CERN / LHCC / 95–71
LHCC / P3
15 December 1995

ALICE

Technical Proposal
for
A Large Ion Collider Experiment
at the
CERN LHC

For this edition obvious typing and typographical errors have been corrected. Thus layout and pagination are slightly different from those of the first, limited, edition.

Cover design by Sylvette Stillebacher

Printed at CERN
December 1995

ISBN 92-9083-077-8

ALICE Collaboration

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

Aligarh, India, Physics Department, Muslim University:
N. Ahmad, T. Ahmad[*)], I. Aziz, M. Irfan, K.A. Siddiqui and M. Zafar.

Athens, Greece, Nuclear and Particle Physics Division, University of Athens:
N. Antoniou, A. Kapogiannis, A.D. Panagiotou and G. Vassiliadis.

Bari, Italy, Politecnico and Sezione INFN:
E. Cantatore, F. Corsi, D. De Venuto, L.M. Galantucci, G. Gramegna and C. Marzocca.

Bari, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
D. Di Bari, D. Elia, R.A. Fini, B. Ghidini, A. Jachołkowski, V. Lenti, V. Manzari, E. Nappi, F. Navach and G. Tomasicchio.

Beijing, China, China Institute of Atomic Energy:
Z. Sun.

Bergen, Norway, Department of Physics, University of Bergen:
E. Andersen, H. Helstrup, A.K. Holme, A. Klovning, O. Maeland, T.F. Thorsteinsen and E. Zabrodin.

Bhubaneswar, India, Institute of Physics:
D. Mahapatra[#)].

Birmingham, United Kingdom, School of Physics and Space Research, University of Birmingham:
I.J. Bloodworth, D. Evans, P. Jovanović, J.B. Kinson, A. Kirk, J. Nelson, O. Villalobos Baillie, M.F. Votruba and R. Zybert.

Bombay, India, Physics Department, Indian Institute of Technology:
R. Varma.

Bratislava, Slovakia, Faculty of Mathematics and Physics, Comenius University[‡)]:
J. Braciník, V. Černý, J. Ftáčnik, V. Hlinka, M. Ivanov, R. Janik, R. Lietava, M. Pikna, J. Pišút, N. Pišútová, B. Sitar, P. Strmeň and I. Szarka.

Budapest, Hungary, KFKI Research Institute for Particle and Nuclear Physics, Academy of Science:
E. Denes, Z. Fodor, T. Kiss, Z. Meggyesi, G. Palla, G. Rubin, J. Sziklai and J. Zimanyi.

Calcutta, India, Saha Institute of Nuclear Physics:
P. Bhattacharya, S. Bose, Sukalyan Chattopadhyay, S. Roy and P. Sen.

Calcutta, India, Variable Energy Cyclotron Centre:
Z. Ahamad, Subhasis Chattopadhyay, M.R. Dutta Mazumdar, M.S. Ganti, D.S. Mukhopadhyay, R.P. Sharma, B. Sinha[1)], M.D. Trivedi and Y.P. Viyogi.

Catania, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
C. Caligiore, G. Cinà, A. Insolia, D. Lo Presti, S. Panebianco, C. Petta, N. Randazzo, S. Reito, G.V. Russo, M. Russo[2] and C. Tuvé.

CERN, Geneva, Switzerland, European Organization for Nuclear Research:
F. Antinori, H. Beker, J.-C. Berset[#], M. Bosteels[#], R. Brun, M. Campbell[#], W. Carena, E. Chesi, J. Christiansen, J. de Groot, R. Divia, C. Engster[#], H.G. Fischer, C. Gregory, B. Goret, G. Harangozo[3], E. Heijne[#], W. Klempt, K. Knudson, J.-C. Lassalle, J.-C. Legrand, L. Leistam, B. Lofstedt[#], C. Lourenço, A. Morsch, M. Mota[4], L. Pigni, F. Piuz, E. Quercigh, J. Raynaud, E. Rosso[#], J.-C. Santiard[#], J. Schukraft, W. Snoeys[#], P. Sonderegger, D. Swoboda, P. Vande Vyvre, A. Vascotto and S. Wenig.

Chandigarh, India, Physics Department, Panjab University:
M.M. Aggarwal, R.K. Bharadwaj, A.K. Bhati and V.S. Bhatia.

Clermont-Ferrand, France, Université Blaise Pascal and IN2P3:
J.-P. Alard, A. Baldit, N. Bastid, J. Castor, Th. Chambon, A. Devaux, P. Dupieux, B. Espagnon, J. Fargeix, P. Force, F. Manso, G. Roche and P. Saturnini.

Copenhagen, Denmark, Niels Bohr Institute:
I. G. Bearden, H. Bøggild, J.J. Gaardhøje and O. Hansen.

Cracow, Poland, H. Niewodniczański Institute of Nuclear Physics, High Energy Physics Laboratory:
J. Bartke, E. Gładysz-Dziaduś, E. Górnicki, M. Kowalski, A. Rybicki and P. Stefański.

Darmstadt, Germany, Gesellschaft für Schwerionenvorschung:
H. Appelshäuser, R. Bock, P. Braun-Munzinger, R. Brockmann, J. Eschke, B. Kolb, U. Lynen, J. Lühning, D. Miskowiec, W.F.J. Müller, J. Pochodzalla, M. Purschke, A. Sandoval, H. Sann, H.R. Schmidt, H. Stelzer, W. von Rüden, D. Vranic and A. Wörner.

Frankfurt, Germany, Institut für Kernphysik, Johann-Wolfgang Göthe Universität:
J. Bächler, C. Borman, M. Fuchs, M. Gaździcki, M. Günter, S. Haider, R. Renfordt, D. Röhrich, G. Roland, R. Stock and H. Ströbele.

Gatchina, Russia, St Petersburg Nuclear Physics Institute[‡]:
O. Kiselev, B. Komkov, V. Nikulin and S. Volkov.

Geneva, Switzerland, Department de Physique Nucléaire et Corpusculaire, Université de Geneve:
A. Angelis, P. Foka, H. Kalechofsky, M. Martin and L. Rosselet.

Heidelberg, Germany, Max-Planck-Institut für Kernphysik:
F. Ceretto, U. Faschingbauer, C. Fuchs, O. Nix and P. Wurm.

Heidelberg, Germany, Physikalisches Institut, Ruprecht-Karls Universität:
R. Baur, A. Drees, P. Glässel, C. Jung, B. Lenkeit, M. Messer, A. Pfeiffer, J. Stachel, Th. Ullrich, C. Voigt and J.P. Wessels.

Jaipur, India, Physics Department, University of Rajasthan:
A. Agnihotri, K.B. Bhalla, A. Bharti, S.K. Gupta, V. Kumar, R. Raniwala and S. Raniwala.

Jammu, India, Physics Department, Jammu University:
P.V.K.S. Baba, S.K. Badyal, A. Bhasin, A. Gupta, V.K. Gupta, L.K. Mangotra, N.K. Rao and S. Sambyal.

JINR, Dubna, Russia, Joint Institute for Nuclear Research:
P.G. Akishin, V.A. Arefiev, V.I. Astakhov, A.A. Baldin, A.M. Baldin, B.V. Batyunya, Yu.A. Belikov, N.A. Blinov, Z.V. Borisovskaya, V.I. Datskov, V.K. Dodokhov, L.G. Efimov, A.G. Fedunov, B.N. Guskov, V.G. Kadyshevsky, V.S. Korolev, A.D. Kovalenko, V.I. Lobanov, A.N. Maksimov, E.A. Matyushevsky, K.V. Mikhailov, D.P. Mikhalev, P.V. Nomokonov, A.N. Parfenov, V.N. Pozdnyakov, I.V. Puzynin, G.P. Reshetnikov, A.E. Senner, S.V. Semashko, G.S. Shabratova, A.V. Shabunov, I.A. Shelaev, Yu.A. Shishov, N.V. Slavin, G.P. Tsvineva, A.S. Vodopianov, O.I. Yuldashev, M.B. Yuldasheva and A.I. Zinchenko;

N. Grigalashvili[5)], A. Mamulashvili[5)], Z. Menteshashvili[5)], M. Nioradze[5)] and Yu. Tevzadze[5)];

K.G. Akhobadze[6)], E.S. Ioramashvili[6)], A.K. Javrishvili[6)], A.V. Kharadze[6)], T.V. Khuskivadze[6)], L.V. Shalamberidze[6)] and N. Shubitidze[6)].

Kharkov, Ukraine, National Scientific Centre ‘Kharkov Institute of Physics and Technology’:
G.L. Bochek, V.F. Boldychev, I.F. Chervonny, V.I. Kulibaba, N.I. Maslov, I.M. Prokhorets and D.G. Pugachev.

Kharkov, Ukraine, Scientific and Technological Research Institute of Instrument Engineering:
V.N. Borshchov, V.V. Dorokhov, A.G. Fyodorov, S.K. Kiprich, A.M. Listratenko and V.E. Starkov.

Kiev, Ukraine, Institute for Theoretical Physics, International Centre of Physics, National Academy of Science of Ukraine:
D.E. Kharzeev, A.I. Siamtomov and H. Zinovyev.

Košice, Slovakia, Institute of Experimental Physics, Slovak Academy of Sciences and Faculty of Science, Šafárik University:
A. Jusko, I. Králik, M. Lupták, G. Martinská, B. Pastirčák, K. Šafařík, L. Šándor, J. Urbán, S. Vokál and J. Vrláková.

Legnaro, Italy, Laboratori Nazionali INFN:
R.A. Ricci.

Lund, Sweden, Division of Cosmic and Subatomic Physics, University of Lund:
L. Carlsson, K. El Chenawi, S. Garpman, H.-A. Gustafsson, A. Oskarsson, I. Otterlund, K. Söderström and E. Stenlund.

Marburg, Germany, Fachbereich Physik, Philipps Universität:
F. Eckhardt, A. Piper and F. Pühlhofer.

Mexico City, Mexico, Centro de Investigación y de Estúdios Avanzados[‡)]:
G. Herrera Corral, E. Linares and H. Mendez.

Minsk, Belarus, Institute for Nuclear Problems, State University:
A. Borisevich, G. Drobyshev, O. Kondratiev, A. Lobko, A. Lopatik, A. Oskin and R. Zouevsky.

Moscow, Russia, Institute for Nuclear Research, Academy of Science:
S.N. Filippov, A.V. Dementyev, Yu.K. Gavrilov, L.G. Gevorkov, M.B. Golubeva, T.A. Golubeva, M.N. Gotra, F.F. Guber, T.L. Karavicheva, A.B. Kurepin, V.D. Laptev, Yu.V. Perepechkin, V.I. Rasin, A.I. Reshetin, K.A. Shileev and N.M. Sobolevsky.

Moscow, Russia, Institute for Theoretical and Experimental Physics:
V.A. Akimov, A.N. Akindinov, S.V. Boyarinov, M.M. Chumakov, J.G. Grishuk, M.M. Kats, S.M. Kisselev, Y.T. Kisselev, V.G. Komlev, S.V. Koulechov, A.N. Martemianov, K.R. Mikhailov, A.A. Nikitine, P.A. Polozov, S.A. Pozdnyakov, V.S. Serov, A.V. Smirnitsky, V.P. Surin, K.G. Voloshin, L.S. Vorobyov, A.V. Yumashev and B.V. Zagreev.

Moscow, Russia, Kurchatov Institute of Atomic Energy:
V. Antonenko, S. Belyaev, R. Cherbatchev, Iou. Doubovik, S. Fokin, M. Ippolitov, K. Karadjev, A. Lebedev, V. Lebedev, V. Manko, G. Mgebrishvili, A. Nianine, S. Nikolaev, O. Patarakin, Iou. Sibiriak, A. Tsvetkov, A. Vasiliev and A. Vinogradov.

Moscow, Russia, Moscow Engineering Physics Institute:
V. Grigoriev, V. Kaplin, I. Linin and V. Loginov.

Münster, Germany, Institut für Kernphysik, Westfälische Wilhelms Universität:
C. Barlag, C. Blume, E.-M. Bohne, D. Bucher, A. Claussen, T. Peitzmann, R. Santo and H. Schlagheck.

Nantes, France, Laboratoire de Physique Subatomique et des Technologies Associées:
J.-P. Cussoneau, C. Drancourt, B. Erazmus, G. Guilloux, H. Gutbrod, P. Lautridou, M. Le Guay, L. Martin, G. Paić[7)] and J. Pluta.

Novosibirsk, Russia, Budker Institute for Nuclear Physics:
A.R. Frolov, Yu.N. Pestov and M.A. Tiunov.

Oak Ridge, U.S.A., Instrumentation and Controls Division, Oak Ridge National Laboratory:
C.L. Britton, W.L. Bryan, M.N. Ericson, M.L. Simpson, J.W. Walker and A.L. Wintenberg.

Orsay, France, Institut de Physique Nucléaire, Université de Paris-Sud:
L. Bimbot, M.P. Comets, M. Denoit, D. Jouan, Y. Le Bornec, J. Pouthas, V. Poux, R. Sellem, N. Willis and Z. Zojceski.

Oslo, Norway, Department of Physics, University of Oslo:
B. Kvamme, G. Løvhøiden, B. Skaali, D. Wormald, B. Wu and J. Yuan.

Padua, Italy, Dipartimento di Fisica dell'Università and Sezione INFN:
A. Andrighetto, N. Carrer, C. Fanin, M. Morando and G. Segato.

Prague, Czech Republic, Institute of Physics, Academy of Science:
J.A. Mareš, M. Nikl, K. Píška, K. Polák, P. Staroba and P. Závada.

Protvino, Russia, Institute for High Energy Physics:
G.V. Khaoustov, Yu.V. Kharlov, V.N. Kolosov, Yu.M. Orlov, V.A. Poliakov, S.A. Polovnikov, S.A. Sadovsky, V.D. Samoilenko, P.A. Semenov and A.E. Sobol.

Rehovot, Israel, Weizmann Institute of Science:
Z. Fraenkel, C. de los Heros, I. Ravinovich, E. Socol, G. Tel-Zur and I. Tserruya.

Řež u Prahy, Czech Republic, Nuclear Physics Institute, Academy of Science:
A. Kugler, M. Pachr, V. Petracek, J. Rak, M. Šumbera, P. Tlustý and V. Wagner.

Rome, Italy, Dipartimento di Fisica dell'Università 'La Sapienza' and Sezione INFN:
S. Di Liberto, M.A. Mazzoni, F. Meddi and D. Prosperi.

Salerno, Italy, Dipartimento di Fisica Teorica e sue Metodologie per le Scienze Applicate, Università di Salerno:
G. Grella, M. Guida, G. Romano, G. Rosa and T. Virgili.

St Petersburg, Russia, Institute for Physics of St Petersburg State University, Mendeleev Institute for Metrology and Meson Scientific Association:
L.Yu. Abramova, V.S. Alexandrov, V.M. Baratov, A.A. Bolonin, M.A. Braun, V.M. Dobulevitch, V.M. Fedorov, G.A. Feofilov, S.F. Gerasimov, O.N. Godisov, S.N. Igolkin, A.A. Kolojvari, I.A. Novikov, S.V. Potapov, O.I. Stolyarov, F.A. Tsimbal, T.A. Tulina, F.F. Valiev, V.V. Vechernin, L.I. Vinogradov, L.F. Vitushkin and M.I. Yudkin.

Strasbourg, France, Centre de Recherches Nucléaires and Laboratoire d'Electronique et de Physique de Systèmes Instrumentaux, Université Louis Pasteur and IN2P3:
M. Ayachi, J.D. Berst, J.P. Blondé, J.-P. Coffin, P. Crochet, G. Gaudiot, G. Guillaume, L. Hebrard, G. Heitz, Y. Hu, F. Jundt, C. Kuhn, C. Maazouzi, A. Michalon, F. Rami, Ch. Roy and L. Tizniti.

Trieste, Italy, Dipartimento di Fisica dell'Università and Sezione INFN[‡)]:
V. Bonvicini, A. Rashevsky and A. Vacchi.

Turin, Italy, Dipartimento di Fisica Sperimentale dell'Università and Sezione INFN:
B. Alessandro, S. Beolé, G. Bonazzola, E. Chiavassa, F. Daudo, N. De Marco, G. Dellacasa, M. Gallio, P. Giubellino, M. Idzik, A. Marzari-Chiesa, M. Masera, G. Mazza, A. Musso, A. Piccotti, L. Ramello, L. Riccati, M.S. Sartori, E. Scomparin and E. Vercellin.

Utrecht, The Netherlands, NIKHEF and Subatomic Physics Department of the Utrecht University:
M. Botje, A. Buijs, A. de Haas, C. de Laat, R. Kamermans, P. Kuijer, W. Lourens, A. van den Brink, G. van der Mark and N. van Eijndhoven.

Warsaw, Poland, Soltan Institute for Nuclear Studies:
K. Karpio, T. Siemarczuk, G. Stefanek, L. Tykarski and G. Wilk.

Wuhan, China, Institute of Particle Physics, Huazhong Normal University:
X. Cai, F. Liu, L. Liu, H. Wang and D. Zhou.

Yerevan, Armenia, Yerevan Physics Institute[‡)]:
R. Avakyan, V. Eganov, A. Grigoryan, H. Gulkanyan, R. Hakobian, I. Keropyan, S. Mehrabian and R. Shahoian.

Zagreb, Croatia, Ruder Bošković Institute:
D. Ferenc, A. Ljubičić Jr. and T. Tustonic.

Spokesperson:	**J. Schukraft,**	E-mail=Jurgen.Schukraft@cern.ch
Deputy:	**H. Gutbrod,**	E-mail=GUTBROD@VXWA80.CERN.CH

Acknowledgements

It is a pleasure to acknowledge the important contributions given to the preparation of this Technical Proposal by: E. Badura, P. Bloch, R. Brown, S. Cittolin, V. Eckardt, K. Eggert, V. Goidadin, A. Grimaldi, F. Hahn, S. Jaquier, V. Klyukhin, A. Marchioro, R. McLaren, J.-F. Michaud, C. Neyer, M. Price, H.G. Ritter, A. Rudge, O. Runolfsson, E. Sabin, R. Schulze, P. Seyboth, G.R. Stevenson, A. Tricomi and E. Vanderbij.
The collaboration wishes to thank all the technical and administrative staff involved during the preparation of this Technical Proposal. In particular, we are grateful for the help given by the CERN ECP-EDO group. We also thank M. Goossens and the staff of the Desktop Publishing Service for their professional help with the editing.

[*)] Deceased.
[#)] Also participating in other LHC experiments.
[‡)] Applying to join ALICE.
[1)] Also at Saha Institute of Nuclear Physics.
[2)] Istituto di Ingegneria Elettrica e Telecomunicazioni dell'Università di Catania and Sezione INFN.
[3)] Visitor at CERN from KFKI-Budapest, Hungary.
[4)] Visitor at CERN from LIP-Lisbon, Portugal.
[5)] High Energy Physics Institute, Tbilisi State University, Tbilisi, Georgia.
[6)] Institute of Physics, Georgian Academy of Sciences, Tbilisi, Georgia.
[7)] On leave of absence from Ruder Bošković Institute.

Contents

List of Figures

1 Introduction and overview

1.1 Ultra-relativistic heavy-ion collisions

The aim of high-energy heavy-ion physics is the study of strongly interacting matter at extreme energy densities (QCD thermodynamics). Statistical QCD predicts that, at sufficiently high density, there will be a transition from hadronic matter to a plasma of deconfined quarks and gluons —a transition which in the early universe took place in the inverse direction some 10^{-5} s after the Big Bang and which might still play a role today in the core of collapsing neutron stars.

The study of the phase diagram of nuclear matter, using methods and concepts from both nuclear and high-energy physics, constitutes a new and interdisciplinary approach to investigating matter and its interactions. It is of interest to explore and test QCD on its natural scale (Λ_{QCD}) and to address the fundamental questions of confinement and chiral-symmetry breaking. Moreover, it is of general relevance in understanding the dynamic nature of phase transitions involving elementary quantum fields, as the QCD phase transition is the only one accessible to laboratory experiments.

1.1.1 The current heavy-ion program

The early exploratory program with light ions (up to ^{32}S projectiles), at CERN and Brookhaven, has established the feasibility of high-energy ion–ion experiments with their abundant particle production. It has shown that high energy densities can indeed be obtained in these reactions and has produced evidence for the onset of new collective phenomena. Some characteristic features of these reactions (e.g. momentum distributions and particle ratios) are already surprisingly close to the ones expected for a 'macroscopic' lump of hadronic matter heated to a temperature of close to 200 MeV.

The recently started experiments with really heavy ions (Au and Pb beams), both at BNL and CERN, should determine to what extent one can actually establish a regime of thermodynamic behaviour. They should reach baryon densities close to or even exceeding the ones in the core of neutron stars, and, if present estimates are correct, could produce more conclusive evidence for the onset of quark deconfinement.

1.1.2 Heavy-ion physics at the LHC

Besides proton–proton collisions, aimed at exploring physics at the TeV scale, both proton–nucleus and nucleus–nucleus collisions are foreseen as part of the initial experimental program at the LHC. With heavy ions at a centre-of-mass energy of about 5.5 TeV/nucleon, the LHC is the only machine which will reach, and even extend, the energy range probed by cosmic ray nucleus–nucleus collisions.

Extrapolating from present results, all parameters relevant to the formation of the Quark–Gluon Plasma (QGP) will be more favourable: the energy density, the size and lifetime of the system, and relaxation times should all improve by a large factor, typically by an order of magnitude compared to Pb–Pb collisions at the SPS [1, 2]. It should then be possible to obtain energy densities well above the deconfinement threshold, and to probe the QGP in its asymptotically free 'ideal gas' form. Unlike at lower energies, the central rapidity region will have a nearly vanishing baryon number density, similar to the state of the early universe. It is probably dominated in the early pre-equilibrium stage by a very dense system of semi-hard partons ('mini-jets'), which would lead to rapid thermalization and extremely high initial temperatures.

Heavy-ion collisions at the LHC are therefore expected to provide a very different, and significantly better, environment for the study of strongly interacting matter than existing accelerators.

1.1.3 The ALICE experiment

The heavy-ion detector ALICE (A Large Ion Collider Experiment) has emerged as a common design from the heavy-ion community currently working at CERN and a number of groups new to this field from both nuclear and high-energy physics. It currently includes about 560 physicists from 63 institutions.

ALICE was first proposed as a central detector in 1993, see Letter of Intent (LoI) [3], and later complemented by an additional forward muon spectrometer designed in 1995 [4]. It is a general-purpose heavy-ion experiment, sensitive to the majority of known observables (including hadrons, electrons, muons and photons), and it will be operational at the start-up of the LHC.

With ALICE we want to measure the flavour content and phase-space distribution, event by event, for a large number of particles whose momenta and masses are of the order of the typical energy scale involved (temperature $\sim \Lambda_{\mathrm{QCD}} \sim 200$ MeV). The experiment is designed to cope with the highest particle multiplicities anticipated for Pb–Pb reactions ($\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$).

In addition to heavy systems (e.g. Pb–Pb), the ALICE Collaboration will study collisions of lower-mass ions, which are a means of varying the energy density [5], and protons (both pp and p–nucleus), which provide reference data for the nucleus–nucleus collisions.

1.2 Physics topics

1.2.1 Heavy-ion reactions

ALICE has been designed as a dedicated heavy-ion experiment and its prime aim is the study of nuclear collisions at LHC. In order to establish and analyse the existence of QCD bulk matter and the QGP, a number of observables have to be studied with ALICE in a systematic and comprehensive way [6]. Some observables are needed to characterize the global features of the state created during the collision in order to constrain theoretical models (e.g. relevant degrees of freedom, size, lifetime, density, dynamic evolution). These observables yield information about the initial conditions and space–time evolution which is necessary in order to interpret a specific signal as a QGP signature or as an indication of new physics. Our strategy is to study a number of these *specific signals* together with *global information* about the events, in the same experiment.

The signals accessible to our detector are described in detail in Chapter 11. They are summarized below according to the aspect of the collision on which they have a bearing (see, for example, Ref. 7 for original references):

- **Initial Conditions:** *global event features* measure the number of colliding nucleons and give information on the energy density obtained.
- **Quark-Gluon Plasma:** *open charm* production will probe the parton kinematics in the very early stage; *prompt photons* can reveal the characteristic thermal radiation from the plasma; the cross-section of *high-p_t hadrons* is sensitive to the energy loss of the partons in the plasma; J/ψ and Υ production probe deconfinement.
- **Phase Transition:** *strangeness production* is sensitive to the large s-quark density expected from (partial) chiral-symmetry restoration in the plasma; *multiplicity fluctuations* are a signature for the critical phenomena at the onset of a phase transition; *particle interferometry* measures the expansion time in the mixed phase, which is expected to be long in the case of a first-order phase transition.
- **Hadronic Matter:** *particle ratios, p_t distributions* and *resonance line-shape parameters* are all sensitive to the dynamic evolution of the hadronic phase; *interferometry* allows measurement of the freeze-out radius of the hadronic fireball.

1.2.2 Non-QGP physics

$\gamma\gamma$ collisions

In addition to the heavy-ion physics programme, the ALICE Collaboration has stated in the LoI its intention to study photon–photon reactions and a number of large cross-section pp processes not accessible to the dedicated pp experiments.

The large flux of virtual photons associated with the Coulomb fields of coherently interacting Pb nuclei will lead to $\gamma\gamma$ luminosities which exceed by orders of magnitude those achieved at existing or future e^+e^- colliders [8, 9]. The physics interest in $\gamma\gamma$ collisions includes a large number of QED and QCD processes, from lepton pairs to hadrons, hadron pairs and jet pairs [10]. While the production of high-mass objects such as a light Higgs ($m \approx 100$ GeV) seems out of reach, even a short run of 10^6 s will produce a copious amount of mesons in the mass range below a few tens of GeV, e.g. a hundred $\eta_b(9366)$, the as yet undiscovered pseudoscalar bottonium, and half a million $\eta_c(2981)$ [9–12].

The two-photon process, being a probe of the quark content of the final state, is of particular interest for meson spectroscopy. The LHC with Pb or Ca beams should permit the study of, for example, hybrid mesons, multiquark states, high-spin states, and glueballs.

ALICE may well provide a very attractive tool for photon–photon collisions. While we clearly recognize this possibility, a substantial upgrade of our detector — in particular to include large-scale coverage with electromagnetic calorimeters for photon detection— is currently beyond the means of our collaboration. However, a study is under way to explore the physics capabilities of the present set-up using exclusively charged final states.

Large cross-section pp physics

In a memorandum to the LHCC [13], the ALICE Collaboration, together with a group of individuals and institutions not part of ALICE, has expressed an interest in studying large cross-section processes in pp collisions (e.g. diffraction, rapidity gap physics, elastic scattering) within the environment of our present detector. However, if the constraints of the LHC experimental programme permit, an independent experiment might be dedicated to such physics [14]; therefore, for the moment the ALICE Collaboration does not actively pursue these topics.

1.3 Design and performance

The ALICE detector has been designed according to the following criteria:

- A single dedicated heavy-ion detector is foreseen at the initial stage of LHC. It has to cover, as completely as possible the full range of sensitive signatures. Its design must be conservative and robust.
- In addition to addressing particular signals which look most promising today, we will search in an unbiased way for qualitative and quantitative differences between pp and nucleus–nucleus collisions in a number of final states (hadrons, leptons, photons).
- The detector will concentrate on physics at mid-rapidity, i.e. on the regions of minimum baryon and

maximum energy density, which are the unique features of LHC compared to the lower energy machines SPS and RHIC.

- An open geometry will facilitate future modifications and upgrades if first physics results suggest focusing on specific signals, selective triggers or larger acceptance.

1.3.1 Experimental conditions

The average event-rate for Pb–Pb collisions at the LHC will be of the order of 10^4 minimum-bias collisions per second. Of these, a few per cent (~ 100 events/s) correspond to the most interesting central collisions, with maximum particle production. The low interaction rates lead naturally to an approach which combines large geometrical acceptance with simple central-collision triggers and a high data-taking rate.

These rates are matched with the dead-time of our slowest detector and the capabilities of our DAQ system. A large collected sample of central collisions will allow an unbiased search for signals, including 'surprises' not predicted by current theories, a feature of particular importance in the rapidly evolving field of relativistic heavy-ion physics. In the initial phase, running times comparable with the present SPS ion program will be adequate to collect sufficient statistics. Therefore, we will base our rate estimates on 10^6 s/year heavy-ion running time ($\sim 10\,\%$ of the total), yielding a few 10^7 central events collected for offline analysis.

The main emphasis will be on data taking at the highest possible energy density, i.e. with Pb beams. In addition, running with one or two intermediate-mass ion species will provide the necessary means to vary the energy density. Running with protons will be required initially to commission the detector and to provide reference data for the nuclear programme. Schemes are under investigation that would allow pp running at luminosities adapted to the rate capability of the detector without any adverse effect on the high-luminosity interaction regions.

A design change made recently to the LHC RF system will permit asymmetric beam collisions, in particular p–nucleus reactions. These will be important for disentangling initial state effects, like nuclear shadowing, from final state interactions in hard probe signals, like jet quenching and quarkonium production. As both luminosity and particle multiplicity in p–A are expected to be between the extremes of pp and A–A, ALICE should be well suited to make use of this new possibility.

1.3.2 Acceptance

The rapidity acceptance has to be large enough to allow some variables to be studied on an event-by-event basis (particle ratios, p_t spectra, HBT radii), meaning typically several thousand reconstructed particles per event. Good efficiency for detecting the decay of particles at $p_t < m$ requires about 2 units in rapidity (for masses above 1–2 GeV) and corresponding coverage in azimuth. In addition, efficient rejection of low-mass Dalitz decays, which is needed for the lepton-pair measurements, can only be approached with full azimuthal coverage. The rapidity coverage of our central detector ($|\eta| < 0.9$) has been chosen as a compromise between acceptance and cost [15].

In order to have some sensitivity to the global event structure and long-range correlations on scales $\Delta y \gg 1$ (e.g. 'bubbles' of QGP, critical phenomena associated with a first-order phase transition), we will measure $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta$ with multiplicity detectors in a large rapidity window, $|\eta| < 4$, outside the central acceptance.

1.3.3 Tracking and momentum resolution

Track finding in the environment of heavy-ion collisions at LHC presents one of the greatest challenges for the detector. Therefore, the design of our tracking system has been driven primarily by the requirement for safe and robust pattern recognition. It uses mostly three-dimensional hit information, with many points and a weak magnetic field.

The momentum cut-off should be as low as possible (< 100 MeV/c) in order to study collective effects associated with large length scales. A low-p_t cut-off is also essential to suppress the soft conversions and Dalitz background in the lepton-pair spectrum. The most stringent limit on momentum resolution in the low-p_t region is imposed by identical particle interferometry, owing to the large source radii and the corresponding narrow correlation enhancement. For intermediate momenta, the mass resolution should be of the order of the natural width of the ω and ϕ in order to maximize the signal-to-background ratio and, more importantly, to study the mass and width of these mesons in the dense medium. At high momenta, the resolution has to be sufficient to measure the spectrum of jets via leading particles. Detection of secondary vertices from strange and, in particular, charm hadrons requires a vertex detector with high resolution at a small distance from the interaction point. Our tracking system, together with a weak (0.2 T) solenoidal field, has been designed to fulfil all of the above requirements.

1.3.4 Particle identification

The momentum range for particle identification (PID) can be limited, for the bulk of the hadronic signals, to a few times the average p_t —more than 97 % of all charged particles are below $p_t = 2$ GeV/c, less than 0.02 % are above $p_t = 5$ GeV/c. π:K:p separation better than $3\,\sigma$ is needed on a track-by-track basis for the abundant soft hadrons in order to study HBT with identified particles, decays (e.g. hyperons and $\phi \to$ KK) and event-by-event particle ratios. In order to measure inclusive particle ratios and p_t spectra in the mini-jet region we need some limited

PID coverage which extends towards higher momenta. The e/π rejection power must be sufficient to reduce the combinatorial background caused by misidentification to below the level of the remaining unrejected Dalitz pairs.

1.3.5 Photon detection

Besides measuring π^0 and η spectra, whose momentum accuracy is better than that for charged particles, the primary goal is the measurement of direct photon spectra. The accuracy of the latter will be determined in the range of interest ($p_{\mathrm{t}} < 10$ GeV/c) by systematic errors in photon-reconstruction efficiency and knowledge of the decay background. A sensitivity to direct photons at the level of 5 % of the decay photons requires a knowledge of the reconstruction efficiency and the π^0 cross-section to ≤ 3 %, and of the η cross-section to ≤ 10 %. An acceptable systematic error can be obtained only at low channel occupancy and therefore requires a calorimeter with small Molière radius at a large distance (≥ 5 m) from the vertex. The minimum size is determined by the acceptance for the 2γ decays of π^0 and η at low p_{t}. The signal-to-background ratio in the two-photon invariant mass spectra improves proportionally to the energy resolution of the calorimeter.

1.4 Detector overview

The overall detector layout of the ALICE experiment is shown in Figs. 1.1 and 1.2.

The central part, which covers $\pm 45°$ ($|\eta| < 0.9$) over the full azimuth, is embedded in a large magnet with a weak solenoidal field. It consists (from the inside out) of an inner tracking system (ITS) with six layers of high-resolution silicon tracking detectors, a cylindrical TPC, and a large area PID array of TOF counters. In addition, there are two small area single-arm detectors: an electromagnetic calorimeter (PHOS) and an array of counters optimized for high-momentum inclusive particle identification (HMPID), made up of either RICH or TOF counters.

The forward muon spectrometer will have an acceptance of 2–10° ($\eta = 2.4-4$). It consists of a complex arrangement of absorbers, a large dipole magnet, and twelve stations of tracking and triggering chambers.

The set-up is completed by a set of zero-degree calorimeters (ZDCs) located far downstream in the machine tunnel and a forward multiplicity detector (FMD) covering a significant fraction of the phase space ($|\eta| < 4$).

1.4.1 Magnet

The optimal choice for our experiment is a large solenoid with a weak field. The choice of a weak and uniform solenoidal field together with continuous tracking in a TPC eases considerably the task of pattern recognition. The field strength —a compromise between momentum resolution, low momentum acceptance and tracking efficiency— will be 0.2 T, allowing full tracking and particle identification down to ~ 100 MeV/c in p_{t}. Lower momenta are covered by the inner tracking system. The inner radius of the magnet must be large enough to accommodate a single-arm electromagnetic calorimeter for prompt photon detection, which must be placed at a distance of ~ 5 m from the vertex because of the particle density.

The magnet of the L3 experiment fulfils all these requirements. It can be left in its present position, centred 0.3 m below the LHC beam height. Some repair and maintenance work will be required at the end of its operation at LEP in order to ensure continued proper functioning in the LHC era.

1.4.2 Inner Tracking System (ITS)

The basic functions of the inner tracker —secondary vertex reconstruction of charm and hyperon decays, particle identification and tracking of low-momentum particles, improvement of the momentum resolution— are achieved with six barrels of high-resolution detectors. The number of layers and their position has been optimized for efficient pattern recognition and impact parameter resolution. Because of the high particle density, the four innermost layers ($r \leq 24$ cm) must be truly two-dimensional devices, e.g. silicon pixel and silicon drift detectors. The outer superlayer at $r \approx 45$ cm will be equipped with double-sided silicon micro-strip detectors. With the exception of the two innermost pixel planes, all layers will have analog readout for independent particle identification via $\mathrm{d}E/\mathrm{d}x$ in the non-relativistic region, which will give the inner tracking system a stand-alone capability as a low-p_{t} particle spectrometer.

1.4.3 Time Projection Chamber (TPC)

The need for efficient and robust tracking has led to the choice of a TPC as the main tracking system. In spite of its drawbacks concerning speed and data volume, we have concluded that only a conventional device and redundant tracking can guarantee reliable performance at up to 8000 charged particles per unit of rapidity. The inner radius of the TPC ($r \approx 90$ cm) is given by the maximum acceptable hit density (0.1 cm^{-2}). The outer radius (250 cm) is determined by the length required for a $\mathrm{d}E/\mathrm{d}x$ resolution of < 7 %. With this resolution, in addition to its tracking function the TPC can serve as a detector for electron identification up to momenta of ~ 2.5 GeV/c. The design of the readout and of the end plates, based on ongoing developments for other heavy-ion TPCs, as well as the choice of the operating gas, are being optimized for good double-track resolution.

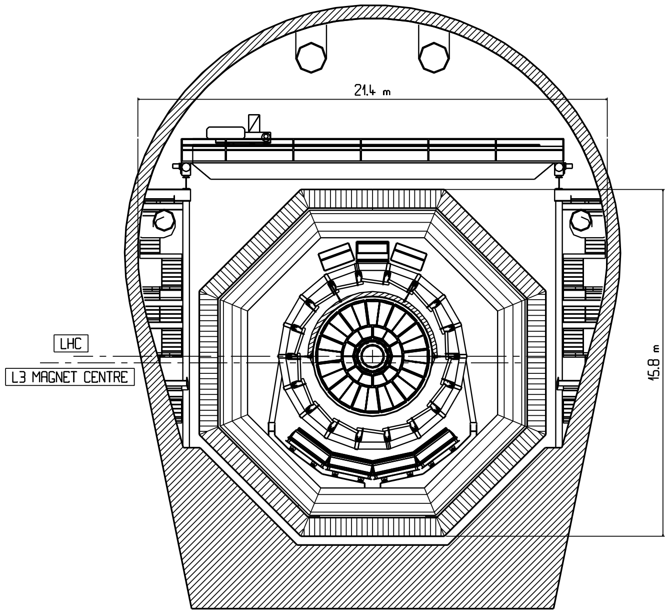

Figure 1.1: Transverse view of the ALICE detector.

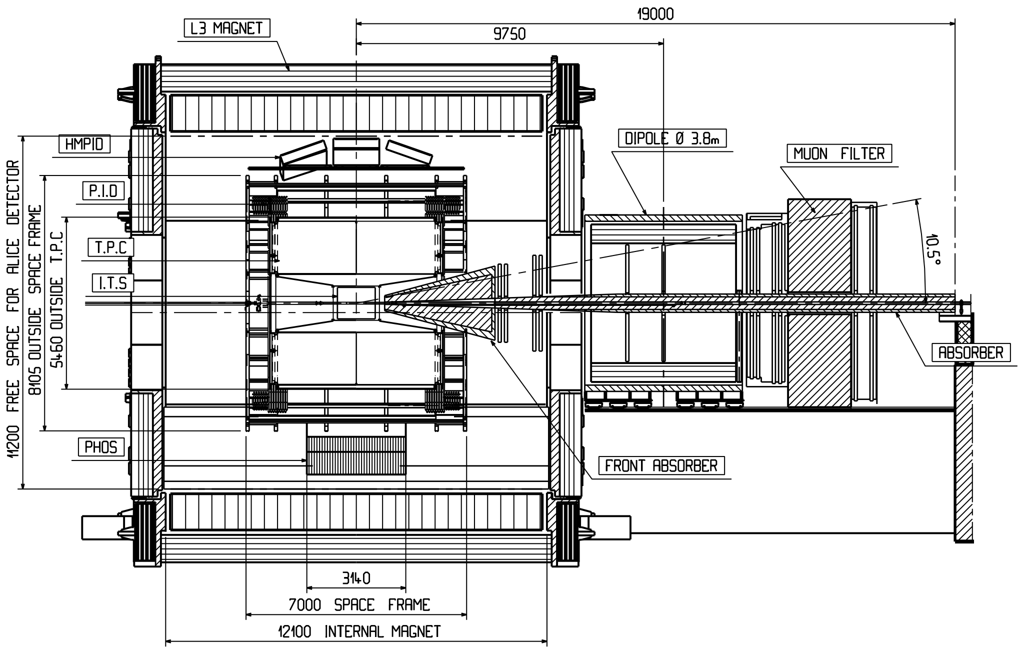

Figure 1.2: Longitudinal view of the ALICE detector.

1.4.4 Particle Identification System (PID)

Two TOF technologies, i.e. Pestov spark counters and parallel plate chambers (PPC), are under study for the large area particle identification array at a radius of about 3.5 m. First prototypes have been operated in a test beam and have shown a timing resolution of less than 50 ps (r.m.s.) for the Pestov counter. The PPCs, which are less demanding in terms of construction and operation, have achieved a resolution better than 200 ps. Due to its superior timing resolution, the Pestov counter will be the preferred choice if it can be shown that large area systems can be reliably manufactured and operated.

A second system, of smaller acceptance and at larger radii, is optimized for the detection of high p_t particles and will extend the accessible momentum range for inclusive particle spectra into the semi-hard region. Our current choice, already matching most of the performance specifications for this system, is a proximity-focusing RICH detector with liquid freon radiator, solid photocathode and pad readout. However, depending on the final performance achieved in our R&D, a TOF solution could also be envisaged for this detector.

1.4.5 Photon Spectrometer (PHOS)

Prompt photons, π^0's and η's will be measured in a single-arm high-resolution electromagnetic calorimeter. It is located 5 m from the vertex and covers 18 m^2 with around 36 000 channels. It has to be built from a material with small Molière radius to reduce the occupancy and sufficient light output to allow readout with silicon photodiodes appropriate for the low photon energies of interest to us (normal PMs are excluded in the magnetic field). $PbWO_4$, a scintillating crystal, is well suited for our application. The detector will be cooled to $-25°$C in order to increase the light output.

1.4.6 Forward muon spectrometer

The forward muon arm, as described in an addendum to the ALICE LoI [4], is designed in order to cover the complete spectrum of heavy quark resonances, i.e. J/ψ, ψ', Υ, Υ' and Υ''. It will measure the decay of these resonances, both in proton–proton and in heavy-ion collisions, with a mass resolution sufficient to separate all states. The first section of the spectrometer is a composite absorber ($\sim 10\,\lambda_{\rm int}$), one metre from the vertex, made with layers of both high and low Z materials to reduce multiple scattering and particle leakage. This is followed by a large dipole magnet with 3 Tm field integral, placed outside the L3 magnet, and 10 planes of thin high-granularity tracking stations. A second absorber ($\sim 11\,\lambda_{\rm int}$ of iron) at the end of the spectrometer and two more tracking planes are used for muon identification and triggering. The spectrometer is shielded throughout its length by a dense tungsten tube of about 60 cm outer diameter, which surrounds the beam pipe.

The muon-arm addendum was accepted by the LHCC in the second half of 1995. In the present document we concentrate on aspects of the muon spectrometer which are relevant to its overall integration and functioning in ALICE (e.g. mechanical integration and installation, services and infrastructure requirements, trigger and data flow, interference of the L3 and muon magnets). Technological choices for the tracking chambers and reliable estimates of background hits induced in the TPC require further optimization of the absorber configuration and, in particular, experimental verification with test beams. A number of tests are being planned or are under way and we expect to submit an addendum to this technical proposal, with a more detailed description of the muon arm, in the second half of 1996.

1.4.7 Large rapidity detectors

A multiplicity counter array close to the interaction region will measure the pseudorapidity distribution of charged particles over a large fraction of the phase space ($|\eta| < 4$). A number of options are being studied, including a set of micro-channel plates surrounding the beam pipe.

Two sets of compact calorimeters will be used to measure and trigger on the impact parameter of the collision. They are made of tungsten and lead with embedded quartz fibres, read out by PMs, and located on both sides in the machine tunnel, about 92 m from the interaction region.

1.4.8 Significant changes since the LoI

Inner tracking system

The number of planes in the ITS and their positions have been optimized in order to increase the performance at constant cost. We have reduced the diameter of the beam pipe and added an innermost sixth layer of silicon pixel detectors, thus increasing the impact parameter resolution by a factor of about two. This permits a direct measurement of open charm production in heavy-ion reactions, of importance both as a signal of the high-temperature early phase of the collision and as normalization for J/ψ production. In addition, the signal-to-background ratio for electron pairs is improved by rejecting semi-leptonic open charm decays.

The outer layers of the ITS have been reduced in size (and therefore cost) while keeping the angular coverage and tracking performance of the LoI layout. This has been made possible by reducing the effects of multiple scattering in the inner walls of the TPC, which are now both thinner and at a smaller radius.

Finally, we have decided on double-sided silicon micro-strip detectors for the two outer layers, instead of a combination of silicon drift (5th layer) and micro-strip gas chambers (MSGC, 6th layer), which was the other option in the Letter of Intent. The difference in cost between these options is marginal, whereas pattern recognition, dE/dx resolution for PID and response to low-p_t inclined tracks are all in favour of the current solution. The ITS is now only made with silicon detectors, which promises an additional saving in terms of services, cooling and mechanics.

PID array

From the four options presented for the large area PID array in the LoI —Pestov, PPC, scintillation TOF counters and RICH detectors— we have retained only Pestov and PPC as potential candidates. While the R&D on all four systems is progressing satisfactorily, we reduced the number of options on the grounds of complexity of the large size system and of resources available within the collaboration. Based on current performance, the Pestov counter would be our preferred solution, with the PPC as backup. The final decision will be taken by the end of 1996 on the basis of tests with large area prototypes.

Performance requirements, technological risks and cost implications are different for the High Momentum PID array (HMPID). The RICH detector, which is currently the most advanced in terms of R&D and testing of large prototypes, seems the most promising solution and is therefore our baseline choice.

Photon Spectrometer

The orientation of PHOS has been changed from a rectangular array (large rapidity, small azimuthal coverage) to a ring segment (small rapidity, large azimuthal coverage) in order to allow a modular construction using identical non-pointing crystals. In addition, the current shape increases the acceptance for low-p_t meson decays and, accordingly, we have decreased the area of PHOS from 25 to 18 m^2.

Multiplicity array

Decreasing the diameter of the beam-pipe to improve the vertex resolution and the introduction of the muon-arm absorber have led to a redesigning of the multiplicity detector. It will now be placed outside the beam-pipe (instead of inside, in a secondary vacuum) and has a reduced acceptance, $-3.5 < \eta < 4.5$ instead of $|\eta| < 5$.

2 The Inner Tracking System (ITS)

2.1 Introduction

The main purpose of the ITS is the detection of secondary vertices (hyperons and charm) and the stand-alone track finding of low-p_t charged particles, down to a p_t of ~ 20 MeV/c for electrons. It has a number of additional roles: improvement of the momentum resolution at large momenta, momentum reconstruction of the low-energy particles and their identification via dE/dx in the $1/\beta^2$ region, and definition of the angles of the tracks for HBT interferometry analysis.

A general view of the inner detector is shown in Fig. 2.1.

The system consists of six cylindrical layers of coordinate-sensitive detectors, covering the central rapidity region ($|\eta| \leq 0.9$) for vertices located within the length of the interaction diamond (2σ), i.e. 10.6 cm along the beam direction (z).

The geometrical dimensions and the technology used in the various layers of the ITS are summarized in Table 2.1.

Table 2.1: Inner Tracking System dimensions.

Layer	Type	r (cm)	$\pm z$ (cm)	Area (m^2)
1	pixel	3.9	12.25	0.06
2	pixel	7.6	16.3	0.17
3	drift	14	21.1	0.37
4	drift	24	29.6	0.89
5	strip	40	45.1	2.27
6	strip	45	50.4	2.85
Total area = 6.61 m^2				

The number and location of layers and the thickness of material are the result of extensive simulation presented in Section 11.3. The design guidelines for the geometry of the ITS take into account the very large track density, the need for minimizing the multiple scattering, the need for good track matching between the ITS and the TPC and the compromise between impact parameter resolution and angular resolution.

The granularity required for the innermost planes, given the expected high multiplicity of charged particle tracks, can only be achieved with silicon micro-pattern detectors with true two-dimensional readout, such as Silicon Pixel Detectors (SPDs) and Silicon Drift Detectors (SDDs).

We shall use both SPDs and SDDs. Silicon pixel detectors are better adapted to the geometry and the assembly constraints of the first two layers at small radii, and assure optimum granularity. In addition, they are fast enough to operate at high rate, and they will be read out in a separate 'fast' data stream (together with the muon spectrometer) to provide the event vertex.

For the third and fourth layers, which cover a larger area with lower track density, we have chosen Si drift detectors, which feature similar performance at lower cost, and in addition provide dE/dx information for particle identification, which is crucial for the lepton pair analysis.

At present, pixel detectors close to the specification of ALICE have been used in the heavy-ion experiment WA97 and the linear SDDs are produced by industry using standard processes.

At larger radii, the requirements in terms of granularity are less stringent, therefore the use of silicon strip detectors, which are a well-proven and cheaper technique, is possible. The large sensitive area needed for the outer layers involves the use of a great number of detectors. In such a case, reliability, ease of operation, and the established capability of industrial production become major points in favour of a more conservative choice.

We propose to use double-sided silicon microstrips with a small stereo angle as they offer the possibility of correlating the pulse height readout from the two sides, to help resolve ambiguities.

The main parameters for each of the three detector types are indicated in Table 2.2: spatial precision, two-track resolution, pixel size, number of channels of an individual detector, total number of electronic channels, dissipated power both in the central region and in the end caps.

The momentum and impact parameter resolutions are dominated by multiple scattering for most of the momentum range of relevance in ALICE. We therefore want to keep the amount of material within the acceptance at an absolute minimum. We aim to keep the radiation length of all passive material (support, electronics, cooling, cables, etc.) at a level not exceeding that of the active detector material (typically 300 μm silicon, ~ 0.3 % of X_0), i.e. at a total thickness of ~ 0.6 % of X_0, averaged over one layer.

The large number of channels in the layers of the ITS requires a large number of connections from the front-end electronics to the detector and to the readout. The requirement for a minimum of material within the acceptance does not allow the use of conventional copper cables near the active surfaces of the detection system. Therefore, new techniques with aluminium ultrasonic bonding on Al/kapton multilayer micro-cables need to be investigated; one of these possibilities has been developed by the ALICE Collaboration.

The different detector technologies require specific front-end electronics. However, we foresee the unification of electronic standards (type of links, protocols) at

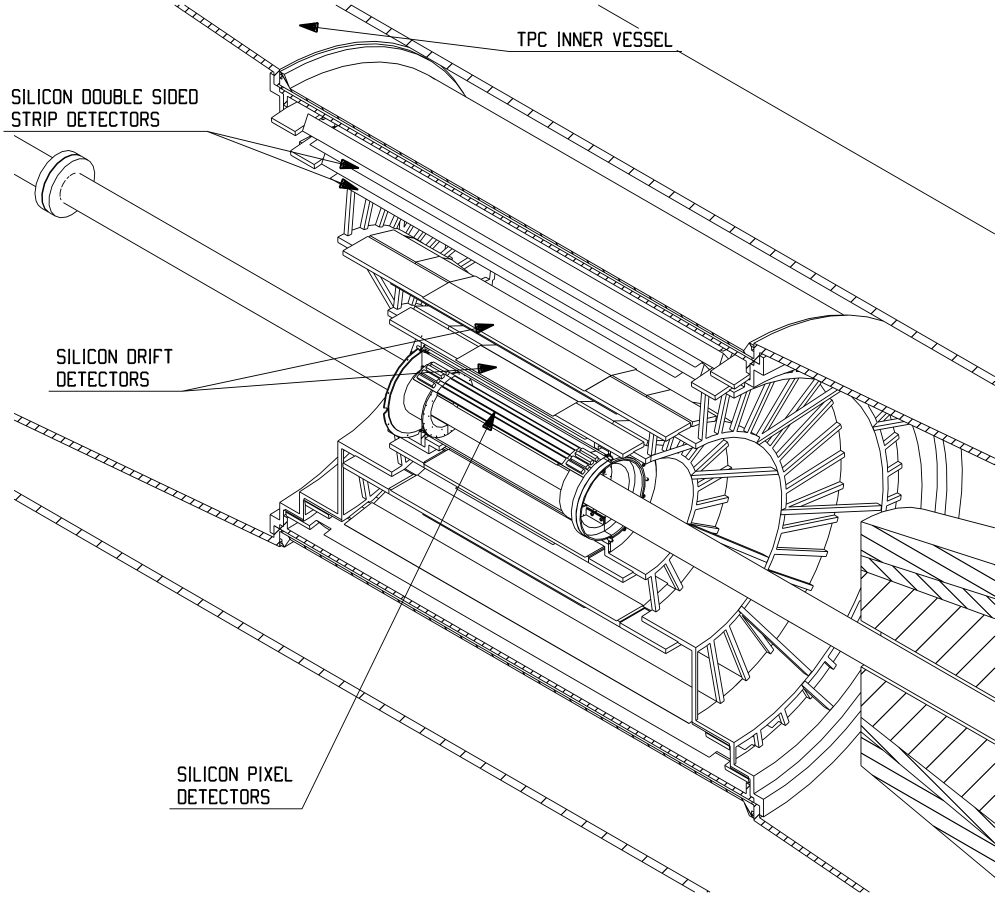

Figure 2.1: General view of the Inner Tracking System.

Table 2.2: Parameters of the various detector types. A module represents a single detector.

Type	Spatial precision (μm)		Two-track resolution (μm)		Cell size (μm)	Module size (mm)	Channels per module	Electr. channels total (k)	Power dissipated	
	$r\varphi$	z	$r\varphi$	z					barrel (W)	end-caps (W)
Pixel	15	90	100	600	50 × 300	13.3 × 81	65536	14420	450	–
Drift	20	30	200	420	–	70 × 87	2 × 384	192	350	1950
Si Strip	30	860	200	7000	95 × 41000	75 × 42	2 × 768	2620	<3000	<1000

the level of optical links, i.e. at the front-end outputs.

The detectors and their front-end electronics will produce a large amount of heat which must be removed while maintaining a very high degree of temperature stability. In particular, SDDs demand temperature stability and monitoring in the 0.1°C range. For these reasons, particular care is being taken in the design of the cooling system and different schemes are considered for the various layers. For SDDs, evaporative cooling at room temperature is at present the proposed solution, but is still under development. For the other layers a water cooling system should be sufficient.

2.2 The silicon pixel layers

2.2.1 Introduction

SPDs have been chosen for the two innermost planes due to their high granularity and excellent precision and two-track resolution. A silicon detector with two-dimensional segmentation combines the advantages of unambiguous two-dimensional readout with the geometrical precision, double-hit resolution, speed, simplicity of calibration and ease of alignment of silicon microstrip detectors. In addition, the high segmentation leads to low individual cell capacitance, resulting in an excellent signal-to-noise ratio at high speed. The price to pay is a large increase in the number of channels. The production of practical silicon pixel devices has been made possible by the continuous progress in CMOS microelectronics, which today allows the integration of complicated functions on miniaturized electronic chips. These can be placed close to the sensor elements, allowing signal processing and information extraction without prohibitive volumes of cables and electronics. Nowadays much of the electronics for the readout of a Si pixel cell can be made to fit within a surface of the size of the detector cell (~ 0.025 mm^2 for the recently built Omega3 chip [1]).

2.2.2 Design criteria

In this section we shall illustrate the main requirements which guide the pixel design.

Very good tracking and vertexing capabilities in the inner tracker can be obtained with two layers of pixels, one at 3.9 cm and another at 7.6 cm from the beam, with a cell size of 50 μm ($r\varphi$) $\times$ 300 μm (z). The differences from the set-up illustrated in the ALICE Letter of Intent [2, 3] (the extra inner layer and the reduced cell size) were introduced to give access to the charm signal with an improved impact parameter resolution (see Section 11.9). Moreover, robustness in low-p_t track finding is strongly enhanced by the presence of the extra plane and by the improved space accuracy of the cells. The presence of two pixel layers also strengthens the possibility of limited stand-alone tracking capabilities for the ITS (namely measurement of the primary vertex position) at the highest event rates envisaged for the Debye screening studies, i.e. with the muon arm, in Ca–Ca collisions.

The Level 0 strobe would arrive at the detectors about 1.2 μs after the passage of the particles. The maximum envisaged interaction rate (in pp and high-luminosity Ca–Ca running) will be about 500 kHz, implying a maximum integration time for the pixel signals of the order of a few hundred nanoseconds. The readout time must be kept to within about 200 μs in order to match that of the muon chambers.

If the readout time is kept below 100 μs for at least one of the two layers, there is the possibility that a multiplicity trigger at Level 2 can be provided with the full information available.

The amount of material, as is the case everywhere else in the inner tracker, should be kept close to 0.6 % X_0 per layer. This requirement is particularly strict for the pixel detectors, as both charm and lepton physics are very sensitive to material close to the vertex.

Such requirements do not seem to be out of reach given the present status of the technology and the remarkable speed of on-going developments in the field of microdetectors. These requirements are different from those of the LHC pp experiments: they are milder in terms of strobe accuracy and radiation hardness. In addition, the total area and the number of channels foreseen are considerably lower than is being proposed for the pp experiments.

On the other hand, ALICE aims to study essentially soft phenomena, placing strict requirements on the tracking efficiency for low-p_t particles and therefore on the permissible amount of material. We have developed a design which is relatively conservative in many aspects, identifying a few areas in which to concentrate our future R&D efforts. We nonetheless intend to follow advances in this field and to be prepared to introduce more advantageous schemes, should they become viable in the future.

We plan to use a hybrid approach, with Pb–Sn solder-bump bonding of detectors to the readout chips and wire-bonding for the connection of the chips to the readout and control bus. This has worked well for the WA97–RD19 SPDs (more details in the next section). We plan to use binary readout since, with 50 μm pitch in $r\varphi$, the errors in track parameter determination, in the momentum range of interest in ALICE, are dominated by multiple scattering rather than by the measurement error. Therefore, a further increase in accuracy does not seem necessary.

The material budget requirements of ALICE are much more stringent than those of WA97. We therefore plan to thin down the detector and the electronics chips and to eliminate the ceramic substrate. We would separate the bussing and support functions, using a thin multilayer circuit as bus and a carbon-fibre structure as support. We plan to modify the column shift-register approach of the WA97–RD19 architecture in order to achieve on-chip data sparse readout. This requires fewer clock cycles, i.e. a lower frequency clock to read out the data in a given time period and eliminates the need to propagate very fast clocks inside the chip, which should allow easier and safer operation, and a reduction in power consumption. An additional advantage of sparse readout is a reduction in the number of bus lines.

2.2.3 Silicon pixel R&D

We base much of our project on the technology developed over the last five years within the framework of the RD19 [4] project. Much experience in this field has been accumulated within the heavy-ion experiment WA97 [5]. This is the first experiment to employ silicon pixel devices and the majority of the people involved in the ALICE pixel group are taking part in it. The WA97 Collaboration studies the production of strange and multi-strange baryons and antibaryons in central Pb–Pb collisions at the SPS. They are currently operating a telescope of seven 5×5 cm^2 planes of silicon pixels (500 000 channels), built in close collaboration with RD19, and have a second pixel telescope in preparation.

The Omega2 detectors

In the WA97–RD19 approach the readout electronics is built on a separate chip from the detector ('hybrid' approach). The readout is binary, and data are 'clocked out' along column shift-registers. The basic building block of the WA97 pixel detectors is the ladder: a matrix of 96 columns and 64 rows of rectangular diodes (pixels), ion-implanted on a high-resistivity (~ 5 kΩ cm) n-type silicon substrate of 300 μm. The dimensions of most cells are $75 \times 500\,\mu\text{m}^2$. Each pixel is connected to the virtual ground of a front-end amplifier on the readout chip by a Pb–Sn solder bump with a diameter of 38 μm. Six chips, each containing 16×64 pre-amplifier cells, are bump-bonded to one ladder.

Each cell contains a low-noise front-end amplifier followed by a comparator with an adjustable threshold, an adjustable delay line and coincidence and memory elements, which store the bit indicating a particle hit during the time frame selected by an external strobe. Permitted values for the threshold range from 4000 to 15 000 electrons, while the internal delay is adjustable between 100 ns and 1 μs.

The rear of the detector wafer is ion-implanted in order to provide a common n^+ ohmic contact without segmentation. A bias voltage of 30 to 70 V is applied to this rear side. Six ladders, spaced by a few millimetres, are glued on a ceramic support which accommodates the bus lines. The lines are isolated from the back planes of the chips by an intermediate ceramic slab glued onto each ladder. Electrical connections between the bus lines and bonding pads on the readout chips are made by ultrasonic wire-bonding. The detector bias connections are made by wire-bonding to the rear side of the ladder. A bias filter network is placed on the ceramic substrate, beside each ladder. A flexible multilayer circuit is wire-bonded to the lines of the ceramic substrate. This flexible circuit provides connection to an intermediate electronics card for the supply, control and data lines of the ceramic substrate.

Detailed studies of the performance of the WA97–RD19 planes have been performed [5–7]. The efficiency is very high, practically 100 % for working cells. The average plane efficiency in WA97 is better than 95 %, inefficiencies being due only to faults. The random trigger noise after masking of 'always-on' channels is better than 10^{-10} spurious hits per pixel per trigger. The measured position accuracy of 24 μm is as expected for this pixel size, i.e. close to pitch/$\sqrt{12}$. In the Pb runs, events were recorded with track densities as high as 5 cm^{-2}, which still corresponds to an occupancy of only 2×10^{-3}. A fully reconstructed event with 40 tracks (no *B* field) is shown in Fig. 2.2.

While the overall performance of the WA97 pixel telescope is more than satisfactory, the preparation and operation of the detectors gave us a detailed understanding of certain critical issues which will need to be taken into account for future applications [7], e.g. the need for systematic tests at all levels in the detector production chain and the need for a better control of the internal delay and threshold settings for the individual chips.

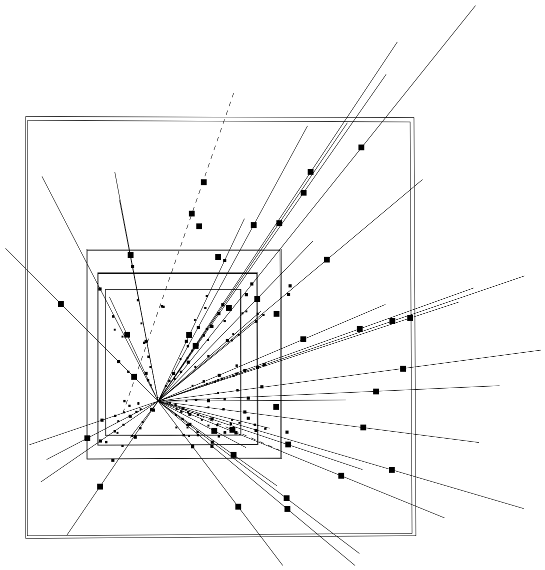

Figure 2.2: A Pb–Pb event with 40 tracks reconstructed in the WA97 pixel telescope (autumn 1994).

The Omega3 detectors

During the course of 1995 a new chip (Omega3) was produced. We plan to use it for building a new telescope, which would enlarge the WA97 phase-space coverage. We changed from the 3 μm SACMOS3 technology of Faselec AG, Zürich, to the 1 μm SACMOS1 process. This allowed us to squeeze enhanced operation (extra test features, better control of the individual cell delay, possibility of a fast column OR output) into a smaller cell $(50 \times 500\,\mu\text{m}^2)$ than for the Omega2 chip. The first bump-bonded Omega3 chips were recently delivered. Preliminary results [1] indicate a large improvement in the timing and threshold precision with respect to Omega2, and better noise figures.

While the Omega2-based detectors were mounted on a single layer ceramic support, we plan to use a multilayer circuit for the Omega3-based detectors. This allows us greater freedom in the design of the support. For example, the current design foresees separate delay bias lines for the individual chips. It also allows us to gain experience in the design and use of thin-layer circuits. We are currently testing the first samples of a multilayer ceramic designed for use with the Omega3 ladders and produced in the CERN PCB workshops. It consists of three layers (ground, supply and a low-density signal layer) obtained by standard thick-layer technology, plus an upper, dense signal layer fabricated using the FODEL process of Du Pont de Nemours SA, Geneva.

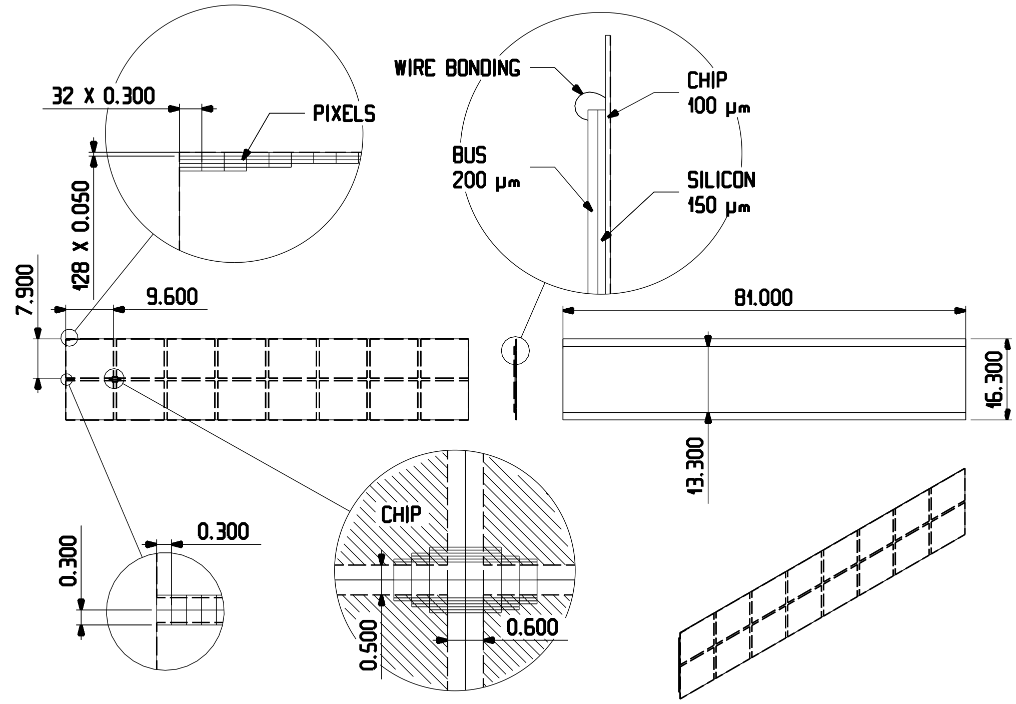

Figure 2.3: The silicon pixel ladder.

At the same time, we are exploring other possibilities for producing thin flexible circuits without ceramics, following the on-going developments in multilayer circuit technology. In the near future these should lead to the possibility of making circuits with a layer thickness below 50 μm, 100 μm pitch lines and via diameters below 100 μm. One possibility we have considered is the use of the DYCOstrate process of Würth Elektronik, Rot am See (Germany), which consists in superimposing several polyimide–Cu layers. The precision of this process allows track pitches close to 100 μm, with vias of 80 μm diameter. At present, however, a pad of 250 μm diameter is required around each via, making their use problematic in regions of small track pitch.

Another interesting process is being studied at CERN in the PCB workshop: building successive layers of polyimide with superimposed sputtered Al layers, which are then etched. The layers can be as thin as 20–30 μm, with 100 μm lines and 50 μm vias. One needs to start with a fairly thick Al layer, which can, however, be removed later. This technique looks promising and we would like to try it in WA97 in the near future. Various other options for the use of multi-chip module technology are being studied by the RD19 Collaboration.

While at present WA97 employs 300 μm thick detectors with 300 μm thick electronics, in ALICE we plan to use 150 μm thick detectors with 100 μm thick chips. R&D towards thinner assemblies has started: in RD19 detectors as thin as 150 μm and Omega2 chips thinned down to 80 μm have already been succesfully tested, including ‘thin detector – thin chip’ combinations, although with poor yield [8, 9]. The manipulation of these fragile assemblies has still to be optimized.

2.2.4 Design of the pixel layers

In this section we describe one possible scheme for the implementation of two layers of silicon pixel detectors in ALICE.

The basic unit is a 81 mm long, 13.3 mm wide and 0.15 mm thick silicon detector ladder (Fig. 2.3), segmented into 256×256 cells, and bump-bonded to 16 readout chips, each equipped for the readout of 4096 cells. The dimensions of most cells are 50×300 μm^2, with larger cells in the stitched regions between chips. The readout chip dimensions are 7.9 mm in $r\varphi$, 9.6 mm in z and 0.1 mm in thickness.

The thin multilayer bus is glued onto the opposite (back plane) side of the detector. The chips protrude laterally in order to allow wire-bonding to the bus (see Fig. 2.3). The z coverage is obtained by lining four and three ladders for the outer and inner layers of pixels, respectively (see Fig. 2.4).

Two multilayer circuits are used for each line of ladders: one for the two ladders on the left and one for the two ladders on the right (for the outer layer), while for

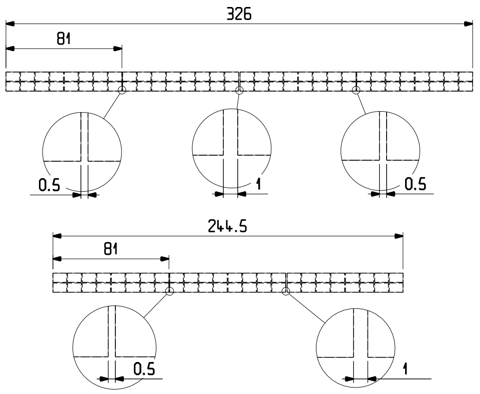

Figure 2.4: Ladder lines for the inner (three ladders) and outer (four ladders) planes.

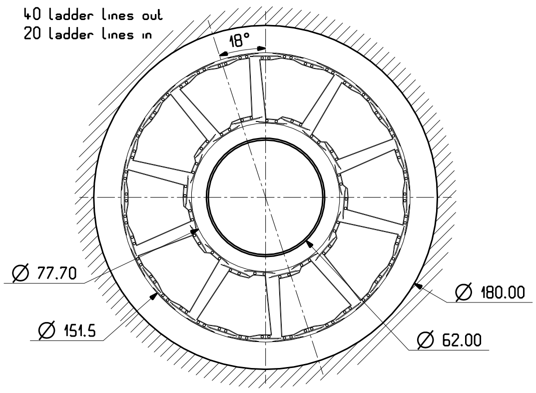

Figure 2.6: Cross-section of the pixel barrel.

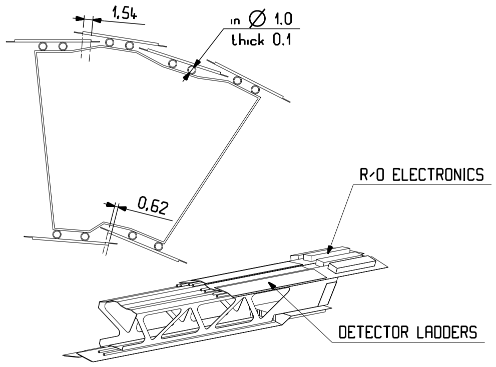

Figure 2.5: View of a sector of the pixel barrel.

the inner layer one of the circuits serves only one layer. Each circuit contains two busses, one for each row of chips, one row consisting of 8 chips from each of two adjacent ladders. In order to avoid further wire bonding or the use of connectors for the bus lines, we plan to extend the multilayer circuits in the z direction and to use them for mounting readout and control electronics. Each bus (about 6.5 mm wide) carries about 60 lines: 12 for the sparsed and binary encoded data, 16 separate enable lines, 16 separate delay bias lines and about a dozen common bias and digital control lines. With a 5 MHz clock, at 1 % occupancy for the inner layer (about 50 hits per chip), 16 chips could be read out in less than 200 μs. For the outer layer, at 3×10^{-3} occupancy (about 15 hits per chip), 16 chips could be read out in about 50 μs.

Mechanical support for the pixel ladders will be provided by a 200 μm thick carbon fibre structure. A carbon fibre sector will support four lines of four ladders each from the outer layer, and two lines of three ladders each from the inner layer, for a total of 22 ladders, as shown in Fig. 2.5.

The approach of using a single structure for the support of both layers has the advantage of minimizing the material for a given rigidity.

Using the current Omega2 figure of 30 μW per cell for the power dissipation, we reach a total power generated by the readout chips of around 450 W, to which should be added the power dissipated on the readout cards. It does not seem possible to extract this much heat with a $\Delta T \lesssim 5°$C using gas cooling, and we therefore plan to integrate water cooling pipes into our design. A simple scheme is shown in Fig. 2.5: the ladders are glued onto a pair of 1 mm cross-section, 100 μm thick tubes. We are also studying the possibility of integrating the coolant flux directly into the carbon fibre structure. In this way, the detector ladders would be positioned directly on the carbon fibre support, which would make it possible to clip the ladders in place instead of gluing them. This would be an advantage for the long-term maintenance of the system.

As shown in Fig. 2.6, 10 sectors are needed to close the barrel. The inner layer is positioned at 3.9 cm and the outer layer at 7.6 cm from the beam line. The main design parameters for the two layers are shown in Table 2.3. The pseudorapidity coverage is the minimum guaranteed for events occurring at the 2σ edge of the interaction region ($\Delta z = 10.6$ cm). The occupancy figures for central events are $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$.

The average thickness seen by a particle traversing the two pixel layers is about 1.2 % of X_0.

Table 2.3: Main design parameters.

	Layer 1	Layer 2
r	3.9 cm	7.6 cm
Δz	24.5 cm	32.6 cm
η coverage	1.34	1.17
No. ladders	60	160
Surface	0.06 m^2	0.17 m^2
No. cells	3.9×10^6	10.5×10^6
Power	120 W	330 W
Occupancy	1.0×10^{-2}	2.7×10^{-3}

2.2.5 Readout and control systems

Each bus serves a row of 16 chips and is equipped with its own set of control electronics, mounted on the flexible bus circuit at the side of the barrel, outside the acceptance, and mostly integrated in one or a few dedicated chips.

Data are zero suppressed on the readout chips and binary encoded addresses of the hit cells are read via a 12-bit bus. Several sparse readout schemes are presently under study [10–12]. Communication between the control electronics and the readout chips is established via several control lines (e.g. DATA-READY, ENCODE, START-OF-READ, END-OF-READ, RESET) common to all the chips, plus 16 individual ENABLE lines. The logic for the generation of the ENABLEs will provide a binary encoded address of 4 bits, which will be merged with the binary encoded address of the cells (12 bits) into a 16-bit word containing the absolute address of a cell in the 16 chips row. The chip row address can then be generated by remote electronics.

The encoded data are stored locally in a memory, ready to be transferred to a central (remote) buffer memory via a fibre optics link. The fibre optics link should be bidirectional, for writing data into the control logic.

About 20 d.c. control lines will be needed, in order to supply the chips with the bias currents generated by a set of DACs (Digital-to-Analog Converters), according to a digital code sent on the bidirectional optical link. Another solution under study is to build the DACs directly on the readout chips. The loading of the digital code could then be carried out serially, which would allow a further reduction in the number of bus lines.

A slow control system with computer independent interlocks will be implemented for monitoring the temperature and the values of the different power supply and bias voltages.

2.2.6 Radiation effects

Estimates of the radiation flux have been performed using FLUKA [13] (see Section 12.5.3). The effect of a 30 mm radius, 0.6 mm thick Be beam pipe was included. For the 10 years running scenario illustrated in Section 12.5.3, the total dose would be about 120 krad (1st layer), and about 40 krad (2nd layer). The neutron fluence should be below 10^{12} n/cm^2 (1st layer) and below 5×10^{11} n/cm^2 (2nd layer). Such values do not seem to pose problems of bulk damage in the detector silicon. As for the electronics, although it has been found that CMOS chips can be damaged already with a dose of a few krad, several of the standard CMOS techniques remain operational up to over 100 krad [14].

We plan to study the effects of irradiation on the Omega3 chips. It is expected that these surface effects, essentially due to charge trapping in the insulator layer of CMOS transistors, should be progressively reduced as the technology moves to smaller structure sizes and thinner insulator layers. We shall keep following closely the on-going developments towards radiation-hard electronics. We also expect to take advantage of the solutions which will be developed for the much more demanding LHC pp applications. The replacement of the inner layer ladders (only 0.06 m^2 in total) after some years of operation is also a possibility.

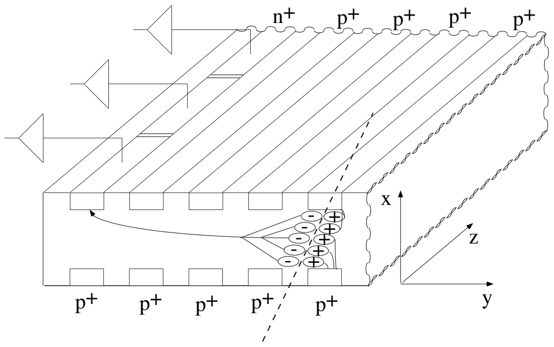

Figure 2.7: The operation principle of a SDD.

2.3 The silicon drift layers

2.3.1 Introduction

SDDs [15, 16], like gaseous drift detectors, exploit the measurement of the transport time of the charge deposited by a traversing particle to localize the impact point in one of the dimensions, thus enhancing resolution and multi-track capability at the expense of speed. They are therefore well suited to this experiment in which very high particle multiplicities are coupled with relatively low event rates.

A linear SDD, shown schematically in Fig. 2.7, has a series of parallel, implanted p$^+$ field strips, connected to a voltage divider on both surfaces of the high-resistivity n-type silicon wafer. The voltage divider can be integrated in the detector substrate itself. The field strips provide the bias voltage to fully deplete the volume of the detector and they generate an electrostatic field parallel to the wafer surface, thus creating a drift region.

Electron-hole pairs are created by charged particles crossing the detector. The holes are collected by the nearest p$^+$ electrode, while the electrons are focused into the middle plane of the detector. The superimposed drift field forces the electrons to drift parallel to the wafer surface towards the edge of the detector where they are collected by an array of anodes composed of n$^+$ pads. The small size of the anode, and hence its small capacitance (~ 50 fF), imply low noise and good energy resolution. Fig. 2.8 shows a simulated shape of the electric field in a SDD, showing the characteristic 'gutter' which directs the electrons towards the anodes.

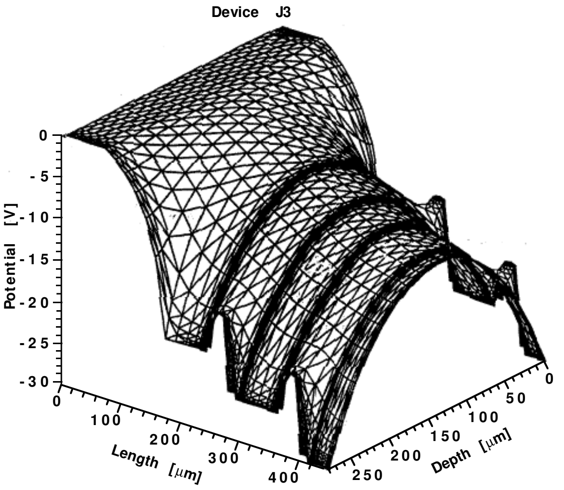

Figure 2.8: Simulated shape of the electric field in a SDD.

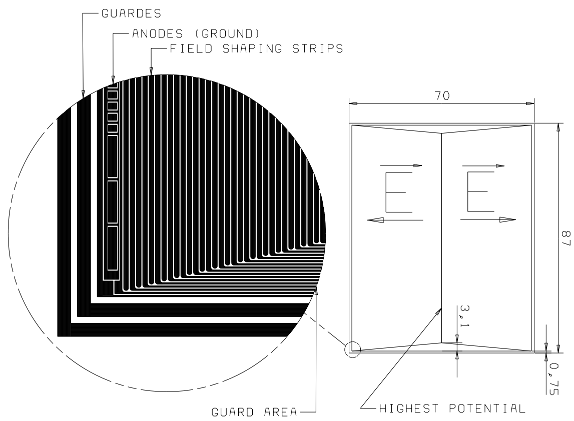

Figure 2.9: Sketch of the two-directional SDD with a blow-up of a corner showing the anode region.

Owing to the diffusion and mutual repulsion during the drift, the electrons reach the anode region with a Gaussian distribution. The coordinate transverse to the drift direction is given by the centroid of the collected charge. The coordinate along the drift direction is measured by the centroid of the signal in the time domain, taking into account the amplifier response.

A space precision of better than 30 μm in both coordinates can be obtained for most of the detector surface. Very close to the anode array the transverse coordinate is measured with less accuracy as the charge is collected on one pad only. A two-track resolution of 200 μm is feasible by performing a waveform analysis [17, 18]. In NA45, a two-track resolution of 180 μm was obtained in the drift direction by applying a double-pulse fit.

Temperature dependence

The drift time depends on the drift velocity of the electrons and therefore on the electron mobility ($v_{\mathrm{drift}} = \mu \cdot \mathrm{E}$). The mobility of the electrons is temperature dependent, being $\mu \sim T^{-2.4}$, so that a temperature variation of $\sim 0.1°$C corresponds in our design, on average, to a ~ 14 μm change of the position of the reconstructed points. This requires an accurate stabilization of the temperature, and periodic calibration of the drift velocity.

We plan to thermostabilize the detectors (i) by an evaporative cooling scheme to drain the heat from the front-end electronics, and (ii) by physically separating the electronics and the detector.

Calibration

To calibrate the detector by charge injection, a MOS injector can be integrated near the central electrode. Another option would be to inject charge by means of an infrared laser, for example, Nd-YAG, for which the absorption length in silicon is about 1 mm. In addition, this method could be used to monitor the position of the detectors in space. A continuous control will be performed using residuals of high-momentum tracks.

2.3.2 Design of the SDD

We foresee a design in which the sensitive area is split into two parts so that electrons drift in opposite directions in the two halves of the detector. Such a detector is shown schematically in Fig. 2.9. The central electrode is the p^+ strip at the highest, negative potential. In this way, the maximum drift length and the required drift field are reduced by a factor of two with respect to a SDD having the anode pad array on one side only. Moreover, the two half-chambers share part of the guard region, limiting dead zones. A similar detector has been built and tested [19] for the UA6 experiment at CERN, and several designs based on the same concept have been developed for the STAR experiment at RHIC and for the ALICE R&D programme.

Following the results of the development programmes, the detectors to be used in ALICE are still evolving but, based on the experience accumulated up to now, we are able to describe here a realistic detector design which can be used to evaluate the system parameters. It will be made from a 5-inch, high-resistivity, neutron-doped silicon wafer of 300 μm thickness. For optimum use of the wafer, each SDD would be 70 mm in the drift direction and 87 mm in the other direction.

The high voltage is brought to the detector in the centre, and the decreasing voltages are generated onto the detector itself by an integrated voltage divider. The last few (typically four) strips nearest to the anodes will be connected independently, in order to modulate the field in the collection region; the voltages for these strips are in any case moderate and can be carried on the same bus as the

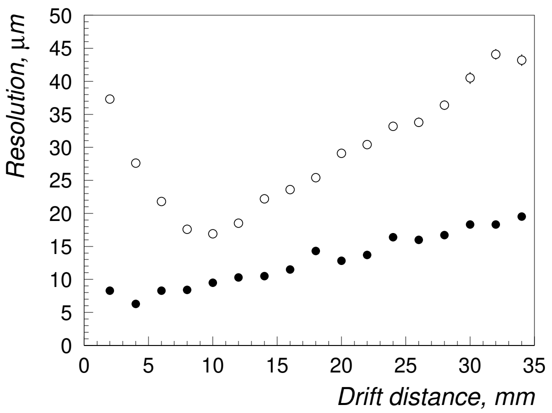

Figure 2.10: SDD resolution, for single tracks, along and perpendicular to the drift direction (full and open circles, respectively) as a function of the distance from the anode.

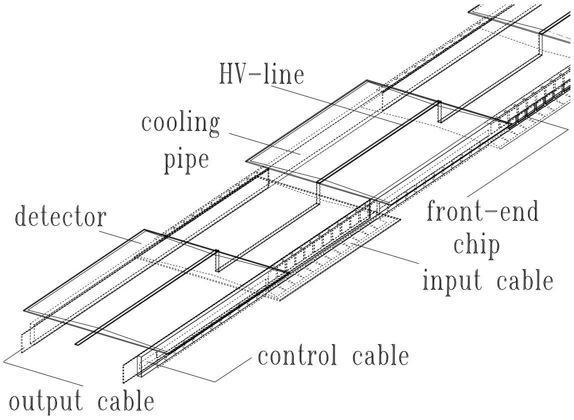

Figure 2.11: SDD assembly scheme.

power supply for the electronics. There are two active regions, from the centre to the edges, with a guard region on both sides to allow the potential to degrade smoothly to the edges. We foresee the need for at least 3.1 mm of guard in the centre, to degrade the full bias voltage to ground in 25 V steps of 40 μm each.

Each detector has 384 anode pads on each edge, 378 facing the full length drift region, with a pitch of 210 μm, and three on each side facing the shorter drift region. The total dead area including all guard areas and the safety edges around the detector will be below 14 %. We plan to work with a central high voltage of 2000 V, i.e. a drift field of ~ 600 V/cm. With this field, we have measured drift velocities of ~ 6.7 μm/ns, which would imply a total drift time of ~ 5 μs. The charge cloud will reach the collecting anode with a r.m.s. radius of ~ 173 μm for the maximum drift length, and 116 μm on average.

We have performed a simulation of the detector performance [20, 21], including pessimistic assumptions about the electronics noise (500 e ENC), and we have obtained an average resolution in Pb–Pb events of 20 μm along the drift direction (which is the relevant coordinate for momentum resolution), and 31 μm in the transverse direction. Because of the effects of diffusion, the resolution is not constant across the detector: the local resolution for single tracks is shown in Fig. 2.10 as a function of the distance of the particle impact point from the anodes.

2.3.3 The SDD ladders

The SDDs will be mounted on linear structures called ladders each holding 5 detectors for layer 3, and 8 detectors for layer 4. The layers will be composed of 14 and 22 ladders, respectively. Each detector will first be assembled together with its front-end electronics as a unit, which we shall refer to as a *module*, which can be fully tested before being mounted on the ladder.

We foresee using lightweight carbon fibre structures with integrated cooling arteries to support the detector modules, front-end electronics, and cables. The modules will be mounted as shown in Fig. 2.11, i.e. alternating on the two sides of the supporting ribs, with the anode side towards them.

The detectors will overlap so that there will be no dead regions. The electronics will be mounted on a small board, which will be positioned along its edge on the cooling channels, i.e. perpendicular to the detector plane. Such an assembly is possible thanks to the use of Al micro-cables which can be welded directly to the detectors and to the electronics. The cables will run on the side of the ribs opposite the electronics. The micro-cables, including signals and power buses, will represent 0.024 % of X_0 for plane 3, and 0.036 % of X_0 for plane 4, having averaged over the full surface of the plane in both cases.

The SDDs will be positioned so that the electrons drift orthogonal to the beam axis and therefore to the magnetic field. The low magnetic field would have a marginal effect on the electron cloud formation and essentially none on the charge transport, as the Lorentz force is compensated by the confining field. The drift plane would just be displaced somewhat relative to the middle plane of the detector.

Backup alternatives with different cooling and assembly schemes are also being studied.

Cooling

We have performed a detailed analysis of different cooling options, taking into consideration efficiency and the possibility of reliable control [22]. We favour the use of an evaporative cooling system, which is very efficient and provides a simple way of controlling the temperature via pressure. Such a system would be even more attractive if further stages of the readout, in addition to the preamplifier, are mounted near the detectors.

A prototype cooling channel of 1.1 m length and 3 mm diameter has been succesully operated to drain, via the evaporation of freon, up to 25 W of power [22]. A full low-mass, reliable, leakless (slightly underpressurized) system based on evaporative cooling is currently being developed.

2.3.4 Silicon drift R&D

Although SDDs are not yet commonplace in high-energy physics experiments, they have been successfully employed in the NA45 experiment with ion beams at the CERN SPS [23, 24], and more recently in the WA93/WA98 experiment.

In NA45, a cylindrical SDD has operated with an average track density of 13 particles cm^{-2} over the whole detector, and up to 0.66 particles mm^{-2} in the central region, values which are much larger than the ones expected during operation in ALICE. The overall efficiency, including detector and electronics, was 96 %, and the homogeneity, again including the effect of both charge collection and electronics gain, was better than 6.5 %.

Within the ALICE Collaboration, a specific R&D effort is in progress, with two main goals: the establishment of SDD production in industry, and the optimization of a detector design which, while being reliable for production in series, would minimize the dead regions and optimize the distribution of heat generation on the detector. Linear SDDs have been successfully produced at BNL and, more recently, the LBL labs in the USA, but for ALICE, given the large number of detectors needed, it is essential to establish industrial production and to carefully optimize the design.

In a dedicated R&D effort (INFN DSI project) we are developing a detector design (shown in Fig. 2.9) suitable for production by industry which follows the same philosophy as that developed for the UA6 [19] experiment at CERN, and the STAR [25] experiment at RHIC.

Uniformity measurements

To establish the possibility of producing large numbers of SDDs, we have first studied the uniformity of the resistivity of neutron-doped n-type high-resistivity silicon [26] from a large batch of 4-inch wafers which we also used for our prototypes [1]. We implanted a matrix of 256 diodes on the wafers, which allowed measurement of the resistivity variations on the millimeter scale over most of the wafer surface both before and after several of the processing steps involved in SDD production, which could, in principle, degrade the wafer quality. We measured variations below 3 % (r.m.s.), which makes us confident that this material will be adequate for the production of large-area SDDs.

[1] ≈ 3 kΩ cm from Wacker.

Field simulations

In parallel, detailed simulations of the detectors have been performed in order to guide the design of several prototypes with different geometries and technological choices. In particular, we studied [27] the shape of the electric field at the junction edges in order to develop methods to reduce local maxima in the field strength, and the field modulation in the collection region.

Industrial prototypes

A set of prototypes has been produced for DSI by CANBERRA, Belgium, which includes a considerable number of different designs and technological options, to enable a detailed study of the possible design issues to be carried out thoroughly. The prototypes include detectors with both mono- and bi-directional drift, with dividers made from polysilicon or p^+ implant resistors, and with different geometries of the voltage dividers and the drift electrodes. All the prototypes were produced with two different oxide thicknesses.

The detectors with mono-directional drift are interesting because they allow complete separation between the drift region and the voltage divider. We also used different values for the integrated resistors which connect contiguous drift electrodes. It is important to test different resistor values, as the minimization of local power dissipation calls for large resistor values, while the linearity and stability of the divider call for low ones, and a compromise must be found depending on the specific technology chosen.

The tests of these prototypes are now in progress, and the first positive results have demonstrated the feasibility of industrial production of SDDs. A collection of the results can be found in Ref. 28.

Specific studies are being performed on the optimization of the charge collection, the draining of surface currents, the reliability and uniformity of the voltage divider, the generation of heat on the detector, and the optimization of the guard regions.

Photographs of one of the first prototypes are shown in Figs. 2.12 and 2.13. In Fig. 2.13 the anodes, the individual contacts to the last field strips which are used to study the optimization of the collection region, and test contacts to the voltage divider which allow the measurement of its linearity are visible.

The integrated voltage dividers show very good linearity: in Fig. 2.14 a typical measurement of the potential distribution of the drift electrodes of a two-directional drift detector is shown. The deviations from the linear fit are all below 0.4 V, and their mean value is 0.14 V for a bias voltage of 512 V.

In Fig. 2.15 the drift time plotted against the distance from the anode, for different drift fields, measured injecting localized charge with an infrared laser is shown. The deviations from the expected linear behaviour are mod-

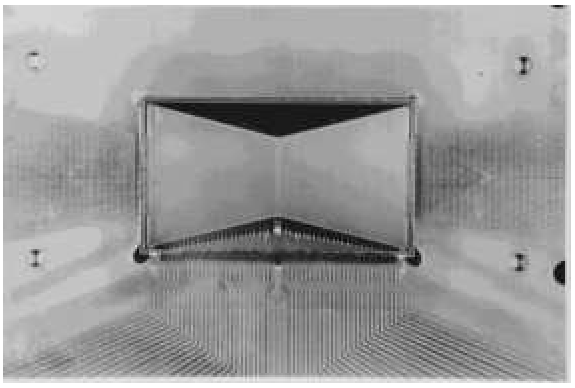

Figure 2.12: Photograph of one of the prototypes produced by CANBERRA for the DSI R&D project.

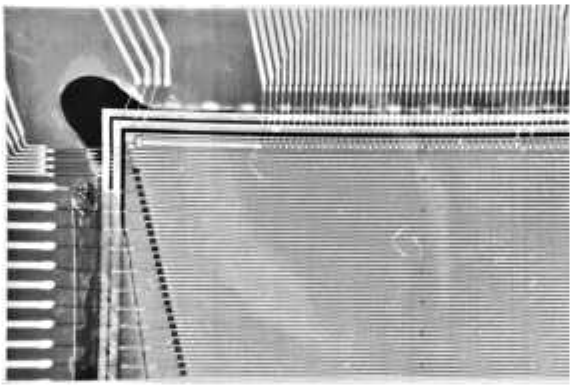

Figure 2.13: Details of the anode region of the CANBERRA prototype.

est up to a field of 670 V/cm, which was the limit of the available power supply. The change in the electron mobility at the different field values reflects the different power dissipation of the divider.

The typical total anode currents on our prototypes are between 30 and 100 nA. The mean value of the current for each individual anode is 0.9 nA at $V_{\rm bias} = 500$ V.

Finally, in Fig. 2.16 a typical measurement of the breakdown voltage between contiguous guard strips for detectors with thin oxide, connected to drift strips biased through the integrated divider is shown. In this case the guard lines have a pitch of 60 μm (gap of 35 μm).

On the basis of the results of these tests, we are currently defining the design of a large area detector, i.e. as large as possible within a 5-inch wafer, optimized for ALICE. The optimization process will involve, in particular, a minimization of the guard areas.

In a parallel effort, production of linear SDDs is being pursued at the Innovation Centre of Micronics Ltd., Prague (ICM). The first detectors on 3-inch wafers were successfully produced in February 1995 using masks and a process [2] developed at BNL for the Generic SDD project [16]. Results obtained with test structures, compared with samples produced at BNL, indicate a similar process quality. Tests of the drift detectors are scheduled for the first half of 1996.

Radiation damage tests of SDDs

Linear SDDs manufactured at ICM, Prague were exposed to fast neutrons at the isochronous cyclotron irradiation facility at the NPI in Řež. The neutron energy spectrum spanned up to 20 MeV with an average energy of 6 MeV. Samples of the SDD were placed at five different distances from the generation target. Neutron fluences spanned the interval from 1.25×10^{13} to 3.5×10^{13} n/cm^2. The measured variation of both the flatband voltage, which is proportional to the positive charge trapped in the surface oxide layer, and the leakage current, which is linked to the bulk generation processes on lattice defects are consistent with similar studies of Si-based detectors [29–31].

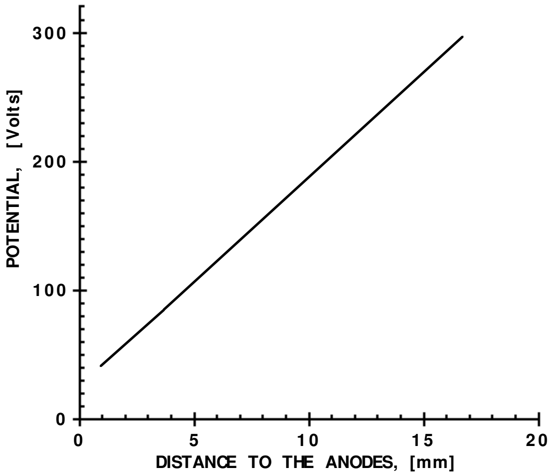

Figure 2.14: Measured potential distribution in the integrated voltage divider.

A study of changes in drift velocity, charge collection efficiency, and linearity is on-going, but the first results show that the detectors are still operational after 3.5×10^{13} n/cm^2. The neutron dose in ALICE at the position of the SDDs in 10 years of running is well below this limit (see Section 12.5.3).

2.3.5 Front-end electronics

Basic design

The signals from the detector are first amplified by a circuit directly bonded to the anode and then transferred to a second chip where they are sampled at a frequency of $\sim$40 MHz and stored in a 256-cell deep analogue memory, hence accommodating drift times up to 6.4 μs. Upon

[2] Courtesy of P. Rehak.

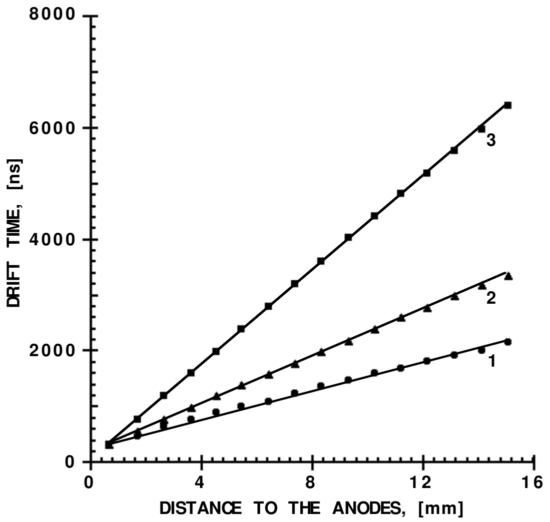

Figure 2.15: Drift time versus distance to the anode, for three different values of the drift field, from a DSI prototype.
Curve 1: $E = 600.1$ V/cm; $\mu_0 = 1368\,\mathrm{cm}^2 \cdot (\mathrm{V}\cdot\mathrm{s})^{-1}$.
Curve 2: $E = 349.8$ V/cm; $\mu_0 = 1372\,\mathrm{cm}^2 \cdot (\mathrm{V}\cdot\mathrm{s})^{-1}$.
Curve 3: $E = 173.6$ V/cm; $\mu_0 = 1388\,\mathrm{cm}^2 \cdot (\mathrm{V}\cdot\mathrm{s})^{-1}$.

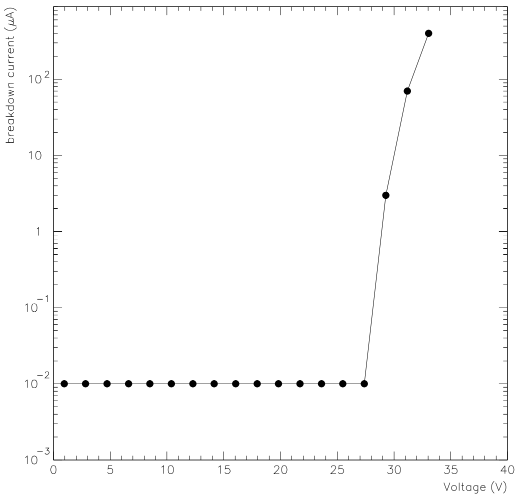

Figure 2.16: Measurement of the breakdown voltage between adjacent guard strips on one DSI prototype.

reception of a trigger the data are digitized, zero suppressed and transmitted to fibre optic links.

The storage unit can be installed either on the same board with the preamplifier directly bonded to it, or on the control cards which are mounted at the end of each ladder. In the latter case the cabling is more cumbersome, and a bipolar preamplifier is necessary in order to drive the cables efficiently with low power. In the first case, the assembly is simpler but additional cooling power is needed and the local hybrids are larger.

For both concepts prototypes have been produced or are planned in the near future, and we foresee having results which would allow a choice to be made within a year. In both cases the material budget remains within the design constraints.

Dynamic range

The resolution for the analog-to-digital conversion is dictated by the need for a dynamic range sufficient to accommodate the large energy deposition of slow particles whilst maintaining good resolution for low charge deposits required by the centroid finding algorithm. The maximum energy deposition that we need to measure corresponds to the the lowest momentum protons which we can reconstruct with reasonable efficiency, i.e. ~ 10 times the average signal released by a Minimum Ionizing Particle. We hope to achieve noise levels well below 500 e of ENC, and therefore we foresee having 8-bit resolution and 10-bit range.

Data volume

In the following, we will call *cluster* the 2-dimensional group of contiguous cells occupied by the signal from the passage of one particle. Given our drift time and sampling speed, the *maximum* size of a cluster is 5×6 memory cells. In a central Pb–Pb event the average number of clusters is ~ 90 per 1/2 detector in plane 3, and about 35 in plane 4.

Two levels of data compression are also being developed, as alternatives or additions to the suppression of zeros after digitization: a simple zero suppression to be performed at the analogue level in order to limit the size and power consumption of the memory, and a further level of data processing to be performed by a Fuzzy Processor (FP) either before or after digitization [32]. In this case, for each cluster the centroid position and the amplitude of the peak would be transferred (35 bits in total). Without any sparse readout, one event would represent 800 kbit per 1/2 detector; this value decreases to 23 kbit (plane 3) if zero suppression is applied, and to 3.2 kbit with signal processing. The two schemes are being evaluated independently and could both be implemented in the final readout. The size and power consumption of these additional circuits will obviously influence the choice of whether to assemble the analog memory on the detector module or at the end of the ladder.

The TPC readout has similar features to the SDD one, and some of the solutions and components, e.g. the analog memory unit, could be used for both detectors, and is being developed in close collaboration with the TPC group.

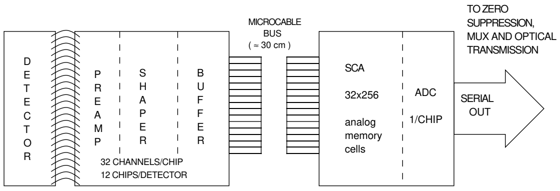

Figure 2.17: Basic scheme of the readout chain for the SDDs.

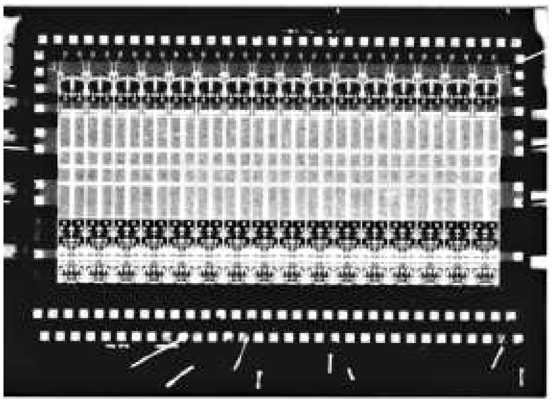

Figure 2.18: Photograph of the OLA chip.

Preamplifier

The basic readout concepts is shown in Fig. 2.17. A bipolar preamplifier-shaper-driver is located on the on-detector board and drives 100 μm pitch differential cables to a second stage. This second stage includes an analog memory in the form of a Switched Capacitor Array (SCA), and an ADC, located on a readout board at the end of the ladder.

The low-noise preamplifier-shaper has a peaking time of <40 ns and a Gaussian shaping. The latter is dictated by the need to optimize two-track resolution, and to improve the precision along the drift direction. The noise level is essentially determined by the capacitance of the connection elements, which is at the level of 0.5 pF.

With these constraints, we have developed a design in bipolar technology which has achieved an ENC of less than 300 electrons, and has a differential output capable of driving a short cable connection to the next stage. This noise level is more than adequate to reach our spatial resolution goal. For this design the power consumption can be contained below 1 mW per channel.

A 32-channel prototype of this design (OLA) has been produced in a small series of several hundred chips. The chip performs very well, with a noise level below 250 electrons, excellent uniformity, and cross-talk below 1 % [33]. A photograph of the chip is shown in Fig. 2.18. The prototype chip is currently being used for the readout of cylindrical silicon drift detectors by the WA98 and NA45 experiments.

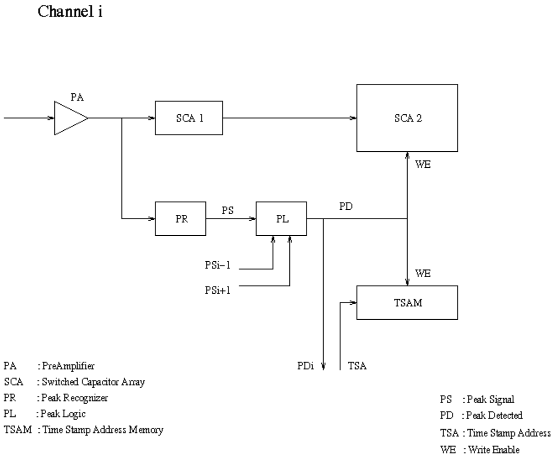

Figure 2.19: Scheme for an analog zero-suppression circuit.

Analog Memory

The short cables will carry the signals to an analog memory mounted on the detector end-plates. This could be an SCA; such an ASIC exists [34] and has been tested with an acquisition rate of 40 MHz. In this case we would require a storage depth of 256 cells. When the full event has been stored on the SCA, data are clocked to the ADC. One modest speed ADC would be sufficient for each SCA (4 MHz or 8 MHz for 32 or 64 channels, respectively, even if no zero suppression is applied at the analog level). Digital data are then transmitted to the readout electronics where zero suppression is performed before transmission on fibre optic links. In this configuration, the power dissipated locally will consist of approximately 100 mW per detector for the voltage divider and $\leq$1.5 mW per channel for the front-end.

If the analog memory was also mounted next to the detector, there would no longer be a need for a robust driver, and a CMOS preamp/shaper would be the chosen solution. This circuit could be integrated on the same chip together with the analog memory, if the interference between the two circuits was tolerable. We are currently developing a design based on the innovative concept of Ref. 35, and we plan to have a prototype ready in the spring of 1996.

Zero suppression

In the following, we briefly describe the on-going R&D programs on analog zero suppression and fuzzy online processing. An extensive description of these concepts can be found in Refs. 36 and 37. In addition, an innovative scheme for data coding has been proposed [38] and is currently under study.

In the zero-suppression circuit currently under study (Fig. 2.19), the signal from each preamplifier is fed to a short (8-cell deep) analog memory (SCA 1). The data are stored for the time necessary for the Peak Recognizer (PR) circuit to identify that a peak has occurred, via derivation and zero-crossing detection in coincidence with a low threshold, and for the Peak Logic (PL) to verify if a peak has occurred in one of the neighbouring channels. The latter is necessary in order to avoid losing the tails of the charge cloud. If a peak is recognized, the data are transferred to the SCA2 and a pointer to the time stamp associated with them is stored in the Time Stamp Address Memory (TSAM). Upon arrival of a trigger signal, the data are sent for processing to the FP described below.

If the FP is not implemented, the chain would consist of a pair of chips: the first one would be mounted next to the detector and would contain the preamplifier, the SCA1, the zero-suppression logic, the SCA2, and an analog multiplexer. The second one would be installed on the ladder-end board and contain the A/D conversion and the time-stamp memory. We estimate that the power consumption of the on-detector chip can be maintained below 5 mW per channel.

Fuzzy online processing

In this design each 1/6 of a detector is connected to four 32-input charge amplifiers (CAS) each driving a 32×256-cell Analog Memory Chip (AMC). An Analog Memory (AM) is made up of four AMCs and works like a 128×256-cell analog matrix.

The arrival of a trigger freezes the AM and activates the Cluster Finder (CF) circuit, which looks for two-dimensional peaks. The CF requires a total of 680 μs to scan the whole memory. For each cluster found, the FP calculates the position of the centroid in both coordinates, which are stored together with the digitized signal amplitude. Such an operation is executed for all the clusters. The 35×35 output memory stores the results of the last event until they are sent to the acquisition system through an Output Buffer.

Simulations are being carried out to study the performance of the system. We have two different possible hardware implementations for the FP: an analog and a digital one. At the moment we are working on both solutions [39,40].

2.4 The silicon strip layers

2.4.1 Introduction

The two outer layers of the ITS are crucial for the connection of tracks from the ITS to the TPC. They also provide $\mathrm{d}E/\mathrm{d}x$ information to assist particle identification for low-momentum particles. Both outer layers of the ITS consist of double-sided silicon strip detectors. Double-sided silicon strip detectors are being used in several high-energy physics experiments; they function reliably and perform well [41]. The detectors can be produced in large quantities in industry and they are available with capacitive read-out and integrated bias resistors.

Within ALICE, we have started some specific R&D [42] to optimize the price/performance ratio for our needs.

At 40 cm from the interaction point the track density, about 0.5 cm^{-2}, allows the use of small stereo angle pairs and reasonable strip lengths with an acceptable occupancy and a relatively small ambiguity. A GEANT simulation [43] of the response of a silicon strip layer was used to investigate the efficiency of the charge matching between the two pulses produced by the same energy deposit in the detector. The total fraction of correctly identified tracks was found to be about 96 %.

2.4.2 The silicon strip detector design

In order to obtain maximum resolution in the bending plane the strips in the ALICE ITS are oriented parallel to the beam line. In this way the resolution in the $r\varphi$ direction is essentially fixed by the pitch of the strips, 95 μm giving a r.m.s. resolution of 27 μm, even without centroid finding. With a stereo angle of 35 mrad between the read-out strips on either side of a single silicon wafer, a resolution of 860 μm along the beam direction is obtained. With a strip length of 41 mm in the outer layer and 40 mm in the adjacent layer, the occupancy of a single strip remains below 2 %. The probability of having ambiguous hits —a property of all projective, one-dimensional read-out schemes— is less than 26 %. As mentioned above, most of these ambiguities can be resolved by matching the pulse heights from both sides of the detector.

Both layers are constructed from modules (see Fig. 2.20), each approximately 75 mm by 42 mm. Each module contains a double-sided detector with front-end electronics, decoupling capacitors, bias circuits and local cabling. Each detector is connected to the electronics in the end-caps by low-mass micro-cables (see Section 2.4.4). The modules are glued to a carbon fibre space-frame (Fig. 2.21), forming the basic ITS ladder structure. Each ladder in the outer layer holds 24 modules and each ladder in the fifth layer holds 22 modules. The outer layer consists of 39 ladders, while the fifth layer consists of 35 ladders.

The front-end electronics is located on top of a heat bridge. The heat bridge is made of carbon fibre with a large thermal conductivity. The heat bridge itself is glued onto an insulating foil, which is in turn glued on top of the detector. The heat produced by the twelve front-end chips per detector is conducted through the 0.3 mm by 12 mm heat bridge to two 2 mm thick aluminium cooling arteries located on the side of the detector. Based on currently available designs the front-end electronics of the approximately 2.6 million channels dissipates less than 3 kW inside the barrel and less then 1 kW at the end-cap. The

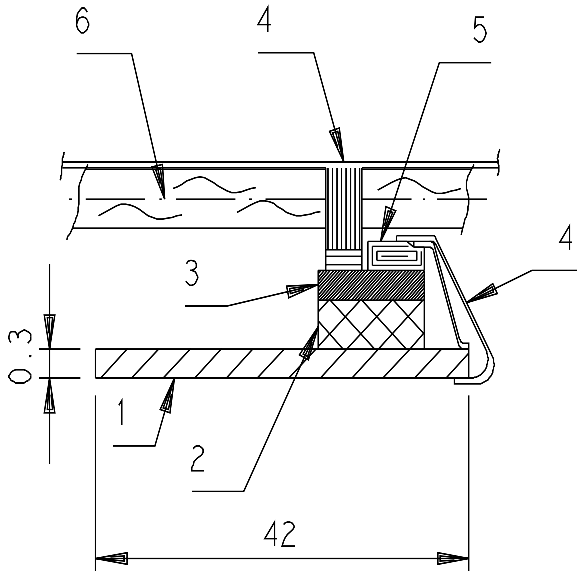

Figure 2.20: Detector module assembly. The electronics (5) for the double-sided detector (1) is located on one side and connected using micro-cables (4). The heat produced in the electronics is transferred to the cooling pipe (6) through a heat-conducting carbon fibre (3). A spacer (2) between the heat-conducting fibre and the detector prevents capacitive coupling with the readout strips on the detector.

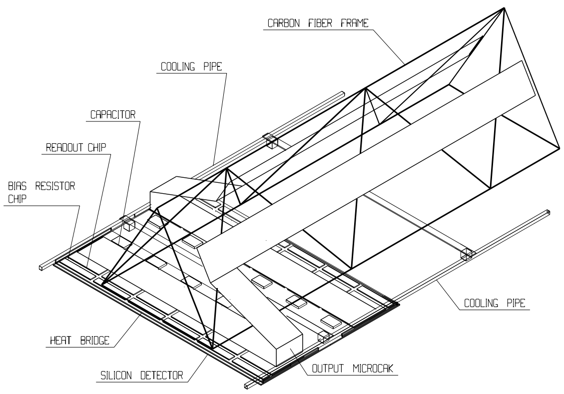

Figure 2.21: Three-dimensional view of the silicon strip assembly. The front-end electronics is connected to the control logic in the end-caps of the ITS through micro-cables. The module shown here is glued on the carbon fibre ladder support.

cooling system, using water at moderate pressure, is designed to limit the temperature to 10 degrees above ambient temperature.

The front-end electronics is connected to the detector using micro-cables. One cable connects the channels on top of the detector to the even channels of the electronics, while another micro-cable connects the channels from the other side of the detector to the odd channels of the electronics. In this way the spacing of the input pads on the electronics (40 μm pitch) is optimally connected to the strips on the detector, which have a pitch of 95 μm. In addition, the decoupling capacitors needed around the front-end electronics are connected using micro-cables. In this way extra material caused by printed circuit boards is avoided. Since all the components on a detector module are glued together and cannot be removed once installed, the electronics and detectors will be pre-tested.

The average thickness of a detector module, including the supporting carbon fibre frame, is less than 0.59 % of X_0. The main contribution to the material is the silicon detector itself (300 μm). The other components, i.e. the heat bridge, the front-end electronics, and the decoupling capacitors, cover only a fraction of the area of the detector. Locally the thickness increases up to 1 % of X_0.

2.4.3 Readout electronics

Existing electronics comes close to the requirements for the silicon strip detectors in the ITS [44]. The ALICE environment requires optimization of the electronic noise because of the particle identification through $\mathrm{d}E/\mathrm{d}x$ measurements, and optimization for power consumption because of the high channel density in large areas.

The current design of the detector modules is based on 128-channel integrated chips. The front-end electronics for the silicon strip detectors must contain a pre-amplifier and shaper circuit, with a shaping time around 700 ns. The noise must remain below 400 electron equivalent noise. The dynamic range needed for $\mathrm{d}E/\mathrm{d}x$ is the same for the SDD and the strips, i.e. 10-bit range, 8-bit resolution. For efficient mounting of the electronics, the pitch for the inputs of the pre-amplifiers should be less than 40 μm, to allow interleaved readout of both sides of a detector. Power consumption must be below 1 mW per channel to allow efficient cooling with a minimum of material.

A new design, based on an existing 32-channel chip of DELPHI (see Fig. 2.22) [45], is being developed for ALICE at CRN-LEPSI. This 128-channel preamplifier shaper circuit —with sample and hold, multiplicity measuring capability, less than 1 mW/channel dissipation— will meet the above-mentioned mechanical and electrical specifications.

A Level 0 trigger signal will be available about 1.2 μ s after the interaction, and in any case no later than 1.5 μs (see Chapter 9). Depending on the final timing of L0, an analog delay of the order of 1 μs may be needed for

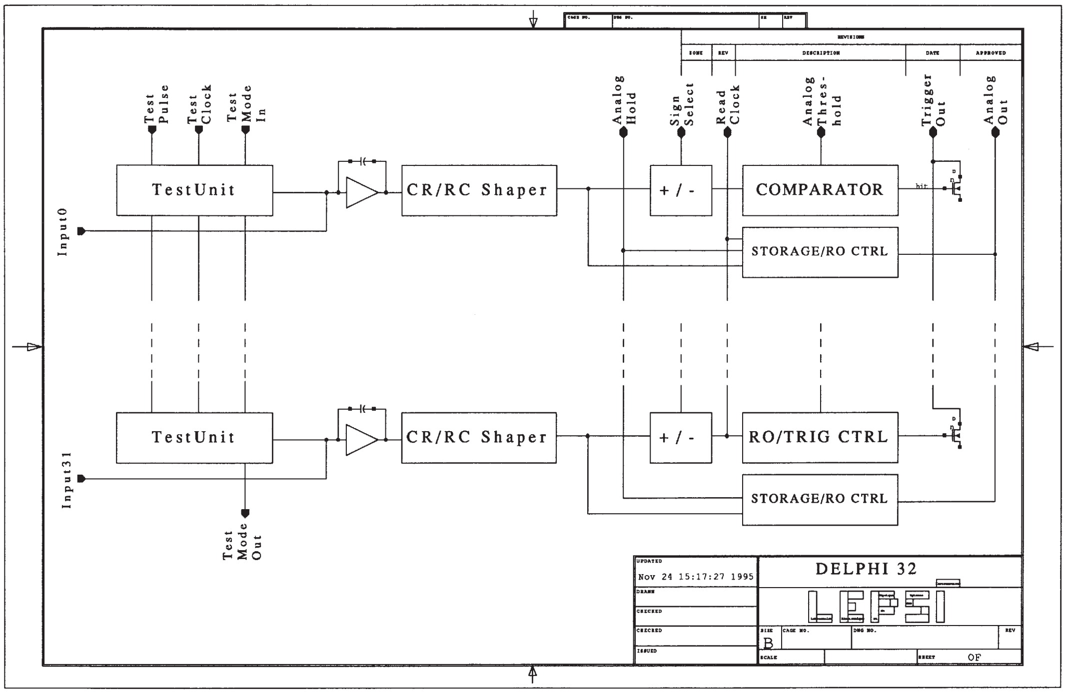

Figure 2.22: 32-channel chip, 100 mV/MIP gain, $100\,e + 10\,e/\mathrm{pF} \cdot C_{\mathrm{Input}}$ noise on a 40 μm pitch.

each channel. The incoming trigger signal will operate a sample and hold circuit, freezing the analog signal until the readout starts.

The stored analog signal is multiplexed to a single output during readout. The outputs from all the chips on a detector are further multiplexed to one cable connected to readout electronics in the end-caps of the ITS. The total readout time of the system must be less than 1 ms. The use of micro-cables for these connections allows a separate connection from each detector to the end-cap electronics without introducing too much material.

The task of the electronics in the end-caps is to provide control signals for the front-end electronics on the detectors, to distribute trigger signals, to collect the analog signals on the multiplexed outputs, to convert from analog to digital, and to encode and transmit zero-suppressed data to the ALICE data acquisition system. The silicon strips will use the standard for transmission to be adopted by ALICE.

Each detector will transmit 1536 analog channels sequentially at a 10 MHz rate to the end-caps, where the conversion is made at the same rate by a 10-bit flash ADC. An attractive option is to use a low power pipeline ADC [46]. Thus, after 154 μs the data is available in digital form. Each ladder in the outer layer of the ITS contains 36 864 channels. Dividing each ladder into two halves and reading out all half-ladders in parallel results in the transport of the data of 18 432 channels to a single end-cap module for the outer layer of silicon strip detectors. In the end-cap modules zero suppression is implemented and a 15-bit address is attached to the resulting data. With 2 % occupancy and 10-bit ADC values the total data from a single end-cap module amounts to an average of 9216 bits.

Since digitization, zero suppression and data transmission can be implemented in a pipeline structure, the data can be available in the data acquisition system within 500 μs after the readout is started. In the data acquisition crate the data from all end-cap modules are combined into a single event structure. With 10-bit ADC values, 22-bit addresses and 2 % occupancy, the combined zero-suppressed data of the two outer silicon strip layers amounts to approximately 1.7 Mbit.

2.4.4 Micro-cable technology

A micro-cable [47] consists of aluminized, 12 μm thick, polyamide layers. The cables can be either single or multilayer. Aluminium traces are formed using a lithographic process, producing multiple conductors on a single layer. Currently the accuracy of the production process allows pitches for the pattern of conductors down to 100 μm.

Aluminium was chosen as a conducting material because of its low density together with a sufficiently

Table 2.4: Mechanical and electrical parameters of existing micro-cables.

Micro-cable parameters	
Kapton base thickness	10–40 μm
Conductive layer thickness	10–40 μm
Protective coating thickness	10–40 μm
Number of layers	1–10
Conductor material	Al
Pad material	Al, solder coating
Minimum track size	50 μm
Minimum track spacing	50 μm
Electric strength of polymer	50 kV/mm
Resistance per unit length	
of polymer	10^{14} Ω/cm
of Al cable	
(100 μm × 12 μm)	0.23 Ω/cm

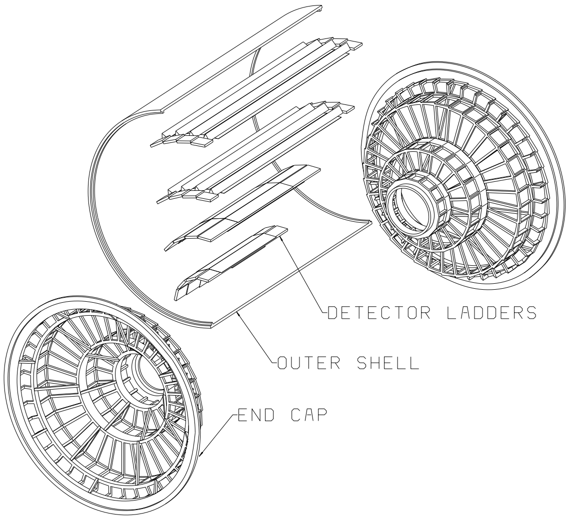

Figure 2.23: Blow-up of the mechanical structure showing separately, the end caps, the outer shell and detector ladders.

high conductivity. The thickness of the aluminium traces is 12 μm, resulting in a resistance per unit length of 0.23 Ω/cm for 100 μm wide traces. Furthermore, aluminium allows reliable ultrasonic bonding of the micro-cables directly onto the electronic components or the silicon detectors, thus avoiding the use of connectors or printed circuit boards. It is even possible to connect different layers of two multilayer cables to each other through ultrasonic bonding. Connections to individual layers of a multilayer cable are made through windows in the polymer base. If a solder coating is deposited on the cable pads soldering is also possible.

The polymer base provides a flexible but strong cable. The cable allows repeated bending up to 90° with a minimum bending radius of 1 mm. However, due to the small size of the individual aluminium conductors and hence the small mechanical strength of a single connection it is nescessary to glue the polymer base of the cables to the device to which it is connected. The breaking force of a bonded junction of a single lead is approximately 5–25 g.

The mechanical and electrical properties of micro-cables (see Table 2.4) allow the construction of all signal cables with a minimum of material. Cabling and decoupling of power lines is investigated separately for each layer of the ITS. Additional investigation of the micro-cable technology is being directed towards reducing the minimum pitch between conductors. In addition, large-scale industrial production and long-term reliability in radiation environments will be investigated.

2.5 Engineering design

The ITS system will have a structure mechanically independent from the other detectors which can be preassembled and aligned accurately before final installation in the experiment [48]. The aim is to produce a lightweight structure with high mechanical rigidity and a very low thermal expansion coefficient which will provide a high degree of thermostabilization and a minimum of matter along the track path. In addition, the mechanical design of the ITS has to take into account the integration of different cooling schemes for the various detectors.

The basic principle of the ITS design is to concentrate the supporting and cooling structures into the end-caps rather than the central region. The detector planes will be mounted on individual space-frame ladders made of carbon fibre. The supporting structure, shown in Fig. 2.23, consists of two end-cap cones where the ladders are fixed, and the outer shell which provides the necessary rigidity to the barrel.

2.5.1 Mechanical design

Support ladders

The ITS detector planes will be mounted on space-frame ladders which are slightly different for the pixel, drift, and strip detectors. The individual designs are governed by the length and relative distance between adjacent planes and by the detector cooling systems.

We have concentrated our efforts on design and tests for the outermost layer, which is the most demanding in terms of rigidity because of its length. A photo of a prototype, full length (1 m) space-frame support made by MESON (St. Petersburg) is shown in Figs. 2.24 and 2.25.

This particular carbon fibre space-frame design provides the best ratio of mass to rigidity because the high elastic modulus of the fibre, which is maximal along the fibre axis, is being used to the full extent.

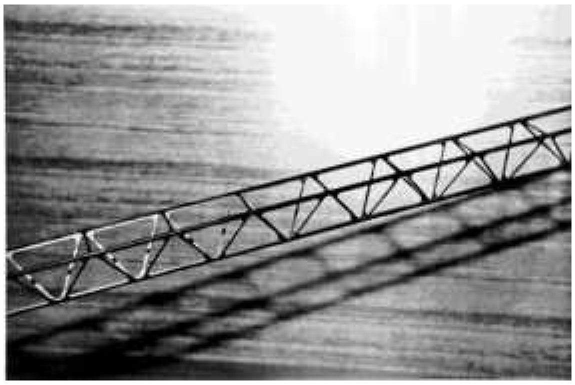

Figure 2.24: Lightweight (22 g) monolythic space-frame ladder for detector support, made from carbon fibre.

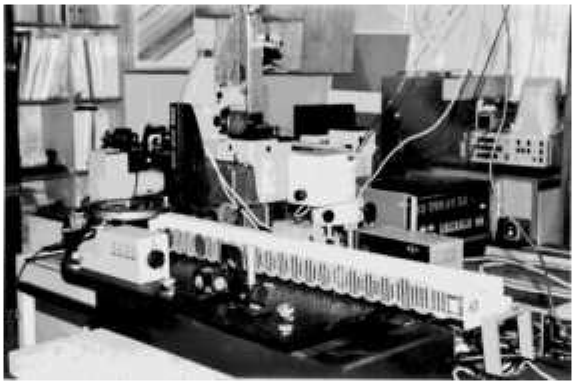

Figure 2.25: Carbon fibre space-frame ladder including silicon wafers, mounted on the test bench.

Finite element analysis (FEA) predicts a gravitational sag of this structure of 20 μm under its own weight (22 g). When adding the weight corresponding to the silicon detectors (about 80 g distributed over 1 m) the sagging increases to 96 μm, as shown in Fig. 2.26.

The FEA calculations have been confirmed by measuring the gravitational deformations with and without the silicon load and for different ladder orientations [49].

By using high modulus carbon fibre it is possible to further decrease the weight and increase rigidity.

A space-frame option suitable for industrial production was designed and produced by STESALIT with a performance very close to the support described above.

ITS support structure

The internal tracker mechanics is divided into two independent parts: an inner part very close to the beam pipe holding the pixel system, and an outer part with the drift and strip layers. The reason for decoupling the two subsystems is to permit the removal of the pixel cylinder

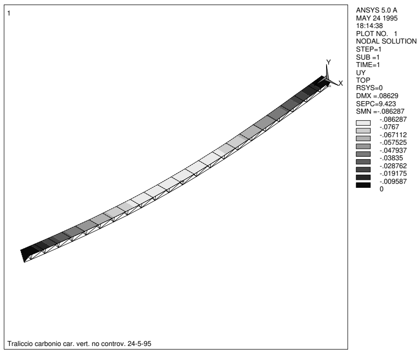

Figure 2.26: FEA calculations of the prototype support ladder with silicon wafers distributed over the full length. The maximum sagging at the centre is 0.086 mm.

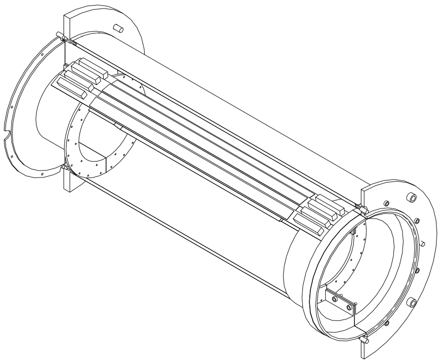

Figure 2.27: Pixel barrel.

without dismounting the complete ITS, for example during installation, or during bake-out of the vacuum chamber.

The pixel detector barrel (see Fig. 2.27) is divided into two clam shell pieces, each consisting of two end-pieces and five support bars. It will be held by the outer barrel by precision fittings (tolerance ≤ 0.1 mm). The pixel cylinder will be firmly fixed only on one side, the other side (towards the muon absorber) will be supported by precision guides, such that the transmission of any mechanical stress from the ITS structure to the pixel support bars can be avoided.

The mechanics for the outer part of the ITS (drift and strip) is shown in Figs. 2.28 and 2.29. The outer shell is made of two thin layers of carbon fibre composite material with a thickness of 0.6 mm each, separated by a standard honeycomb sheet, 10 mm thick. The monolithic end-cap structure, also of carbon fibre composite, consists of four concentric rings (interconnected by radial rods) on which

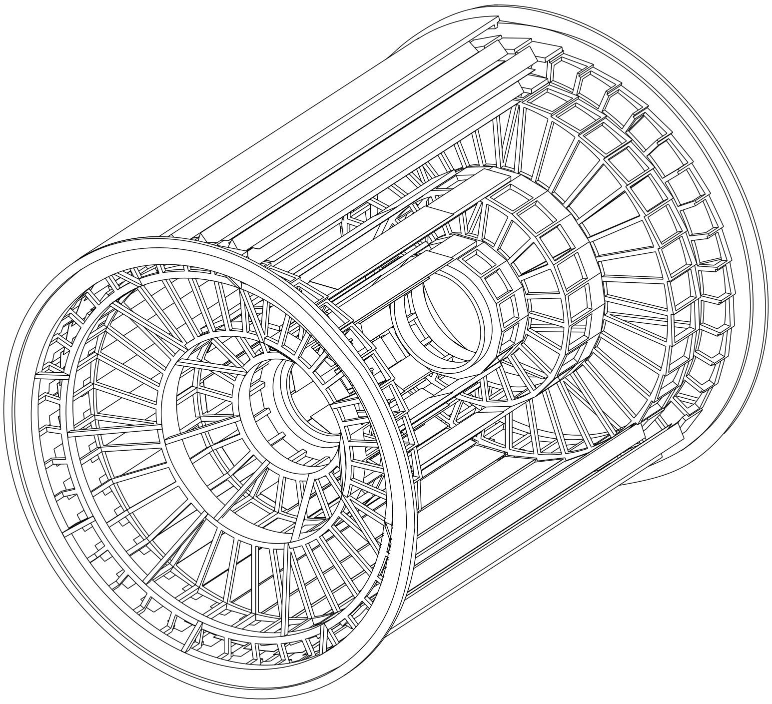

Figure 2.28: The mechanical support of the ITS.

