

OPEN ACCESS

Upgrade of the ALICE Experiment: Letter Of Intent

To cite this article: B Abelev et al and (The ALICE Collaboration) 2014 *J. Phys. G: Nucl. Part. Phys.* **41** 087001


View the [article online](#) for updates and enhancements.

You may also like

- [Overview of results from ALICE](#)
Mateusz Posko and the ALICE Collaboration

- [The ALICE pixel detector upgrade](#)
F. Reidt

- [Design and implementation of detector control system for muon forward tracker at ALICE](#)
K. Yamakawa, A. Augustinus, G. Batigne et al.



CERN-LHCC-2012-012 / LHCC-I-022
September 8, 2012

Upgrade of the ALICE Experiment Letter Of Intent

The ALICE Collaboration

0.1 The ALICE Collaboration

B. Abelev⁷¹, J. Adam³⁷, D. Adamová⁷⁸, A.M. Adare¹²⁹, M.M. Aggarwal⁸², G. Aglieri Rinella³³, M. Agnello^{102, 88}, A.G. Agocs¹²⁸, A. Agostinelli²⁷, Z. Ahmed¹²⁴, N. Ahmad¹⁷, A. Ahmad Masoodi¹⁷, S.A. Ahn⁶⁴, S.U. Ahn^{40, 64}, I. Aimo²², M. Ajaz¹⁵, A. Akindinov⁵⁰, D. Aleksandrov⁹⁴, B. Alessandro¹⁰², A. Alici^{98, 12}, A. Alkin³, E. Almaráz Aviña⁶⁰, J. Alme³⁵, T. Alh³⁹, V. Altin³¹, S. Altinpinar¹⁸, I. Altsybeev¹²⁵, C. Andrei⁷⁴, A. Andronic⁹¹, V. Anguelov⁸⁷, J. Anielski⁵⁸, C. Anson¹⁹, T. Antić⁹², F. Antinori⁹⁹, P. Antonioli⁹⁸, L. Aphecetche¹⁰⁸, H. Appelshäuser⁵⁶, N. Arbor⁶⁷, S. Arcelli²⁷, A. Arend⁵⁶, N. Armesto¹⁶, R. Arnaldi¹⁰², T. Aronsson¹²⁹, I.C. Arsene⁹¹, M. Arslandok⁵⁶, A. Asryan¹²⁵, A. Augustinus³³, R. Averbeck⁹¹, T.C. Awes⁷⁹, J. Åystö⁴², M.D. Azmi^{17, 84}, M. Bach³⁹, A. Badalà¹⁰⁵, Y.W. Baek^{66, 40}, R. Bailhache⁵⁶, R. Bala^{85, 102}, R. Baldini Ferroli¹², A. Baldissari¹⁴, F. Baltasar Dos Santos Pedrosa³³, J. Bán⁵¹, R.C. Baral⁵², R. Barbera²⁶, F. Barile³¹, G.G. Barnaföldi¹²⁸, L.S. Barnby⁹⁶, V. Barret⁶⁶, J. Bartke¹¹², M. Basile²⁷, N. Bastid⁶⁶, S. Basu¹²⁴, B. Bathen⁵⁸, G. Battighe¹⁰⁸, M. Battistin³³, B. Batyunya⁶², J. Baudot⁶¹, C. Baumann⁵⁰, R. Bavontaweepanya¹⁰⁹, I.G. Bearden⁷⁶, H. Beck⁵⁶, N.K. Behera⁴⁴, I. Belikov⁶¹, F. Bellini²⁷, R. Bellwied¹¹⁸, E. Belmont-Moreno⁶⁰, G. Benedi¹²⁸, S. Beole²², I. Berceanu⁷⁴, A. Beruci⁷⁴, Y. Berdnikov⁸⁰, D. Berenyi¹²⁸, M. Berger¹¹¹, A.A.E. Bergognon¹⁰⁸, D. Berzana^{22, 102}, L. Betev³³, A. Bhasin⁸⁵, A.K. Bhati⁸², J. Bhom¹²², L. Bianchi²², N. Bianchi⁶⁸, J. Bielćik³⁷, J. Bielćiková⁷⁸, A. Bilandžić⁷⁶, S. Bjelogrlic⁴⁹, F. Blanco¹¹⁸, F. Blanco¹⁰, D. Blau⁹⁴, C. Blume⁵⁶, M. Boccioli³³, S. Böttger⁵⁵, A. Bogdanov⁷², H. Bøggild⁷⁶, M. Bogolyubsky⁴⁷, L. Boldizsár¹²⁸, M. Bomba³⁸, J. Book⁵⁶, H. Borel¹⁴, A. Borissov¹²⁷, F. Bossu⁸⁴, C. Bortolin³³, J.A. Botelho Direito³³, M. Botje⁷⁷, E. Botta²², E. Braidot⁷⁰, P. Braun-Munzinger⁹¹, M. Bregant¹⁰⁸, T. Breitner⁵⁵, T.A. Broker⁵⁶, T.A. Browning⁸⁹, M. Broz³⁶, R. Brun³³, E. Bruna^{22, 102}, G.E. Bruno³¹, D. Budnikov⁹³, H. Buesching⁵⁶, S. Bufalino^{22, 102}, P. Buncic³³, O. Busch⁸⁷, Z. Buthelezi⁸⁴, D. Caballero Orduna¹²⁹, D. Caffarni^{28, 99}, X. Cai⁷, H. Caines¹²⁹, E. Calvo Villar⁹⁷, P. Camerini²⁴, V. Canoa Roman¹¹, G. Cara Romeo⁹⁸, W. Carena³³, F. Carena³³, N. Carlin Filho¹¹⁵, F. Carminati³³, A. Casanova Díaz⁶⁸, J. Castillo Castellanos¹⁴, J.F. Castillo Hernandez⁹¹, E.A.R. Casula²³, V. Catanescu⁷⁴, T. Caudron³³, C. Cavigchioli³³, C. Ceballos Sanchez⁹, J. Cepila³⁷, P. Cerello¹⁰², B. Chang^{12, 131}, N. Chankhunthot¹⁰⁹, S. Chapeland³³, J.L. Charvet¹⁴, S. Chattopadhyay⁹⁵, S. Chattopadhyay¹²⁴, I. Chawla⁸², M. Cherney⁸¹, C. Cheshkov^{33, 117}, B. Cheynis¹¹⁷, V. Chibante Barroso³³, D.D. Chinellato¹¹⁸, P. Chochula³³, M. Chojnacki^{76, 49}, S. Choudhury¹²⁴, P. Christakoglou⁷⁷, C.H. Christensen⁷⁶, P. Christiansen³², T. Chujo¹²², S.U. Chung⁹⁰, C. Cicali¹⁰¹, L. Cifarelli^{27, 33, 12}, F. Cindolo⁹⁸, J. Cleymans⁸⁴, F. Cocetti¹², F. Colamaria³¹, D. Colella³¹, A. Collu²³, G. Conesa Balbastre⁶⁷, Z. Conesa del Valle³³, M.E. Connors¹²⁹, G. Contin²⁴, J.G. Contreras¹¹, T.M. Cormier¹²⁷, Y. Corrales Morales²², P. Cortese³⁰, I. Cortés Maldonado², M.R. Cosentino⁷⁰, F. Costa³³, M.E. Cotallo¹⁰, E. Crescio¹¹, P. Crocher⁶⁶, E. Cruz Alaniz⁶⁰, R. Cruz Albino¹¹, E. Cuaute⁵⁹, L. Cunqueiro⁶⁸, A. Dainese^{28, 99}, H.H. Dalsgaard⁷⁶, A. Danu⁵⁴, E. Da Riva³³, I. Das⁴⁶, D. Das⁹⁵, S. Das⁴, K. Das⁹⁵, A. Dash¹¹⁶, S. Dash⁴⁴, S. De¹²⁴, G.O.V. de Barros¹¹⁵, A. De Caro^{29, 12}, G. de Cataldo¹⁰⁴, C. Decosse³³, J. de Cuveland³⁹, A. De Falco²³, D. De Gruttola²⁹, H. Delagrange¹⁰⁸, A. Deloff⁷³, N. De Marco¹⁰², E. Dénes¹²⁸, S. De Pasquale²⁹, A. Deppman¹¹⁵, G. D'Erasmo³¹, R. de Rooij⁴⁹, M.A. Diaz Corchero¹⁰, D. Di Bari³¹, T. Dietel⁵⁸, C. Di Giglio³¹, S. Di Liberto¹⁰⁰, A. Di Mauro³³, P. Di Nezza⁶⁸, R. Divià³³, Ø. Djurvsland¹⁸, A. Dobrin^{127, 32}, T. Dobrowolski⁷³, B. Döningus⁹¹, O. Dordic²¹, O. Driga¹⁰⁸, A.K. Dubey¹²⁴, A. Dubla⁴⁹, L. Duxour¹¹⁷, W. Dulinski⁶¹, P. Dupieux⁶⁶, A.K. Dutta Majumdar⁹⁵, D. Elia¹⁰⁴, D. Emschermann⁵⁸, H. Engel⁵⁵, B. Erazmus^{33, 108}, H.A. Erdal³⁵, B. Espagnon⁴⁶, M. Estienne¹⁰⁸, S. Esumi¹²², D. Evans⁹⁶, G. Eyyubova²¹, L. Fabbietti¹¹¹, D. Fabris^{28, 99}, J. Faivre⁶⁷, D. Falchieri²⁷, A. Fantoni⁶⁸, M. Fasel^{91, 87}, R. Fearick⁸⁴, D. Fehlker¹⁸, L. Feldkamp⁵⁸, D. Feleai⁵⁴, A. Feliciello¹⁰², B. Fenton-Olsen⁷⁰, G. Feofilov¹²⁵, J. Ferenczi³⁷, A. Fernández Téllez², A. Ferretti²², A. Festanti²⁸, J. Figiel¹¹², M.A.S. Figueiredo¹¹⁵, S. Filchagin⁹³, D. Finogeev⁴⁸, F.M. Fionda³¹, E.M. Fiore³¹, E. Floratos⁸³, M. Floris³³, S. Foertsch⁸⁴, P. Foka⁹¹, S. Fokin⁹⁴, E. Fragiacomo¹⁰³, A. Francescon^{33, 28}, U. Frankenfeld⁹¹, U. Fuchs³³, C. Furet⁶⁷, M. Fusco Girard²⁹, J.J. Gaardhøje⁷⁶, M. Gagliardi²², A. Gago⁹⁷, M. Gallio²², D.R. Gangadharan¹⁹, P. Ganot⁷⁹, C. Garabatos⁹¹, E. Garcia-Solis¹³, C. Gargiulo³³, I. Garishvili⁷¹, P. Gasik¹¹¹, J. Gerhard³⁹, M. Germain¹⁰⁸, C. Geuna¹⁴, M. Gheata^{54, 33}, A. Gheata³³, B. Ghidini³¹, P. Ghosh¹²⁴, P. Gianotti⁶⁸, M.R. Girard¹²⁶, P. Giubellino³³, E. Gladysz-Dziadus¹¹², P. Glässel⁸⁷, M. Goffe⁶¹, R. Gomez^{114, 11}, M. Gomez Marzoa³³, E.G. Ferreiro¹⁶, L.H. González-Trueba⁶⁰, P. González-Zamora¹⁰, S. Gorbunov³⁹, A. Goswami⁸⁶, S. Gotovac¹¹⁰, L.K. Graczykowski¹²⁶, R. Grajcarek⁸⁷, A. Grelli⁴⁹, C. Grigoras³³, A. Grigoras³³, V. Grigoriev⁷², A. Grigoryan¹, S. Grigoryan⁶², B. Grinyov³, N. Grion¹⁰³, P. Gros³², J.F. Grosse-Oetringhaus³³, J.-Y. Grossiord¹¹⁷, R. Grossiord³³, F. Guerrier⁴⁸, R. Guernane⁶⁷, B. Guerzoni²⁷, M. Guibalbaud¹¹⁷, K. Gulbrandsen⁷⁶, H. Gulkanyan¹, T. Gunji¹²¹, A. Gupta⁸⁵, R. Gupta⁸⁵, R. Haake⁵⁸, Ø. Haaland¹⁸, C. Hadjidakis⁴⁶, M. Haiduc⁵⁴, H. Hamagaki¹²¹, G. Hamar¹²⁸, B.H. Han²⁰,



Content from this work may be used under the terms of the Creative Commons Attribution 3.0

licence. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

L.D. Hanratty⁹⁶, A. Hansen⁷⁶, Z. Harmanová-Tóthová³⁸, J.W. Harris¹²⁹, M. Hartig⁵⁶, A. Harton¹³, D. Hatzifotiadou⁹⁸, S. Hayashi¹²¹, A. Hayrapetyan³³, S.T. Heckel⁵⁶, M. Heide³⁸, H. Helstrup³⁵, A. Herghelegiu⁷⁴, G. Herrera Corral¹¹, N. Herrmann⁸⁷, B.A. Hess¹²³, K.F. Hetland³⁵, B. Hicks¹²⁹, H. Hillemanns³³, B. Hippolyte⁶¹, A. Hoenle¹¹¹, Y. Hori¹²¹, P. Hristov³³, I. Hřivnáčová⁴⁶, C. Hu⁶¹, M. Huang¹⁸, T.J. Humanic¹⁹, D.S. Hwang²⁰, R. Ichou⁶⁶, S. Igolkin³³, P. Ijzermans³³, R. Ilkaev⁹³, I. Ilkiv⁷³, M. Inaba¹²², E. Incani²³, P.G. Innocenti³³, G.M. Innocenti²², M. Ippolito⁹⁴, M. Irfan¹⁷, C. Ivan⁹¹, V. Ivanov⁸⁰, A. Ivanov¹²⁵, M. Ivanov⁹¹, O. Ivanytskyi³, A. Jacholkowski²⁶, P. M. Jacobs⁷⁰, H.J. Jang⁶⁴, M.A. Janik¹²⁶, R. Janik³⁶, P.H.S.Y. Jayarathna¹¹⁸, S. Jena⁴⁴, D.M. Jha¹²⁷, R.T. Jimenez Bustamante⁵⁹, P.G. Jones⁹⁶, H. Jung⁴⁰, A. Junque³³, A. Jusko⁹⁶, A.B. Kaidalov⁵⁰, S. Kalcher³⁹, P. Kalináč⁵¹, T. Kallioski⁴², A. Kalweit⁵⁷, J.H. Kang¹³¹, V. Kaplin⁷², A. Karasu Uysal³³, O. Karavichev⁴⁸, T. Karavicheva⁴⁸, E. Karpechev⁴⁸, A. Kazantsev⁹⁴, U. Kebschull⁵⁵, R. Keidel¹³², B. Kettner¹¹¹, P. Khan⁹⁵, S.A. Khan¹²⁴, M.M. Khan¹⁷, K. H. Khan¹⁵, A. Khanzadeev⁸⁰, Y. Kharlov⁴⁷, B. Kileng³⁵, B. Kim¹³¹, J.S. Kim⁴⁰, J.H. Kim²⁰, D.J. Kim⁴², D.W. Kim⁴⁰, T. Kim¹³¹, S. Kim²⁰, M. Kim⁴⁰, M. Kim¹³¹, S. Kirsch³⁹, I. Kisel³⁹, S. Kiselev⁵⁰, A. Kisel¹²⁶, J.L. Klay⁶, J. Klein⁸⁷, C. Klein-Bösing³⁸, M. Kliemann⁵⁰, A. Kluge³³, M.L. Knichel⁹¹, A.G. Knospe¹¹³, M.K. Köhler⁹¹, C. Kobda¹⁰⁹, T. Kollegger³⁹, A. Kolojvari¹²⁵, M. Kompaniets¹²⁵, V. Kondratyeva⁷², A. Konevskikh⁴⁸, V. Kovalenko¹²⁵, M. Kowalski¹¹², S. Kox⁶⁷, G. Koyithatta Meethaleveedu⁴⁴, J. Kral⁴², I. Králik⁵¹, F. Kramer⁵⁶, A. Kravčáková³⁸, T. Krawutschke⁸⁷, M. Krelina³⁷, M. Kretz³⁹, M. Krivda⁹⁶, M. Krizek⁴², M. Krus³⁷, E. Kryshen⁸⁰, M. Krzewicki⁹¹, Y. Kuchereria⁹⁴, T. Kugathasan³³, C. Kuhn⁶¹, P.G. Kuijer⁷⁷, I. Kulakov⁵⁶, J. Kumar⁴⁴, P. Kurashvili⁷³, A. Kurepin⁴⁸, A.B. Kurepin⁴⁸, A. Kuryakin⁹³, S. Kushpil⁷⁸, V. Kushpil⁷⁸, H. Kvaerno³¹, M.J. Kweon⁸⁷, Y. Kwon¹³¹, P. Ladrón de Guevara⁵⁹, I. Lakomov⁴⁶, R. Langoy¹⁸, S.L. La Pointe⁴⁹, C. Lara⁵⁵, A. Lardeux¹⁰⁸, P. La Rocca²⁶, R. Lea²⁴, M. Lechman³³, K.S. Lee⁴⁴, S.C. Lee⁴⁰, G.R. Lee⁹⁶, I. Legrand³³, J. Lehnert⁵⁶, R. Lemmon¹⁰⁶, M. Lenhard⁹¹, V. Lenti¹⁰⁴, H. León⁶⁰, I. León Monzón¹¹⁴, H. León Vargas⁵⁶, Y. Lesenechal³³, P. Léval¹²⁸, S. Li⁷, J. Lien¹⁸, R. Lietava⁹⁶, S. Lindal²¹, V. Lindenstruth³⁹, C. Lippmann⁹¹, M.A. Lisa¹⁹, H.M. Ljunggren³², P.I. Loenna¹⁸, V.R. Loggins¹²⁷, V. Loguinov⁷², D. Lohner⁸⁷, C. Loizides⁷⁰, K.K. Loo⁴², X. Lopez⁶⁶, E. López Torres⁹, G. Løvhøiden²¹, X.-G. Lu⁸⁷, P. Luettig⁵⁶, M. Lunardon²⁸, J. Luo⁷, G. Luparello⁴⁹, C. Luzzi³³, R. Ma¹²⁹, K. Ma⁷, D.M. Madagodahettige-Don¹¹⁸, A. Maevskaia⁴⁸, M. Mager⁵⁷, D.P. Mahapatra⁵², A. Maire⁸⁷, M. Malaev⁸⁰, I. Maldonado Cervantes⁵⁹, L. Malinina⁶², A. Mal'Kevich³⁰, P. Malzache⁹¹, A. Mamontov⁹³, L. Manceau¹⁰², L. Mangotra⁸⁵, V. Manko⁹⁴, F. Manso⁶⁶, C. Mansuy³³, V. Manzari¹⁰⁴, Y. Mao⁷, A. Mapelli³³, M. Marchisone⁶⁶, J. Mares⁵³, G.V. Margagliotti²⁴, A. Margotti⁹⁸, A. Marin⁹¹, C. Markert¹¹³, M. Marquardt⁵⁶, D. Marras²³, I. Martashvili¹²⁰, N.A. Martin⁹¹, P. Martinengo³³, M.I. Martínez², A. Martínez Davalos⁶⁰, G. Martínez García¹⁰⁸, Y. Martynov³, A. Mas¹⁰⁸, S. Masciocchi⁹¹, M. Masera²², A. Masoni¹⁰¹, L. Massacrier¹⁰⁸, A. Mastroserio³¹, A. Matyja¹¹², C. Mayer¹¹², J. Mazer¹²⁰, G. Mazza²², M.A. Mazzoni¹⁰⁰, F. Meddi²⁵, A. Menchaca-Rocha⁶⁰, J. Mercado Pérez⁸⁷, M. Meres³⁶, Y. Miake¹²², L. Milano²², J. Milosevic²¹, A. Mischke⁴⁹, A.N. Mishra⁸⁰, D. Mišković⁹¹, C. Mitu⁵⁴, S. Mizuno¹²², J. Mlynarz¹²⁷, B. Mohanty¹²⁴, L. Molnar¹²⁸, L. Montaño Zetina¹¹, M. Monteno¹⁰², E. Montes¹⁰, T. Moon¹³¹, M. Morando²⁸, D.A. Moreira De Godoy¹¹⁵, S. Moretto²⁸, A. Morreale⁴², A. Morsch³³, V. Muccifora⁶⁸, E. Mudnic¹¹⁰, S. Muñuri¹²⁴, M. Mukherjee¹²⁴, H. Müller³³, M.G. Munhoz¹¹⁵, S. Murray⁸⁴, L. Musa³³, J. Musinsky⁵¹, A. Musso¹⁰², B.K. Nandi⁴⁴, R. Nania⁹⁸, E. Nappi¹⁰⁴, C. Natrass¹²⁰, T.K. Nayak¹²⁴, S. Nazarenko⁹³, A. Nedosekin⁵⁰, M. Nicassio³¹, M. Niculescu⁵⁴, B.S. Nielsen⁷⁶, T. Niida¹²², S. Nikolaev⁹⁴, V. Nikolic⁹², S. Nikulin⁹⁴, V. Nikulin⁸⁰, B.S. Nilsen⁸¹, M.S. Nilsson²¹, F. Noferini⁹⁸, P. Nomokonov⁶², G. Nooren⁴⁹, N. Novitzky⁴², A. Nyanin⁹⁴, A. Nyatha⁴⁴, C. Nygaard⁷⁶, J. Nystrand¹⁸, A. Ochirow¹²⁵, H. Oeschler⁵⁷, S. Oh¹²⁹, S.K. Oh⁴⁰, J. Oleniacz¹²⁶, A.C. Oliveira Da Silva¹¹⁵, C. Oppedisan¹⁰², A. Ortiz Velasquez³², A. Oskarsson³², P. Ostrowski¹²⁶, J. Otwinowski⁹¹, K. Oyama⁸⁷, K. Ozawa¹²¹, Y. Pachmayer⁸⁷, M. Pachr³⁷, F. Padilla²², P. Paganò²⁹, G. Paic³⁹, F. Painke³⁹, C. Pajares¹⁶, S.K. Pal¹²⁴, A. Palaha⁹⁶, A. Palmeri¹⁰⁵, V. Papikyan¹, G.S. Pappalardo¹⁰⁵, W.J. Park⁹¹, A. Passfeld⁵⁸, B. Pastirák⁵¹, C. Pastore¹⁰⁴, D.I. Patalakha⁴⁷, V. Paticchio¹⁰⁴, B. Paul⁹⁵, A. Pavlinov¹²⁷, T. Pawlak¹²⁶, T. Peitzmann⁴⁹, H. Pereira Da Costa¹⁴, E. Pereira De Oliveira Filho¹¹⁵, D. Peresunko⁹⁴, C.E. Pérez Lara⁷⁷, D. Perini³³, D. Perrino³¹, W. Peryt¹²⁶, A. Pesci⁹⁸, V. Peskov³³, Y. Pestov⁵, P. Petagna³³, V. Petráček³⁷, M. Petran³⁷, M. Petris⁷⁴, P. Petrov⁹⁶, M. Petrovici⁷⁴, C. Petta²⁶, S. Piano¹⁰³, M. Pikna³⁶, P. Pillot¹⁰⁸, O. Pinazza³³, L. Pinsky¹¹⁸, N. Pitz⁵⁶, D.B. Piyarathna¹¹⁸, M. Planinic⁹², M. Ploskon⁷⁰, J. Pluta¹²⁶, T. Pocheptsov⁶², S. Pochybova¹²⁸, P.L.M. Podesta-Lerma¹⁴, M.G. Poghosyan³³, K. Polák⁵³, B. Polichtchouk⁴⁷, W. Poonsawat¹⁰⁹, A. Pop⁷⁴, S. Porteboeuf-Houssais⁶⁶, V. Pospišil³⁷, B. Potukuchi⁸⁵, S.K. Prasad¹²⁷, R. Pregenella⁹⁸, F. Prino¹⁰², C.A. Pruneau¹²⁷, I. Pshenichnov⁴⁸, G. Puddu²³, C. Puggioni²³, V. Punin⁹³, M. Putiš³⁸, J. Putschke¹²⁷, E. Quercigh³³, H. Qvigstad²¹, A. Rachevski¹⁰³, A. Rademakers³³, T.S. Rähä⁴², J. Rak⁴²,

A. Rakotozafindrabe¹⁴, L. Ramello³⁰, A. Ramírez Reyes¹¹, R. Raniwala⁸⁶, S. Raniwala⁸⁶, S.S. Räsänen⁴², B.T. Rascanu⁵⁶, D. Rathee⁸², K.F. Read¹²⁰, J.S. Real⁶⁷, K. Redlich⁷⁵,^{iv}, R.J. Reed¹²⁹, A. Rehman¹⁸, P. Reichelt⁵⁶, M. Reicher⁴⁹, R. Renfordt⁵⁶, A.R. Reolon⁶⁸, A. Reshetin⁴⁸, F. Rettig³⁹, J.-P. Revol³³, K. Reygers⁸⁷, L. Riccati¹⁰², R.A. Ricci⁶⁹, T. Richert³², M. Richter²¹, P. Riedler³³, W. Riegler³³, F. Rigg²⁶,¹⁰⁵, A. Rivetti²², M. Rodriguez Cahantzi², A. Rodriguez Manso⁷⁷, K. Røed¹⁸,²¹, D. Rohr³⁹, D. Röhrl¹⁸, R. Romita⁹¹,¹⁰⁶, F. Ronchetti⁶⁸, P. Rosnet⁶⁶, S. Rossegger³³, A. Rossi³³,²⁸, C. Roy⁶¹, P. Roy⁹⁵, A.J. Rubio Montero¹⁰, R. Rui²⁴, R. Russo²², E. Ryabinkin⁹⁴, A. Rybicki¹¹², S. Sadovsky⁴⁷, K. Šafařík³³, R. Sahoo⁴⁵, P.K. Sahu⁵², J. Saini¹²⁴, H. Sakaguchi⁴³, S. Sakai⁷⁰, D. Sakata¹²², A.B. Saleem¹⁵, C.A. Salgado¹⁶, J. Salzwedel¹⁹, S. Sambyal⁸⁵, V. Samsonov⁸⁰, X. Sanchez Castro⁶¹, L. Sándor⁵¹, A. Sandoval⁶⁰, M. Sano¹²², G. Santagati²⁶, R. Santoro³³,¹², J. Sarkamo⁴², E. Scapparone⁹⁸, F. Scarlassara²⁸, R.P. Scharenberg⁸⁹, C. Schiaua⁷⁴, R. Schicker⁸⁷, C. Schmidt⁹¹, H.R. Schmidt¹²³, S. Schuchmann⁵⁶, J. Schukraft³³, T. Schuster¹²⁹, Y. Schutz³³,¹⁰⁸, K. Schwarz⁹¹, K. Schwindt⁹¹, G. Scioli²⁷, E. Scomparin¹⁰², P.A. Scott⁹⁶, R. Scott¹²⁰, G. Segato²⁸, I. Selyuzhenkov⁹¹, S. Senyukov⁶¹, J. Seo⁹⁰, S. Seri²³, E. Serradilla¹⁰,⁶⁰, A. Sevcenco³⁴, A. Shabetai¹⁰⁸, G. Shabratova⁶², R. Shahoyan³³, N. Sharma⁸²,¹²⁰, S. Sharma⁸⁵, S. Rohm⁸⁵, I. Sgura¹⁰⁴, K. Shigaki⁴³, K. Shitejer⁹, Y. Sibirtiat⁹⁴, E. Sicking⁵⁸, S. Siddhanta¹⁰¹, T. Siemianczuk⁷³, D. Silvermyr⁷⁹, C. Silvestre⁶⁷, G. Simatovic⁵⁹,⁹², G. Simonetti³³, R. Singaraju¹²⁴, R. Singh⁸³, S. Singha¹²⁴,⁷⁵, V. Singhal¹²⁴, B.C. Sinha¹²⁴, T. Sinha⁹⁵, B. Sitar³⁶, M. Sitta³⁰, T.B. Skaal²¹, K. Skjerdal¹⁸, R. Smakal³⁷, N. Smirnov¹²⁹, R.J.M. Snellings⁴⁹, W. Snoeys³³, C. Søgaard⁷⁶,³², R. Soltz⁷¹, H. Son²⁰, J. Song⁹⁰, M. Song¹³¹, C. Soos³³, F. Soramel²⁸, I. Sputowska¹¹², M. Spyropoulou-Stassinakis⁸³, B.K. Srivastava⁸⁹, J. Stachel⁸⁷, I. Stan⁵⁴, I. Stan⁵⁴, G. Stefanek⁷³, M. Steinpreis¹⁹, E. Stenlund³², G. Steyn⁸⁴, J.H. Stiller⁸⁷, D. Stocco¹⁰⁸, M. Stolpovskiy⁴⁷, P. Strmen³⁶, A.A.P. Suade¹¹⁵, M.A. Subieta Vásquez²², T. Sugitate⁴³, C. Suire⁴⁶, R. Sultanov⁵⁰, M. Šumbera⁷⁸, X. Sun⁷, T. Susa⁹², T.J.M. Symons⁷⁰, A. Szanto de Toledo¹¹⁵, I. Szarka³⁶, A. Szczepankiewicz¹¹²,³³, A. Szostak¹⁸, M. Szymański¹²⁶, J. Takahashi¹¹⁶, J.D. Tapia Takaki⁴⁶, A. Tarantola Pelon⁵⁶, A. Tarazona Martínez³³, A. Tauro³³, G. Tejeda Muñoz², A. Telesca³³, C. Terrevoli³¹, J. Thäder⁹¹, D. Thomas⁴⁹, R. Tieulent¹¹⁷, A.R. Timmins¹¹⁸, D. Tlusty³⁷, C. Tobon Marin³³, A. Toia³⁹,²⁸,⁹⁹, H. Torii¹²¹, L. Toscano¹⁰², V. Trubnikov³, D. Truesdale¹⁹, W.H. Trzaska⁴², T. Tsuji¹²¹, A. Tumkin⁹³, R. Turchetta¹⁰⁷, R. Turrisi⁹⁹, T.S. Twetei²¹, J. Ulery⁵⁶, K. Ullaland¹⁸, J. Ulrich⁶³,⁵⁵, A. Uras¹¹⁷, J. Urbán³⁸, G.M. Urciuoli¹⁰⁰, G.L. Usai²³, M. Vajzer³⁷,⁷⁸, M. Vala⁶²,⁵¹, L. Valencia Palomo⁴⁶, S. Vallerio⁸⁷, J. Van Beelen³³, P. Vande Vyvre³³, J. Van Hoornen³³, M. van Leeuwen⁴⁹, L. Vannucci⁶⁹, A. Vargas², R. Varma⁴⁴, M. Vasileiou⁸³, A. Vasilev⁹⁴, V. Vechernin¹²⁵, M. Veldhoen⁴⁹, M. Venaruzzo²⁴, E. Vercellin²², S. Vergara², R. Vernet⁸, M. Verweij⁴⁹, L. Vickovic¹¹⁰, G. Viesti²⁸, J. Viinikainen⁴², Z. Vilakazi⁸⁴, O. Villalobos Baillie⁹⁶, Y. Vinogradov⁹³, A. Vinogradov⁹⁴, L. Vinogradov¹²⁵, T. Virgili²⁹, Y.P. Viyogi¹²⁴, A. Vodopyanov⁶², S. Voloshin¹²⁷, K. Voloshin⁵⁰, G. Volpe³³, B. von Haller³³, I. Vorobyev¹²⁵, D. Vranic⁹¹, J. Vrláková³⁸, B. Vulpescu⁶⁶, A. Vyushin⁹³, B. Wagner¹⁸, V. Wagner³⁷, R. Wan⁷, D. Wang⁷, Y. Wang⁸⁷, M. Wang⁷, D. Wang⁷, K. Watanabe¹²², M. Weber¹¹⁸, J.P. Wessels³³,⁵⁸, U. Westerhoff⁵⁸, J. Wiechula¹²³, J. Wikne²¹, M. Wilde⁵⁸, G. Wilk⁷³, A. Wilk⁵⁸, M.C.S. Williams⁹⁸, B. Windelband⁸⁷, M. Winter⁶¹, L. Xaplanteris Karampatos¹¹³, C.G. Yaldo¹²⁷, Y. Yamaguchi¹²¹, H. Yang¹⁴,⁴⁹, S. Yang¹⁸, P. Yang⁷, S. Yasnopolskiy⁹⁴, J. Yi⁹⁰, Z. Yin⁷, I.-K. Yoo⁹⁰, J. Yoon¹³¹, W. Yu⁵⁶, X. Yuan⁷, I. Yushmanov⁹⁴, V. Zaccolo⁷⁶, C. Zach³⁷, C. Zampolli⁹⁸, S. Zaporozhets⁶², A. Zarochentsev¹²⁵, P. Závada⁵³, N. Zaviyalov⁹³, H. Zbroszczyk¹²⁶, P. Zelnicek⁵⁵, I.S. Zgura⁵⁴, M. Zhalov⁸⁰, H. Zhang⁷, X. Zhang⁷⁰,⁶⁶,⁷, F. Zhou⁷, Y. Zhou⁴⁹, D. Zhou⁷, H. Zhu⁷, J. Zhu⁷, X. Zhu⁷, A. Zichichi²⁷,¹², A. Zimmermann⁸⁷, G. Zinovjev³, Y. Zoccarato¹¹⁷, M. Zynov'yev³, M. Zyzak⁵⁶

Affiliation notes

ⁱ Deceased

ⁱⁱ Also at: M.V.Lomonosov Moscow State University, D.V.Skobeltsyn Institute of Nuclear Physics, Moscow, Russia

ⁱⁱⁱ Also at: University of Belgrade, Faculty of Physics and "Vinča" Institute of Nuclear Sciences, Belgrade, Serbia

^{iv} Also at: Institute of Theoretical Physics, University of Wroclaw, Wroclaw, Poland

Collaboration Institutes

¹ A. I. Alikhanyan National Science Laboratory (Yerevan Physics Institute) Foundation, Yerevan, Armenia

² Benemérita Universidad Autónoma de Puebla, Puebla, Mexico

³ Bogolyubov Institute for Theoretical Physics, Kiev, Ukraine

⁴ Bose Institute, Department of Physics and Centre for Astroparticle Physics and Space Science (CAPSS), Kolkata, India

⁵ Budker Institute for Nuclear Physics, Novosibirsk, Russia

⁶ California Polytechnic State University, San Luis Obispo, California, United States

⁷ Central China Normal University, Wuhan, China

⁸ Centre de Calcul de l'IN2P3, Villeurbanne, France

⁹ Centro de Aplicaciones Tecnológicas y Desarrollo Nuclear (CEADEN), Havana, Cuba

¹⁰ Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

¹¹ Centro de Investigación y de Estudios Avanzados (CINVESTAV), Mexico City and Mérida, Mexico

¹² Centro Fermi - Museo Storico della Fisica e Centro Studi e Ricerche "Enrico Fermi", Rome, Italy

¹³ Chicago State University, Chicago, United States

¹⁴ Commissariat à l'Energie Atomique, IRFU, Saclay, France

¹⁵ COMSATS Institute of Information Technology (CIIT), Islamabad, Pakistan

¹⁶ Departamento de Física de Partículas and IGFAE, Universidad de Santiago de Compostela, Santiago de Compostela, Spain

¹⁷ Department of Physics Aligarh Muslim University, Aligarh, India

¹⁸ Department of Physics and Technology, University of Bergen, Bergen, Norway

¹⁹ Department of Physics, Ohio State University, Columbus, Ohio, United States

²⁰ Department of Physics, Sejong University, Seoul, South Korea

²¹ Department of Physics, University of Oslo, Oslo, Norway

²² Dipartimento di Fisica dell'Università e Sezione INFN, Turin, Italy

²³ Dipartimento di Fisica dell'Università e Sezione INFN, Cagliari, Italy

²⁴ Dipartimento di Fisica dell'Università e Sezione INFN, Trieste, Italy

²⁵ Dipartimento di Fisica dell'Università 'La Sapienza' and Sezione INFN, Rome, Italy

²⁶ Dipartimento di Fisica e Astronomia dell'Università and Sezione INFN, Catania, Italy

²⁷ Dipartimento di Fisica e Astronomia dell'Università and Sezione INFN, Bologna, Italy

²⁸ Dipartimento di Fisica e Astronomia dell'Università and Sezione INFN, Padova, Italy

²⁹ Dipartimento di Fisica 'E.R. Caianiello' dell'Università and Gruppo Collegato INFN, Salerno, Italy

³⁰ Dipartimento di Scienze e Innovazione Tecnologica dell'Università del Piemonte Orientale and Gruppo Collegato INFN, Alessandria, Italy

³¹ Dipartimento Interateneo di Fisica 'M. Merlin' and Sezione INFN, Bari, Italy

³² Division of Experimental High Energy Physics, University of Lund, Lund, Sweden

³³ European Organization for Nuclear Research (CERN), Geneva, Switzerland

³⁴ Fachhochschule Köln, Köln, Germany

³⁵ Faculty of Engineering, Bergen University College, Bergen, Norway

³⁶ Faculty of Mathematics, Physics and Informatics, Comenius University, Bratislava, Slovakia

³⁷ Faculty of Nuclear Sciences and Physical Engineering, Czech Technical University in Prague, Prague, Czech Republic

³⁸ Faculty of Science, P.J. Šafárik University, Košice, Slovakia

³⁹ Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe-Universität Frankfurt, Frankfurt, Germany

⁴⁰ Gangneung-Wonju National University, Gangneung, South Korea

⁴¹ Gauhati University, Department of Physics, Guwahati, India

⁴² Helsinki Institute of Physics (HIP) and University of Jyväskylä, Jyväskylä, Finland

⁴³ Hiroshima University, Hiroshima, Japan

⁴⁴ Indian Institute of Technology Bombay (IIT), Mumbai, India

⁴⁵ Indian Institute of Technology Indore, Indore, India (IITI)

⁴⁶ Institut de Physique Nucléaire d'Orsay (IPNO), Université Paris-Sud, CNRS-IN2P3, Orsay, France

⁴⁷ Institute for High Energy Physics, Protvino, Russia

⁴⁸ Institute for Nuclear Research, Academy of Sciences, Moscow, Russia

⁴⁹ Nikhef, National Institute for Subatomic Physics and Institute for Subatomic Physics of Utrecht University, Utrecht, Netherlands

⁵⁰ Institute for Theoretical and Experimental Physics, Moscow, Russia

⁵¹ Institute of Experimental Physics, Slovak Academy of Sciences, Košice, Slovakia

⁵² Institute of Physics, Bhubaneswar, India

⁵³ Institute of Physics, Academy of Sciences of the Czech Republic, Prague, Czech Republic

⁵⁴ Institute of Space Sciences (ISS), Bucharest, Romania
⁵⁵ Institut für Informatik, Johann Wolfgang Goethe-Universität Frankfurt, Frankfurt, Germany
⁵⁶ Institut für Kernphysik, Johann Wolfgang Goethe-Universität Frankfurt, Frankfurt, Germany
⁵⁷ Institut für Kernphysik, Technische Universität Darmstadt, Darmstadt, Germany
⁵⁸ Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, Münster, Germany
⁵⁹ Instituto de Ciencias Nucleares, Universidad Nacional Autónoma de México, Mexico City, Mexico
⁶⁰ Instituto de Física, Universidad Nacional Autónoma de México, Mexico City, Mexico
⁶¹ Institut Pluridisciplinaire Hubert Curien (IPHC), Université de Strasbourg, CNRS-IN2P3, Strasbourg, France
⁶² Joint Institute for Nuclear Research (JINR), Dubna, Russia
⁶³ Kirchhoff-Institut für Physik, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
⁶⁴ Korea Institute of Science and Technology Information, Daejeon, South Korea
⁶⁵ KTO Karatay University, Konya, Turkey
⁶⁶ Laboratoire de Physique Corpusculaire (LPC), Clermont Université, Université Blaise Pascal, CNRS-IN2P3, Clermont-Ferrand, France
⁶⁷ Laboratoire de Physique Subatomique et de Cosmologie (LPSC), Université Joseph Fourier, CNRS-IN2P3, Institut Polytechnique de Grenoble, Grenoble, France
⁶⁸ Laboratori Nazionali di Frascati, INFN, Frascati, Italy
⁶⁹ Laboratori Nazionali di Legnaro, INFN, Legnaro, Italy
⁷⁰ Lawrence Berkeley National Laboratory, Berkeley, California, United States
⁷¹ Lawrence Livermore National Laboratory, Livermore, California, United States
⁷² Moscow Engineering Physics Institute, Moscow, Russia
⁷³ National Centre for Nuclear Studies, Warsaw, Poland
⁷⁴ National Institute for Physics and Nuclear Engineering, Bucharest, Romania
⁷⁵ National Institute of Science Education and Research, Bhubaneswar, India
⁷⁶ Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark
⁷⁷ Nikhef, National Institute for Subatomic Physics, Amsterdam, Netherlands
⁷⁸ Nuclear Physics Institute, Academy of Sciences of the Czech Republic, Řež u Prahy, Czech Republic
⁷⁹ Oak Ridge National Laboratory, Oak Ridge, Tennessee, United States
⁸⁰ Petersburg Nuclear Physics Institute, Gatchina, Russia
⁸¹ Physics Department, Creighton University, Omaha, Nebraska, United States
⁸² Physics Department, Panjab University, Chandigarh, India
⁸³ Physics Department, University of Athens, Athens, Greece
⁸⁴ Physics Department, University of Cape Town and iThemba LABS, National Research Foundation, Somerset West, South Africa
⁸⁵ Physics Department, University of Jammu, Jammu, India
⁸⁶ Physics Department, University of Rajasthan, Jaipur, India
⁸⁷ Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg, Heidelberg, Germany
⁸⁸ Politecnico di Torino, Turin, Italy
⁸⁹ Purdue University, West Lafayette, Indiana, United States
⁹⁰ Pusan National University, Pusan, South Korea
⁹¹ Research Division and ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung, Darmstadt, Germany
⁹² Rudjer Bošković Institute, Zagreb, Croatia
⁹³ Russian Federal Nuclear Center (VNIIEF), Sarov, Russia
⁹⁴ Russian Research Centre Kurchatov Institute, Moscow, Russia
⁹⁵ Saha Institute of Nuclear Physics, Kolkata, India
⁹⁶ School of Physics and Astronomy, University of Birmingham, Birmingham, United Kingdom
⁹⁷ Sección Física, Departamento de Ciencias, Pontificia Universidad Católica del Perú, Lima, Peru
⁹⁸ Sezione INFN, Bologna, Italy
⁹⁹ Sezione INFN, Padova, Italy
¹⁰⁰ Sezione INFN, Rome, Italy
¹⁰¹ Sezione INFN, Cagliari, Italy
¹⁰² Sezione INFN, Turin, Italy
¹⁰³ Sezione INFN, Trieste, Italy
¹⁰⁴ Sezione INFN, Bari, Italy

¹⁰⁵ Sezione INFN, Catania, Italy
¹⁰⁶ Nuclear Physics Group, STFC Daresbury Laboratory, Daresbury, United Kingdom
¹⁰⁷ STFC Rutherford Appleton Laboratory, Chilton, United Kingdom
¹⁰⁸ SUBATECH, Ecole des Mines de Nantes, Université de Nantes, CNRS-IN2P3, Nantes, France
¹⁰⁹ Suranaree University of Technology, Nakhon Ratchasima, Thailand
¹¹⁰ Technical University of Split FESB, Split, Croatia
¹¹¹ Technische Universität München, Munich, Germany
¹¹² The Henryk Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences, Cracow, Poland
¹¹³ The University of Texas at Austin, Physics Department, Austin, TX, United States
¹¹⁴ Universidad Autónoma de Sinaloa, Culiacán, Mexico
¹¹⁵ Universidade de São Paulo (USP), São Paulo, Brazil
¹¹⁶ Universidade Estadual de Campinas (UNICAMP), Campinas, Brazil
¹¹⁷ Université de Lyon, Université Lyon 1, CNRS/IN2P3, IPN-Lyon, Villeurbanne, France
¹¹⁸ University of Houston, Houston, Texas, United States
¹¹⁹ University of Technology and Austrian Academy of Sciences, Vienna, Austria
¹²⁰ University of Tennessee, Knoxville, Tennessee, United States
¹²¹ University of Tokyo, Tokyo, Japan
¹²² University of Tsukuba, Tsukuba, Japan
¹²³ Eberhard Karls Universität Tübingen, Tübingen, Germany
¹²⁴ Variable Energy Cyclotron Centre, Kolkata, India
¹²⁵ V. Fock Institute for Physics, St. Petersburg State University, St. Petersburg, Russia
¹²⁶ Warsaw University of Technology, Warsaw, Poland
¹²⁷ Wayne State University, Detroit, Michigan, United States
¹²⁸ Wigner Research Centre for Physics, Hungarian Academy of Sciences, Budapest, Hungary
¹²⁹ Yale University, New Haven, Connecticut, United States
¹³⁰ Yildiz Technical University, Istanbul, Turkey
¹³¹ Yonsei University, Seoul, South Korea
¹³² Zentrum für Technologietransfer und Telekommunikation (ZTT), Fachhochschule Worms, Worms, Germany

Contents

0.1 The ALICE Collaboration	iii
1 Introduction	1
1.1 Study of Quark–Gluon Plasma	1
1.2 Proposed Physics Programme	2
1.2.1 Heavy-Flavour Production	3
1.2.2 Production of Quarkonia	4
1.2.3 Low-Mass Dileptons	5
1.2.4 Jet Measurements	6
1.2.5 Heavy Nuclear States	7
1.2.6 Comparison of Physics Reach	7
1.2.7 Comparison with Other LHC Experiments	7
1.3 ALICE Detector Upgrade	9
1.4 Running scenario	10
2 Physics Motivation	13
2.1 Heavy Flavour	13
2.1.1 Heavy-flavour Thermalization, Coalescence, and Possible Thermal Production	14
2.1.2 Heavy-flavour Energy Loss	32
2.1.3 Proton–Proton Running Requirements for Heavy-flavour Reference Data	35
2.2 Quarkonia	38
2.2.1 Motivation	38
2.2.2 Inputs for Simulations	42
2.2.3 J/ψ Yield and the Nuclear Modification Factor	43
2.2.4 J/ψ Elliptic Flow	43
2.2.5 J/ψ Polarization	44
2.2.6 Electromagnetic J/ψ Production	46
2.2.7 $\psi(2S)$ Measurement	47
2.3 Low-Mass Dileptons	48

2.3.1	Scientific Objectives	48
2.3.2	Experimental Challenges	52
2.3.3	Physics Performance	52
2.4	Jets	60
2.4.1	Jet Measurements in ALICE	63
2.4.2	Jet Structure Modifications	65
2.4.3	Photon–Jet Measurements	67
2.4.4	Heavy Flavour Jets	70
2.4.5	Reference Data	71
2.5	Heavy Nuclear States	72
3	Detectors and Readout Electronics	75
3.1	Introduction	75
3.2	Beampipe	76
3.3	ITS Upgrade	76
3.3.1	Current ITS Performance and Limitations	77
3.3.2	ITS Upgrade	79
3.3.3	Technical Implementation	83
3.4	TPC Upgrade	84
3.4.1	Status and Limitations of the Present TPC	84
3.4.2	Upgrade Concept	85
3.4.3	Expected Performance of the Upgraded TPC	85
3.4.4	Technical Implementation	89
3.4.5	R&D	90
3.4.6	Prototype Test	91
3.4.7	Front-End and Readout Electronics	91
3.5	Upgrade of TOF Readout Electronics	92
3.6	Upgrade of TRD Readout Electronics	93
3.7	Muon Spectrometer Upgrade	95
3.7.1	Muon Trigger Upgrade	95
3.7.2	Muon Tracking Upgrade	96
3.8	Other Upgrade Items	98
4	Data Collection and Processing	101
4.1	Introduction	101
4.2	Requirements	101
4.2.1	Event Rates, Event Sizes and Data Throughput	101

4.2.2	Online Data Reduction	102
4.2.3	Online Data Processing Requirements	103
4.2.4	Offline Data Processing and Storage Requirements	104
4.3	Architecture	105
4.3.1	Fast Trigger Processor	105
4.3.2	Dataflow	105
4.3.3	Farm Layout	107
4.3.4	Platform Independence	108
4.3.5	Detector Control System	108
4.3.6	Infrastructure	108
4.4	Research, Prototyping and Development	109
4.4.1	DDL and RORC	109
4.4.2	High-Level Programming of FPGAs	110
4.4.3	FLP and EPN	110
4.4.4	Farm Network	111
4.4.5	Reconstruction Software	113
4.4.6	Parallel HLT Reconstruction on GPU and CPU	117
4.4.7	Analysis of Current AliRoot Performance	118
4.4.8	Plans for the Development of Parallel AliRoot	119
5	Schedule, Cost Estimate and Organization	121
5.1	Schedule	121
5.2	Cost Estimate	125
5.3	Organization	128
References		143

Chapter 1

Introduction

Quantum ChromoDynamics (QCD) is well-established as the gauge theory of strong interactions. However, several of its fundamental aspects are not well-understood at present. There remain important open questions about the parton–hadron transition and the nature of confinement, and about the nature of QCD matter at high temperature. A much deeper insight into the mechanisms underlying chiral-symmetry breaking and the origin of light-quark mass is necessary.

This Letter of Intent (LoI) presents the plans of the ALICE (A Large Ion Collider Experiment [1]) collaboration to extend its physics programme, in order to fully exploit the scientific potential of the Large Hadron Collider (LHC) for fundamental studies of QCD, with the main emphasis on heavy-ion collisions. The proposed enhancement of the ALICE detector performance will enable detailed and quantitative characterization of the high density, high temperature phase of strongly interacting matter, together with the exploration of new phenomena in QCD. In the following we outline the physics motivation for running the LHC with heavy ions at high luminosities and summarize the performance gains expected with the upgraded ALICE detector. With the proposed timeline of initiating high-rate operation after the 2018 Long Shutdown (LS2), the objectives of our upgrade plans will be achieved by collecting data into the mid-2020's.

1.1 Study of Quark–Gluon Plasma

Strongly-interacting matter at very high temperature and density is expected to exist in a state called the Quark–Gluon Plasma (QGP), in which quark and gluon degrees of freedom are liberated, and with properties very different from the hadronic matter we ordinarily find around us. Such conditions of high temperature and density prevailed in the early Universe, a few microseconds after its formation. However, the cosmological QGP epoch is effectively shielded from astronomical observations by the subsequent evolution of the Universe, and the only means to study this fundamental state of matter is via the collision of heavy nuclei in the laboratory. In heavy-ion collisions at ultra-relativistic energies, nuclear matter is heated and brought to values of temperature and density well beyond those required for the creation of QGP, and the same kind of medium that filled the very early Universe is generated for a fleeting instant.

To reach the QGP state, matter has to undergo a QCD phase transition, and that reflects breaking of a fundamental symmetry in the theory. Below the corresponding critical temperature, matter is best described in terms of hadronic degrees of freedom, with quarks and gluons confined in colour-neutral objects. As the temperature rises the hadrons melt (*deconfinement* phase transition) and quarks and gluons are no longer bound into hadrons. The deconfinement phase transition is caused by the breaking of the Z_3 -symmetry (exact symmetry in the limit of pure-gauge QCD) at high temperature, and is illustrated in Figure 2.46 (left) by a marked change of the corresponding order parameter, the expectation value of Polyakov loop.

A second phase transition is connected with the generation of hadronic mass, as a consequence of the presence of a quark–antiquark condensate in the vacuum at low temperature. At high temperatures, the vacuum condensate decreases, and the masses of quarks drop to their bare values during the *chiral* phase transition. This can be seen in Figure 2.46 (right): the expectation value for the density of vacuum quark–antiquark condensate, *i.e.* the order parameter of the chiral transition, exhibits a sudden drop in the vicinity of the critical temperature. The underlying

reason for this phase transition is the restoration of the approximate chiral symmetry (exact symmetry in the limit of zero bare quark masses) of the QCD Lagrangian.

Lattice-QCD calculations predict that the transition between the QGP and hadronic matter at zero baryo-chemical potential is not a sharp phase transition, but rather a smooth cross-over occurring over a wide temperature range (see Figure 2.46). Recent calculations also suggest that the two phase transitions may occur at different critical temperatures [2]. Lattice QCD further predicts that even at several times the transition temperature, the energy density is still about 15 % below that expected for the Stefan–Boltzmann law for an ideal gas of non-interacting quarks and gluons, with very slow convergence to this limit. This indicates that the effective degrees of freedom in a QGP at finite temperature are not bare quarks and gluons, but rather more complex formations whose nature has not yet been understood. The crucial question of effective degrees of freedom in the QGP is also addressed experimentally, and is an area of intense activity in the field at present. However, it remains a profound challenge for experiment and theory, and may have a deep connection to other areas of physics and cosmology.

The experimental demonstration of these phase transitions, the verification of the lattice QCD predictions reflecting the fundamental symmetries of the theory, and a detailed investigation of the properties of strongly interacting matter at high temperature, are the principal aims of the ALICE scientific programme. Precise determination of the QGP properties, including critical temperature, degrees of freedom, speed of sound, and, in general, transport coefficients and equation of state, is the ultimate goal in the field. This would go a long way towards a better understanding of QCD as a genuine multi-particle theory, shedding light on the complex issues of deconfinement and chiral-symmetry restoration.

The theoretical expectation for many years was that the QGP at high temperature is a weakly-interacting gas of quarks and gluons, with the constituents traveling long distances between interactions, relative to the size of a proton. Heavy-ion experiments indicate instead fundamentally different and surprising behaviour of the created matter: the formation of a strongly-coupled plasma with very short mean free path, which exhibits a high degree of collectivity and flows, and which absorbs a significant fraction of high-energy partons propagating through it. Over time, the image of the QGP as an almost-perfect, inviscid liquid emerged from the experimental investigation at both, CERN SPS and BNL RHIC. With the first two years of LHC running, the ALICE collaboration has confirmed this basic picture, observing the creation of hot hadronic matter at unprecedented values of temperatures, densities and volumes. The first results extended the precision and kinematic reach of all significant probes of the QGP that had been measured over the past decade, and new and intriguing phenomena were observed with charm mesons and charmonia.

The observation that the QGP is a near-perfect liquid is based on measurements of inclusive spectra (radial flow) and azimuthal anisotropy of particle production (elliptic flow) for identified hadrons at soft transverse momenta (p_T), up to about 3 GeV/c. This is further confirmed by comparison to model calculations based on viscous relativistic hydrodynamics. Hydrodynamics is a general approach to fluid dynamics, describing the long-wavelength behaviour of a complex system in quasi-equilibrium, and serves as a powerful tool to determine the QGP properties. The good agreement of hydrodynamic calculation with ALICE and other experiments flow measurements provides strong evidence for the formation of a quasi-equilibrated QGP in heavy-ion collisions at LHC energies. The very low ratio of shear-viscosity to entropy-density of the QGP, inferred from the comparison of model calculations and experimental measurements, demonstrates very short mean free path inside this medium composed of strongly-coupled quasi-particle modes, whose origin is the subject of further experimental study.

1.2 Proposed Physics Programme

The LHC provides the optimal experimental conditions to study the QGP, due to the following:

- the net baryon density in the central (mid-rapidity) region is very small, corresponding to the conditions of the early Universe;
- the initial temperature and energy density are the highest achievable in the laboratory;
- the large collision energy ensures an abundance of perturbatively calculable hard QCD processes.

The study of the strongly-interacting state of matter in the second generation of LHC heavy-ion studies following LS2 will focus on rare probes, and the study of their coupling with the medium and hadronization processes. These include heavy-flavour particles, quarkonium states, real and virtual photons, jets and their correlations with other

probes. The cross sections of all these processes are significantly larger at LHC than at previous accelerators. In addition, the interaction with the medium of heavy-flavour probes is better controlled theoretically than the propagation of light partons. All these investigations should involve soft momentum scales, and thus benefit from the ALICE detector strengths: excellent tracking performance in high-multiplicity environment and particle identification over a large momentum range. In most of these studies, the azimuthal anisotropy of different probes will be measured. Major highlights of the proposed programme focus on the following physics questions:

- Study of the thermalization of partons in the QGP, with focus on the massive charm and beauty quarks. Heavy-quark azimuthal-flow anisotropy is especially sensitive to the partonic equation of state. Ultimately, heavy quarks might fully equilibrate and become part of the strongly-coupled medium.
- Study of the low-momentum quarkonium dissociation and, possibly, regeneration pattern, as a probe of deconfinement, and an evaluation of the medium temperature.
- Study of the production of thermal photons and low-mass dileptons emitted by the QGP. This should allow to assess the initial temperature and the equation of state of the medium, as well as to shed light on the chiral nature of the phase transition.
- Study of the in-medium parton energy-loss mechanism that provides both a testing ground for the multi-particle aspects of QCD and a probe of the QGP density. The relevant observables are: jet structure, jet-jet and photon-jet correlations, and jet correlations with high-momentum identified hadrons and heavy-flavour particle production in jets. In particular, it is crucial to characterize the dependencies of energy loss on the parton colour-charge, mass, and energy, as well as on the density of the medium.
- Search for heavy nuclear states such as light multi- Λ hyper-nuclei ${}_{\Lambda\Lambda}^5H$, bound states of $(\Lambda\Lambda)$ or the H dibaryon, $(\Lambda\Lambda)$ bound state, as well as bound states involving multi-strange baryons; a systematic study of light nuclei and anti-nuclei production.

