAL calorimeter will be immersed in a 4 T magnetic field, while the calibration coefficients are found in the test-beams without magnetic field. The magnetic field influences the showering corrections in the transition region between HB and HE; HE and HF. For data taken at high luminosity, the response in HE will be degraded due to radiation damage [164]. For these reasons, the recalibration and monitoring of HCAL towers during run time is required.

Fast monitoring can be achieved from pedestal runs, laser, LED and radioactive source data taken between stores and minimum-bias data during run time.

The Blue-LED system can make a 3% measurement of the response of the electronics for all individual channels. The equivalent test with the UV-laser has a resolution of 4%. The radiation source makes a 2% measurement of the response for the full chain of layers 0 and 9. As the LED and the source information have differing sensitivity to ageing and radiation damage, the combination of the two monitoring signals can be used to diagnose the origin of a change in the signal strength in an individual channel.

The histogramming of tower energy distributions above threshold in minimum bias events provides a high statistics test of the uniformity of the energy scale in ϕ. By tracking the higher order moments of the energy distribution per readout channel, effects of miscalibration can be more rapidly detected than from studying shifts in the mean response alone. The statistics needed to calibrate each tower with a precision better than 2% is collected in 1–2 hours.

The barrel part of HCAL and part of HE are covered by the tracker up to $|\eta|$ <2.4. The towers can be recalibrated with isolated energetic particles from $\tau \to \pi\nu$ in $W \to \tau\nu$ and $Z, \gamma^* \to \tau\tau$ processes [165]. The other possibility is to use isolated energetic particles from QCD-jets [166]. Both these possibilities allow us to calibrate HCAL towers covered by the tracker with an accuracy better than 2% during one month using charged particles with transverse momentum from 15 GeV/c to 70 GeV/c. The calibration method uses the E/p ratio similar to the techniques used in test beams. The regions that are not covered by the tracker ($|\eta|$ >2.4) can be calibrated using E_{T}-balance in di-jet events or γ/Z+jet events, as described in Chapter 11.

5.3.4 Database

The online database (OMDS) contains all values loaded by the control system and data acquisition, channel mapping, and calibration constants. The online-offline (ORCON) database for the HCAL system contains a replica of the minimal calibration information needed for event reconstruction.

The OMDS is subdivided in two: "Configuration" database and "Conditions" database, as described in Section 2.8. The "Configuration" database contains the relationships or mapping for all HCAL detector components: wedges, layers, read-out boxes (RBX), cables, HCAL Trigger (HTR) cards, and test and calibration results for various components *e.g.* RBX, QIE, source types and strength. The "Conditions" database has the slow controls logging, the calibration constants like pedestals, gains, timing information, etc., and the "Configuration" database downloaded to the readout system during the initialization.

The overview of the HCAL controls and readout chain are given in Fig. 5.6 and Fig. 5.20, respectively. These systems record temperature, humidity and other constants useful for

correlation studies of detector/calibration stability. The calibration constants needed for data analysis are stored in the "Conditions" database. A brief breakdown of the number of type of calibration data is listed here:

Channel response [scint-HPD-QIE]	- 1 per channel and/or 18 pixel per channel
(see Sections 5.3.2,5.3.3.)	15 per channel or 1 per layer
QIE calibration [ADC-to-Charge(fC)]	- 128 per channel (32 per gain, mean CAP-ID's)
(see Section 5.3.1.)	512 per channel (128 $\times$ 4 CAP-ID's separately)
Charge-to-GeV [HB,HE,HO]	- 1 per region (assuming good uniformity)
(see Section 5.3.2.)	1 per LUT channel in case of the trigger
Pedestals	- 4 per channel (1 per CAP-ID)

The database information for HF is similar. The spread of the channel response will also be kept for simulation purposes.

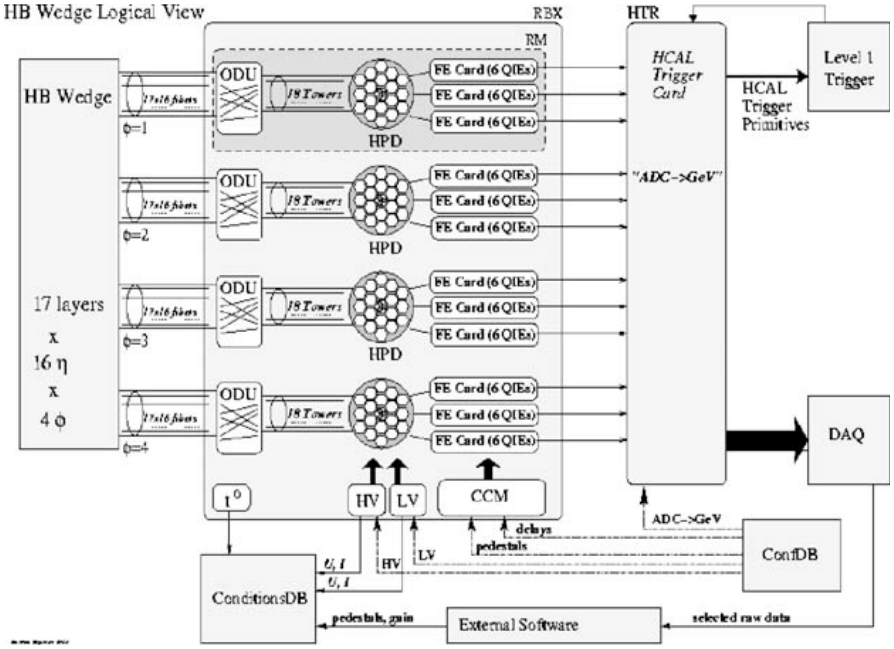


Figure 5.20: Overview of HCAL readout/trigger chain and connections to database.

5.4 HCAL performance in test beams

Test beam data are compared with simulation to validate the modeling of the calorimeter response, resolution and shower shapes for pions and electrons over the energy range of 5–300 GeV. The low energy performance is studied with muons. The CMS Collaboration developed the Object oriented Simulation for CMS Analysis and Reconstruction [31] (OSCAR) framework, based on the GEANT4 tool kit [30], to describe the detector geometry and the passage of particles through the detector material. GEANT4 uses either parametric (LHEP) or microscopic (QGSP) physics models to simulate the particle showers arising from the interaction of particles with matter.

5.4.1 Longitunidal shower shape and magnetic field brightening

A prototype of this calorimeter was made and tested [167] in a particle beam in 1996. This prototype geometry has been simulated using OSCAR (version 2.4.5) based on GEANT (version 4.5.2.p02) [30] and GEANT3 [28]. The data have been analysed in the same way as data from the test beam experiment. The features that are compared here are the longitudinal shower development of single particles, and the effect of a magnetic field in the calorimeter performance.

The test beam prototype of HCAL is based on a hanging file structure, in which absorber plates made out of a special type of brass (composed of 59% copper, 39% zinc, 1% iron, 1% manganese by weight) and varying in thickness from 2 cm to 8 cm are sandwiched with scintillator tiles. There are 28 tiles, mostly 4 mm thick. A detailed description of the 1996 test beam setup and the analysis presented here is available in [167, 168].

Data were recorded with several different particle beams going through the HCAL prototype detector. Electron and pion beams had energies in the 10–300 GeV range, and the muon beam energy was 225 GeV. One set of data were taken with only HCAL in the beam. Another set was taken with ECAL upstream of HCAL in the beam. All the test beam data taken with HCAL only and ECAL+HCAL system were repeated in the presence of a solenoidal magnetic field. The strength of the magnetic field could be increased up to 3 T.

Figures 5.21 shows the total energy measurement for 100 GeV pions with the HCAL setup without ECAL. The two plots show measurements at different magnetic field values: 0.0, 0.75, 1.5 and 3.0 T. The peak position shifts towards higher values and the distribution gets broader as the field value is increased, as predicted by scintillator brightening and showering effects. The behavior in data is reproduced in the Monte Carlo simulation.

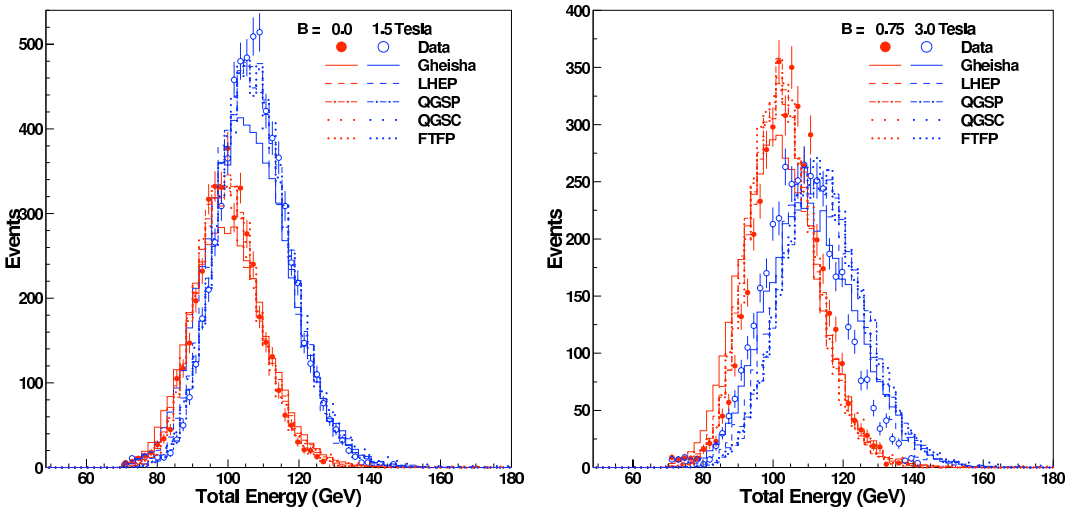


Figure 5.21: Total energy distribution for 100 GeV pions in the HCAL setup without ECAL for different settings of the magnetic field.

Figure 5.22 shows the mean energy response as a function of the B-field value for 100 GeV pion in hadron calorimeter. The increase in the measured energy with B-field, as seen in the data, is well explained by the Monte Carlo simulation models.

Figure 5.23 shows energy resolution for 100 GeV pions in the hadron calorimeter as a function of B-field value. There is a slight increase in the energy resolution with magnetic field

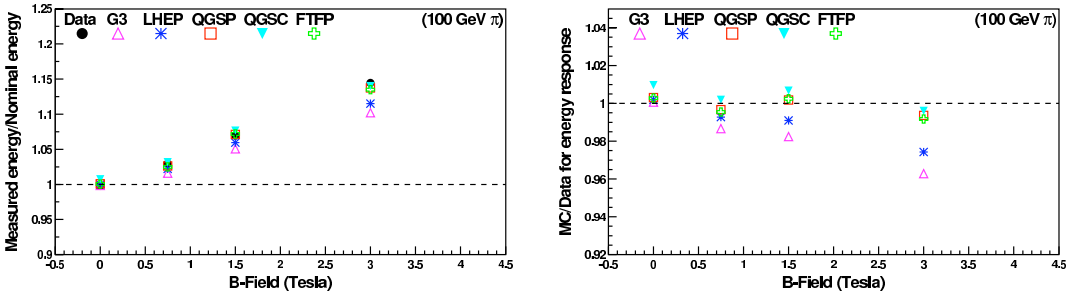


Figure 5.22: Energy response for 100 GeV pions at different field values for HCAL only.

which is well reproduced by the Monte Carlo simulation models.

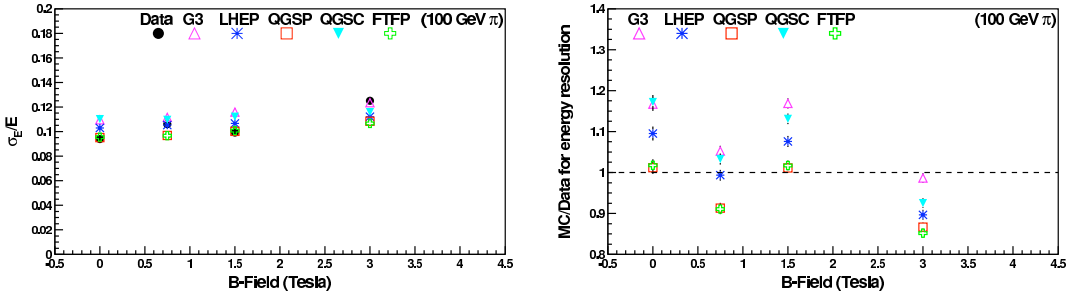


Figure 5.23: Energy resolution for 100 GeV pions at different field values for HCAL only.

Figure 5.24 shows the variation of longitudinal shower profile of 100 GeV pions in the HCAL due to the presence of 0.0 (no field), 1.5 and 3.0 T magnetic fields. The shower profile for test beam data is broader than the Monte Carlo prediction and there is a clear indication that the increase in response is at the shower maximum.

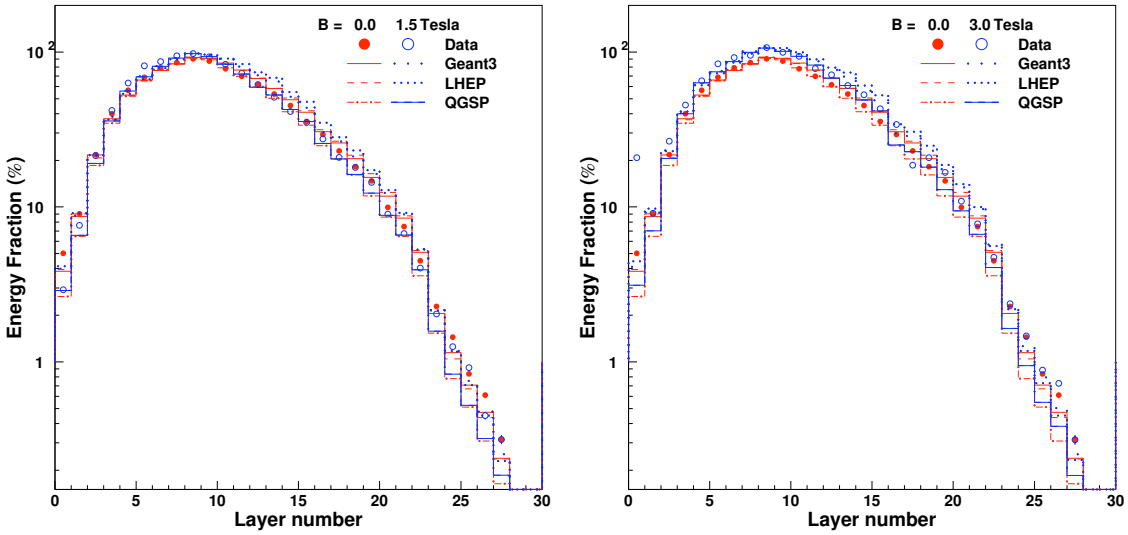


Figure 5.24: Longitudinal shower profile of 100 GeV pions in HCAL in the presence of magnetic field of (left) 0.0 and 1.5 T; (right) 0.0 and 3.0 T.

5.4.2 Response and resolution

The 2002 HCAL test beam (TB02) experiment [169] was a small scale detector setup designed to study the performance of the CMS hadron calorimeters. The detector was exposed to beams of π^- (20, 30, 50, 100, 300 GeV), electrons (20, 30, 50, 100 GeV), and muons (225 GeV) over a large energy range. More than 100 million events were read out with a 29.6 ns period, slightly longer than the 25 ns planned for CMS. A total of 144 Hadron Barrel (HB) channels (two wedges) and 16 Outer Hadron (HO) channels were tested in a detector configuration which included an aluminum slab, representing the solenoid material, and a prototype of the electromagnetic calorimeter (ECAL). The single particle performance measurements and comparison with Monte Carlo simulations are discussed in detail in [170, 168].

The way to determine the "nominal" pion energy distributions was to reject the background events by means of cuts in the HCAL energy versus ECAL energy plane. Figure 5.25 shows the pion energy distributions after cuts in ECAL versus HCAL energy space. Gaussian fits to the energy distributions in Fig. 5.25 show non-Gaussian high energy tails for 20-30 GeV pions as a consequence of the non-compensating nature of the CMS calorimeters. The low energy tails for higher energy pions are a manifestation of energy leakage beyond the HB outer limits.

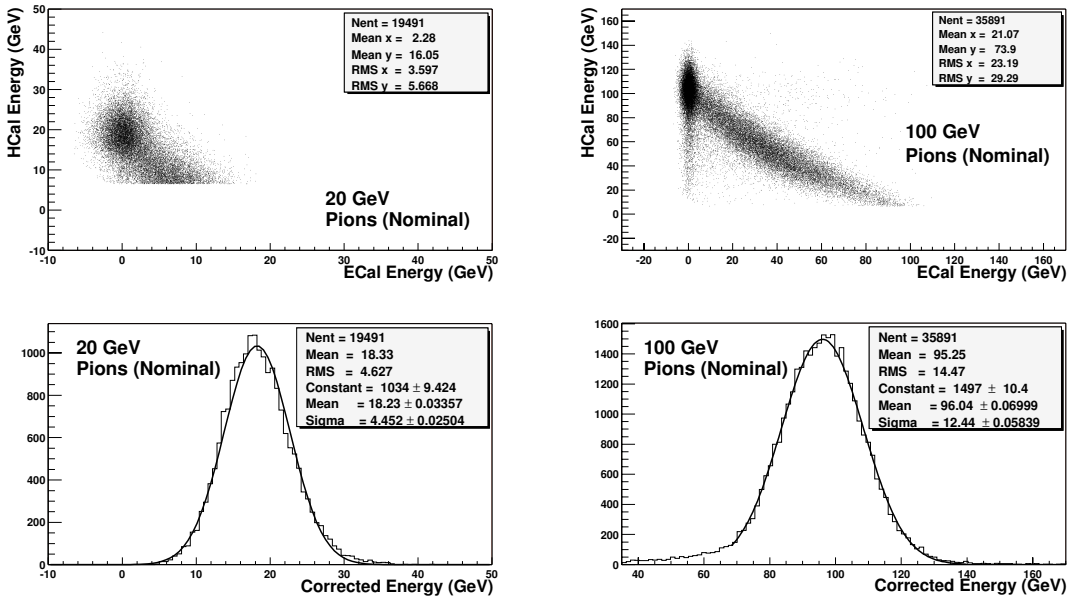


Figure 5.25: Energy distributions of 20 GeV and 100 GeV pions, after the HCAL versus ECAL energy based background cuts were applied. Some signal events are lost at low energies. The non-Gaussian low energy tails are due to energy leakage beyond the HB outer limits.

Figure 5.26 shows the response of the ECAL+HB system to π^- as a function of energy calculated as E_{π^-}/E_i, where E_{π^-} is the mean of Gaussian function fit to the energy distributions and E_i is the initial beam energy. The result from a Monte Carlo simulation, based on OSCAR_2_4_5 using GEANT4.5.2.P02 and LHEP-3.6 or QGSP-2.7, is compared with the TB02 data measurement. The vertical bars are the statistical and systematic uncertainties in the data measurement added in quadrature. For comparison, the simulated linearity plots were scaled up by 5% (LHEP) and 3% (QGSP) to approximately match the data at 300 GeV. Within the experimental uncertainties, there is good agreement between the data and the simula-

tion. While the LHEP (version 3.6) physics list predicts a response growing slightly faster than the measured in data, the one derived from the QGSP (version 2.7) physics list grows slower.

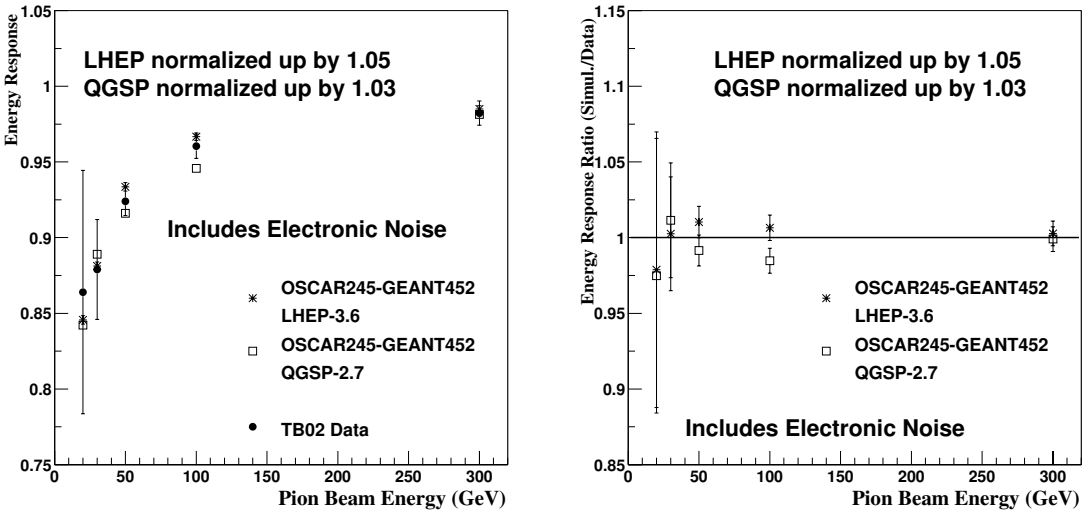


Figure 5.26: In the ECAL+HB configuration for pions, comparison between the simulated and measured energy response, E_{π^-}/E_i, using both LHEP (version 3.6) and QGSP (version 2.7) physics lists.

Figure 5.27 shows the resolution of the ECAL+HB system for π^- as a function of energy calculated as σ/E_{π^-}, where E_{π^-} is the mean and σ the square root of the variance of a Gaussian function fit to the energy distributions. The vertical bars are the statistical and systematic uncertainties in the data measurement added in quadrature. The data measurement is in good agreement with the simulation, for the two physics lists tested, given the high correlation of the TB02 measurement uncertainties point-to-point in pion energy. The LHEP (version 3.6) physics list gives a fractional energy resolution a few percent larger than the QGSP (version 2.7) list. The data results and Monte Carlo prediction are easier to compare in the ratio plot shown in Fig. 5.27.

5.4.3 Low energy e/π response

The 2004 HCAL test beam detector configuration was very similar to that of the 2002 experiment [17]. The beam line was more complex to allow for the very low energy beam (VLE), as well as better particle identification using muon counters, Cerenkov detectors, and wire chambers. The detector was exposed to π^- beams of energies in the 10–300 GeV range, as well as to VLE beams in the 2-9 GeV range. The HCAL readout was arranged in two different configurations. The standard configuration consisted of one channel per tower, giving maximum transverse granularity but no longitudinal segmentation. The second configuration was designed to study longitudinal shower shapes, and the signal for a channel was therefore integrated in η for each ϕ slice. The π/e response ratio was measured as a function of the pion energy in the 5–300 GeV range. Fig. 5.28 shows a comparison to a Monte Carlo simulation based on OSCAR_3_7_0 using GEANT4.6.2.P02. The agreement is good over the full energy range which was tested.

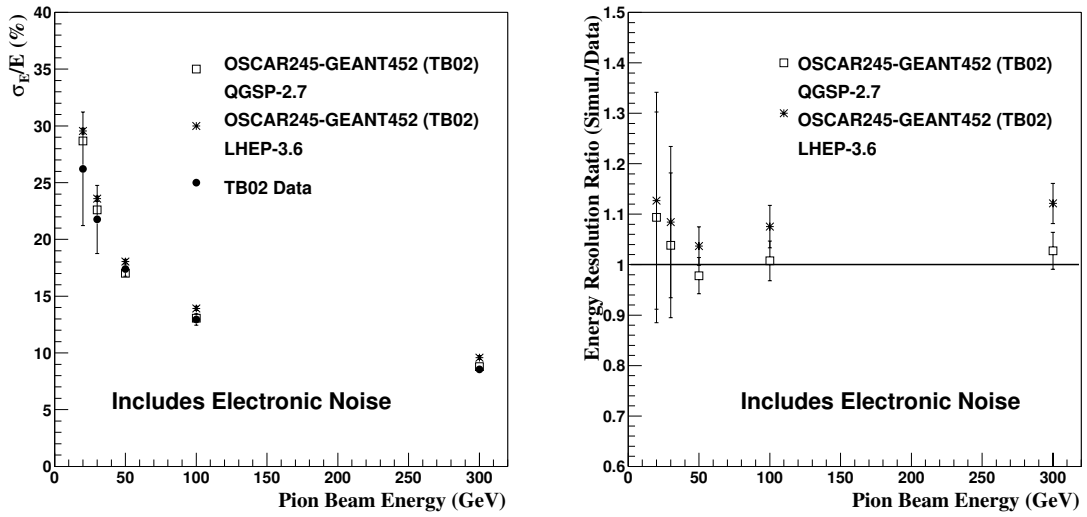


Figure 5.27: In ECAL+HB configuration for pions, comparison between the simulated and measured energy resolutions, σ/E_{π^-}, using both LHEP (version 3.6) and QGSP (version 2.7) physics lists.

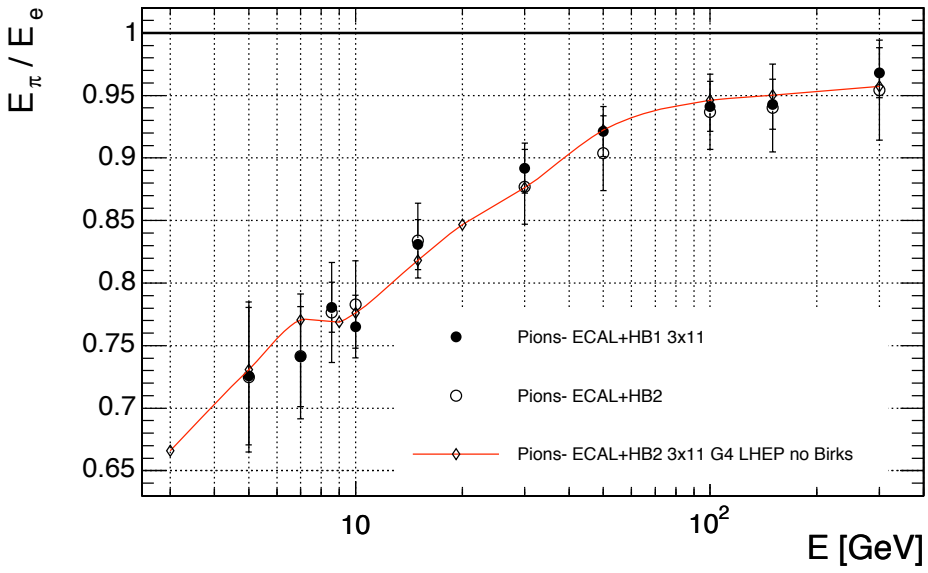


Figure 5.28: Pion response versus energy in the central HCAL barrel. The error bars are statistical and systematic uncertainties added in quadradure.

Figure 5.29 shows the measured mean energy depositions in layers 0–16 of a hadron barrel module for 300 GeV pions. The agreement between data and the Monte Carlo prediction is excellent when using the LHEP (version 3.6) physics list. The QGSP (version 2.8) physics list predicts shorter showers than observed.

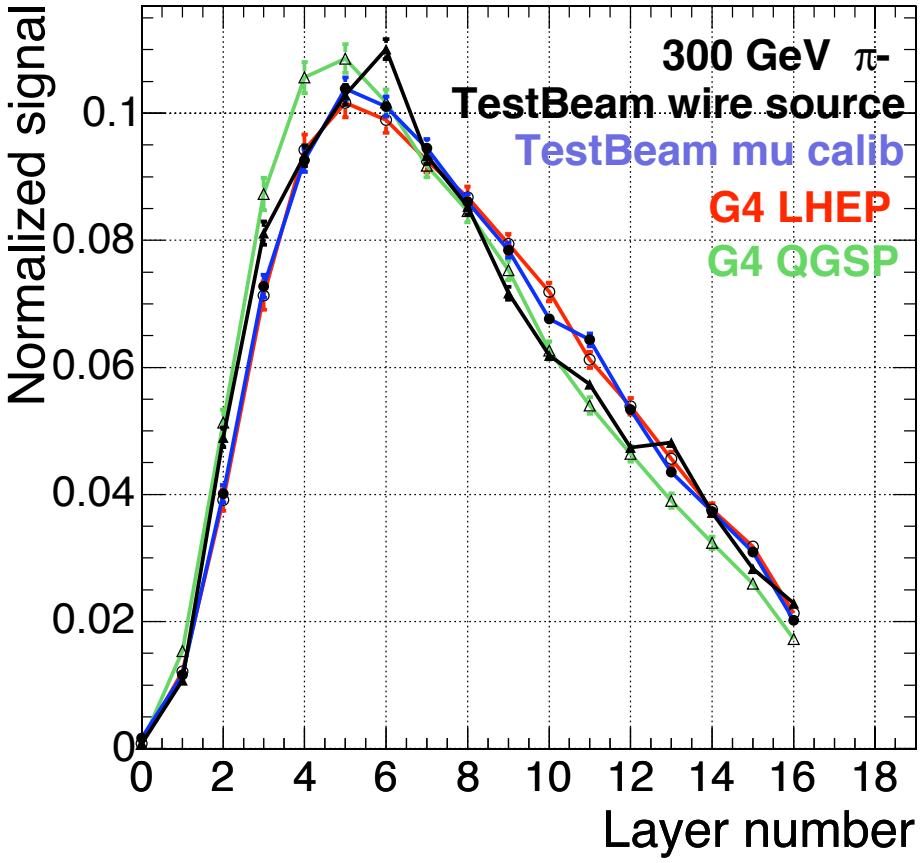


Figure 5.29: Longitudinal shower profile for 300 GeV pions.

5.4.4 Cluster-based response compensation

The response of the full ECAL+HCAL barrel calorimeter to electrons differs from pions in a substantial and energy dependent fashion, as shown in Figure 5.28. This leads to an energy dependent response and a degradation in the pion and, hence, jet resolution. The measured raw energy response and fractional energy resolution for pions as a function of energy are shown in Fig. 5.30. The non-compensated energy resolution is parameterized as $\frac{\sigma}{E} = \frac{120\%}{\sqrt{(E)}} \oplus 6.9\%$. The uniformity and resolution of the pion energy measurement can be substantially improved using clustering techniques that utilize the fine granularity and separate energy measurements of the ECAL and HCAL compartments [171].

The technique in [171] uses test beam data to fit the intrinsic electron to hadron response (e/h) and the average neutral fraction f_0 of the ECAL and HCAL as a function of the raw total calorimeter energy, $\epsilon_E + \epsilon_H$, using the following fitting functions:

$$(e/h)_{(H,E)} = a_{(H,E)}[1 + b_{(H,E)}/(\epsilon_E + \epsilon_H)] \quad (5.1)$$

$$f_0 = c \log(\epsilon_E + \epsilon_H)^d \quad (5.2)$$

for the free parameters: a_H, a_E, b_H, b_E, c and d. The fit is performed minimizing the calorimeter energy resolution as follows:

$$(e/\pi)_{(H,E)} = (e/h)_{(H,E)}/[1 + ((e/h)_{(H,E)} - 1) \cdot f_0] \quad (5.3)$$

$$E = (e/\pi)_E \epsilon_E + (e/\pi)_H \epsilon_H \quad (5.4)$$

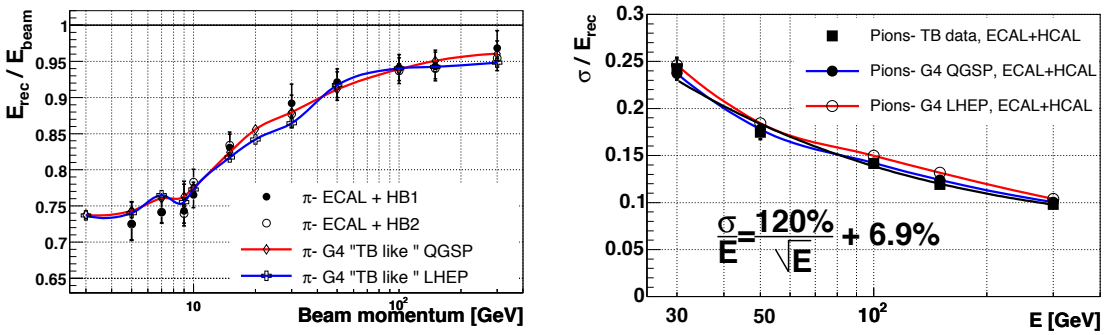


Figure 5.30: Nominal raw energy response and fractional energy resolution as a function of ECAL+HB energy for pions. The fit for the non-compensated energy resolution is shown.

$$\chi^2 \quad = \quad \sum_i (P_0 - E_i)^2/\sigma_i^2 \tag{5.5}$$

where χ^2 is minimized given the known beam momentum P_0 and the computed pion energy E_i for 2004 test beam data taken at 10, 30, 100 and 300 GeV. The improvement in the compensation of the single pion response using the weighting procedure is clearly demonstrated in Figure 5.31 for 15 GeV pions from test beam. The result of this procedure over the 5–300 GeV pion energy range yields an improved pion resolution as shown in Figure 5.32.

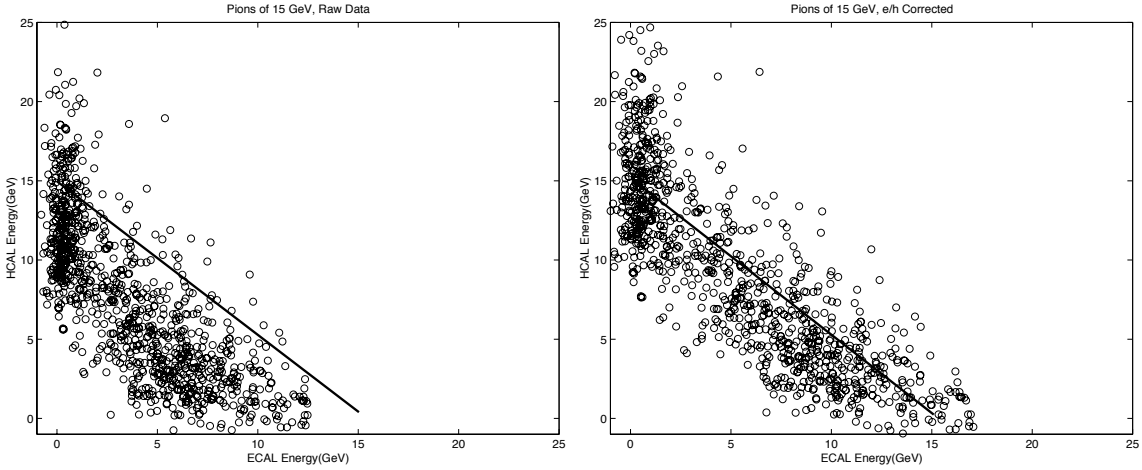


Figure 5.31: Scatter plot of the 15 GeV incident pion beam data. The plots are shown for raw data (left) and for corrected data (right) showing an improved linearity in pion energy response for ECAL+HB.

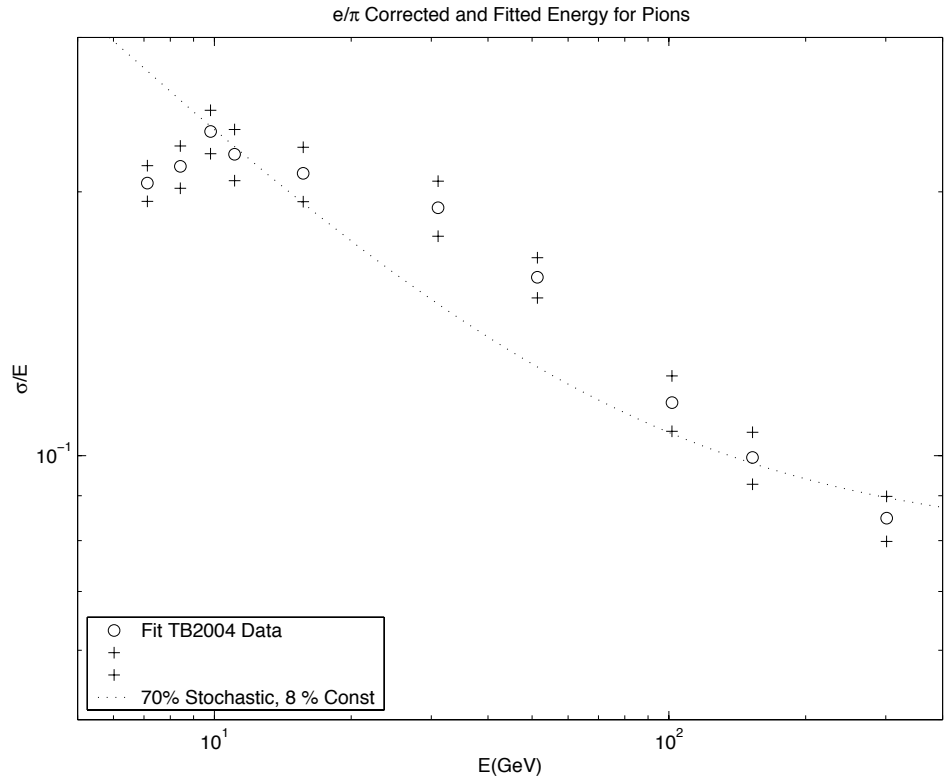


Figure 5.32: Application of the cluster-based weighting technique to the fractional energy resolution as a function of ECAL+HB energy for pions. The uncertainty band on the data points is indicated by the crosses. The test beam data are compared with a resolution curve with a 70% stochastic and 8% constant term.

Chapter 6

Inner Tracking System

6.1 Simulation

6.1.1 Simulated geometry

The layout of the Silicon Strip Tracker (SST) and the Pixel Tracker is described in Section 1.5.5.

The Pixel barrel detector is composed of 3 layers, each layer is divided into 2 half-cylinders, and each half-cylinder is composed of ladders and half-ladders. The ladders and half-ladders provide the support structure and cooling for pixel modules (Fig. CP 11) with each ladder containing 8 modules. To make the layers hermetic in the r-ϕ plane, the half-cylinder structure is designed with half-ladders at each edge. The 2 half-cylinders overlap in the edge regions and therefore provide a small overlap of the sensitive detector area.

The Pixel forward detector is composed of 4 disks, with 2 disks on each side. Each disk is divided into 24 blades with 7 modules of different sizes on each blade. The blades are tilted by 20° resulting in a turbine-like geometry. Both the blades and the modules on the blades overlap in order to provide hermetic coverage.

The Tracker Inner Barrel (TIB) is composed of 4 layers, each divided into 4 parts: the positive and negative sides in z, and upper and lower parts in each side (Fig. CP 12). In these "shells" (parts), the modules are arranged inside "strings." The structure of the strings is different if the strings are positioned in the forward part or the backward part of the internal layer with respect to the external layer. The structure also varies among the different layers. There are 3 modules in each string and there are double-side modules in the first 2 layers.

The Tracker Inner Disk (TID) is comprised of 6 identical disk structures. In each disk the modules are arranged in 3 rings and placed alternatively in the forward and backward parts of the disk.

The Tracker Outer Barrel (TOB) is comprised of 6 layers. Each layer is made of rods with 6 modules arranged inside each rod. The 2 innermost layers have double-sided modules. The rod structure changes according to the layer and depends on the z position.

The Tracker EndCap (TEC) contains 9 disks. Due to their large size, the disks are made of 8 petals. In each petal the modules are arranged in rings. The number of rings in each petal varies from 4 to 7, depending on the position in z. As for the TID, the TEC modules are arranged in rings placed alternatively in the forward and backward sides of the petal.

The CMS Tracker includes both sensitive volumes and non-sensitive ones. Since the tracker requires a large amount of low-voltage power, a large amount of heat needs to be dissipated.

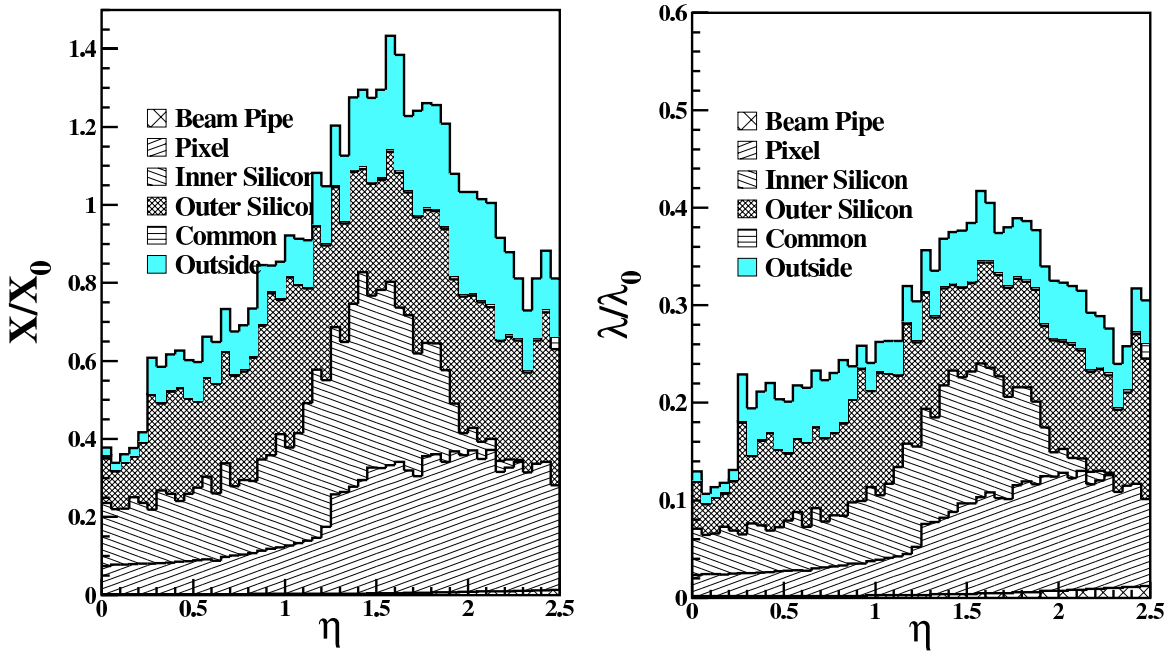


Figure 6.1: Material budget in units of radiation length (left) and in units of interaction length (right) as a function of η for the different subunits.

Therefore a large fraction of the material of the tracker consists of electrical cables and cooling services. Other non-sensitive parts include support structures, electronics, the beam-pipe, and the thermal screen outside the tracker.

The decomposition of the tracker material in terms of radiation lengths and interaction lengths versus η for the different subdetectors is shown in Figure 6.1.

6.1.2 Simulation of the detector response

During GEANT4 (OSCAR) track propagation, the entrance and exit points of particles in the Tracker sensitive volumes are recorded, together with the deposited energy. In GEANT4, low cuts for delta-ray production are used (120 keV and 30 keV for the strip and pixel trackers, respectively) to realistically simulate charge distributions.

The distribution of deposited energy along the track segment inside each sensor volume is estimated by subdividing it into many equal subsegments, small compared with the sensor pitch. Each subsegment is assigned a fraction of the deposited energy using the GEANT4 routine G4UniversalFluctuation, which takes into account Landau fluctuations in thin layers.

The charge from each track subsegment is drifted to the detector surface and simultaneously diffused in the perpendicular plane. The diffusion is assumed to be Gaussian and is proportional to the square-root of the drift length, with the diffusion constant normalized to 2 μm (7 μm) for a 300 μm thick strip (pixel) sensor. The magnitude of the Lorentz drift in a 4 T magnetic field is defined by the drift length and the average Lorentz angle (7° and 23° for the strip and pixel trackers, respectively).

The resulting 1-(2-)dimensional charge distribution is mapped to the Strip (pixel) geometry and the fraction of charge collected by each channel is determined. A list of hit channels for all contributing tracks is formed. If a single channel is hit by more than 1 track, the charge

contributions are added together.

The finite time resolution of the silicon tracker electronics is taken into account by superimposing minimum bias collisions, not only from the current LHC bunch-crossing, but also from a few crossings immediately preceding and following it. The signals from out-of-time particles are scaled-down in size according to the time resolution of the electronics, taken to be a Gaussian of width 12.06 ns. The readout chip of the pixel tracker assigns hits to the correct bunch crossing within a 25 ns time window. This is taken into account in the simulation by accepting only signal and pile-up hits generated in the correct bunch crossing.

All hit channels have a noise contribution added to them that is derived from a Gaussian distribution centred at zero. Minimum ionizing particles crossing 300 μm of silicon on average yield a signal-to-noise ratio of 11 (70) in the strip (pixel) tracker.

To digitize the signal collected by each channel, the charge is multiplied by a gain factor and rounded to the nearest integer, thus simulating ADC digitization. Signals exceeding the ADC range (6 bits in the pixel tracker and 8 bits in the strip tracker) are assigned the maximum allowed ADC value. The Tracker electronics can zero-suppress small signals. In the pixel tracker this is simulated by accepting only signals with a signal-to-noise ratio $S/N > 5$. In the strip tracker, isolated strips are accepted with $S/N > 5$, as are groups of neighbouring strips that all have $S/N > 2$. Furthermore, when a strip failing these criteria lies between 2 successful strips, it is also accepted.

The finite space available in the pixel readout chip limits the size of data buffers and the complexity of circuits. Therefore some data losses will always be present. Additional pixel inefficiencies are expected due to unbonded and noisy pixels. To simulate such inefficiencies, randomly chosen pixels, double-columns, and whole readout chips are erased. The data losses are in general dependent on the CMS trigger rate and the average detector occupancies. They were parameterized according to results obtained from detailed simulations and included in ORCA digitization code. The default inefficiencies used in CMS simulations are 4% for the 4 cm layer at high luminosity and $\leq 1\%$ for all other cases.

6.1.2.1 Tuning of the pixel detector simulation with test-beam data

The pixel detector is the innermost component of the CMS experiment and the most exposed to radiation. At 4 cm from the beam line, the innermost detector layer will be exposed to a fluence of about 3×10^{14} $\mathrm{n_{eq}}$ $\mathrm{cm^{-2}yr^{-1}}$ at the full LHC luminosity. The second and third layers will be exposed to fluences of about 1.2×10^{14} $\mathrm{n_{eq}}$ $\mathrm{cm^{-2}yr^{-1}}$ and 0.6×10^{14} $\mathrm{n_{eq}}$ $\mathrm{cm^{-2}yr^{-1}}$, respectively. In the LHC low-luminosity configuration, the particle fluence at the innermost layer will be about 1×10^{14} $\mathrm{n_{eq}}$ $\mathrm{cm^{-2}yr^{-1}}$.

Particle radiation generates defects in the silicon lattice that can trap charge carriers, therefore the average collected charge is expected to decrease as the detector is exposed. This effect has implications for the detector performance, e.g., the hit detection efficiency. This can be partly compensated by increasing the sensor bias voltage. In addition, in the barrel section the spatial resolution along the r-ϕ coordinates is mainly determined by the Lorentz deflection caused by the 4 T magnetic field. The tangent of the Lorentz angle is linearly proportional to the charge carrier mobility, which itself is a function of the electric field in the sensor bulk. Therefore any change in the bias voltage will change the detector position resolution. Detailed beam-test measurements of the detector response, before and after irradiation, are therefore necessary to optimize the reconstruction algorithms, the detector simulation, and

the operational parameters.

The performance of the pixel sensors was measured in 2003-4 in the H2 beam line of the CERN SPS using 150–225 GeV pions. The beam impact position on the pixel sensor was measured with a beam telescope to a precision of about 1 μm [172]. A detailed description of the experimental setup and of the data reduction can be found in [173]. The measurements were performed on prototype barrel and endcap sensors with 125×125 $\mu\mathrm{m}^2$ pixel cells. A description of both designs can be found elsewhere [174].

Signal, noise, and their ratio measured at different radiation fluences and bias voltages are summarized in Table 6.1 for both barrel and endcap sensors. The signal was defined as the average of the signal distribution recorded by the hit pixel, while noise was the σ of a Gaussian fit applied to the signal distribution recorded during gaps between spills. The measurements were performed with tracks perpendicular to the sensor surface without magnetic field. As discussed above, the bias voltage must be increased after irradiation to collect charge from the full sensor thickness. Large signal-to-noise ratios were recorded even at high fluences.

Table 6.1: Signal, noise, and S/N measured with prototype barrel and endcap sensors of the CMS pixel detector. Measurements are at different irradiation fluences and bias voltages.

Barrel sensors				
Fluence ($\mathrm{n_{eq}\ cm^{-2}}$)	Bias voltage (V)	Signal (ADC counts)	Noise (ADC counts)	S/N
0	150	727.7	11.1	65.2
0.6×10^{14}	150	633.2	8.9	70.9
2×10^{14}	200	654.0	10.9	59.8
6×10^{14}	400	475.6	9.9	48.2
Endcap sensors				
0.6×10^{14}	100	421.7	9.7	43.5
6×10^{14}	450	415.8	9.5	43.8

The hit detection efficiency was measured using tracks at normal incidence to the detector plane. If the collected charge in the pixel predicted by the beam telescope was above a threshold of 5 times the pixel noise, the hit was counted as detected. Pixels cells with high noise or with faulty bump-bond connections were excluded from the analysis. Figure 6.2 shows the hit detection inefficiency as a function of the irradiation fluence for barrel and endcap sensors. For each fluence the bias voltage was set to the minimum value at which charge collection saturated.

For the barrel sensors, the particle detection inefficiency is a few per mil at start-up and increases with the radiation fluence. After a fluence of 6×10^{14} $\mathrm{n_{eq}\ cm^{-2}}$, the inefficiency is still below 2%. The inefficiencies after irradiation are largely due to charge losses along the metal line. By applying a magnetic field parallel to the direction with the larger pixel, side charge carriers are deflected along the direction parallel to the metal line. This implies that the particle detection inefficiency is almost unaffected by the presence of the magnetic field. For the endcap sensors the inefficiency is about 2–4%. This is mostly due to the larger gaps between n+ implants, which are necessary to house the p-stop rings, and is independent of the level of irradiation within the measured range.

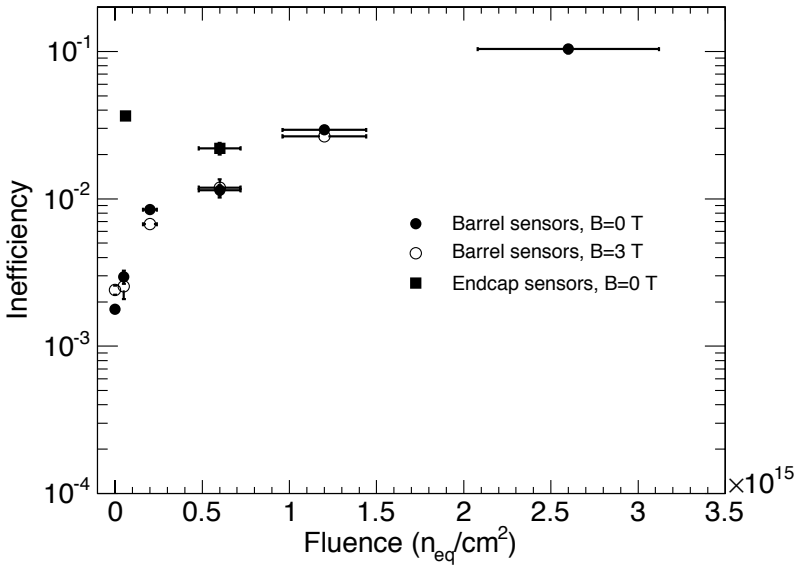


Figure 6.2: Hit detection inefficiency as a function of the irradiation fluence measured with barrel sensors with (empty circles) and without (filled circles) magnetic field. The inefficiency measured with the endcap sensors without magnetic field is represented by the filled squares.

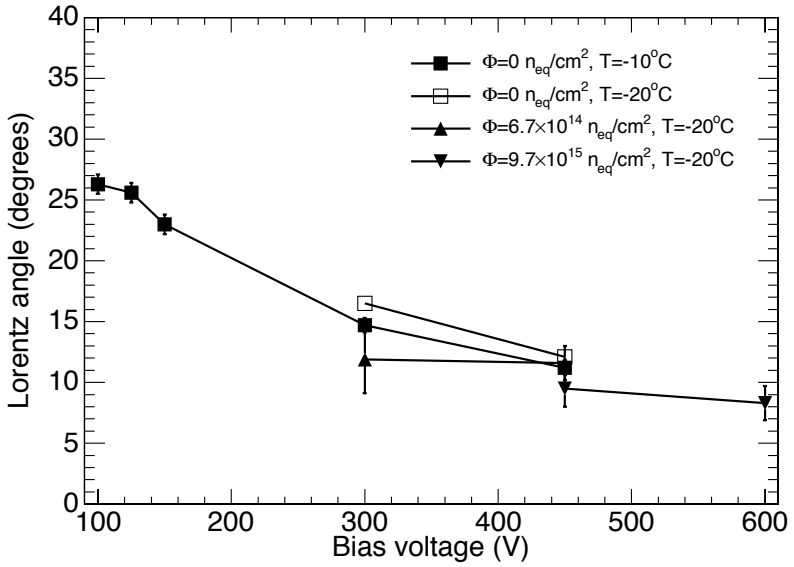


Figure 6.3: Lorentz angle as function of the bias voltage for several fluences and sensor temperature.

The CMS pixel detector will be subject to a 4 T magnetic field. In presence of the magnetic field the Lorentz force deflects the charge carriers from their drift along the electric field lines. The deflection angle, called the Lorentz angle, Θ_L, can be measured with the so-called "grazing angle" technique [175]. The beam crosses the sensor at a shallow angle of 15° and the charge carriers are displaced by the magnetic field parallel to the beam. Figure 6.3 shows

the Lorentz angle extrapolated to 4 T magnetic field (the measurement was done with 3 T) as a function of the bias voltage [176]. The measurements were performed with prototype sensors of the barrel detectors. The Lorentz angle decreases with increasing values of the bias voltage due to the dependence of the charge carrier mobility on the electric field. In addition, the values measured at -20° C are about 2° larger due to the increase of electron mobility at lower temperature. The values measured after heavy irradiation are compatible with the respective values of the unirradiated sensor. However, it must be stressed that after irradiation and in the presence of high leakage Current, the Lorentz angle is not constant across the sensor bulk due to the double-peak structure of the electric field [177]. At the start-up of detector operation, a Lorentz angle of about 26° is expected. After irradiation the bias voltage will be increased to compensate for the trapping of charge carriers, with a consequent decrease of the Lorentz angle.

6.1.2.2 Tuning of the strip tracker simulation with test-beam data

Noise after common-mode noise (CMN) subtraction enters the simulation as the equivalent noise charge (ENC) which is then converted into ADC counts by an adjustable conversion factor. The raw data of a given front-end chip have a common shift from one event to another called common-mode. This shift is subtracted on an event-by-event basis, and the remaining noise is called the "common-mode subtracted noise." The common-mode depends both on the detector electronics and on the environmental conditions, and it is not possible to disentangle them completely. Measurements in a beam-test environment give an estimation on what we will face in the experiment. For TEC modules, the common-mode noise (RMS of the common-mode distribution) was about 1 ADC count. In worse situations, common-mode noise of up to 5 ADC counts has been reported.

Values of the noise measured for the different module types in the strip tracker are listed in Table 6.2. The modules were unirradiated and cooled. The noise scales with the strip length as expected by [178]. Using a linear fit (Fig. 6.4) one can express the noise as

$$\mathrm{ENC\ (peak)} = (38.8 \pm 2.1)\mathrm{e^{-}cm^{-1}} \times \mathrm{L} + (414 \pm 29)\mathrm{e^{-}} \quad (6.1)$$

$$\mathrm{ENC\ (deconv.)} = (51.0 \pm 3.2)\mathrm{e^{-}cm^{-1}} \times \mathrm{L} + (630 \pm 45)\mathrm{e^{-}} \quad (6.2)$$

where the gain was normalized assuming that one MIP leaves $25\,000\ e^{-}$ in a 300 μm sensor. After irradiation of TOB modules with the fluence expected in 10 years of LHC running, the noise at -10° C increased by about 35%, both in peak and deconvolution modes. On irradiated TIB modules at -10° C, the noise showed an increase of about 15% in peak mode, while deconvolution mode stayed unaffected. The simulation should use a conservative 50% increase in noise to account for radiation damage.

The S/N observed in the test beam was 36 (38) for OB1 (OB2) modules in peak mode and 25 (27) for OB1 (OB2) modules in deconvolution mode. In the simulation the default value is 11.

The cluster signal measures the ionization charge collected by adjacent strips. Clusters have typically one or few strips, depending also on the angle between the track and the sensor. The distributions of cluster signals routed by different optical channels are affected by the attenuation in the optical connections, which has large channel-to-channel variations. This spread is partially compensated by selecting a proper output current range of the laser drivers ($\pm$2, $\pm$3, $\pm$4, $\pm$5 mA) eventually reducing the RMS of the distribution of net gains to $\approx 0.2\%$, as

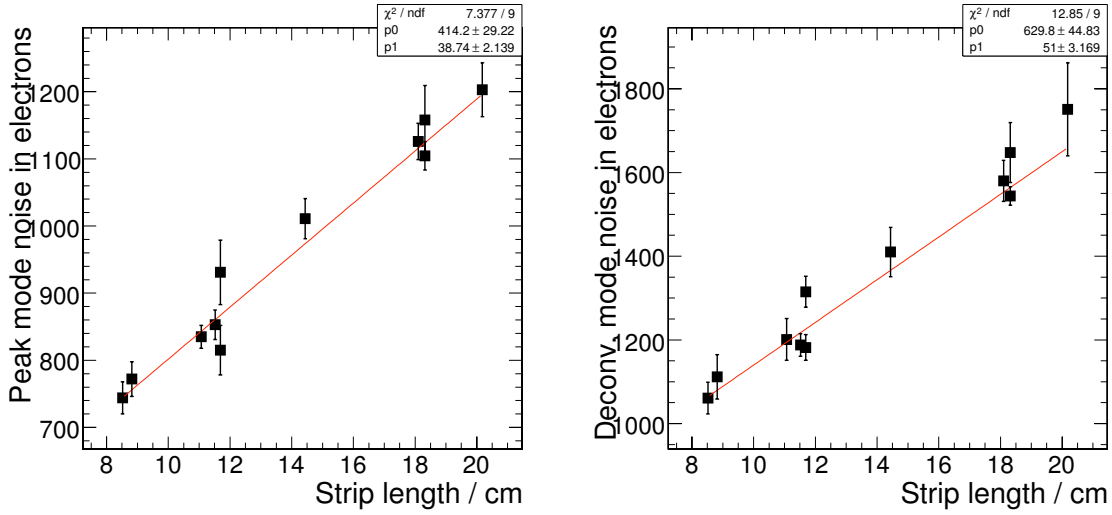


Figure 6.4: Noise after CMN subtraction measured on several module types versus the strip length for 2 APV modes: peak mode (left) and deconvolution mode (right). The results of the linear fits are quoted in Equations (6.1) and (6.2).

Table 6.2: Summary of the common-mode subtracted equivalent noise charge (ENC) for the different module types. The values are given for the APV modes: Peak (inverter on, calibration off) and Deconvolution (inverter on, calibration off).

Module Type	ENC in e^-		Strip Length
	Peak	Deconv.	in mm
OB1	1156±49	1647±78	183.2
OB2	1101±18	1550±23	183.2
IB1	931±48	1315±37	116.9
IB2	815±37	1182±31	116.9
W1TEC	744±24	1061±38	85.2
W2	772±26	1112±53	88.2
W3	835±17	1201±50	110.7
W4	853±22	1188±27	115.2
W5	1011±30	1410±59	144.4
W6	1126±27	1580±49	181.0
W7	1203±40	1751±111	201.8

explained in Section 6.3.2. In the test-beam data, the cluster signal most probable value in ADC counts had an RMS of 12% in peak mode and 10% in deconvolution mode. The simulation assumes a constant conversion of 313 e^- per ADC count, based on the agreed hardware setting.

The interstrip coupling is implemented in the simulation by simply moving fractions of the charge of the hit strip to the neighbours. The fraction of charge put on one neighbour is by default 12%.

Under the influence of a magnetic field, drifting charge is deflected by the Lorentz angle, which can be expressed as

$$\tan\Theta_L = \mu_H\, B = r_H\, \mu\, B \tag{6.3}$$

A parameterization of the drift velocity of electrons in an electric field was suggested in [179] that can be translated into a formula for the mobility:

$$\mu = \frac{\mu_{\rm low}}{(1 + (\frac{\mu_{\rm low} E}{v_{\rm sat}})^{\beta})^{1/\beta}} \tag{6.4}$$

This formula is used in the simulation together with the following parameters applicable for drifting holes:

- Low field mobility $\mu_{\rm low} = A\,(T/300{\rm K})^{-2.5}$. For unirradiated detectors $A = 470.5\,{\rm cm}^2/{\rm Vs}$ and for irradiated detectors it will change to $A = 460\ {\rm cm}^2/{\rm Vs}$ [180]
- Fitting parameter $\beta = 1.213\,(T/300{\rm K})^{0.17}$
- Saturation drift velocity $v_{\rm sat} = (8.37 \times 10^6\ {\rm cm/s})\ (T/300\ {\rm K})^{0.52}$
- Hall scattering factor $r_H = 0.7$

When the detectors are biased with a voltage well above depletion the electric field can be assumed constant through the detector thickness D and the mean electric field is $E=V_{\rm bias}/D$. Finally, it is shown in [180] that this parameterization is in good agreement with data.

The time resolution was studied using the TOB Cosmic Rack. Two parallel modules were exposed to perpendicular muons. The first module (in peak mode) was used to measure the location of the track and to identify the strips hit by the muon on the second module (in deconvolution mode). The signal of the highest strip and its neighbours were measured on the second modules for different delays of its clock and trigger signals using the Phase Locked Loop (PLL) chip. The result is shown in Figure 6.5 where the average pulse height (pedestal subtracted) is plotted as a function of PLL delay. The Gaussian fit on the highest strip gives a σ of 13.75 steps of the PLL delay, or $\sigma = 14.3$ ns. The default value in the simulation was 12 ns. The shape of the pulse depends on the operating conditions and on the settings of the APVs. A σ of 13.2 ns was observed in a previous beam test.

Hit efficiency was studied on the TOB Cosmic Rack. The setup was equipped with 2 double-sided rods (to measure both x and z coordinates) located in the first and last layer with 4 single-sided rods in between. Three of the single-sided rods contained modules with a strip pitch of 183 μm and one module with a strip pitch of 122 μm. Tracking was done using CMS Reconstruction software. First a manual alignment was performed by maximizing the number of 6-hit tracks and minimizing χ^2. The hit-finding efficiency reached with this setup was between 99.2% and 99.8%.

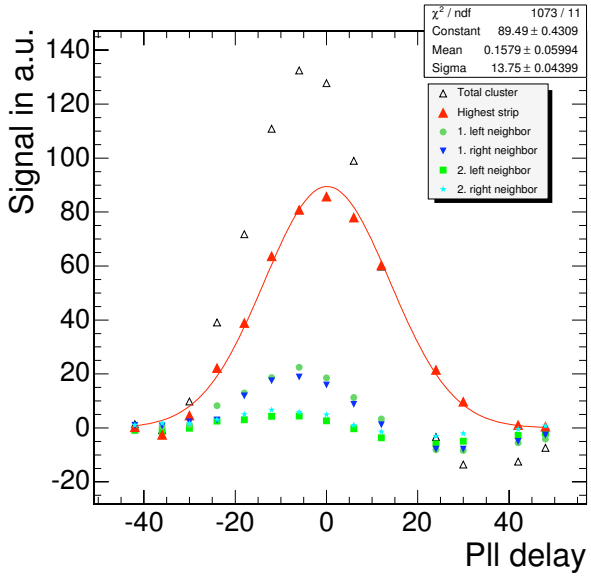


Figure 6.5: Signal shape in deconvolution mode taken with muons for sensors at room temperature. The pulse shape in deconvolution mode is implemented as a Gaussian peak in the simulation. Therefore, the data have been fitted with a Gaussian distribution with σ=14.3 ns.

6.2 Data handling

6.2.1 Databases

Flexible configuring and operation of the tracker hardware and efficient implementation of the software require the use of a number of databases (DB). As described in Section 2.8, the offline Conditions DB, referred to simply as the offline DB throughout this chapter, is accessed by the reconstruction software through an interface that is common to all CMS subdetectors. Data access to the offline DB is needed in nearly all steps of the reconstruction in the Tracker detectors. Three categories of tracker data can be found in the offline DB. The first category corresponds to the data used to configure the readout electronics: examples of such data are the strip pedestals and noise values. Some of these data are also needed by the reconstruction software. The second category consists of non-event data acquired during data taking. These are the data produced by the environmental and power supply monitoring sensors. The third category is represented by the actual calibrations that are computed offline using the event data as well as data from the first 2 categories. Data from the first 2 categories reside in the online Configuration and Conditions DBs, which are specific to each subdetector.

The other tracker-specific DB is the Construction DB. The goals of the Construction DB are

- to track the fabrication and history of components;
- to allow for storage and retrieval of results of the various tests performed on each component; and
- to store the dimensions of each component and its location in the tracker.

The information about detector dimensions and their locations is needed to validate the description of the Tracker as used in the reconstruction software. In addition, it will be necessary to transfer some of the information stored in the Construction DB to the offline and

online DBs.

The status of the DBs specific to the strip and pixel detectors is described as follows.

6.2.1.1 Silicon-strip databases

The CMS Tracker is composed of $200\,000$ components of about 20 different species assembled and tested in 20 laboratories around the world. Each component has to be identified, tested, and assembled into higher level components.

Two implementations of the construction DB exist for the Strip tracker. One of them is used in the construction phase of all the subdetectors (TIB, TID, TOB, and TEC) and in the integration phase of the TEC and TOB. A different implementation is used for the integration of the TIB and TID.

The strip construction DBs contain information on the electrical, Optical, and cooling connections up to the patch panel. In addition, the TIB and TID Construction DB will also contain the precise survey measurements of the locations of all the modules on the support structures. These measurements represent the initial alignment constants, which therefore must be copied to the offline DB.

The strip-tracker online Configuration DB has to store a large number of parameters relevant to the configuration of the control and readout electronics.

- FEC parameters: each full version of the control configuration of the whole strip tracker contains about 1.68×10^6 parameters.
- FED parameters: each one of the 400 strip-tracker FEDs requires some 10^5 parameters. More than 90% of these parameters are represented by the pedestals, noise, and bad-strip flags.

An Oracle DB management system is used to store and handle this information. A versioning system is used to track the history of all changes in the parameter values.
The strip-tracker online DB system will also contain information about all active readout components and their connections. This information is obtained by specialized software procedures (described in Section 6.3.2) and then automatically filled into the DB. A graphical user interface, which allows parameters stored in the DB to be edited by properly authorized users, is also available. During data taking, the strip online DB will also receive data from the sensors for temperature, humidity, voltage, and leakage current that are installed on each module (some are also installed on the FEDs).

All configuration parameters used for a given run and the recorded conditions will have to be transferred to the offline DB since they represent conditions data.

Error situations detected online have to be logged and stored for later use. A first prototype of software architecture for error logging exists and has been recently deployed.

6.2.1.2 Pixel databases

A single pixel-specific DB exists. It has been designed with several objectives in mind, the principal ones being

- To keep track of the large number of components during the complex detector fabrication process, 66×10^6 pixels;
- To monitor the performance of the individual pixel elements;
- To implement quality-check procedures and select components suitable for detector construction;
- To store all information necessary for data analysis.

During the construction phase the pixel DB will reside at Fermilab. The DB will be moved to P5 at CERN, prior to commissioning. An XML-based data transfer mechanism will be used to store data from remote institutions. In designing the database, the principle that data once entered in the database should not be deleted has been adopted. The data can be updated (changed), but not deleted. Each table has associated history information which can be used to monitor the behaviour of various components with time.

The pixel DB will enable tracking of every component used to fabricate the pixel detector. One attribute that deserves special mention is the "pixelinfo" attribute that contains the 4 trim bits and 1 mask bit for each of the approximately 66×10^6 pixels. A mask bit allows a particular pixel to be enabled or disabled while the trim bits allow the threshold of an individual pixel to be adjusted. The trim bits and mask bit for all 4160 pixels in a ROC are encoded and stored as a large binary object. For the pixel detector the configuration data needed to be downloaded totals about 1.2 GB. Pixel information necessary for data analysis includes the gain and pedestal for each pixel. The pixel gain and pedestal are determined offline and are stored in the Conditions DB.

The pixel DB will need different interfaces for data input, data retrieval, and display. While a graphical user interface will be used to access and retrieve certain categories of information into the DB, other types of interfaces will be needed for other data groups, e.g., calibration data. An interface will also be needed for component selection for detector construction.

6.2.2 Data format and size

6.2.2.1 Silicon strip

The FEDs receive analogue data from up to 48 front-end modules, which are digitized, processed, and formatted before being forwarded to the global DAQ. The FED has 4 readout modes, as described below.

- **Zero-Suppressed**: this is the default mode for proton-proton collisions. The raw detector data are zero-suppressed, using firmware algorithms implemented within FPGA devices, so that only useful signal information is transmitted to the DAQ computing farm. The algorithms perform pedestal and common-mode subtraction before identifying signal-above-noise thresholds (by default, a factor of 2 above the noise, except for isolated hits, which require a factor of 5 above the noise). The pedestal and noise calibrations used by the algorithms are described in Section 6.3.2. For an event sample of dijets with transverse momenta in the range of $50-100$ GeV and a luminosity of 10^{34} cm^{-2}s^{-1}, the average FED buffer and event sizes are 1.3 kB and 496 kB, respectively, for an average strip occupancy of 1.2%. A root compression factor of 1.7 is achieved when writing the raw FED buffers to disk.

- **Virgin Raw**: in this mode, the FED performs no processing and only formats the raw detector data. This mode is intended for use during detector commissioning and also facilitates debugging. The FED buffer size is fixed at 49.5 kB.
- **Processed Raw**: the data are pedestal-subtracted and reordered to reflect the detector strip ordering (this is not the case for the APV25 data streams). However, the data are not zero-suppressed and it is envisaged that this mode will be used for heavy-ion collisions for which high occupancies are expected. The FED buffer size is again fixed at 49.5 kB.
- **Scope mode**: in this mode, the FED simply captures the APV25 data stream observed within a given time window (of configurable size). This mode is again used during detector commissioning and is useful for system debugging. The FED buffer size depends on the size of the time window and has a theoretical maximum of 196 kB.

The standard DAQ header and trailer words, as described in Section 2.3.2.1, are appended to the data payload. Additional header words specific to the SST contain information that allow checking of the operational states of the FED and the connected front-end APV25 chips. The SST-specific headers take 2 different formats: a verbose "Full Debug" format and a lightweight "APV Error" format designed for use during physics runs. The data within the FED buffers have to be formatted appropriately (into "digis") to be used by the reconstruction software. For an event sample of dijets with tranverse momenta in the range $50-100$ GeV and a luminosity of 10^{34} cm^{-2}s^{-1}, the time taken to format all FED data from an event is 140 ms on a 2 GHz PC.

6.2.2.2 Pixel FED data format and event size

The pixel FED readout board will read raw data from 36 input links, build events, and send the event packets through via S-LINK [41, 42] connections to the DAQ. During the event building the FED has to reconstruct pixel addresses from 6-level analogue signals. This procedure requires a set of pre-programmed thresholds. The determination of the thresholds will be part of the calibration procedure. Otherwise no processing, except error monitoring, is performed by the pixel FED.

Each FED will send to the DAQ a data packet starting with the standard DAQ header and ending with the standard DAQ trailer described in Section 2.3.2.1. Between the header and the trailer there will be a number of 32-bit data words, each containing the information from 1 pixel. The format of the 32-bit word is:

- 6-bit Link id, defines the input link to the FED (varies within 0–35);
- 5-bit ROC id, defines the ROC within one link (0–23);
- 5-bit Double-Column id, defines the double column within a ROC (0–25);
- 8-bit Pixel id, defines the pixel address within a Double-Column (0–179);
- 8-bit ADC value, the signal amplitude, extracted from a 10-bit ADC.

The FED data volume resulting from the format described above corresponds to $4 \times (3+$ number-of-pixels) bytes. The pixel readout has been arranged in such a way as to approximately fit the requirement of 2.0 kB per FED. For high (low) luminosity pp collisions, the pixel data volume will be about 72 (22) kB per event. The pixel detector average occupancies for the 3 barrel layers are 0.00062, 0.00027, and 0.00016, respectively. In the forward endcaps

the average occupancy is 0.00015, but it varies by a factor of 2 with the radius. The FEDs are arranged to provide uniform average data rates transmitted over the S-LINKs. The average S-LINK rate for the 100 kHz L1A rate at high luminosity will be 190 MB/s, with the highest occupancy FED running at the DAQ limit of 200 MB/s.

When reading data over VME, an alternative "transparent" data format is foreseen in addition to the standard one. This is a raw data format where the ADC values are not interpreted but are just sent out. This format will be used only during the initial level calibration procedure described in Section 6.3.1.

The FED formatted data will be transformed into the internal "digis" format. One pixel digi occupies a 32-bit data word and has the following format:

- 9-bit Column address within a detector module;
- 8-bit Row address within a detector module;
- 8-bit ADC value, the signal amplitude.

When pixel digis are stored on disk they have to be assigned to detector modules. The empty 7-bit field left in the 32-bit word mentioned above is too small to code the module ID. Therefore a header for each detector module is required with the det-unit ID and the number of hit pixels (digis). Assuming that such a header is 64 bits, the total pixel data volume on disk will be 84 (34) kB for high (low) luminosity.

The conversion of the FED data into the digi format is performed by a software module that, making use of cabling information, can assign the FEDs output data to the correct detector modules. A prototype of such a module is now available and its preliminary performance is 10.5(3.5) ms on a 2.8 GHz PC.

6.3 Configuration of the tracker hardware and electronics

Optimum detector performance and efficient track reconstruction requires an appropriately configured, fully synchronous, and calibrated readout system. This is achieved through procedures that allow optimization of the parameter sets uploaded to the digital registers of the various hardware components comprising the tracker readout system. A brief overview of the tracker readout system can be found in Section 1.5.5.3. An overview of the procedures for both the pixel and strip tracker readout systems is given below. Performances are discussed in Volume 2 of this document.

6.3.1 Configuration of the Pixel tracker

The pixel calibration procedures are divided into 2 groups. The first group involves very basic detector parameters that can be determined within the standalone pixel system where no additional information from other CMS subdetectors is required. These calibrations will be run on the local DAQ system which will be based on XDAQ [38, 39, 40] clients. The calibration tasks are

- Verify the DAC settings for all front-end components (ROCs + TBMs).
- Verify the clock's fine phase in the front-end components and in the ADCs (FEDs). The optimization of the fine phase of the 40 MHz clock will be measured by checking the acceptance for low amplitude pulses.

- Check the level thresholds for analogue address recognition and offsets for analogue signals. Verify that the analogue coded pixel addresses are clearly recognizable (see the Tracker TDR [6] for details) with the various address levels being cleanly separated. The analogue pixel charge signal coming from the 10-bit ADC should have the right gain and offset to fit correctly into the 8-bit data field foreseen in the output data packet. After this step we can switch from the transparent mode (the mode in which the FED transmits just the ADC values without interpreting them) to the standard FED mode (the event building mode with pixel addresses correctly decoded).
- Identify noisy and dead pixels. This will be done by a scan where a calibration signal is injected into each pixel.
- Determine the noise and threshold for each pixel and adjust the trim bits. An "S-curve" scan (efficiency versus pulse amplitude) will be made for each pixel. From the shape of the "S-curve" the pixel noise can be derived. The pixel threshold will be adjusted to be, on average, equal to 5 times the average noise. The thresholds will be made uniform by adjusting the 4-bit trim in each pixel cell.
- Calibrate the pedestals and gains for each pixel. This will be measured by checking the response of each pixel to a set of 5–10 calibration signals of different amplitudes. The response will be linear in the middle range from which the gain and pedestal can be derived. Within the non-linear low and high gain ranges the response will be parameterized with a more complex non-linear formula.

All these tasks will be done during the "off-beam" periods using charge injection. The parameters that are calibrated should not vary rapidly in time. At the beginning, soon after detector commissioning, it is planned to run the calibration tasks every few days (2–3). Later, when the detector is more stable, performing the calibration on a weekly basis should be sufficient. Neither the central DAQ nor the global trigger system will be needed to perform these measurements.

The second calibration group involves more advanced parameters which have to be determined using information from other subdetectors. This can be only done in the filter farm during full reconstruction. This group includes

- Local, internal pixel synchronization and global synchronization with the rest of CMS. The fact that pixel tracks can be reconstructed and that they match tracks in the rest of the tracker will be proof that the global pixel synchronization is correct. If this is not the case the L1A delay will be scanned to find the correct bunch crossing. This item will be very important during detector commissioning; during physics data taking this will be monitored.
- Tuning of the high voltage bias settings. This will not be very critical at the beginning when the detector is not yet radiation damaged. Later it will need adjustments but they are expected to vary slowly and smoothly with time. The simple adjustment can be made by scanning the pixel analogue signal corresponding to a MIP as a function of the bias voltage and observing its saturation. A more advanced method uses tracks which pass through the pixel detector at shallow angles (the so-called grazing-angle technique). When the bias is high enough the charge deposited in all pixels along the track should be roughly equal.

None of the measurements mentioned above requires any special control procedures or special triggers.

For most of the pixel calibration tasks the algorithms developed during the pixel module testing will be used. Only small changes will be needed to adopt them to the local DAQ (XDAQ) and filter-farm (ORCA) environments. Some of the procedures will produce large quantities of parameters which will have to be stored in databases.

6.3.2 Configuration of the Silicon Strip Tracker

A series of procedures have been developed to configure, synchronize, and calibrate the Silicon Strip Tracker (SST) readout system and these are primarily concerned with the configuration of the front-end electronics (APV25 and other ancillary chips) and the off-detector FED boards. A brief overview of the procedures is given below.

- **Identifying hardware devices and detector partitions**: Prior to configuration and synchronization of the readout system, all hardware devices within the individual detector partitions (defined as a group of devices sharing a common trigger source) must be identified. Procedures comprising VME crate scans are used to identify all front-end devices (via the FEC and slow control system) and FEDs. All devices are initially configured with default parameter sets. This procedure is performed whenever the off-detector electronics is changed or hardware failures are identified.
- **Identifying the cabling of the optical readout system**: An automated procedure determines the cabling of the optical readout links, a highly desirable feature for a system comprising $\approx 40\,000$ fibres. In addition, the procedure allows checking of the detector partitioning and identification of cabling errors and faulty optical links, and provides a cabling map that is used by the offline reconstruction software. This procedure is performed whenever the off-detector FED boards are changed or a link fails.
- **Synchronization**: Efficient track reconstruction is only possible with a synchronous readout system, as signal magnitude is attenuated by up to 4% for every nanosecond in timing misalignment for the default operating mode of the APV25 chip. Synchronization of the SST to LHC collisions comprises 2 steps.

  First, all the front-end electronics of the SST are relatively synchronized. Relative differences in clock phase at the front end due to the different positions of the front-end modules within the TTC system (that distributes clock and triggers) must be accounted for to synchronize the front-end electronics. This is achieved using information within the APV25 data streams, captured by the FEDs, to measure the differences in clock phase at the front end. This procedure assumes that the fibre lengths of the readout optical links, which differ and therefore also introduce relative phase shifts, are known and accounted for within the FED configuration.

  Second, the phase of the clock distributed to the SST readout system must be adjusted to synchronize the entire front end to LHC collisions. This is achieved through a "fine delay scan" that involves applying a range of phase shifts to the clock to reconstruct the signal pulse shapes sampled into the APV25 pipeline memories. This allows identification of the optimum sampling point and is effec-

tively a measurement of the trigger latency. Corrections for the time-of-flight are also required, which are provided either by simulation or measurements using the aforementioned delay scans. This procedure is likely to be performed infrequently, perhaps weekly to monthly, depending on the system stability.

- **Optimizing the performance of the APV25 chip**: The APV25 chip has several configurable digital registers that allow optimization of various aspects of its operation and performance, including tuning of the analogue pulse shape that is sampled into the on-chip pipeline memory to maximize signal-to-noise and improve bunch-crossing identification, tuning the signal levels to ensure optimal use of the available dynamic range, and tuning the gain of the on-chip multiplexer stage to improve gain matching among readout channels. These operations will be performed with varying frequencies, ranging from daily to monthly.
- **Gain matching and biasing of the optical link system**: This procedure ensures the optimal use of the available dynamic range for each link and optimum gain matching across the entire link system by tuning the parameter set uploaded to the digital registers of the analogue laser drivers on the front-end modules. The laser components are sensitive to changes in the environmental temperature, thus this procedure should be performed whenever there are significant changes in temperature.
- **Calibration constants for the FED**: The FED is a complicated board with many configurable aspects, the most important of which is the population of the pedestal and noise look-up tables that are used by the data zero-suppression algorithms of the FED. Pedestal and noise determination are required prior to every physics run.

The principles for configuring and synchronizing the SST are well understood and prototype implementations exist for all procedures. These procedures have been extensively tested during beam tests and with several prototype systems as large as $1/8$ of a partition (≈ 400 modules). Much of the required software is now implemented within the offline software frameworks (both in the previous ORCA framework with its interface COSINE to the online system, and now in CMSSW). This allows the software to be run using either local DAQ resources, the default operating mode during detector commissioning, or global DAQ resources, which permits distributed analysis using the extensive processing power of the filter farm and fast detector readout speeds.

6.4 Track reconstruction

Track reconstruction in a dense environment needs an efficient search for hits during the pattern recognition stage and a fast propagation of trajectory candidates. In the CMS Tracker, the first task is simplified by the arrangement of sensitive modules in layers that are practically hermetic for a particle originating from the centre of the detector.

The second task uses the fact that the magnetic field is almost constant in a large part of the tracker volume and also that most of the support structure is concentrated on the layers, close to the sensors. During reconstruction the typical step length for propagation of track parameters is on the order of the distance between 2 layers and a helical track model is adequate. For reconstruction purposes the detailed distribution of passive material as used in the simulation is replaced by an attribution of material to layers. This model simplifies the

estimation of energy loss and multiple scattering, which can be done at the position of the sensitive elements without requiring additional propagation steps.

The track reconstruction is decomposed into 5 logical parts:

- Hit reconstruction, which in turn consists of clustering of strips or pixels and estimating a position and its uncertainty.
- Seed generation
- Pattern recognition or trajectory building
- Ambiguity resolution
- Final track fit

These parts are described in the following subsections. The section is completed by a description of the reconstruction performance and a description of track reconstruction in heavy-ion events. More detailed information on track reconstruction can be found in [181].

6.4.1 Tracker clusterization

6.4.1.1 Silicon strip tracker clusterization

Strip tracker cluster reconstruction usually begins with the gain-calibrated zero-suppressed strip data. The gain is expected to vary by up to 20% across APV-pairs and by a few percent among the channels within an APV-pair. During commissioning, the FEDs may be operating in "virgin raw" mode (Section 6.2.2.1), in which case the software must also subtract the pedestals. The gain and pedestals will each require about 20 MB in the offline DB, as will the RMS noise that is required for cluster reconstruction itself. These should be fairly stable over time, with the possible exception of the APV-pair gain, which is temperature dependent. It will be possible to store the gain-calibrated strip data on disk to minimize the need for database access when reprocessing the data.

Clusters are reconstructed in the strip tracker by searching for a seed strip with a signal to noise ratio $S/N > 3$. Nearby strips are included in the cluster if they satisfy $S/N > 2$. Holes are allowed inside clusters only if they are assigned to highly inclined tracks. The total signal size of the cluster must exceed 5 times the square-root of the sum of the RMS-noise-squared of the individual strips inside it.

The cluster position is usually determined from the centroid of the signal heights. The large interchannel coupling in the strip tracker makes it non-trivial to determine the position from the cluster edges, as is done for the pixel tracker. Nonetheless, this is done for very wide clusters (containing at least 4 strips).

The position resolution is parameterized as a quadratic function of the projected track width on the sensor in the plane perpendicular to the strips. The parameters are also dependent on whether the observed cluster width is compatible with the track width and cluster position. Figures 6.6a and 6.6b show the observed resolution in simulated data, together with the fitted parameterization of it, for 2 principal categories of reconstructed clusters.

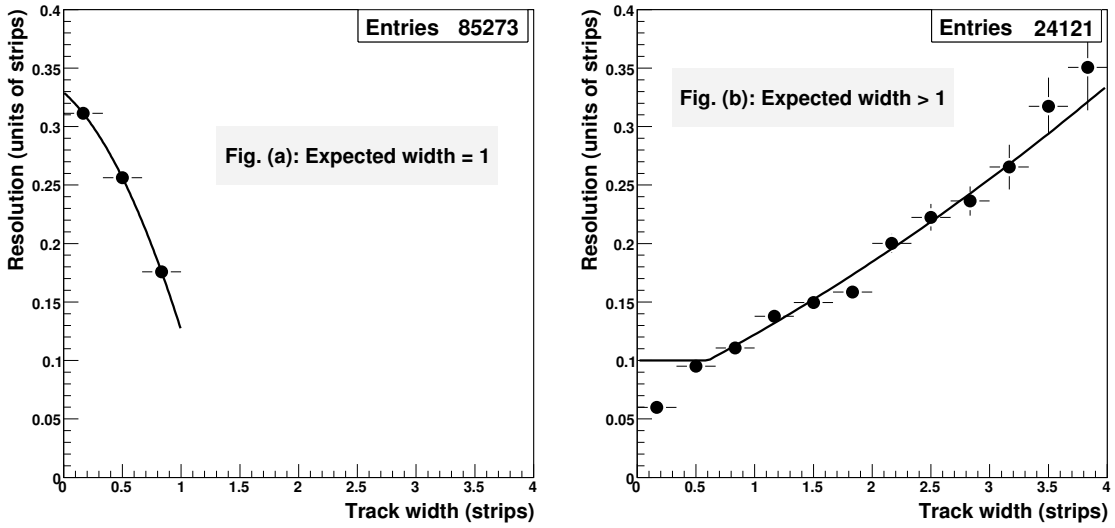


Figure 6.6: Strip tracker cluster position resolution in simulated data as a function of the projected track width. The 2 figures correspond, respectively, to clusters whose expected width (derived from the track angle and reconstructed cluster position) is (a) 1 strip and (b) more than 1 strip.

6.4.1.2 Pixel tracker clusterization

The cluster reconstruction algorithm for the pixel detector starts from a cluster seed, defined as a pixel with $S/N > 6$. It then adds pixels adjacent to the cluster if they have $S/N > 5$, continuing this process until no more adjacent pixels are found. Diagonally adjacent pixels are considered adjacent. Finally the cluster is retained if its total charge has $S/N > 10.1$. The same algorithm is applied to the barrel and forward pixel detectors.

The position of pixel clusters is estimated independently in both dimensions. It is based on the relative charges of the pixels at the edges of the cluster and the associated reconstructed track angle. Depending on the detector module orientation, both the track inclination and the Lorentz shift can contribute to the charge sharing. More details of the procedure together with the formulae used are given in [182].

The algorithm used needs as an input parameter the expected width of the charge distribution collected on the sensor surface. A precise charge width estimate is performed if the impact angles of the particle on the detector are known from the partial track reconstruction, otherwise the track impact angles are estimated from the polar and azimuthal angles of the hit modules with all tracks assumed to originate from the nominal interaction point.

The error on the reconstructed position is estimated from the spatial displacement between simulated and reconstructed hits (residuals) and error parameterization is performed as a function of the cluster size and the rapidity. An additional and more precise error parameterization is performed when the track impact angles are available from the partial or complete track reconstruction. A detailed description of the error and position assignment of pixel hits can be found in [182].

The reconstruction inefficiency is defined as the fraction of simulated hits that do not have any associated reconstructed hit. For the pixel detectors this is below 0.5%. The fraction

of reconstructed hits that is not associated with any simulated hit (ghost hits) is less than 0.01%. Both the reconstruction efficiency and the ghost rate quoted here do not take into account readout inefficiencies.

Typical simulated resolutions, for unirradiated sensors, are better than $15\ \mu\mathrm{m}$ in the barrel in the transverse direction and vary between 15–30 $\mu\mathrm{m}$ for the barrel longitudinal direction and for the endcap disks. The position resolution before and after irradiation was measured in test beams with prototype sensors and a 3 T magnetic field. Along the r-ϕ direction the measured resolution is between 4 and 15 $\mu\mathrm{m}$ for clusters of more than 1 pixel, depending of the impact angle and on the irradiation fluence. To estimate the position resolution in CMS conditions a sensor simulation was implemented which used the electric field profile measured with irradiated sensors and trapping of charge carriers. A detailed description of the simulation program and of the data analysis can be found in [173]. Figure 6.7 shows the spatial resolution along the r-ϕ direction as a function of the angle between the track direction and the normal to the sensor plane. The position resolution before irradiation depends only weakly on the impact angle. After a fluence of $\Phi = 6.7 \times 10^{14}\ \mathrm{n\,cm^{-2}}$, the bias voltage is increased to compensate for the charge losses. Due to the decreased amplitude of the Lorentz deflection, the resolution degrades at large ϕ angles. A bias voltage of 450 V provides the best position resolution at this fluence. At a fluence of $\Phi = 9.7 \times 10^{14}\ \mathrm{n\,cm^{-2}}$ the resolution further deteriorates, however the optimal bias voltage is still below 600 V.

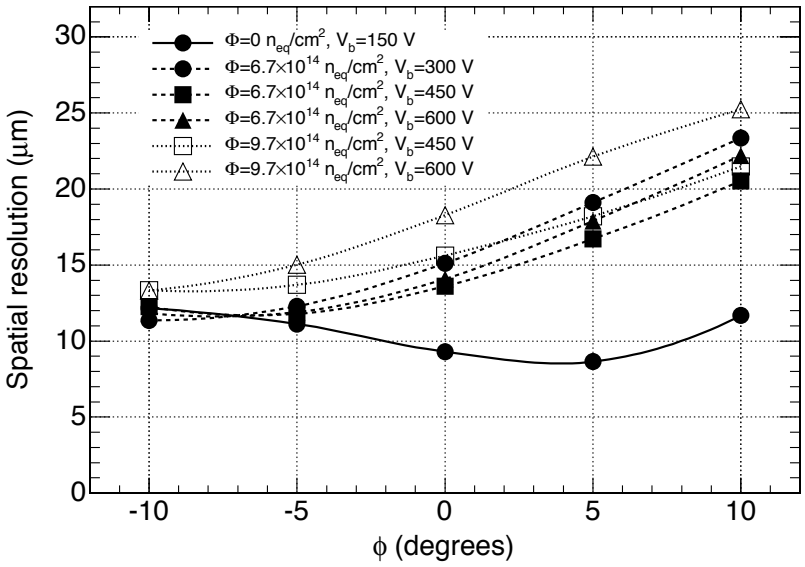


Figure 6.7: Pixel position resolution along the r-ϕ direction as a function of the angle between the track direction and the normal to the sensor plane. The resolution is calculated for an unirradiated sensor (solid line) and for sensors exposed to an irradiation fluence of $\Phi = 6.7 \times 10^{14}\ \mathrm{n\,cm^{-2}}$ (dashed lines) and $\Phi = 9.7 \times 10^{14}\ \mathrm{n\,cm^{-2}}$ (dotted lines). For the irradiated sensors, the position resolution is calculated at different bias voltages.

In the default mode the pixel simulation is tuned to unirradiated sensors, therefore the best position resolution is used in most physics studies.

6.4.2 Seed generation

Seed generation provides initial trajectory candidates for the full track reconstruction. A seed must define initial trajectory parameters and errors. They can be obtained externally to the Tracker, using inputs from other detectors, but the precision of initial trajectory parameters obtained in such a way is, in general, poor. Another way is to construct seeds internally. In this case each seed is composed from the set of reconstructed hits that are supposed to come from 1 charged particle track. Since 5 parameters are needed to start trajectory building, at least 3 hits, or 2 hits and a beam constraint, are necessary. If the beam constraint is used it is removed during the final fit. Hits that are seed constituents are provided by the dedicated reconstruction.

6.4.2.1 Regional seed generation

Although the external information is usually not sufficient for full seed definition it is still useful to constrain the search area for its constituents.

A major improvement in the current implementation is the possibility to do the reconstruction in the region of interest only (regional reconstruction), which is now appearing to be an important requirement for the online software. The key concept for the new regional reconstruction software is called the "TrackingRegion." It is meant to be an abstract definition of the region of interest with kinematical constraints. The TrackingRegion specifies the direction around which the region is defined, the (signed) inverse transverse momentum range, and the allowed position of the track impact point (vertex along beam line and maximum allowed distance from vertex in the transverse plane and along the beam line). Two concrete implementations are provided: "GlobalTrackingRegion" and "RectangularEtaPhiTrackingRegion." The latter allows the direction to be constrained within a given range of η and ϕ. Regional seed generation reconstructs sets of hits which are compatible with a track passing the kinematical requirements of the TrackingRegion.

6.4.2.2 Choice of the layers for hit finding

In the baseline, seeds are defined by pairs of pixel hits. The pixel detector is well suited for seeding purposes because of its low occupancy. Furthermore, pixel hits are the most precise, close to the beam pipe, and have both r-ϕ and z-r coordinate measurements. To assume maximal efficiency for finding hit pairs in each direction, 3 combinations of layer pairs are used.

6.4.2.3 Hit-pair finding

The finding of hit pairs that are constituents for seeds can be decomposed into the following tasks:

- First hit finding—searching for hits in the region of interest. It is straightforward for the GlobalTrackingRegion, where all hits from a given "DetLayer" (ORCA terminology corresponding in the case of the pixel detector to a forward disk or barrel layer) are compatible. For the RectangularEtaPhiTrackingRegion the allowed range for hit positions is predicted analytically. The compatible hits are collected only from the "DetUnits" (corresponding to a module in the case of the pixel detector) compatible with the prediction.

- Second hit finding—for each selected hit, inner DetLayers are searched for a second hit compatible with the first one. The analytical prediction for the hit's position is computed independently in the r-ϕ and r-z planes. Since hits from a given DetLayer are accessed more than once, the result is stored in a cache. The hits in the cache are sorted in ϕ, which allows operation on hit ranges rather than hit-by-hit accessing.

The performance of the hit-pair finding is shown in Figures 6.8-6.10.

6.4.2.4 Conversion of hit pairs to seeds

Seed creation is much more computationally intensive than just the 2 (at least) hits used for its definition. Seed construction involves computation of the "FreeTrajectoryState," construction of an error matrix, and definitions of the first 2 "TrajectoryMeasurements" and a "TrajectoryStateOnSurface." The time to generate a seed is about 0.3 ms/seed.

6.4.2.5 Additional seed generators for specific purposes

Although the time for single seed creation is small, the large number of combinatorial hit pairs may lead to significant and unacceptable time spent in the overall seed generation. To reduce the combinatorics, additional special-case seed generators are provided:

- **Triplet-based seed generation:**
  This involves reconstruction of pixel-hit triplets instead of pairs. Since the pixel detector has only 3 layers, the reconstruction cannot be fully efficient. These seeds can actually be used as tracklets in applications that require high speed, like the first stages of the High Level Trigger. Performance plots for pixel triplets are shown in Figure 6.11.
- **Seed generator from hit pairs with vertex reconstruction:**
  This involves reconstructing primary vertices with triplets of pixel hits. The reconstructed vertices are used as a vertex constraint of the tracking region. A reduction of ghost-hit pairs by a factor of 10 (50) can be obtained for low (high) luminosity. The efficiency depends highly on the event type (60–100%).

Another type of seed generator is the generator based only on hits from the strip tracker. It can be especially useful for the detector at start-up, when the pixel detector will not yet be installed, or installed in an incomplete configuration. The hit-pair finding algorithm described above is applied to pairs of layers of silicon-strip detectors.

The efficiency of the pixel-less seed generation for one particular choice of the seeding layers (first 3 TIB layers, the TID layers, and the innermost 2 rings of the first 3 TEC layers) is shown in Figure 6.12 for a sample of muons with transverse momentum of 1 GeV/c. This figure also shows the global efficiency of the pixel seeding, which is limited by the geometrical acceptance of the pixel detector.

It is also possible to combine the pixel layers with some silicon layers. An example of "mixed" seeding is also shown in Figure 6.12. In this case the 2 innermost rings of the first 3 TEC layers were added to the pixel layers to extend the coverage at high pseudorapidity.

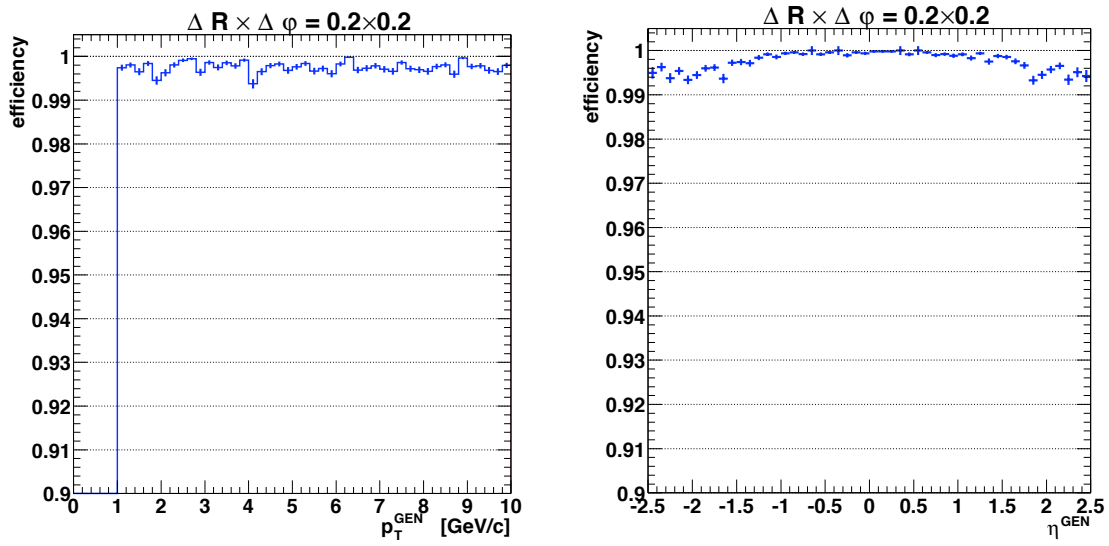


Figure 6.8: The algorithmic hit-pair finding efficiency as a function of (left) generated transverse momentum and (right) pseudorapidity. The flat plateau in (left) results from accommodation of multiple scattering effects. The plot is made for single muon events (region size is $\Delta\eta \times \Delta\varphi = 0.2 \times 0.2$). The geometrical and detector inefficiencies lower the algorithmic efficiencies by 1–2% (luminosity dependent).

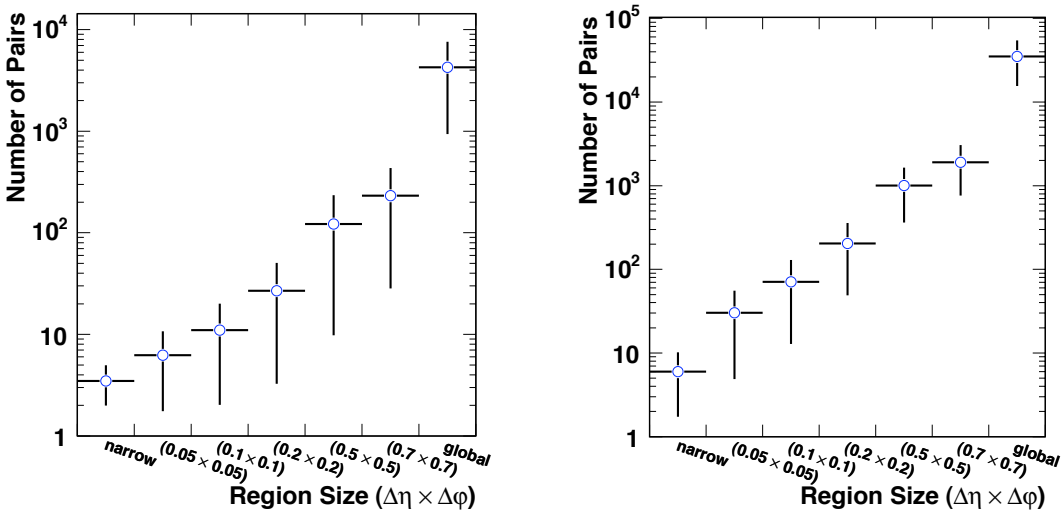


Figure 6.9: The reconstructed number of hit pairs as the function of region size for $h \to ee\mu\mu$ events in the case of (left) low and (right) high luminosity.

6.4.3 Pattern recognition

6.4.3.1 Description of the algorithm

The pattern recognition is based on a combinatorial Kalman filter method. The filter proceeds iteratively from the seed layer, starting from a coarse estimate of the track parameters provided by the seed, and including the information of the successive detection layers one by one. On each layer, i.e., with every new measurement, the track parameters are known

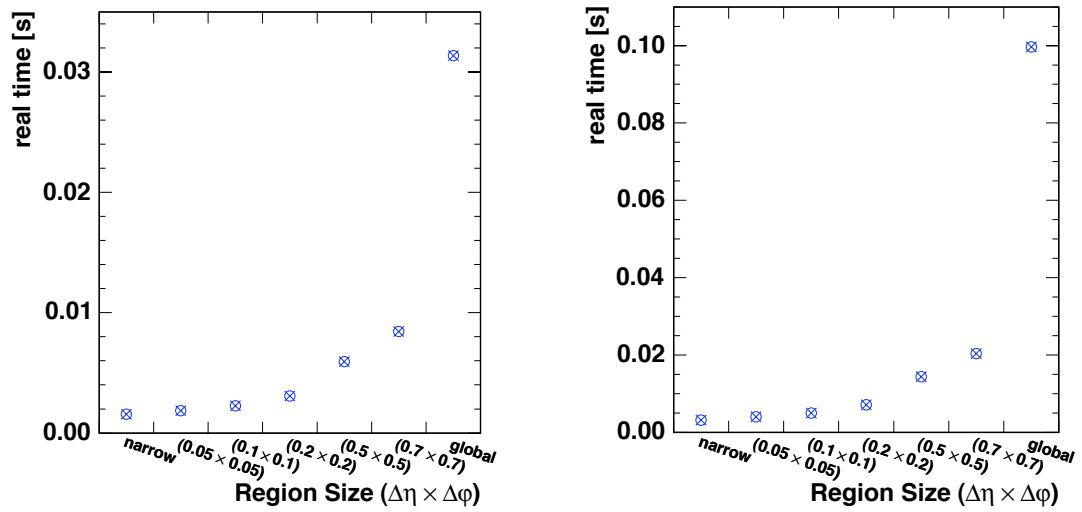


Figure 6.10: The real CPU time (2.8 GHz Xeon) used to reconstruct hit pairs as a function of region size. The sample is $h \rightarrow ee\mu\mu$ events with (left) low and (right) high luminosity pile-up.

with a better precision, up to the last point, where they include the full tracker information.

First, a dedicated *navigation* component determines which layers are compatible with the initial seed trajectory. The trajectory is then extrapolated to these layers according to the equations of motion of a charged particle in a magnetic field, accounting for multiple scattering and energy loss in the traversed material.

Since several hits on the new layer may be compatible with the predicted trajectory, several new trajectory candidates are created, 1 per hit. In addition, 1 additional trajectory candidate is created, in which no measured hit is used, to account for the possibility that the track did not leave any hit on that particular layer. This fake hit is called an "invalid hit."

Each trajectory is then "updated" with the corresponding hit according to the Kalman filter formalism. This update can be seen as a combination of the predicted trajectory state and the hit in a weighted mean, as the weights attributed to the measurement and to the predicted trajectory depend on their respective uncertainties.

All resulting trajectory candidates are then grown in turn to the next compatible layer(s), and the procedure is repeated until either the outermost layer of the tracker is reached or a "stopping condition" is satisfied. In order not to bias the result, all trajectory candidates are grown in parallel. To avoid an exponential increase of the number of trajectory candidates, the total number of candidates is truncated at each layer. To limit the number of combinations, and hence to avoid an exponential increase thereof, only a limited number of these are retained at each step, based on their normalized χ^2 and number of valid and invalid hits.

6.4.3.2 Tunable parameters, regional and partial tracking

This algorithm is configurable through several parameters. Depending on their values, it can provide either a high track finding efficiency, as needed offline, or very fast CPU performance suitable for use in the HLT.

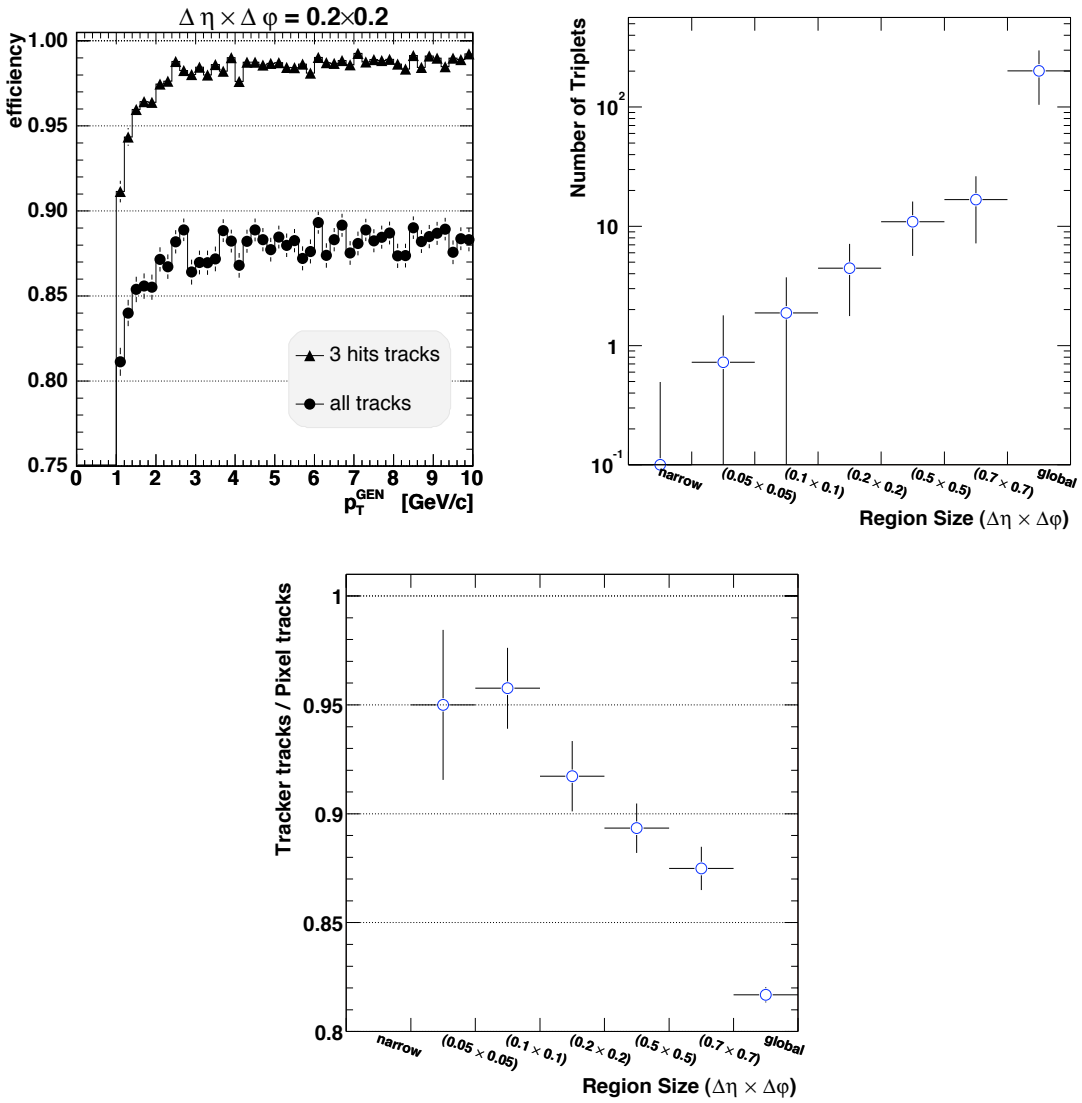


Figure 6.11: The performance of triplet finding: (top left) algorithmic and global efficiency, (top right) reconstructed number of triplets, and (bottom) purity. The efficiency (region size is $\Delta\eta \times \Delta\varphi = 0.2 \times 0.2$) is measured with single muon events, thus quoted efficiency does not take into account a few percent loss due to luminosity dependent readout losses. The purity is defined as the number of reconstructed full tracker tracks to triplets used for seeding. The number of triplets and purity plots are made with muons in $b\bar{b}$ jets in high luminosity environment. In this case the reconstruction region is centered on the jet direction.

These main parameters are (default values, used in the production of DSTs and for the performance plots (unless stated otherwise) are given in brackets) the following:

- the maximum number of candidates that are propagated at each step (5)
- the inclusion of an invalid hit in the list of compatible hits, when the latter is not empty (always include invalid hit)
- the maximum χ^2 of the hits considered compatible with the predicted track state (30)

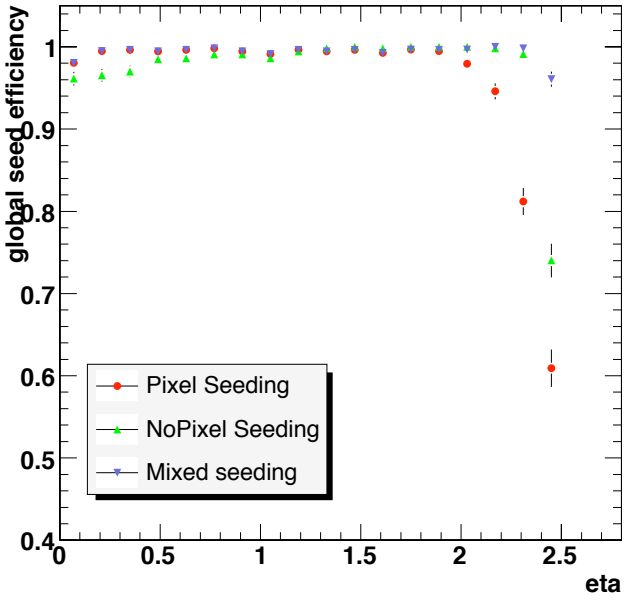


Figure 6.12: The global hit-pair finding efficiency as a function of pseudorapidity for single muons with transverse momenta of 1 GeV/c, shown for pixel-less seeding, pixel seeding, and mixed seeding.

- the maximum number of invalid hits, i.e., crossings of sensitive detectors without a measurement (1)
- the maximum number of consecutive invalid hits (1)
- the minimum transverse momentum (0.9)
- the minimum number of hits per track (5)

In addition, it is possible to specify an arbitrary stopping condition, in which case the pattern recognition is interrupted before the end of the tracker is reached (*partial track reconstruction*). Such a condition is typically used in the HLT, since the required accuracy on track parameters is often reached after 5 or 6 hits, and the continuation of the pattern recognition to 12–13 hits would be a waste of CPU time.

If the track reconstruction is constrained to a region, the constraint affects mostly the seed generation phase. The only constraint that can be applied effectively at the pattern recognition stage is the transverse momentum cut.

6.4.3.3 Combinatorial behaviour of the pattern recognition algorithm

When a trajectory is propagated to a given layer, the uncertainty of the predicted state (Fig. 6.13) has a direct effect on pattern recognition, since it determines the compatibility between the trajectory and nearby hits. The number of compatible hits found on a layer determines the increase of the number of trajectories to be propagated, as the initial trajectories are multiplied by the number of hits found.

In the barrel, a large fraction of the seeds are composed of hits in the first 2 pixel layers, and the seed trajectories are first propagated from the second to the third layer. With the limited information available at that stage, the parameters of the trajectories are still poorly defined, and the uncertainties of the predicted states are quite large. Even though the uncertainty is

large, due to the very fine granularity and the low occupancy of the pixel detectors, only 2 hits are usually found to be compatible, the correct hit and the invalid hit. As a consequence, the contamination of the trajectories formed when the found hits are included is very low.

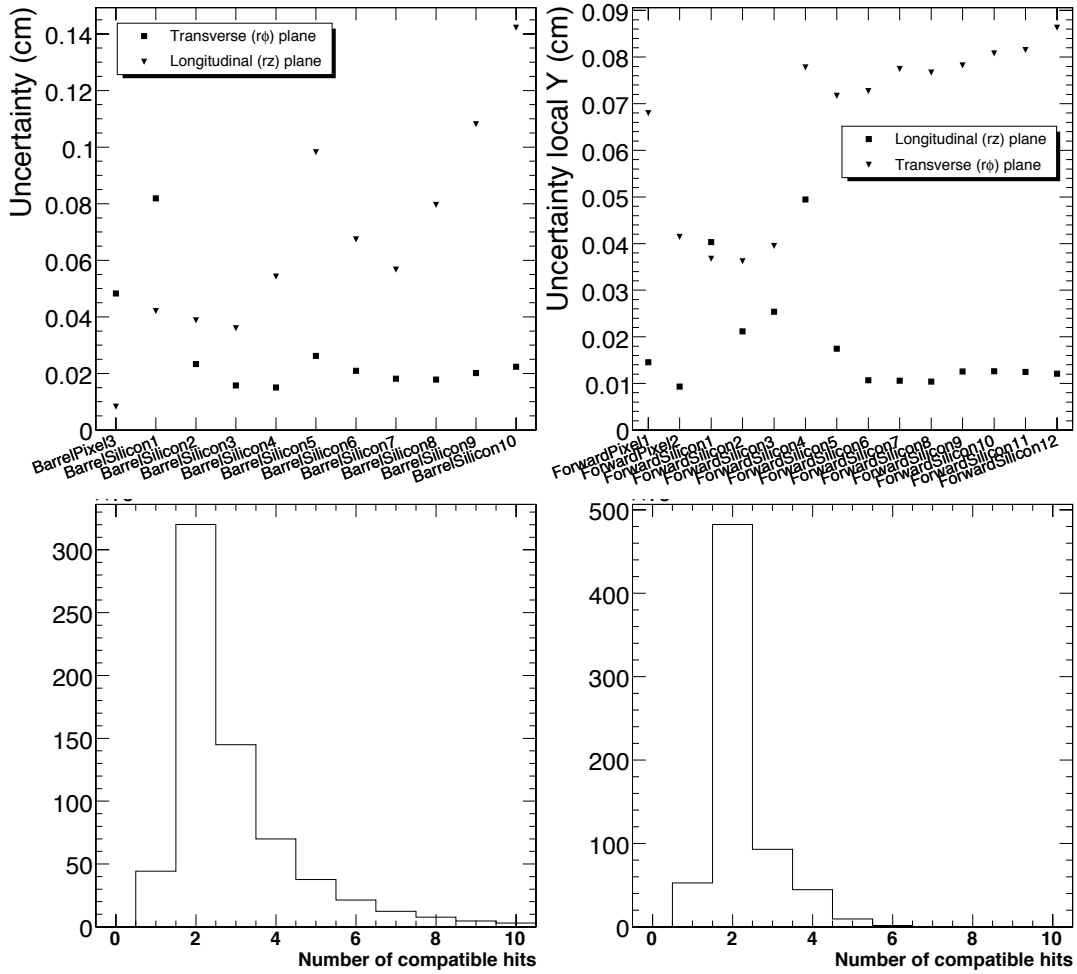


Figure 6.13: Upper figures: Mean uncertainties of the predicted state in the barrel (left) and forward tracker (right) in the transverse (r-ϕ) and longitudinal (r-z) planes. Lower figures: number of compatible hits (including invalid hit(s)) found per trajectory, when leaving the pixel barrel layer 3 (left) and TIB layer 3 (right). Simulated b-jets with transverse momenta between 120 and 170 GeV/c are used, with low luminosity pile-up included.

When these trajectories are propagated to the first TIB layer, the uncertainty on the predicted state increases and the distribution broadens due to the large gap between the 2 layers (approximately 13 cm) and the small lever arm of the initial trajectories (approximately 6 cm). As a strip detector with approximately 10-cm-long strips, and thus with a higher occupancy, is now reached, the number of compatible hits on that layer is larger (Fig. 6.13, bottom-left), and the contamination by spurious hits markedly higher. Nevertheless, when the hit is included, the trajectories are now much better defined, with a larger lever arm. When these are then propagated to the next layer, TIB layer 2, the uncertainties of the predicted states are again reduced, which in turn reduces the number of spurious hits found on that layer and the contamination decreases to a negligible level. Indeed, many of the trajectories containing spurious hits have now been discarded, either because they were not retained for a further

propagation, or because no compatible hits were found when propagated. From there on, for the subsequent propagations through the rest of the barrel, the trajectories are well defined and the contamination stays at the same low level (Fig. 6.13, bottom-right).

While navigation in the barrel is quite easy, it is more complex in the transition region. Due to the difficult geometry, the navigation can point to several layers. For example, for high-η tracks leaving the forward pixel disk 2, all 3 inner disks and the first 3 endcap disks could be compatible with the initial trajectory. When these layers are queried for compatible hits, each of these returns at least the invalid hit. As the propagation distance to some of these layers can be quite large, the uncertainties are comparatively large, and the probability of finding a spurious hit increases. There is as such a much higher number of compatible hits and a higher contamination from spurious hits. Once the trajectory is in the endcap, the navigation is easy again: many of the trajectories containing spurious hits are quickly discarded and the contamination is reduced.

6.4.4 Ambiguity resolution

Ambiguities in track finding arise because a given track may be reconstructed starting from different seeds, or because a given seed may result in more than 1 trajectory candidate. These ambiguities, or mutually exclusive track candidates, must be resolved in order to avoid double counting of tracks.

The ambiguity resolution is based on the fraction of hits that are shared between 2 trajectories. For any pair of track candidates, this fraction is defined in the following way:

$$f_{\text{shared}} = \frac{N_{\text{shared}}^{\text{hits}}}{\min(N_1^{\text{hits}}, N_2^{\text{hits}})},$$

where N_1^{hits} (N_2^{hits}) is the number of hits in the first (second) track candidate. If this fraction exceeds a value of 0.5, the track with the least number of hits is discarded, or, if both tracks have the same number of hits, the track with the highest χ^2 value is discarded.

The ambiguity resolution is applied twice: the first time on all track candidates resulting from a single seed, and the second time on the complete set of track candidates from all seeds.

6.4.5 Track fitting and smoothing

For each trajectory, the building stage results in a collection of hits and in an estimate of the track parameters. However, the full information is only available at the last hit of the trajectory and the estimate can be biased by constraints applied during the seeding stage. Therefore the trajectory is refitted using a least-squares approach, implemented as a combination of a standard Kalman filter and smoother.

The Kalman filter is initialized at the location of the innermost hit with an estimate obtained during seeding. The corresponding covariance matrix is scaled by a large factor in order to avoid any bias. The fit then proceeds in an iterative way through the list of hits. For each valid hit the position estimate is re-evaluated using the current values of the track parameters: information about the angle of incidence increases the precision of the measurement especially in the pixel modules. The track parameters and their covariance matrix are updated with the measurement and the trajectory is propagated to the surface associated with

the next hit. The track parameters and their covariance matrix are modified according to the estimates for energy loss and multiple scattering at the target surface and the sequence is repeated until the last hit is included.

This first filter is complemented with a smoothing stage: a second filter is initialized with the result of the first one—except for the covariance matrix, which is scaled with a large factor—and run backward toward the beam line. At each hit the "updated" parameters of this second filter, which contain all information from the outermost hit up to and including the current hit, are combined with the "predicted" parameters of the first filter, i.e., the information from the innermost hit outward, but excluding the current hit.

This filtering and smoothing procedure yields optimal estimates of the parameters at the surface associated with each hit and, specifically, at the first and the last hit of the trajectory. Estimates on other surfaces, e.g., at the impact point, are then derived by extrapolation from the closest hit.

6.4.6 Track reconstruction performance

6.4.6.1 Efficiency and fake rate

The efficiency of reconstructing single tracks with the combinatorial Kalman filter has been estimated using samples of muons and pions with transverse momenta of 1, 10, and 100 GeV/ c. For this analysis, tracks are reconstructed using seeds in the pixel detector and default settings for the pattern recognition, ambiguity resolution, and fitting stages. Reconstructed tracks are required to have a minimum of 8 hits, with a hit missing in at most 1 layer, and $p_{\mathrm{T}} > 0.8\ \mathrm{GeV}/c$. A track is deemed to be successfully reconstructed if it shares more than 50% of the hits with a simulated track.

To measure the track reconstruction efficiency, 2 definitions of efficiency are used, as in the case of the seed generator. The algorithmic efficiency is the efficiency of reconstructing correctly tracks with $p_{\mathrm{T}} > 0.9\ \mathrm{GeV}/c$ which have simulated hits in at least 8 layers in the tracker, of which at least 2 are in the pixel detector, and which are originating in a region compatible with the seeding requirements. It measures directly the performance of the track reconstruction algorithm. It is essentially the efficiency of the pattern recognition stage, the other stages being fully efficient. The global efficiency is the reconstruction efficiency for *all* tracks with $p_{\mathrm{T}} > 0.9\ \mathrm{GeV}/c$ and with a production vertex inside the beam pipe. In addition to the efficiency of the algorithm, it includes the acceptance, hit efficiency, and any other factor influencing reconstruction. It mainly differs from the algorithmic efficiency in the forward region, with the loss of coverage by the disks, especially in the pixel system. The fake rate is defined as the fraction of reconstructed tracks that could not be associated with a simulated track of $p_{\mathrm{T}} > 0.7\ \mathrm{GeV}/c$.

The efficiency for single particles is shown in Figure 6.14. As the algorithmic efficiency indicates, the pattern recognition is fully efficient for pseudorapidities up to $|\eta| = 2.4$. For the global efficiency (Fig. 6.14), the drop of efficiency in the region $|\eta| < 0.1$ is due to the gaps between the sensors in the ladders of the pixel detector at $z = 0$. The alignment of the gaps in the 3 layers causes some tracks not to have the 2 required pixel hits. At high η, the drop of efficiency is mainly due to the lack of coverage by the 2 pairs of forward pixel disks. For hadrons, the global efficiency is obviously lower than for muons because hadrons interact with the material present in the tracker.

The global and algorithmic efficiencies and the fake rate for tracks in b jets with transverse

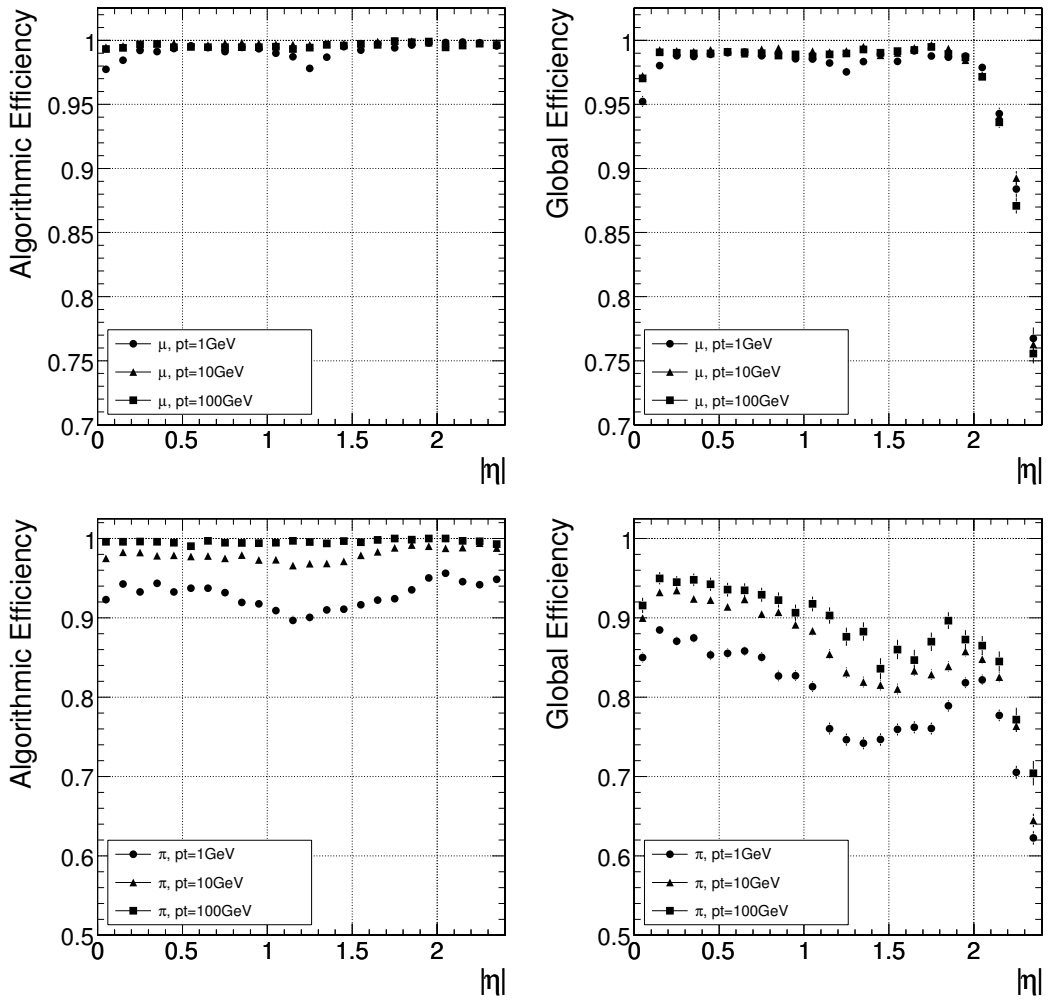


Figure 6.14: Algorithmic (left) and global track reconstruction efficiency (right) for muons (top) and pions (bottom) of transverse momenta of 1, 10 and 100 GeV/c.

momenta in the range 120–170 GeV/c and including low luminosity pile-up are shown in Figure 6.15. The track selection used for this analysis is a very loose one, and the efficiency and fake rate can be tuned by applying additional quality criteria. The most important parameters that are available for such a selection are the number of hits used in the track fit, the number of invalid hits, and the χ^2. As an example, the change of global efficiency and fake rate as a function of the cut on the normalized χ^2 are also shown in Figure 6.15.