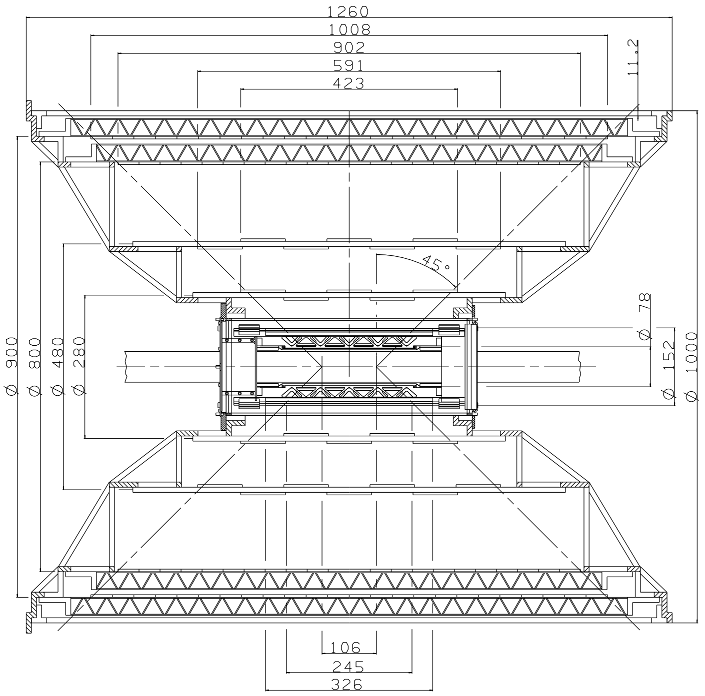

Figure 2.29: Side view of the mechanical structure showing the pixel, the drift, and the strip ladders.

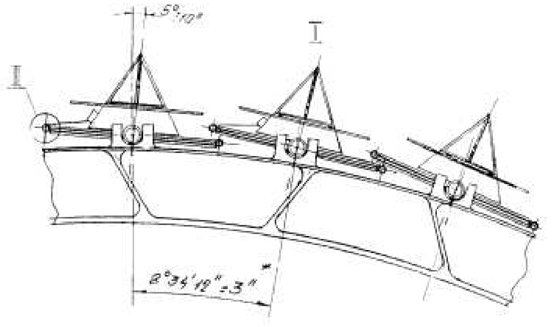

Figure 2.30: Spaceframe mounting with thermoplastic glue.

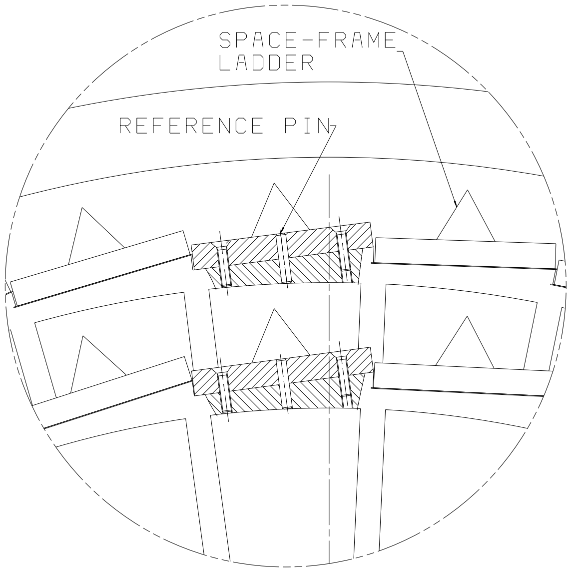

Figure 2.31: Details of the mounting of the spaceframe ladder on the end-cap structure.

the ladders of the four outer detector planes are fixed.

Results from the FEA calculations of the end-cap show a negligible sagging (about 12 μm) of the central part of the ITS under the load of its own weight. Work is in progress to minimize the number of rods necessary for strengthening the structure while maintaining low deformations and stresses in the inner part.

We have studied a number of ways to fix the ends of the detector ladders to the supporting ring. One possibility is thermoplastic gluing, which was used in the ladder prototype tests (see Fig. 2.30). It has the advantage of simplifing the mounting of the ladder on an assembly station with high accuracy, without requiring precision machining.

Another approach is shown in Fig. 2.31. It relies on accurate machining of a concentric ring to obtain flat surfaces, where the ladders can be positioned with reference pins (accuracy about 10 μm). Both methods allow easy repeated dismountings of various elements.

The transparency of the ITS end-caps will allow the use of optical control during and after the ITS assembly.

2.5.2 Alignment and assembly

The final alignment of all the ITS detectors will be performed *in situ* with the ALICE tracking software using residuals of measured and extrapolated tracks. This imposes minimum requirements for the geometrical precision during the initial assembly procedure.

The ladders have to be mounted to within 100 μm with respect to the ITS coordinates. The ITS system has to be positioned, with respect to the TPC, with a precision of 200 μm in the $r\varphi$ direction, and 500 μm in z. The ladders and the detectors will be put together on optically-controlled assembly stations. This allows the Si detectors to be fixed onto the ladders with a relative precision of 10 μm. The individual components of the ITS barrel are then mounted on a specific assembly station as shown in Fig. 2.32.

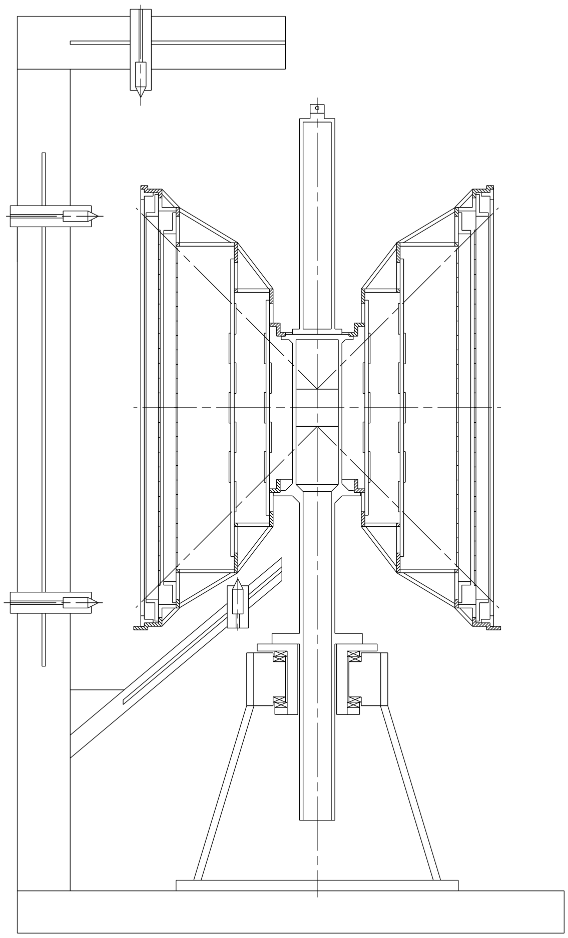

Figure 2.32: Ladder assembly on the work station.

Both end-caps will be mounted onto the axis of a rotating table which is equipped with a micrometric angular control. The SDDs and the Si strip ladders will then be

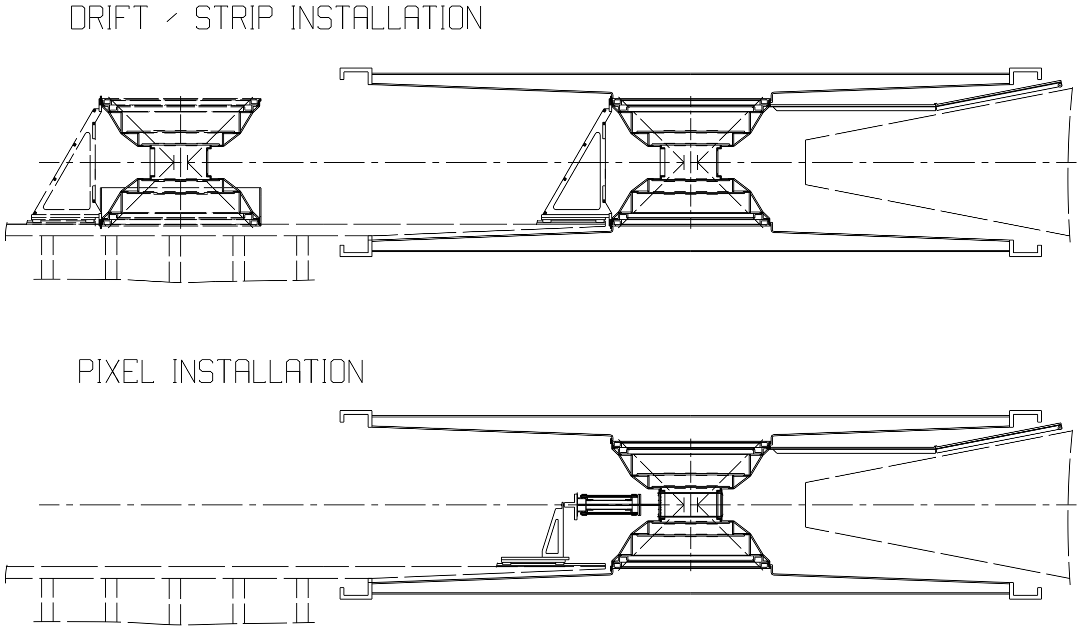

Figure 2.33: Installation of the ITS inside the TPC.

fixed onto the end-caps, starting with the innermost layers. The ITS will be assembled in a vertical position to avoid gravitational sag during assembly.

During the assembly procedure, optical microscopes equipped with CCD cameras will allow to control the position of the ladders.

Finally, the outer shell will be inserted to connect the end flanges mechanically. Then the ITS will be turned and transferred to a dedicated tool for final installation inside the TPC.

During running, the relative alignment and stability will be monitored by optical alignment systems based on semiconductive lasers and position sensitive photo diodes. The development of these alignment systems is still on-going and will benefit from the considerable R&D efforts made by ATLAS and CMS in this field.

2.5.3 Installation of the ITS

The ITS is connected to the TPC inner vessel by two conical end-pieces (see Chapter 3). The contact surface between the ITS and the TPC must be precision machined in order to reach a relative position accuracy ≤ 0.2 mm. Owing to the high rigidity of the ITS system the maximum expected sagging of the overall structure (inner vessel and ITS) under gravitational load is $\leq 20\,\mu$m.

The installation procedure for the ITS is shown in Fig. 2.33. A 3 m long part of the beam pipe in front of the TPC will be removed in order to install the ITS with the help of a temporary support. On the muon absorber side, cables and services will be placed in flexible cable ducts, which would slide along the absorber and allow the cables to reach the end flange of the TPC where the connections will be accessible.

The beam pipe will then be brought back into place. The inner bore of the ITS (with pixels removed) is large enough to install the temporary heating jacket needed for the bakeout of the beam pipe. After removal of the heating jacket, the two halves of the pixel barrel can be mounted around the beam pipe and moved into position using a special installation tool.

To allow this operation the pixel detector can only be cabled from the side opposite the muon absorber.

2.5.4 Cooling system

The ITS system has to drain 3–5 kW of heat while introducing only a minimum amount of material. Different cooling schemes were considered and their efficiency compared for the ITS applications [22,50]. In the present Technical Proposal we use a combination of a rather conservative water cooling system for the two large area outer layers and a new advanced evaporative thermostabilized one for the Si drift layers which requires temperature stability within 0.1°C.

Effective cooling of pixel detectors is at present being studied, in order to minimize the material in the acceptance.

Water cooling and test results

For the two pixel and the two strip layers water cooling is sufficient, as for these detectors temperature gradients of up to 5°C can be accepted.

As described in Section 2.2.4 the pixel ladders are mounted directly onto aluminium tubes (∅ 1 mm), which transport the cooling water. One single tube will cool 4 ladders on the outer layer (corresponding to 7.9 W of produced heat) and 3 ladders (5.9 W) on the inner layer. Allowing for a ΔT of 2°C in the water, this can be achieved by a water flow of 1 cm^3/s, corresponding to a pressure drop $\Delta P = 0.5$ bar. The water in the cooling pipes will add on average 0.04 % of a radiation length for each of the layers. Locally the maximum amount will be up to 0.27 % of a radiation length.

The Si strip detectors produce less than 1 mW per channel of heat, mainly in the front-end electronics. This heat will be transported by a thermoconductive carbon fibre (product under development by AMOCO Performance Product, Inc.) from the electronics to the cooling pipes. The feasibility of the thermoconductive heat drain from the front-end electronics was demonstrated with an ordinary carbon fibre heat bridge and compared with a brass heat bridge in Ref. 22. The test was performed with 2 W of heat per wafer. Temperature gradients between the coolant and the electronics along the Si wafer surface of a few degrees were observed, agreeing well with the calculations.

Each of the 74 Si strip ladders will be equipped with two cooling pipes of 1.8 mm diameter. This diameter was chosen to minimize the total amount of water in the sensitive area of the detector, but still preserve a laminar flow regime. Estimations [22] have shown that the heat produced in the Si strip detector (up to 3 kW) can be removed by the 148 cooling arteries with a flow of 2.1 cm^3/s per pipe and a ΔT of 2°C in the water. Such a flow would require a ΔP of 0.2 bar. The cooling pipes would add locally up to 0.49 % of radiation length to the material budget, while the average amount is 0.02 % per layer.

Evaporative cooling

The cooling of the SDDs is more ambitious than for the other ITS detectors as the temperature has to be controlled to within $\sim$0.1°C. Inside the acceptance region the SDDs will produce 0.3 to 1.9 kW of heat, depending on how much of the front-end electronics is located on the end-caps.

We intend to apply here a cooling scheme which uses the latent evaporative heat of a cooling fluid. Such a two-phase cooling (liquid and gaseous) offers inherent temperature stability, as the absorbed heat will evaporate part of the liquid. For the SDDs cooling we envisage using ozone-safe freons such as C_5F_{12} or C_2F_5I in a closed circuit. These liquids offer a boiling temperature of 25–30°C, close to the ambient temperature. The actual working temperature can be adjusted by operating the system slightly below atmospheric pressure (0.7–0.9 bar for ambient temperature). To achieve a temperature stability of 0.1°C the pressure has to be stabilized to within 10 mbar.

The SDDs will be assembled on 36 ladders, each supported by two carbon fibre tubes which are covered on the inside with an aluminium coating, and carry the cooling fluid. The heat producing front-end electronics will be mounted directly on these tubes.

One channel has to remove 5–30 W of heat by evaporating 0.05 to 0.3 g/s of the cooling freon. The amount of freon located in the acceptance region of the detector is 300 g (0.15 % of a radiation length) for both layers.

A first prototype test was performed with a 1.1 m long cooling artery (∅ 3 mm) fabricated by MESON [22]. A total temperature gradient of 0.8°C was observed when removing 25 W of heat, created by 20 heaters along the artery.

We are now studying ways of increasing the efficiency of the evaporative cooling system by incorporating a phase separation inside the cooling arteries.

2.5.5 Total amount of matter in the ITS

The various contributions to the ITS material budget are listed in Table 2.5. The total amount as seen by a particle at zero rapidity corresponds, on average, to 4.4 % of X_0. Of this, 1.6 % comes from the Si detectors, 0.9 % from the beam pipe and outer shell, and the remaining 1.9 % from the front-end electronics, supports, cables and cooling system.

Table 2.5: Summary of ITS materials.

	X/X_0 (%)
Vacuum pipe (0.6 mm Be)	0.17
Detector	0.35
Electronics	0.28
Bus	0.18
Mechanics, cooling	0.36
Pixel 1+2	1.17
Thermal shield (Mylar, gas)	0.05
Detector	0.73
Electronics	0.07
Cables	0.06
Mechanics, cooling	0.27
Drift 3+4	1.13
Detector	0.64
Electronics	0.10
Cables	0.06
Mechanics,cooling	0.35
Strip 5+6	1.15
Outer shell (carbon fiber, nomex)	0.74
Total	4.41

3 The Time Projection Chamber

3.1 Introduction

The Time Projection Chamber (TPC), the main tracking detector of ALICE, is central to the design of our experiment. With its large acceptance, it will enable us to analyse individual events and perform charged hadron exclusive analysis in the domain $|\eta| < 0.9$. The task of large acceptance tracking is similar to that encountered in e^+e^- colliders (e.g. for W, Z or jet analysis) where TPC tracking was frequently adopted. However, the extreme multiplicities of ion collisions at the LHC set qualitatively and quantitatively new demands as the charged hadron multiplicity in $|\eta| < 0.9$ is of the order of 10^4. Also, two-track separation and momentum resolution have to be kept close to the limits imposed by multiple scattering, in order to measure the two-particle correlation function down to about 5–10 MeV/c momentum difference (see Section 11.5).

These requirements call for continuous tracking with about an order of magnitude higher spatial granularity than, for example, in the ALEPH TPC [1]. Power consumption and electronics costs have to be diminished in parallel (the price per electronic channel for ALICE is expected to decrease by a factor of about 10 compared to ALEPH channel cost).

The NA49 TPCs have satisfied the above demands, for the SPS lead run of 1995, where the track density at high rapidity is comparable to the mid-rapidity densities expected for ALICE. Thus, in principle, the NA49 TPC technique would constitute a conservative, fall-back solution. Its adaptation into a mid-rapidity collider TPC is already under way with the STAR TPC, now under construction for installation at RHIC in 1998. However, a number of ALICE-specific requirements, (e.g. electronics packing density, the need for 1024 slices in time direction instead of 512, the need for zero suppression in the front-end, etc.) make new designs indispensible.

The TPC is a comparatively slow detector, with about 100 μs drift time over the full length of 2.5 m. However, it can cope with the minimum-bias collision rate in Pb–Pb of about 10 kHz, expected for the design luminosity $\mathcal{L} = 10^{27}$ cm^{-2}s^{-1} at the LHC. The large granularity of the TPC (about 5×10^8 pixels) leads to large event sizes and a data volume in the front-end which is far beyond the limit of present data-handling techniques. Zero suppression at the front-end and higher level triggers will reduce the data rate to the order of 1 GByte/s for Pb–Pb interactions, with about 50 events/s written to permanent storage. A high-speed online track reconstruction processor is under development [2] which could increase the event writing speed by up to a factor 5.

3.2 Design considerations

The task of the TPC is track finding, momentum measurement and particle identification by dE/dx. The mean momentum of the particles tracked in the TPC is around 500 MeV/c. The most important design goal for all the components of the ALICE TPC is to achieve a good two-track resolution in order to cope with the high multiplicity at the LHC.

The cylindrical shape of the TPC has been chosen since it is appropriate for a collider experiment. The electrical drift field is constant and points from the central HV plane of the cylinder orthogonally towards the two readout planes at the end-caps. These planes are on ground potential. Field cages of cylindrical shape and of considerable size have been successfully built in the past for TPCs such as ALEPH [1] and DELPHI [3], showing that the technique is well proven. However, the cylindrical geometry has a drawback in the position and two-track resolution in the drift direction if the polar angle (dip angle) is larger than 45°.

The inner radius of the sensitive volume, $R_{\rm i} = 0.88$ m, has been chosen such that the expected particle density on the surface of the cylinder is of the order of 0.1–0.2 particles/cm^2.

From experiments like NA35/NA49 [4] it is known that these particle densities can be safely handled even for dE/dx measurements. The outer radius, $R_{\rm o} = 2.5$ m, has been chosen to achieve a track length such that the dE/dx resolution is of the order of 6–7 %.

The total active length of the TPC, $L_{\rm tot} = 5$ m, is determined by the acceptance of ALICE ($|\Delta\eta| = 0.9$ around mid-rapidity).

A charged particle passing through the gas volume creates electrons by ionization. The electrons drift in the electrical field towards the readout chambers, where they are amplified in the field of the sense wires by a factor of several 10^3. This signal is coupled to the pads which are on ground potential and at a few millimetres distance behind the sense wire-plane. The width of the electrical signal, determinant for the two-track resolution in the azimuthal and drift direction, depends mainly on:

- electronics noise,
- diffusion constant,
- distance between wire-plane and pad-plane,
- drift velocity,
- shaping time.

The crucial parameters that determine the two-track resolution are the diffusion constant, the total drift length, and the drift velocity of the electrons in the gas. All the other parameters are of a technical or geometrical nature.

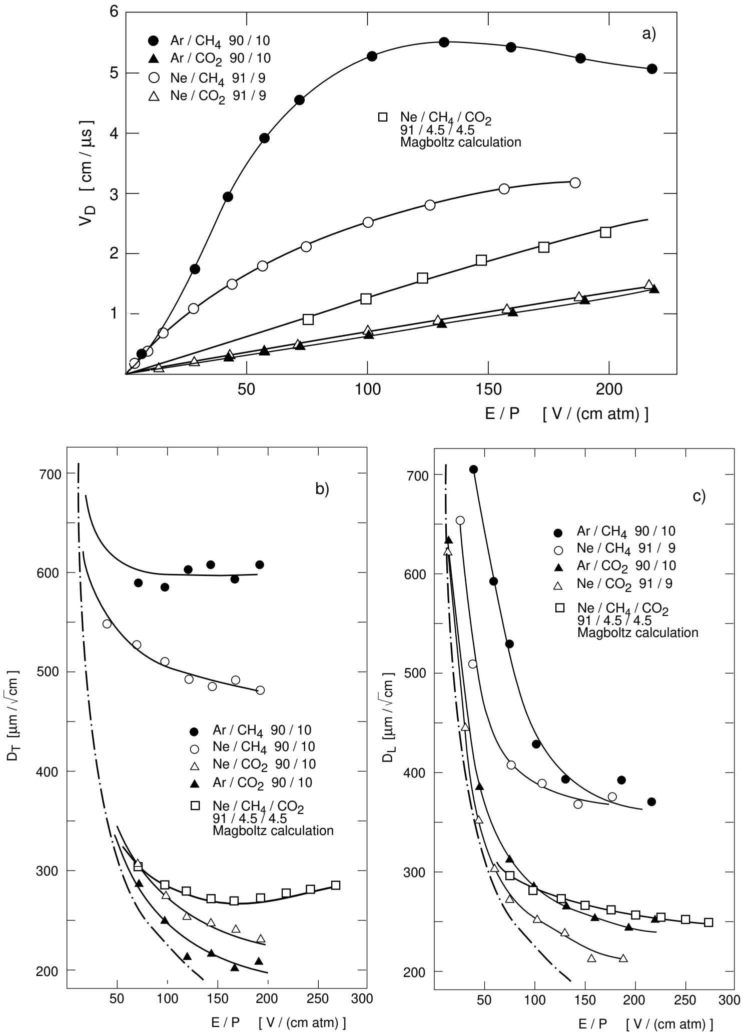

Figure 3.1: Drift velocity and diffusion constant for a number of gas mixtures. The $Ne/CH_4/CO_2$ curves rely on calculations (MAGBOLTZ [6]) which reproduce the available measurements [5]. The dash-dotted line in the lower figures indicates the thermal limit.

Extensive investigations of gas mixtures have been performed within the R&D work of RD32 [2,5]. As a result, mixtures of cool gases which have a small diffusion constant are proposed. Figure 3.1 shows the dependence of the drift velocity and of the diffusion constant on the electrical drift field, for Ar/CO_2 (90/10), Ne/CO_2 (91/9) and other similar mixtures, for comparison.

Some operating experience has now been gained by operating the NA49 TPCs with these gases. For reasons of stability of the electrical drift field, electron absorption, and to enhance the drift velocity, mixtures with an addition of CH_4 will be considered, like (Ar or Ne)/CO_2/CH_4 (91/4.5/4.5), shown in the same figure [6,7].

In order to operate the TPC at interaction rates of up to 10 kHz, the drift time should not exceed 100 μs. Consequently, the drift field will be set at about 400 V/cm, resulting in a drift velocity of 2.5 cm/μs and diffusion constants of about 250 μm/$\sqrt{\text{cm}}$. These values are proposed as a guideline for the current planning.

The preferred gas would be a neon mixture since it has a radiation length higher than argon and, consequently, effects of multiple scattering and space charge are considerably smaller in neon than in argon (see Section 3.11).

The expected contribution to the signal width (cluster width) due to diffusion is about $\sigma \lesssim 3$ mm given the average drift length of 125 cm. Other components which, in addition, influence the signal width are the granularity of the readout chambers in the $r\varphi$ direction (Pad Response Function, PRF) and the sampling frequency in the drift (z) direction. Both are chosen such that their individual contribution to the signal width is also in the order of 3 mm. For the readout chambers these considerations lead to a

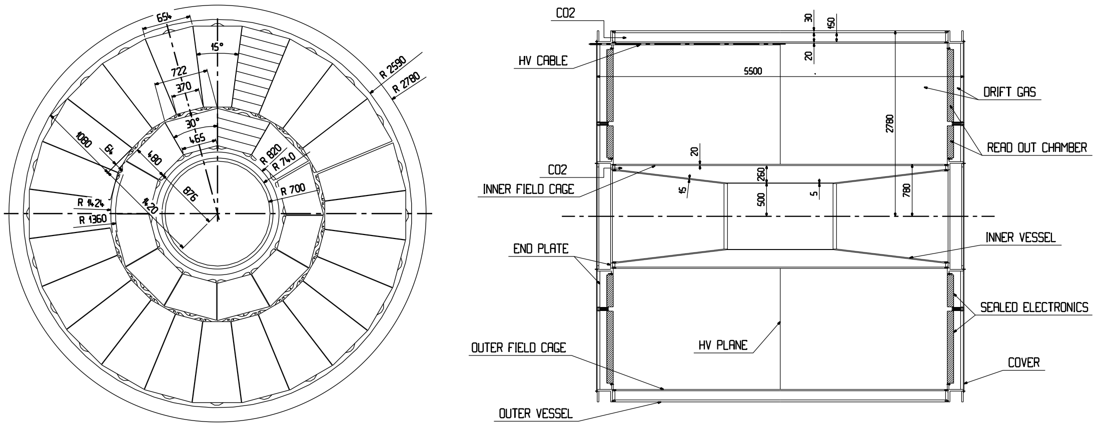

Figure 3.2: Schematic layout of the TPC.

pitch of the pads in the $r\varphi$ direction of about 3.5 mm and to a sampling in the z direction of about 1000 time buckets per track, with a single sample length of ~ 2.5 mm. The justification for the choice of the pad length is discussed in Section 3.11.

In addition, other geometrical effects also contribute to the signal width. These effects stem, for example, from the topology of the tracks with respect to the position of the pads, the pad-wire geometry, and the relative direction between the electrical drift field **E** and the magnetic field **B**. These effects have been taken into account in the evaluation of the performance of the TPC and are discussed in Section 3.11.

The materials for the field cages and the readout chambers have been selected to keep effects of multiple scattering and the production of background at a low level. Carbon-fibre composites will be used for the readout chambers since their radiation length and elastic modulus are large.

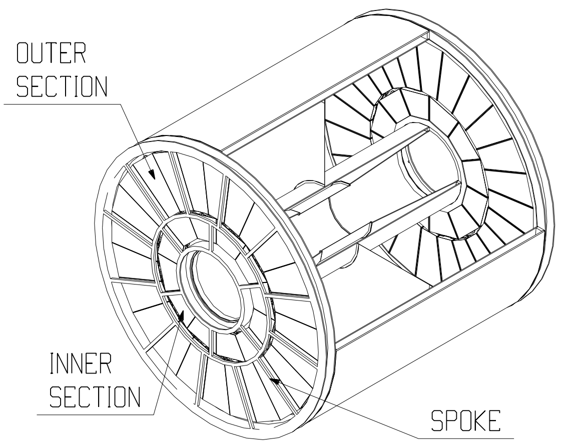

Figure 3.3: General view of the TPC.

3.3 Mechanical structure

The mechanical structure of the TPC will be composed of the outer vessel, the end flanges, and the inner vessel, as can be seen in Fig. 3.2.

The outer vessel has an overall diameter of 5560 mm, a length of 5500 mm and consists of a sandwich structure made up of two skins of carbon-fibre composite material and a core of Aramid fibre honeycomb (NOMEX). The end flanges, supporting the readout chambers, will be assembled by three concentric rings interconnected by radial rods; as shown in Fig. 3.3. They are made from a carbon-fibre composite. The construction of the inner vessel is governed by the following requirements:

- to integrate the conical front-end absorber of the muon arm;
- to provide an electrical ground plane at a sufficient distance from the inner field degrader;
- to put the material as close as possible to the vertex in order to minimize the influence of multiple scattering on the resolution.

Therefore, the inner vessel consists of three parts: two conical end pieces, of the same composition as the outer vessel, which support the ITS, and an extremely thin cylinder, in the acceptance region of the TPC ($|\eta| \leq 0.9$), which provides gas tightness and thermal insulation towards the environment. In this region, an inner diameter of 1 m allows the introduction of the ITS.

The mechanical stiffness of the TPC will be mainly provided by the outer vessel and the end flanges.

In Table 3.1 the thickness of the various components is listed in units of radiation length. As indicated in Fig. 3.3, the two readout planes, situated at both ends of the detector, will be composed of 2×12 inner (S-sector) and 2×24 outer (L-sector) modules. The inner detector plane will provide the data for the reconstruction of the trajectories in the radial range of $0.88 \leq R \leq 1.36$ m, the outer one will be efficient for $1.42 \leq R \leq 2.50$ m, thus leaving

Table 3.1: Thickness of various components of the mechanical structure of the TPC.

Element	Material	X/X_0 (%) at $\eta = 0$
Inner vessel $\lvert\eta\rvert \leq 0.9$	G10, al. Mylar, NOMEX	0.27
Insulating gas	CO_2 (260 mm)	0.14
Inner field cage	G10, al. Mylar, NOMEX	0.52
Drift gas	Ne/CO_2 (90/10)	0.59
Drift gas	Ar/CO_2 (90/10)	1.72
Outer field cage	G10, al. Mylar, NOMEX	0.52
Insulating gas	CO_2 (150 mm)	0.08
Outer cylinder	C fibre, al. Mylar, NOMEX	1.33
Total		3.45 (Ne) – 4.58 (Ar)
Element	Material	X/X_0 (%)
Cone pieces $\lvert\eta\rvert > 0.9$	C fibre, al. Mylar, NOMEX	1.15
End plate	C fibre, al. Mylar, NOMEX	1.72
Pad chambers+elect.		19
End plate cover	C fibre, al. Mylar, NOMEX	0.86

a dead area ring of 6 cm needed to support the chambers.

In order to avoid any transfer of mechanical stress from the end flanges, and minimize the deformations of the pad-planes, each readout module will be held to the end flange by three support pieces. Finite element calculations have shown that the deformation of the end flange under the gravitational load of the readout modules will be less than 0.2 mm.

To achieve this, the TPC has to be supported on the outer diameter of each end flange in the horizontal plane. This will be effected by two arch-shaped support bars, each connected by four thin blades (5×200 mm^2) to the space frame, as seen in Fig. 3.4. This arrangement allows

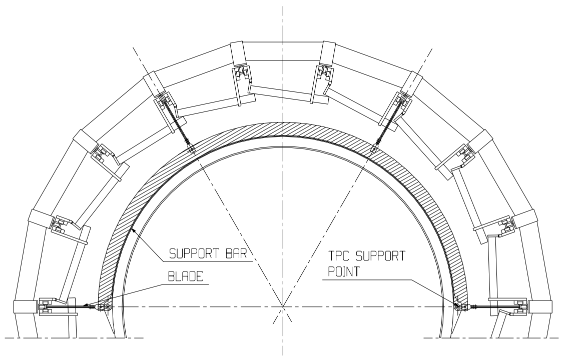

Figure 3.4: Support of the TPC inside the PID.

one to minimize the amount of material for the support and gives the possibility to centre the TPC and the ITS to within 0.2 mm with respect to the beam axis.

The two cylindrical field degraders, described in the next section, will be fixed to the end plates and divide the total space of the TPC vessel (~ 100 m^3) into three gas-tight volumes. The main part will be filled with the drift gas mixture as can be seen in Fig. 3.2. Similar to the STAR design, pure CO_2, flowing through the two other volumes, will provide sufficient electrical strength to hold a 100 kV high tension between the field cages and the inner and outer vessels.

The readout modules will form gas-tight containers to house the front-end electronics, avoiding any contamination of the drift gas; they will be flushed either with pure argon or with the same mixture as the drift gas, but from different supplies.

Since the front-end electronics will incorporate zero suppression and optical-fibre multiplexing, only the trigger, the gating, the supply voltages, and nine optical fibres for each L-sector and four fibres for each S-sector will have to be connected to each readout module by means of gas-tight feedthroughs.

The drift gas mixture has been chosen in order to fulfil different requirements:

- non-flammability
- electron mobility ≥ 6 cm$^2/\mu$s kV
- diffusion coefficient $\leq 250\,\mu$m$/\sqrt{\text{cm}}$

Unfortunately, even for an electric field of 400 V/cm, the drift velocity is not saturated and varies with temperature by 0.4 %/°C; in addition, for a precise dE/dx measurement, a very constant gas gain is required.

In the L3 set-up, temperature differences of up to ±3°C have been measured inside the magnet, in spite of the cooling systems mounted on the detectors. Therefore, it will be necessary to equip the TPC with an appropriate control system in order to provide a temperature stabilization of about 0.1°C over the whole volume, corresponding to a variation of the z-(drift) coordinate by 1 mm in the worst case.

In order to achieve this stability, it is proposed to oper-

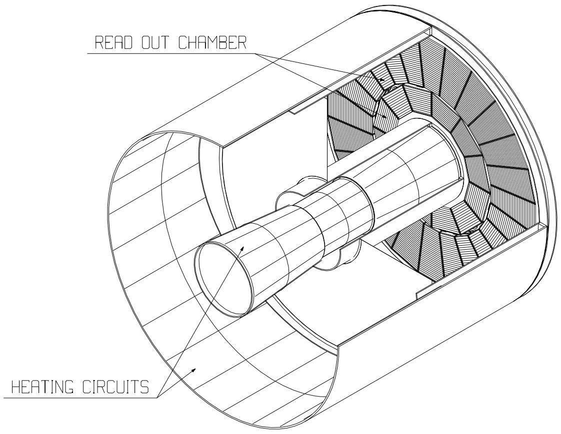

Figure 3.5: Heating circuits on the TPC's inner and outer vessels.

ate the TPC at a working temperature of 25°C, well above the ambient temperature. The surfaces of the inner and outer vessels in contact with the CO_2 will be covered by individually stabilized heating circuits, as seen in Fig. 3.5. The NOMEX honeycomb in the composite structure provides the thermal insulation with respect to the outside environment. Additionally, the readout modules will be equipped with water cooling systems at 25°C, to extract the heat produced by the front-end electronics: 400 W per L-sector and 200 W per S-sector, in total 24 kW.

The main parameters of the ALICE TPC are summarized in Table 3.2.

Table 3.2: TPC dimensions and gas properties.

Sensitive volume:	
Inner/outer radius	0.88 / 2.50 m
Length	5 m
L-sector:	
R_{min} / R_{max}	1.42 / 2.50 m
No. of pads / pad rows	$\sim$ 8200 / 54
S-sector:	
R_{min} / R_{max}	0.88 / 1.36 m
No. of pads / pad rows	$\sim$ 4400 / 24
Pad size (average)	20×3 mm^2
Pad pitch in $r\varphi$	3.5 mm
Gas volume	83 m^3
Gas mixture	(Ar, Ne)/CO_2
Drift velocity	$\sim$ 2.5 cm/μs
Diffusion constants:	
σ_{long}	$\sim 250\ \mu\text{m}/\sqrt{\text{cm}}$
σ_{trans}	$\sim 250\ \mu\text{m}/\sqrt{\text{cm}}$
MIP specific ionization:	
Ar mixture	$\sim 90\ e^-$/cm
Ne mixture	$\sim 45\ e^-$/cm

3.4 Field cage

3.4.1 Mechanical description

The electrical field cage, consisting of the high-voltage central plane and the inner and outer field degraders, will provide a uniform electrical field of up to 400 V/cm, which defines the drift trajectories of the ionization electrons from the charged particle trajectories towards the readout planes.

Owing to the 'slow' drift gas, we foresee a total high voltage of up to 100 kV. The electric field uniformity in the total sensitive volume is fundamental for the accuracy of the particle momenta measurement. Therefore, the central high-voltage electrode and the readout planes have to be flat and strictly parallel to each other. The field degrader cylinders will be equipped with potential defining strips which minimize field distortions at the boundary.

The construction of the outer field degrader cylinders is shown in Fig. 3.6. They will consist of sandwich structures of a 20 mm plastic honeycomb and two thin (0.1 mm) glass-fibre skins. The inside and outside surfaces of each cylinder will be covered by a Mylar foil, equipped on both sides with aluminium strips forming the potential guard rings; the staggering of the two series of parallel rings between the two sides of the Mylar sheet will provide, as in the ALEPH TPC, a complete screening of the non-aluminized areas of the insulating layer. Thus, the field distortions outside the field cage will not penetrate into the sensitive area [8].

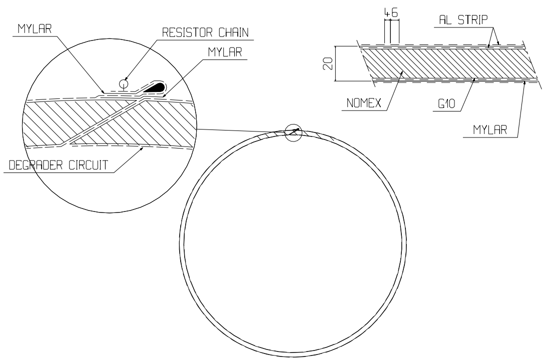

Figure 3.6: Construction of the field degraders.

In order to avoid complicated HV feedthroughs, it is planned to use one single Mylar foil (Fig. 3.6) to cover the inside and the outside of the field degrader. To establish a uniform potential gradient, both the inside and outside strips will be connected, at least at one end of the Mylar ribbon, to a linear high-voltage resistor chain.

The high-voltage central plane, made out of a double-sided aluminized Mylar foil, will be stretched between two conductive rings. The outer ring is attached to the

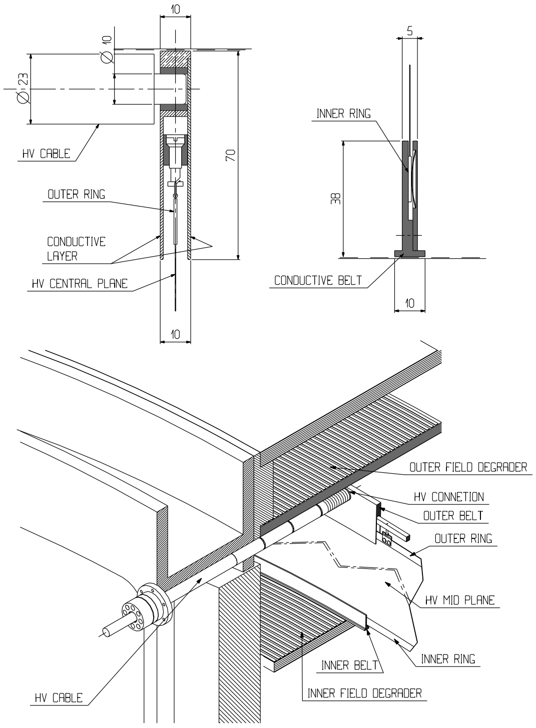

Figure 3.7: Installation of the HV cable and HV mid-plane. The figure is not to scale.

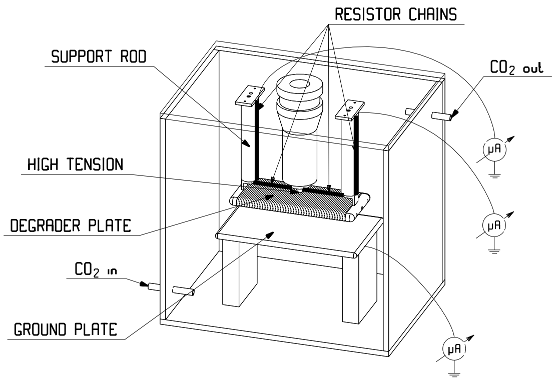

Figure 3.8: Test set-up for the gaseous HV degrader.

middle of the outer field degrader, as shown in Fig. 3.7. The inner ring is connected to a conductive belt, which is attached to the middle strip of the inner degrader cylinder. The high-voltage cable (see Fig. 3.7) will be guided in a degrader tube as in the DELPHI Barrel RICH [9] and connected to the conductive ring, which is mounted on the outer degrader cylinder and supports the high-voltage mid-plane [1].

3.4.2 High-voltage insulation

The field cage will be surrounded by CO_2 to degrade the 100 kV tension from the mid-plane to ground as described in Section 3.3. A gaseous degrader has been chosen because of its low mass. The design is similar to the STAR layout [10]. The distance from the field cage is 150 mm to the outer cylinder and 250 mm to the inner cylinder, corresponding to maximum degradation fields of 6.7 kV/cm and 4 kV/cm, respectively. In order to verify that these fields can be achieved, a test set-up, as shown in Fig. 3.8, has been built. In a box filled with CO_2 a degrader plate of 40×30 cm^2 is mounted on two insulating rods. The degrader plate is equipped with strips similar to the ones later used in the TPC. The HV of up to 100 kV is applied to the mid-strip of the plate. Two resistor chains degrade the central HV with 400 V/cm to each side of the degrader plate, simulating the electrical environment met later in the TPC. From here the potential is degraded to ground potential by two resistor chains fixed on the support rods. An aluminium plate, connected to ground by a microampere-meter, is mounted parallel to the degrader plate opposite to the support rods. The distance between the degrader and the ground plate can be between 50 mm and 250 mm. Special care has been taken to round all corners in the set-up, to avoid any corona discharge.

The electrical current into the ground plate, as well as the currents through the resistor chains, are recorded. First preliminary results show that such a gaseous degrader is able to withstand the fields required for the ALICE TPC. Discharges between the two plates have only been observed at degrading fields as high as 14 kV/cm, whereas at 13 kV/cm the system has proven to be stable without any noticeable corona discharge. The ongoing tests will define, in the near future, the geometrical and mechanical requirements with which the field cage and the outer and inner cylinders will finally be designed and assembled.

3.4.3 R&D and alternatives

We plan to design and construct the central part of the field cage on a 1:1 scale for testing. This cage could then also be equipped with prototype readout chambers to investigate not only the behaviour of the carbon-fibre chamber, but also to test the alternatives to this conservative approach (see Section 3.5.5).

Other TPC geometries with better space-point resolutions than the cylindrical ones and better adapted acceptances for the spherical particle emission pattern at colliders are sometimes discussed as desirable, but no real effort has been made to study these so far. To develop a new type of TPC field cage geometry, e.g. with spherical shape for collider experiments, simulations are planned to investigate the drift behaviour of the electrons in a spherical inner and outer field cage with radial field lines. At the present preliminary stage such a solution cannot be ruled out but requires a dedicated R&D effort.

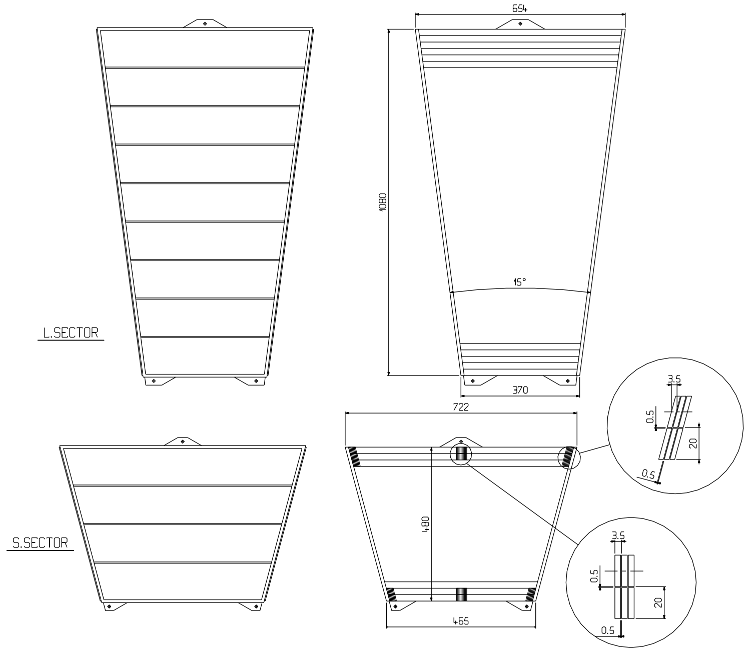

Figure 3.9: Sector frames L- and S-type and the pad row geometry.

3.5 Wire chambers

3.5.1 Mechanical structure

The wire chamber sectors are mounted on each side of the cylinder, on the end-cap wheels, subdivided into outer sectors of 12 spokes and inner sectors of 6 spokes (see Fig. 3.3). Two sectors will be mounted between each two spokes. Each sector is suspended at three points on the wheel. This suspension allows the sectors to be aligned. Alignment procedures do not influence the gas-tightness since the closing flanges are behind the sectors (Section 3.3). Each readout plane is subdivided in two regions along the radial direction, one extending from $R_{\rm min} = 88$ cm to $R_{\rm max} = 136$ cm, the other from $R_{\rm min} = 142$ cm to $R_{\rm max} = 250$ cm, leaving a boundary of 6 cm. The inner ring is composed of 12 small (S-type) sectors, extending 30° in azimuth, the outer one of 24 large (L-type) sectors, of $\Delta\varphi = 15°$ each (see Fig. 3.2). The distance between the sectors, in the azimuthal direction, is about 3 cm. Stiff tracks that fall on such 'cracks' are measured either in the L- or in the S-type sectors, since the two rings are rotated relative to each other by $\Delta\varphi = 7.5°$. The sectors have straight boundaries in the azimuthal and radial directions. Zigzag-type boundaries, as used in the past (ALEPH), have not been considered for ALICE. The complication in the mechanical construction would have outweighed the advantages in accuracy and resolution.

In all the sectors, only the pads are read out, since at the high track density the signals at the wires are ambiguous.

The mechanical strength of a sector is given by a frame of carbon-fibre (see Fig. 3.9) designed such that the forces exerted by the wire tension, together with a slight overpressure of 0.5 mbar of the gas, cause a bend of the pad-plane by about 20 μm.

The pad-plane is glued as a G10 multi-layer board to the carbon-fibre frame.

The readout sectors on both sides, together with the high-voltage plane, have to be parallel to each other within an accuracy of 10^{-4} rad. This will be achieved by controlling the alignment with photogrammetric methods.

3.5.2 Wire grids

There are three wire grids in each sector: the sense/field wire grid, the Frisch (null) grid, and the gating grid. The wire grids and their spacings are shown in Fig. 3.10. In Table 3.3 further information on the wire grids is given. Each wire plane is soldered and glued to a G10 frame that is then fixed to the chamber frame.

Table 3.3: Characteristics of the wire grids.

Grid	Material	Wire spacing	Wire diameter	Tension
Sense	W-Ru	4 mm	20 μm	50 g
Field	Cu-Be	4 mm	125 μm	300 g
Frisch	Cu-Be	1 mm	75 μm	120 g
Gating	Cu-Be	2 mm	75 μm	120 g

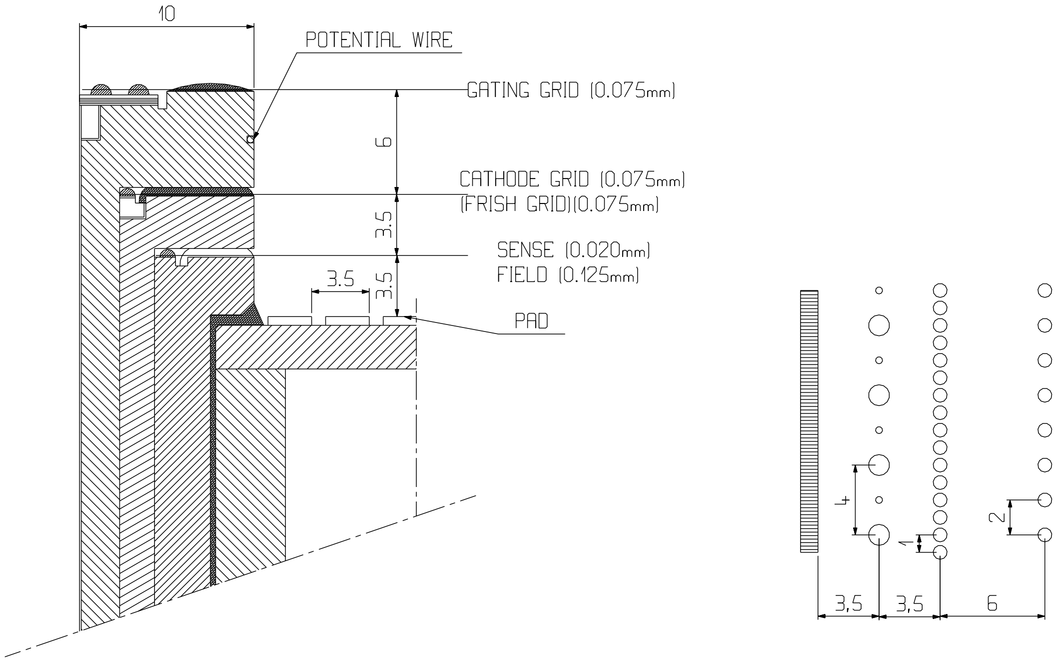

Figure 3.10: Spacing of the wire grids.

3.5.3 Pads

Straight pad rows, as shown in Fig. 3.9, have been chosen. The size of the pad row in the radial direction is 20 mm, the pads have parallelogrammar or rectangular shape, depending on the azimuthal position. The pad length in the radial direction varies slightly with the azimuthal position of the pad, from 19.5 mm to 20.18 mm in the L-sector (Fig. 3.9) and from 19.5 mm to 22.5 mm in the S-sector. In the $r\varphi$ direction the pads have a width of 3 mm and a pitch of 3.5 mm in both sector types.

The L-sectors have 54 pad-rows with an average number of 153 pads/row, in total $\sim$8200 pads per L-sector. The S-sectors have 24 pad-rows with an average of 184 pads/row, in total $\sim$4400 pads per S-sector. In both readout planes together there are in total about 500 000 pads.

Circular pad-rows have also been considered. The improvement in space-point resolution compared to straight pad-rows is 100 μm for the lowest-momentum tracks and vanishes at higher momenta (>500 MeV/c). As the momentum resolution is not affected by a slightly worse space-point accuracy, we have decided on straight pad-rows, which are easier to produce.

3.5.4 Electronics mounting and cooling

The copper pads are engraved on a multi-layer board with one preamp/shaper for every four channels. They are mounted on the backside using the 'flip-chip' technique, so that the preamp/shaper part is as close as possible to the pads, keeping the cross-talk minimal. For reasons of yield it was decided to have four channels/chip. The yield can then be expected to be 90 %, whereas for 8 or even 16 channels/chip, the yield drops down to about 70 %.

The other front-end electronics, mounted on the front-end cards, is connected by micro-connectors and Kapton cables to the pad-multilayer board. It is anticipated that 128 channels will be mounted on one front-end card.

The electronics and cooling lines (Fig. 3.11) are enclosed in a volume separated from the sensitive gas volume of the TPC, filled with the same gas but from a different supply. A detailed discussion of the electronics is given in Section 3.6.

3.5.5 R&D and alternatives

The total area that is used for the electronics is estimated to be about 50 mm^2 per channel. Given the pad size of 60 mm^2, it would in principle be possible to have the electronics directly behind the pads, avoiding the micro-connectors. However, uncertainties as to how much the TPC gas will be heated by the electronics, despite cool-

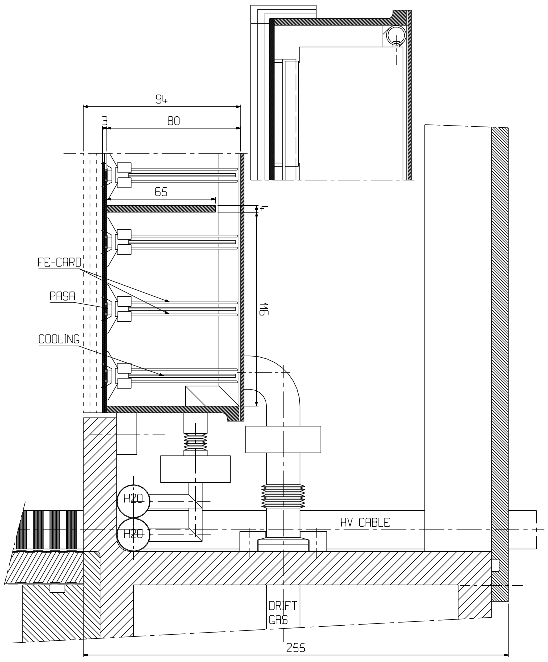

Figure 3.11: Electronics and cooling.

ing, led us to the safer scenario shown in Fig. 3.11, with the electronics decoupled from the pad-plane, except for the preamp/shaper (PASA). This aspect of mounting the electronics, as well as the backup solution with the PASA mounted on the front-end card, will be investigated in the near future.

The proposed wire chamber concept is regarded as a baseline proposal that can be constructed today. However, other concepts might have advantages in production costs and gas amplification. It should be mentioned that a small gas amplification is desirable not only for dE/dx measurements but also to keep the current in the chamber low, resulting in stable working conditions. However, every new concept should be compared with the baseline one regarding performance and cost.

Within the RD32 project [2] two new concepts are under investigation. The concepts of Ring-Cathode Chamber and Micro-Strip Gas Chamber are alternatives to the conventional wire readout chamber. In the following these two developments are discussed.

The Ring-Cathode Chamber

A new readout concept for TPC applications has been developed with the following aims:

1. maximize the fraction of the signal induced on the readout pads;
2. minimize the number of wire planes needed in the chamber construction.

This has been achieved by using a pad structure which completely surrounds the sense wires (see Fig. 3.12). The pads are made of thin (50 μm) Cu-Be sheets aligned vertically to the sense wires which pass through circular holes in the sheets. The basic detector cell is therefore a cylindrical tube around the sense wire. This geometry has the advantages of operational stability and insensitivity to geometrical tolerances. The hole diameter defines the width of the pad response function, and the distance

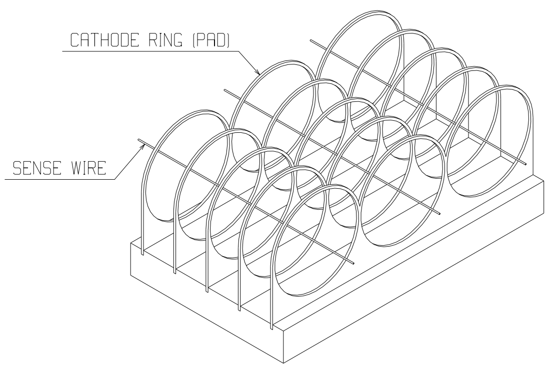

Figure 3.12: Schematic layout of the ring-cathode chamber.

between neighbouring pad sheets along the sense-wires determines the granularity with which the pad response is measured. It is apparent that the size of the induced signal in this pad structure is about equal to the sense-wire signal. This yields a factor of about 4 in signal size as compared to a normal TPC pad-plane readout. This is of importance for LHC applications, where the charge load on the readout chamber is a major source of concern. In addition, both the field wire and the Frisch grid planes can be omitted. One is thus left with only two wire planes, i.e. the sense wires and the gating grid.

In a test chamber the pads are 1.5 mm apart, along the sense wire. The readout granularity is thus well adapted to the width of the pad response function, of about 1.8 mm (r.m.s.). This optimization is possible for any chamber geometry by selecting the pad-sheet distance or by grouping together several pad sheets on one readout channel. The time distribution of pulse height in this structure is much simpler than in the normal TPC readout chambers. It has a well-defined $1/t$ behaviour (at least for the drift time of the TPC) and avoids complex undershoot and overshoot features.

The pad sheets can be produced by standard techniques like chemical etching, electro-erosion or stamping. A prototype chamber of a few square centimetres in size, shown in Fig. 3.12, has been built using chemical etching. The hole diameter for the sense wires is 4 mm, with eight sense wires passing through one pad corresponding to a pad length of 3.3 cm.

The detector response has been studied with α-particles from ^{241}Am and photons from ^{55}Fe sources. The corresponding ionization charges were transferred into the

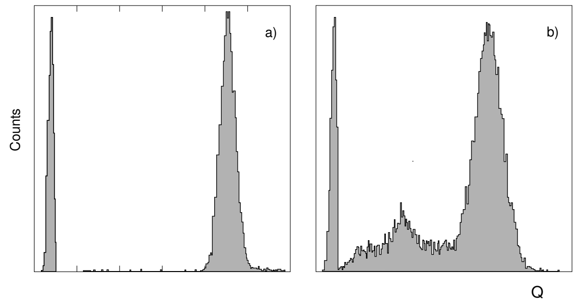

Figure 3.13: Pulse height spectrum, obtained with the ring cathode chamber, for (a) α-particles from ^{241}Am source and (b) 5.9 keV photons from ^{55}Fe.

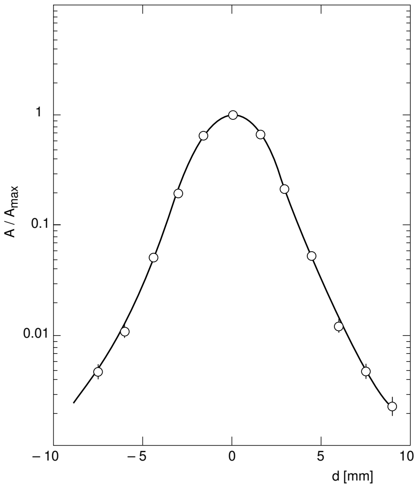

Figure 3.14: Pad response function of the ring-cathode structure. The line helps to guide the eye.

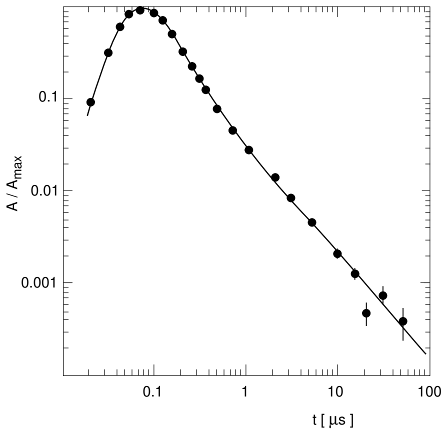

Figure 3.15: Time distribution of the induced cathode signal without tail suppression, showing a $1/t$ dependence above 1 μs.

chamber through a drift channel with 5 cm drift length. The pulse-height distributions shown in Fig. 3.13 demonstrate the good charge collection quality of the detector. It was verified that the total charge in the pad signal (measured over shaping times of the order of 200 ns) is equal to the charge in the anode wire signal. This is to be compared with a charge fraction of only about 25 % found on the pads of standard TPC chambers. The pad response function (uncorrected for source collimation and diffusion) is presented in Fig. 3.14. The data points are well fitted by a Gaussian curve with $\sigma = 1.8$ mm, as expected from the hole diameter of the sense-wire openings. The pulse shape corresponds to the expected $1/t$ behaviour, above 1 μs, as our measurements show in Fig. 3.15.

The Micro-Strip Gas Chamber

A promising alternative to the conventional design of the TPC readout modules is the use of Micro-Strip Gas Chambers (MSGC) [11,12]. Here the multiplication plane of the readout module consists of an anode–cathode structure deposited directly on a glass surface. The shape of this structure is made by a lithographic process. In this way the mechanical layout and the tolerances in the multiplication structure are defined by this lithographic process. This offers many advantages in comparison with wire chambers:

- Short and even curved anodes are possible.
- Gas tightness in the readout modules is simplified.
- The number of grids can be reduced, thus improving the mechanical stability of the readout chambers.

Up to now, two main problems have prevented the large-scale use of a MSGC TPC.

Firstly, the gas gain changes as the electrical properties of the glass deteriorate. This problem can be solved by using electronconductive glass with a bulk resistivity of $\rho \approx 10^{11}\ \Omega\,\mathrm{cm}^{-1}$. Two types of such glass are known: the Pestov glass (see Section 4.2) and Schott S8900. The electronic conductivity in both cases is produced by a ~ 40 % content of Fe-oxide in the glass. Secondly, the pad response function of a MSGC was up to now unknown.

In the context of RD32 a test TPC has been constructed with a $0.45 \times 50 \times 50$ mm^3 MSGC glass plate. The MSGC structure has a pitch of 1 mm and an open cathode (Fig. 3.16) [13]. Owing to the open cathodes, this arrangement gives a backplane-to-anodes signal ratio of 60 % at a glass thickness of 0.45 mm. The pad length is 4 cm, the width is 1 mm, and the pad pitch is 1.5 mm. The 16 pads were evaporated on the back side of the glass. Two pad types were investigated:

1. Pads evaporated onto the back side of the glass.
2. Pads on a PC-board which was mounted on the back side of the glass with an insulating Kapton foil in between.

The TPC was triggered by a β-source. After a drift length of 7 cm the electrons of the track were ampli-

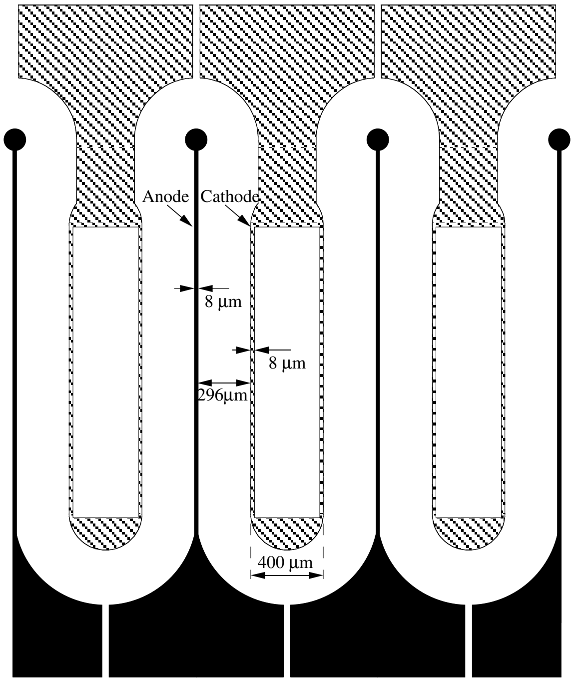

Figure 3.16: Layout of the MSGC structure with a pitch of 1 mm used in the test TPC.

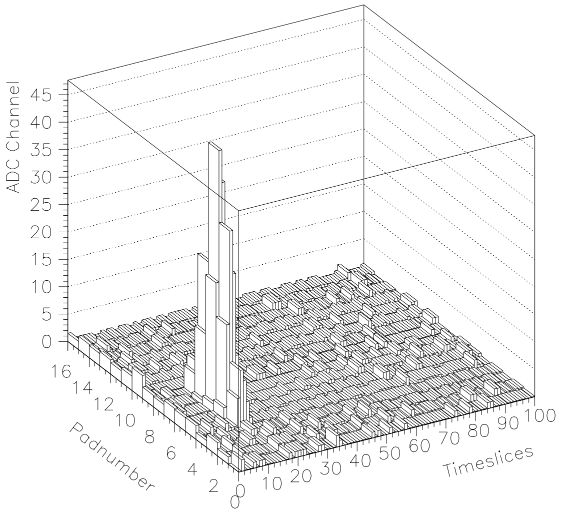

Figure 3.17: Typical response, in the MSGC prototype, for a minimum-ionizing particle.

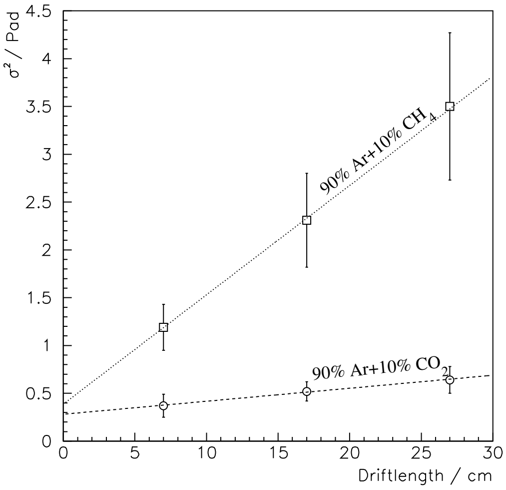

Figure 3.18: The squared cluster width σ^2 in units of pads as a function of the drift length.

fied at the MSGC. The pad signals were read out by the NA35 electronics. The amplification gain in this test was a factor 10 below the maximum possible gain with this MSGC plate. With both methods more than 50 % of the anode signal was observed on the pads, which is sufficient for a TPC. In Fig. 3.17 a typical event from the test TPC is shown. The resulting widths of the pad response function are shown in Fig. 3.18 for two different gas mixtures: Ar/CH_4 and Ar/CO_2 (90/10). The values are 0.8 mm (Ar/CH_4) and 0.93 mm (Ar/CO_2), for a pad pitch of 1.5 mm [14].

These results are based on measurements with a single glass thickness. Investigations of different glass thicknesses are underway, to pin down the value that would lead to a pad response function of $\sigma_0 \approx 3$ mm.

The current ageing measurements are in good agreement with the results of Ref. 15. Ageing effects in MSGCs pose additional conditions on the choice of the drift gas. Gases as proposed here for conventional readout chambers might not be applicable for use in MSGCs. A systematic gas study has also to be performed to substantiate the advantage in replacing conventional readout chambers with MSGCs.

3.6 Electronics

3.6.1 Requirements

In the TPC, the position in the z (drift) direction of a charged particle is determined by measuring the time the electron cloud, generated by primary ionization, drifts towards the sense wire. The r and $r\varphi$ coordinates are derived from the fraction of the signal which is induced on the pad plane, located behind the sense wire plane. For every pad, the charge is integrated and subsequently shaped by a shaping amplifier. The pulse-height spectrum covering the maximal drift time is stored either as analog or digital data. Then, the position of the charge centroid is a measure of the spatial coordinates and the charge integral leads to a specific energy-loss measurement.

The requirements for the front-end electronics are listed in the Table 3.4. Two basic architectures, shown in Fig. 3.19, are being considered for the TPC front-end electronics:

1. Preamp/shaper, analog storage (SCA), ADC, multiplexing;
2. Preamp/shaper, ADC, digital memory, multiplexing.

For the moment, both approaches will be pursued with the same priority.

Table 3.4: Front-end electronics, requirements.

Item	Value
No. of channels	500 k
No. of samples/channel	1000
Noise (r.m.s.)	800 e
Signal-to-noise (MIP)	40 : 1
Signal-to-noise dynamic range	800 : 1
Max. integrated charge	0.2 pC
Input capacitance	< 30 pF
Gain	5 mV/fC
Gain variation	< 10 %
Nonlinearity	< 2 %
Cross-talk	< 0.3 %
Tail correction after 1 μs	< 0.1 %
Tail correction after 10 μs	< 0.2 %
Output peaking time	240 ns
Shaping time variation	< 5 %
ADC range	10 bits (8 nonl.)
Sampling rate	8 MHz
Digitization time per event	0.5 - 2 ms
Total no. of pixels	500 M
After zero suppression in Pb–Pb	~ 50 M
Channels per link	~ 1000
No. of optical links	~ 500

To achieve the necessary rate capability, in both cases the zero suppression will be done before the data transfer.

It is planned to test a sizeable amount of electronics of both kinds in the environment of NA49, where the track density in Pb–Pb collisions is higher than in ALICE, before taking the final decision.

These tests can be done during the heavy-ion running periods of 1996/97.

Drift velocity, drift length, and degree of diffusion determine most of the parameters for the front-end electronics. For the TPC a cool gas mixture with a drift velocity of 2.5 cm/μs has been chosen as the most probable candidate. The average longitudinal diffusion determines the peaking time of the shaper. It is chosen so that the shaping time of 240 ns is comparable to the diffusion width of the pulse in the time domain. From the shaping time, a sampling frequency of approximately 8 MHz can be derived. Thus, the total drift space of 250 cm is divided into

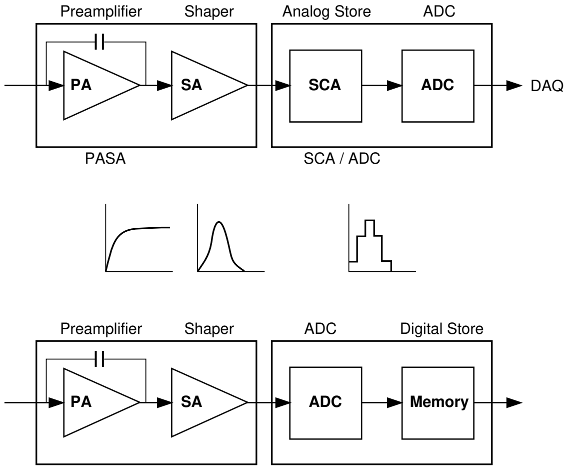

Figure 3.19: Basic front-end architectures.

about 1000 time samples. The dynamic range of the ADC is determined by several factors:

- The signal-to-noise ratio is 40 : 1 for a mean minimum-ionizing particle.
- The ionization of a 280 MeV/c proton is a factor of ~ 10 above the one of a minimum-ionizing particle.
- Owing to the Landau distribution of the ionization, one has to accommodate at least a factor of two for fluctuations.

Those factors add up to a requirement of a 10-bit range for the ADC. A detailed simulation shows that the performance of the TPC does not deteriorate by storing 8-bit words provided that the ADC has a nonlinear response. This can be achieved by a nonlinear ADC followed by a look-up table later in the analysis.

3.6.2 Front-end electronics

Preamp/Shaper

The charge collected on the TPC pads is amplified and integrated by a low-noise, low-input impedance preamplifier. If extensive offline data correction or digital filtering procedures during data transfer are to be avoided, the high channel occupancy combined with the Landau fluctuations calls for very precise tail suppression in the analog input stage. In order to be compatible with the required dE/dx resolution, a suppression to 0.1 % or better of the maximum pulse height, is indeed required.

The pulse shape of pad signals in a TPC geometry is rather complex. It depends on the details of the chamber and pad geometry. The $1/t$ tail behaviour from a closed proportional tube is replaced by a bipolar signal due to the particular motion of positive ions relative to the pad and wire planes. The negative undershoot reaches several per mille of the peak pulse height and falls into the normal drift time regime of a TPC.

The necessary undershoot suppression can be handled with normal pole-filtering circuitry, requiring four poles and four amplifiers to achieve a suppression of $\sim 0.1\,\%$ after 0.7 μs and a shaping time of 200 ns (FWHM).

Within RD32 [2], a gateable preamp/shaper has been developed and is presently being tested. It uses the good high-frequency response of the translinear circuit to cut off the low-frequency domain by a simple combination of differentiation and integration with a minimum of two time constants. The required pulse width is then achieved by a special current-mode integration circuit.

This new layout has been produced as a fully customized IC in bipolar, radiation-hard technology and tests have shown that this design works according to the simulations. This circuit gives a pulse with a shaping time width of 70 ns (FWHM) and has a noise value of about 1500 e (r.m.s.). In a next step the width will be set to 200 ns (FWHM) and a noise value of about 800 e (r.m.s.) is expected. The long-time tail suppression of the delivered chip behaves as expected (Fig. 3.20). We propose to

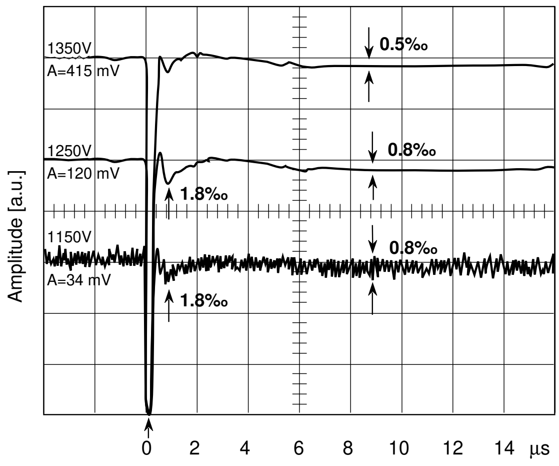

Figure 3.20: Pulse-height signal from the preamp/shaper.

use this new preamp/shaper chip in the baseline design, in the 4 channels/chip version, with the chips mounted directly on the multi-layer board behind the pads.

The size of the chip is 0.7 mm^2 per channel and the power consumption is 4.0 mW/channel ungated (d.c.).

SCA/ADC

The time history of the pad signal is fed to a switched capacitor array (SCA) analog memory that is clocked at approximately 8 MHz.

A new SCA chip was developed, adding a threshold comparator to the input of every channel. This has the advantage that zero suppression can be done on the analog data. Furthermore, the foreseen time bucket occu-

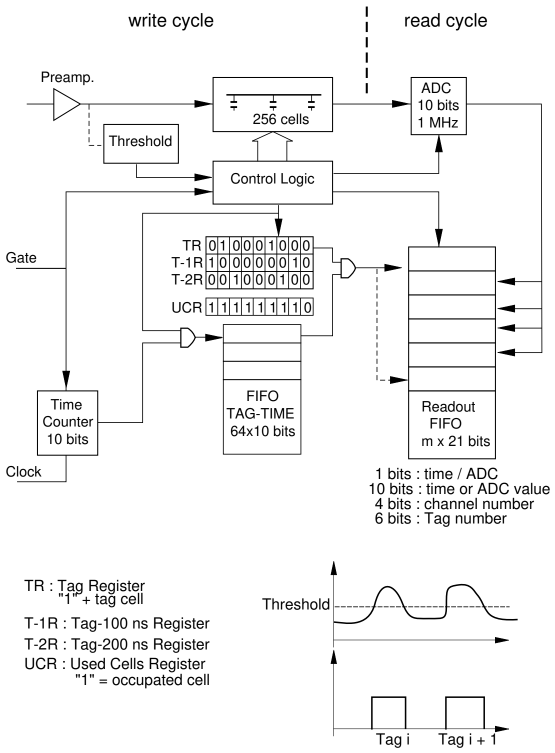

Figure 3.21: Schematics of the SCA and associated logic.

pancy of each TPC pad being only 10 %, the SCA can work with fewer sampling capacitors (256 cells). However, a time stamp must be added to each cluster of time buckets and the level of background, before the integrated charge crosses the threshold, is recorded during the two preceding time slices (2×100 ns). Figure 3.21 shows the layout of one SCA channel with the discriminator, read-in and readout logic. When a valid trigger arrives, a clock is started and during 100 ns the output of the preamp/shaper is sensed by the discriminator and integrated in the first SCA cell. As long as the integrated charge per sample does not cross the threshold, the samples are written in the first three contiguous cells, configured as a circular buffer (Fig. 3.22). When the threshold is crossed, the time value is stored in the FIFO memory as a time stamp and a flag is stored in the corresponding 'Tag Register' cell. The two other cells of the circular buffer are flagged in the 'Tag-100ns register T-1R' and in the 'Tag-200ns register T-2R'. These three cells are now occupied and flagged as such in the UCR register; the following samples above threshold are stored in the next contiguous SCA cells. When the integrated charge per sample falls below the threshold, the next three unoccupied cells of the memory are configured as a circular buffer and the zero skipping cycle begins again. Thus, in addition to the samples above threshold, we store the two preceding samples and hence have information on the pedestal and background values [16]. Existing single-ramp ADCs satisfy

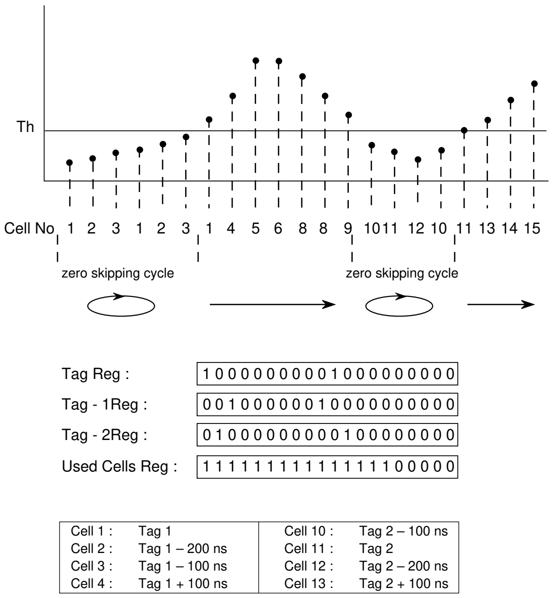

Figure 3.22: Principle of operation of the SCA write cycle, with circular buffer.

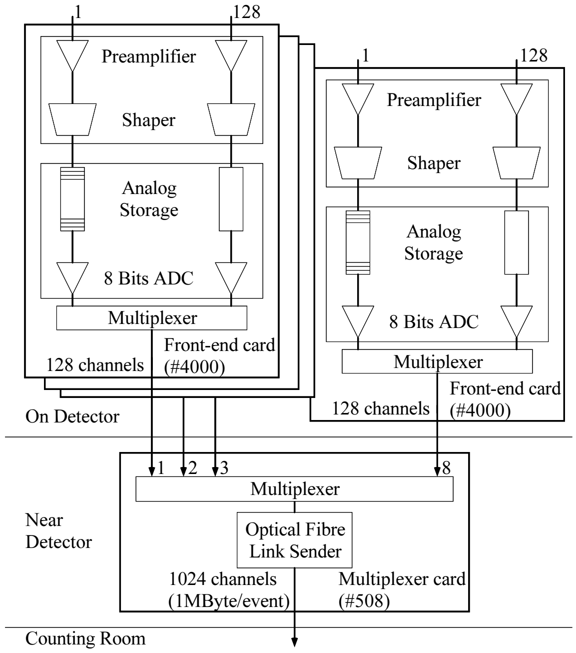

Figure 3.23: Front-end card and multiplexing.

all the requirements of ALICE and could be added to this new design on the basis of one ADC per channel.

The size of a SCA/ADC chip will be about 6.5 mm^2 per channel and its power consumption is about 1.3 mW per channel.

Multiplexing

The data of 128 channels are digitized on one front-end card, the data of eight front-end cards are then transmitted to a readout board for data collection before being transmitted off-chamber by an optical Gbit link (Fig. 3.23). The proposed SCA/ADC scheme will have a 1 MHz ADC with 10-bit accuracy. Assuming 1000 samples of 10 bits and 10 % occupancy in central Pb–Pb, 100 samples of 10 bits will be transferred, on the average, after zero-suppression. The conversion time of these 100 samples amounts to 100 μs. Therefore, the data flow from one front-end card amounts to 100 samples/channel $\times$ 128 channels $\times$ 10 bit = 16 kByte/front-end card. In the first multiplexer stage behind the SCA/ADC the multiplexing time is 20 ns per word, i.e. $320 + 100 = 420\,\mu$s per front-end card. By multiplexing eight front-end cards (see Fig. 3.23), 128 kByte of data have to be processed by the second multiplexer stage, which will take 2.5 ms for the complete event. This gives a maximum data rate of about 500 Hz at this stage. After the second multiplexing step the data are transmitted to an optical link. Existing links are able to transfer 100 MByte/s, well above our needs.

In the pp case, assuming 20 kByte per pp event and 20 events per drift time, the event size will be about 500 kByte. Assuming that the events are equally distributed over 500 fibres, the load per fibre is 1 kByte on average. A trigger rate of 1 kHz would amount to a data flow per link of about 1 MByte/s.

ADC/digital memory

Fast-conversion ADCs of the required dynamic range and precision have become available in the last few years. Conversion times of the order of 100 ns (required in our case) can be achieved with flash-ADCs or with successive approximation pipelined designs. Any of those designs could be a substitute for the SCA/ADC function. A technologically valid choice will have to be made on the basis of availability, noise performance, power consumption, geometric size, and price.

One possible technology choice is a modified version of the ADC of the FERMI project [17]. The FERMI ADC has an architecture that could quite easily be modified to a multichannel version for the ALICE detector. The ADC is designed in such a way that 14 successive approximation converters, on one single chip, work in parallel on one input signal, to achieve a high sampling rate. In case more channels are required with a lower sampling rate, the ADC could be divided into 14 individual channels with a sampling rate of 1/14 of the current sampling rate. It is also possible, for example, to manufacture seven channels with a 1/7 sampling rate.

An ADC with a sampling rate of 100 MHz and eight effective bits is under development. If successful, the modified Fermi ADC would have 14 channels with 7–8 MHz sampling frequency on one chip. Further studies will show whether a 16-channel chip is within reach in the next few years.

In order to compare the digital and analog storage architecture, the figures of merit also include the data storage. Assuming that a memory of 12 kbits is required per channel, and that this memory is implemented in a modern 0.6 μm technology, the following figures for the full system are obtained:

Sampling frequency:	8 MHz
Power/channel:	32 mW
Area/channel:	4.8 mm^2

The FERMI ADC is mentioned here as an example of power consumption and area/channel in such an architecture. The final solution has to be investigated more closely, with further R&D.

In the ADC/digital memory architecture the multiplexing scheme is the same as described before in the paragraph on multiplexing.

3.6.3 Printed circuits and mounting

The cathode board of the TPC will be of the multi-layer type, containing the pad structure on one side of the board. The pads are connected to the other side of the board through gold-plated vias, where four pads are fed into a preamp/shaper. The outputs from four of these are grouped together into a 16-channel micro-connector (Fig. 3.24) that brings the signal to the carrier board containing the rest of the electronics. Eight of these cables,

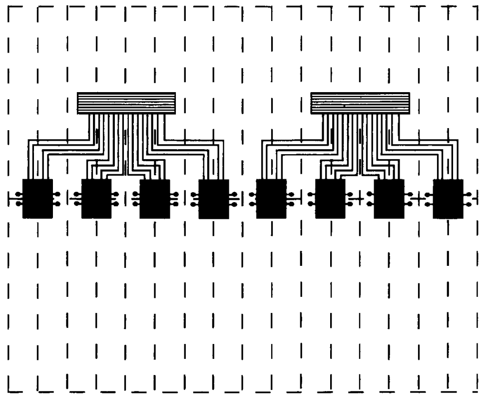

Figure 3.24: Mounting of preamp/shaper and connection to micro-connector.

with 16 channels each, will be connected to one FEE card that carries the electronics of 128 channels. The carrier boards, placed isolated from the active volume of the TPC, will hold a large number of channels, from which a multiplexed signal is transferred to the rest of the system by optical links. The ASICs on the cathode board, as well as on the carrier boards, will be mounted using the flip-chip technique: the chips are glued to the board by conducting glue. After testing, the chips are covered to protect them from dust. The flip-chip technique has opened new possibilities for very dense packing of components. By using the cathode plane as a printed circuit board, which distributes all input and output signals and the d.c. supply, flip-chip mounting could eliminate some intermediate connections between the pads and the digital outputs. The flip-chip technique is still under very intensive development but the initial tests show very promising results concerning reliability as well as serviceability.

The advantages of this mounting concept are:

- It can be applied to large areas, such as the sectors of the TPC.
- It is relatively easy to remove a malfunctioning chip once mounted. This is important from the point of view of costs and serviceability.
- Cooling can be applied directly on the chip.
- Most of the heat-producing electronics is placed isolated from the gaseous volume of the TPC, thus reducing the difficult problem of cooling.

Radiation hardness

Radiation-hard electronics is not normally necessary at the position of the TPC (see Chapter 12). However, radiation-resistive electronics will be preferable as a precaution, whenever the price is not prohibitive. The preamp/shaper samples are currently produced in radiation-hard technology.

3.6.4 R&D projects

The two proposed architectures, with analog and digital storage, have to be investigated in a R&D project before a final decision can be made. We shall design a complete circuitry, including the preamp/shaper, and measure the performance of the circuits on the NA49 TPC, in a region with high particle multiplicity.

Analog optical signal transfer is an interesting alternative to the proposed digitization and digital data transfer scheme. The project in RD23 [18] will be observed closely in the future and, if successful, the analog optical signal transfer could be incorporated in the ALICE TPC readout scheme.

3.7 Calibration

The TPC electronics poses the problem of having about 500 000 electronics channels responding precisely in the same manner to a given number of electrons collected on the sense wires. Gas amplification, coupling of the signal to the pads, linearity of the signal, and timing within the electronics have to be controlled channel by channel and corrected so that the signal variation does not exceed 1 %.

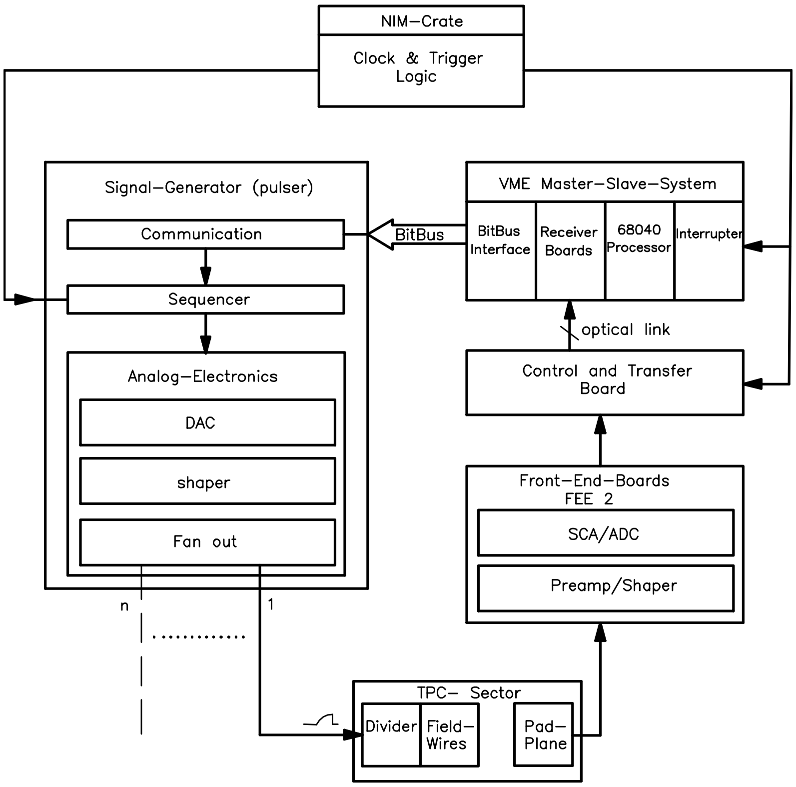

Figure 3.25: Schematics of the field wire pulsing calibration system in NA49.

In the past, several schemes have been developed to deal with these problems (see, for instance, Ref. 1). For the ALICE TPC two schemes, successfully applied in the NA49 experiment, are proposed: the pulsing of the field wire grid and the 'Krypton method' for pulse-height calibration.

Field wire pulsing

With this method the field wire planes are pulsed with a pulser (Fig. 3.25) and the signals coupled to the pads are read out and compared. In the NA49 experiment a 12-bit DAC has been used to set the amplitude, followed by a network of three RC circuits that shape the pulse, simulating as closely as possible the real particle pulses. The same pulser type is proposed for the ALICE TPC. Pulses in the range of 0–0.5 V with a nonlinearity below 0.1 % are sent to the field wire plane of each readout chamber. The stability of the pulses is better than 0.25 %. Each pulser has 25 separate output channels, the pulse-height variation among these channels is 0.04 %. This calibration is repeated on a day-by-day basis. In NA49, a calibration correction factor is determined offline, on a channel-by-channel basis, and applied to the raw data.

For ALICE it is intended to do the calibration corrections online, with VME processors, so that the data written to tape will already be corrected.

Calibration with krypton

With the 'krypton method', the variation of the gas amplification in different readout modules, as well as the variation over the surface of a single chamber, can be measured and compared.

A foil with radioactive ^{83}Rb, produced at the ISOLDE accelerator at CERN, can be placed into the incoming gas stream within a separate cylinder (see Fig. 3.32 on page 50). With a lifetime of 90 days the rubidium decays into ^{83}Kr, this in turn decays with a lifetime of 90 min.

For the 'krypton runs', a small amount of the incoming gas passes through the cylinder that contains the radioactive krypton. The krypton is transported into the sensitive gas volume and decays there. The pulse-height spectrum of the different decay lines is then measured. An example of the pulse-height spectra of two different readout chambers is given in Fig. 3.26. After a few hours all the krypton has decayed and the TPC can operate under clean conditions. This method has already been applied in NA49.

Drift velocity monitor

The constant monitoring of the drift velocity is important for the precise determination of the z (drift) coordinate. Drift velocity monitors will be implemented in the gas inlet and outlet tubes.

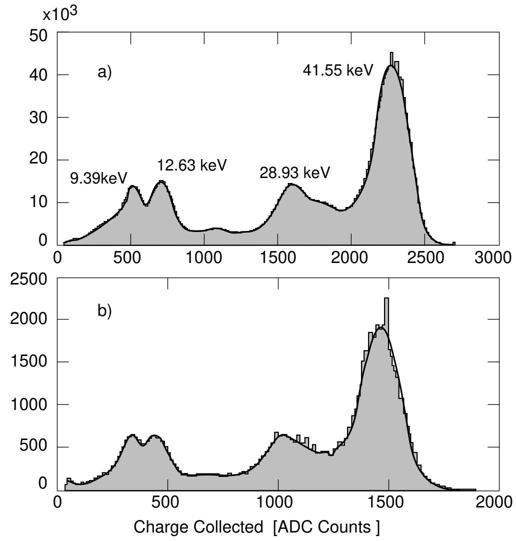

Figure 3.26: Pulse-height spectrum of ^{83}Kr, in different read-out chambers.

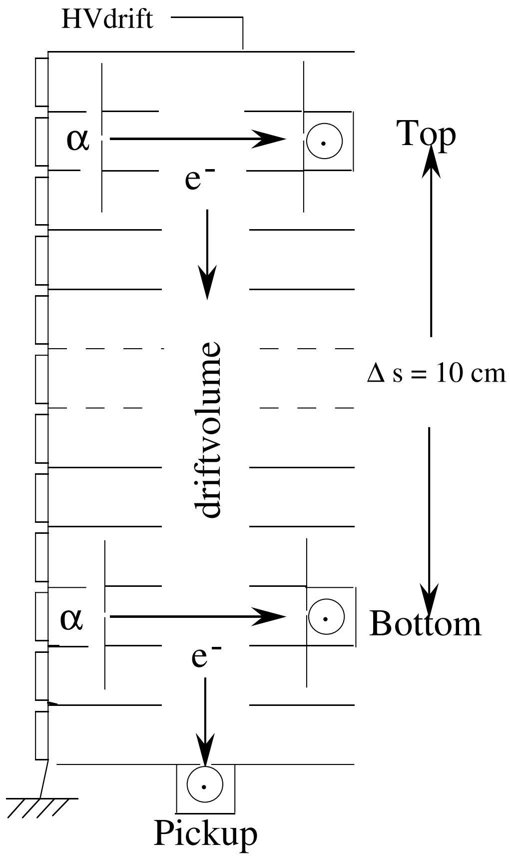

Figure 3.27: Drift velocity monitor.

The principle of the monitor is shown in Fig. 3.27. It consists of two α-sources, two proportional counters, Top and Bottom, and a pick-up counter. Either the Top or the Bottom counter starts the drift time measurement in the drift channel; the pick-up counter stops the measurement. The time distribution of the STOP signals in the pick-up counter is presented in Fig. 3.28. The drift velocity is determined from the time difference between the two peaks and the distance of the two counters Top and Bottom. It has been estimated [19] that the drift velocity can be determined by this method with a precision better than 0.1 %.

3.8 Gating

Design considerations

In order to have the TPC in the sensitive mode only during the drift time, an extra grid of wires, or 'gating grid', is installed between the Frisch grid and the drift region. During the 'closed' state, the potentials $V_G \pm \Delta V$ are placed on alternate wires of the grid, so that the resulting electrostatic field renders the gate opaque to the passage of charged particles. The value of ΔV depends on the magnetic field and, for Ar/CH_4 (90/10), it is less than 100 V for $B < 0.5$ T. The NA49 experience shows that for a cool gas, Ar/CO_2 (90/10), this value is significantly smaller: 20 V without magnetic field and 40 V for $B = 1$ T.

In the transparent state, all wires of the gating grid are kept at V_G, enabling the passage of drifting electrons.

The gating circuit has to provide pulses with amplitudes up to ± 100 V to the gating grid. Oscillations can be avoided by adapting the output impedance to the cable. The rise time is given by the internal capacity of the grid and can be estimated, for the present design of the chambers, to be of the order of 350 ns. The resulting charging time of the gating grid is of the order of 1.6 μs. It can be reduced by 30 % by compensating the capacitance of the chamber with the serial inductance, which provides critical damping in the network. Furthermore, this inductance decreases the pick-up signal induced in the preamplifier electronics.

Gating circuit design

A block diagram of the gating circuit is shown in Fig. 3.29. The application of the commercially available high-speed MOSFET drivers makes it possible to use n-channel power MOSFETs as switching devices. In this way, it is not necessary to have an extra timing of the FET control pulses. A switching speed of 100 ns can be achieved. It has to be less than twice that of the cable delay so that the reflected pulse from the chamber 'sees' the matching impedance on the pulser side of the cable. For the cable with 5 ns/m delay, this gives a minimal length of 10 m. The impedance of the pulser is determined by the resistors R_o (for gate opening), and R_c (for gate closing). Closing of the gate occurs after the pad preamplifiers are disabled, and therefore it is not necessary for R_c to match

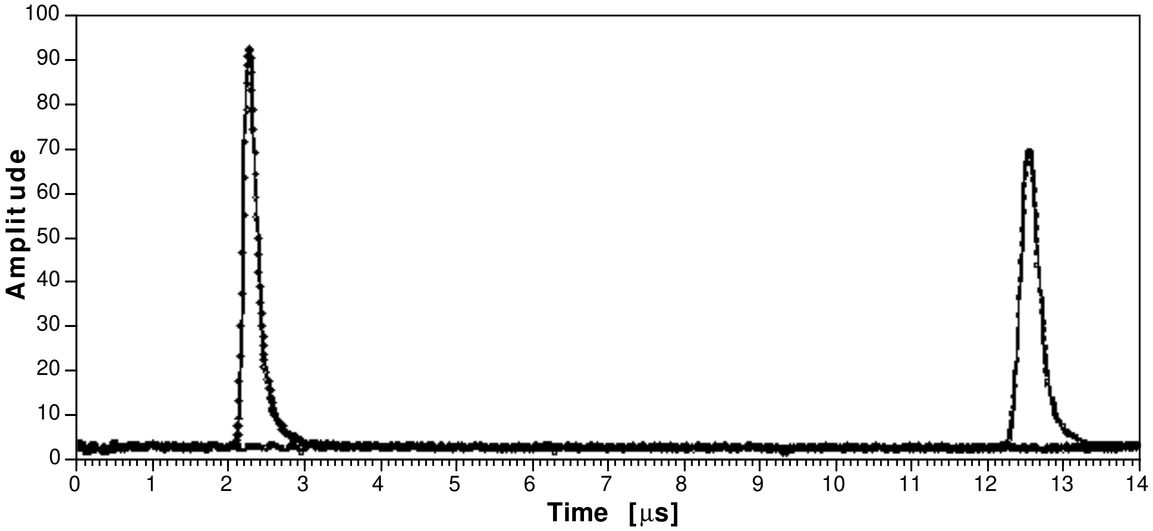

Figure 3.28: Time signal distribution in the pick-up counter.

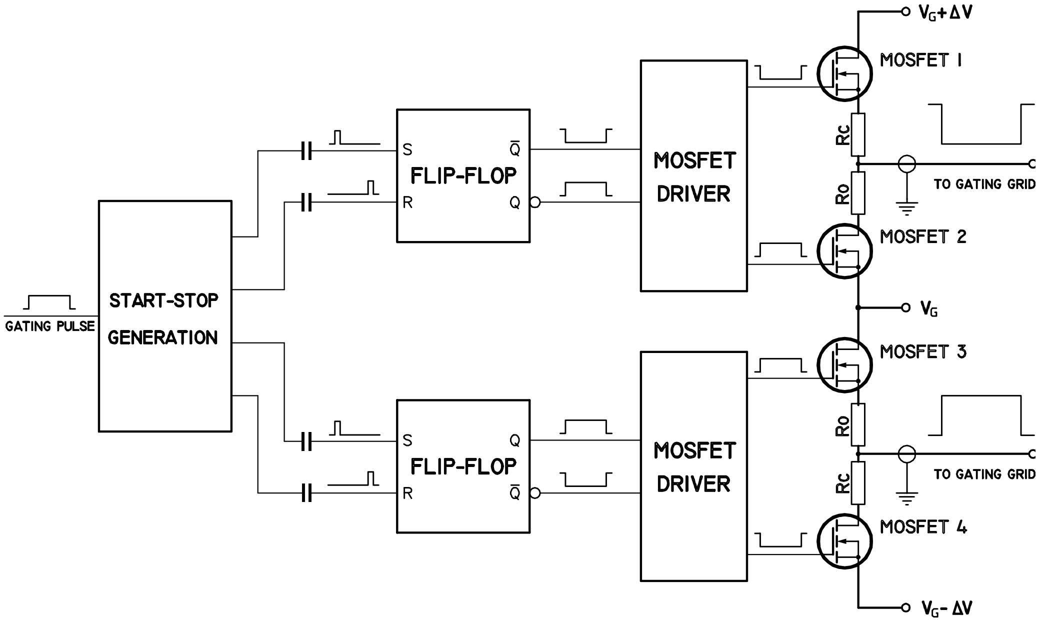

Figure 3.29: Block diagram of the gating circuit.

R_0. It can be larger in order to compensate small difference in opening and closing the FET, which improves the symmetry of the positive and negative pulses. The propagation delay of this pulser can be less than 400 ns so that the TPC gate can be opened within less than 1 μs after the trigger. The minimum pulse duration is determined by the start-stop generation circuit and flip-flop, being less than 1 μs. The maximal pulse duration is unlimited.

Pulsers based on this design have been installed in the TPCs of the NA49 experiment. Results of testing confirm the design performances.

3.9 Laser calibration

The purpose of this system is to:

- allow general tests of the TPC before run time and without special trigger for cosmic rays;
- study systematic effects like distortions with and without magnetic field: since the laser tracks are generated at fixed positions the statistical uncertainty can be made negligible. Laser tracks are not affected by the magnetic field and therefore distortion effects due to inhomogeneous magnetic fields can be deter-

mined;
- determine drift velocities: with a proper distribution of laser tracks over the drift volume, potential problems due to temperature gradients can be studied.

In order to achieve these goals the position of the laser tracks inside the TPC volume has to be known to a precision of better than 1 mm.

For the implementation of the laser system two different options are being considered: one, based on the design of DELPHI is more conventional; the second one, based on the design of STAR, is a more recent development.

The DELPHI design [20] makes use of miniature nitrogen lasers emitting at λ = 337 nm, with a pulse energy of 70 μJ, a pulse length $\sim$0.8 ns (FWHM) and a divergence at the diffraction limit. Because of the large number of wire chambers in ALICE, not every chamber can be equipped with its own laser, like in the DELPHI design. Instead, a configuration is proposed where one inner wire chamber and two outer ones are covered by three laser beams, according to Fig. 3.30. In this way there is always one laser track, parallel to the readout plane, crossing three wire chambers, and two tracks entering the TPC volume with an angle of about 60°. For this arrangement, 24 lasers are needed, generating 72 laser tracks. Since the lasers are mounted directly on the frame of the end-caps, no provision has to be made to control the position and direction of the laser beams interactively to detect possible movements of the TPC relative to the magnet. On the other hand, the coverage of the volume is minimal and only with strongly inclined tracks.

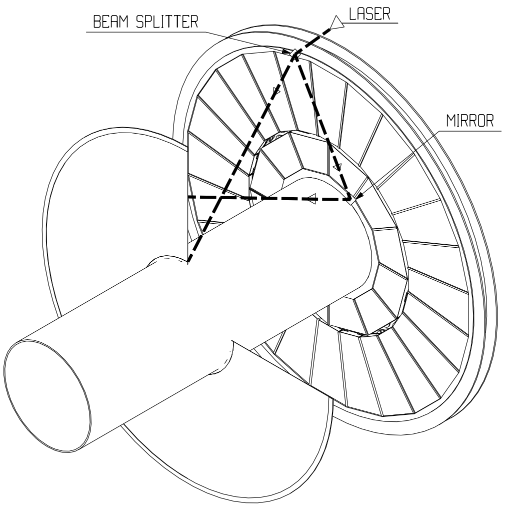

Figure 3.30: Track pattern generated by the nitrogen laser. In total 12 lasers per pad plane are proposed.

To improve on the latter point many more tracks are distributed more evenly over the volume in the STAR design [21]. Two powerful Nd-YAG lasers are used to achieve this, one for each half of the TPC. In order to reach the parallelism necessary for the long beam path, the laser beam is expanded to the rather large diameter of 3 cm. The actual laser tracks are created by a novel design of micro mirrors, with 1 mm diameter, mounted inside the TPC.

The parameters of the commercially available Nd-YAG laser are:

- $\lambda = 266$ nm,
- pulse energy $\sim$ 10 mJ (depends on the model),
- pulse width 10 ns,
- divergence 1.7 $\times$ diffraction limit.

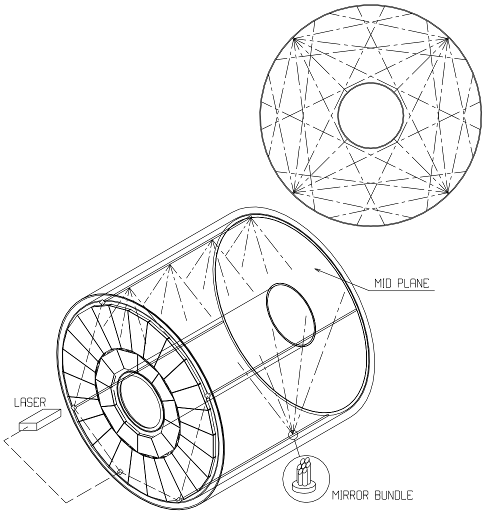

Figure 3.31: Laser arrangement with Nd-YAG laser and micro mirrors shown for one half of the TPC.

A schematic view of the arrangement is shown in Fig. 3.31. The laser beam is first expanded to a diameter of around 3 cm and thereafter transported to the TPC end-cap, after passing active steering and beam diagnostics elements. The beam is split fourfold before entering the sensitive volume, close to the inside wall of the outer field cage, parallel to the beam axis. At four points along the path, bundles of seven glass fibres, with dielectrically coated surfaces cut under 45°, generate a pattern of laser beams. In total, 112 laser tracks are generated per half of the TPC.

Several aspects of this novel laser system are subject to ongoing R&D within STAR:

- production of the 1 mm micro mirrors;
- precise calibration of mirror positioning and direction;
- control and monitoring of the position of the primary laser beam using a well-defined reference line.

This new design seems to provide sufficient coverage of the TPC volume while keeping the investment in lasers at a relatively low level.

3.10 Gas system

The conceptual layout of the gas system is shown in Fig. 3.32 and some parameters are listed in Table 3.5. The proposed system follows the basic ideas of the NA49 TPC gas system but can be made much simpler since only one TPC is involved.

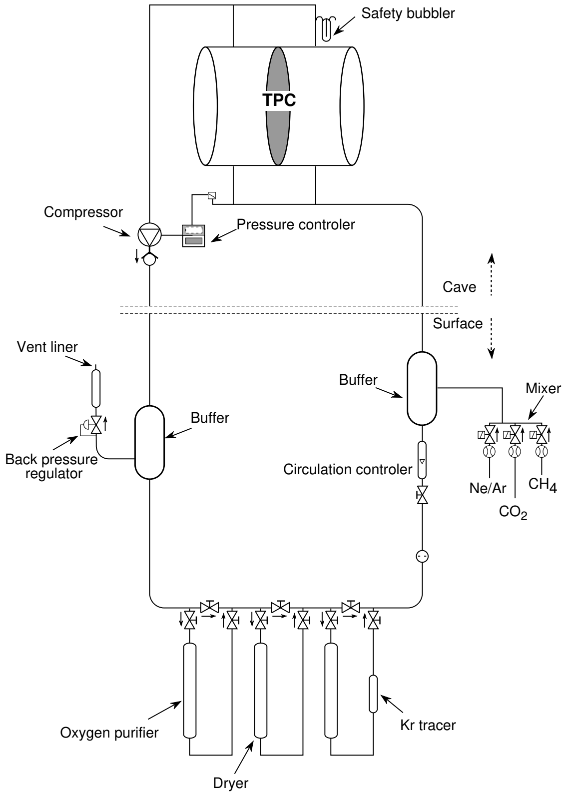

Figure 3.32: Layout of the ALICE TPC gas system.

The gas components, Ar or Ne together with CO_2 and CH_4, will be mixed to the desired gas mixture by mass flow-meter controllers. This can be done with a precision of 0.2 %. The gas is moved in a closed circuit by a compressor under pressure control by sensors and a regulator which changes the speed of the compressor. This recirculation flow is of the order of 20 % vol/h. The constant gas flow passes an oxygen purifier and a dryer to keep the oxygen and water content at a tolerable level (several ppm, see Table 3.5). For calibration purposes (see Section 3.7) the gas flow can be loaded in a bypass with radioactive ^{83}Kr.

During the operation, a certain amount of fresh gas has to be injected via mass flow-meters. Based on NA49 experience, this amount is expected to be of the order of 2 % vol/h (2 m^3/h).

Table 3.5: ALICE TPC gas system

Item	Value
Total volume	84 m^3
Gas	(Ar or Ne)/CO_2/CH_4
Pressure	P_{atm} + (0.5±0.01) mbar
Fresh gas flow	2 % Vol/h
Recirculation flow	20 % Vol/h
Purging flow	20 % Vol/h
O_2 contam.	ca. 3 ppm
H_2O contam.	ca. 3 ppm

This purging gas flow will be released through a back pressure regulator located behind the compressor. A bubbler mounted on the TPC will contribute to a safe operation.

Temperature controls in the inlet and outlet gas tubes, as well as in the TPC gas volume, will provide means for an operation at stable temperature.

In the case of Ne mixtures, one could add a cryosystem to freeze out all the gas components except the Ne [22], in order to reduce the operating costs.

3.11 Expected performance

3.11.1 Space-point accuracy and cluster width

The accuracy of the space-point reconstruction (space-point accuracy) and the cluster width, determining the separation power of two points in space (two-point resolution) have been obtained using a detailed simulation of the clusters generated in the readout module. This includes:

- generation of the space-point where the primary ionization occurred;
- calculation of the energy loss according to the Landau distribution;
- calculation of the total number of ionization electrons;
- calculation of the final position of the electron at the readout chamber, taking diffusion into account;
- calculation of the **E** × **B** effect near the sense wires;
- calculation, electron by electron, of the charge induced on the pad plane and sampling it in time;
- integration of the signal over all electrons and digitization.

All calculations presented in this section have been performed for the Ne/CO_2 (90/10) gas mixture. The parameters used in these simulations, with their numerical values, are listed in Table 3.6. Two-dimensional clusters in

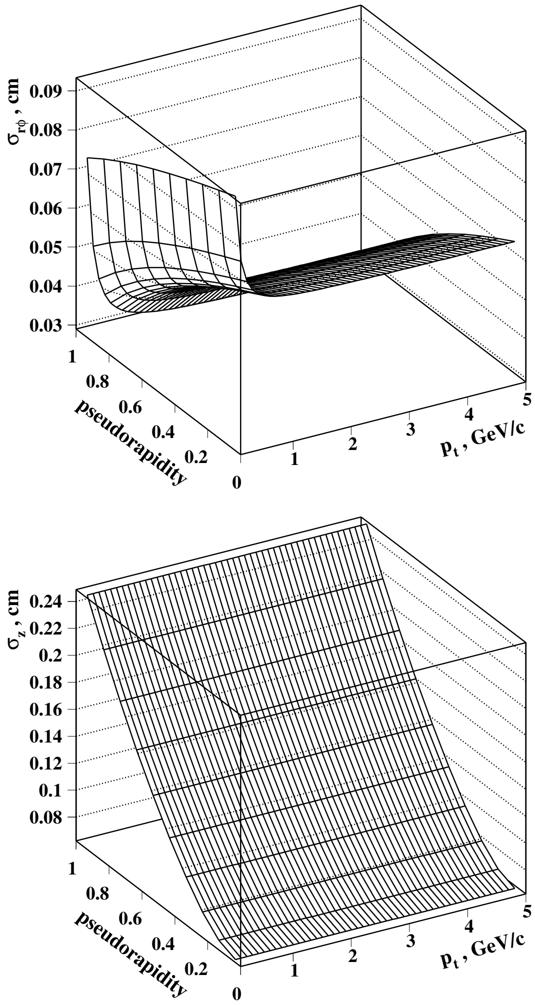

Figure 3.33: Space-point accuracy in $r\varphi$ and z as a function of p_t and η.

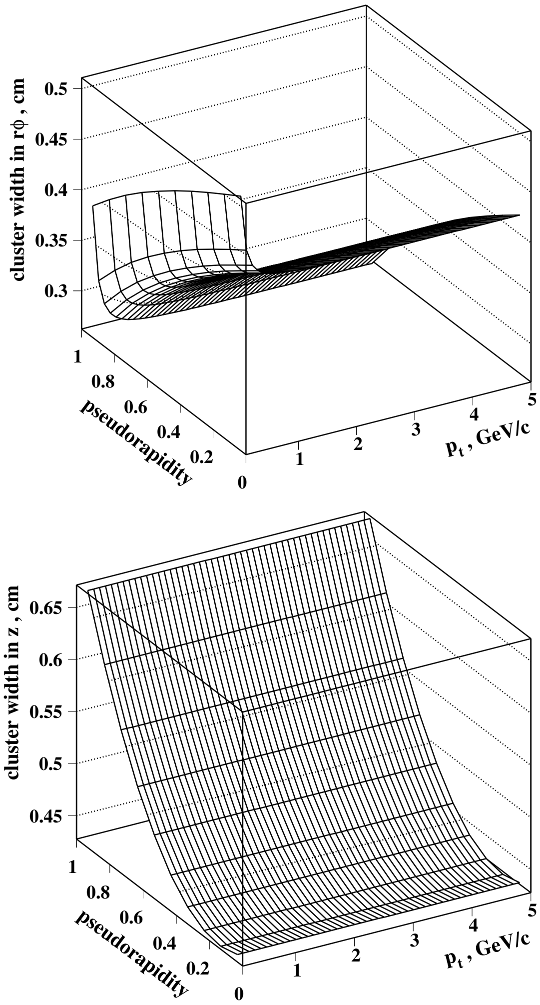

Figure 3.34: Cluster width (r.m.s.) in $r\varphi$ and z as a function of p_t and η.

the pad–time-bucket space have been obtained, which allowed the calculation of the $r\varphi$ and z coordinates of the space-points and thus the determination of the accuracy of their reconstruction and the cluster width.

For further calculations these values have been parametrized as a function of pseudorapidity, η, transverse momentum, p_t, and radius of the pad row, R_{pad}. Plots for $R_{pad} = 175$ cm are shown in Fig. 3.33.

Given this accuracy, a momentum resolution of $\Delta p/p = 1.3\,\%$ for a 500 MeV/c particle is achieved for stand-alone tracking in the TPC. Conservatively, the double-track resolution will correspond to about 2.5 times the cluster width σ, which is shown in Fig. 3.34 as a function of η and p_t.

Table 3.6: Parameter values used in the TPC simulation.

Item	Value
v_{drift}	2.5 cm/μs
$\omega\tau$	0.125
$\sigma_{trans}^{diff} = \sigma_{long}^{diff}$	250 μm/$\sqrt{\text{cm}}$
Pad size	20× 3 mm^2
Output peaking time	240 ns
Sampling time	125 ns
Dynamic range	10 bits
σ_{noise}	1 ADC count
Signal/noise ratio	40 : 1
Electric field	400 V/cm
Magnetic field	0.2 T

This results in a critical distance for two-track resolution of about 1.0 to 1.8 cm in z, the $r\varphi$ resolution being slightly better. Recently, NA49 has developed a method which can help to disentangle even closer tracks [23]. It is based on a cluster unfolding algorithm coupled to global tracking, using the generalized Hough transform. Two Gaussian distributions with the expected width are fitted to the clusters in question, making it possible to resolve hits as close as 1σ. This should significantly improve measurements of a narrow two-particle correlation function.

3.11.2 dE/dx resolution

In the situation of almost straight tracks in the ALICE TPC the dE/dx performance can be judged from what has been observed in the NA35 TPC. Tracks were sampled in an Ar CH_4 (91/9) mixture over a length up to 2.40 m by a maximum of 60 samples, each sample 4 cm long (Fig. 3.35). From these measurements it is concluded that in the case of the ALICE TPC with 1.5 m-long tracks and 75 samples, each sample 2 cm long, the resolution will be 5.5 % for isolated tracks. For tracks in the densely populated region it has been estimated, based on NA35 experience, that dE/dx resolution will be about 6.5–7 %. For Ne-based mixtures [24] the same performance in dE/dx resolution has been measured as in Ar mixtures.

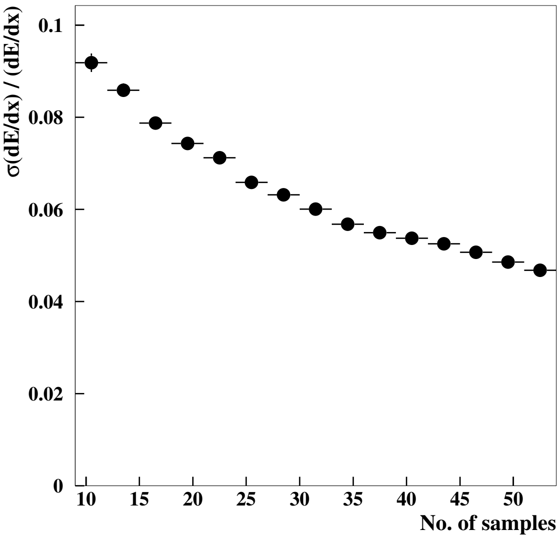

Figure 3.35: dE/dx resolution as a function of the number of samples in the NA35 TPC. Sample length 4 cm, gas Ar/CH_4 (91/9).

3.11.3 Particle identification

The ALICE TPC provides particle identification, in a certain interval of momenta, by dE/dx measurements. For a calculation of the separation power between different particle species, the parametrization of the energy loss proposed by the ALEPH Collaboration [25], fitted to Ar/CH_4 (90/10) data [24, 25], has been used. The relative dE/dx accuracy of 7 %, as discussed in the previous section, was assumed. The results for the ALICE TPC are presented in Fig. 3.36. It can be seen that electrons with a momentum below ~ 5 GeV/c can be identified by the dE/dx measurement in the TPC with significance better than 3σ. Particle identification, for hadrons, is restricted to the low-momentum range, 100–800 MeV/c.

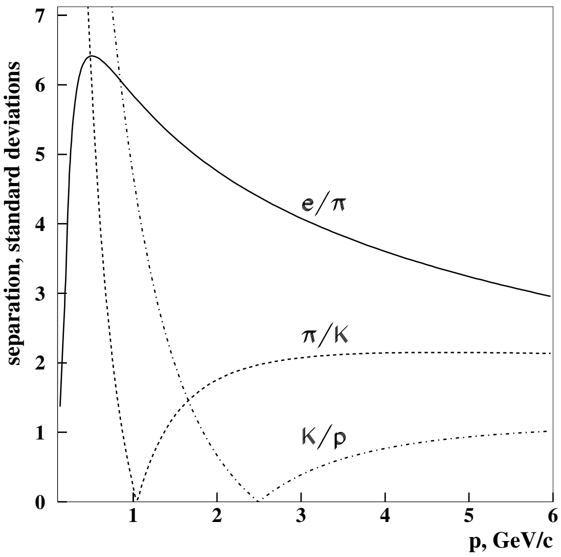

Figure 3.36: Particle separation in the ALICE TPC.

3.11.4 TPC granularity

The TPC granularity in terms of the pad width and sampling frequency has been discussed in Section 3.2. Here the granularity in the radial direction is discussed and the arguments for the choice of the pad length are given. In Fig. 3.37 the space-point resolution and the cluster width in the z (drift) direction are shown for different pad lengths for η close to 0.9. It is seen that both σ_z and the cluster width get significantly worse when increasing the pad length. For low p_t (curved) tracks, a very similar behaviour is observed in the $r\varphi$ plane, for the resolution and cluster width.

In order to measure very narrow two-particle correlation functions, such as expected in Pb–Pb collisions at LHC energy, the best possible double-track resolution is required. In Fig. 3.38 the double-track resolution of the TPC is shown as a function of $Q_{\rm inv} = \sqrt{-(p_1 - p_2)^2}$,

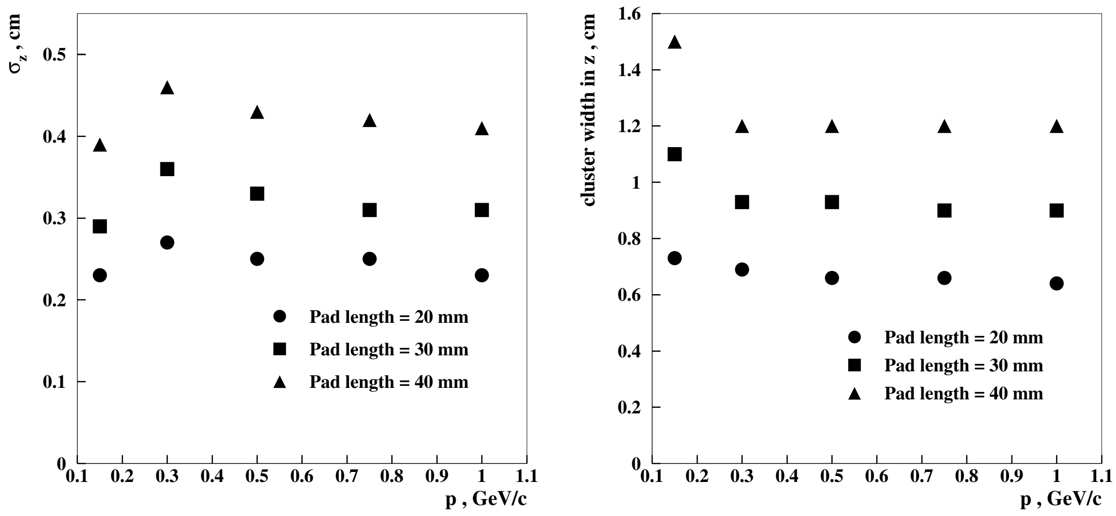

Figure 3.37: Space-point accuracy and cluster width (r.m.s.) in z, as a function of total momentum, for different pad lengths ($\eta \approx 0.9$).

where $p_{1,2}$ are 4-momenta of pions. For low $Q_{\rm inv}$, a good two-point resolution is seen to be necessary, ~ 1 cm both in $r\varphi$ and z. This can only be achieved with short pads. Another argument in favour of short pads is that good

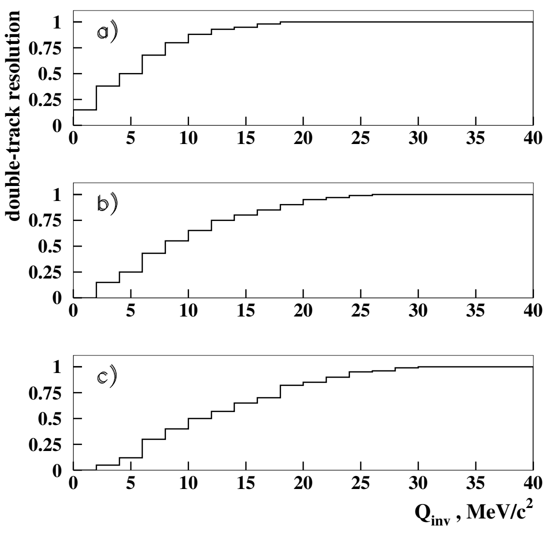

Figure 3.38: Double-track resolution (in the TPC alone) as a function of $Q_{\rm inv}$, for two-point resolution of (a) 1 cm, (b) 2 cm, (c) 3 cm, both in $r\varphi$ and z directions.

dE/dx resolution, and thus good particle separation, requires a large number of samples per track (see Fig. 3.35). For a given geometry, longer pads mean less samples per track and worse particle separation. These considerations lead to the conclusion that the pad length of 20 mm would be the preferable choice.

3.11.5 Space-charge effects

There are two sources of positive charge in the drift volume of the TPC: the ionization caused by the produced particles and the positive-ion feedback from the proportional readout chamber. The latter is reduced by the implementation of the gating grid, while the first one leads to a positive charge build-up in the drift volume. The reason is that the drift velocity of the positive ions is lower by 3–4 orders of magnitude than that for electrons, which results in the permanent presence of the positive charge in the drift volume of the TPC. The extra electric field due to this charge affects the drift of electrons towards the readout chambers.

Distortions caused by the space charge have been calculated through the following steps:

- the charge density $\rho(R,z)$, where R is the radial coordinate of the TPC and z is the coordinate along the drift, has been calculated using the HIJING event generator, tuned to $\mathrm{d}N_{\rm ch}/\mathrm{d}y = 8000$ per unit of rapidity at mid-rapidity, for central Pb–Pb events; the multiplicity of minimum bias events, the major contribution to the creation of the space charge, has been scaled by an appropriate factor;
- the build-up of the positive charge in the drift volume has been calculated for the luminosity $\mathcal{L} = 10^{27}$ cm^{-2}s^{-1}, assuming a drift velocity of positive ions of 540 cm/s;
- the potential $\Phi(R,z)$, and thus the extra electric field, has been found by solving the Poisson equation

$$\nabla^2\Phi(R,z) = \rho(R,z);$$

- distortions of a single electron in radial (ΔR) and azimuthal ($\Delta\xi$) directions have been calculated.

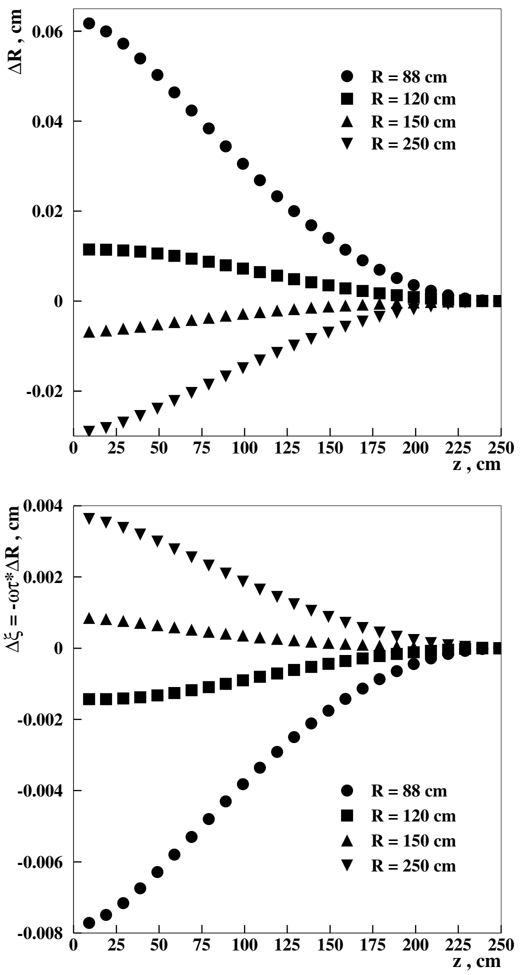

Figure 3.39: Radial (ΔR) and azimuthal ($\Delta\xi$) distortions due to the space charge.

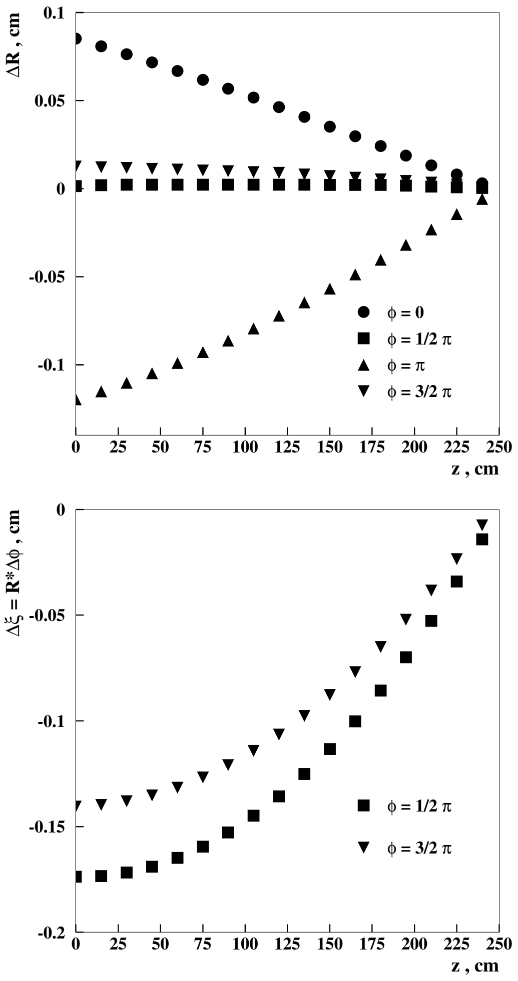

Figure 3.40: $\mathbf{E} \times \mathbf{B}$ distortions in radial (ΔR) and azimuthal ($\Delta\xi$) directions for R_{pad} = 250 cm. The azimuthal distortions for $\varphi = 0$ and $\varphi = \pi$ are very close to those for $\varphi = 3/2\,\pi$ and are not shown here.

In Fig. 3.39 these distortions are shown, as a function of the z-coordinate, for different radial positions. One should notice that, if the track goes exactly along the pad, the shift of the electron in radial direction does not lead to any distortions, while in the general case the distortions due to such a shift are proportional to $\Delta R\ \tan\alpha$, where α is the angle between the track and the pad axis. It can be shown that $\tan\alpha$ is the main factor determining $\sigma_{r\varphi}$. Therefore, the geometry of pads has been chosen to minimize this angle. This, together with the fact that ΔR does not exceed 700 μm (see Fig. 3.39), leads to the conclusion that distortions due to the radial component of the electric field should stay within an acceptable limit of $\sim 150\,\mu$m. In the azimuthal direction, the distortions are well within the overall resolution.

The space-charge effects in argon are stronger than in neon. The specific ionization in argon is twice that for neon (see Table 3.2), which results in two times larger charge density. In addition, the drift velocity of ions in argon is about two times smaller than in neon, which gives two times longer clearing time. Altogether, the space-charge distortions in argon-based mixtures are expected to be about four times larger than in neon-based ones.

3.11.6 E $\times$ B distortions

An electron created by ionization drifts in the electric (**E**) and magnetic (**B**) fields with velocity

$$\mathbf{v}_{\text{drift}} = \frac{\mu}{1+(\omega\tau)^2} \times \left(\mathbf{E} + \omega\tau\,\frac{\mathbf{E}\times\mathbf{B}}{B} + (\omega\tau)^2\,\frac{(\mathbf{E}\cdot\mathbf{E})\cdot\mathbf{B}}{B^2}\right),$$

where μ is the electron mobility in the gas and $\omega\tau$ is the tangent of the Lorentz angle. For both fields uniform and parallel, the drift velocity is:

$$\mathbf{v}_{\mathrm{drift}} = \mu \mathbf{E}.$$

In the L3 magnet the field has a radial component and thus the drift velocity acquires radial and azimuthal components, which leads to distortions in all directions. Radial and azimuthal distortions due to the $\mathbf{E} \times \mathbf{B}$ effect are displayed in Fig. 3.40 as a function of the z-coordinate, for $R_{\mathrm{pad}} = 250$ cm (where inhomogeneities of the magnetic field are the largest) and for different values of φ. Distortions in z are negligible and thus not shown here. The $\mathbf{E} \times \mathbf{B}$ distortions are larger than those caused by the space charge. However, the behaviour of a charged particle in the electric and magnetic field is well known and can be corrected for (as done in NA36 and NA49). The above calculations have been performed using the magnetic field map of the L3 magnet, together with the dipole of the muon arm and the iron plugs on both sides of the L3 solenoid (see Chapter 6, magnet configuration 'CC'). Note also that the ALICE TPC is placed 30 cm above the L3 magnet axis. The radial component of the magnetic field for this configuration, at the edge of the TPC, is ~ 20 G. Without additional iron plugs, the radial distortions are very similar to those shown in Fig. 3.40, whereas the azimuthal ones are larger by a factor ~ 2.

