

ALICE: Physics Performance Report, Volume I

ALICE Collaboration⁵

**F Carminati¹, P Foka², P Giubellino³, A Morsch¹, G Paic⁴, J-P Revol¹,
K Šafařík¹, Y Schutz¹ and U A Wiedemann¹ (editors)**

¹ CERN, Geneva, Switzerland

² GSI, Darmstadt, Germany

³ INFN, Turin, Italy

⁴ CINVESTAV, Mexico City, Mexico

Received 21 July 2004

Published 19 October 2004

Online at stacks.iop.org/JPhysG/30/1517

doi:10.1088/0954-3899/30/11/001

Abstract

ALICE is a general-purpose heavy-ion experiment designed to study the physics of strongly interacting matter and the quark–gluon plasma in nucleus–nucleus collisions at the LHC. It currently includes more than 900 physicists and senior engineers, from both nuclear and high-energy physics, from about 80 institutions in 28 countries.

The experiment was approved in February 1997. The detailed design of the different detector systems has been laid down in a number of Technical Design Reports issued between mid-1998 and the end of 2001 and construction has started for most detectors.

Since the last comprehensive information on detector and physics performance was published in the ALICE Technical Proposal in 1996, the detector as well as simulation, reconstruction and analysis software have undergone significant development. The Physics Performance Report (PPR) will give an updated and comprehensive summary of the current status and performance of the various ALICE subsystems, including updates to the Technical Design Reports, where appropriate, as well as a description of systems which have not been published in a Technical Design Report.

The PPR will be published in two volumes. The current Volume I contains:

1. a short theoretical overview and an extensive reference list concerning the physics topics of interest to ALICE,
2. relevant experimental conditions at the LHC,
3. a short summary and update of the subsystem designs, and
4. a description of the offline framework and Monte Carlo generators.

⁵ A complete listing of members of the ALICE Collaboration and external contributors appears on pages 1742–8.

Volume II, which will be published separately, will contain detailed simulations of combined detector performance, event reconstruction, and analysis of a representative sample of relevant physics observables from global event characteristics to hard processes.

(Some figures in this article are in colour only in the electronic version.)

Contents

1. ALICE physics—theoretical overview	1526
1.1. Introduction	1526
1.1.1. Role of ALICE in the LHC experimental programme	1527
1.1.2. Novel aspects of heavy-ion physics at the LHC	1527
1.1.3. ALICE experimental programme	1528
1.2. Hot and dense partonic matter	1529
1.2.1. QCD-phase diagram	1529
1.2.2. Lattice QCD results	1531
1.2.3. Perturbative finite-temperature field theory	1535
1.2.4. Classical QCD of large colour fields	1538
1.3. Heavy-ion observables in ALICE	1539
1.3.1. Particle multiplicities	1539
1.3.2. Particle spectra	1543
1.3.3. Particle correlations	1546
1.3.4. Fluctuations	1550
1.3.5. Jets	1551
1.3.6. Direct photons	1555
1.3.7. Dileptons	1557
1.3.8. Heavy-quark and quarkonium production	1558
1.4. Proton–proton physics in ALICE	1564
1.4.1. Proton–proton measurements as benchmark for heavy-ion physics	1564
1.4.2. Specific aspects of proton–proton physics in ALICE	1565
1.5. Proton–nucleus physics in ALICE	1569
1.5.1. The motivation for studying pA collisions in ALICE	1569
1.5.2. Nucleon–nucleus collisions and parton distribution functions	1569
1.5.3. Double-parton collisions in proton–nucleus interactions	1571
1.6. Physics of ultra-peripheral heavy-ion collisions	1572
1.7. Contribution of ALICE to cosmic-ray physics	1574
2. LHC experimental conditions	1578
2.1. Running strategy	1578
2.2. A–A collisions	1580
2.2.1. Pb–Pb luminosity limits from detectors	1580
2.2.2. Nominal-luminosity Pb–Pb runs	1580
2.2.3. Alternative Pb–Pb running scenarios	1580
2.2.4. Beam energy	1583
2.2.5. Intermediate-mass ion collisions	1583

2.3. Proton–proton collisions	1584
2.3.1. Standard pp collisions at $\sqrt{s} = 14$ TeV	1584
2.3.2. Dedicated pp-like collisions	1584
2.4. pA collisions	1585
2.5. Machine parameters	1585
2.6. Radiation environment	1588
2.6.1. Background conditions in pp	1588
2.6.2. Dose rates and neutron fluences	1589
2.6.3. Background from thermal neutrons	1591
2.7. Luminosity determination in ALICE	1592
2.7.1. Luminosity monitoring in pp runs	1593
2.7.2. Luminosity monitoring in Pb–Pb runs	1594
3. ALICE detector	1596
3.1. Introduction	1596
3.2. Design considerations	1597
3.3. Layout and infrastructure	1599
3.3.1. Experiment layout	1599
3.3.2. Experimental area	1600
3.3.3. Magnets	1601
3.3.4. Support structures	1601
3.3.5. Beam pipe	1603
3.4. Inner Tracking System (ITS)	1603
3.4.1. Design considerations	1604
3.4.2. ITS layout	1606
3.4.3. Silicon pixel layers	1607
3.4.4. Silicon drift layers	1612
3.4.5. Silicon strip layers	1617
3.5. Time-Projection Chamber (TPC)	1620
3.5.1. Design considerations	1620
3.5.2. Detector layout	1622
3.5.3. Front-end electronics and readout	1625
3.6. Transition-Radiation Detector (TRD)	1627
3.6.1. Design considerations	1627
3.6.2. Detector layout	1628
3.6.3. Front-end electronics and readout	1630
3.7. Time-Of-Flight (TOF) detector	1631
3.7.1. Design considerations	1631
3.7.2. Detector layout	1632
3.7.3. Front-end electronics and readout	1634
3.8. High-Momentum Particle Identification Detector (HMPID)	1637
3.8.1. Design considerations	1637
3.8.2. Detector layout	1637
3.8.3. Front-end electronics and readout	1640
3.9. PHOTon Spectrometer (PHOS)	1641
3.9.1. Design considerations	1641
3.9.2. Detector layout	1642
3.9.3. Front-end electronics and readout	1644

3.10. Forward muon spectrometer	1645
3.10.1. Design considerations	1645
3.10.2. Detector layout	1646
3.10.3. Front-end electronics and readout	1649
3.11. Zero-Degree Calorimeter (ZDC)	1649
3.11.1. Design considerations	1649
3.11.2. Detector layout	1650
3.11.3. Signal transmission and readout	1654
3.12. Photon Multiplicity Detector (PMD)	1654
3.12.1. Design considerations	1654
3.12.2. Detector layout	1655
3.12.3. Front-end electronics and readout	1658
3.13. Forward Multiplicity Detector (FMD)	1658
3.13.1. Design considerations	1658
3.13.2. Detector layout	1658
3.13.3. Front-end electronics and readout	1661
3.14. V0 detector	1664
3.14.1. Design considerations	1664
3.14.2. Detector layout	1664
3.14.3. Front-end electronics and readout	1665
3.15. T0 detector	1666
3.15.1. Design considerations	1666
3.15.2. Detector layout	1666
3.15.3. Front-end electronics and readout	1667
3.16. Cosmic-ray trigger detector	1668
3.16.1. Design considerations	1668
3.16.2. Detector layout	1668
3.16.3. Front-end electronics and readout	1668
3.17. Trigger system	1668
3.17.1. Design considerations	1668
3.17.2. Trigger logic	1670
3.17.3. Trigger inputs and classes	1671
3.17.4. Trigger data	1673
3.17.5. Event rates and rare events	1673
3.18. Data AcQuisition (DAQ) System	1674
3.18.1. Design considerations	1674
3.18.2. Data acquisition system	1677
3.18.3. System flexibility and scalability	1680
3.18.4. Event rates and rare events	1680
3.19. High-Level Trigger (HLT)	1684
3.19.1. Design considerations	1684
3.19.2. System architecture: clustered SMP farm	1688
4. Offline computing and Monte Carlo generators	1691
4.1. Offline framework	1692
4.1.1. Overview of AliRoot framework	1692
4.1.2. Simulation	1697
4.1.3. Reconstruction	1708

4.1.4. Distributed computing and the Grid	1711
4.1.5. Software development environment	1721
4.2. Monte Carlo generators for heavy-ion collisions	1723
4.2.1. HIJING and HIJING parametrization	1724
4.2.2. Dual-Parton Model (DPM)	1726
4.2.3. String-Fusion Model (SFM)	1728
4.2.4. Comparison of results	1728
4.2.5. First results from RHIC	1729
4.2.6. Conclusions	1730
4.2.7. Generators for heavy-flavour production	1731
4.2.8. Other generators	1732
4.3. Technical aspects of pp simulation	1736
4.3.1. PYTHIA	1736
4.3.2. HERWIG	1738
ALICE Collaboration	1742
External Contributors	1748
Acknowledgments	1748
References	1749




Figure I. Layout of the ALICE detector. For the sake of visibility, the HMPID detector is shown in the 12 o'clock position instead of the 2 o'clock position in which it will actually be positioned.

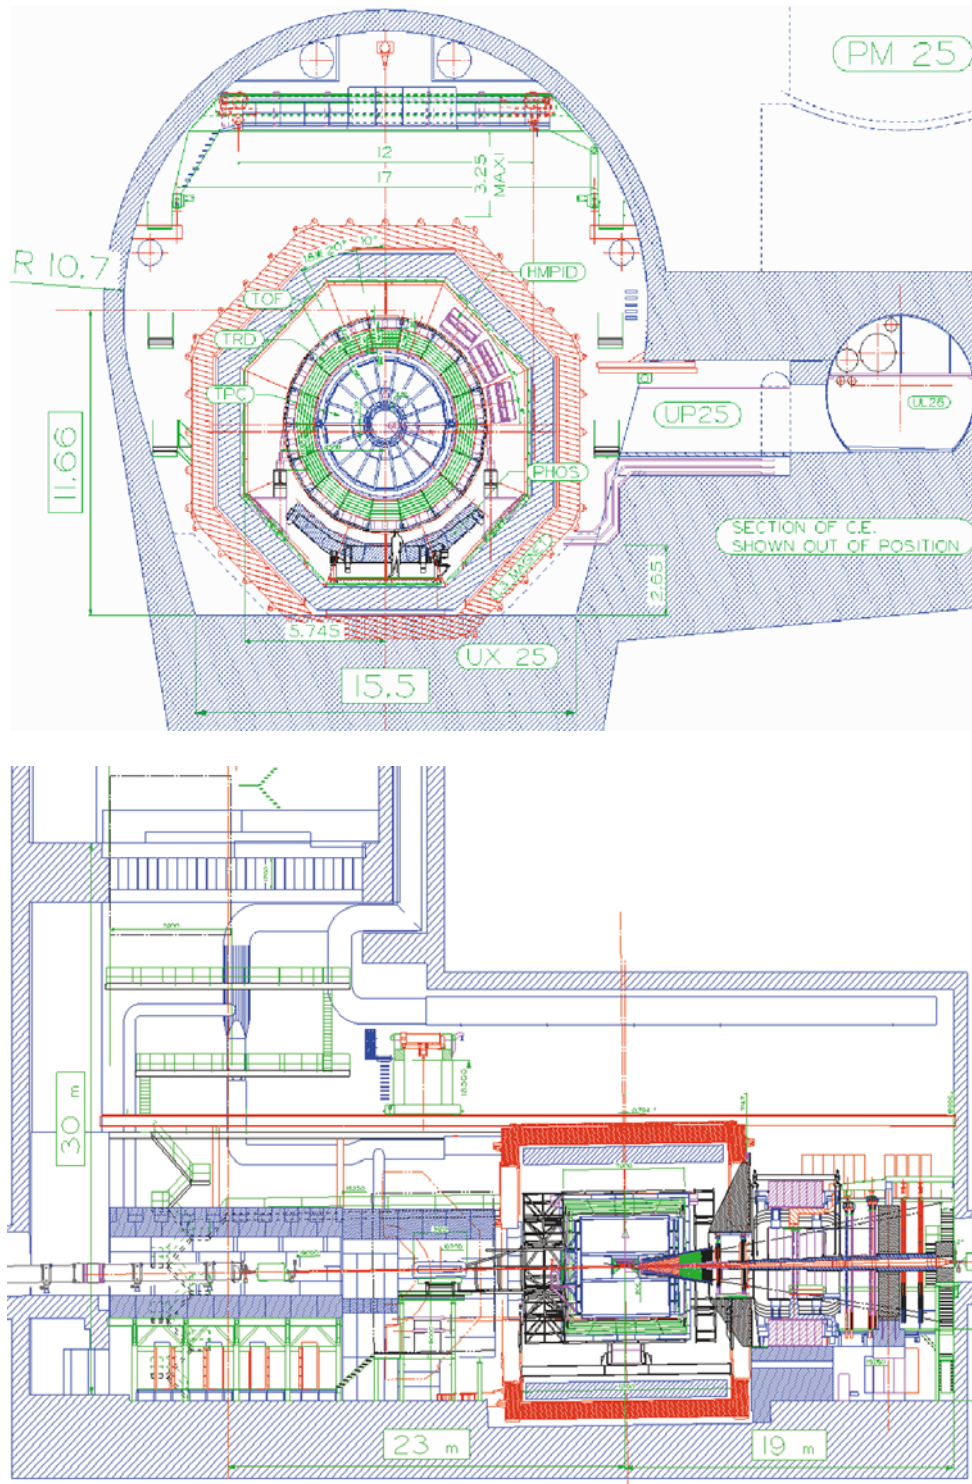


Figure II. Front (top) and side (bottom) view of the ALICE cavern.




Figure III. ALICE detector as described by the geometrical modeller in the simulation software.

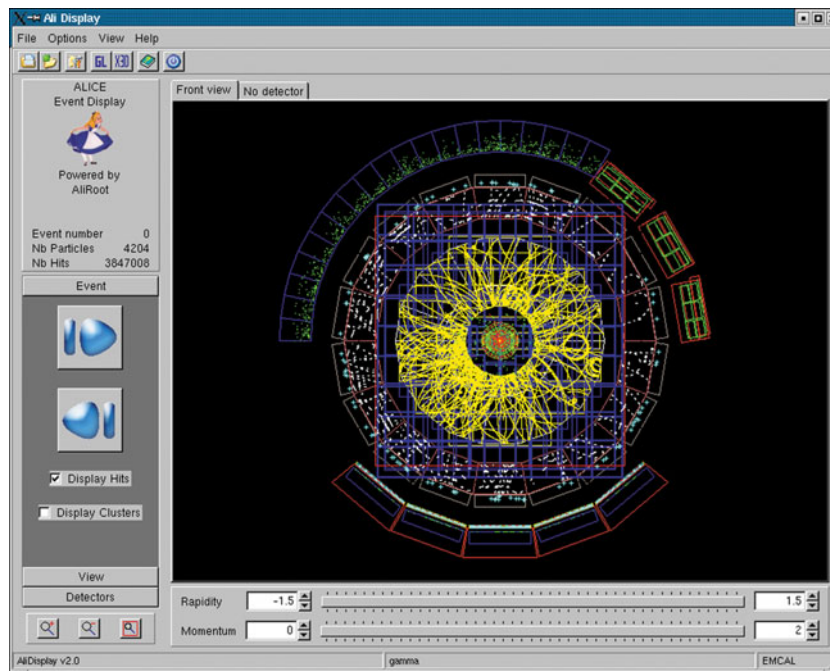


Figure IV. Prototype of the ALICE event display.

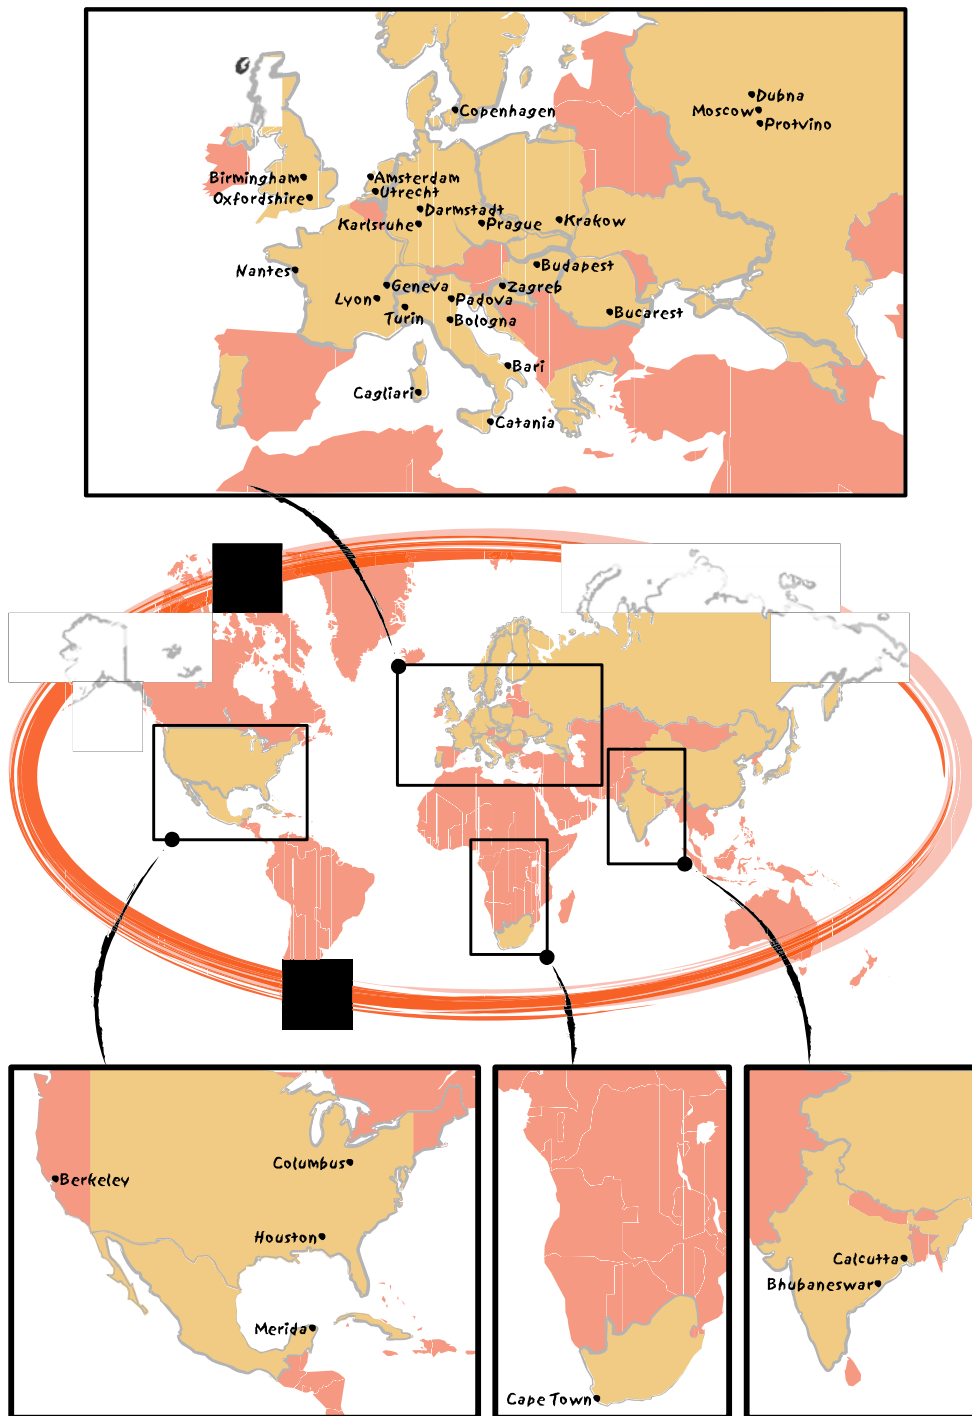


Figure V. World map of the ALICE GRID.

1. ALICE physics—theoretical overview

1.1. Introduction

High-energy physics has established and validated over the last decades a detailed, though still incomplete, theory of elementary particles and their fundamental interactions, called the Standard Model. Applying and extending the Standard Model to complex and dynamically evolving systems of finite size is the aim of ultra-relativistic heavy-ion physics. The focus of heavy-ion physics is to study and understand how collective phenomena and macroscopic properties, involving many degrees of freedom, emerge from the microscopic laws of elementary-particle physics. Specifically, heavy-ion physics addresses these questions in the sector of strong interactions by studying nuclear matter under conditions of extreme density and temperature.

The most striking case of a collective bulk phenomenon predicted by the Standard Model is the occurrence of phase transitions in quantum fields at characteristic energy densities. This affects crucially our current understanding of both the structure of the Standard Model at low energy and of the evolution of the early Universe. According to Big-Bang cosmology, the Universe evolved from an initial state of extreme energy density to its present state through rapid expansion and cooling, thereby traversing the series of phase transitions predicted by the Standard Model. Global features of our Universe, like baryon asymmetry or the large scale structure (galaxy distribution), are believed to be linked to characteristic properties of these phase transitions.

Within the framework of the Standard Model, the appearance of phase transitions involving elementary quantum fields is intrinsically connected to the breaking of fundamental symmetries of nature and thus to the origin of mass. In general, intrinsic symmetries of the theory, which are valid at high-energy densities, are broken below certain critical energy densities. Particle content and particle masses originate as a direct consequence of the symmetry-breaking mechanism. Lattice calculations of Quantum ChromoDynamics (QCD), the theory of strong interactions, predict that at a critical temperature of $\simeq 170$ MeV, corresponding to an energy density of $\varepsilon_c \simeq 1$ GeV fm⁻³, nuclear matter undergoes a phase transition to a deconfined state of quarks and gluons. In addition, chiral symmetry is approximately restored and quark masses are reduced from their large effective values in hadronic matter to their small bare ones.

In ultra-relativistic heavy-ion collisions, one expects to attain energy densities which reach and exceed the critical energy density ε_c , thus making the QCD phase transition the only one predicted by the Standard Model that is within reach of laboratory experiments. The longstanding main objective of heavy-ion physics is to explore the phase diagram of strongly interacting matter, to study the QCD phase transition and the physics of the Quark–Gluon Plasma (QGP) state. However, the system created in heavy-ion collisions undergoes a fast dynamical evolution from the extreme initial conditions to the dilute final hadronic state. The understanding of this fast evolving system is a theoretical challenge which goes far beyond the exploration of equilibrium QCD. It provides the opportunity to further develop and test a combination of concepts from elementary-particle physics, nuclear physics, equilibrium and non-equilibrium thermodynamics, and hydrodynamics in an interdisciplinary approach.

As discussed in this document, a direct link between the predictions of the Standard Model and experimental observables in heavy-ion collisions exists only for a limited number of observables so far. For example, for high-momentum processes, understanding the medium dependence constitutes an open field of research which is rapidly evolving. Therefore, the complexity of many collective aspects and bulk properties of heavy-ion collisions currently requires recourse to effective descriptions. These approaches range from idealized models,

like hydrodynamics, which emerge in a well-defined limit of multiparticle dynamics, up to detailed Monte Carlo simulations that—depending on the model—incorporate different microscopic pictures. Some of the theoretical approaches currently being pursued (e.g. lattice QCD, classical QCD) are directly related to the fundamental QCD Lagrangian, their range of applicability remaining to be determined in an interplay of experiment and theory. Others involve model parameters that are not solely determined by the Standard Model Lagrangian but provide powerful tools to study the origin of collective phenomena. The predictions of these approaches, their uncertainties, and the extent to which their comparison to experimental data can determine the underlying physics of various collision scenarios are discussed in this document.

1.1.1. Role of ALICE in the LHC experimental programme. One of the central problems addressed at the LHC is the connection between phase transitions involving elementary quantum fields, fundamental symmetries of nature and the origin of mass. Theory draws a clear distinction between symmetries of the dynamical laws of nature (i.e. symmetries and particle content of the Lagrangian) and symmetries of the physical state with respect to which these dynamical laws are evaluated (i.e. symmetries of the vacuum or of an excited thermal state). The experimental programme at the LHC addresses both aspects of the symmetry-breaking mechanism through complementary experimental approaches. ATLAS and CMS will search for the Higgs particle, which generates the mass of the electroweak gauge bosons and the bare mass of elementary fermions through spontaneous breaking of the electroweak gauge symmetry. They will also search for supersymmetric particles which are manifestations of a broken intrinsic symmetry between fermions and bosons in extensions of the Standard Model. LHCb, focusing on precision measurements with heavy b quarks, will study CP-symmetry-violating processes. These measure the misalignment between gauge and mass eigenstates which is a natural consequence of electroweak symmetry breaking via the Higgs mechanism. ALICE will study the role of chiral symmetry in the generation of mass in composite particles (hadrons) using heavy-ion collisions to attain high-energy densities over large volumes and long timescales. ALICE will investigate equilibrium as well as non-equilibrium physics of strongly interacting matter in the energy density regime $\varepsilon \simeq 1\text{--}1000 \text{ GeV fm}^{-3}$. In addition, the aim is to gain insight into the physics of parton densities close to phase-space saturation, and their collective dynamical evolution towards hadronization (confinement) in a dense nuclear environment. In this way, one also expects to gain further insight into the structure of the QCD phase diagram and the properties of the QGP phase.

Moreover, all LHC experiments are expected to have an impact on various astrophysical fields. For example, the top LHC energy $\sqrt{s} = 14 \text{ TeV}$ corresponds to 10^{17} eV in the laboratory reference frame and, therefore, LHC may contribute to the understanding of cosmic-ray interactions at the highest energies, especially, to the open question of the composition of primaries in the region around the ‘knee’ ($10^{15}\text{--}10^{16} \text{ eV}$).

1.1.2. Novel aspects of heavy-ion physics at the LHC. The nucleon–nucleon centre-of-mass energy for collisions of the heaviest ions at the LHC ($\sqrt{s} = 5.5 \text{ TeV}$) will exceed that available at RHIC by a factor of about 30, opening up a new physics domain. Historical experience suggests that such a large jump in available energy usually leads to new discoveries. Heavy-ion collisions at the LHC access not only a quantitatively different regime of much higher energy density but also a qualitatively new regime, mainly because:

1. *High-density (saturated) parton distributions determine particle production.* The LHC heavy-ion programme accesses a novel range of Bjorken- x values as shown in figure 1.1 where the relevant x ranges of the highest energies at SPS, RHIC and LHC with the




Figure 1.1. The range of Bjorken x and M^2 relevant for particle production in nucleus–nucleus collisions at the top SPS, RHIC and LHC energies. Lines of constant rapidity are shown for LHC, RHIC and SPS.

heaviest nuclei are compared. ALICE probes a continuous range of x as low as about 10^{-5} , accessing a novel regime where strong nuclear gluon shadowing is expected. The initial density of these low- x gluons is expected to be close to saturation of the available phase space so that important aspects of the subsequent time evolution are governed by classical chromodynamics. Theoretical methods based on this picture (Weizsäcker–Williams fields, classical Yang–Mills dynamics) have been developed in recent years.

2. *Hard processes contribute significantly to the total A – A cross section.* These processes can be calculated using perturbative QCD. In particular, very hard strongly interacting probes, whose attenuation can be used to study the early stages of the collision, are produced at sufficiently high rates for detailed measurements.
3. *Weakly interacting hard probes become accessible.* Direct photons (but in principle also Z^0 and W^\pm bosons) produced in hard processes will provide information about nuclear parton distributions at very high Q^2 .
4. *Parton dynamics dominate the fireball expansion.* The ratio of the lifetime of the QGP state to the time for thermalization is expected to be larger than at RHIC by an order of magnitude, so that parton dynamics will dominate the fireball expansion and the collective features of the hadronic final state.

In the following sections we present a short summary of the current theoretical understanding of hot and dense partonic matter, the main observables measured by ALICE, and how they test the properties of the matter created in heavy-ion collisions.

1.1.3. ALICE experimental programme. In general, to establish experimentally the collective properties of the hot and dense matter created in nucleus–nucleus collisions, both systematics- and luminosity-dominated questions have to be answered at LHC. Thus, ALICE aims firstly at accumulating sufficient integrated luminosity in Pb–Pb collisions at $\sqrt{s} = 5.5$ TeV per nucleon pair, to measure rare processes such as jet transverse-energy spectra up to $E_t \sim 200$ GeV and the pattern of medium induced modifications of bottomium bound states. However, the interpretation of these experimental data relies considerably on a systematic comparison with the same observables measured in proton–proton and proton–nucleus collisions as well as in

collisions of lighter ions. In this way, the phenomena truly indicative of the hot equilibrating matter can be separated from other contributions.

The successful completion of the heavy-ion programme thus requires the study of pp, pA and lighter A–A collisions in order to establish the benchmark processes under the same experimental conditions. In addition, these measurements are interesting in themselves. For example, the study of lighter systems opens up possibilities to study fundamental aspects of the interaction of colour-neutral objects related to non-perturbative strong phenomena, like confinement and hadronic structure. Also, due to its excellent tracking and particle identification capabilities, the ALICE pp and pA programmes complement those of the dedicated pp experiments. A survey of specific ALICE contributions to pp and pA physics will be given towards the end of this section. Finally, we shall also discuss other physics subjects within the capabilities of ALICE like ultra-peripheral collisions and cosmic-ray physics.

1.2. Hot and dense partonic matter

1.2.1. QCD-phase diagram. Even before QCD was established as the fundamental theory of strong interactions, it had been argued that the mass spectrum of resonances produced in hadronic collisions implies some form of critical behaviour at high temperature and/or density [1]. The subsequent formulation of QCD, and the observation that QCD is an asymptotically free theory led to the suggestion that this critical behaviour is related to a phase transition [2]. The existence of a phase transition to a new state of matter, the QGP, at high temperature has been convincingly demonstrated in lattice calculations.

At high temperature T and vanishing chemical potential μ_B (baryon-number density), qualitative aspects of the transition to the QGP are controlled by the chiral symmetry of the QCD Lagrangian [3]. This symmetry exists as an exact global symmetry only in the limit of vanishing quark masses. Since the heavy quarks (charm, bottom, top) are too heavy to play any role in the thermodynamics in the vicinity of the phase transition, the properties of 3-flavour QCD are of great interest. In the massless limit, 3-flavour QCD undergoes a first-order phase transition. However, in nature quarks are not massless. In particular, the strange-quark mass, which is of the order of the phase-transition temperature, plays a decisive role in determining the nature of the transition at vanishing chemical potential. It is still unclear whether the transition shows discontinuities for realistic values of the up-, down- and, strange-quark masses, or whether it is merely a rapid crossover. Lattice calculations suggest that this crossover is rather rapid, taking place in a narrow temperature interval around $T_c \sim 170$ MeV [4].

At low temperature and large values of the chemical potential, the basic properties of hadronic matter can be understood in terms of nearly degenerate, interacting, Fermi gases; nuclear matter consisting of extended nucleons at low density and a degenerate quark gas at high density. Although we have no direct evidence from analytic or numerical calculations within the framework of QCD, all approximate model-based calculations suggest that the transition between cold nuclear matter and the high-density phase is of first order. In the high-density phase there is a remnant attractive interaction among quarks, which perturbatively is described by one-gluon exchange, and in a non-perturbative approach can be described by instanton-motivated models. In a degenerate Fermi gas, an attractive interaction will almost inevitably lead to quark–quark pairing and thus to the formation of a colour superconducting phase [5–7]. Depending on the details of the interaction, one may expect this region of the QCD phase diagram to be subdivided into several different phases [8].

The generic form of the QCD phase diagram is shown in figure 1.2. A first-order phase transition is expected at low temperatures and high densities. If no first-order transition occurs




Figure 1.2. The phase diagram of QCD: the solid lines indicate likely first-order transitions. The dashed line indicates a possible region of a continuous but rapid crossover transition. The open circle gives the second-order critical endpoint of a line of first-order transitions, which may be located closer to the $\mu_B = 0$ axis.

at vanishing chemical potential and high temperature, then the first-order transition must end in a second-order critical point somewhere in the interior of the QCD phase diagram. The study of the phase diagram at small temperature and non-zero chemical potential is known to be notoriously difficult in lattice regularized QCD and it probably will still take some time until accurate results become available in this regime. Some progress has, however, been made in analysing the regime of non-zero chemical potentials at temperatures close to T_c [9–11]. The generic problem of Monte Carlo simulation of QCD with finite chemical potential, which is related to a complex structure of the fermion determinant, can partly be overcome in this regime. Using a two-parameter reweighing method close to T_c [9] first results on the location of the chiral critical point could be obtained. The results for the μ_B -dependence of the phase boundary at smaller values of μ_B are consistent with those obtained from a Taylor expansion of T_c in μ_B around $T_c(\mu_B = 0)$ [10]. The region of applicability of these approaches and uncertainties on the results due to small lattice size and large strange-quark mass is, however, not well established so far.

Figure 1.3 shows the results on the position of the phase boundary that were obtained using the methods indicated above. The dashed line in this figure represents an extrapolation of the leading μ_B^2 order of the Taylor expansion of T_c to larger values of the chemical potential. It is interesting to note that, within (the still large) statistical uncertainties, the energy density along this line is constant and corresponds to $\varepsilon_c \sim 0.6 \text{ GeV fm}^{-3}$, i.e. to the same value as that found from lattice calculations at $\mu_B = 0$. It is thus conceivable that the QCD transition indeed sets in when the energy density reaches a certain critical value. Such an assumption is often used in phenomenological approaches to determine the transition line in the μ_B - T plane.

Figure 1.3 also shows a compilation of the chemical freeze-out parameters, extracted from experimental data that were obtained in a very broad energy range from GSI/SIS through BNL/AGS, CERN/SPS and BNL/RHIC [12–16], together with the freeze-out condition of fixed energy per particle $\simeq 1 \text{ GeV}$. It is interesting to note that at the SPS and RHIC the chemical freeze-out parameters coincide with the critical conditions obtained from lattice calculations.




Figure 1.3. Lattice Monte Carlo results on the QCD-phase boundary [9, 10] shown together with the chemical freeze-out conditions obtained from a statistical analysis of experimental data (open symbols) [12, 13]. The dashed line represents the Lattice Gauge Theory (LGT) results obtained in [10] with the grey band indicating the uncertainty. The filled point represents the endpoint of the crossover transition taken from [9]. The solid line shows the unified freeze-out conditions of fixed energy per particle $\simeq 1$ GeV [12, 13].

1.2.2. Lattice QCD results. The most fascinating aspect of QCD thermodynamics is the theoretically well-supported expectation that strongly interacting matter can exist in different phases. Exploring the qualitatively new features of the QGP and making quantitative predictions of its properties is the central goal of numerical studies of equilibrium thermodynamics of QCD within the framework of lattice-regularized QCD.

Phase transitions are related to large-distance phenomena in a thermal medium. They go along with long-range collective phenomena, spontaneous breaking of global symmetries, and the appearance of long-range order. The sudden formation of condensates and the screening of charges indicate the importance of non-perturbative physics. In order to study such mechanisms in QCD we thus need a calculational approach that is capable of dealing with these non-perturbative aspects of the theory of strong interactions. It is precisely for this purpose that lattice QCD was developed [17]. Lattice QCD provides a first-principles approach to studies of large-distance, non-perturbative aspects of QCD. The discrete space-time lattice that is introduced in this formulation of QCD is a regularization scheme particularly well suited for numerical calculations.

Lattice calculations involve systematic errors which are due to the use of a finite lattice cutoff (lattice spacing $a > 0$) and also arise, for instance, from the use of quark masses which are larger than those in nature. Both sources of systematic errors can, in principle, be eliminated. The required computational effort, however, increases rapidly with decreasing lattice spacing and decreasing quark-mass values. With improved computer resources, these systematic errors have indeed been reduced considerably and they are expected to be reduced further in the coming years.

Transition temperature. The phase transition to the QGP is quantitatively best studied in QCD thermodynamics on the lattice. The order of the transition as well as the value of the critical temperature depend on the number of flavours and the values of the quark masses. The current estimates of the transition temperature obtained in calculations which use different discretization schemes in the fermion sector (improved staggered or Wilson fermions), as well as in the gauge-field sector, agree quite well [18, 19]. Taking into account statistical and systematic errors, which arise from the extrapolation to the chiral limit and the not yet fully analysed discretization errors, one finds $T_c = (175 \pm 15)$ MeV in the chiral limit of 2-flavour QCD and a 20 MeV smaller value for 3-flavour QCD. First studies of QCD with two light-quark flavours and a heavier- (strange-)quark flavour indicate that the transition temperature for the physically realized quark-mass spectrum is close to the 2-flavour value. The influence of a small chemical potential on the transition temperature has also been estimated and it has been found that the effect is small for typical chemical potentials characterizing the freeze-out conditions at RHIC ($\mu_B \simeq 50$ MeV) [9–11]. The influence of a non-zero chemical potential will thus be even less important at LHC energies.

Phase transition or crossover. Although the transition is of second order in the chiral limit of 2-flavour QCD, and of first order for 3-flavour QCD, it is likely to be only a rapid crossover in the case of the physically realized quark-mass spectrum. The crossover, however, takes place in a narrow temperature interval, which makes the transition between the hadronic and plasma phases still well localized. This is reflected in a rapid rise of energy density (ε) in the vicinity of the crossover temperature. The pressure, however, rises more slowly above T_c , which indicates that significant non-perturbative effects are to be expected at least up to temperatures $T \simeq (2-3)T_c$. In fact, the observed deviations from the Stefan–Boltzmann limit of an ideal quark–gluon gas are of the order of 15% even at $T \simeq 5T_c$, which hints at a complex structure of quasi-particle excitations in the plasma phase. The delayed rise of the pressure also has consequences for the velocity of sound, $c_s^2 = dp/d\varepsilon$, in the plasma phase. While the ideal gas value $c_s^2 = 1/3$ is almost reached at $T = 2T_c$, the velocity of sound reduces to $c_s^2 \simeq 0.1$ at T_c [20]. In a hydrodynamic description of the expansion of a hot and dense medium this will lead to a slow-down of the expansion for temperatures in the vicinity of the transition temperature.

Equation of state. Recent results of a calculation of p/T^4 [21] are shown in figure 1.4. The corresponding energy densities for 2- and 3-flavour QCD with light quarks are shown on the right-hand side of this figure. A contribution to ε/T^4 , which is directly proportional to the quark masses and thus vanishes in the massless limit, has been ignored in these curves. Since the strange quarks have a mass $m_s \sim T_c$ they will not contribute to the thermodynamics close to T_c but will do so at higher temperatures. Heavier quarks will not contribute in the temperature range accessible in present and near-future heavy-ion experiments. Bulk thermodynamic observables of QCD with a realistic quark-mass spectrum thus will essentially be given by massless 2-flavour QCD close to T_c and will rapidly switch over to the thermodynamics of massless 3-flavour QCD in the plasma phase. This is indicated by the crosses appearing in the right part of figure 1.4. Details of this picture will certainly change when calculations with smaller quark masses and realistic strange-quark masses are performed, in the future, closer to the continuum limit ($a \rightarrow 0$). The basic features of figure 1.4, however, are quite insensitive to changes in the quark mass, have been verified in different lattice fermion formulations, and reproduce well-established results already found in the heavy-quark mass limit [21].




Figure 1.4. The pressure (a) and energy density (b), in QCD with 0, 2 and 3 degenerate quark flavours as well as with two light and a heavier (strange) quark. The $n_f \neq 0$ calculations were performed on a $N_\tau = 4$ lattice using improved gauge and staggered fermion actions. In the case of the SU(3) pure-gauge theory the continuum extrapolated result is shown. The arrows on the right-side ordinates show the value of the Stefan–Boltzmann limit for an ideal quark–gluon gas.

Critical energy density. As can be seen in figure 1.4, the energy density reaches about half the infinite temperature ideal gas value at T_c . The current estimate for the energy density at the transition temperature thus is $\epsilon_c \simeq (6 \pm 2)T_c^4$. Although the transition is only a crossover, this happens in a narrow temperature interval and the energy density changes by $\Delta\epsilon/T_c^4 \simeq 8$ in a temperature interval of only about 20 MeV. This is sometimes interpreted as the latent heat of the transition. When converting these numbers into physical units it becomes obvious that the largest uncertainty on the value of ϵ_c or similarly on $\Delta\epsilon$ still arises from the 10% uncertainty on T_c . Based on T_c for 2-flavour QCD one estimates $\epsilon_c \simeq 0.3\text{--}1.3 \text{ GeV fm}^{-3}$. In the next few years it is expected that the error on T_c can be reduced by a factor 2–3, which will improve the estimate of ϵ_c considerably.

Critical fluctuations. It is evident from the temperature dependence of the energy density and pressure that rapid changes occur in a narrow temperature interval, even when the transition to the plasma phase is only a crossover phenomenon. This leads to large correlation lengths or a rapid rise in susceptibilities, which in turn become visible as large statistical fluctuations. These might be detectable experimentally through the event-by-event analysis of fluctuations in particle yields. Of particular interest are studies of the baryon-number susceptibility [22, 23] which can be linked to charge fluctuations in heavy-ion collisions. In fact, the functional dependence of the baryon-number susceptibility closely follows the energy density shown in figure 1.4. A further enhancement of fluctuations is expected to occur in the vicinity of the tricritical point which is expected to exist in the μ_B – T plane of the QCD phase diagram as an endpoint of the line of first-order transitions at large baryon chemical potential. Although the location of this endpoint is not yet well known [9], it is expected that lattice calculations will allow one to estimate its location and can provide an analysis of fluctuations in its neighbourhood during the coming few years.

Screened heavy-quark potential and heavy-quark bound states. Although chiral-symmetry breaking and confinement are related to different non-perturbative aspects of the QCD vacuum, both phenomena are closely connected at T_c . The transition from the low-temperature hadronic phase to the QGP phase is characterized by sudden changes in the non-perturbative vacuum structure, like a sudden decrease of the chiral condensate, as well as a drastic change of the




Figure 1.5. (a) The heavy-quark free energy in units of the square root of the string tension σ versus the quark–antiquark separation. The different points correspond to calculations at T/T_c values used in (b). The band shows the normalized Cornell potential $V(r) = -\alpha/r + r\sigma$ with $\alpha = 0.25 \pm 0.05$. (b) The estimate of the dissociation energy defined as $\Delta V \equiv \lim_{r \rightarrow \infty} V(r) - V(0.5/\sqrt{\sigma})$.

heavy-quark free energy (heavy-quark potential). This characterizes the QCD transition as chiral-symmetry restoring as well as deconfining. This in turn has consequences for the in-medium properties of both light- and heavy-quark bound states. The heavy-quark free energy, which is commonly referred to as the heavy-quark potential, starts to show an appreciable temperature dependence for $T > 0.6T_c$. With increasing temperature it becomes easier to separate heavy quarks to infinite distance. Already at $T \simeq 0.9T_c$, the free-energy difference for a heavy-quark pair separated by a distance similar to the J/ψ radius ($r_\psi \sim 0.2$ fm) and a pair separated to infinity is only 500 MeV, which is compatible with the average thermal energy of a gluon ($\sim 3T_c$). The $c\bar{c}$ bound states are thus expected to dissolve close to T_c . Some results for the heavy-quark free energy and the estimated dissociation energies are shown in figure 1.5 [18]. These results have been used in phenomenological discussions of suppression mechanisms for heavy-quark bound states. A more direct, unambiguous analysis, however, should be possible on the lattice through the calculation of spectral functions [24]. These techniques are currently being developed and did provide first results in the light- and heavy-quark sector of QCD. These suggest that the J/ψ bound states persist as almost temperature-independent resonance up to $T \simeq 1.5T_c$ while the excited χ_c -states dissolve already at T_c . For larger temperatures strong modifications of the spectral functions in the J/ψ channel are observed. Understanding the detailed pattern of dissolution of the bound state, however, requires further studies.

Chiral-symmetry restoration and modification of in-medium hadron masses. Chiral-symmetry restoration affects the light-meson spectrum. This is clearly seen in lattice calculations of hadron correlation functions in different quantum-number channels [25]. Hadronic correlation functions constructed from light-quark flavours show a drastically different behaviour below and above T_c . This suggests the disappearance of bound states in these quantum-number channels above T_c . Moreover, it provides evidence for the restoration of chiral symmetry at T_c , and a strong reduction of the explicit breaking of the axial $U_A(1)$ symmetry [26, 27]. The latter is effectively restored at least for $T > 1.5T_c$. These results, which previously have been deduced from the behaviour of hadronic screening masses and susceptibilities, are now being confirmed by direct studies of spectral functions [28]. Spectral functions give direct access to the temperature dependence of pole masses, the hadronic widths of these states, and thermal modifications of the continuum. Presumably, a detailed analysis of the vector-meson spectral functions will be most relevant for experimental studies of

in-medium properties of hadrons. At present such studies are performed within the context of quenched QCD [29]. They show that above T_c no vector-meson bound state (e.g. ρ) exists. At present there is, however, no indication for a significant shift of the pole mass below T_c . This may change when similar calculations are performed with light dynamical quarks. The present studies of the vector-meson spectral function also indicate an enhancement over the leading-order perturbative spectral function at intermediate energies and a suppression at lower energies [29]. This has direct consequences on the experimentally accessible thermal dilepton rates.

1.2.3. Perturbative finite-temperature field theory. Lattice calculations show that for temperatures beyond a few hundred MeV, QCD enters a deconfined phase in which, unlike in ordinary hadronic matter, the fundamental fields of QCD—the quarks and gluons—are the dominant degrees of freedom, and the fundamental symmetries are explicit. Although not so far supported by lattice calculations, some form of deconfinement is also expected at several times ordinary nuclear-matter density, as indicated in figures 1.2 and 1.3.

The motivation for considering QGP as a weakly coupled system is, of course, asymptotic freedom, which suggests that the effective coupling g used in thermodynamical calculations should be small if either the temperature T or the baryon chemical potential μ_B is high enough. For instance, if the temperature is the largest scale then $\alpha_s(\bar{\mu}) \equiv g^2/4\pi \propto 1/\ln(\bar{\mu}/\Lambda_{\text{QCD}})$ with typically $\bar{\mu} \simeq 2\pi T$. For $T > 300$ MeV, i.e. well above the critical temperature T_c , the coupling is indeed reasonably small, $\alpha_s < 0.3$.

Extensive weak-coupling calculations have been performed (see [30] for recent reviews and references). These calculations complement the lattice results. Lattice calculations in four dimensions cannot be extended to very high temperatures, currently not beyond $5T_c$, while weak coupling methods are expected to work better at higher temperatures. Moreover, beyond what is currently accessible via lattice calculations, weak-coupling methods allow dynamical quantities and non-equilibrium evolution to be studied.

The picture of the QGP which emerges from these weak-coupling calculations is in many respects very similar to that of an ordinary electromagnetic plasma in the ultra-relativistic regime with, however, specific effects related to the non-Abelian gauge symmetry. To zeroth order in an expansion in powers of g , the QGP is a gas of non-interacting quarks and gluons. In the presence of interactions, plasma constituents with momenta of the order of the temperature ($k \sim T$) become massive quasi-particles, with ‘thermal masses’ of order gT . At small momenta ($k \ll T$) interactions generate collective modes which can be described in terms of classical fields. A hierarchy of scales and degrees of freedom emerges that allows construction of effective theories appropriate to these various degrees of freedom. Once the effective theories are known, they can also be used to describe non-perturbative phenomena.

Scales and degrees of freedom. At high temperature, degrees of freedom with typical momenta of order T are called the *hard* plasma particles. Other important degrees of freedom, collective excitations with typical momenta $\lesssim gT$, are *soft*. Interactions affect hard and soft excitations differently. Gauge fields couple to the excitations through covariant derivatives, $D_x = \partial_x + igA(x)$, so the effect of the interactions depends on the momentum of the excitation and the magnitude of the gauge field. In thermal equilibrium, the strength of the gauge fields is determined by the magnitude of their thermal fluctuations $\bar{A} \equiv \sqrt{\langle A^2(t, \mathbf{x}) \rangle}$ with $\langle A^2 \rangle \approx \int (d^3k/(2\pi)^3) (N_k/E_k)$, where $N_k = 1/(e^{E_k/T} - 1)$ is the Bose–Einstein occupation number. The plasma particles are themselves thermal fluctuations with energy $E_k \sim k \sim T$, so that $\langle A^2 \rangle_T \sim T^2$. The associated electric (or magnetic) field fluctuations $\langle (\partial A)^2 \rangle_T \sim T^4$ are the dominant contributions to the plasma energy density. The main effect of these dominant

thermal fluctuations on hard particles can be calculated perturbatively since $g\bar{A}_T \sim gT \ll k$, generating thermal masses of order gT for both quarks and gluons.

Soft excitations, on the other hand, have momentum $k \sim gT$ and thus are non-perturbatively affected by the hard thermal fluctuations. In this case, the kinetic term $\partial_x \sim gT$ and $g\bar{A}_T$ become comparable. Indeed, the scale gT is that at which collective phenomena develop. The emergence of the Debye screening length, $m_D^{-1} \sim (gT)^{-1}$, which sets the range of the (chromo)electric interactions, is one of the simplest examples of such a phenomenon.

Thermal fluctuations also develop at the soft scale $gT \ll T$ and for these fluctuations the occupation numbers N_k are large, $N_k \sim T/E_k \sim 1/g$, allowing their description in terms of classical fields. Moreover, the large occupation number compensates partially for the reduction in the phase space and implies that $g\bar{A}_{gT} \sim g^{3/2}T$ is still of higher order than the kinetic term $\partial_x \sim gT$. This allows one to calculate the self-interactions of soft modes with $k \sim gT$ in an expansion in powers of \sqrt{g} , leading, for example, to a characteristic g^3 contribution to the pressure as discussed below.

At the *ultra-soft* scale $k \sim g^2T$ the unscreened magnetic fluctuations play a dominant role so that $g\bar{A}_{g^2T} \sim g^2T$ is now of the same order as the ultra-soft kinetic term $\partial_x \sim g^2T$. These fluctuations are no longer perturbative. This is the origin of the breakdown of perturbation theory in high-temperature QCD. A more detailed analysis reveals that the fluctuations at scale g^2T come from the zero Matsubara frequency and correspond therefore to those of a three-dimensional theory of static fields.

We thus observe that the perturbative expansion parameter in high-temperature QCD is in fact not $\mathcal{O}(\alpha_s)$ as naively expected but $\mathcal{O}(\sqrt{\alpha_s})$ or $\mathcal{O}(1)$, depending on the sensitivity of the observable in question to soft physics. Let us turn to some specific examples.

Perturbative thermodynamic calculations. The weak coupling expansion of the free energy is presently known [31] to order $\alpha_s^{5/2}$, or g^5 . However, in spite of the high order, the series show a deceptively poor convergence except for coupling constants as low as $\alpha_s < 0.05$ corresponding to temperatures as high as $> 10^5 T_c$. The convergence problem is already obvious for the lowest orders which read, for a pure SU(3)

$$\frac{P}{P_0} = 1 - \frac{15}{4} \left(\frac{\alpha_s}{\pi} \right) + 30 \left(\frac{\alpha_s}{\pi} \right)^{3/2} + \mathcal{O}(\alpha_s^2), \quad P_0 = 8 \frac{\pi^2 T^4}{45}, \quad (1.1)$$

where P_0 is the ideal-gas pressure. The large coefficient of the g^3 term makes its contribution to the pressure comparable to that of the order g^2 when $g \sim 0.8$, or $\alpha_s \sim 0.05$. Larger couplings make the pressure larger than that of the ideal gas, P_0 , in obvious contradiction with the lattice results [32] shown in figure 1.6(b).

There has been significant recent activity to try to improve on this convergence by developing refined resummation schemes. This goes under the name of Hard Thermal Loop (HTL) resummations. The idea is to borrow the well-developed effective theory for the soft ($k \sim gT$) collective excitations [33–35] and try to apply it also for the hard particles, giving the dominant contribution to the thermodynamics. This is motivated by the attractive physical picture involved in the HTL theory, suggesting a description of the plasma in terms of weakly interacting hard and soft *quasi-particles*. The former are massive excitations with masses $m_D \sim gT$, while the latter are either on-shell collective excitations or virtual quanta exchanged in the interactions between the hard particles. Unfortunately, this framework cannot be made fully systematic theoretically when applied to thermodynamics. Thus different approaches within the same general framework lead to differing predictions [36, 37].

Figure 1.6(a) displays the results obtained in [37] for the entropy of pure-gauge SU(3) theory with different approximations, which describe the contributions of quasi-particles to




Figure 1.6. Entropy (a) and pressure (b) relative to their ideal gas values for a pure-gauge SU(3) Yang–Mills theory. The curves in (a) give the range of uncertainties of the NLA and HTL calculations. The curves in (b) present results obtained in HTL approximation, the approximately self-consistent NLA [37], the three-dimensional reduction [39] and the four-dimensional lattice data [32].

different accuracy. In contrast to the HTL approximation, the next-to-leading approximation (NLA) also accounts fully for the higher order g^3 perturbative contributions. Most importantly, both approximations show good agreement with lattice data for $T > 2.5T_c$. Improving the approximation to NLA leads only to a moderate change, i.e. this scheme results in a stable well-controlled approximation, in contrast to ordinary perturbation theory. More generally, this illustrates the complementarity of lattice and resummed perturbative calculations. The latter are controlled in the regime above $5T_c$, not accessible by lattice calculations. The extension of this method to a plasma including quarks (in particular, to non-zero baryon density) presents no difficulty. Indeed, quark-number susceptibilities have recently been computed [37] and compared with lattice results [23].

At the same time, there are other important observables, such as colour-electric and colour-magnetic screening lengths, which are not well described by this approach due to their stronger sensitivity to non-perturbative collective phenomena. Thus, it is a fact that we still lack a reliable and fully comprehensive analytic description of the QGP phase.

Dimensional reduction. There is also a non-perturbative first-principles method to compute the contribution of the soft fields to the thermodynamics. This uses a combination of perturbation theory (for the hard modes) and lattice calculations (for the soft ones), known as dimensional reduction (see [38, 39] and references therein). First, an effective three-dimensional theory for the static Matsubara modes is constructed which is then studied numerically on a three-dimensional lattice. In figure 1.6(b), the pressure obtained in this way [39] is compared to the results of the analytic resummation in [37] and to the four-dimensional lattice results; the latter stop at $5T_c$. The results of [39] depend upon an undetermined parameter e_0 because of the (so far) incomplete matching between the four- and three-dimensional theories. The fairly good agreement with the best estimates of [39] ($e_0 = 10$) supports the assumption that the strong interactions of the soft modes tend to cancel in thermodynamical observables, such as the pressure. At the same time, non-perturbative strong interactions manifest themselves very clearly in all screening lengths—but these can still be evaluated within the simple dimensionally reduced effective theory, rather than the full QCD [40].

Non-static and non-equilibrium observables and kinetic theory. From the physical point of view, the most prominent non-static observables studied with perturbation theory in the QGP are photon and dilepton production rates, as well as various transport coefficients. For a

discussion on photon rates, see section 1.3.6 and for the current status and references on transport coefficients, see [41]. In addition to these specific observables, there has also been an attempt to make a general theoretical framework which would address non-equilibrium properties in the QGP in order to assess whether, for instance, such a plasma can thermalize in the timescale available in ultra-relativistic heavy-ion collisions. Such a framework could be provided by a kinetic theory for the QGP but so far progress is limited to situations close to equilibrium [30]. In that limit it is known that collective phenomena at the scale gT contained in the HTL theory are reproduced by the kinetic theory which can also account for the collective phenomena at the ultra-soft scale g^2T , once collision terms are included [30, 42].

1.2.4. Classical QCD of large colour fields. The bulk of multiparticle production at central rapidities in A–A collisions at the LHC will be from partons in the nuclear wavefunctions with momentum fractions $x \leq 10^{-3}$. Partons with these very small x values interact coherently over distances much larger than the radius of the nucleus. In the centre-of-mass frame, a parton from one of the nuclei will coherently interact with several partons in the other nucleus. From these considerations alone, it is unlikely that the perturbative QCD collinear factorization approach resulting in a convolution of probabilities will suffice for describing multiparticle production at the LHC. Understanding the initial conditions for heavy-ion collisions therefore requires understanding scattering at the level of amplitudes rather than probabilities. In other words, we must understand the properties of the nuclear wavefunction at small x .

One may anticipate that, like most problems involving quantum mechanical coherence, this is a very difficult problem to solve. There are, however, several remarkable features of small- x wavefunctions that simplify the problem. First, perturbative QCD predicts that parton distributions grow very rapidly at small x . At sufficiently small x , the parton distributions saturate; namely, when the density of partons is such that they overlap in the transverse plane, their repulsive interactions are sufficiently strong to limit further growth to be at most logarithmic [43]. The occupation number of gluons is large and is proportional to $1/\alpha_s$ [44] while the occupation number of sea quarks remains small even at small x . Second, the large parton density per unit transverse area provides a scale, the saturation scale $Q_s(x)$, which grows like $Q_s^2 \propto A^{1/3}/x^\delta$ for $A \rightarrow \infty$ at $x \rightarrow 0$ and where $\delta \sim 0.3$ at HERA energies [45]. Large nuclei thus provide an amplifying factor—one needs to go to much smaller x in a nucleon to have a saturation scale comparable to that in a nucleus. One may estimate that $Q_s^2 \sim 2\text{--}3 \text{ GeV}^2$ at the LHC [46–48]. The physics of small- x parton distributions can be formulated in an effective field theory (EFT), where the saturation scale Q_s appears as the only scale in the problem [44]. The coupling must therefore run as a function of this scale, and since $Q_s^2 \gg \Lambda_{\text{QCD}}^2$ at the LHC, $\alpha_s \ll 1$. Note that the small coupling also ensures that the occupation number is large.

Thus even though the physics of small- x partons is non-perturbative, it can be studied with weak coupling [49–51]. This is analogous to many systems in condensed-matter physics. In particular, the physics of small- x wave functions is similar to that of a spin glass [50, 52]. Further, in this state the mean transverse momentum of the partons is of order Q_s , and their occupation numbers are large, of order $1/\alpha_s$ as in a condensate. Hence, partons in the nuclear wavefunction form a colour glass condensate (CGC) [51, 52]. The CGC displays remarkable Wilsonian renormalization group properties which are currently a hot topic of research [50, 51, 53].

Returning to nuclear collisions, the initial conditions correspond to the melting of the CGC in the collision. The classical effective theory for the scattering of two CGCs can be formulated [54] and classical multiparticle production computed. No analytical solutions of this problem exist (for a recent attempt, see [55]) but the classical problem can be solved numerically [56].

Since Q_s (and the nuclear size R) are the only dimensional scales in the problem, all properties of the initial state can be expressed in terms of Q_s (and functions of the dimensionless product $Q_s R$) [57]. The classical simulations of the melting of the CGC can only describe the very early instants of a nuclear collision with times of the order $1/Q_s$. At times of order $1/\alpha_s^{3/2} Q_s$, the occupation number of the produced partons becomes small and the classical results are unreliable [48]. The subsequent evolution of the system can then be described by transport methods, with initial conditions given by the classical simulations [58]. Although there is a very interesting recent theoretical estimate of thermalization in this scenario [48], the problem of whether such a system thermalizes remains.

The colour-glass picture of nuclear collisions has been applied to study particle production at RHIC energies. It is claimed that variation of distributions with centrality, rapidity, and energy can be understood in this picture [59, 60]. The variation of p_t -distributions with centrality also appear to be consistent with the CGC [61]. These studies do not assume thermalization but assume parton-hadron duality in computing distributions. The CGC approach is a quickly evolving field of research, which at the moment relies on the assumption that dynamical rescattering during the expansion does not affect predictions.

1.3. Heavy-ion observables in ALICE

The aim of this section is to give a theoretical overview of the observables that can be measured in ALICE. Whenever results from pp and pA collisions are needed for the interpretation of A–A data, they are discussed in this section, too. Other aspects of pp and pA physics accessible in ALICE are addressed in sections 1.4 and 1.5.

1.3.1. Particle multiplicities. The average charged-particle multiplicity per rapidity unit (rapidity density dN_{ch}/dy) is the most fundamental ‘day-one’ observable. On the theoretical side, it fixes a global property of the medium produced in the collision. Since it is related to the attained energy density, it enters the calculation of most other observables. Another important ‘day-one’ observable is the total transverse energy, per rapidity unit at mid-rapidity. It determines how much of the total initial longitudinal energy is converted to transverse debris of QCD matter. On the experimental side, the particle multiplicity fixes the main unknown in the detector performance; the charged-particle multiplicity per unit rapidity largely determines the accuracy with which many observables can be measured.

Despite their fundamental theoretical and experimental importance, there is no first-principles calculation of these observables starting from the QCD Lagrangian. Both observables are dominated by soft non-perturbative QCD and the relevant processes must be modelled using the new, large-scale, $R_A \approx A^{1/3}$ fm.

These difficulties are reflected in the recent theoretical discussion of the expected particle multiplicity in heavy-ion collisions. Before the start-up of RHIC, extrapolations from SPS measurements at $\sqrt{s} = 20$ to 200 GeV varied widely [62], mostly overestimating the result. Even after the RHIC data, extrapolations over more than an order of magnitude in \sqrt{s} , from 200 to 5500 GeV, are still difficult. Based on this experience and the first RHIC data, this section summarizes the current theoretical expectations of the particle multiplicity and total transverse energy.

Multiplicity in proton–proton collisions. Understanding the multiplicity in pp collisions is a prerequisite for the study of multiplicity in A–A. The inclusive hadron rapidity density in $pp \rightarrow hX$ is defined by

$$\rho_h(y) = \frac{1}{\sigma_{\text{in}}^{\text{pp}}(s)} \int_0^{p_t^{\text{max}}} d^2 p_t \frac{d\sigma^{\text{pp} \rightarrow hX}}{dy d^2 p_t}, \quad (1.2)$$




Figure 1.7. Charged-particle rapidity density per participant pair as a function of centre-of-mass energy for A–A and pp collisions. Experimental data for pp collisions are overlaid by the solid line, equation (1.3) [equations (1.4) and (1.5) practically coincide in this range]. The dashed line is a fit $0.68 \ln(\sqrt{s}/0.68)$ to all the nuclear data. The dotted curve is $0.7+0.028 \ln^2 s$. It provides a good fit to data below and including RHIC, and predicts $N_{\text{ch}} = 9 \times 170 = 1500$ at LHC. The long dashed line is an extrapolation to LHC energies using the saturation model with the definition of Q_s as in equation (1.12) [64].

where $\sigma_{\text{in}}^{\text{pp}}(s)$ is the pp inelastic cross section. The energy dependence of $\sigma_{\text{in}}^{\text{pp}}(s)$ which starts growing above $\sqrt{s} = 20$ GeV is poorly known [63]. At high energy the dependence can be parametrized by a power $\sqrt{s}^{0.14}$ or either $\sim \ln s$ or $\sim \ln^2 s$. The hadron rapidity density $\rho_{\text{h}}(y)$ also grows. It can be parametrized for charged particles by expressions like (\sqrt{s} in GeV):

$$\rho_{\text{ch}}(y=0) \approx 0.049 \ln^2 \sqrt{s} + 0.046 \ln \sqrt{s} + 0.96, \quad (1.3)$$

$$\approx 0.092 \ln^2 \sqrt{s} - 0.50 \ln \sqrt{s} + 2.5, \quad (1.4)$$

$$\approx 0.6 \ln(\sqrt{s}/1.88). \quad (1.5)$$

Thus at SPS energies, $\sqrt{s} = 20$ GeV, the charged-particle rapidity density, $\rho_{\text{ch}}(y)$, at $y \approx 0$ is about 2, at RHIC energies about 2.5 and extrapolates to about 5.0 at ALICE energies (see figure 1.7). We reiterate that these numbers cannot be derived from the QCD Lagrangian.