Below we outline the basic physics motivation for these measurements; the details and the corresponding performance studies are described in Chapter 2.

1.2.1 Heavy-Flavour Production

High-precision measurements of charm and beauty production in heavy-ion collisions at the LHC is one of the principal physics motivations for the upgrade of the ALICE detector and the high-luminosity running of the experiment. The envisaged measurements of open heavy-flavour production will make possible to precisely determine important parameters of the strongly interacting matter that are not accessible with the present experimental setup, see Section 2.1 for details. Two topics that need high-precision and high-statistics measurements are proposed to study:

- thermalization of heavy quarks in the medium, by determination of the baryon-to-meson ratio for charm and for beauty particles, the azimuthal-flow anisotropy for charm mesons and baryons and beauty particles, and the possible in-medium thermal production of charm quarks;
- parton mass and colour-charge dependence of in-medium energy loss, by measuring the p_T -dependencies of the nuclear modification factors separately for D and B mesons, and comparing them with those for light-flavour particles.

These two topics are closely connected: the in-medium heavy-quark energy loss lowers the momenta of heavy quarks, they may thermalize in the system, and thus participate in the collective flow dynamics. The simultaneous observation of the two phenomena opens the possibility for determination of the heavy-flavour transport coefficients. Heavy-flavour production plays a special role in heavy-ion physics: it provides a calibrated probe (input p_T spectra calculable from perturbative QCD), and, in addition, this probe is well-tagged (identified), from production till observation, which enables a unique access to its interactions in the QGP, also in the low- and intermediate- p_T regime.

The ALICE collaboration already presented the first results on D-meson production in heavy-ion collisions. To address the physics questions mentioned above, these measurements have to be extended to a lower p_T , include charm baryons (possibly also charm-strange baryons) and beauty particles. The capability of studying yields and spectra of particles containing heavy quarks is given by the performance of secondary-vertex isolation close to

the primary-interaction vertex. Charm production is measured by the reconstruction of exclusive hadronic decays using topological selection of a secondary vertex. Particle identification for charged hadrons is needed to reduce the very large backgrounds in heavy-ion collisions, especially at low transverse momentum. In addition, charm and beauty can be tagged in semi-leptonic decays, detecting electrons and muons. Therefore, the excellent particle-identification capabilities of the ALICE detector have to be preserved. However, important physics topics, such as the study of heavy-flavour baryons or of open heavy-flavour hadrons with more than one heavy quark, are beyond the capability of the present detector.

Preliminary measurements of the elliptic-flow coefficient v_2 (the amplitude of second order harmonic of the azimuthal distribution) for different D mesons in 30–50 % central Pb-Pb collisions were obtained by the ALICE experiment; the result is very intriguing, because it suggests that D mesons in the p_T range 2–6 GeV/c may indeed take part in the collective flow. Models taking into account heavy-quark transport in the medium, with various implementations of the quark–medium interaction, predict a large v_2 for D mesons at low transverse momenta that should become accessible with the upgraded ALICE detector. The low- p_T B mesons v_2 , not in the reach of present experiments, is predicted to be substantially smaller than that for D mesons. Such difference in the azimuthal anisotropy at intermediate p_T is inherent to the QCD interaction mechanisms, and would thus serve as an important test of our understanding of the nature of matter formed in heavy-ion collisions.

The study of heavy-flavour energy loss has a particular interest, because gluon radiation from heavy quarks at small angles is predicted to be suppressed, in comparison to the case of light partons. Moreover, at LHC energies, high- p_T light-flavour hadrons are dominantly produced in gluon fragmentation, and gluons presumably lose more energy than quarks, due to their larger colour charge. Consequently, the clear prediction for the hierarchy of energy loss in strongly interacting matter is: gluons lose more energy than charm quarks, and the latter lose more energy than beauty quarks. Experimentally, this should be investigated by comparing the nuclear modification factors as a function of p_T of light-flavour hadrons, of charm particles, and that of beauty particles. The first measurement of the D-meson nuclear modification factor for p_T above 2 GeV/c has been already published by the ALICE collaboration. In order to access p_T down to zero, the improvement in vertexing capabilities is mandatory. This is even more true for a precision measurement of the B-meson nuclear modification factor, where, in addition, a substantial increase in event statistics is necessary.

1.2.2 Production of Quarkonia

Charmonium is the first hadron for which a clear mechanism of suppression in QGP was proposed, based on the colour-charge analogue of Debye screening. Because of difficulties to explain the observed suppression pattern, especially the $\psi(2S)$ production, alternative models were proposed. The statistical hadronization model was motivated by the observation that the $\psi(2S)$ -to-J/ψ production ratio has the value corresponding to that obtained at the chemical freeze-out temperature determined for other hadrons. In this model, the charm quarks produced in the initial hard collisions thermalize in the QGP and are distributed into hadrons at chemical freeze-out. Charmonium states are produced, together with all other hadrons, only at chemical freeze-out. The predictions of this model depend on the amount of available charm quarks, *i.e.* on the charm-production cross section, which we aim to measure with high precision.

Another model proposes kinetic recombination of c and \bar{c} quarks in the QGP as an alternative charmonium production mechanism. In this model a continuous dissociation and regeneration of charmonium takes place in the QGP over its entire lifetime. Besides the charm-production cross section, the input parameters of this model are the time dependence of the temperature, as well as other relevant cross sections and assumptions on the melting scenarios of charmonium states. Important observables, like the production yields and elliptic flow as a function of p_T and rapidity, are calculated within this kinetic transport model.

The measurement of the production of different charmonium states in Pb-Pb collisions at the LHC should provide a definitive answer on the question of J/ψ production mechanism in the QGP. Details on the expected performance are given in Section 2.2. A clear p_T -dependent pattern of J/ψ suppression, seen in the first ALICE measurements and being in good agreement with the high- p_T CMS results, strongly suggests that a (re)generation mechanism plays a significant role in low- p_T (≤ 3 GeV/c) J/ψ production at LHC energies. Indeed, both the statistical hadronization and kinetic transport model predictions do explain these first observations. Concerning $\psi(2S)$ production, preliminary results presented by the ALICE and CMS collaborations show some tensions, but are inconclusive due to the large uncertainties. However, already the first charmonium results from the LHC indicates the importance of the measurement down to zero p_T for the understanding of underlying production mechanism. In addition, large rapidity coverage of the ALICE measurements (and complementary to that of CMS) will impose

further constraints to relevant models. Statistically significant measurements of different charmonium states is also mandatory, and $\psi(2S)$ is the prime example calling for high statistics. A possibility of χ_c -production measurement in heavy-ion collisions with the upgraded ALICE detector is under investigation.

The J/ψ produced by the recombination of $c\bar{c}$ pairs in later stages of the collisions would inherit the elliptic flow of the charm quarks in the QGP. In this respect, like for the open heavy-flavour particles, the measurement of quarkonium elliptic flow is especially promising to complement the measurements of yields and nuclear modification factors. ALICE recently reported the first measurement of $J/\psi v_2$ in the p_T range $0 < p_T < 10$ GeV/ c at forward rapidities in Pb–Pb collisions, and a hint for non-zero elliptic flow was observed. With the upgraded detector, such a measurement will be possible on a qualitatively new precision level. Other measurements that will benefit from the ALICE upgrade are J/ψ polarization and J/ψ production in very low p_T (< 300 MeV/ c), where an excess, which may be attributed to J/ψ photo-production, is observed.

Given the very large energy of the collisions at LHC, an abundant production of charm and beauty quark–antiquark pairs is expected in the initial hard-scattering processes (about 80 and 3, respectively, per central Pb–Pb collision). The LHC has opened up the measurement of the Υ family in Pb–Pb collisions, where a suppression of the excited states has been observed by CMS. Still, the density of charm quarks is more than one order of magnitude larger than that of beauty, and therefore the behaviour of charmonium is expected to be completely different from that of Υ states. A detailed study of the Υ states will be performed by the ATLAS and CMS collaborations, ALICE will complement these measurements, especially in the forward rapidity region.

1.2.3 Low-Mass Dileptons

Electromagnetic radiation is produced during all stages of the heavy-ion collision, and, being detected either as a real photon or as a dilepton pair, it brings information about the entire system evolution, since the detected particles do not interact strongly with the medium. The measurement of low-mass dilepton production gives an insight into the bulk properties and the space–time evolution of the hot and dense QCD matter formed in ultra-relativistic heavy-ion collisions, and provides an access to the hadronic excitation spectrum in the medium. Comprehensive measurements of low-mass dileptons in heavy-ion collisions at the LHC, described in Section 2.3, will allow to study the following topics:

- The masses of the light-quark particles are connected with spontaneous breaking of chiral symmetry of QCD. The theory predicts that this fundamental symmetry is restored at high temperature, leading to substantial distortions of the vector and axial-vector spectral functions. Such modifications, in particular for the ρ meson, should be observable in dilepton spectra.
- The temperature reached by the system can be assessed by measuring the dilepton invariant-mass and p_T spectra. The study of low-mass dileptons also allows an estimate of real direct-photon production which is complementary to direct real-photon measurements.
- The lifetime of the system and its overall space–time evolution can be inferred from low-mass dilepton measurements. The possibility to disentangle early and late contributions makes the evolution of collectivity and the fundamental properties related to it, such as transport coefficients, viscosity, and the equation of state potentially accessible.

The dilepton invariant-mass spectrum contains information about the relevant degrees of freedom of the system and their dependence on temperature and density. At masses below 1 GeV/ c^2 the spectrum is dominated by the contribution of the light vector resonances, in particular of the ρ meson. As a consequence of the strong coupling between the ρ resonance and the hot and dense hadronic matter, close to the phase-transition boundary the ρ -meson spectral function strongly broadens. This behaviour, predicted by multi-particle QCD theory, was experimentally demonstrated in heavy-ion collisions at CERN SPS.

Recent lattice QCD calculations indicate that the critical temperature of the deconfinement phase transition may be higher than that of the chiral phase transition (see Figure 2.46), in that case the chiral symmetry will remain restored still in a hot hadronic resonance gas. As a consequence of the restoration chiral symmetry the vector and axial-vector spectral functions are modified, ultimately leading to a degeneracy of the two. While the latter is not measurable experimentally, the observation of the vector spectral-function modification is essential to reveal chiral-symmetry restoration, using QCD sum rules and constraints from lattice QCD. A precise measurement of the low-mass dilepton spectrum in heavy-ion collisions constitutes the only known means to assess experimentally the nature of the chiral phase transition.

The spectrum of real direct-photons can be inferred by the extrapolation of measurements of very low-mass dileptons to zero mass; this method is complementary to that using direct measurement and is sufficiently accurate for p_T above 1 GeV/c (when measuring dilepton masses down to about 200 MeV/c²). The virtual-photon measurement has the advantage of much less physical background than the direct-photon measurement, although the yield of virtual photons is suppressed compared to real ones. This extrapolation gives access to the direct-photon spectrum in the p_T region (1–5 GeV/c) where the thermal contribution is expected to dominate. The temperature at early stage of the system is accessible by dilepton invariant-mass spectrum at larger masses (1.5–2.5 GeV/c²).

The equation of state, *i.e.* the relation between the pressure and the temperature, is one of the most fundamental characteristics of strongly interacting matter. It determines the expansion history of the early Universe and has obvious consequences for the space–time evolution of the matter formed in heavy-ion collisions. The equation of state is a basic input to hydrodynamical models which give a good description of the observed collective flow. Dilepton pairs are emitted at all stages of the collision, causing a complex collectivity pattern when studied as a function of the invariant mass and p_T . In heavy-ion collisions at the LHC, significantly higher initial temperatures are reached than at previous accelerators, entailing even more pronounced contributions from the QGP to dilepton production at high masses. A systematic and detailed investigation of radial and elliptic flow as a function of the invariant mass and p_T will give access to the evolution of collectivity at different stages of the collision, and will provide a unique experimental handle on the equation of state of partonic matter.

We aim to measure low-mass dileptons with the upgraded ALICE detector exploiting the e^+e^- channel. The measurement of e^+e^- pairs in Pb–Pb collisions at LHC energies is experimentally very challenging. The requirement for the acceptance is to reach dilepton invariant masses and transverse momenta values as low as the critical temperature, *i.e.* $T_c \simeq 150$ MeV. This implies electron detection down to p_T in the range 0.1–0.2 GeV/c. To increase the acceptance for low- p_T electrons in this measurement, we plan to decrease the field of the main ALICE solenoidal magnet to 0.2 T (from the nominal value 0.5 T), this means that a special Pb–Pb run will be necessary with such settings. The main challenge in this measurement and analysis is the rejection of different backgrounds, such as electrons from Dalitz decays, charm decays and from photon conversions in detector material. To achieve this goal, we will rely on: improved tracking at very low momenta, enhanced vertexing capabilities suppressing the photon-conversion and charm-decay backgrounds, and low material-budget tracker reducing the conversion probability. Additional gain comes from an increased statistics due to the high-rate capability of the upgraded ALICE detector. This will be a unique measurement in the highest-energy heavy-ion collisions, giving access to the initial temperature, the partonic equation of state, and nature of chiral phase transition at vanishing baryon density.

1.2.4 Jet Measurements

The main motivation for measuring jets in heavy-ion collision is to map out the energy loss of hard scattered partons in the QGP, and thereby access the properties of the strongly interacting medium. On their way out of the medium, partons interact with the QGP, losing energy through both radiative and elastic channels with a magnitude of the energy loss strongly dependent on the density of colour charges in the medium. The observed energy loss is usually referred to as *jet quenching*. In effect, the energy loss softens the fragmentation function of the jet, resulting in an enhancement of low-momentum hadron multiplicity and a suppression of high- p_T hadrons. A precise measurement of jet-quenching effects has the potential to probe the medium at the hottest, densest stage of the collision. Specific to ALICE, following studies, discussed in Section 2.4, will benefit from the upgrade programme:

- ALICE will measure particle-identified fragmentation functions and their in-medium modifications over a wide momentum range, exploiting its particle identification systems. Such a study, unique to ALICE, will shed light on the thermalization of fast quarks and gluons, as well as on the response of the medium to large local energy deposits by partons.
- Heavy-flavour production in jets, especially at low z (fraction of original parton momentum) is another domain where ALICE will bring a unique contribution to jet studies, with its excellent performance for heavy-flavour particle detection at low p_T .
- In photon-jet correlations, the precise measurement of photons, will calibrate the energy of the partons recoiling against them, opening the possibility of detailed jet investigations down to relatively low transverse energies ($E_T \approx 20$ GeV). Moreover, a photon-tagged jet, is predominantly a quark jet, unlike for the unbiased jet sample, which is dominated by gluon jets; this provides another way to study the energy-loss hierarchy mentioned above.

- Excellent low-momentum tracking will help ALICE to investigate the fate of the energy quenched in the medium, which in the current measurements is usually hard to distinguish inside the underlying event and its fluctuations.

Precise measurements of the jet structure and of its modification in terms of energy flow, of the modification of jet-energy patterns, and of the broadening of jets due to interactions in matter promises to offer new, fundamental insights into the underlying theory. The ALICE strength is that it provides tools for differential studies of the jet structure in two approaches: track reconstruction and separation down to low momentum; and particle identification including heavy-flavour. The benefits of the high-rate upgrade will be also substantial for the photon-jet correlations studies, which are usually limited by the small cross section of electromagnetic processes and the challenging extraction of isolated photons in heavy-ion collisions.

1.2.5 Heavy Nuclear States

Another area where the ALICE measurements are unique is the search for exotic objects produced in heavy-ion collisions. The ALICE detector has superior particle identification capability combining the measurements of different detectors and using various techniques. For example, ALICE already reported the successful search for ${}^4\text{He}$ antinuclei. With the high-rate upgrade, the inspection of as many as 10^{10} Pb–Pb events should become feasible; this will allow a systematic study of the production of nuclei and antinuclei and bring within reach the detection of light multi- hyper-nuclei, such as ${}_{\Lambda\Lambda}^5\text{H}$.

Other exotic objects to search for include bound states of ($\Lambda\Lambda$) or the H dibaryon, a possible (Λn) bound state, as well as bound states involving multi-strange baryons. The important issue to note here is that, using the ALICE apparatus, in addition to search for the existence of these states, we can also study their decay properties. The presence of strangeness in these objects, making them more flavour symmetric than ordinary matter, may increase their stability. Such observations would give access to completely new information on hyper-nuclei and other heavy-nuclear bound states. The production rate estimates for different states are given in Section 2.5.

1.2.6 Comparison of Physics Reach

The improvement in the physics reach for different observables is summarized in Table 1.1. The comparison is presented between two scenarios (see also Section 1.3):

- the approved ALICE programme, which corresponds to delivered integrated luminosity of 1 nb^{-1} , however, for observables relying on minimum-bias data (such as heavy-flavour production and low- p_T dielectrons) only 1/10 of this is usable due to the TPC rate limitation;
- the proposed ALICE upgrade with assumed integrated luminosity of 10 nb^{-1} , fully used for minimum-bias data collection, and a special low-magnetic-field run with 3 nb^{-1} of integrated luminosity for low- p_T dielectron study.

Table 1.1 illustrates that, with the ALICE upgrade, in some cases (D meson from B decays R_{AA} , D meson elliptic flow) we will move from an observation to the precision measurement, and other signals (dielectrons) will become accessible.

1.2.7 Comparison with Other LHC Experiments

The ALICE upgrade and the proposed physics programme build on the demonstrated strengths of the current ALICE detector configuration:

- excellent tracking performance, especially at low momenta (down to $p_T \simeq 150 \text{ MeV}/c$);
- efficient secondary vertex reconstruction, in particular for heavy-flavour particle decays: the resolution in the distance of closest approach between a track and the interaction vertex in the transverse projection at $p_T = 1 \text{ GeV}/c$ is $\sigma_{r\phi}^{\text{1GeV}} \simeq 60 \mu\text{m}$;
- very low material budget of the tracking system, the thickness of the Inner Tracking System is 7–8 % of radiation length for normal-incident particles (*i.e.* 1.2–1.3 % per tracking layer), the overall thickness including the TPC is $\simeq 10\text{--}11\%$;

Table 1.1: Comparison of the physics reach, minimum accessible p_T and relative statistical uncertainty, for selected observables between the approved scenario (1 nb^{-1} of delivered integrated luminosity, out of which 0.1 nb^{-1} is used for minimum-bias data collection) and the proposed upgrade (10 nb^{-1} of integrated luminosity, fully exploited in minimum-bias data recording).

Observable	Approved		Upgrade	
	p_T^{Amin} (GeV/c)	statistical uncertainty	p_T^{Umin} (GeV/c)	statistical uncertainty
Heavy Flavour				
D meson R_{AA}	1	10 % at p_T^{Amin}	0	0.3 % at p_T^{Amin}
D meson from B decays R_{AA}	3	30 % at p_T^{Amin}	2	1 % at p_T^{Amin}
D meson elliptic flow ($v_2 = 0.2$)	1	50 % at p_T^{Amin}	0	2.5 % at p_T^{Amin}
D from B elliptic flow ($v_2 = 0.1$)		not accessible	2	20 % at p_T^{Umin}
Charm baryon-to-meson ratio		not accessible	2	15 % at p_T^{Umin}
D_s meson R_{AA}	4	15 % at p_T^{Amin}	1	1 % at p_T^{Amin}
Charmonia				
$J/\psi R_{\text{AA}}$ (forward rapidity)	0	1 % at 1 GeV/c	0	0.3 % at 1 GeV/c
$J/\psi R_{\text{AA}}$ (mid-rapidity)	0	5 % at 1 GeV/c	0	0.5 % at 1 GeV/c
J/ψ elliptic flow ($v_2 = 0.1$)	0	15 % at 2 GeV/c	0	5 % at 2 GeV/c
$\psi(2S)$ yield	0	30 %	0	10 %
Dielectrons				
Temperature (intermediate mass)		not accessible		10 %
Elliptic flow ($v_2 = 0.1$)		not accessible		10 %
Low-mass spectral function		not accessible	0.3	20 %
Heavy Nuclear States				
Hyper(anti)nuclei ${}^4_{\Lambda}\text{H}$ yield		35 %		3.5 %
Hyper(anti)nuclei ${}^4_{\Lambda\Lambda}\text{H}$ yield		not accessible		20 %

- particle identification using various techniques, such as specific ionization-energy loss (dE/dx), time-of-flight, transition radiation, Čerenkov radiation, to separate different particles up to a few GeV/c track by track, and up to a few tens of GeV/c statistically.

The upgraded ALICE detector will significantly improve most of these performance figures, except for charged particle identification, which will be preserved. To assess differences in the physics reach between ALICE and the other experiments participating in the LHC heavy-ion programme, we compare the tracking and particle-identification performances.

Both the ATLAS and CMS tracking systems are optimized with different requirements than those of ALICE, they have significantly larger acceptance in pseudorapidity and consequently more material in the tracking volume. As their primary physics goals rely on the tracking of high- p_T particles, the detectors use higher magnetic field, and do not concentrate on the low- p_T region. Nevertheless, after they corresponding upgrades, they will be at p_T around 1 GeV/c close in the tracking performance to the present ALICE tracker.

The ATLAS experiment with the insertable B-layer pixel detector [3] will achieve the resolution in the distance of closest approach $\sigma_{r\phi}^{\text{1GeV}} \simeq 65 \mu\text{m}$. This layer has a thickness of $\simeq 1.4\%$ of radiation length in the central part (at zero pseudorapidity), and it is placed at similar radial distance from the interaction point as the current innermost ALICE layer. The thickness of the silicon part of the ATLAS tracker at mid-rapidity is $\simeq 20\%$ and together with Transition Radiation Tracker $\simeq 40\%$ of radiation length.

After the upgrade [4] (phase 1) of the CMS pixel detector, the resolution in the distance of closest approach will be $\sigma_{r\phi}^{\text{1GeV}} = 60\text{--}65 \mu\text{m}$ at pseudorapidities $\eta = 0.0\text{--}1.0$. The four-layer pixel detector will have thickness $\simeq 7\%$

of radiation length (*i.e.* about 1.7 % per layer) in the central part ($\eta \sim 0$), and its innermost-layer radial position corresponds to that of the present ALICE. The CMS collaboration published the measurement of pions down to $p_T = 100$ MeV/c [5]. The standalone track-finding performance is evaluated down to $p_T = 300$ MeV/c with less than 10 % of fake track contamination.

Both ATLAS and CMS reported [5, 6] their first results on charged particle identification in pp collisions using dE/dx measurement in silicon detectors. From these results we estimate the upper momentum limits for statistical separation, which can be eventually achieved, to be $\simeq 0.9$ GeV/c and $\simeq 1.8$ GeV/c, for π -K and K-p separation, respectively. The track-by-track separation, needed for extraction of heavy-flavour signals at low- p_T , will be limited to significantly lower values.

The expected performance for the low-momentum tracking and vertexing (p_T below a few GeV/c) of the ATLAS and CMS detectors after upgrades is close to that of the present ALICE detector. Their particle-identification performance for charged hadrons, which is not their priority, is limited to narrower momentum region compared to that of ALICE. Therefore, the physics-performance comparison for probes relying on low-momentum tracking and vertexing will be similar to the comparison of the present and the upgraded ALICE detectors; and for those relying in addition on hadron particle identification, the ALICE experiment has better performance already at present.

In addition, many of the proposed physics topics (heavy flavours, low-mass dileptons, heavy nuclear states) rely on the inspection of high event statistics, without a dedicated trigger. With the high-rate upgrade, ALICE will be able to inspect practically all heavy-ion collisions. This is very different compared to the strategy adopted by the other experiments: to select the events by means of a trigger and to reduce significantly the event rate sent to the permanent storage. Despite the smaller acceptance, for untriggerable observables ALICE will have the advantage of larger statistics.

The physics reach with low- p_T heavy-flavour particles, such as D_s and Λ_c , and low-mass low- p_T dileptons is thus uniquely accessible by the upgraded ALICE detector. The other experiments will measure these observables at higher p_T (above 5–6 GeV/c).

Quarkonium production will be studied by all three experiments: ALICE will measure charmonia down to zero p_T at mid-rapidity and in the forward-rapidity region; ATLAS and CMS will cover central rapidities up to ~ 2.5 at higher p_T (above 3–6 GeV/c). CMS will do an effort to lower the charmonium p_T cut-off in the Pb-Pb collisions in the rapidity range 1.5–2.5, as already demonstrated in pp collisions [7], and this would complement the rapidity covered by the ALICE detector. The Υ family will be measured by ATLAS and CMS, and complemented by the ALICE measurement at forward rapidities.

In jet physics, both ATLAS and CMS have a clear advantage of significantly larger acceptance and calorimetry coverage. ALICE will contribute to heavy-ion jet measurements by studying particle-identified (including heavy-flavour) fragmentation functions.

The particle-identification capabilities together with the high event statistics make ALICE ideally suited to the search for heavy nuclear states. The other experiments may also contribute to this search at low momenta.

1.3 ALICE Detector Upgrade

To achieve the goals described above high statistics and high precision measurements are required, which will give access to the rare physics channels needed to understand the dynamics of the condensed phase of QCD. Many of these measurements will involve complex probes at low transverse momentum, where traditional methods for triggering will not be applicable. Therefore, the ALICE collaboration is planning to upgrade the current detector by enhancing its low-momentum vertexing and tracking capability, and allowing data taking at substantially higher rates.

The upgrade strategy is formulated under the assumption that, after the second long shutdown in 2018, the LHC will progressively increase its luminosity with Pb beams eventually reaching an interaction rate of about 50 kHz, *i.e.* instantaneous luminosities of $L = 6 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$. In the proposed plan the ALICE detector is modified such that all interactions will be inspected. ALICE will then be in a position to accumulate 10 nb^{-1} of Pb-Pb collisions inspecting about 10^{11} interactions. This is the minimum needed to address the proposed physics programme with focus on rare probes both at low and high transverse momenta as well as on the multi-dimensional analysis of such probes with respect to centrality, event plane, multi-particle correlations, and so on.

High statistics Pb–Pb measurements will have to be accompanied by precision measurements with pp and p–Pb collisions to provide a quantitative base for comparison with results from Pb–Pb collisions. Those are furthermore crucial to understand initial state modifications in nuclei, such as shadowing or non-linear QCD evolution, possibly leading to gluon saturation, which are important for the correct interpretation of final state effects due to the QGP and the influence of initial conditions on observables like elliptic flow. In order to collect a sample of pp reference data with a statistical significance comparable to that of the Pb–Pb data, the required integrated luminosity for pp collisions is estimated to about 6 pb^{-1} , taking into account smaller combinatorial background in most of the cases. This results in a pp data taking at an event readout rate of 200 kHz over a period of a few months, which would allow to fulfill this programme with one special run at the pp centre-of-mass energy equal to that of Pb–Pb per nucleon pair. Another option under discussion is the utilization of lighter ion collisions, which the LHC could produce at higher luminosities. For Ar–Ar collisions an increase of one order of magnitude could potentially be reached.

On the basis of the above considerations, the ALICE collaboration proposes a strategy for upgrading the ALICE detector to be able to make full use of a high-luminosity LHC ($L = 6 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$) for Pb–Pb. The planned upgrades will preserve the excellent particle identification capability while enhancing the vertexing and tracking capabilities, and include:

- A new, high-resolution, low-material-thickness Inner Tracking System (ITS). With this new detector the resolution of the distance of closest approach between a track and the primary vertex will be improved by a factor of about 3, and the standalone ITS tracking performance will be significantly enhanced. An overview of the new ITS is given in Section 3.3, while for a detailed discussion we refer to the Conceptual Design Report [8].
- Upgrade of the Time-Projection Chamber (TPC) with replacement of the readout multi-wire chambers with GEM (Gas Electron Multiplier) detectors and new pipelined readout electronics. This is needed in order to operate the TPC in ungated mode and make its readout dead-time free. Section 3.4 gives an overview of the TPC upgrade concept, the technical implementation and the ongoing R&D activities.
- Upgrade of the readout electronics of: Transition-Radiation Detector (TRD), Time-Of-Flight (TOF) detector, PHOton Spectrometer (PHOS) and Muon spectrometer. This upgrade will allow for high rate data taking with these detectors and is discussed in Sections 3.5 to 3.8.
- Upgrade of the forward trigger detectors and trigger system for high rate operation, which is discussed in Section 3.8.
- Upgrade of the Online Systems: High-Level Trigger (HLT), data acquisition (DAQ) and trigger system, to adapt for high rates. The rate for heavy-ion events handled by the online systems up to permanent data storage should be increased up to 50 kHz corresponding to roughly two orders of magnitude with respect to the present system. A description of the upgrade of the data collection and processing systems is given in Chapter 4.
- Upgrade of the offline data processing software to cope with the reconstruction and analysis of a much larger number of events. The strategy for increasing the offline event processing rate up to two order of magnitude is outlined in (see Chapter 4).

1.4 Running scenario

The approved running scenario, which assumes 1 nb^{-1} , corresponds to event statistics about 10^{10} interactions. Without the high-rate upgrade, ALICE would be able to inspect all interactions for rare probes which can be selected with a trigger, such as high- p_T jets or dimuons. On the other hand, ALICE would be limited to a statistics of 10^9 interactions for most of the relevant probes, those which cannot be selected by a trigger and require an inspection of all collisions. For these probes, the high-rate upgrade and an integrated luminosity of 10 nb^{-1} will allow ALICE to collect a statistics hundred times larger compared to the approved running scenario, while for rare probes the increase will be a factor of ten.

During the 2011 LHC Pb–Pb operation the luminosity has reached $\mathcal{L} = 5 \times 10^{26} \text{ cm}^{-2}\text{s}^{-1}$, twice the design luminosity at the energy of $\sqrt{s_{\text{NN}}} = 2.76 \text{ TeV}$. These results give confidence that about twice design peak luminosity, i.e. levels of $\mathcal{L} = 2 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$, can be reached at full energy after LS1 for a bunch spacing of 200 ns [9]. Decrease of β^* from 0.5 m to 0.4 m and 50 ns bunch spacing with the same bunch parameters would already result in

peak luminosities in excess of $\mathcal{L} = 7.2 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$. Possible additional handles for increasing luminosity are the optimization of emittance and increase of bunch charge. The dispersion suppressor installation in IR2 during LS2 should allow the ALICE luminosity to be pushed to these levels without quenching magnets [9]. Assuming a stable beam duration of 10^6 s during the yearly heavy ion period of one month, together with a ratio of average to peak luminosity of 40%, the above example would lead to an integrated luminosity of 2.85 nb^{-1} per year.

The Pb-Pb program presented in this LOI would therefore require 3.5 months of ALICE operation to arrive at the 10 nb^{-1} of integrated luminosity. One month of a special run at low fields for the di-lepton measurement will be needed. The required 6 pb^{-1} of pp reference data at the equivalent Pb nucleon energy can be integrated in a one month period at a leveled luminosity of $6 \times 10^{30} \text{ cm}^{-1}\text{s}^{-1}$. The p-Pb program can be achieved by running for half a month at a leveled luminosity of $10^{29} \text{ cm}^{-1}\text{s}^{-1}$ which integrates 50 nb^{-1} . A run with lower mass nuclei (e.g. Ar) could be considered in addition, if a physics case for it would emerge.

Based on the above numbers, a possible running scenario for the operation of the upgraded ALICE detector could be the following:

- 2019: Pb-Pb 2.85 nb^{-1}
- 2020: Pb-Pb 2.85 nb^{-1} at low magnetic field
- 2021: pp reference run
- 2022: LS3
- 2023: LS3
- 2024: Pb-Pb 2.85 nb^{-1}
- 2025: 50% Pb-Pb 1.42 nb^{-1} + 50% p-Pb 50 nb^{-1}
- 2026: Pb-Pb 2.85 nb^{-1}

In this document we have assumed a peak luminosity of $6 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$ and average luminosity of $2.4 \times 10^{27} \text{ cm}^{-2}\text{s}^{-1}$ for benchmarking the performance of the upgraded ALICE detector.

Chapter 2

Physics Motivation

In this Chapter we outline the physics motivations for the upgrade of the ALICE experiment. The main part, which define the requirements for the detector and running scenario, constitutes the first three sections, where the studies of physics performance of:

- heavy-flavour production,
- production of quarkonia,
- low-mass dileptons

are presented. In the remaining two sections additional physics topics, for which the upgraded ALICE detector can bring significant contributions, are described. These are: jet measurements and search for heavy nuclear states.

2.1 Heavy Flavour

Heavy quarks play a special role in heavy-ion physics because they constitute a tagged (identified) probe (from production to observation), which enables a unique access to their interactions in the QGP. This allows us to gain microscopic insights into the transport properties of the medium. Heavy-flavour particles may be thought of as “Brownian motion” markers, the kinematical distributions of which (especially in momentum and azimuthal angle) reflect their reinteraction history. The large mass makes complete thermalization very difficult, thus most likely preserving a “memory” of the interaction history. As an additional benefit, the large-mass limit allows for simplifications in the theoretical treatment via controlled approximations (such as potential approaches), which improves the accuracy in extracting microscopic information on QCD matter from experiment (for a review, see e.g. [10]). Moreover, in QCD as in QED, the interactions of non-relativistic and semi-relativistic particles in a medium are dominated by elastic collisions, while in the ultra-relativistic limit radiation takes over. At which momentum scale this transition occurs in QCD is currently an open, yet fundamental question, whose answer is required for a proper interpretation of hard- and heavy-flavour probes. In fact, heavy quarks provide a unique tool to map out the momentum scale of this transition, since they can be tagged throughout all momentum ranges, from low to intermediate, to high p_T .

The two main open questions concerning heavy-flavour interactions with the QGP medium —and the corresponding experimental handles— are:

- Thermalization and hadronization of heavy quarks in the medium, which can be studied by measuring the baryon/meson ratio for charm (Λ_c/D) and for beauty (Λ_b/B), the azimuthal anisotropy v_2 for charm mesons and baryons, and the possible in-medium thermal production of charm quarks.
- Heavy-quark in-medium energy loss and its mass dependence, which can be addressed by measuring the nuclear modification factors R_{AA} of the p_T distributions of D and B mesons separately in a wide momentum range.

These two topics are presented separately in the following. However, it is important to emphasize that they are closely related. The high-momentum heavy quarks quenched by in-medium energy loss are shifted towards low momentum and may ultimately thermalize in the system, through QCD interaction mechanisms that are essentially the same as those responsible for the energy loss, and participate in the collective expansion dynamics. Therefore, the simultaneous experimental investigation and theoretical understanding of the thermalization-related observables and of the energy-loss-related observables constitute an unique opportunity for the characterization of the QGP properties, in particular of the heavy-flavour transport coefficient. In this scope, heavy quarks are particularly well suited, because they preserve their identity (mass and flavour) while interacting with the medium.

Presently, at mid-rapidity, the capability of studying yields and spectra of particles containing heavy quarks is mainly provided by the ITS. Charm production can be studied through meson and baryon hadronic decays ($D^0 \rightarrow K^-\pi^+$, $D^+ \rightarrow K^-\pi^+\pi^+$, $D_s^+ \rightarrow K^+K^-\pi^+$, $D^{*+} \rightarrow D^0\pi^+$, $D^0 \rightarrow K^-\pi^+\pi^-\pi^+$, and $\Lambda_c^+ \rightarrow pK^-\pi^+$) via topological selection of secondary decay vertices displaced from the primary vertex. Particle identification is required to reduce the very large backgrounds in heavy-ion collisions. It is provided mainly by the TPC and TOF detectors. In addition, charm and beauty can be tagged in semi-leptonic decays $D, B \rightarrow e+X$. Electron identification is provided by the TPC, TOF, TRD and EMCAL detectors. However, in order to obtain the B -decay electron yields, a residual component of electrons from charm decays has to be subtracted statistically, implying a significant systematic uncertainty, especially in the low p_T region (below 5 GeV/c). Important physics topics such as the study of beauty baryons or of hadrons with more than one heavy quark are beyond the capability of the current detector, and the performance for charm baryons will be much worse than for charm mesons, given that the most abundantly produced baryons (Λ_c) have a mean proper decay length of only 60 μm , to be compared with the 120–300 μm of D mesons.

The physics motivation related to heavy-quark thermalization will be discussed in Section 2.1.1, along with the summary of the simulation studies for the measurements of low- p_T D meson production, low- p_T beauty production, and Λ_c production (Section 2.1.1.1). Projected sensitivity will be presented for the study of the charm baryon/meson ratio and of the charm and beauty elliptic flow at low transverse momentum (Section 2.1.1.2).

The physics motivation related to heavy-quark energy loss will be discussed in Section 2.1.2, together with an example of the expected performance for the measurement of the relative suppression of B and D mesons at low transverse momentum (Section 2.1.2.1).

The requirements for the proton–proton running (centre-of-mass energy, statistics, rates), needed to collect reference data for the heavy-ion studies, will be treated in Section 2.1.3.

2.1.1 Heavy-flavour Thermalization, Coalescence, and Possible Thermal Production

Given the very large energy of the collisions at LHC, an abundant production of $c\bar{c}$ and $b\bar{b}$ pairs is expected in the initial hard-scattering processes (about 80 and 3, respectively, per central Pb-Pb collision at $\sqrt{s_{\text{NN}}} = 5.5$ TeV, are estimated from NLO QCD calculations [11] with nuclear PDF modification [12] and binary collision scaling).

Heavy quarks lose energy while traversing the hot and dense medium and, at sufficiently low transverse momentum, they can thermalize in the medium itself. Reinteractions will reflect in the p_T spectra and, in particular, in the elliptic flow in semi-central collisions.

Elliptic flow, measured in non-central collisions, provides the most direct evidence of the collective hydrodynamical behaviour of the medium. During the collision, the two nuclei overlap in an elliptically-shaped region, the short axis of which lies in the reaction plane. The expansion, driven by pressure gradients, translates the space asymmetry into a momentum anisotropy. Anisotropy is detected by measuring the momentum-dependent azimuthal distribution $d^2N/dp_T d\Delta\phi$ of the produced particles with respect to the reaction plane. The second coefficient, v_2 , of the Fourier expansion of the azimuthal distribution quantifies the anisotropy.

The distribution of heavy quarks at the moment of their creation is isotropic in azimuthal direction. Therefore, a non-zero v_2 value of heavy-flavour mesons can only originate from interactions between QGP constituents and the heavy quarks.

A preliminary measurement of the D^0 , D^+ and D^{*+} v_2 coefficient in semi-central Pb-Pb collisions was obtained by ALICE using about 10^7 events from the 2011 run (see Figure 2.1). The result is very intriguing: the non-zero v_2 (3σ significance for D^0 in 2–6 GeV/c) suggests that D mesons may take part in the flow dynamics.

Several models [13–15] that include heavy-quark transport in the medium, with various implementations of the




Figure 2.1: Preliminary measurement of D^0 , D^+ and D^{*+} v_2 in 30–50% Pb–Pb collisions at the LHC by ALICE, with 10^7 events from the 2011 run.




Figure 2.2: Model predictions for the azimuthal anisotropy parameter v_2 of D and B mesons (or J/ψ from B decays) in Pb–Pb collisions at LHC energies: top, from left to right [13, 14], bottom [15].

quark–medium interaction, predict a large v_2 (up to 0.2 in semi-central collisions) for D mesons at low momentum, as shown in Figure 2.2. At low p_T , v_2 for B mesons is predicted to be substantially smaller than for D mesons. This is a consequence of the smaller mass of the charm quarks, which is more easily influenced by the v_2 of the light plasma particles with which they collide during the expansion. This difference of v_2 between beauty and




Figure 2.3: Constituent quark scaling of meson and baryon v_2 as measured by STAR in Au–Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV [16]. Left: v_2 vs. p_T . Right: v_2/n_q vs. p_T/n_q .

charm quarks at intermediate p_T is inherent of the QCD interaction mechanisms and can, thus, serve as test of our understanding of this mechanism in an extended QGP.

Measurements of the v_2 of D mesons and B mesons (or J/ψ / D from B decays) starting at p_T of about 2 GeV/c with absolute uncertainties well below 0.02 for D and of at most 0.05 for B are required in order to achieve sensitivity to the predicted difference between charm and beauty. On the basis of the results obtained for D mesons with 10^7 events from the 2011 Pb–Pb run, with absolute statistical uncertainties of about 0.10, we expect to reduce the statistical uncertainties to about 0.05 with the data of the 2015–16 runs (3 times larger integrated luminosity). However, a measurement for beauty down to 2 GeV/c will not be possible, for ALICE due to the lack of statistics, and for CMS due to the p_T cut of about 6 GeV/c for non-prompt J/ψ imposed by the muon trigger p_T threshold.

At RHIC, v_2 of identified particles scales with $m_T - m_0$ for $p_T < 1\text{--}2$ GeV/c [17]. The same behaviour is observed also at LHC [18]. At higher p_T , v_2 depends strongly on the number of constituent quarks n_q , so that v_2/n_q scales with $p_T/n_q \simeq m_T/n_q$. This is visible in Figure 2.3, where the Λ and K_S^0 data extend up to about 5 GeV/c [16]. This observation points to a partonic degree of freedom in the initial state. It can be naturally explained with a coalescence model [19] for hadronization, where the flow of constituent quarks add up, so that: $v_2^{\text{meson}}(p_T) = 2v_2^q(p_T/2)$ and $v_2^{\text{baryon}}(p_T) = 3v_2^q(p_T/3)$. Preliminary results from ALICE on proton and pion v_2 up to $p_T \approx 9$ GeV/c indicate that this scaling holds at LHC energies only within about 20% [20], thus challenging the interpretation based on coalescence models.

A very stringent test of the level of in-medium thermalization of charm and beauty quarks would be to verify whether this universal scaling continues to hold also for heavy-flavour mesons and baryons. This requires measuring the D meson and Λ_c elliptic flow in the range between 2 and 5 GeV/c, where the baryon/meson (Λ/K) separation measured at RHIC is most pronounced.

The production of charm and beauty baryons has also a particular interest to assess the thermalization and the mechanisms of hadronization of heavy flavours in the medium, because it was predicted that their production could be significantly enhanced in nuclear collisions [25]. For light flavour and strange baryons, such an enhancement was indeed observed at intermediate transverse momenta at RHIC and at LHC [21, 22], as shown in Figure 2.4.

Within thermal models, it is assumed that charm and beauty hadrons are produced during the QGP hadronization and they are in thermal equilibrium with the medium. Within coalescence models, partons produced in hard scatterings can combine with quarks and anti-quarks in the QGP to form hadrons. The resulting baryons from coalescence will have their strongest contribution in a momentum range intermediate between those from independent fragmentation and from thermal production. In addition, contrary to thermal models, coalescence models consider also the possibility of recombination of a heavy quark with di-quarks present in the QGP. It was suggested that this could lead to a rather significant enhancement of the Λ_c with respect to thermal models [25], where the relative




Figure 2.4: Λ/K_s^0 ratio vs. transverse momentum as measured by ALICE at LHC and by STAR at RHIC [21, 22].

abundance of particles depends only on their masses.

Figure 2.5 shows the $\Lambda_c/(D^0 + D^+)$ ratio in central (0–10%) Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV from a model with charm-quark transport in the QGP and hadronization via coalescence (based on [23]). The ratio reaches a maximum of 0.16 at p_T of about 4 GeV/c, and the enhancement with respect to pp collisions is of a factor about 1.5.

Some coalescence models consider also the possibility of recombination of a heavy quark with di-quarks present in the QGP, which would lead to larger enhancements. As an example of the size of the expected effect, we report in Figure 2.6 the Λ_c/D^0 and Λ_b/B^0 enhancements as a function of transverse momentum for central Au–Au collisions at RHIC ($\sqrt{s_{NN}} = 200$ GeV) [24]. The enhancement is of up to 7–10 for both baryon/meson ratios, and it is maximum at p_T of about 2 GeV/c for charm and about 4 GeV/c for beauty. For strange quarks (Figure 2.4) the position of the maximum is higher by about 0.5 GeV/c at LHC with respect to RHIC energy. Therefore, one can expect the maximum of Λ_c/D^0 to be at p_T of about 3 GeV/c at LHC energy. This demands for a measurement of Λ_c production in central Pb–Pb collisions starting from p_T of 2–3 GeV/c.

Complementary to the Λ_c measurement is the study of the production of the strange charm meson D_s^+ . Due to the abundance of strange quarks in the QGP, D_s^+ production is expected to be enhanced at low momentum, if charm hadrons form predominantly by in-medium hadronization of charm quarks, independently of the specific hadronization mechanism (see e.g. [23, 26, 27]). In particular, the enhancement of D_s^+/D in Pb–Pb vs. pp collisions quantifies how much of the strangeness content of the QGP is picked up by the charm; such an enhancement therefore quantifies the coalescence contribution in charm-quark hadronization. The maximal enhancement effect may reach up to a factor of 2. Therefore, an accuracy of at least 20–30% in the D_s^+ measurement is required for an evidence of the effect and of 10% for its quantitative study.

Also the comparison of D^0 (or D^+) to D_s^+ v_2 [23, 27] is very interesting. D mesons formed during the medium expansion are expected to couple (through their light-quark content) to the bulk hadronic medium [28] and thus gain additional v_2 from the hadronic phase. Since this effect would be smaller for strange than for up and down quarks, a v_2 -splitting of D^0 vs. D_s^+ is expected. The effect would be of 20–30% relative difference in v_2 [23, 27, 28], thus requiring, for observation, an experimental accuracy on each v_2 of better than 10%.

A first preliminary measurement of D_s^+ production in $4 < p_T < 12$ GeV/c in central Pb–Pb collisions was carried out with 16×10^6 events from the 2011 Pb run [29]. Figure 2.7 (left) shows the nuclear modification factor $R_{AA}(p_T) = (dN_{AA}/dp_T)/(\langle T_{AA} \rangle d\sigma_{pp}/dp_T)$ for D_s^+ compared to the average R_{AA} of non-strange D mesons (this measurement will be discussed in more detail in the next section, in the context of heavy-flavour energy loss). The D_s^+ nuclear modification factor shows similar suppression as for D^0 , D^+ and D^{*+} for $p_T > 8$ GeV/c. For lower transverse momenta, the central values of the measurement show an increase, however the large statistical ($\approx 35\%$) and systematic ($\approx 60\%$) uncertainties do not allow to draw a conclusion on a possible D_s^+/D enhancement. Figure 2.7 (right) shows the results of a calculation [27] that includes heavy-quark transport in the QGP and hadronization via coalescence, obtaining an enhancement of D_s^+/D by about a factor of 2 in the interval $0 < p_T <$




Figure 2.5: $\Lambda_c/(D^0 + D^+)$ ratio in central Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV from a calculation with charm-quark transport in the QGP and coalescence (based on [23]).




Figure 2.6: Heavy flavour baryon/meson enhancement as a function of p_T in Au–Au collisions at $\sqrt{s_{NN}} = 200$ GeV [24]. Left: Λ_c/D^0 . Right: Λ_b/\bar{B}^0 .

4 GeV/ c .

If the initial temperature of the high-density system produced in central collisions is large enough, a thermal production of $c\bar{c}$ pairs may occur, leading to a measurable increase of the total production yields of charm particles [30–33].

For example, thermal charm production is studied at next-to-leading-order in QCD in Ref. [30]. It is found that the total yield of charm production may be enhanced by 30 (80)%, if initial temperatures of 700 (750) MeV are considered and a charm quark mass value of 1.3 GeV/ c^2 is used (see Figure 2.8, left). However, the enhancement would be marginal if the initial temperature is of the order of 600 MeV or if the charm-quark mass is 1.5 GeV/ c^2 or higher.

In-medium $c\bar{c}$ production is studied within a partonic transport model in Ref. [31]. As shown in Figure 2.8 (right), the relevance of the effect is found to depend strongly on the initial gluon density assumed for the deconfined system (initial conditions from PYTHIA, Color Glass Condensate or mini-jets) and on the charm quark mass value.

In conclusion, while the yield of possible thermal charm production depends on specific model parameters, it is




Figure 2.7: Left: R_{AA} of D_s^+ and non-strange D mesons in central Pb-Pb collisions (0–7.5%) at $\sqrt{s_{NN}} = 2.76$ TeV [29]. Right: the same observables from a calculation with charm-quark transport in the QGP and coalescence [27].




Figure 2.8: Thermal charm production in Pb-Pb collisions at top LHC energy, as a function of the proper time, for two different models. Left: $c\bar{c}$ yields per unit of rapidity at mid-rapidity from [30]; total $c\bar{c}$ yields from [31].

clear that its observation would provide key information on the initial temperature and density of the deconfined plasma. From the experimental point of view, this effect has to be searched as an enhancement of the total D meson production yields per binary collision in Pb-Pb with respect to pp collisions, or in central with respect to peripheral Pb-Pb collisions. It is, therefore, essential to reconstruct D meson decays down to zero p_T . In the comparison of pp and Pb-Pb, the initial-state modification has to be taken into account. This latter can be understood by studying p-Pb collisions.

The measurement of the total charm production will also provide the natural normalization for the study of medium effects on charmonium production (J/ψ and ψ'). The crucial advantage of this normalization is that the initial-state effects are mostly common between open charm and charmonium and, thus, cancel out in the ratio $R_{AA}^{J/\psi}/R_{AA}^D$. In addition, the total charm production yield in Pb-Pb collisions is an ingredient in the model calculations of charmonium regeneration in the QGP [34] (see Section 2.2). At present, the large uncertainty on the estimate of this yield from perturbative QCD calculations translates into the main limitation in the comparison of statistical models predictions with J/ψ suppression data [35] (see Figure 2.31). Only a precise measurement of D meson

production down to zero transverse momentum can provide a significant improvement in this comparison.

2.1.1.1 *Simulation Studies: Low- p_T Charm and Beauty, Charm Baryons and Charm-Strange Mesons*

A detailed presentation of the performance studies for heavy-flavour detection with the ALICE upgrades is given in Chapter 2 of the ITS Upgrade CDR [36]. In this section we summarize the main results. We consider the following benchmark analyses:

- Charm meson production via $D^0 \rightarrow K^-\pi^+$;
- Beauty production via $B \rightarrow D^0 (\rightarrow K^-\pi^+)$;
- Charm baryon production via $\Lambda_c \rightarrow p K^-\pi^+$.

Other analyses have been started and are in progress:

- D_s production via $D_s^+ \rightarrow K^+ K^-\pi^+$, also described in this section;
- Beauty production via $B \rightarrow D^0\pi^+$, see [36];
- Beauty production via $B \rightarrow J/\psi (\rightarrow e^+e^-)$, see [37];
- Beauty production via $B \rightarrow e^+$, see [37].

For all the studies, a fast simulation scheme has been employed. It is based on Monte Carlo productions including the detailed geometry and response of the current ALICE detector setup. The impact of the ITS upgrade is obtained by recomputing reconstructed track parameters by means of a scaling of the residuals of the impact parameters in $r\varphi$ and z , as well as of the transverse momentum, with respect to their true values. These are known from the generated particle kinematics. The scaling factors are the ratios of the upgrade/current resolutions on these variables.

For the performance studies, the following baseline configuration for the upgraded ITS (see detailed description in Chapter 3) has been considered: 7 pixel layers with radii from 2.2 to 43 cm instrumented with pixel detectors with an intrinsic resolution $(\sigma_{r\varphi}, \sigma_z) = (4, 4) \mu\text{m}$, and with a radiation length of $0.3\% X_0$ per layer. Figure 3.5 in Chapter 3 shows the corresponding impact parameter resolutions for charged pions as a function of transverse momentum for the upgraded ITS compared to those of the current ITS. The improvement in the pointing precision is of about a factor 3 in the transverse plane ($r\varphi$) and a factor 6 in the longitudinal direction (z).

Charm Mesons: $D^0 \rightarrow K^-\pi^+$

The $D^0 \rightarrow K^-\pi^+$ reconstruction can be considered as a benchmark for all the D meson studies. Here we present a comparison of the performance achievable in Pb–Pb collisions with the current and upgraded ITS configurations. The comparison is made for the centrality class 0–20%, which is the same as was used for the first R_{AA} measurement [38].

The resolutions on the reconstructed position of the $D^0 \rightarrow K^-\pi^+$ decay vertex are shown in Figure 2.9. The simulation study was performed using the same cut values for the two ITS configurations, in order to single out the effect of the improved tracking resolutions. The cuts were fixed to values close to those used for the 2010 Pb–Pb data analysis. Particle identification information from the TPC and TOF detectors was used as in the current data analysis. In particular, the kaon identification up to a momentum of about 2 GeV/c provides a reduction by a factor of about 3 of the combinatorial background at low $D^0 p_T$. The selected signal (raw) yield was obtained by multiplying the corrected $D^0 p_T$ spectrum dN/dp_T measured with the 2010 data [38], by the efficiency calculated from the simulation for the current or upgraded ITS configuration. The background yield in the D^0 mass region was scaled in the simulation so as to match, for the current ITS case, the value measured in the data. The same background scaling factor was then used also for the upgraded ITS case.

The left panel of Figure 2.10 shows the comparison for the signal-to-background ratio. In general the efficiency of the signal selection in the current and upgraded ITS is comparable. On the other hand, the background rejection improves by a factor of 6 for $p_T > 2 \text{ GeV}/c$ and by a factor 25 for $p_T < 2 \text{ GeV}/c$, so that a strong increase of the signal-to-background ratio is obtained.




Figure 2.9: $D^0 \rightarrow K^- \pi^+$ secondary vertex position resolutions for current and upgrade scenarios: x (left) and z (right) coordinates.




Figure 2.10: $D^0 \rightarrow K^- \pi^+$: comparison of the signal-to-background ratio (left) and significance (right) obtained for the current and upgraded ITS.

The right panel of Figure 2.10 shows the significance normalized to the number of events for both the present and upgraded ITS. For $p_T < 2 \text{ GeV}/c$ the significance is affected by the present uncertainties in the expected nuclear modification factor in Pb–Pb, and in the background as extrapolated from $p_T > 2 \text{ GeV}/c$, taking into account also the uncertainty in the pion nuclear modification factor (error box).

Considering an integrated luminosity of 10 nb^{-1} , the number of central events in the class 0–20% would be 1.7×10^{10} . Under these assumptions, the significance is several hundreds at any p_T . For comparison, the significance in the 2010 run with the present setup, corresponding to 3×10^6 central events, is 8–10 for $p_T > 2 \text{ GeV}/c$ (see Figure 2.10) and smaller than 3 for $p_T < 2 \text{ GeV}/c$. At very low p_T , even considering the uncertainties in the nuclear modification factor and in the background that enter in the estimation of the significance, the measurement will be quite precise and p_T binning finer than that considered in this study (like $0.5 \text{ GeV}/c$) will be possible.

We estimated the systematic uncertainties that we expect on the measurement of prompt D^0 production in the $D^0 \rightarrow K^- \pi^+$ channel after the upgrade of the ITS on the basis of those evaluated for the measurement performed with 3×10^6 Pb–Pb events in the centrality class 0–20% [38]. A reduction of the systematic uncertainties can derive from the following features:




Figure 2.11: Left: Nuclear modification factor of prompt D^0 mesons in central Pb–Pb for $L_{\text{int}} = 10 \text{ nb}^{-1}$ with statistical and systematic uncertainties. Right: Contributions to the relative systematic uncertainty.

- Improved spatial and momentum resolutions. These will allow for a reduction of the background yield, for the capability of performing a more p_T -differential measurement, and for an increase of the selection efficiency avoiding the selection of D mesons in the tails of the cut variables, where a precise description of the real shape in MC simulations is generally harder.
- Higher statistics. In [38], the evaluation of systematic uncertainties was affected by the interplay of pure systematic effects and statistical fluctuations due to the limited statistics available; the higher statistics will allow for a detailed study of systematic effects, unaffected by statistical fluctuations. In this respect, the D^0 could be used as a benchmark also for other heavy-flavour analyses (e.g. D_s , Λ_c reconstruction).

A strong improvement in the accuracy of the measurement will come from the direct measurement of the fraction of prompt and secondary (from B decay) D mesons. In [38] the fraction of secondary D^0 from B decays, f_{sec} was estimated relying on the FONLL [39] prediction of beauty production and by assuming the nuclear modification factor of B mesons R_{AA}^B to be in the range $1/3 < R_{AA}^{\text{from } B} / R_{AA}^{\text{prompt } D} < 3$. Both the improved resolution and the higher statistics will allow to estimate f_{sec} using data-driven methods. In particular, for $p_T > 2 \text{ GeV}/c$, the method based on the fit of the impact parameter distribution of the candidate, described in [36], will be used. As described in the next paragraph, the statistical and systematic uncertainties on the fraction of prompt D^0 are expected to be $< 1\%$ and $< 5\%$, respectively.