4 Particle Identification

4.1 Requirements for the PID

Particle identification (PID) over a large part of the phase space and for many different particles is an important design feature of ALICE. It is of crucial importance for a number of signals (e.g. flavour composition and chemical equilibrium, lepton pairs, strangelets) and extremely useful for most others (e.g. momentum spectra, HBT, charm, jet quenching). We have two detector systems dedicated exclusively to PID, a TOF array optimized for large acceptance and average momenta and a small system specialized on higher momenta. In addition, both of our tracking systems (ITS and TPC) provide energy loss information which is vital for electron ID and hadrons of low momentum.

The PID requirements can be divided into three categories:

Event-by-event hadron identification
In order to measure the particle composition and momentum distribution of stable hadrons (π, K, p, light nuclei) in single events, the momentum range for PID can be limited to a few times the average p_t. Even at the highest multiplicities, the statistics in single events are sufficient only for transverse momenta up to about 2 GeV/c. A large area coverage is however required, as the significance of most single-event observables depends strongly on the number of particles within the acceptance. The relevant PID techniques, which cover the full central acceptance of ALICE ($|\eta| < 0.9$), are energy loss measurement (dE/dx) in ITS and TPC in the non-relativistic region, and TOF for higher momenta.

Inclusive hadron identification
The momentum region above 1.5–2 GeV/c, which is accessible only to inclusive measurements, is covered by a specialized PID detector, the HMPID. It will extend most hadronic observables from the 'average hadrons' measured event-by-event into the semi-hard regime (e.g. particle composition and momentum spectra, flow, HBT). Our current choice is a RICH detector covering about 5 % of the central acceptance.

Electron identification
In general, electron identification poses requirements which are very different from those of hadron identification. While all types of stable hadrons have a similar abundance (π:K:p $\approx$ 3:1:1 for $p_t > 1$ GeV/c), direct electrons are very rare (e/$\pi \approx 10^{-3}$–10^{-4} at $p_t > 0.5$ GeV/c). Because of a number of constraints, in particular concerning available space, material thickness and cost, ALICE has no dedicated detector for electron ID. We mainly rely on dE/dx in the TPC, which provides sufficient pion rejection up to about 2.5 GeV/c. Electron ID is complemented at low momenta ($p < 140$ MeV/c) by dE/dx in the ITS and at intermediate momenta (around 200 MeV/c) by TOF.

4.1.1 Requirements for large area TOF

The PID capabilities of ITS and TPC have been described in previous chapters. The design specifications for the large area TOF are as follows:

- For hadrons, it should cover the momentum range from about 0.5 GeV/c (upper limit for dE/dx in the TPC for K/π separation) to about 2 GeV/c (statistics limit in single events). Our target is a system resolution of 100 ps, i.e. including all other sources of timing errors in addition to the detector resolution itself. This resolution would guarantee a 3σ K/π-separation up to $p_{3\sigma} = 2.1$ GeV/c (TOF at $r = 3.5$ m). A resolution of 150 ps would still be acceptable ($p_{3\sigma} = 1.7$ GeV/c), whereas 200 ps or worse ($p_{3\sigma} < 1.4$ GeV/c) would leave a substantial fraction of the momentum spectrum unmeasured.
- For electrons, it has to cover the region where dE/dx does not separate electrons and pions ($140 < p < 200$ MeV/c). Whereas timing resolution is of less importance in this case, the momentum cut-off imposed by the magnetic field limits the maximum distance of the TOF array to ~ 3.5 m.

The overall inefficiency (geometrical coverage, intrinsic response or pile-up) should be small, i.e. of the order of 20 % or less, for both the electron and the hadron analysis. Therefore we want to keep the occupancy below 10 % at the highest anticipated multiplicity, implying more than 10^5 independent channels.

The technological choices for the barrel are Pestov counters as the baseline design and Parallel Plate Counters (PPC) as the fall-back solution. The present status of the Pestov detectors as well as the uncertainties in the overall performance of a 100 m^2 detector make it essential to keep an alternative design of lower resolution but still acceptable for the physics goals. A final choice between the two technologies will be made before the end of 1996, based on test results and operational experience gained with large prototypes.

4.1.2 Requirements for the HMPID

The requirements and optimization for the HMPID are different from those for the barrel PID. It should identify hadrons starting at around 1 GeV and cover the largest possible momentum range. As it will measure inclusive

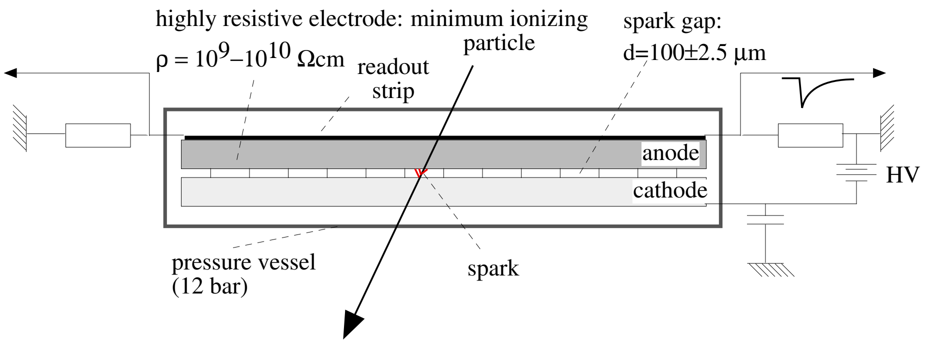

Figure 4.1: Schematic layout of the Pestov spark counter.

spectra only, efficiency and geometrical acceptance are less crucial than performance. In addition, it can be placed a long way from the vertex, where particle densities are smaller.

Our current choice for the HMPID is a proximity focusing RICH detector, optimized for high momentum hadron identification. It can provide 3σ K/π separation up to 3.4 GeV/c.

If we can show within the next two years that a system resolution of < 100 ps is feasible, a high performance TOF array could become competitive.

4.2 TOF by Pestov spark counters

4.2.1 Introduction

The Pestov spark counter [1–6], schematically shown in Fig. 4.1, is a single-gap, gaseous parallel-plate detector working in the streamer/spark mode. It makes use of a semi-conductive anode and of a special gas mixture that keeps the discharge local, thus increasing the rate capability. The time resolution improves with decreasing gap size. For the ALICE TOF a gap of 100 μm between the electrodes will give sufficient pulse amplitude (necessary for the front-end readout electronics) and the desired time resolution. With this configuration, a resolution as good as 25 ps (r.m.s.) is possible at twice the threshold voltage (~ 6 kV). The counter operates at a pressure of 12 bar, which yields 4–5 primary electrons from a minimum ionizing particle (MIP). The spark discharges only a limited area of the semi-conductive anode (~ 1–2 mm^2), while the remaining area of the anode remains sensitive. The counter is read out via strip lines, which allows a two-dimensional reconstruction of the spark location. The longitudinal position is deduced from the difference between the signal arrival times at the two ends of the strip, while the transverse position is obtained from the centre of gravity of the strips exited.

The Pestov spark counter features a number of intrinsic properties which make this type of detector an attractive choice for a large-scale TOF system at high particle densities:

- time resolution up to 25 ps (depending on operating conditions);
- rate capability better than 2 kHz/cm^2 (the counting rate in ALICE is $<$50 Hz/cm^2 for Pb–Pb and < 2 Hz/cm^2 in pp);
- lifetime larger than 10^9–10^{10} sparks/cm^2 (corresponding to more than 20 years running);
- an overall thickness of about 10 % of a radiation length per module;
- an intrinsic efficiency better than 96 %;
- a sensitive area only 100 μm deep, i.e. no increase in multiplicity due to inclined tracks;
- a signal output large enough to work without preamplifiers.

However, no experience exists with large-scale systems or long-term operation, and the requirements in terms of tolerances and material properties are very high.

4.2.2 Design considerations

Strip versus pad readout

In the ALICE LoI [7], a pad readout geometry was proposed in order to reduce the number of electronic channels for a given occupancy (particles per detector cell) compared to a strip readout. With strip readout, on average three adjacent strips will respond to each particle hit, therefore each ‘detector cell’ requires six electronic channels (three on either side). A readout based on individual large pads would require only one channel per cell. For the pad sizes under consideration (2.5×4 cm^2) no substantial degradation of time resolution would result from geometrical signal transit times alone (< 50 ps).

The signal propagation and cross-talk within the detector has been investigated with the electromagnetic simulation program MAFIA [8]. For the pad geometry, the MAFIA simulations yield substantial cross-talk between adjacent pads, i.e. for each centrally excited pad the sum of all adjacent cells reaches at least 50 % of the central amplitude. Furthermore, the signal shape and rise time strongly depend on the location of the spark, as the same signal can travel along different paths to the readout chip.

Both the simulation and experimental tests with a strip readout show negligible cross-talk (beyond charge sharing) and loss of signal rise time. In addition, the position of a spark can be detected accurately ($\leq 300\,\mu$m) in one coordinate by charge sharing and in the other by the difference between the signal arrival times at the two ends of each strip.

Based on these simulations and tests, we have opted for a strip readout.

Granularity

The charge sharing between channels is Gaussian and on average three strips will respond for a pitch of 2.54 mm. With a double-threshold discriminator (see Section 4.4), the timing resolution has been measured to be approximately the same for all three strips. Therefore two particles can be separated even if they hit adjacent channels, i.e. the average transverse cell size is two strips. With a strip length of 30 cm and a strip width of 2.54 mm, the total cell size is 15.2 cm^2, where a 'cell' is defined as the area within which double hits cannot be resolved. In central Pb–Pb collisions ($\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$), with 75 particles per m^2 at a radius of 3.4 m, the double-hit probability is $\sim 10\,\%$. The TOF barrel will consist of $\sim 95\,000$ cells with 380 000 electronic readout channels (each cell consisting of two strips read out at both ends).

The length of an individual counter will be 120 cm. Shorter counters would increase the number of connections and material due to flanges, whereas longer counters will pose problems for standard industrial equipment (e.g. for glass polishing, multilayer board production or sputtering).

4.2.3 Geometric efficiency

The geometric efficiency is given by the dead areas of the counter itself and gaps in the barrel structure where particles may escape. Each individual counter has dead areas around the spacers (about 1 % of the area) and along the edges of the electrodes. The latter are compensated for by overlapping adjacent counters. The decrease in efficiency due to the gap between detector sub-modules (see Figs. 4.4 and 4.5 on page 62) will be $< 2\,\%$.

In our current design, the TOF counters do not cover an area of ~ 10 m^2 in front of PHOS (i.e. $< 8\,\%$ of the central acceptance) in order to limit the amount of material in front of the electromagnetic calorimeter. However, we are investigating whether special modules with slightly lower geometrical efficiency, but $< 10\,\%$ average radiation length, could be introduced in this area.

4.2.4 The TOF detector module

Mechanical construction

The basic Pestov detector module is shown in Fig. 4.2. It consists of two electrodes of 1.2 m length, 40.68 mm width and 2 mm thickness. The gap between the electrodes is kept at 100 μm by small spacers. The two electrodes are embedded in a plastic support structure with integrated gas channels which maintain a constant gas flow through the gap. On top of the semi-conductive anode is a readout board with strip lines, which picks up the signal and transports it to the front-end electronics mounted on the back of the board. The pressure vessel, which houses the complete counter including readout, is made out of an aluminum tube of 50 mm diameter and has a wall thickness of 1 mm. A cap is welded onto the tube on one side; the other side is connected to a large, multi-tube flange. Several of these tubes have been tested at pressures of up to 90 bar without rupturing.

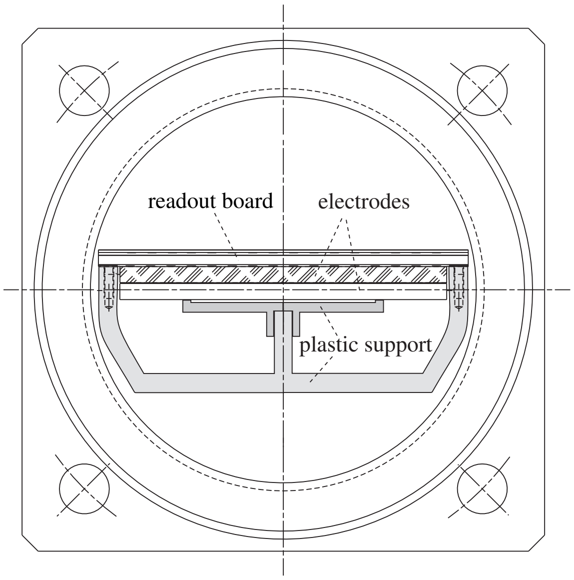

Figure 4.2: Cross-section of a single detector module.

Counter assembly

The anode is made from semi-conductive glass (resistivity $\sim 10^9$ Ωcm) and optically polished on both sides. The edges are specially shaped ($y = x^4$) to compensate for the enhanced field due to the sharp edge of the opposite electrode. Several technologies and materials have been tested for producing a high-quality cathode surface. The most

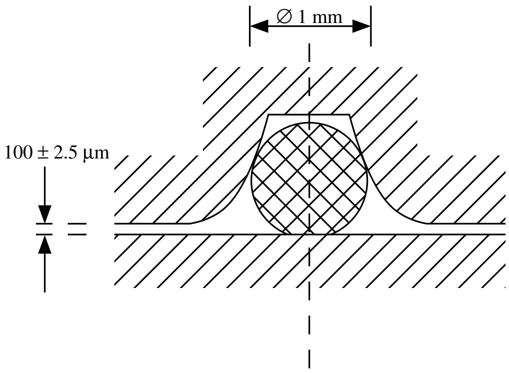

Figure 4.3: Cross-section of a spacer which maintains the gap of 100 μm between the electrodes.

promising seems to be sputtering of a copper film on standard float glass.

A grid ($d = 22$ mm) of spacers (see Fig. 4.3) is manufactured into the semi-conductive anode in order to keep the sagitta due to the electrostatic pressure (6.4 N/cm^2 at 6 kV) below ~ 0.2 % of the gap distance. These spacers, glass balls of 1 mm diameter, are glued into depressions of 0.9 mm depth and are adjusted by polishing to a gap width of 100 μm. The edges of the depressions are also x^4-shaped to avoid electrical breakdown [9].

The plastic support structure, which can be produced by extrusion, is shown in Fig. 4.2. The cathode is placed on the upper, thinner part of the support (0.4 mm plastic), which acts as a spring and ensures gas tightness by pressing the anode and cathode against the readout board.

The readout board, which is glued into the end flange, acts as a gas seal between the gas channels and forces the gas flow through the gap. At the same time it gives the system, together with the plastic support, mechanical stability. This board contains both the strip lines to pick up the signal from a spark as well as the readout lines which carry the signal out of the pressure vessel. The FE electronics will be housed inside the pressure tube, thus avoiding high-frequency feed-throughs.

The radiation thickness of a single detector module (see Table 4.1) is 9.7 % of a radiation length, excluding the end flange region. A particle will pass about 5 % of a radiation length before reaching the spark gap.

4.2.5 Sub-module and space frame

A TOF sub-module consists of 86 tubes arranged orthogonally to the beam axis and covering an area of 125×350 cm^2, corresponding to an azimuthal section of $\varphi = 20°$ and a polar coverage of $\theta = 45°$, i.e. 36 sub-modules approximate a barrel of radius 3.4 m.

All 86 tubes of one sub-module (Fig. 4.4) are connected to a common flange, which houses the gas lines to the individual tubes. The electrodes in the tubes are tilted such that a projective geometry is obtained, ensuring almost complete geometrical coverage also for inclined tracks. The weight of a single Pestov detector will be 1.5 kg; a sub-module of 86 detectors including flange will weigh about 250 kg.

Table 4.1: Contributions to the radiation thickness of a single detector module

Component	Thickness (mm)	$X/X_0(\%)$
Cathode (glass)	2.0	1.6
Anode (glass)	2.0	2.2
Support (plast.)	$\langle 5 \rangle$	1.5
Tube (Al)	$\langle 3 \rangle$	3.4
Board (G10)	2.0	1.0
Total		9.7

The individual sub-modules are supported by a space frame of 7 m length consisting of 4 rings of 4.31 m outer and 3.95 m inner radius (approximated by a polygon with 18 corners, see Fig. 4.5). The distance between the two inner rings is 2.4 m in order not to shadow the PHOS calorimeter placed behind. Each of the TOF sub-modules can be loaded individually for installation (or repair) onto a rail system.

4.2.6 The gas system

The standard gas mixture used at present consists of 76.9 % argon, 20.0 % isobutane, 2.5 % ethylene and 0.6 % divinyl. Due to the low partial pressure, isobutane and divinyl gas cannot be premixed and compressed for transportation. Therefore, the gas will be mixed at low pressure (≤ 2.5 bar) on-site at the surface area, as shown in Fig. 4.6. The premixed gas is compressed in two steps to 12 bar and pumped to the sub-modules. Each sub-module is equipped with its own circulation pump, and pressure, flow and temperature control. The total quantity of flammable gas in the underground area is ~ 50 kg.

4.2.7 Mounting of front-end electronics

We plan to group the readout electronics for 16 channels directly on the multilayer readout board inside the pressure tube. Each detector has four sections with 16 strip lines of 30 cm length. All strip lines are readout at both ends, i.e. a single detector will have 8 groups of 16 readout channels. Each of these subgroups includes discriminators, TDCs and QDCs and the necessary auxiliary electronics for readout. Because of the large pulse height of a Pestov counter, no amplification is required. To match the high resolution of the counters the electronic time resolution (jitter, walk) has to better than 50 ps. The electronics, which is largely the same for all TOF options, is described in Section 4.4. The electronics can be cooled by passive

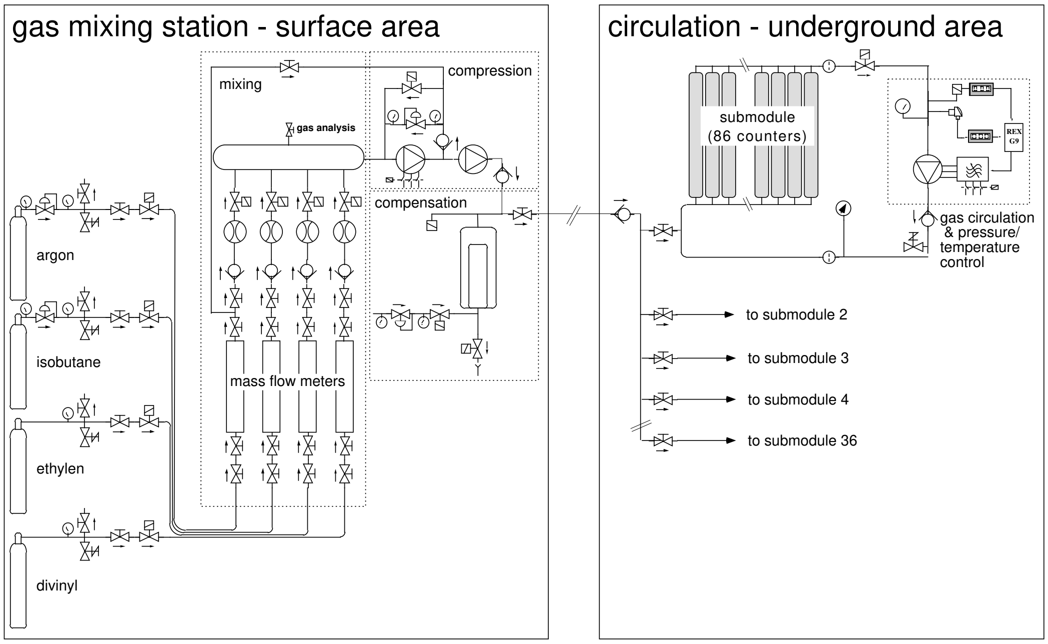

Figure 4.6: Layout of the gas mixing stations in the surface and underground areas.

heat exchange with the surrounding Al pressure tube. Due to the high pressure of 12 bar, cooling by convection is sufficient to remove at least 50 mW/channel. The tube itself will be cooled by water.

4.2.8 Counter conditioning

An initial conditioning procedure must be applied to each spark counter after assembly. It involves the irradiation of the counters with a high-intensity radioactive source at slowly increasing voltage. During this procedure the surface conditions on the counter electrodes change and the location of the discharge improves. The process can be easily controlled by the width of the measured charge distribution. After $\sim 10^6$ sparks/cm^2 on the counter surface (about 3 days of continuous irradiation), the amplitude distribution stabilizes and the counter is ready for use. At this stage,the quality of the counter assembly can also be verified: a counter is acceptable if its background rate is at the level of cosmic rays. The counter characteristics obtained after the conditioning are preserved during shut-down periods of at least 6 months.

4.2.9 Calibration procedure

The plateau curve gives information about the spark counter status. To obtain the plateau curve for a 1 m long counter, a special trigger, taken from the cathode, can be used. A special calibration of amplitude channels is not foreseen since an amplitude-dependent walk correction is not possible in a Pestov counter and ultimate transverse position resolution (via charge sharing) is not needed in ALICE. Pedestals and rough gain corrections can easily be taken in parallel with data taking. A periodic calibration of the TDCs has to be performed to compensate drifts due to variation of temperature and voltage. This will be done periodically either with cosmic rays or with the actual data and takes a maximum of a few hours.

4.2.10 R&D results

This type of counter has been pioneered by Yu. Pestov at BNIP Novosibirsk [1–6]. Over the last two years several prototypes have been built in collaboration with Yu. Pestov in order to develop a simplified design of similar performance which is suitable for mass production and large-scale applications. Our prototypes have as an essential new feature a strip line readout enabling a good transverse position resolution in addition to the excellent timing properties.

A first prototype, with an active size of 40.68×540.8 mm^2, reached satisfactory results in terms of time resolution ($\sigma_{\text{time}} \lesssim 70$ ps) [10]. Its design, however, is not suitable for mass production. A new prototype, which has been tested recently, avoids coupling of the strip lines on the back of the anode to the readout board via springs. Instead, the signal is picked up directly by a multilayer board, which itself has strip lines printed on it. One side of the strip line is fed back through a second layer of the board, i.e. both ends of the strip lines are read out at the same side of the counter. A newly developed double-threshold discriminator is mounted directly onto the mul-

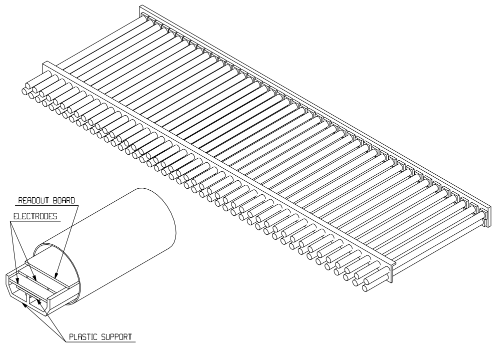

Figure 4.4: Schematic view of a sub-module consisting of a common flange with integrated gas distribution system connected to 86 individual detectors.

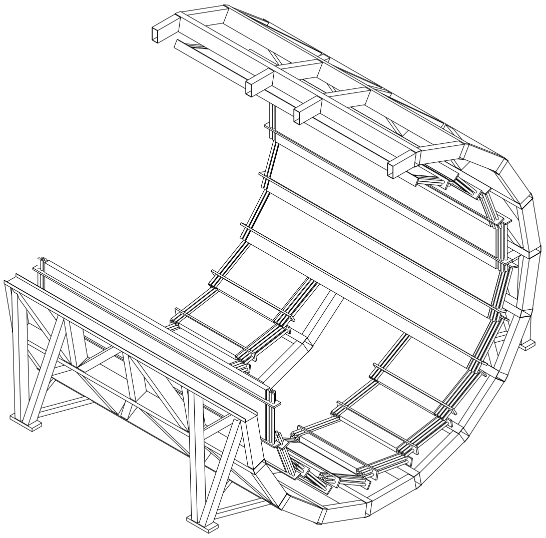

Figure 4.5: Space frame support structure for the TOF counters. The area in front of the PHOS is left free.

tilayer board, thus minimizing the degradation of signal rise time due to cables. This new design already includes many features, although at a prototype level, which are envisaged for the TOF barrel as described in Section 4.2.4.

Position resolution

The counter provides position resolution in two dimensions. The position resolution $\sigma_{\rm long}$ of a discharge in the strip direction, i.e. longitudinally, is limited by the resolution of the electronics, and here mainly by the binning of the time digitizer. The walk of the discriminator is of less importance, since the signals seen at the two ends of a strip come from the same spark and do not fluctuate against each other. For a TDC with 50 ps binning, as used in our tests, and a signal velocity of 14 cm/ns, one can expect $\sigma_{\rm long} \approx 2$ mm as a lower limit. A value of $\sigma_{\rm long} \leq 2.5$ mm has indeed been reached (see Fig. 4.7).

The transverse position resolution $\sigma_{\rm trans}$ is based on charge sharing between adjacent strips (centre of gravity method). With a 2 mm anode and 2.54 mm strips, the signal will be distributed on average over 2–3 channels. The resolution obtained from the coincidence of two spark counters is $\sigma_{\rm trans} \approx 320\,\mu$m, as shown in Fig. 4.8.

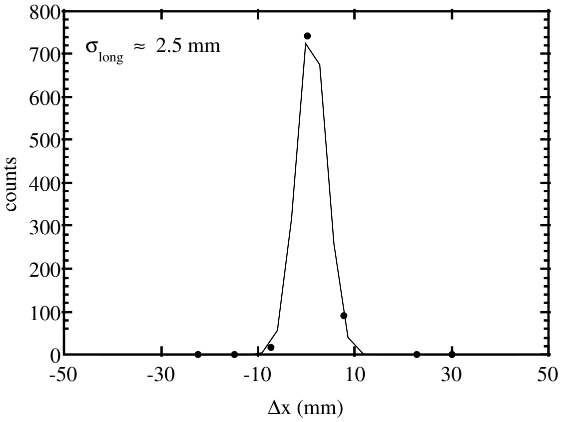

Figure 4.7: Longitudinal position resolution, obtained by measuring signal arrival times in a TDC with 50 ps binning.

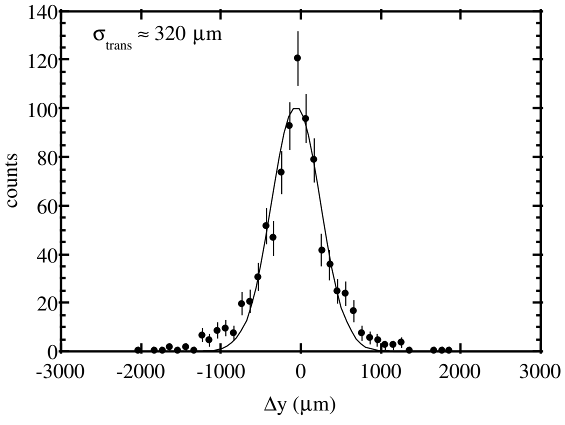

Figure 4.8: Transverse position resolution in the Pestov counter.

Efficiency

The efficiency of the Pestov spark counters was measured by comparing the counting rate of two detectors in a test beam. At 4.9 kV, the intrinsic detection efficiency is 96 %, which is in agreement with previous results from Novosibirsk [11].

Time resolution

With first prototypes we reached a time resolution of $\sigma_{\text{time}} \approx 70$ ps at 4.4 kV. These results were limited by the available electronics, in particular by the walk of the standard leading-edge discriminators used in this test. With a new type of discriminator (described in Section 4.4), which was specially developed for the Pestov counters, we now regularly reach resolutions of the order of 50 ps or better. One example with $\sigma_{\text{time}} \approx 40$ ps (FWHM = 133 ps for two counters in coincidence) is shown in Fig. 4.9, where the counters were operated at 4.9 kV. This performance compares very well with other TOF technologies and easily surpasses our target specification of < 100 ps.

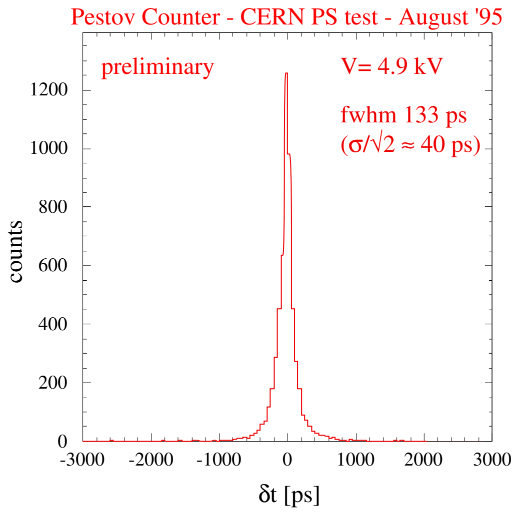

Figure 4.9: Time coincidence spectrum between two Pestov counters operated at 4.9 kV.

Ageing tests

In order to assess the ageing properties of Pestov counters, we first made a measurement of the long-term behaviour over a test period of ~90 days. The HV was raised slowly over one week from the threshold of 3 kV to the onset of the plateau at 4 kV, where it was kept constant for about 1 month and then increased to 5 kV. At 4 kV and below the counter was permanently exposed to a strong γ-source for conditioning. At 5 kV the source was removed and the counter was operated for one month.

Figure 4.10 shows the plateau curve of the counter. The filled circles indicate the counting rate with γ-source irradiation as a function of voltage, the squares show the background rate without source. A slight increase of background discharges sets in at 4.3 kV. After subtraction of these background discharges a saturation of the rate, i.e. a plateau, can be seen, starting at around 4 kV (open circles).

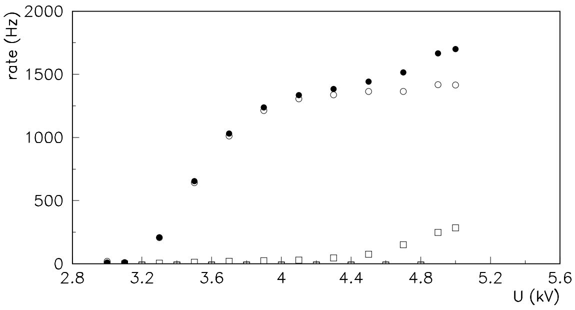

Figure 4.10: Rate from a γ-source (closed circles) as a function of high voltage ('plateau curve'). The open squares show the background obtained without source, the open circles the plateau curve after background subtraction.

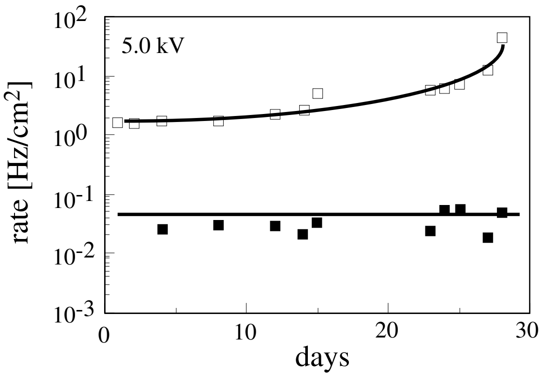

Figure 4.11: Background rate of the counter at 5 kV as a function of time (solid squares: good area, open squares: area with surface defects).

Figure 4.11 shows the background rate at 5 kV over a period of 1 month. The background rate slowly increased with time from 0.3 to 1 kHz and then suddenly to 8 kHz.

The cathode used for this test was produced via direct evaporation, i.e. the substrate was exposed directly to copper vapour. It appears that direct evaporation leads to a small number of defects, i.e. copper pins deposited on the surface, possibly by 'micro-explosions' in the heated copper container. Using the two-dimensional position resolution of the detector to reconstruct the location of background discharges, we were able to investigate both the stability in 'good' areas, i.e. with no defects, as well as at the location of surface defects.

The initial slow increase of the background rate with time is explained by the formation of new spots close to an original 'bad' spot at a copper pin. Approximately 10 % of the counter area in this test was affected. In the other regions, the rate remained at a very low level of about 0.05 Hz/cm^2 during the whole test. The sudden increase in the rate at the very end of the test occurred at the gas inlet of the counter, where no background discharges were present at the beginning. The most likely cause is too high a gas flow in the gap between the electrodes, which causes turbulence in the counter. Turbulence appeared at the counter gas inlet because the cross-section of the gas channel has sharp edges, and also at 'bad' points due to collision of the gas flow and the shock waves from discharges.

Turbulence has been avoided in more recent prototypes simply by reducing the gas flow. Surface defects are more problematic because they can act as 'seeds' for new background discharges. We expect that a different technology for copper film deposition via sputtering and rigorous quality control will solve this problem in the future.

4.2.11 R&D programme

While efficiency, position and time resolution of the new prototype have been well established during beam tests and meet all our requirements, the long-term stability and the lifetime of the counter remain the subject of further R&D. We will concentrate on optimizing the gas mixture, as the ageing of the counter depends crucially on its composition.

To ensure a high-quality cathode surface in a mass production process, a number of production technologies have to be explored as well as means of quality control.

Production of semi-conductive glass is carried out in industry. Test meltings of small quantities have been done out both at LENZOS, St. Petersburg and SCHOTT, Mainz, with positive results. An attempt to produce large quantities of glass at SCHOTT was partly successful, as it could be demonstrated that large amounts (≥ 600 l) could be melted and conditioned to the required resistivity. To be solved, however, remains the problem of manufacturing flat sheets (the float process is excluded for reasons of principle) of less than 4 mm thickness from the melt, thus avoiding costly cutting from large blocks.

4.3 TOF by Parallel Plate Chambers (PPC)

4.3.1 Introduction

A Parallel Plate Chamber (PPC) [12, 13] is a single-gap gaseous detector working in the avalanche mode (Fig. 4.12). It consists of two planar metallized electrodes made of ceramic or plastic kept a fixed distance apart (0.5–2 mm) by spacers. Typical sizes range from 30×30 mm^2 up to 60×60 mm^2, the total detector thickness is of the order of a few mm. A high and uniform electric field (4–6 kV/mm) between the electrodes yields immediate amplification of the primary ionization. The motion of electrons in the avalanche induces a very fast signal (rise time ~ 1 ns) on the electrodes [14]. A gas gain of the order of $\sim 10^4$ can be obtained without sparking [12].

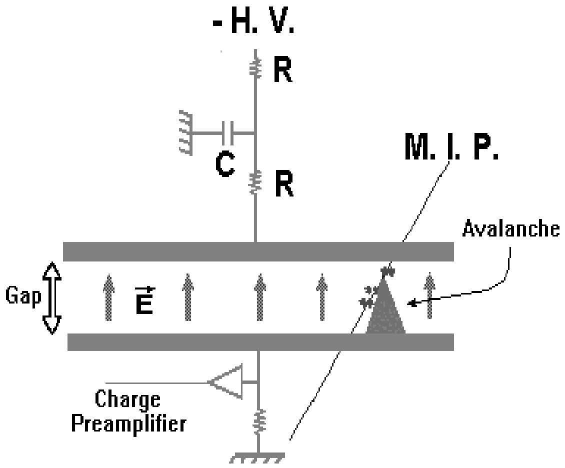

Figure 4.12: Schematic drawing of a PPC, also showing high-voltage and readout circuits.

4.3.2 Simulation of the PPC signal

Monte Carlo simulations of the gas discharge inside a PPC have been performed to predict the amplitude spectrum and the intrinsic time resolution. The initial charge produced in the PPC volume by an ionizing particle grows exponentially with the first Townsend coefficient which crucially depends on the electric field strength, the gas mixture and the pressure [15]. Secondary processes such as photo-ionization and ionization by metastable atoms or molecules have not been taken into account since their influence is not well understood. The parameter values used in the simulation are taken both from the literature [16] and from test measurements. Figure 4.13 shows the simulated PPC time resolution as a function of electric field strength for pure isobutane at atmospheric pressure.

These simulations predict an intrinsic time resolution of ~ 100 ps (excluding contribution from the readout electronics) at almost 100 % efficiency, meeting the target specifications for ALICE.

4.3.3 Cell design

Mechanical design

Contrary to standard PPCs, we currently intend to use a 'double-gap' for each individual PPC cell, with two consecutive gas-gaps (Fig. 4.14). Compared to a single gap (of the same total gap width), this design should provide equivalent or better detection efficiency for MIPs but possibly with somewhat better timing resolution. Tests of this concept are under way but have not yet yielded conclusive results.

We are currently investigating both electrodes made from thin ceramic plates and from plastic (polisulfone and policarbonate).

A PPC cell consists of three parallel electrodes (1 mm

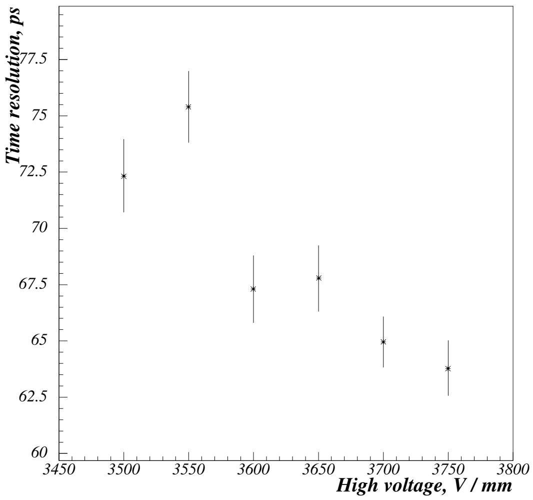

Figure 4.13: Monte Carlo simulation of the PPC time resolution as a function of voltage for pure isobutane at atmospheric pressure.

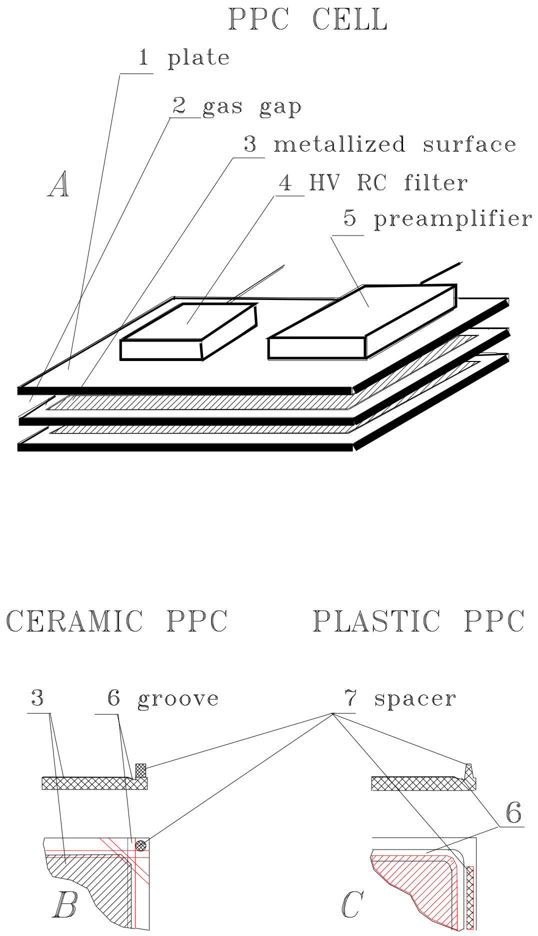

Figure 4.14: a) Schematic drawing of double-gap PPC cell; b) edge region of a ceramic electrode; c) edge region of a plastic electrode.

thick for ceramic and 1–2 mm thick for plastic) separated by spacers which determine the distance between the electrodes and the two gas-gaps with an accuracy of better than 5 μm. The gap will be between 1 and 2 mm, depending on further optimization. The inner surface of the electrode is mechanically polished (ceramic) or stamped (plastic), also with an accuracy of about 5 μm, and then metallized by vapour deposition (mainly chromium) to obtain a conductive electrode of 1–2 μm thickness. The dimensions of a complete cell will be 50×50 mm^2 with an active (metallized) area of 45×45 mm^2.

The two gaps form a single detector channel, i.e. high voltage and a fast-current preamplifier are common to both and placed on one side of the PPC outer surface. The opposite side is kept free for assembling several cells in a TOF module.

Ceramic electrodes

Ceramic PPCs of various designs have been widely used in the past in R&D programmes for CMS and ALICE at CERN and ITEP [12, 17]. The majority of their characteristics is well known and described in RD37 reports and ALICE Notes [18].

Details of the ceramic PPC cell foreseen for ALICE are shown in Fig. 4.14b. The electrodes and spacers are made of alumina (Al_2O_3) substrata, rigid enough to be processed to the required accuracy by a polishing machine. Grooves with rounded profiles, metallized up to half of their width, are made near the edges of the electrodes in order to prevent breakdown due to the increased electric field in this region.

Plastic electrodes

Plastic PPCs are much more novel detectors, but potentially cheaper and simpler to fabricate and assemble. The plastic electrodes (see Fig. 4.14c for details) are stamped at very high pressure (~ 100 bar) in order to obtain the necessary surface quality. The spacers are produced as part of the electrodes in a single processing step. A special low-temperature metal evaporation technique has to be employed to prevent deformation of the plastic.

Additional details and drawings of ceramic and plastic PPCs used during the R&D can be found in [19].

4.3.4 Module and barrel design

Simulations

In order to optimize detection efficiency and timing resolution of the PPC barrel, we have studied a multilayer TOF array with several consecutive detector planes (one

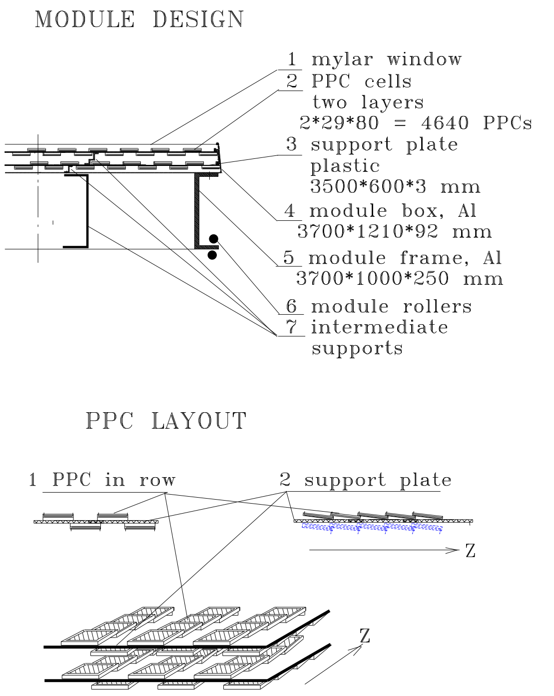

Figure 4.15: Cut view of a PPC module. The PPC TOF barrel will consist of two layers of PPC cells, assembled in a single gas volume.

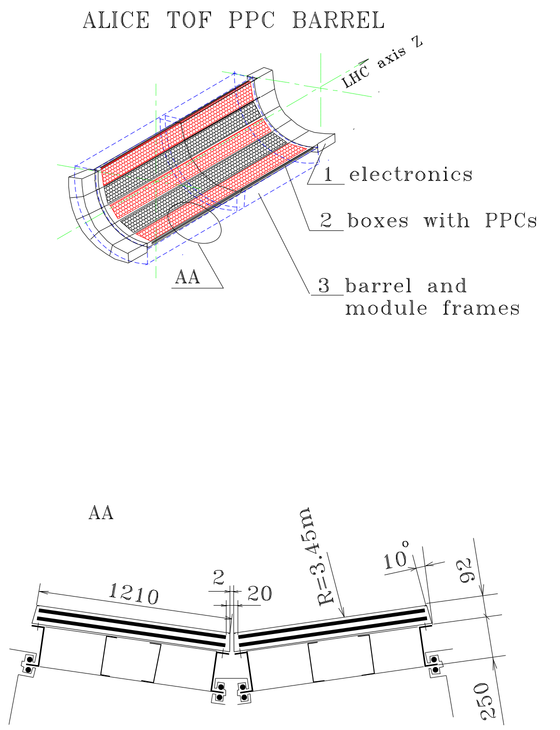

Figure 4.16: PPC barrel assembly. Most of the electronics will be housed in separate boxes at the end of each module.

to four planes). The number of electronics channels —the main cost factor for the PPCs— was kept constant at a value of 160 000 (corresponding to $\sim 10\,\%$ occupancy in the case of a single layer). With this constraint, cells in a n-plane configuration are n times larger than in a single-plane design.

For a single layer TOF barrel, the efficiency for particle detection is directly given by the product of the detector efficiency ($\sim 90\,\%$, including intrinsic PPC efficiency, inactive areas, etc.) and the single hit probability (also $\sim 90\,\%$ for $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$). For two consecutive layers, the geometrical and intrinsic efficiencies of the PPCs are less critical, with appropriate staggering, and the overall probability that a particle is detected in at least one layer reaches more than 90 %.

Unlike for the high resolution Pestov counter, sources of timing error other than intrinsic sources (e.g. electronics, path length corrections, etc.) are less important for the PPCs, and multiple timing measurements are largely uncorrelated. Doubling the number of planes will therefore also improve the time resolution, in principle by up to a factor of $\sqrt{2}$ at very low occupancy. At higher multiplicities, where a fraction of the particles has only one timing measurement, the gain is less and on average about 25 % at $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$.

Based on these considerations and on Monte Carlo simulations, we have decided to built two planes of PPC detectors, thus significantly improving the performance while keeping the complexity of construction and data analysis at a manageable level.

Mechanical design

Figure 4.15 shows the design of a TOF module with two layers of PPCs. Each module, consisting of a frame, a gas-tight box with PPC cells and an additional box containing front-end electronics, covers 3.7 m in length (3.5 m active), 1.21 m in width and weighs approximately 150 kg. If necessary, a few modules could be made shorter to reduce the amount of material in front of the PHOS and HMPID.

The support frame, made of aluminium alloy, is intended to ensure module rigidity. The gas box (2 mm Al side and back walls) is mounted onto this support; it is closed at the front (facing the vertex) by a thin Mylar window.

The preamplifier and high-voltage RC filters are mounted directly on the cell and connected in parallel to both gas gaps. High voltage and signal cables as well as d.c. current supply wires are situated between the PPC layers.

Each PPC layer is formed of rows of PPC cells mounted on both sides of a thin plastic support-plate. Both the cells within a row and the rows on either side of the plastic support are overlapping so as to ensure hermetic coverage (see Fig. 4.15). A module will contain 4640 cells ($29 \times 80 = 2320$ cells in each layer); the total number of PPC cells in the barrel is therefore about 167 000.

The module box has wedge-shaped edges in the longitudinal direction ensuring mutual overlap of adjacent modules. Layers of PPC cells in the module gas box are shifted accordingly in order to avoid creating a gap between the modules.

The total amount of matter in the PPC system, excluding the module frame, is equivalent to about 2 g/cm^2 on average, corresponding to 5 % of a radiation length or 1 % of an interaction length.

The PPC barrel consists of 36 identical modules, placed at an average distance of 3.5 m from the vertex (see Fig. 4.16). Modules are individually mounted on a space frame, identical to the one described previously for the Pestov counters, by means of rollers to facilitate installation and maintenance.

4.3.5 Front-end electronics

Layout and requirements

A schematic diagram of an electronic channel for the PPC [17] is shown in Fig. 4.17. The preamplifier (and in one of the schemes under consideration also the main amplifier) is mounted directly behind each individual PPC cell. It will be made from discrete components in surface mount technology, which helps to reduce cross-talk. Discriminators, TDCs, ADCs and auxiliary electronics will be housed in a separate box at the end of each module (see Fig. 4.16), allowing easy access and cooling. This part of the electronics could be identical to all TOF options (with the possible exception of the discriminators) and is described in Section 4.4.

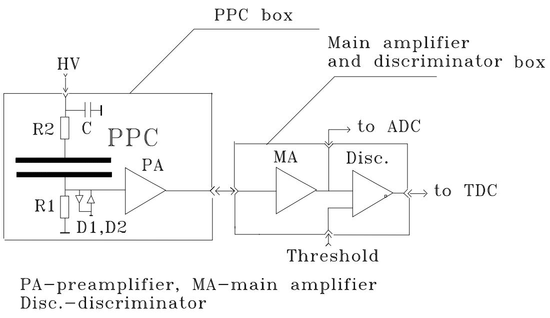

Figure 4.17: Electronic channel for PPC.

The basic requirements for the amplifier chain, which are specific to the PPC, are a fast rise-time for good timing and a large signal-to-noise ratio. The latter is of particular importance for efficiency and time resolution since the time uncertainty depends on this ratio linearly. The target figures for a time resolution of 100–200 ps are listed in Table 4.2.

Table 4.2: Requirements for the PPC amplifiers.

Gain	7–10 V/pC
Rise time	better than 3 ns
Noise	less than 900 e (r.m.s.)

Preamplifier

We have constructed and tested three different types of hybrid preamplifier for use in the PPC electronics:

- AMP1, which uses the common base Radeka scheme, is described in [20]. The main part of the amplifier is a grounded base scheme which provides it with low noise and low input impedance. It makes use of a new range of high-speed, low-noise transistors (Hitachi 2SC4995).
- AMP2, with the preamplifier and the main amplifier implemented together on one card. The scheme, presented in Fig. 4.18, includes a grounded base on the input and differential output signals. Stability is provided by feedback.
- AMP3, made with the help of industry in Russia, uses a cascode scheme on its input as described in [21]. The cascode circuit has a current shunt feedback. Only single polarity power voltage is applied, the supply current is 13.3 mA. The same Hitachi transistors as those described in the first case are used in the scheme.

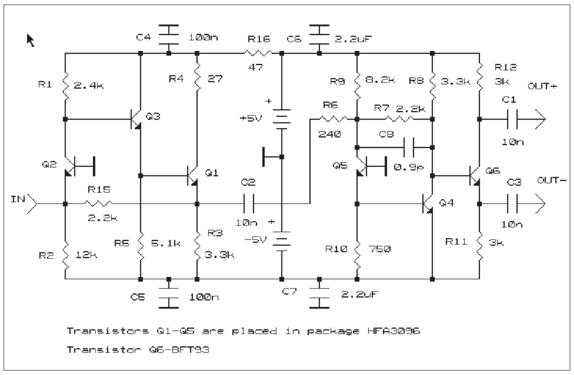

Figure 4.18: PPC amplifier chain (type AMP2).

The parameters of the three different types of amplifier are listed in Table 4.3.

Main amplifier

Additional amplification is achieved by means of the main amplifier, which can be placed some distance from the PPC. So far we have used a commercial amplifier chip (INA-02184) with parameters as listed in Table 4.4. It

Table 4.3: Parameters of the preamplifiers used in PPC test.

Type	AMP1	AMP2	AMP3
Gain	280 mV/pC	10 V/pC	1.2 V/pC
Rise time	2.0 ns	2.2 ns	1.5 ns
Power	140 mW	55 mW	78 mW
Noise (r.m.s.)	1100 *e*	1400 *e*	1200 *e*

Table 4.4: Parameters of the main amplifier used with PPCs.

Gain (at 1 GHz)	28 dB
Noise factor	2 dB
Input (output) impedances	50 Ω

provides total gains of 8 and 25 V/pC in connection with AMP1 and AMP3, respectively (AMP2 has sufficient gain already).

Discriminator

In the PPC test we have been using mainly Constant Fraction Discriminators (CFD), both models available from industry and a hybrid design made at ITEP. Given the pulse shape of a PPC and the expected timing resolution of a single counter (150–200 ps), both constant fraction and constant threshold discriminators could be envisaged, although a low threshold is essential for good timing resolution. So far we have concentrated on a discriminator design for the Pestov counter, which is more challenging, but could eventually also be used for a PPC TOF solution (see Section 4.4).

Cooling and power supply

The power consumption of the preamplifier, located in the PPC gas volume, will be some 400 W per module (distributed over 4 m^2), which does not require any forced cooling inside the module. The total power of about 15–30 kW, depending on the type of preamplifier used and including the electronics at the ends, will be removed by water cooling of the module frames and the electronics boxes.

Power (low and high voltage) will be supplied to individual groups of some 50 cells. The HV current will be monitored and limited for each of these groups by a single optron-based HV current trip circuit.

4.3.6 Gas system

To provide stable operation at fairly high electric field (a few kV/cm), the gas used in the PPC should not be transparent to its own UV light in order to avoid secondary avalanches in the gas volume. A large primary ionization is needed to improve the signal-to-noise ratio required for good timing.

We have investigated different gases for PPCs, among them monogases, such as pure isobutane or carbon dioxide, and double- or triple-component gas mixtures based on isobutane, Ar and CO_2 with a small quantity of freons and/or CF_4. The best efficiency and time resolution have been achieved for mixtures such as Ar/isobutane/CF_3Br (70:30:0.5). We propose to use a similar mixture, although the R&D on gases will continue: in particular, we will investigate the possible replacement of freon by another quencher (CF_4). Inflammable isobutane can be used in ALICE because of the relatively small volume.

The requirements for the PPC gas system are comparatively modest. The gas purity is not very crucial for PPCs because of the small drift and commercially available gases are adequate. A filter system (including electrostatic filters) is foreseen to remove dust. Effects of ageing due to polymer deposition on the anode surface are expected to be much smaller than in wire chambers because of the large surface electrodes. The total gas volume of the TOF system is approximately 8 m^3 and the gas flow will be around 100 l/h.

4.3.7 R&D programme

Test results

R&D on PPCs is currently performed for radiation-resistant calorimetry in the framework of the CERN RD37 programme inside CMS and for TOF application inside ALICE [17].

PPC detectors based on ceramic electrodes have been tested mainly at the 10 GeV ITEP proton synchrotron with secondary beams of 2–5 GeV/*c* momentum. The time resolution and amplitude distribution were measured for PPCs with different gap sizes (1, 1.5 and 2 mm), different gas mixtures (see above) and for different readout electronics. Fast current preamplifiers of different types (see above and Refs. 22, 23), with rise times of 2–7 ns, ENC noise from 1000 to 5000 *e* and 50–100 mV/μA amplification were used in the measurements.

The timing resolution could be improved from about 500 ps in the initial tests to better than 200 ps with the latest prototypes. As an example, signal amplitude distribution and TOF resolution obtained with a 1 mm double-gap PPC filled with an argon/isobutane mixture are shown in Fig. 4.19.

Future plans

While PPCs are comparatively easy to build and operate, further R&D is needed in a number of areas, in particular optimization of the gas mixture, pressurized operation (1–2 bar) and the design of the preamplifier. Little experience exists so far for plastic PPCs, and their performance, long term stability and production technologies have still

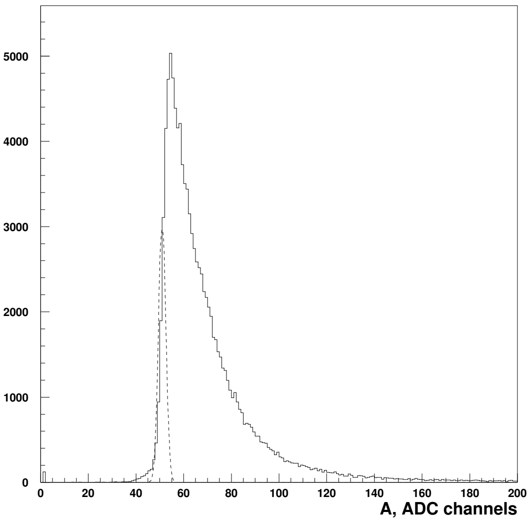

(a) Amplitude distribution (pedestal is shown as a dashed line).

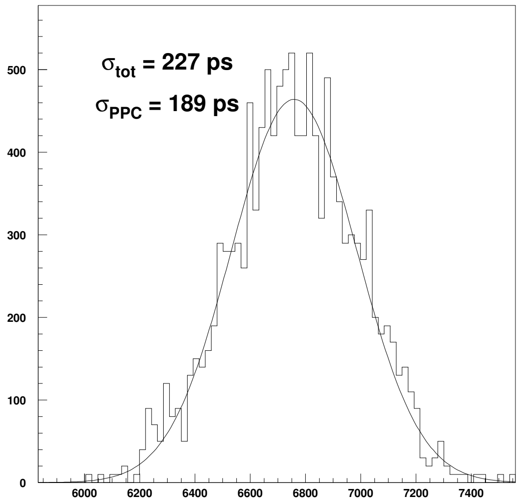

(b) Timing distribution.

Figure 4.19: Results from measurements of a PPC filled with an argon/isobutane mixture.

to be evaluated. A large area prototype is in production and will be tested in 1996 to study cross-talk and system aspects.

4.4 TOF readout electronics

The front-end electronics described below (TDC, QDC and possibly discriminators) are common to all three TOF options. However, only the Pestov spark counter requires an electronic resolution better than 50 ps to match its intrinsic performance. In the case of the PPC or the scintillator TOF system, the requirements for the front-end electronics are less demanding. The amplifier chain needed for the PPCs has been described previously.

Discrete prototype versions of most of the components have been built and tested with Pestov counters in the beam under realistic conditions. First integrated versions will be ready for evaluation by spring 1996.

Requirements for the readout electronics

The expected low rate in ALICE requires no pipelining for the TOF electronics. In general, we will use either analogue delays, slow shaping or a track-and-hold logic until the arrival of a trigger or strobe signal. The TDC is started on arrival of a particle and stopped by the beam clock. The subsequent readout depends on the TDC scheme used and is described below.

The information volume per channel is 8 bits for the QDC and 8 and 12 bits for the TDC for the fine and coarse ranges, respectively. Zero suppression for both the QDC and the TDC is performed via digital masks on the TDC content. The address uses 20 bits, i.e. 48 bits in total are generated for each channel containing a hit. The data will be multiplexed and transferred to front-end digital crates via an optical link using the standard to be adopted for ALICE (see Chapter 10).

4.4.1 The double-threshold discriminator

For the Pestov counter a fast leading-edge discriminator is essential, since the excellent timing information is contained in the first 300 ps of the pulse. Moreover, due to after-pulses, the trailing part of the signal varies and may even have secondary maxima. Standard leading-edge discriminators limit the time resolution to ~ 70 ps, as no amplitude-dependent walk correction is possible with the irregular pulse shapes. We envisage therefore a double-threshold discriminator (DTD) [24] with hardware extrapolation to t_0 which has the following characteristics:

- To extract time information, the discriminator extrapolates to the zero crossing of the leading edge of the pulse. This is equivalent to the compensation of a time jitter caused by rise time or amplitude variations.

- The response of the discriminator is independent of the pulse shape or the pulse length. In particular, after-pulses will not influence the proper functioning.

- The thresholds are low enough to allow for signal amplitudes of about 100 mV.

- The range for the rise time of the input signal is 250 ps to 1 ns.

Functional description

A circuit diagram of the double-threshold discriminator is shown in Fig. 4.20. It consists of two independent discriminators with different thresholds. When the rising edge of the input signal passes one of the two thresholds, the corresponding comparator switches to the low state. The voltage on the capacitor starts to decrease linearly as the capacitor is discharged by one of the current sources.

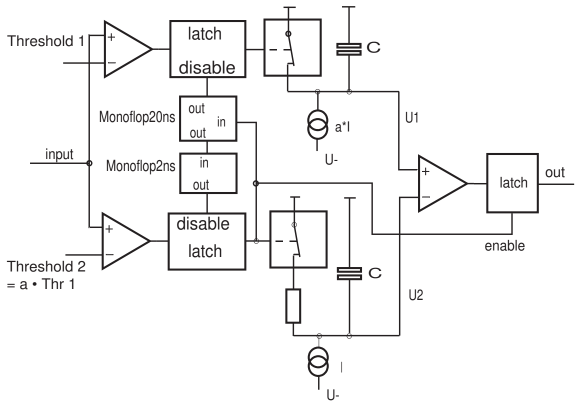

Figure 4.20: Circuit diagram of the double-threshold discriminator.

The individual components are adjusted such that there is an offset between the initial voltages U1 and U2 of the two capacitors, determined by a resistor, and a factor of two difference in the time constant for discharge. Because of this difference in starting value and discharge speed, the voltages of both capacitors will be equal some time after the arrival of the signal. This intersection happens after a fixed and constant time delay with respect to the extrapolated zero crossing of the input signal and is sensed by a third comparator. Its output is then used in the TDC for the actual TOF measurement.

Test results

In order to test the properties of this discriminator concept, we measured the output delay as a function of the signal shape with a pulse generator for rise times between 250 ps and 1 ns. The results, obtained by using a fast oscilloscope (1.5 GHz bandwidth), are shown in Fig. 4.21. For an ideal discriminator the output delay should be independent of the signal rise time. The actual performance is very close to achieving this, with a peak-to-peak variation of only 20 ps and a r.m.s. of 6 ps.

Discriminators for PPC and scintillator options

Both the scintillator and PPC options could use standard solutions for the discriminator, but the double-threshold discriminator might still be an attractive option. It requires no offline corrections and its performance is certainly adequate. For the PPC, the parameters of the DTD have to be matched to the shaping time of the preamplifier.

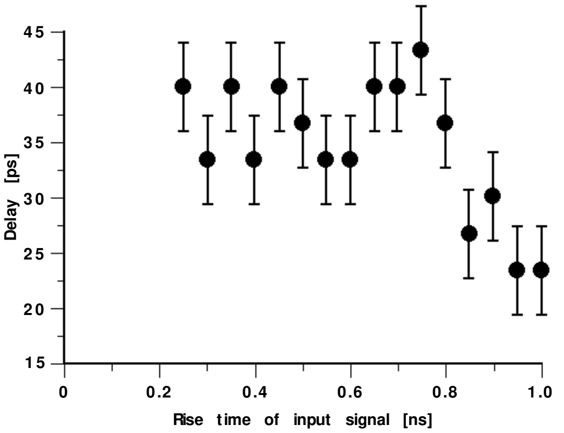

Figure 4.21: Output delay of the double-threshold discriminator as a function of input-signal rise time.

4.4.2 TDC options

The TOF measurement in the ALICE experiment requires a r.m.s. time resolution from the front-end electronics of better than 100 ps for the PPC and scintillator options and better than 50 ps for the Pestov spark counter. The time range should be at least 40 ns and preferably much larger. Because of the very high number of channels, a large degree of integration is desirable. We are currently evaluating two concepts for the TDC, an analogue and a digital version. Given our timing requirements, the real performance or potential problems, like cross-talk, can only be assessed on an actual chip. The final choice will therefore be made on the bases of VLSI chips, which are currently in the prototype stage.

Analogue Time-Amplitude-Converter (TAC)

In a TAC, the voltage difference produced by a ramp generator between the arrival of a start and a stop signal is used to measure time intervals. The present design employs a new scheme for the ramp generator, which is shown in Fig. 4.22 [25].

A start signal sets a flip-flop which is connected to a chain of 16 digital gates. A rising edge propagates through the daisy chain with a velocity determined by the gate delay, switching the outputs from low (0 V) to high voltage. All outputs are connected via equal resitors to a common node, where the voltage consequently increases in equal steps.

The node is also conected to an analogue low-pass filter of 5$^{\text{th}}$ order, which linearizes the voltage step function and charges a capacitor at its output. Simulations have shown

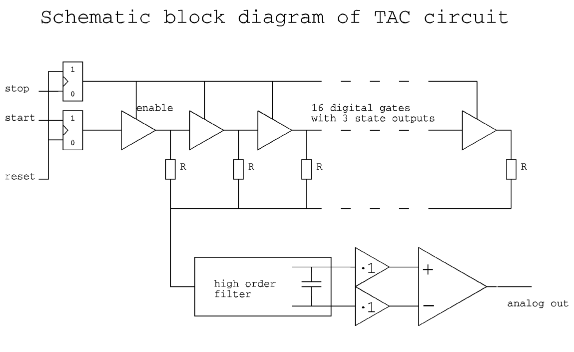

Figure 4.22: Block diagram of the analogue TAC option.

that the actual design of this filter is not very critical and a number of solutions have been identified.

When the stop signal arrives, the gates are switched to high impedance and the capacitor at the filter output is disconnected. Its voltage, available for several μs, its proportional to the time delay between start and stop signals and will be read out in an ADC.

During a nonlinear region of 7 ns after start, the stop is disabled leading to a dead-time of one clock period (25 ns). In general, a TAC is sensitive to noise on the power line, but as a significant current (25 mA) flows only during the time between start and stop, the the power line can easily be kept stable. Cross-talk effects should be reduced in this design as any charge injection will disappear within a few ns. A good linearity requires careful matching of all components, in particular the resistor values. A feasibility study of a VLSI version of this TAC has shown that a r.m.s. resolution of better than 50 ps should be possible. Test samples with slightly different designs to evaluate cross-talk and performance will be obtained in early 1996.

Digital Time-To-Digital Converter (TDC)

The second option under development is a high resolution, highly integrated, and fully digital TDC. It will make use of an array of Delay Locked Loops (DLL) [26].

A block diagram of the digital TDC circuit is shown in Fig. 4.23. The basis of the time measurement is a basic gate delay. Since this delay is very dependent on the process, supply voltage and temperature, voltage-controlled delay elements are used as a part of a DLL locked to the system clock. This results in a self-calibrating TDC. In the DLL array scheme, the resolution of the time measurement is better than the delay of a basic delay cell due to the skew introduced between the DLLs in the array by a time-shifting DLL with a slightly bigger delay per cell. All the DLLs are synchronized to the system clock.

Previous work demonstrating this concept, in a 1 μm CMOS process using a clock frequency of 65 MHz, has already shown a r.m.s. resolution of 76 ps. Using a standard 0.5 μm CMOS process and a clock frequency of

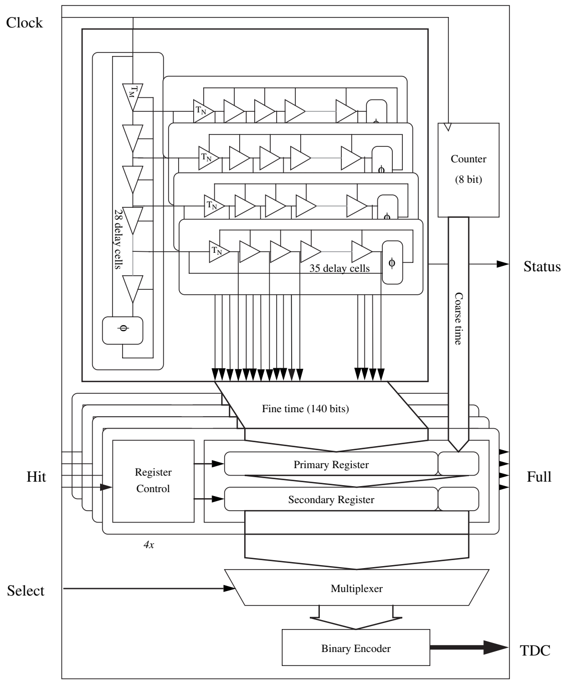

Figure 4.23: Block diagram of the digital TDC option.

80 MHz, a time-bin size of 90 ps (r.m.s. resolution of 26 ps) can be obtained, therefore avoiding the use of exotic processes or very high clock frequencies, and the corresponding penalty in terms of power consumption and system cost.

The dynamic range of the measurement is readily expandable by a coarse time counter that counts clock cycles. Since both fine and coarse time measurements are synchronized to the same clock, no ambiguity is generated in the resulting measured time. The time measurements are performed relative to the bunch crossing (system clock) therefore the coarse time counter can also be used for bunch crossing identification. The time resolution required is not compatible with the use of the slow-edged, single-ended TTL inputs in the reference (clock); and signal (hit) paths. Noise insensitive differential inputs will therefore be used. Each channel consists of a double register where the measurement is stored when a hit is detected. The use of the secondary register to store the measurement before it is read out greatly increases the double-pulse resolution.

In this architecture, the number of channels integrated on a single chip can be large, since it does not affect the array functionality. The readout logic is common to all the channels on the chip. In the final chip, once we have a better understanding of the experiment readout requirements, different readout logic and buffering schemes can be introduced, ensuring smooth integration of the TDC chip into the data acquisition system. The use of a CMOS process and the availability of digital standard cell libraries make the integration of complex digital readout circuitry with the more sensitive fully-customized array of DLLs straightforward. In the proposed circuit, 16 channels are included. The dynamic range of the measurements is expanded to 51.2 μs by the introduction of a 12-bit coarse time counter and the r.m.s. resolution will reach the 25 ps range by dividing the 80 MHz clock period into 140 time bins, using an array of 4 DLLs.

4.4.3 Charge measurement

While for the Pestov spark counter the charge measurement will be used to determine the particle's position across the strips by the centre of gravity method, the PPC and scintillator detectors will use the charge measurement for additional (offline) slewing correction of the time walk.

A first prototype unit, using SMD technology, is based on the 8-input 8-bit commercial ADC MAX 155, with an encoding time of 4 μs for each channel. The detector signals from the strips are integrated by fast operational amplifiers (AD 828). The output signals of the integrators are proportional to the collected charge and have a decay time constant of about 1–2 μs. These signals are fed to the inputs of the ADCs which track-and-hold until the strobe signal arrives from the trigger after about 1.2 μs. The back edge of the strobe signal (the duration of the strobe signal is 150–200 ns) starts the encoding of the signals channel-by-channel in parallel in all ADCs.

A prototype has been built in the CAMAC standard as a 32-channel QDC. Tests with a generator showed that the spread of codes for all 32 channels was not more than 1 bit.

We are also considering an option with fast gating instead of operational amplifiers. It would be less sensitive to noise and cross-talk, but needs an analogue delay unit of about 2 μs at each input.

4.5 Trigger and timing distribution

The Trigger, Timing and Control system (TTC) has to

- deliver a stop signal for the TDCs derived from, for example, the 40 MHz bunch-crossing clock, with a (pulse-by-pulse) timing jitter smaller than 50 ps;
- transmit the trigger and strobe signals to initiate the readout or to deliver a fast clear signal;
- transmit control information, i.e. set-up information for the front-end chips as well as event-related information (event number, bunch-crossing number, etc.).

A TTC system under consideration for ALICE is the optoelectronic distribution system being developed within the framework of DRDC projects RD12 and RD27. The RD12/RD27 system uses relatively few high-power lasers which are distributed in a tree-like structure via passive optical couplers to several thousand receivers. The fibre receiver cards perform optical-to-electric conversion, deskewing (optional) and fan-out to the front-end controllers, from where the individual TDCs are addressed.

Slow, long-term variations of the clock signal over several hours can be corrected using the actual data. However, our requirements for pulse-to-pulse stability are very strict. At present, 50 ps jitter has been reached in the laboratory [27], but appears to be the limit of present technology.

4.6 High-momentum PID by RICH

4.6.1 Introduction

The detector presented here is devoted to the identification of particles above the momentum range to be achieved by the TOF barrel at r = 3.5 m. The proposed RICH high-momentum PID (HMPID) array has a total area of ~ 10 m^2 and is located 4.7 m from the beam to benefit from the lowest possible particle density.

RICH detectors have been used for particle identification over the last decade as an alternative solution to TOF. Several large experiments have adopted this method, e.g. OMEGA, DELPHI, CRID. The special feature of the ALICE RICH is identification of particles reaching the detector in a high-multiplicity environment (50 particles per m^2) and in a broad range of momenta.

The detectors mentioned above have in common the technique of converting photons in a wide gap (62 mm in DELPHI) filled with TMAE as photoconverter and drifting these electrons, like in a TPC, towards the photodetector. Such a design can handle the event rates found in these experiments and would, in principle, also be feasible at ALICE at the expected rate of 10 kHz in Pb–Pb mode. However, the high density of particles together with background hits generated by, for example, neutrons in the gas gap, make pattern recognition capability a prime requirement for our application. Therefore the ALICE RICH should have the thinnest converting layer possible in order to provide quasi 2-dimensional event patterns.

Such a requirement can be fulfilled with the fast-RICH layout [28], where the conversion gap can be restricted to a few millimetres at the cost of detecting the photoelectrons at a cathode segmented into small pads. The availability of highly-integrated electronics at reasonable cost has made such a pixel-like readout possible. In addition, due to the fast collection of the photoelectrons (within tens of nanoseconds), this layout allows operation at much higher rates than Pb–Pb at ALICE, e.g. pp collisions.

At present, the CERES pixel RICH [29] is operating successfully with heavy-ion beams at the CERN SPS, using TMAE as a photoconverter.

Since 1989, a novel photoconverter has been investigated, consisting of a thin film of CsI operated in re-

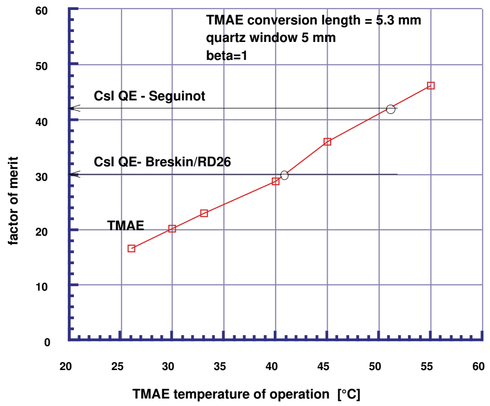

Figure 4.24: Comparison of the factors of merit of a fast RICH operated with CsI and TMAE in a photoconversion gap of 5.3 mm and a 5 mm thick quartz window. Arrows refer to: (1) average CsI QE, (2) best CsI QE. In the case of TMAE, the factor is plotted against the temperature at which a fraction 1/e of the photons are converted in the 5.3 mm gap. Identical factors are obtained at 42^{o}C for (1) and 52^{o}C for (2).

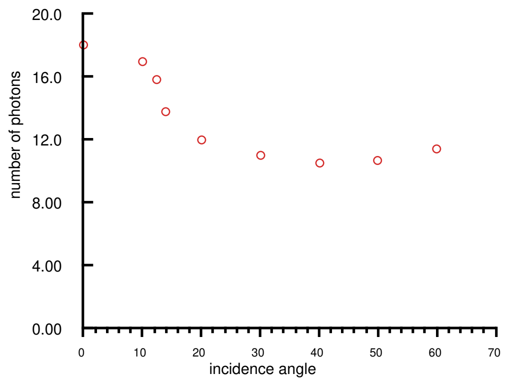

Figure 4.25: Calculated number of photoelectrons per ring as a function of incidence angle of a particle (using the measured CsI QE, 5 mm thick quartz window and including all photon losses).

flective mode on large areas of photocathodes with high quantum efficiency [30]. This solution considerably eases the construction, eliminating the use of TMAE and high-temperature operation, as well as the need for a quartz window at the entrance of the photodetector. Considering the present TMAE and CsI quantum efficiency and a fast-RICH geometry with a 4 mm conversion gap, Fig. 4.24 indicates at what temperature a TMAE detector has to be operated to achieve an equivalent factor of merit.

The properties of this novel technique have been under evaluation since 1992 in the CERN RD26 project [31] and at many external institutions [32].

4.6.2 Intrinsic particle separation

The particle separation in a RICH detector is determined by a large number of factors: the refractive index of the radiating medium, the chromaticity of the radiator, the range of wavelengths accepted by the photon detector, the focusing arrangement of the detector, the pixel size of the detector, the number of detected photoelectrons and, finally, by the angle of incidence of the particles on the detector. Given the requirements for the HMPID, we have adopted a proximity focusing RICH with a C_6F_{14} radiator and a pad cathode segmented into pads of 8×8 mm^2. The RICH modules will be oriented in such a way as to optimize the resolution with respect to the average angle of incidence.

Lower momentum limit

In principle, a Čerenkov counter can positively identify particles above the threshold for photon emission. In a high-multiplicity environment, however, the momentum range for identification is also set by the requirements of pattern recognition. The number of detected photoelectrons should be above a certain threshold value for the pattern recognition to be effective. This threshold translates into a lower momentum cut-off since the number of detected photoelectrons in the pattern is proportional to $\sin^2\theta$.

In addition, the detection and resolution power is dependent on the angle of incidence of the particles detected. In Fig. 4.25 we show the dependence of the number of detected photoelectrons on the angle of incidence [33]. If this angle is too large then a fraction of the emitted Čerenkov photons do not enter the photodetector, and the Čerenkov patterns tend to be very large and difficult to recognize, especially in a high-density environment. In practice, this feature can limit the lowest accessible momenta when particles are bent by the magnetic field.

In our case the effective low threshold value is ~ 0.9 GeV/c, assuming a charged multiplicity of 8000 per unit of rapidity and a HMPID at 4.7 m.

Upper momentum limit

At the upper end of the momentum range the constraints are different. The resolution is mainly determined by the noise from the charged-particle background, and by the lever arm chosen for the proximity focusing. A larger lever arm results in larger patterns, allowing better angular resolution and hence better particle separation, although at the cost of a reduced acceptance at the edges of

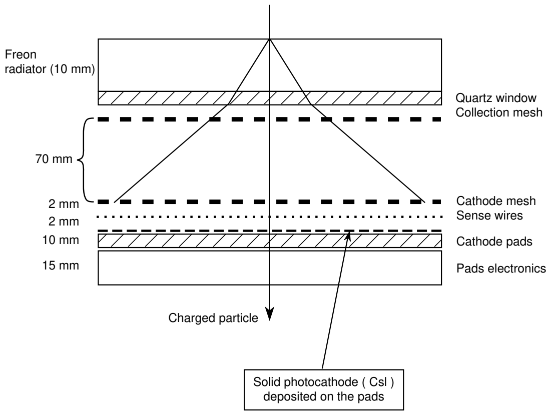

Figure 4.26: Schematic layout of the fast CsI-RICH.

the modules. Consequently two optimizations are possible; one aimed at a maximum geometrical efficiency and the other at the highest upper momentum limit for particle separation. For the HMPID, the latter is of more importance.

Detection of minimum ionizing particles

The RICH detector also acts as a charged-particle tracker. The minimum ionizing particles passing through the MWPC are recorded and, due to the difference in induced charge on the cathode for MIPs and photoelectrons, unambiguously identified in the analogue readout of the detector.

Tracking in the TPC stops at 2500 mm radius and the particles endure multiple scattering in the TPC outside wall (1.9 % of X_0). However, at the positions of the RICH and in the momentum range covered by the detector the multiple scattering should contribute to a tracking error which is much smaller than the mean distance of particle impacts. Therefore it is possible to uniquely determine the track and its impact on the detector with a precision that is in all directions better than 1 mm.

4.6.3 The fast CsI-RICH

The layout under consideration, sketched in Fig. 4.26, has a proximity focusing geometry. It is composed of

- A symmetric MWPC as photodetector, with 2 mm anode to cathode distance, 4 mm wire spacing and pad size of 8×8 mm^2. The CsI film is deposited on the pad cathode.

- A distance frame creating the proximity gap necessary to obtain suitable ring radii (100–150 mm at $\beta = 1$), equipped with a wire electrode suitably polarized to collect the electrons deposited in the proximity gap by the charged particles.

- A radiator array using liquid C_6F_{14} with a quartz window for containment.

The chamber gas, e.g. methane, is circulated throughout the whole detector at room temperature and atmospheric pressure. The integrated pad front-end electronics, chosen to operate in an analogue multiplexed mode, is implemented at the back of the pad cathode.

The CsI photocathode (PC)

The success in operating large CsI PCs depends on the achieved QE and on the stability of the CsI film under various ageing agents. For four years, these features have been investigated following two main lines (for recent reviews see [34, 35]): laboratory studies of the basic CsI properties made on small samples illuminated with monochromatic UV lamps [36], and studies of large PCs in fast-RICH prototypes by means of Čerenkov effect in beam tests [37–41]. These have enabled a technology for the production of large films of good performance to be developed, which are summarized below.

Quantum efficiency: Fig. 4.27 shows a CsI QE curve obtained from a large number of laboratory samples, thus taken as a reference curve. It has been demonstrated [42] that, in order for this result to be reproducible, the film should be submitted to a heat treatment, consisting in heating the film after deposition at 60°C under vacuum for 8–10 h.

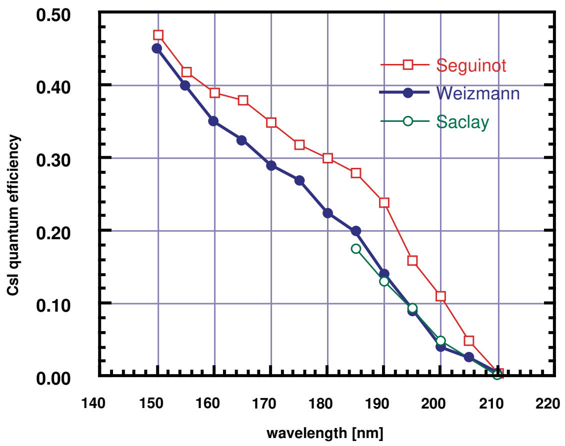

Figure 4.27: CsI quantum efficiency plotted against wavelength from a monochromatic UV beam.

Substrate: It has been observed that the good QE achieved on small substrates, which are easy to clean and polish, was degraded when the large areas of printed circuit board (PCB) necessary for pad cathodes were used. A number of micro-analysis studies of the films [43] have pointed out some important features to be satisfied by the substrate, such as absence of contact between CsI and copper, good polishing, etc. A standard technology has been developed for producing adequate PCBs by pre-polishing the copper layer and covering it with a thick nickel layer.

Operation of CsI PCs under gas: It has been demonstrated [44] that the QE of Fig. 4.27 is achieved in methane at 740 torr when the reduced field value E/p at the PC surface is ≥ 2 kV/cm$\times$torr, as is the case for us. For other gases (iC_4H_{10}, Ar, C_2H_6), a slight reduction of the electron yield ($\leq$10–20 %) was observed at the same E/p value, due to the back-scattering of electrons on chamber gas atoms. This implies that argon-based mixtures could be used to reduce the total flammable gas load [38].

Processing of the CsI film: Tests have shown that a standard thermal evaporation of CsI under a vacuum of 10^{-6} torr achieves the necessary quality provided certain precautions are taken, such as pre-melting of the CsI load, thorough cleaning and outgassing under vacuum. A multi-source evaporation technique has been developed, making the evaporation of large films of homogeneous thickness a simple and routine procedure [31].

Exposure to air: CsI is known to be hygroscopic. The QE drop induced by exposure to air has been measured on small and large PCs [34,38,39]. A one hour exposure was found to have no effect whereas the QE decreased by 30 % after a 15 h exposure (Fig. 4.28). This observation allows the transfer of a PC from the evaporation vessel to the detector to take place in air within a short but manageable time. Furthermore, promising results have recently been obtained [45] with protective layers (e.g. MgF_2) deposited on CsI films.

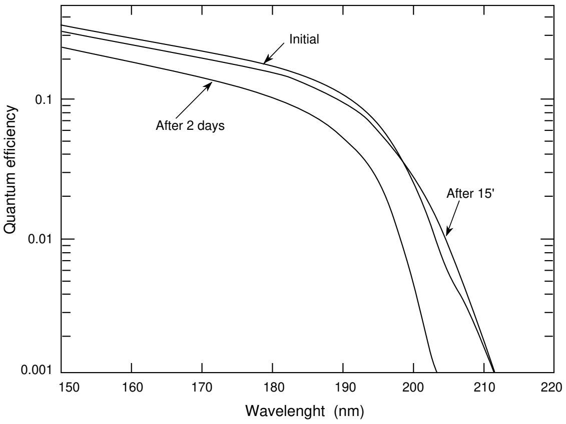

Figure 4.28: Influence of exposure to air on the CsI QE.

Ageing: Ageing of a CsI film can be induced by direct photon exposure and by ionic bombardment from the avalanches in a MWPC. The former case is never reached in the ALICE detector. Measurements are still in progress on the ion bombardment because there are still large discrepancies between measurements: the same QE drop to 80 % is observed after exposure to an integrated charge of 1.5×10^{-4} C/cm^2 [34] and 20×10^{-4} C/cm^2 [39,40]. However, for a gas gain of 5.0×10^5, the less conservative dose (1.5×10^{-4} C/cm^2) is reached after 10^{10} central events with a density of 100 particles and 1000 photoelectrons per square metre. This figure is much larger than the 10^7 central collisions expected in one year of operation for ALICE.

Long term stability: It is possible that long-term degradation of the CsI QE can still take place, as a result of unidentified slower reactions initiated by air/gas exposure or bombardments. Such long-term stability is being tested by keeping PCs under gas flow, exposed to a moderate irradiation for long periods (8–12 months) and periodically evaluated. Such a test [46] has shown (Fig. 4.29 and 4.30) a good stability over 8 months; a large PC has been under periodic beam test for about one year, showing equally good behaviour. However, more information is required, using exposure levels likely to be met in the experiment.

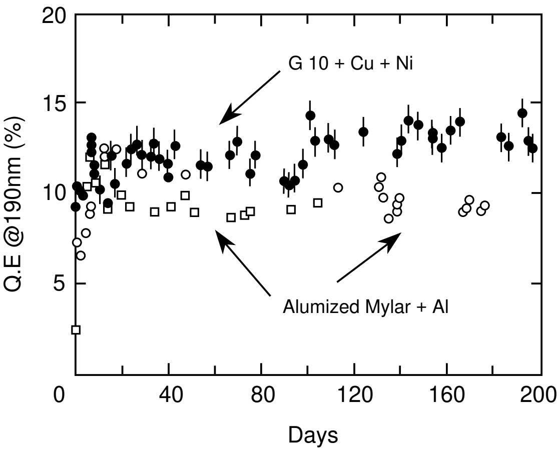

Figure 4.29: Ageing of a CsI PC exposed to small radiation doses and kept under gas flow.

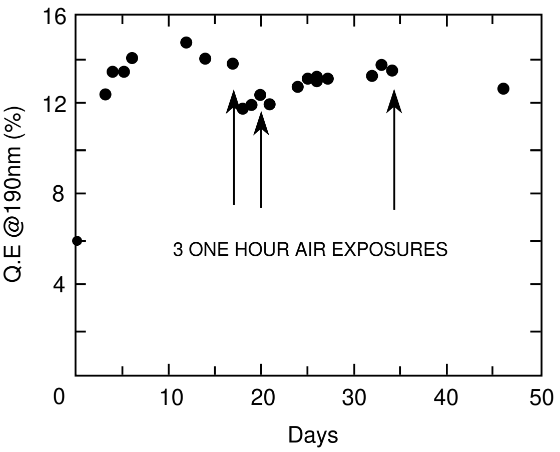

Figure 4.30: Long term stability after air exposure.

The photon detector

The photon detector aims to detect very low-level signals induced on the pads by single electron avalanches. A limitation on the chamber gain is imposed by the mechanism of photon feedback, where photons emitted by a primary avalanche re-emit electrons at the PC, possibly in a diverging manner. In order to obtain the maximum signal, we have taken advantage of the low event-rate at ALICE by using a front-end preamplification with a long integrating time (700 ns), allowing conservative MWPC geometrical parameters. A 2 mm gap, which keeps the sensitive volume sufficiently thin, is easy to maintain on large-area modules. It also allows a wide enough induction spread on a pad size of 8×8 mm^2 to perform accurate location by centroid measurements (≤ 1 mm r.m.s.) when using low-noise analogue front-end electronics.

For 3 years, prototype detectors have been operated at test beams at a total gain of $2–3 \times 10^5$ and at a rate of up to 1 kHz per cm of wire, showing neither problems of stability nor degradation, and keeping the photon feedback at a sufficiently low level (see below).

Different gas mixtures were tested, i.e. CH_4, CH_4/iC_4H_{10} (90/10), Ar/iC_4H_{10} (80/20), with equally good results.

The radiator

For the momentum range of interest, the radiator showing the best performance is liquid C_6F_{14} with a thickness of up to 12 mm. It has a good UV transparency and a low chromaticity, permitting a satisfactory Čerenkov angle resolution. The radiator is closed by a quartz window for containment. This compound is widely used in existing large RICH systems, like DELPHI and CRID, where a lot of experience has been gained.

4.6.4 Performance of the fast CsI-RICH

We now describe the results obtained with a fast CsI-RICH using single particles at a test beam and present the status of the pattern recognition of very dense events.

Beam test results

A total of 22 large PCs were studied [37, 38], ending with the production of four PCs of sizes 50×50 cm^2 (1 PC), 80×30 cm^2 (2 PCs), 30×30 cm^2 (1 PC) made of polished PCBs covered with nickel/gold, all subject to the heat treatment mentioned above. These PCs repeatedly provided the best QE performance. The tests were performed at the PS/T11 line using a 1.0–3.5 GeV/c proton/pion beam. The analysis of analogue pad ring patterns requires a specific programme to evaluate the fraction of Čerenkov photoelectrons in the presence of background photon sources. A Monte Carlo program was developed to simulate the full response of the detector to photon generation, transport and electron detection, allowing the QE of the PCs to be evaluated. The main results are summarized below.

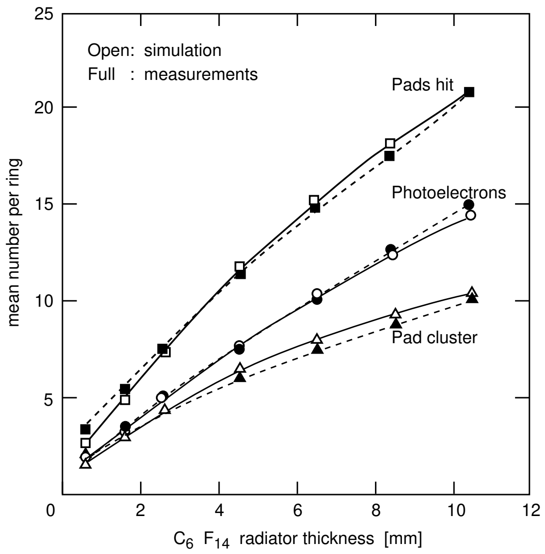

Figure 4.31: Test beam results obtained with a 30×30 cm^2 PC (C_6F_{14} radiator, pions at 3.0 GeV/c, 100 mm ring radius).

Photoelectron yield and resolutions: In Fig. 4.31, we notice that 15 photoelectrons per ring are detected for a 10 mm thick C_6F_{14} radiator crossed at normal incidence by 3.0 GeV/c pions. An average of 20–24 pads are recorded per ring, corresponding to a mean cluster size of 1.5 pads per photoelectron [35, 38].

The gas was methane. A 8 mrad resolution per Čerenkov photon is measured with a 100 mm ring radius as seen in Fig. 4.32. Simple technical modifications, e.g. grid transparencies, electronic noise and a larger proximity gap will further improve the number of detected photoelectrons and the angular resolution.

CsI QE in a fast RICH: Figure 4.33 [38] shows that the QE measured on the 4 PCs in a fast RICH reaches 80–90 % of the reference QE curve. Such QE already satisfies the ALICE PID in the high momentum range, as seen in the next paragraph. However, in order to improve the separation at low momentum, that is in the region where the number of created photons is smaller, a further increase in QE would be welcome. We believe that further development could increase the QE of CsI, notably in the wavelength region above 190 nm.

Photon feedback evaluation and stability: In open MWPC geometry, the photon feedback yield has been studied [35, 38] with UV lamps and charged particles as a function of gas gain and composition. Well-quenched mixtures keep this yield at an acceptable level at a gas gain of $2–3 \times 10^5$ where the single electron detection efficiency is ~ 85 %. The gas mixtures tested were: CH_4, CH_4/iC_4H_{10} (90/10), Ar/iC_4H_{10} (80/20). It was concluded that 3 out of the 15 photoelectrons quoted above

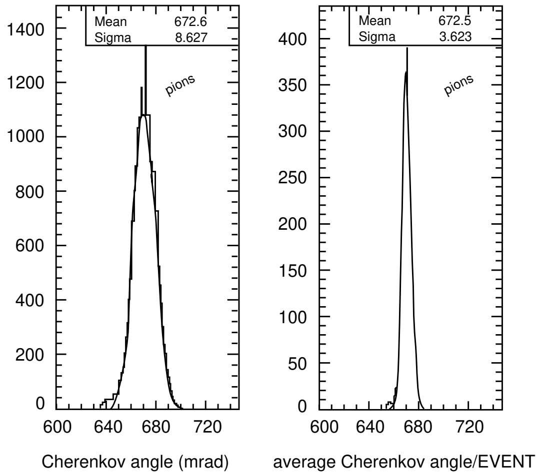

Figure 4.32: Čerenkov angle resolution: (1) single photon, (2) Čerenkov rings obtained with a 30×30 cm^2 PC (10 mm thick C_6F_{14} radiator, pions of 3 GeV/c, 100 mm ring radius).

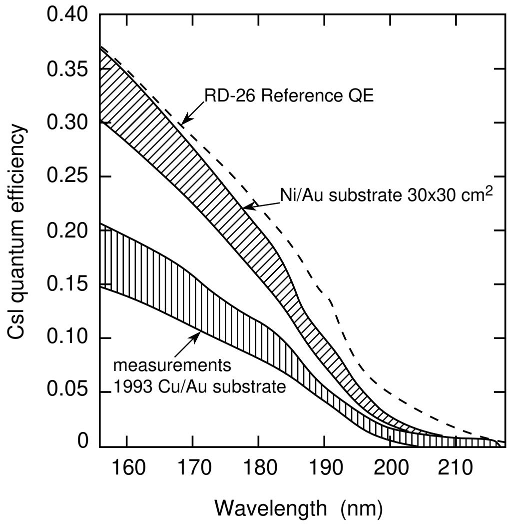

Figure 4.33: Comparison of the CsI QE measured with a fast RICH for a 30×30 cm^2 PC and UV monochromatic lamps.

originate from photon feedback. However, since the conversion occurs close to the site of the primary emission due to the large cluster size, in fact they reinforce the primary signal overlapping it, without deteriorating the angular resolution. The chamber operation has been found to be of excellent stability.

Background: The total number of background hits recorded outside the Čerenkov fiducial zone is shown in Fig. 4.34. In methane, this corresponds to an average of 0.5 hits per event (particle and ring), originating also from Čerenkov photons due to delta rays created in the quartz window during the passage of the particles [35, 38].

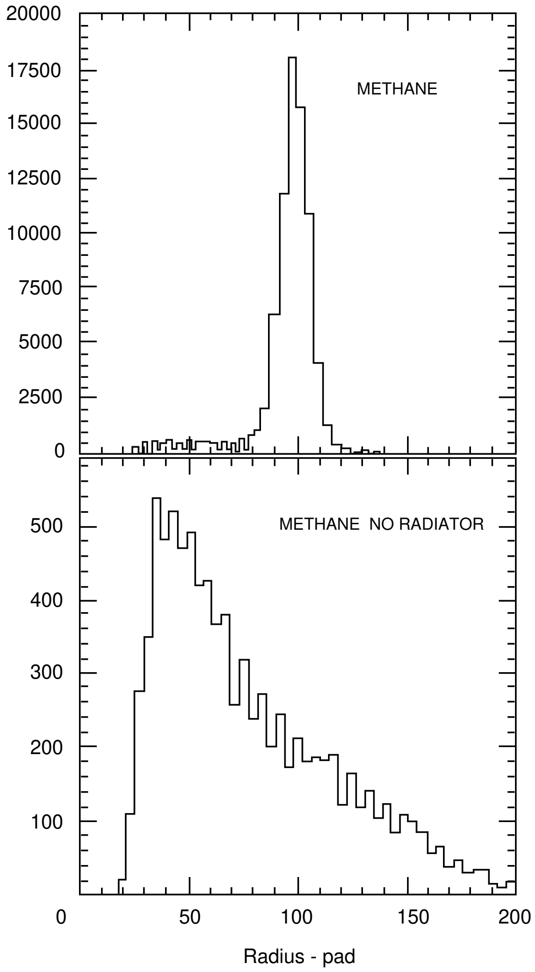

Figure 4.34: Background level due to secondary photon emission.

Pattern recognition of very dense events

Efficient pattern recognition of events with 50 to 100 particles per m^2 is the second key feature to be demonstrated. Since it is not yet possible to test a large enough fast CsI-RICH with such very dense events at an ion beam, the results presented were obtained by Monte Carlo simulation, and by superimposed real events.

A special pattern recognition method was applied to simulated high-density events (50 particles per m^2) [47]. The simulated events were constructed both from computer-generated patterns (using a SHAKER momentum distribution) and by the superposition of events measured with the prototype RICH at the PS in a 3 GeV/c beam of pions and protons. The overlap was achieved by translating the single-ring events in x and y coordinates randomly over the area of the imaginary module of 48×48 cm^2, and simply superimposing the raw pad data, thus creating an image that corresponds to a high-density

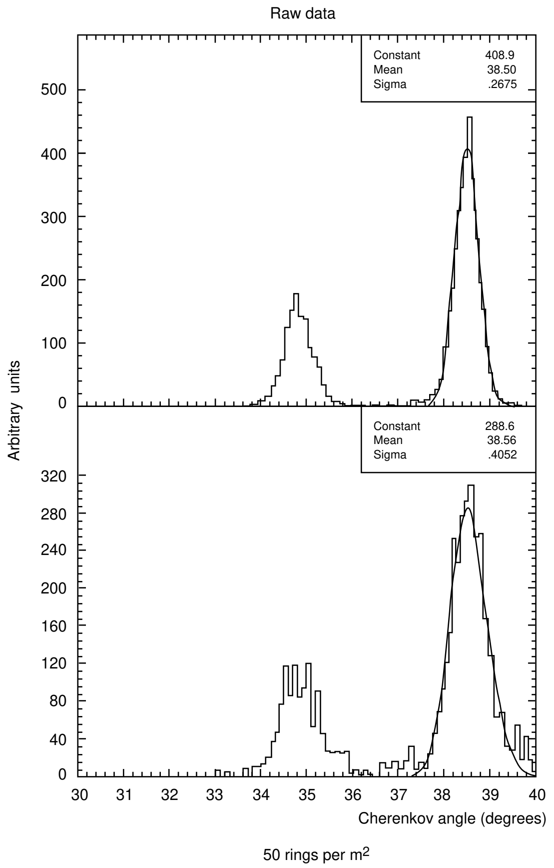

Figure 4.35: Separation obtained with pattern recognition of single events (top) and superimposed experimental events (bottom), simulating particle densities of 50 m^{-2}.

event. It is important that no selection was made in single events, except that the proton/pion ratio was artificially increased using the TOF information from the beam set-up.

The results of the pattern recognition are shown in Fig. 4.35 together with a fit to the pion distribution for the case of 50 particles per m^2. The 3σ π/K separation obtained in this simulation corresponds to 2.7 and 2.4 GeV/c for cuts with efficiencies for kaons of 81 % and 92 %, respectively. Note that the present data refer to $\beta = 1$ rings of 10 cm radius.

4.6.5 Electronics

Front-end electronics

The system will represent about 1.5×10^5 channels with a 5–8 % occupancy per event. The ASIC–VLSI pad electronics under development at RD26 is based on an analogue multiplexed readout scheme with zero suppression on the chip. The final chip (16 channels) is referred to as DIGITPLEX [48]. As shown in Fig. 4.36a, it is composed of the following functional blocks:

- the GASSIPLEX [49] for signal amplification and shaping;
- the CRIAD2 [50], an 8-bit, two-step, four-range analog-to-digital converter;
- the DILOGIC digital circuitry for local storage of the data after zero suppression.

These three elements are developed separately. The GASSIPLEX, currently available in two technologies (1.5 and 0.7 μm), is being used on the RICH prototypes. The CRIAD2 and DILOGIC chips will be produced in autumn 1996. Before starting the development of the full chip, a discrete version will be available at the beginning of 1996, referred to as MCM-DIGITPLEX (see Fig. 4.36b). The GASSIPLEX is a 16-channel analogue processor with improved shaping and a switchable filter fitted to the input current shape delivered by a MWPC to restore the d.c.-baseline level to 1 % after 3 μs. The peaking time is used as an analogue delay to provide triggers with temporary storage for the signal amplitudes by means of a TRACK/HOLD (T/H) external signal before the multiplexing sequence is activated. Table 4.5 summarizes the characteristics the GASSIPLEX.

Table 4.5: Main performance characteristics of the GASSIPLEX.

Noise (r.m.s.)	450 e at 0 pF
Noise slope	$\leq$10 e
Dynamic range	$\pm$ 2 V
Conversion gain	10 mV/fC
Baseline recovery	$\leq$1 % after 3 μs
Peaking time	300–800 ns
Power consumption/ch	$\leq$5 mW
Analog readout speed	10 MHz

The CRIAD2 [50] has been developed for the RD2-S.I.T. detector. The CRIAD2, an improved version with four 8-bit ranges, is under development at the MEAD company in Lausanne and works on the same principle as the CRIAD: a 2-bit flash ADC determines one of the 4 ranges of conversion of the 7 time-interleaved charge-ADCs which are required to perform the equivalent flash data throughput of the system. Table 4.6 summarizes the characteristics and specifications of the CRIAD.

First-level data handling

This task is performed by the DILOGIC circuitry located on the detector modules. This circuitry can be operated in two modes: (i) pedestal mode, where all channels are sequentially read out to create (externally) pedestal and threshold maps loaded alternately on the chip, and (ii) zero suppression mode where digitized data are compared with threshold maps, pedestals being subtracted if

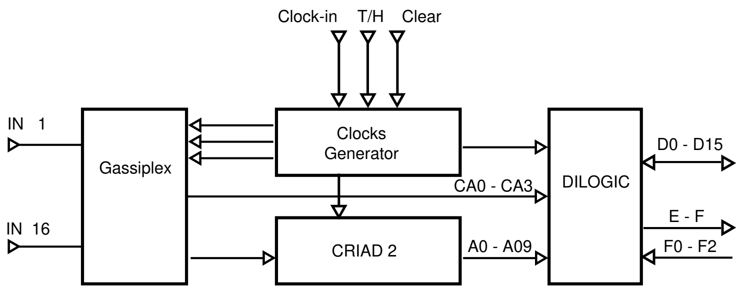

(a) DIGITPLEX: 16-channel chip composed of the analog block, GASSIPLEX, the 8-bit ADC, CRIAD2 and the logic array DILOGIC.

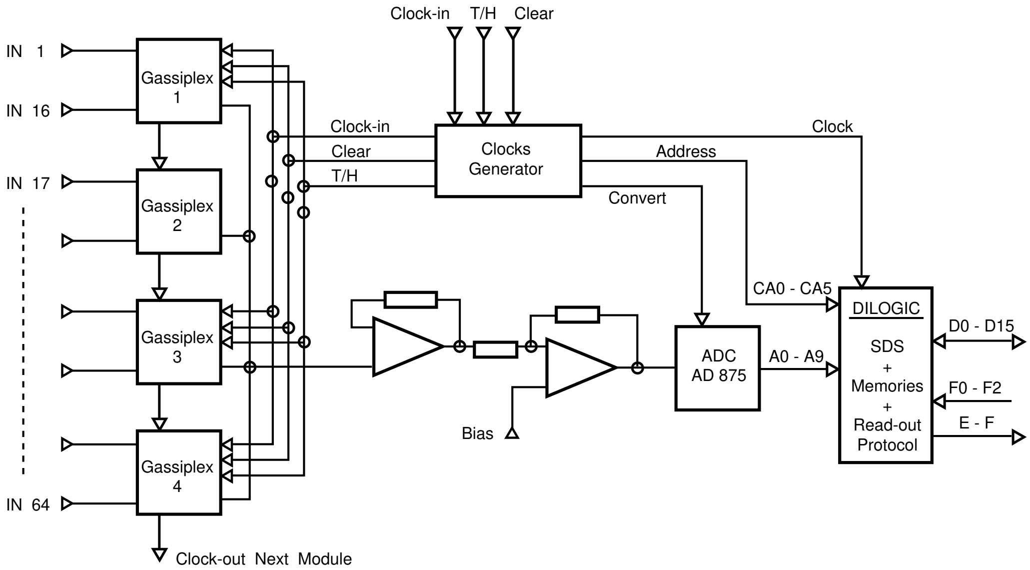

(b) MCM–DIGITPLEX: discrete version of DIGITPLEX allowing to read out up to 4 front-end analog blocks (64 channels).

Figure 4.36: Schematic diagram of the pad readout electronics.

Table 4.6: Specifications of the CRIAD chip.

Dynamic range	2048 mV
Sub-ranges	128, 256, 1024, 2048 mV
Number of bits	$8+2$
Accuracy	12 bits(1st range)
Conv. speed	35 MHz
Consumption	$\leq$50 mW
Baseline recovery	$\leq$1 % after 3 μs
Peaking time	300–800 ns
Power / channel	$\leq$5 mW
readout speed	10 MHz

data are larger than threshold. A bit-map of the event and the analogue value of the good channels is stored in FIFO memories, and finally an input/output connection is made with the read out through a dedicated protocol. A flag marks empty chips, thus allowing sparse data scan mode.

At best a complete readout operation is performed in 3–4 μs with DIGITPLEX (16 MPX steps). However, since other detectors will also use either multiplexed schemes (pixel) or long occupation time (TPC), our multiplexing sequence could be extended to the same value. In this case, the MCM scheme would be implemented, implying only a few CRIAD and DILOGIC elements associated to analogue GASSIPLEX front-end chips.

Trigger requirements

The delay for zero level trigger (L0 trigger) is given by the peaking time of the pad signals, that is 0.8–1.0 μs,

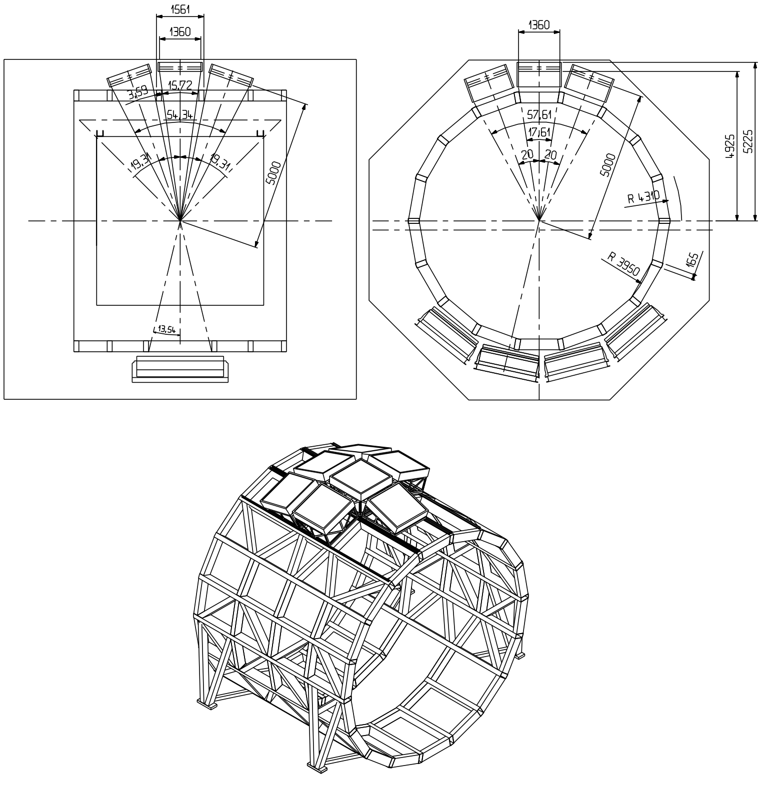

Figure 4.37: Layout of the RICH HMPID for high p_t PID.

by which time at least a pre-T/H signal has to be sent to the system. In fact, this choice avoids the use of analogue pipelining, which is expensive and has high power consumption. To be in time, a pre-trigger signal related to minimum bias events is required within about 1 μs. This signal has to be generated and transported to the system, remaining in the zone. The L0 trigger provides the necessary time to obtain from the control room a higher-level trigger validating a central collision which prolongs the HOLD signal and allows the readout sequence to take place. In the event of a negative decision, the HOLD is released, resulting in a dead time of a few μs.

Interface to data acquisition, pedestal/data modes, and calibration

Two kinds of data tranfer mode are foreseen: the pedestal mode, where the whole system is read out, and the data mode handling a small percentage of the channels. A channel corresponds to 1 and 3 bytes in pedestal and data modes, respectively. As seen in the next section, the whole detector is subdivided into modules of about 1.69 m^2 area, that is about 21 000 pad channels. In turn, a module is segmented into 72 blocks of 288 channels. Operating a series/parallel readout at 20 MHz clock rate allows block transfers into buffer memories in about 20 μs per block in pedestal mode. For a smaller event size in data mode, this figure could even be reduced by the sparse mode readout.

4.6.6 Engineering design

The HMPID will be located a long way from the vertex, at $r = 4.7$ m, in order to reduce the particle density and hence improve the performance.

The design of the RICH HMPID

To satisfy an incident angle ≤ 10–12° over the polar emission angle range, the HMPID is segmented into 7 modules, tilted according to their distance from the interaction vertex. As they are located behind the TOF barrel, care is taken to avoid positioning the RICH modules in the shadow of the supporting elements. Figure 4.37 shows the implementation of 7 modules, covering ~ 1 unit of pseudorapidity, and satisfying coverage requirements for all parameters relevant to correlation studies at high momenta.

The modules are supported by the space frame, used also by the other detectors, and can be slid into position parallel to the beam direction.

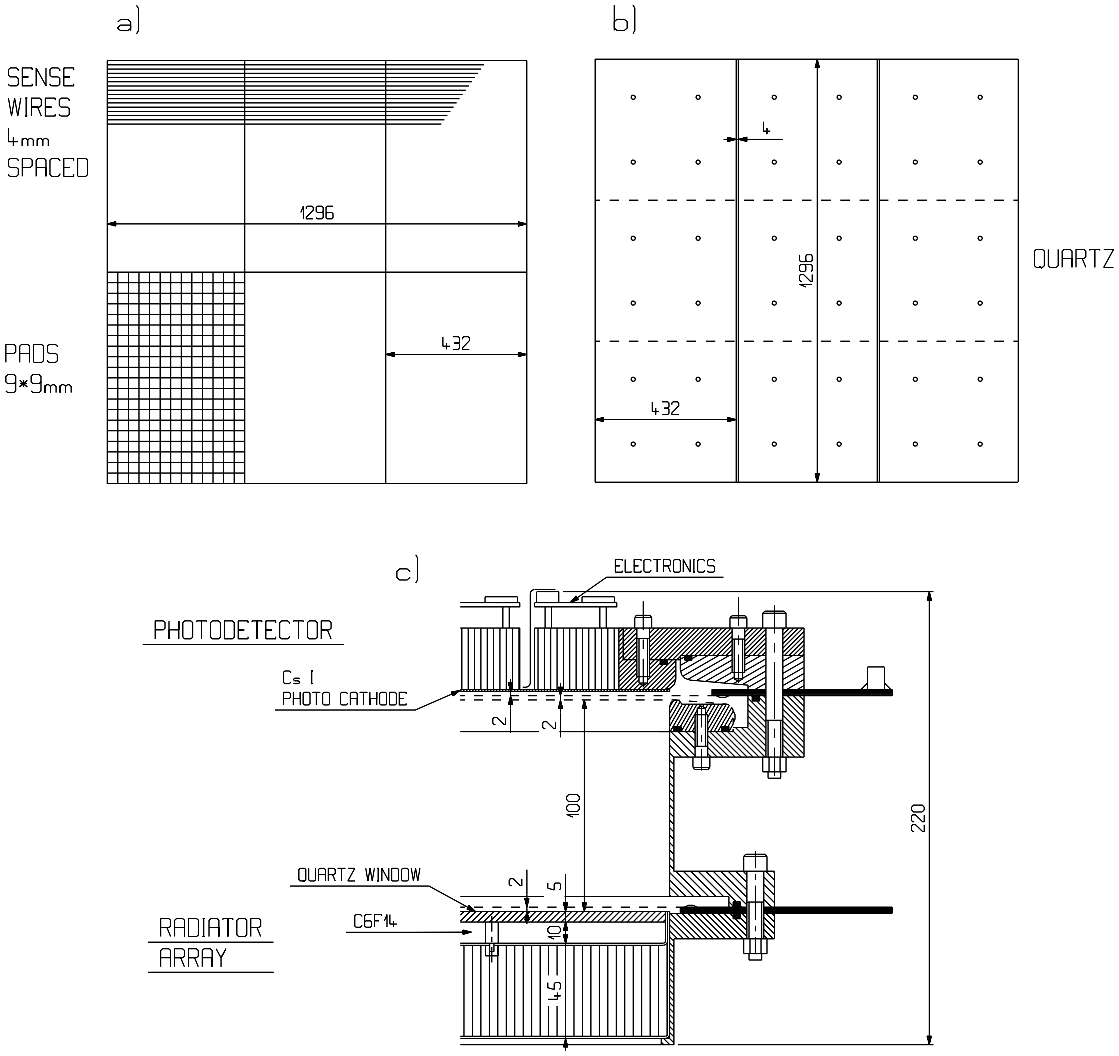

Figure 4.38: Layout of a RICH HMPID module. Schematic view of the sensitive area: a) pad and wires of the photodetector; b) radiator segmentation and assembly; c) cross-section of the RICH module.

The design of a RICH module

The layout of the RICH module is shown in Fig. 4.38. Its size is defined by the maximum feasible size of the quartz window containing the liquid radiator (about 50 × 50 cm^2). The fact that rings are truncated close to the chamber edges leads to a mean area per module of about 1.7 m^2.

The radiator

The radiator design is based on the technology in use at the DELPHI liquid barrel RICH [51]. Succesfully operated for several years, it permits a safe choice of materials guaranteeing long-term operation with stable and optimal UV transparency.

The radiator array is based on a stiff composite panel supporting three independent cavities of about 35 × 140 cm^2 circulated with C_6F_{14}. The window is assembled by gluing three quartz plates following the DELPHI procedure. The plate is supported by several inserts, enabling it to withstand the liquid weight and the hydrostatic pressure of a liquid column 40 cm high with a thickness of 5 mm and a sagitta $\leq$0.1 mm. Inefficient zones are created at the joins between cavities and the inserts, amounting to 2 % of the sensitive area.

Engineering studies and tests are in progress in order to optimize the inserts' locations for a minimum thickness of the quartz plate. In order to guarantee a long-term UV transparency performance, the liquid has to be circulated according to the scheme in use at DELPHI [52,53].

The photodetector and electronics

The photodetector is a square MWPC composed of an assembly of frames made of materials of low water absorption capacity, which are easy to outgas under vacuum (peraluman, carbon fibre).

The CsI pad cathode, made of 4 or 6 elements, is built following the composite sandwich panel technique, ensuring stiffness with a minimum amount of material (Fig. 4.39). The expected flatness tolerance is $\pm 75\ \mu$m.

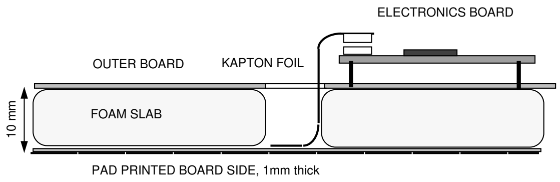

Figure 4.39: Diagram of the construction of a pad-cathode panel.

The pad PCB described in Section 4.6.3 is the internal skin of the panel, the external skin being the pad electronics. A quantity of material equivalent to 3 % of a radiation length, including electronics, is currently achieved. Automated techniques for soldering pad connections, gluing, etc, are standard.

The front-end pad electronics is made of VLSI chips mounted on boards easily accessible at the back plane of a module. A few additional cards will be mounted on the sides to provide an optical data link to the control room. The power dissipation is expected to be 80 W/m^2. As discussed in Section 4.6.3, a standard evaporation station is necessary for CsI film deposition on the pad panel.

Two frames support the sense and cathode wires, 20 and 100 μm diameter, respectively. These are followed by the proximity gap frame, supporting the last wire collection-electrode. Viton seals ensure gas tightness. Automated wiring stations will also be necessary.

Auxiliary services

A system is foreseen to independently supply each module with gas, CH_4 or CH_4/iC_4H_{10} (90/10), and nitrogen for purging. To protect the CsI PCs, the modules have to be kept under continuous gas flow (e.g. with nitrogen) between running periods. The gas system will have purification elements (oxysorb and hydrosorb cartridges), and flow impurities and control.

Slow control

Control and associated safety systems are planned for

- chamber gas circulation (automatic valve operations, inlet and outlet flow control, O_2, H_2O, C_6F_{14} trace monitoring, temperatures);
- radiator array (temperature and differential pressure);
- low-voltage supplies and electronics;
- high voltage;
- alignment.

4.6.7 R&D programme

The principle of operation of large-area fast CsI-RICH, including basic production technologies, has been proved after three years of R&D. However, the long-term stability of large detectors using this new technology under the expected operating conditions has still to be demonstrated. With regard to basic research, two areas call for further R&D. First, any possible improvement in the QE performance to reduce the separation power in the critical low momentum range; and second, a physical understanding of the ageing mechanisms. The pattern recognition capability has been extensively simulated, but the results require confirmation for heavy-ion beams at large multiplicities.

A 1 m^2 prototype is under construction which will be exposed next year to high multiplicity events in a heavy-ion line and to complete a programme to study ageing under real conditions. Then, a full-size module and complete electronics could be developed in 1997. Engineering studies should be pursued at the same time on the radiator and the large-scale production of the photodetectors.

4.7 High-momentum PID with scintillators

Scintillators read out by photo-multipliers were mentioned as an option for the barrel PID array [54] in the ALICE LoI, but are no longer considered, mainly for reasons of cost. However, for the small area of the HMPID, they can be considered as backup in case of problems with the RICH detector and if, in addition, the barrel PID does not reach a system resolution of 100 ps.

The use of fast organic scintillators has a number of advantages, in particular regarding count rate capability, light yield, and the fact that delicate handling before and during measurements is not necessary. The time resolution can be better than 100 ps, depending on the primary light yield and the transmission properties of the optical coupling.

4.7.1 Design of the tile module

In the LoI, we considered a ‘grid’ option with long scintillator slats, arranged in four layers and read out with light pipes into a multi-channel PMT. However, tests have shown that not enough light can be transmitted in this arrangement to achieve a good timing resolution. Therefore we only consider a 'tile' geometry for the HMPID.

The tile geometry is based on scintillator tiles of area 4×4 cm^2 (10 % occupancy at $dN_{ch}/dy = 8000$), each coupled to a fine-mesh photo-multiplier tube (PMT). PMTs need to be of the fine-mesh type because they will have to operate in a ~ 0.2 T magnetic field (up to 0.5 T in special runs).

The tiles are mounted on modular aluminium support trays; each module contains 1140 tiles, giving a total of 7980 channels, assembled with the PMT axis at 90° with respect to the beam direction. The total sensitive area is 12.7 m^2, arranged in 7 modules of 1.82 m^2. The overall size of a single module, including the frame structure, will be 1.4×1.4 m^2. The geometry of the detector is the same as planned for the RICH option for the HMPID, as described in Section 4.6.6.

The module consists of a frame with a number of wedges, which form the supporting plane for the scintillators. Each wedge supports the tiles, their PMTs and the front-end electronics. The Mylar which covers the scintillators separates each tile in this plane from adjacent tiles. Tiles and PMTs are fixed to the module by means of a system of retainers.

The main parameters of the TOF scintillator tile option are summarized in Table 4.7.

Table 4.7: Summary of the Scintillator TOF parameters.

Number of tiles	7980
Tile size	$4 \times 4 \times 1.5$ cm^3
Module size	1.4×1.4 m^2
Tiles per module	1140
Material	NE102A or Pilot-U
PMT	fine-mesh tubes

4.7.2 R&D on TOF with scintillators

Test measurements have been carried out using both cosmic rays and the 7 GeV proton beam of the AGS at BNL (USA) [54]. We used NE102A scintillator slats ($45 \times 20 \times 7$ mm^3, 0.6 s rise time, 2.4 s decay time) covered by aluminized Mylar and coupled to a R1398 Hamamatsu PMT. Two detectors in coincidence provide the start and stop signals.

The resolution measured at the AGS with two scintillators in coincidence, each of 20 mm thickness, is shown in Fig. 4.40. The best result obtained was about 95 ps.

The main requirements for a suitable fine-mesh PMT are low cost, dimensions of about 1 inch per channel, QE of around 25 %, gain of 10^6 at $B = 0$, transit time spread < 0.5 ns. An R&D project has been started, in collaboration with Russian industry, to develop a low-cost PMT

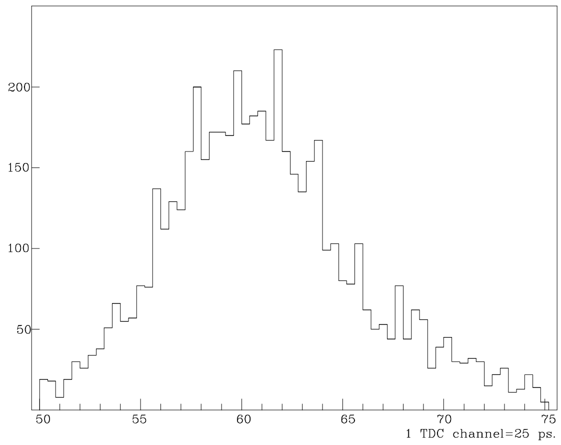

Figure 4.40: Time resolution measured at the AGS with two scintillators in coincidence. The horizontal axis is in TDC channels (25 ps per bin).

meeting these specifications [55]. First prototypes are expected in early 1996.

The high voltage for the PMTs will be supplied using d.c./d.c. converters. We could either use a dedicated unit for each channel close to the PMT base or use a converter capable of supplying voltage to 10 different PMT channels. The latter is foreseen for the high-voltage distribution system of the PHENIX detector at RHIC/BNL.

5 The Photon Spectrometer (PHOS)

5.1 Design considerations

The PHOS detector has been designed to search for direct photons. Their production and p_t dependence reflect, to a large extent, the initial conditions which will occur in heavy-ion collisions at the LHC.

The second goal of the PHOS is to measure π^0 and η production at the highest momenta, where the momentum resolution is an order of magnitude better (at 25 GeV) than for charged particles measured in the tracking detectors. The spectrum of high momentum ('leading') particles gives information about the propagation of jets in the dense medium created during the collision ('jet quenching').

The direct photons are embedded in a large background of photons from hadron decays (mainly π^0 and η). The ratio of direct to decay photons is of the order of 5 % in the most conservative estimate, but could well be considerably larger up to a few 10 % (see Section 11.7). We expect to achieve a sensitivity ≤ 5 %, dominated by systematic errors mainly in the γ and π^0 reconstruction efficiencies. If the yield of direct photons has a stronger than linear dependence with multiplicity, the sensitivity can be increased by comparing central and peripheral collisions, because then systematic errors cancel to some extent. A close to quadratic dependence with multiplicity could be expected e.g. for thermal photons.

In order to reach the required sensitivity, we have to measure the rates and the p_t spectra of photons, π^0 and η mesons in the same apparatus. The expected extremely high multiplicity environment implies a high segmentation of the calorimeter, the largest possible distance to the vertex, and use of a very dense active medium with the smallest possible Molière radius.

In order to have adequate space resolution and to allow separation of overlapping showers, the transverse cell size should be of the order of the Molière radius R_M [1]. Little shower overlap is needed in order to ensure reliable reconstruction of photons and mesons.

The particle occupancy should not exceed ~ 3 % (which corresponds to ~ 15 % cell occupancy) in order to be able to determine the photon reconstruction efficiency with high accuracy. At this occupancy, as shown by the SPS heavy-ion experiment WA80 [2], the photon reconstruction efficiency can be measured to an accuracy of ~ 2 %.

A material with a Molière radius of ~ 2 cm is needed for the PHOS spectrometer (at $r = 4.6$ m) in order to keep the occupancy within the required range in central Pb–Pb reactions (up to $dN_{ch}/dy = 8000$).

In addition to photons, the PHOS also responds to charged hadrons. It is therefore necessary to include a charged-particle veto detector in front of the PHOS detector (see Section 5.2.5).

Similarly, the PHOS detects K^0_L, n and $\bar{n}$. All of them are produced abundantly at the LHC. The most disturbing in terms of energy deposited in the PHOS are the antineutrons. Neutral hadron rejection (see Section 11.7) can be achieved by a cut on the shower width, which operates at all energies, and/or by a cut on time of flight. Indeed, over a distance of 4.6 m, a particle of Lorentz factor 2.8 will arrive 1 ns later than a photon. This covers a relevant energy range, especially for antineutrons. Therefore an option of nanosecond time-of-flight measurement in PHOS is currently being investigated, and will be described below.

Finally, the PHOS detector must be able to operate in the full L3 magnetic field up to ~ 0.5 T, and it should be compact enough to be integrated into the ALICE set-up.

5.1.1 Lead-tungstate crystal

In our LoI, we have identified lead tungstate ($PbWO_4$) crystals as the most promising material for the PHOS detector. Today, it is still the only material which fulfils our requirements concerning Molière radius, light yield and cost. This crystal has also been chosen by the CMS Collaboration for their electromagnetic calorimeter [3]. However, the requirements of the two experiments are rather different.

The possibility of using this crystal for electromagnetic calorimetry has been studied in recent years [4–9] and very significant progress has been made towards its use in large calorimeters.

Table 5.1: $PbWO_4$ properties.

Density	8.28 g/cm^3
Radiation length	0.89 cm
Interaction length	19.5 cm
Molière radius	2.0 cm
Melting point	1123°C
Hardness	4 Moh
Refractive index along z axis ($\lambda = 632$ nm)	2.16
Hygroscopicity	none
Chemical activity	inert

Its physical and chemical properties are given in Table 5.1.

Lead tungstate is a scintillating monocrystal with a rather complex emission spectrum, including two broad bands at 440 nm and 530 nm. The decay time constant has a mean value of about 10 ns at room temperature.

The light yield of lead tungstate crystals at room temperature is ~ 5 % of that of BGO, i.e. 50–80 photons/MeV for a ~ 20 cm long crystal. However, the light yield is constantly being improved and the most recent crystals produced for CMS have reached more than 100 photons/MeV. As shown in Fig. 5.1, a rather good longitudinal response homogeneity was measured for a crystal coupled to a photomultiplier and exposed to a ^{60}Co source, which was displaced along the crystal.

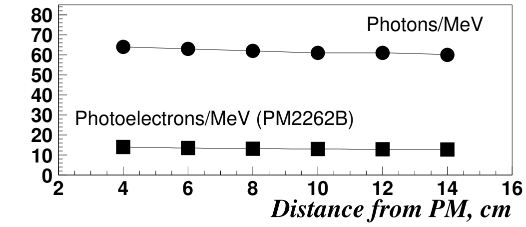

Figure 5.1: Light yield of a $PbWO_4$ crystal and the resulting signal of the photomultiplier.

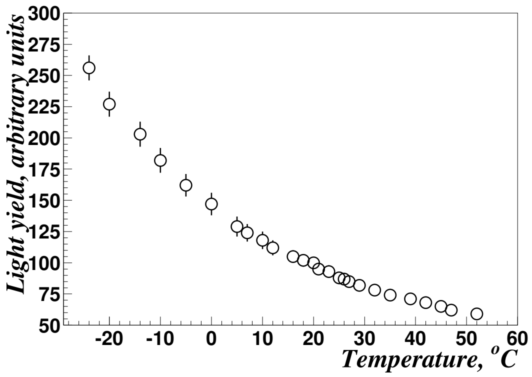

Figure 5.2: Light yield of a $PbWO_4$ crystal vs. temperature, normalized to 100 at 20°C.

The light yield depends strongly on the temperature. The temperature coefficient is ~ -2 % per °C at 20°C and increases to ~ -4 % at -25°C, as shown in Fig. 5.2 (from Ref. 7).

The temperature dependence offers the possibility to increase the intrinsic light yield of the crystal. At a working temperature of -25°C, the light yield will be enlarged by about a factor 2.5 and furthermore the electronics noise of the photon detector will be reduced. Both effects will lead to improved energy resolution. The feasibility of this method was demonstrated in the first test with a cooled lead tungstate matrix (see Section 5.3.1). We have chosen -25°C as the working temperature for PHOS, which can be obtained easily and economically. This temperature should be stabilized with ~ 0.1°C precision.

Lead tungstate ($PbWO_4$) is grown from a 50 %–50 % mixture of lead oxide (PbO) and tungsten oxide (WO_3). Homogeneous crystals can be grown to the required size by the Czochralski technique in platinum crucibles. The total time necessary to produce a ~ 20 cm long ingot is ~ 130 hours. A large-scale crystal production technology exists, e.g. in Russia and in the Ukraine.

Considerable variations in optical properties are induced even by small changes in the crystal production process. The crystal growing and annealing technology is being optimized in close collaboration with industry and research institutes specialized in crystal growing technology.