Multiplicity in A–A collisions. It is possible to estimate the multiplicity in A–A collisions at the LHC using dimensional arguments: present theoretical estimates of the multiplicity assume the existence of a dynamically determined *saturation scale* Q_s , which is basically the transverse density of all particles produced within one unit of rapidity

$$\frac{N}{R_A^2} = Q_s^2, \quad (1.6)$$

where $R_A = A^{1/3}$ fm and all constants, factors of α_s , etc are assumed to be 1. Then

$$N = Q_s^2 R_A^2 = (2 \text{ GeV} (200)^{1/3} \text{ fm})^2 = 3500 \quad (1.7)$$

if $Q_s = 2 \text{ GeV}$ at the LHC.

However, this value for Q_s is undetermined. Microscopic information is needed for a more reliable estimate. Let us assume that all particles are produced by hard sub-collisions. Then, in a central A–A collision, at scale Q_s ,

$$N = \frac{A^2}{R_A^2} \frac{1}{Q_s^2} = Q_s^2 R_A^2, \quad (1.8)$$

where the factor A^2/R_A^2 is the nuclear overlap function (T_{A-A} at zero impact parameter) and the subprocess cross section is $\approx 1/Q_s^2$. Then by taking the square root, we find, using $R_A = A^{1/3}$ fm ($0.2 \text{ fm GeV} = 1$),

$$Q_s^2 R_A^2 = A, \quad \rightarrow \quad Q_s = 0.2 A^{1/6} \text{ GeV}. \quad (1.9)$$

Putting this result back into expression (1.8), one has

$$N = A = Q_s^2 R_A^2. \quad (1.10)$$

This equation could also be taken as the starting point of this dimensional exercise, written as

$$Axg(x, Q_s^2) = Q_s^2 R_A^2. \quad (1.11)$$

Note that $N \sim A$, not $\sim A^{4/3}$, as expected for independent nucleon–nucleon collisions at a fixed, A-independent, scale.

However, for something quantitatively useful, it is necessary to include more about the magnitude of Q_s and its energy dependence. For example, instead of expression (1.9), a quantitatively more accurate expression could be (in GeV units)

$$Q_s = 0.2 A^{1/6} \sqrt{s}^b, \quad b \approx 0.2 \quad (1.12)$$

which feeds directly into the energy dependence of N , a necessity for LHC predictions.

Theoretical models of the LHC multiplicity basically aim to determine the constants missing from this dimensional argument.

Estimates of multiplicity. The aim of studying heavy-ion collisions is to discover qualitatively new effects, ascribed to the new scale, $\sim A^{1/3}$ fm, and not observed in pp collisions. It is thus appropriate also to compare the A–A multiplicity with that of pp collisions. It has appeared that the most illuminating way is to compute the number of participants, N_{part} , in the collision. This is a phenomenological quantity (see, for example, [59], which for pp is 2, for central A–A collisions, about $2A$, and which also can be estimated as a function of impact parameter b . A plot of $N_{\text{ch}}(0.5N_{\text{part}})$, where $N_{\text{ch}} \equiv \rho_{\text{ch}}(y=0)$, then gives a quantitative measure of how efficient A–A collisions are for each sub-collision. RHIC data are plotted in this way in figure 1.7.

The outcome thus is that while pp collisions at $\sqrt{s} = 200 \text{ GeV}$ produce 2.5 charged particles per collision, A–A ($A_{\text{eff}} = 170$, taking into account centrality cuts) collisions produce 3.8 particles. This is a 50% increase relative to pp, quite a sizeable effect.

One example [64] of predictions for N_{ch}^{A-A} at larger s , up to $\sqrt{s} = 5.5 \text{ TeV}$, is also shown in figure 1.7. This particular prediction involves a power-like behaviour, $\sim \sqrt{s}^{0.38}$, and implies that while pp collisions produce 5 particles, A–A collisions produce 13 particles (always per




Figure 1.8. Data and predictions of the saturation model for charged-particle multiplicity per unit pseudo-rapidity near $\eta = 0$ as a function of centre-of-mass energy of the nucleon–nucleon system. Note that experimental points measured at the same energy have been slightly displaced for visibility.

unit y at $y = 0$). At this higher energy the increase thus is predicted to be 160%. In absolute numbers an approximate prediction is $\approx 13A_{\text{eff}} = 13 \times 170 \approx 2200$ —although this number is still affected by a set of corrections (see below).

The crucial issue here is whether the behaviour is power-like with a reasonably large exponent or only $\sim \ln \sqrt{s}$ or $\sim \ln^2 \sqrt{s}$ like in pp. The range within RHIC is small and that from RHIC to LHC is large (see figure 1.7) so there is a lot of room for error. Within RHIC one cannot distinguish between even a large power-like $\sqrt{s}^{0.38}$ and a single $\ln s$ dependence.

A very simple extrapolation one may attempt is to fit $N_{\text{ch}}(\sqrt{s}) \sim \ln s$ to $N_{\text{ch}}(130) = 555$, $N_{\text{ch}}(200) = 650$: $N_{\text{ch}} = -518 + 220.5 \ln \sqrt{s}$. This would give $N_{\text{ch,LHC}} = 1381$.

There are several different factors which have to be taken into account when comparing theoretical to experimentally observed numbers:

- Experiments measure the charged multiplicity, theory usually gives the total one which has to be reduced by a factor somewhat less than $2/3$ to get the charged one.
- Experiments may measure only the pseudo-rapidity distribution, theoretically the rapidity distribution is obtained. A precise relation can be formulated only when the double-differential y, p_t distributions and mass spectrum are known.
- A ‘central collision’ in practice is defined only up to some small non-centrality and $A_{\text{eff}} < A$.
- There may be many model-dependent uncertainties: theory may give multiplicity at formation (at time of about $0.1 \text{ fm } c^{-1}$) and this is modified during the evolution of the system (up to a time of more than $10 \text{ fm } c^{-1}$).

For the model of [64] these correction effects are taken into account [65] and lead to the numbers for the pseudo-rapidity density shown in figure 1.8.

In [66] $dN_{\text{ch}}/d\eta$ is computed in a two-component soft and semihard string model; here the density at $\eta = 0$ is between 2600 and 3200, depending on assumptions, somewhat larger than in [65].

We repeat that there is no way of giving a quantitative estimate of the accuracy of these numbers (nor of those from any of alternative models). Even if the model in [64] predicted the RHIC numbers well, it is still a model and its uncertainties are unknown. The first LHC events will give an answer.

1.3.2. Particle spectra. The bulk of the particles emitted in a heavy-ion collision are soft hadrons with transverse kinetic energies $m_t - m_0 < 2 \text{ GeV}$, which decouple from the collision region in the late hadronic freeze-out stage of the evolution. The main motivation for studying these observables is that parameters characterizing the freeze-out distributions constrain the dynamical evolution and thus yield indirect information about the early stages of the collision [14, 67–72]. They provide information on the freeze-out temperature and chemical potential, radial-flow velocity, directed and elliptic flow coefficients, size parameters as extracted from two-particle correlations, event-by-event fluctuations and correlations of hadron momenta and yields. Furthermore, there are theoretical arguments showing that different aspects of the measured final hadronic momentum distributions are determined at different times [69–76]. This suggests that the final momentum distributions (shapes, normalizations, and correlations) provide detailed information about the time evolution of the collision fireball [72]. The current understanding is that the chemical composition of the observed hadronic system is already fixed at hadronization, i.e. at the point where hadrons first appear from the initial hot and dense partonic system [14, 69, 70]. Certain fluctuations from event to event in the chemical composition may point back to even earlier times [77, 78]. At LHC energies, the elliptic flow coefficient [79–90], which describes the anisotropy of the transverse-momentum spectra relative to the reaction plane in non-central heavy-ion collisions, is expected to saturate well before hadronization, during the QGP stage [86–88], providing information about its equation of state [79, 83, 85–87, 90]. In the following, we discuss this picture in more detail.

Chemical and kinetic freeze-out. As first observed at the AGS and SPS [71, 72, 91] and now confirmed at RHIC [89, 92], the shapes and the normalizations of the hadron momentum spectra reflect two different, late, stages of the collision. The total yields, reflecting the particle abundances and thus the chemical composition of the exploding fireball, are frozen almost directly at hadronization and are, at most, very weakly affected by hadronic rescattering. The spectral shapes, on the other hand, reflect a much lower temperature in combination with strong collective transverse flow, at least some fraction of which is generated by strong quasi-elastic rescattering among the hadrons. The rescattering is facilitated by the existence of strongly scattering resonances such as the ρ , K^* , Δ and other short-lived baryon resonances [93]. The role of these resonances in the late-stage rescattering can be assessed through direct measurements. The majority of all particles show a distribution consistent with formation at T_c [16] although one might expect exceptions for wide resonances [94]. In general, what is measured is not the original resonance yield at hadronization but the lower yield of resonances just before kinetic freeze-out. On the other hand, when reconstructing resonances from decay products that do *not* rescatter, such as e^+e^- or $\mu^+\mu^-$ pairs, one can probe their decays during the early part of the hadronic rescattering stage. In the case of the short-lived ρ mesons, a large fraction of which decay before kinetic hadronic rescattering has ceased, the invariant-mass distribution was found to be significantly broadened by collisions [95, 96]. It will be interesting to contrast hadronic and leptonic decays of resonances such as $\phi \rightarrow K^+K^-$ and $\phi \rightarrow l^+l^-$ at the same centrality to better understand the hadronic rescattering stage.

Temperature and collective flow. The separation of thermal motion and collective flow in the analysis of soft-hadron spectra requires a thermal model analysis based on a hydrodynamic approach or on simple hydrodynamically motivated parametrizations. In such an analysis one assumes longitudinal boost invariance and the measured transverse-mass spectra are compared to either the following form or simplified approximations thereof [97, 98]:

$$\frac{dN_i}{dy m_t dm_t d\varphi_p} = \frac{2g_i}{(2\pi)^3} \sum_{n=1}^{\infty} (\mp 1)^{n+1} \int r dr d\varphi_s \tau_f \exp [n(\mu_i + \gamma_{\perp} \vec{v}_{\perp} \cdot \vec{p}_t)/T] \times \left[m_t K_1(n\beta_{\perp}) - (\vec{p}_t \cdot \vec{\nabla}_{\perp} \tau_f) K_0(n\beta_{\perp}) \right]. \quad (1.13)$$

Here g_i is the spin–isospin degeneracy factor for particle species i , $\mu_i(\vec{r})$ is its chemical potential as a function of the transverse distance \vec{r} from the collision axis, $T(\vec{r})$ and $\tau_f(\vec{r})$ are the common temperature and longitudinal proper time at freeze-out, $\vec{v}_{\perp}(\vec{r})$ is the transverse fluid velocity with $\gamma_{\perp} = (1 - v_{\perp}^2)^{-1/2}$, and $\beta_{\perp} \equiv \gamma_{\perp} m_t / T$. The sum arises from the expanding thermal distribution, either Bose–Einstein (–) or Fermi–Dirac (+). Except for pions, the Boltzmann approximation to the thermal distribution is adequate and only the first term in the expansion is used. Transverse mass m_t and momentum \vec{p}_t are related by the particle mass m_i as $m_t = (m_i^2 + p_t^2)^{1/2}$. The dependencies on the transverse position \vec{r} (in particular the angle φ_s relative to the reaction plane) describe arbitrary transverse density and velocity profiles which become important for studying momentum anisotropies such as elliptic flow. In this general case, the integrals must be carried out numerically. For angle-averaged spectra, however, the angular integral can be calculated analytically. To further simplify the remaining radial integral, constant chemical potentials, temperature and freeze-out proper time are often assumed, along with a simple linear parametrization for the transverse fluid factor $\gamma_{\perp} v_{\perp} \equiv \sinh \eta_{\perp}$ where $\eta_{\perp}(\vec{r})$ is the transverse-flow rapidity.

Strong constraints on the energy density and pressure in the early, pre-hadronic, QGP stage arise from the anisotropies of the generated transverse flow. To exploit these constraints, it is necessary to check if the momentum spectra can be simultaneously described with a common freeze-out temperature and transverse-flow profile according to equation (13). In the Boltzmann approximation, the value of the chemical potential affects only the normalizations, not the shapes of the distribution. Thus, this analysis tests the *kinetic* thermal equilibrium and is insensitive to whether or not the system is in chemical equilibrium. For a single particle species, the temperature and transverse flow extracted from such a fit are strongly correlated: the spectra can be made flatter by increasing either the temperature or the transverse flow. This ambiguity in T and v_{\perp} can be removed by simultaneously fitting spectra for different mass particles and searching for a common overlap region in the T – v_{\perp} plane. If overlap exists, the collective flow picture with common freeze-out works. Alternatively, the T – v_{\perp} ambiguity can be resolved by combining the spectral shapes with information from two-particle correlations (see the following subsection).

With sufficiently accurate identified-hadron spectra, it is in general possible to determine T and v_{\perp} unambiguously because for the same temperature and flow the spectra of particles with different masses have different shapes. The rest mass plays no role for relativistic transverse momenta, thus for $m_t \gtrsim 2m_i$ the inverse slope of all hadrons should be given by the same ‘blue-shifted temperature’ [97, 98]:

$$T_{\text{slope}} = T \sqrt{\frac{1 + \langle v_{\perp} \rangle}{1 - \langle v_{\perp} \rangle}}. \quad (1.14)$$

For non-relativistic transverse momenta, $m_t < 1.5m_i$, the collective flow contributes to the measured momentum a term $\sim m_i \langle v_{\perp} \rangle$ which flattens the spectrum in direct proportion to the

rest mass [97, 99]:

$$T_{\text{slope},i} = T + \frac{1}{2}m_i \langle v_{\perp} \rangle^2. \quad (1.15)$$

When connecting this flatter slope at low m_t with the steeper one from equation (1.14) at high m_t , one obtains a spectrum of ‘shoulder-arm’ shape where the shoulder becomes more and more prominent as the hadron mass increases. To observe this change in slope, an experiment must cover the entire low- p_t range, from the smallest possible p_t to about 3 GeV for the heaviest hadrons. At significantly larger p_t the thermal picture begins to give way to semihard physics and the slopes can no longer be interpreted in a collective flow picture.

Hadron yields and chemical composition. The measured hadron yields or their ratios are controlled by the chemical potentials μ_i and temperature T . Owing to flow effects, the transverse spectra are neither straight exponentials nor simple Bessel functions and can thus not easily be extrapolated to low p_t without introducing systematic uncertainties. A good determination of the yields thus requires good low- p_t coverage. It is then possible to check whether the resulting chemical potentials satisfy the expected relations for a hadronic system in chemical equilibrium subject to certain constraints such as vanishing net strangeness, charge and baryon number conservation, and possibly an under- or over-saturation of total strangeness [100]. The temperature required to satisfy these relations turns out not to be the same as the one required to fit the shapes of the momentum spectra [12, 98, 100, 101]. The particle yields freeze out once inelastic processes cease and because of the smaller cross sections, this happens significantly before resonant quasi-elastic collisions which modify the shape of the momentum spectra cease. The degree to which chemical equilibrium among the hadrons is established provides important constraints on the microscopic chemical reaction processes and their timescales. If it is the case, as already observed at the SPS and RHIC [12, 92, 102], that the chemical freeze-out temperature coincides with the critical temperature predicted from lattice QCD, the observed equilibrium cannot have been dynamically generated via hadronic rescattering due to the short timescale but must have been established by the hadronization process itself: “hadrons were born into chemical equilibrium” [14, 69, 70, 103]. The degree of strangeness saturation and the absence of canonical correlations among hadrons with conserved quantum numbers then give further clues about the nature of the pre-hadronic state from which the hadrons emerged.

Elliptic flow and early pressure. In non-central heavy-ion collisions, the initial overlap zone between the colliding nuclei is spatially deformed. If the matter produced in the reaction zone rescatters efficiently, this spatial anisotropy gets transferred to momentum space and the initial, locally isotropic, momentum distribution develops anisotropies. This will not happen, however, in the absence of rescattering and the transfer of the initial spatial anisotropy to momentum space will be inefficient if rescattering among the produced particles is weak. The momentum anisotropies lead to a dependence of the transverse-momentum distribution on the emission angle relative to the reaction plane and can be quantified by the coefficients of an azimuthal Fourier decomposition of the distribution:

$$\frac{dN(b)}{dy m_t dm_t d\varphi_p} = \frac{1}{2\pi} \frac{dN(b)}{dy m_t dm_t} (1 + 2v_1(p_t, y; b) \cos(\varphi_p) + 2v_2(p_t, y; b) \cos(2\varphi_p) + \dots). \quad (1.16)$$

At mid-rapidity ($y = 0$), the direct flow v_1 vanishes by symmetry. The first non-zero coefficient is the *elliptic-flow coefficient* $v_2(p_t; b)$.

Our present understanding of how anisotropic flow develops strongly suggests that, for a given initial spatial deformation of the reaction zone, the largest momentum anisotropies are obtained in the hydrodynamic limit [89], i.e. the limit of infinitely fast rescattering, leading to

instantaneous local thermal equilibrium distributions. What is important is that the energy–momentum tensor assumes an approximate ideal fluid form with an equilibrium equation of state. The anisotropic pressure gradients then create the flow anisotropies.

As the elliptic flow develops, the early strong spatial deformation decreases since the matter begins to expand more rapidly in the initially short direction than in the long one [81–83]. As the spatial deformation disappears, the build-up of flow anisotropies from pressure gradients ceases and the elliptic flow saturates [83]. This saturation requires a finite time which is controlled by the collision geometry but unrelated to the time needed for the reaction zone to hadronize or freeze-out. If the initial energy density is sufficiently high, elliptic flow saturates before hadronization sets in. RHIC data [104–107] show that in semi-central Au–Au collisions the measured elliptic flow saturates the hydrodynamic upper limit for transverse momenta up to 2 GeV, requiring very early thermalization and pressure build-up [83, 89]. At the time when thermalization must have set in to obtain the observed elliptic flow, the energy densities in the successful hydrodynamic model calculations are about an order of magnitude above the hadronization transition. This suggests that the pressure for the expansion comes from a thermalized QGP [89].

Early pressure diagnostics via anisotropic flow measurements will also be important at the LHC where the initial energy densities are so high that all elliptic flow should have been generated before the reaction zone hadronizes [83]. Since all hadron species should then see the same anisotropic flow, their v_2 coefficients should obey simple hydrodynamic relations [85]. A systematic comparison of v_2 for different types of hadrons, fixing the freeze-out conditions from the azimuthally averaged spectra, may then reveal the stiffness of the equation of state [85], i.e. the ratio of pressure to energy density, during the stage when the anisotropy developed, i.e. in the QGP.

1.3.3. Particle correlations. The unique feature which distinguishes heavy-ion collisions from those of simpler systems is the collective behaviour of matter under extreme conditions and their large space–time extent. The evolution of the system is driven by the properties of the strongly interacting bulk matter and thus it is sensitive to the equation of state. The resulting space–time picture at the time of its breakup gives access to the pressure gradients developed during this evolution. Thus the extraction of the fireball’s final state is essential in order to determine its dynamics. This final-state information is most directly accessed by interferometric methods that measure the size of the fireball, its expansion, and its phase-space density. They also carry information about the timing of hadronization or the sequence of emission of individual species. For the latest reviews see [108–110].

At high multiplicity, the two main effects that lead to measurable correlations are (i) wave function quantum symmetrization (QS) or anti-symmetrization due to Bose–Einstein or Fermi–Dirac statistics respectively [111–113] and (ii) final-state interactions (FSI) between the produced particles, either Coulomb or strong interaction [114–118]. Moreover, global energy–momentum or charge conservation can introduce correlations which, however, are less important at high multiplicities [116]. Conservation laws which correlate a produced charge–anti-charge pair and are accessible via balance functions [119] are of particular interest for heavy-ion physics. Also, more exotic effects may show up in correlation measurements. For example, it has been suggested [120] that a medium-induced change of the pion mass can lead to a novel type of back-to-back correlation in the case of sudden freeze-out.

The measured quantity is the two-particle correlation function defined as

$$C(p_1, p_2) = \frac{d^6 N}{dp_1^3 dp_2^3} \bigg/ \frac{d^3 N}{dp_1^3} \frac{d^3 N}{dp_2^3}. \quad (1.17)$$

In heavy-ion collisions one can neglect kinematic constraints as well as most of the dynamical correlations and construct the reference distribution in the denominator by mixing the particles from different events, normalizing the correlation function to unity at sufficiently large relative velocities. For ALICE it is justified to assume (except for events with large phase-space density fluctuations) that the correlation of two particles emitted with a small relative velocity is influenced by the effects of their mutual QS and FSI only and that the momentum dependence of the one-particle emission probabilities is unimportant when varying the particle four-momenta p_1 and p_2 by the amount characteristic for the correlation due to QS and FSI (smoothness assumption).

The formalism is simplest for spin-0 particles. Since the measurement requires good statistics, pions and kaons are most often used. If one can correct for the effect of FSI [118, 121], quantum statistics leads to the following relation between the correlation function and the source distribution [108, 122]:

$$C(K, q) = 1 + \frac{|\int d^4x S(x, K) \exp(ix \cdot q)|^2}{\int d^4x S(x, K + q/2) \int d^4y S(y, K - q/2)}, \quad (1.18)$$

where $K = (p_1 + p_2)/2$, $q = p_1 - p_2$ and x is the space-time four-vector. Here $S(x, K)$ is the emission function, the Wigner density of the source normalized to the particle number. Classically it can be understood as a probability to emit a particle of momentum K from the position x . The aim of particle interferometry is to learn as much as possible about this emission function.

The emission function $S(x, K)$ cannot be reconstructed uniquely from experimental data on $C(K, q)$. Firstly, since the correlation function depends on the square of Fourier transform of $S(x, K)$, the phase of $S(x, K)$ remains inaccessible—three-particle correlations are needed for its determination [123]. Secondly, the momenta of the detected particles satisfy the on-shell constraint $q \cdot K = 0$. This implies that only three of the four components of q can be determined independently. Usually one adopts the three spatial components in the so-called *out-side-long* coordinate frame: *longitudinal* (z -axis) direction parallel to the beams, *outward* (x -axis) given by the direction of transverse pair momentum, and *sideward* perpendicular to the above two (in the direction of the y -axis).

Quantum-statistical effects are important if the particles are close in phase space, $\Delta x \cdot q \sim \hbar$. Hence, the inverse widths of the correlation function in q give information about the corresponding spatial extensions $\sim 1/q$ of the particle emitting source. These widths determine independent combinations of three-dimensional spatial and temporal information. This makes it important to measure all three independent q components. In the frequently adopted Gaussian parametrization of $C(K, q)$, the widths are parametrized in terms of the radius parameters R_{ij} , ($i, j = 0, s, l$) [124–126], sometimes called HBT (Hanbury-Brown–Twiss) radii:

$$C(K, q) - 1 = \lambda(K) \exp[-q_0^2 R_0^2(K) - q_s^2 R_s^2(K) - q_l^2 R_l^2(K) - 2q_0 q_l R_{0l}^2(K)]. \quad (1.19)$$

This form holds for azimuthally symmetric sources, i.e. for central collisions, while for collisions at finite impact parameter all three cross-terms should be included [127].

A non-Gaussian shape of the correlation function may render the use of equation (1.19) more problematic. In such a case, the shape of the relative distance distribution can be reconstructed model-independently by imaging techniques [128].

Intercept parameter. In practice, the measured correlation function does not reach the ideal value of 2 for spin-0 particles of vanishing relative momentum q . The ideal value is different for other spins, e.g. it is 1/2 for unpolarized protons. In equation (1.19) this is parametrized

by the intercept parameter $\lambda(K)$ which may be less than 1 for several reasons. Experimental effects like particle misidentification and uncertainties in the corrections for FSI may strongly affect the value of $\lambda(K)$. In the theoretical literature, a further reduction of $\lambda(K)$ is mainly attributed to two effects. First, decay products of long-lived resonances are typically emitted far from the collision region. A contribution from such a large source to the correlator will be visible as a very narrow peak of $C(K,q)$, beyond the momentum resolution of the detector [129–133], thus leading to an apparent reduction. Secondly, coherent particle production, e.g. stemming from the decay of a disoriented chiral condensate (DCC) [134], can lead to a reduced value of λ . These two effects could be distinguished by measuring two- and many-particle correlations [109, 123, 135]. A method for extracting the coherent component of particle radiation by analysing two-particle correlation functions of like- and unlike-charged particles was proposed in [136]. As for a method to exploit multiparticle correlations, the simplest case of three-particle correlations may not be enough to do this, but the differences are more pronounced for higher n 's [137]. In section 6 of the PPR Volume II, an evaluation of the highest n reachable by ALICE will be provided.

Space-time information from the radius parameters and the dynamics. The dependence of the correlation radii $R_{ij}(K)$ on the average pair momentum K contains information about the dynamical expansion of the collision region. For strong dynamical expansion as expected at LHC, the radius parameters do not measure the entire source size, but the so-called regions of homogeneity [138]. Generally, the length of homogeneity determines the spatial distance over which the emission probability of a given momentum drops by a factor of two. In a thermal model, a region of homogeneity is the part of the fireball which approximately co-moves with the momentum K and in which the velocity spread is sufficiently small in order for the emitted particles to show Bose–Einstein enhancement. Since for an expanding fireball the velocity is given by the position, this mechanism leads to x – p position–momentum correlations. Such correlations appear not only for collective hydrodynamic expansion but also for particles produced in decays of resonances [130, 133] or colour strings [139, 140]. By measuring the correlation radii as a function of pair momentum K different parts of the source are scanned. The K -dependence of the radius parameters corresponds to the variation of the homogeneity lengths with momentum. Thus dynamics plays a rather crucial role in determining the size of correlation radii. In order to access it, a measurement over a very wide range of transverse momentum K is needed. In particular, in combination with single-particle spectra it allows the freeze-out temperature and transverse flow to be disentangled [99, 141, 142]. This works because the transverse homogeneity length grows with temperature and decreases with stronger flow gradient, while the single-particle spectrum gets flatter in response to both effects. Thus the flow–temperature correlation which leads to the same spectrum is different from that leading to the same m_t dependence of R_s . By measuring both these correlations, the temperature and transverse flow can be extracted.

Correlation measurements give information about the region where particles scatter for the last time. This is often simplified into a three-dimensional, so-called freeze-out hyper-surface. In general, however, particles escape when their mean free path is bigger than the homogeneity length and this may happen over a finite four-volume [143, 144].

The studies of the dynamics can be extended to non-central collisions [127]. Simulations of non-central collisions at SPS and RHIC show that elliptic flow builds up very early and is thus rather sensitive to the equation of state and carries information about the phase transition [83, 86]. Combining elliptic flow determined from single-particle spectra with azimuthally sensitive correlations will probe the influence of flow on the geometric shape and provide strong constraints on the dynamical evolution of the reaction zone.

Lifetime of the collision. Access to the emission duration may be achieved by measuring the difference $R_0^2 - R_s^2$ [99, 112, 125, 126, 145], although in general spatial information enters this difference as well. A large value of $R_0^2 - R_s^2$ was proposed as a signal for QGP formation, resulting from the softness of the equation of state at the phase-transition point [146]. SPS and RHIC data, however, do not show such a signal [106, 147, 148]. Instead, they suggest an extremely rapid expansion followed by a short decoupling period consistent with a hard equation of state. A satisfactory understanding has not yet been reached and theoretical expectations for the LHC are currently still open. In any case, a direct measurement of $R_0^2 - R_s^2$ or R_0/R_s is expected to constrain models of plasma formation. While the temporal contribution to $R_0^2 - R_s^2$ is expected to be the strongest at higher K_t , geometrical contributions to $R_0^2 - R_s^2$ increase with K_t , too. Thus, a measurement over a large range in K_t is necessary to gain insight into the relative contributions from spatial and temporal effects.

Moreover, the longitudinal radius parameter R_l gives access to the total lifetime of the system. At the LHC, where the collision region is expected to show Bjorken scaling [149] over many units in rapidity, the longitudinal radius should be directly proportional to this lifetime [138].

Final-state interactions. Momentum correlations due to FSI are often treated as an undesirable effect which has to be subtracted from the data [118, 121, 147]. Nevertheless, these correlations can also give access to temporal information. They carry space-time information about the source because of the dependence of the Coulomb wave function on the relative distance. In particular, non-identical particle correlations can be used to access the relative differences of production points by taking the ratio of two correlation functions with opposite orientation of their relative momentum vector q . Depending on the choice of this axis the respective space-time direction is probed. In particular, the correlation asymmetry in the outward direction is determined by the difference of emission times of the species used [150] and by the intensity of the transverse collective flow [151].

In addition, strong two-body interactions, for example between $\Lambda\Lambda$, are hardly known but can be determined by measuring the corresponding correlation function [152]. In a large system, the homogeneity radius is much bigger than the effective radius of the interaction, thus facilitating the extraction of the latter.

Phase-space density and expectations for LHC. It is unclear how large the correlation radii will be at the LHC. If the fireball volume at freeze-out grows proportionally to the multiplicity density, then the product $R_0 R_s R_l$ or $R_s^2 R_l$ increases linearly with dN/dy . This increase may in general be different for the different radii. The multiplicity estimates of section 1.3 and the value $dN_{ch}/dy \sim 650$ for RHIC lead to an average increase of the correlation radii by a factor $[(dN_{ch}/dy)_{LHC}/(dN_{ch}/dy)_{RHIC}]^{1/3} \approx 1.3-1.75$. This would result in radii of the order of 8–12 fm. However, larger radii (e.g. $R_l \simeq 15$ fm) cannot be firmly ruled out at present.

It is, however, conceivable that the freeze-out particle density is not a universal quantity [144, 153, 154]. The average phase-space density of pions was determined at the AGS [155, 156] and the SPS [154, 157, 158]. Moreover, preliminary data from RHIC indicate an increase in phase-space density at freeze-out. In general, if the phase-space density at freeze-out grows, correlation radii may increase by less than the factor $[(dN_{ch}/dy)_{LHC}/(dN_{ch}/dy)_{RHIC}]^{1/3}$ estimated above.

Should rather small HBT radii be observed at LHC, then one would expect the larger phase-space density at freeze-out to result in multi-boson effects. These may be observed by the simultaneous appearance of a dip in the correlation function at non-vanishing q occurring in

event samples with fixed multiplicity, a decrease of the intercept λ , and a low- p_t enhancement in the single-particle spectrum [159, 160].

Spin correlations. The correlation measurements discussed above can be generalized to correlations of particle spins using for the determination of spin either asymmetric weak particle decays [161] or particle scattering [162]; the correlations of protons from decays of two Λ 's being the simplest example. This technique does not require the construction of an uncorrelated reference sample entering the denominator in equation (1.17). Thus, it also serves as a consistency check for standard correlation measurements.

1.3.4. Fluctuations. Any physical quantity measured in an experiment is subject to fluctuations. In general, these fluctuations depend on the properties of the system under study and may reveal important information about the system. The most efficient way to address fluctuations of a system created in a heavy-ion collision is via the study of event-by-event fluctuations, where a given observable is measured on an event-by-event basis and the fluctuations are studied over the ensemble of events. In most cases, namely when the fluctuations are Gaussian [163], this analysis is equivalent to the measurement of two-particle correlations in the same region of acceptance. Consequently, fluctuations contain information about the two-point functions of the system, which in turn determine the response of the system to external perturbations.

In the framework of statistical physics, which appears to describe the bulk properties of heavy-ion collisions up to RHIC energies, fluctuations measure the so-called susceptibilities of the system. These susceptibilities determine the response of the system to external forces. For example, by measuring fluctuations of the net electric charge in a given rapidity interval, one obtains information on how this (sub)system would respond to the application of an external (static) electric field. In other words, by measuring fluctuations one gains access to the same fundamental properties of the system as ‘table-top’ experiments dealing with macroscopic probes. In the latter case, of course, fluctuation measurements would be impossible.

In addition, the study of fluctuations may reveal information beyond the thermodynamic properties of the system. As the system expands, fluctuations may be frozen in early and thus provide information about the properties of the system prior to its thermal freeze-out. A well-known example is that of the fluctuations in the cosmic wave background radiation, as first observed by COBE [164].

The field of event-by-event fluctuation is relatively new to heavy-ion physics and ideas and approaches are just being developed. So far, most of the analysis has concentrated on transverse-momentum and charge fluctuations.

Transverse-momentum fluctuations should be sensitive to temperature–energy fluctuations [165]. These in turn provide a measure of the heat capacity of the system [166]

$$\langle(\delta T)^2\rangle = \langle T^2\rangle - \langle T\rangle^2 = \frac{T^2}{C_V}. \quad (1.20)$$

Since the QCD phase transition is associated with a maximum of the specific heat, the temperature fluctuations should exhibit a minimum in the excitation function. It has also been argued [167] that these fluctuations may provide a signal for the long-range fluctuations associated with the tricritical point of the QCD phase diagram. In the vicinity of the critical point the transverse-momentum fluctuations should increase, leading to a maximum of the fluctuations in the excitation function.

Charge fluctuations [77, 78], on the other hand, are sensitive to the fractional charges carried by the quarks. Therefore, if an equilibrated partonic phase has been reached in these

collisions, the charge fluctuations per entropy would be about a factor of 2–3 smaller than in a hadronic scenario. One proposed observable is the fluctuation of ratio of positively charged to negatively charged particles,

$$D = \left\langle \left(\frac{N^+}{N^-} \right) \right\rangle \simeq 4 \frac{\langle (\delta Q)^2 \rangle}{\langle N_{\text{ch}} \rangle} \sim \frac{\langle (\delta Q)^2 \rangle}{S}. \quad (1.21)$$

More observables have been proposed in the mean time [168] and it remains to be seen which is the most useful one. Given a sufficiently large rapidity acceptance, it is expected that charge fluctuations should be frozen in [78, 169] and thus provide a hadronic signal of the partonic phase. This is similar to the fluctuations of the cosmic microwave background.

As the field of event-by-event studies progresses, there will be, and have already been, new ideas, such as fluctuations due to bubble formation [170] and disoriented chiral condensate (DCC) phenomena [172, 171]. The present situation of this very young field is described in [173, 174].

1.3.5. Jets. A hard hadronic collision at high energy may be pictured in the following way. Partons distributed in the hadronic projectiles are involved in a hard scattering with a large transfer of energy–momentum, whereas the non-colliding remnants of the incoming hadrons initiate what is usually called the ‘underlying event’. Then, the energetic coloured partons produced by the hard subprocess undergo a cascade of branchings which degrade their energies and momenta, as they escape from each other. Finally, the end points of this branching process and the remnants of the incoming projectiles fragment into a number of colourless hadrons during the so-called hadronization stage.

The hadronic final state may be analysed in various ways, depending on which information is sought. It may be partitioned into a number of clusters of hadrons, called ‘jets’. The standard way of defining jets in pp experiments is based on some suitably defined jet algorithm involving a calorimetric criterion [175]. The jet definition used fixes the size of a jet, gives the rule for the clustering of the hadrons, and assigns a transverse energy E_t , a pseudo-rapidity η , and an azimuthal angle φ to each jet. Roughly speaking, jet clustering aims at tracing back through the showering and hadronization stages to the hard partons which have initiated them, while the inner properties of the jets, such as the transverse-energy profile inside a jet, give information on the way the mother-parton’s transverse energy has been shared during the stage of showering.

This picture becomes more complicated due to the dependence of the jet multiplicity on the jet algorithm used, and the difficulties in establishing a correspondence between jets and the hard primary partons. In particular, if the hard process generates a large number of high- E_t partons, the products of the fragmentations of these partons may overlap. This is more likely for nucleus–nucleus collisions, in which the products of several hard collisions between nucleons of both nuclei are piled up incoherently. Also, this picture may not hold for jets whose E_t , though larger than the hadronic scale, is much smaller than the hadronic centre-of-mass energy \sqrt{s} , e.g. $E_t \lesssim 10$ GeV at the LHC. Such jets cannot be understood solely in terms of the fragmentation of partons of comparable E_t produced in a hard subprocess. These so-called minijets also receive a contribution from the dynamics of the underlying events, which in nucleus–nucleus collisions have a substantial transverse activity.

The hadronic final state may also be ‘projected’ so as to study the inclusive production of individual hadrons of identified species. In this case, most of the details of the hadronic final state become irrelevant, and the relevant information about the effects of showering and fragmentation leading to the hadron species considered is encoded in inclusive fragmentation functions of the species of partons into the considered hadron, which are much simpler objects than a full portrait of the final-state event.

In a Pb–Pb collision, calculations based on the QCD improved parton model, and treating the high-energy collision of two heavy nuclei as an incoherent superposition of high-energy collisions between nucleons from each nucleus, give an idea of what would be the enormous hadronic activity at large E_t , in the absence of nuclear modifications in the final state. In this framework, the inclusive cross section for producing a E_t parton takes the form

$$\frac{d\sigma}{dE_t d\eta d\varphi} = \sum_{ab} \int_{x_a^{\min}}^1 dx_a \int_{x_b^{\min}}^1 dx_b f_{a/A}(x_a, Q^2) f_{b/B}(x_b, Q^2) \frac{d\hat{\sigma}_{ab}}{dE_t d\eta d\varphi}, \quad (1.22)$$

where the $f_{i/X}$ are the parton distributions inside the nucleus X and $d\hat{\sigma}_{ab}/(dE_t d\eta d\varphi)$ is the perturbatively calculated hard-scattering cross section. These calculations predict as much as ~ 30 hard partons per event with $E_t \geq 10$ GeV, and still 3×10^{-3} hard partons per events with $E_t \geq 100$ GeV. Besides, the cross section for minijet production above 5 GeV transverse energy is $\sim 30\%$ of the total inelastic cross section, suggesting that, in contrast to RHIC, hard processes will contribute significantly to global observables and may play an important role in understanding the thermalization and thermal evolution of the system [48]. Moreover, the large jet cross section amounts, within the ALICE central acceptance ($|\eta| < 0.9$), to $\approx 10^6$ jets with $E_t > 100$ GeV per month of running at average design luminosity ($5 \times 10^{26} \text{ cm}^{-2} \text{ s}^{-1}$).

The high- E_t partons produced in the initial stage of a nucleus–nucleus collision are actually expected to undergo multiple scattering inside the collision region prior to fragmenting into hadrons. This multiple scattering is expected to induce modifications of the properties of the produced jets or individual particles, which probe the properties of the medium produced in the collision region. This is the main motivation for studying jet and individual hadron inclusive production in nucleus–nucleus collisions. The strategy is to identify these medium-induced modifications that characterize the hot and dense matter in the initial stage of the collision, by comparing the cross sections for the corresponding observables in A–A and benchmark pp collisions at the same centre-of-mass energy. An accurate understanding of jet and individual hadron inclusive production in pp is therefore quite important in order that this strategy be successful. In this respect, the LHC will open a new kinematic regime, in which the pp collisions involve features which are not well understood yet. Therefore, the ALICE experimental programme will also involve specific studies on jet and high- p_t particle production in pp collisions.

Jets in proton–proton collisions. The standard jet definitions used in pp experiments rely on a calorimetric criterion. The challenge for ALICE, where no extensive hadronic calorimetry is available so far, is therefore to define and construct jets out of tracking measurements. Recent progress in this respect has been reported by the CDF Collaboration at the Fermilab Tevatron. CDF has extensively studied the properties of jets in $p\bar{p}$ collisions [176] in the low-energy region (up to 50 GeV of charged-particle energy and in $|\eta| < 1$) measuring only the charged particles in the jets. In minimum-bias data they observed jets with total charged-particle momenta of $\sim 2 \text{ GeV } c^{-1}$, with two charged particles on average, while for jets up to $50 \text{ GeV } c^{-1} \sim 10$ charged particles were detected. The QCD Monte Carlo models describe quite well jet observables such as the multiplicity distribution of charged particles within the leading jet (i.e. the jet with the largest energy in the event), the radial distributions of charged particles, and their transverse momentum. Surprisingly the agreement of the Monte Carlo calculation with the data is as good at $5 \text{ GeV } c^{-1}$ as at $50 \text{ GeV } c^{-1}$. ALICE will reconstruct charged particles with sufficient accuracy to allow similar measurements up to charged particle momenta of $\sim 100 \text{ GeV } c^{-1}$. In addition, charged-particle identification up to $5 \text{ GeV } c^{-1}$ for protons will allow detailed comparisons with the QCD Monte Carlo models, thus providing

a benchmark for comparison of the fragmentation function of similar energy jets in heavy-ion collisions where medium effects are expected.

Jet production in hadronic collisions with a momentum transfer of the order of the centre-of-mass energy, in which case the process is characterized by a single large-energy scale factor, is successfully described by perturbative QCD. In this regime, higher orders in the perturbative expansion of the elementary interaction, i.e. in the QCD coupling constant α_s , are related directly to processes where an increasing number of large- E_t jets are produced. However, the simplest perturbative scheme is no longer adequate when two or more different scale factors become relevant. An interesting case where two scales play an important role is when the centre-of-mass energy of the partonic process, although much larger than the hadronic scale, is nevertheless very small as compared with the centre-of-mass energy of the hadronic interaction. In this kinematical regime the typical final state is characterized by the presence of many jets and, while the partonic collision is still well described by the conventional perturbative approach, the overall semihard component of the hadronic interaction acquires a much richer structure. A recent analysis performed by CDF [176] of the jet evolution and underlying event in $p\bar{p}$ collisions shows that, while the leading jet is fairly well described by Monte Carlo models based on perturbative QCD in a wide kinematical range, the models fail to describe correctly the next few jets at a relatively low p_t .

Jets in A–A collisions. Bjorken [177] stressed as early as 1982 that a ‘high- p_t quark or gluon might lose tens of GeV of its initial transverse momentum while plowing through quark–gluon plasma produced in its local environment’. While Bjorken’s estimates based on collisional energy loss had to be revised (for subsequent studies of collisional energy loss see [178, 179]), Gyulassy, Plümer and Wang [180] suggested that the dominant energy-loss mechanism is radiative rather than collisional. The work of Baier, Dokshitzer, Mueller, Peigné and Schiff finally identified the dominant radiative mechanism: it is not the direct analogue of the Abelian bremsstrahlung radiation but a genuine non-Abelian effect, namely, gluon rescattering [181]. To quantify this jet ‘quenching effect’, several groups [182–185] calculated in recent years the leading order in $1/E$ of the complete non-Abelian quantum interference of gluon emission. These approaches differ in details, especially (i) in resumming the effect of either a small fixed number [184, 185] or of arbitrarily many scattering centres [182–184]; (ii) in working either for an arbitrary finite [184, 185] or in the limit of very large [183, 184] in-medium path-length, and (iii) in kinematical assumptions which favour the relative importance of multiple small-angle scattering [182–184] or single large-angle scattering [185] in the medium. Still, these formalisms are just specific limiting cases of the same formalism [184]. For a review of the current state of the art, see [186].

The main conclusions of these studies are: (i) non-Abelian parton energy losses grow quadratically, $\propto L^2$, with the in-medium path-length [183] thus being very sensitive to the geometry of the collision region; (ii) p_t -broadening and energy losses of a jet are determined by the same parameter, a kinetic transport coefficient [183] which characterizes the amount of transverse momentum transferred to the hard parton per unit path-length and is directly related to the phase-space density attained in the initial stage of the collision; (iii) the interplay between destructive interference and p_t -broadening results in a characteristic, non-monotonic dependence of energy loss on the jet opening angle [184]; (iv) for low jet energies, such as $p_t < 10$ GeV, calculations suggest a characteristic $\ln E$ dependence of the medium modifications [185]. Numerical estimates of the jet-quenching effect depend sensitively on the scattering properties of the medium. Most studies conclude that high- p_t partons can lose up to ~ 20 – 80% of their initial energy prior to hadronization. First data from RHIC [187–189] are discussed in terms of these calculations.

In addition to the above-mentioned multiple scattering eikonal approaches, there are complementary calculations [190] which consider only one additional hard scattering described in perturbative QCD in terms of four-point parton correlation functions of the nucleus. This approach was developed [191] and used extensively in pA collisions (e.g. for the nuclear dependence of high- p_t Drell–Yan production) where rigorously proven factorization theorems [191] ensure its applicability. It was recently extended to the description of medium modification of parton fragmentation functions. Although the potential of this approach is not yet fully explored, first studies confirm some of the main features of the multiple-scattering approach and, in particular, the characteristic L^2 -dependence of the energy loss [190].

Parton-energy loss prior to hadronization is expected to affect essentially all high- p_t hadronic observables, as well as leptonic decay products of hadronic bound states. For a hadronic spectrum in nucleus–nucleus collisions, such energy loss effects can be accounted for by modifying the typical factorized expression for the cross section [192]:

$$E_h \frac{dN_h}{d^3p} = T_{A-A} \sum_{abcd} \int dx_a dx_b f_{a/A}(x_a, Q^2) f_{b/A}(x_b, Q^2) \frac{d\sigma^{ab \rightarrow cd}}{d\hat{t}} \times \int d\varepsilon P(\varepsilon) \frac{1}{1-\varepsilon} \frac{D_{h/c}[x_c/(1-\varepsilon), Q^2]}{\pi x_c}. \quad (1.23)$$

Here, $D_{h/c}$ is the fragmentation function and $f_{a/A}$ the corresponding structure function for partons a in nucleus A . $P(\varepsilon)$ denotes the probability that a fraction ε of the initial energy of the hard parton is lost due to gluon radiation because of multiple scattering in the medium. The importance of the probability distribution $P(\varepsilon)$ is emphasized in [193] in a slightly different setting. A numerical routine to incorporate $P(\varepsilon)$ in event generator studies is available [194]. The ε -integration of $D_{h/c}$ with weight $P(\varepsilon)$ can be considered as a medium-modified fragmentation function.

In general, parton-energy loss cannot be reduced to the discussion of medium-modified fragmentation functions that apply for the leading hadron only. For example, to study the medium modification of the inclusive jet cross section, equation (1.22), one has to calculate how much energy is radiated outside a given jet cone. The relevant angular dependence of the medium induced gluon-radiation spectrum is known [195, 196], but its application to equation (1.22) is not yet available.

In order to get to numerical predictions, early theoretical studies made significant approximations: energy loss was often taken to be linear in the medium path-length, and spectra were often shifted by a constant energy loss without taking the p_t -dependence of hadronic spectra into account. Such studies are of limited use only, since they neglect: (i) that the dependence of energy loss on the in-medium path-length is largely quadratic [183] with further non-linear features due to small finite size effects [195] and (ii) that, in particular, due to the statistics of multiple-gluon emission, the effect of the partonic-energy loss probability $P(\varepsilon)$ [193, 194, 197] is significantly different from a shift. Nevertheless, these earlier studies identify qualitatively the most promising observables for the study of parton-energy loss:

- Reduction in the yield of high- p_t particles [180].
- Particle ratios at high p_t . Owing to their different colour representation, hard gluons are expected to lose approximately a factor two more energy than hard quarks. Depending on the relative contribution of gluon fragmentation, this modifies the ratio of hadronic species. For example, the ratio \bar{p}/p is expected to decrease.
- Dependence of hadronic spectra on the nuclear geometry [198]. Jet-quenching phenomena are expected to show a strong and characteristic dependence on the impact parameter of the collision as well as on their orientation with respect to the reaction plane.

- Attenuation and p_t -broadening of all types of ‘true’ jet phenomena. This includes the total energy radiated within a jet cone, the broadening of the transverse-energy distribution in the jet [199, 200], and jet tagging [192] by determining the hadronic activity in the direction opposite to a hard photon or Z^0 .
- Parton energy loss manifests itself as a softening and distortion of the jet multiplicity distribution. A full characterization of the jet fragmentation function thus requires a measurement of both the soft and hard components of the jet.
- Medium-induced modification of heavy quark fragmentation functions [201–204] reflected in the p_t -spectra of open charm and bottom.

1.3.6. Direct photons. Before discussing photon production in nucleus–nucleus collisions, we briefly review what is known about the corresponding topics as benchmark in pp collisions.

Photons in proton–proton collisions. Although the production of photons at large transverse momentum in proton–proton collisions has been extensively studied over the last 20 years, both theoretically and experimentally, no good agreement between experiment and theory has yet been achieved. The rate of production is essentially proportional to the gluon distribution in the proton which can be probed directly by looking at the transverse-momentum dependence of the photon spectrum. At the LHC two new features appear:

- At not too high p_t (up to 20 GeV), the production of a photon by bremsstrahlung from a gluon is dominant (the photon is radiated from the quark content in the gluon). However, little is known about gluon fragmentation into a photon [205]. Indeed, the current uncertainty in the fragmentation functions into a photon results in a factor 2 uncertainty in the prompt photon rate for $p_t < 20$ GeV.
- For photon production at small transverse momentum, the region of small $x = 2p_t/\sqrt{s}$ (for $p_t = 2$ GeV, this corresponds to $x = 3 \times 10^{-4}$ at $\sqrt{s} = 14$ GeV) is relevant where next-to-leading logarithm calculations [206] become insufficient and the only recently explored ‘recoil resummation’ [207] is needed.

There is no consensus on the interpretation of the present data (from $\sqrt{s} = 17$ GeV to 2 TeV) within QCD. Some groups advocate the introduction of ‘intrinsic transverse’ momentum to fit the data [208]. However, the model dependence and lack of predictive power of this approach have been criticized [209]. Data in the small- p_t domain at LHC may contribute to clarifying this issue and its relation to the recoil resummation [207].

Since the bremsstrahlung production of photons from final-state quarks remains important at high p_t (about 40% at $p_t = 50$ GeV) and is not precisely known, the extrapolation from lower energy data to LHC energies is complicated at high- p_t values, too. This will render, for example, the calibration of the energy of a recoiling jet via the photon-energy measurement a difficult task.

An almost unlimited number of correlation functions such as photon–jet, photon–hadron and photon–photon can also be studied. In this way, many of the inputs entering the theoretical calculations, in particular the fragmentation functions [210], can be tested. As an example, photon–photon correlations between a directly produced very high- p_t photon and a low- p_t one produced by bremsstrahlung of a recoiling quark or gluon will allow us to measure the fragmentation functions of partons into photons. It provides the necessary checks on the same quantities extracted from single-photon production at low p_t .

In summary, a detailed study of prompt-photon production in proton–proton collisions at the LHC is needed (i) to constrain non-perturbative fragmentation parameters in a new

kinematical range, (ii) to test new theoretical ideas on the necessary resummation to be performed at low x and (iii) to obtain benchmark cross sections for the study of A–A collisions.

Photons in A–A collisions. In 1976 Feinberg [211] suggested that the production of photons could be used to probe the dynamics of strong interactions in hadronic collisions: basically the idea was that, during the scattering process, many photons would be radiated off the hadron constituents which collided. More than 25 years later, this picture remains essentially valid, although it has become much more complex since several mechanisms are at work in the different temporal stages of ultra-relativistic heavy-ion collisions. Indeed, owing to their small electromagnetic coupling, photons once produced do not interact with the surrounding matter and thus probe the state of matter at the time of their production. The production of photons in the different stages of a heavy-ion collision can be summarized qualitatively as follows:

1. Early in the collision, so-called ‘prompt’ photons are produced by parton–parton scattering in the primary nucleus–nucleus collisions, as in the nucleon–nucleon case. For large enough values of the photon transverse momentum p_t , this process can be calculated in perturbative QCD. Although their rate decreases as an inverse power of p_t , photons up to several hundred GeV are expected to be detected at LHC. An important background to direct photon production is the decay $\pi^0 \rightarrow \gamma\gamma$, produced in similar hard partonic processes.
2. In the following stage of the collision, a bubble of quark–gluon plasma is expected to be formed with a temperature of up to 1 GeV. Photons are radiated off the quarks which undergo collisions with other quarks and gluons in the thermal bath. The energy spectrum of these photons is exponentially suppressed but should extend up to several GeV.
3. The plasma expands and cools. At the critical temperature ($T_c = 150\text{--}200$ MeV), a hadronic phase is formed. In this phase, photons can be produced in the scattering of π 's, ρ 's, ω 's, and so on or in resonance decays. This mechanism continues until the resonances cease to interact, i.e. until the freeze-out temperature of the order of 100 MeV is reached. Photons produced in this stage will have an energy in the range from a few hundred MeV to several GeV.
4. Finally, after freeze-out, further photons can be produced by the decay of π^0 's, η 's and higher resonances. Their energy lies in the range of up to a few hundred MeV.

Photons produced by π^0 decays, either from the primary collisions or final-state contribution, constitute a large ‘reducible’ background to ‘direct’ photon production over the whole p_t range. They can be subtracted from the data. On the other hand, the prompt photons of phase (1) provide an irreducible background and a precise estimate of their rate (e.g. via comparison to the pp benchmark) is needed to extract the rate of thermal photons. These thermal photons are emitted during phases (2) and (3). In general, for the study of thermal effects, one can distinguish two types of observables corresponding to two different kinematical regimes:

- For photon energies in the few GeV range, an excess in the observed inclusive rate, after subtraction of the background from decay photons, over the rate predicted for phase (1) would be a clear indication of a thermalized medium. Here, the signal relies essentially on understanding the normalization of the photon rate and, to a lesser extent, on the shape of the distribution.
- For higher photon energies, above 20 GeV, the production mechanism is necessarily of type (1) since the exponentially damped thermal production has become irrelevant at these energies. Such hard photons are produced in association with a recoil jet, the fragmentation

properties of which are affected by its interaction with the hot matter, as discussed in section 1.3.5. The comparison of photon–hadron correlations in heavy-ion and in nucleon–nucleon collisions can be used to probe this effect. Similarly two-photon correlation functions would be interesting in a certain kinematical domain where the softer photon is produced by bremsstrahlung from a quark. Here, thermal effects will show up in the shape of the correlation functions.

Concerning low-energy photons, the background from π^0 decays is large; this makes it difficult to extract the direct photon signal. However, since this prompt-photon signal arises partly from a bremsstrahlung process, it will be affected by the hot medium where the hard parton will lose energy before radiating the photon. First studies indicate that this reduces the prompt-photon background thus making the extraction of the thermal-photon rate easier [212]. Recent theoretical progress on the calculation of the thermal-photon production rate in a quark–gluon plasma was made within finite temperature QCD, especially in the Hard Thermal Loop (HTL) effective theory formulation [34, 213]. The mechanisms included here are Compton and annihilation processes [214], bremsstrahlung, the off-shell annihilation or annihilation with scattering which has no equivalent at zero temperature and dominates the rate below a few GeV [215], as well as the Landau–Pomeranchuk–Migdal rescattering effect [216]. The extension of this formalism to the non-equilibrium case [217] and to small-mass lepton pairs [218] at large momentum has been considered. The latter process is interesting since theoretically it gives access to the same dynamical information as the rate of real photons, while experimentally the background is different.

Past experience with photon physics in hadronic collisions has shown, from fixed-target to Tevatron energies, that the extrapolation of rates from one experiment to another, or from one energy to another is not obvious. Reliable tests of thermal-photon production will be obtained only when comparing data collected in proton–proton collisions (no thermal effects), proton–nucleus collisions (nuclear effects but no thermal effects) and nucleus–nucleus collisions (both nuclear and thermal effects present). As seen above, a variety of observables (rates, correlations) are at our disposal to disentangle thermal effects from standard QCD effects.

1.3.7. Dileptons. Dilepton production is an important tool for measuring the temperature and the dynamical properties of the matter produced in a relativistic heavy-ion collision. Lepton pairs are emitted throughout the evolution of the system, as is the case for photons, and the stages of dilepton production are analogous to those of photons. Specifically, these are prompt contributions from hard nucleon–nucleon collisions, thermal radiation from the QGP and hot hadronic phases as well as final-state meson decays after freeze-out.

The prompt contribution to the continuum in the dilepton mass range above pair mass $M \sim 2$ GeV is dominated by semileptonic decays of heavy-flavour mesons and by the Drell-Yan process [219]. These leptons originate from hard scatterings so that their rates can be calculated in perturbative QCD.

Since more than one heavy-quark pair will be produced in each central A–A collision, heavy-quark decays will dominate the lepton pair continuum between the J/ψ and Z^0 peaks [220]. Although heavy quark energy loss in the medium is currently expected to be small [201], it could nevertheless significantly alter the shape of the dilepton continuum arising from heavy quark decays [202, 203]. Additional effects are expected from nuclear shadowing [221] which modifies the parton distribution functions in the nucleus, affecting the heavy quark production mechanism. The PHENIX Collaboration at RHIC has already demonstrated that charm can be measured using single leptons [222]. At higher lepton p_t , bottom decays will dominate the single lepton rate. The charm and bottom rates can also be measured

in $e\mu$ coincidence studies. Thus, although the heavy quark cross sections still suffer from theoretical uncertainties, measurements of high mass lepton pairs and large p_t single leptons can reduce these uncertainties as well as provide valuable information about heavy quark energy loss and the nuclear gluon distribution.

In the high mass part of the lepton pair spectrum, the $Z^0 \rightarrow l^+l^-$ rate ($l = e, \mu$) is sensitive to the nuclear quark and antiquark densities at higher values of Q^2 than previously available [223]. High Q^2 pA measurements will probe the x range around 0.015 at $y = 0$. This range of x has already been probed in fixed target nuclear deep-inelastic scattering so that a measurement at the LHC will constrain the evolution of the nuclear quark distributions over a large lever arm in Q^2 . This knowledge will aid in studies of jet quenching with a Z^0 opposite a jet in A–A collisions [220].

The thermal production rates for both dileptons and photons are based on the imaginary part of the electromagnetic current–current correlation function [224] in hot and dense matter. The pertinent emissivity of lepton pairs, which can be regarded as decays of virtual (time-like) photons, from equilibrated matter in A–A collisions is suppressed by an additional power of $\alpha = 1/137$ compared to real photons. However, dilepton observables offer distinct advantages over those of photons. In particular, lepton pairs carry an additional variable, the pair invariant mass, which encodes dynamical information on the vector excitations of the matter.

At masses above ~ 1.5 GeV thermal radiation is very sensitive to temperature variations and thus expected to originate from the early hot phases with a rather structureless emission rate determined by perturbative $q\bar{q}$ annihilation. The physics objective [225, 226] is then similar to that of the photon case, i.e. the discrimination of thermal QGP radiation from the large prompt background [219]. Note, however, that in contrast to thermal photons, the leading term in dilepton emission is of order 1, rather than order α_s for photons, and thus, in principle, under better theoretical control.

At low masses, less than 1.5 GeV, thermal dilepton spectra are dominated by radiation from the hot hadronic phase [227, 228]. Here, the electromagnetic current is saturated by the light vector mesons (ρ , ω and ϕ), allowing direct access to their in-medium modifications. The ρ plays a key role since its e^+e^- decay width is a factor of ~ 10 (5) larger than the ω (ϕ). In addition, the ρ has a well-defined partner under SU(2) chiral transformations, the $a_1(1260)$. The approach towards restoration of chiral symmetry at T_c therefore requires the spectral distributions in the corresponding vector and axial-vector channels to become degenerate. How this degeneracy occurs is one of the crucial questions related to the chiral phase transition. The possibilities range from both ρ and a_1 masses dropping to (almost) zero, the so-called Brown–Rho scaling conjecture [229], to a complete melting of the resonance structures due to intense rescattering in the hot and dense hadronic environment [230], or scenarios with rather stable resonance structures [231, 232]. Clearly, low-mass dilepton observables such as the invariant mass and pair transverse momentum spectra are necessary to find an answer. Despite their relatively low rates, dilepton decays of the narrower ω and ϕ mesons within the fireball can also provide useful information on their in-medium line-shapes [228]. The low mass background primarily consists of π^0 , η and ω Dalitz decays which can be rather reliably inferred from hadronic production systematics.

1.3.8. Heavy-quark and quarkonium production. Heavy quarks (charm and bottom) provide sensitive probes of the collision dynamics at both short and long timescales. On the one hand, heavy-quark production is an intrinsically perturbative phenomenon which takes place on a timescale of the order of $1/m_Q$ where m_Q is the heavy-quark mass. On the other hand, the long lifetime of charm and bottom quarks allows them to live through the thermalization phase of the plasma and to possibly be affected by its presence. In addition, the high temperature of

the plasma may give rise to thermal production of heavy-quark pairs. Finally, heavy-quark–antiquark pairs can form quarkonium states which are bound by energies of the order of a few hundred MeV. These binding energies are comparable in size with the mean energies ($\sim 3T_c$) of the plasma, implying a large probability for quarkonium breakup.

The ability to extract information about the plasma from the features of heavy-quark and quarkonium production in A–A collisions relies on our understanding of production in pp and pA interactions. Typical observables of interest include total production rates, transverse-momentum distributions and kinematical correlations between the heavy quark and antiquark. The total cross section is the most inclusive quantity, and in nucleon–nucleon collisions it is mostly determined by the gluon density of the proton. The knowledge of the gluon density in a nucleus A should then fix the total rate of heavy-quark production in pA collisions, as well as the perturbative, short-distance component of the heavy-quark rate in A–A collisions. Additional contributions, coming, for example, from the high-energy tails of the plasma spectra, should then manifest themselves through differences between the total rates observed in A–A collisions and those estimated in perturbation theory using nuclear parton densities. Given the large theoretical uncertainties, however, a direct measurement of the total cross sections in both pp and A–A will be absolutely necessary. Unfortunately no data are available on total charm cross sections in high-energy hadronic collisions. Trigger thresholds and detector configurations prevent the Tevatron experiments from carrying out such a measurement. The ability of ALICE to monitor the charm and bottom spectra in the region of very low transverse momentum, where the bulk of the cross section lies, will then be essential. One would ideally like to normalize the heavy-quark rates in A–A collisions with those in pp collisions at the same centre-of-mass energy per nucleon. In the absence of a pp run at 5.5 TeV, one should interpolate the results obtained at 14 TeV. This interpolation is expected to be rather reliable.

Studies of the heavy-quark spectra should allow a more detailed understanding of the non-perturbative component expected to be concentrated at very small transverse momenta since the high- p_t tails are exponentially suppressed relative to the perturbative component. At the same time, the behaviour of the spectra at large transverse momenta should display the energy loss suffered by a perturbatively produced quark during its motion through the plasma. Measuring the heavy quark p_t spectrum in pp collisions will then be crucial in order to separate the effects induced by the nuclear medium in A–A collisions. Additional observables should complement and complete this study. For example, it would be interesting to be able to monitor the soft-particle activity in a cone surrounding the heavy quark, activity which should be increased in the nuclear case due to the interactions of the high- p_t massive quark with the medium.