Figure 2.11 (left) shows the nuclear modification factor of prompt D^0 with the systematic and the statistical uncertainties expected for 8.5×10^9 central (0–10%) Pb–Pb events, corresponding to $L_{\text{int}} = 10 \text{ nb}^{-1}$. The right panel shows the p_T dependence of the various contributions to the relative systematic uncertainty. Here, it is assumed that the systematic uncertainties in pp have the same values as in Pb–Pb and that the two are uncorrelated (which is a conservative assumption), thus the relative systematic uncertainties on R_{AA} are $\sqrt{2}$ times larger than the Pb–Pb ones.

B Meson Production via Displaced D^0

Most of the B meson decay channels include a D^0 (\bar{D}^0) particle (the branching ratio $B \rightarrow D^0$ is about 60% [40]). The kaon and pion tracks coming from secondary D^0 decays are, on average, more displaced from the primary vertex than those coming from the decay of a prompt D^0 . This is due to the relatively long lifetime of B mesons ($c\tau \approx 460\text{--}490 \mu\text{m}$). Therefore, the selection applied on reconstructed D^0 candidates, optimized to prefer secondary vertices displaced from the primary vertex, further enhances the secondary-to-prompt ratio of reconstructed D^0 up to typical values around 15% even for $p_T < 5 \text{ GeV}/c$ [41].




Figure 2.12: Left: Impact parameter distributions for prompt and secondary (from B decays) D^0 obtained with the current and upgraded ITS configurations in the transverse momentum range $2 < p_T < 3 \text{ GeV}/c$. Right: Relative statistical uncertainty on the fraction of D^0 mesons from B decays in central Pb-Pb collisions (three cases of integrated luminosity, with 1.7×10^{10} events corresponding to 10 nb^{-1}), with the upgraded ITS, for two values of the input fraction of non-prompt D^0 mesons (0.1 and 0.2).




Figure 2.13: Left: Relative statistical uncertainty on $f_{\text{feed-down}}$ expected with the current ITS in the *no high-rate* scenario (0.1 nb^{-1}), and with the upgraded ITS in both the *no high-rate* and *high-rate* (10 nb^{-1}) scenarios (for details see text). A realistic p_T dependence was considered for $f_{\text{feed-down}}$. Right: p_T dependence of the different sources of systematic uncertainty on $f_{\text{feed-down}}$.

The fraction of *prompt* and *displaced* D^0 mesons can be measured by exploiting the different shapes of the impact parameter distributions of primary and secondary mesons. This approach has been already used in pp collisions by the CDF Collaboration [42] to measure the production of prompt D mesons at $\sqrt{s} = 1.96 \text{ TeV}$. It has also been used by the LHCb Collaboration [43] to measure the production of B mesons at $\sqrt{s} = 7 \text{ TeV}$ at forward rapidity.

The main factors determining the performance achievable with this method are the available statistics, the uncertainty deriving from the subtraction of the background impact parameter distribution, and the resolution of the D meson impact parameter, which is directly related to the resolution of the track position in the vicinity of the primary vertex. The left panel of Figure 2.12 shows the impact parameter distribution for prompt and secondary D^0 in $2 < p_T < 3 \text{ GeV}/c$ as obtained from a heavy-flavour enriched MC simulation of pp collisions at $\sqrt{s} = 7 \text{ TeV}$

produced with the PYTHIA generator. The resolution on the impact parameter of the D meson improves by a factor of 2, as shown by the reduced width of the distribution for prompt D^0 (see also [36]). For example, the probability to reconstruct an impact parameter of $100 \mu\text{m}$ for prompt D^0 decreases by a factor larger than 3 in the upgraded ITS case compared to the current ITS case. Such a decrease allows for a much better separation of the two components with smaller systematic uncertainties.

An estimation of the performance for the measurement of beauty production in central Pb–Pb collisions using the fraction of non-prompt D^0 mesons was carried out, starting from the simulation results on the impact parameter resolution, the $D^0 S/B$ ratio, and the expected D^0 signal statistics. For the latter, we have considered a sample of 10^8 , 10^9 , or 1.7×10^{10} events in the centrality class 0–20% (the last value corresponds to an integrated luminosity of 10 nb^{-1}). An additional input parameter of the study is the value of the fraction of non-prompt D mesons, for which we have considered two cases, 0.10 and 0.20, for illustration, and then used the p_T dependence observed in proton–proton data. The relative statistical uncertainty on the fraction of D^0 mesons from B decays is shown in the right panel of Figure 2.12. Since the statistical uncertainty on the measurement of the total D^0 yields is expected to be of the order of 1%, the values shown in the figure coincide in practice with the relative statistical uncertainty on the measurement of beauty production. The results are very promising, with a statistical uncertainty smaller than 10% down to $D^0 p_T$ of $1 \text{ GeV}/c$, for 10 nb^{-1} . Since the mass difference of B and D mesons is larger than $3 \text{ GeV}/c^2$, the decay of a B meson with transverse momentum below $1 \text{ GeV}/c$ can yield a D meson with $p_T > 2 \text{ GeV}/c$. Therefore, the measurement will give access to B mesons with p_T down to almost 0.

Figure 2.13 (left) shows the comparison of the performance expected with the current ITS in the *no high-rate* scenario (0.1 nb^{-1} , i.e. 1.7×10^8 events in 0–20%), and with the upgraded ITS in both the *no high-rate* and *high-rate* (10 nb^{-1} , i.e. 1.7×10^{10} events in 0–20%) scenarios. The statistical uncertainty in the *no-high rate* scenario, reported in the left-hand panel, does not allow to extract $f_{\text{feed-down}}$ with the current ITS and limits the feasibility to $p_T > 3 \text{ GeV}/c$, with a 10% uncertainty, also in the upgrade case. Conversely, in the *high-rate* scenario the statistical precision allows to perform this measurement down to $p_T = 1 \text{ GeV}/c$. However, the applicability of this method below $2 \text{ GeV}/c$ is still under study and might be limited by the short decay length of low- p_T B mesons in the laboratory (small Lorentz boost).

Systematic uncertainties were studied by biasing the input ingredients/parameters of the fit procedure in the simulation, and evaluated by comparing the mean and the sigma of the residual distributions with and without a given bias. The following sources were considered:

- Yield extraction: the subtraction of the impact parameter distribution of the background in the D^0 mass signal region is done using the impact parameter distribution in the side bands, normalized to the amount of background in the signal region. While the statistical uncertainties on the signal and background yields recovered from the mass fit will be negligible for $p_T > 2 \text{ GeV}/c$, the expected $\delta S/S \sim 1\text{--}2\%$ systematic uncertainty on the signal can influence the impact parameter fit. This effect was studied by varying the amount of background subtracted, considering that $\delta B/B = \delta S/S \times S/B$, utilizing the values of the signal-to-background ratio expected after the upgrade of the ITS.
- Monte Carlo resolution and Monte Carlo p_T shape: in the fit procedure, the template distributions of the reconstructed impact parameter of prompt D^0 and of the true impact parameter of secondary D^0 are fitted and the fit functions are used as ingredients in the convolution performed in the final fit, which is therefore sensitive to their shape. A discrepancy between the actual distributions and the Monte Carlo ones can derive from a not precise description of the spatial resolutions, due for instance to misalignment, and from a difference of the B and D meson p_T spectra in Monte Carlo and data. The latter can be relevant because the spatial resolution is p_T dependent and because the exponential shape of the distribution of the true impact parameter of secondary D^0 is sensitive to the decay length of the parent B meson. In order to study the effect of a different track spatial resolution in data and Monte Carlo, we used template histograms obtained from a simulation in which the track impact parameter resolution was deteriorated by 10% at low p_T and by $\sim 30\%$ at high p_T to simulate the data impact parameter distributions, and we used the standard distributions as input for the fit.

The p_T dependence of each uncertainty considered is reported in the right-hand panel of Figure 2.13.




Figure 2.14: Signal-to-background ratio (left panel) and significance per event (right panel) for $\Lambda_c \rightarrow pK^-\pi^+$ in Pb-Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV, for current and upgraded ITS with 0.3% x/X_0 layer thickness. The results are shown for 0–5% central collisions and peripheral (70–80%) collisions.

Charm Baryons: $\Lambda_c \rightarrow pK^-\pi^+$

The most promising heavy-flavour baryon measurement is the decay of the Λ_c^+ into three charged prongs (p , K^- and π^+) with a B.R. of about 5.0% [44]. In order to identify the decay vertex, a very high resolution is needed because of the short mean proper decay length of the Λ_c ($c\tau \approx 60 \mu\text{m}$ [44]). Therefore, an improvement of the resolution would allow a much cleaner separation of its decay point (the secondary vertex) from the interaction point (primary vertex) with respect to the current ITS.

The decay channel of $\Lambda_c \rightarrow pK^-\pi^+$ is studied by analyzing the invariant mass of fully-reconstructed three-prong decays, selected by applying topological cuts and particle identification criteria, namely the proton and kaon identification using the TPC and TOF information.

Presently a Λ_c signal is visible in the $pK^-\pi^+$ invariant mass distribution obtained from a data sample of 1.9×10^8 proton–proton events at $\sqrt{s} = 7$ TeV collected with the ALICE minimum bias trigger. However, the significance is only ~ 5 due to the limited efficiency for background rejection with the current ITS.

For the performance study in pp collisions with the upgraded ITS, pp events produced with the PYTHIA generator were used. In this way, for the ITS upgrade scenario we obtain for Λ_c baryons with $p_T > 3$ GeV/c a significance of about 12. The increase of the signal statistics is $\sim 50\%$ (less stringent cuts can be used for the upgrade case) and the increase of the signal-to-background ratio is more than a factor of 5.

The most dramatic improvement is obtained for Pb–Pb collisions. The Λ_c signal could not be observed with the 2010 Pb–Pb data sample, because of the very large combinatorial background. We have studied the selection strategy using a dedicated simulation sample of about 10^4 central (0–10%) Pb–Pb events at $\sqrt{s_{NN}} = 5.5$ TeV produced with the HIJING event generator. In each event, 20 $\Lambda_c \rightarrow pK^-\pi^+$ decays were added, in order to obtain a high-statistics signal sample allowing for the study of the selection cuts. The signal per event was then rescaled to the expected $dN/dy = 1.4$ (see Ref. [36]) and also the branching ratio (5%) normalization was accounted for.

The performance in non-central collisions was estimated by scaling the signal and background yields with centrality and evaluating then the signal-to-background ratio and the significance (see Ref. [36] for more details). The results, in terms of S/B ratio and significance per event, are reported in Figure 2.14 as a function of p_T for different centrality classes and for minimum bias collisions. For the most central collisions, the signal-to-background improves by a factor 400 (in $2 < p_T < 4$ GeV/c) from current to upgraded ITS. The significance also improves by a factor 5–10 in all p_T intervals above 2 GeV/c. The performance in peripheral collisions is better, in terms of significance, by about a factor of 2. The signal-to-background ratio becomes quite large in peripheral collisions.

In the following we will consider an integrated luminosity of 10 nb^{-1} , which corresponds to 1.7×10^{10} central events (using the centrality class 0–20%). For this event sample, the significance is 1.3×10^5 times the values reported in the figure: 7, 40 and 53 in the p_T bins 2–4 GeV/c, 4–6 GeV/c and 6–8 GeV/c respectively. For




Figure 2.15: Comparison of statistical precision for different cases of ITS performance and integrated luminosity: current and upgraded ITS; *high-rate* scenario, i.e. 50 kHz readout rate that provides 10 nb^{-1} ; *no high-rate* scenario, i.e. 500 Hz readout rate that provides 0.1 nb^{-1} .

the peripheral collisions the significance is 20 in $2 < p_T < 4 \text{ GeV}/c$ for 1.7×10^{10} events (in a centrality class like 70–90% or 60–80%). Figure 2.15 shows the comparison of statistical precisions for the current and upgraded ITS in two scenarios of integrated luminosity: 10 nb^{-1} (*high-rate* capability: 50 kHz) and 0.1 nb^{-1} (*no high-rate* capability: 500 Hz, current TPC limit). Since the integrated luminosity differs by a factor 100 in the two scenarios, the precision differs by a factor 10. With the upgraded ITS in the *no high-rate* scenario, the precision is marginal in the whole p_T range. In the *high-rate* scenario, instead, Λ_c production is measurable down to a transverse momentum of $2 \text{ GeV}/c$ in central collisions. In addition, a measurement of the Λ_c elliptic flow is feasible for $p_T > 4 \text{ GeV}/c$ in semi-central collisions (e.g. 30–50%) and for $p_T > 2 \text{ GeV}/c$ in peripheral collisions (e.g. 60–80%) [36].

The signal-to-background ratio is expected to be very small ($\sim 10^{-4}$ in $2\text{--}4 \text{ GeV}/c$). Therefore, the signal yield extraction from the invariant mass distributions will require a dedicated strategy to minimize the systematic effects from background fluctuations and correlated backgrounds. This aspect is discussed in detail in the ITS Upgrade CDR [36]. In the following, a summary of this study is reported.

Two main sources of background have been considered: the combinations of three uncorrelated tracks and the combinations of two or three tracks coming from a common particle decay (correlated background). The uncorrelated background accounts for a large fraction of the total background. This contribution can be estimated precisely using the event mixing technique, and subtracted from the invariant mass distribution prior to the fit that is used to extract the signal yield. The correlated background can feature structures in the invariant mass distribution that can affect the signal extraction, if they are comparable with the position and width of the Λ_c signal. In this first study, the following cases were considered:

- 2 tracks from a common D meson decay (in particular D^+ , D^{*+} , D_s^+ and D^0 were considered separately) and 1 uncorrelated track;
- 2 tracks from a common Λ_c decay and 1 uncorrelated track;
- 2 tracks from a common resonance decay (K^{0*} , Δ^{++} , $\Lambda(1520)$, ρ , η , η' , ω , ϕ were considered) and 1 uncorrelated track;
- 3 tracks from a common D meson decay.




Figure 2.16: $[pK\pi]$ invariant mass distributions of the correlated background candidates (4–6 GeV/c) with 2 (left panel) and 3 (right panel) tracks from a common D meson decay. The dashed lines show the position of the Λ_c mass and a $\pm 3\sigma$ range.

Table 2.1: Breakdown of the systematic uncertainty of the various Λ_c measurements (the symbol \oplus indicates quadratic sum).

Contribution	Systematic uncertainty (%)					
	$N_{\Lambda_c}^{\text{Pb-Pb}}$	$N_{\Lambda_c}^{\text{pp}}$	R_{AA}	$(\Lambda_c/D^0)_{\text{Pb-Pb}}$	$(\Lambda_c/D^0)_{\text{pp}}$	$\frac{(\Lambda_c/D^0)_{\text{Pb-Pb}}}{(\Lambda_c/D^0)_{\text{pp}}}$
Tracking efficiency	12	12	$12 \oplus 12$	4	4	$4 \oplus 4$
PID	7	7	$7 \oplus 7$	5	5	$5 \oplus 5$
Yield extraction	5	5	$5 \oplus 5$	$5 \oplus 1$	$5 \oplus 1$	$5 \oplus 1 \oplus 5 \oplus 1$
MC correction	8	8	$8 \oplus 8$	$8 \oplus 8$	$8 \oplus 8$	$8 \oplus 8 \oplus 8 \oplus 8$
MC p_{T} shape	2	2	$2 \oplus 2$	$2 \oplus 2$	$2 \oplus 2$	$2 \oplus 2 \oplus 2 \oplus 2$
Feed-down	20	10	17	$20 \oplus 3$	$10 \oplus 3$	$17 \oplus 3 \oplus 3$
B.R.	26	26	0	$26 \oplus 1$	$26 \oplus 1$	0
Total (excl. B.R.)	26	26	30	25	18	27
Total (incl. B.R.)	37	37	30	36	32	27

In the simulated sample, the charm hadrons (D and Λ_c) were forced to decay in decay channels with only charged hadrons in the final state, because these are more likely to yield a large invariant mass, close to the Λ_c mass of $2.286 \text{ GeV}/c^2$, when the $[pK\pi]$ mass hypotheses are used for the three tracks. These decay channels have typical branching ratios of the order of 5–10%. It was found that for all background sources with 2 correlated tracks the $[pK\pi]$ invariant mass distribution has a smooth shape without structures in the Λ_c mass region. An example, for the case of 2 tracks from a common D meson decay, is shown in Figure 2.16 (left). The Λ_c mass region is displayed as a $2.286 \text{ GeV}/c^2 \pm 3\sigma$ band (where $\sigma = 8 \text{ MeV}/c^2$ is the expected mass resolution for this p_{T} interval). On the other hand, the last of the listed background sources (3 tracks from a common D meson decay) presents a structure roughly centered at the Λ_c mass (Figure 2.16, right). However, a Gaussian fit indicates that the structure is about 10 times broader than the Λ_c signal distribution ($\approx 70 \text{ MeV}/c^2$ vs. $8 \text{ MeV}/c^2$). In addition, it was estimated that this background component has a total yield 10 times smaller than that of the Λ_c signal. This large suppression factor is most likely induced by the PID selection: one of the three tracks is required to be compatible with a proton; since D mesons do not decay to protons, at low and intermediate p_{T} this selection is very efficient in the rejection of this kind of background. In summary, the study indicates that correlated background sources are not expected to introduce a systematic bias for the Λ_c signal extraction.

Since Λ_c production was not measured in Pb–Pb so far, the estimation of the systematic uncertainty is made on the basis of the “experience” with D meson measurements and of the uncertainties expected for the D^0 case (described in the previous section). In the following, the different contributions are discussed, first for the Λ_c cross section, and then for the R_{AA} and the Λ_c/D^0 ratio. A summary is given in Table 2.1.

- Yield extraction: after the subtraction of the uncorrelated background with event-mixing, the correlated contributions are left. In the previous section it was shown that these are not expected to affect significantly the signal extraction via an invariant mass fit. Therefore, we assume that the yield extraction has the same systematic uncertainty as for the D^0 analysis.
- Tacking efficiency: due to the presence of three prongs we assumed three times the single track uncertainty (4%).
- PID efficiency: for this analysis tight PID cuts are applied for two tracks, the kaon and the proton. We, therefore, assume twice the systematic uncertainty expected for the D^0 analysis, where tight cuts are applied on one track (kaon).
- Selection efficiency: this correction has turned out to be of the same order for D^0 , D^+ and D^{*+} analyses; we assume that it will be similar for the Λ_c .
- Feed-down from Λ_b decays: we assume a systematic uncertainty similar that on D mesons with the current data [38]. Due to the low signal-to-background ratio, it is unlikely that, even in the high-rate scenario, the feed-down contribution can be estimated using the impact parameter fit technique as planned for D mesons (see previous section). Therefore, we plan to proceed as it is done on the current data for the D analyses, that is, to rely on an hypothesis on the Λ_b production. For the Pb–Pb case, this implies an assumption on the $\Lambda_b R_{AA}$. If there will not be any Λ_b measurement available in Pb–Pb, we will base our hypothesis on the measurement of the B R_{AA} , varying the Λ_b/B ratio in a reasonable range, e. g. $1 < R_{AA}(\Lambda_b)/R_{AA}(B) < 2$ (similar to that of the $\Lambda/K^0 R_{AA}$ ratio). If the fraction of secondary Λ_c is kept at the level of 10–15%, similar to the current values of D analyses, it is likely that the overall uncertainty is contained within about 20%.
- Branching ratio (B.R.): this is the uncertainty from [44].

In the case of ratios ($\Lambda_c R_{AA}$ and Λ_c/D^0), some contributions to the systematic uncertainty cancel. In particular, does the tracking efficiency in case of the Λ_c/D^0 ratio for two of the three tracks. This is not assumed in the case of R_{AA} because of the different collision environments. In the R_{AA} and the double ratio $\frac{(\Lambda_c/D^0)_{\text{Pb-Pb}}}{(\Lambda_c/D^0)_{\text{pp}}}$, the branching ratio uncertainties cancel.

Ongoing Studies: Reconstruction of D_s^+ Mesons

As discussed in detail at the beginning of this section, the study of the D_s^+ production in heavy-ion collisions allows one to measure the effects of the strangeness enhancement on charmed-strange mesons, which is predicted to play a role in the intermediate transverse momentum region (up to a factor 2 enhancement of D_s^+/D) [23, 26, 27]. D_s^+ mesons and their antiparticles ($c\tau \approx 150 \mu\text{m}$) are reconstructed in ALICE in the decay chain $D_s^+ \rightarrow \phi\pi^+$ (and its charge conjugate) followed by $\phi \rightarrow K^-K^+$. The branching ratio of the chain $D_s^+ \rightarrow \phi\pi^+ \rightarrow K^-K^+\pi^+$ is $2.28 \pm 0.12\%$ [44].

In order to reduce the large combinatorial background, D_s^+ candidates are selected according to topological cuts, e.g. decay length, invariant mass of the reconstructed ϕ meson and the cosine of the pointing angle, which is the angle between the reconstructed D_s^+ meson momentum and the line connecting the primary and secondary vertex. A crucial role in the background rejection is played by the particle identification, which is based on measurement of the specific energy loss, dE/dx , from the TPC and the time-of-flight from the TOF detector. Compatibility cuts are applied on the difference between the measured signals and those expected for a pion or a kaon. The PID selection provides a background rejection of a factor 20–30 at low p_T , as shown in Figure 2.17 for the p_T intervals 3–4 and 4–6 GeV/c .

The D_s^+ p_T -differential cross section was measured in pp collisions at $\sqrt{s} = 7 \text{ TeV}$ in the p_T range from 2 to 12 GeV/c with a statistical significance of 4–5 depending on the transverse momentum [45]. Preliminary results for Pb–Pb collisions at 2.76 TeV in the centrality range 0–7.5% show a significance of 3–4 in the interval 4–12 GeV/c [29].

In order to address the performance for D_s^+ measurements with *high-rate* capabilities and the upgraded ITS, we have carried out two first studies, described in the following:

- Effect of *high-rate* capabilities: extrapolation of the performance on the 2011 Pb–Pb data sample to the target integrated luminosity after LS2.




Figure 2.17: Background distributions of D_s^+ candidates in two p_T intervals, 3–4 and 4–6 GeV/c with and without PID selection applied.

Table 2.2: Statistical significances expected for the D_s^+ measurement for central Pb–Pb events (0–7.5%), for the 2011 sample and scaled to two scenarios of integrated luminosity.

Integrated luminosity	Significance		
	p_T interval (GeV/c)		
	4–6	6–8	8–12
$28 \mu\text{b}^{-1}$ (Pb–Pb 2011) [29]	4.0 ± 1.0	3.3 ± 0.9	4.2 ± 0.9
0.1 nb^{-1} (in 1 GeV/c bins)	7	5	5
10 nb^{-1} (in 1 GeV/c bins)	70	50	50

– Effect of upgraded ITS: preliminary study based on simulation with the *Hybrid* approach.

Effect of high-rate capabilities

We present here an estimate of the expected statistical significance for the two scenarios of integrated luminosity: 0.1 nb^{-1} (*no high-rate*) and 10 nb^{-1} (*high-rate*). The values have been obtained by scaling with the square root of the number of the events the significances of the 2011 Pb–Pb data.

In Table 2.2 the statistical significance in the centrality range 0–7.5% is presented for different p_T intervals. The calculation assumes the signal extraction performed in bins of 1 GeV/c . The scaled significances are 5–7 and 50–70, for the two cases of integrated luminosity, equivalent to a statistical uncertainty associated to the measurement of about 20% in the *no-high rate* scenario and of about 2% in the *high-rate* scenario.

Recent theoretical calculations [23] indicate that the measurement of D_s^+ elliptic flow can provide interesting information concerning the effects of diffusion in the hadronic phase. In these calculations, the v_2 of D_s^+ mesons is predicted to be reduced by 30–40% in the intermediate p_T region with respect to that of non-strange D mesons. Also in this case, in order to appreciate this difference, a very precise measurement is required. We provide here an estimation of the statistical significances expected for the v_2 measurements starting from the first preliminary results of signal extraction in the centrality range 20–40% obtained on the 2011 Pb–Pb data sample. In the two scenarios of integrated luminosity (0.1 and 10 nb^{-1}), the significances expected in the p_T range 3–5 GeV/c are about 7–8 and 70–80, respectively, for p_T bins of 1 GeV/c . The values for the *no high-rate case* are not sufficient to perform a φ -differential analysis and obtain a precise measurement of v_2 .

We conclude that both the analyses, R_{AA} and v_2 , would largely benefit from the high statistics provided by the




Figure 2.18: Left: example of $[\text{KK}\pi]$ invariant mass distribution in a charm-enriched MC sample analyzed with the *Hybrid* approach. Right: D_s^+ statistical significance and signal-to-background ratio with upgraded ITS divided by the values obtained with the current ITS, as a function of p_T .

high-rate upgrade of the ALICE central barrel.

Effect of upgraded ITS with improved resolution

The performance in terms of significance and S/B ratio has been studied with the current and upgraded ITS configurations, applied via the *Hybrid* approach. HIJING simulations for Pb-Pb collisions in the centrality range 0–10% were used, with additional D_s^+ signals embedded to increase the statistics. An example of $[\text{KK}\pi]$ invariant mass distribution in the range $2 < p_T < 3$ GeV/c is shown in Figure 2.18 (left). In right-hand panel of the same figure, the significance and the signal-to-background ratio with upgraded ITS divided by the values obtained with the current ITS are presented in the p_T range 2–24 GeV/c. A large increase of both significance and S/B is observed in the low-intermediate momentum region with the upgraded detector. This result is particularly promising, since in this p_T range a very precise measurement is required to distinguish between different theoretical models. The study will be soon extended down to zero p_T .

2.1.1.2 Physics Performance: Charm Baryon/Meson, Heavy-Flavour Flow

The Λ_c/D^0 Ratio

In the previous section we have shown that Λ_c production can be measured in central and peripheral collisions for $p_T > 2$ GeV/c. Figure 2.19 (left) presents the expected statistical uncertainties for the measurement of the Λ_c/D^0 ratio using 1.7×10^{10} central Pb-Pb collisions (0–20%), corresponding to an integrated luminosity of 10 nb^{-1} . The points are drawn on a line that captures the trend and magnitude of the Λ_c/K_S^0 ratio. The proton–proton expectation from the PYTHIA 6.4.21 generator [46] is also shown.

Figure 2.19 (right) shows the enhancement of the Λ_c/D ratio in central Pb-Pb (0–20% for $L_{\text{int}} = 10 \text{ nb}^{-1}$) with respect to pp collisions. Statistical and systematic uncertainties (from Table 2.1) are shown. It is assumed that the statistical uncertainties for the D^0 measurements (see Figure 2.11-left) and for the Λ_c measurement in pp (see discussion in Section 2.1.3) are negligible with respect to those for the Λ_c measurement in Pb-Pb. The points are drawn on a line that captures the trend and magnitude of the Λ_c/K_S^0 double-ratio. The two model calculations presented in Figures 2.5 [27] and 2.6 [24] are shown to illustrate the expected sensitivity of the measurement.




Figure 2.19: Left: estimated statistical uncertainties on the measurement of the Λ_c/D^0 ratio using 1.7×10^{10} central Pb–Pb collisions (0–20%), corresponding to an integrated luminosity of 10 nb^{-1} . The points are drawn on a line that captures the trend and magnitude of the Λ/K^0_s ratio (see Figure 2.4). The pp expectation from the PYTHIA 6.4.21 generator [46] is also shown. Right: enhancement of the Λ_c/D ratio in central Pb–Pb with respect to pp collisions. Two model calculations [24, 27] are also shown.




Figure 2.20: Estimated statistical uncertainties on v_2 of prompt and secondary D^0 mesons for $1.7 \cdot 10^{10}$ events (left) in the 30–50% centrality class, which correspond to 10 nb^{-1} , and for $1.7 \cdot 10^8$ events (right), which correspond to about 0.1 nb^{-1} .

Heavy-flavour Elliptic Flow

We have estimated the expected precision on the measurement of v_2 for prompt and secondary D mesons with the upgraded ITS and with an integrated luminosity of 10 nb^{-1} . We have scaled the statistical uncertainties obtained from the simulation studies (Section 2.1.1.1), considering that the significance/event is the same for central (0–20%) and semi-central events (e.g. 30–50%). This feature is observed in the D^0 analysis from 2010 data [38].

The v_2 extraction can be performed using the raw signal yields in two large intervals of azimuthal angle φ with respect to the Event Plane (EP) direction Ψ_{EP} , determined for each collision [47]: $[-\pi/4 < \Delta\varphi < \pi/4] \cup [3\pi/4 < \Delta\varphi < 5\pi/4]$ (in-plane) and $[\pi/4 < \Delta\varphi < 3\pi/4] \cup [5\pi/4 < \Delta\varphi < 7\pi/4]$ (out-of-plane). Given the in-plane and out-of-plane yields, N_{in} and N_{out} , one has $v_2 = (\pi/4) \cdot (N_{\text{in}} - N_{\text{out}}) / (N_{\text{in}} + N_{\text{out}})$. In order to measure separately v_2 for prompt (charm) and secondary (beauty) D mesons, the prompt fraction will be determined for the in-plane and




Figure 2.21: Left: Prediction for R_{AA} vs p_T for D mesons (blue) and B mesons (red) in Pb–Pb collisions at the LHC, from radiative + collisions energy loss (DGHW) [50]. Right: Mass dependence of B meson R_{AA} in Pb–Pb collisions at the LHC from radiative energy loss (ASW) [51].

out-of-plane signal using the D^0 impact parameter fit, as described in Section 2.1.1.1.

We have estimated the statistical uncertainties on v_2 considering that the relative statistical uncertainties on N_{in} and N_{out} are $\sqrt{2}/(1 \pm v_2)$ times larger, respectively, than those on the total raw yield and assuming $v_2(p_T)$ for prompt and secondary D mesons as in the predictions of the BAMPS model [14], shown in the top-right panel of Figure 2.2 (the v_2 for non-prompt J/ψ is used for D^0 from B decays). The resulting statistical uncertainties are reported in Figure 2.20 for $1.7 \cdot 10^{10}$ events in the 30–50% centrality class, which correspond to 10 nb^{-1} (*high-rate*), and for $1.7 \cdot 10^8$ events, which correspond to about 0.1 nb^{-1} (*no high-rate*). The systematic uncertainties can be expected to be rather small, since most of them are common for the N_{in} and N_{out} raw yields and cancel in the v_2 ratio.

2.1.2 Heavy-flavour Energy Loss

One of the distinctive features of the hot and dense medium formed in heavy ion collisions is the strong energy loss induced on the hard partons that are produced in the initial hard scattering processes. Parton energy loss is thought to be dominated by gluon radiation, but also elastic collisions with the medium gluons should play an important role. The investigation of heavy-flavour energy loss has a particular interest. Indeed, gluon radiation from heavy quarks is predicted to be suppressed, with respect to the case of light partons, at angles smaller than the quark energy-to-mass ratio (dead cone effect) [48, 49]. Moreover, light-flavour hadrons are dominantly produced at LHC energies by hard fragmenting gluons, which lose more energy due to their stronger coupling to the medium. Thus, one has the prediction for the energy loss $\Delta E_g > \Delta E_c > \Delta E_b$. Experimentally, the energy loss in heavy ion collisions can be investigated as a function of transverse momentum via the nuclear modification factor R_{AA} .

Theoretical models based on perturbative QCD with the inclusion of radiative parton energy loss predict for charm mesons a suppression factor of 3–5 and a significantly smaller suppression for B mesons (see left panel of Figure 2.21 [50]). The mass dependence of energy loss is more pronounced for the beauty, as seen in the right panel of Figure 2.21 [51]. Mass and colour charge dependence can be investigated experimentally with the heavy-to-light ratios $R_{D/h} = R_{AA}^D(p_T)/R_{AA}^h(p_T)$ and $R_{B/h}$ [51].

Even more interesting is the ratio $R_{B/D} = R_{AA}^B(p_T)/R_{AA}^D(p_T)$, which is shown in the left panel of Figure 2.22 [51]. Here, the ratio shows a strong deviation from one in particular at moderately low p_T , decreasing only rather slowly for increasing p_T . String theory models inspired by the AdS/CFT correspondence have been able to describe qualitatively a number of aspects of the heavy ion collisions phenomenology (see e.g. [53]). Predictions for the ratio $R_{D/B} = 1/R_{B/D}$ are shown in the right panel of Figure 2.22 [52]. This ratio has also the advantage of magnifying the differences in the mass and p_T dependence of pQCD and AdS/CFT models. While different settings of pQCD models or AdS/CFT models, respectively, lead to different results, the two classes of models yield largely different predictions irrespectively of their input parameters.




Figure 2.22: Predictions for the comparison of charm and beauty R_{AA} in Pb–Pb collisions at the LHC. Left: R_B^B/R_{AA}^D vs p_T from radiative energy loss (ASW) [51]. Right: $R_{AA}^c(p_T)/R_{AA}^D(p_T)$ vs p_T from radiative energy loss (pQCD) and AdS/CFT [52].




Figure 2.23: Heavy-flavour nuclear modification factors at mid-rapidity in central Pb–Pb collisions at the LHC measured by ALICE. Left: D meson R_{AA} in the centrality classes 0–20% (2010 data) [38] and 0–7.5% (2011 data). Right: R_{AA} for the heavy-flavour decay electrons and heavy-flavour decay muons [54] in the centrality class 0–10%.

Using data from 2010 and 2011 runs at the LHC with Pb–Pb collisions at 2.76 TeV, ALICE performed the first measurement of the D meson R_{AA} , which is shown in Figure 2.23. With 2010 data the analysis was restricted to $p_T > 2$ GeV/c [38]. It was possible to go down to 1 GeV/c using higher-statistics data from the 2011 Pb–Pb run, but reaching zero transverse momentum seems to be precluded with the current setup, due to the huge background level.

The other key measurement is a precision measurement of beauty energy loss via R_{AA} with coverage down to low p_T . In ALICE, beauty production is accessed at mid-rapidity via $B \rightarrow e + X$, using electrons selected on the basis of their displacement from the primary vertex. In the present setup, this will be the only way to measure it. At low p_T , the residual component of electrons from charm decays has to be subtracted statistically, implying a significant systematic uncertainty.




Figure 2.24: Nuclear modification factor of D^0 from b-hadron decay for (left) 8.5×10^7 (no high-rate scenario) and (right) 8.5×10^9 (high-rate scenario) Pb-Pb events in the centrality range 0–10%.




Figure 2.25: Ratio of the nuclear modification factors of prompt D mesons and D mesons from b-hadron decay in the no high-rate (left) and high-rate (right) scenarios for Pb-Pb collisions in the centrality range 0–10%.

The R_{AA} for the heavy-flavour decay electron spectrum at mid-rapidity and for the heavy-flavour decay muon spectrum at forward rapidity [54] are shown in Figure 2.23 (right).

Another possibility is to measure beauty via the J/ψ decay, tagging a secondary vertex with two leptons attached. A first measurement in Pb-Pb was performed by the CMS Collaboration [55]. This measurement is limited to the region $J/\psi p_T > 6.5 \text{ GeV}/c$, due to the large momentum cut on the decay muons.

2.1.2.1 Physics Performance: Comparison of B and D Nuclear Modification Factors

The expected performance on the nuclear modification factor of D^0 from B decays is shown in Figure 2.24 for the no high-rate (left) and high-rate (right) scenarios. The central values were obtained by scaling the nuclear modification factor of the prompt D^0 by the ratio predicted by the energy loss models shown in Figure 2.22

(left). The systematic uncertainties are the same as for prompt D mesons (Figure 2.11-right), with the feed-down uncertainty replaced by that estimated for $f_{\text{feed-down}}$ (Figure 2.13-right), multiplied by $\sqrt{2}$ (for pp and Pb–Pb). In the *high-rate* scenario, the measurement uncertainty is completely dominated by systematic uncertainties in the whole p_T range. In the *no high-rate* scenario, systematic uncertainties are larger than the statistical ones at low p_T , while for $p_T > 5 \text{ GeV}/c$ they have a comparable or smaller size.

The ratio of the nuclear modification factors of D^0 from B and prompt D^0 is shown in Figure 2.25 for the *no high-rate* (left) and *high-rate* (right) scenarios. In the former, the statistical uncertainties are very similar to the systematic uncertainties in essentially all p_T intervals. This indicates that this measurement would benefit substantially from the high-rate capabilities, because an integrated luminosity of at least 1 nb^{-1} (8.5×10^8 events in the class 0–10%) would make the statistical uncertainties subdominant.

2.1.3 Proton–Proton Running Requirements for Heavy-flavour Reference Data

We have studied the requirements for the proton–proton reference data to be used to investigate QGP-effects in heavy-flavour production. In particular:

- the beauty-to-charm suppression ratio $R_{\text{AA}}^{\text{feed-down}}/R_{\text{AA}}^{\text{prompt D}}$ requires the measurement of the production cross section of prompt and secondary D mesons;
- the search for the enhancement of the baryon-to-meson ratio for charm requires, in addition, the measurement of the Λ_c production cross section.

The requirements in terms of both integrated luminosity and centre-of-mass energy are assessed.

2.1.3.1 pp Integrated Luminosity

We start from the consideration that the statistical uncertainty on the pp reference has to be negligible with respect to that of the Pb–Pb measurement, e.g. $\sqrt{2}$ times smaller (so that the combined relative statistical uncertainty is about 20% larger than the Pb–Pb uncertainty). Since the relative statistical uncertainty is the inverse of the statistical significance $1/\mathcal{S} = \sqrt{S+B}/S$, the requirement is:

$$\mathcal{S}_{\text{pp}} = \sqrt{2} \cdot \mathcal{S}_{\text{Pb–Pb}}. \quad (2.1)$$

This condition leads to different requirements in terms of statistics, depending on the S/B ratio, thus on the background level. The two extreme cases are the following.

- *High-background measurements* (this is the case for charm production measurement, particularly at low transverse momentum): $S \ll B$, thus $\mathcal{S} = S/\sqrt{B}$, and the condition becomes

$$S_{\text{pp}}/\sqrt{B_{\text{pp}}} = \sqrt{2} \cdot S_{\text{Pb–Pb}}/\sqrt{B_{\text{Pb–Pb}}}$$

thus

$$N_{\text{pp}} = 2 \cdot N_{\text{Pb–Pb}} \cdot [(\mathcal{S}/\sqrt{N})_{\text{Pb–Pb}}/(\mathcal{S}/\sqrt{N})_{\text{pp}}]^2$$

with N the number of events in pp and Pb–Pb (in a given centrality class, e.g. 0–20%).

- *Background-free measurements* (this is the case e.g. for jets): $S \gg B$, thus $\mathcal{S} = \sqrt{S}/S = 1/\sqrt{S}$, and the condition becomes

$$S_{\text{pp}} = 2 \cdot S_{\text{Pb–Pb}}$$

thus

$$N_{\text{pp}} = 2 \cdot N_{\text{Pb–Pb}} \cdot [(S/\text{ev})_{\text{Pb–Pb}}/(S/\text{ev})_{\text{pp}}] \approx 2 \cdot N_{\text{Pb–Pb}} \cdot \langle N_{\text{coll}} \rangle$$

assuming approximate scaling of the signal per event with the number of binary nucleon–nucleon collisions; for central Pb–Pb events ($\langle N_{\text{coll}} \rangle \simeq 1500$):

$$N_{\text{pp}} \simeq 3 \cdot 10^3 \cdot N_{\text{Pb–Pb, 0–10\%}}.$$




Figure 2.26: Significance/event for D^0 reconstruction in pp collisions. Left: \mathcal{S}/\sqrt{N} for pp collisions for current ITS (from [41]) and upgraded ITS. Right: \mathcal{S}/\sqrt{N} for upgraded ITS in central Pb-Pb collisions (from Figure 2.10, right) and in pp collisions.

For the D meson production measurements performed with pp and Pb-Pb data collected in 2010 [38, 41], the ratio of $(\mathcal{S}/\sqrt{N})_{\text{Pb-Pb}}/(\mathcal{S}/\sqrt{N})_{\text{pp}}$ was about 5–7 depending on p_T (considering Pb-Pb collisions in the centrality class 0–20%). Therefore, the condition above would give $N_{\text{pp}} = 25\text{--}50 N_{\text{Pb-Pb}}$. The analysis was carried out with 300 million pp events and 3 million central Pb-Pb events, thus a ratio of 100. Indeed, the statistical uncertainty on the pp reference was much smaller (by a factor 2 typically) than that on Pb-Pb.

In the following, we estimate the required reference statistics for low- p_T charm measurements (*High-background case*)¹.

For the case of $D^0 \rightarrow K^- \pi^+$ reconstruction, we have estimated the value of \mathcal{S}/\sqrt{N} (significance per event) for pp collisions with the upgraded ITS using the following procedure.

- We started from the value of \mathcal{S}/\sqrt{N} obtained on pp data at $\sqrt{s} = 7$ TeV from the 2010 run [41].
- We assumed that the signal/event is the same with the upgraded ITS. This is justified by the observation that the efficiency is the same with current and upgraded ITS, for the same cut values (see Figure 2.18 of [37]).
- We assumed that the background/event is reduced, with the upgraded ITS, in pp by the same factor as in Pb-Pb, i.e. a factor about 5 (see Figure 2.10, left).

The left-panel of Figure 2.26 shows \mathcal{S}/\sqrt{N} for pp collisions for current ITS (from [41]) and upgraded ITS. The right-panel of the same figure shows \mathcal{S}/\sqrt{N} for upgraded ITS in central Pb-Pb collisions (from Figure 2.10, right) and in pp collisions. The ratio $(\mathcal{S}/\sqrt{N})_{\text{Pb-Pb}}/(\mathcal{S}/\sqrt{N})_{\text{pp}}$ is about 5. Therefore, the required number of pp events is about 50 times the number of central Pb-Pb events that will be used for the D^0 analysis.



If the Pb-Pb analysis will be carried out with 10 nb^{-1} (about $8 \cdot 10^9$ events in the centrality class 0–10%), the required reference pp sample is about $4 \cdot 10^{11}$ events, corresponding to $L_{\text{int,pp}} \approx 6 \text{ pb}^{-1}$. This sample can be collected in 10^6 s at 400 kHz, for example. If the D^0 analysis in Pb-Pb will be carried out with 2 nb^{-1} (about $1.5 \cdot 10^9$ events in the centrality class 0–10%), the required reference pp sample is about $8 \cdot 10^{10}$ events, corresponding to $L_{\text{int,pp}} \approx 1.2 \text{ pb}^{-1}$. This sample can be collected in 10^6 s at 80 kHz, for example.

For the case of $\Lambda_c \rightarrow p K^- \pi^+$ reconstruction, since the measurement was not yet carried with pp data, we base our considerations on different arguments. The possible enhancement of the ratio Λ_c/D will have to be studied in Pb-

¹For completeness, we work out the numbers also for the *Background-free case*: the condition $N_{\text{pp}} \simeq 3 \cdot 10^3 \cdot N_{\text{Pb-Pb, 0–10\%}}$ is equivalent to $L_{\text{int,pp}} \sigma_{\text{pp}}^{\text{inel}} \simeq 3 \cdot 10^3 \cdot L_{\text{int,Pb-Pb}} \cdot 10\% \cdot \sigma_{\text{Pb-Pb}}^{\text{hadronic}}$. Thus, using $\sigma_{\text{pp}}^{\text{inel}} = 70 \text{ mb}$ and $\sigma_{\text{Pb-Pb}}^{\text{hadronic}} = 7 \text{ b}$, $L_{\text{int,pp}} = 3 \cdot 10^4 L_{\text{int,Pb-Pb}}$. Hence, $L_{\text{int,pp}} = 300 \text{ pb}^{-1}$ for $L_{\text{int,Pb-Pb}} = 10 \text{ nb}^{-1}$. For example, this luminosity can be integrated in 10^7 s at a rate of 2 MHz.

Table 2.3: Summary of required proton–proton integrated luminosity.

Measurement	Pb–Pb L_{int}	pp L_{int}
D^0	10 nb^{-1}	6 pb^{-1}
D^0	2 nb^{-1}	1.2 pb^{-1}
Λ_c	10 nb^{-1}	0.6 pb^{-1}

**Figure 2.27:** Scaling factors for D (left) and B (right) mesons from $\sqrt{s} = 14$ to 5.5 TeV using FONLL.**Figure 2.28:** Relative scaling uncertainty for D (left) and B (right) mesons.

Pb collisions as a function of centrality, in comparison with the pp reference. Since the significance/event for Λ_c reconstruction is expected to be larger for more peripheral collisions (see Figure 2.14), the case of the most peripheral centrality class (e.g. 60–80%) dictates the required precision of the pp reference. We assume, conservatively, that the significance/event in pp will be the same as in peripheral Pb–Pb collisions, $(\mathcal{S}/\sqrt{N})_{\text{Pb–Pb,periph.}}/(\mathcal{S}/\sqrt{N})_{\text{pp}} = 1$. Thus, the required pp statistics is $N_{\text{pp}} = 2 \cdot N_{\text{Pb–Pb}}$. With $L_{\text{int,Pb–Pb}} = 10 \text{ nb}^{-1}$, the number of events in a centrality class covering 20% of the hadronic cross section is $1.7 \cdot 10^{10}$. Therefore, $N_{\text{pp}} \approx 4 \cdot 10^{10}$, corresponding to $L_{\text{int,pp}} \approx 0.6 \text{ pb}^{-1}$. This sample can be collected in 10^6 s at 40 kHz, for example.

The integrated luminosity requirements are summarized in Table 2.3.

2.1.3.2 pp Centre-of-Mass Energy

After the second long shut-down (LS2), the LHC is foreseen to accelerate the hadron beams at the nominal energy of $Z/A \cdot 7 \text{ TeV}$, and the collision centre-of-mass energy will be $\sqrt{s} = 14 \text{ TeV}$ for pp and $\sqrt{s_{\text{NN}}} = 5.5 \text{ TeV}$ for Pb–Pb. If the pp reference data are collected at 14 TeV, the charm and beauty production cross sections have to be scaled to the Pb–Pb energy, in order to define the nuclear modification factors $R_{\text{AA}}(p_T)$. The scaling factors can be obtained using perturbative QCD calculations, like FONLL [39], as it was already done for the first measurement of the D meson R_{AA} [38]. The definition of the scaling factor and of its theoretical uncertainty are described in detail in [56] for the case of $\sqrt{s} = 7 \text{ TeV}$ and 2.76 TeV. The uncertainty is constructed from the envelope of the

scaling factors obtained by varying the values of the perturbative scales (for factorization μ_F and renormalization μ_R) independently within the ranges $0.5m_T < \mu_{F,R} < 2m_T$ (m_T is the transverse mass of the heavy quark), with the constraint $0.5 < \mu_F/\mu_R < 2$. Figure 2.27 shows these envelopes for D mesons (left) and B mesons (right). The corresponding relative uncertainties are shown in Figure 2.28. For the case of D mesons at low p_T , the uncertainty is larger than 50% ($p_T < 2 \text{ GeV}/c$). This would be the dominant uncertainty in the measurement of the D meson nuclear modification factor and of the beauty-to-charm double-ratio $R_{AA}^{\text{feed-down}}/R_{AA}^{\text{prompt D}}$ at low p_T (see Figure 2.25). Therefore, the reference pp data should be collected at the same centre-of-mass energy as the Pb–Pb data.

2.2 Quarkonia

2.2.1 Motivation

Among the various suggested probes of deconfinement, charmonium plays a distinctive role. It is the first hadron for which a clear mechanism of suppression in QGP was proposed early on, based on the color analogue of Debye screening [57] (see [58] for a recent review). A suppression compared to pp collisions was observed in pA collisions, which was understood as a destruction of the pre-resonant $c\bar{c}$ state by the nucleons of the colliding nuclei. Beyond this "cold nuclear matter" effect, the measurements in Pb–Pb at the SPS ($\sqrt{s_{NN}} = 17.3 \text{ GeV}$) and in Au–Au at RHIC ($\sqrt{s_{NN}} = 200 \text{ GeV}$) demonstrated an "anomalous" suppression, attributed to the dense and hot QCD matter. The theoretical interpretation was a sequential suppression [59] of various charmonium states as a function of energy density (temperature) of QGP. In this picture, J/ψ survives in QGP and only the contribution from $\psi(2S)$ and χ_c charmonium states (leading to about 10% and 30% of the J/ψ yield, respectively) vanishes as a consequence of melting of $\psi(2S)$ and $\chi_{c1,2}$ states. As appealing as this interpretation is, it was questioned by the finite $\psi(2S)$ yield measured at the SPS [60] and also later by the re-evaluation of cold nuclear matter suppression derived from p–A measurements [61].

The $\psi(2S)$ yield relative to the yield of J/ψ was noted earlier [62, 63] to correspond, for central collisions, to the chemical freeze-out temperature derived from fits to other hadron abundances. This led to the idea of statistical hadronization of charm quarks in nucleus-nucleus collisions [63]. In this model, the charm quarks produced in initial hard collisions thermalize in QGP and are "distributed" into hadrons at chemical freeze-out. All charmonium states are assumed to be fully melted (or, more precisely, not formed at all) in QGP and produced, together with all other hadrons, exclusively at chemical freeze-out. This model has gained support from experimental data at RHIC [64, 65] (see also a recent review [66]). The model predicted a notably different pattern of J/ψ production at the LHC. Depending on the charm production cross section, even an enhanced production relative to pp collisions could be expected at the LHC [65] (see [67] for further predictions).

Proposed at the same time as the statistical hadronization model, the idea of kinetic recombination of charm and anti-charm quarks in QGP [68] is an alternative quarkonium production mechanism. In this model (see [69, 70] for recent results), a continuous dissociation and regeneration of charmonium takes place in the QGP over its entire lifetime. Besides the charm production cross section, this model has as input parameters the time dependence of the temperature in the fireball as well as relevant cross sections and assumptions on the melting scenarios of charmonium states. Important observables, like the transverse momentum (p_T) dependence of production yields and of elliptic flow can be calculated within the transport (kinetic) models. It predicted a rather small J/ψ regeneration component at RHIC and a sizable one at the LHC.

The measurement of J/ψ production in Pb–Pb collisions at the LHC was expected to provide a definitive answer on the question of (re)generation. The first data, measured at high- p_T [55, 71] showed a pronounced suppression. Subsequently, the data [55] also showed a J/ψ suppression of the same magnitude as that of open-charm hadrons [72], as shown in Figure 2.29. This may indicate that the high- p_T charm quarks that form either D or J/ψ mesons had the same dynamics including a thermalization stage and a late hadronization.

A clear difference to the measurement at RHIC [73, 74] was seen in the first LHC measurement of the overall (inclusive in p_T) production [75]. In this case, where low- p_T J/ψ production is dominant, less nuclear suppression (larger R_{AA} value) is seen at the LHC both at forward rapidity [76] and at mid-rapidity [77], Figure 2.30. Corroborated with the CMS data [55], this points to an enhanced production of J/ψ at low- p_T compared to high- p_T .

At the LHC, the estimated energy density is at least a factor of 3 larger than at RHIC [78], leading to an initial temperature above the one required for J/ψ dissociation. Therefore, one can conclude that the production mechanism of J/ψ and charmonium in general at the LHC is determined (to a large extent) by regeneration in QGP or by




Figure 2.29: Centrality dependence of the nuclear modification factor for D mesons and J/ψ at high- p_T at the LHC. The lines indicate the effect expected due to production in the corona of the colliding nuclei where cross sections in pp collisions are used.




Figure 2.30: Centrality dependence of the nuclear modification factor for inclusive J/ψ production. The ALICE measurement (preliminary) at the LHC is compared to the PHENIX data at RHIC. The two panels show the data at forward rapidity (left) and at mid-rapidity (right).

generation at chemical freeze-out.

Indeed, both the statistical hadronization [67] and transport [69, 70] models reproduce the data [75], as seen in Figure 2.31. Based on these observations, the J/ψ production can be considered a probe of QGP as initially proposed [57], but may not be a ‘‘thermometer’’ of the medium [59]. Within the statistical model, the charmonium states become probes of the phase boundary between QGP and hadron phase. This extends with a heavy quark the family of quarks employed for the determination of the hadronization temperature (via the conjectured connection to the chemical freeze-out temperature extracted from fits of statistical model calculations to hadron abundances).

The transverse momentum dependence of the nuclear modification factor, shown in Figure 2.32, is, at LHC, dramatically different than the one measured at RHIC. At low- p_T the nuclear suppression is significantly reduced (i.e. larger R_{AA}) at the LHC [76] compared to RHIC [73]. Transport model calculations reproduce the data quantitatively, as can be seen in Figure 2.32 (with model of ref. [70]). In current models [69, 70], about half of the low- p_T J/ψ yield in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV is produced by the recombination of charm quarks in QGP, while the rest is due to primordial production.

Both the kinetic and the statistical hadronization models require thermalization of the charm quarks in QGP. As a consequence they will follow the collective behavior of the bulk QGP and their flow will be reflected in that of




Figure 2.31: Centrality dependence of the nuclear modification factor for J/ψ at all momenta in comparison to theoretical models, for forward rapidity (left panel) and mid-rapidity (right panel).




Figure 2.32: Transverse momentum dependence of the nuclear modification factor for J/ψ at forward rapidity. Left panel: the preliminary ALICE data at the LHC, together with CMS data, are compared to measurements at RHIC by PHENIX. Right panel: the preliminary ALICE data in comparison with transport model calculations.

charmed hadrons and quarkonia. Indeed, elliptic flow of J/ψ at LHC energies was predicted within a transport model [69]. The first measurement at the LHC [79], shown in Figure 2.33, provides a tantalizing hint of a finite elliptic flow of J/ψ (the significance of non-zero flow for the centrality range 20-60% is 2.3σ in the p_T range 2-4 GeV/c). The preliminary ALICE data is consistent with the expectation from transport models, but precision data is needed in order to be able to extract information on the QGP properties and on the amount of J/ψ produced via regeneration.

We note that, at RHIC energies, the preliminary measurement by the STAR collaboration shows a J/ψ elliptic flow in Au-Au collisions at $\sqrt{s_{\text{NN}}} = 200 \text{ GeV}$ [80] consistent with zero.

The picture outlined above for the charmonium production, extracted by comparison of data and model predictions, remains to be tested by precision measurements at the LHC which further constrain model calculations. In particular, both the statistical and the transport models employ as input parameter the $c\bar{c}$ production cross section $\sigma_{c\bar{c}}$. The sensitivity of the calculations on $\sigma_{c\bar{c}}$ is rather large (see Figure 2.31). A precision measurement of $\sigma_{c\bar{c}}$ in Pb-Pb collisions, within reach with the proposed ALICE upgrade, will place an important constraint to models.

It is important to emphasize that the two competing models discussed above, despite providing similar predictions, have rather different underlying physics. While in the statistical model the hadronization is a process in which all quark flavors take part concurrently, in the kinetic model J/ψ survives as a hadron in the hot medium dominated by deconfined gluons and light quarks. In the statistical model all charmonium states are generated exclusively at hadronization, while in the kinetic model only at most half of the J/ψ yield originates from deconfined charm and anti-charm quarks. Discriminating the two pictures implies providing an answer to fundamental questions related




Figure 2.33: Elliptic flow of J/ψ as a function of the transverse momentum in 20-60% Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. The ALICE data is compared with parton transport model predictions.

to the fate of hadrons in a hot medium [81]. High-quality J/ψ data over broad ranges in p_T and y and a precision measurement of $\psi(2S)$ and χ_c charmonium states could allow this.

The LHC has opened up the measurement of the bottomonium family in Pb-Pb collisions, where a significant suppression of the excited states of Υ has been observed [82, 83]. The understanding and implications of these measurements is expected to evolve in the coming years, complementing the charmonium sector. The quality of the existing data in the bottomonium sector is already very good, due to rather straightforward triggering criteria. Further precision is expected from the measurements by the CMS collaboration. The ALICE measurements at forward y will contribute to a significant extension of the y range, complementing the range around mid-rapidity covered by CMS.




Figure 2.34: The centrality dependence of the production of $\psi(2S)$ relative to J/ψ . The double ratio in Pb-Pb relative to pp collisions at the LHC is shown for CMS data ($p_T > 3$ GeV/c) and ALICE data.

The production of $\psi(2S)$ is often studied relative to J/ψ . This relative production in Pb-Pb collisions normalized to the value in pp collisions gives the ratio of the nuclear modification factors, R_{AA} , of $\psi(2S)$ and J/ψ :

$$R = \frac{N_{\text{Pb-Pb}}^{\psi(2S)} / N_{\text{Pb-Pb}}^{J/\psi}}{N_{\text{pp}}^{\psi(2S)} / N_{\text{pp}}^{J/\psi}} = \frac{R_{\text{AA}}^{\psi(2S)}}{R_{\text{AA}}^{J/\psi}}, \quad (2.2)$$

where N denote production yields. This double ratio is shown in Figure 2.34. The measurement by the CMS collaboration [84], performed for $3 < p_T < 30$ GeV/c shows an enhancement of the $\psi(2S)$ production with respect




Figure 2.35: Energy dependence of production cross section for $\psi(2S)$ (left) and χ_c (right) states relative to J/ψ . The experimental data in pp and pA collisions are contrasted with the expectation from the statistical model in AA collisions.

to J/ψ in $Pb-Pb$ collisions compared to pp collisions. The preliminary ALICE data [85] seem to disfavor such an enhancement. Clearly, more precise data is needed in order to conclude. At SPS, the NA50 data [60] (dominated by low- p_T production) showed a double ratio well below unity.