5.1.2 Crystal length

Our primary choice for light readout are silicon PIN diodes. However, PIN diodes respond to charged particles traversing them, and the energy resolution can be spoiled by this shower punchthrough effect. Figure 5.3 shows the response to 10 GeV electrons for crystal lengths of 16 and 20 cm. The tail visible in the upper part of the spectrum (channel number $\geq$ 1200) is caused by punchthrough, i.e. energy deposited in the diode by charged particles rather than by light. For the photons of interest to ALICE ($p_t \leq 10$ GeV$/c$), the effect is acceptably small at a crystal length of 20 cm. It would be eliminated altogether by placing the PIN diodes in front of the crystals. The coupling to the readout electronics would be more complicated, but the required resolution could be obtained with shorter crystals. However, in this case backsplash from the crystals into the PIN diode is of concern and will have to be studied.

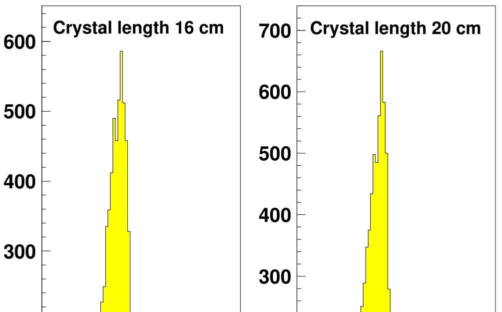

Figure 5.3: Energy spectra from $PbWO_4$ crystals with different lengths irradiated with 10 GeV electrons.

Longitudinal segmentation of the detectors is known [10] to improve essentially the separation of neighbouring pairs of showers (both photon–photon and photon–hadron), without significant degradation of the energy resolution. However, the detector module structure becomes more complicated, the charged-particle veto detector will be somewhat farther from the PHOS spectrometer, and the number of the photo-readout channels will increase by a factor of 2.

Thus, 20 cm long single crystals with transverse size $2.2 \times 2.2\,\mathrm{cm}^2$ have been chosen as a PHOS baseline. A final decision concerning longitudinal segmentation and location of the PIN diodes will be taken only after extensive simulations and beam tests foreseen for 1996.

5.1.3 Acceptance

The π^0 and η yields and spectra will be measured via their two-photon decays. The measurement has to be done with special care, i.e. high and well-known accuracy, in the 1–10 GeV/c momentum range, to determine precisely the background for the measurement of direct photons in the same p_t range.

The maximal p_t for separating the two photons at the distance of PHOS is $\sim$50 GeV/c. The minimal p_t depends, for a given distance $r = 4.6$ m from intersection point to detector, on the size as well as on the geometrical configuration of the detector. Indeed, for γ–γ decays, there is a minimal opening angle $\psi_\mathrm{min} \simeq 2 \times m/p_\mathrm{t}$, which is also the most probable one. For a detector of given area and rectangular shape, it will depend on the length-to-width ratio L/W. A larger L/W ratio will increase the smallest accepted ψ_min, and hence, the geometrical acceptance at low p_t. However, the ratio of the detector perimeter to its area also increases, thus decreasing the plateau acceptance.

This behaviour is demonstrated in Fig. 5.4, where the geometrical acceptance for π^0 and η are plotted vs. p_t for various values of L/W.

The geometrical acceptance has been calculated by means of a Monte Carlo simulation. The only parameters entering this simulation are the detector boundaries (geometry) and the threshold for photon detection. A plane detector of 18 m^2 total area located at 4.6 m from the intersection point, and a photon energy threshold of 250 MeV were assumed. As can be seen from the figure, the minimal p_t depends only slightly on the L/W ratio for π^0, but rather markedly for η mesons.

The L/W ratio has been chosen to be $\sim$4. At this ratio the minimal accepted p_t is $\sim$0.5 GeV/c for both π^0 and η. In spite of the varying incidence angles of the photons, the planar geometry of the PHOS hardly affects the energy and space resolution. A study by the PHENIX Collaboration [11] has shown little change in energy resolution up to 20°, for a highly segmented lead glass calorimeter exposed to a 1–4 GeV electron beam. Spatial resolution is important for the π^0 and η mass resolution. For small incidence angles the mass resolution does not degrade substantially, e.g. the π^0 mass resolution of $\sim$10 MeV in the

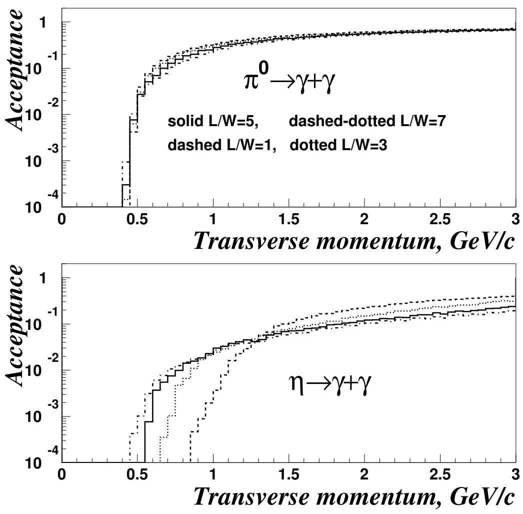

Figure 5.4: Geometrical acceptance for π^0 and η vs. p_t for a rectangular detector of 18 m^2 total area and various length-to-width ratios L/W.

WA80 experiment does not depend on the incident angle up to the maximum of 7°.

We therefore plan to divide the PHOS detector azimuthally into four roughly quadratic planar sub-units, each facing the interaction point, so that incidence angles will nowhere exceed $\sim$13°. This simple planar geometry facilitates the production of crystals, mechanics, and support structures. We will, however, continue to investigate the implications of this geometry on double shower separation and cost reduction; therefore a pointing geometry, with crystals shaped like truncated pyramids, is not yet ruled out.

The π^0 and η acceptance for this configuration, assuming a 3 cm space between sub-units for the support structures, are shown in Fig. 5.5.

5.1.4 Energy and position resolution

The sensitivity to direct photons depends on the accuracy of the measured π^0 and η yields, which in turn is determined by the energy and space (i.e. opening angle) resolution of the PHOS.

The π^0 and η yields will be extracted from the $\gamma\gamma$ invariant mass spectrum, on top of a large combinatorial background.

The signal-to-background ratio S/B is inversely proportional to the mass resolution σ_M. A better mass resolution entails higher statistical significance and lower systematic errors for the meson yields and spectra.

The mass resolution is given by the following expres-

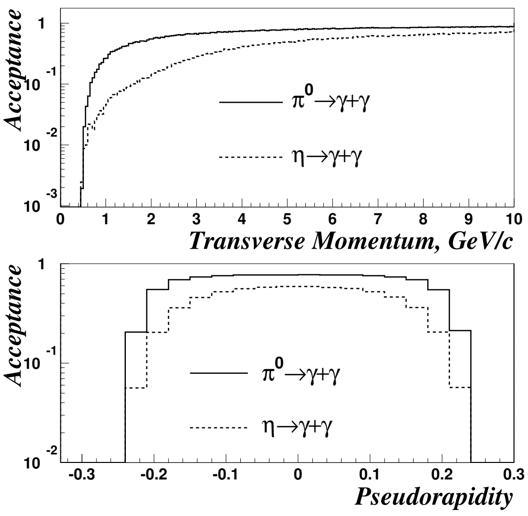

Figure 5.5: The π^0 and η acceptance for the chosen geometrical configuration of the PHOS spectrometer.

sion:

$$\frac{\sigma_M}{M} = \frac{1}{2}\sqrt{\frac{\sigma^2_{E_1}}{E_1^2} + \frac{\sigma^2_{E_2}}{E_2^2} + \frac{\sigma^2_{\psi}}{\tan^2(\psi/2)}}$$

where E_1 and E_2 are the photon energies and ψ is the opening angle between the photons, whose error is given by the position error.

The energy resolution can be parametrized as (E in GeV):

$$\frac{\sigma_E}{E} = \sqrt{\frac{a^2}{E} + \frac{b^2}{E^2} + c^2}$$

where a is the stochastic term, b is determined by the readout noise, and the constant term c is due to detector and readout inhomogeneity and to the calibration error. Since PHOS will measure relatively low-energy photons ($E_\gamma \leq 10$ GeV), the first two terms are more important than the third [12].

The stochastic and constant terms a and c have been measured to be around $a = 3\,\%$ and $c < 1\,\%$ [13], with some of the more recent results as low as $a = 2.8\,\%$ and $c < 0.5\,\%$ [14].

If the electronic noise level can be kept as low as ~ 15 MeV per channel, then the contributions of the stochastic and noise terms to the photon energy resolution would be approximately equal at 1 GeV. Since on average ~ 5 channels will be fired by one photon, one has $b \approx \sqrt{5} \times 15$ MeV ≈ 33 MeV (b increases only logarithmically with E).

With a light yield of ~ 100 photons/MeV at 20°C one could expect up to ~ 250 photons/MeV at a temperature of -25°C. PIN photodiodes with an active area of 1×1 cm^2 and quantum efficiency of 70 % would then give 36 electrons/MeV. A 15 MeV r.m.s. noise would correspond to ~ 600 electrons, which is feasible even at room temperature with existing PIN photodiodes and preamplifier (see below).

Cooling of the electronics, which is placed close to the crystals, as well as ongoing R&D for PIN photodiode and preamplifier optimization will lead to further improvement of the noise term.

The statistical accuracy of π^0 and η yield measurements is determined by the ratio $S/\sqrt{B}$, where S is the neutral meson signal (yield) and B is the combinatorial background under the neutral meson peak in the invariant mass spectrum. Higher $M_{\gamma\gamma}$ resolution will lead to better significance for a given number of events.

We have evaluated the number of events needed to reach a given accuracy for the meson cross-section with the SHAKER generator and a GEANT detector simulation. The space resolution was parametrized as $\sigma_x = 50$ mm $\times$ σ_E/E. The parameters, $a = 3\,\%$ and $c = 1\,\%$, were taken from measurements of the $PbWO_4$ energy and space resolution, done with photomultipliers in an electron beam test at high energies (10, 25 and 100 GeV) [13]. The readout noise parameter b was given four values: $b =$ 30, 60, 90 and 120 MeV.

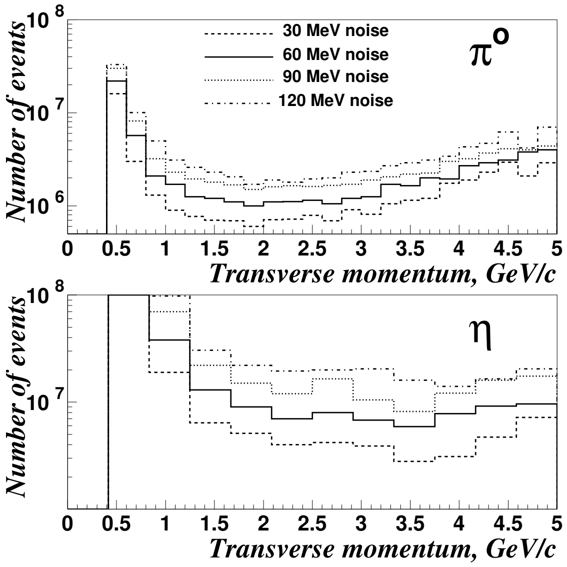

Figure 5.6: Total number of events necessary for achieving a statistical accuracy of 1 % per 200 MeV/c bin in the π^0 yield (top figure), and 10 % per 400 MeV/c bin in the η yield (bottom figure), plotted vs. transverse momentum.

The results of the simulation are presented in Fig. 5.6. It shows, for each p_t bin, the total number of events which must be accumulated if one wants to achieve 1 % statistical accuracy in the π^0 yield and 10 % accuracy in the η yield (see Section 11.7).

In ALICE we expect to collect about 5×10^7 central Pb–Pb events per year. Even for a noise level as high as

120 MeV, the π^0 yield can be extracted with better than 1 % statistical accuracy. On the other hand, the parameter b should be ≤ 60 MeV to obtain the desired accuracy of 10 % in the η meson yield.

5.2 Engineering design

5.2.1 Mechanical structure

The basic element of the PHOS detector is a $PbWO_4$ crystal of size $2.2 \times 2.2 \times 20$ cm^3, wrapped in Tyvek and coated with black shrinktube. A PIN diode is glued onto the outer face of the crystal.

Every 64 crystals arranged in a 8×8 square will form a module of size $\sim 18 \times 18 \times 48$ cm^3, including light detectors and preamplifiers. The module will have 0.5 mm thick walls made of a folded aluminium plate fixed to the aluminium support. The PHOS module is shown in Fig. 5.7.

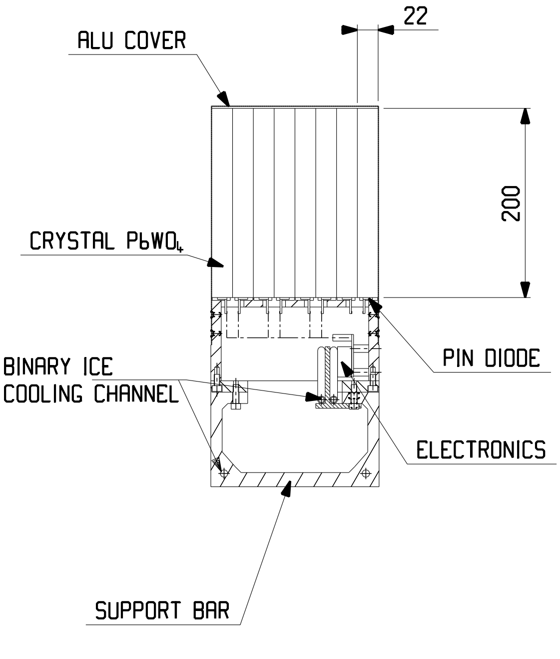

Figure 5.7: PHOS module.

Thirteen modules will be mounted in line on a 2.314 m long, U-shaped aluminium bar parallel to the beam axis. Eleven bars, positioned alongside each other, will form an independent sub-unit, of approximately square shape, called a cradle. The PHOS detector will consist of a total of 36 608 crystals, grouped into four such cradles (see Fig. 5.8).

The cradles will be completely independent of each other, thermo-insulated, and mounted on rails to allow them to be moved in and out of the ALICE set-up. Adjacent cradles will be separated by gaps of ~ 3 cm, taken up by the mechanical support structures.

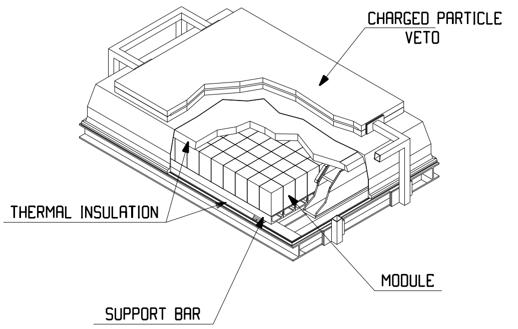

Figure 5.8: Detailed view of modules and bars inside a cradle.

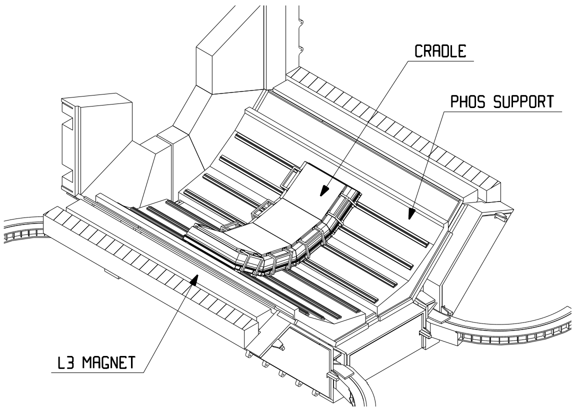

Figure 5.9: Layout of the PHOS detector.

The layout of the PHOS detector is shown in Fig. 5.9. It will be positioned on the bottom of the ALICE set-up, and cover half a unit in pseudorapidity, $-0.27 \leq \eta \leq 0.27$, and 106° in azimuthal angle φ. Its total area will be ~ 18 m^2.

Each of the four cradles will be positioned with its front face pointing towards the vertex as shown in Fig. 5.9, implying a maximum angle of incidence for photons of $\sim 13°$. The cradles will be positioned in a semicircle around φ (length to width $L/W \approx 4$).

A charged-particle veto detector will be mounted in front of each cradle. The distance from the intersection point will be 4.41 m to the veto detectors, and 4.60 m to the front face of the $PbWO_4$ modules.

The heat-producing electronics will therefore be separated from the modules and mounted inside the bars. The bars will be cooled with binary ice. Only the PIN diodes, which have very small heat output (≤ 2 μW per channel), will be in close contact with the $PbWO_4$ crystals (see Fig. 5.7). Thermoinsulation of each cradle as a whole will help to keep the inside temperature stable around -25°C with ~ 0.1°C accuracy. Binary ice was chosen as a coolant because it makes it possible to maintain low temperature with high precision. The cooling system will be designed for a 5 kW heat dissipation.

5.2.2 Photodetector

All detector sub-systems of the ALICE experiment are to be operated inside a magnetic field of up to ~ 0.5 T. This requires a readout which is able to operate in such a field. Two options for the photodetector are still under consideration.

The PIN photodiode with amplifier is being considered as the best candidate. It is compact and insensitive to magnetic fields. The present design makes use of 10×10 mm^2 PIN diodes, but can easily incorporate larger area diodes. The choice will mainly depend on the development of future costs.

The use of PIN diodes presents, however, certain disadvantages. They are sensitive to charged particles, so even a small shower leakage can worsen the energy resolution. The noise level is higher than with vacuum photodetectors. Furthermore, the PIN diode has no intrinsic amplification, so that the readout line has to include a preamplifier and an amplifier.

The alternative option would be a vacuum phototetrode (or photopentode), which is much less sensitive to magnetic fields than the standard photomultiplier. However, partial sensitivity to magnetic field still exists; phototetrodes are also more expensive than PIN diodes.

Different photodetectors will be tested at CERN in 1996, and the final decision will be based on these results.

5.2.3 Calibration and monitoring

In order to keep PHOS energy and space resolution at the required level, the detectors and readout electronics have to be calibrated and monitored.

Each crystal of the PHOS will be calibrated before installation with a relatively low-energy (few GeV) electron test beam at several electron energies. The calibration will be carried out in parallel with the crystal production, i.e. each cradle will be calibrated separately after assembly.

The calibration will need a total of ~ 12 months of test beam time distributed over 3 years (see Section 12.1).

Additional *in situ* calibration will be done during data-taking using crystal amplitude spectra. The truncated mean of the amplitude spectrum of each crystal can be used for relative calibration, and the π^0 peak position for absolute calibration.

Monitoring will be carried out constantly during data-taking.

The following values will be monitored:

- The temperature of the lead tungstate crystals with $\sim 0.1^\circ$C accuracy. One measurement per module (64 crystals).
- The gain of the readout electronics. For this purpose each preamplifier will have a calibration input. One measurement per crystal.
- Transparency of the lead tungstate crystals and their optical contacts with the PIN diodes. This value will be monitored by calculating the truncated mean of the spectrum of each crystal for a certain time (typically 5–20 minutes) and comparing it with previous measurements. One value per crystal.

If the PIN diodes are chosen as the photodetectors, there will be no need for gain monitoring. However, if the vacuum phototetrodes are chosen for light detection, additional monitoring systems will be needed.

5.2.4 Readout electronics

Several options for the readout electronics are under consideration, depending on the ability to deliver a Level 1 (L1) trigger signal in time and the need for TOF information from PHOS. While currently it looks feasible to generate a trigger sufficiently fast for L1 signals, the timing and cable delays have to be better understood before we can be completely confident (see Chapter 9). The choice concerning TOF readout will be made after completion of ongoing simulations concerning antineutron rejection (see Section 11.7).

The option including timing and L1 trigger information is presented in this section.

The readout electronics will consist of a multichannel, custom integrated circuit, board-level discrete electronics, and programmable devices. A block diagram of the system is shown in Fig. 5.10. The custom integrated circuit will provide energy, timing, trigger sums, and ADCs. The board level electronics will include a serial interface circuit, a board controller, and a data readout interface circuit which includes communication to the local data collection modules and L1 trigger system. Many of the same functional blocks were developed for the WA98 lead-glass calorimeter which is described in detail in Refs. 15 and 16.

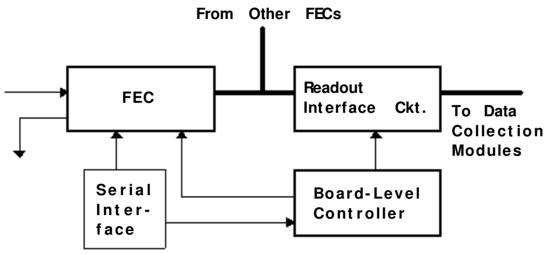

Figure 5.10: System block diagram.

Front-End Chip

A block diagram of the front-end chip (FEC) is shown in Fig. 5.11. The function of the FEC is to generate energy, timing, and trigger sum signals from the signal provided by the charge-sensitive preamplifier described in Section 5.3.4. The shaping amplifier is used to improve the signal-to-noise ratio of the input signal before digitization. Figure 5.12 shows the expected noise as a function

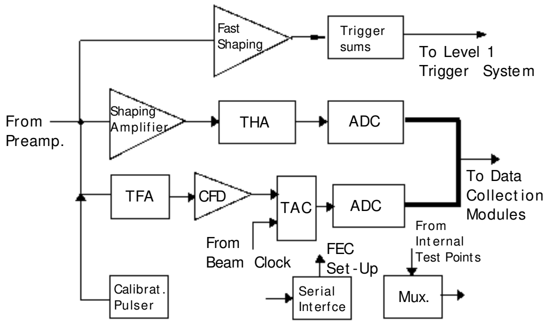

Figure 5.11: Block diagram of Front-End Chip for the option with timing.

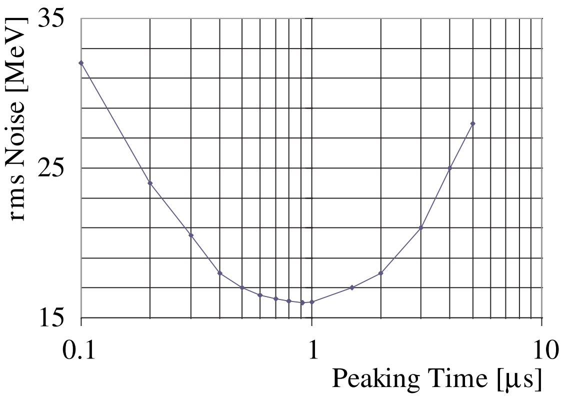

Figure 5.12: Equivalent input noise vs. shaping amplifier peaking time.

of shaping amplifier peaking time. For a detector capacitance of 70 pF, a peaking time of approximately 1 μs allows for an acceptable output noise level of ~ 15 MeV (r.m.s.) without generating excessive amounts of dead-time.

The peak signal from the shaping amplifier could be obtained by taking multiple samples near the peak and reconstructing the pulse shape, or by stretching the peak signal with a track-and-hold amplifier. While the former method simplifies the complexity of the circuit somewhat, the data bandwidth for the energy channels would increase by a factor of 3–5. The track-and-hold (T/H) amplifier proposed for this system has recently been developed for use with germanium strip detectors [17]. The performance of this circuit is summarized in Table 5.2. In addition to holding the peak value of the shaped signal, the T/H amplifier generates a peak detect signal that will be used in the zero-suppression logic.

The charge-sensitive output of the preamplifier is not suitable for performing the timing measurement. Two approaches are being considered for developing an adequate timing signal, i.e. use of a timing filter amplifier (TFA) or use of a separate timing signal generated by the preamplifier.

Table 5.2: T/H amplifier performance results.

Dynamic range	100 : 1 (25 mV min)
Power consumption	0.75 mW
Drop rate	8.25 mV/ms
INL	$< \pm 0.05$ %
Pedestal voltage	< 35 mV

The TFA approach has the advantage of requiring only one input pin on the FEC for the timing and energy signals. However, at the low end of the energy range, the timing jitter will be determined by the risetime of the signal. To obtain sub-nanosecond timing resolution, a fast timing output from the preamplifier may be required.

Two topologies are possible for the timing discriminator: leading-edge (LE) and constant-fraction (CFD). Low-walk, integrated, CMOS implementations of CFDs have recently been developed [18–20]. Although the walk from a CFD might require offline correction, the correction would be much smaller and less sensitive to energy channel noise than the large slew correction required for an LE discriminator. The discriminator threshold will be adjustable on each chip via the serial interface.

The time-to-amplitude converter (TAC) will have a gated charge integrator topology. All the TACs will use the beam clock as a stop signal while the discriminator outputs will start each TAC conversion. This common-stop approach keeps the TAC dead-time to a minimum and reduces the amount of clock feedthrough to the TAC output.

There will be one ADC channel for each TAC and energy output to allow for parallel conversion. The ADC will be of the 11-bit Wilkinson type. The conversion clock will be disabled, except during the conversion, to limit power consumption and clock-edge coupling during the signal acquisition. A prototype version of the ADC has been evaluated and is shown as a block diagram in Fig. 5.13. Early test results obtained with the prototype are shown in Fig. 5.14 as well as Table 5.3.

Table 5.3: ADC Performance Results.

Clock	100 MHz
INL	0.15 %
Channel Width	1.5 (r.m.s.)
DNL	$\sim$ 1/2 LSB

Non-overlapping 2×2 energy sums (the sums of the signals in four adjacent detectors) will be formed on each FEC. These 2×2 sums will be used to form overlapping 4×4 sums that will be sent to the L1 trigger system. To be useful in the L1 trigger, these sums will be formed from quickly peaking signals rather than the energy signals. Since they are overlapping, one 4×4 sum will be collected on each chip for each 2×2 sum formed on the

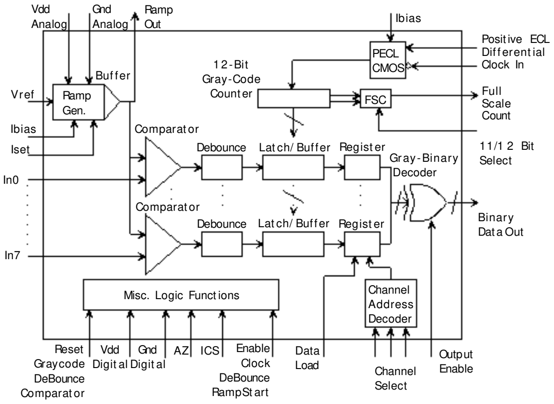

Figure 5.13: Block diagram of ADC.

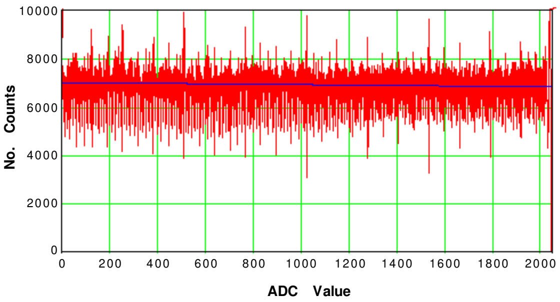

Figure 5.14: ADC differential nonlinearity.

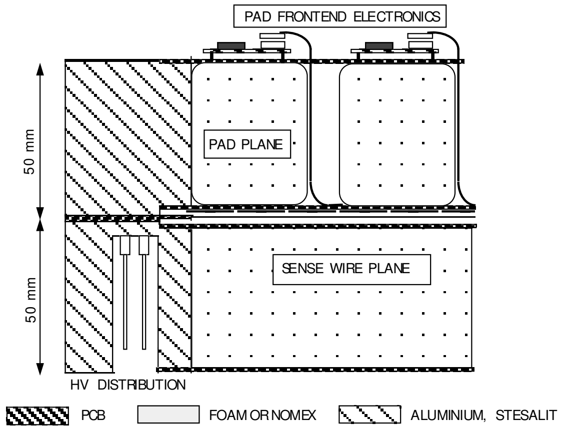

Figure 5.15: Veto detector structure.

chip. This will require that copies of 2×2 sums be communicated across chip and board boundaries to collect all of the 4×4 sums. A differential current mode communication scheme will be used to preserve speed and limit ground-loop possibilities.

A calibration pulser will be included on the front-end chip that will exercise all the energy and timing channels. This feature helps in the pre-installation testing of the chips and boards and is also useful in calibration and testing of installed boards. The amplitude of the pulser and the pattern of channels to be exercised will be controlled on each chip via the serial interface. A multiplexer that allows the probing of an important internal signal on the FEC will also be included for diagnostics purposes. The multiplexer will be controlled via the serial interface.

Board-level electronics

The board-level electronics handle the communications between the readout system and the L1 trigger, the data collection modules, and the master timing and control circuits. It is advantageous to amortize these functions over as many channels as possible. In this sub-system we propose that one circuit board contains 2 supermodules or 128 channels. It is anticipated that the board-level functions can be carried out by a few programmable logic devices and a small amount of supporting circuitry.

The data readout interface circuit handles the transfer of digitized data from the FEC chip up to the data collection modules. This entails the proper sequencing of the data, header and trailer generation, and control of the data bus.

The board-level controller (BLC) controls the T/H amplifier and the ADC. The BLC receives the L1 trigger signal, starts the conversion, and transmits a ready signal to the data readout interface circuit after the conversion is complete. The T/H amplifier is then cleared in preparation for the next signal.

The serial interface circuit is used to receive set-up and control instructions such as discriminator thresholds and calibration pulser parameters. Operational-mode selection (data-taking, calibration, diagnostics, etc.) will be transmitted to the BLC via the serial interface.

Using WA98 calorimeter readout electronics [15,16] as a guide, the power consumption for a 128 channel board is expected to be ~ 18 W. The board will have direct thermal contact to the binary ice cooling system.

5.2.5 Charged-Particle Veto

The requirements on the Charged Particle Veto (CPV) system for the PHOS detector are the following:

- large-area modules of $\sim 230 \times 196$ cm^2,
- close to 100 % detection efficiency for charged particles,
- two-dimensional readout with localization accuracy of 3-4 mm,
- minimum dead-zone,
- minimum amount of material, ≤ 5 % of X_0, electronics included.

The CPV can be made using known techniques of MWPC chambers with pad readout and does not present any major design difficulties.

Detector layout

As shown in Fig. 5.15, the MWPC is composed of two panels, made of a composite structure of two thin printed-circuit layers (1 mm PCB) glued on both faces of a light panel (foam of 30 g/l density) surrounded by a thin frame. Such a structure allows for the production of large-area panels of adequate flatness (50–100 μm) since the thickness of the foam plate has a stiffness able to stand the stress from the sense wires. The MWPC geometry is chosen as follows: 5 mm for the distance from anode to cathodes and 4 mm between sense wires (∅ 25 μm). Given the size of a module, support lines or guirlands have to be added in order to keep the distance between sense wires and cathodes constant.

Two-dimensional readout is achieved by segmenting one of the PCB into pads and detecting the charge induced by the avalanches at the sense wires. Every pad is equipped with an electronic chain via connections as seen in Fig. 5.15. The electronics components are placed on thin boards situated on the back face of the panel. The use of integrated components reduces the amount of material to a minimum. Such a scheme, also proposed for the HMPID RICH, has been extensively tested on various prototypes.

Each PHOS cradle will be equipped with a separate CPV detector, which will be placed on top of the cradle, see Fig. 5.8 on page 89.

Electronics and spatial resolution

The best resolution can be achieved with an analog readout scheme by calculating the centre of gravity of the charges induced on several adjacent pads (cluster). We intend to use a pad size of 2×2 cm^2 and perform centroid finding only on the side zone of the pads, where clusters will be produced. The central part of the pad (~ 1 cm^2) will give a statistical resolution of about 10 mm/$\sqrt{12}$. This scheme will provide the required accuracy of about 3–4 mm (r.m.s.).

An integrated multiplexed analog electronics has been developed within the framework of the RICH project based on the DIGITPLEX chip or the MCM (multi-chip) version (see Section 4.6). A 2×2 cm^2 pad size implies 2500 electronics channels/m^2.

Using digital electronics would give a coarser resolution, directly related to the size of the pad and would necessitate a new electronics development.

5.3 R&D

5.3.1 Light yield at low temperature

The light yield of the lead tungstate crystals is known to be strongly dependent on the temperature, however, so far, most tests have been done on small samples or single crystals. We have therefore tested two complete arrays of 3×3 single crystals ($2 \times 2 \times 16$ cm^3) at different temperatures with a particle beam. One array was built from the crystals produced by the Monocrystal Institute in Kharkov, Ukraine, the other by the Bogoroditsk factory in Russia. All detectors have been calibrated in the X1 test beam at CERN using electrons with energies of 3, 5, 7, 10 and 15 GeV. Calibration was done at room temperature as well as at a temperature of -24°C, using PIN diodes as the photodetectors. The experimental test set-up is shown in Fig. 5.16.

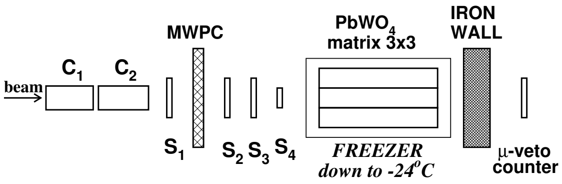

Figure 5.16: Test set-up.

The electron energy was determined as the sum of the energy depositions in individual modules for selected events in which electrons hit a 1×1 cm^2 area in the centre of the central module. The electron signal increases by a factor of ~ 3 (slightly higher but still compatible with Ref. 7) when lowering the temperature to -24°C. At the same time, the energy resolution improved by nearly a factor of 1.5–2 compared with the energy resolution at room temperature, due in part also to a reduction in the electronics noise. Good linearity of the detector response was found both at room temperature and at -24°C. The data are shown in Fig. 5.17, where the energy dependence of the peak positions and σ_E/E are shown for both tem-

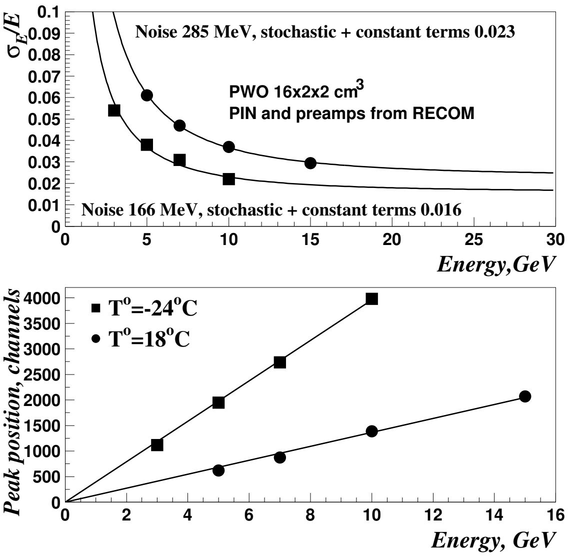

Figure 5.17: σ_E/E (top figure) and electron peak position (bottom figure) vs. electron energy.

peratures. It should be noted that in this measurement the electronics pick-up noise was a limiting factor for the resolution.

5.3.2 Test of timing resolution

Simulations show that antineutrons will imitate photons in the p_t range of interest and appear as background in the photon yield (see Section 11.7). With time-of-flight information this background can be subtracted in the offline analysis.

The beam test investigated the timing performance of $PbWO_4$ and PIN diodes, coupled with different types of Charge-Sensitive Preamplifiers (CSP). A fast shaper (ORTEC474) was used in the test. For comparison, time-resolution measurements were also made with a vacuum phototetrode (Hamamatsu 2259).

The results of time-resolution measurements with electrons at an energy of 5 GeV are shown in Table 5.4. While the intrinsic time resolution of $PbWO_4$ seems extremely good (< 500 ps), the resolution reached with the CSPs does not yet meet our requirements of < 1 ns.

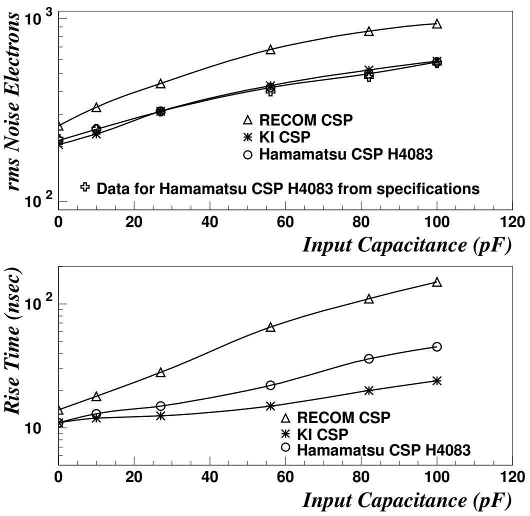

Figure 5.18: Dependence of the noise level and the rise time on the input capacitance for a number of preamplifiers (RECOM, KI, Hamamatsu).

However, none of the CSPs is optimized for TOF measurements and all of them had a long pulse rise-time (see Fig. 5.18). The development of a fast CSP is now in progress in collaboration with the Centre for Future Technologies and Instruments, Moscow.

5.3.3 PIN diode tests

First generation PIN photodiodes (from RECOM, Moscow) have been tested in an electron beam at 10 GeV energy. For comparison, the Hamamatsu S3590-01 PIN diodes were included in the test. Both diodes had the same active area of 10×10 mm^2. The PIN diodes were coupled to the same crystal of 16 cm length by means of optical grease (Dow-Corning Q3067). The same Hamamatsu H4083 charge-sensitive preamplifier was used during all tests. The gain of the electronic chain was monitored with a calibrated pulse generator. The results of the tests are shown in Table 5.5.

Table 5.4: Preliminary results of the TOF resolution measurements with different types of CSP.

CSP	σ (ns)
Recom	7.8
KI CSP	3.0
Bergen CSP	4.2
Phototetrode	< 0.5

Table 5.5: Position of peaks from 10 GeV electrons for different types of PIN diodes.

PIN diode type		Peak position
Hamamatsu	# 1	998
	# 2	940
RECOM	# 1	753
	# 2	829

At present, the sensitivity of the RECOM PIN diodes to the lead tungstate crystal light is ~20 % lower than that of the HAMAMATSU PIN diodes. In 1996 R&D will be performed in order to optimize the RECOM technology.

5.3.4 Preamplifier tests

Two different designs (RECOM and RRC 'Kurchatov Institute') were made in the Kurchatov Institute with slightly different circuit diagrams. The first one has lower power consumption and the estimated production cost is lower. The circuit diagram of the second one is shown in Fig. 5.19.

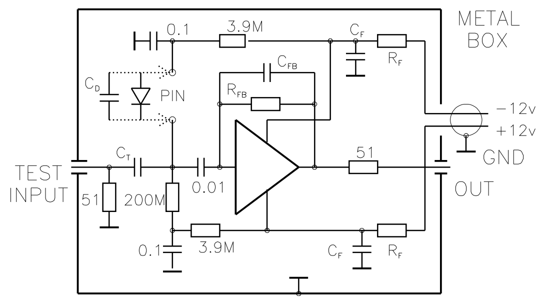

Figure 5.19: CSP circuit diagram.

A comparison of noise level and rise time (including, for comparison, a commercial CSP) is shown in Figs. 5.18 and 5.20. The measurements were done with equivalent capacitance and with different PIN diodes connected to the preamplifier input.

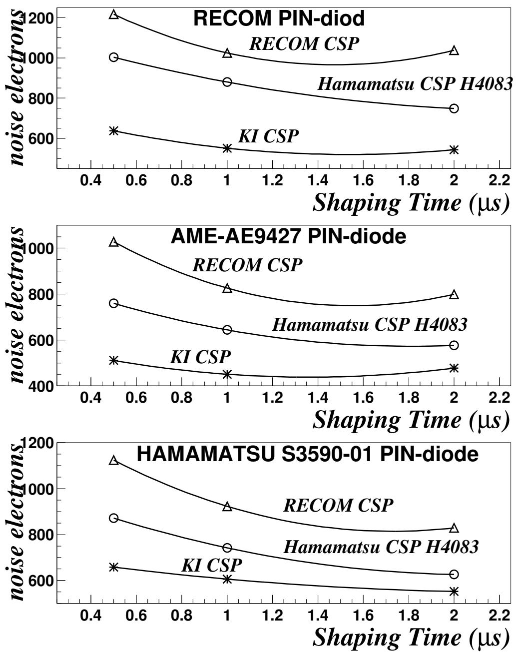

Figure 5.20: Noise level vs. shaping time for different combinations of preamplifier and a PIN photodiode.

All the CSPs tested include a field effect transistor in the input cascade of the CSP, which should provide for a low noise level [21]. Initial noise levels of the CSP are defined by the quality of the components in the first cascade, gate leakage current of FET, and the CSP production quality. In order to decrease the influence of pickup noise on the CSP parameters, it will be necessary to shield both the PIN diode and the CSP.

The PIN diodes and CSPs will work at $-25°$C, which should reduce the electronic noise level. According to Ref. 22, a decrease of the temperature by $50°$C reduces the noise level by a factor of 2–3.

6 Muon Arm Integration

6.1 Introduction

The production and suppression of heavy quarkonium resonances via Debye screening is one of the most eminent QGP signatures which needs to be convincingly addressed as part of the LHC heavy-ion programme. The LHCC therefore asked the ALICE Collaboration to extend its baseline programme by including a measurement of J/ψ and ϒ production.

The collision of Pb ions at the LHC will lead to energy densities —over a sizeable volume— well beyond the threshold of 1–3 GeV fm^{-3} predicted for a transition to QGP. One of the most promising signatures for the existence of this new phase is the suppression of heavy quark resonances (J/ψ, ψ′, ϒ, ϒ′, ϒ″) [1, 2], signals which are accessible through dimuon production.

The measurement of a suppression alone is, however, not an unambiguous proof for the existence of the QGP. It could also be explained by other processes like absorption in a dense hadronic medium or shadowing of the nuclear structure functions. A distinction between these processes is only possible if one measures the *relative suppression* between the *different onium states* and, in addition, relative to an *unsuppressed reference* process like inclusive heavy quark production (beauty or charm). The mass resolution has to be as good as 1 % in order to separate the different resonance states. It is also mandatory to measure the suppression as a function of the transverse momentum including the low-p_t region. In addition, within ALICE we can *correlate* the suppression pattern measured in the muon arm with many other signals which are observed in the central detectors on an *event-by-event basis*.

Muon identification in the LHC environment is only feasible for muon momenta above about 4 GeV/*c*. In the forward direction the muons are Lorentz-boosted and it is possible to measure charmonium production down to low p_t. Furthermore, the background caused by pion and kaon decays is reduced due to the larger momenta involved and also to the lower particle multiplicity per unit of rapidity.

Therefore we have proposed a forward muon spectrometer which would meet all the above-mentioned requirements. Its layout and expected performance are described in an addendum to the ALICE Letter of Intent [3]. The muon spectrometer was accepted in the second half of 1995 with the recommendation to proceed with this upgrade towards the technical proposal stage.

The technical design of the muon spectrometer is not yet on the same level as that of the central ALICE detectors. In two crucial areas we need further detailed simulations and, in particular, *experimental verification* of these simulations, before we can proceed with more specific solutions concerning the components of the muon arm (detectors, magnet, trigger, etc.):

- The first item is the arrangement and design of the absorbers which shield the forward detectors against the enormous particle flux in central Pb–Pb collisions (> 10 000 charged particles in a forward cone of 15°). A good knowledge of the particle spectrum leaking from these absorbers is crucial for both the occupancy in the muon chambers, which will define granularity and technological choices, and for the rapidity coverage at small angles, which will determine the statistical significance of the results. This point is of relevance for the muon arm only, i.e. largely independent of the other detectors in ALICE.

- The second item concerns back scattering of charged and neutral particles from the absorbers into the central detectors, in particular into the TPC. This point is crucial for the proper functioning of the central detectors, and will have an impact on the positioning of the front absorber (i.e. the distance to the TPC) and on our running scenario.

In the present technical proposal, we will address only those aspects of the muon arm which have implications on the performance and design of the central ALICE detectors. Besides preliminary studies on *back scattering*, these are: the *mechanical integration* of the muon detector, the influence of the additional *magnetic dipole field* on the central solenoidal field, muon *trigger rates*, and *data acquisition*. These aspects are partially dealt with in the corresponding chapters (Trigger: Chapter 9; DAQ: Chapter 10; Integration: Chapter 12) and partially below.

We expect to submit an addendum to this Technical Proposal with a detailed description of the muon arm in the second half of 1996.

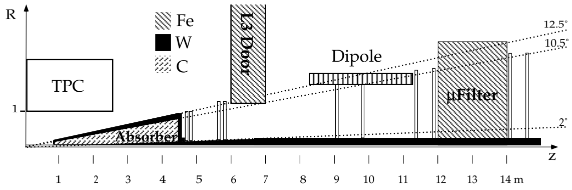

Figure 6.1: Principle layout of the ALICE muon arm.

6.2 Mechanical integration

The principal layout of the muon arm is shown in Fig. 6.1. It follows the same basic concept as presented in the Let-

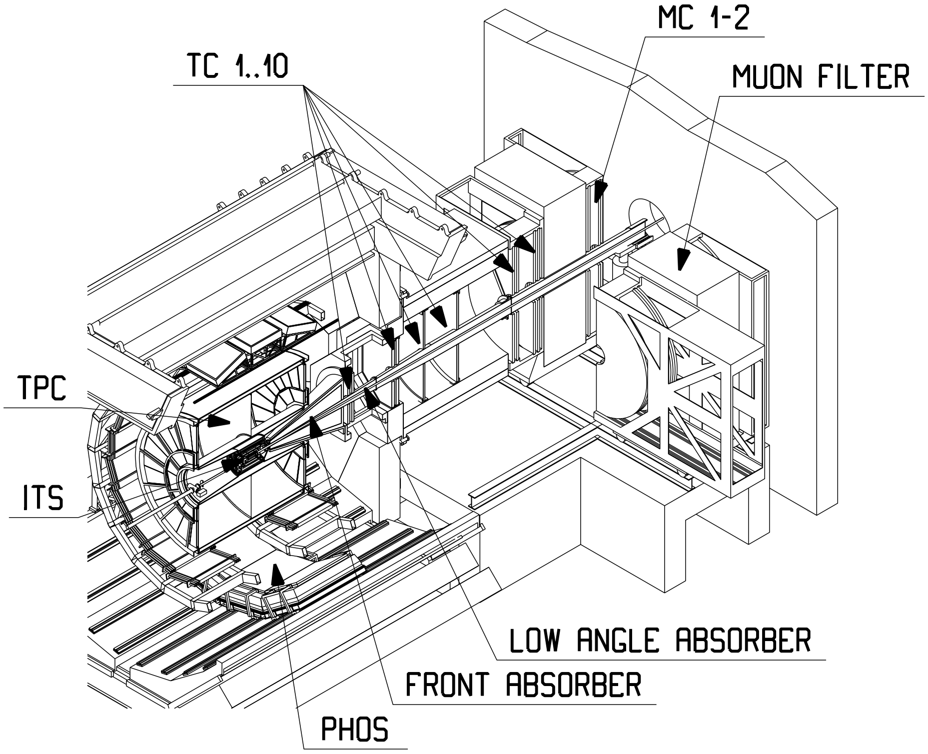

Figure 6.2: General view of the muon arm support structures and its integration into ALICE.

ter of Intent addendum [3]. The outer angular acceptance of the spectrometer has been decreased from 12° to 10.5°, the inner acceptance remains at 2° (corresponding to η = 2.4–4). This allows the dipole magnet to be placed just outside the L3 magnet door at a distance of 9.75 m, between the centre of the dipole and the intersection point, and avoids an expensive rebuilding of the L3 magnet door.

The elements of the muon arm —front absorber (30 t), small-angle absorber (80 t), dipole magnet (850 t), chambers and the muon filter (600 t)— are all placed on a common rail system on the far side of the L3 magnet, as shown in Fig. 6.2 and 6.3.

The front absorber consists of a conical carbon absorber covering the angular acceptance of the muon arm. It is placed 900 mm from the intersection point and has a total length of 3.5 m. To reduce back scattering and particle leakage, the absorber is covered by a 10 mm layer of tungsten at the front end of the cone (close to the ITS) and in the sector between 10.5° and 12.5° which faces the TPC. A tungsten cover of 100 mm thickness at the back end absorbs most of the low energetic electrons created in the absorber. An additional ring of 100 mm of tungsten is added to the 2° cone to improve the shielding against particles from the beam pipe.

The front absorber will be supported by a stainless steel structure attached to the yoke of the dipole magnet, i.e. the absorber can be retracted by moving back the dipole. The space between the octagonal opening of the L3 magnet door and the absorber support is filled with soft iron in order to decouple more efficiently the two magnetic fields, and to improve the field quality inside the L3 magnet (see below).

The small-angle absorber shields the tracking chambers from the large flux of particles emerging from the beam pipe. Close to the vertex it surrounds the beam pipe in the form of a cone with an opening angle of 2°, while for $z > 5.5$ m it forms a cylinder of 58 cm diameter. The total length of the absorber will be 18.1 m. The preferred material is tungsten, because of its small nuclear interaction length, but we envisage using a cheaper material like copper or lead for the parts of the absorber which are far away from the chambers, and therefore less critical.

To ensure mechanical stability, the small-angle absorber will be supported by a steel tube which contains the absorber material. The inner diameter of this tube is 120 mm and encloses the beam pipe and a permanently installed heating jacket. The small-angle absorber will be divided into two parts. The first part is imbedded into the front absorber and supported by a structure attached to the

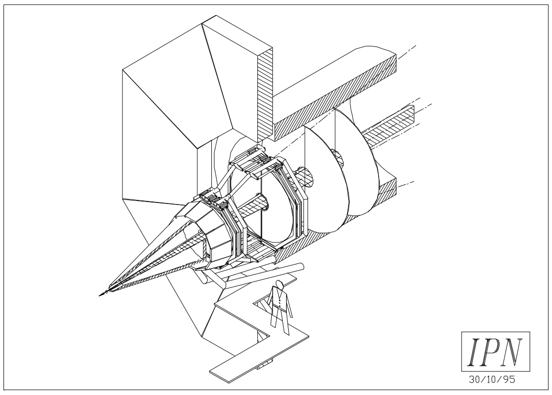

Figure 6.3: Detailed view of the muon arm absorber support structure.

dipole magnet. The second part is fixed on one side also to the dipole, and on the other to a support located in front of the low-β quadrupoles.

The muon filter, made from iron, has a width and height of 5 m and a depth of 2 m. It is placed behind the dipole magnet at a distance of 12 m from the intersection. It is divided into two pieces, which can be moved perpendicular to the beam axis. This allows the retraction of the dipole back to the wall of the experimental hall (see also Chapter 12).

The muon tracking chambers are divided into three groups: four planes (TC1–TC4) between the front absorber and the dipole, two planes (TC5 and TC6) inside the magnet, and four planes (TC7-TC10) between the dipole and the muon filter. The total sensitive area is about $100\,\mathrm{m}^2$. All chambers are divided into two half planes to ease maintenance and installation.

TC1–TC4 are attached to the front absorber support structure. Access holes in the structure will allow some maintenance work on the front-end electronics without opening the spectrometer. The chambers can be dismounted while keeping the front absorber in place.

TC5 and TC6 are supported inside the cryostat of the dipole. The support will allow the chambers to be rotated in order to access the electronics. However, for major interventions, the chambers will have to be retracted near the end of the magnet.

TC7–TC10 are mounted onto a cradle system attached to the vertex side of the muon filter. On one side, the cradles are enlarged to allow the retraction of the respective chambers from the beam pipe. This provides easy access to enter into the magnet to service chambers TC5 and TC6.

The trigger chambers will also be mounted into cradles attached to the far side of the muon filter. Thus the necessary space to retract the dipole is provided by simply opening the muon filter.

6.3 Impact of the dipole field

6.3.1 Introduction

In order to reach a mass resolution of $\sim 1\,\%$ to separate the Υ' and Υ'' resonances, the forward dipole must produce an integrated field of the order of 3 Tm. We have proposed a 1 T superconducting dipole magnet with a magnetic length of 3 m and an aperture of 3.8 m [4]. If the magnet is placed directly behind the L3 magnet door the angular acceptance is $10.5°$. While this magnet is still subject to further optimization, the concept of the forward muon arm depends strongly on the coupling between its magnet and the central solenoid. In particular, the forces on the L3 doors and the field quality in the region of the TPC have been of some concern.

A 3D calculation of the ALICE magnets using the TOSCA program [5] was used to estimate the impact of the muon dipole on the L3 magnetic field in the region of the TPC and on the mechanical forces on the doors of that magnet [6]. The results led to an improvement of the field homogeneity by moving the centre of the dipole magnet to $z = 9.75$ m and by modifying the doors of L3. In this configuration, neither the coupling of the two magnets nor the field quality in the TPC region pose any particular problem.

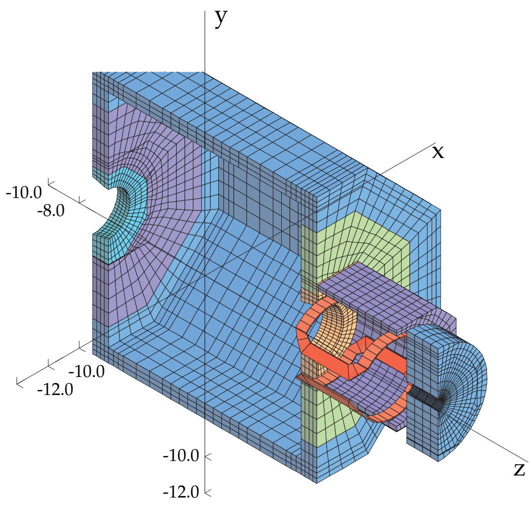

Figure 6.4: Magnet geometry used for field calculations.

6.3.2 Computer model geometry

The geometrical and magnetical complexity of the ALICE detector system made it necessary to model only one half of the apparatus (employing symmetry with respect to the $y-z$-plane, see Fig. 6.4). The iron yoke of the L3 solenoid has been subdivided into an octagonal cylinder around the coil, the 'door' volume as installed at present and an additional 'plug', which fills the present opening of the 'doors'

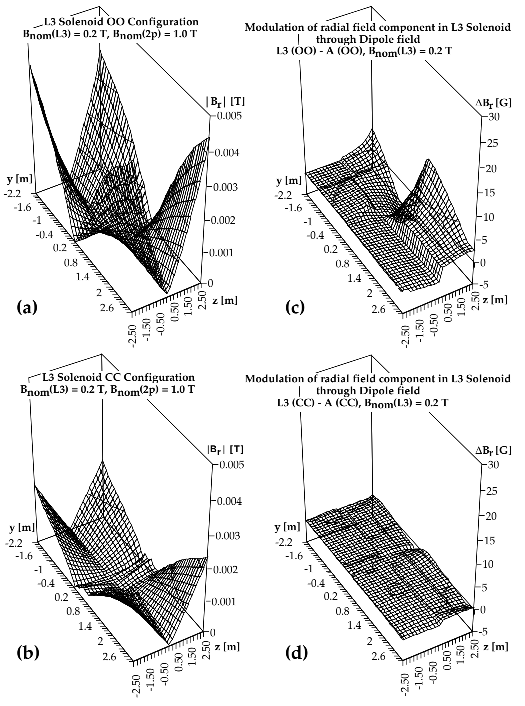

Figure 6.5: Radial field component B_r in the plane $-2.5 < z < 2.5\,\mathrm{m}, -2.2 < y < 2.8\,\mathrm{m}$ which is occupied by the TPC. Note that the L3 magnet is displaced by 30 cm in the y direction with respect to the nominal LHC axis. (a) OO Configuration and (b) CC Configuration for nominal solenoidal field of 0.2 T plus dipole field of 1 T. Change in the B_r component of L3 introduced by the additional dipole field for (c) OO Configuration and (d) the CC Configuration.

outside the muon arm acceptance (i.e. above 12.5°). An improvement of the magnetic decoupling between both magnets is obtained by inserting the iron plugs and by filling the 10 cm gap between the L3 door and the dipole-magnet yoke.

Calculations were carried out for the configurations listed in Table 6.1. The results from the calculations with the asymmetric configuration OC can be found in [6] and are not discussed further. All configurations have been computed at full L3 field and at reduced field level, i.e. 0.5 T and 0.2 T, respectively.

Table 6.1: Model configurations of the ALICE magnet system.

Suffix	Description
OO	L3 doors as presently installed
OC	Iron plug in door at dipole magnet side
CC	Iron plug in door at both sides

6.3.3 Magnetic field in the L3 solenoid

Field calculations have been carried out for the magnet configurations and field values indicated in Table 6.2. The calculations of the L3 magnet at the nominal field of 0.5 T have been compared with the measured L3 field map values and are in excellent agreement (better than 10 G in the region relevant for the TPC).

The values in Table 6.2 show the radial field component B_r and the total deviation of the axial field component B_z from the nominal values ($B_z = 0.5$ or 0.2 T) throughout the volume of the TPC. The conical plugs in the L3 magnet doors improve the field homogeneity by a factor of two, independent of the muon magnet.

Figure 6.5 shows a map of the radial field and the modulations caused by the dipole. The distortions decrease very rapidly towards the centre of the solenoid. In both the OO and CC configurations the dipole field is completely negligible.

When comparing the configurations OO and CC (without and with modified doors), an important decrease by about a factor of two can be noticed in both radial and axial field deviations close to the circumference of the TPC. This improvement of the field homogeneity reduces the track distortions in the TPC (see Section 3.11.6), and will therefore be carried out independent of the final choice for the muon magnet. A remeasuring of the field in the L3 solenoid will however be required after installation of the plugs.

6.3.4 Magnetic forces

Forces induced by the electromagnetic fields of the two magnets were calculated in order to estimate the required retaining structure for the dipole magnet and to ensure that the forces on the L3 doors are manageable. The results are summarized in Table 6.3, where the coordinate system is the one shown in Fig. 6.4. Each of the L3 magnet doors is divided into two halves in order to open the magnetic volume, and the forces given refer to one of these half doors only.

The forces calculated for the configuration 'OO L3 only' correspond to the present situation of the L3 solenoid and are in good agreement with the measured values. The asymmetry of the forces between the two doors in the second line ('CC + Dip (0 T)') is caused by the presence of the additional iron of the dipole magnet yoke.

The field of the dipole magnet reduces the axial force F_z on door 2 by more than a factor of 2. When 'plugs' are inserted in addition, both axial forces at the dipole side of L3 cancel to some extent.

The reduction of the L3 field to 0.2 T inverses the axial force on the L3 doors next to the dipole magnet, pulling each door towards the dipole with approximately 200 t. By inserting the 'plugs' this force is then amplified by a factor of 2.

Table 6.2: Maximum deviation in the TPC region from an ideal solenoidal field in the axial (B_z) and radial (B_r) field components (in Gauss). The values with the muon magnet switched off (at 0 T) correspond closely to the current situation within L3.

L3 Solenoid field (T)	Muon dipole field (T)	OO configuration (L3 doors as is)		CC configuration (Cone 12.5° both sides)	
		ΔB_r (G)	ΔB_z (G)	ΔB_r (G)	ΔB_z (G)
0.5	0.0	117	197	65	85
0.5	1.0	122	190	67	85
0.2	0.0	47	79	26	34
0.2	1.0	55	75	27	34

Table 6.3: Expected forces on the doors of the L3 solenoid with and without the additional dipole magnet.

L3 Solenoid field (T)	Configuration	Forces on L3 doors (t)			
		L3 door 1 (Shaftside)		L3 door 2 (Dipoleside)	
		F_y	F_z	F_y	F_z
0.5	OO (L3 only)	1	517	1	−525
0.5	CC + Dipole (0 T)	−0.4	532	−6	−465
0.5	OO + Dipole (1 T)	−0.1	517	−129	−228
0.5	CC + Dipole (1 T)	2	531	−57	−28
0.2	OO + Dipole (1 T)	−8	82	−62	190
0.2	CC + Dipole (1 T)	−3	85	−19	381

The axial forces are comparable to the weight of the doors and are therefore acceptable. When operating at reduced field (0.2 T), however, special precautions should be taken to block the doors at the dipole magnet side, i.e. by closing the gap between the door and the dipole yoke.

6.3.5 Magnetic field in the muon dipole

The field integral of the muon dipole between the vertex and the muon filter has been calculated over the angular acceptance of the spectrometer ($-10.5° < \theta < 10.5°$). Table 6.4 shows the integral of the vertical field component B_y in this volume at reduced L3 field. For the largest angle the integration stops at the centre of the dipole magnet. The difference between the positive and negative angles is a result of the displacement between the nominal centres of the LHC and LEP.

Table 6.4: Integrals of the vertical field over the acceptance of the muon spectrometer.

Angle (degrees)	$\int B_y\,dz$ (Tm)	
	L3 doors: OO	L3 doors: CC
10.5	1.86	1.70
8.5	3.81	3.57
2.0	4.08	3.82
0.0	4.09	3.83
−2.0	4.12	3.86
−8.5	3.98	3.71
−10.5	2.13	1.90

6.4 Back scattering into the TPC

6.4.1 Introduction

The task of the hadron absorber —placed at a distance of 90 cm from the interaction point— is to reduce the huge primary particle flux from Pb–Pb collisions by at least two orders of magnitudes. The shape and material composition of the present design is based on shower simulations performed with GEANT and FLUKA (see Ref. 3 and references therein).

The simulations predict that it is possible to reduce the charged particle flux into the muon chamber behind the absorber to values between $\sim 10^{-2}\,\mathrm{cm}^{-2}$ per event, close to the beam pipe, and $\sim 10^{-3}\,\mathrm{cm}^{-2}$, at larger radii.

However, our simulations also show that the absorber radiates a considerable flux of photons and neutrons through its large outer surface which is oriented towards the TPC. The TPC volume starts at a radial distance of about 1 m, extending ±2.5 m along the beam. The number of charged particles entering the TPC from the absorber with at least a few tens of MeV kinetic energy, i.e. particles that would leave recognizable tracks, is small compared to primary particles from the vertex, and increases the occupancy in the TPC by only about 10 %. Such a minor increase would not present any problem and is far below the incertainty inherent in the estimates of the primary particle multiplicity at the LHC.

However, the effect of neutrons and very soft photons cannot be estimated sufficiently reliably by means of Monte Carlo simulations. This background will, in most cases, not leave distinct tracks in the TPC but might contribute to a background of single hits or increase the chamber current. Therefore a test of the performance of a TPC close to an absorber was performed with the NA49 experiment at the SPS. FLUKA simulations of both the NA49 test and the ALICE set-up were then used to scale the experimental results to the LHC environment.

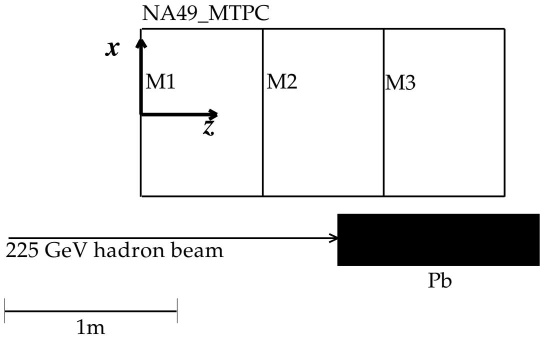

Figure 6.6: Set-up of the NA49 Absorber/TPC test, seen from above.

6.4.2 The NA49 Absorber/TPC test

In the NA49 test a 225 GeV/c hadron beam was dumped on a lead block of dimension 30 cm × 30 cm × 120 cm. As illustrated in Fig. 6.6, the block was placed at a distance of 10 cm from the NA49 Main TPC (MTPC) and centered vertically. An Ar/CH_4 mixture (95/5 per volume) was used as the chamber gas. Data were taken from the three TPC modules closest to the lead block (M1, M2, and M3). Electrons drift along the y direction towards the readout planes, located on top of the modules in the $x - z$ plane.

Table 6.5 shows, for each module, the measured anode current as a function of the beam intensity (number of particles per spill of 2.6 s) and for three different settings: OPEN (the gate was open during the duration of the spill), PULSED (the gate was opened for 52μs 13 times during the spill), and CLOSED (the gate was kept closed). By increasing the beam intensity the energy dumped into the absorber was varied between 2×10^3 and 1×10^5 TeV/s. In the range of 2×10^3 and 2×10^4 TeV/s the anode current of the module closest to the absorber (M3) increased from 0.47μA to 3.7μA when the gate was permanently open.

6.4.3 Nature of the Anode Current

We have checked that the anode currents I_1, I_2, and I_3 are directly proportional to the total pulse height Q read out

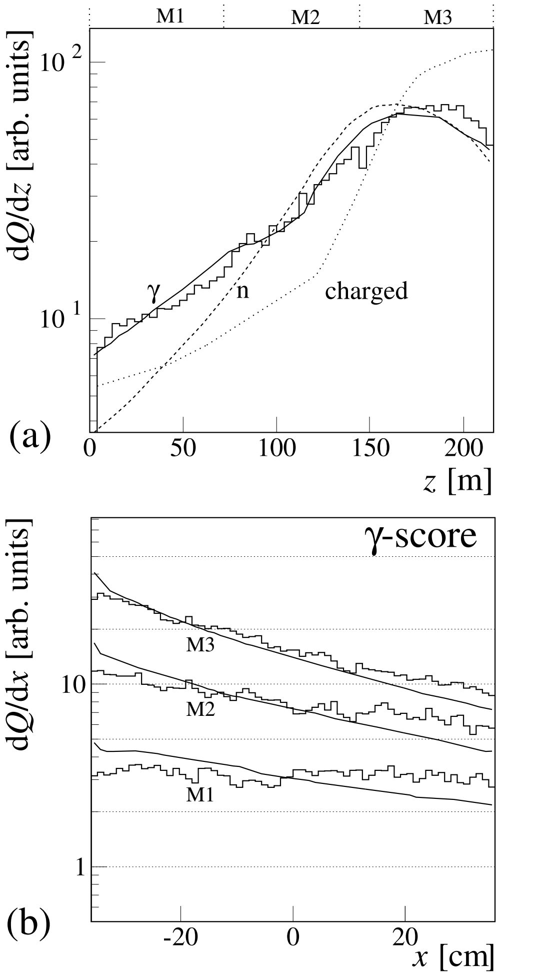

Figure 6.7: Measured charge distribution in the TPC along z (upper) and along x (lower) averaged over several thousands of hadron–Pb interactions. The distribution of the flux × track-length scores (normalized to the same integral) as simulated by FLUKA are shown for photons (solid), neutrons (dashed) and charged particles (dotted) for comparison.

in each of the modules during the gate opening time via the electronics chain. The proportions between the three modules are $Q_3 : Q_2 : Q_1 = 5.3 : 2.5 : 1$ in terms of collected charge, and $I_3 : I_2 : I_1 = 5.3 : 2.2 : 1$ in terms of currents. The anode currents shown in Table 6.5 are directly proportional to the power dumped in the absorber, i.e. to the number of particles/spill. In addition, the ratio of currents between the OPEN and PULSED mode are given to a good approximation by the ratio of the spill length to the gate opening time ($13 \times 52\mu$s$/2.6$ s $= 6.5 \times 10^{-4}$). The measured ratio is (6–8) 10^{-4} at intensities where the current difference between the pulsed and open gate settings is larger than the measurement error (some nA). These relations are as expected, independent from any simulation, and show that the currents are indeed a good indication

Table 6.5: Anode currents measured in the NA49 Absorber/TPC test.

Intensity (Particles/Spill)	Dumped Power (TeV/s)	Gate	I_1 (nA)	I_2 (nA)	I_3 (nA)
22k	1.9×10^3	p	1	2	3.5
		o	88	230	470
		c	1	1	2
110k	9.6×10^3	p	3	6	11
		o	320	850	1230
		c	3	6	10
220k	1.9×10^4	p	5	12	22
		o	700	1700	3700
		c	5	11	19
520k	4.5×10^4	p	14	30	58
		c	13	30	52
1250k	1.1×10^5	p	29	67	141
		c	27	65	140

Table 6.6: Comparison of simulated particle scores and currents measured in modules M3, M2 and M1.

Score	Proportion M3:M2:M1
NA49 Absorber test	
Charge	5.3 : 2.5 : 1
Current	5.3 : 2.2 : 1
FLUKA Simulation	
Flux × track-length	
neutrons	9 : 5 : 1
gammas	5.1 : 2.5 : 1
charged	13 : 2.2 : 1

for the load and occupancy in the TPC.

A FLUKA simulation of the rate of particles leaving the absorber and entering the TPC modules shows that per 100 TeV of absorbed energy 1.6×10^3 charged particles, 1.8×10^4 photons, and 3.6×10^5 neutrons enter the three TPC modules. In order to estimate how these fluxes are related to the currents produced in the modules we score in each module the particle flux times the track length within the module. This score should be proportional to the probability of leaving a hit for all particle types that have either a low probability of interaction or are not stopped inside the TPC.

The proportions of these scores are compared in Table 6.6 with the experimental results for total charge and current per module. One can see that the γ-score reproduces best the charge proportions between the three modules. Charged particle and neutron scores coincide within approximately a factor of 2.5. The same result is obtained by comparing the measured differential charge distributions to the corresponding distributions of the scores (Fig. 6.7).

Further experimental input comes from the ratio between the currents measured with the OPEN and CLOSED gate settings. This ratio varies between 90 and 230 in the data. In the simulation, the CLOSED setting can be evaluated by scoring the flux times the track-length in the region between the sense and gate wires. The ratio between this score and that obtained in the entire volume of the modules (gate open) amounts to approximately 1:150.

We conclude that the ratio of currents between modules as well as between the OPEN and CLOSED gate is within a factor of 2–3 proportional to the flux × track-length score, as obtained from a FLUKA simulation. As these ratios are non-trivial —besides geometry they involve also the angular distribution of particles as calculated with FLUKA— we can have some confidence that the simulation can be used to extrapolate from the NA49 test to the ALICE muon arm environment.

6.4.4 Extrapolation to ALICE

The result of an extrapolation of the anode current for ALICE is shown in Table 6.7. The first line compares the mean kinetic energies of the particles emerging from the NA49 test absorber and the ALICE absorber, respectively. For all particle types the energies are similar at the SPS and the LHC.

Next we compare the flux × track-length scores for charged particles, neutrons and gammas of the NA49 test (Module 3, 220k/spill, OPEN gate) with the corresponding scores for the ALICE set-up. Note that the differences between the numbers are combined effects of the different geometries and absorbed power. The latter amounts to $1.9 \times 10^4\,\mathrm{TeV/s}$ for the NA49 test and to $6 \times 10^3\,\mathrm{TeV/s}$ for ALICE at a trigger rate of 60 Hz (central Pb–Pb). We note that the ratio between the ALICE and NA49 scores is

Table 6.7: Extrapolation to ALICE.

	NA49 Test	ALICE
Mean kinetic energy (MeV) charged : n : γ	150 : 2 : 2	140 : 1.8 : 1.9
Power (TeV/s)	19000	6000
Flux $\times$ track-length score ($\mathrm{cm\,s^{-1}}$) charged : n : γ	$2.3 \times 10^7 : 2.3 \times 10^8 : 4.8 \times 10^9$	$3.4 \times 10^7 : 8.4 \times 10^8 : 2.1 \times 10^9$
Integrated Current (nA)	5400	13500
Sense wire surface ($\mathrm{m^2}$)	0.5	16
Current per unit area ($\mathrm{nA\,m^{-2}}$)	7400	1750–450 for R = (1–2.5) m

largest for gammas. We use the gamma score, which will give the most conservative extrapolation and also reproduces best the proportions between the module currents.

Scaling the measured anode current with the ratio of the gamma scores yields an expected total current of 13.5 μA for ALICE. From a distribution of the scores as a function of the radial distance to the beam axis it is found that the current per unit area varies between 1.7 μA/m^2 and 0.45 μA/m^2 for the smallest (1 m) and largest (2.5 m) radius, respectively.

The same extrapolation can be used to estimate the current in the region between the sense wires and gate wires, i.e. the region which is not made insensitive by the gate. The integrated current from this region amounts to only 2.5 % of the total current in the ALICE TPC.

The estimated current in the ALICE TPC produced by back scattered particles is of the same order as the current produced by primary particles from the vertex ($\sim$5 μA).

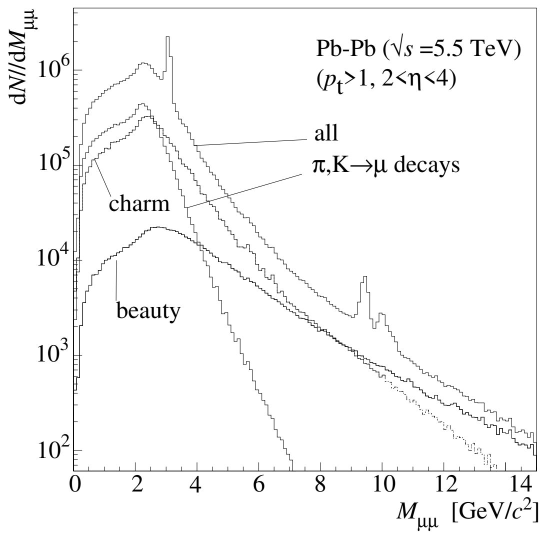

Figure 6.8: Simulated dimuon mass spectrum in central Pb–Pb collisions, showing separately the contribution from decays, charm and beauty-decays.

This first test therefore increases our confidence that the absorber can be placed close to the currently foreseen position without impairing the operation of the central detectors in ALICE. However, further tests with a more realistic absorber configuration are foreseen, and the background of hits (as opposed to currents) still needs to be evaluated.

The charge produced by neutral particles from the absorber will be significantly lower than estimated above if Ne is used instead of Ar as the chamber gas. Neon has a smaller Z and hence a lower $\gamma A \rightarrow e^{\pm}X$ cross-section. The dominant process in which neutrons with a kinetic energy of some MeV produce charges is by knocking out of protons. This favours the choice of a hydrogen-free gas.

6.5 Dimuon trigger

6.5.1 Heavy quark production at the LHC

The cross-section for heavy quark production has been extrapolated from data up to Tevatron energies using QCD calculations and structure functions as described in [3, 7, 8].

Figure 6.8 illustrates the main features of the dimuon mass spectrum as seen after the absorber (i.e. with certain momentum cuts). The aim of the measurement is to extract J/ψ and ψ′ in the 3 GeV range, and the ϒ family and the $b\bar{b}$ production in the 9 GeV region and above.

6.5.2 Trigger concept

The thickness of the absorber close to the interaction region (10 interaction lengths) was minimized in order to preserve a good momentum resolution. For muon identification and trigger purposes a 2 m iron absorber (11 interaction lengths) is placed behind the magnet. The momentum cut-off for muons due to energy loss amounts to 4 GeV/c.

The trigger system is based on two logical layers of dedicated trigger stations, e.g. drift tubes or resistive plate chambers, which are placed behind the iron absorber and separated by about 0.5 m. The trigger makes use of the coincidences between these two stations and is based on

Table 6.8: Average number of muons reaching the trigger chambers per event for central heavy ion collisions. The momentum cut due to energy loss amounts to 4 GeV/c.

Collision	p_t^{min}(GeV/c)	Υ	J/ψ	bb	cc	Decays	Total
Pb–Pb	no cut	0.0002	0.04	0.1	1.5	10.5	13
	1	0.0002	0.03	0.07	0.24	0.24	0.55
	3	0.0002	0.0015	0.01	0.01	0.002	0.03
Ca–Ca	1	0.000014	0.003	0.004	0.015	0.04	0.06
pp	1	10^{-7}	2×10^{-5}	3×10^{-5}	10^{-4}	2×10^{-4}	3×10^{-4}

the track position and angle. The angle between a straight line to the vertex and the muon track measures the momentum with a precision which depends on the ratio of the multiple scattering angle, and the deflection angle due to the magnetic field. The trigger processors will be placed close to the chamber electronics.

The muon trigger and its relation with the general ALICE trigger and DAQ are described in detail in the corresponding chapters (Chapter 9 and Chapter 10). Essentially, it has to signal at Level 1 (after ~ 1.2 μs) if two penetrating tracks are seen in the trigger chambers with a momentum above ~ 1 GeV ('soft' $\mu\mu$) or above ~ 3 GeV ('hard' $\mu\mu$). For Level 2 (4–50 μs), it calculates the invariant mass of all candidate muon pairs in order to select events in the mass region of the J/ψ. The general trigger and DAQ layout will ensure that the rare dimuon events are given high priority and are read out with little dead time.

6.5.3 Expected trigger rates

Table 6.8 shows a break-up of the expected muon rates for Pb–Pb, Ca–Ca and pp collisions. The effect of tracks formed by random combinations of punch-through background hits has not yet been taken into account. In Pb–Pb collisions, ~ 13 muons per event originating from physics processes reach the trigger chambers. The dominant source are π, K $\rightarrow \mu$ decays (10.5) followed by semileptonic charm decays (1.5).

Table 6.9: $\mu\mu$ trigger rates for central Pb–Pb, Ca–Ca and pp collisions.

	Pb–Pb	Ca–Ca	pp
Conditions			
$\mathcal{L}$ (cm^{-2} s^{-1})	10^{27}	10^{29}	10^{30}
Minimum bias (Hz)	8000	3×10^5	10^5
σ_{tot} (b)	8	3	0.1
Central / Min. bias	10 %	10 %	100 %
$\mu\mu$ trigger rate			
$p_t > 1$ GeV/c (Hz)	~ 240	~ 100	~ 2
$p_t > 3$ GeV/c (Hz)	$\sim$ 0.5	$\sim$ 2	~ 0.05

Multiple scattering limits the momentum resolution obtained with the trigger chambers alone. Therefore, as shown in Ref. 3, transverse momentum cuts alone do not sufficiently reduce the trigger rate in Pb–Pb collisions without compromising the J/ψ and ψ′ efficiencies. However, a combined transverse momentum and mass cut reduces the rate by approximately an order of magnitude for Pb–Pb and more than two orders of magnitude in Ca–Ca collisions. The resulting trigger rates for Pb–Pb, Ca–Ca and pp collisions are shown in Table 6.9. For the high mass region (Υ family) a transverse momentum cut of 3 GeV/c reduces the trigger rates to below 1 Hz.

7 Other forward detectors

7.1 Zero-Degree Calorimeters (ZDC)

7.1.1 Design considerations

The observable most directly related to the collision geometry, or impact parameter, is the energy carried away by the non-interacting beam nucleons (spectators), which will be measured in a zero-degree calorimeter. The LHC dipole D1 (Fig. 7.1), which separates the two colliding beams in space, will also separate the spectator neutrons and protons, which have a different charge-to-mass ratio than the ion beams. It is then possible to detect these neutrons and protons in two distinct calorimeters, placed in front of dipole D2 (Fig. 7.1).

Figure 7.2 shows the positions of the two calorimeters, together with the beam pipes, in a plane perpendicular to the accelerator axis; it also shows the distribution of both the neutron and proton spectators, taking into account the beam crossing angle. The distance between the calorimeters and the interaction point is about 92 m.

The calorimeters will detect not only the spectators but also some particles produced in the collision. This pile-up can be measured and subtracted by using a third calorimeter, identical to the proton one, but placed in the bending direction of the negative particles. According to present simulations, this correction seems to be only a few per cent. If more detailed calculations, including background from the machine elements between the vertex and the ZDCs and using different event generators, confirm this estimate, the third calorimeter might actually not be needed.

Two identical sets of calorimeters will be placed on both sides of the interaction point to improve the resolution of the impact parameter and to study fluctuations in the number of spectators in order to better determine the collision geometry; moreover, this correlated information will be used to remove the background from beam–gas interactions.

Nuclear fragments with a charge-to-mass ratio close to the one of the incoming beams remain inside the beam pipe and cannot be detected in the current set-up. The effectiveness of the ZDC detector as an impact parameter trigger therefore depends, to some extent, on the production of nuclear fragments in central collisions. Data are available only at low energies of about 1 GeV/nucleon from Bevalac [1] and GSI [2], and from emulsion experiments for minimum bias interactions at the SPS heavy-ion beam [3]. The expectation that few fragments will survive in central collisions will be verified with Pb beams at CERN and possibly at RHIC, where plans to install similar devices are currently being discussed.

Finally, we are considering adding a small e.m. calorimeter in front of the neutron one to measure the energy carried by photons.

7.1.2 Detector design

The ZDC calorimeters are of the ‘spaghetti’ type, with quartz fibres as the active material instead of the conventional scintillating fibres. The use of quartz fibres is imposed by their radiation hardness [4], as the calorimeters will work in a high radiation environment. In a minimum bias collision, an average number of 80 neutrons will hit the calorimeter. With GEANT simulation we have estimated that the dose deposited in the neutron calorimeter is of the order of 10^4 Gy per day, at a luminosity $\mathcal{L} = 10^{27}\ \mathrm{cm^{-2}s^{-1}}$.

In this type of calorimeter, shower particles produce light through the Čerenkov effect, generating a very short signal (FWHM of the order of 5 ns). The fraction of the hadronic showers which produces Čerenkov light has a transverse dimension of ~ 7 mm at 200 GeV, much smaller than in conventional calorimeters, which guarantees a good energy resolution with small calorimeter dimensions [5]. Quartz fibre calorimeters are, in addition, intrinsically insensitive to radio-activation background, which produces particles below the Čerenkov threshold.

The neutron calorimeters (ZN)

The very limited space between the beam pipes (see Fig. 7.2) makes the use of high-density shower material necessary for the neutron calorimeter. We will use a tungsten alloy (Triamet with 97 % W, density 17.2 g/cm^3), which can be machined easily.

The proposed dimensions are $8 \times 8 \times 100\ \mathrm{cm}^3$; the transverse dimensions exactly fit the free space between the beam pipes, and the length corresponds to 10 λ_{int}. The pseudorapidity interval covered by the calorimeter is $\eta \geq 8.43$ and the shower leakage for a neutron of 2.7 TeV is ~ 20 %.

A compromise has to be found between the resolution and the number of fibres. The fibre-to-absorber filling ratio of 1/40 (quartz to tungsten volume ratio) has been decided from simulation; from Fig. 7.3 it is clear that below that value the resolution starts to deteriorate very rapidly.

One can see from Fig. 7.4 that the energy resolution depends on the orientation of the fibre axis with respect to the incident particle direction. Even if the fibres are at around 45° to maximize the detector light output, the resolution does not deteriorate much for fibres along the beam direction, and in this case it is much easier to couple them to the photomultipliers. The proposed neutron calorimeter structure can be seen in Fig. 7.5.

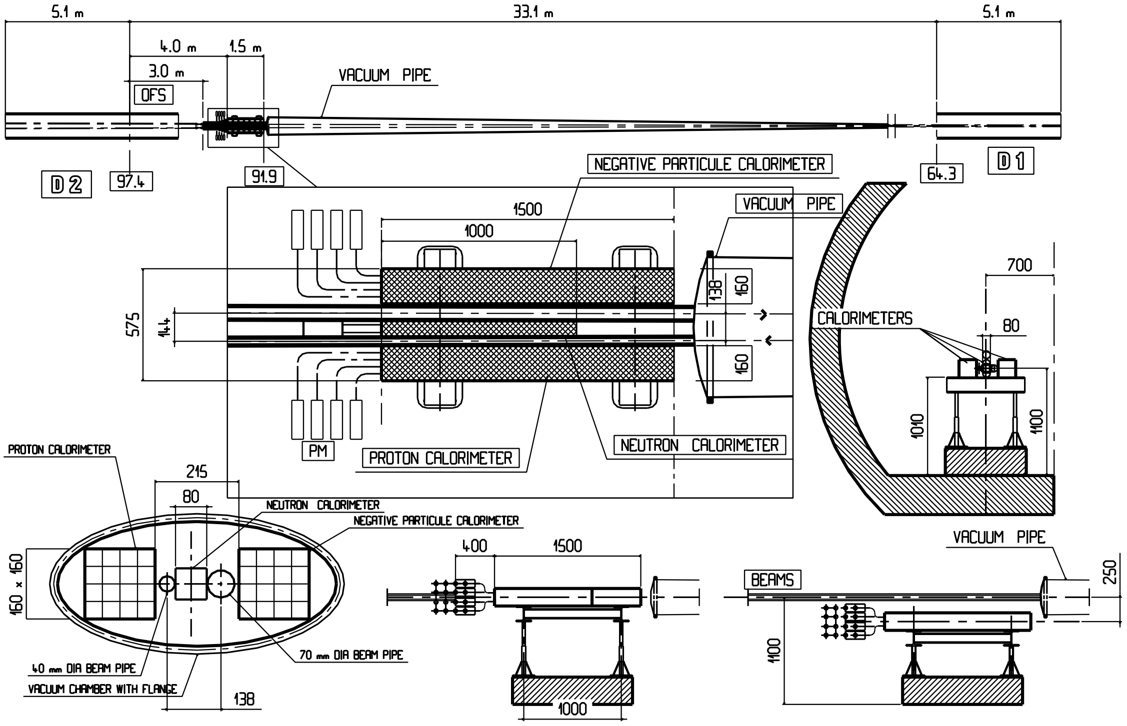

Figure 7.1: Layout of the ZDC detectors.

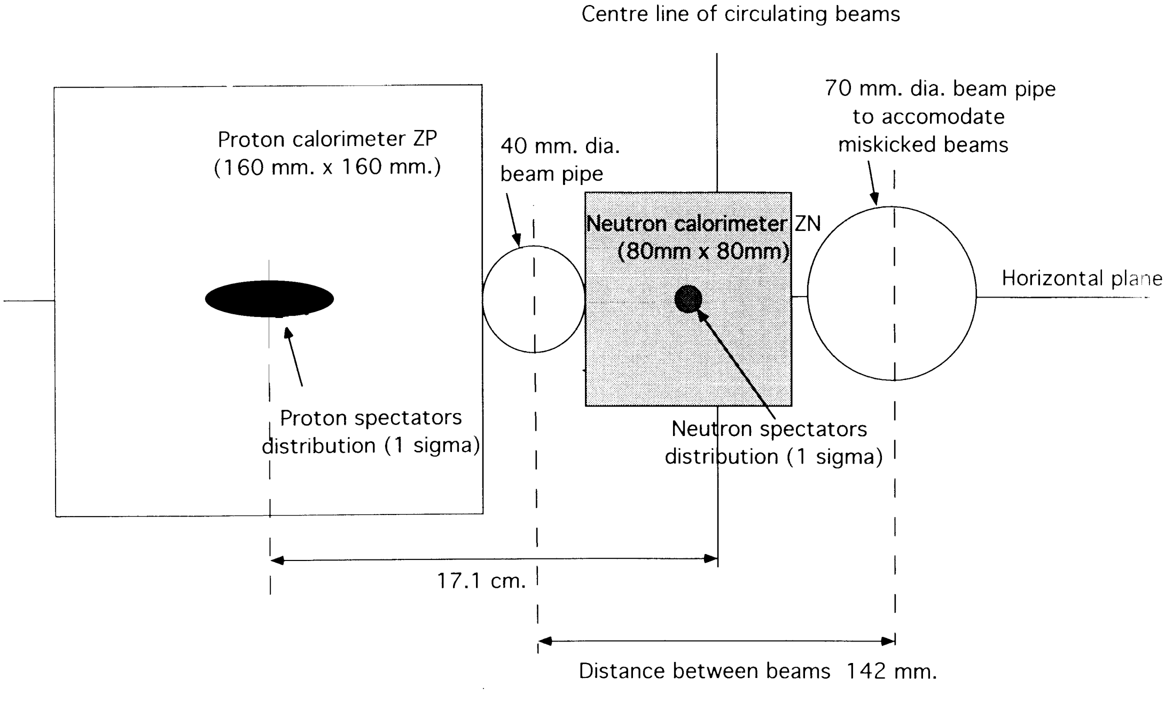

Figure 7.2: Positions and beam pipe geometry of the ZDC detectors.

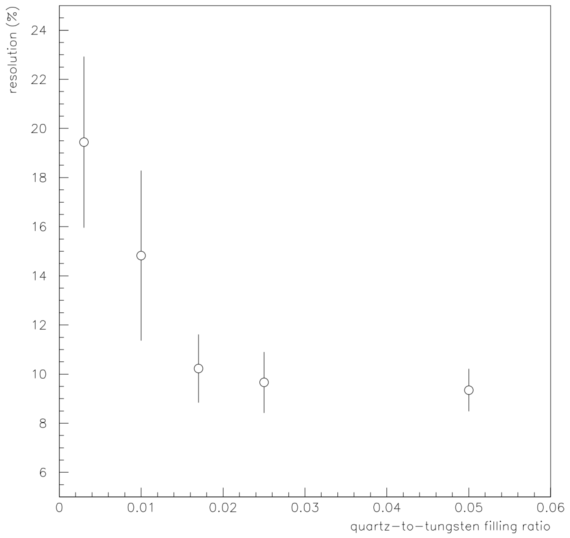

Figure 7.3: Energy resolution of the neutron calorimeter (ZN), as a function of the quartz to tungsten ratio, for 2.7 TeV protons.

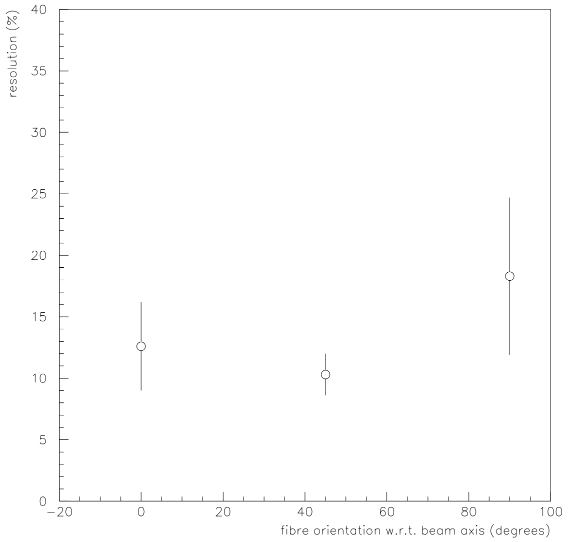

Figure 7.4: Neutron calorimeter resolution as, a function of the angle of fibres with respect to the beam axis, for 2.7 TeV neutrons.

The proton calorimeters (ZP)

The distribution of the energy detected by the proton calorimeter, as a function of the distance from the collider axis, is shown in Fig. 7.6 for the spectator protons as well as for the particles produced in the collision (background). The dimensions of the proton calorimeter are less constrained than those of the neutron one. We can therefore use materials that are less expensive than tungsten, e.g. copper or brass. The dimensions are $16 \times 16 \times 150$ cm^3, the length still corresponding to

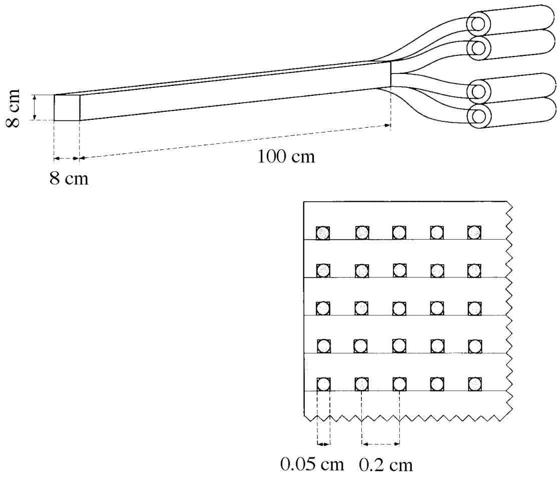

Figure 7.5: Neutron calorimeter structure (top), and front view showing the equally spaced quartz fibres (bottom).

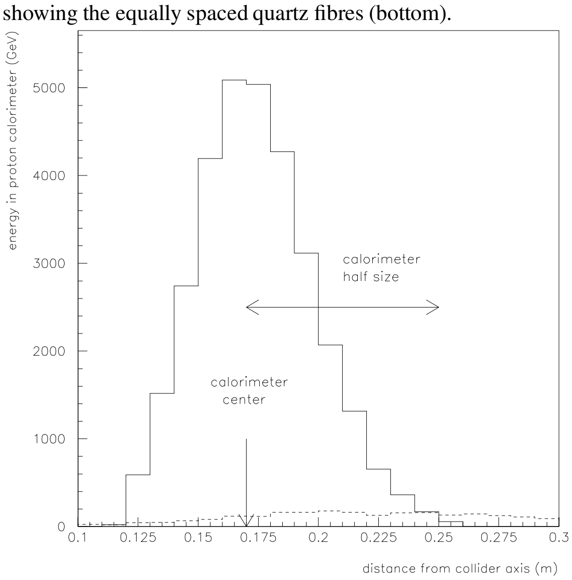

Figure 7.6: Distribution of the energy carried by spectator protons (full line) and by produced particles (dashed line) detected by the proton calorimeter, as a function of the distance from the collider axis.

10 λ_{int}. Concerning the quartz-to-absorber filling we intend to use a volume ratio of 1/80. In this case ~ 8 km of quartz fibres will be used. The shower leakage for a proton of 2.7 TeV is $\sim 15\,\%$.

Mechanics

The total volume of the ZDC will be 6400 cm^3 for the neutron one and 38 400 cm^3 for the proton calorimeter, the weights being approximately 128 kg and 340 kg, respectively. The volume and weight come essentially from the passive material, while the mechanics for the assembly, the optics insulation, and the photomultiplier support will contribute very little to the total weight. The calorimeters can be placed together on a single mechanical support table.

In order to avoid unnecessary radiation damage, one needs to remove the ZDC from the beam plane during LHC proton operation, by moving the supporting table vertically by about half a metre. As the detectors will not reconstruct the coordinates of each detected particle, but only the centroid of their spatial distribution, a positioning accuracy of 0.5 mm will be sufficient.

Detector readout

In order to read out the neutron (proton) calorimeter, we will need ten (sixteen) photomultipliers, even with some redundancy. The photomultipliers must have a low gain (10^5) and a very fast rise time. The signals from the PMs will be sent directly to the trigger system through fast cables or, if necessary for timing purposes, through air-core cables. The signals will then be amplified and sent to the ADCs. As we plan to use single nucleon signals to monitor the stability (see below), we need a large dynamic range (13-bit ADC or two ADCs with different gains).

Calibration procedure

To calibrate the calorimeters before final assembly we need an SPS ion beam of low intensity, possibly at two different energies or with the same energy per nucleon but with two different ion types. Once the calorimeters are mounted in the LHC tunnel we will monitor the overall response by means of the e.m. dissociation which will cause the emission of some nucleons in the forward direction. Simulations indicate that the proposed energy resolution will be good enough to separate the peaks corresponding to one or two nucleons, and that the spatial distribution of the emitted nucleons on the front of the neutron calorimeter will have a width of less than 1 cm (r.m.s.) (see also Fig. 7.6).

The gain of the photomultipliers and the transparency of the fibres will be continuously monitored by a calibrated pulsed light source on the front of the calorimeters.

7.1.3 Performance

The energy detected by each ZDC calorimeter is correlated to the impact parameter. The width of this correlation is determined from event-by-event fluctuations of the number of nucleon spectators and by fluctuations due to the particles produced in the pseudorapidity interval covered by the calorimeters. As these intrinsic fluctuations are of the order of 20 % (see Fig. 7.7), the experimental resolution does not have to be better than 5–10 %.

The energy resolution of the proposed calorimeters was simulated with the GEANT package. To speed up the calculation the light response for each crossing of a fibre by a shower particle was taken from a look-up table [6]. For a nucleon of 2.7 TeV entering the neutron calorimeter, we find a resolution of about 9 % with a quartz-to-absorber volume ratio of 1/40. The simulation programs were checked with experimental data from the NA50 experiment (see below) and from a series of tests performed with hadron and electron beams on a prototype. The results of the programs are in good agreement with the experimental data. The performance of the ZDC calorimeters in the central-event trigger is summarized in Section 9.3.

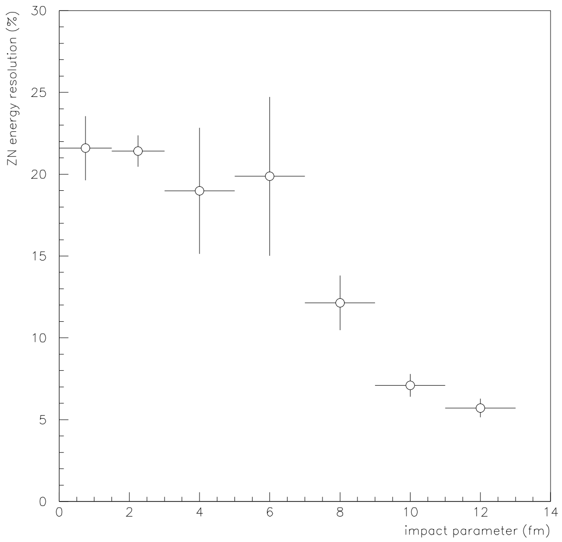

Figure 7.7: Event-by-event fluctuations of the energy emitted in the acceptance of the ZN calorimeter as a function of the impact parameter, as predicted by HIJING.

7.1.4 R&D results

A calorimeter similar to the neutron one proposed for ALICE is at present operational in the NA50 experiment [7]. It is made of tantalum slabs which are grooved for the insertion of quartz fibres (HCG-M365-U), manufactured by Ensign-Bickford. These fibres have quartz cladding, the core diameter is 0.36 mm, and the numerical aperture is 0.22. The dimensions of the NA50 calorimeter are $5 \times 5 \times 65$ cm^3, the angle of the fibres with the beam direction is 0°, and the quartz-to-absorber volume ratio is 1/17. The calorimeter is placed inside the beam dump of the NA50 experiment and stops the 33 TeV lead ion beam. Last year, during the lead ion run period, it worked with a beam intensity of 1.2×10^7 ions per second. At the end of the run period we observed a reduction of the signal amplitude of less than 10 %. The dose received by the NA50 ZDC during this period amounts to about 1 Grad, one order of magnitude bigger than an equivalent run period at the LHC. The back of the calorimeter was divided into four quadrants, each equipped with four Philips XP2242B

PMs; the duration of the signals after 200 m of 50 Ω fast cables (of type C-50-11-1) was 5 ns (FWHM), as shown in Fig. 7.8. In the NA50 experiment the ZDC continuously monitors the beam position, and is able to reconstruct the beam hit position with an uncertainty of less than 0.5 mm.

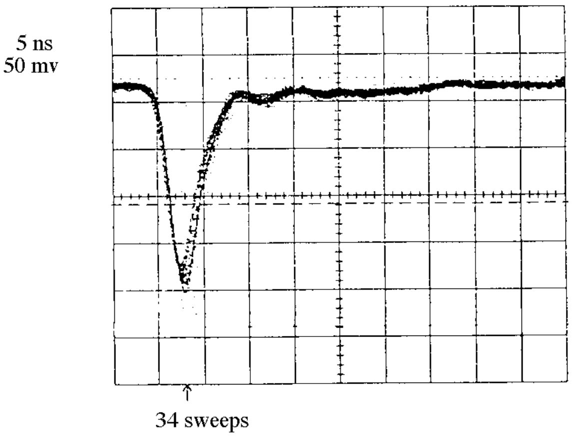

Figure 7.8: Signal shape from one PM of the NA50 calorimeter after 200 m of fast cable for a pion beam of 200 GeV.

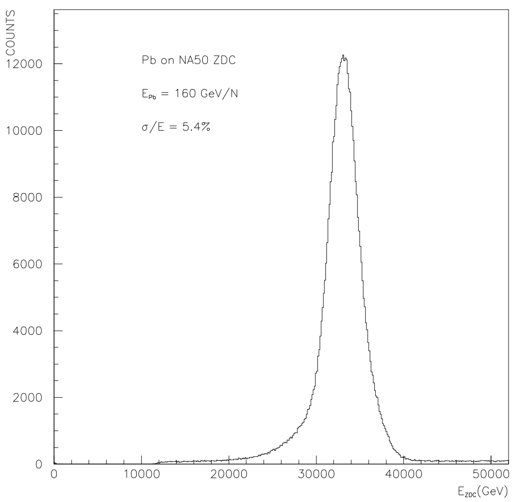

Figure 7.9: Energy resolution of the NA50 calorimeter for a lead ion of 33 TeV.

The energy resolution for the lead ion beam is $\sim 5.4\,\%$ (see Fig. 7.9), as expected from simulations, and we are therefore confident about the results predicted for the ALICE calorimeters.

The experimental results obtained with the NA50 calorimeter show that the quartz fibre calorimeters work well and are suitable for the ALICE experiment. Further studies are in progress to test fibres with a smaller diameter in order to reduce costs and the occurrence of fibre breakages; new photon detectors like the Hamamatsu metal package types R5600 and R5900 will also be tested because their small size simplifies the mechanical supports.

Together with the R&D devoted to the improvement of the detectors, an experiment to measure the fragments produced in lead ion collisions at SPS energies should be foreseen in the forthcoming years, to obtain information about the yield of nuclear fragments as a function of beam energy and centrality.

7.2 The Forward Multiplicity Detector (FMD)

7.2.1 Introduction

The purpose of the Forward Multiplicity Detectors (FMD) is to measure $dN/d\eta$ in the rapidity region outside the central acceptance and to provide information for the trigger system, Level 0 (L0) and Level 1 (L1), at the earliest possible moment.

We have considered two possible detector types for use in the FMD: Micro Channel Plate and Si multipad.

Micro Channel Plate (MCP) detectors

Two or three lead glass MCPs, with a channel diameter of $\sim$6–15 μm, are stacked together in a chevron geometry inside a vacuum container. The charge, produced by a traversing MIP and amplified by a factor of up to 10^6 in the lead glass channels, is collected on a pad plane situated $\sim$0.5 mm behind the MCP [8, 9].

In the environment of heavy-ion collisions the multiplicity information can be derived by analogue summing of the charge collected on the anode pads [10]. In pp collisions the multiplicity is obtained by counting individual hits (digital summing).

Owing to the excellent timing properties [11] (resolution $\sim$50 ps), the MCP detector will also provide information for the TOF system and the trigger: event time (T_0) determination, measurement of the vertex z-coordinate, identification of the beam–gas interactions, and pile-up protection for the slower detectors.

Si multipad detectors

These detectors are widely used for multiplicity measurements in heavy-ion experiments [12, 13]. The FMD could use disks formed by separate wafers with a simple pad geometry. The front-end electronics would be similar to the preamplifier chip developed for the ITS Si strip detectors (see Chapter 2).

The main limitation when using Si detectors in the FMD is their poor timing characterstics. For both online vertex determination and for the TOF T_0 start signal

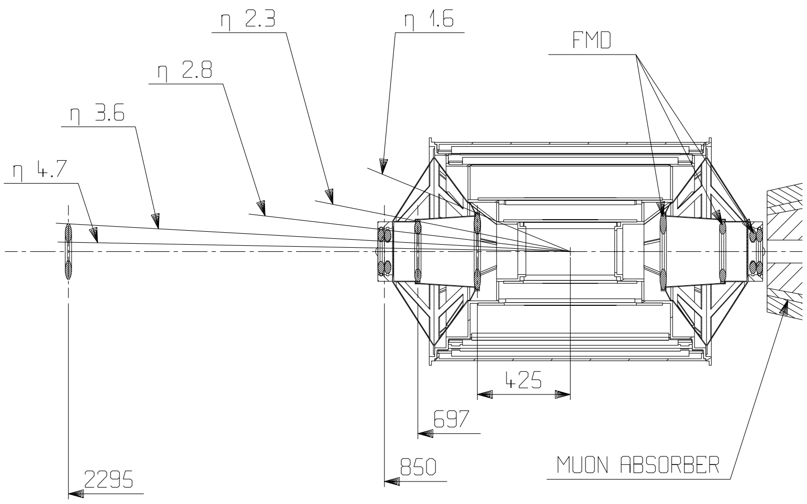

Figure 7.10: Location and support structure of the Forward Multiplicity Detectors (FMD).

we need a timing resolution of significantly better than 100 ps. This is not possible with silicon detectors, whose output signals have a rise time of at least 20–40 ns. However, for multiplicity measurements they are well suited and we might use them in combination with MCPs, e.g. in the position of Disk 1 (see Fig. 7.10). Disk 1 has the largest area and is of limited use for timing owing to spreads in path length and signal transit time as well as the small Lorentz boost of particles in its acceptance ($\beta < 1$).

7.2.2 Design considerations

Geometry

The geometrical design is similar for both FMD options. The rapidity coverage is achieved by a set of multipad disks, located symetrically on each side of the interaction region close to the beam pipe. The general layout of the FMD is shown in Fig. 7.10. This arrangement makes an efficient use of the available space and allows the installation of the ITS pixel detector and the beam-pipe heating jacket.

Multiplicity and timing resolution

The physics programme of ALICE foresees the use of different colliding beams (pp, Ca–Ca, Pb–Pb). Therefore, the dynamic range of the FMD output signal should not be less than 10^3, in order to cover multiplicities from tens (in pp runs) to thousands (Pb–Pb runs) per unit of pseudorapidity. The multiplicity resolution for heavy-ion collisions, where the occupancy per pad is much larger than one, should be below or comparable to the Poisson fluctuations of the number of particles hitting each pad.

The time resolution of each MCP disk is ~ 50 ps, which will provide a precise T_0 signal. The vertex location along the beam can be determined with independent measurements made by each pair of MCP disks to an accuracy of ~ 2 cm at the trigger level.

Radiation hardness

The long-term stability of the MCPs depends on the total accumulated charge at the MCP output: after about 0.1 C/cm^2, corresponding to 10^{12} particles per cm^2 with a gain of 10^6, the gain decreases by about 50 % [14]. As gain losses can be compensated to a large extent by an increase of the bias voltage, we do not anticipate any degradation of the efficiency during at least 10 years of operation.

7.2.3 Design of the FMD

Detector

The FMD is segmented into seven disks which surround the beam pipe at distances of between ~ 42 and 225 cm from the vertex (see Fig. 7.10). Together they will cover the pseudorapidity range from 1.6 to 3.6 on the muon arm side and from 1.5 to 4.7 on the opposite hemisphere. Most pads in the FMD disks will cover in θ a pseudorapidity range of 0.1, and in φ a range of 30°, i.e. 120 pads will cover one unit of pseudorapidity. The pad size of the two largest disks (at $z = 42$ cm) is smaller by a factor of two (i.e. $\Delta\varphi = 15°$) in order to reduce degradation of the resolution from signal transit time spread (see Fig. 7.11).

The size (inner and outer radius r_i, r_o), location in z, and the number of pads for each of the disks is summarized in Table 7.1. In total 780 pads will be read out.

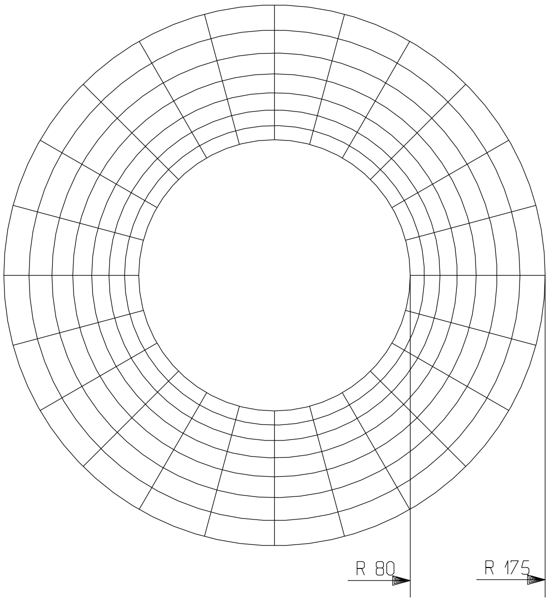

Figure 7.11: Geometry of the pad readout plane of the FMD disks located at $z = 42.5$ cm.

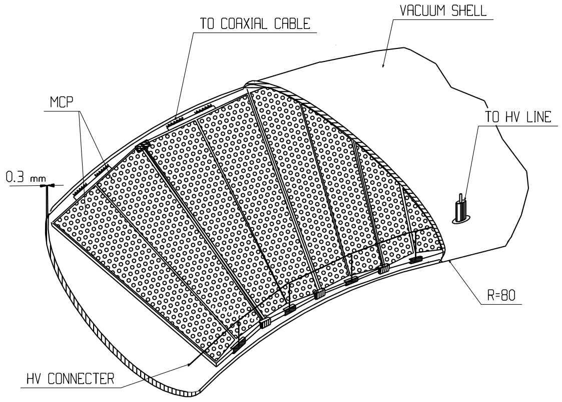

Figure 7.12: MCP disk positioned inside a thin wall vacuum chamber.

Table 7.1: Characteristics of the FMD disks (Disk 4 on one side only).

	Rapidity coverage	$r_\mathrm{i} - r_\mathrm{o}$ (mm)	Pads	z (cm)
Disk 1	1.6–2.3	80–175	168	42.5
Disk 2	2.3–2.8	80–140	60	69.7
Disk 3	2.8–3.6	45–105	96	85.0
Disk 4	3.6–4.7	42–130	132	229.5

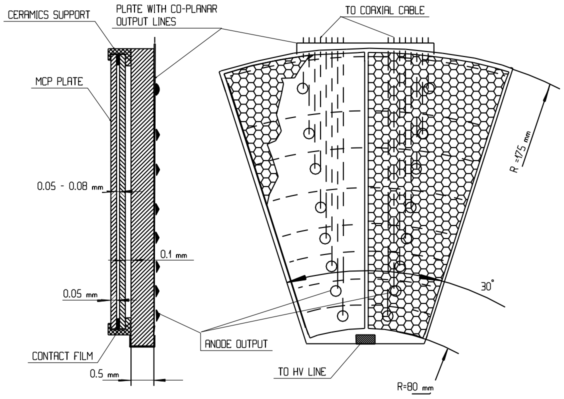

Figure 7.13: Design of a FMD segment.

The FMD disks are divided into two half-rings in order to install them around the beam pipe; a module is shown in Fig. 7.12. Each module has a ceramic backplane equipped with 6 or 12 trapezoidally-shaped MCP segments. The alumina ceramics backplane of 500 μm thickness supports the MCP sectors and integrates the pad readout on the front and signal transmission lines on the back. This layout will introduce a dead area of about 3.5 % caused by inefficiencies at the edges of the MCPs. Details of the largest MCP segment are shown in Fig. 7.13; the other disks are similar in design. The sectors are mounted and tested separately before being assembled in the FMD.

All FMD modules are independent, vacuum sealed, and equipped with a getter pump which will be activated after the module is baked at 150°C. The vacuum shell will be made of stainless steel (0.2–0.3 mm) or glass (1–1.5 mm). High-frequency high-vacuum feedthroughs are used to extract the signals. The chambers are supported around the beam pipe with a lightweight carbon fibre structure (see Fig. 7.10). We foresee the possibility of moving apart the two halves of the MCP disk closest to the muon arm remotely, in order to have more free space during the beam pipe installation.

The time resolution of the MCP requires the use of transmission lines, feedthroughs, coaxial connectors and front-end inputs with a bandwidth of $\sim$5 GHz to avoid any distortion of the signal rise time. We plan to use non-dispersive transmission lines [15,16] for this purpose. The transmission lines have to be isochronous, i.e. all delays to the front-end preamplifiers or summators must be the same within 10 ps (3 mm accuracy of the cable length). The transmission lines are coupled to standard coaxial 50 Ω feedthroughs welded onto the stainless steel shell (or fused into the glass container).

The readout signal, with a pulse length of 1 ns, will reach the preamplifiers through a 10–30 cm long cable. The integration of the front-end electronics and FMD module are presently under study.

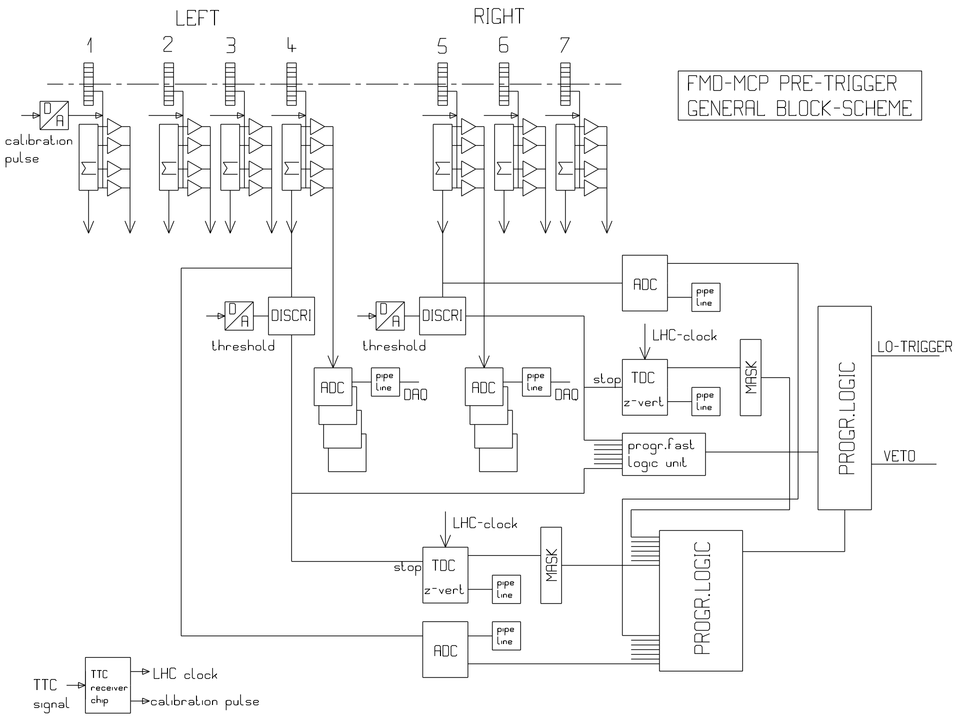

Figure 7.14: Block diagram for the FMD and pre-trigger electronics for seven multipad disks.

Readout electronics

The general scheme of the FMD electronics is shown in Fig. 7.14. The charge collected from each pad will be preamplified and converted in an 8-bit nonlinear ADC. Fast ADCs, presently under investigation at JINR, achieve a conversion time of below 50 ns. The digitized signals will be delayed in a pipeline clocked by the LHC clock and read out in the event of a L1 trigger. The pipeline has to be approximately 100 cells deep, corresponding to the fixed latency of the synchronous L1 trigger. The full pipeline can be read out for monitoring and testing purposes whenever neccessary. The data flow (0.4 MByte/s) from each side of the detector (4 FMD disks on the right and 3 disks on the left) will be transmitted by the detector data link (DDL) and collected in the local data concentrator (LDC) after a L2 trigger (see Chapter 10 for details).

A part of the pad readout signal is split by a capacity of 1–5 pF. The signals obtained from a complete (or half) MCP disk will be added in isochronous passive summators ($\delta t < 10$ ps) [17]. The output signal, containing T_0 and the multiplicity summed over a large number of readout pads, will be transmitted from the detector region to the local trigger crate situated close to the L3 magnet. Special fast transmission lines, using optical fibres, are under investigation [18, 19].

Calibration

During a calibration run, the time and trigger control (TTC) system [20] will be used to inject calibration pulses, correlated with the 40 MHz LHC clock, into the front-end electronics in order to calibrate the gain of the FMD electronics and to tune the timing of the trigger. The gain of the MCP detectors must be checked during a normal running period, using amplitude spectra obtained with the ADCs.

7.2.4 L0 trigger

In the local trigger crate the averaged signals from the individual FMD disks are converted by FADCs to extract a global multiplicity per disk. The signals will also be discriminated and fed into high-resolution TDCs to obtain the average event time (T_0 signal) for each disk. The discriminator output will also allow interactions to be monitored on the basis of individual bunch crossings.

The information from all the FMD disks will create the L0 trigger in programmable fast-logic units. The multiplicity information will be used to select centrality (impact parameter) in heavy-ion reactions or minimum bias interactions during pp operation. The timing information, together with a left–right symmetry requirement, will be used to reject beam–gas interactions, and to obtain a vertex selection. The trigger decision must be fast (< 300 ns), and the electronics must provide a new decision every

25 ns.

If a new interaction is detected after an L2 trigger and during the drift time of the slower detectors (TPC, SDD), a veto signal will be sent to the trigger system.

In addition, we envisage pipelining the data forming the L0 trigger for each bunch crossing, in order to monitor the trigger (see Chapter 9).

Details of the L0 trigger hardware are presently under study [21].

7.2.5 Multiplicity determination

The multiplicity of pp events will be evaluated by counting the number of hit pads in the FMD. This is feasible, since the noise of the MCP is only about 1 Hz per cm^2 of sensitive area.

In high-multiplicity events, during Pb–Pb runs, the multiplicity will be unfolded from the total charge collected from the individual pads using the known single particle detector response function. This method has been successfully applied in heavy-ion experiments with Si multiplicity counters [12, 13].

The multiplicity measurement in MCPs can be of a quality similar to that as in silicon detectors [10]. Figure 7.15 shows the simulated multiplicity resolution ($\sigma/M_{\rm p}$ in per cent) as a function of multiplicity M_p per pad. These simulatations make use of the appropriate single particle response functions of Si pad and MCP detectors (see Fig. 7.16).

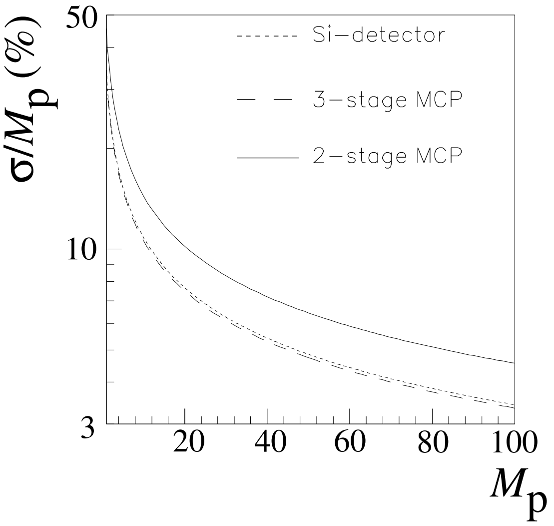

Figure 7.15: Multiplicity resolution as a function of multiplicity M_p for three types of single particle response function (Si detector, chevron MCP set-up, and three-stage MCP set-up).

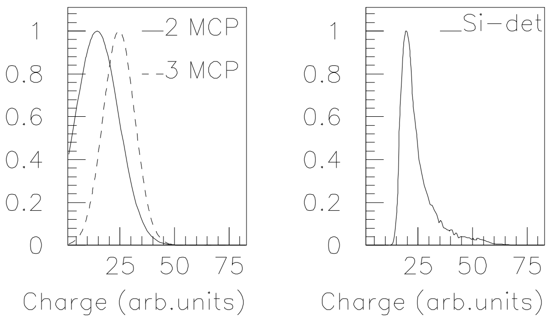

Figure 7.16: Single particle response functions for two-stage MCP, a three-stage MCP, and a Si detector.

Table 7.2: MCP prototype parameters

Material	$(SiO_2)PbO$
Density	3.89 g/cm^3
Sensitive area	43 × 63 mm^2
Thickness	770 μm
Channel diameter	15 μm
Channel pitch	17 μm
Channel bias angle	7°
Resistance	$6.4 \times 10^8\ \Omega$
Gain at 1.14 kV	1.6×10^4

Using the three-stage MCP set-up one can get a resolution of ∼3.5 % at high multiplicities, which is identical to the performance calculated for the Si detector. For the two-stage MCP arrangement, one can expect to determine multiplicities with an accuracy of $\sigma = 5 - 10\,\%$ for M_p between 20 and 100 charged particles per pad. For all cases, the resolution is better than $1/\sqrt{M_p}$ (the statistical significance) which is relevant for example in the search for structures in $dN_{\rm ch}/d\eta$.

7.2.6 R&D and tests

We have tested several rectangular FMD prototypes with a sensitive area of 43 × 63 mm^2. Two MCP plates, stacked in a chevron geometry, have been assembled in front of 3 × 4 pads with an area of 13 × 15 mm^2 (see Table 7.2). The pads and the coplanar transmission lines were integrated into a plate of alumina ceramics in order to cope with the high vacuum and high-frequency bandwidth requirements [8].

An efficiency of (80 ± 10) % for MIPs was obtained in the first measurements [22], in good agreement with the 90 % efficiency published in Ref. 23.

The measurements of the time resolution were performed with 800 MeV/c pion and proton beams [24]. Two detectors, separated by a distance of 34 mm, were placed in a beamline inside a vacuum chamber. A FWHM of 285 ps was obtained for the coincidence time distribution, corresponding to $\sigma = 85$ ps for a single MCP detector.

In order to test the effect of magnetic fields on efficiency and amplification, we placed commercially available MCPs with 6 μm channel diameter and 300 μm thickness, in a 1 T field [25]. The MCPs were irradiated by a 1 mCi Cs γ-source placed 5 cm away. Pulses from the MCP anode were amplified by a charge-sensitive preamplifier and converted in an ADC. The MCP performance did not deteriorate for either perpendicular or parallel orientation in the magnetic field.

The timing resolution has so far only been estimated by simulating electron propagation inside an MCP placed perpendicular to a magnetic field of maximum value 1 T [25]. These calculations predict that the time resolution should be effected at most by 20–30 ps.

We plan to produce a prototype sector module of the FMD to test the UHF performance of the readout (cross-talk, uniformity of the isochronous transmitting lines, fast summators) and to study the multiplicity resolution in a beam of relativistic heavy ions.

7.3 Detectors under study

In addition to the baseline design described above, we are currently evaluating two additional detector elements. Their physics goals, performance, and construction will be described in detail in an addendum, if and when the collaboration feels confident about their need and feasibility. Here we outline briefly the current status.

7.3.1 Scintillator barrel

Requirements

The two main trigger detectors of ALICE, the ZDC and the FMD, are specially designed and well suited to trigger for minimum bias or central events in heavy-ion collisions. These detectors do not need to have acceptance around mid-rapidity because in ion collisions particle production is correlated via the impact parameter in all regions of phase space. However, in pp reactions a certain bias at low multiplicity might occur because the ZDC cannot be used, and the FMD does not cover the acceptance of our central detectors.

In addition, neither the ZDC nor the FMD can be used to trigger on cosmic ray muons. Such a trigger has proven to be extremely useful for the LEP detectors, in particular during the running-in phase, for testing, calibration and alignment purposes of large volume tracking devices. We are currently investigating whether our large area TOF array (Pestov or PPC counters) can provide a trigger signal of sufficient quality for selecting cosmic ray tracks crossing the experiment. However, they consist of about 10^5 individual counter elements and are not designed to select single tracks. We have therefore made some preliminary investigations into using a scintillator barrel (ScB) (for details see Ref. 26) —also employed by most LEP detectors— whose task would be to provide trigger coverage in the central region and for cosmic rays.

Mechanical design

The ScB should cover the central range, $|\eta| < 0.9$, where the main ALICE detectors will be installed. The length of the ScB along the beam axis is 5.6 m, its inner radius is 2.75 m, and the outer radius is 2.95 m. It will be placed between the TPC and the PID system.

A general view of the scintillator barrel is presented in Fig. 7.17. It consists of 24 supermodules. The granularity of the current design has been made sufficiently large to provide a trigger both for low multiplicity events (cosmics, pp minimum bias) as well as for central heavy-ion interactions. Each supermodule (1.5×2.75 m^2, see Fig. 7.18) includes two layers of overlapping scintillation tiles (140 tiles in the supermodule, and 3360 tiles in the full barrel). The tile sizes in the external and internal layers are $0.6 \times 17 \times 18.5$ cm^3 and $0.6 \times 34 \times 37$ cm^3, respectively. The tiles in the supermodules are supported by a honeycomb structure. The transverse radiation length of the barrel is 3 % X_0.

Four tiles in the external layer of the supermodule overlap one tile from the internal layer. The light from the tiles goes through optical fibres to multi-anode PMs (e.g. two Hamamatsu H5828 PMs for each supermodule, see Fig. 7.18), which can operate in a magnetic field of up to 1 T parallel to the PM axis. We have measured a light yield of more than 10 photo-electrons per crossing particle, resulting in a detection efficiency close to 100 %. The signals from the overlapping tiles are fed into coincidence circuits to suppress the PM noise.

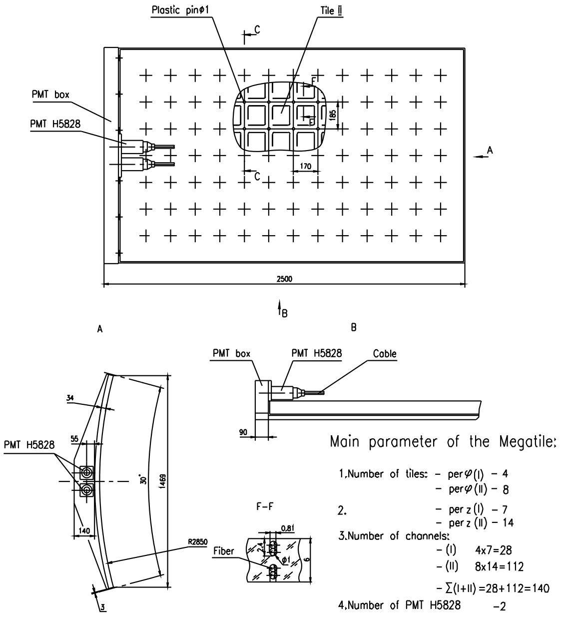

Figure 7.18: A supermodule of the scintillator barrel.

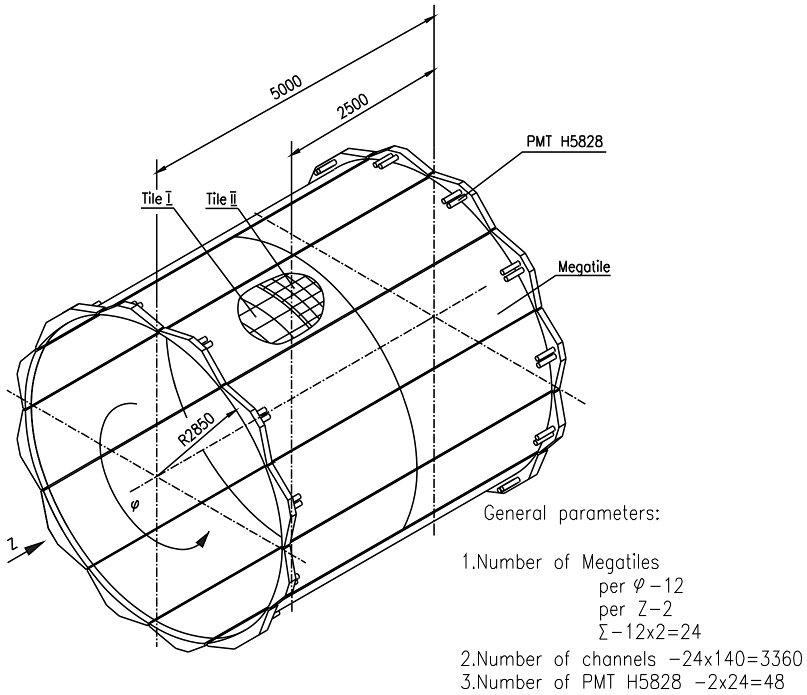

Figure 7.17: General view of the scintillator barrel.

Electronics

The trigger electronics of the ScB can operate in two different modes, digital or analogue. In the first case, for low multiplicity events, we count the total number of hits in the barrel. For high multiplicity events, where the occupancy will be much larger than one, we perform an analogue sum of pulse heights, which will be proportional to the total number of particles traversing the ScB. In both cases a trigger signal, requiring a minimum number of charged particles in the acceptance, can be generated within 200 ns.

Trigger performance

The performance of the ScB in pp and heavy-ion runs can be found in Ref. 26. Here we only summarize how it functions as a cosmic ray trigger.

The intensity of the charged penetrating component of cosmic rays (mainly muons) crossing a horizontal area is 1.3×10^2 m^{-2} s^{-1}, at sea level. The rate in the ScB will be of the order of 4×10^3 s^{-1}. Taking into account that testing runs can take as long as two months (5×10^6 s), we can expect up to 2×10^{10} cosmic ray events per run.

The cosmic ray trigger would work asynchronously. The logic of the trigger should be based on 'hit' signals (H-signals) from 24 supermodules of the ScB. A H-signal from one of the 12 supermodules in the upper hemisphere of the barrel can be used as the master signal to open a 40 ns time gate. In the most simple case the master signal should be in coincidence with an H-signal from one of the 12 supermodules in the lower hemisphere during the gating time. The trigger rate is $\sim 4 \times 10^3$ s^{-1} in this case, while the noise rate will be negligible ($\sim$ 2 events/s for a PMT noise of 10^4 Hz per channel).

7.3.2 Photon Multiplicity Detector (PMD)

Requirements

Most major heavy-ion experiments and essentially all general-purpose collider detectors employ large acceptance electromagnetic and/or hadronic calorimeters for triggering and event characterization. In ALICE, calorimetry around mid-rapidity would cover a very large surface, as tracking and PID detectors have to extend to radii of almost 4 m, in order to cope with the large multiplicities. Therefore, we rely on charged particle multiplicity measurements which are very well correlated with the transverse energy, at least on average.