Less than 1% of all heavy-quark pairs form quarkonium bound states. Therefore, we do not expect phenomena such as quarkonium suppression to have any impact on the open heavy-quark rates and spectra. Conversely, any phenomenon which changes the inclusive heavy-quark production rates will affect the quarkonium yield. The study of correlations between the properties of open heavy-quarks and quarkonia spectra provides an independent probe of the dynamics of the high-density medium. Additional production of quark–antiquark pairs in the plasma will increase the chances of quarkonium formation, but we expect this increase to be limited to the very low- p_t region. At large p_t , where we expect the perturbative production mechanisms to dominate, one should observe the phenomenon of quarkonium suppression. As a result, the dependence of the quarkonium yield on the transverse momentum, normalized to expectations in absence of the nuclear medium or plasma, is expected to discriminate between different models of quarkonium production. The correlation with the momentum dependence of open heavy-quark rates should disentangle the effects of enhancement in the production of

quark–antiquark pairs from the effects of quarkonium suppression. Once again, the ability to cover a significant range of transverse momenta, and in particular the very low- p_t region, for both open quark and quarkonium states, is essential. For normalization purposes, and to help clarifying several puzzles in charmonium production still left unsolved at the Tevatron, a run in pp and pA mode is also mandatory. ALICE is probably the only LHC experiment with access to the low- p_t region.

Heavy-quark production. The heavy-quark cross section has been calculated to next-to-leading order (NLO), by several groups [233–235]. Only the calculation in [234, 235] deals with the full differential production of $Q\bar{Q}$ pairs. Beyond NLO, a complete calculation has not been done so far. Near threshold, the heavy-quark cross section has been resummed to leading logarithm [236–238] and next-to-leading logarithm [239–241]. There are some uncertainties about how the resummed cross section is obtained. This can be avoided by sticking to a fixed-order calculation, see [242] for a recent discussion. A fixed-order calculation of the total cross section to next-to-next-to-leading order and next-to-next-to-leading logarithm has become available, but is also only applicable near threshold [243]. Charm and bottom production in heavy-ion collisions at the LHC is far above production threshold so that the most relevant calculations remain next-to-leading order.

At any order, the partonic cross section may be expressed in terms of dimensionless scaling functions $f_{ij}^{(k,l)}$ that depend only on the variable η [243]:

$$\hat{\sigma}_{ij}(\hat{s}, m_Q^2, \mu^2) = \frac{\alpha_s^2(\mu^2)}{m^2} \sum_{k=0}^{\infty} (4\pi\alpha_s(\mu^2))^k \sum_{l=0}^k f_{ij}^{(k,l)}(\eta) \ln^l \left(\frac{\mu^2}{m_Q^2} \right), \quad (1.24)$$

where \hat{s} is the partonic centre-of-mass energy squared, m_Q the heavy-quark mass, μ the renormalization scale and $\eta = \hat{s}/4m_Q^2 - 1$. The summation is expanded in powers of α_s , with $k = 0$ corresponding to the $\mathcal{O}(\alpha_s^2)$ Born cross section while $k = 1$ corresponds to the $\mathcal{O}(\alpha_s^3)$ NLO cross section. It is only when $k \geq 1$ that the renormalization scale dependence enters the calculation (aside from the scale dependence of α_s which is already present at leading order) since when $k = 1$ and $l = 1$, the logarithm $\ln(\mu^2/m_Q^2)$ appears. Note that equation (25) assumes that the renormalization and factorization scales, μ_R and μ_F respectively, are equal, as is also assumed in global analyses of the parton densities.

For the total charm and bottom production cross sections at the LHC, NLO estimates can be obtained in two ways. One can start from values of m_Q and $\mu_R = \mu_F = \mu \propto m_Q$ that give agreement with the bulk of the available data and then extrapolate to the energy of the LHC [244, 245]. At 5.5 TeV, using the MRST HO [246] parton densities, the pp result is 6.3 mb for charm and 0.19 mb for bottom [245]. Alternatively, one changes m_Q , μ_R , and μ_F independently to provide upper and lower bounds on the predicted cross sections [247, 248]. Using modern parton densities, the range of predictions in the second method varies by up to a factor of four, 4–15 mb for charm and from 0.08 to 0.34 mb for bottom. Thus, while the central values obtained from both approaches are consistent, the latter method gives a more conservative error estimate. At variance with these approaches are simulations with the PYTHIA event generator [249] which gives cross sections larger than the NLO evaluations, ~ 20 mb for charm and ~ 0.5 mb for bottom [249].

Beyond the measurement of total cross sections, the transverse-momentum and rapidity dependence of the expected quark and $Q\bar{Q}$ -pair distributions is important both to determine the rate within a finite acceptance and to have a baseline expectation against which to compare nuclear and quark–gluon plasma effects. The NLO calculations of heavy-quark production

in pp interactions include fragmentation of the heavy quarks [250] which degrades the quark momentum in the final-state hadron and an intrinsic transverse-momentum kick, $\langle p_t^2 \rangle_{pp} = 1 \text{ GeV}^2$ [247] which broadens the transverse momentum. Nuclear effects (shadowing and p_t broadening) have been included in the NLO code [245]. More generally, shadowing is reduced with increasing Q^2 and x while p_t broadening increases proportionally to the number of collisions. Shadowing as parametrized by EKS98 [251] results in a 35% reduction of the total $c\bar{c}$ production at 5.5 TeV but only in a 15% reduction in the $b\bar{b}$ rate [245]. Alternatively, the current knowledge of shadowing effects and in particular of the nuclear gluon distribution can be improved by comparing heavy quark production in pp and pA collisions at the same energy [252].

It is important to have an accurate measure of the charm and bottom cross sections for several reasons. Heavy-quark decays are expected to dominate the lepton-pair continuum up to the mass of the Z^0 [202, 219, 220]. Thus the Drell–Yan yield and any thermal dilepton production will essentially be hidden by the heavy-quark decay contributions [219]. The shape of the charm and bottom contributions to this continuum could be significantly altered by heavy-quark energy loss [202, 203]. So far, the amount of the energy lost by heavy quarks is unknown. While a number of calculations have been made of the collisional loss in a quark–gluon plasma [177–179, 253], only recently calculations of radiative loss have been applied to heavy quarks [201–203, 254]. The radiative loss can be rather large, $dE/dx \sim -5 \text{ GeV fm}^{-1}$ for a 10 GeV heavy quark, and increasing with energy, but the collisional loss is smaller, $dE/dx \sim -(1-2) \text{ GeV fm}^{-1}$, and nearly independent of energy [254]. If the loss is large enough, it may be possible to extract a thermal dilepton yield if it cannot be determined by other means [255]. Heavy-quark production in a quark–gluon plasma has also been predicted [256–258]. This additional yield can only be determined if the A–A rate can be accurately predicted. Finally, the total charm rate would be a useful reference for J/ψ production since an enhancement of the J/ψ to total charm ratio over the extrapolated pp yield has been predicted in a number of models [259–266].

Parton energy loss does not reduce the number of $Q\bar{Q}$ pairs produced but only changes their momentum. However, an effective reduction in the observed heavy-quark yield can be expected in a finite acceptance detector because fewer leptons from the subsequent decays of the heavy quarks will pass kinematic cuts such as a minimum lepton p_t . While there is no firm knowledge about the magnitude of heavy-quark energy loss, heavy quarks are expected to lose less energy than light quarks because of the ‘dead-cone effect’ [201].

In addition to the hard primary production, secondary $Q\bar{Q}$ production in the quark–gluon plasma has been considered [256–258]. At high temperatures, thermal charm production could be considerable since m_c is only 20–50% larger than the highest predicted temperature at the LHC. The thermal yield from a plasma of massless quarks and gluons is probably not comparable with initial production. These thermal charm pairs would have lower invariant masses than the initial $c\bar{c}$ pairs and would thus be much more susceptible to lepton- p_t cuts. Thermal bottom production is not likely because m_b is nearly a factor of five greater than the highest expected temperature.

Lastly, the total heavy-quark yield would be a valuable reference against which to compare quarkonium production since some models predict that the J/ψ to total charm ratio will be enhanced considerably in a plasma. For this comparison, a count of the total charm yield, open charm as well as charmonium, is necessary to obtain an accurate comparison. This will be virtually impossible given finite acceptances and the inability to detect some charmonium states such as χ_c in the high-multiplicity environment. Ideally, charm pairs should be detected because assumptions about unreconstructed charm have led to overestimates of the total charm cross section in previous experiments [267].

Quarkonium production. In the quark–gluon plasma, quarkonium suppression is expected to occur owing to the shielding of the $c\bar{c}$ binding potential by colour screening, leading to the breakup of the quarkonium states, first the ψ' and χ_c , and finally the J/ψ itself [268, 269]. For the much higher energies of nuclear collisions at the LHC, the Υ family will also be copiously produced and any possible Υ suppression pattern can be studied in detail [270].

In order to understand quarkonium suppression, a good estimate of the expected yields without medium effects is needed. However, significant uncertainties remain in the description of quarkonium production in nucleon–nucleon collisions. We start by reviewing two approaches that have been used to describe quarkonium production phenomenologically—the colour evaporation model (CEM) [271] and non-relativistic QCD (NRQCD) [272].

In the CEM, the $Q\bar{Q}$ pair neutralizes its colour by interaction with the collision-induced colour field—‘colour evaporation’. The Q and the \bar{Q} either combine with light quarks to produce heavy-flavoured hadrons or bind with each other in a quarkonium state. The additional energy needed to produce heavy-flavoured hadrons is obtained non-perturbatively from the colour field in the interaction region. The yield of all quarkonium states may be only a small fraction of the total $Q\bar{Q}$ cross section below the heavy-hadron threshold, $2m_H$. At leading order, the production cross section of quarkonium state C is

$$\sigma_C^{\text{CEM}} = F_C \sum_{i,j} \int_{4m_Q^2}^{4m_H^2} d\hat{s} \int_0^1 dx_1 dx_2 f_{i/A}(x_1, \mu^2) f_{j/B}(x_2, \mu^2) \hat{\sigma}_{ij}(\hat{s}) \delta(\hat{s} - x_1 x_2 s), \quad (1.25)$$

where $ij = q\bar{q}$ or gg and $\hat{\sigma}_{ij}(\hat{s})$ is the $ij \rightarrow Q\bar{Q}$ subprocess cross section. The factor F_C is independent of the kinematics, differs for each state and depends on m_Q , the scale μ in the strong coupling constant α_s , and the parton densities. The CEM was taken to next-to-leading order (NLO) using exclusive $Q\bar{Q}$ hadroproduction [235] to obtain the energy, x -, and p_t -dependence of quarkonium production [23, 273]. In the CEM, $gg \rightarrow g(g^* \rightarrow Q\bar{Q})$, incorporated at NLO, is similar to models of $g^* \rightarrow \Upsilon$ fragmentation [274]. By including this splitting, the CEM provides a good description of the quarkonium- p_t distributions at the Tevatron. The CEM cross sections were calculated with the MRST HO [246] parton densities using $m_c = \mu/2 = 1.2 \text{ GeV}$ and $m_b = \mu = 4.75 \text{ GeV}$, the same values used for the NLO evaluations of the heavy-quark cross sections [245]. The values of F_C for charmonium and bottomium have been calculated from fits to the J/ψ and Υ data combined with relative cross sections and branching ratios, see [270, 275] for details. The results, with and without shadowing using the EKS98 parametrization [251], are shown in table 1.1.

NRQCD describes quarkonium production as an expansion in powers of v , the relative $Q\bar{Q}$ velocity. The NRQCD matrix elements specify the initial angular momentum and colour of the produced $Q\bar{Q}$ pairs. Thus NRQCD goes beyond the leading colour singlet state to include colour octet production. The production cross section of quarkonium state C in NRQCD is

$$\sigma_C^{\text{NRQCD}} = \sum_{i,j} \sum_n \int_0^1 dx_1 dx_2 f_{i/A}(x_1, \mu^2) f_{j/B}(x_2, \mu^2) C_{Q\bar{Q}[n]}^{ij}(\mu^2) \langle \mathcal{O}_n^C \rangle, \quad (1.26)$$

where the partonic cross section is the product of perturbative expansion coefficients, $C_{Q\bar{Q}[n]}^{ij}(\mu^2)$, and non-perturbative parameters describing the hadronization, $\langle \mathcal{O}_n^C \rangle$. The parameters determined by Beneke and Rothstein for fixed-target hadro-production using the CTEQ3L parton densities [276] with $m_c = 1.5 \text{ GeV}$, $m_b = 4.9 \text{ GeV}$ and $\mu = 2m_Q$ [277] are used here. Since the parameters $\langle \mathcal{O}_n^C \rangle$ are fitted to the LO calculation with a LO set of parton densities, no further K factor is required. Different NRQCD parameters were obtained from comparison [278] to unpolarized high- p_t quarkonium production at the Tevatron [279]. The calculations describe these data rather well but fail to explain the polarization data [280]. The NRQCD results are also shown in table 1.1.

Table 1.1. The total quarkonium cross sections per nucleon calculated in the CEM and NRQCD models for Pb–Pb collisions at $\sqrt{s} = 5.5$ TeV with and without EKS98 shadowing. No nuclear effects, neither absorption nor suppression, are included so that the results are scaled directly up from pp collisions at the same energy. The direct production cross sections are given unless otherwise indicated. The total J/ψ and Υ cross sections include feed-down from the higher resonances. The χ_c cross section to J/ψ is given while the sums of the 1P and 2P χ_b cross sections, without decays to any of the S states, are given.

Resonance	CEM		NRQCD	
	σ (μb)	σ_{EKS98} (μb)	σ (μb)	σ_{EKS98} (μb)
Total J/ψ	30.5	18.9	83.1	50.7
J/ψ	19.0	11.7	48.1	28.6
ψ'	4.3	2.6	13.6	8.2
$\chi_c \rightarrow J/\psi$	9.1	5.6	27.6	17.7
Total Υ	0.36	0.28	0.64	0.56
Υ	0.19	0.15	0.28	0.21
Υ'	0.12	0.09	0.015	0.012
Υ''	0.072	0.056	0.012	0.011
$\sum_J \chi_{b_j}(1\text{P})$	0.39	0.32	1.35	1.06
$\sum_J \chi_{b_j}(2\text{P})$	0.30	0.24	1.44	1.14

To obtain the quarkonium yields at central collisions ($b = 0$) for a one-month, 10^6 s, run at the LHC, the cross sections in table 1.1 should be multiplied by the number of nucleon collisions, $\sigma_{\text{NN}}^{\text{inel}} T_{\text{PbPb}}(0) = 1824$ for $\sigma_{\text{NN}}^{\text{inel}} = 60$ mb and by the integrated Pb–Pb luminosity, $\mathcal{L}_{\text{int}}^{\text{PbPb}} = 10^3 \mu\text{b}^{-1}$. For other centralities, see [281]. There are thus $\simeq 10^8$ J/ψ and $\simeq 5 \times 10^5$ Υ resonances in a one-month run. These rates are sufficiently copious for high-statistics measurements at the LHC before suppression is accounted for. Therefore, quarkonium suppression should be measurable to fairly high accuracy, even for the Υ family.

Even though the basic idea of quarkonium suppression by colour screening has been known for some time and the J/ψ suppression pattern studied in detail at the SPS, a better understanding of quarkonium break-up has recently emerged owing to lattice QCD [275, 282]. New potential model studies based on lattice QCD with dynamical quarks show that ψ' , χ_c , $\chi_b(2\text{P})$, and Υ'' suppression are not necessarily characteristic of quark–gluon plasma formation but of the decrease of the string tension as chiral-symmetry restoration is approached. Thus suppression of these states does not depend on deconfinement. The more tightly bound quarkonium states with lower masses, the J/ψ , Υ , $\chi_b(1\text{P})$, and the Υ' , survive as T_c is approached and need a deconfined medium for dissociation. The J/ψ , Υ' and $\chi_b(1\text{P})$ should dissociate already at $T \sim 1.1T_c$. The most tightly bound quarkonium state, the Υ , will not be suppressed below $T \sim 2.3T_c$. According to estimates of the initial temperatures at the LHC, Υ suppression should indeed be possible [270].

In addition to colour screening, the quarkonium states will also be subject to nuclear absorption as well as possible breakup by secondary hadrons. These effects were studied extensively at SPS energies [68, 283–291]. However, at the much higher energies of the LHC, the J/ψ is expected to be formed far away from the nucleons, thereby reducing absorption effects.

At the LHC, the initial nucleon–nucleon collisions may not be the only source of quarkonium production. Regeneration of quarkonium, either in the plasma phase [259–266] or in the hadron phase [292, 293], could counter the effects of suppression, ultimately leading to enhanced quarkonium production. In the plasma phase, there are two basic approaches: statistical and dynamical coalescence. Both of these approaches depend on being able to

measure the quarkonium rate relative to total $Q\bar{Q}$ production. The first calculations in the statistical approach assumed an equilibrated fireball in a grand canonical ensemble [259]. This approach could be reasonable at the high energies of the LHC where the number of produced $c\bar{c}$ pairs is large, but at lower energies charm conservation is required since a $c\bar{c}$ pair is not produced in every event. More recent calculations assumed a canonical ensemble only for charm production [260, 263–265]. A factor of 20 J/ψ enhancement is predicted at the LHC [260]. The dynamical coalescence model assumes that some of the produced $Q\bar{Q}$ pairs, even those that otherwise would not do so, can also form quarkonium. The model includes the rapidity differences $|\Delta y|$ between the Q and \bar{Q} and shows that the larger the rapidity difference, the smaller the enhancement. Assuming 200 $c\bar{c}$ pairs at LHC and $|\Delta y| = 1$, the ratio $\langle J/\psi \rangle / (\text{charm})$ increases by a factor of ~ 60 in central collisions [266].

Much smaller enhancements are predicted for secondary quarkonium production in the hadron gas, particularly for the J/ψ where the additional production is either small (between 20 and 60%) [292] or about a factor of 2 [293]. Larger enhancements may be expected for the ψ' [292]. The predictions depend strongly on the scattering cross section of the J/ψ with π and ρ , typically not more than 1–2 mb [288].

It is worthwhile noting that secondary production will be at lower centre-of-mass energies than the initial nucleon–nucleon collisions. Thus the production kinematics will be different, leading to narrower rapidity and p_t distributions. Secondary quarkonium could be separated from the primary quarkonium, subject to suppression, by appropriate kinematic cuts. Such cuts will also be useful for separating initial J/ψ 's from those produced in B-meson decays.

These secondary production models should be testable already at RHIC where enhancements by factors of 2–3 are expected from coalescence [262, 265]. Hard scatterings of produced particles are related to the idea of crosstalk between unrelated interactions [294]. Important crosstalk effects were predicted in e^+e^- collisions at LEP [294] but were not observed. If secondary quarkonium production is found, it would indicate the relevance of such effects.

1.4. Proton–proton physics in ALICE

ALICE has several features that make it an important contributor to proton–proton physics at the LHC. Its design allows particle identification over a broad momentum range, powerful tracking with good resolution from 100 MeV c^{-1} to 100 GeV c^{-1} , and excellent determination of secondary vertices. These, combined with a low material thickness and a low magnetic field, will provide unique information about low- p_t phenomena in pp collisions at the LHC [295]. Such studies will in particular help to understand the underlying event and minimum-bias event pile-up properties. This is of interest since the latter constitute a major part of the background in searches for rare high- p_t processes in the dedicated pp experiments ATLAS and CMS. In the following subsections, we review the importance of studying proton–proton collisions with ALICE, both as a benchmark for the understanding of the heavy-ion collisions and as a means to explore proton–proton physics in a new energy domain. Obviously, entering a new energy domain has the attraction of the unknown, i.e. a discovery potential which is certainly not exhausted by the following list.

1.4.1. Proton–proton measurements as benchmark for heavy-ion physics. Most of the heavy-ion observables reviewed in section 1.3 require pp measurements of the same observables for comparison. This is important in order to identify the genuine collective effects in A–A collisions and to separate them from phenomena present already in pp collisions. A non-exhaustive list of observables to be studied for such purposes is presented below.

- Particle multiplicities: differences in particle multiplicities between pp and A–A are related to the features of parton distributions in the nucleon with respect to those in nuclei (shadowing) and to the onset of saturation phenomena occurring at small- x as discussed in section 1.3.1.
- Jet fragmentation functions: model calculations of medium-induced parton-energy loss predict a modification (softening) of the jet fragmentation functions as detailed in section 1.3.5.
- Slopes of transverse-mass distributions: the comparison of slopes in A–A collisions with those in pp allows one to determine the collective effects such as transverse flow present in A–A and absent in pp, as described in section 1.3.2.
- Particle yields and ratios: particle ratios are indicative of the chemical equilibration achieved in A–A collisions and should be compared to those in pp collisions, see section 1.3.2.
- Ratios of momentum spectra: the ratios of transverse momentum spectra at sufficiently high momenta allow one to discriminate between the different partonic-energy losses of quarks and gluons discussed in section 1.3.5.
- Strangeness enhancement: strange particle production exhibits a very regular behaviour in pp collisions between 10 and 1800 GeV, with an almost constant ratio between newly produced s and u quarks. On the other hand, a strangeness enhancement is observed in heavy-ion collisions at rather low centre-of-mass energies between 2 and 10 GeV. In particular, the K^+/π^+ ratio becomes more than twice as large as in pp collisions and then decreases again towards RHIC energies [296–298]. Therefore, the comparison of strangeness production in A–A and pp collisions at the LHC at comparable centre-of-mass energies per nucleon pair is particularly interesting. Changes in this ratio are indicative of new production mechanisms, as provided, for example, by new collective effects or by the significant contribution of jet fragments to total multiplicity.
- Heavy-quark and quarkonium production cross sections: the signals of possible suppression or enhancement of heavy-quarkonium production, as well as parton-energy losses, have to be evaluated with respect to the pp yields measured in the same experiment. In addition, these yields are not well established and must be determined more precisely.
- Dilepton spectra: dilepton production from resonance decays yields information on in-medium modifications in A–A collisions. The determination of the details of the effect relies on comparisons to smaller systems and to pp collisions.
- Photon spectra: the pp photon-energy spectrum is needed to calibrate photon production in order to estimate the background to the thermal photon production in heavy-ion collisions. Reference values for the γ -jet cross sections in pp collisions are also important.

For the observables mentioned above the dominant error is often due to the systematics. To minimize the systematic errors from comparison with the baseline measurements, it is mandatory that the observables in A–A and pp collisions be measured in the same detector set-up.

1.4.2. Specific aspects of proton–proton physics in ALICE. In addition to the benchmark role emphasized in the previous subsection, the study of pp collisions in ALICE addresses some genuinely important aspects of pp physics. It includes, in particular, the exploration of a novel range of energies and of Bjorken- x values accessible at the LHC, see figure 1.9. More generally, the ALICE pp programme aims at studying non-perturbative strong-coupling phenomena related to confinement and hadronic structure. The cross sections relevant for this study range from 100 mb for the total cross section, to 60 mb for non-diffractive inelastic processes, 12 mb for diffractive processes, and down to less than 1 mb for bottom-quark production, see figure 1.9(a). The main contribution will be in the low transverse-momentum

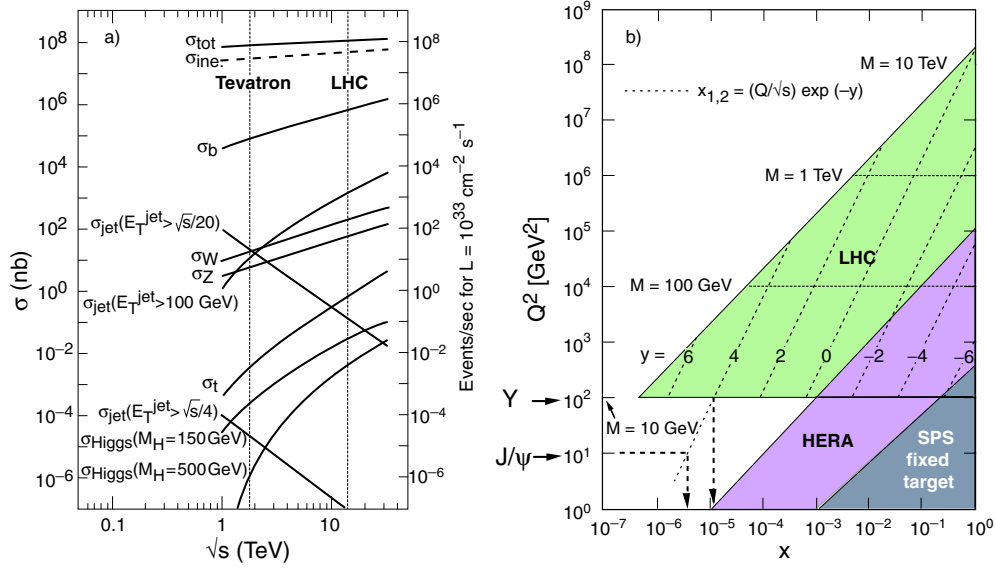


Figure 1.9. (a) Proton–proton cross sections (left-hand scale) and proton–proton collision rates for a luminosity of $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (right-hand scale) as a function of centre-of-mass energy; (b) Parton kinematic domain in Q^2 versus x for pp collisions at the LHC, compared to HERA and fixed-target experiments. The regions reached by the ALICE Υ and J/ψ measurements are also shown.

domain for which the ALICE detector was optimized. We review here some specific issues in pp physics where ALICE is unique or at least competitive with other LHC experiments.

Particle multiplicity. A simple scaling law for the \sqrt{s} dependence was proposed by Feynman who predicted a logarithmic increase of the average particle multiplicity with \sqrt{s} . However, as already shown in section 1.3.1, the best fit to the pp and $p\bar{p}$ data is given by a quadratic polynomial in $\ln s$. This non-linear dependence of N_{ch} on $\ln s$ suggests that the Feynman scaling is only approximately valid.

The general behaviour of multiplicity distributions in pp collisions in full phase space is uncertain. This makes it inaccurate to extrapolate to higher energies or to full phase space for distributions measured in limited rapidity intervals only.

The measurement of the multiplicity distributions and their comparison to KNO scaling and to various QCD-inspired multi-modal distributions, as well as the search for substructures in a model-independent way, are of relevance for distinguishing between relative contributions of soft and semihard components in minimum-bias pp interactions [299–301].

Moreover, pp collisions at LHC energies will measure very high-multiplicity events with charged-particle rapidity densities at mid-rapidity in the range 60–70, with good statistics, i.e. 10 times the average. This event class may give access to initial states where new physics such as high-density effects and saturation phenomena set in. Also, local fluctuations of multiplicity distributions in momentum space and related scaling properties (intermittent behaviour) might be a possible signature of a phase transition to QGP [302]. This makes it interesting to study such multiplicity fluctuations in pp collisions.

Particle spectra. ALICE will measure charged-particle spectra in rapidity and transverse momentum for a wide variety of particle species including π , p, \bar{p} , Λ , K, etc. The interest in

identified particles containing strangeness, charm or bottom will be discussed separately below. In addition, this allows one to study the minijet contribution, by determining the hardening of p_t spectra and various correlations between particles with higher p_t . In particular, the correlation between mean p_t and multiplicity is of interest [302]. Furthermore, the same correlation can be studied separately for identified particles, as was already done at the Tevatron [303, 304] where differences between correlations for different particles were observed.

Strangeness production. The energy dependence of strange-particle production exhibits a smooth behaviour over a large range, from tens of GeV up to 1.8 TeV in centre-of-mass energy. The recent attempts to determine the Wroblewski ratio (i.e. the suppression factor for strange quark production with respect to u and d quarks) using the statistical model of hadronization have given a value of about 0.25, independent of energy [305]. On the other hand, the correlation between the mean kaon transverse momentum and the charged-particle multiplicity observed at the Tevatron [303, 304] is significantly stronger than that for pions and is not understood. ALICE at the LHC may access considerably increased charged-particle densities and will better control transverse-momentum measurements owing to its unique low- p_t cutoff.

Baryon-number transfer in rapidity. The rapidity distribution of baryon number in hadronic collisions is not understood. In some models, the baryon number of the incident proton is transferred to a more central rapidity region by diquark exchange. This mechanism is attenuated exponentially with the rapidity interval over which the baryon charge is moved. Alternatively, baryon number flow can be due to a purely gluonic mechanism. This is accounted for in models where the three valence quarks of the proton fragment independently but are joined by strings to a baryonic gluon field configuration at more central rapidity, the so-called string junction. There exist different model estimates for the importance of this mechanism to baryon-number transfer. If one assumes that the exchange probability is independent of the rapidity interval (as for gluons) [306, 307], then the baryon number will be transferred to the central-rapidity region dominantly by gluons. On the other hand, the assumption that the string junction–antijunction Regge trajectory has an intercept 0.5 leads to a decrease of this effect with energy [308], by a factor 5 at the LHC. The ALICE detector, with its particle-identification capabilities, is ideally suited to clarify this issue with abundant baryon statistics in several channels (p , \bar{p} , Λ , $\bar{\Lambda}$) in the central-rapidity region.

Correlations. Two-particle correlations have been traditionally studied in pp multiparticle production in order to gain insight into the dynamics of high-energy collisions via a unified description of correlations and multiplicity distributions. In particular, there is interest in measuring the energy dependence, and the dependence on particle type, of the two-particle pseudo-rapidity correlations.

Heavy-flavour production. The b-production cross section measured at the CERN SPS collider and the Fermilab Tevatron lies about a factor 2 above the predictions in both cases. Also, data taken at HERA and LEP require better understanding. Data from the LHC may help in clarifying these issues. In particular, in comparison to other high-energy experiments, ALICE can measure heavy-flavour production down to very low p_t because of its unique low transverse-momentum cutoff for particle detection. This can be achieved by using inclusive large impact-parameter lepton detection, and by reconstructing exclusive charm-meson decays at relatively low p_t . As a result, the measurement of the total heavy-flavour cross section will require a smaller extrapolation and will thus show improved precision.

Jet studies. The measurement of jet production in pp collisions is an important benchmark for understanding nucleus–nucleus collisions, as discussed in section 1.3.5. In addition, the ALICE experimental programme aims at characterizing events with several jets at relatively low p_t . Observables of interest include:

- The semihard cross sections, measured by counting all events with at least one jet produced above some given E_t . This can be related to the probability of not having any hard interaction at all in an inelastic event.
- The relative rates of production of 1, 2 and 3 jets as a function of the lower E_t cutoff.
- The measurement of double-parton collisions and their distinction from the leading QCD $2 \rightarrow 4$ process. This gives a way to characterize two-body parton correlations in the nucleon projectile [309, 310].

It is also expected that ALICE will be able to study jet fragmentation in a unique way thanks to its ability to identify particles and measure their properties in a very high density environment which is relevant to jet topologies.

Photon production. The pp aspects of photon production are discussed in section 1.3.6.

Diffraction physics. Interest in diffractive physics ranges from understanding Pomeron exchange within Regge theory to small- x phenomena [311–314]. Observables at the LHC which could improve our understanding of diffractive physics include the study of the elastic and total proton–proton cross section, multiparticle production, and many others. Here, we restrict the discussion to observables accessible to ALICE which has coverage in the central pseudo-rapidity region ($|\eta| < 1.5$), supplemented in the forward pseudo-rapidity region by a forward muon arm (up to $\eta = 4$), forward multiplicity detector (up to $\eta = 5.1$), and a small so-called zero-degree calorimeter. ALICE should be able to observe central-region events with large-rapidity gaps as well as very low- x phenomena (down to 10^{-6}) such as those accessible by Drell–Yan muon pair production. The single diffraction-dissociation cross section is predicted to increase slowly (proportional to $\ln s$) but to reach large values (11–13 mb [315]) at LHC energy, while the double diffraction dissociation is predicted to be of the same magnitude. Therefore, diffractive processes will have sizeable effects. Investigation of the structure of the final hadronic state (with particle identification) produced in diffractive processes can provide important information on the mechanism of high-energy hadronic interactions. The interesting class of hard diffractive processes, however, in which a hard interaction takes place in the system of produced hadrons, is expected to show up mainly at far forward rapidity. Its study would require an upgrade of the ALICE detector.

Double-parton collisions. Increasing the centre-of-mass energy increases the parton fluxes in pp collisions. Therefore, at very high energies, multiple-parton collisions become increasingly important. The first measurement of double-parton collisions was performed at the Tevatron [316, 317] selecting final states with three jets and a photon. These results indicate non-trivial correlations of the proton structure in transverse space [309, 310]. The structure of the proton appears to be much richer than the independent superposition of single-parton distribution functions accessible by deep-inelastic scattering.

Estimates for LHC suggest a significant cross section for double-parton collisions into final states with four jets, and three jets and a photon, even in the case of charm and bottom heavy-flavour jets. The more conventional leading QCD $2 \rightarrow 4$ processes will have a sub-leading

contribution at LHC. For instance, double-parton scatterings dominate over the leading QCD $2 \rightarrow 4$ process by possibly more than an order of magnitude for the case of $c\bar{c}c\bar{c}$ final states, in the low transverse-momentum region. Moreover, CDF data [316, 317] point to a peculiar pattern of events where the dispersion in the number of partonic interactions is high, i.e. one observes strong fluctuations in the number of produced jets. The large values for the double-parton cross sections at LHC may allow one to identify also triple- and quadruple-parton collision processes.

In addition, the study of multiple-parton collisions in pp interactions is interesting because unitarity constraints relate its cross section to the total pp cross section.

1.5. Proton–nucleus physics in ALICE

1.5.1. The motivation for studying pA collisions in ALICE. Most generally, the uncertainties in the initial conditions for heavy-ion collisions at the LHC translate into significant ambiguities for separating initial- from final-state medium effects. A main motivation for a nucleon–nucleus run is to alter the relative importance of initial- and final-state medium effects, thus allowing for a better separation between them. pA (or dA) collisions will provide, together with the pp results, a compulsory benchmark for the interpretation of the A–A data.

In more detail, the observables measured in pA collisions are mostly sensitive to final-state rescattering and energy loss of partons traversing the nuclear medium. This is important for comparing the radiative energy loss of a fast parton in hot and in cold nuclear matter. Also, the interpretation of open-charm and charmonium data requires input from the measurement of the production and suppression mechanisms in pA. In addition, pA collisions allow one to test a key assumption of the interaction of nucleons with nuclei, namely that in the hadronic multiple scattering scheme, the incident proton interacts with one target nucleon at a time. This is the main assumption of the Glauber model which may be questioned in view of the large inelastic cross section at the LHC. In summary, the interpretation of A–A data requires both the study of pp collisions to measure the properties of the nucleon–nucleon interaction and the study of pA collisions to understand the interaction of hadrons with cold nuclear matter.

The study of pA collisions at LHC provides opportunities to improve our knowledge of nuclear parton distribution functions and to go beyond the knowledge of single-parton distribution functions with the measurement of double-parton collisions, as discussed below.

1.5.2. Nucleon–nucleus collisions and parton distribution functions. Hard processes such as heavy-flavour production and high- p_t particle production will play a central role in the physics of ALICE. In the QCD-improved parton model, they are described by convoluting the perturbatively calculated partonic cross sections with parton distribution functions (PDFs) for the description of the initial state and with fragmentation functions for the description of the hadronization process in the final state. Precise knowledge of PDFs is thus essential for a quantitative understanding of these cross sections. Perturbative QCD predicts only the scale dependence of PDFs but not their (non-perturbative) input values which have to be extracted from experimental data. In practice, the initial distributions for proton PDFs at Q_0^2 are evolved using the DGLAP equations to larger Q^2 and fitted to available data. The main input comes from deep-inelastic scattering (DIS) data, in particular from HERA data for the small- x region. In this way, several groups (MRST [318], CTEQ [319], GRV [320]) have developed parametrizations.

PDFs for nuclei differ from those of free protons on account of several nuclear effects, namely shadowing at small values of x , anti-shadowing at intermediate x , and the EMC effect and Fermi motion at large x . The definition of a nuclear PDF set according to the same




Figure 1.10. Ratios of gluon distribution functions for Pb relative to the gluon distribution for proton using different models at $Q^2 = 5 \text{ GeV}^2$, see text for more details.

procedure as for protons is complicated due to the lack of experimental data in the perturbative region ($Q^2 \gtrsim 1 \text{ GeV}^2$) at small x . Available experimental information comes mainly from DIS and Drell–Yan data in the range $0.005 \lesssim x < 1$ for $Q^2 \gtrsim 1 \text{ GeV}^2$ [321]. Up to now, only two analyses of this type exist: EKS98 [251] (see also [322]) and HKM [323]. They differ mainly in the sets of experimental data used; Drell–Yan [324] and Q^2 -dependent DIS data [325] were used in EKS98 but not in HKM. PDFs were also computed in several models which tend to disagree where no experimental constraints are available. This is particularly pronounced for gluon PDFs which are only weakly constrained since they do not enter the measured F_2 structure function at leading order. The nuclear data on processes mainly driven by gluons, such as open heavy flavour, are very poor. The situation is summarized in figure 1.10, where the ratio of the gluon distribution in a Pb nucleus over the gluon distribution in a proton (at $Q^2 = 5 \text{ GeV}^2$) is computed by different methods: ‘saturation plus Glauber’ (Armesto [326]); perturbative QCD (Sarcevic [327]); from diffraction data using the Gribov model (Frankfurt [328]); the new parametrization included in HIJING (new HIJING [329]); and the only two presently available DGLAP analyses with nuclei (EKS98 [251] and HKM [323]).

Despite the large uncertainties in the gluon PDF, some constraints [330, 331] exist from experiment. The Q^2 -evolution at small x is mainly driven by gluons in the DGLAP equations. Using this fact, it is possible to constraint the gluon PDF by comparison with the existing experimental data on Q^2 -dependence of the F_2 ratios of Sn over C. The result, see [331], is that, in the framework of leading-twist leading-order DGLAP evolution, gluon shadowing is constrained to be of the same order as that of quarks (the one directly measured in F_2). In particular, very strong gluon shadowing such as the one in the new HIJING parametrization, is ruled out in this framework. This constraint is valid for $x \gtrsim 0.01$ (the range of NMC data). For smaller values of x more data in DIS with nuclei and/or pA collisions at very high energy are needed.

Measurements of pA collisions allow us to reduce the uncertainties in the nuclear parton distribution functions and in particular in the gluon structure function. An experimental handle is provided by the quark and gluon initiated processes $q\bar{q} \rightarrow Z^0, W^\pm$ and $gg \rightarrow J/\psi, \Upsilon$. In principle, measuring the yields in pPb relative to pp collisions at 9 TeV allows us to explore values of $x \approx M/\sqrt{s} \simeq 10^{-2}$ for Z^0 and W^\pm as well as $x \simeq 10^{-3}$ for Υ and $x \simeq 10^{-4}$ for J/ψ production at $y = 0$. Moreover, the measurement of prompt photons, produced by quark–gluon fusion, gives access to the structure functions at $x \approx 2p_t/\sqrt{s}$. This probe is particularly interesting because it is only sensitive to the initial conditions (i.e. the parton densities) and not to the surrounding nuclear environment.

1.5.3. Double-parton collisions in proton–nucleus interactions. Multi-parton interactions explore the hadron at different space points simultaneously [332]. Hadron–nucleus collisions [333] substantially enhance multi-parton interactions because of the increased flux of incoming partons. Moreover, the rate of multi-parton processes can be varied with the nuclear targets at fixed centre-of-mass energy and with the same cuts in the produced final state. The simplest case is double-parton scattering. Under the simplifying assumption that non-additive corrections to the nuclear parton distributions are negligible, the inclusive cross section for two identical parton processes is

$$\sigma_D^A = \frac{1}{2} \int \Gamma_N(x_1, x_2; r) \hat{\sigma}(x_1, x'_1) \hat{\sigma}(x_2, x'_2) \Gamma_A(x'_1, x'_2; r) dx_1 dx'_1 dx_2 dx'_2 d^2r. \quad (1.27)$$

The double-parton distributions $\Gamma(x, x'; r)$ for the nucleon (N) and nucleus (A) depend on the momentum fractions x, x' of the two partons, on their transverse distance r , on the parton species, and on their virtualities. For simplicity the latter dependence is not denoted explicitly. In pA interactions, the two partons entering $\Gamma_A(x, x'; r)$ originate from either the same or different nucleons in the nucleus. The cross section σ_D^A is correspondingly expressed as the sum of two terms, σ_1^D and σ_2^D . The component of the nuclear double-parton distribution originating from a single nucleon differs from the double-parton scattering cross section σ_D in pp only by the enhanced normalization. This normalization is given by the nuclear thickness function $T(b)$ integrated over impact parameter b ,

$$\sigma_1^D = \sigma_D \int d^2b T(b) = A \sigma_D. \quad (1.28)$$

For the second case, where two different nucleons are involved in a double-collision process, the corresponding cross section is

$$\sigma_2^D = \frac{1}{2} \int G_N(x_1, x_2) \hat{\sigma}(x_1, x'_1) \hat{\sigma}(x_2, x'_2) G_N(x'_1) G_N(x'_2) dx_1 dx'_1 dx_2 dx'_2 \int d^2b T^2(b). \quad (1.29)$$

This equation is derived assuming factorization of the hard interactions in the double-parton scattering process and neglecting the hadron scale as compared to the nuclear scale. The nuclear parton flux, at a given transverse coordinate b , is represented by the product $T(b)G_N(x')$, where $G_N(x')$ is the nuclear parton distribution divided by the atomic mass number A . In equation (1.29) the nuclear parton flux is squared since the interaction takes place with two different target nucleons. The difference between the transverse coordinates of the two nuclear target partons, $b + r/2$ and $b - r/2$, is neglected, since the nuclear density does not change in the transverse scale $\langle r \rangle \approx R_N \ll R_A$. With this approximation, the integration over r in equation (1.27) involves only Γ_N , and the cross section depends on the dimensionless quantity $G_N(x_1, x_2) = \int d^2r \Gamma_N(x_1, x_2; r)$. In contrast to σ_1^D , no scale factor related to the nucleon transverse size enters the σ_2^D term. The correct dimension for σ_2^D is provided by the nuclear thickness function which gives the normalization to the nuclear parton flux.

Remarkably, in pA interactions the presence of a large transverse scale, the nuclear radius, allows one to separate the longitudinal- and transverse-correlation effects in the hadron structure. The two components of σ_D^A , σ_1^D and σ_2^D , can be separated on account of their different dependence on the atomic mass number of the target. In the kinematical regime where non-additive contributions to the nuclear parton distributions are negligible, the A -dependence of the two terms does not change with the momentum fractions and with the virtuality scale of the partonic interactions and one has $\sigma_1^D \sim A$ and $\sigma_2^D \sim A^{1.5}$. Here, the nuclear surface effects lead to a faster dependence of $\int T^2(b)d^2b$ on A than the naive expectation $\sim A^{4/3}$. The A -dependence of σ_1^D is the same as that of the single-scattering inclusive cross section and, in fact, characterizes all parton processes which can be treated in impulse approximation. In particular, this A -dependence allows one to disentangle contributions from $2 \rightarrow 4$ parton processes which constitute a background to the double-scattering term. The identical A -dependence of σ_1^D and other processes such as those producing two large p_t partons, allow for the separation of σ_1^D and σ_2^D even in the kinematical regime in which non-additive effects to the parton distributions are important. Moreover, while non-additive effects to the nuclear parton distributions may reduce the nuclear densities and the cross section by a factor $\simeq 0.7$ at small x , the enhancement factor $A^{1.5}$ of σ_2^D due to interactions with different target nucleons, is a much larger effect. Hence, the σ_2^D term may constitute 70% of the cross section σ_D^A in a collision with a heavy nucleus.

The comparison of the two terms, σ_1^D and σ_2^D , will allow one to determine the average transverse distance between two partons in the hadron structure as a function of their momentum fractions x_1 and x_2 . This is a check of the factorization approximation to multiple-parton collisions, which is implicitly assumed in all present considerations. pA interactions thus provide a novel point of reference, namely the A -dependence, to gain insight into the dynamics of hadronic reaction.

1.6. Physics of ultra-peripheral heavy-ion collisions

The fast-moving nuclei in relativistic heavy-ion collisions are surrounded by strong electromagnetic fields. This is due to the coherent action of all the protons, as well as to the Lorentz contraction $\gamma \simeq 2900$ of the electromagnetic field at high energies. These strong fields are a useful source of quasi-real photons [334–336] and therefore make possible the study of photon–photon and photon–nucleus collisions (nuclear excitation as well as photon–hadron or photon–Pomeron processes) in very peripheral collisions [337] (see figure 1.11). Only collisions where the impact parameter b between the two nuclei is larger than the sum of the nuclear radii are useful for such studies. In more central collisions the high particle multiplicity of the hadronic interaction will completely mask the electromagnetic processes. One characteristic of peripheral collisions is therefore that either one or both ions leave the interaction region either intact or only weakly excited (see figure 1.11).

Photon–photon collisions The achievable luminosity for photon–photon collisions is large up to an invariant mass range of about 200 GeV. Therefore, very peripheral collisions extend the invariant-mass (photon energy) range which is currently being studied at LEP. This opens the possibility for photon–photon processes either at higher invariant masses or with a larger luminosity than what has been possible up to now. A few processes of interest are listed below:

- Meson spectroscopy of charm and bottom $C = +1$ mesons can be studied with photon–photon collisions. Also, meson and baryon pair production is accessible.




Figure 1.11. Owing to the large electromagnetic fields, relativistic heavy ions are a strong source of quasi-real photons. These can be used for photon–photon (a), and photon–nucleus (b) collisions, where even higher order processes, with more than one photon exchange, are possible (c).

- Exotic mesons like glueball or hybrids may be measured. Glueball candidates would reveal themselves since their decays to $\gamma\gamma$ are highly suppressed. At LEP, a deviation from the universal Regge parametrization has been found in the total cross section $\gamma\gamma \rightarrow$ hadrons. Although a wide coverage in rapidity is desirable, such studies are possible also with ALICE.
- The accessible invariant-mass range also allows one to search for electroweak processes like W^+W^- production or even new physics which, however, would require a luminosity upgrade.

Photon–nucleus interactions. For photon–hadron collisions, the photon has an energy range up to 500 TeV in the rest frame of the ion, extending the energy range beyond the one reached at HERA. This covers the range of low-energy nuclear structure physics up to the diffractive production of the Υ . At low energies new phenomena can be studied since the strong electromagnetic fields lead to the exchange of several photons and multiple excitations.

Photonuclear excitation. The photonuclear excitation cross section is of the order of 215 b, exceeding by far the geometrical cross section due to the direct nuclear overlap. It is responsible for about half of the total beam loss in the collisions, see [338] and section 2, along with the bound-free pair production, another electromagnetic process.

Since the exchange of more than one photon is very likely, it is even possible that both colliding nuclei may disintegrate in a single event. This mutual electromagnetic dissociation is currently used at RHIC to measure the beam luminosity by measuring high-energy neutrons in each ZDC emitted in the forward–backward directions [339].

Even higher-order processes contribute substantially to the mutual excitations [340, 341]. The excitation of the double and probably also the triple giant dipole resonance should be observable in the ALICE ZDC [341].

Photon–hadron interaction. The diffractive production of the ρ -meson has been observed at the STAR detector at RHIC [342]. In ALICE, one can measure J/ψ and even Υ production with much larger rates than those at HERA. This is of special interest due to the transition from the non-perturbative to the perturbative regime of the Pomeron with increasing quark masses. In addition, photo-production on a nuclear target can be studied. Especially J/Ψ production allows the study of colour transparency in nuclei.

Semi-coherent processes. These are another possibility for photon–hadron processes, where one ion serves as the source of the quasi-real photon and the other ion as the target. Of course the target ion will in general be broken up in this case. These processes may allow access to the partonic structure of the heavy ions at small Bjorken x . Photon–gluon fusion, especially with the production of heavy quarks, has been proposed as a way to study the gluon distribution

function inside the nucleus which is expected to saturate at sufficiently small x . Incoherent vector meson production, which is proportional to the square of the gluon distribution function and, therefore, much more sensitive to medium modification, has been studied recently as well. Another option, which is currently under study, is the scattering of a single lepton (e or μ) under large angle. Together with the measurement of a jet in the opposite direction, this can be used to study deep-inelastic lepton scattering within the framework of the ‘equivalent muon approximation’. This also allows the quark distribution function to be measured.

Purely diffractive processes. These are another type of very peripheral collisions. They are an interesting topic in themselves but are also a possible background for photon–photon and photon–hadron processes. Studies for specific processes have shown that double-diffractive processes are dominant for the lightest ions, as well as for protons, whereas for Pb–Pb collisions they should only represent a background. This opens the possibility to study, for example, mesons in both Pomeron–Pomeron and photon–photon processes in order to assess their gluon content.

In summary, very peripheral collisions and their accompanying photon–photon and photon–hadron processes are always present in heavy-ion collisions. In order to study them one has to find a way to tag and record them within the existing experimental set-up. The task is not an easy one since the characteristics of these collisions are quite different from those of central collisions. Although sometimes called ‘silent collisions’, due to the rather small multiplicities, the range of physics that can be covered makes them of sufficient interest for study at the LHC [343]. In any case, the soft-photon cloud produced in the LHC vacuum pipe at each beam crossing by these processes contributes to background and has to be understood.

1.7. Contribution of ALICE to cosmic-ray physics

Cosmic rays are particles originating outside the Earth atmosphere either in our galaxy or, for the very high-energy cosmic rays, in other galaxies. The energy of cosmic rays which reach our planet spans many decades from a few GeV per particle up to at least 10^{20} eV. Whereas at low energies (up to TeV) these particles are mainly protons with a small percentage of He nuclei the nuclear contribution starts to increase at higher energies. Helium nuclei are always present. Above 10^{14} eV the C–N–O group shows an increasing contribution; at even higher energies Mg, Si, and nuclei up to Fe appear.

The main observables in cosmic-ray physics are the energy spectrum (figure 1.12) and the evolution of the composition of the primaries with energy. These observables convey information on the process of particle acceleration at the sources and on their propagation through the interstellar medium. Measurements up to energies of $\simeq 10^{14}$ eV are called ‘direct’, since they can be directly performed by balloon-borne detectors in the atmosphere or by satellites. While the energy spectrum is well established, the composition of cosmic rays above energies of 10^{14} eV is still unsolved and is the subject of much active research. The ‘knee’ of the cosmic-ray energy spectrum at $\simeq 3 \times 10^{15}$ eV where the slope of the spectrum steepens (figure 1.12), is of particular interest. It is an open question whether this knee has an astrophysical origin or whether it is due to a change in the properties of the hadronic interactions.

At higher energies, in the region of the knee and above, only ‘indirect’ detections are possible because of the extremely low flux. Ground based arrays, so-called ‘extensive air-shower arrays’, which are located at the Earth surface or underground, detect the shower of particles created by the interaction of the primary cosmic ray with a nucleus of the Earth atmosphere. The large collecting surface of these arrays allows the detection of a number of events of extremely high energies. However, one of the basic difficulties of these detectors




Figure 1.12. Cosmic-ray flux as a function of energy, showing the knee structure appearing between Tevatron and LHC energies.

in the reconstruction of the primary energy and mass originates from the required detailed understanding of the interaction mechanism of the showering particles and the modelling of the shower development and its fluctuations in the 10 interaction length thick atmosphere. In fact, new phenomena in very forward high-energy hadronic interactions, such as coherent pion production, disoriented chiral-condensate states or heavy-flavour production can significantly influence the hadronic cascade and hence the observables at the ground level. This may well be the cause for the conflicting results about the particle composition among various experiments. Particle production, both at high energies and in the forward direction, can today only be estimated by model based extrapolation of accelerator data. The interpretation of the cosmic-ray data, in particular the identification of the primary cosmic ray, depends crucially on these models. Indeed, there are presently no accelerator data for particle production at very small forward angles and in the relevant energy region around the cosmic-ray knee. While it may be that some of the models will converge on a common interpretation, they may still




Figure 1.13. The ALEPH event with the highest multiplicity (about 150 tracks, in half of the TPC) in two different views: transverse (left-hand side), and along the shower axis (right-hand side). Reconstructed tracks are shown as lines.

be wrong. Via a better understanding of hadronic interactions of proton and nucleons at very high energies, ALICE may contribute to the clarification of this question and constrain the models significantly. In the past years many interaction models such as VENUS, QGSJET, DPMJET, SIBYLL and others have been extensively used by the cosmic-ray community (for a review see [344]). The general approach in the creation of a model is to tune the free parameters comparing the simulations with experimental data coming from $p\bar{p}$, ep and heavy-ion collisions at the maximum energy available from the accelerator experiments. At present, data exist up to energies of $\sim 10^{15}$ eV. ALICE will extend the range in the centre-of-mass energy up to $\sqrt{s} = 14$ TeV which corresponds to a proton of $E \simeq 10^{17}$ eV interacting with a fixed target. This lies above the knee of the energy spectrum of cosmic rays, see figure 1.12. The LHC will allow to extend the study of the A -dependence in A - A interactions to higher energies. Of particular interest is the study of nitrogen–nitrogen interactions, since nitrogen and oxygen are the most abundant nuclei in the atmosphere and since nitrogen nuclei are also present in high-energy ($E \geq 10^{15}$ eV) cosmic rays.

ALICE can also be used as an active cosmic-ray experiment. In the past, the use of large underground high-energy physics experiments for cosmic ray studies has been suggested by several groups [345]. The L3 experiment has established a cosmic-ray experimental programme, L3–Cosmics, with the principal aim to measure the inclusive cosmic-ray muon spectrum between 20 and 2000 GeV, in the context of the current interest in neutrino oscillations and the pressure for a more precise calculation of the muon-neutrino spectrum. The utility of LEP experiments based data for cosmic-ray studies has also been explored by the CosmoLep group [346] in an analysis of multi-muon events recorded by the ALEPH experiment, using triggers during normal e^+e^- data taking. Intriguing high-multiplicity events have been observed with the ALEPH detector [347]. The five highest-multiplicity events are presently unexplained and will require further studies with larger statistics and bigger detectors. The muon bundle with the highest density of about 150 muons within a detection area of 8 m^2 is shown in figure 1.13.

The underground location of the ALICE experiment with the 29 m overburden is ideal for muon based experiments: the electromagnetic and hadronic components of the air showers are fully absorbed by the overburden and the muon momentum cutoff is around 15 GeV. This is in contrast to deep underground experiment, such as MACRO [348] where the momentum




Figure 1.14. CORSIKA Monte Carlo simulations of the muon density at ground level for proton and iron induced showers of various energies as a function of the radial distance from the shower centre.

cutoff is of the order of TeV, as well as to surface experiments such as KASCADE [349] with a muon cutoff of the order of 1 GeV. The ALICE TPC offers the opportunity to have magnetic analysis over a large volume which will provide a precise determination of the muon directions as well as their momenta up to order of TeV. The fine granularity of the ALICE TPC permits the measurement of a high density of muon tracks, the so-called muon bundles. These muon bundles contain complementary information compared to the usual measurement of the more calorimetric electromagnetic air-shower component. They allow to study the characteristics of the first interactions at the top of the atmosphere. Muons, because they originate from the decay of hadrons, and do not multiply but only slowly lose energy by ionization as they transverse the atmosphere, are particularly useful in this regard. Furthermore, muons of different energies provide somewhat different traces of the shower development. The optimal muon energy is around 50–100 GeV. At these energies the parent pions have a decay mean free path comparable to their interaction mean free path at the top of the atmosphere and therefore a sizeable fraction of the muons are born at this height reflecting some properties of the primary interactions. An important aspect is the size of the muon bundles. The radial muon distribution is shown in figure 1.14. Compared to protons, iron induced muon bundles exhibit a flatter radial distribution and contain twice as many muons. The ability of the ALICE TPC to measure the momenta, even at large muon densities, and to study the radial distribution of muon bundles with large statistics will be a powerful tool to disentangle various uncertainties inherent in the models and in the particle composition and perhaps to make unforeseen discoveries. The size of the detector area, the spatial resolution and the pattern recognition in complicated high-density events play a decisive role in the study of multi-muon bundles. With its excellent angular resolution, in particular in the presence of many parallel muons, the ALICE TPC can also be used as a telescope to look for point sources in our galaxy and possibly beyond.

It should be stressed that in order to efficiently trace back the cosmic-ray air-shower development and to probe details of the primary cosmic-ray composition or the forward

production processes, the maximum set of observables should be used. In principle, a precise measurement of the muon bundles would be sufficient to determine the primary energy and composition if the interaction characteristics would be known. But in combination with an independent determination of the shower characteristics at the ground level, the identification of the primaries is drastically improved. This has been demonstrated in the CORAL proposal [350]. Compared to protons, heavy-nuclei interactions with air molecules take place higher up in the atmosphere and produce more hadrons with lower energies. As a consequence, the electromagnetic component develops faster in the upper atmosphere resulting in more hadron decays into muons and in a larger electromagnetic attenuation in the lower atmosphere. It is very significant that these two effects, the increase in muon number and decrease in electron number in heavy-nuclei-initiated showers are in opposite direction, which helps in magnifying the observable differences between showers initiated by heavy or lighter nuclei.

It is advantageous that an operating surface air-shower array ($50 \times 50 \text{ m}^2$), containing 40 counters with a typical distance between the counters of 8 m, already exists on the flat top of a building almost vertically above the ALICE collision point. In the future, this array might be extended with additional existing counters. A typical air shower of a few 10^{15} eV (the knee region) deposits most of its energy within a distance of about 30 m from the shower core. With a counter grid structure of about 10 m we should locate the shower core to better than 5 m. Knowing the shower angle from the underground muons with good precision we can extrapolate the shower core in the underground area and hence measure the radial muon distribution even if the shower core is not falling into the acceptance of the ALICE detector.

In summary, the ability to combine the information from the electromagnetic shower array on the top with precise muon measurements in the ALICE TPC, even in dense muon bundles, will provide a powerful tool for the study of high-energy cosmic-ray air showers in the energy range around the knee and above and should allow to determine at least statistically the nature of the primary cosmic rays.

2. LHC experimental conditions

In section 1 we summarized the observables relevant to the physics goals of ALICE. Whether ALICE can reach these goals depends not only on the performance of the detectors, but also on the number of events which can be collected, the number of collision systems which can be studied, and last but not least on the background conditions at the LHC. The number of collected events depends on the nature and the energy of the beam, the luminosity, and the running time. Therefore, we first describe in this section our proposed running strategy, followed by a discussion of the required luminosity for each collision system including luminosity limitations of both the detector and the LHC machine complex. Then we describe the expected background conditions at the LHC. In the last section, the ALICE luminosity monitoring and measurement is outlined.

2.1. Running strategy

A comprehensive heavy-ion programme at the LHC, like the SPS and RHIC programmes, will be based on two components: colliding the largest available nuclei at the highest possible energy and a systematic study of different collision systems (pp, pA, A–A) and of different beam energies. As the number of possible combinations of collision systems and energies is very large, continuous updating of priorities will be required as data become available in order to optimize the physics output. We have therefore divided the ALICE programme into two phases: an initial programme with priorities based on our current theoretical understanding

and the results from the SPS and RHIC, and a later stage with a number of options whose relative importance is likely to become clear only after the first LHC data have been analysed.

The LHC is expected to run essentially in the same yearly mode as the SPS, starting with several months of pp running followed at the end of each year by several weeks of heavy-ion collisions. For rate estimates, all LHC experiments use an effective time per year of 10^7 s for pp and 10^6 s for heavy-ion operation.

ALICE will take its first data with pp collisions, because the LHC will be commissioned with proton beams, but also because pp physics is an integral part of the ALICE programme. ALICE will in fact require pp running throughout its time of operation: during the initial few years longer periods to both commission the detector and to take pp physics data, and later in the programme shorter periods to start up and calibrate the detector prior to each heavy-ion period.

Pb–Pb collisions, which provide the highest energy density, are foreseen immediately after the end of the first pp run. Even if of short duration and at low luminosity, this first physics pilot run will already provide a wealth of information on global event properties and large cross section observables, as was the case during the very successful commissioning of RHIC. For low cross section observables, in particular the hard processes, which are one of the main focuses at the LHC, some further 1–2 years of Pb–Pb runs at the highest possible luminosity should provide sufficient statistics. One period of pPb running is required early on, most likely in the third year of LHC operation. This will provide reference data and will allow us to determine nuclear modifications of the nucleon structure functions (shadowing), which is necessary for the interpretation of the Pb–Pb data. The best way to study energy density dependencies is to use lower-mass ion systems, which can be selected from a list of candidates suitable for the CERN ion source. We plan to study first Ar–Ar collisions over a period of 1–2 years. This initial ALICE programme, which has been discussed and endorsed by the LHCC, is summarized below:

- Regular pp runs at $\sqrt{s} = 14$ TeV.
- Initial heavy-ion programme:
 - Pb–Pb physics pilot run;
 - 1–2 years Pb–Pb;
 - 1 year pPb-like collisions (pPb, dPb or α Pb);
 - 1–2 years Ar–Ar.

For the later phase, we have considered a number of running options, the relative importance of which will depend on the initial results. For a direct comparison of the Pb–Pb and pp data, a dedicated pp run at $\sqrt{s} = 5.5$ TeV corresponding to the nucleon–nucleon centre-of-mass energy for nominal Pb–Pb running is desirable. A more complete energy density scan would require additional intermediate-mass ion runs. To map out the A -dependence, further pA runs with different nuclei could be necessary. Additional Pb–Pb runs at lower energy would allow us to measure an energy excitation function and to connect to the RHIC results. Finally, some rare processes limited by statistics in the early runs could require additional high-energy Pb–Pb running. The list of running options for the second phase is summarized below:

- Later options, some of them depending on the outcome of the initial data analysis:
 - dedicated pp or pp-like (dd or $\alpha\alpha$) collisions at $\sqrt{s_{NN}} = 5.5$ TeV;
 - possibly another intermediate-mass A–A system (N–N, O–O, Kr–Kr or Sn–Sn). N–N and O–O are of interest for the study of cosmic rays, as N and O are the most abundant nuclei in the atmosphere;
 - possibly another pA (dA, α A) system;

- possibly lower energy Pb–Pb runs;
- further high-energy Pb–Pb runs to increase the statistics of rare events.

2.2. A–A collisions

2.2.1. Pb–Pb luminosity limits from detectors. The luminosity for Pb–Pb collisions is an important issue, because there are limitations coming from both the detector and the accelerator. We start with the limitations due to the ALICE detector. The limits are coming from the main tracking device, the Time-Projection Chamber (TPC), and the forward muon spectrometer, which are considered separately.