The relative production of $\psi(2S)$ and χ_c charmonium states with respect to J/ψ is weakly-dependent (or even independent [86]) of charm production cross section employed in models for $Pb-Pb$ collisions. Such a measurement will be a crucial test of the picture outlined above.

In Figure 2.35 we show a compilation of data for the production cross section for $\psi(2S)$ and $\chi_{c1,2}$ states relative to J/ψ in pp and pA collisions [86]. The relative production cross sections of both $\psi(2S)$ and χ_c ($R_{\chi_c} = \sum_{j=1}^2 \sigma(\chi_{cj}) Br(\chi_{cj} \rightarrow J/\psi \gamma) / \sigma(J/\psi)$) states measured in $pp(\bar{p})$ and $p(\pi)A$ collisions cannot be described within the statistical model. This is in sharp contrast to the only existing measurement in central nucleus-nucleus collisions to date, performed for the $\psi(2S)$ at the SPS by the NA50 experiment [60], which is well described by the model.

The centrality dependence of the nuclear modification factor for charmonium states predicted within the transport model of Zhao and Rapp [70] is shown in Figure 2.36. The ratio of the nuclear modification factors of $\psi(2S)$ and χ_c to that of J/ψ is also shown in Figure 2.36, together with the prediction of the statistical hadronization model. The difference between the predictions of the two models for this double ratio is substantial. The high-precision data envisaged within the upgrade proposal will allow us to discriminate between models (see Section 2.2.7 below).

2.2.2 Inputs for Simulations

The following estimates of the performance of the ALICE measurement of J/ψ and $\psi(2S)$ in the Central Barrel ($|y| < 0.9$, measurement done in the e^+e^- channel) and in the Muon Spectrometer ($2.5 < y < 4.0$, in the $\mu^+\mu^-$ channel) are based on our current data in $Pb-Pb$ collisions at $\sqrt{s_{NN}} = 2.76$ TeV. The increase of the charmonium production cross section at $\sqrt{s_{NN}} = 5.5$ TeV will provide a better measurement accuracy than the one estimated in our present study. In the Central Barrel, J/ψ is currently measured utilizing only the TPC electron identification. The important contribution of the Transition Radiation Detector (TRD) for the electron identification is considered based on its present performance. The measured transverse momentum spectrum [85] is also used as input.

The measurement performance at mid-rapidity and at forward rapidities by ALICE have been estimated for integrated luminosities of 1 nb^{-1} and 10 nb^{-1} , corresponding to the data sample expected before and after the upgrade, respectively. For the Central Barrel, the impossibility of triggering (for low- p_T J/ψ), enforces the usage of minimum bias data. In this case, the limitations imposed by the current data readout speed are considered, while for the Muon Spectrometer the muon trigger allows an efficient sampling of all interactions. We provide our estimates




Figure 2.36: The centrality dependence of the nuclear modification factor for charmonium states predicted within the transport model. In the plot of the ratios of nuclear modification factors of $\psi(2S)$ and χ_c to that of J/ψ , shown in the lower right panel, the predictions of the statistical model are also included. The thinner lines are for the case of shadowing considered for the charm production cross section.

for the statistical errors and consider that the systematic errors can ultimately be decreased to match the statistical ones. A data sample in pp collisions of several pb^{-1} is needed in order to achieve a statistical significance better than that of the envisaged measurements with 10 nb^{-1} in Pb-Pb.

We focus below on the measurement of J/ψ and $\psi(2S)$. The detection of $\chi_{c1,2}$, possible in the Central Barrel via the photon reconstruction through conversions, needs a careful separate study. The challenge of this measurement, a first of this type in heavy-ion collisions, is embodied by the roughly 1000 $\chi_{c1,2}$ events expected to be reconstructed in 10 nb^{-1} Pb-Pb data.

2.2.3 J/ψ Yield and the Nuclear Modification Factor

From the data shown in Section 2.2.1, it is clear that ALICE will be able to perform a very good measurement of inclusive J/ψ yields and nuclear modification factor as a function of p_T and y with the Muon Spectrometer with the data collected prior to LS2. This is the case also for the Central Barrel where the usage of the complete TRD for electron identification leads to a good accuracy of the measurement up to about $6 \text{ GeV}/c$, see Figure 2.37.

An important capability of the Central Barrel is the measurement in Pb-Pb of J/ψ from B hadrons decays, covered in the open beauty section, which can only be realized with the full capability of the upgrade.

2.2.4 J/ψ Elliptic Flow

The expected absolute statistical error of the elliptic flow (v_2) measurement is shown in Figure 2.38 (left panel) for the Muon Spectrometer data. The measurement with 1 nb^{-1} data will improve significantly the present situation, but it is ultimately the upgrade data that will provide a precision measurement. The $J/\psi v_2$ could be measured with




Figure 2.37: The statistical accuracy of the J/ψ yield measurement in the Central Barrel as a function of transverse momentum for three centrality classes. The full symbols are for electron identification employing only the TPC, the open ones for including also TRD.

high precision as a function of rapidity, as illustrated in Figure 2.38 for the rapidity range $3.7 \leq y \leq 4.0$ (where the detector acceptance and the production cross sections are the lowest). The data collected in the upgrade scenario may even allow, via comparisons to transport models, to indirectly quantify the degree of B hadron flow via the measurement of $J/\psi v_2$ in the p_T range 4-6 GeV/c , where the model calculations (see Figure 2.33 above) show sensitivity to B hadron flow.

At mid-rapidity, the measurement of elliptic flow of J/ψ is only possible with the 10 nb^{-1} Pb-Pb data expected with the upgrade. As shown in Figure 2.38 (right panel), a measurement with a good significance over a broad range in p_T (up to 5-6 GeV/c) can be achieved only with the usage of the TRD electron identification. In this case, a direct measurement of elliptic flow of J/ψ from B hadrons could become also feasible.

2.2.5 J/ψ Polarization

Due to the J/ψ spin state, the distribution of its decay products can be expressed in its general form as [87]:

$$W(\theta, \phi) = \frac{1}{3 + \lambda_\theta} (1 + \lambda_\theta \cos^2 \theta + \lambda_\phi \sin^2 \theta \cos 2\phi + \lambda_{\theta\phi} \sin 2\theta \cos \phi) \quad (2.3)$$

where θ and ϕ are the polar and azimuthal angles, respectively, and λ parameters describe the spin state of the J/ψ in a given reference frame. The first measurement of the polarization parameters for inclusive J/ψ production in pp collisions at 7 TeV [88], carried out by ALICE in the kinematical region $2.5 < y < 4$, $2 < p_T < 8 \text{ GeV}/c$, showed that λ_θ and λ_ϕ are consistent with zero, in both the helicity and Collins-Soper reference frames.

It has been argued that the QGP is expected to screen the non-perturbative physics in the J/ψ production and consequently quarkonia escaping from the plasma should possess polarization as predicted by QCD models [89]. The measurement of the J/ψ polarization parameters in Pb-Pb collisions at low- p_T is a challenging measurement that ALICE will perform with the envisaged luminosity of 10 nb^{-1} . It is expected that statistical errors on λ_θ (see




Figure 2.38: The absolute statistical error of the elliptic flow of J/ψ as a function of transverse momentum for the measurement with the Muon Spectrometer (left panel, centrality range 20-60%) and for the Central Barrel (right panel, centrality range 10-40%).




Figure 2.39: Absolute statistical error of the J/ψ polarization parameters λ_θ (left panel) and λ_ϕ (right panel) as a function of centrality, measured with the Muon Spectrometer for 1 nb^{-1} and 10 nb^{-1} .

Figure 2.39) of about 0.02 will be reached in case of the Muon Spectrometer data for such integrated luminosities. A comparable, albeit slightly worse, accuracy is expected for the Central Barrel data.

2.2.6 Electromagnetic J/ψ Production

Exclusive vector meson production in ultra peripheral heavy-ion collisions (UPC) is expected to probe the nuclear gluon distribution [90], for which there is considerable uncertainty in the low- x region. In UPC the nuclei are separated by impact parameters larger than the sum of their radii and therefore hadronic interactions are strongly suppressed. The ALICE collaboration has made the first measurement at the LHC of J/ψ photo-production in Pb-Pb UPC at $\sqrt{s_{NN}} = 2.76$ TeV. The J/ψ is measured via its dimuon decay in the forward rapidity region with the Muon Spectrometer for events where the hadronic activity is required to be minimal. The analysis is based on an event sample corresponding to an integrated luminosity of $55 \mu\text{b}^{-1}$. The measured cross section for coherent J/ψ production is in good agreement with theoretical models for J/ψ production that include nuclear gluon shadowing.

Further measurements will be needed to improve the sensitivity to the gluon distribution function in nuclei. The measurement of J/ψ photo-production at mid-rapidity will allow us to constrain the gluon shadowing at Bjorken- x values of about $\sim 10^{-3}$. In addition, the precision of the measurement at forward rapidities will be dramatically improved with an integrated luminosity of 1 nb^{-1} . About 4000 photo-produced J/ψ in the p_T range 0–300 MeV/c would be then collected, allowing for a statistical precision in the range 2–4%. The systematic uncertainties are expected to match the ones obtained in pp and Pb-Pb hadronic collisions, being around 5–6%. The relative error would then be reduced by a factor of 5 with respect to this first measurement. With an integrated luminosity of 10 nb^{-1} , the photo-produced $\psi(2S)$ could be measured for the first time in heavy ion collisions: about 400 $\psi(2S)$ events are expected to be detected with a S/B ratio around 1, that is a statistical error of about 8%.

Recently, ALICE has measured an excess of the J/ψ yield at low p_T , below 300 MeV/c, in nuclear Pb-Pb collisions at LHC energies. The physical origin of this excess, shown in Figure 2.40 for two different centrality bins, is yet to be understood and further studied with more precise data. As it was suggested by Nystrand [91], J/ψ photo-production could explain the observed excess. Indeed the J/ψ photo-production cross section could be comparable to the hadro-production one at LHC energies. In this respect, the observed excess would result from the coherent photo-production of J/ψ during the nuclear collision. Such an observation will open fundamental questions on how the coherent photo-production could take place since the system will further interact due to the hadronic interaction at similar time scales. Moreover there would be the question about how does the photo-produced J/ψ interact with the QGP. In this respect, it will be very interesting to study with better precision the low- p_T excess as a function of centrality and, in particular, if this excess is also present in the most central collisions. In Figure 2.41 we show the relative statistical error of the expected measurement of the the low- p_T excess at LHC energies as a function of centrality. In this estimate, we have assumed that the low- p_T excess is centrality-independent. The J/ψ low- p_T excess can be measured in the most central collisions for an integrated luminosity of 10 nb^{-1} .




Figure 2.40: Low- p_T excess observed in the Pb-Pb at $\sqrt{s_{NN}} = 2.76$ TeV in the centrality intervals 70–90% (left) and 50–70% (right). The solid line is based on scaling from pp collisions.




Figure 2.41: Centrality dependence of the relative statistical error of the low p_T J/ψ yield excess measured in Pb-Pb collisions at LHC energies.

2.2.7 $\psi(2S)$ Measurement

Due to the lower production cross section and the smaller branching ratio into dileptons, the measurement of $\psi(2S)$ is much more difficult compared to J/ψ . The $\psi(2S)$ production is measured in nucleus-nucleus collisions only at the SPS [60] and is described by the statistical hadronization model, see Section 2.2.1. Two scenarios are considered for our estimates of the measurements in ALICE: production yield as predicted by the statistical model and as in pp collisions, scaled to Pb-Pb with the number of binary collisions (N_{coll}).

In Figure 2.42 the estimated statistical error of the $\psi(2S)$ measurement in the Muon Spectrometer is shown as a function of centrality for an integrated luminosity of 1 nb^{-1} and 10 nb^{-1} . The full upgrade potential allows for a precision measurement even for the relatively low production expected in case of the thermal model scenario.




Figure 2.42: The estimated relative statistical error of the $\psi(2S)$ measurement in the Muon Spectrometer as a function of centrality for an integrated luminosity of 1 nb^{-1} and 10 nb^{-1} . Two scenarios are considered: the statistical model prediction (left panel) pp scaling (right panel).

In the dielectron channel the measurement is more challenging and can be achieved with good significance only with the 10 nb^{-1} Pb-Pb data expected with the full upgrade, see Figure 2.43. Such a measurement will allow, as for the case of the dimuon channel, to disentangle between a statistical production at the phase boundary and production during the QGP lifetime.




Figure 2.43: Centrality dependence of the expected statistical error of the $\psi(2S)$ measurement in the Central Barrel with 10 nb^{-1} Pb-Pb data in the hypothesis of pp scaling and of the statistical model production (left panel). In the right panel these errors are placed in the context of model expectations for the ratio of nuclear modification factors of ψ' and J/ψ .

2.3 Low-Mass Dileptons

In this section we discuss the feasibility and scientific potential of a detailed study of low-mass dilepton production in e^+e^- with ALICE. The upgrades of the ITS and the TPC as main tracking and PID devices in the ALICE central barrel enable a high statistics measurement of e^+e^- -pairs at mid-rapidity in Pb-Pb collisions. Specifically, the measurements benefit from the low material budget and the enhanced low- p_T tracking capabilities of the new ITS to extend the measurement into the relevant low pair-mass (M_{ee}) and low pair transverse momentum ($p_{T,ee}$) region, and to suppress combinatorial backgrounds from photon conversions and π^0 Dalitz decays. Moreover, the excellent secondary vertex resolution of the ITS will allow to separate contributions from semi-leptonic charm decays, which constitute a major physics background to the prompt radiation.

2.3.1 Scientific Objectives

The measurement of low-mass dilepton production provides access to the bulk properties and the space-time evolution of the hot and dense QCD matter formed in ultra-relativistic heavy-ion collisions, and reveals microscopic properties such as the relevant degrees of freedom and the hadronic excitation spectrum in medium (see [92] for a recent review). Electromagnetic (EM) radiation is produced at all stages of the collision, and since leptons couple only weakly to the surrounding medium, their spectrum retains information of the entire system evolution. This is contrary to the measurement of final state hadrons which interact with the system until freeze-out, and therefore carry only indirect information on the early stages of the collision.

The fundamental questions to be addressed by a comprehensive measurement of low-mass dileptons in heavy-ion collisions at the LHC are the following:

- The generation of hadron masses is driven by the spontaneous breaking of chiral symmetry of QCD in the vacuum. Lattice QCD predicts that this fundamental symmetry is restored at finite T , leading to substantial modifications of the vector and axialvector spectral functions. Such modifications, in particular of the ρ -meson, can be inferred from low-mass dilepton spectra, which constitute the only experimental observable proposed so far to characterize the chiral properties of hot QCD.
- Dilepton production is intimately related to the temperature of the system at all stages of the collision. The temperature dependence is encoded in the transverse momentum and invariant mass spectra. Unlike transverse momentum, invariant mass is not subject to blue-shift in collectively expanding systems and therefore most directly related to temperature. The measurement of low-mass dileptons can also be related to real direct photon production, providing a complementary approach [93] which is less prone to physical backgrounds than real photon measurements from conversions and calorimetry.




Figure 2.44: Temperature distribution of the thermal dilepton yield in different invariant mass intervals (from [94]).

- The space-time evolution of the system, in particular its lifetime can be extracted from low-mass dilepton measurements. The potential to disentangle early from late contributions gives detailed access to the evolution of collectivity and the fundamental properties related to it, such as transport coefficients, viscosity, and the equation of state.

The electromagnetic radiation spectrum emitted in the course of a heavy-ion collision can be expressed by the space-time integral over the differential thermal production rate. The production rate depends on a thermal Bose-Einstein weight f^{BE} and a non-trivial electromagnetic spectral function $\text{Im}\Pi_{EM}$:

$$\frac{dN_{ee}}{d^4x d^4q} = -\frac{\alpha^2}{\pi^3 M_{ee}^2} f^{BE}(q_0, T) \text{Im}\Pi_{EM}(M_{ee}, q, \mu_B, T). \quad (2.4)$$

At large masses ($M_{ee} > 1 \text{ GeV}/c^2$) the spectrum is most sensitive to the initial temperature because the thermal weight leads to a dominance of radiation from the hot early stage. At masses $M_{ee} < 1 \text{ GeV}/c^2$ the emission spectrum is dominated by radiation from hot and dense hadronic matter at or below T_c , which arises as a consequence of the larger volume at late times contributing to the radiation yield and compensating for the smaller thermal weight. This interplay between decreasing thermal weights and increasing space-time volume as the system evolves allows for a detailed investigation of various stages of the collision by proper inspection of different invariant mass windows [94] (see Figure 2.44).

2.3.1.1 Study of Chiral Symmetry Restoration

The electromagnetic spectral function encodes the relevant microscopic degrees of freedom of the system and their dependence on temperature and density. At low masses ($M_{ee} < 1 \text{ GeV}/c^2$) the spectral function is saturated by the light vector resonances, in particular the ρ -meson. Due to its strong coupling to the hot and dense hadronic medium the ρ spectral function features strong broadening and becomes essentially structureless in the vicinity of the phase transition. This ‘melting’ behavior of the ρ spectral function, derived from many-body theory [95], has been observed in In-In at the CERN-SPS, where the low-mass dimuon excess spectrum measured by NA60 [96–98] could be well described assuming this scenario [99] (see Figure 2.45).

The unique potential of low-mass dilepton data to assess in-medium properties of hadronic spectral functions provides means to study chiral symmetry restoration in heavy-ion collisions. Lattice QCD calculations suggest that most of the reduction of the chiral condensate occurs in a temperature regime where the order parameter of deconfinement is still small, giving rise to the expectation that chiral symmetry restoration is manifest in a hot hadronic resonance gas. Recent calculations show that about 60% of the thermal radiation in the low-mass region




Figure 2.45: Low-mass dimuon excess spectrum in In-In measured by NA60, compared to a calculation assuming strong ρ -meson broadening in a hot and dense hadronic medium (from [94]).




Figure 2.46: Lattice QCD calculations of the order parameters of the deconfinement (left) and the chiral (right) phase transition as a function of temperature [2].

($0.15 < M_{ee} < 0.6 \text{ GeV}/c^2$) is emitted in the temperature range $T = 125 - 180 \text{ MeV}$ [94], which corresponds to the chiral transition region indicated by lattice QCD [2], see Figure 2.46. The outstanding feature of chiral symmetry restoration is the degeneracy of the vector and axialvector spectral functions. While the axialvector spectral function is not accessible experimentally, the modification of the vector spectral function is instrumental to provide a connection to chiral symmetry restoration via Weinberg and QCD sum rules [100–102] and constraints from lattice QCD. A precise measurement of the low mass dilepton spectrum in heavy-ion collisions constitutes the only known means to assess the nature of the chiral phase transition experimentally and is therefore of paramount importance for the understanding of excited QCD matter. Conceptually, a measurement at the LHC is most favorable because a rigorous theoretical evaluation is substantially enhanced by the applicability of lattice QCD at $\mu_B = 0$.

2.3.1.2 Early Temperature

At masses $M_{ee} > 1 \text{ GeV}/c^2$ the description of the spectral function becomes 'dual' in the sense that hadronic and partonic degrees of freedom lead to the same structureless spectral function. Moreover, the thermal Bose-Einstein weight suppresses contributions from late stages, such that the invariant mass dependence of the thermal dilepton yield is very sensitive to the early temperature of the system (see also Figure 2.44). However, a precise




Figure 2.47: Effective slope parameter T_{eff} from exponential fits to the dimuon p_{T} excess spectra measured by NA60 [96].

measurement of the thermal yield in this mass window ($M_{\text{ee}} > 1 \text{ GeV}/c^2$) requires a good understanding of the contribution from correlated semi-leptonic open charm decays (see below).

The spectrum of real direct photons ($M = 0$) has long been considered a most sensitive probe for the source temperature. The measurement of dileptons at very low invariant masses ($0.15 < M_{\text{ee}} < 0.3 \text{ GeV}/c^2$) and finite transverse momentum ($p_{\text{T},\text{ee}} > 1 \text{ GeV}/c$) provides access to the real direct photon spectrum, by extrapolation of the virtual photon yield to $M = 0$ [93]. This method to infer the direct photon yield is complementary to calorimetry of real photons and conversion measurements. Although the virtual direct photon yield is suppressed by one order in α , the virtual photon measurement is very powerful because it suffers much less from physical background than the calorimetric measurement. This allows to measure the direct photon yield in the low transverse momentum region ($1 < p_{\text{T},\gamma} < 5 \text{ GeV}/c$) where a large thermal contribution is expected.

2.3.1.3 Space-Time Evolution and Equation of State

The equation of state, i.e. the pressure as a function of the temperature is one of the most fundamental characteristics of strongly interacting matter. It governs the expansion properties of the early universe and has obvious phenomenological consequences for the space-time evolution of heavy-ion collisions. The equation of state is a basic ingredient to hydrodynamical models which yield a good description of the observed phenomenon of collective flow. Conceptually, collective flow patterns of final-state hadrons arise from integration over the pressure fields acting at all stages of the collision, and differentiation of partonic and hadronic contributions is typically ambiguous. Electromagnetic probes like dilepton pairs are emitted at all stages of the collision, giving rise to a complex collectivity pattern if studied as a function of invariant mass [103]. Systematic analysis of radial and elliptic flow in different invariant mass windows gives access to the evolution of collectivity at different times during the collision, and may therefore provide stringent constraints on the equation of state [99].

This connection has been demonstrated on the basis of NA60 results from In-In collisions at the SPS [96, 97]. Figure 2.47 shows the 'effective temperature' coefficient T_{eff} extracted from the transverse momentum spectra of dimuon pairs. A clear increase is observed below the ρ -meson mass which is consistent with previous observations of radial flow patterns in the p_{T} spectra of light hadrons. The appearance of radial flow is commonly attributed to late collision stages, which govern the emission spectrum at low masses. At higher masses, T_{eff} decreases which may be attributed to increasing contributions from the partonic stage, where radial flow is less developed. In heavy-ion collisions at the LHC significantly higher initial temperatures are reached, giving rise to even more pronounced contributions from the QGP to dilepton production at high masses. A detailed investigation of radial and elliptic flow as a function of invariant mass will thus provide a unique experimental handle on the equation of state of partonic matter.

2.3.2 Experimental Challenges

The measurement of low-mass e^+e^- -pairs in Pb-Pb collisions at LHC energies poses major experimental challenges. A powerful assessment of the electromagnetic radiation spectrum requires acceptance for dilepton pairs at invariant masses and transverse momenta as low as $M_{ee} \approx p_{T,ee} \approx T$. This implies electron detection down to $p_T = 0.1 - 0.2 \text{ GeV}/c$. Since the production rates of thermal dileptons are low (suppressed by $\propto \alpha^2$) a very good electron identification is mandatory to suppress combinatorial background from hadronic contaminations. Moreover, electrons from π^0 Dalitz decays and photon conversions (mainly from $\pi^0 \rightarrow \gamma\gamma$) form a substantial combinatorial background. This demands for a low material budget before the first active detector layer and offline strategies to detect e^+e^- -pairs from photon conversions and Dalitz pairs for further rejection. The large combinatorial background prevents also a straight-forward online trigger scheme. Consequently, high data rates need to be handled by the detectors and the readout systems.

These experimental requirements will be fulfilled after the luminosity increase of the LHC, providing Pb-Pb collisions at 50 kHz, and a substantial upgrade of the ALICE detector and readout systems. Specifically, running the TPC in continuous mode will allow to record about $2.5 \cdot 10^9$ central (0–10%) Pb-Pb collisions, corresponding to about $\int Ldt = 3 \text{ nb}^{-1}$, within a typical running period dedicated for the low-mass dilepton measurement, where the ALICE solenoid will be operated at a reduced magnetic field of $B = 0.2 \text{ T}$ (see below). This is to be compared to $2.5 \cdot 10^7$ central Pb-Pb collisions without TPC upgrade. Moreover, the enhanced low- p_T tracking capability of the ITS will allow to track electrons down to $p_T \geq 0.05 \text{ GeV}/c$. This will substantially improve the reconstruction efficiency of photon conversions and Dalitz pairs for combinatorial background suppression.

The improved secondary vertex resolution of the new ITS will also enable efficient tagging of electrons from semi-leptonic charm decays, thus allowing to separate the displaced charm contribution from prompt dileptons. This is of particular importance since a measurement of the thermal dilepton yield requires a precise knowledge of the hadronic contributions. These include dileptons from correlated charm decays, which form a substantial physical background to the pair yield and dominate the mass spectrum at $M_{ee} > 1.1 \text{ GeV}/c^2$. Suppression of electrons from charm via their impact parameter distribution to the primary vertex will be significantly improved with the new ITS and thus reduce the systematic uncertainties on the thermal excess yield related to the subtraction of the charm contribution from the inclusive dilepton yield.

The main tracking and PID detector for electrons is the TPC, providing particle identification via the measurement of the specific energy loss by ionization dE/dx . Additional PID information is required from the TOF systems to suppress hadrons at $p < 3 \text{ GeV}/c$. Operation at the nominal magnetic field of $B = 0.5 \text{ T}$ implies that soft particles do not reach the TOF, leading to a drastic efficiency loss for $p_T < 0.4 \text{ GeV}/c$. Therefore a dedicated running period is foreseen at a reduced magnetic field of $B = 0.2 \text{ T}$ for the measurement of low-mass dileptons. This will improve the acceptance for low- p_T electrons in TOF.

The dE/dx measurement as function of p for all charged particles in the TPC is shown in the upper panels of Figure 2.48 at $B = 0.5 \text{ T}$ (left) and $B = 0.2 \text{ T}$ (right). The same distributions after selection of a typical 3σ band on electrons in TOF is shown in the lower panels. Inclusion of the TOF system into the electron identification scheme leads to a significant purification of the electron sample. The reduction of the magnetic field to $B = 0.2 \text{ T}$ improves considerably the acceptance for soft electrons.

The improvement of low- p_T electron reconstruction allows to explore the most interesting kinematic range at low M_{ee} and $p_{T,ee}$. This is illustrated in Figure 2.49, where the dilepton acceptance as a function of M_{ee} and $p_{T,ee}$ is shown for a tracking and PID scheme using ITS, TPC and TOF at $B = 0.5 \text{ T}$ (left) and $B = 0.2 \text{ T}$ (right). While the low- p_T cut-off at $B = 0.5 \text{ T}$ causes a dramatic loss of acceptance at low M_{ee} and $p_{T,ee}$, the acceptance is dramatically improved at $B = 0.2 \text{ T}$. We therefore foresee a dedicated low field run with $\int Ldt = 3 \text{ nb}^{-1}$ at $B = 0.2 \text{ T}$ for the measurement of low-mass dileptons.

2.3.3 Physics Performance

To assess the anticipated significance of a low-mass dilepton measurement with the upgraded ALICE detector in Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5 \text{ TeV}$, a detailed simulation has been performed. In the following, we consider three experimental scenarios:

Scenario 1: Current ITS, low rate

The present ALICE setup including the current ITS and the existing TPC, operated at a magnetic field of




Figure 2.48: Upper row: TPC dE/dx as a function of momentum at $B = 0.5$ T (left) and $B = 0.2$ T (right), plotted in units of the dE/dx resolution after subtraction of the Bethe-Bloch curve for electrons. Lower row: same after application of a 3σ selection of electrons in TOF.

$B = 0.2$ T is assumed. In this scenario, we expect in a typical one-month heavy-ion run $2.5 \cdot 10^7$ and $5 \cdot 10^7$ Pb–Pb collisions at 0–10% and 40–60% centrality, respectively.

Scenario 2: New ITS, low rate

The ALICE setup including a new ITS system with 7 pixel layers, but the existing TPC, operated at $B = 0.2$ T is assumed. Due to the rate limitations of the existing TPC the anticipated event sample in one year of heavy-ion running is the same as before, i.e. $2.5 \cdot 10^7$ and $5 \cdot 10^7$ Pb–Pb collisions at 0–10% and 40–60%, respectively.

Scenario 3: New ITS, high rate

The ALICE setup including the new ITS with 7 pixel layers and a new TPC with continuous readout, operated at $B = 0.2$ T is assumed. The upgrade of the TPC increases the readout rate by about a factor 100, allowing to collect all of the 50 kHz delivered by the LHC after luminosity upgrade. The number of expected events is then $2.5 \cdot 10^9$ Pb–Pb collisions at 0–10%, and $5 \cdot 10^9$ Pb–Pb collisions at 40–60%.

2.3.3.1 Detector Performance

The tracking efficiency for single electrons in $|\eta| < 0.84$ at $B = 0.2$ T is calculated using a GEANT3 implementation of the current ALICE detector setup. Electron reconstruction includes charged particle tracking in the ITS




Figure 2.49: Acceptance for $e^+ e^-$ -pairs from PYTHIA at $B = 0.5$ T (left) and $B = 0.2$ T (right).




Figure 2.50: . Left: Combined ITS-TPC-TOF efficiencies for electrons in $|\eta| < 0.84$ at $B = 0.2$ T for the current ITS (blue circles) and the new ITS (green rectangles) as a function of p_T . Also shown for comparison is the efficiency with current ITS at $B = 0.5$ T (open circles). Right: Standalone tracking efficiency of the current and the new ITS system.

and the TPC. Additionally, a TOF signal within $\pm 3\sigma$ of the TOF resolution around the nominal electron position is required. PID information is also included from the TPC, where a $\pm 3\sigma$ band of the dE/dx resolution around the nominal electron peak position is applied (see also Figure 2.48). Electron candidates are required to have an associated hit in the inner-most layer of the ITS to suppress electrons from conversions in the detector material. The resulting reconstruction efficiency for single electrons is parametrized as a function of p_T and used for further fast simulations employing different event generators, see left panel of Figure 2.50. Also shown for comparison is the electron efficiency with the current ITS at $B = 0.5$ T.

For scenarios including the new ITS system, the electron reconstruction efficiency as determined for the current ITS has been scaled by the p_T -dependent ratio of the ITS standalone efficiencies for the new ITS (7 pixel layers) over that for the current ITS. The standalone efficiencies of the current and the new ITS are shown in the right panel of Figure 2.50. It is furthermore assumed that the performance of the upgraded TPC in terms of tracking efficiency and dE/dx resolution is the same as for the current TPC. The resulting efficiency for electron tracks with new ITS at $B = 0.2$ T is also shown in Figure 2.50 (left panel).

An important feature of the new ITS system is its improved capability to separate prompt from displaced electrons. Prompt electrons emerge from all thermal sources and the decays of the light mesons (ρ , ω , ϕ , and the Dalitz




Figure 2.51: Efficiency for electrons from charm (magenta) and conversions (blue) as a function of the efficiency for prompt (primary) electrons, for current ITS (open symbols) and new ITS (full symbols).

decays of π^0 , η , η' , ω). Displaced electrons include those from semi-leptonic decays of hadrons with charm and conversions in the detector material. The separation is based on the distance-of-closest-approach (DCA) to the main interaction vertex, which can be measured with significantly improved resolution with the new ITS. Figure 2.51 shows the efficiency of displaced electrons from conversions and charm decays for the current and the new ITS system, as a function of the prompt efficiency. While the separation of displaced electrons from charm is significantly improved with the new ITS, there is only a small improvement for conversions. The reason is that the new ITS system requires also a new beam pipe with smaller radius, which is the main converter before the first detector layer. The smaller lever arm to the main vertex compensates the improved intrinsic resolution of the new ITS. On the other hand, a tight cut on the DCA suppresses electrons from charm over prompt electrons by more than a factor of two, implying a suppression of pairs from correlated charm decays by about a factor five over prompt pairs. This improved capability of the new ITS is one of the key features of the ALICE upgrade which render possible a measurement of thermal radiation at the LHC despite the large physical background from charm, as will be demonstrated below.

2.3.3.2 Signal Generation

A realistic physics input into the simulation of the expected dilepton signal is mandatory in order to perform a solid investigation of the physics performance. To this end, we compute the dilepton signal composed of the contributions listed below. The signal is calculated for central (0–10%) and semi-central (40–60%) Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV.

Hadronic cocktail

The hadronic cocktail includes contributions from the decays of light pseudoscalar and vector mesons. The yield is adjusted to charged particle densities $\langle dN_{\text{ch}}/d\eta \rangle = 1750$ and $\langle dN_{\text{ch}}/d\eta \rangle = 248$ in 0–10% and 40–60% most central Pb–Pb at $\sqrt{s_{\text{NN}}} = 5.5$ TeV, respectively. The particle ratios and the spectral shapes are extrapolated from existing heavy-ion data at lower energies.

Charm

The contribution from correlated semi-leptonic charm decays is based on calculations from the PYTHIA event generator. A total charm cross section of $\sigma_{\text{cc}} = 7.55$ mb in pp at $\sqrt{s} = 5.5$ TeV has been used which is derived by interpolation of existing measurements [41, 104–106]. Around mid-rapidity, a differential cross section $d\sigma_{\text{cc}}/dy = 1.34$ mb is used, also taken from interpolation. The charm yields are scaled by $\langle N_{\text{coll}} \rangle = 1625$ (125) for Pb–Pb collisions at 0–10% (40–60%) centrality.




Figure 2.52: The dilepton signal in the 40–60% (left) and 0–10% (right) most central Pb–Pb collisions at $\sqrt{s_{\text{NN}}}=5.5$ TeV. Shown are contributions from hadronic decays, charm, and thermal radiation from the hadronic and the QGP phase (see text).

Thermal radiation

The calculation of thermal radiation is based on a hadronic many-body approach [95] and perturbative emission rates to model thermal dilepton radiation from the hadronic phase and the QGP, including medium-modified spectral functions and a realistic space-time evolution [99]. This approach has proven to provide a quantitative description of dilepton data over a wide range of collision energies [94]. The calculations are performed for the charged particle densities anticipated in 0–10% and 40–60% most central Pb–Pb collisions at $\sqrt{s_{\text{NN}}}=5.5$ TeV, i.e. for $\langle dN_{\text{ch}}/d\eta \rangle = 1750$ and 248, respectively².

The resulting dilepton signal distributions in central and semi-central collisions are shown in Figure 2.52. With the exception of the π^0 , ω and ϕ mass regions, the yield is dominated by the contribution from correlated charm decays. Thermal radiation from the hadronic phase dominates over that from QGP in the low-mass window, giving access to chiral symmetry restoration. In turn, QGP radiation outshines the hadronic contribution at $M_{\text{ee}} > 1 \text{ GeV}/c^2$, where information on the early temperature can be extracted. However, in all mass regions an extraction of thermal radiation requires careful subtraction of the contributions from hadronic decays and, in particular, charm.

2.3.3.3 Combinatorial Background

The measured raw dilepton yield is dominated by combinatorial pairs of electrons and positrons, which arise from random combinations of tracks from uncorrelated decays, mainly π^0 –Dalitz, and from conversions. This combinatorial background contribution can be estimated by like-sign pair combinations or pairs from mixed events, and subtracted from the unlike-sign distribution. However, a small signal-to-background ratio (S/B) limits the statistical significance of the signal and imposes a systematic uncertainty. To minimize the contribution of combinatorial pairs, leptons from conversions and Dalitz decays must be identified and removed from the sample. To this end, all electrons and positrons from an event are combined to pairs and rejected if they form a 'close pair' with small invariant mass ($M_{\text{ee}} < 0.05 \text{ GeV}/c^2$) and small opening angle ($\Theta_{\text{ee}} < 0.1$). Since leptons from conversions or π^0 –Dalitz decays typically have small momentum, the capability of Dalitz and conversion rejection is much enhanced by the improved ITS efficiency for standalone tracks with low p_T . Even if not used in the signal sample, ITS standalone tracks are instrumental for rejection if they form a close pair with a signal track.

To estimate the combinatorial background, S/B , and the corresponding significance $1/\sqrt{N_{\text{ev}}} \cdot S/\sqrt{S+B}$ we studied PYTHIA pp events. The events are processed by the GEANT3 model of the ALICE detector to provide a realistic contribution from conversions. Parametrized efficiencies (see Figure 2.50) for electron and positron tracks are applied. Tracks are rejected if they form a close pair with any of the other reconstructed lepton tracks, including

²R. Rapp, private communication




Figure 2.53: Signal and background distributions (left), S/B (middle) and significance per event (right) in 0–10% (upper row) and 40–60% (lower row) Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV.

standalone tracks which are only found in the ITS. In order to mimic the combinatorial background in Pb–Pb, a number of pp events are overlaid to match the corresponding $\langle dN_{\text{ch}}/d\eta \rangle$ in semi-central or central Pb–Pb collisions.

Examples for the expected combinatorial background in central and semi-central Pb–Pb collisions for the current and new ITS systems are shown in Figure 2.53. The new ITS improves S/B by about a factor 2 with respect to the current ITS system.

2.3.3.4 Systematic Errors

The dominant sources of systematic uncertainties on the dilepton measurement arise from the large combinatorial and physical backgrounds.

The relative systematic uncertainty on the combinatorial background $\Delta B/B$ propagates into the extracted inclusive dilepton signal error as $\Delta S/S = \Delta B/B \cdot B/S$. For the studies below we assume conservatively $\Delta B/B = 0.0025$ which corresponds to the number quoted by the PHENIX collaboration in a similar analysis [107]. The systematic error on S is calculated bin-by-bin based on the results for S/B as shown in Figure 2.53.

Measurement of the thermal excess yield implies a precise subtraction of the hadronic cocktail and the contribution from charm from the inclusive dilepton yield. For the hadronic cocktail we assume a relative uncertainty of 10%. We note that ALICE has the unique capability to measure π^0 and η down to low p_T via conversions [108] and can therefore constrain the uncertainty on the hadronic cocktail by data from the same experiment.

For the subtraction of the charm contribution a relative systematic uncertainty of 20% is assumed. Also here we note that the yield of correlated dilepton pairs from charm decays can be well constrained by inversion of the DCA cuts, which enhances the charm contribution, and by the exclusive measurement of charmed hadrons at low p_T with complementary methods in ALICE.

2.3.3.5 Results

In the following, the results of the physics performance study described before are discussed. Figure 2.54 (left) shows the inclusive e^+e^- invariant mass spectrum in the 0–10% most central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV




Figure 2.54: Inclusive e^+e^- invariant mass spectrum (left) and excess spectrum (right) for 0–10% most central Pb–Pb collisions at $\sqrt{s_{NN}}=5.5$ TeV in Scenario 1 (current ITS, $2.5 \cdot 10^7$ events). No tight DCA cuts are applied. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution.




Figure 2.55: Inclusive e^+e^- invariant mass spectrum (left) and excess spectrum (right) for 0–10% most central Pb–Pb collisions at $\sqrt{s_{NN}}=5.5$ TeV in Scenario 1 (current ITS, $2.5 \cdot 10^7$ events). Tight DCA cuts are applied. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution.

in Scenario 1, i.e. current ITS and $2.5 \cdot 10^7$ events. No particular DCA cuts are applied to reject displaced electrons. The same spectrum after subtraction of the hadronic cocktail and the charm contribution (the ‘excess spectrum’) is shown in the right panel of Figure 2.54. The low-mass region $M_{ee} < 1 \text{ GeV}/c^2$ is dominated by systematic uncertainties related to the subtraction of the combinatorial background. In the mass region $M_{ee} > 1 \text{ GeV}/c^2$, the systematic uncertainties from the charm subtraction do not allow quantitative analysis of the thermal radiation spectrum.

The DCA resolution of the current ITS allows for some limited suppression of displaced electrons (see also Figure 2.51). In the left panel of Figure 2.55, the inclusive e^+e^- in Scenario 1 is shown after application of tight DCA cuts. The relative contribution from charm can be suppressed by about a factor 2 (compare to Figure 2.54,




Figure 2.56: Inclusive e^+e^- invariant mass spectrum (left) and excess spectrum (right) for 0–10% most central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV in Scenario 2 (new ITS, $2.5 \cdot 10^7$ events). Tight DCA cuts are applied. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution .

left), at the expense of an additional loss in statistics. In the right panel of Figure 2.55, the corresponding excess spectrum is shown which indicates improved systematic uncertainties from charm subtraction, but still large errors from combinatorial background and insufficient statistics.

Figure 2.56 shows the inclusive e^+e^- invariant mass spectrum (left panel) and the excess spectrum (right panel) in 0–10% most central Pb–Pb collisions in Scenario 2 (new ITS, $2.5 \cdot 10^7$ events). Tight DCA cuts to reject displaced electrons are applied. The enhanced low- p_T tracking capability of the new ITS leads to significantly improved rejection of combinatorial background, and consequently reduced systematic uncertainties, as compared to the current ITS system (see Figure 2.55). Further reduction of systematic uncertainties related to charm subtraction is also achieved. However, the statistical limitations of the measurement would not allow for a quantitative analysis of the thermal dilepton excess.

A key element of the ALICE upgrade strategy is therefore a concept for a continuously operated TPC, allowing to record the full collision rate of 50 kHz anticipated at the LHC after 2020. This would imply an increase by about a factor 100 in recorded luminosity, leading to about $2.5 \cdot 10^9$ central Pb–Pb events in a typical year of heavy-ion running (Scenario 3). The expected inclusive e^+e^- invariant mass spectrum and the excess spectrum in Scenario 3 are shown in Figure 2.57. The statistical precision achievable in this high-rate scenario would allow for a detailed and differential investigation of dilepton production.

Transverse momentum spectra provide valuable information on the degree of collectivity of the dilepton sources. Moreover, their mass dependence is related to collectivity in different stages of the collision, and may give access to partonic flow, if measured at sufficiently high masses. The combination of thermal and collective motion is encoded in the effective temperature parameter T_{eff} , which can be extracted from an exponential fit to the excess $p_{T,\text{ee}}$ spectra in intervals of invariant mass.

Examples for mass-dependent transverse momentum spectra are shown in Figure 2.58. With the new ITS and an upgraded high-rate TPC (Scenario 3, right panels in Figure 2.58) a detailed measurement of the $p_{T,\text{ee}}$ spectra is possible, allowing an extraction of T_{eff} with a statistical precision on the level of $\sigma(T_{\text{eff}})/T_{\text{eff}} \approx 1\%$.

We note that the dominant source of systematic errors at low M_{ee} and $p_{T,\text{ee}}$ arises from the uncertainty on the combinatorial background $\Delta B/B = 0.0025$ assumed in this study. For future measurements, we aim to decrease the systematic uncertainty to $\Delta B/B = 0.001$, based on the large acceptance and excellent tracking capabilities of the upgraded ALICE detector, and the high statistics data set that is anticipated.

A more detailed investigation of dilepton collectivity and a possible access to the partonic equation of state is expected from a measurement of the elliptic flow coefficient v_2 as a function of M_{ee} . To this end, the physics




Figure 2.57: Inclusive e^+e^- invariant mass spectrum (left) and excess spectrum (right) for 0–10% most central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV in Scenario 3 (new ITS, $2.5 \cdot 10^9$ events). Tight DCA cuts are applied. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution.

performance study is extended to semi-central collisions, where elliptic flow is most pronounced. The e^+e^- invariant mass excess spectra in semi-central (40–60%) Pb–Pb collisions is shown in Figure 2.59 for Scenario 1 (left) and Scenario 3 (right). Note that relative systematic uncertainties in semi-central collisions are smaller than in central collisions, due to larger S/B (Figure 2.53) and a smaller relative contribution from charm due to $\langle N_{\text{coll}} \rangle$ scaling. The absolute statistical uncertainties on v_2 as a function of M_{ee} are shown in Figure 2.60 for Scenario 1 and Scenario 3. After the high-rate upgrade and with the new ITS, invariant-mass dependent v_2 measurements with absolute statistical uncertainty of order $\sigma(v_2) \approx 0.01\text{--}0.02$ can be achieved.

Information on the early temperature of the system can be derived from the invariant-mass dependence of the dilepton yield at masses $M_{ee} > 1 \text{ GeV}/c^2$. To quantify the sensitivity of the anticipated measurement we employ an exponential fit, $dN_{ee}/dM_{ee} \propto \exp(-M_{ee}/T_{\text{fit}})$, to the simulated spectra in the invariant mass region $1 < M_{ee} < 1.5 \text{ GeV}/c^2$. The fit parameter T_{fit} is compared to T_{real} which is derived from the same fit to the thermal input spectrum. The ratio $T_{\text{fit}}/T_{\text{real}}$ for Pb–Pb collisions at 0–10% and 40–60% centrality is shown in Figure 2.61. Only the high-rate scenario with new ITS (Scenario 3) allows a quantitative extraction of the slope parameter in the relevant M_{ee} range, with statistical and systematic uncertainties in the range of 10–20%.

In conclusion, the measurement of e^+e^- production in central and semi-central Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV provides unique experimental access to modifications of the vector spectral function and restoration of chiral symmetry, the early temperature of the system, and the equation of state of partonic matter. With the upgrade of the ITS and TPC detectors such measurements become feasible. The new ITS detector will allow for a significant suppression of combinatorial background, and a separation of prompt from displaced electrons, the latter mainly from correlated charm. We demonstrated that this leads to a significant reduction of the main sources of systematic uncertainties. Moreover, the upgrade of the TPC with GEM readout will allow for continuous operation, making possible to record Pb–Pb collisions at a rate of 50 kHz. This improvement in statistical accuracy will enable a multi-differential analysis of the dilepton excess as a function of M_{ee} , $p_{T,ee}$, and the orientation to the reaction plane. We have shown that observables that parametrize the corresponding dependencies can be extracted with unprecedented precision.

2.4 Jets

The main motivation for measuring jets in heavy-ion collision is to map out the properties of the created medium via its interaction with hard scattered partons. Hard scatterings ($Q^2 \gg (2 \text{ GeV}/c)^2$) occur in the early reaction phase ($\tau \ll 1 \text{ fm}/c$), well before the formation of a hot and dense medium and enable in principle the *tomographic* study of the medium. The basis of this approach is that the initial production of hard scattered partons is well




Figure 2.58: Transverse momentum e^+e^- excess spectra in 0–10% most central Pb–Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV in intervals of invariant mass for Scenario 1 (current ITS, $5 \cdot 10^7$ events, left) and Scenario 3 (new ITS, $5 \cdot 10^9$ events, right). Tight DCA cuts are applied. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution.

defined and also calculable in perturbative QCD, which can be tested in the *vacuum* case of jet measurements in proton–proton. In heavy-ion collisions, the medium modification of hard probes has been first observed at RHIC in single inclusive hadron production and particle correlations, where the particle production in central Au–Au collisions with $\sqrt{s_{NN}} = 200$ GeV at high p_T and the jet-like correlations are significantly suppressed compared to proton–proton (*jet quenching*) [109, 110].

In practice several distinct differences of jet tomography to the familiar medical X-ray imaging exist, which put limitations on the direct, quantitative tomographic interpretation:

- The probed medium itself expands, depending on initial conditions and its hydrodynamic properties.
- The origin of the probe is only known on average.
- The (partonic) probe cannot be observed as a free particle, hence no direct attenuation can be defined.
- Several mechanisms of the parton-medium interaction exist, which in practice also can occur in parallel, e.g. elastic and radiative energy loss.




Figure 2.59: e^+e^- excess spectra in 40–60% centrality Pb–Pb collisions at $\sqrt{s_{NN}} = 5.5$ TeV in Scenario 1 (left panel) and Scenario 3 (right panel). Tight DCA cuts are applied. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution.




Figure 2.60: Expected absolute statistical uncertainty of the elliptic flow coefficient v_2 of the e^+e^- excess spectrum as a function of M_{ee} . Results are shown for Pb–Pb collisions at 40–60% centrality in Scenario 1 (current ITS, $5 \cdot 10^7$ events, left panel) and Scenario 3 (new ITS, $5 \cdot 10^9$ events, right panel). Tight DCA cuts are applied.

Thus, the measurement of hard probes and the modified fragmentation process into observable hadrons provides not only access to the mechanisms of partonic energy loss, it also puts additional and complementary constraints on the hydrodynamic evolution of the system and its initial conditions (see e.g. [111]). Furthermore, the presence of the underlying event in heavy-ion collisions and its structure has a direct impact on jet reconstruction and thereby the measured jet observables such as the differential jet yield, jet shape, longitudinal and transverse fragmentation. These effects need to be carefully separated from the true medium modification of the parton fragmentation.

The advantage of the ALICE detector in this context is that it provides the measurement of jets with a minimal bias, in a sense that it allows jet reconstruction and background characterization on the individual (charged) particle level due to its excellent track separation as well as high and uniform efficiency from high (> 100 GeV/c) down to low momentum (150 MeV/c). From there on biases can be gradually introduced to study the evolution of jet observables under different constraints, such as minimum particle p_T , recoil jets off a certain trigger particle type




Figure 2.61: Expected relative uncertainty on the extraction of the T parameter from a fit to the invariant mass excess spectrum in $1.1 < M_{ee} < 1.5 \text{ GeV}/c^2$ (see text). The results are shown for Scenarios 1, 2, and 3 with tight DCA cuts, and for 0–10% and 40–60% event centrality. Error bars show the statistical uncertainties. The green boxes show the systematic uncertainties from the combinatorial background subtraction, the magenta boxes indicate systematic errors related to the subtraction of the cocktail and charm contribution.

or topology etc.

The proposed upgrade will enable additional, unique contributions of ALICE to the differential study of medium modification of jet probes via three major techniques: direct reconstruction of jets and jet structure observables, (identified) particle – jet correlations, (identified) particle – particle correlations. It will extend the minimum-bias study of jet observables since no hardware selection of jet event is needed. Hence, the focus can be on the detailed characterization of the in-medium fragmentation of light quark, heavy quark and gluon jets and the exploration of modified jet properties in a wide parton momentum range down to $p_T^{\text{jet}} \approx 10 \text{ GeV}/c$, where the interplay between soft and hard processes is largest, but at the same time the separation of the underlying event impact on the measured jet observables is most challenging.

2.4.1 Jet Measurements in ALICE

Jet Reconstruction

Jets are reconstructed combining information from charged and neutral particle measurements. By explicitly measuring essentially all (charged) jet constituents with large efficiency down to very low p_T ($p_T^{\text{tracks}} > 150 \text{ MeV}/c$, $E_T^{\text{clus}} > 150 \text{ MeV}$) ALICE jet measurements provide important complementary information to calorimeter based jet measurements and provide a systematically different approach in separating the impact of the underlying event from the medium modification of jet observables.

Charged particle momentum vectors are measured with the central tracking detectors, the Time Projection Chamber (TPC) and the Inner Tracking System (ITS) covering the full azimuth and $|\eta| < 0.9$. Additional space-point information and PID are provided by TOF and TRD. Energy and direction of neutral particles are currently mea-

sured with the Pb-scintillator sampling ElectroMagnetic Calorimeter (EMCal), covering 1/3 of the azimuth and $|\eta| < 0.7$, it will be complemented in the near future by the DCAL covering 60° opposite in φ to the EMCal [112]. For jet reconstruction the anti- k_T algorithm from the FastJet package [113] with resolution parameters R varying between 0.2 and 0.4 is used. The jet 4-momentum vector is calculated using the boost invariant p_T recombination scheme. Analyses with charged jets using only tracking information and fully reconstructed jets including EMCal information have been performed.

Jet-by-jet we correct for the energy contribution from charged particles to the energy measured with EMCal and the contribution from the underlying event. The sum of momenta of charged tracks matching the EMCal clusters from the cluster energy is subtracted resetting negative values to zero. In Pb–Pb collisions, one has also to subtract the contribution of the Underlying Event (UE) from the reconstructed jet p_T . The summed p_T from the background is calculated as the product of mean momentum density ρ and the jet area A^{jet} , where ρ is determined using the k_T -algorithm [114] via $\rho = \text{median}\left(\frac{p_T^{\text{jet}}}{A^{\text{jet}}}\right)$.

Further corrections can only be applied on the raw spectrum bin-by-bin or via unfolding techniques. These corrections comprise the unmeasured p_T from neutrons and K_L^0 , as well as the tracking efficiency and the corresponding jet-by-jet fluctuations of these quantities. In Pb–Pb collisions, one also has to correct for the smearing of the spectra induced by the UE energy fluctuations. This smearing is quantified by δp_T , the difference between the UE corrected summed p_T and the true jet p_T : $\delta p_T = (p_T^{\text{rec}} - \rho A^{\text{jet}}) - p_T^{\text{true}}$.

A data driven method to determine the distribution of δp_T consists in embedding different objects into measured Pb–Pb collisions [115]. These objects can be single high- p_T tracks, jets or random cones. The distribution is almost Gaussian with enhanced tails towards positive differences, mainly due to pile-up of jets in the same jet area. For a resolution parameter $R = 0.2$ the width (σ) of the Gaussian amounts to 6.2 GeV/c (4.5 GeV/c) summing neutral and charged p_T (for tracks only).

Jet Spectra: pp Collisions

Figure 2.62 (left) shows the jet p_T spectrum measured in pp collisions at $\sqrt{s} = 2.76$ TeV with $R = 0.4$. The jet energy scale uncertainty amounts to 4% and it is mainly due to uncertainties on the missing neutral energy, the tracking efficiency and energy double counting. The jet p_T resolution $\Delta p_T/p_T$ amounts to 20% and is dominated by the jet energy scale fluctuations, tracking and EMCal resolutions. The measured spectrum is compared to NLO pQCD calculations and PYTHIA8 and we find a good agreement. The measurement itself represents an import reference for our Pb–Pb measurements.

Ratios of differential cross-sections measured with different R inform about the energy distribution within a jet (jet shape). Note that for jets reconstructed with the anti- k_T algorithm, R is to a good approximation the radius of a circular jet area. Figure 2.62 (right) shows the measured ratio for $R = 0.2$ and 0.4. The ratio rises with p_T^{jet} illustrating the well known fact that higher p_T jets are more collimated.

Jet Suppression in Pb–Pb

First measurements of jet spectra in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV have been performed using charged particle information only. In order to study the modifications of the Pb–Pb spectra with respect to an incoherent superposition of colliding nucleons the nuclear modification factor $R_{\text{AA}}^{\text{jet}}$ is studied. As a reference we use the spectra from pp collisions at the same centre of mass energy simulated by the PYTHIA MC [46, 116]. The results for the highest centrality bin (0–10%) and the lowest one (50–80%) for $R = 0.2$ are compared in Figure 2.63 (left). A strong jet suppression qualitatively and quantitatively similar to the R_{AA} of inclusive hadrons is observed in the most central collisions.

A possible effect of jet quenching is the redistribution of the radiated energy within the jet cone leading to jet shape modifications. As mentioned earlier the ratios of differential cross-sections measured with different R can inform about these jet shape modifications. Figure 2.63 (right) shows the measured ratio for $R = 0.2$ and 0.3. Despite the low momentum cut-off for the jet constituents of only 150 MeV/c, no modifications of the jets measured in central Pb–Pb collisions with respect to more peripheral collisions or the PYTHIA pp reference are observed within the still large experimental uncertainties.




Figure 2.62: Left: ALICE measurement of the differential cross section of fully reconstructed jets in pp collisions at $\sqrt{s} = 2.76$ TeV compared to PYTHIA8 and NLO pQCD. Right: Ratio of jet cross sections reconstructed with different R in pp collisions at $\sqrt{s} = 2.76$ TeV compared to PYTHIA8 and NLO pQCD.

Hadron–Jet Correlations

In case of leading particles the surface bias is strong in central Pb–Pb collisions and the average in-medium path length is small, conversely, the recoiling parton is biased towards higher in-medium path length. To make use of this effect we study the conditional jet yield requiring a trigger hadron back-to-back with respect to the jet axis. There are additional advantages of this method. No bias on the fragmentation of the recoiling jet is present. Moreover, the requirement of a correlated high- p_T hadron tags hard scatterings and suppresses the combinatorial, fake jet background at low jet p_T .

2.4.2 Jet Structure Modifications

ATLAS and CMS have performed the first analysis of the structure of high energy jets in central Pb–Pb collisions with data recorded in 2010. It has been shown that within experimental uncertainties the remnant jet after quenching has an unmodified structure and the radiated energy results in low- p_T particles found far from the jet axis [117–119]. Using the 10 times larger data sample recorded in 2011 it was possible to refine these results measuring the jet shape and fragmentation function of jets with $E_T > 100$ GeV using particles with p_T down to 1 GeV/c [120,121]. Modifications of the differential energy distribution (dE/dR) within the jet cone up to 40% with respect to the pp reference at large radius ($R = 0.3$) and a modification of the fragmentation function for ($R < 0.3$) in the p_T range 1–3 GeV/c have been reported. The modification of the fragmentation function corresponds to about one additional particle per jet produced in this kinematic region.