6.4.6.2 Resolution

Five parameters are chosen to describe a track: d_0, z_0, ϕ, $\cot\theta$, and the transverse momentum p_T. The track parameters are defined at the point of closest approach of the track to the beam axis (called the impact point); d_0 and z_0 hence measure the coordinate of the impact point in the transverse and longitudinal plane ($d_0 = y_0 \cos\phi - x_0 \sin\phi$, where x_0 and y_0 are the transverse coordinates of the impact point). The azimuthal angle of the momentum vector of the track, ϕ, is taken at the impact point, and θ is the polar angle. Figure 6.16 shows the resolution of the 5 track parameters for samples of single muons with p_T of 1, 10, and 100 GeV/c.

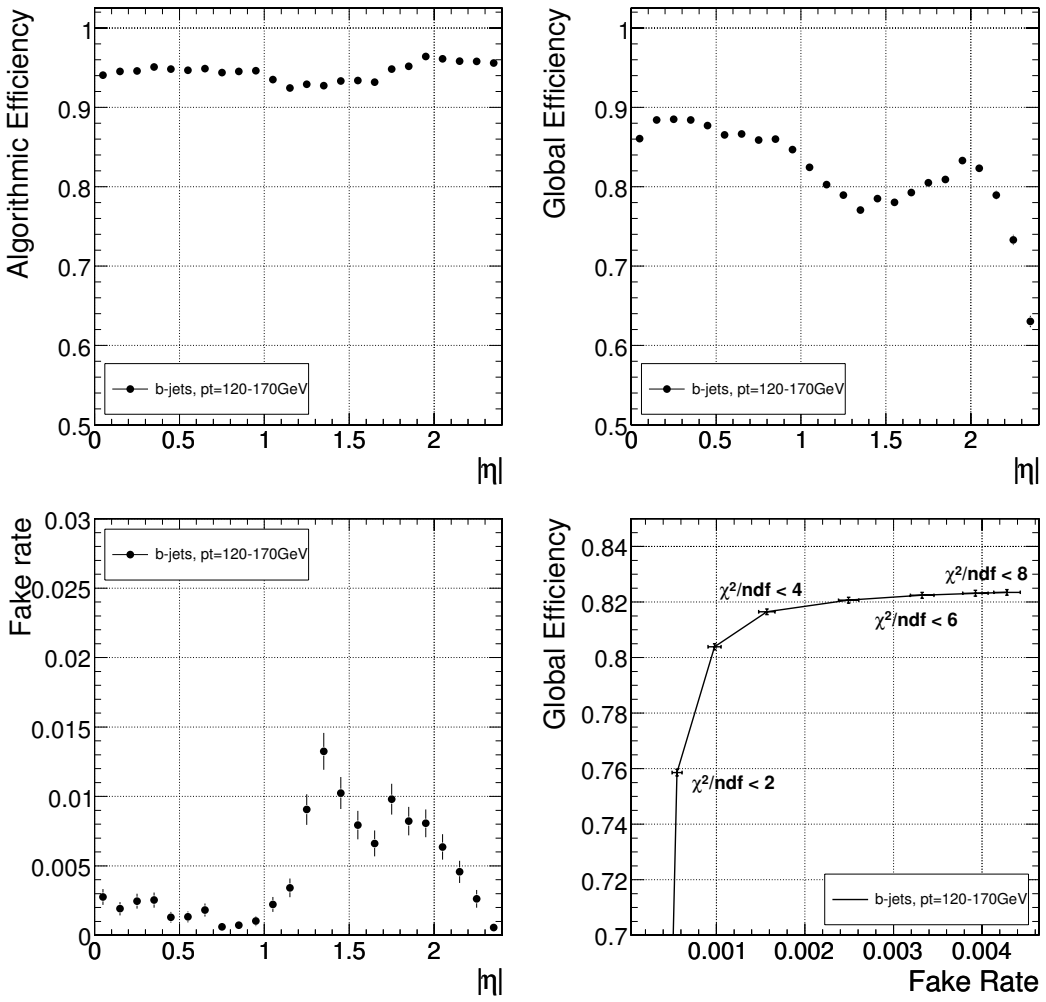


Figure 6.15: Global track reconstruction efficiency for b jets with transverse momenta between 120 and 170 GeV/c, including low luminosity pile-up. Upper row: algorithmic and global efficiency. Lower row: fake rate and evolution of the average efficiency and fake rates for different cuts on the normalized χ^2 of the track.

The resolution of the transverse momentum is shown in Figure 6.16a. At high momentum (100 GeV/c), the resolution is around 1–2% up to a pseudorapidity of $|\eta| = 1.6$; for higher values of $|\eta|$ the lever arm of the measurement is reduced. The degradation around $|\eta| = 1.0$ is due to the gap between the barrel and the endcap disks and the degradation beyond $|\eta| = 1.2$ is due to the lower hit resolution of the last hits of the track measured in the TEC ring 7 with respect to the hit resolution in the TOB layers 5 and 6. At a transverse momentum of 100 GeV/c, the material in the tracker accounts for between 20% and 30% of the transverse momentum resolution; at lower momenta, the resolution is dominated by multiple scattering and its distribution reflects the amount of material traversed by the track. The resolutions of the transverse and longitudinal impact parameters d_0 and z_0 are shown in Figures 6.16d and 6.16e. At high momentum, the d_0 resolution is fairly constant and is dominated by the hit resolution of the first hit in the pixel detector. At lower momenta, the d_0 resolution is progressively degraded by multiple scattering, until the latter becomes dominant. The z_0 resolution of high momentum tracks is also dominated by the hit resolution of the first

pixel hit, with multiple scattering dominating at low momentum. The improvement of the z_0 resolution up to a pseudorapidity of $|\eta| = 0.5$ can be attributed to the fact that in the barrel, as the angle with which the tracks cross the pixel layers increases the clusters become wider, improving the pixel-hit resolution. The average reduced χ^2 is close to 1 as shown in Figure 6.16f.

6.4.7 Fast track reconstruction with pixel detector only

As explained in Section 6.4.2, triplets of hits found in 3 different layers of the pixel detector (pixel triplets) may be used to create a collection of seeds for offline track reconstruction. In certain cases where speed is an important requirement and the highest possible efficiencies are not a key issue—like HLT b-jet and τ tagging, HLT muon isolation—the hit triplets can be converted into tracks without further propagation into the tracker.

The pattern recognition performance for pixel triplets is shown in Figure 6.11. Evidently, requiring 3 hits out of 3—together with the beam or vertex constraint—is a sufficient condition to yield good pattern recognition results. Depending on the definition of the region and luminosity, approximately 90% of the triplets corresponds to real tracks. Of course, the requirement that 3 hits are found in 3 layers lays a constraint on the efficiency. Whereas the algorithmic efficiency—the probability of finding the correct triplet if at least 3 hits have been found—is close to 100%, the global efficiency is 80–90% (luminosity dependent), due to geometrical coverage, pixel inefficiencies, and readout losses.

The determination of the track parameters and errors is optimized for speed. Nevertheless, tracks created from hit triplets reach a acceptable resolution. The main limitation comes from the small lever arm of the pixel detector. In Figure 6.17 the transverse impact parameter resolution is shown versus transverse momentum. Clearly, there is a rapid degradation of the transverse momentum resolution toward large transverse momenta. The uncertainty of the transverse momentum determination affects the extrapolation to the beam line. The transverse impact parameter resolution versus transverse momentum and for 2 pseudorapidity bins is shown in the right plot of Figure 6.17.

The most important characteristic of the pixel-only track finder is its timing performance. For the regional cases the timing does not exceed 20 ms/event (high luminosity, 2.4 GHz Xeon) and reaches about 110 ms/event for global reconstruction. Given the favourable balance between tracking performance and CPU requirements, tracks based on pixel hit triplets form a valuable asset in the first stages of the High-Level Trigger.

6.4.8 Track reconstruction in heavy-ion collisions

The extremely high particle density in heavy-ion collisions of up to 3000 charged particles per unit rapidity in central events leads to a very high detector occupancy in the silicon-strip detectors.

6.4.8.1 Modifications to the algorithm

To make the track reconstruction algorithm robust against the combinatorial problem resulting from the high particle density the default track reconstruction procedure is modified as follows:

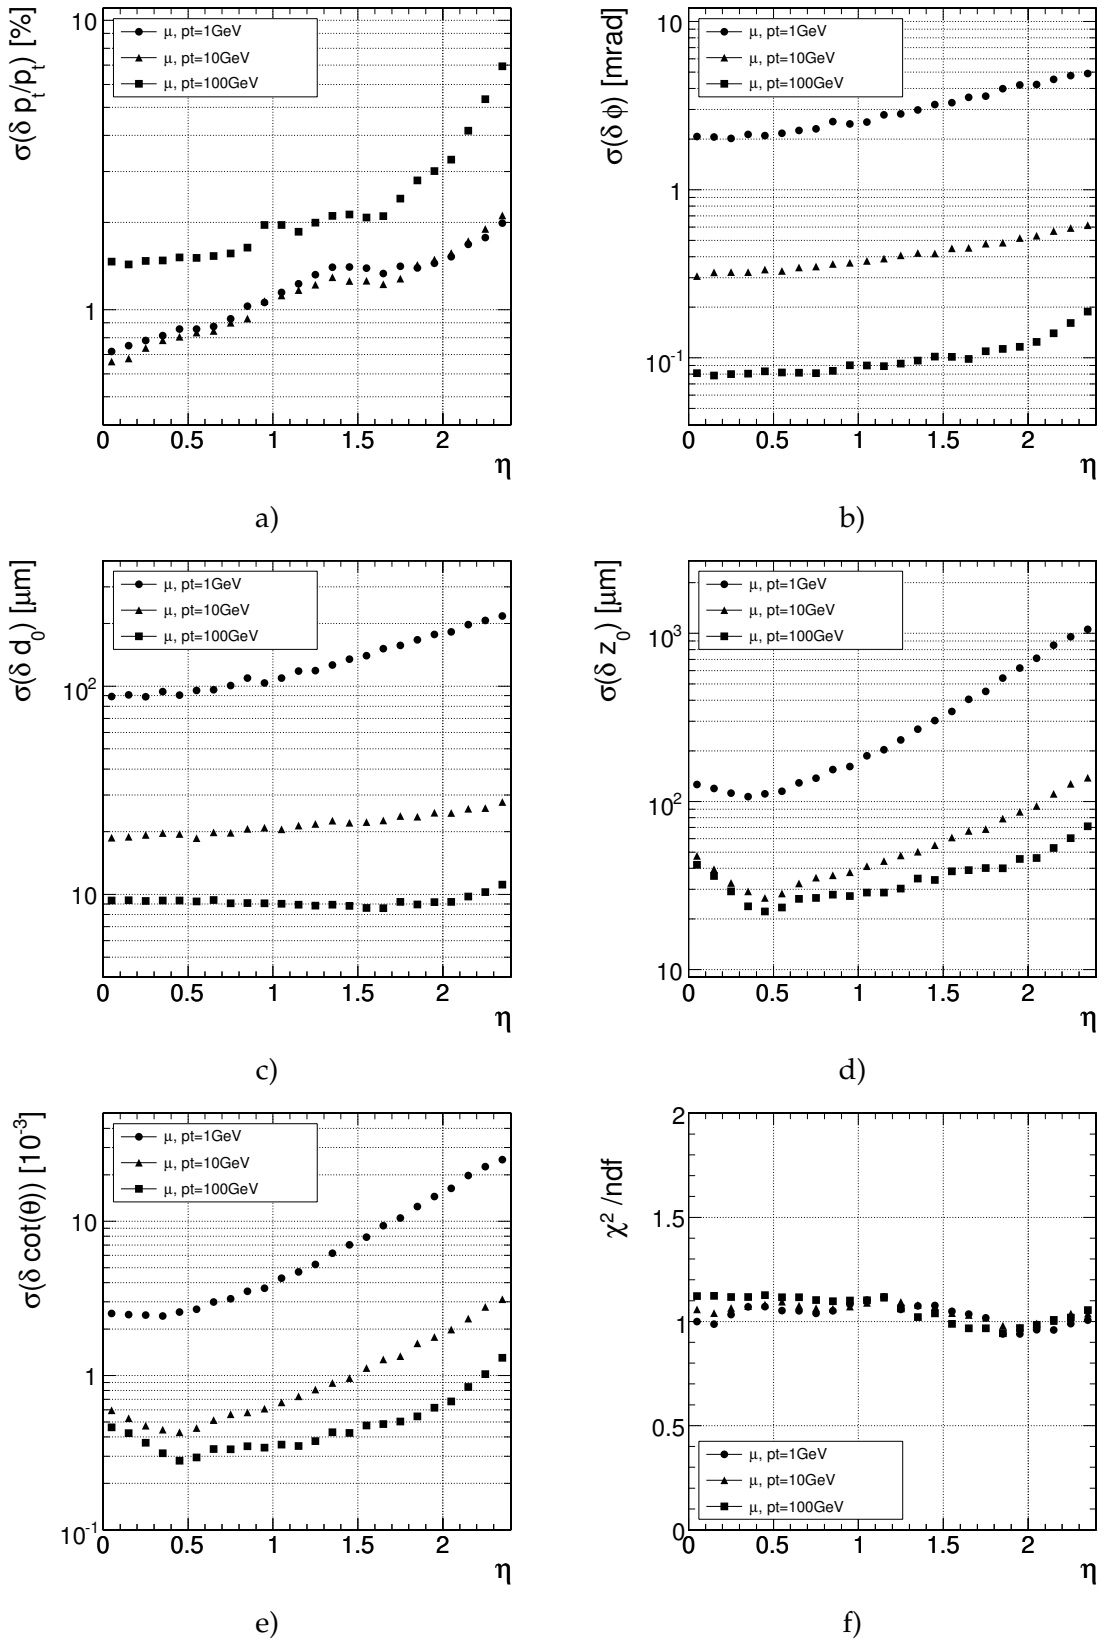


Figure 6.16: Resolution of the 5 track parameters for single muons with transverse momenta of 1, 10 and 100 GeV/c: a) transverse momentum, b) ϕ, c) transverse, d) longitudinal impact parameter, e) $\cot\theta$, and f) reduced χ^2.

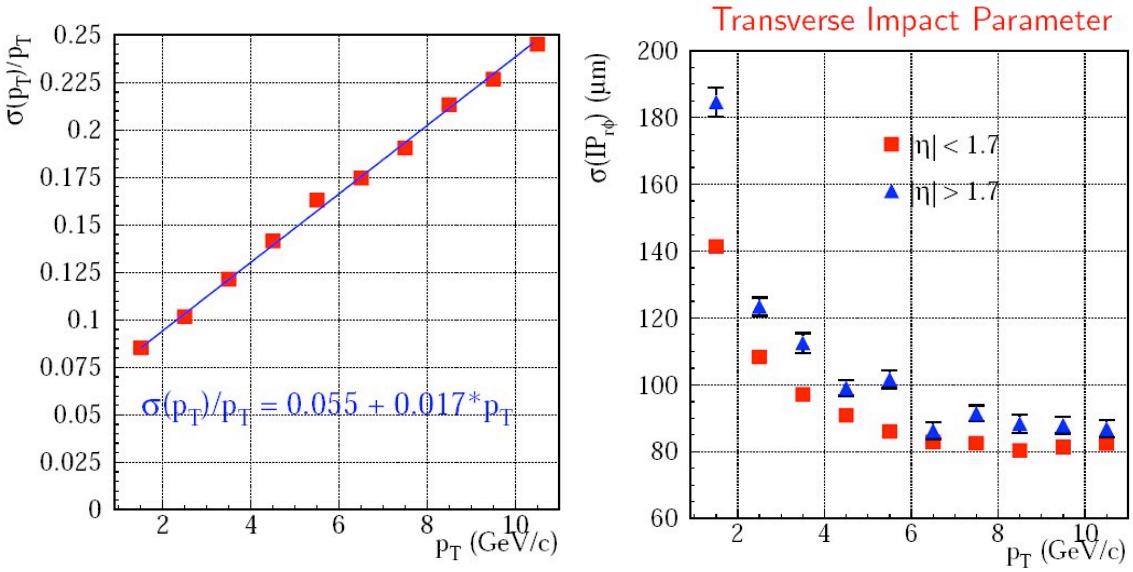


Figure 6.17: Transverse momentum (left) and transverse impact parameter (right) resolution for single muon tracks, created directly from hit triplets.

- The seeding of the track reconstruction relies on 3 hit combinations in the pixel detectors to achieve more precise initial estimates of the track parameters.
- Merged hits are recognized in the silicon-strip detectors by comparing the found cluster width with the width expected from the angle of the trajectory to the detector surface. An error proportional to the cluster width is assigned to merged hits.
- In the final smoothing step hits in the double silicon-strip layers are split and treated as separate hits.

The number of fake tracks in the data sample can be controlled by imposing constraints on the quality of the reconstructed tracks. The reconstruction quality is addressed by the number of reconstructed hits on the track, the χ^2-probability of the track fit, and the compatibility of the track with the event vertex. The corresponding distributions are shown in Figure 6.18.

6.4.8.2 Performance

The performance of the track reconstruction algorithms in heavy-ion collisions was evaluated using a data sample generated using the HYDJET event generator [183] with parameters tuned to yield a charged particle density at midrapidity dN/dy of about 3000.

With the modifications mentioned above, a high algorithmic reconstruction efficiency can be achieved in central heavy-ion collisions while retaining a very low fake rate. Figure 6.19 shows the track reconstruction efficiency and fake rate as a function of transverse momentum in the barrel region of the tracker for 2 sets of quality cuts imposed on the reconstructed tracks.

The momentum and impact parameter resolution achieved in heavy-ion collisions (Fig. 6.20) is comparable to the resolution in low-occupancy p-p events.

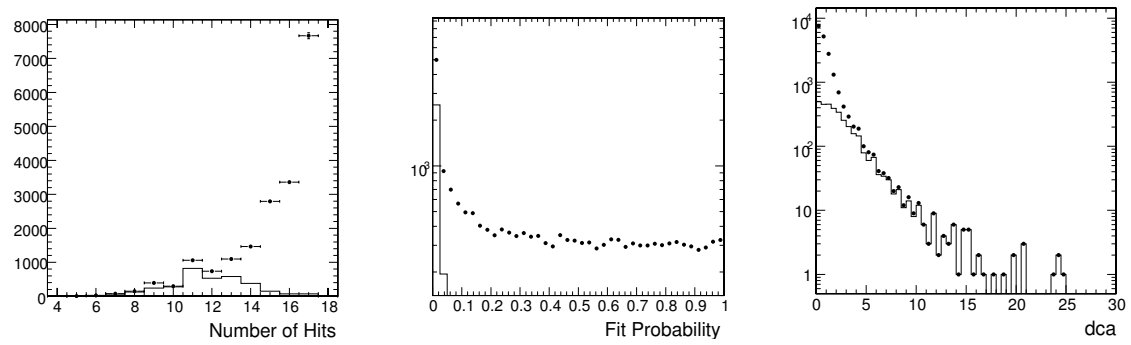


Figure 6.18: Track quality measures for all reconstructed tracks (dots) and for fake tracks contained in the sample (histogram). Left panel: Number of hits per reconstructed track. Hits on stereo layers are counted as 2 separate hits. Central panel: χ^2 probability distribution for reconstructed tracks. Right panel: Distance of closest approach to the primary vertex normalized by the error.

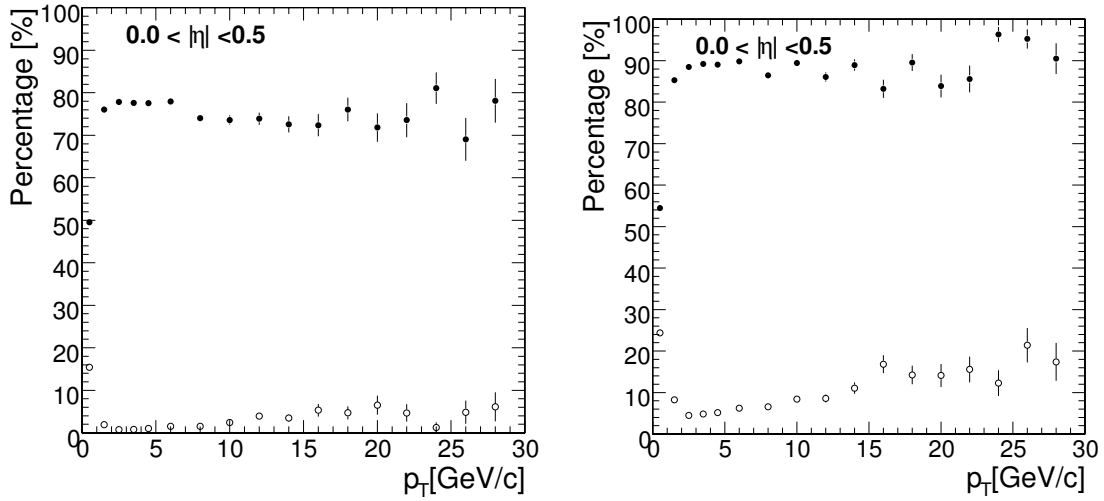


Figure 6.19: Reconstruction efficiency (filled circles) and fake rate (open circles) as a function of transverse momentum in the barrel region of the tracker for central Pb-Pb collisions with a charged particle density of $dN/dy \approx 3000$. Left panel: Track quality cuts optimized for low fake rate (number of hits > 12, fit probability > 0.01, and dca < 3). Right panel: Track quality cuts optimized for high efficiency (number of hits > 12).

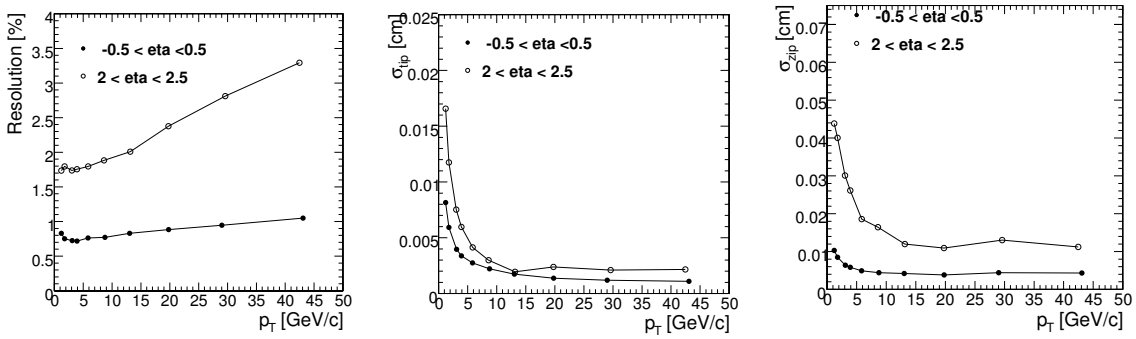


Figure 6.20: Track parameter resolution achieved in heavy-ion events in the barrel region (filled circles) and in the forward region (open circles). Left panel: Transverse momentum resolution. Central panel: Transverse impact parameter resolution. Right panel: Longitudinal impact parameter resolution.

6.4.8.3 High occupancy effects in detector hardware and readout

To demonstrate the feasibility of reconstructing charged particles in the high occupancy environment of heavy-ion collisions, a detailed study was performed using a full detector simulation and reconstruction while checking at each stage of the readout chain of the various detector components that the readout electronics and buffers can cope with the high hit density. All detector components were found to be fully functional, however 2 effects were identified that could lead to a potential loss of reconstruction efficiency.

The first effect is due to the Common Mode Noise (CMN) subtraction in the silicon-strip tracker. The tracker data are likely to be subject to common-mode variations, whose magnitude is not yet known. In the present hardware design, optimized for p-p collisions, the FED estimates the common-mode offset event-by-event for each APV. The offset is estimated by calculating the median ADC value of the data on the 128 strips read by the APV. At high occupancy this simple algorithm introduces a false common mode offset that is dependent on the detector occupancy, which in turn leads to an inefficiency in the hit finding. This inefficiency can be fully recovered with a more sophisticated common-mode offset estimation. Currently more sophisticated and robust algorithms for the CMN subtraction that could be implemented in the FED firmware are under evaluation. If achieving sufficiently robust performance requires a CPU-intensive algorithm that cannot be implemented in the FED firmware, the low interaction rate for Pb-Pb collisions of 8 kHz allows the detector to be read in non-zero-suppressed mode. In that case, the CMN subtraction could be done in the HLT farm, where more CPU power would be available.

The second effect influencing the reconstruction efficiency are highly ionizing particles. The high particle density in heavy-ion collisions also results in a high probability of hadronic interactions with the detector material. This results in a large charge deposit in the active detector volume. High amounts of charge will saturate the dynamic range of the readout electronics resulting in a loss of reconstructed hits. This effect is commonly referred to as the HIP effect and has been extensively studied for p-p interactions [184]. The HIP effect in the silicon tracker can be simulated in the ORCA reconstruction package, based on a parameterization of test-beam results. The inefficiency due to this effect is estimated by comparing the same data sample reconstructed with and without simulating this effect. An overall loss of 3–5% reconstruction efficiency is observed for central Pb-Pb collisions.

Table 6.3: The mean number of simulated tracks with $p_{\mathrm{T}} > 0.7$ GeV/c per vertex, the mean number of reconstructed tracks associated to the simulated tracks, and the average transverse momentum of the reconstructed tracks.

Sample	# SimTracks	# RecTracks	Mean p_{T} (GeV/c)
$B_s \to J/\psi\ \phi$ (Secondary Vertex)	4	4	3.5
$B_s \to J/\psi\ \phi$ (Primary Vertex)	28.9	12.4	1.6
$H \to \gamma\gamma$, gg fusion	52.1	23.2	2.0
Drell-Yan $\mu^+\mu^-$	50.6	22.7	7.2
$t\bar{t}H$, $H(m_H = 120\ \mathrm{GeV}/c^2) \to b\bar{b}$	79.3	44.3	4.0

6.5 Vertex reconstruction

Vertex reconstruction usually involves 2 steps, vertex finding and vertex fitting. Vertex finding involves grouping tracks into vertex candidates. The vertex-finding algorithms can be very different depending on the physics case (primary or secondary vertex finding, reconstruction of exclusive decays, etc.). Vertex fitting involves determining the best estimate of the vertex parameters (position, covariance matrix, and track parameters constrained by the vertex position and their covariances) for a given set of tracks, as well as indicators of the fit quality (total χ^2, number of degrees of freedom, or track weights).

Since some vertex-finding algorithms use vertex-fitting algorithms to associate tracks with a given vertex candidate, vertex-fitting algorithms are described first. On-line and off-line primary vertex reconstruction are described next. Then off-line secondary vertex finding algorithms are discussed. An evaluation of the impact of tracker misalignment on vertex reconstruction is given. Finally, the strategy to measure the beam spot position at the High-Level Trigger is described.

Vertex reconstruction is studied in several different physics channels that are representative of the event kinematics expected at the LHC. To study primary vertices, the following channels are chosen: $B_s \to J/\psi\ \phi$, $H \to \gamma\gamma$ (produced through gg and vector-boson fusion), $t\bar{t}H$ with $H(m_H = 120\ \mathrm{GeV}/c^2) \to b\bar{b}$, Drell-Yan $\mu^+\mu^-$ production with $\sqrt{s} = 115\ \mathrm{GeV}/c$, u dijet events with $50\ \mathrm{GeV} < E_{\mathrm{T}}^{\mathrm{jet}} < 100\ \mathrm{GeV}$, b-jet events with $30\ \mathrm{GeV} < E_{\mathrm{T}}^{\mathrm{jet}} < 50\ \mathrm{GeV}$ where at least 1 b-jet is within the tracker acceptance, and $t\bar{t}$ events.

The average track multiplicities and track transverse momenta for some of the samples studied are summarized in Table 6.3. The difference between the simulated and reconstructed track multiplicities is due to the tracker acceptance ($|\eta| < 2.5$), the p_{T} requirement applied during track seed generation (0.9 GeV/c by default), and the 15% track reconstruction inefficiency for hadrons arising from nuclear interactions in the tracker. All samples are simulated with low luminosity pile-up ($\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$).

The secondary vertex of the decay $B_s \to J/\psi\ \phi$ is also studied, since a precise determination of the decay position is essential for both the selection and for the measurement of the parameters of interest. In this case, exactly 4 tracks with relatively low transverse momentum are fitted to the vertex. These tracks are usually selected using only kinematic requirements, and do not rely on a vertex-finding technique such as those used for the primary vertex. A cut on the transverse momentum of the muons is made during Monte Carlo generation. Therefore, only decays for which both muons have a transverse momentum above 2 GeV are retained.

6.5.1 Vertex fitting

6.5.1.1 Algorithms

Several vertex-fitting algorithms have been implemented and studied, differing mainly in their sensitivity to outlying tracks, that is, either mismeasured tracks or tracks from another vertex. The implemented algorithms can be divided into linear (least-squares) and nonlinear algorithms, the latter being more robust with respect to outlying tracks. In least-squares algorithms all tracks are used with equal (unit) weight. Nonlinear algorithms are able to down-weight or discard tracks, making them less sensitive to outliers.

The most often used algorithm for vertex fitting is the well-known Kalman filter (KVF) [185]. It is mathematically equivalent to a global least-squares minimization, which is the optimal estimator when the measurements are Gaussian and the fitted parameters depend linearly on those measurements. The filter can also compute an improved estimate of the track momenta, using both the vertex and the resulting track-to-track covariance matrices as constraints [186].

A robust filter introduced for the first time in CMS is the Adaptive Vertex Fitter (AVF) [187, 188, 189]. It is an iterative re-weighted fit which down-weights tracks according to their reduced (χ^2) distance from the vertex. The weights vary from one iteration to the next, until the fit converges. This algorithm has the advantage that the weights can be fractional (soft assignment). No prior estimate of the track weights or of the fraction of outliers is needed.

The Trimmed Kalman Fitter (TKF) is the conventional robust version of the Kalman vertex fitter, where tracks incompatible with the vertex are removed one by one from the vertex, starting with the least compatible track [186]. It is a hard-assignment, iterative fit.

6.5.1.2 Performance

The performance of the filters has been assessed using the samples described above. Only reconstructed tracks matched to simulated tracks that were produced in the relevant decay are used in the fit, without any selection cuts. As such, only mismeasured tracks can potentially be included in the fit, with contamination from tracks from another vertex having been eliminated.

The main statistical properties of the fits are summarized in Table 6.4. The *average weight* is the average of the sum of the track weights normalized by the number of initial tracks associated to each vertex. For the Kalman filter, all tracks are used with unit weight, while for the Trimmed Kalman Fitter, the weight of a discarded track is zero, and the average weight is equivalent to the average fraction of tracks remaining. The failure rate is the fraction of vertex fits where either the fit did not converge, or the sum of the weights of all the tracks was smaller than 1.

For vertices fitted with the Kalman filter, a large number of fits have a low χ^2-probability, typically below 0.01. These have been shown to be due to non-Gaussian tails in the distribution of the track parameter errors [190]. In the case of the decay $B_s \to J/\psi\ \phi$, approximately 10% of the fits have a χ^2-probability below 0.01. For $t\bar{t}$ vertices, however, this number can be as high as 80%. For the nonlinear filters, a "pseudo-χ^2" variable can be defined in analogy with the Kalman filter. This pseudo-χ^2 will obviously not be χ^2-distributed. The distributions of the pseudo-χ^2-probability do not feature the peak at low values observed in the

Table 6.4: The main statistical properties of the 3 vertex fitting algorithms, determined using different data samples, assuming a perfectly aligned tracker. The average computation time per fit and failure rate are also given.

Filter	Average χ^2/ndf	Average $P(\chi^2)$	Average Weight	Average Time (ms)	Failure Rate (%)
$B_s \to J/\psi\ \phi$ - secondary vertex					
KVF	1.32	0.46	1	1.2	0
AVF	0.97	0.52	0.927	3.8	0.025
TKF	0.73	0.61	0.927	2.1	0.074
$B_s \to J/\psi\ \phi$ - primary vertex					
KVF	1.29	0.47	1	3.3	0
AVF	0.79	0.68	0.929	12	0.01
TKF	0.68	0.77	0.934	7.7	0.1
$H \to \gamma\gamma$ (GF) - primary vertex					
KVF	1.42	0.45	1	7.4	0
AVF	0.74	0.8	0.929	32	0.01
TKF	0.65	0.88	0.935	24	0.091
Drell-Yan $\mu^+\mu^-$ - primary vertex					
KVF	1.78	0.31	1	8	0.01
AVF	0.83	0.72	0.91	33	0.04
TKF	0.7	0.84	0.913	27	0.12
$t\bar{t}H, m(H) = 120\ \mathrm{GeV}/c^2$ - primary vertex					
KVF	2.05	0.21	1	13	0
AVF	0.77	0.87	0.905	54	0
TKF	0.67	0.95	0.911	60	0.01

case of the Kalman filter. By construction, tracks with a large χ^2-distance to the vertex, and hence a large contribution to the pseudo-χ^2, are down-weighted, shifting the distribution of the pseudo-χ^2-probability to higher values. The average weight and the average number of discarded tracks for the Adaptive filter and the Trimmed Kalman Fitter are very similar, even though the former uses a soft track-assignment and the latter a hard track-assignment. Since, on average, the same number of tracks remain associated to the vertex, the estimated position from the 2 filters is nearly equal. The Trimmed Kalman Fitter exhibits the highest failure rate, on the order of 0.1%.

The Kalman filter is the fastest of the 3 filters, as the nonlinear filters have to perform more iterations. The Trimmed Kalman Fitter is faster than the Adaptive filter for track multiplicities below approximately 30, and slower for larger multiplicities, indicating different algorithm complexities.

The estimated vertex position and associated errors for a perfectly aligned tracker are summarized in Table 6.5. To characterize the residual distributions, the mean (or *bias*) and the standard deviation (or *resolution*) of a Gaussian fitted to the distribution are given. The contributions from the tails have been estimated by determining the half-widths of the symmetric intervals covering 95% of the distribution (*95% coverage*) and are also quoted. In case of a Gaussian distribution, the 95% coverage corresponds to twice the standard deviation. The standardized residual (or pull) distributions are also fitted with a Gaussian distribution, and the standard deviation is quoted as the pull value. The distribution of pulls is well described

Table 6.5: The resolution, 95% coverage, and pull of the x and z-coordinates of the reconstructed vertex from the 3 vertex-fitting algorithms, determined using different data samples, assuming a perfectly aligned tracker. The biases are compatible with zero and are not shown.

Filter	x-coordinate			z-coordinate		
	σ (μm)	95% Cov. (μm)	Pull	σ (μm)	95% Cov. (μm)	Pull
$B_s \to J/\psi\ \phi$ - secondary vertex						
KVF	54.8	164	1.08	73.8	471	1.08
AVF	53.6	155	1.02	73	440	1.02
TKF	54	174	1.04	75	502	1.05
$B_s \to J/\psi\ \phi$ - primary vertex						
KVF	44.1	176	1.11	54.3	224	1.07
AVF	38.4	94.9	0.94	48.7	140	0.94
TKF	39.4	98.7	0.97	49.5	144	0.95
$H \to \gamma\gamma$ - primary vertex						
KVF	28.1	124	1.11	34	152	1.06
AVF	22.1	73.7	0.9	29.2	106	0.9
TKF	23	74.9	0.93	29.6	111	0.92
Drell-Yan $\mu^+\mu^-$ - primary vertex						
KVF	15.5	77.1	1.51	26.5	119	1.48
AVF	12.7	39.2	1.21	22.5	60.4	1.18
TKF	13.6	39.6	1.21	23	62.5	1.18
$t\bar{t}H, m(H) = 120\ \mathrm{GeV}/c^2$ - primary vertex						
KVF	14	118	1.51	17.9	122	1.46
AVF	9.55	21.1	0.99	13	30.3	1
TKF	9.87	21.7	1.01	13.3	31.7	1.02

by a single Gaussian distribution.

The resolutions of the reconstructed vertex improve with the number of tracks and their average transverse momentum. The residual and the pull distributions for the z-coordinates exhibit a significantly larger tail contribution than the corresponding distributions for the transverse coordinates. Comparing the different filters, the nonlinear filters are more precise than the Kalman filter, especially in high track-multiplicity scenarios, as the outliers can be more easily identified and the vertex better defined by the remaining tracks:

- In the case of $t\bar{t}H$ events, the improvement on the resolution is large (on the order of 30%), and the tails are significantly reduced. The estimated vertex covariance matrix in $t\bar{t}H$ events is much more accurate with the robust filters, the pull distributions being almost perfectly Gaussian with a standard deviation very close to 1.
- In the fit of B_s decay vertices, where only 4 tracks are fitted, little improvement can be seen on the estimated positions with respect to those obtained with the Kalman filter, with only a slight improvement of the coverage. The estimated covariance matrix is slightly more accurate than with the Kalman filter.

In the case of the decay $B_s \to J/\psi\ \phi$, the improvement of the resolution of the primary vertex provided by robust filters yields an improvement in the reconstructed proper decay length. For the Kalman filter, the resolution is $40.6\ \mu$m and the 95% coverage is $129\ \mu$m, while for the

Adaptive Vertex Fitter, the resolution is $37.4\ \mu\mathrm{m}$ and the 95% coverage is $100\ \mu\mathrm{m}$.

6.5.2 The kinematic vertex fit

The kinematic vertex fit has been described extensively in [191]. The goal of a kinematic fit is to improve the resolution of experimental measurements by introducing constraints derived from physics laws into the minimization problem. It is based on a least-squares minimization, where constraints are introduced using the Lagrange multipliers technique. This technique was chosen because it allows modelling of the constraints by linear equations, easing minimization. However, the constraints must always be satisfied exactly.

The $B_s \rightarrow J/\psi\ \phi$ signal sample is used here, without applying trigger or selection cuts. The most important parameters for a full analysis are the mass of the reconstructed B_s and the position of its decay vertex.

The performance of the kinematic fit should be gauged with respect to the results obtained using the Kalman filter, where the mass of the B_s candidate is calculated at the secondary vertex from the tracks refitted with the vertex constraint. The distribution of the reconstructed mass residual thus obtained is shown in Figure 6.21 (left). When fitting this distribution with a Gaussian, the mean is approximately $15\ \mathrm{MeV}/c^2$ higher than the world-average mass of the B_s, and the standard deviation is about $35\ \mathrm{MeV}/c^2$, with a 95% coverage of $97\ \mathrm{MeV}/c^2$. The shift and the asymmetry of the distribution are attributed to the inhomogeneity of the magnetic field in the tracker endcaps, which is only partly taken into account during track reconstruction.

With the kinematic fit, the parameters of the B_s candidate are calculated after fitting the 2 muon and 2 kaon tracks to the same vertex, requiring the invariant mass of the 2 muons to be equal to the mass of the J/Ψ. A constraint on the invariant mass of the 2 kaons cannot be introduced in the Lagrange multipliers technique, since the experimental resolution on the invariant mass of the 2 tracks (around $3\ \mathrm{MeV}/c^2$) is smaller than the width of the ϕ ($4.26 \pm 0.05\ \mathrm{MeV}/c^2$ [119]). The distribution of the reconstructed mass residuals is shown in Figure 6.21 (right). The bias on the reconstructed mass is reduced to $3.8\ \mathrm{MeV}/c^2$, the resolution to $14.0\ \mathrm{MeV}/c^2$, and the 95% coverage to $42\ \mathrm{MeV}/c^2$.

The vertex coordinates are not affected by this constraint.

6.5.3 Primary-vertex finding with Pixel Tracks

The primary-vertex finding based on pixel hits provides the HLT with a primary-vertex position measurement. This measurement is subsequently used for HLT track seeding and in most HLT analyses. The algorithms must therefore be fast and reasonably precise. To satisfy these requirements, primary-vertex finding is reduced to a 1-dimensional search along the $z-$axis.

Pixel hit triplets are collected and the pixel "tracklet" parameters are evaluated using a helix approximation, as described elsewhere [192]. The tracklets are then grouped into primary-vertex candidates.

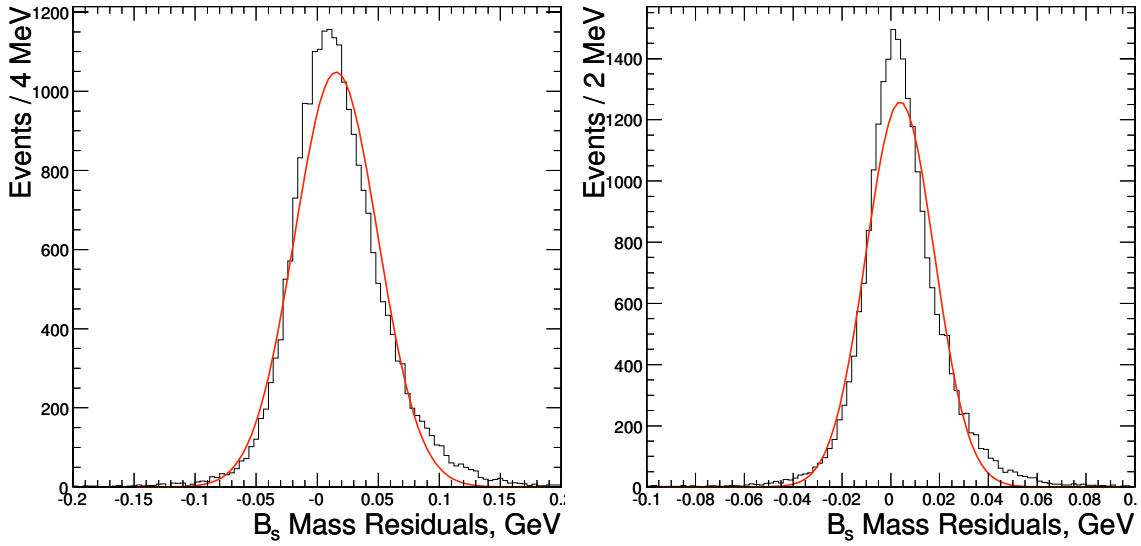


Figure 6.21: Residual of the reconstructed mass of the B_s of the Kalman filter (left) and the kinematic fit (right).

6.5.3.1 Algorithms

Two primary-vertex finding algorithms were implemented and tested. The Histogramming Method progressively merges tracklets that are deemed to be close together, according to the z-coordinate of their point of closest approach to the beam line (z_{IP}). The Divisive Method looks for large z_{IP} intervals without tracks in order to divide the z axis into several regions. An average primary-vertex position is computed from all tracks in each region, and tracks not compatible with that average position are discarded. The discarded tracks are then used to make new vertex candidates. The procedure iterates and stops when all tracks are found to be compatible with the corresponding primary vertex positions.

After vertex finding, the vertex candidates are sorted in decreasing order of the sum of the p_{T}^2 of their associated tracklets. An upper p_{T} limit is set at $10\ \mathrm{GeV}/c$ to account for the imprecision of the p_{T} measurement for pixel tracklets. The signal primary vertex is in most samples the first in the ordered list. A detailed description of the algorithms can be found elsewhere [193].

6.5.3.2 Analysis

Among the primary-vertex candidates, the *closest* primary vertex is defined as that closest in z to the simulated signal primary vertex and the *tagged* primary vertex as the one with the largest p_{T} sum.

The efficiency, ε, for finding a primary vertex (tagged or closest) is defined with respect to vertices reconstructed inside a window of 500 μm around the position of the signal primary vertex in the Monte Carlo simulation.

6.5.3.3 Primary-vertex finding efficiencies and resolutions

The primary-vertex finding efficiencies achieved using the 2 different algorithms are given in Table 6.6, assuming low-luminosity pile-up.

Table 6.6: The efficiencies for primary-vertex finding using the Histogramming and Divisive algorithms assuming low-luminosity pile-up (see text for definitions of ε and $\varepsilon_{\mathrm{tag}}$).

	Histogramming		Divisive	
	ε	$\varepsilon_{\mathrm{tag}}$	ε	$\varepsilon_{\mathrm{tag}}$
u jets; $50 < E_{\mathrm{T}} < 100$ GeV	0.97	0.90	0.99	0.94
b jets; $30 < E_{\mathrm{T}} < 50$ GeV	0.97	0.89	1.00	0.96
$H(115\ \mathrm{GeV}/c^2) \rightarrow \gamma\gamma$, g fusion	0.89	0.75	0.94	0.80
$B_s^0 \rightarrow J/\psi\ \phi$, primary vertex	0.81	0.61	0.97	0.78
$t\bar{t}$	0.99	0.98	1.00	1.00
$t\bar{t}H$, $H(120\ \mathrm{GeV}/c^2) \rightarrow b\bar{b}$	1.00	0.99	1.00	1.00

The Divisive method shows generally better primary vertex efficiencies compared to the Histogramming method, due to the fact that the former takes into account the compatibility of the tracks with the estimate of the vertex instead of considering only the compatibility between adjacent tracks. For most of the event samples the primary vertex is recovered with an efficiency close to 100%. The efficiencies of primary-vertex finding are significantly below 100% for events like $H(m_H = 115\ \mathrm{GeV}/c^2) \rightarrow \gamma\gamma$, where the low charged particle multiplicity does not always allow the signal primary vertex to be distinguished from pile-up interactions. At high luminosity, the inefficiency roughly doubles compared to the results of the Divisive algorithm at low luminosity [194].

The primary vertex z-resolutions achieved for both algorithms are given in Table 6.7, under the assumption of low-luminosity pile-up. A resolution better than 50 μm is achieved for most event samples, which is sufficient for the HLT. Here also the Divisive Method performs better because of the rejection of incompatible tracks.

Table 6.7: The spatial resolution of the z-coordinate of the primary vertex from the Histogramming and Divisive algorithms, at low luminosity. Dashes (–) denote insufficient statistics.

	Histogramming		Divisive	
	σ (μm)	95% Cov. (μm)	σ (μm)	95% Cov. (μm)
u jets; $50 < E_{\mathrm{T}} < 100$ GeV	41.5	81.1	33.2	62.6
b jets; $30 < E_{\mathrm{T}} < 50$ GeV	50.4	187.8	44.4	99.5
$H(115\ \mathrm{GeV}/c^2) \rightarrow \gamma\gamma$ (g fusion)	51.1	–	39.1	–
$B_s^0 \rightarrow J/\psi\ \phi$	69.4	–	74.8	–
$t\bar{t}$	40.1	126.5	32.5	78.9
$t\bar{t}H$, $H(120\ \mathrm{GeV}/c^2) \rightarrow b\bar{b}$	38.6	104.2	28.7	64.7

6.5.4 Primary-vertex finding using fully reconstructed tracks

Primary-vertex finding using fully reconstructed tracks provides a precise estimation of the vertex position and of the vertex position covariance matrix, as well as a list of tracks associated to the primary vertex.

6.5.4.1 Algorithm

The default primary-vertex finder performs primary-vertex reconstruction using all tracks reconstructed in the event. It resembles the Divisive method used for primary-vertex reconstruction in the pixel detector. The 4 main steps are

- track preselection, based on their distance of closest approach to the beam (by default, the transverse impact parameter significance, the ratio of the transverse impact parameter divided by its uncertainty, is required to be < 3) and their p_{T} ($p_{\mathrm{T}} > 1.5$ GeV/c);
- formation of clusters of tracks, based on the z-coordinate of their point of closest approach with respect to the beam line (tracks closer than 1 mm are grouped together);
- a fit of a primary-vertex candidate for each of these clusters, discarding tracks incompatible with the candidate vertex (the default compatibility cut is at 5%);
- the exclusion of poor fits (χ^2 probability $< 1\%$) and vertices incompatible with the beam line (probability $< 1\%$).

All the compatibility probabilities are computed assuming Gaussian resolutions. The compatibility with the beam axis is computed assuming a Gaussian beam spot with a width of 15 μm in x and y.

After vertex finding, the vertex candidates are sorted in decreasing order of the sum of the p_{T}^2 of the associated tracks. A detailed description of the algorithm can be found elsewhere [195].

6.5.4.2 Analysis

The primary-vertex finding algorithm is evaluated according to the following criteria:

1. the efficiency of finding a primary vertex candidate in a window of ± 500 μm around the Monte Carlo signal primary vertex, ε;
2. the efficiency of finding the vertex with the largest $\sum p_{\mathrm{T}}^2$ within ± 500 μm around the Monte Carlo signal primary vertex, ε_{tag};
3. the position resolution, tails, biases, and pulls as defined in Section 6.5.1.

6.5.4.3 Primary-vertex finding efficiency

The primary-vertex finding efficiency obtained from fully reconstructed tracks is given in Table 6.8. Primary vertices with a large number of harder tracks are easier to reconstruct. The lower efficiency compared to the pixel primary-vertex finding is mostly due to the p_{T} cut applied during track preselection. When the p_{T} cut is lowered from 1.5 GeV/c to 0.7 GeV/c, the vertex finding efficiency increases significantly, e.g., up to $(96.3 \pm 0.3)\%$ in the $B_s^0 \to J/\psi\, \phi$ channel.

6.5.4.4 Primary-vertex resolutions

The precisions of the x- and z- coordinates of the reconstructed primary vertex are given in Tables 6.9 and 6.10. Unlike the algorithms discussed in Section 6.5.1, track selection is

Table 6.8: The efficiency of primary-vertex finding with fully reconstructed tracks, assuming a perfectly aligned tracker and low luminosity pile-up (see text for definitions of ε and ε_{tag}).

	ε	ε_{tag}
$H(150GeV/c^2) \rightarrow ZZ \rightarrow 4e$	0.969 ± 0.002	0.9626 ± 0.003
$H(115GeV/c^2) \rightarrow \gamma\gamma$, g fusion	0.86 ± 0.008	0.76 ± 0.01
$H(115GeV/c^2) \rightarrow \gamma\gamma$, VB fusion	0.904 ± 0.005	0.848 ± 0.007
$B_s^0 \rightarrow J/\psi\ \phi$, primary vertex	0.835 ± 0.005	0.663 ± 0.007
b-jets; $30 < E_{\mathrm{T}} < 50GeV$	0.959 ± 0.003	0.92 ± 0.008
$t\bar{t}$	0.987 ± 0.002	0.979 ± 0.002
Drell-Yan 2μ	0.940 ± 0.004	0.930 ± 0.004
$t\bar{t}H$	0.993 ± 0.001	0.989 ± 0.002

performed by the primary-vertex finding algorithm, and the Monte Carlo simulation information is not used.

Table 6.9: The resolutions, tails (95% coverage), and pulls for the primary-vertex x-coordinate, assuming a perfectly aligned tracker and low luminosity pile-up. The biases are compatible with zero.

x-coordinate			
	σ (μm)	95% coverage (μm)	Pull width
$H(150GeV/c^2) \rightarrow ZZ \rightarrow 4e$	17	45	1.36
$H(115GeV/c^2) \rightarrow \gamma\gamma$, g fusion	25	79	0.91
$H(115GeV/c^2) \rightarrow \gamma\gamma$, VB fusion	20	74	0.97
$B_s^0 \rightarrow J/\psi\ \phi$, primary vertex	44	121	1.14
b-jets; $30 < E_{\mathrm{T}} < 50GeV$	24	67	1.06
$t\bar{t}$	13	34	1.16
Drell-Yan 2μ	13	46	1.10
$t\bar{t}H$	10	27	1.19

The position resolutions improve with the number of tracks associated to the vertex and with their average p_{T}. The results are comparable to those determined using a Monte-Carlo-guided vertex fit (Section 6.5.1), sometimes even better as a result of the rejection of incompatible tracks. Thus, the algorithm assigns tracks to the primary vertex with sufficient efficiency and purity such that the intrinsic primary-vertex resolution is not degraded. Lowering the p_{T} cut leads to an improvement in the resolution, e.g., from 45 μm to 37 μm in x and y, and from 64 μm to 48 μm in z for the $B_s^0 \rightarrow J/\psi\ \phi$ channel.

6.5.5 Influence of tracker misalignment on vertex reconstruction

To gauge the robustness of the vertex reconstruction methods with respect to the misalignment of the tracker, 2 misalignment scenarios (short- and long-term) have been tested. These scenarios are described in more detail in Section 6.6.4. The outcome of the tracker misalignment studies are described in detail elsewhere [186, 195].

Table 6.10: The resolution, tails (95% coverage), and pulls for the primary-vertex coordinate along the beam axis, assuming a perfectly aligned tracker and low luminosity pile-up. The biases are compatible with zero.

z-coordinate			
	σ (μm)	95% coverage (μm)	Pull width
$H(150GeV/c^2) \to ZZ \to 4e$	21	64	1.13
$H(115GeV/c^2) \to \gamma\gamma$, g fusion	32	142	0.92
$H(115GeV/c^2) \to \gamma\gamma$, VB fusion	31	134	0.94
$B_s^0 \to J/\psi\ \phi$, primary vertex	65	229	1.26
b-jets; $30 < E_{\mathrm{T}} < 50GeV$	31	108	1.07
$t\bar{t}$	18	51	1.16
Drell-Yan 2μ	25	76	1.14
$t\bar{t}H$	14	35	1.12

Table 6.11: The primary vertex finding efficiency ε, as well as the resolution, 95% coverage, bias, and pull of the x, y, z coordinates, for both the short- and long-term tracker alignment scenarios described in Section 6.6.4, assuming low luminosity pile-up.

	ε	$\sigma_{x,y}$	C95$_{x,y}$	Bias$_x$	Bias$_y$	Pull$_{x,y}$	σ_z	C95$_z$	Bias$_z$	Pull$_z$
Short-term tracker alignment										
$B_s^0 \to J/\psi\ \phi$	0.825	51	128	-5.8	12	1.16	67	235	-2.6	1.28
$t\bar{t}H$	0.958	18	47	2.4	16	1.48	22	66	3.6	1.41
DY	0.914	24	62	1.6	16	1.23	35	93	3.7	1.19
Long-term tracker alignment										
$B_s^0 \to J/\psi\ \phi$	0.826	51	127	-10	11	1.16	67	242	-7.8	1.29
$t\bar{t}H$	0.960	17	47	-9.5	11	1.46	21	58	-4.1	1.38
DY	0.916	22	59	-8.9	11	1.28	33	92	-4.7	1.23

6.5.5.1 Primary-vertex finding efficiency

The effect of tracker misalignment on the primary-vertex finding efficiency is small. A maximum decrease of 3.5% is observed among the samples studied. The reason for this decrease is the increase in the fraction of vertices which fail the selection cut related to the compatibility of the vertex with the beam line.

6.5.5.2 Resolution

The resolutions, fraction of tails, biases, and pull widths for primary vertices are summarized in Table 6.11. The position resolution is significantly affected, as might be expected. The degradation of the resolution is of the order of 6-8 μm in both the short-term and long-term alignment scenarios. Indeed, in both scenarios, the pixel detector is supposed to be aligned with the same precision.

In the long-term misalignment scenario, the primary vertex biases in the 3 coordinates are consistent with the residual shifts of the pixel half-barrels $(\Delta x, \Delta y, \Delta z)$: the innermost track hits are thought to be located at a position $-(\Delta x, \Delta y, \Delta z)$ away from their true position, and

since about half of the tracks cross each half-barrel, the observed biases are approximately the mean of the residual shifts of the 2 half-barrels. In the short-term scenario, the misalignment of the silicon strip tracker is 10 times worse, which further increases the primary vertex bias. The effect is larger for harder events, since the constraint on high-momentum tracks from hits in the silicon strip tracker is stronger (less multiple scattering).

6.5.6 Secondary-vertex reconstruction

6.5.6.1 Inclusive secondary-vertex reconstruction in jets

The challenge of inclusive secondary-vertex reconstruction in jets is that no a priori assumption about the vertex is available. The main focus lies in identifying the tracks that come from a secondary vertex, whereas the main task for exclusive reconstruction is to fit a preselected set of tracks. A good estimation of secondary-vertex parameters is important, e.g., for b-tagging.

Most vertex finders are sensitive to primary and secondary vertices, so a vertex filter is needed to select only the secondary-vertex candidates. The discrimination is based on the distance of a vertex to the beam-line or to an already reconstructed primary vertex.

In most cases b-hadrons produce a tertiary vertex because the decay chain proceeds via charm production (the b-c-decay chain). If tracks coming from a tertiary vertex are also used to fit the secondary vertex the measured flight distance is shifted to a higher value.

A detailed description of the CMS secondary vertex reconstruction algorithms and their performance can be found in [196]. The main results are described below.

6.5.6.2 Algorithms

Trimmed Kalman Vertex Finder (TKVF): The TKVF searches for vertex candidates among the input set of tracks in an iterative way. During the first iteration, a Trimmed Kalman vertex fitter (Section 6.5.1) is applied to the complete input set of tracks, yielding as outputs a vertex candidate and a set of tracks that are incompatible with that vertex candidate. During the subsequent iterations, the same procedure is applied to the set of incompatible tracks identified in previous iterations.

The TKVF is sensitive to primary and secondary vertices, so a vertex filter is used to select secondary vertex candidates. The vertex filter uses the following cuts on the vertices:

- The distance from the primary vertex to the secondary vertex in the transverse plane has to exceed 100 μm, but must not exceed 2 cm.
- The distance from the primary vertex to the secondary vertex in the transverse plane divided by its uncertainty has to be greater than 3: $L_{\mathrm{T}}/\sigma_{L_{\mathrm{T}}} > 3$.
- The total invariant mass of the tracks associated to the vertex must be smaller than 6.5 GeV/c^2.

Tertiary Vertex Track Finder TVTF: This vertex algorithm first uses the TKVF and its vertex filter before trying to find additional tracks coming from a b-c-decay chain. It looks for tracks that are close to the flight-line of a possible b-hadron. These additional tracks are not used to

refit the vertex, but rather to get a more complete reconstruction of the vertex. This can be useful, for example, to discriminate between b- and c-jets.

The p_{T} of the tracks must be greater than 1.0 GeV/c. For jet reconstruction, an iterative cone algorithm with a cone size of 0.5 was used. In the barrel, only jets with $|\eta_{jet}| < 1.4$ are investigated, and in the endcaps only jets with $1.4 < |\eta_{jet}| < 2.5 - \frac{2}{3} \cdot 0.5$ are used. The upper limit for the endcaps was chosen such that the jet is fully contained within the fiducial volume of the tracking detector.

Definitions:

Rate of secondary-vertex finding: The number of jets with a reconstructed secondary vertex divided by the number of used jets.
Efficiency of secondary-vertex finding: The number of jets with a reconstructed secondary vertex with more than 50% tracks coming from the b- or c-hadron divided by the total number of jets. All tracks coming from the b-c-decay chain are considered as good.
Purity of secondary-vertex finding: The efficiency divided by rate.
Track association efficiency: The number of b- or c-tracks associated to the vertex divided by the number of b- or c-tracks associated with the jet.
Track association purity: The number of b- or c-tracks associated to the vertex divided by the total number of tracks associated to the vertex.

6.5.6.3 Performance

Table 6.12 shows the secondary-vertex parameters determined using the TKVF for weakly decaying b- and c-hadrons with a simulated p_{T} of 20–70 GeV/c, in the 3 tracker alignment scenarios previously described. The simulated distance to the beam line of the vertex must be between 100 $\mu\mathrm{m}$ and 2 cm. The 68.3% coverage, the standard deviations of a double-Gaussian fit (from -1 mm to $+1$ mm for the barrel and from -3.5 mm to $+3.5$ mm for the endcaps), and the fraction of vertices covered by the first Gaussian are given. The 3D angle in space is defined as the angle between 2 lines: 1 from the simulated primary vertex to the simulated secondary vertex and the other from the simulated primary vertex to the reconstructed secondary vertex. The resolutions for the flight distance and the 3D angle in space are shown for b-hadrons, pure b-hadrons (all used tracks must be directly from the weakly decaying b-hadron), and c-hadrons in Figures 6.22 and 6.23.

Not all tracks that are used to fit the secondary vertex are coming directly from the weakly deacying b- or c-hadron. Some tracks are from the primary vertex, from pile-up events, or, in case of the weakly decaying b-hadrons, from the tertiary vertex. Because of this, the standard deviation of the pulls is greater than 1; only in the case of pure b-vertices is the standard deviation approximately 1. The resolution from the TVTF is approximately the same as from the TKVF because the additional tracks are not used to refit the vertex, but the efficiencies for vertex finding and track association are higher. Table 6.13 shows the parameters that differ between the 2 algorithms.

Table 6.14 shows the dependency of b-vertex parameters on the simulated distance to the beam line. Table 6.15 shows the vertex parameters for b-, c-, and light-quark-jets jets that lie in the p_{T} range 30–80 GeV/c for jets in the barrel. Here no cut on the simulated distance of

Table 6.12: The efficiency and purity for vertex finding and track association, as well as the resolutions from the TKVF for b- and c-hadrons that lie in the p_{T} range 20–70 GeV/c, in the 3 tracker alignment scenarios previously described.

tracker alignment	perfect				short term		long term	
	Barrel		Endcaps		Barrel		Barrel	
	b	c	b	c	b	c	b	c
secondary vertex finding								
rate (%)	70.3	29.3	62.9	25.3	67.7	23.7	68.0	24.8
efficiency (%)	63.3	20.9	54.9	15.2	61.4	17.3	61.7	18.2
purity (%)	90.0	71.5	87.3	60.1	90.8	73.2	90.6	73.3
track association								
efficiency (%)	75.6	78.3	70.1	70.1	77.5	79.2	77.3	79.2
purity (%)	89.6	76.3	86.9	68.1	89.8	76.7	89.8	77.0
resolutions								
68.3% coverage flight distance (μm)	765	550	2275	1850	800	665	795	630
SD (narrow) of flight distance (μm)	203	254	584	696	224	271	225	273
SD (broad) of flight distance (μm)	581	594	1835	1684	590	631	610	630
fraction of narrow Gaussian (%)	37.4	35.2	36.1	34.6	34.3	25.2	35.3	27.3
68.3% coverage angle in space (mrad)	15.5	8.0	7.5	4.0	19.0	12.5	18.0	11.0

Table 6.13: The efficiency and purity for vertex finding and track association from the TVTF for b- and c-hadrons that lie in the p_{T} range 20–70 GeV/c, in the 3 tracker alignment scenarios previously described.

tracker alignment	perfect				short term		long term	
	Barrel		Endcaps		Barrel		Barrel	
	b	c	b	c	b	c	b	c
secondary vertex finding								
efficiency (%)	63.7	21.1	55.2	15.4	61.7	17.4	62.0	18.2
purity (%)	90.7	72.1	87.8	60.8	91.1	73.2	91.1	73.5
track association								
efficiency (%)	79.5	79.6	73.5	71.4	80.5	80.2	80.5	80.3
purity (%)	89.5	76.2	86.6	68.0	89.6	76.6	89.6	76.8

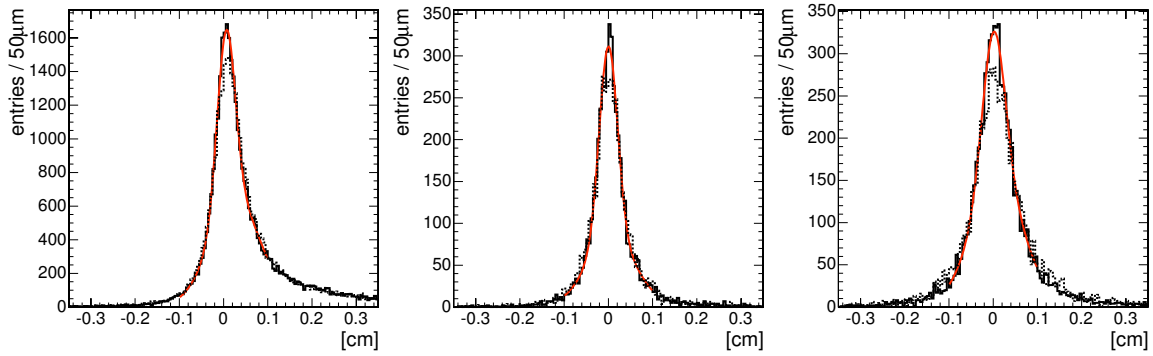


Figure 6.22: The resolution of the flight distances for b-, pure b-, and c-vertices in the barrel. The full histogram indicates the distribution assuming perfect alignment, while the dashed histogram represents the distribution from the short-term alignment scenario. A fit to the former is shown by the solid line.

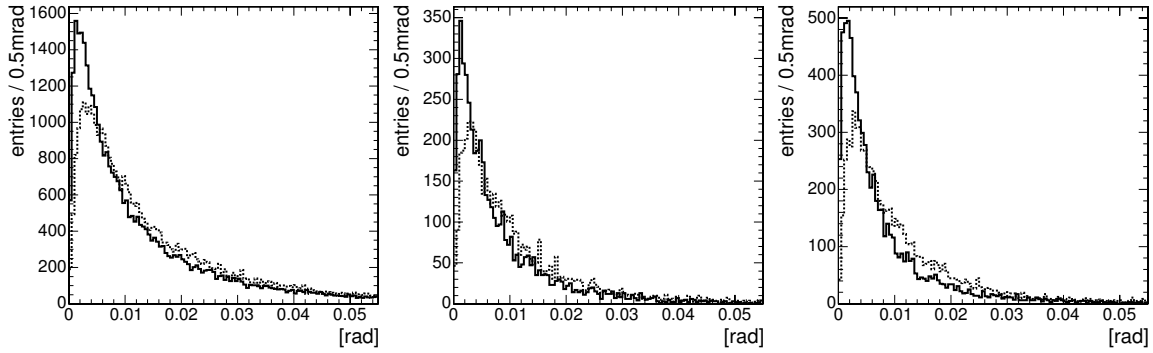


Figure 6.23: The resolution of the angle in space for b-, pure b-, and c-vertices in the barrel. The full histogram indicates the distribution assuming perfect alignment, while the broken histogram represents the distribution from the short-term alignment scenario.

the vertex to the beam line or on the p_{T} of the simulated hadron is required. The gain in the number of tracks per secondary vertex from the TVTF is much higher for b-jets than for other jets.

6.5.6.4 Determination of the beam spot

The position of the beam spot in the experimental area with respect to the CMS detector is expected to vary with time. Rather detailed estimates for the variations on various time scales have been drawn up by the experimental-area interface group:

- During the beam coast (which should take of the order of 10 hours) a very small variation is expected: smaller than 20% of the beam width of 16 μm, i.e., a few microns. This variation is significantly smaller than the best resolution expected on single tracks. Assuming that the detector position is stable to the same level on this relatively short time-scale, short-term variations throughout the physics coast can be safely ignored.
- After initial early-day runs—at very low luminosity—where the beams collide head-on, the crossing angle of the beams will be fixed to 285 μrad. The beam

Table 6.14: The vertex finding rate, efficiency, and purity and the track efficiency and purity for b-hadrons that lie in the p_{T} range 20–70 GeV/c for b-hadrons in the barrel as a function of the simulated distance to the beam line.

	100 μm–500 μm	500 μm–2 mm	2–5 mm	5 mm–2 cm
vertex finding rate	31.1%	63.9%	81.2%	84.2%
vertex finding efficiency	24.7%	56.1%	73.8%	78.7%
vertex finding purity	79.3%	87.8%	90.8%	93.4%
track efficiency	64.2%	75.6%	76.9%	75.9%
track purity	78.7%	85.3%	91.0%	94.4%
SV tracks used in PV	2.22	1.15	0.36	0.10

Table 6.15: Vertex finding rate, number of secondary vertices per jet with a reconstructed secondary vertex, and number of tracks per secondary vertex for jets that lie in the p_{T} range 30–80 GeV/c for jets in the barrel.

	b-jets	c-jets	uds-jets	b-jets	c-jets	uds-jets
p_{T} of jet	30-50 GeV			50-80 GeV		
vertex finding rate	49.5%	17.2%	1.00%	64.5%	21.6%	1.26%
number of SVs per jet with a reconstructed SV	1.09	1.03	1.01	1.18	1.05	1.02
number of tracks per secondary vertex (TKVF)	3.35	2.77	2.54	3.81	2.79	2.39
number of tracks per secondary vertex (TVTF)	3.50	2.82	2.57	4.02	2.86	2.41

inclination is expected to be extremely stable throughout the beam coast and between fills. Time-dependent deviations of the detector symmetry axis with respect to the nominal z-axis may occur.

- The position of the centre of the collision area is expected to be reproducible from one fill to the next only at the level of 1 mm.

Thus, while short-term variations are negligible during the beam coast, large variations occur from one fill to the next.

6.5.6.4.1 Requirements Given the stable beam position over long periods of time (10 hours) a large number of events is available at the offline stage. Therefore, a very precise reconstruction of the 3-dimensional luminosity profile of the interaction point is possible.

In the online environment (HLT) the following procedure is foreseen: At start-of-fill a very fast (on the order of 1 s) reconstruction of the beam spot is performed. This information is fed back to the accelerator. Provided the precision of this procedure is comparable to the beam size (i.e., 15 μm), the online algorithms—for example b-tagging algorithms—can use the nominal position with negligible loss of precision. The speed requirement implies that the beam-spot determination gives a precise, unbiased result with a limited number of events. Moreover, to render the beam-spot determination independent of the correct functioning of the high-level or even Level-1 triggers, the method should be independent of the exact type of events fed into it.

Of course, the beam-spot determination should be efficient over the full range of positions, i.e., even for a macroscopic displacement on the order of a millimetre.

In the following sections, the track-based $d_0\phi$ method and its performance on small numbers of minimum-bias events is evaluated.

6.5.6.4.2 Algorithm The $d_0\phi$ algorithm has been used successfully in previous experiments [197]. For each track the transverse impact parameter d_0—the shortest distance in the transverse plane between the track and the nominal beam axis—is determined. This distance is then signed according to the angular momentum convention. The displaced beam spot will result in a net offset in the impact parameter distribution which depends on the orientation of the track. If the track momentum is along the vector $\vec{V}$ linking the nominal and true beam positions in the transverse plane, no net impact parameter is found. For tracks perpendicular to $\vec{V}$, the displacement is maximum. Plotting d_0 versus ϕ (Fig. 6.24), a typical sine shape is observed. A fit with a sine allows determining the 2 polar coordinates of the beam spot that fix its position in the transverse plane.

In case the detector symmetry axis and the beam line are not parallel, the $d_0\phi$ method may be applied in slices along the z-axis. However, the beam inclinations are expected to be stable to the required precision over long periods of time, even from one fill to the next. Therefore, it is expected that such a procedure will not be necessary in the quasi-online environment.

6.5.6.4.3 Performance To study the performance of the $d_0\phi$ method, a small sample of minimum bias events with displaced beam spot has been generated. The displacement amounted to 1 cm in x and 5 mm in y. No pile-up is included in the current analysis. No

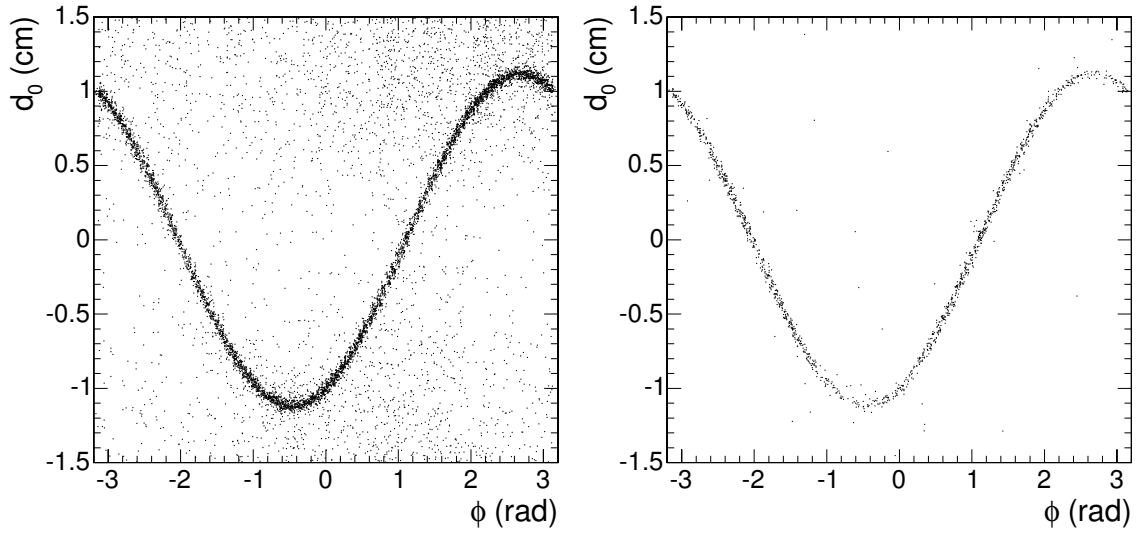


Figure 6.24: Scatter plot of signed transverse impact parameter with respect to the nominal beam spot vs. track azimuthal angle at impact point. At left, all tracks are shown. At right, only tracks satisfying the track quality requirements discussed in the text are taken into account. Result for 1000 minimum bias events without pile-up.

requirement is made on the trigger response to these events. These events with very few stiff tracks correspond to the worst case: studying these events, all assumptions on the trigger are avoided. Of course, if the (Level-1) trigger is available, the number of tracks can be made to increase rapidly and the algorithm converges with even fewer events.

The tracks are reconstructed using the standard CMS combinatorial track finder. Only 1 parameter was changed with respect to the default operation of this track finder: the maximum distance to the nominal (0,0) beam spot allowed in the pixel seed generator is increased from 2 mm to 2 cm.

A number of track quality requirements are applied:

- at least 8 hits in the silicon tracker
- at least 2 hits in the pixel detector
- track fit $\chi^2 < 5$
- transverse impact parameter error $\sigma_{d_0} < 150\ \mu\mathrm{m}$

The wider tolerance (2 cm) on the track origin causes an increased fake rate and a reduction of the efficiency by 20%. The application of additional track quality requirements effectively reduces the fake rate to the level of a few percent.

The polar coordinates of the beam spot are inferred from a simple least-squares fit with a sine function to the shape in Figure 6.24. Outliers—due to fake tracks or conversions—might gain a large weight in the fit and significantly bias the result. Therefore, tracks giving entries that are far (> 3 mm) from the fit are removed. It is found that only 2 or 3 iterations are needed to reach a stable result.

To evaluate the performance of this simple algorithm, several hundred small samples are created. For each of these the reconstructed beam spot is compared to the simulated value. The resolution of the algorithm is evaluated as the width of the residual distribution of polar

coordinates d_0 and ϕ. For samples of 500 events (containing on the order of 1000 tracks) a resolution of 8 μm is found for d_0 and of 0.7 mrad for ϕ, corresponding to an error of approximately 8 μm in the x- and y-coordinates.

In conclusion, with a minor change of 1 of its parameters, the standard combinatorial track finder provides efficient reconstruction of tracks that originate from a vertex displaced by up to 1 cm. The $d_0\phi$ method yields a beam-spot measurement with an error that is sufficiently small compared to the transverse beam size that can be obtained using only 1000 tracks. Thus, 500 minimum bias events suffice to realize the fast beam-spot measurement needed to steer the beams to their nominal position at start-of-fill.

6.6 Alignment

The large number of independent silicon sensors (about $15\,000$) and their excellent resolution of 10–50 μm make the alignment of the CMS strip and pixel trackers a complex and challenging task. The residual alignment uncertainties[1] should not lead to a significant degradation of the intrinsic Tracker resolution. Therefore, the required accuracy of the alignment has to be at least equal to, but ideally significantly better than, the ideal track parameter resolution. Certain physics requirements, such as the $W^{\pm}$ boson mass measurement, place even more stringent constraints on the alignment precision. To achieve a desired precision of $\sigma(M_W) < 15$–20 MeV, the momentum scale has to be known to an accuracy of 0.020–0.025%. This implies the absolute knowledge of detector positions to be known with better than 10 μm uncertainty in the r-ϕ plane. This level of accuracy can only be reached with a track-based alignment procedure.

However, track-based alignment cannot be used as the only method for alignment since it relies on tracks and, in turn, pattern recognition for single track reconstruction already requires the position and orientation of tracking devices to be known with an accuracy higher than the mechanical placement precision of the individual tracking subdetectors. Therefore, the alignment has to be performed in 3 steps: 1) measurement of placement and its precision during assembly of tracking devices, e.g., from photogrammetry and detector position survey measurements; 2) measurement of relative positions of subdetectors using the Laser Alignment System (LAS); and 3) track-based alignment.

6.6.1 General alignment strategy

During the ongoing assembly of the CMS Tracker, positions and orientations of the Tracker components (silicon sensors as well as support structures) are being measured and stored in databases with the help of, e.g., coordinate-measurement machines or photogrammetry. If such measurements can be transformed, directly or indirectly, into detector positions, they will provide a correction to the otherwise assumed ideal Tracker geometry. If no complete measurement of all silicon detectors is available, because only samples have been taken, the standard deviation of the measurements can be used as an estimate of the placement uncertainty. The error of the track hit position on a corresponding module will be enlarged by this value, leading to an improved efficiency during initial track reconstruction.

[1] Unless stated otherwise, the term "alignment uncertainty" refers to the accuracy to which the absolute position and orientation of individual detector elements are known with respect to the global CMS coordinate system.

After assembly, the LAS will provide the positions of the Tracker subdetectors TIB, TOB, and TEC (not including the pixel detector and the TID) with respect to each other with a precision of about $100\ \mu\mathrm{m}$, as described in Section 6.6.3. In order to make efficient pattern recognition for the track reconstruction possible at CMS start-up, it is sufficient that the individual positions of the silicon sensors are known to about $100\ \mu\mathrm{m}$. This can be achieved with a combination of survey and LAS measurements.

The initial track-based alignment procedure will significantly benefit from a good prior knowledge of the positions and orientations of the Tracker substructures. In theory, global alignment uncertainties of $\approx 10\ \mu\mathrm{m}$ or better can already be achieved with a sample of tracks obtained from only few days of running at low luminosity. In practice, however, the time dependent movements of individual detector components, as well as the systematic uncertainties involved in the track-based alignment algorithm, will be the limiting factors for the Tracker alignment. A combination of the LAS and the track-based alignment will be used to monitor the time-dependence of the Tracker alignment due to shutdown operations, dry-out effects, or thermal stress during operation.

It is foreseen that the pixel detector will only be installed in CMS during the shutdown after the 2007 LHC pilot physics run. Therefore, 2 different scenarios have to be considered for track-based alignment, where the pixel detector is either installed in CMS or not. In case the pixel tracker is installed, the first task will be its precise internal and external alignment, so that it can then define the reference system for the subsequent alignment of the strip tracker. The pixel alignment will be performed using tracks that have been reconstructed using pixel hits only (and a vertex constraint) to avoid bias originating from the not yet aligned strip tracker. In the second step, strip tracker alignment will be carried out using the pixel detector as a reference. In case the pixel tracker is not installed (e.g., during the 2007 pilot run), a standalone alignment of the strip tracker has to be performed. This will most likely proceed by first aligning the innermost double-sided layers of TIB and TOB, which will then define the reference system for the outer single-sided layers of the barrel Tracker (and similarly for the endcaps). However, it is currently not expected that the integrated luminosity delivered in 2007 will be sufficient for a precise alignment of the strip detector. Therefore, it will have to be continued in 2008.

Once the Tracker is aligned, it will define the reference system relative to which the positions of the calorimeter modules and muon detectors will be adjusted. In particular, for the optimal reconstruction of muon tracks with high p_{T}, it is important to combine information from the Tracker and the muon detectors. In order to study the impact of both Muon and Tracker misalignment on the reconstruction of muon tracks, the tools described in Section 6.6.4 are used to simulate a realistic environment that includes the expected alignment uncertainties for the individual tracking devices. On the other hand, the alignment of the ECAL modules with respect to the Tracker can be studied by comparing track impact point and calorimeter cluster position for electrons, for example.

6.6.2 Alignment data samples

During nominal LHC collider operation, the data sets ideally suitable for track-based alignment are $W^{\pm} \to \mu^{\pm}\nu$ and $Z^0 \to \mu^+\mu^-$. Already for low luminosity running approximately $20\,000\ Z^0 \to \mu^+\mu^-$ and $100\,000\ W^{\pm} \to \mu^{\pm}\nu$ events are selected per day after HLT [8]. It has been estimated that roughly 1 to 2 million tracks possess the statistical power to carry out a full track-based alignment for the entire Tracker. Therefore, assuming the low luminosity

running scenario and only utilizing these well-reconstructed muon tracks, roughly 1 to 2 weeks of data taking are sufficient to collect the required track sample for the alignment of the entire Tracker.

Furthermore, the Z^0 mass constraint can be exploited not only to correlate the detector parts that are not crossed by a single collision track (e.g., the 2 endcaps), but also to tightly constrain the momentum scale of the muon tracks. The invariant mass constraint could also be applied to $J/\psi \rightarrow \mu^+\mu^-$ decays. However, in contrast to the Z^0 decay, muon tracks from J/ψ have lower momentum, so that further studies are needed to determine their usefulness for the alignment.

Apart from isolated muon tracks, it might also be possible to use well-measured tracks from minimum bias events for the track-based alignment. This option is currently under investigation. In particular, for the first data-taking period where the machine luminosity is expected to be very small, minimum bias events might be the only source providing a sufficient amount of tracks for a comprehensive alignment of the Tracker at start-up.

For the beginning of data taking, but possibly also as a permanent complement to the alignment information from collision tracks and hardware alignment systems, other event samples like cosmic muons as well as beam gas and beam halo muons might become very useful. A first feasibility study suggests that the abundance as well as the complementary topology of these samples are very attractive for Tracker alignment. The TOTEM [198] T1 telescopes can be used to trigger on beam gas and beam halo events, while the CMS muon trigger can be configured to trigger on cosmic muons. With an expected rate of 35 Hz of cosmic muon tracks with $E_\mu > 10$ GeV leaving at least 1 hit in the Tracker barrel, and assuming a uniform distribution of the tracks [199], roughly 700 cosmic muons will cross the entire barrel part of the CMS Tracker per hour. Cosmic muons could thus be used to 1) improve the alignment uncertainties and LAS starting values in the Tracker barrel region (important for pattern recognition); 2) perform an initial barrel alignment prior to the start-up of the LHC; and 3) generally serve as a tool for gaining first operational experience (dead modules, data readout, etc.). On the other hand, beam-halo muons originating from beam background traverse the CMS detector horizontally and thus again provide complementary information for the alignment. First studies of such an event sample suggest that the rate and energy distribution of beam halo muons is beneficial for the Tracker alignment [199].

6.6.3 Laser Alignment System

The LAS uses infrared laser beams to monitor the positions of selected detector modules of the strip tracker and of special alignment sensors in the Muon system[2]. Therefore it operates globally on the larger Tracker composite structures (TIB, TOB, and TEC disks) and cannot determine the position of individual modules. The goal of the system is to generate alignment information on a continuous basis, providing geometry reconstruction of the Tracker substructures at the level of 100 μm, which is mandatory for pattern recognition and for the High Level Trigger (HLT). In addition, possible Tracker structure movements can be monitored at the level of 10 μm, providing additional input for the track-based alignment.

The LAS design is illustrated in Figure 6.25. Each Tracker endcap (TEC) uses in total 16 beams, distributed in ϕ and crossing all 9 TEC disks in rings 4 and 6, which are used for the internal alignment of the TEC disks. The other 8 beams are foreseen to align the TIB, the

[2]It is important to note that no laser alignment is foreseen for the pixel tracker or for the TID.

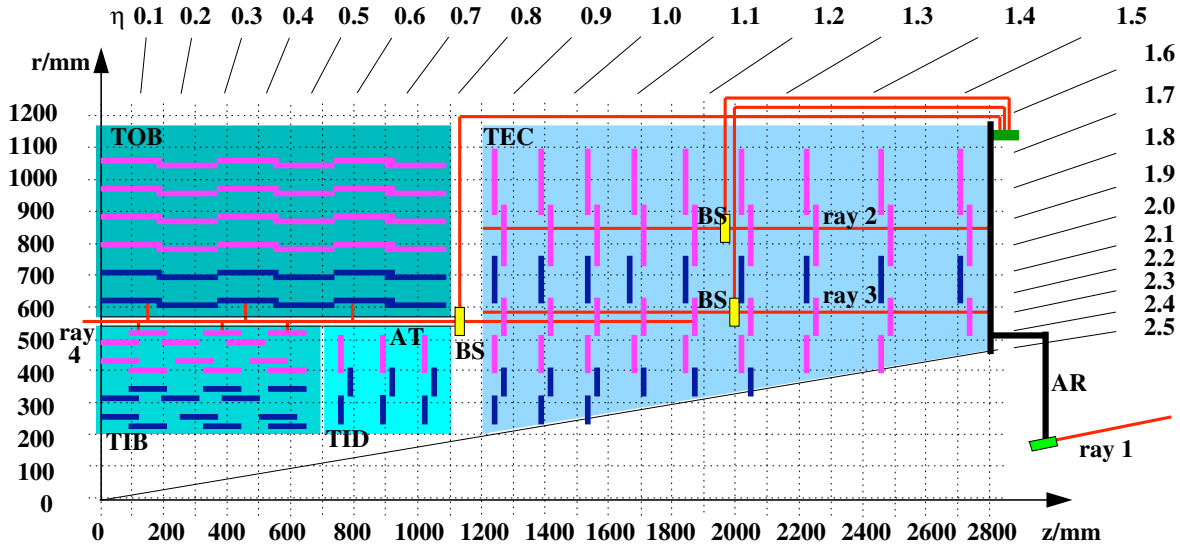


Figure 6.25: Overview of the CMS Laser Alignment System. The laser beams are distributed by Beam Splitters (BS) and Alignment Tubes (AT). The link to the muon system is implemented on the Alignment Rings (AR) that are connected to the Tracker Back Disks.

Table 6.16: Results of a LAS alignment simulation in a benchmark scenario. Azimuthal angle ϕ corrections are given (in mrad) for 8 TEC disks after LAS alignment and compared with the input values.

	Fitted Correction	Error	Input Value
Disk 0	−0.505	± 0.021	−0.500
Disk 1	−10.796	± 0.024	−10.800
Disk 2	13.192	± 0.022	13.200
Disk 3	−7.704	± 0.026	−7.700
Disk 4	21.199	± 0.026	21.200
Disk 5	16.699	± 0.023	16.700
Disk 6	−9.999	± 0.027	−10.000
Disk 7	15.909	± 0.028	15.900

TOB, and both Tracker endcaps with respect to each other. Finally, there is a link to the Muon system, which is established by 12 laser beams (6 on each side) with precise positions and orientations in the Tracker coordinate system.

To measure the intersection point of the laser beams with the silicon modules with high accuracy, a good signal-to-noise (S/N) ratio is required for the signals on the modules. The signals induced by the laser beams on the silicon modules decrease in height as the beams penetrate through the silicon layers. To get optimal signals on all layers, the lasers fire a sequence of pulses with increasing intensities, which are optimized for a given silicon layer. To further improve the S/N ratio, several triggers per intensity are taken and the signals are averaged. In total, a few hundred triggers are needed to get a full picture of the alignment of the Tracker structure. Since the trigger rate for the alignment system is around 100 Hz, this will take only a few seconds. Such snapshots will be taken at regular intervals.

The LAS is foreseen to operate both in dedicated runs and during physics data taking. The trigger type is written into the FED header and thus the alignment events can be identified. They are treated as ordinary events by the event builder, written into the calibration data

Table 6.17: Expected RMS values for Δx, Δy, Δz, and R_z for layers and disks (after laser alignment). These values are used for the "First Data Taking" scenario.

	Δx (μm)	Δy (μm)	Δz (μm)	R_z (μrad)	LAS available
Pixel Barrel	10	10	10	10	no
Pixel Endcap	5	5	5	5	no
Strip Inner Barrel	100	100	500	90	yes
Strip Outer Barrel	70	70	500	60	yes
Strip Inner Disk	400	400	400	100	no
Strip Endcap	60	60	500	45	yes

stream, and a copy is sent to the online laser alignment process for further evaluation.

In the data analysis, the first step is to accumulate all events belonging to one snapshot and to determine for each silicon module the events with optimal signal, which are then averaged. Next, a Gaussian is fitted to the laser profiles. From the fits, the actual positions and uncertainties for all considered planes are calculated. Accumulating these positions delivers the input for the alignment procedure, which should finally give the geometrical corrections for the different Tracker substructures with respect to each other.

The procedure is applicable to both simulation and real data within the CMS reconstruction framework. In order to use real data and simulated events within the same program packages the LAS has been simulated within the GEANT4 [30] framework. This allows the usage of the standard CMS software with only marginal changes. Laser beams are generated using optical photons taking into account the absorption, reflection, and refraction of silicon. The photons are partially absorbed in the silicon and produce simulated hits around the interaction point of the laser beams and the silicon micro-strip modules. To determine the position of the beam spots on the modules, the hits are first clusterized, before Gaussian profiles are fitted to the simulated distributions.

Table 6.16 shows the results of a simulation of the LAS alignment performance in a benchmark scenario in which 9 TEC disks (one endcap) have been shifted both in x and z and rotated by an azimuthal angle ϕ around the beam axis. The agreement between the fitted corrections and the input values is well within the LAS specification.

6.6.4 Simulation of misalignment

The displacement of silicon sensors from their expected position in the Tracker is one of the largest potential sources of tracking uncertainties. To study the impact of Tracker misalignment on track and vertex reconstruction in concrete physics analysis channels (see Volume 2), as well as to study track-based alignment algorithms, a realistic model of misalignment effects has been implemented within the standard CMS reconstruction software.

The displacement of detector modules is implemented after detector simulation at the reconstruction level using a dedicated software tool, which is able to move and rotate all Tracker parts (individual modules as well as rods, layers, half-barrels, etc.). Hits on Tracker sensors are generated according to the ideal detector geometry, and the geometrical shift and rotation of the sensor are introduced afterward. Within the quoted uncertainties, the misalignments are applied using a flat distribution for mechanical constraints, whereas a Gaussian distribution is used for laser and track-based alignment results. To achieve a reasonable χ^2

Table 6.18: Mounting precisions (in μm) used in the misalignment simulation.

	Pixel		Silicon Strip			
	Barrel	Endcap	Inner Barrel	Outer Barrel	Inner Disk	Endcap
First Data Taking Scenario						
Modules	13	2.5	200	100	100	50
Ladders/Rods/Rings/Petals	5	5	200	100	300	100
Long Term Scenario						
Modules	13	2.5	20	10	10	5
Ladders/Rods/Rings/Petals	5	5	20	10	30	10

distribution in the track fit, the hit position error is increased by adding an additional error in quadrature that reflects the size of the assumed misalignment (alignment position error).

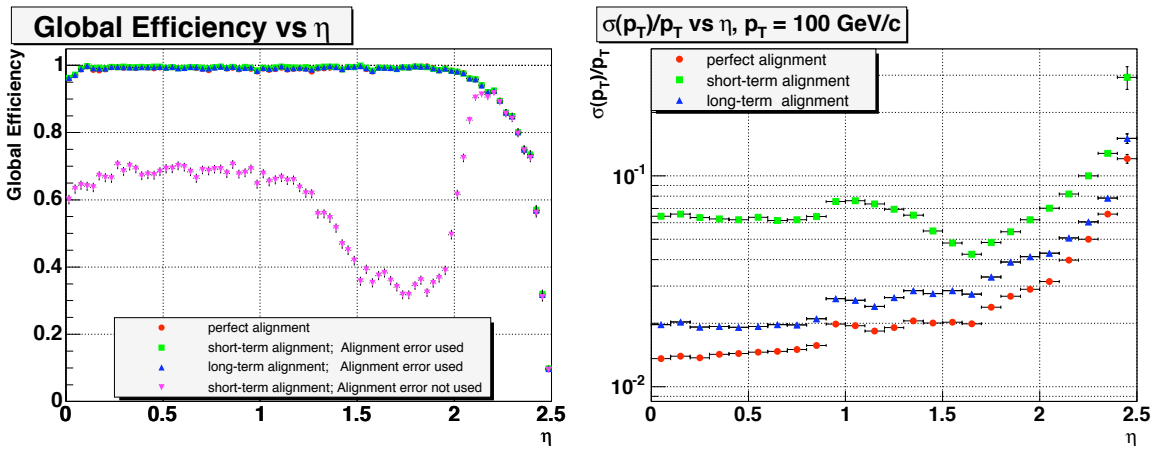
Two different default misalignment scenarios have been developed, which can be easily used but also modified if needed, as explained in [94]. The "First Data Taking" scenario corresponds to the expected misalignment conditions during the first data taking up to a few hundred pb^{-1}, whereas the "Long Term" scenario corresponds to the expected level of misalignment after the first full track-based alignment has been carried out.

- **First Data Taking scenario:** This scenario is supposed to resemble the expected conditions during the first data taking of CMS (few $100\ \mathrm{pb}^{-1}$ of accumulated luminosity). It assumes the availability of the LAS for the alignment of the larger structures of the strip tracker, that the pixel detector has been aligned to a reasonable level using tracks, and that survey results from the Tracker construction, e.g., from photogrammetry, are available. It is further assumed that no track-based alignment of silicon-strip detectors is possible, owing to insufficient high p_{T} tracks. Based on the experience from other experiments, the track-based alignment of the pixel detector would have reduced its placement uncertainties by a factor of 10. Table 6.17 summarizes the expected uncertainty in the detector positions for the various Tracker subsystems after laser and track-based pixel detector alignment.

  In the misalignment simulation, a Gaussian random number generator with a fixed seed is used to generate 1 random number per entry in Table 6.17, centred around 0 and with its width set to the stated uncertainty. This number is then used as a correlated shift (or rotation) for all corresponding structures.
- **Long Term scenario:** It is assumed that after the first few fb^{-1} of data have been accumulated, a first complete track-based alignment down to the sensor level has been carried out, resulting in an overall alignment uncertainty of the strip tracker of $\approx 20\ \mu\mathrm{m}$.

The scenarios perform realistic displacements of individual sensors according to currently estimated placement uncertainties, listed in Table 6.18. These misalignment scenarios will be updated when new information about the performance of the alignment procedures and/or placement uncertainties becomes available. As soon as the actual alignment corrections for the real CMS Tracker become available, they will be used for a realistic time-dependent simulation of the Tracker geometry.

As an illustration of the successful implementation and use of these misalignment scenarios, Figure 6.26 shows the effects of misalignment on track-finding efficiency and transverse



(a) (b)

Figure 6.26: Impact of misalignment on (a) the track finding efficiency and (b) the transverse momentum resolution as a function of η for single muons.

momentum resolution for single muons [195]. The global track finding efficiency is close to unity for $|\eta| < 2$ for all misalignment scenarios, provided the alignment position error is taken into account. If not, the efficiency is significantly reduced, which is illustrated for the short term scenario. The dip in the distribution in the range $1.2 < |\eta| < 2.0$ is due to tracks passing through the TID, which has large alignment uncertainties due to the missing LAS (Table 6.17). For $|\eta| > 2.2$ the inclusion of the alignment position error does not improve the efficiency due to the large track extrapolation uncertainties involved in the very forward direction.

6.6.5 Track-based alignment

Track-based alignment was shown to be the optimal method for the alignment of large tracking detectors in previous particle physics experiments. However, it represents a major challenge at CMS because of the enormous number of degrees of freedom involved: Considering at least 6 alignment parameters (3 shifts, 3 rotations, and eventually tilts or sags) for each of the $\mathcal{O}(15\,000)$ silicon sensors, one has to consider a problem with $\mathcal{O}(100\,000)$ unknowns. Moreover, the full covariance matrix is of size $\mathcal{O}(10^{10})$.

Track-based alignment algorithms have been implemented within the standard CMS reconstruction software through a common software framework, which provides management and input/output of alignment parameters and correlations, and allows the alignment of individual modules as well as of composed structures (e.g., barrels, rods) in a straightforward way by providing the necessary derivatives. In addition, a track refit at the DST level (using the hits attached to the persistent tracks without repeating pattern recognition) has been implemented to improve processing time.

In CMS, 3 different track-based alignment algorithms are considered for track-based alignment, some having been established at other HEP experiments, others newly developed. In the following sections, the main features and initial results of these algorithms are summarized.

6.6.5.1 Hits and Impact Points algorithm

An iterative alignment algorithm using the Hits and Impact Points (HIP method) has been developed [200]. The algorithm is able to determine up to 6 alignment parameters for individual sensors. It involves an analytic formula of the hit residuals as a function of the alignment parameters of N selected individual modules. A χ^2 function, which depends on the alignment parameters, is constructed from the residuals. The minimizing procedure for the $6N$ alignment parameters involves a block diagonal $6N \times 6N$ Jacobian matrix whose inversion reverts to the inversion of individual 6×6 matrix blocks so that inversion of very large matrices is avoided. The implementation of the method allows to fix a subset of parameters for all or for chosen modules.

In the HIP method the alignment parameters are updated only after accumulating all the selected track and hit entries. This approach has the following benefits: 1) it does not easily converge to a local minimum and 2) the module inter-correlations become automatically (or implicitly) included in the procedure. The iteration in the method involves consecutive cycles of performing the alignment and refitting the particle tracks, until no improvement in the track reconstruction is obtained.

An alternative implementation of the algorithm is designed to align composite detector structures for a common rotation and translation [201], for example pixel ladders or layers. The composite alignment involves only 6 parameters of the composite object, and therefore a rather small number of tracks is sufficient to carry out alignment already in the beginning of data taking.

The HIP algorithm has been designed especially in view of the pixel alignment. For this purpose, tracks with relatively low p_{T} can be used, since the multiple scattering and curvature effects are smaller near the beam-line. In addition, the primary-vertex constraint helps: the tracks from the primary vertex will be reconstructed to originate from a single point. This constraint affects the residuals and hence also the alignment parameters. The composite alignment alternative is suitable also for aligning the strip tracker and the Muon system with respect to each other using high p_{T} muons.

The HIP algorithm has been used [201] for the alignment of the pixel barrel modules using the First Data Taking misalignment scenario (enlarged by a factor 10 in order to resemble pre-start-up conditions; only translations). The pixel endcaps and the strip tracker are not misaligned. The procedure has been iterated 10 times using $200\,000$ simulated $Z^0 \rightarrow \mu^+\mu^-$ events. Figure 6.27 shows the residuals (difference of true and estimated misalignment in global coordinates) in x (top), y (middle), and z (bottom). On the left the residuals are shown as a function of iteration. On the right, the residuals are projected for the initial misalignment (I=0) and after 1, 5, and 10 iterations. The convergence is very good, with RMS values of 7 μm (23 μm) for the x and y (z) coordinates.

6.6.5.2 Kalman filter algorithm

A method for global alignment using charged tracks can be derived from the Kalman filter. The method is iterative, so that the alignment parameters to be estimated are updated after each track. It can be formulated in such a way that no large matrices have to be inverted. The update formulae for the alignment parameters and for the associated covariance matrix can be found in [202]. In order to achieve a global alignment, the update is not restricted to the detector elements that are crossed by the track. On the other hand, it need not be applied to

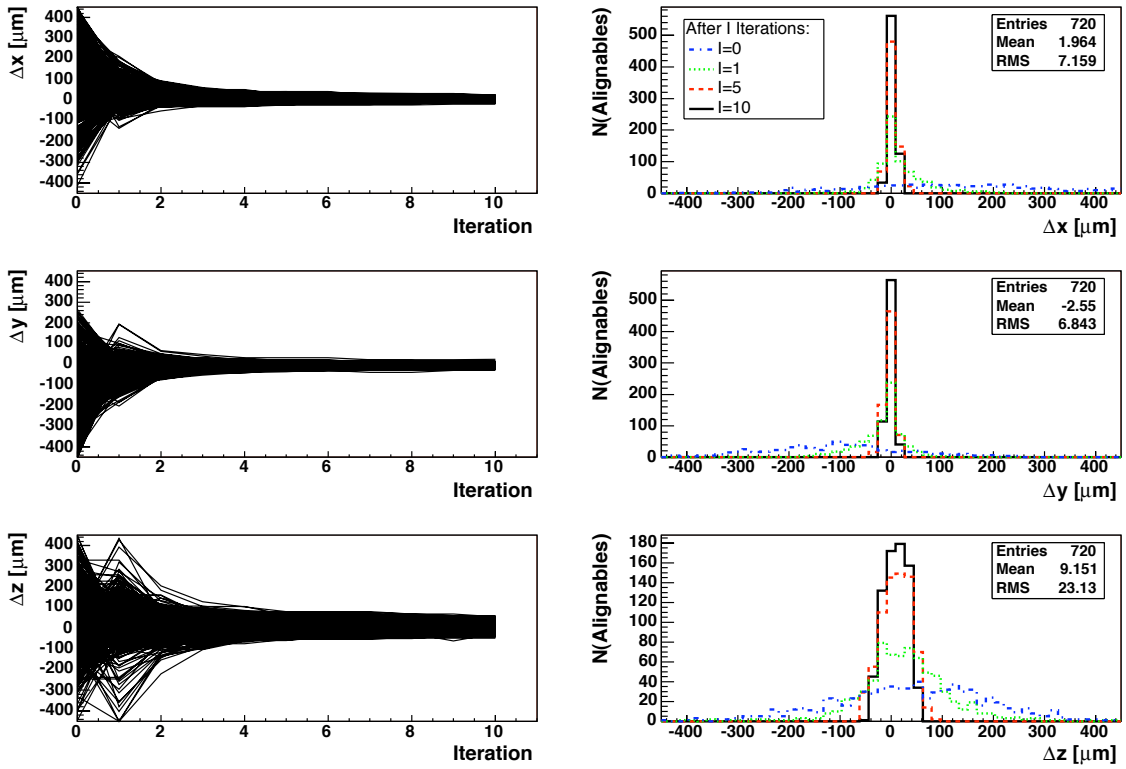


Figure 6.27: Alignment of all 720 pixel barrel modules using the HIP Algorithm. The residuals in global coordinates are shown as a function of iteration (left) and projected for 0, 1, 5, and 10 iterations (right).

all detector elements; rather only those detector elements are updated that have significant correlations with the ones in the current track. This requires some bookkeeping, but allows one to reduce the computational load to an acceptable level. In practice, a list L_i is attached to each detector element i, containing those detector elements that have significant correlations with i. This list may contain only i itself in the beginning and grows as more tracks are processed. In the current implementation, the list management is based on a metrics defining a distance $d(i, j)$ between any 2 detector elements. A detector element is kept in the list L_i only if its distance to i does not exceed an upper bound. Details about the metrics can be found in [202].

It is possible to use prior information about the alignment obtained from mechanical and laser measurements. The position of certain detector elements can be fixed by giving them a large prior weight (small prior uncertainty). A requirement that several detector elements move simultaneously can be enforced by large prior correlations. The formalism can be applied to the alignment of either individual modules or larger structures. It can also be extended to deal with geometrically and kinematically constrained track pairs (or even track bundles) rather than single tracks.

In the beginning, the covariance matrix $\mathbf{D}$ of the alignment parameters is in general block-diagonal and contains the parameters prior uncertainties. If required, it may also contain prior correlations among different detector elements. When a track is processed, the correlation lists L_i of all detector elements i crossed by the track are scanned and, if necessary, extended. Only the detector elements contained in the combined list $L = \bigcup_i L_i$ are updated

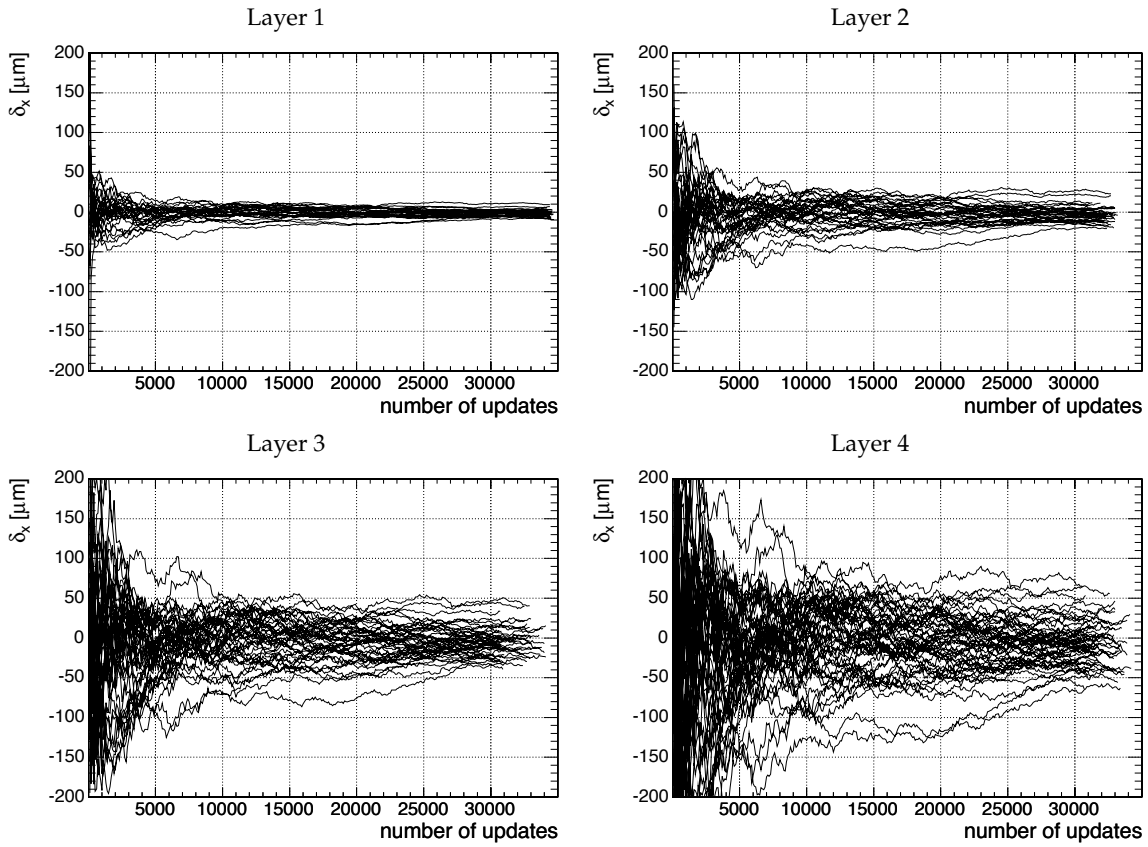
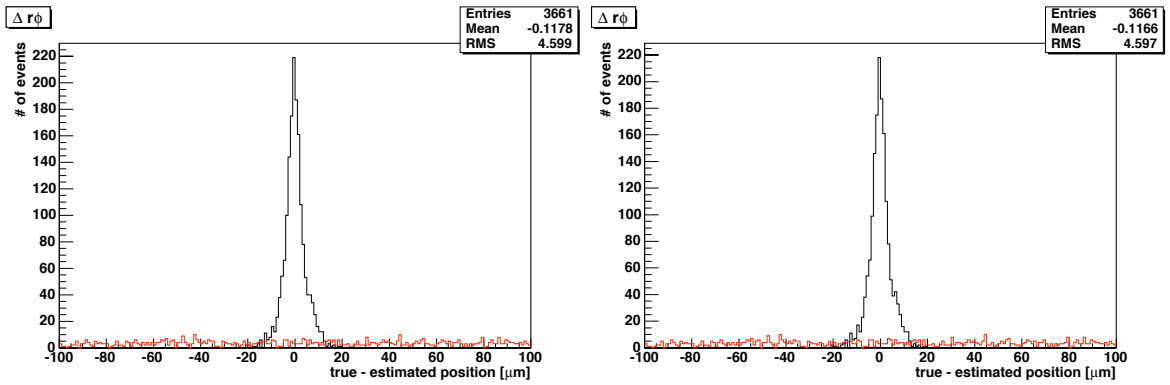


Figure 6.28: Kalman filter algorithm: Evolution of the differences between the estimated and the true local x-shifts for the wheel-like setup after $100\,000$ processed tracks. Here only correlations between modules with $d(i,j) \leq 6$ were considered. The pixel detector serves as the reference frame.

by the Kalman filter. The computational complexity of the parameter update is therefore proportional to the size $|L|$ of L, and the complexity of the covariance matrix update is proportional to $|L|^2$. Restricting the size of the lists L_i without sacrificing precision is therefore of crucial importance.

The algorithm was implemented in ORCA and studied in 2 small subsets of the silicon tracker: a telescope-like section of the inner and outer barrels, and a wheel-like subset of the inner barrel consisting of 156 modules in 4 layers. The tracks used were muons with a p_{T} of $100\ \mathrm{GeV}/c$, generated by the particle gun (inline single-particle simulation) with Gaussian smearing of the hits. While the pixel modules were kept fixed and used as the reference frame, all silicon-strip modules were shifted by a random shift in the local x- and y-direction with a standard deviation of $\sigma = 100\,\mu$m. The upper bound on the distances on the update lists was set to 6. Some results of the wheel-like setup are presented in Figure 6.28. The figure shows the evolution of the differences between true and estimated x-shifts, separately for each of the 4 inner barrel layers. A total of $100\,000$ tracks were processed. The speed of convergence clearly depends on the layer, being fastest for the innermost one, as expected. More results from a toy detector and from the 2 subsets of the silicon tracker can be found in [202].

Although the method has been shown to work in principle, clearly more detailed studies in a larger setup are required. The misalignment will be extended to include rotations. The list



(a) (b)

Figure 6.29: Comparison of (a) MILLEPEDE (matrix inversion) and (b) MILLEPEDE II (iterative solution). The residuals are presented before (dotted) and after (solid) the alignment. The residual distribution before alignment is non-flat, with an RMS of around 200 mm. It appears flat in the histograms because it is truncated to the range shown.

management will have to be tuned on realistically simulated data, and alternative ways of limiting the range of the update need to be explored. As convergence in the outer layers of the inner barrel and in the outer barrel tends to be slow, the algorithm will be extended to deal with kinematically constrained muon pairs from Z^0 or J/ψ decays.

6.6.5.3 MILLEPEDE

algorithm

6.6.5.3.1 MILLEPEDE MILLEPEDE [93, 203] is a well established and robust program package for alignment which has been used and tested successfully at several high energy physics experiments, for example at H1 [204], CDF [205], HERA-b, LHC-b, and others. Being a non-iterative method, it has been shown that it can improve the final alignment precision considerably compared with other algorithms. MILLEPEDE is a linear least-squares algorithm. Such algorithms have proven to be well suited for alignment problems since they are stable, fast, and accurate and can take into account correlations among parameters.

MILLEPEDE distinguishes between global parameters that are common to all data, namely the parameters describing the positions of the detectors, and local parameters, present only in a subset of the data. Track parameters are local parameters as they are specific to a single event. MILLEPEDE performs an overall least-squares fit of the data, fitting all global and local parameters simultaneously. Making use of the special structure of the least-squares matrix in such a fit, the problem is reduced to a matrix equation for the global parameters only. For N global parameters this amounts to an equation with a symmetric $N \times N$ matrix. In the previous version of MILLEPEDE, a solution was found by inverting the $N \times N$ matrix. However, due to CPU and memory constraints, this method can only be used up to $N = 5000$–$10\,000$. Computing time increases with N^3 while memory goes with N^2. The alignment of the CMS Tracker exceeds this limit by at least an order of magnitude, hence new methods had to be developed to cope with the solution of such a system of linear equations. A new version, MILLEPEDE II, was developed, which offers different solution methods, and is applicable for N much larger than $10\,000$.

6.6.5.3.2 MILLEPEDE-II In MILLEPEDE II, the 2 tasks of the program are split: Accumulation of data (MILLE) and solving the set of linear equations (PEDE). The advantage of this procedure is that once a dataset has been defined for alignment, one can test the solution under changing conditions. In addition to the matrix inversion and a diagonalization method, a new method for the solution of very large matrix equations is implemented. This minimum residual method applicable for sparse matrices determines a good solution by iteration in acceptable time even for large N.

As mentioned before, to align the CMS Tracker at the sensor level, the matrix is of the order 10^{10} and matrix inversion is not a viable solution. Diagonalization has the advantage that the eigenvalues provide extra information like the degree of correlation between parameters. However, this method is also rather CPU intensive. The third method, an iterative solution for sparse matrices, is most promising for the CMS Tracker. Here the fact is used that the matrix containing the alignment parameters contains many zero-elements due to detector elements that are not linked to each other via common tracks. However, for comparison and testing of the robustness of the methods, it is useful to compare the results of the various methods.

MILLEPEDE II has been successfully interfaced with the ORCA framework and the alignment of parts of the CMS Tracker has been carried out using different scenarios. For details on these studies see [206]. As an example, Figure 6.29 shows hit residuals for $r\phi$ for the inversion and the iterative method. Here each individual detector of the Tracker was misaligned using the standard CMS procedure. The alignment procedure was carried out in the central region ($|\eta| < 0.9$) of the strip tracker using 1.8 million $Z^0 \rightarrow \mu^+\mu^-$ events. The 3 pixel layers and the outermost barrel layer were kept fixed, which amounted to ≈ 8400 translation parameters to be aligned. The RMS values for both methods are virtually identical ($4.6\ \mu\mathrm{m}$), underlining that the new iterative method performs well.

6.6.5.4 Summary

Only with track-based alignment will it be possible to align the strip and pixel detectors to a level where the residual alignment uncertainties do not significantly degrade their intrinsic resolution. As a starting point, a common software infrastructure has been developed that allows the implementation and finally comparison of several statistical algorithms capable of handling the complexity of the CMS Tracker. Presently, CMS possesses the full software infrastructure to not only realistically misalign the tracking devices, but also to accommodate the use of any statistical alignment algorithm. Three different algorithms, HIP, Kalman filter, and MILLEPEDE, are in the process of being fully implemented and validated. Furthermore, comprehensive tests are ongoing to demonstrate that the considered algorithms exhibit the principal capability to align at least parts of the silicon tracker under realistic conditions.

The next phase will be the development of a realistic alignment strategy of the full Tracker that also takes into account different start-up scenarios (with and without the pixel tracker, see Section 6.6.1). The results of these further studies will be documented in Volume 2.

In addition, the impact of systematic uncertainties on track-based alignment arising from the description of the Tracker material as well as the magnetic field need to be evaluated. In particular, procedures need to be established that allow the material, magnetic field, and alignment effects to be disentangled in the track reconstruction.

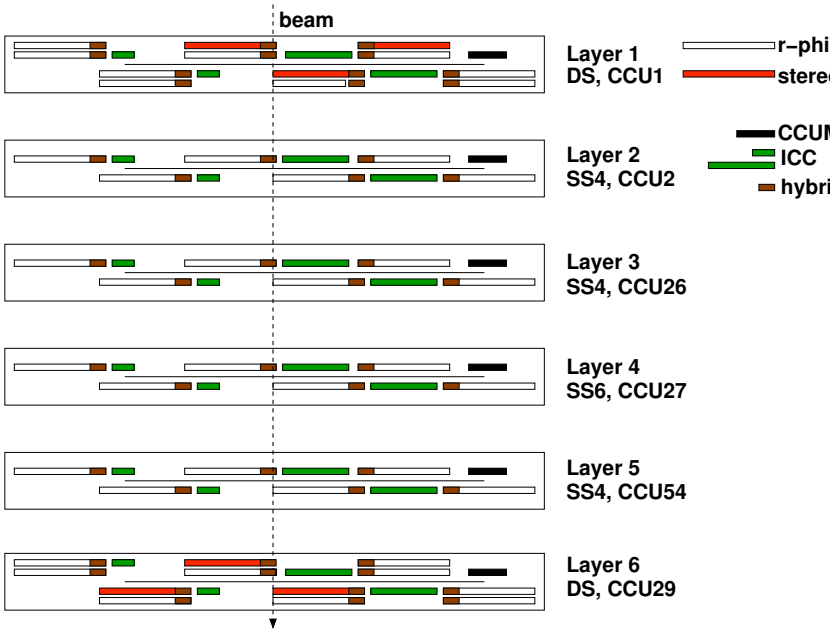


Figure 6.30: Layout of the strip tracker outer barrel (TOB) test-beam setup used for alignment.

6.6.6 Alignment of a test-beam setup

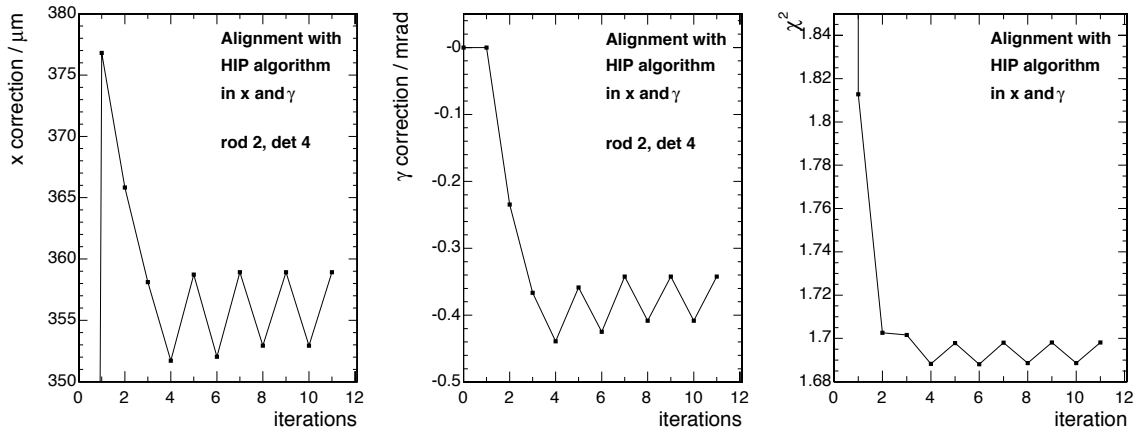
While it is crucial to verify the capability of the aforementioned alignment algorithms to align the full CMS Tracker with simulated data, the results [207] described in this section complement the verification effort with an independent approach, using test-beam data on a small-scale system.

In October 2004, the TOB group recorded data using a 120 GeV pion and a 70–120 GeV muon beam with a dedicated test setup, comprising 48 silicon-strip modules on 6 TOB rods in a configuration shown in Figure 6.30. Since no magnetic field was present, particle trajectories are almost straight tracks which are only mildly affected by multiple scattering. The beam size was about $8 \times 5\ \mu\mathrm{m}^2$ for the pion beam, and the acceptance region for the much larger muon beam was constrained by the trigger scintillator size of about $10 \times 10\ \mathrm{cm}^2$. The setup was adjusted with respect to the beam such that the beam hit the overlap region between 2 adjacent modules (detectors 3 and 4 in each rod). The beam was perpendicular to the strips of the detectors, and in the other direction the beam formed an angle of 94° with respect to the surface. The data were used for track reconstruction and alignment, requiring the adoption of the existing CMS tracking algorithms and alignment software framework for the case of this special geometry and track model.

During alignment, the positions and orientations of the rods in layers 1 and 6 (Fig. 6.30) have been kept fixed, and alignment corrections are only computed for modules 3 and 4 in layers 2 to 5. Results for the HIP and MILLEPEDE algorithms are given in Table 6.19, compared with a manual alignment obtained from track residuals. In addition to the 1-dimensional results where only x, the direction perpendicular to the strips, was aligned, the HIP algorithm was also used to carry out a simultaneous alignment of both x and γ, where γ defines the rotation angle around the axis normal to the sensors. The results of this 2-dimensional alignment procedure are also listed in Table 6.19. The different track-based alignment procedures agree well among themselves and are also compatible with the results

Table 6.19: Alignment corrections obtained from test-beam data. Good agreement among the different algorithms is observed, as well as a small improvement in convergence for the 2-dimensional HIP-2D algorithm.

	manual	HIP-1D	Millepede	HIP-2D	
	x (μm)	x (μm)	x (μm)	x (μm)	γ (mrad)
Rod 2					
Detector 3	-105	-105±2	-101±4	-114±6	-0.12±0.08
Detector 4	363	380±7	379±17	356±13	-0.37±0.18
Rod 3					
Detector 3	-454	-466±2	-457±4	-466±6	-0.00±0.08
Detector 4	-99	-61±7	-96±15	-77±13	-0.26±0.19
Rod 4					
Detector 3	-935	-946±2	-938±6	-954±4	-0.11±0.06
Detector 4	-579	-541±6	-544±16	-532±9	0.22±0.14
Rod 5					
Detector 3	-457	-470±2	-467±4	-479±4	-0.13±0.05
Detector 4	-141	-80±7	-91±17	-67±9	0.27±0.15
mean track χ^2	1.75	1.72	1.73	1.69	



(a) (b) (c)

Figure 6.31: Alignment of test-beam data. Shown are the convergences of (a) x, (b) γ, and (c) the mean χ^2 value. The small oscillations are caused by changes in the track sample.

of the simple manual alignment.

The convergence of the shift x and of the rotation with angle γ for a particular detector using the HIP algorithm with 2 degrees of freedom are shown in Figure 6.31 together with the convergence of the χ^2 values. The overall convergence is quite similar. The 2-dimensional alignment converges to a slightly better χ^2 value, since the correlation between x and γ is taken into account. The alignment precision is limited by the amount of data as well as the small pion-beam size. However, the results prove that the alignment software implementation is working as expected.

6.7 Data quality monitoring

The goal of the data quality monitoring (DQM) for the Tracker is to use the collected data to monitor the performance of the detectors to ensure the smooth running of the system and

that good quality data are being taken. It should also ensure that any possible problem is identified efficiently at a very early stage in the data acquisition chain so that actions can be taken promptly to resolve it. The system is based on the CMS DQM framework discussed in Section 2.9.

6.7.1 Quantities to monitor

Monitoring will be performed at different levels of data processing and reconstruction. Accordingly, the "monitorable quantities" for the pixel and strip detectors in the tracker can be classified into different categories as given below.

- Raw Data from FEDs: inconsistency of channel addresses, bad synchronization of the readout chips, and error flags produced by the FEDs have to be handled and monitored by inspecting the event header of the FEDs.
- Individual channel output: during the Virgin Raw Data mode (Section 6.2.2.1) of strip detector FED operation, the pedestals, noise, and the map of dead and noisy channels will be monitored.
- Raw Hits : the raw charge collected in the detectors, and its position and distribution will be monitored at this level. The frequency, as well as the average charge per pixel or strip, will give an indication of dead or noisy channels.
- Reconstructed Clusters: collected charge in neighboring pixels or strips above threshold are merged into clusters by the reconstruction software and the hit position is calculated. The distributions of the cluster position, cluster size, and cluster charge will be monitored for individual detectors. Since the charge released in silicon varies with the track length, i.e., the incident angle of the particle, the charge collected for single-strip clusters, where almost normal incidence is ensured, will be considered. However, this quantity, as well as the cluster size, is best studied by considering only the clusters attached to tracks, where the track inclination is precisely known.
- Tracks: a number of quantities will be monitored once the tracks are reconstructed (Section 6.4). The number of hits per reconstructed track, χ^2 of the fit, impact parameter, etc. will be monitored to ensure the overall quality of tracks. Events with high momentum isolated tracks are preferable here to ensure reconstruction quality. The cluster width with track incidence angle, and the mean and r.m.s of hit residuals will be used to monitor individual detector performance.

  A few specific items, namely the pixel track seeds (pairs or triplets), pixel tracks, primary vertices from pixel tracks, and impact parameter of the pixel tracks with respect to the primary vertex will be monitored exclusively in the pixel system.

  Due to radiation damage the silicon sensor response is expected to change during the detector operation. A number of quantities—for example Lorentz deflection, signal trapping (Section 6.1.2.1), and inter-channel couplings—will be monitored to guarantee a precise determination of the hit position.

The tracker has a complex structure with $\approx 17\,000$ detector modules arranged in its mechanical structure. To facilitate the navigation, monitorables are arranged in folders reflecting the hierarchical structure. At each level of this logical structure there is a folder containing some header information and the corresponding monitorables. The header for a detector module contains detailed information of all the electrical, optical, and cooling connections and the

detector module identification from the configuration database. This information will enable easy access to the detector modules common to a given control, readout, or cooling circuit.

The environmental information of the tracker consisting of the status of the modules, the power supplies, and the cooling system will constitute a large volume of data not correlated with event data. Part of this information will be utilized in the Detector Control System (DCS, see [8]) logic and will be monitored to ensure reliable running of the tracker. The environmental data are integrated in the DCS in various ways. The temperature and humidity measurements are available via Detector Control Unit (DCU) chips on the front-end module (detector hybrid) and on the Communication and Control Unit (CCU) (see [6]).

The DCUs also read the low voltages and leakage currents of the detectors. The power supply status is monitored and controlled via standard OPC connections to the DCS, realized in the CMS standard PVSS SCADA software. Panels for navigation in the power supply hierarchy, as well as display of the tracker geometry with colour codes associated to the finite state machine of the power conditions, will help to diagnose problems easily. The DCS will have to access the online Configuration and Conditions databases described in Section 6.2.1. The start-up setting of the power supply will be loaded from the Configuration database and the conditions during the run will be written to the Conditions database with appropriate time stamps. In this way, the Data Quality Monitoring applications will be able to correlate event data with environmental data.

6.7.2 CPU estimate and requirements

The monitorables as described in the previous section are booked, filled, collected, and displayed using the common CMS DQM framework described in Section 2.9.2. The booking and filling are performed by "source" applications that have access to the tracker geometry and the detected hardware information (Section 6.3.2) as well as the actual event data. These applications are potentially capable of running the full CMS reconstruction software.

The "source" applications can run directly on the Filter Farm or on dedicated processors accessing event data buffers output from the Filter Farm. The choice depends on several factors like the CPU needed to compute the monitorable quantities, the necessary input event rate, and the acceptable delay by which the monitored information can be computed and made available. Estimates of the CPU usage per event needed to compute a full representative set of monitorable quantities for the entire tracker are presented in Table 6.20. The set of monitorable elements that has been considered involves some basic quantities for each of the $\approx 17\,000$ tracker modules in the strip and pixel tracker, as described in Section 6.7.1. The overall monitorable data volume, computed as the total number of bins multiplied by 4 Bytes, is also shown in Table 6.20. All numbers refer to a high luminosity run and have been obtained by studying a sample of simulated dijet events with transverse momentum between 50 and 100 GeV. The CPU measurements have been performed on a 1 GHz, Pentium III machine with 1 GB RAM.

The numbers in Table 6.20 are obviously driven by the histograms that require 1 bin per channel. The production of higher level histograms, like those with 1 bin per detector, are orders of magnitude less demanding.

Table 6.20: Input data type, overall number of events required, CPU usage (per event), and monitorable data volume for a full representative set of tracker monitoring tasks. Numbers refer to a high luminosity run and have been obtained by studying a sample of simulated dijet events with transverse momentum between 50 and 100 GeV.