The TPC limits the maximum usable luminosity because of event pile-up during the $88 \mu\text{s}$ drift time. At a luminosity of $10^{27} \text{ cm}^{-2} \text{ s}^{-1}$, the pile-up probability is 76% for a hadronic interaction cross section of 8 b. However, since the average particle multiplicity amounts to only 20% of the maximum multiplicity and since only partial events overlap, the average increase of the track multiplicity due to pile-up is about 13% for central collisions. Locally the track density increases by up to 26%. It is conceivable that the TPC can be operated at luminosities above $10^{27} \text{ cm}^{-2} \text{ s}^{-1}$, in particular, if the multiplicity turns out to be low compared to the maximum of 8000 per unit of pseudo-rapidity considered for the design of ALICE. However, the gain in rate has to be weighed against the loss in tracking performance.

The luminosity limitation of the muon spectrometer comes from the maximum acceptable illumination of the trigger chambers, Resistive-Plate Chambers (RPC), of $50\text{--}100 \text{ Hz cm}^{-2}$. This limit corresponds to a maximum usable luminosity of $(2\text{--}4) \times 10^{28} \text{ cm}^{-2} \text{ s}^{-1}$. As will be shown in the following section, the limitations from the machine are stronger than those coming from the ALICE detector, which justifies the choice of detector technologies.

2.2.2. Nominal-luminosity Pb–Pb runs. For a given initial luminosity (\mathcal{L}_0), optimal operation requires a maximization of the time-averaged luminosity $\langle \mathcal{L} \rangle$ given by

$$\langle \mathcal{L} \rangle(t) = \frac{1}{t + T_f} \int_{t_{\text{set-up}}}^t dt' \mathcal{L}(t'). \quad (2.1)$$

Here T_f is the filling time and $t_{\text{set-up}}$ the experiment set-up time, i.e. the time during which the collider is operating but the experiment is not yet taking data. The high electromagnetic cross section for removing Pb ions from the beam of $\simeq 500 \text{ b}$ is the main limit for the luminosity lifetime. The lifetime depends on how many experiments are active during the ion runs. Using $\beta^* = 0.5 \text{ m}$ which requires the smallest beam intensity, the half-lives for an initial luminosity $\mathcal{L}_0 = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ are 6.7 h for one, 3.7 h for two, and 2.7 h for three experiments (see [351] but using updated cross sections from [352]).

The luminosity and time-averaged luminosity relative to $\mathcal{L}_0 = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ are shown in figure 2.1 for nominal operation with Pb–Pb collisions in one, two and three interaction regions. A filling time of 3 h and experiment set-up time of 20 min are assumed. The average luminosity has a plateau at around 8 h. At this time the average luminosities are $0.44 \mathcal{L}_0$, $0.35 \mathcal{L}_0$ and $0.29 \mathcal{L}_0$ for one, two and three experiments, respectively.

The beam position monitors (BPMs) in the LHC require a minimum charge per bunch in order to function properly. With the present design this corresponds to about one third of the initial intensity. As can be seen from figure 2.1 this limit is not reached. However, the gap between nominal operation and BPM intensity limit is very narrow.

2.2.3. Alternative Pb–Pb running scenarios. For a fixed \mathcal{L}_0 the luminosity lifetime is proportional to the beam intensity. Hence, longer lifetimes can be obtained in high- β^* high-intensity runs. The luminosity optimization is constrained by three effects: intra-beam




Figure 2.1. Luminosity, average luminosity relative to $\mathcal{L}_0 = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ and beam intensity relative to $I_0 = 6.8 \times 10^7$ ions per bunch for nominal operation with Pb–Pb collisions in one (solid line), two (dashed line) and three (dotted line) interaction regions.

scattering, longitudinal emittance growth and the maximum beam intensity. Intra-beam scattering plays a role because the emittance growth time is inversely proportional to the beam intensity. The relative longitudinal emittance growth must be less than a factor of 2 to avoid beam losses. Finally, the beam intensity is constrained by space-charge limits in the SPS.




Figure 2.2. Luminosity and average luminosity relative to \mathcal{L}_0 for β^* -tuning operation with Pb–Pb collisions in one (solid line), two (dashed line) and three (dotted line) interaction regions. This scenario assumes an intensity of 10^8 Pb ions per bunch, the highest envisaged intensity.

At the nominal initial luminosity, running at $\beta^* = 2$ m and consequently twice the nominal beam intensity, the average luminosity increases by only $\simeq 10\%$. A further, and much more important, increase of the average luminosity is possible by β^* -tuning, decreasing β^* during the run to keep the luminosity constant until the minimum β^* of 0.5 m is reached. The maximum possible bunch intensity corresponds to about 1.4 times the nominal intensity. Using β^* -tuning with this intensity could increase the average luminosity by approximately a factor of 1.6 as shown in figure 2.2.

β^* -tuning is also very important for operation with the nominal initial beam intensity. It can be shown that using β^* -tuning the same average luminosity as anticipated for an initial luminosity of $\mathcal{L}_0 = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ can be reached with a lower initial luminosity: $0.7\mathcal{L}_0$, $0.55\mathcal{L}_0$ and $0.5\mathcal{L}_0$ for one, two and three experiments, respectively [353]. This has two important advantages. First, a constant luminosity during the run makes it easier to handle space-charge effects in the TPC. Second, operation at a lower initial luminosity is safer for the machine, because the quench limit (see below) might be close to the nominal \mathcal{L}_0 .

Since the ALICE measurement of Υ production in the forward muon spectrometer is statistics limited, ALICE would take advantage of dedicated, high-luminosity Pb–Pb runs. The same scheme as described above, high-intensity and β^* -tuning, can be applied. It is




Figure 2.3. Variation of the produced energy density with collision system. The energy density has been calculated using the Bjorken formula $\varepsilon = 160 \text{ MeV fm}^{-3} A^{-2/3} dN_{\text{ch}}/dy$, with maximum charged-particle multiplicities of 6000, 1200, 230 and 6.5 for central Pb–Pb, Ar–Ar, O–O, and pp collisions, respectively. The bands show the approximate range covered by changing the impact parameter. The upper bound corresponds to central collisions and the lower bound was calculated using the minimum-bias charged-particle multiplicity. In pp collisions the energy density can be increased by up to a factor 10 compared to the average value shown on the figure selecting high-multiplicity events.

currently assumed that the maximum Pb–Pb luminosity is limited to $(0.5\text{--}1) \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ by the maximum allowed beam-pipe heating in the cold machine elements of the dispersion suppressor, the quench limit. The energy deposition in these elements comes from Pb ions lost in electromagnetic interactions [351, 354].

2.2.4. Beam energy. Short runs of a few days each at lower beam energies may be requested to study the energy as well as the Bjorken- x dependence of global hadronic event features. These runs at the lowest possible beam energy will help bridge the gap in energy between the top RHIC energy of $\sqrt{s_{\text{NN}}} = 200 \text{ GeV}$ and the LHC. More complete energy excitation functions in observables such as multiplicity, transverse-momentum spectra, particle ratios, will then be obtained.

2.2.5. Intermediate-mass ion collisions. To vary the energy density, ALICE will study at least one intermediate-mass system during the first five years of operation. The currently envisaged choice is Ar–Ar. Figure 2.3 shows the approximate energy-density bands covered by varying the impact parameter for several systems. Using medium central collisions in several systems to vary the energy density has the advantage of keeping the collision geometry similar, thus simplifying the interpretation of the data. Other options suitable for LHC (Sn–Sn, Kr–Kr, N–N, O–O) will possibly be required in later runs. The final choice will depend on the physics outcome.

Two maximum luminosities will be considered for Ar–Ar interactions: $\mathcal{L} = 2.8 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ to match the maximum Pb–Pb rates at which we can run with the TPC and $\mathcal{L} = 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$ to maximize the heavy vector-meson rate in the dimuon decay channel. The corresponding luminosities for low- and high-luminosity O–O runs are 5.5×10^{27} and $2 \times 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$, respectively.

Table 2.1. Maximum bunch intensities allowed by space-charge effects in the SPS and the corresponding maximum and average luminosities for one (two, three) experiments participating in the run. The limits for Ar–Ar and O–O are preliminary estimates.

System	Ions per bunch	\mathcal{L}_0 ($\text{cm}^{-2} \text{s}^{-1}$)	$\langle \mathcal{L} \rangle / \mathcal{L}_0$
Pb–Pb	7.0×10^7	1.0×10^{27}	0.44 (0.35, 0.29)
Ar–Ar	5.5×10^8	0.6×10^{29}	0.64 (0.62, 0.60)
O–O	1.0×10^9	2.0×10^{29}	0.73 (0.70, 0.67)

The same limitations from space-charge effects in the SPS for Pb–Pb also apply to intermediate-mass ions. Presently, only the limits for Pb–Pb collisions have been studied in detail. Scaling from the Pb–Pb values results in the scenario presented in table 2.1. Only 60% of the required high Ar–Ar luminosity can be delivered by the machine while the required high O–O luminosity is at the limit.

2.3. Proton–proton collisions

2.3.1. Standard pp collisions at $\sqrt{s} = 14 \text{ TeV}$. Since pp collisions are needed to obtain reference data and are in addition of intrinsic interest, they are an integral and important part of the heavy-ion physics programme [355, 356]. Moreover, pp runs will provide low multiplicity, thus simpler, data to commission and calibrate the components of the ALICE detector. Hence, they are needed during the whole period of ALICE operation, both initially as well as in later years for shorter periods prior to every heavy-ion run.

The pp runs will be in parallel with the other experiments but at reduced luminosities in our interaction region (IP2). In order to keep the pile-up in the TPC and Silicon Drift Detectors (SDDs) at an acceptable level, the luminosity during pp runs has to be limited to $\simeq 5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$ [355], corresponding to an interaction rate of $\simeq 200 \text{ kHz}$. At this rate, we record on average 30 overlapping events, i.e. 97% of the data volume corresponds to unusable partial events. This factor 30 increase in data volume has obvious negative consequences both in terms of data storage and offline computing requirements as well as reducing the physics performance of the central barrel detectors because of increased occupancy. While the High-Level Trigger (HLT) is expected to be able to remove pile-up events online to a significant extent, the optimal detector operation and physics performance with the TPC, i.e. no pile-up, is at $\mathcal{L} \simeq 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$. For the muon spectrometer the highest acceptable luminosity of about $5 \times 10^{31} \text{ cm}^{-2} \text{ s}^{-1}$ is set by the RPC illumination limit.

ALICE will request proton–proton operation in IP2 with both the maximum acceptable luminosity ($\mathcal{L} \simeq 5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$) in order to maximize integrated luminosity for rare processes as well as with lower luminosity ($\simeq 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$) to collect statistics for large cross section observables and global event properties at optimum DAQ bandwidth and detector performance.

Depending on the beam intensity and emittance, the luminosity reduction can be obtained either by running with a higher β^* , 50–200 m, for a reduction by factors of 100–400, or with displaced beams. This will be discussed in more detail in section 2.5.

2.3.2. Dedicated pp-like collisions. Optionally, pp-like collisions close to the heavy-ion nucleon–nucleon centre-of-mass energy, $\sqrt{s_{\text{NN}}} = 5.5\text{--}7 \text{ TeV}$, like low-energy pp, dd or $\alpha\alpha$, might be needed for further reference data, such as resolving ambiguities in Monte Carlo interpolations from $\sqrt{s} = 14 \text{ TeV}$ to the A–A nucleon–nucleon centre-of-mass energies.

Collisions of very low-mass ions, such as d or α , can be treated in most cases as systems of independent nucleons. The centre-of-mass energy per nucleon pair is 7 TeV in the LHC nominal configuration. These collisions have the additional advantage of having factors of 4 (dd) and 16 ($\alpha\alpha$) higher hard cross sections as compared to pp collisions thereby compensating for the shorter runs.

Limiting the interaction rate to 200 kHz, as in standard pp collisions, the required luminosities are 1.1×10^{30} and $6.2 \times 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$ for dd and $\alpha\alpha$, respectively. The net gain in the hard interaction rate, taking into account the lower luminosities, is a factor of 1.5 in dd and a factor of 3.3 in $\alpha\alpha$.

2.4. pA collisions

In addition to pp reference data, pA collisions are also needed for reference data and for studying gluon shadowing. They are thus an integral part of the heavy-ion physics programme.

The optimal luminosity in pA runs is rate and not lifetime limited [357]. Constraints exist from the source and the injection. However, since the luminosity is proportional to the product of the two beam intensities and since high-intensity proton beams are available, the ion-beam intensity can be optimized accordingly. Limiting the collision rate to 200 kHz, as in pp collisions, leads to luminosities of 1.1×10^{29} and $3 \times 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$ for pPb and pAr collisions, respectively.

For asymmetric collision systems like pA, the Z/A ratio is very different for the two beams leading to different nominal momenta, 7 TeV for protons and about 3.5 TeV for ions. One consequence is that the centre-of-mass energy per nucleon pair given by

$$\sqrt{s_{\text{NN}}} = 14 \text{ TeV} \times \sqrt{\frac{Z_1 Z_2}{A_1 A_2}} \quad (2.2)$$

is higher in pA collisions ($14 \text{ TeV} \times \sqrt{Z/A}$) compared to A–A ones ($14 \text{ TeV} \times Z/A$).

Another consequence is that the centre-of-mass system of the collision moves with respect to the laboratory system. This movement corresponds to a shift of the central rapidity Δy observed in the laboratory system

$$\Delta y = 0.5 \ln \frac{Z_1 A_2}{Z_2 A_1}. \quad (2.3)$$

Alternatively, dA or αA collisions may have the advantage that they are more symmetric in Z/A . Hence, the centre-of-mass energy per nucleon pair is closer to that of heavy-ion collisions: 6.2 TeV (6.6 TeV) for dPb (dAr) collisions as compared to 8.8 TeV (9.4 TeV) for pPb (pAr) ones. Furthermore, the central rapidity is less shifted: 0.12 (0.05) for dPb (dAr) collisions as compared to 0.47 (0.40) for pPb (pAr) ones. In any case, since the ALICE experiment is asymmetric in the beam direction, both pA and Ap runs are planned. The luminosity limits are summarized in table 2.2.

2.5. Machine parameters

In table 2.3, the LHC machine parameters for pp and Pb–Pb operation are summarized.

Luminous region. The size of the luminous region, vertex spread, in the approximation of Gaussian shaped, bunches is given by

$$\sigma_v = \frac{\sigma_1}{\sqrt{2}} F. \quad (2.4)$$

Table 2.2. Maximum nucleon–nucleon centre-of-mass energies, rapidity shifts, geometric cross sections, and lower and upper limits on luminosities for several different symmetric and asymmetric systems.

System	$\sqrt{s_{NN_{max}}}$ (TeV)	Δy	σ_{geom} (b)	\mathcal{L}_{low} (cm ⁻² s ⁻¹)	\mathcal{L}_{high} (cm ⁻² s ⁻¹)
Pb–Pb	5.5	0	7.7	1.0×10^{27}	
Ar–Ar	6.3	0	2.7	2.8×10^{27}	1.0×10^{29}
O–O	7.0	0	1.4	5.5×10^{27}	2.0×10^{29}
N–N	7.0	0	1.3	5.9×10^{27}	2.2×10^{29}
$\alpha\alpha$	7.0	0	0.34	6.2×10^{29}	
dd	7.0	0	0.19	1.1×10^{30}	
pp	14.0	0	0.07	1.0×10^{29}	5.0×10^{30}
pPb	8.8	0.47	1.9	1.1×10^{29}	
pAr	9.4	0.40	0.72	3.0×10^{29}	
pO	9.9	0.35	0.39	5.4×10^{29}	
dPb	6.2	0.12	2.6	8.1×10^{28}	
dAr	6.6	0.05	1.1	1.9×10^{29}	
dO	7.0	0.00	0.66	3.2×10^{29}	
α Pb	6.2	0.12	2.75	7.7×10^{28}	
α Ar	6.6	0.05	1.22	1.7×10^{29}	
α O	7.0	0.00	0.76	2.8×10^{29}	

Table 2.3. LHC machine parameters for pp and Pb–Pb runs for ALICE.

	pp	Pb–Pb
Energy per nucleon (TeV)	7	2.76
β at the IP: β^* (m)	10	0.5
R.m.s. beam radius at IP: σ_t (μ m)	71 ^a	15.9
R.m.s. bunch length: σ_l (cm)	7.7	7.7
Vertical crossing half-angle (μ rad) for pos. (neg.) μ -spectr. dipole polarization	150 (150)	150 (100)
No. of bunches	2808	592
Bunch spacing (ns)	24.95	99.8
Initial number of particles per bunch	1.1×10^{11}	7.0×10^7
Initial luminosity (cm ⁻² s ⁻¹)	$< 5 \times 10^{30}$	10^{27} ^b

^a For low-intensity runs β^* could be 0.5 m and $\sigma_t = 15.9 \mu$ m as in Pb–Pb.

^b Early operation will be with 62 bunches and $\beta^* = 1$ m, which yields an initial luminosity of 5.4×10^{25} cm⁻² s⁻¹.

The factor F is due to the finite crossing angle ϕ and depends on the ratio between longitudinal, σ_l , and transverse, σ_t , beam sizes

$$F = \frac{1}{\sqrt{1 + (\tan(\phi/2)\sigma_l/\sigma_t)^2}}. \quad (2.5)$$

Its value is 0.81 for $\beta^* = 0.5$ m and 0.99 for $\beta^* = 10$ m resulting in initial vertex spreads of 4.3 and 5.4 cm for Pb–Pb and pp runs, respectively.

At injection the shape of the bunches is close to a truncated \cos^2 -distribution. During the coast, longitudinal intra-beam scattering and RF noise increase the longitudinal beam size (bunch length) and change its shape progressively to a Gaussian. For pp collisions this increase is dominated by the influence of RF noise and is expected to be 30% after a 10 h coast [358]. In Pb–Pb collisions the increase is dominated by intra-beam scattering. The expected increase is




Figure 2.4. Location of beam elements for the interaction region IR2.




Figure 2.5. Optical functions for Ring 1 close to IR2 for $\beta^* = 10$ m: horizontal β -function, β_x (solid line), vertical β -function, β_y (dashed line) and horizontal dispersion, D_x (dotted line).

currently under study [359]. The increase of the luminous region proportional to the increase of the bunch length at $\phi = 0$ is, for finite crossing-angles, partially compensated by the decrease of F . For a 30% increase of the bunch length, this decrease amounts to 10% for Pb–Pb and 1% for pp.

In summary, assuming that the bunch length increases by 30% during the run, the size, σ_v of the luminous region increases from 5.4 cm to 7.0 cm (4.3–5 cm) in pp (Pb–Pb) collisions. The Pb–Pb value is important for ALICE, because of the limited luminosity available for this collision system. In case the moderate increase of bunch length is confirmed, no loss of events due to the vertex spread is expected.

Luminosity reduction in pp runs. In the present nominal design, the IR2 optics will work in the pp-collision mode for a $\beta^* = 10$ m as shown in figures 2.4 and 2.5. Thus the luminosity is reduced by only a factor of 20 compared to the dedicated high-luminosity experiments, not enough to reach the desired luminosity of about $5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$. An additional reduction can be obtained by displacing the two beam centres at the interaction point by a distance d .

Table 2.4. Early and nominal LHC pp operation parameters.

Scenario	Number of bunches	Protons per bunch	Bunch spacing (ns)	Normalized emittance (nm)	\mathcal{L} ($\text{cm}^{-2} \text{s}^{-1}$)
Nominal	2808	1.15×10^{11}	24.95	3.75	1.0×10^{34}
Early operation	1188	0.40×10^{11}	74.85	3.75	5.0×10^{32}

The reduction factor f_R obtained in this way is

$$f_R = \exp\left(-\frac{d^2}{4\sigma^2}\right), \quad (2.6)$$

where σ is the transverse beam size. We need $d/\sigma = 4.3$ in order to reduce the luminosity to $5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$ at full bunch intensity and $d/\sigma = 5.8$ for a reduction to $10^{29} \text{ cm}^{-2} \text{ s}^{-1}$ at initial LHC running.

During the LHC running-in phase, beam intensities and luminosities are lower than their nominal values. Although the present design values (table 2.4) do not allow for such a scenario, it could well be that during some limited time the beam intensities are so low that defocussing is enough to reach the desired luminosity. Since we expect to obtain the most stable running conditions without beam displacement we asked the SL division to calculate the maximum for a range of beam intensities. At present $\beta^* \simeq 100 \text{ m}$ yielding a factor of 200 reduction does not seem impossible.

2.6. Radiation environment

2.6.1. Background conditions in pp

2.6.1.1. Background from beam–gas collisions in the experimental region. The dynamic gas pressure inside the experimental region, $\pm 20 \text{ m}$ around the interaction point, has been simulated [360]. The considered dynamic effects comprise proton-beam-induced phenomena: ion, electron and photon stimulated molecular desorption, ion-induced desorption instability, and electron cloud build-up. It has been shown that electron-cloud build-up is the dominant effect [361]. The simulations assume a stainless-steel beam pipe with TiZrV sputtered Non-Evaporable Getter (NEG) pump over the whole surface and ion pumps at about $\pm 20 \text{ m}$ from the interaction point.

The resulting hydrogen-equivalent gas density in pp operation amounts to $\simeq 2 \times 10^{14}$ molecules m^{-3} . For Pb operation the gas density is expected to be two orders of magnitude lower. At full current, the gas density in proton operation leads to a beam–gas interaction rate of 12 kHz m^{-1} , i.e. 500 kHz integrated over the ALICE experimental region, to be compared to the pp collision rate of 200 kHz .

2.6.1.2. Background from IR2 straight section. The dynamic gas densities in the Dispersion Suppressor (DS) and the Long Straight Section (LSS) of insertion region IR2, excluding the beam pipe within the experimental region, $\pm 20 \text{ m}$, have been estimated for beam optics version 6.3 and for the first years of running [362]. The dynamic gas density at LHC is determined by ion, photon, and electron stimulated desorption. The latter two change with absorbed dose, so that results have been presented for a specific run scenario during the first years of LHC operation. The results have been obtained using pessimistic assumptions, i.e. no NEG pumping




Figure 2.6. Residual H₂-equivalent density distribution for IR2 in the first three years of LHC operation as a function of the distance from interaction point.

in the elements at room temperature and a factor of 100 overestimate of the out-gassing rate inside the injection beam stopper (TDI). After discussing the issue with the vacuum group, it has been agreed that the realistic reference for background simulation should be the third-year scenario presented in [362] with the gas pressure inside the TDI scaled down by a factor of 100, shown in figure 2.6.

In a more recent calculation [363] the residual gas densities have been estimated using the VAcuum Stability COde (VASCO) [364] for two different scenarios: the machine start-up and after machine conditioning. Machine conditioning using electron scrubbing [365] is expected to lower the secondary electron emission avoiding electron multiplication. The residual gas densities are estimated to be a few 10^{14} molecules m^{-3} at machine start-up and about two orders of magnitude lower after machine conditioning. These values will decrease during each operational period since desorption yields will reduce with particle bombardment. Measurements of electron induced desorption yields are ongoing and will allow the refinement of the parameters used for these estimates.

To be conservative we present here calculations using the previous estimates (figure 2.6), which lie between those for the running-in scenario and the period after conditioning. Based on these gas pressure calculations, the fluxes of secondary particles induced by the proton losses upstream and downstream of IR2 were estimated [366]. Azhgirey *et al* [366] have provided to the ALICE Collaboration an output file containing the particle types and momentum vectors of particles entering the experimental region. The particle statistical weight has been normalized to one proton–gas interaction per second and per metre. Knowing the distance to the interaction point for primary proton–gas collisions, we are able to perform background simulations for the ALICE experiment for different gas–pressure distributions. The muon and hadron fluxes at the tunnel entrance as a function of the distance to the beam pipe are shown in figure 2.7. The maximum muon flux is about 3.5 Hz cm^{-2} whereas the hadron flux has a wide plateau of about 30 Hz cm^{-2} and a maximum of 1 MHz cm^{-2} inside the beam pipe. Hadrons outside the beam pipe, $r = 3 \text{ cm}$, will be shielded. Particles that enter the experimental area inside the beam pipe are by far more dangerous since they will shower further downstream closer to the ALICE central detectors.

2.6.2. Dose rates and neutron fluences. We shall take data with proton and light- or heavy-ion beams for different time periods and different luminosities. The radiation load on the various




Figure 2.7. Muon flux (a) and hadron fluxes (b), (c) at the tunnel entrance as a function of the distance from the beam axis; (c) an enlarged view of (b) for $R < 10$ cm.

parts of the detectors must therefore be calculated for a combination of beam conditions. There are three major sources of radiation in ALICE:

- particles produced at the interaction point in planned collisions;
- beam losses due to mis-injection since ALICE is located near the injection point;
- beam–gas interactions in pp operation.

Table 2.5. Operation scenario for a 10-year run period, where $\langle \mathcal{L} \rangle$ is mean luminosity, and σ_{inel} is the inelastic cross section. One year of pp run corresponds to 10^7 s and 1 year of heavy-ion run corresponds to 10^6 s.

	pp	Ar–Ar	Ar–Ar	Pb–Pb	dPb
$\langle \mathcal{L} \rangle$ ($\text{cm}^{-2} \text{s}^{-1}$)	3×10^{30}	3×10^{27}	10^{29}	10^{27}	8×10^{28}
σ_{inel} (mb)	70	3000	3000	8000	2600
Rate (s^{-1})	2×10^5	9×10^3	3×10^5	8×10^3	2×10^5
Runtime (s)	10^8	1.0×10^6	2.0×10^6	5×10^6	2×10^6
Events	2×10^{13}	9×10^9	6×10^{11}	4×10^{10}	4×10^{11}
Particles per event	100	2400	2400	14 200	500
N_{tot}	2.1×10^{15}	2.2×10^{13}	1.4×10^{15}	5.7×10^{14}	2×10^{14}

It has been shown that beam–gas interactions and beam losses contribute, respectively, 10% and 1% to the total dose and neutron fluence [368]. The dominant contribution comes from particles produced at the interaction point.

In table 2.5 we present a typical scenario for 10 years of operation. It consists of 10 years of pp operation and 5 years each of Pb–Pb and Ar–Ar (or dPb) runs. The Ar–Ar runs are split into low- and high-luminosity runs. The table shows the average luminosity, the number of expected minimum-bias collisions and the total primary particle multiplicity per event as well as the total number of produced particles N_{tot} . The latter number suggests the relative contributions to the total dose. In this scenario, 80% of the radiation dose is from pp and Ar–Ar operation.

In order to estimate the overall radiation load on the ALICE detector, typical volumes (scoring regions) were identified, generally the location of the detector electronics: Inner Tracking System (ITS) comprising Silicon Pixel Detector (SPD), Silicon Drift Detector (SDD) and Silicon Strip Detector (SSD); the Time-Projection Chamber (TPC) in the region where the electronics is situated; the Transition-Radiation Detector (TRD), the Time-Of-Flight detector (TOF); the High-Momentum Particle Identification Detector (HMPID) using a Ring Imaging Cherenkov counter (RICH) and four electronics rack locations (RackLoc 1–4). The racks are located in the following areas: RackLoc1 is on platforms situated inside the experimental area, 12 m above the floor level at the side of the L3 magnet; RackLoc2 is on the floor of the experimental area UX25 on both sides (RB24, RB26) of the L3 magnet; RackLoc3 is in four levels of counting rooms in the main access shaft PX24; and RackLoc4 is in the shielding plug PX24.

The deposited energy and the total neutron fluence, including thermal neutrons, were simulated with the FLUKA Monte Carlo code [369] for each of the selected volumes and beams. The geometry includes the main structural elements and a detailed description of the absorbers. The DPMJET-II model [370] was used for the simulation of the pp interactions while the nucleus–nucleus collisions were generated with HIJING [371].

The results are shown in table 2.6 with all values normalized to the 10-year scenario above. The pixel detectors close to the beam receive the highest dose, up to 200 krad (2 kGy) and 1×10^{12} neutrons cm^{-2} , over 10 years. Since the radiation load scales approximately with $1/r^2$, doses for the other subsystems can also be estimated from this table. In particular, the regions of the forward detectors, V0 and T0, close to the beam pipe will receive doses and neutron fluences similar to those of the pixel detector.

2.6.3. Background from thermal neutrons. In Pb–Pb collisions at a luminosity of $10^{27} \text{cm}^{-2} \text{s}^{-1}$, the primary particle production rate is 2×10^8 Hz. Many of these particles produce

Table 2.6. Doses and neutron fluences in detectors and electronic racks.

System	Radius (cm)	Dose (Gy)	Neutron fluence (cm ⁻²)
SPD1	3.9	2.2×10^3	8.0×10^{11}
SPD2	7.6	5.1×10^2	5.6×10^{11}
SDD1	14	1.9×10^2	4.5×10^{11}
SDD2	24	1.0×10^2	4.2×10^{11}
SSD1	40	4.0×10^1	4.1×10^{11}
SSD2	45	2.6×10^1	4.1×10^{11}
TPC (in)	78	1.3×10^1	3.6×10^{11}
TPC (out)	278	2.0×10^0	2.4×10^{11}
TRD	320	1.6×10^0	1.5×10^{11}
PID	350	1.1×10^0	1.0×10^{11}
HMPID	490	5.0×10^{-1}	8.0×10^{10}
RackLoc1		5.6×10^{-1}	8.4×10^7
RackLoc2		3.8×10^{-1}	1.5×10^6
RackLoc3		2.2×10^{-6}	3.5×10^3
RackLoc4		7.8×10^{-6}	9.2×10^3

secondaries through hadronic and electromagnetic cascades in the absorbers and structural elements of ALICE. In particular, neutrons are copiously produced. Highly energetic neutrons lose energy in subsequent scatterings to finally produce a gas of thermal neutrons. Each minimum-bias event produces approximately 50 thermal neutrons per m² at a radius of 5 m in ALICE central region.

Since the average time between two subsequent Pb–Pb collision, $\simeq 100 \mu\text{s}$, is much larger than the decay time of the neutron signal, we do not expect any substantial build-up of the thermal neutron fluency in the ALICE experimental regions. Only detectors using materials with high capture cross sections could experience an increase of event-uncorrelated background.

We have identified one such detector, the TRD, which uses Xe as a drift gas. The capture cross section for low-energetic neutrons has resonance peaks up to 50 kb for some of the isotopes, followed by a multi-gamma de-excitation cascade. These gammas in turn can create low-energy electrons as they Compton scatter in the Xe gas, create secondary electrons, or convert into electron–positron pairs. This will be a source of event-uncorrelated background in the TRD.

An intensive study was performed to estimate the level of this background [372]. As a result, the number of particles per TRD layer is expected to increase by not more than 10%. Given the particular topology of the hit pattern produced by low-energetic electrons spiralling along the beam direction, this background contribution cannot be neglected.

2.7. Luminosity determination in ALICE

The luminosity \mathcal{L} is a quantity which relates the rate R of a process to its cross section σ

$$R = \mathcal{L}\sigma. \quad (2.7)$$

It is entirely defined by the characteristics of the colliding beams at the interaction point

$$\mathcal{L} = fN_b \frac{N^2}{2\pi(\sigma_1^2 + \sigma_2^2)} F, \quad (2.8)$$

where f is the revolution frequency, N_b the number of bunches, N the number of particles per bunch and $\sigma_{1,2}$ the transverse beam sizes of the two beams; symmetry in the transverse plane




Figure 2.8. Charged-particle pseudo-rapidity distribution for pp interactions at $\sqrt{s} = 14$ TeV. Solid line: inelastic non-diffractive interactions; dashed line: diffractive interactions. The vertical lines indicate the acceptance of the forward detectors V0 (solid lines) and T0 (dashed lines).

at the interaction point has been assumed. The reduction factor F is due to the finite crossing angle and is given by equation (2.5).

Thus, measuring the beam parameters is one way of determining the luminosity. The accuracy, $\simeq 10\%$, is limited by the extrapolation of σ_t from measurements of beam profiles elsewhere to the interaction point.

Another way consists in measuring the rate of a well-known process. Through the optical theorem, one obtains a relation between the rate of elastic events, R_{el} , in the forward direction, i.e. momentum transfer $\sqrt{t} = 0$, and the total rate of pp interactions R_{tot} :

$$\mathcal{L} = \frac{(1 + \rho^2)}{16\pi} \frac{R_{tot}^2}{(dR_{el}/dt)_{t=0}}, \quad (2.9)$$

where ρ is the ratio of the real to imaginary parts of the elastic-scattering forward amplitude. The measurement of the elastic rate at zero momentum transfer requires a specialized very forward experiment. It uses a dedicated beam optics with low beam divergence at the interaction point [373]. At the LHC, the TOTEM Collaboration [374] will perform such a measurement. For heavy-ion collisions there exist two cross sections that can be calculated with sufficient precision: the total hadronic cross section and the cross section of mutual electromagnetic dissociation. In the following sections we will discuss the luminosity determination with ALICE for pp and A–A collisions separately.

2.7.1. Luminosity monitoring in pp runs. The TOTEM experiment [374] will measure the total cross section σ_{tot} at the LHC. In ALICE, we will measure and monitor the luminosity by measuring a fraction $R = Acc \cdot R_{tot}$ of the rate of inelastic interaction. Then luminosity

$$\mathcal{L} = \frac{R}{Acc \cdot \sigma_{inel}}, \quad (2.10)$$

where Acc is the detector acceptance. In reality the inelastic rate is the sum of the rates of the inelastic non-diffractive (R_{nd}), the single-diffractive (R_{sd}), and the double-diffractive (R_{dd}) processes. The detector acceptance is different for each of these processes (figure 2.8). Monte Carlo simulations to determine the trigger efficiency can be tuned to the angular track distributions measured by TOTEM. Experience from the Tevatron shows that the error on the




Figure 2.9. Contribution of the SD and NSD pp inelastic events on the V0 triggering efficiency as a function of the minimum number of cells required from the V0A array for the V0A \otimes V0C coincidence. Here, $N_{\text{cut}}^{\text{right}} = 1$.

acceptance can be reduced to a few per cent and that the total uncertainty on the measured luminosity is dominated by the error on $\delta\sigma_{\text{inel}}/\sigma_{\text{inel}} \simeq 5\%$.

During high-intensity proton runs, the background from beam–gas interactions could be so high that the pp trigger has to be validated by a V0 trigger. The T0 represents a back-up solution. Having a much smaller angular acceptance than the V0 it can trigger on only 70% of the inelastic non-diffractive interactions and on 10% of the diffractive events.

Luminosity monitoring with the V0. The V0 minimum bias efficiency can be obtained from MC simulation by evaluating the fraction of inelastic events detected by the coincidence between V0A and V0C (V0A \otimes V0C). At least one MIP through each array is required. An additional selection on the V0A and V0C time-of-flight difference rejects events from beam–gas interactions. However, simulations show that this does not lead to a complete suppression and a low threshold on the multiplicity provided by each V0 array is necessary to improve the beam–gas event rejection capability. Since the multiplicity resolution of the individual V0 channels is not good enough to allow cuts on multiplicity, only a low threshold on the number of fired cells ($N_{\text{cut}}^{\text{left}}$, $N_{\text{cut}}^{\text{right}}$) in the V0L and V0R, respectively, can be applied to reduce the background.

Figure 2.9 shows the V0 trigger efficiency distributions for single-diffractive (SD) and non-single-diffractive (NSD) pp events as a function of a series of $N_{\text{cut}}^{\text{left}}$ values with $N_{\text{cut}}^{\text{right}} = 1$. Due to the limited rapidity coverage there is no efficiency for elastic scattering. If N_{cut} values larger than 1 have to be applied, the efficiency, $\varepsilon_{\text{inel}}$, will decrease and, consequently, the uncertainty on the luminosity value will increase. Table 2.7 shows values of the resulting trigger efficiency for $N_{\text{cut}}^{\text{left}} = 1\text{--}6$ and $N_{\text{cut}}^{\text{right}} = 1\text{--}2$ applied to left and right arrays, respectively.

2.7.2. Luminosity monitoring in Pb–Pb runs. Two interaction cross sections known with reasonable accuracy can be used for luminosity determination in heavy-ion collisions. The

Table 2.7. V0 triggering efficiencies for the signal (ϵ_{inel}) as a function of N_{cut} on left and right cells required for the VOA \otimes V0C coincidence in pp reaction at 7 TeV per beam. $N_{\text{cut}}^{\text{left}} = N_{\text{cut}}^{\text{right}} = 1$ are conditions for the minimum bias trigger.

		ϵ_{inel} (%)					
$N_{\text{cut}}^{\text{left}}$		1	2	3	4	5	6
$N_{\text{cut}}^{\text{right}}$	1	86	80	76	70	66	60
	2	82	78	74	69	65	59

total hadronic cross section σ_{had} is mainly given by the geometry of the colliding nuclei and is known with an accuracy better than 10%. Hence, measuring the hadronic interaction rate R_{had} will allow us to calculate the luminosity, $\mathcal{L} = R_{\text{had}}/\sigma_{\text{had}}$. Of course, the measurement is only sensitive up to some maximum impact parameter, thus there will be an additional small systematic error from extrapolating to the total rate. The second possibility is to measure the rate of electromagnetic dissociation.

Using electromagnetic dissociation to measure and monitor the luminosity in heavy-ion colliders has been proposed in [375]. According to this method, the rate of mutual electromagnetic dissociation events R_{m}^{ED} measured by means of Zero-Degree Calorimeters (ZDCs) provides the luminosity value:

$$\mathcal{L} = \frac{R_{\text{m}}^{\text{ED}}}{\sigma_{\text{m}}^{\text{ED}}}. \quad (2.11)$$

The method relies on the accuracy with which the mutual electromagnetic dissociation cross section $\sigma_{\text{m}}^{\text{ED}}$ is computed from theory. Simultaneous forward–backward single-neutron emission from each of the collision partners provides a clear signal of the mutual dissociation process. For the most part this process proceeds through the absorption of virtual photons, also called equivalent photons, emitted by collision partners which is followed by the excitation and subsequent decay of the Giant Dipole Resonances (GDRs) in both of the colliding nuclei. In heavy nuclei, like Au or Pb, the single-neutron emission channel (1n) is the main mechanism of GDR decay. The same basic idea has been adopted for luminosity measurement and monitoring in the ALICE ZDC. We refer the reader to [376] for an extended technical description of the ALICE ZDC.

2.7.2.1. Cross sections predicted by the RELDIS model. In [377, 378], the RELDIS (RELativistic-DISsociation model) results were tested by studying their sensitivity to the variation of input data and parameters. Corrections for the Next-to-Leading Order (NLO) processes of mutual dissociation were also taken into account. In addition to GDR excitation, quasi-deuteron absorption, Δ photo-excitation, and multiple-pion photo-production were considered in different regions of equivalent photon energy, E_{γ} . A good description of the existing data on single dissociation of lead and gold nuclei at the CERN SPS was obtained [379, 380].

In particular, the photo-nuclear cross sections for specific neutron emission channels were calculated by two different models of photo-nuclear reactions, the GNASH code [381] and the RELDIS code itself, see table 2.8 and [377, 378] for details. In addition, in the latter model two different probabilities of direct neutron emission in the 1n channel, $P_{\text{n}}^{\text{dir}}$ were assumed.

The model also defines the cumulative value, Low Multiplicity Neutron (LMN) emission cross section,

$$\begin{aligned} \sigma_{\text{m}}^{\text{ED}}(\text{LMN}) &= \sigma_{\text{m}}^{\text{ED}}(1\text{nX} | 1\text{nY}) + \sigma_{\text{m}}^{\text{ED}}(1\text{nX} | 2\text{nY}) \\ &+ \sigma_{\text{m}}^{\text{ED}}(2\text{nX} | 1\text{nY}) + \sigma_{\text{m}}^{\text{ED}}(2\text{nX} | 2\text{nY}) \end{aligned} \quad (2.12)$$

Table 2.8. Sensitivity of the mutual electromagnetic dissociation cross sections to the equivalent photon-energy range, the probability of direct neutron emission in the 1n channel, P_n^{dir} , the photo-nuclear cross sections, and the NLO corrections. The results obtained with the GNASH and RELDIS codes are given for 5.5 TeV per nucleon pair Pb–Pb collisions. The recommended values are those in the last column. The prediction of [375] for $\sigma_m^{\text{ED}}(1n|1n)$ is given for comparison.

	Cross section (mb)				
	$E_\gamma \leq 24$ MeV		$E_\gamma \leq 140$ MeV		Full range of E_γ
	LO		LO		LO + NLO
	RELDIS $P_n^{\text{dir}} = 0$	GNASH	RELDIS $P_n^{\text{dir}} = 0$	RELDIS $P_n^{\text{dir}} = 0$	RELDIS $P_n^{\text{dir}} = 0.26$
$\sigma_m^{\text{ED}}(1nX 1nY)$	519 533 [375]	488	544	727	805
$\sigma_m^{\text{ED}}(1nX 2nY) + \sigma_m^{\text{ED}}(2nX 1nY)$	154	220	217	525	496
$\sigma_m^{\text{ED}}(2nX 2nY)$	11	24	22	96	77
$\sigma_m^{\text{ED}}(\text{LMN})$	684	732	783	1348	1378

where X and Y denote any particle other than a neutron. The LMN cross section was calculated for several regions of equivalent photon energy in the GDR region, $E_\gamma < 24$ MeV, for the GDR and quasi-deuteron absorption region, $E_\gamma < 140$ MeV, and for the full range of E_γ .

The uncertainty in the 1n–1n correlated emission cross section, $\sigma_m^{\text{ED}}(1nX|1nY)$, is up to 10% (see table 2.8), mainly due to the uncertainty in the measured photo-neutron cross section. However, the uncertainty is reduced to $\sim 2\%$ if the sum of the one- and two-neutron emission channels, $\sigma_m^{\text{ED}}(\text{LMN})$, is considered. As shown in [377, 378], $\sigma_m^{\text{ED}}(\text{LMN})$ is also more stable with respect to other input parameters compared to $\sigma_m^{\text{ED}}(1nX|1nY)$ and other cross sections. Therefore, $\sigma_m^{\text{ED}}(\text{LMN})$ represents a cumulative neutron emission rate which should be used for the luminosity measurement at heavy-ion colliders.

3. ALICE detector

3.1. Introduction

The ALICE experiment is described in the ALICE Technical Proposal (TP) [382] and its two Addenda [383, 384]. The detector systems are described in detail in Technical Design Reports (TDR) [385–397]. In this section, we give an overview of the present status of the ALICE detectors, trigger, data-acquisition and high-level trigger systems. It includes lists of the main parameters, updated with the latest design adjustments and recent results obtained with prototypes. Since most systems have already entered the construction phase, this section provides a description of the ALICE experiment in its final layout shown in the figures I and II.

ALICE is a general-purpose experiment whose detectors measure and identify mid-rapidity hadrons, leptons and photons produced in the interaction. A unique design, with very different optimization than the one selected for the dedicated pp experiments at LHC, has been adopted for ALICE. This results from the requirements to track and identify particles from very low (~ 100 MeV c^{-1}) up to fairly high (~ 100 GeV c^{-1}) p_t , to reconstruct short-lived particles such as hyperons, D and B mesons, and to perform these tasks in an environment with large charged-particle multiplicities, up to 8000 charged particles per rapidity unit at mid-rapidity. The detection and identification of muons are performed with a dedicated spectrometer,

including a large warm dipole magnet and covering a domain of large rapidities¹ ($-4.0 \leq \eta \leq -2.4$). Hadrons, electrons and photons are detected and identified in the central rapidity region ($-0.9 \leq \eta \leq 0.9$) by a complex system of detectors immersed in a moderate (0.5 T) magnetic field. Tracking relies on a set of high-granularity detectors: an Inner Tracking System (ITS) consisting of six layers of silicon detectors, a large-volume Time-Projection Chamber (TPC) and a high-granularity Transition-Radiation Detector (TRD). Particle identification in the central region is performed by measuring energy loss in the tracking detectors, transition radiation in the TRD, Time Of Flight (TOF) with a high-resolution array, Cherenkov radiation with a High-Momentum Particle Identification Detector (HMPID), and photons with a crystal PHOTon Spectrometer (PHOS). Additional detectors located at large rapidities complete the central detection system to characterize the event and to provide the interaction trigger. They cover a wide acceptance ($-3.4 \leq \eta \leq 5.1$) for the measurement of charged particles and triggering (Forward Multiplicity Detector—FMD, V0 and T0 detectors), and a narrow domain at large rapidities ($2.3 \leq \eta \leq 3.5$) for photon multiplicity measurement (Photon Multiplicity Detector—PMD), and the coverage of the beams' rapidity to measure spectator nucleons in heavy-ion collisions (Zero-Degree Calorimeters—ZDC).

The physics programme of ALICE includes running periods with ions lighter than lead in order to study the energy-density dependence of the measured phenomena, as detailed in section 2.2. It also includes data taking during the pp runs and dedicated proton–nucleus runs to provide reference data and to address a specific pp physics programme (see for details section 1). Issues concerning pp and light-ion running modes will be addressed, whenever relevant, for individual detector subsystems.

3.2. Design considerations

Theoretically founded predictions for the multiplicity in central Pb–Pb collisions at the LHC range at present from 2000 to 6000 charged particles per rapidity unit at mid-rapidity, while extrapolations from RHIC data point at values of about 3500, see section 1. The ALICE detectors are designed to cope with multiplicities up to 8000 charged particles per rapidity unit, a value which ensures a comfortable safety margin.

The event rate for Pb–Pb collisions at the LHC nominal luminosity of $10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ will be of about 8000 minimum-bias collisions per second, of which only about 5% correspond to the most central ones, see section 2. This low interaction rate plays a crucial role in the design of the experiment, since it enables us to use slow but high-granularity detectors, such as the time-projection chamber and the silicon drift detectors.

The detector acceptance must be sufficiently large to enable the study, on an event-by-event basis, of particle ratios, p_t spectra, and HBT (Hanbury–Brown–Twiss) correlations, see section 1. This implies tracking several thousand particles in every event. In addition, a coverage of about two units of rapidity, together with an adequate azimuthal coverage, is required to detect the decay products of low-momentum particles ($p_t < m$ for $m > 1\text{--}2 \text{ GeV}$). A similar acceptance is necessary to collect a statistically significant sample of Υ in the dielectron channel (a few 10^3 with the expected integrated luminosity). The coverage of the central detectors, $|\eta| < 0.9$ and full azimuth, is a compromise between these requirements,

¹ In ALICE the coordinate axis system is a right-handed orthogonal Cartesian system with the point of origin at the beam interaction point. The axes are defined as follows: x -axis is perpendicular to the mean beam direction, aligned with the local horizontal and pointing to the accelerator centre; y -axis is perpendicular to the x -axis and to the mean beam direction, pointing upward; z -axis is parallel to the mean beam direction. Hence the positive z -axis is pointing in the direction opposite to the muon spectrometer. This convention, which is coherent with other LHC experiments, is different from the one adopted in the previous ALICE documents.

cost, and practical implementation. To gain in sensitivity on the global event structure, charged-particle multiplicity will be measured in a larger rapidity domain ($-3.4 < \eta < 5.1$).

The design of the tracking system has primarily been driven by the requirement for safe and robust track finding. The detectors use mostly three-dimensional hit information and continuous tracking with many points in a moderate magnetic field. The field strength is a compromise between momentum resolution, acceptance at low-momentum, and tracking and trigger efficiency. The momentum cutoff should be as low as possible ($< 100 \text{ MeV } c^{-1}$), in order to study collective effects associated with large length scales and to detect the decay products of low- p_t hyperons. A low- p_t cutoff is also mandatory to reject the soft conversion and Dalitz background in the lepton-pair spectrum. At high p_t the magnetic field determines the momentum resolution, which is essential for the study of jet quenching and high- p_t leptons. It also determines the effectiveness of the online selection of high-momentum electrons, which is performed in the TRD trigger processors by an online sagitta measurement. The ideal choice for hadronic physics, maximizing reconstruction efficiency, would be around 0.2 T, while for the high- p_t observables the maximum field the magnet can produce, 0.5 T, would be the best choice. Since the high- p_t observables are the ones which are limited by statistics, ALICE will run at the higher field for the majority of the time. The most stringent requirement on momentum resolution in the low- p_t region is imposed by identical-particle interferometry. In the intermediate energy regime, the mass resolution should be of the order of the natural width of the ω and ϕ meson, in order to study in particular the mass and width of these mesons in the dense medium. At high momenta, a reasonably good resolution up to $100 \text{ GeV } c^{-1}$ is essential for jet-quenching studies, since it has to be sufficient to measure the leading particles of jets up to jet momenta above a few hundreds of $\text{GeV } c^{-1}$. The resolution must also be sufficient to separate the different states of the Υ family. The detection of hyperons, and even more of D and B mesons, requires in addition a high-resolution vertex detector close to the beam pipe.

Particle identification with a resolving power better than 3σ , for π , K and p, is needed on a track-by-track basis to measure HBT correlations with identified particles, to identify hyperons, vector mesons ($\phi \rightarrow K^+K^-$) and heavy-flavour mesons through their hadronic decay, and to measure particle ratios on an event-by-event basis. A resolving power better than 2σ will be sufficient to construct from statistical analysis inclusive particle ratios and p_t spectra for higher momenta. Hadron identification with full acceptance and 3σ resolving power in the ALICE central detector is provided up to the momenta with which the bulk of hadrons are produced, $p_t \leq 2.5 \text{ GeV } c^{-1}$, while at higher momenta the measurement will be provided with reduced acceptance and/or reduced resolving power. The e/π rejection must allow on one hand the reduction of the electron-pair combinatorial background, due to misidentified pions, below the background level due to leftover Dalitz pairs. On the other hand, it must allow the extension of identification to the high p_t electrons created in the decay of Υ states.

Direct photons must be measured among the overwhelming background of decay photons. At low p_t ($p_t < 10 \text{ GeV } c^{-1}$) systematic errors on the inclusive photon spectrum are mainly determined by the accuracy with which neutral mesons (π^0 and η) are identified. At larger p_t values direct photons and neutral mesons must be identified to enable jet physics analysis. These requirements can be fulfilled with a high-granularity (spatial resolution) and high-resolution (energy resolution) calorimeter covering a limited acceptance around central rapidity.

The muon spectrometer is designed to measure the heavy-quark resonances spectrum and identify J/ψ and ψ' , Υ , Υ' and Υ'' particles. Production rates and decay geometry impose an acceptance which covers the full azimuth and more than one unit in rapidity. The LHC environment prevents muon identification with momenta below about $4 \text{ GeV } c^{-1}$. This points

to the large rapidity region, where muons are Lorentz boosted, for the measurement of low- p_t charmonia. The rapidity coverage ($-4.0 \leq \eta \leq -2.4$) was chosen as a compromise between requirements and costs. Having the muon spectrometer and the central detectors covering distinct and non-overlapping rapidity domains has the advantage of enabling independent optimizations of the two detection systems. To resolve the various states, the mass resolution must be better than 70 MeV at the J/ψ (3 GeV) and 100 MeV at the Υ (10 GeV). This sets the requirements on the tracking resolution and magnetic bending power. To operate the tracking chambers at tolerable occupancies, a hadron absorber must shield the chambers from the high hadron multiplicity generated in Pb–Pb collisions and secondary particles generated in the beam pipe. The absorber-material composition must be selected to efficiently absorb hadrons, while limiting muon scattering to preserve mass resolution.

A minimum-bias trigger must provide a fast signal to the slower detectors with good efficiency at all multiplicities, and therefore with large acceptance and dynamic range, and be able to reject beam–gas interactions. This can be achieved with a set of segmented scintillator counters placed on both sides of the interaction point at large rapidities. The significance of high- p_t observables is limited by statistics. To increase the sensitivity of the experiment to these observables, dedicated trigger systems that will select candidate events in the muon spectrometer and in the TRD are necessary.

The ALICE data-acquisition system will have to be able to cope with an unprecedented data rate, defined in particular by the very large data volume from the high-granularity tracking detectors. Estimates of the final throughput vary depending on the degree of online data reduction, event size, and trigger selectivity and composition, but in general an overall data flow in the neighbourhood of 1.25 GB s^{-1} is expected, and has been taken as the ALICE DAQ design guideline.

3.3. Layout and infrastructure

3.3.1. Experiment layout. The ALICE experiment, shown in figure 3.1, consists of a central detector system, covering mid-rapidity ($|\eta| \leq 0.9$) over the full azimuth, and several forward systems. The central system is installed inside a large solenoidal magnet which generates a magnetic field of $\leq 0.5 \text{ T}$. The central system includes, from the interaction vertex to the outside, six layers of high-resolution silicon detectors (Inner Tracking System—ITS), the main tracking system of the experiment (Time-Projection Chamber—TPC), a transition radiation detector for electron identification (Transition-Radiation Detector—TRD), and a particle identification array (Time-Of-Flight—TOF). The central system is complemented by two small-area detectors: an array of ring-imaging Cherenkov detectors ($|\eta| \leq 0.6$, 57.6° azimuthal coverage) for the identification of high-momentum particles (High-Momentum Particle Identification Detector—HMPID), and an electromagnetic calorimeter ($|\eta| \leq 0.12$, 100° azimuthal coverage) consisting of arrays of high-density crystals (PHOTon Spectrometer—PHOS). The large rapidity systems include a muon spectrometer ($-4.0 \leq \eta \leq -2.4$, on the RB26 side of the solenoid), a photon counting detector (Photon Multiplicity Detector—PMD, on the opposite side), an ensemble of multiplicity detectors (Forward Multiplicity Detector—FMD) covering the large rapidity region (up to $\eta = 5.1$). A system of scintillators and quartz counters (T0 and V0) will provide fast trigger signals, and two sets of neutron and hadron calorimeters, located at 0° and about 115 m away from the interaction vertex, will measure the impact parameter (Zero-Degree Calorimeter—ZDC). An absorber positioned very close to the vertex shields the muon spectrometer. The spectrometer consists of a dipole magnet, five tracking stations, an iron wall (muon filter) to absorb remaining hadrons, and two trigger stations behind the muon filter.




Figure 3.1. Longitudinal view of the ALICE detector.




Figure 3.2. Point-2 surface buildings.

3.3.2. Experimental area. The ALICE experiment is located at Intersection Point (IP) 2 of the LHC machine. Figure 3.2 shows the surface buildings at Point 2. The ALICE detector is housed in an underground cavern with the intersection point 45 m below ground level. The overburden amounts to 29 m of rock at the top of the cavern. The detector operations are controlled and supervised from the ALICE control room in the SX2 building. A large part of the electronics for the ALICE detector is housed in four floors of counting rooms, which are suspended inside the UX24 shaft, see figure 3.3. The counting rooms are separated from the experiment cavern by a concrete shielding plug. Parts of the gas systems for the detector are installed on top of the plug. The primary gas supply comes from a special surface building, SGX2. Other services, such as electricity, cooling and ventilation, are distributed from the




Figure 3.3. Point-2 underground area.

corresponding service buildings at the surface. Building SXL2 has been used for construction and pre-installation of detector components. Personnel access to the cavern is through the two shafts UX24 and PM25. Equipment for the experiment reaches the cavern through the SX2 building and the UX24 shaft.


Racks for trigger, electronics, gas systems and services such as cooling, ventilation, and power distribution are located on the cavern floor and along the cavern walls. A concrete beam shield is located between the L3 solenoid and the RB24 section of the LHC tunnel. The laser for the TPC has been placed below this beam shield. The ventilation system circulates $20\,000\text{ m}^3\text{ h}^{-1}$ of air through the UX cavern under normal operating conditions and keeps it at a temperature to within $\pm 2^\circ\text{C}$ at a relative humidity of $\sim 50\%$.

3.3.3. Magnets. The ALICE experiment uses two large magnets. The central part of the detector is enclosed in the solenoid magnet constructed for the L3 experiment at LEP, with an internal length of 12 m and a radius of 5 m. The nominal field of the solenoid is 0.5 T. The diameter of the axial holes in the magnet ‘doors’ has been reduced in order to improve the magnetic field homogeneity in the volume of the TPC. An improvement by a factor two has been achieved compared to the L3 situation. The field variations in the volume of the detectors, up to 2.5 m in radius and ± 2.5 m along the axis around the centre, are below 2% of the nominal field value. A large warm dipole magnet with resistive coils and a horizontal field perpendicular to the beam axis is used for the muon spectrometer. The field integral in the forward direction is 3 T m. The tracking stations are situated up to 14 m from the IP in a conical acceptance volume covering a region in rapidity $-4.0 \leq \eta \leq -2.4$. The polarity of the magnetic fields in both magnets can be reversed within a short time. Table 3.1 lists the magnetic field and electrical power of the two magnets.

3.3.4. Support structures. The TPC, TRD and TOF detectors are supported inside the L3 solenoid by the so-called space frame, see figure 3.4. The space frame is a cylindrical, metallic structure of about 7 m length and with a diameter of 8.5 m. The HMPID is mounted on a

Table 3.1. Electrical power and magnetic field of the magnets.

Magnet	Magnetic field	Electrical power (MW)
L3 solenoid	0.5 T	4.0
Muon dipole	3.0 T m	3.8

**Figure 3.4.** Support structures inside the L3 solenoid.

cradle, which in turn is attached to the space frame. The PHOS has its own support frame independent of the space frame. The total weight of detectors on the space frame is about 77 t, see table 3.2.

Provisions are made to incorporate at a later stage an electromagnetic calorimeter outside of the TOF detector. Additional structures on both sides of the space frame carry services and cables, and serve as access platforms. On the RB26 side of the space frame, the *back frame* supports services for the detectors and a 20 cm thick aluminium ring with inner and outer diameters of 2 and 3.6 m respectively. This aluminium disc serves as particle absorber in order to reduce the background in the muon arm. For the weight of these elements, see table 3.3.

On the RB24 side the *baby space frame* serves the same function together with the *mini space frame*, which is supported on the baby space frame and the laser platform outside the L3 solenoid. The mini frame carries the TPC services, supports the PMD detector, and provides access through the central hole in the L3 doors to the inside of the L3 solenoid. The weights of the equipment on the baby space frame are listed in table 3.4.

Table 3.2. Weight of detectors on the space frame.

Load on space frame	Weight (kg)
TPC	20 000
TRD	26 000
TOF	28 000
HMPID	3 500
Total	77 500

Table 3.3. Weight of equipment on back frame.

Load on back frame	Weight (kg)
TPC services	2 800
TRD services	7 000
TOF services	4 200
Aluminium ring	3 800
Total	17 800

Table 3.4. Weight of equipment on baby space frame.

Load on baby space frame	Weight (kg)
Mini frame	10 000
TRD services	7 000
TOF services	4 200
Total	21 200

These structures rest on about 12 m long rails, which span the length of the L3 solenoid. There is one pair of rails for the PHOS, a second pair for the space frames, and a third pair for the future electromagnetic calorimeter. The rails in turn are supported on the end pieces of the magnet yoke. The rails with removable extensions on the RB24 side allow the displacement of the space frames and detectors away from the intersection point for installation and maintenance. The support structures inside the L3 solenoid are produced from non-magnetic materials, either stainless steel or aluminium, to avoid distortions of the magnetic field. The space frame and the attached frames on both sides follow the 18-fold segmentation of the TPC. The beams of these structures fall as much as possible into the shadow of the insensitive regions between the readout chambers of the TPC.

3.3.5. Beam pipe. The vacuum pipe consists of a 0.8 mm thick, straight beryllium tube in the region from 3.5 m from the IP on the RB24 side to 0.4 m from the IP on the RB26 side. The outer diameter of the beam pipe is 59.6 mm. Outside this region the vacuum pipe is either a copper (RB24 side) or stainless-steel (RB26 side) tube.

3.4. Inner Tracking System (ITS)

The Inner Tracking System (ITS) consists of six cylindrical layers of silicon detectors, located at radii, $r = 4, 7, 15, 24, 39$ and 44 cm. It covers the rapidity range of $|\eta| < 0.9$ for all vertices located within the length of the interaction diamond ($\pm 1\sigma$), i.e. 10.6 cm along the beam direction. The number, position and segmentation of the layers are optimized for efficient track

finding and high impact-parameter resolution. In particular, the outer radius is determined by the necessity to match tracks with those from the Time-Projection Chamber (TPC), and the inner radius is the minimum allowed by the radius of the beam pipe (3 cm). The first layer has a more extended coverage ($|\eta| < 1.98$) to provide, together with the Forward Multiplicity Detectors (FMD), a continuous coverage in rapidity for the measurement of charged-particles multiplicity.

Because of the high particle density, up to $80 \text{ particles cm}^{-2}$, and to achieve the required impact parameter resolution, pixel detectors have been chosen for the innermost two layers, and silicon drift detectors for the following two layers. The outer two layers, where the track densities are below $1 \text{ particle cm}^{-2}$, will be equipped with double-sided silicon micro-strip detectors. With the exception of the two innermost pixel planes, all layers will have analogue readout for particle identification via dE/dx measurement in the non-relativistic ($1/\beta^2$) region. This will give the ITS a stand-alone capability as a low- p_T particle spectrometer.

3.4.1. Design considerations. The tasks of the ITS are:

- to localize the primary vertex with a resolution better than $100 \mu\text{m}$;
- to reconstruct the secondary vertices from decays of hyperons and D and B mesons;
- to track and identify particles with momentum below 100 MeV ;
- to improve the momentum and angle resolution for the high- p_t particles which also traverse the TPC;
- to reconstruct, albeit with limited momentum resolution, particles traversing dead regions of the TPC.

The ITS contributes to the global tracking of ALICE by improving the momentum and angle resolution obtained by the TPC. This is beneficial for practically all physics topics addressed by the ALICE experiment:

- Global event features are studied by measuring multiplicity distributions and inclusive particle spectra.
- For the study of resonance production (ρ , ω and ϕ), and in particular the behaviour of the mass and width of mesons in a dense medium, the momentum resolution is a key requirement. Precision on mass measurement must be comparable to, or better than, the natural width of the resonances to observe possible changes of their parameters caused by the predicted chiral symmetry restoration.
- Mass measurements for heavy-flavour states with better resolution improves the signal-to-background ratio in the study of heavy-quarkonia suppression, such as J/ψ and Υ .
- Improved momentum resolution enhances the performance in the observation of jet production and jet quenching, i.e. the energy loss of partons in strongly interacting dense matter. The coverage of the TPC inactive regions improves the ALICE capability to detect all particles in the jet cone.

For these studies, the PID capabilities of the ITS in the non-relativistic region are also of great help. Low-momentum particles (around $100 \text{ MeV } c^{-1}$) are detectable only by the ITS. This is of interest in itself, because it widens the momentum range for the measurement of particle spectra, important to study collective effects associated with large length scales. In addition, keeping the p_t cutoff low is essential to suppress soft γ conversions as well as the Dalitz background in the electron-pair spectrum.

In addition to the improved momentum resolution, the ITS provides an excellent double-hit resolution enabling the separation of tracks with close momenta.

Table 3.5. ITS material budget traversed by straight tracks perpendicularly to the detector surface. Units are percentages of radiation length.

Detector	Pixel		Drift		Strip	
	Inner	Outer	Inner	Outer	Inner	Outer
Layer	1.0	1.0	1.1	1.1	0.81	0.83
Thermal shield/support		0.36	0.29		0.42	
Total			7.0			

The requirement on angular resolution is determined by the need to precisely measure the position of the primary vertex and of the first points on the tracks. The accuracy of the measurement depends on the precision, position and thickness of the first ITS layer. Its design has been optimized to provide accurate measurements of track impact parameter enabling the identification on a statistical basis of secondary vertices of charm and beauty mesons and on a track-by-track basis the identification of hyperons.