The aim of ALICE is to extend these measurements to lower constituent p_T and lower jet p_T . The sample size required to perform measurements of this kind has been discussed in the ALICE-PPR [122]. The basic assumptions used for the simulations, in particular, the large out-of-cone radiation have shown to be correct. It has been argued that $O(10^4)$ jets per p_T bin should be sufficient to perform such measurements down to the lowest particle p_T . For the approved running scenario the jet p_T reach for this analysis is about 150 GeV/c for charged jets and about 200 GeV for the full jets within EMCal/DCAL as shown in Figure 2.64 (left).

Using the existing measurements some of the assumptions can be refined and combined. They strongly motivate the need for an increase of the sample size by an order of magnitude:

- The strength of ALICE is particle identification. So far no modification of the particle composition of jets has been observed. Performing jet structure analysis for identified particles needs an increase of the sample size corresponding to the particle fraction and PID efficiency.
- Large energy imbalance of di-jets have been reported by ATLAS and CMS. Hence, stronger jet structure modifications are expected for the jet opposite to a trigger jet. The requirement of having two jets within




Figure 2.63: Left: Nuclear modification factor for jets reconstructed from charged particle in central and peripheral Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, the reference is obtained from PYTHIA simulation at the same \sqrt{s} . Right: Ratio of reconstructed charged jets with different cone radii in central and peripheral Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV compared to PYTHIA.




Figure 2.64: Jet yield and number of produced jet-pairs above (leading) p_T -threshold for 0-10% most central Pb–Pb events. The yields are obtained by geometry (T_{AA}) scaled PYTHIA8 simulations, no quenching effects have been considered. Left: After the upgrade. Right: Before the upgrade and with TPC rate limitations in case of the charged jet reconstruction, i.e. no additional triggering on charged jets.

our η -acceptance decreases the yield by approximately an order of magnitude (see Figure 2.64).

- The same applies to hadron–jet correlations.
- An excellent control of the systematic uncertainties related to the background subtraction is needed at high R and low p_T . In particular flow modulations have to be taken into account. Hence all analysis should be performed for several bins in jet direction with respect to the reaction plane.
- The unfolding of jet spectra for high R and low p_T cuts need an increased lever arm at high jet p_T to compensate for the increased background fluctuations.

2.4.3 Photon–Jet Measurements

Direct photons, defined as all photons not originating from hadronic decays, provide together with dileptons one of the most versatile tools to study the medium created in heavy-ion collision. Depending on the momentum scale they probe different stages of the reaction, since, once produced, they escape the strongly interacting medium basically unaffected. At low transverse momentum the spectrum of direct photons is dominated by thermal production and its measurement can provide the temperature averaged over the evolution of the medium [123], while direct photons at large transverse momentum originate from early hard parton scatterings. The production rate of these photons provides a direct control of the rate of initial hard scatterings and the validity of scaling hard processes by the number of binary collisions N_{coll} when comparing proton–proton to nucleus–nucleus collisions, e.g. via the nuclear modification factor R_{AA} which is unity for high p_{T} direct photons [124, 125] and other color neutral probes such as the Z^0 [71, 126].

At leading order (LO) hard direct photons are produced via quark–antiquark annihilation ($q\bar{q} \rightarrow \gamma g$) or quark–gluon Compton scattering ($qg \rightarrow \gamma q$), where the latter is the dominant process for $p_{\text{T}}^{\gamma} < 50 \text{ GeV}/c$ with $\sigma_{\gamma q}/\sigma_{\gamma g} \gtrsim 8$. Thus measuring the momentum of a prompt photon together with the recoiling jet provides directly the momentum scale of the hard scattering and is the ideal probe for (quark) energy loss and medium modification of the (quark) fragmentation process. At next-to-leading order (NLO) an additional source of photons is given by the photon radiation off a scattered quark (bremsstrahlung), in this case the hard photon is accompanied at the near side by a fragmenting quark, and no longer isolated.

Due to their smaller production cross section, photon–jet measurements cannot cover the same kinematical range as fully reconstructed dijets, but photon–jet or photon–hadron correlation measurements, allow to explore the region of parton $p_{\text{T}} < 30 \text{ GeV}/c$ that is not well accessible in Pb–Pb with full jet reconstruction. The advantage compared to the measurements at higher p_{T} is the study of hard processes in the transition region, where the interplay and between hard processes and the soft medium is largest. The excellent particle identification and low p_{T} tracking capabilities of ALICE are a pre-requisite in understanding this interplay, in particular to address the questions of modified jet chemistry for light flavoured quarks picking up flavour from the bulk via recombination or coalescence, and its path length dependence.

The measurement directly connects to the triggered, identified hadron–jet correlation studies. In contrast to hadron triggers, the direct photon trigger is not affected by parton energy loss and imposes only little geometrical bias, which can be further reduced after selection of LO photons via isolation cuts. These photons should be unaffected by the medium and exhibit no azimuthal asymmetry ($v_2^{\gamma, \text{LO}} = 0$), providing an unbiased constraint on the direction of the recoiling quark traveling through the medium. A further differentiation will be provided by the study of the recoiling (quark) jet composition. Due to the large, fluctuating underlying event, the separation of fragmentation (NLO) photons needs to be done on a statistical basis. It will provide extra information on the energy loss of the accompanying parton depending on the time of photon emission. I.e. early radiation from an unmodified parton vs. late radiation and medium induced photon bremsstrahlung.

Above $p_{\text{T}} \approx 30 \text{ GeV}/c$ the study of photon–jet production and fragmentation pattern bridges the measurements to the region where full jet reconstruction in central heavy-ion collisions is no longer limited by fluctuations of the measured jet momentum due to the underlying event. It provides the necessary information to evaluate biases in full jet reconstruction in heavy-ion collisions and the complementary measurement of quark fragmentation to the gluon–jet dominated inclusive jet sample.

The major challenge in the measurement of direct photons is their separation from the background of decay photons, which is dominated by π^0 and η decays. The extraction of the direct photon signal is often done on a statistical basis, e.g. by comparing the yield of inclusive photons γ_{inel} to the expectation from a hadronic decay cocktail γ_{decay} . Since π^0 ’s provide the largest contribution to the decay background and also share many systematic uncertainties with the inclusive photon measurement, the measured inclusive and calculated decay photon yields are conveniently normalized to the π^0 yield in the double ratio:

$$R^{\gamma} = \frac{N^{\gamma}/N^{\pi^0}|_{\text{meas}}}{N^{\gamma}/N^{\pi^0}|_{\text{cocktail}}}. \quad (2.5)$$

A direct photon signal is seen as an excess of R^{γ} above unity, and $(R^{\gamma} - 1)$ provides directly the signal-to-background ratio. It is also the essential input for each statistical measurement of direct photon observables, e.g. for the separation of the direct γ –hadron correlation function and the successful extraction of the quark fragmentation




Figure 2.65: Ratio of photons per π^0 as measured by ALICE in the 0-40% most central Pb-Pb collisions at 2.76 TeV, compared to the expectation from hadronic decays. The expected signal from a NLO pQCD calculation of the direct photon yield scaled by the number of binary collisions and combined with the measured π^0 yield is shown as blue band.

function as demonstrated by the PHENIX-Experiment in pp and Au-Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV [127, 128].

In ALICE, two experimental approaches to photon reconstruction are pursued, one is via the electromagnetic calorimeters in ALICE which are discussed in detail in [122, 129], the other is via the tracking of e^+e^- from photon conversions in the material of the inner detectors. The two methods are largely independent and carry distinctly different types of systematic uncertainty, e.g. the opposite dependence of momentum/energy resolution for the tracking and in the calorimetric measurement, or the different contamination and purity of the direct photon signal. As demonstrated in the combined ALICE neutral meson result in proton-proton collisions [130], this approach allows to keep the total systematic uncertainty small over a wide momentum range (below 5% in the case of the π^0 measurement). For direct photons, first preliminary results on the measurement of the double ratio Pb-Pb are shown in Figure 2.65. In central heavy-ion collisions the photon background from hadronic decays is significantly suppressed and the observed direct photon signal above 5 GeV/c shows a good agreement with the expectation from scaled pQCD reference, which is indicated as blue band. The extraction of a direct photon-jet (hadron) correlation is currently limited by statistics but already at this early stage the systematic uncertainty can be as low as 7.5%. Below 5 GeV/c , additional sources contribute to the direct photon yield and become visible as an excess above unity. These can be thermal radiation from the QGP (see Chapter 2.3), but also photon radiation from jet-medium interactions.

The use of isolation cuts around a direct photon candidate improves the signal-to-background ratio since it reduces the contribution from high p_T photons produced either in a hadronic decay, directly in the fragmentation, or via a higher order processes, e.g. bremsstrahlung off a hard scattered quark. This has been done successfully down to $p_T = 20$ GeV/c by the CMS experiment [125] demonstrating the expected consistency of isolated direct photons with the scaled reference measurement in proton-proton. Isolation cuts in the dense background of heavy-ion collisions are sensitive to the same background fluctuation as the jet reconstruction which hinders their direct application on an event-by-event basis at lower p_T . The ALICE approach to a precision measurement of these fluctuations on particle basis will allow for a statistical separation of the LO and NLO photons in this regime.

The focus of the photon-jet measurements with the upgraded ALICE-experiment will be to extend the direct photon-jet measurements to the lowest possible transverse momentum with two independent methods to

- Separate the medium effects on NLO photons on a statistical basis, via isolation cuts.
- Study the composition and medium modification of the fragmentation function and the possible coalescence of quark jets recoiling off a LO photon.

To estimate the prospects of these measurements in central Pb-Pb collisions at 5.5 TeV, the expected direct photon




Figure 2.66: Left: Signal-to-background ratio of direct photons to expected decay photons. The suppression of the decay background in heavy-ion collisions has been taken into account by scaling the PYTHIA8 hadron decay simulation with the measured hadron R_{AA} where available. Right: Projected $n\sigma$ separation of direct photon signal in central Pb-Pb collisions at 5.5 TeV in the double ratio. Systematical (7%) and statistical uncertainties have been combined. The measured hadron R_{AA} in central Pb-Pb collisions at 2.76 TeV has been used to scale the decay background.

signal-to-background ratio with the known sources of direct and decay photons is investigated:

$$\frac{N_{\text{direct}}^{\gamma}}{N_{\text{decay}}^{\gamma}} = R^{\gamma} - 1. \quad (2.6)$$

Here the ratio is derived from PYTHIA8 simulations. The direct photon yields are consistent with NLO pQCD calculations above $p_T = 10$ GeV/c, after applying a K -factor of two, which also appears in the comparison of LO and NLO pQCD calculations. For the π^0 and other hadron decays contributing to the inclusive photon yield, the PYTHIA cross sections is used directly. This choice is motivated by the NLO-description of the measured charged particle and π^0 cross sections at LHC energies with deviations of up to a factor of two [130, 131], while at the same time the various PYTHIA tunes provide an agreement within 30%. For the case of heavy-ion collisions an additional suppression is applied to the hadronic decay background based on the measured nuclear modification factor in Pb-Pb collisions at 2.76 TeV [131, 132]. The uncertainty in the nuclear PDF has not been considered here for the direct photon production, the effect on the hadron production is covered by the data driven scaling approach via the measured R_{AA} . The approach can be verified with measured data at $\sqrt{s_{NN}} = 0.2$ and 2.76 TeV as shown in Figure 2.66 (left).

To assess the benefits of the proposed upgrade for the photon measurements we explore the yield of the direct photon signal for various trigger scenarios (triggered T, and minimum bias MB), before (17) and after the upgrade (19+) and detection methods. The evaluated detector configurations are:

- **PCT1** Photon conversions using tracking with the ITS and TPC within $|\eta| < 0.9$. The photon conversion probability in the material ($X/X_0 = 11.4 \pm 0.5\%$) is the dominating factor in the photon detection efficiency. The reduced number of detected photons is partially compensated by larger acceptance compared to the calorimeters. There is currently no online triggering on conversions, so this method benefits most from an increased MB rate.
- **PCT2** Photon conversions using tracking with the ITS and TPC within $|\eta| < 0.9$, with reduced material




Figure 2.67: Direct photon yield and number of detectable γ -jets above p_T -threshold for 0-10% most central Pb-Pb events. The yields are obtained by geometry (T_{AA}) scaled PYTHIA8 simulations. Left: Upgrade scenario. Right: Without upgrade and with TPC rate limitations. In the photon conversion technique reduced reconstruction efficiency due to material thickness has been considered.

budget after the upgrade in the inner detectors from 7.3% [133] to 2.1% (total $X/X_0 \approx 6.2\%$).

- **Calo** the current setup of electromagnetic calorimeters, with the EMCAL and PHOS modules plus the approved DCAL upgrade of the ALICE calorimetry [112], all can act as trigger detector on photons, which here is assumed to be fully efficient.
- $|\eta| < 0.9$ For comparison purposes we assume all photons are detected within $|\eta| < 0.9$.

As seen in Figure 2.66 (right) the measurement is dominated by the systematical uncertainty, which is for the purpose of this study set constant to 7%, a value already achieved in the intermediate p_T region of preliminary direct photon measurement using the conversion technique. With the assumed systematical uncertainty the separation of the prompt direct photon signal is feasible down to lowest p_T . However, a sufficient statistics is also pre-requisite for control of the systematical uncertainties at a level below 10% for all p_T . A further reduction will be achieved by the complementary measurement of direct photons with two different methods. As shown in Figure 2.67 (left) the proposed upgrade will significantly improve the measurement of direct photons via conversions yielding two complementary ways to measure direct photon and photon-jet correlations with sufficient statistics for differential studies out to more than 50 GeV/c in direct photon p_T . Note that only in the upgrade scenario there is sufficient overlap between direct photon-jet and dijet measurements in the region above 40 GeV/c to allow for detailed comparison of jet fragmentation biases and jet reconstruction in the region where the combinatorial jet rate is small.

2.4.4 Heavy Flavour Jets

If radiative energy loss is the main parton energy loss mechanism, heavy quarks are expected to lose less energy than light quarks, because of the so-called dead cone effect which limits the radiation to larger angle when the velocity of the quark is below the velocity of light [134]. First results on the nuclear modification factor of heavy flavour electrons and open charm at RHIC and LHC [72, 135] tend to show a smaller difference between light and heavy hadron suppression than was initially expected.

The measurement of charm production in jets may bring qualitatively new information about heavy flavour energy loss. One can distinguish two areas of interest:

- **Production mechanism; gluon splitting** At LHC, a substantial fraction of charm is expected to be produced in NLO processes, via gluon splitting. It is likely that the energy loss in this process is different (depending on the lifetime of the gluon and the separation between the charm and anti-charm quark) than for directly produced charm quarks. The gluon splitting process mainly contributes charmed hadrons with a low momentum fraction in the jets.
- **Direct measurement of energy loss** Directly produced charm quarks are expected to emerge as leading hadrons in heavy-flavour jets. The modification of the momentum distribution of fragments at large momentum fraction is a measure of the gluon radiation off a heavy quark.




Figure 2.68: Correlation of reconstructed D^* mesons with hadrons in $\sqrt{s} = 7$ TeV p+p collisions. Left: Measured with ALICE Right: PYTHIA simulation separated for D^* originating from charm and bottom quarks.

Recent measurements of charm in jets by the ATLAS collaboration [136] already shed light on the importance of gluon splitting for charm production. In p+p collisions the momentum distribution of D^* mesons in reconstructed jets with $25 < p_T < 70$ GeV/c has been compared compared to expectations from POWHEG+PYTHIA. At low momentum fraction $z = p_D/E_{jet}$ more D^* mesons are observed than expected by the calculation. This suggests that charm production via gluon splitting is more important than was thought so far. Measurements of azimuthal correlations between B -mesons with $p_T > 15$ GeV by CMS [137] lead to a qualitatively similar conclusion for B -mesons. A preliminary ALICE study of D^* -hadron correlations in proton–proton-collisions down to 3 GeV/c is shown in Figure 2.68 (left). This complements the kinematic reach of the ATLAS and CMS results and exemplifies a reference measurement for the study of gluon splitting induced charm production and tagging of heavy-flavour jets in heavy-ion collisions. As shown in Figure 2.68 (right) from MC simulations, the dominating contribution in this kinematic region is not from directly produced charm and beauty quarks. However, it is also seen that the fragmentation of heavy quarks into high- z charm leaves only little correlated hadronic activity on the near side, which allows to tag jets and study the unbiased heavy quark-fragmentation at the recoil side.

We aim to measure charm production in jets in heavy-ion collisions, ideally over a large range of the fragmentation momentum fraction z , to gain insight in the energy loss of high-momentum leading charm quarks as well as the importance of gluon splitting and the associated parton energy loss for lower momentum open charm particles reconstructed in jets. The improved reconstruction of the various heavy-flavour decay channels for the proposed upgrade will greatly enhance the performance for tagging jets that contain heavy flavour in a broad z -range.

2.4.5 Reference Data

The quantification of medium effects on jet fragmentation requires the measurement of reference distributions in the QCD vacuum. As discussed above the main objectives of the jet measurements in heavy-ion collisions are the detailed study of jet fragmentation of low to intermediate p_T jets of different flavour, where the need for high statistics is driven by the requirement of a sufficiently large lever arm in p_T^{jet} to unfold the increasing impact of jet fluctuations towards lower p_T , the differential study of jets, e.g. with respect to the event reaction plane and an extraction of low p_T fragmentation products in a large background of the underlying heavy-ion event. These conditions do not apply to the vacuum case, here jet reconstruction is well defined and no differentiation between different propagation paths and collision geometries is needed. Thus, as reference we only aim for a more modest jet- p_T reach of more than $\approx 10^4$ jets above 120 GeV/c for charged and full jet reconstruction. As seen in Figure 2.69 (left), this is reached with the integrated luminosity of 6/nb that is aimed for the reference measurement of heavy-flavour hadrons (see Section 2.1.3). In the case of direct photon and direct photon-jet measurements, the extraction of a signal without isolation cuts will be limited to the region above ≈ 20 GeV/c




Figure 2.69: Left: Projected measured jet and dijet yields Right: Expected measured photon yields and γ -jets above p_T -threshold for proton–proton-collisions at $\sqrt{s} = 5.5$ TeV for an integrated luminosity of 6 pb. In the photon conversion technique the change of material thickness after the upgrade has been considered.




Figure 2.70: Anti- ${}^4\text{He}$ candidates identified via TPC dE/dx measurements in the full 2011 Pb–Pb statistics. The inset shows the m^2/z^2 when applied on this pre-selected data. The 10 anti-alphas clearly identified by TPC and TOF are indicated as red dots.

due to systematical uncertainties of the measurement and the worse S/N in proton–proton-collisions as seen in Figure 2.66 (right). Here, the main reference measurement will be provided by the calorimetric method alone as indicated in Figure 2.69.

2.5 Heavy Nuclear States

Another area where ALICE measurements are unique is the search for and possible spectroscopy of exotic objects produced in Pb–Pb collisions. The resolving power already of the current setup can be seen from measurements with the TPC and TOF which together with the ITS and TRD comprise the core detectors for track reconstruction, momentum determination, and particle identification.

Figure 2.70 shows the superb particle-species separation via the specific energy loss in the TPC. An example application of this particle identification (PID) is given by the recent ALICE measurement of anti- ${}^4\text{He}$ nuclei.

Table 2.4: Expected yields of exotica per 10^{10} central collisions into the acceptance of the ALICE central barrel. The numbers include an 8% efficiency per detected baryon.

Particle	Yield
Anti-alpha ${}^4\bar{\text{He}}$	3.0×10^4
Anti-hypertriton ${}^3\bar{\Lambda}\bar{H}$ ($\bar{\Lambda}\bar{p}\bar{n}$)	3.0×10^5
${}^4\bar{\Lambda}\bar{H}$ ($\bar{\Lambda}\bar{p}\bar{n}\bar{n}$)	8.0×10^2
${}^5\bar{\Lambda}\bar{H}$ ($\bar{\Lambda}\bar{p}\bar{n}\bar{n}\bar{n}$)	3.0
${}^4\bar{\Lambda}\bar{\Lambda}\bar{H}$ ($\bar{\Lambda}\bar{\Lambda}\bar{p}\bar{n}$)	3.4×10^1
${}^5\bar{\Lambda}\bar{\Lambda}\bar{H}$ ($\bar{\Lambda}\bar{\Lambda}\bar{p}\bar{n}\bar{n}$)	0.2
H-Dibaryon ($\Lambda\Lambda$)	5.0×10^6
$\Xi\Xi$	1.5×10^5
$\Lambda\bar{n}$	8.0×10^7

An offline trigger was used to select the tracks shown in Fig. 2.70, and, with further use of the TOF detector for mass identification, 10 (out of roughly 23 million Pb–Pb events) candidates were isolated, as shown in the insert. With the upgraded ALICE detector, inspection of 10^{10} central Pb–Pb collisions should be feasible. This should allow a systematic study of anti-nucleus production and bring into reach measurements on the lightest multi- Λ hypernuclei such as ${}^5_{\Lambda\Lambda}\bar{H}$ [138]. Furthermore, we plan to search for possible dibaryons and anti-dibaryons with strangeness [139, 140].

The energy dependence and expected yields for some of these objects have been computed in the framework of the statistical hadronization model. The result is shown in Fig. 2.71. We also include here in Table 2.4 the resulting yields for 10^{10} central collisions, from which the expected physics reach can be read-off. Note that for these measurements the statistical errors are by far the dominant uncertainties.

Other exotic objects to search for include bound states of ($\Lambda\Lambda$) or the H-Dibaryon, a possible ($\Lambda\bar{n}$) bound state. The possible existence of these objects can already be investigated with the present ALICE setup. We show here recent results for the H-Dibaryon and the ($\Lambda\bar{n}$) bound state on basis of about 13.8 million Pb–Pb events in the centrality of 0–80% taken with the ALICE apparatus in 2010. The reconstructed invariant mass is shown in Fig. 2.72. From the utilized data there is no evidence for a signal, neither for the H-Dibaryon nor the ($\Lambda\bar{n}$) bound state. The figure also show the expected signal, for the H-Dibaryon for two possible bound states (at masses of $2.21\text{ GeV}/c$ and $2.23\text{ GeV}/c$) and a possible ($\Lambda\bar{n}$) signal. With roughly a factor of 1000 more statistics obtainable with the upgraded ALICE apparatus the long-standing issue of the possible existence of these bound states can finally be settled once and for all. Should such states exist, then measurements of lifetimes and decay schemes come into reach.

A very interesting area where new results can be expected with the upgraded ALICE apparatus is the search for and possible study of bound states involving multi-strange baryons. This could open up a completely new chapter in di-baryon spectroscopy [139, 140]. The search for dibaryons like ($\Lambda\Xi$) or ($\Xi\Xi$) can be performed via invariant mass analysis and the ALICE detector is the ideal apparatus to search for such exotic bound states. Furthermore, strangeness is produced copiously at LHC energy, so that even $\Xi\Xi$ states are produced with significant yields, see Table 2.4. In addition, both particles and anti-particles can be measured, which will give an improved handle on systematic errors. Should none of these states be bound (a result which would already be very important for our understanding of hadron spectroscopy) we plan to investigate the interaction among strange and multi-strange baryons by studying, e.g., $\Lambda\Xi$ - or $\Xi\Xi$ -correlations. First results on $\bar{\Lambda} - \Lambda$ correlations have already been presented by ALICE [141] and indicate the power of this approach.

In summary, the upgraded detector can be used for investigations of multi-baryon states which were so far inaccessible.




Figure 2.71: Energy dependence of the yield of various (anti)-hypernuclei relative to the Λ yield. The calculations are for central Pb–Pb collisions, in the framework of the statistical hadronization model [138].




Figure 2.72: Invariant mass of $\Lambda + p + \pi^-$ (left) and the invariant mass for $\bar{d} + \pi^+$ (right). The figures also show the expected signal using thermal model predictions and the estimated Acceptance x Efficiency for two different possible bound states of the H-Dibaryon, respectively the expected signal for the $\bar{\Lambda}\Lambda$ bound state.

Chapter 3

Detectors and Readout Electronics

3.1 Introduction

As discussed in Chapter 2, the main physics topics addressed by the proposed upgrade require the measurement of heavy flavour hadrons, quarkonia, and low-mass dileptons at low transverse momenta. These measurements are characterized by a very small signal-over-background ratio, which calls for large statistics. In addition, the large background makes the application of triggering techniques impossible or very inefficient, as for example for the heavy-flavour case discussed in Chapter 2 of [8]. Moreover, these measurements require a significant improvement of the pointing resolution and tracking efficiency, while preserving the excellent particle identification capabilities of the current ALICE detector.

On the basis of the above considerations, substantiated by quantitative studies illustrated in Chapter 2, the experimental approach taken by ALICE is to read out all Pb–Pb events at an interaction rate of 50 kHz and to significantly improve vertexing and tracking capabilities at low transverse momentum. This implies building a new Inner Tracking System (ITS), a major upgrade of the TPC detector and a modification of the readout electronics of other detectors to comply with the high readout rate. It implies also a major upgrade of the online and offline systems.

More precisely the proposed upgrade consists of:

- A new beampipe with smaller diameter;
- A new, high-resolution, low-material ITS;
- Upgrade of the TPC consisting in the replacement of the wire chambers with GEM detectors and new pipelined readout electronics;
- Upgrade of the readout electronics of TRD, TOF, PHOS and Muon Spectrometer for high rate operation;
- Upgrade of the forward trigger detectors;
- Upgrade of the online systems and offline reconstruction and analysis framework.

This chapter deals with the description of the upgrade of the detectors and readout electronics and is organized as follows. Sections 3.2 and 3.3 give a short description of the specifications of the new beampipe and an overview of the new ITS detector. The upgrade of the TPC is described in Section 3.4, while the upgrade of the readout electronics of the TOF detector is discussed in Section 3.5. The upgrade items covered by the first four sections have been studied in the context of the measurement of heavy-flavour and low-mass dielectrons. The study of quarkonia at mid-rapidity and forward rapidities, requires a modification of the readout electronics of the TRD detector, which is discussed in Section 3.6, and new readout electronics for the Muon Spectrometer, which is discussed in Section 3.7. Plans for the upgrade of other detector systems are discussed in Section 3.8. The upgrade of the online systems, and of the offline reconstruction and analysis framework, are discussed in the following chapter. The specifications for the readout of the ITS, TPC, TRD, TOF and Muon Spectrometer and their mode of operations are listed in Table 3.1. It should be noticed that the ITS, TPC and TOF detectors will be able to record all minimum bias Pb–Pb interactions. The information of the other detectors, whenever available, will be

Table 3.1: Readout rates of the main ALICE detectors after the upgrade.

Detector	Max R/O rate (kHz) (pp and Pb–Pb)
ITS	continuous
TPC	continuous
TOF	200 – 400
TRD	27 – 100
EMCal	50
Muon	5

combined at the event building stage in the online systems. In the following we describe the modes of operation envisaged for the main physics observables.

- For the measurement of heavy-flavour and low-mass dielectrons, the key detectors are the ITS, TPC and TOF. ITS and TPC will be read out continuously while the TOF can be read out up to a rate in the range 200–400 kHz both in Pb–Pb and pp.
- For the measurement of quarkonia at mid-rapidity, the TRD will be read out with a minimum-bias trigger at a maximum rate of 27 kHz or 100 kHz in Pb–Pb, depending on whether the raw ADC data or the track let parameters are read out. The information of the TRD will be combined to the one of the ITS, TPC and TOF.
- For the measurement of quarkonia at forward rapidities, the Muon Spectrometer will be read out on its own trigger at a maximum rate of 5 kHz. The information of the ITS, which is used for the reconstruction of the primary vertex, will be combined to the Muon Spectrometer data in the online system.

For the study of jets, the EMCal will be read out either upon the minimum bias trigger, up to 50 kHz, or on its own trigger and its information combined to the one of the other central detectors at the event building stage.

3.2 Beampipe

The reduction of the beampipe diameter in the center of the ALICE detector is one of the main ingredients in view of improved impact parameter resolution. Current studies indicate that it should be possible to arrive at a beampipe inner radius of 17.2 mm, to be compared to the present value of 29 mm. Estimates for the linear sum of fabrication tolerance, survey precision and alignment uncertainties amount to 5.1 mm, resulting in a minimum clearance of 12.1 mm radius with respect to the nominal beamline. The LHC aperture is quoted in terms of the so called n_1 parameter which is a function of this mechanical clearance as well as the position along the beam line due to the varying beta function. The 12.1 mm clearance corresponds to an aperture of $n_1 = 13.9$ at the ALICE interaction point for nominal injection optics. Beyond a distance of 2 m from the interaction point the minimal aperture requirement of $n_1 > 10$ is however violated and a larger beampipe radius is foreseen. In the current layout we therefore assume a conical or cylindrical beampipe with larger diameter beyond a distance of 1 m from the IP. The wall thickness of the central Beryllium beampipe section is assumed to be 0.8 mm. Using a smaller value of 0.5 mm is challenging due to possible issues with gas tightness and mechanical stability. The wall thickness and material for the beampipe section beyond 1 m from the IP will be decided by background simulations and practical considerations. Since the machine parameters for the high luminosity heavy ion running are not yet defined and since ALICE wants to operate this detector configuration also during the HL-LHC era after LS3, a conservative number of 19.2 mm for the beampipe inner radius is assumed until further studies of LHC optics confirm the feasibility of the 17.2 mm radius.

3.3 ITS Upgrade

A detailed discussion of the physics motivation, detector requirements and technical implementation of the ITS upgrade is presented in a conceptual design report [8]. In this section we give first an overview of the performance and limitations of the current ITS and then we shall discuss the conceptual design of the new ITS and its simulated performance. For a detailed discussion of the technical implementation of the proposed detector and the ongoing R&D activities we refer to Chapters 4 and 5 of [8].

Table 3.2: Characteristics of the six ITS layers, the beam-pipe and the thermal shields.

Layer / Type	r [cm]	$\pm z$ [cm]	Number of modules	Active area per module $r\phi \times z$ [mm 2]	Intrinsic resolution [μm]	Material budget X/X_0 [%]
				$r\phi$	z	
Beam pipe	2.94	-	-	-	-	0.22
1 / SPD	3.9	14.1	80	12.8×70.7	12 100	1.14
2 / SPD	7.6	14.1	160	12.8×70.7	12 100	1.14
Th. shield	11.5	-	-	-	-	0.65
3 / SDD	15.0	22.2	84	70.2×75.3	35 25	1.13
4 / SDD	23.9	29.7	176	70.2×75.3	35 25	1.26
Th. shield	31.0	-	-	-	-	0.65
5 / SSD	38.0	43.1	748	73×40	20 830	0.83
6 / SSD	43.0	48.9	950	73×40	20 830	0.83

3.3.1 Current ITS Performance and Limitations

The present ITS consists of six cylindrical layers of silicon detectors placed coaxially around the beam pipe. They are located at radii between 39 mm and 430 mm and cover the pseudo-rapidity range $|\eta| < 0.9$ for vertices located within ± 60 mm with respect to the nominal interaction point. Within the boundaries set by technological limitations and available funds, the number, position and segmentation of the layers were optimized to achieve a high precision in the determination of the charged particle distance of closest approach (dca) to the primary vertex and efficient track finding in combination with the TPC. The inner radius is the minimum allowed by the radius of the beampipe. The outer radius is determined by the necessity to match tracks with those from the TPC. As it will be illustrated in the following, optimizing the detector geometry to achieve a higher tracking efficiency at very low transverse momentum, based on the ITS standalone tracking, would lead to an alternative configuration including a larger number of layers and different radii. The first layer has a more extended pseudo-rapidity coverage ($|\eta| < 1.98$) which, together with the Forward Multiplicity Detectors (FMD), provides continuous coverage for the measurement of charged particle multiplicity.

As a result of the high particle density (the current system is designed for up to 100 particles per cm^2 for Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV), and in order to achieve the required dca resolution, the first two layers of the ITS are made of Silicon Pixel Detectors (SPD) and the two middle ones are made of Silicon Drift Detectors (SDD). The two outer layers, where the track density has fallen to one particle per cm^2 , are equipped with double-sided Silicon micro-Strip Detectors (SSD). The four outer layers have analogue readout and therefore can be used for particle identification via dE/dx measurement in the non-relativistic ($1/\beta^2$) region. Based on the technologies available at that time, all detector elements were optimized to minimize their radiation length, achieving between 0.8 % and 1.3 % of X_0 per layer. Another 1.3 % of X_0 comes from the thermal shields and supports installed between SPD and SDD barrels and between SDD and SSD barrels, thus making the total material budget for perpendicular tracks equal to 7.63 % of X_0 . The geometrical parameters of the layers (radial position, length along beam axis, number of modules, spatial resolution) and the material budget are summarized in Table 3.2.

Figure 3.1 shows the impact parameter resolution as a function of transverse momentum. The resolution in the bending plane ($r\phi$), which is mainly determined by the ITS, is plotted for ITS stand-alone tracks for Pb-Pb data and Monte Carlo simulations with realistic residual misalignment. The impact parameter resolution in the longitudinal direction (z), is plotted for the ITS-TPC combined tracking. Owing to its low material budget, the impact parameter resolution at low p_T ($\leq 1 \text{ GeV}/c$) of the ITS is the best among the LHC experiments. Nevertheless, in order to access the measurement of charmed mesons and baryons at very low transverse momentum, an improvement of at least a factor 3 is required.

In the left panel of Figure 3.2 the relative p_T resolution is shown as a function of p_T for ITS stand-alone and ITS-TPC combined tracks, which have been selected requiring 6 ITS clusters per track. The p_T resolution provided by the ITS stand-alone is about 6 % for tracks with $p_T < 2 \text{ GeV}/c$. Due to the smaller lever arm and the limited number of points, this is worse by about an order of magnitude with respect to the ITS-TPC combined tracks.

In the right panel of Figure 3.2 the ITS stand-alone tracking efficiency is shown as a function of p_T for Monte Carlo Pb-Pb central collisions. The efficiency is defined as the number of “good” refitted tracks (i.e. tracks with ≥ 3 associated clusters and without any fake associated cluster from another track) divided by the number of




Figure 3.1: Impact parameter resolution of present ITS versus p_T , for the $r\phi$ (left, Pb–Pb data and MC) and z (right, pp MC simulation [142]) components. Reconstructed tracks have been selected requiring that all ITS layers have a cluster associated to the track.




Figure 3.2: Left panel: Monte Carlo simulation of ITS stand-alone p_T resolution compared to ITS+TPC combined resolution [142]. Right panel: ITS stand-alone tracking efficiency as a function of p_T obtained in a Monte Carlo simulation of central PbPb collisions without any ITS dead module.

“trackable” particles. A “trackable” particle has been defined as a particle with at least three reconstructed clusters in three different ITS layers and at least one of the reconstructed points on one of the three innermost layers. Figure 3.2 shows the tracking reconstruction efficiency for the ITS detector with all modules 100 % efficient and for the detector configuration during the 2010 Pb–Pb run.

The four layers equipped with drift and strip detectors provide a measurement of the specific energy loss, dE/dx , which can be used for particle identification. For each track, the dE/dx is calculated using a truncated mean. Figure 3.3 shows the truncated mean dE/dx for a sample of ITS stand-alone tracks along with a parametrization of the most probable value [143] based on the Bethe–Bloch formula.

A crucial limitation of the present ITS detector is given by its poor readout rate capabilities. Table 3.3 lists the readout time and maximum event readout rate under the assumption of no back pressure from the DAQ network. It should be noted that the readout time of the three ITS detectors depends only marginally on the detector occupancy and, therefore, is very similar for p–p and Pb–Pb events. From this table, we conclude that the ITS can run up to a maximum of about 1 kHz (with 100% dead time) irrespective of the detector occupancy. For the physics observables aimed at by the upgrade physics programme, for which selective triggers cannot be applied, this rate limitation will restrict ALICE to use only a small fraction of the full Pb–Pb collision rate. It will also prevent the




Figure 3.3: Specific energy loss dE/dx as a function of the momentum measured with the ITS stand-alone in Pb–Pb collisions. The solid lines are a parametrization of the detector response based on the Bethe–Bloch formula.

Table 3.3: Readout time and maximum rate, assuming 100% dead time, of the ITS sub-detectors.

Detector	R/O time (μs)	Max. rate (Hz)
SPD	296	3300
SDD	1023	985
SSD	310	3265

collection of the required reference data in pp collisions. Clearly the present ITS is totally inadequate to fulfill the required rate capabilities envisaged for the ALICE long-term plans discussed in the previous chapters.

Finally, the impossibility to access the present ITS detector for maintenance and repair interventions during the yearly LHC shutdowns represents a major limitation for sustaining high data quality. In the context of an upgrade, the rapid accessibility to the detector will be set as a priority.

3.3.2 ITS Upgrade

In this section, the key features of the ITS upgrade will be discussed and compared to the present ITS.

- **First detection layer closer to the beam line.** At present the radial distance of the first layer from the beamline is 39 mm. As discussed in Section 3.2, the installation of a new beampipe with an outer radius of 20 mm is considered a conservative possibility. The installation of such a beampipe would enable the first detection layer to be located at a radius of 22 mm. The baseline for the wall thickness of the beampipe in this document is 0.8 mm.
- **Reduction of material budget.** Reducing the material budget of the first detection layer is particularly important for improving the impact parameter resolution ϵ reducing the probability for photon conversion. In general, reducing the overall material budget allows the tracking performance and momentum resolution to be significantly improved. The use of Monolithic Active Pixel Sensors (MAPS) allows the silicon material budget per layer to be reduced by a factor of 7 in comparison to the present ITS (50 μm instead of 350 μm). MAPS detectors feature a very low capacitance at the input of the analogue front-end. This fact together with a careful optimization of the analogue front-end timing specifications and readout architecture allow reducing the power density by a factor of 2 while increasing the pixel density by a factor of 50. The lower power consumption and a highly optimized scheme for the distribution of the electrical power and signals allow reducing the material budget of the electrical power and signal cables by a factor of 5. Mechanics, cooling and other detector elements can also be improved when compared to the present ITS design. Combining all these new elements, it should be possible to build a detector with a radiation length

of 0.3% of X_0 per layer or better.

An example of the feasibility of such a design is represented by the STAR HFT detector [144]. Detailed studies on material budget and prototypes of detector ladders developed for the ITS upgrade, which show the feasibility of such a design, are discussed in Chapter 5 of [8].

Detector layers based on hybrid pixel technology, with the same optimization of timing specifications and readout architecture mentioned above, would feature a higher radiation length (0.5% X_0). This would still imply a significant improvement of the performance as compared to the present ITS.

- **Geometry and segmentation.** The studies presented in this document are based on a detector consisting of seven concentric cylindrical layers covering a radial extension from 22 mm to 430 mm with respect to the beamline. The physics studies of the benchmark channels presented in Chapter 2 are based on the assumption that all layers are segmented into pixels with dimensions of $20 \times 20 \mu\text{m}^2$. However, the detector performance in terms of impact parameter resolution and standalone tracking efficiency will not change significantly, in case the outer four layers would have a much lower granularity as for example with the silicon micro-strip detectors.
- **Measurement of energy loss.** Whether the physics goals targeted in this LoI require the new detector to preserve PID capabilities by measuring the ionization in the silicon layers is still under study. If this will turn out to be an important requirement, this could be achieved also in the scenario where all detector layers are implemented with monolithic pixel technology. In this case the performance of the detector would be slightly reduced with respect to the present one. Nevertheless it would maintain a 3σ separation power up to 600 MeV/c for pion to kaon and up to 1 GeV/c for proton to kaon. It should be noticed that all physics performance studies presented in Chapter 2 do not assume PID in the silicon tracker.
- **Readout time.** As we have seen in Section 3.3.1, the present ITS features a maximum readout rate of 1 kHz. The new detector aims to read the data related to each individual interaction up to a rate of 50 kHz for Pb–Pb collisions and several hundred kHz for pp collisions. The new ITS readout will support a "continuous readout" and "triggered readout" modes. In "continuous readout" mode, the information of the pixel matrices is shipped off-detector continuously regardless of the occurrence of an interaction. The association of pixel hits to an event is done at the reconstruction level in the online system, as explained in Chapter 4. In "triggered readout" mode, only the hits integrated over a time window that includes the event will be shipped off-detector. The actual value of the integration time window is part of the R&D and is expected to be between 5 μs and a maximum of 32 μs .

In the following section we shall illustrate that a new silicon tracker featuring the characteristics listed above will enable the track position resolution at the primary vertex to be improved by a factor of 3 or larger. The standalone tracking efficiency at low transverse momentum would dramatically improve with respect to what can be presently achieved by combining the information of the ITS and the TPC. The relative momentum resolution of the ITS standalone would be about 2% up to 2 GeV/c and remain below 3% up to 20 GeV/c.

In summary, a fully upgraded ITS will consist of 7 layers of silicon tracking detectors. The number of layouts under study has been reduced to two to minimize the complexity of the system in terms of number of different components and designs.

- **Layout 1:** The layers will be made entirely of silicon pixel detectors. All layers will provide the same intrinsic resolution of 4 μm .
- **Layout 2:** The 3 inner layers will be made of silicon pixel detectors, followed by 4 layers of double sided silicon strip detectors. This layout will provide a better particle identification.

Two basic technology choices are considered for the ITS pixel detector: hybrid silicon pixel detectors and MAPS. In Chapter 4 of [8] we illustrate the key R&D areas for both pixel detector technologies, for the new Silicon micro-Strip Detector and for the system aspects. Table 3.4 summarizes the key parameters for the two detector configurations.

The number of layers and their radial positions have been determined taking into account the available space between the new beam pipe and the outermost radius of the current ITS. The outcome of the simulations indicates that an improved tracking efficiency and p_T resolution in stand-alone mode is obtained by grouping the layers in an innermost triplet, an intermediate pair and an outermost pair as indicated in Figure 3.4.




Figure 3.4: Schematic layout of the upgraded ITS.

Table 3.4: Characteristics of the two upgrade layouts. The numbers in brackets refer to microstrip detectors:

Layer / Type	r [cm]	$\pm z$ [cm]	Intrinsic resolution [μm]	Material budget X/X_0 [%]
			$r\phi$	z
Beam pipe	2.0	-	-	0.22
1 / new pixel	2.2	11.2	4	0.30
2 / new pixel	2.8	12.1	4	0.30
3 / new pixel	3.6	13.4	4	0.30
4 / new pixel (strip)	20.0	39.0	4(20)	0.30 (0.83)
5 / new pixel (strip)	22.0	41.8	4(20)	0.30 (0.83)
6 / new pixel (strip)	41.0	71.2	4(20)	0.30 (0.83)
7 / new pixel (strip)	43.0	74.3	4(20)	0.30 (0.83)

The performance of the upgraded and the current ITS is compared in Figures 3.5, 3.6 and 3.7 for both ITS stand-alone and ITS-TPC combined tracking.

Figure 3.5 shows the pointing resolution to the vertex for charged pions. Left and right panels show the ITS stand-alone and ITS-TPC combined tracking mode, respectively. Both the $r\phi$ and z components are shown in the same plots for the pixel (Layout 1) and the combined pixel/strip (Layout 2) configurations (see Table 3.4). As an example, at a p_T of 400 MeV/c, an improvement by a factor of 3 and 5 is achieved for the $r\phi$ and z components, respectively. It should be noticed that for the present ALICE set-up the ITS-TPC combined tracking provides at high p_T a sizable improvement with respect to the ITS stand-alone tracking. Conversely, in the case of the upgraded ITS, adding the information from the TPC does not yield any further improvement.

Figure 3.6 illustrates the improvements in terms of the transverse momentum resolution. For the ITS stand-alone tracking mode, the upgraded ITS yields a dramatic improvement. It should be noticed that the p_T resolution in the stand-alone mode benefits significantly from the intrinsic resolution of the outer layers, and in particular at high p_T . In the ITS-TPC combined tracking mode (Figure 3.6) at low p_T (below 400 MeV/c) both upgraded ITS configurations would improve the resolution due to the reduction of the material budget of the innermost layers. At high p_T the full pixel configuration provides a better resolution.

Finally, Figure 3.7 shows the tracking efficiency for the ITS stand-alone tracking and the ITS-TPC combined tracking (right panel). An impressive improvement is obtained with the upgraded ITS configuration for $p_T < 1$ GeV/c. At p_T values above 200 MeV/c the two upgrade layouts yield comparable results, while below 200 MeV/c Layout 1




Figure 3.5: Pointing resolution to the vertex of charged pions as a function of the transverse momentum for the current ITS and the upgraded ITS. Left panel: ITS stand-alone tracking; right panel: ITS-TPC combined tracking.




Figure 3.6: Momentum resolution as a function of transverse momentum for charged pions for the current ITS and the upgraded ITS. The results for the ITS stand-alone and ITS-TPC combined tracking are shown in the left and right panel, respectively.




Figure 3.7: Tracking efficiency of charged pions for the current and upgraded ITS for the ITS stand-alone (left panel) and ITS-TPC combined (right panel) tracking.

features a higher tracking efficiency.

Concerning the performance of the upgraded detector in terms of PID, three detector configurations with different numbers and thickness of the layers have been studied:

- a) **4 layers of $300\ \mu\text{m}$ thickness.** The performance of the present ITS in pp collisions was used to extract the specific energy loss information. A noise of 700 electrons and a charge collection efficiency of 95 % were

assumed in the simulation as well as an 11-bit ADC with a dynamic range of 20 MIP.

- b) **7 layers of $15 \mu\text{m}$ thickness.** The dimension of the pixels is assumed to be $20 \times 20 \mu\text{m}^2$ and an analogue readout has been considered using an 8-bit ADC.
- c) **3 layers of $100 \mu\text{m}$ thickness + 4 layers of $300 \mu\text{m}$ thickness.** A Time-Over-Threshold (TOT) readout has been considered for the 3 layers of $100 \mu\text{m}$ thickness. A signal over noise ratio of 46, as reported in [145], and an 8-bit clock counter with a dynamic range of about 20 MIP have been assumed. For the 4 layers of $300 \mu\text{m}$ thickness, the same readout as in configuration a) has been taken into account. For all layers the charge collection efficiency has been fixed at 95%.

The comparison of the three configurations for the particle species separation is reported in Figure 3.8. For configuration a) it is possible to have a 3 sigma separation of kaons from protons with momentum lower than $1.2 \text{ GeV}/c$ and pions with momentum lower than $0.7 \text{ GeV}/c$. This performance is very similar to the one obtained with configuration c). A worsening in the separation capabilities is found, as expected, for configuration b). The latter has been studied assuming the detector performance of the MIMOSA chip [146]. In this simulation the charge sharing and the collection efficiency have been introduced. A threshold equal to 120 electrons and a Gaussian noise of 20 electrons have also been taken into account.

3.3.3 Technical Implementation

As discussed in the previous section, two option are being considered for the realization of the new ALICE silicon tracker: Layout 1, which consist of 7 layers of silicon pixel detectors, and Layout 2, which consists of 3 inner layers of silicon pixel detectors and 4 outer layers of silicon micro-strip detectors. A number of technologies could potentially fulfill the requirements of the inner layers. The technologies can be divided into two main categories: monolithic and hybrid silicon pixel detectors. Concerning the micro-strip detectors, the layout of the current ITS represent the basis for the new design. In order to cope with the expected occupancy levels at smaller radii a new design with reduced strip length is foreseen. In the following we shall give an overview of the technologies that are being considered for the pixel detectors and the ongoing R&D activities.

Hybrid pixel detectors In these detectors the front-end chip and the sensor are produced on two different wafers and then connected using bump bonds. Present bump bonding techniques are limited to a pitch of $30\text{--}50 \mu\text{m}$. However, the recently introduced Cu-pillar technology may reduce substantially this limit in the near future. In the current state-of-the-art pixel detectors, the silicon represents one of the main contributions to the material budget. The silicon thickness of the present ITS detector is $150 \mu\text{m}$ and $200 \mu\text{m}$ for the front-end chip and the sensor respectively. The target thickness of the front-end chip and the sensor of the new ITS is set to $50 \mu\text{m}$ and $100 \mu\text{m}$, respectively. These requirements are unprecedented and necessitate a detailed study of the silicon wafer thinning process. A first trial using dummy components with the layout of the actual ALICE pixels has




Figure 3.8: Pion to kaon separation (black circles) and proton to kaon separation (red triangles) in unit of sigma in the case of 4 layers of $300 \mu\text{m}$ (left panel), 7 layers of $15 \mu\text{m}$ (central panel) and 3 layers of $100 \mu\text{m}$ + 4 layers of $300 \mu\text{m}$ (right panel) silicon detectors. The horizontal lines correspond to a 3 sigma separation.

been started successfully with IZM¹. Assemblies consisting of 5 chips, with 50 μm thickness, connected to a 70 mm long sensor tile, with 100 μm thickness, have been produced. Currently, the sensors and the bump bonding represent the main cost factors for the production of hybrid pixel detectors and prevent their application to large surfaces as those of the outer four layers.

Monolithic pixel detectors Monolithic pixel sensors use as detection volume the p-type epitaxial layer grown on the highly p-doped silicon substrate during standard microelectronics CMOS process. The epitaxial layer typically has a thickness of 10–18 μm , thus the most probable signal generated by a MIP is of the order of $\approx 10^3$ electrons assuming an average production of 80 e-h pairs/ μm . In the standard monolithic sensor the epitaxial volume is only partially depleted and the charge generated by an ionizing particle is collected by regularly implanted NWELLS, through a mechanism that combines thermal diffusion and drift. CMOS pixel sensors typically feature pixel dimensions of $20 \times 20 \mu\text{m}^2$. They allow the integration of the signal processing circuitry on the same substrate as the sensitive volume. Therefore they offer a significant reduction in cost since only CMOS wafers are used. In recent years CMOS sensor technology has reached an adequate level of maturity to be chosen to equip the vertex detector of the STAR experiment [147], which is currently under construction. Current state-of-the-art monolithic pixels only use NMOS transistors in the pixel cell. The NWELLS needed for PMOS transistors fabrication would act as competing charge collection centers with respect to the main diode. Therefore, the front-end cell has to be kept very simple and a serial readout scheme is needed. The use of a deep PWELL underneath the PMOS devices acts as a shield for the generated charges and allows to implement more complex electronics in the pixel cell. Monolithic pixels with a hybrid-like front-end, embedding preamplifier, shaper and discriminator inside the pixel cell, have been recently demonstrated using the INMAPS process [148] at TOWER/JAZZ². The R&D on monolithic pixel detectors specific for the ITS aims at developing a monolithic pixel chip with an integration time smaller than 32 μs in the TOWER/JAZZ technology and verify their radiation resistance to the levels expected for the innermost layer of the ITS. First prototypes designed in 2011 have been characterized. The results show that the technology and architectures adopted are suitable for the ITS application.

For a detailed discussion of the detector conceptual design, the technologies under consideration, the mechanical support structure and cooling, we refer to the ITS Upgrade Conceptual Design Report [8].

3.4 TPC Upgrade

In the following we briefly describe the status and the limitations of the present TPC. The performance of an upgraded TPC with continuous readout is presented and the main directions of R&D to develop a technical implementation using GEMs are discussed.

3.4.1 Status and Limitations of the Present TPC

The ALICE TPC is the main charged particle tracking and PID device of the ALICE detector. The TPC was optimized to provide high precision momentum and dE/dx measurements in the unprecedented charged-particle densities in central Pb-Pb collisions at the LHC. Detailed descriptions of the ALICE TPC can be found in [149, 150].

The TPC employs a cylindrical drift field cage and two readout planes on the end caps. The active volume covers $87 < r < 247 \text{ cm}$ and $-250 < z < 250 \text{ cm}$ in radial and longitudinal directions, respectively. A central high voltage electrode at $z = 0$ divides the active drift volume of about 90 m^3 into two halves and provides the electric drift field of 400 V/cm. The maximum electron drift time for a Ne-CO₂ (90-10) gas mixture is about 100 μs . The neon-based mixture was chosen to keep multiple scattering and space charge effects at a minimum. At the same time, Ne-CO₂ (90-10) is well suited for operation in a moderate magnetic field of $B = 0.5 \text{ T}$ due to its low diffusion coefficient ($D_t \approx D_l \approx 220 \mu\text{m}/\sqrt{\text{cm}}$).

The readout planes are subdivided into 18 trapezoidal sectors, each covering 20° in azimuth. In radial direction, each sector is further subdivided into an inner and an outer readout chamber (IROC and OROC). The design of the readout chambers is based on Multi-Wire Proportional Chamber (MWPC) technology with cathode pad readout. The size of the readout pads is chosen to minimize the detector occupancy and optimize the position resolution and two-track separation. The pad size is $4 \times 7.5 \text{ mm}^2(r\phi \times r)$ in the IROCs, and 6×10 and $6 \times 15 \text{ mm}^2$ in the OROCs,

¹Frauenhofer Institute for Reliability and Microintegration (IZM), Dept. High Density Interconnect and Wafer Level Packaging, Gustav-Meyer-Allee 25, D-13355 Berlin

²Specialty Silicon Foundry, Ramat Gavriel Industrial Park Migdal Haemek, 23105, ISRAEL.

depending on the radial position. This results in a total of 557,568 readout channels, and an active readout area of about 32 m^2 .

The readout chambers employ a commonly used scheme of wire grids, including a plane of anode wires, cathode wires, and a gating grid. In the absence of a trigger, the gating grid wires are biased by a voltage that prevents electrons to enter the amplification region and ions from preceding events to penetrate into the drift volume.

The maximum drift time of ions from the amplification region to the gating grid is about $180 \mu\text{s}$. In order to provide efficient ion gating, the gating grid must be closed for $180 \mu\text{s}$ after the end of the electron drift. Additionally, a $100 \mu\text{s}$ post protection is applied to avoid event pile-up. This sets a first limitation of about 3.5 kHz to the maximum readout rate of the present TPC. Another limitation is given by the readout electronics that is capable of sustaining a maximum rate of minimum-bias events of 520 Hz.

Operating the TPC in ungated mode, i.e. leaving the gating grid continuously open, would result in a severe buildup of space charge in the drift volume due to back drifting ions. Owing to their low mobility ($4 \text{ cm}^2/\text{Vs}$ for Ne^+ in Ne), ions need 156 ms to reach the central electrode. At a minimum bias event rate of 50 kHz this would lead to a sustained ion accumulation from several thousand collision events piling up in the drift volume. Under such conditions, the space charge would distort the electron drift paths in such a way that a meaningful reconstruction of the particle trajectories would no longer be possible.

In addition, the anticipated particle rates at the wire chambers would reach 100 kHz/cm², for which space charge effects in the amplification region result in a few percent gain drop, thus deteriorating the dE/dx performance of the detector.

3.4.2 Upgrade Concept

Operating the TPC at a Pb–Pb collision rate of 50 kHz requires a scenario where the present limitations imposed by the operation of the gating grid can be overcome, and the TPC can be continuously read out. To this end, the present MWPC based readout chambers will be replaced by GEM detectors which feature intrinsic ion blocking without additional gating and exhibit excellent rate capabilities. This implies also the replacement of the existing front-end electronics to accommodate for the inverted signal polarity and to provide the full flexibility of a continuous readout. The existing field cage as well as most of the services of the present TPC will remain.

3.4.3 Expected Performance of the Upgraded TPC

The goal is to operate the TPC in an ungated mode while keeping the ion space charge induced drift field distortions at a tolerable level. In particular, the excellent tracking and PID capabilities as provided by the present system must be preserved. While detailed simulations of the full detector response will be documented in a future Technical Design Report, the basic performance features of a GEM-based upgrade of the ALICE TPC will be briefly discussed in the following.

3.4.3.1 Drift Field Distortions

The intrinsic capability of GEMs to prevent ions from the amplification region to flow back into the drift volume is based on the operation in asymmetric field configurations. In order to limit the drift field distortions, a maximum number of 10 ions released into the TPC drift volume for each electron entering the multiplication region is required as shown in the following. The space charge density and drift distortions for a number of 5 and 10 back drifting ions, were calculated for a 50 kHz Pb–Pb interaction rate [151]. The space charge density reaches a few tens of fC/cm³ at the smallest r and close to the central electrode (Fig. 3.9). The resulting shifts dr , $d\phi$ and dz of the electron arrival point are shown in Fig. 3.10 as a function of the initial z and r position where the electron was created. The shifts reach up to 8 cm at small r for tracks at full drift length and need to be corrected to the level of the intrinsic spatial resolution of the TPC, i.e. by about two orders of magnitude. This is the order of magnitude of the space points corrections applied for example in the TPC of the STAR experiment at RHIC. At a typical gas gain of 2000 the above requirements translate into an ion back flow probability of less than 0.5%, defined by the ratio of the ion current to the drift electrode and the electron current to the pads.

Previous experience during LHC operation shows that the luminosity is extremely stable and varies only on the timescale of the beam life time. This suggests accordingly that also the space charge distribution in the TPC changes very slowly with time. Temporal variations of the space charge distribution can be precisely monitored via the measurement of the interaction rate as well as the readout chamber anode currents. Anyway, the corresponding




Figure 3.9: Space charge density in the TPC drift volume for 5 (left) and 10 (right) back drifting ions per primary electron.

distortions of the drift field can be precisely measured and corrected for using the current TPC laser system and the data themselves. These are well established techniques in the ALICE and STAR experiments.