We are currently investigating whether some information about transverse energy could be obtained with a thin electromagnetic pre-shower detector placed slightly outside our central acceptance. It would enable us to measure and select E_t, and to investigate energy flow or event-by-event fluctuations in the ratio of neutral to charged particles. The PHOS e.m. calorimeter, whilst well suited to measure inclusive photon spectra, is only of limited use in this context because of its small acceptance ($\sim 10\,\%$ of one unit of rapidity). Measuring the photon multiplicity

N_γ in single events together with the charged particle multiplicity N_{ch} becomes possible at high multiplicities where statistical fluctuations are small. At the LHC, in central Pb–Pb collisions, $dN_{ch}/d\eta$ is at least 2000, even at $\eta \approx 4$. The N_γ/N_{ch} distribution for central Pb–Pb collisions from the HIJING event generator is shown in Fig. 7.19. In normal hadronic reactions, where approximately equal amounts of positive, negative and neutral pions are produced, N_γ/N_{ch} is of the order of one. As can be seen in the figure, statistical fluctuations in this ratio amount to only a few per cent in heavy-ion reactions when measured over a sufficiently large phase space ($\Delta y > 1$). One can therefore distinguish, on an event-by-event basis, exotic processes with different values of N_γ/N_{ch}, as predicted, for example, for isospin fluctuations in the QPG phase transition (disordered chiral condensates) or hinted at in cosmic ray events (CENTAUROs).

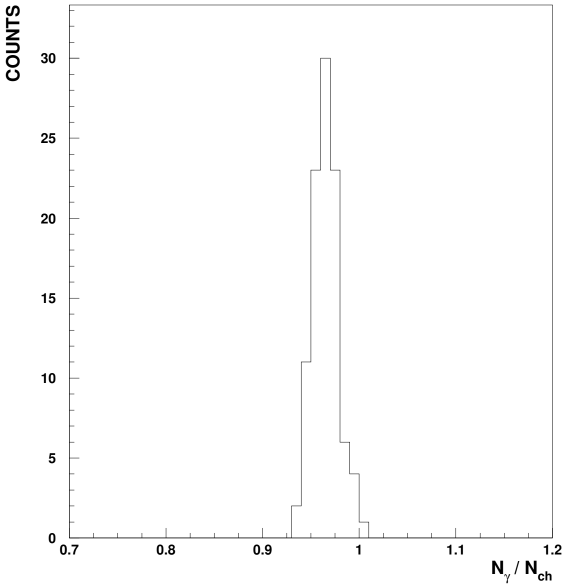

Figure 7.19: N_γ/N_{ch} for central Pb–Pb collisions from the HIJING event generator in the region $1.5 \leq \eta \leq 3.5$.

Detector design

In order to measure the ratio N_γ/N_{ch} and E_t^{em} on an event-by-event basis, a very fine granularity pre-shower Photon Multiplicity Detector (PMD), covering the pseudorapidity region $1.5 \leq \eta \leq 3.5$, could be added to ALICE [27]. The PMD will have a charged particle veto (CPV) in front of it. The part of the detector covering the region $1.5 \leq \eta \leq 2.5$ will be housed inside the magnet, about 5.5 m from the vertex. With a weight of 5 t and a thickness of less than 50 cm, it could be fixed to the door of the L3 magnet. The second part, weighing 6 t and covering the region $2.5 \leq \eta \leq 3.5$, could be positioned about 18 m from the vertex. The total area covered by the PMD amounts to about 40 m^2.

Other options, for example a single PMD detector at a suitable distance from the vertex, are also being considered.

The basic operation principles of a pre-shower PMD are described in Ref. 28. The PMD, similar in design to existing devices used in WA93 and WA98, consists of a lead converter followed by a radiation-sensitive medium divided into a matrix of pads. The CPV detector in front of it will be used for distinguishing photons from charged particles. Charged particles, which are needed to construct the ratio N_γ/N_{ch}, will also be measured within the same acceptance using the FMD (see Section 7.2).

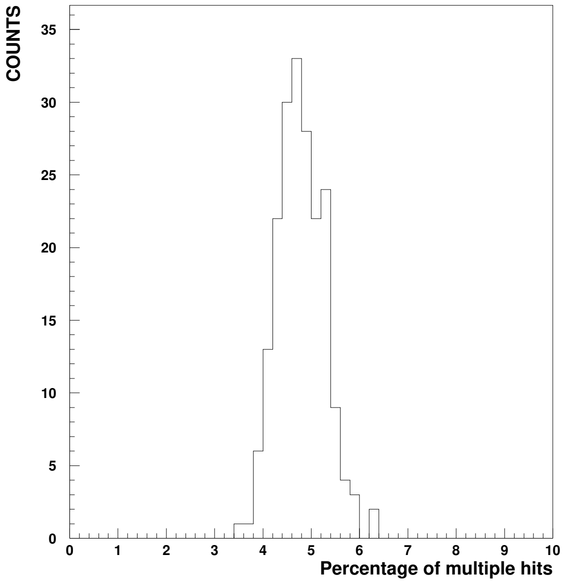

Figure 7.20: Multiple hit probability in a PMD with 12 mm^2 pads placed 18 m from the vertex, covering the region $2.5 \leq \eta \leq 3.5$.

The response of the PMD has been studied with Monte Carlo simulations using HIJING and GEANT [27]. As an example, Fig. 7.20 shows the pad occupancy for a WA93-like PMD [28] with a 3 X_0 thick converter and 3 mm thick plastic scintillator pads of 12 mm^2 size, located 18 m from the vertex. The pad occupancy, defined here as the ratio of the number of pads to the total number of charged particles and photons hitting the detector, should be a few per cent, in order to avoid shower overlap.

The criteria for selecting the technology for a pre-shower detector are the following: (a) MIP and photon separation for energies as little as 0.5 GeV; (b) noise level (detector + readout) ≤ 0.2 MIP; (c) large dynamic range and good linearity; (d) fast response to be used in the first-level trigger; and (e) stable operation at rates of up to 5 kHz/cm^2. Two options are presently under investigation.

Silicon pads – The first option is a silicon pad detector, similar to the CMS pre-shower detector [29]. The main disadvantage of silicon, besides cost, is its sensitivity to radiation damage, which might be a minor problem in ALICE. The use of low-resistivity silicon (1–1.5 kΩ cm), which is both cheaper and less sensitive to radiation, will be investigated. The readout electronics developed by the RD36 Collaboration [30] could easily be adapted for the PMD as the pixel areas for both the CMS and the ALICE pre-shower detectors are similar.

Gaseous Detectors – The other option is to use gaseous proportional chambers, either in a configuration similar to the drift tube pre-shower detector of NOMAD or using novel schemes such as cellular proportional chambers with wire readout. The same technology could be used both for the CPV and the pre-shower part. In addition to rate limitations, the main problem of gaseous detectors is the increased shower size (due for example to long-range delta rays), which limits the achievable granularity. Individual cells surrounded by metal walls, which can stop or rescatter the low-energy electrons, could help in this respect. Preliminary simulations have shown that in a chamber with hexagonal cells and using a gas mixture of 80 % Ar and 20 % CO_2, cell walls of 0.5 mm Al reduce the mean number of hit pads from 3.4 to 2.0. Further optimization and hardware tests will be carried out.

N_γ/N_{ch} measurement

In order to count the number of photons and charged particles, a clustering algorithm is used to identify hits in the PMD. Hadrons are rejected by applying a threshold equivalent to 3 MIPs on the total cluster pulse height. At this threshold level the average photon detection efficiency is 65 % for central events and 75 % for peripheral events (including the conversion probability of 85 %). The number of hadrons not identified by this cut can be estimated using information from the CPV detector in front of the PMD.

The performance of the ALICE PMD can be extrapolated from the simulations and beam test results of the WA93 PMD [28] which is now well understood. In addition, the following improvements are foreseen for ALICE:

- The photon counting efficiency should be higher because of the finer granularity (occupancy $< 5\,\%$ in ALICE, compared with $\sim 10\,\%$ for WA93). The use of a veto, not present in WA93, allows a lower pulse height cut to reject charged particles, which should further improve the photon detection efficiency.

- The systematic error in the estimated number of photons should considerably decrease because of event-by-event subtraction of the background, instead of using a mean value obtained from Monte Carlo simulation. The statistical error will also decrease because of the very large multiplicity of photons.

E_{t} measurement

The pre-shower PMD will also provide an estimate of transverse e.m. energy, $E_{\mathrm{t}}^{\mathrm{em}}$. Although there are large fluctuations in the fraction of energy actually deposited by each single particle in a pre-shower detector, the sum over many particles is a good measure of the total transverse e.m. energy.

As seen in Fig. 7.21, which shows the ratio of measured over true e.m. energy, $E_{\mathrm{t}}^{\mathrm{em}}$ can be extracted from the PMD with an accuracy of the order of a few per cent. Even for differential distributions, i.e. $\mathrm{d}E_{\mathrm{t}}/\mathrm{d}\eta$, the resolution stays at about 10 % for $\delta\eta = 0.1$. This resolution is sufficient for studying event shapes in heavy-ion reactions, in a similar way as with more traditional full absorption calorimeters.

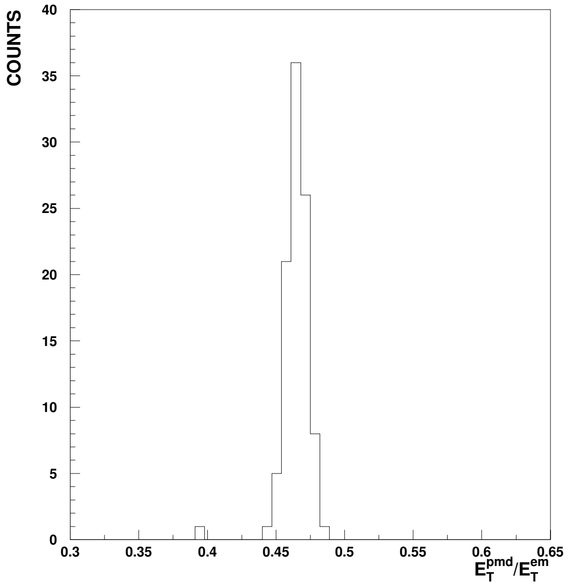

Figure 7.21: The ratio of $E_{\mathrm{t}}^{\mathrm{PMD}}$ to $E_{\mathrm{t}}^{\mathrm{em}}$ for the PMD covering the region $1.5 \leq \eta \leq 2.5$.

8 Luminosity considerations

8.1 Introduction

Achievable luminosities at the LHC filled with ions depend critically on the ion accumulation and injection schemes [1]. Since runs with different ion masses are of interest, these schemes have to be optimized taking into account space charge effects during accumulation and transfers between the machines involved. For example, a bunch distance of 125 ns gives the highest luminosity for lead whereas a time structure of 25 ns is optimal for calcium beams.

In a number of cases the luminosity is limited by the short beam-lifetime, a consequence of the large cross-sections for electromagnetic processes, which either disintegrate the ions or change their charge due to electron capture (see, however, Section 8.5). These electromagnetic processes are strongly dependent on the ion charge and hence much larger luminosities can be expected for lighter ions.

The luminosity ($\mathcal{L}$) at the LHC is given by the number of bunches (k) and the luminosity per bunch ($\mathcal{L}_{\text{bunch}}$) which depends on the number of particles per bunch (N) and the transverse beam radius (σ_t), both varying with time [2, 3],

$$\mathcal{L} = k\,\mathcal{L}_{\text{bunch}} \propto \frac{N^2}{\sigma_t^2}\ .$$

The number of particles per bunch decreases due to interactions of the particles in the crossing points (beam–beam) and to interactions with the remaining gas in the vacuum chamber (beam–gas) [4]. The transverse beam size (σ_t) grows due to the Intra-Beam Scattering (IBS). With heavy ions in the LHC, the cross-sections for beam–beam interactions are huge, exceeding 300 barn for Pb–Pb collisions. Compared to this, beam–gas interactions are small and hence are neglected in the luminosity lifetime calculations.

The beam decay time (t_{bb}) is inversely proportional to the initial number of ions per bunch (N_0) the number of interaction points (n_x) and the total cross-section (σ_{tot}) but is independent of the number of bunches (k):

$$N(t) = \frac{N_0}{1 + t/t_{\text{bb}}}$$

$$t_{\text{bb}} = \frac{N_0}{\sigma_{\text{tot}}\,\mathcal{L}^0_{\text{bunch}}\,n_x} \propto \frac{1}{\sigma_{\text{tot}}\,N_0\,n_x}$$

The luminosity half lifetime ($t_{1/2}$) is related to the beam lifetime through

$$t_{1/2} = (\sqrt{2} - 1)\ t_{\text{bb}}\ .$$

The strong increase of the total ion cross-section as the atomic number of the ion increases is limiting the maximum achievable luminosity.

8.2 Beam–beam interactions

Three main processes contribute to beam–beam interactions in which the ions are lost from the beam:

- Hadronic nuclear interactions [5].
- Electromagnetic dissociation where an ion is excited by a γ–nucleon interaction and subsequently decays [6–8]. The largest contribution to this process comes from the excitation of the giant dipole resonance which decays mainly by particle emission or fission.
- Two-photon process creating an electron–positron pair, followed by the subsequent capture (mainly K-shell) of the electron by an ion [6, 9].

The nature of the collision process, electromagnetic or nuclear, depends upon the value of the impact parameter. The hadronic nuclear interaction is dominant when the colliding ions overlap. Hence, the nuclear interaction rate increases approximately with the geometrical overlap of the two ions. A reasonable description of present data is provided by the hard-sphere overlap parametrization [10],

$$\sigma_{\text{nucl}} = \sigma_0 \left(A_1^{1/3} + A_2^{1/3} - \text{b}\right)^2 ,$$

where A_1 and A_2 are the mass numbers of the two ions, $\text{b} \approx 1.3$ and $\sigma_0 \approx 70$ mb. A more refined description is the semi-classical optical model [11], which also provides analytical expressions for the nuclear cross-section.

The electromagnetic processes, in contrast, involve impact parameters much larger than the range of nuclear forces. Their relative importance rapidly increases with the charge Z of the ions. The electromagnetic interaction of the peripheral colliding ions can be described by the reaction of equivalent photons (Weizäcker–Williams method of virtual quanta [7]). Simple A and Z dependencies for the individual cross-sections have been derived for the processes listed below [6]. We use these relations in order to calculate the cross-sections for ions lighter than Pb. Normalizing to the Pb–Pb results [8, 9] gives:

- Electromagnetic Dissociation at the Giant Dipole Resonance

$$\sigma_{\text{GDR}} \approx 67.5\,\frac{(A-Z)Z^3}{A^{2/3}}\ (\mu\text{b})\ .$$

- Electromagnetic Dissociation above the Giant Dipole Resonance

$$\sigma_{\text{ED}} \approx 108\,A^{0.9}Z^2\ (\mu\text{b})\ .$$

Table 8.2: Standard LHC Parameters for the Heavy Ion (Pb–Pb) Mode.

Atomic Number of Pb	Z	82
Mass Number of Pb	A	208
Initial Luminosity	$\mathcal{L}_0$	$0.85 \times 10^{27}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$
Initial Number of Ions/Bunch	N_0	6.3×10^7
Number of Bunches	k	560
Bunch Occupation	k/k_0	0.848
Revolution Frequency	f	11.2 kHz
Bunch Interval	Δt	124.75 ns
Ion Energy	E	7 TeV·Z
γ	E/m_p	7516·Z/A
β at Interaction Point	$\beta^\star$	0.5 m
Transverse Normalized Emittance	ε_m	4.75 μm·Z/A
Longitudinal Emittance	$\varepsilon_l = 4\pi(\sigma_s/c)\sigma_E E$	2.5 eVs·Z
Relative r.m.s. Energy Spread	σ_E	1.14×10^{-4}
Crossing Angle	Φ	$< 100\mu$rad
Average β	$\hat{\beta}$	62 m
Average Dispersion	$\hat{D}$	1.29 m
Luminosity Half Lifetime	$t_{1/2}$	10 h

Table 8.1: Cross-sections for ion-ion interaction (in barn).

Ion	Nucl.	e.m. Dissoc.	e^- Capture	Total
Pb^{208}	7.8	222	101	336
Nb^{97}	4.9	23.4	3.1	31.4
Ca^{40}	3.1	2.1	4×10^{-2}	5.2
O^{16}	1.5	0.12	8×10^{-5}	1.6
He^4	0.35	2×10^{-3}	6×10^{-9}	0.35
p	0.105	–	5×10^{-11}	0.105

- e^- capture

$$\sigma_{e^+e^-} \approx 2.1 \times 10^{-6}\,\frac{Z^8}{e^{2\pi\alpha Z}-1}\,(\mu\mathrm{b})\,.$$

In Fig. 8.1 the total cross-sections and the cross-sections for the individual processes are plotted as a function of the ion charge using A and Z of some stable nuclei. The nuclear cross-sections calculated with the semi-classical optical model [11] using nuclear radii from Ref. 12 are shown as points for pp, He–He, O–O, Ca–Ca, Nb–Nb, and Pb–Pb. The individual ion–ion cross-sections are shown in Table 8.1.

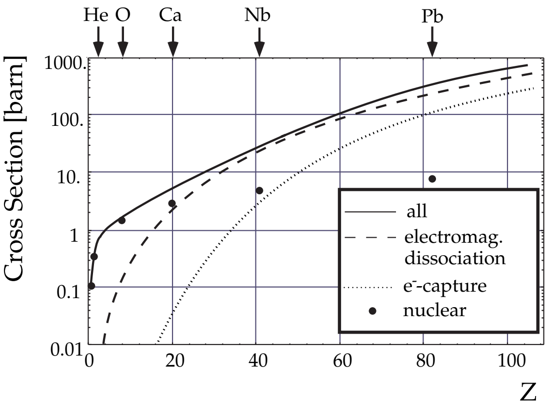

Figure 8.1: Z-dependence of ion–ion partial cross-sections for selected ion pairs: total cross-section (solid line), electromagnetic dissociation (dashed line), pair production electron capture (dotted line), and nuclear interactions (solid points).

8.3 Intra-beam scattering

An additional effect is the Intra-Beam Scattering (IBS) which expands the beam dimensions during the run. The scattering cross-section is proportional to the square of the classical ion radius, i.e. $\propto Z^4/A^2$. Remembering that all the parameters in the IBS calculations (longitudinal and transverse emittances, energy factor) have their own A and Z dependencies, in the high-energy limit the IBS beam growth rate scales in good approximation with NZA [13, 14].

The emittance increase after a sufficiently small time step Δt can be approximated as

$$\varepsilon_{x,z,l}(t+\Delta t) = \left(1 + \frac{\Delta t}{\tau_{x,z,l}(\sigma, N, \ldots)}\right)\varepsilon_{x,z,l}(t)$$

Here, $\varepsilon_{x,z,l}$ are the transverse (x, z) and longitudinal (l) emittances, and $1/\tau_{x,z,l}$ are the respective emittance growth rates characterizing the process. They are themselves functions of the beam dimensions and the number of particles per bunch.

Inserting the standard LHC parameters [15] for Pb–Pb collisions, as listed in Table 8.2, obtained the initial transverse emittance growth time is for arbitrary A, Z and N [16]:

$$\tau_{x,y} = 26\,\mathrm{h}\,\frac{208}{A} \times \frac{82}{Z} \times \frac{9.4 \times 10^7}{N}\ .$$

The parameters determining the growth-time change with the collider running-time due to particle losses and the intra-beam scattering process itself. This leads to a relatively fast blow-up at the beginning of the run and a slower growth at the end. To calculate the time evolution of the luminosity, one has to find a self-consistent solution of the particle loss equation, including the beam radius growth $\sigma_\mathrm{t}(t,N)$:

$$\frac{1}{N}\frac{\mathrm{d}N}{\mathrm{d}t} = -\frac{1}{t_\mathrm{bb}}\frac{N}{N_0} = -\mathrm{const.}\frac{N(t)}{\sigma_\mathrm{t}^2(t,N(t))}\ .$$

8.4 Average luminosity for various ion species

The number of signal events accumulated during a given collider running time is given by the integrated luminosity (average luminosity times running time). A realistic calculation of the average luminosity for a running cycle of the LHC must take into account the filling time (t_f) of the machine and the waiting time (t_w) when the beams are colliding but experiments are not yet operational [17]. Thus, the average luminosity for a collider operating in physics conditions for a time t_r is

$$\frac{\langle\mathcal{L}\rangle}{\mathcal{L}_0} = \frac{1}{t_\mathrm{f}+t_\mathrm{r}}\int_{t_\mathrm{w}}^{t_\mathrm{r}}\frac{\mathcal{L}(t)}{\mathcal{L}_0}\mathrm{d}t\ .$$

We used $t_\mathrm{f} = 4$ h and $t_\mathrm{w} = 1/3$ h as standard parameters. The parameters for heavy-ion collisions used in the following calculations are derived from those listed in Table 8.2 [15]. As mentioned above and indicated in this table, some of the parameters change with Z or Z/A. Figure 8.2 shows the time evolution of the normalized luminosity $\mathcal{L}(t)/\mathcal{L}_0$ and the average luminosity $\langle\mathcal{L}\rangle/\mathcal{L}_0$ for the ions Pb (A = 208, Z = 82), Ca (A = 40, Z = 20) and He (A = 4, Z = 2). The number of ions per bunch was tuned such that the luminosity half-lifetime amounts to 10 h for one experiment, which corresponds to $\langle\mathcal{L}\rangle/\mathcal{L}_0 = 45\,\%$ after 20 h running time. In order to illustrate the relative importance of beam–beam interactions and IBS, we have plotted the time evolution with and without IBS. In the case of Pb, the luminosity decrease is dominated by the beam–beam interaction. For ions with Z in the range 10 to 40, such as Ca, IBS has its maximum influence, while for light ions $Z < 10$ its influence is very small.

As shown previously, the luminosity lifetime due to beam–beam losses strongly decreases with increasing atomic number and is inversely proportional to the number of ions per bunch. Thus, by fixing the luminosity lifetime the number of ions per bunch can be increased for lighter ions. Since the luminosity is proportional to the square of the number of ions per bunch, a significant increase in luminosity is expected. In order to gain a comprehensive understanding of the effects of changing

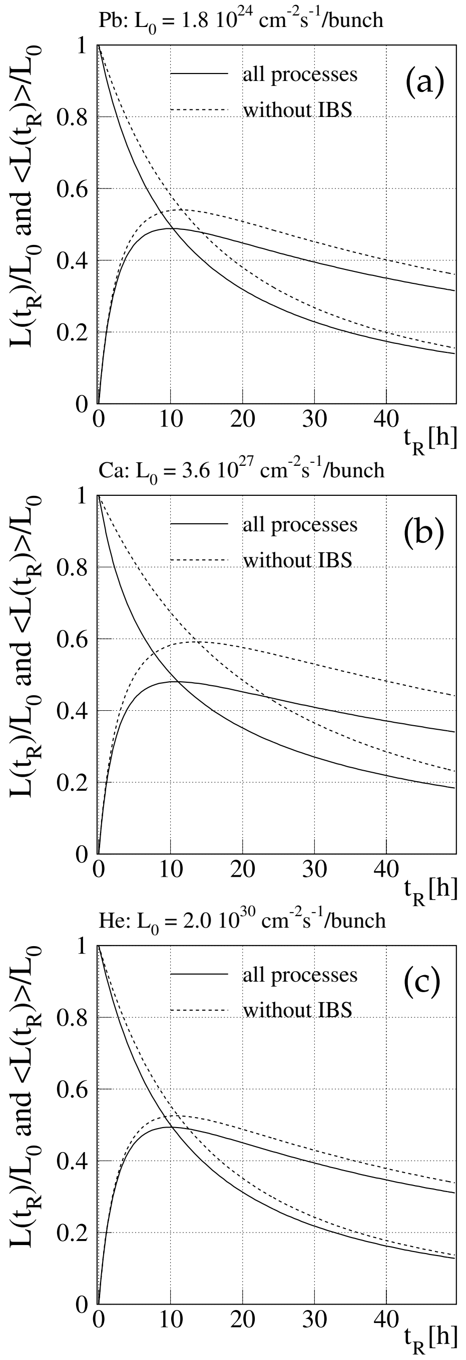

Figure 8.2: Luminosity and average luminosity as a function of the collider running time for (a) Pb–Pb (b) Ca–Ca and (c) He–He collisions.

Table 8.3: Initial luminosities (upper values, in units of $cm^{-2}s^{-1}$) and number of ions per beam (lower values) at a fixed luminosity lifetime ($t_{1/2}$ = 10 h). The average luminosity will reach 45 % of the initital value for a running time of 20 hours.

Ion	1 Exp./25 ns	1 Exp./125 ns	2 Exp./25 ns	2 Exp./125 ns	3 Exp./25 ns	3 Exp./125 ns
p	6.9×10^{34} 7.5×10^{14}	1.3×10^{34} 1.3×10^{14}	1.9×10^{34} 3.9×10^{14}	3.6×10^{33} 6.8×10^{13}	8.7×10^{33} 2.8×10^{14}	1.6×10^{33} 4.6×10^{13}
He	5.5×10^{33} 2.1×10^{14}	1.2×10^{33} 4.5×10^{13}	1.7×10^{33} 1.2×10^{14}	3.6×10^{32} 2.5×10^{13}	8.4×10^{32} 8.2×10^{13}	1.8×10^{32} 1.8×10^{13}
O	1.4×10^{32} 3.4×10^{13}	3.0×10^{31} 7.2×10^{12}	5.3×10^{31} 2.1×10^{13}	1.1×10^{31} 4.4×10^{12}	2.8×10^{31} 1.5×10^{13}	5.9×10^{30} 3.2×10^{12}
Ca	8.7×10^{30} 8.3×10^{12}	1.9×10^{30} 1.8×10^{12}	3.6×10^{30} 5.4×10^{12}	7.7×10^{29} 1.2×10^{12}	2.0×10^{30} 4.0×10^{12}	4.3×10^{29} 8.6×10^{11}
Nb	3.0×10^{29} 1.6×10^{12}	6.5×10^{28} 3.3×10^{11}	1.2×10^{29} 9.9×10^{11}	2.6×10^{28} 2.1×10^{11}	6.8×10^{28} 7.4×10^{11}	1.5×10^{28} 1.6×10^{11}
Pb		9.9×10^{26} 4.1×10^{10}		3.3×10^{26} 2.4×10^{10}		1.7×10^{26} 1.7×10^{10}

to lighter ion species, we show in Fig. 8.3 the luminosity half lifetime for different ion species as a function of the initial number of ions per bunch and the corresponding initial luminosity per bunch for either one or two experiments. The horizontal displacements of the curves measure the increases in luminosity for a fixed luminosity lifetime. Assuming a lifetime of 10 h and 1 experiment, the initial luminosity can be increased by a factor of 2000 if Ca instead of Pb ions are used. Doubling the number of experiments decreases the luminosity by a factor of 3.0 for Pb and 2.5 for Ca, reflecting the reduced importance of the beam–beam interaction for lighter ion species.

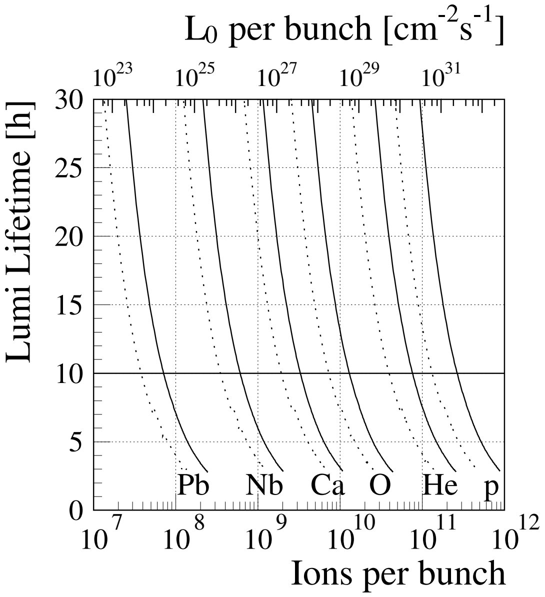

Figure 8.3: The luminosity half lifetime as a function of the initial number of ions per bunch (initial luminosity) for one (full-line) and two (dotted-line) experiments.

Table 8.3 gives an overview of the initial luminosities and total number of ions per beam at a fixed luminosity lifetime ($t_{1/2}$ = 10 h) for all ion species considered. The luminosity lifetime of 10 h corresponds to an average luminosity of 45 % for a running time of 20 h. In addition to the values for the nominal Pb–Pb bunch spacing of 125 ns (k = 560), we quote the corresponding numbers for a 25 ns bunch spacing (k = 2870), a possible scenario for lighter ions. As discussed in the previous section with respect to the luminosity lifetime, much higher intensities could be envisaged for lighter ions compared to those for Pb ions. However, it remains to be demonstrated that such intensities can be accommodated along the whole chain of injectors. Keeping in mind that in the standard injection scheme retained for the Pb ions (including accumulation in LEAR) [1] the space charge effect already approaches its limits, it becomes obvious that the problem would be even worse for increased intensities of light ions; such a scheme cannot be envisaged. Both the space charge problems and the accumulation of ions lighter than Pb are studied in detail in Ref. 18. This report also presents a new scheme which consists in injecting the ions in LEAR with a low charge state and accelerating them to an energy high enough so they can be fully stripped just after extraction. In order to avoid expensive modifications to the machine, LEAR would work in a single bunch mode and the required accumulation would then take place in the PS. As mentioned in Ref. 18, this new scheme requires further investigation; however, it looks promising and could allow high intensity runs with light ions at the LHC.

8.5 Constant luminosity scenario

As explained above, the upper limit on the achievable luminosity with Pb is given by the constraint of a minimum acceptable beam lifetime. This limit can be raised by

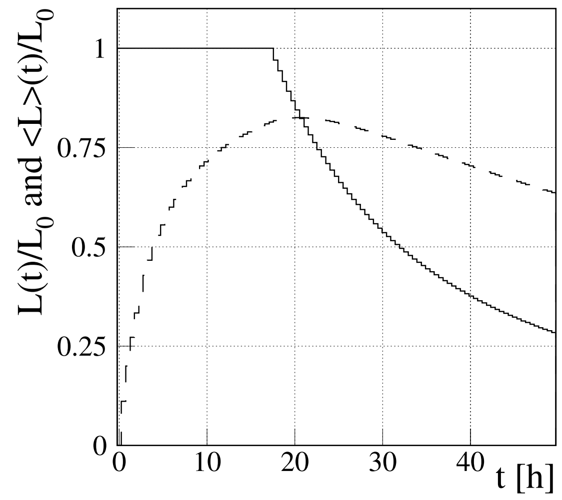

Figure 8.4: Luminosity and average luminosity as a function of the collider running time for Pb–Pb collisions in a variable $\beta^\star$ scheme. The initial luminosity was fixed 3×10^{27} cm^{-2} s^{-1} with $N_0 = 2\times10^{11}$. By reducing $\beta^\star$ the luminosity remains constant until $\beta^\star_{\rm min} = 0.5$ m is reached.

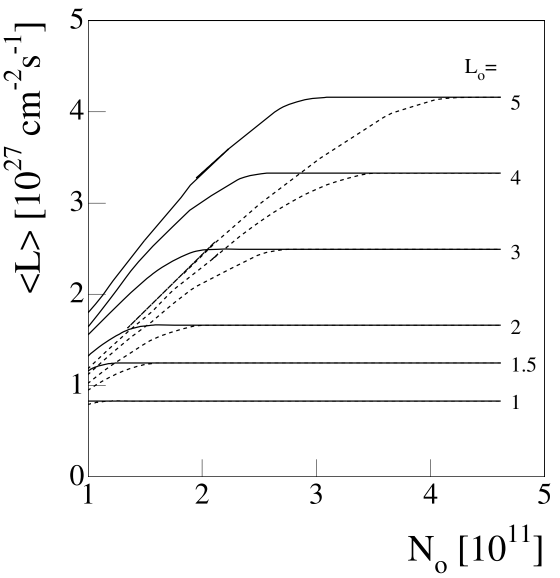

Figure 8.5: The average luminosity as a function of the initial number of ions for Pb–Pb collisions, running with one experiment (solid line) and with two experiments (dashed line).

about an order of magnitude if the beams can be filled with more ions and if $\beta^\star$ at the intersection point can be adiabatically lowered during the run.

For a given heavy ion species, i.e. fixed beam–beam cross-section, the luminosity lifetime is proportional to $N_0/\mathcal{L}^0_{\rm bunch}$, where N_0 is the initial number of ions per bunch and $\mathcal{L}^0_{\rm bunch}$ the initial luminosity per bunch. Thus, for a given luminosity obtained with a larger $\beta^\star$ the luminosity lifetime rises in proportion to the number of particles, and the average luminosity can be increased. It is then possible to reduce $\beta^\star$ continuously during the run down to its minimal value of 0.5 m in such a way that the luminosity remains constant, at least initially. According to preliminary investigations, such a continuous slow decrease of $\beta^\star$ should pose no particular problems. It has obvious advantages for the running conditions of the experiments (trigger rates, space charge in the TPC, etc.) and considerably increases the average luminosity. Of course, the proposed scenario has its own limitations. The maximum luminosity is now more severely limited by the number of ions that can be accumulated and transferred into the LHC. In addition, intra-beam scattering, which is proportional to the number of particles per bunch, will dominate the luminosity lifetime for very high N_0, leading to a fast emittance blow-up at the start of the run.

To demonstrate the advantage of a constant luminosity with larger ion-beam intensities we consider the case of one experiment with Pb beams (Table 8.3). If we increase the initial number of ions per beam from 3.5×10^{10} to 2×10^{11} and the initial luminosity to 3×10^{27} cm^{-2} s^{-1} at $\beta^\star = 4.9$ m, the luminosity lifetime increases by approximately a factor of 2. Furthermore, it is now possible to reduce $\beta^\star$ during the run while keeping the luminosity constant. The evolution of luminosity and average luminosity for this scenario is shown in Fig. 8.4. After a running time of 20 h the average luminosity amounts to 2.5×10^{27} cm^{-2} s^{-1} for $\beta^\star_{\rm min} = 0.5$ m, seven times larger than for the scenario in Table 8.3.

The average luminosity for Pb–Pb collisions after 20 h collider running time, as a function of the initial conditions $\mathcal{L}_0$ and N_0, is plotted in Fig. 8.5 for $\beta^\star_{\rm min} = 0.5$ m for one and two experiments. For every initial luminosity there exists an optimal initial number of heavy ions per bunch, beyond which the average luminosity does not increase further. This number corresponds to a situation where the luminosity can be kept constant over the full running period with $\beta^\star_{\rm min} = 0.5$ m at the end of the run. In the standard scenario, where $\beta^\star$ stays constant and the initial parameters are given by fixing the luminosity lifetime, the initial luminosity decreases with the square of the number of experiments. For the scenario considered here, the influence of the number of experiments depends on the actual choice of running conditions, but in general it is weaker than in the standard scenario.

The corresponding curves for Ca runs are shown in Fig. 8.6. Similar gains to those for Pb can be obtained.

8.6 Runs with protons

For comparison and normalization ALICE will frequently run with proton beams, in parallel with the other experiments, at reduced luminosities of around 10^{30} cm^{-2} s^{-1} or less. Two methods of reducing the luminosity by more

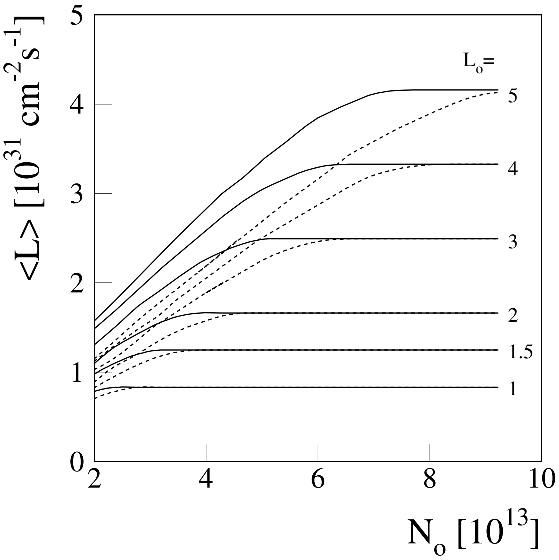

Figure 8.6: The average luminosity as a function of the initial number of ions for Ca–Ca collisions running with one experiment (solid line) and with two experiments (dashed line).

than three orders of magnitude in one Intersection Point (IP) are considered. One is the displacement of the beams such that only the tails collide. Since the beam tails can be subject to large fluctuations, the luminosity might be unstable. However, as described in Ref. 19, this scheme seems to be a viable one.

A far superior alternative would be to increase the transverse dimensions of the beams by increasing the β-value at the IP. In a standard insertion two ranges of β-values can be achieved by powering the insertion magnets differently: one covering 0.5–10 m and the other 300–1000 m. This gives sufficient flexibility in the luminosity adjustment. However, at large β-values the crossing angle has to be enlarged from 0.1 mrad to about 0.4 mrad to avoid strong long-range tune shift.

In the latest design scheme of the LHC there is one beam-accelerating superconducting RF cavity for each beam. This provides a most welcome opportunity to collide different ion species up to the extreme of p–Pb collisions.

8.7 Beam induced background

The beam-related background will scale with the beam current rather than with the interaction rate in the interaction region and is therefore potentially most dangerous during pp running, when we limit the luminosity in our crossing point to 10^{30} cm^{-2} s^{-1}. The following estimates are therefore made for pp runs, assuming the design current of 530 mA per proton beam, which will be reached after a few years of operation.

Beam–gas interactions

Beam–gas interactions can be vetoed at the Level 0 (L0) trigger by our multiplicity array, which will determine the event vertex by TOF to within a few cm (see Chapter 9). Therefore, we only need to consider random coincidences between a true pp interaction and an upstream event. The detector most affected will be the TPC because of its long memory time (± 100 μs) and because of the large area (~ 16 m^2) presented to the background particles, which arrive at small angles relative to the beam.

The rate of beam–gas interactions in the warm straight sections ($50 < z < 200$ m) is expected to be about 200 Hz per metre per beam and is mainly due to protons interacting on oxygen and carbon nuclei. This will lead to a total rate of beam–gas interactions —with produced particles heading towards the detector— of $< 10^5$ Hz. However, most of the straight section is very well shielded, mainly by machine elements, shielding of the tunnel entrance and, on one side, by the muon-arm absorbers. A calculation done for the outer CMS muon chambers [20] shows that in practice only muons from interactions far downstream ($>$ 20 m) reach the experiment. The flux of muons from heavy quark (b, c) and π/K decays is $\sim 5 \times 10^{-4}$/m^2 per interaction at the location of our TPC ($r = 1$–2 m). This corresponds to a muon flux of less than 1000 muons per second through the area covered by the TPC, or < 0.2 muons per TPC memory time. Therefore, this source of background causes no problem, with a large safety margin.

The small unshielded region (± 20 m) close to the detector is more problematic. The vacuum in this region is determined by the pumping rate and the beam pipe treatment, i.e. bake-out. With a reasonable amount of pumping, which presents no problem in our open geometry, the vacuum is slightly better than in the warm straight sections (see Section 12.5.1) and we expect about 50 Hz/m per beam at the design current. This will lead to a beam–gas interaction rate of 2×10^3 Hz, or < 0.5 interactions per TPC memory time. For interactions ~ 20 m away from the normal vertex, about 5 charged particles will cross the TPC per interaction and a comparable number from photon conversions or interactions in the beam pipe. This is a worst-case scenario, as beam–gas interactions closer to the beam crossing will have fewer tracks at the large angles required to reach the TPC.

The TPC and our tracking algorithms can easily cope with such a small number of background tracks in pp collisions as they are designed to track more than 10 000 particles in Pb–Pb. In addition, these tracks are almost parallel to the beam and can easily be distinguished from true vertex tracks. Therefore, we do not expect that beam–gas interactions will pose a problem for the central ALICE detectors. Their influence on the muon chambers and muon trigger will be investigated in the context of the muon arm technical proposal.

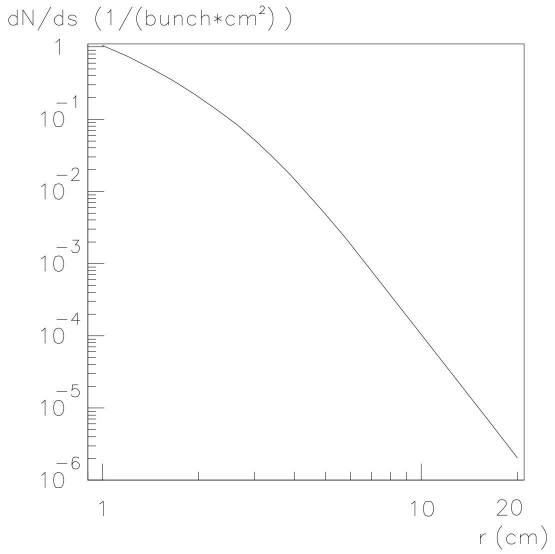

Figure 8.7: Total number of QED electrons passing through a detector perpendicular to the beam axis per bunch crossing and cm^2, as a function of radial position r (Pb–Pb, $\mathcal{L} = 2 \times 10^{27}$ cm^{-2}s^{-1}, $B = 0.2$ T, $z = 10$ cm).

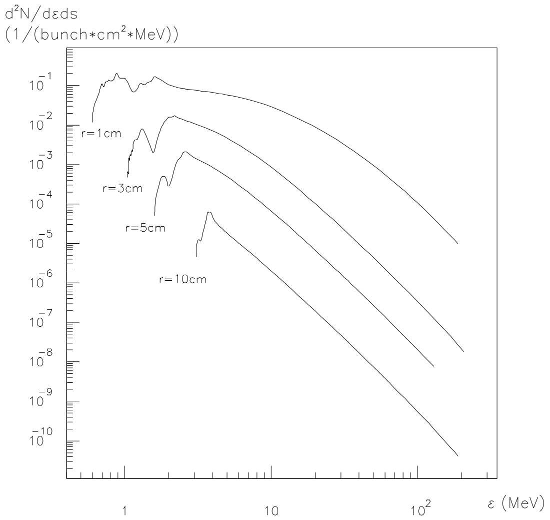

Figure 8.8: QED electron energy distribution (per bunch crossing, cm^2 and MeV) on a detector perpendicular to the beam axis at $z = 10$ cm for different radial positions r.

Beam losses

Local beam losses in any of the machine elements are a potentially dangerous source of background. They can occur either in the separation dipoles, the focusing triplet or the small angle absorber, where the apertures are small and the β oscillation is large. In particular the small angle absorber is a major concern for the high luminosity pp experiments. However, this element, designed to protect the low beta quadrupoles from energy deposit by beam–beam interaction secondaries, might not be necessary or have a large aperture in our low luminosity insertion. In case we run at very high $\beta^\star$ (see above), the effect of a larger crossing angle in ALICE is compensated by the smaller focusing, leaving sufficient clearence ($>20\,\sigma$) in the dipole and quadrupole magnets.

QED electron pairs

A background unique to very high energy ion beams is due to the large cross-section for electron pair creation via electromagnetic interactions (QED pair production). While the rates are large, most of these electrons are produced at small angles with respect to the beam and with very small transverse momenta.

The flux of QED pairs in Pb–Pb collisions at design luminosity (2×10^{27} cm^{-2} s^{-1}) is shown in Figs. 8.7 and 8.8 [21]. The results are shown as the number of electron hits per cm^2 and per bunch crossing as a function of distance to the beam. The actual number of random background hits from QED pairs in our detectors depends on position, area and integration time (memory time or gate opening), but is negligible in all cases. It ranges from < 10 in the silicon drift layers of the ITS (drift time 6 μs) to < 1 in the TPC (drift time 100 μs).

8.8 Running scenario

In Table 8.4 our running scenario is summarized for pp, light ion (e.g. Ca–Ca) and heavy ion (e.g. Pb–Pb) collisions. The table shows average luminosities, $\langle\mathcal{L}\rangle$, minimum-bias interaction rates and minimum bias charged-particle densities, $dN_{ch}/d\eta$ at $\eta = 0$, calculated with PYTHIA for pp and HIJING for nuclear collisions (see also Section 11.2 and note that $dN_{ch}/d\eta$ is about 20 % lower than dN_{ch}/dy at $y = 0$). Asymmetric collision systems (e.g. p–A) are not mentioned as the expected luminosities are not yet known.

Table 8.4: Average luminosity $\langle\mathcal{L}\rangle$ (in units of cm^{-2} s^{-1}), minimum bias event rate and minimum bias charged particle multiplicity $dN/d\eta$ in pp, low and high luminosity Ca–Ca, and Pb–Pb runs.

	pp	Ca–Ca	Ca–Ca	Pb–Pb
$\langle\mathcal{L}\rangle$	10^{30}	2.7×10^{27}	10^{29}	10^{27}
Rate (Hz)	10^5	3×10^5	8000	8000
$dN_{ch}/d\eta$	7.2	225	225	1500

8.8.1 pp running

In order to keep the pile-up in the TPC and SDDs at an acceptable level, the luminosity during pp runs has to be limited to about 10^{30} cm^{-2} s^{-1}, corresponding to 20 minimum bias pp events during the TPC drift time. This pile-up will be acceptable for pattern recognition, as the track density is still more than an order of magnitude below that in central Pb–Pb events. The number of pp events collected per year will be two orders of magnitude larger than with heavy ions (a factor of ten due to the longer running time and a second factor of ten due to the increased data taking rate), which will largely compensate the smaller particle multiplicity in pp, leading to comparable statistical errors in both cases [19,22,23].

8.8.2 Ca–Ca running

With light ions, we plan to run at two different luminosities in order to optimize the statistics both for the central ALICE detectors (including muon arm) and for the muon arm in stand-alone mode [23]. Reasonable pile-up conditions in the TPC require the event rate to be below 10^4 Hz, which is reached at a luminosity of 2.7×10^{27} cm^{-2} s^{-1} (see Chapter 9). Separate high luminosity runs, where we operate only the muon spectrometer and the pixel detectors (for vertex determination), are foreseen in order to increase the statistics needed for the Debye screening measurement.

8.8.3 Pb–Pb running

As our measurement of Debye screening will be limited by statistics, we request the highest possible luminosity for Pb–Pb collisions. At 10^{27} cm^{-2} s^{-1}, the number of accepted events in the TPC (i.e. events without pile-up) will be at its maximum (see Chapter 9). Such an average luminosity is a factor of two greater than in the current design and will require tuning of the focusing as described above.

If substantially higher luminosities could be reached, the Pb runs would also be split into high and low luminosity parts, as we have described for Ca, with an almost linear increase of statistics with luminosity, which is very beneficial for the muon programme [23].

9 The trigger

9.1 Introduction

The ALICE trigger system has to operate in both the ion–ion and proton–proton collision mode. In the ion–ion mode its principal function is to enhance centrality. It should be capable of selecting a wide range of centralities chosen either according to multiplicity values at mid to forward rapidities using the Forward Multiplicity Detectors (FMD), or according to forward energy using the Zero Degree Calorimeters (ZDC). The trigger should work for a variety of different ion species. It is also envisaged that certain types of event, for example those in which a high p_t dimuon pair is produced, could be given priority in the event stream.

In the proton–proton mode, the trigger is intended to select minimum bias events, again with the possibility of defining a 'priority stream' of rarer event types, such as dimuon events.

The trigger system is also required to provide past–future protection for the ALICE detectors. The TPC is the slowest of these, with a drift time of $\sim 100\,\mu$s, though the Si drift detectors, with a drift time of $\sim 5\,\mu$s, also require some care. In general, a pile-up is not tolerable in the ion–ion mode, owing to the very large multiplicities of the events. For this reason, the trigger must ensure that for each sub-detector no second event occurs in a period given by the read out/drift time of the detector, either before or after the event of interest, i.e. in a $200\,\mu$s interval for the TPC. In contrast, in the proton–proton mode, up to 20 superimposed events in the TPC can be tolerated.

ALICE is required to operate with widely differing luminosities. The mean luminosity in Pb–Pb is expected to be $\mathcal{L} = 10^{27}$ cm^{-2} s^{-1}, and $\mathcal{L} = 2.7 \times 10^{27}$ cm^{-2} s^{-1} in Ca–Ca collisions. The Ca–Ca luminosity has been chosen in order to give the same interaction rate (about 8000 Hz) in Ca–Ca as in Pb–Pb collisions, the limiting factor being the pile-up in the TPC.

Substantially higher luminosities, up to around $\mathcal{L} = 10^{30}$ cm^{-2} s^{-1}, are achievable in Ca–Ca collisions. It will not be possible to use the TPC at such high luminosities, nor is it required for the majority of our physics goals. However, Debye screening studies with the muon arm will be statistics limited. We therefore plan to collect dimuon events at a high rate in Ca–Ca interactions by reading out the dimuon arm in conjunction with a limited set of 'fast' central detectors, including the pixel layers of the ITS, to define the vertex position. The trigger accommodates this mode of operation; for every event flags are set to indicate which detectors are to take part in the data acquisition.

The bunch crossing rate in the proton–proton mode is 25 ns, while in the Pb–Pb mode it is 125 ns. The trigger electronics will be designed to run with the LHC 40.08 MHz clock in order to be able to process the proton–proton mode events on a crossing-by-crossing basis. This allows full flexibility for the high rate Ca–Ca running described above, and allows past–future protection to be implemented in a flexible way.

The ALICE trigger is foreseen to work on three levels as shown in Fig. 9.1. Both the Level 0 (L0) and Level 1 (L1) decisions have a fixed latency. The purpose of the L0 trigger, based on the FMD, is to signal an interaction at the earliest possible time, i.e. after about $1.2\,\mu$s. The L0 trigger is used by some of the detectors to strobe the front-end electronics. L0 is not very selective and runs at a rate of between $< 10^4$ up to a few times 10^4 Hz, depending on the beam type (pp, Ca–Ca, Pb–Pb) and beam luminosity. Its rate is close to the minimum bias rate, ensuring that most of the available luminosity is preserved for higher level triggers to select the events of interest.

The L1 decision, also synchronous, is made at about $2.4\ \mu$s, and leads to a considerable reduction in the rate (≤ 1000 Hz). Its selection is based on information from the muon system (two muon candidates above a p_t threshold) and on the centrality detectors, FMD and ZDC. At this point all the remaining detectors are strobed. The L1 trigger is also used by the data acquisition system to allow readout to start from the detector front-end electronics.

Owing to the more complicated correlations required by the dimuon system, the fastest dimuon decision comes at L1. Dimuon triggers are rare and therefore the strict centrality requirements are not applied. For this reason L0 has loose centrality requirements, which are tightened at L1.

The $100\ \mu$s drift time of the TPC is a period during which the Level 2 (L2) decisions can be made. More selective algorithms can be applied, e.g. a mass cut on the dimuon system or a clustering analysis on the FMD data. Other algorithms could be added, each with their own processor, provided the algorithms are completed in less than $100\,\mu$s. The latency of the L2 trigger is not fixed because, in general, the algorithms will take a variable length of time.

After the L2 trigger, the data are all read out into the data acquisition system. Further event filtering and/or data reduction can be performed after event-building, where a processor has access to all the data. This stage is not considered part of the trigger and is described in Chapter 10.

9.2 Central collisions

The principal objective of the ALICE experiment is to study the hot dense medium formed in 'central' heavy-ion collisions. Central collisions correspond to relatively

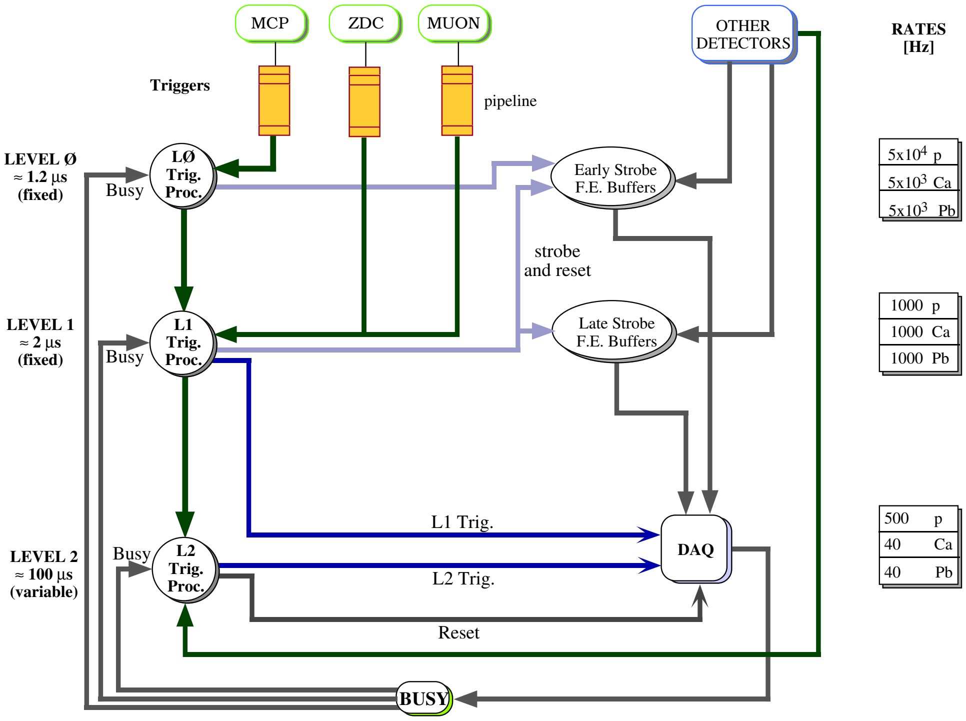

Figure 9.1: Trigger block diagram. L2 rates are for central triggers with TPC readout.

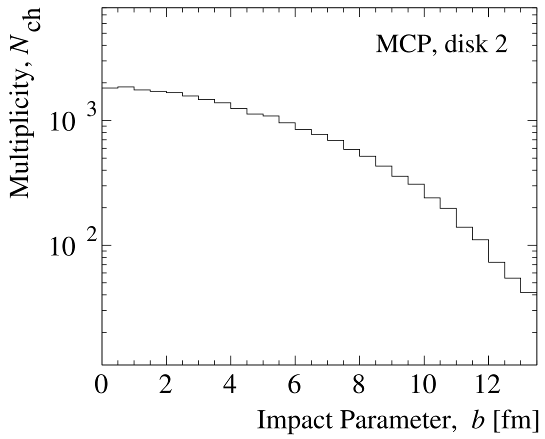

Figure 9.2: Measured multiplicity in Disk 2 of the FMD–MCP plotted against impact parameter (QGSFM Model).

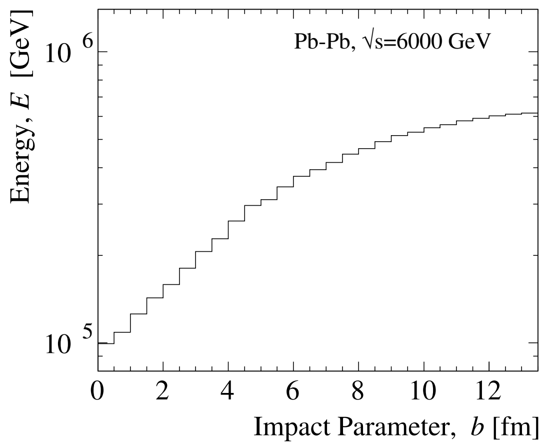

Figure 9.3: Forward energy in the ZDC plotted against impact parameter. (QGSFM Model.)

small impact parameters. Although the impact parameter cannot be determined directly, there are several measurable variables which are well correlated with it: multiplicities and/or transverse energies at mid-rapidity increase at small impact parameter, and forward energy decreases. The ALICE detector is well suited for multiplicity-based triggers, and the forward energy can be measured using the ZDC. Figure 9.2 shows the variation of measured multiplicity with the impact parameter according to the QGSFM model [1, 2], and Fig. 9.3 the variation of forward energy with the impact parameter.

Central collisions occur relatively frequently, and therefore the trigger does not need to be very selective. As an example, we show how one might select the most central 10 % of the cross-section. While in practice this would be based on multiplicity requirements on several disks, we illustrate a selection based on disk 2 alone, as shown in Fig. 9.2. The cross-section requirement corresponds to a a multiplicity cut $\mu_{\mathrm{FMD}} > 1270$ for Disk 2, and from Fig. 9.2 this selects events with impact parameter $b < 4$ fm. The impact parameter cut can be refined using the ZDC information, as discussed in Section 9.3.2.

9.3 Trigger detectors

9.3.1 The FMD

The preferred solution for the FMD is to use Micro-Channel Plate (MCP) detectors, which are symmetrically placed on both sides of the collision region. On each side, there are three FMD disks, each segmented into about one hundred pads, and covering between 0.5 and 0.8 units of pseudorapidity. In addition, there is a fourth disk on the side opposite the dimuon arm. The modules are independently sealed under vacuum (see Section 7.2).

The FMD detectors provide a multiplicity signal and a timing signal, allowing the vertex position to be determined.

The analogue pulse heights obtained from the FMD detectors are proportional to the number of hits in them. They give precise timing information, with a 1 ns width and a 100–200 ps leading edge rise time. These signals will be split to allow their use both in the L0 trigger and via a fast ADC for data recording.

In the trigger we use the pulse height analogue sum from the segments of a given detector to make selections based on multiplicity, and the time difference between corresponding left and right pads to find the vertex position. In the pp mode, where the occupancy of the pads is low, a simple count determines the multiplicity. It should be possible to determine the time difference at the trigger level to a precision of about 100 ps, corresponding to a vertex spread of ~ 3 cm (the offline time resolution is expected to be better than 50 ps).

The timing precision requirements place some constraints on the positioning of the electronics, as long cables will degrade the resolution offered by the fast signals. The front-end electronics is supposed to be positioned not far from the surface of the beam pipe. However, it is possible to add about one metre of cable in front of the preamplifiers without noticeable degradation of the timing properties. This allows some flexibility in the positioning of the preamplifiers and reduces the amount of matter in the interaction region. Also the local FMD trigger electronics (discussed in Chapter 7) should be placed as close as possible to the detector, i.e. in the experimental area (see Fig. 9.6 on page 133).

Alternatively, transmission via an optolink may offer a way to transmit the signals to the counting room with acceptable timing precision, and this solution is also being investigated.

9.3.2 The ZDC

A measurement of the energy carried away by the spectator nucleons can also provide an indication of the centrality of the collision. In the ALICE set-up the proton and neutron spectators are detected in two distinct calorimeters (ZDCs) located in front of the LHC D2 magnet 92 m from the interaction point. There will be a set of calorimeters on both sides of the intersection point. The effectiveness of the ZDC calorimeters is based on the assumption that, at least for the central collisions, the colliding ions will disintegrate fully into protons and neutrons. An experiment with emulsions using ^{32}S at 200 GeV/nucleon and triggering on minimum bias events has shown that in 16 % of the events there were no fragments with a charge greater than two [3].

Using the HIJING [4] event generator to simulate a minimum bias sample of events, one can look at the energy detected by the proposed calorimeters. The beam divergency, the Fermi momentum of the nucleons in the lead ions, the background due to the particles produced in the collisions and detected by the ZDCs, and the experimental resolution of the calorimeters, are all taken into account in the simulation. One can see from Fig. 9.3 that the impact parameter is correlated with the sum of the energies measured by the neutron and proton calorimeters (both sides). The width of the correlation is mainly due to the event-by-event fluctuations of the energy carried by the produced particles, and to the fluctuations in the number of the spectator nucleons. Even at zero impact parameter, there is a non-negligible amount of energy, corresponding to about 25 equivalent spectators per lead ion. A similar correlation can be found if one uses the VENUS [5] event generator.

By selecting different energy thresholds, it is possible to calculate how many events are accepted out of a minimum bias sample. As an example, Fig. 9.4 shows the impact parameter distribution of the events accepted when the upper thresholds on the sum of the energies from the ZDCs are 200, 400, 600 and 800 TeV. Dividing these spectra by the minimum bias distribution, one obtains Fig. 9.5, where the trigger yield with respect to the impact parameter is shown for the four selected energy cuts.

The signals from all the PMs of the calorimeters will be transmitted to the central electronic barrack. The signals from each side of the interaction point will be summed up and then the two analogue sums will be added and discriminated for a preselected energy interval. It is foreseen that the ZDC centrality trigger will be available in the electronic barrack after $\sim 1.3\ \mu$s. The analogue signals from all the PMs will be digitized in VME ADCs in the counting room, using as a gate for the ADCs the L1 trigger decision.

We expect that the neutron calorimeter will be hit most of the time by a few nucleons due to the e.m. dissociation (see Chapter 8). The lower threshold could be tuned to

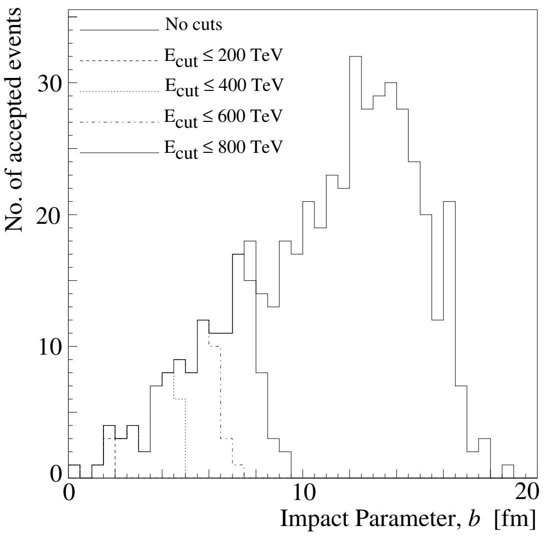

Figure 9.4: Impact parameter distributions for events satisfying different zero degree energy cuts: 200, 400, 600 and 800 TeV (HIJING Model).

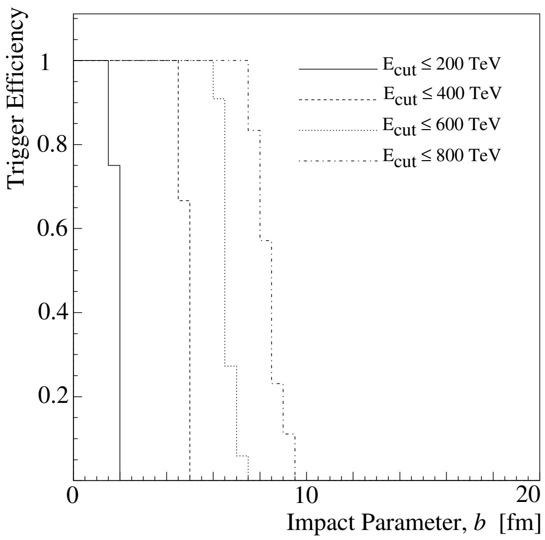

Figure 9.5: Impact parameter distributions for events satisfying zero degree energy cuts of 200, 400, 600, and 800 TeV after normalization to the overall minimum bias impact parameter distribution (HIJING Model).

reject this kind of events. Moreover, we expect that there will be less than 200 interactions per second due to the collisions of beam with the residual gas inside the pipe. These events, although not all of them produce signals in the ZDC detectors, can be reduced using the two summed signals from both sides of the interaction point and require a certain degree of symmetry.

It is foreseen that the ZDC will be used in the ion–ion mode only. In the pp mode a minimum bias trigger is required, and the very limited solid angle of the ZDC means it does not contribute any useful information.

9.4 Dimuon triggers

ALICE plans to study the dimuon spectrum for masses up to the Υ range. In this range events are very rare and special care must be taken to ensure that as few as possible are lost as a result of dead time or event pile-up. While the dimuon spectrometer itself is not part of this technical proposal, in the following we describe the conceptional ideas concerning the muon trigger and rate estimates based on the dimuon addendum to the ALICE Letter of Intent. As at this stage the rate estimates are preliminary, both the trigger and DAQ system include a large safety margin which could accommodate significantly higher rates.

Dimuon triggers form a different category from those we have considered so far. Multiplicity and zero degree energy detectors consider the global features of an event, which are also strongly correlated. In contrast, the dimuon trigger identifies a specific feature of an event (two penetrating tracks in the muon spectrometer) which is generally rare and less strongly correlated with centrality.

The first level dimuon trigger selects two penetrating tracks above a certain momentum threshold. A first trigger decision can be made in about 600 ns, based on correlations between hits in two planes of the trigger chambers. Two different correlation matrices select low or high p_t muon candidates. The low p_t cut (at ~ 1 GeV/c) selects events with $M(\mu\mu) \gtrsim 1.5$ GeV, and gives access to J/ψ events. Our current estimate for this rate is about 240 Hz. The high p_t class ($p_t \gtrsim 3$ GeV/c) yields high mass dimuon candidates in the Υ range. The rate is very low (less than 1 Hz in Pb–Pb collisions, and much less in Ca–Ca and pp collisions). The first level dimuon trigger takes part in the general L1 trigger. Events giving a 'high p_t ' trigger at this stage are flagged as high-priority events.

The second level dimuon decision involves a computation of the effective mass of the dimuon system. Its aim is to enhance J/ψ events by eliminating the low-mass dimuon events still included after L1 (high-mass events in the Υ region are already at a very low rate after L1). The algorithm studies all combinations of hits and selects those compatible with two tracks having an effective mass above a given threshold. The computation time depends on the number of combinations to be considered and can take up to about 50 μs for Pb–Pb events. The total rate for dimuon events read out together with the TPC, has to be limited to 40 Hz central events plus a comparable number of peripheral events. Such a rate can be accommodated by the DAQ bandwidth and opens the possibility of correlating dimuon physics with other observables measured in the rest of the detector. A second 'fast' data stream for a

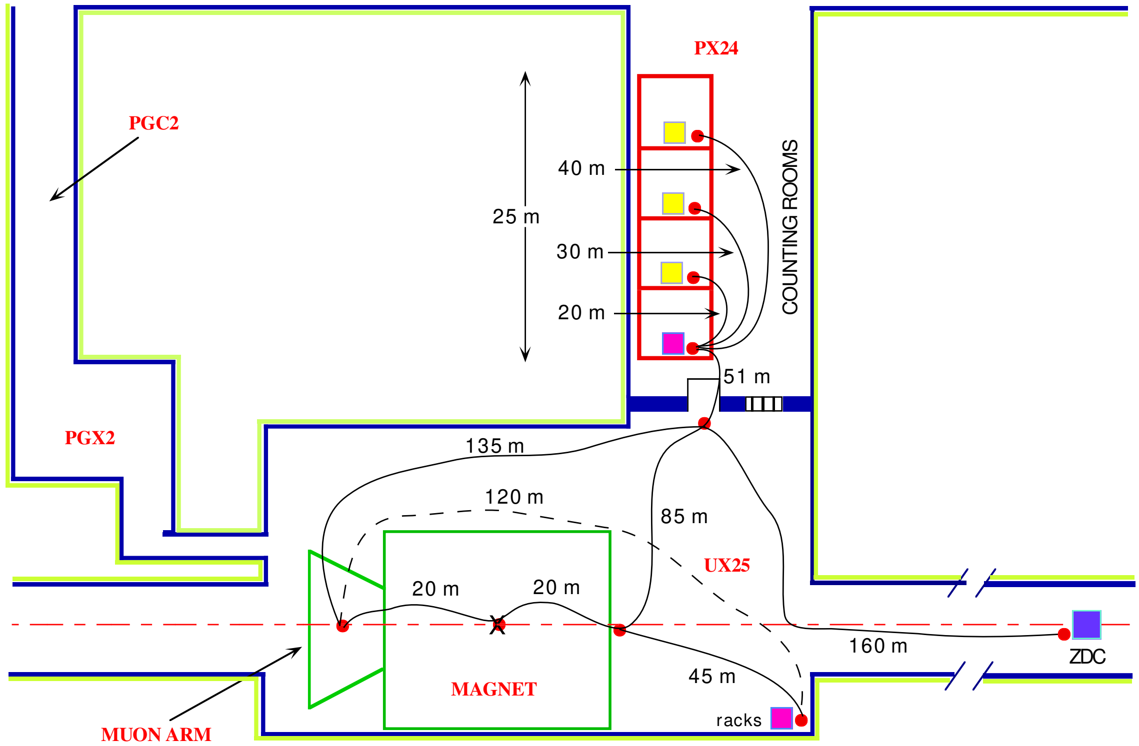

Figure 9.6: Schematic layout of the I2 pit, showing the estimated cable lengths between different parts of the ALICE detector and possible locations for the trigger electronics.

reduced set of detectors could accommodate dimuon triggers up to a rate of 1 kHz if the L2 rate exceeds the above limit.

In order to study correlations with other variables, it is desirable to read out the dimuon system in conjunction with the full ALICE detector. However, owing to the high probability of pile-up in the TPC, a fraction of the events satisfying the L2 dimuon trigger will have unusable TPC data. For this reason (i) events with a L2 dimuon trigger and no pile-up in the TPC are given the highest priority in the DAQ, and (ii) if the TPC data are not usable a reduced subset of detectors is still read out together with the dimuon arm. The readout dead time in this reduced set of 'fast' sub-detectors is kept below 200 μs, to be compared with 2 ms for the full detector. In this way the dead time overheads are kept to a minimum and are never greater than 20 %.

9.5 Level 0 and Level 1 trigger

The L0 and L1 triggers are both synchronous, fixed-latency triggers. They are separated owing to the conflicting requirements of an early trigger decision to strobe front-end electronics, based on a track-and-hold system (Si microstrips, RICH, PHOS), and the need to have the trigger in an area accessible during running and near the DAQ system in order to simplify the interface between them.

We currently plan to place the L0 trigger electronics in the underground area, close to the detector, as shown in Fig. 9.6. It will be inaccessible during running time, and special care will have to be taken to make the hardware and logic as simple and reliable as possible. We will have to foresee some redundancy and full remote control and diagnostics capabilities. The L1 electronics will have to be placed in the lowest counting room (see Fig. 9.6).

An alternative solution, also being studied, would be to combine levels 0 and 1 and place the electronics in the second access shaft at PGX2. In this case, both the L0 and L1 triggers are accessible, but communication with the DAQ is made more complicated. This solution is not discussed further here.

The L0 trigger is based on the FMD detector alone. This detector supplies information on several features of an interaction. These are summarized below for an ion–ion collision:

1. The interaction vertex is close to the nominal collision point,

2. the forward–backward distribution of tracks is consistent with a colliding beam interaction, and

3. the measured multiplicity is above a given threshold.

Details about how this is done are given in Chapter 7. Basically, the vertex definition is made from the timing asymmetry, while points 2 and 3 (beam-gas rejection, centrality) are carried out by measuring the multiplicity in each FMD disk. All three operations can be performed in $\sim$ 150 ns, with a new decision available every 25 ns.

No strong centrality condition is made at L0, as non-central events giving dimuon triggers are also required. At L1 both dimuon and ZDC information is available, and at this point more stringent centrality requirements can be made. The strong correlation between energy in the ZDC and multiplicity in the FMD allows checks on the efficiency of the trigger to be made.

In the case of L2 dimuon triggers, a 'high priority' trigger is issued, as discussed in Section 9.4. Other trigger types will be subject to downscaling. At L1 the TPC is gated, and it is important to ensure that the frequency of the gates does not exceed a maximum of about 1 kHz. All detectors which do not receive a strobe at L0 receive one at L1. The advantage of the extra latency is a 5–10 fold decrease in rate, the result of the tighter centrality cut, and the arrival of dimuon information.

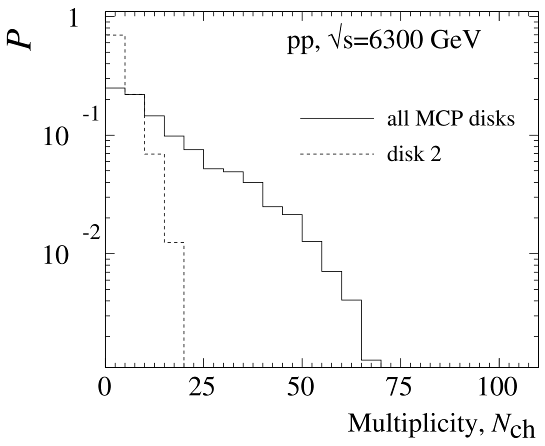

Figure 9.7: Minimum bias pp multiplicity distribution in the FMD system. Full line: total multiplicity summed over all disks. Dashed line: multiplicity in Disk 2 (QGSFM model).

The charged particle multiplicities will be much lower for minimum bias pp events than in heavy-ion interactions. Figure 9.7 shows the minimum bias pp multiplicity distribution in the FMD system as predicted by the QGSFM model. The mean multiplicity per unit of rapidity is given by $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y \approx 8$. The noise level in the MCP detectors is about 1 Hz/cm^2, i.e. extremely low with respect to the expected interaction rate of 10^5 Hz. These detectors can therefore be used in single particle counting mode at very low multiplicities.

The estimated contributions to the L1 trigger latency are given in Table 9.1. The times allowed for the electronics decisions include a safety margin. However, the major part of the delay comes from the transmission time in cables. Only approximate cable lengths can be given at present, but the estimates shown are very conservative. The L0 trigger, which involves the FMDs only and is located closer to the detector, has a latency of about 1.2 μs.

Table 9.1: Contributions to the L1 latency in nanoseconds, assuming the electronics is placed in the counting room. The sub-detectors considered are the FMD, the dimuon trigger chambers (DM), and the ZDC.

Component	FMD	Dimuon	ZDC
TOF Particle	0	50	300
F.E. Electronics	300	600	50
Cables TO	624	744	1244
Cables BACK	780	780	780
Decision time	200	200	200
Total (Central Detector)	1904	2374	2574
Total (Dimuon arm)	2179	2649	2849

Table 9.1 gives the times at which a trigger decision based on one sub-detector alone could be given to a piece of FE electronics situated (i) at the centre of the ALICE detector, and (ii) at the dimuon arm. The timing includes the time of flight from the collision point to the detector, the transit time in the front-end electronics, the transmission time to the central trigger electronics and back to the centre of the detector along the cables, and the decision time for the central trigger. The estimates given are conservative and have a safety marigin of at least 20 %. From this table one can see, for example, that for central detectors a L1 decision time involving the dimuon arm and the ZDC, could be given in about 2.4 μs.

It is crucial that sufficient time be allowed for the L1 latency, and for this reason a margin of safety has been provided. Besides making conservative estimates, the front-end electronics for the sub-detectors must be able to work for a range of possible latencies. Table 9.2 lists the expected latency for each trigger level, and also the maximum values which must be foreseen by the sub-detector electronics.

The L1 trigger electronics will be fully pipelined in

Table 9.2: Calculated trigger latencies for each trigger level. In order to protect against unforeseen increases in the latency values, sub-detectors are required to accommodate values in the range shown in the table.

Trigger Level	Estimate (μs)	Range (μs)
L0	1.2	1.0–1.5
L1	2.4	2.0–3.0
L2	Variable	4–100

order to be able to deal with bunch crossings with a 25 ns spacing on a crossing-by-crossing basis. While not strictly necessary for the rates expected in ALICE, the number of channels is small and the use of pipelined trigger detectors means that past-future protection can be handled in a flexible way. A similar approach is being taken by the STAR Collaboration at RHIC [6, 7]. The use of pipelined trigger electronics is straightforward for the ZDC, which provides four signals which can be digitized close to the L1 trigger electronics; for the FMD system, studies are underway. For the dimuon system, the electronics architecture has not been decided yet. A modified version of the system used by the ATLAS local muon trigger electronics [8] would clearly satisfy the timing requirements.

Good diagnostic capabilities are regarded as essential for the trigger allowing efficient start-up and rapid attention to subsequent problems which may arise. In order to do this, buffer space will be made available to allow:

i) full information about the crossing selected for the trigger, and

ii) condensed information about each crossing in the $100\,\mu$s interval before and after the selected collision,

to be recorded for each trigger sub-detector. The status of all the sub-detectors (live, busy, offline) will also be recorded. This allows the trigger decision, trigger efficiency and effectiveness of the past–future protection to be monitored. While the data volume per crossing is small, the requirement that this be kept for every bunch crossing in a 200 μs interval leads to considerable data volume per event. However, in general it is not necessary to accumulate large statistics to detect trigger errors. Therefore, in order to limit the number of data links required for the readout of the trigger system, only a sample of events will be read out in this manner. This could either be a random sample taken from the data stream, or those events trapping a specific error condition.

9.6 Level 2 trigger

The L1 trigger operates about 2.4 μs after an interaction has taken place. However, during the drift time of the TPC ($\sim 100\ \mu$s), there is a considerable length of time available for further manipulation of the data extracted from the trigger sub-detectors (including those which have not taken part in the earlier trigger levels, e.g. PHOS). During this time resets may also arise as a result of the of the past–future protection.

Besides rejecting pile-up events, L2 will enrich the subsample of events with full readout with dimuon events in the mass region of the J/ψ. This is currently the only task foreseen. However, as time evolves, other trigger algorithms could be evaluated for application in ALICE. For example, we could consider the use of a moments analysis algorithm to select events with unusual multiplicity structures [9].

The more complex data manipulations envisaged for the L2 trigger imply that the trigger decisions can no longer be synchronous. Instead a pool of parallel trigger processors will be prepared, one for each trigger algorithm. An important restriction is that the processors will not have access to the full data transmitted to the DAQ. Instead, participating sub-detectors must prepare special trigger data streams for use in the L2 trigger processors. The trigger processors may reside near the central trigger electronics or, if appropriate, near the local trigger electronics. For example, it may be more efficient to put the first and second level dimuon trigger electronics together, since they access the same data.

9.7 Overall trigger layout

The overall layout of the trigger is shown in Fig. 9.8, which indicates where the data available to the trigger system come from at each level. In order to satisfy the constraints put on the maximum trigger rates at L1 and L2 (not more than 10^3 Hz for the TPC strobe, not more than 80 Hz for readout with TPC) programmable scaling factors will be implemented at both L1 and L2 for each trigger class. The allocation of scaling factors can thus be used to define priority streams. For example, trigger classes involving dimuons, which are rare in all modes of operation, and especially in pp, would not be downscaled. In addition, such rare events can be recovered from those events 'spoiled' by pile-up, as described below.

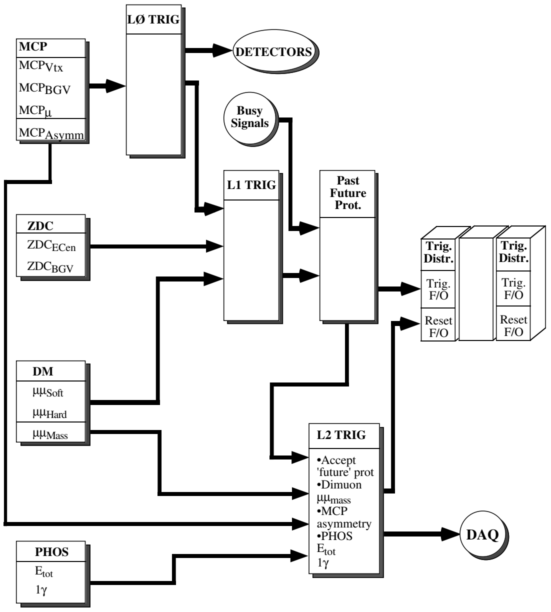

Figure 9.8: Overall Trigger Logic Diagram.

Table 9.3: Expected rates (in Hz), dead time and TPC pile-up for different modes of operation.

	Pb–Pb	Ca–Ca		pp
$\mathcal{L}$ (cm^{-2}s^{-1})	10^{27}	2.7×10^{27}	10^{29}	10^{30}
Int. Rate	8000	8000	3×10^5	10^5
L0 Rate	5000	5000	3×10^4	5×10^4
L1 Rate	1000	1000	~ 100	1000
L2 Rate				
TPC $\mu\mu$ central	40	~ 3	–	–
TPC $\mu\mu$ MB	40	38	–	–
TPC no $\mu\mu$	–	37	–	500
no TPC $\mu\mu$	120	~ 4	30	< 0.1
high p_t $\mu\mu$	0.18	$\sim 10^{-3}$	0.1	$\sim 10^{-3}$
Dead time (%)	11	7	10	15
TPC pile-up (%)	63	63	–	–

Pile-up is a major consideration for this experiment. In order to keep track of pile-up, L0 and L1 decisions are first made without dead time, crossing-by-crossing, and ignoring pile-up. The results are fed into a past–future protection unit, which also monitors dead time from the DAQ system. This unit keeps track of pile-up for each sub-detector in ALICE, and determines whether it is possible to take an event or not. Initially, readout is attempted using the largest system of sub-detectors available. For example, if the trigger system detects a rare high p_t dimuon event, it will attempt to read out the full detector; if this is not possible because of pile-up in the TPC, then readout using a reduced set of sub-detectors, which would include the dimuon arm and the pixel layers of the ITS, will be attempted. The event is rejected only if even the minimum set of sub-detectors is not available, for example because of DAQ-imposed dead time.

The past–future protection unit can also modify event classes. For example, if readout is attempted for a 'rare' event with all detectors and a second event is detected before the end of the TPC drift time, the event is reclassified and read out with the reduced sub-detector set.

The output from the past–future protection unit is a trigger level and a trigger word. This identifies the event class and in principle determines what each sub-detector should do. This information is fed to a series of programmable trigger distribution modules, one for each sub-detector, which produces appropriate signals for the sub-detector given the trigger class. These modules would be placed in a crate close to the central trigger crate. The RD12 TTC distribution system [10] is an attractive choice for performing the communication between this crate and the sub-detector front-end systems. Although not all the features of the system are exploited in our experiment, much attention has been given to reliability, and its use also in other LHC detectors mean's that it would be widely supported at CERN.

9.8 Rates

The rates expected in the different modes of operation (Pb–Pb, Ca–Ca and pp) for ALICE represent a balance between two rather different modes of operation. One aim of the system is to select central events corresponding to the most central $\sim 10\,\%$ of the cross-section. This is not a large reduction from the minimum bias rate, and is not very demanding to achieve. It is more difficult to balance this with the requirements of the dimuon system, as this trigger selects a far smaller fraction of the total cross-section, particularly in pp mode.

In the following, the expected rates for Pb–Pb, Ca–Ca and pp interactions will be discussed in detail; the results are summarized in Table 9.3. The dimuon rates will be estimated in more detail in the dimuon addendum to the technical proposal and the values given here are preliminary.

In contrast, the rates listed in Table 10.2 of Chapter 10 are the *maximum* rates which can be handled by the data acquisition system. The DAQ can accommodate a higher rate of 'fast-readout' triggers (i.e. no TPC) than is listed here, up to a maximum readout rate of 1 kHz.

The pile-up probability estimates assume that interactions producing less than 1 % additional tracks in the central acceptance region are tolerable.

Pb–Pb rates

A luminosity $\mathcal{L} = 10^{27}$ cm^{-2} s^{-1} gives rise to a raw event rate of 8 kHz. The L0 trigger does not make stringent requirements on centrality in order to allow as large a fraction of events as possible to be available to the dimuon trigger. It selects events with at least ~ 100 tracks in the central acceptance.

At L1 the information from the dimuon system is available, and the ZDC data allow a more complete centrality decision to be made. The centrality cut gives a rate of 1000 Hz, which far exceeds the capacity of the data acquisition system. This can be reduced either by making further selections (i.e. with the dimuon trigger) or simply by the application of a scaling factor.

Turning to the dimuon system, the L1 'low-p_t' dimuon trigger ($p_t \gtrsim 1$ GeV/c) gives a rate of 240 Hz, of which 120 Hz come from central triggers and 120 Hz from non-central ones. The L1 'high-p_t' dimuon trigger ($p_t \gtrsim 3$ GeV/c) gives a rate of 0.18 Hz. It will be classified as a high priority trigger, but has a negligible effect on the overall rates.

The steep fall of the dimuon mass spectrum with increasing mass means that most of the events selected by the 'low-p_t' trigger are below the J/ψ region. The L2 dimuon trigger enriches J/ψ by imposing a mass cut at around 2.5 GeV/c^2. The dimuon L2 rate could be as low as ~ 30 Hz (central plus non-central) of which ~ 10 Hz would satisfy the TPC past–future protection requirement. In this case we could select further central events with an associated L1 dimuon flag to make up a total central rate with the TPC of 40 Hz.

In addition a sample of non-central events will be read out with the TPC at a rate of 40 Hz. This combination of central and peripheral events saturates the allocated bandwidth for readout with the TPC. The remaining dimuon triggers would be read out using the 'fast readout' system. The overall dead time (including DAQ) with these rates is about 7 %.

pp rates

The raw interaction rate in pp interactions will be about 10^5 Hz at a luminosity of $\mathcal{L} = 10^{30}$ cm^{-2}s^{-1}. The luminosity in our interaction region is determined by the luminosity requirements of the pp experiments (see Chapter 8).

In order to produce data for comparison with central Pb–Pb events no special selections are required for the barrel detectors. As we need of the order of 500 Hz events to collect sufficient statistics, this interaction rate is about an order of magnitude larger than what is needed for a large part of the physics programme.

The dimuon requirements are different. The very low rates for dimuon triggers mean that the DAQ overheads are negligible, and the methods described in Chapter 10 ensure that these can be treated as 'priority triggers'. However, only events which give L0 triggers can be used for dimuon analysis, and this limits the effective luminosity, which is given by the L0 rate.

The optimum solution, given the constraints on the location of the trigger electronics —L0 near the detector, L1 near the DAQ— would be to strobe the dimuon arm and the pixel system at the highest possible rate allowed by the front-end electronics and use the 'fast readout' system to collect dimuon triggers. The maximum trigger rate with the current design of pixel detectors is around 5×10^4 Hz. Other detectors not involved in the dimuon stream will be triggered with an L0 rate of $< 10^4$ Hz in order to avoid excessive gating of the TPC.

A possible way to access the full pp luminosity would be to include the dimuon information in the L0 trigger. This is excluded with the current proposed locations of the trigger electronics. Other possible locations of the L0 electronics and/or more optimal cabling of the links from the dimuon system to the L0 trigger are being investigated. We will also investigate whether a higher L0 rate could be made manageable in the pixel detectors.

Ca–Ca rates

Two operating luminosites are being considered for Ca–Ca interactions: $\mathcal{L} = 2.7 \times 10^{27}$ cm^{-2}s^{-1}, to match the Pb–Pb centrality rates, and $\mathcal{L} = 10^{29}$ cm^{-2}s^{-1} to maximize the useful dimuon rates. The considerations for low luminosity Ca–Ca running are similar to those for Pb–Pb running, except that the dimuon rates will be two orders of magnitude smaller and there is therefore no need for a second level dimuon trigger.

It will not be possible to use the TPC in high luminosity Ca–Ca running, as all events will have pile-up. Therefore, the main purpose of this mode would be to record dimuon events with high statistics. All events would be read out with the 'reduced set' of sub-detectors. Here the raw interaction rate is 3×10^5 Hz, so a standard centrality cut would give an L0 rate of 3×10^4 Hz, close to the maximum pixel strobe rate. This means that all central events are available for dimuon analysis, and a final event selection rate of 30 Hz (mainly low p_t $\mu\mu$ triggers) would be achieved. This rate could be doubled if the dimuon p_t triggers were available at L0 (see above), as in this case the non-central events would also be accessible.

9.9 Future options

The trigger system described above can cover the requirements of the baseline ALICE detector. An interesting further option involving existing detectors would be to have a fast trigger output from PHOS, with a decision time of at most a few hundred nanoseconds, which could participate in the L1 trigger. Some further work is required for the PHOS front-end system to arrange this. This trigger would be suitable for selecting high p_t photons. The photon multiplicity in pp interactions is low and such a trigger would greatly enhance the sample of high p_t photons available for analysis.

Extra inputs are foreseen in the L0 and L1 trigger units to allow new trigger detectors and L2 processors to be added easily. A possibility being investigated would be to make the L0 and L1 trigger processors identical. This would require them to be able to both generate and receive trigger words as part of their function. Such a solution would allow the L1 trigger to be scalable, and allow for flexibility in the provision of spare modules. There could be some increase in the latency overheads which would require further investigation, but this could be small compared to the overall latency. Even without this feature, it should be possible to accommodate all the additional trigger detectors currently foreseen.

Among the proposed detectors which are not part of the baseline design, the Scintillator Barrel (SCB, see Section 7.3) is of particular interest as it would give the opportunity to implement a cosmic ray trigger. This would be of great help during the ALICE installation phase, and allow regular monitoring of efficiencies during shutdown periods when the detector is accessible.

10 Data Acquisition, Control and Computing

10.1 Requirements

The ALICE data acquisition and reduction system is designed to be flexible enough to accommodate the different run periods of the experiment and the various types of physics and triggers investigated.

The heavy-ion run will last for a few weeks per year only and will require the largest possible bandwidth to permanent storage. The rest of the running period will be used to acquire data produced by pp interactions which produce five times less data. During pp running, a small fraction of the computing infrastructure will be used for data acquisition and the majority of the computing power will perform data analysis on the data acquired during the heavy-ion run.

The relatively short heavy-ion run, of the order of 10 % of a year determines two of the main features of our DAQ:

- In order to collect a sufficient number of events for physics analysis, the DAQ system has to be designed with a very large bandwidth. Heavy-ion experiments at RHIC, which runs dedicated with nuclear beams throughout most of the year, can afford an order of magnitude lower instantaneous rates and still accumulate similar statistics.
- As we have roughly ten times more time to analyse the data than to collect them, we can try to economize on computing power by performing CPU-intense operations (event filtering, data reconstruction) offline rather than in real time. However, data storage and data analysis costs have to be globally optimized.

During the heavy-ion run, the DAQ system acquires two different types of events. The first one consists of the huge volume of data produced in Pb–Pb interactions when all detectors are read out. In this mode, the experiment will produce a low rate of very large events (see Table 10.1). It combines 40 Hz of central events (between 33 and 39 MByte) and 40 Hz of minimum-bias events with an average size of 10 MByte. The combined data traffic corresponds to 50 Hz of central events of up to 40 MByte and this estimation is often refered to in the rest of this chapter. This data volume has to be reduced before mass storage by at least a factor 2.

The second one concerns the physics of high-mass muon pairs, when only the muon arm, the pixel planes (for vertex finding), and the trigger detectors are read out. While for these triggers, which are interposed with the previous ones, the rate is much higher (up to 1 kHz), the data throughput is nevertheless only 20 % of the total.

It is essential that the trigger and DAQ systems be capable of handling both types of traffic at the required rate. Moreover, several detectors will be read out by both types of triggers, but the two event classes should be stored on separate permanent storage volumes. The dead-time generated by the DAQ system (transfer of data from the front-end electronics to the first DAQ memory) should not exceed 10 % of the total time.

For data taking, groups of detectors must therefore be read out independently. This is also the case during the installation and the commissioning periods of the experiment.

Table 10.2 summarizes the different modes of running with the corresponding data traffics. In summary, the DAQ system must be capable of multiplexing several data sources to several data destinations at an aggregate bandwidth of up to 2.5 GByte/s and reduce it to 1.25 GByte/s before storing it on mass-storage. Most of the resources must, in addition, be easily reconfigurable from data-taking mode to data analysis mode.

10.2 DAQ architecture

10.2.1 General architecture

The functions of the DAQ and control system are:

- to implement a readout selection according to the trigger type;
- to transport the data from the front-end electronics to the computing room;
- to perform the zero suppression (unless done in the front-end electronics) and assemble the data in sub-event data blocks;
- to assemble the sub-event data blocks into full events;
- to provide enough computing power to perform the adequate level of preprocessing in order to reduce the data volume;
- to record the data onto permanent storage;
- to monitor the quality of the data and detect malfunctions as quickly as possible;
- to manage smoothly several types of data traffic either in parallel (central and dimuon events) or consecutively (heavy-ions and pp run);
- to have a control and presentation system for the experiment in a unified and manageable human interface.

Table 10.1: Event size of the sub-detectors for central Pb–Pb interactions. The trigger system, while only having few channels, has a considerable size (115 kByte) because it accumulates some information during the time preceding (and following) the actual event.

Detector	Subdector or option	Number of channels	Event size (MByte)	Max. readout time (μs)
ITS	Si pixel	14.4×10^6	0.140	200
	Si drift	192×10^3	1.500	≤ 1000
	Si strips	2.6×10^6	0.160	≤ 1000
TPC		500×10^3	30.0–36.0	2000
PID	TOF Pestov	380×10^3	0.440	≤ 1000
	TOF PPC	160×10^3	0.150	≤ 1000
HMPID	RICH	145×10^3	0.120	≤ 1000
	Scint. TOF	8×10^3	0.005	≤ 1000
PHOS		37×10^3	0.022	≤ 1000
Muon		1.0×10^6	0.100	200
Trigger		700	0.115	200
Total			$\sim$ 33–39	

Table 10.2: Event readout time, rate and size for the different types of beams and trigger.

Beam type	Trigger type	Detectors	Readout time (μs)	Maximum event rate in DAQ (Hz)	Event size (MByte)
Pb–Pb	Central	All	2000	40	33–39
	Min. bias	All	2000	40	0–39 (Average 10)
	Dimuon	Muon, pixel, trigger	200	1000	0.250
Ca–Ca	Central	All	1000	150	5–6
	Min. bias	All	2000	150	0–6 (Average 1.5)
	Dimuon	Muon, pixel, trigger	200	1000	0.100
pp	Min. bias	All	200	1000	0.500

The DAQ architecture is based on a network of high-speed links linking all the data sources and the data destinations through a switch, as shown in Fig. 10.1. This architecture provides the required flexibility and scalability to run in very different modes and seems the most promising solution today. Several technologies (ATM, fibre channel, HiPPI, SCI) are or will soon be available to implement a switching network to connect data sources and data destinations with an aggregate bandwidth of up to 2.5 GByte/s. The respective advantages and disadvantages of each technology are detailed later. However, this comparison of the relative merits of each of them should not be over-emphasized. The acceptance of a given technology by the telecommunication and computing industries is also (if not more) important. The development of the technology will allow us to buy most of the components of this network. The future work will consist in designing the overall system after extensive simulation and prototyping. We will then be able to choose the technology suitable to providing a solution at the best price and with the required industrial support and adequate lifetime.

In the case of ALICE, the huge requirements for permanent storage may require a complex link between the

Table 10.3: ALICE DAQ parameters.

No. Detector Data Links	$\sim$ 700 at 100 MByte/s
No. Read-Out Receive Cards	$\sim$ 175
No. Front-End Digital Crates	32
Event Builder Switch	32×32
Event Builder bandwidth	2.50 GByte/s
Data Storage Switch bandwidth	1.25 GByte/s
Data Storage bandwidth	1.25 GByte/s

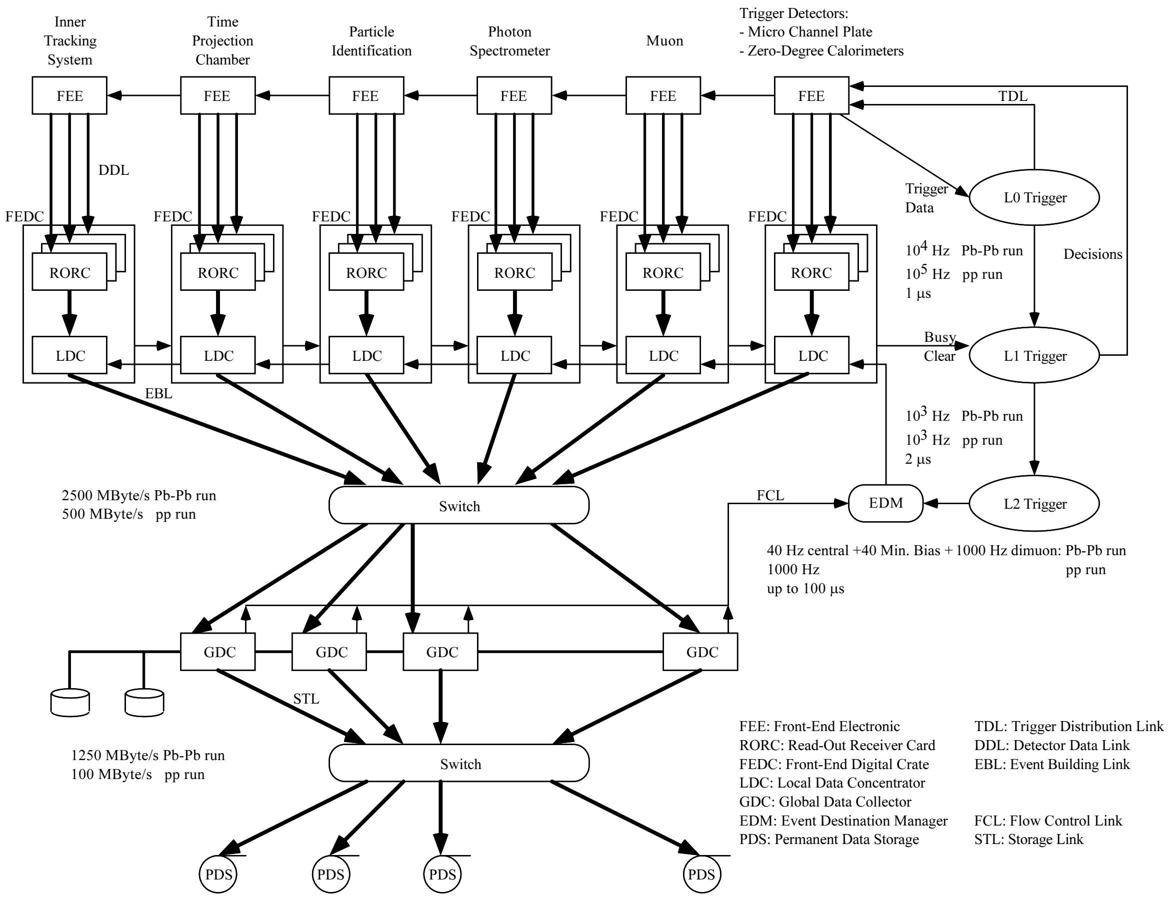

Figure 10.1: Overall DAQ architecture.

computer and the mass storage for reasons of reliability and flexibility. This may be obtained by loops connecting groups of computers and tapes, or by a switch as indicated in Fig. 10.1. If the same technology is chosen both switches could physically be merged into one device.

10.2.2 DAQ components

Table 10.3 summarizes the parameters of the ALICE DAQ system; its individual components are described in the following.

Front-end electronics

The functions of the front-end electronics include amplification, shaping, buffering, multiplexing, sparse readout, digitization, pedestal subtraction, data compression, and data transport. This processing is implemented in the analog and/or the digital electronic chain. These front-end electronics functions and their possible implementation are specific for different sub-detectors and have been described in the corresponding previous chapters. However, the interface of these electronics sub-systems to the trigger and the DAQ systems is similar and is based on the same functional model.

Front-End Digital Crates

The architecture of the Front-End Digital Crates (FEDC) system is shown in Fig. 10.2. The data collected by the front-end electronics will be transferred by the Detector Data Links (DDL) to ReadOut Receiver Cards (RORCs) located in the FEDC. The collaboration will try to standardize on a few types of transfer links according to the needs of the different sub-detectors, as described in Refs. 1 and 2. Current estimates on a possible distribution of the electronics of the sub-detectors show that there is a need for a digital link able to support around 100 MByte/s of data. Preprocessing of the data (zero suppression, address insertion, etc.) may be done in the front-end. Event fragments from a number of RORCs are transferred to the Local Data Concentrator (LDC) where they are merged into sub-events. The data processing applied depends very much on the sub-detector. Most of the data reduction will be done in the front-end electronics or in the front-end digital crates to save on the required bandwidth of the

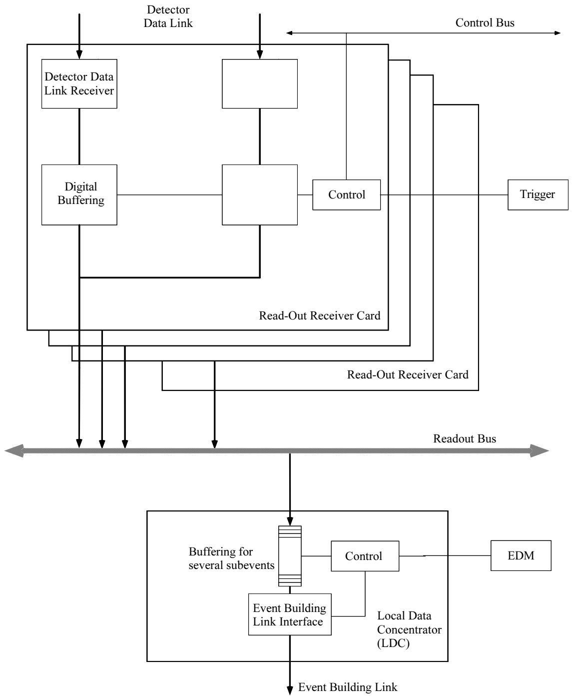

Figure 10.2: Front-End Digital Crates architecture.

event-building network. However, the data reduction applied in the FEDC or before in the acquisition chain can only be done on sub-events. The data processing which requires a full event, such as TPC pre-tracking or, to a lesser extent, cluster finding, will be done in the Global Data Collector (GDC), after the event-building.

Local Data Concentrators

The Local Data Concentrator (LDC) has several functions in the overall system:

- buffering of the sub-events until they have been transferred to the GDC;
- collection of data from the different RORCs in the FEDC and sub-event building;
- data formatting;
- participation in the event building and distribution protocol.

The LDCs receive regular messages from the Event Destination Manager (EDM) with the trigger information and the event building information. The latter contains the GDC which will be used to build the event referenced in the trigger information. The role of the LDC is to collect the information for the event, format it and send it to the destination GDC when required by the protocol. A central part of the LDC will be a dual-port type memory with facilities for fast memory management, capable of operating at 100 MByte/s sustained throughput.

Event building links and switch

A switching network interconnects the LDCs and the GDCs. This network can be implemented using ATM [3], Fibre Channel [4], HiPPI [5] or SCI [6]. This network may also constitute the physical layer to run the event building and distribution protocol (see Section 10.2.3). As mentioned before, the ALICE total bandwidth is or will soon be available from various technologies. Two arguments have to be considered for the choice of technology: whether it is adapted to the ALICE requirements, and if the industrial support is such that most of the components can be bought off the shelf. The different technologies mentioned above tend to be used where they are best-suited. The ATM is clearly ideal for mixed traffic and bandwidth reservations. Its applications have a large scope from WAN and telecommunication to LAN. The Fibre Channel seems to find its place for I/O device connections with the Arbitrated Loop (AL). Several of the largest suppliers of magnetic disks offer Fibre Channel as a possible disk attachment [7, 8]. Two other technologies are targeted at the same problem: Serial Bus [9] and SSA [10]. Finally, SCI may be used in large memory-mapped systems and multi-processor machines.

Global Data Collectors

The Global Data Collectors (GDCs) constitute the destination of the event building network. The GDC receives the sub-events from several LDCs and builds a full event from these fragments. The role of the GDC is to check the integrity of the data, the presence of the sub-events needed for this type of event, the data formatting, and the recording onto permanent storage. The continuous progress in processing power and memory size will provide extremely powerful computers. The exact computer architecture (farm of workstations or symmetric multi-processor machines) will be chosen at the last possible moment to profit from industrial progress. Commodity computer products will most likely continue to increase performance at constant cost. We will follow closely this evolution and see how it influences the best computing architecture. Owing to the huge bandwidth needed for the Permanent Data Storage (PDS), the events will be preprocessed in the GDC. Currently our target figure for data reduction at this stage is a factor of at least two (from 2.5 to 1.25 GByte/s for Pb–Pb events), mainly to be achieved by cluster finding or pre-tracking in the TPC (see Section 10.6.1). This requires great confidence in the event reconstruction software as some raw data has to be discarded at this stage. It is anticipated that all the raw data will be recorded during the first running periods. While gaining experience and confidence in the software it should become possible to record only the reconstructed data. The need for this online reconstruction will also be influenced by the evolution of permanent data storage technologies.

Table 10.4: Storage technologies summary.

Technology	Bandwidth MByte/s	Device cost kCHF	MByte/s cost kCHF	Capacity (uncomp.) GByte/cart.	Cartridge cost CHF/cart.	Media cost CHF/GByte
Low-end						
DAT	0.5	2.0	4.0	8	30	3.8
DEC DLT 4000	1.5	9.0	6.0	20	140	7.0
Exabyte 8500	0.5	3.4	6.8	5	11	2.2
Midrange						
IBM MagStar	9	50	5.6	10	60	6.0
STK Redwood	11	150	13.6	50	120	2.5
High-end						
SONY	32	300	9.4	100	250	2.5
AMPEX	15	230	15.3	165	840	5.0
CREO	3			1000		
Interferometrics	256	420	1.6	800	2880	3.6

Event Destination Manager

The Event Destination Manager (EDM) is responsible for receiving the request for events to be built from the trigger, and for each event decides in which GDC the event will be built. This mode of operation will allow one to handle the different traffic types in a common way. At any time during the acquisition of data, it will be possible to specify to the EDM which GDC can be used for each trigger type. Together with the availability status received regularly from the GDCs, the EDM will decide for each trigger which GDC is to be used according to the actual status. This scheme allows for easy and fast system reconfiguration when a GDC or a PDS is broken, taken out from, or added to the system. The exact definition of this protocol and the physical layer, Flow Control Link (FCL), on which it will run is not yet defined and is the subject of future investigations.

Permanent Data Storage

The ALICE raw data Permanent Data Storage (PDS) must combine very high bandwidth per device to limit the number of devices, and high capacity per volume to keep tape handling and storage of the $\sim$ 1000 TByte written during one heavy-ion run at a reasonable level.

The combination of these two characteristics restricts our choice for the time being to a limited set of technologies and devices that are reviewed hereafter.

The PDS is therefore the biggest technological challenge that must be tackled by the ALICE DAQ system. We have, however, assumed in the following that the PDS will be acquired from industry and requires little or no hardware work to be integrated into our system.

We expect that the data will be processed and reduced at least once before distribution to the collaboration. Processed data can then be made accessible via either network or removable storage media (‘shipping’). In the latter case, different media could be selected for raw data storage and for distribution allowing an independent optimization of technical performance and ‘exportability’.

A radically different solution for data storage could emerge from new techniques such as holographic storage [11], molecular-scale storage [12] or usage of the Kerr effect to read magnetic tapes optically [13]. However, these techniques have not yet produced any commercial product and cannot be considered as alternatives at this stage.

Finally we want to mention the High Performance Storage System (HPSS), a large consortium lead by the National Storage Lab, in the United States [14]. HPSS is planning to provide a scalable storage architecture supporting by the end of the century the following features: GByte/s data transfer, TByte data files, and PByte total data storage. These features match the ALICE requirements well and we shall follow the progress of this effort closely.

The current data storage market can be divided into three categories [15]. The characteristics of some typical devices of each of these categories are shown in Table 10.4.

- Low-end systems are commodity products driven by the PC and workstation markets. The key requirements are standard form factors, multiple sources, and low entry cost. This market is progressing the fastest in terms both of volume and of achieved performances.

- Midrange systems are driven by computer centre needs. The key requirements are data integrity, reliability under high load, and backward compatibility. The performances of these systems are changing very slowly.

- High-end systems are dedicated to the recording of scientific data [16–18] or special industrial applications such as oil prospecting. This market is not well served by the computer industry. Several alternatives have been developed such as derivations of professional video devices.

The evolution of these different categories is difficult to predict. However, several different sources indicate the following tendencies. Low-end systems evolve quickly. The predicted grow is of a factor four every six years [19]. Midrange systems grow much more slowly. Recent experience indicates a factor three in bandwidth over the last 10 years [20]. High-end systems are driven by few applications and their evolution cannot be predicted [21]. Today's market size is too small to justify new investments from industrial companies. The future of high-end systems is also strongly linked to the evolution of HDTV. The predictions for the HDTV technology forecast promising performances (up to 150 MByte/s for the D6 standard [20]), but the future of the HDTV standard is uncertain.

In summary, the permanent storage for ALICE could be implemented today with a reasonable number of high-end systems, but at a high cost. The midrange systems would require too many devices to achieve the required bandwidth. By the time a decision has to be taken, midrange systems will probably become a valid alternative, and, given their much faster evolution, even low-end systems could become viable.

10.2.3 DAQ software

Event building and distribution protocol

For each event selected, the trigger sends a message to the EDM. The EDM selects the GDC which will build this event and broadcasts the trigger information and the destination address to all the LDCs. The LDC will perform the readout of the data in the same front-end crate, format these data and send them to the destination GDC. The complete event will be transferred fragment by fragment to the GDC. This will create some contention between all the LDCs trying to access the same GDC destination. The event-building and distribution protocol will be in charge of optimizing the traffic and minimizing these contentions.

Control and online monitoring

The DAQ control system will control and monitor the correct functioning of the detectors, of the front-end, and of the DAQ system itself. It will also monitor the quality of the data. The DAQ control system should be able to exchange information with the experiment control system. Integration of DAQ control and overall experiment control seems attractive and will be investigated.

Data recording and data access

The ALICE raw event data will most probably be recorded and processed in a sequential manner. However, several applications will require databases such as event cataloguing, detector and calibration data, etc. The ALICE Collaboration is participating in the RD45 research project at CERN ('A Persistent Object Manager for HEP') [22].

As object-oriented software technology evolves, it is very likely that the analysis procedures and packages used in the LHC era – and probably well before that – will use object-oriented languages like `C++`. However, the *persistency* of objects (storing them on a medium and later accessing them) is not yet properly addressed by current object-oriented languages. The RD45 project evaluates whether existing commercial object database packages can be used by the LHC experiments to store 'objects' in an object-oriented programming environment.

The ALICE Collaboration will participate in this effort and will start to use the object-oriented database technology in pilot projects.

10.2.4 DAQ prototyping

The ALICE DAQ system must handle a very high bandwidth and must also be easy to reconfigure for different types of running. A prototyping and evaluation programme has been started to ensure that these goals can be reached, and that the most cost-effective system for the lifetime of the experiment is attained. The programme will encompass items listed below.

- Testing and evaluation of readout modules, transmission links, switches, and processors.
- Verification that detector-specific modules conform to DAQ specifications.
- Test benches for evaluating technology-specific DAQ components and system architecture.
- Development of test-specific hardware.
- Comparison of system performance with simulation results.
- Trigger, readout control, and synchronization protocols.
- Sub-event and global event building.
- Development and testing of readout and monitoring software.
- Online data reduction/processing.
- Error handling.

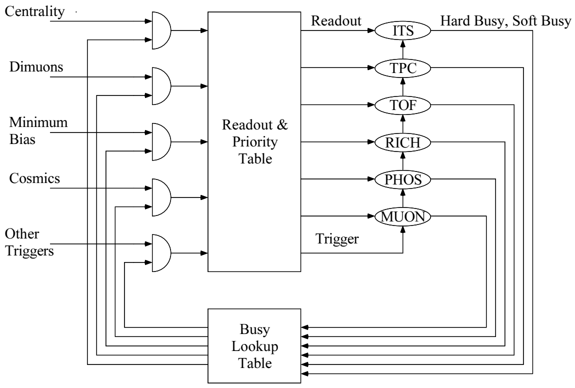

Figure 10.3: DAQ trigger interface.

- General testing and verification strategies for the integration of the ALICE DAQ system.
- Detector control and monitoring.

10.3 Trigger – DAQ interface

The rates of various triggers in different running conditions are discussed in detail in Chapter 9.

10.3.1 Parallel triggers

The trigger system may provide a number of different triggers to the DAQ system and the DAQ provides a differentiated busy status to the trigger system. Different triggers may require different detector readouts, e.g. a dimuon trigger may request only the dimuon apparatus and the silicon pixels, whereas a central trigger requests the readout of all detectors. A trigger is transmitted to the DAQ only if all requested detectors are free for the next event. The check can be done with minimal delay in a memory look-up unit as sketched in Fig. 10.3.

A trigger signal is fanned out to all sub-detector units, even though they may not participate in the readout, in order to allow them to account for the trigger number and event serialization. The TTC system as described in Ref. 23 is considered a likely solution for the trigger distribution.

A separate signal causes the actual readout of the detector, i.e. the transfer of data from the detector to the RORC in the FEDC crate over the DDL. This readout has to be performed fast enough to avoid excessive dead-time.

The role of the LDC is to empty the RORC buffers, to assemble the data blocks into sub-events and send them to the GDC used for the event building.

In case the LDC buffer is not emptied, the corresponding 'busy' signal will not be cleared and no more trigger involving this detector will be possible. In this scheme it is actually possible that different sub-detectors are simultaneously performing the readout for different triggers.

10.3.2 Trigger levels

The system allows different readout times for different triggers depending upon the amount of information to be transferred as indicated in Table 10.2. This maximum readout time is limited to reduce the DAQ generated dead-time to the order of 10 %.

The readout is started by the trigger L1. The trigger L1 is then followed by a trigger L2 or by an abort. In the latter case, data may already have been transferred into the RORC, which has to be able to discard data belonging to aborted events.

Decisions requiring more time or access to more data than available to the trigger L2 will be made by a filtering software running on full events in the GDC.

10.3.3 Trigger priorities

The trigger and DAQ system allows one to set different priorities for different trigger types. The dead-time is reduced to a minimum for triggers which are considered as privileged while allowing other triggers with high rate or large data volume to utilize the total bandwidth of the DAQ system.

Each readout system provides two 'busy' signals. The 'hard busy' is raised by the trigger itself during the readout and the data transfer from the detector to the FEDC. A completely full FIFO will suspend the current transfer and hence maintain this signal. No trigger (even a privileged one) requiring any particular sub-detector in this state can be accepted. The 'busy' signal is cleared by the RORCs at the end of the data transfer.

Completely full FIFOs should occur rarely, as the high volume/rate triggers will already respect the 'soft busy' signal which is asserted when the RORC FIFO gets half full, hence leaving enough space for a number of privileged triggers.

Therefore, the memory look-up unit, shown in Fig. 10.3, has two inputs for every sub-detector and decides whether a given trigger has to respect only the hard or also the soft busy of a required sub-detector.

10.3.4 Trigger – Front-End – DAQ interface

The interface between the trigger and the DAQ consists of a few signals [24], which can be sent separately to each sub-detector.

- Trigger to DAQ:
  1. Trigger: broadcasts next trigger number (common to all sub-detectors);
  2. Readout: causes data readout;
  3. Abort: higher level trigger which stops digitization and readout.
- DAQ to trigger:

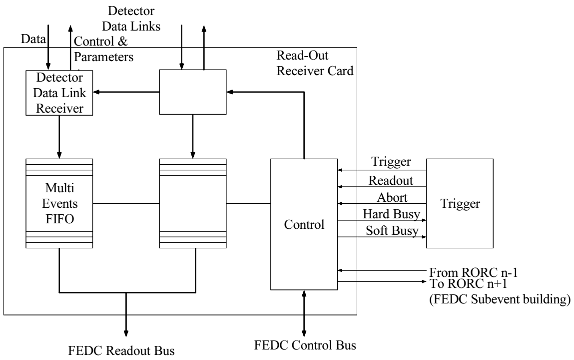

Figure 10.4: Detector Data Link interface to RORC.

1. Soft Busy: indicates that only a high-priority event can be read out;
2. Hard Busy: indicates ongoing digitization or data transfer.

10.3.5 Software

As shown in Fig. 10.4, all readout is hardware-driven up to the link receiver cards in the LDC. A program in the LDC processor needs only to empty the RORC FIFOs regularly, but asynchronously with respect to the triggers. It serializes the data coming from all its input links into one sub-event.

Based on the decision of a central EDM or following alternative distributed schemes [25], it transfers each sub-event to a GDC across the link. Reference 26 describes these mechanisms in detail.

We will study scenarios where special, privileged, low-rate triggers (e.g. dimuon triggers) are sent to a separate PDS output stream for easy and fast access.

10.4 DAQ simulation

10.4.1 Introduction

The ALICE DAQ architecture is modelled and simulated in order to provide information on the behaviour of the system as well as requirements for topology, technology, link bandwidth, switch latency, etc. [27].

The model has been used for the preparation of this Technical Proposal. However, the very long time-scale of the preparation, the construction, and the lifetime of ALICE detector requires having a tool to evaluate the DAQ architecture and the technologies foreseen to implement it. This simulation environment will therefore be further developed during the next phases of the experiment's preparation.

10.4.2 The ALSIM program

A standard ALICE DAQ SIMulation environment (ALSIM) [28, 29] has been developed. ALSIM is written in the MODSIM II [30] language and follows an object-oriented approach which provides a flexible and open simulation environment. Models of components can easily be added and various configurations can be chosen dynamically by means of description files. An interactive tool has been developed to generate the description files. A graphics version of ALSIM is also under development.

Generic models of DAQ components have been developed such as data generators, memories, links, and switches. Each component has parameters which can be configured by the user, such as the memory size, the link latency and bandwidth, and the switch latency. The description file describes the system topologies and the details in each DAQ component.

The simulation provides results on the event-building rate, the buffer occupancy, the event-building latency, the link latency, etc.

10.4.3 Results

Even though several technologies have already been built into the ALSIM code (ATM, Fibre Channel and SCI) and much simulation work has been carried out [31], it is too early at this stage to give results for comparison of the different technologies. We therefore want to be as general as possible when simulating the ALICE DAQ system and the event-building algorithms. The generic model in the ALSIM code is designed for such purposes.

The model that has been extensively simulated corresponds to a simplified version of the global architecture as seen in Fig. 10.1 on page 141. The parameters which were used are summarized in Table 10.5.

The aim of the simulations was to see how the system performance depends on different values of some basic input parameters.

We have examined the behaviour of a system with 32 LDCs and 32 GDCs connected to the switch. The L1 trigger rate was set to 300 Hz and followed an exponential distribution. The average event size varied between 30 and 58 MByte for central triggers and between 150 and 300 kByte for the dimuon triggers. The event size followed a Gaussian distribution, but only an interval around the mean value was accepted to simulate realistic heavy-ion events. This interval was defined by a range factor. The primary goal of the simulation has been to investigate the performance of the switch-based event-building system. Therefore, the throughput of the DDLs was artificially limited to reach the requirement of 50 central events per second. The actual throughput of the system was then determined by the throughput and the communication latency of the event-building links between the LDCs and the GDCs. For the event-building links, the throughput and the communication latency were set to realistic val-

Table 10.5: Simulation parameters

Parameter		Values
Central event size	(MByte)	30, 35, 42*, 45, 48, 51, 58
Dimuon event size	(kByte)	150, 300
Event size range factor		1.00, 1.05, 1.10*, 1.15, 1.20, 2.00
Sub-event size deviation	(% of mean value)	0, 5*, 10, 15, 20, 25, 30, 40, 50, 60
LDC buffer size	(MByte)	20, 24, 32, 40, 42, 48, 56, 64, 72, 80*, 120
GDC buffer size	(MByte)	128
Event-building link throughput	(MByte/s)	50, 60, 70, 80*, 90, 100
Event-building link latency	(ms)	0.50, 1.0*, 1.50, 2.5, 3.5, 5.0, 7.0

* Stands for default value.

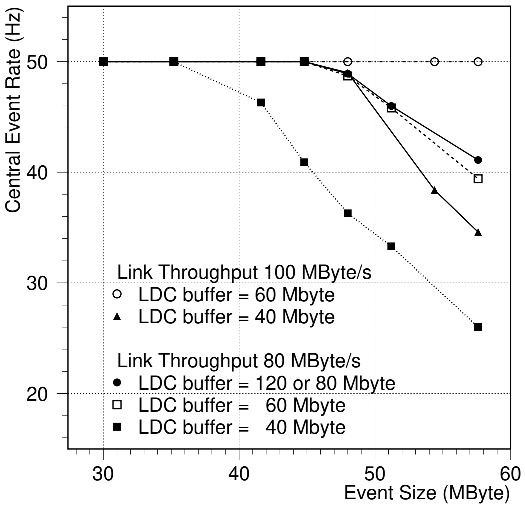

Figure 10.5: Central event rate as a function of the event size at different LDC buffer sizes.

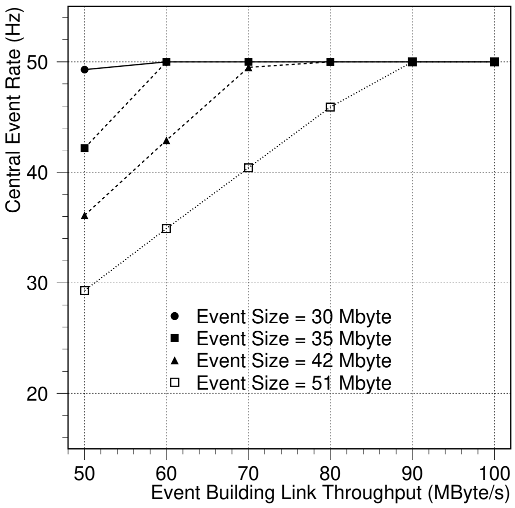

Figure 10.6: Central event rate as a function of the event-building link throughput.

ues compatible with measurements [32] done on a existing switch-based network. These measurements include not only the physical transfer time but also the interrupt latency of an operating system and the networking software.

Figure 10.5 shows how the central event rate depends on the event size for different LDC buffer sizes. At 120, 80 and 60 MByte LDC buffer, an event rate of 50 Hz is attained for event sizes up to around 46 MByte. Above the critical event size, the event rate decreases proportionally to the event size. This shows that with reasonable buffer size, the required performance can be obtained from the system.

Figure 10.6 shows the central event rate as a function of the event-building link throughput for different event sizes. The critical throughput is about 70 MByte/s if the link communication latency is 1 ms. It shows that, for the expected central event size, the required performances can be obtained with a system using 70 % of the theoretical bandwidth of 100 MByte/s.

We have also simulated how the event rate depends upon the event-building link communication latency. We have observed that the rate of central trigger events is stable at 50 Hz, with a communication latency lower than 2.5 ms, and decreases for larger latencies. This relative independence of the ALICE event rate on the switch latency is due to the large size of the ALICE central events.

In Fig. 10.7 we see how the event rate depends on the event-building link latency when we consider dimuon events. The simulations are performed under the theoretical condition that the central event will get high priority of the link and that the dimuon trigger rate is very high in order to show clearly the limiting factors for the dimuon event traffic. We show curves for both 150 kByte and

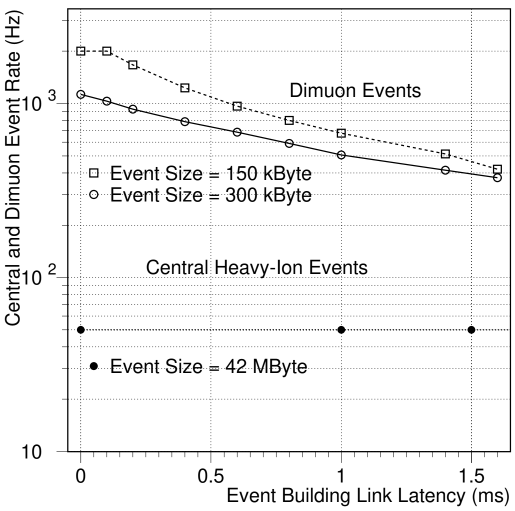

Figure 10.7: Central and dimuon event rate as a function of the event-building link latency.

300 kByte dimuon events. It is notable that the requirement for link latency is proportionally more strict in the dimuon case than in the central event case owing to the higher event rate and the shorter data size.

In summary, the current choice of parameters makes us confident that the architecture proposed can fulfil the requirements of cumulated central and dimuon event traffic with an event-building link giving a throughput of around 80 MByte/s. The dimuon event traffic requires a latency smaller than 1 ms.

10.5 Control

10.5.1 Control architecture

It is foreseen to integrate all the control functions of the ALICE detector in a large system which will control and monitor the behaviour of the experiment and will allow it to run with the minimum of people on shift. The main requirements for this control system are the following:

- Support distributed and heterogeneous environments;
- Collect and present the status of the experiment;
- Store/Retrieve the experiment's set-up configurations;
- Store/Retrieve detector monitoring values;
- Fulfil the safety requirements.

The control and the monitoring of ALICE will be based on a hierarchy of interconnected computers. This hierarchy will consist of several layers: overall system control, subsystems, components, and sub-components. Each subsystem, such as every individual detector or the DAQ system, will be a unit which must be controllable independently. This will be of particular importance during the system installation and commissioning.

The network used for the control will be independent from the data acquisition readout network which will be purely dedicated to the transport of data and of dataflow control protocols.

10.5.2 Detector control

Each detector will need a sub-system to control and monitor the various parameters related to its operation. A database will be used to store and retrieve these parameters and to archive their evolution. The exact requirements will be established by every detector group.

The possibility of using industrial software in this area has been identified and has been the object of the RD38 project [33]. The exact area covered by these products and their integration with customized application will require particular attention.

10.5.3 Run control

The ALICE run control will include the control of the DAQ system, the use of databases, and graphics software to present the huge amount of complex information. These needs are partially covered by the tools needed for the detector control. However, the run control also requires a set of tools to:

- Decompose the experiment into logical system, subsystem, and components and describe their static behaviour. It is expected that this description could be stored in a database.
- Describe the dynamic behaviour of the experiment and operate systems or sub-systems in a flexible and independent way.
- Detect and recover from errors by possibly using software to help in the diagnosis and repair of faults.

10.6 Offline computing

There will be no clear distinction, neither in hardware nor in software, between online and offline computing in the ALICE experiment. Because of the smaller load in proton running, we plan to reuse for offline processing a part of the computing power used online during heavy-ion data taking. Modules of the reconstruction software capable of data reduction could be integrated, after offline testing, into the online processing. It is therefore desirable to use from the beginning a common approach (methods, technology, tools, etc.). The Object-Oriented (OO) methodology will be followed for the software development.

A common detector and calibration database will be used at all stages for online, reconstruction and data analysis, as well as for simulation. We will start to use the OO database technology as soon as possible, during the evaluation phase of the RD45 project [22].

10.6.1 Reconstruction requirements

Traditionally, HEP experiments wrote raw data onto permanent storage and performed the analysis of these data offline. The total data volume usually increases after the first processing steps. The need for ALICE to store raw data is so large that particular precautions should be taken in order to move away from this tradition.

In order to reduce the raw data volume, the best way would be to perform part of the reconstruction online, mainly the TPC data which constitutes more than 90 % of the data (see Table 10.1 on page 140).

A central Pb–Pb event (at $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$) with about 12 000 charged tracks in the acceptance will produce between 30 and 36 MByte of data in the TPC after suppression of zeros and hits outside $\pm 45°$ (assuming for each track hits on all 75 pad rows consisting of 4.5–5 pads and 5.5–6 time slices each, and including two bytes for addressing each time slice).

There are several ways to represent this information in a more compact form. We could either reconstruct the hit parameters (cluster finding) or proceed up to the track parameters (track finding and track fitting). In the first case, by reserving 8 bytes for each hit (including hit coordinates, total amplitude and cluster width information), we estimate the TPC data volume after cluster finding to be about 7 MByte. In the second case, assuming 80 bytes for each track (which includes the track parameters and their covariance matrix), the minimum TPC data volume could be as low as 1 MByte.

At the present stage it is difficult to state which fraction of the reconstruction can be performed online. However, it is obvious from the numbers above that there is ample room for data reduction and that we could benefit from any online (pre-)processing of the TPC data.

We are currently studying different types of algorithms for track finding (see also Section 11.3) which can be divided into two categories. First there is the classical approach that starts with cluster finding and hit reconstruction and then proceeds for actual track finding. Then we have global methods which do not necessarily include the hit reconstruction and, therefore, are not adequate for the first type of reduction mentioned above. However, they could be faster for track finding.

Usually, most of the basic track finding is performed within a fraction of the reconstruction code and processing time. The remaining processing time is spent in several further passes which improve track finding efficiency, double track resolution and precision by the merging of track segments, cluster deconvolution, etc. Such improvements are necessary at least for a part of the physics programme (e.g. HBT correlations) but not necessarily for the whole. Therefore, some of the physics programme which needs large statistics (e.g. di-leptons) could be performed with a fast reconstruction program of reduced performance.

In conclusion, we envisage to have two separate TPC reconstruction programs. The fast one will process online up to the hit reconstruction or track reconstruction depending upon the algorithm chosen, the processing power available, and the performance of the program itself. The second one would use the data in the optimal way which probably can be done only offline (in the traditional meaning). The second stream will also be used for calibration and for verification and development of the algorithm used in the fast stream.

10.6.2 Simulation software

The detector response is simulated by a package currently based on GEANT 3.21 (see also Chapter 11). With our complex and large events, we may be close to the limits of the current simulation environment. The RD44 project [34] aims at developing a simulation toolkit based on OO technology (GEANT 4), which might be better adapted to our needs. We are participating in the user requirement specification of this R&D project and plan to evaluate the resulting packages. We will progressively implement ALICE detector simulation using this new technology.