The following factors were taken into consideration for the design of the ITS:

- *Acceptance:* The tracking system must have a sufficiently large rapidity acceptance to study, on event-by-event basis, particle ratios, p_t -spectra and particle correlations. According to current predictions, several thousands of particles will be emitted per heavy-ion collision within the rapidity range ($|\eta| < 0.9$) covered by the tracking system. Such a rapidity window is also necessary for the detection with a good efficiency of the decay of large mass, low transverse-momentum particles. An efficient rejection of low-mass Dalitz decays can only be implemented if the detector provides full azimuthal coverage. To extend the rapidity coverage of the multiplicity measurement, the first pixel layer has a pseudo-rapidity coverage ($|\eta| < 1.98$ for interactions taking place at $z = 0$).
- *dE/dx measurement:* The ITS contributes to particle identification through the measurement of specific energy loss. To apply a truncated-mean method, a minimum of four measurements are necessary, which implies that four layers out of the six need analogue readout. The dynamic range of the analogue readout must be large enough to provide dE/dx measurement for low-momentum, highly ionizing particles, down to the lowest momentum at which tracks can still be reconstructed with a reasonable ($>20\%$) probability.
- *Material budget:* The values of momentum and impact parameter resolution for particles with small transverse momenta are dominated by multiple scattering effects in the material of the detector. Therefore the amount of material in the active volume has to be reduced to a minimum. However, silicon detectors used to measure ionization densities (drift and strips) must have a minimum thickness of approximately $300 \mu\text{m}$ to provide reasonable signal-to-noise ratio. In addition, detectors must overlap to cover entirely the solid angle. Taking into account the incidence angles of tracks, the detectors effective thickness amounts to 0.4% of X_0 . The thickness of additional material in the active volume, i.e. electronics, cabling, support structure and cooling system, is limited to a comparable effective thickness (table 3.5). As discussed in section 5, the relative momentum resolution, which can be achieved with ITS, is better than 2% for pions with momentum between $100 \text{ MeV } c^{-1}$ and $3 \text{ GeV } c^{-1}$.
- *Spatial precision and granularity:* The granularity of the detectors is selected to cope with a maximum track density of 8000 tracks per unit of rapidity, the upper limit of current theoretical predictions. The ITS would therefore detect simultaneously more than 15 000 tracks. Keeping the system occupancy low, at the level of a few per cent, requires several millions of effective cells in each layer of the ITS. The resolution of the




Figure 3.5. Layout of the ITS.

impact-parameter measurement is determined by the spatial resolution of the ITS detectors. For charmed-particle detection the impact-parameter resolution must be better than $100 \mu\text{m}$ in the $r\varphi$ plane. Therefore, the ITS detectors have a spatial resolution of the order of a few tens of μm , with the best precision ($12 \mu\text{m}$) for the detectors closest to the primary vertex. In addition, for momenta larger than $3 \text{ GeV } c^{-1}$, relevant for the detection of the decay products of charmed mesons and high-mass quarkonia, the spatial precision of the ITS becomes an essential element of the momentum resolution. This requirement is met by all layers of the ITS with a point resolution in the bending plane about one order of magnitude better than that of the TPC, which in turn provides many more points.

- **Radiation levels:** Detailed calculations of the radiation level in ALICE are presented in section 2.6.2. For the ITS, the total dose received during the expected lifetime of the experiment varies from a few krad (tens of Gy) for the outer parts of the ITS to about 220 krad (2.2 kGy) for the inner parts (table 2.6). Each of the subdetectors is designed to withstand the ionizing radiation doses expected during ten years of operation. The neutron fluence is approximately $5 \times 10^{11} \text{ cm}^{-2}$ throughout the ITS, which does not cause significant damage to the detectors or the associated electronics. Where necessary, the components used in the ITS design were tested for their radiation hardness to levels exceeding significantly the expected doses.
- **Readout rate:** The ALICE system will be used in two basically different readout configurations, operated simultaneously with two different triggers. The centrality trigger activates the readout of the whole of ALICE, in particular all layers of the ITS, while the trigger of the muon arm activates the readout of a subset of fast readout detectors, including the two inner layers of the ITS. Therefore the readout time for the pixel detectors is set at less than $400 \mu\text{s}$.

3.4.2. ITS layout. The ITS (figure 3.5) 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. The detectors and front-end electronics are held by lightweight carbon-fibre structures.

The geometrical dimensions and the technology used in the various layers of the ITS are summarized in table 3.6. The granularity required for the innermost layers, is achieved with silicon micro-pattern detectors with true two-dimensional readout: Silicon Pixel Detectors (SPD) and Silicon Drift Detectors (SDD). At larger radii, the requirements in terms of granularity are less stringent, therefore double-sided Silicon Strip Detectors (SSD) with a

Table 3.6. Dimensions of the ITS detectors (active areas).

Layer	Type	r (cm)	$\pm z$ (cm)	Area (m^2)	Ladders	Lad./stave	Det./ladder	Channels
1	Pixel	3.9	14.1	0.07	80	4	1	3 276 800
2	Pixel	7.6	14.1	0.14	160	4	1	6 553 600
3	Drift	15.0	22.2	0.42	14	–	6	43 008
4	Drift	23.9	29.7	0.89	22	–	8	90 112
5	Strip	37.8/38.4	43.1	2.09	34	–	22	1 148 928
6	Strip	42.8/43.4	48.9	2.68	38	–	25	1 459 200
Total area				6.28				

Table 3.7. Parameters of the various detector types. A module represents a single sensor element.

Parameter	Silicon pixel	Silicon drift	Silicon strip
Spatial precision $r\varphi$ (μm)	12	38	20
Spatial precision z (μm)	100	28	830
Two track resolution $r\varphi$ (μm)	100	200	300
Two track resolution z (μm)	850	600	2400
Cell size (μm^2)	50×425	150×300	$95 \times 40\,000$
Active area per module (mm^2)	12.8×69.6	72.5×75.3	73×40
Readout channels per module	40 960	2×256	2×768
Total number of modules	240	260	1698
Total number of readout channels (k)	9835	133	2608
Total number of cells (M)	9.84	23	2.6
Average occupancy (inner layer) (%)	2.1	2.5	4
Average occupancy (outer layer) (%)	0.6	1.0	3.3
Power dissipation in barrel (W)	1500	1060	1100
Power dissipation end-cap (W)	500	1750	1500

small stereo angle are used. Double-sided micro-strips have been selected rather than single-sided ones because they introduce less material in the active volume. In addition, they offer the possibility to correlate the pulse height readout from the two sides, thus helping to resolve ambiguities inherent in the use of detectors with projective readout. The main parameters for each of the three detector types are summarized in table 3.7.

3.4.3. Silicon pixel layers

3.4.3.1. Design considerations. The two innermost layers of the ITS are fundamental elements for the determination of the position of the primary vertex as well as for measurement of the impact parameter of secondary tracks originating from the weak decays of strange, charm, and beauty particles. They will operate in a region where the track density could be as high as $80 \text{ tracks cm}^{-2}$. To cope with such densities, detectors of high precision and granularity are mandatory. In addition, the detector must be able to operate in a relatively high radiation environment: in the case of the inner layer, the integrated levels (10 years, standard running scenario) of total dose and fluence are estimated to be 220 krad (2.2 kGy) and $10^{12} \text{ n cm}^{-2}$ (1 MeV neutron equivalent), respectively.

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, a high segmentation leads naturally to a low individual diode capacitance, resulting in an excellent signal-to-noise ratio at high speed.

These considerations led us to select silicon pixel detectors for the two innermost layers of ALICE (SPD).

ALICE will operate with two types of level-2 (L2, final) triggers. The first one, selecting the event centrality, will run at a frequency of about 40 Hz and will trigger the readout of all ALICE detectors (central detectors and muon arm). For these events, the information provided by the two layers of the SPD barrel will be combined, in the global tracking process, with the information provided by the other ITS detectors and by the TPC. The second type of L2 trigger is based on the detection of di-muons in the muon arm. It will run at a frequency of about 1 kHz and will trigger the readout only of the SPD and of the muon arm. For these events, the information provided by the SPD will be used in a stand-alone mode to determine the position of the primary interaction vertex.

The ALICE SPD will employ hybrid silicon pixel detectors, consisting of a two-dimensional matrix (sensor ladder) of reverse-biased silicon detector diodes bump-bonded to readout chips. Each diode is connected through a conductive solder bump to a contact on the readout chip corresponding to the input of an electronics readout cell. The readout is binary: a threshold is applied to the pre-amplified and shaped signal and each cell outputs a logical 1 if the threshold is surpassed. This technique has already been successfully applied in the WA97 and NA57 experiments at CERN [401].

The development of the ALICE SPD has taken advantage of state-of-the-art technologies: a six-metal-layer CMOS deep submicron process ($0.25\ \mu\text{m}$ feature size) with radiation-tolerant layout design [402] for the readout chip and fine-pitch flip-chip techniques for bump-bonding the chips to the sensors.

Other key issues in the SPD are the material budget, cooling, transfer of signals from the front-end to the readout racks and controls. Substantial progress has been made in all these areas.

3.4.3.2. Detector layout. The basic building block of the ALICE SPD is a ladder. It consists of a silicon-sensor matrix bump-bonded to five front-end chips. The sensor matrix consists of 256×160 cells, each measuring $50\ \mu\text{m}$ in the $r\phi$ direction by $425\ \mu\text{m}$ in the z direction. Longer sensor cells are used in the boundary region to assure coverage between readout chips. Each ladder has a sensitive area of $12.8\ \text{mm}$ ($r\phi$) \times $69.6\ \text{mm}$ (z). Each front-end chip contains the electronics for the readout of a sub-matrix of 256 ($r\phi$) \times 32 (z) detector cells. The thickness of the sensor is $200\ \mu\text{m}$. It is the smallest one which can be achieved on 5-in wafers with an affordable yield. The target thickness of the readout chip is $150\ \mu\text{m}$. This requires the readout wafers to be thinned after bump deposition, before bump-bonding. The total silicon budget is therefore of the order of $350\ \mu\text{m}$.

Two ladders are mounted along the z direction to form a $144\ \text{mm}$ long half-stave (figure 3.6).

The half-staves are mounted on a carbon-fibre support and cooling sector (figure 3.7).

Each sector supports six staves: two from the inner layer and four from the outer layer. Ten sectors are then mounted together around the beam pipe to close the full barrel. In total, there will be 60 staves, 240 ladders, 1200 chips and 9.8×10^6 cells. The staves of the inner (outer) SPD layer are located at an average distance of about $3.9\ \text{cm}$ ($7.6\ \text{cm}$) from the beam axis (figure 3.7, bottom).

The data/control bus lines and the power/ground planes are implemented in an aluminium/polyimide multi-layer flex (pixel bus). This flex is a major technical challenge. The bus is glued to the sensors and wire-bonded to the readout chips. The readout of the two ladders on each half-stave is controlled by the digital PILOT ASIC and other auxiliary ASICs (the analogue PILOT and the gigabit optical link driver GOL) on a multi-chip-module (MCM)




Figure 3.6. Half-stave assembly—test prototype.





Figure 3.7. Top, carbon fibre support sectors assembly; bottom, carbon fibre support of the Si-pixel staves.




Figure 3.8. MCM prototype in ceramic multilayer.

located at the end of the half-stave. MCM prototypes have been developed and tested in FR4 and in ceramics (figure 3.8).

The technology chosen for production is sequential build up (SBU). The interface of the front-end electronics to the remote readout electronics (Router) in the counting room is done

Table 3.8. Main specifications of the ALICE SPD front-end chip.

Cell size	$50 \mu\text{m} (r\varphi) \times 425 \mu\text{m} (z)$
Number of cells	$256 (r\varphi) \times 32 (z)$
Minimum threshold	$1000e$
Threshold uniformity	$200e$
L1 latency	Up to $51 \mu\text{s}$
Operating clock frequency	10 MHz
Radiation tolerance	In excess of 100 kGy

through optical fibres. The digital PILOT chip performs the multiplexing of the outgoing data. The GOL chip serializes the outgoing data stream and drives a laser diode located in the optical module. This module also contains 2 PIN photo-diodes for (a) the incoming LHC 40 MHz clock and (b) the trigger and JTAG configuration data. The receiver amplifiers are located in the digital PILOT [403] which controls the pixel chips.

The power dissipated in the front-end electronics will be about 1.5 kW. The relatively high power density, together with the small mass of the detector, require a very efficient cooling system. Prototyping work has shown that good results can be obtained with evaporative cooling using C_4F_{10} . The sectors are equipped with cooling tubes embedded in the sector support and running underneath the staves (one per stave). The heat transfer from the front-end chips is assured with grease of good thermal conductivity.

To avoid radiation of heat towards the SDD layers, which are very sensitive to temperature changes, an Al-coated carbon-fibre external shield surrounds the SPD barrel.

The average material traversed by a straight track perpendicular to the beam line crossing the SPD barrel corresponds to about $2\% X_0$.

3.4.3.3. Front-end electronics—Fast-OR. The pixel readout chip (ALICE1LHCb) is a mixed-signal ASIC for the readout of 8192 pixels [404].

Each pixel cell contains a preamplifier-shaper with leakage current compensation, followed by a discriminator. A signal above threshold results in a logical 1 which is propagated through a delay line during the $6 \mu\text{s}$ latency time until the arrival of the L1 trigger signal.

A four-hit-deep front-end buffer on each cell allows derandomization of the event arrival times. Upon arrival of the L1 trigger, the logical level present at the end of the delay line is stored in the first available buffer location.

All parameters of the pixel chip are programmable. The on-chip global registers include 42 8-bit DACs that adjust current and voltage bias references, L1 trigger delay, global threshold voltage, and leakage compensation. In each pixel cell a 3-bit register allows individual tuning of the threshold; there is also provision to enable the test pulse input and to mask the cell. All configuration parameters are controlled through the JTAG bus via the digital PILOT chip.

The front-end pixel chip has been designed in standard $0.25 \mu\text{m}$ CMOS using a radiation-tolerant layout technique (enclosed gate geometry). The design has proven to be insensitive to total ionization dose (TID) exceeding 10 Mrad (100 kGy). The cross section for single-event effects (SEU) has been measured to be approximately $3 \times 10^{-16} \text{ cm}^2$, corresponding to an upset rate of 0.1 bit h^{-1} in the full barrel for Pb–Pb collisions at nominal luminosity.

The specifications of the ALICE SPD front-end chip are summarized in table 3.8.

The outputs of the discriminators in the pixel cells of the ALICE1LHCb chip provide a fast-OR digital pulse when one or more pixels are hit on the chip. The fast-OR is an invaluable tool in testing and allows implementing a unique triggering capability in the SPD. Provision has




Figure 3.9. Fast-OR processing.

been made in the front-electronics to output the fast-OR signals off-detector so that they can be processed and contribute to the L0 trigger. This is particularly interesting in the case of events with very low multiplicities in pp runs. However, certain constraints deriving from internal design have to be taken into account. The fast-OR is synchronous with the 10 MHz pixel system clock, hence the signal is integrated over 100 ns or 4 bunch-crossings. Thus the pixel trigger signal cannot identify the bunch-crossing in which the corresponding collision occurred but can only point to a string of consecutive bunch crossings. Furthermore, if a collision occurs only one to three bunch-crossings before the rising edge of the 10 MHz system clock, the fast-OR might be recorded in the following system clock cycle, or in the one afterwards. Therefore the fast-OR signal response time jitters up to 100 ns which enlarges the uncertainty window to up to eight consecutive bunch-crossings or 200 ns. To limit this ambiguity as much as possible when all fast-OR signals of 120 half-staves are merged together and also to minimize the latency, all half-staves are synchronized to an integration period which covers the same bunch-crossings. Using the ALICE trigger system, the bunch-crossing ambiguity can be resolved by performing a logical AND between the pixel trigger signal and the ALICE T0 detector signal.

On each half-stave, the 10 fast-OR signals on the bus are merged in the outgoing data stream of the pixel detector in the PILOT chip [403] and transferred every 100 ns via the optical link (figure 3.9).

The digital PILOT design has been modified to implement this additional functionality.

The optical fibres carrying the outgoing data stream will be equipped with two-way optical splitters; readout and trigger data are extracted at the readout crate and at the trigger crate, respectively. The off-detector pixel trigger electronics will include 120 G-Link-compatible optical receivers and deserializer modules.

Trigger data should reach the trigger crate within ~ 800 ns [405] in order to meet the latency requirements of the L0 trigger electronics. Short connections and minimal processing overhead are essential. The processor will be based on fast FPGAs and will feed the ALICE trigger crate directly.

First estimates indicate that with processing limited to the strict minimum, and with the expected length of the optical fibres (43 m), the pixel L0 trigger signal can only be delivered within about 900 ns, that is approximately 100 ns later than the prescribed latency. The implications are being investigated.

The off-detector fast-OR electronics and optical components are at this moment regarded as a future upgrade, for which funding and manpower have to be found.

3.4.3.4. Readout. Upon arrival of the second level trigger (L2), the data contained in the front-end buffer locations corresponding to the first (oldest) L1 trigger are loaded onto the

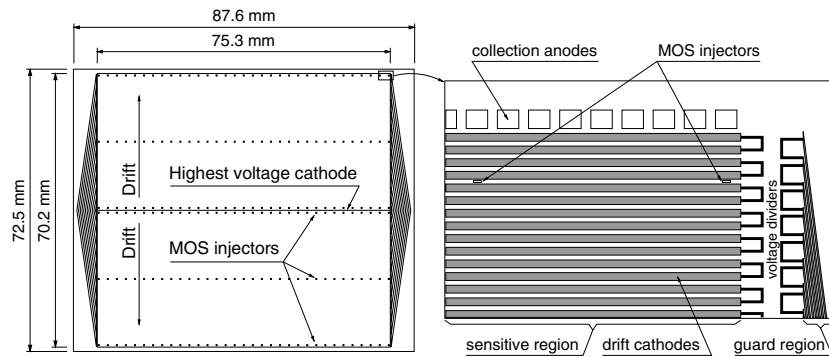


Figure 3.10. Layout of the ALICE SDD. The sensitive area is split into two drift regions by the central, highest voltage, cathode. Each drift region has one row of 256 collection anodes and three rows of 33 point-like MOS charge injectors for monitoring the drift velocity. Drift and guard regions have independent built-in voltage dividers.

output shift registers. Then, for each chip, the data from the 256 rows of cells are shifted out during 256 cycles of a 10 MHz clock. At each cycle, a 32-bit word containing the hit pattern from one chip row is output on the 32-bit stave data bus where it is read out by the PILOT chip. One pixel chip is read out in $25.6 \mu\text{s}$. The 10 chips of two ladders (one half-stave) are read out sequentially in a total time of about $256 \mu\text{s}$. This is also the readout time of the full SPD since the 120 half-staves are read out in parallel. The dead time introduced by the readout of the SPD is estimated to be less than 10% in the worst case, corresponding to ALICE running with Ar–Ar beams at high luminosity, with an L1 rate of 2.5 kHz.

At the receiving end of the optical link, the data are converted and deserialized in the Router module. Zero suppression and a second level of multiplexing are performed on the data which are then transferred to the DAQ on 20 DDL optical links.

3.4.4. Silicon drift layers. The Silicon Drift Detectors (SDD) will equip the two intermediate layers of the ITS [388], where the charged particle density is expected to reach up to 7 cm^{-2} . They have a very good multitrack capability and provide two out of the four dE/dx samples needed for the ITS particle identification.

3.4.4.1. Detector layout. The ALICE SDDs will be produced from very homogeneous high-resistivity ($3 \text{ k}\Omega \text{ cm}$) $300 \mu\text{m}$ thick Neutron Transmutation Doped (NTD) silicon [406]. As shown in figure 3.10, they have a sensitive area of $70.17 \times 75.26 \text{ mm}^2$ and a total area of $72.50 \times 87.59 \text{ mm}^2$. The sensitive area is split into two drift regions by the central cathode strip to which a nominal bias of -2.4 kV is applied. In each drift region, and on both detector surfaces, 291 p^+ cathode strips, with $120 \mu\text{m}$ pitch, fully deplete the detector volume and generate a drift field parallel to the wafer surface. To keep the biasing of the collection region independent on the drift voltage, a second bias supply of -40 V is added. The degrading of the high voltage to the zero potential of the detector boundary is implemented by two insensitive guard regions biased by 145 cathode strips with $32 \mu\text{m}$ pitch. To improve the detector reliability, all the drift and guard regions have their own built-in voltage dividers. Their total power dissipation is 1 W per detector and will be removed by an appropriate air circulation system. Each drift region has 256 collection anodes with $294 \mu\text{m}$ pitch and three rows of 33 point-like ($20 \times 100 \mu\text{m}^2$) MOS charge injectors to monitor the drift velocity which

Table 3.9. The main characteristics of the ALICE silicon drift detectors.

Sensitive area	$70.17 \times 75.26 \text{ mm}^2$
Total area	$72.50 \times 87.59 \text{ mm}^2$
Collection anodes (readout channels)	2×256
Anode pitch	$294 \mu\text{m}$
Nominal operating voltage	-2.4 kV
Nominal bias of the collection region	-40 V
Nominal drift velocity	$8.1 \mu\text{m ns}^{-1}$
Nominal maximum drift time	$4.3 \mu\text{s}$
Cell size at nominal drift velocity	$294 \times 202 \mu\text{m}^2$
Cells per detector at nominal drift velocity	$2 \times 256 \times 174$
Total number of cells (260 SDDs)	$23 \times 16 \times 10^6$
Average resolution along the drift ($r\varphi$)	$35 \mu\text{m}$
Average resolution along the anode (z)	$25 \mu\text{m}$
Detection efficiency	99.5%
Average double-track resolution at 70% efficiency	$700 \mu\text{m}$

depends on temperature: $v_{\text{drift}} \propto T^{-2.4}$ [407]. They will be triggered at regular intervals during the long gap between the LHC orbits.

At the nominal bias voltage of -2.4 kV the drift velocity is $8.1 \mu\text{m ns}^{-1}$. Since the front-end electronics samples the signal of each anode at a frequency of 40.08 MHz , the size of the sensitive element (cell) is $294 \times 202 \mu\text{m}^2$, corresponding to 89.1×10^3 cells per detector, which are readout by 512 channels.

The space precision along the drift direction ($r\varphi$), as obtained during beam tests of full-size prototypes, is better than $38 \mu\text{m}$ over the whole detector surface. The precision along the anode axis (z) is better than $30 \mu\text{m}$ over 94% of the detector surface and reaches $60 \mu\text{m}$ close to the anodes, where a smaller fraction of clusters affect more than one anode. The average values are 35 and $25 \mu\text{m}$ respectively [408]. The detection efficiency is larger than 99.5% for amplitude thresholds as high as 10 times the electronic noise.

The double-track resolution was calculated using the detailed simulation program, included in AliRoot, after its preliminary tuning based on the single-track results from beam tests [409]. Due to charge diffusion during the drift process, the double track resolution is a function of the drift time for a given separation efficiency. The relative distance at which two clusters are disentangled with a 70% efficiency grows almost linearly from $600 \mu\text{m}$ near the anodes to $800 \mu\text{m}$ at the maximum drift distance [410]. The main parameters of the ALICE SDD are summarized in table 3.9.

SDD ladders. The SDDs are mounted on linear structures called ladders. There are 14 ladders with six detectors on layer 3, and 22 ladders with eight detectors on layer 4. Detectors and ladders are assembled (figure 3.11) to have an overlap of the sensitive areas larger than $580 \mu\text{m}$ in both $r\varphi$ and z directions. This ensures full angular coverage for vertices located in the interaction diamond, $\pm\sigma = 10.6 \text{ cm}$, and for $35 \text{ MeV } c^{-1} \leq p_t \leq 2.8 \text{ GeV } c^{-1}$.

The main geometrical parameters of the SDD layers and ladders are summarized in table 3.10. The ladder space frame is a lightweight triangular truss made of Carbon-Fibre Reinforced Plastic (CFRP) and has a protective coating against humidity absorption. The sagging of the space frame under the load of the ladder components applied to the mid point of the truss was evaluated with a Finite Element Analysis (FEA) model. When the load acts perpendicularly to the detector plane the sagging is $21 \mu\text{m}$, when it acts parallel to the detector plane the sagging is $7 \mu\text{m}$. The sag measured on prototypes of the layer-4 ladders is




Figure 3.11. The SDDs are mounted at different radii in both r_z and r_ϕ planes to obtain the full coverage in the acceptance region. Units are millimetres.

Table 3.10. The main parameters of the ALICE SDD layers and ladders.

	Layer 3	Layer 4
Detectors per ladder	6	8
Ladders per layer	14	22
Detectors per layer	84	176
Ladder sensitive half-length (cm)	22.16	29.64
Ladder length (cm)	45.56	60.52
Average layer radius (cm)	15.03	23.91
Ladder space-frame weight (g)	11	15
Weight of ladder components (g)	87	121

$33.3 \pm 1.5 \mu\text{m}$. Sags smaller than $40 \mu\text{m}$ are accepted. The difference with the simulations is explained by uncertainties in the characteristics of the carbon fibre pre-pregs and by the fact that in the simulations the ladder is a ‘restrained beam’ while, during the tests, it is a quasi-restrained beam, since the rotational degree of freedom transverse to the ladder is not well satisfied (quasi-hinge constraint).

The ladders are assembled on a CFRP structure made of a cylinder, two cones and four support rings (figure 3.12). The cones provide the links to the outer Silicon Strip Detector barrel and have windows for the passage of the SDD services. The support rings are mechanically fixed to the cones and bear reference ruby spheres for the ladder positioning. The accuracy of the re-positioning of the ladders was measured to be better than $10 \mu\text{m}$ full width.

The detectors are attached to the space frame using ryton pins and have their anode rows parallel to the ladder longitudinal (z) axis. The front-end electronics, assembled on two hybrid circuits, one per anode row, is glued with a thermo-conductive compound on rigid carbon fibre heat-exchangers, which in turn are clipped to the cooling pipes running along the ladder structure (figure 3.13). The pipes are made of phynox and have a $40 \mu\text{m}$ wall thickness and a 2 mm inner diameter. The coolant is demineralized water. The connections between the detectors, the front-end hybrids and the end-ladder boards, and the connection for the detector biasing are all assured with flexible aluminium–polyimide micro-cables that are Tape Automatic Bonded (TAB). These micro-cables have typical polyimide and aluminium thicknesses of $12\text{--}20 \mu\text{m}$ except for the $400 \mu\text{m}$ thick high-voltage cable. Each detector will

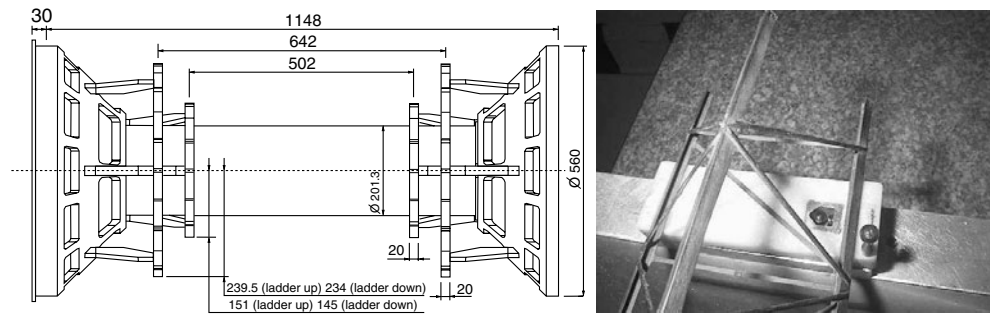


Figure 3.12. Left: support structure of the SDD ladders; units are millimetres. Right: ladder foot with its V-shaped nut and the reference ruby sphere which is pressed against the nut by a spring. The ruby sphere glued onto the ladder foot is used to verify the ladder position once the assembly of a layer is completed.




Figure 3.13. Cut-view of the ladder cross section showing the front-end micro-cables and the cooling tube.

be first assembled with its front-end electronics, the high-voltage connections and the related end-ladder boards as a unit, called an SDD module, which shall be fully tested before mounting it on the ladder. When a ladder is assembled and the hybrids are not yet clipped to the cooling pipes, the positions of the detectors will be measured with respect to the reference ruby spheres glued to the ladder feet.

The end-ladder boards provide the interface between the SDD module and the external sub-systems: DAQ, trigger, Detector Control System (DCS), and low- and high-voltage supplies. Each hybrid circuit has its own low-voltage board, carrying the rad-hard low-voltage regulators, the LVDS (Low-Voltage Differential Signalling) signal-receivers and the interface with the DCS. Each detector has its own high-voltage board, containing the filtering of the high-voltage bias, the driver of the MOS injector lines and, possibly, an external voltage divider. Since the power dissipation per unit length is much higher in the end-ladder than in the ladder, the end-ladder has an independent water-based liquid cooling circuit.

The material budget amounts to $X/X_0 = 1.13\%$ per SDD layer, to which 0.29% has to be added for the cylinder of the supporting structure; these values have been calculated for normal incidence tracks.

3.4.4.2. Front-end electronics and readout. The SDD front-end electronics is based on three ASICs. The first one, PASCAL assembled on the front-end hybrid, contains three functional blocks: preamplifier, analogue storage and Analogue-to-Digital Converter (ADC) [411]. The second integrated circuit, AMBRA, also on the hybrid, is a digital four-event buffer which performs data derandomization, baseline equalization on an anode-by-anode basis and a 10- to 8-bit non-linear data compression. AMBRA also sends the data to the third ASIC, CARLOS, which is a zero-suppressor and data-compressor mounted in one of the end-ladder boards [412, 413]. The three ASICs have been designed using radiation-tolerant layout technique (enclosed gate geometry) based on a commercial deep-submicron process ($0.25\ \mu\text{m}$). The PASCAL prototype designed with this technology has proved to be insensitive to total ionization dose up to 300 kGy. The average power dissipation of each PASCAL–AMBRA front-end channel is estimated to be about 6 mW.

The signal generated by an SDD anode feeds the PASCAL transimpedance amplifier-shaper which has a peaking time of about 40 ns and a dynamic range of 32 fC (the charge released by an 8-MIP particle hitting near the anode). The amplifier output is sampled at 40.08 MHz by a ring analogue memory with 256 cells per anode. This is the mode of operation in the idle state of the front-end. On a run-by-run basis, PASCAL can be programmed to use half of this nominal frequency for the sampling, thus reducing the the amount of data and, therefore, the sub-system dead-time. Analysis of beam test data and simulation have shown that the cost in terms of both spatial resolution and double track-resolution is negligible. When an L0 trigger is received, the SDD BUSY is immediately set and, after a programmable delay which accounts for the L0 latency ($1.2\ \mu\text{s}$) and the maximum detector drift time ($\sim 5\ \mu\text{s}$), the analogue memories are frozen. The BUSY being still set, their contents are then digitized by a set of 10-bit linear successive-approximation ADCs which write the data into one of the free AMBRA buffers. The advantages of a front-end A/D conversion are the noise immunity during signal transmission, and the possibility of inserting a multiple-event buffer to derandomize the data and, therefore, to slow down the transfer rate to the DAQ system. This greatly reduces the material budget of the cabling. The digitization lasts for about $230\ \mu\text{s}$ ($120\ \mu\text{s}$ when the half-frequency mode is programmed) and can be aborted by the absence of the L1 trigger or by the arrival of an L2-reject signal; in both cases, the front-end electronics reset the SDD BUSY and returns to the idle state within 100 ns. On the successful completion of the analogue-to-digital conversion the SDD BUSY is reset if at least one buffer is still available in the AMBRAs. As soon as the conversion is completed, all the AMBRAs transmit the data in parallel to the CARLOS chips on the end-ladders, an operation which takes 1.24 ms ($0.62\ \text{ms}$ when the half-frequency mode is programmed). By means of a two-dimensional two-thresholds algorithm and with no additional dead time, the CARLOS chips reduce the SDD event size from the raw 22.1 MB by more than one order of magnitude. They also embody the trigger information in the data flow, format data and feed the GOL ASICs which in turn drive the optical links. In the counting room, several dedicated VME boards, CARLOS-rx, concentrate the GOL output fibres into 12 DDL (Detector Data Link) [414] channels. CARLOS-rx also downloads in parallel the configuration parameters received over the DDL to the ladder electronics it controls, and monitors the error-flag words embedded in the data flow by the CARLOS chips in order to signal potential Single-Event Upsets (SEU) on the ladder electronics.

To allow the full testability of the readout electronics at the board and system levels, the three ASICs embody a JTAG standard interface. In this way it is possible to test each chip after the various assembly stages. The same interface is used to download control information into the chips before the data taking.

The SDD dead time was evaluated with a simulation program taking into account the ADC conversion time, the AMBRA transmission time, the past–future protection of the TPC and




Figure 3.14. SSD detection module.

the availability of AMBRA buffers. With an interaction rate of 1700 Hz and with input trigger rates of 1360 Hz (L0), 570 Hz (L1) and 400 Hz (L2 accepted) the SDD are busy 10.5% of the time and the rates of trigger signals delivered to the SDD are: 1110 Hz (L0), 563 Hz (L1) and 351 Hz (L2a). The fourth AMBRA buffer is used 1.77% of the time.

SDD occupancy and event size. The average cell-occupancies were calculated on the basis of simulations performed within the ALICE offline framework AliRoot [398], using a drift velocity of $7.3 \mu\text{m ns}^{-1}$ corresponding to 192 detector cells per anode. For central Pb–Pb collisions a charged-particle multiplicity density of $dN_{\text{ch}}/d\eta_{\eta=0} = 8000$ was assumed. The simulation of the response of the electronics channels was performed using $\text{r.m.s.}_{\text{noise}} = 1.6$ ADC counts ($\text{ENC} = 320e$) and a safe threshold of twice the r.m.s. of the noise. For one Pb–Pb central event the obtained average occupancies for layers 3 and 4 are 6.46 and 3.06% respectively. The average values obtained for 1000 pp events are 2.08×10^{-4} and 1.01×10^{-4} for layers 3 and 4 respectively. Secondary particles are included in both cases.

From these occupancy values SDD event sizes of 1 MB and of 3.4 kB can be deduced for central Pb–Pb and pp interactions respectively. To obtain the real event size we must add 130 kB for monitoring the baseline and the noise of the channels, and then increase the result by 0.04% for the cluster addressing and by an additional 6% for the data formatting overhead. This gives 1.21 MB and 141 kB for Pb–Pb and pp, respectively. An empty event requires about 138 kB.

3.4.5. Silicon strip layers

3.4.5.1. Design considerations. The outer layers of the ITS are crucial for the connection of tracks from the TPC to the ITS. They also provide dE/dx information to assist particle identification for low-momentum particles. Both outer layers consist of double-sided Silicon Strip Detectors (SSD), mounted on carbon-fibre support structures identical to the ones which support the SDD. The detection module consists of one detector connected to two hybrids (see figure 3.14) featuring two layers of aluminium tracks [415]. Each hybrid carries six front-end chips (HAL25). The system is optimized for low mass in order to minimize multiple scattering.

Table 3.11. SSD system parameters.

Sensor active area	$73 \times 40 \text{ mm}^2$
Sensor total area	$75 \times 42 \text{ mm}^2$
Number of strips per sensor	2×768
Pitch of sensors on a ladder	39.1 mm
Strip pitch on a sensor	95 μm
Strip orientation p side	7.5 mrad
Strip orientation n side	27.5 mrad
Spatial precision $r\varphi$	20 μm
Spatial precision z	820 μm
Two track resolution $r\varphi$	300 μm
Two track resolution z	2400 μm
Radius layer 5 (lowest/highest)	378/384 mm
Radius layer 6 (lowest/highest)	428/434 mm
Number of ladders layer 5	34
Number of ladders layer 6	38
Modules per ladder layer 5	22
Modules per ladder layer 6	25
Number of modules layer 5	748
Number of modules layer 6	950
Material budget SSD cone	$0.28X_0$
Material budget per SSD layer	$0.81X_0$ (layer 5), $0.83X_0$ (layer 6)

3.4.5.2. Detector layout. The SSD sensors are double-sided strip detectors with a 35 mrad stereo angle. This small stereo angle has been selected to limit the number of ambiguities resulting from the expected high particle densities. Each detector has an active area of $73 \times 40 \text{ mm}^2$. The strips on one side of the sensor are oriented at an angle of 7.5 mrad with respect to the beam direction and with an angle of 27.5 mrad with respect to the beam direction on the other side. Mounting the sensors with the n or p side facing the interaction region in layers 5 and 6 respectively results in four different orientations of the strips with respect to the beam direction. This reduces the fake track probability significantly, resulting in much more robust tracking. For layer 5(6), the p(n)-side is facing the beamline. The active area on each sensor is surrounded by bias and guard rings which occupy 1 mm along each side of the sensor. The sensors are mounted with the strips nearly parallel to the magnetic field, so that the best position resolution is obtained in the bending direction. The modules are cooled using thin wall (40 μm) phynox tubes [416]. The main parameters of the system are summarized in table 3.11.

3.4.5.3. Front-end electronics and readout. Twelve ASIC front-end chips (HAL25) [417] are connected to each sensor. The HAL25 contains 128 identical channels. It amplifies, shapes and stores analogue signals from the sensor and provides an interface to the readout system.

The pre-amplifier converts the charge input from the sensor into an analogue voltage step whose magnitude is a function of the charge. It uses an integrator circuit whose shaping time is adjustable between 1.4 and 2.2 μs . Set-up and test of the HAL25 are performed using programmable registers accessed via an on-chip JTAG control interface. The analogue signals are stored in a sample-hold circuit, controlled by an external HOLD signal. This HOLD signal is derived from the L0 trigger signal, adequately delayed to match the shaping time of the HAL25.

Once the analogue signals are stored they can be read out serially, at a maximum specified speed of 10 MS s^{-1} through an analogue multiplexer provided in the HAL25. A fast clear

Table 3.12. SSD system parameters.

Number of HAL25 ASIC	20 376
Total number of channels	2 608 128
Number of DDL	16
Maximum hit rate per strip	50 Hz
Maximum L0 rate	75 kHz
Maximum L1 rate	6 kHz
Maximum L2 rate at 10% occupancy	1 kHz
Data volume/event at 10% occupancy	1 MB

input allows abortion of the readout sequence when the event is rejected at the L1 trigger level. The linear range of the analogue output corresponds to approximately 13 MIP.

Electronics located at the end of the ladders [418] will distribute the power to the front-end chips. It will protect each individual module against latch-up by shutting down the power supply in case of too large power supply currents. The end-cap will provide an interface to the detector control system to signal a power trip. The detector control system will be able to reset the power supply to the module and to reload the front-end chips through the JTAG interface. The end-cap electronics will also buffer all signals to and from the front-end module. In addition it isolates the front-end chips, which are connected to the sensor bias voltage, from the rest of the readout system which has a separate ground. The Front-End ReadOut Modules (FEROM) are located in eight crates besides the L3 magnet. Its distance to the end-cap electronics is about 25 m. This system controls the readout of the front-end chips and it digitizes the analogue data from all modules in parallel. The sensor itself is very fast, producing analogue signals of only a few tens of nanoseconds width. Therefore the sensor does not impose significant limitations on the trigger rate. However, the various other components in the system limit the event rate and trigger rates in several ways. The average hit rate per HAL25 input should be limited to less than 50 s^{-1} because of the time needed to discharge the integrator.

Upon receipt of an L0 trigger the HOLD signal is applied to the HAL25 chips. The past-future protections for this trigger should be $10 \mu\text{s}$ (past) and $3 \mu\text{s}$ (future) if pile-up of events in the front-end electronics is to be avoided. The HOLD signal is removed from the front-end chips if the event is rejected at the L1 trigger level. The SSD system will not accept additional L0 triggers until $1 \mu\text{s}$ after the L1 reject decision. In addition the L0 (HOLD) rate must be less than 75 kHz.

Upon receipt of an L1 trigger the digitization of the analogue signals stored in the front-end chips is started. This conversion takes a fixed amount of time, $160 \mu\text{s}$. During this time, the system cannot accept additional L0 or L1 triggers. However, if an L2-reject signal is received then the readout will be aborted within $1 \mu\text{s}$. After this time the system can accept new L0 triggers.

If an L2-accept trigger is received then the digitization is completed. Thus, $165 \mu\text{s}$ after the corresponding L0 trigger the system is able to receive new (L0) triggers. Digitized and zero-suppressed data are stored in a multiple-event buffer, allowing a new L0–L1–L2 cycle to start before all data are transferred to the DAQ system. For each signal (after zero-suppression) 4 B of data are transferred to the DAQ system. The average data transfer rate is limited to 800 MB s^{-1} , corresponding to approximately 20 MS s^{-1} . Each FEROM is connected to the DAQ system through two optical links (DDL). The busy signal generated by the FEROM will automatically limit the L2 rate without consequences for the data integrity. The readout system parameters are summarized in table 3.12.

3.5. Time-Projection Chamber (TPC)

3.5.1. Design considerations. The Time-Projection Chamber (TPC) [393] is the main tracking detector of the ALICE central barrel and, together with the other central barrel detectors has to provide charged-particle momentum measurements with good two-track separation, particle identification, and vertex determination. The phase space covered by the TPC ranges in pseudo-rapidity $|\eta| < 0.9$ (up to $|\eta| \sim 1.5$ for tracks with reduced track length and momentum resolution); in p_t up to $100 \text{ GeV } c^{-1}$ is reached with good momentum resolution. In addition, data from the central barrel detectors will be used to generate a fast online High-Level Trigger (HLT) for the selection of low cross-section signals.

All these requirements need to be fulfilled at the Pb–Pb design luminosity, corresponding to an interaction rate of 8 kHz, of which about 10% are to be considered as central collisions. For these we assume the extreme multiplicity of $dN_{\text{ch}}/d\eta = 8000$, resulting in 20 000 charged primary and secondary tracks in the acceptance, an unprecedented track density for a TPC. These extreme multiplicities set new demands on the design which were addressed by extensive R&D activities; test beam results show a good performance even at the highest anticipated multiplicities, see for details [393, 424]. Careful optimization of the TPC design finally resulted in maximum occupancies (defined as the ratio of the number of readout pads and time bins above threshold to all pads and time bins) of about 40% at the innermost radius and 15% at the outermost radius. Substantial improvements in the tracking software were necessary to achieve adequate tracking efficiency under such harsh conditions; tracking efficiencies $\geq 90\%$ were obtained for primary tracks at the time of the TPC TDR [393], and they have improved further recently as will be described in section 5 of the PPR Volume II.

For proton–proton runs, the memory time of the TPC is the limiting factor for the luminosity due to the $\sim 90 \mu\text{s}$ drift time. At a pp luminosity of about $5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$, with a corresponding interaction rate of about 350 kHz, ‘past’ and ‘future’ tracks from an average of 60 pp interactions are detected together with the triggered event; the detected multiplicity corresponds to about 30 minimum-bias pp events. The total occupancy, however, is lower by more than an order of magnitude than in Pb–Pb collisions, since the average pp multiplicity is about a factor 10^3 lower than the Pb–Pb multiplicity for central collisions [382]. Tracks from pile-up events can be eliminated because of their wrong vertex pointing.

3.5.1.1. Hadronic observables. The TPC is the main detector for the study of hadronic observables in both heavy-ion and pp collisions. Hadronic measurements give information on the flavour composition of the particle-emitting source via the spectroscopy of strange and multi-strange hadrons, on its space–time evolution and extent at freeze-out via single- and two-particle spectra and correlations, and on event-by-event fluctuations.

Correlation observables place the highest demands on relative momentum and two-track resolution. For event-by-event analyses, large rapidity and p_t acceptance are essential for the study of space–time fluctuations of the decomposing fireball. For a detailed analysis of kaon spectra and the kaon-to-pion ratio on an event-by-event basis, more than 100 analysed kaons are required, supporting the need for large acceptance and good particle identification. In addition, event-plane reconstruction and flow studies require close to 2π azimuthal acceptance.

Hard probes, such as heavy quarkonia, charmed and beauty particles, and high- p_t jets, require very good momentum resolution at high momenta, which has to be achieved with the help of other tracking detectors. A large acceptance is also beneficial because of low production cross sections.

Specific requirements on the TPC for hadronic observables are the following:

- **Momentum resolution:** For the study of soft hadronic observables, momentum resolution on the level of 1% is required for momenta as low as $100 \text{ MeV } c^{-1}$. On the other hand, hard

probes set the requirements for the high- p_t region. As will be discussed in section 5 of the PPR Volume II, the momentum resolution for low-momentum tracks (between $100 \text{ MeV } c^{-1}$ and $1 \text{ GeV } c^{-1}$) reconstructed in the TPC is between 1 and 2%, depending on the magnetic field setting. To measure higher momenta it is necessary to use the TPC in combination with the other tracking detectors (ITS and TRD). Using these detectors we can achieve about 10% momentum resolution for tracks with p_t of $100 \text{ GeV } c^{-1}$ at 0.5 T magnetic field which opens the possibility to extend the ALICE physics programme into the high- p_t region.

- *Two-track resolution:* In order to measure two-particle correlations efficiently (e.g. HBT, see section 1) we have to be able to detect tracks with very small momentum differences. The required resolution on the momentum difference depends on the actual source size, and has to be about $5 \text{ MeV } c^{-1}$ or better for a source size of 20 fm. For the ‘side’ and ‘long’ components of the momentum difference (see section 1.3.3) the achieved resolution is much better than what is required. However, for the ‘out’ component, especially at higher p_t (about $1 \text{ GeV } c^{-1}$), we will have to run at 0.5 T magnetic field in order to satisfy this requirement.
- *dE/dx resolution:* Particle identification by dE/dx measurement in the low-momentum ($1/\beta^2$) region is achieved in certain momentum intervals, where the expected ionization for different particle types is well separated. In order to cover the whole momentum range, up to a few $\text{GeV } c^{-1}$, we therefore need additional measurements (e.g. TOF detector). For even higher momenta, up to a few tens of $\text{GeV } c^{-1}$ (i.e. in the relativistic rise region), we will be able, at least on a statistical basis, to separate different hadron species if the dE/dx resolution is better than 7% [382]. As will be discussed in section 5 of the PPR Volume II, according to the present simulations, the resolution of the ionization measurements depends on the particle density, and a value of 6.9% is reached for the extreme multiplicities.
- *Track matching:* Efficient track matching with the ITS detector is necessary to measure the track impact parameter at the interaction point and for secondary vertex reconstruction. In addition, matching with other tracking detectors will improve the momentum resolution significantly relative to stand-alone TPC reconstruction for tracks with momentum above a few $\text{GeV } c^{-1}$ (e.g. an improvement by a factor 5 is expected for $10 \text{ GeV } c^{-1}$ tracks). In order to increase the matching efficiency the thickness of the material between the detectors (i.e. the TPC field-cage and containment vessels) has to be kept to a minimum.
- *Azimuthal coverage:* Full azimuthal coverage is necessary for global analysis of the event, such as the determination of the event plane, or flow analysis. Other signals, especially those limited by statistics, will benefit from a large azimuthal acceptance.

3.5.1.2. Leptonic observables. The physics objectives of the ALICE central barrel [382] have been extended with the addition of the Transition Radiation Detector (TRD [384]). As a consequence, the performance and corresponding design criteria for the TPC, and for the other central barrel detectors, had to be reassessed and optimized taking into account the requirements for electron physics. In particular, the TPC design was optimized to provide the largest possible acceptance for full-length, high- p_t tracks, in order to ensure significant statistics and good momentum resolution for high-mass and high- p_t electron pairs. Therefore, the inactive areas between the readout chambers have been aligned towards the centre of the TPC.

Electrons identified by the central barrel tracking detectors whose impact parameters are determined using the ITS will be used to measure charmed- and beauty-particle production. Moreover, impact-parameter measurement can be used to separate directly produced J/ψ mesons from those produced in B-decays.

Specific requirements on the TPC for electron physics are as follows:

- *Tracking efficiency:* The tracking efficiency for tracks with $p_t > 1 \text{ GeV } c^{-1}$ that enter the TRD should be as large as possible. As will be discussed in section 5 of the PPR Volume II, combined track finding in the central barrel detectors (ITS–TPC–TRD) has an efficiency well above 90% for these momenta.
- *Momentum resolution:* The momentum resolution for electrons with transverse momentum around $5 \text{ GeV } c^{-1}$ should be better than 1.5% in order to keep the electron-pair mass resolution below 1% and thereby be able to resolve the members of the Υ family. This resolution is achieved using all the barrel tracking detectors (ITS–TPC–TRD) in combination and when running at 0.5 T magnetic field.
- *dE/dx resolution:* To achieve the required pion rejection factor (better than 200 at 90% electron identification efficiency) for momenta above $1 \text{ GeV } c^{-1}$, the TRD electron identification has to be complemented by the TPC dE/dx measurement, especially in the lower part of the momentum range, up to $3 \text{ GeV } c^{-1}$. In order to satisfy this requirement, the TPC must provide a dE/dx resolution better than 10% in the high-multiplicity environment of central Pb–Pb collisions. As mentioned above, the expected precision for ionization measurement is significantly better than that required here. Simulation studies [384, 395, 419] show that the electron identification capability is within the requirements of the present ALICE dielectron physics programme and will be reached with the present design.
- *Rate capability:* To inspect and track electron candidates identified in the TRD, the TPC should be operated at central collision rates of up to 200 Hz. While there is not much operational experience with large TPCs at these rates, the current load on the readout chambers is not excessive. Simulations have shown that, at this rate, the space charge due to the ion feed-back during gate-open time starts to be comparable to the space charge due to the ionization in the TPC drift volume itself. The true space-charge limit, however, might eventually depend on the background conditions due to the beam quality. Offline corrections for the space charge are expected to recover part of the resolution loss and should thus extend this limit.

For pp runs with $5 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity, the space charge due to both ionization in the drift volume and to the ion feed-back during gate-open time, is about one order of magnitude lower than for Pb–Pb, and thus not seen as a problem. Trigger rates of up to 1 kHz seem to be realistic.

- *High-Level Trigger (HLT):* Given the discussed rate limitations, rare processes like J/ψ and Υ production need a trigger to achieve useful statistics. The HLT (see section 3.19 and remarks below) is intended for this task. It is designed to find high-momentum electron tracks in the TRD and match them to the TPC tracks. The recorded data volume can be reduced using the ‘region-of-interest’ option of the trigger, reading out only the sectors of the TPC containing the data about the ‘interesting’, high- p_t , tracks (see section 3.17).

3.5.2. Detector layout. The TPC design is ‘conventional’ in overall structure but innovative in many details. The TPC layout is shown in figure 3.15 and a synopsis of its main parameters is presented in table 3.13. The TPC is cylindrical in shape and has an inner radius of about 85 cm, an outer radius of about 250 cm, and an overall length along the beam direction of 500 cm.

The detector is made of a large cylindrical field cage, filled with 88 m^3 of Ne/CO₂ (90%/10%), which is needed to transport the primary electrons over a distance of up to 2.5 m on either side of the central electrode to the end-plates. Multi-wire proportional chambers with cathode pad readout are mounted into 18 trapezoidal sectors of each end-plate.




Figure 3.15. TPC schematic layout.

The main aspects of the design are listed below.

- *Acceptance:* The overall acceptance covered by the TPC is $|\eta| < 0.9$ for full radial track length and matches that of the ITS, TRD, and TOF detectors; for reduced track length (and poorer momentum resolution), an acceptance up to about $|\eta| = 1.5$ is accessible.
- *Material budget:* The material budget of the TPC is kept as low as possible to ensure minimal multiple scattering and secondary particle production. Thus both the field cage and drift gas are made of materials with small radiation length. The TPC material is about 3.5% of a radiation length for tracks with normal incidence.
- *Field cage:* The field cage is based on a design with a central high-voltage electrode and two opposite axial potential dividers which create a highly uniform electrostatic field in the common gas volume, see figure 3.15. The central electrode is a single stretched Mylar foil to satisfy the requirement of minimal material near 90° relative to the beam direction. The potential of the drift region is defined by Mylar strips wound around 18 inner and outer support rods which also contain the resistive potential dividers. The rods are aligned with the dead zones in between the readout chambers. Because of the Ne/CO₂ (90%/10%) gas mixture used in the TPC, the field cage will have to be operated at very high-voltage gradients, of about 400 V cm^{-1} , with a high voltage of 100 kV at the central electrode which results in a maximum drift time of about $90 \mu\text{s}$.

An insulating gas envelope of CO₂ in containment vessels surrounds the field cage. The field cage and containment volume are each constructed from two concentric cylinders, sealed by the end-plate on either side. To provide high structural integrity against gravitational and thermal loads while keeping the material budget low, composite materials were chosen. Hence the mechanical stability and precision is guaranteed to be about $250 \mu\text{m}$.

More details on the design and extensive prototyping are discussed in section 3.1 of the TPC TDR [393]; details on test beam results can be found in [420, 421].

- *Drift gas:* The drift gas Ne/CO₂ (90%/10%) is optimized for drift speed, low diffusion, low radiation length and hence low multiple scattering, small space-charge effect, and ageing properties. Mixtures containing CH₄ and CF₄ were rejected due to their ageing properties.

Table 3.13. Synopsis of TPC parameters.

Pseudo-rapidity coverage	$-0.9 < \eta < 0.9$ for full radial track length $-1.5 < \eta < 1.5$ for 1/3 radial track length
Azimuthal coverage	2π
Radial position (active volume)	$845 < r < 2466$ mm
Radial size of vessel	$780 < r < 2780$ mm
Length (active volume)	5000 mm
Segmentation in φ	18 sectors
Segmentation in r	Two chambers per sector
Segmentation in z	Central membrane, readout on two end-plates
Total number of readout chambers	$2 \times 2 \times 18 = 72$
Inner readout chamber geometry	Trapezoidal, $848 < r < 1320$ mm active area
Pad size	4×7.5 mm ($\varphi \times r$)
Pad rows	63
Total pads	5504
Outer readout chamber geometry	Trapezoidal, $1346 < r < 2466$ mm active area
Pad size	6×10 and 6×15 mm ($\varphi \times r$)
Pad rows	$64+32 = 96$ (small and large pads)
Total pads	$4864+5120 = 9984$ (small and large pads)
Detector gas	Ne/CO ₂ 90/10
Gas volume	88 m ³
Drift length	2×2500 mm
Drift field	400 V cm ⁻¹
Drift velocity	2.84 cm μ s ⁻¹
Maximum drift time	88 μ s
Total HV	100 kV
Diffusion	$D_L = D_T = 220$ μ m cm ^{-1/2}
Material budget	$X/X_0 = 3.5$ to 5% for $0 < \eta < 0.9$
Front-End Cards (FEC)	121 per sector $\times 36 = 4356$
Readout Control Unit (RCU) scheme	6 per sector, 18 to 25 FEC per RCU
Total RCUs	216
Total pads — readout channels	557 568
Pad occupancy (for $dN/dy = 8\,000$)	40–15% (inner/outer radius)
Pad occupancy (for pp)	$5\text{--}2 \times 10^{-4}$ (inner/outer radius)
Event size (for $dN/dy = 8000$)	~ 60 MB
Event size (for pp)	$\sim 1\text{--}2$ MB depending on pile-up
Data rate limit	400 Hz Pb–Pb minimum bias events
Trigger rate limits	200 Hz Pb–Pb central events 1000 Hz proton–proton events
ADC	10 bit
Sampling frequency	5.7–11.4 MHz
Time samples	500–1000
Conversion gain	6 ADC counts fC ⁻¹
Position resolution (σ)	
In $r\varphi$	1100–800 μ m (inner/outer radii)
In z	1250–1100 μ m
dE/dx resolution	
Isolated tracks	5.5%
$dN/dy = 8000$	6.9%

The drawback of Ne/CO₂ is that this mixture is a ‘cold’ gas, with a steep dependence of drift velocity on temperature. For this reason, the TPC is aiming for a thermal stability with $\Delta T \leq 0.1$ K in the drift volume over the running period.

- *Readout chambers:* The readout chambers instrument the two end-caps of the TPC cylinder with an overall active area of 32.5 m². The chambers are multi-wire proportional chambers with cathode pad readout. Because of the radial dependence of the track density, the readout is segmented radially into two readout chambers with slightly different wire geometry adapted to the varying pad sizes mentioned below. The radial distance of the active area is from 84.1 to 132.1 cm (and from 134.6 to 246.6 cm) for the inner (and outer) chamber, respectively. The inactive areas between neighbouring inner chambers are aligned with those between neighbouring outer chambers. Such an arrangement optimizes the momentum precision for detected high-momentum tracks but has the drawback of creating cracks in the acceptance—in about 10% of the azimuthal angle the detector is non-sensitive. The readout chambers are made of standard wire planes, i.e. they consist of a grid of anode wires above the pad plane, a cathode wire plane, and a gating grid.

To keep the occupancy as low as possible and to ensure the necessary dE/dx and position resolution, there are about 560 000 readout pads of three different sizes: 4×7.5 mm² in the inner chambers, 6×10 and 6×15 mm² in the outer chambers.

Further details are discussed in section 4.1.2 of the TPC TDR [393] and in [422, 423].

- *Gating:* The readout chambers are normally closed by a gating grid for electrons coming from the drift volume and they are opened only by the L1 trigger ($6.5 \mu\text{s}$ after the collision) for the duration of one drift-time interval, i.e. of about $90 \mu\text{s}$ (see section 3.17). This helps to prevent space charge due to positive ions from drifting back from the multiplication region for non-triggered interactions and background.
- *Laser calibration:* A laser system with some hundred straight tracks in all regions of the drift space will allow precise position inter-calibration for the readout chambers and monitoring of temperature and space-charge distortions.

3.5.3. Front-end electronics and readout. The front-end electronics has to read out the charge detected by about 560 000 pads located on the readout chambers at the TPC end-caps. These chambers deliver on their pads a current signal with a fast rise time (less than 1 ns), and a long tail due to the motion of the positive ions. The amplitude, which is different for the different pad sizes, has a typical value of $7 \mu\text{A}$. The signal is delivered on the detector impedance, to a very good approximation a pure capacitance of the order of a few pF. The requirements and design specification for the front-end and readout electronics are discussed in section 5 of the TPC TDR [393].

A single readout channel is comprised of three basic functional units: (i) a charge sensitive amplifier/shaper (PASA); (ii) a 10-bit 10-MSPS low-power ADC; and (iii) a digital circuit that contains a shortening filter for the tail cancellation, baseline subtraction and zero-suppression circuits, and a multiple-event buffer.

The charge collected on the TPC pads is amplified and integrated by a low-input impedance amplifier. It is based on a Charge Sensitive Amplifier (CSA) followed by a semi-Gaussian pulse shaper of the fourth order. These analogue functions are realized by a custom integrated circuit (PASA), implemented in a CMOS technology $0.35 \mu\text{m}$, which will contain 16 channels with a power consumption/channel of 12 mW. The circuit has a conversion gain of 12 mV fC^{-1} and an output dynamic range of 2 V with a linearity of 1%. It produces a pulse with a rise time of 120 ns and a shaping time (FWHM) of 190 ns. The single channel has a noise value (r.m.s.) below $1000e$ and a channel-to-channel cross-talk below -60 dB. Immediately after the PASA, a 10-bit pipelined ADC (one per channel) samples the signal at a rate of 5–6 MHz.

The digitized signal is then processed by a set of circuits that perform the baseline subtraction, tail cancellation, zero-suppression, formatting and buffering. The ADC and the digital circuits are contained in a single chip named ALTRO (ALice Tpc ReadOut) [425, 426]. The ALTRO chip integrates 16 channels, each of them consisting of a 10-bit, 30-MSPS ADC, a pipelined Digital Processor and a multi-acquisition Data Memory. When a L1 trigger is received, a predefined number of samples (acquisition) are temporarily stored in a data memory. Upon L2 trigger arrival the latest acquisition is frozen, otherwise it will be overwritten by the next acquisition. The Digital Processor, running at the sampling frequency, implements several algorithms that are used to condition and shape the signal. After digitization, the Baseline Correction Unit I is able to perform channel-to-channel gain equalization and to correct for possible non-linearity and baseline drift of the input signal. It is also able to adjust dc levels and to remove systematic spurious signals by subtracting a pattern stored in a dedicated memory. The next processing block is an 18-bit, fixed-point arithmetic, third order Tail Cancellation Filter [427]. The latter is able to suppress the signal tail, within $1 \mu\text{s}$ after the pulse peak, with an accuracy of 1 LSB. Since the coefficients of this filter are fully programmable, the circuit is able to cancel a wide range of signal-tail shapes. Moreover, these coefficients can be set independently for each channel and are re-configurable. This feature allows a constant quality of the output signal regardless of ageing effects on the detector and channel-to-channel fluctuations. The subsequent processing block, Baseline Correction Unit II, applies a baseline correction scheme based on a moving average filter. This scheme removes non-systematic perturbations of the baseline that are superimposed on the signal. At the output of this block, the signal baseline is constant with an accuracy of 1 LSB. Such accuracy allows an efficient zero-suppression procedure, which discards all data below a programmable threshold, except for a specified number of pre- and post-samples around each pulse. This produces a limited number of non-zero data packets, thus reducing the overall data volume. Each data packet is formatted with its time stamp and size information in such a way that reconstruction is possible afterwards. The output of the Data Processor is sent to a Data Memory of 5 kB, able to store up to 8 full acquisitions. The data can be read out from the chip at a maximum speed of 60 MHz through a 40-bit wide bus, yielding a total bandwidth of 300 MB s^{-1} . Moreover, the readout speed and the ADC sampling frequency are independent. Therefore, the readout frequency does not depend on the bandwidth of the input signal being acquired. The ALTRO chip is implemented in the ST 0.25 μm HCMOS-7 process.

The complete readout chain is contained in the Front-End Cards (FEC) [428], which are plugged in crates integrated in the service support wheel. Each FEC contains 128 channels and is connected to the cathode plane by means of six flexible cables. A number of FECs (up to 25) are controlled by a Readout Control Unit (RCU) [429], which interfaces the FECs to the DAQ, the trigger, and the Detector Control System (DCS). The RCU broadcasts the trigger information to the individual FEC modules and controls the readout procedure. Both functions are implemented via a custom bus, based on low-voltage signalling technology (GTL), which provides a data bandwidth of 200 MB s^{-1} . The interfacing of the RCU modules to the Trigger and to the DAQ follows the standard data acquisition architecture of the experiment. In summary, for each of the 36 TPC sectors, the front-end electronics consists of 121 FECs, six RCUs, and six Detector Data Links (DDL).

The data rate capabilities of the TPC readout are designed to allow transfer of 200 central or 400 minimum-bias Pb–Pb events s^{-1} for the extreme multiplicity case. The recorded data volume can be reduced using the ‘region-of-interest’ option of the trigger, reading out only a few sectors of the TPC.