3.4.3.2 Position and Momentum Resolution

Precise determination of the position of an arriving charge cluster in the TPC is achieved by proper sharing of the charge on two or more adjacent readout pads. For clusters from high momentum tracks with small inclination angle, charge sharing is achieved by the spread of the charge cloud due to diffusion in the gas and the intrinsic response width of the MWPC, the Pad Response Function (PRF). Since GEM detectors do not feature a sizable intrinsic PRF, they would require a higher pad granularity. It should be noted that a pad plane design with high granularity in about 10% of the pad rows will suffice to achieve a tracking resolution comparable to the present design. In a microscopic simulation of the ALICE TPC it was studied whether diffusion in the drift gas alone is sufficient to provide adequate charge sharing when GEM detectors are combined with the present readout pad geometry. For the following results, a vanishing intrinsic PRF is assumed for the GEM case.

Fig. 3.11 shows the result for the position resolution $\sigma_{r\phi}$ in the magnetic bending plane for straight tracks with MWPC and GEM readout. At large drift distances, similar resolution can be achieved for GEM and MWPC, indicating that diffusion is efficient to provide sufficient charge sharing also with GEMs. This is not the case when the drift distance is very short, in particular for the 6×10 and $6 \times 15 \text{ mm}^2$ pads at $z > 200 \text{ cm}$, where a marked deterioration of the resolution is observed for GEMs. This can be understood as due to the progressive appearance of single-pad clusters in the GEM case, as illustrated in Fig. 3.12. It should be noted, however, that short drift distances correspond partially to regions outside the nominal acceptance of the ALICE central barrel ($|\eta| < 0.9$), as indicated by arrows in Fig. 3.11. This implies that, even with the present pad granularity, the position resolution with GEMs detectors would decrease by no more than 20% for tracks in $|\eta| < 0.75$.

The expected momentum resolution for the MWPC and GEM scenarios is shown in Fig. 3.13. In comparison to MWPC readout, a deterioration by about 20% is observed for GEMs when the momentum is determined by a fit to the TPC space points only, or when the fit is constrained by the position of the interaction vertex. However, the difference becomes significantly smaller when global tracks are considered, where tracking information from TPC and ITS is combined. In this case, the difference in momentum resolution for GEM and MWPC readout is less than 10% for high p_T tracks in $|\eta| < 0.9$. As indicated above, such a difference can be eliminated with the design of an optimized pad plane.

3.4.3.3 dE/dx Resolution

The excellent PID capability over a large momentum range via the measurement of the specific ionization dE/dx is a key feature of the ALICE TPC. In events with low particle multiplicity, a dE/dx resolution better than 5% has been achieved which is consistent with simulations [149] and empirical parametrizations [152, 153]. GEM detectors feature a local energy resolution which is similar to that of a MWPC and therefore impose no principal limitation on the achievable dE/dx resolution in a large-scale detector. The characterization of the temporal stability of the gas gain in GEMs is being addressed by the ongoing R&D activities.




Figure 3.10: Drift distortions due to space charge in the TPC drift volume for 5 (left) and 10 (right) back drifting ions per primary electron (see text).




Figure 3.11: Space point resolution in $r\phi$ as function of z for GEM (blue) and MWPC (red) readout with $4 \times 7.5 \text{ mm}^2$ (left), $6 \times 10 \text{ mm}^2$ (middle) and $6 \times 15 \text{ mm}^2$ (right) pads.




Figure 3.12: Relative frequency of one-, two-, three-, and four-pad clusters for GEM (top row) and MWPC (bottom row) readout with $4 \times 7.5 \text{ mm}^2$ (left), $6 \times 10 \text{ mm}^2$ (middle) and $6 \times 15 \text{ mm}^2$ (right) pads.




Figure 3.13: Momentum resolution as a function of $1/p_T$ for MWPC (left) and GEM readout (right). Open squares are for TPC information only, closed squares show track fits including the vertex point, and open circles show the result for combined fits to TPC and ITS tracks.

3.4.4 Technical Implementation

Significant R&D efforts for the application of GEM detectors in high-rate drift detectors, such as the ILD TPC at ILC [154] and the PANDA TPC at FAIR [155], have been conducted in the past years. Large scale GEM detectors are also foreseen for the upgrade of the CMS forward muon system [156]. In particular, the R&D work pursued for the PANDA TPC serves as a starting point for the implementation of the ALICE TPC GEM upgrade because gain and rate requirements are similar, and in both cases particular emphasis lies on the challenges connected to continuous readout and ion back flow suppression as well as on PID via dE/dx . Moreover, the PANDA TPC employs the same drift gas and drift field as the ALICE TPC.

For the upgrade of the ALICE TPC a triple GEM scheme is foreseen to provide effective gas amplification by a factor of about 2000. At a typical noise level (ENC) of $750 e^-$ [150] this corresponds to a signal-to-noise ratio of about 20 for minimum ionizing particles. Such relatively low gains can be safely provided by triple GEM structures. The expected rate of $10^{-10} \text{ Cmm}^{-2}\text{s}^{-1}$ during operation of the ALICE TPC at 50 kHz in Pb-Pb is low compared to laboratory tests, where stable operation has been demonstrated at much higher rates and gains. Similarly, the integrated charge of 1 mC/mm^2 in 10^7 's of operation in Pb-Pb is a few orders of magnitude less than what has already been accumulated by GEM detectors in COMPASS and LHCb without notable aging effects.

The technical implementation of the GEM based upgrade of the ALICE TPC will benefit from new technologies presently being developed for other large upgrade projects like the CMS forward muon system. The new single mask technology [157] allows production of GEM foils which are much larger than those commonly produced in double mask technology. This is very attractive for applications of GEMs in large area detector modules like the trapezoidal ALICE TPC readout chambers (IROC $390 \times (280\text{-}460) \text{ mm}^2$, OROC $1000 \times (468\text{-}860) \text{ mm}^2$). The use of single mask GEMs is also foreseen for the large area modules of the CMS forward muon system. In addition, a new assembly procedure for GEM detectors, the no-stretch-no-spacer (NS2) technique, is presently being developed and tested by CMS. The future large area GEM applications in LHC experiments (1000 m^2 for the CMS forward muon system, 100 m^2 for the ALICE TPC) demand cost-effective GEM foil production in industry. The RD51 collaboration is presently pursuing an industrialization program for GEM production with companies in Korea and Italy. First tests of double and single GEM foils manufactured by Newflex Technology³ showed promising results.

³NEW FLEX TECHNOLOGY CO.,LTD., Korea




Figure 3.14: This figure illustrates the ion back flow measurements carried out on a triple GEM operated at the indicated voltages. The current on the drift electrode, the current on the anode pads and their ratio (ion back flow) are shown as a function of the transfer field between the second and the third GEM foil.

3.4.5 R&D

3.4.5.1 Gas Studies

Argon based gas mixtures are most commonly used for the operation of GEM detectors. However, due to the large primary ionization and the low ion mobility ($1.7 \text{ cm}^2/\text{Vs}$ for Ar^+ in Ar) the use of argon would increase the space charge density by a factor 2.5-5. Therefore, a neon based gas mixture is also the baseline for the upgraded TPC. A dedicated R&D effort is started to optimize the gas mixture, in particular with respect to drift velocity. In Ne-CO₂ (90-10), the maximum drift time of $100 \mu\text{s}$ will lead to an overlap of 5 minimum bias events on average at a collision rate of 50 kHz, corresponding to about 5,000 charged particles per pseudo-rapidity unit. This is still less than the design goal of the ALICE TPC, which has been optimized for charged particles densities as high as $dN_{\text{ch}}/d\eta = 8,000$ [149]. Nevertheless, online reconstruction would benefit from a faster drift gas and, therefore, reduced pile-up. Neon based gas mixtures containing CF₄ promise to feature electron drift velocities that are larger than those in Ne-CO₂ (90-10) by about a factor 3 at comparable drift field. Open questions concerning electron and ion drift properties, gas gain and operational stability, aging, electronegativity and chemical reactivity in mixtures containing CF₄ remain to be settled by further R&D.

3.4.5.2 Ion Back Flow

Previous studies with Ar-based gas mixtures showed that ion back flow probabilities below 1% at zero magnetic field can be achieved. To reach the design goal of 0.5% further optimization of the GEM geometry and potential settings is required. Moreover, most of the previous measurements which were performed with argon based mixtures need to be conducted with neon. Specific R&D on this issue is being carried out with support from the RD51 Collaboration. As an example of the ongoing R&D, the ion back flow for a small TPC-like structure with a standard triple GEM arrangement and the usual Ar-based mixture is shown in Fig. 3.14. It shows the ion back flow as a function of one of the transfer fields, which varies between 0.25 and 0.35%. Considering the corresponding gas gain, which varies between 4000 and 2000, shows that a number of less than 10 back-drifting ions is already achieved. Further measurements with the nominal neon mixture and with foils with a lower optical transparency are in progress.

3.4.5.3 Gain Stability

Temporal variations of the gas gain of GEMs due to rate-dependent charge-up effects have been observed in laboratory tests on double mask GEMs [158]. This could imply a particular challenge for the gain calibration. Possibilities to keep this effect to a minimum include the use of cylindrical instead of double conical GEM holes, minimization of the insulating surface between the GEM hole and the metal layers as well as the use of water addition to the drift gas. Laboratory tests are in progress to confirm these findings under ALICE-specific conditions and to investigate the behavior of single mask foils.

3.4.6 Prototype Test

A full size prototype test based on a spare IROC of the ALICE TPC equipped with a triple GEM readout scheme is presently being prepared. The prototype will be operated near the beampipe in the ALICE cavern during the p-Pb running period in 2012. It will employ a small test field cage which is presently used for laboratory tests. The goal of this test is to prove the feasibility of GEM readout coupled to an existing TPC readout chamber body and its operational stability under LHC conditions. In addition, a precise measurement of the dE/dx capabilities will be carried out at the PS, using the same single-mask IROC prototype. In this test beam the prototype will be exposed to identified electrons and pions at momenta of 2-6 GeV/c and its response measured.

3.4.7 Front-End and Readout Electronics

The replacement of multi-wire proportional chambers with GEM detectors and the requirement of a continuous readout, call for the development of new front-end and readout electronics. In the current TPC the readout of the signals is done by 557568 pads that form the cathode plane of the multi-wire proportional chambers. The signals from the pads are passed to 4356 Front-end Cards (FECs), located 7 cm away from the pad plane. In the FECs a custom charge sensitive amplifier, named PASA (PreAmplifier ShAper), transforms the charge signal induced in the pads into a differential semi-Gaussian voltage signal that is fed to the input of the ALTRO (ALice Tpc Read Out) chip. Each ALTRO contains 16 channels that concurrently digitize and process the input signals. Upon arrival of a first-level trigger, the data stream corresponding to the detector drift time ($\leq 100 \mu\text{s}$) is stored in a memory. When a second-level trigger is received, the latest event data stream is either frozen in the data memory, until its complete readout takes place, or discarded. The readout can take place at any time at a speed of up to 200 MByte/s through a backplane bus linking the FECs to the Readout Control Unit (RCU), which interfaces them to the DAQ, the Trigger and the Detector Control System.

The current system is not suitable for the upgraded TPC for two main reasons. First the shaping-amplifier reads signals with a positive polarity, while its dynamic range for negative polarity, as those produced by GEM detectors, is limited to about 100 times the noise and, therefore, not adequate to accommodate the TPC signals (dynamic range 900:1). The second limitation is related to the readout scheme. The ALTRO chip does not support a continuous readout. It is either sampling the TPC signals and recording them in a local memory or transferring them off detector. The two operations cannot occur concurrently.

In the context of the R&D for the Linear Collider TPC, a new chip has been developed, named Super-ALTRO [159] that overcomes the first limitation. The Super-ALTRO is a further evolution of the ALTRO chip. It combines in a single silicon die the PASA and the ALTRO functionality and includes more processing functions. In particular the PASA is programmable in terms of gain and shaping time and can operate with both positive and negative polarities. The chip has been fabricated in a 130 nm CMOS technology. A series of 1000 circuits (16000 channels) is used to instrument a TPC prototype developed in the context of the EUDET project (Detector R&D towards the International Linear Collider). The block diagram and a picture of the chip are shown in Figure 3.15.

The new readout system for the upgraded TPC can be built on the basis of the current TPC readout and the Super-ALTRO chip. However, a new version of the Super-ALTRO chip, that supports a continuous readout scheme, is needed. In this concern we would like to mention that the continuous readout might become a limitation for the integration of the PASA and digital chip in the same substrate. Whether or not this is achievable should be verified with measurements done on the Super-ALTRO chip. Independently of this, we can reasonably assume that 32 channels can be integrated in a chip. About 16 chips (32 channels \times 16 = 512 channels) would be connected with a point-to-point serial link to an RCU. The latter would interface the FEE to the DAQ with a 10 Gbit/s link. The way chips are grouped in printed circuit boards depends on the physical layout of the pad plane and support structure.




Figure 3.15: Block diagram (left side) and picture (right side) of the Super-ALTRO prototype chip.

3.5 Upgrade of TOF Readout Electronics

The operation of the TOF detector at Pb–Pb collision rates of 50 kHz defines new running conditions in terms of charged particle flux on the Multigap Resistive Plate Chambers (MRPCs) as well as new requirements on data rates in the readout electronics.

During the 2011 heavy-ion run with a maximum instantaneous luminosity of $4.3 \times 10^{26} \text{ Hz/cm}^2$ (3.4 kHz interaction rate), a charged particle flux of 7 Hz/cm^2 was measured by the TOF. At an interaction rate of 50 kHz a flux of 100 Hz/cm^2 is therefore expected. Previous measurements [160, 161] of the intrinsic rate capability of the MRPCs showed that this rate can be easily sustained.

The current TOF readout can sustain readout rates of tens of kHz, so for the ALICE upgrade program the main aim will be to further increase the readout rate capability for both Pb–Pb and pp interactions. The current readout system is shown in Fig. 3.16. Each of the 18 super-modules is read out by 4 electronics crates. In each crate 1 DRM collects the data of 10 TRMs, each of them including 30 HPTDC chips. The pipelined internal architecture of the HPTDC chips coupled with readout buffers in the TRMs allows a zero deadtime readout.

The readout is done in three distinct phases:

- HPTDC readout: transfer of data from the HPTDC internal buffers to TRM memories
- VME readout: data transfer from the 10 TRM cards to the DRM via the VME bus
- DAQ readout: data transfer from the DRM to the DAQ via the DDL link

Upon L1 trigger arrival ($\approx 7 \mu\text{s}$ after collision), the readout of the HPTDC starts, and upon L2 arrival ($\approx 100 \mu\text{s}$ after collision) the TRMs are read out.




Figure 3.16: Current readout TOF scheme and its main components analysed in the text.

The TOF readout time depends on the multiplicity per event. The average hit multiplicity measured by the TOF for minimum bias Pb–Pb collisions at $\sqrt{s_{\text{NN}}}=2.76 \text{ TeV}$ is 1500 hits/event. Scaling this number to $\sqrt{s_{\text{NN}}}=5.5 \text{ TeV}$, we expect 2000 hits/event.

According to the specification manual [162], the HPTDC is able to handle a maximum input trigger rate of 500 kHz. The first limit on the TOF readout rate is the time needed to transfer the data from the HPTDCs to the TRMs. The transfer of an 'empty' event takes 2.25 μ s, which is determined by the communication protocol between TRM and HPTDCs, and the readout of the data adds 0.15 μ s assuming 2000 hits/event. This results in a limit on the readout rate for pp of about 400 kHz and even lower for Pb–Pb according to the multiplicity. Since the HPTDCs and TRMs are located inside the TOF supermodules in a region which is not easily accessible, any modification would require removal of the TOF from the ALICE detector and a partial disassembly of the modules, which is not feasible within the installation time available for the upgrade in 2018. Moreover, such an upgrade would go beyond the requirements defined by the ALICE upgrade strategy. The readout rate of 400 kHz is therefore the maximum that can be achieved by the TOF and the upgrade of the rest of the electronics chain is intended to eliminate all further restrictions on this rate.

The VME readout is currently limited to a throughput of 40 MByte/s. As shown in Figure 3.17 this already allows a maximum readout of about 80 kHz for the expected average Pb–Pb hit multiplicity of 2000. At this hit multiplicity the overhead of the data format amounts to 2/3 of the total event size, so the maximum readout rate for pp collisions, where the hit multiplicity is much lower, is limited to about 160 kHz. By upgrading the rmware it should be possible to reduce the volume of the readout control words as well as to increase the data transfer rate up to 160 or even 320 MByte/s. This would therefore remove all limits on the readout rate from the VME. Further studies are planned to prove this approach and the maximum capability of the HPTDC.

The DRM modules inside the TOF crates ship the data to the DAQ via the DDL link. These modules will be reengineered using a more performant FPGA chip, a new version of the DDL link and probably a new version of the TTCrx chip. With this upgrade the readout to the DAQ should also be able to cope with the maximum readout rate that can be handled by the HPTDC chip.

For a more detailed discussion of the rate limitations of the current TOF readout electronics and plans for its upgrade we refer to [163].




Figure 3.17: TOF readout rate as function of the VME bandwidth for four different running conditions as described in the legend.

3.6 Upgrade of TRD Readout Electronics

The TRD has originally been designed to cope with a Pb–Pb interaction rate of 10 kHz and a maximum track multiplicity of 8000 charged particles per rapidity unit, which is four times larger than the one actually observed, and to sustain a maximum readout rate of few hundred Hz. The performance of the TRD readout chain and its potential upgrade depends on two major functional units, the Front-End Electronics (FEE) and the Global Tracking Unit (GTU). The FEE is mounted on each TRD chamber inside the 18 supermodules and provides digitization as well as two stages of data processing and readout (see Figure 3.18) in a dedicated ASIC (TRAP). Each FEE event readout sequence includes about 4 μ s of detector drift time and processing time. The first stage is designed to have low latency, typically finishing within 5 μ s after the interaction, while its readout is limited to four 32-bit words per MCM (group of 18 pads). It currently provides tracklet and PID information for the TRD L1 trigger. The second stage has no limitations on the number of words transferred per event, but an overhead of 8 μ s due to

communication handshakes in the readout. It is currently used to readout full zero suppressed ADC data. The FEE provides no multi-event buffering and remains busy until the full processing and readout sequence is completed or the reject of a higher trigger level aborts the sequence. Readout of the FEE is done with 1044 optical links (2 per chamber) with an effective total data rate of 250 GByte/s.




Figure 3.18: Processing and readout sequence of the TRD FEE.

Data is transferred to the GTU, a set of electronics modules located in 3 racks inside the ALICE cavern. The GTU provides the TRD contribution to the L1 trigger based on online p_T reconstruction and electron identification. It also implements the interface to the DAQ system using one DDL per sector, providing a total effective data rate of 3.6 GByte/s for the readout of L2 triggered events. With a trigger mix of predominantly central triggers, for which each supermodule produces on average 370 kBByte/event, the system can handle a rate of accepted L2 triggers of 540 Hz.

The major limitation of the current GTU, which prevents from running at higher rates, is the readout bandwidth to the DAQ system. In order to eliminate the readout limitation set by the output bandwidth of the GTU system, a substantial hardware upgrade is required. In the following we discuss the maximum achievable readout rates assuming that the GTU system does not represent any bottleneck.

Three readout modes are considered.

- **ADC Raw Data Readout:** this readout mode uses the full FEE readout sequence and is the one currently used. Assuming minimum bias Pb–Pb events with an average TRD raw event size of 3 MByte/event results in an average readout time of $36.5 \mu s$. This leads to a maximum readout rate of 27.4 kHz at 100% dead time and a corresponding data throughput of 82 GByte/s.
- **Reduced Data Readout:** the processing capabilities of the TRAP device are used to read out only processed information such as tracklet fit parameters, charge sums and other PID-related variables, but no ADC raw data. This mode uses the second readout stage, which provides flexibility in the amount of the read out data and allows to optimize the reduced data set for offline performance. With a reduction factor of 4 a maximum readout rate of 51 kHz at 100% deadtime is possible with a corresponding data throughput of 34 GB/s.
- **Tracklet Readout:** a further significant rate increase can be achieved by omitting the second readout stage with its $8 \mu s$ overhead. Only the data words of the first readout stage are used to transfer tracklet information and to some extent PID information. Performance losses with this very restricted data set have to be evaluated carefully. For the tracklet readout mode about 25% of the available words are used for minimum bias Pb–Pb events. The readout time of $5 \mu s$ leads to a maximum event rate of 200 kHz and a data volume of 54 GByte/s. For a Pb–Pb interaction rate of 50 kHz the data volume is therefore 13.4 GByte/s. The FEE readout time is not dominated by data transfer, which accounts only for about 25% in case of minimum bias Pb–Pb data. The maximum rate will therefore not strongly depend on the data volume. Also for pp data taking the maximum readout rate will be mostly determined by the sum of the drift time and processing time.

An additional limitation to the maximum trigger rate originates from the power budget, since significant parts of the electronics of the detector are normally in a sleep mode and are activated upon trigger arrival. The TRD has been operated at trigger rates up to 80 kHz. For operation at higher rates, the actual limitation set by the LV system has to be verified.

3.7 Muon Spectrometer Upgrade

Quarkonium, open heavy flavour and vector meson production in heavy ion collisions at the LHC are the main physics topics addressed by the ALICE Muon Spectrometer. The spectrometer covers a pseudo-rapidity range of $2.5 < \eta < 4$ and is unique at the LHC to study charmonium down to zero transverse momentum and vector mesons at low transverse momentum. In the next years, J/ψ production will be studied in detail, in particular the nuclear modification factor and azimuthal asymmetry as a function of the rapidity and transverse momentum will be measured.

As described in Chapter 2, the 50 kHz Pb-Pb collision rate after LS2 will open the possibility to study the ψ' production with equivalent precision. In order to cope with the increased rate, an upgrade of the front-end and readout electronics for the muon detectors is needed.

The Muon Spectrometer is composed by two sub-systems: The Muon Tracking based on multi-wire proportional chambers with cathode pad readout, the so called Cathode Pad Chambers (CPCs), and the Muon Trigger based on Resistive Plate Chambers (RPCs). The upgrade of these systems is described in the following.

3.7.1 Muon Trigger Upgrade

The muon trigger system is based on four planes of RPC detectors, arranged in two stations one meter apart from each other. The coordinates measured in these two stations are combined to determine the muon p_T .

3.7.1.1 Expected Rates and Limitations

The expected counting rates for the RPC detectors are given in Table 3.5 for pp and Pb-Pb. These values are extrapolated from the present measurements [164]. The counting rate in pp does not account for the beam induced background. We see that mean counting rates of up to 50 Hz/cm² are expected.

2 MHz pp		50 kHz Pb-Pb	
Single counting rate mean - max	Mean Current per RPC	Single counting rate mean - max	Mean Current per RPC
50 - 150 Hz/cm ²	100 μ A	25 - 45 Hz/cm ²	50 μ A

Table 3.5: Expected counting rates of the RPCs for pp and Pb-Pb.

The expected muon trigger rates are given in Table 3.6 for pp and Pb-Pb, as extrapolated from the present muon trigger performance. The typical cocktail of triggers, which is presently used, includes DiMuon Unlike-Sign (US) and Like-Sign (LS) with $p_T > 1$ GeV/c, Single Muon with $p_T > 4$ GeV/c and a sample of Single Muon with $p_T > 1$ GeV/c. The maximum expected rate of this cocktail does not exceed 5 kHz for a Pb-Pb collision rate of 50 kHz.

	2 MHz pp		50 kHz Pb-Pb	
MTR p_T cut	4 GeV/c	1 GeV/c	4 GeV/c	1 GeV/c
Single Muon	800 Hz	7 kHz	1.5 kHz	10 kHz
DiMuon US + LS		400 Hz		3 kHz

Table 3.6: Expected muon trigger rates for pp and Pb-Pb

In view of the these rates, an upgrade of the RPC front-end and readout electronics is needed. The RPC detectors are currently operated in 'maxi avalanche' mode [165], with a front-end electronics (ADULT ASIC [166]) that does not include any amplification stage. In this mode, the mean charge is of the order of 100 pC. On the basis of the studies carried out in the past, this sets a limit on the instantaneous counting rate of <50 Hz/cm². This number is very close to the expected rates after LS2 and leaves no margin. Another limit comes from our R&D on RPC aging: the safe operation of the detectors cannot be guaranteed for an accumulated dose larger than 50 mC/cm² (500 Mhits/cm²). This condition translates into ten months of data taking (10^7 s) at 50 Hz/cm². These two limitations are not compatible with a safe operation of the RPCs in the conditions expected after the upgrade. One possible solution to overcome these limitations is the operation of the RPCs in 'genuine avalanche' mode, as

presently used by ATLAS and CMS. In this mode, the avalanche charge and the rate limitations are reduced by a factor of about 10. However, operating the RPCs in this ‘genuine avalanche’ mode requires a change of front-end electronics.

The second serious limitation of the present Muon Trigger system is the readout speed. The event readout time is 110 μ s. Considering the increased trigger rates for the upgrade, such a readout time would cause an unacceptable dead-time of more than 30 %. Hence, the readout speed must be increased which requires changing part of the readout electronics.

3.7.1.2 *Front-End Electronics Upgrade*

The RPCs are read out by segmented strips distributed across their entire surface. This choice was necessary to limit the tracking ambiguities in the high multiplicity environment of Pb-Pb collisions. In order to limit the cluster size, i.e. the number of strips fired for a single track, the RPCs were originally planned to be operated in streamer mode. The large signals in this operational mode allowed to have a front-end without any gain. In order to limit possible aging effects, the ‘maxi-avalanche’ mode was finally chosen for the operation in ALICE, which is still compatible with the present front-end electronics and provides the required performance for the current ALICE detector. However, the adoption of the ‘genuine avalanche’ mode for the upgrade requires an amplification stage in the front-end electronics.

As a first step it is foreseen to build a front-end card prototype, based on commercial fast amplifiers, to perform measurements of signal to noise level, fast and total charge, spatial and timing resolution and detection efficiency. In parallel we will develop an ASIC (called FEERIC, for Front-End Electronics Rapid Integrated Circuit) based on existing circuits [167, 168].

3.7.1.3 *Readout Electronics Upgrade*

The modification of a part of the readout electronics is necessary in order to reduce the readout time to less than 20 μ s per event. In the present scheme the data are stored, upon a L0 trigger, in the buffers of the decision electronics boards (234 Local, 16 Regional and 1 Global board). Upon L1 reception, the data are transferred to the buffers of the two DARC readout boards via optical fibers. The data of the 16 Regional crates, housing 1 Regional and 16 Local boards each, are transferred in parallel into the DARC. The transfer time to the DARC is dominated by the readout of the Local boards which takes 40 μ s.

Presently all Local boards are read out at each event, while in the worst case of Pb-Pb central collisions, the mean number of Local boards delivering a positive decision on Single Muon is of the order of 2 (out of 234 possible). Therefore one possibility of improvement is to read out only the Local boards which deliver a positive decision. This would reduce the transfer time to the DARC to few μ s and bring down the dead-time to less than 10 %. For this purpose, it is foreseen to completely change the Regional and DARC boards. New functionalities (e.g. multi-event buffering in the DARC) could also be added. In this scenario, the Local boards, which are the most numerous, would be left unchanged.

3.7.2 *Muon Tracking Upgrade*

The muon tracking system consists of ten planes of tracking chambers grouped in pairs into five stations. The CPCs are wire chambers operated with a gas mixture of 80 % Argon and 20 % CO₂. The chambers provide fine granularity to cope with the high multiplicity and to limit the occupancy to less than 5 %. Low material budget (about 3 % of X₀ per chamber) and good spatial resolution (about 100 μ m in the bending plane) are achieved with these detectors.

3.7.2.1 *Muon Tracking Performance and Limitations*

The design of the tracking system electronics was driven by two main requirements: to read 1.08 millions channels at a rate up to 1 kHz and to allow a spatial resolution of better than 100 μ m in each station. This electronics is divided in three parts: the front-end board called MANU (MANas NUmerica), the readout system called CROCUS (Cluster Read Out Concentrator Unit System) and the interface with the ALICE central trigger called TCI (Trigger Crocus Interface) [169].

The front-end MANU boards were built as thin as possible, to minimize the material budget, and were carefully designed to achieve a noise figure of 1000 electrons [170, 171]. Each board houses four 16-channel chips

called MANAS (Multiplexed ANAlogic Signal processor) including the following functionalities: charge amplifier, shaper and track&hold. The MANAS signals are fed into two 12-bit ADCs, which in turn are read out by the MARC (Muon Arm Readout Chip). This chip allows the communication with the CROCUS and performs zero suppression. About 17000 of these MANU cards are used to read the tracking system. The PATCH buses (Protocol for Alice Tracking Chamber) provide the connection between the MANU boards and the CROCUS crate.

Each muon station is read out by four CROCUS crates, leading to a total of 20 CROCUS crates. During the data acquisition phase, the main tasks of the CROCUS are to concentrate and to format the data from the chambers, to transfer them towards the DAQ and to dispatch the trigger signals. The CROCUS crates also control the front-end electronics and calibration processes. All these tasks are done through a DSP and FPGA farm. A CROCUS reads up to 50 PATCH buses. The present minimum read-out time is 200 μ s, which corresponds to a maximum readout rate of 1kHz at 20% dead-time.

The trigger signal, coming from the Central Trigger Processor (CTP) is distributed to the front-end through the CROCUS by the TCI. The main goals of the TCI are to decode the trigger signal, to generate the L1 reject and to manage the busy signals of all the CROCUS crates.

3.7.2.2 Front-End Electronics Upgrade

The goal of the upgrade consists in increasing the readout rate to a minimum of 5 kHz while keeping the dead-time below 10%. This means that the total readout time will have to be reduced below 20 μ s in case no multi-event buffering is implemented in the front-end electronics. This will anyway require a replacement of the MANU card. The baseline option is to develop a new MANU card with an architecture similar to the present one. The ADCs, which in the present version work at 1 MHz, and therefore take 32 μ s for the digitisation, will be replaced by 10 MHz ADCs. This will reduce the digitisation time to 3.2 μ s. It should be noted that also the MANAS and MARC chips will have to be replaced by new ASICs with the same or similar functionalities.

3.7.2.3 Read-Out Electronics Upgrade

In order to achieve the goal of a maximum readout time below 20 μ s, also the CROCUS readout system has to be replaced. Preliminary studies indicate that the upgrade requirements could be fulfilled by a new readout system based on state-of-the-art FPGAs. The design of the new readout system will benefit from the work which is being done for the development of a similar readout unit for another experiment. The general idea is to keep the same readout scheme and cabling infrastructure between the detectors and the CROCUS. Therefore, the new CROCUS will have also 5 readout boards (see Figure 3.19) based on a SPARTAN FPGA. These 5 boards will be read out by a concentrator board (CRT) based on a VIRTEX 6 FPGA. In order to de-randomize the data transferred to the DAQ, this system will implement a multi-event buffering able to store several central Pb-Pb collisions. The trigger signals will be still managed by the TCI. The new CROCUS system combined to the new MANU will allow readout times below 20 μ s.




Figure 3.19: New CROCUS design proposal.

3.8 Other Upgrade Items

Forward Trigger Detectors

The forward region of the present ALICE detector hosts the V0 and T0 detectors. The V0 detector is based on plastic scintillators and shows a time resolution of better than 1 ns. It provides a highly efficient interaction trigger, a selective centrality trigger for Pb-Pb collisions and it is also used for rejection of beam-gas interactions. The T0 detector is based on Quartz Cherenkov technology with a time resolution of better than 50 ps. It provides the t_0 for the TOF detector and selects the primary vertex position with centimeter precision at the trigger level. Which of these functionalities has to be preserved for the upgraded ALICE detector is subject to further studies. The goal is to have a single detector system that can provide all the required functionality.

The Zero Degree Calorimeters (ZDCs) are located in the LHC tunnel at a distance of 112.5 m from the ALICE interaction point. These detectors will be upgraded in order to provide all data for the 50 kHz Pb-Pb interaction rate upon an interaction trigger. Currently the readout rate is limited to 8 kHz. Whether the upgrade of the readout electronics is sufficient or whether the photon detectors or other detector elements also have to be changed, is subject to further studies.

PHOS Detector

PHOS (PHOton Spectrometer) is an electromagnetic calorimeter of high granularity consisting of 17280 detection channels of lead-tungstate crystals (PbWO_4) of $2.2 \times 2.2 \times 18 \text{ cm}^3$ dimensions, coupled to large-area Avalanche Photo Diodes (APDs) with low-noise preamplifiers. It covers approximately a quarter of a unit in pseudo-rapidity ($-0.12 \leq \eta \leq 0.12$) and 100° in azimuthal angle. Its total area is about 8 m^2 , and the total crystal volume is about 1.5 m^3 . The PHOS electromagnetic calorimeter is optimized for measuring photons (of about 0.5–10 GeV/c), π_0 (of about 1–10 GeV/c) and η mesons (of about 2–10 GeV/c).

Studies on an upgrade of the front-end electronics in order to improve the performance for the measurement of direct photons are ongoing. The energy of these photons is expected to be below 5 GeV. In order to reduce the background from anti-neutrons in this energy range, an improved time resolution of 1 ns should allow a rejection by a Time of Flight measurement. The physics performance of such an upgrade as well as the need for an increase of the PHOS readout rate capability are currently under study and a decision on the upgrade plans will be taken at a later date.

EMCal Detector

The chosen technology is a layered Pb-scintillator sampling calorimeter with a longitudinal pitch of 1.44 mm Pb and 1.76 mm scintillator with longitudinal wavelength shifting light collection. The full detector spans $\eta = -0.7$ to $\eta = 0.7$ with an azimuthal acceptance of 110° . The detector is segmented into 12672 towers, each of which is approximately projective in η and ϕ to the interaction vertex. An extension of this EMCal (so called DCal) will increase azimuthal coverage by 70° and will be installed during LS1. Also the readout electronics will be upgraded. The current performance with new readout in an ALICE DAQ test stand has shown a top readout rate of about 25 kHz for Pb-Pb event sizes. With additional firmware modifications (reading Low Gain channels only when the High Gain channel has saturated) we expect to double the readout rate. With multi-event buffering, EMCal should be able to run at rates up to about 50 kHz with reasonable busy time. Therefore, the EMCal can be read out either upon the minimum bias trigger, up to 50 kHz, or on its own trigger.

HMPID Detector

The HMPID is a proximity focusing RICH detector based on MWPCs equipped with pad segmented photocathodes, activated with CsI. It uses a liquid C6F14 as Cherenkov radiator. The detector has an active area of 11 m^2 , it covers one unit in pseudorapidity ($|\eta| < 0.5$) and 5% of TPC acceptance.

With 3σ separation the HMPID can identify protons and kaons in the momentum range $1.5 \text{ GeV}/c < p < 3 \text{ GeV}/c$, and protons in the range $1.5 \text{ GeV}/c < p < 5 \text{ GeV}/c$.

Its excellent PID capabilities can be exploited for physics and to constrain the charged hadron identification by the dE/dx measurement of the TPC in the overlapping momentum range.

At present, the maximum readout rate is 2.5 kHz. It is possible to increase this rate significantly by parallelizing the readout cycles and replacing the DDL and RORC cards with those of the new generation (DDL3 and RORC3). De-

pending on the maximum achievable rate, the HMPID readout will be downscaled accordingly and the information will be combined to the one of other central detectors at the event building stage.

Chapter 4

Data Collection and Processing

4.1 Introduction

As discussed in the previous chapters, the physics programme addressed by the ALICE upgrade requires a large number of events to be recorded. The rate for heavy-ion events handled by the online systems up to permanent data storage should be increased up to 50 kHz corresponding to roughly two orders of magnitude, in respect to the present system. The raw data will be processed and compressed by the online systems but no event will be discarded. The online data reconstruction and compression will therefore be instrumental in keeping the total data throughput within an envelope compatible with the available online and offline resources. One of the roles of the online systems will be to perform many parts of what is commonly called the offline processing. This integration will require a common software framework and this requirement motivates the common chapter of the Letter of Intent dedicated to both data collection and processing.

This chapter includes a description of the upgrade of the data collection and processing systems of ALICE and is organized as follows. The Section 4.2 lists the requirements in terms of event rates, event sizes and data throughput that the systems will have to handle. It also describes precisely the steps of data compression that will be applied to the data to reduce the input peak data throughput of 1 TByte/s to an average recorded data output of 12 GByte/s. It finally gives an estimate of the computing power needed to perform the offline data processing. The architecture of the data collection is described in section 4.3. The present and future research and development are discussed in Section 4.4.

4.2 Requirements

For many of the discussed physics observables, no satisfactory event selection by a hardware trigger system is possible in heavy-ion collisions. This therefore sets the new main requirement for the ALICE online systems: a dramatic increase of the rate at which events will be inspected, compressed and recorded.

The general strategy is to read out the ITS, TPC, TRD and TOF detectors at the foreseen interaction rate of 50 kHz and ship this related data to the online systems. The baseline scenario for the TRD is to read out the tracklet data at 50 kHz. However, the online system has been dimensioned to support the fallback solution of reading out the full ADC data at a reduced rate which imply however a larger data throughput. The EMCAL and Muon system will be read out upon their own trigger at a lower rate.

The online systems should be designed to keep the nominal performance, even in case of noise or background larger than anticipated, and scale to twice this performance in case of higher interaction rates.

4.2.1 Event Rates, Event Sizes and Data Throughput

The architecture of the online systems should allow an event rate of 50 kHz, corresponding to the top instantaneous luminosity. It should also be able to scale by a factor two to deal with an increased event rate if needed. In the following we will refer to a maximum rate of 50 kHz. The average rate during a fill will be of the order of 20 kHz. Considering the LHC duty factor of 40 % the effective data taking time during one heavy-ion period of one month

corresponds to 1×10^6 s. This will result in a total of 2×10^{10} minimum-bias events (MB) recorded during the yearly heavy-ion run.

The Table 4.1 presents the event size of all detectors as being sent to the online systems (after zero suppression) and as being recorded (after data compression by online reconstruction). The same table summarizes the estimated overall data throughput at the input of the online systems and the peak and average rates to mass storage. The peak rate has to be sustained in the whole online systems up to the local data storage whereas the average rate has to be sustained all the way to the data recording in the CERN computing center. In the latter case, the data storage at the experimental area is providing the buffer required, and is making use of the 40% efficiency mentioned previously to reduce the data rate to the computing centre.

Table 4.1: Expected event sizes and data rates.

Detector	Event Size (MByte)	Input to online systems (GByte/s)	Peak output to local data storage (GByte/s)	Average output to computing center (GByte/s)
	After zero suppression	After data compression		
TPC	20.0	1.0	1000	50.0
TRD	1.6	0.2	81.5	10.0
ITS	0.8	0.2	40	10.0
Others	0.5	0.25	25	12.5
Total	22.9	1.65	1146.5	82.5
				13.2

With the estimated event size of roughly 23 MByte, this would result into a global data throughput of approximately 1150 GByte/s for the online systems and 460 GByte/s to the mass storage. A further data reduction is therefore a key element of the overall upgrade feasibility and will be discussed in more detail in the next section.

4.2.2 Online Data Reduction

The first steps towards an online data reduction have already been taken in recent years, with the current ALICE High Level Trigger (HLT). We will report on these briefly here, as they form the basis for the much more advanced online data reduction which is mandatory for the upgrade.

The HLT receives the full detector raw data in the current system. In order to reduce the data rate from the detector read-out bandwidth to the ALICE Data Acquisition system (DAQ) bandwidth, two possibilities were foreseen in the original design. The first is the use as a trigger, i.e. the selection of events or regions of interest. The second possibility is the replacement of some of the raw data with results of the online reconstruction. The latter has the advantage of increasing the statistics for all physics analysis, but does however set more stringent requirements on the quality of the online reconstruction. As can be seen from Table 4.1, the TPC is the detector with the biggest data volume in ALICE and has been the main focus of the online data reduction.

The TPC Front-End Electronics (FEE) performs the first step of data reduction with zero suppression, which reduces the data volume to about 20 MByte for minimum bias Pb–Pb events. The raw data can not be suppressed much further, an optimized entropy encoding with Huffman code allows for a further factor 2 of data reduction while simple run-length encoding techniques result in even smaller factors. Any data compression above this factor has thus to be based on a (partial) event reconstruction.

The first step of the TPC reconstruction is the cluster finding, i.e. the identification of charge clusters in the TPC drift volume. Each cluster is characterized by its position, width and charge. A cluster finding algorithm has been implemented on the Field-Programmable Gate Arrays (FPGAs) of the HLT Read-Out Receiver Cards (H-RORCs). Detailed performance studies before the 2011 heavy-ion run showed a performance comparable to the offline implementation. During the 2011 heavy-ion run, the TPC raw data was therefore discarded by the DAQ, except for a small fraction for QA purposes. This raw data was then replaced by the output of the HLT reconstruction. HLT clusters are now routinely used as input for the TPC offline reconstruction.

The data volume is minimized by optimizing the size of each cluster property to the one required by the intrinsic detector resolution. This is followed by the application of standard, loss-less, entropy encoding with a Huffman code. By using additional information, some of the cluster variables are transformed to reduce the entropy of the resulting distribution. The best example uses the fact that the cluster finding is done pad-row by pad-row, i.e.

the difference in pad-row between two following clusters is zero or one in most cases. By storing the pad-row difference instead of the pad-row number directly, the entropy of the distribution is reduced from about 6 to 1.1. Overall this allows for a data reduction by a factor 5 to 7. This has been demonstrated to work well as it has been the running mode of ALICE during the 2011 heavy-ion run and 2012 pp run.

For the upgrade a further data reduction factor is needed. The HLT also performs a complete track reconstruction of the event, the details of which are discussed in section 4.4.6. After the track reconstruction not all clusters in the TPC are associated to tracks. Studies with both the offline and online reconstruction show that half of the clusters are created from primary and secondary particles of the event, while the other half is coming from other sources. Storing only the clusters belonging to the tracks would increase the overall data reduction factor already above 10.

It should be noted that the online track reconstruction will only be used for the selection of the charge clusters to be stored on tape. Therefore, the requirements on the online calibration of the TPC drift field and velocity are mild and well within the accuracy presently achieved.

The reconstructed tracks and the cluster-track association also allow for further data reduction. The offline TPC reconstruction uses several of the cluster parameters to select the clusters added to a track. Moving some or all of these selection criteria to the HLT would reduce the number of clusters. Knowing which clusters belong to a track also enables more advanced transformation schemes to optimize the parameter distributions for entropy encoding, as well as the possible replacement of some of the individual cluster parameters by track-based properties. We estimate the further reduction potential by these steps to be on the order of two.

Table 4.2 summarizes the data reduction steps for the TPC. Applying all the steps just as described reduces the data volume by a factor greater than 20, shrinking it to less than 1 MByte per Pb–Pb minimum bias event. Based on the experience gained with the TPC, similar studies have started for the other detectors in ALICE. Preliminary results have been used to calculate the event sizes presented in Table 4.1.

Table 4.2: TPC Event size and data reduction factors performed in the FEE and the HLT.

Data Format	Data reduction factor	Event size (MByte)
Raw data	1	700
Zero Suppression (FEE)	35	20
Clustering (HLT)	5-7	~3
Remove clusters not associated to relevant tracks (HLT)	2	1.5
Data format optimization (HLT)	2-3	< 1

Based on online detector calibration and full event reconstruction, a more aggressive scenario could be envisaged. Most of the parts of what is traditionally called offline processing would then be done online. This would offer the exciting possibility to have as output of the online systems a dataset which is comparable to the output of the traditional first physics reconstruction pass of the data. This would only require the addition of a modest amount of software and computing resources to the online systems. The subsequent offline processing would be reduced to the creation of the analysis datasets, which would then significantly reduce the time between data-taking and physics analysis as compared to the present system. However, in this document we will not elaborate further on this scenario.

4.2.3 Online Data Processing Requirements

We base our estimate of the required computing resources on the experience gained with the current HLT and offline systems. The computing cluster of the present HLT has more than 2500 cores distributed over 200 nodes, supplemented by additional infrastructure nodes. In addition to Central Processing Unit (CPU) cores, the available processing capacity has been significantly enhanced by special processing devices. As described above, the TPC cluster finding runs on the H-RORC FPGAs, using in total 108 Xilinx Virtex 4 LX40. Each of the FPGAs is able to handle the full data rate of two incoming Detector Data Link (DDLs), each with a nominal bandwidth of 200 MByte/s. Typically two H-RORCs are installed in a node. In comparison, the same bandwidth would require on the order of 40 CPU cores/DDL when using the software implementation of the algorithm, which would then require additional 8000 cores in the HLT: i.e. a factor 3-4 more than currently installed.

Large parts of the TPC track reconstruction run on Graphical Processor Units (GPUs), with currently 32 NVIDIA GTX480 and 32 NVIDIA GTX580 installed. The use of these GPUs reduces the number of nodes needed by a factor 2 to 3. All nodes are connected by an Infiniband QDR network in a fat-tree layout. With this configuration, the HLT is able to handle the maximum rate of the current TPC read-out electronics of up to 520 Hz minimum bias and 310 Hz central Pb–Pb events.

The data rates after the upgrade are expected to be two orders of magnitude higher. For the FPGA based cluster finder the development of the C-RORC, described in section 4.4.1, demonstrates that the required performance is within reach with state-of-the-art FPGAs. To estimate the future processing power we extrapolate the current number of cores/GPUs in the HLT system to the new data rates: 250,000 cores and 6,400 GPUs would be needed with todays technology and linear scaling of the current HLT computing requirements with the event rate. Considering the evolution of computing equipment since several decades and following Moore’s law [172–174] which predicts a doubling of the number of transistors every two years, D. House predicted a doubling of the computing performance every 18 months. This results in an expected increase of the performance of the computing systems by a factor ($2^{(6 \times 12)/18} = 16$) in the six years until the upgrade. However, one has to note that the single core performance does not increase anymore, but rather the number of cores per CPU. The electronics chip makers and computer manufacturers anticipate 100 cores/CPU by 2018.

Assuming a similar performance gain for GPUs and dual-CPU nodes, we estimate about 1,250 nodes for the full online processing. This estimate is based on predictions provided by industry which will evolve with time. It includes thus an inaccuracy of a factor roughly equal to two. The real evolution of new computing products during the coming years will be closely monitored and the R&D will establish the actual benefit of this evolution for our application. Other issues that will be addressed by the R&D are: the future evolution of the memory and I/O bandwidth, the data storage capacity and access speed and their adequacy with the increase of computing power.

We note that to reach the target event rate, the performance of the current HLT and DAQ data transport frameworks has to be increased from the current maximum rate of 2 kHz in the HLT to 50 kHz and above; therefore a linear scaling of the required processing power has been assumed. Pre-processing of the data, for example in the case of the TPC cluster finding, was assumed in the calculation of the number of nodes. This will require additional processing power, which is expected to be provided by the FLPs as the first layer of nodes in the system. These FLPs are introduced in section 4.3.2. The currently favoured implementation will follow the HLT approach and uses FPGAs on the read-out receiver boards. The CPU cores and possibly additional computing devices such as GPUs on these systems might also be used for first computing tasks like applying local corrections. The number of nodes also relies on an efficient use of the available cores. Based on the experience gained with the HLT tracker, this will require a redesign of significant parts of the reconstruction software, to stay within a reasonable latency for the online systems.

The numbers presented so far are based on a traditional event based data reconstruction, which would rely on collision timing information provided by trigger detectors. However at 50 kHz interaction rate, the average number of events visible in the TPC drift time of up to 100 μ s will be close to 5, increasing the data volume per TPC event by a of factor 2 to 4, which would require significant additional computing resources. To avoid this additional overhead, we are considering moving to a time window based reconstruction strategy. For example, the data continuously read out by the TPC in a time window much larger than the drift time would be fully reconstructed by the online systems. The splitting of the data into events would only be done at this stage, also utilizing the timing and trigger information. The subsequent offline reconstruction steps as well as the physics analysis could then be based on the traditional event based approach.

4.2.4 Offline Data Processing and Storage Requirements

The requirement to run at a peak frequency of 50 kHz translates into an average frequency of ~ 20 kHz during a complete fill of the LHC. One month of Pb–Pb running will result in the acquisition of $\sim 2 \times 10^{10}$ events, which is two orders of magnitude more than what we have collected in our 2011 Pb–Pb run. This has been reconstructed in two months with 10^4 cores. This means that, in order to reconstruct the data produced by one Pb–Pb run after the upgrade within two months, $\sim 10^6$ cores of current performance would be needed. It seems reasonable to suggest that in the six years from now till the upgrade, the performance of the computing systems will increase by a factor of about 16 as shown in the previous section. We will have to increase the performance of our code of at least a factor of 6 in order to cope with the requirements. We believe that this additional improvement can come from an optimization of the current code, leading to a better exploitation of the modern computing architectures and possibly from the usage of General-Purpose GPUs (GPGPUs) and similar devices.

The number of simulated events needed for physics analysis is on the order of one tenth of the data events. The CPU time needed for one simulated event is about one order of magnitude larger than for a real data event, the CPU needs for the simulation are therefore estimated to be equivalent to the reconstruction of all HI data. This estimate is based on the R&D towards a faster simulation framework (see section 4.4.5.2) exploiting parallelization and GPUs for improved transport.

The rewriting of a large code such as AliRoot is certainly a major task. However there is already a substantial amount of experience within the collaboration in dealing with parallel systems. We will describe below the experience we have with parallelism at different levels.

The current estimate of the total storage needs over the complete data taking period (1 run pp, 3 runs Pb-Pb, 1 run p-Pb) include:

- The simulated and reconstructed data on disk amounts to 110 PB out of which 70 PB are for simulated data.
- The raw data of real events on tape for a total of 330 PB.

The evolution of the offline data storage needs on tape is an increase of roughly one order of magnitude compared to the present needs. This corresponds to a yearly increase of 50 needs on disk will increase by a factor 3, well below the evolution of the capacity of commodity hard disks drives since 30 years. In summary, the ALICE needs will evolve at a pace equivalent or lower than most of the common applications and of the cheapest available technologies.

4.3 Architecture

4.3.1 Fast Trigger Processor

The online architecture of the ALICE upgrade will include a Fast Trigger Processor (FTP) which will generate two triggers levels (L0 and L1). The FTP will accommodate both the continuous readout (used by the ITS and the TPC) and the triggered readout (used by the other detectors) for MB rates of up to a few hundreds of kHz in pp and 100 kHz in PbPb. A detailed review of the needs has been started within the collaboration.

4.3.2 Dataflow

The present architecture of the DAQ and HLT systems is based on a symmetric interface to the detector read-out [175]. A copy of the same data is transferred to both systems. This architecture has proven to be effective and flexible by providing several HLT modes, reflecting an increasing usage of the HLT from the commissioning up to the operation periods. This architecture could be reproduced for the upgrade but at the cost of an input bandwidth into the systems, which would correspond to twice the detector data throughput. Given the very high bandwidth considered for the upgrade, another architecture consisting of a common computer farm shared by the DAQ and HLT systems will be adopted.

The online architecture of the ALICE upgrade is based on a combination of continuous and triggered read-out:

- Two detectors (ITS and TPC) will implement a continuous read-out. It is the most efficient way to transfer the data to the online systems when the detector integration time is larger than the average inter-events arrival time. For example, at the maximum rate of 50 kHz, there will be close to 5 events simultaneously visible in the TPC. However, the LHC clock and the L0 trigger will be distributed to the ITS and TPC electronics for test purposes, data taking at lower rate and in order to tag the data with a common time reference.
- The TOF and TRD detectors will use a triggered read-out using the L0 trigger.
- A delayed trigger will be available for slower detectors.

The FTP will decide for each interaction the set of detectors being read out, in order to accommodate the detectors being operated at a lower rate (EMCAL, Muon system).

The architecture of the data collection and processing system is based on a pipeline that starts with the data collection on the detector and ends with the recording of compressed data on the local storage system at the experimental area.

As indicated in Table 4.1, the detector read out at 50 kHz requires a bandwidth of 1.15 TByte/s or 9 Tbit/s. The detector read-out is the only part of the system for which it is mandatory to deploy the capability to handle the 100 kHz from the beginning. In order to take into account the construction constraints of detector read-out and provide some headroom for the link protocol, the total capacity is estimated to be of the order of 25 Tbit/s corresponding to 2500 detector links at 10 Gbit/s.

This DDL will be the 3rd generation (DDL3) of the present optical link used in ALICE. It transfers data from the detector read-out electronics to a PC adapter which will also be the 3rd generation of Read-Out Receiver Card (RORC3) used by ALICE. The RORC3 will compress the data on the fly in the FPGA that is used to ship the data from the detector to the PC memory. The data traffic is therefore reduced to the strict minimum.

The data acquisition and processing will then be carried out by a large processor farm, based on CPU and GPU. The compressed detector data are transferred by the RORC3 into the memory of a *First Level Processor* (FLP), which might then also perform further localized reconstruction steps. The full event is then assembled and reconstructed in one *Event-building and Processing Node* (EPN) performing the final data compression and the data recording on the local data storage. The FLPs and EPNs are connected by the farm network while the EPNs access the local data storage through the storage network. These two networks are functionally separate but could be implemented by the same physical devices: this is the option which has been used later in this document.

The read-out architecture described above requires a profound redesign of the DAQ and HLT systems. The overall new architecture is presented in Figure 4.1. The computing farm will implement buffering and processing ca-




Figure 4.1: The upgraded DAQ and HLT systems.

pacities matching the nominal ALICE needs while being fully deployed. However, it has to handle unforeseen overloading or operate at reduced performance in case of a staged deployment. The latter is likely to happen during the commissioning phase given that, due to the large size of the full DAQ and HLT systems, there is a clear advantage of purchasing computing equipment as late as possible. The proposed architecture does not foresee any kind of back-pressure for the detectors being read out continuously. In the case of overloading, the data will therefore be discarded by the FLPs on the basis of its internal buffer usage. The FLPs buffer could become full, for example if there is a sudden input overload or because only a fraction of the EPNs is installed. If such a condition

is met, old data would be discarded and the buffering be made available for the new data transferred continuously from the detectors.

4.3.3 Farm Layout

The online computing fabric for the upgraded ALICE is an integrated farm performing the detector read-out, the data buffering, the event-building and the data processing. The event-building and the data processing are performed in two steps: first for the data of parts of each detector and then for the data of the complete event. The online computing fabric will consist of approximately 1,500 computing nodes of 2 different categories, corresponding to these two steps.

The 250 FLPs have four different functions with specific requirements:

- The FLPs will interface the online fabric to the detectors. The FLPs will receive the data from the detectors over up to 12 optical links at 10 Gbit/s for an aggregate bandwidth of up to 120 Gbit/s. In order to limit the footprint as much as possible, the goal is to concentrate all optical links on a single card requiring a single Peripheral Component Interconnect Express (PCIe) slot. The bandwidth needed requires this slot to be at least a PCIe Gen 3 16 lanes.
- The data will be processed on the fly by the RORC3 before being transferred to the FLP memory. This first data processing step includes the clusterisation.
- The data from the different DDL3s will then be combined to form a common event (for FLPs interfacing triggered detectors) or a common time window (for detectors with a continuous read-out). The FLPs will buffer the data until they have been sent to an EPN. The memory needed to buffer all the raw data for a period of 1 s corresponds to 15 GByte of RAM per FLP. The data from all FLPs related to a given event or time window will be sent to the same EPN.
- The FLPs will send the data related to events or time window to the EPNs. The sustained output bandwidth of each FLP can be up to 12 Gbit/s and can therefore be handled by one port of a commercial network at 40 Gbit/s or two ports at 10 Gbit/s. The processing power of the FLPs will be based on the combination of the FPGA of the optical link interface and its own CPUs and could possibly be enlarged by adding further additional processing devices as GPUs if needed. The number of FLPs is mainly fixed by the input I/O bandwidth needed and the number of I/O slots needed to interface to the 2500 optical links.

The 1250 EPNs have three main functions:

- The second step of data processing is the event reconstruction, which will allow to assign each cluster to a track and finally to an interaction. This step has therefore to be performed before the event building for data produced by detectors read out continuously.
- The event building can then be performed: it will associate all the data and clusters relative to the same primary interaction using trigger and timing information.
- The possibility of tagging events according to physics signatures.

The EPNs have a network data traffic (input bandwidth of 7.2 Gbit/s and output bandwidth of 200 Mbit/s) which can be handled by a single port of a commercial network at 10 Gbit/s. The EPNs are equipped with GPUs. The number of EPNs and their processing power have been estimated by scaling the present HLT processing power to the future needs, and taking into account the expected performance gain of CPUs and GPUs over the next few years.