Monitorable	Number of events	CPU Required	Data Volume
Raw Hits	1000	Digi building + 0.03 s	1 MB
Clusters	10000	Cluster reconstruction + 0.4 s	50 MB
Tracks	1000	Track reconstruction	1 MB

6.7.3 Event-display-based monitoring

A specialized event-display software tool, called the "Tracker Map" has been developed [208] for monitoring purposes. The tool allows collections of data associated with individual Tracker modules (strip and pixel) to be viewed in a single 2-dimensional image that shows all modules arranged in their mechanical structures (Fig. 6.32): disks for the endcaps and flattened layers for the barrels. The Tracker Map can be used in an interactive mode that provides zooming functionality. In this mode it is possible to get finer details, which can go down to the level of the individual module channels. The basic format in which the information of the Tracker Map is stored is an XML file.

Three main use cases of the Tracker Map are foreseen:

1. A dashboard display in the control room. In this case a static image, produced by specialized monitoring applications with data access, is regularly produced and displayed on a screen. The image can give a fast visual feedback to the operator and can be used to complement more sophisticated applications for automatic error recognition and alarms. Of course this mode of operation can also be used during the data processing at Tier 1 or Tier 2 centers.

2. An efficient way to store and transfer monitoring data about the Tracker. In this case the Tracker Map is used in its original format (XML file). The Map can be visualized through any web browser once the freely available Javascript software is installed on the machine that runs the browser. It should be noted that any other information accessible via the Web, such as histograms, can be linked to the individual modules in the Map by registering their Web address on the file. Zooming functionality is available in this mode and each module on the image is clickable. Clicking a module causes the object attached to it to be accessed and visualized. The same file can be made available on the Web and therefore be accessible from anywhere in the world. In this way, information about the status of the Tracker can immediately be made available to experts or any CMS collaborator.

3. A graphical interface fully integrated with monitoring applications that have online or offline access to event or non-event data. An implementation of this mode of operation has been achieved with the IGUANA software project (Section 2.10). In this case the Tracker Map image actually works as an event display, since the image is updated every time new events are available.

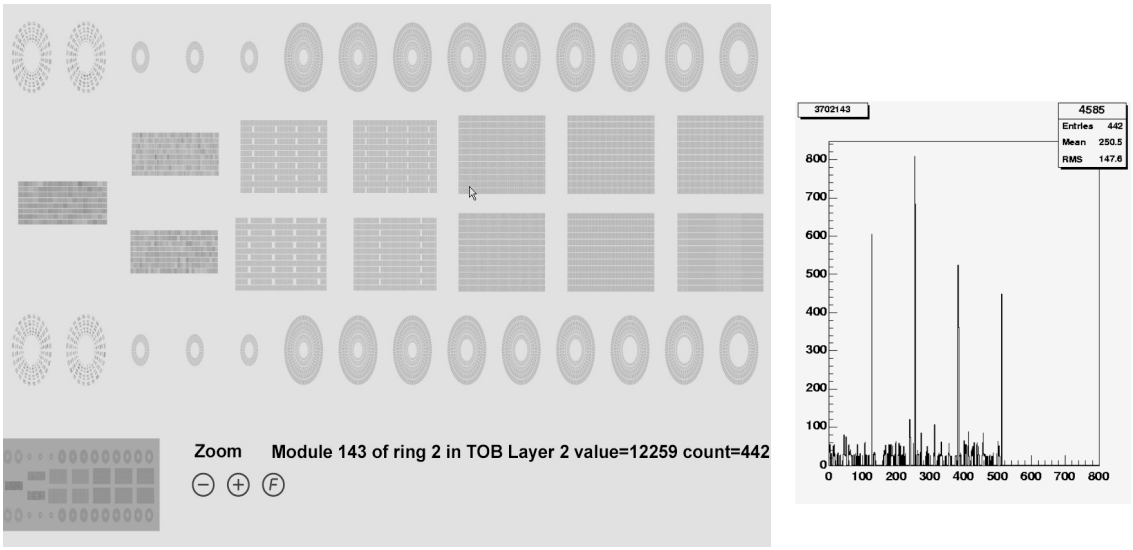


Figure 6.32: Example of a Tracker Map image. The mean ADC counts in each module is depicted in the gray scale image to the left. Clicking on one of the modules causes the detailed occupancy histogram of that module to be shown on the right.

Chapter 7

Forward Detectors

7.1 Overview

The central detector of the CMS experiment has an acceptance in pseudorapidity η, of roughly $|\eta| < 2.5$ for tracking information and $|\eta| < 5$ for calorimeter information. Figure 7.1[209, 210] shows the expected pseudorapidity distribution of the charged particles and of the energy flow at the LHC, demonstrating that with an acceptance limited to $|\eta| < 5$ most of the energy in the collision will not be detected.

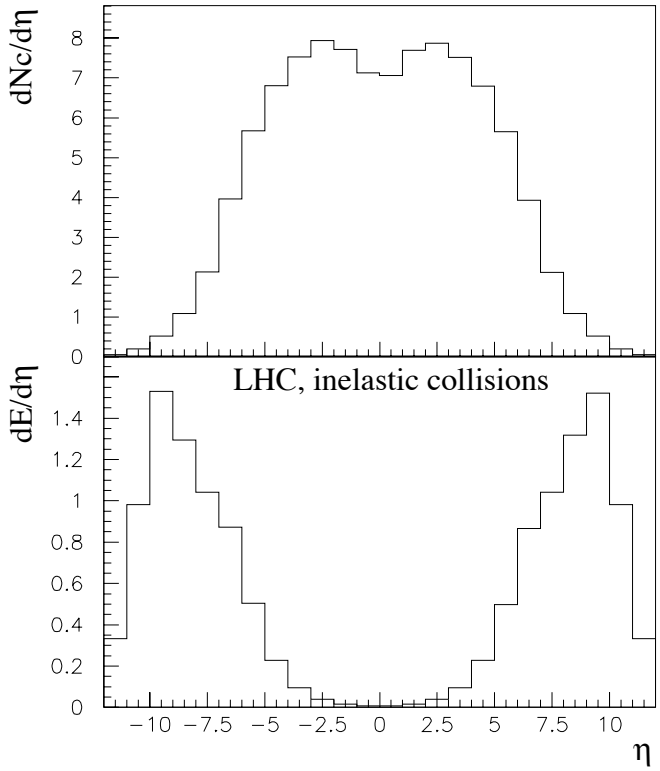


Figure 7.1: Pseudorapidity distribution of the charged particles and of the energy flow at the LHC, taken from [209, 210]. The energy is in units of TeV.

Presently there are 2 proposals to extend the coverage in the forward region:

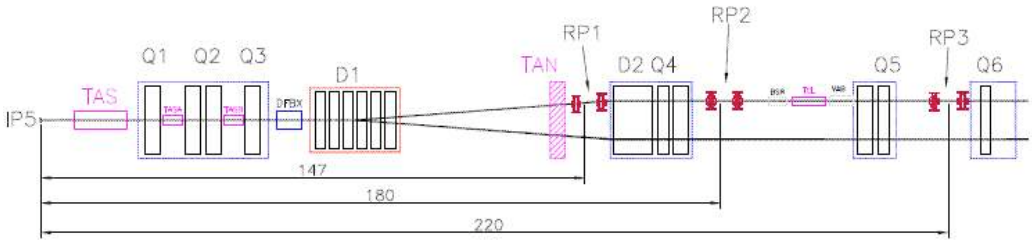


Figure 7.2: The LHC beamline and the Roman Pots of TOTEM at 3 different locations. Distances are measured in meters. The detectors at 180 m are optional.

- Add 2 calorimeters on either side of the interaction region which will cover higher $|\eta|$ values, called CASTOR ($5.1 < |\eta| < 6.5$) and the Zero Degree Calorimeter (ZDC). Both have an electromagnetic and hadronic section. These calorimeters are of interest for measurements in pp, pA and AA collisions (where A is a heavy ion).
- Capitalize on the opportunity to have common runs with the TOTEM experiment, which uses the same interaction region as CMS (IP5). This common physics programme is presently under discussion within the 2 Collaborations and is investigated by CMS/TOTEM working groups.

The TOTEM experiment [198, 211] is an approved experiment which will measure the pp elastic cross section as a function of t, the square of the exchanged four-momentum, the total cross section with a precision of approximately 1%, and diffractive dissociation at $\sqrt{s}$ = 14 TeV. The TOTEM experimental set-up consists of 2 tracking telescopes T1 and T2, as well as Roman Pot (RP) stations, 1 on either side of IP5. The T1 and T2 telescopes consist of CSC (Cathode Strip Chambers) and GEM (Gas Electron Multipliers) chambers respectively, and will detect charged particles in the η regions $3.2 < |\eta| < 5$ and $5 < |\eta| < 6.6$.

The TOTEM RP stations will be placed at a distance of ± 147 m and ± 220 m from IP5 (Fig.7.2). Each station will consist of 2 units, 2.5 m and 4 m apart, each with 1 horizontally and 2 vertically movable pots equipped with silicon strip detectors. These stations can measure protons with a momentum loss $\xi = \Delta p/p$ in the range $0.02 < \xi < 0.2$ for the nominal collision optics. For other optics with larger β^*, and hence lower luminosity, much smaller values of ξ can be reached.

Detectors at a distance of 420 m, in the cryogenic region of the LHC ring, are currently being considered by the FP420 project [212]. They would provide a coverage of $0.002 < \xi < 0.02$, complementary to that of the TOTEM detectors, for the high luminosity, low $\beta*$ optics.

The TOTEM detectors can provide input data to the Global Trigger of the CMS Level-1 trigger. Track finding in T1 and T2 (combined coverage $3.2 < |\eta| < 6.6$) for triggering purposes is optimized to select beam-beam events with charged tracks that point back to the IP and reject beam-gas and beam-halo events which have tracks that do not.

The Roman Pot detectors of TOTEM aim to detect the protons in diffractive interactions of the type $pp \rightarrow p + X$ and $pp \rightarrow p + X + p$. When used in conjunction with the central CMS detector interesting phenomena such as hard diffractive scattering can be studied, where the

system X can consist of jets, W, Z bosons, high E_{T} photons, top quark pairs or even the Higgs particle, as discussed recently in [213, 214].

The combination of T2 and CASTOR will allow the study of phenomena at lower Bjorken-x than otherwise reachable. Drell-Yan measurements will enable the parton distributions to be probed down to $x \approx 10^{-6}$. The energy and particle flows in the forward region are also of prime interest for tuning Monte Carlo simulation programs used in cosmic ray studies. CASTOR is designed especially to hunt for "strangelets" in AA collisions, which are characterized by very atypical fluctuations in hadronic showers.

The prime goal of the ZDC is to measure the centrality in AA collisions. The so-called ultra-peripheral events can also be tagged. In pp interactions it will allow the study of events with charge exchange and consequently a forward high energy neutron. Its ability to see low energy (≈ 50 GeV) photons is important for exclusive diffractive studies.

The full physics program with these detectors will be outlined in a forthcoming Report. In the present document, we will only discuss the forward detectors of CMS, and refer the reader to the TOTEM TDR for the detailed description of the detectors of the TOTEM experiment.

7.2 The Castor Calorimeter

7.2.1 Description of the calorimeter

7.2.1.1 Mechanical sesign

CASTOR is an electromagnetic/hadronic calorimeter, azimuthally symmetric around the beam and divided into 16 sectors ($\phi = 22.5°$). It is also longitudinally segmented into 12 slices, the so called Reading Units (RU), in order to observe and measure the propagation of hadronic cascades (showers) along its depth. This characteristic is particularly necessary for detecting penetrating cascade particles.

The calorimeter is a Cherenkov-light device, consisting of successive layers of tungsten plates, as absorber, and fused silica (quartz) plates, as active medium. The plates are inclined at $45°$ with respect to the impinging particles in order to maximize the Cherenkov light collected. The light is produced by the passage of the charged particles in the shower (primarily e^+, e^-) through the quartz. The light reaching the top of the quartz plates by internal reflection is collected in the RUs along the depth of the calorimeter and transmitted by aircore lightguides onto the radiation-hard photomultiplier tubes (PMTs).

The Castor calorimeter comprises 2 electromagnetic samplings with 5 mm W plates sandwiching 2 mm quartz plates and a hadronic part that has up to 10 samplings with 10 mm W plates sandwiching 4 mm quartz plates. The total depth of the EM section is about 22 radiation lengths (X_0) and the total depth is 10.3 interaction lengths (λ_{I}). It is situated in the collar shielding at the very forward region of CMS, starting at 1437 cm from the interaction point, as shown in Figure 7.3.

The pseudorapidity range covered is $5.3 < |\eta| < 6.5$ for the EM-section and $5.15 < |\eta| < 6.4$ for the hadronic section. This η-coverage closes hermetically the CMS pseudorapidity range over almost 13 units.

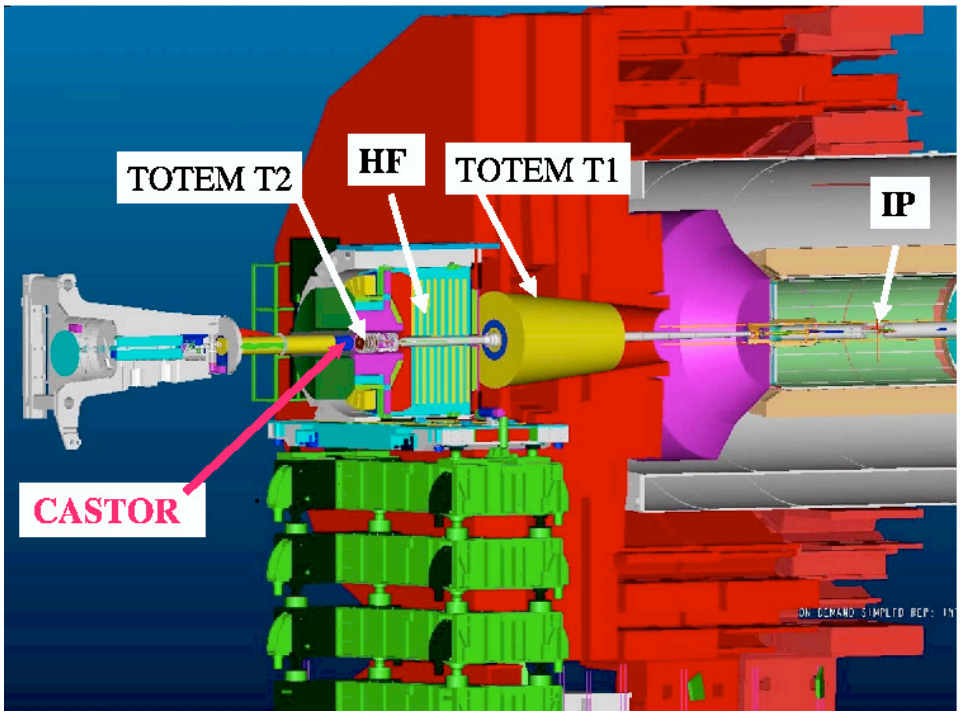


Figure 7.3: Schematics of the CMS forward region.

The current plan is to build the calorimeter in 2 stages:

- Stage I is for pp physics and will consist of the EM and (6) RUs of the hadronic section, leading to a total of $\approx 6.6\lambda_{\mathrm{I}}$. The aim is to complete and install stage I for the pilot run or at the latest by the first physics run.
- Stage II is for heavy-ion physics and will have the additional (4) RUs of the hadronic section. The stage should be completed in 2008-09.

Figures 7.4 and 7.5 show views of the mechanical design of the calorimeter.

7.2.1.2 Calibration and simulation

The 16×12 RUs of the calorimeter are calibrated using the blue-light laser from the HF calorimeter. The light reaches the top of the light guide of each RU via quartz fiber and shines light onto the top of the quartz plates. The reflected light is read by the PMT. For monitoring radiation damage of the quartz plates, red light will also be used in comparison, since it remains unaffected.

The CASTOR calorimeter is implemented in the GEANT4 simulation program of CMS. The calorimeter geometry, as used in the Monte Carlo simulations for this report, is however somewhat different from the one described above. Its depth is of 10.2 λ_{I} and is divided into 18 RUs. All other characteristics are the same. Figure 7.6 shows a general view of the CASTOR geometry, implemented in the simulation and visualized with the IGUANA visualization package (Section 2.10).

7.2.2 Beam test results

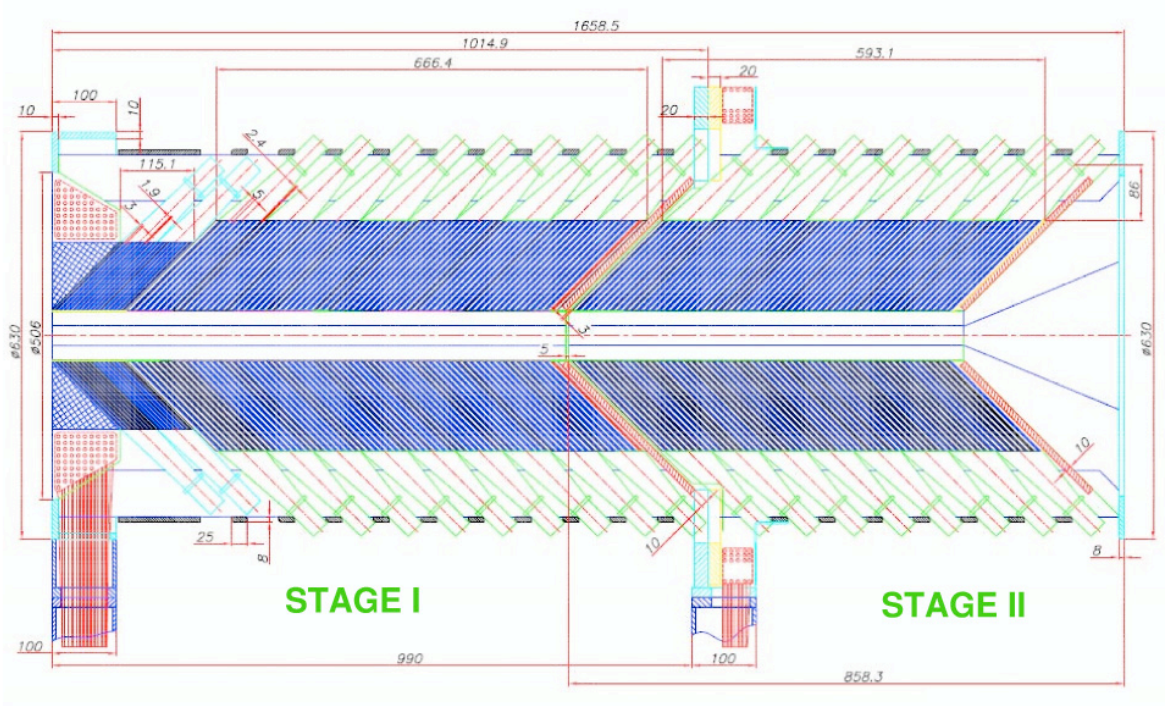


Figure 7.4: View of the mechanical design of the calorimeter showing the 2 stages.

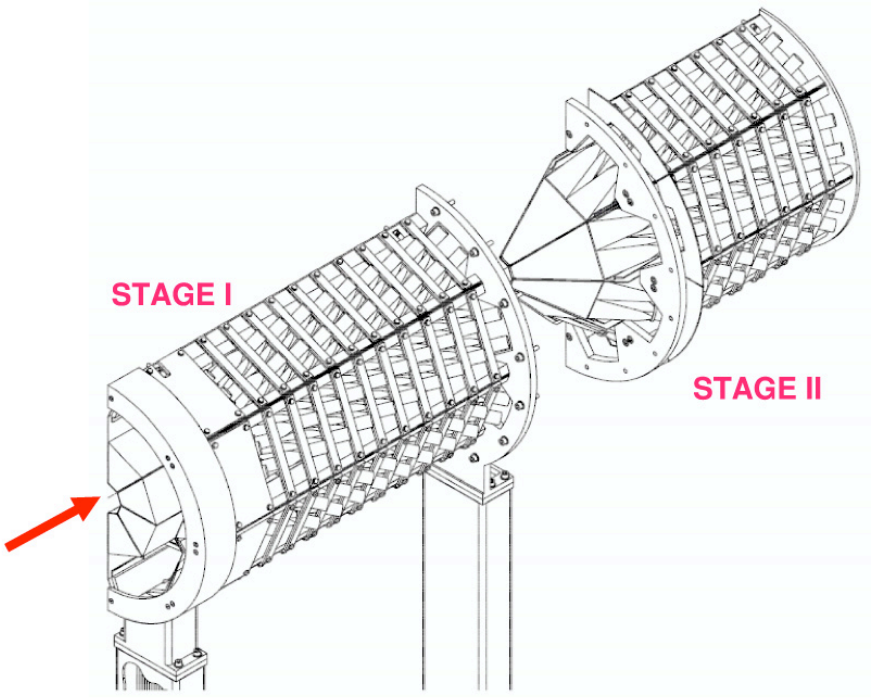


Figure 7.5: View of the mechanical design of the calorimeter showing the 2 stages.

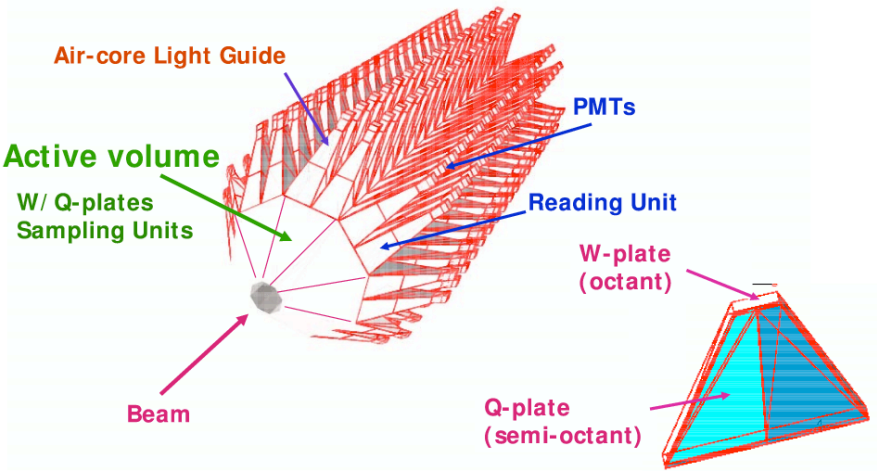


Figure 7.6: General view of CASTOR as visualized with IGUANA.

7.2.2.1 Prototype I

The first prototype was an electromagnetic calorimeter with 4 octant-shaped sectors (Fig. 7.7). It was built to study several aspects of the design: the type of quartz medium (fibres vs. plate), the reflecting medium in the air-core light guide (glass mirror vs multilayer stack), the reading device (avalanche photodiodes, APDs, vs. PMTs). In addition, we studied the energy response and the energy resolution for these different options.

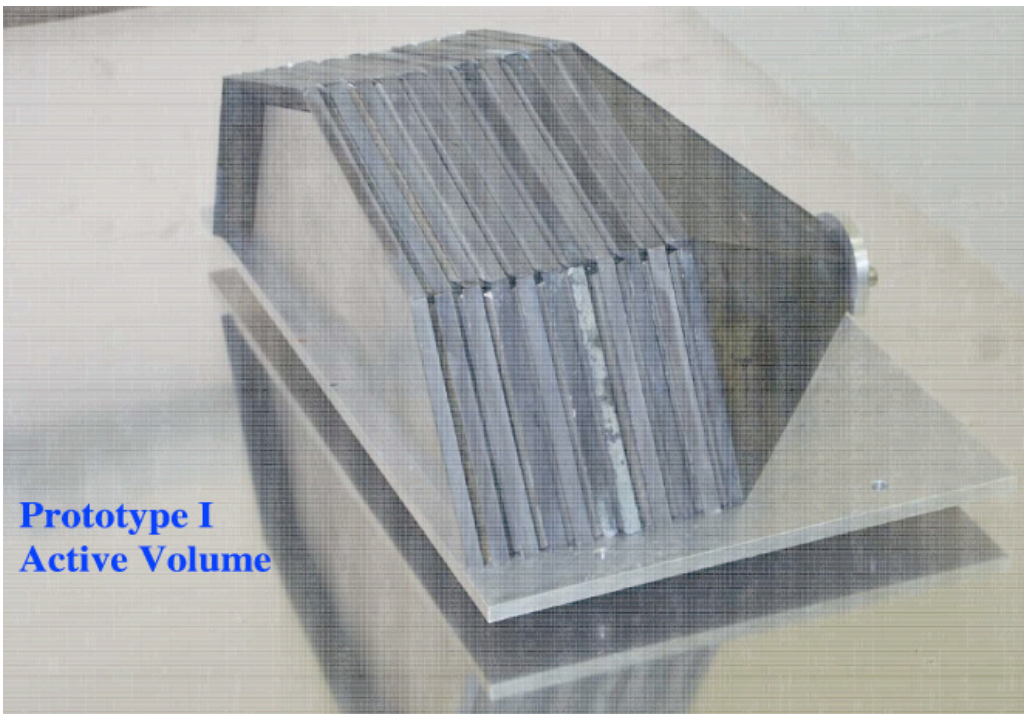


Figure 7.7: Photo of the active volume (tungsten and quartz media) of CASTOR Prototype I.

Each of the 4 azimuthal sectors (octants) consisted of 10 pairs of tungsten/quartz planes, totaling 23.7 radiation lengths, or 0.83 interaction lengths.

7.2.2.2 Beam test of Prototype I

The beam tests of prototype I took place in June 2003 at the beam line H4 of the SPS accelerator at CERN. The prototype was positioned on a moving table with horizontal (left-right) and vertical degrees of freedom with respect to the electron beam. Electron beams of several energies were used. A pair of wire chambers, positioned upstream, defined the incident position of the beam particles.

Several combinations were tested involving quartz plates or quartz fibres, two different reflectors in the light guides, and different types of light-reading devices . These are illustrated in Figure 7.8.

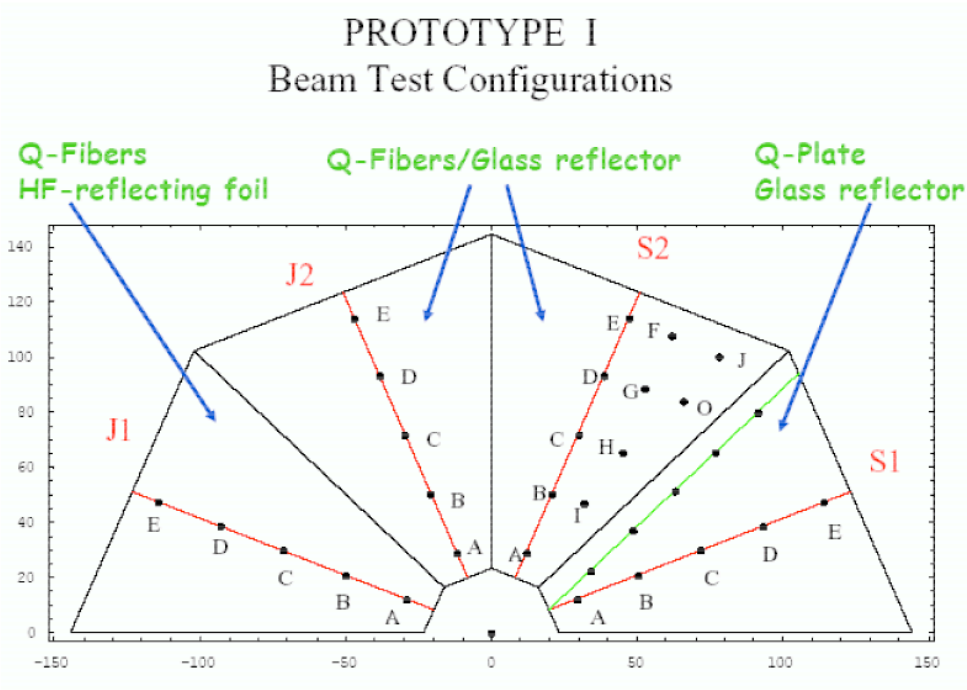


Figure 7.8: Different options investigated in the beam test of Prototype I.

7.2.2.2.1 Energy response To study the linearity of response with energy and the relative energy resolution, electrons were fired at the central points C (Fig. 7.8) in the different azimuthal sectors. Electron beams of energy 20, 40, 80, 100, 150 and 200 GeV (Fig. 7.9) were used. The distributions of the signal amplitudes are well fitted by a Gaussian function.

7.2.2.2.2 Energy resolution For the relative energy resolution of the calorimeter, the ratio σ/E vs. E was fitted to 2 different expressions [215, 216]:

$$\frac{\sigma}{E} = p_0 + \frac{p_1}{\sqrt{E}} \tag{7.1}$$

$$\frac{\sigma}{E} = p_0 \oplus \frac{p_1}{\sqrt{E}} \oplus \frac{p_2}{E} \tag{7.2}$$

where $\oplus$ denotes sum in quadrature.

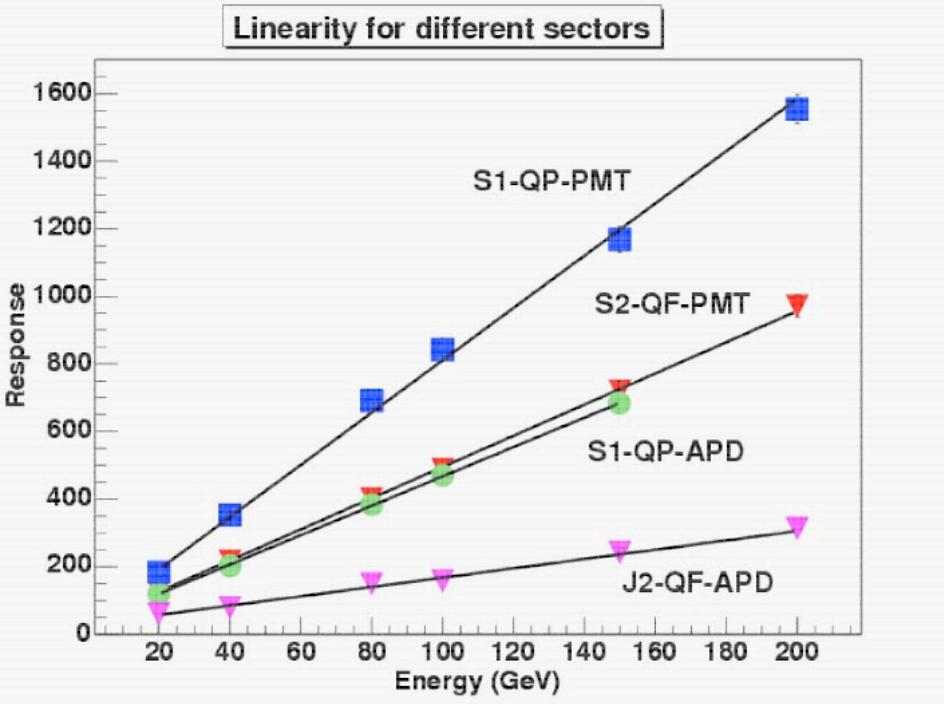


Figure 7.9: Response Linearity in sectors: S1-(Philips PMT or APD2), S2-(Philips PMT), J2-(APD1). Calorimeter response in ADC channels vs. electron energy in GeV.

The energy resolution for the various options studied is shown in Figure 7.10. Generally, both formulae fit the data satisfactorily. The constant term p_0 is close to zero for all options. Also the stochastic term p_1, being less than 38% for the S1 sector, seems to be reasonable, when compared to $p_1 = (36.2 \pm 0.2)\%$, obtained in [217] for a calorimeter prototype of similar geometry and technology. The readout by photodiodes requires the p_2 term in the fit due to the sizable electronics noise, measured to be 1.25 GeV and 4.5 GeV for Advanced Photonic APD and Hamamatsu APD, respectively. It should be noted that the APDs are very sensitive to both voltage and temperature changes, but in this test there was no such stabilization.

7.2.2.2.3 Uniformity of response The uniformity of the calorimeter response was studied using 80 GeV electrons incident of sectors J1, J2 and S1 with Hamamatsu PMT readout. Also studied was energy sharing between neighboring sectors. The response at points A-E in the middle of the sectors and at the border of S1 was also studied. The energy resolution for these sectors is compared in Figure 7.11. The light output is highest in the S1 sector, and it is practically the same for the central and for the border points. It depends weakly on the position of the impact point: for S1 (J1, J2) a weak decrease (increase) of the calorimeter response with distance R from the calorimeter center is observed.

7.2.2.3 Prototype II

The results of the beam test analysis of prototype I prompted the construction of a second prototype using quartz plates, large area APDs as well as PMTs, and Dupont reflective foil. A new geometry of the semi-octant was tested.

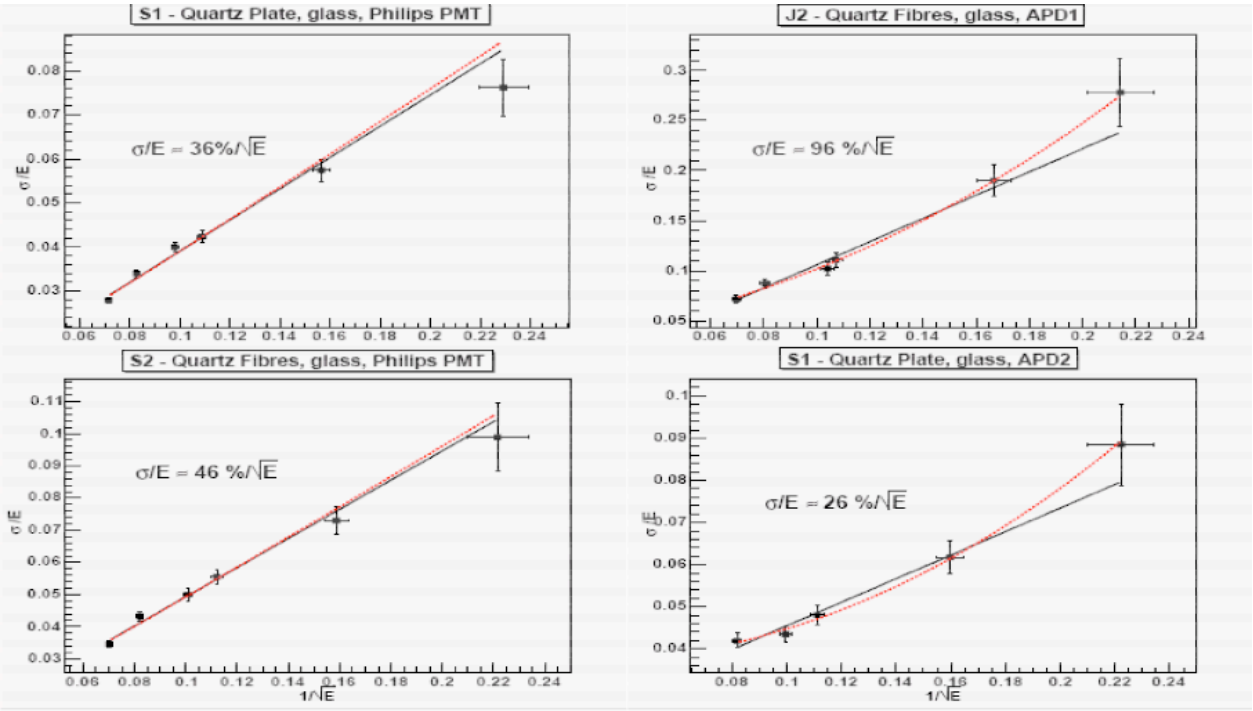


Figure 7.10: Energy resolution in different sectors. Two fits are shown: $\sigma/E = p_0 + p_1/\sqrt{E}$ (black line); $\sigma/E = p_0 \oplus p_1/\sqrt{E} \oplus p_2/E$ (red dotted line).

Prototype II comprised electromagnetic and hadronic sections, with 1 octant equipped. The EM section was further divided into 2 semi-octant sectors and was longitudinally segmented into 2 samplings. The hadronic section was longitudinally segmented into 4 samplings. The calorimeter was constructed from layers of tungsten plates as absorber and fused silica quartz plates as active medium. The light was collected in sections along the length of the calorimeter and focused by air-core light guides onto APDs and PMTs. In the EM section, the tungsten/quartz plate thickness was 3 mm/1.5 mm and each of the 2 samplings had a thickness of 13.4 X_0 (0.536 λ_I). In the hadronic section, the tungsten/quartz plate thickness was 5 mm/2 mm and each sampling has a thickness of 0.796 λ_I. The plates are inclined at 45 degrees.

The photodetectors were arranged in a matrix of 4 or 6 APDs and PMTs. Figure 7.12 shows a picture of the fully assembled prototype II positioned on the moving table at the H2 beam line of the SPS at CERN.

7.2.2.4 Beam test of Prototype II

The beam test of prototype II took place in October 2004 at the H2 beam line of the SPS at CERN. Electron, hadron (π^-) and muon beams at several energies were used. A telescope of scintillator detectors and wire chambers were upstream of the prototype, giving precise information on the position of each particle hitting the calorimeter.

Electron beams of energy 20–200 GeV were used in the test. Together with electrons, muons were also detected, as shown in a typical spectrum, Figure 7.13. The electron peak is almost

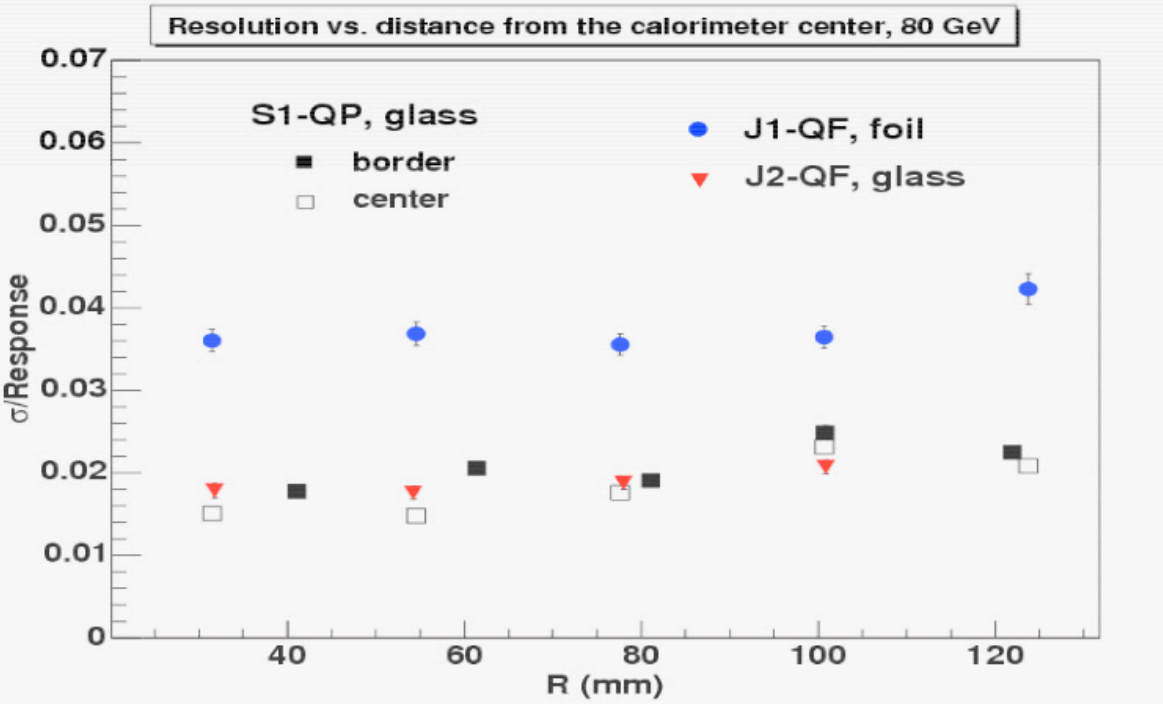


Figure 7.11: Resolution vs. distance from the calorimeter center for 80 GeV electrons in several points (AE) of J2, J1 and S1 sectors, all connected to Hamamatsu PMTs.



Figure 7.12: Assembled prototype II on the moving table in the H2 beam line.

always well fitted with a Gaussian curve.

7.2.2.4.1 Linearity of response Figure 7.14 shows the linearity of the energy response using an electron beam for different photo-devices: groups of 4 and 6 APDs and the Hamamatsu PMT.

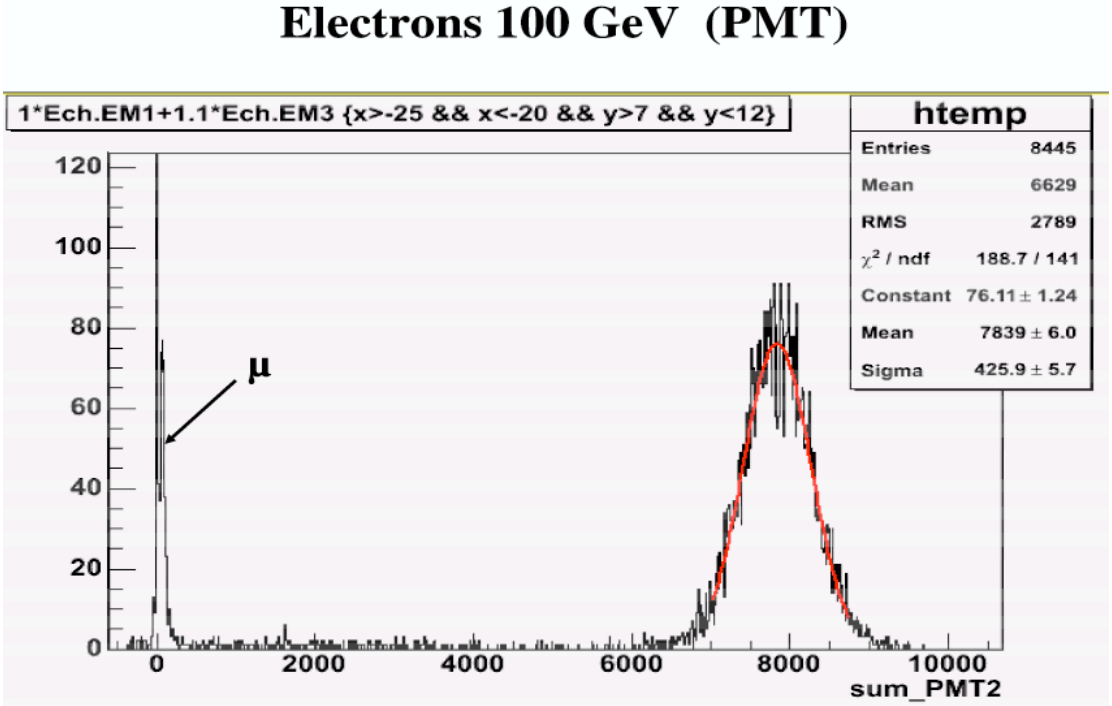


Figure 7.13: Typical energy spectrum with 100 GeV electron beam.

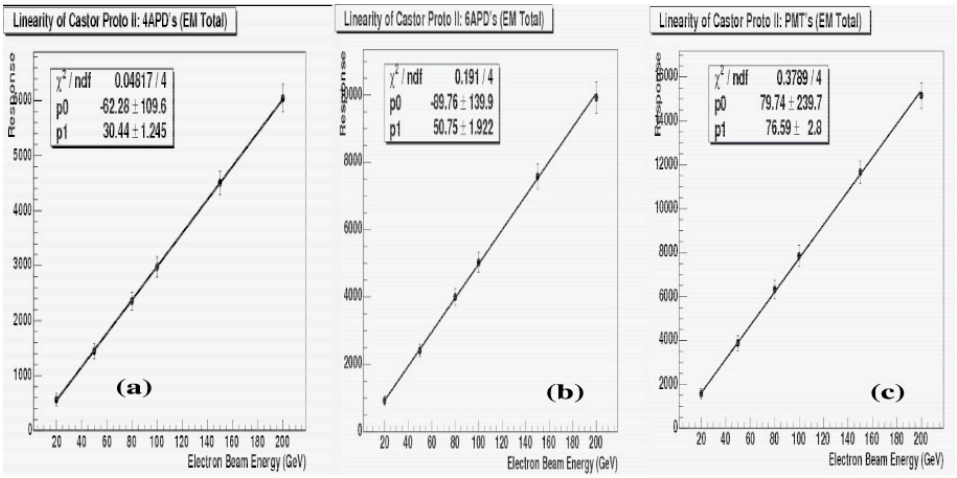


Figure 7.14: Linearity of the energy response with electrons at several energies and for different readout devices.

7.2.2.4.2 Energy resolution The energy resolution was obtained for the 3 readout devices – 4 and 6 APDs, and the PMT. Figure 7.15 shows the comparison of the resolution for 4 APDs and the PMT.

7.2.2.4.3 Uniformity of response and position resolution The response for different points of impact of electrons was studied using an 100 GeV electron beam and 4 APDs as photodetectors. Figure 7.16a shows the result of this scan and the characteristic energy response of the 2 channels, as the beam moves from one to the other. The derivative of the

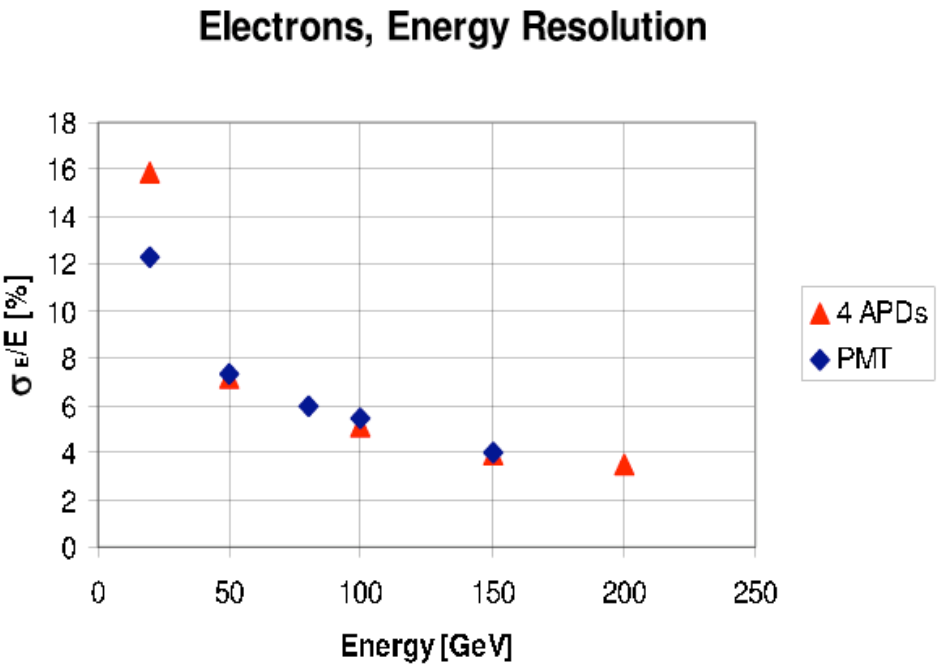


Figure 7.15: Comparison of the energy resolution with the PMT and 4 APDs as the readout.

sigmoid response curve gives the width of the EM shower transverse size (Fig.7.16b). We obtain a FWHM of ≈2 mm, in agreement with Monte Carlo simulations. This means that 95% of the EM energy is contained within a cylinder of a radius of approximately 5 mm.

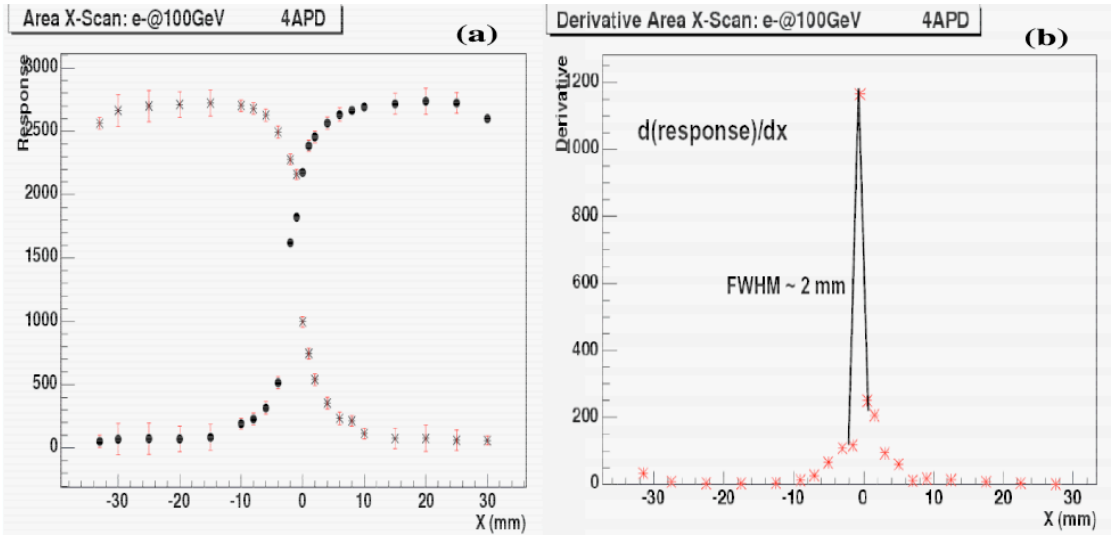


Figure 7.16: Spatial response (a) and resolution (b) of the EM section for 100 GeV electrons.

7.3 The Zero Degree Calorimeter

7.3.1 Description of the calorimeter

The ZDCs need to be compact, fast, highly radiation resistant and with good energy and time resolution. To accomplish this it is planned to use tungsten as an absorber and collect the signal from Cherenkov light emitted by relativistic charged particles in quartz fibres. This is the basis of the hadronic forward calorimeter in CMS. A similar design has proved very robust at RHIC.

The design requirements for the ZDC are:

1. width < 9.6 cm, length < 100 cm;
2. energy resolution sufficient to resolve the 1 neutron peak;
3. very high radiation tolerance;
4. low sensitivity to induced radioactivity;
5. rate capability above 50 kHz (for Ar-Ar);
6. vertex resolution through timing of few cm, i.e. $\sigma_{\rm t} \approx 100$ ps.

The zero degree calorimeters for CMS will follow closely the successful RHIC design [218, 219]. Figure 7.17 shows the design of the RHIC ZDC modules. These are Cherenkov calorimeters which sample only the high energy core of the induced showers. The use of tungsten provides sufficient density to contain the transverse spread of the signal to within a few cm of the initial trajectory. Quartz fibres provide high radiation tolerance. To accommodate the higher energy of LHC, we plan to increase the depth of the calorimeter from 6 to 8 interaction lengths. Cherenkov calorimeters are almost immune to induced radiation since they are only sensitive to high velocity charged particles. Since the typical energies will be 25 times larger at the LHC than at RHIC, a large amount of Cherenkov light is produced. Therefore the calorimeter can be made more compact, by using 1 cm or 1.5 cm plates instead of the 0.5 cm plates used at RHIC and still have adequate photostatistics. The proposed layout of the detector is shown in Figure 7.18.

A GEANT4 based simulation of the ZDC is available. Figure 7.19 shows a simulation of the ZDC response to 1 TeV neutrons. The peak has a width of $\approx 12\%$. Figure 7.20 shows the shower simulation of a single neutron with energy of a 1 TeV in the ZDC.

The location and configuration of the zero degree calorimeters for CMS follows closely the successful RHIC design [218, 219] adapted to accommodate the different accelerator geometry. To measure forward going neutrons the calorimeter needs to be located at the end of the straight section surrounding the interaction point between the 2 pipes containing the counter-circulating beams. In CMS this occurs 140 m from the vertex at the so-called "pair of pants." This area is inside the Neutral Beam Absorber device (TAN) which shields the superconducting magnets from synchrotron radiation produced by the beams. Figure 7.21 shows a picture of the TAN.

The TAN surrounds the beam splitting area with copper absorbers. To allow measurement of the forward going energy, the design of the TAN includes a detector cavity precisely at zero degrees. During pp runs part of the cavity will be filled with a shower maximum gas detector for measuring luminosity and to help with tuning the accelerator [220]. We propose to use the vacant space within the TAN for the ZDC.

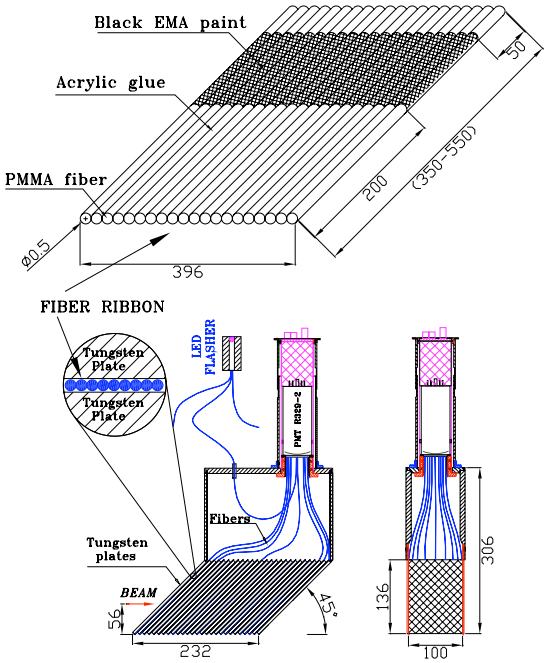


Figure 7.17: Mechanical design of the RHIC ZDC modules. Dimensions shown are in mm. We propose to use 3 such modules for the hadronic section and 1 smaller module with vertical plates for the electromagnetic section.

The detector cavity inside the TAN is 9.6 cm wide and 100 cm long. The strongest constraint on the design of the ZDC is that it must fit inside the TAN. A large space is available above the beam center line but only 6.7 cm is available below. The short length of the cavity precludes our copying the ALICE design [221] which has a total length of 140 cm. Therefore we propose to follow the design of the RHIC ZDCs [219] and read out the calorimeter from the top. The geometry of the TAN precludes the measurement of the spectator protons since they will be bent into the absorber. However our experience at RHIC shows that measurements of the spectator neutrons and charged particles at mid-rapidity will allow us to measure the number of participant nucleons, $N_{\rm part}$, to within 5% for $N_{\rm part} \approx 80$, dropping to 3% for the most central collisions [222].

As mentioned above, the TAN will also house an instantaneous pp luminosity monitor, similar to that used at HERA. The luminosity monitor will have a depth of 10 cm along the z axis and will need to have an absorber in front of it to convert high energy photons into electrons. This absorber will be the electromagnetic section of the ZDC. With a depth of 10 cm it should be possible to make this 22 radiation lengths deep. The remaining 80 cm behind the ZDC will be used for the hadronic section. This will consist of 3 modules each 1.6 interaction lengths thick. The tungsten plates will be tilted at 45^o in order to equalize, as best as

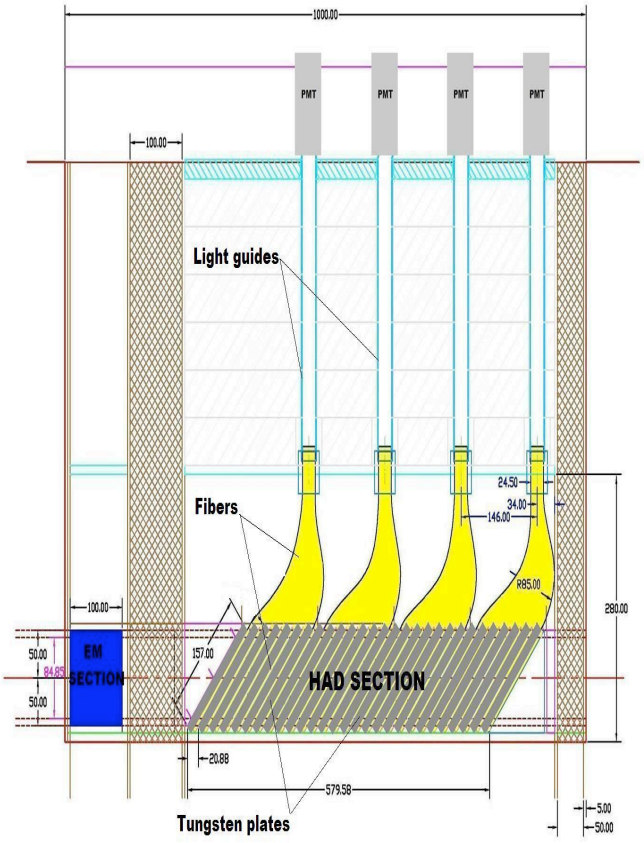


Figure 7.18: Side view of the ZDC showing the electromagnetic and hadronic sections with the luminosity monitor in between.

possible, the response of the detector to electromagnetic and hadronic energy. For the electromagnetic section the plates will be oriented vertically. The pp monitor is not expected to significantly degrade the energy resolution of the ZDC and may well enhance it by serving as a shower maximum detector. The ZDC and luminosity detectors are now included in the CMS software framework.

There has been considerable interest within the HCAL group on using the ZDCs and CASTOR to improve the missing energy resolution of CMS. For this reason we have decided to use the ZDCs for the first 2 low luminosity pp runs. This will require the use of quartz/quartz fibres.

The radiation load of 300 kGy/year of Pb-Pb runs for the ZDCs is about 1/3 of that expected for the forward hadron calorimeter, HF, in the region $4.5 < \eta < 5.0$. During the pp runs the power dissipated in the TAN will reach a maximum of 10 W/Kg, i.e. 10 Gy/s [220]. Thus in a year of pp run we would expect a dose of about 180 MGy or 18 GRad. Hence it should be possible to use the same fibres as used in the HF. These are so called "Quartz-Plastic" or QP fibres which have a quartz core and a plastic cladding. However one of the advantages of CMS is the ability of the DAQ to take the full luminosity for pA and light ion collisions. For these systems the radiation load could be much higher and so it seems prudent to use

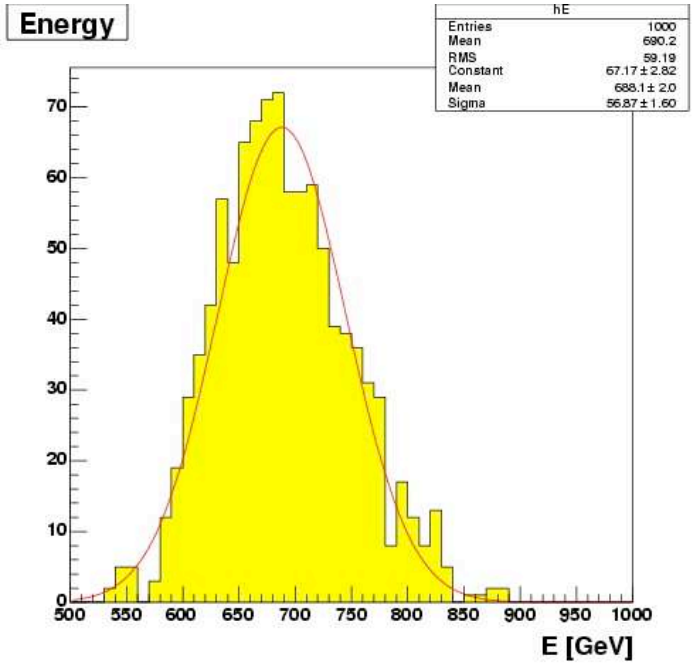


Figure 7.19: Monte Carlo simulation of the ZDC response to 1 TeV neutrons.

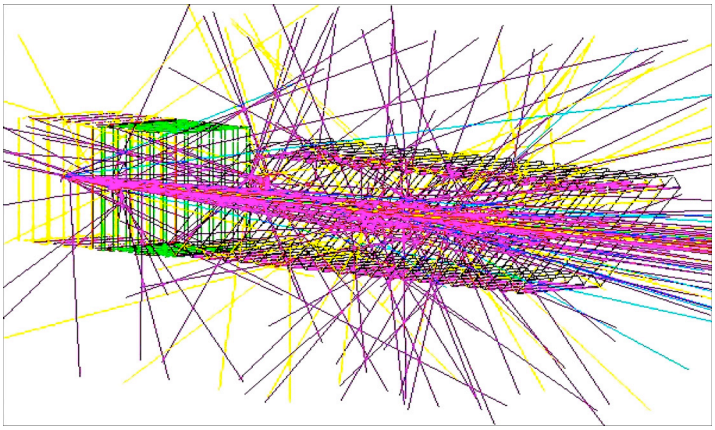


Figure 7.20: Visualization of the ZDC response to one 1 TeV neutron shower.

fibres that have the maximum possible radiation hardness. This would imply using the more expensive "Quartz-Quartz" fibres.

To minimize risks we plan to use the 8 stage Hamamatsu R7525HA, the same tubes as for the HF. These have been studied very extensively in high radiation environments. They are optimized to match with the CMS readout system. A large dynamic range is accomplished through a multi-range technique.

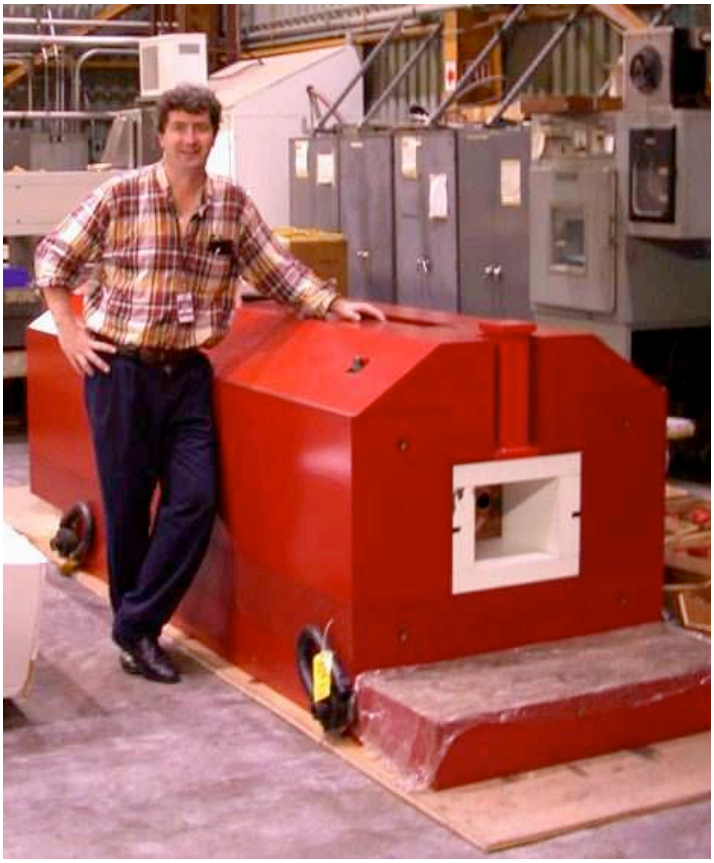


Figure 7.21: Close up of the TAN showing the cavity which will hold the ZDC.

7.3.2 Calibration

The best way to calibrate the ZDCs is by observation of the 1 and 2 neutron peaks from the heavy ion physics runs. Figure 7.22 shows these peaks in a spectrum of $AuAu$ collisions at RHIC [218]. At RHIC the stability of the system is also monitored by using LEDs to inject a known amount of light into the ZDC photomultipliers. It should be possible to get some information about the radiation damage to the fibres by injecting light into the bottom of the fibres themselves.

7.3.3 Trigger and DAQ

The plan is to copy the HF readout chain, but the high radiation levels and the large dynamic range required cause some complications. To avoid damaging the electronics we will transfer the signal along fast hard core cables to the counting room. There the signal will be split 3-ways for timing and low and high gain energy measurements. The ZDCs can then provide the following information to the trigger system:

1. Vertex Position: $c(T_{\mathrm{left}} - T_{\mathrm{right}})$;
2. Timing: $(T_{\mathrm{left}} + T_{\mathrm{right}})/2$;
3. Hadronic Interaction Trigger: $(E_{\mathrm{left}} \approx E_{\mathrm{right}})$;

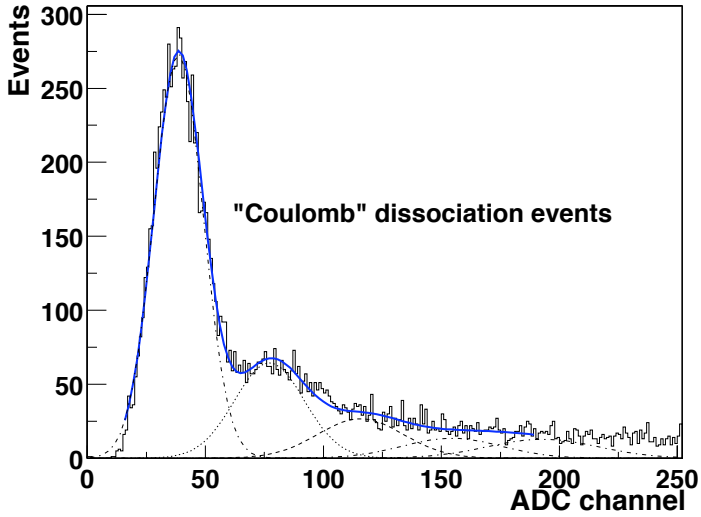


Figure 7.22: ZDC spectrum for electromagnetic Au-Au collisions with $\sqrt{S_{nn}}$=130 GeV.

4. Electromagnetic Interaction Trigger: $(E_{\rm left} \gg E_{\rm right})$ or $(E_{\rm right} \gg E_{\rm left})$.

Initial discussions [223] indicate that it should be possible to use the ZDCs at Level 1.

7.3.4 Heat dissipation

We might expect a heat load from radiation of $\approx$ 200 W. We plan to remove this heat by bonding the tungsten to copper side plates. Figure 7.23 shows results of a test in which heat was injected into an insulated steel plate, with similar conductivity as tungsten. The temperature was then measured as a function of time at 12 points around the plate. The only cooling for the plate was provided by 2 copper strips that extended outside the insulation into the air. From these tests we expect a maximum temperature rise of 20°C.

7.3.5 Performance

We expect the ZDC hadronic section to have a resolution of 10% for 2.7 TeV neutrons. The electromagnetic section is designed to reconstruct 50 GeV photons. At this energy the resolution is 11%. As the photon energy increases to 100 GeV, the shower penetrates deeper into the calorimeter causing some leakage into the luminosity monitor. Since the showers are very narrow and the electromagnetic section is divided into 5 sections horizontally we can get position resolution of ≈ 6 mm. This will be very useful for measuring the beam crossing angle. Finally, given the timing resolution of order 100 ps we can make vertex selection with a resolution of 3 cm. This can be done in the level 1 trigger.

7.4 Trigger for common CMS/TOTEM runs

The trigger is a key component in capitalizing on the physics potential provided by combining data from CMS and TOTEM[1]. The feasibility of a special forward-detectors trigger

[1]This section gives a brief overview of results from a common study performed by CMS and TOTEM

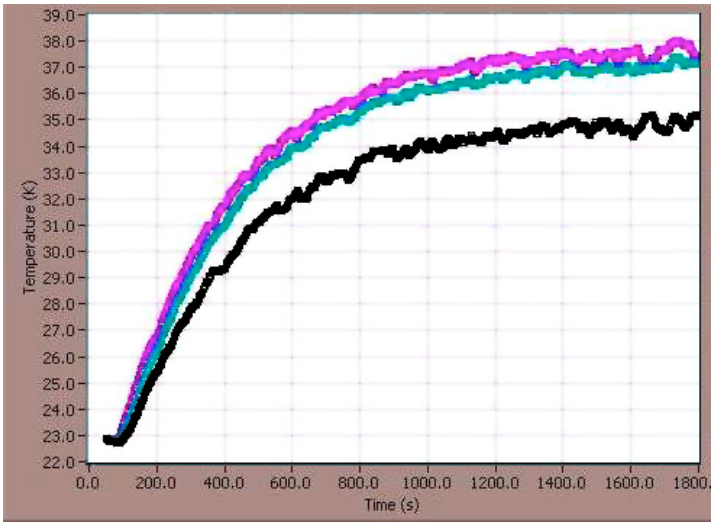


Figure 7.23: Temperature as a function of time for points at various distances along the x-axis of the heated plate. The plate is cooler near the copper strips at the side.

stream, with Level-1 and HLT output target rates of $\mathcal{O}(1)$ kHz and $\mathcal{O}(1)$ Hz respectively, will be discussed in detail in the Volume 2 of the Physics TDR, along with the potential of the already foreseen CMS Level-1 trigger streams. The aim is to be able to select events in which the interaction is mediated by the exchange of the vacuum quantum numbers – yielding large rapidity gaps in the final-state hadron activity as well as final-state protons carrying a large fraction of the incoming beam momentum.

The forward-detectors trigger stream combines the information of the central CMS detector with that from detectors further downstream of the CMS interaction point (IP), notably the TOTEM T1 and T2 tracker telescopes as well as the TOTEM Roman Pot detectors up to 220 m downstream of CMS [211]. Information from TOTEM will be available to the CMS Level-1 trigger. We also consider the detectors at a distance of 420 m, in the cryogenic region of the LHC ring, currently being studied by the FP420 project [212]. Additional information from detectors in the forward region will be available from CMS, namely CASTOR and the ZDC, discussed in the previous sections, but this information was not used here.

The advantages of a trigger stream that combines CMS and TOTEM quantities include:

- Thresholds can be used for the quantities measured by the central CMS detector that are substantially lower than in conditions based on CMS alone. This is crucial for retaining central exclusive and diffractive events in the Level-1 trigger, since these events typically have energy depositions below the standard Level-1 trigger thresholds.
- The TOTEM T1 and T2 tracker telescopes can be used to veto events with activity within their rapidity range; this is important to implement a large-rapidity gap selection in the trigger.
- At the HLT, a powerful selection condition for central exclusive events is provided by the kinematic correlation between the missing mass evaluated from the protons momenta and the mass measured in the central CMS detector.
- The TOTEM tracker telescopes T1 and T2 can also be used in a minimum bias

trigger. This trigger would be fully independent of the central CMS detector and hence would be ideal for cross-calibration purposes. In addition, since T1 and T2 provide track information, they could trigger events that depose too little transverse energy in the CMS calorimeters to pass the calorimeter noise thresholds.

Central exclusive production of a Higgs with mass close to the current exclusion limit is a prime example of a reaction that the forward trigger stream can retain: a Standard Model Higgs with 120 GeV/c^2 mass decays preferably into two b-quarks. The relatively low transverse momenta of the resulting 2 jets necessitate Level-1 jet E_{T} thresholds as low as 40 GeV. Thresholds that low would result in a Level-1 trigger rate of more than 50 kHz, due to the QCD background, essentially saturating the available output bandwidth.

The output rate of a dijet Level-1 trigger condition with thresholds of 40 GeV per jet can be kept at an acceptable level of order 1 kHz, in the absence of pile-up, by either using the TOTEM T1 and T2 detectors as vetoes or by requiring that a proton be seen in the TOTEM RP detectors at 220 m on one side of the IP. This gives a sufficient reduction of the QCD background event rate. At higher luminosities, up to $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, where pile-up is present, it is necessary to combine the single-sided 220 m condition with conditions based on event topology and on H_{T}, the scalar sum of all Level-1 jet E_{T} values; it is also necessary to impose a ξ cut at Level-1 — an option actively being pursued by the TOTEM collaboration at the time of writing. For a detailed discussion, the reader is referred to [224]. These Level-1 trigger conditions result in signal efficiencies of about 15%. This is a substantial increase with respect to the approximatley 10% efficiency achievable by selecting events with a muon in the final state (i.e. exploiting the relatively muon-rich final state of b decays), the only condition possible within the given bandwidth limits if using CMS alone.

Chapter 8

Luminosity Measurement

8.1 Introduction and overview

8.1.1 Relevant definitions

Instantaneous luminosity is defined in the usual way as

$$R = \mathcal{L}\sigma \tag{8.1}$$

where R is the observed event rate for a process with cross section σ. Integrated luminosity will be denoted by L. We define *interactions* to be inelastic proton-proton collisions (hard-core and diffractive scatterings) leading to significant energy depositions in the main parts of the CMS detector. When running at design luminosity, there will in be (≈ 25) interactions for each bunch crossing (BX).

8.1.2 Goals and requirements

The luminosity measurement is used to monitor the LHC's performance in real time and to provide an overall normalization for physics analyses. The design goal for the real time measurement is to determine the average luminosity with a 1% statistical accuracy in 0.1 s. For offline analyses, the design goal is a systematic accuracy of 5%, although every reasonable effort will be made to produce a more accurate result. Both of these requirements must be met over a very large range of luminosities, extending from roughly $10^{28}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ to $10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and possibly beyond.

In addition to providing average luminosity measurements in real time and integrated luminosity values for offline analyses, the luminosity system will produce bunch-by-bunch luminosities useful for accelerator diagnostics and potentially also for accurate modeling of underlying event backgrounds. [1]

Other important and desirable features of the luminosity system include a capability for "always on" operation and a bookkeeping system that is robust and easy to use. Always-on operation means that luminosity information should be available for real-time monitoring of the LHC, whether or not the main CMS DAQ is operational. The bookkeeping system should be such that corrections to the luminosity information can readily be made (for example if more accurate values are obtained through offline analysis). Moreover, users carrying out physics analyses should be able to determine the luminosity corresponding to their samples in a straightforward way.

[1] If the bunch-to-bunch variations in the luminosity are large, the underlying event backgrounds are no longer accurately described by simple Poisson statistics.

8.1.3 General strategy

The normalization of physics analyses will ultimately depend on careful measurements of known cross sections such as the pp total cross section or the production rates for Ws and Zs.

Since the precision pp total cross section measurement can only be done at $\mathcal{L} = 10^{28}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ (see Section 8.4.1 below) and requires a special configuration of the machine optics, it is not suitable for real-time monitoring. Measurement of the production rate for vector bosons can be done at all luminosities, but requires that all elements of the detector be operating and well understood. Moreover, the statistical accuracy of this method will fall well short of what is required to meet the 1% in 0.1 s goal.

The real time monitoring function must therefore be based on the measurement of high cross sections using comparatively simple hardware. These measurements will not only provide real time information during detector operations, but will also be used to "interpolate" between precision measurements to provide statistically precise luminosity information for small samples.

Multiple techniques capable of providing suitable luminosity information in real time have been identified. One technique employs signals from the forward hadron calorimeter (HF) while another, called the Pixel Luminosity Telescope (PLT), uses a set of purpose-built particle tracking telescopes based on single-crystal diamond pixel detectors. At this writing, the PLT has not been formally approved, but is under study. Yet another, which is being pursued by the LHC machine group, employs a Fast Ionization Calorimeters (FICs) positioned at ± 140 m from the IP, at the point of transition between a single and a combined beam pipe. The FICs measure the flux of 0° $\approx$TeV-scale neutrons that are produced in pp collisions.

The HF and PLT techniques, which will be an important focus of this chapter, provide complementary information, which will afford a redundant measure of the luminosity. Since it is based on simple counting, the PLT is expected to provided excellent linearity over a wide range of luminosities. It is, however, potentially subject to saturation at the highest luminosities. The HF presents more of a challenge in terms of linearity, but should continue to provide meaningful information at luminosities well beyond the LHC design luminosity.

In addition to these primary systems, we envisage a number of other luminosity measurements that can be derived from CMS components whose main function is something other than determining the luminosity. For example, the global calorimeter trigger may be capable of providing real-time luminosity information. Similarly, the main pixel detector can be used offline to count primary interaction vertices, which should provide a redundant cross check on other methods.

8.2 LHC beam parameters

The bunch crossing rate at the LHC is 40.08 MHz. An orbit of the machine comprises 3564 bunches, 2808 of which have real collisions when running at design luminosity. The time for a single orbit is 89.92 μs. An important secondary function of the luminosity system is to measure the bunch-by-bunch luminosity and to confirm that collisions are restricted to the proper bunches. Assuming an inelastic and diffractive interaction cross section of $\sigma = 80$ mb[2], and an effective bunch crossing rate of $f_{\mathrm{BX}} = (2808/3564) \times 40$ MHz, the

[2]If one includes elastic scattering, the total cross section is expected to be about 110 mb.

number of interactions per bunch crossing will be

$$\mu = \frac{\sigma \mathcal{L}}{f_{\mathrm{BX}}} \simeq 25 \tag{8.2}$$

at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Techniques for measuring the luminosity must therefore be able to contend with many overlapping interactions.

8.3 Online techniques

8.3.1 HF

8.3.1.1 Description

The forward hadron calorimeter, or HF, covers the pseudorapidity $3 < |\eta| < 5$ and is described in more detail in Chapter 5. It consists of quartz fibres embedded in a steel matrix. The signals from the HF result from Cherenkov light emitted in the fibres in response to charged particles. The fibres are viewed by phototubes situated on the downstream face. Since most of the visible energy is carried by relativistic particles—i.e., electrons—the calorimeter is most sensitive to the electromagnetic component of the hadronic showers. Thus the measured transverse shower size is determined by the Moliere radius rather than the nuclear interaction length.

Each HF endcap is divided into 36 segments in azimuth and 12 segments in η—i.e., each *physical tower* subtends an angular region of $\Delta\eta \times \Delta\phi \approx 0.175 \times 0.175$. In addition, crude longitudinal segmentation is achieved through the use of long fibres that run from the front face of the HF to the phototube readout at the back end and short fibres that cover only the rear part of the modules. The HF-based luminosity measurement to be described is based solely on the long fibres. For incident electrons, the light yield is 0.25 p.e./GeV.

Signals from the HF PMTs are digitized on a bunch-by-bunch basis using QIE chips, which employ a nonlinear scale to achieve a large dynamic range. The QIE outputs are routed to a set of 18 HCAL Trigger and Readout (HTR) boards, each of which services 48 HF physical towers.

8.3.1.2 Technique

8.3.1.2.1 Concept Two methods for extracting a real-time luminosity signal from the HF have been studied. The first is based on "zero counting." The number of interactions, n, in a given bunch crossing is distributed according to Poisson statistics

$$p(n;\mu) = \mu^n \frac{e^{-\mu}}{n!} \tag{8.3}$$

where μ is the mean number of interactions, given by Eq. 8.2 above. At very low luminosities μ is much less than 1 and is approximately equal to the fraction of bunch crossings that contain interactions. Determining this fraction is relatively straightforward using the HF, since nearly all interactions produce summed E_{T} signals well above the noise—see Fig. 8.1.

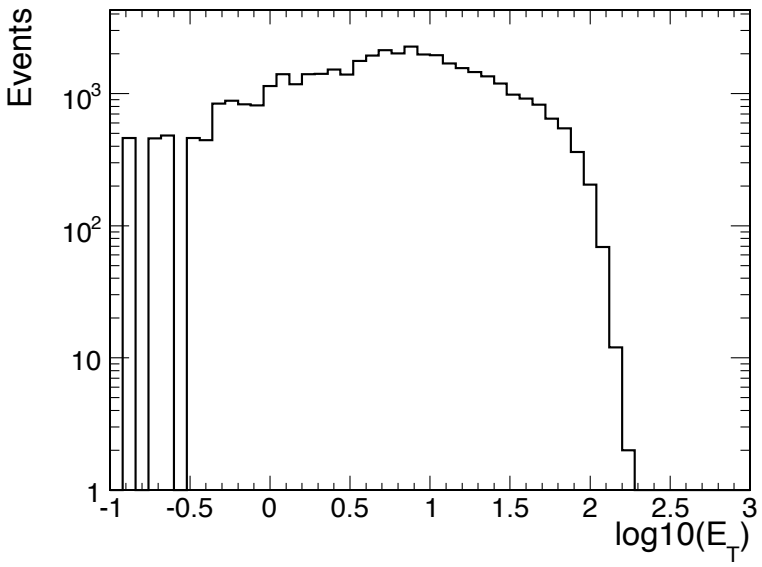


Figure 8.1: E_{T} spectrum of single interactions in the HF. Most events lie in the region $E_{\mathrm{T}} > 1$ GeV.

Even at a luminosity that is only 10% of design, the mean number of interactions per bunch crossing is ≈ 2.5, meaning that $\approx 70\%$ of bunch crossings have 2 or more interactions. Since one cannot reliably distinguish between single and multiple interactions, it is not possible to count interactions in a straightforward way. Fig. 8.1 does, however, show that in a large fraction of cases, one can distinguish between zero and one or more interactions. The mean number of interactions can thus be determined by inverting Eq. 8.3 to obtain

$$\mu = -\ln p(0), \tag{8.4}$$

which is known as "zero counting."

In principle, Eq. 8.4 can be used to infer μ at arbitrarily large luminosities, but for $\mu \gg 1$, the fraction of zeroes becomes very small, a condition here referred to as "zero starvation." We anticipate that well before a zero-counting method succumbs to a loss of statistical power associated with zero starvation, its systematic uncertainties will become unmanagable. We have thus adopted the rule of thumb that the fraction of zeroes should be at least 1%, corresponding to an upper limit on μ of 4.6. Thus, the total E_{T} signal from the HF will not provide a useful luminosity signal at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. If, however, a *single* $\Delta\phi \times \Delta\eta = 0.175 \times 0.175$ physical tower is used, even at design luminosity the mean number of hits above threshold ($E_{\mathrm{T}} \approx 0.1$ GeV) is roughly $\mu \approx 1$, and zero counting is a viable method. Since the occupancies of all HF physical towers are roughly the same, the statistical power of the technique can be increased by averaging the results from all such towers. In effect, one makes 864 (quasi) independent measurements each bunch crossing.

A second method exploits the linear relationship between the total E_{T} deposited in the HF and the number of interactions and thus the luminosity. Since the HF is very far forward, the maximum E_{T} is kinematically limited to a few hundred GeV, which explains the sharp upper edge in E_{T} visible in Fig. 8.1. This suppresses the large statistical fluctuations that can accompany unbounded power-law distributions.

8.3.1.2.2 Hardware implementation Since the HF and its readout is already part of the CMS baseline design, only a small amount of extra hardware is needed to derive a luminosity signal from it. In particular, a "mezzanine board," called the HF Luminosity Transmitter (HLX) card mounted in each of the 18 HTR boards is used to collect and reduce the HF QIE data for transmission via 100 Mbps ethernet to a router, which further concentrates the data to a single stream that is passed to a dedicated luminosity computer. This is shown in block diagram form in Fig. 8.2.

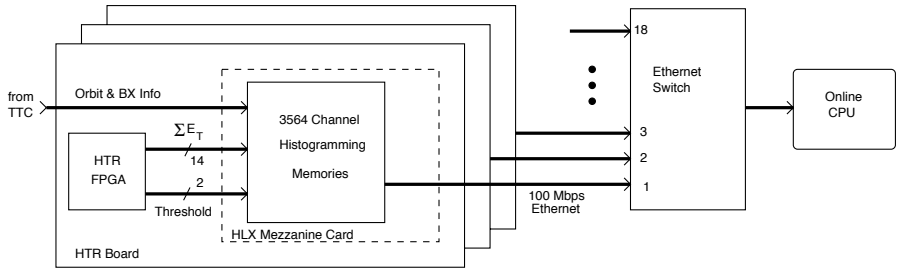


Figure 8.2: Block diagram of the HF luminosity readout.

The HLX boards have the same footprint and form factor as the Synchronization and Link Boards (SLBs) that provide an interface between the ECAL and HCAL readouts and the Regional Calorimeter Trigger system. Each HLX site is connected to the 2 main HTR FPGAs by 36 lines per FPGA. These lines can be driven at up to an 80 MHz cycle rate. At 80 MHz, 1 line must be reserved for framing, effectively leaving 70 bits of data available to describe the state of the 24 HF channels.

Two bits of status information encode the following four possible states for each tower: disabled, enabled-below-threshold, over-threshold-1, over-threshold-2. Inside the HLX, this input will be used to create 4 histograms. Each histogram will have 3564 bins, one for each bunch in the orbit. Each bin will be 2 bytes, so the total data size for a set of histograms is $28\,512$ bytes. An additional 14 bits will be allocated to a running E_{T} sum, which will be used to implement the second method described above.

The baseline design is to add the results from all 48 channels into a single set of histograms (effectively summing the number of zeroes across all channels in a module) the histograms would take at least $2^{16} \div 48 = 1365$, orbits or about 120 ms to overflow (in the worst case, 48 zeroes are accumulated during every bunch crossing). Leaving a safety margin, we will transmit the data to the router at 10 Hz.

8.3.1.3 Expected performance

The zero-counting method has been simulated using a sample of PYTHIA [45] minimum bias interactions. The response of the HF was simulated using OSCAR (Section 2.5), which is based on GEANT4 [30]. The output of OSCAR was analysed using a simple ORCA module that read hit data from the OSCAR output and produced an ntuple that was analyzed using a ROOT script.

Bunch crossings for a given luminosity are simulated by drawing a number of interactions according to Eq. (8.3). The energy depositions (expressed in photoelectron units) for the

next n interactions are then overlaid in an array where each element corresponds to 1 of the 864 physical towers in the HF. These "energy sums" are then smeared to take into account the statistics at the first dynode, scaled by the PMT gain, smeared again by the electronic noise at the QIE input, and quantized according to the QIE ADC calibration. In a final step, the energy values thus obtained are converted to E_{T} values. The $\sin\theta$ values used in this conversion range from 0.016 to 0.10 for $|\eta| = 5$ and $|\eta| = 3$, respectively.

The sum for each physical tower is then compared to a threshold of 6 ADC counts, which corresponds to 2 photoelectrons, or about 8 GeV of incident energy. This translates to an E_{T} ranging from 0.13 GeV for the innermost ring of the HF that is used to $E_{\mathrm{T}} = 0.75$ GeV for the outermost ring.

The zero fraction was computed using

$$p(0) = \frac{\#\mathrm{zeroes}}{\#\text{ of opportunities}} = \frac{\sum \#\text{ of towers with } E_{\mathrm{T}} < E_{\mathrm{thresh}}}{792 \times N_{\mathrm{BX}}} \tag{8.5}$$

where $2 \times 11 \times 36 = 792$ is the number of towers in the forward and backward HFs and the sum extends over all towers and all bunch crossing, except for the innermost and outmost rings. The value of $p(n)$ thus obtained is then converted to a mean tower occupancy (μ) using Eq. (8.4).

In the ideal case, the resulting value of μ should be proportional to the luminosity over a wide range. Fig. 8.3 shows that this is indeed the case. The upper panel shows the observed value of μ plotted as a function of input luminosity. The straight line plotted with the points is not a fit, but rather is chosen to go through the origin and the observed value of μ at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. In the lower panel, the same data are plotted, but in this case the points represent the fractional deviation from the average response.

A slight sub-linearity associated with zero starvation is observed at the highest luminosities. The seemingly simple expedient of raising the E_{T} threshold results in more zeroes, but introduces a superlinearity of different origin. In particular, for thresholds that are too high, the probability of 2 low-energy depositions conspiring to exceed the threshold competes with that for a single large deposition. Since the former mechanism scales with the luminosity squared, a nonlinearity results.

Additional runs of the Monte Carlo show that acceptable linearity is obtained with thresholds between 3 and 8 ADC counts. Although this is a workable range, and it is in principle possible to correct for nonlinearities, any such correction would be Monte Carlo dependent. It is thus desirable to employ alternative algorithms as a cross check. To that end, we have studied the use of a simple linear sum over scalar E_{T}.

Fourteen-bit E_{T} values from the HTR boards with a least count of 0.0625 GeV[3] are summed on a bunch-by-bunch basis. Since the scalar E_{T} deposited in the HF is proportional to the number of interactions, the average E_{T} per bunch is proportional to the luminosity. Fig. 8.4 shows that good linearity is achieved over the full range of luminosities.

8.3.2 Pixel Luminosity Telescope

[3]This least count value, which is 4 times smaller than that used elsewhere in the trigger, is required to avoid nonlinearities in the $\langle E_{\mathrm{T}}\rangle/\mathrm{BX}$ vs. luminosity function.

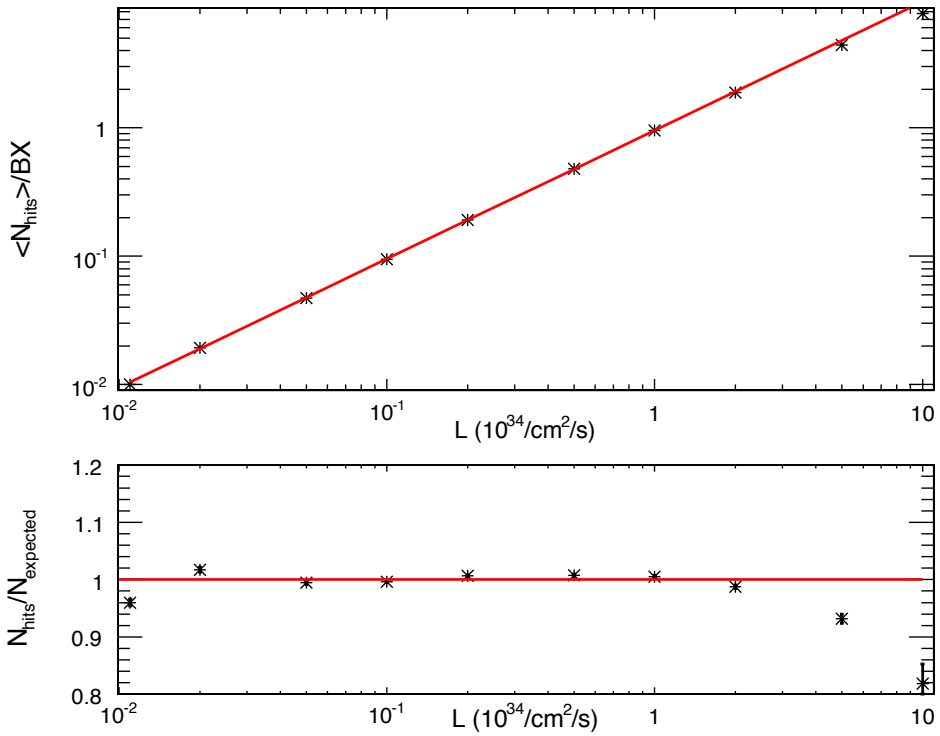


Figure 8.3: Linearity test of HF zero-counting technique. The upper panel shows the mean physical-tower occupancy inferred from counting the number of zeroes vs. the luminosity expressed in units of the LHC design luminosity of $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The lower panel shows the same data plotted relative to what is expected for a linear response. See text for additional details.

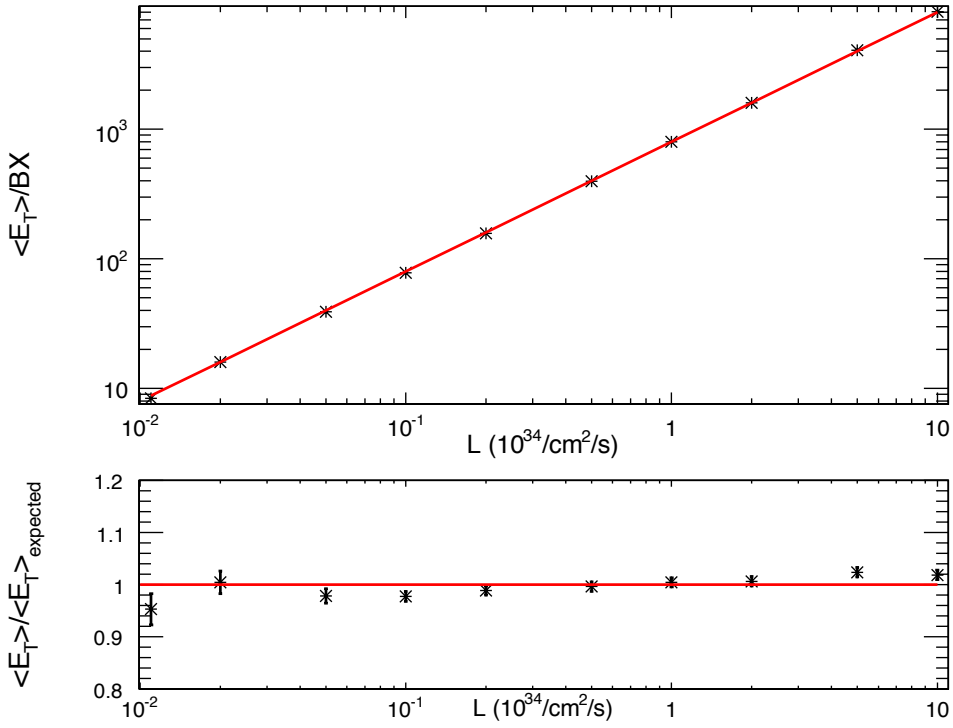


Figure 8.4: Linearity test of the HF E_{T} sum technique. The upper panel shows the average E_{T} per bunch crossing vs. the luminosity expressed in units of the LHC design luminosity of $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The lower panel shows the same data plotted relative to what is expected for a linear response. See text for additional details.

8.3.2.1 Description

The Pixel Luminosity Telescope (PLT) is a proposed, dedicated device for luminosity measurement that is currently under consideration by CMS. The PLT is designed to make fast, stable and precise determinations of the bunch-by-bunch luminosity and of the location of the interaction point centroid. It is comprised of arrays of small-angle telescopes located on either side of the interaction region. Each telescope consists of 3 planes of single crystal diamond sensors bump-bonded to a CMS pixel readout chip. For each bunch crossing, the number of particles traversing the telescope arrays is determined by forming a coincidence of the 3 planes in each telescope from fast out signals provided by the pixel readout chip.

The PLT has an array of 8 telescopes uniformly distributed in ϕ on each side of the interaction region. Each telescope consists of 3 equally-spaced sensor planes. The sensitive area of each plane, 8×8 mm^2, is determined by the active area of the pixel readout chip. In order to maximize the solid angle acceptance, the telescopes are placed as close to the interaction region as possible. They will be located in z between 1.65 m and 1.85 m with 10 cm spacing between sensor planes. Radially, they with be located 10 mm from the beam pipe with the sensor planes at $z = 1.65$ m located between $r = 41$ mm and $r = 49$ mm. The telescopes will be projective at an angle of 1.56° to the interaction point corresponding to a rapidity of $\eta = 4.3$. Figure 8.5 shows the location of a PLT array. The minimum z location is determined by the beam pipe support rods and collar at $z = 1.6$ m. The PLT will be located outside of this so that the telescopes can be installed and serviced without the necessity of removing the beampipe support.

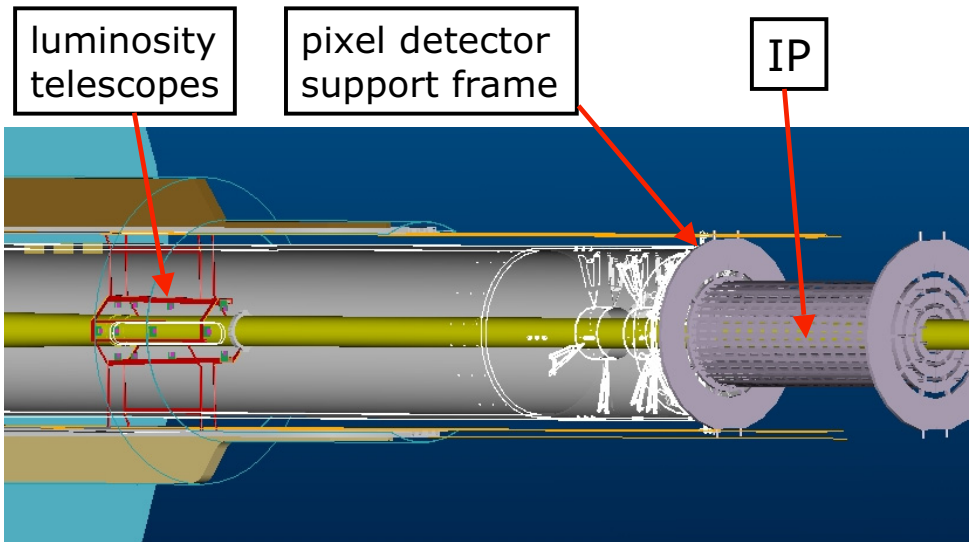


Figure 8.5: Location of a PLT array within CMS.

Each telescope plane consists of a sensor configured with a pixel pattern electrode and bump bonded to the CMS pixel readout chip. The reason for using pixels is to reduce the capacitance of each channel and, thereby, the inherent serial noise that otherwise would make readout of the sensor in the bunch crossing time of 25 ns prohibitive. With a pixellated sensor the per channel capacitance is reduced by several orders of magnitude and such that the

signal-to-noise ratio will be greater than 20:1 for 25 ns readout. The CMS readout chip [225] provides a fast out signal for each bunch crossing with an output level corresponding to the number of double columns with hit pixels. These signals, 1 for for each telescope plane, will be sent over standard CMS optical fibre analogue links to the electronics room where a three-fold coincidence for each telescope will be determined.

Counting of the fast three-fold coincidence signals constitutes the primary measurement of the bunch-by-bunch luminosity. In addition, since the sensors will be bonded to pixel readout chips, the full pixel information with column and row addresses and pulse height of the hit pixels can also be obtained. The full pixel readout of a group of several chips requires 1 to 2 μs and cannot be obtained for every bunch crossing but can be readout for every CMS L1A trigger, approximately a 30 kHz rate. This less frequent but more detailed information will allow full examination of the telescope hit patterns. Any anomalies seen in the fast coincidence signal data can be investigated and, with tracking information, the location of any "hot" sources of particles, e.g., beam halo impacting the beam pipe, can be determined. Furthermore, this tracking capability will allow the interaction point centroid to be determined. Since the entire readout and control chain from detector to the VME flash ADC module will be essentially identical to that for the CMS pixel detector, this capability can be obtained with modest additional cost or effort.

Single crystal Chemical Vapor Deposited (CVD) diamond [226] is used for the sensor material because of its superiority to silicon in terms of radiation hardness and its lack of need for cooling. Since they are located at a small angle, the charged particle fluence on the telescopes will be large, a few $\times 10^{14}$ cm^{-2} per year at full design luminosity, comparable to that of the inner layer of the barrel pixel detector. Measurements of CVD diamond[227, 228] indicate that the diamond signal amplitude decreases by at most 20% to 30% with no increase in leakage current at fluences of 2×10^{15} cm^{-2} equivalent to several years at full LHC luminosity. Another decisive advantage of diamond over silicon is the lack of a need to cool the diamond sensors. Unlike silicon, that must be cooled to -5° to -10° C in order to maintain operation at high radiation, diamond performs well and maintains negligible leakage current at room temperature and above.

Single crystal CVD diamond has a narrow Landau pulse height distribution that is well separated from zero as seen in figure 8.6 which shows the signal pulse height distribution measured for ^{90}Sr β particles incident on a single crystal diamond sensor of 480 μm thickness. The average pulse height for this diamond corresponds to a collected signal of $18\,100$ electrons, well above the noise level for pixel electronics. Because of the large separation from zero, the single crystal diamond sensors will have an essentially 100% detection efficiency. This is important in order to ensure that the efficiency is well determined and is stable as required for a high precision luminosity measurement. For the PLT, the diamond sensor thickness will be 400 μm yielding an average signal of $15\,100$ electrons for a normally incident minimum ionizing particle.

8.3.2.2 Expected performance

The primary function of the PLT is to provide a prompt and precise determination of the relative bunch-by-bunch luminosity. In addition to providing a real time measurement of the relative luminosity, this measurement when calibrated with an absolute measurement of the luminosity as obtained by TOTEM or by W and Z production will provide a high precision measurement of the absolute luminosity. The PLT also provides a real time determination

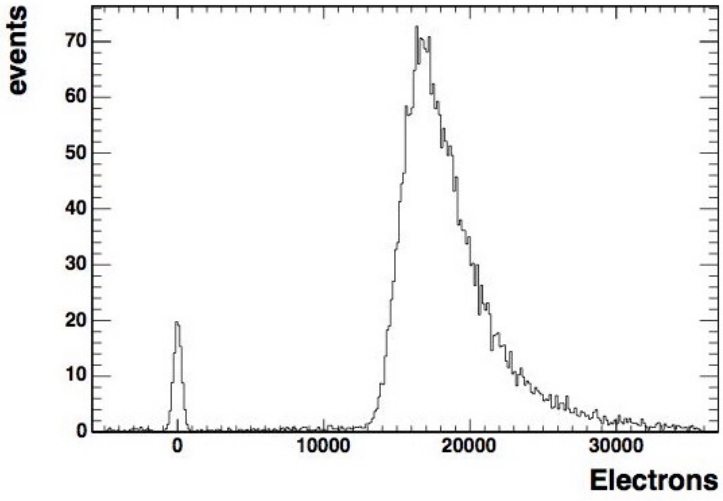


Figure 8.6: Response of a 480 μm thick single crystal CVD diamond to ^{90}Sr β's. The pedestals events are due to particles that fired the trigger scintillator but that missed the diamond sensor.

of the relative location of the interaction point centroid. It will also be able to detect and quantify the particle fluxes due to beam halo and hot spots along the beam pipe. Simulations of the PLT detector have been made to determine: the rates of particle tracks expected in the telescopes, the backgrounds expected from particle interactions in the beam pipe and other CMS components, the uniformity of acceptance around the interaction point region, the resolution achievable on the interaction point centroid from particle track extrapolation and the ability to detect beam pipe hot spots and beam halo.

The particle rates from minimum bias interactions in the PLT were simulated using PYTHIA (version 6.227). With the CMS magnetic field on but without the beam pipe or other CMS detector components, the number of three-fold coincidences per telescope per bunch crossing is 0.043 at a luminosity of 10^{33} cm^{-2}s^{-1}, 10% of design luminosity. These coincidences arise from tracks originating from the pp interaction. With a total of 16 telescopes, the net particle detection rate will be 0.69 per bunch crossing at 10^{33} cm^{-2}s^{-1} luminosity. With this rate, the relative luminosity of each of the 2835 filled bunch crossings within an orbit can be determined to a precision of 1.1% every second (about 10^4 orbits). When the beam pipe and other CMS components are included in the simulation, the three-fold coincidence rate increases by about 10% to 0.048 per bunch crossing. This increase is due almost entirely to photons that interact in the beam pipe at around $z = 1.0$ m. These are photons from the decay of π^0's produced in the pp interactions that are directed at the solid angle of a telescope. When traversing the beam pipe at a 1.5° angle, these photons see about 3.0 cm of beryllium corresponding to 8.5% of a radiation length and some are converted. The increased three-fold coincidence is, thus, due to pp production and will therefore scale with the luminosity.

In order for the luminosity measurement to be precise and stable it is important for the acceptance of the telescope array to be uniform over the interaction region. The luminosity measurement will then be insensitive to any small movement of the interaction point centroid. The acceptance was determined by shifting the interaction point in r, and z and determining the rate of the three-fold coincidences observed at each location. These rates were

then normalized to the rate obtained wsith the interactiion point at $r = 0$, $z = 0$. Fig. 8.7 shows this normalized rate as a function of r and z. The z-projection at $r = 0$ is shown in Fig. 8.8. The acceptance is flat to 1% out to $\pm$ 320 mm. The r-projections at $z = 0$ are shown in Figs. 8.9 a) and b) for the cases in which the interaction point is offset in a radial direction toward 1 of the 8 telescopes and in a radial direction exactly between 2 of the telescopes, respectively. The acceptance is 1% flat out to $r = 12$ mm, in the former, and to $r = 8$ mm, in the latter.

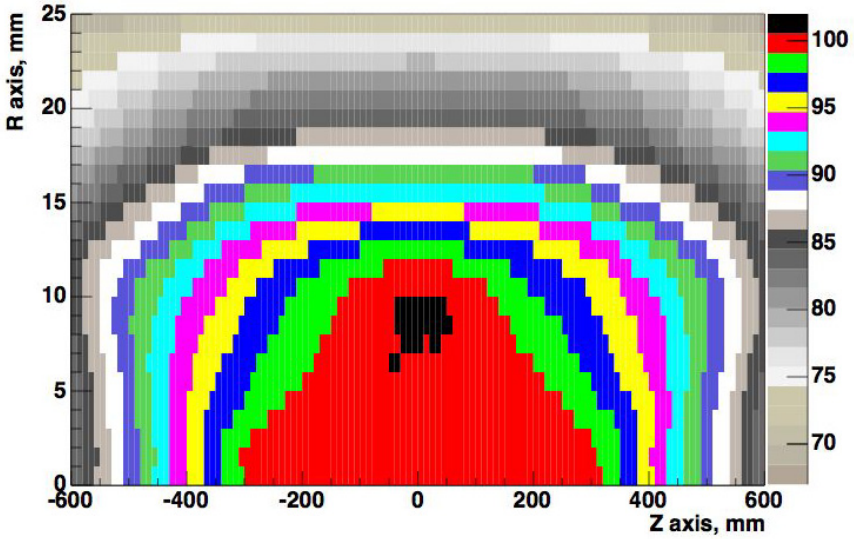


Figure 8.7: Track acceptance as a function of the position of the interaction point normalized to $r = 0$, $z = 0$.

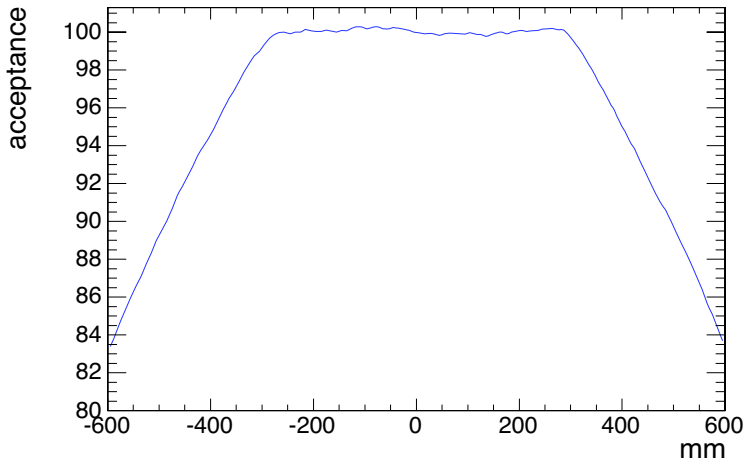


Figure 8.8: Track acceptance as a function of z at $r = 0$.

Another important function of the PLT is to determine the relative location of the interaction point centroid and to monitor any possible drifts on the time scale of a second. For this measurement, the addresses and pulse heights of those pixels above threshold, recorded for every CMS Level 1 trigger, will be used. Extrapolation of tracks to the interaction point will be affected by the bending of the particles in the 4 T magnetic field. For a given telescope, the

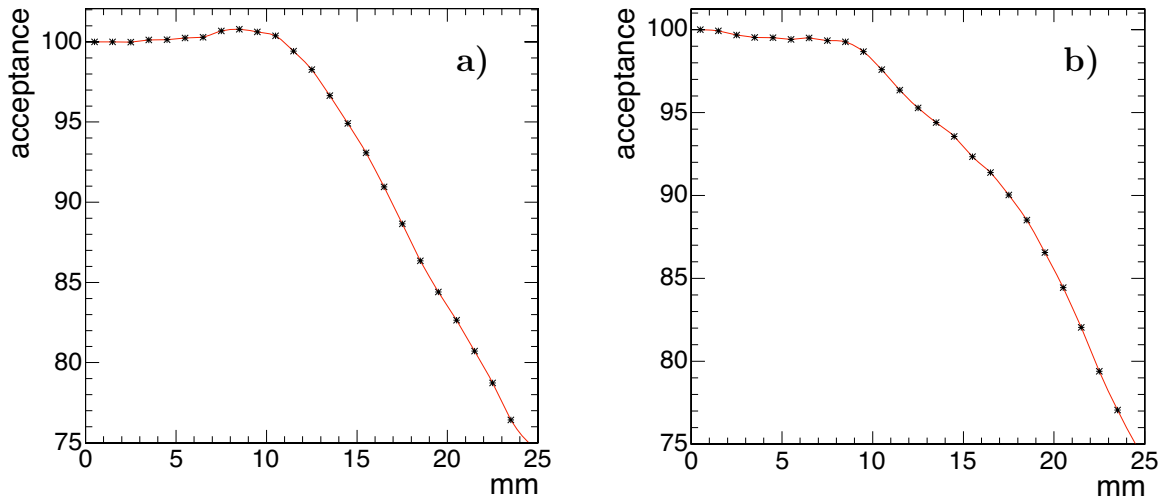


Figure 8.9: Track acceptance as a function of r at $z = 0$ for a) r in direction of 1 of the telescopes and b) r exactly between 2 of the telecopes.

effect of this curvature will be smallest in the plane determined by z and the radial direction of the telescope and will be an order of magnitude larger in the plane determined by z and the radial direction orthogonal to the telescope. For example, for the telescope at $\phi = 0$ the sagitta in the x-z plane is about 1 μm while the sagitta in the y-z plane is about 20 μm for a 10 GeV/c particle. Since the sensor planes will be essentially perpendicular to the particle tracks, there will be little charge sharing among pixels and the hit spatial resolution will be the digital resolution of 30 to 40 μm set by the pixel size of 100×150 μm ^{2}. Since even for low momentum, 10 GeV/c, tracks the sagitta is too small to be determined with the spatial resolution of the telescope, the tracks will be linearly extrapolated to the interaction point. Due to the variation of curvature with momentum and the longitudinal spread of the breams, the extrapolation point of the tracks will be smeared. Fig. 8.10 shows the extrapolation resolution in the radial direction. It has a fit standard deviation of $\sigma = 2.4$ mm. The longitudinal spread of the beam, $\sigma = 7.5$ cm, is the dominant contribution to the width of this distribution. For a 30 kHz trigger rate, more than 10^3 tracks will be recorded within 1 second at a luminosity of 10^{33} cm^{-2} s^{-1}. Four telescopes, 2 on each side of interaction region, will be used in the measurement of the x (y) interaction point location. This will allow determination of the transverse location of the interaction centroid to 35 μm in both x and y.

Because of the shallow track angle, approximately 1.5°, the RMS of the extrapolated distribution in z is 8.8 cm, 40 times greater than for the transverse distribution and comparable to the spread due to the longitudinal bunch length RMS of 7.5 cm. Using all 16 telescopes, the z interaction point centroid can be located to a precision of 720 μm within 1 second again at a luminosity of 10^{33} cm^{-2} s^{-1}.

For beam diagnostic purposes, the PLT has been designed to be sensitive to beam halo. With a 20 cm total length for the telescope, there is a 32% overlap of the first and last telescope planes for horizontal tracks making the telescopes sensitive to beam halo tracks. Their fraction can be determined by measuring the fraction of horizontal tracks seen in the full pixel data. These beam halo tracks will contribute to the three-fold fast coincidence measurement but once their fraction is determined in the pixel readout data they can be subtracted. The

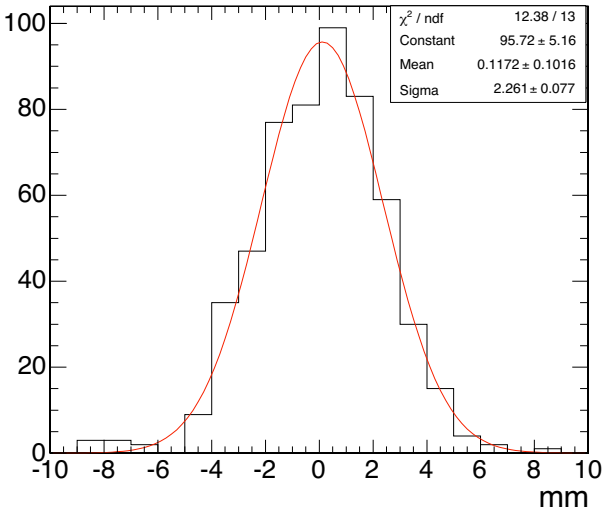


Figure 8.10: Transverse distribution of tracks linearly extrapolated to $z = 0$.

pixel readout chip has the capability of allowing individual pixels to be masked removing them from both the pixel readout and from the fast out signal. If it is discovered that the fraction of beam halo tracks is large and compromises the luminosity measurement, the acceptance for horizontal tracks can be reduced by masking out appropriate pixels.

In summary, on a time scale of 1 second at a luminosity of 10^{33} cm^{-2} s^{-1}, 10% of design, the PLT will perform the following.

- Determine the bunch-by-bunch luminosity to a precision of 1.1%.
- Determine the transverse interaction point centroid to a precision of 35 μm.
- Determine the longitudinal interaction point centroid to a precision of 720 μm.
- Locate possible beam pipe hot spots within ± 1.5 m of the interaction point.
- Monitor the amount of beam halo.

8.4 Offline and normalization techniques

8.4.1 TOTEM

The TOTEM experiment [198, 211], further described in Chapter 7, will measure the total pp cross section and study elastic and diffractive processes at the LHC. Elements of TOTEM will be situated in the far forward regions of CMS. Of particular relevance to the CMS luminosity determination will be a precision ($\approx 1\%$) measurement of the total pp cross section using the so-called luminosity independent method. This approach uses the optical theorem, which holds that the total cross section can be related to the elastic cross section extrapolated to $t = 0$ via the relation

$$\sigma_{\text{tot}} = \frac{16\pi}{(1+\rho^2)} \frac{(\mathrm{d}N_{\text{el}}/dt)_{t=0}}{N_{\text{el}} + N_{\text{inel}}}, \tag{8.6}$$

where ρ is the ratio of the real to the imaginary part of the forward scattering amplitude. The extrapolation of the elastic cross section to $t = 0$ will be made using roman pots placed far (150 200 m) from the IP.

The error on $\sigma_{\rm tot}$ is expected to be of order 1%, coming mainly from the uncertainty in the inelastic cross section measurement. Since CMS and TOTEM will operate simultaneously, it will be possible to arrive at an absolute normalization for the real time techniques described above. However, TOTEM will operate at low luminosity [4] with different machine optics, which will introduce an additional uncertainty in the calibration of the real-time methods at design luminosity.

8.4.2 W and Z rate measurements

Since the pp total cross section measurement described above can only be carried out at low luminosity, using it to normalize the luminosity at $\mathcal{L} = 10^{34}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ will involve a sizeable extrapolation. Moreover, it will not be possible to detect and correct possible long-term drifts in the normalization constants for these systems. For these reasons, it will be useful to have a normalization technique that is based on production data taking.

The production rate for Ws and Zs provides just such a "standard candle." Moreover, the rates are high enough [5] datasets of only several minutes can be normalized with good statistical accuracy, even when realistic trigger and reconstruction efficiencies are taken into account.

The main challenge will be to control the uncertainties associated with theoretical estimates of the cross section and the modeling of the detector acceptance. Frixione and Mangano have argued that uncertainties associated with the interplay between detector acceptance and the parton density functions (PDFs) can be controlled at the few percent level [229]. Moreover, since PDFs are needed to relate parton-level cross sections to the cross sections actually observed in pp collisions, uncertainties in the PDFs represent an irreducible uncertainty in many comparisons between theory and experiment, even if the pp cross sections are perfectly measured. Indeed Dittmar, Pauss, and Zürcher have argued that rapidity distributions in W and Z production can be used to determine the x distributions and consequently the corresponding "parton luminosity," and that this raises the possibility of $\approx \pm 1\%$ accuracy in the determination of certain parton-parton cross sections[230].

8.5 Sources of systematic effects

Both HF luminosity measurements involve the raw rate for energy (or hit) depositions. There is little possibility to carry out event selection, since at design luminosity, each "event" already represents the superposition of several interactions. This leaves the HF techniques vulnerable to non-pp interaction backgrounds, notably "beam-halo" and beam-gas interactions.

Beam gas interactions are thought to be negligible, but this will depend on the vacuum conditions in the ring and will only be verified through special single-beam runs.

[4] For the total cross section measurement, the LHC will operate with 43 bunches at a luminosity of $10^{28}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$, which corresponds to a luminosity per bunch 15 000 times lower than design.

[5] At design luminosity, the raw rate for $Z \to \ell^+\ell^-$ is 30 Hz.

Beam halo in the HF was estimated using a model for beam halo[231]. The resulting background was found to be negligible (at the 10^{-5} level) using the default parameters provided in the model. Beam-halo estimates are, however, notoriously difficult and a quantitative estimate will require single-beam running.

Although the PLT is also subject to beam-gas and beam-halo backgrounds, by using the its fine-grain readout capability an estimate of these backgrounds should be possible even without single-beam operation of the LHC.

8.6 Luminosity monitoring, reporting, and logging

An important design requirement is that the luminosity system must be capable of providing luminosity information whether or not the main CMS DAQ is operational. To satisfy this requirement, an autonomous DAQ system will be used to acquire and analyze real time luminosity information from the HF (via the HLX boards) and the PLT. This system will serve luminosity information to various consumers—e.g., the LHC control room, a publicly availably web page, and the database system used to log luminosity information.

A carefully thought out and reliable luminosity database will be an essential element of the luminosity reporting system. To avoid unneccessary development costs and to facilitate long-term maintenance, we will use the same database planned for maintaining calibration constants in CMS. Preliminary discussions with the CMS database group indicate that the demands of the luminosity system will be modest.

Our current plans call for an approach similar to that employed at the Tevatron, where runs were broken into small segments of order 1 minute duration[6]. These run sections are short enough that the luminosity can be taken to be constant throughout. Moreover, trigger prescaler factors will be allowed to change only on run segment boundaries. To accommodate "on-the-fly" changes, run segments will comprise an integer number of LHC orbits and will start and end during the abort gap. Scaler and luminosity tallies from the real time systems will be written for each run segment.

8.6.1 Luminosity accounting

The cross section for a given process, j, is given by

$$\sigma_j = \frac{N_j^{\text{yield}}}{\epsilon_j L_0}. \tag{8.7}$$

where N_j^{yield} is the number of events selected after background subtraction, ϵ_j is the detection efficiency of the complete selection chain (trigger and offline), and L_0 is the integrated luminosity, assuming no dead time or prescaler losses. Multiplying and dividing by the number of Level 1 Accepts (L1As) for the relevant trigger channel, N_j^{L1A}, we obtain

$$\sigma_j = \frac{1}{\epsilon_j \mathcal{L}_0} \left(\frac{N_j^{\text{yield}}}{N_j^{\text{L1A}}} \right) N_j^{\text{L1A}} \equiv \frac{1}{\epsilon_j \mathcal{L}_0} \rho_j N_j^{\text{L1A}} \tag{8.8}$$

[6]CDF calls these segments "run sections," and DØ calls them "luminosity blocks."

where $\rho_j \equiv N_j^{\rm yield}/N_j^{\rm L1A}$ is the fraction of L1As that end up in the final event sample. Note that ρ_j can be estimated from *any unbiased sample* of L1As. Thus prescaler and dead time losses need not be explicitly taken into account, although it would certainly be advisable to do so as a redundant cross check. Note that since the rejection of events in the high level trigger (HLT) is not unbiased (there is by design a strong bias against background events), rejection factors in the HLT must be tracked and logged. Put another way, ρ_j can be factored into 2 terms—i.e., $\rho_j = \rho_j^{\rm HLT}\rho_j^{\rm offline}$—and it is only possible to determine the second term from the events recorded in mass storage.

Calculating the cross section for any process thus requires knowledge of L_0, which will be determined from: (i) the online monitors, calibrated using the pp total cross section and/or the yield of Ws and Zs; (ii) the raw number of L1As ($N_j^{\rm L1A}$), which will be determined by reading the scalers; [7] (iii) ρ_j determined as outline above; and (iv) ϵ_j the trigger and reconstruction efficiency, which will be determined via the usual Monte Carlo and data based studies.

The analysis above assumes that just 1 trigger provides acceptance for the process under study. That assumption simplifies things for purpose of illustration, but is not strictly necessary. The analysis can be extended to multiple triggers with overlapping acceptance.

[7] In the CMS trigger, the raw number of L1As will be determined from reading the prescale factor, and the number of counts at the prescaler output. There is potentially another small dead time associated with Level 1 trigger rules, which will be accounted for separately.

Chapter 9

Muons

9.1 Muon reconstruction

The muon reconstruction software is able to perform reconstruction in the muon system and the silicon tracker. The software has been designed using the concept of *regional reconstruction* in order to allow its use in both the offline reconstruction and the High-Level Trigger [8] (online event selection).

Employing regional reconstruction results in very small parts of the detector actually needing to be reconstructed before a physics object is validated. As an example, the amount of information needed to reconstruct a muon track in the silicon tracker represents less than a few per cent of the total tracker data volume. The software does not reconstruct tracks in the entire tracker, but only in that part which can possibly be involved in the reconstruction of a charged particle track compatible with the hits in the muon chambers. This results in savings on the overall CPU power needed to process the events. The method depends strongly on the identification of a good "seed," providing initial values of the 5 trajectory parameters and their errors, that can start the reconstruction with high efficiency and reliability. In the online environment these seeds are provided by the Level-1 Trigger system [7].

The muon reconstruction algorithm used by the HLT is seeded by the muon candidates found by the Level-1 muon trigger, including those candidates that did not necessarily lead to a Level-1 trigger accept. These seeds define a region of interest in the muon system, in which local reconstruction is performed. For offline reconstruction a different seed-generation algorithm has been developed, which performs local reconstruction in the entire muon system and uses patterns of segments reconstructed in the CSC and/or DT chambers as initial seeds. Muon reconstruction is performed in 3 stages: local reconstruction (local-pattern recognition), standalone reconstruction and global reconstruction. Starting from a seed, the chambers compatible with the seed are identified and local reconstruction is performed only in these chambers. Standalone muon reconstruction uses only information from the muon system, while global-muon reconstruction uses also silicon tracker hits. The HLT standalone and global reconstruction are called Level-2 and Level-3 reconstruction, respectively.

9.1.1 Standalone muon reconstruction

The standalone/Level-2 muon reconstruction uses only data from the muon detectors—the silicon tracker (Chapter 6) is not used. Both tracking detectors (DT and CSC) and RPCs participate in the reconstruction. Despite the coarser spatial resolution, the RPCs complement the tracking chambers, especially where the geometrical coverage is problematic, mostly in the barrel-endcap overlap region.

The reconstruction starts with the track segments from the muon chambers obtained by the local reconstruction (Section 3.5). The state vectors (track position, momentum, and direction) associated with the segments found in the innermost chambers are used to seed the muon trajectories, working from inside out, using the Kalman-filter technique [232]. The predicted state vector at the next measurement surface is compared with existing measurements and updated accordingly. In the barrel DT chambers, reconstructed track segments are used as measurements in the Kalman-filter procedure. In the endcap CSC chambers, where the magnetic field is inhomogeneous, the individual reconstructed constituents (three-dimensional hits) of the segments are used instead. Reconstructed hits from the RPC chambers are also included. A suitable χ^2 cut is applied in order to reject bad hits, mostly due to showering, delta rays and pair production. In case no matching hits (or segments) are found, e.g. due to detector inefficiencies, geometrical cracks, or hard showering, the search is continued in the next station. The state is propagated from one station to the next using the GEANE package [233], which takes into account the muon energy loss in the material, the effect of multiple scattering, and the nonuniform magnetic field in the muon system. The track parameters and the corresponding errors are updated at each step. The procedure is iterated until the outermost measurement surface of the muon system is reached. A backward Kalman filter is then applied, working from outside in, and the track parameters are defined at the innermost muon station. Finally, the track is extrapolated to the nominal interaction point (defined by the beam-spot size: $\sigma_{xy} = 15\,\mu$m and $\sigma_z = 5.3\,$cm) and a vertex-constrained fit to the track parameters is performed.

9.1.2 Global muon reconstruction

The global/Level-3 muon reconstruction consists in extending the muon trajectories to include hits in the silicon tracker (silicon strip and silicon pixel detectors). Starting from a standalone reconstructed muon, the muon trajectory is extrapolated from the innermost muon station to the outer tracker surface, taking into account the muon energy loss in the material and the effect of multiple scattering. The GEANE package is currently used for the propagation through the steel, the coil and the calorimeters. Silicon layers [6] compatible with the muon trajectory are then determined, and a region of interest within them is defined in which to perform regional track reconstruction. The determination of the region of interest is based on the track parameters and their corresponding uncertainties of the extrapolated muon trajectory, obtained with the assumption that the muon originates from the interaction point. The definition of the region of interest has a strong impact on the reconstruction efficiency, fake rate, and CPU reconstruction time: well measured muons are reconstructed faster and with higher efficiency than poorly measured ones.

Inside the region of interest, initial candidates for the muon trajectory (regional seeds) are built from pairs of reconstructed hits. The 2 hits forming a seed must come from 2 different tracker layers, and all combinations of compatible pixel and double-sided silicon strip layers are used in order to achieve high efficiency. In addition, a relaxed beam-spot constraint is applied to track candidates above a given transverse momentum threshold to obtain initial trajectory parameters. Starting from the regional seeds, a track-reconstruction algorithm, based on the Kalman-filter technique, is used to reconstruct tracks inside the selected region of interest. The track-reconstruction algorithm consists of the following steps: trajectory building (seeded pattern recognition), trajectory cleaning (resolution of ambiguities) and trajectory smoothing (final fit). In the first step, the trajectory builder transforms each seed into a set of trajectories. Starting from the innermost layer, the trajectory is propagated to the next

tracker reachable layer, and updated with compatible measurements found on that layer. In the second step, the trajectory cleaner resolves ambiguities between multiple trajectories that may result from a single seed on the basis of the number of hits and the χ^2 of the track fit. In the final step, all reconstructed tracks are fitted once again, without a beam-spot constraint, using the hits in the muon chambers from the original standalone reconstruction together with the hits in the silicon tracker. To resolve possible ambiguities a second cleaning step is performed which selects the final muon candidates on the basis of a χ^2 cut.

The selected trajectories are then refitted excluding measurements (hits or segments) with high χ^2 values in muon stations with high hit occupancy. In addition the trajectories are refitted using only silicon tracker hits plus hits in the innermost muon station (excluding RPC hits) and the χ^2 probability of the fit is compared with the χ^2 probability of the tracker-only trajectory in order to detect muon bremsstrahlung or any kind of significant energy loss of the muon before the first muon station. This procedure is important for the accurate momentum reconstruction of very high-p_{T} (TeV) muons. The precise reconstruction of these objects is very challenging because of catastrophic energy loss and severe electromagnetic showers in the muon system.