After zero-suppression and data encoding, the event size from the TPC for a central Pb–Pb collision of extreme multiplicity will be about 60 MB. In order to increase the physics potential

of ALICE especially on jet and electron physics, rare signals like dielectron pair candidates have to be enriched to readout rates of 100–200 Hz. Therefore an ‘intelligent’ readout is under development via an HLT processor farm, which will operate on the raw data shipped via optical links to the ALICE counting house. The HLT will allow near lossless data compression, selective readout of electron candidates identified by the TRD (the ‘region-of-interest’ option of the trigger), as well as online track finding and possibly even tracking of the whole TPC. Details of the intelligent readout including strategies for its implementation are discussed in section 3.19.

3.6. Transition-Radiation Detector (TRD)

3.6.1. Design considerations. The main goal of the ALICE Transition-Radiation Detector (TRD) [395] is to provide electron identification in the central barrel for momenta greater than $1 \text{ GeV } c^{-1}$, where the pion rejection capability through energy loss measurement in the TPC is no longer sufficient. As a consequence, the addition of the TRD [384] significantly expands the ALICE physics objectives [382, 383]. The TRD will provide, along with data from the TPC and ITS, sufficient electron identification to measure the production of light and heavy vector-meson resonances and the dilepton continuum in Pb–Pb and pp collisions.

In addition, the electron identification provided by the TPC and TRD for $p_t > 1 \text{ GeV } c^{-1}$ can be used, in conjunction with the impact-parameter determination of electron tracks in the ITS, to measure open charm and open beauty produced in the collisions. A similar technique can be used to separate directly produced J/ψ mesons from those produced in B-decays. These secondary J/ψ 's could potentially mask the expected J/ψ yield modification due to quark–gluon plasma formation; their isolation is, therefore, of crucial importance for such measurements. Furthermore, since the TRD is a fast tracker, it can be used as an efficient trigger for high transverse momentum electrons. Such a trigger would considerably enhance the recorded Υ yields in the high-mass part of the dilepton continuum as well as high- p_t J/ψ .

These physics requirements have driven the following design considerations [395]:

- *The required pion rejection capability* is driven mostly by the J/ψ measurement and its p_t dependence. As outlined in the Addendum 2 to the ALICE proposal [384], the goal is an increase in pion rejection by a factor of 100 for electron momenta above $3 \text{ GeV } c^{-1}$. While the requirement for the Υ is less stringent, the light vector mesons ρ , ω , and ϕ as well as the dielectron continuum between the J/ψ and the Υ are only accessible with this level of rejection.
- *The required momentum resolution* is determined by the matching to the TPC. The momentum resolution requirements for the central barrel are fulfilled by combining TPC and ITS reaching e.g. a mass resolution of $100 \text{ MeV } c^{-2}$ at the Υ for $B = 0.4 \text{ T}$. The function that the TRD needs to fulfil is to add electron identification. This goal can be reached by having a pointing capability from the TRD to the TPC with an accuracy of a fraction of a TPC pad. The TRD will provide a momentum resolution of 5% at $5 \text{ GeV } c^{-1}$ leading to a pointing accuracy of 30% of the pad width. Such accuracy enables unambiguous matching with exception of very close hits. At the trigger level, good momentum resolution leads to a sharper threshold and a smaller probability of fake tracks but no strict requirement can be derived from this.
- *The thickness of the TRD* in radiation lengths must be minimized since material generates additional background, mainly from photon conversion, and increases the pixel occupancy. Also, electron energy loss due to Bremsstrahlung removes electrons from the resonance reconstruction sample.




Figure 3.16. Cut through the ALICE TRD with the TPC inside.

- *The granularity* of the TRD is driven in the bend direction by the required momentum resolution and along the beam direction by the required capability to identify and track electrons efficiently at the highest possible multiplicity. To not drastically affect the reconstructed pair signal, the detector has been designed for 80% single-track efficiency. The pads have an area of about 6 cm^2 to attain the desired efficiency.
- *Occupancy:* The assumption of $dN_{\text{ch}}/d\eta = 8000$ leads to a readout-pixel occupancy in central collisions of about 34% in the TRD (including secondary particles) for the selected pad size. The detector is designed to function at this occupancy.

3.6.2. Detector layout. The physics requirements and design considerations listed above led to the present design of the TRD which is depicted in figure 3.16. The main aspects of the design are summarized below. For a quick overview, see the synopsis in table 3.14. Cylindrical coordinates with the origin at the beam intersection point and with the positive z -axis pointing away from the muon arm are used. The angle φ is then also the deflection angle in the magnetic field. Since the TRD chambers are flat and not on a cylindrical surface it is often more convenient, when discussing processes and resolutions inside a given chamber, to use Cartesian coordinates. In this case we keep the same z -axis, y is the direction of the wires and of the deflection in the magnetic field and x is the electron drift direction.

The coverage in pseudo-rapidity matches the coverage of the other central barrel detectors, $|\eta| \leq 0.9$. The TRD fills the radial space between the TPC and the TOF detectors. For quality of electron identification, the TRD consists of six individual layers. To match the azimuthal segmentation of the TPC, there are 18 sectors. There is a 5-fold segmentation along beam direction, z . In total there are $18 \times 5 \times 6 = 540$ detector modules.

Each module consists of a radiator of 4.8 cm thickness, a multi-wire proportional readout chamber, and the front-end electronics for this chamber. The signal induced on the cathode

Table 3.14. Synopsis of TRD parameters.

Pseudo-rapidity coverage	$-0.9 < \eta < 0.9$
Azimuthal coverage	2π
Radial position	$2.9 < r < 3.7$ m
Length	Up to 7.0 m
Azimuthal segmentation	18-fold
Radial segmentation	Six layers
Longitudinal segmentation	5-fold
Total number of modules	540
Largest module	117×147 cm ²
Active detector area	736 m ²
Radiator	Fibres/foam sandwich, 4.8 cm per layer
Radial detector thickness	$X/X_0 = 15\%$
Module segmentation in φ	144
Module segmentation in z	12–16
Typical pad size	0.7×8.8 cm ⁻² = 6.2 cm ²
Number of pads	1.16×10^6
Detector gas	Xe/CO ₂ (85%/15%)
Gas volume	27.2 m ³
Depth of drift region	3 cm
Depth of amplification region	0.7 cm
Nominal magnetic field	0.4 T
Drift field	0.7 kV cm ⁻¹
Drift velocity	1.5 cm μ s ⁻¹
Longitudinal diffusion	$D_L = 250$ μ m cm ^{-1/2}
Transverse diffusion	$D_T = 180$ μ m cm ^{-1/2}
Lorentz angle	8°
Number of readout channels	1.16×10^6
Time samples in r (drift)	20
Number of readout pixels	2.32×10^7
ADC	10 bit, 10 MHz
Number of multi-chip modules	71 928
Number of readout boards	4108
Pad occupancy for $dN_{ch}/d\eta = 8000$	34%
Pad occupancy in pp	2×10^{-4}
Space-point resolution at 1 GeV c^{-1}	
in $r\varphi$	400 (600) μ m for $dN_{ch}/d\eta = 2000$ ($dN_{ch}/d\eta = 8000$)
in z	2 mm (offline)
Momentum resolution	$\delta p/p = 2.5\% \oplus 0.5\%$ (0.8%) p for $dN_{ch}/d\eta = 2000$ ($dN_{ch}/d\eta = 8000$)
Pion suppression at 90% electron efficiency and $p \geq 3$ GeV c^{-1}	Better than 100
Event size for $dN_{ch}/d\eta = 8000$	11 MB (cf text)
Event size for pp	6 kB
Trigger rate limits for minimum-bias events	100 kHz
Trigger rate limits for pp	100 kHz (cf text)

pads is read out. Each chamber has 144 pads in the direction of the amplification wires, $r\varphi$, and between 12 and 16 pad rows in the z direction. The pads have a typical area of 6–7 cm² and cover a total active area of about 736 m² with 1.16×10^6 readout channels. Prototypes have been shown to perform according to the specifications given in table 3.14 [430].

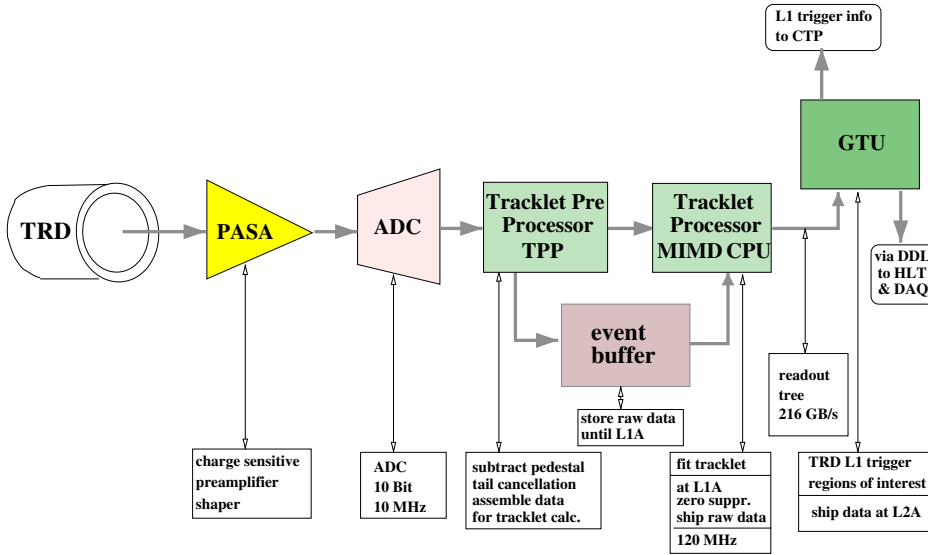


Figure 3.17. Block diagram of the TRD front-end electronics.

The gas mixture in the readout chambers is Xe/CO₂ (85%/15%). Each readout chamber consists of a drift region of 3.0 cm separated by cathode wires from an amplification region of 0.7 cm. The drift time is 2.0 μs , requiring a drift velocity of 1.5 cm μs^{-1} . The nominal drift velocity will be reached with an electric field of 0.7 kV cm⁻¹. In this gas mixture a minimum ionizing particle liberates 275 electrons cm⁻¹. The gas gain will be of the order of 5000. The properties of the proposed gas mixture have been extensively tested in measurements [431]. The induced signal at the cathode pad plane will be sampled in 20 time intervals spaced 1.5 mm or 100 ns apart. Diffusion is negligible, see table 3.14. At a magnetic field of 0.4 T, the Lorentz angle is 8°.

At extreme multiplicity ($dN_{ch}/d\eta = 8000$) the pixel occupancy will be 34%. A space point resolution in bend direction of 400 μm can be achieved for low multiplicity ($dN_{ch}/d\eta = 2000$) at $p_t = 1 \text{ GeV } c^{-1}$. For full multiplicity, this resolution is degraded to 600 μm after unfolding. The momentum resolution of the TRD in stand-alone mode is determined by a constant term of 2.5% and a linear term of 0.5% per GeV c^{-1} . The linear term is degraded to 0.8% for full multiplicity. The event sizes shown in table 3.14 refer to the data size shipped to the Global Tracking Unit. Depending on its operation the event size to DAQ may be significantly less. The quoted trigger rate limitations are determined by the requirement not to have overlapping events and the currently foreseen low-voltage power budget.

3.6.3. Front-end electronics and readout. A block diagram of the main elements of the readout electronics of the TRD is shown in figure 3.17. All components apart from the Global Tracking Unit (GTU) are implemented in two custom ASICs and are directly mounted on the detector in order to optimize latency for the trigger decision. The analogue ASIC contains the pre-amplifier, shaper, and output driver. The other ASIC is a mixed analogue/digital design containing the ADC, tracklet pre-processor, event buffer, and tracklet processor. Both chips are mounted on a so-called Multi-Chip Module (MCM) serving 18 pads. The pre-amplifier and shaper circuit has a conversion gain of 12 mV fC⁻¹, a shaping time of 120 ns, an input dynamic range of 164 fC, and provides a maximum differential output of $\pm 1 \text{ V}$. A single channel has a noise value (r.m.s.) below 1000e.

The ADC is a full custom $0.18\ \mu\text{m}$ design of a 10-bit 10 MHz differential ADC with a power consumption of about 5 mW. It is followed by a digital filter in order to compensate tails of the pulses due to slow ion drift and tails from the electronics, i.e. the time response function. Without such a filter significant distortions in the position measurement would be incurred depending on the history of pulse heights. These position measurements are used to reconstruct the short track pieces (tracklets) inside the drift region.

The output of the digital filters is directly fed into the tracklet pre-processor. Here, all relevant sums are calculated which are subsequently used by the tracklet processor in order to calculate the position and inclination of tracklets, i.e. the track pieces contained in the drift region. Up to four candidate tracklets per MCM are shipped to the GTU, where tracklets from the individual layers are combined and matching cuts as well as momentum cuts and more involved cuts (such as cuts on the invariant mass of pairs) can be applied. The GTU will be implemented in FPGAs close to the Central Trigger Processor (CTP). All computations are finished and the trigger decision of the TRD is sent to the CTP $6\ \mu\text{s}$ after the interaction.

3.7. Time-Of-Flight (TOF) detector

3.7.1. Design considerations. The Time-Of-Flight (TOF) detector of ALICE [394] is a large area array that covers the central pseudo-rapidity region ($|\eta| \lesssim 0.9$) for Particle IDentification (PID) in the intermediate momentum range (from 0.2 to $2.5\ \text{GeV}\ c^{-1}$). Since the majority of the produced charged particles is emitted in this range, the performance of such a detector is of crucial importance for the experiment [382]. The measurement and identification of charged particles in the intermediate momentum range will provide observables which can be used to probe the nature and dynamical evolution of the system produced in ultra-relativistic heavy-ion collisions at LHC energies.

The TOF, coupled with the ITS and TPC for track and vertex reconstruction and for dE/dx measurements in the low-momentum range (up to about $0.5\ \text{GeV}\ c^{-1}$), will provide event-by-event identification of large samples of pions, kaons, and protons. The TOF-identified particles will be used to study relevant hadronic observables on a single-event basis. In addition, at the inclusive level, identified kaons will allow invariant mass studies, in particular the detection of open charm states and the ϕ meson.

A large-coverage, powerful TOF detector, operating efficiently in extreme multiplicity conditions, should have an excellent intrinsic response and an overall occupancy not exceeding the 10–15% level at the highest expected charged-particle density ($dN_{\text{ch}}/d\eta = 8000$). This implies a design with more than 10^5 independent TOF channels.

Since a large area has to be covered, a gaseous detector is the only choice. In the framework of the LAA project at CERN an intensive R&D programme has shown that the best solution for the TOF detector is the Multi-gap Resistive-Plate Chamber (MRPC) [432, 433].

The key aspect of these chambers is that the electric field is high and uniform over the whole sensitive gaseous volume of the detector. Any ionization produced by a traversing charged particle will immediately start a gas avalanche process which will eventually generate the observed signals on the pick-up electrodes. There is no drift time associated with the movement of the electrons to a region of high electric field. Thus the time jitter of these devices is caused by the fluctuations in the growth of the avalanche.

The main advantages of the MRPC technology with respect to other parallel-plate chamber designs (Pestov counter, PPC) are that:

- it operates at atmospheric pressure;
- the signal is the analogue sum of signals from many gaps, so there is no late tail and the charge spectrum is not of an exponential shape—it has a peak well separated from zero;




Figure 3.18. Cross section of a 10-gap double-stack MRPC strip.

- the resistive plates quench the streamers so there are no sparks, thus high-gain operation becomes possible;
- the construction technique is in general rather simple and makes use of commercially available materials.

The latest tests of several MRPC multicell strip prototypes built with a double-stack structure (figure 3.18) show that these devices can reach an intrinsic time resolution better than about 40 ps and an efficiency close to 100%.

3.7.2. Detector layout. The detector covers a cylindrical surface of polar acceptance $|\theta - 90^\circ| < 45^\circ$. It has a modular structure corresponding to 18 sectors in φ (the azimuthal angle) and to five segments in z (the longitudinal coordinate along the beam axis).

The whole device is inscribed in a cylindrical shell with an internal radius of 370 cm and an external one of 399 cm. In terms of material, the whole device thickness corresponds to 20% of a radiation length.

The basic unit of the TOF system is an MRPC strip 1220 mm long and 130 mm wide, with an active area of $1200 \times 74 \text{ mm}^2$ subdivided into pads of size $35 \times 25 \text{ mm}^2$. An overall view of a strip is shown in figure 3.19. The strips are placed inside gas-tight modules (which also act as Faraday cages) and are positioned transversely to the beam direction. Five modules of three different types are needed to cover the full cylinder along the z direction. They all have the same structure and width ($\sim 128 \text{ cm}$) but differ in length. The actual dimensions are defined in such a way that the joining areas of the modules are aligned with the dead areas of the other detectors projected from the interaction point, thus creating a configuration of minimal



Figure 3.19. Photograph of an MRPC strip during construction.

disturbance for the external detectors. The central module is 1.17 m long, the intermediate ones are 1.57 m long, and the external ones are 1.77 m long. The overall TOF barrel length is ~ 745 cm.

The general design has been conceived taking into account the results of the simulations, the feasibility of the proposed solution, the performances of the detector, and the need to keep the dead area inside the module to a minimum.

The granularity of the TOF detector is dictated by the requirement of identifying, on an event-by-event basis, as many charged particles as possible reaching the detector, even at the highest expected charged-particle multiplicity density ($dN_{\text{ch}}/d\eta = 8000$). Detailed simulation studies have shown that, with the chosen pad size of $3.5 \times 2.5 \text{ cm}^2$ and the tilted-strip geometry, the occupancy of the detector is $\sim 16\%$, at the highest charged-particle density (including secondary particles) with a magnetic field of 0.2 T. Lower values of the particle density or higher values of the magnetic field improve, as expected, the occupancy ($\sim 7\%$ for $dN_{\text{ch}}/d\eta = 4000$ and $B = 0.2 \text{ T}$; $\sim 13\%$ for $dN_{\text{ch}}/d\eta = 8000$ and $B = 0.4 \text{ T}$).

In particular, it is important to minimize the transversal path of the incident particles through the strip chambers. This will reduce the number of very oblique transversal paths that can create a sharing effect of the signal among adjacent pads, thereby increasing the occupancy and the time jitter of the detected signals. To overcome this effect, a special positioning of the strips has been envisaged. Their angle with respect to the axis of the cylinder is progressively increasing from 0° in the central part ($\theta = 90^\circ$) of the detector to 45° in the extreme part of

the external module ($\theta = 45^\circ$ and 135°). This arrangement makes the median zone of a strip perpendicular to a radius coming from the interaction point, thus minimizing the impact angle of the incoming particles with respect to the normal direction. To avoid dead areas, adjacent strips have been overlapped inside the modules so that the edge of the active area of one pad is aligned with the edge of the next one. This gives the possibility of creating a full active area with no geometrical dead zones. The modules have been designed in such a way as to avoid any loss of the sensitive area, along the z -axis. The only dead area is due to the unavoidable presence of the supporting space-frame structure.

Every module of the TOF detector consists of a group of MRPC strips (15 in the central, 19 in the intermediate, 19 in the external modules) closed inside an aluminium box that defines and seals the gas volume and supports the external front-end electronics and services. A honeycomb plate is the backbone of the module and gives the necessary mechanical stiffness to the system. The rolling system that allows the insertion of the module into the space-frame rails is connected to this supporting element. The honeycomb plate is 15 mm thick including two aluminium skins, each 1 mm thick. A cover connected to the backbone by means of screws and standard sealing O-Ring closes the total gas volume; it has to withstand the pressure required for gas circulation inside the chamber (up to 3 mbar). The cover is made of fibreglass 3 mm thick, reinforced with protruded ribs, molded according to the final design. An aluminium foil of 0.3 mm is glued to the inside surface of the cover in order to create an electromagnetic shielding. Inside the gas volume, fixed perpendicularly to the honeycomb plate, there are two 5-mm-thick aluminium plates to which the MRPC strips are attached. The system turns out to be very simple and allows for a fast insertion of the strips at whatever angle is needed. Holes that accommodate feed-through for the signal cables coming from the readout pads, the HV connectors and the gas inlet and outlet, are machined into the honeycomb plate. The signal feed-through consists of a PCB having on one side connectors receiving the cables coming from the strips inside the gas volume and, on the other side, connectors that accommodate the front-end electronics cards, see figure 3.20. The PCB is glued to the honeycomb plate via special fibreglass tubes which are inserted into the holes. The volume containing the electronics cards, input and output cables, water cooling pipes and radiators is closed by a cover. This allows access to the electronics with no disturbance to the active part of a module.

The complete TOF system consists of 90 modules. The five modules in a row, located inside a supermodule for each of the 18 sectors, are kept in position by two rails fixed to the space frame, see figure 3.21. Three sliding bushes are fixed to the module body permitting the insertion of the modules into the supporting structure from either side.

An overview of the TOF detector parameters is shown in table 3.15.

Typical results obtained at the CERN/PS T10 test-beam with a 10-gap double-stack MRPC are shown in figure 3.22 for both the efficiency and the time resolution. Similar results have been obtained over a sample of about forty MRPC strips. Therefore it is confirmed that the detector has almost full efficiency and an intrinsic time resolution better than 40 ps.

Extensive tests [434] performed at the CERN Gamma Irradiation Facility (GIF) have shown that the MRPC strips have a rate capability far in excess of the 50 Hz cm^{-2} maximum expected rate at the ALICE experiment.

A recent review of the TOF detector is available in [435].

3.7.3. Front-end electronics and readout. The front-end electronics for the TOF must comply with the basic characteristics of the MRPC detector, i.e. very fast and differential signals from the anode and cathode readout pads and intrinsic time resolution better than 40 ps. During the R&D phase of the detector a solution based on commercial components (a very fast amplifier but not differential at input and a comparator) was adopted. In 2002 an ASIC chip ('NINO')




Figure 3.20. Schematic of a TOF module; the strips are installed inside the gas volume and the FEA cards are plugged onto the interface card.




Figure 3.21. TOF sector (supermodule), consisting of 5 modules inside the space frame.

in CMOS $0.25\ \mu\text{m}$ technology has been developed and tested at PS in 2003 with excellent results, thus allowing a substantial reduction of the power dissipation (now at about 30 mW per channel). All the stages (input amplifier and comparator stages) of the eight-channel ASIC are fully differential, its input impedance and capacitance are matched to the transmission line and the peaking time is smaller than 1 ns.

Table 3.15. Overview of TOF parameters.

Pseudo-rapidity coverage	$-0.9 < \eta < 0.9$
Azimuthal coverage	2π
Radial position	$3.70 < r < 3.99$ m
Length	7.45 m
Segmentation in φ	18-fold
Segmentation in z	5-fold
Total number of modules	90
Central module (A)	117×128 cm ²
Intermediate module (B)	157×128 cm ²
External module (C)	177×128 cm ²
Detector active area	141 m ²
Detector thickness radially	$X/X_0 = 20\%$
Number of MRPC strips per module	15 (A), 19 (B), 19 (C)
Number of readout pads per MRPC strip	96
Module segmentation in φ	48 pads
Module segmentation in z	30 (A), 38 (B), 38 (C) pads
Readout pad geometry	3.5×2.5 cm ²
Total number of MRPC strips	1638
Total number of readout pads	157 248
Detector gas	C ₂ H ₂ F ₄ (90%),i-C ₄ F ₁₀ (5%),SF ₆ (5%)
Gas volume	16 m ³
Total flow rate	2.7 m ³ h ⁻¹
Working pressure	<3 mbar
Fresh gas flow rate	0.027 m ³ h ⁻¹
Number of readout channels	157 248
Number of front-end analogue chips (8-ch)	19 656
Number of front-end boards	6552
Number of HPTDC chips (8-ch, 24.4 ps bin width)	19 656
Number of HPTDC readout boards (TRM)	684
Number of readout boards (DRM) and crates	72
Occupancy for $dN_{ch}/d\eta = 8000$	13% ($B = 0.4$ T), 16% ($B = 0.2$ T)
Occupancy for pp	6×10^{-4}
π , K identification (with contamination <10%)	0.2–2.5 GeV c^{-1}
p identification (with contamination <10%)	0.4–4.5 GeV c^{-1}
e identification in pp (with contamination <10%)	0.1–0.5 GeV c^{-1}
Event size for $dN_{ch}/d\eta = 8000$	100 kB
Event size for pp	<1 kB

The basic Front-End Analogue card (FEA) contains three ASIC chips (24 channels) with a common threshold regulation; it is connected to the PCB interface cards on top of the honeycomb support plate of a module. The LVDS output signals, routed along the two sides of the supermodule (1 per sector), carry the information of the hit time (leading edge) and of the Time-Over-Threshold (TOT), related to the input charge, that is needed for the time-slewing correction.

The readout electronics, located in custom crates at both ends of a supermodule, consists of the TRM (TDC Readout Module) and DRM (Data Readout Module) cards. The TRM card houses the HPTDC (High Performance TDC) [436] eight-channel chips that, for the TOF, are used in the very high-resolution mode (24.4 ps bin width). Each TRM card contains 30 HPTDC chips, i.e. 240 channels, corresponding to the readout pads of 2.5 MRPC strips. The




Figure 3.22. Efficiency and time resolution of a 10-gap double-stack MRPC strip as a function of the HV value. Measurements in ten different positions (pads) along and across the strip are shown.

DRM card is the TOF interface to the ALICE DAQ system, it reads and encodes the data from the TRM cards and sends them to the DAQ via the DDL optical link. The DRM card receives the trigger information (L1, L2a, L2r) from the CTP via the TTCrx chip and performs a slow-control function with a dedicated CPU.

3.8. High-Momentum Particle Identification Detector (HMPID)

3.8.1. Design considerations. The High-Momentum Particle Identification Detector (HMPID) [385], is dedicated to inclusive measurements of identified hadrons for $p_t > 1 \text{ GeV } c^{-1}$. The HMPID was designed as a single-arm array with an acceptance of 5% of the central barrel phase space. The geometry of the detector was optimized with respect to particle yields at high- p_t in both pp and heavy-ion collisions at LHC energies, and with respect to the large opening angle (corresponding to small effective size particle emitting sources) required for two-particle correlation measurements. HMPID will enhance the PID capability of the ALICE experiment by enabling identification of particles beyond the momentum interval attainable through energy loss (in ITS and TPC) and time-of-flight measurements (in TOF). The detector was optimized to extend the useful range for π/K and K/p discrimination, on a track-by-track basis, up to 3 and 5 $\text{GeV } c^{-1}$ respectively.

3.8.2. Detector layout. The HMPID is based on proximity-focusing Ring Imaging Cherenkov (RICH) counters and consists of seven modules of about $1.5 \times 1.5 \text{ m}^2$ each, mounted in an independent support cradle [385, 437]. The cradle will be fixed to the space frame in the two o'clock position (figure 3.23).

The radiator, which defines the momentum range covered by the HMPID, is a 15 mm thick layer of low chromaticity C_6F_{14} (perfluorohexane) liquid with an index of refraction of $n = 1.2989$ at $\lambda = 175 \text{ nm}$ corresponding to $\beta_{\min} = 0.77$ (i.e. a momentum threshold $p_{\text{th}} = 1.21 m$, where m is the particle mass).

Cherenkov photons, emitted when a fast charged particle traverses the radiator, are detected by a photon counter (figure 3.24) which exploits the novel technology of a thin layer of CsI deposited onto the pad cathode of a Multi-Wire Pad Chamber (MWPC). The




Figure 3.23. Axonometric view of the HMPID with cradle and space frame.




Figure 3.24. Working principle of a RICH detector employing CsI thin films deposited onto the cathode plane of a MWPC. The Cherenkov cone refracts out of the liquid radiator of C_6F_{14} and expands in the proximity volume of CH_4 before reaching the MWPC photon detector. Electrons released by ionizing particles in the proximity gap are prevented to enter the MWPC volume by a positive polarization of the collection electrode close to the radiator.

HMPID, with its surface of about 10 m^2 , represents the largest scale application of this technique [438].

A relevant feature of the adopted design for the HMPID photo-detector is the open geometry, which suppresses the electrode elements that in the ‘old-fashion’ TMAE photo-detectors were specifically implemented to prevent spurious avalanches from feedback photons.

Each HMPID module consists of a stack of five independent frames made out of a stabilized aluminium alloy (peralumin) which minimizes out-gassing into the chamber active volume.

The photo-detector, filled with pure methane and operated at ambient temperature and pressure, is closed on one side by an end-flange which supports six independent CsI photocathode boards, $64 \times 40\text{ cm}^2$ large, segmented into pads of $8 \times 8.4\text{ mm}^2$. On the opposite side, a honeycomb panel supports three C_6F_{14} radiator vessels placed at a distance of 80 mm from the anode wire plane. The gas tightness is ensured by FM soft O-rings placed in grooves in the chamber frames.

The assembly procedure of module elements fixes the wire-chamber gap (2 mm anode-cathode distance), with a precision of $50\text{ }\mu\text{m}$, and the parallelism between radiator trays and photon detector within an accuracy better than $100\text{ }\mu\text{m}$. The MWPC anode plane is made of gold plated tungsten-rhenium anode wires, $20\text{ }\mu\text{m}$ in diameter, spaced 4.2 mm apart. They are soldered on a G-10 printed board with a precision of 0.1 mm and a tension of 50 g. On both edges of the anode plane, thicker guard wires resist the boundary discontinuity of the electrostatic field. A support line structure in Macor is implemented between the pad cathode and the anode plane to ensure stability of the sensitive wires against electrostatic force.

Cathode and collection wire planes are made out of $100\text{ }\mu\text{m}$ diameter gold plated Cu/Be wires stretched at 200 g per wire by using crimping pins.

Stiffness and flatness of pad cathodes is obtained by gluing together two multilayer printed circuit boards. The photo-converter is a 300 nm thick layer of CsI deposited onto printed circuit boards with a copper layer. The boards are accurately pre-polished through mechanical and chemical treatments, and covered by electrolytic deposition first with a thick nickel layer and then by a thinner layer of gold.

The liquid radiator containers consist of trays of $1330 \times 413\text{ mm}^2$ made out of glass-ceramic material (NEOCERAM), thermally compatible (thermal coefficient $0.5 \times 10^{-6}\text{ }^\circ\text{C}^{-1}$) with the fused silica plates used as UV-transparent windows. The thickness and size of the tray’s elements have been carefully optimized by investigating the best compromise between the detector total radiation length and the perfluorohexane hydrostatic pressure: the quartz window is 5 mm thick, while the NEOCERAM base plate is 4 mm thick.

To withstand the hydrostatic pressure, thirty cylindrical spacers are glued to the NEOCERAM bottom plate on one side and the quartz window on the other side. Spacers consist of fused silica rods with a diameter of 10 mm placed in three rows of 10 equidistantly spaced elements.

Radiator trays are supported by a stiff composite panel, consisting of a 50 mm thick layer of Rohacell sandwiched between two thin 0.5 mm layers of aluminium. Connections to the liquid radiator inlet and outlet pipework are obtained by gluing flexible stainless steel bellow to opposite edges of the NEOCERAM tray, the outlet (inlet) always being at the highest (lowest) location.

Experimental results motivated the choice of the radiator thickness. In particular, with the help of Monte Carlo simulation and test beam measurements, it was demonstrated [439] that a radiator thickness of 15 mm will enable operating the detector at a lower gain, keeping unchanged the performance and improving the stability of operation. The higher photon yield

coupled with the improved QE of the recently produced photo-cathodes will be used to tune the detector gain adapting it to the specific running conditions. Lower gain can be used in ion–ion collisions, minimizing the photon feedbacks and MIPs contribution to the occupancy, while in proton–proton runs, where the much lower track density eases the pattern recognition, higher gain can be used to improve the efficiency.

A positive voltage of 2050 V, applied to the anodes, while cathodes are grounded, provides a total gas gain of almost 10^5 .

A liquid circulation system is required to purify C_6F_{14} , fill and empty at a constant flow the twenty-one radiator trays independently, remotely and safely. By considering the inaccessibility of the detector during the run and the fragility of the radiator trays, a system based on a gravity flow principle has been chosen owing to its safe nature. Since C_6F_{14} is not available in a high-purity grade form, filters are implemented in the circulation system in order to remove contaminants (mainly water and oxygen) and achieve the best transparency in the UV region where the RICH detector operates. The total HMPID gas volume of 1.4 m^3 of pure CH_4 is split into the seven independent RICH modules of 0.2 m^3 by supplying each of them individually with a gas flow rate providing five volume changes per day and maintaining an operating pressure of some mbar above ambient pressure.

The parameters of the detector are summarized in table 3.16.

3.8.3. Front-end electronics and readout.

Front-end electronics. The front-end electronics is based on two dedicated ASIC chips, GASSIPLEX [440] and DILOGIC [441], successfully developed in the framework of the HMPID project in the ALCATEL-MIETEC $0.7\text{ }\mu\text{m}$ technology.

The GASSIPLEX-07 chip is a 16-channel analogue multiplexed low-noise signal processor working in TRACK&HOLD mode. It features a dedicated filter to compensate for the long ion drift tail, a semi-Gaussian shaper and internal protection against discharges. The noise on detector is found to be $1000e$ r.m.s. In the HMPID application the GASSIPLEX analogue output will be presented to the input of a commercial 12-bit ADC (AD9220ARS). The multiplexing level will be 3 chips (48 channels) per ADC at a frequency of 10 MHz (tunable) to cope with the maximum expected interaction rate, 20 MHz during proton–proton collisions. The digitization and zero-suppression time does not depend on the occupancy, being determined only by the multiplexing rate. The foreseen highest luminosity ($5 \times 10^{30}\text{ cm}^{-2}\text{ s}^{-1}$, interaction rate of about 200 kHz) in proton–proton interactions will be acceptable for the HMPID since the front-end electronics has a baseline recovery better than 0.5% after $5\text{ }\mu\text{s}$. Moreover, the very low multiplicity per event makes event overlap negligible.

Readout. The DILOGIC chip is a sparse data scan readout processor providing zero suppression and pedestal subtraction with individual threshold and pedestal values for up to 64 channels. Several chips can be daisy-chained on the same 18-bit output bus. Asynchronous read/write operations are allowed. Data are readout via the standard ALICE optical link (DDL).

The readout time after L2 arrival (for 20% occupancy) will be of the order of $200\text{ }\mu\text{s}$.

Since the momentum information is vital to exploit the HMPID detector only events for which the TPC information is available are of interest and the HMPID can perfectly cope with the readout rates foreseen for the TPC. Although the readout of the anode wires is not foreseen, we are investigating the possibility to contribute a signal to the Level 0 trigger, based on a simple wire readout. The motivation for such a trigger is the very low particle multiplicity

Table 3.16. Synopsis of HMPID parameters.

Pseudo-rapidity coverage	$-0.6 < \eta < 0.6$
Azimuthal coverage (rad)	57.61°
Radial position	5 m
Segmentation in φ	3-fold
Segmentation in z	3-fold
Total number of modules	7
Detector active area	10 m^2
Detector thickness radially	$X/X_0 = 18\%$
Radiator thickness	15 mm
Radiator medium	Liquid C_6F_{14}
Refractive index	1.2989 at 175 nm
Independent radiator trays	3 per module (21 total)
β threshold	0.77
Detector gas	CH_4
Gas volume	1.4 m^3
MWPC anode–cathode gap	2 mm
Operating voltage	2050 V
Photon converter	Cesium Iodide (CsI)
CsI thickness	300 nm
Quantum Efficiency (QE)	$> 25\%$ at 175 nm
Number of Photo-Cathodes (PC)	42
Number of pads per PC	3840 (161 280 total)
Pad size	$8.0 \times 8.4 \text{ mm}^2 = 67.2 \text{ mm}^2$
Front-end chip (GASSIPLEX-07-3)	10 080
Peaking time	$1.2 \mu\text{s}$
Readout chip (DILOGIC-3)	3360
Front-end cards	3360
Readout cards	672
ADC (12-bit)	3360
Power consumption per module	500 W (total 3.5 kW)
Number of DDL	14
Multiplexing frequency	10 MHz (maximum)
Multiplexing time	$5 \mu\text{s}$ at 10 MHz
Readout time (for 12% occupancy)	$< 300 \mu\text{s}$
Event size (for 12% occupancy)	$< 0.1 \text{ MB}$
Number of readout channels	161 280
Occupancy ($dN_{\text{ch}}/d\eta = 8000$)	12%

(compared to ion–ion) in proton–proton run that coupled with the small HMPID acceptance (5% of the TPC) greatly reduce the useful particle yield in the HMPID detector. Preliminary investigations have shown that the implementation of such a trigger is possible at a reasonable cost without modification of the present mechanics. The effectiveness of a trigger based on a single wire plane is, nevertheless, still to be worked out.

3.9. PHOTon Spectrometer (PHOS)

3.9.1. Design considerations. The PHOTon Spectrometer (PHOS, [386]) is a high-resolution electromagnetic spectrometer which will detect electromagnetic particles in a limited acceptance domain at central rapidity and provide photon identification as well as neutral

mesons identification through the two-photon decay channel. The main physics objectives are the following:

- Testing thermal and dynamical properties of the initial phase of the collision, in particular the initial temperature and space–time dimensions of the hot zone, through measurement of direct single-photon and diphoton spectra and Bose–Einstein correlations of direct photons.
- Investigating jet quenching as a probe of deconfinement, through measurement of high- p_t π^0 spectrum, and identifying jets through γ -jet and jet–jet correlations measurements.

The principal requirements on PHOS include the ability to identify photons, discriminate direct photons from decay photons and perform momentum measurements over a wide dynamic range with high energy and spatial resolutions.

Photon identification. Photons must be identified and discriminated against charged hadrons with a high efficiency. This is achieved by applying three redundant discrimination criteria (see section 5 of the PPR Volume II). (i) A high granularity segmentation so that showers induced by impinging particles develop over many adjacent cells. Topology analysis of the shower will be used to discriminate electromagnetic and hadronic showers (see section 5 of the PPR Volume II and [442]). (ii) The measurement of the time of flight with resolutions of a few ns will provide means to discriminate photons and baryons (specially useful for neutron and antineutron discrimination). (iii) The addition of a charged-particle detector will enable to veto impacts from charged particles (electrons and charge hadrons). The topology analysis will also help to discriminate high-momentum photons and high-momentum π^0 's which decay into photons emitted in a cone too small to generate two distinct showers in PHOS.

High-energy resolution. The high-energy resolution needed to achieve π^0 identification through invariant mass analysis of the decay photons is achieved by using scintillator material of adequate thickness ($10X_0$) and which provide high photon–electrons yield. The light output must be readout by low-noise photodetectors and processed by low-noise front-end electronics.

High spatial resolution. The spatial resolution required for the invariant-mass analysis and HBT measurements is achieved through a highly granular segmentation of the spectrometer. The particle induced shower spreads over several cells, allowing precise reconstruction of the impact point by calculating the centre of gravity of the shower. The segmentation must be slightly below the Molière radius of the material. The spatial resolution is further determined by the distance of the spectrometer to the interaction point, the larger the distance the better the resolution.

Large dynamical range. Large dynamic range is achieved by selecting appropriate detector thickness to minimize shower leakage for the highest particle-energies without deterioration of energy resolution for lowest particle-energies due to light attenuation along the detector thickness. Using low-noise and high-gain photodetectors based on avalanche photo-diodes, which are insensitive to shower particles leaking out of the material, is required.

3.9.2. Detector layout. PHOS is a single-arm high-resolution high-granularity electromagnetic spectrometer including a highly segmented ElectroMagnetic CAlorimeter (EMCA) and a Charged-particle Veto (CPV) detector. A synopsis of PHOS parameters is given in table 3.17. PHOS is subdivided into five independent EMCA+CPV units, named PHOS modules, positioned on the bottom of the ALICE setup at a distance of 460 cm from the interaction point. It will cover approximately a quarter of a unit in pseudo-rapidity, $-0.12 \leq \eta \leq 0.12$, and 100° in azimuthal angle. Its total area will be $\sim 8 \text{ m}^2$.

Table 3.17. Synopsis of PHOS parameters.

Coverage in pseudo-rapidity	$-0.12 \leq \eta \leq 0.12$
Coverage in azimuthal angle	$\Delta\phi = 100^\circ$
Distance to interaction point	4600 mm
Modularity	Five modules of 3584 crystals
<i>EMCA</i>	
Material	Lead-tungstate crystals (PWO)
Crystal dimensions	$22 \times 22 \times 180 \text{ mm}^3$
Depth in radiation length:	20
Number of crystals	17 920
Segmentation	3584 crystals per module
Total area	8 m^2
Crystal volume	1.5 m^3
Total crystal weight	12.5 t
Operating temperature	-25°C
<i>CPV</i>	
Gas	80% Ar/20% CO ₂
Thickness	$0.5X_0$
Active area	$1.8 \text{ m}^2 \times 14 \text{ mm}$ per module
Wire diameter:	$30 \mu\text{m}$
Number of wires per module	256
Wire pitch	5.65 mm
Pad size	$22 \times 10.5 \text{ mm}^2$
Pad inter-distance	0.6 mm
Number of pads per module	7168

3.9.2.1. ElectroMagnetic Calorimeter (EMCA). Each EMCA module is segmented into 3584 detection channels arranged in 56 rows of 64 channels. The detection channel consists of a $22 \times 22 \times 180 \text{ mm}^3$ lead-tungstate crystal, PbWO₄ (PWO), coupled to a $5 \times 5 \text{ mm}^2$ Avalanche Photo-Diode (APD) which signal is processed by a low-noise preamplifier [443]. The total number of crystals in PHOS is 17 920 representing a total volume of $\sim 1.5 \text{ m}^3$. The main mechanical assembly units in a module is the *crystal strip unit* consisting of eight *crystal detector units* forming 1/8 of a row. The APD and the preamplifier are integrated in a common body glued onto the end face of the crystal with optically transparent glue of a high refractive index.

To significantly (by about a factor of 3) increase the light yield of the PWO crystals (temperature coefficient $\sim -2\%$ per $^\circ\text{C}$), the EMCA modules will be operated at a temperature of -25°C . The temperature will be stabilized with a precision of $\sim 0.3^\circ\text{C}$. To this purpose, the EMCA module is subdivided by thermo-insulation into a ‘cold’ and ‘warm’ volume. The crystal strips will be located in the ‘cold’ volume, whereas the readout electronics will be located outside this volume. All six sides of the ‘cold’ volume will be equipped with cooling panels, and heat is removed by a liquid coolant (hydrofluoroether) pumped through the channels of these panels. Temperature monitoring will be provided by means of a temperature measurement system, based on resistive temperature sensors of thickness $30\text{--}50 \mu\text{m}$, which will be inserted in the gap between crystals.

A monitoring system using Light Emitting Diodes (LED) and stable current generators will monitor every EMCA detection channel. The system consists of Master Modules (MM) and Control Modules (CM). The MM (one per PHOS module) are located in the pit in the same VME crate as the PHOS trigger electronics. For each EMCA module there are 16 CM

Table 3.18. Synopsis of PHOS front-end electronic parameters.

Least count energy single channel	5–10 MeV
Dynamic range	100 GeV
Energy channels	‘High’ and ‘low’ gains
Timing resolution	Around 1 ns at 1–2 GeV
Trigger	L0, L1
Max channel counting rate	
in Pb–Pb	1 kHz
in pp	10 Hz
APD gain control	Individual bias setting

boards, located in the ‘cold’ volume of the EMCA modules, directly on top of the crystals. Each board, placed on a 15 mm thick NOMEX plate, is equipped with a 16×14 LED matrix and the control and decoding circuits.

3.9.2.2. Charged-Particle Veto (CPV) detector. The CPV detector equipping each PHOS module is a Multi-Wire Proportional Chamber (MWPC) with cathode-pad readout [444, 445]. Its charged-particle detection efficiency is better than 99%. The spatial precision of the reconstructed impact point is about 1.6 mm in both directions.

The CPV consists of five separate modules placed on top of the EMCA modules at a distance of about 5 mm. The material budget is less than 5% of X_0 . The active area of 14 mm separating the anode from the cathode is filled with a gas mixture is 80% Ar/20% CO₂ at a pressure slightly (1 mbar) above atmospheric pressure. 256 anode wires, 30 μm in diameter, are stretched 7 mm above the cathode with a pitch of 5.65 mm. They are oriented along the direction of the L3 magnetic field. The cathode plane is segmented into 7168 of $22 \times 10.5 \text{ mm}^2$ pads with an inter-pad distance of 0.6 mm. The largest dimension is aligned along the wires. The total sensitive area of the CPV module is equal to about 1.8 m².

3.9.3. Front-end electronics and readout.

3.9.3.1. EMCA front-end electronics. The EMCA electronic chain includes energy digitization, timing for TOF discrimination, and trigger logic for generating L0 and L1 triggers to ALICE. To cover the required large dynamic range up to 100 GeV, each energy shaper channel supplies two outputs with ‘low’ and ‘high’ amplification, digitized in separate ADCs. The gains of the APDs are equalized by means of a control system where the bias is set individually for each APD. The preamplifier is integrated with the APD and mounted on the crystal in the cold volume. The energy digitization, timing, trigger, readout, and control electronics are mounted on cards placed in the ‘warm’ volume of the PHOS module.

Each of the PHOS modules comprises 3584 crystals. The number of electronic channels per crystal is 4: for dual gain, energy, timing and trigger logic. In addition there is an individual control of the bias for each APD. The total number of electronic channels for the five PHOS modules is 89 600. These electronics, plus the readout controllers for moving the event data to the ALICE DAQ system, is contained in the ‘warm’ volume under the crystal compartment. This volume measures $145 \times 131 \text{ cm}^2$ with a height of 21.5 cm available for the electronics.

The main physics requirements for the front-end electronics are summarized in table 3.18.

The timing resolution is sufficient for TOF discrimination against low-energy antineutrons. The physics trigger can operate in L0 mode with 600 ns decision latency, or in L1 mode. The mode and trigger level are programmable.

The main parameters of the preamplifier are: sensitivity 1 V pC^{-1} , max. input charge 8 pC, ENC around $400e$ for $C_{\text{APD}} = 100 \text{ pF}$, and power dissipation 110 mW. The shaping

amplifier has a time constant of $1.6 \mu\text{s}$. The electronic noise in the output signal corresponds to around 5 MeV. The shaper also supplies a fast energy signal for timing and the trigger logic.

The segmentation of the front-end electronics for a PHOS module is as follows: 112 FEE boards each processing the energy and timing from 32 crystals, eight trigger boards (TRU) each covering a matrix of 28×16 crystals, and two readout controller units (RCUs). The trigger boards also serve as buffer and router for the event data stream from 14 FEE boards to the RCU. The interconnection between TRU and the FEE boards is the GTL cable bus.

The trigger algorithm calculates in parallel all 4×4 partial energy sums of the digitized fast energy signals from the trigger matrix crystals. This algorithm is implemented in a FPGA. The FPGA is programmable via the RCU.

To equalize the gain of the energy channels, the APD bias control system regulates the bias voltage of each of the APDs with an accuracy of 0.5 V.

3.9.3.2. CPV front-end electronics. The CPV pad electronics is the one used for the HMPID. Three kinds of electronics cards are used. The amplifier shapers for 48 pad channels are organized like the three GASSIPLEX cards, placed on top of the pad plane. 32 series of two Multi-Chip Module (MCM) cards and one Column-Memory and Read-Write (CMRW) protocol card are mounted on the Bus card, installed at the periphery of the pad plane, i.e. outside of the sensitive area of the CPV module.

3.9.3.3. Readout. The readout is based on the concept and modules developed for the TPC detector. A readout controller unit (RCU) on the detector transfers formatted event data over the ALICE digital data link (DDL) to the ALICE DAQ system. The RCU operations are executed by means of firmware in the on-board FPGA. The firmware includes a processor core for handling the Ethernet connection to the ALICE detector control system.

The data and control interfaces between the RCU and the TRUs of the front-end electronics are implemented by means of the GTL cable bus.

3.10. Forward muon spectrometer

3.10.1. Design considerations. Hard, penetrating probes, such as heavy-quarkonia states, are an essential tool for probing the early and hot stage of heavy-ion collisions. At LHC energies, energy densities high enough to melt the $\Upsilon(1s)$ will be reached. Moreover, production mechanisms other than hard scattering might play a role [446, 447]. Since these additional mechanisms strongly depend on charm multiplicity, measurements of open charm and open beauty are of crucial importance (the latter also represents a potential normalization for bottomium). The complete spectrum of heavy quark vector mesons (i.e. J/ψ , ψ' , Υ , Υ' and Υ''), as well as the ϕ meson, will be measured in the $\mu^+\mu^-$ decay channel by the ALICE muon spectrometer. The simultaneous measurement of all the quarkonia species with the same apparatus will allow a direct comparison of their production rate as a function of different parameters such as transverse momentum and collision centrality. In addition to vector mesons, also the unlike-sign dimuon continuum up to masses around $10 \text{ GeV } c^{-2}$ will be studied. Since at LHC energies the continuum is expected to be dominated by muons from the semi-leptonic decay of open charm and open beauty, it will also be possible to study the production of open (heavy) flavours with the muon spectrometer. Heavy-flavour production in the region $-2.5 < \eta < -1$ will be accessible through measurement of $e-\mu$ coincidences, where the muon is detected by the muon spectrometer and the electron by the TRD.

As discussed in section 2, the muon spectrometer will participate in the general ALICE data taking for Pb–Pb collisions at the machine-limited luminosity $\mathcal{L} = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$. The situation is different for intermediate-mass ion collisions (e.g. Ar–Ar) where the luminosity

limitations from the machine are less severe. In this case, beside a general ALICE run at low luminosity $\mathcal{L} = 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$, to match the TPC rate capability, a high luminosity one $\mathcal{L} = 10^{29} \text{ cm}^{-2} \text{ s}^{-1}$ is also foreseen, to improve the Υ statistics. For the high-luminosity run, the muon spectrometer will take data together with a limited number of ALICE detectors (ZDC, ITS Pixel, PMD, T0, V0 and FMD) able to sustain the event rate. These detectors allow the determination of the collision centrality.

The main design criteria of the spectrometer are driven by the following considerations:

- *High multitrack capability:* The tracking detectors of the spectrometer must be able to handle the high particle multiplicity.
- *Large acceptance:* As the accuracy of dimuon measurements is statistics limited (at least for the Υ family), the spectrometer geometrical acceptance must be as large as possible.
- *Low- p_t acceptance:* For direct J/ψ production it is necessary to have a large acceptance at low p_t (at high p_t a large fraction of J/ψ 's is produced via b-decay [448]).
- *Forward region:* Muon identification in the heavy-ion environment is only feasible for muon momenta above $\sim 4 \text{ GeV } c^{-1}$ because of the large amount of material (absorber) required to reduce the flux of hadrons. Hence, measurement of low- p_t charmonia is possible only at small angles where muons are Lorentz boosted.
- *Invariant-mass resolution:* A resolution of 70 (100) $\text{MeV } c^{-2}$ in the 3 (10) $\text{GeV } c^{-2}$ dimuon invariant-mass region is needed to resolve the J/ψ and ψ' (Υ , Υ' and Υ'') peaks. This requirement determines the bending strength of the spectrometer magnet as well as the spatial resolution of the muon tracking system. It also imposes the minimization of multiple scattering and a careful optimization of the absorber.
- *Trigger:* The spectrometer has to be equipped with a selective dimuon trigger system to match the maximum trigger rate of about 1 kHz handled by the DAQ.

3.10.2. Detector layout. The muon spectrometer is designed to detect muons in the polar angular range $2\text{--}9^\circ$. This interval, a compromise between acceptance and detector cost, corresponds to the pseudo-rapidity range of $-4.0 \leq \eta \leq -2.5$.

The spectrometer consists of the following components:

- a passive front absorber to absorb hadrons and photons from the interaction vertex;
- a high-granularity tracking system of 10 detection planes;
- a large dipole magnet;
- a passive muon filter wall, followed by four planes of trigger chambers;
- an inner beam shield to protect the chambers from particles and secondaries produced at large rapidities.

The main challenge for the ALICE muon spectrometer results from the high particle multiplicity per event rather than from the rate. Great care has been taken both in the design of the absorbers (which have to provide strong absorption of the hadron flux coming from the interaction vertex) and of the detectors (which must be able to sustain the remaining high multiplicity). In order to optimize the spectrometer layout, simulations with FLUKA [449], C25 [450] and GEANT3 [400] have been carried out. The particle yields predicted by the HIJING [399] event generator and multiplied by an extra (safety) factor of two have been used as input. The main parameters of the muon spectrometer are summarized in table 3.19. It is worthwhile to note that the muon spectrometer relies on the V0 detector as fast trigger to make the system more robust against background from beam–gas interactions (in particular for pp). A High-Level Trigger (HLT) for dimuons will reduce, by a factor four to five, the need in bandwidth and data storage.

Table 3.19. Summary of the main characteristics of the muon spectrometer.

Muons detection		
Polar, azimuthal angle coverage	$2^\circ \leq \theta \leq 9^\circ, 2\pi$	
Minimum muon momentum	$4 \text{ GeV } c^{-1}$	
Resonance detection		
Pseudo-rapidity coverage	J/ψ	Υ
Transverse momentum range	$-4.0 \leq \eta \leq -2.5$	$-4.0 \leq \eta \leq -2.5$
Mass resolution	$0 \leq p_t$	$0 \leq p_t$
	70 MeV	100 MeV
Front absorber		
Longitudinal position (from IP)	$-5030 \text{ mm} \leq z \leq -900 \text{ mm}$	
Total thickness (materials)	$\sim 10\lambda$ (carbon–concrete–steel)	
Dipole magnet		
Nominal magnetic field, field integral	0.7 T, 3 T m	
Free gap between poles	2.972–3.956 m	
Overall magnet length	4.97 m	
Longitudinal position (from IP)	$-z = 9.87 \text{ m}$ (centre of the dipole yoke)	
Tracking chambers		
Number of stations, number of planes per station	5, 2	
Longitudinal position of stations	$-z = 5357, 6860, 9830, 12\,920, 14\,221 \text{ mm}$	
Anode–cathode gap (equal to wire pitch)	2.1 mm for st. 1; 2.5 mm for st. 2–5	
Gas mixture	80% Ar/20% CO ₂	
Pad size st. 1 (bending plane)	$4 \times 6, 4 \times 12, 4 \times 24 \text{ mm}^2$	
Pad size st. 2 (bending plane)	$5 \times 7.5, 5 \times 15, 5 \times 30 \text{ mm}^2$	
Pad size st. 3, 4 and 5 (bending plane)	$5 \times 25, 5 \times 50, 5 \times 100 \text{ mm}^2$	
Max. hit density st. 1–5 (central Pb–Pb $\times 2$)	$5.0, 2.1, 0.7, 0.5, 0.6 \cdot 10^{-2} \text{ hits cm}^{-2}$	
Spatial resolution (bending plane)	$\simeq 70 \mu\text{m}$	
Tracking electronics		
Total number of FEE channels	1.09×10^6	
Shaping amplifier peaking time	1.2 μs	
Trigger chambers		
Number of stations, number of planes per station	2, 2	
Longitudinal position of stations	$-z = 16\,120, 17\,120 \text{ mm}$	
Total number of RPCs, total active surface	72, $\sim 150 \text{ m}^2$	
Gas gap	single, 2 mm	
Electrode material and resistivity	Bakelite TM , $\rho = 2\text{--}4 \times 10^9 \Omega \text{ cm}$	
Gas mixture	Ar/C ₂ H ₂ F ₄ /i-butane/SF ₆ ratio 49/40/7/1	
Pitch of readout strips (bending plane)	10.6, 21.2, 42.5 mm (for trigger st. 1)	
Max. strip occupancy bend. (non bend.) plane	3%(10%) in central Pb–Pb	
Maximum hit rate on RPCs	3 (40) Hz cm ⁻² in Pb–Pb (Ar–Ar)	
Trigger electronics		
Total number of FEE channels	2.1×10^4	
Number of local trigger cards	234 + 2	

Front absorber, beam shield and muon filter. The front absorber is located inside the L3 magnet. The fiducial volume of the absorber is made predominantly out of carbon and concrete to limit small-angle scattering and energy loss by traversing muons. At the same time, the front-end absorber is designed to protect other ALICE detectors from secondaries produced within the absorbing material itself. The spectrometer is shielded throughout its length by a dense absorber tube surrounding the beam pipe. The tube (beam shield) is made of tungsten, lead

and stainless steel. It has an open geometry to reduce background particle interaction along the length of the spectrometer. While the front absorber and the beam shield are sufficient to protect the tracking chambers, additional protection is needed for the trigger chambers. For this reason, an iron wall about 1 m thick (muon filter) is placed after the last tracking chamber, in front of the first trigger chamber. Front absorber and muon filter stop muons with momentum less than $4 \text{ GeV } c^{-1}$.

Dipole magnet. A dipole magnet with resistive coils is placed about 7 m from the interaction point, outside the L3 magnet. It has a free gap between poles of $\sim 3.5 \text{ m}$, a $\sim 9 \text{ m}$ high yoke, weighting about $\sim 850 \text{ t}$. Its magnetic strength, defined by the requirements on mass resolution, is $B_{\text{nom}} = 0.7 \text{ T}$, and the field integral between interaction point and muon filter is 3 T m .

Tracking chambers. The design of the tracking system is driven by two main requirements: a spatial resolution of about $100 \mu\text{m}$, and the capability to operate in a high-particle multiplicity environment. For central Pb–Pb collisions, a few hundred particles are expected to hit the muon chambers, with a maximum hit density of about $5 \times 10^{-2} \text{ cm}^{-2}$. Moreover, the tracking system has to cover a total area of about 100 m^2 . All these requirements can be fulfilled by the use of cathode pad chambers. They are arranged in five stations: two are placed before, one inside and two after the dipole magnet. Each station is made of two chamber planes. Each chamber has two cathode planes, which are both read out to provide two-dimensional hit information. The first station needs to be flush with the absorber to measure the exit points of the muons as precisely as possible. To keep the occupancy at a 5% level, a large segmentation of the readout pads is needed. For instance, pads as small as $4.2 \times 6 \text{ mm}^2$ are needed for the region of the first station close to the beam pipe, where the highest multiplicity is expected. Since the hit density decreases with the distance from the beam, larger pads are used at larger radii. This enables keeping the total number of channels to about 1 million.

Multiple scattering of the muons in the chambers is minimized by using composite materials (e.g. carbon fibre). The chamber thickness corresponds to about $0.03X_0$. Because of the different size of the stations (ranging from few square metres for station 1 to more than 30 m^2 for station 5), two different designs have been adopted. The first two stations are based on a quadrant structure, with the readout electronics distributed on their surface. For the other stations, a slat architecture has been chosen. The maximum size of the slat is $40 \times 280 \text{ cm}^2$ and the electronics is implemented on the side of the slats. The slats overlap to avoid dead zones on the detector.

The position of the tracking chambers is monitored by a rather sophisticated system of optical lines inspired to the RASNIK [451] concept. The main components of each optical line are a IR LED, a lens and a CCD or CMOS camera. The relative positions of the different chambers are monitored with an accuracy better than $20 \mu\text{m}$ by means of about 100 optical lines, while 160 lines are used to monitor the planarity of the chambers. Monitoring of the position of chamber 1 (chamber 9) with respect to the ITS (ALICE cavern) is also foreseen.

Extensive tests on small-sized and full-sized tracking chamber prototypes (both for quadrant and slat architecture) have shown that these detectors achieve the required performances. In particular, space resolutions of the order of $70 \mu\text{m}$ have been measured.

Trigger chambers. In central Pb–Pb collisions, about eight low- p_t muons from π and K decays are expected to be detected per event in the spectrometer. To reduce to an acceptable level the probability of triggering on events where these low- p_t muons are not accompanied by the high- p_t ones emitted in the decay of heavy quarkonia (or in the semi-leptonic decay of open charm and beauty), a p_t cut has to be applied at the trigger level on each individual muon.

A dimuon trigger signal is issued when at least two tracks above a predefined p_t threshold are detected in an event. According to simulation results, a ‘low’- p_t cut ($1 \text{ GeV } c^{-1}$) will be used for J/ψ and a ‘high’ one ($2 \text{ GeV } c^{-1}$) for Υ selection. To perform the p_t selection, a position-sensitive trigger detector with space resolution better than 1 cm is required. This resolution is achieved by Resistive-Plate Chambers (RPCs) operated in streamer mode [452]. The trigger system consists of four RPC planes arranged in two stations, one metre apart from each other, placed behind the muon filter. The total active area is about 150 m^2 . The RPC electrodes are made of low-resistivity Bakelite ($\rho \sim 3 \times 10^9 \Omega \text{ cm}$), to attain the needed rate capability (maximum expected value about 40 Hz cm^{-2} for Ar–Ar high-luminosity runs, see below). To improve the smoothness of the electrode surface, these are coated with linseed oil. Extensive tests were carried out to study the long-term behaviour of small-sized RPC prototypes. It was shown that RPCs are able to tolerate several ALICE-years of data taking with heavy-ion beams [453]. The x – y coordinates of the RPC hits are read out by segmented strips with pitch and length increasing with their distance from the beam axis.

3.10.3. Front-end electronics and readout. The front-end electronics is based, for all the tracking stations, on a 16-channel chip (MANAS) including the following functionalities: charge amplifier, filter, shaper and track & hold. It also includes a 16 to 1 analogue multiplexer. The channels of four of these chips are fed into a 12-bit ADCs, read out by the MARC chip which includes zero suppression. This chain is mounted on a front-end boards (MANUs). About 17 000 MANU cards are necessary to treat the 1.08 millions channels of the tracking system. Up to 26 MANUs are connected (via PATCH bus) to the translator board which allows the data transfer to the Concentrator ReadOut Cluster Unit System (CROCUS). Each chamber is read out by two CROCUS, leading to a total number of 20 CROCUS. The main tasks of the CROCUS are to concentrate data from the chambers, to ship them to the DAQ and to perform control of the front-end electronics, including calibration and dispatching of the trigger signals.

The RPCs are equipped with dual-threshold front-end discriminators [454] (ADULT), adapted to the timing properties of the detector and reach the time resolution (1–2 ns) necessary for the identification of the bunch crossing. From the discriminators, the signals are sent to the trigger electronics (local trigger cards, based on programmable circuits working in pipeline mode) where the coordinates measured in the first and second stations are compared to determine the muon p_t . Thanks to the short decision time (600–700 ns) of the electronics, the dimuon trigger participates in the ALICE L0 trigger.

Finally, we note that, according to simulations, the total dimuon trigger rate (low- and high- p_t cut) is expected not to exceed 1 kHz, both in Pb–Pb and Ar–Ar collisions. This rate is compatible with the ALICE DAQ bandwidth.