The computer resources necessary for full simulation of a single central event are very large, up to a few hours on an HP 735/125 MHz using up to 200 MByte of operating memory. Therefore we have to rely both on a 'fast' simulation program, based on parametrizations, and on a 'slow' one, very detailed for special physics cases and for verification of the 'fast' program. We estimate the storage volume needed for simulation to be about 20 % of that for real data.

10.6.3 Data analysis

The technology used today in data analysis will probably not be capable of handling efficiently the data volume produced by the ALICE detector. The first reconstruction pass, which will require only sequential access, should reduce the data volume by one order of magnitude. After this reduction we will still have about 100 TByte of data per year, which needs to be accessed and updated in different random ways. This has to be done both with short response time in interactive mode, which can be expensive in terms of hardware (fast disks, caching, parallel access, etc.), and in batch mode, using less resources but the same data model. Physics data analysis needs a diversity of retrieval functions, like queries for a few full events satisfying some criteria or access to a small number of attributes for all events. It will be desirable to store the reconstructed event data using the same OO database as for detector and calibration data thus allowing a consistent and transparent interface for all data types.

It is not obvious that commercially available database systems, even if they would be able to handle such data volumes, would do it efficiently in this case. The same is true for other data analysis tools, like histogramming, visualization, and graphical packages. Until now commercially available packages do not meet the requirements for physics data analysis. Unlike the simulation environment and databases, there is so far no centrally coordinated project dedicated to data analysis tools which involves OO technology. Therefore, in addition to evaluating existing systems, we plan to put some effort into the development of suitable tools.

10.7 Software engineering

Software activities inside any large collaboration will span a large number of applications, e.g. DAQ, online control, simulation, reconstruction and data analysis. The production of software for ALICE started a few years ago and will continue for more than a decade. The geographical distribution of the members of the collaboration and their unequal access to well-performing hardware and software tools complicate this task enormously. The use of formal methods and the utilization of the Computer-Aided Software Engineering (CASE) tools should help deal with this diversity.

Compatible approaches will be established for the analysis, design and development of the DAQ, simulation, reconstruction, and data analysis software. This has to take into account the large existing investment in simulation software. We believe that we will benefit from a comprehensive approach to all software tasks not only during the development stage, but also during exploitation and maintenance.

Training

The present evolution in software technology towards OO methods will be a major change in HEP computing and needs sufficient training to minimize the transition period. In addition to the existing and well-organized software courses, we feel the need for specialized training, in very small groups or even individually, for physicists with experience in HEP algorithms. We are preparing training sessions on the simulation code and on the geometrical database for ALICE.

In this context one has to change progressively the main programming language into one supporting OO technology. The obvious choice today is `C++` as the language widely used and supported. Despite the fact that most of the code in ALICE software has been written in `FORTRAN` there is a growing experience in our collaboration with `C++` coming from running heavy-ion experiments.

Methods and tools

The field of methods and tools is evolving rapidly and this constitutes a major problem for a long-term project. However, the OO methodology is emerging now not only in research institutes but also in industry. The recent unification of two major analysis and design methods (Booch and OMT) is an encouraging sign of convergence in the CASE industry.

We intend to apply a method and use supporting software tools in the framework of the prototyping work and in the R&D projects where ALICE is involved. This will include tools for defining user requirements, for the analysis and the design of the applications, and for the software development and code management. We plan to decide within the collaboration on specific tools by the end of 1996.

10.8 Computing infrastructure

The majority of the ALICE team has moved over the past two years from mainframes to workstations for interactive as well as for batch work. The main interactive platform is a UNIX workgroup server at CERN, which will be upgraded and merged with the NA49 server. For batch work we have used the CSF system at CERN and private workstations. Because of the large event size, our requirements in operating memory exceed what is usually available in general services. Therefore we will need a dedicated SHIFT service with large memory.

We are also using computing clusters at GSI Darmstadt and Utrecht University and numerous workstations in many of the collaborating institutes.

10.8.1 Organization of data flow and computing

Currently we foresee to record and store raw data sequentially on magnetic tapes close to the experiment on the I2 site, using about 30 to 50 recording devices in parallel. The amount of data recorded in one year will be ~ 1000 TByte, mostly accumulated during the few weeks of heavy-ion running. Unlike a central data-recording system, this solution will avoid a very high bandwidth connection to the experimental site, which would be practically idle for most of the year. Moreover, we will install in the GDCs located on the site a substantial part of our computing power, which we plan to use, together with the recording devices, during the rest of the year for the first reconstruction pass. We currently envisage to have up to 75 % of our CPU capacity and some TByte of disk space for staging on the experimental site.

The reduction factor in data volume after the first processing stage will be about one order of magnitude (see above). The resulting data volume of 100 TByte per year will be sent over the network to the computer centre for data recording. Here we envisage to have a random access disk space which can accommodate one years production, a back-up tape robot, and the rest of the computing power for further analysis and simulation. This is motivated by the better local and wide-area connectivity of the computer centre.

10.8.2 Computing power

The batch computing power is the dominant fraction of the computing infrastructure. We have estimated the ALICE needs by a comparison with NA49 and RHIC experiments similar to ours. A full NA49 event with 2000 tracks is processed in 10 minutes of an HP 735/125 MHz computer, which corresponds to a processing power of 32 CERN Units (CU).

Taking into account some differences between the current NA49 analysis and the one envisaged for ALICE, and scaling with the number of tracks to be recorded on average by ALICE per year, we estimate the total processing power needed for the reconstruction to be about 75 kCU. Adding 50 % for simulation and 15 % for other data analysis, we will need a total CPU power of 125 kCU ($\sim$ 220 GFlops).

Besides the CPU power, we are also concerned about the operating memory size required because of the large event size. We will benefit from the experience of other heavy-ion experiments in handling large event sizes [35]. Nevertheless, we estimate the necessary operating memory per processor for our batch activities to be about 1 GByte.

Resources for interactive work are today of the order of 30 CU for 10 users. If we assume an average of 150 interactive users at experiment start-up, this will require less than 1 kCU.

In Table 10.6 we show our current estimate and profile for computing resources needed for ALICE. To evaluate the number of CPUs we have assumed a doubling of processor power every two years and a processor lifetime of four years.

Table 10.6: Evolution of computing power and disk space.

Year	CPU power (kCU)	Number of CPUs	Disk space (TByte)
1996	0.1	3	0.03
1997	0.3	6	0.07
1998	0.6	9	0.1
1999	1.0	12	0.15
2000	1.6	13	0.3
2001	3.2	16	0.7
2002	6.4	23	3
2003	23	51	12
2004	63	100	55
2005	91	122	82
2006	125	138	110

10.8.3 Data storage

The raw data will be stored on tape volumes and kept ready for occasional re-reading on the experimental site, where we foresee a tape vault for some tens of PByte. This store may be organized in two levels, one of which should be close to the experiment with a capacity for one year of data. We do not envisage any physical exportation of raw data, therefore the exportability of media used for raw data recording will not be a restriction. A random access disk space of up to 10 TByte will also be necessary on the experimental site for a variety of functions, e.g. staging, detector and calibration data bases, test data recording, monitoring, etc.

In the offline computer centre we need to have about 100 TByte of random access disk space. A large part of this space will be dedicated to reconstructed data received from the experimental site. We will need here sequential access for data checking and further processing as well as random access for physics analysis. Local and wide-area network access have to be foreseen. In addition, a tape robot with a capacity of the order of some 1000 TByte is required mainly for data back-up and storage. A first estimate of disk space evolution is also shown in Table 10.6.

10.8.4 Networking

The work of dispersed collaborations relies to a greater and greater extent on effective networking. ALICE has similar needs to other experiments for standard networking such as interactive traffic, mailing, file transfer, etc.

In addition, the proposed organization of data flow requires a high-capacity link between I2 and the computer centre with a bandwidth of about 100 MByte/s. This will also allow us to use the large CPU power available on the experimental site, e.g. for simulation.

The need for external networking will also depend on the detailed computing model for ALICE and on the evolution of storage technology. Another crucial factor will be the evolution of telecommunication costs, especially after the foreseen deregulation in Europe. In order to allow for secondary data processing using computing power in external collaborating institutions on data sub-samples of reasonable size ($\sim$ 1 % of annual production), a link with a throughput of 1 GByte/s or more (i.e. two orders of magnitude above today's standards) will have to be economically advantageous.

10.8.5 Electronic mail and W3

Electronic mail is extensively used in the collaboration. The numerous mailing lists are extracted from a database of collaboration members. These lists are then converted into the different formats of the various platforms used in ALICE.

The World-Wide Web will most likely become the major tool for sharing and distributing information between the members of the collaboration. An adequate network infrastructure should be provided to and from CERN.

11 Physics performance

11.1 Signals and observables

In order to establish and analyse the existence of QCD bulk matter and the Quark–Gluon Plasma (QGP), a number of different observables have to be studied, including in particular global information about the events and signatures specific to the occurence of deconfined matter. The signals accessible to our detector are listed below, and some more detailed performance calculations for selected topics are given in the following sections.

11.1.1 Global event features

Global event features (particle and transverse-energy flow) are primarily a way to determine and select the geometry of nuclear reactions, i.e. the impact parameter, overlap volume, and number of constituents (nucleons, quarks) participating in the interaction. In addition, they indirectly reflect some of the underlying dynamics of nuclear collisions (nuclear stopping, primary particle production mechanism, rescattering), and specify the initial conditions (e.g. the energy and entropy density). The observable most directly related to the *collision geometry* is the energy carried away by the non-interacting beam nucleons (spectators), which will be measured in a forward zero-degree calorimeter (ZDC, see Section 9.3.2). The shape of the *pseudorapidity distribution* $dN/d\eta$ will distinguish different longitudinal expansion scenarios. *Structures* in rapidity or azimuth, caused, for example, by non-statistical fluctuations or collective phenomena, will be investigated event by event with a large acceptance multiplicity detector ($|\eta| < 4$) (see Sections 9.3.1 and 7.2.5).

11.1.2 Transverse momentum spectra

The transverse-momentum spectra can be divided roughly into three distinct regions —low (< 0.2 GeV/c), intermediate (0.2–2 GeV/c), and high (> 2 GeV/c)— where different processes are expected to be relevant.

a) Low momentum region
By the uncertainty principle, collective effects which occur over distances (or times) of 1–10 fm are associated in general with small momenta (200–20 MeV/c). One particular example is an excess of low-p_t pions produced coherently in the decay of a disoriented chiral condensate or a classical pion field. All these pions would have similar soft momenta and non-statistical fluctuations in the charged-to-neutral ratio. A modification of the pion dispersion relation in matter (quasi-pions) or a non-zero chemical potential (pion condensate) would lead to an enhancement of all pion species at low p_t. In order to have access to physics at long distances, the inner tracking system has a stand-alone capability as a spectrometer with momentum measurement and particle identification for soft particles ($p_t < 120$ MeV/c) which do not reach the outer layers (see Sections 11.3 and 11.4.1).

b) Intermediate momentum region
The intermediate momentum region, which includes the majority of all produced particles, will give information about the nature and dynamical evolution of the system. With up to 10 000 charged particles in the central rapidity acceptance, the *shape of the* p_t *distribution* can be measured for pions, kaons (~ 1000), and even protons (several hundred) on a single-event basis (see Section 11.4.3). The *average* p_t will be measured with high accuracy (~ 1 % for pions) and therefore individual events can be assigned a 'temperature', per particle type, which can be correlated with other observables. Thermodynamical instabilities during phase transition resulting in 'explosive' deflagration of the matter could lead to large event-by-event *fluctuations and extremely high temperatures*. Inclusive, high-statistics measurements of different particle types will allow the investigation of *expansion dynamics* and *collective flow phenomena*.

c) High momentum region
The propagation of high-p_t partons arising from an initial-state hard scattering is sensitive to the properties of the surrounding medium. In particular, the *energy loss of fast quarks and gluons* could be different in a dense hadronic system and a QGP (*jet quenching*). Experimentally, jets in nuclear collisions at the LHC cannot be measured well for p_t below 50–100 GeV/c [1] because of statistical fluctuations in the soft energy of the underlying event. However, hadrons at high p_t provide an alternative way to measure this effect. Similar to the jets themselves, single-particle spectra at high p_t (> 2–5 GeV/c) reflect the energy loss of partons (before fragmentation) and this results in striking differences in the shape of the spectra predicted for nucleus–nucleus and pp collisions. The high momentum region, which will be populated by only a few particles per event, will be measured in inclusive spectra, in particular in the region covered by the small area PID array (see Section 11.4.3).

11.1.3 Flavour composition

The abundance of particle species (hadrons/quarks) in chemical equilibrium is determined completely by a few thermodynamical quantities, i.e. temperature, mass, and chemical potential. The presence of (at least partial and local) equilibrium is crucial if we want to describe heavy-ion collisions in terms of QCD thermodynamics. This must therefore be verified by measuring the densities of a sufficient number of different particle species. Strange quarks are of particular importance because their production is favoured in a QGP, since their mass will be reduced from the constituent (~ 500 MeV) to the current value (~ 150 MeV) (partial chiral symmetry restoration). During the dynamic evolution the particle composition will change, but each particle species provides a unique 'clock' which characterizes the system on a different time scale. Reactions with large cross-sections at thermal energies (small relaxation times), such as K production, will stay in equilibrium until freeze-out. Heavier particles, in particular Ξ and Ω, are difficult to produce in thermal collisions (long relaxation time) and their abundance might not change significantly during the short lifetime of the hadron gas (< 10 fm/c). Multi-strange hyperons therefore seem to be very good signatures for the large strangeness density expected in the QGP. If extremely high temperatures should be reached, even charm quarks will be produced in secondary collisions during the pre-equilibrium or thermal stages, a feature probably unique to ion collisions at the LHC. The detection of hyperons and in particular charmed mesons requires a vertex detector close to the beam pipe (see Sections 11.8 and 11.9).

We will measure simultaneously a large number of stable particles and resonances (π, η, ω, ϕ, p, K, Λ, Ξ, Ω, D mesons, and light nuclei) as a function of charged-particle density and p_t. This will allow us to test *thermalization*, will constrain the *dynamic evolution scenarios*, and should also yield information about the *phase transition* itself. Some particle ratios can be accurately measured on an event-by-event basis to look for fluctuations or to correlate them with other observables.

11.1.4 Correlations and fluctuations

One of the unique features of heavy-ion collisions at the LHC is the possibility to measure a large number of observables on an *event-by-event* basis: impact parameter (via ZDC), dN/dy, particle ratios and p_t spectra ($\pi, \gamma, K, p, \Lambda$) and, of particular importance, size and lifetime from interferometry. To make full use of this opportunity is one of the important design considerations for our detector because it will allow the study of correlations and non-statistical fluctuations which would otherwise be washed out when averaging over many events. Such fluctuations are, in general, associated with *critical phenomena* in the vicinity of a phase transition, and lead to local or global differences in the events, even for similar initial conditions.

Structures in dN/dη, which can be measured over ~ 8 units of rapidity in our detector, have been predicted to arise from the hadronization of the QGP. *Multiparticle correlations* are a sensitive tool for studying correlations over many length scales and will be carried out in three dimensions (y, p_t, φ) in the central region [2]. A large imbalance between hadronic and electromagnetic energy, as observed in the still unexplained CENTAURO events in cosmic-ray exposures, will be observable if present at mid-rapidity, using the single-arm e.m. calorimeter.

Bose–Einstein correlations between identical particles should give access to the space–time history of the particle-emitting source. High-statistics data will be available at the LHC to do a differential Hanbury-Brown–Twiss (HBT) analysis (R_L, R_{Tside}, R_{Tout}, see Section 11.5) as a function of dN/dy and p_t. The correlation for a number of particle types can be investigated ($\pi^\pm, K^\pm, p, \overline{p}, K^0_S$), which will yield different radii if they freeze out at different times. This comprehensive analysis will provide detailed information on the *freeze-out geometry*, on the *expansion dynamics* in the hadronic phase and, possibly, on the existence and nature of the *phase transition*. Moreover, if the transition is of first order, the long lifetime in the mixed phase will lead to large values of R_L.

The large radii (of the order of ~ 20 fm) and the long lifetimes ($\gtrsim 10$ fm/c) expected at the LHC place demands on the detector in terms of momentum and two-track resolution. We will be able to measure, on an event-by-event basis, radii up to 15 fm, and inclusively up to 30–40 fm (see Section 11.5.3).

No primary nuclear fragments are expected at mid-rapidity at the LHC. Deuterons and heavier nuclei result from the *coalescence* of baryons; their abundance is governed by the density and lifetime of the hadronic phase. The d/p ratio will therefore give information complementary to the HBT analysis.

11.1.5 Prompt photons

Electromagnetic probes are a unique tool for studying the early and hot phases of nuclear collisions because the produced photons and lepton pairs escape from the finite system without re-interaction. The *thermal radiation* from the QGP and mixed phase could be observable in the medium-p_t range (around 1–3 GeV/c). Direct photons from the *pre-equilibrium* phase (3–6 GeV/c) contain information about the parton dynamics in the very early stages, about the equilibration times, and about the transition from perturbative to non-perturbative phenomena. At still higher momenta, direct photons originate mainly from *initial hard processes* and can be used to extract structure functions in nuclei (in particular from p–A reactions).

Because of the large background arising from hadronic decays (mainly π^0 and η), a precise measurement of the rates and p_t spectra of these mesons is vital in order to extract the prompt photon component. The predicted signal of thermal photons is of the order of a few per cent of the total inclusive photon yield, depending on initial conditions and on the shape of the meson's p_t spectra. It will, in general, be more favourable at high initial-energy density and, indeed, might only be observable at the large densities predicted for the LHC. We aim at a sensitivity, limited by systematic errors, of $\sim 5\,\%$ for the measurement of prompt photons in the energy range 1–10 GeV (see Section 11.7).

11.1.6 Lepton pairs

The physics interest in 'continuum' lepton pairs is identical to the one mentioned above for photons, i.e. thermal radiation at low masses, pre-equilibrium, and hard processes (Drell–Yan, open charm production) at high masses. In addition, the *vector mesons* (ρ, ω, ϕ, J/ψ, ϒ) can be observed in their leptonic decay modes. The lighter mesons, with lifetimes of the order of the expansion time scale, will partially decay during the evolution of the system. Their properties (mass, width, branching ratios), observable through the leptonic decay, should *change in dense matter* owing, for example, to 'collision broadening', 'induced radiation' and other in-medium effects. In the vicinity of the phase transition, partial restoration of chiral symmetry will lead to additional shifts of the mass. The ϕ meson, only 30 MeV above the KK threshold, is expected to be extremely sensitive. A few MeV shift of the in-medium kaon mass would drastically alter the relative branching fractions of $\phi \rightarrow$ KK and $\phi \rightarrow e^+e^-$. The width of the ω might reach 50–60 MeV, i.e. more than five times the normal value. The magnitude and even the sign of these effects are difficult to predict with current theories (mainly chiral perturbation theory), but will be experimentally unambiguous if a change of *resonance line shapes* is observed (see Sections 11.4.3 and 11.6).

The *suppression of heavy quarkonia resonances* via Debye screening is an important tool for diagnosing the early stages of the QGP. A full 'spectral analysis' of the J/ψ and ϒ families as a function of energy density and transverse momentum should distinguish QGP from hadronic suppression scenarios.

Because of the large combinatorial physics background from Dalitz decays, continuum lepton pairs will not be easily accessible at high multiplicities. However, we will measure ω and ϕ production, via their e^+e^- decays, with a mass resolution of $\Delta m/m \approx 1\,\%$ (see Section 11.6). The rates will also be sufficient to measure inclusive J/ψ production at low p_t in the electron channel at mid-rapidity. The forward muon arm ($\eta = 2.4 - 4$) was specially designed to measure the complete spectrum of heavy quarkonia resonances, i.e. J/ψ, ψ′, ϒ, ϒ′ and ϒ″, via their muonic decay, with a mass resolution sufficient to separate all resonance states (see Ref. 3 and Chapter 6).

11.1.7 Strangelets

Long-lifetime massive objects of strange matter, composed of roughly equal numbers of up, down and strange quarks, could be produced in heavy-ion collisions, either in the form of elementary *strangelets* or as Meta-stable Exotic Multi-hyperonic Objects (MEMOs). Strangelets, besides being of interest for cosmology as a source of dark matter, would be a unique signature for QGP formation, and are currently being searched for both at the SPS and the AGS. Although the baryon free region at LHC mid-rapidity might be unfavourable for their production, we have investigated the performance of our detector with respect to this signal, as an example of its sensitivity to recognize less conventional signals (see Section 11.10). A study is also underway to estimate sub-threshold production of hypothetical superheavy objects (see Ref. 4).

11.2 Event generators

In view of detector feasability and performance studies, most of the calculations presented in the current document concerning nucleus–nucleus collisions were performed using a phase space generator program (SHAKER) with easily adjustable particle ratios, multiplicities and p_t spectra. Essentially all calculations were performed with $dN_{ch}/dy = 8000$ for charged particles at mid-rapidity, which is close to the maximum rapidity density anticipated (see below).

In this section we present a short comparison of SHAKER with various commonly used models concerning global features relevant for detector and DAQ design, i.e. particle density and the transverse momentum spectra. The particle density directly influences the detector granularity and therefore the DAQ design, via the number of readout channels and the data volume. In addition, it also affects the performance of the pattern recognition algorithms.

In the simulation two main event classes were used: a central event sample and a peripheral event sample consisting, respectively, of 5 % and 30 % of the geometrical cross-section.

Calculations concerning pp minimum biased studies have also been performed, using the DTUJET [5] and PYTHIA [6] event generators. These generators are not described in this context, but a detailed writeup of the comparison results of both pp and nucleus–nucleus event generators can be found elsewhere [7].

11.2.1 The various models

SHAKER

The SHAKER program is a simple central rapidity phase space event generator developed for the simulation of LHC heavy-ion events [7]. All event information is managed according to the Lund conventions within the JETSET 7.3 program [8]. The main input parameter is the charged particle density at central rapidity, according to which the various types of particles are generated (π^0, η, π^+, p, K^+, K^0 and their antiparticles). The rapidity distributions are flat between fixed limits. The transverse momentum distributions are those measured at the Tevatron [9], with the exception of that of the η particle, for which the Tevatron π^0 distribution has been used applying m_t-scaling [10]. Particle ratios are fixed by isospin invariance ($\pi^0/\pi^+, K^0/K^+$), by the measured Tevatron values ($K^+/\pi^+, p/\pi^+$) [9], or by m_t-scaling of the measured asymptotic values (η/π^0). In the generation of the central event sample discussed here, a rapidity window $-1.5 \leq y \leq 1.5$ and a flat charged particle density of $dN_{ch}/dy = 8000$ were used.

VENUS

This model [11] for particle production in nucleus–nucleus collisions is based on the Gribov–Regge theory of multiple pomeron exchange and classical relativistic string dynamics. VENUS (Very Energetic NUclear Scattering) includes the participation of (anti)quarks in the colour exchange mechanism to form strings, the possibility of diquark break-up leading to double-strings, a sophisticated fragmentation procedure, a resonance table including, for example, all nucleon resonances up to 2 GeV, and a covariant treatment of secondary interactions. In the analysis presented here, VENUS 4.12 was used.

HIJING

This model combines a QCD-inspired model for jet production [12] with the Lund model [13] for jet fragmentation. The formulation of HIJING (Heavy Ion Jet INteraction Generator) was guided by the Lund FRITIOF [14] and Dual Parton Model [15] for soft nucleus–nucleus reactions at intermediate energies ($\sqrt{s} \approx 20$ GeV/nucleon) and the successful implementation of perturbative QCD processes in the PYTHIA model [6] for hadronic collisions. To study nuclear effects, nuclear shadowing of parton structure functions and a schematic model of final-state interactions of high p_t jets in terms of an effective energy-loss parameter (dE/dz) have been included.

HIJING does not include secondary interactions. The event samples used in the analysis described here were produced using HIJING 1.31.

DPMJET–II

The DPMJET–II event generator [16] is based on the Dual Parton Model [15]. The Dual Parton Model provides a unique picture not only for hadron–hadron interactions, but also for the description of particle production in hadron–nucleus and nucleus–nucleus collisions at high energies. The two-component Dual Parton Model [17] is used with multiple soft chains and multiple minijets at each elementary interaction. The minijets are calculated using parton structure functions which are consistent with the HERA measurements at small x. Within this model the high-energy projectile undergoes a multiple scattering process as formulated in Glauber's approach; particle production is achieved by the fragmentation of colourless parton–parton chains constructed from the quark content of the interacting hadrons. DPMJET–II includes the cascading of secondaries within the target which is suppressed by the formation time concept. The excitation energy of the remaining target and projectile nuclei is calculated and used in the nuclear evaporation included in the model. The main use of the model so far was for simulating hadron–nucleus and nucleus–nucleus collisions in the cosmic ray cascade. It has been checked recently that DPMJET–II also works for heavy-ion collisions up to LHC energies.

SFM

The String Fusion Model (SFM) event generator [18] is based on the parton picture of strong interactions, connected with the structure and properties of multipomeron vertices. This picture presents a high-energy hadron or nucleus collision as a single interaction between two clouds of moving partons. The number of interacting partons, which defines the number of newly produced strings, increases with the energy and atomic numbers of the colliding nuclei, while the parton–parton interaction cross-section remains the same. In Pb–Pb (pp) collisions at the LHC, the model predicts the formation of 4015 (15.4) strings with a density of 28.5 (2.2) fm^{-2}. Therefore, high densities can lead to the overlapping and interaction of the strings. The strings can fuse with a cross-section comparable to the one of the parton–parton interactions. A colour charge at the end of the fused strings is obtained according to the standard SU(3) colour composition law. Strings decay in pairs of coloured partons which neutralize the colour string field. The high colour charge of fused strings leads to an increase of baryon–antibaryon pairs and strange particles, and thus to a reduction of the total multiplicity of the reaction.

The SFM event generator can be used in two different modes: with string fusion and without string fusion. No rescattering is included in the current version.

11.2.2 Comparison of the spectra

As mentioned before, our comparison of the various models will be focused primarily on the particle density and transverse momentum spectra. In Fig. 11.1 the various multiplicity distributions for the central event samples are shown.

The charged particle densities $dN_{ch}/d\eta$ at mid-rapidity for VENUS, SHAKER, HIJING, DPMJET-II and SFM (without and with fusion) were found to be 6800, 6200, 5200, 3800, 3700 and 1400, respectively. Note that $dN_{ch}/d\eta$ is typically some 20 % lower at midrapidity than dN_{ch}/dy. We see that SHAKER, with our current input parameters, overestimates nearly all of the predicted particle multiplicities at mid-rapidity. The SFM generator is seen to yield relatively low particle densities around mid-rapidity, whereas the VENUS particle densities are found to be the highest ones.

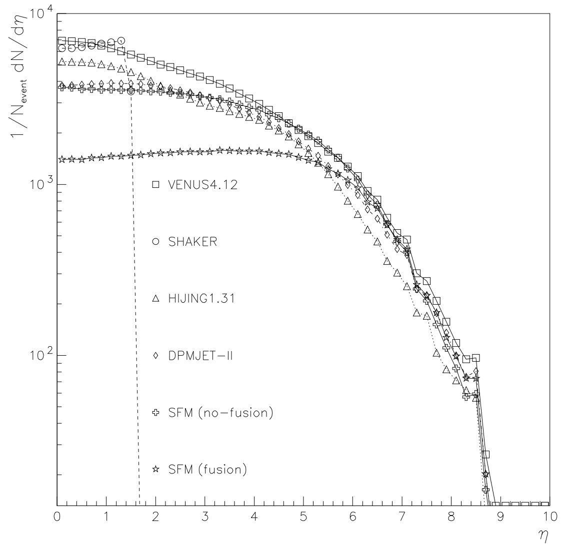

Figure 11.1: Charged particle multiplicities for central Pb–Pb events at a beam energy of 3 TeV per nucleon. The generators have been ordered by decreasing multiplicity at $\eta = 0$.

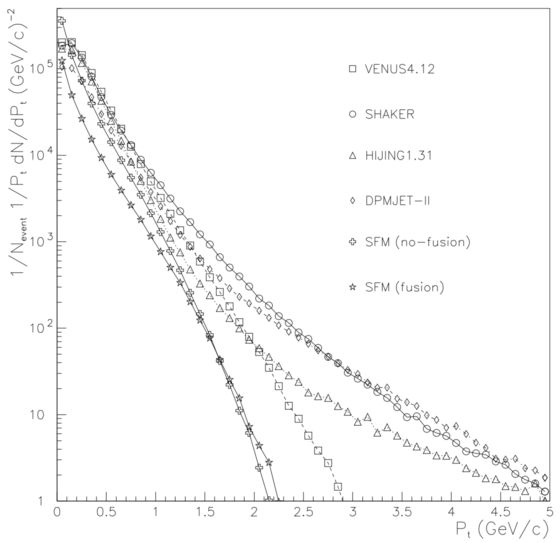

Figure 11.2: Charged particle transverse momentum distributions in the pseudo-rapidity interval $-1.5 \leq \eta \leq 1.5$ for central Pb–Pb events at a beam energy of 3 TeV per nucleon.

The maximum number of charged particles within the acceptance of the barrel detectors (i.e. $|\eta| \leq 0.9$) is about 12 000 for VENUS and about 11 000 for SHAKER.

The various transverse momentum distributions for the central event samples are presented in Fig. 11.2.

The mean p_t values for SHAKER, VENUS, HIJING, DPMJET-II and SFM (with and without fusion) were found to be 525, 442, 432, 560, 403 and 352 MeV/c, respectively. Here it is seen that for $p_t \leq 1$ GeV/c (i.e. the region around the mean p_t values) most models agree with each other. The SHAKER data reproduce the the other models rather well. For $p_t > 1$ GeV/c the various models start to differ substantially and the SFM generator produces a more soft spectrum than the other models.

Table 11.1: Averaged material budget in the tracking volume at normal incidence, as used in the simulations.

	Item	Radius (cm)	$X/X_0(\%)$
ITS	Beam pipe	3.0	0.17
ITS	Si pixel 1	3.9	0.64
ITS	Si pixel 2	7.3	0.64
ITS	Si drift 1	14	0.64
ITS	Si drift 2	24	0.64
ITS	Si strip 1	40	0.64
ITS	Si strip 2	45	0.64
ITS	Outer shell	50	1.00
ITS	Air	3–52	0.16
TPC	Inner vessel	52	0.27
TPC	CO_2	52–78	0.14
TPC	Inner cage	78	0.52
TPC	Working gas	78–250	0.56
	Total	0–250	6.66

11.3 Tracking performance

The performance of the tracking system has been evaluated using the GEANT 3.21 package [19]. A detailed description of the ALICE simulation package can be found in Ref. 20.

All results presented below were obtained assuming a particle density of $dN_{ch}/dy = 8000$ and a material distribution as given in Table 11.1. In addition to the different sub-detectors we have also included mechanical supports, cooling systems, vacuum chamber, electronics and wiring, according to our present detector description. In some cases (e.g. ITS) we have used a detailed description taking into account a realistic distribution of materials, while for other sub-detectors (e.g. TPC) only simplified geometries with average material distributions have been used.

We have included all relevant physics processes in this simulation (e.g. multiple scattering, energy loss, bremsstrahlung, etc.). The efficiencies presented below are normalized to 'recognizable' tracks (unless otherwise noted), i.e. particles that leave the acceptance, are stopped in the material or particles that decay are not counted in the efficiency.

In several of the performance studies presented below we have used a 'fast' program based on parametrizations of the full detector simulation.

11.3.1 Tracking requirements

The ALICE tracking system must have an excellent track recognition capability up to the highest envisaged track densities in Pb–Pb central collisions. In order to have sufficient statistics in a single event for event-by-event analysis and for rare and/or large background signals (e.g. lepton pairs, charm, strangelets), we aim at 90 % or better track-finding efficiency above $p_t > 100$ MeV$/c$. Fake tracks ('ghosts') must be kept as low as possible. In addition, we require some track-finding capabilities for low-p_t tracks, in particular for electrons and positrons, down to a p_t of 20–30 MeV$/c$, in order to suppress the background in the electron pair analysis (see Section 11.6 for details). We also need a good double-track resolution to be able to measure large radii with particle interferometry (HBT, see Section 11.5).

For most of the observables we have to reconstruct low-momentum particles (from 100 MeV$/c$ up to 1 GeV/c). In order to achieve the required effective mass resolution for different particles and resonances (e.g. hyperons, D mesons, J/ψ, ρ, ω, ϕ) and a sufficient precision in the HBT variables, the relative momentum resolution should be about 1 % for average momenta ($\sim$ 0.5 GeV/c) and should not exceed 3 % for the higher momentum range ($\sim$ 10 GeV/c). The angular precision should be of the order of 1 mrad. The detection of open charm particles and, to a smaller extent, of hyperons requires an excellent impact parameter resolution, significantly better than 100 μm for the charm measurements.

The resolution for low-momentum tracks is dominated by multiple scattering and not by the position resolution. The multiple scattering contribution to the relative momentum error depends on the magnetic field B and on the tracking length l according to $(\Delta p/p)_{\text{scat}} \propto 1/(B \times l)$. In contrast, the dependence of the position measurement is $(\Delta p/p)_{\text{meas}} \propto 1/(B \times l^2)$. Track finding is more efficient in a low magnetic field, in particular at small momenta, as the track curvature is smaller. These considerations, amongst others, have led us to choose a low magnetic field ($B = 0.2$ T) and to track over a large distance ($l = 2.5$ m) with two 'light' (low material) tracking detectors:

- TPC: a 'continuous' tracker which has to assure a high track-finding efficiency, even at large track densities;
- ITS: a vertex detector which is used to improve the momentum resolution and to detect secondary vertices.

The track-finding performance and track parameter resolution will be described in the following section.

11.3.2 Track finding

Tracking strategy

The track finding in ALICE uses the following approach:

- track finding in the TPC;
- 'guided' track finding in the ITS, i.e. following the tracks found in the TPC back to the ITS;
- stand-alone track finding in the ITS for low-momentum tracks, i.e. those which are not found or do not reach the TPC.

We have also tested a different strategy, similar to the one adopted by the STAR Collaboration [21]. It is based on separate tracking in the TPC and in the vertex detector, followed by track matching in the minimum error plane. Our study (see below) has shown a strong p_t dependence of the efficiency for the stand-alone ITS track finding at high track densities. In this case the overall tracking efficiency is dominated by the ITS track finding, which is insufficient at low momenta.

The choice of the track-finding strategy has an influence on the arrangement of the TPC–ITS interface. A consequence of our strategy is that we should concentrate the material between the two detectors as close as possible to the ITS. Moreover, this improves the momentum resolution.

We have prototyped various track-finding methods for the ALICE detector, based on experience gained in the SPS heavy-ion experiments WA97 [22] and NA49 [23] which have shown that a good tracking performance can be achieved in a high-multiplicity environment very similar to the one in ALICE. The methods used can be divided into two categories: local methods and global methods [1].

The local methods follow the tracks by collecting the hits step by step, after initial selection of track candidates. Typical examples of this type of algorithms are the Kalman filter [24] and other track-following methods [25].

The global methods are based on a global transformation of measured information for all hits (or part of them if pre-tracking was used). Examples, which we are evaluating, are the Hough–Radon transformation [26] (see Ref. 27 for application in ALICE) and neural network algorithms [28].

The advantage of the global compared to the local methods is in their speed, because they significantly reduce combinatorics. In high track density conditions, in particular, the local methods search branches off as soon as the information available at each step is not sufficient to take a definitive decision. However, up until now we have found the local methods to be more robust and efficient.

We plan to continue the study of track-finding algorithms before taking any decision for our reconstruction software. In the following we present the results obtained by a local method (track following) for the TPC and connection to the ITS, whereas the global method (neural network) is used for the ITS stand-alone track finding.

[1] In some cases a combination of the two methods is used.

Vertex finding

Vertex finding can be achieved prior to tracking by a simple correlation of hits in the two Si pixel detectors. Figure 11.3 illustrates this correlation. Using the average vertex position and a hit in the first pixel layer we predict the impact in the second pixel layer. Using this predicted position we define a window in $r\varphi$ in the second pixel plane. Figure 11.3 shows the correlation between z coordinates of each hit in the first plane with all hits in the corresponding $r\varphi$ window in the second plane. A correlation band can clearly be seen, and this determines the z position of the vertex. The procedure is repeated for the $r\varphi$ coordinate. The precision of this method depends on the track multiplicity and can be parametrized as:

$$\sigma_{r\varphi} = 100\,\mu\mathrm{m}/\sqrt{\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta} + 6\,\mu\mathrm{m}$$

$$\sigma_{z} = 250\,\mu\mathrm{m}/\sqrt{\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}\eta} + 7.5\,\mu\mathrm{m}$$

in $r\varphi$ and z directions, respectively.

While a tight vertex constraint is in general not used during track finding, the stand-alone vertexing capability of the two pixel planes is important during high-luminosity running for the muon arm. In this mode, most other detectors are not operational and a good vertex is needed to achieve the required dimuon mass resolution.

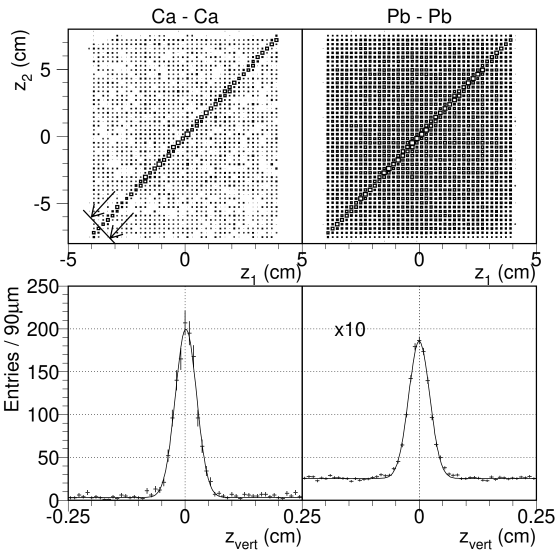

Figure 11.3: Correlation between the z coordinates of the hits in the two pixel layers in a restricted $r\varphi$ window (top). In the bottom part the diagonal projection of the correlation band, which determines the vertex coordinate, is shown. The amplitude of the correlation and the background is strongly multiplicity dependent. The examples shown correspond to the Ca–Ca and Pb–Pb multiplicity densities.

TPC track finding

The TPC track finding starts from the outermost pad rows where the hit density is lowest and where the precision is less affected by the limited double-track resolution. From initial hit triplets we extrapolate track candidates inwards and add the hits which are closest to the predicted positions within some defined limits. A maximum of two subsequent pad rows without hits can be bridged. At $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$, 83 % of the hits are 'singles', i.e. separated by more than the double-hit resolution. We allow for a hit to be shared by up to four tracks. We use three passes, the second with a wider road and the third starting closer to the interaction point.

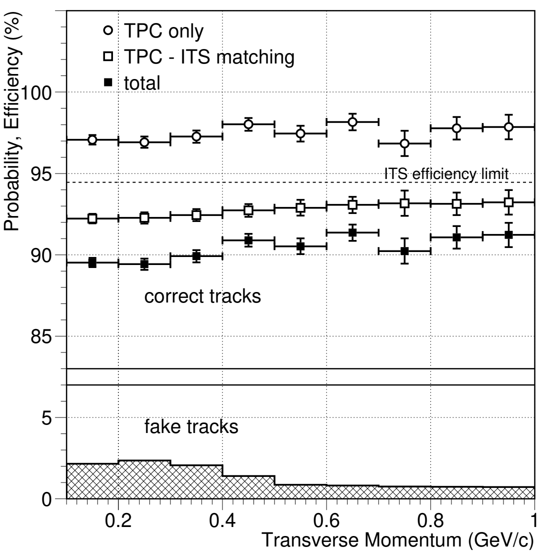

Figure 11.4: Track-finding efficiency and fake-track probability as a function of p_{t}. The contributions to the efficiency from the TPC track finding and the matching to ITS are shown separately. The results were obtained by the track following method.

The track finding in the TPC achieves an efficiency of close to 97 % for p_{t} as low as 100 MeV$/c$ (see Fig. 11.4). The efficiency seems to be independ on p_{t} and η. The fraction of fake tracks is below 10^{-3}. This excellent performance confirms the remarkable tracking qualities of TPC detectors, even in an environment as hostile as ion collisions at the LHC.

TPC–ITS matching

The tracks found in the TPC are ordered according to their momenta and followed back towards the ITS, starting with the highest momentum track. A vertex constraint is imposed only at this stage. A road is defined around the predicted trajectory to search for a consistent set of hits in the ITS. The precision of the extrapolation from TPC to ITS is determined largely by multiple scattering in the material between the two detectors. Therefore, the width of the road is momentum dependent, and hence we start with a small search area and then continue with a larger and larger one as we proceed to lower momentum tracks. Whenever an ambiguity arises in one of the two outermost layers, it is resolved by choosing the hit combination with the smallest amplitude difference on the two sides of the detector (as discussed in Section 2.4). Whenever a successful TPC–ITS match is found, the assigned hits are removed, thereby progressively lowering the hit density.

Several tracking passes are made with different constraints. The first one requires hits in all the ITS layers. Further passes have looser vertex constraints and require less hits. In the final pass the hits already assigned are reused, thus allowing for hit sharing. The final requirements are:

- at least three hits in the two double-sided microstrip detectors;
- at least one hit in the two silicon drift detectors;
- hits in both pixel detectors.

As we have assumed an intrinsic efficiency for the pixel layers of 98 % and for the other ITS sub-detectors of 95 % (which includes geometrical coverage, broken channels, etc.) the above conditions can be fulfilled by at most 94.5 % of all tracks.

The overall track-finding efficiency and the contribution from each step is shown in Fig. 11.4. The efficiencies are almost constant over the full p_t interval. The small decrease at low p_t (a few per cent) is due to a poorer extrapolation from TPC to ITS for low-p_t tracks. At low p_t, the fake-track probability is largest, but still only about 2 %. A track is regarded as 'fake' if a single wrong hit has been assigned, i.e. many 'fake' tracks are only slightly mismeasured real particles.

We do not observe any η dependence for the efficiency or for the fake tracks. The average efficiency is about 90 %, where 2.5 % is lost in TPC track finding and 7.5 % is lost in the matching to the ITS. Only 2 % of the matching loss can be attributed to the tracking algorithm, the rest (5.5 %) is caused by the intrinsic inefficiencies of the ITS detectors.

Low-momentum tracks

After the removal of all of the ITS hits assigned to tracks, we proceed with the stand-alone finding for tracks with $p_t < 100$ MeV$/c$ in the ITS. The method, based on a neural network algorithm [29], is similar to the one adopted by the CMS Collaboration [1].

Figure 11.5 shows the track-finding efficiency for electrons with p_t between 20–100 MeV$/c$ together with the fraction of fake tracks. We reach on average an efficiency of 63 % for electron finding (84 % for $p_t > 30$ MeV$/c$)

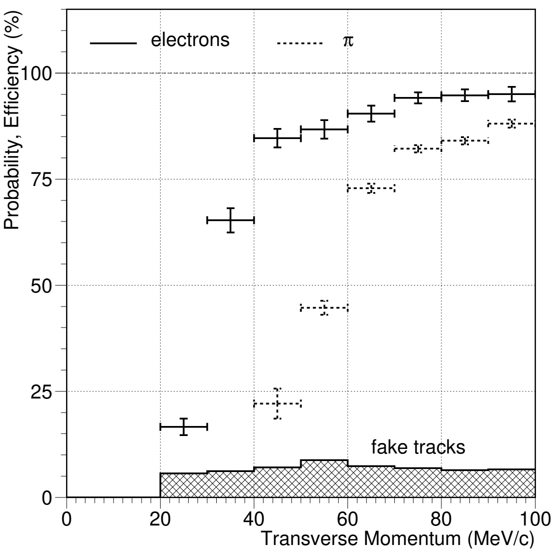

Figure 11.5: Track-finding efficiency for low-momentum electrons and π mesons and fake-track probability as a function of p_t. The results are obtained by the stand-alone ITS track finding using the neural network algorithm.

with a 7 % fake-track contamination. The efficiency for π mesons is similar down to a p_t of about 60 MeV$/c$. However, for $p_t < 60$ MeV$/c$, the probability of finding π mesons decreases more rapidly because of ionization energy losses in the material of the ITS (see below).

Robustness of the track finding

In order to check the robustness of our track finding, we have tested a number of different alternatives and varied some of the parameters, in particular the amount of material in the ITS and the particle multiplicity.

Figure 11.6 shows the track-finding efficiency and the fake-track probability as a function of (i) the material thickness between the TPC and the ITS and (ii) the multiplicity. The number of fake tracks as well as the tracking inefficiency are directly proportional to the multiplicity, but both are still small even at twice the normal particle density.

In one of the tests we postpone the resolution of space point ambiguities in the two double-sided Si microstrip layers until the guidance from the TPC is received. The efficiency is worse at low p_t by less than 10 %, but with this algorithm we do not depend on the pulse-height matching capability of the silicon strips.

We have performed a preliminary study of the stand-alone track finding in the ITS at full track density — assuming no 'help' from the TPC— and with increased magnetic field (needed in this case for good momentum resolution). The track-finding efficiency at higher p_t is still good at about 90 %. Below 400 MeV$/c$ it decreases

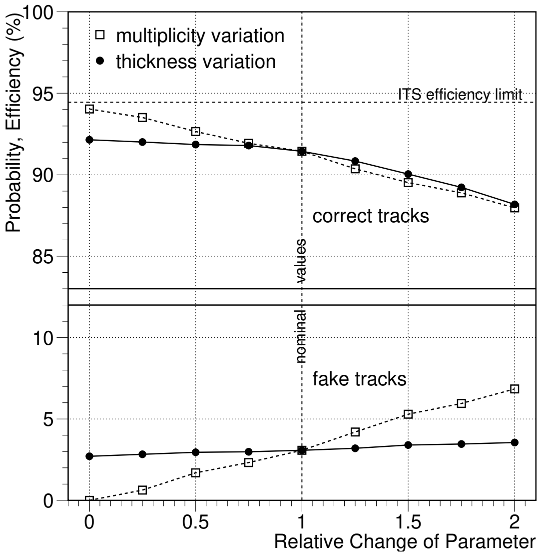

Figure 11.6: Variation of the efficiency (top) in the 'guided' ITS track finding for pions with $p_t \approx 100$ MeV/c as a function of: (i) the amount of material between the TPC and the ITS; (ii) the track multiplicity. The value '1' corresponds to nominal parameters (see text), and (e.g.) '1.5' to a 50 % increase in thickness (full circles), or multiplicity (open squares). The lower set of curves show the fraction of fake tracks in the ITS.

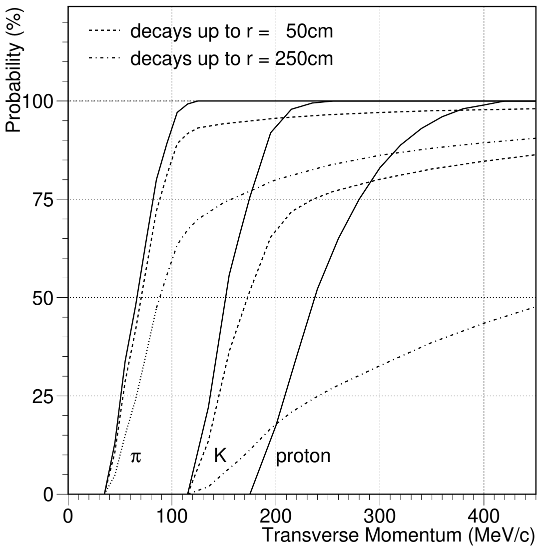

Figure 11.7: Fraction of pions, kaons and protons entering the TPC. The full lines show the effect of the ionization energy losses. For mesons, decays are included up to the end of ITS (r = 50 cm, dashed lines) or up to the end of TPC (r = 250 cm, dashed-dotted lines).

gradually, reaching a value of 75 % at $p_t = 100$ MeV/c. At lower track densities, $dN_{ch}/dy \approx 1000$ (e.g. in Ca–Ca collisions), this loss of efficiency at low p_t disappears.

11.3.3 Tracking resolutions

The performance of the tracking system has been evaluated with the detector characteristics as described previously (see Section 2.1 for the ITS and Section 3.11 for the TPC). In addition to the intrinsic resolution, we have added *linearly* a residual misalignment of 5 μm for the ITS and of 10 μm for the TPC. The alignment error in the radial direction was assumed to be 10 μm for the ITS and 75 μm for the TPC, respectively.

Momentum resolution

The momentum resolution is strongly affected by the multiple scattering over most of the p_t range of interest to us. An additional problem at very low momenta is ionization energy loss and fluctuations thereof. Therefore, one of the key design considerations was to keep the amount of material in the tracking volume as low as possible. Table 11.1 on page 157 summarizes the thicknesses seen by particles in the tracking volume at $\eta = 0$.

In Fig. 11.7 we show the probability of different hadrons penetrating the material of the ITS. Efficient detection of π mesons is only possible at $p_t > 80$ MeV/c, for K mesons at $p_t > 180$ MeV/c and for protons at $p_t > 300$ MeV/c. Near these thresholds, the fluctuations in energy losses have to be taken into account.

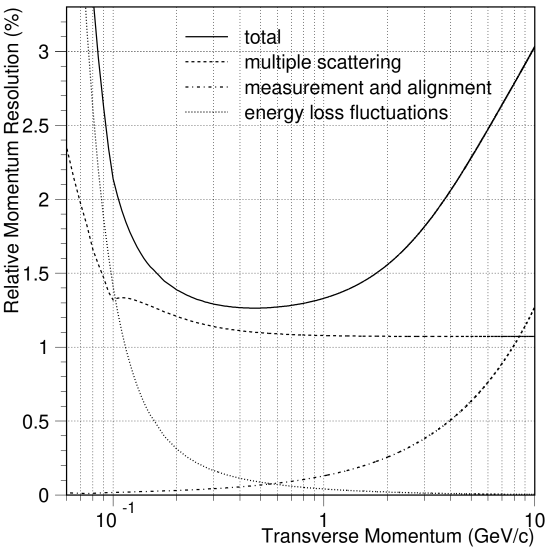

Figure 11.8: Momentum resolution for π mesons as a function of p_t integrated over η in the tracking acceptance. The contributions due to multiple scattering, measurement/alignment precision and energy-loss fluctuations are shown separately.

Figure 11.8 shows the overall detector momentum resolution for π mesons as a function of p_t and the contributions due to multiple scattering, measurement precision and alignment, and the fluctuation in energy loss [2]. Over most of the momentum range the resolution is dominated by multiple scattering, which increases with $1/\beta$ at low p_t. This velocity dependence is responsible for the difference between various particle species visible in Fig. 11.9.

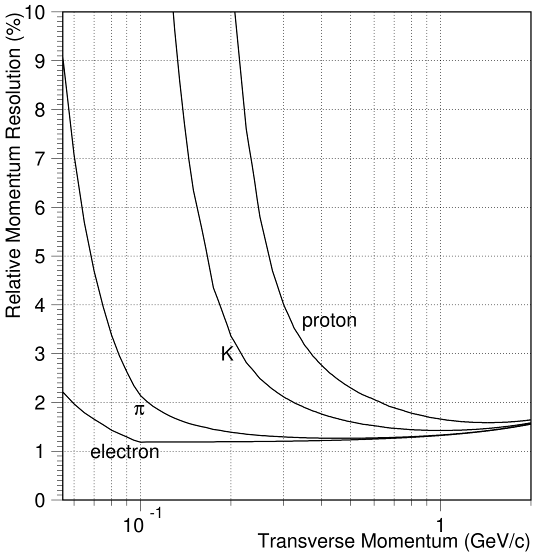

Figure 11.9: Momentum resolution for electrons, π mesons, K mesons and protons as a function of p_t near to the detection thresholds.

The momentum resolution for pions depends only weakly on η, as shown in Fig. 11.10, and is insensitive to the vertex resolution within our p_t range.

Angular resolution

The angular resolution, which is shown in Fig. 11.11, depends on the vertex precision and hence on the multiplicity (see above). Angles are determined largely by the vertex position and the first measured point in the pixel layer, since here multiple scattering effects are small. Therefore the angular resolution is worse in pp events than in high multiplicity heavy-ion events. Without vertex information, the angular resolution is worse, but independent of the multiplicity. For ions heavier than Ca we reach the limit determined by the precision of the Si pixel detector and the angular resolution does not improve further.

Mass resolution

The mass resolution for a number of decay channels of interest to ALICE are summarized in Table 11.2. These resolutions are used in other parts of this document to evaluate physics performance.

[2] The decomposition into various sources is not a 'well-defined' procedure, because each of them has a completely different correlation matrix and thus one has to use different estimators. Therefore the total resolution is not the quadratic sum of the individual contributions.

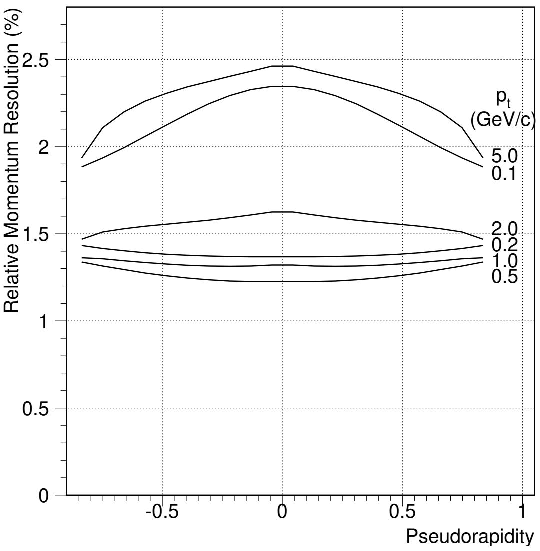

Figure 11.10: Pseudorapidity dependence of the momentum resolution for π mesons at different p_t.

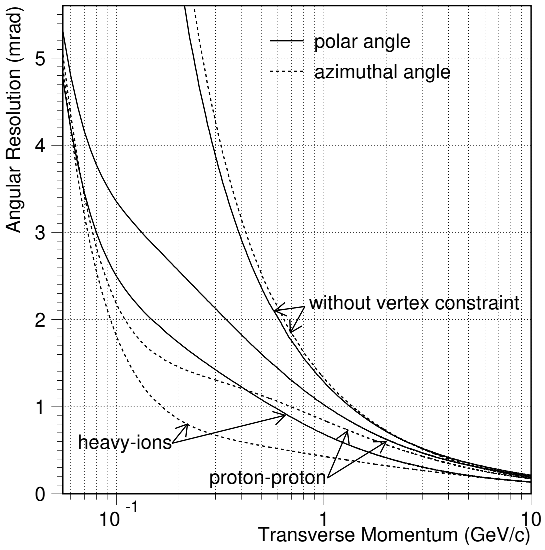

Figure 11.11: Resolution of polar and azimuthal angles for π mesons as a function of p_t for heavy-ion and pp collisions. For comparison, the resolution without vertex constraint is shown.

Impact-parameter resolution

In order to detect unstable particles and secondary vertices, we require good impact-parameter resolution. The p_t dependence of the impact-parameter resolution is shown in Fig. 11.12 for different particle species and for the $r\varphi$ and z directions. The resolution is more than sufficient for the hyperon detection (Section 11.8) down to low momenta. In the bending direction at p_t above

600 MeV$/c$ the impact-parameter resolution is better than 100 μm for both π and K mesons, which will allow for open charm particle detection (Section 11.9).

Table 11.2: Mass resolution for some decay channels. The natural width is shown for comparison.

Decay channel	σ_m (MeV)	Width (MeV)
$\rho \to e^+e^-$	10	151
$\omega \to e^+e^-$	10	8.4
$\phi \to e^+e^-$	12	4.4
$J/\psi \to e^+e^-$	35	0.09
$\Upsilon \to e^+e^-$	155	0.05
$\phi \to K^+K^-$	2.4	4.4
$K^0_S \to \pi^+\pi^-$	7.8	0
$\Lambda \to p\pi^-$	5.1	0
$D^0 \to K^-\pi^+$	15	0
$D^+ \to K^-\pi^+\pi^+$	12	0

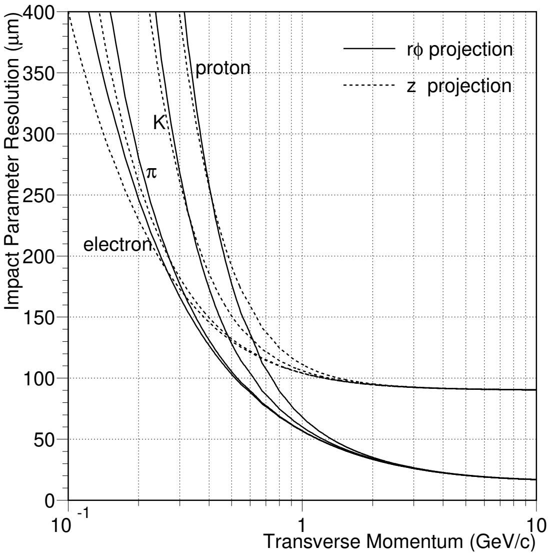

Figure 11.12: Impact parameter precision for electrons, π mesons, K mesons and protons as a function of p_t. The bending and non-bending projections are shown separately.

Stand-alone resolution of ITS and TPC

The momentum resolution for electrons with p_t = 20–100 MeV$/c$ using only the ITS is $\Delta p/p = 7.8\,\%$ independent of p_t. If the ITS should be used as a stand-alone tracker in a wider momentum range, we will have to increase the magnetic field. The estimate for the π meson momentum resolution in this mode ($B = 0.5$ T) is shown in Fig. 11.13.

We have also estimated the precision using only the TPC, at nominal field, and including the corresponding

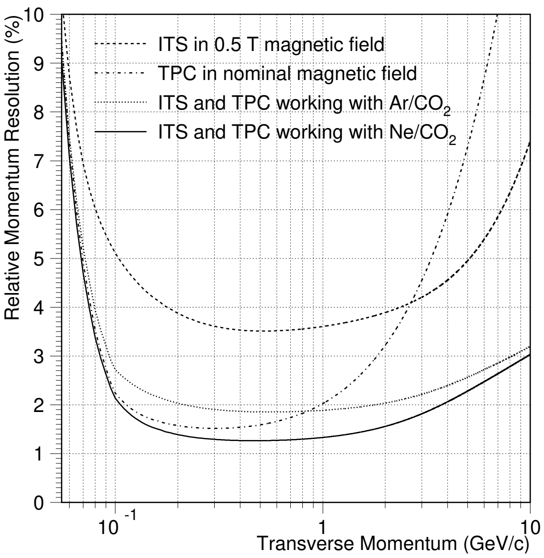

Figure 11.13: Momentum resolution for π mesons as a function of p_t using both ITS and TPC (for two gas mixtures); using *only* the ITS but at a higher magnetic field ($B = 0.5$ T); using *only* the TPC (at nominal field).

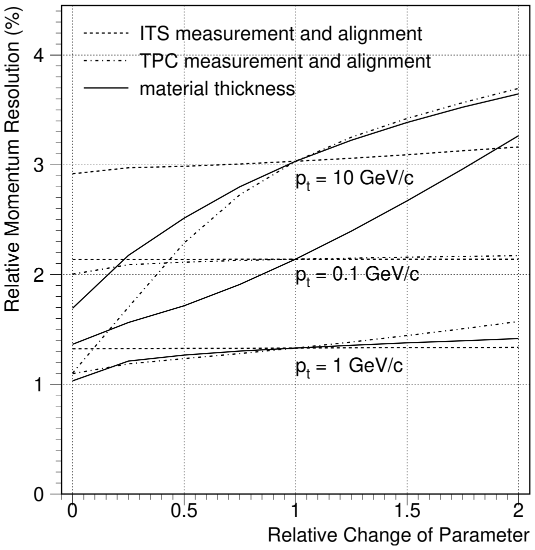

Figure 11.14: Sensitivity of the momentum resolution for pions to a variation of (i) the measurement/alignment precision for ITS and TPC and (ii) the material thickness.

vertex resolution. When only one of the two tracking detectors is used, the resolutions are significantly worse than for the combined system.

Figure 11.14 summarizes the influence of variations of the measurement/alignment precision and the material thicknesses (excluding gases) on the momentum resolution. Unlike for track finding, the resolution of the tracking system is clearly dependent on material thickness and resolution in the TPC, but less so in the ITS.

11.4 Particle identification

As mentioned in the introduction to Chapter 4, particle identification (PID) is an essential asset of ALICE, and is required or at least beneficial for the largest part of our physics programme [3]. Over the large acceptance in the barrel region, PID is achieved with dE/dx measurements in ITS and TPC at low momenta [4] (0.1–0.7 GeV/c), and with TOF at higher momenta (0.5–2 GeV/c). A specialized PID array (HMPID) increases the accessible range to 3.4 GeV/c.

In this section, we discuss the PID performance of these detectors, both in terms of efficiency for positive identification and in terms of remaining contamination by other particle species. At the end, we mention a few examples of physics observables which make use of PID and are not covered in other parts of this document.

11.4.1 PID with ITS, TPC and TOF

dE/dx in the ITS

The measurement of the energy loss in thin silicon detectors can be used for particle identification in the non-relativistic ($1/\beta^2$) region. Together with tracking and momentum reconstruction, the analogue pulse-height measurement in the outer four layers makes the ITS an independent particle spectrometer for low-momentum tracks which do not reach the outer detectors because of decay, energy loss, or curling.

The energy loss in one of the 300 μm silicon layers of the ITS is shown in Fig. 11.15, obtained with a full detector simulation [30] using GEANT and the HIJING generator. The figure corresponds to five central Pb–Pb events with a mean rapidity density of charged particles dN_{ch}/d$y \simeq 6000$ at $y = 0$, and includes conversions and other secondary particle production but no decays.

The mean number of dE/dx measurements useful for PID information is smaller than the number of layers due to overlaps of tracks, inefficiencies in the detection (intrinsic and geometric, $\sim$ 4–5 % per plane) and loss of particles which leave the acceptance. The percentages of tracks with 4, 3, 2 and 1 good measurements are 66 %, 23 %, 8 % and 3 %, respectively. Only tracks with 3 or 4 good hits are considered in the dE/dx analysis, resulting in an efficiency of $\sim$ 90 % before proper PID is applied.

The raw pulse-height information of the silicon detectors has to be corrected for geometrical effects (i.e. path length in the silicon), as well as for detector specific effects. In particular, electronic noise can deteriorate the

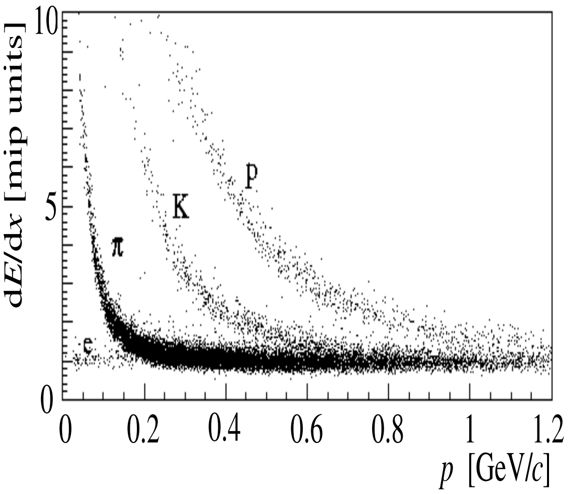

Figure 11.15: Simulated momentum dependence of the energy loss in 300 μm silicon for e, π, K, p.

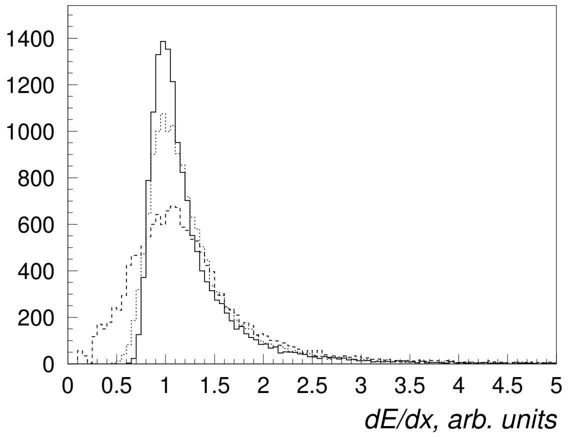

Figure 11.16: Simulated energy loss distribution for particles in the silicon drift detector without noise (full), including noise (dashed), and after applying the correction procedure (dotted).

dE/dx resolution in the drift detectors (and, to a smaller extent the strips), as the charge of a single hit is, in general, collected in a number of anodes. Depending on the drift distance, a variable fraction of the charge can be lost below the noise threshold. This effect, as well as the improvement in resolution resulting from a correction procedure [30], is shown in Fig. 11.16.

The truncated mean dE/dx distribution is shown for a number of cases in Fig. 11.17 for an idealized situation (single particles, no noise). The results of the full simulation described above, and including our expected momentum resolution, are only slightly worse. When averaging the truncated mean [5] of all tracks with $p_{\mathrm{t}} \geq$ 20 MeV/c , the r.m.s. of the fully simulated distribution is 16 % and a Gaussian fit results in $\sigma = 12$ %, compared

[3] Even the photon analysis makes use of p($\overline{\mathrm{p}}$) identification to subtract the $\overline{\mathrm{n}}$ contamination.

[4] Upper limits correspond to 3σ K/π separation.

[5] Two-of-four for four hits and one-of-three for three hits.

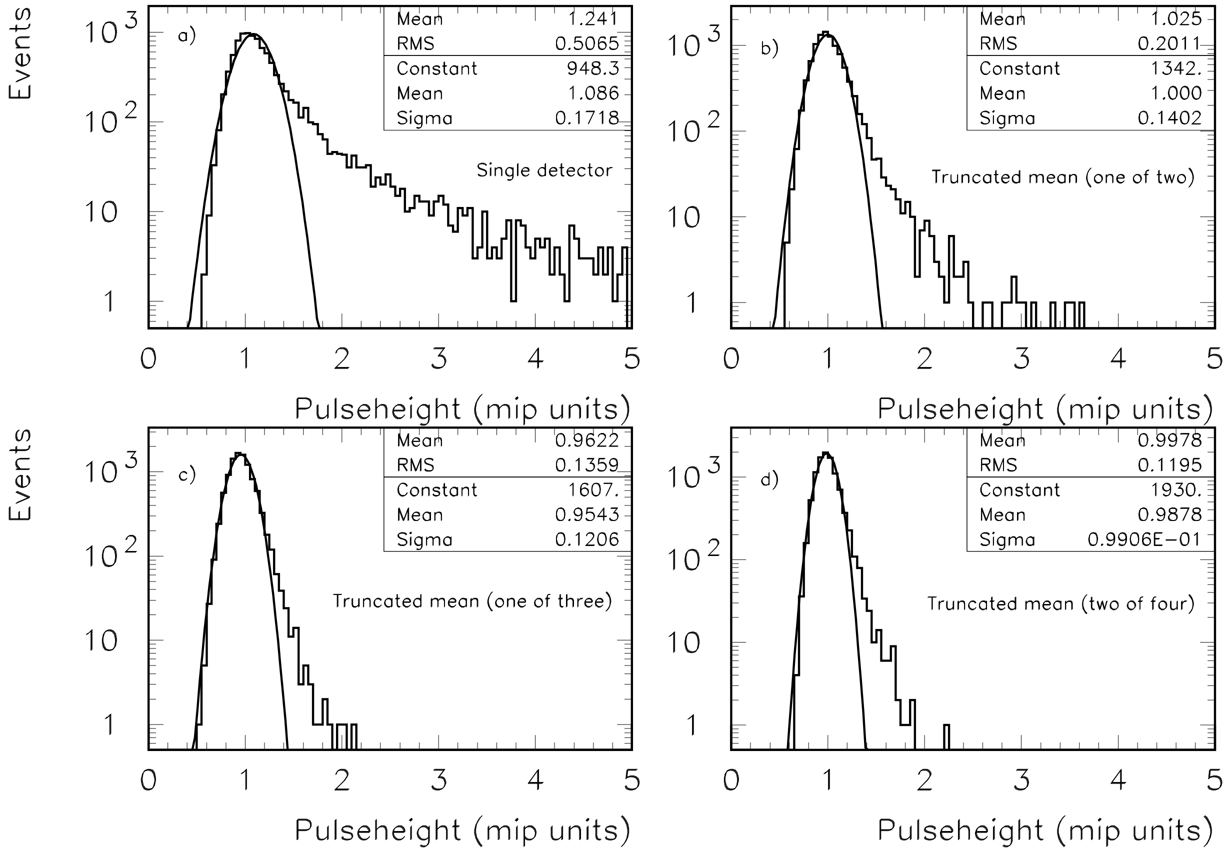

Figure 11.17: Simulated pulse-height distributions in silicon detectors, normalized to 1 at the most probable value (single tracks, no noise). The parameters of a Gaussian fit (full lines) around the most probable value are shown in the plots. a) Single detector response to minimum ionizing particles. b-d) Truncated mean using the *m* lowest of *n* measurements (b: one of two, c: one of three, d: two of four).

to r.m.s. = 12–14 % and $\sigma = 10$–12 % for the ideal case shown in Fig. 11.17.

The upper momentum limits for a rejection factor above 100 between different particle species are listed in Table 11.3 for a cut with 95 % signal efficiency. As can be seen in Fig. 11.15, the particle separation improves rapidly towards lower momenta, e.g. the e/π rejection exceeds 1000 at a momentum of about 130 MeV$/c$.

Table 11.3: Upper momentum limits (MeV$/c$) for a rejection factor above 100 at 95 % efficiency between different particles.

e/π	π/K	K/p
145	540	890

dE/dx in the TPC

The dE/dx performance of the TPC has been described in Section 3.11. We expect a dE/dx resolution for the truncated mean of $\sigma = 5.5$ % for single tracks and of $\sigma \leq 7$ % for high-multiplicity events. The response of the TPC has been parametrized in the simulations as a Gaussian with this width (i.e. $\sigma = 7$ %), while taking the mean ionization loss as a function of momentum from LEP measurements. The dE/dx distribution, normalized to 1 for minimum ionizing particles, is shown in Fig. 11.18. The TPC, on its own, provides a 3σ separation for K/π up to 0.7 GeV$/c$ and for K/p up to 1.4 Gev/c.

TOF performance

The performance of a TOF system for particle ID does not only depend on the intrinsic timing resolution of the detector, but on the system resolution. In our design, this includes the following contributions (see also Sections 4.4 and 4.5):

- error in the track length estimate (∼ 30 ps);
- error in the track extrapolation from TPC to TOF (20 ps);
- error due to the momentum resolution (∼ 20 ps);

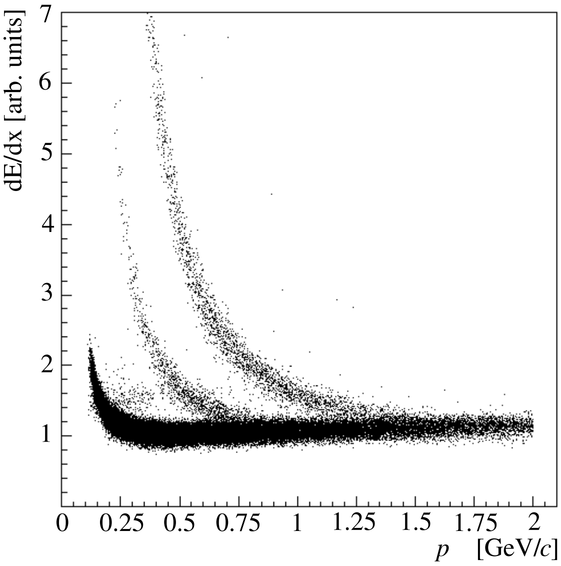

Figure 11.18: Energy loss distribution for particles in the TPC.

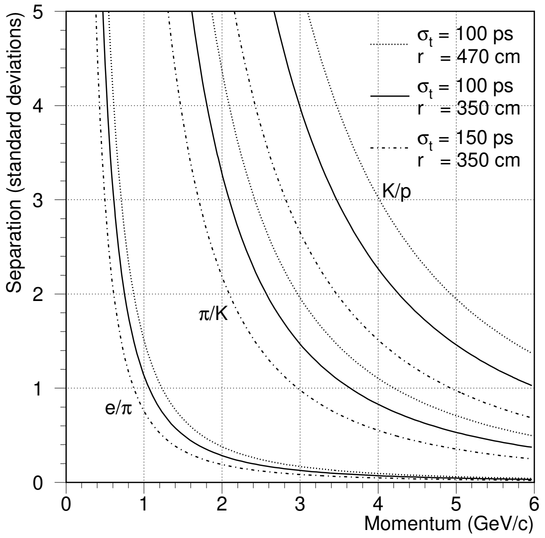

Figure 11.19: Particle separation with TOF.

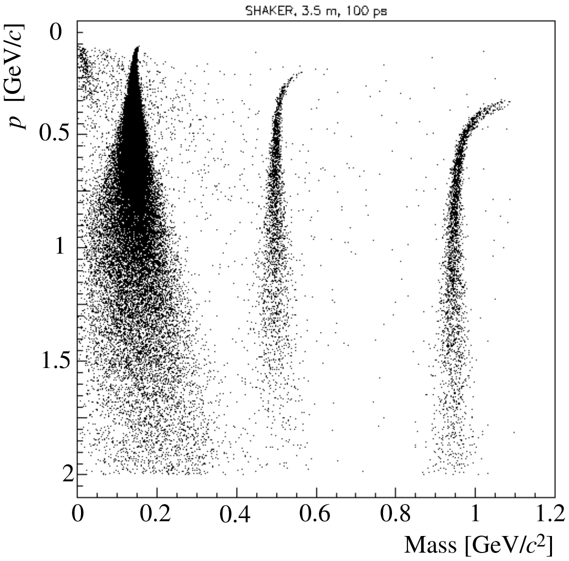

Figure 11.20: Mass separation obtained from the TOF barrel ($\sigma = 100$ ps) for electrons, pions, kaons and protons as a function of momentum.

- digitization error in the TDC (25 ps);
- error in the start time signal (TTC) (50 ps).

The different TOF techniques which have been investigated in ALICE should reach intrinsic resolutions of ~ 50 ps for the Pestov counter, ~ 80 ps for the scintillator and ~ 140 ps for the two-layer PPC array [6]. Including the other sources of timing error, we expect system resolutions of ~ 90 ps (Pestov), ~ 110 ps (scintillator) and ~ 160 ps (PPC), respectively.

The nominal performance of a TOF array with a system resolution of 100 ps or 150 ps, placed at $r = 3.5$ m from the vertex, is shown in Fig. 11.19 for particles at $\theta = 90°$. For comparison, we have also included a TOF with 100 ps at $r = 4.7$ m, corresponding to the location of the HMPID [7].

The PID efficiency and the contamination in the spectra have been evaluated in combination with tracking in the TPC in a full detector simulation. All processes (decays, secondary interactions, energy loss, multiple scattering, etc.) were simulated with GEANT for central Pb–Pb SHAKER events. Included are further TPC tracking efficiency, matching of the extrapolated track with a hit in the TOF array, intrinsic efficiency of the TOF counters (~ 95 %) and losses due to double hits in the TOF (~ 10 % at $dN_{ch}/dy = 8000$).

Figure 11.20 shows the mass as calculated from TOF as a function of the momentum measured in the TPC for each individual track. The electron, pion, kaon and proton bands can be clearly identified. The diffuse noise visible outside these bands arises from wrongly associated hits, non-vertex tracks, decays and other sources of background. The curvature at low p_t reflects the energy loss of particles in the material in front of the TOF.

For both resolutions considered (full line: 100 ps; dashed line: 150 ps), the PID efficiency, including all losses mentioned above, is shown in the upper part of Fig. 11.21, and the contamination by other particle species or noise in the lower part. The drop in efficiency at low momentum is due mainly to particles that do not reach the TOF barrel because of curling in the magnetic field, decay,

[6] The timing resolution for the double-layer PPC array is better than the intrinsic resolution of an individual PPC cell by about $\sqrt{2}$ for single particles and by ~ 25 % in high-multiplicity events.

[7] Equivalent to a TOF resolution of 75 ps at 3.5 m.

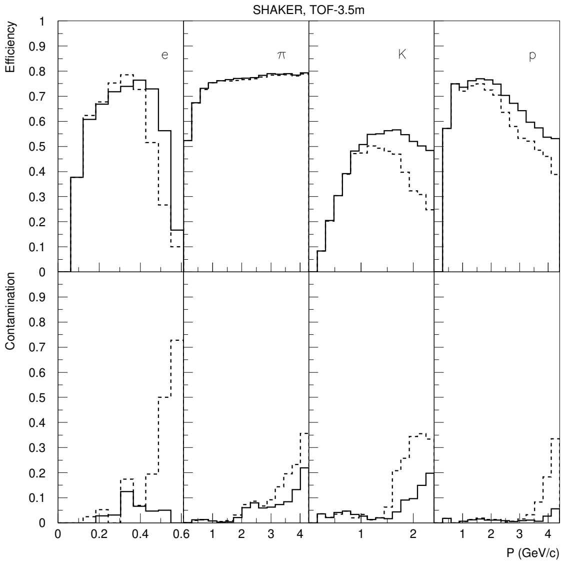

Figure 11.21: Efficiency of identification (upper part) and corresponding contamination (lower part) for different species as a function of momentum for 100 ps (solid line) and 150 ps (dashed line) TOF resolution.

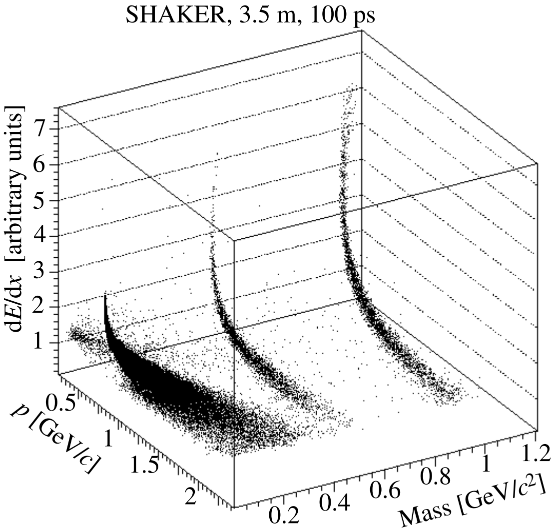

Figure 11.22: Simulated energy loss distribution in the TPC versus mass and momentum, as calculated from TOF.

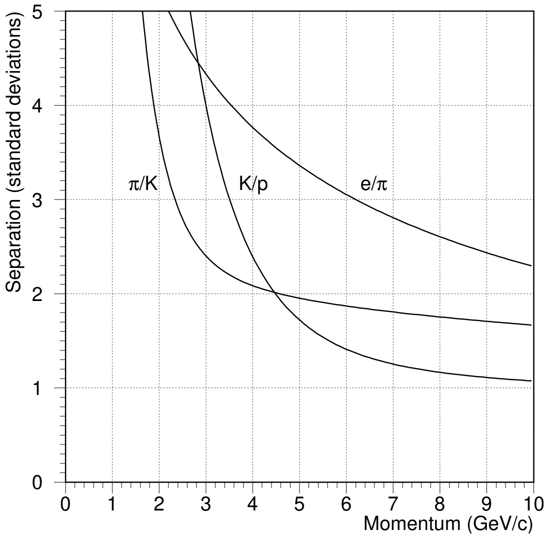

Figure 11.23: Combined particle separation with TPC ($\sigma = 7\,\%$) and TOF ($\sigma = 100$ ps) at $r = 3.5$ m.

or energy loss. In this momentum range, ITS and TPC will be more efficient for PID because they are located at smaller radii. The loss of efficiency at higher momenta is due to the specific set of cuts applied in this case, which places emphasis on low contamination at the expense of efficiency.

In Table 11.4 we list the corresponding upper momentum limits for less than 10 % and less than < 20 % contamination for kaons and protons. These values reflect the particle yields of pions, kaons and protons as generated by SHAKER (π:K:p = 1:0.1:0.06 for $|\eta| < 0.9$, integrated over p_t) and are thus model dependent. For comparison we show the nominal values corresponding to 3σ separation, which are taken from Fig. 11.19.

Combined PID performance

The combined information from dE/dx in the TPC and mass as calculated with TOF is shown in Fig. 11.22. The separation (in standard deviations) expected when using both TPC and TOF is shown in Fig. 11.23 for $\sigma_{\text{TOF}} =$ 100 ps. There is little or no additional gain in rejection at the 3σ separation level (~ 200 MeV/c for K/π, ~ 0 for K/p). However, the TPC information might still be helpful to test the consistency of a PID assignment in the TOF and reduce the background of noise or mis-measured particles.

So far, we have not made full use of the joint performance of ITS, TPC and TOF for hadron ID, as TOF alone already provides a large discrimination power for $p_t > 200$ MeV/c. Currently, dE/dx identification in the ITS and TPC is only used in the electron analysis, described in Section 11.6.

In the future, we intend to also use these detectors for hadron ID, in particular at low momenta and for kaons which decay before reaching the TOF array. These two detectors should compensate a large fraction of the efficiency loss below 800 MeV/c visible in Fig. 11.21, thus increasing the statistics available for event-by-event analysis. The ITS will give access to pion spectra at very low momenta, between 50 and 100 MeV/c, which are of interest in studying collective phenomena (as low momenta

Table 11.4: Momentum limits in GeV/c for 3σ separation, less than 10 % and less than 20 % contamination, for the barrel PID with two time resolutions and for the two HMPID options.

	Barrel PID						HMPID					
	100 ps			150 ps			RICH			Scintillator		
	3σ	10 %	20 %	3σ	10 %	20 %	3σ	10 %	20 %	3σ	10 %	20 %
π/K	2.1	2.0	2.3	1.7	1.6	1.7	3.4	3.5	3.7	2.4	2.6	2.8
K/p	3.5	4.5	4.7	2.8	3.6	3.8	6.5	6.7	7.0	4.0	5.0	5.3

are related by the uncertainty principle to a large length-scale) or coherent pion emission from chiral condensates (see Section 11.11).

11.4.2 HMPID performance

The HMPID RICH detector was simulated using the performance figures achieved in the pion/proton beam tests described in Section 4.6.

The simulation was fully analyzed with a pattern recognition program (see Section 4.6) applying the same criteria as in the TOF analysis, i.e. minimizing the contamination at the expense of efficiency. For the HMPID, efficiency is less crucial than for the barrel PID as it will be used to measure inclusive spectra only.

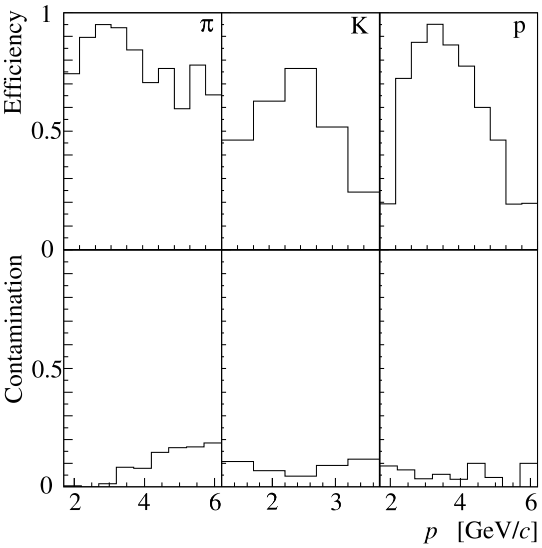

Figure 11.24: Efficiency of identification (upper part) and corresponding contamination (lower part) for pions, kaons and protons in the HMPID.

Figure 11.24 shows the efficiency and contamination for pions, kaons and protons in the high momentum range. The results pertain to a RICH with asymptotic rings of 15 cm radius and for particles at normal incidence. The 3σ limit and the momenta corresponding to 10 % and 20 % contamination, respectively, are listed in Table 11.4.

Calculations have also been performed for a high-momentum PID array of TOF counters with an overall resolution of 100 ps. The results show that identification of kaons and protons up to 2.8 GeV/c and 5.3 GeV/c, respectively can be achieved (see Table 11.4.).

11.4.3 Examples for physics with PID

We will now give some examples of physics observables requiring hadron identification; others can be found in subsequent sections of this chapter.

Momentum spectra and transverse expansion

In a simplified description of heavy-ion reactions, the temperature at freeze-out of a hadronic system in thermal equilibrium could be directly measured via the slope of the p_t spectra. However, the actual momentum distribution will be significantly distorted because of resonance decays —mainly at low p_t — and by transverse flow, which changes the slope at high momenta [8].

It has been suggested [31] that it is possible to determine both the flow velocity and the temperature of the system at freeze-out by measuring the transverse mass slopes simultaneously for several particle species in a sufficiently large range of m_t.

In ALICE, we will be able to measure inclusively a large number of stable particles and resonances, and a number of them event-by-event (π, K, p, possibly Λ). The p_t spectra of identified $\pi^{\pm}$, $K^{\pm}$ and p($\overline{\rm p}$) (both charges added), as measured for a single central Pb–Pb SHAKER event [9] in ALICE, are shown in Fig. 11.25. The simulation includes decays and all other sources of inefficiency; the number of particles generated and identified within the acceptance of ALICE is listed in Table 11.5. Both the statistics in single events and the momentum range covered by PID are sufficient to study not only 'average' slopes, discussed in Section 11.11, but also the shape of the distribution up to about 2 GeV/c (see also Section 11.1.2).

[8] Above ~ 2 GeV/c, hard scattering will also have to be taken into account.

[9] The standard SHAKER ratio of K/π = 0.1 has been doubled in this simulation to the thermal equilibrium value expected for $T = 150$ MeV.

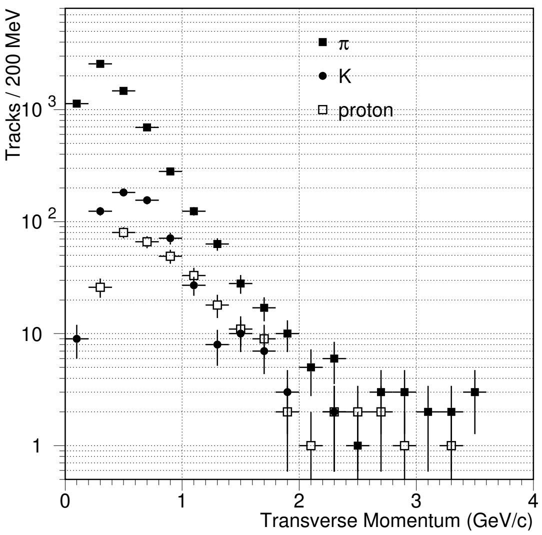

Figure 11.25: Momentum distribution for a single event (SHAKER, dN_{ch}/dy = 8000) after tracking and PID reconstruction.

Table 11.5: Number of charged pions, kaons and protons generated with SHAKER (K/π = 0.2) and identified in ALICE.

	π	K	p
generated	9000	1800	480
identified	6500	530	300

Particle ratios

The abundance of different particle species can be a measure of the degree of chemical equilibrium at freeze-out (see Section 11.1.3). As an example, we will discuss K/π ratios measured event-by-event. This ratio (or more generally the strangeness content of the event) could be sensitive to the presence of a phase transition or to the chemical equilibrium reached in the hadronic phase (depending on the evolution scenario).

The accuracy of single-event observables will decrease typically with the square root of the number of measured particles. Taking the number of identified particles, as listed in Table 11.5, a K/π ratio of order 0.2 can be measured event-by-event with a typical error of 4–5 % in ALICE.

While the K/π ratio in $p\bar{p}$ reactions at the Tevatron reaches ~ 0.15 at high multiplicity, the equilibrium ratio in a hadron gas will be ~ 0.2. Given our accuracy, this difference corresponds to 5–6 σ, and therefore events could be ordered into different distinct classes according to their strangeness content and analysed for correlated signals.

In a hadronic gas at equilibrium, the K/π ratio will depend on the freeze-out temperature [32]. (K/π = 0.203, 0.214, 0.221 for T = 140, 150, 160 MeV, respectively). This temperature, which is expected to be between 150 and 200 MeV, could be measured via the K/π ratio to within about 10 MeV, event-by-event.

φ mesons

The identification of kaons, in a large momentum interval, is also important for the identification of φ mesons through their decay into two kaons. The measurement of the yield of the φ (a $s\bar{s}$ pair) sets more stringent constraints on the origin of the observed flavour composition compared to the K/π ratio. Shifts of the φ meson mass and changes in its line-shape in a hot hadronic medium have been predicted [33] and are possibly observed in Si–Au collisions at the AGS [34]. The production rate and the partial decay width into kaons may change as a consequence of partial chiral symmetry restoration [35,36]. The φ decay will be observed in ALICE both in $\phi \to e^+e^-$ and in $\phi \to K^+K^-$, and the ratio of these two channels might be very sensitive to changes in parton or kaon masses. In addition, this ratio is independent of production cross-sections.

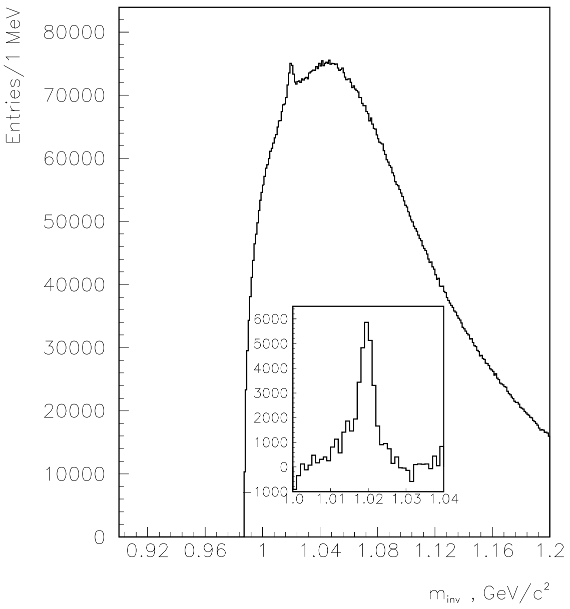

Figure 11.26: The effective mass spectrum of K^+K^- pairs for 5000 SHAKER events. The insert shows an enlarged view of the φ resonance.

We have simulated the expected experimental spectrum of K^+K^- pairs using standard SHAKER events (K/π = 0.1) and by applying cuts on the angle between the two kaons ($\cos\theta \geq 0.84$) and on the transverse momentum ($p_t \geq 0.5$ GeV$/c$). The effective mass spectrum, including experimental resolution ($\sigma = 2.4$ MeV) and natural width ($\Gamma = 4.4$ MeV), is shown in Fig. 11.26. Our resolution is of the order of the natural width and will allow any change in position or shape of the resonance to be observed.

The signal-to-background ratio (S/B) is 0.07 in a ± 3 MeV bin around the central mass, and the significance is $S/\sqrt{B} \approx 27$ for 5000 events (2 minutes data taking).

Jet quenching

The energy loss of partons in a strongly interacting medium is different in hadronic matter than in a QGP. This 'jet quenching' will change the yield and p_t slope of hard 'leading' particles (see Section 11.1.2).

It has been suggested that the difference between p and $\overline{\text{p}}$ production at high $p_t > 3$ GeV$/c$ is particularly sensitive to this effect [37]. At high p_t (larger Bjorken x), valence quarks contribute more to the cross-section than sea quarks and a quark jet is more likely to produce a leading proton than an antiproton, while a gluon jet should have the same probability for both. Therefore protons and antiprotons can be used to some extent to tag quark and gluon jets, respectively.

The difference in the p and $\overline{\text{p}}$ spectra, as predicted by HIJING and including jet quenching, is shown in Fig. 11.27. The spectrum, corresponding to 5×10^4 Pb–Pb events, contains 10 000 protons and 3500 antiprotons identified in the HMPID (including all efficiencies) in the momentum range 3–6 GeV/c.

Owing to their larger colour charge, gluons should lose more energy in a medium than quarks and therefore one should see a stronger suppression of antiprotons at large p_t due to jet quenching. The effect can be measured by comparing p and $\overline{\text{p}}$ spectra in pp reactions, where no quenching is expected, to the same ratio in heavy-ion collisions. The difference is predicted to be strong, a factor of two in the yield ratio p/$\overline{\text{p}}$ between pp and Pb–Pb above 3 GeV/c, and can be easily observed in the HMPID.

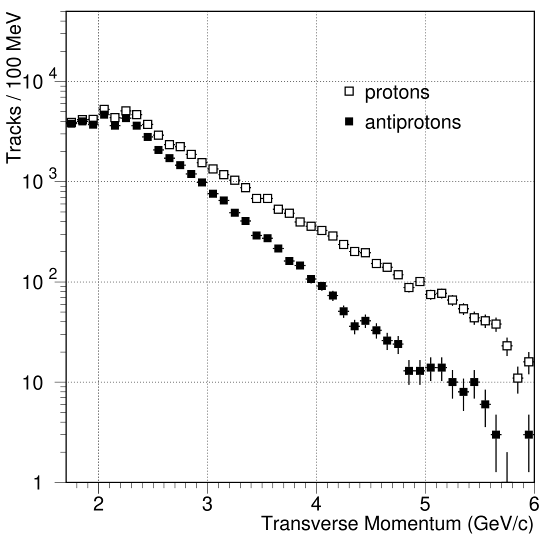

Figure 11.27: Momentum distribution of protons and antiprotons identified in the HMPID for 5×10^4 HIJING events.

The very high multiplicity at the LHC will enable us to study pion correlations in a single event with good accuracy. The information on the space–time evolution for every single event can then be correlated with other characteristics of the same event (see Section 11.11).

The good particle identification and the large acceptance in ALICE will allow correlation studies for a number of combinations, i.e.:

- identical bosons (pions and kaons),
- identical fermions (protons, possibly Λs), and
- non-identical particles.

Identical bosons and fermions will be influenced by Bose–Einstein and Fermi–Dirac statistics, respectively, and by Coulomb and nuclear final-state interactions.

The correction of the experimental correlation functions for the Coulomb interaction presents an important difficulty at very high particle density because of multiparticle interactions [48, 49]. The correction method can be based partially on measured correlations of oppositely charged particles (pions and kaons), which are mainly influenced by the Coulomb effect [50]. Note that the K^0_S particle will not experience the Coulomb interaction.

The two-proton correlations are dominated by Coulomb and strong interactions. The intensity of the enhancement in the correlation function around $\simeq 20$ MeV/c, which arises due to strong interaction, is related to the inverse effective source size [51].

It appears that the correlations of non-identical particles are directly sensitive to the time of particle emission and thus can complement in an important way the standard interferometry measurements (see below).

11.5 Particle correlations

11.5.1 Introduction

The study of particle interferometry —or of particle correlations in general— is of great interest in heavy-ion reactions because in expanding systems correlations contain information about the dynamical evolution, e.g. proper *time of decoupling* ('freeze-out'), *duration* of particle emission for different particle species, and the presence of *collective flows* [38–47]. The interferometry measurements are of particular importance because the decoupling time and intensity of transversal flows should be closely related to the QGP formation and to the latent heat of the phase transition [38, 41, 42]. Good particle identification and large acceptance (both in rapidity and transverse momentum) are necessary to perform such studies. The possibility to study correlations in a wide range of transverse momenta (up to 3 GeV/c) is of particular importance.

11.5.2 Detector performance

The correlation measurements are strongly dependent on a number of experimental characteristics, in particular *acceptance, PID, momentum resolution* and *two-track separation* [10]. The ALICE experiment is well-suited for correlation studies in all respects, and in the following its performance will be illustrated with a number of examples.

Acceptance and particle identification

The ALICE experiment has a full coverage in the azimuthal angle, and a large acceptance in rapidity. The acceptance in the transverse momentum for different particle species is determined by the particle identification performance.

A number of aspects in correlation studies require such a large acceptance:

The maximal measurable Q_L, which has to be larger than the correlation width, depends on the available rapidity interval. To perform an unbiased study of Q_L —the longitudinal component of the correlation function— for a wide set of different colliding systems, including proton–proton collisions, it is necessary that the acceptance in rapidity be sufficiently wide. Likewise, the acceptance in azimuth and p_t are relevant for the sideward and outward components, respectively (Q_{Tside}, Q_{Tout}).

The acceptance of the HMPID detector was optimized to accommodate the wide correlation functions which are expected at large transverse momentum (e.g. corresponding to effective sizes of only 1 fm/c). Figure 11.28 shows the two-kaon correlation function in the interval $2 < p_t < 3$ GeV/c for 2×10^6 Pb–Pb SHAKER events. A Gaussian shape was used for the correlation function with the intercept $\lambda = 0.5$ and with all three sizes equal to $R = 1$ fm.

A large acceptance also improves the correlation studies for particles of opposite charge or different mass (which populate different parts of the detector even for identical momenta) as well as for single-event pion interferometry.

The particle identification performance has been discussed in Section 11.4. Charged kaons will be identified for $p_t < 2$ GeV/c in the full azimuthal acceptance, and up to 3.5 GeV/c in the restricted acceptance of the HMPID detector. The particle identification in the HMPID detector will serve for the analysis of semi-hard and early-time phenomena, which are expected to leave signatures at high transverse momentum. Protons will be identified in the momentum range $0.3 < p_t < 4.5$ GeV/c.

K^0_S particles will be identified from the V^0 decays, and one can expect about 60 measured K^0_S particles per event.

[10] The two-track separation is less relevant for correlations of non-identical particles.

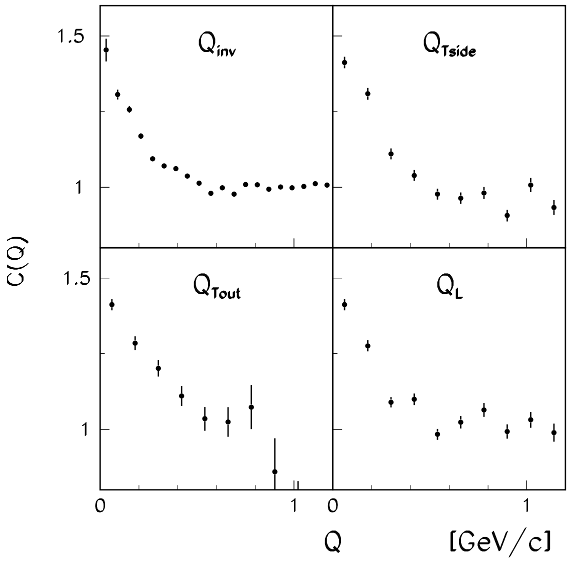

Figure 11.28: Two-kaon correlation function in the interval $2 < p_t < 3$ GeV/c, for kaons identified in the HMPID detector. The sample consists of 2×10^6 Pb–Pb SHAKER events, $\lambda = 0.5$, and $R = 1$ fm.

Two-particle momentum resolution and two-track separation

The resolution of the interferometric two-particle momentum difference variables Q_{Tside}, Q_{Tout} and Q_L was evaluated for pairs of charged particles, applying the same method that was used in Section 11.3, but in addition taking into account:

- the loss of overlapping clusters in the tracking system;
- the inhomogeneous distribution of material in the detector;
- the loss of tracks corresponding to double-hits in the PID barrel.

A partial overlap of tracks in the tracking system causes a loss in momentum resolution. This effect is particularly important for the studies of correlations of identical particles, because two tracks with a very small relative momentum will have a considerable overlap in space.

Figure 11.29 shows the effect of the experimental momentum resolution —taking track overlap into account— as the difference ΔQ between the true and the measured relative momentum components in the region of the Bose–Einstein correlation peak (for pion pairs with $Q_{inv} < 50$ MeV/c). Figure 11.30 shows the absolute value of the momentum error as a function of Q_{inv}; each entry in the scatter plot represents one particle pair. The corresponding r.m.s. values of the resolution functions are listed in Table 11.6.

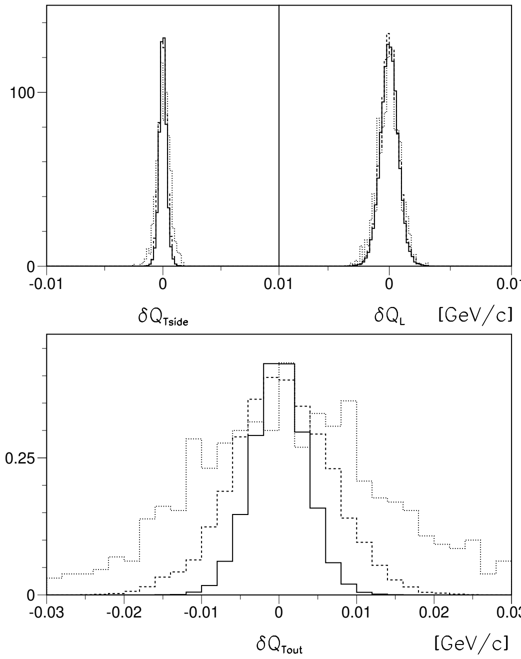

Figure 11.29: Momentum resolution for pions of the individual interferometric momentum components (Q_{Tside}, Q_{Tout} and Q_{L}) for three intervals in transverse momentum: $0.1 < p_{\mathrm{t}} < 0.3$ GeV/c (full lines), $0.3 < p_{\mathrm{t}} < 0.6$ GeV/c (dashed lines), and $p_{\mathrm{t}} > 0.6$ GeV/c (dotted lines).

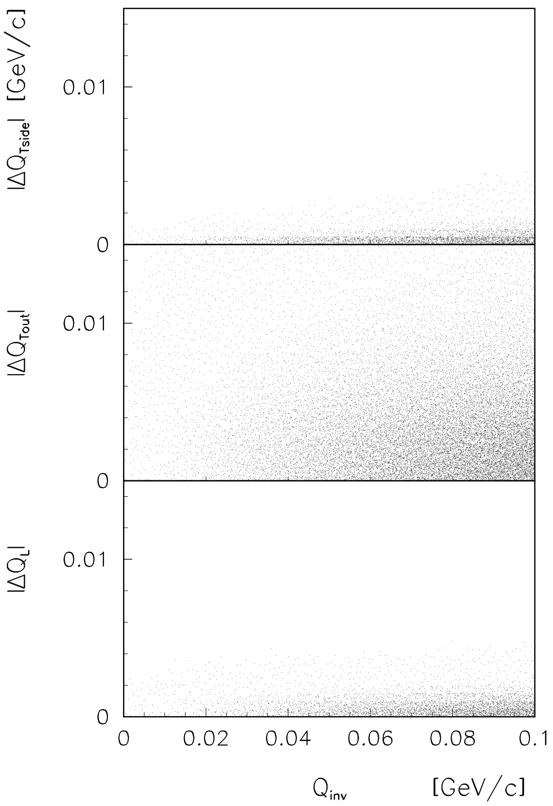

Figure 11.30: Measurement error in GeV/c of the relative momentum components for individual particle pairs as a function of Q_{inv}.

Table 11.6: Resolution values (in MeV/c) for $\pi^+\pi^+$ pairs with $Q_{\mathrm{inv}} < 50$ MeV/c.

p_{t} range (MeV/c)	Resolution (r.m.s.)		
	Q_{Tside}	Q_{Tout}	Q_{L}
$100 < p_{\mathrm{t}} < 300$	0.4	3.4	0.8
$300 < p_{\mathrm{t}} < 600$	0.4	6.4	0.8
$600 < p_{\mathrm{t}}$	0.5	11.0	1.0

Table 11.7: Resolution values (in MeV/c) for pairs of different particles with $Q_{\mathrm{inv}} < 50$ MeV/c.

	Resolution (r.m.s.)			
	Q_{Tside}	Q_{Tout}	Q_L	Q_{inv}
$\pi^+\pi^+$	0.4	3.8	0.8	1.3
K^+K^+	0.5	9.5	2.3	4.2
$K^0_S K^0_S$	3.9	7.8	3.4	4.0
pp	0.7	13.0	4.3	8.0
π^+K^-	0.7	4.1	1.1	1.9
π^+p	0.7	4.2	1.2	2.1
K^+p	1.0	8.3	3.2	4.8

The r.m.s. values of the resolution functions for different particles, without a p_{t} cut, are listed in Table 11.7.

For a small relative momentum of a particle pair, the resolution for Q_{Tout} is mostly determined by the resolution of the track curvature, while for the other two components, Q_{Tside} and Q_{L}, it is mostly determined by the angular resolution, which is very good in ALICE.

The maximal expected effective source size at the LHC is about 20 fm, corresponding to a correlation function with a r.m.s. width of 10 MeV/c. Compared to this width, the resolution for Q_{Tside} and for Q_{L} is excellent, and still acceptable for Q_{Tout}. We should be able not only to estimate the width of the correlation function, up to radii of $\sim$30–40 fm, but also to study in detail its shape, in the whole transverse-momentum range of interest.

The effect of the actual distribution of material in our detector, which is non-uniform, has been found to be small.

Figure 11.31 shows the fraction of lost hits for close track pairs. Hits were assumed to be lost when their distance was smaller than 1σ of the local cluster size [52]. Only pairs of a small relative momentum suffer from con-

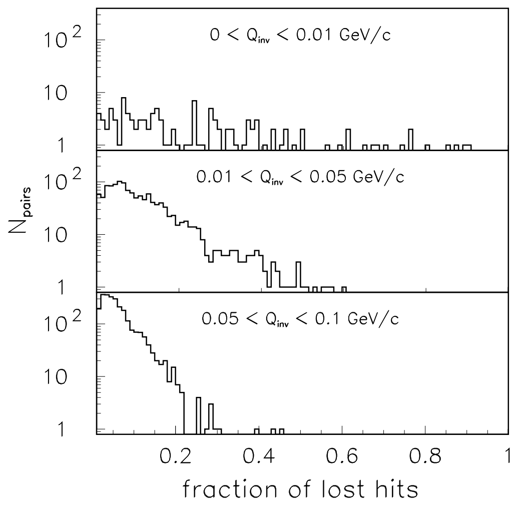

Figure 11.31: Fraction of lost hits along the tracks for close pairs, in three intervals of Q_{inv}.

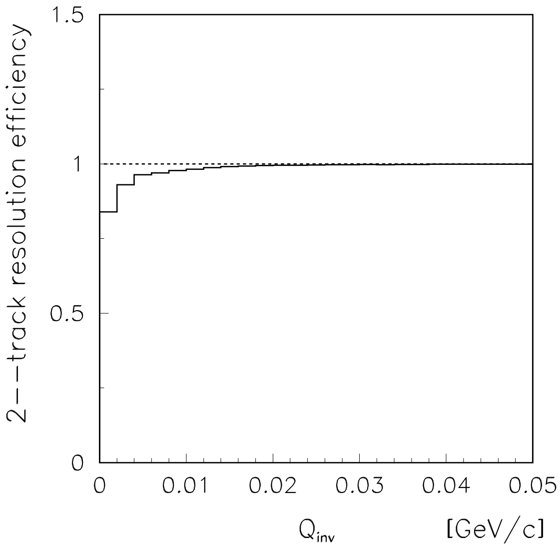

Figure 11.32: Two-track efficiency for pairs of charged kaons in the PID barrel.

siderable overlap.

The efficiency of identifying two close tracks in the PID barrel is presented in Fig. 11.32 for charged kaons as a function of Q_{inv}. The fine granularity of the PID detector provides a good two-track separation down to very small values of Q_{inv}.

11.5.3 Single-event pion interferometry

In general, the statistical accuracy $\delta R/R$ for the measured effective sizes (R_{Tside}, R_{Tout} and R_{L}) depends on the following parameters:

- on the number of measured events ($\delta R/R \sim (\sqrt{N_{\mathrm{events}}})^{-1}$) ;
- on the phase space density ($\delta R/R \sim (\mathrm{d}N/\mathrm{d}y)^{-1}$) ;
- on the acceptance ($\delta R/R \sim (\sqrt{\Delta y \Delta \varphi})^{-1}$) ;
- on the track reconstruction efficiency;
- on the effective sizes;
- on the correlation intensity.

The dependence of the accuracy on all listed parameters was studied by a special simulation procedure, described in Ref. 53, for a rapidity density of charged particles $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$. The shape of the simulated correlation function was taken to be Gaussian in all three components with $\lambda = 0.5$ and $R = R_{\mathrm{Tside}} = R_{\mathrm{Tout}} = R_{\mathrm{L}}$. The simulated correlation functions, for different effective sizes and for samples consisting of different numbers of events, were finally fitted in the three-dimensional relative momentum space by Gaussian functions. The results for two effective sizes are shown in Fig. 11.33 as a function of the number of analysed events.

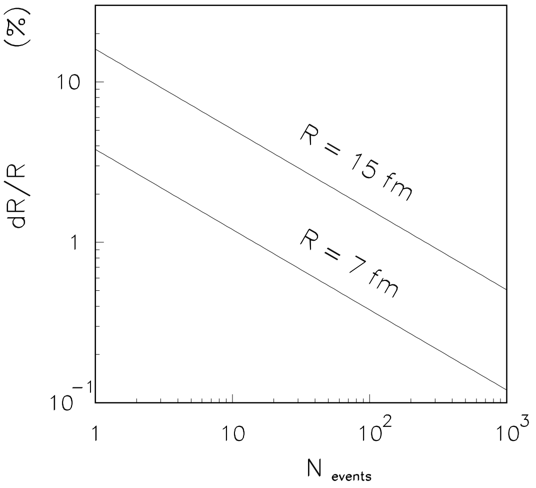

Figure 11.33: Measurement accuracy for the effective interferometric sizes $R = 7$ fm and $R = 15$ fm as a function of the number of analysed events.

Single-event pion interferometry, i.e. one analysed event, is feasible with a relative error of $\sim 20\,\%$ up to effective sizes of 15 fm at the highest particle density. Since

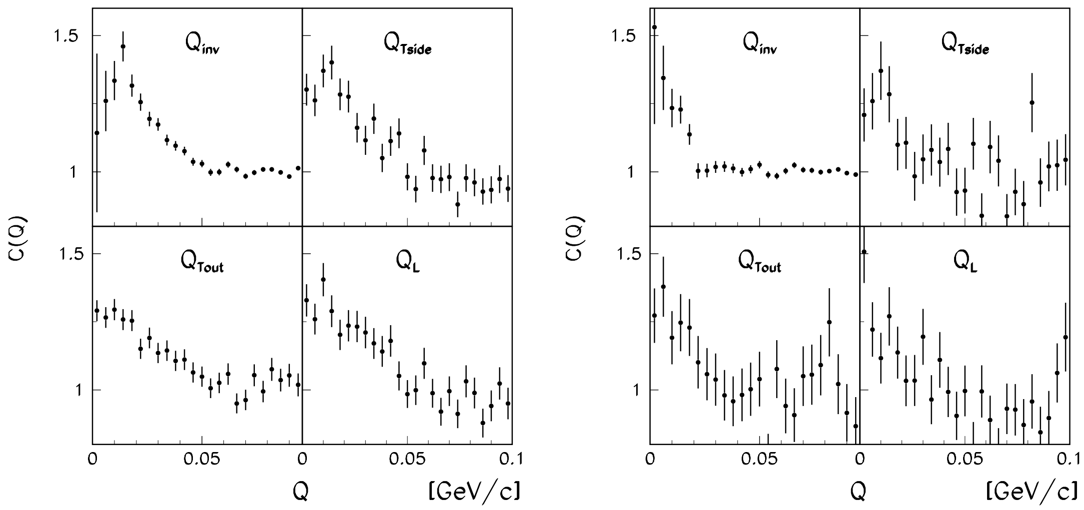

Figure 11.34: Two examples for the simulated single-event correlation functions with the three effective sizes set to $R = 7$ fm (left) and $R = 15$ fm (right).

pions of both charges are available, two independent measurements may be done for each event, thus reducing the statistical errors.

Two examples of the single-event correlation functions are shown in Fig. 11.34 for $R = 7$ fm and $R = 15$ fm ($\lambda = 0.5$, $R = R_{\mathrm{Tside}} = R_{\mathrm{Tout}} = R_{\mathrm{L}}$).

11.5.4 Correlations of non-identical particles

Identical-particle interferometry yields information on the *relative* space–time distances between the emission points of particles. In order to get *absolute* values, such as the decoupling proper time or the duration of particle emission, model-dependent assumptions have to be made. In contrast, the correlations of non-identical particles appear to be directly sensitive to the *time difference* or 'delay' between the emission of two different particle types [54].

It has been shown recently that correlations —induced by Coulomb and nuclear final-state interactions— of two non-identical particles of types 1 (e.g. pions) and 2 (e.g. kaons) allow one to measure directly the difference of the mean emission times $\langle \Delta t \rangle = \langle t_1 - t_2 \rangle$, including its sign [55]. This way one can determine in a model-independent way which particle types ($\mathrm{K}^+, \mathrm{K}^-, \pi^+, \pi^-, \mathrm{p}\ldots$) are emitted in the early or later stages of the freeze-out. Note that kaons, for example, are expected to decouple earlier than pions because of their larger mean free path.

To study the sequence of emission, correlation functions of two non-identical particles should be constructed separately for angles smaller and larger than 90° between the momentum $\mathbf{k}^* = \mathbf{p}_1^* = \mathbf{Q}/2$ and the total pair velocity $\mathbf{v}$. The corresponding correlation functions are denoted as $R_+(k^*)$ and $R_-(k^*)$. Depending on $\langle \Delta t \rangle$, their ratio R_+/R_- should show a peak or a dip in the region of small k^* and approach 1 both at $k^* \to 0$ and $k^* \to \infty$.

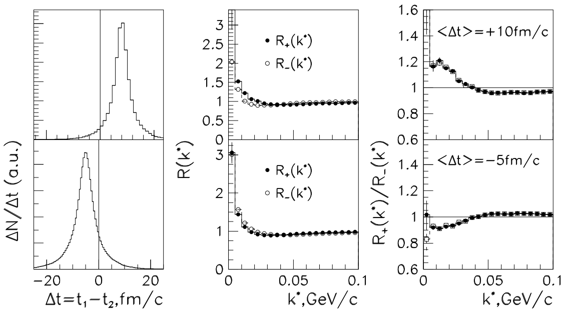

Figure 11.35: Left: Time difference between K^+ and K^- emission ($\Delta t = t_1 - t_2$) simulated by VENUS with $\langle \Delta t \rangle = +10, -5$ fm/c introduced *ad hoc*. Middle: Corresponding correlation functions R_+ (for $\mathbf{v} \cdot \mathbf{k}^* \geq 0$) and R_- (for $\mathbf{v} \cdot \mathbf{k}^* < 0$) and their ratios (right) calculated for $\mathrm{K}^+\mathrm{K}^-$ pairs. Positive values of $\langle \Delta t \rangle$ means that K^+ are, on average, emitted after K^-. Ratios corrected for experimental resolution are represented by open symbols.

As shown in Figs. 11.35 and 11.36, the ratio R_+/R_- can be measured for different pairs of non-identical particles even if $\langle \Delta t \rangle$ is of the order of a few fm/c only. The correlation functions of the pairs $\pi^+\mathrm{K}^+$, $\pi^+\mathrm{K}^-$, $\pi^+\mathrm{p}$, $\pi^-\mathrm{p}$ have been calculated for one simulated event; i.e. they can be studied event-by-event to look for fluctuations arising, for examples, from the phase transition.

The simulations were performed using VENUS (ver-

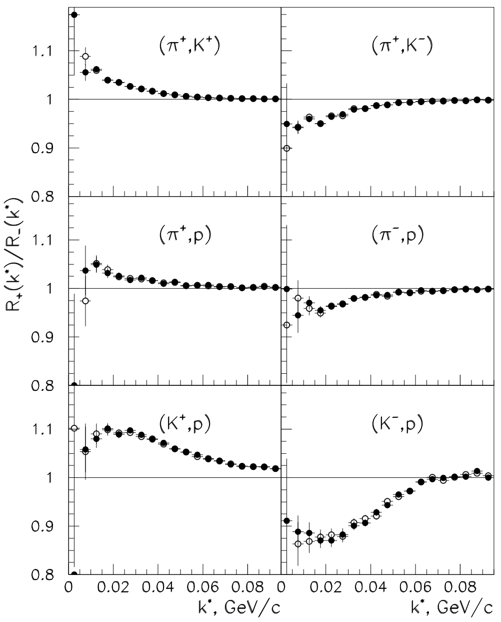

Figure 11.36: The same ratios R_+/R_- as in Fig. 11.35 calculated with $\langle \Delta t \rangle = -10$ fm/c, for different pairs of non-identical particles, taking into account (open symbols) and neglecting (full symbols) the effect of experimental resolution.

sion 5.14, see Section 11.2), which provides space–time coordinates of freeze-out points. The calculation of correlation functions takes into account the effects of quantum statistics and final-state interactions (Coulomb and nuclear) as well as the experimental resolution (see Table 11.7 on page 172).

11.6 Electron pairs

11.6.1 Introduction

Leptons, once produced in the course of a heavy-ion reaction, are not affected by final-state interactions in the subsequent evolution and can therefore relate information about the early phases of the collision. Of particular interest are the leptonic decays of vector mesons. The lighter mesons (ρ, ω, ϕ) are expected to change their properties (mass, width, decay pattern) in dense matter as a consequence of collective interactions ('in-medium modification') [56] and of chiral symmetry restoration in the vicinity of the phase transition [57]. The production of the heavier mesons (J/ψ and above) in heavy-ion collisions will be suppressed with respect to the one in pp reactions in a deconfining medium via Debye screening [58]. In ALICE, we will measure the lighter mesons (ρ, ω, ϕ, J/ψ) in the electron channel with the central detectors, the heavier ones (J/ψ and Υ families) in the muon channel with the forward muon spectrometer.

11.6.2 Detector requirements

ALICE has no detector system dedicated exclusively to the measurement of electrons. Electrons are tracked and identified in the same detectors as hadrons, i.e. the ITS, the TPC, and the PID-array. This approach, dictated largely by financial considerations, is possible only by requiring very stringent performance specifications for all detectors and in addition by restricting the momentum range for electrons to $p < 2.5$ GeV/c. In particular, we need a very high efficiency to track and identify electrons in order to recognize and eliminate the background from Dalitz decays and conversions. The design of the ITS has been optimized in this respect, in particular for low-p_{t} electrons (20 to 100 MeV/c) which do not reach or can not be identified in the outer detectors. In addition, the amount of material in the ITS the and beam pipe are minimized in order to reduce conversions.

For electron identification, the most problematic area is at a momentum of around 200 MeV/c, where the TPC can not discriminate between electrons and pions (see Section 11.4). This area can be bridged from below with the $\mathrm{d}E/\mathrm{d}x$ measurement in the ITS and from above by the TOF, which starts to become effective at a momentum of 140 MeV/c at its nominal radial distance.

The high particle multiplicities at the LHC pose an additional problem to the electron measurement at the mass and p_{t} of interest to us. There will be several hundred electrons per event within the acceptance (mainly π^0 Dalitz decays and conversions), compared to a few from interesting meson decays. Triggering on the signal is impossible, as there are many unlike-sign combinations per event, even in very narrow regions around the meson masses. We therefore need to analyse a large number of events, typically more than 10^7. This makes the electron analysis one of the most demanding ones in terms of statistics and establishes our need to collect of the order of 5×10^7 events per run.

11.6.3 Physics generator

The simulations were performed using the SHAKER event generator, which includes a large number of particle decays relevant to the electron analysis. For the simulations presented in the LoI, we had included the Dalitz decays of π, η, η', ω, and ϕ as background and the leptonic two-body decays of ρ, ω, ϕ and J/ψ as well as the Drell–Yan continuum as sources [59]. For the present calculations, we have added photon conversions in the detector and semileptonic charm decays as additional background [60].

For heavy flavour production we used the model developed for the dimuon spectrometer (see Chapter 6), i.e. PYTHIA and JETSET for production and fragmentation and recent cross-section calculations [61] scaled from pp collisions with $A^2/2$ (the factor two accounts for nuclear effects such as shadowing). This leads to a rate of 6×10^{-3} J/$\psi \rightarrow e^+e^-$ decays and 15 $c\bar{c}$ pairs per unit of rapidity in a central Pb–Pb event.

11.6.4 Detector simulation

The detector response is parametrized in terms of tracking efficiency, PID efficiency, and pion rejection. The tracking efficiency for electrons in the ITS and TPC is 90 % above $p_t = 100$ MeV/c (Section 11.3.2). At lower momenta, tracks are reconstructed in the ITS only and the corresponding efficiency decreases to 50 % at $p_t = 30$ MeV/c (Section 11.3.2, Fig. 11.5 on page 160). The momentum resolution takes into account measurement errors, multiple scattering, energy loss and bremsstrahlung (Section 11.3.3).

Electron identification at very low momenta is achieved via $\mathrm{d}E/\mathrm{d}x$ in the ITS alone ($p < 100$ MeV/c) or in combination with the TPC ($100 < p < 140$ MeV/c) with an efficiency of more than 95 % and a pion rejection factor above 1000 (Section 11.4). In the momentum range $140 < p < 200$ MeV/c the TPC can not distinguish electrons and pions (the energy-loss curves cross at around 140 MeV/c) but the TOF starts to become effective above 160 MeV/c, as the minimum momentum required to reach the TOF at $r = 3.5$ m is about 110 MeV/c. The TOF pion rejection is excellent (> 1000); however, not all particles can be identified in the TOF because of geometrical inefficiency and multiple hits in the counters. For particles without a PID tag in the TOF, $\mathrm{d}E/\mathrm{d}x$ in the ITS can still be used, however with lower efficiency (80 %) and lower rejection (about 10). The pion contamination reaches a significant but still acceptable level of about 20 % in this momentum bin.

Between 200 and about 600 MeV/c, both TPC and TOF are used to identify electrons, and, finally, at higher momenta, only the TPC remains (see Fig. 11.23 on page 167).

11.6.5 Analysis cuts

The invariant-mass spectrum in the acceptance before any cuts is shown in Fig. 11.37 for a rapidity density of $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$ (after rejecting conversion electrons). The signal-to-background ratio will be inversely proportional to the multiplicity density, i.e. a factor of four better for the lowest multiplicity anticipated in central Pb–Pb collisions ($\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 2000$).

In order to improve on the signal-to-background ratio, the electron analysis proceeds in three steps which pose different requirements in terms of electron efficiency and pion rejection.

In the first step, conversions are recognized and removed from the electron sample. By requiring a hit in the first ITS layer (which therefore has to have a high efficiency of ≥ 98 %), photon conversions are restricted to the beam pipe (0.17 % X_0) and the active area of the first pixel layer (0.15 % X_0), which, after folding with the energy-dependent conversion cross-section, corresponds to ~ 0.2 % conversion probability. Almost 80 % of these can be recognized by reconstructing the secondary vertex (conversion point). The remaining conversions are well below the level of electrons from Dalitz decays.

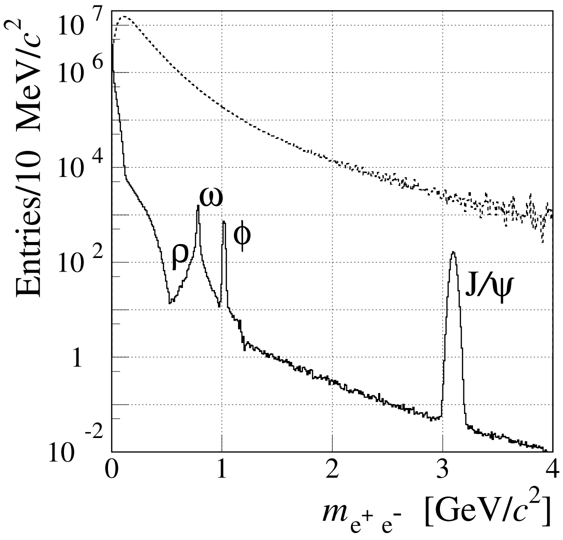

Figure 11.37: Invariant-mass distribution for electrons before any cuts. The dashed line shows the mass spectrum of all combinations (i.e. essentially combinatorial background) at a multiplicity of $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 8000$, the full line shows the mass spectrum of physical particles (sources) as generated in the Monte Carlo simulation.

In the next step, we try to identify and eliminate low-mass pairs ($m_{ee} < 100$ MeV) originating from hadronic decays with the following cuts:

- All electrons forming an unlike-sign pair with any other electron of $m_{ee} < 100$ MeV are discarded if the p_t of both electrons exceeds $9 \cdot m_{ee} + 50$ MeV/c. In this kinematic regime, Dalitz pairs can be clearly identified with very little background.

- Next, all electrons forming an unlike-sign pair with any other electron of $m_{ee} < 100$ MeV are discarded if the opening angle between the electrons fulfils the condition $\cos\theta > 1 - 0.0005 \cdot m_{ee}$ (MeV) or if $m_{ee} < 30$ MeV.

The low-mass pair rejection decreases the background from Dalitz decays by approximately an order of magnitude and has a signal efficiency of 35–45 %, depending on the mass. Owing to the presence of 'fake' electrons (misidentified pions) and prompt electrons from charm

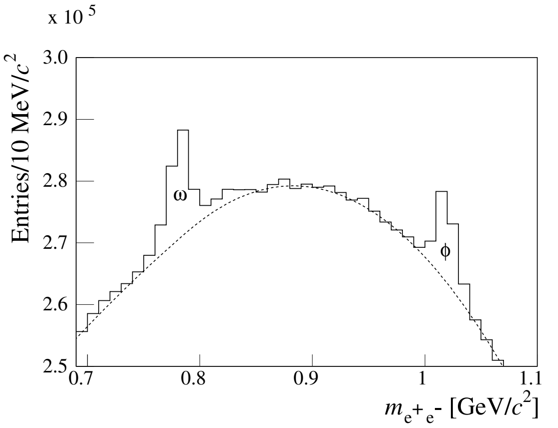

Figure 11.38: Invariant-mass distributions, after cuts, in the ω/ϕ region for 5×10^7 central events, including tracking efficiency and momentum resolution. Electrons and positrons have been selected from the momentum region $0.4 < p < 2$ GeV/c.

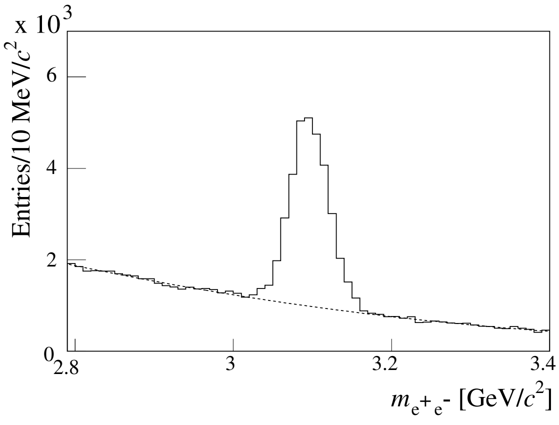

Figure 11.39: Invariant-mass distributions, after cuts, in the J/ψ mass region for 5×10^7 central events, including tracking efficiency and momentum resolution. Electrons and positrons have been selected from the momentum region $1 < p < 2.5$ GeV/c.

decays, the signal-to-background ratio increases, however, only by about a factor of 3.5. In this step, the electron efficiency needs to be very high down to the lowest possible momenta, as the final background will be roughly proportional to the square of the *inefficiency* to find both partners of the low-mass pair. A pion contamination in the electron sample of 10–20 % is acceptable as it enters only via the formation of a random low-mass combination with a real electron.

Finally, we remove *single* electrons which are outside the fiducial area of $90 \pm 40°$ or have a p_t outside the range $0.4 < p < 2$ GeV/c for the ω, ρ, ϕ selection and outside the range $1 < p < 2.5$ GeV/c for the J/ψ selection. The background contribution from semileptonic charm decays can be reduced by requiring that the electron track point to the main vertex. About 57 % of the electrons from charm decays can be rejected with a vertex cut which has 80 % efficiency for the signal.

The fiducial cut reduces the number of Dalitz pairs where one partner is outside the experimental acceptance, the lower p_t cut is effective because 3-body Dalitz decays have a much lower average momentum than the signal electrons. The pion rejection is now more important because the e/π ratio at high p_t drops to a few times 10^{-4}, whereas electron efficiency is less crucial then before; it now enters linearly into the significance $S/\sqrt{B}$. In this step we therefore can tighten our PID cuts. Nevertheless, the pion rejection in the TPC starts to get insufficient at $p > 2.5$ GeV/c.

11.6.6 Results

The invariant mass spectra for electron–positron pairs passing all our cuts are shown in Figs. 11.38 and 11.39 for the ω/ϕ and J/ψ region, respectively. They correspond to a statistics of 5×10^7 central events and no cut has been applied on the transverse momentum of the electron–positron pair.

The number of signal events (S) for different p_t bins in a 2σ region around the central mass bins together with the signal/background ratio (S/B) and the significance ($S/\sqrt{S+B}$) are shown in Tables 11.8 and 11.9.