9.1.3 Performance

9.1.3.1 Ideal detector

This subsection documents the *ideal* offline reconstruction performance, i.e., without taking into account effects from misalignment, miscalibration, neutron background or pile-up. Simulated single-muon samples with both fixed and continuous values of p_{T} (between 10 GeV/c and 1 TeV/c), and with a flat distribution in η and ϕ are used for this study. For the dimuon reconstruction performance studies, the following event samples were used: $J/\psi \to \mu^+\mu^-$, $Z \to \mu^+\mu^-$, $Z' \to \mu^+\mu^-$ (with Z' masses of 1, 1.5, 3 and 5 TeV/c^2) and various Drell-Yan dimuon samples generated with mass cut-offs from 70 to 1000 GeV/c^2.

Figures 9.1a and 9.1b show the reconstruction efficiency as a function of pseudorapidity for various values of p_{T}. Simulated single-muon samples were used to obtain the results shown in Figs. 9.1–9.4. The reconstruction efficiency achieved is typically 95–99%, except in the pseudorapidity regions around $|\eta| = 0.25$ and $|\eta| = 0.8$ (the regions between 2 DT wheels) and around $|\eta| = 1.2$ (the transition region between the DT and CSC systems), where the efficiency drops. (Figure 1.6 displays the geometric structure of the muon system.)

Figure 9.2 shows the transverse momentum resolution as a function of pseudorapidity for muons reconstructed with the standalone reconstruction algorithm (a) and the global reconstruction algorithm (b). The transverse momentum resolution was obtained by a Gaussian to the distribution of the quantity

$$\frac{q^{\mathrm{rec}}/p_{\mathrm{T}}^{\mathrm{rec}} - q^{\mathrm{gen}}/p_{\mathrm{T}}^{\mathrm{gen}}}{q^{\mathrm{gen}}/p_{\mathrm{T}}^{\mathrm{gen}}}, \tag{9.1}$$

where q is the charge and $p_{\mathrm{T}}^{\mathrm{gen}}$ and $p_{\mathrm{T}}^{\mathrm{rec}}$ are the generated and reconstructed transverse momenta, respectively.

The standalone muon momentum precision is essentially determined by the measurement in the transverse plane of the muon bending angle at the exit of the 4 T coil, taking the interaction point as the origin of the muon. This measurement is dominated by multiple scattering

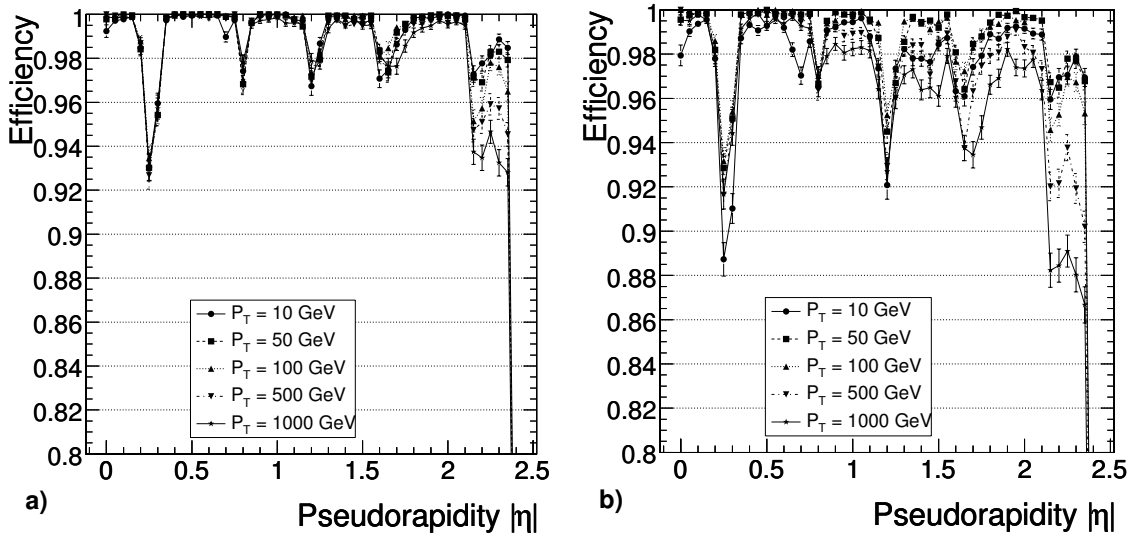


Figure 9.1: Muon reconstruction efficiency as a function of pseudorapidity for various values of p_{T}. a) Standalone reconstruction (using only hits from the muon system with a vertex constraint); b) Global reconstruction (using hits from the muon system and the tracker).

in the material before the first muon station up to p_{T} values of 200 GeV/c, when the chamber spatial resolution starts to dominate. For low-momentum muons, the best momentum resolution is given by the resolution obtained in the silicon tracker. Using measurements from the silicon tracker in addition to the muon system substantially improves the p_{T}-resolution (a factor of 10) compared to the resolution obtained by the standalone muon reconstruction.

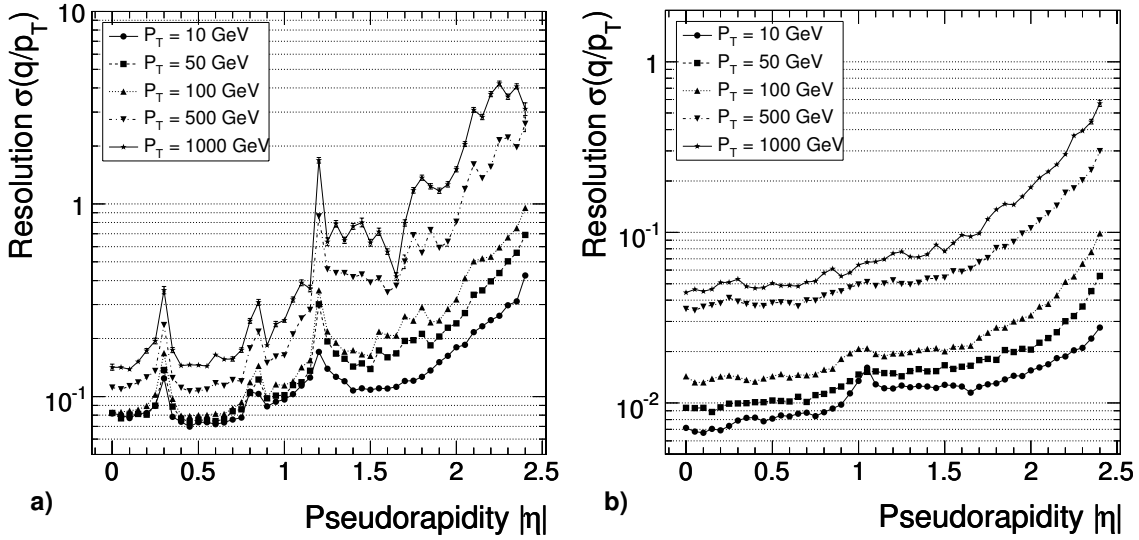


Figure 9.2: The q/p_{T} resolution for standalone (a) and globally (b) reconstructed muons (combined muon system and silicon tracker) as a function of pseudorapidity.

Figure 9.3 shows a comparison of the momentum resolution obtained from standalone and global reconstruction as a function of momentum. Two values of η have been chosen to illustrate the effect in the barrel and the endcap regions.

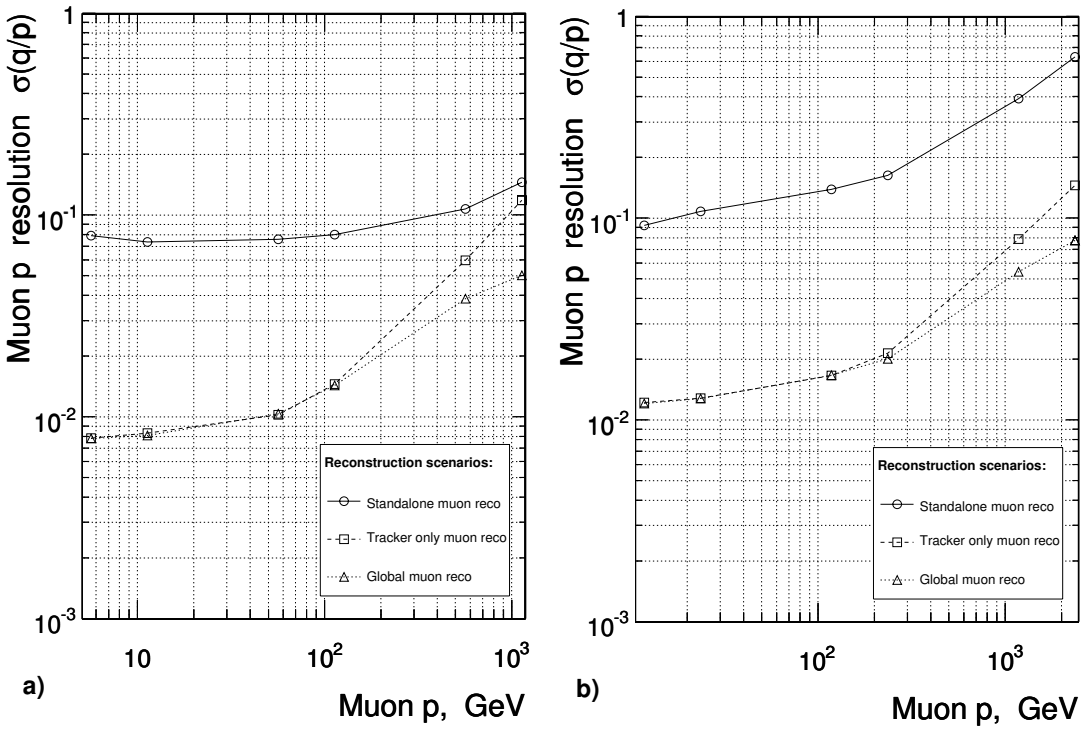


Figure 9.3: Resolution of $(1/p)$ versus p for standalone, global and tracker-only reconstruction. a) Barrel, $\eta = 0.5$; b) Endcap, $\eta = 1.5$.

In addition to an accurate measurement of its momentum, it is important that a reconstructed muon has the correct charge assignment. The probability to reconstruct muons with misassigned charge is shown in Fig. 9.4 as a function of pseudorapidity for various values of p_{T} for both standalone and global reconstruction. When the measurements from the silicon tracker are included, the percentage of misassigned charges for muons with $p_{\mathrm{T}} = 100$ GeV$/c$ is less than 0.1%.

9.1.3.2 Realistic detector

The muon reconstruction performance, taking into account effects from detector misalignment and event pile-up, is discussed in this Section. Events are produced using the full detector simulation package (OSCAR), which simulates the particle propagation and interactions through the detector (Section 2.5). A detailed description of the detector geometry, detector materials, and the magnetic field was used. The effect of pile-up has been taken into account assuming a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Misalignment of muon chambers is introduced using a dedicated software package that provides the means to modify the simulated measurements after digitization. The package allows different misalignment displacements to be run, using the same set of digitized data. It also provides a description of 2 predefined misalignment scenarios, called the "First-Data-Taking Scenario" and the "Long-Term Scenario." The first one corresponds to the alignment at the beginning of data taking, while the second one describes the situation when all alignment procedures have enough data to obtain a full set of alignment constants. A detailed description of the alignment scenarios can be found in Section 3.2.2.

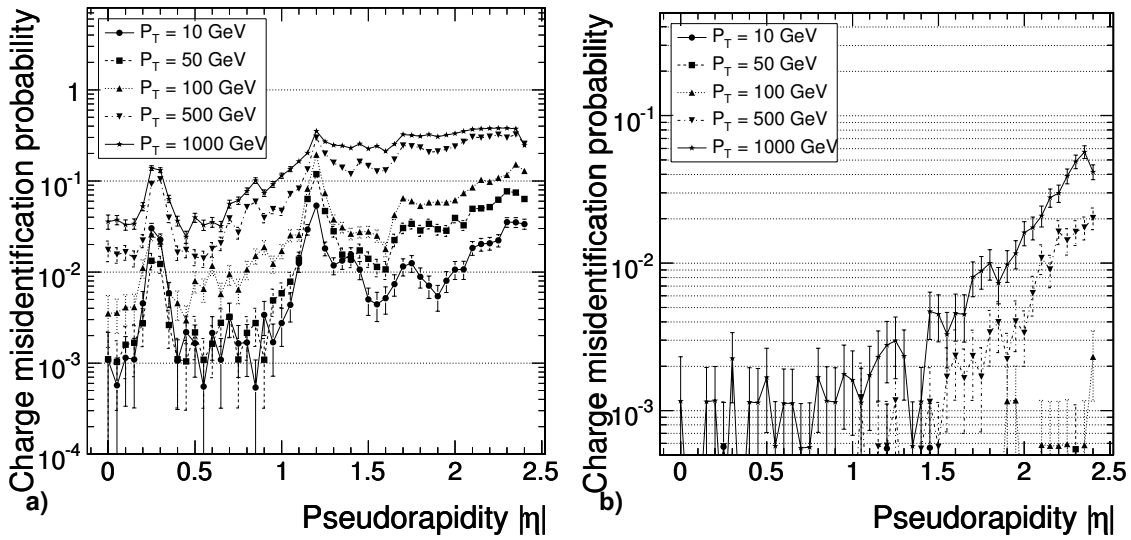


Figure 9.4: Muon charge-misidentification probability as a function of pseudorapidity for a) standalone and b) global reconstruction.

Figures 9.5, 9.6 and 9.7 show the effect of misalignment for standalone reconstruction. The effect of misalignment is demonstrated in terms of reconstruction efficiency, transverse momentum resolution and charge misassignment. The results were obtained using simulated single-muon samples. Figures 9.8, 9.9 and 9.10 show the same distributions for global muon reconstruction, using hits from the muon system and the silicon tracker.

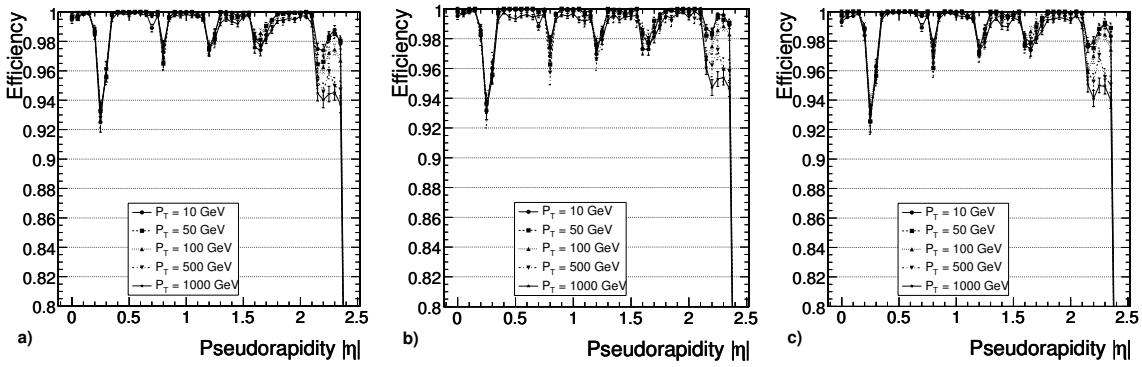


Figure 9.5: Standalone muon reconstruction efficiency as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the "First-Data-Taking Scenario;" and c) aligned according to the "Long-Term Scenario."

Figure 9.11 shows the dimuon reconstruction efficiency for various physics channels assuming different detector misalignment scenarios. The efficiencies shown are offline reconstruction efficiencies obtained from events having 2 muons within the muon detector acceptance. No requirement on the charges of the 2 muons was imposed in the reconstruction of these dimuon resonances. Figure 9.12 displays the corresponding dimuon invariant mass resolution. Samples of $J/\psi \rightarrow \mu^+\mu^-$ from $B_S^0 \rightarrow J/\psi\ \phi$, $Z^0 \rightarrow \mu^+\mu^-$, $Z' \rightarrow \mu^+\mu^-$ and Drell-Yan events were used to obtain these results.

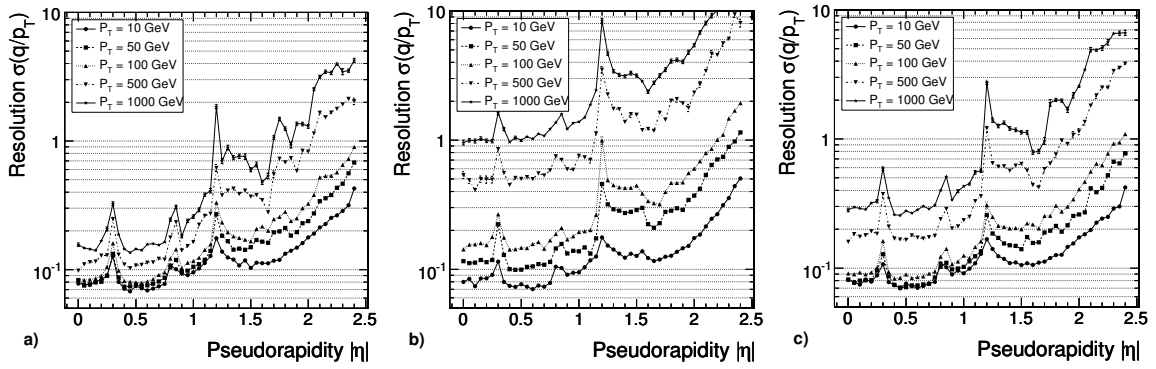


Figure 9.6: (q/p_{T}) resolution for standalone reconstructed muons (using only hits from the muon system with a vertex constraint) as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

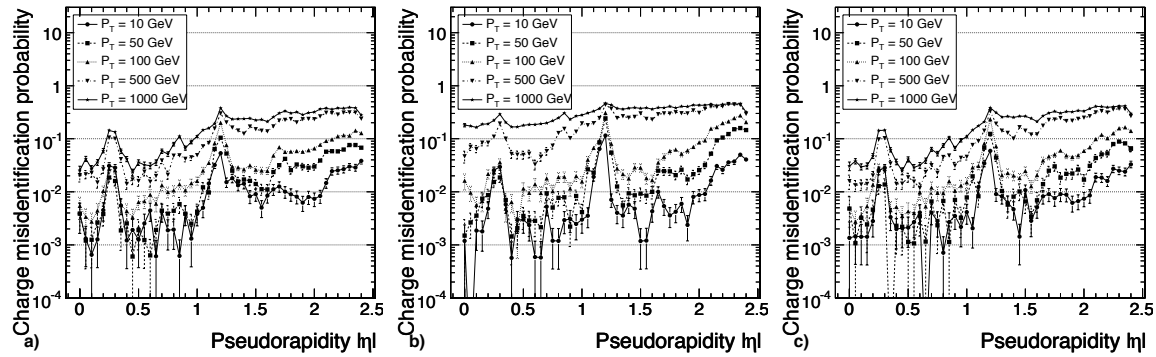


Figure 9.7: Muon charge misidentification probability for standalone reconstruction as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

Figure 9.13 shows the reconstructed invariant mass distributions for $J/\psi \rightarrow \mu^+\mu^-$ from $B_S^0 \rightarrow J/\psi\ \phi$ decays for different misalignment scenarios. Figure 9.14 displays the same distributions for $Z' \rightarrow \mu^+\mu^-$ decays, assuming a Z' mass of 1 TeV/c^2.

9.1.4 Cosmic-ray and beam-halo muon reconstruction

Beam-halo and cosmic-ray muons are 2 sources of particles that do not originate from LHC collisions but can nevertheless be found in the CMS detector. Whereas beam-halo muons move predominantly along the beam axis, cosmic-ray muons typically traverse the detector from top to bottom. While such muons have to be considered as sources of background, they can be employed for many purposes, for instance alignment, calibration, and the understanding of trigger and reconstruction efficiencies.

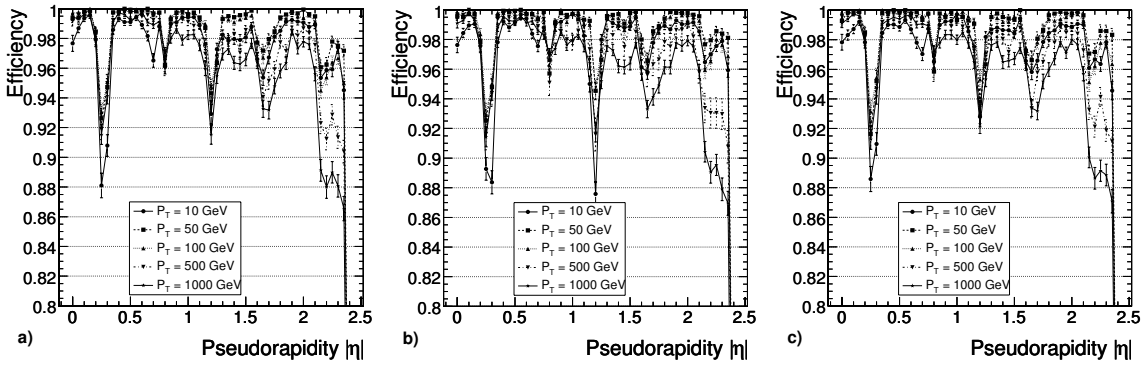


Figure 9.8: Global muon reconstruction efficiency as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

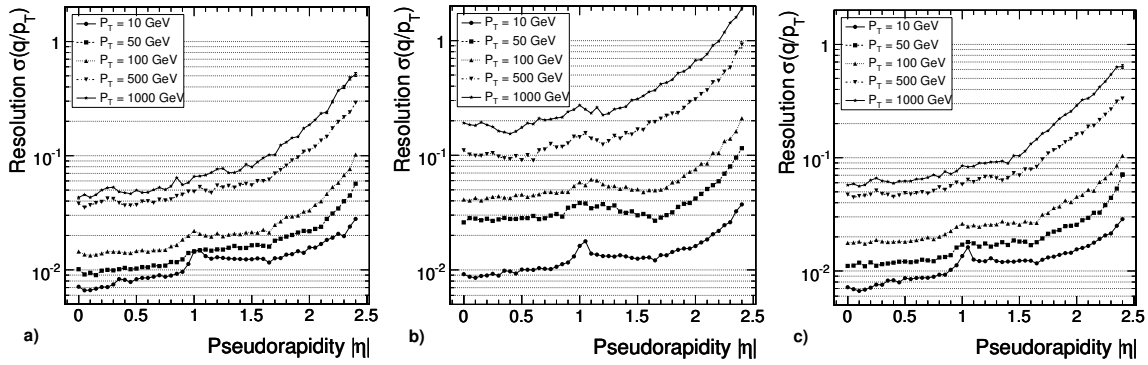


Figure 9.9: (q/p_{T}) resolution for global reconstructed muons (using only hits from the muon system with a vertex constraint) as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

9.1.4.1 Event generation and detector simulation

The simulation of beam-halo and cosmic-ray muons, described in detail in [199], is performed in two steps. The muons are first generated with $E(\mu^{\pm}) > 10$ GeV, and they are then passed through the CMS detector simulation.

Cosmic-ray muons are generated at the surface of the earth according to measurements described in [119], and are then propagated to the surface of the CMS detector, taking into account the energy loss in the material below the surface of the earth.

Beam-halo muons are generated according to kinematic distributions that have been provided by the LHC simulation group [234]. It should be stressed that the rates given in [234] are valid only for one specific set of LHC parameters, which is obsolete now. The absolute rates given here are thus to be understood as rough estimates. At the entrances of the LHC tunnel, the muons are passed to the detector simulation, which includes the simulation of the forward shielding.

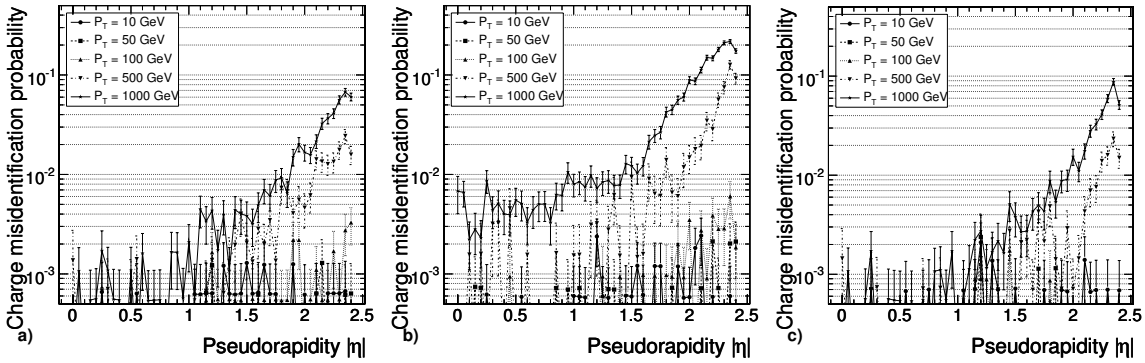


Figure 9.10: Muon charge misidentification probability in global reconstruction as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long Term Scenario.

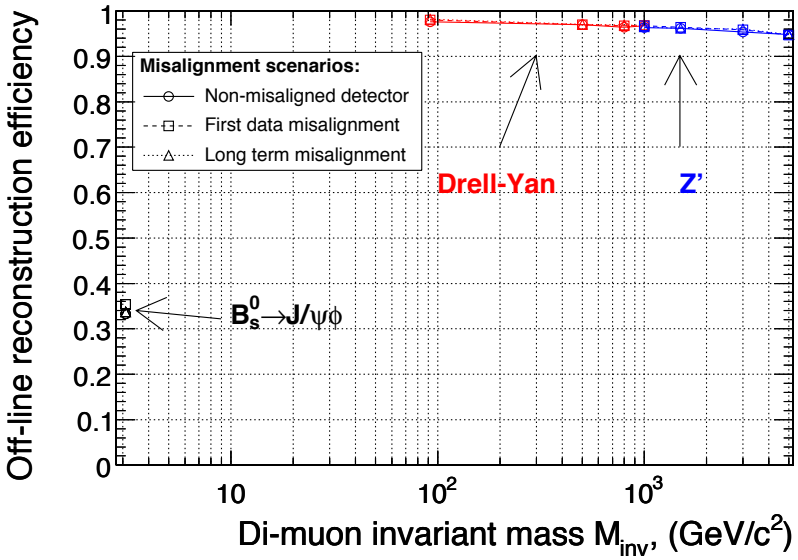


Figure 9.11: Dimuon reconstruction efficiency for different physics channels:: $J/\psi \rightarrow \mu^+\mu^-$ from $B_S^0 \rightarrow J/\psi\ \phi$, $Z^0 \rightarrow \mu^+\mu^-$ and $Z' \rightarrow \mu^+\mu^-$, and for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking scenario; c) aligned according to the Long-Term Scenario.

9.1.4.2 Level-1 trigger and standalone reconstruction

Two scenarios are discussed here, a "single-beam run" to study beam-halo muons and a "cosmic run" for cosmic-ray muons in the CMS cavern. Since an event consists of a single muon, a Level-1 trigger requirement is sufficient to control the event rate. In order to maximize the Level-1 rate, the look-up tables (LUT) for the CSC and DT detectors are widened as much as possible, which allows the triggering of muons with at least 2 hits and almost no angular constraint.

Because beam-halo and cosmic-ray muons enter the CMS detector from outside and do not

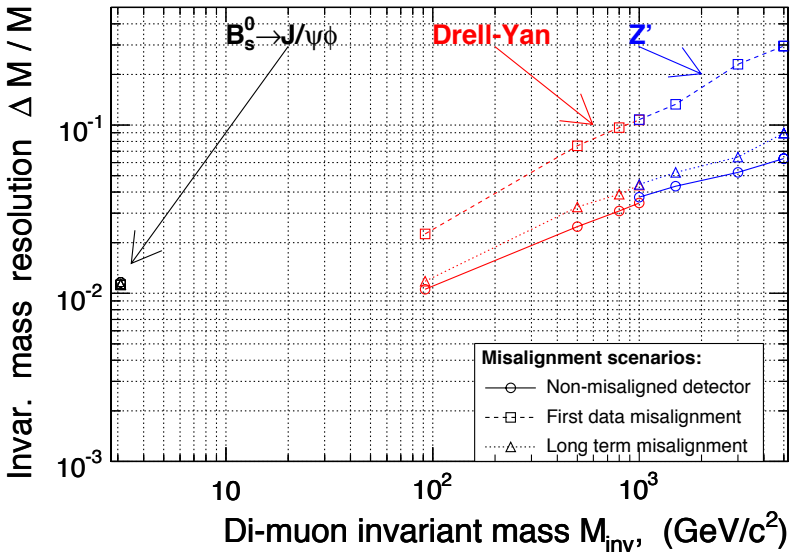


Figure 9.12: Dimuon invariant mass resolution for different physics channels: $J/\psi \to \mu^+\mu^-$ from $B_S^0 \to J/\psi\ \phi$, $Z^0 \to \mu^+\mu^-$ and $Z' \to \mu^+\mu^-$, and for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

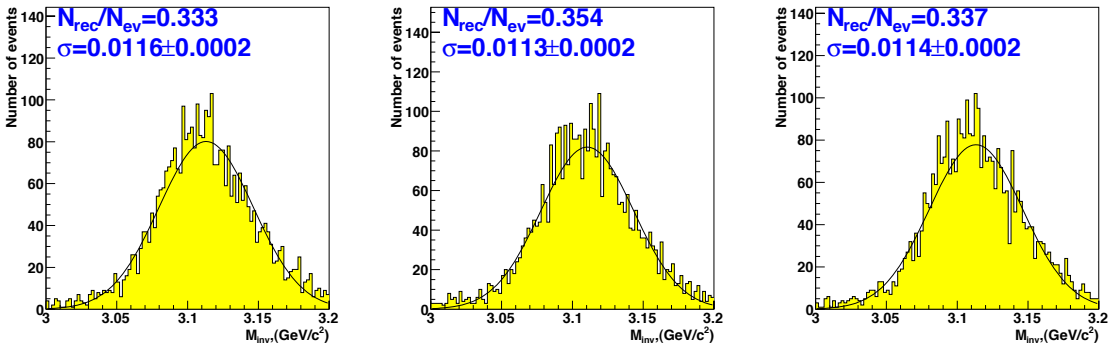


Figure 9.13: Reconstructed invariant mass distributions for $J/\psi \to \mu^+\mu^-$ from $B_S^0 \to J/\psi\ \phi$ decays for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

come from the interaction point, the time-of-flight patterns are substantially different from muons produced in pp collisions. The different timing requires a different synchronization of the detector. This is done separately for beam-halo and cosmic-ray muons by assigning an average, constant time-of-flight offset to each muon chamber.

Whereas most of the beam-halo muons are expected to have the same bunch structure as the proton beam itself, the bunch crossing assignment is rather artificial for the case of cosmic-ray muons. The Level-1 trigger rate increases by about 16% when the neighbouring bunches are added to the bunch with the highest rate.

The standalone (SA) muon reconstruction algorithm turns out to be flexible enough to per-

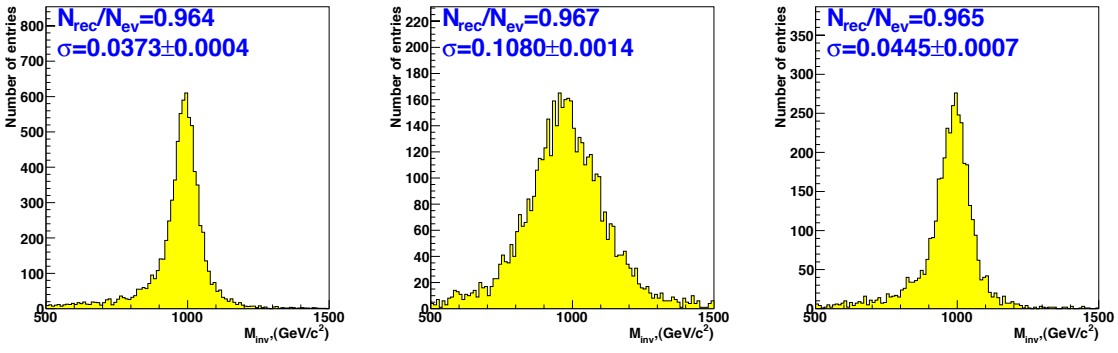


Figure 9.14: Reconstructed invariant mass for $Z' \rightarrow \mu^+\mu^-$ decays for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

form well on cosmic-ray muons. For beam-halo muons, the constraint on the pseudorapidity has to be ignored. At least 3 reconstructed hits are required for all SA muon tracks.

9.1.4.3 Kinematic properties and rate estimates

An important property of both beam-halo and cosmic-ray muons is that, if they are energetic enough, they can lead to 2 SA muon tracks. Beam-halo muons, as shown in Fig. 9.15, move almost parallel to the beam axis and lead to tracks in the forward muon detectors. When 2 SA muon tracks come from 1 beam-halo muon, one track is reconstructed on each side of the detector. This is useful for inter-alignment of the 2 halves of the forward muon chambers.

If only 1 SA muon track is found, it is most probably on the side from which the beam-halo muon enters the CMS detector. The opposite side is much less populated because many muons are stopped in the detector material, mainly the iron yoke, the magnet and the calorimeters, on their way through CMS.

The kinematic properties of cosmic-ray muons are again different, as illustrated in Fig. 9.16. A cosmic-ray muon typically traverses the CMS detector from top to bottom, and the majority of the tracks are reconstructed in the barrel region.

The majority of tracks are reconstructed in the barrel region. There are more single SA muon tracks reconstructed on the top than the bottom, due to the energy loss in the detector material. Cosmic-ray muons with 2 reconstructed tracks connect the upper half with the lower half of the CMS detector, including the inner tracker in cases where the muons are central enough.

In general, events with 2 SA muon tracks are useful for studying trigger and reconstruction efficiencies, energy loss in the detector material, and alignment on a large scale. Events with only 1 SA muon track are expected to add precision to the alignment of neighbouring muons chambers. The number of triggered events of each type are summarized in Table 9.1. It should be kept in mind that all rates depend on the trigger conditions and the track requirements.

In summary, beam-halo and cosmic-ray muons can be triggered and reconstructed with the CMS muon system. Both sources of muons are useful for studying and understanding the

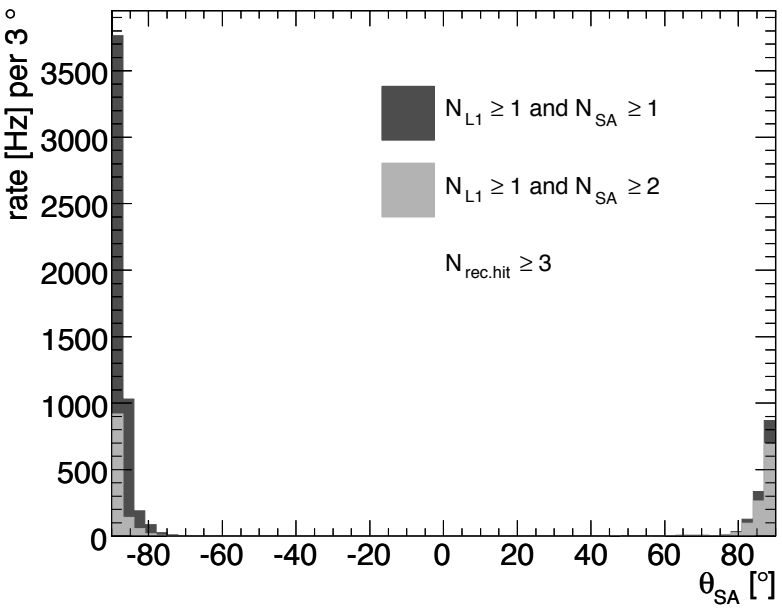


Figure 9.15: Rates of SA muon tracks as a function of θ_{SA} for beam-halo muons accepted by the Level-1 trigger. The beam-halo muons come from the $-z$ ($-\theta_{\mathrm{SA}}$) side and move to the $+z$ ($+\theta_{\mathrm{SA}}$) side.

Table 9.1: Event rates for beam-halo and cosmic-ray muons. All rates are for events accepted by the Level-1 trigger. Events with reconstructed SA muons are divided into 2 categories: events with exactly 1 SA muon track and events with exactly two SA muon tracks.

event rate	Level-1	Level-1 with 1 SA track	Level-1 with 2 SA tracks
beam-halo muons	6600 Hz	4300 Hz	1100 Hz
cosmic-ray muons	680 Hz	330 Hz	140 Hz

performance of the detector, especially the muon trigger system. Furthermore, these muons can help align and calibrate the CMS detector. A comprehensive study with many more details is given in [118].

9.1.5 Dimuon reconstruction in heavy-ion collisions

The reconstruction of muons in the heavy-ion environment requires specialized algorithms due to the high hit occupancy in the tracker [235]. In addition, it is important to extend the detector acceptance to the lowest-p_{T} muons. The relatively low expected luminosity in heavy-ion collisions allows a less stringent trigger criteria as compared to the pp case (Section 3.4.7).

The muon tracks in the muon chambers are used to identify hits in the outer tracker that can form the starting points (seeds) for the matching muon candidate tracks. The initial momentum value of the muon candidate is taken from the trigger momentum estimation. If

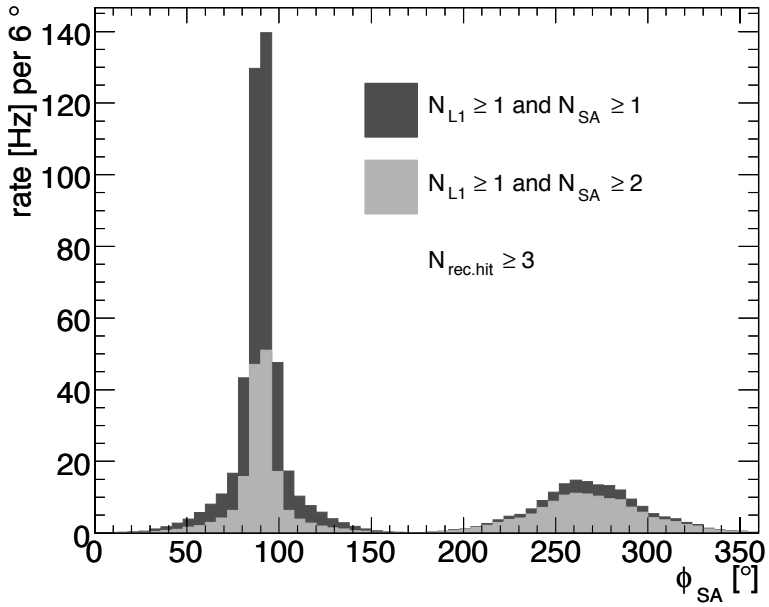


Figure 9.16: Rates of SA muon tracks as a function of ϕ_{SA} for cosmic-ray muons accepted by the Level-1 trigger. The maximum rate is at the top of the detector, which corresponds to $\phi_{\mathrm{SA}} = 90°$. The bump around $\phi_{\mathrm{SA}} = 270°$ is at the bottom of the CMS detector.

the candidate has both Level-1 and Level-2 momentum estimations, the momentum obtained by the Level-2 trigger is used. Only hits located within a search window are considered as seeds.

The track searching and fitting is done in several passes and is performed simultaneously for both muons forming the dimuon pair. During the first pass, both muon track candidates are required to have hits in all the outer silicon and pixel layers. For the successful candidates, a Kalman-fitting procedure is performed inside the tracker volume for both muons of the pair. Muon candidates that satisfy a χ^2 quality criterion are propagated back to the muon stations and are ordered according to their distance to the muon track segments. For the surviving pairs, a common vertex is calculated using a Principal Vertex Finder. In the final step, the pair that has the best quality (the smallest value of $\chi^2_{\mu_1} \times \chi^2_{\mu_2} \times \chi^2_{\mathrm{vertex}}$ and the smallest distance to the muon track segments) is chosen. If there are no pairs that satisfy the above criteria, the propagation pass is repeated, this time allowing one missing tracker layer for 1 of the muons. The selection procedure is terminated if either a pair is found or all possible combination of hits with 1 missing layer is considered. The detailed description of the track-finding procedure is given in [236].

9.2 Muon identification

Standard muon track reconstruction, as described in Section 9.1, starts from the muon system and combines standalone reconstructed muon tracks with tracks reconstructed in the silicon tracker. The muon identification algorithm provides an additional tool for selecting muon candidates using an inside-out algorithm that is by design complementary to the stan-

dard muon reconstruction. The algorithm starts with all reconstructed tracks from the silicon inner tracking detector and attempts to quantify a muon compatibility for each (including those tracks with no matching standalone tracks in the muon detectors). Muon identification also takes advantage of information from other detectors not used in reconstruction, such as associated energy deposits in the calorimeters for determining the compatibility of a given track with the muon hypothesis. In addition, the algorithm is able to make use of hit information from individual layers of the muon detectors, even in cases where those hits are not associated with a reconstructed, standalone muon track. Lower-p_{T} muons, for example, that range out within the iron yoke before depositing hits in the outer muon detector layers are difficult to reconstruct with the standalone muon track reconstruction. However, this class of muons can potentially be identified offline via this algorithm by matching the hits found in the inner muon detector layers with reconstructed silicon tracks and examining the associated calorimeter energy deposits to see if they are compatible with those from a minimum-ionizing particle.

The muon identification algorithm is useful for physics analyses involving event signatures that contain multiple muons (e.g., $H \rightarrow ZZ^* \rightarrow \mu^+\mu^-\mu^+\mu^-$), where the reconstruction conditions can be relaxed for some muons. If the muon identification algorithm is used to identify one or more of the muons in these channels rather than requiring all muons to be fully reconstructed, the overall selection efficiency for these events is increased. The muon identification algorithm can be particularly useful in cases where one has a soft p_{T} spectrum for the lowest-energy muon in the event since these muons are more likely to range out in the detector and not be fully reconstructed. Even for event topologies in which all of the final state muons are fully reconstructed, the information provided by the identification algorithm helps in understanding muon fake rates and solving potential ambiguities at the analysis level. In addition, the algorithm will be useful during commissioning for picking out a class of potential, unreconstructed muon candidates for studying and improving the reconstruction software.

9.2.1 Algorithm

The basic idea of the muon identification algorithm is to extrapolate each reconstructed silicon track outward to its most probable location within each detector of interest (ECAL, HCAL, HO, muon system). The track extrapolation takes into account the magnetic field and energy loss via the GEANE software package [233]. The algorithm then searches in a cone around that most probable location for any signals that might be associated with the track of interest. The cone sizes used for making these searches are different for each detector and are parameterized as functions of the p_{T} and η of the silicon track. Optimization of cone sizes is based on studies of simulated event samples. After collecting the associated signals from each detector, the algorithm determines a combined compatibility value corresponding to how well the observed signals fit with the hypothesis that the silicon track used to seed the algorithm is produced by a muon. The assignment of the compatibility value is somewhat arbitrary and based primarily on studies of various simulated event samples. The compatibility variable output by the algorithm is constructed so that different physics analyses can make different cuts on the minimum value required to identify muons. In choosing this threshold, each analysis can independently optimize the trade off between muon selection efficiency and the background fake rate for muons.

9.2.1.1 Identification in the calorimeters

The minimum-ionizing signature of muons in the calorimeters can be used to identify reconstructed tracks produced by muons. The identification algorithm looks for energy deposits associated with candidate tracks in both the electromagnetic and hadron calorimeters and compares the observed energies with those expected for a minimum-ionizing particle. In the central region of the detector ($|\eta| < 1.2$), the algorithm also checks if a signal is observed in the additional outer hadron calorimeter (HO), a scintillator layer located just outside of the superconducting solenoid in front of the inner muon chambers. Particles that enter this layer, having passed through the material in the calorimeter modules and magnet, are more likely to be minimum-ionizing. In addition, any excess energy deposit in cells from sources such as pile-up events should have less of an effect in the outer layers. There should, therefore, be a relatively lower probability of misidentification in the outer layers of the calorimeters, especially in the HO.

The muon identification algorithm extrapolates the seed track to its most likely location within each calorimeter volume. In the case of the HO layer, the location of the extrapolated track at the innermost muon chambers is used. The algorithm determines the total energy associated with the track candidate by summing the energies in towers whose centers are within some distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ of the extrapolated track position. The value of ΔR is optimized separately for each calorimeter to account for differences in tower granularities and to ensure that a big enough search cone is used to collect all of the energy associated with muons that either share their energy in adjacent cells or undergo significant multiple scattering. Optimization of cone sizes is also done independently for 3 separate bins in track η: barrel ($|\eta| < 0.8$), overlap ($0.8 < |\eta| < 1.2$), and endcap ($|\eta| > 1.2$) to account for differences in the calorimeter modules used in these regions. In order to avoid collecting energies not directly associated with the candidate track, the ΔR values used by the algorithm in each region are chosen to be as small as possible. However, they are not so small that a high collection efficiency for the energies associated with real muon tracks is not maintained. For the finely segmented electromagnetic calorimeter modules, studies of simulated single-muon event samples advocate optimal ΔR values between 0.02 and 0.03 (all energy deposited in 1 crystal). In the case of the more coarsely segmented hadron calorimeter modules and the HO scintillator layer, bigger ΔR values between 0.08 and 0.13 are found to be the optimal choice.

The energy thresholds for observing the minimum-ionizing calorimeter signals are also important. For the offline algorithm, the single-tower thresholds are made as low as possible so that the small minimum-ionizing signal can be observed, especially in the case of the signals from the electromagnetic calorimeter. The minimum-ionizing energy expected to be deposited in the crystals is on the order of 300 MeV in the barrel region and 400 MeV in the endcap. Due to the fact that a high threshold on the order of 500 MeV is applied during the online readout of the endcap crystals, it is not always possible to observe the minimum-ionizing signal in the forward-electromagnetic calorimeter modules. A much lower threshold of about 60 MeV is applied during the online readout of the barrel crystals, which means that there is a much higher efficiency for observing the minimum-ionizing signal in the central electromagnetic calorimeter modules. The difference is illustrated in Fig. 9.17. The top and bottom figures show the electromagnetic energy returned from the muon identification algorithm for single $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ muon tracks in the barrel and endcap regions of the detector. In the endcap region, there are a substantial number of cases where the algorithm finds no electromagnetic energy associated with the track. For the hadron calorimeter modules, the online

readout thresholds are well below the threshold of the expected minimum-ionizing signal. Therefore, an offline single-tower threshold on the order of 700 MeV is applied within the identification algorithm for forming the hadron calorimeter energy sum. This threshold is applied to remove potentially noisy towers.

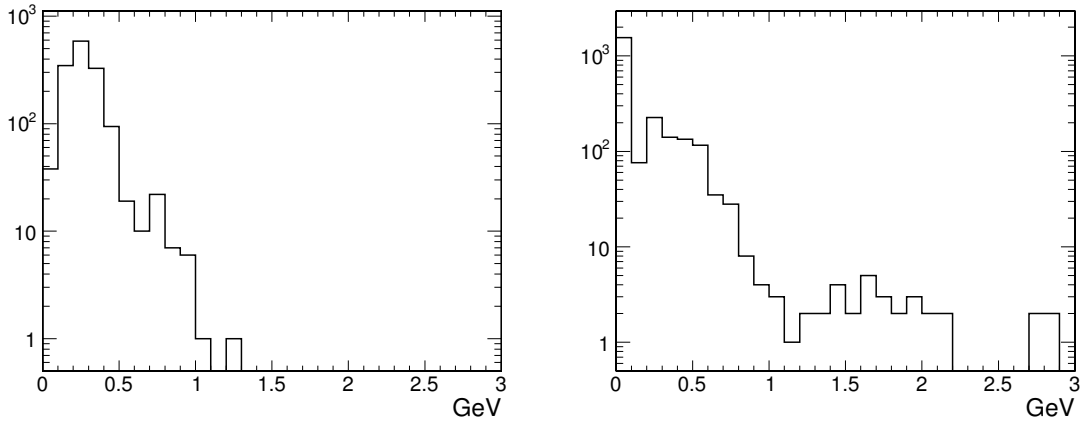


Figure 9.17: Electromagnetic calorimeter energy associated with single $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ muon tracks as returned by the muon identification algorithm for the barrel (left) and endcap (right) regions of the detector. The significant number of zeroes returned for muons in the endcap region is due to the high online thresholds applied in the readout of the endcap crystals.

Based on the measured energies associated with candidate tracks, the muon identification algorithm calculates a compatibility value between 0 and 1 that attempts to describe how consistent these energies are with respect to what one expects for a muon. The compatibility value is obtained from a three-dimensional likelihood function of the form

$$\frac{P_S(x) \cdot P_S(y) \cdot P_S(z)}{P_S(x) \cdot P_S(y) \cdot P_S(z) + P_B(x) \cdot P_B(y) \cdot P_B(z)}\ , \tag{9.2}$$

where P_S and P_B are the signal and background probabilities as functions of the measured energies in the electromagnetic calorimeter (x), the hadron calorimeter (y), and the HO scintillator layer (z). The signal and background probability distributions are obtained from simulated samples of single-muons and pions, respectively. Independent distributions are obtained for the barrel, overlap, and endcap regions of the detector and for different track p_{T} ranges. Because there is no HO scintillator layer in the endcap region of the detector, the $P_S(z)$ and $P_B(z)$ functions for this region are set to one. The resulting compatibility values for tracks associated with simulated $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ single muons and pions are shown in Fig. 9.18 for the barrel and endcap regions of the detector.

9.2.1.2 Muon detector identification

Reconstructed hit information from the muon detectors can be used to identify muon candidates even in cases where the standalone muon reconstruction fails. Low-p_{T} muons, in particular, can be bent back in the magnetic return field toward the interaction region or range out in the magnet yoke. In these cases, no hits are produced in the outermost layers of the muon detectors, and consequently standalone reconstruction is less efficient. Conversely,

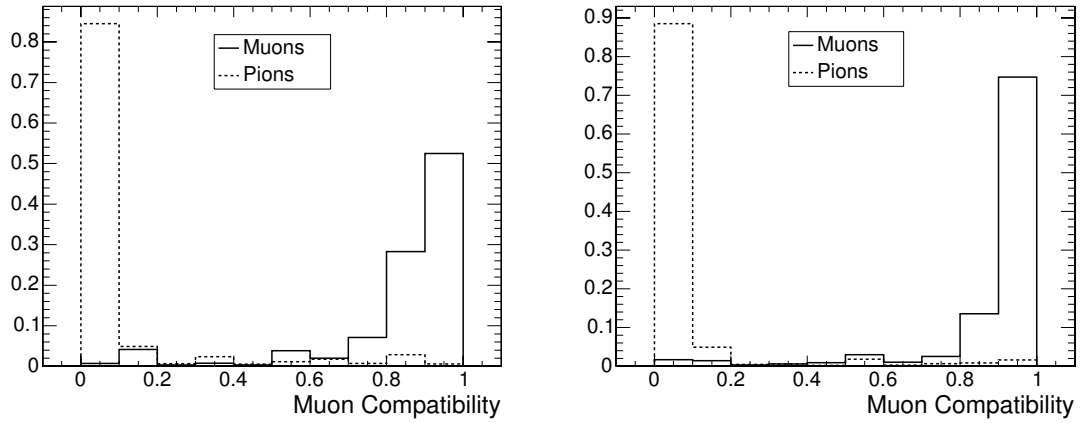


Figure 9.18: Muon compatibility values for tracks in the silicon tracker associated with simulated $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ single muons (solid lines) and pions (dashed lines) in the barrel (left) and endcap (right) regions of the detector. The compatibility values shown here are based solely on observed energy deposits in the calorimeters matched with the candidate track.

muon detector hits coming from even a single inner layer that can be associated with a track reconstructed in the silicon tracker (thereafter simply called a "track") are potentially helpful in determining the muon compatibility of that track.

To incorporate muon detector hit information into the muon identification algorithm, tracks are extrapolated sequentially into each layer of the muon detector. A search road is defined to look for hits and reconstructed segments within the layer that may be compatible with the extrapolated trajectory. For each potential matching hit or segment, the algorithm calculates a χ^2 measuring the compatibility of the position and direction (for segments) with those of the extrapolated track at the corresponding muon detector layer. The algorithm defines hits and segments to be associated with the track if the corresponding χ^2 is below a programmable threshold. Using this definition, the algorithm outputs a list of matching hits and segments for each reconstructed track.

Using this list, the identification algorithm calculates a second muon compatibility value for each track based solely on muon detector hit information. The weight assigned to a single matching segment or hit depends on the dimensional content. Drift tube hits, for example, provide matching information in only 1 dimension. Drift tube chamber segments, on the other hand, contain matching information in 2 or even 4 dimensions if hits from all superlayers are incorporated. CSC and RPC hits provide slightly better matching in 2 dimensions, while CSC segments also provide ideal matching in all 4 dimensions. Matching hits and segments are also weighted by layer. Outer-layer hits and segments are more likely to be associated with a real muon since more material must be traversed to reach these layers, while inner-layer hits are more likely to be associated with punch-through of hadrons.

For each track, the algorithm assigns a compatibility value between 0 and 1 based on the number of matching muon detector segments and hits. The calculation is designed so that tracks matching to the best possible combination of hits and/or segments in each of the muon detector layers have a corresponding compatibility value of one. Of course, based on this definition, tracks with an associated compatibility value near 1 are almost sure to be globally reconstructed as muons. In the case of muon identification, therefore, cuts on the

compatibility value designed to recover muon reconstruction inefficiencies will need to be set well below this maximum value. Details of the algorithm are described in Ref. [237].

Based on the algorithm described above, the muon compatibility value distributions for tracks associated with simulated $p_{\mathrm{T}} = 10\,\mathrm{GeV}/c$ single muons and pions are shown in Fig. 9.19 for the barrel and endcap regions of the detector.

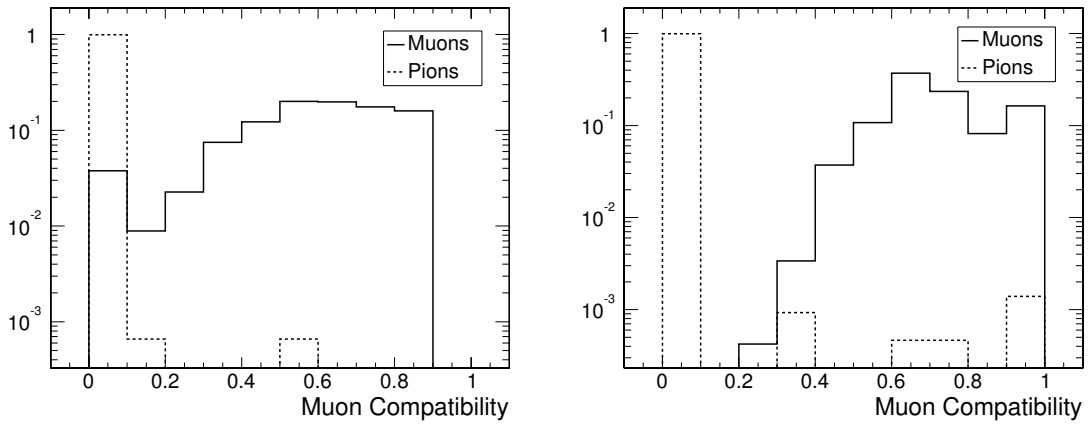


Figure 9.19: Muon compatibility values for tracks associated with simulated $p_{\mathrm{T}} = 10\,\mathrm{GeV}/c$ single muons (solid lines) and pions (dashed lines) in the barrel (left) and endcap (right) regions of the detector. The compatibility values shown here are based solely on matched hits and segments in the muon detectors.

9.2.2 Performance

The performance of the muon identification algorithm was studied on several simulated event samples. For these studies, a track is considered to be identified as a muon if its associated compatibility value based on energy deposits in the calorimeters is above 0.8 and its associated compatibility value based on hit and segment information in the muon detectors is greater than 0.4. These cut values are programmable and in general require optimization for specific algorithm applications. Here, the same cut values are used on several different types of event samples to provide a baseline overview of algorithm performance. Low-p_{T} muons are of particular interest since in a significant fraction of cases these muons do not produce hits in the outer muon detector layers and are therefore more difficult to reconstruct using the standard standalone algorithm. The algorithm was run on a sample of $50\,000$ simulated single-muon events. The muons were generated flat in ϕ and η (from -2.5 to 2.5) with a single-p_{T} value of $5\,\mathrm{GeV}/c$. The single muons were generated on top of simulated pile-up events corresponding to a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The efficiency distributions for both global muon reconstruction and global muon reconstruction plus muon identification are shown in Fig. 9.20 (left), as a function of the simulated muon pseudorapidity. For the events in this sample, the efficiency for reconstructing single, low-p_{T} muons is 68.7%. Adding muon identification as an additional selection mechanism increases the overall selection efficiency for the muons in this sample by roughly 15% (for a combined efficiency of 78.6%). The muon identification algorithm was also run on all of the non-muon tracks in these events to investigate the rate at which the algorithm misidentified these tracks as muon candidates. The fake rate as a function of track η is shown in Fig. 9.20 (right). Only

30 out of over $250\,000$ non-muon tracks in these events are identified as muon candidates by the algorithm.

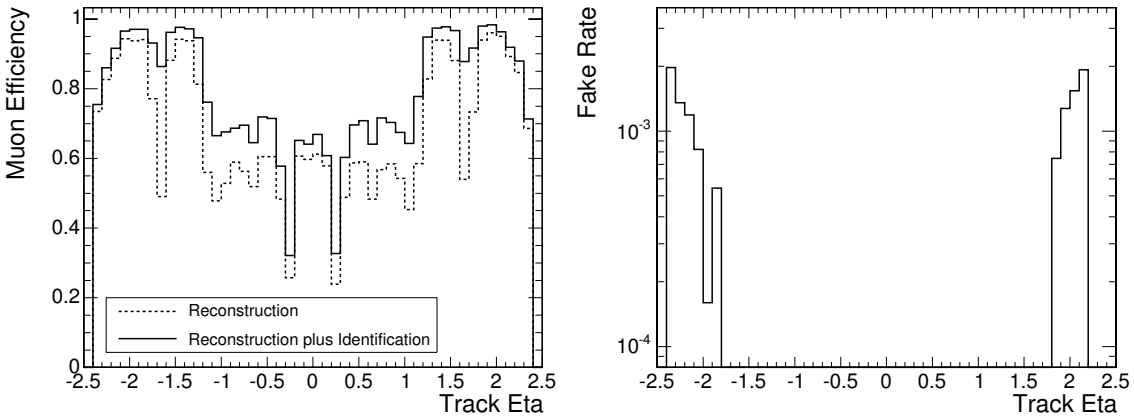


Figure 9.20: Efficiencies for global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines), as a function of pseudorapidity for $p_{\mathrm{T}} = 5\ \mathrm{GeV}/c$ single muons (left). Fake rate at which non-muon tracks in the same events are tagged as muon candidates using the muon identification algorithm also as a function of pseudorapidity (right). The observation of fake rates equal to 0 in the central part of the detector ($|\eta| < 1.8$) is due to insufficient statistics.

The muon identification algorithm was also run on a sample of simulated $H \rightarrow WW \rightarrow \mu^+\mu^-\nu\overline{\nu}$ events generated using a Higgs boson mass of $200\ \mathrm{GeV}/c^2$. The events include pile-up events corresponding to a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Due to the low overall production cross section for this process, observing a signal in the data requires maintaining high event selection efficiency. Since the transverse momentum of the lower-p_{T} muon is observed to be less than $10\ \mathrm{GeV}/c$ in roughly 20% of these events, muon identification can potentially be a useful tool for increasing the overall selection efficiency in this channel. Because the algorithm has not yet been optimized for selecting very low-p_{T} muons, this study is restricted to the subset of events where the p_{T} of both muon candidates is greater than $5\ \mathrm{GeV}/c$. The distributions in Fig. 9.21 (left) show the selection efficiency for events in this channel as a function of the simulated p_{T} of the lower-p_{T} muon, using both global reconstruction and global reconstruction plus muon identification. From these distributions, a net gain on the order of 5% is observed in the overall selection efficiency incorporating the muon identification algorithm. The rate at which non-muon tracks in these events are tagged as muons by the identification algorithm was also measured. The fake rate as a function of track p_{T} is shown in Fig. 9.21 (right). Out of over $750\,000$ non-muon tracks in these events, less than 125 are identified as muon candidates using the algorithm.

Low-p_{T} muons are also produced in the semi-leptonic decays of heavy quarks. The tagging of muons within jets (soft-lepton tagging) is therefore a useful tool in selecting jets originating from bottom quark decays (an important requirement for many analyses, and discussed further in Section 12.2.4). In order to see if the muon identification tool might be useful in improving the efficiency for soft-lepton tagging in bottom quark jets, the algorithm was run on a simulated sample of bottom quark jets with generated p_{T} values between 50 and $80\ \mathrm{GeV}/c^2$. The events were generated on top of simulated pile-up events corresponding to a

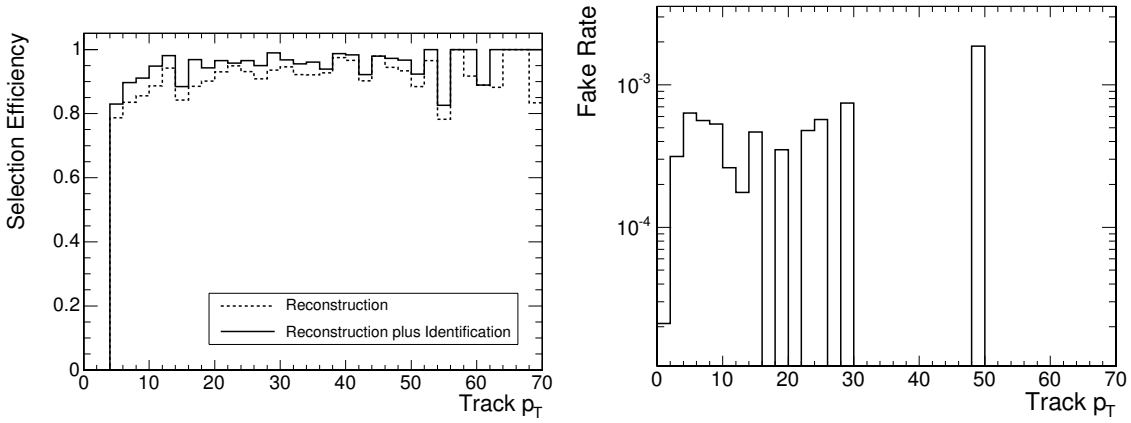


Figure 9.21: $H \to WW \to \mu^+\mu^-\nu\overline{\nu}$ event selection efficiencies using global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines), as a function of the p_{T} of the lower-p_{T} muon in the events (left). Fake rate at which non-muon tracks in the same events are tagged as muon candidates using the muon identification algorithm as a function of track p_{T} (right). The observation of fake rates equal to 0 for the majority of p_{T} values above 15 GeV/c is due to insufficient statistics.

luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The distributions shown in Fig. 9.22 are efficiency curves for both global muon reconstruction and global muon reconstruction plus muon identification for the muons produced in the semi-leptonic decays of the bottom quarks in these events, as functions of the simulated muon p_{T} and pseudorapidity. The efficiencies shown here are for muons with a minimum p_{T} of 5 GeV/c. The study is restricted to this range since the muon identification algorithm has not yet been optimized for selecting very low-p_{T} muons. The global reconstruction efficiency for these muons is 71% and incorporating the muon identification algorithm increases the overall selection efficiency to 84%. To be useful in soft-lepton tagging, the associated fake rate of the algorithm (the rate at which non-muon tracks within jets are identified as muons) must also be low since jets originating from both light and heavy quarks, as well as those from gluons, contain many non-muon tracks. The algorithm fake rate is studied using a sample of reconstructed pion tracks found within the subset of bottom quark jets from our simulated sample that contain no prompt muons. The efficiencies for tagging these pion tracks as muons using global muon reconstruction and global muon reconstruction plus muon identification are shown in Fig. 9.23 as functions of the simulated pion p_{T} and pseudorapidity. The measured fake rate at which the pion tracks within bottom quark jets are globally reconstructed as muons is 0.17%. Adding muon identification as an additional selection tool, increases the muon fake rate for pion tracks to roughly 0.20%. The small increase in the observed fake rate indicates that the muon identification tool can be potentially useful in soft lepton bottom quark tagging.

9.3 Muon isolation

9.3.1 Principles of the muon-isolation algorithms

Muons selected by a transverse-momentum cut in the muon trigger are mostly real muons. In the momentum range relevant for triggering (p_{T} threshold in the range 10–30 GeV/c), the

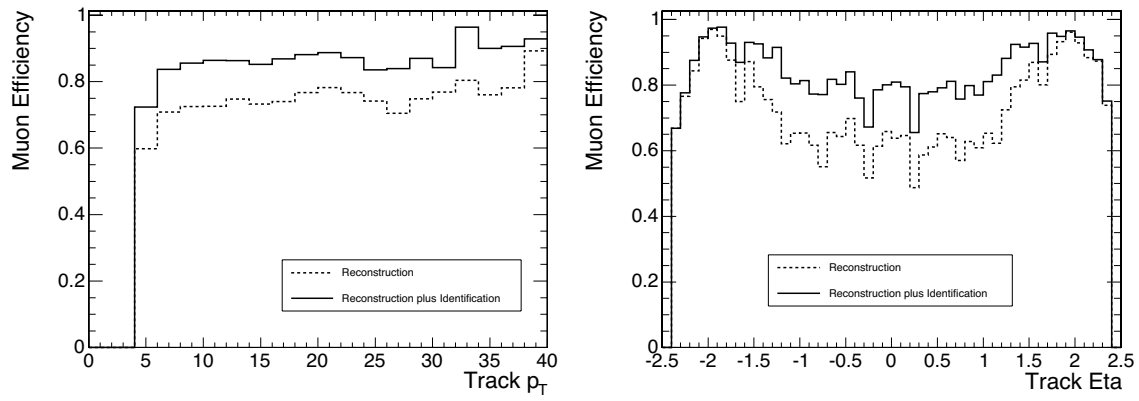


Figure 9.22: Efficiencies for global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines) for prompt muons contained within bottom quark jets, as a function of muon p_{T} (left) and pseudorapidity (right).

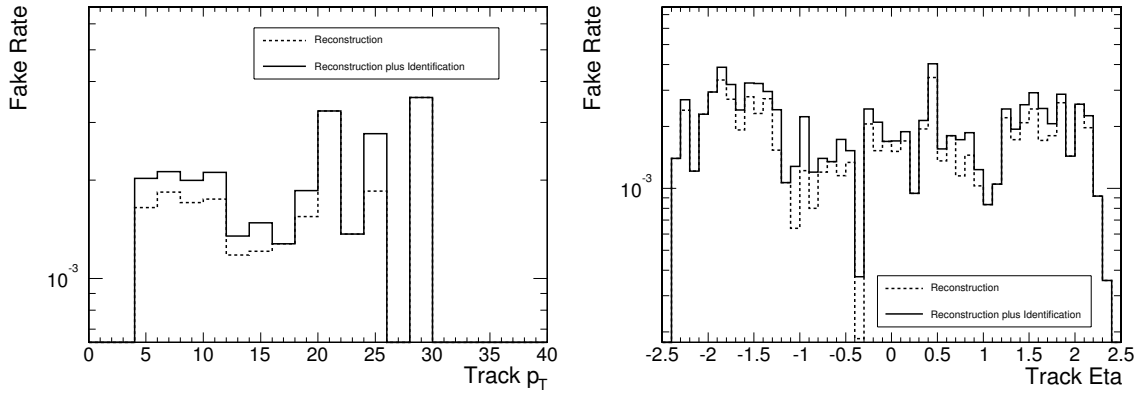


Figure 9.23: Fake rates for global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines) for pions contained within bottom quark jets, as a function of pion p_{T} (left) and pseudorapidity (right).

main sources of muons are from b and c decays. Another important contribution, mostly for low-p_{T} muons, is given by muonic K and π decays. All of these muons are produced in (usually soft) jets and are thus accompanied by nearby particles. Only for p_{T} above approximately $30\ \mathrm{GeV}/c$ do muons from W and Z decays become dominant. Muons from W, Z, and other heavy objects are accompanied only by particles from pile-up and by uncorrelated particles from the underlying event. Muon isolation is a tool to distinguish between the muons produced in jets and those coming from the decays of heavy objects.

The isolation algorithms that have been developed rely on the comparison of the total energy deposited in a cone around the muon with a predefined threshold. The deposit can be transverse energy in a calorimeter or the sum of transverse momenta of reconstructed charged-particle tracks. The cone axis is chosen according to the muon direction with a procedure that is tailored to the specific properties of each algorithm. The geometrical definition of the cone is given by the condition $\Delta R \leq \Delta R_{\mathrm{MAX}}$, where $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$, with $\Delta\eta$ and $\Delta\phi$ being the distances in pseudorapidity and azimuthal angle between the deposit and the cone axis, respectively. The muon itself contributes to the energy measurement inside the cone. This contribution (called the veto value) can be subtracted to improve the discriminating power of the isolation algorithm. A schematic illustration of the isolation cone is shown in Fig. 9.24.

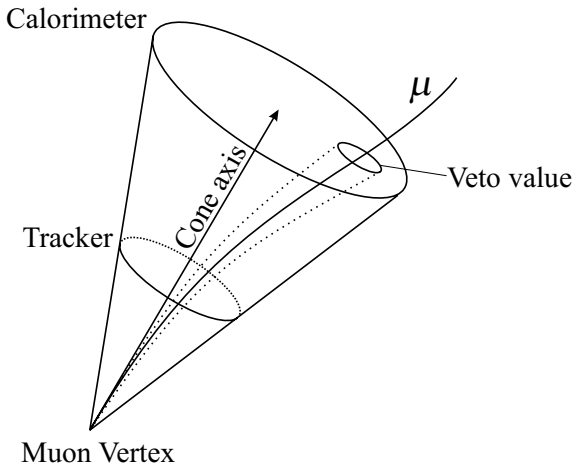


Figure 9.24: Schematic illustration of the isolation cone. The muon direction at the vertex defines the cone axis. The energy deposit ($\sum p_{\mathrm{T}}, \sum E_T^{\mathrm{weighted}}$) in the cone is computed, and the muon contribution is removed by excluding the small area around the muon (the "veto value") from the cone. Comparison of the deposit in the cone with a predefined threshold determines the muon isolation.

9.3.2 Online event selection

For the purposes of the High-Level Trigger, 3 isolation techniques were developed. Their detailed description can be found in [8, 238].

9.3.2.1 Calorimeter isolation

The calorimeter isolation is based on a standard technique of summing the calorimeter transverse energy in a cone around the muon direction. This method can be used with the stan-

dalone reconstruction at Level-2 and with fully-reconstructed Level-3 muons. The energy deposit in the cone is defined as the weighted sum of the transverse electromagnetic ($E_{\mathrm{T}}^{\mathrm{ECAL}}$) and hadron calorimeter energies ($E_{\mathrm{T}}^{\mathrm{HCAL}}$), using $E_{\mathrm{T}} = \alpha E_{\mathrm{T}}^{\mathrm{ECAL}} + E_{\mathrm{T}}^{\mathrm{HCAL}}$, with $\alpha = 1.5$, which reflects the better discrimination performance of the electromagnetic calorimeter. The veto cone is defined by the extrapolation of the muon to the ECAL/HCAL boundary (Fig. 9.25). Additional cuts against noise are used at the level of contributing crystals and towers. Typical energy-deposit distributions are shown in the left-most plots of Fig. 9.26.

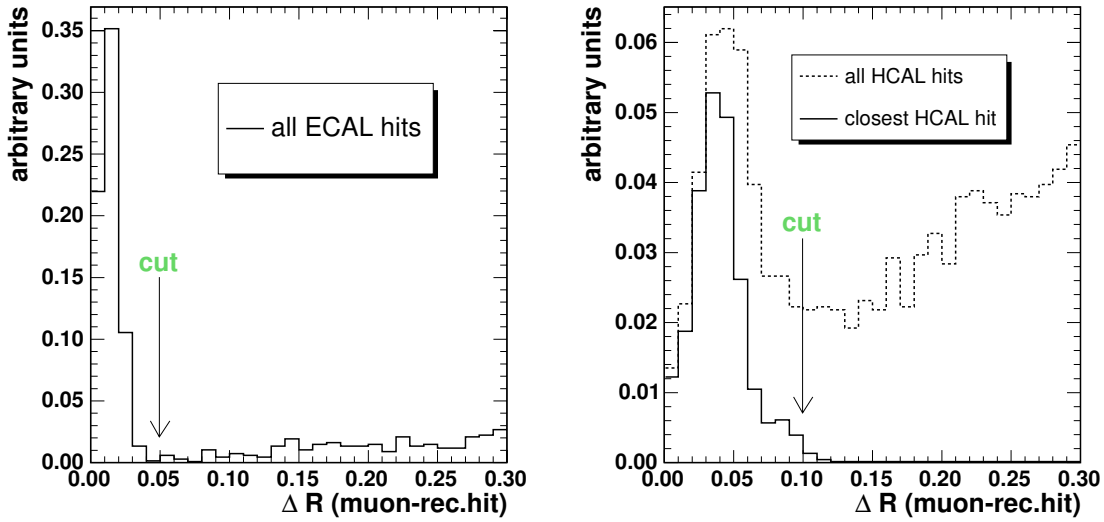


Figure 9.25: Distribution of the distance between a reconstructed muon and its energy deposit in the ECAL (left) and HCAL (right). The muon position is taken at the boundary between the ECAL and HCAL. The arrows indicate the size of the veto cone.

9.3.2.2 Pixel and full-tracker isolation

The pixel and full-tracker isolation algorithms are based on the scalar sum of p_{T} from all tracks reconstructed in a cone around the direction of the Level-3 muon at the vertex, neglecting the contribution from the muon itself. In both algorithms, the tracks are reconstructed using regional tracking.

In the case of the pixel isolation, the track candidates are searched for only in the pixel detector by dedicated pixel-reconstruction software. Such a reconstruction is based on 3-hit tracks. It is optimized for speed but, due to the small number of hits, has very poor momentum resolution and limited track-finding efficiency ($\approx 90\%$). For the full-tracker reconstruction, a Kalman-Filter fitting procedure is used at the price of longer CPU execution time. For both methods, the isolation cone defines the tracking region, with the vertex constraint coming from the isolated muon. Thus, the influence of tracks from pile-up is minimized. Since the direction resolution is good for both algorithms, the veto cone is very small ($\Delta R \approx 0.015$). A comparison of the energy-deposit distributions for the full-tracker and the pixel isolation is presented in the centre and right-most plots of Fig. 9.26, respectively.

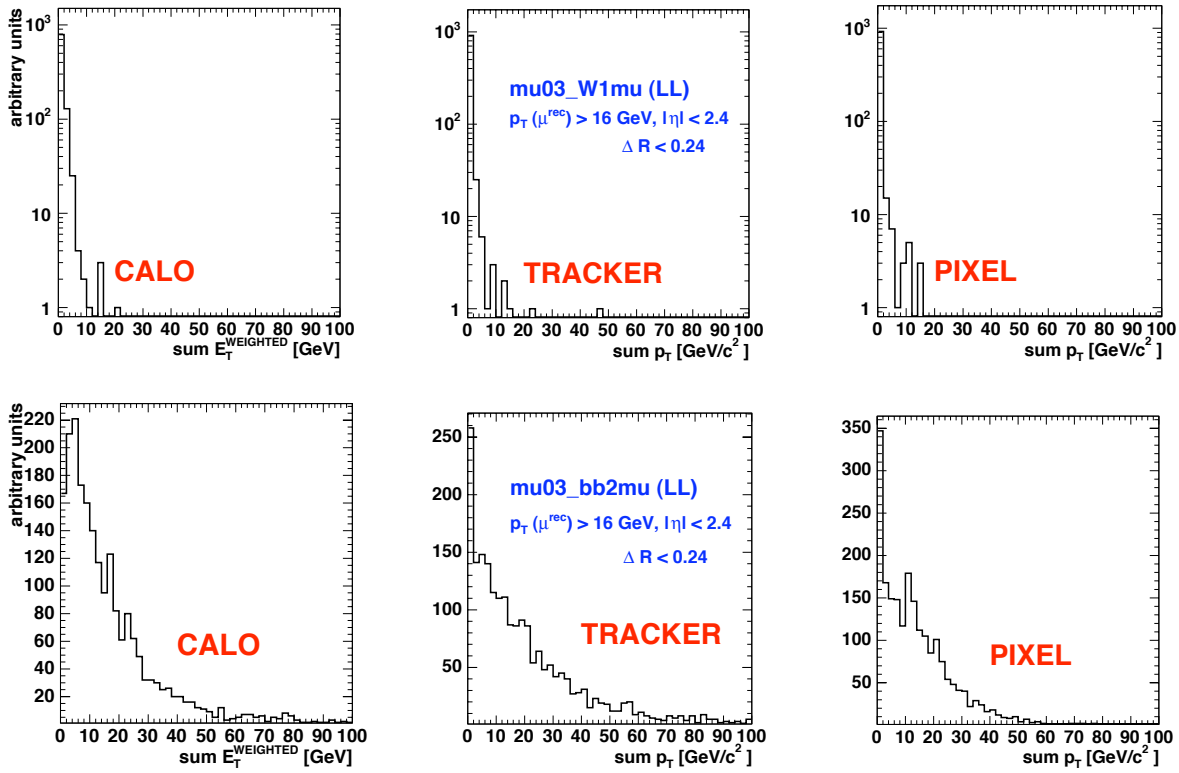


Figure 9.26: Distributions of the energy deposited in a $\Delta R = 0.24$ isolation cone for the calorimeter (left), the full tracker (centre), and the pixel tracker only (right). The difference in energy deposits between a typical source of isolated muons ($W \rightarrow \mu X$, top) and non-isolated muons ($b\bar{b} \rightarrow 2\mu X$, bottom) is well visible. The distributions are for the low-luminosity case.

9.3.2.3 Optimization and performance

The first step in the optimization of the isolation algorithm is the determination of the transverse energy and momentum thresholds as a function of cone size. Since some of the detector effects (noise, efficiency) may depend on pseudorapidity, the thresholds are evaluated in bins of η. The values of the thresholds are obtained using muons from $W \rightarrow \mu\nu$ decays (the "reference signal"), a typical source of isolated muons at the HLT level. For each combination of a cone size and an η bin, the threshold on the transverse energy or momentum in the cone is determined such that the efficiency of the reference signal is equal to some target efficiency (the "nominal efficiency"); however, for some choices of cone sizes, this is not always possible and the reference signal efficiency is somewhat larger than the chosen nominal efficiency.

The next step is the choice of a particular cone size (sizes considered are listed in Table 9.2). It depends on the background process that has to be rejected. For purposes of the HLT, $b\bar{b} \rightarrow \mu X$ events ("reference background") were used. The dependence of the background rejection on cone size is presented in Fig. 9.27 for several choices of the nominal efficiency. The efficiency of the HLT algorithms after the cone-size optimization is shown in Fig. 9.28 for the reference signal (left) and background (right) events, as a function of the nominal efficiency.

Table 9.2: The relationship between the cone-size indices and the ΔR values used in the optimization procedure.

Cone Index	ΔR_{MAX}	Cone Index	ΔR_{MAX}
1	0.02	8	0.28
2	0.045	9	0.32
3	0.09	10	0.38
4	0.13	11	0.45
5	0.17	12	0.5
6	0.2	13	0.6
7	0.24	14	0.7

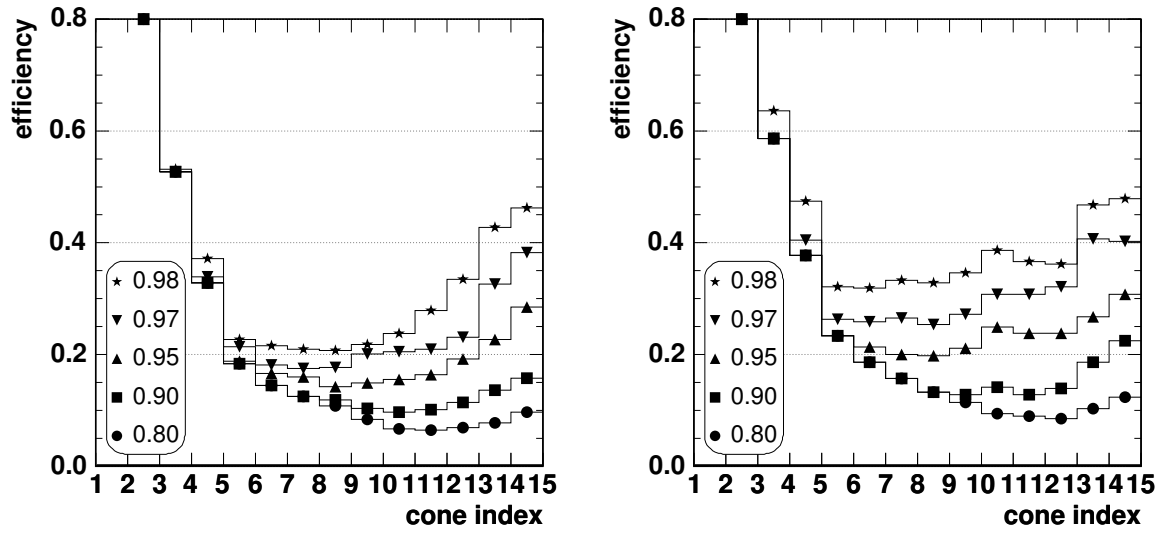


Figure 9.27: An example of an optimization of cone size: the efficiency for a background sample is shown as a function of the cone-size index for different predefined nominal efficiencies. The ΔR values corresponding to the different cone-size indices (horizontal axis) are given in Table 9.2. The left plot is for the full-tracker isolation in the low-luminosity case, the right plot is done with the pixel-based isolation for the high-luminosity environment.

9.3.3 Example of the isolation procedure applied offline

An offline algorithm should be optimized for a given signal and its most difficult background. For example, a calorimeter-based isolation requirement applied to high-p_{T} muons is bound to be inefficient because such muons are likely to radiate photons. Conversely, the isolation algorithm based on the distance between the muon and nearby jets does not efficiently reject background muons in soft jets that are likely to pass the HLT isolation. The aim of the study presented here is to show the possibility of further suppressing minimum-bias events that pass the HLT isolation chain by using additional offline isolation requirements. The HLT reference signal ($W \rightarrow \mu\nu X$) and background ($b\bar{b} \rightarrow \mu X$) processes are used, as before, for optimization.

The goal of the implemented offline algorithm is to combine precisely reconstructed charged particles measured in the tracker with neutral particles measured by the calorimeters. Since CPU processing time is less of an issue, the algorithms can benefit from global reconstruction in order to reduce pile-up contributions to the isolation cone.

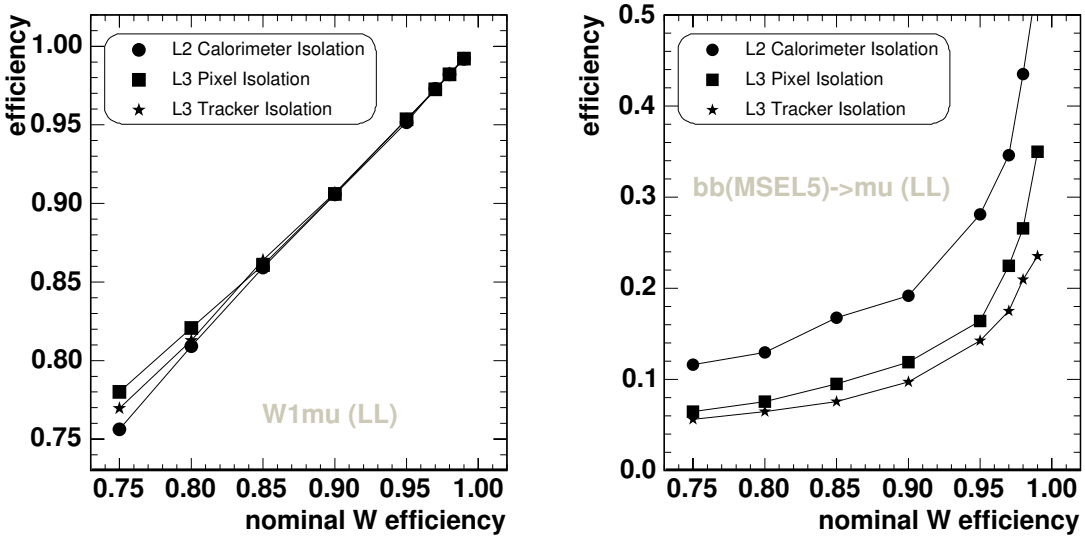


Figure 9.28: Comparison of the performance of 3 types of isolation algorithms at the HLT level. The isolation efficiency is plotted as a function of the nominal efficiency for signal events (left) and background events (right).

Particle trajectories from the global track reconstruction are propagated to the calorimeters. The muon energy deposit is identified and removed (by setting a veto cone) and is not used further in the computation. The deposit in the isolation cone is taken as the sum of the transverse momenta of reconstructed charged particles and the calorimeter transverse energy, both inside the isolation cone, from which the charged-particle deposit from the global reconstruction is subtracted. The distribution of deposits is shown in Fig. 9.29. The subtraction of the charged particle contribution from the calorimeter energy deposit is illustrated in Fig. 9.30.

The procedure to find the optimal threshold and cone size is the same as for the HLT. The results are shown in Fig. 9.31. The best background rejection that preserves 97% of the signal events (after already passing the HLT selection) is for a cone size of 0.28. A rejection factor of approximately 2 is obtained.

9.3.4 Sensitivity of the muon isolation cut efficiency to the underlying event uncertainties

The uncertainties in predicting the muon isolation cut efficiency were studied in the context of the $H \rightarrow ZZ \rightarrow 4\mu$ analysis by varying the PYTHIA parameters responsible for the simulation of the underlying event (UE) [239]. The isolation cut efficiency per muon can vary as much as $\pm 5\%$ due to uncertainties in the UE model (depending strongly on the energy deposit threshold), as shown in Fig. 9.32 (left). This results in a 4-muon isolation cut efficiency for the $ZZ \rightarrow 4\mu$ background of $(78 \pm 6)\%$. To eliminate these large uncertainties, the isolation cut efficiency can be calibrated from data using Z-inclusive events ($Z \rightarrow 2\mu$) and a random-cone technique, as displayed in Fig. 9.32 (right). This procedure eliminates the uncertainties associated with a poor understanding of the UE physics. There might be at most a $\approx 2\%$ systematic shift in the 4-muon isolation cut efficiencies using this method. This shift could be corrected for in principle, but it does not appear to be necessary since

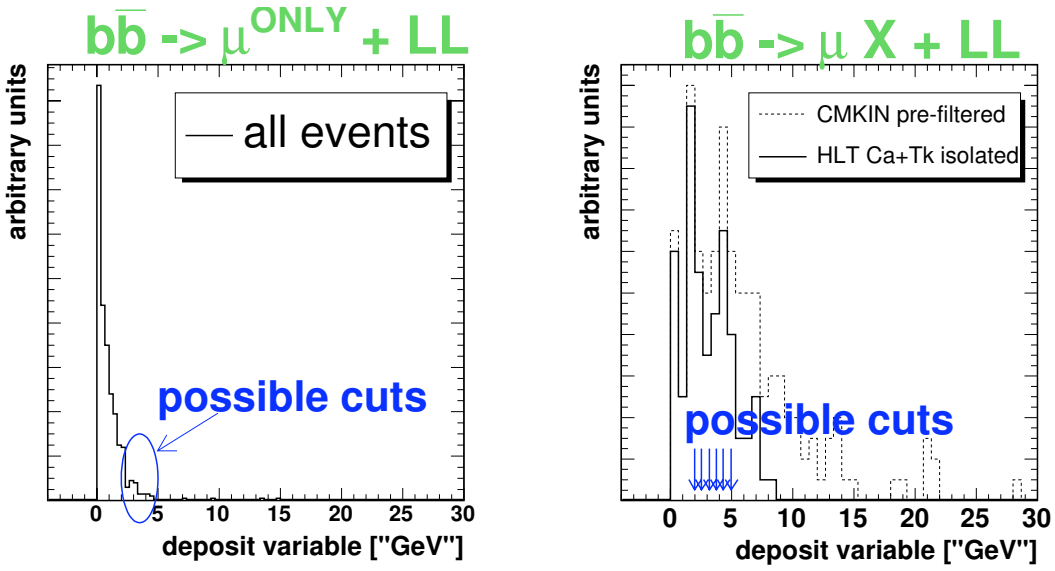


Figure 9.29: Energy deposits found offline in a $\Delta R = 0.24$ cone for isolated muons (left) and muons in $b\bar{b}$ jets (right). Only events passing the online isolation cut are used. Possible values for the energy deposit cut are marked on both plots.

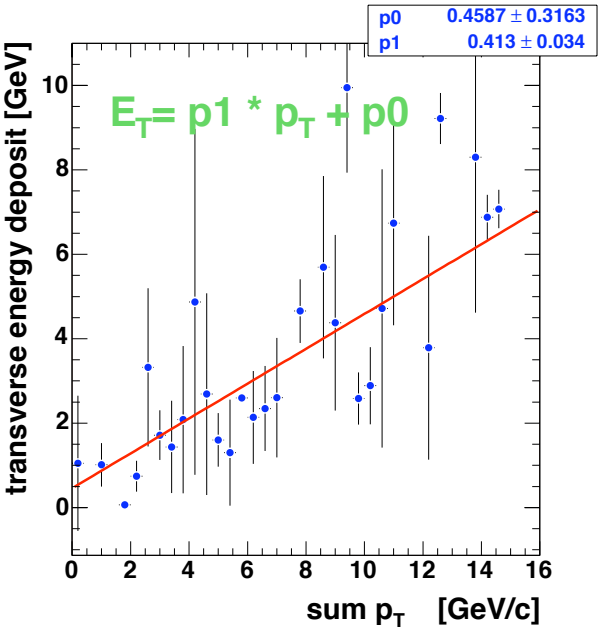


Figure 9.30: The correspondence between the measured transverse energy (vertical axis) and the total transverse momenta of all charged particles (horizontal axis) contributing to a deposit in the calorimeters. For the purpose of this figure, only charged particles were simulated. The cone size used was 0.24.

it is already small compared to the expected statistical uncertainties. These results and the analysis technique may be of interest to other analyses relying on lepton isolation cuts.

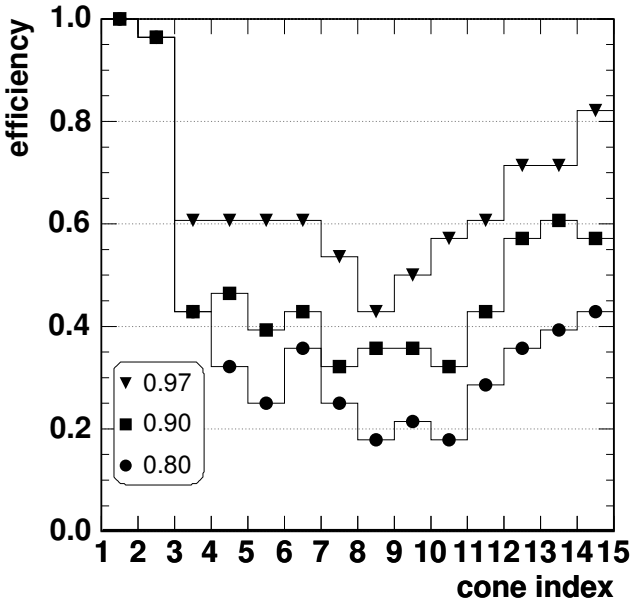


Figure 9.31: Performance of an offline isolation algorithm, in addition to the online selection. The efficiency for the background sample is shown as a function of the cone index for predefined nominal efficiencies in the low-luminosity case. The ΔR value for each cone index (horizontal axis) is given in Table 9.2.

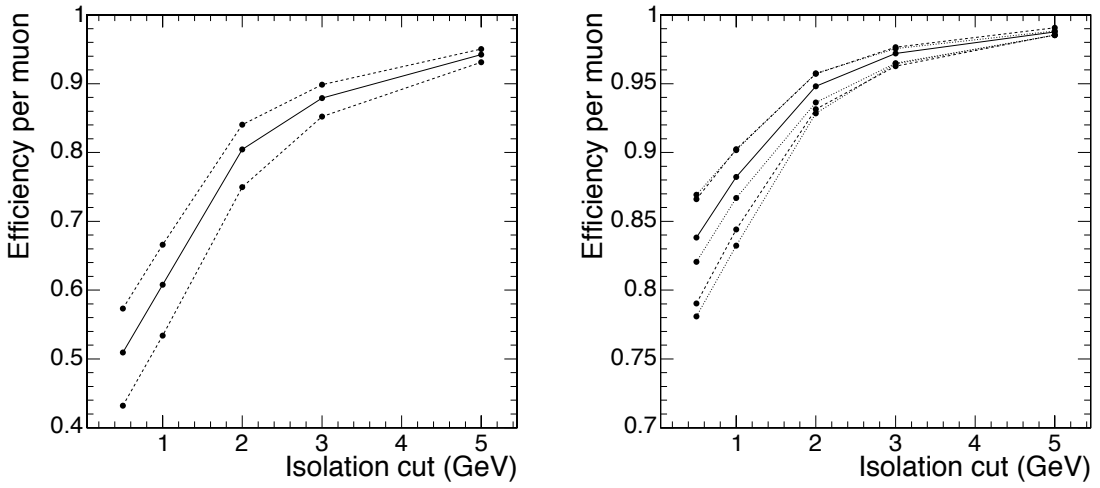


Figure 9.32: Left: the muon isolation cut efficiency for the least-isolated muon from $H \to 4\mu$ events. The middle solid line is for the default Multiple Interaction (MI) p_{T} "cut-off" value, the upper dashed line is for a -3σ variation in the p_{T} cut-off value, and the lower dashed line is for a $+3\sigma$ variation. Right: the muon isolation cut efficiency found using a random-cone technique in ZZ (solid and dashed lines) and Z-inclusive (dotted lines) events. The meaning of the middle, upper and lower lines is the same as for the left plot.

9.4 High-Level trigger

A large fraction of the expected LHC physics is characterized by the presence of high-p_{T} leptons, and CMS has been designed and optimized, in particular, to reconstruct muons. Muons provide a very clean signature; very high efficiency can be reached with inclusive-muon triggers based on the selection of 1 or more muons above a given p_{T} threshold. While this threshold should be low enough not to lose efficiency, it must be high enough to reduce the trigger rate to an acceptable level.

The online selection of muons within the High-Level Trigger (HLT) system is performed in 2 steps. First, to confirm the Level-1 decision, a Level-2 muon reconstruction is performed, i.e., muons are reconstructed using only information from the muon system and with the Level-1 candidates as a seed. The details of this reconstruction algorithm are described in Section 9.1.1. Secondly, the Level-2 p_{T} threshold reduces the rate enough to then allow a Level-3 reconstruction using the full tracker information and seeded by the Level-2 candidates. The Level-3 muon reconstruction algorithm is described in Section 9.1.2. Additional rate reductions can be achieved, if needed, by requiring muons to be isolated (Section 9.3). Already after the Level-2 selection, the calorimeter isolation can be required (Section 9.3.2.1), and the tracker and pixel isolation can be required after the Level-3 selection.

After each reconstruction step, a set of selection cuts is applied. These selection cuts are different for single and dimuon triggers. The details of the selection cuts are described in [8]. The p_{T} thresholds are set to achieve a sufficient rate reduction. As discussed in [8], for a luminosity of $2\times10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ an inclusive single-muon trigger rate of 25 Hz can be achieved with a p_{T} threshold of $19\ \mathrm{GeV}/c$ (including isolation). For dimuon triggers, a typical p_{T} threshold is $7\ \mathrm{GeV}/c$, resulting in a rate of 4 Hz.

9.4.1 Performance

A detailed performance study of single and dimuon high-level triggers can be found in [8]. However, since the time this study was performed, the CMS simulation and reconstruction software has been significantly changed. Hence, some of the findings presented in [8] were cross-checked with the results obtained using the most recent version of the CMS simulation and reconstruction software. The overall efficiency for muons to pass the Level-1 through Level-3 single-muon trigger criteria cumulatively as a function of η is shown in Fig. 9.33. Muons were generated without pile-up and with flat distributions in the intervals $5 < p_{\mathrm{T}} < 100\ \mathrm{GeV}/c$ and $|\eta| < 2.1$. Figure 9.34 shows the single-muon HLT rates without applying muons isolation cuts for a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. A sample of 10^6 minimum-bias events was used to calculate these rates. The single-muon rate at Level-2 is still dominated by feed-through of low-p_{T} muons. The improved Level-3 momentum resolution allows these muons to be discarded. Despite the fact that completely different event samples and a different version of the CMS simulation and reconstruction software were used, the result are in agreement with those in [8].

9.4.2 High-Level trigger for muons in heavy-ion collisions

The High-Level Trigger will be used to select dimuons from quarkonia produced in heavy-ion collisions. The J/ψs and Υs produced with low p_{T} are expected to provide information on the state of the hot nuclear matter produced in these collision. The relatively low luminosity of $L = 10^{27}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ allows any low quality muon in the CSC trigger system (a muon

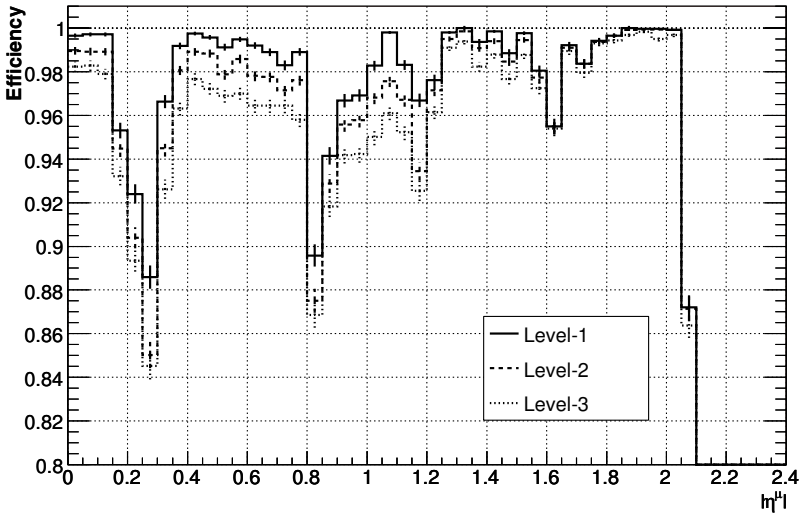


Figure 9.33: Cumulative efficiency for single muons to pass Level-1 (solid), Level-2 (dashed), and Level-3 (dot-dashed) trigger levels as a function of the generated muon pseudorapidity. No thresholds on p_{T} are applied. The vertical axis is zero-suppressed.

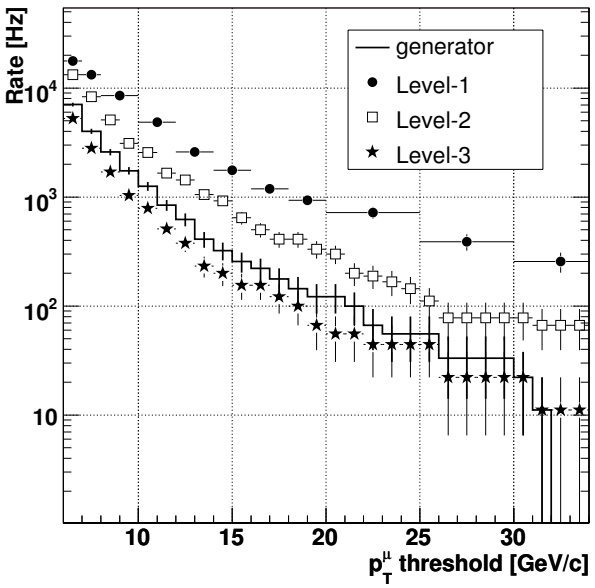


Figure 9.34: HLT single-muon rates after Level-1, Level-2, and Level-3 triggers, as a function of the p_{T} threshold for a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The rates are obtained without applying isolation cuts.

without a segment in ME1, as described in Section 3.4.3) to be sufficient to generate a Level-1 accept. Such muons passing the looser requirements at Level-1 are denoted hereafter "Open Level-1" muons (OL1), and are considered in the heavy-ion HLT trigger decision. The final trigger decision is delayed until the detailed reconstruction is performed at Level-2 and Level-3 of the HLT.

Different configurations of the Level-1 and Level-2 trigger flags were considered to maximize the selection efficiency for J/ψ and Υ. No cuts on any kinematic variables (p_{T}, η, ϕ) and no isolation cuts in the calorimeter are applied, apart from those caused by the detector acceptance.

The information obtained from Level-1 and Level-2 processing is used to estimate the muon momentum and its direction based on muon chamber information only. The Level-2 processing is applied on the output from Level-1 (including both "high quality" Level-1 muons and OL1 muons). If Level-2 fails, the information from Level-1 is used to estimate the muon momentum. For the case when Level-2 fails in the *barrel* and the momentum of the OL1 muon is undetermined, then the p_{T} of the OL1 muon is set to $3.5\ \mathrm{GeV}/c$. If Level-2 fails in the *endcap* and the momentum of the OL1 muon is undetermined, then the p_{T} of the OL1 is fixed to $1\ \mathrm{GeV}/c$. In both cases, the momentum components are recalculated according to the Level-1 muon direction.

The chosen strategy requires either 2 opposite-sign Level-1 or 2 opposite-sign Level-2 muon candidates.

The use of OL1 muons as the Level-2 input favors the survival of low-p_{T} J/ψs; the selection efficiency is $\approx 1\%$, as compared to 0.41% with the more stringent Level-1 usage without such low quality muons. For Υs, the acceptances are 21% for the OL1/Level-2 combination and 16% for the more stringent Level-1/Level-2 combination. The η and p_{T} distributions for accepted J/ψs and Υs are shown in Figs. 9.35, 9.36, 9.37 and 9.38. The results obtained using OL1 as input to Level-2 are shown with a solid-line histogram and those using Level-1 as input to Level-2 are shown with the shaded histogram.

While the trigger criteria are optimized for muons from decaying quarkonia, they will also be used to select same-sign dimuon pairs needed to estimate the uncorrelated background [240].

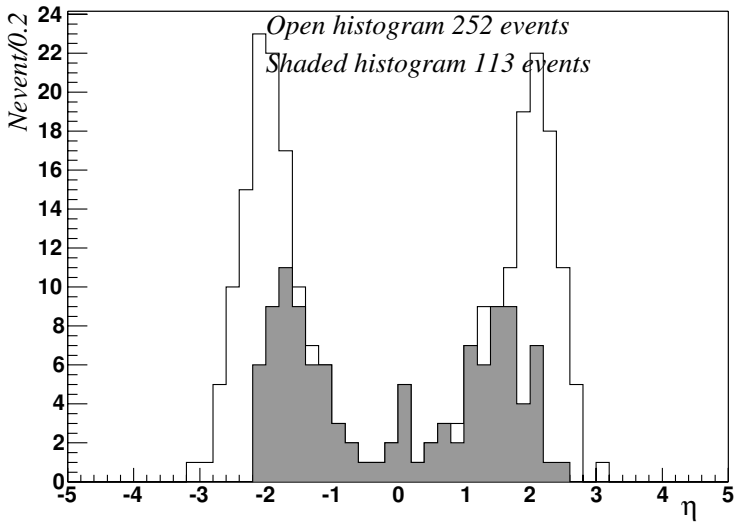


Figure 9.35: The η distribution of J/ψs with both muons triggered by the "Open Level-1" and Level-2 trigger combination (solid-line histogram) as compared with the more stringent Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required. See text for details.

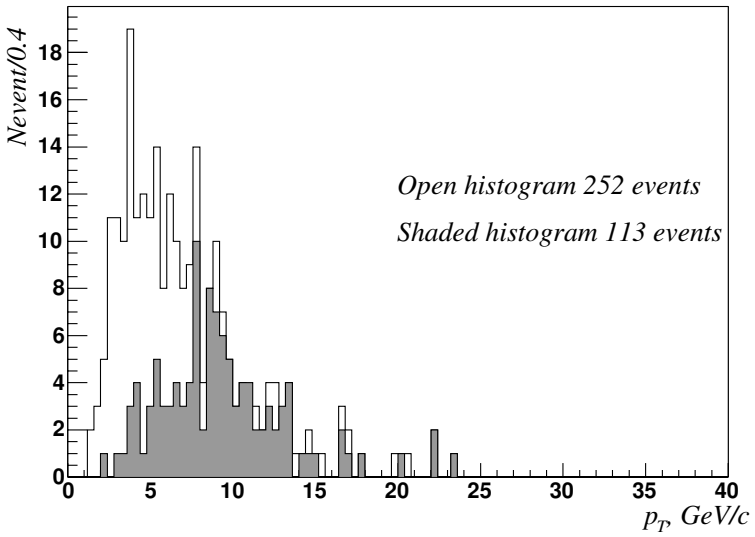


Figure 9.36: The p_{T} distribution of J/ψs with both muons triggered by the Open Level-1 and Level-2 trigger (solid-line histogram) as compared with the more stringent Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required.

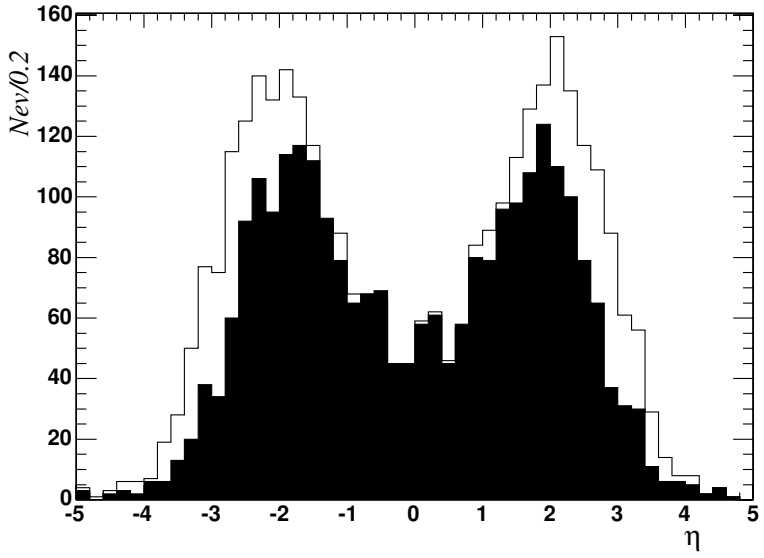


Figure 9.37: The η distribution of Υs with both muons triggered by the Open Level-1 and Level-2 trigger (solid-line histogram) as compared with the Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required.

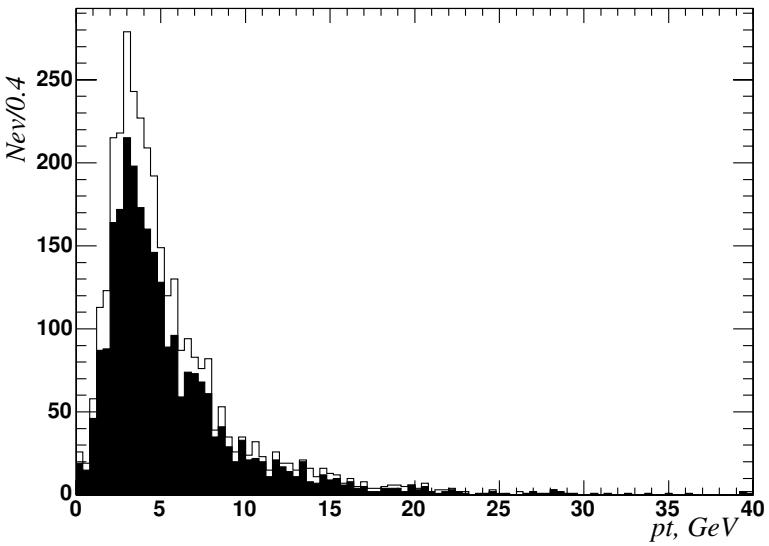


Figure 9.38: The p_{T} distribution of Υs with both muons triggered by the Open Level-1 and Level-2 trigger (solid-line histogram) as compared with the Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required.

Chapter 10

Electrons and Photons

10.1 ECAL clustering and superclustering

Electron and photon showers deposit their energy in several crystals in the ECAL. Approximately 94% of the incident energy of a single electron or photon is contained in 3×3 crystals, and 97% in 5×5 crystals. Summing the energy measured in such fixed arrays gives the best performance for unconverted photons, or for electrons in the test beam. The presence in CMS of material in front of the calorimeter results in bremsstrahlung and photon conversions. Because of the strong magnetic field the energy reaching the calorimeter is spread in ϕ. The spread energy is clustered by building a cluster of clusters, a "supercluster," which is extended in ϕ. Details of the superclustering algorithms used have been given in previous publications [8], but for convenience of reference they are reviewed briefly in the first subsection.

10.1.1 Superclustering algorithms

For single showers, such as those produced by unconverted photons, or those produced by electrons in test beam conditions, energy sums of fixed arrays of crystals give better energy resolution performance than energy sums of crystals collected dynamically according to a cluster or "bump" finding algorithm. The Hybrid algorithm uses the η-ϕ geometry of the barrel crystals to exploit the knowledge of the lateral shower shape in the η direction (taking a fixed bar of 3 or 5 crystals in η), while searching dynamically for separated energy in the ϕ direction.

The Hybrid algorithm was originally designed to reconstruct relatively high energy electrons in the barrel, but has been recently tuned to allow efficient reconstruction of electron showers down to $p_{\mathrm{T}} = 5\ \mathrm{GeV}/c$ (see Section 10.4.1). However when looking for small deposits of energy in individual clusters, for example when making a calorimetric isolation cut, the basic clusters of the Island algorithm are more appropriate objects to work with.

The Island algorithm starts by a search for seeds which are defined as crystals with an energy above a certain threshold. Starting from the seed position, adjacent crystals are examined, scanning first in ϕ and then in η. Along each scan line crystals are added to the cluster until a rise in energy or crystal that has not been read out is encountered. In much the same way as energy is clustered at the level of calorimeter cells or crystals, non-overlapping Island clusters can in turn be clustered into superclusters. The procedure is seeded by searching for the most energetic cluster and then collecting all the other nearby clusters in a very narrow η-window, and much wider ϕ-window.

10.1.2 Algorithmic energy corrections

There are some sources of variation in the clustered energy for which corrections can be made:

- The fraction of energy in a fixed array cluster varies as function of the shower position with respect to the cluster boundary—the highest fraction is contained if the shower is perfectly centred. This can be described as a variation of "local containment"
- There are large losses due to rear leakage for showers close to the barrel inter-module, and inter-supermodule, borders. At these borders there are cracks which contain negligible material, and which are aligned at 3° to the average direction of shower propagation. Close to these cracks there is a considerable considerable reduction of the effective depth of the calorimeter.
- The spread of energy due to showering in the tracker material (bremsstrahlung and photon conversion) and the behavior of the superclustering algorithms which are designed to collect this spread energy.

A further correction needed in the handling of simulated events results from the fact that the average fraction of energy of an incoming electron or photon contained in a fixed array of crystals (5×5, for example) varies a little in the barrel as function of η, despite the fact that the crystals are all of very similar size. This variation, of ≈0.7% over the length of the barrel is mainly due to a variation of shower leakage as a function of η. The stepped front surface of the ECAL towards the ends of the barrel exposes the sides of the crystals to a significant depth, leading to some lateral leakage from the front of the ECAL. As described in Section 4.4, the calibration constants are defined with a selected sample of non-radiating electrons measured in a 5×5 crystal array. Such calibration naturally includes this overall average containment factor.

The local containment variation can be corrected for by measuring the centrality of the cluster using the energy deposit pattern. A possible measure is the energy ratio E_1/E_9, where E_1 is the energy in the central crystal and E_9 the sum of energies in the central 3×3 crystal array. In test beam data it has been found that better results can be obtained by measuring the position separately in each lateral dimension. As described in Section 4.3.2.2 a logarithm of the ratio of energies on either side of a line defined by the crystal edge has been used for this purpose.

There are 18 supermodules in ϕ in each half barrel. These supermodules each comprise 4 modules. At both the ϕ boundaries between the supermodules and the η boundaries between the modules there is a 6 mm space between adjacent crystals. Particles entering the calorimeter near these borders encounter a reduced depth of material and deposit less of their energy in the crystals. A correction can be made based on the measurement of the shower position. The position measurement is made from $\log(E_1/E_2)$, where E_1 is the energy in one side of the border and E_2 on the other side.

The energy from an electron or converted photon area is deposited over a larger area than an unconverted photon. The clustering algorithms recover the energy by extending the clustering region in ϕ. A systematic variation of the fraction of energy clustered as a function of number of crystals in the cluster is observed. There is also a systematic variation of the fraction of energy reconstructed as a function of the shower spread measured by E_9/E_{SC} where E_9 is the energy in the central 3×3 array and E_{SC} the energy is the super cluster. More elabo-

rate corrections which can be applied to electrons and photons are described in Sections 10.4 and 10.3, respectively.

10.1.3 Position measurement

A simple position measurement of the shower can be obtained by calculating the energy-weighted mean position of the crystals in the cluster. To increase hermeticity the ECAL is designed so that the crystals do not point exactly to the nominal interaction vertex. Therefore, the lateral position of the crystal axis depends on depth, and it is necessary to define the depth to obtain the crystal position in the other (lateral) coordinates. The depth definition is parameterized as $A[B+\log(E)]$ where the parameters are different for electrons and photons as photons penetrate deeper in the crystals before showering starts.

Since the energy density decreases approximately exponentially with lateral distance from the shower core, the simple energy weighted mean position of the crystals in a cluster gives a position biased towards the core of the show. Much of the bias can be avoided, and a better position estimate obtained, by taking a weighted mean calculated using the logarithm of the crystal energy [241].

$$x = \frac{\Sigma x_i \cdot W_i}{\Sigma W_i} \tag{10.1}$$

where x_i is the position of crystal i and W_i is the weight of the crystal defined as

$$W_i = W_0 + \log \frac{E_i}{\Sigma E_j} \tag{10.2}$$

.

The position of a super cluster is calculated by making the energy-weighted mean of positions of its component clusters. More details are given in [160].

10.1.4 Corrections for saturated channels

Very energetic electrons and photons saturate the range of the Multi Gain Pre-Amplifier (MGPA) ADC. This saturation occurs at at 1.7 TeV in the barrel and 3.0 TeV in the end-caps. A number of possible techniques are available to reconstruct the energy of a shower containing a crystal for which saturation has occurred. Electron and photon single particle samples have been used to study a method of estimating the energy contained in the central, maximum energy, crystal (E_1) of a 5×5 cluster in the barrel [242]. The fraction of the cluster energy contained in the central crystal, E_1/E_{25}, varies as a function of the position of the shower with respect to the crystals.

The X and Y position variables are defined as the following:

$$X = \ln(\frac{\Sigma_{\text{right}}}{\Sigma_{\text{left}}}) \quad , \quad Y = \ln(\frac{\Sigma_{\text{top}}}{\Sigma_{\text{bottom}}}) \tag{10.3}$$

where Σ_{left} (Σ_{right}) is the sum of the energy deposit in the 10 crystals located at the left (right) side of the central crystal and Σ_{top} (Σ_{bottom}) in the 10 top (bottom) crystals. In each of the $|X|$ and $|Y|$ bins, the value of E_1/E_{25} is determined by taking the mean of a Gaussian fit. Figure 10.1a) shows these mean values as a function of $|X|$ and $|Y|$. A two-dimensional parameterization of this distribution is made using the function:

$$F(X,Y) = a + b(e^{cX^2} + e^{dY^2}) + eX^2 + fY^2. \tag{10.4}$$

The distribution is largely energy independent (see Fig. 10.1b)), and can be used to estimate the energy in saturated crystals.

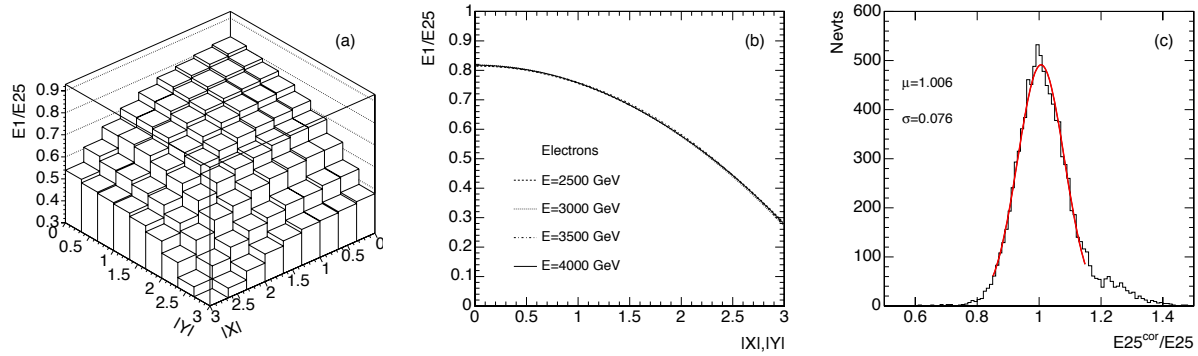


Figure 10.1: a) Distribution of the E_1/E_{25} ratio as a function of the variables $|X|$ and $|Y|$ for electrons with $E = 3500$ GeV; b) Fitted function (see Eq. 10.4) for E_1/E_{25} versus position, for $|X|$ = $|Y|$, for electrons with $E = 2000, 2500, 3000, 3500$, and 4000 GeV; c) E_{25}^{cor}/E_{25} for electrons with E =4000 GeV

Figure 10.1c) shows the estimated value E_{25} using the correction method for the saturation divided by the actual energy in E_{25}, for $E = 4000$ GeV electrons. The resolution of the method is found to be about 8% for electrons and 9% for photons.

10.2 Electron and photon trigger

This section gives a brief overview of electron and photon triggers. All events passing the Level-1 trigger are read out and processed in the High Level Trigger (HLT) farm. Four electron and photon streams are output to mass storage: single and double electrons and photons.

10.2.1 Level-1 trigger

Three electromagnetic triggers are used at Level-1. A single isolated trigger, a double isolated trigger and a double trigger at a slightly higher threshold with the trigger isolation requirement removed. The $E_{\rm T}$ thresholds used for these 3 triggers at $\mathcal{L} = 2 \times 10^{33}$ cm^{-2} s^{-1} are the same as those used for the studies reported in [8], and correspond to approximately 50% efficiency for electrons and photons with $E_{\rm T}$ of 23, 12 and 19 GeV respectively.

With respect to the details and rates given in [8] there has been only one significant change. For Level-1 triggers which saturate the calorimeter trigger $E_{\rm T}$ scale ($E_{\rm T} = 63.5$ GeV) none of the Level-1 trigger cuts (isolation, h/e, "fine grain veto" [7]) are applied, and the candidates are treated as if they were isolated. This strategy increases the efficiency for high $E_{\rm T}$ electrons and photons at a small cost in Level-1 trigger rate.

Figure 10.2 shows the Level-1 trigger rate, mainly due to jet events, and the efficiency for 120 GeV/c^2 mass Higgs events decaying into 4 electrons and 2 photons. In this plot and in all the following tables the signal efficiencies are computed for events where at least 2 of the generated electrons or the photons (one in the case of $W \rightarrow e\nu$ events) lie inside the fiducial

acceptance of the electromagnetic calorimeter ($|\eta|$ <1.4442 or 1.566< $|\eta|$ <2.5). In addition a lower cut, of 5 GeV/c, is applied on the p_{T} of the electrons and photons in case of the processes $H \rightarrow ZZ \rightarrow e^+e^-e^+e^-$ and $Z \rightarrow e^+e^-$, and a cut of 15 GeV/c for $H \rightarrow \gamma\gamma$, and 25 GeV/$c$ for $W \rightarrow e\nu$.

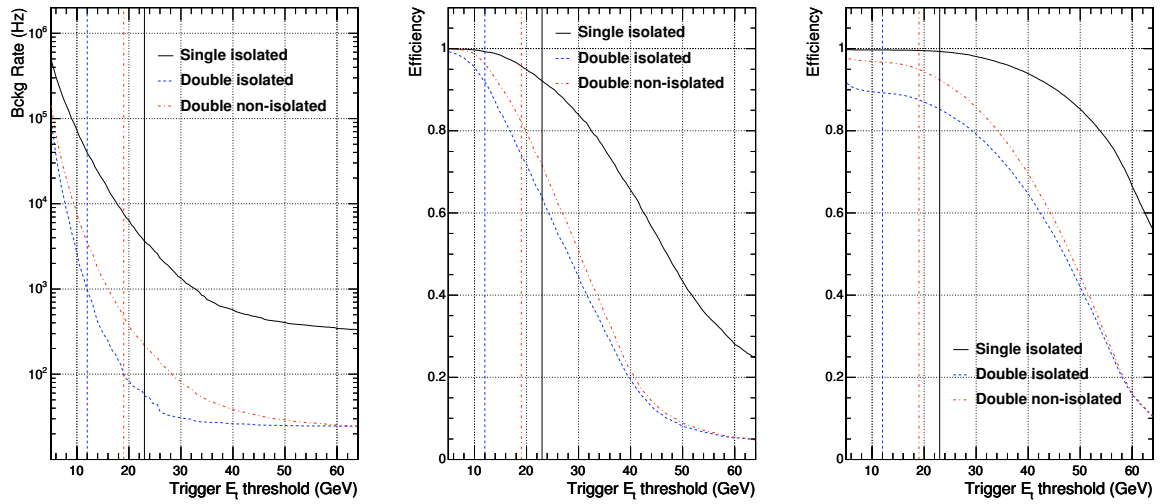


Figure 10.2: Level-1 trigger rate for jet events (left), efficiency for H→4e events (middle) and efficiency for $H \rightarrow \gamma\gamma$ events (right) as function of the trigger thresholds.

Table 10.1 shows the Level-1 trigger rates at low luminosity ($\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$). These rates, and all the HLT rates for the jet background to electrons and photons given in the following pages have been estimated for jet events simulated with PYTHIA [45] with $\hat{p}_{\mathrm{T}}$ larger than 20 GeV/c (where $\hat{p}_{\mathrm{T}}$ is the parton p_{T} in the $2 \rightarrow 2$ scattering process in PYTHIA).

Tables 10.2 shows Level-1 trigger efficiencies for benchmark photon and electron signal channels.

Table 10.1: Level-1 trigger rates at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Single isolated	Double isolated	Double non-isolated	Total
3400 Hz	1010 Hz	420 Hz	4440 Hz

Table 10.2: Level-1 trigger efficiency at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Signal Process	Single isolated	Double isolated	Double non-isolated	Total
$H \rightarrow \gamma\gamma$ ($M_H = 120$ GeV/c^2)	99.3%	89.2%	94.7%	99.7%
$H \rightarrow ZZ \rightarrow e^+e^-e^+e^-$ ($M_H = 120$ GeV/c^2)	90.8%	89.5%	79.5%	96.5%
$Z \rightarrow e^+e^-$	93.5%	81.5%	85.1%	97.1%
$W \rightarrow e\nu$	89.8%	2.7%	2.0%	90.0%

10.2.2 HLT selection of electron and photon candidates

The HLT selection of electrons and photons in the High Level Trigger proceeds in 3 steps. The first step, Level-2, uses only calorimeter information. The energy in the ECAL is clustered to form superclusters and an E_{T} threshold is applied on the corrected supercluster energy.

The second step, Level-2.5, uses the information from the pixel detectors. Hits are sought consistent with the hypothesis that the supercluster belongs to an electron. If the required hits are not found the candidate is rejected as an electron. Conversely all candidates are further considered for the photon triggers regardless of pixel information. In the final step, Level-3, the full tracker information is used. Tracks are used for isolation, and the electron tracks are matched in momentum and position with the electromagnetic supercluster.

Rates and efficiencies of the HLT reconstruction and selection are given in [8].

We recently carried out a further detailed optimization of the High Level Trigger. The changes with respect to the earlier study are the following:

- the strict restriction to the ECAL fiducial region has been relaxed, so that the region of the barrel/endcap transition (1.4442$<|\eta|<$1.5660) is no longer excluded; the corresponding trigger rate increase is of the order of a few per cent;
- the Level-2 threshold cuts have been lowered by about 2 GeV;
- the Level-1 single trigger is now allowed to populate the double electron and photon streams—double electron/photon triggers are allowed where 1 of the electrons/photons does not match a Level-1 trigger;
- the isolation cuts at Level-3 have been refined, giving higher efficiencies or lower rates; in the case of the double photon stream this allows a lowering of the thresholds.

10.2.2.1 Level-2 selection

The first step of the HLT selection consists in the verification of the Level-1 trigger decision and it is common for electrons and photons. The 50% efficiency E_{T} thresholds are applied to reconstructed supercluster matched in pseudo-rapidity and azimuthal angle with the triggering towers and the pseudo-rapidity of the supercluster is required to be less than 2.5. The E_{T} thresholds are 23, 12 and 19 GeV for the single isolated, double isolated and double non-isolated Level-1 triggers respectively.

10.2.2.2 HLT selection for electrons

The High level trigger selection for electrons is based on the following requirements:

- matching of the supercluster with pixel hits;
- requiring electron isolation;
- for the single trigger e/p matching between the energy reconstructed in the ECAL and the corresponding track momentum measured in the tracker.

The most effective isolation variables are found to be isolation in the hadron calorimeter and isolation in the tracker. In the hadron calorimeter, the energy collected in a cone around the supercluster is required to be below a certain threshold. The optimum cone size for HLT has

been found to be $\Delta R < 0.15$. For the single electron stream a maximum energy of 3 GeV is allowed within the cone region while for the double electron streams this cut is loosened to 9 GeV.

Track isolation is defined by the sum of $|p_{\mathrm{T}}|$ of the tracks that are found in the region between 2 cones around the electron candidate direction. An inner cone of 0.02 radius is used to exclude the electron track while for the outer cone a radius of 0.2 is used. Only tracks with a reconstructed $z-$impact parameter consistent with the electron candidate $z-$impact parameter ($|z^{e}_{ip} - z^{tk}_{ip}| < 0.1$ cm) and with a transverse energy above 1.5 GeV are considered. The $|p_{\mathrm{T}}|$ sum of of these tracks, divided by the electron p_{T}, is required to be less than 0.06 for the single stream and 0.4 for the double stream. In the case of the double stream triggers the cuts are applied on both electron candidates.

The effect of the hadron calorimeter isolation and of the track isolation in the benchmark channel $H \rightarrow ZZ \rightarrow e^{+}e^{-}e^{+}e^{-}$ for the double isolated stream are shown in Fig. 10.3.