3.11. Zero-Degree Calorimeter (ZDC)

3.11.1. Design considerations. The observable most directly related to the geometry of the collision is the number of participant nucleons, which can be estimated by measuring the energy carried in the forward direction (at zero degree relative to the beam direction) by non-interacting (spectator) nucleons. The zero degree forward energy decreases with increasing centrality. Spectator nucleons will be detected in ALICE by means of Zero-Degree Calorimeters (ZDC). In the ideal case, in which all spectators are detected, the estimate of the number of participants is deduced from the measurement as

$$E_{\text{ZDC}}(\text{TeV}) = 2.76 N_{\text{spectators}}, \quad N_{\text{participants}} = A - N_{\text{spectators}},$$

where 2.76 TeV is the nucleon energy of the Pb beam at LHC. Such a simple estimate can however not be used at a collider since not all the spectator nucleons can be detected.




Figure 3.25. Schematic view of the beam line from ALICE interaction point (IP2) till the ZDC location.




Figure 3.26. Cross section of the beam line 115 m from the IP.

3.11.2. Detector layout. In the ALICE experiment, the ZDCs will be placed at 116 m from the Interaction Point (IP), where the distance between beam pipes (~ 8 cm) allows insertion of a detector (see figure 3.25).

At this distance, spectator protons are spatially separated from neutrons by the magnetic elements of the LHC beam line. Therefore two distinct detectors will be used: one for spectator neutrons, placed at zero degrees relative to the LHC axis, and one for spectator protons, placed externally to the outgoing beam pipe on the side where positive particles are deflected (see figure 3.26).

The quartz fibres calorimetry [455] technique has been adopted for the ALICE ZDC. The shower generated by incident particles in a dense absorber (the so-called 'passive' material) produces Cherenkov radiation in quartz fibres ('active' material) interspersed in the absorber. This technique fulfils two fundamental requirements. Firstly, due to the small amount of space available (particularly for the neutron calorimeter), the detectors need to be compact and therefore a very dense passive material must be used for the absorber to contain the shower. Secondly, the ZDC will operate in a very high radiation environment (about 10^4 Gy day $^{-1}$ is

Table 3.20. Dimensions and main characteristics of absorber and quartz fibres for neutron and proton calorimeters.

	ZN	ZP
Dimensions (cm ³)	7.04 × 7.04 × 100	12 × 22.4 × 150
Absorber	Tungsten alloy	Brass
ρ_{absorber} (g cm ⁻³)	17.61	8.48
Fibre core diameter (μm)	365	550
Fibre spacing (mm)	1.6	4
Filling ratio	1/22	1/65

the estimated dose for the neutron calorimeter at a luminosity of $10^{27} \text{ cm}^{-2} \text{ s}^{-1}$). Radiation hardness is guaranteed by radiation resistance of quartz fibres [456, 457]. Furthermore, Cherenkov effect has two more advantages: it provides a very fast signal due to the intrinsic speed of the emission process and a very low sensitivity to induced radioactivation thanks to its threshold behaviour.

3.11.2.1. Hadron calorimeters. The neutron ZDC (ZN) has the most severe geometrical constraint and therefore the detector transverse dimension should not be greater than 7 cm (~ 1 cm is left for the box containing the detector). For this reason a very dense passive material must be used to maximize containment of showers generated by neutrons of 2.76 TeV.

For the proton calorimeter (ZP) there are no such stringent space constraints and, moreover, spectator protons spot has a large spatial distribution [387, 459].

A larger detector made of a less dense material can therefore be used [460]. Another important characteristic for these detectors is the fibre spacing that should not be greater than the radiation length of the absorber to avoid electron absorption in the passive material. Table 3.20 contains a summary of the main characteristics for the two detectors (further details can be found in [387]).

The energy resolution of the ZDCs is a fundamental parameter in the design of the devices. The physics performance of the detector for the measurement of the centrality of the collision is in fact directly related to the resolution on the number of spectator nucleons which hit the calorimeters' front faces. Therefore a good energy resolution is a necessary prerequisite for a reliable estimation of the centrality variables.

Simulations performed with the GEANT3.21 package [400] for a single nucleon of 2.76 TeV impinging in the centre of the front face gives a resolution value equal respectively to $(8.3 \pm 0.2)\%$ for the neutron calorimeter and to $(9.1 \pm 0.2)\%$ for the proton one (figure 3.27).

As will be shown in detail in sections 5 and 6 of the PPR Volume II, with such an energy resolution one can estimate the impact parameter of the collision with a resolution of the order of 1 fm [462].

Finally, it is well known that the time response of Cherenkov calorimeters is extremely fast. The limiting factor in speed and rate capability of the quartz fibre calorimeters is essentially due to the time response of the light detectors and to the cable length. Recent beam tests on the first ZN calorimeter have shown that the FWHM of the signal of the detector, at the output of a Philips XP2020 PMT, is of the order of 5 ns. We estimate that after 120 m of low-loss cable, as foreseen for ALICE, a FWHM of about 10 ns can be expected. Given the expected collision rate of 8 kHz for Pb–Pb hadronic collisions, we do not foresee any pile-up in the ZDCs. Even the much higher electromagnetic collision rate, which is expected to be of the order of 0.1 MHz, should not pose any particular problem.




Figure 3.27. Resolution of neutron (left) and proton (right) calorimeters for a single nucleon of 2.7 TeV impinging in the centre of detector front faces. The resolution is calculated from the superimposed Gaussian fit.

3.11.2.2. Electromagnetic calorimeter. The ZDC project includes a forward electromagnetic calorimeter to improve the centrality trigger. It is designed to measure, event by event, the energy of particles emitted at forward rapidities, essentially photons generated from π^0 decays.

The detection technique employed for the electromagnetic calorimeter is the same as the one used for the hadronic calorimeters. The most important difference consists in the choice of the angle of the fibres relative to incoming particles. Fibres are oriented at 45° , while for the hadronic calorimeters they are at 0° . This choice maximizes the detector response, since Cherenkov light production has a pronounced peak around 45° [458]. The electromagnetic calorimeter is made of lead, with quartz fibres sandwiched in layers between the absorber plates. Two consecutive planes of fibres are separated by a lead thickness of 3 mm which, due to the 45° inclination, results in a total thickness of 4.24 mm seen by incident particles. Fibre cores have a diameter of $550 \mu\text{m}$ and the active to passive volume ratio is about 1/11. The resulting calorimeter dimensions are $7 \times 7 \times 21 \text{ cm}^3$, the total absorber length corresponding to about 30 radiation lengths. This condition assures the total containment of showers generated by participants [461].

In the ZDC Technical Design Report [387], the proposed position for the detector was at ZDC location (116 m from the IP) above the neutron calorimeter, just outside the cone of spectator neutrons. The pseudo-rapidity range covered by the two detectors, placed on both sides relative to IP, is between 7.8 and 9.2 m.

Studying the acceptance of the detector and analysing the correlation between electromagnetic calorimeter and ZDC energies (figure 3.29) it has emerged that the proposed location of the electromagnetic calorimeter is far from being optimal.

A way to increase detector acceptance and to obtain a narrower correlation is to move the electromagnetic calorimeters closer to the IP, before any magnetic elements of the LHC beam line. In this way, shifting the detector acceptance to more central rapidities, the electromagnetic calorimeter can detect a significant number of particles produced in the interaction. In the original position the optics elements between the IP and the ZDC location drastically reduce the intercepted solid angle and therefore the number of detectable secondaries. The measurement would be substantially dominated by the high level of fluctuations in number of detected photons. The increase in detected energy obtained with the shift of detector acceptance leads to an improvement of the correlation between the measured energy and the impact parameter. Therefore the new position chosen for the calorimeters is at about 7 m from the IP, just before the coils of the first compensator dipole. Two devices will be placed on the same side relative




Figure 3.28. Energy detected by two zero-degree electromagnetic calorimeters as a function of the impact parameter.




Figure 3.29. Correlation between energy measured by one ZEM and one ZDC set (left) for ZEM located at 116 m from IP and (right) for the new proposed position at 10 m from IP.

to the IP, on the side opposite to the dimuon arm. In this configuration the electromagnetic devices cover the pseudo-rapidity range $4.8 < \eta < 5.7$.

The correlation between energy measured by the two forward electromagnetic calorimeters (called ZEM since they belong to the ZDC project) and the impact parameter is shown in figure 3.28.

The correlation between the total energy detected in the electromagnetic calorimeters in the new position and the energy measured by ZDC is shown in figure 3.29. The improvement in the definition of the correlation is evident.

To estimate the energy resolution that can be achieved with the electromagnetic forward calorimeter for the typical ALICE range of energies in the rapidity domain $4.8 < \eta < 5.7$, simulations have been performed using HIJING [399] event generator and GEANT3 [400] detector response simulation package.

The results of the simulations show that in central collisions ($b < 2$ fm) the total incident energy on the two electromagnetic calorimeters is about 7 TeV, while in more peripheral interactions ($b \sim 10$ fm) it is of the order of 1.5 TeV. Considering the resolution experimentally measured for a prototype on an SPS test beam,

$$\frac{\sigma}{E} = \frac{0.69}{\sqrt{E(\text{GeV})}},$$

an energy resolution $< 1\%$ for central collisions is obtained. This value increases up to 1.8% in the most peripheral events. This result ensures that the designed detector can be successfully adopted at ALICE for the measurement of secondaries at forward rapidity.

3.11.3. Signal transmission and readout. The quartz fibres emerging from the rear face of the calorimeters are gathered together in bundles, which are routed towards the photomultipliers, located 50 cm far from the end of the calorimeter outside the beam horizontal plane. The fibre bundle terminates in a rectangular box coupled to the photocathode. The signal is transmitted through long low-loss coaxial cables (165 m for the right set of ZDCs and 195 m for the left one) from the photomultipliers to the lowest counting room, where the trigger logic is built. The signals from all the photomultipliers of all the hadronic calorimeters are summed, in order to obtain a signal proportional to the total number of spectators coming out from the interaction. Three levels of discrimination will be applied to provide three Level 1 (L1) triggers, defining different centrality intervals: the first will select the most central events (10% of the total inelastic cross section), the second the semi-central events (15% of the total inelastic cross section) and the third the minimum bias events. The ZDCs cannot provide a Level 0 (L0) trigger since they are located too far from the interaction point. Finally, each analogue signal from the photomultipliers will be sent to commercial ADC modules. When a L0 trigger is issued, the ZDC electronics will start to convert the signals and make them available for the DAQ if a positive L1 trigger is received.

3.12. Photon Multiplicity Detector (PMD)

The Photon Multiplicity Detector (PMD) [391, 392] is a preshower detector that measures the multiplicity and spatial (η - ϕ) distribution of photons on an event-by-event basis in the forward region of ALICE [382]. The PMD addresses physics issues related to event-by-event fluctuations, flow and formation of Disoriented Chiral Condensates (DCC) and provides estimates of transverse electromagnetic energy and the reaction plane on an event-by-event basis [391].

3.12.1. Design considerations. Measurement of photon multiplicity in the high particle density environment of ALICE requires optimization of granularity and converter thickness so that overlap of photon showers is minimum. To minimize the occupancy, it is also necessary that charged particles are essentially confined to one cell. In addition efficient photon hadron discrimination technique is required to improve the efficiency, the photon reconstruction and the purity of the photon sample.

The PMD consists of two identical planes of detectors with a $3X_0$ thick lead converter in between them as shown schematically in figure 3.30. The front detector plane is used for vetoing charged-particle hits. The detector plane behind the converter is the preshower plane which registers hits from both photons and charged hadrons.

The choice of detector technology is based on the following considerations:

- the active volume of the detector should be thin and very close to the converter so that the transverse spread of the shower is minimized;
- low-energy δ -electrons should be prevented from travelling to nearby cells and causing significant crosstalk among adjacent channels;
- the detector material (gas) should be insensitive to neutrons; in a hydrogenous medium neutrons tend to produce large signals due to recoil protons, which can mimic a photon signal;




Figure 3.30. Cross section of the PMD (schematic only) showing the veto plane, lead converter and the preshower plane. SS is the support plate on which lead plates and the detectors will be mounted.

- the technology should permit the construction of large chambers;
- charged particles should be confined preferably to one cell so that the occupancy does not increase significantly. If the signal spreads to neighbouring cells then there is a significant probability of vetoing nearby photons.

A novel design of a gas proportional counter having honeycomb structure and wire readout has been developed to meet the above design criteria. The honeycomb cells are physically isolated from each other by thin metallic walls to contain δ -rays. The metallic wall of the honeycomb forms the common cathode and is kept at a large negative potential. The individual anode wires in the cells are kept at ground potential and connected to the readout electronics. The honeycomb geometry provides almost circular equipotentials within a cell and facilitates close packing of large arrays. Further field shaping has been achieved by extending the cathode onto the inside of the printed circuit boards covering the honeycomb. The cross section of a cell is shown in figure 3.31. The minimum anode–cathode distance is close to 0.1 cm. This ‘extended cathode’ design provides almost uniform charge collection within a cell independent of the point of impact of ionizing radiation.

The operating gas in the chamber is selected to be a mixture of Ar and CO₂ because of its insensitivity to neutrons and its non-inflammable nature. The efficiency for charged-particle detection is almost uniform within the cell and approaches 96%. The signal of individual charged particles is almost always confined to one cell. On the preshower plane the transverse spread of the shower is close to that given by simulation [391, 392]. This ensures that the actual occupancy in the detector will be close to that given by simulation results.

3.12.2. Detector layout. The PMD (figure 3.32) will be installed at 360 cm from the interaction point, on the opposite side of the forward muon spectrometer, covering the region $2.3 \leq \eta \leq 3.5$. This region has been selected to minimize the effect of upstream material such as the beam pipe and the structural components of TPC and ITS. Important parameters of the PMD are listed in table 3.21.

The PMD consists of four supermodules, two of each kind. The supermodule is a gas-tight enclosure. There are six identical unit modules in each supermodule. The unit modules (figure 3.33) are separated among themselves by a thin 100 μm kapton strip supported on




Figure 3.31. A schematic diagram of the cross section of a unit cell. The diameters of the insulation circle and the wire support are shown. The hexagonal section depicts a region of the PCB which provides cathode extension.




Figure 3.32. Layout of the PMD showing four supermodules. Each supermodule has six unit modules. The detector has full azimuthal coverage in the region $2.3 \leq \eta \leq 3.5$. The inner hole is 22 cm \times 20 cm.


a 0.3 mm thick FR4 sheet for rigidity. This provides high voltage isolation among the unit modules.

The honeycomb, made of copper sheet, has 48×96 cells covered on both sides by printed circuit boards (PCBs). The inner parts of the PCB provide extensions to the honeycomb cathode. The upper part of the top PCB has tracks from 4 rows and 8 columns of the cell centres (connecting anode wires) merging into 32-pin connectors. After wire insertion and soldering, the PCBs form gas-tight planes. Two 32-pin connectors are joined together by a flexible kapton cable which connects to the front-end electronics (FEE) board.

The PMD electronics will be cooled by convection using chilled air. For efficient cooling each half of the detector will have suitable enclosures made of polythene sheets with zip

Table 3.21. Summary of design and operating parameters of the PMD.

Pseudorapidity coverage	$2.3 \leq \eta \leq 3.5$
Azimuthal coverage	2π
Distance from vertex	361.5 cm
Detector active area	2 m^2
Detector weight	1200 kg
Number of planes	two (Veto and Preshower)
Converter	$3X_0$ lead
Hexagonal cell cross section	0.22 cm^2
Hexagonal cell depth (gas thickness)	0.5 cm
Detector gas	Ar/CO ₂ (70%/30%)
Operating voltage	-1400 V
Charged-particle detection efficiency	96%
Number of supermodules per plane	4
Number of unit modules per supermodule	6
Number of cells in a unit module	4608
Number of HV channels	48
Total number of cells	221 184
Number of FEE boards	3456
Number of CROCUS crates	4
Number of DDL channels	4
Cell occupancy for $dN_{\text{ch}}/d\eta = 8000$ for veto plane	13%
Cell occupancy for $dN_{\text{ch}}/d\eta = 8000$ for preshower plane	28%
Average photon reconstruction efficiency	54%
Average purity of photon sample	65%
Event size for $dN_{\text{ch}}/d\eta = 8000$	0.12 MB
Average time for readout of event	$100 \mu\text{s}$

**Figure 3.33.** Components of a unit module: (1) top PCB, (2) 32-pin connector, (3) HV isolation and gas seal, (4) copper honeycomb, (5) bottom PCB showing islands of insulation circles.

fasteners on both sides to have easy access for servicing the electronics. To reduce vibration due to flowing air, baffles will be used at the entry points of chilled air.

The PMD is designed in two equal halves. Each half has independent cooling, gas supply and readout accessories. The suspension mechanism is made entirely of non-magnetic material and has only manual controls.

3.12.3. Front-end electronics and readout. The front-end electronics and readout of the PMD is based on the design of the electronics for the tracking chambers of the muon spectrometer [391]. The detector signals will be processed using the MANAS chip which handles 16 channels providing multiplexed analogue signals. The FEE board uses four MANAS chips, two ADCs and a MARC chip to handle 64 channels of the detector. A set of FEE boards will be read out using Digital Signal Processors (DSP). All the FEE boards in a row in a super-module are connected to a single long PCB which provides mechanical support to the FEE boards, supplies low-voltage power, carries the data and control signals between the MARC and the DSP and forms part of the link-bus for the DSP.

The DSPs are handled through a cluster readout system CROCUS, developed for the tracking chambers of the muon spectrometer. Each DSP has five link-buses. The number of FEE boards in a link-bus varies from 9 to 36 depending on the occupancy of the cells. The maximum occupancy in a link bus is around 250, which results in a dead time of about $100 \mu\text{s}$. A CROCUS crate handles 10 DSPs. The PMD readout will require four such crates and four DDL links.

3.13. Forward Multiplicity Detector (FMD)

3.13.1. Design considerations. The main functionality of the silicon strip Forward Multiplicity Detector (FMD) is to provide (offline) charged-particle multiplicity information in the pseudo-rapidity range $-3.4 < \eta < -1.7$ and $1.7 < \eta < 5.1$.

The FMD will allow for the study of multiplicity fluctuations on an event-by-event basis and for flow analysis (relying on the azimuthal segmentation) in the considered pseudo-rapidity range. Together with the pixel system of the ITS, the FMD will provide early charged particle multiplicity distributions for all collision types in the range $-3.4 < \eta < 5.1$. Overlap between the various rings and with the ITS inner pixel layer provides redundancy and important checks of analysis procedures.

The average number of hits for very central Pb–Pb collision, assuming the extreme charged-particle multiplicity density of $dN_{\text{ch}}/d\eta \approx 8000$, and including background from secondary interactions, will be less than 3 charged particles per strip for all channels. The majority of channels will on average be traversed by about one charged particle per central event. For these, multiplicity information may be obtained by comparing the number of occupied and empty channels. In general, however, multiplicity information will be obtained by measuring the energy deposition in each channel and relating this to the number of charged particles. The readout time of the system does not allow the FMD to deliver a multiplicity trigger at L0 or L1.

3.13.2. Detector layout. The FMD consists of 51 200 silicon strip channels distributed over five ring counters of two types with each 20 or 40 sectors in azimuthal angle, respectively. Each sector will be read out independently and contains 512 or 256 detector strips at constant radius.

The individual Si sensors will be manufactured out of $300 \mu\text{m}$ thick 15 cm diameter Si wafers, each covering two sectors. The layout permits the coverage of the desired pseudo-rapidity range with two designs of ring counters (so-called inner and outer ring).




Figure 3.34. Layout of Si1 (inner) and Si1 (outer) rings of the FMD on the muon absorber side of the IP. The figure also shows the location of the T0 and V0 detectors.

Figure 3.34 shows the arrangement for the Si detectors closest to the IP on the muon absorber side where the available space is most restrained. Coverage from $\eta = -3.4$ to -1.7 is obtained with an inner and an outer ring (labelled Si1 inner and Si1 outer). The design assures a pseudo-rapidity overlap between the two rings of $\eta = 0.2$.

Coverage in the range $1.7 < \eta < 3.7$, on the opposite side of the IP, is similarly obtained with an inner ring and outer ring of identical design (labelled Si2 inner and Si2 outer). Coverage of the high pseudo-rapidities in the range $3.7 < \eta < 5.1$ is achieved by a fifth ring, of the ‘inner’ design (labelled Si3) placed at 340 cm from the IP.

Table 3.22. The distance z , from the detector to the IP, the inner and outer radii, and the resulting pseudo-rapidity coverage of each ring.

Ring	z (cm)	R_{in} (cm)	R_{out} (cm)	η coverage
Si1 outer	-75.2	15.4	28.4	$-2.29 < \eta < -1.70$
Si1 inner	-62.8	4.2	17.2	$-3.40 < \eta < -2.01$
Si2 outer	75.2	15.4	28.4	$1.70 < \eta < 2.29$
Si2 inner	83.4	4.2	17.2	$2.28 < \eta < 3.68$
Si3	340.0	4.2	17.2	$3.68 < \eta < 5.09$




Figure 3.35. Pseudo-rapidity coverage of the five rings of the FMD. The coverage of the two pixel layers of the ITS is also shown.

Table 3.22 summarizes the positions of the various FMD rings and the corresponding pseudo-rapidity coverages. The combined coverage of the FMD is shown in figure 3.35. Also shown in this figure is the coverage of the ITS pixel layers. The design ensures, together with the ITS inner pixel layer, a full pseudo-rapidity coverage in the range $-3.4 < \eta < 5.1$, and an overlap between the FMD and ITS pixel system of about $\eta = 0.2$. There is no overlap with the second ITS pixel layer (also indicated in figure 3.35).

The ability to extract exact charged particle multiplicity information depends on the level of multiple hits on individual detector strips and is deteriorated in the presence of secondary particles generated in material between the IP and the detector. The expected flux of charged particles per cm^2 at the positions of the various FMD modules is displayed in figure 3.36.

The figures were obtained from Monte Carlo calculations with the standard AliRoot 3.06.02 using as input central Pb–Pb events generated with HIJING [399]. Each panel shows the primary charged particle flux at the position of an FMD ring, together with the expected flux of secondary charged particles originating in the beam pipe, the ITS detector and its services, the T0 and V0 detectors, and the material of the muon absorber and various mechanical structures.

The choice of segmentation of the FMD is driven by the requirement to keep the average number of hits per strip well below 2–3 particles for most strips in order to enable an accurate multiplicity reconstruction based on total energy deposition. When taking into account fluctuations away from the average multiplicities, it follows that the front-end electronics must be adapted to handle a maximum signal deposition corresponding to about 20 MIPs. The results from the Monte Carlo simulations translate this maximum hit multiplicity into limits on the strip area to about 0.03 cm^2 at the innermost radii and about 0.3 cm^2 at the outer radii. We have chosen to subdivide each detector sensor into two azimuthal sectors and use a radial segmentation fine enough to meet these maximum areas. From these considerations, we arrive at the detector segmentation given in table 3.23.

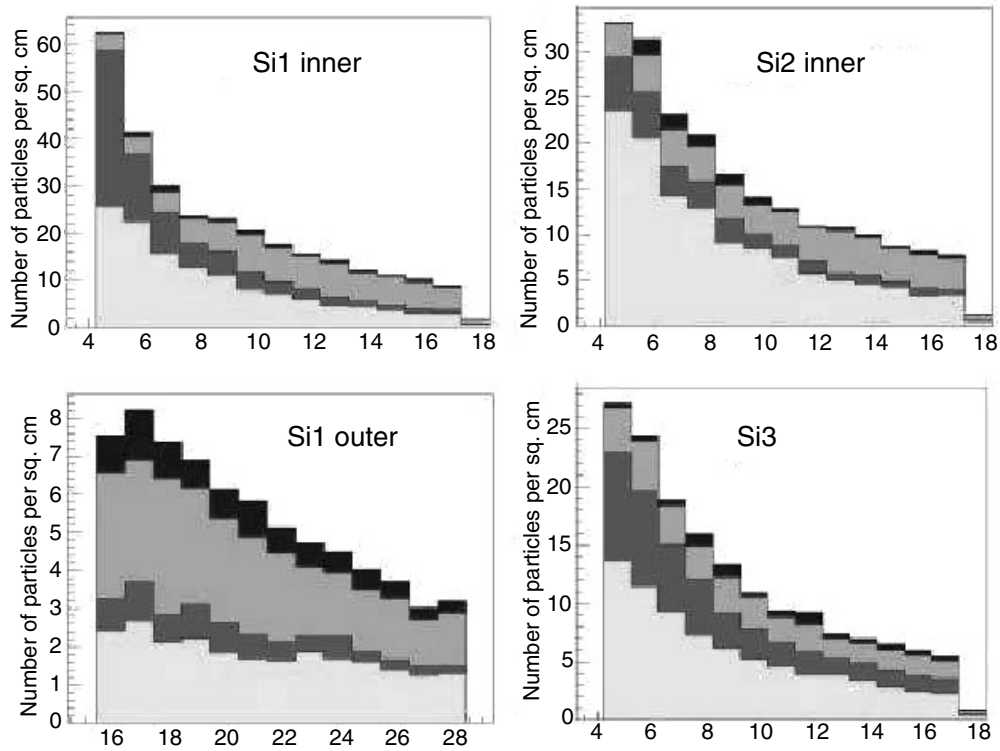


Figure 3.36. Flux of charged particles at the position of four of the five FMD rings simulated using central Pb–Pb events, as function of the radius from the centre of the beam. In each graph, primary charged particles are shown at the bottom, followed upward by the expected background of secondary particles originating from the beam pipe, the ITS detector and services, and the TO and V0 detectors including various support structures.

Table 3.23. Physical dimensions of Si segments and strips, together with the average number of charged particles hitting each strip in central Pb–Pb collisions.

	Radial coverage (cm)	Particle flux (cm^{-2})	Azimuthal sectors	Radial strips	Strip area (cm^2)	Average number of hits
Si1 inner	4.2–17.2	10–65	20	512	0.03–0.14	2.2–1.4
Si1 outer	15.4–28.4	3–8	40	256	0.12–0.23	1.0–0.7
Si2 inner	4.2–17.2	8–35	20	512	0.03–0.14	1.2–1.1
Si2 outer	15.4–28.4	3–8	40	256	0.12–0.23	1.0–0.7
Si3	4.2–17.2	6–27	20	512	0.03–0.14	0.9–0.8

3.13.3. Front-end electronics and readout. The role of the Front-End (FE) electronics is to amplify and shape the signals from the Si strips. To assure the best possible signal-to-noise ratio, the amplifiers must be located close to the detector. This part of the electronics will be highly integrated circuits mounted on a hybrid board onto which the Si sensor itself will be firmly attached. Strips will be bonded to the hybrid. In addition to the electrical functions, the hybrid substrate serves as mechanical support for the Si sensor and the corresponding FE amplifier chips, as shown in figure 3.37.




Figure 3.37. Proposed assembly of an inner Si ring showing mechanical support plates and hybrid cards with each one Si sensor. The sensors will be glued to the hybrid cards and the strips wire bonded to the hybrids. The glued segments are attached to the support plate by screws on the feet. Shown on the hybrid cards are FE chips along the two radial edges, other electronic components and connectors for the cables. The support plates will also carry the FMD Digitizer boards on the back side.

Table 3.24. The total number of strip channels per ring and the number of front-end (FE) preamplifier chips (128 channels per chip are listed). Also shown is the number of ALTRO readout chips for each ring and the number of digitizer cards envisaged to carry these. One readout controller unit (RCU) is foreseen at each side of the IP.

	FE channels	FE chips	ALTRO chips	FMD digitizers	RCU modules
Si1 inner	10 240	80	6	2	1
Si1 outer	10 240	80	6	2	
Si2 inner	10 240	80	6	2	1
Si2 outer	10 240	80	6	2	
Si3	10 240	80	6	2	
Total system	51 200	400	30	10	2

Figure 3.38 shows how the electronics chain is foreseen to be built and table 3.24 summarizes the number of electronics channels and modules in the system. An electronics board mounted on each half-ring assembly carries ADC chips for the digitization of the signal amplitude of each strip and multiple-event buffers. This digitizer board also serves as readout controller for the FE chips and the distribution of power and L0 trigger for a half-ring. The digitized signals are read out by one readout controller unit (FMD-RCU) at each side of ALICE which also interfaces to the DAQ, DCS and trigger systems.

The preamplifier–shaper integrated circuit that was chosen is the VA1'' chip which is a radiation hard and 128 channel version of the well proven VA1 amplifier. The characteristics of this amplifier is a good match to the FMD strip sensors. It is a low-noise amplifier (r.m.s. noise of 1.5% MIP for 25 pF input capacitance) with a 1–2 μ s speaking time and a dynamic range

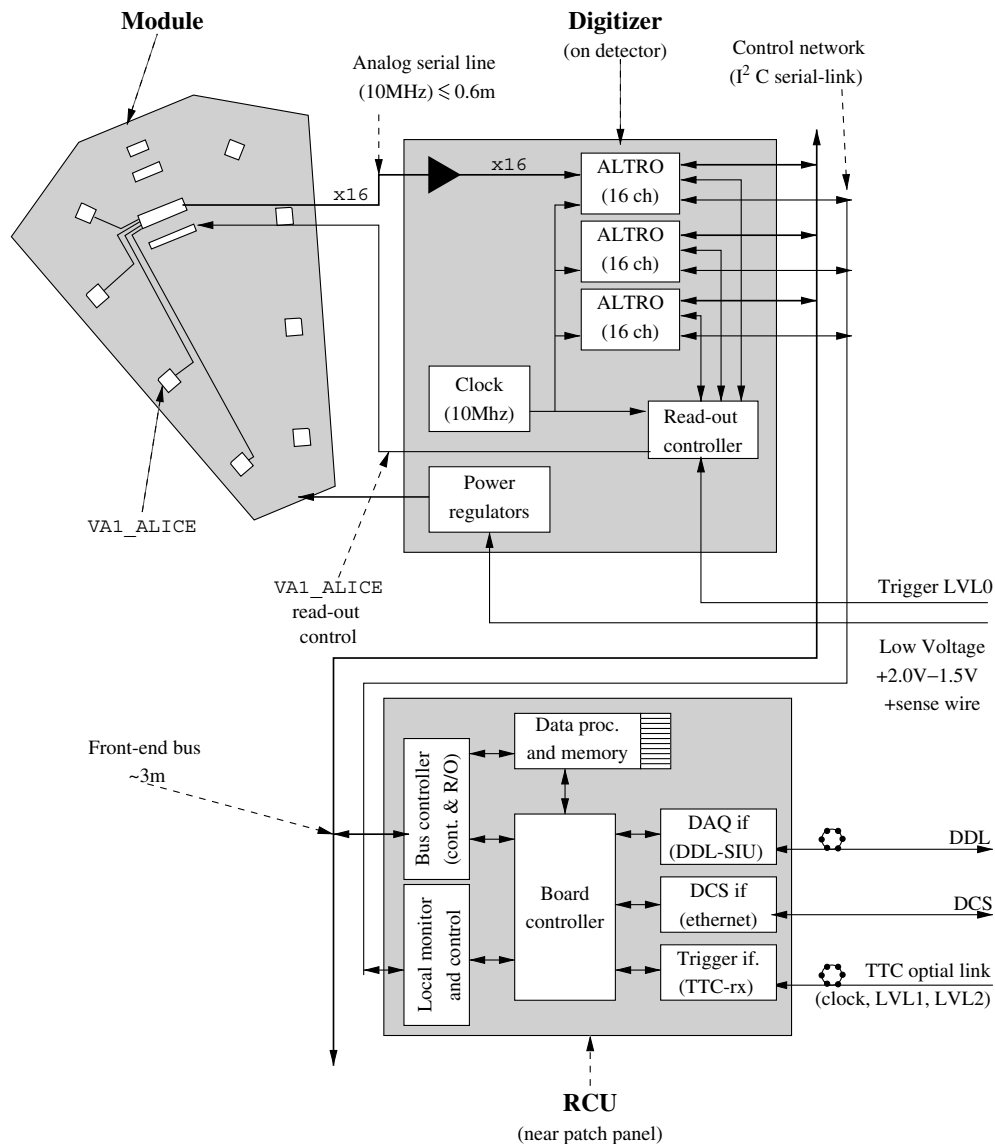


Figure 3.38. Schematics of the full electronics chain of the FMD detectors. A silicon detector segment with its hybrid is shown, followed by a FMD Digitizer card with one ADC channel per FE amplifier chip, the FMD Readout Controller Unit (FMD-RCU) and the DDL, DCS and Trigger links.

of 0–20 MIPS. The chip incorporates a sample-and-hold circuit for each of the 128 channels and a multiplexed readout at a speed of 10 MHz, resulting in a full readout time of $12.8 \mu\text{s}$.

Each FE chip is read out by a separate ADC channel, running at the same clock frequency of 10 MHz as the FE multiplexing rate. The 16-channel ALTRO circuit developed for the TPC detector will be used for this on-detector digitization. Its 10-bit resolution is a good match for the level of FE amplifier noise given above.

The use of the ALTRO essentially allows to copy the electronics of the TPC for the remainder of the FMD readout and DCS systems. Thus, the FMD-RCU unit is almost identical

Table 3.25. V0A and V0C arrays. Pseudo-rapidity coverage and angular acceptance (in deg) of the rings.

Ring	V0A		V0C	
	η_{\max}/η_{\min}	$\theta_{\min}/\theta_{\max}$	η_{\max}/η_{\min}	$(\pi-\theta)_{\min}/(\pi-\theta)_{\max}$
1	5.1/4.5	0.7/1.3	-3.7/-3.2	2.8/4.7
2	4.5/3.9	1.3/2.3	-3.2/-2.7	4.7/7.7
3	3.9/3.4	2.3/3.8	-2.7/-2.2	7.7/12.5
4	3.4/2.8	3.8/6.9	-2.2/-1.7	12.5/20.1

to the similar unit in the TPC, as are the links to the experiment-wide DAQ, DCS and Trigger systems.

The ALTRO incorporates a multiple-event buffer on the chip and the possibility of zero suppression. This suppression feature will not be useful for the high multiplicity central Pb–Pb events where the event size will be 60–100 kB per event without any packing (10–16 bits per channel). However, for more peripheral collisions or in pp running, the zero-suppression feature might be exploited and result in considerable reduction of event sizes.

3.14. V0 detector

3.14.1. Design considerations. The V0 detector [396] has multiple roles. It provides:

- a minimum bias trigger for the central barrel detectors;
- two centrality triggers in Pb–Pb collisions;
- a centrality indicator;
- a control of the luminosity (see section 2);
- a validation signal for the muon trigger [464] to filter background in pp mode.

Special care must be taken to minimize the background due to the location of the V0 detector. Indeed, the presence of important material volumes (beam pipe, front absorber, FMD, T0, ITS services) in front of the V0 arrays will generate an important number of secondaries (mainly electrons) which will affect physical information about the number of charged particles. The efficiency of minimum bias triggering and the multiplicity measurement will be strongly modified by this secondary particle production. Beam–gas interactions will be another source of background. It will provide triggers which have to be identified and eliminated. This background is particularly important in pp runs. Measuring the time-of-flight difference between two detectors located on each side of the interaction point will enable to identify these background events. The V0 detector must therefore provide signal charge and time-of-flight measurement capabilities.

3.14.2. Detector layout. The V0 detector [463] is made of two arrays (V0A and V0C) located asymmetrically on each side of the interaction point. The V0A is located 340 cm from the vertex, on the side opposite to the muon spectrometer. The V0C is fixed at the front face of the front absorber, 90 cm from the vertex.

The V0A/V0C are segmented (table 3.25, figure 3.39) into 32 elementary counters distributed in four rings. Each ring covers 0.4–0.6 unit of pseudo-rapidity. The rings are divided into eight sectors of 45°. The elementary counter consists of scintillator material with embedded WaveLength Shifting (WLS) fibres. The light from the WLS is collected by clear




Figure 3.39. Segmentation of the V0A/V0C arrays.

fibres and transported to PhotoMultiplier (PM) installed at 3–5 m from the detectors, inside the L3 magnet. The time resolution of each individual counter will be better than 1 ns.

The origin of background from secondaries is mainly the beam-pipe section equipped with bellows and the flange in front of the detector. The V0C array is particularly exposed to these secondaries. This background will largely modify the hit distribution [396], the inner ring being much more affected than the outer ones. In pp mode, the efficiency for the detection of at least one charged particle detected in both V0 arrays is about 77% when no secondary particle effects are taken into account. It rises up to 82% when the real environment effects are introduced [396]. Triggering with either the left or the right V0 would increase the efficiency by about 5%. Unfortunately, the large background from beam–gas interactions prohibits the use of this trigger configuration. The measurement of the relative time-of-flight between the two detector arrays will be necessary to help the identification of beam–gas interactions. When the beam–gas collision is localized outside the V0A–V0C interval, the time difference is about 6 ns larger than the one given by the beam–beam interactions. When the beam–gas collision is localized inside the V0A–V0C interval, only an offline analysis will allow to disentangle this background from the beam–beam interaction [396].

Obviously, for Pb–Pb collisions where the multiplicity of charged particles is much larger, the trigger efficiency is close to 100%. Moreover, the contribution from beam–gas background can be neglected.

3.14.3. Front-end electronics and readout. The front-end electronics is, at the time of writing, still in a development stage. Therefore, only the general principles can be described. It will provide a minimum-bias trigger for all reactions and two centrality triggers in the case of Pb–Pb collisions. A fast decision on the minimum cells fired in the collision can be applied if necessary. Individual integrated charges from PMs and time between signals and bunch clock will be digitized for the offline analysis. The electronics system will provide a dynamical range of about 1000 for each V0 channel. The important background will not affect the timing performances of the V0, neither for the minimum bias trigger, nor for the centrality triggers. The central and semi-central triggers are defined by a discrimination level on the added integrated charges supplied by all the PMs. Despite the contribution of background in the collected charge, it will be possible to define these two triggering levels in ion–ion mode.




Figure 3.40. The layout of T0 detector arrays inside ALICE.

3.15. T0 detector

3.15.1. Design considerations. The T0 detector [465] has to perform the following functions:

1. To generate a T0 signal for the TOF detector. This timing signal corresponds to the real time of the collision (plus a fixed time delay) and is independent on the position of the vertex. The required precision of the T0 signal is about 50 ps (r.m.s.).
2. To measure the vertex position (with a precision ± 1.5 cm) for each interaction and to provide a L0 trigger when the position is within the preset values. This will discriminate against beam–gas interactions.
3. To provide an early ‘wake-up’ signal to TRD, prior to L0.
4. To measure the particle multiplicity and generate one of the three possible trigger signals: $T0_{\text{min-bias}}$, $T0_{\text{semi-central}}$, or $T0_{\text{central}}$.

Since the T0 detector generates the earliest L0 trigger signals, they must be generated online without the possibility of any offline corrections. The dead time of the detector should be less than the bunch-crossing period in pp collisions (25 ns).

3.15.2. Detector layout. The detector consists of two arrays of Cherenkov counters, 12 counters per array. Each Cherenkov counter is based on a Russian made fine-mesh photomultiplier tube FEU-187, 30 mm in diameter, 45 mm long optically coupled to a quartz radiator 30 mm in diameter and 30 mm thick. One of the arrays, labelled $T0_R$ in figure 3.40, is placed 70 cm from the nominal vertex. Such a small distance had to be chosen because of the space constraints imposed by the front cone of the muon absorber and other forward detectors. The pseudo-rapidity range of $T0_R$ is $2.9 \leq \eta \leq 3.3$. On the opposite side the distance of the left array (labelled $T0_L$ in figure 3.40) is about 350 cm—comfortably far from the congested central region. $T0_L$ is grouped together with the other forward detectors (FMD, V0, and PMD) and covers the pseudo-rapidity range of $-5 \leq \eta \leq -4.5$. In the radial (transverse) direction both T0 arrays are placed as close to the beam pipe as possible to maximize triggering efficiency.

The triggering efficiency of the detector for minimum bias pp collisions, estimated by Monte Carlo simulations, is about 48% for all inelastic processes. The efficiencies of $T0_R$ and $T0_L$ for pp inelastic processes are 60 and 67%, respectively. The triggering efficiency in heavy ion collisions is, due to the high multiplicities, practically 100%. The basic parameters of the T0 detector are summarized in table 3.26.

Table 3.26. Overview of the T0 detector parameters.

Parameters	Left array	Right array
z -position (cm)	-350	+70
Number of Cherenkov counters	12	12
Pseudo-rapidity coverage	$-5 \leq \eta \leq -4.5$	$2.9 \leq \eta \leq 3.3$
Detector active area (cm ²)	84.78	84.78
Efficiency with beam pipe (%)	67	60
Efficiency with beam pipe of both arrays in co-incidence (%)		48
Number of physical readout channels		56
Time resolution (ps)		37
Vertex position resolution (cm)		1.3

3.15.3. Front-end electronics and readout. The T0 electronics consists of the front-end electronics, located close to the arrays, inside the so-called shoe boxes, and the main T0 electronics, placed inside the T0 electronic rack, already outside of the L3 magnet, close to the racks assigned to ALICE trigger electronics. All amplitude and timing signals will be digitized and stored by ALICE DAQ.

The right and left arrays generate $T0_R$ and $T0_L$ timing signals correspondingly which are fed directly to the TRD, to be used as a pre-trigger ‘wake-up’ signal. Due to the length of the connecting cables the coincidence requirement for the ‘wake-up’ is implemented not in the T0 electronic rack but directly at the TRD.

The T0 timing signal is generated by a mean timer. The position of the T0 signal on the time axis is equal to $(T0_R + T0_L)/2 + T_{\text{delay}}$, where T_{delay} is the fixed delay of the analogue time meaner. Thus the T0 signal gives a start signal for the TOF detector independent on the position of the vertex.

The position of the vertex is measured as $T0_L - T0_R$ and this value is fed to a digital discriminator with preset upper and lower limits thus providing the $T0_{\text{vertex}}$ trigger signal.

$T0_{\text{min-bias}}$, $T0_{\text{semi-central}}$ and $T0_{\text{central}}$ multiplicity trigger signals are generated by three discriminator levels applied to the linear sum of the amplitudes from all the detectors in the array.

In order to achieve very high online time resolution of about 50 ps the T0 detector must be equipped with sophisticated fast timing electronics [466]. The most crucial elements of T0 electronics are: the timing discriminators, the mean timer, and the module generating $T0_{\text{vertex}}$ signal for the trigger. All these components should have dead time below 25 ns, good stability and long-term reliability. From the discriminators a small time walk (< 25 ps) and a very broad dynamic range of the input amplitudes (about 1 : 500) are required. The prototypes of these modules have been designed, built and tested using a picosecond diode laser [467], and PS beams at CERN.

The prototype of the timing constant fraction discriminator had the time walk ± 25 ps in the full dynamic range (1 : 500). The parameters of the mean timer and the $T0_{\text{vertex}}$ module were measured in beam test using two Cherenkov counters emulating the two arrays, and an additional START Cherenkov counter. The time resolution of each counter was measured to be about $\sigma = 37$ ps. The position of the mean timer output signal was stable on the time axis within ± 10 ps. The vertex position resolution of the trigger module was also very good ($\sigma = 1.3$ cm) in the full range of allowed positions.

The readout electronics consists of the TRM (TDC Readout Module) and DRM (Data Readout Module) cards. The TRM card houses the HPTDC 8-channel chips. Both modules are quite similar to those used in TOF 3.7.3. All in all 56 physical channels will be stored.

In addition to time and amplitude for each of the 24 PMTs, the main signals produced by T0 electronics will also be included: interaction time, vertex, $T0_R$, $T0_L$, linear sum of amplitudes and 3 multiplicity signals. To use TOF readout scheme while allowing 25 ns repetition rate required for T0 signals, they have to be demultiplexed. This increases the number of the needed readout channels accordingly.

3.16. Cosmic-ray trigger detector

3.16.1. Design considerations. The cosmic-ray trigger for ALICE will be part of ACORDE (A COsmic Ray DEtector for ALICE) which together with some other ALICE tracking detectors, will provide precise information on cosmic rays with primary energies around 10^{15} – 10^{17} eV [468].

The Cosmic-Ray Trigger (CRT) system will provide a fast L0 trigger signal to the central trigger processor, when atmospheric muons impinge upon the ALICE detector. The signal will be useful for calibration, alignment and performance of several ALICE tracking detectors, mainly the TPC and ITS. The cosmic-ray trigger signal will be capable to deliver a signal before and during the operation of the LHC beam. The typical rate for single atmospheric muons crossing the ALICE cavern will be less than 3 – 4 Hz m^{-2} [469]. The rate for multi-muon events will be lower (less than 0.04 Hz m^{-2} [469]) but sufficient for the study of these events provided that one can trigger and store tracking information from cosmic muons in parallel to ALICE normal data taking with colliding beams. The energy threshold of cosmic muons arriving to the ALICE hall is approximately 17 GeV, while the upper energy limit for reconstructed muons will be less than 2 TeV, depending of the magnetic field intensity (up to 0.5 T).

3.16.2. Detector layout. The CRT system consist of an array of plastic scintillator counters placed on the top sides of the ALICE magnet. The available plastic scintillator material to build the array was previously used in the DELPHI experiment [470]. The current layout of the cosmic trigger on the top face of the ALICE magnet consists of 20 scintillator modules (covering the central part of the top face) arranged perpendicularly (i.e. with the counters running parallel to the shorter side of the face). Figure 3.41 shows a schematic representation of this design. There is a possibility to cover also two other faces, depending on physics needs.

Each module consists of scintillators with 1.88×0.2 m² effective area, arranged in a doublet configuration. Each doublet consists of two superimposed scintillator counters. See figure 3.42 for a schematic representation of a doublet and the set-up used to estimate the performance of these type of counters.

With this set-up we achieve a uniform efficiency higher than 90% along the whole module.

3.16.3. Front-end electronics and readout. The trigger system is based on four coincidence modules and one control module. Each coincidence module will analyse the signals from 30 scintillators. If one of the coincidence modules reports a hit, then the control module will generate a level zero trigger signal within a time window of 450 ns, shorter than the 500 ns upper limit, required by the CTP system. The coincidence and control modules will be in synchronization with the LHC clock. Block diagrams of the electronics trigger system are shown in figure 3.43. The readout system for this scintillator counter array is under development.

3.17. Trigger system

3.17.1. Design considerations. The ALICE trigger is designed to select events displaying a variety of different features at rates which can be scaled down to suit physics requirements and the restrictions imposed by the bandwidth of the DAQ and the HLT. The challenge for the




Figure 3.41. CRT scintillator array on top of the ALICE magnet.




Figure 3.42. Schematic view of a doublet of scintillator counters and the couple of paddles (hodoscope) used to measure the counter detection efficiency. The thickness of these plastic scintillators is 9 mm.




Figure 3.43. Preliminary design of the CRT logic electronics.

ALICE trigger is to make optimum use of the component detectors, which are busy for widely different periods following a valid trigger, and to perform trigger selections in a way which is optimized for several different running modes: ion (Pb–Pb and several lighter species) and pp, having a range of a factor 30 in counting rate.

The first response from the trigger system has to be fast ($\sim 1.2 \mu\text{s}$) to suit the detector requirements. The principal design requirement for the tracking detectors is to be able to cope with the large multiplicities in Pb–Pb collisions (interaction rate 8 kHz at nominal luminosity). In some cases (e.g. those detectors using the GASSIPLEX front-end chip) this requires a strobe to be sent at $1.2 \mu\text{s}$. This has led to the requirement that the ‘fast’ part of the trigger be split into two levels: a Level 0 (L0) signal which reaches detectors at $1.2 \mu\text{s}$, but which is too fast to receive all the trigger inputs, and a Level 1 (L1) signal sent at $6.5 \mu\text{s}$ which picks up all remaining fast inputs.

Another feature of the ALICE environment is that the high multiplicities make events containing more than one central collision unreconstructable. For this reason, *past–future protection* is an important part of the ALICE trigger. A final level of the trigger (Level 2, L2) waits for the end of the past–future protection interval ($88 \mu\text{s}$) to verify that the event can be taken. This interval can also be used for running trigger algorithms, though at present there are no candidate algorithms.

3.17.2. Trigger logic. The number of trigger inputs and classes required for ALICE (24 L0 inputs, 20 L1 inputs, 6 L2 inputs, 50 trigger classes) means that the trigger logic requires some justification, since a simple enumeration of all outcomes (look-up table approach) is not feasible. An investigation of the trigger conditions actually required for ALICE shows that these can be accommodated by logic involving three states of the inputs (*asserted, negated, not relevant*), coupled by ANDs. Negated inputs are rarely used, so they are available for six out of the fifty trigger classes, the rest allowing the two requirements *asserted* and *not relevant* only. We define a trigger class in terms of the logical condition demanded for the inputs, the set of detectors required for readout (a maximum of six combinations, called detector *clusters*, can be defined at any one time) the past–future protection requirements (which are in fact determined by the detector set) and the scaling factor to be applied.

Table 3.27 gives a list of trigger inputs currently envisaged for each trigger level, and table 3.28 the trigger classes for Pb–Pb interactions together with their corresponding trigger conditions.

Table 3.27. List of trigger inputs for Pb–Pb and pp interactions.

Number	L0 (Pb–Pb)	L0 (pp)	L1 (Pb–Pb)
1	V0 minimum bias	V0 minimum bias	TRD unlike e pair high p_T
2	V0 semi-central	V0 high multiplicity	TRD like e pair high p_T
3	V0 central	V0 beam gas	TRD jet low p_T
4	V0 beam gas	T0 right	TRD jet high p_T
5	T0 vertex	T0 left	TRD electron
6	PHOS MB	T0 vertex	TRD hadron low p_T
7	PHOS jet low p_T	PHOS MB	TRD hadron high p_T
8	PHOS jet high p_T	PHOS jet low p_T	ZDC 1
9	EMCAL MB	PHOS jet high p_T	ZDC 2
10	EMCAL jet high p_T	EMCAL MB	ZDC 3
11	EMCAL jet med p_T	EMCAL jet high p_T	ZDC special
12	EMCAL jet low p_T	EMCAL jet med p_T	Topological 1
13	Cosmic Telescope	EMCAL jet low p_T	Topological 2
14	DM like high p_T	Cosmic Telescope	
15	DM unlike high p_T	DM like high p_T	
16	DM like low p_T	DM unlike high p_T	
17	DM unlike low p_T	DM like low p_T	
18	DM single	DM unlike low p_T	
19	TRD pre-trigger	DM single	
20		TRD pre-trigger	
21			
22			
23			
24			

Note that the ‘TRD pre-trigger’ input fulfils a different function from the other trigger inputs. It is formed by copies of the T0 and V0 inputs sent in parallel to the TRD, where they serve to ‘wake-up’ the TRD electronics, which otherwise in a standby state. If the TRD pre-trigger is sent, it means that the TRD is ready to run triggers, or to read out; if not, it is not. Therefore, the confirmatory TRD pre-trigger signal (interaction detected AND TRD in ready state) is a prerequisite for any class for which TRD triggers or readout are required.

3.17.3. Trigger inputs and classes. The trigger inputs and classes were introduced in the previous section. Trigger inputs are pulses provided by the trigger detectors so as to be synchronized to the LHC clock cycle, as distributed by the TTC system.

Trigger inputs are sent using twisted pair cables, and are put in time by delays in the CTP input circuits. The trigger outputs are returned to the detectors using coaxial cables for L0 (in order to reduce the latency) and the RD12 TTC system (using optical fibre) for the other levels. Trigger data are transmitted with each successful L1 and L2 trigger, as described in the next section.

The expected trigger inputs are listed in table 3.27. Table 3.28 lists the trigger classes, indicating which trigger inputs are required for each class. A trigger *class* collects together a number of different features: the inputs required at each level, the cluster required and the downscaling factors. The cluster determines the past–future protection requirements, which are those of the detector in the cluster.

The purpose of the past–future protection circuit is to ensure that the events selected for readout are not spoiled by pile-up. As pile-up will be frequent, or, in some cases, inevitable at expected LHC luminosities, the pile-up protection must do more than simply reject events if

Table 3.28. List of trigger classes with trigger conditions.

Number	Description	Condition
1	MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} [ZDC ₁] _{L1}
2	SC	[T0.V0 _{SC} .TRD _{pre}] _{L0} [ZDC ₂] _{L1}
3	CE	[T0.V0 _{CE} .TRD _{pre}] _{L0} [ZDC ₃] _{L1}
4	DM _{unlike,highp_T} .TPC.MB	[T0.V0 _{MB} .DM _{unlike,highp_T} .TRD _{pre}] _{L0} [ZDC ₁] _{L1}
5	DM _{unlike,highp_T} .TPC.SC	[T0.V0 _{SC} .DM _{unlike,highp_T} .TRD _{pre}] _{L0} [ZDC ₂] _{L1}
6	DM _{unlike,highp_T} .No TPC.MB	[T0.V0 _{MB} .DM _{unlike,highp_T}] _{L0} [ZDC ₁] _{L1}
7	DM _{unlike,lowp_T} .No TPC.MB	[T0.V0 _{MB} .DM _{unlike,lowp_T}] _{L0} [ZDC ₁] _{L1}
8	DM _{unlike,lowp_T} .No TPC.SC	[T0.V0 _{SC} .DM _{unlike,lowp_T}] _{L0} [ZDC ₂] _{L1}
9	DM _{like,highp_T} .TPC.MB	[T0.V0 _{MB} .DM _{like,highp_T} .TRD _{pre}] _{L0} [ZDC ₁] _{L1}
10	DM _{like,highp_T} .TPC.SC	[T0.V0 _{SC} .DM _{like,highp_T} .TRD _{pre}] _{L0} [ZDC ₂] _{L1}
11	DM _{like,highp_T} .No TPC.MB	[T0.V0 _{MB} .DM _{like,highp_T}] _{L0} [ZDC ₁] _{L1}
12	DM _{like,lowp_T} .No TPC.MB	[T0.V0 _{MB} .DM _{like,lowp_T}] _{L0} [ZDC ₁] _{L1}
13	DM _{like,lowp_T} .No TPC.SC	[T0.V0 _{SC} .DM _{like,lowp_T}] _{L0} [ZDC ₂] _{L1}
14	DM _{single} .TRD _e .MB	[T0.V0 _{MB} .DM _{si} .TRD _{pre}] _{L0} .[TRD _e .ZDC ₁] _{L1}
15	DM _{single} .TRD _e .SC	[T0.V0 _{SC} .DM _{si} .TRD _{pre}] _{L0} .[TRD _e .ZDC ₂] _{L1}
16	TRD _e .MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} .[TRD _e .ZDC ₁] _{L1}
17	TRD _{lowp_T} .MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} .[TRD _{lowp_T} .ZDC ₁] _{L1}
18	TRD _{highp_T} .MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} .[TRD _{highp_T} .ZDC ₁] _{L1}
19	TRD _{unlike,highp_T} .MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} .[TRD _{unlike,highp_T} .ZDC ₁] _{L1}
20	TRD _{unlike,highp_T} .SC	[T0.V0 _{SC} .TRD _{pre}] _{L0} .[TRD _{unlike,highp_T} .ZDC ₂] _{L1}
21	TRD _{like,highp_T} .MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} .[TRD _{like,highp_T} .ZDC ₁] _{L1}
22	TRD _{like,highp_T} .SC	[T0.V0 _{SC} .TRD _{pre}] _{L0} .[TRD _{like,highp_T} .ZDC ₂] _{L1}
23	TRD _{jet,highp_T} .SC	[T0.V0 _{SC} .TRD _{pre}] _{L0} .[TRD _{jet,highp_T} .ZDC ₁] _{L1}
24	TRD _{jet,lowp_T} .MB	[T0.V0 _{MB} .TRD _{pre}] _{L0} .[TRD _{jet,lowp_T} .ZDC ₁] _{L1}
25	TRD _{jet,lowp_T} .SC	[T0.V0 _{SC} .TRD _{pre}] _{L0} .[TRD _{jet,lowp_T} .ZDC ₂] _{L1}
26	PHOS _{highp_T} .MB	[T0.V0 _{MB} .PHOS _{highp_T} .TRD _{pre}] _{L0} [ZDC ₁] _{L1}
27	PHOS _{lowp_T} .MB	[T0.V0 _{MB} .PHOS _{lowp_T} .TRD _{pre}] _{L0} [ZDC ₁] _{L1}
28	PHOS _{lowp_T} .SC	[T0.V0 _{SC} .PHOS _{lowp_T} .TRD _{pre}] _{L0} [ZDC ₂] _{L1}
29	PHOS Standalone	[T0.V0 _{MB} .PHOS _{MB}] _{L0} [ZDC ₁] _{L1}
30	EMCAL _{jet,highp_T} .MB	[T0.V0 _{MB} .EMCAL _{jet,highp_T}] _{L0} [ZDC ₁] _{L1}
31	EMCAL _{jet,medp_T} .MB	[T0.V0 _{MB} .EMCAL _{jet,medp_T}] _{L0} [ZDC ₁] _{L1}
32	EMCAL _{jet,lowp_T} .MB	[T0.V0 _{MB} .EMCAL _{jet,lowp_T}] _{L0} [ZDC ₁] _{L1}
33	EMCAL _{jet,lowp_T} .SC	[T0.V0 _{SC} .EMCAL _{jet,lowp_T}] _{L0} [ZDC ₂] _{L1}
34	ZDC _{diss}	[BX] _{L0} .[ZDC _{spe}] _{L1}
35	cosmic	[BX.cosmic.telescope] _{L0}
36	beam gas	[T0 _{beamgas}] _{L0}

any other event occurs within a specified time window. The circuit classifies events into two groups, peripheral and semi-central, and specifies separate limits for each. The time windows are set to be equal to the sensitive windows for the detectors, i.e. the (nominal) maximum time over which two signals from two different events could be registered in a single firing of the detector. Thus, for instance, the sensitive window for the TPC is $\pm 88 \mu\text{s}$, so for clusters in which the TPC is included a reasonable past–future protection specification might be to demand no more than four additional peripheral events and no additional semi-central events to take place in an interval of twice $88 \mu\text{s}$ centred on the event under consideration.

A slightly different approach is required for proton–proton interactions, where the increased luminosity makes pile-up a certainty, but where the multiplicities are much lower than in ion–ion interactions, and therefore greater degrees of pile-up are tolerable. Here, what is important is to ensure that all detectors in a cluster read out in a satisfactory manner. For

example, pile-up in the ITS is more serious for reconstruction than pile-up in the TPC. For this reason, in this case the past–future protection does not select on multiplicity, but offers checks in two different time windows (e.g. $\pm 10 \mu\text{s}$ for the ITS and $\pm 88 \mu\text{s}$ for the TPC) so as to ensure that pile-up is not excessive in either.

3.17.4. Trigger data. The Central Trigger Processor (CTP) must be able to process triggers for all trigger clusters concurrently, and therefore has a tight time budget for collecting and distributing data. For this reason these data are kept to a minimum, so, for example, all information concerning how a trigger signal was generated must be obtained from the detector which produces it. Several types of data recorded concerning the operation of the trigger, as listed below.

Event data. For each accepted event, information is sent to each detector specifying the event identifier (orbit number and bunch crossing number), trigger type (physics trigger, software trigger, calibration trigger), detector set to be read out (given as a cluster list for physics data) and the list of active trigger classes. These data are sent in the L2a message, which goes to each TTCrx on each detector.

Interaction record. All interactions that occur in a given orbit are recorded (classified into centrality categories: peripheral and semi-central). The principal purpose of this record is to aid pile-up correction. Out-of-time events recorded in the TPC will have different apparent origins from triggered events owing to their different drift times. Their apparent origin can be computed using the interaction record, thus aiding their removal from the event data. The data are sent as special records in the data stream.

Scalers. A large number of scalers are recorded. In particular, all inputs are counted, and for each trigger class the number of events passing each stage of the trigger (L0, L1, L2), after physics selections, if any, and again after the corresponding BUSY and past–future protection constraints have been applied. In addition, for each detector cluster the number of interactions not vetoed by its BUSY or past–future protection for the cluster are counted, allowing evaluation of cross sections by normalization to the total cross section.

Snapshots. For completeness, we note that it will also be possible to run ‘snapshots’ in which all steps in the trigger (input and output patterns, busy status, etc) are recorded on a bunch crossing by bunch crossing basis for a period of about 30 ms. The purpose of this facility is principally for diagnostic tests of the CTP itself.

In summary, the trigger information allows the performance of the trigger, the event rates and the details of the trigger selection for each event to be followed.

3.17.5. Event rates and rare events. A detailed simulation of the trigger and data acquisition system has been developed. Some results are given in the data acquisition section. The question of L2 rates is governed by the downscaling factors chosen, as the physics interest has to be matched to the bandwidth for writing data to permanent storage. A problem which has been identified is that where rare trigger classes compete with more common (frequently occurring) ones, the common trigger classes tend to take all available resources, thus suppressing the rare classes. In order to solve this problem restrictions must be made on the common classes in order to share the resources better, for example by limiting the time fraction available for recording common classes. Various strategies to implement this have been studied. The strategy adopted is to block the most high rate classes for some of the time, in order to ensure that there is always enough buffer space to allow rarer, less high rate classes to be taken. The details, given

in the DAQ section, show that the most sensitive part of the system comes in the buffering in the LDCs, where adjustments can be made on a timescale of about a second. However, in order to reach this state the front end buffering in the different detectors has been studied and increased where necessary in order to avoid these being the critical parameters, since reaction times for a dynamic approach to front-end buffering would be prohibitively short.

3.18. Data Acquisition (DAQ) System

3.18.1. Design considerations. As presented in the previous sections, ALICE will study several physics topics using different beam conditions. A large number of trigger classes will be used to select and characterize the events. These trigger classes belong to two broad categories depending on whether they are frequent or rare.

Several triggers (central, semi-central and minimum bias) are so frequent that the limiting factor is the performance of the data acquisition system. These triggers will use a very large fraction of the total data acquisition bandwidth. On the other hand, rare triggers such as dimuon or dielectron events, use less bandwidth and are limited by the detector livetime and the luminosity. The experiment will use the data-taking periods in the most efficient way by acquiring data for several observables concurrently following different scenarios. The task of the ALICE Trigger, Data Acquisition (DAQ) and High-Level Trigger (HLT) systems is therefore to select interesting physics events, to provide an efficient access to these events for the execution of high-level trigger algorithms and finally to archive the data to permanent data storage for later analysis.

The trigger and DAQ systems have been designed to give different observables a fair share of the trigger and DAQ resources with respect to DAQ bandwidth for frequent triggers and detector lifetime for rare triggers. The trigger and DAQ systems must balance the capacity to record central collisions which generate large events with the ability to acquire the largest possible fraction of rare events.

In the ALICE TP [382, 383], it was estimated that a bandwidth of 1.25 GB s^{-1} to mass storage would provide adequate physics statistics. While the current estimate of event sizes, trigger rates and, in particular, the number of physics observables (and therefore trigger classes) has increased considerably, it is still possible to satisfy the physics requirements with the same total bandwidth by using a combination of increased trigger selectivity, data compression and partial readout [471]. This bandwidth is consistent with constraints imposed by technology, cost and storage capacity, including the cost of the media needed to store the data and the computing needed to reconstruct and analyse these data. The Tier-0 of the LHC Computing Grid project that will be assembled at CERN will provide a bandwidth to mass storage higher than the needs of ALICE for data acquisition. However, efforts will be continued to reduce the bandwidth without jeopardizing the physics performance.