The importance of excellent mass resolution (~ 1 %) is evident from Figs. 11.38 and 11.39, as it reduces both the statistical as well as the systematic error when extracting the yield of resonances. The spectra contain approximately 40 000 ω, 25 000 ϕ and 20 000 J/ψ mesons. Also the number of ρ's is large (40 000), but the signal is hardly visible owing to the huge background and the large width of the resonance.

For the low-mass mesons, Dalitz decays represent the dominant background. The influence of electrons from charm and fake electrons is, however, significant and for transverse momenta above 1.5 GeV/c they gain equal importance. Fake electrons dominate the background of the J/ψ especially for high transverse momenta.

For the ω and ϕ signals the significance corresponds to a 3 % and to a 10 % error, respectively, on the production cross-section. Owing to the softer p_t distribution of the background the errors decrease with increasing p_t. The signal-to-background ratios are low (between 2 and 10 %) but sufficiently broad guard bands exist to estimate the background under the signals. The significance of the ρ-signal is marginal at this multiplicity.

The S/B of the J/ψ signal ranges between 2 and 3 and the significance amounts to approximately 50 standard deviations, corresponding to a 2 % error on the measured production cross-section for each transverse-momentum bin. The production cross-section times branching ratio of the ψ′ is expected to be a factor 50 lower than that of the J/ψ. Integrating over all p_t bins the S/B ratio is 5 %

Table 11.8: Signal, signal-to-background ratio and significance for ω and ϕ in a 2σ region around the central mass bin for 5×10^7 central Pb–Pb events.

p_t (GeV/c)	ω S	ω S/B	ω $S/\sqrt{S+B}$	ϕ S	ϕ S/B	ϕ $S/\sqrt{S+B}$
0.50–0.75	6 500	0.02	11	4 700	0.02	10
0.75–1.00	7 000	0.03	14	4 000	0.03	10
1.00–1.50	13 000	0.05	25	7 400	0.04	17
1.50–2.50	12 000	0.12	36	7 600	0.1	25

and the significance for the ψ′ amounts to 5 standard deviations.

Table 11.9: Signal, signal-to-background ratio and significance for the J/ψ in a 2σ region around the central mass bin for 5×10^7 central Pb–Pb events.

p_t (GeV/c)	S	S/B	$S/\sqrt{S+B}$
0.50–0.75	4 300	2.3	55
0.75–1.00	4 800	2.9	60
1.00–1.50	7 400	2.7	74
1.50–2.50	4 700	1.6	54

11.7 Photon detection

11.7.1 Physics goals

The ALICE Photon Spectrometer (PHOS) gives access to a number of signals involving photons. These are in particular:

Direct photons: Photons are amongst the few signals from the early stages of the collision, including QGP, which are undistorted as they leave the collision zone without re-interaction. Thermal direct photons are expected to be emitted by hot dense matter of any kind, and to be measurable in the $1 \leq p_t \leq 3$ GeV/c range [62–65]. Their yield and spectral shape contain information on the space–time and temperature evolution of the reaction and should be sensitive to the presence of a long-lived phase transition.

Direct photons at somewhat higher p_t may carry information on parton dynamics at the pre-equilibrium stage. Finally, direct photons are also expected from initial-state hard QCD processes.

π^0 and η p_t spectra: Besides being the major background to the prompt photon analysis, the transverse momentum spectra of neutral mesons are of interest in particular in the very high-p_t region, where the resolution of the photon spectrometer is superior to the momentum resolution of the tracking system for charged particles. High-p_t ('leading') particles reflect the jet cross-section and can therefore indirectly measure the energy loss of partons. The energy loss of partons in a hadronic medium is expected to be different from the one in deconfined matter [66], which could lead to very significant differences in the yield of high-p_t particles between pp and heavy-ion reactions ('jet quenching') [67].

$\pi^0/(\pi^+ + \pi^-)$ ratio: The relative production of charged and neutral pions in strong interactions respects isospin symmetry, i.e. all pion species are equally abundant. However, during the phase transition, topological defects (called 'Disordered Chiral Condensates') could lead to very different ratios of charged to neutral pions [68] (see also Section 11.11). These condensates might be related to the production of Centauro events observed in cosmic ray experiments [69].

A number of calculations have been performed to estimate the yield of thermal photons at the LHC [70–73]. The predicted signal of thermal photons is strongly model dependent, ranging from less than 5 % of the inclusive photon spectrum to rates which are comparable to or even exceed the background of decay photons.

At the LHC, pertubative QCD processes (semi-hard scattering, minijets) play an important role and probably dominate particle production in nuclear collisions. This rather recent recognition has led to a substantial upward revision of the expected thermal photon yield (see e.g. [72, 74]).

Some theoretical estimates [70, 72, 73] for direct and decay photons produced in Pb–Pb collisions at the LHC are shown in Fig. 11.40.

11.7.2 Reconstruction of inclusive photon spectra

Each photon and electron, and about 50 % of the hadrons will produce showers in the PHOS. The electromagnetic showers are narrow ($\simeq 5$ cells at 1 GeV) and contain almost the full energy. Hadronic showers are wide and will fire more cells, but the total deposited energy E_{dep} will not exceed $\sim 1/3$ of their energy.

At high multiplicity, showers will frequently overlap. The photon reconstruction has therefore to make use of shower-finding and unfolding algorithms. Two different methods of photon reconstruction are under development:

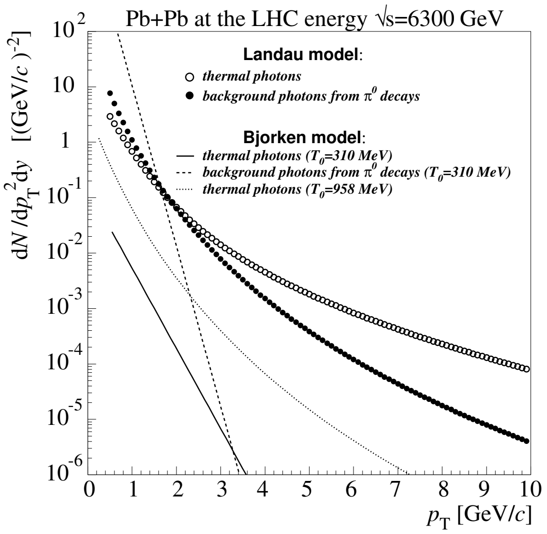

Figure 11.40: Direct and background photon p_t spectra for central Pb–Pb collisions at LHC for different scenarios.

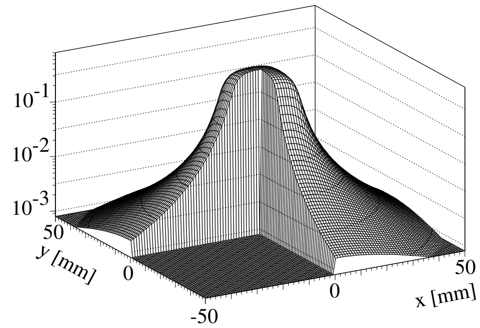

Figure 11.41: Shower profile function for the PHOS detector.

the shower profile method and the cluster method.

Shower profile method

This version of the event reconstruction program is based on methods developed for the NA12 (GAMS Collaboration) experiment at the CERN SPS [75]. It recognizes and reconstructs overlapping photons down to very small shower separation by comparing the energy distribution in PHOS cell clusters with the energy distribution expected for single- or multiple-photon hits [76].

The finite granularity of the detector is taken into account, in the case of single photons, by a 'shower profile function'. This function describes the normalized energy distribution in the cluster cells as a function of the primary photon coordinate within a cell. For incident angles below 10°, this spatial shower profile is energy independent [77].

The shower profile for the PHOS prototype, as measured in a test beam at the SPS, is shown in Fig. 11.41. The width of this function is close to the Molière radius in PHOS, i.e. overlapping showers should be resolvable for relative distances exceeding $R_M \approx 2$ cm.

The photon reconstruction program has been tuned for energies from 0.5 to 5 GeV, the most likely range where thermal photons might be measurable. Photon showers were generated using the measured shower profile function and including the energy resolution of each cell [78].

The unfolding efficiency is shown for the case of two showers with energies of 1 and 2 GeV, respectively, in Fig. 11.42 and for the case of three equidistant showers of 1 GeV each in Fig. 11.42. Overlapping showers with energies of 1 GeV and more can be resolved at distances above 2 cm. The shower-resolving efficiency increases with shower energy and decreases with the ratio of shower energies, since a low-energy shower may be hidden by an energy fluctuation of the high-energy neighbouring shower.

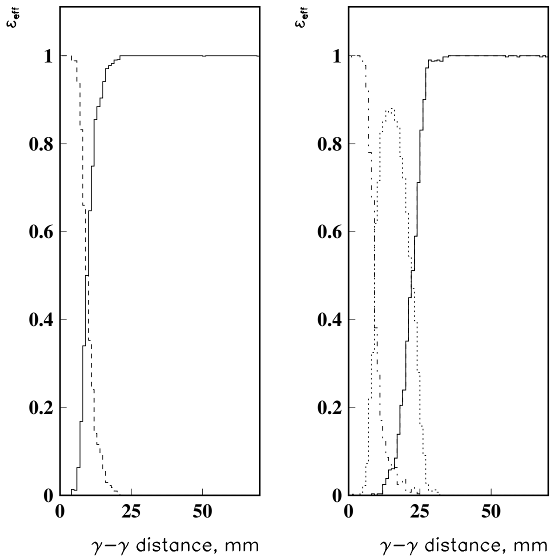

Figure 11.42: Left: For two overlapping showers of 1 and 2 GeV as function of relative distance (in mm): Probability to reconstruct correctly both showers (solid line) or to merge the two photons (dotted line); Right: For three equidistant showers of 1 GeV each: Probability to reconstruct all three photons (solid line), two photons (dotted line), or one single shower (dashed-dotted line).

The reconstruction program has been tested with SHAKER events using a parametrization of the measured shower profile function and applying an energy threshold of 30 MeV per cell. Charged hadron showers are not yet included in the simulation.

More than 30 % of the photons belong to clusters with two or more overlapping showers. Nevertheless, the program reconstructs the original photon spectrum very efficiently. This can be seen in Fig. 11.43, which shows that the ratio of reconstructed over generated distributions is very close to 1 in the region between 400 MeV up to at least 3 GeV. Slight deviations from 1 at low energy reflect smearing of the spectrum because of energy resolution.

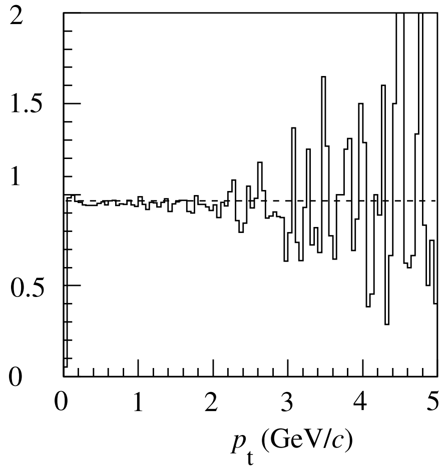

Figure 11.43: Ratio of reconstructed over generated photon distributions (shower profile method).

Cluster method

The second photon reconstruction program is in use in the WA80, WA93 and WA98 experiments [79]. There, a finely granulated photon spectrometer of 10 080 lead glass modules measures photons in heavy-ion reactions in an environment very similar to that of the LHC.

It uses a recursive cluster-finding and cluster-unfolding procedure which recognizes overlapping showers only if they are separated by a local minimum. While the two-photon resolution is therefore limited, this algorithm is simple and robust, and its efficiency can be very well measured by means of event mixing. The method is described in detail in Ref. 80.

The method was tested with the LoI design of PHOS, i.e. a rectangular array of 50 445 $PbWO_4$ crystals of size $2.2 \times 2.2 \times 18$ cm located 6 m from the interaction point. For this configuration a full detector simulation (GEANT 3.15) was performed for central Pb–Pb (SHAKER generator) to evaluate the photon reconstruction efficiency, shown in Fig. 11.44. This 'efficiency' can be larger than one because photons can be shifted towards higher energy by shower overlap.

11.7.3 Rejection of charged particles

Charged particles can be rejected, in principle, with the help of the charged-particle veto (CPV) detector, described in Chapter 5. However, some fraction of charged particles will always escape detection in the CPV —in particular at high multiplicity— and has to be subtracted

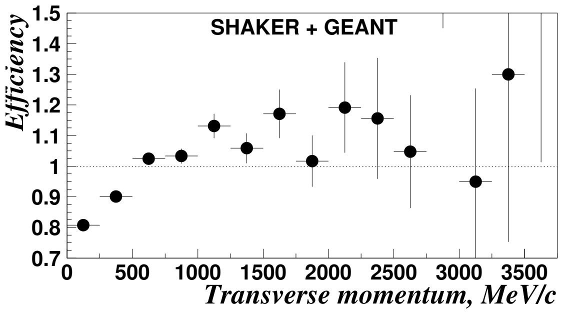

Figure 11.44: Photon registration efficiency for the cluster method.

from the photon spectrum. This can be done, via the measured response of the PHOS for showers tagged in the CPV as charged particles, provided the CPV efficiency is known with good accuracy.

For very high occupancy, were the accidental overlap between photons and charged particles can be substantial, it might be more advantageous *not* to veto showers individually but rather to subtract the full charged particle contribution from the inclusive energy spectrum. This method is used in the WA80 data analysis. An example of the separation of charged from neutral particles is shown in Fig. 11.45 in the region of the minimum-ionizing particle peak. This method does not impose any strict requirements on the position resolution or the charged-particle veto efficiency.

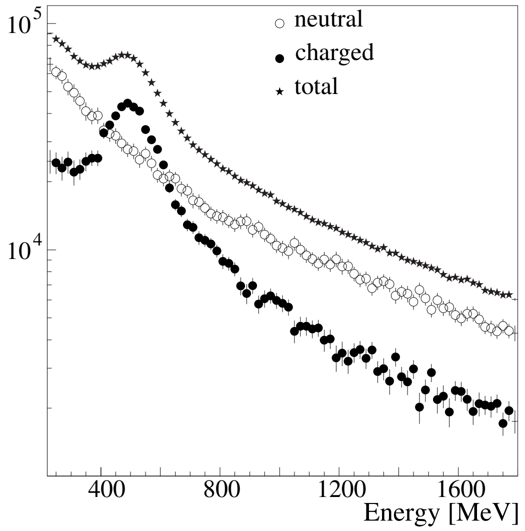

Figure 11.45: Separation of neutral and charged particle contributions to the measured energy spectrum in WA80. Total spectrum (stars), response to charged particles tagged by a veto counter (full circles), neutral particle spectrum (total − charged, open circles).

11.7.4 Rejection of neutral hadrons

K^0_L's, neutrons and antineutrons are not rejected by the charged particle veto and are, therefore, photon candidates. The PHOS will detect them with a probability of $\sim 50\,\%$ and will record a fraction of the total energy (K^0_L) or kinetic energy (neutron). Antineutrons can annihilate in the detector and deposit correspondingly more energy.

The importance of this background is illustrated in Fig. 11.46, which shows the fraction of all neutral showers caused by (anti)neutrons as a function of energy $E_{\rm dep}$ measured in the calorimeter.

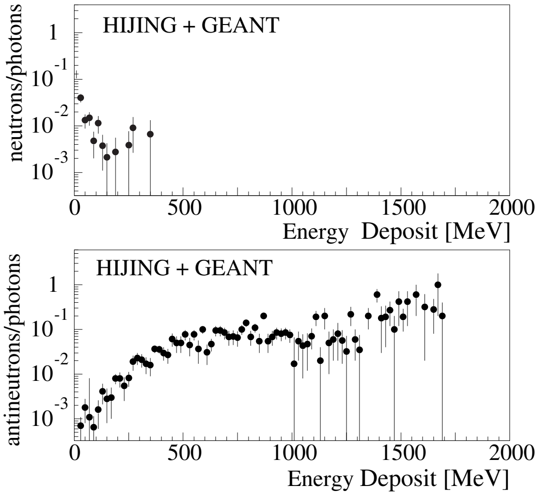

Figure 11.46: n/γ and $\bar{n}/\gamma$ ratios, simulated with HIJING and using the GEANT detector response.

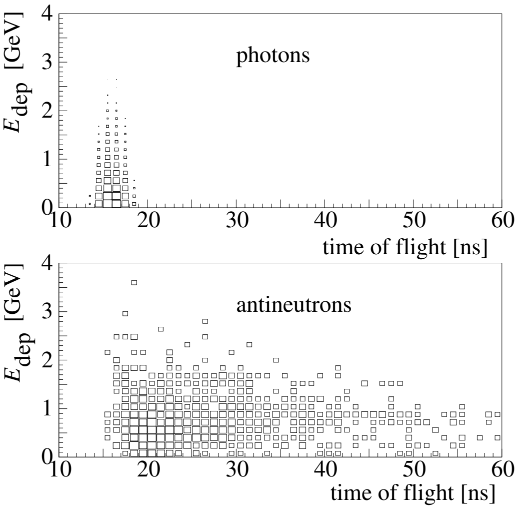

Figure 11.47: Deposited energy $E_{\rm dep}$ vs. TOF for photons and antineutrons (TOF resolution 1 ns).

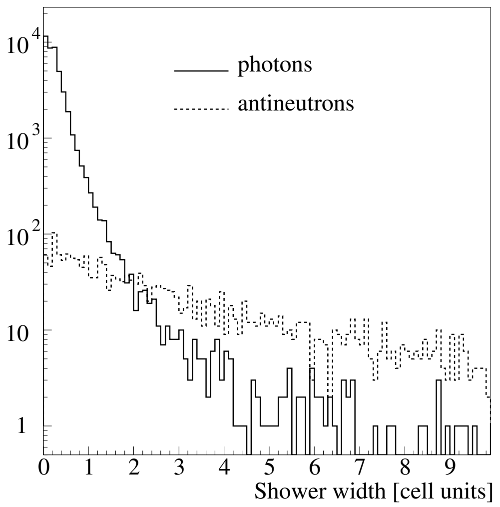

Figure 11.48: Shower width distributions for photons and antineutrons (GEANT simulation).

As expected, the n/γ ratio is rather small, while antineutrons contribute substantially to the photon spectrum in the $p_{\rm t}$ range of interest.

Neutral hadrons can be separated from photons either by time of flight (TOF) or by a shower shape analysis. The shower width is much larger for hadrons than for photons and can discriminate between the two at all except the lowest energies.

Figure 11.47 compares the two-dimensional ($E_{\rm dep}$ vs. TOF) spectra of photons and antineutrons, assuming a timing resolution of 1 ns.

A time-of-flight cut with 99 % photon efficiency would reject $\sim$ 50–85 % of the antineutrons (depending on their $p_{\rm t}$).

The second method uses the difference in shower widths, defined as the larger r.m.s., between the two transverse directions, of the shower profile. Figure 11.48 shows the shower width distributions for photons and antineutrons.

Figure 11.49 shows the fraction of rejected antineutrons vs. the fraction of lost photons for different cuts on the shower width. It is possible to reject more than 60 % of the antineutrons while keeping 99 % of the photons by applying a cut on the shower width at 1.8 cell sizes. After this cut, the ratio $\bar{n}/\gamma$ is reduced to below 10 % (see Fig. 11.50).

This method works equally well at all $p_{\rm t}$. The time-of-flight method would add further improvement at intermediate $E_{\rm dep}$; its main disadvantage is the extra cost required to achieve a good time resolution ($\sim$ 1 ns).

The remaining antineutrons can be subtracted inclusively from the photon spectrum. The antineutron flux

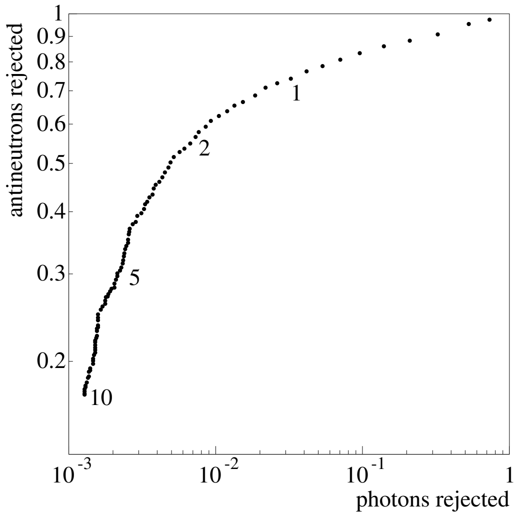

Figure 11.49: Fraction of rejected antineutrons vs. fraction of rejected photons for different shower width cuts. Numbers on the figure correspond to shower width cuts, measured in units of cell size.

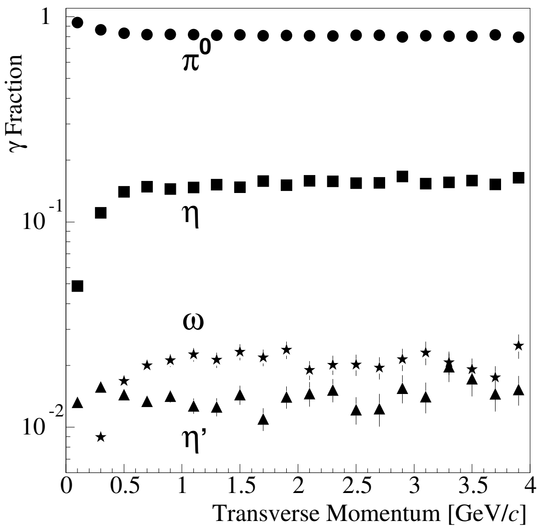

Figure 11.51: Fractions of photons produced by various meson decays, normalized to the total number of decay photons. Neutral meson yields are calculated from the π^0 spectrum assuming transverse mass scaling.

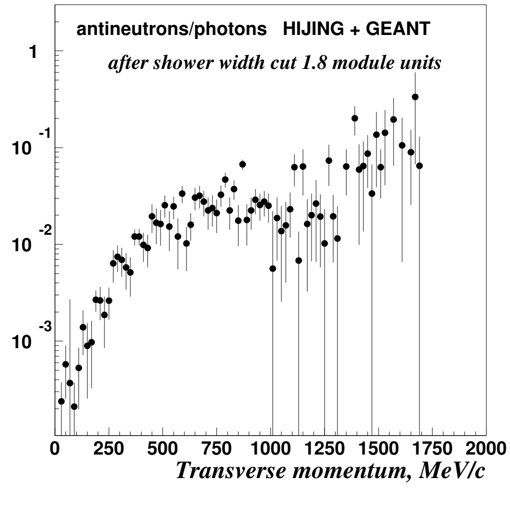

Figure 11.50: $\bar{n}/\gamma$ ratio with a cut on shower width at 1.8.

at the LHC around midrapidity is equal to the proton and antiproton flux, which will be independently measured in ALICE. The PHOS response to antineutrons can then be either simulated or measured (by comparing p and $\bar{p}$ response) and subtracted from the measured photon spectrum.

11.7.5 Momentum spectra of π^0 and η

Inclusive π^0 and η distributions will be extracted from the two-photon invariant mass spectra for different centrality classes.

With the chosen geometry, the low-p_t detection limit is ~ 0.5 GeV/c for both π^0 and η mesons (see Section 5.1).

The upper p_t detection limit is ~ 50–60 GeV/c for the π^0. The two photon showers are in this case separated by 2 cm; they overlap and form a single cluster whose energy profile is distinctly elongated and cannot be faked by the circular profile of a single photon shower.

Charged particle rates (mainly $\pi^\pm$) were computed up to $p_t = 10$ GeV/c using two models, HIJING and the Parton Cascade model [81], and extrapolated to 60 GeV/c using the (probably optimistic) assumption $dN/dp_t^2 \propto p_t^{-(4-6)}$. In 10^7 central Pb–Pb collisions, we could expect between 0.6 and 70 π^0's per 1 GeV/c p_t bin measured in PHOS at 60 GeV/c, depending on model and p_t dependence [11].

11.7.6 Extraction of direct photon spectra

Direct photon yields and p_t spectra will be determined on a statistical basis for different centrality classes by subtracting from the observed inclusive γ spectrum $\gamma_{obs}(p_t)$ the decay photon spectrum. The latter is calculated with a Monte Carlo program using the measured π^0 and η distributions. Two other sources have to be added, i.e. ω and η' (see Fig. 11.51). Their yield and spectra will be obtained from the π^0 spectrum by assuming transverse mass scaling.

Once the spectrum of background photons expected

[11] Extrapolating with a parameterization given by UA1 (CERN/PPE 94-77) yields about 10 π^0's above 25 GeV/c.

from the measured meson decays $\gamma_{\mathrm{bkg}}(p_{\mathrm{t}})$ has been calculated, one can, in principle, extract the direct photon excess using $\gamma_{\mathrm{dir}}(p_{\mathrm{t}}) = \gamma_{\mathrm{obs}}(p_{\mathrm{t}}) - \gamma_{\mathrm{bkg}}(p_{\mathrm{t}})$.

Since the large majority of photons observed at a given p_{t} originate from the decay of π^0's, it is more convenient to work with the ratio

$$\gamma_{\mathrm{obs}}/\gamma_{\mathrm{bkg}} = 1 + \gamma_{\mathrm{dir}}/\gamma_{\mathrm{bkg}} \,,$$

which gives the fraction of photons observed relative to the expected decay background and should have a value of 1 if there are no excess photons.

However, if the direct photon yield at the LHC turns out to be at the limit of detectability, then it would still be possible to measure the excess in $\gamma_{\mathrm{dir}}/\gamma_{\mathrm{bkg}}$ in central collisions over the same ratio in peripheral collisions, assuming that in the latter case few direct photons are radiated, and making use of the fact that the systematic errors largely cancel in the comparison.

11.7.7 Photon reconstruction efficiency

In a high-multiplicity environment one of the major sources of systematic errors are shower overlaps. Photons could be lost or change their energy or coordinates, which would lead to a distortion of the measured photon spectrum.

The distortion of the photon spectrum caused by the shower overlaps can be found for various centrality classes by adding either real or simulated test photons to real events (one test photon per event) and then re-extracting them. In this way the probability of a photon changing energy or being lost due to overlap with another shower can be calculated [82,83]. Comparison of efficiencies calculated with different methods and cuts permits an estimate of systematic errors.

Results of the simulation for ALICE Simulations (see Section 11.7.6) show that the PHOS spectrometer will detect ~ 1000 charged and neutral particles in a single central collision; i.e. the cell occupancy will be ~ 10–$12\,\%$.

The ratio of the reconstructed photon spectrum to the generated photon spectrum is shown in Fig. 11.52, using the cluster method mentioned earlier. Within statistical errors, the calculated efficiency (see Fig. 11.44 on page 180) permits a reconstruction of the true photon spectrum.

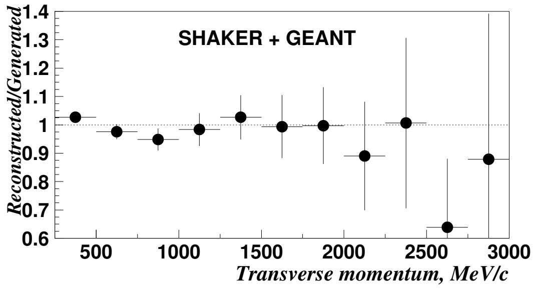

Figure 11.52: Ratio of reconstructed and generated photon spectra obtained in a simulation for the PHOS detector.

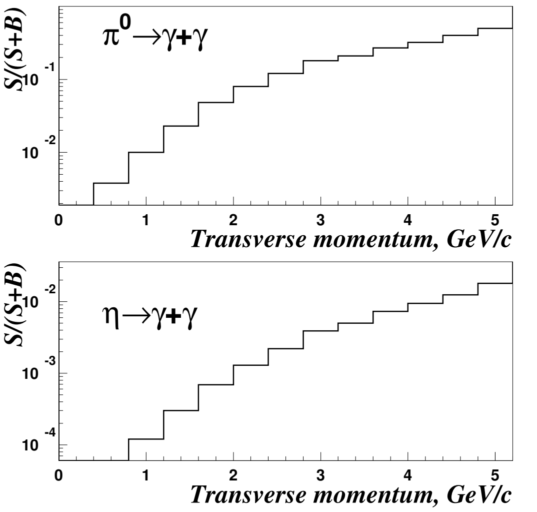

Figure 11.53: Signal-to-(signal + background) ratio vs. transverse momentum.

11.7.8 π^0 and η yields and reconstruction efficiencies

In relativistic nucleus–nucleus collisions, the number of particles produced in each collision is very high, and so is the combinatorial background (B) in the two-photon invariant mass spectrum. The signal fraction ($S/(S+B)$) for the π^0 and η mesons in central Pb–Pb collisions is shown in Fig. 11.53.

Because of the large background, the meson yields can be determined only inclusively. In order to extract the π^0 (or η) yield, it is necessary to reconstruct with very high precision the combinatorial background in the relevant mass region.

This can be done with the event mixing technique, as used, for example, for the WA80 data analysis [84, 85]. The π^0 and η reconstruction efficiencies are evaluated with a method similar to the one described previously for photons, i.e. by adding simulated test π^0's or η mesons to real events.

It was demonstrated in the WA80 data analysis that the π^0 yield can be determined with the systematic error $\leq 5\,\%$ at $S/B \leq 1\,\%$.

11.7.9 Systematic errors

We will be able to collect sufficient events from heavy-ion collisions so that the measurement of direct photons will be limited by systematic rather than statistical errors over most of the relevant p_t range. While statistical errors can be reasonably estimated, a good *a priori* estimate of systematic errors by Monte Carlo simulation is difficult.

Systematic errors can be of two categories. Some parameters (e.g. the linearity of the detector) are known only with a certain precision. However, the range allowed for these parameters can be constrained with the actual data (e.g. via the position of the π^0 peak as a function of p_t), and therefore the possible systematic error is reduced to a statistical error. These errors are reasonably well controllable and will in general decrease with increasing statistics.

The second category includes assumptions which can not be tested with the data (e.g. the ω/π ratio). One can estimate the resulting uncertainty only via an 'educated guess', but in this case the 'confidence' in the error estimate is not well defined. We will therefore try to relate most systematic errors to quantities which can be controlled with the actual data (e.g. measure the η rather then use models). A detailed list of systematic errors expected in ALICE can be found in Ref. 86. Below we give a brief update and compare in particular with a similar analysis [82] performed within the WA80 and WA93 experiments. The PHENIX Collaboration has evaluated the performance of a detector similar to PHOS for RHIC and estimates a systematic error of below 4 % [87].

The following sources of systematic errors are possible.

- **Photon reconstruction efficiency:** Error in WA80 $\sim 2\,\%$. This error can be estimated by comparing efficiencies calculated using different methods (see Section 11.7.7). Since the mean cell occupancy for ALICE will be lower than in WA80 by almost a factor of two, one could expect that this error will be $\leq 2\,\%$ for central collisions and $\leq 1\,\%$ for peripheral collisions.

- **Charged and neutral particles:** Error in WA80 $\sim 1\,\%$. In ALICE, this error should be comparable or smaller, since the charged-particle veto detector will be situated closer to the photon spectrometer.

- e^+e^- **conversions:** Error in WA80 $\leq 1\,\%$. This error can be reduced to a minimum if the exact amount of matter between the interaction point and the PHOS is known.

- π^0 **acceptance:** Error in WA80 $\sim 0.5\,\%$. This error is determined by uncertainties in geometry and can be reduced to a minimum if the PHOS geometry is measured with good accuracy.

- π^0 **efficiency:** Error in WA80 ~ 4–$5\,\%$ (central); ~ 2–$3\,\%$ (peripheral). This is the largest systematic error. It can be estimated by applying different cuts to the photons used for invariant mass calculation (e.g. energy cut, CPV cut, shower width cut, one maximum in a cluster cut, etc.). Different methods of background normalization and different invariant mass ranges for summing up of π^0's should also be used. Again the smaller cell occupancy in ALICE, in between central and peripheral collisions in WA80, should lead to a smaller error of about ~ 3–$4\,\%$.

- **Other neutral mesons:**. Error in WA80 $\leq 1.5\,\%$. This error is determined mostly by the knowledge of the η meson yield. Since we plan to measure η yield with 10 % accuracy —and η mesons contribute $\sim 15\,\%$ of all decay photons— the uncertainty in the final result should be $\leq 1.5\,\%$.

- π^0**'s from strange particle decays:** Error in WA80 $\sim 0.1\,\%$. Since strange particles have a long lifetime, they may decay into π^0's far from the interaction point and are therefore reconstructed with the wrong mass, leading to an underestimate of the π^0 cross-section. A simulation [88] shows that this effect is very small.

- **Nonlinearity of the detector:** Error in WA80 ~ 1–$2\,\%$. The nonlinearity has to be controlled with data by measuring the π^0 peak position for different p_t.

- **Antineutrons**. As was shown in Section 11.7.4, the neutron contribution to the photon spectrum is negligible, and the antineutron contribution can be reduced to $\leq 10\,\%$. The remaining error is determined by how well the yield and detector response can be extracted from p and $\overline{\mathrm{p}}$ spectra. This should be possible to much better than 10 %, leading to a residual systematic error of $\leq 1\,\%$.

Adding all error in quadrature, the total systematic error in the direct photon determination in central Pb–Pb collisions will be $\sim 5\,\%$.

11.8 Hyperon detection

The production of strange particles should be sensitive to the presence and features of the QGP phase, as the abundance of strange quarks should be higher in a QGP than in the final states of normal hadronic interactions. Heavy particles and those with multiple strangeness are of particular interest, because they are difficult to produce in thermal collisions and hence their abundances might not change significantly during the lifetime of a hadron gas. Therefore the yields of multi-strange hyperons are potentially a signature for enhanced strangeness density at an early stage of the collision [89].

11.8.1 V^0 and cascade detection

We shall detect hyperons —with $c\tau$ of the order of centimetres— by reconstructing the decay products in the tracking system and by finding secondary vertices near to the interaction point. The impact parameter resolution is more than sufficient for this purpose.

Currently we search for secondary vertices only well within the beam pipe. Therefore our new ITS lay-out, with a smaller beam pipe diameter and a better impact parameter resolution than the LoI design, has also reduced the fiducial volume for V^0 finding as compared to our previous studies [90].

In order to find tracks of large impact parameter, we relax the vertex constraint during the final pass of the 'guided' track finding in the ITS, and no constraint is imposed during the track reconstruction. To search for secondary vertices, we combine one large impact parameter track (more then 1 mm) with a track of opposite sign, regardless of its impact parameter. PID information is not needed, as the backgrounds are sufficiently small. A combination is retained if:

- the closest distance between the two tracks is below 100 μm,
- the secondary vertex is inside the fiducial volume, defined as the region between $0.5 < r < 2.5$ cm, and
- the V^0 momentum points well back to the primary vertex.

Applying these cuts we typically end up with a background of a few per cent of the V^0 signal. The effective mass distribution for a Λ mass assignment for 10 central Pb–Pb events is shown in Fig. 11.54.

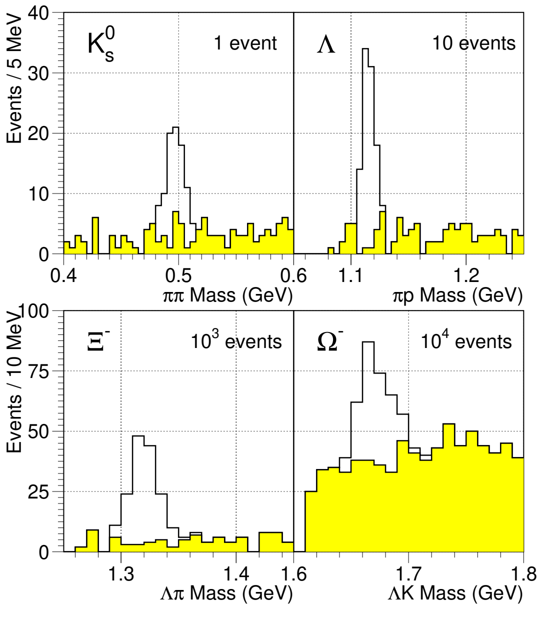

Figure 11.54: Effective mass distribution for K^0_S, Λ, Ξ^- and Ω^- candidates from 1, 10, 10^3 and 10^4 central Pb–Pb events, respectively.

When searching for cascades (Ξ^- and Ω^-), the last cut is relaxed. We still require backward pointing, but with a much larger window. The Λ candidates found in the previous step are combined with negative large impact parameter tracks (and $\overline{\Lambda}$ candidates with positive tracks) and the reconstructed cascade should point back to the primary vertex. Although a full helix extrapolation has been used, for all practical purposes a linear extrapolation is sufficient. The background for the cascades is of the order of a few per cent for Ξ^- and of about 30 % for Ω^- (see Fig. 11.54).

11.8.2 Expected rates

We investigated reconstruction efficiencies for the following decay chains:

- $K^0_S \to \pi^+\pi^-$
- $\Lambda \to \pi^-p$
- $\Xi^- \to \pi^-\Lambda \to \pi^-\pi^-p$
- $\Omega^- \to K^-\Lambda \to K^-\pi^-p$

Geometrical acceptance times branching ratio for each of the above decay chains is shown in Fig. 11.55 as a function of p_t and rapidity. In order to get the total number of reconstructed particles, we have to convolute the accep-

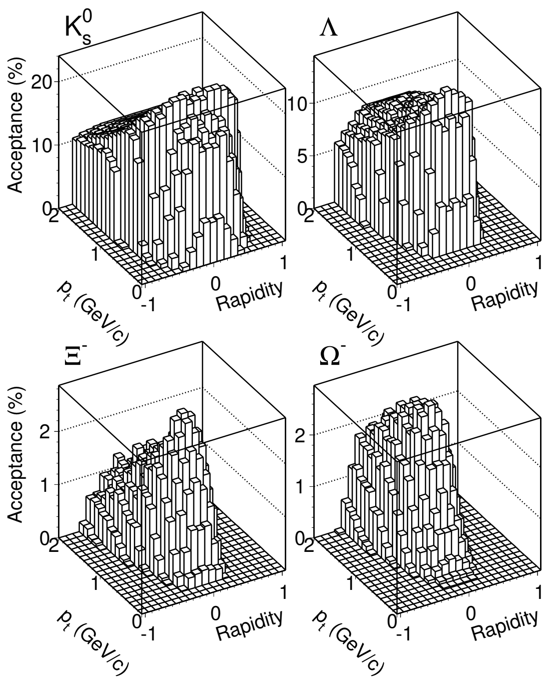

Figure 11.55: Acceptance as a function of rapidity and p_t for K^0_S, Λ, Ξ^- and Ω^- detection. The decay branching ratios for detected channels are included in the acceptance.

Table 11.10: Integrated acceptance (relative to production in $\Delta y = \pm 1$), reconstruction efficiency, production ratio relative to charged pions and number of reconstructed decays for V^0 particles and cascades in central Pb–Pb. The branching ratio (BR) is included in the acceptance.

Particle → decay	BR (%)	Acceptance (%)	Efficiency (%)	Produced/π (%)	Detected/event
$K^0_S \to \pi^+\pi^-$	69	12	66 ± 5	5	63 ± 5
$\Lambda \to \pi^- p$	64	4.2	61 ± 6	2	8.3 ± 0.8
$\Xi^- \to \pi^-\Lambda \to \pi^-\pi^- p$	64	0.58	47 ± 7	0.3	0.13 ± 0.02
$\Omega^- \to K^-\Lambda \to K^-\pi^- p$	44	0.48	45 ± 7	0.05	0.017 ± 0.004

tance and the reconstruction efficiency with the primary production rate. The results, integrated over p_t and acceptance, are summarized in Table 11.10. Hyperons were generated flat in rapidity and with momentum distributions following m_t-scaling. The integrated acceptance is normalized to the overall production in two units of rapidity

The expected multiplicities for K^0_S and Λ detected in one event is sufficient for event-by-event analysis. The statistics accumulated for cascades in 10^7 central Pb–Pb events (dN_{ch}/dy = 8000) will be in the range 10^5 to 10^6 (for Ω^- and Ξ^-, respectively), which will allow a high-statistics inclusive analysis in sub-samples of 10^5 to 10^6 events.

11.9 Open charm detection

11.9.1 Physics motivation

Owing to its large mass, charm can only be produced in the early, hotter stages of the collision. The probability of subsequent annihilation is very low, so that essentially all produced charm ends up in the final state, mostly in the form of open charm hadrons. Open charm is therefore an especially interesting probe of the dynamics of high-energy nucleus–nucleus collisions. Enhancement of charm production in nucleus–nucleus collisions has been predicted for some time [91]; in addition to initial-state hard scattering processes, charm could also be produced via parton–parton rescattering during the pre-equilibrium (thermalization) phase.

The degree of enhancement is expected to be sensitive to the initial value of the energy density and to the duration of the thermalization phase [91–93]. Predicted values for the ratio of total to initial production range from about two [92] to about ten [93]. Moreover, the pre-equilibrium production in central collisions is expected to depend on the mass number of the colliding nuclei more strongly (A^2) than the initial production ($A^{4/3}$) [92, 93]; the deviation from $A^{4/3}$ behaviour should provide a direct measurement of the importance of rescattering production. Even once thermal equilibrium is reached, the temperature could be sufficiently large to allow charm production to continue during the plasma phase itself, i.e. 'thermal' production [94].

11.9.2 Experimental approach

Lower mass charm hadrons can only decay through weak processes, and travel a measurable distance before decay ($c\tau \simeq 100$ μm). The fraction of decays involving only charged products is large, allowing direct detection of charm decays by computing the invariant mass of completely reconstructed secondary vertex topologies. The key parameter which measures the ability of a detector to separate secondary vertices from the primary interaction vertex is the impact parameter resolution. The impact parameter is defined as the distance in space between the backward projection of the reconstructed track and the primary vertex. Charm decay tracks have typical impact parameters of the order of 100 μm.

With the design described in the ALICE LoI (one Si pixel layer at 7.5 cm from the beam) the impact parameter resolution would also be about 100 μm. In order to separate the few charm decay tracks from the large number of primary tracks, we need an impact parameter resolution of about 50 μm, which can be obtained by adding an extra layer of Si pixels at $r \simeq 4$ cm (see Fig. 11.12 on page 163). This enables us to have sufficient precision to perform full vertexing, which is necessary to suppress the large combinatorial background.

11.9.3 Simulation studies

Assuming no enhancement effects, one can estimate the yield of charm pairs created at central rapidity in Pb–Pb collisions by considering only the 'initial' production extrapolated from pp collisions. Using $c\bar{c}$ cross-sections calculated to next-to-leading-order with HERA structure functions [61], we have estimated for central Pb–Pb events at the LHC a rate of 15–40 $c\bar{c}$ pairs per unit of rapidity in the central region (the spread is due to the use of two different sets of structure functions). In the following, all calculations are done assuming the lowest prediction, i.e. 15 $c\bar{c}$ pairs per unit of rapidity.

The most important source of background is given by those primary tracks which acquire an impact parameter due to multiple scattering. Primary tracks have been simulated by generating SHAKER events, and by smearing the track parameters according to measurement and multiple scattering errors. The main contribution to the impact parameter resolution comes from multiple scattering in the beam pipe and in the first pixel layer. This has been simulated in detail applying the complete Molière theory.

In the exclusive approach with secondary vertex reconstruction and invariant mass calculation, charm decays themselves represent an important source of background. Therefore, we have also included all the charm decay products in the background sample.

The contribution of other sources of background (hyperon, kaon and pion decays, secondary interactions at the beginning of the tracker and fake tracks) has been simulated by adding a flat background of impact parameter tracks. We have estimated that the number of background tracks from these sources, which have impact parameter below 1 cm, is $\sim 1\,\%$ of the number of primary tracks. Increasing this amount up to $3\,\%$ has no effect on the results. In the analysis we require tracks to be reconstructed with points in all six ITS layers, in order to minimize the number of poorly measured tracks.

11.9.4 D^0 detection

D^0 mesons have a mean lifetime of 0.415 ps ($c\tau = 124.4\ \mu$m). Their decay branching ratio into $K^-\pi^+$ is about 4 %. This is the simplest charm decay channel, and therefore a very good benchmark.

Using the lowest estimate for the charm production rate mentioned above, one expects 0.37 D^0 (or $\overline{D^0}$) per event where D^0 (or $\overline{D^0}$) decays into $K^-\pi^+$ (or the charge conjugate channel), with both products in the acceptance of the central ALICE detector. The total number of combinations of a charged kaon and a charged pion of opposite sign is of the order of 10^7 per event, so that the initial signal over background (S/B) is about 4×10^{-8}. Kinematic cuts on invariant mass and decay angle increase the S/B to about 2×10^{-5}. To improve on this value, one needs secondary vertex detection capabilities.

We have generated D^0 mesons with a flat rapidity distribution and a p_t distribution calculated using the same set of structure functions as for rate calculations. The momentum of the decay products was smeared in the same way as for the background sample described above.

Figure 11.56 shows the distribution of the $r\varphi$ projection of the impact parameter for pions with $p_t > 500$ MeV$/c$ coming from the D^0 sample as well as from the background sample (tracks with $p_t < 500$ MeV$/c$ are eliminated by the final cuts). The impact parameter distribution from the D^0 sample is roughly twice as wide as the one from the background sample.

The $r\varphi$ impact parameters are given a sign in order to distinguish tracks passing on opposite sides with respect to the primary vertex. In order to have a good separation of the D^0, both impact parameters have to be large. As a rough momentum balance requirement, we also demand opposite signs for the two impact parameters.

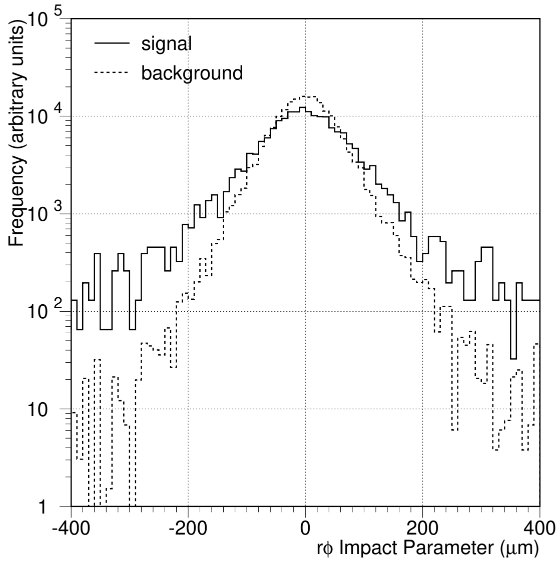

Figure 11.56: $r\varphi$ impact parameter distribution for pions coming from the signal (solid) and from the background (dashed) sample. The two distributions are normalized to the same number of entries.

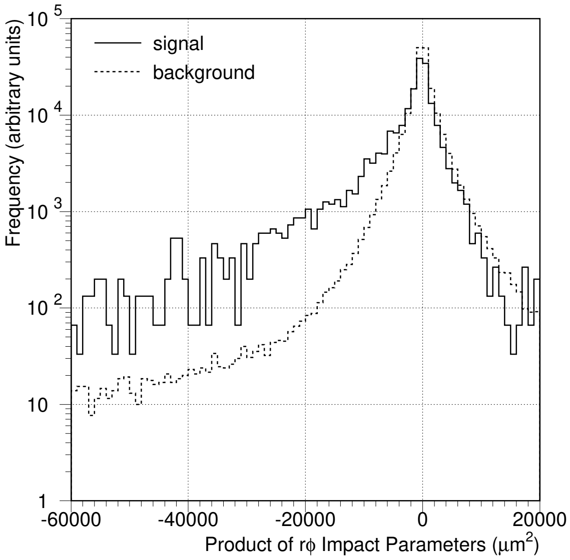

Figure 11.57: Product of the pion and kaon $r\varphi$ impact parameters for signal (solid) and background (dashed) combinations. The two distributions are normalized to the same number of entries.

The two above conditions can be combined by requiring that the product of the two impact parameters be large and negative. Figure 11.57 shows the distribution of the products of the kaon and pion impact parameters for the signal and for the background samples, normalized to the same number of entries. As can be seen from the figure, the S/B can be easily improved by a factor of about 10 by cutting on this variable.

A more stringent pointing condition can be imposed by a cut on the angle between the momentum vector of the D^0 candidate, and the line connecting the primary vertex to the candidate decay vertex. This pointing angle θ_{point} should ideally be zero for the signal, and essentially randomly distributed for the background. The cosine of this angle is shown in Fig. 11.58. Restricting the selection to values of $\cos\theta_{\mathrm{point}}$ close to 1 would also give, by itself, a background rejection of an order of magnitude. The long tail to the left of the signal is due to decays close to the primary vertex, for which the line of flight is poorly measured.

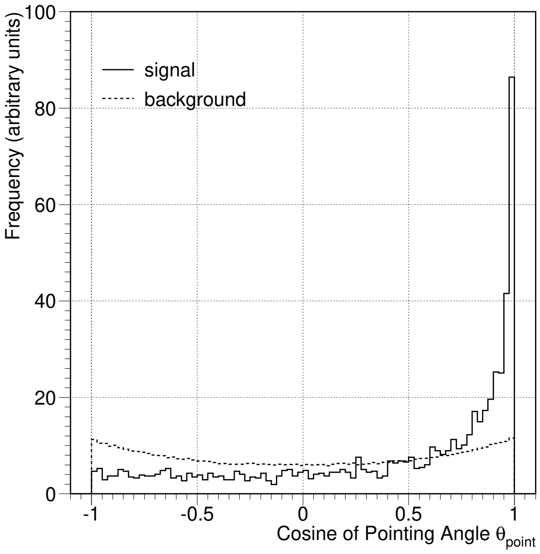

Figure 11.58: Cosine of the pointing angle for signal (solid) and background (dashed) combinations. The two distributions are normalized to the same number of entries.

While separation from the primary vertex and good pointing are essentially uncorrelated for the background, the line of flight can be measured with good precision for a signal combination that is well separated from the primary vertex. Therefore a combination of the two cuts can give a very large gain in S/B. This is seen in Fig. 11.59, which shows the scatter plot of $\cos\theta_{\mathrm{point}}$ versus the product of the kaon and pion impact parameters. Selecting combinations for which the product of the impact parameters is less than $-2\times10^4\mu\mathrm{m}^2$ and $\cos\theta_{\mathrm{point}} > 0.95$ leaves 4.6×10^{-3} signal and 0.21 background combinations per event, i.e. a S/B ratio of 2.2 %.

For a sample of 10^7 events (corresponding to about five days of running at 50 % data taking efficiency), the significance would be $S/\sqrt{B} = 32$, which provides us with a good margin to accomodate unavoidable differences between simulations and reality. The corresponding mass plot is shown in Fig. 11.60, before and after background

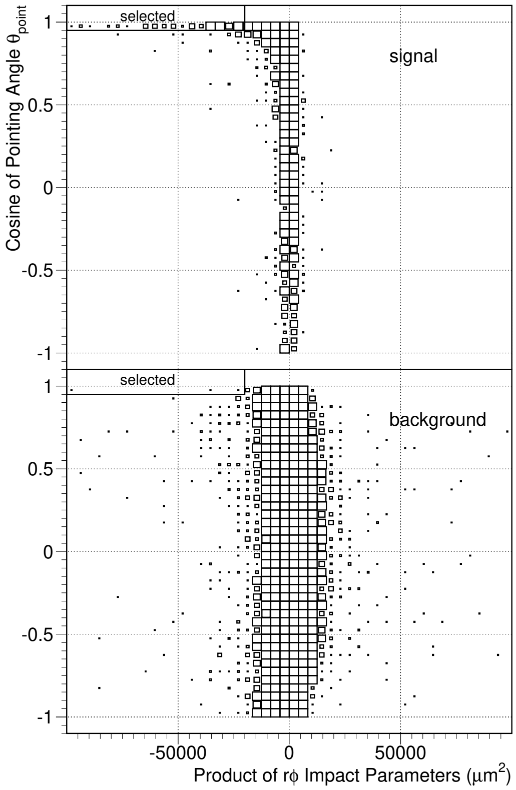

Figure 11.59: $\cos\theta_{\mathrm{point}}$ versus the product of the kaon and pion impact parameters for signal (top) and background (bottom) combinations.

subtraction.

The above numbers include the effects of reconstruction and particle identification inefficiencies. The significance is proportional to the single track efficiency $\varepsilon_{\mathrm{tr}}$ and is rather weakly dependent on the kaon identification efficiency η_{K} and on the $\pi\to\mathrm{K}$ misidentification probability $\delta_{\pi\to\mathrm{K}}$:

$$\frac{S}{\sqrt{B}} \propto \frac{\varepsilon_{\mathrm{tr}}\,\eta_{\mathrm{K}}}{\sqrt{\eta_{\mathrm{K}} + \frac{N_\pi}{N_{\mathrm{K}}}\delta_{\pi\to\mathrm{K}}}},$$

We assumed $\eta_{\mathrm{K}} = 0.4$ (including a probability of 0.55 for the survival of charged kaons up to the PID), $\delta_{\pi\to\mathrm{K}} = 0.008$, $N_\pi/N_{\mathrm{K}} = 5$ and $\varepsilon_{\mathrm{tr}} = 0.85$ (this corresponds to restricting the selection to tracks with hits in all six ITS layers).

In the extreme case of no particle identification the significance would be worse by a factor $\simeq 2.6$.

11.9.5 D^+ detection

Charged D mesons have a larger mean lifetime than neutrals ($c\tau = 317\ \mu$m). In addition, the branching ratio into the simplest decay channel involving only charged

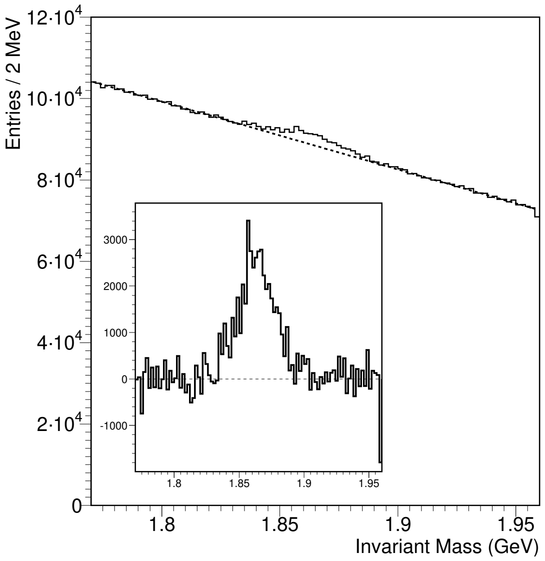

Figure 11.60: $K^-\pi^+$ (+c.c.) invariant mass distribution for 10^7 events. The same distribution after background subtraction in shown in the inset.

products ($K^-\pi^+\pi^+$) is larger than for D^0 (9 % compared with 4 %). On the other hand, the combinatorial background for this channel is significantly larger than for the $D^0 \to K^-\pi^+$ channel. Using the same rate assumptions as for the D^0 study, one expects 0.61 D^+ (or D^-) per event decaying into the $K^-\pi^+\pi^+$ (or charge conjugate) channel. The initial number of background combinations is about 2.5×10^{10} per central Pb–Pb event, but most of these can be eliminated before attempting to reconstruct vertices, using impact parameter and invariant mass cuts.

The selection strategy is very similar to the one used for the D^0 study. We select combinations of a kaon of a given sign and two pions of the opposite sign ($p_t > 500\,\mathrm{MeV}/c$). The sum of the squares of the $r\varphi$ projections of the impact parameters for the three tracks has to be larger than $3 \times 10^5\,\mu\mathrm{m}^2$. At this point the number of combinations per event has been reduced to 8×10^5, it is now possible to go to the more time-consuming vertexing step. We find the point of closest approach in space of the three decay tracks and retain only those candidates for which the sum of the squares of the distances from this point is less than $3.5 \times 10^4\,\mu\mathrm{m}^2$. Finally, two additional cuts are applied: the $\cos\theta_{\mathrm{point}}$ defined as above has to be larger than 0.99756 ($\theta_{\mathrm{point}} < 4^\circ$) and the separation between the primary and secondary vertices has to be larger than $600\,\mu$m. After this selection we end up with 2.8×10^{-3} signal and 0.09 background combinations per event, including reconstruction and PID inefficiencies, corresponding to a S/B ratio of 3.1 %. For a sample of 10^7 central Pb–Pb events the significance would be $S/\sqrt{B} = 30$. Here, too, in the absence of kaon identification, the significance would deteriorate by a factor 2.6.

11.10 Strangelets

11.10.1 Strangelet production at the LHC

QGP formation should result in an enhanced production of strange quarks and it has been speculated that droplets of strange matter (i.e. strangelets) could be formed in heavy-ion collisions (for a review see Ref. 95 and references quoted in Ref. 96). Strange matter may also appear in the form of exotic multi-hypernuclei (MEMOs) resulting from the coalescence of several strange baryons (e.g. Λ's). The formation of MEMOs does not require a QGP phase. Although those two forms of exotic matter would be nearly indistinguishable in their decay modes, the observation of either would clearly be of great interest.

In order to assess the capability of the experiment to recognize less conventional or unusual signals, we have investigated sensitivity to strangelet production. A similar analysis could be done for other hypothetical objects like 'free quarks' (Z = 1/3 or 2/3) or magnetic monopoles.

In heavy-ion reactions strangelets and MEMOs might be found in the final state as objects with baryon number $A \approx 2$–40, Z/A ratio ranging from ~ -0.5 up to $+0.5$, and fraction of strangeness within $f_s \approx 0.5$–1.5. Strangelets should be created preferentially in a region with large net baryon density. The phase space covered by ALICE ($-0.9 \leq \eta \leq 0.9$) is characterized by a low net baryon density and a chemical potential $\mu_B \approx 0$, thus conditions not favourable for strangelet formation (as opposed to strangelet production at large rapidities [98]). Recent theoretical developments [97] suggest, however, that strangelets could be produced also at the LHC as a result of local fluctuations in the net baryon number [12].

Strangelets and MEMOs could be stable or metastable objects, and their stability, lifetime, and decay modes are strongly parameter dependent [96].

Strangelets (S) may be

i) unstable ($\tau < 10^{-20}$ s), in which case they decay via hyperon emission (Λ, Σ, Ξ) and meso-nucleonic strong interaction processes [96]:

$$S \to S' + N + \pi$$
$$S \to S' + N$$

ii) metastable ($\tau < 10^{-4}$ s), in which case they decay via weak interaction processes:

$$S \to S' + N$$
$$S \to S' + \pi$$
$$S \to S' + e + \nu$$

iii) stable ($\tau > 10^{-4}$ s).

[12] As the average baryon and strangeness density at midrapidity is zero, strangelets and *anti*-strangelets would be produced in equal numbers.

11.10.2 Strangelet detection in ALICE

The search for strangelets within ALICE will depend on lifetime and decay modes.

i) If these objects are stable or long-lived metastable ($> 10^{-7}$ s), they will pass through the complete detector (assuming $\beta \approx 0.1$) and can be seen as tracks with unusual charge-to-mass ratio.

ii) If they are metastable with a lifetime below (10^{-7} s), they may be found, in principle, by identification of the decay products (π, K, p, Λ, ...) and secondary vertex reconstruction.

iii) If they are unstable and decay immediately ($< 10^{-20}$ s), identification is very difficult. However, correlations of decay products like Λ, K, $\overline{\mathrm{K}}$, π^- may still provide some information even in this case.

We have so far examined the sensitivity for a strangelet search in ALICE for cases (i) and (iii) [96], and the results are summarized below.

Stable or long-lived strangelets

In general, strangelets will have some non-integer value for the charge-to-mass ratio and can therefore be identified via $\mathrm{d}E/\mathrm{d}x$ and/or time of flight (TOF) versus momentum per charge (p/Z). As an example, we consider strangelets with $Z = 1$ and $Z = 2$ and a mass between 6 and 15 GeV (i.e. $|Z/A| < 0.3$). This mass range is of particular interest as lower mass strangelets are less stable (see Ref. 96) while heavier objects have lower production cross-section.

In the next section, S_i ($i = 1$–6) refer to strangelets with the characteristics shown in Table 11.11.

Table 11.11: Strangelet characteristics.

S_i	Z	Mass (GeV)
1	1	6
2	1	10
3	1	15
4	2	6
5	2	10
6	2	15

The mass stopping power $(1/\rho)\,\mathrm{d}E/\mathrm{d}x$, calculated analytically with the formula of Bethe–Bloch is shown in Figs. 11.61 and 11.62 for the TPC and the ITS, respectively. Assuming a 7 % and 10 % resolution in the TPC and ITS, one sees that strangelets may be resolved (with $> 3\sigma$ separation) from all other products over a large domain of p/Z. The additional use of TOF information would significantly improve the separation [96]; however, the TOF measurement has significant non-Gaussian tails (see Section 11.4) which might limit its usefulness in the search for very rare events.

We have evaluated the probability for normal stable particles to contaminate the strangelet energy-loss (E_S) distribution for S_1 in the TPC. Figure 11.63 shows the probability (as a function of p/Z) of different particles having an energy loss in the domain $\langle E_\mathrm{S}\rangle \pm 3\sigma$ by assum-

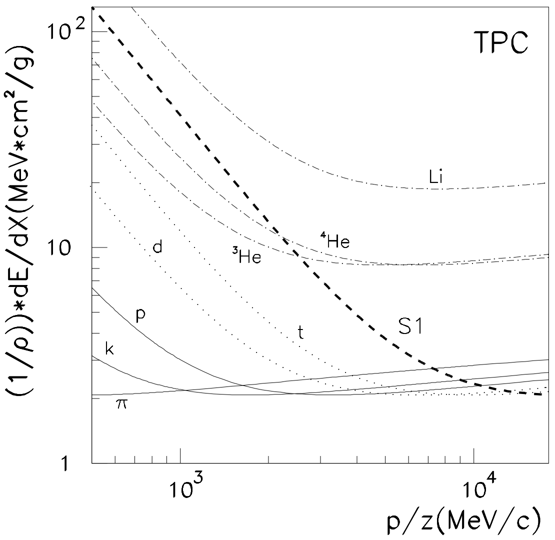

Figure 11.61: Mass stopping power $(1/\rho)\,\mathrm{d}E/\mathrm{d}x$ as a function of p/Z in the TPC for various particles and a strangelet S_1 ($Z = 1$, mass 6 GeV).

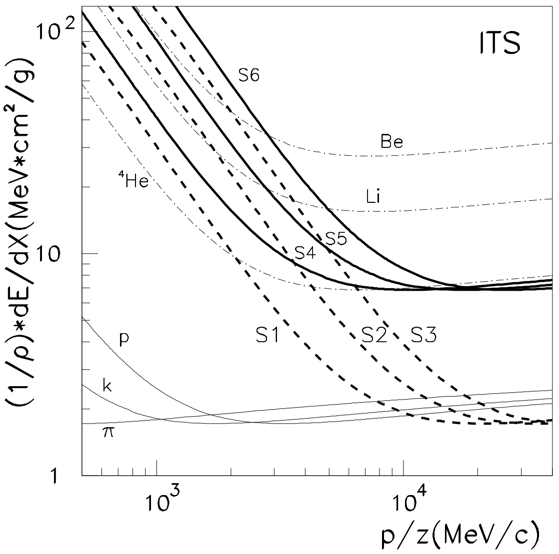

Figure 11.62: Mass stopping power $(1/\rho)\,\mathrm{d}E/\mathrm{d}x$ as a function of p/Z in the ITS for various particles and strangelets ($\mathrm{S}_1 - \mathrm{S}_6$ defined in Table 11.11).

Table 11.12: Yields of different particles for 5×10^7 central Pb–Pb events.

π	K	p	d	t	^{3}He	^{4}He	^{6}Li
5×10^{11}	1.25×10^{11}	10^{11}	5×10^8	5×10^6	5×10^6	5×10^4	5×10^2

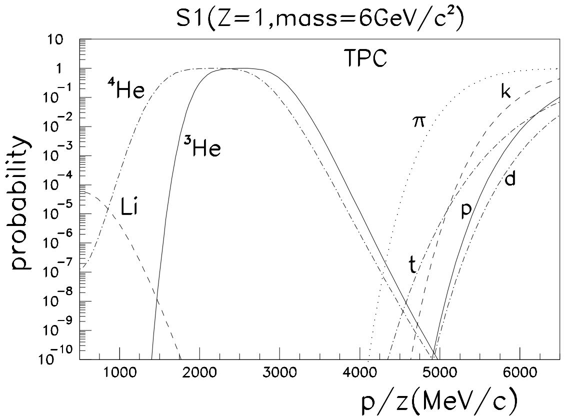

Figure 11.63: Probability as a function of p/Z for different particles in the TPC to have an energy loss in the domain $\langle E_{S_1} \rangle \pm 3\sigma$, where E_{S_1} is the strangelet S_1 energy loss.

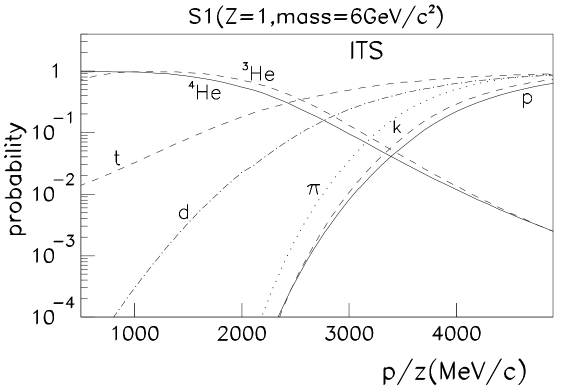

Figure 11.64: Probability as a function of p/Z for different particles in the ITS to contaminate the strangelet S_1 energy-loss domain ($\langle E_{S_1} \rangle \pm 3\sigma$).

ing a purely Gaussian shape with $\sigma = 7\,\%$ for the dE/dx distribution. Below about 5 GeV/c, it is possible to discriminate strangelets reliably from protons and deuterons, below 4 GeV/c from π, kaons and tritons. However, He isotopes start to contaminate the distribution below 4 GeV/c.

The same probability distribution for the ITS is shown in Fig. 11.64. Because the dE/dx measurement in the ITS has a non-Gaussian shape (Landau distribution), the results are worse, but might be still useful as a consistency check of the TPC energy-loss measurement.

The total number of particles that can fake a strangelet signal is shown in Table 11.12 for 5×10^7 central Pb–Pb events collected in ALICE in 10^6 s. SHAKER has been used to calculate π/K/p yields and the Quantum Statistical Model (QSM) [99] for light nuclei. All momentum distributions (including the one for strangelets) were assumed to be thermal.

The maximum number ($N_{\max}$, 99 % confidence level) of contaminating particles, i.e. particles that fall within the energy-loss range of S_1, is shown as a function of p/Z in Fig. 11.65 for 5×10^7 analysed events. In order to identify strangelets positively, a minimum number N_{Smin} has to be produced, such that N_{Smin} is larger than the number $N_{\max}$ of contaminating particles. The p/Z domains where strangelet identification is possible in the TPC ($N_{\mathrm{Smin}} - N_{\max} > 0$) are shown in Fig. 11.66 for different values of strangelet production, i.e. 10^{-1} to 10^{-4} strangelets per event within the acceptance. Stable strangelets are identifiable for a production probability per event of order 10^{-4} or more.

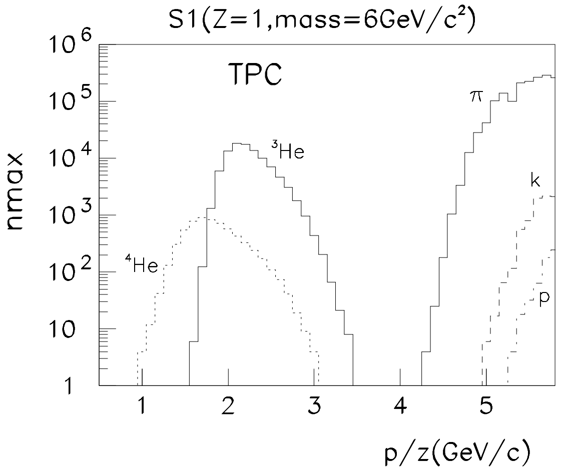

Figure 11.65: Maximum number ($N_{\max}$) of contaminating particles in the TPC as a function of p/Z, according to the probability values shown in Fig. 11.63 and the production rates given in Table 11.12, assuming thermal distributions.

The above limits depend on the assumption of a purely Gaussian shape of the dE/dx measurement, i.e. that the analysis is dominated by statistical rather then systematic errors. The ALEPH dE/dx resolution has shown no significant departure from a Gaussian curve down to the lowest measured level of 10^{-3} ($\sim 3\,\sigma$) [100]. Below this limit, systematic errors in the dE/dx measurement are difficult to predict and the final sensitivity can only be assessed

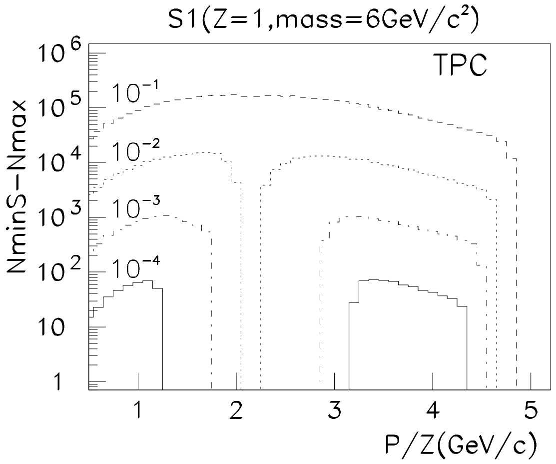

Figure 11.66: Difference between the minimum number (N_{Smin}) of strangelets S_1 and the maximum number of contaminating particles (N_{max}) as a function of p/Z for different assumptions on the production rate of S_1 per event (10^{-1} to 10^{-4}) for 5×10^7 central Pb–Pb events.

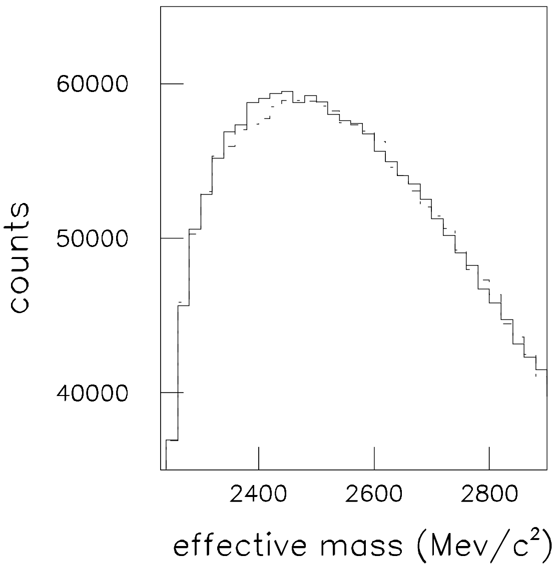

Figure 11.67: Simulated dilambda signal + combinatorial background ($S+B$) for 10^5 central Pb–Pb events (solid histogram). The dashed histogram represents the background obtained by fitting the ($S+B$) mass distribution in the region outside the dilambda peak.

with data from the actual experiment. The above limits would not be affected if the distribution stays Gaussian down to a level of 10^{-6}.

Results for the ITS which can be used only as a complement of the TPC are less favourable since the minimum measurable production rate increases to 10^{-2}.

Unstable strangelets

We have investigated the possibility to recognize unstable strangelets for the simplest possible object, i.e. the H-dibaryon or dilambda which can be viewed as a uuddss strangelet or a $\Lambda\Lambda$ MEMO. We used a hypothetical mass and width for the dilambda of $m = 2400$ MeV and $\sigma = 40$ MeV and a dominant decay mode into two Λ's. The efficiency necessary to reconstruct the $\Lambda \rightarrow \mathrm{p}+\pi$ decays has been taken from the simulations described in Section 11.8. From the roughly 250 Λ's produced per event within the acceptance, about eight can be reconstructed in the $\mathrm{p}+\pi$ decay channel.

The invariant mass distribution of all $\Lambda\Lambda$ pairs is shown in Fig. 11.67 for 10^5 events. The simulation includes the dilambda with a production multiplicity of 1.17 per event. The dilambda signal is visible in Fig. 11.68 after subtraction of the combinatorial background, whose shape can be estimated either by a fitting procedure [96] or by event mixing.

For a total statistics of 5×10^7 events, the dilambda particle could be observed if it is produced at a rate of about 10^{-2} to 10^{-1} per event within the acceptance.

Further details about simulation and extraction of the dilambda, including the treatment of statistical errors, can be found in Ref. 96.

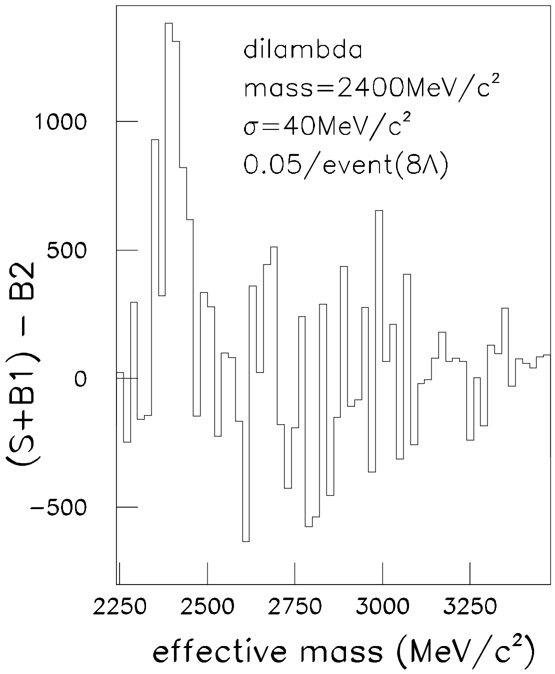

Figure 11.68: Background subtracted dilambda invariant mass distribution. The peak observed around 2400 MeV would correspond to the effective mass of the dilambda.

11.11 Event-by-event analysis

11.11.1 Introduction

Analysis of properties of a single high-energy collision and their variations from collision to collision is called event-by-event analysis. In this analysis the maximum available information concerning the dynamics of high-energy collisions can be used, in contrast to the inclusive analysis where only information on the 'average' collision is extracted.

An event-by-event analysis has been sucessfully used throughout the history of high-energy physics. It originates from the exclusive studies of the first data on proton–proton interactions using cloud and bubble chambers. It led to the discovery of fundamentaly new phenomena like the jet structure of high-energy hadronic and e^+e^- collisions by use of collective event observables, e.g. sphericity, planarity or thrust. The obvious basic requirement of the event-by-event physics is large phase-space acceptance of the experiment. Extremely high multiplicities and novel physical phenomena expected in ultrarelativistic nucleus–nucleus collisions open new dimensions for the event-by-event analysis.

The statistical resolution of the measurement of a single event observable is typically inversely proportional to $\sqrt{N}$, where N is the number of particles (or particle n-tuplets in the correlation analysis) used for the analysis in the single event. It is therefore clear that high-energy nucleus–nucleus collisions with extremely high multiplicities of produced particles allow a precise determination of many properties of a single event [101]. The study of directional flow due to mean field interactions in dense hadronic matter [102] is a well-known example.

If the collision energy is close to the critical energy at which Quark–Gluon Plasma (QGP) is created in an 'average' collision, the event-by-event analysis can be used to separate plasma from non-plasma events [101]. These two classes of events can occur at fixed collision energy due to statistical fluctuations in the nucleation of QGP [103] or its creation due to a percolation effect [104]. They can be distinguished by typical plasma signals like strangeness and entropy enhancement.

This large multiplicity will give us the possibility to:

- construct single-event particle spectra (study not only first but also higher moments of a distribution);
- determine thermodynamic quantities (temperature, chemical potentials) in a single event;
- use the Central Limit Theorem to study fluctuations of single-event properties in a model-independent way;
- select event subsamples on the basis of single-event properties for further 'semi-inclusive' analysis in each subsample.

11.11.2 Requirements for event-by-event analysis

The search for event-by-event fluctuations needs a good selection on the impact parameter to reduce variations in the collision geometry. This should be done without introducing a systematic bias in the studied observables. The best known experimental method for the 'unbiased' reduction of geometrical fluctuations is a selection of collisions according to the energy deposited in the fragmentation regions of the colliding nuclei. The description of the forward-energy trigger used in ALICE for the selection of central events is given in Section 9.3.

As discussed in the next section, there are various possible origins of event-by-event fluctuations in nucleus–nucleus collisions at LHC energies. In order to limit to a minimum the potential freedom in the interpretation of the data, it will be necessary to study the observed effects as a function of the collision energy and the volume of the created matter. The energy dependence can be obtained by comparison with the corresponding results obtained at the CERN SPS ($\sqrt{s} \approx 20$ GeV) and the BNL RHIC ($\sqrt{s} \approx 200$ GeV). The volume dependence, however, can only be obtained by studing collisions of nuclei with different mass numbers at the LHC.

11.11.3 Physics of event-by-event analysis

The physics of nucleus–nucleus collisions at the LHC will be dominated by studies of the pre-equilibrium partonic system created at the early stage of the collision, its evolution towards locally equilibrated QGP, and the properties of its transition to hadronic degrees of freedom. This also determines physics of event-by-event analysis in ALICE, as briefly discussed below. Further examples of event-by-event analysis can be found in Ref. 105.

Inhomogeneity of the early stage

Large fluctuations observed in $p(\overline{p})+p$ interactions at high energies (e.g. due to hadronic cross-section fluctuations) [106] may lead to a dependence of single-particle momentum spectra on the particle multiplicity in a collision [107], as observed, for example, at the Tevatron [9]. These fluctuations in the elementary processes should lead to an inhomogeneous structure of the system created at the early stage of the nucleus–nucleus collision. This inhomogeneity may, or may not, survive at the final state depending on the dynamics of the system evolution and its volume, controlled by the sizes of the colliding nuclei.

The internal inhomogeneity of the system created in central nucleus–nucleus collisions can be studied by the event-by-event analysis of the fluctuations of the sum-variable, defined as [108]

$$Z = \sum_{i=1}^{N} z_i \ ,$$

where $z = p_t - \langle p_t \rangle$ and summation runs over all particles detected in a single event. These fluctuations are studied in terms of the parameter α:

$$\alpha(\mathrm{AA}) = \frac{\sqrt{\frac{\langle Z^2 \rangle_{\mathrm{AA}}}{\langle N \rangle_{\mathrm{AA}}}} - \sqrt{\langle z^2 \rangle_{\mathrm{AA}}}}{\sqrt{\frac{\langle Z^2 \rangle_{\mathrm{NN}}}{\langle N \rangle_{\mathrm{NN}}}} - \sqrt{\langle z^2 \rangle_{\mathrm{NN}}}} ,$$

where indices NN and AA refer to values obtained in nucleon–nucleon interactions (or collisions of light nuclei) and central nucleus–nucleus collisions, respectively. By definition, $\alpha = 1$ for AA = NN. When the initial inhomogenity of the system disappears, the α parameter approaches zero [108].

This method requires data on nucleon–nucleon interactions (or collisions of light nuclei) as reference data with large inelasticity fluctuations and data on collisions of intermediate and heavy nuclei taken using the same experimental set-up. It was estimated that about 10^5 events in each sample are enough for the analysis [108], provided that the magnitude of the effect at LHC energies is similar to that observed at lower energies.

Quark–Gluon Plasma instabilities

A characteristic property of the QGP system is the possibility to observe collective effects due to colour degrees of freedom [109, 110]. Such effects can occur provided that the local colour neutrality (locally vanishing colour charge and colour current) is violated by random colour fluctuations. These fluctuations are strongly damped for a system close to thermodynamic equilibrium. However, at the early stage of high-energy nuclear collisions, the momentum distribution of partons is expected to be strongly anisotropic, far from equilibrium. In this case, as in the electromagnetic plasma [111], one expects the existence of unstable plasma modes with amplitudes exponentially growing with time. The dynamics of such a system is dominated by mean field interactions. The system behaves in a collective way, which should be reflected in event-by-event fluctuations. Note that the effects should be large (in contrast to nuclear flow) in central nucleus–nucleus collisions.

First quantitative estimates show that the plasma instabilities due to colour collective effects can, in fact, occur in heavy-ion collisions at LHC energies [110]. They should be observed as a directional flow of energy and/or particle number, studied event-by-event. Their systematic study could allow for extraction of basic properties of the QGP created at the early stage of the collision.

The plasma instabilities should cause a directional flow localized in rapidity which can be detected by the analysis of azimuthal event-by-event fluctuations. Full coverage in the azimuthal angle of the ALICE detector at midrapidity allows these detailed studies. The statistical methods for the analysis of collective flow from azimuthal correlations are well developed [112].

Transition from QGP to hadronic degrees of freedom

The expanding and cooling QGP is finally converted into hadronic degrees of freedom. This transition takes place at the critical temperature $T_C \approx 150$ MeV [113] (or below if supercooling of plasma happens). There is speculation that the evolution of the chirally symmetric QGP state to the ground state —where the chiral symmetry is broken— proceeds through formation of domains with classical field configurations (chiral condensate) [114–118]. These classical coherent pion fields are aligned in isospin leading to an imbalance between charged and neutral pions when decaying. The effect can be observed in non-statistical event-by-event fluctuations of the phase-space population of charged and/or neutral pions. The second characteristic feature of chiral condensates is coherent pion emission. This should lead to deviations from the Bose–Einstein correlations and to an enhancement in the pion spectra at very low p_t (< 100 MeV/c). The characteristic size of the domains which determine the magnitude of the observable fluctuations is still under debate [119].

Decay of domains with classical field configurations should be observed by large variations of the charged-to-neutral pion ratio and in modifications to the Bose–Einstein correlations. The ability of the ALICE detector to measure isospin fluctuations is discussed in Section 11.7, where the physics related to photon detection is presented. The analysis of the phase-space fluctuations of identical pions is done using methods developed for particle interferometry studies as described in Section 11.5.

'Thermodynamic' fluctuations

The large number of particles produced in high-energy nucleus–nucleus collisions, as well as the existing data on particle production, motivate the use of thermodynamic models. For central Pb–Pb collisions at the LHC the thermodynamic parameters of the created matter, temperature and chemical potentials, can be determined not only as averages over an event sample but also for each single collision.

It was recently suggested [120] that in this case the event-by-event temperature fluctuations can be related to the total heat capacity, C, of the matter by a basic relation of thermodynamics:

$$C^{-1} = \frac{V(T)}{T^2} ,$$

where T and $V(T)$ are a single-event temperature and its variance. Thus, the measurement of event-by-event temperature fluctuations gives an estimation of the heat capacity, an important parameter of the QGP.

Temperature fluctuations

Event-by-event fluctuations due to QGP instabilities and/or 'thermodynamic' fluctuations can be observed by studing fluctuations in transverse momentum (p_t) or transverse mass ($m_t = \sqrt{m^2 + p_t^2}$) spectra of produced particles. The m_t spectra of pions (one charge sign only) produced in a single 'average' Pb–Pb collision ($\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 2000$) as registered by the central detectors, are shown in Fig. 11.69. The spectra were simulated assuming that particles are emitted independently and that the m_t distribution is of the form:

$$\rho(m_t) = A \cdot m_t \cdot \mathrm{e}^{-m_t/T} ,$$

where A is the normalization factor and T is the slope parameter related to the event 'temperature'. The spectra for two different values of slope parameter ($T = 200$ and 250 MeV) are plotted in Fig. 11.69, serving as an illustration of the resolution of the ALICE detector with respect to single-event spectra.

The average transverse mass in the event is defined as:

$$\overline{m_t} = \frac{1}{N} \sum_{i=1}^{N} m_t^i ,$$

where N is the number of measured particles in the event, and m_t^i is the transverse mass of particle i.

From the Central Limit Theorem it follows that $\overline{m_t}$ is distributed according to a Gaussian distribution provided that the particles are emitted independently and N is fixed and large enough [121]. The mean of this distribution is equal to

$$\langle \overline{m_t} \rangle = \langle m_t \rangle ,$$

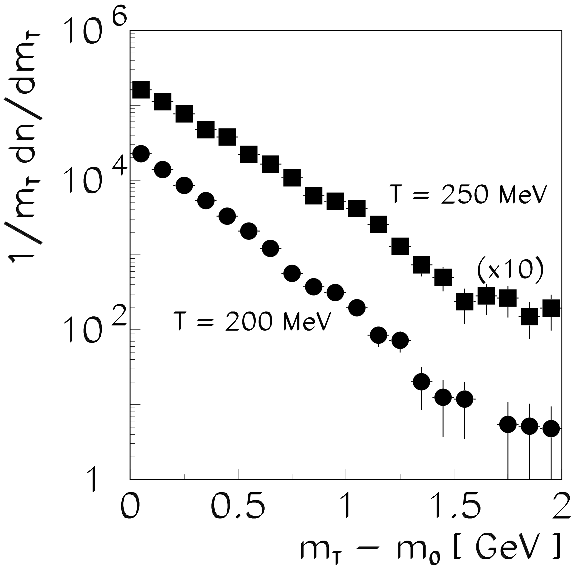

Figure 11.69: Single-event transverse mass spectra for 'average' Pb–Pb collisions ($\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y = 2000$) at the LHC as measured by ALICE for two different values of the slope parameter.

where $\langle m_t \rangle$ is the mean of the inclusive distribution. The variance of the $\overline{m_t}$ distribution is equal to:

$$V(\overline{m_t}) = \frac{V(m_t)}{N-1} ,$$

where $V(m_t)$ is the variance of the inclusive distribution. Thus, the form and the parameters of $\overline{m_t}$ distribution are completely defined by the Central Limit Theorem and the inclusive distribution. If the number of particles varies in the selected sample of events and the inclusive distribution changes with the multiplicity, we can define a normalized variable:

$$x = \frac{\overline{m_t}^{(N)} - \langle m_t \rangle^{(N)}}{\dfrac{V^{(N)}(m_t)}{N-1}} ,$$

where the index (N) tells us that the quantity should be calculated using events with the fixed particle multiplicty N. The x variable defined in the above way should have a Gaussian distribution with mean equal to zero and variance equal to unity, independent of the particle multiplicity.

Thus, the search for non-statistical fluctuations, e.g. due to the collective flow caused by QGP instabilities, can be done by studying the deviations of the x distribution from the Gaussian one with the parameters defined entirely by the inclusive spectrum. Of course, this kind of model-independent analysis, possible due to the large particle multiplicities, can be performed for any other relevant observable measured in ALICE.

Assuming the functional form $\rho(m_t)$ of the m_t spectrum, the slope parameter T of a single-event spectrum is related to the $\overline{m_t}$:

$$T = \frac{1}{4} \cdot \left(\overline{m_t} - 2 \cdot m + \sqrt{4 \cdot m \cdot (\overline{m_t} - m) + \overline{m_t}^2} \right) .$$

The variance of the T distribution can be estimated as:

$$V(T) = \left(\frac{\mathrm{d}T}{\mathrm{d}\overline{m_t}} \right)^2 \cdot V(\overline{m_t}) ,$$

where $V(\overline{m_t})$ is the variance of the $\overline{m_t}$ distribution:

$$V(\overline{m_t}) = \frac{1}{N-1} \cdot \left(\frac{m^3 + 3m^2T + 6mT^2 + 6T^3}{m+T} - \frac{(m^2 + 2mT + 2T^2)^2}{(m+T)^2} \right) .$$

For events with about 8×10^3 measured charged particles in central Pb–Pb and a value of the slope parameter $T = 200$ MeV, one estimates the typical resolution in T as

$$\sqrt{V(T)} \simeq 1 \text{ MeV} .$$

Therefore, different dynamical scenarios which lead to 'temperatures' differences of a few per cent [120] should be distinguishable in the event-by-event analysis. The selection of subsamples of events characterized, for example, e.g. by high values of the slope parameter, for further semi-inclusive analysis, will be possible.

12 Implementation

12.1 Test beam requirements

With the exception of the forward muon spectrometer, most of the particles within the acceptance of ALICE will have low or moderate momenta in the range between a few tens of MeV/c and 10 GeV/c. Consequently, we would be able to perform a large proportion of R&D and prototype tests at the energies available in the PS. If our requirements for low energy beams can be met, we could limit the use of the SPS, which is in heavy demand, to the minimum needed for a number of specific tests. In addition, we have made use of outside facilities available to us (in particular ITEP and DUBNA), and will continue to do so in the future, as far as possible.

12.1.1 Low-energy single particles

Many of the current R&D tests have been performed in the T11 zone at the PS. In the coming years, we intend to test a variety of small prototypes, in particular silicon drift and strip detectors, TOF counters (Pestov and PPC) and RICH prototypes. In addition, we intend to perform low-energy scans with the $PbWO_4$ crystals. The current beam line (limited to 3.5 GeV/c) and the available floor and counting room space are no longer adequate for this programme. In order to limit the influence of multiple scattering in tests of the silicon tracker and to cover the complete energy range relevant for the e.m. crystal calorimeter, the beam energy has to be increased to 10 GeV/c. More space (both in the test area and in the counting room) is needed to install and keep in place mechanical structures, services (e.g. gas systems) and electronics for the large number of different sub-detectors under test. We have started and will continue to equip the area with a common infrastructure for beam counters, tracking chambers, precision silicon telescope, DAQ, etc. This increases significantly the efficiency of the tests, but requires an area dedicated to ALICE.

12.1.2 High-energy proton beams

The PS is inadequate for testing the muon chambers and muon absorber: muon and hadron beams in the energy range 10 to $>$ 200 GeV will be required. In particular, the shape and composition of the absorber are crucial for the functioning of the muon arm, and a number of tests will be necessary in order to verify punch-through, side leakage and back-splash into the TPC region. So far, few experimental data exist concerning the side leakage of low-energy particles from a geometry like the one foreseen for our beam shield. While the R&D planning for the muon arm is still at an early stage, we anticipate a need for 6 weeks per year of proton beam time, starting with first tests of 4 weeks duration in 1996, e.g. in X1.

Currently, the $PbWO_4$ crystals are tested in the X1 beam line. Until the high-energy PS beam is available, we want to continue to test here for about 4 weeks per year. Later, some tests with higher energies (10–100 GeV/c electrons and pions) will still be needed at the SPS, about 2 weeks per year.

12.1.3 High-energy heavy-ion beams

The most severe challenge for ALICE is the high particle-density in ion collisions at the LHC. In order to test the performance of larger prototypes in a realistic environment, we have to make use of the SPS heavy-ion beam. Although the multiplicity $\mathrm{d}N_{\mathrm{ch}}/\mathrm{d}y$ at the SPS is an order of magnitude below that expected at LHC, we can still obtain comparable track densities (particles per cm^2) by choosing an appropriate distance from the vertex, in the fixed target geometry. Double-track resolution, pattern recognition and PID performance have to be tested in such an environment for the ITS, TPC, RICH/TOF and PHOS detectors.

In particular, in order to test PID capabilities, we will need a set-up resembling as closely as possible a sector of the ALICE detector, including at least a magnet, tracking devices and PID detectors over a solid angle large enough to contain several tens of charged particles in central Pb–Pb collisions.

12.1.4 Calibration

The only detector which needs extended periods of calibration is the $PbWO_4$ crystal calorimeter: a total of about 50 weeks will be required during 1999–2002. The calibration will be performed in the test area at the PS.

12.1.5 PS test-beam requests

Detailed requests concerning the ALICE PS test requirements (low-energy single particles) have been discussed with the relevant PS staff. A test area adequate for our purposes should be created, possibly in the course of a rearrangement under discussion in the context of other PS requests. This area should be available for use to ALICE in 1997. During 1996 we could, with certain limitations, remain in the current environment at T11, using it as an interim area dedicated to ALICE.

The beam time required at the PS will be about 20 weeks per year (see Table 12.1), including PHOS calibration, which will start around 1999.

Table 12.1: Test-beam requests (in weeks per year) at the PS, the West Area of the SPS (beam line X1), and the North Area (hadron and ion beams indicated separately for NA).

Beam	1996	1997	1998	1999	2000	2001	2002	2003
PS	20	20	20	20	20	20	20	10
X1	8	4	–	–	–	–	–	–
NA	–	4+2	7+3	7+3	7+3	7+3	6	4

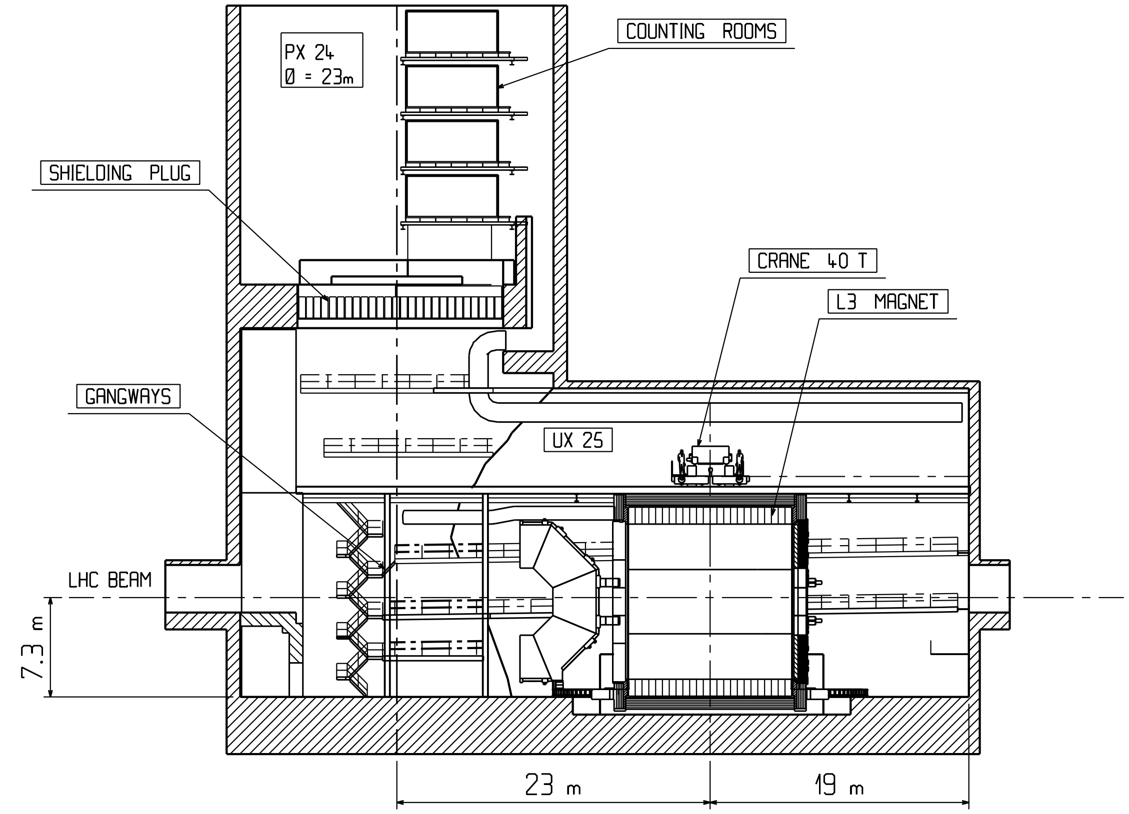

Figure 12.1: General layout of the existing underground structures at Point 2, showing the L3 magnet and counting rooms.

12.1.6 SPS test-beam requests

For high-energy and high-multiplicity tests, we will need a test area which can be served by high-energy hadrons as well as ions. If the primary beams in the west area are to be closed by the end of 1996, then only the NA beam lines H2, H4 and H8 are suitable. We request a test area in one of these lines, to be available for 1997. We should be able to group all muon-arm related tests — as well as the high-energy/high-density tests needed for the ITS, PID and PHOS— in this area. The beam time needed will be around 7 weeks per year with hadron and electron beams and at least 3–4 weeks per year during the scheduled ion periods. We will have to equip the area with a magnet and tracking devices capable of handling the multiplicity of Pb–Pb interactions.

The ALICE SPS test area and some of the equipment (magnet, silicon tracking telescope, DAQ) could be shared with a continuation of the WA97 heavy-ion experiment which has been proposed (SPSLC 95–43, SPSLC/M557). We have an understanding with the proponents of this experiment (most of whom are members of the ALICE Collaboration) that for a sufficient proportion of the ion-beam time, test activities would have priority over physics data taking. In addition, the silicon telescope would closely resemble a sector of the ITS, using pixel, drift and silicon microstrip prototypes developed for ALICE. Such a symbiosis would continue the very successful collaboration between R&D and physics experiments with, however, a clear shift in emphasis towards ALICE R&D.

In addition, over the next few years we plan to have some test activity around NA49 in the H2 beam line. The emphasis of NA49 will be on physics data taking, but

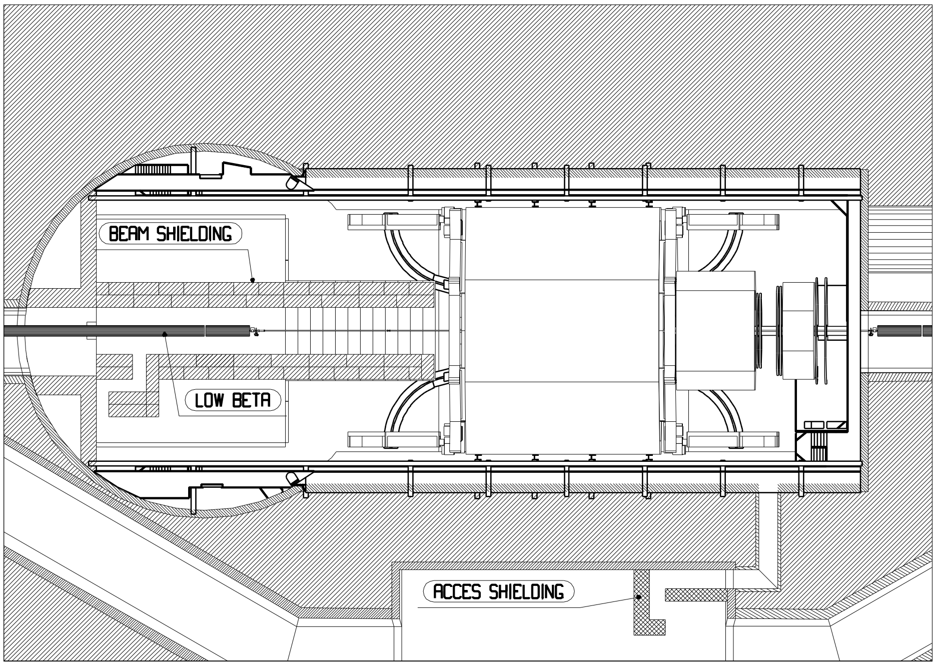

Figure 12.2: Shielding structure placed around the beam line.

some prototypes tested for ALICE preparation could also be used to enhance the physics capabilities of NA49. Currently it is envisaged that Pestov TOF walls of order 1 m^2 will be added to NA49 around 1998, depending on R&D progress and the final decision on PID in ALICE. TPC-related tests could also make use of the extensive infrastructure of and experience gained from this experiment.

12.2 Underground area and surface facilities

12.2.1 The Point 2 experimental area

The experimental areas of the LHC are being designed to the requirements of the detectors which make up the LHC experimental programme. The location of all LHC experiments has been discussed in several technical meetings and by the LHC committee. A decision has been taken to recommend installation of ALICE at Point 2 [1].

The Point 2 experimental area was designed for the L3 experiment. The main access shaft, 23 m in diameter, provides a 15×15 m^2 opening for the installation of magnets and detector units (see Fig. 12.1). The experimental cavern is 21.4 m in diameter and equipped with a 40 t crane. The crane runs very close to the L3 magnet (clearance ~ 0.5 m) and therefore limits the possibility of passing any larger objects across the magnet.

The L3 magnet provides a 11.6 m long and 11.2 m diameter 0.5 T solenoidal field. The end-caps have a door-like construction and can not be used to support the detector units. However, the door frame, the so-called 'crown', can be used to support the ALICE central detectors.

The counting rooms are installed in the 23 m access shaft, where they occupy about half of the shaft surface, leaving an installation passage of 15×7 m^2. The counting rooms are separated from the experimental area by a large concrete shielding plug. Most of the counting-room installations can conveniently be used by the ALICE experiment.

12.2.2 Modifications to Point 2

The Point 2 experimental area is ideally suited for the ALICE detector. Essentially no modifications would be necessary to the experimental cavern or the surface zone. However, the more severe radiation environment of the LHC will require reinforced radiation shielding and changes to the present access. The reinforcement of the shielding is not related to the ALICE experiment, since even without any experimental installation at Point 2 the shielding would have to be upgraded to LHC standards.

According to the dose calculations reported in Section 12.3.8, both the plug in the 23 m access shaft and the wall separating UX25 (the experimental area) from US25 (the machine bypass area) include sections with insufficient shielding. Since it is relatively difficult to modify or increase these sections, it is proposed to add a shielding layer close to the beam line as indicated in Fig. 12.2. It is only necessary to shield the beam-line section on the access shaft side (PX24) of the magnet, since the muon arm side is self-shielded by the absorber.

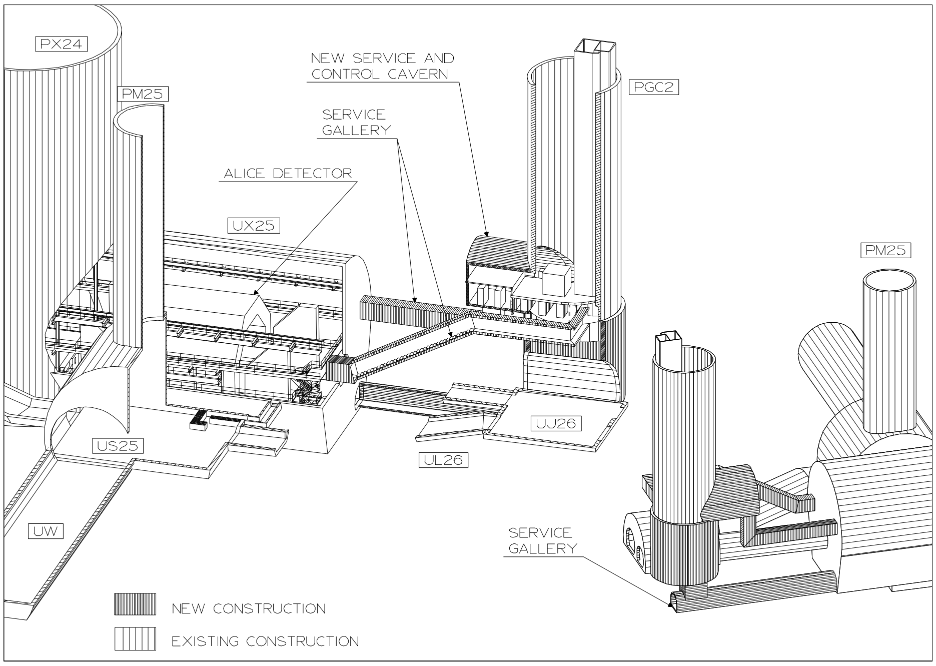

Figure 12.3: The layout of the PGC2 access shaft. New constructions are shown in black.

At present a second access to the Point 2 experimental area is provided via a chicane in the wall separating UX25 and US25. This access is placed at beam height on the non-shielded side of the L3 magnet and will, therefore, have to be modified. The modification can be made by simply adding a second 'leg' to the existing chicane inside the US25. In order to provide installation and services access to the muon arm and at the same time an additional emergency exit, it is proposed to construct a new access tunnel between the existing PGC2 access shaft and the back wall of the cavern, as indicated in Fig. 12.3.

The re-use of the PGC2 access shaft provides a cost-effective way of ensuring adequate and safe working conditions in the ALICE experimental area. In order to limit the modifications and simplify the radiation shielding one could prohibit use of the access shaft during LHC operation. The PGC2 shaft would be equipped with a standard lift and staircase module, and a free opening of $3.2 \times 4\,\mathrm{m}^2$ for services and detector installation. A more elaborate modification could include a permanently accessible underground service cavern at the base of the PGC2 shaft. The necessity of such a modification can only be evaluated once the design of the muon spectrometer has been finalized.

The capacities of the infrastructure installations at Point 2, such as cranes, power, cooling and ventilation systems all meet the requirements of the ALICE experiment. The 40 t crane in the experimental area is of limited use for the ALICE detector and will essentially be used to move the shielding units. It is, therefore, proposed to install a new 16 t crane, which would have a 2.5 m clearance over the L3 magnet.

12.2.3 The surface zone

The PID detector, the dipole magnet and the muon absorber are either too large or too heavy to transport by road and must, therefore, be assembled and tested at Point 2. The present surface zone at Point 2 includes sufficient assembly hall space and storage areas (the 1350 m^2 SXL2 assembly hall and the SXS storage platform) to meet the requirements of ALICE and no new hall construction will be necessary (see Fig. 12.4). The access shaft is covered by a 1790 m^2 building (SX2) equipped with a 63 t gantry crane. The ALICE Collaboration also requests the use of the SA building, which can be conveniently converted into a 'clean-room' area.

The ALICE data taking rate will require the removal of a large number of tapes from a tape storage area to the online computer room installed in one of the four counting-room modules. The tape storage area (about 500 m^2) will be set up inside the SX2 building, as close as possible to the counting rooms.

A new building would have to be constructed over the PGC2 shaft, since the present 'rain-cover' is not adequate for frequent use of PGC2. The building would have a surface area of $15 \times 25\,\mathrm{m}^2$ to accommodate an unloading area and would need to be equipped with a 16 t crane.

According to the ALICE construction schedule (see Section 13.1), it would be necessary to provide some hall space at Point 2 early in 1999, in order to receive the first elements of the ALICE detector.

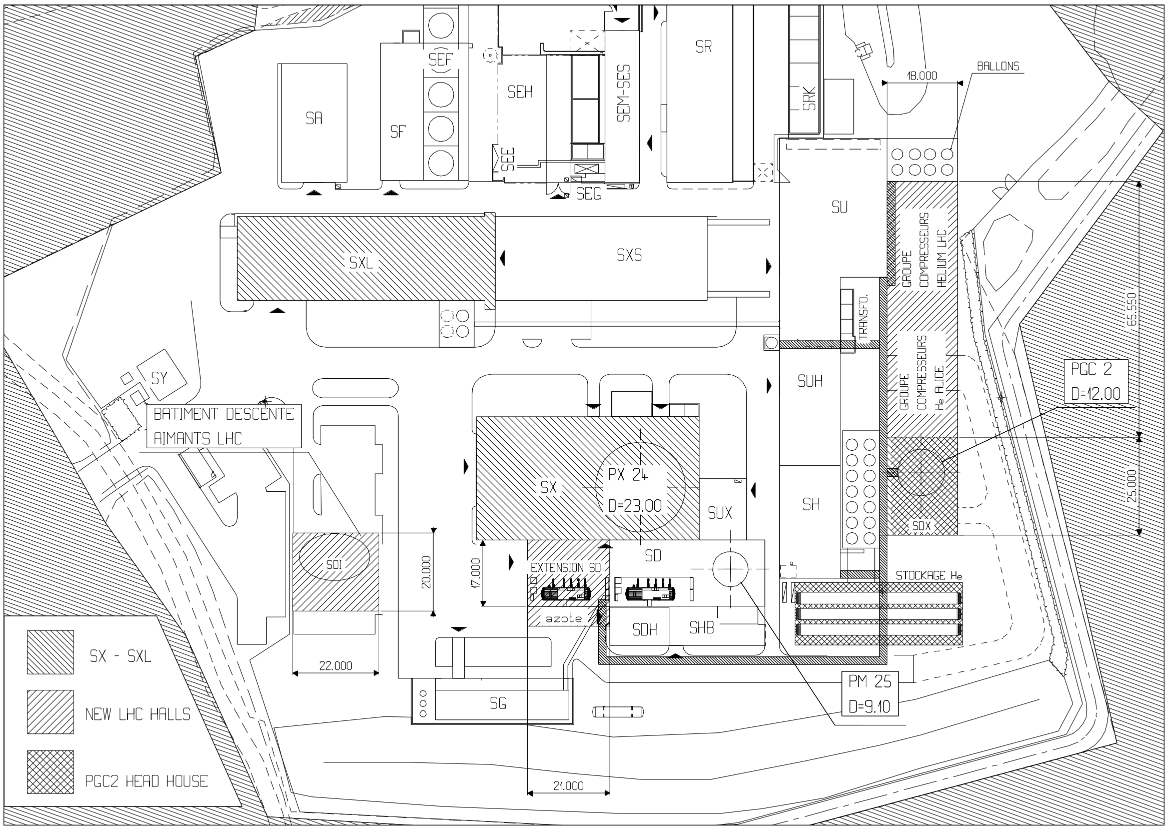

Figure 12.4: The Point 2 surface zone indicating the new building covering the PGC2 access shaft.

Table 12.2: Dismantling of the L3 experiment.

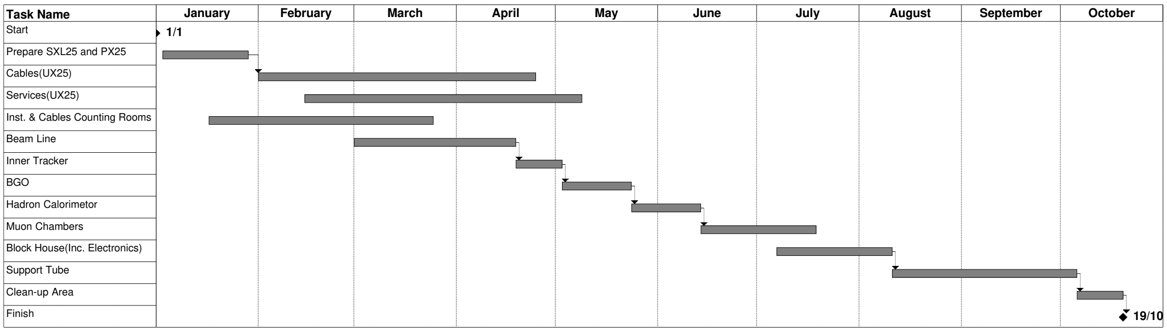

12.3 Integration and assembly

12.3.1 Dismantling the L3 detector

As mentioned above, the ALICE detector will be installed in the present L3 experimental area using as much as possible of the L3 infrastructure and, in particular, the L3 magnet and counting-room installations. Clearly, the removal of the L3 detector systems can only take place once the LEP operation has ended and the L3 Collaboration has taken the necessary precautions to safeguard all relevant information for the final data analyses.

It has been estimated that it will take 10 months to remove the L3 detectors, leaving the L3 magnet and the existing counting rooms in place (see Table 12.2). It is proposed to cut the central 32 m long support tube into 5 sections, which would allow each section to be removed by the existing crane in the SX2 hall. The hadron calorimeter would also be separated into 5 separate wheels, which could be brought to the surface using the existing cranes.

At this stage it is not envisaged to save the L3 cable installation, since it is considered unlikely that the lengths

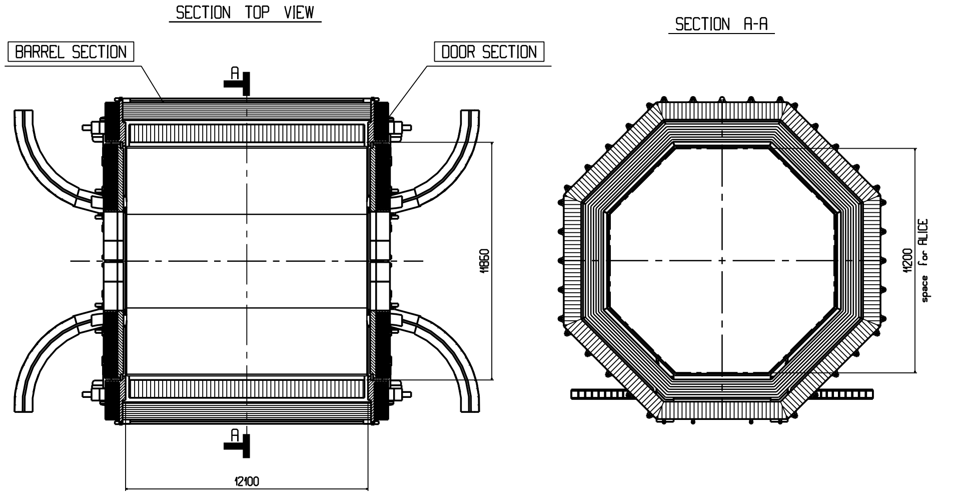

Figure 12.5: The L3 magnet.

and types of cable would be suitable for ALICE. When more detail is available about the cable requirements of ALICE an attempt will be made to save as many cables as possible. Most of the fixed gas and hydraulic pipe installations can be reused by the ALICE detector.

12.3.2 The L3 magnet

The L3 magnet has a central field of 0.5 T, developed at a rated current of 30 kA, and consists of:

- A magnetic yoke made of low carbon steel divided into a barrel and two poles. Each pole includes two doors to give access to the magnetic volume.
- An aluminium coil cooled by 2 independent circuits welded onto the outer and inner surfaces.
- A power supply connected to the magnet by 6 water-cooled busbars.
- A polarity inverter, integrated with the power supply.
- A water-cooled thermal shield placed on the inner surface to protect the magnetic volume from the heat flux of the coil.
- A monitoring device situated in the magnet control room (close to the existing counting rooms).

The main parameters of the L3 magnet are listed in Table 12.3.

Table 12.3: Main parameters of the L3 magnet.

Inner radius of coil	5930 mm
Width of conductor	890 mm
Outer radius of yoke	7900 mm
Total length of coil	11 900 mm
Total length of yoke	14 100 mm
Total height across flats	15 800 mm
Electrical power (at 0.5 T field)	4.2 MW
Central field	0.5 T
Stored energy	150 MJ
Rated current	30 kA
Current density in conductor	55.5 A/cm^2
Demineralized cooling water	150 m^3/h
Coil weight	1100 t
Yoke weight	6700 t

The magnet structure

The magnet structure surrounding the coil is inclined by 1.39%, following the LEP beam line, and is supported by a concrete cradle integrated in the UX25 hall foundations. The magnet consists of a barrel and two poles as shown in Fig. 12.5.

The barrel part consists of 96 bars arranged in an octagonal geometry. The 12.4 m long bars (cross-section 865 × 440 mm^2) are made up of 40 mm-thick steel plates tack-welded together to form individual masses of 36 t. There are eight corner pieces, designed to transmit forces from one wall to another. The two poles of the magnetic structure are made of self-supporting skeletons (open frames) comprising 1080 t of steel pieces made of welded heavy plates, giving the required mechanical rigidity and serving as a support and as a reference frame for installing the filling material.

The filling material (5600 t in total) provides the mass needed for the magnetic flux return, both in the two poles and in the return yoke (barrel). The filling material is made up of 50 mm-thick steel plates cut to the desired shape and tack-welded together to form individual masses of 15 t. Rails are welded onto each side of the open frames to guide the pack of plates making up the filling material.

Each pole consists of two parts, the crown and the double doors. Both of these parts are made of open frames bolted (and sometimes also welded) together and positioned with expansion keys. The crown forms a complete ring and each door forms a half ring. The frame elements of the doors are also joined by welds, made at the CERN site after assembly. Sixty-three axial tie bars connect the two poles in order to obtain the required axial rigidity of the magnetic structure, to stabilize the crown with respect to the magnetic forces, to help maintain the open frames during assembly and to press the two poles firmly against the barrel. These tie bars are 13 m long and are hydraulically pre-stressed at 50 t each.

Each half door, which when filled has a mass of ~ 340 t, is placed on a grease pad system installed under the centre of gravity and can pivot on large hinges. When opening the doors, the grease pad moves on a bridge-like structure, which is part of the crown, spanning the assembly pit.

Conventional hinges connect each door to the crown. These hinges can be mechanically disconnected from the doors, using self-locking hydraulic jacks to prevent overstressing the hinge pins due to magnetic forces.

The magnet coil

The coil has a form of a solenoid with 168 octagonal turns, each with a conducting section of 540 cm^2. The 168 turns are electrically connected in series. The turns are built from prefabricated components (sectors, half sectors). The components are assembled using an electron-beam welding procedure and distributed into 28 packages of 6 turns, each with a mass of about 40 t. The packages are electrically connected at their upper parts by connecting bars welded manually (MIG process). Each connecting bar has a dedicated cooling circuit (27 in total). Figure 12.6 shows the basic layout of the coil.

On the inner and outer surfaces, two cooling circuits, each consisting of a hexagonal aluminium tube are welded directly onto the turns. The tubes are connected to each other by an aluminium tube of ∅ 20/25 mm. Since a hydraulic circuit (6 turns) includes 3 turns, each package has 2 external and 2 internal cooling circuits (112 in total). Each hydraulic circuit is independently connected to a main water manifold (16 in total) via a 1 m long insulated passage in the iron yoke. The manifolds are placed on each side of the magnet. Each water circuit includes measurement of three quantities:

- flow (25 l/min for the coil, 14 l/min for the connecting bars);
- pressure (15 bar inlet, 3.5 bar outlet);
- temperature.

The coil is subject to a 50°C temperature rise during operation, causing a 20 mm axial dilatation. The coil packages are mounted on bronze plates in order to slide on the coil rails which are fixed to the yoke barrel.

The coil is mounted between fixed stops on the Jura-crown side and elastic air cushions on the Pit-crown side. These air cushions create a 320 t restoring force on the coil and act with a constant 40 t force on each coil octant.

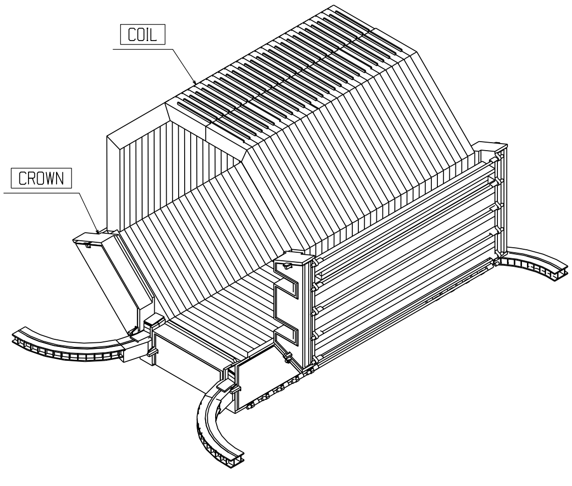

Figure 12.6: Basic layout of the coil.

The magnet power supply and busbars

The coil is powered by a 30 kA power supply. The thyristor power converter is installed in building SX25 and is powered from the CERN 18 kV network. The busbars are connected to the output terminals power supply via polarity inverters.

The busbars are made of rigid aluminium tubes stretching from SX25 to the underground hall UX25. The system comprises 6 positive and 6 negative conductors. Each pair of 6 bars is connected in series to form two hydraulic cooling circuits. The 6 positive conductors are connected in parallel via flexible shunts to the input piece and the 6 negative conductors to the output piece. Each input/output piece is bolted to the coil and has its own cooling circuit. The 4 hydraulic circuits are connected to the main magnet manifold.

The thermal shield

Most of the detectors inserted in the magnet are temperature sensitive. Therefore, an active thermal shield is inserted between the magnet coil and the first layer of detectors to make the magnetic volume thermally independent of the magnet. Experience from the L3 detector shows that the temperature will remain stable within the magnet volume to within 3–4°C, with the exception of some 'hot-spots', where the local temperature deviation can reach 5–6°C.

The shield surrounds the 1000 m^3 of the magnet volume and is directly attached to the 168 turns of the main magnet coil by insulating bolts. The cooling jacket is protected by a 50 mm continuous insulation sheet. The in-

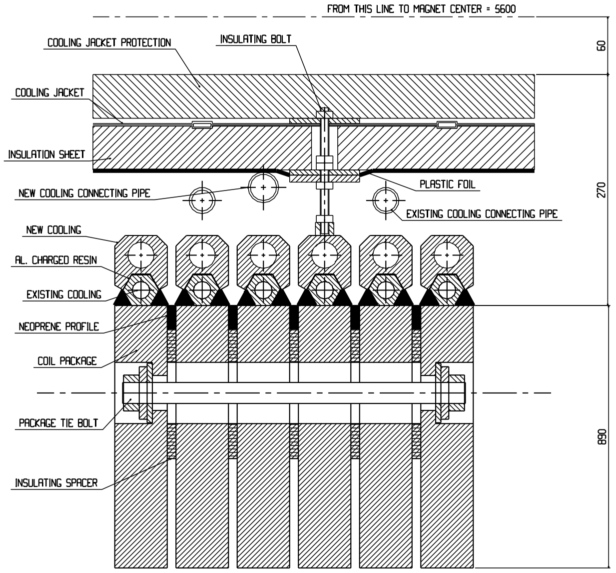

Figure 12.7: Layout of the new cooling circuit.

sulation on the top part of the system is protected against water loss by a thin sheet of plastic.

The active part of the system is made of a heat exchanger manufactured by rolling 2 aluminium plates homogeneously together. The inner embossing was created by tracing the circuit on one plate with a carbon paste. No welding or gluing was required.

The assembled heat exchanger has a total surface area of 450 m^2. The 192 plates of the heat shield are connected in 16 blocks of 12 units each. Each unit forms one cooling circuit (16 in total). The circuits have an independent manifold with a water inlet temperature of 15°C and a pressure of 8 bars.

Because of the high mass of the magnet coil and the risk of water circuit failure, an emergency circuit is installed, which the event operates in two independent modes. The first is the automatic replacement of the main pump by an auxiliary pump and the second is the discharge by gravity of a reserve of 20 m^3 (about 4 h) of demineralized water in the heat shield.

12.3.3 Modifying the L3 magnet

The main concerns regarding a further 15 years operation of the L3 magnet are all connected with the cooling system. The weak points in the cooling circuit are the 4000 welds connecting the hexagonal pipes (MIG process) and the welds of the cooling system pipework which cover the connection between the turns and the manifold (TIG process). Corrosion can be caused by the cooling fluid (demineralized water) at specific points in the welding joints (cracks, pores or cavities). To prevent this, precautions have been taken in the selection of materials, the cooling fluid is tested periodically to maintain a constant pH, and as far as possible the water circulation is not interrupted during shut-down periods.

In 1991, during modification work on the busbar system, extensive corrosion was found close to one of the mechanical flanges in the water-cooling circuit. In one place the tube had been perforated by the corrosion process.

Investigations have been launched to ascertain the condition of the existing cooling circuits, which have been operating since August 1988. It is planned to set up a pilot test station to fully analyse the ageing of the cooling system. In parallel with this, a study has been made of the addition of a new cooling circuit to the inner surface of the coil. The new circuit will only be used if one or more inner or outer coil circuits fail. The final decision to add the new cooling circuit can only be made once the L3 detector has been removed and a detailed *in situ* investigation of the cooling system has been made.

One solution is to weld a new conductor on top of the old one, but this would imply a large heat transfer (250°C) to the neoprene profile, which is compressed between the coil turns and the adjacent Vetresite insulating spacers. We will investigate whether these materials can stand such high temperatures. Another solution would be to fix the new conductor to the old one with a metal charged resin which has a good thermal conductivity (see Fig. 12.7). Tests are being performed to choose a suitable resin and a prototype (scale 1:1) will be developed to determine the feasibility and the appropriate technical solution. Tests will be carried out with the prototype to ascertain the possibility of cooling the coil using only this additional circuit.

This new cooling system, consisting of 5000 m of extruded aluminium section, will be connected in parallel to form two series of three-coil turns. Therefore each coil package will have one cooling circuit (28 circuits in total). A pipework system will connect these circuits to a new manifold. To be able to assemble this new cooling circuit, the existing pipework system and the heat shield will have to be modified.

We are currently also investigating the possibility of reducing the opening in the magnet doors to the minimum (i.e. 12.5°) needed for the acceptance of the muon arm. The additional iron plugs, with a total weight of about 100 t, would effectively reduce the radial field components in the L3 magnet by a factor of two in the TPC region [2] and therefore reduce the track distortions.

12.3.4 Maintenance of the L3 magnet

The L3 magnet system has been operated for 6 years (about 25 000 h) at full power and is regularly serviced and maintained in order to ensure its long term reliability.

Nevertheless, and in addition to the modifications discussed above, the magnet will have to undergo extensive maintenance after the end of LEP operation, during the installation shut-down, in order to guarantee its continued operation for the ALICE experiment. The power supply, the cooling circuit, the interlock system and the magnet control system will have to be serviced and exchanged or upgraded if necessary.

Coil short circuit

A coil short-circuit occurred accidentally during installation, probably as a result of welding electrodes which were trapped between the barrel and the outside coil, and is located approximately between octant 2 and 3. In order to reach this part of the coil about 36 iron beams would have to be removed. The risk and cost of such an operation cannot be justified and, therefore, is not proposed by the ALICE Collaboration.

Coil manifold

The coil manifold will require a complete revision. In 1994, during the installation of the toroid magnet cooling system, it was discovered that the check valves installed in every coil circuit no longer maintained the water pressure. The internal mechanism had started to show signs of ageing.

During shut-down periods it is frequently necessary to stop the circulation of the cooling water for relatively long periods, for maintenance or modification. At start-up, either several flow meters remain blocked or the flow indication is altered. In the first case we have to disassemble and clean them (this involves partially draining the manifold) and in the second case we have to readjust the setting.

The busbars

The busbars were built using an aluminium tube (∅ 140/120 mm) from Russia. To reduce the internal cross-section, in order to increase the water flow rate, special plastic pipes were introduced along the whole length and held in position with conical wedge.

During the 1991 shut-down all busbar columns were drained in order to install individual retaining valves. After the modification it was not possible to restore the water circulation to one of the cooling circuits. It was found that one column was blocked and that only reverse flow was possible.

12.3.5 Installation

The ALICE experiment can be divided into three detector sub-structures, all of which place different requirements on the installation and integration arrangements: the ‘light’ detectors (ITS, TPC and PID), the heavy PHOS detector and the muon arm (see Fig. 12.8).

The overall installation procedure has the following sequence (see also Fig. 12.9 a–f):

1. installation of a temporary beam structure through the L3 magnet;
2. transportation of the dipole magnet and the muon filter through the L3 magnet;
3. installation of the large support beams and the PHOS detector;
4. installation of the central support frame and the ITS, TPC and PID detectors;
5. installation of vacuum chamber and absorbers;
6. installation of beam shielding.

The dipole muon magnet and the muon filter are placed on a common rail system at the far side of the L3 magnet (opposite PX24). They can be installed as complete units via the 23 m access shaft and transported through the L3 magnet. The weight of the proposed magnet is about 850 t, which makes it necessary to assemble and test it in one of the Point 2 assembly halls. The complete magnet can then be installed in the experimental area using a temporary crane. The present SX2 building is already equipped with a removable roof over the access pit, which permits the temporary crane to operate outside the building.

The front absorber, the first 6 planes of muon chambers (TC1 to TC6), and 7 m of the small angle absorber, are supported by the dipole magnet. The remaining part of the absorber is supported by two separate pillars. TC7 to TC10, MC1 and MC2 are all supported by a mobile frame.

The dipole magnet is placed outside the L3 magnet door at a distance of 9750 mm from the magnet centre, which avoids a costly rebuilding of the magnet door. The PHOS detector must be supported by 12 m long beams attached to the door supports (the ‘crown’) of the L3 mag-

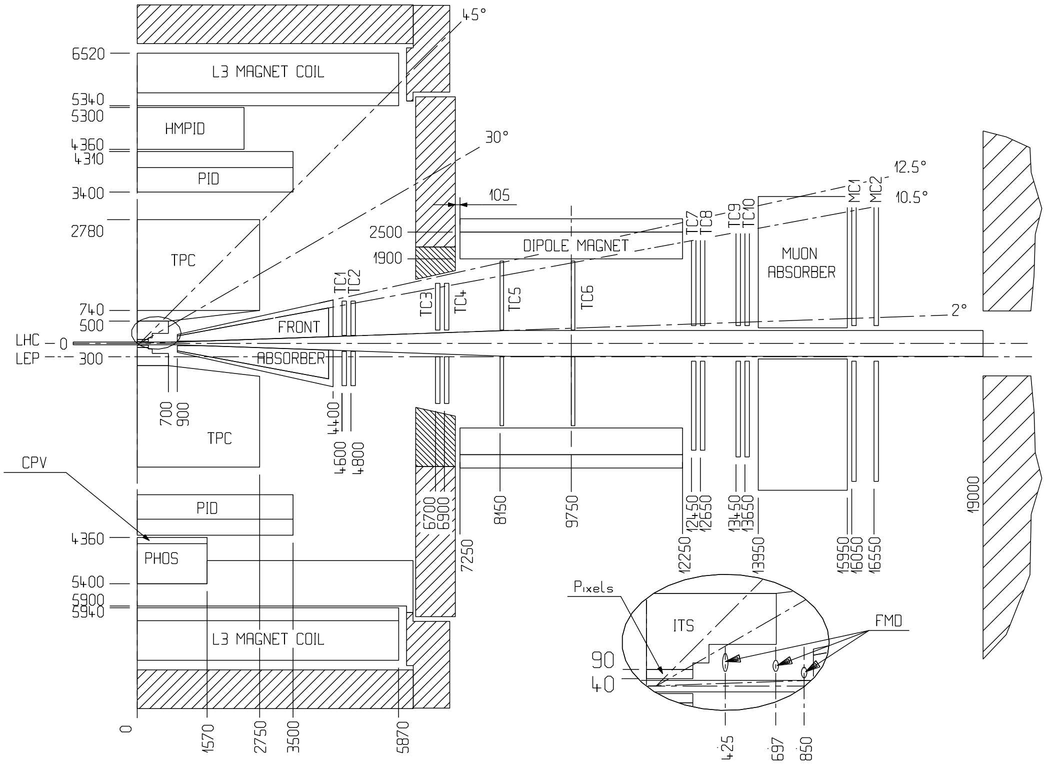

Figure 12.8: Conceptual layout of the ALICE detector with a forward muon spectrometer.

net, since the coil of the L3 magnet does not permit any intermediate support points inside the magnet volume. The PHOS detector is divided into four identical sub-units, each weighing 10 t. The sub-units will be placed on rail extensions outside the magnet and pushed into place using hydraulic pistons.