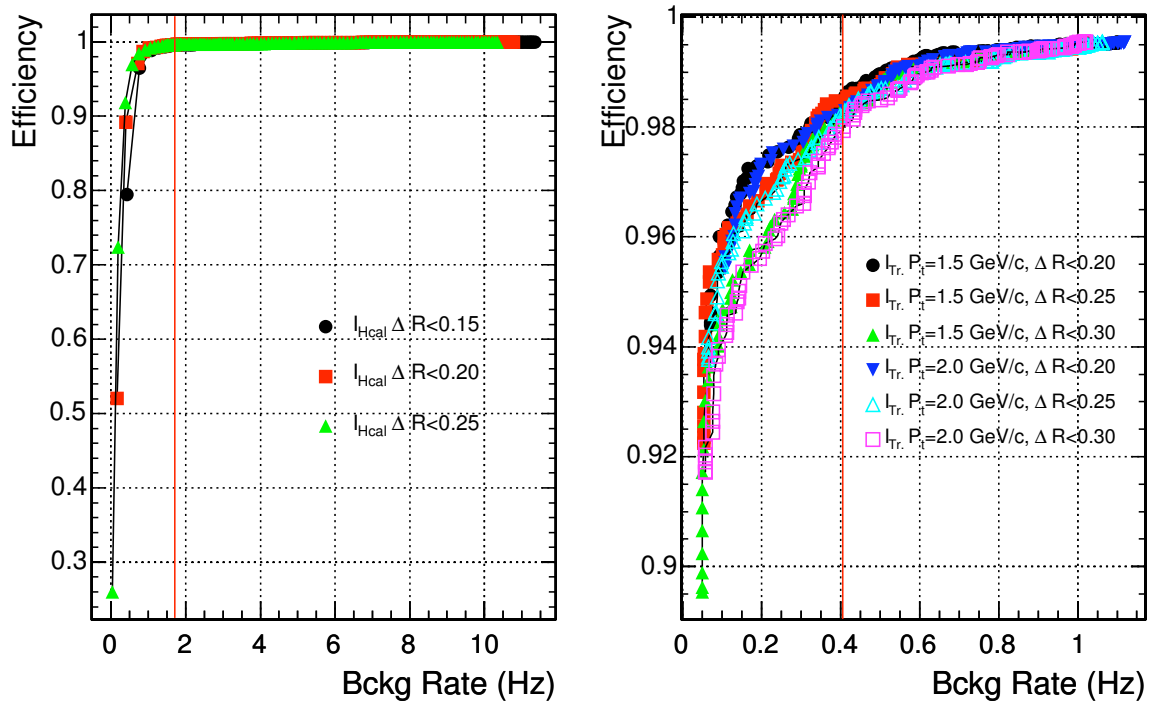


Figure 10.3: Effect of varying the HCAL isolation (left) and track isolation (right) cuts on the background rate and signal efficiency for the process $H \rightarrow ZZ \rightarrow e^{+}e^{-}e^{+}e^{-}$ with $M_H = 120$ GeV$/c^2$ in the double isolated electron stream. In the track isolation plot, an upper cut of 3 GeV on the hadron calorimeter energy within $\Delta R < 0.15$ is applied.

Finally in the single electron stream an additional cut on the E/p of the electron candidate is applied by requiring it to be less than 1.5 in the barrel and 2.45 in the endcaps.

A summary of the selection cuts applied in the HLT is given in Table 10.3.

Table 10.3: HLT cuts for the electron streams.

Variable	Single electron	Double electron
$\|\eta\|$	< 2.5	< 2.5
E_{T}	> 26 GeV	> 12 GeV
Hcal Isolation	< 3 GeV	< 9 GeV
Track Isolation	< 0.06	< 0.4
E/p (Barrel)	< 1.5	–
E/p (Endcaps)	< 2.45	–

10.2.2.3 HLT selection for photons

Photons have a large background coming from high E_{T} particles decaying into photons, such as π^0 and η. Isolation requirements are the most effective way to reduce this background. In the benchmark signal process $H \rightarrow \gamma\gamma$ the 2 photons are isolated while in the reducible background processes such as $pp \rightarrow$jets and $pp \rightarrow\gamma$+jet at least 1 of the photon candidates comes from the fragmentation of a jet and is not well isolated.

The other handle that can be used to reduce the trigger rate is to increase the E_{T} cuts on the candidate photons.

The isolation variables that are used at trigger level are:

- number of tracks with p_{T} larger than 1.5 GeV/c inside a cone with $\Delta R < 0.3$ around the photon candidate;
- total E_{T} of all island basic clusters with $\Delta R < 0.3$ around the photon candidate, excluding those belonging to the supercluster itself;
- total transverse energies of HCAL towers within $\Delta R < 0.3$ around the photon candidate.

We apply different cuts on these variables for the single and double triggers with looser cuts being applied on the double trigger. The E_{T} thresholds for the single trigger is 80 GeV while for the double is 30 GeV for the highest E_{T} photon and 20 GeV for the second highest E_{T}. In the case of the double stream triggers the cuts are applied on both photon candidates.

A summary of all HLT cuts is reported in Table 10.4. We would like to point out that the cut on the track isolation for the double trigger allows a maximum of 2 charged tracks inside the isolation cone, and thus does not reject early conversions for which 2 tracks are found in the direction of the photon.

Table 10.5 shows the rates output to archival storage for the 4 electron and photon streams at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Tables 10.6 and 10.7 show trigger efficiencies for benchmark photon signal and electron signal channels after the HLT. Note that the double photon HLT trigger is also very efficient for double electron processes, given the fact that in order to keep conversions no veto is applied on matching tracks. This feature could be exploited, particularly in the early stages of the CMS data taking, to verify the cluster-track matching requirements of the electron HLT.

As can be seen from Table 10.5 the HLT trigger rate for the single photon selection is less than 3.5 Hz of which approximately 2 Hz are due to the γ + jet process that is irreducible. From

Table 10.4: HLT cuts for the single and double photon streams.

Variable	Single photon	Double photon
$\|\eta\|$	< 2.5	< 2.5
E_{T}	> 80 GeV	$> 30, 20$ GeV
Track isolation	$= 0$	< 3
HCAL isolation (Barrel)	< 6 GeV	< 8 GeV
HCAL isolation (Endcaps)	< 4 GeV	< 6 GeV
ECAL Isolation	< 1.5 GeV	< 2.5 GeV

Table 10.5: Output rates from HLT for electron and photon streams at $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

	Signal		Background		Total
Single electron (E_{T} >26 GeV)	$W \to e\nu$ $Z \to e^+e^-$	9.8 Hz 1.3 Hz	Jets	9.4 Hz	21 Hz
Double electron ($E^1_{\mathrm{T}}, E^2_{\mathrm{T}}$ >12 GeV)	$Z \to e^+e^-$	1.1 Hz	Jets	0.8 Hz	1.8 Hz
Single photon (E_{T} >80 GeV)	γ + jet	2.1 Hz	Jets	1.4 Hz	3.5 Hz
Double photon (E^1_{T} >30, E^2_{T} >20 GeV)		≈0 Hz	Jets γ + jet	1.9 Hz 0.4 Hz	2.3 Hz
Total:		13.3 Hz		13.9Hz	27.2 Hz

Table 10.6: HLT efficiency for the single and double electron triggers for $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Signal Process	Single electron	Double electron	Total
$W \to e\nu$	68.0%	–	68.0%
$Z \to e^+e^-$	81.2%	76.7%	89.5%
$H \to ZZ \to e^+e^-e^+e^-$ ($M_H = 120\ \mathrm{GeV}/c^2$)	76.9%	83.2%	90.2%

Table 10.7: HLT efficiency of the single and double photon triggers for $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Signal Process	Single photon	Double photon	Total
$H \to \gamma\gamma$	14.0%	86.9%	88.4%
$Z \to e^+e^-$	–	70.0%	70.0%
$H \to ZZ \to e^+e^-e^+e^-$ ($M_H = 120\ \mathrm{GeV}/c^2$)	–	59.6%	59.6%

this it is understood that the single photon trigger E_{T} threshold of 80 GeV cannot be reduced without substantially increasing the final HLT output rate. On the other hand it would be useful to select γ + jet events even at lower E_{T} for different purposes. The HLT trigger rate for the single photon selection becomes approximately 400 Hz for an E_{T} cut of 23 GeV. Analogously the HLT trigger rate for the double photon selection is approximately 20 Hz when the E_{T} cut at 12 (19) GeV is relaxed for the isolated (non-isolated) trigger. This reduction would allow the study of the $pp \rightarrow \gamma\gamma$ process that is the main irreducible background to the $H \rightarrow \gamma\gamma$ search. Implementation of pre-scaled HLT triggers with these low thresholds is foreseen, with pre-scaling factors defined as function of E_{T} in such a way that the differential rate with respect to E_{T} is constant. These pre-scaled triggers would increase the final HLT output by (typically) less than 1 Hz.

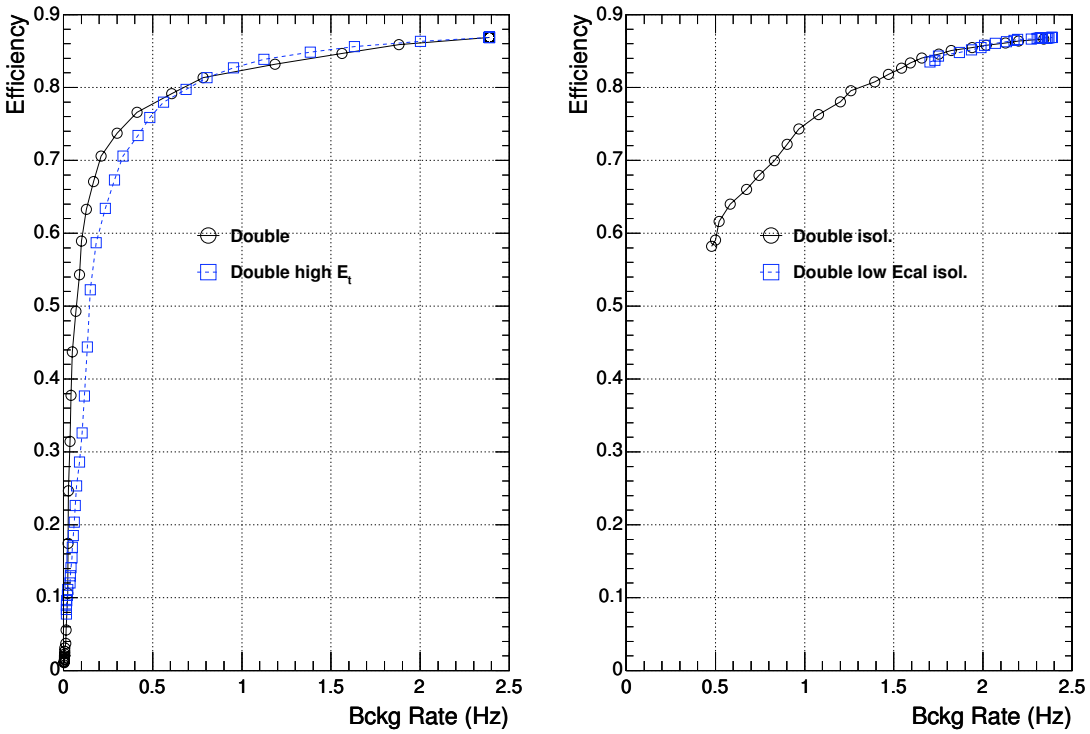


Figure 10.4: Effect of tightening the E_{T} thresholds (left) and ECAL isolation requirement (right) on the $H \rightarrow \gamma\gamma$ signal efficiency versus background rate for the double photon streams.

10.2.2.4 Further HLT tuning

The final Level-1 and HLT thresholds will be defined based on the actual luminosity of LHC. The HLT selection presented here is relatively loose and there is space to further reduce the background rate without compromising the signal efficiency if necessary. As an example Figure 10.4 shows the effect of the hadron calorimeter isolation and of the track isolation. The E_{T} cut varies from 30 (20) to 100 GeV in steps of 2 GeV for the highest (second highest) E_{T} candidate while the ECAL isolation cut varies from 2.5 down to 0 GeV in steps of 0.1 GeV. It is clear from the figure that the rate can be reduced by a factor 3 or more with a reduction of signal efficiency of less than 10% by applying cuts that are still looser than the final analysis cuts.

10.3 Photon reconstruction and selection

This section describes the reconstruction and selection of high p_{T} prompt photons. Because of the very high background to photons from jets, the reconstruction and selection of prompt photons for $H \rightarrow \gamma\gamma$, or for single photon studies, is rather naturally restricted to triggering photons, and rather sharply demarcated from the reconstruction of low p_{T}, non-triggering, neutral energy deposits to be used, for example, in energy-flow jet reconstruction algorithms. The triggering photons selected by the HLT (Section 10.2.2) thus provide the starting point for offline reconstruction and selection of prompt photons. These selected photons already have transverse energy above the HLT selection thresholds and have passed the isolation criteria of the Level-1 and HLT selection.

10.3.1 Photon energy reconstruction

The photon energy is reconstructed by summing the energy deposited in a cluster of crystals. The clusters are 5×5 crystal arrays, or the superclusters found by the hybrid algorithm, in barrel, or the island algorithm, in end-caps. The clustering algorithms have been discussed in Section 10.1. The choice of clustering algorithm depends on the spread of the deposited energy. For compact energy deposits, originating mainly from the single showers of unconverted photons, the best energy measurement is achieved using a 5×5 crystal array, while superclustering algorithms provide better measurement for multiple showers originating from conversions where the bending of the electron and positron tracks, and their radiation in tracker material, spreads the energy over a larger area of the calorimeter.

A very convenient measure of the lateral spread of energy deposition is provided by the $R9$ variable, defined as the ratio of energy contained in a 3×3 array of crystals (centred on the crystal with the highest deposited energy) to the total supercluster energy. Values approaching unity are obtained for unconverted photons, or for photon that have converted very close to the ECAL. Smaller values are obtained for increasing distances of the conversion vertex from the ECAL. Figure 10.3.1 shows the photon energy distribution for 2 bins in the $R9$ variable: the first bin, $R9$ >0.943, contains the best 70% of photons, and the other bin, $R9$ <0.943, the worst 30%. The difference in energy resolution in the 2 bins is striking. These are barrel photons in $H \rightarrow \gamma\gamma$ events, fully simulated with $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ pile-up, and passing the HLT selection cuts.

The precision of photon energy measurement varies due to a number of factors. Shower spread due to early conversion has been mentioned. The shower energy is another obvious factor. Showers close to the module borders have their energy resolution degraded. An estimate of the photon energy measurement error can be provided by classifying the showers and parameterizing the error as a function of relevant quantities. The error estimate can then be used in the analysis of photon signals to supplement other selection variables.

An example of the derivation of such an error estimator is shown in Figure 10.6. The figure shows the energy resolution for barrel photon showers with $R9$ >0.943, reconstructed in a 5×5 crystal array, as a function of energy, together with the fitted parametrization. The energy resolution has been obtained from the energy distribution at a given energy by measuring the "effective sigma," σ_{eff}, half the width of the smallest window which contains 68.3% of the distribution. This measure takes account, in a useful way, of any non-Gaussian tail. The resolution is dominated by contributions for constant term, stochastic term, and noise which have been added to the GEANT simulation.

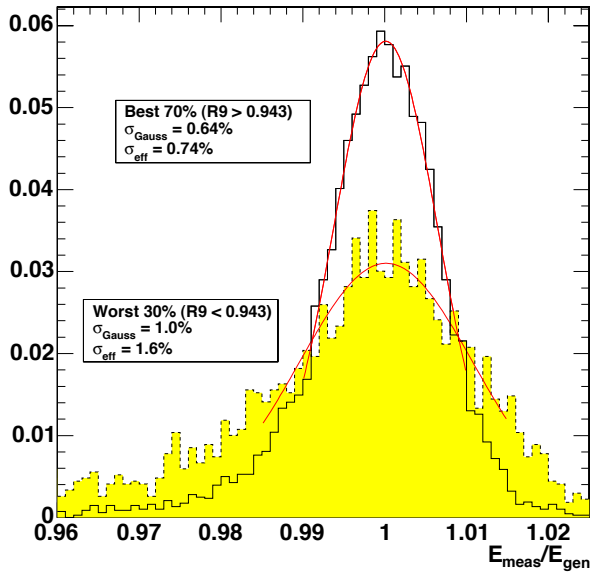


Figure 10.5: Reconstructed photon energy normalized to true energy, $E_{\mathrm{meas}}/E_{\mathrm{true}}$, for 2 different bins in the $R9$ variable.

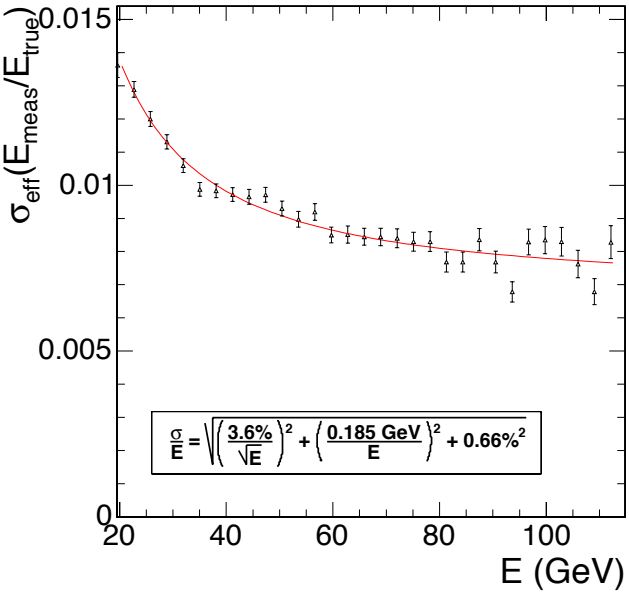


Figure 10.6: Energy resolution for barrel photon showers with $R9$ >0.943, reconstructed in a 5×5 crystal array, as a function of energy, together with the fitted parametrization.

10.3.2 Photon isolation

Isolation requirements are one of the most important tools needed to reduce the background to prompt photons, both at trigger level and at analysis level. In our reference signal process $H \rightarrow \gamma\gamma$ the 2 photons are isolated while in the reducible background processes such as $pp \rightarrow$ jets and $pp \rightarrow \gamma$ + jets at least 1 of the photon candidates comes from the fragmentation of a jet and is not well isolated. Accompanying these high E_{T} electromagnetic showers there

are usually additional particles and hence hadronic energy near the showers.

Fake photon signals due to jets can usually be rejected by looking for additional energetic particles in a cone around the reconstructed ECAL cluster. Charged pions and kaons can be detected in the tracker or in the calorimeter. Neutral pions and other particles decaying to photons can be detected in the ECAL. The hadron calorimeter may be important for detecting charged particles not efficiently reconstructed in the tracker, particularly at high pseudo-rapidity, or particles like neutrons or K_L^0.

The size of the veto cones should be optimized. Since energy may be spread out in the calorimeter due to showering and magnetic field, cones in the hadron calorimeter, for example, should likely be larger than in the tracker.

The basic isolation variables considered are based on charged tracks reconstructed in the tracker, electromagnetic energy deposits observed in the ECAL, and hadronic energy deposits in the HCAL. We also make some attempt to combine the different isolation information in an optimal way. Further details can be found in [243].

10.3.2.1 Monte Carlo simulation samples for the isolation study

We use a sample of simulated gluon-gluon fusion $H \rightarrow \gamma\gamma$ events ($M_H = 120\ \mathrm{GeV}/c^2$) as a source of signal photons.

For background jets, we use a sample of γ + jet events and select the "fake" reconstructed photon shower based on Monte Carlo truth information. That is, we make sure that the highest energy reconstructed ECAL cluster matches with a generator level prompt photon and that the second highest energy ECAL cluster does not. The one that does not match with a photon is then used as an example of jet background.

We exclude clusters with $|\eta| > 2.5$ and in the gap between the barrel and endcap ($1.4442 < |\eta| < 1.566$). To study photons and jets relevant to the Higgs analysis, we also require $E_{\mathrm{T1}} > 40$ GeV, $E_{\mathrm{T2}} > 25$ GeV and $105 < M_H < 155\ \mathrm{GeV}/c^2$. Photon efficiencies and background rejection factors are then computed relative to this pre-selection. All figures refer to the second highest E_{T} photon candidate and, if not specified, further require that $E_{\mathrm{T2}} > 40$ GeV.

10.3.2.2 Performance estimation

In order to estimate the performance of a given isolation variable we plot the signal efficiency versus the ratio of signal efficiency to background efficiency. This is equivalent to a plot of signal efficiency versus signal purity. The x-axis is also closely related to the background rejection factor. The larger the signal efficiency and purity, the better the performance of the analysis. We can compare the different variables, cone sizes, and thresholds by looking at these performance plots.

The lines in these plots will become vertical when an increasingly tighter cut begins to reject the same fraction of signal and background. Most of the performance curves for single variables simply end at some point just before all signal and background are rejected by further tightening the cut.

10.3.2.3 Tracker variables

Three types of isolation variables are considered in the tracker.

- The sum of track p_{T} for tracks with p_{T} above a given threshold in a cone or between 2 cones of size ΔR around the ECAL cluster. For example "Σp_{T} all tracks $\Delta R < 0.3$" is the sum of track p_{T} for in a cone of 0.3.
- The angle, from the ECAL cluster, to the n^{th} nearest track above some p_{T} threshold. For example, "ΔR 1^{st} track, $p_{\mathrm{T}} > 1.5$" is the angle to the first track above 1.5 GeV/c.
- N_{tk}, the number of tracks with above a given threshold, within a cone around the ECAL cluster ("same vertex" indicates that only the tracks that come from the chosen primary vertex are counted). This variable is an integer and therefore does not offer a continuous spectrum.

In the 3 cases the separation between the track and photon direction is measured by $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. Different cone sizes have been considered and tracks have been used when their p_{T} exceeds a given threshold $p_{\mathrm{T}}^{\mathrm{min}}$. In the following this threshold, unless otherwise specified, is set at 1.5 GeV/c that seems to be an optimal value for most variables.

If all reconstructed tracks are used the efficiency is reduced by the presence of pile-up. It is possible to consider only the tracks consistent with the primary vertex assigned to the event. In case of the background to $H \rightarrow \gamma\gamma$ the correct vertex can be identified in approximately 80% of the cases at low luminosity.

Figure 10.7 shows the efficiency versus rejection for the tracker related isolation variables. It is apparent that these variables are much less effective in the endcaps due to the larger tracker inefficiencies.

It should be noted that the angle to the first track, represented by the stars in the figure, is performing very well but shows an instant loss of about 4% in signal efficiency. This is due to tracks at very small angle with respect to the signal photon that are produced in early conversions. This signal efficiency could be recovered by reconstructing the conversions and removing the tracks from the analysis. This has not yet been fully implemented.

10.3.2.4 ECAL variables

The electromagnetic calorimeter can be used to see if the photon candidate is accompanied by π^0 or η particles. In addition it may be used to determine whether the ECAL cluster that is the photon candidate was due to a π^0 that decays to 2 photons or due to a single isolated photon. In some cases, the cluster may even be due to 2 particles that decayed into photons giving a broader shower than would a single photon.

To study isolation in the ECAL we used basic clusters and superclusters reconstructed with the island algorithm even though the superclusters in the barrel were reconstructed using the hybrid algorithm. Island basic clusters were found to be significantly better for this purpose.

Four types of variables were studied:

- the sum of E_{T} from basic ECAL clusters, in a cone of radius ΔR around the candidate supercluster, excluding all the energy in the supercluster;
- The sum of E_{T} from basic clusters between 2 cones around the candidate super-

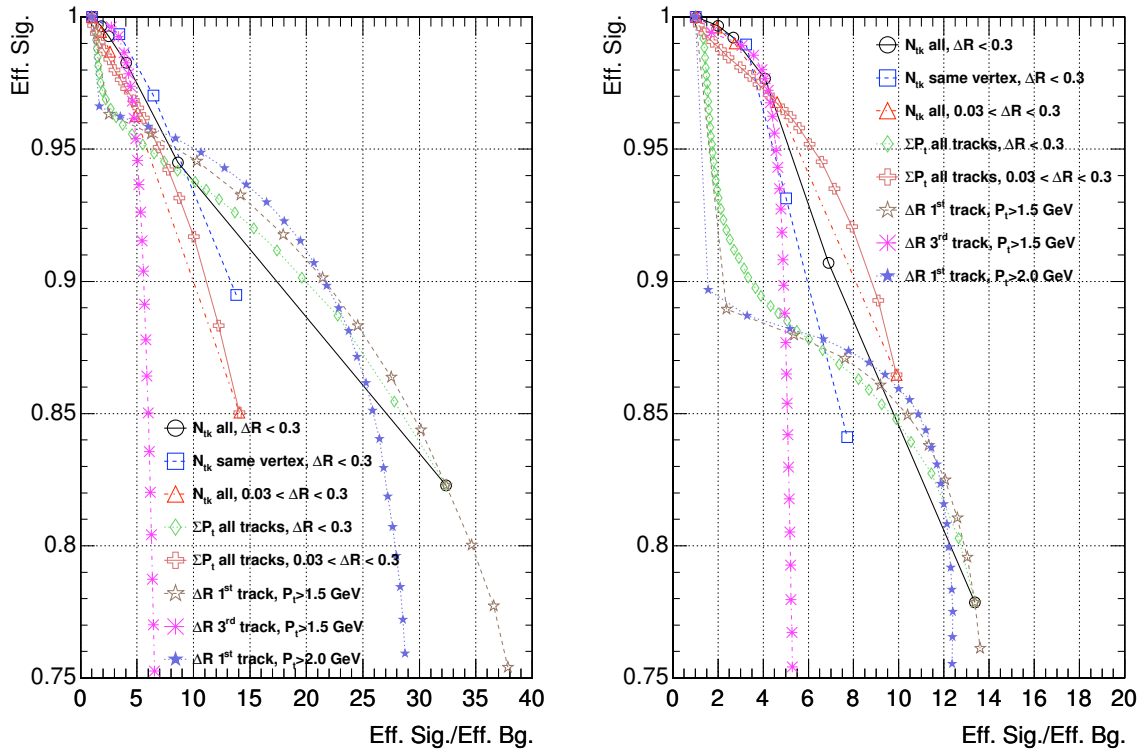


Figure 10.7: Performance plots of some of the best tracker isolation variables. The left hand plot is for the barrel and right hand plot for the endcaps.

cluster ($\Delta R_{\min} < \Delta R < \Delta R_{\max}$);

- The simple shower shape variable R_9, defined as the fraction of the supercluster energy found inside the 3 by 3 array of crystals centred on the highest energy crystal;
- The legacy isolation variable from the electron analysis "ORCA iso" which is related to variables of type 2 but excludes energy in an η band centred on the cluster which might come from conversion and bending in the magnetic field.

Some of the best variables of the various types are evaluated in Figure 10.8.

We can notice how the different methods of defining the isolation transverse energy behave in the barrel and endcaps: in the endcaps the variables using inner and outer cones are much more powerful in all cases while also in the barrel they seem to yield the largest background rejection factors, probably by beginning to "look inside" the supercluster.

10.3.2.5 Hadron calorimeter variables

The hadron calorimeter offers information that is partially redundant with tracker measurements. Charged pions and kaons measured in the tracker will deposit much of their energy in the HCAL. Long lived neutral particles like K_L and neutrons may only be seen in the HCAL and also the energy from charged hadrons that are not reconstructed may only be detected in the HCAL.

Figure 10.9 shows the performance of variables that simply use the sum of HCAL E_T in a

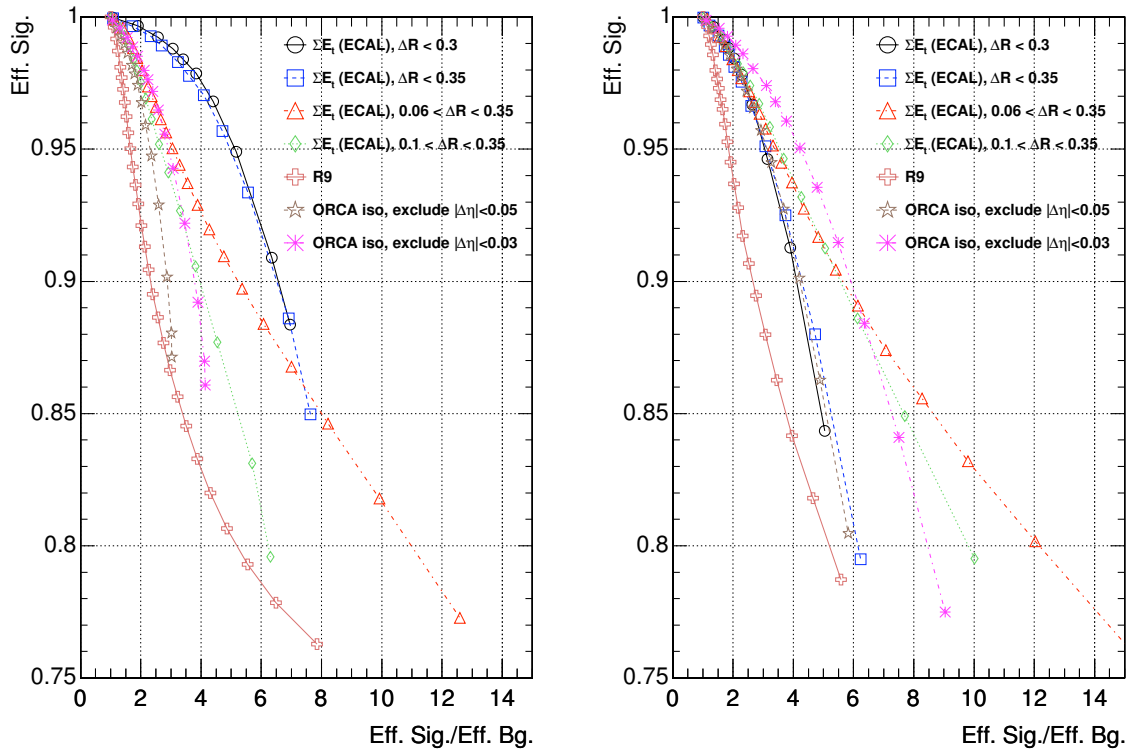


Figure 10.8: Performance plots of some of the best ECAL isolation variables in barrel (left) and in endcap (right).

cone around the candidate. They do not differ very much nor do they provide large background rejection factors. The performance is better in the endcaps and can complement there the reduced power of the tracker related variables.

The "ORCA H/E" variable, defined as the HCAL energy in the tower immediately behind the supercluster divided by the ECAL supercluster energy, seems to be significantly worse than the simple sums.

10.3.2.6 Combination of detectors

There are various ways how the information from different detectors can be combined to significantly improve background rejection. We have added some of the E_{T} variables presented in previous sections to make more performing isolation variables. This is justified by the fact that the particle content of the jets will vary and hence the jet could deposit energy in any layer of the detector and not in the others. Sensible combinations would be (Tracker + ECAL), (ECAL + HCAL) and (Tracker + ECAL + HCAL). In principle, variable multiplicative factors could be employed in the sum.

Figure 10.10 shows all 3 types of sums. These reach a background rejection of approximately 70 at 75% photon efficiency in both the barrel and the endcap. The performance is better than any single detector variable.

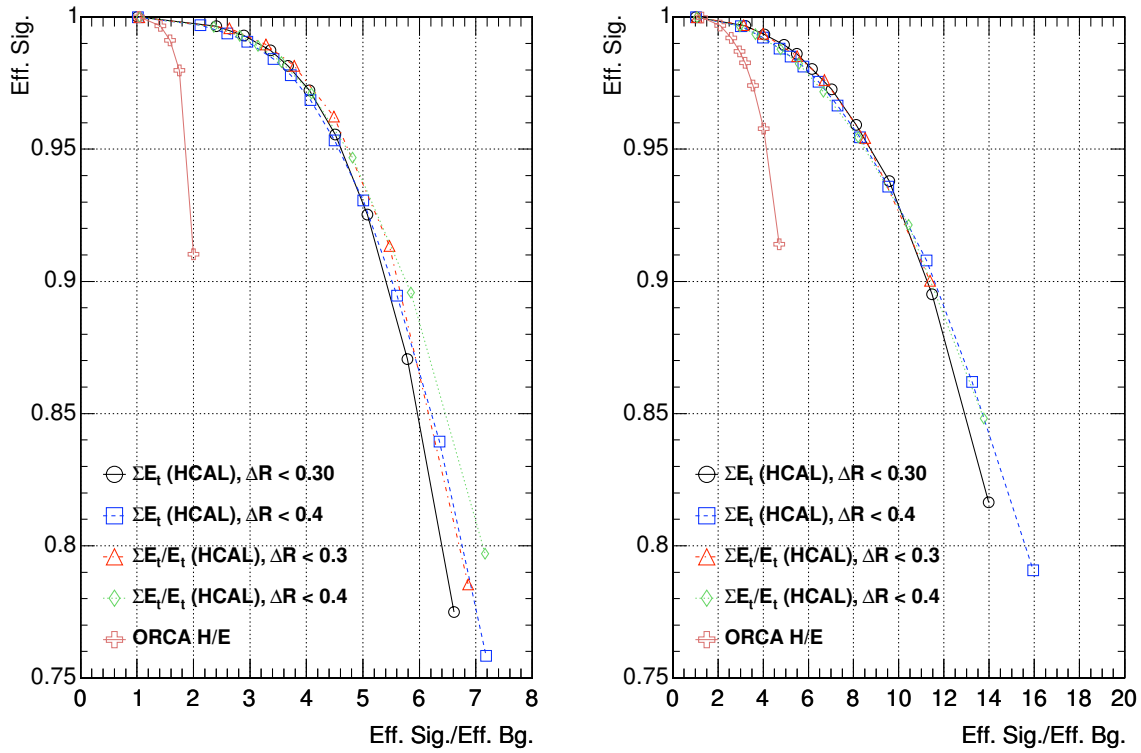


Figure 10.9: Performance plots of some of the best HCAL isolation variables in barrel (left) and in endcap (right).

10.3.2.7 Optimized combination of cuts

All isolation variables shown are highly correlated, nevertheless higher rejection power for the same signal efficiency can be obtained by applying cuts on different variables or with multivariate analysis techniques. The final optimization of the isolation requirements must be specific to a given analysis, that is: specific to the particular signal and background.

By using a set of cuts on different variables we found that similar background rejection, at the level needed for the $H \rightarrow \gamma\gamma$ search, can be achieved at 10–15% higher signal efficiency than by using E_{T} sums.

10.3.2.8 Combining isolation variables with a neural net

A neural net can be an effective way of combining the information from a correlated set of input variables into 1 output variable. The neural net is trained to put signal events near a value of 1 and background events near a value of zero. This output variable then provides a continuous variable that we can cut on to improve the signal to background ratio. There are several other methods that could be employed to do this but the results will probably not differ very much from an optimized neural net.

We have used a neural net program to optimize a single isolation variable from the input variables studied above. Various combinations of inputs have been tried following our understanding of the isolation analysis. While further optimization is possible, the current results are very good and we think close to the overall optimum.

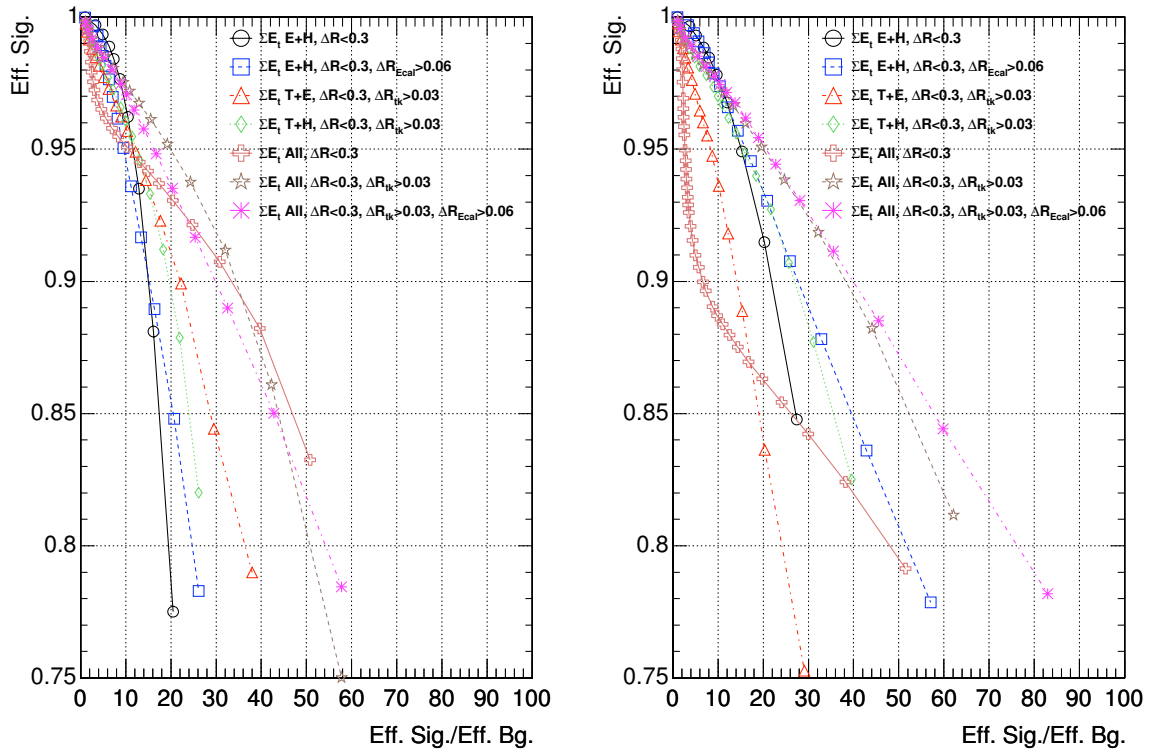


Figure 10.10: Performance plots of some of the best combined isolation variables in barrel (left) and in endcap (right).

We used the following isolation variables as inputs to the neural network:

- the angle to the first track above an E_{T} of 2 GeV;
- the ECAL transverse energy in a cone of radius 0.3 excluding the candidate supercluster energy;
- the shower shape variable R_9;
- E_{T} sum in the HCAL in a cone of of radius 0.35;
- the sum of E_{T} in the tracker in a rather narrow cone of 0.2.

Figure 10.11 shows the performance plots of the neural network output variable in the barrel and in the endcaps. We can see the important contribution of the R_9 variable for large rejection factors as well as the relevance of HCAL isolation in the endcaps.

At high photon efficiency the performances are significantly better than those corresponding to simple isolation variables shown above, reaching rejection ratios of 50 at 90% photon efficiency. For the $H \rightarrow \gamma\gamma$ analysis, we are also interested in very large rejection factors. At 50% photon efficiency, the signal to background efficiency ratio reaches 200. Similar results have also been obtained by applying sequential cuts on the previously defined individual variables.

We should also note that none of these neural net results uses any of the kinematic variables such as the supercluster E_{T}. The Neural Net output can then be easily used in conjunction with the di-photon mass and other kinematic variables to optimize the signal detection performance.

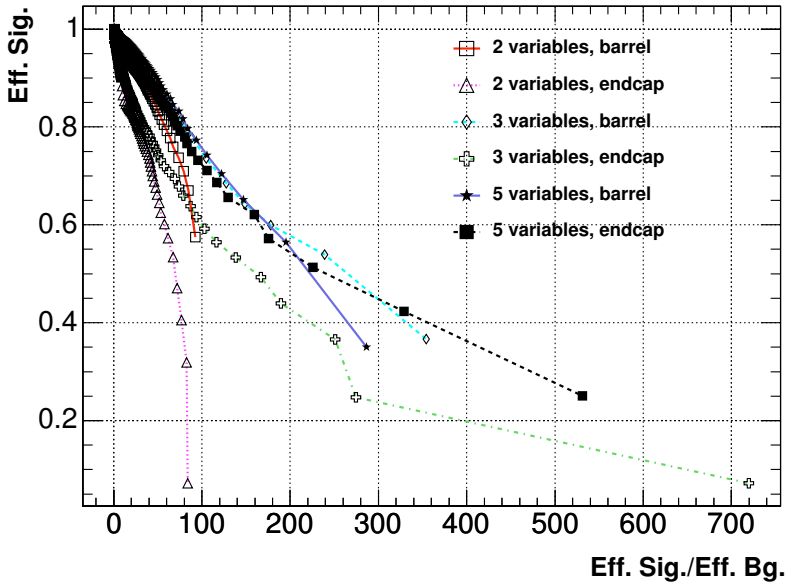


Figure 10.11: Performance plots of the neural network variable in barrel and in endcaps. Results are shown for the basic 2 input variables from tracker and ECAL, for the addition of the R_9 input to make 3 variables, and for the addition of HCAL isolation and another tracker variable to make 5 input variables.

10.3.3 Converted photons

A large number of photons originating from the primary interaction vertex convert in the tracker material. Figure 10.12 shows the integral distribution of the fraction of photons converted as a function of radius, measured in a fully simulated sample of $H \rightarrow \gamma\gamma$ events, for 3 representative values of $|\eta|$: near the centre of the barrel ($|\eta| = 0.2$), in the middle of 1 half barrel ($|\eta| = 0.7$), and towards the end of the barrel of the ECAL ($|\eta| = 1.3$).

Identification of converted photons allows a better choice of energy clustering algorithm. Reconstructed tracks of converted photon candidates can be used to provide π^0/γ separation. The reconstructed tracks can also be used to locate the primary vertex coordinate along the beam line, for use in $H \rightarrow \gamma\gamma$ mass reconstruction. Reconstruction of secondary, conversion, vertices from the tracks also allows mapping of the distribution of tracker material ("radiography").

10.3.3.1 Converted photon tracking

Track reconstruction can be divided into 4 separate steps: a) trajectory seed generation; b) trajectory building; c) trajectory cleaning for resolving ambiguities; and d) trajectory smoothing (i.e., the final track fit). Full details of CMS track reconstruction are given in Section 6.4.

Tracking of the electrons from a converted photon requires a dedicated seed finding technique [244]. Energy deposits in the ECAL are used as a starting point for inward track seed finding. All basic clusters included in the supercluster and all other basic clusters less than $\Delta\eta = 0.015$ and $\Delta\phi = 0.25$ from the supercluster position are used to build track seeds. For each cluster both positive and negative charge hypotheses are tried. The ϕ range is chosen taking into account the maximum separation at the ECAL front face between tracks from

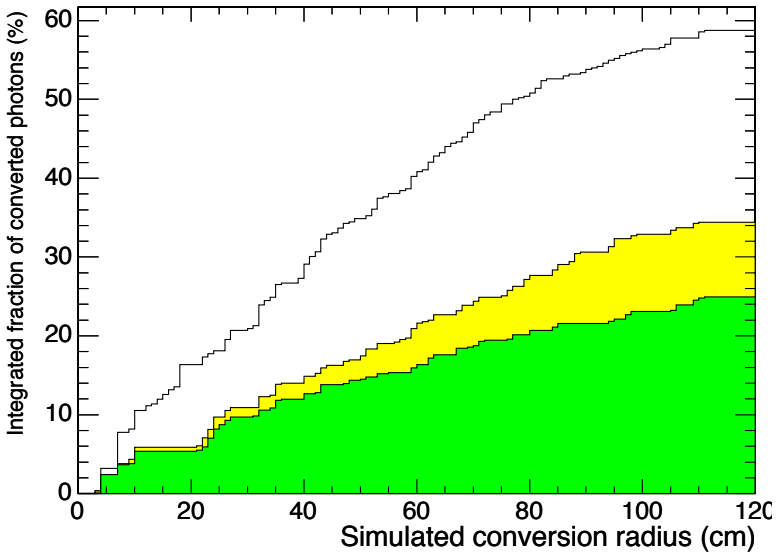


Figure 10.12: Fraction of photons from $H \rightarrow \gamma\gamma$ events converted as a function of radius. The 3 histograms correspond to slices of width $\Delta\eta = 0.1$ around $|\eta| = 0.2$ (dark grey), $|\eta| = 0.7$ (light grey) and $|\eta| = 1.3$ (empty histogram).

converted photons found in the simulated data. The η range is small because track pairs from converted do not open in the longitudinal plane.

The initial momentum estimate is obtained from the basic cluster energy. Reconstructed hits compatible with the trajectory hypothesis are sought in the 2 outermost layers of the tracker. In the (r, ϕ) plane a maximum separation of $\Delta\phi = 0.015$ between the trajectory and the hits is allowed while in the (r, z) plane the hits are required to lie along a straight line from the primary vertex to the basic cluster, with the spread in the position of the primary interaction vertex ($\pm 15cm$) taken into account.

If a compatible hit is found the predicted track state is then propagated inward to find a second compatible hit. If none is found the next innermost layer is searched. All pairs of hits found are collected and input to the trajectory builder, where compatible reconstructed hits are added layer by layer moving inward. Redundant track trajectories are identified from the number of shared reconstructed hits: if 2 candidates have a fraction of shared hits larger than 50%, the trajectory with the largest number of hits is retained. The trajectory with the smallest χ^2 is retained if 2 tracks happen to have the same number of reconstructed hits.

This inward tracking procedure results in a list of tracks of which only the opposite charge pair with the largest number of reconstructed hits is retained. Each of the 2 tracks is used in turn, independently from the other, as a starting point for seed and track finding of the parent track in the pair with an outward tracking procedure. A given inward track is propagated further inside the Tracker, to cover the possibility that the innermost hit has not been reached in the first step. At each inward propagation step the assumption is made that the innermost hit is at the conversion vertex and a search for the parent track of the pair is performed. Each ECAL basic cluster compatible in ϕ with the crossing point is used to predict the momentum and the curvature of the track path. Pairs of reconstructed hits compatible with the predicted track trajectory are then sought moving outward in the 2 following layers. The pairs of hits

so found are used as seeds and fed to the trajectory builder for outward trajectory building. Trajectories are then cleaned and finally smoothed with the backward Kalman fit to obtain the parameters of the tracks at the vertex.

The average radiation energy loss electrons experience while traversing the tracker material is described by the Bethe-Heitler formula. Within the Kalman filter the radiation energy loss is taken into account by correcting the track momentum by an amount corresponding to the predicted mean value of the energy loss and by increasing, at each propagation step, the track momentum variance by adding the predicted variance of the energy loss under the assumption that its distribution is Gaussian.

At the end of the tracking procedure there can be multiple track pairs for a single ECAL supercluster. A further criterion must be applied to choose the best candidate at the analysis stage.

10.3.3.2 Reconstruction efficiency

The algorithmic efficiency can be defined as the probability of reconstructing tracks that have enough hits to be reconstructed. Photons converting late in the tracker are largely indistinguishable from unconverted photons. The shower shape in the ECAL is compact and they do not leave a sufficient number of hits in the tracker for their tracks to be reconstructed.

Converted photons with 2 simulated tracks, each with at least 3 simulated hits and matching in (η, ϕ) a supercluster with transverse energy above 25 GeV are used. The requirement on the minimum number of simulated hits is equivalent to requiring $r_{\text{conv}} < 85$ cm (barrel) and $z_{\text{conv}} < 210$ cm (endcap). Simulated tracks were considered down to $p_{\text{T}} = 0.3$ GeV/c.

Each of these photons can have more than one track or track pair reconstructed. If at least one reconstructed track or track pair (correctly matching the e^+e^- simulated tracks) is found the converted photon is considered as reconstructed. A distinction is made, however, between those candidates which have only 1 track and those which have 2 tracks reconstructed. If more than one matching track or track pair is found, with at least one match being a track pair, the pair is counted for the efficiency calculation.

It is to be expected that the two-track reconstruction efficiency in the ECAL barrel decreases with the radius, r, of the conversion point since there are fewer hits to be found. Asymmetric conversions, on the other hand, where one track takes most of the photon transverse momentum, are less efficiently reconstructed when the conversion occurs early in the detector since very low p_{T} tracks bend more in the magnetic field, and may not even reach the ECAL. The combination of these 2 effects leads to the reconstruction efficiency seen in Figure 10.13 a), which shows the reconstruction efficiency measured on a sample of photons with fixed $p_{\text{T}} = 35$ GeV/c, in the absence of other particles.

In the region corresponding to the ECAL barrel, the inefficiency is mainly due to poor seed finding and track reconstruction in the delicate tracker barrel-endcap transition region, around $\eta = 1$.

Similar considerations apply to the ECAL endcap region, where the z coordinate is more appropriate to measure the position of the conversion point. The efficiency is lower at small z values because 1 of the 2 tracks is easily lost and it drops at large z values because less hits are available for tracking (Fig. 10.13 b). The inefficiency in the endcaps is mainly due to the reduced tracker coverage at very large $|\eta|$ and to the overall delicate range $1 < |\eta| < 1.5$

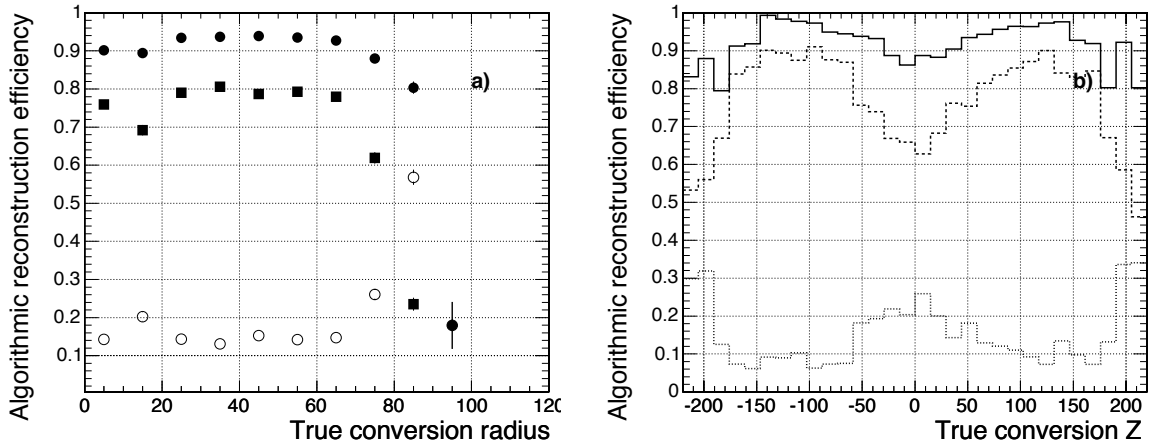


Figure 10.13: Reconstruction efficiency measured for photons with fixed $p_{\mathrm{T}} = 35$ GeV/c as a function of the simulated conversion-point position. a) Versus r for $|\eta| \leq 1.4442$. Total (solid circles), two-tracks (squares) and single track (empty circles); b) versus z for $|\eta| \geq 1.566$. Total (continuous line), two-tracks (dashed line) and single track (dotted line).

where the ECAL barrel matches a mixture of Tracker barrel and endcap. The one- and two-track reconstruction efficiencies as a function of photon η are shown in Figure 10.14).

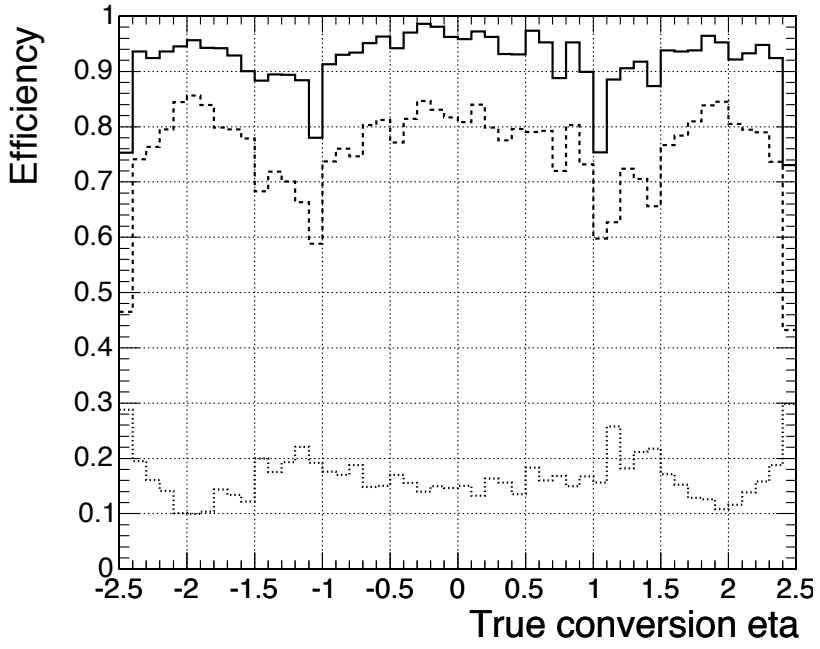


Figure 10.14: Reconstruction efficiency measured from photons with fixed $p_{\mathrm{T}} = 35$ GeV/c in the whole η range. Total (continuous line), two-tracks (dashed line) and single track (dotted line).

The p_{T} value chosen is representative of photons from Higgs boson decay. The average efficiency for full two-track reconstruction over the whole detector is 75±0.3%. If single tracks

are also taken into account the average efficiency reaches 91±0.2%. The track transverse momentum resolution is 6% (9%) in barrel (endcap) respectively (Fig. 10.15).

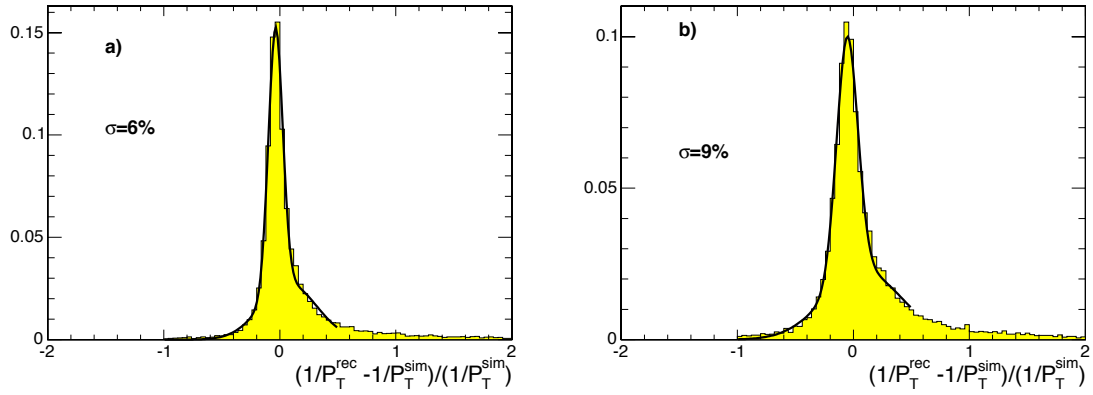


Figure 10.15: Track transverse momentum resolution (measured on a track-by-track basis) for candidates reconstructed in the range a) $|\eta| < 1.4442$ and b) $1.556 < |\eta| < 2.5$.

10.3.3.3 Converted photon vertexing

Each pair of opposite charge tracks is used for vertex finding. An unconstrained vertex fit is performed, which is based on the Kalman formalism (see Section 6.5.1). The pair of tracks used in the vertex fit are refitted with the vertex constraint and all their parameters re-evaluated. The achieved conversion vertex position resolution is shown in Figure 10.16.

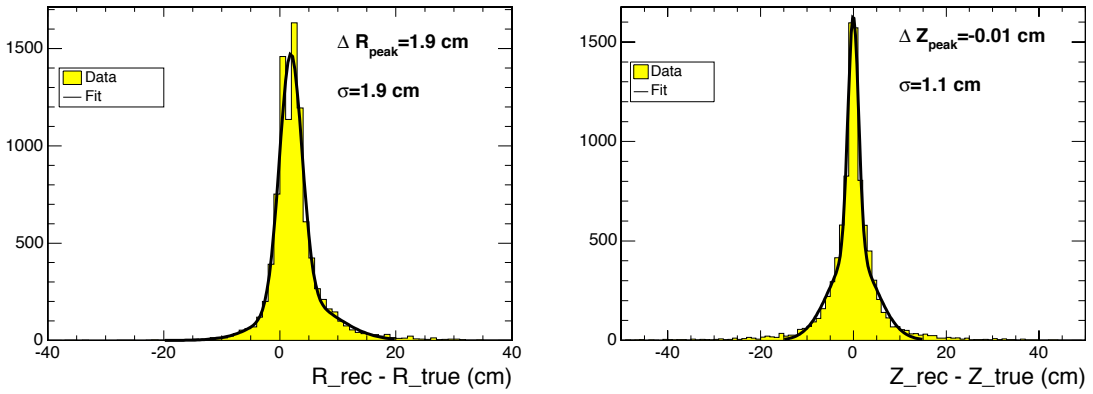


Figure 10.16: Conversion vertex position resolution: radius and longitudinal coordinate.

10.3.3.4 Fake rate in $H \to \gamma\gamma$ events

The fake rate is defined here as the number of two-track candidates reconstructed in unconverted photons divided by the total number of two-track candidates. When multiple two-track candidates are found per each supercluster the one with the smallest invariant mass, calculated after vertex fitting, is chosen. Fake pairs can be caused by hits produced in the pile-up or underlying event. The great majority ($\approx 80\%$) of fakes are found beyond

$|\eta| > 1.5$. In single photons the fake rate is about 0.1% while in $H \rightarrow \gamma\gamma$ events (simulated with pile-up) it is about $(5 \pm 0.3)\%$ averaged over the whole detector.

The transverse momentum measured from tracks in fakes candidates, however, is much smaller than the transverse energy collected in the ECAL supercluster. This is shown in Figure 10.17. A cut on $p_{\mathrm{T}}(\mathrm{tracks})/E_{\mathrm{T}}(\mathrm{ECAL})$ >0.3 (for example) brings the fake rate in single photons down to 10^{-5} and in $H \rightarrow \gamma\gamma$ events with low luminosity pile-up to the level of $0.9 \pm 0.1\%$.

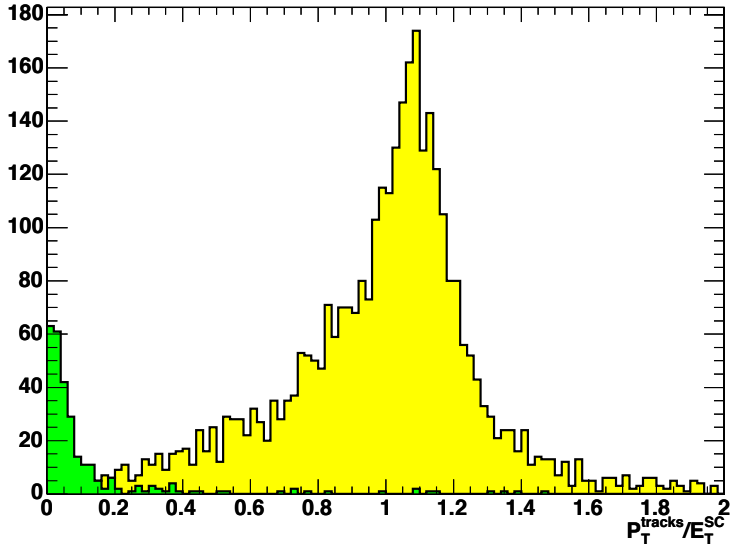


Figure 10.17: $p_{\mathrm{T}}(\mathrm{tracks})/E_{\mathrm{T}}(\mathrm{ECAL}$ for fake (dark grey) and for true (light grey) converted photons in $H \rightarrow \gamma\gamma$ events with low luminosity pile-up ($\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$).

10.3.3.5 Primary vertex z-coordinate from converted photons

The z_{PV} coordinate can be expressed in terms of the reconstructed photon momentum from tracks, the radial ($r^{\mathrm{rec}}_{\mathrm{conv}}$) and longitudinal ($z^{\mathrm{rec}}_{\mathrm{conv}}$) reconstructed vertex coordinates as $z_{\mathrm{PV}} = z^{\mathrm{rec}}_{\mathrm{conv}} - r^{\mathrm{rec}}_{\mathrm{conv}} \times \cot(\theta_{\vec{P}_{\mathrm{tracks}}})$. For locating the primary vertex only converted photons with an accurately reconstructed pseudorapidity are of interest. Photons converting in the outer parts of the tracker are not useful for this purpose. In Figure 10.18 (Fig. 10.19) the achievable accuracy on z_{PV} is shown for $H \rightarrow \gamma\gamma$ events with low luminosity pile-up for conversions reconstructed in the range $|\eta| < 1.4442$ ($1.556 < |\eta| < 2.5$). From left to right the results from converted photons in the pixels, inner tracker and outer tracker are shown.

10.3.4 π^0 rejection

A significant background to the $H \rightarrow \gamma\gamma$ process originates from π^0 in jets, which fake single isolated photons. In the barrel region the fine granularity of the ECAL crystals can be used for $\gamma - \pi^0$ separation. In the endcap the preshower detector, a position-sensitive device with high granularity, aids $\gamma - \pi^0$ discrimination. Algorithms to achieve such separation have been studied for many years. If unconverted photon candidates are selected, rejection factors of more than 2 against π^0 can be achieved for an efficiency of 90%. However the use of this

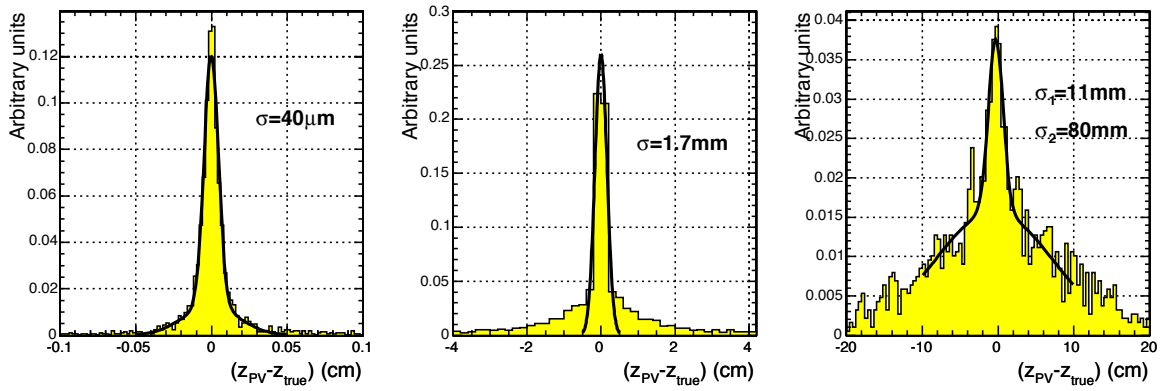


Figure 10.18: Longitudinal coordinate of the primary interaction vertex: accuracy achievable in the Ecal barrel pseudorapidity range. From left to right for $r_{\rm conv} < 15$ cm, $15 < r_{\rm conv} < 58$ cm and $r_{\rm conv} > 58$ cm.

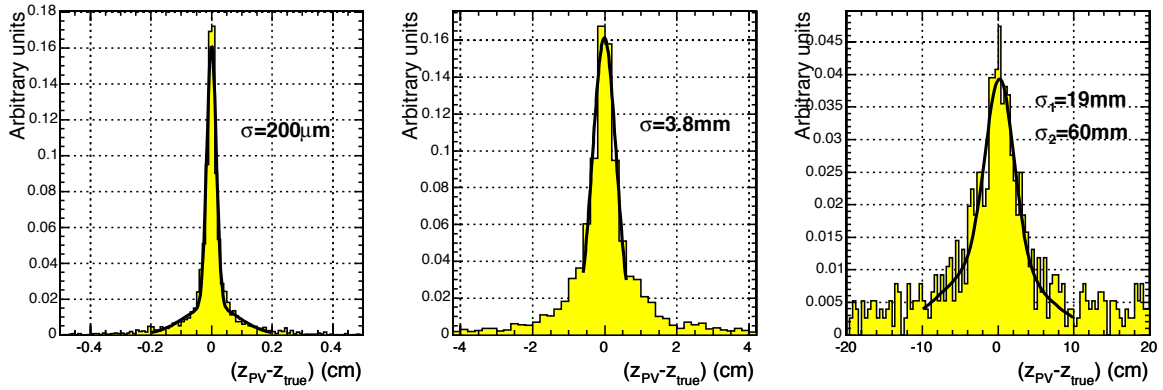


Figure 10.19: Longitudinal coordinate of the primary interaction vertex: accuracy achievable in the Ecal endcap pseudorapidity range. From left to right for $|z_{\rm conv}| < 70$ cm, $70 < |z_{\rm conv}| < 160$ cm and $|z_{\rm conv}| > 160$ cm.

potential requires the integration and optimization of the algorithms in a wider framework where converted and unconverted photon candidates are treated separately.

10.3.5 Offpointing photons

Exploiting the transverse granularity of ECAL it is possible to discriminate neutral isolated particles which do not point to the interaction vertex, typically photons, and reconstruct their impact direction. The strategy to do so relies on the asymmetry of the lateral shower shape. The algorithm consists of 2 steps:

- firstly the two-dimensional projection of the lateral shower shape is characterized by the direction of its major axis (α is the angle between the axis and the η direction) and its degree of asymmetry, Δ.
- secondly, if the asymmetry is above a certain threshold (see Table 10.8), the angle β between the direction of the impinging particle and the plane tangent to ECAL

in the impact point is determined, employing the direct proportionality between β and Δ.

The performances of the algorithm are given in Table 10.8.

Table 10.8: For a given rejection of pointing photons, the efficiency for non-pointing is given for different tilts (β). Intervals are given where, for a given β, there is a slight dependence on different combinations of $\Delta\phi$ and $\Delta\theta$.

Pointing γ rejection	Cut: $\Delta >$	eff. $\beta = 0.2$	eff. $\beta = 0.4$	eff. $\beta >$0.4
90%	0.31	0.32	$\approx$1	$\approx$1
99%	0.40	0.08	0.8-0.9	$\approx$1
99.9%	0.43	0.02-0.04	0.75-0.85	$\approx$1

The precision attainable have been estimated to be $\delta_\alpha = 0.15$ rad and $\delta_\beta = 0.2$ rad.

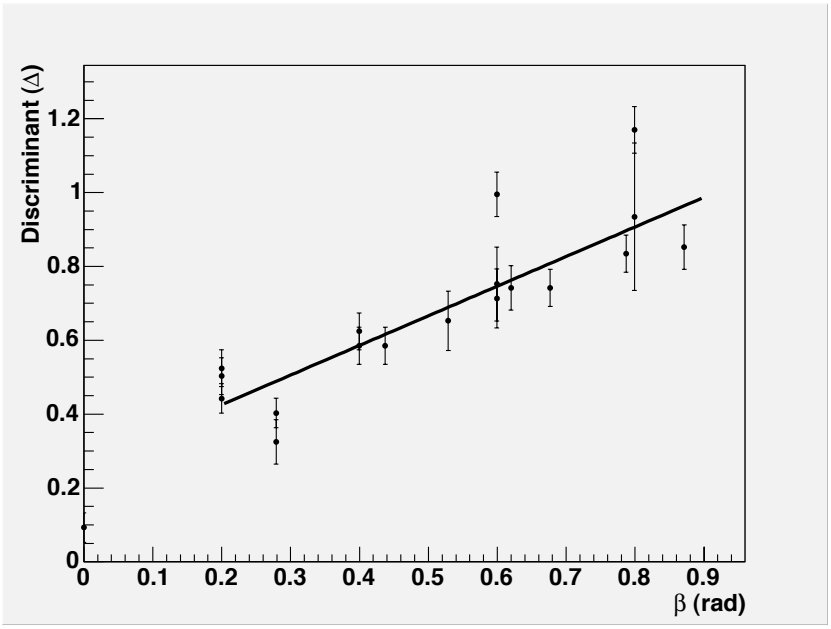


Figure 10.20: The correlation between β and the shower asymmetry parameter Δ.

10.4 Electron reconstruction and selection

This section focuses on the identification and four-momentum measurement of electrons originating from a primary interaction vertex. A primary electron is composed of a single track emerging from the interaction vertex and matched to an electromagnetic supercluster. Such electrons are measured and identified in the angular range $|\eta| < 2.5$ viewed from the nominal vertex and for transverse momenta $p_{\mathrm{T}}^{e} \geq 5$ GeV/c. The emphasis in this section is put on electrons at low and moderate p_{T}^{e}, in a kinematic range relevant for Standard Model Higgs decays.

The electron reconstruction in CMS is hampered by the amount of tracker material which is discretely distributed in front of the ECAL. Electrons traversing the silicon layers of the pixel

and inner tracker detectors radiate collections of bremsstrahlung photons and, in presence of the 4 T solenoidal $\vec{B}$ field, the energy reaches the ECAL spread in ϕ. The spread is p_{T}^{e} dependent. The amount of bremsstrahlung emitted when integrating along the electron trajectory can be very large. It depends on average on the amount of tracker material traversed which varies strongly with η, rising from $\approx 0.3X_0$ at central pseudorapidities to $\approx 1.5X_0$ toward the edge of the barrel ($|\eta| \simeq 1.5$), and falling back to $0.7X_0$ in the endcaps at $|\eta| = 2.5$ As an illustration for electrons of $p_{\mathrm{T}}^{e} = 10$ GeV/c, considering a uniform $\eta - \phi$ distribution, about 50% of the electrons radiate away more than 50% of their energy before reaching the surface of the ECAL. In about 10% of the cases, more than 95% of the initial electron energy is radiated !

The electron measurements can be further complicated by the conversion of secondary photons in the tracker material, which might lead to "showering" patterns and entail energy lost in the tracker material. Soft secondary electrons from e^+e^- pairs get partly trapped in the external $\vec{B}$ field and loose most, or all, of their energy before reaching the ECAL. The photon conversion probability depends on the material budget and thus varies with η. The energy lost can be large. For instance, it amounts on average to 7% of the initial energy for electrons of $p_{\mathrm{T}}^{e} = 10$ GeV/c at $|\eta| \simeq 1.5$.

The bremsstrahlung emission pattern along the electron trajectory can fluctuate considerably on event-by-event basis. The amplitude of the fluctuations increases on average with increasing tracker material budget. This in general introduces largely non-Gaussian event-by-event fluctuations which affect the energy measurement in the ECAL and the momentum measurement in the tracker, as well as the electron identification observables. Sophisticated clustering algorithms (Section 10.4.1) and track reconstruction method (Section 10.4.2) are used to better cope with these effects. Distinct track-supercluster patterns (electron "classes") are recognized (Section 10.4.3) for the energy corrections (Section 10.4.4) and error weighting in the track-supercluster combination (Section 10.4.5). The different "classes" influence the electron isolation (Section 10.4.6) and identification (Section 10.4.7) strategies.

10.4.1 Electron clustering

The building of electron objects is initiated by the presence of electromagnetic superclusters.

The basic "hybrid" and "island" superclustering algorithms used for electrons in the ECAL barrel and endcaps respectively, have been described in detail already in [8] and [160]. An essential purpose of the superclustering algorithms is to deal with the energy spread in ϕ. The collection of bremsstrahlung photons is necessary to minimize the cluster containment variations for the measurement of the primary electron energy.

The ECAL superclustering algorithms have been recently tuned to better cope with the characteristic energy deposition patterns of low p_{T} electrons. The minimal E_{T} threshold for the basic seed cluster of superclusters has been lowered from the previous default of $E_{\mathrm{T}}^{\mathrm{seed}} = 4$ GeV down to $E_{\mathrm{T}}^{\mathrm{seed}} = 1$ GeV. This leads to considerable improvement of the efficiency for reconstructing a supercluster for low p_{T}^{e}. Considering e^+e^- pairs and integrating over the acceptance in η, the efficiency with $E_{\mathrm{T}}^{\mathrm{seed}} = 1$ GeV reaches $\geq 99\%$ for $p_{\mathrm{T}}^{e} \geq 7$ GeV/c, compared to an original efficiency with $E_{\mathrm{T}}^{\mathrm{seed}} = 4$ GeV varying from about 65% for $p_{\mathrm{T}}^{e} = 7$ GeV/c to about 93% for $p_{\mathrm{T}}^{e} = 10$ GeV/c.

Having lowered the supercluster seed threshold, there is a tendency for extra basic clusters, caused e.g. by radiated photons with $p_{\mathrm{T}}^{\gamma} > 1$ GeV/c, to remain separate and form their own

supercluster. For example, with the original ϕ roads for bremsstrahlung recovery extending to ± 0.2 rad in the endcaps and ± 10 crystals (i.e. about ± 0.17 rad) in the barrel, about 8% of back-to-back e^+e^- pairs at $p_{\mathrm{T}}^e \geq 7$ GeV/c gave three superclusters. To better collect the bremsstrahlung and reduce, well below 1%, the probability to find a number of superclusters in excess of the number of isolated electrons, the ϕ roads have been increased to ± 0.3 rad in the endcaps and ± 17 crystals (i.e. about ± 0.3 rad) in the barrel.

ECAL superclusters are use to drive the finding of pixel seeds for the primary electron tracks. The new tuning of the clustering algorithm parameters for bremsstrahlung collection has, in turn, contributed to an improvement of the cluster-driven pixel matching efficiency.

10.4.2 Electron tracking

10.4.2.1 Seed finding

The first step of track reconstruction consists of finding seeds in the pixel detector. For electrons, who suffer from radiative losses, this requires dedicated strategies to preserve efficiency while keeping control of fake track rates. The situation is most challenging for low p_{T} electron tracks.

The electromagnetic superclusters drive the finding of the first 2 track hits ("seeds") in the pixel detector. Such a supercluster-driven pixel seed finding strategy for the tagging of primary electron-like objects has been successfully developed in CMS for robust applications at the High Level Trigger (HLT) [8], where fast and drastic reduction of fake background rates is a key issue. The supercluster and pixel matching takes advantage of the fact that the energy weighted average impact point of the electron and associated bremsstrahlung photons, as calculated using information from the supercluster in the ECAL, coincides (assuming a successful collection of photons) with the impact point that would have been measured for a non-radiating electron of the same initial momentum.

Hits in the pixel layers are predicted by propagation of the energy weighted mean position of the supercluster backward through the magnetic field under both charge hypotheses towards the pixel detector. A first compatible hit is then looked for in the innermost (barrel) pixel layer within a loose $\Delta\phi$ window and loose Δz interval When a first compatible hit is found a new estimate for the z coordinate of the primary track vertex is calculated combining the pixel hit found and calorimetry information in the Rz plane. The predicted trajectory is then propagated to look for a second pixel hit in the next pixel layer(s), within some narrower $\Delta\phi$ and Δz windows. More details on the supercluster-driven pixel matching algorithm can be found in [8]. The requirements for the search of the first and second pixel hits have been loosened with respect to the HLT to recover electron detection efficiency at low p_{T}. The new threshold values are given in Table 10.9.

With this loosening of the matching criteria, the efficiency of the pixel finding for single electrons, averaged over the full ECAL barrel and endcaps η range, reaches about 90% for electrons at $p_{\mathrm{T}}^e = 10$ GeV/c.

The 2 pixel hits found serve as seeds for the building and fitting of electron tracks in the Silicon Tracker Detectors.

Table 10.9: Seeding of electron tracks. Allowed $\Delta\phi$ and Δz search windows for the 1$^{\text{st}}$ and 2$^{\text{nd}}$ pixel hits of the supercluster-driven seed finding algorithm, for HLT and offline reconstruction.

Observable	Requirements	
	HLT level	Offline Reconstruction
$\Delta\phi$ - 1$^{\text{st}}$ pixel hit	40 mrad	200 mrad
	e^- : $[-25\,\text{mrad}, +15\,\text{mrad}]$	
	e^+ : $[-15\,\text{mrad}, +25\,\text{mrad}]$	
Δz - 1$^{\text{st}}$ pixel hit	15 cm	15 cm
$\Delta\phi$ - 2$^{\text{nd}}$ pixel hit	2 mrad	10 mrad
Δz - 2$^{\text{nd}}$ pixel hit	± 0.05 cm	± 0.07 cm

10.4.2.2 Track building and fitting

Dedicated track reconstruction strategies are required for electrons as they are affected by non-Gaussian fluctuations due to bremsstrahlung emission along their trajectory.

The default track reconstruction method in CMS relies on a simple Kalman Filter (KF) algorithm. This has been shown [8] to be usable for high p_{T} electrons (e.g. for applications at HLT), provided that a tight χ^2 cut is used in the trajectory building. In this way, the emphasis is put on the early stages of the electron track evolution which contains the most significant information on the electron initial momentum and direction from the primary vertex. A different approach is needed for low p_{T} electrons and in general for detailed electron analysis.

The simple KF is equivalent to a global least-square minimization based on a linear model for the track evolution and Gaussian fluctuations. The assumption of considering all random fluctuations of track parameters due to material effects as Gaussian, while roughly valid when the dispersion is mainly caused by multiple scattering processes, clearly fails for the energy radiated in the thin layers of the tracker material. A nonlinear filter approach such as the Gaussian Sum Filter (GSF) which had been studied for track reconstruction in CMS [190] can therefore in principle better describe the propagation of electrons. The GSF is a nonlinear generalization of the KF, in which the distributions of all state vectors and errors are Gaussian mixtures. The weights of the components of the mixture depend on the measurements.

A new optimization of electron GSF track reconstruction has been performed recently and is described in detail in [245]. This makes use of a specific Bethe Heitler modeling of the electron energy losses during track building when collecting matching hits in successive silicon layers, with a tolerance of 1 layer without hits. As can be inferred from Fig. 10.21a, the electron track hits can thus be efficiently collected up to the end of the tracking volume.

The performances for transverse momentum reconstruction is illustrated for low p_{T}^e in Fig. 10.21b. A best measurement of the electron p_{T} at the primary vertex is obtained from the most probable values of the mixture of Gaussians characterizing the track state at each tracker layer. More detailed discussions on the electron GSF tracks and measurement performances at the primary vertex can be found in [245]. When compared to the simple KF approach mentioned above, the GSF tracks are found to provide a very similar momentum resolution and slightly improved η and ϕ measurements from $p_{\text{T}}^e = 5$ GeV/c up to at least 30 GeV/c.

The efficiency of electron track reconstruction is shown as function of p_{T}^e and η in Fig. 10.22. The efficiency is found to be above 90%, except at very low p_{T}^e (below about 10 GeV/c), in

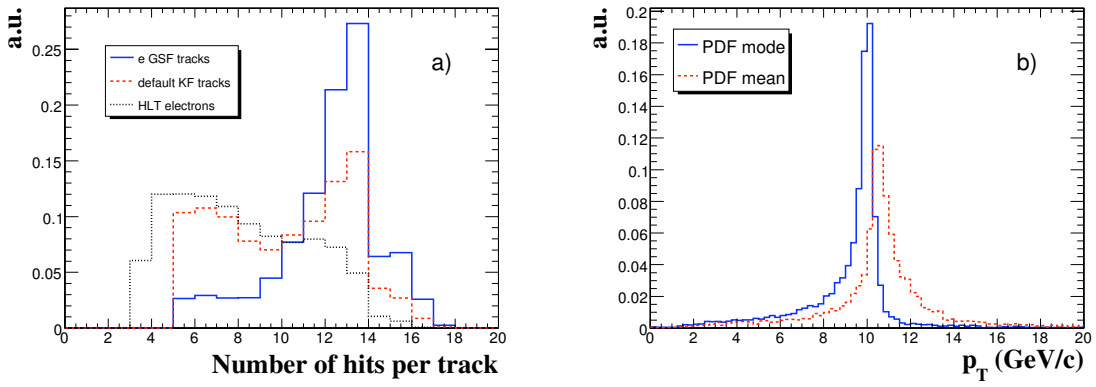


Figure 10.21: Track reconstruction performances for electrons of $p_{\mathrm{T}}^{e} = 10\ \mathrm{GeV}/c$: a) number of collected hits for tracks reconstructed using the Gaussian Sum Filter (solid histogram), the default Kalman Filter (dashed histogram) and the HLT electron Kalman Filter (dash-dotted histogram); b) reconstructed p_{T} for Gaussian Sum Filter tracks as obtained from the most probable value (full histogram) and the weighted mean (dashed histogram) of the Gaussian mixture.

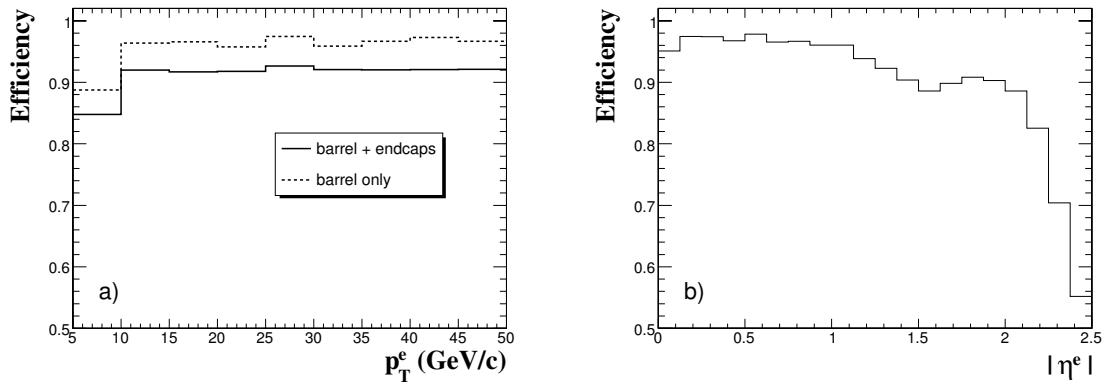


Figure 10.22: Electron track reconstruction efficiency : a) as a function of p_T^e, averaged over the full η range of the ECAL barrel and endcaps (full line), or the range of the ECAL barrel only (dotted line); b) as a function of $|\eta^e|$ for a uniform distribution in p_T^e between 5 and $50\ \mathrm{GeV}/c$. The performances combine effects from supercluster-driven pixel seed matching and a Gaussian Sum Filter trajectory building and fitting.

the transition region between the ECAL barrel and endcap (around $|\eta| = 1.5$), and towards the edge of the acceptance. at largest η.

One of the great benefit of the GSF tracks comes from the combined facts that hits are collected efficiently along the full trajectory through the tracker volume, and that meaningful track parameter errors are available at both track ends. Thus, a good estimation of the electron track parameters at ECAL entrance is made available. But most importantly, the fractional amount of momentum carried away by bremsstrahlung photons can be evaluated from the outermost and innermost track parameters, each obtained e.g. using the weighted mean of Gaussian mixtures [245]. This is illustrated in Fig. 10.23.

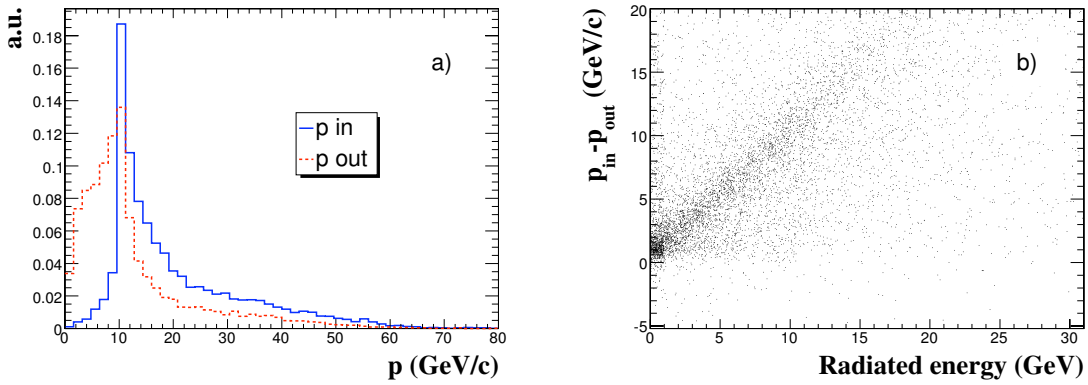


Figure 10.23: Sensitivity of the Gaussian Sum Filter electron track reconstruction to radiated bremsstrahlung energy using the weighted mean of the Gaussian mixture: a) momentum $p_{\rm in}$ evaluated at the vertex (full histogram) and $p_{\rm out}$ at the position of the outermost hit (dashed-dotted histogram) for $p_T^e = 10$ GeV/c; b) difference $p_{\rm in} - p_{\rm out}$ against the total energy radiated in the tracker volume.

This will find major applications in distinguishing various electron patterns, to improve electron energy measurements and electron identification. In particular, classes of "low-radiating" electrons can be defined in a novel way and used for ECAL inter-calibration and absolute calibration purposes (Section 4.4).

10.4.3 Electron classes

Tracking and calorimetry observables can be found which are, on event-by-event, sensitive to the integral amount of radiated energy, to the bremsstrahlung emission pattern along the electron trajectory, or to effects induced by energy lost in the tracker material. Using such observables for the separation of the electron population into distinct classes, will allow to better disentangle the non-Gaussian sources of fluctuations of the ECAL supercluster energy containment and tracker momentum measurement. Distinct track-supercluster patterns will, in general but especially at low $p_{\rm T}^e$, imply different energy-momentum measurement errors (Section 10.4.4) and different electron identification performances (Section 10.4.7).

A set of electron classes was introduced in [245]. A useful measurement of the amount of bremsstrahlung is obtained by the relative difference between the momentum measured at the last point ($p_{\rm out}$) on the electron track and the momentum measured at the origin ($p_{\rm in}$). The ϕ match between the reconstructed track and the supercluster is sensitive to the

bremsstrahlung collection. The matching between the total energy $E_{\rm sc}$ collected by the supercluster with the momentum measured at the track origin is sensitive to the energy lost in the tracker material. These and other observables are used to separate "well measured" and "badly" measured electrons in a sequence described in the following.

The most precisely measured electrons fall in the category:

- **golden electrons.** This class represents electrons least affected by radiation emission, with a reconstructed track well matching the supercluster and a well behaved supercluster pattern. It is defined as:
  - A supercluster constituted by a single "seed" cluster (i.e. without observed bremsstrahlung subcluster);
  - A measured bremsstrahlung fraction $f_{\rm brem} = (p_{\rm in} - p_{\rm out})/p_{\rm in}$ below 0.2;
  - A ϕ matching between the supercluster position and the track extrapolation from last point within $\pm$ 0.15 rad;
  - An $E_{\rm sc}/p_{\rm in}$ value in excess of 0.9.

The "golden" electrons are dominantly truly "low radiating" electrons for which all measurement issues greatly simplifies. The usage of a subset of golden electrons originating from high statistics physics processes like single W and single Z production at the LHC has been considered for in situ energy calibration of the ECAL crystals (Section 4.4).

But other classes of electrons can also be very well measured and identified. On event-by-event, this is irrespective of the total amount of bremsstrahlung radiated away. Electrons which fail to satisfy the golden electron requirements might be accepted as:

- **big brem electrons.** This class contains electrons with a good energy-momentum matching between the supercluster and the electron track at origin, a well behaved supercluster pattern, and no evidence of effects from secondary photon conversion despite a very large measured bremsstrahlung fraction. Electrons for which all the bremsstrahlung is radiated in a single step, either very early or very late when crossing tracker silicon layers, can fall in this category. The class is defined as:
  - A supercluster constituted by a single "seed" cluster;
  - A $f_{\rm brem}$ above 0.5;
  - An $E_{\rm sc}/p_{\rm in}$ value between 0.9 and 1.1.

A complementary set of electrons, still with a good energy-momentum (at origin) matching, but which fails some criteria for golden and big brem can fall in the category:

- **narrow electrons.** In this intermediate class, electrons have a significantly large bremsstrahlung fraction but not has high as for big brem, a rather well behaved supercluster (i.e. the bremsstrahlung has been merged in 1 cluster), but, as for big brem, a relaxed track-supercluster geometrical matching. It is defined as:
  - A supercluster constituted by a single "seed" cluster;
  - An $E_{\rm sc}/p_{\rm in}$ value between 0.9 and 1.1;
  - A measured bremsstrahlung fraction and/or a ϕ matching outside the range of golden and big brem electrons.

Finally, the remaining electrons ("bad") electrons are classified as:

- **showering electrons.** This class contains electrons which failed to enter any of the above classes. This will be very likely for instance in cases of secondary conversion of some early radiated bremsstrahlung for electrons having radiated large fraction of their initial energy. The class includes electron supercluster patterns involving an identified bremsstrahlung subcluster, or cases where a bad energy-momentum E/p matching is observed.

The fraction of the initial electron population falling in the different classes as a function of $|\eta|$ is shown in Fig. 10.24a for a sample of single isolated electrons. The shape of the distribution for the showering class clearly reflects the material budget η distribution. Conversely, the probability to observe golden electrons is, as expected, anti-correlated with the material budget. Overall, about 20% of the electrons are observed in the golden class while showering electrons represent more than 50% of the population. For simplicity, the electrons impacting in the vicinity of ECAL inter-module or supermodule borders of the barrel have not been included here. Such electrons are collected in a separate class and are susceptible of requir-

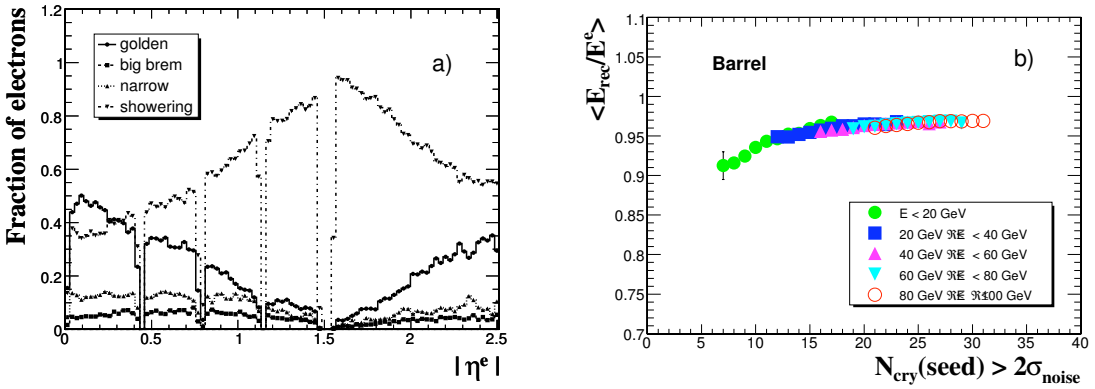


Figure 10.24: Classes of track-supercluster patterns and energy scale corrections for electrons uniformly distributed in energy between 5 and 100 GeV : a) fraction of population in the different classes as a function of $|\eta^e|$; b) mean of the raw reconstructed supercluster energy $E_{\rm rec}$ normalized to the generated electron energy $E^{\rm e}$ as a function of the number of crystals of the seed cluster having a signal above $2\sigma_{noise}$, in the η range of the ECAL barrel.

ing specific treatment. Whether or not an electron depositing energy near a "crack" can be measured with precision is evaluated by an algorithmic procedure described elsewhere [8].

10.4.4 Electron energy scale corrections

The electron energy measurement provided by the superclustering algorithm is obtained by simply adding the contributions from the energy deposits in each crystal. The ϕ spread of the energy deposition in the ECAL, and hence the energy containment of the supercluster, will in general depend on the pattern of bremsstrahlung emission and on $p_{\rm T}^e$. The variations of the supercluster energy containment associated to bremsstrahlung collection can be correlated with the volume of the shower integration, as measured approximately by the number of contributing crystals. The Fig. 10.24b shows the energy response in the ECAL barrel as function of the number of crystals in the seed cluster of the supercluster which have energy above 2 standard deviations of the electronic noise. Electrons from golden, big brem and narrow

classes all follow the same universal dependence. More details are provided in [245] where similar results are obtained for electrons in the η range of the ECAL endcaps. There, the electron energy is obtained by adding energy deposited in the preshower detector to the corrected ECAL supercluster. In the case of showering electrons for which the bremsstrahlung recovery is not fully efficient and which are affected by a significant fraction of energy lost in the tracker material, the correction based on $f(N_{cry})$ must be complemented by a residual parametrized η-dependent correction.

An example of the reconstructed energies for each electron classes, before and after corrections, is shown in Fig. 10.25 for electrons measured in the ECAL barrel and endcap and with generated energy E^e in the range 5 to 100 GeV.

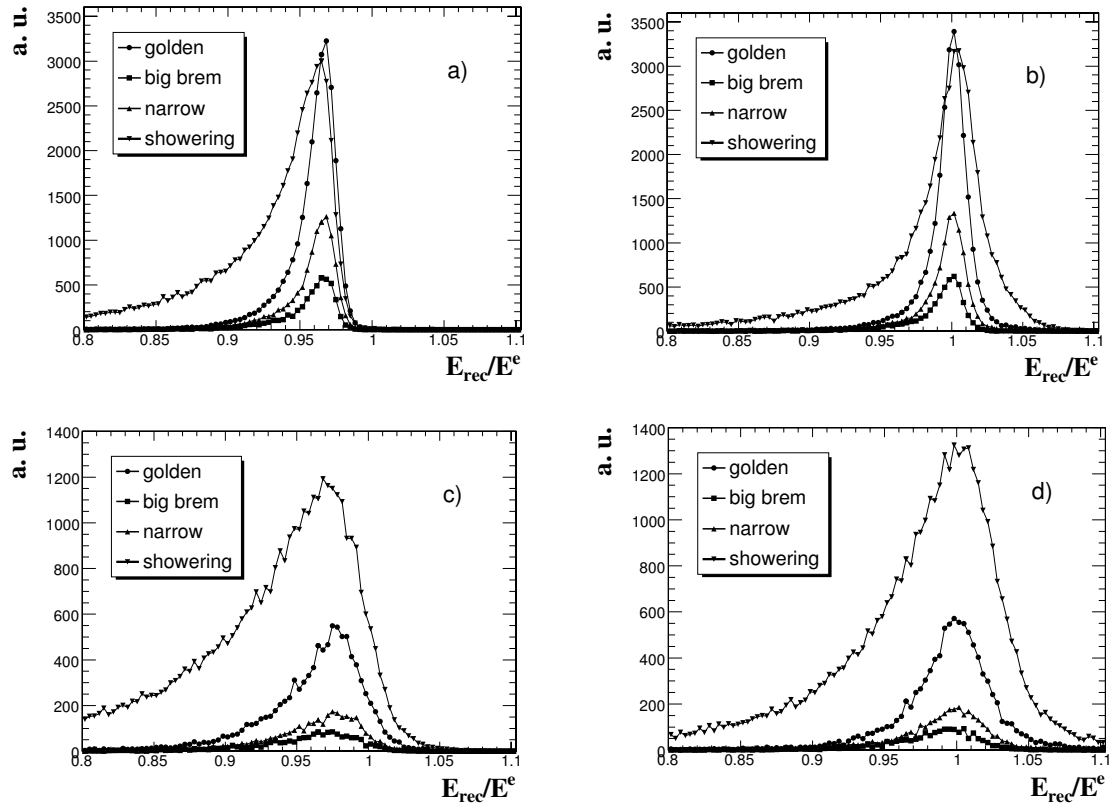


Figure 10.25: Distributions of the (a)(c) raw and (b)(d) fully corrected reconstructed energy $E_{\rm rec}$ normalized to the generated energy E^e, for electrons uniformly distributed in energy between 5 and 100 GeV. The distributions are shown separately for the η range of (a)(b) the ECAL barrel and (c)(d) ECAL endcaps, and for the different track-supercluster patterns. In the range of the endcaps, $E_{\rm rec}$ includes the energy deposited in the preshower detectors.

Besides the proper normalization of the energy response, the corrected energy distributions are slightly narrower, more symmetric and Gaussian than the uncorrected one. Overall, the effective RMS resolution improves by about 5% in the η range of the barrel 10% in the range of the endcaps. Different energy measurement errors are obtained for the different electron classes [245].

10.4.5 Track-cluster combination

10.4.5.1 Track-cluster matching

In a "pre-selection" step, a loose geometrical and energy-momentum matching is imposed between the reconstructed electron track and the corresponding supercluster. This complements the loose track-supercluster matching which is implicitly built in the supercluster-driven pixel seed finding described in section 10.4.2.1. The "electronicity" of the candidate electrons is further re-enforced by setting an upper threshold on the fraction of the electron supercluster energy collected in the hadron calorimeter.

Thus, electron candidates can therefore be pre-selected as:

- a reconstructed electron track initiated by the reconstruction of a supercluster in the ECAL matched with hits in the pixel detector;
- an energy-momentum matching between the supercluster and the track, $E_{\rm rec}/p_{\rm in} < 3$;
- an η geometrical matching $|\Delta\eta_{\rm in}| = |\eta_{sc} - \eta^{\rm extrap.}_{\rm track}| < 0.1$, where η_{sc} is the supercluster η position and $\eta^{\rm extrap.}_{\rm track}$ is the track pseudorapidity at the closest position to the supercluster position;
- a ϕ geometrical matching $|\Delta\phi_{\rm in}| = |\phi_{sc} - \phi^{\rm extrap.}_{\rm track}| < 0.1$, where ϕ_{sc} is the supercluster ϕ position and $\phi^{\rm extrap.}_{\rm track}$ is the track ϕ position at the closest position to the supercluster position;
- a ratio of the energy deposited in the HCAL tower just behind the electromagnetic seed cluster over the energy of the seed cluster $H/E < 0.2$.

The efficiency of such pre-selection, can be evaluated for illustration from a sample of isolated electrons obtained from the final state in single production of Standard Model Higgs bosons decaying via $H \to ZZ^{(*)} \to 4e$. It reaches a plateau around 93% for central η values. The pre-selection entails a drop of efficiency of about 3 to 5% depending on η when compared to electron track reconstruction efficiencies alone. A more severe drop of efficiency of about 20% occurs around $|\eta| \simeq 1.5$ due to a deterioration of the track-supercluster matching in the transition range between the ECAL barrel and endcaps.

10.4.5.2 Combination of energy and momentum measurements

The corrected energy measurement $E_{\rm rec}$ provided by electromagnetic calorimeter can be combined with the tracker momentum measurement $p_{\rm rec}$ to improve the estimate of the electron momentum at the interaction vertex.

The improvement is expected to come both from the opposite behaviour with E^e (or $p^e_{\rm T}$) of the intrinsic calorimetry (σ_E/E) and tracking ($\sigma_{p_{\rm T}}/p_{\rm T}$) fractional resolutions, and from the fact that $p_{\rm rec}$ $E_{\rm rec}$ are differently affected by the bremsstrahlung radiation. This latter point is illustrated in Fig. 10.26 for electrons in the η range of the ECAL barrel. Similar plots are obtained in the η range of the endcaps [245]. Cases with $E_{\rm rec}/p_{\rm rec} > 1$ are almost always due to an underestimation of the track momentum. An early emission of a large amount of radiation along the electron trajectory is likely to lead to an underestimate of the initial momentum by a track reconstruction algorithm. Cases with $E_{\rm rec}/p_{\rm rec} < 1$ can be due either to an underestimate of the true energy by the calorimetry measurement (e.g. due to

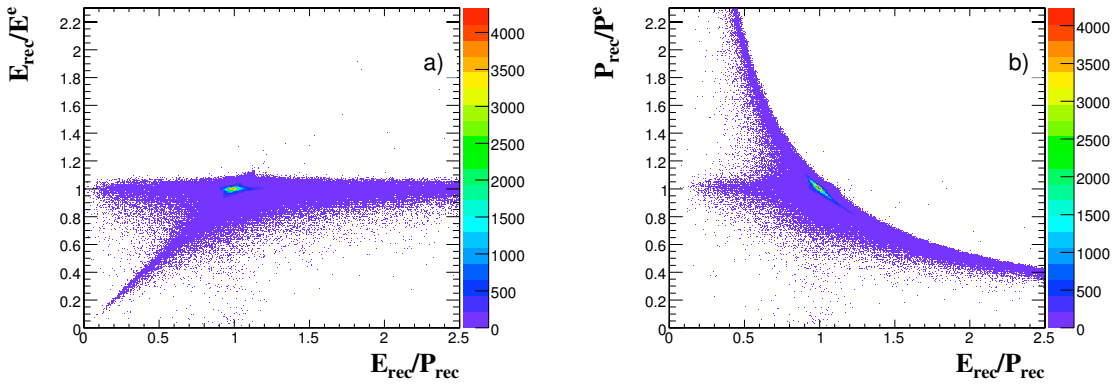


Figure 10.26: Correlations between ECAL energy and tracker momentum measurements in the η range of the barrel: a) corrected supercluster energy $E_{\rm rec}$ normalized to the generated energy E^e as a function of $E_{\rm rec}/P_{\rm rec}$; b) reconstructed track momentum at vertex $P_{\rm rec}$ normalized to the generated momentum P^e as a function of $E_{\rm rec}/P_{\rm rec}$.

energy lost from the conversion of secondary photons) or to an overestimate of the true initial track momentum. Otherwise, the tracking and calorimetry measurements are always both providing a good measurement when $E_{\rm rec}/p_{\rm rec} \simeq 1$.

For electrons $p_{\rm T}^e$ below $\approx$ 30 GeV/c the different electron classes determine the weight of the tracking and calorimetry information for the combined measurement of the energy-momentum. The energy resolution observed for different electron classes is illustrated in Fig. 10.27a. The best precision is obtained with golden and narrow electrons. The resolution for big brem electrons is only slightly degraded. In contrast, the showering electrons suffer from a severe deterioration of the resolution due to fluctuations of the energy lost in the tracking material and fluctuations of the supercluster energy containment.

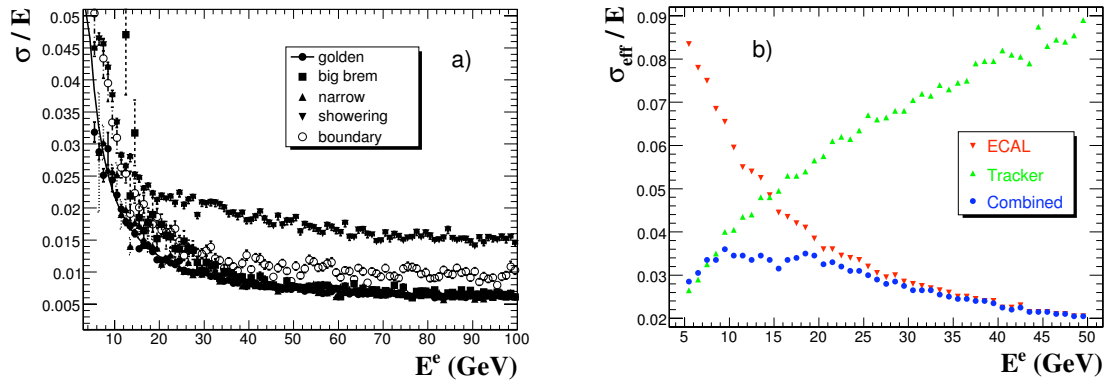


Figure 10.27: Precision on the electron energy measurements: a) fractional resolution (Gaussian fits) as a function of generated energy E^e for different track-supercluster patterns; b) fractional resolution (effective RMS) as a function of generated energy E^e measured with the ECAL supercluster (downward arrows), the track (upward arrows), and the combined track-supercluster (circles) .

The electron energy is finally determined by combining the corrected electron supercluster measurement E, the electron track momentum p, and their associated errors. In cases when $|E/p - 1| < 2\sigma_{E/p}$, it is is taken as the weighted mean of E and p with weights defined as the normalized inverse of the variance of each measurement. Otherwise, E alone is used when $E/p > 1 + 2\sigma_{E/p}$ or for $E > 15$ GeV when $E/p < 1 - 2\sigma_{E/p}$, while p alone is used for $E < 15$ GeV, when $E/p < 1 - 2\sigma_{E/p}$. The effective RMS resolution obtained from combined track-supercluster is illustrated in Fig. 10.27b and compared to the separate ECAL and tracker measurements. The precision is clearly improved with the combined estimate for energies between about 10 and 25 GeV.

10.4.6 Electron isolation

Lepton isolation can be imposed as a simple and powerful means of suppressing QCD background in multi-lepton physics channels at the LHC.

For electrons in CMS, the simplest and most powerful isolation criteria is obtained from tracks originating from a common (primary) vertex. Using track measurements at a primary vertex for the electron isolation avoids the complication due to the severe external bremsstrahlung, photon conversion, and early showering in the tracker material. It moreover allows to postpone the question of the identification of possible internal bremsstrahlung photons, associated to final state electrons, which might appear in the ECAL within an isolation cone. Track-based electron isolation must normally be complemented by electron identification requirements.

As an example of simple track isolation criteria, all tracks having $p_{\mathrm{T}} > 1.5$ GeV/c are considered in an $\eta - \phi$ isolation cone of radius $R_{cone} = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ centred on the reconstructed electron. The tracks are further required to have $|\Delta IP_L| < 0.1$ cm, where ΔIP_L is the difference between the longitudinal impact parameter and the z position of the primary vertex. The electron isolation variable is then defined as the sum of the p_{T} of all the tracks satisfying these requirements (excluding the electron track), divided by the electron p_{T}.

The performances of such a track-based electron isolation is illustrated in Fig. 10.28a by considering the suppression of the $t\bar{t}$ background to the Standard Model Higgs boson in the channel $H \rightarrow ZZ^* \rightarrow e^+e^-e^+e^-$. The observation of 4 isolated electrons is imposed. The efficiency for the Higgs signal after preselection is plotted as a function of the rejection obtained against the $t\bar{t}$ background.

A further requirement can be imposed on the normalized transverse impact parameter (IP_T/σ_{IP_T}) to further refine the selection of primary electrons. This is illustrated in Fig. 10.28b again for the Higgs channel in 4 electrons compared to the $t\bar{t}$ and $Zb\bar{b}$ backgrounds when involving electrons from semi-leptonic b decays.

10.4.7 Electron identification

Electron identification makes use of a complete set of "electronicity" estimators. These are combined to establish full compatibility of the observations with expectations from single electrons. The performance (efficiency, rejection power, purity) of this identification depends of course on the degree of isolation imposed on the electron candidates and on the nature of the considered background. It also depends on the quality requirements imposed on the electron objects themselves. The identification and rejection performances will in general

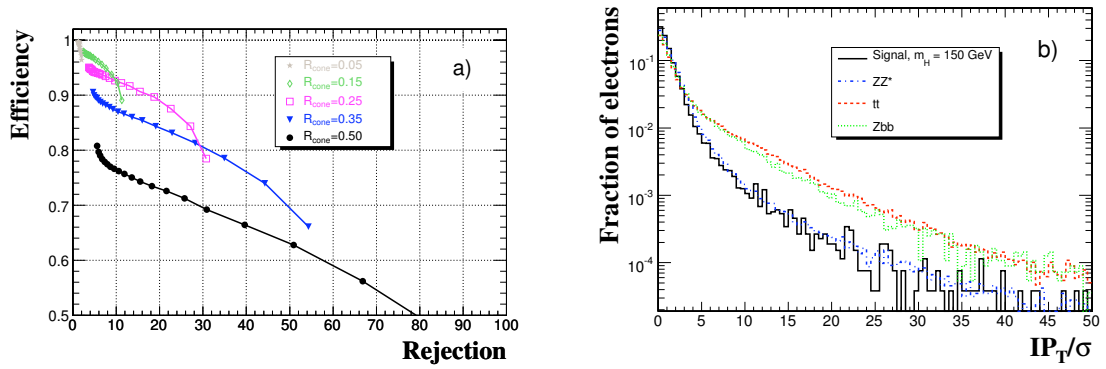


Figure 10.28: Selection of isolated primary electrons. a) Signal efficiency versus background rejection power from electron track isolation alone for SM Higgs $4e$ signal and $t\bar{t}$ background, for various track isolation parameters and cone sizes. b) Distribution of the normalized transverse impact parameter I_p/σ_{I_p} for electrons from the SM Higgs boson decaying in $H \to ZZ^* \to e^+e^-e^+e^-$ and for electrons from b jets in $t\bar{t} \to W^+W^-b\bar{b}$ and $Zb\bar{b}$ background.

vary, in particular at low p_{T}^e, with the electron class. The "well measured" and the "badly measured" electrons are likely to be differently affected by possible fake background sources. Finally, the distinction between multi-clusters and single cluster electron patterns is expected to play an important role in the separation of electron from "fake" electrons in QCD jets formed by overlapping particles

Figure 10.29 shows some typical distributions of electron identification variables exhibiting different characteristics for different electrons classes. They are shown both for a sample of real isolated electrons and for "fake" electrons from a sample QCD di-jets enriched in events which would pass the Level 1 electromagnetic trigger of CMS [245]. From such distributions, one arrives at a prescription on the use of identification variables depending on the different electron classes:

- the $\sigma_{\eta\eta}$ shape variable is discriminating for all electron classes;
- H/E and $\Delta\eta_{\mathrm{in}}$ and $\Delta\phi_{\mathrm{in}}$ with track parameters at the vertex and using supercluster energy and η and ϕ positions are discriminating for all electron classes, with a slight loss of discriminating power for Showering electrons;
- $\sigma_{\phi\phi}$ and Σ_9/Σ_{25} shower shape variables involving ϕ projection gives discriminating power for all but Showering electrons;
- E/p_{in} with track parameters at the vertex and using supercluster energy and E/p_{out} with track momentum from the outermost track position and energy from the electron sub-cluster are more discriminating for the Showering class;
- $\Delta\phi_{\mathrm{out}}$ with track momentum from the outermost track position and energy from the electron sub-cluster discriminating only for Golden electrons.

Using for instance a simple sequence of cuts adapted to each pre-selected electron class, an absolute overall jet background fake rate efficiency at the level of 6×10^{-4} is obtained while keeping efficiency on electrons of p_T from 5 to 50 GeV/c at the level of 90%.

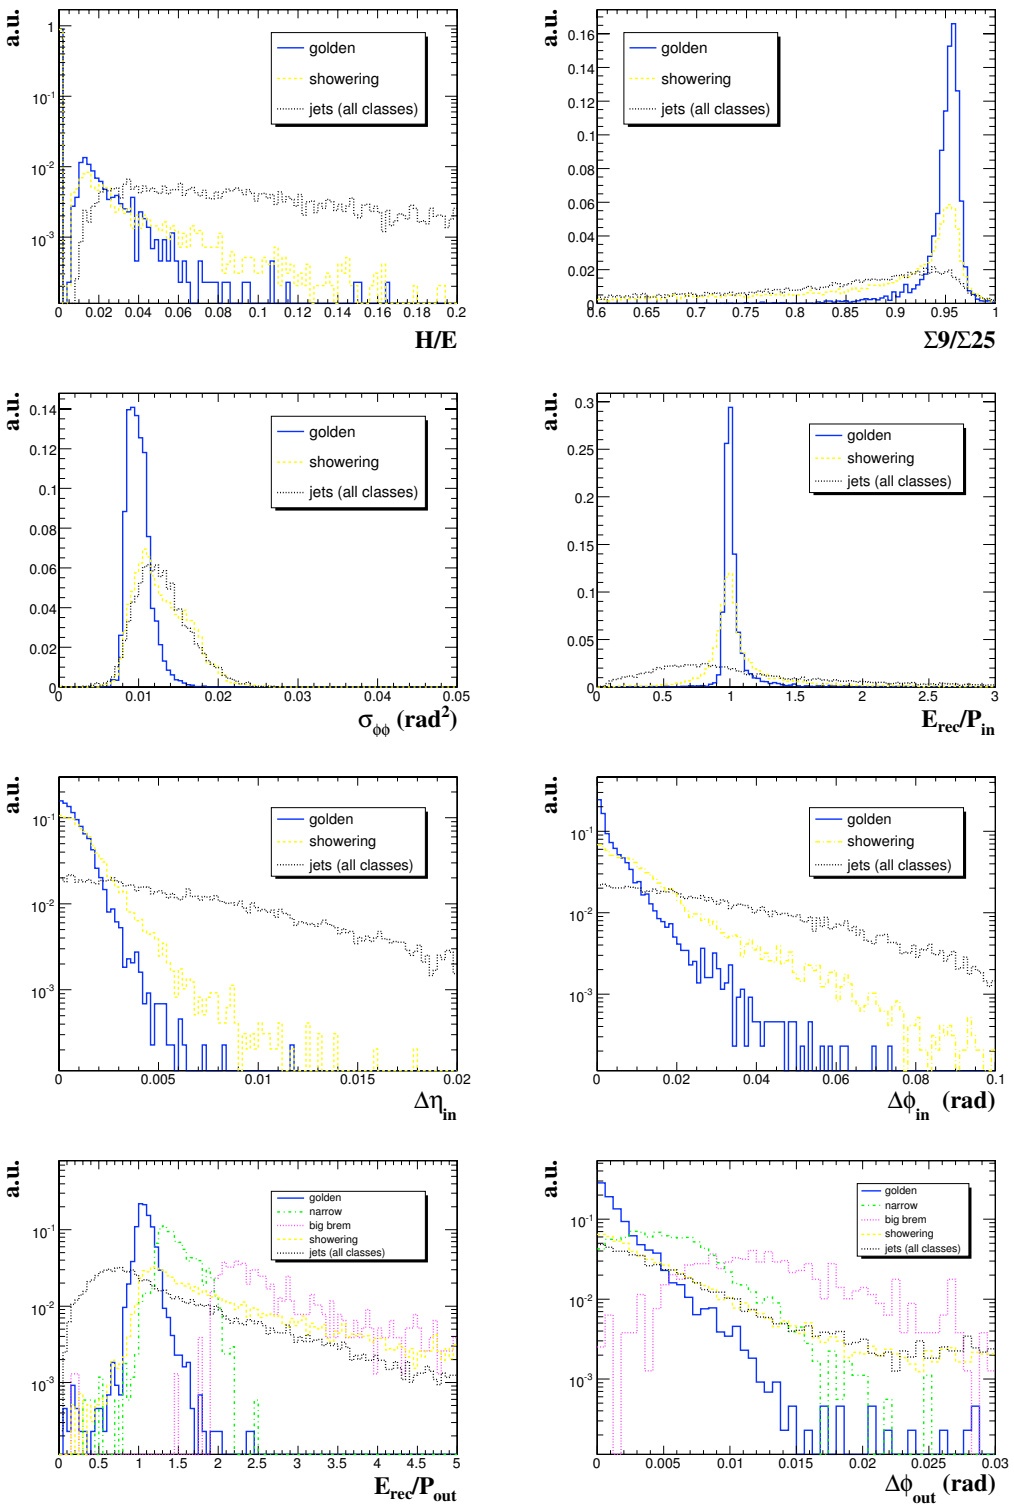


Figure 10.29: Distributions of electron identification observables in the η range of the ECAL barrel . The distributions of shower shape and track-supercluster matching observables are shown for different classes of electrons. Electrons uniformly distributed in p_T^e between 5 and 50 GeV/c are compared to "fake" electron candidates found in jets from a QCD di-jet background generated in the range $25 < \hat{p}_T < 170$ GeV/c.

Chapter 11

Jets and Missing Transverse Energy

The huge QCD cross section (Fig. 11.1) ensures that jets will dominate high-p_{T} physics at the LHC. Jets will not only provide a benchmark for understanding the detector, but will also serve as an important tool in the search for physics beyond the standard model. Event signatures for SUSY, Higgs boson production, compositeness, and other new physics processes require accurate reconstruction and measurement of jets coming from high-p_{T} quarks and gluons [246]–[249]. The problems with associating a jet measured in a calorimeter with a scattered parton is an old, persistent problem in hadron collisions [250]–[256]. Jet energy resolution and linearity are key factors in separating signal events from backgrounds. Missing transverse energy resolution, which historically has played an important role in the discoveries of the W boson and the top quark and the search for new phenomena at hadron colliders, is closely related to the calorimeter jet energy response.

A detailed description of the calorimeter response, including pulse shape, digitization and zero suppression have been simulated with high statistics. Monte Carlo simulation samples of fully simulated events at the detector level have been used to study jet and missing transverse energy response. The calibration procedures described in Chapter 5, together with these simulations, give an idea of the jet and missing transverse energy performance at startup. The commissioning phase of the detector will have a tremendous impact on understanding the calorimeter reponse and will help to refine plans for data-driven calibrations and jet energy scale determination.

11.1 Tower definition and thresholds

Readout cells in HCAL are arranged in a tower pattern in η, ϕ space, projective to the nominal interaction point. The cells in the barrel region have segmentation of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$, becoming progressively larger in the endcap and forward regions. Since the ECAL granularity is much finer than HCAL, calorimeter towers (ECAL plus HCAL) are formed by addition of signals in η, ϕ bins corresponding to individual HCAL cells. In total there are 4176 such towers, which when unfolded, may be represented in a familiar "lego" plot (Fig. 11.2).

The towers are used as input to several jet clustering algorithms. The energy associated with a tower is calculated as the sum of all contributing readout cells which pass the online zero-suppression threshold and any additional offline software thresholds. For the purpose of jet clustering, the towers are treated as massless particles, with the energy given by the tower energy, and the direction defined by the interaction point and the center of the tower.

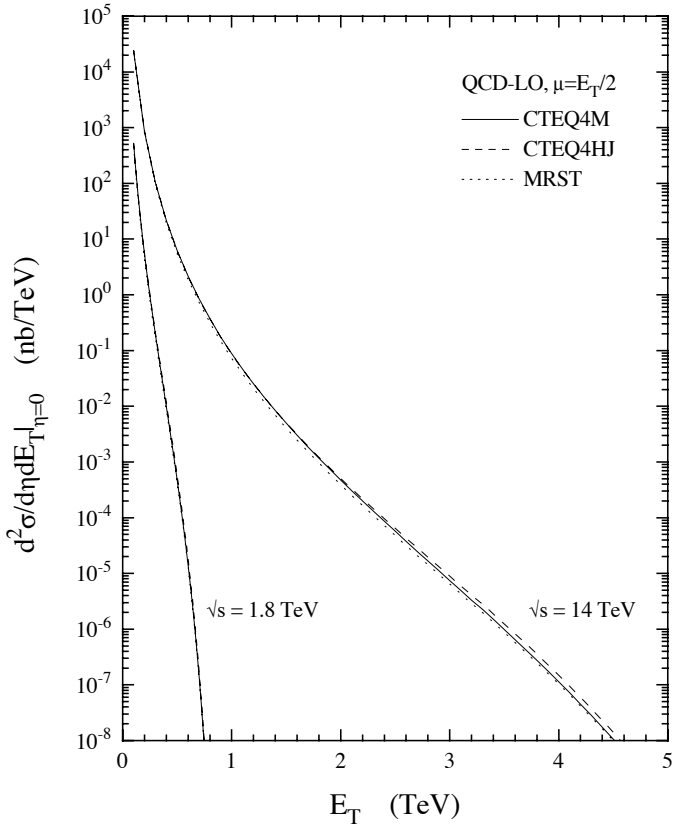


Figure 11.1: Inclusive jet cross section vs. jet transverse energy at the LHC compared with the Tevatron [257]. The cross section is 7 orders of magnitude greater at the LHC than at the Tevatron kinematic limit, and the luminosity will also be more than 2 orders of magnitude greater.

Optimum performance of higher-level objects reconstructed from calorimeter towers requires careful selection of these inputs because calorimeter noise contributions can have significant impact on the reconstruction of low-E_{T} jets. Various schemes of suppressing contributions of noise and pile-up to jet energies have been studied in detail [258], based on simulation of calorimeter response as implemented in ORCA. These studies include variation of thresholds on the towers as well as the individual cells which constitute towers.

Usually either a transverse energy cut $E_{\mathrm{T}} > 0.5$ GeV or $E_{\mathrm{T}} > 0.5$ GeV and energy cut $E > 0.8$ GeV (scheme T) were applied to all towers used in jet reconstruction. While both approaches give similar results for high-E_{T} jets, the latter scheme eliminates more noise in the central η region, and is hence preferred for jet reconstruction at low E_{T}.

To further refine the noise rejection, energy thresholds applied to individual cells were investigated. It is a natural choice as the noise contribution depends on the type of the calorimeter compartment. HCAL cells show discrete ADC readout patterns (Fig. 5.8), with 1 ADC count corresponding to approximately 250 MeV in HB and HO, and 400 MeV HE. This quantization must be taken into account when setting thresholds. Three schemes referred to as A, B, and C in Table 11.1 were designed to increment the thresholds on HCAL cells in steps of an ADC count. As an example, the scheme A retains 1.4 GeV of noise in the $R = 0.5$ cone at

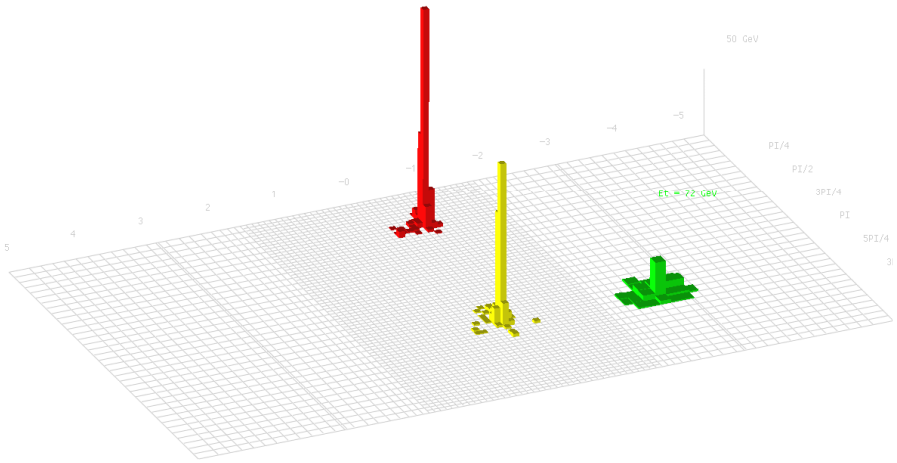


Figure 11.2: The η, ϕ segmentation of the CMS hadron calorimeter.

$\eta = 0$. In comparison, the tower-threshold scheme T has 0.6 GeV of noise, and 2.8 GeV more loss of jet energy on average in the QCD jet sample generated with $\hat{p}_{\mathrm{T}} > 40$ GeV/c (where $\hat{p}_{\mathrm{T}}$ is the parton p_{T} in the $2 \to 2$ scattering process in PYTHIA).

Table 11.1: Cell energy threshold schemes. Values are given in GeV. ΣEB (ΣEE) are thresholds on the total EB (EE) energy of all contributing ECAL crystals in a tower. NIC is the average noise-in-cone energy and JEL is the energy loss in a jet relative to scheme A for $R = 0.5$. NIC is measured in a noise sample and JEL is measured in a QCD40 sample without noise or pile-up included in the simulation. ΣE_{T} is the total transverse energy in the calorimeter.

Scheme	Thresholds					NIC	JEL	NIC	JEL	ΣE_{T}	ΣE_{T}
	HB	HO	HE	ΣEB	ΣEE	$\eta \simeq 0$	$\eta \simeq 0$	$\eta \simeq 2$	$\eta \simeq 2$	noise	QCD
A	0.7	0.85	0.9	0.2	0.45	1.4	–	1.1	–	28	168
B	0.9	1.1	1.4	0.2	0.45	0.3	1.0	0.4	2.7	6	162
C	1.2	1.3	1.8	0.2	0.45	0.2	1.9	0.3	5.2	4	158

The individual cell energy threshold schemes A and B decrease the noise to a manageable level while losing minimal amount of real energy. These 2 schemes should be used for noise suppression online or in future jet reconstruction and physics studies.

To compare jet reconstruction efficiencies for different threshold schemes, the minimum value of reconstructed jet E_{T} has been adjusted to give 50% efficiency at 20 GeV for each scheme separately; this procedure is intended to compensate for differences in the calibrations. These (uncorrected) jet thresholds are shown in Fig. 11.3 a as a function of generated jet E_{T} where the efficiency condition is imposed. The resulting efficiency curves (Fig. 11.3 b) are remarkably similar in shape indicating that the corresponding jet resolutions for the different threshold schemes are similar.

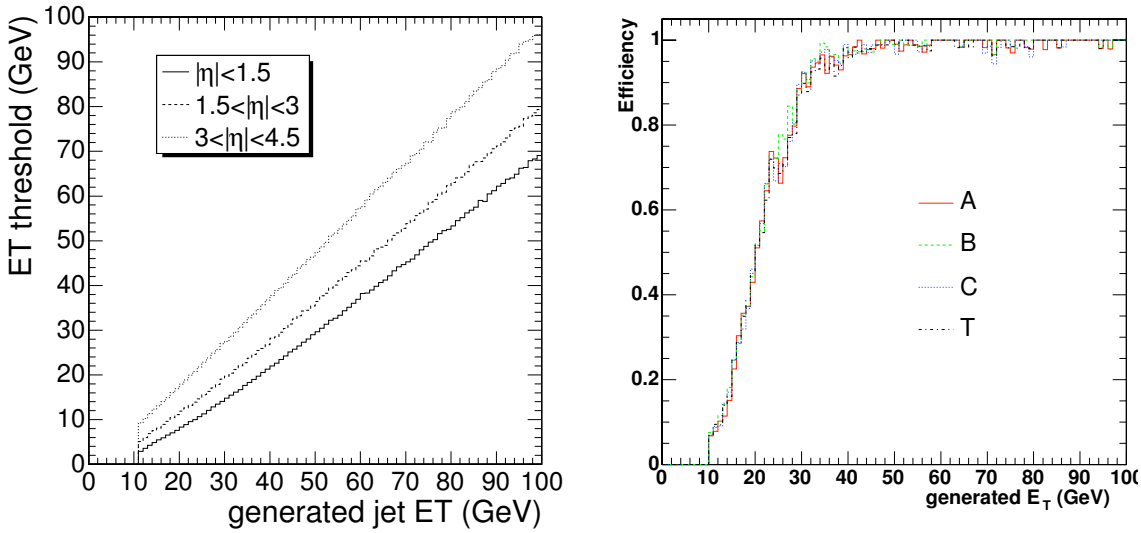


Figure 11.3: Left) Jet E_{T} threshold needed to reconstruct jets with 50% efficiency for the tower threshold scheme (T) with $E > 0.8$ GeV and $E_{\mathrm{T}} > 0.5$ GeV. Right) Jet efficiency curves for different threshold schemes and central η.

11.2 Jet algorithms

The first jet algorithms for hadron physics were simple cones [250, 259]. Over the last two decades, clustering techniques have greatly improved in sophistication. Three principal jet reconstruction algorithms have been coded and studied for CMS: the iterative cone [260], the midpoint cone [261] and the inclusive k_{T} jet algorithm [262, 263]. The midpoint-cone and k_{T} algorithms are widely used in offline analysis in current hadron collider experiments, while the iterative cone algorithm is simpler and faster and commonly used for jet reconstruction in software-based trigger systems.