The present ALICE data storage is dimensioned to support an internal peak bandwidth of up to 8 GByte/s and a sustained migration bandwidth to the computing centre of 4 Gbyte/s. Already today it would be technically feasible to increase these bandwidths to 20 GByte/s. We also note that, presently, the CERN computing centre is using one third of its global bandwidth to tape to handle the ALICE data traffic. The proposed ALICE upgrade is designed to lead to a peak output to local data storage of 82.5 GByte/s and an average output to the computing center of 13.2 GB/s. The rapid evolution of the data storage technology sustained by the ever increasing needs of both the commodity and commercial markets will allow to implement the required upgrade of the data storage with commercial components.

4.3.4 Platform Independence

While computer speed doubles every 18 months, the increase of clock frequency has plateaued since few years and the increase in performance will mostly come by an enhanced processing parallelism at several levels. This argument will be described in detail in Section 4.4.5. Moreover, after two decades of domination by a single hardware architecture (Intel X86), the computing market is evolving toward different computing platforms.

In order to exploit the parallel hardware being offered by vendors, the code has to be adapted to the hardware to a large extent. This is even more true if we consider devices such as the Intel Many Integrated Core (MIC) [176] recently commercialized under the name of Xeon Phi, the different flavors of GPGPUs (see for instance [177]) or even more special purpose devices such as the Texas Instruments DaVinci [178] DSP, which is being also considered for general purpose programming. While on one side we cannot ignore the potential performance improvements offered by these devices, not to mention the reduced energy and space footprint, harnessing the performance gain will require to "program around the hardware". Platform dependent software makes problematic the migration to another device, or even to the next generation of the same device. The lack of one common parallel programming standard makes this problem even more acute.

We believe that, at a coarse parallelism level, a generic thread parallel programming model (using pthreads or ROOT threads) can be applied in order to divide the work amongst the different cores. Coarser levels of parallelism can be handled via interprocess communication standards such as the Message Passing Interface (MPI). At this level the code can be maintained reasonably portable.

Our preliminary investigations of finer levels of parallelism indicate that it will be necessary to maintain different versions of the computational kernels. This may even require implementing different algorithms, for example one for CPUs and one for GPGPUs. We will investigate emerging standards such as OpenMP [179] and OpenACC [180], but we believe that these will only partially alleviate the need to develop different versions of the code. While this will certainly constitute an overhead, the cross-testing of the different implementations will be a very stringent quality check for our code. It will ensure that whatever the evolution of the hardware will be, we will always have a version that is "close enough" to the new hardware to make porting feasible. First experience has been gained with the current HLT cluster finder and tracking algorithms. The effective use of additional computing devices required new reconstruction algorithms. These algorithms had to be implemented at the same time as CPU versions for use during the offline simulation and subsequent reconstruction. A very good consistency between the different versions had been achieved through this.

4.3.5 Detector Control System

The present Detector Control System (DCS) of ALICE will also need to be upgraded. The DCS will have to integrate new detectors and interfaces to new classes of equipment. New electronics that will be adopted by the detectors for the upgrade will require the development and implementation of new interface standards. The current plan is to keep the same SCADA (supervisory control and data acquisition) system and migrate to its next generation, the WinCC Open Architecture. Also all operating systems will be upgraded and the cluster will be upgraded. The benefits of this migration will be the following:

- Improved performance due to parallel processing and load balancing.
- Enhanced availability through a better fault tolerance.
- Object-oriented databases and graphics
- Easier import and export of data

This migration will require a substantial adaptation or redesign of many software components. At the same time new operational tools will be developed and, in order to further improve efficiency, the DCS will be prepared for fully automated operation and error recovery. The DCS hardware will also require an adaptation to more recent standards for the computing and PLC systems.

4.3.6 Infrastructure

The infrastructure (counting rooms, racks, cooling, power, etc) needed for the upgraded online and DCS systems will reuse as much as possible of the present infrastructure used by the DAQ, HLT and DCS systems after some modifications performed during the LS2. The computing equipment used for the upgraded online system should

be commercial off-the-shelf equipment. The present strategy for the cooling of this equipment is therefore to use conventional techniques as presently used in the ALICE online and DCS systems. It consists of racks with forced air flow racks equipped with water-cooled doors. Given the important increase of computing capacity, the thermal capacity presently installed will most probably be insufficient and will have to be adapted. The cost of adapting the cooling capacity and other infrastructure items such as the power distribution is foreseen in the online systems budget.

4.4 Research, Prototyping and Development

4.4.1 DDL and RORC

The development of a new Common Read-Out Receiver Card (C-RORC also known as RORC2) has been started as a common R&D project of Data Acquisition (DAQ) and High-Level Trigger (HLT). This project aims to replace the current RORCs with an up-to-date hardware architecture with increased number of links, link speed and data processing capabilities while maintaining compatibility with the current read-out chain. The update of the RORCs is crucial especially for HLT operation because the current HLT RORCs use PCI-X as the interface to the host machines, which is no longer supported due to recent server hardware changes. The C-RORC is also one intermediate step towards the final RORC and DDL used for the ALICE upgrade scheduled for the LS2.

The new C-RORC is equipped with a Xilinx Virtex-6 FPGA, an eight lane PCI-Express Gen2 interface to the host machine and twelve optical links of the second generation (DDL2) realised with three Quad Small Form-factor Pluggable (QSFP) transceivers. Twelve links have been chosen in order to cover a complete TPC segment consisting of six DDL links with a single board, while still being able to send a copy of the data to the High-Level Trigger. The serial transceivers in the FPGA allow the optical links to be operated with the current read-out link speed while enabling upgrades of the link rates up to 6.6 Gbit/s. A sketch of the C-RORC hardware is shown in Figure 4.2. First prototypes of the C-RORC board are scheduled for 2012. While the board is in production, the FPGA firmware development has been started with evaluation boards with comparable hardware but reduced link densities. This prototyping platform consists of a Virtex-6 based PCIe board, with an add-on board for the optical connectivity with QSFP sockets. Evaluation of the optical parts with different QSFP transceivers and cabling solutions are ongoing. The firmware development for this new RORC will include functional modules of




Figure 4.2: C-RORC schematic layout

the present one. This is currently being ported to the new hardware, including the TPC cluster finding algorithm.

The throughput performance of the PCIe link has been tested with different host architectures and DMA channel configurations. Figure 4.3 shows a plot of the performance test results for single channel DMA operation from the FPGA into host memory for two successive Intel host architectures (the so-called Nehalem and Sandy Bridge). The software side on the host machine consists of custom kernel module with scatter-gather based memory allocation for DMA operation. The measured transfer rates fulfill all expectations and comply with the requirements for a twelve channel read-out. Once all the key issues such as data transfer and processing power of the C-RORC prototype will have been fully demonstrated, the RORC3 will be developed based on the experience of the C-RORC hardware and of the present RORCs firmware.




Figure 4.3: DMA performance test results of one PCIe Gen2 x8 slot used by the C-RORC. There is a substantial performance gain over the whole range of block size from the Nehalem to the Sandy Bridge for a single channel read-out. The Sandy Bridge also improves substantially the performance for smaller block sizes compared to the Nehalem with 12 channels.

4.4.2 High-Level Programming of FPGAs

However the issue of programming FPGA with complex programming languages still remains. FPGAs have a long history in data processing for experiments in HEP covering the handling of low level protocols, up to online processing and first event building tasks (e.g. clustering). Every new FPGA generation comes with an increased device size, and as a consequence a larger number of more complex algorithms can be implemented in hardware. Examples include already existing hardware cluster finders and algorithms for future data compression. Until now, these algorithms are described using low level hardware description languages like Very-high-speed integrated circuits Hardware Description Language (VHDL) or Verilog. Low level languages have been proven to be well suited for the description of interface blocks like PCIe, DRAM controllers or serial optical links. However, development is expensive in these languages for processing algorithms based on dataflows. Complex pipeline architectures with hundreds of stages easily lead to code that is hard to read, while optimized processing modules with differing latencies cannot be integrated easily. Maintenance and modification of this kind of code is a complex task. In recent years frameworks for pipeline generation have become available, which generate an optimized pipeline architecture from a high-level dataflow description and simplify development and maintenance of the code. Furthermore, several examples have shown that code generated from a high-level framework produces better results than hand-written code in a low-level language. Utilization of these techniques in HEP can reduce the design effort of developing firmware dramatically while producing more efficient hardware. Our aim for the next major upgrade period is to investigate the benefits for FPGA firmware when using a higher level framework and compare the results to manually-written hardware descriptions.

4.4.3 FLP and EPN

The nodes of the online data collection and computing system proposed for the ALICE upgrade will require a substantial input-output bandwidth and processing power. We review here the foreseen needs and discuss how they will be addressed by the present and future evolution of the industry.

Some of the nodes will require an input-output bandwidth of more than 100 Gbit/s, such as the FLPs for the TPC. The PCIe is the current de-facto standard as computer expansion bus. Since its introduction in 2003, several generations have been released and used by all the large computer manufacturers. The main characteristics of all the generations are shown in Table 4.3.

The present PCs are equipped with PCIe Gen 2. The Figure 4.3 indicates that the actual performance reached by this generation of machines is 3.2 GByte/s for a 8 lanes slot for a theoretical maximum of 4 GByte/s. The first computing products equipped with PCIe Gen 3 have been released in the 2nd quarter of 2012. The PCIe Gen 4 has


Table 4.3: Characteristics of the successive releases of the PCI Express bus.

PCI Express Version	Lane Bandwidth MByte/s	Total Bandwidth 8/16 lane slot (GByte/s)	Date of introduction into commercial products
Gen 1	250	2/4	2003
Gen 2	500	4/8	2006
Gen 3	1000	8/16	2012
Gen 4	2000	16/32	2016 (expected)

been released by the standardisation body in November 2011 and is expected to equip the PC from 2016. These options (PCIe Gen 3 and 4) have therefore been used for the dimensioning of the system.

As described above, all the nodes on the farm, both FLPs and EPNs, will have a substantial processing load. A first proposal of processing unit is associated to each type of processing. It is proposed to use some programmable logic for the cluster finder. This programmable logic will be combined with the DDL3 link receivers inside the RORC3. The tracking in the EPN will use a combination of cores (simple or complex) and of GPUs.

The present generation of PCs is based on the Sandy Bridge chip which is the latest generation of the complex-core chip of Intel (see Figure 4.4). Each processor includes 8 cores, 4 ports to the DDR3 memory and 5 slots of PCIe Gen3 8 lanes. Two socket nodes with this type of processor are adequate as FLPs when equipped with a RORC3. They can also be used as EPNs when equipped with a GPU.

**Figure 4.4:** Block diagram of a PC based on the Sandy Bridge

Another key issue for such a large system is its footprint in terms of space and power. In order to reduce the space footprint, the RORC3 will include 12 DDL3 receiver links. A single PCIe slot will therefore be needed in the FLPs which can therefore be 1U rack-mount PCs. The standard GPU boards take 1 slot in the EPNs, which can be implemented as 1U rack-mount PC or also more compact but more expensive blade system.

4.4.4 Farm Network

The requirements for the future ALICE online computing network fabric mainly concern the number of nodes involved and the aggregate bandwidth of the network. The network fabric must be able to interconnect all the FLPs and EPNs, amounting to up to 1500 nodes. As indicated above, the network needs can be addressed by a single or a double network connection at 40 or 10 Gbit/s respectively. The aggregate bandwidth that the network will have to handle during the most demanding period of the Pb–Pb run is of the order of 2.4 Tbit/s and concerns two types of data traffic:

- The bandwidth between FLPs and EPNs for data processing and event building has to match a data throughput that corresponds to the minimum bias event rate of 50 kHz. This generates a data traffic of approximately 1.8 Tbit/s.

- The output from the EPNs, which represents approximately 0.6 Tbit/s, corresponds to the data actually recorded.

All the FLPs and EPNs will therefore be linked by a commercial high-performance network at 10 or 40 Gbit/s. The present DAQ and HLT networks are based on two different technologies: Ethernet and Infiniband. Devices (routers and switches) of the same two technologies exist which would fulfill the upgrade requirements.

These devices can be combined in two different ways:

- either as a *fat tree* [181] network combining a large central *director switch* and several *edge switches*,
- or as *spine and leaf network* [182] which tries to reduce the vulnerability of the network to a single component by distributing the central core network over several spine switches.

These two arrangements are independent of the technology and could both be used for Infiniband or Ethernet. Two examples of network implementation are given hereafter using the two technologies currently used in ALICE.

The first example of implementation of the ALICE network fabric is based on the Infiniband technology with a link speed of 40 Gbit/s. Some *director switches* already available on the market allow for the combination of hundreds of ports into a single switch. In order to provide thousands of ports, a fat tree network is assembled combining a large central *director switch* and several *edge switches*. The example of network shown in Figure 4.5 combines:

- one large *director switch* of type Mellanox IS5200 [183] with up to 216 ports at 40 Gbit/s and a total throughput of 17.3 Tbit/s;
- 48 *edge switches* of type Mellanox SX6025 [184] with up to 36 ports at up to 56 Gbit/s and a throughput of 4.03 Tbit/s.

Each of the *edge switches* provides connections to 32 nodes and delivers an up-link bandwidth of 80 Gbit/s. This configuration can support up to 3072 nodes for a maximum bandwidth of 7.7 Tbit/s.




Figure 4.5: The online computing network fabric implemented with Infiniband.

The largest model of *director switch* from the same manufacturer (IS5600 [185], 648 ports and 51.8 Tbit/s) allows to assemble larger configurations of up to 5184 ports and 25.9 Tbit/s. This technology would address the previously mentioned need to design an architecture which would be able to scale up to an interaction rate of up to 100 kHz (twice the nominal value).

The second example of implementation of the ALICE network fabric is based on the Ethernet technology with a link speed of 10 Gbit/s. Some *director switches* already available on the market allow combining hundreds of ports in a single switch. In order to provide thousands of ports, a spine and leaf layout has been used. In the example of network shown in Figure 4.6, the spine and leaf types of switches are actually the same piece of hardware:

- 4 spine switches of type DELL Z9000 [186] with up to 32 ports at 40 Gbit/s and a total throughput of 2.56 Tbit/s and

- 24 leaf switches of type DELL Z9000 with up to 128 ports at 10 Gbit/s.

Each of the leaf switch provides connections to 75 nodes and delivers an up-link bandwidth of 160 Gbit/s. This configuration can support up to 3072 nodes for a maximum bandwidth of 7.7 Tbit/s. Using a 3-stage *Clos network*:




Figure 4.6: The online computing network fabric implemented with Ethernet.

work [187] with the same switch, large configurations can be assembled with up to 8192 ports and 80 Tbit/s. The scalability requirement is also addressed with the Ethernet technology. These two examples show that the network required by the ALICE upgrade is already feasible today with a scalability of more than twice the number of nodes and the maximum bandwidth.

4.4.5 Reconstruction Software

In order to meet the requirements of the data processing architecture described in the previous section, we aim at designing a new common software framework for the HLT and offline code. There will be a single version of the AliRoot code, offering offline-quality performance with a computational efficiency comparable to the current HLT code. The calibration and reconstruction of the data will have to run very efficiently on a highly parallel system, with several levels of parallelism of different granularities. This will require the code to be optimized for the newly emerging parallel architectures.

This section is organized as follows. First we describe the current opportunities and experience of parallelisation of the AliRoot code, including reconstruction, simulation and analysis. Then we present a characterization of the present code in terms of CPU instructions to evaluate the optimization opportunities. Lastly, we present the plans to port the AliRoot code on the future architectures to cope with the challenges of the data-flow schema presented above.

4.4.5.1 Parallel Reconstruction

The ALICE Offline reconstruction code part of AliRoot has been parallelised in order to run it on the CERN Analysis Facility (CAF) [188].¹ Parallelisation at the event level is achieved using the PROOF [189] system and this code is now an integral part of the standard AliRoot code. Therefore any version of AliRoot can be run in parallel on any standard ALICE Analysis Facility (AAF).

The main motivation for this development was to provide fast feedback via full online reconstruction to assess data quality and detector performance beyond what could be done via the online Quality Assurance framework. It is also used to tune the code and several reconstruction parameters before full data reconstruction.

One of the major problems with parallel processing is load balancing. A 'direct' approach would make the total task no faster than the slowest individual task. To avoid this, each worker consumes its own data first. Then, in

¹Similar computing farms have been established at major ALICE computing centres such as CCIN2P3 at Lyon, Kosice, JINR in Russia, Subatech at Nantes, Torino and KISTI/GSDC in Korea.

order to ensure load balancing, it can consume data from another worker, accessing them over the local fabric till all data is consumed.

The processing rate in events per second and MByte per second achieved with this setup are shown in Figure 4.7. This system shows a good scaling with the number of workers. As shown, the performance is quite linear even




Figure 4.7: Dependence of the events rate per second and the MByte processed per second with the number of workers.

with a very small amount of files, showing a strong deviation from linearity only in case of an I/O bound example (reconstruction with ITS only) and a very small number of files (3). In this case the I/O is mostly non-local as tens of workers read three files, and most of the time is spent in reading. Prompt CAF reconstruction has been very important during the early phases of data taking. This coarse grain parallelisation, which has not been fully optimised, has required very few changes at the level of AliRoot. This shows that the intrinsic parallelism of the reconstruction code can be easily exploited, at least at the event level.

A finer granularity, thread-level version of the parallel AliRoot reconstruction has been developed in the framework of a doctoral thesis [190]. The prototype was developed to explore the issue of thread-safety in the ROOT and AliRoot frameworks, by using semi-automatic code transformation. A code transformation tool was developed to perform static analysis and source-code rewriting based on the Low Level Virtual Machine (LLVM) front-end Clang [191], following an approach derived from the Geant4-MT project [192].

With this method each thread contains a whole event reconstruction procedure, providing a reduction of the data processing time by a total factor of 3.6 for p-p raw events and 5.6 for lead-lead data with 8 threads on an Intel Westmere CPU as shown in Figure 4.8. The Cint interpreter is shared and the virtual memory consumption for a single thread is reduced by 370 MByte. For comparison, a single process has a footprint of around 770 MByte of virtual memory for p-p reconstruction and 1 GByte for lead-lead reconstruction. Since each thread is using its own set of files, up to 200 MByte are additionally allocated per thread, resulting in around 600 MByte for each additional lead-lead reconstruction thread.

Performance gains have shown to be reasonably good for a test environment with less than 10 threads. Further gains can be obtained by additional manual intervention, reducing lock contention in ROOT objects, using atomic operations and improving the algorithms used to optimise memory locality. Memory consumption can still be optimised e.g. by sharing the detector geometry between similar threads.

4.4.5.2 Parallel Simulation

Event level parallelism has been exercised by sharing the code and data structures [193] via OS (Operating System) page sharing mechanisms such as KSM (Kernel Samepage Merging). This first parallel approach contains few critical sections: retrieving the track to be transported, reordering the kinematics, renumbering the stack and hits, then finally filling the output structure.

This approach implies a small penalty due to memory-bus contention between threads, but exhibits an almost linear scaling of the memory increase due to the fact that the GEANT data structures are not shared. Sharing of




Figure 4.8: Speed-up versus memory consumption for a lead-lead reconstruction, expressed as the ratio to memory usage in a single thread.

the data structures in this prototype results in using less memory than independent copies of the full AliROOT. Only 400 MByte of resident memory and 900 MByte of virtual are necessary for each process, which represents a significant reduction in memory usage.

Another outcome of this exercise was a study of the performance that could be achieved depending on the size of the output. It became clear that the more frequent dumps have to be made from the memory buffers to file, the bigger the penalty since I/O is a sequential process. In Figure 4.9 we can see the results obtained with up to 4 threads for a 2 GByte output of the simulation.




Figure 4.9: Processing time (left) and resident memory usage (right) for 2 GByte output. Resident memory per thread is 135 MByte and virtual 175 MByte.

To move one step further in the optimisation of the simulation, we started a new prototype with the goal to review the transport procedure from a parallel perspective, starting with a realistic geometry but only few simple physics processes. Preliminary work was done to fully parallelise and optimise the ROOT geometry modeller, moving all state-dependent data structures into a navigation utility class. The multi-threaded version of the geometry is part of the production release of ROOT.

To improve data locality, the new transport strategy will profit from the cached cross sections for the current material and a given type of particle. The transport of all particles will be done at the same time in a given detector region (like those produced by an EM shower) rather than following each individual particle within the entire detector. In addition, one could profit for resolving physics processes or geometry queries for a vector of particles instead of a single one. This does not only offer the possibility to parallelise loops, but in many cases allows for compiler-specific optimisation and vector-wise handling of floating point operations.

Figure 4.10 presents the performance of the prototype in its current state. Tests have been performed on a 12 core machine allowing for up to 24 threads (12 native and 12 hyper-threads). The prototype shows a good scalability up to 11 worker threads and 1 dispatcher. A concurrency efficiency analysis shows that the main overhead comes from the transition states triggered by communicating baskets during the garbage collections. This will be easily compensated in future by balancing the number of transport threads and switching them to other types of tasks, like digitisation or I/O.




Figure 4.10: Performance of the new transport prototype. The test machine has 12 native cores and the speed-up was computed by dividing the real time for running with a single thread by the time taken by N parallel threads.

The prototype currently uses toy models for energy deposition, scattering and ionisation processes, which will soon be extended to EM processes that will allow realistic production of showers. The performance is expected to improve by an important factor when the physics simulation will become more realistic and reduce the level of inter-thread communication. The knowledge achieved along the process of developing the prototype will be used for improving gradually the ALICE simulation.

4.4.5.3 Parallel Analysis

The ALICE analysis framework was designed to be modular with respect to generic “analysis modules”, which are executed sequentially after the current event is being loaded into memory. An ALICE analysis module is seen as an independent task that shares with others the main event loop steered by the framework.

While this analysis model can be naturally parallelised at file and event levels on computing clusters or batch farms, finer parallelisation is intrinsically limited by the sequential nature of the processing phases involved. The output of analysis is typically a set of histograms or trees and there is usually no correlation between different outputs since the input events are in general independent for a given sub-sample. The caveat in this schema is the need to get the full output corresponding to the totality of the input data-set from the partial outputs. This merging procedure is by definition sequential but can be parallelised as $\log(N)$.

The wall time fractions for different processing stages suggests a large potential gain if the event processing phase is parallelised, in case of CPU bound analysis trains. This comes however with a caveat: any kind of splitting of the output data structures involve an extra merging phase having a time weight proportional to the level of parallelism used while processing. This imposes additional constraints on the output format to support parallel filling and asynchronous merging into the final result.

Parallelisation of the event processing phase is maybe the most natural and easy to implement. Here we have to keep in mind that the data structures for storing the output need to be adapted to a multi-threaded environment from the side of ROOT.

The merging phase is maybe the hardest to implement in a fully parallel environment. An efficient parallel flow needs that all phases are asynchronous and that any phase can start working ahead to minimise the finishing tail. For this reason, the output ROOT structure would need to support parallel asynchronous filling with a minimum or

no overhead.

The observations above hint for possible strategies for pushing the ALICE analysis beyond the embarrassing parallelism model. The efforts in dealing with the input, processing and output phases can be done independently since they can have immediate impact on the analysis efficiency. The steering of input events and reshuffling of output structures have to be certainly reviewed in collaboration with the ROOT team.

4.4.6 Parallel HLT Reconstruction on GPU and CPU

The track recognition and fitting is the most time-consuming step during the event reconstruction. Its optimization has been one of the main R&D activities in the current HLT project, both to reduce the latency in the system as well as to reduce the overall cost of the system.

As we discussed earlier in this document, CPU clock speed has stagnated over the last years while the number of cores keeps increasing. Good parallelization was thus one of the main design objectives of the ALICE HLT TPC tracker algorithm. The first two steps of the tracker perform a combinatorial search for short track candidates based on the Cellular Automaton principle. The steps in detail are:

- Neighbors finder: for each cluster the neighboring rows are searched for the pair of clusters such that these three clusters form the best straight line. Such connections are stored and in the following called link.
- Evolution step: reciprocal links are determined and stored, all the other links are removed. Chains of consecutive links build the initial tracklets: short track candidates which are used as basis for the next tracking steps.
- Tracklet constructor: track parameters are fitted to each tracklet applying the Kalman filter. Using these parameters the trajectory is extrapolated to adjacent rows. New clusters close to the extrapolated location are searched for and, if a χ^2 criterion is met, added to the tracklet. The new cluster is incorporated in the parameter fit and the process is repeated until no new clusters are found in at least five consecutive rows.
- Tracklet selector: it is possible that multiple tracklets have overlapping parts. In such cases only the longer tracklet is stored, the shorter removed. If thereafter the tracklet passes a cut on χ^2 and on the number of clusters, it is finally stored as a track.

Each of these steps can be processed in a parallel way. For instance the neighbors finder can search for the links for each cluster independently and in parallel, the tracklet constructor can fit and extrapolate many tracklets at the same time, and so forth.

Besides the above described steps of the algorithm, two more simple data reformatting steps are needed for the implementation: an initialization step that prepares special data structures such as a grid for fast search, and a tracklet output step.

The parallel architecture of the algorithm allowed for an easy adaptation of the tracker to GPUs. All of the main steps of the tracker have been ported to the GPU, with the initialization and output step still executed on the CPU. A pipelined processing ensures good GPU utilization: while the GPU performs the tracking for sector i the DMA engine can transfer sector $i-1$ and the processor can preprocess sector $i-2$. Transfer back to the host and post-processing are pipelined accordingly. An improved multi-threaded version of the pipeline has been developed, which processes initialization and output in a round-robin fashion on multiple CPU cores in parallel, with currently three cores used for this task.

Figure 4.11 shows a comparison of the CPU and GPU versions of the tracker. On the left the time needed for the tracking steps for a single TPC sector and on the right the total tracking time. The CPU version uses a parallelization over sectors to enable a fair comparison. As can be seen from the section in the right of the figure, overall the GPU outperforms the CPU by a factor of three.

The HLT computing system was equipped with GPUs in all computing nodes, which frees the CPU cores in the nodes for additional tasks like vertexing or trigger algorithms. The use of GPUs allows reducing the number of computing nodes in the system by a factor of three, which results in significant cost savings.




Figure 4.11: Comparison of HLT TPC tracking on CPU and GPU.

4.4.7 Analysis of Current AliRoot Performance

In order to be able to quantify the potential performance increase that could be obtained by an optimisation of the AliRoot code on modern CPUs, we have analysed the performance of some typical AliRoot applications and compared them with well optimised applications on modern CPUs. The machine where we run the benchmarks is 6 core Intel i7-3930K (3.2GHz) with 16 GByte of RAM. The results we obtained are shown in Table 4.4.

In this table we report the instructions per CPU cycle, the load and store operations from memory, the stalls, i.e. the cycles where the CPU has been waiting for resources, such as I/O or memory access, and the branches, i.e. `if` statement and equivalent constructs.

In comparison to a well optimised code such as BLAS (Basic Linear Algebra Subroutines) or the ROOT linear algebra algorithms, we can see a potential gain of a factor 2 in instructions per cycle. A better reorganization of the data layout could provide an optimization of the load and store operations, with a corresponding gain in performance.

Table 4.4: Performance comparison for various applications on a 6 core Intel i7-3930K (3.2GHz) with 16 GByte RAM.

	Instructions per cycle	Load and store	Stalls	Branches
complex*16 blas	2.0	34.4%	9.6%	13.2%
root stressLinear	2.0	38.4%	12.4%	19.4%
Analysis train	1.1	60.1%	56.4%	20.9%
Simulation (PbPb)	1.0	53.3%	34.5%	17.4%
Reconstruction (PbPb)	1.1	55.0%	36.6%	17.0%

Some recent work on porting AliRoot and similar code on GPUs [194] shows clearly that, with an opportune optimization of the algorithms, it is possible to achieve a substantial acceleration at the micro-instruction level even for very complex codes such as the TPC tracking. This optimization alone would not yield appreciable results if it is not accompanied by a coarser level optimization of the work streams, aimed at reducing or eliminating all synchronization points. Nevertheless this optimization at the level of micro-instructions would allow to much better follow the “technology wave”. It is very likely that the parallelism opportunities included in the code could be exploited by future compilers.

4.4.8 Plans for the Development of Parallel AliRoot

The development of a parallel version of AliRoot will take advantage from the substantial experience gained with parallel programming both in Offline, Online and HLT described above.

The first step will be to establish a working group made by members of HLT, Offline and DAQ to revise the current structure of AliRoot and of the HLT reconstruction. A reconstruction prototype will then be developed during 2013 with the aim to reproduce the current AliRoot / HLT functionality and performance in the reconstruction of both pp and Pb–Pb events by the end of that year. This prototype should feature coarse level parallelism (process and thread-level) with minimal synchronisation and efficient load balancing.

During the development of the prototype, several micro-parallel version of the computational kernels will be implemented on CPU, GPGPUs and MICs will be evaluated. This will be done with OpenMP, OpenACC, CUDA and OpenCL.

One important step here will be to reinforce our testing framework, via a better integration with the existing *cdash* tools, part of the *cmake* framework that we are already using for our code development environment. Immediate testing of code compilation, coverage and even certain tests can be automated, and the developers and responsible persons of different packages will be immediately informed in case of errors or problems. This will introduce a necessary improvement of the quality assurance of the newly developed code which will be very important during an intense code development phase.

Once the framework has been defined, detector experts from each detector will be associated to the work to develop the reconstruction code for their detector in close collaboration with the core offline and HLT groups.

The problem of the I/O should be worked out in close collaboration with the ROOT team that is developing a parallel I/O subsystem for ROOT.

Chapter 5

Schedule, Cost Estimate and Organization

5.1 Schedule

The ALICE upgrade is planned to be in operation after Long Shutdown 2 (LS2) and has a programme that will extend into the HL-LHC era after Long Shutdown 3 (LS3). The current LHC schedule foresees LS2 to take place in 2018 and LS3 to start in 2022. The upgrade of the TPC defines the installation strategy and schedule. The installation of the new TPC readout chambers has to take place on the surface in the SXL2 cleanroom at P2, and therefore the TPC must be removed from the cavern during this time. In order to maximize the time available for the installation of the new TPC readout chambers and electronics, the removal of the TPC will take place at the beginning of the shutdown period. The reinstallation takes place towards the end of this shutdown period, in a time which is still compatible with the subsequent ITS installation and commissioning. While the TPC is on the surface, the upgrade of the other detectors, related services and infrastructure work will also be performed. We currently envisage two installation scenarios (see Fig. 5.1). The two plans are based on different assumptions for the duration of LS2. Scenario 1 anticipates a duration of 18 months, while Scenario 2 has a time frame of 14 months.

- Installation Scenario 1 (see Fig. 5.2). The complete upgrade takes place in an LS2 of at least 18 months. After removal of the shielding, compensator magnet, Miniframe (a large support structure carrying ITS and TPC services), ITS and beampipe, the TPC arrives in the cleanroom 3 months after the start of the shutdown. The replacement of the readout chambers with GEM detectors and the recommissioning with cosmic rays on the surface lasts for 10 months. During this time the upgrade of the other detectors and the required service modifications are being carried out in the cavern. The reinstallation of the TPC and the new beampipe in the cavern takes 2 months. This leaves 3 months for installation and commissioning of the new ITS detector. The time sharing between the TPC work on the surface and ITS installation can be optimized according to the specific needs. Some contingency is included in the individual activities.
- Installation Scenario 2 (see Fig. 5.3). The TPC is removed during the 'End of Year Technical Stop' 2016/2017, ALICE does not operate in 2017 and the upgrade installation happens in an LS2 of at least 14 months. A temporary beampipe is installed inside the ALICE detector for the 2017 LHC run. The advantage of this scenario is that the TPC can be in the cleanroom for 16 months and there is an explicit contingency of 2 months included in the installation schedule. There is also more flexibility for the sharing of TPC upgrade and ITS installation time. In order to achieve the 1 nb^{-1} of integrated heavy ion luminosity before LS2, this scenario requires an extended heavy ion run in 2016.

Project timelines for the ITS, TPC, TOF, Muon Tracking, Muon Trigger and Online Systems upgrade are given in Tables 5.1 to 5.7.




Figure 5.1: LHC schedule, Heavy Ion periods (HI) and the two upgrade installation scenarios.




Figure 5.2: Installation scenario 1: the entire upgrade takes place in an LS2 of 18 months.




Figure 5.3: Installation scenario 2: the upgrade is shared between the End of Year Technical Stop 2016/2017 and an LS2 of 14 months.

Table 5.1: ITS upgrade timeline

Year	Activity
2012 – 2014	
2012	R&D finalization of detector specifications evaluation of detector technologies (radiation and beam tests) first prototypes of sensors, ASICs, and ladders (demonstrators)
2013	selection of technologies and full validation engineering design for sensors, ASICs, ladders, data links engineering design for support mechanics and services Technical Design Report
2014	final design and validation
2015 – 2018	
2015 – 2016	Construction and Installation production, construction and test of detector modules
2017	assembly and pre-commissioning in clean room
2018	installation in the cavern

Table 5.2: TPC upgrade timeline

Year	Activity
2012 – 2013	
Sep 2012	R&D test of IROC prototype with GEM readout at PS verification of dE/dx resolution
Jan 2013	test of IROC prototype with GEM readout in ALICE cavern verification of operational stability in LHC conditions
2012 – 2013	tests of small prototypes with different GEM parameters evaluation ion back flow and gain stability
2013	studies of drift properties for gases with CF4 Technical Design Report
2014 – 2016	
	GEM production QA and characterization of GEM foils preparation and tests of GEM foil stacks
2013 – 2016	
2013 – 2014	FEE development ASIC development (shaping amplifier, digitization and zero suppression)
2015	development of front-end cards, pre-series production and qualification
2016	mass production and testing
2016 – 2019	
2016-2017	Scenario 1 construction of new ROCs mounting of GEMs on new ROCs and tests
2018	removal of the TPC from the cavern
2019	installation of new ROCs on the TPC and commissioning TPC installation in the cavern
2017 – 2019	
2017	Scenario 2 removal of the TPC from the cavern removal of the ROCs and mounting of GEMs
2018	installation of ROCs and TPC commissioning TPC installation in the cavern

Table 5.3: TOF upgrade timeline

Year	Activity
2013 – 2014	R&D
2013	DRM2dev prototype (demonstrator) simulation studies to assess TOF trigger upgrade needs DRM prototypes and firmware developments
2015	DRM2.0 version (subject to TTCrx successor development) trigger card prototypes for communication with FTP
2016 – 2018	Construction and Installation
2016	procurement and start of production
2017	DRM2 and Trigger Card production
2018	installation in the cavern

Table 5.4: Muon Tracking upgrade timeline

Year	Activity
2013 – 2015	R&D
2013	CROCUS and MANU prototypes
2014	CROCUS pre-production
2015	MANU pre-production
2015 – 2018	Construction and Installation
2015	CROCUS production
2016 – 2017	MANU production
2018	installation in the cavern

Table 5.5: Muon Trigger upgrade timeline

Year	Activity
2013 – 2015	R&D
2013	front-end and regional boards prototypes
2015	DARC card prototype
2014 – 2018	Construction and Installation
2014	Regional Boards production and front-end pre-production
2015	front-end electronics production
2016	front-end electronics production and DARC card production
2018	installation in the cavern

Table 5.6: PHOS upgrade timeline

Year	Activity
2014 – 2016	R&D
2014	system requirements document
2015	design and development of new FEC modification of DCal TRU for PHOS
2016	validation of new FEC and TRU units system ready for production
2016 – 2018	Construction and Installation
2016	electronics production
2017	installation on the PHOS at the surface
2018	installation in the cavern

Table 5.7: Online Systems upgrade timeline

Year	Activity
2012 – 2014	strategy definition and R&D definition of the strategy for common DAQ, HLT and Offline framework R&D on key hardware, firmware and software technologies
2013 – 2016	simulation, demonstrators and prototypes, choice of technologies simulation of trigger and data-flow architecture development of demonstrators and prototypes for the key hardware and firmware technologies development of the new common software framework selection of technologies used in production
2017 – 2020	production and procurement, staged deployment production of the hardware developed by the projects market surveys, tendering and procurement of commercial equipment full deployment of readout part (DDL3 and FLPs) staged deployment of processing part (EPNs, network and data storage) according to accelerator luminosity.

5.2 Cost Estimate

The costs associated to the activities described in this LoI should be understood in terms of CORE values: material costs, excluding Institutes personnel costs and basic infrastructures. CORE estimates include detector components and production costs, as well as industrial or outsourced manpower for production, installation and integration. Table 5.8 shows the total estimated costs of all upgrade items and Common Projects. Table 5.9 shows a preliminary spending profile for the upgrade projects.

As part of this, it is expected that each Project will prepare an addendum to the ALICE Memorandum of Understanding (MoU) [195]. This Addendum will describe the obligations of each Institute involved in the construction, the commissioning and the operation of the system, as well as the value of the deliverables of each Institute. The Addendum will be signed by each Funding Agency partner in the Project. It will include the Project organization, a plan including the main Projects milestones and a cost breakdown for each deliverable.

A separate Addendum to the MoU for the Common Projects will include a detailed cost breakdown and the planning. It will also describe the obligations of each Institute which are full members of the ALICE Collaboration, in terms of cost sharing of the Common Projects. As per the 'Maintenance and Operation category A' costs [196], the obligations of the Institutes and their Funding Agencies, towards the Common Projects will be proportional to the total number of scientists signing ALICE publications. The Common Projects amount to $\leq 15\%$ for a total upgrade cost of 36 MCHF and $\leq 13\%$ in case the total upgrade cost amounts to 42 MCHF. Contributions to the Common Projects will be made in two ways:

- as cash payments to a dedicated common fund account at CERN managed by the ALICE Resources Coordinator;
- as in-kind contribution by taking responsibility for a Common Project item or activity, in agreement with the ALICE Management.

A cost breakdown for the major upgrade projects is presented in the following.

Table 5.8: Cost estimates for the ALICE Upgrade

System	ALICE Upgrade	Cost[MCHF]
ITS option 1 ^a		10.0
ITS option 2 ^b		16.0
TPC ^c		5.5
TRD		0.6
TOF		0.7
PHOS ^d		0.8
Muon Tracking		1.6
Muon Trigger		0.5
Online Systems ^e		9.3
Offline		0.5
Trigger Detectors		1.0
Common Projects		5.5
Total with ITS option 1		36.0
Total with ITS option 2		42.0

^a ITS option 1: 7 layers of monolithic pixels^b ITS option 2: 3 layers (inner) of hybrid pixels and 4 layers (outer) of strips^c TPC: if the TPC is removed during the technical stop 2016/2017 then the estimated cost is 4.5 MCHF since existing readout chamber are reused^d The PHOS upgrade is not yet approved by the ALICE Collaboration. Physics performance studies ongoing.^e Online Systems: the option for a rate of 100 kHz is estimated with a total cost of 15.1 MCHF**Table 5.9:** Spending profile for construction and installation in MCHF

Systems	2013	2014	2015	2016	2017	2018	2019	2020	Total [MCHF]
ITS Option 1		4.00	4.00	1.40	0.60				10.00
ITS Option 2		6.40	6.40	2.30	0.90				16.00
TPC Scenario 1	0.10	0.50	1.20	3.00	0.70				5.50
TPC Scenario 2	0.10	0.30	0.80	2.60	0.70				4.50
TOF	0.02	0.04	0.06	0.30	0.30				0.72
PHOS		0.30		0.40	0.12				0.82
Muon Tracking	0.04	0.40	0.52	0.60					1.56
Muon Trigger	0.18	0.17	0.19						0.54
Online Systems			3.10	2.90	3.30				9.30
Online Systems @ 100 kHz			3.10	2.90	3.60	5.50			15.10
Common Projects: beampipe			1.50						1.50
Common Projects		0.80	1.20	2.00					4.00

ITS Cost Estimate

At the present stage several technologies are being considered for the implementation of the ITS detector. The technology that will be adopted for the implementation of the different detector layers, will be chosen to best suit the detector requirements within the boundaries set by the available funds and the project time line. In terms of performance and cost the use of monolithic pixel detectors would be suitable for all layers, while the hybrid pixel detectors may become prohibitive as the radius increases due to cost reasons. However, a limitation to the use of monolithic pixel detectors for the innermost layers may arise from the level of radiation that can be tolerated by the selected technology. Moreover, it could turn out that a few microstrip layers at large radii is mandatory for particle identification if the charge signal delivered by the thin sensor layer of monolithic pixel detectors does not provide enough resolution. Based on the above considerations it can be reasonably assumed that the ITS will be based at most on two silicon detector technologies.

In the following we give a cost estimate for the two design options under study:

- Option 1: 7 layers of monolithic pixel detectors, which would provide excellent standalone tracking efficiency and p_T resolution, but limited PID capability.
- Option 2: a combination of 3 innermost layers of hybrid pixel detectors and 4 outermost layers of strip detectors, which would provide better PID capability but worse standalone tracking efficiency and momentum resolution.

The cost breakdown for the two detectors options is reported in Table 5.10.

TPC Cost Estimate

Table 5.11 describes the cost breakdown for the TPC upgrade in two different installation scenarios. The cost estimate is based on the cost for the construction of the current TPC chambers as well as existing detectors, which include GEM readout. Depending on the final LHC and shutdown schedule we follow two scenarios. In scenario 1 we assume that the chamber bodies have to be rebuilt to accommodate a complete exchange of the readout chambers in 2018. In scenario 2, the TPC will be removed already in 2016/2017, making possible a timely dismantling and reuse of the existing readout chamber bodies.

TOF Cost Estimate

Table 5.12 shows the cost estimate for the upgrade of the TOF electronics. It includes the replacement of the modules which contain the DDL connection (called Data Readout Module, DRM, for a total of 72 units) with new ones containing the new DDL2 or DDL3. A possible upgrade of the TOF central trigger module (CTTM) is also included.

TRD Cost Estimate

It includes the replacement of the current GTU system with an enhanced bandwidth version.

PHOS cost estimate

Table 5.13 shows the cost estimate for the upgrade of the PHOS electronics. A preliminary study, which is used as basis for the cost estimate, indicates that in order to perform the readout of Pb–Pb interactions at 50 kHz the entire electronics system (front-end, readout and trigger) of the PHOS detector has to be replaced. The studies on the physics benefit that would result from this upgrade are not yet completed (see also Section 3.8 of Chapter 3) and a decision on upgrading PHOS is therefore postponed.

Muon Spectrometer Cost Estimate

Table 5.14 presents the cost breakdown for the upgrade of the Muon Spectrometer electronics. The cost of the CROCUS and the MANU cards is based on existing units. For the front-end electronics of the Muon Trigger, the estimate is based on similar systems in ATLAS and CMS. For the Muon Trigger readout, the estimate is based on the cost of the current system.

Online Systems Cost Estimate

The Online Systems will be built from a combination of custom and commercial equipment. A new version of the custom units (DDL2 and RORC2) are being developed and will be used after LS1. They are prototypes of the following generation (DDL3 and RORC3) to be used after LS2. These prototypes are used to assess the technology and as basis for the cost estimate. Establishing contacts with the major computing companies will lead to an evaluation of the performance evolution, and the price of commercial equipment at the time of procurement. In Table 5.15 we give a cost estimate for an online system able to read out, compress and record the data resulting from a rate of 50 kHz of minimum bias heavy-ion interactions. In a second stage, if required, the system will be able to scale to 100 kHz. The price estimate for this rate increase is also indicated.

Offline Cost Estimate

The computing upgrade will include two aspects:

- Contingent on the detector and online systems upgrade. This will include simulation, development of code for calibration and alignment of the new detectors elements and their integration in the offline framework.
- Development of a new reconstruction and analysis framework for efficient running on multicore and GPU processors.

The cost indicated in Table 5.8 represents the investment needed for the offline central services at CERN.

Common Projects Cost Estimate

The cost for the Common Projects is summarized in Table 5.16. The design effort for detector and services integration assumes 2.5 FTEs for a period of 3 years leading up to LS2. For the installation period we assume a team of 5 service contract technicians supporting the installation team during 1.5 years. The installation of new services including cabling work, electrical infrastructure, new general cooling and ventilation infrastructure amounts to 1.5MCHF. The latter includes the installation of new services as well as the removal of numerous services that will be replaced by the new ones. Access and support structures including installation tools and structure modifications amount to 0.5MCHF. The cost for the beampipe reflects the current estimate as defined in the workpackage agreement with the CERN vacuum group.

Trigger Detectors Cost Estimate

Several ALICE groups have expressed interest in the development of the new Trigger Detectors and are currently developing a project organization. Given the more limited size of the project, this does not introduce delays in the overall project. A cost estimate of 1 MCHF, which is based on the cost of the current TO, VO and ZDC detectors, is included in the overall upgrade cost.

5.3 Organization

As previously mentioned in section 5.2, each project will prepare an addendum to the MoU that describes the project organization and obligations of each participating institute and their funding agency. The list of institutes that have expressed an interest in participating in the development of the various upgrade projects is reported in tables 5.17 to 5.23. It should be noticed that all proposed upgrades are framed within the organization of existing ALICE projects. The institutes, which are currently participating in the projects that are part of the upgrade programme, have all endorsed the long-term operation of the current detectors and their corresponding upgrade. In addition, several new institutes have joined the ALICE upgrade effort and are in the process of becoming ALICE members.

All institutes, which are full member of the Collaboration, will contribute to the Common Projects as already mentioned in Section 5.2. Table 5.24 shows all the institutes member of the ALICE Collaboration as of August 2012.

All institutes taking part in the ITS project have specific expertise and past experience in the development, construction and running of silicon trackers. Several institutes have long standing expertise in ASIC design, construction of detector ladder and support mechanics, manufacturing of composite materials, integration and characterization of

Table 5.10: Cost estimates for the ALICE Inner Tracking System

Item	ITS option 1 ^a	Cost[MCHF]
Silicon	2.0	
Thinning, dicing and mounting	1.0	
MCM	1.0	
Kapton cables	1.0	
Cables, connectors and patch panels	2.0	
Mechanics and cooling	1.5	
On-detector low-voltage regulation	0.5	
Power supply	1.0	
Total	10.0	

Item	ITS option 2 ^b	Cost[MCHF]
Double side detector	8.0	
Front-end electronics	4.0	
Readout	0.8	
Assembly tools	0.8	
Cables and optical fibers	0.8	
Cooling	0.2	
Mechanics	0.6	
Power supply	0.8	
Total	16.0	

^a ITS option 1 = 7 layers of monolithic pixels^b ITS option 2 = 3 layers (inner) of hybrid pixels and 4 layers (outer) of strips

Table 5.11: Cost estimates for the Time Projection Chamber

Item	TPC scenario 1 ^a	Cost[MCHF]
GEM foils	0.5	
Chamber bodies, pad planes	1.0	
Frames, components	0.1	
Assembly and installation tooling	0.2	
HV system	0.2	
Services	0.5	
Front-end electronics	3.0	
Total	5.5	

Item	TPC scenario 2 ^b	Cost[MCHF]
GEM foils	0.5	
Frames, components	0.1	
Assembly and installation tooling	0.2	
HV system	0.2	
Services	0.5	
Front-end electronics	3.0	
Total	4.5	

^a TPC scenario 1: complete exchange of the readout chambers takes place in 2018 and chamber bodies have to be rebuilt

^b TPC scenario 2: TPC removed during Technical Stop 2016/2017 and existing readout chamber are reused.

Table 5.12: Cost estimates for the Time Of Flight

Item	TOF	Cost[MCHF]
DRM	0.60	
Trigger card CTTM	0.12	
Total	0.72	

Table 5.13: Cost estimates for the Photon Spectrometer

Item	PHOS [*]	Cost[MCHF]
Front-End Cards	0.70	
Trigger Region Units	0.10	
Readout Controller	0.02	
Total	0.82	

^{*} The PHOS upgrade is not yet approved by the ALICE Collaboration. Physics performance studies ongoing.

Table 5.14: Cost estimates for the Muon Spectrometer

Item	Muon Spectrometer	Cost[MCHF]
CROCUS		0.24
MANU		1.32
MTR front-end electronics		0.24
Readout		0.18
RPC replacement		0.12
Total		2.10

Table 5.15: Cost estimates for the Online Systems

Item	Online Systems	Cost[MCHF]
DDL fibres		0.9
EPN		4.1
FLP and CRORC		0.9
Infrastructure		1.3
Networks		0.8
Servers		0.5
Storage		0.6
Central DCS		0.2
Total[*]		9.3

^{*} The option for a rate increase to 100 kHz is estimated with a total cost of 15.1 MCHF.

Table 5.16: Cost estimates for the Common Projects

Item	Common Projects	Cost[MCHF]
Design and engineering		1.0
Installation Manpower		1.0
Services		1.5
Beam Pipe		1.5
Access and Support structures		0.5
Total		5.5

hybrid pixel and microstrip detectors. The world leading experts for the development of monolithic pixel detectors are also part of the ITS collaboration.

Most of the institutes involved in the TPC upgrade have participated in the construction of the current TPC. Additionally, a number of institutes with expertise in the construction of GEM detectors and their associated readout electronics have joined the TPC Collaboration.

All institutes taking part in the TOF upgrade project have already been involved in the design, construction, assembly and commissioning of the detector. In particular all the readout electronics was developed and produced in a joint effort among INFN personnel and external firms.

The upgrade of the Muon Spectrometer is under study by the same institutes that participated in the construction of the current detectors and the related electronics. The upgrade of the Muon Trigger will be carried out by Clermont-Ferrand, Nantes and Torino. They have expertise in electronics, readout and RPC detectors respectively. For the Muon Tracking, Orsay, which designed and built the present CROCUS, will carry out its upgrade. Concerning the front-end MANU electronics, besides Orsay and Cagliari, which have developed the present MANU electronics, other members of the Muon Spectrometer Collaboration have expressed their interest in participating in the upgrade.

All institutes taking part in the online and offline projects upgrade have specific expertise and past experience in the development, production and operation of large online and offline computing systems.

Table 5.17: Expression of interest to the ITS detector upgrade, subject to funding

Country Funding Agency	City	Institute
CERN	Geneva	European Organization for Nuclear Research
China	Wuhan	Central China Normal University (CCNU)
Czech Republic	Řež u Prahy	Nuclear Physics Institute of the ASCR
France IN2P3	Strasbourg	Institut Pluridisciplinaire Hubert Curien (IPHC), Université de Strasbourg, CNRS-IN2P3
Italy INFN	Bari	Sezione INFN e Dipartimento dell'Università e del Politecnico di Bari
Italy INFN	Cagliari	Sezione INFN e Dipartimento di Fisica dell'Università di Cagliari
Italy INFN	Catania	Sezione INFN e Dipartimento di Fisica dell'Università di Catania
Italy INFN	Frascati	Laboratori Nazionali di Frascati (LNF)
Italy INFN	Legnaro	Laboratori Nazionali di Legnaro (LNL)
Italy INFN	Padova	Sezione INFN e Dipartimento di Fisica e Astronomia dell'Università di Padova
Italy INFN	Roma	Sezione INFN e Dipartimento di Fisica dell'Università "La Sapienza" di Roma
Italy INFN	Torino	Sezione INFN e Dipartimento di Fisica dell'Università di Torino
Italy INFN	Trieste	Sezione INFN e Dipartimento di Fisica dell'Università di Trieste
Korea	Seoul	Yonsei University
Pakistan	Islamabad	Faculty of Sciences, COMSATS, Institute of Information Technology
Russia	St. Petersburg	Institute of Physics, St. Petersburg State University
Slovakia	Košice	Slovak Academy of Sciences, IEP
Thailand	Nakhon Ratchasima	Suranaree University of Technology
UK STFC	Birmingham	University of Birmingham
UK STFC	Warrington	STFC Daresbury Laboratory
UK STFC	Chilton	Rutherford Appleton Laboratory
Ukraine	Kharkov	Ukrainian Academy of Sciences, KIPT-KFTI
Ukraine	Kharkov	Scientific Research Technological Institute of Instrument Engineering SRTIE
Ukraine	Kiev	Bogolyubov Institute for Theoretical Physics (BITP)

Table 5.18: Expression of interest to the TPC upgrade, subject to funding

Country Funding Agency	City	Institute
Denmark	Copenhagen	Niels Bohr Institute, University of Copenhagen
Finland	Helsinki	Helsinki Institute of Physics
Germany BMBF	Frankfurt	Institut für Kernphysik, Johann Wolfgang Goethe-Universität Frankfurt
Germany BMBF	Heidelberg	Physikalisches Institut, Ruprecht-Karls Universität Heidelberg
Germany BMBF	Munich	Physik Department, Technische Universität München
Germany BMBF	Munich	Excellence Cluster 'Universe', Technische Universität München
Germany BMBF	Tübingen	Physikalisches Institut, Eberhard Karls Universität Tübingen
Germany BMBF	Worms	FH Worms, Worms
Germany GSI	Darmstadt	Research Division and ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung
Japan	Tokyo	University of Tokyo
Mexico	Mexico City	Instituto de Ciencias Nucleares, Universidad Nacional Autónoma de México
Norway	Bergen / Tonsberg	Department of Physics, University of Bergen, Vestfold University College, Tonsberg
Norway	Bergen	Faculty of Engineering, Bergen University College
Poland	Cracow	The Henryk Niewodniczanski Institute of Nuclear Physics, Polish Academy of Science
Romania	Bucharest	National Institute for Physics and Nuclear Engineering
Slovakia	Bratislava	Faculty of Mathematics, Physics and Informatics, Comenius University
Sweden	Lund	Division of Experimental High Energy Physics, University of Lund
USA DOE	New Haven	Yale University, New Haven

Table 5.19: Expression of interest to the TOF upgrade, subject to funding

Country Funding Agency	City	Institute
Italy INFN	Bologna	Sezione INFN and Dipartimento di Fisica dell'Università di Bologna
Italy INFN	Salerno	Sezione INFN and Dipartimento di Fisica dell'Università di Salerno
Italy Centro Fermi	Rome	Centro Studi e Ricerche e Museo Storico della Fisica "Enrico Fermi"
Rep. of Korea NRF	Gangneung	Gangneung-Wonju National University

Table 5.20: Expression of interest to the TRD upgrade, subject to funding

Country Funding Agency	City	Institute
Germany BMBF	Heidelberg	Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg
Germany BMBF	Frankfurt	Institut für Kernphysik, Johann Wolfgang Goethe-Universität Frankfurt
Germany BMBF	Frankfurt	Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe-Universität
Germany BMBF	Münster	Institut für Kernphysik, Westfälische Wilhelms-Universität Münster
Germany GSI	Darmstadt	Research Division and ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung

Table 5.21: Expression of interest to the Muon Spectrometer upgrade, subject to funding

Country Funding Agency	City	Institute
Armenia	Yerevan	Yerevan Physics Institute
France CEA	Saclay	Commissariat à l'Energie Atomique, IRFU
France IN2P3	Orsay	Institut de Physique Nucléaire, Université de Paris-Sud, IN2P3-CNRS
France IN2P3	Clermont-Ferrand	Laboratoire de Physique Corpusculaire, Université Blaise Pascal Clermont-Ferrand II, IN2P3-CNRS
France IN2P3	Lyon	Université de Lyon, CNRS/IN2P3, Institut de Physique Nucléaire de Lyon
France IN2P3	Nantes	SUBATECH, Ecole des Mines de Nantes, Université de Nantes, IN2P3-CNRS
India	Kolkata	Saha Institute of Nuclear Physics
Italy INFN	Cagliari	Sezione INFN and Dipartimento dell'Università di Cagliari
Italy INFN	Torino	Sezione INFN and Dipartimento dell'Università di Torino
Russia	Gatchina	Petersburg Nuclear Physics Institute
South Africa	Cape Town / Somerset West	Physics Department, University of Cape Town and iThemba LABS, National Research Foundation

Table 5.22: Expression of interest to the PHOS upgrade, subject to funding

Country Funding Agency	City	Institute
China	Wuhan	Central China Normal University
Czech Republic	Prague	Institute of Physics, Academy of Sciences of the Czech Republic
Japan	Hiroshima	Hiroshima University
Norway	Bergen	University of Bergen, Department of Physics and Technology
Norway	Oslo	University of Oslo, Department of Physics
Poland	Warsaw	Soltan Institute for Nuclear Studies
Russia	Moscow	National Research Center - Kurchatov Institute
Russia	Protvino	Institute for High Energy Physics
Russia	Sarov	Russian Federal Nuclear Center
Russia	Dubna	Joint Institute for Nuclear Research

Table 5.23: Expression of interest to the upgrade of the Online Systems, subject to funding

Country Funding Agency	City	Institute
CERN	Geneva	European Organization for Nuclear Research
Croatia	Split	University of Split
Croatia	Zagreb	Rudjer Bošković Institute
Croatia	Zagreb	Zagreb University
Germany BMBF	Frankfurt	Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe-Universität
Germany BMBF	Frankfurt	Institut für Informatik, Johann Wolfgang Goethe-Universität
Hungary	Budapest	Wigner Research Centre for Physics
India	Jammu	Physics Department, Jammu University
India	Mumbai	Indian Institute of Technology (ITT) Bombay
Poland	Warsaw	University of Technology
Slovakia	Košice	Košice Technical University
Turkey	Karatay	Karatay University

Table 5.24: List of Institutes members of the ALICE Collaboration as of 15th of October 2012.