The ITS, TPC and PID will all be supported by a common frame (space frame), which can be constructed and tested on the surface and brought into the L3 magnet as a complete unit. The space frame will be placed on a temporary rail system outside the magnet and pushed onto a permanent rail system inside the magnet, where the complete structure can be easily aligned. Since the ITS is supported directly by the TPC, it will also be possible to install the ITS detector separately. The space frame for the ITS, TPC and PID will be placed on a dedicated rail system integrated with the PHOS support beams. Calculations show that a movement of 3–4 mm will be transmitted to the space frame of the central detector during the removal of the PHOS units.

In this scenario the large and heavy units of the muon arm are trapped behind the L3 magnet and cannot be removed without removing all the central detectors. Installation of the muon arm on the access shaft side would make it very difficult to remove the central detectors and is, therefore, not at present considered to be an option.

12.3.6 Installation schedule

An overall schedule for the installation of the ALICE detector in the Point 2 experimental area has been developed according to the ALICE installation scenario presented in the previous section (see Table 12.4). The installation schedule is compatible with a 3 year shut-down. It includes the decommissioning and removal of the L3 experiment and a 3 month period for comprehensive field measurement of the combined L3 magnet and new dipole system.

12.3.7 Installation of ZDC

The ZDC will be installed in front of the D2 machine magnet at a distance of 93.1 m from the Point 2 intersection region, as indicated in Fig. 12.10. The calorimeter units will be placed on a vertically mobile support, which will make it possible to lower the complete ZDC by about 50 cm during injection and pp operation. In order to reduce the amount of material in front of the detector, it is proposed to install a conical, oval shaped vacuum chamber between the D1 magnet and the ZDC detector. Preliminary calculations show that the windows in front of the neutral detector and in front of the charged detectors would have a thickness of 2 and 4 mm of stainless steel, respectively.

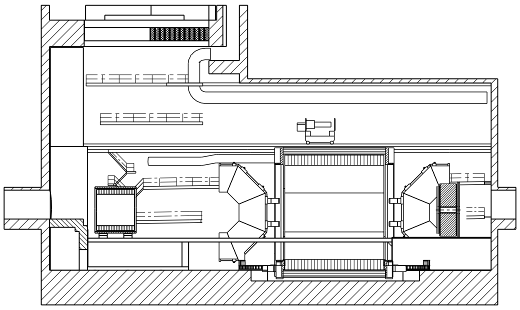

(a) Installation of temporary beams for the transport of the dipole magnet and the muon filter through the L3 magnet.

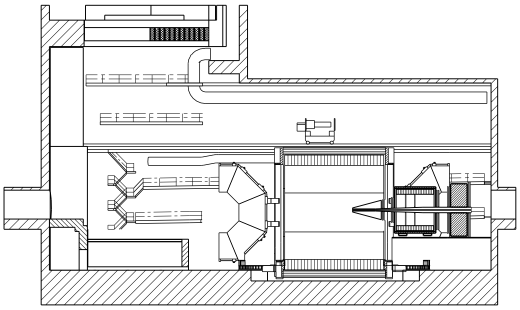

(b) Muon arm in place.

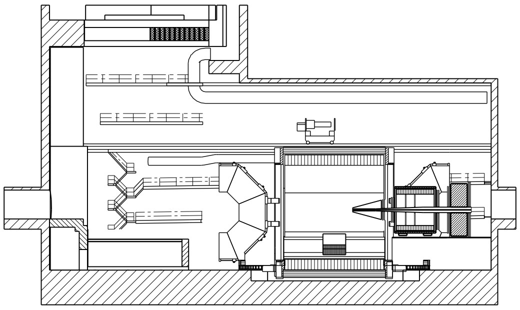

(c) Installation of the PHOS detector.

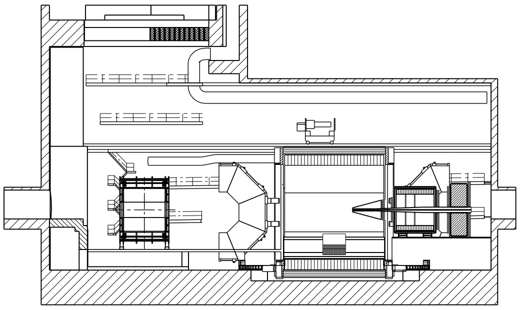

(d) Installation of the support frame carrying the ITS, TPC and PID.

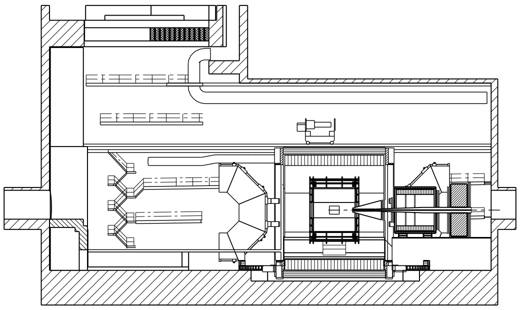

(e) Installation of vacuum chamber and absorbers.

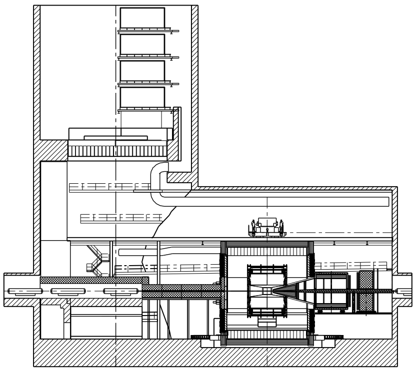

(f) Installation of beam shielding.

Figure 12.9: Overall installation sequence.

Task Name	Duration
CIVIL ENGINEERING	62w
L3 DECOMMISSIONING	52w
L3 DISMANTLING	40w
MUON ARM RAILS BASE	6w
LOW ß SUPPORT STRUCTURE	5w
MECHANICAL MODIFICATIONS	10w
COOLING SYSTEM UPGRADE	10w
MAGNET TEST	4w
ASSEMBLY IN SXL25	10w
CONNECT SERVICES	12w
TEST IN SXL25	10w
OPEN MAGNET DOORS	3d
INST. MUON ARM RAILS	2w
INST. TEMPORARY RAILS	1w
POSITION 1000t CRANE	1w
INST. MUON FILTER	2w
INST. DIPOLE MAGNET	2w
REMOVE TEMP. RAILS	1w
REMOVE 1000t CRANE	1w
CLOSE DOORS PGC2 SIDE	2d
COMPLETE MUON ARM RAILS	1w
POS. MAGNET & MUON FILTER	1w
POSITION MUON FILTER	2w
DESIGN MEASURING DEVICE	20w
TEST MEASURING DEVICE	8w
INST. IN MAGNET	2w
CLOSE DOORS(PX24 SIDE)	2d
FIELD MAPPING	10w
OPEN DOORS(PX24 SIDE)	2d
REMOVE MEASURING DEVICE	1w
INST. CRYO. COMPRESSORS	6w
INST. CRYO. LINES TO SXL	6w
INST. COLD BOX IN SXL	2w
TEST CRYO. SYSTEM IN SXL	4w
INST. CRYO. LINES IN PGC2	8w
TRANSFER COLD BOX TO PGC2	3w
INST. CRYO LINES TO MAGNET	2w
TEST SYSTEM	4w
INST.SPACE FRAME RAILS	4w
INST. TEMP. RAILS	1w
INST. PHOS	4w
INST. SPACE FRAME	2w
REMOVE TEMP. RAILS	1w
INTERNAL SERVICES	3w
CLOSE DOORS	2d
RADIATION SHIELDING	1w
EXTERNAL SERVICES	4w
OPEN MUON ABSORBER	2d
ROLL MAGNET BACK	3d
INST. SMALL ANGLE ABS.	3w
INST. TRACK CHAMBERS TC1/6	3w
CLOSE MAGNET	2d
INST. TRACK CHAMBERS TC7/10	2w
INST. MUON CHAMBERS	3w
CONNECT SERVICES	5w
TESTS AND RUNNING IN	22w

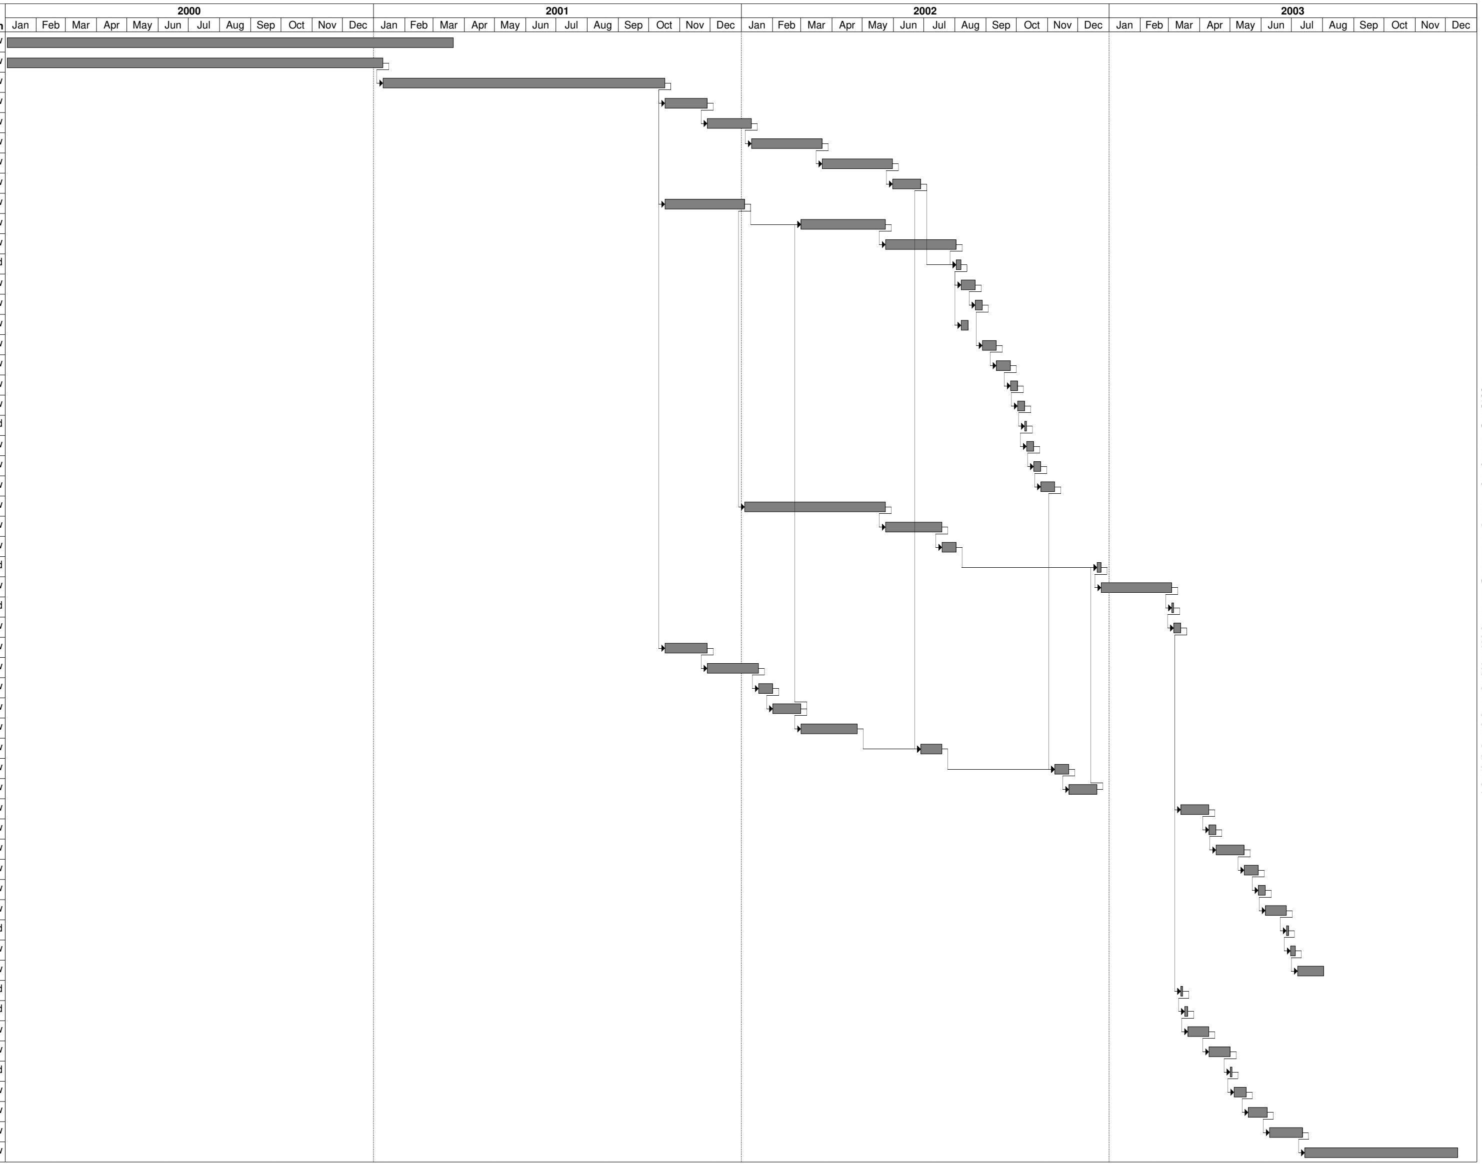

Table 12.4: The ALICE installation schedule.

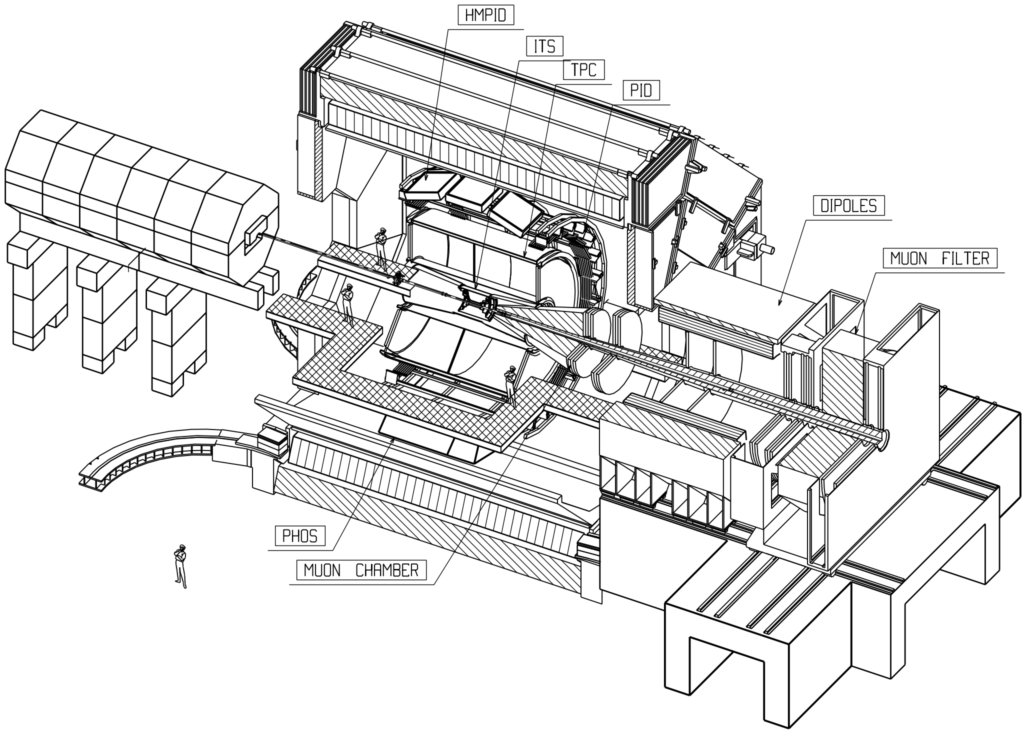

Figure 12.11: Gangway system inside the L3 magnet.

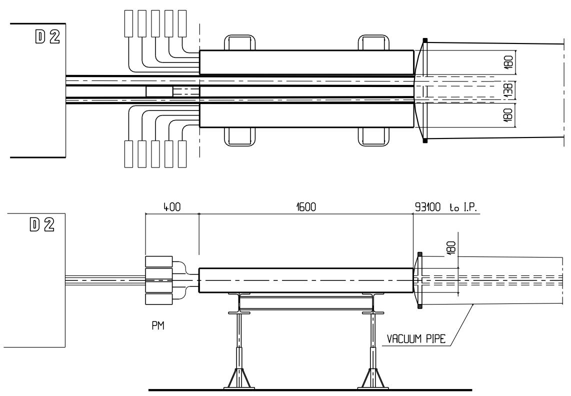

Figure 12.10: Layout of the ZDC installed in front of the D2 machine magnet.

12.3.8 Radiation shielding

The upgrade of the radiation shielding at Point 2 can be carried out in two ways: by increasing the shielding at the weak spots, and by placing a dedicated shielding around the beam line.

Placing the shielding around the beam line would interfere with the movement of the magnet doors, but is still preferable to reinforcing the present concrete plug in the 23 m shaft.

The thickness of the beam line shielding has been estimated (see Section 12.5.4) to be 1 m of concrete. The total weight of the shielding is about 2000 t. It would be constructed using standard CERN size blocks and a limited number of specially made concrete modules (about 300 t), which could be removed and stored inside the experimental area within one day. The shielding over the low-β magnets would be semi-permanent.

12.4 Access, maintenance and operation

12.4.1 Access to sub-detectors

In general, access to the different parts and subdetectors of ALICE is favourable. Since the central detectors of ALICE do not extend out to the coil, sufficient space (about 80 cm) is left for the installation of a gangway system (non-magnetic) inside the L3 magnet volume as shown in Fig. 12.11. A chicane arrangement in the beam shielding would allow immediate access to the gangway system via the large opening in the magnet door (on the access-shaft side). This would even permit rapid access to the first five planes of tracking chambers of the muon arm and convenient access to the end planes of the TPC and PID detectors. Because of the front absorber, the ITS would only be accessible from one side. The PHOS is completely obstructed by the PID detector, but sufficient

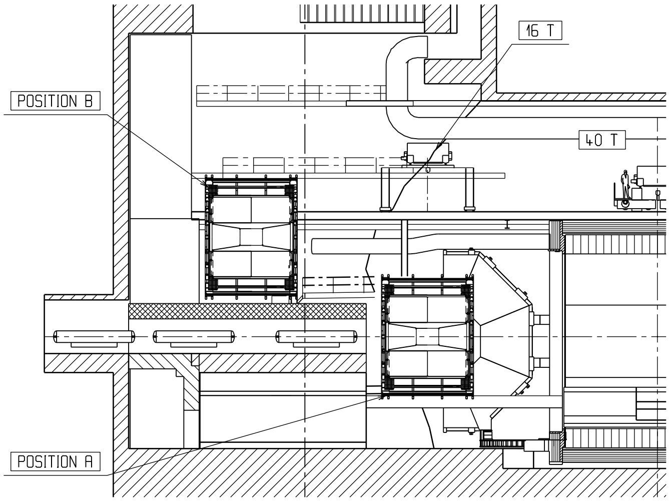

Figure 12.12: Removing the space frame with the central detectors over the low-β magnets.

space is available in between the PHOS support beams to reach the PHOS detector.

Dismantling the ITS detector

Removing the ITS detector for maintenance or repair will necessitate breaking the vacuum chamber. Since the front absorber prohibits access to one side of the ITS, a section (about 2 m) of the vacuum chamber on the 'free side' will be removed. A support tube will then be inserted through the ITS and attached to the absorber, to support the thin central vacuum chamber. The ITS can then be disconnected and removed from the TPC. This procedure assumes that the Pixel and FMD detectors are already removed, as described in the ITS chapter.

Dismantling the TPC/PID system

The shielding and beam line must be removed over a 7 m long section. Temporary rails are used to move out the common support structure. The complete structure can then be brought to the surface without removing the low-β magnets (see Fig. 12.12). The doors of the SX2 building would not allow the complete structure to be rolled out of the building, but a standard crane would be able to remove it via the 12×12 m^2 roof access panel.

Removing the front absorber section

Since the absorber is attached to the dipole magnet, it is removed by displacing the magnet by about 6 m. In this position the overhead crane can remove the absorber and the magnet can be reinstalled in its operating position. In order to make space for the withdrawal of the magnet, the muon filter is divided into two parts, which can be separated and moved perpendicularly to the beam axis (see Fig. 12.13 a–d). The section of the small angle absorber inside the magnet does not have to be removed since it is attached to the magnet yoke. However, the remaining part of the absorber, between the dipole and the low-β magnets, must be removed before the dipole magnet can be rolled back.

The sequence assumes that the Pixel detector has been removed and that the vacuum chamber is disconnected in front of the absorber on the 'opposite' side of the ITS. When the absorber is rolled back, a 1.5 m long section of the beam pipe will pass through the ITS, supported by a temporary support tube (for a more detailed description of the vacuum chamber see Section 12.5.2).

12.4.2 Electronic racks

It is proposed to keep the racks for the readout, control and monitoring in the existing counting rooms which are located in the PX24 shaft, and to install a new counting room at the base of the PGC2 shaft for the muon arm extension. There are a total of 4 barracks housed in the vertical shaft of the pit PX24, containing the following number of racks:

- Level 1: 32 racks (closest accessible rack position)
- Level 2: 22 racks
- Level 3: control room, consoles only
- Level 4: 30 racks

We plan to add a Level 5 counting room on top of the Level 4 one. This was already foreseen in the original design of the L3 counting rooms and will therefore not require any structural changes. In addition to these permanently accessible counting rooms, the present L3 blockhouse is equipped with 64 racks and it is also possible to place about 20 racks close to the muon arm. The L0 trigger electronics for ALICE might similarly be placed close to the detector in the experimental hall. However, these racks will not be accessible during LHC operation.

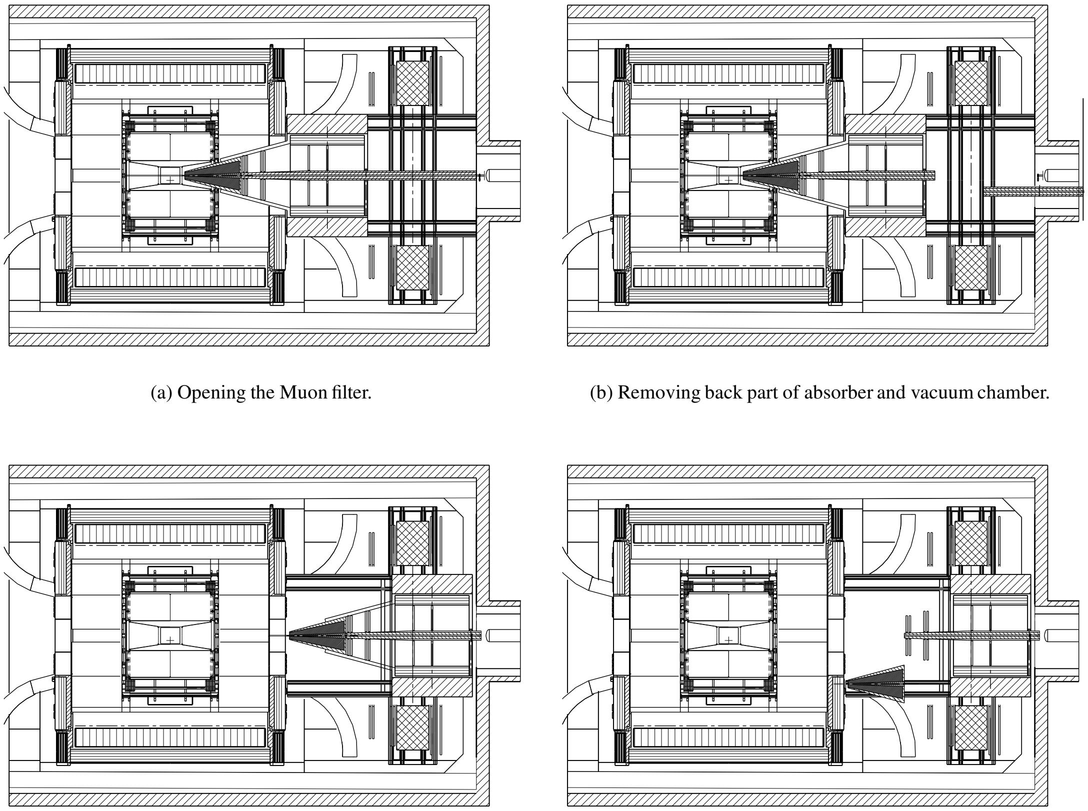

(a) Opening the Muon filter.

(b) Removing back part of absorber and vacuum chamber.

(c) Breaking the vacuum chamber and translation of the dipole magnet and absorber by 6 m

(d) Removal of front absorber.

Figure 12.13: Removal of front absorber section.

The difference in cable length between Level 1 and Level 4 is 40 m. A typical cable length from Level 1 to the centre of the L3 magnet is 100 m. About 10 additional racks could be installed at the base of pit PGC2 to house the muon-arm extension electronics, as well as control systems for the dipole magnet and gas circuits. The cable length from the muon arm to the PGC2 control room could be as little as 75 m, but a link between Level 1 and PGC2 could not be less than 250 m.

Cooling of the electronics is achieved by local water-cooled heat exchangers. Air will be forced vertically upwards through heat exchangers and crates, finally being returned to the base of the rack via ducts incorporated in the side walls. Consequently, electronics crates must allow for vertical air flow. Cooling water will be supplied at 13°C from the 'mixed' water supply. It assumes an average heat dissipation of 9 kW per rack. Each rack will be equipped with a safety monitoring system, including temperature sensors and smoke detectors.

Cable routing will essentially be as it is at present, but some modifications to the cable chicanes might be necessary in order to allow for additional radiation protection. All cables installed for ALICE will conform to the CERN safety instruction, IS 23, in respect of being halogen free and low smoke-producing.

12.4.3 Cooling

Heat removal from the densely packed front-end electronics and control racks is essential for proper operation of the ALICE detector. It is estimated that some 940 kW will be dissipated by the experiment, excluding the L3 solenoid magnet.

The complete water-cooling loop will comprise five circuits: untreated (raw) water, primary water, chilled water, mixed water and demineralized water. The different cooling circuits and the expected heat dissipation in the ALICE detector systems are summarized in Fig. 12.14.

The primary water will be cooled using the existing cooling towers, where the majority of the heat will be removed by evaporation. This primary water circuit will be used to cool the refrigeration plant which produces chilled water and to supply cooling water for the cryogenic com-

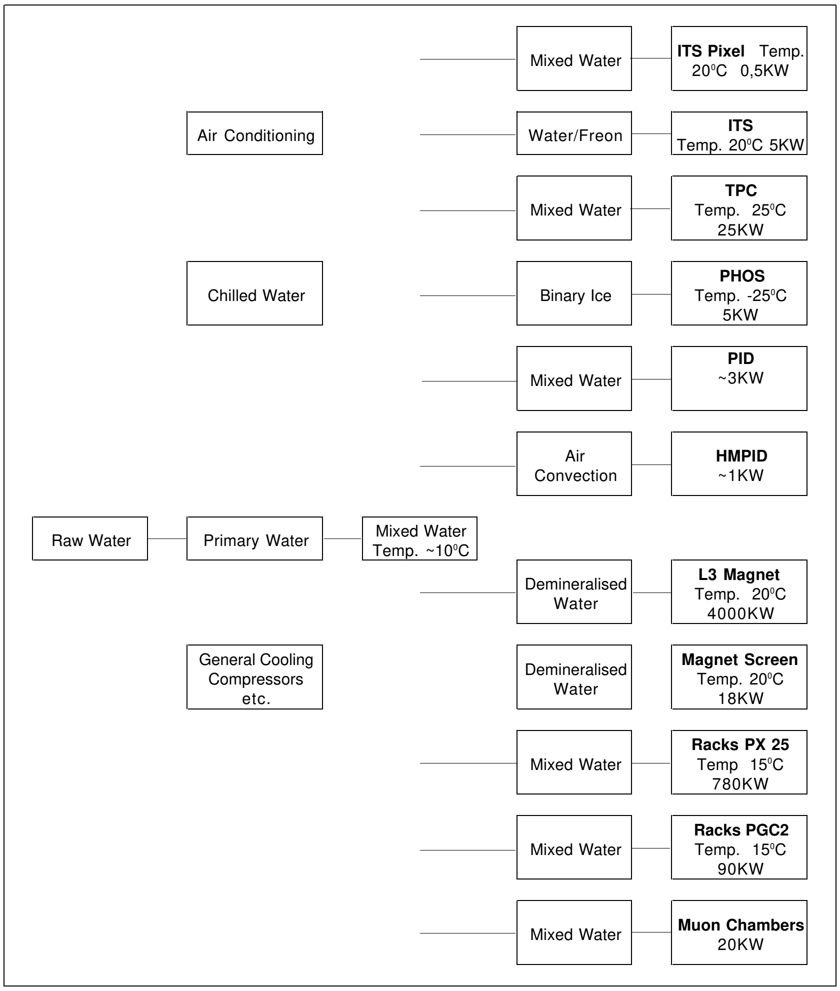

Figure 12.14: Summary of the expected heat dissipation in the ALICE detector.

pressors. Air conditioners will use chilled water directly. This circuit will also be used for cooling the mixed-water circuit.

ITS – The ITS will work at an operating temperature of 20°C. The heat produced by the pixel layer (500 W), the Si strip detector (3 kW) and the front-end electronics (2 kW) will be removed by water. For the Si drift detector we require a temperature stability of ± 0.1°C. Here it is envisaged to use evaporative cooling using Perfluorocarbon (C_5F_{12}) in special cooling arteries along the detector.

TPC – The front-end electronics of the TPC, situated at the end flanges, will produce a 25 kW heat dissipation. It will be cooled by water taken directly from the mixed-water circuit and adjusted to the operating temperature of 25°C by simple bypass temperature control.

PHOS – The PHOS must be maintained at -25°C with a temperature stability of ± 0.1°C. The heat dissipated by the PHOS is expected to be approximately 5 kW. To prevent condensation, the sub-detector will be installed in an insulated enclosure. The PHOS will use either binary ice or some other suitable refrigeration system to reach the required operating temperature.

Racks – Electronic racks will be cooled directly from the mixed-water circuit, as is the present system. It is estimated that, on average, 9 kW will be dissipated in each rack. A ventilation system is incorporated in each rack to remove any excess heat not removed by the water-cooled exchangers and also to improve temperature uniformity. Each rack will be equipped with a safety monitoring system, including temperature sensors and a smoke detection system.

The L3 magnet – The 4 MW dissipated by the L3 magnet at full field is removed by the demineralized-water circuit, which has an inlet temperature varying between 20 and 25°C. In addition, there is a thermal shield around the internal surface of the magnet coil to protect the internal sub-detectors from the heat produced by the coil. It is capable of removing 18 kW at an operating temperature of 20°C and is also connected to the demineralized-water circuit.

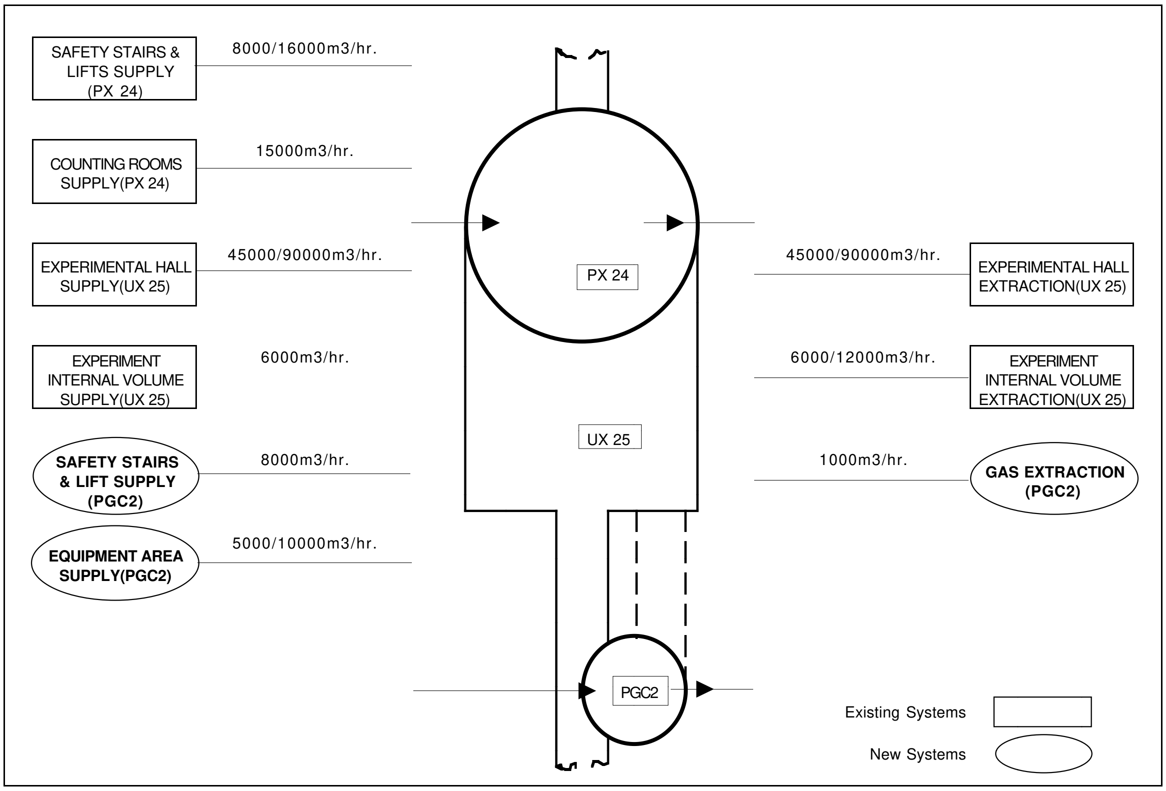

Figure 12.15: Overall capacity of the ALICE ventilation system.

Dipole magnet – It is proposed to connect the muon arm dipole magnet to the existing helium supply and recovery system at Point 2 to fulfil the helium gas requirements. The compressor will be located in SUH25, primary water will be used to remove the 500 kW generated. Stainless steel pipes connect the helium supply and return to the cold box, which will be situated at the base of pit PGC2, as close as possible to the experiment.

12.4.4 Ventilation

Point 2 of LEP is already equipped with a ventilation system both underground and aboveground. All systems operate with 100% renewal of fresh air. Additional filtration of the extracted air or limited recycling might also be required by the increased radiation levels.

Hall UX 25

In normal operation fresh air is supplied to the underground experimental area at the rate of 45 000 m^3/h. This air is treated (17°C and with a 10°C dew point) and is injected into the experimental hall through four low-level diffusers, which help to eliminate draughts.

Used air is extracted from the experimental hall at high level and at the same rate. However, before it is ejected into the atmosphere it is passed through filters to remove any possible contamination of the dust (US Air Filter Institute NBS-1 standards).

In the event of fire or gas alert, the throughput of the ventilation systems is doubled to 90 000 m^3/h, this additional flow not being passed through any filter system. Figure 12.15 gives a summary of existing and new ventilation systems foreseen for ALICE.

Inside the L3 magnet volume

In addition to the hall ventilation there exists a dedicated supply and extraction system for the requirements of the volume inside the L3 magnet. This system will be re-used in the ALICE configuration to remove residual heat not removed by the water-cooling circuits, and also to limit the build-up of flammable gases in the event of small leaks. The normal air flow is 6000 m^3/h. The inlet temperature is 18°C with a dew point of 10°C. A branch of this system extracts air from the gas distribution racks which are situated on the plug in PX24. Under normal running conditions the extracted air is treated before being ejected into the atmosphere. This is not the case when the air flow is doubled.

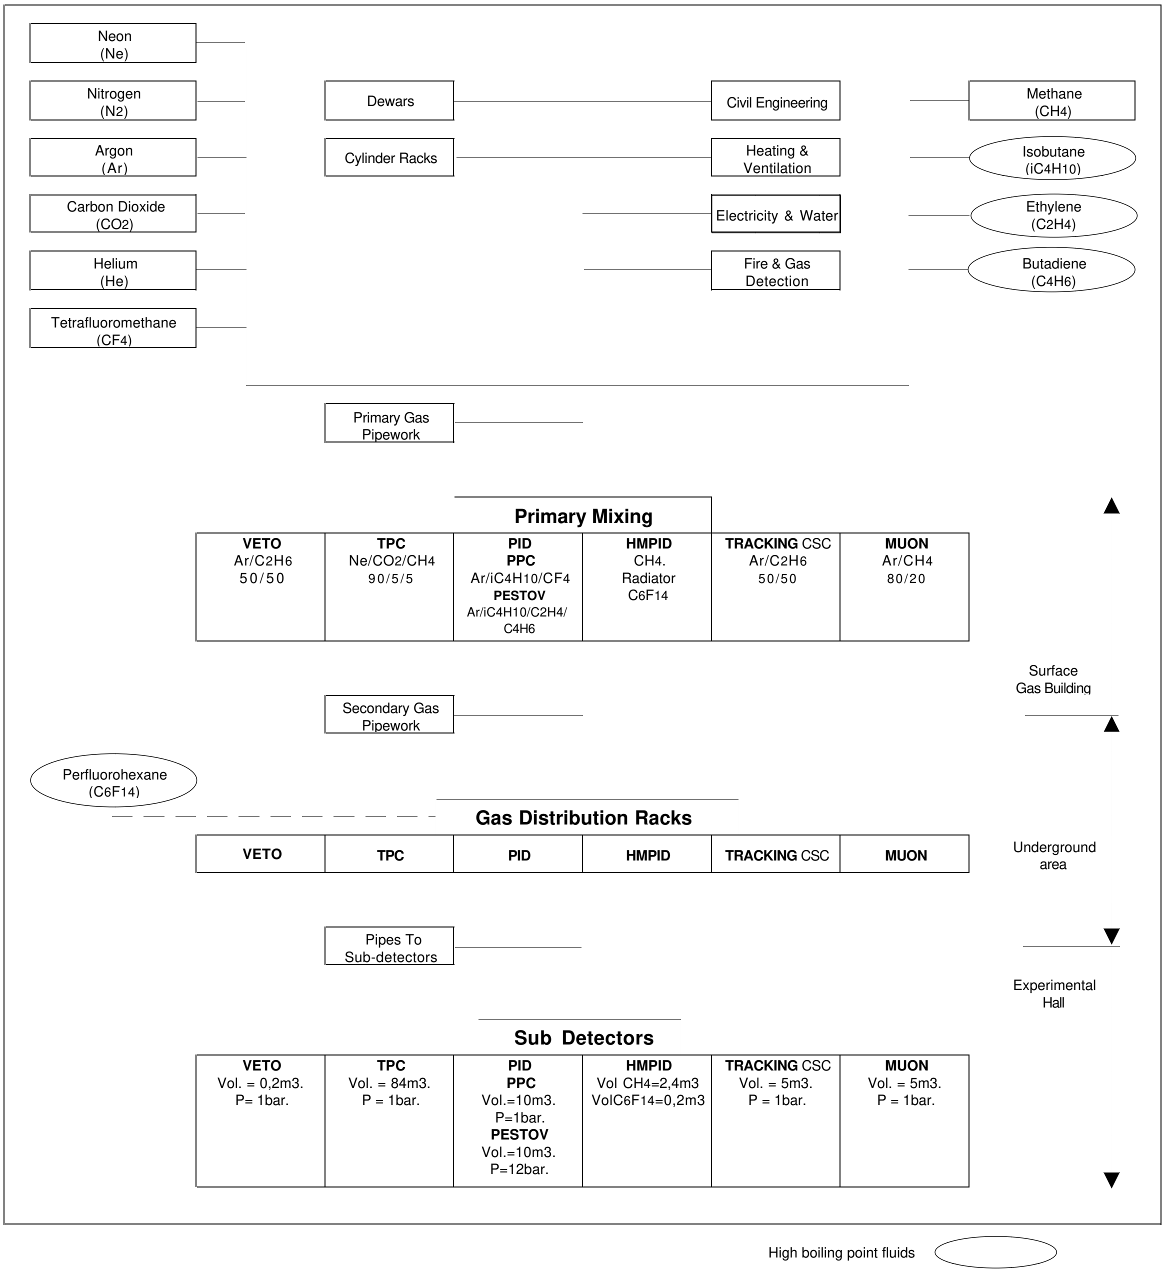

Figure 12.16: Summary of ALICE gas requirements.

Counting rooms PX24

To maintain a normal temperature in the counting rooms, both for non rack-mounted equipment (e.g. computers) and personnel, fresh treated air is supplied at 15 000 m^3/h.

Safety

For safety purposes each stair and lift shaft is pressurized with fresh air at the rate of 8000 m^3/h. A smoke extraction system capable of removing 10 000 m^3/h is installed at a high level in UX25. Smoke extractors are also located in the roof of SX25 to remove smoke in the event of fire.

Additional requirements

New heating and ventilation systems will be required for the area at the base of the new access shaft PGC2, where the electronic racks, gas and cryogenic equipment will be located for the muon-arm extension. The stairwell and lift shaft will require a supply of 8000 m^3/h. It is estimated that the remaining PGC2 shaft and working areas require

5000–10 000 m^3/h. In addition a small air extraction unit (1000 m^3/h) will be needed to remove flammable exhaust gases originating from the gas control rack.

12.4.5 Electrical power supply

The power supply currently installed at Point 2 available to the L3 experiment is the following: two transformers (4 MVA each) working in parallel supplying the L3 magnet; one transformer (2 MVA) supplying the counting rooms and electronics.

In addition, the General Service supply circuit provides power in the experimental area, assembly halls and for conventional infrastructure. The power distribution system in the counting rooms, the experimental area and the assembly halls will not need to be replaced or modified. A new distribution system, but without additional power capacity, will have to be installed in the PGC2 shaft and for the muon-arm extension.

12.4.6 Gas systems

The gas systems for the ALICE experiment will use much of the existing installation located at LEP Point 2. In total five sub-detectors will require gas mixtures. The gas requirements are summarized in Fig. 12.16. Proper attention will be paid to the handling of the gases and special recovery systems will be installed to ensure the correct disposal of used products.

Gas building

The existing gas building situated at LEP Point 2 will be re-used for the gas systems of the ALICE experiment. It is not envisaged to carry out any major modification to the building, although some minor changes cannot be avoided.

Storage and primary distribution

Gas storage and primary distribution will be located in the existing gas building on the surface and will be sufficient for the requirements of the ALICE experiment. All gases will be delivered either in the liquefied state, as is the case of argon and nitrogen, or in pressurized cylinders for carbon dioxide, ethane, etc.

The gases will flow in fixed pipes, which will be equipped with purging systems and pressure relief safety valves. All flammable gas systems will have an automatic shut-off valve installed in the pipeline, which will close in the event of a the gas alarm being triggered.

The Gas Systems

These will be of the closed circuit re-circulating type, similar to the ones in operation now. The primary gases will be fed from the storage area at medium pressure into the mixing units located in the temperature-controlled 'mixing room' of the surface gas building.

The mixing of the specified gases for each sub-detector will be done automatically by computer control of mass flow meters. The mixtures obtained will be monitored and checked against certified pre-mixed bottles. The electronic equipment and computers needed for the mixing and regulation of the gas systems will be housed in the instrument room of the gas building.

Online monitoring of the gas systems will check for impurities such as oxygen or moisture. Purifiers and driers will be installed in the circuits to eliminate these impurities if required. The monitors will also be connected to an alarm system designed to protect the chambers by cutting off the gas supply and high tension if required.

The gas building will supply the gases to the distribution racks located on the shielding plug in PX24 and in PGC2 (for the muon arm). All piping will be made of stainless steel.

These distribution racks will contain the recirculating compressors, pressure control loops and safety shut-off valves, which will isolate the detector gas volume in the event of power failure, or of the gas alarm being triggered. All systems will be equipped with a 'bubblers' system. A system for purging the detectors will be installed in both the gas-mixing and distribution racks. A dedicated nitrogen gas circuit will be installed for this purpose.

From the distribution racks a limited number of main distribution pipes will bring the gases to the experiment, where a manifold will distribute to each individual detector unit.

A flow regulator will be installed in each supply line to monitor the gas flow through each set of chambers. The return gas flow from the chambers to the gas distribution rack will follow the same path as the supply lines.

12.5 Detector – machine interface

The interface between the LHC and the ALICE detector does not present any particular difficulties, since maintenance and complete dismantling of the ALICE sub-detectors can be carried out without removing any machine elements (see Section 12.4.1). The machine elements, which penetrate into the experimental area, will be supported on a concrete platform (part of the present L3 blockhouse structure) to give the necessary stability.

12.5.1 Vacuum

The vacuum requirements in the ALICE insertion are given by the expected beam–gas event rate generated by the 530 mA proton beam. An average vacuum pressure of 2.5×10^{-11} torr (see Fig. 12.17) in the ± 20 m intersection region would generate a background event-rate of about 60 Hz/m, which is fully compatible with the overall ALICE event-rate considerations.

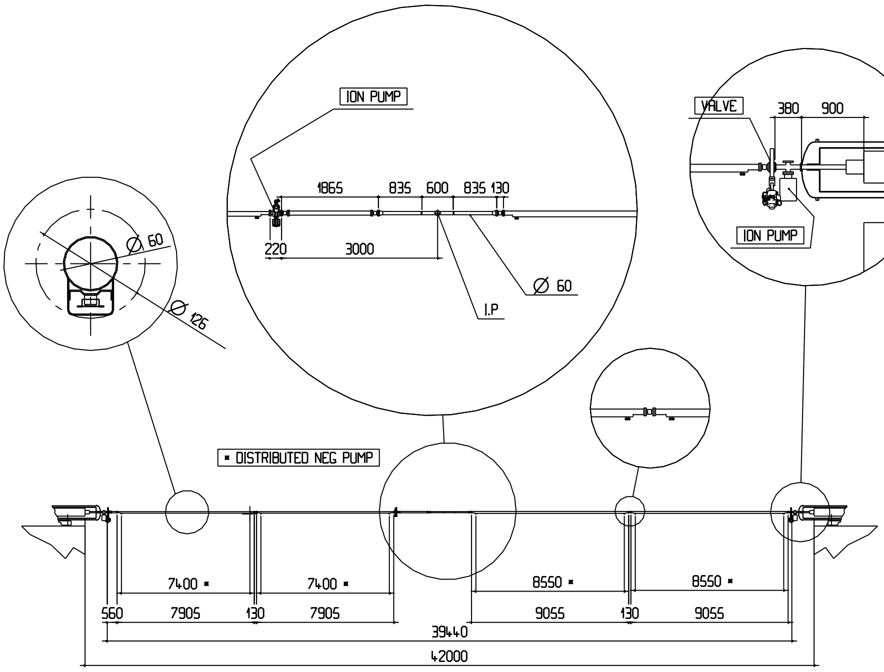

Figure 12.18: Overall layout of the ALICE vacuum chamber.

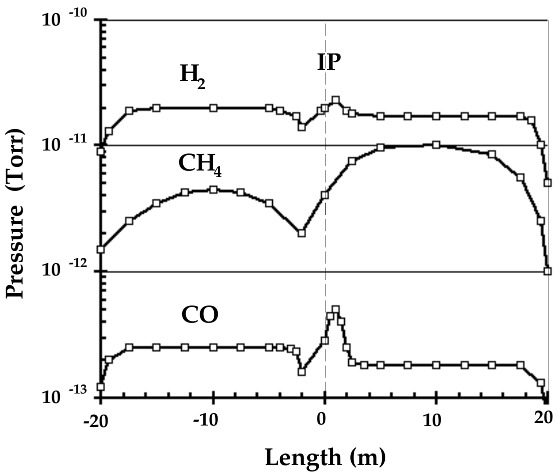

Figure 12.17: Vacuum pressures in the ALICE intersection region.

In order to achieve the required vacuum pressure, the ALICE vacuum chamber will be equipped with a NEG pumping system, covering the entire intersection region apart from the central section [3]. In addition, ion pumps will be placed just in front of the last insertion quadrupoles and close to the central section. The central ion-pump will be placed as far as possible away from the intersection point, in order not to obstruct the access to the ITS detector. Alternatively it can also be placed on the absorber side, integrated into the front absorber block.

It will be necessary to make a complete bake-out of the ALICE insertion chamber. The forward sections (from about 1 m to 19 m) of the vacuum chamber can be equipped with permanent heating jackets. The central section inside the ITS detector will require a special removable heating jacket, as described in the ITS chapter.

12.5.2 Beam pipe

The beam pipe is divided into a central thin section and two forward sections (see Fig. 12.18) where the mechanical design of the forward vacuum chambers is of less importance to the ALICE detector. The central section of the beam pipe is a beryllium cylinder, 30 mm in radius (inner radius) and 0.6 mm thick, extending up to $z = \pm 0.30$ m. The beryllium tube is welded to an aluminium alloy section with the same diameter and a wall thickness of about 1.2 mm. The aluminium section extends on one side up to the front of the absorber ($z = 0.90$ m). The non-absorber side is more complicated, since provisions have to be made to facilitate the removal of the Pixel detectors, as well as the ITS detector. The aluminium section is extended to $z = 1.13$ m, where it is attached to the forward-section vacuum chamber by a standard vacuum flange. An ion pump is situated at $z = 3.0$ m, and a flange placed at $z = 5.0$ m permits the removal of a 3 m long section of the

Table 12.6: Radiation doses (in Gy). The last column gives the total dose received over ten years and the respective neutron fluences (in neutrons per cm^2).

Detector	pp (Gy)	Ca–Ca (Gy)	Pb–Pb (Gy)	Total (Gy)	Total n/cm^2
Pixel 1	250	700	230	1180	$< 10^{12}$
Pixel 2	76	220	71	367	$< 5 \times 10^{11}$
TPC	0.1	0.2	0.08	0.4	$< 10^{10}$
TC1	0.1	0.2	0.04	0.3	$< 5 \times 10^{11}$
TC3	0.04	0.06	0.02	0.1	$< 10^{11}$
PID	0.01	0.04	0.02	0.07	$< 10^{10}$

Table 12.5: Running time, average luminosity $\langle \mathcal{L} \rangle$, number of minimum-bias collisions and total charge-particle multiplicity per minimum-bias interaction N_{ch}.

Type	Time (s)	$\langle \mathcal{L} \rangle$ (cm^{-2}s^{-1})	Events	N_{ch}
pp	10^8	10^{30}	10^{13}	100
Ca–Ca	2.5×10^6	3×10^{27}	2×10^{10}	2 400
Ca–Ca	2.5×10^6	10^{29}	8×10^{11}	2 400
Pb–Pb	5.0×10^6	10^{27}	4×10^{10}	14 200

vacuum chamber and provides the possibility of removing the ITS detector.

The vacuum chamber inside the absorber should fit the inner aperture of the absorber (∅ 126 mm), but can in all other respects be optimized according to LHC machine requirements (wall thickness and effective aperture). For reasons of symmetry identical designs are used in the two forward sections.

12.5.3 Radiation environment

We will take data with proton and light- or heavy-ion beams for different time periods and different luminosities. The radiation load on the various parts of the detectors must therefore be calculated for a combination of different beam conditions. In Table 12.5 we list a typical running scenario for ten years of operation (see also Section 8.8). It includes ten years of pp operation and five years each of heavy- and light-ion running, the latter being split equally into low and high luminosity parts. The table shows running time, average luminosity $\langle \mathcal{L} \rangle$ (in cm^{-2} s^{-1}), number of minimum-bias collisions during this time and the total charged-particle multiplicity per minimum-bias interaction N_{ch} in pp, low and high luminosity Ca–Ca, and Pb–Pb runs.

In order to estimate the overall radiation load on the ALICE detector, six typical volumes were identified, mainly in places where detector electronics will be situated:

- silicon Pixel layers ($r = 4$ cm and $r = 7.5$ cm);
- TPC electronics (muon arm side, disc at $z = 270$ cm, $r = 80$–120 cm);
- muon chambers TC1 and TC3 (discs below 12.5°, $z = 465$ cm and $z = 675$ cm, respectively);
- PID electronics (cylinder at $r = 380$ cm).

The deposited energy and the total neutron fluency, including thermal neutrons, were simulated with the FLUKA-95 Monte Carlo code [4] for each of the selected volumes and the different beams [5]. The geometry includes the L3 magnet and the muon absorber. The DTUJET model [6] was used for the simulation of the proton–proton interactions, and the nucleus–nucleus collisions were generated with HIJING.

The results are shown in Table 12.6, where all values have been normalized to the ten-year running scenario mentioned above. The pixel detectors close to the beam receive the highest dose, up to 100 krad (~ 1000 Gy) and $< 10^{12}$ neutrons/cm^2, over ten years. As the radiation load scales approximately with $1/r^2$, doses for other detectors can also be estimated from this table.

12.5.4 Radiation shielding

The radiation shielding of the Point 2 experimental area was adopted to the LEP radiation situation, i.e. concrete walls of about 1.0 m thickness separate the radiation controlled zone from the permanently accessible areas. The LHC will require a more stringent arrangement and the shielding thickness should be increased to protect from a full beam loss.

The FLUKA code [4] was used to simulate the hadronic and electromagnetic cascades generated by the loss of a 530 mA proton beam inside the Point 2 experimental area, where an additional shielding of 1 m of concrete is added around the beam pipe on the access shaft side of the L3 magnet. Several different beam-loss scenarios were investigated and the least-protected loss point was found to be situated between the new dipole magnet and the muon filter. The dose map for this situation is shown in Fig. 12.19. The dose area in the counting room has a maximum value of 40 mSv, which is below the CERN limit of 50 mSv [7] imposed for a non-repetitive loss situation.

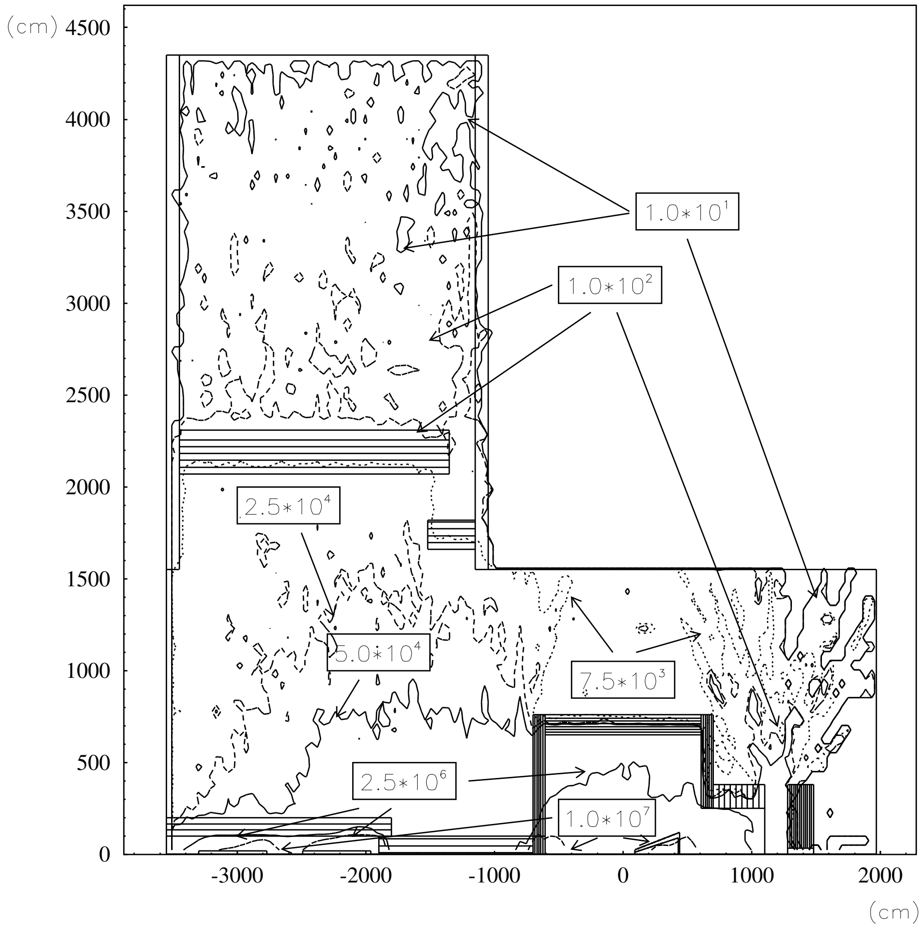

Figure 12.19: Dose map, in mSv, following a loss of a 530 mA beam.

The dose in the counting room area could be further reduced by adding more shielding around the beam line. Increasing the shielding from 1 to 2 m is possible without interfering with the ALICE detector. The additional shielding would however be rather costly. The final shielding situation in the Point 2 experimental area will also depend on the different radiation protection measures associated with the LHC machine.

12.6 Safety

The ALICE Collaboration is determined to provide, in a combined effort with the CERN–TIS Division, the highest possible level of safety during the construction, installation and operation of the ALICE detector system. The principal aim is to ensure the safety of all personnel involved in the project, but also to ensure, as far as possible, minimal or no material damage to detector equipment. The overall safety policy for ALICE will strive to identify potential safety hazards at an early stage by performing safety analyses during the design and construction phases.

12.6.1 General safety considerations

To reduce safety hazards, the ALICE detector avoids large quantities of flammable gases, cryogenic liquids and material which gives off toxic fumes in the event of fire. The use of plastic and other non-metallic materials will conform to the TIS instructions IS 23 and IS 41. The passive safety arrangements in the experimental area (radiation shielding, access and escape exits, cranes and lift-systems) will conform strictly to CERN requirements or industrial standards. No derogation will be requested by the ALICE Collaboration. The active safety system will include hard-wired alarms connected to the fire brigade, forced ventilation of the L3 magnet volume and specialized fire extinguishing systems. Potentially the most dangerous period will be the time of installation. In order to guarantee that direct action can be taken with regard to everyday safety hazards, the ALICE Collaboration will request that a safety inspector be present during the assembly and installation operations.

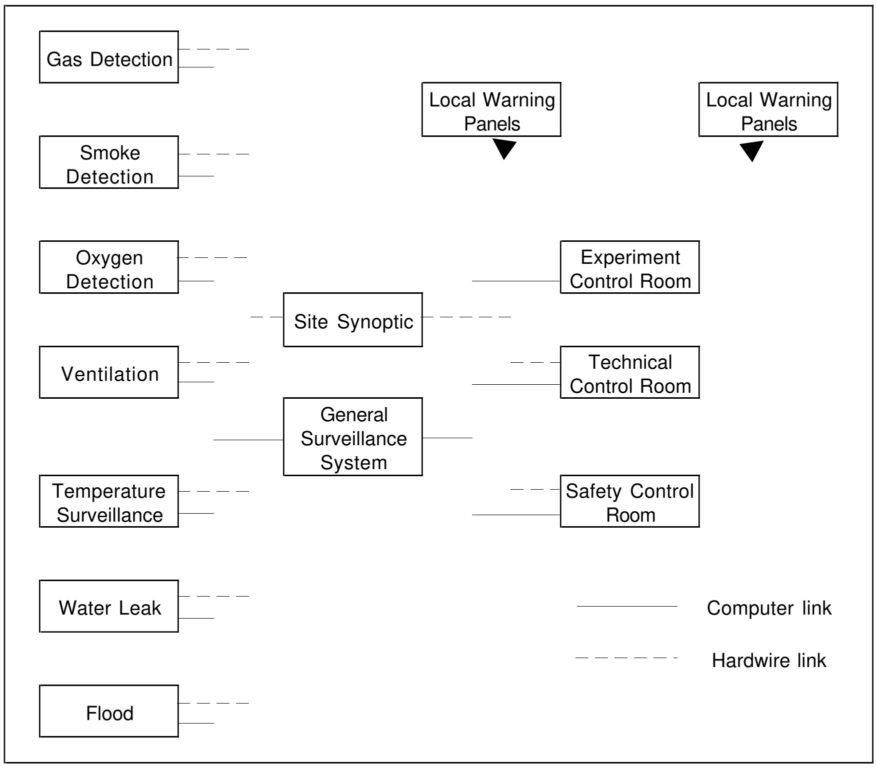

Figure 12.20: Principal flow chart of the ALICE safety network.

12.6.2 Safety Systems

The General Surveillance System (GSS) in operation at Point 2 will be adapted to the new situation with the implementation of the ALICE experiment. It will be used to monitor the equipment in the ALICE experiment for any anomalies which could affect the safety of personnel or equipment. The main task of the GSS is to transmit all level 3 alerts directly to the fire brigade. It will also distribute and display information relevant to the overall safety status and, in the case of a hazardous situation, automatically place all circuits and equipment in a safe state.

The information to be handled by the GSS will come from detection equipment (gas, smoke detectors, etc.) placed in the experimental area close to the various detector units. The specialized detector equipment will also provide local alarms, which can be transmitted both by computer link and by hardwire to the various control rooms, as indicated in Fig. 12.20.

Processing and transmission of information

The GSS processes the information according to the following levels of alert:

- Alert level 0 – Information (in service, etc.).
- Alert level 1 – Technical fault (as soon as possible).
- Alert level 2 – Technical fault (immediate intervention).
- Alert level 3 – Full alarm (immediate intervention by fire brigade).

This information will be displayed in the ALICE control room. The level 3 alarm will be transmitted to the fire brigade, the LHC main control room, the CERN technical control room and the Point 2 local guard house SY2.

Gas detection

All zones where flammable gases are to be used, will be equipped with a flammable-gas detection system. Two systems are currently in use in the existing LEP Point 2 area:

- The standard industrial pelister type, which is used to monitor the surface gas building, the gas racks located on the plug in PX24 and also in the underground experimental area UX25.
- The SDN air sampling system, which extracts air from critical volumes near the gas-filled detectors.

Smoke Detection

It is proposed to continue using the existing SDN systems for fire detection in the underground areas. If smoke is detected inside the L3 magnet volume (response time about 2 minutes), the fire will be extinguished by either injecting pressurized nitrogen from a surface storage tank into the detector volume, or by a water fog system.

The counting rooms and electronic racks are at present protected by a halon fire extinguishing system. This system is activated by a SDN system installed in individual racks and counting rooms. It is proposed to further reduce the fire risk in electronic racks by eliminating, as far as possible, the amount of flammable material in the racks and providing very early smoke/temperature detection in each rack, which would directly cut off the electrical power in the event of an alert. Standard smoke detection systems, like those already installed, will continue to be used in the surface buildings.

Oxygen deficiency

Oxygen deficiency monitors will be installed in the confined volumes of the ALICE experiment (inside the magnet) and in the underground passageways.

Ventilation

The ventilation system will be in continuous operation. In the event of either a gas or smoke alert it will be switched into the high-speed mode. In the event of a general power failure at Point 2, the ventilation systems will be powered by the standby diesel generator.

Temperature surveillance

Temperature sensors will monitor the air volumes close to the detector in order to maintain a stable ambient temperature.

Water leak

The water-cooling circuits will be under constant surveillance. Alarms will automatically be transmitted to the CERN maintenance service in the event of malfunction.

Flood

Large water infiltration or leaks are collected in a sump below the experimental area floor level, from where it can be pumped to the surface.

Access

Access to the ALICE experimental area will be prohibited during machine operation. The present access system must be modified and converted to LHC standards.

Safety problems related to the Pestov Counter

As discussed in section 4.2, the Pestov counter will work with a gas mixture of Ar (76.9 %), CH_4 (20 %), C_2H_4 (2.5 %) and C_4H_6, at an operating pressure of 12 bar. It consists of 36 modules with a volume of 0.2 m^3 each. All the modules will be tested before installation to withstand an overpressure of 18 bar [8]. Additionally, they will be equipped with separation valves, mass flow meters, and pressure gauges on the inlet and outlet. If during operation any anomaly in the gas flow is recorded or the SDN system detects the presence of flammable gas in the surrounding air, the appropriate module will be isolated. The gas content will thereafter be vented into the exhaust and the module will be flushed with N_2. In this way, leakage of flammable gas into the volume inside the L3 magnet is minimized.

13 Planning and organization

13.1 Schedule

The present schedule for the LHC foresees the shutdown of LEP operation at the end of 1999, and the first physics runs starting in 2004. Leaving one year for the final decommissioning of the L3 experiment, the dismantling of the L3 detectors could start at the beginning of 2001. The experimental area could then be cleared for ALICE by October 2001.

The complete installation schedule covering the period between LEP shutdown and the commissioning of ALICE is described in Section 12.3.

This scenario sets the time scale for the ALICE experiment. A global schedule for design, prototyping, construction, and installation of all the ALICE sub-systems is presented in Fig. 13.1.

All sub-detectors will require at least two more years of intensive R&D, detailed design, and initial tendering. This also applies to the integration studies of the whole detector. Most detectors will have to start being constructed in 1997 or 1998.

It is foreseen to pre-assemble all the barrel detectors in the SLX2 hall at Point 2. This will allow early preparation of the different detector services at the surface and the installation and assembly procedure to be analyzed and corrected before integrating the detectors into the L3 magnet. The pre-assembly stage will last about one year, and should therefore start at the beginning of 2002.

With this schedule, the final installation of the equipment into the UX25 experimental area would start at the beginning of 2003, and proceed rapidly, leaving the second half of 2003 for commissioning the complete experiment.

13.2 Cost estimates

The overall cost estimate for the ALICE experiment is summarized in Table 13.1, taking into account the recommendations made by the LHC Cost Review Committee (CORE) [1]. The cost has been based, whenever possible, on quotes from industry. However, in a number of areas, in particular electronics and computing, large extrapolations are necessary.

A cost range for some of the items is listed in Table 13.1, reflecting options where a decision cannot yet be taken. For the TPC this corresponds to the choice of front-end electronics (flash ADCs or SCA analog memory), which is estimated to be more expensive when using the analog memory.

As described in Chapter 4, the preferred solution for the TOF particle identification array consists of Pestov spark counters. The present uncertainties connected with building and operating a large TOF system of spark counters make it necessary to keep the alternative design based on the cheaper PPC detectors, as a back-up for about one more year.

Table 13.1: Cost estimates for the ALICE experiment.

Detector System	Cost (kCHF)
Inner tracking system	20 943
Time projection chamber	13 062–14 160
Barrel PID	8696–13 496
High momentum PID	2728
Photon spectrometer	13 405
Charged particle veto	564
Forward multiplicity detector	766
Zero-degree calorimeters	1779
Muon arm	30 000
Data acquisition and trigger	11 004
Software and offline links	5000
Infrastructure and installation	8034–9034
Total detector cost	115 981–122 879

We have estimated the costs of the infrastructure based on LEP experience, assuming that much of the existing equipment in Point 2 can be re-used. Some costs have been included for modifications to the L3 magnet. A major uncertainty persists about the cooling system of this magnet, which might have to be renovated at an estimated cost of 1 MCHF. A final decision about this intervention might be possible only after the shutdown of LEP (see Section 12.3).

The muon arm has not yet been studied with the same level of detail as the other sub-systems. Its cost, estimated currently as 30 MCHF, has to be regarded as preliminary. A detailed cost estimate will be presented in the muon addendum to the ALICE technical proposal in the second half of 1996.

The total cost of the ALICE detector, including the Muon Arm, as presented in this Technical Proposal is between 116 and 123 MCHF. A detailed breakdown of the cost estimate is available and will be presented to CORE.

13.3 Responsibilities

A preliminary sharing of responsibilities for the different sub-systems of ALICE is indicated in Table 13.2. The responsibilities will include all technical and financial aspects of the project, from R&D and design, to construc-

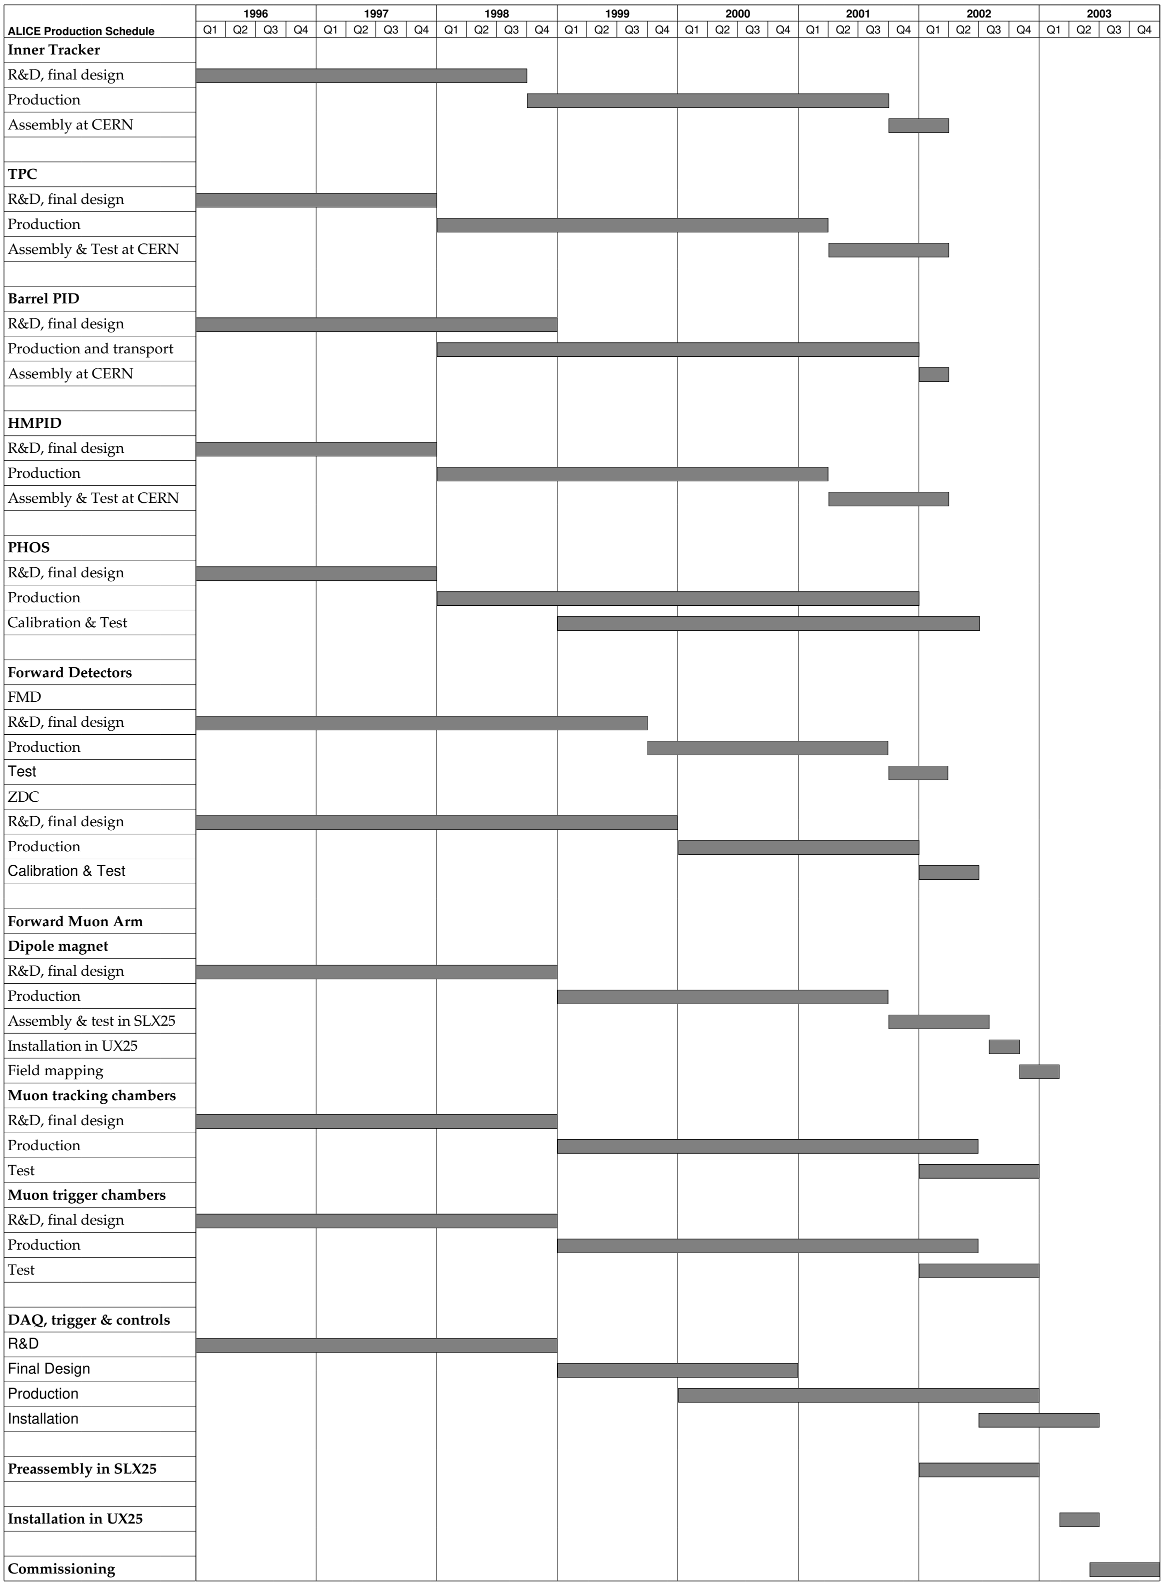

Figure 13.1: Overall schedule for construction and installation of the ALICE experiment.

Table 13.2: Responsibilities.

ITS	
Pixels	Bari-Politecnico, CERN, Padua, Rome, Salerno
Drift	Catania, Heidelberg-MPI, Rez, Turin, Trieste
Strips	Kharkov, Kiev, Nantes, Strasbourg, Turin, Utrecht-Nikhef
Mechanics, Cool. and Align.	Padua, St Petersburg, Strasbourg, Turin, Utrecht-Nikhef
TPC	Bratislava, CERN, Cracow, Frankfurt, GSI, Lund
Barrel PID	
either TOF-Pestov	JINR, GSI, Marburg, Moscow-ITEP, Moscow-MePHI, Novosibirsk
or TOF-PPC	CERN, JINR, Moscow-ITEP, Moscow-MePHI
High Momentum PID	
either RICH	Bari-University, CERN, Zagreb, Moscow-INR
or TOF-Scintillator	Catania, JINR
PHOS	Bergen, Heidelberg-University, Moscow-Kurchatov, Münster, Protvino, Prague
ZDC	Turin
MCP	Moscow-Kurchatov, St Petersburg
DAQ	Budapest, CERN, Oslo
Trigger	Birmingham, Košice
Muons	CERN, JINR, Clermont-Ferrand, Gatchina, Moscow-Kurchatov, Moscow-INR, Nantes, Orsay
Options	
PMD	India
Scintillator Barrel	Protvino

tion, assembly and operation of the equipment.

The sharing is based on the interest expressed by the institutes, taking into account expertise, resources and ongoing R&D programmes. The assignment of responsibilities is neither complete nor final, and will be subject to negotiations with the respective funding agencies. A Memorandum of Understanding is expected to be established during 1996.

The ALICE Collaboration will regard a number of items as 'Common Projects' to which all institutes are expected to contribute. The list presently under discussion represents about one quarter of the overall costs.

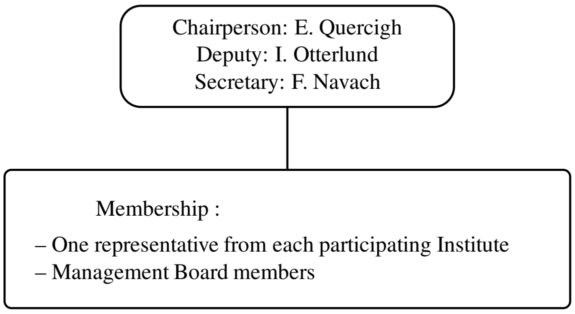

Figure 13.2: Collaboration Board.

13.4 Organization

The formal organization of ALICE as agreed by the Collaboration in 1994 [2] consists of a Collaboration Board (CB), a Management Board (MB) and a Technical Board (TB). Its present structure and composition is shown in Figs. 13.2 and 13.3.

This organization has been adopted for the technical proposal phase, and all nominations and elections are subject to confirmation in the course of 1996.

The CB is the governing body; it defines the scope of the ALICE project and takes decisions concerning modifications and relative priorities during its execution. Each institute, regardless of its size, has one representative in the CB; the right of vote, however, is restricted to institutes with at least three members. The CB is chaired by the elected Chairperson of the Collaboration, who also represents the CB in the Management Board. He ensures that the MB acts in conformity with the CB directives.

The MB supervises the progress of the experiment, along the lines defined by the CB. Decisions taken by the MB must be endorsed by the CB. For important matters, the MB submits a proposal to the CB, for approval. The

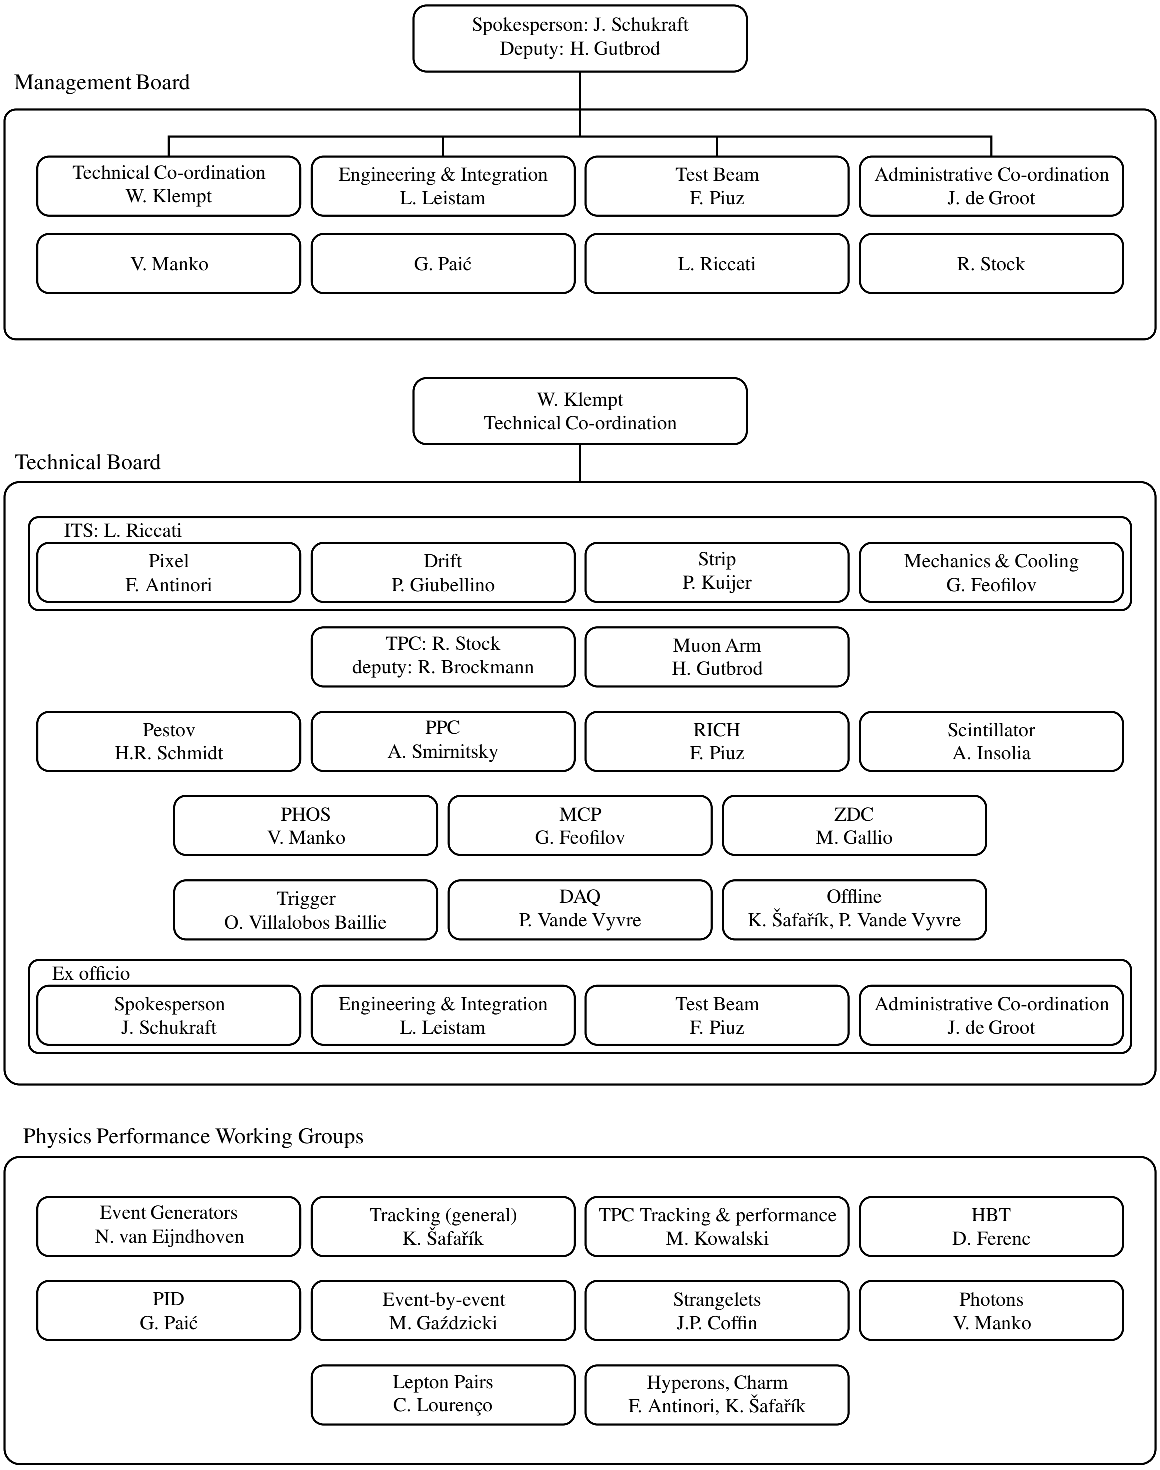

Figure 13.3: ALICE Organization.

MB appoints the Project Leaders and manages the human and financial resources of the Collaboration.

The design and co-ordination of the experiment is discussed and agreed upon in the TB. The TB, chaired by the Technical Co-ordinator, includes the Project Leaders and a number of *ex officio* or invited members.

The Spokesperson and Deputies are elected by the CB. The Spokesperson directs the ALICE project and has the overall responsibility towards the CB.

An important role in the management of the Collaboration is played by the Co-ordinators (Technical, Engineering & Integration and Administrative). The Technical Co-ordinator monitors and co-ordinates the activities of all the projects, and is responsible for technical planning, scheduling and the overall technical performance of the ALICE detectors. The Engineering & Integration Co-ordinator deals with issues concerning the integration of the sub-detectors, safety, infrastructure, experimental area, services, machine interface and the installation of the detectors. The Administrative Co-ordinator is responsible for the management of the common resources of the Collaboration.

The Project Leaders ensure the construction of specific sub-systems and, in agreement with the Technical Co-ordinator and all institutes involved, define and distribute the tasks to be carried out within their project.

References

Chapter 1

[1] H. Satz, Nucl. Phys. **A544** (1992) 371c.
[2] J. Schukraft, Nucl. Phys. **A583** (1995) 673c.
[3] ALICE Collaboration, Letter of Intent, CERN/LHCC 93–16.
[4] ALICE Collaboration, Addendum to the Letter of Intent, CERN/LHCC 95–24.
[5] F. Karsch and H. Satz, Z. Phys. **C51** (1991) 209.
[6] L. Van Hove, Nucl. Phys. **A461** (1987) 3c.
[7] H.R. Schmidt and J. Schukraft, J. Phys. **G19** (1993) 1705.
[8] C. Berger and W. Wagner, Phys. Rep. **146** (1987) 1.
[9] S. Sadovsky, Internal Note ALICE 93–7.
[10] G. Baur, Z. Phys. **C54** (1992) 419.
[11] G. Baur and L.G. Ferreira Filho, Nucl. Phys. **A518** (1990) 786.
[12] J.K. Bienlein, Crystal Ball, DESY 92–083 (1992).
[13] ALICE Collaboration, Large cross-section pp physics, gamma-gamma collisions and color screening with the ALICE detector, CERN/LHCC 95–01.
[14] Y. Takahashi et al., CERN/LHCC 95–51.
[15] ALICE Collaboration, Answers to LHCC Questions, CERN/LHCC 93–59.

Chapter 2

[1] M. Campbell et al., in Proc. Nuclear Science Symposium, San Francisco, October 1995, to be published.
[2] ALICE Collaboration, Letter of Intent, CERN/LHCC 93–16.
[3] F. Meddi, Internal Note ALICE 92–02.
[4] RD19 1994 Status Report, CERN DRDC/94–51, and references therein.
[5] F. Antinori et al., in Proc. Quark Matter '95, Monterey, January 1995, Nucl. Phys. **A590** (1995) 139c.
[6] F. Antinori et al., Nucl. Instrum. Methods **A360** (1995) 91.
[7] F. Antinori et al., in Proc. Strangeness in Hadronic Matter, Tucson, January 1995, AIP Press, 67.
[8] E.H.M. Heijne et al., Nucl. Instrum. Methods **A348** (1994) 399.
[9] E.H.M. Heijne, private communication.
[10] E. Chesi and E. Cantatore, Internal Note ALICE 95–36.
[11] J.J. Jaeger and J. Waisbard, Internal Note ALICE 95–37.
[12] L. Gevorkov and V. Laptev, Internal Note ALICE 95–38.
[13] A. Fasso et al., in Proc. Workshop on Simulating Accelerator Radiation Environments, Santa Fe, January 1993 (LA–12835–C);
A. Fasso and A. Ferrari, private communication.
[14] E. Beuville et al., Nucl. Instrum. Methods **A288** (1990) 68.
[15] E. Gatti and P. Rehak, in Proc. DPF Workshop on Collider Detectors, LBL–15973 UC–37, Conf. 830224, 1983.
[16] P. Rehak et al., Nucl. Instrum. Methods **A248** (1986) 367.
[17] M. Clemen et al., Nucl. Instrum. Methods **A316** (1992) 283.
[18] E. Gatti et al., Nucl. Instrum. Methods **A274** (1989) 469.
[19] A. Vacchi et al., Nucl. Instrum. Methods **A306** (1991) 187.
[20] B. Batunya et al., Internal Note ALICE 94–11.
[21] B. Batunya et al., Internal Note ALICE 94–31.
[22] O. Godisov et al., Workshop on Electronics and Detector Cooling (WELDEC), October 1994, Lausanne.
[23] W. Chen et. al., IEEE Trans. Nucl. Sci. **39** (1992) 619.
[24] U. Faschingbauer et. al., in Proc. 7th European Symposium on Semiconductor Detectors, Schloss Elmau, Germany, 7–10 May 1995.
[25] STAR Collaboration, Conceptual Design Report for the Solenoidal Tracker at RHIC.

[26] S. Beolé et al., in Proc. 6th Pisa Meeting on Frontier Detectors for Frontier Physics, Pisa, 1994, to be published in Nucl. Instrum. Methods.
[27] G. Gramegna et al., Nucl. Instrum. Methods **A360** (1995) 110.
[28] S. Beolé et al., in Proc. 7th European Symposium on Semiconductor Detectors, Schloss Elmau, Germany, 1995.
[29] V. Petracek et al., Internal Note ALICE 95–2.
[30] T. Angelescu, A. Mihul et al., Radiation damage in silicon detectors UBpub.-EPPG/phys. 1, (1993); UBpub.-EPPG/phys. 21, (1995);
J.E. Gover, J.R. Srour, Basic radiation effects in nuclear power electronics technology, American Nuclear Society, 1985, preprint SAND85–0766.
[31] S. Sotthibandhu, Radiation damage studies of silicon detectors, EX-RALT-025 SW 9505.
[32] G.V. Russo et al., in Proc. Fourth Annual Conference on Fuzzy Theory & Tech., Wrightsville, 28 September 1995, to be published.
[33] W. Dabrowski et al., Proc. Int. Conference on Advanced Technology and Particle Physics, Como, October 1994, to be published in Suppl. of Nucl. Phys. B.
[34] S. Kleinfelder, IEEE Trans. Nucl. Sci. **37** (1990) 1230.
[35] F. Anghinolfi et al., CERN/ECP 95–10 (1995).
[36] G.V. Russo et al., Internal Note ALICE 95–25.
[37] G.C. Bonazzola et al., Internal Note ALICE 95–53.
[38] L. Gevorkov et al., Internal Note ALICE 95–21.
[39] M. Masetti et al., AIHENP 1995, April 1995.
[40] V. Catania et al., IEEE Trans. Fuzzy Syst. **2** (1994) 9 3.
[41] A.S. Schwartz, Nucl. Instrum. Methods **A342** (1994) 218.
[42] G. Bochek et al., Internal Note ALICE 95–54.
[43] M. Monteno, Internal Note ALICE 94–13.
[44] O. Toker et al., Nucl. Instrum. Methods **A340** (1994) 572.
[45] Y. Hu et al., and Nucl. Instrum. Methods **A361** (1985) 568.
[46] L. Gevorkov et al., CERN/LHCC/95–56, and Proc. First Workshop on Electronics for LHC Experiments, Lisbon, 11–15 September 1995, to be published.
[47] V. Borchtchev et al., Internal Note ALICE 94–28.
[48] G. Feofilov et al., in Proc. International Workshop on Advanced Materials for High Precision Detectors, 28-30 September 1994, Archamps.
[49] L.Yu. Abramova et al., Internal Note ALICE 95–15.
[50] V.M. Dobulevitch et al., Internal Note ALICE 92–27.

Chapter 3

[1] W. B. Atwood et al., Nucl. Instrum. Methods **A306** (1991) 446;
The ALEPH Handbook, ALEPH 89–77.
[2] RD32 Proposal CERN/DRDC 92–32;
RD32 Status Report CERN/DRDC 94–10.
[3] P. Antilogus et al., Nucl. Instrum. Methods Phys. Res. **A283** (1989) 567.
[4] V. Eckardt et al., Nucl. Instrum. Methods **A315** (1992) 33.
[5] T. Alber et al., Nucl. Instrum. Methods **A349** (1994) 56.
[6] S.F. Biagi, Nucl. Instrum. Methods **A283** (1989) 716.
[7] V. Eckardt and S. Schönfelder, private communication.
[8] M. Price, CERN/EF/1485H/MP/em.
[9] W. Adam et al., Nucl. Instrum. Methods **A343** (1994) 68.
[10] STAR Technical Design Report, PUB–5347.
[11] A. Oed, Nucl. Instrum. Methods **A263** (1988) 351.
[12] M. Konrad, Diploma thesis, TUM Munich, 1994.
[13] A. Oed, Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[14] U. Frankenfeld, Diploma thesis, Frankfurt Univ., 1995.
[15] L. Alunni et al., Nucl. Instrum. Methods **A348** (1994) 344.
[16] M. Ayachi, J. D. Berst, J.-P. Blonde, J.-P. Coffin, F. Jundt, C. Kuhn and F. Rami, LEPSI preprint 95–12.
[17] Status Report on the FERMI project, CERN/LHCC 95–28;

LERB Status Report RD16.
[18] RD23 Status Report CERN/DRDC 93–35.
[19] C. Markert, Diploma thesis, Frankfurt Univ., 1995.
[20] C. Brand et al., Nucl. Instrum. Methods **A283** (1989) 567.
[21] A. Lebedev et al., Internal Note STAR, 1995.
[22] M. Bosteels, private communication.
[23] M. Fuchs, Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[24] I. Lehrhaus et al., IEEE Trans. Nucl. Sci. **30** (1983) 50.
[25] W. Blum, L. Rolandi, Particle Detection with Drift Chambers, Springer-Verlag, Berlin and Heidelberg, 1993.

Chapter 4

[1] Yu.N. Pestov, Proc. 4th San Miniato Topical Seminar, World Scientific, 1991, 156.
[2] A.R. Frolov, Yu.N. Pestov and V.V. Primachek, Nucl.Instrum. Methods **A307** (1991) 497.
[3] Yu.N. Pestov, Nucl. Instrum. Methods **A265** (1988) 150.
[4] W.B. Atwood et al., Nucl. Instrum. Methods **206** (1983) 99.
[5] Yu.N. Pestov, Nucl. Instrum. Methods **196** (1982) 45.
[6] V.V. Parkhomchuck, Yu.N. Pestov and N.V. Petrovykh, Nucl. Instrum. Methods **93** (1971) 269.
[7] ALICE Collaboration, Letter of Intent, CERN/LHCC 93–16.
[8] T. Weiland, Part. Accel. **15** (1984) 245.
[9] Yu. Pestov, M.A. Tiunov, preprint BNIP 94–77 (1994).
[10] E. Badura et al., preprint GSI–95–74 (1995).
[11] V.D. Laptev and Yu. Pestov, Pribory i Tekhnika Eksperimenta, **6** (1975) 41.
[12] Yu. Galaktionov et al., Nucl. Instrum. Methods **A317** (1992) 116.
[13] J. Christiansen, Z. Angew., Phys. **4** (1956) 46;
R. Bouclier et al., Report CERN–NP 71–12;
A. Peisert, Nucl. Instrum. Methods **A217** (1983) 229.
[14] A. Arefiev et al., preprint INFN–DFF 173–10, 1992.
[15] A. J. Davies and C. J. Evans, CERN 73–10.
[16] A. Sharma and F. Sauli, CERN/PPE 93–50;
A. Zarubin, Nucl. Instrum. Methods **A283** (1989), 411.
[17] V. Akimov et al., Nucl. Instrum. Methods **A344** (1994) 120.
[18] A. Smirnitsky, Internal Note ALICE 92–15;
F. Navach et al., Internal Note ALICE 93–36;
ALICE Collaboration, Internal Document ALICE/MIN 94–02, 94–07, and 95–04.
[19] ALICE Collaboration, Internal Document ALICE/MIN 95–03.
[20] A. Rudge and P. Weilhammer, Proc. 3rd Int. Conference on Electronics for Future Colliders, Chestnut Ridge, NY, 1993, 181, edited by G.J. Blanar and R.L. Summer.
[21] A. Martemyanov, WASA MEMO, Uppsala, Sweden, 1994.
[22] R. D'Alessandro, INFN Firenze preprint DFF 185/4/1993.
[23] R. D'Alessandro et al., ITEP preprint 51–95 (to be published);
R. D'Alessandro et al., Proc. First Workshop on Electronics for LHC Experiments, Lisbon, 11–15 September 1995, to be published.
[24] A.R. Frolov, T.V. Oslopova, Yu. Pestov, Nucl. Instrum. Methods **A356** (1994) 497.
[25] C. Neyer and R. Schulze, Internal Note ALICE 94–07.
[26] J. Christiansen, preprints CERN/ECP 93–23 and CERN/ECP 94–25.
[27] B. Taylor (RD12 Collaboration), Private communication.
[28] R. Arnold et al., Nucl. Instrum. Methods **A273** (1988) 466.
[29] R. Baur et al., Nucl. Instrum. Methods **A343** (1994) 87.
[30] J. Seguinot et al., Nucl. Instrum. Methods **A297** (1990) 133.
[31] RD26 status report CERN/DRDC 94–49, 1994.
[32] HADES Collaboration, Proposal for a high-acceptance di-electron spectrometer, GSI, Darmstadt, 1994;
HERA–B Collaboration, Proposal for an experiment to study CP violation in the B system using an internal target at the HERA proton ring, DESY–PRC 94/02.
[33] A. Amadon, Etude de faisabilité d'un détecteur fast-RICH pour identifier les particules à l'usine B de SLAC, thèse à l'Université PARIS XI Orsay, 1995.

[34] A. Breskin, WIS–95/31/July–PH, Weizmann Institute, and Proc. of the RICH'95 Workshop, Uppsala, 1995, to be published in Nucl. Instrum. Methods.
[35] F. Piuz, Proc. of the RICH'95 Workshop, Uppsala, 1995, to be published in Nucl. Instrum. Methods.
[36] Ph. Miné, Nucl. Instrum. Methods **A343** (1994) 99;
G. Malamud et al., Nucl. Instrum. Methods **A343** (1994) 121;
A. Breskin, Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[37] A. Braem et al., Nucl. Instrum. Methods **A343** (1994) 172;
J. Almeida et al., Nucl. Instrum. Methods **A348** (1994) 216.
[38] J. Almeida et al., CERN/PPE/95–63, and Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[39] C. Lu et al., Princeton/HEP/95–4;
C. Lu et al., Princeton/HEP/94–10.
[40] M. Staric et al., Nucl. Instrum. Methods **A343** (1994) 192;
S. Korpar et al., Proc. of the RICH'95 Workshop, Uppsala, 1995, to be published in Nucl. Instrum. Methods.
[41] J. Friese and N. Lockyers, Proc. of the RICH'95 Workshop, Uppsala, 1995, to be published in Nucl. Instrum. Methods.
[42] A. Buzulutskov et al., Heat enhancement of the photoyield from CsI, NaI and CuI photocathodes, WIS–95/37/Aug–PH, Weizmann Institute, submitted to Nucl. Instrum. Methods.
[43] J. Almeida et al., Nucl. Instrum. Methods **A361** (1995) 524;
J. Almeida et al., Microanalysis surface studies and photoemission properties of CsI photocathodes, WIS–95/14/March–PH, Weizmann Institute, and Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[44] A. Breskin et al., WIS–95/4/Feb–PH, Weizmann Institute, and Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[45] A. Buzulutskov et al., WIS–95/32/July–PH, Weizmann Institute, and in Proc. of the RICH'95 Workshop, Uppsala, 1995, to be published in Nucl. Instrum. Methods.
[46] R. Aleksan et al., Quantum efficiency and ageing tests of CsI photocathode, Saclay DAPNIA, in preparation.
[47] A. Di Mauro et al., Nucl. Instrum. Methods **A343** (1994) 284.
[48] J.C. Santiard, CERN/ECP/MIC–1995.
[49] W. Beusch et al., CERN/ECP/MIC Internal Note No.4, 1993, to be published in Nucl. Instrum. Methods.
[50] F. Anghinolfi et al., Development of an eight-bits ADC based on charge redistribution technique, CERN/ECP/MIC, to be published in Nucl. Instrum. Methods.
[51] P. Aarnio et al., Nucl. Instrum. Methods **A303** (1991) 233.
[52] W. Adam et al., Nucl. Instrum. Methods **A338** (1994) 284.
[53] G. Lenzen et al., Nucl. Instrum. Methods **A343** (1994) 268.
[54] G. Cinà et al., Internal Note ALICE 95–51;
S. Albergo et al., Nucl. Instrum. Methods **A362** (1995) 423.
[55] Memorandum of Understanding CERN–INFN–JINR on the R&D activities concerning the realization and test for new scintillation counters for the TOF system foreseen for the ALICE experiment at LHC.

Chapter 5

[1] F. Binon et al., Nucl. Instrum. Methods **188** (1981) 507.
[2] R. Albrecht et al., Limits on production of direct photons in 200 A GeV ^{32}S+Au collisions, to be published in Phys. Rev. Lett.
[3] CMS Collaboration, Technical Proposal, CERN/LHCC 94–38.
[4] W. van Loo, Phys. Stat. Sol. **28** (1975) 227.
[5] W. Moses and S.E. Derenzo, IEEE Trans. Nucl. Sci. **36** (1989) 173.
[6] M. Kobayashi et al., Nucl. Instrum. Methods **A333** (1993) 429.
[7] V.G. Baryshevsky et al., Nucl. Instrum. Methods **A322** (1992) 231.
[8] P. Lecoq et al., Lead tungstate scintillators for LHC e.m. calorimetry, CMS TN/94–308 (1994).
[9] Yu. Borodenko et al., Int. Conf. on Nuclear Tracking and Radiation Measurements, Amsterdam, 1995.
[10] K.-H. Kampert et al., A high resolution BGO calorimeter with longitudinal segmentation, Münster University preprint IKP-MS-94/0301(1994).
[11] PHENIX Collaboration, Conceptual Design Report (1993), 9.
[12] ALICE Collaboration, Answers to LHCC Questions, Internal Note ALICE 93–33.

[13] A. Fyodorov et al., Internal Note ALICE 94–29.
[14] G. Alexeev et al., Nucl. Instrum. Methods **A364** (1995) 307.
[15] A.L. Wintenberg et al., Proc. IEEE Nuclear Science Symposium **1** (1994) 493.
[16] A.L. Wintenberg et al., Proc. Int. Conf. on Electronics for Future Colliders, LeCroy Corp., May 1994.
[17] M.N. Ericson et al., to be published in IEEE Trans. Nucl. Sci.
[18] M.L. Simpson et al., presented at the 1995 IEEE Nuclear Science Symposium.
[19] M.L. Simpson et al., to be published in IEEE Trans. Nucl. Sci.
[20] D.M. Binkley et al., IEEE Trans. Nucl. Sci. **41** (1994) 4.
[21] V. Radeka, IEEE Trans. Nucl. Sci, **15** (1968) 268.
[22] G. Blanar et al., Nucl. Instrum. Methods **203** (1982) 213.

Chapter 6

[1] F. Karsch, Proc. of the Large Hadron Collider Workshop, Aachen, Germany, 1990, vol. II, ECFA 90–133, CERN 90–10, 1141.
[2] T. Matsui and H. Satz, Phys. Lett. **178B** (1986) 416.
[3] ALICE Collaboration, Addendum to the Letter of Intent, CERN/LHCC 95–24.
[4] Z.V. Borisovskaya et al., Internal Note ALICE 95–08.
[5] Vector Fields Limited, 24 Bankside, Kidlington, Oxford OX5 1JE, England.
[6] V.I. Klyukhin et al., Internal Note ALICE 95–46.
[7] K. Eggert and A. Morsch, AT Group Report 95–01 (DI) and Internal Note ALICE 95–05.
[8] D. Jouan and D.C. Zhou, Internal Note ALICE 95–04.

Chapter 7

[1] S. Wang et al., Phys. Rev. Lett. **74** (1994) 2646.
[2] W.F. Muller et al., GSI Nachrichten 03–95.
[3] A. Dabrowska et al., Phys. Rev. **D47** (1993) 1751.
[4] P. Gorodetsky et al. Nucl. Instrum. Methods **A361** (1995) 161.
[5] G. Anzivino et al., Nucl. Instrum. Methods **A357** (1995) 369.
[6] G. Anzivino et al., Nucl. Instrum. Methods **A357** (1995) 380.
[7] E. Chiavassa et al., Proc. of the Wire Chamber Conference, Vienna, 1995, to be published in Nucl. Instrum. Methods.
[8] G.A. Feofilov et al., Internal Note ALICE 95–07.
[9] A. Baldin et al., Nucl. Instrum. Methods **A323** (1992) 439–444.
[10] A. Kolojvari et al., Internal Note ALICE 93–15.
[11] G.A. Feofilov and F.F. Valiev, Internal Note ALICE 92–9.
[12] R. Brockmann, Proc. Conf. of Physics and Astrophysics of Quark-Gluon Plasma, World Scientific, Singapore, 1988, 617;
P. Giubellino et al., Nucl. Instrum. Methods **A279** (1989) 259.
[13] T.F. Gynzel et al., preprint HD–PY–91/6, 20 November 1991.
[14] Hamamatsu Corporation, Technical Information, September 1991.
[15] G.I. Veselov et al., Microelectronic UHF devices, Vishaya Shkola, Moscow, 1988, 68, (in Russian).
[16] Design and Calculations for the Microstrip Lines, ed. I.S. Kovalev, Moscow, Sov.Radio, 1974, 180, (in Russian).
[17] G. Feofilov et al., Internal Note ALICE 95–39.
[18] I.A. Novikov, Advance Programm on CLEO, Proc. OELS'95, Baltimore, 22–26 May 1995, 86.
[19] I.A. Novikov and A.L. Dmitriev, The spectral multi-demultiplexer on the holographic optical wedge element basis for the fiber optics communication systems, to be published in Jour. of Applied Optics.
[20] B.G. Taylor, CERN/LHCC/95–56 (1995) 180 (see also http://www.cern.ch/TTC/intro.html).
[21] L.G. Efimov and O. Villalobos-Baillie, Internal Note ALICE 95–40.
[22] F.F. Valiev et al., JINR Rapid Communications No. 4 **50**-91, Dubna, 1991, 27.
[23] K. Oba et al., IEEE Trans. Nucl. Sci., **28** (1981) 705–711.
[24] A.A. Baldin et al., JINR Rapid Communications, No. 2 **65**-94, 33, and Internal Note ALICE 94–05.
[25] F. Valiev et al., Internal Note ALICE 95–22.

[26] S. Sadovsky, Internal Note ALICE 95–45.
[27] Y. P. Viyogi et al., Internal Note ALICE 95–20.
[28] M.M. Aggarwal et al., VECC/EXP/95–02, to be published in Nucl. Instrum. Methods Phys. Res.
[29] CMS Collaboration, Technical Proposal, CERN/LHCC 94–38.
[30] RD36 Status Report, CERN/DRDC/94–47.

Chapter 8

[1] P. Lefèvre, CERN–PS/95–09 (DI) and LHC Note 324;
P. Lefèvre and D. Möhl, CERN–PS/93–45 (DI) and LHC Note 257.
[2] D. Brandt, SPS/AMS/Note/89–6, LHC/Note 87, (1989) and references therein;
D. Brandt, E. Brouzet and J. Gareyte, LHC Note 208.
[3] K. Hübner and E. Keil, IEEE Trans. Nucl. Sci., **32** (1985) 1632;
E. Fischer, CERN–ISR/67–16 (VA).
[4] A.G. Mathewson, Private communication.
[5] J. Bartke, J. Mod. Phys. **A4** (1989), 1319.
[6] G. Baur and C.A. Bertulani, Nucl. Phys. **A505** (1989) 835, and references therein.
[7] J.D. Jackson, Classical electrodynamics, 2nd ed. Wiley, New York,, Chapter 15, 619;
D. Olson et al., Phys. Rev. **C24** (1981) 1529.
[8] M. Vidovic, M. Greiner and G. Soff, Phys. Rev. **C48** (1993) 2011.
[9] G. Baur and C.A. Bertulani, Phys. Rep. **163** (1988) 299.
[10] H.L. Bradt and B. Peters, Phys. Rev. **77** (1950) 54.
[11] P.J. Karol, Phys. Rev. **C11** (1975) 1203.
[12] G. Fey, Z. Phys. **256** (1973) 401.
[13] L. Evans and B. Zotter, CERN–SPS/80–15.
[14] L. Vos, CERN–SPS/87–3 (AMS), LHC Note 64.
[15] The LHC Study Group, LHC: The Large Hadron Collider Conceptual Design, CERN–AC/95–05 (LHC).
[16] D. Brandt, K. Eggert and A. Morsch, Luminosity considerations for different ion species, CERN–AT/94–05 (DI), LHC Note 264.
[17] K. Eggert, K. Honkavaara and A. Morsch, Luminosity considerations for the LHC, CERN–AT/95–04 (DI), LHC Note 263.
[18] P. Lefèvre and D. Möhl, CERN–PS/93–62 (DI) and LHC Note 259.
[19] ALICE collaboration, LHCC/94-15.
[20] G. Wrochna (CMS Collaboration), Private communication.
[21] A. Kolojvari, Internal Note ALICE 93–01.
[22] ALICE Collaboration, Answers to LHCC Questions, Internal Note ALICE 93–33.
[23] ALICE Collaboration, Addendum to the Letter of Intent, CERN/LHCC 95–24.

Chapter 9

[1] L.V. Bravina, N.S. Amelin and L.P. Csernai, KFKI–1993–11/A, 52–58.
[2] N.S. Amelin, M.A. Braun and C. Pajares, Phys. Lett. **B306** (1993) 312 and Z. Phys **C63** (1994) 507.
[3] A. Dabrowska et al. Phys. Rev. **D47** (1993) 1751.
[4] X.N. Wang and M. Gyulassy, Phys. Rev. **D44** (1991) 3501.
[5] K. Werner, Phys. Rep. **232** (1993) 87.
[6] K. Marks et al., Trigger Level 0 Architecture, STAR Collaboration Internal Document.
[7] V. Lindenstruth et al., Trigger Level 1 and 2 Architecture, STAR Collaboration Internal Document.
[8] R. Cardarelli et al., CERN/LHCC/95–56 (1995) 213.
[9] T.A. Trainor, Power Spectrum Analysis with Non-uniform Bin Size and Variable Bin Number, University of Washington Nuclear Physics Laboratory preprint, September 1994.
[10] B.G. Taylor, CERN/LHCC/95–56 (1995) 180.

Chapter 10

[1] RD23 Status Report, CERN/DRDC 94–38.

[2] O. Boyle, R. McLaren and E. van der Bij, Internal Note ALICE 95–35.
[3] RD31 Status Report, CERN/DRDC 95–14.
[4] G. Harangozo and G. Rubin, Internal Note ALICE 95–26.
[5] RD13 Technical Note 70 (1993).
[6] RD24 Status Report CERN/DRDC 95–42.
[7] T. Sutton et al., Fibre Channel: The Digital Highway Made Practical, Seagate Technical note (1994).
[8] Quantum press release, 17 April 1995, Milpitas, CA.
[9] M. Teener, New Technology, Proc. IEEE P1394 Serial Bus, Apple Computer, Cupertino, CA (1993), to be published.
[10] The Serial Storage Architecture SSA, SSA Project Office, IBM Storage Systems Division, San Jose, CA (1994).
[11] J.F. Heanue et al., Volume holographic storage and retrieval of digital data science, Science **265** (1994) 749.
[12] Technical Directions for the 21st Century, Chapter 2, Storage Technology Trends, NASA (March 1993).
[13] J.-C. Lehureau and J.-P. Castera, Proc. 13th IEEE Symp. on Mass Storage Systems, Annecy, 1994, 127.
[14] K.L. Kliewer, to be published in Proc. of the Computing in High Energy Physics Conference, Rio de Janeiro, 1995.
[15] R.I. Patel et al., The emerging role of the flexible media, Proc. IEEE, **81** (1993) 4.
[16] H.F. Hinteregger, Technical Article MI-0001, The National Media Laboratory, St. Paul, MN (1993).
[17] H.F. Hinteregger et al., IEEE Trans. Magn. **27** (1991) 3.
[18] The VLBA Tape Recorder, Interferometrics Inc., Leesbur Pike, Vienna, VA (1995).
[19] Meeting at CERN between DIGITAL and ALICE, January 1995.
[20] Storage Technology Assessment - Final Report, Technical Report RE-0016, The National Media Laboratory, St. Paul, MN (1994).
[21] Meeting at CERN between SONY, ALICE and NA49 (June 1995).
[22] RD45 Proposal, CERN/DRDC 94–50.
[23] B. Taylor, Proc. First Workshop on Electronics for LHC Experiments, Lisbon, 11–15 September 1995, to be published.
[24] H. Beker, Internal Note ALICE 95–42.
[25] D.C. Doughty, Proc. 1st Int. Data Acquisition Conference, Fermilab, IL, Fermilab-Conf-95–54, (1995) 56.
[26] H. Beker et al., to be published in Proc. of the Computing in High Energy Physics Conference, Rio de Janeiro, 1995.
[27] B. Wu, B. Kvamme and B. Skaali, Internal Note ALICE 94–34.
[28] B. Wu, B. Skaali and P. Vande Vyvre, presented at the 9th Conference on Real-Time Computer Applications in Nuclear, Particle and Plasma Physics (RT95), East Lansing, Minn., 23–26 May 1995, to be published in IEEE Trans. on Nuclear Science.
[29] B. Wu et al., Internal Note ALICE 95–31.
[30] MODSIM II, The language for object-oriented programming, reference manual, CACI Products Company, La Jolla, CA (1991).
[31] G. Harangozo et al., Internal Note ALICE 95–33.
[32] RD11 Technical Note 94–22 (1994).
[33] RD38 Proposal, CERN/DRDC 93–50.
[34] RD44 Proposal, CERN/DRDC 94–29.
[35] R. Zybert and P. Buncic, to be published in Proc. of the Computing in High Energy Physics Conference, Rio de Janeiro, 1995.

Chapter 11

[1] CMS Collaboration, Technical Proposal, CERN/LHCC 94–38.
[2] J. Bartke et al., Internal Note ALICE 93–12.
[3] ALICE Collaboration, Addendum to the Letter of Intent, CERN/LHCC 95–24.
[4] A. Kurepin, Internal Notes ALICE 93–03 and ALICE 95–47.
[5] P. Aurenche et al., Comp. Phys. Commun. **83** (1994) 107.
[6] H.-U. Bengtsson and T. Sjöstrand, Comp. Phys. Commun. **46** (1987) 43.
[7] F. Antinori, Internal Note ALICE 93–9.
N. van Eijndhoven et al., Internal Note ALICE 95–32.
[8] T. Sjöstrand, Comp. Phys. Commun. **39** (1986) 347.

T. Sjöstrand and M. Bengtsson, Comp. Phys. Commun. **43** (1987) 367.
[9] T. Alexopoulos et al., Phys. Rev. Lett. **64** (1990) 991.
[10] M. Bourquin and J. M. Gaillard, Nucl. Phys. **B114** (1976) 334.
[11] K. Werner, Phys. Rep. **232** (1993) 87.
[12] X. N. Wang and M. Gyulassy, Phys. Rev. **D 44** (1991) 3501.
[13] B. Andersson et al., Phys. Rep. **97** (1983) 31.
[14] B. Andersson et al., Nucl. Phys. **B281** (1987) 289;
B. Nilsson-Almqvist and E. Stenlund, Comp. Phys. Comm. **43** (1987) 387.
[15] A. Capella et al., Z. Phys. **C3** (1980) 329;
J. Ranft, Phys. Lett. **188B** (1987) 379;
J. Ranft, Phys. Rev. **D37** (1988) 1842;
A. Capella et al., Phys. Rep. **236** (1994) 227.
[16] J. Ranft, Phys. Rev. **D51** (1995) 64.
[17] P. Aurenche et al., Phys. Rev. **D46** (1992) 92;
F. W. Bopp et al., Phys. Rev. **D49** (1994) 3236.
[18] N. S. Amelin et al., Phys. Lett. **B306** (1993) 312;
N. S. Amelin et al., Z. Phys. **C63** (1994) 507.
[19] R. Brun et al., GEANT3, CERN program library Q123.
[20] N. van Eijndhoven, Internal Notes ALICE 93–05, ALICE 93–06 and ALICE 95–44.
[21] W.K. Wilson (STAR Collaboration), Proc. of pre-Conference Workshop Quark Matter 95, UCRL-ID-121571, 55.
[22] F. Antinori (WA97 Collaboration), Proc. of Quark Matter 95, Nucl. Phys. **A590** (1995) 139c.
[23] M. Gaździcki (NA49 Collaboration), Proc. of Quark Matter 95, Nucl. Phys. **A590** (1995) 197c.
[24] P. Billoir, Nucl. Instrum. Methods **A225** (1984) 352;
R. Frühwirth, Nucl. Instrum. Methods **A262** (1987) 444;
P. Billoir and S. Qian, Nucl. Instrum. Methods **A294** (1990) 219 and **A295** (1990) 492.
[25] M. Regler and R. Frühwirth, Proc. Advanced Study Institute on Techniques and Concepts in High Energy Phys., St. Croix, 1988.
[26] P.V.C. Hough, U.S. Patent 3069654, 1962;
R.O. Duda and P.E. Hart, Commun. ACM **15** (1972) 11;
B.H. Ballard, Pattern Recognition **13** (1981) 111.
[27] J. Blom et al., in Proc. Conference on Computing in High Energy Physics, Berkeley, 1994, 497, and Internal Note ALICE 93-32;
G.A. Feofilov et al., Internal Note ALICE 94–21.
[28] B. Denby, Comput. Phys. Commun. **49** (1988) 429;
C. Peterson, Nucl. Instrum. Methods **A279** (1989) 537;
M. Gyulassy and M. Harlander, Comput. Phys. Commun. **66** (1991) 31.
[29] G. Stimpfl-Abele and L. Garrido, Comput. Phys. Commun. **64** (1991) 46;
B. Batyunya et al., Internal Note ALICE 95-50.
[30] B. Batyunya and A. Zinchenko, Internal Notes ALICE 94–11 and ALICE 94–31.
[31] Schnedermann et al., Phys. Rev. **C48** (1993) 2462c.
[32] J. Cleymans and H. Satz, Zeit. Phys. **C57** (1993) 135.
[33] M. Asakawa, C.M. Ko, Nucl. Phys. **A572** (1994) 732.
[34] Y. Wang et al., Nucl. Phys. **A590** (1995) 539c.
[35] R.D. Pisarski and F. Wilczek, Phys. Rev. **D29** (1984) 338.
[36] E.V. Shuryak, Nucl. Phys. **A525** (1991) 3c.
[37] X.-N. Wang and M. Gyulassy, Phys. Rev. **D44** (1991) 3501;
X.-N. Wang, Private communication.
[38] S. Pratt, Phys. Rev. Lett. **53** (1984) 1219.
S. Pratt, Phys. Rev. **D33** (1986) 1314.
[39] B. Andersson and W. Hofmann, Phys. Lett. **B169** (1986) 364.
[40] K. Kolehmainen and M. Gyulassy, Phys. Lett. **B180** (1986) 203.
[41] G. Bertsch, M. Gong and M. Tohyama, Phys. Rev. **C37** (1988) 1896.
[42] A.N. Makhlin and Yu.M. Sinyukov, Yad. Fiz. **46** (1987) 637;
A.N. Makhlin and Yu.M. Sinyukov, Z. Phys. **C39** (1988) 69;
Yu.M. Sinyukov, Nucl. Phys. **A498** (1989) 151c;
Yu.M. Sinyukov and B. Lörstad, Z. Phys. **C61** (1994) 587.

[43] Y. Hama and S.S. Padula, Phys. Rev. **D37** (1988) 3237.
[44] T. Csörgö and S. Pratt, Workshop on Relativistic Heavy Ion Physics, Central Research Institute Budapest, KFKI–1991-28/A, 75, and LUND preprint LU TP 91–10;
S. Pratt, T. Csörgö and J. Zimanyi, Phys. Rev. **C42** (1990) 2646.
[45] U. Mayer, E. Schnedermann and U. Heinz, Phys. Lett. **B294** (1992) 69.
[46] B.R. Schlei et al., Phys. Lett. **B293** (1992) 275;
J. Bolz et al., Phys. Lett. **B300** (1993) 404.
[47] Yu.M. Sinyukov, Phys. Lett. **B356** (1995) 525.
[48] J.G. Cramer, Phys. Rev. **C43** (1991) 2798.
[49] D.V. Anchiskin, W.A. Zajc and G.M. Zinovjev, preprint JYFL No.14/95, Department of Physics, University of Jyvaskyla.
[50] D. Ferenc et al., (NA35 Collaboration) Proc. of the Int. Symposium on Quantum Interferometry Studies in High-Energy Nuclear Collisions, Hiroshima, Japan, April 1995.
[51] D.H. Boal, C. Gelbke and B.K. Jennings, Rev. Mod. Phys. **62** (1990) 553.
[52] M. Fuchs, private communication, to be published.
[53] D. Ferenc, Internal Note ALICE 93–14.
[54] B. Erazmus et al., Internal Note ALICE 95–43
[55] R. Lednicky et al., Internal Note ALICE 95–48, to be published in Phys. Lett. B (1995).
[56] E.V. Shuryak, Phys. Rev. **D42** (1990) 1764;
E.V. Shuryak, preprint SUNY-NTG/91/11 (1991).
[57] F. Karsch, K. Redlich and L. Turko, Z. Phys. **C60** (1995) 519;
R. Pisarski, Phys. Lett. **B110** (1982) 155;
G.E. Brown and M. Rho, Phys. Rev. Lett. **66** (1991) 2720.
[58] T. Matsui and H. Satz, Phys. Lett. **B178** (1986) 416.
[59] J. Böhm, Internal Note ALICE 93–10.
[60] B. Batyuya et al., Internal Note ALICE 95–24.
[61] R.V. Gavai et al., preprint GSI-94–76.
[62] E.V. Shuryak, Phys. Lett. **B78** (1978) 15.
[63] R. Hwa and K. Kajantie, Phys. Rev. **D32** (1985) 1109.
[64] M. Kataja et al., Phys. Rev. **D34** (1986) 2755.
[65] M. Neubert, Z. Phys. **C42** (1989) 231.
[66] M. Pluemer et al., Nucl. Phys. **A590** (1995) 511c.
[67] X.N. Wang and M. Gyulassy, Phys. Rev. Lett. **68** (1992) 1480.
[68] J.D. Bjorken, K.K. Kowalski and C.C. Taylor, 'Baked Alaska', Proc. Les Rencontres de Physique de la Vallée d'Aoste, edited by M. Greco, Editions Frontières (1993).
[69] L.T. Baradzei et al., Nucl. Phys. **B370** (1992) 365;
A.V. Kulik et al., preprint IHEP 85–17, Serpukhov, 1985.
[70] P.V. Ruuskanen, Proc. of the Large Hadron Collider Workshop, Aachen, Germany, 1990, vol. II, ECFA 90–133, CERN 90–10, 1166.
[71] R. Baier et al., Z. Phys. **C53** (1992) 433.
[72] J. Alam et al., Phys. Rev. **D48** (1993) 1117.
[73] Y.A. Tarasov and V.G. Antonenko, Z. Phys. **C66** (1995) 215.
[74] E.V. Shuryak, Phys. Rev. Lett. **70** (1993) 2241.
[75] D. Alde et al., preprint CERN–EP/87–28.
[76] A.A. Lednev, preprint IHEP 93–153, Protvino, 1993.
[77] A.A. Lednev, preprint IHEP 93–152, Protvino, 1993.
[78] S.A. Sadovsky, preprint IHEP 85–35, Serpukhov, 1985.
[79] R. Albrecht et al., Nucl. Phys. **A590** (1995) 81c;
R. Santo et al., Nucl. Phys. **A566** (1994) 61c.
[80] F. Berger et al., Nucl. Instrum. Methods **A321** (1992) 152.
[81] K. Kinder-Geiger, Phys. Rep. **258** (1995) 237.
[82] R. Albrecht et al., Limits on Production of Direct Photons in 200 A GeV ^{32}S + Au Collisions, to be published in Phys. Rev. Lett.
[83] A. Lebedev et al., preprint of Kurchatov Institute, to be published.
[84] A. Lebedev et al., Nucl. Phys. **A566** (1994) 355c.
[85] S. Fokin et al., GSI scientific report 1992, March 1993, 71.
[86] ALICE Collaboration, Answers to LHCC Questions, Internal Note ALICE 93–33.

[87] P. Braun-Munzinger et al., PHENIX note 79 (1993).
[88] R. Albrecht et al., Z. Phys. **C51** (1991) 1.
[89] P. Koch, B. Müller and J. Rafelski, Phys. Rep. **142** (1986) 167.
[90] K. Šafařík, Internal Note ALICE 95–14, and Proc. of pre-Conference Workshop Quark Matter 95, UCRL-ID-121571, 55.
[91] E.V. Shuryak, Sov. J. Nucl. Phys. **28** (1978) 408.
[92] B. Müller and X.-N. Wang, Phys. Rev. Lett. **68** (1992) 2437.
[93] K. Kinder-Geiger, Phys. Rev. **D48** (1993) 4129.
[94] A. Shor, Phys. Lett. **B215** (1988) 375.
[95] W. Greiner, H. Stöcker and A. Gallmann, 'The nuclear equation of state', NATO ASI series B: Physics, Vol. 335 (1994), Plenum Press, New York.
[96] J.P. Coffin et al., Internal Note ALICE 95–49.
[97] C. Spieles, L. Gerland, H. Stöcker, W. Greiner, C. Greiner, C. Kuhn and J.P. Coffin, submitted to Phys. Rev. Lett.
[98] E. Gladysz-Dziadus and A. Panagiotou, Internal Note ALICE 95–18.
[99] D. Hahn and H. Stöcker, Nucl. Phys. **A476** (1988) 718;
J. Konopka, private communication.
[100] D. Buskulic et al., Z. Phys. **C66** (1995) 355.
[101] R. Stock, Proc. of a NATO Advanced Research Workshop on the 'Hot Hadronic Matter: Theory and Experiment', June 27–July 1, 1994, Divonne, France, edited by J. Letessier, H. Gutbrod and J. Rafelski.
[102] H.A. Gustafsson et al., Phys. Rev. Lett. **52** (1984) 1590;
R. Renfordt et al., Phys. Rev. Lett. **53** (1984) 763.
[103] A.P. Vischer, Nucl. Phys. **A590** (1995) 565c;
J.I. Kapusta and A.P. Vischer, University of Minnesota Report, NUC-MINN-95/19-T (1995).
[104] K. Werner, Nucl. Phys. **A566** (1994) 477c.
[105] ALICE Collaboration, Answers to LHCC Questions, CERN/LHCC 93–59.
[106] G. Baym, L.L. Frankfurt and M. Strikman, Nucl. Phys. **A566** (1994) 149c.
[107] A.I. Golokhvastov, Z. Phys. **C26** (1984) 469.
[108] M. Gaździcki and St. Mrówczyński, Z. Phys. **C54** (1992) 127.
[109] St. Mrówczyński, Phys. Lett. **B315** (1993) 118.
[110] St. Mrówczyński, Phys. Rev. **C49** (1994) 2191.
[111] N.A. Krall and A.W. Trivelpiece, Principles of Plasma Physics, McGraw-Hill, New York, 1973.
[112] J.-Y. Ollitrault, Nucl. Phys. **A590** (1995) 561c.
[113] F. Karsch and E. Laermann, Rep. Prog. Phys. **56** (1993) 1347.
[114] A.A. Anselm, Phys. Lett. **B217** (1989) 169;
A.A. Anselm and M.G. Ryskin, Phys. Lett. **B266** (1991) 482.
[115] J.P. Blaizot and A. Krzywicki, Phys. Rev. **D46** (1992) 246.
[116] J.D. Bjorken, Acta Phys. Pol. **B23** (1992) 637.
[117] K. Rajagopal and F. Wilczek, Nucl. Phys. **B379** (1993) 395 and **B404** (1993) 577.
[118] L.P. Csernai and I.M. Mishustin, Phys. Rev. Lett. **74** (1995) 5005.
[119] Proc. 11th Int. Conference on Ultra-Relativistic Nucleus–Nucleus Collisions, Monterey, 1995, edited by A.M. Poskanzer, J.W. Harris, and L.S. Schroeder, Nucl. Phys. **A590** (1995).
[120] L. Stodolsky, Phys. Rev. Lett. **75** (1995) 1044.
[121] M. Gaździcki, St. Mrówczyński and W. Retyk, Warsaw University Report IFD/4/1991 (1991), unpublished.

Chapter 12

[1] LHCC recommendations, CERN–LHCC 94–18.
[2] V.I. Klyukhin, Internal Note ALICE 95–46.
[3] The Alice Vacuum System, LHC Project Note 19, CERN-AT-Division, November 1995.
[4] A. Fasso et al., Proc. IV Int. Conference on Calorimeters and their Applications, La Biodola, Isola d'Elba, September 21–26 1993, World Scientific, Singapore 1994, 493.
[5] G. Shabratova, Internal Note ALICE 95–41.
[6] K. Hahn and J. Ranft, Phys. Rev. **D41** (1990) 1463.
[7] CERN/TIS–RP/IR/95–04, LHC Note 309, January 1995.
[8] CERN–TIS Safety Document D2, 1988.

Chapter 13

[1] CORE report, CERN/LHCC 93–44 and Internal Note ALICE 93–31.
[2] ALICE Organization, Internal Note ALICE 94–01.