The jet algorithms may be used with one of two recombination schemes for adding the constituents. In the energy scheme, constituents are simply added as four-vectors. This produces massive jets. In the E_{T} scheme, massless jets are produced by equating the jet transverse momentum to the ΣE_{T} of the constituents and then fixing the direction of the jet in one of two ways: 1) $\sin\theta = \Sigma E_{\mathrm{T}}/E$ where E is the jet energy (usually used with cone algorithms), or 2) $\eta = \Sigma E_{\mathrm{T}i}\eta_i/\Sigma E_{\mathrm{T}}$ and $\phi = \Sigma E_{\mathrm{T}i}\phi_i/\Sigma E_{\mathrm{T}}$ (usually used with the k_{T} algorithm). In all cases the jet E_{T} is equal to $p_{\mathrm{T}}c$.

The inclusive k_{T} algorithm merges, in each iteration step, input objects into possible final jets and so the new jet quantities, the jet direction and energy, have to be calculated directly during the clustering. The cone jet algorithms, iterative and midpoint, group the input objects together as an intermediate stage and the final determination of the jet quantities (recombination) is done in one step at the end of the jet finding.

11.2.1 Iterative cone

In the iterative cone algorithm, an E_{T}-ordered list of input objects (particles or calorimeter towers) is created. A cone of size R in η,ϕ space is cast around the input object having the

largest transverse energy above a specified seed threshold. The objects inside the cone are used to calculate a "proto-jet" direction and energy using the E_{T} scheme. The computed direction is used to seed a new proto-jet. The procedure is repeated until the energy of the proto-jet changes by less than 1% between iterations and the direction of the proto-jet changes by $\Delta R < 0.01$. When a stable proto-jet is found, all objects in the proto-jet are removed from the list of input objects and the stable proto-jet is added to the list of jets. The whole procedure is repeated until the list contains no more objects with an E_{T} above the seed threshold. The cone size and the seed threshold are parameters of the algorithm. When the algorithm is terminated, a different recombination scheme may be applied to jet constituents to define the final jet kinematic properties.

11.2.2 Midpoint cone

The midpont-cone algorithm was designed to facilitate the splitting and merging of jets. The midpoint-cone algorithm also uses an iterative procedure to find stable cones (proto-jets) starting from the cones around objects with an E_{T} above a seed threshold. In contrast to the iterative cone algorithm described above, no object is removed from the input list. This can result in overlapping proto-jets (a single input object may belong to several proto-jets). To ensure the collinear and infrared safety of the algorithm, a second iteration of the list of stable jets is done. For every pair of proto-jets that are closer than the cone diameter, a *midpoint* is calculated as the direction of the combined momentum. These midpoints are then used as additional seeds to find more proto-jets. When all proto-jets are found, the splitting and merging procedure is applied, starting with the highest E_{T} proto-jet. If the proto-jet does not share objects with other proto-jets, it is defined as a jet and removed from the proto-jet list. Otherwise, the transverse energy shared with the highest E_{T} neighbor proto-jet is compared to the total transverse energy of this neighbor proto-jet. If the fraction is greater than f (typically 50%) the proto-jets are merged, otherwise the shared objects are individually assigned to the proto-jet that is closest in η,ϕ space. The procedure is repeated, again always starting with the highest E_{T} proto-jet, until no proto-jets are left. This algorithm implements the energy scheme to calculate the proto-jet properties but a different recombination scheme may be used for the final jet. The parameters of the algorithm include a seed threshold, a cone radius, a threshold f on the shared energy fraction for jet merging, and also a maximum number of proto-jets that are used to calculate midpoints.

11.2.3 Inclusive k_{T} algorithm

The inclusive k_{T} jet algorithm is a cluster-based jet algorithm. The cluster procedure starts with a list of input objects, stable particles or calorimeter cells. For each object i and each pair (i, j) the following distances are calculated:

$$\begin{aligned} d_i &= (E_{\mathrm{T},i})^2 R^2, \\ d_{ij} &= \min\{E^2_{\mathrm{T},i}\,,\, E^2_{\mathrm{T},j}\} R^2_{ij} \quad \text{with} \quad R^2_{ij} = (\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2, \end{aligned}$$

where R^2 is a dimensionless parameter normally set to unity [261]. The algorithm searches for the smallest d_i or d_{ij}. If a value of type d_{ij} is the smallest, the corresponding objects i and j are removed from the list of input objects. They are merged using one of the recombination schemes listed below and filled as one new object into the list of input objects. If a distance of type d_i is the smallest, then the corresponding object i is removed from the list of input

objects and filled into the list of final jets. The procedure is repeated until all objects are included in jets. The algorithm successively merges objects which have a distance $R_{ij} < R$. It follows that $R_{ij} > R$ for all final jets i and j.

11.3 Monte Carlo corrections

The jet response was studied with fully simulated QCD dijet events over the range $0 < \hat{p}_{\mathrm{T}} < 4000$ GeV/c [264]. Jets were reconstructed using the iterative cone algorithm ($R = 0.5$), midpoint cone, and cluster-based k_{T} techniques using the E_{T} scheme. Comparisons between Monte Carlo simulation particle-level and reconstructed jets were made by applying the same jet algorithm to stable particles (excluding neutrinos and muons) and calorimeter cells, respectively. A matching criterion, based on the distance $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} < 0.2$, was used to associate Monte Carlo particle-level and reconstructed jets.

The data were divided into η bins where the ratio of reconstructed jet transverse energy ($E_{\mathrm{T}}^{\mathrm{rec}}$) to the Monte Carlo particle-level jet transverse energy ($E_{\mathrm{T}}^{\mathrm{MC}}$), as a function of $E_{\mathrm{T}}^{\mathrm{MC}}$ was fit using an iterative procedure [264]. Figure 11.4 shows the ratio R_{jet} as a function of pseudorapidity for different generated jet p_{T} before Monte Carlo corrections.

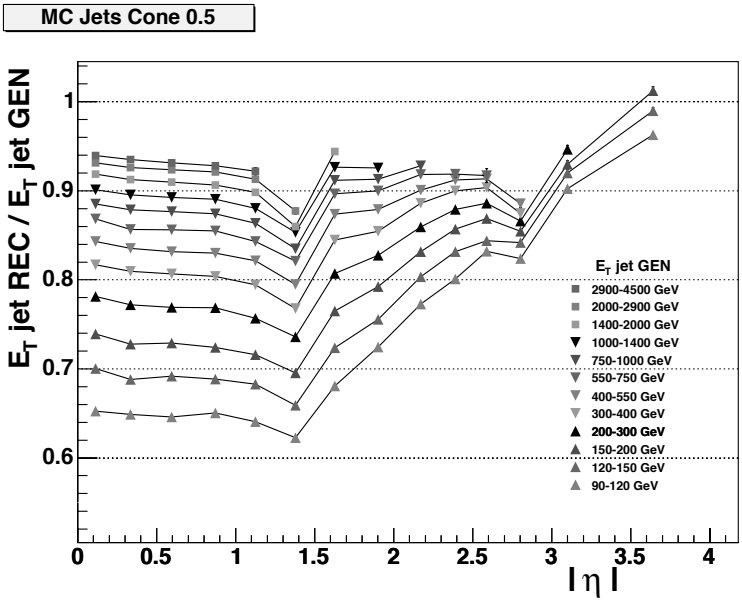


Figure 11.4: The ratio of the reconstructed jet transverse energy $E_{\mathrm{T}}^{\mathrm{rec}}$ to the generated transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$ as a function of generated jet η for jets with different $E_{\mathrm{T}}^{\mathrm{MC}}$ reconstructed by the iterative cone $R = 0.5$ algorithm before Monte Carlo corrections.

11.4 Jet resolution

The jet resolution was determined from a sample of QCD dijet events with parton transverse momenta ($\hat{p}_{\mathrm{T}}$) in the range 0–4000 GeV/c generated with PYTHIA (version 6.226). The events were fully simulated, digitized, and reconstructed assuming low luminosity conditions ($\mathcal{L} = 2 \times 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$) [265]. The events were divided into 21 bins of $\hat{p}_{\mathrm{T}}$ with a statis-

tics of 10^4 events per $\hat{p}_{\mathrm{T}}$ bin. All jets reconstructed in these events are included in the resolution fits. For the purpose of evaluating the linearity of the jet response, particle-level jets were reconstructed from all stable particles (excluding neutrinos and muons) using all 3 jet algorithms: the iterative cone algorithm with a cone size $R = 0.5$, the cluster-based k_{T} algorithm, and the midpoint-cone algorithm. The E_{T} recombination scheme was used. The particle-level jets are required to have $|\eta| < 5$, corresponding to the full η coverage of the calorimeters. A matching criterion based on the distance $R = 0.2$ in η, ϕ space was used to associate particle-level and reconstructed jets axes.

The reconstructed jet transverse energy ($E_{\mathrm{T}}^{\mathrm{rec}}$) was compared to the MC generated transverse energy ($E_{\mathrm{T}}^{\mathrm{MC}}$). The distribution of $E_{\mathrm{T}}^{\mathrm{rec}}/E_{\mathrm{T}}^{\mathrm{MC}}$ was fit to obtain the resolution as shown in Figure 11.5. The lower value of $E_{\mathrm{T}}^{\mathrm{rec}}$ compared to $E_{\mathrm{T}}^{\mathrm{MC}}$ is due in part to the fact that the ECAL is calibrated on photons, whereas a substantial amount of jet energy deposited in the ECAL arises from charged pions. For the lowest energy jets, the distribution of of $E_{\mathrm{T}}^{\mathrm{rec}}/E_{\mathrm{T}}^{\mathrm{MC}}$ is asymmetric and a fit is done in the vicinity of the peak position. The results of these fits provide an MC jet correction function that may be applied to reconstructed jets.

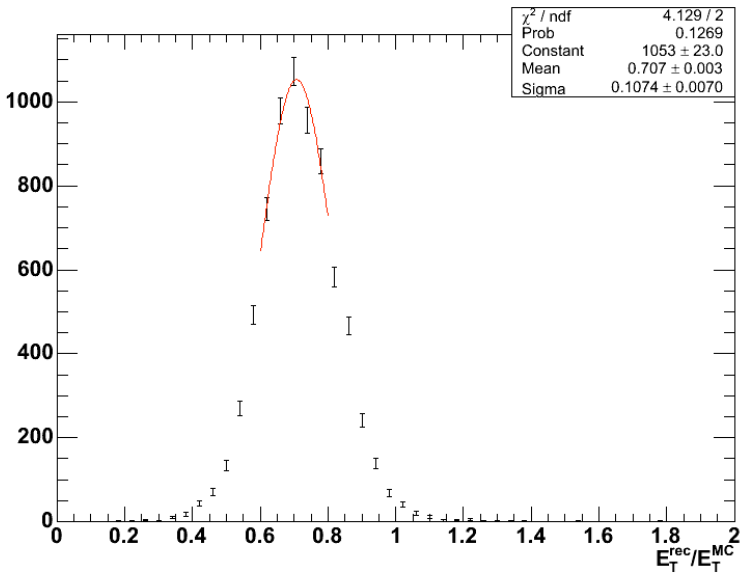


Figure 11.5: Distribution of reconstructed jet transverse energy ($E_{\mathrm{T}}^{\mathrm{rec}} = p_{\mathrm{T}}^{\mathrm{rec}}c$) divided by particle-level generated jet transverse energy ($E_{\mathrm{T}}^{\mathrm{MC}} = p_{\mathrm{T}}^{\mathrm{MC}}c$) for generated jets in the range $105\,\mathrm{GeV} < E_{\mathrm{T}}^{\mathrm{MC}} < 115\,\mathrm{GeV}$. The jets are reconstructed with the iterative cone $R = 0.5$ algorithm.

The resolution plots were fitted with the following functional form:

$$\frac{\sigma\left(\frac{E_{\mathrm{T}}^{\mathrm{rec}}}{E_{\mathrm{T}}^{\mathrm{MC}}}\right)}{\langle\frac{E_{\mathrm{T}}^{\mathrm{rec}}}{E_{\mathrm{T}}^{\mathrm{MC}}}\rangle} = \frac{a}{E_{\mathrm{T}}^{\mathrm{MC}}} \oplus \frac{b}{\sqrt{E_{\mathrm{T}}^{\mathrm{MC}}}} \oplus c \tag{11.1}$$

where the first term is due to fixed energy fluctuations in the cone from electronics noise, pile-up and underlying event energy, the second term comes from the stochastic response of the calorimeter measurements and the last term is the constant term from residual non-

uniformities and non-linearities in the detector response. The fits were done down to a transverse energies of 30 GeV in the barrel and endcap and 20 GeV in the forward region.

The resulting jet resolution for jets with $|\eta| < 1.4$ reconstructed with the iterative cone $R = 0.5$ algorithm is shown in Figure 11.6. The resolution curves for the barrel, endcap and forward regions are shown in Figure 11.7. The resolution curves on the measurement of ϕ and η of the jets for the barrel, endcap and very forward regions are shown in Figs. 11.8 and 11.9.

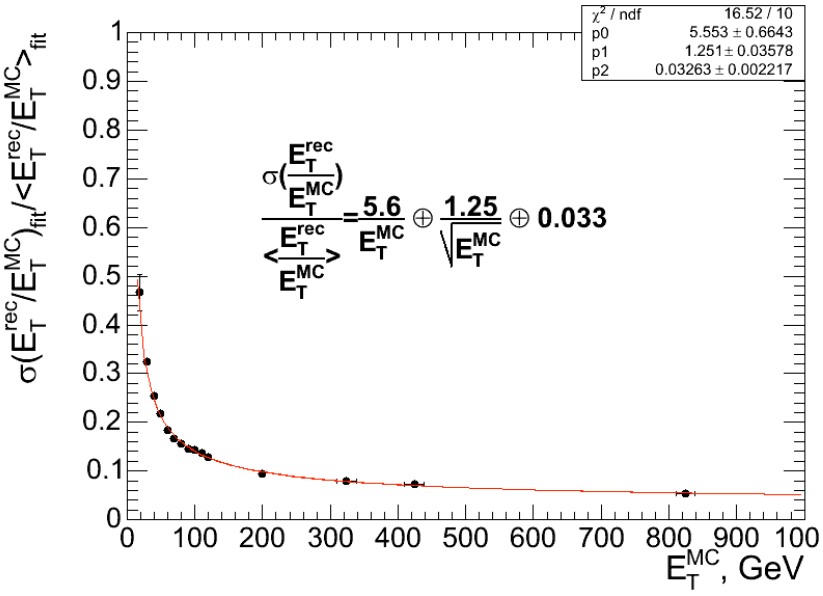


Figure 11.6: The jet transverse energy resolution as a function of the generated jet transverse energy for barrel jets ($|\eta| < 1.4$). The cuts $E_{\mathrm{T}} > 0.5$ GeV and $E > 0.8$ GeV are used. The distance between generated and reconstructed jets is $\Delta R < 0.2$. The Monte Carlo jet calibration has been applied.

11.5 Missing transverse energy

Beginning with UA1 [266], all major detectors at hadron colliders have been designed to cover as much solid angle as practically possible with calorimetry [3]. The primary motivation of this is to provide as complete of a picture as possible of the event, including the presence of one or more energetic neutrinos or other weakly-interacting stable particles though apparent missing energy. Energetic particles produced in the direction of the beam pipe make it impossible to directly measure missing energy longitudinal to the beam direction, however, the transverse energy balance can be measured with an accuracy good enough to help establish a physics signature involving one or more non-interacting particles. The W boson was discovered and its mass determined to 3% with just 6 events due to the ability of UA1 to infer the presence of 40 GeV neutrinos with a resolution of a few GeV [267]. Since the time of the W discovery, measurement of missing transverse energy has been a major tool in the search for new phenomena at hadron colliders [268]–[273].

Measurement of missing transverse energy vector ($\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$) at the LHC will be complicated by the presence of pile-up collisions. In CMS, measurement of $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ will be further degraded by

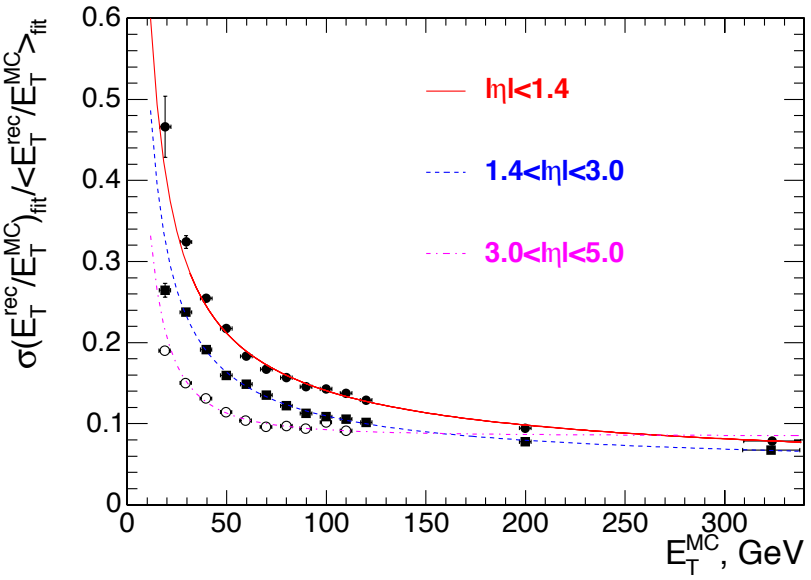


Figure 11.7: The jet transverse energy resolution as a function of the generated jet transverse energy for barrel jets ($|\eta| < 1.4$), endcap jets ($1.4 < |\eta| < 3.0$) and very forward jets ($3.0 < |\eta| < 5.0$). The jets are reconstructed with the iterative cone $R = 0.5$ algorithm. The cuts $E_{\mathrm{T}} > 0.5$ GeV and $E > 0.8$ GeV are used. The distance between generated and reconstructed jets is required to be $\Delta R < 0.2$. The Monte Carlo jet calibration has been applied.

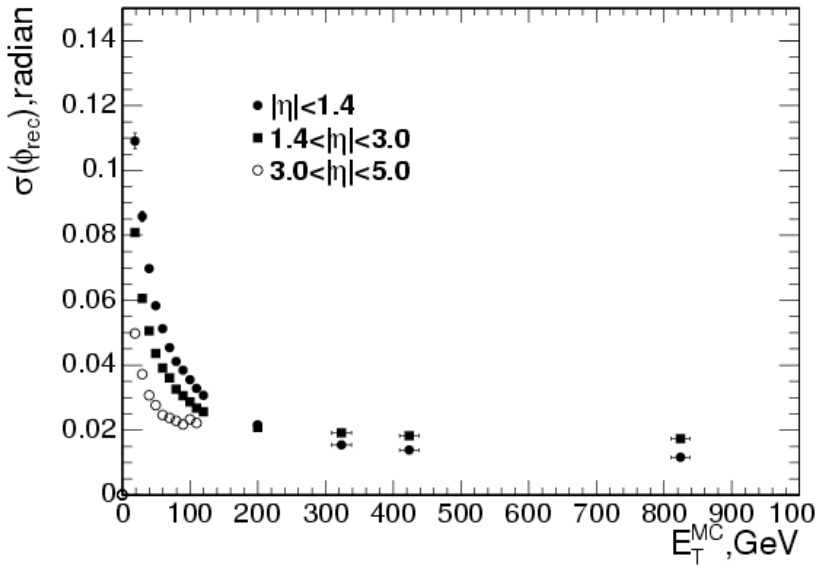


Figure 11.8: The jet ϕ angular resolution as a function of the generated jet transverse energy for barrel jets ($|\eta| < 1.4$), endcap jets ($1.4 < |\eta| < 3.0$) and very forward jets ($3.0 < |\eta| < 5.0$). The cuts $E_{\mathrm{T}} > 0.5$ GeV and $E > 0.8$ GeV are used. The distance between generated and reconstructed jets is required to be $\Delta R < 0.2$.

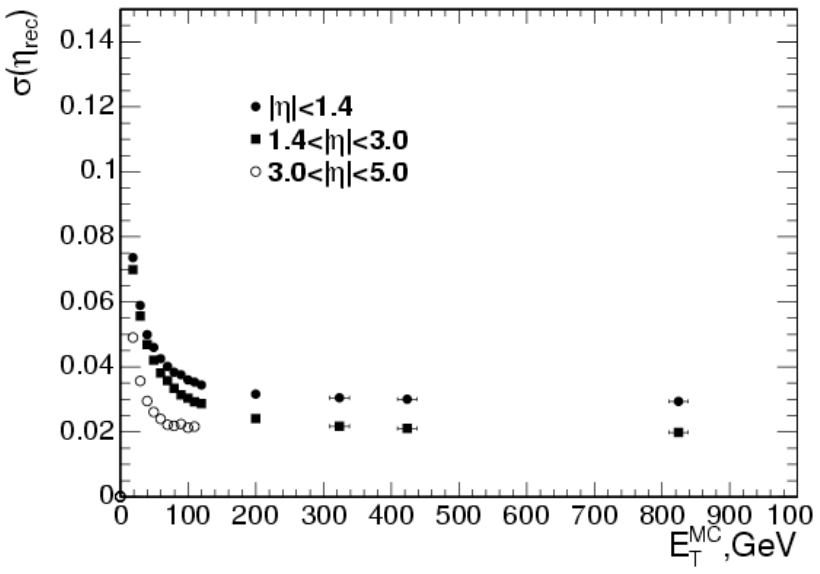


Figure 11.9: The jet η resolution as a function of the generated jet transverse energy for barrel jets ($|\eta| < 1.4$), endcap jets ($1.4 < |\eta| < 3.0$) and very forward jets ($3.0 < |\eta| < 5.0$). The cuts $E_{\mathrm{T}} > 0.5$ GeV and $E > 0.8$ GeV are used. The distance between generated and reconstructed jets is required to be $\Delta R < 0.2$.

the difference between photon and pion response in the combined ECAL plus HCAL detectors and by the bending of tracks by the 4 T magnetic field. On the other hand, the excellent cell segmentation, hermeticity, and good forward coverage of CMS will help measurement of $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$. In spite of all these detector subtleties, the $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ resolution in CMS is expected to be dominated by calorimeter resolution as discussed in Section 11.4 [274].

11.5.1 Comparison with UA1 and CDF

The missing transverse energy vector is calculated by summing individual calorimeter towers having energy E_n, polar angle θ_n and azimuthal angle ϕ_n:

$$\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}} = \Sigma(E_n \sin\theta_n \cos\phi_n \hat{\mathbf{i}} + E_n \sin\theta_n \sin\phi_n \hat{\mathbf{j}}) = E_x^{\mathrm{miss}} \hat{\mathbf{i}} + E_y^{\mathrm{miss}} \hat{\mathbf{j}} \,.$$

Reconstructed muons are taken into account by replacing the expected calorimeter deposit (about 4 GeV) with the reconstructed track p_{T}. It was observed in UA1 that the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution was dependent on the overall activity of the event, characterized by the scalar sum of transverse energy in all calorimeter cells (ΣE_{T}). The resolution is observed to follow the form $\sigma = C\sqrt{\Sigma E_{\mathrm{T}}}$ GeV$^{1/2}$ where C is a constant that depends directly on the calorimeter jet resolution and ΣE_{T} is the scalar sum of the transverse energy of all calorimeter hits. For UA1, the calorimeter (jet) resolution was approximately $0.8\sqrt{E_{\mathrm{T}}}$ GeV$^{1/2}$ which let to an observed distribution of x- or y- components, E_x^{miss} or E_y^{miss}, in minimum bias events that was Gaussian with zero mean and standard deviation $\sigma_x = \sigma_y = 0.4\sqrt{\Sigma E_{\mathrm{T}}}$ GeV$^{1/2}$ [267]. The CDF experiment, which has a scintillating tile geometry similar to CMS (and a completely different magnetic field configuration compared to UA1, solenoid vs. dipole), observed $\sigma_x = 0.47\sqrt{\Sigma E_{\mathrm{T}}}$ GeV$^{1/2}$ in Run I [275]. From the UA1 and CDF results and the CMS

jet resolution ($1.25\sqrt{E_{\rm T}}$ GeV$^{1/2}$) explained in Section 11.4, one may expect an $E_{\rm T}^{\rm miss}$ resolution in CMS of $\sigma_x \approx (0.6-0.7)\sqrt{\Sigma E_{\rm T}}$ GeV$^{1/2}$ for minimum bias events with no pile-up, when dominated by the shower fluctuations and not electronic noise [274].

11.5.2 Resolution in minimum bias events

Figure 11.10 shows the Monte Carlo simulation distributions of $\Sigma E_{\rm T}$ and $E_x^{\rm miss}$ expected in CMS from a high statistics sample of fully reconstructed minimum bias events. In order to compare the CMS result to what may be expected from previous experiments, it is essential to separate out the stochastic part from the electronic noise, the latter of which can greatly effect the observed $\Sigma E_{\rm T}$. For a given set of calorimeter thresholds, there corresponds an offset, i.e., a minimum value of $\Sigma E_{\rm T}$ that will be recorded, which may be seen to be about 150 GeV in Fig. 11.10. The $E_{\rm T}^{\rm miss}$ resolution shown in Fig. 11.10 is 6.1 GeV in excellent agreement with expectations based on a stochastic contribution of 4.8 GeV from calorimeter resolution and 3.8 GeV from electronic noise. It is important to note that at this stage, we are forming the vector $\mathbf{E}_{\rm T}^{\rm miss} = E_x\hat{\mathbf{i}} + E_y\hat{\mathbf{j}}$ using ECAL cells calibrated for photons and HCAL cells calibration for hadrons, appropriate perhaps for understanding the detector response to first collisions. It is believed that making use of energy flow techniques such as the charged track corrections described in Section 11.8 will ultimately improve the $\mathbf{E}_{\rm T}^{\rm miss}$ resolution.

11.5.3 Missing transverse energy resolution in QCD events

While the minimum bias events provide a good check for understanding the calorimeters, a major background to any potential signal containing missing transverse energy will come from QCD. The observed $\mathbf{E}_{\rm T}^{\rm miss}$ resolution is degraded in the presence of event pile-up which results in an increase of overall activity observed in the calorimeters. Even at low luminosity ($\mathcal{L} = 2 \times 10^{33}{\rm cm}^{-2}{\rm s}^{-1}$), there will be an average of 3.5 fully inelastic pile-up events per 25 ns beam crossing (5.0 events including diffractive processes). To study the detector $\mathbf{E}_{\rm T}^{\rm miss}$ resolution in events with hard collisions, 3×10^6 QCD events were generated and fully reconstructed. The events correspond to parton transverse momenta ($\hat{p}_{\rm T}$) ranging from 0–4000 GeV/c. The QCD events were all generated with low luminosity pile-up.

The QCD events with the softest collisions, $0 < \hat{p}_{\rm T} < 15$ GeV, were used to make a connection with the resolution studies performed with the minimum bias sample. Figure 11.11 shows the observed $\Sigma E_{\rm T}$ and $E_x^{\rm miss}$ distributions in these soft QCD events. The observed missing transverse energy resolution of 9.9 GeV is comparable to that from minimum bias events with the addition of pile-up.

Figure 11.12 shows the reconstructed $E_{\rm T}^{\rm miss}$ resolution vs. observed $\Sigma E_{\rm T}$ for both minimum bias events (open circles) and soft QCD events having $0 < \hat{p}_{\rm T} < 15$ GeV/c (squares). Low-luminosity pile-up is included in both cases. In the region of overlap between the minimum bias and soft QCD samples (near $\Sigma E_{\rm T}$ = 250 GeV), the reconstructed transverse energy balance is in excellent agreement.

The reconstructed $\Sigma E_{\rm T}$ for QCD events is shown in Fig. 11.13. The value of $\Sigma E_{\rm T}$ is seen to range from about 500 GeV at $\hat{p}_{\rm T} \approx 65$ GeV/c to 1 TeV at $\hat{p}_{\rm T} \approx 340$ GeV/c to 1.7 TeV at $\hat{p}_{\rm T} \approx 700$ GeV/c. It is well known that $\mathbf{E}_{\rm T}^{\rm miss}$ resolution degrades in very active events compared to the ideal case of minimum bias events; for example, the average reconstructed missing transverse energy observed in UA1 jet events was $\langle E_{\rm T}^{\rm miss}\rangle = 0.7\sqrt{\Sigma E_{\rm T}}$ GeV$^{1/2}$ [268]

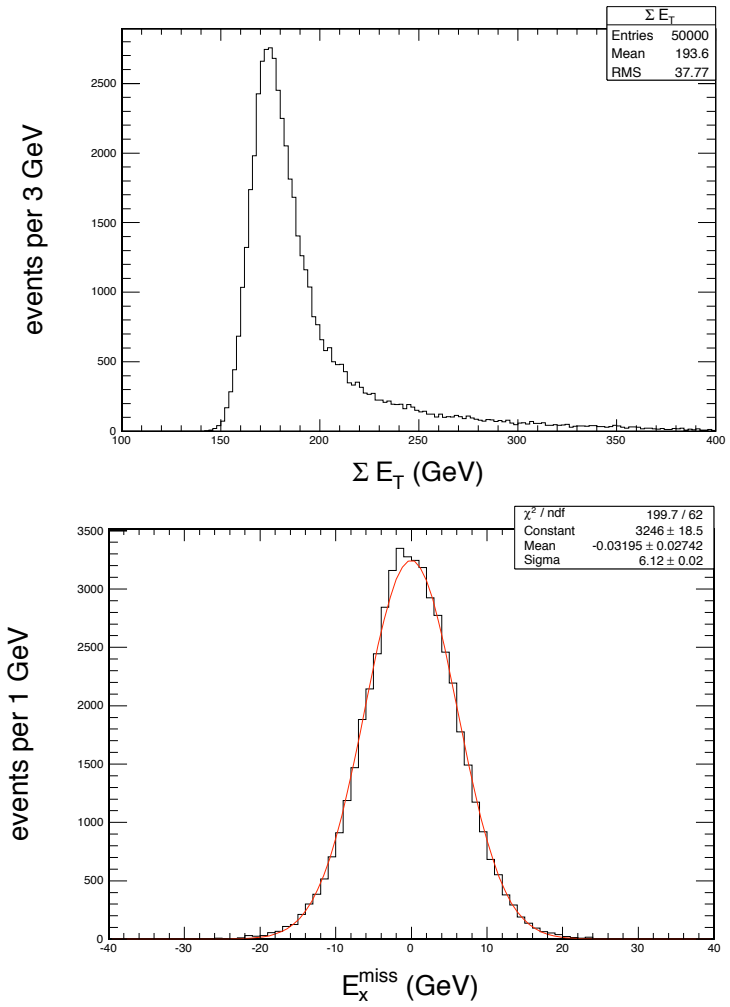


Figure 11.10: Top) Distribution of ΣE_{T} and Bottom) E_x^{miss} for minimum bias events with no pile-up. The resolution is 6.1 GeV in agreement with expectations based on a stochastic term of $0.65\sqrt{\Sigma E_{\mathrm{T}}}$ GeV$^{1/2}$ and a noise contribution of 3.8 GeV.

about 25% more than observed in minimum bias when scaled with ΣE_{T}. In hard-scatter QCD events, the distributions of E_x are also no longer perfectly Gaussian. In this case, the standard deviation of the E_x^{miss} distribution is used as the measure of resolution. Figure 11.14 shows the resolution for QCD events as a function of reconstructed ΣE_{T}. Note the agreement at ΣE_{T} = 500 GeV with Fig. 11.12 which gave $\sigma_x \approx 12$ GeV. The resolution at larger values of ΣE_{T}, however, follows a steeper path which approximately scales from the UA1 result. A fit to the resolution gives $\sigma^2 = (3.8\ \mathrm{GeV})^2 + (0.97\ \mathrm{GeV}^{1/2})^2 \Sigma E_{\mathrm{T}} + (0.012 \Sigma E_{\mathrm{T}})^2$.

The observed $E_{\mathrm{T}}^{\mathrm{miss}}$ balance is directly related to the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution. The observed $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution, of course, by its construction has a one-sided tail. A small $\hat{p}_{\mathrm{T}}$ interval can create a wide spectrum of $E_{\mathrm{T}}^{\mathrm{miss}}$ at both generator and detector level. The reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ for different intervals of parton-level generator $\hat{p}_{\mathrm{T}}$ are shown in Fig. 11.15. The reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ in QCD events (with pile-up) is shown in Fig. 11.16. At ΣE_{T} = 1700 GeV, which corresponds to $p_{\mathrm{T}} \approx 700$ GeV/c jets, an average $E_{\mathrm{T}}^{\mathrm{miss}}$ of about 50 GeV is observed. This number is consistent with measurement of such jets with a resolution of $1.25\ \mathrm{GeV}^{1/2}\sqrt{E_{\mathrm{T}}}$. A fit to the reconstructed missing transverse energy gives $(E_{\mathrm{T}}^{\mathrm{miss}})^2 = (5.4\ \mathrm{GeV})^2 + (1.23\ \mathrm{GeV}^{1/2})^2 \Sigma E_{\mathrm{T}} +$

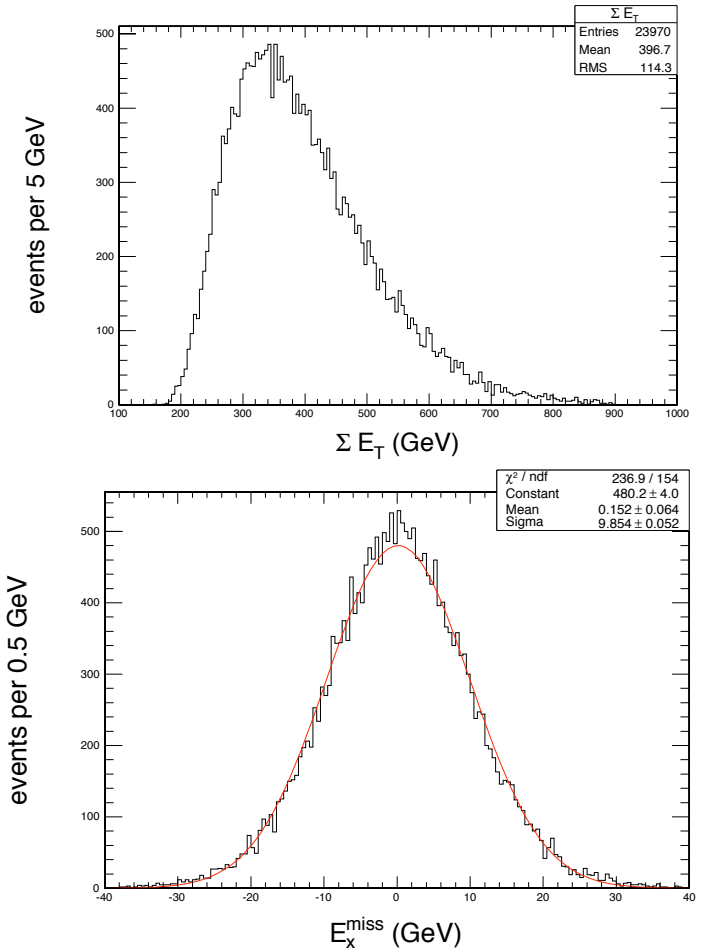


Figure 11.11: Top) Distribution of $\Sigma E_{\rm T}$ and Bottom) $E_x^{\rm miss}$ for soft QCD events ($0 < \hat{p}_{\rm T} < 15{\rm GeV}$) with pile-up. The resolution is 9.9 GeV in agreement with expectations based the study with minimum bias events.

$(0.019\Sigma E_{\rm T})^2$.

As an additional check of detector $\mathbf{E}_{\rm T}^{\rm miss}$ performance, one may look at the resolution in a direction orthogonal to the jet axis. By choosing this direction, the observed resolution is independent of the reconstructed jet resolution, but rather is dominated by the underlying event and pile-up activity. This distribution is shown in Fig. 11.17 vs. $\Sigma E_{\rm T}$. Once again, as expected, the resolution is comparable to that observed in soft collisions (Fig. 11.11). It is believed that the jet calibration described in the following section will serve as the basis for more sophisticated $\mathbf{E}_{\rm T}^{\rm miss}$ corrections and that ultimately energy flow techniques that account for calorimeter nonlinearities, magnetic field, and good charged particle resolution from the tracker can improve $\mathbf{E}_{\rm T}^{\rm miss}$ resolution.

11.5.4 Corrections to $\mathbf{E}_{\rm T}^{\rm miss}$

For QCD events, where reconstructed $\mathbf{E}_{\rm T}^{\rm miss}$ is largely an artifact of detector response, it is observed that correcting the jet energies (as described in the next section) does not significantly improve the missing transverse energy resolution except in the case where the leading jets have significantly different rapidities (and thus, much different energies). In events which

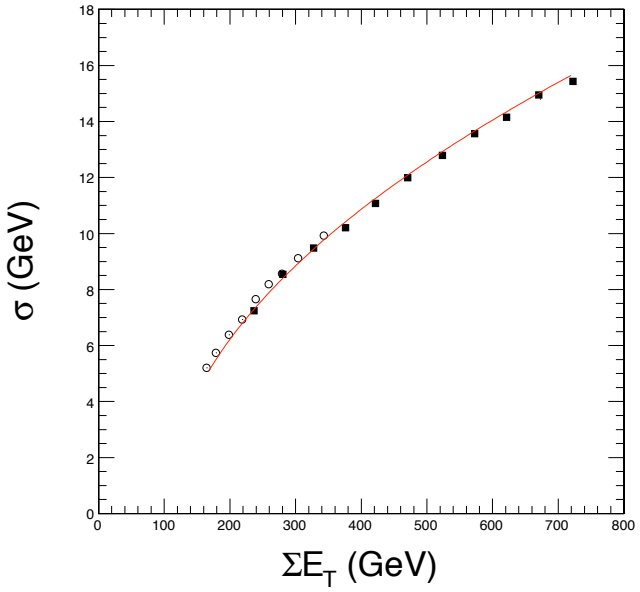


Figure 11.12: Missing transverse energy resolution vs. ΣE_{T} for QCD soft events, $0 < \hat{p}_{\mathrm{T}} < 15\ \mathrm{GeV}/c$, (squares) and minimum bias events (open circles). Low-luminosity pile-up is included in both cases.

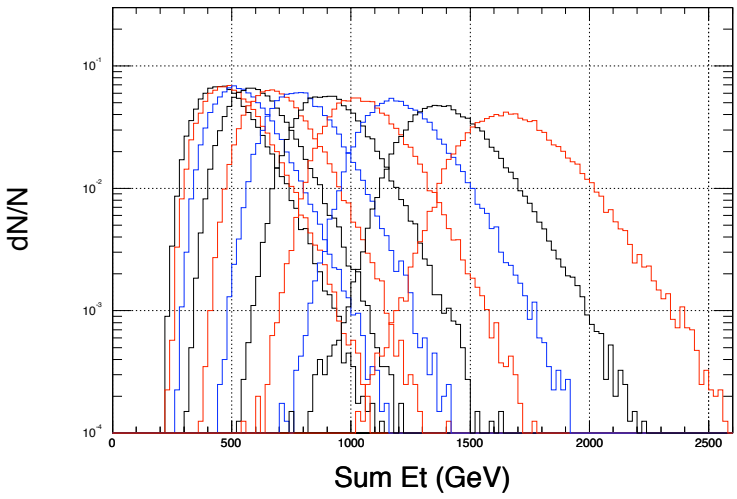


Figure 11.13: Distribution of ΣE_{T} observed in the detector for QCD events corresponding to $\hat{p}_{\mathrm{T}}$ ranges (from left to right) of 20-30, 30-50, 50-80, 80-120, 120-170, 170-230, 230-300, 300-380, 380-470, 470-600 and 600-800 GeV/c.

have a large true $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$, its reconstruction is underestimated due to nonlinearity of the calorimeter. In this case, a jet energy correction serves to help calibrate (linearize) the missing transverse energy scale [274].

A sample of $t\bar{t}$ events was used to investigate a number of corrections to measurement of $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ in events with one or more energetic neutrinos (i.e., genuine $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$). Figure 11.18 shows the resolution before and after jet corrections for inclusive $t\bar{t}$ events vs. reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$. For large values of reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$, we observe an improvement in resolution of nearly

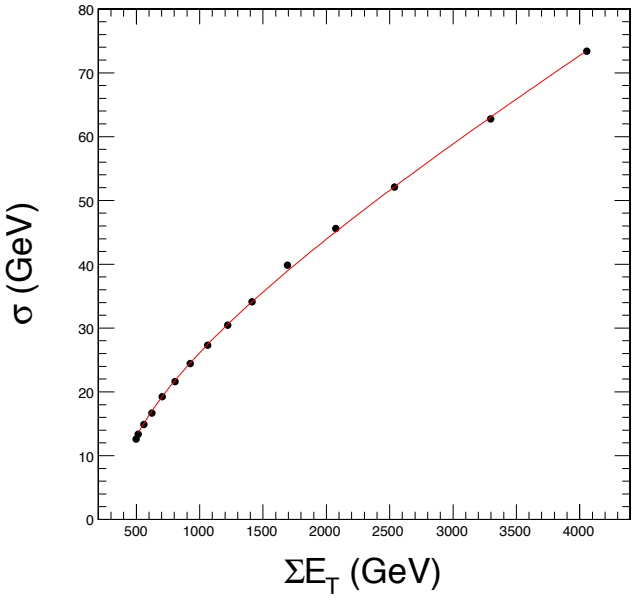


Figure 11.14: Missing transverse energy resolution vs. ΣE_{T} for QCD events with pile-up. The fit gives a stochastic contribution of $0.97\sqrt{\Sigma E_{\mathrm{T}}}$.

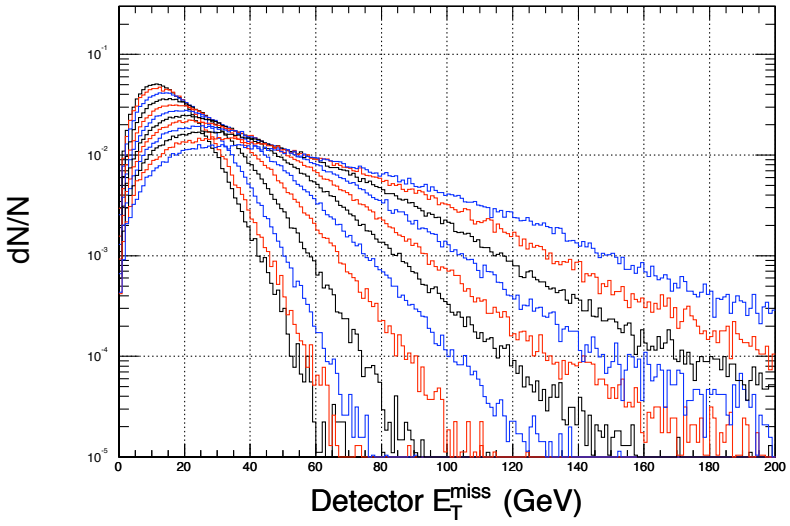


Figure 11.15: Missing transverse energy spectra in QCD samples that correspond to parton-level $\hat{p}_{\mathrm{T}}$ ranges of (from left to right) 20-30, 30-50, 50-80, 80-120, 120-170, 170-230, 230-300, 300-380, 380-470, 470-600, 600-800 and 800-1000 GeV/c.

15%.

Figure 11.18 shows the error in the reconstructed missing transverse energy scale before and after jet corrections for inclusive $t\bar{t}$ events vs. reconstructed reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$. As anticipated, the jet corrections (by design) also bring back the true missing transverse energy scale by correcting for calorimeter particle response.

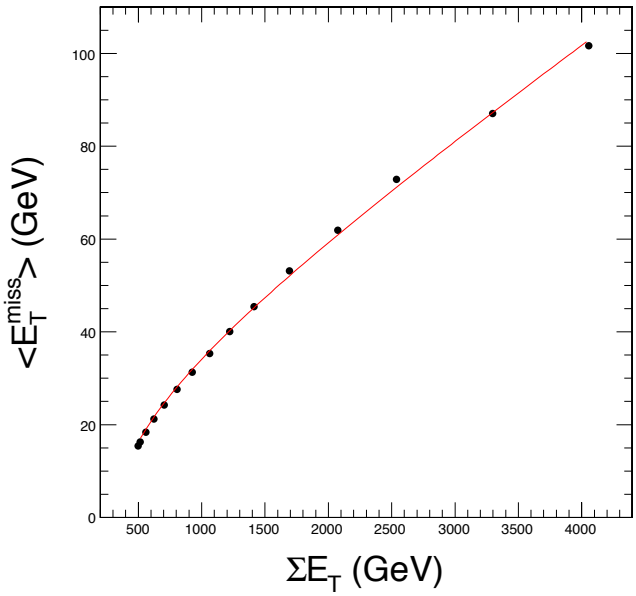


Figure 11.16: Average reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ vs. ΣE_{T} for QCD events with pile-up. The fit gives a stochastic contribution of $1.23\sqrt{\Sigma E_{\mathrm{T}}}$.

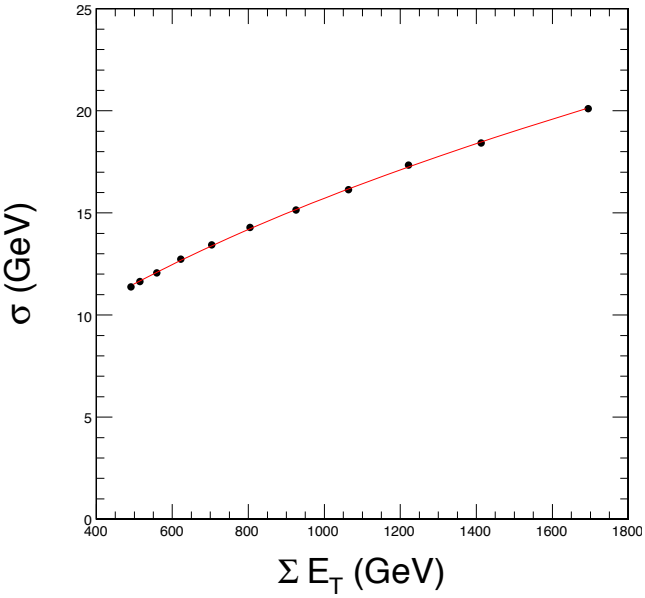


Figure 11.17: Missing transverse energy resolution in a direction orthogonal to the jet axis vs. ΣE_{T} of the entire event. The resolution is comparable to that observed in soft collisions.

11.5.5 Angular resolution

Figure 11.20 shows the angular resolution (ϕ) of the reconstructed missing transverse energy direction from inclusive $t\bar{t}$ events as a function of reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$, before and after the corrections described in Section 11.5.4. One may see that for low values of $E_{\mathrm{T}}^{\mathrm{miss}}$, the angular resolution is comparable to jet size, while for larger values of $E_{\mathrm{T}}^{\mathrm{miss}}$ the angular resolution approaches that of the calorimeter tower size.

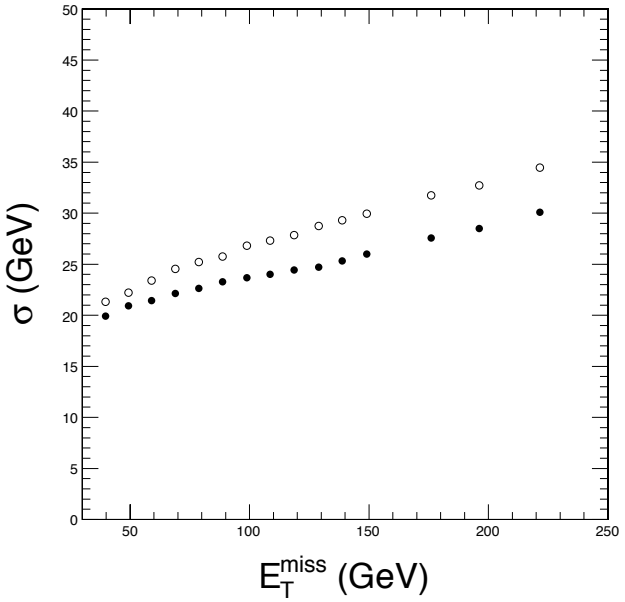


Figure 11.18: Missing transverse energy resolution before (open circles) and after jet corrections (filled circles) for inclusive $t\bar{t}$ events vs. reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$.

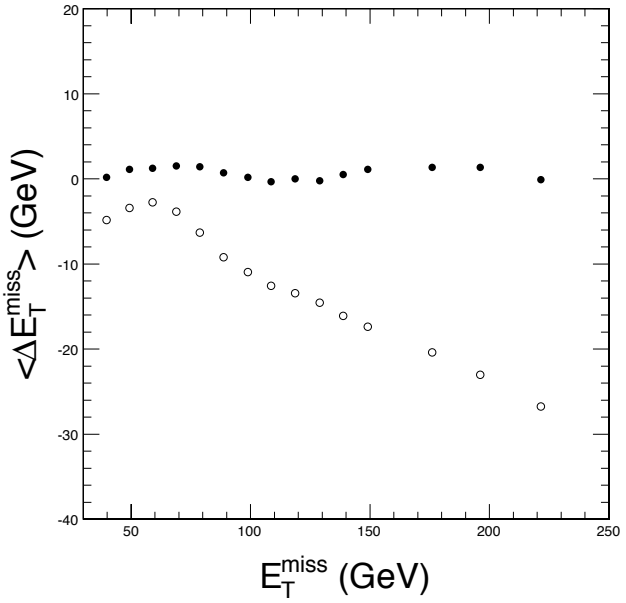


Figure 11.19: Error in the reconstructed missing transverse energy scale before (open circles) and after jet corrections (filled circles) for inclusive $t\bar{t}$ events vs. reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$.

11.6 Jet calibration

Jet calibration takes place in 2 steps: 1) a reconstructed jet is corrected to particle-level and 2) the particle-level jet is corrected to parton-level, depending on the parton type assumed in the analysis. The particle-level calibration corrects the reconstructed jet energy to equal the energy of particles in a jet from the hard scatter, independently clustered by the same algorithm and matched to the reconstructed jet. The parton-level calibration corrects the

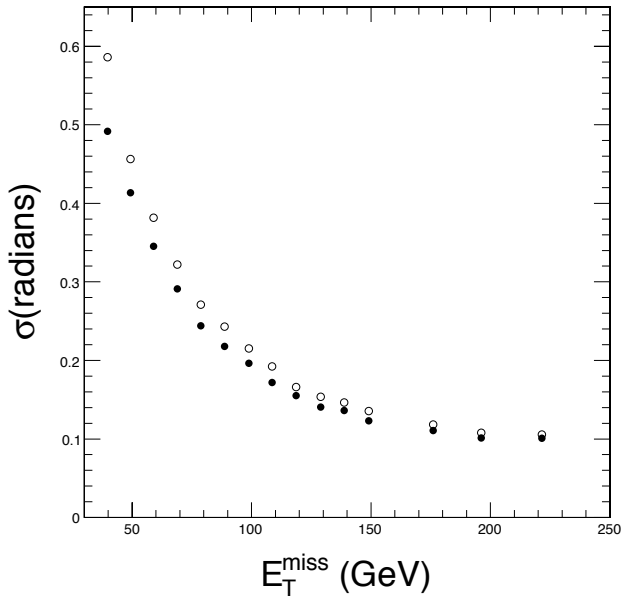


Figure 11.20: Error in the direction of the vector $\mathbf{E}_{\mathrm{T}}^{\mathrm{miss}}$ as a function of reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ in inclusive $t\bar{t}$ events before (open circles) and after corrections (filled circles).

energy of a particle-level corrected jet to the energy of the parton that originated the jet. Both steps may be combined into a single correction from the reconstructed jet to the parton-level jet (Section 11.6.3).

The particle-level calibration has 2 components: offset and response. The offset component results from multiple interactions in the event's bunch crossing, pile-up from interactions from neighboring bunch crossings, the underlying event, and any residual electronic noise after calorimeter thresholds are applied. The particle-level calibration subtracts the average offset from the reconstructed jet. The response component results from nonlinear response of the calorimeter to hadrons, differences in response among the calorimeter regions in η, lower response of cracks between calorimeters, and from the different particles contributing to the independently clustered particle-level and reconstructed jets due to magnetic field and shower spreading effects. The particle-level calibration corrects a reconstructed jet at a given p_{T} and η to the average equivalent particle-level jet.

The parton-level calibration accounts for the particles from the originating parton that are not included in the particle-level jet. This correction depends on the hadronization model used and the type of originating parton.

Jet calibrations needs to be derived for each jet algorithm and set of algorithm parameters and also depend on the p_{T} spectrum of the process. The jet calibration provides an average correction for jets reconstructed under a particular configuration and for a particular process.

11.6.1 Data-driven calibration strategy

The MC calibration technique described in Section 11.4 together with the radioactive source calibration and test beam measurements will provide a starting point for understanding the initial calorimeter calibration. The data, however, will provide a number of invaluable tools which may be used to facilitate the overall calibration procedure [276, 277]. The first stage

of the calibration will check the radioactive source calibration at the tower level, while the second stage will check the calibration of jets.

The following procedures have been identified for verifying the calorimeter tower calibration:

- Measure noise with beam-crossing triggers to check and adjust thresholds.
- Take data without zero-suppression to study the electronic noise offset.
- Check and adjust phi symmetry with minimum bias triggers.
- Use isolated muons from W decays to compare the tower-to-tower response to radioactive source source measurements and test beam muons.
- Compare isolated high p_{T} charged tracks with test beam data.

The following procedures will be used to check the calibration of jets.

- Measure the effect of pile-up on clustering algorithms and thresholds.
- Use p_{T} balance in QCD dijet events as described in Section 11.6.2 to calibrate the jet energy scale vs. η and verify the resolution.
- Use p_{T} balance in γ+jet events as described in Section 11.6.3 to calibrate the absolute energy scale.
- Use W mass fitting in tagged $\mathrm{t\bar{t}}$ events as described in Section 11.6.5 to check and fine tune the jet energy scale.

We expect that a 5% overall uncertainty in the jet energy scale can be achieved before including the W mass fit. The W mass constraint discussed in Section 11.6.5 could reduce this uncertainty to 3% for low p_{T} jets.

11.6.2 Dijet balancing

Transverse momentum balance in QCD dijet events is a proven technique to measure relative jet response and resolution from data. The results can be used to calibrate and test the full CMS simulation.

Events are selected having 1 of the 2 leading jets in the region $|\eta| < 1$ (the "barrel" jet). The other leading jet (the "probe" jet) may be at any value of η. The dijet p_{T} is defined by

$$p_{\mathrm{T}}^{\mathrm{dijet}} = (p_{\mathrm{T}}^{\mathrm{probe}} + p_{\mathrm{T}}^{\mathrm{barrel}})/2,$$

We divide the data into bins of probe jet η and measure the dijet balance (b) defined by

$$b = (p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{barrel}})/p_{\mathrm{T}}^{\mathrm{dijet}}$$

The dijet relative response, defined as the fractional difference between the jet response in the probe region and the jet response for $|\eta| < 1$, is then $2\bar{b}/(2-\bar{b})$, where $\bar{b}$ is the mean value of the dijet balance distribution. The dijet relative response as a function of probe jet η is shown in Figure 11.21. A significant η dependence is expected due to tower geometry and other instrumental effects, however, the distribution flattens when MC corrections are applied.

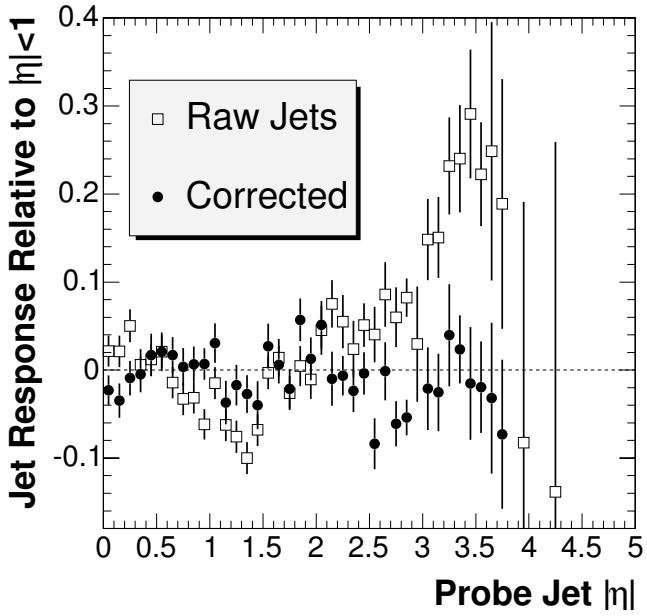


Figure 11.21: Jet response as a function of $|\eta|$ for 120 GeV/c <dijet p_{T} < 250 GeV/c. The graphs compares raw jets (open boxes) with corrected jets (solid circles). The indicated level of precision may be obtained on 1 hour of data taking.

Dijet balancing can be used to derive calibrations as a function of η based solely on the data. The errors shown on the response in Figure 11.21 correspond to a QCD sample of approximately 10^4 events. If an efficient trigger can be deployed for a p_{T} threshold of 120 GeV/c prescaled to an HLT rate of 2.5 Hz, calibration measurements with the precision shown in Figure 11.21 may be made from 1 hour of data taking. One day of data taking would be enough to calibrate the relative response of the detector to jets with a statistical error of 0.5% in the barrel and 2% in the endcap. These data could be further used to monitor the stability of jet response versus η, and provide daily calibrations to HLT triggers that require stable and uniform jet response.

The jet resolution can be measured from the RMS (σ_B) of the dijet balance. Each of the 2 leading jets contributes to the RMS, so the single jet resolution is given by $\sigma_B/\sqrt{2}$. Hard QCD radiation, generally manifested as extra jets in the event, broadens the resolution and creates non-Gaussian tails. We reduce the effects of QCD radiation to negligible levels by selecting events in which there are not any additional jets with $p_{\mathrm{T}} > 0.1 p_{\mathrm{T}}^{\mathrm{dijet}}$. To measure the jet resolution in the barrel we require that both of the leading jets have $|\eta| < 1$. Figure 11.22 shows the barrel jet resolution measured with dijet balancing as a function of dijet p_{T}, which is comparable for both reconstructed jets and corrected jets. The statistics on the measured resolution are what can be expected for 10^5 events in the indicated dijet p_{T} range.

11.6.3 γ+jet events

Apart from higher-order initial-state effects, the direct photon produced from Compton ($qg \rightarrow q + \gamma$) and annihilation ($q\bar{q} \rightarrow g + \gamma$) processes has a transverse momentum that is balanced by the jet. The high resolution (≈1%) of the electromagnetic calorimeter provides an accurate measurement of the photons and is the basis of the jet calibration procedure. The primary

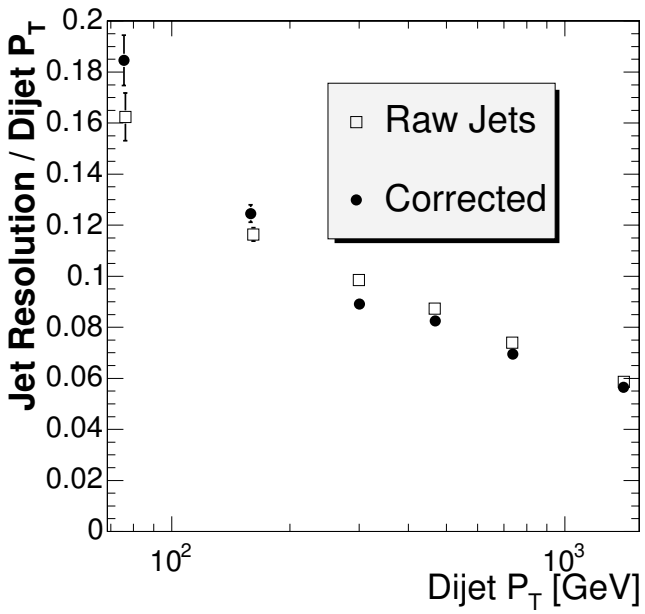


Figure 11.22: Jet resolution for $|\eta| < 1$ as a function of dijet p_{T}.

complications of this calibration procedure come from initial-state radiative corrections and a background of QCD dijet events where one jet is misidentified as a photon in the calorimeter [278].

The measured observable $k_{\mathrm{jet}} \equiv p_{\mathrm{T}}^{jet}/p_{\mathrm{T}}^{\gamma}$ provides an approximate value for the true parton-level calibration of the jet given by $k_{\mathrm{jet}}^{\mathrm{true}} \equiv p_{\mathrm{T}}^{jet}/p_{\mathrm{T}}^{\mathrm{parton}}$. The calibration constant given by $k_{\mathrm{jet}}^{\mathrm{true}}$ is the inverse of the correction factor needed to convert the measured transverse momentum of the jet to the transverse momentum of an initial parton. In the presence of initial state radiation, the transverse momentum balance of the photon and the parton is broken, leading to a two-dimensional distribution in $p_{\mathrm{T}}^{\mathrm{parton}}$ and p_{T}^{γ} (Fig.11.23a). The correlation is symmetric along the line $p_{\mathrm{T}}^{\gamma} = p_{\mathrm{T}}^{\mathrm{parton}}$. Thus the $p_{\mathrm{T}}^{\gamma} - P_{\mathrm{T}}^{\mathrm{parton}}$ balance is preserved by statistically averaging over events with a fixed sum in the transverse momentum of the photon and the parton.

Calibration coefficients are determined directly in bins of p_{T}^{γ}, however, the p_{T} balance of the γ+parton system is broken in this case. From Figure11.23b, projecting a slice of the $p_{\mathrm{T}}^{\mathrm{parton}}$ distribution for p_{T}^{γ} =constant shows a strongly asymmetric distribution with $\langle p_{\mathrm{T}}^{\mathrm{parton}} \rangle < p_{\mathrm{T}}^{\gamma}$. For the measurement of the transverse momentum of the parton, the value k_{jet} will contain an error from initial state radiation corresponding to $\Delta = k_{\mathrm{jet}} - k_{\mathrm{jet}}^{\mathrm{true}} = p_{\mathrm{T}}^{\mathrm{parton}}/p_{\mathrm{T}}^{\gamma} - 1$. This error is significant (6.3% for $p_{\mathrm{T}}^{\gamma} = 100$ GeV), but may be essentially eliminated by defining calibration coefficients k_{jet} to correspond to the peak of the $p_{\mathrm{T}}^{\mathrm{jet}}/p_{\mathrm{T}}^{\gamma}$ spectrum.

Fully simulated γ+jet events are generated with low luminosity ($\mathcal{L} = 2 \times 10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$) and reconstructed using the iterative cone and cluster-based k_T-algorithms. Quantum chromodynamics dijet events in which one jet is misidentified as photon provide the main background to the calibration sample.

The selection of events at the detector level was done with tight cuts on photon isolation ($E_{T\,\gamma}^{\mathrm{isol}} < 5$ GeV), the angle between the photon and the jet ($\Delta\phi_{\gamma,\mathrm{jet}} > 172^{\circ}$) and on transverse

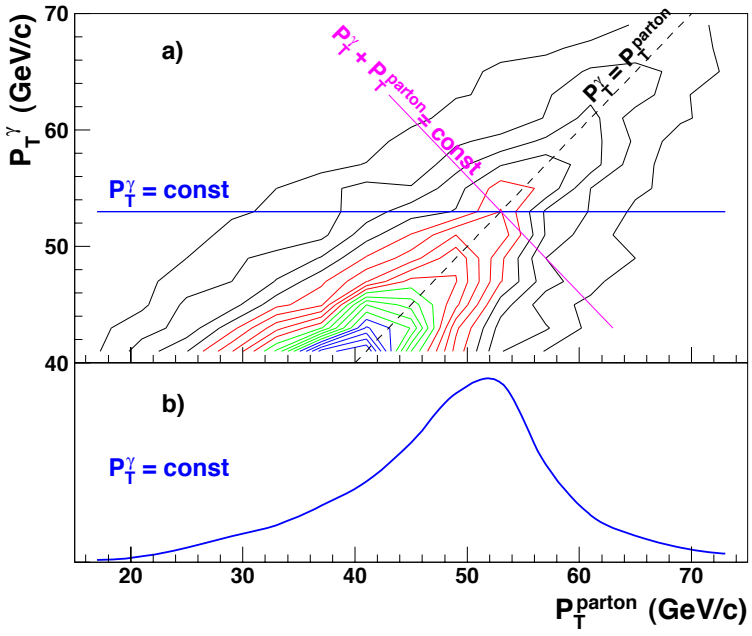


Figure 11.23: a) Distribution of the 2D correlation between the photon and parton transverse momenta, and b) the parton transverse momentum spectrum for a fixed photon transverse momentum in events with direct photons.

energy of additional jets in event ($E_T^{\rm jet2} < 20$ GeV). As a measure of the photon isolation, the value of $E_{T\gamma}^{\rm isol}$ is defined to be the scalar sum of the transverse energy in calorimeter cells within a cone of radius $R = 0.7$ in η, ϕ space with respect to the direction of the parton and outside a central array of 7×7 crystals in the electromagnetic calorimeter. The sum was computed for cells above a threshold of 0.36 and 1.8 GeV for the barrel and endcap of the electromagnetic calorimeters, respectively, and above 2 GeV for the hadron calorimeter. A cut on $E_{T\gamma}^{\rm isol}$ defined in this way gave a large background suppression while maintaining reasonably high signal efficiency(approximately 50%). Figure 11.24 shows signal to background ratio as a function of photon transverse energy. For $E_T^\gamma > 150$ GeV the background is suppressed well below the signal level, while for $E_T^\gamma < 40$ GeV, the background dominates the signal.

Figure 11.25 shows the predicted values for the calibration coefficients and their true values for quark jets and QCD jets, using the iterative cone jet algorithm ($R = 0.5$) and a threshold cut on the transverse energy of calorimeter towers, $E_{\rm T} > 0.5$ GeV. Depending on the algorithm, algorithm parameters and calorimeter cell thresholds, there is a corresponding steepness to the $E_{\rm T}$ dependence of the calibration coefficients and a spread of their values for quark and gluon jets. A strong $E_{\rm T}$ dependence and large spread in response contributes to the error on the calibration.

Figure 11.26 illustrates the sensitivity of calibration errors to the choice of jet algorithm. Quark jets are less sensitive to the jet algorithm compared to gluon jets. For QCD jets, the $k_{\rm T}$ clustering algorithm produces smaller errors than the iterative cone. It should be noted

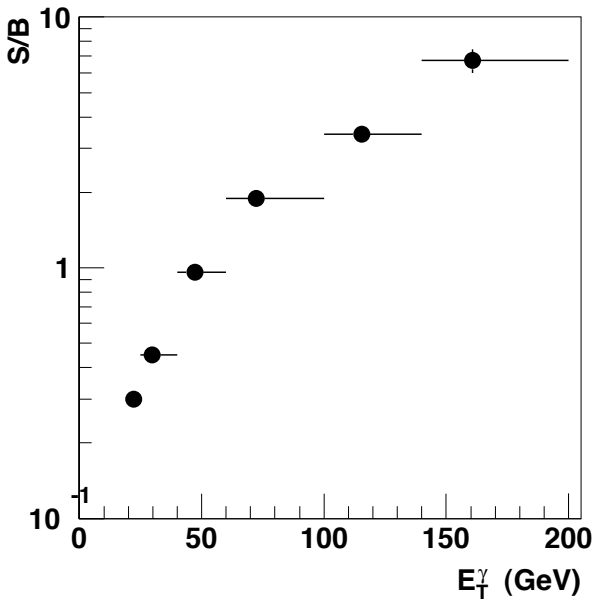


Figure 11.24: The ratio of signal to background after event selection cuts $E_{T\gamma}^{\rm isol} < 5$ GeV, $\Delta\phi_{\gamma,\rm jet} > 172°$, and $E_T^{\rm jet2} < 20$ GeV.

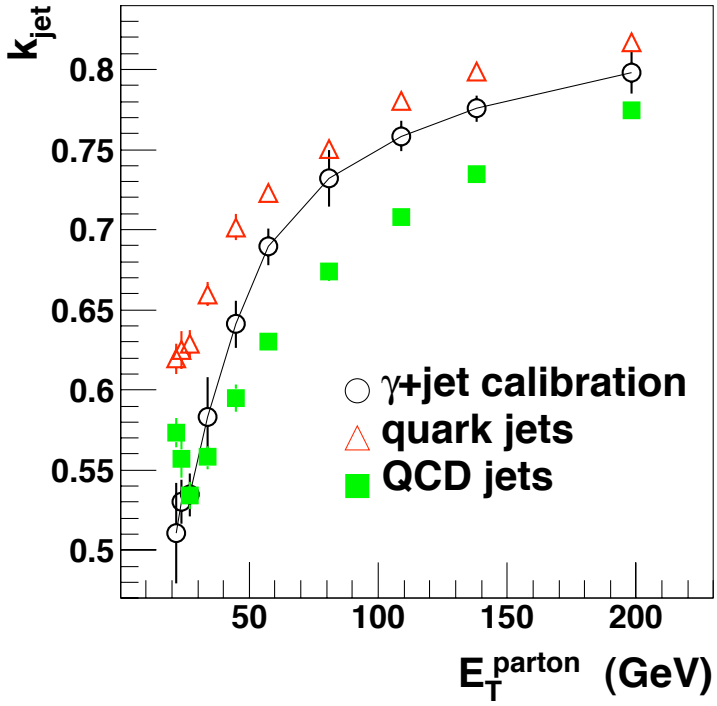


Figure 11.25: The predicted values of calibration coefficients (circles) and their true values for quark (triangles), QCD jet (squares) for the iterative cone algorithm for $|\eta_{\rm jet}| < 1.5$.

that the range of errors on the jet calibration coefficients do not characterize the quality of jet

algorithm. Comparing the effect of thresholds on the calorimeter cell readings, it is found that the lower thresholds yield the most uniform calibration coefficients.

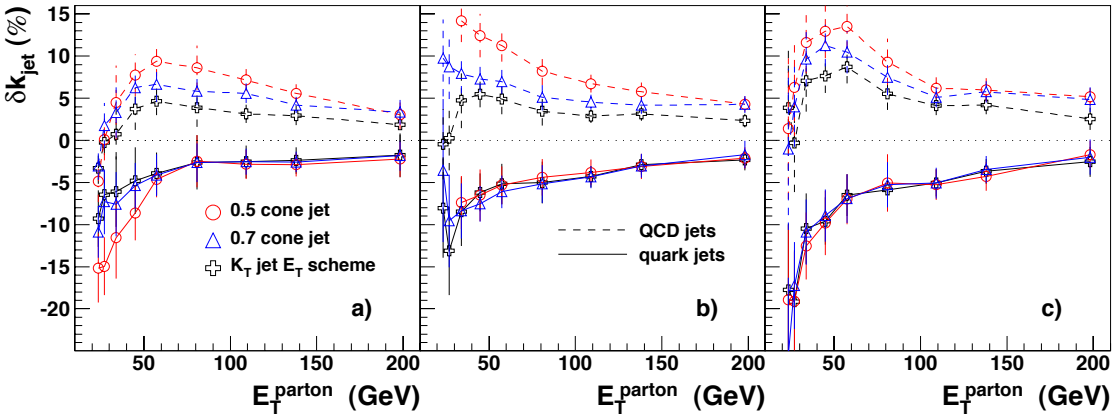


Figure 11.26: Relative systematic errors on the calibration of quark (solid lines) and QCD jets (dashed lines) for the iterative cone algorithm with cone radii of $R = 0.5$ (circles) and $R = 0.7$ (triangles) and for the k_{T}-cluster algorithm using the E_{T} scheme (crosses) for the following thresholds on calorimeter cells: a) $E_T^{\mathrm{tower}} > 0.5$ GeV, b) $E^{\mathrm{tower}} > 1$ GeV, and c) $E^{\mathrm{tower}} > 1.5$ GeV.

11.6.4 Parton-level corrections

In the CDF experiment, it has been shown that total energy of particles in various cones in the vicinity of a parton is well simulated by PYTHIA, enabling the parton energy scale to be corrected to the particle-level jet energy scale via Monte Carlo simulation derived correction factors ($k_{\mathrm{ptcl}} \equiv p_{\mathrm{T}}^{\mathrm{jet}}/p_{\mathrm{T}}^{\mathrm{parton}}$). Figure 11.27 shows these corrections as a function of quark or gluon p_{T}. For quark jets and a cone radius of $R = 0.7$, the correction to parton energy is insignificant. The η dependence of the parton corrections are shown in Figure 11.28.

The parton level corrections are different for light quark and gluon jets. The parton level correction for the mixture of quark and gluons (corresponding to QCD sample, for instance) can be obtained with expression

$$K_{q+g\ \mathrm{mixture}} = \frac{f \times K_q + K_g}{f+1}, \tag{11.2}$$

where f is the ratio of the number of quarks and gluons in the dataset. The parton level correction for the any channel can be derived from the correction for quarks and gluons, supposing the relative ratio of quarks and gluons. The determination of the relative ratio of quarks and gluons in the data is expected to be difficult and would lead to additional systematic errors on the jet energy scale. The b jet has different fragmentation compared with that from both light quarks and gluons and, thus, requires a special set of corrections. The first estimation of the parton level corrections will be obtained from γ+jet and W$\rightarrow$ jet + jet (via $t\bar{t}$) samples. In each of these samples, parton balance can be measured with reasonably high accuracy. In applying the calibration, however, one needs to take into account the differences between quark and gluon jets.

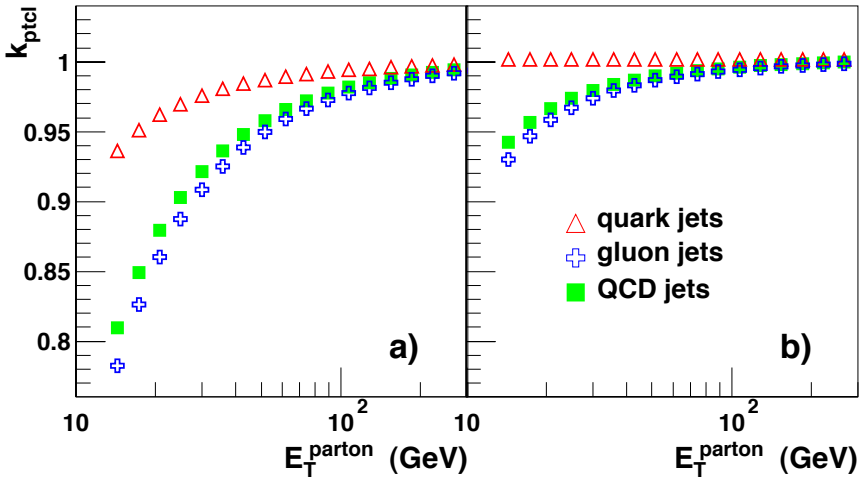


Figure 11.27: Ratio of transverse momenta of particle-level jets to the transverse momenta of the initial partons for QCD, quark and gluon jets collected in cones of (a) $R = 0.5$ and (b) $R = 0.7$ at $|\eta_{\text{jet}}| < 1.5$ as a function of transverse parton energy. These corrections are computed with PYTHIA (version 6.214).

11.6.5 Jet energy scale calibration using the W boson mass constraint in top quark events

In the search for new physics, or when aiming for precise measurements, the knowledge of the absolute energy scale of reconstructed jets originating from quarks is crucial. As demonstrated in [264] one can invert the process and determine, via the well-measured W boson mass, the absolute energy scale of reconstructed jets from the decay $W \to q\bar{q}$. An estimate is made of the precision that can be obtained on the absolute jet energy scale using the hadronic decaying W bosons in a selected sample of $t\bar{t} \to bW\bar{b}W \to bq\bar{q}\bar{b}\mu\nu_\mu$ events.

A study of this calibration technique using 3×10^6 inclusive $t\bar{t}$ events has been performed [279]. With a Next-to-Leading Order cross-section of about 800 pb, this reflects 3.75 fb^{-1} of integrated luminosity. It was found that the W+jets background contribution was negligible after the event selection. The inclusive single muon trigger is applied as described in the DAQ TDR [8].

As an example, the jets in the final state are reconstructed with the Iterative Cone algorithm using an opening angle of ΔR = 0.5. Seeds for the cones where selected from all towers above a pseudorapidity dependent energy threshold determined from the average underlying event energy deposits. An initial jet calibration has been applied as taken from a Monte Carlo study.

In order to discriminate between jets originating from the heavy b-quarks compared to the light quarks, a b-tag probability was constructed from the combined b-tag discriminant variable (Section 12.2.3). The lepton is reconstructed and identified using the methods described in [280]. A combined likelihood ratio of several observables is determined for each muon in the final state in order to enhance the purity of choosing the correct lepton from the leptonic W decay. The muon having the largest combined likelihood ratio value is taken as the hard lepton of interest.

A simple pre-selection was applied on the event requiring at least four jets with pseudora-

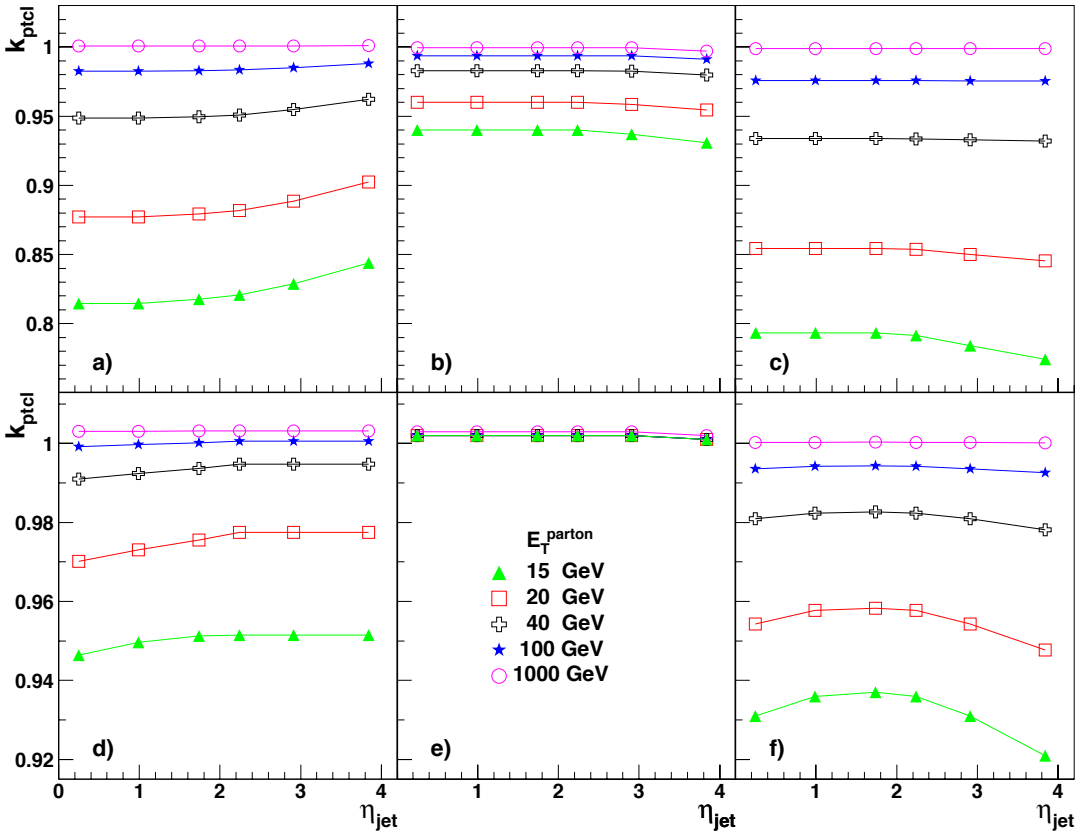


Figure 11.28: Ratio of transverse momenta of particle-level jets to the transverse momenta of the initial partons for QCD, quark and gluon jets collected in cones of (a) $R = 0.5$ and (b) $R = 0.7$ as a function of jet pseudorapidity for the different energies of parton. These corrections are computed with PYTHIA (version 6.214).

pidities in the range of the tracker or $|\eta| < 2.4$ and a raw or not calibrated E_{T} above 10 GeV. The jets must have a flight direction through the tracker to allow for a proper performance of the b-tagging algorithm. At least one muon is required within the tracker acceptance of $|\eta| < 2.4$ and with a combined likelihood ratio value large than 0.01.

The event is required to have exactly 4 jets with a calibrated p_{T} above 30 GeV/c. Exactly two of these four jets need to have a b-tag probability larger than 60%, the remaining two jets should not exceed a probability value of 30% . The latter two are assigned to the W boson decay, resulting in an efficiency of 80% for choosing the correct jet combination. It is also required that the cones of these four jets do not overlap in (η, ϕ) space. The reconstructed hard lepton has to have a transverse momentum p_{T} exceeding 20 GeV/c. The two jets assigned to the W boson are required to have a p_{T} smaller than 120 GeV/c and this two jet system together with one of the two b-tagged jets should have a top quark mass below 350 GeV/c^2. In order to be more robust the angle between these two light quark jets should exceed 1 radian. For 1fb^{-1} and in the W boson mass window between 0 and 160 GeV/c^2 this event selection results in 713 expected signal events and 152 expected $t\bar{t}$ events with different decay channels.

From the light quark jets the W boson mass can be determined. This is shown in Figure 11.29. The invariant W mass spectrum can be fitted with a simple Gaussian function $G(m_W, \sigma)$

and the mean value m_W can be taken as an estimate of the reconstructed W boson mass. It is observed that the estimated W boson mass does not agree with the world average, $M_W = 80.426 \pm 0.034$ GeV/c^2 [119].

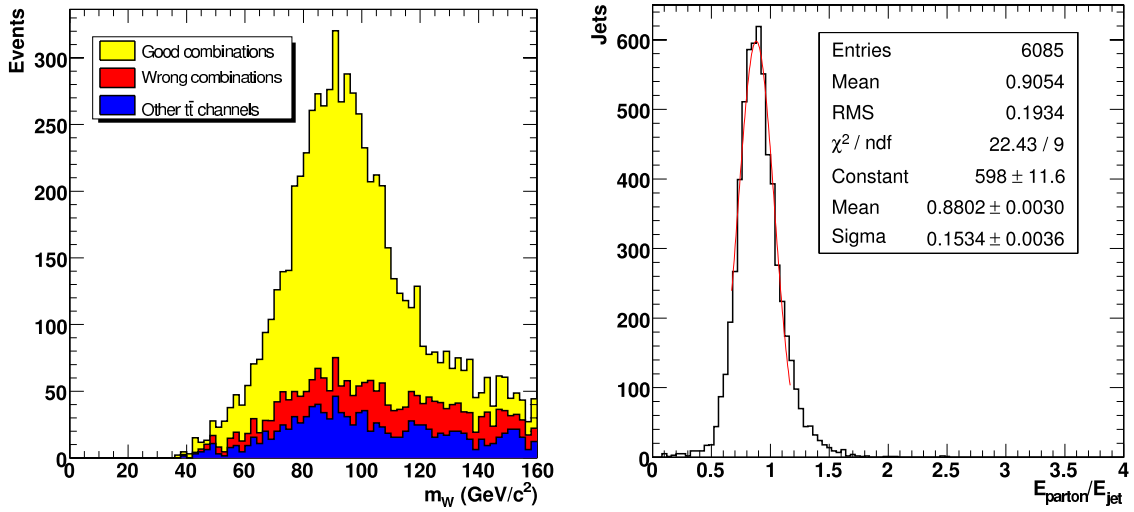


Figure 11.29: Distribution of the W boson mass including combinatorial and process backgrounds, together with the distribution of $E_{\text{parton}}/E_{\text{jet}}$ obtained from simulation truth. The background samples are rescaled to the integrated luminosity of the signal sample.

The W mass spectrum can also be constructed from jets which obtain an extra relative calibration shift ΔC (in %) on their energy scale, $E_{\text{new}} = (1 + \Delta C)E_{\text{jet}}$. The direction of the jet is kept invariant for this rescaling procedure, while the magnitude of its momentum is rescaled in order to keep the E/p ratio of the jet invariant. Therefore one can fit several m_W spectra with different values of ΔC. The value of ΔC which results in a fitted m_W in agreement with the precise world average M_W is the best estimate of the correction needed on the absolute jet energy scale for jets originating from light flavoured partons. Using this method on `MCJet` pre-calibrated jets a value of $\Delta C = -12.9 \pm 0.4\%$ is obtained. The result is obtained by rescaling the event samples to an integrated luminosity equal to the one of the signal sample.

From the selected signal events one can identify the light quark jets with an angular criterion of $\Omega_{jet,parton} < 0.2$ rad where Ω is the angle between the vectors of the jet and the parton in the (θ, ϕ) metric. From these jets the reconstructed and pre-calibrated energy E_{jet} can be compared to the generated parton energy E_{parton}, as shown in Figure 11.29. The average value of $(E_{\text{parton}}/E_{\text{jet}}) - 1$ can be taken as the true energy scale correction to be estimated, denoted as ΔC_{true}. This value is found to be ΔC_{true} = -12.0% . The bias can therefore be defined as $\Delta C_{\text{bias}} = \Delta C - \Delta C_{\text{true}}$ = -0.9 $\pm$ 0.5 % . The resulting bias could be induced by the fact that the jet energy scale correction depends on the kinematics of the $W \to q\bar{q}'$ decays used. Therefore extensions of this inclusive method should be studied, for example by estimating the measured jet energy scale correction as a function of the p_{T} of the jets.

Although the trigger requirements are not the same and there are other reconstruction aspects, it is safe to assume that the available statistics can almost be doubled when including

also the $W \to e\nu_e$ decay. Using this doubling of data, the uncertainty on the estimation of ΔC can be rescaled to an integrated luminosity of 1 fb^{-1} resulting in an expected statistical uncertainty of 0.6%. The jets in the selected sample used to obtain this result have an average transverse energy, E_{T}, equal to 57 GeV and an average energy, E, equal to 79 GeV.

The influence of pile-up is found to be ΔC(PU-noPU) $= -3.1\%$, estimated by using simulated signal event samples with and without low luminosity pile-up collisions. This is the average effect on the energy scale on each reconstructed jet from the inclusion of pile-up collisions in-time with the bunch crossing containing the hard-Q^2 event. With the advent of efficient pile-up subtraction methods, the magnitude of this effect should decrease. The influence of the combinatorial and process background is estimated from a W boson mass spectrum with and without including the background. For both the difference in ΔC is found to be on the level of 0.1% . Systematics arising from the applied b-tagging algorithm will basically change this combinatorial and process background contribution and the effect is therefore partially included in the systematic uncertainties described. The total systematic uncertainty on ΔC is conservatively estimated to be equal to 3.1% and completely dominated by the effect of pile-up.

A similar method using both the W boson mass and the top quark mass constraints can be exploited to estimate the absolute jet energy scale of jets originating from b-quarks. On an event-by-event basis one can force the 2 light quark jets from the W decay to obtain the world average measured W boson mass. The energy scale of the b quark can then be adapted to fit the world average measured top quark mass.

11.7 Association of jets with the signal vertex

One of the most experimentally challenging aspects of LHC physics is the issue of in time pile-up, where minimum-bias interactions can introduce additional jets into the event. A technique to efficiently associate jets with the signal vertex is needed to reduce the large luminosity-dependent instrumental background from pile-up jets. The CMS tracker can provide this information by associating reconstructed charged tracks to primary vertices and to jets. To associate a jet with a signal vertex, defined by a lepton or other high-p_{T} tracks of interest, a significant fraction of the tracks in the jet should originate from this primary vertex.

Two approaches have been investigated to test whether jets originate from the signal vertex or not. They differ in the use of the calorimeter jet energy measurement.

- Method A:
  1) The variable α_{jet} is determined for each jet as the ratio of the sum of transverse momenta of all tracks found inside the jet cone and having the same vertex as the signal (by a z matching residual) to the transverse momentum of the jet as measured by the calorimeters:

$$\alpha_{\mathrm{jet}} = \frac{\sum p_{\mathrm{T}}\ (\text{track in cone})}{E_{\mathrm{T}}\ (\text{jet})}. \tag{11.3}$$

  2) If $\alpha_{\mathrm{jet}} > \alpha_0$, this jet is assigned to the signal vertex.
- Method B:
  1) The variable β_{jet} is determined for each jet as the ratio of the sum of transverse

momenta of all tracks found inside the jet cone and having the same vertex as the signal to the transverse momentum of all tracks inside the jet cone:

$$\beta_{\text{jet}} = \frac{\sum p_{\text{T}} \text{ (track associated with jet)}}{\sum p_{\text{T}} \text{ (track in cone)}} \tag{11.4}$$

2) If $\beta_{\text{jet}} > \beta_0$, this jet is assigned to the signal vertex.

11.8 Jet energy correction using charged tracks

It has been shown that precision charged particle tracking may be used to significantly improve and linearize the jet energy resolution [252, 253]. Energy flow techniques have also been successfully used to improve jet measurements at HERA and LEP [281, 282]). Nevertheless, such corrections can be quite complex and are expected to evolve with time as the detector response is better understood.

An initial jet correction procedure has been identified for CMS [283, 284] which corrects for tracks swept out of the jet cone by the 4 T magnetic field and replaces the expected calorimeter response for those charged tracks with the corresponding energy determined from the more accurate tracker momentum measurement. The expected calorimeter response for such charged tracks is based on test beam measurements and Monte Carlo simulations. This "calorimetry-plus-tracker" correction procedure increases the jet energy since the (underestimated) response from the calorimetry to charged hadrons is replaced with the corresponding track momenta, as measured in the tracker, as well as adding the out-of-cone energy (from charged tracks).

Samples of QCD dijet events, simulated in different bins of $\hat{p}_{\text{T}}$ were simulated with PYTHIA and digitized without simulating the underlying event and without simulating pile-up [284]. Jets are found at the generation and reconstruction levels using the iterative cone algorithm with $R = 0.5$. A comparison was then made of the reconstructed jet with the Monte Carlo generated jet. The dependence of the jet resolution on the Monte Carlo generated jet energy (for jets generated with $|\eta| < 0.3$) is shown in Figure 11.30 for "calorimetry" jets. For comparison, the dependence of the ratio of reconstructed jet energy to Monte Carlo generated jet versus Monte Carlo generated jet energy (again, for jets generated with $|\eta| < 0.3$) is shown in Figure 11.31 for "calorimetry-plus-tracker" jets. The linearity of the response is seen to improve by including track corrections to the jets. Indeed, approximately half of the resolution improvement is a result of recovering the energy from out-of-cone charged tracks.

To illustrate the expected improvement from using the track corrections, Z' bosons ($m = 120\ \text{GeV}/c^2$) decaying into light quarks were fully simulated and digitized with low-luminosity pile-up. The Z' mass was reconstructed from the 2 leading jets. Figure 11.32 shows the ratio of the reconstructed to the generated Z' mass with, and without, the track corrections. The ratio is 0.88 ± 0.12. Including charged-track corrections, the ratio improves to 1.01 ± 0.12, correcting the mass scale, and improving the resolution by about 10%.

11.9 Jet reconstruction in heavy-ion collisions

Jets from heavy-ion collisions at the LHC may provide information about parton propagation through a quark-gluon plasma. The main challenge with jet reconstruction in heavy-ion

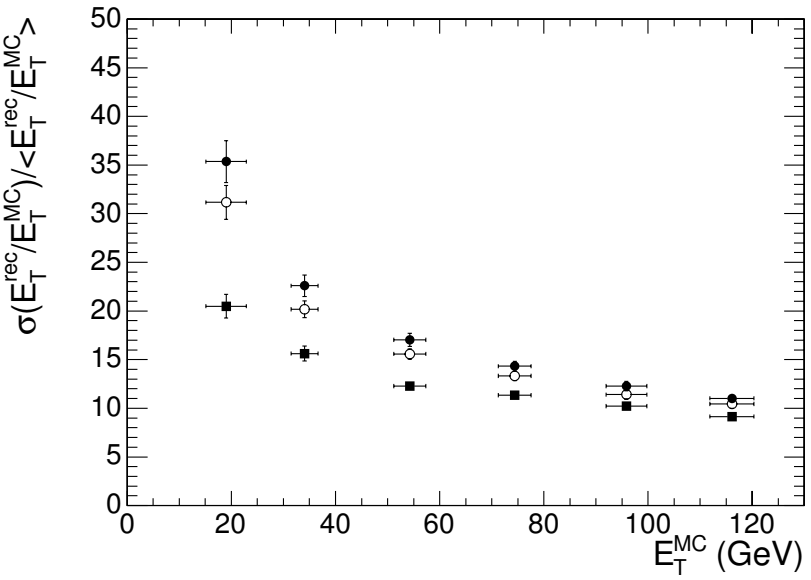


Figure 11.30: The jet E_{T} resolution as a function of generated jet E_{T} for reconstruction with calorimeter only (full circles), addition of out-of-cone tracks (open circles), and further charged track response corrections (full squares).

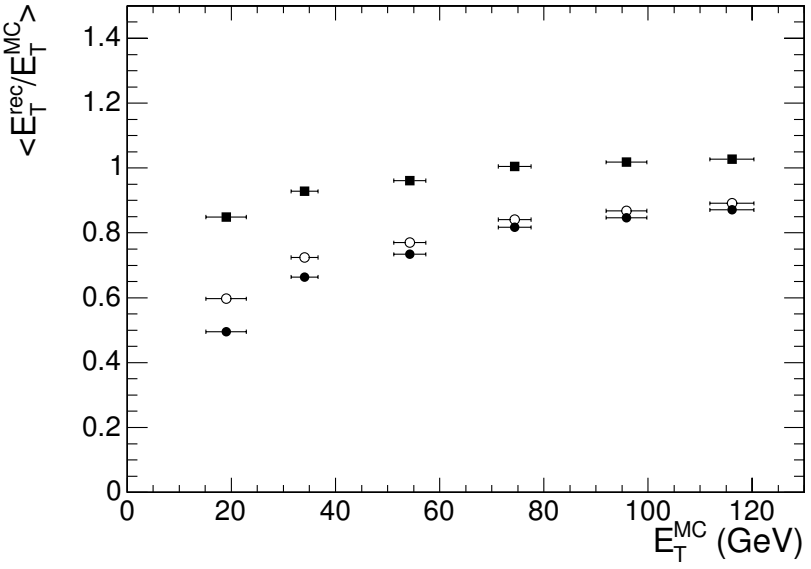


Figure 11.31: Ratio of reconstructed to generated jet E_{T} as a function of the generated jet E_{T} for reconstruction with calorimeter only (full circles), addition of out-of-cone tracks (open circles), and further charged track response corrections (full squares).

collisions is the subtraction of background arising from fluctuations of the transverse energy flow due to the large multiplicity of secondary particles in the event [285]. Predictions vary from 1400 to 8000 charged particles per unit of rapidity in central Pb-Pb collisions at the LHC.

The original jet finding algorithm was developed for reconstructing hard jets (E_{T} of order 100 GeV) in heavy-ion collisions, where particles produced in a typical collision deposit transverse energy up to 10 GeV in every calorimeter tower. The algorithm allows subtraction of

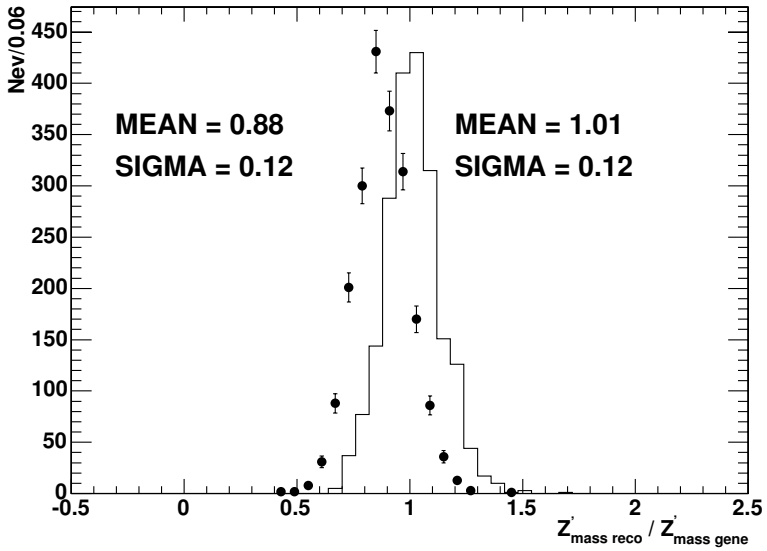


Figure 11.32: Ratio of the reconstructed to the generated Z' mass with calorimetry only (solid circles) and with track corrections (histogram).

the background energy due to the underlying event energy flow and identification of the hard jets on an event by event basis. In this study, jet finding is done with the background subtraction algorithm currently implemented in the CMS reconstruction software (ORCA).

11.9.1 Background subtraction algorithm

In heavy-ion collisions at the LHC, fully formed jets are expected to be visible for the first time in nuclear collisions with transverse energies well over 100 GeV. A jet finding algorithm has been developed to search for clusters above the average transverse energy flow.

- The average transverse energy of tower i ($\langle E_{\mathrm{T}}^{i} \rangle$) and its dispersion ($\sigma_i$) are calculated.
- All tower transverse energies are recalculated by subtracting the average tower energy plus a factor k times its dispersion ($E_{\mathrm{T}}^{i} - \langle E_{\mathrm{T}}^{i} \rangle - k\sigma_i$). If the value of the transverse tower energy after subtraction becomes negative, it is set to zero.
- Using the corrected tower energies, jets are found with the iterative cone algorithm.
- The average tower energies and dispersions are recalculated again using only towers outside of the jets. The original tower energies are used in this calculation.
- All tower energies are recalculated by subtracting the revised average tower energy plus k times its dispersion. If the value of the transverse tower energy after subtraction becomes negative, it is set to zero.
- Using the revised tower energies, jets are found with the iterative cone algorithm one more time.

A factor $k = 1$ is used in this study to compensate the positive bias in the reconstructed jet energy due to suppression of towers with the negative energy. This scheme gives an

approximation for the reconstructed jet energy in Pb-Pb close to the energy obtained in pp.

11.9.2 Performance of jet reconstruction

Dijet events from pp collisions generated with PYTHIA are superimposed on 5.5 TeV Pb-Pb events which are generated with the HIJING Monte-Carlo generator using the default setting (quenching on) with $dN_{\rm ch}/dy|_{y=0} = 5000$. The combined events were fully digitized.

Jet reconstruction is studied using the background subtraction algorithm described above. The threshold for the reconstructed jet energy is set at 30 GeV. Only the highest $E_{\rm T}$ jet in an event is used for further analysis. The correlation between the reconstructed and generated jet $E_{\rm T}$ for Pb-Pb and pp events is shown in Figure 11.33. On average, the measured jet energy in Pb-Pb collisions is the same as that in pp. In other words, the background subtraction algorithm gives a reconstructed jet energy which is the same for pp (without background) and for Pb-Pb (with background) interactions.

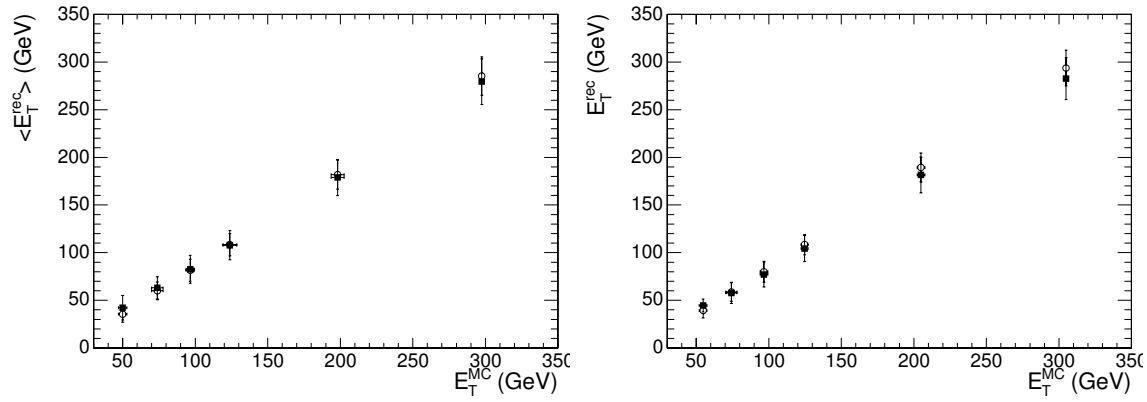


Figure 11.33: The correlation between the average reconstructed and the generated jet transverse energies in Pb-Pb (full squares) and pp (open circles) events in the barrel (left) and endcap (right). The error bars are the dispersion of jet energy distribution.

The jet energy resolution is defined as $\sigma(E_{\rm T}^{\rm reco}/E_{\rm T}^{\rm MC})/\langle E_{\rm T}^{\rm reco}/E_{\rm T}^{\rm MC}\rangle$, where $E_{\rm T}^{\rm reco}$ and $E_{\rm T}^{\rm MC}$ are the reconstructed and generated jet transverse energies. Figure 11.34 shows the jet resolution as a function of $E_{\rm T}^{\rm MC}$. The jet energy resolution is degraded by approximately 20% in high multiplicity central Pb-Pb collisions compared to pp.

Since the azimuthal angle and the rapidity distributions of jets is of particular interest for jet quenching observables in heavy-ion collisions, spatial resolution is important. For 100 GeV jets, the η and ϕ resolutions are 0.028 and 0.032 correspondingly.

The subtraction procedure allows the identification and measurement of jets in heavy-ion collisions using only CMS calorimeters with very high efficiency and purity. Jets reconstructed in central Pb-Pb collisions with $E_{\rm T} > 30$ GeV, that are within $\Delta R < 0.25$ around the direction of the generated MC jet are considered as true QCD jets. The efficiency of jet reconstruction in central Pb-Pb collisions is estimated as the fraction of events with such true QCD jets among all the generated events. The efficiency of finding a true QCD jet is nearly 100% for $E_{\rm T} = 75(100)$ GeV jets in the barrel (endcap) region.

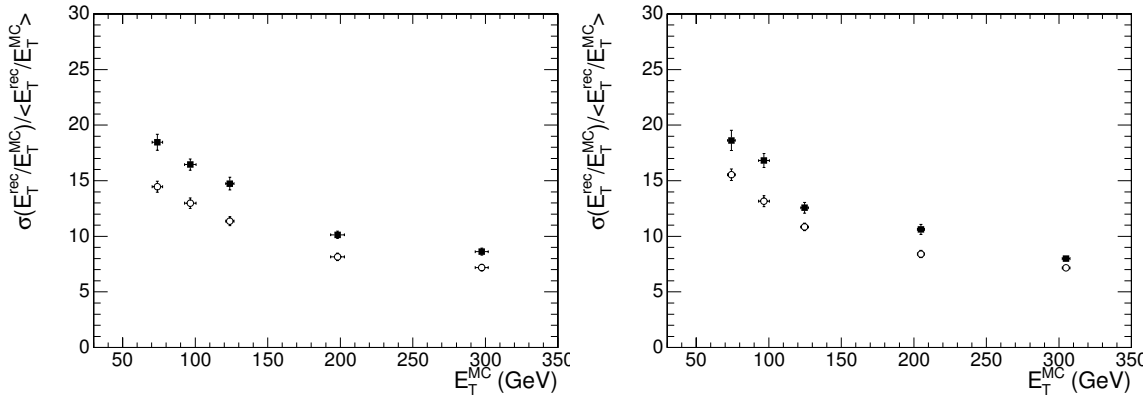


Figure 11.34: Jet energy resolution in Pb-Pb (full squares) and pp (open circles) events in the barrel (left) and endcap (right) regions.

The purity of the reconstructed jet sample is defined as the number of events with true QCD jets divided by the number of events with at least 1 reconstructed jet (fake or real) with transverse energy above 30 GeV. Beyond E_{T} = 50 GeV (75 GeV) jets, the purity is nearly 100% for the barrel (endcap).

Chapter 12

b and τ Tagging

12.1 τ tagging tools

Good τ lepton identification is very important for analyses at the LHC that involve searches for Higgs bosons or evidence of SUSY. We describe a number of methods to identify hadronic decays of τ leptons, so-called "τ jets," based on τ-lepton and τ-jet properties such as lifetime, mass, small charged particle multiplicity, and the collimation and isolation of τ-decay products. The optimal combination of these methods depends on the physics channel under consideration. In general, the primary requirement for τ-jet identification, which can be augmented by other τ-tagging methods, is the isolation of a collimated group of charged particles in the tracker. Further details about the τ-tagging methods and High-Level Trigger strategy descibed here can be found in [286].

The performance of the τ-tagging methods has been evaluated using a sample of simulated single τ jets with transverse momentum ${p_{\mathrm{T}}^{\tau\,\mathrm{jet}}} > 30\ \mathrm{GeV}/c$ distributed uniformly in ϕ and in pseudorapidity up to $|\eta| = 2.2$. This sample consists only of τ jets and does not contain any pp collision products coming from the underlying event. Low momentum τ jets (${p_{\mathrm{T}}^{\tau\,\mathrm{jet}}} < 30\ \mathrm{GeV}/c$) have not been considered. Dijet events generated by PYTHIA with $\hat{p}_{\mathrm{T}}$ between 30 and 170 GeV/c (where $\hat{p}_{\mathrm{T}}$ is the parton p_{T} in the $2 \rightarrow 2$ scattering process in PYTHIA) are used to obtain rejection factors against quark and gluon jets using the τ-tagging methods. Two leading E_{T} Monte Carlo simulated jets in the dijet sample separated in η-ϕ space by a distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ >1.5 and with $|\eta| < 2.1$ are propagated through the τ-identification criteria. The τ-tagging performance is evaluated as a function of the true transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$ and the pseudorapidity of the jet. The true τ-jet transverse energy is defined as the energy of the τ lepton decay products excluding the neutrino. The true energy of the QCD jet is the energy of the Monte Carlo jet found using the cone algorithm (described in Section 11.2) with a cone size 0.5. The cone algorithm uses all stable particles, excluding neutrinos and muons, as input.

12.1.1 Tau properties relevant to τ-jet reconstruction and identification

The τ lepton decays hadronically 65% of the time, producing a τ jet. This is a jet-like cluster in the calorimeter containing a relatively small number of charged and neutral hadrons. When the p_{T} of the τ jet is large compared to the τ mass, these hadrons have relatively small momentum in the plane transverse to the τ-jet axis. In 77% of hadronic τ decays, the τ jet consists of only one charged hadron and a number of π^0s (so-called "one-prong" decays). These features mean that hadronic τ decays produce narrow jets in the calorimeter.

Figure 12.1 shows the ratio $r = E_{\mathrm{T}}^{\mathrm{reco}}/E_{\mathrm{T}}^{\mathrm{MC}}$ as a function of the reconstruction cone size for

the different bins of $E_T^{\rm MC}$, where $E_{\rm T}^{\rm reco}$ is the transverse energy reconstructed in the calorimeter with the iterative cone algorithm. The thresholds on calorimeter towers, which are inputs to the jet finder, were set at $E_{\rm T} = 0.5$ GeV and $E = 0.8$ GeV. The values of r in Figure 12.1 are normalized to unity at a cone size of 0.6. Figure 12.2 shows the transverse energy resolution

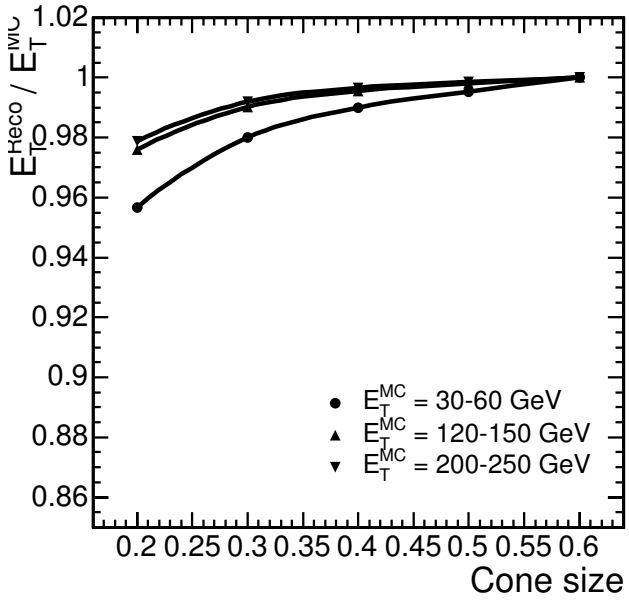


Figure 12.1: The ratio $r = E_{\rm T}^{\rm reco} / E_{\rm T}^{\rm MC}$ as a function of the reconstruction cone size for different bins of $E_{\rm T}^{\rm MC}$. The values of r were normalized to unity at a cone size of 0.6.

of τ-jets as a function of reconstruction cone size for different bins of $E_{\rm T}^{\rm MC}$. For Figure 12.1

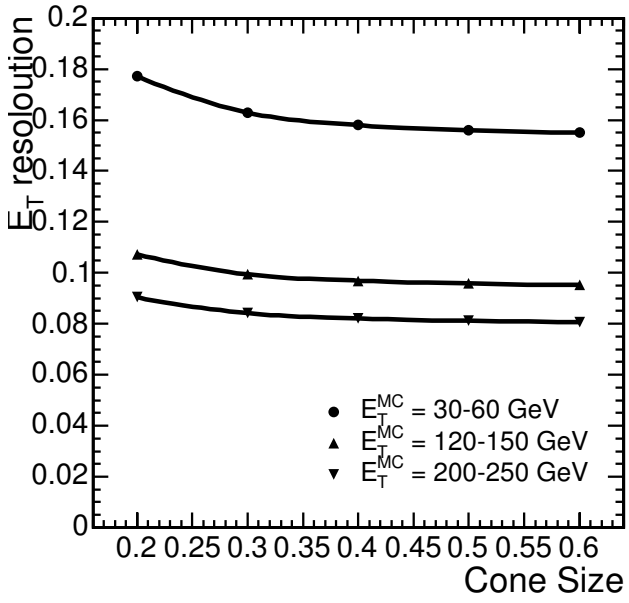


Figure 12.2: The transverse energy resolution of τ jet as a function of the reconstructed cone size for the different bins of $E_{\rm T}^{\rm MC}$.

a cone size of 0.4 for τ-jet reconstruction with the calorimeter was chosen since it contains a

large fraction of the τ-jet energy (more than 98 %). A cone size smaller than 0.4 leads to a degradation of the τ-jet energy resolution as can be seen in Figure 12.2. Figure 12.3 shows the difference between the generated and reconstructed values of the τ-jet direction components, ϕ (left plot) and η (right plot), where the reconstructed jet direction is determined using calorimeter information. These values are shown for 3 ranges of the generated τ-jet energy. The charge of the τ lepton is positive. For jets with E_{T} between 40 and 60 GeV, the

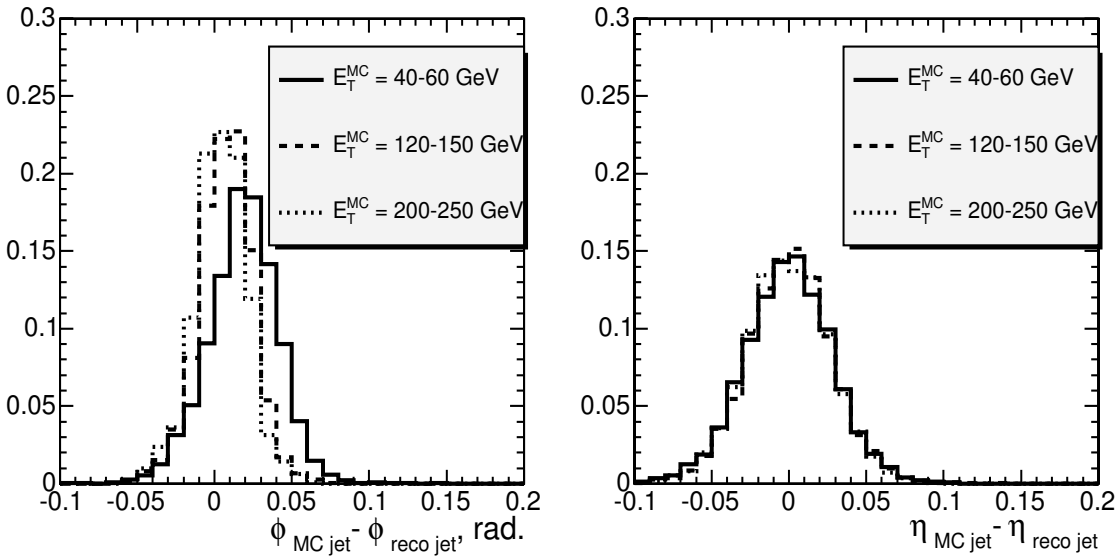


Figure 12.3: The difference between the true and reconstructed values of the τ-jet direction components, ϕ (left plot) and η (right plot), where the reconstructed jet direction is determined using calorimeter information. Three τ-jet energy ranges are shown which have been determined using Monte Carlo truth information. All τ leptons used to make these plots were positively charged.

4 T magnetic field leads to a systematic shift of ≈ 0.02 rad in the ϕ component of the reconstructed τ-jet direction. The shift is reduced for jets with higher E_{T}. The resolution in η is slightly worse than in ϕ and does not depend on E_{T} between 40 and 250 GeV.

Identification of τ-jets requires matching the jet axis determined by the calorimeter with the jet axis determined by charged particle tracking. Figure 12.4 shows the distance, ΔR, in η-ϕ space, between the direction, at the primary vertex, of the leading p_{T} track reconstructed with the silicon tracker and the direction of the τ jet reconstructed with the calorimeter for 3 bins of the true τ-jet transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$.

A cut of 3 GeV/c was applied to the p_{T} of the leading track. Both "one-prong" and "three-prong" (decays to 3 charged particles) τ decays are included in the plot. The value of ΔR does not exceed 0.1 for the range of $E_{\mathrm{T}}^{\mathrm{MC}}$ considered.

In the case of three-prong τ decays with 3 charged particles in the final state, the particles are produced within a narrow cone. Figure 12.5 shows the maximum distance ΔR in η-ϕ space between the leading p_{T} charged particle and the other 2 charged particles in three-prong τ decays, for 3 bins of the generated τ-jet transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$.

The τ-lepton's proper lifetime ($c\tau = 87.11\ \mu\mathrm{m}$) and its mass ($m_\tau = 1.78\ \mathrm{GeV}/c^2$) are used in τ-jet tagging. To utilize the lifetime, either the track impact parameter or the decay vertex (for

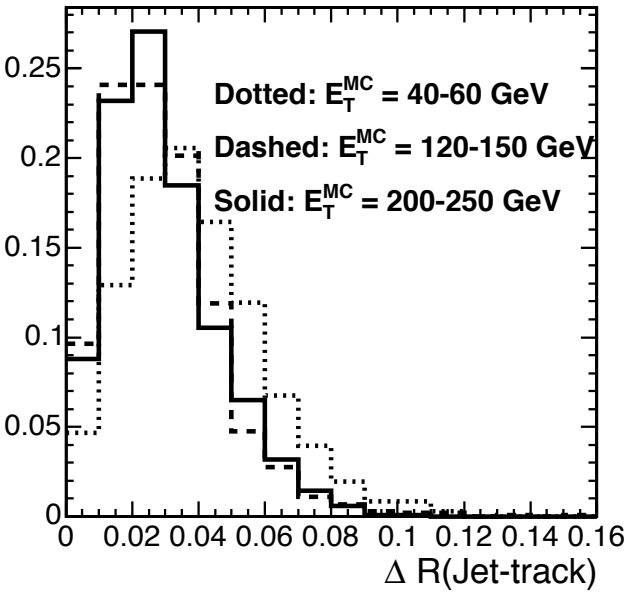


Figure 12.4: The distance ΔR in η-ϕ space between the leading p_{T} track reconstructed with the tracker and the direction of the τ jet reconstructed with the calorimeter for 3 bins of the true τ-jet transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$. A cut of 3 GeV/c was applied to the p_{T} of the leading track. Both the one- and the three-prong τ decays are included.

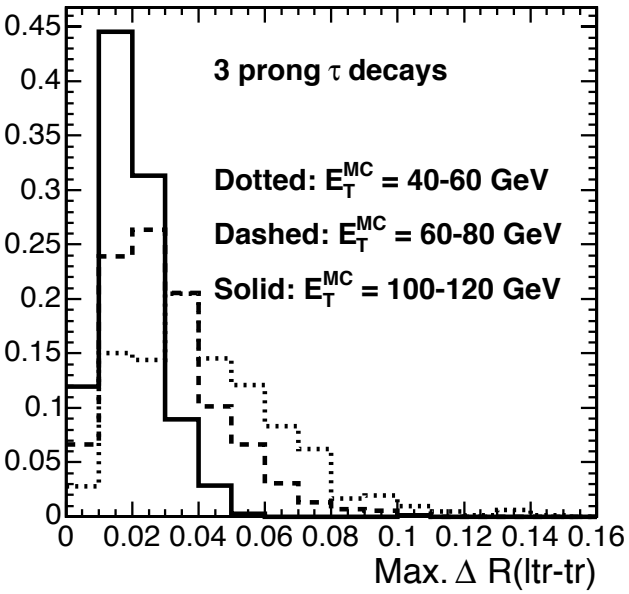


Figure 12.5: The maximum distance ΔR in η-ϕ space between the leading p_{T} charged particle and the other 2 charged particles in the three-prong τ decays for 3 bins of the true τ-jet transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$.

three-prong decays) may be used. An effective mass is formed using tracks and calorimeter clusters.

12.1.2 Methods for τ tagging and performance

12.1.2.1 ECAL isolation

Hadronic τ decays produce localized energy deposits in the electromagnetic calorimeter. Several variables have been examined to quantify this feature and to make use of it for τ tagging and QCD jet rejection [287, 288]. The electromagnetic isolation parameter $P_{\rm isol}$ defined as

$$P_{\rm isol} = \sum_{\Delta R<0.40} E_{\rm T} - \sum_{\Delta R<0.13} E_{\rm T} \tag{12.1}$$

was found to provide the best efficiency for QCD-jet rejection. The sums run over transverse energy deposits in the electromagnetic calorimeter, and ΔR is the distance in η-ϕ space from the τ-jet axis. Jets with $P_{\rm isol} < P_{\rm isol}^{\rm cut}$ are τ candidates.

Figure 12.6 shows the τ jet identification efficiency of the ECAL isolation criterion, as a function of $E_{\rm T}^{\rm MC}$ (left plot) and $|\eta^{\rm MC}|$ (right plot) for $P_{\rm isol}^{\rm cut}$ = 5 GeV/c. The efficiency is shown separately according to the different final states of the hadronic decays of the τ lepton. The dominant mode with 1 charged pion and 1 π^0 is represented by the dashed line. Only a small ($\simeq$ 5 %) variation with $E_{\rm T}^{\rm MC}$ is observed over a large region of transverse energies from 30 to 300 GeV. The efficiency variation versus pseudorapidity for τ decays with π^0's in the final state follows the variation in the tracker material budget in front of the ECAL. This is due to electrons and positrons from photon conversions in the tracker material, which contaminate the ECAL isolation region. Figure 12.7 shows the efficiency of the electromagnetic isolation

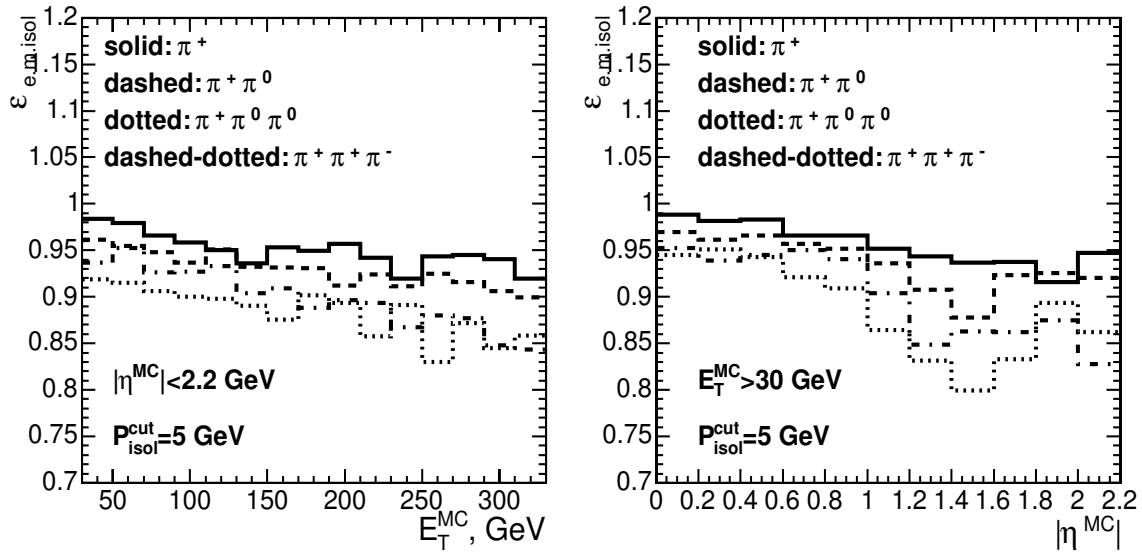


Figure 12.6: The efficiency of the ECAL isolation for τ jets as a function of $E_{\rm T}^{\rm MC}$ (left plot) and $|\eta^{\rm MC}|$ (right plot) for $P_{\rm isol}^{\rm cut}$ = 5 GeV/c. The efficiency is shown separately for different final states of hadronic decays of τ lepton.

criterion, for τ jets and QCD jets in different bins of true transverse energy when the value of $P_{\rm isol}^{\rm cut}$ is varied. It can be concluded from Figure 12.7 that the ECAL isolation requirement can provide a rejection factor of $\simeq$ 5 against high $E_{\rm T}$ QCD jets ($>$ 80 GeV), while maintaining a τ jet efficiency of greater than 80 %.

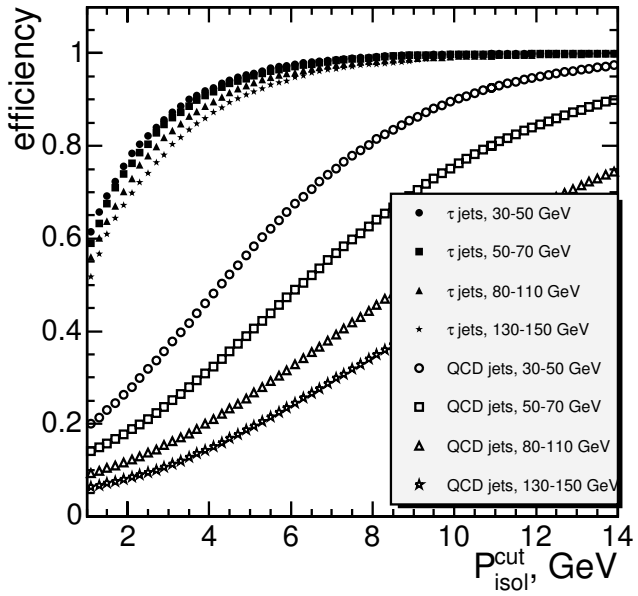


Figure 12.7: The efficiency of the electromagnetic isolation criterion, for τ jets and QCD jets in the different bins of the true transverse energy when the value of $P_{\rm isol}^{\rm cut}$ is varied.

An isolation requirement based on hadron calorimeter information results in worse performance. In addition, it introduces systematic uncertainties which are larger than those for the ECAL isolation because the energy scale is not known as precisely.

12.1.2.2 Tracker Isolation

The principle of τ-jet identification using tracker isolation is shown in Figure 12.8. The direction of the τ jet is defined by the axis of the calorimeter jet. Tracks above a threshold of $p_{\rm T}^{\rm m}$ and inside the matching cone, of $R_{\rm m}$, around the calorimeter jet direction are considered in the search for signal tracks. The leading track ($\rm tr_1$ in Figure 12.8) is defined to be the track with the highest $p_{\rm T}$. Any other tracks in the narrow signal cone $R_{\rm S}$ around $\rm tr_1$ and with a z-impact parameter $z_{\rm tr}$ close to the z-impact parameter of the leading track $z_{\rm tr}^{\rm ltr}$ ($|z_{\rm tr} - z_{\rm tr}^{\rm ltr}| < \Delta z_{\rm tr}$) are assumed to come from the τ decay. Tracks with $|z_{\rm tr} - z_{\rm tr}^{\rm ltr}| < \Delta z_{\rm tr}$ and transverse momentum above a threshold of $p_{\rm T}^{\rm i}$ are then reconstructed inside a larger cone of size $R_{\rm i}$. If no tracks are found in the $R_{\rm i}$ cone, except for the ones which are already in the $R_{\rm S}$ cone, the isolation criterion is satisfied.