(*) Associate Institute

Country	Town	Institute	Representative
Armenia	Yerevan	Yerevan Physics Institute	A. Grigoryan
Brazil	Campinas / São Paulo	Universidade Estadual de Campinas (UNICAMP) and Universidade de São Paulo (USP)	J. Takahashi, A. Szanto de Toledo
CERN, Switzerland	Geneva	European Organization for Nuclear Research (CERN)	W. Carena
Chile	Talca (*)	Universidad de Talca, School of Bioinformatics Engineering, Centre of Bioinformatics and Molecular Simulations	S. Guinoz-Molinos
China	Beijing	China Institute of Atomic Energy	X. Li
	Wuhan	Central China Normal University (CCNU)	D. Zhou
	Wuhan (*)	Hua-Zhong University of Science & Technology (HUST)	Q. Li
Croatia	Split	Technical University of Split FESB	S. Gotovac
	Zagreb	Rudjer Bošković Institute and Zagreb University	T. Anticic
Czech Republic	Prague	Faculty of Nuclear Sciences and Physical Engineering, Czech Technical University in Prague	V. Petracek
	Prague	Institute of Physics, Academy of Sciences of the Czech Republic	P. Zavada
	Řež u Prahy	Nuclear Physics Institute, Academy of Sciences of the Czech Republic	M. Sumbera
Denmark	Copenhagen	Niels Bohr Institute, University of Copenhagen	J.-J. Gaardhøje
Egypt	Cairo (*)	Academy of Scientific Research and Technology (ASRT)	A. Tawfik
Finland	Helsinki / Jyväskylä	Helsinki Institute of Physics (HIP) and University of Jyväskylä	J. Rak
France	Clermont-Ferrand	Laboratoire de Physique Corpusculaire (LPC), Clermont Université, Université Blaise Pascal, CNRS-IN2P3	P. Dupieux
	Grenoble	Laboratoire de Physique Subatomique et de Cosmologie (LPSC), Université Joseph Fourier, CNRS-IN2P3, Institut Polytechnique de Grenoble	C. Furget
	Lyon	Centre de Calcul IN2P3	R. Vernet
	Lyon	Université de Lyon, CNRS/IN2P3, Institut de Physique Nucléaire de Lyon	R. Tieulent
	Nantes	SUBATECH, Ecole des Mines de Nantes, Université de Nantes, CNRS-IN2P3	G. Martinez-Garcia
	Orsay	Institut de Physique Nucléaire d'Orsay (IPNO), Université Paris-Sud, CNRS-IN2P3	B. Espagnon
	Saclay	Commissariat à l'Energie Atomique, IRFU	A. Baldissari
	Strasbourg	Institut Pluridisciplinaire Hubert Curien (IPHC), Université de Strasbourg, CNRS-IN2P3	C. Kuhn
Germany	Darmstadt	Institut für Kernphysik, Technische Universität Darmstadt	H. Oeschler
	Darmstadt	Research Division and ExtreMe Matter Institute EMMI, GSI Helmholtzzentrum für Schwerionenforschung	S. Masciocchi
	Frankfurt	Frankfurt Institute for Advanced Studies, Johann Wolfgang Goethe-Universität Frankfurt	V. Lindenstruth
	Frankfurt	Institut für Informatik, Johann Wolfgang Goethe-Universität, Frankfurt	U. Kebschull
	Frankfurt	Institut für Kernphysik, Johann Wolfgang Goethe-Universität Frankfurt	H. Appelshäuser
	Heidelberg	Physikalisches Institut, Ruprecht-Karls-Universität Heidelberg	J. Stachel
	Köln	Fachhochschule Köln	T. Krawutschke
	Munich (*)	Physik Department, Technische Universität München	B. Ketzer
	Munich (*)	Excellence Cluster 'Universe', Technische Universität München	L. Fabbietti
	Münster	Institut für Kernphysik, Westfälische Wilhelms-Universität Münster	C. Klein-Boesing
	Tübingen	Eberhard Karls Universität Tübingen	H.-R. Schmidt
	Worms	Zentrum für Technologietransfer und Telekommunikation (ZTT), Fachhochschule Worms	R. Keidel
Greece	Athens	Physics Department, University of Athens	M. Spyropoulou-Stassinaki
Hungary	Budapest	Wigner Research Centre for Physics, Hungarian Academy of Sciences	P. Levai
India	Aligarh	Department of Physics Aligarh Muslim University	M. Irfan
	Bhubaneswar	Institute of Physics	D. Mahapatra
	Chandigarh	Physics Department, Panjab University	M. Aggarwal
	Guwahati	Gauhati University	B. Bhattacharjee
	Indore	Indian Institute of Technology Indore (IIT)	R. Sahoo
	Jaipur	Physics Department, University of Rajasthan	S. Raniwala
	Jammu	Physics Department, University of Jammu	A. Bahsin
	Kolkata	Bose Institute	S. Raha

Country	Town	Institute	Representative
India	Kolkata	Saha Institute of Nuclear Physics	S. Chattopadhyay
	Kolkata	Variable Energy Cyclotron Centre	T. Nayak
	Mumbai (*)	Bhabha Atomic Research Center (BARC)	V. B. Chandratre
	Mumbai	Indian Institute of Technology Bombay (IIT)	B. Nandi
Italy	Alessandria	Dipartimento di Scienze e Tecnologie Avanzate dell'Università del Piemonte Orientale and Gruppo Collegato INFN	L. Ramello
	Bari	Dipartimento Interateneo di Fisica 'M. Merlin' and Sezione INFN	E. Nappi
	Bologna	Dipartimento di Fisica dell'Università and Sezione INFN	R. Nania
	Cagliari	Dipartimento di Fisica dell'Università and Sezione INFN	A. Masoni
	Catania	Dipartimento di Fisica e Astronomia dell'Università and Sezione INFN	F. Riggio
	Frascati	Laboratori Nazionali di Frascati, INFN	N. Bianchi
	Legnaro	Laboratori Nazionali di Legnaro, INFN	L. Vannucci
	Padova	Dipartimento di Fisica e Astronomia dell'Università and Sezione INFN	F. Antinori
	Rome	Dipartimento di Fisica dell'Università 'La Sapienza' and Sezione INFN	A. Mazzoni
	Rome	Centro Fermi - Centro Studi e Ricerche e Museo Storico della Fisica "Enrico Fermi"	L. Cifarelli
Japan	Salerno	Dipartimento di Fisica 'E.R. Caianiello' dell'Università and Gruppo Collegato INFN	S. De Pasquale
	Trieste	Dipartimento di Fisica dell'Università and Sezione INFN	G.-V. Margagliotti
	Turin	Dipartimento di Fisica dell'Università and Sezione INFN	M. Gallio
	Hiroshima	Hiroshima University	T. Sugitate
	Tokyo	University of Tokyo	H. Hamagaki
	Tsukuba	University of Tsukuba	Y. Miake
JINR, Russia	Wako-Shi (*)	RIKEN	H. Enyo
	Dubna	Joint Institute for Nuclear Research (JINR)	A. Vodopianov
Mexico	Culiacán	Universidad de Autónoma de Sinaloa	I. Leon Monzon
	Mexico City	Instituto de Ciencias Nucleares, Universidad Nacional Autónoma de México	G. Paic
	Mexico City	Instituto de Física, Universidad Nacional Autónoma de México	A. Menchaca Rocha
	Mérida	Centro de Investigación y de Estudios Avanzados (CINVESTAV)	G. Herrera Corral
Netherlands	Puebla	Benemérita Universidad Autónoma de Puebla	A. Fernandez Tellez
	Amsterdam / Utrecht	NIKHEF and Institute for Subatomic Physics, Utrecht University	T. Peitzmann
	Bergen / Tonsberg	Department of Physics and Technology, University of Bergen, Vestfold University College, Tonsberg	D. Roehrich
Norway	Bergen	Faculty of Engineering, Bergen University College	H. Helstrup
	Oslo	Department of Physics, University of Oslo	T. Tveten
Pakistan	Islamabad	COMSATS Institute of Information Technology (CIIT), Islamabad	A. Bhatti
Peru	Lima	Sección Física, Departamento de Ciencias, Pontificia Universidad Católica del Perú	A. Gago Medina
Poland	Cracow	The Henryk Niewodniczanski Institute of Nuclear Physics, Polish Academy of Sciences	M. Kowalski
	Warsaw	Soltan Institute for Nuclear Studies	T. Siemarczuk
	Warsaw	Warsaw University of Technology	J. Pluta
Rep. of Korea	Daejeon	Korea Institute of Science and Technology Information	H.-J. Jang
	Gangneung	Gangneung-Wonju National University	D.-W. Kim
	Pusan	Pusan National University	I.-K. Yoo
	Seoul	Department of Physics, Sejong University	S. Kim
	Seoul	Yonsei University	Y. Kwon
Romania	Bucharest	Institute of Space Sciences (ISS)	D. Felea
	Bucharest	National Institute for Physics and Nuclear Engineering	M. Petrovici
Russia	Gatchina	Petersburg Nuclear Physics Institute	V. Samsonov
	Moscow	Institute for Nuclear Research, Academy of Sciences	A. Kurepin
	Moscow	Institute for Theoretical and Experimental Physics	A. Akindinov
	Moscow	Moscow Engineering Physics Institute	V. Grigoriev
	Moscow	Russian Research Centre Kurchatov Institute	V. Manko
	Novosibirsk	Budker Institute for Nuclear Physics	Y. Pestov
	Protvino	Institute for High Energy Physics	S. Sadovski
	Sarov	Russian Federal Nuclear Center (VNIIEF)	N. Zavyalov
	St. Petersburg	V. Fock Institute for Physics, St. Petersburg State University	G. Feofilov

Country	Town	Institute	Representative
Slovakia	Bratislava	Faculty of Mathematics, Physics and Informatics, Comenius University	B. Sitar
	Košice (*) Košice	Košice Technical University Institute of Experimental Physics, Slovak Academy of Sciences and Faculty of Science, P.J. Šafárik University	J. Jadlovský L. Sandor
South Africa	Cape Town / Somerset West	Physics Department, University of Cape Town and iThemba LABS, National Research Foundation	J. Cleymans
Spain/Cuba	Madrid / Santiago de Compostela / La Havana	CIEMAT, Universidad de Santiago de Compostela and CEADEN	E. Montes Prado, C. Pajares
Sweden	Lund	Division of Experimental High Energy Physics, University of Lund	A. Oskarsson
Thailand	Nakhon Ratchasima	Suranaree University of Technology	C. Kobdaj
Turkey	Istanbul (*) Konya	Yildiz Technical University KTO Karatay University	M. Subasi A. Okatan
Ukraine	Kharkov (*)	Kharkov Institute of Physics and Technology (KIPT)	N. Maslov
	Kharkov	Scientific Research Technological Institute of Instrument Engineering	V. Borshchov
United Kingdom	Kiev	Bogolyubov Institute for Theoretical Physics	O. Borysov
	Birmingham	School of Physics and Astronomy, University of Birmingham	D. Evans
United States	Daresbury	STFC Daresbury Laboratory	R.-C. Lemmon
	Chilton (*)	STFC Rutherford Appleton Laboratory	R. Turchetta
United States	Austin, TX	The University of Texas at Austin, Physics Department	C. Markert
	Berkeley, CA	Lawrence Berkeley National Laboratory	P. Jacobs
	Chicago, IL	Chicago State University	E. Garcia Solis
	Columbus, OH	Department of Physics, Ohio State University	T. Humanic
	Columbus, OH	Ohio Supercomputer Center OSC	D. Johnson
	Detroit, MI	Wayne State University	T. Cormier
	Houston, TX	University of Houston	L. Pinsky
	Knoxville, TN	University of Tennessee	K. Read
	Livermore, CA	Lawrence Livermore National Laboratory	R. Soltz
	New Haven, CT	Yale University	H. Caines
	Oak Ridge, TN	Oak Ridge National Laboratory	T. Awes
	Omaha, NE	Physics Department, Creighton University	M. Cherney
	S. Luis Obispo, CA	California Polytechnic State University	J. Klay
	West Lafayette, IN	Purdue University	R. Scharenberg

Acronyms**A**

ADC	Analog to Digital Converter
ADULT	A DUaL Threshold ASIC
AliROOT	ALICE software framework based on ROOT
ALTR0	ALICE TPC ReadOut chip
APD	Avalanche Photo Diode
ASIC	Application Specific Integrated Circuit

C

CMOS	Complementary Metal-Oxide-Semiconductor
CPC	Cathode Pad Chamber
CPU	Central Processing Unit
CROCUS	Cluster Read Out Concentrator Unit System
CRT	ConcentRaFor board
CTTM	Central Topology Trigger Module

D

DAQ	Data Acquisition System
DARC	DAta Readout Card
DCal	Dijet Calorimeter
DDL	Detector Data Link
DMA	Direct Memory Access
DRM	Data Readout Module

E

EMCal	Electro-Magnetic Calorimeter
ENC	Equivalent Noise Charge

F

FEC	Front-End Card
FEE	Front-End Electronics
FEERIC	Front-End Electronics Rapid Integrated Circuit
FMD	Forward Multiplicity Detector
FPGA	Field Programmable Gate Array

G

GEM	Gas Electron Multiplier
GPU	Graphical Processor Unit
GP GPU	General Purpose GPU
GTU	Global Trigger Unit

H

HFT	Heavy Flavor Tracker
HI	Heavy Ion
HL-LHC	High Luminosity LHC
HLT	High-Level Trigger
HMPID	High Momentum Particle IDentification detector
HPTDC	High Performance Time to Digital Converter
HV	High Voltage
H-RORC	HLT RORC

I

ILC	International Linear Collider
ILD	International Linear collider Detector
I/O	Input–Output
ITS	Inner Tracking System
IROC	Inner ReadOut Chamber

L

L0	Level-0 (trigger)
L1	Level-1 (trigger)
LHC	Large Hadron Collider
LS	Like Sign
LS1	Long Shutdown 1
LS2	Long Shutdown 2
LTM	Local Trigger Module
LV	Low Voltage

M

MANAS	Multiplexed ANAlogic Signal processor
MANU	MANas NUmerique
MAPS	Monolithic Active Pixel Sensors
MARC	Muon Arm Readout Chip
MCM	Multi Chip Module
MIP	Minimum Ionising Particle
MoU	Memorandum of Understanding
MRPC	Multigap Resistive Plate Chamber
MWPC	Multi Wire Proportional Chamber

O

OROC	Outer ReadOut Chamber
------	-----------------------

P

PASA	PreAmplifier ShAper
PATCH	Protocol for Alice Tracking CHambers
PC	Personal Computer
PCIe	Peripheral Component Interconnect express
PHOS	PHOton Spectrometer
PID	Particle IDentification
PRF	Pad Response Function
PS	Proton Synchrotron

Q

QA	Quality Assurance
QGP	Quark-Gluon Plasma

R

RAM	Random Access Memory
RCU	Readout Control Unit
RHIC	Relativistic Heavy Ion Collider
RICH	Ring Image CHerenkov detector
ROC	Read Out Chamber
ROOT	OO software framework
RORC	Read-Out Received Card
RPC	Resistive Plate Chamber

S

SDD	Silicon Strip Detector
SPD	Silicon Pixel Detector
SSD	Silicon Strip Detector
S/N	Signal-to-Noise ratio
STAR	Solenoidal Tracker At RHIC

T

TCI	Trigger Crocus Interface
TOF	Time Of Flight detector
TPC	Time Projection Chamber
TRAP	TRAcklet Processor chip
TRD	Transition Radiation Detector
TRM	TDC Readout Module
TTC	Timing Trigger and Control
TTCrx	TTC Receiver ASIC

U

US	Unlike Sign
----	-------------

V

VME	Versa Module Eurocard
VHDL	VHSIC Hardware Description Language
VHSIC	Very-High Speed Integrated Circuit

Z

ZDC	Zero Degree Calorimeter
-----	-------------------------

References

- [1] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, Journal of Instrumentation **3**, S08002 (2008), doi:10.1088/1748-0221/3/08/S08002.
- [2] S. e. a. Borsnyi, *Is there still any mystery in lattice QCD? Results with physical masses in the continuum limit III*, Journal of High Energy Physics **2010**, 1 (2010), 10.1007/JHEP09(2010)073.
- [3] M. Capeans *et al.*, *ATLAS Insertable B-Layer Technical Design Report*, CERN-LHCC-2010-013. ATLAS-TDR-019, 2010.
- [4] CMS Collaboration, *Technical Proposal for the Upgrade of the CMS detector through 2020*, CERN-LHCC-2011-006. LHCC-P-004, 2011.
- [5] CMS Collaboration, S. Chatrchyan *et al.*, *Study of the inclusive production of charged pions, kaons, and protons in pp collisions at $\sqrt{s} = 0.9, 2.76$, and 7 TeV*, Eur.Phys.J. **C72**, 2164 (2012), doi:10.1140/epjc/s10052-012-2164-x.
- [6] ATLAS Collaboration, *dE/dx measurement in the ATLAS Pixel Detector and its use for particle identification*, (2011).
- [7] CMS Collaboration, V. Khachatryan *et al.*, *Prompt and non-prompt J/ψ production in pp collisions at $\sqrt{s} = 7$ TeV*, Eur.Phys.J. **C71**, 1575 (2011), arXiv:1011.4193, doi:10.1140/epjc/s10052-011-1575-8.
- [8] ALICE Collaboration, CERN Report No. CERN-LHCC-2012-005. LHCC-G-159, 2012 (unpublished).
- [9] J. Jowett, *LHC heavy ion programme, CERN Council Open Symposium on European Strategy for Particle Physics*, 2012.
- [10] R. Rapp and v. H. Hendrik, in *Quark-Gluon Plasma 4* (R.C. Hwa and X.N. Wang, eds.), *World Scientific (Singapore)*, Phys.Lett. **B664**, 111 (2010), arXiv:0903.1096.
- [11] M. Mangano, P. Nason and G. Ridolfi, *Heavy-quark correlations in hadron collisions at next-to-leading order*, Nuclear Physics B **373**, 295 (1992), doi:10.1016/0550-3213(92)90435-E.
- [12] K. Eskola, H. Paukkunen and C. Salgado, *EPS09: A New Generation of NLO and LO Nuclear Parton Distribution Functions*, JHEP **0904**, 065 (2009), arXiv:0902.4154, doi:10.1088/1126-6708/2009/04/065.
- [13] J. Aichelin, P. Gossiaux and T. Gousset, *Radiative and Collisional Energy Loss of Heavy Quarks in Deconfined Matter*, 1201.4192, arXiv:1201.4192.
- [14] J. Uphoff, O. Fochler, Z. Xu and C. Greiner, *Open heavy flavor in Pb+Pb collisions at $\sqrt{s} = 2.76$ TeV within a transport model*, 1205.4945, arXiv:1205.4945.
- [15] V. Greco, H. van Hees and R. Rapp, *Heavy-quark kinetics at RHIC and LHC*, 0709.4452, arXiv:0709.4452.
- [16] STAR Collaboration, J. Adams *et al.*, *Azimuthal anisotropy in Au+Au collisions at $s(NN)^{**}(1/2) = 200$ GeV*, Phys.Rev. **C72**, 014904 (2005), arXiv:nucl-ex/0409033, doi:10.1103/PhysRevC.72.014904.
- [17] B. Abelev *et al.* (STAR Collaboration), *Mass, quark-number, and \sqrt{s} dependence of the second and fourth flow harmonics in ultra-relativistic nucleus-nucleus collisions*, Phys.Rev. **C75**, 054906 (2007), arXiv:nucl-ex/0701010, doi:10.1103/PhysRevC.75.054906.
- [18] M. Krzewicki *et al.* (ALICE Collaboration), *Elliptic and triangular flow of identified particles at ALICE*, 1107.0080, arXiv nucl-ex/1107.0080, 2011., arXiv:1107.0080.
- [19] V. Greco, C. Ko and R. Rapp, *Quark coalescence for charmed mesons in ultrarelativistic heavy-ion collisions*, Physics Letters B **595**, 202 (2004), doi:10.1016/j.physletb.2004.06.064.
- [20] F. Noferini *et al.* (ALICE Collaboration), *Quark Matter 2012*, Washington D.C, August 2012, http://qm2012.bnl.gov, [link].
- [21] I. Belikov *et al.* (ALICE Collaboration), *K_s^0 and Λ production in Pb-Pb collisions with the ALICE experiment*, 1109.4807, arXiv hep-ex/1109.4807, 2011., arXiv:1109.4807.
- [22] M. Aggarwal *et al.* (STAR Collaboration), *Strange and Multi-strange Particle Production in Au+Au Collisions at $\sqrt{s_{NN}} = 62.4$ GeV*, Phys.Rev. **C83**, 024901 (2011), arXiv:1010.0142, doi:10.1103/PhysRevC.83.024901.
- [23] M. He, R. J. Fries and R. Rapp, *Heavy-Quark Diffusion and Hadronization in Quark-Gluon Plasma*, Phys.Rev. **C86**, 014903 (2012), arXiv:1106.6006, doi:10.1103/PhysRevC.86.014903.
- [24] Y. Oh, C. M. Ko, S. H. Lee and S. Yasui, *Heavy baryon/meson ratios in relativistic heavy ion collisions*,

Phys.Rev. **C79**, 044905 (2009), arXiv:0901.1382, doi:10.1103/PhysRevC.79.044905.

[25] S. Lee, K. Ohnishi, S. Yasui, I. Yoo and C. Ko, Λ_c Enhancement from Strongly Coupled Quark-Gluon Plasma, Phys. Rev. Lett. **100**, 222301 (2008), doi:10.1103/PhysRevLett.100.222301.

[26] I. Kuznetsova and J. Rafelski, Heavy flavor hadrons in statistical hadronization of strangeness-rich QGP, Eur.Phys.J. **C51**, 113 (2007), arXiv:hep-ph/0607203, doi:10.1140/epjc/s10052-007-0268-9.

[27] M. He, R. J. Fries and R. Rapp, D_s -Meson as Quantitative Probe of Diffusion and Hadronization in Nuclear Collisions, 1204.4442, arXiv:1204.4442.

[28] M. He, R. J. Fries and R. Rapp, Thermal Relaxation of Charm in Hadronic Matter, Phys.Lett. **B701**, 445 (2011), arXiv:1103.6279, doi:10.1016/j.physletb.2011.06.019.

[29] G. Innocenti et al. (ALICE Collaboration), Quark Matter 2012, Washington D.C, August 2012, http://qm2012.bnl.gov, [link].

[30] B.-W. Zhang, C.-M. Ko and W. Liu, Thermal charm production in a quark-gluon plasma in Pb-Pb collisions at $s^{**}(1/2)/(NN) = 5.5$ TeV, Phys.Rev. **C77**, 024901 (2008), 8 pages, 8 figures, arXiv:0709.1684, doi:10.1103/PhysRevC.77.024901.

[31] J. Uphoff, O. Fochler, Z. Xu and C. Greiner, Heavy quark production at RHIC and LHC within a partonic transport model, Phys.Rev. **C82**, 044906 (2010), arXiv:1003.4200, doi:10.1103/PhysRevC.82.044906.

[32] P. Levai and R. Vogt, Thermal charm production by massive gluons and quarks, Phys.Rev. **C56**, 2707 (1997), 19 pages in Latex and 9 Postscript figures Report-no: LBNL-39984, arXiv:hep-ph/9704360, doi:10.1103/PhysRevC.56.2707.

[33] P. Braun-Munzinger and K. Redlich, Charmonium production from the secondary collisions at LHC energy, Eur.Phys.J. **C16**, 519 (2000), ps file 11 p, arXiv:hep-ph/0001008, doi:10.1007/s100520000356.

[34] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, Evidence for charmonium generation at the phase boundary in ultra-relativistic nuclear collisions, Phys.Lett. **B652**, 259 (2007), arXiv:nucl-th/0701079, doi:10.1016/j.physletb.2007.07.036.

[35] ALICE Collaboration, B. Abelev et al., J/ψ production at low transverse momentum in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, 1202.1383, arXiv:1202.1383.

[36] ALICE Collaboration, CERN Report No. CERN-LHCC-2012-013. LHCC-P-005, 2012 (unpublished).

[37] ALICE Collaboration, Conceptual Design Report for the Upgrade of the ALICE ITS, CDR-0, version released on the 6th of March (<http://aliceinfo.cern.ch/ITSUpgrade/node/9>), 2012.

[38] ALICE Collaboration, B. Abelev et al., Suppression of high transverse momentum D mesons in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV, 1203.2160, arXiv:1203.2160.

[39] M. Cacciari, M. Greco and P. Nason, The p_t spectrum in heavy flavor hadroproduction, JHEP **9805**, 007 (1998), arXiv:hep-ph/9803400.

[40] K. Nakamura et al., Review of Particle Physics, Journal of Physics G **37** (2010).

[41] ALICE Collaboration, B. Abelev et al., Measurement of charm production at central rapidity in proton-proton collisions at $\sqrt{s} = 7$ TeV, JHEP **1201**, 128 (2012), arXiv:1111.1553, doi:10.1007/JHEP01(2012)128.

[42] A. Abulencia et al. (CDF Collaboration), Measurement of the ratios of branching fractions $B(B_s^0 \rightarrow D_s^-\pi^+)/B(B^0 \rightarrow D^-\pi^+)$ and $B(B^+ \rightarrow \bar{D}^0\pi^+)/B(B^0 \rightarrow D^-\pi^+)$, Phys.Rev.Lett. **96**, 191801 (2006), arXiv:hep-ex/0508014, doi:10.1103/PhysRevLett.96.191801.

[43] S. Ricciardi et al. (LHCb Collaboration), Studies of b -hadron decays to charming final states at LHCb, 1110.3249, arXiv hep-ex/1110.3249, 2011., arXiv:1110.3249.

[44] K. Hagiwara et al., Review of Particle Physics, Physical Review D **66**, 010001+ (2002).

[45] ALICE Collaboration, B. Abelev et al., D_s meson production at central rapidity in proton-proton collisions at $\sqrt{s} = 7$ TeV, 1208.1948, arXiv:1208.1948.

[46] T. Sjöstrand, S. Mrenna and P. Skands, PYTHIA 6.4 Physics and Manual, JHEP **05**, 026 (2006), arXiv:hep-ph/0603175.

[47] for the ALICE Collaboration, C. Bianchin, Measurement of $D^0 v_2$ in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with ALICE at the LHC, 1111.6886, arXiv:1111.6886.

[48] Y. Dokshitzer and D. Kharzeev, Heavy-quark colorimetry of QCD matter, Phys. Lett. B **519**, 199 (2001), doi:10.1016/S0370-2693(01)01130-3.

[49] N. Armesto, C. Salgado and U. Wiedemann, *Medium-induced gluon radiation off massive quarks fills the dead cone*, Phys. Rev. D **69**, 114003 (2004), doi:10.1103/PhysRevD.69.114003.

[50] S. Wicks, W. Horowitz, M. Djordjevic and M. Gyulassy, *Heavy quark jet quenching with collisional plus radiative energy loss and path length fluctuations*, Nucl.Phys. **A783**, 493 (2007), arXiv:nucl-th/0701063, doi:10.1016/j.nuclphysa.2006.11.102.

[51] N. Armesto, A. Dainese, C. Salgado and U. Wiedemann, *Testing the color charge and mass dependence of parton energy loss with heavy-to-light ratios at BNL RHIC and CERN LHC*, Phys. Rev. D **71**, 054027 (2005), doi:10.1103/PhysRevD.71.054027.

[52] W. Horowitz and M. Gyulassy, *Heavy quark jet tomography of Pb + Pb at LHC: AdS/CFT drag or pQCD energy loss?*, Phys.Lett. **B666**, 320 (2008), arXiv:0706.2336, doi:10.1016/j.physletb.2008.04.065.

[53] J. Albacete, Y. Kovchegov and A. Taliotis, *Modeling Heavy Ion Collisions in AdS/CFT*, JHEP **0807**, 100 (2008), arXiv:0805.2927, doi:10.1088/1126-6708/2008/07/100.

[54] ALICE Collaboration, B. Abelev *et al.*, *Production of muons from heavy flavour decays at forward rapidity in pp and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, 1205.6443, arXiv:1205.6443.

[55] CMS Collaboration, S. Chatrchyan *et al.*, *Suppression of non-prompt J/ ψ , prompt J/ ψ , and Y(1S) in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, JHEP **1205**, 063 (2012), arXiv:1201.5069, doi:10.1007/JHEP05(2012)063.

[56] R. Averbeck *et al.*, *Reference Heavy Flavour Cross Sections in pp Collisions at $\sqrt{s} = 2.76$ TeV, using a pQCD-Driven \sqrt{s} -Scaling of ALICE Measurements at $\sqrt{s} = 7$ TeV*, 1107.3243, arXiv:1107.3243.

[57] T. Matsui and H. Satz, *J/ ψ Suppression by Quark-Gluon Plasma Formation*, Phys.Lett. **B178**, 416 (1986), doi:10.1016/0370-2693(86)91404-8.

[58] L. Kluberg and H. Satz, *Color Deconfinement and Charmonium Production in Nuclear Collisions*, 0901.3831, arXiv:0901.3831.

[59] F. Karsch, D. Kharzeev and H. Satz, *Sequential charmonium dissociation*, Phys.Lett. **B637**, 75 (2006), arXiv:hep-ph/0512239, doi:10.1016/j.physletb.2006.03.078.

[60] NA50 Collaboration, B. Alessandro *et al.*, *ψ' production in Pb-Pb collisions at 158 GeV/nucleon*, Eur.Phys.J. **C49**, 559 (2007), arXiv:nucl-ex/0612013, doi:10.1140/epjc/s10052-006-0153-y.

[61] NA60 Collaboration, R. Arnaldi *et al.*, *J/ ψ production in proton-nucleus collisions at 158 and 400 GeV*, Phys.Lett. **B706**, 263 (2012), arXiv:1004.5523, doi:10.1016/j.physletb.2011.11.042.

[62] H. Sorge, E. V. Shuryak and I. Zahed, *ψ'/ψ ratio in nucleus-nucleus collisions: A Measure for the chiral symmetry restoration temperature?*, Phys.Rev.Lett. **79**, 2775 (1997), arXiv:hep-ph/9705329, doi:10.1103/PhysRevLett.79.2775.

[63] P. Braun-Munzinger and J. Stachel, *(Non)thermal aspects of charmonium production and a new look at J/ ψ suppression*, Phys.Lett. **B490**, 196 (2000), arXiv:nucl-th/0007059, doi:10.1016/S0370-2693(00)00991-6.

[64] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, *Statistical hadronization of heavy quarks in ultra-relativistic nucleus-nucleus collisions*, Nucl.Phys. **A789**, 334 (2007), arXiv:nucl-th/0611023, doi:10.1016/j.nuclphysa.2007.02.013.

[65] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, *Evidence for charmonium generation at the phase boundary in ultra-relativistic nuclear collisions*, Phys.Lett. **B652**, 259 (2007), arXiv:nucl-th/0701079, doi:10.1016/j.physletb.2007.07.036.

[66] P. Braun-Munzinger and J. Stachel, *Charmonium from Statistical Hadronization of Heavy Quarks: A Probe for Deconfinement in the Quark-Gluon Plasma*, 0901.2500, arXiv:0901.2500.

[67] A. Andronic, P. Braun-Munzinger, K. Redlich and J. Stachel, *The thermal model on the verge of the ultimate test: particle production in Pb-Pb collisions at the LHC*, J.Phys.G **G38**, 124081 (2011), arXiv:1106.6321, doi:10.1088/0954-3899/38/12/124081.

[68] R. L. Thews, M. Schroedter and J. Rafelski, *Enhanced J/ ψ production in deconfined quark matter*, Phys.Rev. **C63**, 054905 (2001), arXiv:hep-ph/0007323, doi:10.1103/PhysRevC.63.054905.

[69] Y.-p. Liu, Z. Qu, N. Xu and P.-f. Zhuang, *J/ ψ Transverse Momentum Distribution in High Energy Nuclear Collisions at RHIC*, Phys.Lett. **B678**, 72 (2009), arXiv:0901.2757, doi:10.1016/j.physletb.2009.06.006.

[70] X. Zhao and R. Rapp, *Medium Modifications and Production of Charmonia at LHC*, Nucl.Phys. **A859**, 114

(2011), arXiv:1102.2194, doi:10.1016/j.nuclphysa.2011.05.001.

[71] Atlas Collaboration, G. Aad *et al.*, *Measurement of the centrality dependence of J/ψ yields and observation of Z production in lead-lead collisions with the ATLAS detector at the LHC*, Phys.Lett. **B697**, 294 (2011), arXiv:1012.5419, doi:10.1016/j.physletb.2011.02.006.

[72] ALICE Collaboration, B. Abelev *et al.*, *Suppression of high transverse momentum D mesons in central Pb - Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, 1203.2160, arXiv:1203.2160.

[73] PHENIX Collaboration, A. Adare *et al.*, *J/ψ Production vs Centrality, Transverse Momentum, and Rapidity in Au + Au Collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys.Rev.Lett. **98**, 232301 (2007), arXiv:nucl-ex/0611020, doi:10.1103/PhysRevLett.98.232301.

[74] PHENIX Collaboration, A. Adare *et al.*, *J/ψ suppression at forward rapidity in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys.Rev. **C84**, 054912 (2011), arXiv:1103.6269, doi:10.1103/PhysRevC.84.054912.

[75] ALICE Collaboration, B. Abelev *et al.*, *J/ψ suppression at forward rapidity in Pb - Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys.Rev.Lett. **109**, 072301 (2012), arXiv:1202.1383.

[76] ALICE collaboration, C. Suire, *Charmonia production in ALICE*, Hard Probes (2012), arXiv:1208.5601.

[77] ALICE collaboration, J. Wiechula, *Nuclear modification of J/ψ production at central rapidity in Pb - Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Hard Probes 2012 (2012), arXiv:1208.5401.

[78] CMS Collaboration, S. Chatrchyan *et al.*, *Measurement of the pseudorapidity and centrality dependence of the transverse energy density in Pb Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, 1205.2488, arXiv:1205.2488.

[79] ALICE collaboration, L. Massacrier, *J/ψ elliptic flow measurement in Pb - Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV at forward rapidity with the ALICE experiment*, Hard Probes (2012), arXiv:1208.5401.

[80] STAR Collaboration, Z. Tang, *J/ψ production and correlation in p - p and Au + Au collisions at STAR*, J.Phys.G **G38**, 124107 (2011), arXiv:1107.0532, doi:10.1088/0954-3899/38/12/124107.

[81] M. Laine, *News on hadrons in a hot medium*, 1108.5965, arXiv:1108.5965.

[82] CMS Collaboration, S. Chatrchyan *et al.*, *Indications of suppression of excited Υ states in Pb Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys.Rev.Lett. **107**, 052302 (2011), arXiv:1105.4894, doi:10.1103/PhysRevLett.107.052302.

[83] CMS Collaboration, S. Chatrchyan *et al.*, *Observation of sequential Υ suppression in Pb Pb collisions*, 1208.2826, arXiv:1208.2826.

[84] CMS Collaboration, S. Chatrchyan *et al.*, *Measurement of the $\psi(2S)$ meson in Pb Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Analysis Note CMS-PAS-HIN-12-007 (2012).

[85] ALICE collaboration, R. Arnaldi, *J/ψ production at forward rapidity in Pb - Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Quark Matter (2012).

[86] A. Andronic, F. Beutler, P. Braun-Munzinger, K. Redlich and J. Stachel, *Statistical hadronization of heavy flavor quarks in elementary collisions: Successes and failures*, Phys.Lett. **B678**, 350 (2009), arXiv:0904.1368, doi:10.1016/j.physletb.2009.06.051.

[87] P. Faccioli, C. Lourenco, J. Seixas and H. K. Wohri, *Towards the experimental clarification of quarkonium polarization*, Eur.Phys.J. **C69**, 657 (2010), arXiv:1006.2738, doi:10.1140/epjc/s10052-010-1420-5.

[88] ALICE Collaboration, B. Abelev *et al.*, *J/ψ polarization in p - p collisions at $\sqrt{s}=7$ TeV*, Phys.Rev.Lett. **108**, 082001 (2012), arXiv:1111.1630, doi:10.1103/PhysRevLett.108.082001.

[89] B. Ioffe and D. Kharzeev, *Quarkonium polarization in heavy ion collisions as a possible signature of the quark gluon plasma*, Phys.Rev. **C68**, 061902 (2003), arXiv:hep-ph/0306176, doi:10.1103/PhysRevC.68.061902.

[90] V. Rebyakova, M. Strikman and M. Zhalov, *Coherent ρ and J/ψ photoproduction in ultraperipheral processes with electromagnetic dissociation of heavy ions at RHIC and LHC*, Phys.Lett. **B710**, 647 (2012), arXiv:1109.0737, doi:10.1016/j.physletb.2012.03.041.

[91] J. Nystrand, ReteQuarkonii Workshop 2010 (2010).

[92] R. Rapp, J. Wambach and H. van Hees, *The Chiral Restoration Transition of QCD and Low Mass Dileptons*, 0901.3289, arXiv:0901.3289.

[93] PHENIX Collaboration, A. e. a. Adare, *Enhanced Production of Direct Photons in Au + Au Collisions at $\sqrt{s_{NN}} = 200$ GeV and Implications for the Initial Temperature*, Phys. Rev. Lett. **104**, 132301 (2010), doi:10.1103/PhysRevLett.104.132301.

[94] R. Rapp, *Theory of Soft Electromagnetic Emission in Heavy-Ion Collisions*, Acta Phys.Polon. **B42**, 2823 (2011), arXiv:1110.4345, doi:10.5506/APhysPolB.42.2823.

[95] R. Rapp and J. Wambach, *Low mass dileptons at the CERN SPS: Evidence for chiral restoration?*, Eur.Phys.J. **A6**, 415 (1999), arXiv:hep-ph/9907502, doi:10.1007/s100500050364.

[96] NA60 Collaboration, R. e. a. Arnaldi, *Evidence for the production of thermal-like muon pairs with masses above 1-GeV/c^{**2} in 158-A-GeV Indium-Indium Collisions*, Eur.Phys.J. **C59**, 607 (2009), arXiv:0810.3204, doi:10.1140/epjc/s10052-008-0857-2.

[97] NA60 Collaboration, R. e. a. Arnaldi, *NA60 results on thermal dimuons*, Eur.Phys.J. **C61**, 711 (2009), arXiv:0812.3053, doi:10.1140/epjc/s10052-009-0878-5.

[98] NA60 Collaboration, S. Damjanovic, R. Shahoyan and H. J. Specht, *NA60: In hot pursuit of thermal dileptons*, CERN Cour. **49N9**, 31 (2009).

[99] H. van Hees and R. Rapp, *Dilepton Radiation at the CERN Super Proton Synchrotron*, Nucl.Phys. **A806**, 339 (2008), arXiv:0711.3444, doi:10.1016/j.nuclphysa.2008.03.009.

[100] S. Weinberg, *Precise Relations between the Spectra of Vector and Axial-Vector Mesons*, Phys. Rev. Lett. **18**, 507 (1967), doi:10.1103/PhysRevLett.18.507.

[101] T. Das, V. S. Mathur and S. Okubo, *Low-Energy Theorem in the Radiative Decays of Charged Pions*, Phys. Rev. Lett. **19**, 859 (1967), doi:10.1103/PhysRevLett.19.859.

[102] M. A. Shifman, A. Vainshtein and V. I. Zakharov, *QCD and Resonance Physics. Sum Rules*, Nucl.Phys. **B147**, 385 (1979), doi:10.1016/0550-3213(79)90022-1.

[103] R. Chatterjee, D. K. Srivastava, U. Heinz and C. Gale, *Elliptic flow of thermal dileptons in relativistic nuclear collisions*, Phys. Rev. C **75**, 054909 (2007), doi:10.1103/PhysRevC.75.054909.

[104] ALICE Collaboration, B. Abelev, *Measurement of charm production at central rapidity in proton-proton collisions at sqrt(s) = 2.76 TeV*, 1205.4007, arXiv:1205.4007.

[105] CERN Report No. ATL-PHYS-PUB-2011-012, 2011 (unpublished).

[106] *Prompt charm production in pp collisions at $\sqrt{s} = 7$ TeV*, (2010).

[107] PHENIX Collaboration, A. Adare *et al.*, *Detailed measurement of the e^+e^- pair continuum in $p + p$ and $Au+Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV and implications for direct photon production*, Phys.Rev. **C81**, 034911 (2010), arXiv:0912.0244, doi:10.1103/PhysRevC.81.034911.

[108] ALICE Collaboration, B. e. a. Abelev, *Neutral pion and η meson production in proton-proton collisions at $\sqrt{s} = 0.9$ TeV and $\sqrt{s} = 7$ TeV*, 1205.5724, arXiv:1205.5724.

[109] STAR, C. Adler *et al.*, *Disappearance of back-to-back high p_T hadron correlations in central $Au+Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. Lett. **90**, 082302 (2003), arXiv:nucl-ex/0210033, doi:10.1103/PhysRevLett.90.082302.

[110] PHENIX, S. S. Adler *et al.*, *Suppressed π^0 Production at Large Transverse Momentum in Central $Au+Au$ Collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. Lett. **91**, 072301 (2003), arXiv:nucl-ex/0304022, doi:10.1103/PhysRevLett.91.072301.

[111] T. Renk, *Constraining the Physics of Jet Quenching*, Phys.Rev. **C85**, 044903 (2012), arXiv:1112.2503.

[112] J. Allen *et al.*, *ALICE DCal: An Addendum to the EMCal Technical Design Report Di-Jet and Hadron-Jet correlation measurements in ALICE*, (2010).

[113] M. Cacciari, G. P. Salam and G. Soyez, *FastJet User Manual*, Eur.Phys.J. **C72**, 1896 (2012), arXiv:1111.6097, doi:10.1140/epjc/s10052-012-1896-2.

[114] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, Phys. Lett. **B659**, 119 (2008), arXiv:0707.1378, doi:10.1016/j.physletb.2007.09.077.

[115] ALICE, B. Abelev *et al.*, *Measurement of Event Background Fluctuations for Charged Particle Jet Reconstruction in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, JHEP **1203**, 053 (2012), arXiv:1201.2423.

[116] T. Sjostrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178**, 852 (2008), arXiv:0710.3820, doi:10.1016/j.cpc.2008.01.036.

[117] Atlas Collaboration, G. Aad *et al.*, *Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS Detector at the LHC*, Phys.Rev.Lett. **105**, 252303 (2010), arXiv:1011.6182.

[118] CMS Collaboration, S. Chatrchyan *et al.*, *Observation and studies of jet quenching in PbPb collisions at nucleon-nucleon center-of-mass energy = 2.76 TeV*, Phys. Rev. **C84**, 024906 (2011), arXiv:1102.1957,

doi:10.1103/PhysRevC.84.024906.

[119] CMS Collaboration, S. Chatrchyan *et al.*, *Measurement of jet fragmentation into charged particles in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, 1205.5872, arXiv:1205.5872.

[120] *Measurement of inclusive jet charged particle fragmentation functions in Pb+Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector*, (2012).

[121] *Detailed Characterization of Jets in Heavy Ion Collisions Using Jet Shapes and Jet Fragmentation Functions*, (2012).

[122] ALICE, B. Alessandro *et al.*, *ALICE: Physics performance report, volume II*, J. Phys. **G32**, 1295 (2006), doi:10.1088/0954-3899/32/10/001.

[123] PHENIX, A. Adare *et al.*, *Enhanced production of direct photons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. Lett. **104**, 132301 (2010), arXiv:0804.4168, doi:10.1103/PhysRevLett.104.132301.

[124] PHENIX, S. S. Adler *et al.*, *Centrality Dependence of Direct Photon Production in $\sqrt{s_{NN}} = 200$ GeV Au+Au collisions*, Phys. Rev. Lett. **94**, 232301 (2005), arXiv:nucl-ex/0503003, doi:10.1103/PhysRevLett.94.232301.

[125] CMS Collaboration, S. Chatrchyan *et al.*, *Measurement of isolated photon production in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, 1201.3093, arXiv:1201.3093.

[126] CMS Collaboration, S. Chatrchyan *et al.*, *Study of Z boson production in PbPb collisions at nucleon-nucleon centre of mass energy = 2.76 TeV*, Phys. Rev. Lett. **106**, 212301 (2011), arXiv:1102.5435, doi:10.1103/PhysRevLett.106.212301.

[127] PHENIX, A. Adare *et al.*, *High p_T direct photon and π^0 triggered azimuthal jet correlations and measurement of k_T for isolated direct photons in p + p collisions at $\sqrt{s} = 200$ GeV*, Phys. Rev. **D82**, 072001 (2010), arXiv:arXiv:1006.1347, doi:10.1103/PhysRevD.82.072001.

[128] PHENIX Collaboration, N. Grau, *Probing nuclear matter with jets and gamma-hadron correlations: Results from PHENIX*, J. Phys. G **G38**, 124090 (2011).

[129] U. Abeysekara *et al.*, *ALICE EMCAL Physics Performance Report*, 1008.0413, arXiv:1008.0413.

[130] ALICE, B. Abelev *et al.*, *Neutral pion and η meson production in proton-proton collisions at $\sqrt{s} = 0.9$ TeV and $\sqrt{s} = 7$ TeV*, 1205.5724, arXiv:1205.5724.

[131] CMS Collaboration, S. Chatrchyan *et al.*, *Study of high- p_T charged particle suppression in PbPb compared to pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Eur. Phys. J. **C72**, 1945 (2012), arXiv:1202.2554.

[132] ALICE, K. Aamodt *et al.*, *Suppression of Charged Particle Production at Large Transverse Momentum in Central Pb-Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Lett. **B696**, 30 (2011), arXiv:1012.1004, doi:10.1016/j.physletb.2010.12.020.

[133] ALICE, K. Aamodt *et al.*, *The ALICE experiment at the CERN LHC*, JINST **0803**, S08002 (2008), doi:10.1088/1748-0221/3/08/S08002.

[134] Y. L. Dokshitzer and D. E. Kharzeev, *Heavy quark colorimetry of QCD matter*, Phys. Lett. **B519**, 199 (2001), arXiv:hep-ph/0106202, doi:10.1016/S0370-2693(01)01130-3.

[135] PHENIX, S. S. Adler *et al.*, *Nuclear modification of electron spectra and implications for heavy quark energy loss in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. Lett. **96**, 032301 (2006), arXiv:nucl-ex/0510047, doi:10.1103/PhysRevLett.96.032301.

[136] ATLAS Collaboration, G. Aad *et al.*, *Measurement of $D^{*+/-}$ meson production in jets from pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector CERN-LHC*, Phys. Rev. **D85**, 052005 (2012), arXiv:1112.4432, doi:10.1103/PhysRevD.85.052005.

[137] CMS Collaboration, V. Khachatryan *et al.*, *Measurement of $B\bar{B}$ Angular Correlations based on Secondary Vertex Reconstruction at $\sqrt{s} = 7$ TeV*, JHEP **1103**, 136 (2011), arXiv:1102.3194, doi:10.1007/JHEP03(2011)136.

[138] A. Andronic, P. Braun-Munzinger, J. Stachel and H. Stocker, *Production of light nuclei, hypernuclei and their antiparticles in relativistic nuclear collisions*, Phys. Lett. **B697**, 203 (2011), arXiv:1010.2995, doi:10.1016/j.physletb.2011.01.053.

[139] J. Schaffner-Bielich, R. Mattiello and H. Sorge, *Dibaryons with strangeness: Their weak nonleptonic decay using $SU(3)$ symmetry and how to find them in relativistic heavy ion collisions*, Phys. Rev. Lett. **84**, 4305 (2000), arXiv:nucl-th/9908043, doi:10.1103/PhysRevLett.84.4305.

[140] J. Haidenbauer and U.-G. Mei [*LATIN SMALL LETTER SHARP S; e*] ner, *The proton-antiproton mass threshold structure in $\psi(3686)$ radiative decay revisited*, 1208.3343, arXiv:1208.3343.

[141] ALICE Collabortaion, ALICE Collaboration, *Overview of recent ALICE results*, 2012.

[142] E. Crescio et al., *Performance of the ITS stand-alone tracker in pp collisions*, ALICE-INT-2009-046, 2009.

[143] PHOBOS Collaboration, *Identified hadron transverse momentum spectra in Au+Au collisions at $\sqrt{s_{NN}} = 62.4$ GeV*, Phys. Rev. C **75**, 024910 (2007), doi:10.1103/PhysRevC.75.024910.

[144] STAR Collaboration, *The STAR Heavy Flavor Tracker, Conceptual Design Report*, submitted to DOE at CD1 review, 2009.

[145] ATLAS Pixel Collaboration, *ATLAS pixel detector: Running experience*, PoS VERTEX 2009 001.

[146] M. Deveaux et al., *Radiation tolerance of a column parallel CMOS sensor with high resistivity epitaxial layer*, Journal of Instrumentation **6** (2011), doi:10.1088/1748-0221/6/02/C02004.

[147] L. Greiner et al., *A MAPS based vertex detector for the STAR experiment at RHIC*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **650**, 68 (2011), doi:10.1016/j.nima.2010.12.006.

[148] M. Stanitzki, *Advanced monolithic active pixel sensors for tracking, vertexing and calorimetry with full CMOS capability*, Nucl. Instr. and Meth. A (2010), doi:10.1016/j.nima.2010.11.166.

[149] ALICE TPC Collaboration, CERN Report No. CERN-LHCC-2000-001, 2000 (unpublished).

[150] J. Alme et al., *The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events*, Nuclear Instruments and Methods in Physics Research **622**, 316 (2010), doi:10.1016/j.nima.2010.04.042.

[151] S. Rossegger, B. Schnizer and W. Riegler, *Analytical solutions for space charge fields in TPC drift volumes*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **632**, 52 (2011), doi:DOI: 10.1016/j.nima.2010.12.213.

[152] J. Lehraus, *Progress in particle identification by ionization sampling*, Nucl. Instrum. Meth. Phys. Res. **217** (1983) 43.

[153] W. W. M. Allison and J. H. Cobb, *Relativistic Charged Particle Identification by Energy Loss*, Annual Review of Nuclear and Particle Science **30**, 253 (1980), doi:10.1146/annurev.ns.30.120180.001345.

[154] ILD Concept Group - Linear Collider Collaboration, T. Abe et al., *The International Large Detector: Letter of Intent*, 1006.3396, arXiv:1006.3396.

[155] F. et al., *A Continuously Running High-Rate GEM-TPC for*, Nuclear Physics B - Proceedings Supplements **215**, 278 (2011), *Proceedings of the 12th Topical Seminar on Innovative Particle and Radiation Detectors (IPRD10)*; doi:10.1016/j.nuclphysbps.2011.04.030.

[156] C. Collaboration, CERN Report No. CERN-LHCC-2011-006, LHCC-P-004, 2011 (unpublished).

[157] M. Villa et al., *Progress on large area GEMs*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **628**, 182 (2011), *VCI 2010*; *Proceedings of the 12th International Vienna Conference on Instrumentation*; doi:10.1016/j.nima.2010.06.312.

[158] M. Altunbas et al., *Construction, test and commissioning of the triple-GEM tracking detector for COM-PASS*, Nucl.Instrum.Meth. **A490**, 177 (2002), doi:10.1016/S0168-9002(02)00910-5.

[159] P. Aspell et al., *Super-ALTRO 16: a Front-End System on Chip for DSP based Readout of Gaseous Detectors*, IEEE Trans. Nucl. Sci., ID: TNS-00459-2012 (2012).

[160] A. Akindinov et al., Nucl. Instr. and Methods A **490** (2002) 58 .

[161] A. Akindinov et al., Nucl. Instr. and Methods A **579** (2007) 979 .

[162] J. Christiansen, *HPTDC High Performance Time to Digital Converter, Version 2.2 for HPTDC version 1.3*.

[163] P. Antonioli et al., *A first study of the TOF electronics upgrade*, ALICE-INT-2012-007.

[164] F. Bossu, PoS (RPC2012) 059 (2012).

[165] R. Arnaldi et al., Nucl. Phys. Proc. Suppl. **158** (2006) 149 .

[166] R. Arnaldi et al., IEEE Trans. Nucl. Sci. **52** (2005) 1176 (2005).

[167] F. Giannini et al., NIM **A432** (1999) 440 .

[168] M. Abbrescia et al., NIM **A456** (2000) 143 .

[169] P. Courtat et al., *The electronics of ALICE dimuon tracking chambers*, ALICE-INT-2004-026 .

[170] ALICE collaboration - *Technical Design Report of the Dimuon Forward Spectrometer*, CERN/LHCC 99-22 .

[171] *ALICE collaboration - Addendum to the Technical Design Report of the Dimuon Forward Spectrometer*, CERN/LHCC 2000-046 .

[172] G. Moore, *Cramming More Components onto Integrated Circuits*, Electronics Magazine **38** (1965).

[173] G. Moore, *Progress in Digital Integrated Electronics*, IEEE, IEDM Tech Digest , 11 (1975).

[174] G. Moore, *Excerpts from a Conversation with Gordon Moore: Moore's Law*, Video Excerpts, 2005, Intel Corporation.

[175] A. Collaboration, *The ALICE Trigger, Data Acquisition, High-Level Trigger, and Control System Technical Design Report* (CERN-LHCC-2003-062, 2003).

[176] Intel, *Many Integrated Core (MIC) Architecture - Advanced*, <http://www.intel.com/content/www/us/en/architecture-and-technology/many-integrated-core/intel-many-integrated-core-architecture.html>.

[177] Wikipedia, *General-purpose computing on graphics processing units*, <http://en.wikipedia.org/wiki/GPGPU>.

[178] T. Instruments, *DaVinci video processor*, http://www.ti.com/lsds/ti/dsp/platform/davinci/whats_new.page.

[179] OpenMP, *The OpenMP API specification for parallel programming*, <http://openmp.org/wp>.

[180] OpenACC, *Directives for Accelerators*, <http://www.openacc-standard.org>.

[181] C. Leiserson, *Fat-trees: universal networks for hardware-efficient supercomputing*, IEEE Trans. Comput. **34**, 892 (1985).

[182] N. J. Lippis, *A simpler data center fabric emerges for the age of massively scalable data centers*.

[183] Mellanox, *IS5200 216-Port InfiniBand Chassis Switch*, <http://www.mellanox.com>.

[184] Mellanox, *SX6025 36-port 56Gb/s InfiniBand Switch Systems*, <http://www.mellanox.com>.

[185] Mellanox, *IS5600 648-Port InfiniBand Chassis Switch*, <http://www.mellanox.com>.

[186] Dell, *Dell Force10 Z9000 Data Center Core Switch*, <http://www.dell.com>.

[187] C. Clos, *A study of non-blocking switching networks*, Bell System Technical Journal **32**, 406 (1953).

[188] ALICE, *The CERN Analysis Facility*, <http://alisoft.cern.ch>, 2012.

[189] *The Parallel ROOT Facility*, <http://root.cern.ch/proof>, 2012.

[190] S. Lohn, *Parallelization of the AliRoot event reconstruction by performing a semi-automatic source-code transformation*, in *Proceedings of the Computing in High Energy Physics (CHEP12)*, NY City, New York, USA, 2012.

[191] V. A. Chris Lattner, *LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation*, in *Proceedings of the International Symposium on Code Generation and Optimization (CGO'04)*, Palo Alto, California, 2004.

[192] J. A. X. Dong, G. Cooperman, *Multithreaded Geant4: Semi-automatic Transformation into Scalable Thread-Parallel Software*, 2010, Euro-Par.

[193] M. Tadel and F. Carminati, *Parallelization of ALICE simulation - a jump through the looking-glass*, J. Phys.: Conf. Ser. **219 032024** (2010), Proceedings of the 17th International Conference on Computing in High Energy and Nuclear Physics (CHEP09).

[194] M. Al-Turany, *Optimization of HEP codes on GPUs*, The European Physical Journal Plus **126**, 1 (2011), Springer Berlin / Heidelberg.

[195] ALICE Collaboration, *Memorandum of Understanding for Collaboration in the Construction of the ALICE Detector*, CERN RRB-D 00-14 (2000).

[196] ALICE Collaboration, *Memorandum of Understanding for Maintenance and Operation of the ALICE Detector*, CERN-RRB-2002-034 (2002).

Acknowledgements

The Collaboration wishes to thank the following persons for their contribution to the preparation of this LOI:
M. He, R. J. Fries and R. Rapp, Cyclotron Institute and Department of Physics and Astronomy, Texas A&M University; M. Cacciari, LPTHE Paris and CNRS; L. Ropelewski, E. Oliveri and R. Veenhof, CERN (PH-DT); H. Mugnier, J. Rousset, et P. Chalmet of MIND Microtechnologies in Archamps, and the Departement de la Haute Savoie for supporting MIND to collaborate on monolithic detectors;