Figure 12.9 shows the tracker isolation efficiency for the τ jets (left plot) and QCD jets (right plot) as a function of the isolation cone $R_{\rm i}$ for 2 values of the signal cone $R_{\rm S} = 0.07$ (full symbols) and $R_{\rm S} = 0.04$ (open symbols). In order of decreasing efficiency, the symbols correspond to $E_{\rm T}^{\rm MC}$ bins of 130–150, 80–110, 50–70 and 30–50 GeV. The remaining tracker isolation parameters are: $R_{\rm m} = 0.1$, $p_{\rm T}^{\rm i} = 1$ GeV/c, $\Delta z_{\rm tr}$=2 mm. The $p_{\rm T}$ of the leading track was required to be greater than 6 GeV/c. Tracks were reconstructed with the combinatorial track finder algorithm, requiring at least 8 hits per track and a normalized χ^2 <10. Jets were reconstructed in the calorimeter with the iterative cone algorithm using a cone size of 0.4. The reconstructed QCD jets should match the 2 leading $E_{\rm T}$ Monte Carlo simulated jets. The matching criterion requires that the distance between the reconstructed and the Monte Carlo jet axis in the η-ϕ space is less than 0.2. It can be seen from Figure 12.9 that the

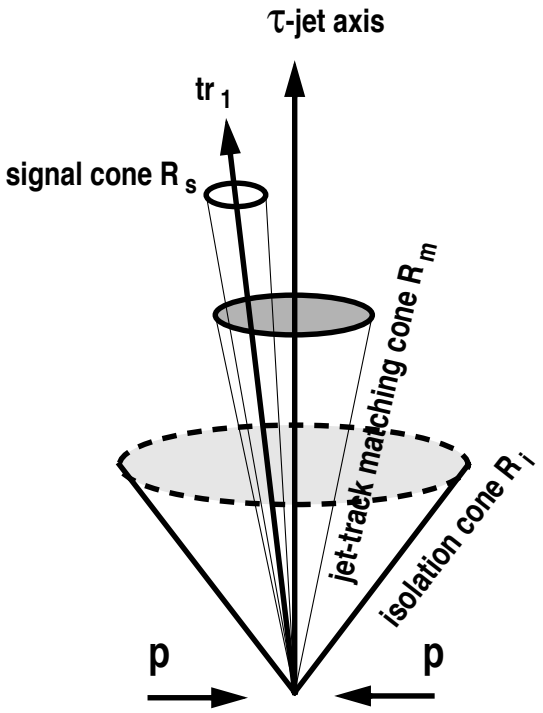


Figure 12.8: Sketch of the basic principle of τ-jet identification using the tracker isolation.

tracker isolation selection can provide a rejection factor of more than 10 for QCD jets with an efficiency of above 70% for selecting τ jets. The inefficiencies of each step of the tracker isolation algorithm for selecting τ jets, are shown in Table 12.1. Data are given for τ jets in 2 bins with $E_{\mathrm{T}}^{\mathrm{MC}}$ ranges of, 30–50 and 130–150 GeV.

Table 12.1: The inefficiencies of each step in the tracker isolation algorithm for the τ jets in 2 bins of $E_{\mathrm{T}}^{\mathrm{MC}}$, 30–50 and 130–150 GeV.

$E_{\mathrm{T}}^{\mathrm{MC}}$ (GeV)	$\geq$1 track in cone 0.4	leading track with $p_{\mathrm{T}} > 6$ GeV/c, R_{m}=0.1	isolation R_{S}=0.07	isolation R_{S}=0.04
30-50	7.7%	10.2%	5.2%	14.2%
130-150	4.8%	2.6%	1.0%	2.5%

τ jets consist mainly of one-prong or three-prong decays, with 1 or 3 charged particles in the final state. The one-prong decays make up $\simeq$ 77 % of all hadronic decay modes of the τ lepton. The tracker isolation requirement is therefore naturally followed by a requirement for there to be only 1 or 3 reconstructed tracks in the signal cone. Table 12.2 shows the efficiency of the track counting requirement for τ and QCD jets in different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$. The sizes of the isolation and the signal cones were $R_{\mathrm{i}} = 0.4$ and $R_{\mathrm{S}} = 0.07$. It can be concluded that the track counting criterion does not improve the suppression of QCD jet selection for events which pass the tracker isolation criteria.

A number of tagging methods which can be applied following the isolation criteria are discussed in the following subsections. The events are preselected using the tracker isolation criteria with parameters similar to the ones used by the High Level trigger: R_{m}=0.1, $R_{\mathrm{S}} = 0.07$, $R_{\mathrm{i}} = 0.4$, $p_{\mathrm{T}}^{\mathrm{i}}$ = 1 GeV/c and $\Delta z_{\mathrm{tr}} = 2$ mm. Either 1 or 3 tracks are required in the signal cone

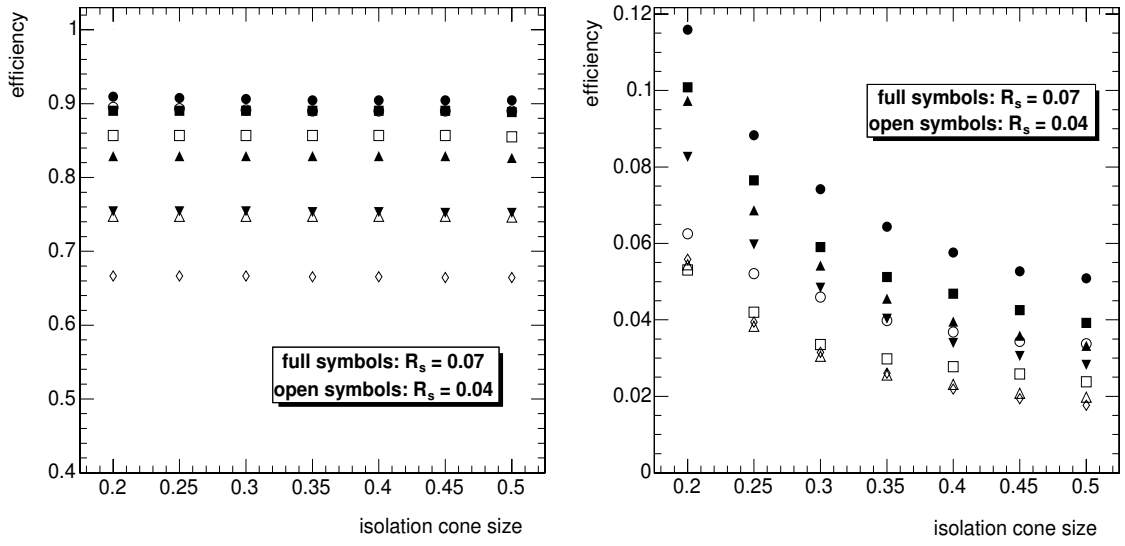


Figure 12.9: The tracker isolation efficiency for τ jets (left plot) and QCD jets (right plot) as a function of the isolation cone R_{i} for 2 values of the signal cone R_{S}=0.07 (full symbols) and R_{S}=0.04 (open symbols). In order of decreasing efficiency the symbols correspond to $E_{\mathrm{T}}^{\mathrm{MC}}$ bins of 130–150, 80–110, 50–70 and 30–50 GeV. The remaining tracker isolation parameters are: R_{m}=0.1, $p_{\mathrm{T}}^{\mathrm{i}}$=1 GeV/$c$, Δz_{tr}=2 mm and the leading track $p_{\mathrm{T}} > 6$ GeV/c.

Table 12.2: The efficiency of the track counting requirement for τ and QCD jets in different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$.

QCD jets; $E_{\mathrm{T}}^{\mathrm{MC}}$ (GeV)	30–50	50–70	80–110	130–150
1 track	63 %	72 %	69 %	60 %
3 tracks	7 %	9 %	9 %	13 %
1 or 3 tracks	70 %	81 %	78 %	73 %
τ jets; $E_{\mathrm{T}}^{\mathrm{MC}}$ (GeV)	30–50	50–70	80–110	130–150
1 track	81 %	77 %	71 %	70 %
3 tracks	10 %	16 %	16 %	20 %
1 or 3 tracks	91 %	93 %	87 %	90 %

and the p_{T} of the leading track is required to be greater than 10 GeV/c. Jets are reconstructed in the calorimeter with the iterative cone algorithm using cone size of 0.4. The reconstructed QCD jets should match the two leading E_{T} Monte Carlo jets as described above.

12.1.2.3 Tagging by impact parameter

This method is similar to the one used in b-jet tagging (Section 12.2.2): the jet must have a minimum number of tracks with a significant impact parameter. Only τ jets with either 1 or 3 tracks in the signal cone are considered, since they have already passed the tracker isolation and the track counting criteria described in the previous section. The use of an impact parameter requirement to tag τ jets is not expected to be as selective as its equivalent use in b-jet tagging. This is because the proper lifetime of τs is 5 times shorter than that of b hadrons and the charged particle multiplicity in hadronic τ decays is also lower than in b-hadron decays. The τ-jet tagging performance for jets with 1 track in the signal cone is presented. The performance for jets with 3 tracks in the signal cone was found to be inefficient for QCD-jet rejection.

Tracks are reconstructed using the combinatorial track finder algorithm requiring at least 8 hits per track and a normalized χ^2 <10. An upper cut of 300 μm on the transverse impact parameter was applied to reject QCD jets where fake tracks have large impact parameters. It is predominantly fake tracks that satisfy the tracker isolation requirement and the requirement to have only one track in the signal cone. This leads to a large tail in the impact parameter distribution. The fake tracks are built up from hits belonging to a number of different charged particles. The fraction of QCD jets containing fake tracks increases with increasing jet E_{T} and η. The optimization of the upper cut on the transverse impact parameter is shown in Figure 12.10(left plot). It displays the efficiency for selecting τ and QCD jets in different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ when the upper cut on the transverse impact parameter is increased starting from 100 μm (highest efficiency point in each set) in steps of 100 μm. At 300 μm the efficiency for selecting τ jets is greater than 95%.

Figure 12.10 (right plot) shows the performance of τ-jet tagging with the transverse impact parameter requirement included. The efficiency for selecting τ and QCD jets is shown for different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ when the lower cut on the significance of the unsigned transverse impact parameter (the impact parameter divided by its uncertainty) is increased, starting from 0, in steps of 1. An upper cut of 300 μm on the transverse impact parameter is also applied and its efficiency is included. The rejection factor for QCD jets is larger at high E_{T} while the efficiency for selecting τ jets is almost independent of E_{T}.

The performance of the τ-jet tagging algorithm using a 3D impact parameter is found to be better than when using a 2D (transverse) impact parameter, assuming a 60 μm resolution on the z coordinate of the primary vertex position in τ jet events. For QCD multi-jet events the z coordinate of the primary vertex position was found using the pixel vertex finder algorithm. Figure 12.11 shows the efficiency for selecting τ and QCD jets for different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ when the lower cut on the significance of the unsigned 3D impact parameter is increased, starting from 0, in steps of 1. An upper cut of 1 mm on the 3D impact parameter is also applied and its efficiency is included.

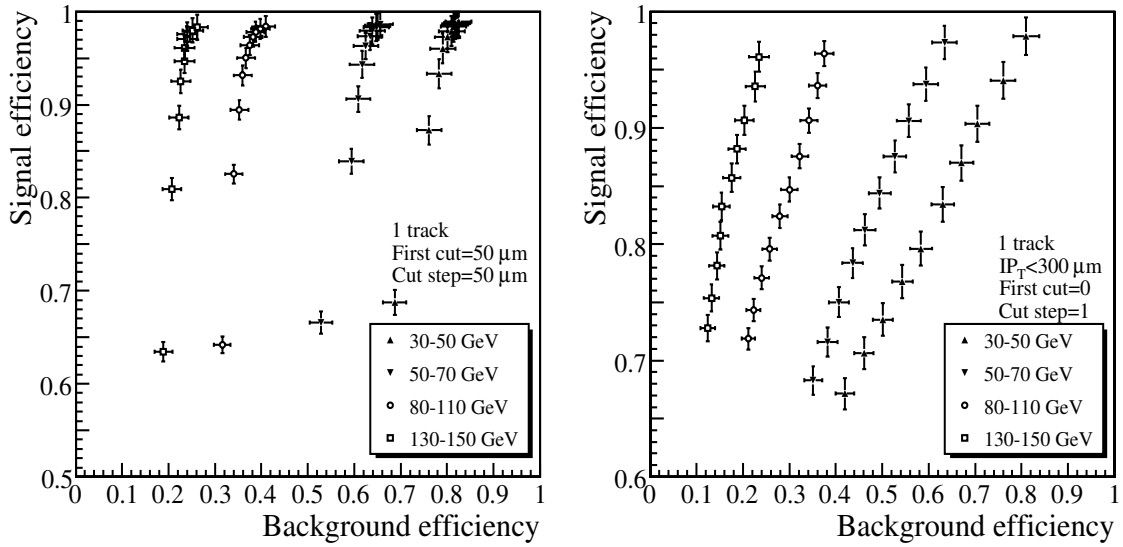


Figure 12.10: Left plot: the efficiency for selecting τ and QCD jets in different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ when the upper cut on transverse impact parameter is increased, starting from 100 μm (highest efficiency point in each set). in steps of 100 μm. Right plot: the efficiency for selecting τ and QCD jets in different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ when the lower cut on the significance of the unsigned transverse impact parameter is increased starting from 0 (highest efficiency point in each set) in steps of 1. An upper cut 300 μm on the transverse impact parameter is also applied and its efficiency is included.

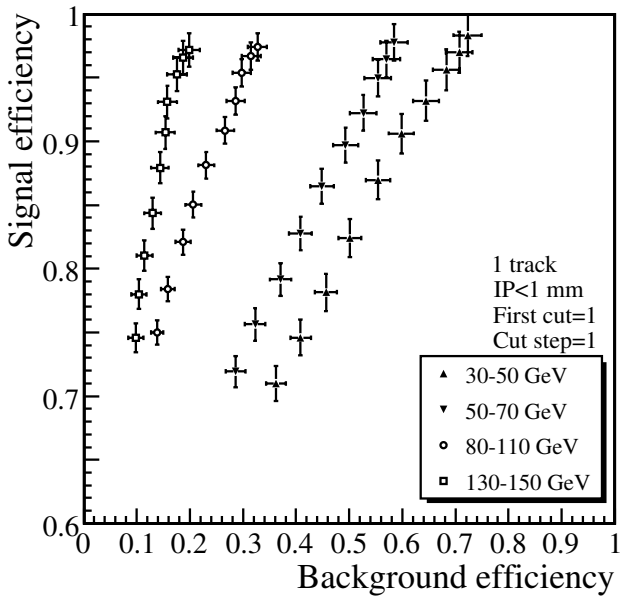


Figure 12.11: The efficiency for selecting τ and QCD jets for different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ when the lower cut on the significance of the unsigned 3D impact parameter is increased, starting from 0, in steps of 1. An upper cut of 1 mm on the unsigned 3D impact parameter is also applied and its efficiency is included.

12.1.2.4 Tagging with flight path

The life-time of the τ lepton ($c\tau = 87.11\ \mu\mathrm{m}$) allows reconstruction of the decay vertices of three- and five-prong τ decays that are well separated from the signal primary vertex. The τ tagging algorithm requires that the flight path between the signal primary vertex and the decay vertex of τ lepton exceed some minimum length. Tracks used by the vertex fitter are required to be in the signal cone. The performance of tagging with the flight path requirement was evaluated for three-prong τ decays when 3 reconstructed tracks are found in the signal cone. Table 12.3 shows the fraction of τ jets with a certain number of the reconstructed tracks in the signal cone for one- and three-prong τ decays. For three-prong τ decays the

Table 12.3: The fraction of τ leptons with the indicated numbers of the reconstructed tracks in the signal cone for one- and three-prong τ decays

	1 track	2 tracks	3 tracks	>3 tracks
1-prong τ	88.4 $\pm$ 0.3 %	6.1 $\pm$ 0.1 %	4.1 $\pm$ 0.1 %	1.4 $\pm$ 0.1 %
3-prong τ	8.6 $\pm$ 0.1 %	16.1 $\pm$ 0.2 %	63.2 $\pm$ 0.4 %	12.1 $\pm$ 0.2 %

probability to reconstruct 3 tracks in the signal cone was found to be $\simeq 63\%$.

Figure 12.12 (left plot) shows the reconstructed transverse flight path of the τ lepton for 2 intervals of the true transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$ <150 GeV and $E_{\mathrm{T}}^{\mathrm{MC}} = 150$–420 GeV. The fake secondary vertices reconstructed in the first and second barrel layers of the Pixel detector are visible as bumps at $\simeq 40$ mm and 70 mm. The contamination of the fake vertices is larger for more energetic τ jets where the 3 charged particles are more collimated. The 3 charged particles from the high E_{T} τ jets can produce overlapping clusters at a pixel layer,

thus leading to the reconstruction of 1 hit for all 3 tracks, and forcing the vertex fitter to reconstruct the vertex position at the location of the pixel layer.

Figure 12.12 (right plot) shows the reconstructed transverse flight path for QCD jets generated with $\hat{p}_{\mathrm{T}} = 80$–$120$ GeV$/c$. The jets with $E_{\mathrm{T}}^{\mathrm{MC}} > 50$ GeV were required to pass the tracker isolation requirement and to have 3 reconstructed tracks in the signal cone. Fake vertices produce the bump at $\simeq 40$ mm in the first barrel pixel layer as well as the tail. The gray histogram shows the contribution of c and b quark jets in the sample. The content of c and b

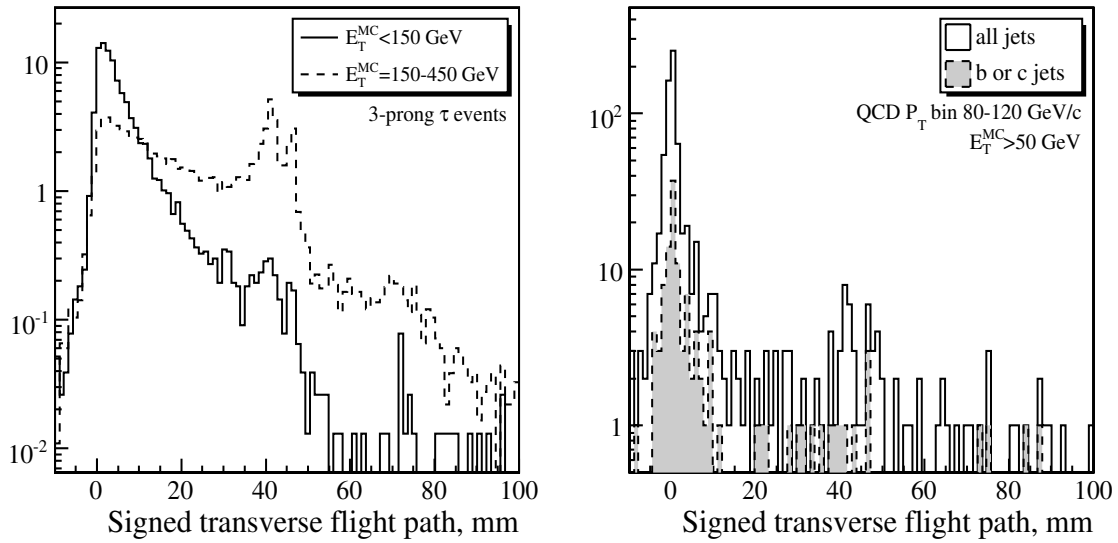


Figure 12.12: The signed transverse flight path for τ (left) and QCD (right) jets passing the tracker isolation requirement. Three reconstructed tracks are required to be in the signal cone. See text for more details.

jets is presented in Table 12.4 for different intervals of $E_{\mathrm{T}}^{\mathrm{MC}}$.

Table 12.4: The fraction of c and b jets in QCD jets in different bins of $E_{\mathrm{T}}^{\mathrm{MC}}$ after the tracker isolation and the requirement to have 3 reconstructed tracks in the signal cone.

$\mathrm{E_T}$ bin (GeV)	fraction of c jets (%)	fraction of b jets (%)
30-50	13.1 ± 2.1	3.1 ± 1.0
50-70	12.3 ± 1.6	4.4 ± 1.0
80-110	13.4 ± 1.4	2.7 ± 0.6
130-150	12.4 ± 1.9	3.0 ± 1.0

The z coordinate of the primary vertex position is found with the pixel vertex finder for QCD multi-jet events. For single τ jets the Monte Carlo value of the primary vertex z coordinate is smeared with the resolution of 60 μm. It is found that the error in the flight path is completely dominated by the secondary vertex resolution. The Kalman, Adaptive, and Principal vertex fitters are used for secondary vertex reconstruction, and all are found to provide the same accuracy as the secondary vertex position. In the plane transverse to the τ-jet axis, the resolution is between 18 and 25 μm and is independent of the τ energy for $E_{\mathrm{T}}^{\mathrm{MC}}$ between 30

and 300 GeV. The resolution in the direction parallel to the τ-jet axis ($\sigma_{\rm long}$) depends on the jet $E_{\rm T}$. The resolution $\sigma_{\rm long}$ is shown in Figure 12.13 (left plot) as a function of the $E_{\rm T}^{\rm MC}$ for 3 vertex fitters. The resolution increases from 500 μm to $\simeq 1.5$ mm when the transverse τ-jet energy is increased from 30 to 250 GeV.

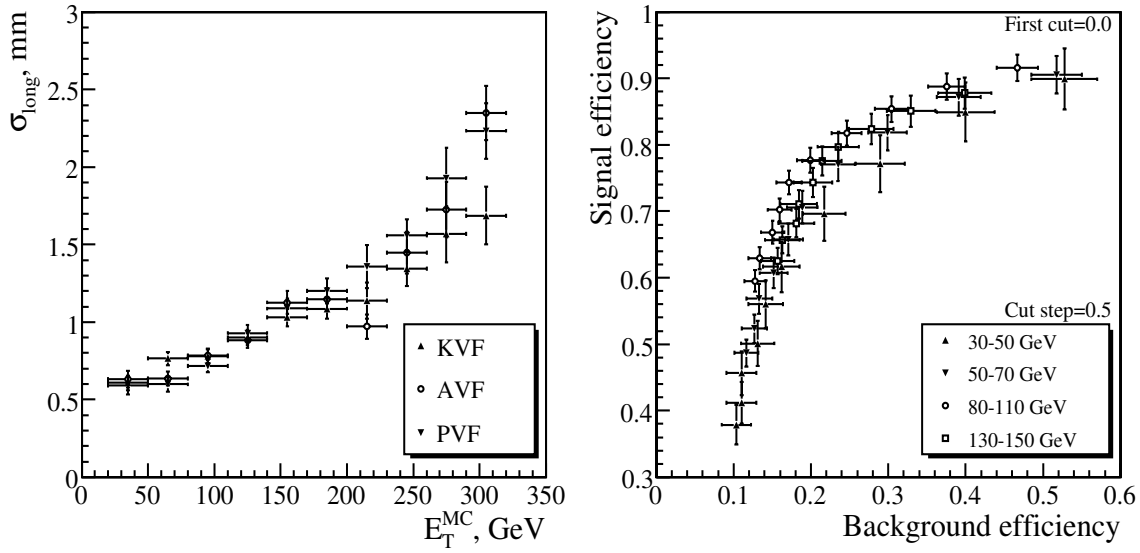


Figure 12.13: Left plot: the secondary vertex resolution in μm for τ jets in the direction parallel to the τ-jet axis as a function of τ-jet $E_{\rm T}^{\rm MC}$ for Kalman, Adaptive, and Principal vertex fitters. Right plot: the efficiency for the τ and QCD jets in bins of $E_{\rm T}^{\rm MC}$ when the lower cutoff on the significance of the signed 3D flight path is varied. An upper cut of 35 mm on the transverse flight path is used and its efficiency is included.

The performance of the τ-jet tagging algorithm with the flight path requirement applied was evaluated using a Kalman vertex fitter. Figure 12.13 (right plot) shows the efficiency for the τ and QCD jets in bins of $E_{\rm T}^{\rm MC}$ when the lower cutoff on the significance of the signed 3D flight path is varied. The upper cut of 35 mm on the transverse flight path was used to suppress fake vertices; the efficiency of this cut is included in the performance plot. It can be concluded that the rejection factor of 5 can be achieved with an efficiency of 70–80% for jets between 30 and 150 GeV.

12.1.2.5 Tagging with mass

In this section we discuss the τ-tagging method with a constraint on the reconstructed mass of the τ-jet, $M_{\tau\ {\rm jet}}$. After the tracker isolation criterion is applied, the τ-jet mass is reconstructed from the momentum of the tracks in the signal cone and the energy of the clusters in the electromagnetic calorimeter that lie within a certain cone in the η-ϕ space around the calorimeter jet axis.

Figure 12.14 shows scatter plots of the transverse energy of the electromagnetic clusters, $E_{\rm T}^{\rm em}$, and the distance $\Delta R_{\rm jet}$ in the η-ϕ space between the calorimeter jet axis and the clusters for the τ jets (left plot) and the QCD jets (right plot), with $E_{\rm T}^{\rm MC}$ between 30 and 150 GeV. The $\Delta R_{\rm jet}$ and $E_{\rm T}^{\rm em}$ are strongly correlated, so a constraint on one variable will sculpt the other one. The clusters are closer to the jet axis for the τ jets than for the QCD jets and for both τ

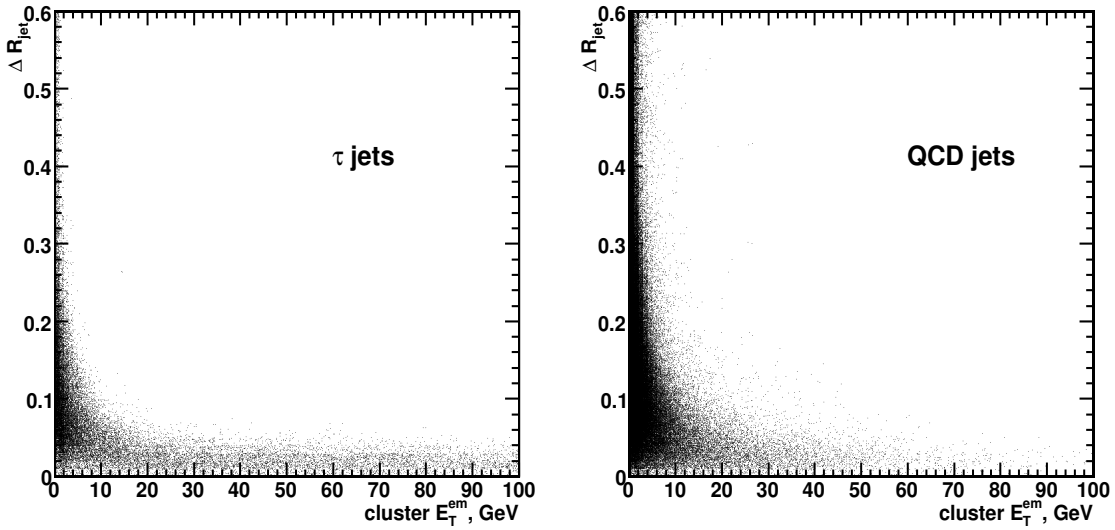


Figure 12.14: Scatter plots of $\Delta R_{\rm jet}$, the distance in η-ϕ space between the calorimeter jet axis and the cluster, versus the transverse energy of electromagnetic clusters, $E_{\rm T}^{\rm em}$, for τ jets (left) and for QCD jets (right) for $E_{\rm T}^{\rm MC}$ between 30 and 150 GeV.

and QCD jets the energetic clusters with $E_{\rm T}^{\rm em} > 10$ GeV are located mostly within the cone of size 0.1 around the calorimeter jet axis. A cone size of 0.4 was found to be optimal for τ tagging performance. A smaller cone size reduces the ability to distinguish τ and QCD jets with the constraint on the $M_{\tau\ {\rm jet}}$ value.

The $M_{\tau\ {\rm jet}}$ calculated from the tracks in the signal cone and the clusters with $\Delta R_{\rm jet} < 0.4$ reveals a very broad distribution with a long tail, as shown in Figure 12.15 (dashed line). This tail is due to double counting, when the clusters in the ECAL produced by the charged particles are included in the $M_{\tau\ {\rm jet}}$ calculation. These clusters are rejected by track-cluster matching analysis, in which the cluster used for the mass calculation must be separated from the track impact point on the ECAL surface by the distance $\Delta R_{\rm track}$ >0.08 m. As expected, this constraint improves the distribution, as shown by the solid line in Figure 12.15. The large peak at zero in $M_{\tau\ {\rm jet}}$ is due to the single track events with no ECAL clusters satisfying the constraints on $\Delta R_{\rm jet}$ and $\Delta R_{\rm track}$.

Distributions of $M_{\tau\ {\rm jet}}$ for different $E_{\rm T}^{\rm MC}$ bins of τ and QCD jets are shown in Figures 12.16 and 12.17 where the τ jets are represented by a solid line and the QCD jets by a dashed line. The selection efficiency of the cut $M_{\tau\ {\rm jet}} < 2.5\ {\rm GeV}/c^2$ is presented in Table 12.5 for τ and QCD jets in different intervals of $E_{\rm T}^{\rm MC}$. The efficiency for the τ jets is almost independent of the jet $E_{\rm T}$, while there is very strong dependence for the QCD jets: the higher jet $E_{\rm T}$ the greater the jet rejection.

12.1.2.6 Rejection of electrons

An electron which originates from the primary signal vertex will satisfy all of the τ tagging criteria described above except for the impact parameter tagging criteria. To suppress the electron-τ-jet misidentification rate in offline analysis, we require that some minimum hadronic energy should be found in the HCAL. Two methods were evaluated and found to

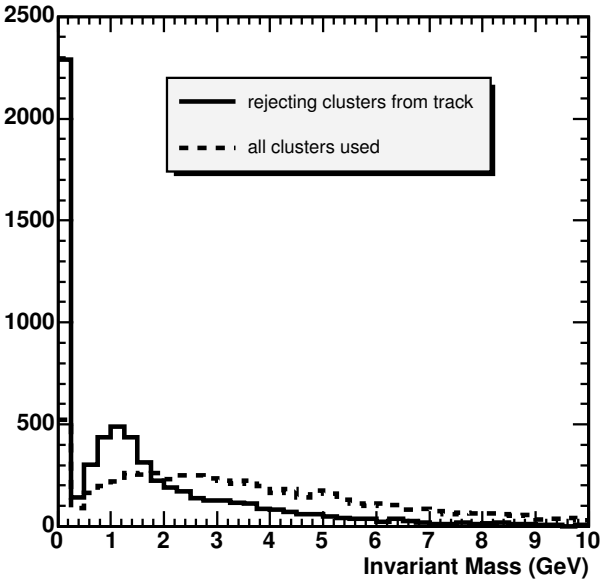


Figure 12.15: The $M_{\tau\ \mathrm{jet}}$ distribution for τ jets when all clusters with $\Delta R_{\mathrm{jet}} < 0.4$ are taken (dashed line) and when clusters not matched with tracks ($\Delta R_{\mathrm{track}}$ >0.08 m) are used (solid line).

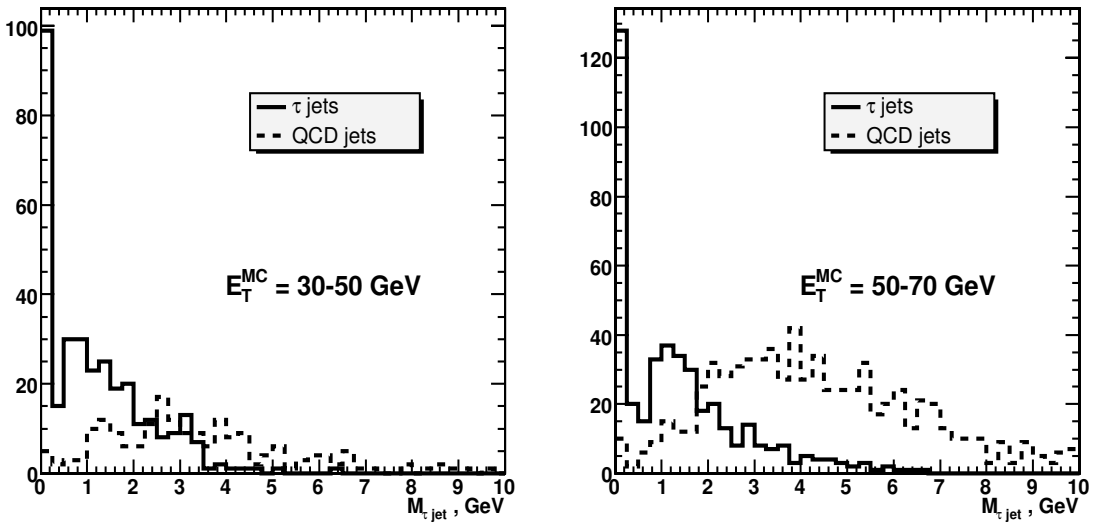


Figure 12.16: The $M_{\tau\ \mathrm{jet}}$ distributions for τ jets (solid line) and QCD jets (dashed line) in the $E_{\mathrm{T}}^{\mathrm{MC}}$ bins of 30–50 GeV (left plot) and 50–70 GeV (right plot).

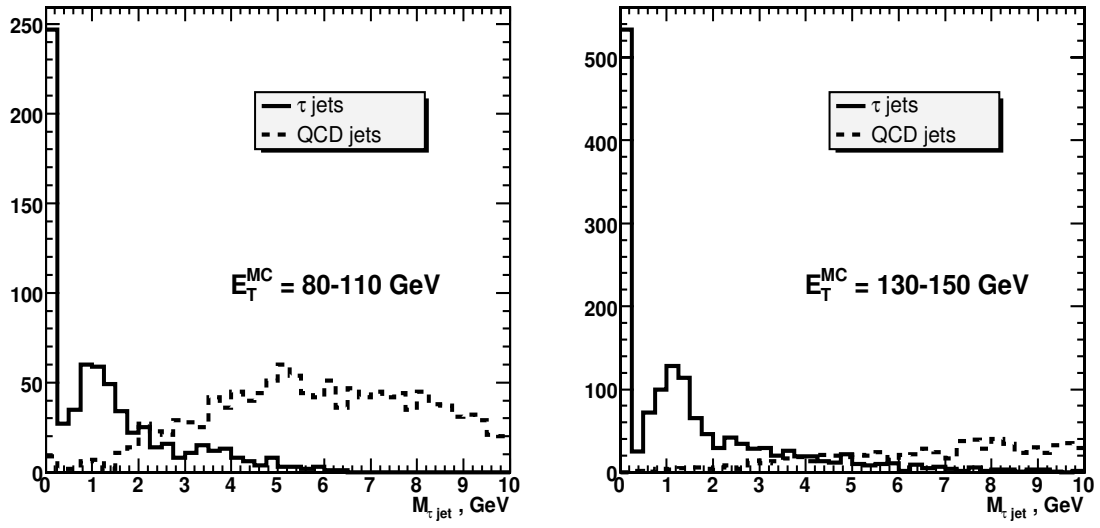


Figure 12.17: The $M_{\tau\ \mathrm{jet}}$ distributions for τ jets (solid line) and QCD jets (dashed line) in the $E_{\mathrm{T}}^{\mathrm{MC}}$ bins of 80–110 GeV (left plot) and 130–150 GeV (right plot).

Table 12.5: The selection efficiency of the cut $M_{\tau\ \mathrm{jet}} < 2.5\ \mathrm{GeV}/c^2$ for τ and QCD jets in different intervals of $E_{\mathrm{T}}^{\mathrm{MC}}$.

$E_{\mathrm{T}}^{\mathrm{MC}}$ bins (GeV)	30–50	50–70	80–110	130–150
Eff. for τ jets, %	86.32	82.27	83.02	80.76
Eff. for QCD jets, %	33.67	19.16	6.05	2.47

provide similar performance. In the first method, electron rejection is achieved by using a lower cutoff on the transverse energy of the hottest (maximum E_{T}) HCAL tower belonging to the reconstructed jet. In the second method the lower cutoff is applied to the value of $E_{\mathrm{T}}^{\mathrm{hadr}}/p_{\mathrm{T}}^{\mathrm{ltr}}$, where $E_{\mathrm{T}}^{\mathrm{hadr}}$ is the transverse energy of the τ jet measured in the HCAL only and $p_{\mathrm{T}}^{\mathrm{ltr}}$ is the transverse momentum of the leading track measured in the tracker. The performance of the first method is presented below.

Figure 12.18 shows the transverse energy of the hottest HCAL tower belonging to the reconstructed jet for an electron of $p_{\mathrm{T}} = 35$ GeV/c reconstructed as a jet (solid line) and τ jets in 2 ranges of the true τ-jet transverse energy, 40–60 GeV (dashed line) and 100–140 GeV (dotted line). The cut on the measured transverse momentum of the leading track in the τ jet, $p_{\mathrm{T}}^{\mathrm{ltr}} > 10$ GeV/c, was applied. The histograms in Figure 12.18 are normalized to unity, thus showing the probability distribution per 1 GeV bin. Table 12.6 shows the efficiency of the

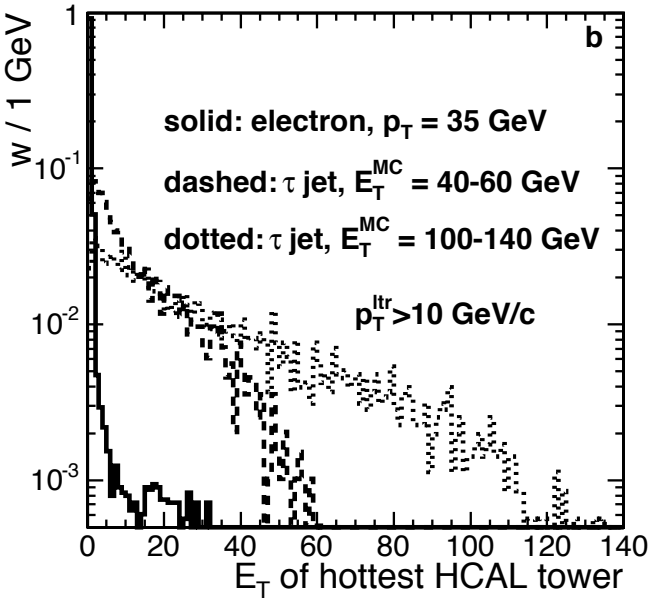


Figure 12.18: The transverse energy of the hottest (maximal E_{T}) HCAL tower belonging to the reconstructed jet. Solid line—electron of $p_{\mathrm{T}} = 35$ GeV/c reconstructed as a jet. Dashed (dotted) line—τ jet in the range true τ-jet transverse energy of 40–60 GeV (100–140 GeV). A cut on the measured transverse momentum of the leading track in τ jet, $p_{\mathrm{T}}^{\mathrm{ltr}}$ >10 GeV/c, was applied. All histograms are normalized to unit area.

cut on the transverse energy of the hottest HCAL tower belonging to the jet for an electron of p_{T} =35 GeV/c reconstructed as a jet and the τ jet in 2 ranges of the true transverse energy of the τ jet.

Muon-τ jet misidentification is not considered, since average energy losses of the muon in the calorimeter are of order of a few GeV, therefore well below the lowest threshold on the E_{T} of τ jet used in the physics analyses (15–30 GeV).

12.1.3 High Level Trigger

ECAL and tracker isolation requirements are used in the High Level Trigger for τ-jet identification. The performance of τ tagging in the HLT was evaluated for the most difficult case of triggering on the decay of the MSSM neutral Higgs boson into 2 τ leptons when

Table 12.6: Efficiency of the cut on the transverse energy of the hottest HCAL tower (maximal E_{T}) belonging to the jet for the electron of $p_{\mathrm{T}} = 35$ GeV/c reconstructed as a jet and τ jet in 2 ranges of the true transverse energy of the τ jet and different cuts on transverse momentum of the leading track ($p_{\mathrm{T}}^{\mathrm{ltr}}$) in τ jet.

cut	electron	τ jet E_{T} 40–60 GeV		τ jet E_{T} 100–140 GeV	
		$p_{\mathrm{T}}^{\mathrm{ltr}} > 10$ GeV	$p_{\mathrm{T}}^{\mathrm{ltr}} > 25$ GeV	$p_{\mathrm{T}}^{\mathrm{ltr}} > 10$ GeV	$p_{\mathrm{T}}^{\mathrm{ltr}} > 25$ GeV
>1 GeV	0.08	0.936	0.971	0.977	0.991
>2 GeV	0.03	0.854	0.917	0.942	0.969

both τ's decay hadronically, thus producing 2 τ jets in the final state. In the High Level Trigger the double τ-jet identification is needed to suppress the rate from single or double Level-1 Tau trigger [287]. Table 12.7 shows the QCD multi-jet background rate in kHz at L=2×10^{33}s^{-1}cm^{-2} for the Level-1 single, double, and single or double Tau triggers with the single (double) trigger threshold 93 (66) GeV optimized in [8, 289]. The rate is shown for different intervals of generated QCD background transverse momentum ($\hat{p}_{\mathrm{T}}$ bins) between 30 and 300 GeV/c. The 3 $\hat{p}_{\mathrm{T}}$ bins in the interval between 50 and 170 GeV/c give the dominant (> 90 %) contribution to the rate. Therefore these 3 bins are used to evaluate the rejection factor obtained in the HLT with the double τ-jet tagging. The signal efficiency of the HLT selection was evaluated for 2 masses of the MSSM neutral Higgs boson, 200 and 500 GeV/c^2, produced in association with a $b\bar{b}$ quark pair.

Table 12.7: The QCD multi-jet background rate in kHz at L=2×10^{33}s^{-1}cm^{-2} for the Level-1 single, double, and single or double Tau triggers with single (double) trigger threshold 93 (66) GeV. The rate is shown for different intervals of generated QCD background transverse momentum ($\hat{p}_{\mathrm{T}}$ bins).

$\hat{p}_{\mathrm{T}}$ (GeV/c)	cross section (fb)	Rate (kHz)		
		single Tau	double Tau	single or double Tau
30-50	1.56×10^{11}	0.04	0.08	0.12
50-80	2.09×10^{10}	0.59	0.70	1.19
80-120	2.94×10^{9}	1.32	0.75	1.65
120-170	5.00×10^{8}	0.46	0.16	0.48
170-230	1.01×10^{8}	0.10	0.03	0.10
230-300	2.39×10^{8}	0.02	0.007	0.021
total rate		2.53	1.73	3.56

In the High Level Trigger the 2 jets are reconstructed with the calorimeter in regions given by the first and the second Level-1 Tau jets. If the second Level-1 Tau jet does not exist in the Global Level-1 calorimeter trigger output, the jet is reconstructed in the region of the first Level-1 Central jet. For signal events the 2 jets selected in this way have good purity, 97 % for the first and 82 % for the second jet; moreover the purity does not depend on the Higgs boson mass between 200 and 800 GeV/c^2.

The next step in the HLT selection is the τ tagging of these 2 jets. Two different approaches are investigated:

- ECAL isolation, followed by tracker isolation with the tracks reconstructed using only the pixel detector.
- Tracker isolation with regional track reconstruction using both pixel and silicon strip tracker layers.

The first approach is fast and gives good performance as far as the isolation algorithm is concerned. It is therefore the preferred approach for decays with 2 taus in the final state (e.g. $A/H \rightarrow \tau\tau$) where the isolation is sufficient to reach the required background rejection factor. The second approach is slower but gives a much more accurate estimate of the momenta of the tracks. It is therefore useful in channels like charged Higgs boson decay to a τ lepton and neutrino. The HLT selection of events with high missing E_{T} must be complemented by the τ tracker isolation and a tight cut on the momentum of the leading p_{T} track in the signal cone [8].

The other tagging methods previously discussed (track counting, impact parameter, flight path, invariant mass) are not used at the High Level trigger since the isolation alone and the cut on the p_{T} of the leading track in the signal cone proved capable of rejecting the QCD background to an acceptable level.

More details on the logic of the trigger system can be found in [8, 287, 290, 291].

12.1.3.1 The τ selection based on ECAL and pixel isolation

The ECAL plus pixel-track isolation in the High Level Trigger will be referred to as the **Calo+Pxl** Tau trigger. In this approach, the Level-1 rate is first suppressed by a factor of $\simeq 3$ by applying the ECAL isolation to the first jet only. Figure 12.19 shows the efficiency of the ECAL isolation for signal and QCD multi-jet background as a function of the ECAL isolation parameter cut $P_{\mathrm{isol}}^{\mathrm{cut}}$. The efficiency is shown for events which pass the Level-1 single or double Tau trigger. A rejection factor of 3 can be achieved with $P_{\mathrm{isol}}^{\mathrm{cut}} = 5\ \mathrm{GeV}/c$. The remaining background rate is suppressed by the tracker isolation requirement, using only the information from the Pixel detector. The isolation is applied to both jets.

With 3 pixel hits one can reconstruct tracks using the pixel detector only; such tracks are called pixel-tracks. The algorithm used to find tracks was explained in detail in [292], here only the most important details will be given. Pixel hit pairs from the first 2 layers (barrel+barrel or barrel+endcap) are matched in r-ϕ and z-r to establish track candidates. The cuts are optimized for a minimum track p_{T} of 1 GeV/c. Valid pixel pairs are matched with a third pixel hit forming pixel-tracks. The momentum of the pixel-tracks is then reconstructed by making use of those 3 pixel hits without a primary vertex constraint. The number of fake pixel-tracks in the isolation cone was found to be very low (3-4%).

Using pixel-tracks, a list of primary vertices is formed at z values where several tracks cross the z axis. Only primary vertices with at least 3 valid tracks are kept and the position of each vertex is estimated as the mean value of the z impact parameters of all tracks assigned to it [292].

Pixel-tracks are then used by the isolation algorithm which was described in Section 12.1.2.2. The leading p_{T} track found in the cone $R_{\mathrm{M}} = 0.1$ around the calorimeter jet axis must have a transverse momentum $p_{\mathrm{T}}^{\mathrm{LT}}$ higher than 3 GeV/c. The leading track from the first jet defines the primary vertex. This selection is very pure, corresponding to the position of the true Monte Carlo vertex in 99% of all cases. All tracks used for the jet tagging must be associated

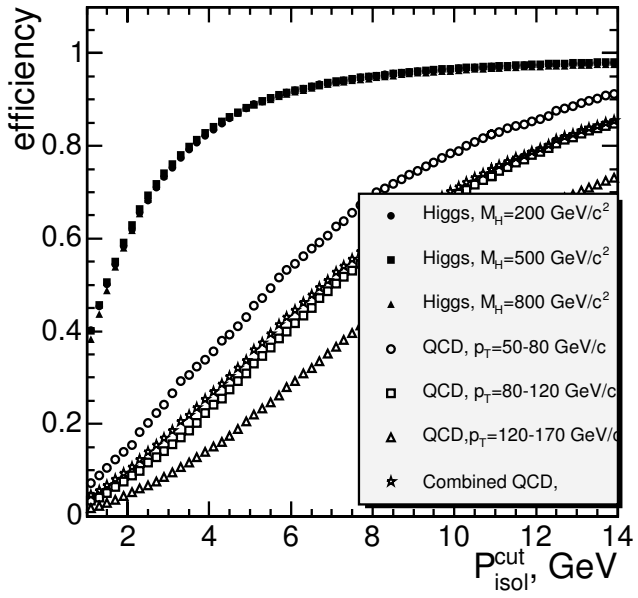


Figure 12.19: The efficiency of the ECAL isolation in the High Level Trigger for the signal and the QCD multi-jet background as a function of the cut $P_{\rm isol}^{\rm cut}$ on the ECAL isolation parameter.

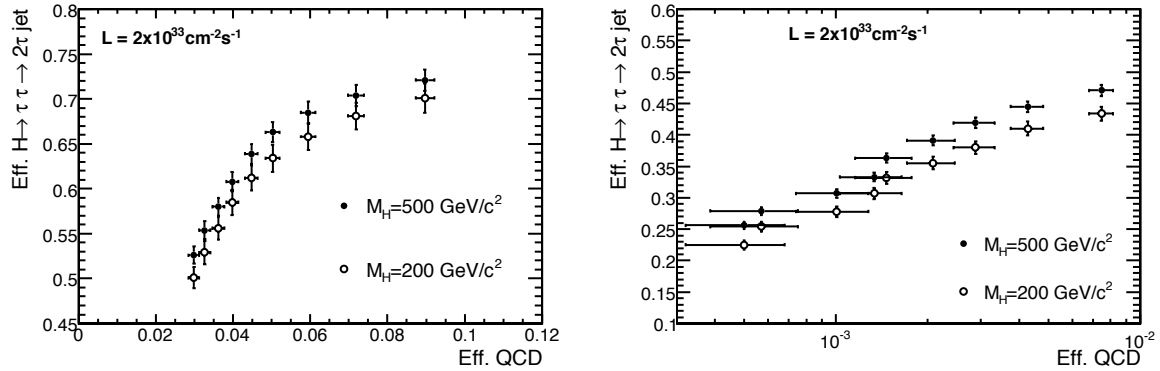


Figure 12.20: Efficiency of **Calo+Pxl** trigger applied to the first jet (left plot) and to both jets (right plot) for signal events versus efficiency for QCD multi-jet events. Two Higgs boson masses of $M_{\rm H}$=200 and 500 GeV/c^2 are shown. The isolation cone is varied from 0.2 to 0.6 in steps of 0.05; the signal cone is 0.07; the matching cone is 0.1; and the $p_{\rm T}$ of the leading tracks must exceed 3 GeV/c.

with this vertex. In addition all tracks considered in the selection of the second jet must also be associated with the same primary vertex that was used for the first jet. Therefore tracks from other vertices are simply ignored.

The signal cone $R_{\rm S}$ was set to 0.07 and the isolation cone $R_{\rm i}$ was varied as a free parameter to adjust the trigger rate. Figure 12.20 shows the performance of the **Calo+Pxl** trigger. The selection efficiency plotted on both axes is defined relative to the events that passed the Level-1 single or double Tau trigger. Figure 12.20 (left plot) shows the performance for the first jet and (right plot) the performance of the double jet tagging. The different points correspond to different sizes of the isolation cone $R_{\rm i}$ which varies between 0.2 and 0.6. The

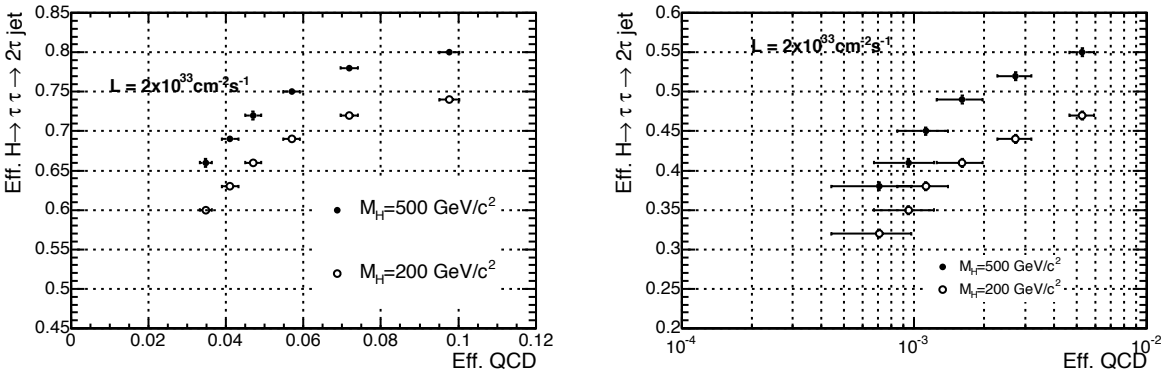


Figure 12.21: Efficiency of **Trk Tau** trigger applied to the first jet (left plot) and to both jets (right plot) for signal events versus efficiency for QCD multi-jet background. Two masses of the Higgs boson, $M_{\rm H} = 200$ and 500 GeV/c^2, are considered. The isolation cone is increased from 0.2 to 0.45 in steps of 0.05, the signal cone is 0.07, the matching cone is 0.1, and the $p_{\rm T}$ of the leading tracks must exceed 6 GeV/c.

required suppression of the QCD multi-jet background of about 10^{-3} can be achieved with the $R_{\rm i}$ between 0.45 and 0.50, with the signal efficiency being 0.30–0.33.

A similar study was performed previously [8], but with somewhat different conditions. In particular, that study used a) a simpler multiple interaction model in PYTHIA; b) a different detector simulation model with smaller electronic noise in the HCAL; c) a different strategy to search for the second τ-jet candidate after Level-1 trigger; and d) preselection cuts at the generator level for the signal events. Allowing for some differences, the comparisons are satisfactory. For a Higgs boson of $M_{\rm H} = 500$ GeV/c^2, for instance, the efficiency for the first tau jet was 0.68 in a previous study [8] as compared to 0.67 obtained here for the same value of $R_{\rm i} = 0.35$.

12.1.3.2 The τ selection based on silicon tracker isolation

The algorithm described in this paragraph will be referred to as the **"Trk Tau"** trigger. Due to time limitations in the HLT it is not possible to perform a full tracker reconstruction of the whole event right after the Level-1 Tau trigger. It is possible, however, to read and reconstruct a selected part of the tracker data. The **Trk Tau** trigger only reconstructs those tracks confined to the restricted regions of the interest ("regional tracking"), defined by the cone around the calorimeter jet direction. The primary signal vertex needed in the **Trk Tau** trigger is obtained using only the pixel detector in order to ensure fast reconstruction. Once the signal vertex is found, the regional track reconstruction begins. The tracker isolation algorithm is then applied using tracks with a z-impact parameter close to the z position of the signal vertex.

Trk Tau trigger performance.

Figure 12.21 shows the **Trk Tau** trigger performance in terms of the signal vs background efficiency for the events passing the Level-1 single or double Tau trigger. The left plot shows the efficiency for the first jet, while the right plot shows the double jet tagging efficiency. The different points correspond to the different size of the isolation cone $R_{\rm i}$ as it is increased from 0.2 to 0.45 in steps of 0.05. The matching cone is set to 0.1, the signal cone is 0.07. The leading

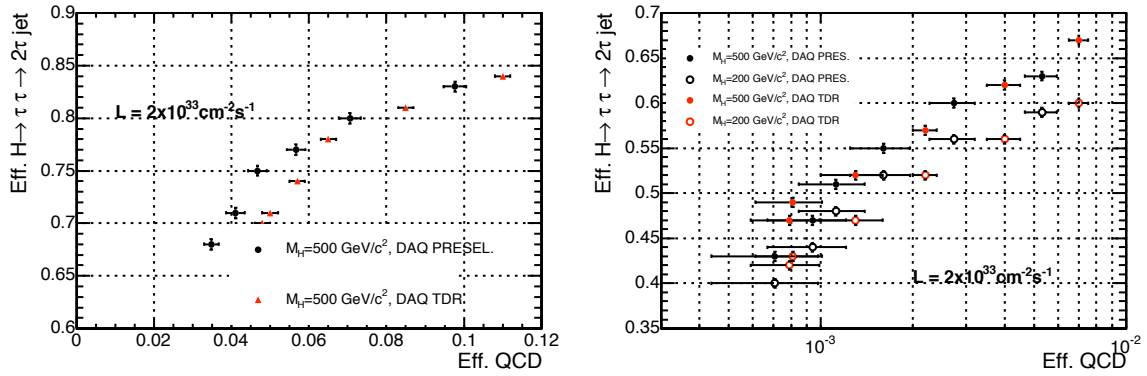


Figure 12.22: Efficiency of **Trk Tau** trigger applied to the first jet (left plot) and to the both jets (right plot) for signal events versus efficiency for QCD multi-jet background, using DAQ TDR preselections for the signal. Two masses of the Higgs boson, M_{H}=200 and 500 GeV/c^2, are considered. The isolation cone is varied from 0.2 to 0.45, the signal cone is 0.07, the matching cone is 0.1 and the $\mathrm{p_T}$ of the leading tracks must exceed 6 GeV/c. The DAQ TDR results [8] are also presented in the plots.

track momentum $p_{\mathrm{T}}^{\mathrm{LT}}$ must exceed 6 GeV/c. The rejection factor of $\simeq 10^3$ against the QCD multi-jet background can be achieved with R_{i} around 0.40.

Comparison with DAQ TDR

The main differences between the DAQ TDR [8] and new studies are mentioned in section 12.1.3.1. Another source of difference is a new implementation of regional seeding and the primary vertex reconstruction algorithm. In order to compare with the DAQ TDR results [8], the same preselections as in DAQ TDR at the generation level for the signal events $H/A \rightarrow \tau\tau \rightarrow$ jet jet were applied: $p_{\mathrm{T}}^{\tau\ \mathrm{jet}} > 45$ GeV/c and $|\eta^{\tau\ \mathrm{jet}}| <$2.4.

Figure 12.22 compares the **Trk Tau** trigger performance obtained at the present time and the DAQ TDR studies [8]. The left plot shows the efficiency for the first jet, while the right plot shows the double jet tagging efficiency. The different points correspond to different sizes of the isolation cone R_{i} varied between 0.2 and 0.45 in steps of 0.05. The efficiency for a Higgs boson with a mass 200 GeV/c^2 is in good agreement with the DAQ TDR results, while for a mass of 500 GeV/c^2 it is $\simeq$ 5 % lower. In the case of the double jet tagging, this can be explained, in particular, by $\simeq$ 8% lower purity of the second jet in the present study. The rejection factor for QCD multi-jet events is compatible, within less than 1.5σ, with the old results. The disagreement with the leftmost point (R_{i}=0.45) is due to the small size of the track reconstruction cone used in the DAQ TDR.

12.1.4 Calibration and tagging efficiency

12.1.4.1 Tau jet energy scale and calibration with calorimeter

The τ jet energy measurement with the calorimeter requires smaller energy corrections than "normal" QCD jets. The reasons are that, first, the average transverse momentum of the charged hadrons is higher, and second, the fraction of the electromagnetic energy in the τ jet brought by π^0s is larger. Figure 12.23 shows the τ-jet energy scale, the ratio $r = E_{\mathrm{T}}^{\mathrm{reco}}/E_{\mathrm{T}}^{\mathrm{MC}}$,

as a function of the $E_{\mathrm{T}}^{\mathrm{MC}}$ and τ-jet pseudorapidity for different final states of hadronic decays of τ lepton. Tau jets were reconstructed by the iterative cone algorithm with a cone size of 0.4. The thresholds on the calorimeter towers were set to $E_{\mathrm{T}} = 0.5\,\mathrm{GeV}$ and $E = 0.8\,\mathrm{GeV}$. One can

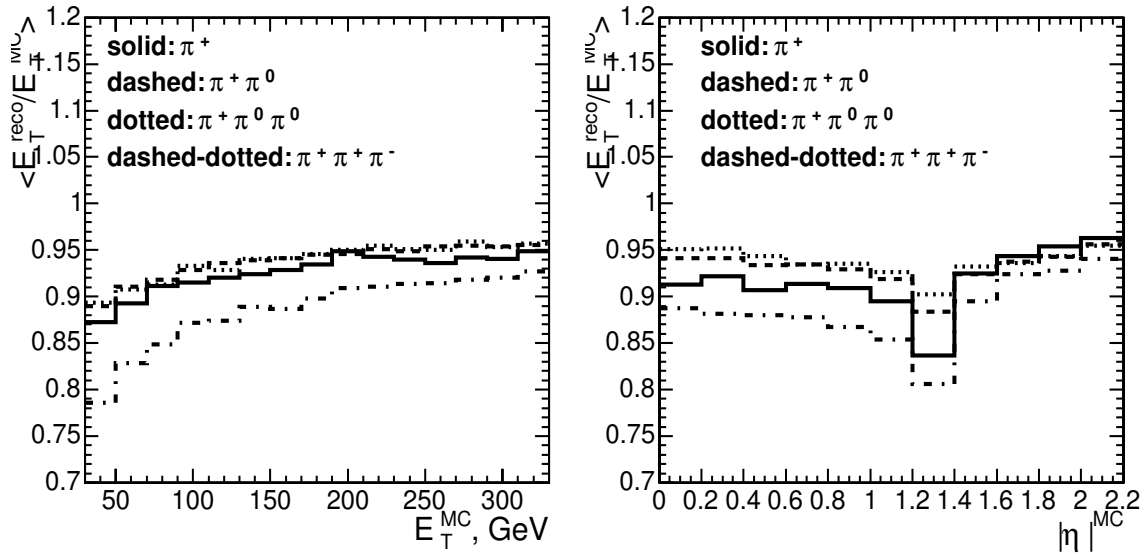


Figure 12.23: The ratio $r = E_{\mathrm{T}}^{\mathrm{reco}}/E_{\mathrm{T}}^{\mathrm{MC}}$ as a function of the $E_{\mathrm{T}}^{\mathrm{MC}}$ (left) and τ-jet pseudorapidity (right) for different final states of hadronic decays of τ lepton.

see that the τ jet formed by the 3 charged pions has the lowest response in the calorimeter in comparison to the jet containing only 1 charged and 1 or 2 neutral pions for the same $E_{\mathrm{T}}^{\mathrm{MC}}$. The drop in the response at the pseudorapidity $\simeq 1.4$ is due to the instrumentation gap between the barrel and the endcap calorimeters.

The energy correction function was obtained from the parameterization of the $E_{\mathrm{T}}^{\mathrm{MC}}$ and $|\eta^{\mathrm{MC}}|$ dependence of the ratio $r = E_{\mathrm{T}}^{\mathrm{reco}}/E_{\mathrm{T}}^{\mathrm{MC}}$. Figure 12.24 shows the $E_{\mathrm{T}}^{\mathrm{reco}}/E_{\mathrm{T}}^{\mathrm{MC}}$ ratio before and after the energy corrections were applied. The τ-jet transverse energy resolution after energy corrections can be parameterized with formula: $\sigma(E_{\mathrm{T}})/E_{\mathrm{T}} = a/E_{\mathrm{T}} \oplus b$, where $a = 0.883$ GeV and $b = 0.058$ for τ jets with E_{T} between 30 and 300 GeV and pseudorapidity less than 2.2.

The γ-plus-jet events where the jet passes the τ-identification criteria and thus becomes the "tau-like" jet can be used to setup the initial τ-jet energy scale from the real data. In the following we present the preliminary results. Figure 12.25 shows the ratio $r = E_{\mathrm{T}}^{\mathrm{reco}}/E_{\mathrm{T}}^{\mathrm{MC}}$ for the unpreselected QCD jets (dashed-dotted line) , "tau-like" QCD jets (dashed line) and the real τ jets (solid line). Both the QCD and the τ jets were reconstructed in the calorimeter with a cone size of 0.4. The same cone size was used to evaluate the true transverse energy $E_{\mathrm{T}}^{\mathrm{MC}}$ of the Monte Carlo QCD jets. The τ-jet identification includes the ECAL and the tracker isolation with the parameters $P_{\mathrm{isol}}^{\mathrm{cut}} = 5\ \mathrm{GeV}/c$, $R_{\mathrm{i}} = 0.4$, and $R_{\mathrm{S}} = 0.07$. The 1 or 3 tracks were required to be in the signal cone and a cut $p_{\mathrm{T}} > 10\ \mathrm{GeV}/c$ on the momentum of the leading track was applied. One can see that "tau-like" QCD jets produce a higher calorimeter response in comparison to the non-preselected QCD jets and it is only 5–10% smaller than the response of the real τ jets. More studies are needed to understand the sources of the remaining difference and the calibration uncertainties.

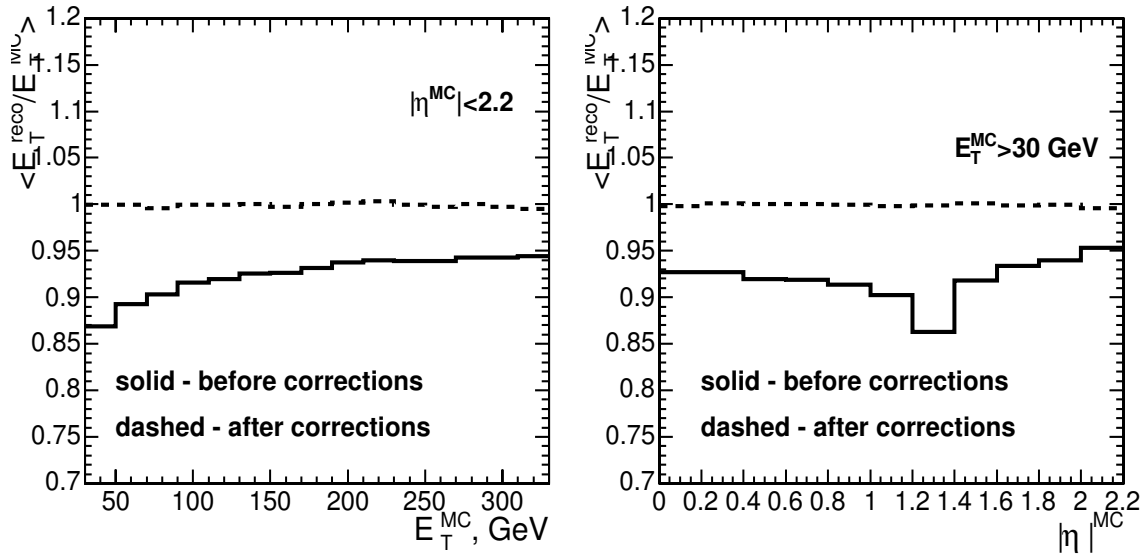


Figure 12.24: The ratio $r = E_{\mathrm{T}}^{\mathrm{reco}} / E_{\mathrm{T}}^{\mathrm{MC}}$ as a function of the $E_{\mathrm{T}}^{\mathrm{MC}}$ (left) and τ-jet pseudorapidity (right) before and after the energy corrections were applied.

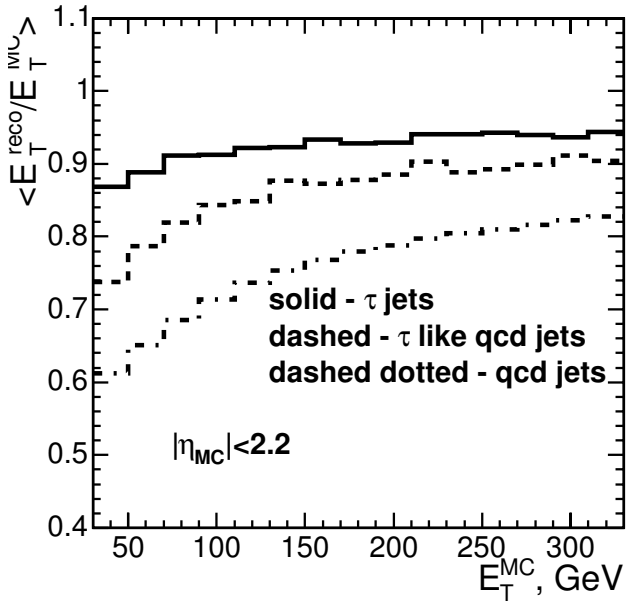


Figure 12.25: The ratio $r = E_{\mathrm{T}}^{\mathrm{reco}} / E_{\mathrm{T}}^{\mathrm{MC}}$ for the not preselected QCD jets (dashed-dotted line) , "tau-like" QCD jets (dashed line) and the real τ jets (solid line) as a function of $E_{\mathrm{T}}^{\mathrm{MC}}$. More detailed explanation can be found in the text.

Another method to evaluate the τ-jet energy scale with the data is to use $Z \to \tau\tau \to \ell + \mathrm{jet}$ events and reconstruct the Z mass peak. This method, however has 2 disadvantages: the background contamination and the uncertainty of the missing E_{T} measurement.

12.1.4.2 Measurement of $\mathrm{jet} \to \tau$ misidentification from the data

The measurement of the $\mathrm{jet} \to \tau$ misidentification rate can be determined with the γ-plus-jet events used for the calorimeter calibration. About 10^5 such events are expected in a 10 fb^{-1} data sample for each E_{T}^{γ} bin of size $0.1 \times E_{\mathrm{T}}^{\gamma}$, with $|\eta^{\mathrm{jet}}| < 3$ and E_{T}^{γ} in the interval between 30 and 300 GeV (see Section 11.6.3). The mistagging rate can then be equated to the fraction of events in which the jet passed the τ-jet identification criteria. Taking into account the jet rejection factor with, for example, the tracker isolation and the mass tagging (evaluated from Figure 12.9 (right plot) for $R_{\mathrm{i}} = 0.4$ and $R_{\mathrm{S}} = 0.04$, and using Table 12.5) one would expect 4-10% uncertainty in the estimated mistag rate per energy bin in the jet E_{T} interval of 30-150 GeV with a 10 fb^{-1} data sample.

12.1.4.3 Measurement of τ tagging efficiency from the data

The τ-tagging efficiency can be evaluated (and compared with the Monte Carlo simulation) from the ratio of $Z \to \tau\tau \to \mu + \mathrm{jet}$ and $Z \to \mu\mu$ events selected by the single muon trigger stream. It assumes that the reconstruction efficiency of the second muon in the $Z \to \mu\mu$ events is known. The preliminary estimates were obtained based on the search for MSSM $H/\mathrm{A} \to \tau\tau \to \mu + \mathrm{jet}$ channel described in [293]. The $Z \to \tau\tau \to \mu + \mathrm{jet}$ event selections are the same as used in [293], but without the b tagging and the jet veto. The systematic uncertainty related to the selection cuts which use the calorimeter information: cuts on $m_{\mathrm{T}}(\ell, E_{\mathrm{T}}^{\mathrm{miss}})$, $E_{\mathrm{T}}^{\tau\,\mathrm{jet}}$, and E_{T} of reconstructed neutrinos were taken into account, as well as the uncertainty of the background evaluation.

With a 30 fb^{-1} data sample the total uncertainty of the τ-tagging efficiency is expected to be between 4 and 5% .

12.2 b-tagging tools

12.2.1 General features

Many physics channels such as events containing top quarks, Higgs bosons, or Supersymmetric particles, produce b jets in the final state which need to be distinguished from more copious backgrounds containing only light flavoured jets. The top quark, for example, decays almost exclusively into a W boson and a b-quark, and similarly, for a low mass Standard Model Higgs boson ($m_H \leq 135\ \mathrm{GeV}/c^2$) the dominant decay is to a pair of b quarks, $H \to b\bar{b}$.

Inclusive tagging of b jets, as opposed to lighter flavour jets, mainly relies upon relatively distinct properties of b-hadrons such as large proper lifetime ($\tau \approx 1.5$ ps, $c\tau \approx 450\ \mu\mathrm{m}$), large mass, decays to final states with high charged track multiplicities (on average 5 charged tracks), relatively large semileptonic branching ratios (in about 20% of the cases, b hadrons decay into muons or electrons), and a hard fragmentation function.

This chapter describes the algorithms currently available in CMS for b-tagging. The performance of the algorithms is evaluated in terms of b-tagging efficiency versus background rejection for various simulated data samples. If not explicitly stated otherwise, studies are done for a complete and perfectly aligned detector. Effects of an imperfectly aligned detector are discussed in Section 12.2.7.

Algorithms for b tagging can be applied both offline (i.e. for an end-user analysis) and at the level of the High-Level Trigger (HLT) [8]. In the latter case, CPU execution time is a critical issue and so fully reconstructed input objects cannot always be used.

12.2.1.1 Reconstructed objects used as input

The b-tagging algorithms rely upon the reconstruction of lower level physics objects. At the LHC, b-tagging typically will be applied to jets and so accurate jet reconstruction algorithms are necessary. The jet direction is typically taken to approximate the original flight path of the b-hadron, although more sophisticated inputs can be used in some algorithms to improve upon the b-hadron direction resolution. The jet reconstruction algorithms available in CMS are described in Chapter 11. For the studies described in this section, an iterative cone algorithm with a cone size of 0.5 has been used. A jet energy calibration with correction factors obtained from the Monte Carlo simulation has been applied.

Most of the b-hadron properties used for b-tagging are exploited using charged particle tracks because only tracking detectors offer the spatial resolution needed to detect, for example, the significant decay length of b-hadrons. Efficient track reconstruction, and in particular precise spatial reconstruction close to the interaction point, is thus the key ingredient for almost all b-tagging algorithms. Track finding is performed using the Kalman filter technique that is described in Section 6.4. The following track selection cuts are common to all b-tagging algorithms described here (algorithm specific cuts are listed in the corresponding sections):

- At least 8 reconstructed hits in total (pixel and silicon strip detectors).
- At least 2 reconstructed hits in the pixel detectors.
- Transverse momentum $p_{\mathrm{T}} > 1$ GeV/c.
- χ^2/ndf of the track fit < 10.
- Transverse impact parameter with respect to the reconstructed primary vertex $<$ 2 mm to reject charged particle tracks having their origin from sources showing much larger displacement from the primary vertex (e.g. V^0 decays, photon conversions and nuclear interactions in the beampipe or the first layers of the pixel detector). To first order, the impact parameter is invariant under boosts of the b-hadron.

For b tagging, the track measurement precision close to the interaction point as given by the impact parameter resolution what is most relevant. The track impact parameter resolution for charged particle tracks is shown in Figure 6.16. For b tagging, a physics motivated lifetime-based definition of the sign of track impact parameters is used: the impact parameter is considered positive if the track is reconstructed to originate downstream from the primary vertex with respect to the jet direction, negative otherwise. The association of tracks to jets is performed via a cone based $\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ distance criterion (with ΔR(jet-track) between 0.3 and 0.4 as a typical cut).

To measure a flight path, or to detect displaced tracks which do not have their origin at the primary event vertex, the precise reconstruction of the coordinates of the primary event vertex is crucial. The b-tagging algorithms described later in this section use the primary event vertex reconstructed globally in the event (Sections 6.5.3 and 6.5.4). The resolution depends on the event topology (e.g. the number of charged particles at the primary vertex) and is typically between 10–40 μm in the transverse plane and 15–50 μm in the z-direction.

Details can be found in Sections 6.5.3.3 and 6.5.4.4.

12.2.1.2 Definitions

The efficiency ϵ_q to tag a certain flavour of jet as a b jet (b-tagging efficiency for b jets, misidentification or "mistag" efficiency for non-b jets) is defined as:

$$\epsilon_q = \frac{\text{Number of jets of flavour } q \text{ tagged as } b}{\text{Number of jets of flavour } q}. \tag{12.2}$$

In the simulation studies reported here, the true flavour of a reconstructed jet has been determined by analyzing the parton content in a cone around the jet direction, where the assignment of a parton flavour to the jet follows a physics-based definition. A reconstructed jet is matched to the initial parton from the primary process (e.g. in QCD jet events these initial partons are the ones produced in the hard interaction, in $t\bar{t}$ events the initial partons are the quarks from the top decays and hadronic W decays) if it is within a cone of radius $\Delta R < 0.3$. Gluon jets and quark jets with b- or c-quarks from gluon splitting among their shower products thus get the flavour of the gluon or quark. To ensure an unambiguous assignment, jets are rejected if more than one initial parton fulfills this requirement.

One can also follow an "algorithmic" definition analysing the final state partons after the shower evolution and assigning the parton flavour that most likely determines the structure of a jet. Here, a jet originating from a gluon that has split into a $b\bar{b}$ pair would get the flavour of the b-quark assigned, because from the algorithm point of view the jet looks in many aspects as a genuine b-jet. The main difference in these definitions is for gluon jets, where the non-negligible splitting rates into $b\bar{b}$ and $c\bar{c}$ pairs impose a serious limitation for b-tagging (depending on the energy of the gluon, these splitting rates are typically at the level of a few percent).

For these studies, simulated samples of QCD jets have been used. Jets in these samples cover a wide range in transverse momentum (from about 30 GeV/c to 200 GeV/c) and the full geometrical acceptance of the tracking detector in pseudorapidity ($|\eta| < 2.4$). Semileptonic $t\bar{t}$ events have also been studied.

The b-tagging performance for specific physics channels, including effects such as overlapping jets in dense hadronic environments, is presented in detail for several benchmark physics analyses in Volume 2 of this document.

12.2.2 Track impact parameter based tags

The track impact parameters of charged particles as measured in the tracking detector provide a powerful means to determine whether a track originates from the primary interaction vertex or from the decay of a particle that has travelled a significant distance from the primary interaction vertex. The impact parameter can be computed in the transverse plane or in three dimensions. Three-dimensional impact parameters yield better overall b-tagging performance and so will be used in what follows. Impact parameters are signed in such a way that a positive impact parameter is consistent with a track from a decay of a particle at a point with some positive displacement from the primary vertex along the direction of the jet. On the contrary, a negative impact parameter corresponds to a decay point "behind" the primary vertex. Negative impact parameters arise most frequently as a result of finite

detector resolution which results in mismeasured track parameters. Other sources of large negative impact parameters include badly reconstructed primary vertex positions, pattern recognition failures that result in fake tracks, and multiple scattering of particles in material at small radii such as the beam pipe or first pixel layer. The efficiency for b-tagging jets which do not originate from c- or b-hadron decays can therefore be measured in data by means of the negative side of the distribution of the track impact parameter significance (the ratio of the track impact parameter to its uncertainty), after removal of charged tracks from reconstructed long-lived V^0's (e.g. K_s) are removed. Two different methods that exploit track impact parameters to tag b-jets are described in the following.

12.2.2.1 Track counting b-tagging

The track counting b-tagging method is sufficiently fast and robust that it can be applied both at HLT as well as offline. The method works as follows. For each selected track in a jet, the impact parameter significance is computed. The jet is tagged if the number of tracks with an impact parameter significance exceeding a given cut is greater than a value which can be optimised for various purposes as discussed below.

Tracks are selected for use in the tagging algorithm if, in addition to the common selection described above, the distance of closest approach of the track to the jet axis is less than 0.07 cm.

Tracks are ordered in decreasing signed impact parameter significance. The significance of the track impact parameter of the N^{th} track thus serves as the discriminator for this algorithm. Figure 12.26 shows the distribution of this variable for $N = 2$.

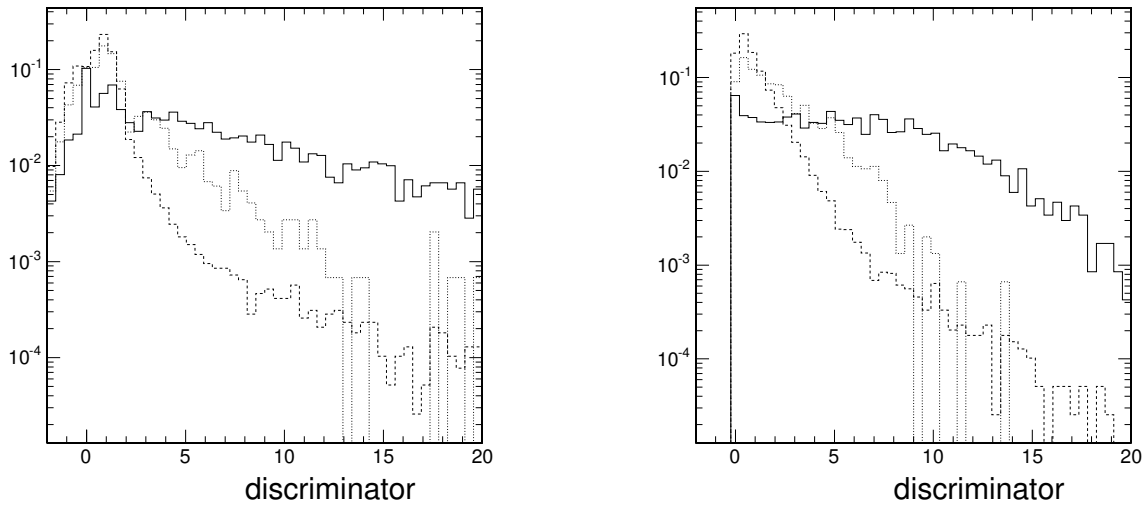


Figure 12.26: Distribution of the discriminator for the track counting algorithm ($N = 2$) (left) and the track probability algorithm (right) for b-jets (solid), c-jets (dotted) and $udsg$-jets (dashed).

12.2.2.2 Probability b-tagging

The probability method is also based on track impact parameters: it entails computing the compatibility of a set of tracks with having come from the primary vertex [294]. The jet

probability estimation is performed through several steps which are described in detail in this section: for each track, the probability of coming from the primary vertex is computed and these probabilities are combined together to provide the jet probability. The track probability is computed using, as calibration, the negative part of the distribution of the impact parameter significance. The negative part is used for this purpose since it is made mainly of primary vertex tracks, as noted above. This approach has the advantage of allowing the detection of the largest range of decay topologies in an inclusive way.

Tracks with negative impact parameter are used to extract the resolution function R , which plays the role of the impact parameter significance distribution. This function is needed to extract the probability $P(x)$ of a track having impact parameter significance x to come from the primary vertex.

In order to take into account the tails of the resolution function and avoid ad-hoc parameterizations, the resolution function is histogrammed with sufficient precision and the probability computed by numerically integrating the normalized histogram.

The track reconstruction quality is related to the momentum and the number of hits in the different types of detectors involved: the number of pixel hits and the position of the innermost pixel layer are crucial for the impact parameter measurement.

The (signed) probability for a track to have come from the primary vertex is defined as the integral of the resolution function multiplied by the sign of the significance S:

$$P_{tr}(S) = \mathrm{sgn}(S) \cdot \int_{|S|}^{\infty} R(x)dx \tag{12.3}$$

With this definition, the distribution of P_{tr} is flat between -1 and 1 for tracks coming from the primary vertex, since the impact parameter significance distribution for the vast majority of tracks from the primary vertex is described by the resolution function (apart from some residual lifetime effects). It is positive and concentrated near 0 for tracks with large impact parameter significance.

The jet probability for a jet containing N tracks is defined as the confidence level for a group of N tracks, none of which originate from the decay of a long-lived particle, would produce the observed, or any less likely value of track probability. The expression of this probability is the following:

$$P_{jet} = \prod \cdot \sum_{j=0}^{N-1} \frac{(-\ln \prod)^j}{j!} \tag{12.4}$$

where

$$\prod = \prod_{i=0}^{N} \hat{P}_i. \tag{12.5}$$

$\hat{P}_i$ is introduced in order to also take into account negative track impact parameters, and is defined as $\hat{P}_i = P/2$ for positive, and $\hat{P}_i = 1 + P/2$ for negative impact parameters tracks.

Figure 12.26 shows the distribution of $-\log P_{jet}$ for b- and $udsg$-jets using three-dimensional impact parameters. The distribution is close to 0 for light quark jets and has a wider distribution for b-jets, providing good discrimination.

12.2.3 Combined secondary vertex tag

12.2.3.1 Algorithm description

The combined b-tagging algorithm is based on the reconstruction of the secondary decay vertex of the weakly decaying b-hadron. Different topological and kinematic vertex variables are combined together with track impact parameter significances into a discriminating variable to distinguish b-quark jets from non-b jets. The main components of the algorithm and its performance are described in the following. A detailed description of the algorithm and its performance can be found in [295].

Secondary vertices are reconstructed in an inclusive way inside the jet using the Trimmed Kalman Vertex Finder described in Section 6.5.1. This algorithm starts with all tracks and successively rejects outliers which then are used to reconstruct additional vertices. Since the algorithm uses not only the presence of a secondary vertex, but also topological and kinematic variables related to the vertex, it is desirable to reconstruct the decay vertex as completely as possible, to increase the discriminating power of these topological and kinematic variables. The performance of the inclusive secondary vertex finding in jets is described in Section 6.5.6.

The following cuts are applied to the resulting vertices to select secondary vertex candidates coming from b-hadron decays:

- The distance from the primary vertex to the secondary vertex in the transverse plane has to exceed 100 μm and must not exceed 2.5 cm.
- The distance from the primary vertex to the secondary vertex in the transverse plane divided by its error has to be greater than 3: $L_t/\sigma_{L_t} > 3$.
- The invariant mass of charged particles associated to the vertex must not exceed 6.5 GeV/c^2.
- The vertex must not be compatible with a V^0 decay.

Based on the result of the secondary vertex reconstruction and selection, 3 vertex categories are defined:

1. **"RecoVertex"**: At least one secondary vertex candidate has been reconstructed and passed the selection criteria. All tracks from all accepted vertices are used for the computation of the vertex related variables if there is more than one accepted secondary vertex.
2. **"PseudoVertex"**: If no reconstructed secondary vertex candidate has been found, a so-called PseudoVertex is created from charged particle tracks not being compatible with the primary event vertex, having a signed transverse impact parameter significance greater than 2, if at least 2 such tracks are present in the jet.
3. **"NoVertex"**: If neither 1) or 2) are fulfilled.

The distribution of the categories for the different jet flavours are shown in Figure 12.27. It can be seen that the presence of a secondary vertex alone is already discriminating between b-quark jets and other jets.

All accepted track impact parameter significances enter into the final discriminator. In order to improve the suppression of charm quark jets, the following quantity is defined in addition: The tracks are sorted in decreasing order of their impact parameter significances. The invariant mass is computed for tracks 1 to n. If this mass exceeds a threshold value related to the mass of charm hadrons, the impact parameter significance of the track moving the n-track mass above this threshold is added to the discriminator. The threshold value is set to be 1.5 GeV/c^2. It is lower than typical charm hadron masses because neutral particles are missed and not all charged particles from the decay are reconstructed and accepted.

12.2.3.2 Input variables and combined discriminator

To get optimal performance, several topological and kinematic variables are combined. The selection of variables entering into the combination depends on the vertex category. Category 1, where a reconstructed secondary vertex has been found, is the most powerful one. For this category, the following variables are defined:

- The invariant mass of charged particles associated to the secondary vertex (shown in Figure 12.27).
- The multiplicity of charged particles associated to the secondary vertex.
- The distance from the primary vertex to the secondary vertex in the transverse plane divided by its error.
- The energy of charged particles associated to the secondary vertex divided by the energy of all charged particles associated to the jet.
- The rapidities of charged particle tracks associated to the secondary vertex with respect to the jet direction: $y = \frac{1}{2} \cdot \ln(\frac{E+p_\parallel}{E-p_\parallel})$. This variable enters for n tracks, where n is the secondary vertex multiplicity.
- The track impact parameter significances in the transverse plane.
- The track impact parameter significance of the first track exceeding the charm threshold as described above.

For the other categories a subset of these variables is used (see [295] for details).

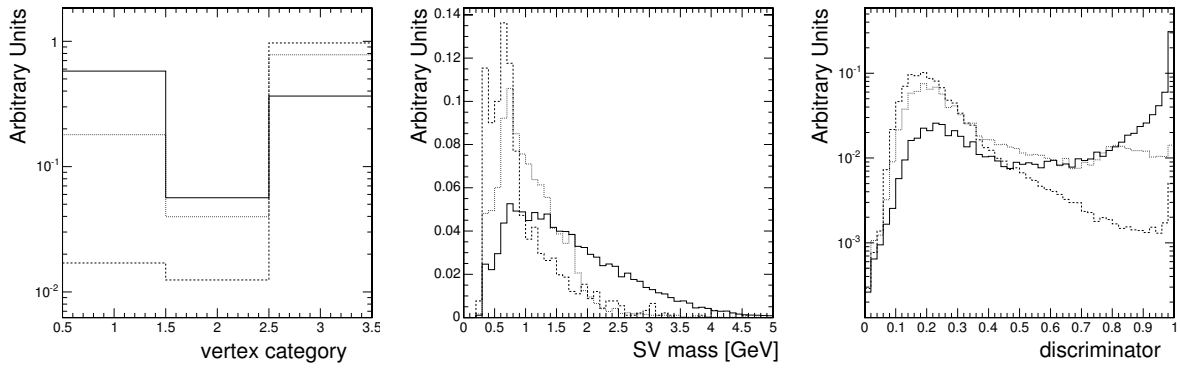


Figure 12.27: Distribution of the vertex category (left), reconstructed secondary vertex mass (middle) and combined discriminator (right) for b-jets (solid), c-jets (dotted) and uds-jets (dashed, $udsg$-jets for the rightmost plot) as obtained from a sample of QCD events for jets with underlying parton momenta between 50 and 80 GeV/c.

The variables described above are combined into a single discriminating variable using a Likelihood ratio technique. Since for most of the variables c-jets and $udsg$-jets look significantly different, the Likelihood ratio contains 2 parts for the discrimination against these 2 backgrounds. With the quantities

$$\mathcal{L}^{b,c,q} = f^{b,c,q}(\alpha) \times \prod_i f^{b,c,q}_\alpha(x_i)$$

the variable d is defined as:

$$d = f_{BG}(c) \times \frac{\mathcal{L}^b}{\mathcal{L}^b+\mathcal{L}^c} + f_{BG}(q) \times \frac{\mathcal{L}^b}{\mathcal{L}^b+\mathcal{L}^q}.$$

where α denotes the vertex category ($\alpha = 1, 2, 3$) as defined above, x_i are the individual variables, q stands for u,d,s-quark jets and gluon jets, $f_{BG}(c, q)$ is the expected prior for the c- and q-content in non-b jets ($f_{BG}(c) + f_{BG}(q) = 1$), $f^{b,c,q}(\alpha)$ is the probability for flavour b, c, and q to fall into category α and $f^{b,c,q}_\alpha(x_i)$ is the probability density function for variable x_i for category α and flavour b, c , q. The quantity d is used as final discriminator. The distribution for the discriminator is shown in Figure 12.27

The probability density functions are extracted from simulated events. They depend on the transverse jet energy and pseudorapidity.

12.2.4 Soft lepton tags

12.2.4.1 Algorithm description

The soft lepton b-tagging algorithm is based on the relatively high inclusive b-quark decay branching ratio to electrons and muons which is about 19% for each lepton family if both direct and cascade decays are taken into account. The key element that is required in order to take advantage of this situation is the identification of leptons among the tracks associated to each jet. To increase the purity of the selection, additional cuts are applied to parameters associated with these leptons. The main components of the algorithm and its performance are described in the following. A detailed description can be found in [296].

12.2.4.2 Electron identification

Electron identification is based on the track extrapolation from the tracker to the calorimeters and a detailed analysis of the energy deposits in the region around the track in both ECAL and HCAL. Reconstructed tracks satisfying the criteria $p_{\mathrm{T}}^{\mathrm{trk}} > 2\ \mathrm{GeV}/c$ and $|\eta^{\mathrm{trk}}| < 1.2$ are considered. The extrapolated tracks are matched with ECAL clusters. The information from the extrapolated tracks and matched ECAL clusters is used for electron candidate selection. The discriminating variables are listed in Table 12.8. These variables are combined into a single discriminating variable using a neural network technique. The training of a neural network has been performed using these variables on a signal sample consisting of 87 000 fully leptonic $t\bar{t}$ events and 91 000 semileptonic $t\bar{t}$ events, and on a background sample consisting of QCD dijets with a p_{T} in the hard process from 50 to 230 GeV/c. Efficiency and misidentification rates for the current implementation of the soft non-isolated electron identification are shown in Figure 12.28.

Table 12.8: List of variables used to identify electrons, where E_{seed}, $E_{N\times N}$ and E_{ECAL} are the energy deposit in the cluster seed, the maximal energy deposit in a square of $N \times N$ ECAL crystals around the seed and total energy of the cluster. p_{trk} is the extrapolated track momentum upon entering the ECAL. E_{HCAL} is the sum of the energy deposits in the HCAL towers next to the ECAL cluster.

Measurement	variables
covariance of cluster energy distribution	$\sigma_{\eta\eta}$, $\sigma_{\eta\phi}$, $\sigma_{\phi\phi}$.
repartition of cluster energy	$\frac{E_{\mathrm{seed}}}{E_{2\times2}}$, $\frac{E_{\mathrm{seed}}}{E_{3\times3}}$, $\frac{E_{3\times3}-E_{\mathrm{seed}}}{E_{5\times5}-E_{\mathrm{seed}}}$, $\frac{E_{2\times2}}{E_{5\times5}}$, $\frac{E_{\mathrm{ECAL}}}{E_{\mathrm{HCAL}}+E_{\mathrm{ECAL}}}$
cluster energy and track momentum ratio	$\frac{E_{\mathrm{ECAL}}-p_{\mathrm{trk}}}{E_{\mathrm{ECAL}}+p_{\mathrm{trk}}}$

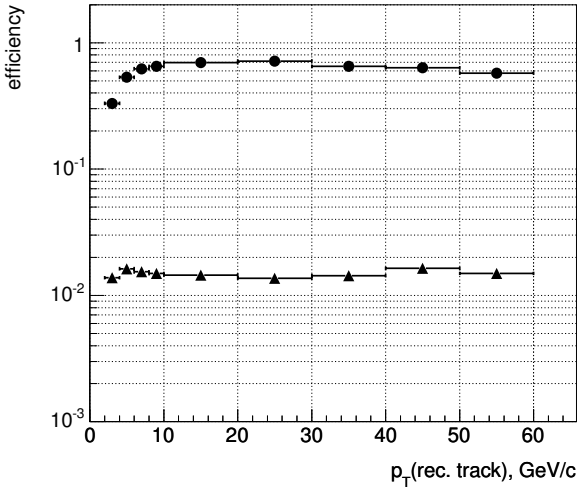


Figure 12.28: Performance of the soft non-isolated electron identification algorithm: efficiency (circles) and misidentification rate (triangles).

12.2.4.3 Muon identification

Muon identification can be performed in two fundamentally different ways:

- Looking for reconstructed muons (9.1) in a region around the jet axis, which is assumed to be the best approximation to the original quark direction.
- Extrapolating the tracks associated to each jet outward through the calorimeters and the magnet to the muon chambers, looking for compatible tracks, track segments, or even energy deposits in the electromagnetic and hadron calorimeters (9.2).

The current analysis is performed following the first approach. However, to allow for both types of identification to be used in a consistent framework, a trivial muon identification is implemented, matching charged tracks inside the jet to muons reconstructed using the algorithm described in Section 9.1. Matching is done requiring at least 70% of a track's reconstructed hits to be in common with the muon track.

Muon reconstruction starting from the muon system is limited to muons with a transverse momentum $p_{\mathrm{T}} > 3.5\ \mathrm{GeV}/c$. With the development of dedicated muon identification, the

second approach will be exploited with the aim to increase the muon identification efficiency.

12.2.4.4 Lepton b-tagging

Jet b-tagging purity is increased through the use of an adaptive nonlinear discriminant function, i.e. a feed-forward neural network, applied to some of the lepton and jet parameters. The desired trade-off between efficiency and purity can be achieved by varying the cut on the network's output. The current implementation gives higher purity, with correspondingly lower efficiency, as the cut value is increased in the range $[0.0, 1.0]$.

The parameters currently used for the lepton and the jet are:

- The lepton transverse momentum relative to the lepton-excluded jet axis, calculated as the p_{T}-weighted average of the other charged tracks inside the jet.
- The ratio of the lepton momentum and the jet energy.
- The significance of the three dimensional lepton impact parameter (Section 12.2.2).
- The pseudoangular distance in the η-ϕ plane between the lepton and the lepton-excluded jet axis.

Additionally, different neural networks or neural network parameters can be used for the different jet energy and pseudorapidity regions and to take lepton flavour into account.

The neural network for electron b-tagging has been trained on the QCD and $t\bar{t}$ samples described in Section 12.2.4.2.

Further improvements in the performance of the neural network can be achieved by taking into account the various possible origins of the lepton:

- The prompt decay of a bottom hadron.
- The prompt decay of a charm hadron, originating from a bottom hadron.
- The prompt decay of a primary charm hadron.
- The decay of light hadrons coming either from the primary hadronization or from the decay of heavier hadrons.

The neural network for soft muon b tagging has been trained on flavour-enriched dijet samples to give an approximation of the *a posteriori* Bayesian probability for a muon to belong to each of these classes [297]. Since the *a priori* class probabilities, i.e., the relative abundance of bottom, charm, and light jets, in actual applications will generally be quite different from those in the training samples, the neural network outputs should be scaled and normalized accordingly.

The use of Bayesian *a posteriori* probabilities has three advantages over simpler nonlinear discriminants. First, assigning a pattern to the most probable class minimizes misidentification probabilities [298]. Second, the use of (approximate) probabilities makes it easier to combine the soft lepton tag with other b-tagging algorithms. Last, having a separate output for each class makes it possible to apply different trade-offs between the bottom jet tagging efficiency, charm jet mistag rate, and light jet mistag rate.

A parallel study is being performed, using the same tagging variables except impact parameter significance. Such an algorithm will be directly applicable to a staging scenario for the pixel detectors, and is expected to be more robust against detector misalignment. It will

also be statistically independent from the other b-tagging algorithm, and thus more easily combined with them.

12.2.5 Performance of b-tagging algorithms

In this section, the offline performance of the b-tagging algorithms described in previous sections is presented. All event samples used for these studies have been simulated with

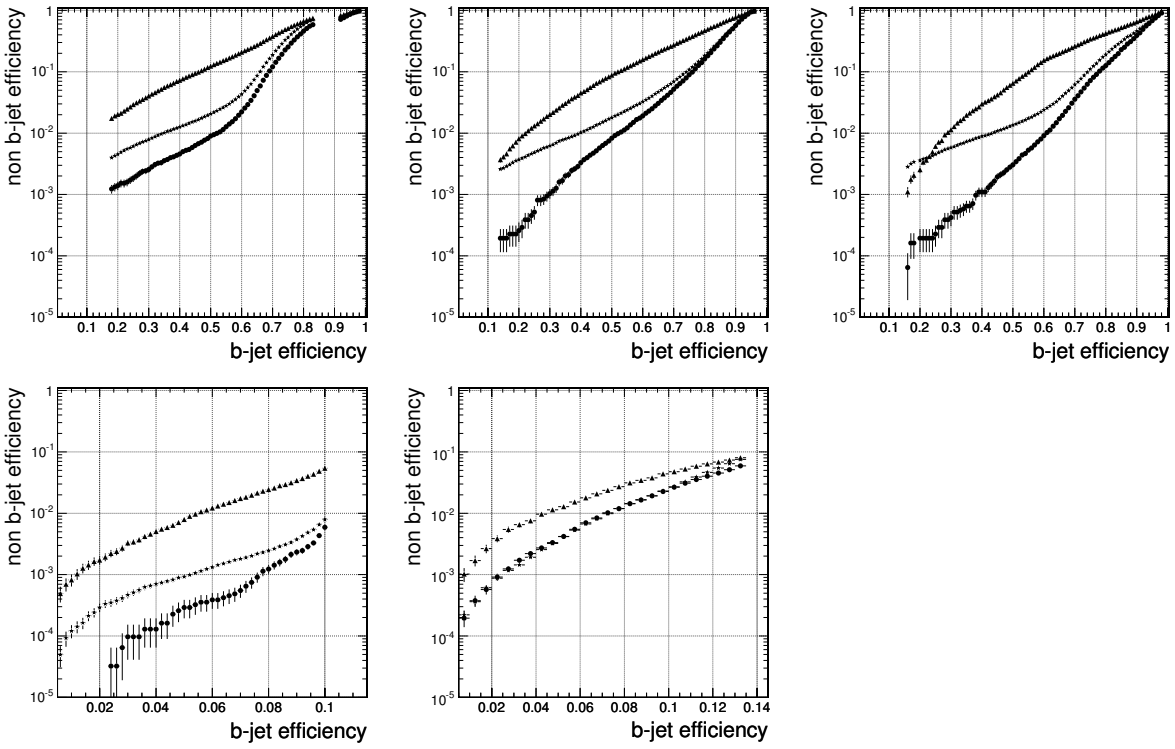


Figure 12.29: Non-b jet mistagging efficiency versus b-jet tagging efficiency for c-jets (triangles), uds-jets (circles) and g-jets (stars) for the track counting algorithm (top left), the track probability algorithm (top middle) and the combined secondary vertex algorithm (top right). For these algorithms, jets in a QCD sample with transverse jet momenta between $50\ \mathrm{GeV}/c$ and $80\ \mathrm{GeV}/c$ in the barrel region ($|\eta| < 1.4$) of the detector have been used. For the soft muon algorithm (bottom left) and the soft electron algorithm (bottom middle), fully leptonic and semileptonic $t\bar{t}$ events as well as QCD jets with transverse momenta between $50\ \mathrm{GeV}/c$ and $230\ \mathrm{GeV}/c$ in the barrel region ($|\eta| < 1.4$) of the detector have been used.

pile-up events corresponding to a luminosity of $\mathcal{L} = 2 \times 10^{33}\ \mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

The figures shown in the following have been obtained by scanning the cut on each of the discriminators used by the different algorithms.

Figure 12.29 shows the non-b jet mistag efficiency versus the b-jet tagging efficiency for the lifetime based algorithms (track counting, track probability, combined secondary vertex) and the soft lepton based algorithms. For the lifetime based algorithms, the performance is shown in the barrel part of the detector ($|\eta| < 1.4$) for QCD jets with transverse momenta between $50\ \mathrm{GeV}/c$ and $80\ \mathrm{GeV}/c$. This momentum range is characteristic of many relevant physics processes and is a domain in which good b-tagging performance can be obtained.

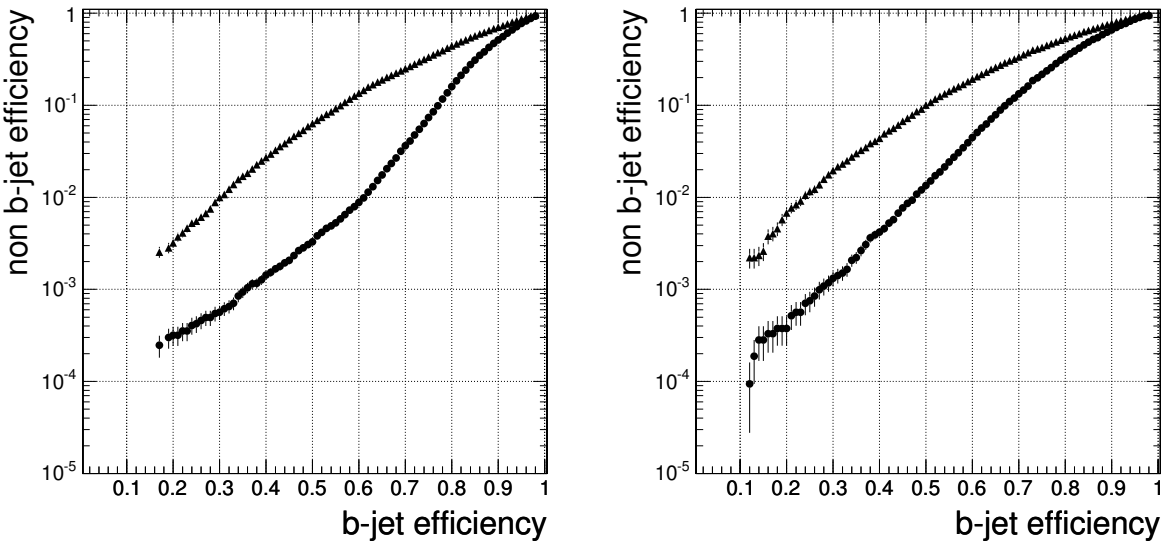


Figure 12.30: Non-b jet mistagging efficiency versus b-jet tagging efficiency for c-jets (triangles) and uds-jets (circles) as obtained in a sample of semi-leptonic $t\bar{t}$ events for jets with transverse momentum larger than $30\ \mathrm{GeV}/c$ in the barrel region ($|\eta| < 1.4$, left) and forward region ($1.4 < |\eta| < 2.4$, right) of the detector for the combined secondary vertex algorithm.

For the soft lepton based algorithms, fully leptonic and semileptonic $t\bar{t}$ events as well as QCD jets with transverse momenta between $50\ \mathrm{GeV}/c$ and $230\ \mathrm{GeV}/c$ in the barrel part of the detector ($|\eta| < 1.4$) have been used. It can be clearly seen that the splitting of gluons

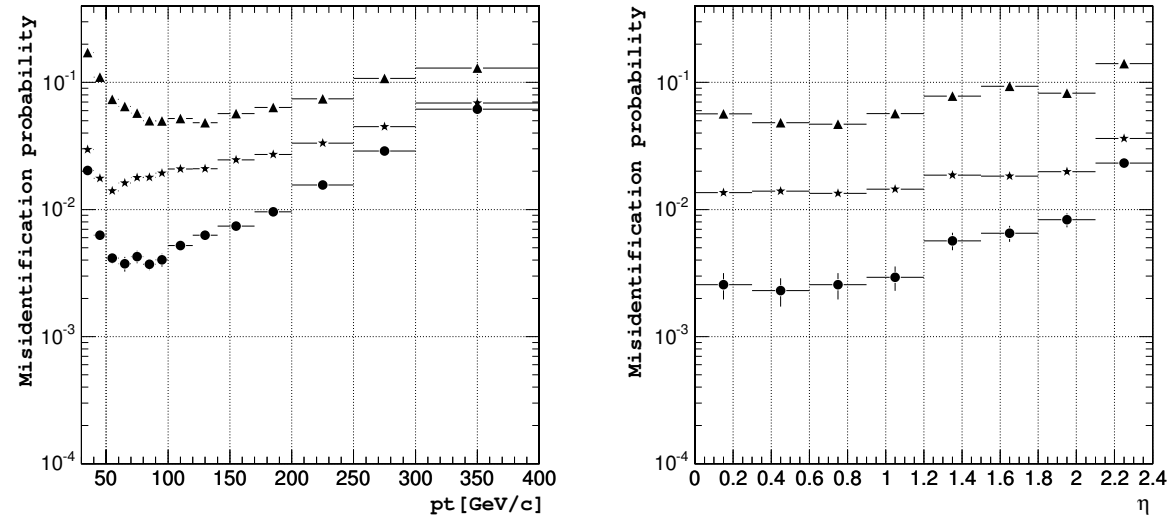


Figure 12.31: The non-b jet mistagging efficiencies for a fixed b-jet tagging efficiency of 0.5 as a function of the jet transverse momentum (left, jets with $|\eta| < 2.4$ have been considered; the transverse momentum shown on the x-axis is the one of the underlying initial parton) and pseudorapidity (right, jets with transverse momenta between $50\ \mathrm{GeV}/c$ and $80\ \mathrm{GeV}/c$ have been considered) for c-jets (triangles), uds-jets (circles) and gluon jets (stars) obtained for jets in an event sample of QCD events for the combined secondary vertex algorithm.

into $b\bar{b}$ and $c\bar{c}$ pairs limits the rejection power as compared with light quark jets (u, d, and

s). For comparison, the performance in a sample of semi-leptonic $t\bar{t}$ events is shown in Figure 12.30 for the combined secondary vertex algorithm for jets with transverse momentum larger than 30 GeV/c. This transverse momentum cut corresponds to a typical requirement in $t\bar{t}$ analyses.

The dependence of the b-tagging performance on jet transverse momentum (transverse momentum of the underlying initial parton) and pseudorapidity are illustrated in Figure 12.31 for the combined secondary vertex algorithm and the soft muon algorithm. Displayed are the non-b jet mistagging efficiencies for a fixed b-tagging efficiency of 0.5. All lifetime based algorithms show a similar behaviour.

The main factors limiting rejection in the case of light quark and gluon jets come from V^0 decays, interactions in the detector material, gluon splitting into $c\bar{c}$ and $b\bar{b}$ pairs and limited detector resolution. In the case of c-jets the rejection is obviously limited by the lifetime of charmed hadrons giving signatures similar to b-hadron decays.

As can be seen in Figure 12.31, the optimum performance is obtained in the central region of the detector and for transverse jet momenta of about 60–90 GeV/c. The performance degrades at larger pseudorapidities mainly because of the increase in material traversed and degraded detector resolution. The slow degradation for larger transverse momenta is caused by an increased splitting rate of gluons to heavy quark pairs, an increased track multiplicity from fragmentation and the increasingly difficult pattern recognition associated with dense jets. The steep fall of performance toward lower transverse momenta is mainly due to increased multiple scattering that results in worse separation of primary and secondary vertices, and more limited jet reconstruction.

12.2.6 HLT b tag

The identification of b-jets in the final state is an important tool during the offline selection for a large range of physics channels—from top physics to MSSM Higgs searches. But also in the online selection, b-tagging may provide a useful handle to reject background events while maintaining a high signal efficiency.

On top of the performance requirements relevant for the offline selection of b-jets, a few additional constraints exist that make the use of b-tagging in the online selection less than straightforward:

- The time constraints in the High-Level Trigger are quite restrictive for the notoriously slow process of track reconstruction. There are several solutions to this problem. An important gain is made by combining the b-tagging information with another requirement that offers fast rejection of a large fraction of events. Furthermore, the development of faster tracking algorithms (pixel-only tracking, regional and conditional tracking) can result in significant time gains.
- In an online environment, robustness and calibration independence of the algorithm become much more important. Therefore, only the simplest algorithms are considered here. Special attention is paid to the beam spot determination.

In the following section, the b-tagging performance of 2 fast algorithms is discussed. A more detailed write-up can be found in reference [299].

12.2.6.1 Algorithm description

In the following subsections 2 fast b-tagging algorithms are described. The 2 algorithms represent 2 different optimizations with respect to performance versus speed. They are envisaged to run in series: the first algorithm (referred to as Level-2.5 from now on) is used to quickly reject those events that definitely do not containing b-jets or whose b-jets are not easily taggable. The second algorithm (labelled Level-3) makes use of the full tracking power of the CMS detector for the final trigger decision.

12.2.6.2 HLT – Level-2.5

This algorithm is optimized to work relatively early in the HLT, where the rate can be as high as 1 kHz. Therefore, the emphasis is on speed, rather than on precision. Full track reconstruction cannot be sustained at this level. Instead, fast track reconstruction is performed within the pixel detector. A detailed description of pixel-only track finding is found in Section 6.4.7 of this document and reference [192].

Due to the granularity and precision of the pixel detector, the pattern recognition with only 3 space points in dense environments of high E_{T} jets is very successful. Importantly, the fake rate is controlled to ≈ 10 %. A tight constraint on the compatibility of the tracks with the longitudinal position of the primary vertex yields an improvement to the level of a few percent. Of course, the efficiency for finding hit triplets in 3 pixel layers depends strongly on the hit finding efficiency of the layers. Pixel layer efficiency loss mechanisms are described in Section 6.1.2.

A drawback of the use of pixel-only tracks is the small lever arm and consequently poor transverse momentum resolution. This causes a severe deterioration of the sensitivity for the transverse impact parameter: the resolution ranges from ≈ 150 μm for 1 GeV/c tracks to an asymptotic value of 80 μm from about 10 GeV/c, a considerable degradation with respect to the results for the full tracker shown in Figures 6.16. The longitudinal impact parameter is much less degraded with respect to tracks reconstructed in the full inner tracker: the resolution ranges from ≈ 120 μm for 1 GeV/c tracks to an asymptotic value of ≈ 70 μm beyond ≈ 10 GeV/c.

The tracks based on triplets of pixel hits are sufficiently pure and accurate that meaningful b-tagging can be performed.

Use of a three-dimensional impact parameter requires the primary vertex of the event to be precisely reconstructed. In the transverse plane, the algorithms rely on the beam spot position to be accurately known (Section 6.5.6.4). Several algorithms for event-by-event determination of the primary vertex based on globally reconstructed pixel-hit triplets are described in Section 6.5.3. For the event topologies of interest for a b-jets trigger—moderate to hard dijet production—the divisive algorithm yields a primary vertex finding and selection efficiency greater than 95 % ($\hat{p}_{\mathrm{T}} > 50$ GeV/c), reaching 99% for $\hat{p}_{\mathrm{T}} > 120$ GeV/c. The resolution for the primary vertex position along the beam line is better than 35 μm for $\hat{p}_{\mathrm{T}} > 50$ GeV/c.

As the timing overhead from the b-tagging algorithm is negligible, the standard b-tagging algorithms can be used at this point.

The track counting algorithm, described in Section 12.2.2.1, offers very good performance. Moreover, the algorithm is simple and does not require calibration (except of course the track error determination).

The number of jets for which the b-tagging is invoked is a free parameter of the algorithm. In the default implementation discussed here, the OR of the 2 leading jets (in E_{T}) is taken. The parameters of the b-tagging algorithm are chosen so that a rejection of a factor 5 is obtained with respect to Level-2.

12.2.6.3 HLT – Level-3

For the events that pass the initial selection based on pixel tracks, tracks are reconstructed in regions-of-interest around jets identified as b-jets in the previous step. At this late stage in the HLT—with a much reduced rate—the full-fledged offline combinatorial Kalman filter track finding algorithm is applied.

The region is constructed by setting limits on the direction with respect to the jet axis ($\Delta\eta < 0.4$, $\Delta\phi < 0.4$), on the minimum transverse momentum ($p_{\mathrm{T}} > 1\ \mathrm{GeV}/c$) and requiring compatibility with the primary vertex determined at Level-2.5. Only seeds within this region are reconstructed. The standard combinatorial Kalman filter track finder then follows the seeds through the tracker volume. Thus, tracks of nearly "offline" quality are obtained. A slight reduction toward the edges of the region—due to the finite resolution of the seeds—is observed.

This step is optimized to achieve a rejection of a factor 4 with respect to the previous level.

12.2.6.4 b-tagging Performance

The performance of the b-tagging algorithms is measured on the 2 leading jets (in an E_{T} ordered list) in a sample of events simulating the hadronic decay of top pairs in low-luminosity running conditions.

For events with true b-quark content, the results of the 2 b-tagging algorithms are extremely correlated: events that do not satisfy the softer requirement at Level-2.5 in general would not have triggered Level-3. For the benchmark channel studied, fully hadronic decays of top pairs, the signal efficiency is virtually unchanged if Level-2.5 is removed.

12.2.6.5 Timing

The time constraint in the HLT is a crucial factor in the choice of the algorithm. The CPU requirement of global track reconstruction in the full tracker is prohibitive. Timing measurements indicate that the CPU load of global pixel triplet reconstruction—as used for the Level-2.5 step and the determination of the primary vertex—is well within the available budget. Full-fledged Kalman-filter based reconstruction of the tracks in a region around a large E_{T} jet takes approximately 300 ms on a 1 GHz CPU, provided the z-position of the primary vertex is constrained.

While it is clear that track reconstruction poses a heavy CPU load, in the case of a b-jets trigger a large gain is obtained by early rejection of events. The b-jets High-Level Trigger is invoked at a rate of 20 kHz. However, the tracker information is not accessed in 98% of these events as the Level-2 jet selection reduces the number of events by a factor 50.

For the majority of the events for which b-tagging has to be performed (400 Hz), only pixel reconstruction is invoked. Only for 1 in 5 events (80 Hz) are tracks reconstructed in the

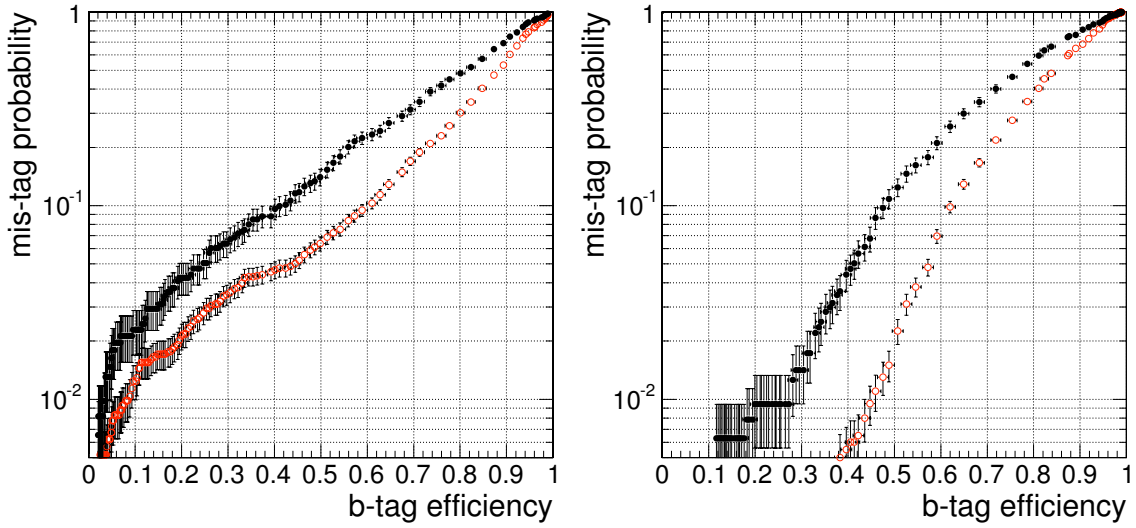


Figure 12.32: Regional (leftmost figure) and pixel-only (rightmost figure) reconstruction b-tagging efficiency and light jet rejection for the track counting algorithm on a sample of hadronically decaying top pairs. The curves are obtained by scanning over the signed 3D IP significance for the N^{th} track in the jet (ordered in order of decreasing significance), where $N = 2$ for pixel-only reconstruction and $N = 3$ for regional reconstruction. The curve for b-jets versus light jets is shown with open markers, while that for b-jets versus c jets has filled markers.

full tracker (and only certain regions). Thus, even though the track reconstruction may be relatively slow, its weight on the time employed for the average event is very small.

12.2.6.6 Performance for different scenarios

The b-tagging performance depends crucially on those measurements that are closest to the interaction vertex. In scenarios where the pixel detector is not present, b-tagging is essentially reduced to the use of leptons.

Moreover, the lifetime based algorithms rely heavily on the correctness of the track parameter error estimates. Provided that the error estimates are reasonably accurate, the b-jet trigger should not be very sensitive to the envisaged misalignment of the pixel detector.

Finally, the impact parameter in the transverse plane is determined with respect to the nominal beam position. The algorithms thus rely on an accurate measurement of the beam spot to be available, even in the HLT environment. This issue is further discussed in Section 6.5.6.4.

12.2.7 Robustness of the performance

The performance studies shown for the different algorithms in the previous sections have been carried out under the assumption of a complete and perfectly aligned detector. However, in reality, the alignment will always be somewhat imperfect and at the beginning of data taking some detector components may be missing. The effect of a missing layer of the pixel detectors will be studied in Volume 2 of this document. The focus here will be on the performance that can be expected for an imperfectly aligned tracking detector. Since the pixel and silicon strip detectors are extremely precise devices with spatial resolutions of about

10 μm for the innermost layers, alignment of tracker modules with comparable precision is particularly challenging.

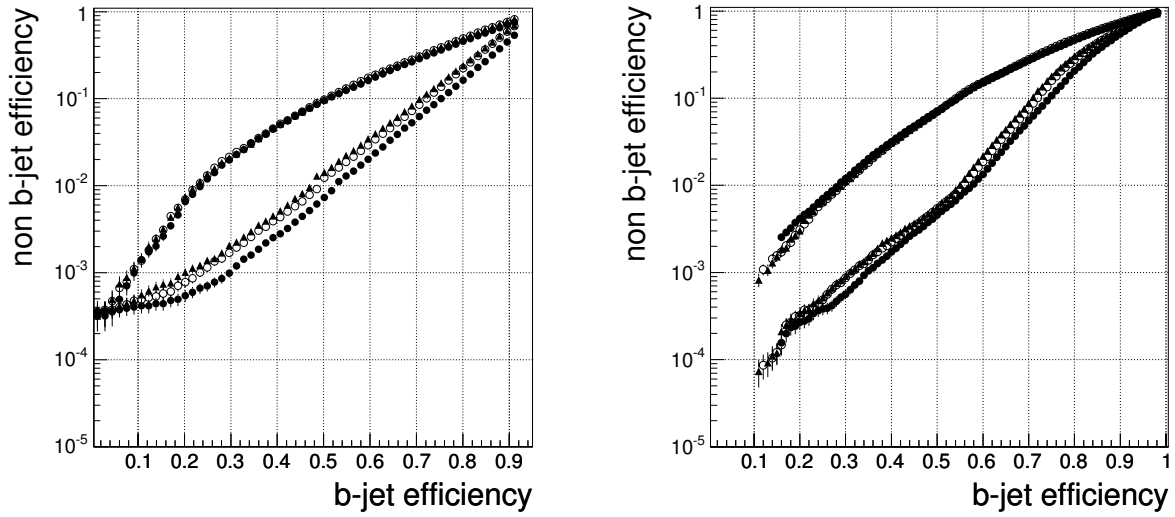


Figure 12.33: Non-b jet mistagging efficiencies versus b-jet tagging efficiency comparing the different misalignment scenarios for the track probability algorithm (left) and the combined secondary vertex algorithm (right). The scenarios considered are: perfect alignment (full circles), long term (open circles) and first data taking (triangles). The upper and lower set of curves correspond to c-jets and uds-jets, respectively. A semi-leptonic $t\bar{t}$ sample was used for these studies.

Two misalignment scenarios (Section 6.6.4) have been defined, a "First Data Taking Scenario" (accuracy as expected from mounting precision, the laser system and some track based alignment for the pixel detectors as expected after an integrated luminosity of less than 1 fb^{-1}) and a "Long Term Scenario" (ultimate precision as expected from the track based alignment after a period of data taking with sufficient integrated luminosity, about 10 fb^{-1}). The assumed accuracy for the Long Term Scenario is typically about a factor 10 better than for the First Data Taking Scenario, apart for the pixel detectors where similar alignment precisions of about 10 μm are assumed for both scenarios. The alignment errors are added in quadrature to the nominal hit errors. Details of the alignment procedures and misalignment scenarios can be found in Section 6.6. Figure 12.33 shows the performance of the track probability algorithm and combined secondary vertex algorithm comparing a perfectly aligned detector and misaligned detectors according to the 2 scenarios described above. For these studies, calibration functions (e.g. as needed for the track probability and the combined secondary vertex algorithm) are same as those used for the perfectly aligned detector. The degradation of the performance for light quark (u, d, and s) and gluon jets is very similar for the two misalignment scenarios. The reason for this behaviour is that impact parameter resolution, which is dominated by the pixel detectors, is much more important than momentum resolution for which the contribution from the strip detectors is also significant. As mentioned above, both misalignment scenarios use similar alignment precision for the pixel detectors. Furthermore, hadrons in jets have typical momenta of a few GeV/c and the track impact parameter error resulting from multiple scattering is bigger than the alignment error assumed here. Therefore, comparing to the scenario of a perfectly aligned detector, the loss in performance is acceptable. It can be seen that there is essentially no effect on the mistagging

efficiency of c-jets if the b-tagging efficiency is kept constant. This can be explained by the fact, that for a large fraction of c-jets tagged as b-jets, the real lifetime of charm hadrons is the origin of the mistag. Variables sensitive to the lifetimes of heavy hadrons thus vary in a very similar way for c- and b-jets when applying the various misalignment scenarios.

12.2.8 Performance studies with data

12.2.8.1 Offline calibration strategies

Due to the high energy and luminosity of the LHC, a completely new method to calibrate b-tagging algorithms can be exploited [300]. Indeed, during the first years of operation, the LHC will produce more than 8 million $t\bar{t}$ pairs each year. With $t \to W^{\pm}b \approx 100\%$, these $t\bar{t}$ events can be used to isolate a highly enriched b-jet sample that can then be exploited to calibrate b-jet identification algorithms.

Two methods to extract a b-jet enriched jet sample in $t\bar{t}$ events are presented. The two methods yield three samples that can be used to measure b-tagging performance. A method based on a likelihood ratio to perform this measurement is presented, along with estimates of the most important systematic uncertainties. The expected precision of the measurement of the b-jet identification performance is estimated and serves as direct input to physics analyses that rely on the b-tagging algorithms considered. [294, 295, 296].

Selection of a b-enriched jet sample
To measure b-tagging efficiencies, a jet sample with an enriched b-jet content needs to be selected. For this purpose, $t\bar{t}$ decay channels in which one or both of the W bosons decay leptonically were investigated with GEANT4 simulated Monte Carlo events. These channels are of particular interest, as the presence of an isolated energetic lepton can be triggered upon with high efficiency. In the *semileptonic decay channel* $t\bar{t} \to b\bar{b}qq'\ell\nu$, events are required to have either an isolated muon or electron, four jets, and missing transverse momentum. To fight the large $W^{\pm}+$ jets background an extra requirement of 1 b-tagged jet is imposed. In the *fully leptonic decay channel* $t\bar{t} \to b\bar{b}e\nu_e\mu\nu_\mu$, 1 muon and 1 electron of opposite charge, 2 jets and missing transverse momentum are required for an event to be selected. These selection criteria suppress the large $Z+$ jets background, and to a lesser extent the WW and ZW backgrounds.

In the case of the semileptonic events, the 12 possible ways of associating jets to partons from the two top quarks were each considered in a kinematic fit that imposed $W^{\pm}$ and top mass constraints [301]. The best jet pairing was chosen by combining the fit results with several other observables into a combined likelihood ratio. In the case of fully leptonic events, the kinematics are underconstrained, but several event observables can still be combined into a similar likelihood function.

In both cases, a cut can be made on the corresponding likelihood ratio to select well reconstructed events. The sample of jets associated with $t \to b\ell\nu$ decays in the semileptonic case, and the jet sample from both jets in the fully leptonic case, can be enriched in b-content with such a cut. In Figure 12.34 the b-jet purity for the different jet samples is shown as a function of the cut on the likelihood ratios.

Systematic uncertainties in the determination of sample purity
In order to ascertain the degree of confidence in the Monte Carlo b-jet sample purity, various

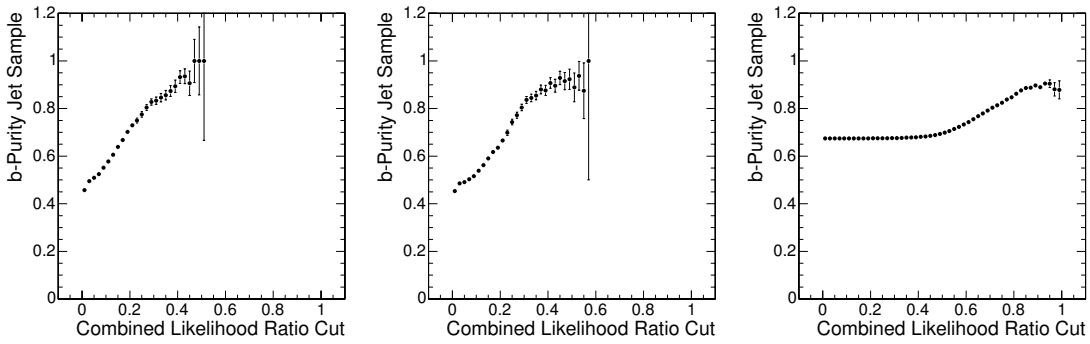


Figure 12.34: The b-jet purity of the different jet samples (semileptonic with muon, semileptonic with electron and fully leptonic), as a function of the cut on the corresponding combined likelihood ratio.

systematic uncertainties in the purity were evaluated [302, 300]. For both decay channels, the most relevant contributions come from uncertainties in the Monte Carlo modelling of initial and final state gluon radiation and uncertainties in the signal and background cross sections. In semileptonic events, due to the demand for a b-tagged jet in the event selection, the uncertainty on the b-tag efficiency itself results in an additional small systematic uncertainty in the sample purity. In Figure 12.35 the various contributions to the systematic uncertainties and their combined total are shown for the 3 samples (semileptonic decays with muons, semileptonic decays with electrons, and fully leptonic decays) as a function of the cut on the combined likelihood ratio. Uncertainties from the parton distribution functions, underlying event, pile-up, luminosity, light- and b-quark fragmentation do not contribute significantly. Additionally, for semileptonic events, the jet energy scale and top mass uncertainty were checked and found to be negligible.

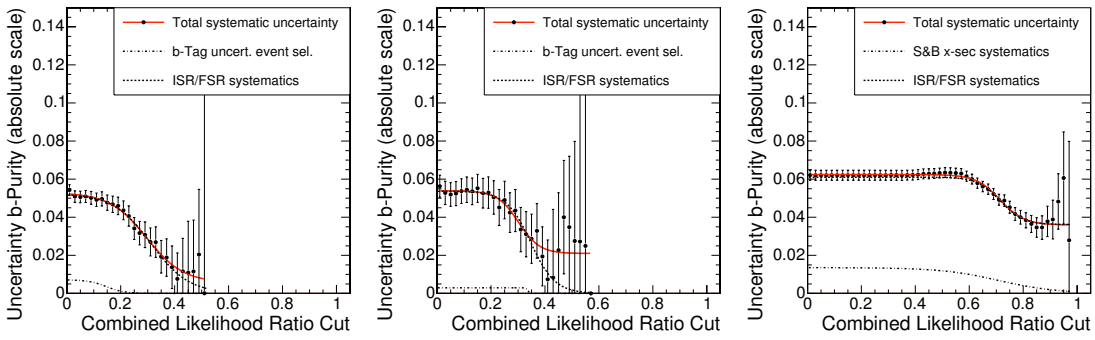


Figure 12.35: Systematic uncertainties on the b-purity of the different jet samples (semileptonic with muon, semileptonic with electron and fully leptonic), as a function of the cut on the corresponding combined likelihood ratio.

Measurement of b-tagging efficiencies in data

For both jet samples the same technique can be used to measure the b-tagging efficiencies. Defining x_b as the b-purity of the jet sample, previously estimated from Monte Carlo, and x_{tag} as the measured fraction of jets tagged by the b-tagging algorithm, then the b-tagging efficiency is

$$\epsilon_b = \left[x_{\text{tag}} - \epsilon_{\text{o}}(1 - x_b)\right] / x_b,$$

where $\epsilon_{\rm o}$ is the mistag rate for non-b jets in the sample. This mistag rate is determined from the simulated data here but can of course be measured with real data when it becomes available.

It is possible to perform the measurement of ϵ_b for different cuts on the likelihood ratio. As the value chosen for the likelihood ratio cut increases, the systematic uncertainty falls but is offset by a rise in the statistical uncertainty. Hence, there exists a point of minimum total uncertainty and this is the point where the estimation of ϵ_b is performed for a specific sample.

Sample independent calibration
Identification efficiencies for b jets depend on many sample-specific jet parameter distributions, of which the most important are the transverse energy $E_{\rm T}$ and the pseudorapidity η of the jet. To make the calibration of the b-tagging algorithms sample-independent, measurements of efficiencies were performed as a function of these 2 parameters.

In Figure 12.36 the combined result from the 3 samples (semileptonic decays with muons, semileptonic decays with electrons, and fully leptonic decays) of this $E_{\rm T}$-dependent measurement of ϵ_b is shown for jets in the barrel ($|\eta| < 1.5$) and endcap ($|\eta| > 1.5$) regions respectively, for $1\ \mathrm{fb}^{-1}$ of data. The b-tagging was performed using the combined secondary vertex discriminant (Section 12.2.3). The working point chosen for these studies corresponds to an average b-tagging efficiency of about 50%, and mistag probabilities for light quark jets of less than 1% and for charm quark jets of about 6%. Figure 12.37 shows the corresponding absolute and relative uncertainty expectations, together with an extrapolation, assuming unchanged systematic uncertainties, to $10\ \mathrm{fb}^{-1}$ of integrated luminosity. The relative uncertainty on the b-efficiency measurement is expected to be about 6% (4%) in the barrel and 10% (5%) in the endcaps for $1\ \mathrm{fb}^{-1}$ ($10\ \mathrm{fb}^{-1}$) of integrated luminosity.

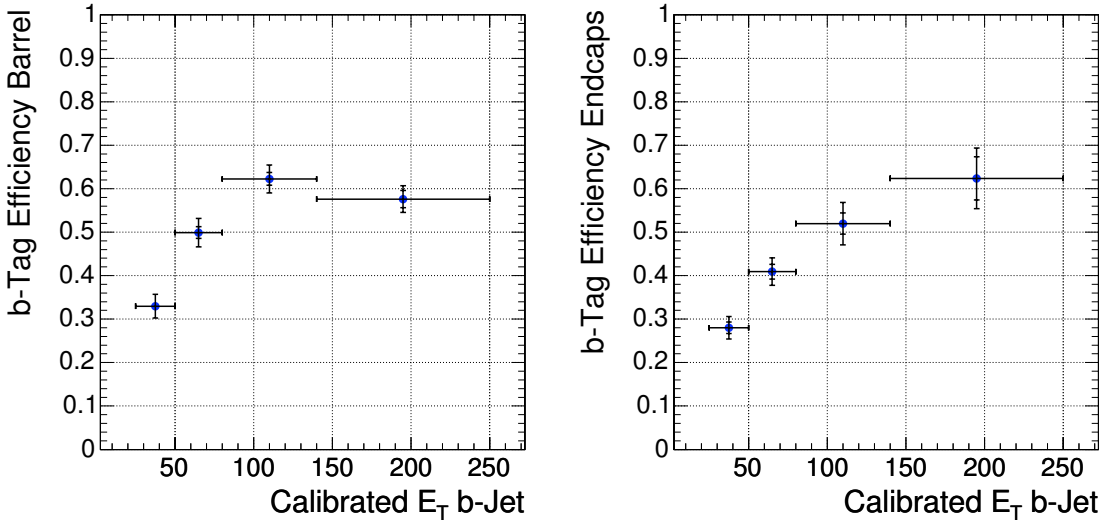


Figure 12.36: b-tagging efficiency measurement, as a function of the jet $E_{\rm T}$ for jets in the barrel ($|\eta| < 1.5$) and endcaps ($|\eta| > 1.5$).

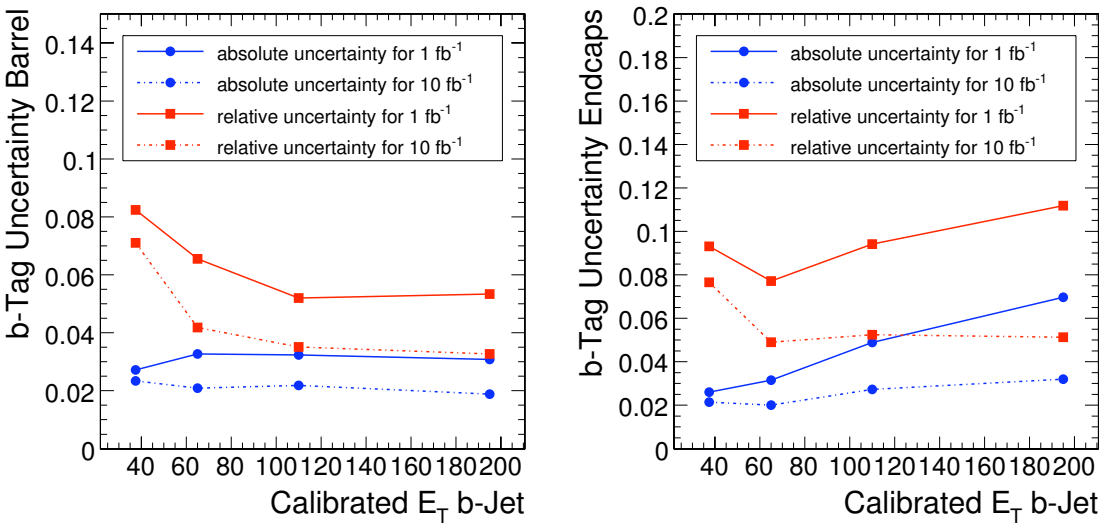


Figure 12.37: Uncertainty on the b-tag efficiency, as a function of the jet E_{T} for jets in the barrel ($|\eta| < 1.5$) and endcaps ($|\eta| > 1.5$) for 1 fb^{-1} and 10 fb^{-1} of integrated luminosity.

12.2.8.2 Other possibilities

An alternative method to measure b-tagging efficiencies makes use of the ratio of single and double tagged events in a sample where 2 b-jets are present in each event. Fully leptonic dec$t\bar{t}$ events can also be used with this technique. Other possibilities include the use of the fully leptonic $e + e$ and $\mu + \mu$ samples with the requirement of at least one b-tagged jet for greater background rejection, jet samples from $Z + jets$, $Zb\bar{b}$, exploiting almost uncorrelated taggers, etc. Finally, for c-jet misidentification rates, a simultaneous determination with b-tagging efficiencies can be envisaged in semileptonic $t\bar{t}$ events.

Glossary

ADC Analog to Digital Converter
AdS Anti de Sitter space
AFEB Anode Front-End Board, CSC system
ALCT Anode Local Charged Track trigger primitive, CSC system
ALEPH An experiment at LEP
ALICE A Large Ion Collider Experiment at the LHC
ALPGEN Monte Carlo event generator for multi-parton processes in hadronic collisions
AOD Analysis Object Data - a compact event format for physics analysis
APD Avalanche Photo-Diode
API Application Programming Interface
APV Analogue Pipeline (Voltage mode), front-end readout chip of Tracker
ASIC Application Specific Integrated Circuit
ATLAS A Toroidal LHC ApparatuS experiment

BLOB Binary Large OBject
BMU Barrel Muon system
BR Branching Ratio
BTI Bunch and Track Identifier trigger primitive, DT system
BU Builder Unit
BX Bunch Crossing
BXN Bunch Crossing Number

CAD Computer-Aided Design
CASTOR Calorimeter in the forward region of CMS
CC Cosmic Challenge
CCS Clock and Control System
CDF Collider Detector Facility experiment at the FNAL Tevatron
CFEB Cathode Front-End Board of the CSC system
CLCT Cathode Local Charged Track trigger primitive, CSC system
CLHEP Class Library for HEP
CMKIN CMS Kinematics Package (legacy Fortran)
CMN Common Mode Noise
CMS Compact Muon Solenoid experiment
CMSIM CMS Simulation Package (legacy Fortran)
CMSSW CMS Software framework
COSINE Consistent Online Software INtegration Environment, project integrating online with offline software
CPT Computing, Physics, TriDAS and software projects of CMS
CRack Cosmic Rack, a set of TOB rods
CPU Central Processing Unit
COBRA Coherent Object-oriented Base for Reconstruction, Analysis and simulation (Framework)
CORBA Common Object Request Broker Architecture
COCOA CMS Object-oriented Code for optical Alignment
CompHEP Monte Carlo event generator for high-energy physics collisions
CPT Computing, Physics, TriDAS and Software projects
CRC Cyclic Redundancy Check error detection
CSC Cathode Strip Chamber muon system
CVS Concurrent Versions System

DØ Experiment at the FNAL Tevatron
D2S Data to Surface
DAC Digital to Analog Converter
DAQ Data Acquisition

DBMS Database Management System
DCC Data Concentrator Card
DCS Detector Control System
DCU Detector Control Unit
DDD Detector Description Database
DDL Data Description Language
DDU Detector Dependent Unit in DAQ system
DELPHI An experiment at LEP
DESY Deutsches Elektronen SYnchrotron laboratory, Hamburg
Digi Digitisation (of detector hit)
DMB DAQ MotherBoard of CSC system
DOH Digital Opto-Hybrid
DQM Data Quality Monitoring
DST Data Summary Tape - a compact event format
DT Drift Tube muon system
DY Drell-Yan

EB Electromagnetic Calorimeter (Barrel)
ECAL Electromagnetic Calorimeter
ED Extra Dimensions
EDM Event Data Model
EDMS Engineering Database Management System
EGEE Enabling Grids for e-science in Europe (a Grid project)
EE Electromagnetic Calorimeter (Endcap)
ELMB Embedded Local Monitoring Board (ECAL)
EM Electromagnetic
EMDB Equipment Management DataBase
EMU Endcap Muon system
ENC Equivalent Noise Charge
ES Endcap preShower detector
ESNET Energy Science Network (in the USA)
ESS ECAL Safety System
EVB EVent Builder
EVF Event Filter Farm
EVM Event Manager
EW ElectroWeak
Express Line Online stream for events requiring high priority and low latency offline processing

FAMOS CMS Fast Simulation
FE Front-End
FEC Front-End Card
FED Front-End Driver
FENIX ECAL front-end readout ASIC
FEVT Event format comprising the union of RAW and RECO data
FMM Fast Merging Module
FIFO First In First Out buffer
FLUKA Computer program for hadron shower calculations
FNAL Fermi National Accelerator Laboratory, USA
FPGA Field Programmable Gate Array
FSR Final State Radiation
FU Filter Unit

Garfield Computer program for electromagnetic field calculation
Gb Gigabit (10^9 bits)
GB Gigabyte (10^9 bytes)
GCALOR Computer program for hadron shower calculations
GCT Global Calorimeter Trigger (L1)
GEANT Detector simulation framework and toolkit
GIF Gamma Irradiation Facility
GMR Global Muon Reconstructor
GMSB Gauge Mediated Symmetry Breaking
GMT Global Muon Trigger (L1)
GOL Gigabit Optical Link
Grid Infrastructure for distributed computing
GSF Gaussian Sum Filter
GT Global Trigger (L1)
GUI Graphical User Interface
GUT Grand Unified Theory

H1 An experiment at the DESY HERA collider
H2 Beamline at CERN
HAD Hadronic
HCAL Hadron Calorimeter
HB Hadron Calorimeter (Barrel)
HE Hadron Calorimeter (Endcap)
HEP High Energy Physics
HEPEVT HEP Event (generated event format)
HERA Electron-proton collider at DESY
HERWIG Hadron Emission Reactions With Interfering Gluons, a monte Calro event generator for high-energy physics collisions
HF Hadron Calorimeter (Forward)
HG High Gain

HI Heavy Ion(s)
HIJING Heavy Ion Jet INteraction Generator, Monte Carlo event generator for heavy-ion collisions
HIP Hits and Impact Point alignment method, also Highly Ionizing Particle
HLT High-Level Trigger
HM Humidity Monitoring
HO Hadron Calorimeter (Outer Barrel)
HPD Hybrid Photo-Diode
HTML Hypertext Mark-up Language
HTR HCAL Trigger and Readout
HV High Voltage

ICC Internal Calibration Circuit, preshower
IGUANA Interactive Graphics for User ANAlysis - used for the CMS Event Display Package
I/O Input/Output
IOMC Input/Output Monte Carlo service
IOV Interval Of Validity
IP Impact Parameter, also Impact Point or Internet Protocol
ISR Initial State Radiation, also Intersecting Storage Ring collider at CERN
IT Information Technology

JES Jet Energy Scale

Kalman Filter Computational method for fitting tracks
kb kilobit (10^3 bits)
kB kilobytes (10^3 bytes)

L1 Level-1 hardware-based trigger
L1A Level-1 Accept
L3 An experiment at LEP
LAN Local Area Network
LAS Laser Alignment System
LCG LHC Computing Grid (a common computing project)
LCT Local Charged Track trigger primitive of CSC system
LED Large Extra Dimenstions, also Light Emitting Diode
LEP Large Electron Positron collider at CERN
LG Low Gain
LHC Large Hadron Collider
LHCb Large Hadron Collider Beauty experiment
LHCC LHC (review) Committee
LHEP Physics model of GEANT4
LL Leading Logarithm, also Log Likelihood
LO Leading Order calculation
LOI Letter Of Intent
LPC LHC Physics Center, Fermilab
LS Like-Sign
LSB Least Significant Bit
LSP Lightest Supersymmetric Particle
LUT Look-Up Table
LV Low Voltage
LVDS Low Voltage Differential Signaling
LVRB Low Voltage Regulator Board

MA Module Alignment
MAB Module Alignment of Barrel
Mb Megabit (10^6 bits)
MB Muon system (Barrel), also Mother Board or Megabyte (10^6 bytes)
MC Monte Carlo simulation program/technique, also Mini-Crate of DT system
ME Muon system (Endcap), also Matrix Element or Monitoring Element
MET Missing Transverse Energy
metadata Data describing characteristics of other data
MGPA Multiple Gain Pre-Amplifier chip, ECAL
MILLEPEDE Algorithm for tracker alignment
MIP Minimum Ionizing Particle
MPC Muon Port Card, CSC L1 trigger
MSUGRA Minimal SUper GRAvity model of supersymmetry
MSSM Minimal SuperSymmetric Model
MT Mean Time
MTCC Magnet Test Cosmic Challenge

ndf number of degrees of freedom
NLO Next-to-Leading Order calculation
NN Neural Network
NNLO Next-to-Next-to-Leading Order calculation
NS Numbering Scheme

O2O Online to Offline
ODBMS Object Database Management System
OMDS Online Master Data Storage
Online Stream Grouping of events (Primary Datasets) to simplify online data management
OO Object Oriented
OPAL An experiment at LEP
ORA Object Relational database
ORCA Object-oriented Reconstruction for CMS Analysis
ORCOF Offline ReConstruction OFfline subset, conditions database
ORCON Offline ReConstruction ONline subset, conditions database
OS Opposite-Sign, also Operating System
OSCAR Object-oriented Simulation for CMS Analysis and Reconstruction
Oval Software testing tool

P5 Point 5 collision area of LHC
PAW Physics Analysis Workstation (legacy interactive analysis application
PACE Preshower front end ASIC
PACT PAttern Comparator Trigger, RPC system
PB Petabyte (10^{15} bytes)
PC Personal Computer
PD Pixel Detector
PDF Parton Density Function, also Probability Distribution Function
PLC Programmable Logic Controller (ECAL)
PLD Programmable Logic Device
PLT Pixel Luminosity Telescope
POOL Persistency software from LCG
Primary Dataset Grouping of events according to physics (trigger) criteria
provenance The record of how a particular reconstruction result was obtained
PRS Physics Reconstruction and Selection groups
PS Proton Synchrotron, also Parton Showers
PTM Precision Temperature Monitoring, ECAL
PV Primary Vertex
PYTHIA Monte Carlo event generator for high-energy physics collisions

QCD Quantum Chromodynamics
QED Quantum Electrodynamics
QGSP Physics model of GEANT4
QPLL Quartz Phase-Locked Loop

RAW Event format from the online containing full detector and trigger data
RECO Event format for reconstructed objects such as tracks, vertices, jets, etc.
RecHit Reconstructed hit in a detector element
RHIC Relativistic Heavy Ion Collider (at Brookhaven, USA)
RMS Root Mean Square
ROB ReadOut Board, DT system
ROC ReadOut Chip, pixels
ROOT An object-oriented data analysis framework
ROS ReadOut Server board, DT system
RPC Resistive Plate Chamber muon system

Savannah Issue tracking system
SCA Switched Capacitor Array buffer, CSC system
SCADA Supervisory Control And Data Acquisition
SCRAM Software Configuration, Release And Management tool
SFM SubFarm Manager
Skim Subset of events selected from a larger set
SLB Sychronization and Link Board
SLT Soft Lepton Tag
SM Standard Model, also Super-Module (ECAL) or Storage Manager (DAQ)
SMB System Mother Board
S/N Signal to Noise ratio
SOAP Simple Object Access Protocol
SPS Super Proton Synchrotron collider at CERN
SRP Selective Readout Processor
SS Same-Sign
SST Silicon Strip Tracker
STL Standard Template Library
SUSY SUperSYmmetry
SV Secondary Vertex
SX5 Surface hall at Point 5 for CMS

T1, T2 Tracking telescopes of TOTEM
TAG Event index information such as run/event number, trigger bits, etc.
Tb Terabit (10^{12} bits)
TB Terabyte (10^{12} bytes)
TCC Trigger Concentrator Card
TCP Transmission Control Protocol
TDC Time to Digital Converter
TDR Technical Design Report
TEC Tracker EndCap
TF Track-Finder, muon L1 trigger
TIB Tracker Inner Barrel
TID Tracker Inner Disks
TMB Trigger MotherBoard, CSC L1 trigger
TOB Tracker Outer Barrel
TOTEM Separate experiment at P5 for forward physics
TPD Tracker Pixel Detector
TPG Trigger Primitive Generator
TRACO TRAck COrrelator, DT L1 trigger
TriDAS Trigger and Data Acquisition project
TRLB Token Ring Link Board
TS Trigger Server, DT L1 trigger
TTC Trigger Timing and Control
TTS Trigger Throttling System

UA1 An experiment at the CERN SPS collider
UA2 An experiment at the CERN SPS collider
UE Underlying Event
UED Universal Extra Dimensions
USC55 Underground Service Cavern at Point 5 for CMS
UXC55 Underground eXperimental Cavern at Point 5 for CMS

VBF Vector Boson Fusion
VBT Vacuum phototriode
VFE Very Front End
VPT Vacuum PhotoTriode

WAN Wide Area Network
WLS WaveLength Shifting
WWW World Wide Web

XDAQ Software framework for CMS Data Acquisition
XML eXtensible Markup Language

YB Yoke (Barrel)
YE Yoke (Endcap)

ZDC Zero Degree Calorimeter
ZEUS An experiment at the DESY HERA collider

References

Notes:

a) CMS Notes are available at `http://cmsdoc.cern.ch/docall.shtml` unless otherwise noted.

b) References marked **doi** should be prefixed with `http:dx.doi.org/`.

---

[1] CMS Collaboration, "The Compact Muon Solenoid Technical Proposal," *CERN/LHCC* **94-38** (1994). LHCC/P1.

[2] CMS Collaboration, "The Magnet Project Technical Design Report," *CERN/LHCC* **97-010** (1997). CMS TDR 1.

[3] CMS Collaboration, "The Hadron Calorimeter Technical Design Report," *CERN/LHCC* **97-031** (1997). CMS TDR 2.

[4] CMS Collaboration, "The Muon Project Technical Design Report," *CERN/LHCC* **97-32** (1997). CMS TDR 3.

[5] CMS Collaboration, "The Electromagnetic Calorimeter Technical Design Report," *CERN/LHCC* **97-033** (1997). CMS TDR 4, Addendum CERN/LHCC 2002-027.

[6] CMS Collaboration, "The Tracker Project Technical Design Report," *CERN/LHCC* **98-006** (1998). CMS TDR 5, Addendum CERN/LHCC 2000-016.

[7] CMS Collaboration, "The TriDAS Project Technical Design Report, Volume 1: The Trigger Systems," *CERN/LHCC* **2000-38** (2000). CMS TDR 6.1.

[8] CMS Collaboration, "The TriDAS Project Technical Design Report, Volume 2: Data Acquisition and High-Level Trigger," *CERN/LHCC* **2002-26** (2002). CMS TDR 6.2.

[9] R. Bailey and P. Collier, "Standard Filling Schemes for Various LHC Operation Modes," *LHC Project Note* **323** (2003).

[10] M. Lamont, "Estimates of Annual Proton Doses in the LHC," *LHC Project Note* **375** (2005).

[11] R. Bailey, "Overall Strategy for Early Luminosity Operation with Protons," *CERN EDMS Document* **LHC-OP-BCP-0001** (2004).

[12] R. Rattazzi, "Transplanckian collisions at future accelerators," `arXiv:hep-ph/0205265`. Presented at the 37th Rencontres de Moriond on Electroweak Interactions and Unified Theories, Les Arcs, France, 9–16 Mar.

[13] T. S. Virdee, "Detectors at LHC," *Phys. Rept.* **403-404** (2004) 401–434. doi:10.1016/j.physrep.2004.08.026.

[14] M. G. Albrow et al., "A uranium scintillator calorimeter with plastic-fibre readout," *Nucl. Instrum. and Methods* **A256** (1987). doi:10.1016/0168-9002(87)91035-7.

[15] V. I. Kryshkin and A. I. Ronzhin, "An optical fiber readout for scintillator calorimeters," *Nucl. Instrum. and Methods* **A247** (1986). doi:10.1016/0168-9002(86)90420-1.

[16] G. W. Foster, J. Freeman, and R. Hagstrom, "Scintillating tile/fiber calorimetry development at FNAL," *Nucl. Phys. Proceedings Supplements* **B23** (1991). doi:10.1016/0920-5632(91)90035-D.

[17] CMS HCAL, "Design, Performance, and Calibration of CMS Hadron-Barrel Calorimeter Wedges," *submitted to Nucl. Instrum. and Methods* (2005).

[18] CMS HCAL, "Design, Performance, and Calibration of CMS Hadron-Endcap Calorimeter Wedges," *submitted to Nucl. Instrum. and Methods* (2005).

[19] L. Borrello, A. Messineo, E. Focardi, and A. Macchiolo, "Sensor design for the CMS Silicon Strip Tracker," *CMS Note* **2003-020** (2002).

[20] M. French et al., "Design and Results from the APV25, a Deep Sub-micron CMOS Front-End Chip for the CMS Tracker," *Nucl. Instrum. and Methods* **A534** (2001) 359–365. doi:10.1016/S0168-9002(01)00589-7.

[21] J. Troska et al., "Optical Readout and Control Systems for the CMS Tracker," *IEEE Trans. Nucl. Sci.* **50** (2003) 1067–1072. doi:10.1109/TNS.2003.815124.

[22] J. Coughlan et al., "The CMS Tracker Front-End Driver," in *Proceedings of the 9th Workshop on Electronics for LHC Experiments*. Amsterdam, 29 September – 3 October, 2003. Available at http://doc.cern.ch//archive/cernrep/2003/2003-006/p255.pdf.

[23] F. Drouhin et al., "The CERN CMS Silicon Strip Tracker Control System," *CMS Conference Report* **2004/032** (2004). Presented at ROMA 2004, NSS MIC SNPS and RTSD, Roma, Italy, October 18-21.

[24] E. Noah et al., "Qualification of the CMS Tracker Control Link Digital Optohybrid," in *Proceedings of the 10th Workshop on Electronics for LHC Experiments*. Boston, Massachusetts, 13-17 September, 2004. Published on http://indico.cern.ch.

[25] CMS Collaboration, "The Computing Project Technical Design Report," *CERN/LHCC* **2005-23** (2005). CMS TDR 7.

[26] V. Innocente, L. Silvestris, and D. Stickland, "CMS software architecture: Software framework, services and persistency in high level trigger, reconstruction and analysis," *Computer Physics Communications* **140** (2001) 31–44. doi:10.1016/S0010-4655(01)00253-3.

[27] "ORCA: CMS Reconstruction Package." Site located at http://cmsdoc.cern.ch/orca.

[28] R. Brun, F. Bruyant, M. Maire, A. C. McPherson, and P. Zanarini, "GEANT3,". CERN Program Library W5013: CERN-DD/EE/84-1.

[29] "CMSIM: CMS Simulation Package (obsolete)." Available at http://cmsdoc.cern.ch/cmsim/cmsim.html.

[30] **GEANT4** Collaboration, S. Agostinelli et al., "GEANT4: A simulation toolkit," *Nucl. Instrum. and Methods* **A506** (2003) 250–303. doi:10.1016/S0168-9002(03)01368-8.

[31] "OSCAR: CMS Simulation Package Home Page." http://cmsdoc.cern.ch/oscar.

[32] "CMS Software home page." https://cmsdoc.cern.ch/cms/cpt/Software/html/General/.

[33] CMS Collaboration, "Computing Technical Proposal," *CERN/LHCC* **1996-45** (1996).

[34] "LCG Project Application Home Page." http://lcgapp.cern.ch/.

[35] "CDF Framework." Information available at http://www-cdf.fnal.gov/upgrades/computing/projects/framework.

[36] "DØ Framework." Information available at http://chep2000.pd.infn.it/short_p/spa_a230.pdf.

[37] "POOL Persistence Framework." Information available at http://lcgapp.cern.ch/project/persist/.

[38] V. Brigljevic et al., "Using XDAQ in Application Scenarios of the CMS Experiment," in *Proceedings of the CHEP 03 Conference*. La Jolla, California, 24-28 March, 2003. arXiv:hep-ex/0305076.

[39] J. Gutleber et al., "Clustered Data Acquisition for the CMS experiment," in *Proceedings of the CHEP 01 Conference*. Beijing, China, 3-7 September, 2001.

[40] "XDAQ home page." http://xdaq.web.cern.ch/.

[41] O. Boyle, R. McLaren, and E. van der Bij, "The S-LINK Interface Specification." http://hsi.web.cern.ch/HSI/s-link.

[42] A. Racz, R. McLaren, and E. van der Bij, "The S-LINK 64-bit Extension Specification: S-LINK64." https://edms.cern.ch/file/249683/2/slink64_v20.pdf.

[43] M. Case et al., "Detector Description Domain Architecture and Data Model," *CMS Note* **2001-057** (2001).

[44] Z. Xie, "POOL The Persistency Framework for LHC - New developments and CMS applications." Presented at Frontier Science 2005: New Frontiers in Subnuclear Physics, September 12-17, Milan, Italy, 2005. Available at http://cmsdoc.cern.ch/~xiezhen/poolora.ps.

[45] T. Sjostrand, L. Lonnblad, and S. Mrenna, "PYTHIA 6.2: Physics and manual," arXiv:hep-ph/0108264.

[46] "Herwig Web Site." http://hepwww.rl.ac.uk/theory/seymour/herwig/.

[47] "Sherpa Web Page." http://www.physik.tu-dresden.de/~krauss/hep/.

[48] P. Arce and M. Wadhwa, "Deviation in Matter of 45 GeV Muons in GEANT 3 and GEANT 4, A Comparison with L3 Data," *CMS Note* **2000-016** (2000).

[49] G. Grindhammer and S. Peters, "The parametrized simulation of electromagnetic showers in homogeneous and sampling calorimeters," *Int. Conf. of Monte-Carlo Simulation in High Energy and Nuclear Physics, Tallahassee, Florida, USA* (1993) arXiv:hep-ex/0001020.

[50] G. Grindhammer, M. Rudowicz, and S. Peters, "The Fast Simulation of Electromagnetic and Hadronic Showers," *Nucl. Instrum. and Methods* **A290** (1990) 469. doi:10.1016/0168-9002(90)90566-O.

[51] "Oracle Relational Database Management System home page." http://www.oracle.com.

[52] "Core Libraries and Services Project." Information available at http://seal.cern.ch.

[53] S. Kosyakov et al., "Frontier: High Performance Database Access Using Standard Web Components," in *Proceedings of the CHEP 04 Conference*. Interlaken Switzerland, September 27^{th} - October 1^{st}, 2004. Published on http://indico.cern.ch.

[54] "Apache Tomcat." Information available at http://www.apache.org.

[55] "Squid home page." http://www.squid-cache.org.

[56] "IGUANA Web Site." Located at http://iguana.web.cern.ch/iguana/.

[57] G. Alverson, G. Eulisse, S. Muzaffar, I. Osborne, L. Taylor, and L. A. Tuura, "IGUANA Architecture, Framework and Toolkit for Interactive Graphics," *ECONF* **C0303241** (2003) MOLT008, arXiv:cs.se/0306042.

[58] G. Alverson, I. Osborne, L. Taylor, and L. Tuura, "A Coherent and Non-Invasive Open Analysis Architecture and Framework with Applications in CMS," in *Proceedings of Computing in High Energy (CHEP 2001)*, H.S. Chen, ed. Science Press, Beijing/New York, Beijing, 3-7 Sept, 2001.

[59] "Coin3D : a high-level 3D graphics toolkit." Information available at http://www.coin3d.org/.

[60] "Qt : a comprehensive C++ application development framework." Information available at http://www.trolltech.com/.

[61] G. Alverson, G. Eulisse, S. Muzaffar, I. Osborne, L. Taylor, and L. A. Tuura, "IGUANA: A high-performance 2D and 3D visualisation system," *Nucl. Instrum. and Methods* **A534** (2004) 143–146. doi:10.1016/j.nima.2004.07.036.

[62] "ROOT : an Object Oriented Data Analysis Framework." Information available at http://root.cern.ch/.

[63] "HippoDraw : a highly interactive data analysis environment." Information available at http://www.slac.stanford.edu/grp/ek/hippodraw/.

[64] "Qt-ROOT home page." http://www-linux.gsi.de/~go4/qtroot/html/qtroot.html.

[65] "The Concurrent Version System (CVS)." Information available at http://www.nongnu.org/cvs/.

[66] "SCRAM." Information available at http://cmsdoc.cern.ch/Releases/SCRAM/current/doc/tex/manual/SCRAM-manual.pdf.

[67] "NICOS manager for nightly builds." Information available at http://www.usatlas.bnl.gov/computing/software/nicos/.

[68] G. Eulisse et al., "A coherent environment of software improvement tools for CMS," *Nucl. Instrum. and Methods* **A534** (2004) 138–142. doi:10.1016/j.nima.2004.07.035.

[69] "Savannah home page." http://savannah.gnu.org/.

[70] "CMSSW Release Schedule." Information available at https://uimon.cern.ch/twiki/bin/view/CMS/ReleaseSchedule.

[71] "Software Developement Tools Homepage." https://uimon.cern.ch/twiki/bin/view/CMS/SoftwareDevelopementToolsSubTask.

[72] "WebSRT." Login page located at http://advance.colorado.edu/cgi-bin/cgiwrap/drjohn/WebSrtLogin.

[73] "xcmsi home page." http://cmsdoc.cern.ch/cms/oo/repos_standalone/download/index.php.

[74] "Unit test." Wikipedia entry located at http://en.wikipedia.org/wiki/Unit_testing.

[75] "Spi - CppUnit." Information available at http://spi.cern.ch/Components/UnitTesting/UserDocumentation/Web/HowTo-TestFramework-CppUnit.html.

[76] "LCG QA Test Coverage Information." Information available at http://spi.cern.ch/qa/qatestcov.html.

[77] "Oval Web Site." http://spi.cern.ch/testing/HowTo-TestFramework-Oval.html.

[78] M. Aguilar Benitez et al., "Construction and test of the final CMS Barrel Drift Tube Muon chamber prototype," *Nucl. Instrum. and Methods* **A480** (2002) 658. doi:10.1016/S0168-9002(01)01227-X.

[79] C. Albajar et al., "Test beam analysis of the first CMS drift tube muon chamber," *Nucl. Instrum. and Methods* **A525** (2004) 465. doi:10.1016/j.nima.2004.01.080.

[80] F. Cavallo et al., "Test of MB3 Muon Barrel Drift Chamber with Cosmic Rays," *CMS Note* **2003/017** (2003).

[81] M. Abbrescia et al., "Beam test results on double-gap resistive plate chambers proposed for CMS experiment," *Nucl. Instrum. and Methods* **A414** (1998) 135. doi:10.1016/S0168-9002(98)00571-3.

[82] P. Arce et al., "Bunched beam test of the CMS drift tubes local muon trigger," *Nucl. Instrum. and Methods* **A534** (2004) 441. doi:10.1016/j.nima.2004.06.169.

[83] J. Christiansen, "A data driven high resolution Time-to-Digital Converter," in *Proceedings of the 6th Workshop on electronics for LHC experiments*. Cracow, Poland, 11-15 September, 2000. Available in CERN-LHCC-2000-041.

[84] M. Husejko et al., "Readout system for CMS RPC Muon Trigger," *Proceedings of SPIE* **5775** (2005) 121–130. doi:10.1117/12.610600.

[85] R. Breedon et al., "Performance and Radiation Testing of a Low-Noise Switched Capacitor Array for the CMS Endcap Muon System," in *Proceedings of the Sixth Workshop on Electronics for LHC Experiments*, pp. 187–191. Crakow, Poland, September, 2000. Available in CERN-LHCC-2000-041.

[86] T. Ferguson et al., "Anode front-end electronics for the cathode strip chambers of the CMS Endcap Muon detector," *Nucl. Instrum. and Methods* **A539** (2005) 386. doi:10.1016/j.nima.2004.09.045.

[87] P. Arce et al., "Performance test of the CMS link alignment system," *Nucl. Instrum. and Methods* **A492** (2002). doi:10.1016/S0168-9002(02)01412-2.

[88] M. Benettoni et al., "Study of the Internal Alignment of the CMS Muon Barrel Drift Chambers using Cosmic Ray Tracks," *CMS Note* **2004/001** (2004).

[89] J. Fuchs et al., "Barrel Muon Chambers survey and alignment DataBase," 2005. Information available at http://kismalac.phys.klte.hu/cgi-bin/Survey/indexpage.pl.

[90] Z. Szillasi et al., "CMS Barrel Yoke YB+2," *CMS-SY-UR* **0019** (2005).

[91] P. Arce and A. López-Virto, "CMS Object Oriented Code for Optical Alignment (COCOA)," *CMS IN* **2002/060** (2002).

[92] A. Calderón et al., "Muon System Alignment with Tracks," *CMS Note* **2006/016** (2006).

[93] V. Blobel and C. Kleinwort, "A New Method for High-Precision Aligment of Track Detectors," in *Proceedings of the Conference on Advanced Statistical Techniques in Particle Physics*. Durham, England, June, 2002. arXiv:hep-ex/0208021.

[94] I. Belotelov et al., "Simulation of Misalignment Scenarios for CMS Tracking Devices," *CMS Note* **2006/008** (2006).

[95] I. Belotelov et al., "Influence of Misalignment Scenarios on Muon Reconstruction," *CMS Note* **2006/017** (2006).

[96] R. Arcidiacono et al., "The final prototype of the Fast Merging Module (FMM) for readout status processing in CMS DAQ," in *Proceedings of the 10th Workshop on electronics for LHC experiments*. Boston, Massachusetts, 24-28 March, 2004. Published on http://indico.cern.ch.

[97] V. Aerts et al., "Status and Perspectives of Detector Databases in the CMS Experiment at the LHC," *CMS Note* **2004/026** (2004).

[98] P. Garcia-Abia, "DT costruction database." Information available at http://oraweb03.cern.ch:9000/pls/cms_mb_prod/chambers.main.

[99] "CSC databases." Information available at http://cmsdoc.cern.ch/cms/CSC/CERN/db.html.

[100] "CSC Test Beam database." Information available at http://oraweb03.cern.ch:9000/pls/cms_emu_cern.pro/run_log1.top_page.

[101] "Barrel RPC Construction database." Information available at http://webcms.ba.infn.it/rpc.

[102] "EndCap RPC Construction database." Information available at http://forwardrpc.cern.ch/cms_forward_rpc/production_db/Main_db.htm.

[103] J. Puerta-Pelayo, M. C. Fouz, and P. Garcia-Abia, "Parametrization of the Response of the Muon Barrel Drift Tubes," *CMS Note* **2005/018** (2005).

[104] R. Veenhof, "Garfield, a Drift Chamber simulation program user's guide," *CERN Program Library W5050* (1994).

[105] R. Wilkinson and P. T. Cox, "Simulating the Endcap Muon CSC System in ORCA," *CMS Note* **2001/013** (2003).

[106] E. Gatti et al., "Analysis of the position resolution in centroid measurements in MWPC," *Nucl. Instrum. and Methods* **A188** (1981) 327–346. doi:10.1016/0029-554X(81)90513-9.

[107] N. Amapane et al., "Comparison of DT testbeam results on local track reconstruction with the OSCAR + ORCA simulation," *CMS Note* **2006/009** (2006).

[108] S. Durkin, T. Ferguson, and N. Terentiev, "Validation of the Simulation for the CMS Endcap Muon CSC Front-end Electronics," *CMS Note* **2006/015** (2006).

[109] I. Belotelov et al., "Electromagnetic Secondaries and Punchthrough Effects in CMS ME1/1," *CMS Note* **2006/xxx** (2006).

[110] M. Abbrescia et al., "Cosmic Ray Tests of Double-gap Resistive Plate Chambers for the CMS Experiment," *Nucl. Instrum. and Methods* **A550** (2005) 116. doi:10.1016/j.nima.2005.06.074.

[111] R. Kinnunen et al., "Impact of Muon Trigger Coverage on Physics," *CMS Note* **1998/020** (1998).

[112] D. Acosta et al., "Design features and test results of the CMS endcap muon chambers," *Nucl. Instrum. and Methods* **A494** (2002) 504. doi:10.1016/S0168-9002(02)01540-1.

[113] D. Acosta et al., "Development and test of a prototype regional track finder for the Level-1 trigger of the cathode strip chamber muon system of CMS," *Nucl. Instrum. and Methods* **A496** (2003) 64. doi:10.1016/S0168-9002(02)01622-4.

[114] H. Sakulin, "Improved di-muon trigger rate simulation methods for the CMS experiment at LHC," *J. Phys. G: Nucl. Part. Phys.* **30** (2004) N45–N55. doi:10.1088/0954-3899/30/12/N02.

[115] R. Wilkinson et al., "Simulating the Endcap Muon CSC System in ORCA," *CMS Note* **2001-013** (2001).

[116] G. Bruno et al., "Local reconstruction in the muon detectors," *CMS Note* **2002-043** (2002).

[117] P. Zotto et al., "Fine syncronization of the muon Drift Tubes Local Trigger," *CMS Note* **2006/002** (2005).

[118] E. Barberis et al., "Trigger and Reconstruction Studies with Beam Halo and Cosmic Muons," *CMS Note* **2006/012** (2006).

[119] **Particle Data Group** Collaboration, S. Eidelman et al., "Review of particle physics," *Phys. Lett.* **B592** (2004) 1.

[120] A. A. Annenkov, M. V. Korzhik, and P. Lecoq, "Lead tungstate scintillation material," *Nucl. Instrum. and Methods* **A490** (2002) 30–50. doi:10.1016/S0168-9002(02)00916-6.

[121] X. D. Qu, L. Y. Zhang, and R. Y. Zhu, "Radiation induced color centers and light monitoring for lead tungstate crystals," *IEEE Trans. Nucl. Sci.* **47** (2000) 1741–1747. doi:10.1109/TNS.2000.914439.

[122] D. Graham and C. Seez, "Simulation of Longitudinal Light Collection Uniformity in $PbWO_4$ crystals," *CMS Note* **1996/002** (1996).

[123] R. Y. Zhu, "Radiation damage in scintillating crystals," *Nucl. Instrum. and Methods* **A413** (1998) 297–311. doi:10.1016/S0168-9002(98)00498-7.

[124] M. Huhtinen, P. Lecomte, D. Luckey, F. Nessi-Tedaldi, and F. Pauss, "High-energy proton induced damage in $PbWO_4$ calorimeter crystals," *Nucl. Instrum. and Methods* **A545** (2005) 63–87. doi:10.1016/j.nima.2005.01.304.

[125] K. Deiters et al., "Double screening tests of the CMS ECAL avalanche photodiodes," *Nucl. Instrum. and Methods* **A543** (2005) 549–558. doi:10.1016/j.nima.2005.01.088.

[126] P. Wertelaers et al., "CMS ECAL EDR-4, Volume 2: ECAL Preshower," *CMS-GE-ER* **0002** (2000).

[127] M. Raymond et al., "The MGPA Electromagnetic Calorimeter Readout Chip for CMS." Proceedings of the 9th Workshop on Electronics for the LHC Experiments, 2003. Available in CERN-LHCC-2003-055.

[128] G. Minderico et al., "A CMOS low power, quad channel, 12 bit, 40 Ms/s pipelined ADC for application in particle physics calorimetry." Proceedings of the 9th Workshop on Electronics for the LHC Experiments, 2003. Available in CERN-LHCC-2003-055.

[129] K. Kloukinas et al., "Charaterization and production testing of a quad 12 bit 40 Ms/s A/D converter with automatic digital range selection for calorimetry." Proceedings of 11th Workshop on Electronics for the LHC and FutureExperiments, Heidelberg Sept. 12-16, 2005, 2005. PowerPoint slides available on http: indico.cern.ch.

[130] M. Hansen, "The New Read Out Architecture for the CMS ECAL." Proceedings of the 9th Workshop on Electronics for the LHC Experiments, 2003. Available in CERN-LHCC-2003-055.

[131] R. Alemany et al., "CMS ECAL Off-Detector Electronics," in *Proceedings of the 11th International Conference On Calorimetry In High Energy Physics, CALOR 2004*. Perugia, Italy, March, 2004. Slides available at `http://www.pg.infn.it/calor2004/program/pres/wednesday_morning/alemany.pdf`.

[132] R. Alemany et al., "Overview of the ECAL Off-Detector Electronics of the CMS Experiment." Presented at IEEE Nuclear Science Symposium, Roma, Italy, Oct. 18-21, 2004, 2004. `doi:10.1109/NSSMIC.2004.1462385`.

[133] P. Paganini et al., "Tests of the boards generating the CMS ECAL Trigger Primitives: from the On Detector electronics to the Off-Detector electronics system," in *Proceedings of the 10th Workshop on electronics for LHC experiments*. Boston, Massachusetts, 13-17 September, 2004. Published on http://indico.cern.ch.

[134] N. Almeida et al., "Calorimeter Trigger Synchronization in CMS, Implementation and Test System," in *Proceedings of the 10th Workshop on electronics for LHC experiments*. Boston, Massachusetts, 13-17 September, 2004. Published on http://indico.cern.ch.

[135] N. Almeida et al., "Data Concentrator Card and Test System for the CMS ECAL Readout." Proceedings of the 9th Workshop on Electronics for the LHC Experiments, 2003. Available in CERN-LHCC-2003-055.

[136] J. C. Silva et al., "Test results of the Data Concentrator Card of the CMS Electromagnetic Calorimeter Readout System," in *Proceedings of the 10th Workshop on electronics for LHC experiments*. Boston, Massachusetts, 13-17 September, 2004. Published on http://indico.cern.ch.

[137] N. Almeida et al., "The Selective Readout Processor for the CMS Electromagnetic Calorimeter," *IEEE Trans. Nucl. Sci.* **52** (2005), no. 3, 772–777. `doi:10.1109/TNS.2005.850946`.

[138] W. Bialas et al., "The Readout, Fast Control and Powering Architecture for the CMS Preshower." Presented at 11th Workshop on electronics for LHC experiments, Heidelberg, Germany, 2005. PowerPoint slides available at http: indico.cern.ch.

[139] P. Aspell et al., "PACE3: A large dynamic range analogue memory ASIC assembly designed for the readout of silicon sensors in the LHC CMS Preshower," in *Proceedings of the 10th Workshop on electronics for LHC experiments*. Boston, Massachusetts, 13-17 September, 2004. Published on http://indico.cern.ch.

[140] K. Kloukinas et al., "Kchip: a radiation tolerent digital data concentrator chip for the CMS Preshower detector." Presented at 9th Workshop on electronics for LHC experiments, Amsterdam, Netherlands, CERN-LHCC-2003-055, 2003.

[141] P. Paganini and I. van Vulpen, "Pulse Amplitude Reconstruction in the CMS ECAL using the Weights Method," *CMS Note* **2004/025** (2004).

[142] R. Brunelière and A. Zabi, "Reconstruction of the CMS electromagnetic calorimeter signal amplitude," *CMS Note* **2006/xxx** (2006).

[143] E. Auffray et al., "Performance of ACCOS, an Automatic Crystal quality Control System for the PWO crystals of the CMS calorimeter," *Nucl. Instrum. and Methods* **A456** (2001). doi:10.1016/S0168-9002(00)00665-3.

[144] E. Auffray et al., "Cross-Calibration of Two Automatic Quality Control Systems for the CMS ECAL Crystals," *CMS Note* **2003/003** (2003).

[145] L. M. Barone et al., "Improvement on $PbWO_4$ Crystal Inter-calibration Precision from Light Yield Measurements at the INFN-ENEA Regional Center," *CMS Rapid Note* **2004/003** (2004).

[146] F. Cavallari et al., "Relative Light Yield comparison between laboratory and testbeam data for CMS ECAL $PbWO_4$ crystals," *CMS Rapid Note* **2004/002** (2004).

[147] L. M. Barone et al., "Correlation between the light yield and the longitudinal transmission in $PBWO_4$ crystals and the impact on the precision of the crystals inter-calibration," *CMS Rapid Note* **2004/005** (2004).

[148] W. Bertl et al., "Feasibility of Intercalibration of CMS ECAL Supermodules with Cosmic Rays," *CMS Note* **2004/036** (2004).

[149] M. Bonesini et al., "Inter-calibration of the CMS electromagnetic calorimeter with cosmic rays before installation," *CMS Note* **2005/023** (2005).

[150] D. Futyan and C. Seez, "Intercalibration of ECAL Crystals in Phi Using Symmetry of Energy Deposition," *J. Phys. G: Nucl. Part. Phys.* **29** (2003) 1299–1326. doi:10.1088/0954-3899/29/6/326.

[151] D. Futyan, "Intercalibration of the CMS Electromagnetic Calorimeter Using Jet Trigger Events," *CMS Note* **2004/007** (2004).

[152] L. Agostino et al., "Intercalibration of the CMS Electromagnetic Calorimeter with Isolated Electrons," *CMS Note* **2006/021** (2006).

[153] A. Favara et al., "Calibration of the L3 BGO Calorimeter Using an RFQ Accelerator," *Nucl. Phys. B (Proc. Supp.)* **78** (1999) 465–470. doi:10.1016/S0920-5632(99)00588-5.

[154] R. Paramatti and P. Meridiani, "Use of $Z \to e^+e^-$ events for ECAL calibration," *CMS Note* **2006/xxx** (2006).

[155] I. Evangelou, "CMS Preshower in-situ absolute calibration." Presented at the 9th ICATPP Conference, Villa-Olmo, Como, Italy, 2005. Microsoft Word version available at http://pcams01.mib.infn.it/ICATPP9th_2005/Calorimetry/Evangelou.doc.

[156] P. Bloch et al., "Performance of Si sensors irradiated to 5×10^{14} n/cm^2," *Nucl. Instrum. and Methods* **A517** (2004) 121–127. doi:10.1016/j.nima.2003.09.058.

[157] L. Zhang et al., "Performance of the Monitoring Light Source for the CMS Lead Tungstate Crystal Calorimeter," *IEEE Trans. Nucl. Sci.* **52** (2005). doi:10.1109/TNS.2005.852661.

[158] A. Van Lysebetten and P. Verrecchia, "Performance and measurements of the light monitoring system for CMS-ECAL from 2002 test beam data," *CMS Rapid Note* **2004/001** (2004).

[159] P. Bonamy et al., "The ECAL calibration: Use of the light monitoring system Version 2.0," *CMS Note* **1998/013** (1998).

[160] E. Meschi et al., "Electron Reconstruction in the CMS Electromagnetic Calorimeter," *CMS Note* **2001/034** (2001).

[161] P. Adzic et al., "Results of the first performance tests of the CMS electromagnetic calorimeter," *Eur. Phys. J. C* **44, Supplement 2** (2006) 1–10. doi:10.1140/epjcd/s2005-02-011-3.

[162] S. Banerjee, "Simulation Geometry for the CMS HCAL," *CMS Note* **2005/016** (2005).

[163] CMS HCAL, "Synchronization and Timing in the CMS Hadron Calorimeter," *submitted to Nucl. Instrum. Methods* (2005).

[164] I. Golutvin et al., "Simulation of Radiation Damage in HE Scintillating Tiles and Pion Energy Resolution After 10 Years of LHC Operation," *CMS Note* **2002-013** (2002).

[165] D. Denegri, R. Kinnunen, and A. Nikitenko, "Study of Calorimeter Calibration with Tau's in CMS," *CMS Note* **1997-039** (1997).

[166] R. Kinnunen and A. Nikitenko, "Study of calorimeter calibration with pions from jets in CMS," *CMS Note* **1997-097** (1997).

[167] **CMS-HCAL** Collaboration, V. V. Abramov et al., "Studies of the response of the prototype CMS hadron calorimeter, including magnetic field effects, to pion, electron, and muon beams," *Nucl. Instrum. and Methods* **A457** (2001) 75–100, arXiv:hep-ex/0007045. doi:10.1016/S0168-9002(00)00711-7.

[168] S. Banerjee and S. Banerjee, "Validation of Geant4 Using 1986 CMS HCAL Test Beam Data," *CMS Note* **2004-032** (2004).

[169] V. D. Elvira, "Validation of Geant4 Physics Using the CMS HCAL Test Beam 2002 Experiment," *CMS Note* **2004-021** (2004).

[170] V. D. Elvira, "Measurement of the Pion Energy Response and Resolution in the Hadronic Barrel Calorimeter using CMS HCAL Test Beam 2002 Data," *CMS Note* **2004-020** (2004).

[171] D. Green, "Calibration of the CMS calorimeters," *Fermi Note* **FN-0704** (2001).

[172] C. Amsler et al., "A high resolution silicon beam telescope," *Nucl. Instrum. and Methods* **A480** (2002) 501–507. doi:10.1016/S0168-9002(01)01241-4.

[173] A. Dorokhov, "Performance of radiation hard pixel sensors for the CMS experiment". PhD thesis, Universität Zürich, 2005.

[174] T. Rohe et al., "Position dependence of charge collection in prototype sensors for the CMS pixel detector," *IEEE Trans. Nucl. Sci.* **51** (2004) 1150–1157. doi:10.1109/TNS.2004.829487.

[175] B. Henrich and R. Kaufmann, "Lorentz-angle in irradiated silicon," *Nucl. Instrum. and Methods* **A477** (2002) 304–307. doi:10.1016/S0168-9002(01)01865-4.

[176] A. Dorokhov et al., "Tests of silicon sensors for the CMS pixel detector," *Nucl. Instrum. and Methods* **A530** (2004) 71–76, arXiv:physics/0311050. doi:10.1016/j.nima.2004.05.050.

[177] A. Dorokhov et al., "Electric field measurement in heavily irradiated pixel sensors," arXiv:physics/0412036.

[178] S. Braibant et al., "Investigation of design parameters for radiation hard silicon microstrip detectors," *Nucl. Instrum. and Methods* **A485** (2002) 343–361. doi:10.1016/S0168-9002(01)02120-9.

[179] C. Canali et al., "Electron and Hole Drift Velocity Measurements in Silicon and Their Empirical Relation to Electric Field and Temperature," *IEEE Trans. on Electron Devices* **22** (1975) 1045–1047.

[180] V. Bartsch et al., "An algorithm for calculating the Lorentz angle in silicon detectors," *Nucl. Instrum. and Methods* **A497** (2003) 389–396. doi:10.1016/S0168-9002(02)01801-6.

[181] W. Adam et al., "Track reconstruction in the CMS tracker," *CMS Note* **2006/xxx** (2006).

[182] S. Cucciarelli, D. Kotlinski, and T. Todorov, "Position Determination of Pixel Hits," *CMS Note* **2002-049** (2002).

[183] I. Lokhtin, "The HYDJET fast event generator." See http://lokhtin.home.cern.ch/lokhtin/hydro/hydjet.html.

[184] **CMS Tracker** Collaboration, W. Adam et al., "The effect of highly ionising particles on the CMS silicon strip tracker," *Nucl. Instrum. and Methods* **A543** (2005) 463–482. doi:10.1016/j.nima.2004.11.049.

[185] R. Frühwirth, P. Kubinec, W. Mitaroff, and M. Regler, "Vertex reconstruction and track bundling at the LEP collider using robust algorithms," *Comput. Phys. Commun.* **96** (1996) 189–208. doi:10.1016/0010-4655(96)00040-9.

[186] T. Speer, K. Prokofiev, R. Frühwirth, W. Waltenberger, and P. Vanlaer, "Vertex Fitting in the CMS Tracker," *CMS Note* **2006/xxx** (2006).

[187] W. Waltenberger, "Development of Vertex Finding and Vertex Fitting Algorithms for CMS". PhD thesis, Technischen Universität Wien, 2004.

[188] R. Frühwirth, W. Waltenberger, K. Prokofiev, T. Speer, and P. Vanlaer, "New developments in vertex reconstruction for CMS," *Nucl. Instrum. and Methods* **A502** (2003) 699–701. doi:10.1016/S0168-9002(03)00548-5.

[189] J. D'Hondt, P. Vanlaer, R. Frühwirth, and W. Waltenberger, "Sensitivity of robust vertex fitting algorithms," *IEEE Trans. Nucl. Sci.* **51** (2004) 2037–2044. doi:10.1109/TNS.2004.832296.

[190] R. Frühwirth and T. Speer, "A Gaussian-sum filter for vertex reconstruction," *Nucl. Instrum. and Methods* **A534** (2004) 217–221. doi:10.1016/j.nima.2004.07.090.

[191] K. Prokofiev and T. Speer, "A kinematic fit and a decay chain reconstruction library," *CERN* **2005-002** (2005). Proc. of the 2004 Conference for Computing in High-Energy and Nuclear Physics (CHEP 04), Interlaken, Switzerland, 2004, available at http://indico.cern.ch.

[192] S. Cucciarelli, M. Konecki, D. Kotlinski, and T. Todorov, "Track Reconstruction, primary vertex finding and seed generation with the Pixel Detector," *CMS Note* **2006/026** (2006).

[193] S. Cucciarelli, M. Konecki, D. Kotlinski, and T. Todorov, "Track-Parameter Evaluation and Primary-Vertex Finding with the Pixel Detector," *CMS Note* **2003-026** (2003).

[194] S. Cucciarelli, "The Performance of the CMS Pixel Detector and the Primary Vertex Finding," *CMS CR* **2003/057** (2003). Presented at Vertex 2003, Lake Windemere, Cumbria, UK, September 14-19.

[195] P. Vanlaer, V. Barbone, N. De Filippis, T. Speer, O. Buchmueller, and F.-P. Schilling, "Impact of Tracker Misalignment on Track and Vertex Reconstruction," *CMS Note* **2006/xxx** (2006).

[196] T. Muller, C. Piasecki, G. Quast, and C. Weiser, "Inclusive secondary vertex reconstruction in jets," *CMS Note* **2006/xxx** (2006).

[197] **CDF** Collaboration, F. Abe et al., "Evidence for top quark production in $p - \bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV," *Phys. Rev.* **D50** (1994) 2966. doi:10.1103/PhysRevD.50.2966.

[198] TOTEM Collaboration, "TOTEM Technical Design Report," *CERN/LHCC* **2004-002** (2004).

[199] V. Drollinger, "Simulation of Beam Halo and Cosmic Muons," *CMS Note* **2005/012** (2005).

[200] V. Karimäki, A. Heikkinen, T. Lampén, and T. Lindén, "Sensor Alignment by Tracks," *CMS CR* **2003/022** (2003). Presented at CHEP 2003, La Jolla.

[201] V. Karimäki, T. Lampén, and F.-P. Schilling, "The HIP Algorithm for Track Based Alignment and its Application to the CMS Pixel Detector," *CMS Note* **2006/018** (2006).

[202] R. Frühwirth, E. Widl, and W. Adam, "A Kalman Filter for Track-based Alignment," *CMS Note* **2006/022** (2006).

[203] V. Blobel, "MILLEPEDE program description and code." See http://www.desy.de/~blobel/wwwmille.html.

[204] V. Blobel, "Private communication," 1997.

[205] R. Mcnulty, T. Shears, and A. Skiba, "A Procedure for Software Alignment of the CDF Silicon System," *CDF/DOC/Tracking/57000* (2001). Postscript version available at http://hep.ph.liv.ac.uk/cdf/Liverpool/silicon/alignment/align.ps.

[206] P. Schleper, G. Steinbrück, and M. Stoye, "Software Alignment of the CMS Tracker using MILLEPEDE II," *CMS Note* **2006/011** (2006).

[207] T. Lampén, V. Karimäki, S. Saarinen, and O. Buchmüller, "Alignment of the Cosmic Rack with the Hits and Impact Points Algorithm," *CMS Note* **2006/006** (2006).

[208] M. S. Mennea, A. Regano, and G. Zito, "CMS Tracker Visualization," *CMS Note* **2004-009** (2004).

[209] **FELIX** Collaboration, E. Lippmaa et al., "FELIX: A full acceptance detector at the LHC," *CERN-LHCC* **97-45** (1997).

[210] A. Ageev et al., "A Full acceptance detector at the LHC (FELIX)," *J. Phys.* **G28** (2002) R117–R215. doi:10.1088/0954-3899/28/6/201.

[211] TOTEM Collaboration, "Addendum to the TOTEM-TDR," *CERN/LHCC* **2004-020** (2004).

[212] **FP420** Collaboration, M. Albrow et al., "FP420: A proposal to investigate the feasibility of installing proton tagging detectors in the 420 m region of the LHC," *CERN/LHCC* **2005-025** (2005).

[213] A. De Roeck, V. A. Khoze, A. D. Martin, R. Orava, and M. G. Ryskin, "Ways to detect a light Higgs boson at the LHC," *Eur. Phys. J.* **C25** (2002) 391–403, arXiv:hep-ph/0207042. doi:10.1007/s10052-002-1032-9.

[214] B. E. Cox et al., "Detecting the standard model Higgs boson in the W W decay channel using forward proton tagging at the LHC," arXiv:hep-ph/0505240.

[215] M. Livan, V. Vercesi, and R. Wigmans, "Scintillating fiber calorimetry," *CERN Report* **95-02** (1995).

[216] G. Anzivino et al., "Results on lead/scintillating fibres calorimetry," *Nucl. Instrum. and Methods* **A357** (1995) 350–362. doi:10.1016/0168-9002(94)01706-9.

[217] G. Anzivino et al., "Recent developments in quartz fibre calorimetry," *Nucl. Instrum. and Methods* **A357** (1995) 369. doi:10.1016/0168-9002(94)01705-0.

[218] M. Chiu, A. Denisov, E. Garcia, J. Katzy, and S. White, "Measurement of mutual Coulomb dissociation in s(NN)**(1/2) = 130-GeV Au + Au collisions at RHIC," *Phys. Rev. Lett.* **89** (2002) 012302, arXiv:nucl-ex/0109018. doi:10.1103/PhysRevLett.89.012302.

[219] C. Adler et al., "The RHIC zero degree calorimeters," *Nucl. Instrum. and Methods* **A470** (2001) 488–499, arXiv:nucl-ex/0008005. doi:10.1016/S0168-9002(01)00627-1.

[220] W. C. Turner et al., "Status report on the development of instrumentation for bunch by bunch measurement and optimisation of luminosity in the LHC," *LBNL* **45549** (2000).

[221] R. Arnaldi et al., "Performances of zero degree calorimeters for the ALICE experiment," *Nucl. Instrum. and Methods* **A456** (2001) 248–258. doi:10.1016/S0168-9002(00)00586-6.

[222] **PHENIX** Collaboration, A. Denisov, "Common event characterization in the RHIC experiments," *Nucl. Phys.* **A698** (2002) 551–554. doi:10.1016/S0375-9474(01)01426-9.

[223] W. Smith. *private communication*.

[224] M. Arneodo et al., "Diffractive Higgs: CMS/TOTEM Level-1 Trigger Studies," in *Proceedings of the HERA-LHC Workshop*. CERN/DESY, January, 2005. Available at http://www.desy.de/~heralhc/proceedings/wg4arneodo.pdf.

[225] M. Barbero et al., "Design and test of the CMS pixel readout chip," *Nucl. Instrum. and Methods* **A517** (2004) 349. doi:10.1016/j.nima.2003.09.043.

[226] J. Isberg, J. Hammersberg, et al., "High carrier mobility in single-crystal plasma-deposited diamond," *Science* **297** (2002) 1670. doi:10.1126/science.1074374.

[227] **RD42** Collaboration, W. Adam et al., "New developments in CVD diamond for detector applications," *Eur. Phys. J. C* **33, Supplement 1** (2004) S1014. doi:10.1140/epjcd/s2004-03-1798-6.

[228] **RD42** Collaboration, W. Adam et al., "Radiation tolerance of CVD diamond detectors for pions and protons," *Nucl. Instrum. and Methods* **A476** (2002) 686. doi:10.1016/S0168-9002(01)01666-7.

[229] S. Frixione and M. Mangano, "How accurately can we measure the W cross section?," *JHEP* **0405** (2004) 056. doi:10.1088/1126-6708/2004/05/056.

[230] M. Dittmar, F. Pauss, and D. Zürcher, "Towards a precise parton luminosity determination at the CERN LHC," *Phys. Rev.* **D56** (1997) 7284. doi:10.1103/PhysRevD.56.7284.

[231] V. Drollinger, "The Beam Halo Generator Manual with Validation Plots," **2005/004** (2005).

[232] R. Frühwirth, "Application of Kalman Filtering to Track and Vertex Fitting," *Nucl. Instrum. and Methods* **A262** (1987) 444. doi:10.1016/0168-9002(87)90887-4.

[233] V. Innocente, M. Maire, and E. Nagy, "GEANE: Average Tracking and Error Propagation Package," *CERN Program Library, IT-ASD W5013-E* (1991).

[234] I. Azhgirey et al., "Cascade simulations for the machine induced background study in the IR1 of the LHC," *CERN LHC Project Note* **324** (2003).

[235] O. Kodolova, I. Tomalin, and P. Yepes, "Expected Data Rates from the Silicon Strip Tracker," *CMS Note* **2002-047** (2002).

[236] M. Bedjidian et al., "Dimuon Reconstruction in Heavy Ion Collisions Using a Detailed Description of CMS Geometry," *CMS Note* **1999/004** (1999).

[237] E. James, Y. Maravin, and N. Neumeister, "Muon Identification in CMS," *CMS Note* **2006/010** (2006).

[238] N. Amapane, M. Fierro, and M. Konecki, "High Level Trigger Algorithms for Muon Isolation," *CMS Note* **2002/040** (2002).

[239] S. Abdullin et al., "Sensitivity of the muon isolation cut efficiency to the underlying event uncertainties," *CMS Note* **2006/xxx** (2006).

[240] G. Baur et al., "Heavy Ion Physics Programme in CMS," *CMS Note* **2000/060** (2000).

[241] T. C. Awes et al., "A Simple Method of Shower Localization and Identification in Laterally Segmented Calorimeters," *Nucl. Instrum. and Methods* **A311** (1992). doi:10.1016/0168-9002(92)90858-2.

[242] B. Clerbaux et al., "TeV electron and photon saturation studies," *CMS Note* **2006/004** (2006).

[243] M. Pieri et al., "Distinguishing Isolated Photons from Jets," *CMS Note* **2006/007** (2006).

[244] N. Marinelli, "Track finding and identification of converted photons," *CMS Note* **2006/005** (2006).

[245] S. Baffioni et al., "Electron Reconstruction in CMS," *CMS Note* **2006/xxx** (2006).

[246] E. Eichten, I. Hinchliffe, K. D. Lane, and C. Quigg, "Super Collider Physics," *Rev. Mod. Phys.* **56** (1984) 579–707. doi:10.1103/RevModPhys.56.579.

[247] M. Carena and H. E. Haber, "Higgs boson theory and phenomenology," *Prog. Part. Nucl. Phys.* **50** (2003) 63–152, arXiv:hep-ph/0208209. doi:10.1016/S0146-6410(02)00177-1.

[248] J. G. Branson et al., "High transverse momentum physics at the Large Hadron Collider: The ATLAS and CMS Collaborations," *Eur. Phys. J. direct* **C4** (2002) N1.

[249] J. W. Rohlf, "Physics with jets at the LHC," *Acta Phys. Polon.* **B36** (2005) 469–479. Available at http://th-www.if.uj.edu.pl/acta/vol36/pdf/v36p0469.pdf.

[250] C. Bromberg et al., "Observation of the Production of Jets of Particles at High Transverse Momentum and Comparison with Inclusive Single Particle Reactions," *Phys. Rev. Lett.* **38** (1977) 1447. doi:10.1103/PhysRevLett.38.1447.

[251] C. Bromberg et al., "Experimental Tests of Quantum Chromodynamics in High P(T) Jet Production in 200-Gev/c Hadron - Proton Collisions," *Phys. Rev. Lett.* **43** (1979) 565. doi:10.1103/PhysRevLett.43.565.

[252] J. W. Rohlf, "Jet Production in High-Energy Hadron-Proton Collisions". PhD thesis, California Institute of Technology, 1979. FERMILAB-THESIS-1979-20, see also [253].

[253] C. Bromberg et al., "Jet Production in High-Energy Hadron - Proton Collisions," *Nucl. Phys.* **B171** (1980) 1–37. Thesis of J. Rohlf, see [252]. doi:10.1016/0550-3213(80)90357-0.

[254] **UA1** Collaboration, G. Arnison et al., "Angular Distributions and Structure Functions from Two Jet Events at the CERN Sps p anti-p Collider," *Phys. Lett.* **B136** (1984) 294. doi:10.1016/0370-2693(84)91164-X.

[255] **UA1** Collaboration, G. Arnison et al., "Angular Distributions for High Mass Jet Pairs and a Limit on the Energy Scale of Compositeness for Quarks from the CERN p anti-p Collider," *Phys. Lett.* **B177** (1986) 244. doi:10.1016/0370-2693(86)91065-8.

[256] **UA2** Collaboration, P. Bagnaia et al., "Measurement of Very Large Transverse Momentum Jet Production at the CERN anti-p p Collider," *Phys. Lett.* **B138** (1984) 430. doi:10.1016/0370-2693(84)91935-X.

[257] J. Huston, "LHC Guide to Parton Distribution Functions," *ATL-PHYS* **99-008** (1999).

[258] R. Demina et al., "Calorimeter Cell Energy Thresholds for Jet Reconstruction in CMS," *CMS Note* **2006/020** (2006).

[259] **UA1** Collaboration, G. Arnison et al., "Hadronic Jet Production at the CERN Proton - Anti-Proton Collider," *Phys. Lett.* **B132** (1983) 214. doi:10.1016/0370-2693(83)90254-X.

[260] S. V. Chekanov, "Jet algorithms: A mini review," arXiv:hep-ph/0211298.

[261] G. C. Blazey et al., "Run II jet physics," arXiv:hep-ex/0005012.

[262] J. M. Butterworth, J. P. Couchman, B. E. Cox, and B. M. Waugh, "KtJet: A C++ implementation of the K(T) clustering algorithm," *Comput. Phys. Commun.* **153** (2003) 85–96, arXiv:hep-ph/0210022. doi:10.1016/S0010-4655(03)00156-5.

[263] S. D. Ellis and D. E. Soper, "Successive combination jet algorithm for hadron collisions," *Phys. Rev.* **D48** (1993) 3160–3166, arXiv:hep-ph/9305266. doi:10.1103/PhysRevD.48.3160.

[264] A. Heister et al., "Measurement of Jets with the CMS Detector at the LHC," *CMS Note* **2006/xxx** (2006).

[265] O. Kodolova, "Jet Energy Measurements in CMS," *CMS CR* **2005-019** (2005). Presented at HCP 2005, Hadron Collider Physics Symposium, Les Diablerets, Switzerland, 4–6 July.

[266] A. Astbury et al., "A 4π solid-angle detector for the SPS used as a proton-antiproton collider at a centre-of-mass energy of 540 GeV," *CERN/SPSC* **78-06** (1978).

[267] G. Arnison et al., "Experimental Observation of Isolated Large Transverse Energy Electrons With Associated Missing Energy at $\sqrt{s}$ = 540 GeV," *Phys. Lett.* **122B** (1983) 103–116. doi:10.1016/0370-2693(83)91177-2.

[268] G. Arnison et al., "Events With Large Missing Transverse Energy at the CERN Collider: I. W $\to \tau\nu$ Decay and Test of $\tau - \mu - \mathrm{e}$ Universality at $Q^2 = m_{\mathrm{W}^2}$," *Phys. Lett.* **185B** (1987) 233–240. doi:10.1016/0370-2693(87)91561-9.

[269] G. Arnison et al., "Events With Large Missing Transverse Energy at the CERN Collider: II. Search for the Decays of $\mathrm{W}^{\pm}$ Into Heavy Leptons and of Z^0 Into Non-Interacting Particles," *Phys. Lett.* **185B** (1987) 241–248. doi:10.1016/0370-2693(87)91562-0.

[270] G. Arnison et al., "Events With Large Missing Transverse Energy at the CERN Collider: III. Mass Limits on Supersymmetric Particles," *Phys. Lett.* **198B** (1987) 261–270. doi:10.1016/0370-2693(87)91509-7.

[271] **D0** Collaboration, S. Abachi et al., "Search for squarks and gluinos in p anti-p collisions at S**(1/2) = 1.8-TeV," *Phys. Rev. Lett.* **75** (1995) 618–623. doi:10.1103/PhysRevLett.75.618.

[272] **CDF** Collaboration, F. Abe et al., "Search for gluinos and squarks at the Fermilab Tevatron collider," *Phys. Rev.* **D56** (1997) 1357–1362. doi:10.1103/PhysRevD.56.R1357.

[273] M. Spiropulu, "A blind search for supersymmetry in p anti-p collisions at s**(1/2) = 1.8-TeV using the missing energy plus multijet channel". PhD thesis, Harvard University, 2000. FERMILAB-THESIS-2000-16.

[274] H. Haifeng et al., "Measurement of Missing Transverse Energy with the CMS Detector at the LHC," *CMS Note* **2006/xxx** (2006).

[275] F. Abe et al., "Measurement of the W-boson mass," *Phys. Rev. Lett.* **65** (1990) 2243. doi:10.1103/PhysRevLett.65.2243.

[276] J. Adelman et al., "Determination of the jet energy scale at the Collider Detector at Fermilab," arXiv:hep-ex/0510047.

[277] **D0** Collaboration, B. Abbott et al., "Determination of the absolute jet energy scale in the D0 calorimeters," *Nucl. Instrum. and Methods* **A424** (1999) 352–394. doi:10.1016/S0168-9002(98)01368-0.

[278] V. F. Konoplyanikov et al., "Study of the background for the 'Gamma+Jet' Process in the CMS Detector at Low Luminosity," *Phys. Part. Nucl. Lett.* **2** (2005) 45–53.

[279] J. D'Hondt, S. Lowette, and S. Kasselmann, "Light quark jet energy scale calibration using the W mass constraint in single-leptonic ttbar events," *CMS Note* **2006/025** (2006).

[280] J. D'Hondt, S. Lowette, and J. Heynick, "Electron and muon reconstruction in single leptonic $t\bar{t}$ events," *CMS Note* **2006/024** (2006).

[281] **ZEUS** Collaboration, J. Breitweg et al., "Measurement of jet shapes in photoproduction at HERA," *Eur. Phys. J.* **C2** (1998) 61–75, arXiv:hep-ex/9710002. doi:10.1007/s100500050091.

[282] **ALEPH** Collaboration, D. Buskulic et al., "Performance of the ALEPH detector at LEP," *Nucl. Instrum. and Methods* **A360** (1995) 481–506. doi:10.1016/0168-9002(95)00138-7.

[283] D. Green et al., "Energy Flow Objects and Usage of Tracks for Energy Measurement in CMS," *CMS Note* **2002/036** (2002).

[284] O. Kodolova et al., "Jet Energy Correction with Charged Particle Tracks in CMS," *Eur. Phys. J.* **C40 S1** (2005) 33–42. doi:10.1140/epjcd/s2005-02-004-2.

[285] O. L. Kodolova, "Measurements in Very High Multiplicity Conditions," *Phys. Part. Nucl.* **35** (2004) 675–691.

[286] G. Bagliesi et al., "Tau jet reconstruction and tagging at High Level Trigger and off-line," *CMS Note* **2006/xxx** (2006).

[287] S. Eno, S. Dasu, W. Smith, R. Kinnunen, and A. Nikitenko, "A Study of a First and Second Level Tau Trigger," *CMS Note* **2000-055** (2000).

[288] D. Denegri, R. Kinnunen, and A. Nikitenko, "Study of Calorimeter Calibration with Tau's in CMS," *CMS Note* **1997-039** (1997).

[289] R. Kinnunen and A. Nikitenko, "Study of $H \to \tau\tau$ with Hadronic Tau Decays in CMS," *CMS Note* **2003-006** (2003).

[290] D. Kotlinski, A. Nikitenko, and R. Kinnunen, "Study of a Level-3 Tau Trigger with the Pixel Detector," *CMS Note* **2001-017** (2001).

[291] G. Bagliesi, S. Gennai, and G. Sguazzoni, "A L2 Trigger for Tau Hadronic Decays with Tracker Isolation in the Low Luminosity Scenario," *CMS Note* **2002-018** (2002).

[292] S. Cucciarelli, M. Konecki, D. Kotlinski, and T. Todorov, "Track-Parameter Evaluation and Primary-Vertex Finding with the Pixel Detector," *CMS Note* **2003-026** (2003).

[293] R. Kinnunen, S. Lehti, F. Moortgat, A. Nikitenko, and M. Spira, "Measurement of the $H/A \to \tau\tau$ cross section and possible constraints on $\tan\beta$," *CMS Note* **2004-027** (2004).

[294] A. Rizzi, F. Palla, and G. Segneri, "Track impact parameter based b-tagging with CMS," *CMS Note* **2006/019** (2006).

[295] C. Weiser, "A Combined Secondary Vertex Based B-Tagging Algorithm in CMS," *CMS Note* **2006/014** (2006).

[296] A. Bocci, P. Demin, et al., "Tagging b jets with electrons and muons at CMS," *CMS Note* **2006/xxx** (2006).

[297] M. D. Richard and R. P. Lippman, "Neural Network Classifiers Estimates Bayesian a posteriori Probabilities," *Neural Computation* **3** (1991), no. 4, 461–483.

[298] C. M. Bishop, "Neural Networks for Pattern Recognition". Oxford University Press, 1995.

[299] M. Vos and F. Palla, "B-tagging in the High Level Trigger," *CMS Note* **2006/xxx** (2006).

[300] S. Lowette, J. D'Hondt, J. Heyninck, and P. Vanlaer, "Offline Calibration of b-Jet Identification Efficiencies," *CMS Note* **2006/013** (2006).

[301] J. D'Hondt et al., "Fitting of Event Topologies with External Kinematic Constraints in CMS," *CMS Note* **2006/023** (2006).

[302] P. Bartalini, R. Chierici, and A. De Roeck, "Guidelines for the Estimation of Theoretical Uncertainties at the LHC," *CMS Note* **2005-013** (2005).

Colour Plates

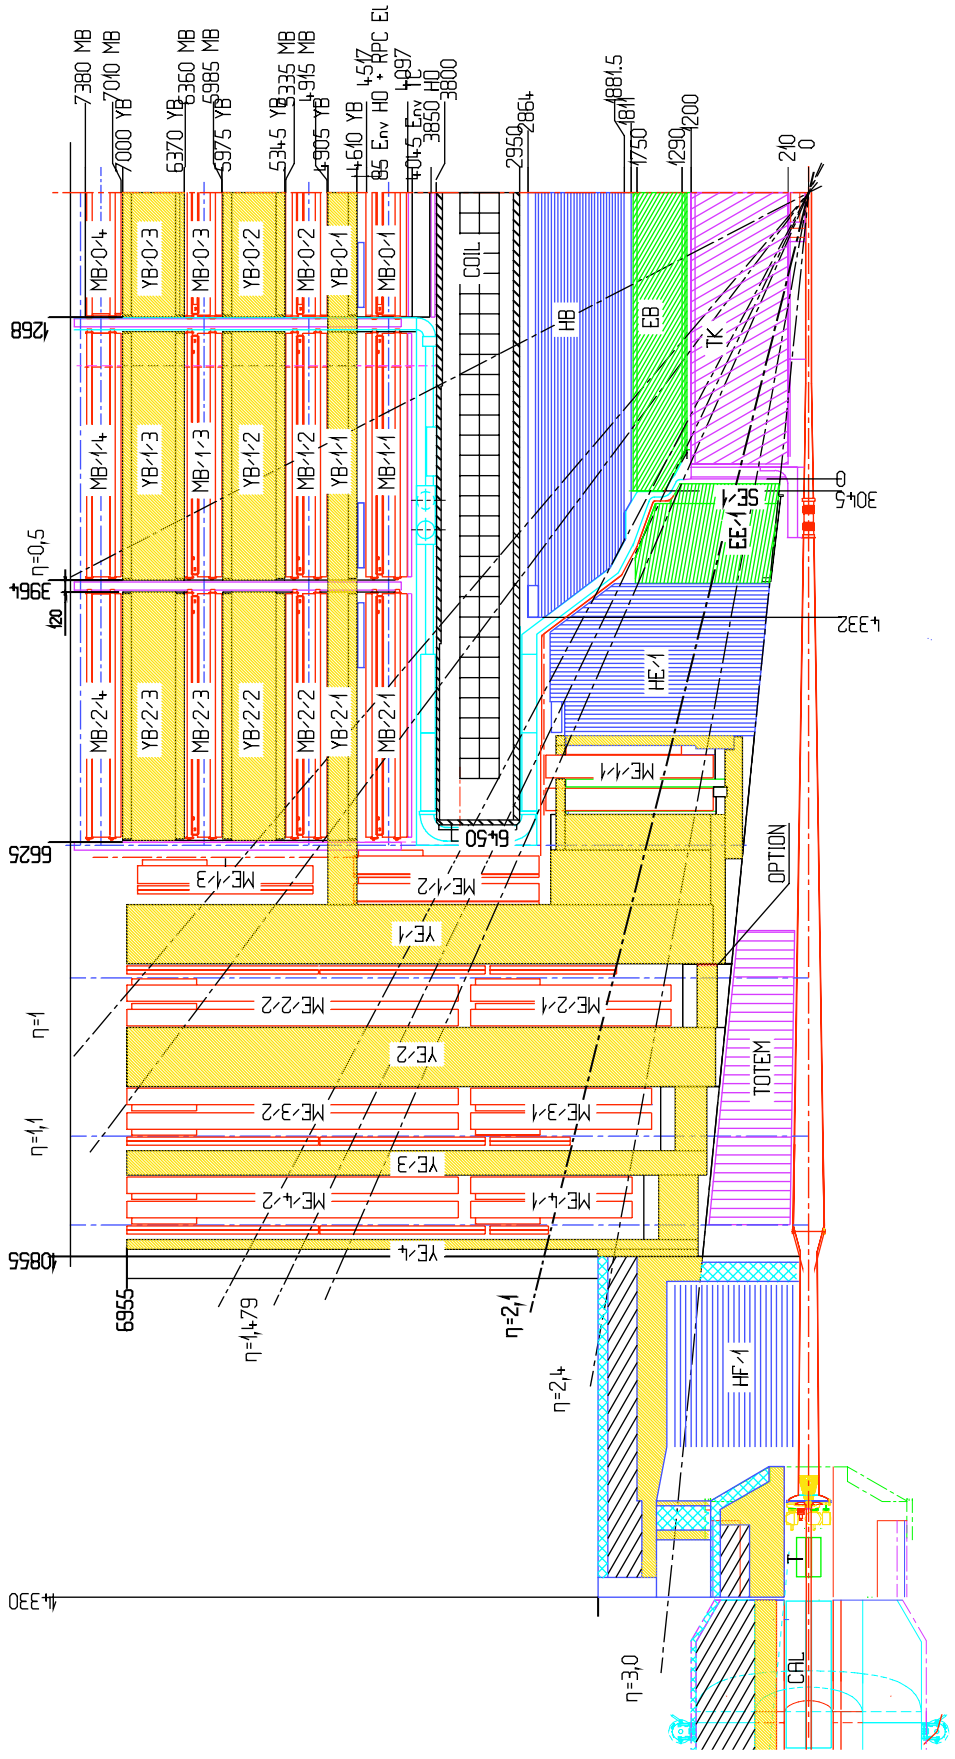


Figure CP 1: One quarter longitudinal view of the CMS Experiment. Dimensions are in units of mm.

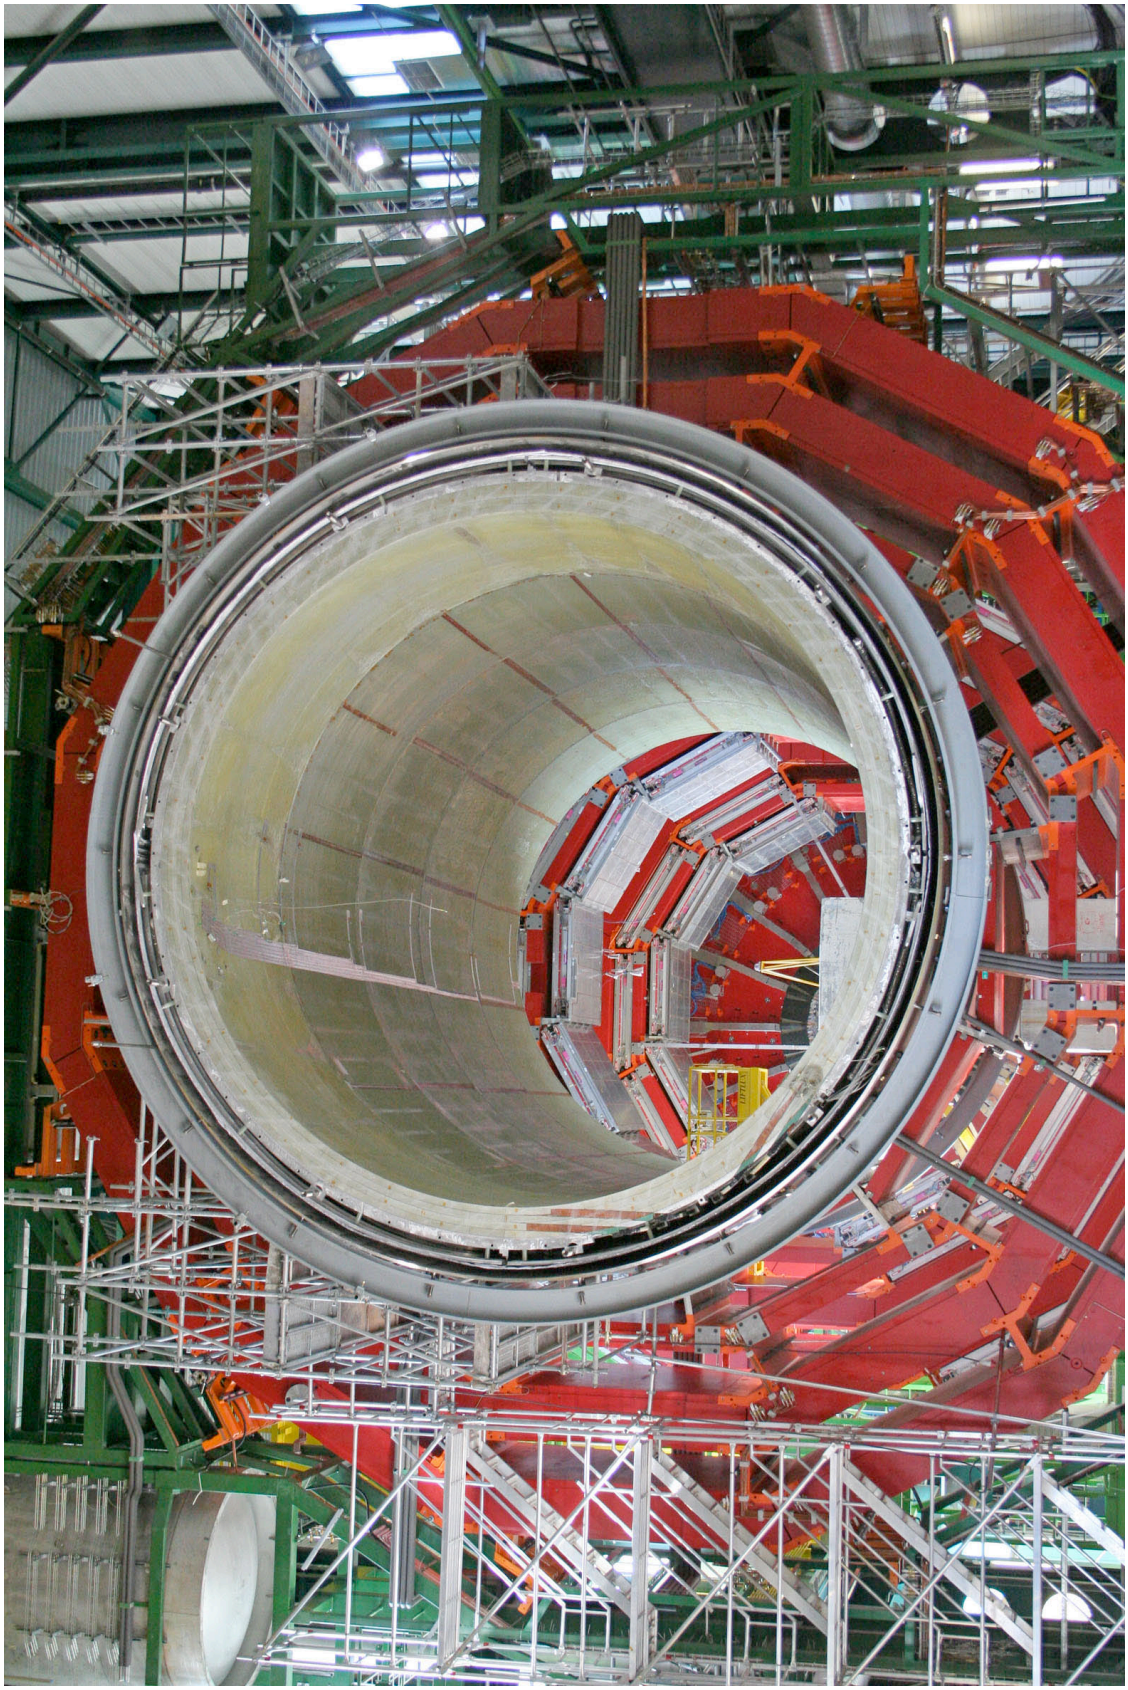


Figure CP 2: Assembly of the yoke of the CMS magnet. The outer vacuum tank and the coil can be seen.

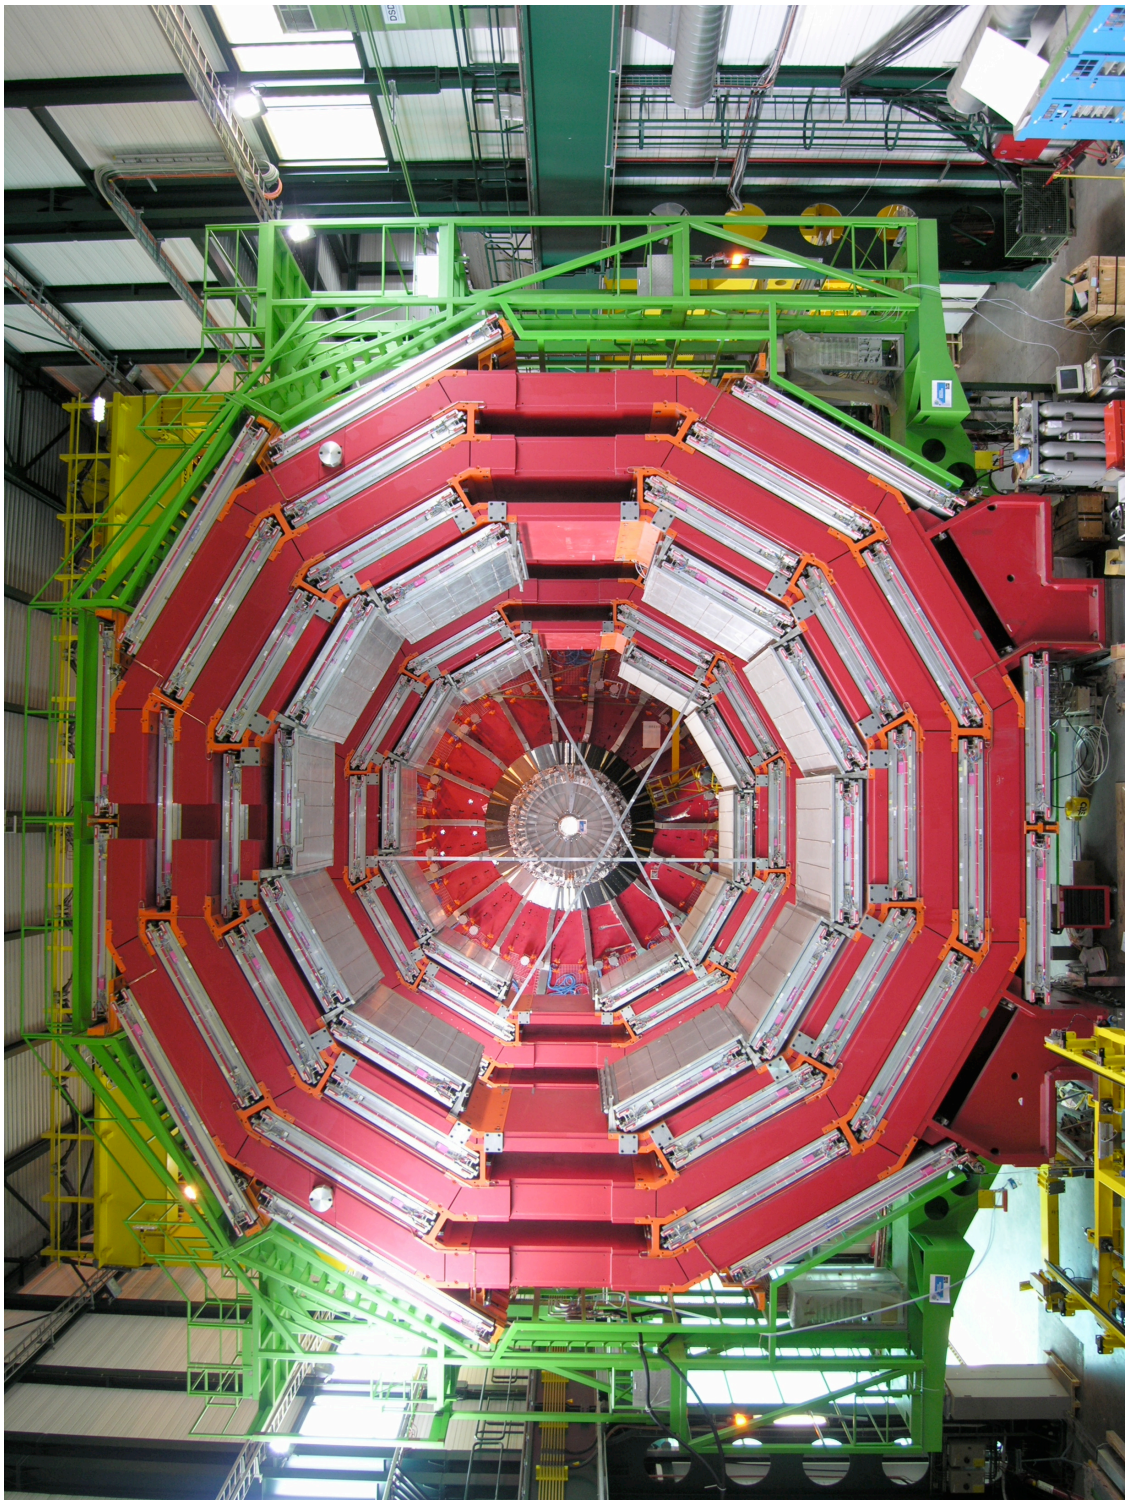


Figure CP 3: Installation of DT chambers into the YB+1 wheel.

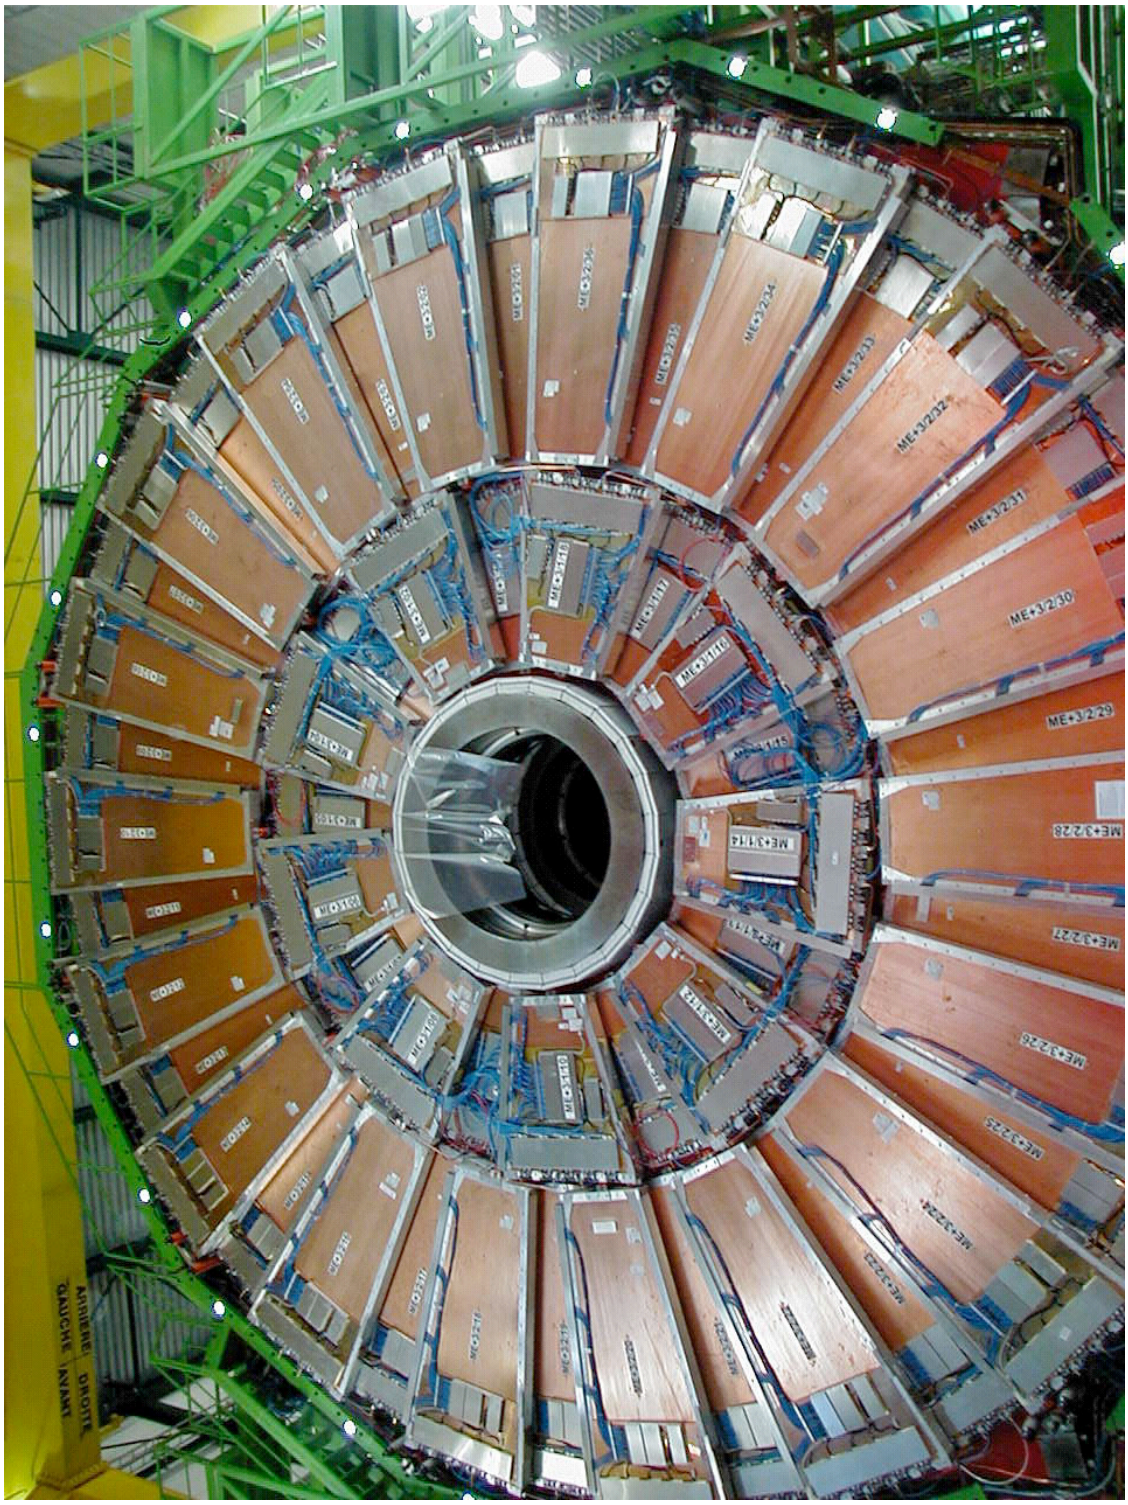


Figure CP 4: Installation of CSC chambers onto the YE+2 disk.

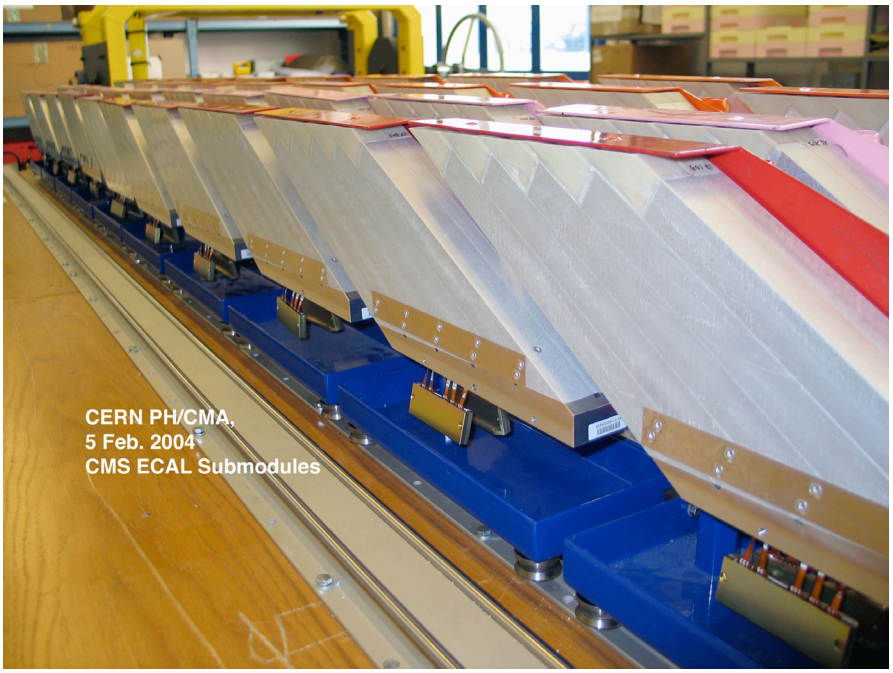


Figure CP 5: Submodules of the ECAL barrel in assemby jigs. Each submodule consists of 10 crystals in a glass-fibre alveola structure. Two APDs are glued to the rear of each crystal and the flexible kapton ribbons from the APDs, terminated with connectors, can be seen hanging down from the back of the submodules.



Figure CP 6: The 60 mm thick aluminium backplate of an ECAL endcap Dee held on its assembly and installation frame. Four mock supercrystals (5×5 crystals) are shown attached in position along with all 138 positional spacers.

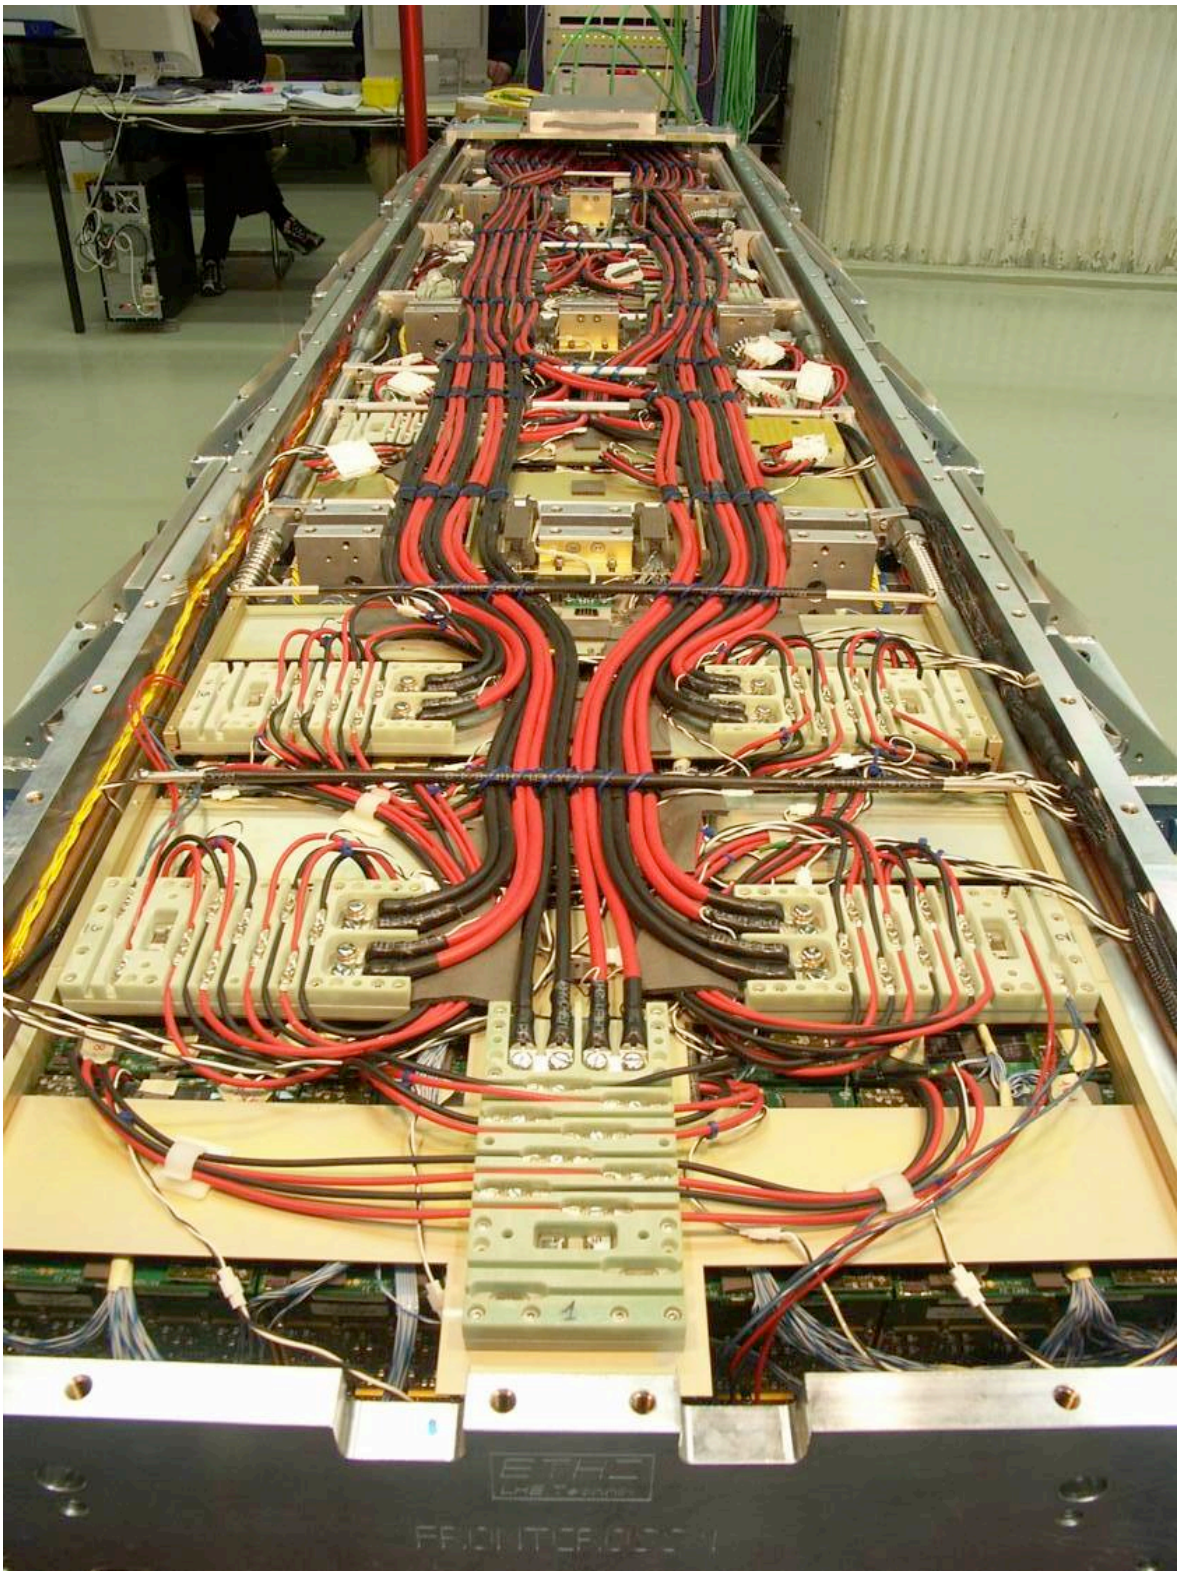


Figure CP 7: The back of an ECAL supermodule fully equipped with cooling pipes and electronics, before installation of the closing backplate.



Figure CP 8: Assembled HB in the surface hall at SX5.

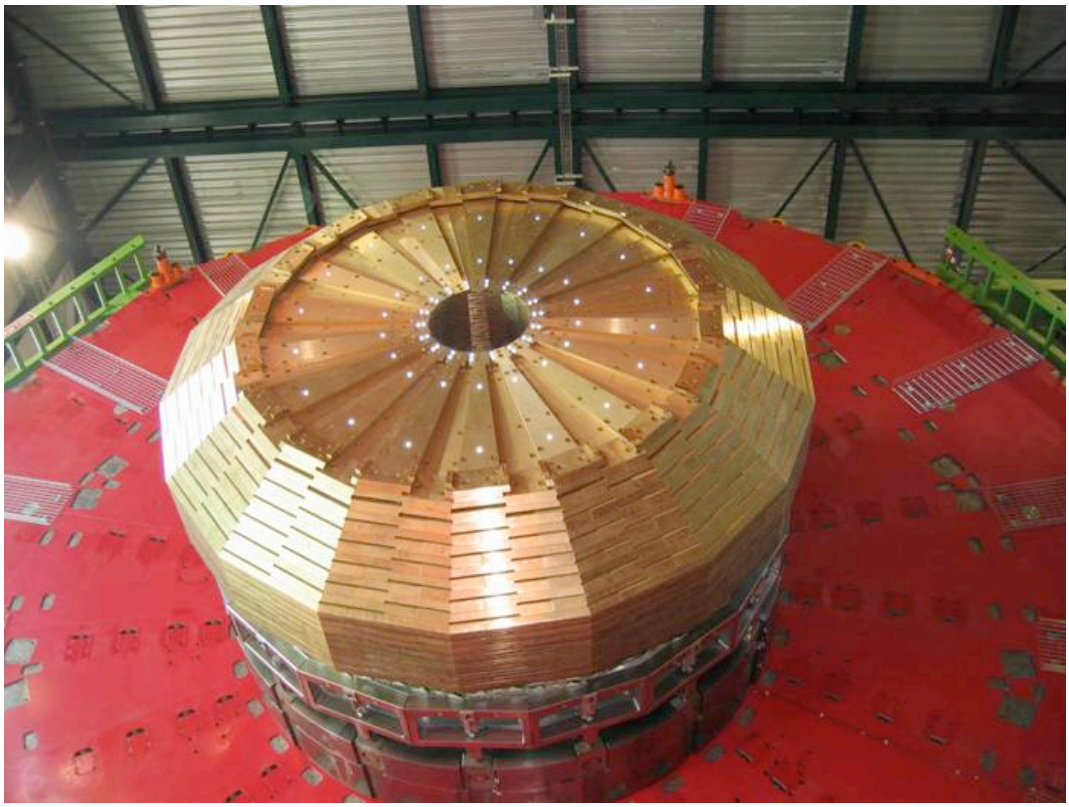


Figure CP 9: Assembled HE in the surface hall at SX5.



Figure CP 10: Assembled HF modules in building 186 at CERN.

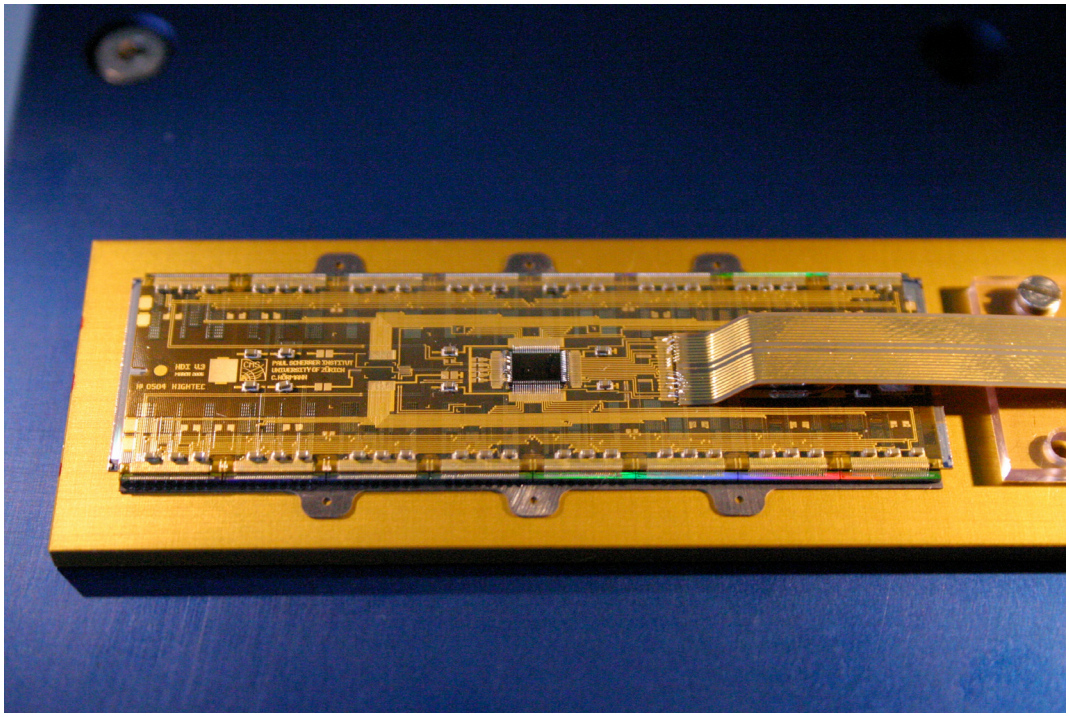


Figure CP 11: Photo of pixel module

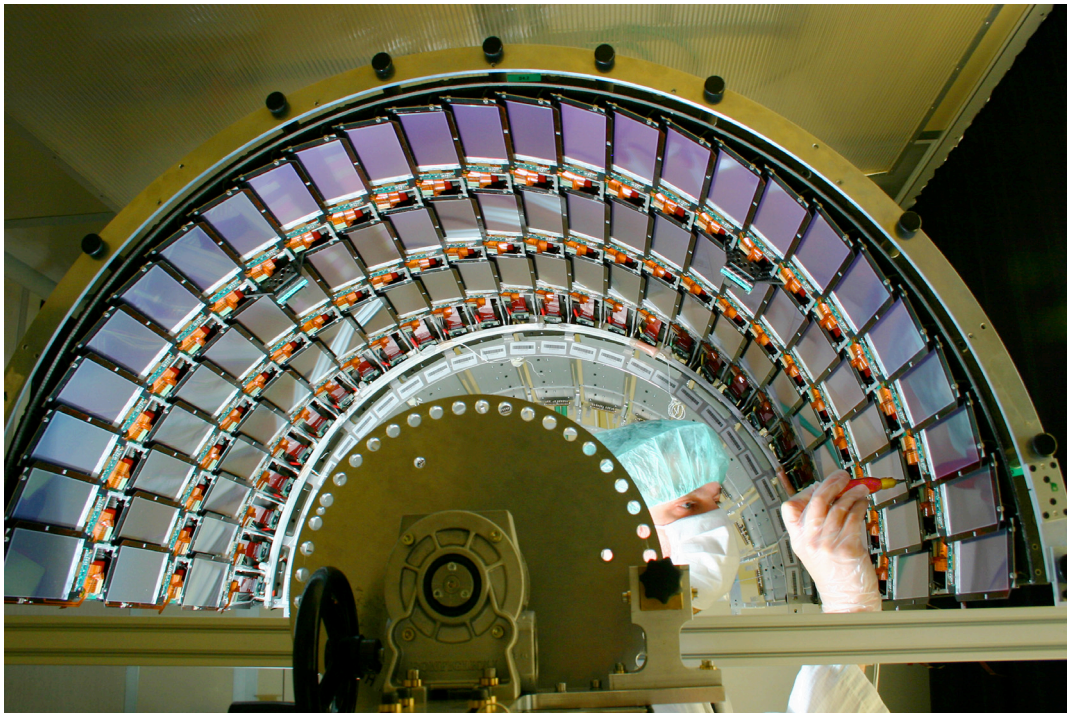


Figure CP 12: Half of the Layer 4+ shell for the Tracker Inner Barrel.

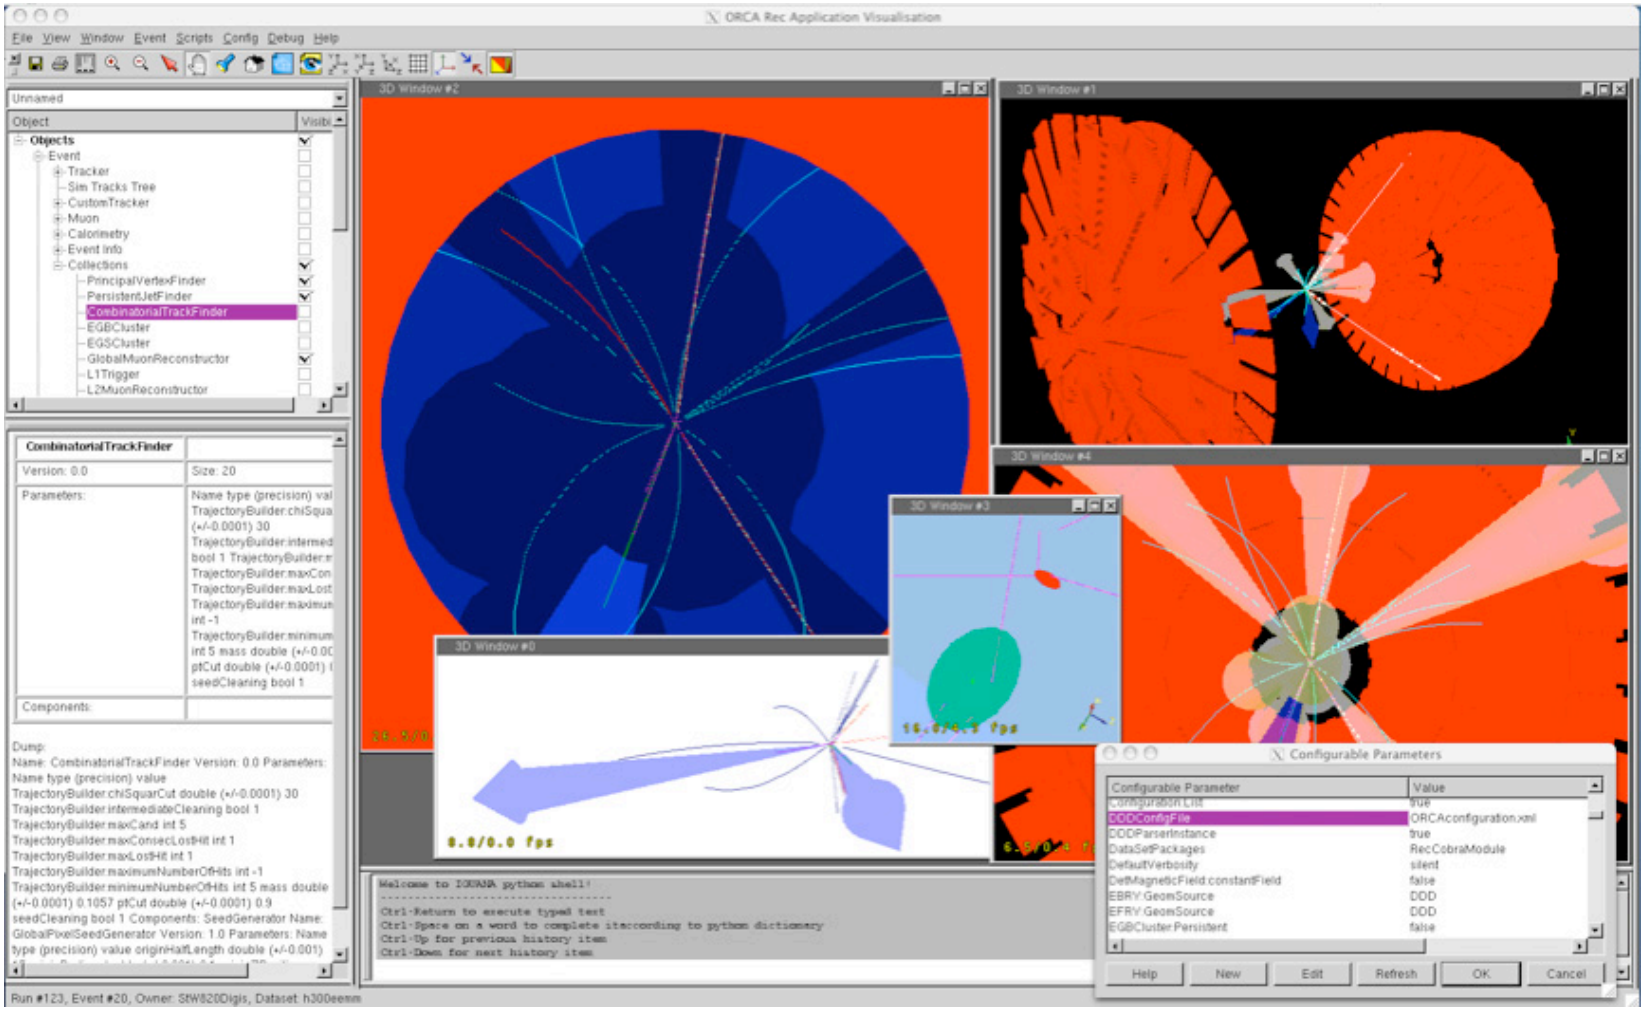


Figure CP 13: Screen shot of the interactive IGUANA-based CMS detector and event display, showing reconstructed objects such as tracks, jets, and vertices in 2D and 3D views.

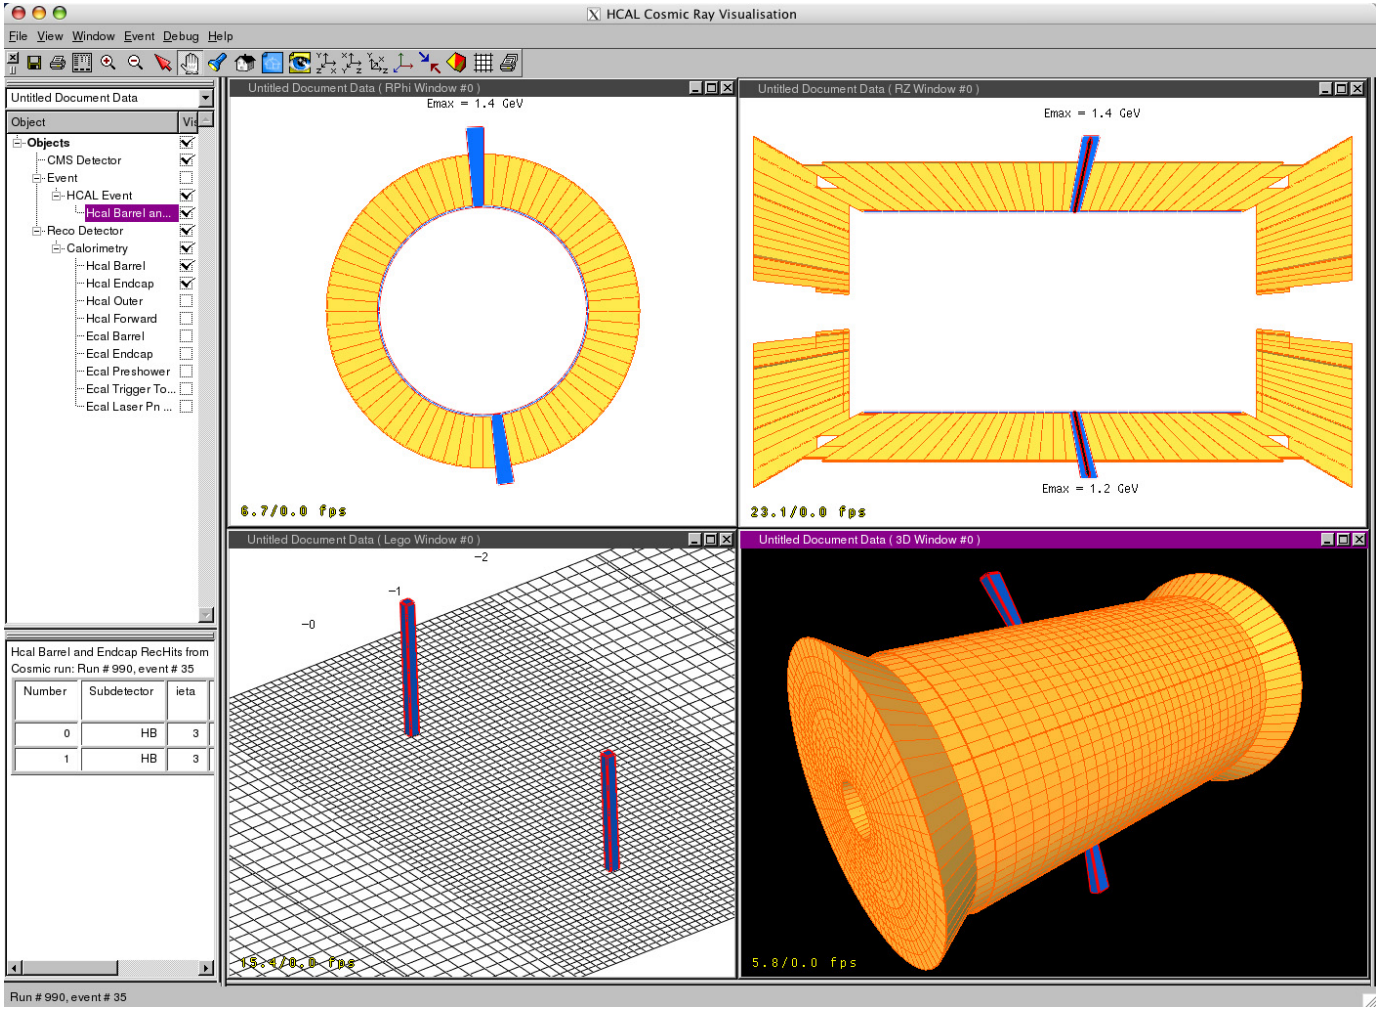


Figure CP 14: Screen shot of the interactive IGUANA-based detector and event display deployed at SX5. The cosmic ray data has been taken in the HCAL half-barrel, then calibrated and reconstructed. The sub-windows show 2D and 3D views of the detector and a 2D tower histogram ("lego plot") of deposited energy.