3.18.1.1. Data volume with Pb–Pb beams

Event size. The event size resulting from heavy-ion collisions is directly proportional to the multiplicity. A maximum of 8000 tracks per pseudo-rapidity unit at mid-rapidity for central Pb–Pb collisions has been considered here. The event size is a function of the centrality of the interaction. The size of central events is estimated for the most central collisions. The size of semi-central has been assumed to be 50% and the size of minimum-bias events has been estimated to be 25% of the central event size, on average.

Based on the present physics simulations, the central event size in Pb–Pb interactions is 86.5 MB before data compression. The size for each sub-system is shown in the second column of table 3.29.

Table 3.29. Portion of the event size from each ALICE sub-detector for pp minimum-bias and Pb–Pb central events.

Detector	pp (kB)	Pb–Pb (MB)
ITS Pixel		0.140
ITS Drift	1.8	1.500
ITS Strips		0.160
TPC	2450.0	75.900
TRD	11.1	8.000
TOF		0.180
PHOS		0.020
HMPID		0.120
MUON		0.150
PMD		0.120
Trigger		0.120
Total	2500	86.500

Studies of simulated ALICE data show that a compression factor of 2 can be obtained without loss of physics information [472, 473]. Several additional steps of data reduction are envisaged for the Time Projection Chamber including cluster finding and data compression based on local tracklets or full reconstruction. These algorithms should be applied on real data only when sufficient experience has been gained. Background suppression has been estimated to reduce the muon data volume by a factor of 4. The data throughput can also be reduced by performing a selective readout (regions of interest) or by applying a High-Level Trigger (HLT). Partial readout is described as the readout of 2–4 sectors of the TPC and the TRD. We have assumed a value of three sectors. For dielectron triggers the HLT is estimated to reduce the rate by a factor of 10. These compression and reduction factors have been used for the estimate of event size in the different scenarios described hereafter.

Event rate. An estimate of the total bandwidth needed in the DAQ requires the evaluation of the data throughput for each type of trigger. The number of events needed to accumulate enough statistics in a one-year period of data taking is up to 10^7 events for minimum bias hadronic physics, a few 10^7 events for hadronic charm decays, jet, and electron physics, and at least 10^9 events for dimuon physics. The Pb–Pb runs will last only a few weeks per year, corresponding to an effective time of 10^6 s. The needed rates are then of the order of a few Hz for hadronic physics, a few tens of Hz for minimum bias, hadronic charm, jet, and electrons, and a few hundred Hz for dimuons for the L2 triggers.

Data-taking scenarios in Pb–Pb. We present some typical data-taking scenarios (table 3.30) used to evaluate the global performance requirements for the DAQ system. We assume that the data are produced under identical data-taking conditions. While reasonable values have been chosen as far as possible for the key parameters, the HLT functionality is included only very schematically in these scenarios; the reader is referred to the appropriate section in this document for a more realistic description of the HLT capabilities.

Scenario 1. In the first scenario, 2×10^7 minimum-bias, 2×10^7 central, and 6.5×10^8 dimuon triggers are collected. With the minimum-bias and central triggers, all the ALICE detectors are read out. With the dimuon trigger, only the muon arm, the Inner Tracking System pixel, the PHOTon Spectrometer, the Photon Multiplicity Detector, the Forward Multiplicity

Table 3.30. Data-taking scenarios with Pb–Pb beams.

	Scenario 1		Scenario 2		Scenario 3		Scenario 4	
	Rates (Hz)		Rates (Hz)		Rates (Hz)		Rates (Hz)	
	Maximum	DAQ	Level 2	DAQ	Level 2	DAQ	Level 2	DAQ
Central	10^3	20	10	10	20	20	20	20
Minimum-bias	10^4	20	10	10	20	20	20	20
Dielectron			100	100	200	20	200	20
Dimuon	1000	650	1600	1600	1600	1600	1600	1600
Total throughput (MB s^{-1})		1250	1400	1400	700			

Detector, and the trigger detectors are read out. This scenario corresponds to data acquisition at the maximum possible rate without further selection or modification of the data, resulting in a total throughput of 1.2 GB s^{-1} to mass storage. Data compression is assumed to be performed by the detector readout electronics before transmission to the DAQ. If the data compression were applied at a later stage, the data throughput would increase by a factor of 2.

Scenario 2. Scenario 2 is similar to scenario 1, but it includes the dielectron trigger and it distinguishes two categories of muon triggers. The dielectron triggers can be acquired at a rate of 20 Hz with the readout of all detectors. For the physics that will be studied with muon events, the full ITS readout would allow better event characterization and correlation studies. However, the event size, $1.1 \text{ MB event}^{-1}$ after data compression, for muon and ITS readout would take up an unacceptably large fraction of the available bandwidth, 0.7 GB s^{-1} . Therefore, two different categories of dimuon triggers are defined. The low- p_t triggers involve the readout of the muon arm and the ITS pixel. The high- p_t triggers also include the readout of the other two ITS layers: the drift and the strips. This addition results in a total throughput of 2.5 GB s^{-1} , higher than the agreed maximum of 1.25 GB s^{-1} . Reducing the number of minimum-bias and central events foreseen in scenario 1 by a factor of two is necessary to reduce the global throughput to 1.4 GB s^{-1} .

Scenario 3. In scenario 3, the HLT system is used to trigger dielectron events online. The raw rate for events to be processed by the HLT can be increased by a factor of 10 given a total of 2×10^8 events, limited by the L1 TPC gating rate. Thanks to this reduction of the electron triggers, the original rates of minimum-bias and central events foreseen in scenario 1 (2×10^7) can be restored almost within the total bandwidth budget. The total throughput is of the order of 1.4 GB s^{-1} . In this scenario, the HLT algorithms require a copy of all the raw data.

Scenario 4. Scenario 4 is similar to scenario 3 but includes more advanced HLT capabilities and also includes the introduction of several new requirements: the rejection of electron events after inspection of the full event and background rejection for the muon data are needed, data compression based on an online tracking is introduced for the minimum-bias and central TPC data. The archived data are not the original raw data but data compressed following this tracking model. The event-building throughput, 0.7 GB s^{-1} , assumes this data compression can be performed on sub-events. Were this not the case, the event-building would increase the rate to 2.5 GB s^{-1} . The benefit of the HLT and these two online filters is to reduce the total throughput to 0.7 GB s^{-1} . It would also be possible to increase the number of useful minimum-bias and central events or to keep some reserve for future extensions.

3.18.1.2. Data volume with pp beams

Event size. The estimated TPC event size for one pp collision is of the order of 60 kB. On average, 40 events will pile up during the drift time of the TPC [474] which limits the luminosity to $3 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$.

The tracks of the events corresponding to interactions occurring before and after the trigger will only be partially included in the TPC data volume. The total data volume therefore corresponds to 20 complete events. The total event size is estimated to be of the order of 2.5 MB, as indicated in the first column of table 3.29. This event size is estimated assuming a coding scheme for the TPC data adapted to low occupancies and without data compression. Online tracking by the HLT will select tracks belonging to the interesting interaction. This online pile-up suppression could reduce the event size by about a factor of 20 (see HLT section).

Event rate. The statistics required for proton physics is of the order of 10^9 events. The event rate will vary with the available luminosity. The requirement is therefore to be able to collect the data at the maximum possible rate for the ALICE detectors: $1000 \text{ events s}^{-1}$ and at an average of $100 \text{ events s}^{-1}$.

3.18.2. Data acquisition system

System architecture. A short description of the architecture of the data acquisition system is presented here. An exhaustive report is available in [397].

The architecture of the data acquisition is shown in figure 3.44. The detectors receive the trigger signals and the associated information from the Central Trigger Processor (CTP), through a dedicated Local Trigger Unit (LTU) interfaced to a Timing, Trigger and Control (TTC) system. The readout electronics of the detectors is interfaced to the ALICE-standard Detector Data Links (DDL). The data produced by the detectors (event fragments) are injected on the DDLs using the same standard protocol. The fact that all the detectors use the DDL is one of the major architectural features of the ALICE DAQ.

At the receiving side of the DDLs there are PCI-based electronic modules, called 'DAQ Readout Receiver Card' (D-RORC). The D-RORCs are hosted by the front-end machines (commodity PCs), called Local Data Concentrators (LDCs). Each LDC can handle one or more D-RORCs. The D-RORCs perform concurrent and autonomous DMA transfers into the LDCs' memory, with minimal software intervention. The event fragments originated by the various D-RORCs are logically assembled into sub-events in the LDCs.

The CTP receives a busy signal from each detector. This signal can be generated either in the detector electronics or from all the D-RORCs of a detector. The CTP also receives a signal from the DAQ enabling or disabling the most common triggers. It is used to increase the acceptance of rare triggers by reducing the detector dead-time. This signal is generated from the buffer occupancy in all the LDCs.

The role of the LDCs is to ship the sub-events to a farm of machines (also commodity PCs) called Global Data Collectors (GDCs), where the whole events are built (from all the sub-events pertaining to the same trigger). Another major architectural feature of the ALICE DAQ is the event builder, which is based upon an event-building network. The sub-event distribution is driven by the LDCs, which decide the destination of each sub-event. This decision is taken by each LDC independently from the others (no communication between the LDCs is necessary), but it is synchronized among them by a data-driven algorithm, designed to share fairly the load on the GDCs. The Event-Destination Manager (EDM) broadcasts information about the availability of the GDCs to all LDCs. The event-building network does

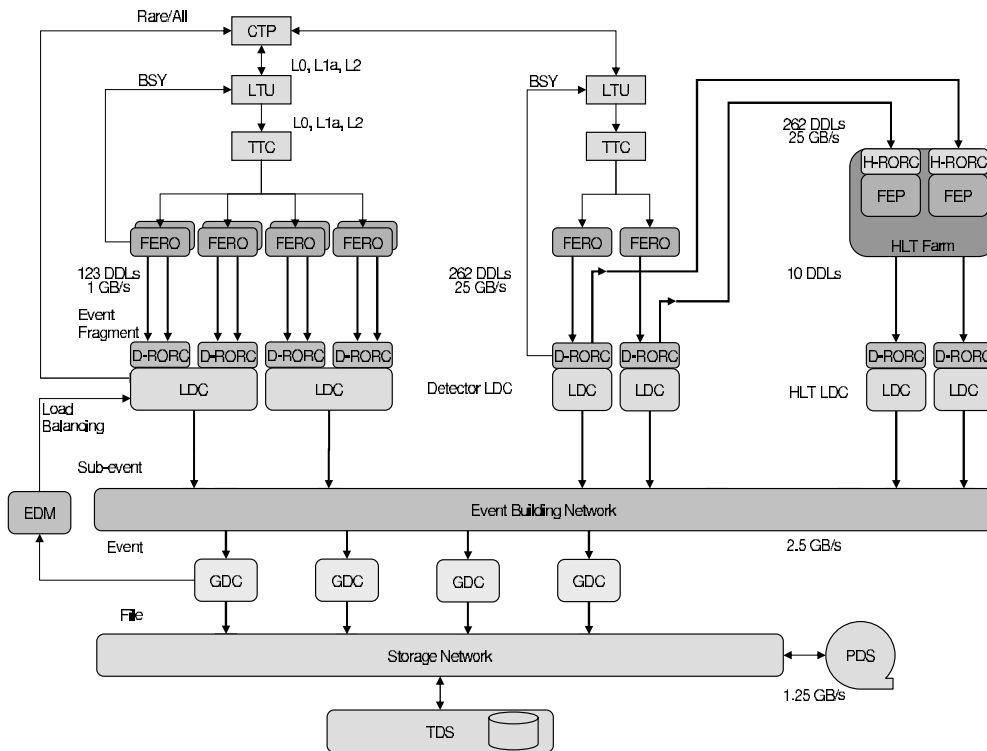


Figure 3.44. The overall architecture of the ALICE DAQ and the interface to the HLT system.

not take part in the decision; it is a standard communication network (commodity equipment) supporting the well-established TCP/IP protocol. The role of the GDCs is to collect the sub-events and assemble them into whole events. The GDCs also feed the recording system with the events that eventually end up in Permanent Data Storage (PDS).

The HLT system receives a copy of all the raw data; generated data and decisions are transferred to dedicated LDCs.

Data transfer. The Detector Data Link (DDL) is the common hardware and protocol interface between the front-end electronics and the DAQ system. The transmission medium is a pair of optical fibres linking sites several hundred metres apart. The DDL is able to transmit data in both directions with a sustained bandwidth of 200 MB s^{-1} . It can be used to send commands and bulky information (typically pedestals or other calibration constants) to the detectors. The DDL hardware and software is described in several ALICE notes [475, 476].

The D-RORC is the readout board of the DDL link. It interfaces the DDL to the PCI bus of the host LDC computer. In addition, the D-RORC grants to the LDC full control of all the features of the DDL. The D-RORC technology is based on a FPGA. The card is able to transfer the event fragments from the link directly into the PC memory at 200 MB s^{-1} , without on-board buffering; this bandwidth fulfils (and far exceeds) the original requirement of the ALICE data-acquisition system. The DMA transfer is carried out by the firmware in co-operation with the readout software running on the LDC. In this closely coupled operation, the role of the software is limited to the bookkeeping of the data-block descriptors. This approach allows sustained autonomous DMA with little CPU assistance and minimal software overhead. The

data formatting introduced by the DDL allows the D-RORC to separate the event fragments in memory, without interrupting the host CPU. Performance measurements and long-term tests show that the system fully exploits the PCI bandwidth and that the transfer rate is largely independent of the block size.

The D-RORC includes two DDL interfaces. For the detectors whose data are analysed by the HLT system, one of the DDL interfaces is used to transfer a copy of the raw data to the HLT computers.

Two RORC prototypes have been built based on the VME and PCI standards [477, 478]. These prototypes have been used to develop the DAQ system to integrate the DAQ with the detector electronics.

Event building. The sub-events prepared by the LDCs are transferred to one GDC where the full event can be assembled. The event building is managed by the Event Building and Distribution System (EBDS). The EBDS distributed protocol runs on all the machines, LDCs and GDCs, participating as data sources or destinations. The goals of the EBDS are to synchronize all the LDCs concerning the GDC used as destination and to balance the loads on the different GDCs. The aim is to keep up with the data-flow while keeping the EBDS protocol overhead as low as possible.

We have used, up to now, the de-facto standard TCP/IP as the transport mechanism for the EBDS protocol and for the data transfer. This protocol is the standard to which the networking industry has converged and it runs on a wide range of hardware layers. Ethernet is the dominant, ubiquitous Local Area Network (LAN) technology. We therefore use this technology to build our prototypes and it has become the baseline choice. However, as long as we can rely on a standard protocol such as TCP/IP, the final choice of the hardware layer used in the event-building network could be modified if a better technology would become available.

DATE: the DAQ software framework. The development of such a large DAQ system by several teams requires clearly defined interfaces between the components of the system. The integration of these components into the system should be started early enough. Furthermore, the long development and production periods of the experiment is inevitably longer than a single generation of some of the components which must be updated before and during the data-taking period. This evolution has to be planned from the start.

It was thus necessary to develop a software framework able to surround the development work of the final DAQ system as well as to support the present detector activities of the experiment: the Data Acquisition and Test Environment (DATE) [479]. One production version of the system is released at time and constitutes the reference framework into which the prototypes are integrated. It is also used for the functional and performance tests of the DAQ system.

The DATE framework is a distributed, process-oriented system. It is designed to run on Unix platforms connected by an IP-capable network and sharing a common file system such as NFS. It uses the standard Unix system tools available for process synchronization and data transmission. The system characterizes different functions:

- The LDC collects event fragments and reassembles them into sub-events. The LDC is also capable of local data recording, if used in stand-alone mode, and online sub-event monitoring.
- The GDC puts together all the sub-events pertaining to the same physics event, builds the full events and sends them to the mass storage system. The GDC is capable of online event monitoring.

- The DATE controls and synchronizes the processes running in the LDCs and the GDCs. It can run on an LDC, a GDC or another computer.
- The monitoring programs receive data from the LDC or GDC streams. They can be executed on any LDC, GDC or any other machine accessible via the network.

The DATE software is being used for large-scale tests of the DAQ and Offline systems called the ‘ALICE Data Challenges’ [480].

3.18.3. System flexibility and scalability. The physics requirements for the DAQ system are still in evolution. These requirements will evolve with improved physics simulations and with new ideas of high-level triggers. There are also important uncertainties in some of the key parameters needed to evaluate the physics requirements. First, there is still an uncertainty in the statistics (and therefore the event rate) needed for measurements such as hadronic charm, which might have important consequences on the necessary DAQ bandwidth. There is also some uncertainty in the multiplicity in heavy-ion collisions. The event size of central events grows linearly with multiplicity. This uncertainty will only be removed after the first ion run. The consequences of these uncertainties on the design of the DAQ system have therefore been reviewed. We have evaluated below the hard limits of the current DAQ system design and its potential scalability to address requirements in evolution.

Currently, the only hard limit present in the design of the DAQ system is the number of DDLs used to read out the TPC and the TRD. The number of links has been fixed so that these two detectors can be completely read out at a maximum rate of 200 Hz. The next stages in the data flow chain (computers and event-building network) are completely scalable and their number can be adapted to the needs. No architectural change would be needed in the DAQ system to scale to a total throughput varying by a factor 4 or 5. This flexibility also authorizes a staged deployment of the system and the implementation of new scenarios on short notice.

3.18.4. Event rates and rare events

Introduction. The full data flow behaviour is modelled, using realistic interaction rates and event size distributions of different trigger classes. The model consists of the Central Trigger Processor (CTP), the detectors, the DAQ, the High-Level Trigger (HLT) and the Permanent Data Storage (PDS) [481–483]. This model is based on the public-domain tool Ptolemy [484].

Each event is assigned a track multiplicity by extrapolating to LHC energies the measured SPS and RHIC multiplicities. Events are characterized by different trigger classes, peripheral (PE), semi-central (SC) or central (CE), according to their multiplicity. In addition, an event can also be of the electron (EL) or muon (MU) trigger class based on an assigned probability.

Extensive simulations of ALICE have been performed using this model and have shown the ability of the whole experiment to function at the desired rate and bandwidth performances. It has also shown potential problems with the way rare triggers (electrons and muons) are handled. The problem arises inevitably whenever the data volume sent over the DDLs (as given by the trigger mix and event size) exceeds the data volume that can be processed by the DAQ–HLT system. In this case, the buffers in the intermediate stages of the DAQ (GDCs, LDCs, or RORCs) get full and by the action of the back-pressure this is propagated up to the detectors that become busy almost all the time. As soon as the busy is released, the frequent triggers are obviously the most likely to fire and to put the detector into busy state again.




Figure 3.45. L2 rates as a function of time. The horizontal lines at the bottom indicate the periods when the PE, SC and CE events are blocked.

The problem is to keep the detector livetime as high as possible in order to catch as many rare triggers as possible and to keep the back-end part of the DAQ system (event building and mass storage) as busy as possible, i.e. as close as possible to the maximum installed bandwidth.

The solution is to include a rate control mechanism such that no rare events (electron and muon) are lost due to back-pressure, while at the same time as many frequent event (peripheral, semi-central or central) are accepted as possible with the installed bandwidth.

This rate control is based on a back-pressure from the DAQ that detects the amount of data filled in each computer, and informs the trigger whenever this level reaches a high fraction of the buffer size. When this happens, the frequent events, as the main bandwidth contributors and events of lower priority, would then be blocked by the trigger.

Figure 3.45 gives the L2 rates as a function of time when the LDC high level is set to 0.3, the low level is at 0.2, and the downscaling factor for PE, SC and CE events is set at 0.2. The periodic behaviour of the PE, MB, and CE events is apparent, and so is the constant rate of the EL and MU events. Much larger rates than previously are achieved, with 22% of EL events (it was 2% before) and 78% of MU events (it was 22% before) accepted. This is a large performance improvement and is close to the maximal rates possible, after past-future (P/F) rejection. The PE, MB, and CE events fill up the remaining bandwidth to the PDS.

Figure 3.46 shows the LDC occupancy as a function of time, from which it is clear that the LDCs are never full, and hence there is no loss due to LDC back-pressure. Figure 3.47 shows the fraction of time a detector is busy. The events lost due to busy detectors are only caused by detector readout dead times and by transfer limitations through the detector buffer chain. These are hardware properties of the detector and cannot be influenced by the DAQ or CTP.

Robustness of the rate control mechanisms. The system performance was tested for additional LDC high/low level combinations (with the high level going up to 0.4) and for downscaling factors ranging from 0 to 1, with stable L2 rates (no losses due to LDC back-pressure) achieved.




Figure 3.46. Fraction of the LDC buffer filled as a function of time.




Figure 3.47. Fraction of time a detector is busy as a function of time.

Further, simulations were made where the data distribution across the TPC is not uniform, but instead systematically changes by up to 50%, with no problems occurring. Varying interaction rates and physics conditions were investigated, with the results given in table 3.31. Again, no LDC back-pressure losses were seen. The much lower fractional acceptances of EL and MU events for the higher interaction rates are entirely caused by larger rejection factors due to P/F protection (see figure 3.48).

Table 3.31. L2 rates of EL and MUON events for different interaction rates and physics conditions with LDC high/low levels at 0.3/0.2. The percentages in the brackets give the fraction of events accepted after P/F criteria are satisfied. Unless otherwise stated, the EL rate is 800 Hz, the MUON rate is 700 Hz, a partial EL readout of the TPC and TRD is used, and the downscaling factor is set at 0.2.

	4000 Hz 400 Hz EL 350 Hz MUON	8000 Hz	8000 Hz downscale 0.0	8000 Hz downscale 0.1
EL	175 Hz (88%)	170 Hz (89%)	180 Hz (94%)	175 Hz (91%)
MUON	302 Hz (91%)	550 Hz (88%)	570 Hz (90%)	560 Hz (89%)
	8000 Hz $dN/d\eta = 4000$	8000 Hz Full EL readout	8000 Hz 1300 Hz MUON	8000 Hz ITS drift 0.8 ms
EL	183 Hz (95%)	175 Hz (92%)	170 Hz (89%)	170 Hz (89%)
MUON	570 Hz (90%)	570 Hz (90%)	980 Hz (84%)	550 Hz (88%)

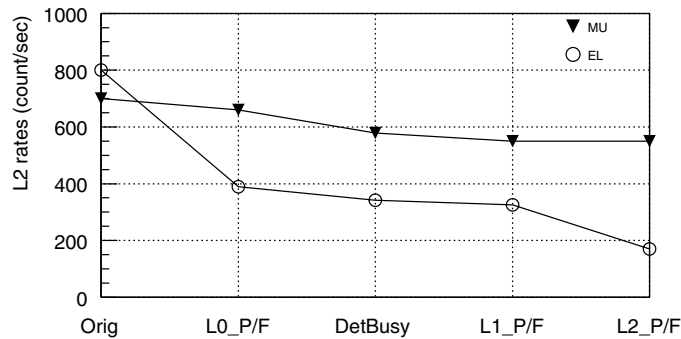


Figure 3.48. Source of event rejection with strict past–future protection.

The robustness of the LDC feedback mechanism indicates that it does not need fine tuning for different interaction rates, physics conditions, and detector performances.

P/F condition relaxation. PE events comprise 65% of the total, even though they contain only up to 20% of the maximum number of tracks. If the offline analysis could successfully process events containing up to one PE P/F violating interaction, then a much larger fraction of high priority events could be accepted. Table 3.32 gives the EL and MUON L2 rates, and shows that indeed 38% of EL events are this way accepted (as opposed to 21% with no P/F violation). However, the amount of MU events is mostly unchanged, since even though fewer of them are lost due to P/F consideration, the larger number of EL events accepted makes the shared detectors between these two trigger classes busy for a longer time. Also, for the same reason the greater rates of PE, SC, and CE events in the time periods when they are taken slightly reduces the MU events accepted. Downscaling by a further amount, for example by 0.1, could increase the final MU rates, but only by a negligible amount.

In summary, an effective and robust method to maximize the L2 rates of rare and interesting events can be achieved by controlling the rates of peripheral, semi-peripheral and central events with feedback from the LDC occupancy and with downscaling. No difficulty in either messaging the information from the LDCs, or the stability of the system with different interaction rates or envisioned physics is anticipated.

Table 3.32. L2 rates when one peripheral event may violate past–future protection.

	Initial rate	L2 rate
EL rate	800 Hz	300 Hz
MUON rate	700 Hz	575 Hz

3.19. High-Level Trigger (HLT)

3.19.1. Design considerations. The measurements in the ALICE experiment cover basic observables such as particle production and correlations as well as probes such as open charm and beauty, quarkonia production, direct photons and jets. Most of the probes are rare signals requiring the exploitation of all the available luminosity, which makes it necessary to sift through all interactions at a rate of about 8 kHz for Pb–Pb. The first level trigger systems of the muon spectrometer and the TRD select muon and electron candidates quite efficiently, but the remaining data rate is still far higher than what can be transferred to the permanent storage system. The High-Level Trigger (HLT) system [485] rejects fake events by sharpening the momentum cut in the muon spectrometer [486], by improving the particle identification, momentum resolution and impact parameter determination for TRD electron candidates in the TPC or by filtering out low momentum tracks in the TPC. ALICE would not be able to acquire sufficient statistics for many rare probes without the HLT system.

Online processing is necessary in order to select interesting events and subevents or to compress the data efficiently by modelling techniques. Detector information is fed into the system via the DDL at rates of about 20 GB s^{-1} , the output into the event builder should be less than about 1 GB s^{-1} [393]. Processing the raw data at a bandwidth of $10\text{--}20 \text{ GB s}^{-1}$ or even higher (the peak rate into the front-end electronics is 6 TB s^{-1}) requires a massive parallel computing system equipped with FPGA co-processors for hardware acceleration.

HLT functionality. The physics functionality of the HLT system can be summarized by the following three categories:

- *Trigger:* Accept or reject events based on detailed online analysis of physics observables, e.g. verify dielectron candidates after TRD L1 acceptance, sharpen dimuon transverse momentum cut or identify jets.
- *Select:* Select relevant parts of the event or regions of interest, e.g. remove pile-up in pp, readout jet regions or filter out low momentum tracks.
- *Compress:* Reduce the amount of data required to encode the event information of the triggered and selected events as far as possible without losing physics information.

The triggered event or selected regions of interest are finally compressed and then submitted to the event builders for permanent storage. In case of HLT reject the event is discarded. An event is required to be triggered by L0 and accepted by both L1 and L2 in order to reach the HLT. The main difference between HLT and the lower trigger layers is the fact that there is no detector requirement imposing a maximum HLT latency, which is basically defined by trigger rate and available buffer space. Further the HLT is the trigger layer allowing to combine the information of all ALICE subdetectors at one place.

A key component of the proposed system is the ability to process the detector raw data performing track pattern recognition in real-time. The HLT system is designed to utilize the information from all ALICE subdetectors including the TPC and the fast detectors, such as the TRD and muon spectrometer. The system will be flexible enough to be expanded to include

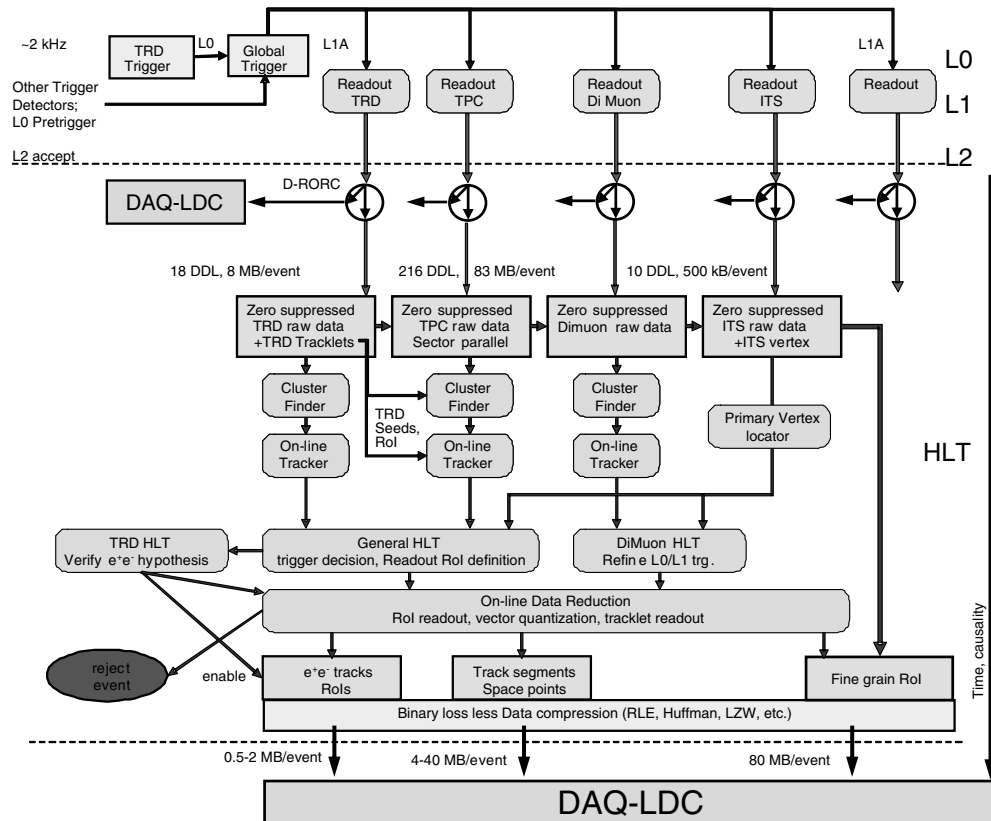


Figure 3.49. Join up HLT functionality and integration in the ALICE trigger hierarchy.

other tracking devices. It allows forwarding zero-suppressed events into the DAQ data stream on an event-by-event and detector-by-detector basis for diagnostic purposes. Furthermore it implements all necessary means to monitor both accepted and rejected events online.

Figure 3.49 outlines the overall functionality of the HLT system within the ALICE trigger hierarchy. Basically all subdetectors submit their data into the HLT upon a L2 accept. In case of a L2 reject the event is terminated within the detectors front-end. Each event can take different paths within the HLT system. Some processing will be contingent upon the outcome of other analyses, such as the verification of TRD trigger candidates. The HLT involves an online analysis of all participating detectors and allows to implement physics triggers based upon derived physics observables. After triggering an event and possibly selecting its regions of interest each subevent is subjected to a detector optimized data compression engine in order to minimize the data recording rate, allowing for a higher event rate at any given taping speed.

Fast pattern recognition. All HLT operations, including (sub)event reconstruction and data modelling, require a fast online pattern recognition [487]. In general, there are different approaches in track reconstruction (see figure 3.50):

1. Sequential feature extraction: A cluster finder searches for local maxima in the raw ADC-data. If an isolated cluster is found, the centroid position in the time and pad directions is calculated. In the case of overlapping clusters simple unfolding (splitting) procedures

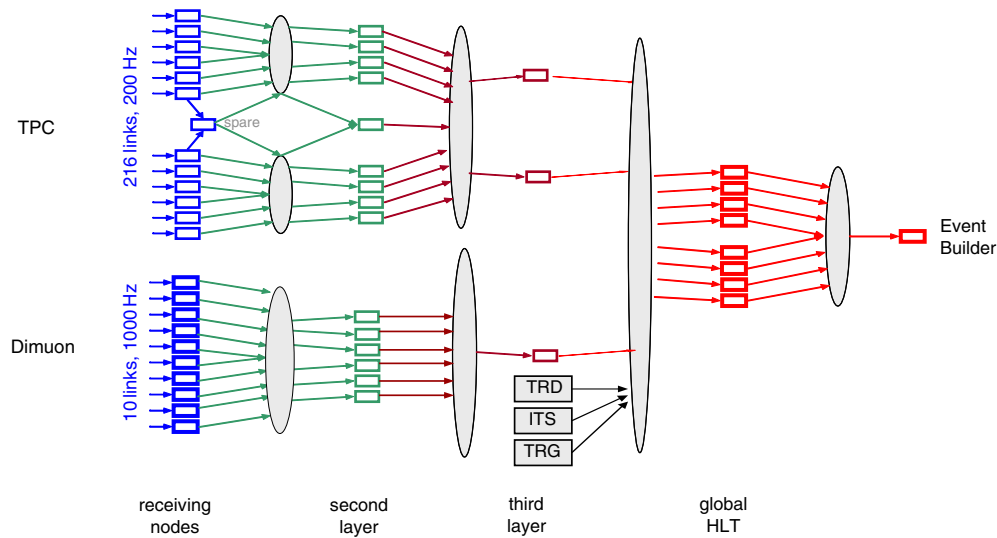


Figure 3.50. HLT operations.

separate charge distributions. Due to the missing track model, there are no cluster models for the ‘two-or-more’ charge distributions, therefore the centroid determination is biased. The coordinates together with the position of the pad row define a space point. Corrections for distortions and calibrations are applied to space points. The list of space points is given to the track finder. The track finder has no access to the raw data. This approach is suited for low occupancy (less than 20%), e.g. for the ALICE pp program and for low multiplicity Pb–Pb events of less than $dN/d\eta = 4000$.

2. Iterative feature extraction: Track segment finding is performed on the raw data level. Based on the list of found track segments, clusters are reconstructed along the track. Since the parameters of the track creating the cluster are known, the cluster shape can be modelled. Based on the additional knowledge of tracks passing nearby and therefore possibly contributing charge to the cluster, overlapping clusters can be de-convoluted. This approach will be used for the Pb–Pb program (if the multiplicity in the central detectors turns out to be higher than $dN/d\eta = 4000$).
3. Simultaneous pattern recognition and track fitting: The tasks for track reconstruction are track finding, track fitting and track matching between detector modules. The Kalman filtering method provides the means to do pattern recognition and track fitting simultaneously, and also in a natural way to prepare for the track matching. The advantages of the Kalman filtering approach make it the method of choice for the offline (final) reconstruction of the tracks. However, there are two main disadvantages of the Kalman filter, which play a major role in a high occupancy environment such as that of the ALICE TPC. The first disadvantage is that one has to reconstruct clusters before applying the Kalman filter procedure. Since cluster finding cannot be done without knowledge of the track parameters, a simple pattern recognition step has to precede the Kalman filter. The other problem with the Kalman filter tracking is that it relies essentially on good seeds to start a stable filtering procedure. Again, a prior pattern recognition step would deliver good seeds. It is planned to apply Kalman filtering for selected tracks, for merging TPC-sector tracks and for extrapolating tracks from the TPC into ITS and TRD.

Fast cluster finder. The input for the cluster finder is a list of above-threshold time-bin sequences for each pad. The algorithm basically matches these sequences by looking at neighbouring pads. In order to reduce memory access it loops over the list of sequences for each pad only once. The cluster finder does everything in one loop: cluster finding, centroid calculation and deconvolution. It loops over sequences of one pad, calculates the weighted mean for each sequence and tries to match the results with the weighted mean of sequences of the last pad (only two pads are stored internally). The centres of each sequence are used to calculate ‘on-the-fly’ the weighted mean in both pad and time direction of the cluster. If deconvolution is switched on, the algorithm searches for peaks in each sequence in the time direction, and for local minima of the sequence charge (sum over all ADC values of one sequence) in the pad direction and simply applies the same method as described for each peak, whereas the sequences are cut at the minimum between two peaks.

Hough transformation. The Hough transformation is a standard tool in image analysis that allows recognition of global patterns in an image space by recognition of local patterns (ideally a point) in a transformed parameter space. The basic idea is to find curves that can be parametrized in a suitable parameter space. In its original form one determines a curve in parameter space for a signal corresponding to all possible tracks with a given parametric form, to which it possibly could belong to. All such curves for the different signals are drawn in parameter space. This space is then discretized and entries are histogrammed—one divides parameter space into boxes and counts the number of curves in each box. If the peaks in the histogram exceed a given threshold, the corresponding parameter values define a potential track [409, 489, 490].

Data rate reduction. Data rate reduction can be achieved in different ways:

1. Generation and application of a software trigger capable of reducing the input data stream by selecting entire events.
2. Reduction in the size of the event data by selecting event dependent regions of interest (ROIs). By analysing the tracking information or by e.g. utilizing TRD information, such regions of interest can be defined. Recording only summary information and raw data of the ROIs (e.g. electron tracks) can reduce the data volume significantly.
3. Reduction in the size of the event data can also be achieved by compression techniques. General loss-less or slightly lossy methods can compress TPC data by factors 2–3. Using data modelling techniques can further reduce the data volume, and storing only quantized differences to a data model, i.e. a cluster and a local track model results in reduction factors of about 10 [487].

Event rate reduction—trigger. Based on results for example from TRD analysis and online TPC tracking information, the (sub)event building can be aborted. The HLT system is responsible for deriving such a trigger decision. The relevant tracks are identified, e.g. by the TRD, and a processing command is sent to the appropriate sector processors of the TPC. The HLT processor receives this information via its network from the TRD global trigger. This can easily be implemented at a rate of a few hundred Hz without presenting any particular technological challenges. The HLT system reconstructs tracks in the tracking chambers, combines track information coming from different detectors and finally decides to accept/reject the event.

Volume reduction—compression. The raw data volume per event can be reduced by compression techniques or by selective or partial readout.

Selective or partial readout: Subevents or ROIs can be defined on the basis of rough tracking information, including PID (particle identification) supplied by the TRD. All raw data inside these regions are recorded whereas all other data are dropped. For the lepton measurements the data volume can be reduced to candidate e^+e^- tracks, which would yield a few tracks per event.

Data compression. General loss-less or slightly lossy methods like entropy coders and vector quantizers can compress tracking detector data only by factors 2–3 [491]. The choice of the data model is of utmost importance in order to achieve high compression factors. All correlations in the data have to be incorporated into the model. Precise knowledge of the detector performance, i.e. analogue noise of the detector and the quantization noise, is necessary for creating a minimum information model of the data. Real-time pattern recognition and feature extraction are the input to the model. The data modelling scheme is based on the fact that the information content of the TPC are tracks, which can be represented by models of clusters and track parameters. The best compression method is to find a good model for the raw data and to transform the data into an efficient representation. Information is stored as model parameters and (small) deviations from the model. The results are coded using Huffman and vector quantization algorithms. The relevant pieces of information given by a tracking detector are the local track parameters and the parameters of the clusters belonging to this track segment [487].

3.19.2. System architecture: clustered SMP farm. The requirements of the HLT processing as outlined above are significant. First estimates indicate that about 1000 PC type processors will be required at ALICE turn-on, assuming today's algorithms and taking Moore's law into account. The best price/performance is achieved today using twin processor units.

Clustered SMP nodes, based on off-the-shelf PCs, and connected by a high bandwidth, low overhead network, can provide the necessary computing power for online pattern recognition and data compression. The interface to the custom detector front-end electronic systems is implemented by custom PCI interface cards (ReadOut Receiver Cards, RORC) housing the ALICE optical detector links. Therefore the implementation of an HLT RORC turns an HLT processing node into an HLT front-end processor receiving the RAW detector data.

Figure 3.51 outlines the HLT architecture [494]. The HLT front-end processor is a commercial off-the-shelf computer, such as a PC. It hosts the PCI RORC, which is plugged directly into its internal peripheral PCI bus, which is adopted by basically any computer to date. The main memory of the computer functions as the event buffer in this case. A further important advantage is the fact that the hardware could be purchased much later that would be possible for custom-designed components. In this respect, the project will profit much better from the constantly developing markets. As shown below, the architecture of an appropriate HLT PCI RORC is relatively simple. The RORC (figure 3.52) interfaces the optical fibre, receiving data from the front-end electronics of the detector, to the receiver nodes of the HLT system. Given the availability of an FPGA within the HLT front-end processor some of the I/O intense HLT processing will be off-loaded into those FPGAs, turning this device into an FPGA co-processor [492].

3.19.2.1. FPGA co-processor. Standard microcomputers and FPGAs are combined to form hybrid computers in which the FPGAs act as flexible co-processors. A typical architecture consists of a 64 bit/66 MHz PCI board containing the FPGA, local memory and fast I/O ports plugged into an off-the-shelf PC. A co-processor can, for instance, perform data compression at one point in time, and then be instantly reconfigured for pattern recognition. In general, the

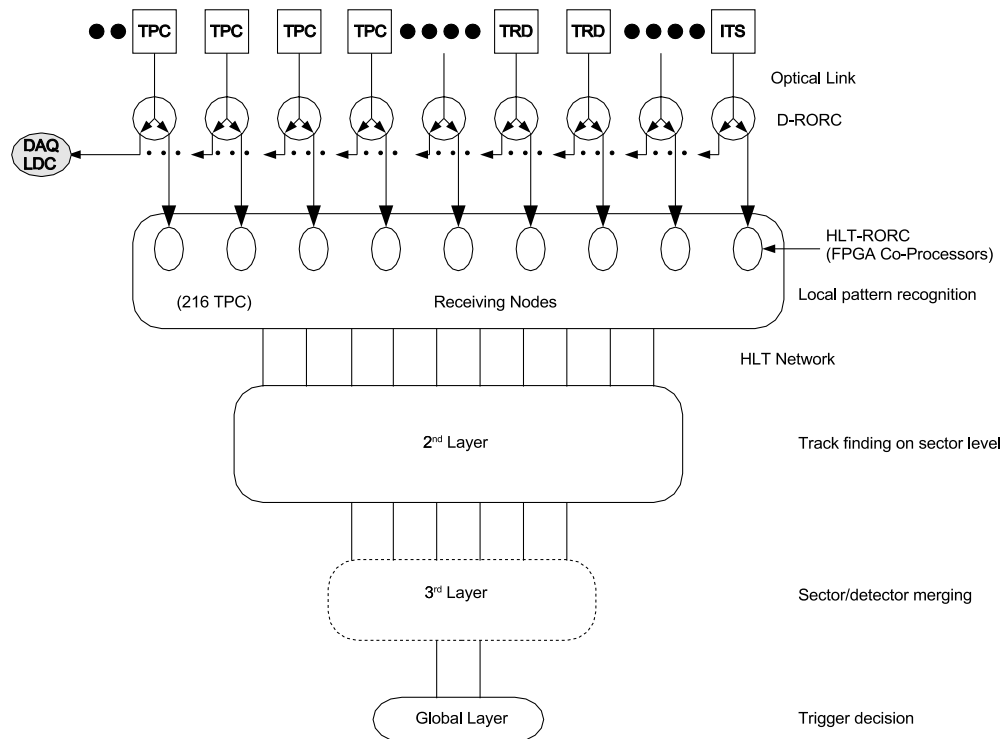


Figure 3.51. HLT architecture.

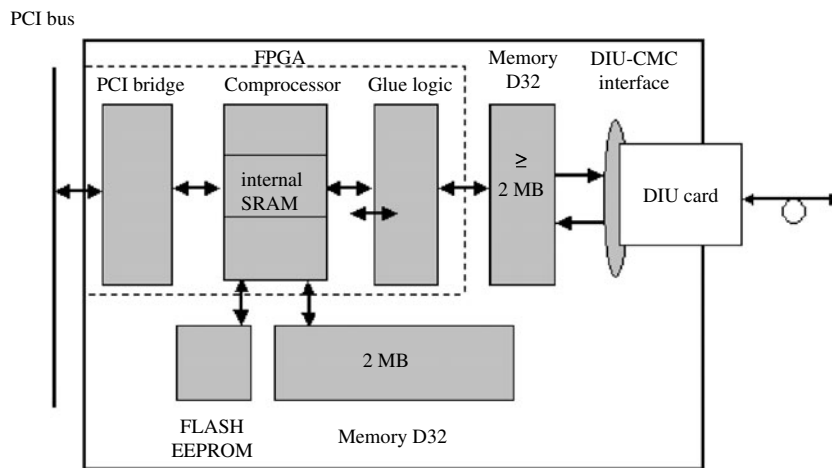


Figure 3.52. HLT-RORC-FPGA co-processor.

co-processor performs computation intensive, repetitive parts of an application, which would be slower on the microprocessor—so-called ‘hardware acceleration’. The above features make a FPGA co-processor a key element in the ALICE High Level Trigger. A fast cluster finder has been developed [493] and a Hough transformation tracker are currently being implemented on the FPGA co-processor.




Figure 3.53. Benchmarking of network technologies.

3.19.2.2. Network topology. The HLT layer of the ALICE experiment will make use of the inherent granularity of the tracking detector readout. The main tracking detector, TPC, consists of 36 sectors, each sector being divided into 6 subsectors. The data from each subsector are transferred from the detector into 216 receiver nodes of the PC farm via optical fibres. The receiver processors are all interconnected by a hierarchical network. Each sector is processed in parallel, results are then merged on a higher level, first locally in 108 sector and 72 sextet nodes and finally in 12 event nodes which have access to the full detector data.

Data flow analysis shows that the integrated data flow of the HLT may well exceed 20 GB s^{-1} . Therefore the network overhead has to be taken well into account, in order to ensure enough available processing on the farm. Figure 3.53 shows measurements of the networking overhead [495]. Both, throughput and overhead, are measured for the two protocols. Note that exactly the same hardware is used in both cases. When comparing the TCP/IP with the novel Scheduled Transfer Protocol (STP), the achieved bandwidth and the involved CPU overhead differ dramatically. With an overhead of up to 50% of the CPU performance on both sides of a communication with TCP/IP, a cluster requiring large data rates will scale very poorly. Other measurements report comparable results. One important observation is that the processing overhead does not scale linearly with the CPU performance as this I/O processing, by definition, is I/O-bound, and thus does not scale with Moore's law. Therefore, faster processors may not have a smaller overhead fraction.

3.19.2.3. Farm hierarchy. The hierarchy of the parallel farm is adapted both to the parallelism in the data and to the complexity of the pattern recognition (see figure 3.51). The first layer of nodes receives data from the detector and performs the pre-processing task, i.e. cluster and track seed finding on the subsector level. The next two levels of nodes exploit the local neighbourhood: track segment finding on a sector level. Finally all local results are collected from the sectors and combined on a global level: track segment merging and final track fitting.

3.19.2.4. Farm specifications. The HLT is a large-scale compute-cluster, which has a scale similar to a Tier-1 regional centre of the HEP Data GRID. Given standard failure models and the experience of computer centres operating large-scale clusters, it is obvious that a strong effort has to be made regarding the reliability and availability. In particular, during running period, failures have to be minimized. On the other hand, standard techniques focusing on failure avoidance typically increase the cost significantly. The HLT architecture will therefore

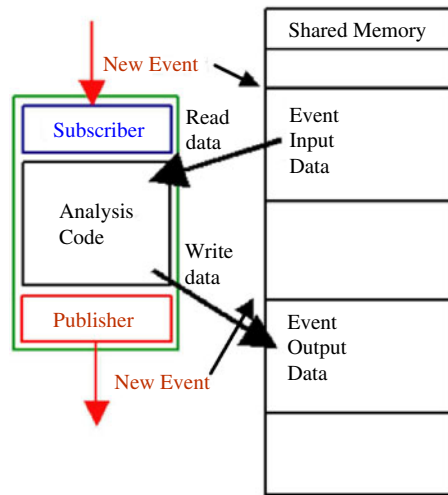


Figure 3.54. The HLT communication primitives.

avoid any single point of failures and implement means to repair or work-around any given failure as autonomously as possible. A maximum of remote diagnostics and control will be provided. Any node may fail at any time without the loss of a single event.

3.19.2.5. Interprocess communication. In order to implement an interface between the different readout steps, producing and consuming data objects, a generic framework has been developed [496], which takes into account a number of general principles, including the avoidance of any unnecessary copying of data. Figure 3.54 sketches the principle. Each analysis module within the HLT consumes data and produces derived output data within its memory. A class of communication primitives subscribes to data producers downstream and publishes the results of the given analysis module upstream. In all cases no event data is copied unless it has to be shipped to another node. The analysis jobs themselves are data driven, executing the appropriate analysis routines. To pass an event to another subscriber and receive it back as completed requires very little overhead in the microsecond range on today's computers. This architecture allows the adoption of any network architecture and topology. It further allows reconfiguration while the system is running without losing a single event.

4. Offline computing and Monte Carlo generators

This section describes the software environment which was used to produce the results presented in this volume and in PPR Volume II. In addition, we wish to put into perspective the main decisions of the offline project for the development and organization of the software. First a general description of AliRoot, the ALICE offline framework, is given, starting with a short historical background followed by a description of the simulation, reconstruction, and analysis architecture. Then the event generators commonly used for heavy-ion and pp interactions are summarized.

4.1. Offline framework

The final objective of the offline framework is to reconstruct and analyse the physics data coming from simulated and real interactions. Also, the optimization of the detector design relies in a reliable simulation chain. The AliRoot framework was used to perform simulation studies for the Technical Design Reports of all ALICE detectors and optimize their design. In the context of this document, it is used to evaluate the physics performance of the full ALICE detector. In addition, this process allow us to assess the functionality of our framework towards the final goal of extracting physics from the data.

4.1.1. Overview of AliRoot framework.

4.1.1.1. Strategy for AliRoot development. The development of the ALICE offline framework started in 1998. At that time, the size and complexity of the LHC computing environment had already been recognized. Since early 1990s it had been clear that the new generation of programs would be based on Object-Oriented (OO) techniques [497]. A number of projects were started, to replace the FORTRAN CERNLIB [498], including PAW [499] and GEANT3 [500], with OO products. At the time when the ALICE offline decided the move to C++, these products were not developed enough for production runs.

The other LHC experiments decided to maintain two lines of software development, one based on the old FORTRAN CERNLIB, dedicated to detector and physics performance studies, and one based on the new generation of CERN/IT products to be used at a later stage. ALICE decided not to adopt this strategy because the offline team was small and could not afford two lines of development. Moreover, training the physicists to C++ is a complex task that could have taken several years. If we had maintained a working FORTRAN environment, they would have probably never adopted the new generation software. In addition, the pressure of the real users on the developers of the new OO environment was mandatory; without such a pressure the products under development might not correspond to the real needs. In ALICE, the offline team developing software tools and the physicists developing software to evaluate the detector performance are in a single group. While improving communication, this puts a lot of pressure on the software developers to provide working tools while developing the offline system. In this context, rapid and complete conversion to OO/C++ was supported by the ALICE physics community under the condition that a working environment at least as good as GEANT3, CERNLIB and PAW could be provided in a short time. After a short but intense evaluation period, the ALICE offline team concluded that one such framework existed, namely the ROOT system [501], which is now the de-facto standard of HEP software and is in widespread use at CERN.

The decision to adopt the ROOT framework was taken in November 1998 and the development of the ALICE offline framework, AliRoot [502], started immediately, making full use of all the ROOT potential. The process that led us to the fundamental technical choices has created a tightly knit offline team with one single line of development. The move to OO was completed successfully resulting in one single framework, entirely written in C++, with some external programs (hidden to the users) still in FORTRAN. All ALICE users adopted the AliRoot framework and the Technical Design Reports were written based on simulation studies performed using it while the system was in continuous development. This choice allowed us to address both the immediate needs and the long-term goals of the experiment with the same framework, which evolves into the final one, with the support and participation of all ALICE users.




Figure 4.1. The ROOT framework and its application to HEP.

4.1.1.2. ROOT framework. A framework is a set of software tools that enables data processing. For example the FORTRAN CERNLIB was a toolkit to build a framework. PAW was the first example of integration of tools into a coherent ensemble specifically dedicated to data analysis. ROOT [501] is a new generation, Object Oriented framework for large-scale data-handling applications. The ROOT framework, schematically shown in figure 4.1, provides an environment for the development of software packages for event generation, detector simulation, event reconstruction, data acquisition and a complete data analysis framework including all PAW features.

ROOT is written in C++ and offers, among other features, integrated I/O with class schema evolution, an efficient hierarchical object store with a complete set of object containers, C++ as a scripting language and a C++ interpreter, advanced statistical analysis tools (multidimensional histogramming, several commonly used mathematical function, random number generators, multi-parametric fit, minimization procedures, cluster finding algorithms etc), hyperized documentation tools and advanced visualization tools.

The user interacts with ROOT via a graphical user interface, the command line or batch scripts. The command and scripting language is C++ (using the interpreter) and large scripts can be compiled and dynamically linked.

ROOT presents a coherent and well integrated set of classes that easily inter-operate via an object bus provided by the interpreter dictionaries (these provide extensive Run Time Type Information, RTTI, of each object active in the system). This makes ROOT an ideal basic infrastructure component on which an experiment's complete data handling chain can be built: from DAQ, using the client/server and shared memory classes, to database, analysis, and data presentation.

The backbone of the ROOT architecture is a layered class hierarchy with, currently, around 650 classes grouped in about 40 libraries divided into 20 categories (about 800 000 lines of code). This hierarchy is organized in a single-rooted class library, that is, most of the classes inherit from a common base class `TObject`. While this organization is not very commonly found in C++, it has proved to be well suited for our needs (and indeed for almost all successful class libraries: Java, Smalltalk, MFC, etc). It has enabled the implementation of some essential infrastructure inherited by all descendants of `TObject`. However, we can also have classes not inheriting from `TObject` when appropriate (e.g. classes that are used as built-in types, like

TString). A ROOT-based program links explicitly with only a few core libraries. At run-time the system dynamically loads additional libraries depending on its needs.

One of the key components of ROOT is the CINT C/C++ interpreter. The ROOT system embeds CINT to be able to execute C++ scripts and C++ command line input. CINT also provides ROOT with extensive RTTI capabilities. CINT covers 95% of ANSI C and about 90% of C++. CINT is complete enough to be able to interpret its own 80 000 lines of C and to let the interpreted interpreter interpret a small program. The advantage of a C/C++ interpreter is that it allows for fast prototyping since it eliminates the typical time-consuming edit/compile/link cycle. Once a script or program is finished, it can be compiled with a standard C/C++ compiler to machine code for full machine performance. Since CINT is very efficient (for example, for/while loops are byte-code compiled on the fly), it is quite possible to run small programs in the interpreter. In most cases, CINT outperforms other interpreters like Perl and Python.

The ROOT system has been ported to all known Unix variants (including many different C++ compilers) and to Windows 9x/NT/2000/XP.

ROOT is currently widely used in Particle and Nuclear physics, notably in all major US labs (FermiLab, Brookhaven, SLAC), and most European labs (CERN, DESY, GSI). Although initially developed in the context of Particle and Nuclear physics it can be equally well used in other fields where large amounts of data need to be processed, like astronomy, biology, genetics, finance, insurance, pharmaceutical, etc.

4.1.1.3. ROOT in ALICE. ROOT offers a number of important elements (summarized in section 4.1.1.2 and detailed in [501]) which were exploited in AliRoot and were the basis for the successful migration of the FORTRAN users to OO/C++ programming.

The ROOT system can be seamlessly extended with user classes and libraries that become effectively part of the system. These libraries are loaded dynamically and the contained classes share the same services of the native ROOT classes, including object browsing, I/O, dictionary and so on.

ROOT can be driven by scripts having access to the full functionality of the classes. For production, it eliminates the need for configuration files in special formats, since the parameters can be entered via the setters of the classes to be initialized. The system is also particularly effective for fast prototyping. Developers can work with a script without concerns about compilation and link in a single ROOT interactive session. When the code is validated, it can be compiled and made available via shared libraries, and the process can restart. Since one single language is used, there is no need for an external description of the classes or a translation of the algorithms. This lowers the threshold for new C++ users. This has been one of the major enabling factors in the migration of the users to OO/C++ environment.

The ROOT system is now being interfaced with the emerging Grid middleware in general and, in particular, with the ALICE-developed AliEn system [503]. In conjunction with the PROOF [504] system, which extends ROOT capabilities on parallel computing systems and clusters, this will provide a distributed parallel computing platform for large-scale production and analysis.

4.1.1.4. AliRoot framework. The role of a framework is shown schematically in figure 4.2. Data are generated via simulation programs, (i.e. Monte Carlo event generators and detector response simulation packages) and are then transformed into data representing the detector response. The data produced by the event generators contain the full information about the generated particles (i.e. PID and momentum). As these events are processed via the simulation

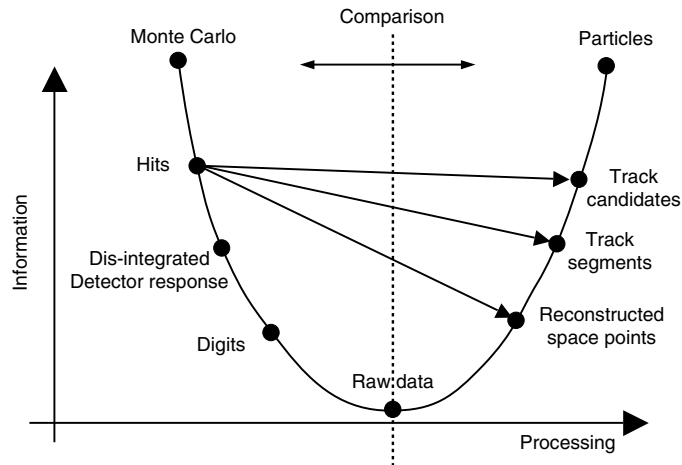


Figure 4.2. Data processing framework.

chain, the information is disintegrated and reduced to that generated by particles when crossing a detector (as detailed in section 4.1.2.3). Finally, ‘raw data’ are produced from simulated events. At this point, the reconstruction and analysis chain is activated. This takes as input raw data, real or simulated. The reconstruction algorithms have to reconstruct the full information about the particles trajectory and mass starting, for example, with space points and track segments in the case of tracking detectors. To evaluate the software and detector performance, simulated events are processed through the whole cycle and finally the reconstructed particles are compared to the Monte Carlo generated ones.

The user can intervene in this cycle provided by the framework to implement his/her own analysis of the data or to replace any part of it with his/her own code. I/O and user interfaces are part of the framework, as are data visualization and analysis tools and all procedures that are considered of general enough interest to be introduced into the framework. The scope of the framework evolves with time as the needs and understanding of the physics community evolves.

The basic principles that have guided us in the design of the AliRoot framework are re-usability and modularity. There are almost as many definitions of these concepts as there are programmers. However, for our purpose, we adopt an operative heuristic definition that expresses our objective to minimize the amount of user code unused or rewritten and maximize the participation of the physicists in the development of the code.

Modularity allows replacement of parts of our system with minimal or no impact on the rest. Not every part of our system is expected to be replaced. Therefore we are aiming at modularity targeted to those elements that we intend to change. For example, we require the ability to change the event generator or the transport Monte Carlo without affecting the user code. There are elements that we do not plan to interchange, but rather to evolve in collaboration with their authors such as the ROOT I/O subsystem or the ROOT User Interface (UI), and therefore no effort is made to make our framework modular with respect to these. Whenever an element has to be modular in the sense above, we define an abstract interface to it. The codes from the different detectors are independent so that different detector groups can work concurrently on the system while minimizing the interference. We understand and accept the risk that at some point the need may arise to make modular a component that was

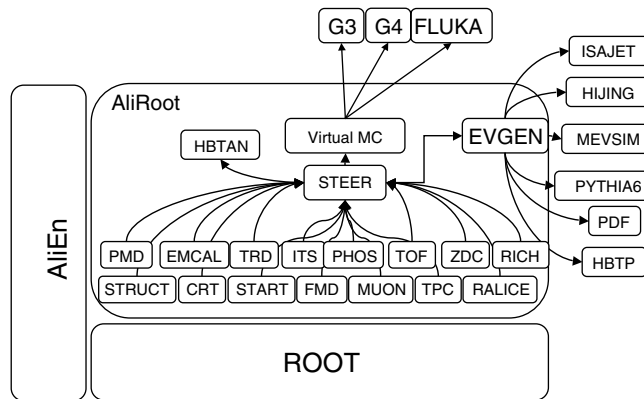


Figure 4.3. Schematic view of the AliRoot framework.

not designed to be so. For these cases, we have developed a strategy that can handle design changes in production code.

Re-usability is the protection of the investment made by the programming physicists of ALICE. The code embodies great scientific knowledge and experience and is thus a precious resource. We preserve this investment by designing a modular system in the sense above and by making sure that we maintain the maximum amount of backward compatibility while evolving our system. This naturally generates requirements on the underlying framework, prompting developments such as, for example, the introduction of the automatic schema evolution in ROOT.

The AliRoot framework is schematically shown in figure 4.3. The STEER module provides steering, run management, interface classes, and base classes. The detectors are independent modules that contain the code for simulation and reconstruction while the analysis code is progressively added. Detector response simulation can be performed via different transport codes like GEANT3 [500], GEANT4 [505], and, shortly, FLUKA [506]. The user can decide which one to load and thanks to the Virtual Monte Carlo abstract interface, detailed in section 4.1.2.3 and [507], the transition from GEANT3 to GEANT4 and to FLUKA is seamless, since no user code has to be changed; only a different shared library has to be loaded. The user code is all in C++, including the geometry definition.

The same technique is used to access different event generators. The event generators are accessed via a virtual interface, called `AliGenerator`, that allows the loading of different generators at run time. Most of the generators are in FORTRAN but the combination of dynamically loadable libraries and C++ wrapper classes implementing the virtual generator interface makes this completely transparent to the users.

In the AliRoot framework, detector modules are independent. Often one module needs the data of another one and this can result in a complicated network of dependencies with the additional problem of data ownership. To avoid this, we have developed a white-board approach where shared data are posted to a ROOT shared memory area called folder and are accessible to all modules. This breaks the N^2 dependencies between modules and replaces them with N references to the white-board. Data are read and posted to the folders, thanks to the automatic ROOT I/O mapping of folders, or explicitly written from the folders to disk.

While the combination of class dictionary and folders allows interactive browsing of the data and object model, such functionality is normally missing for the class methods. To obviate this problem and allow generic programming and interactive browsing of procedures,

the ROOT system introduces the concept of tasks. Tasks are generic procedures that can be invoked and browsed interactively. In AliRoot tasks are used to describe reconstruction and analysis procedures.

4.1.2. Simulation

4.1.2.1. Event generators. Heavy-ion collisions produce a very large number of particles in the final state. This is a challenge for the reconstruction and analysis algorithms which require a predictive and precise simulation of the detector response.

The ALICE experiment was designed when the highest nucleon–nucleon centre-of-mass energy heavy-ion interactions was at 20 GeV per nucleon–nucleon pair at CERN SPS, i.e. a factor of about 300 less than the LHC energy. Model predictions, discussed in section 1.3.1, for the particle multiplicity in Pb–Pb collisions at LHC vary from 1400 to 8000 charged particles per rapidity unit at mid-rapidity. In summer 2000 the RHIC collider came online. The RHIC data seem to suggest that the LHC multiplicity will be on the lower side of the predictions. However, the RHIC top energy of 200 GeV per nucleon–nucleon pair is still 30 times less than the LHC energy. The extrapolation is so large that both the hardware and software of ALICE had to be designed to cope with the highest predicted multiplicity. On the other hand, we have to use different generators for the primary interaction, since their predictions are quite different at LHC energies. The simulations of physical processes are confronted with several problems:

- Existing event generators give different predictions for the expected particle multiplicity, p_t and rapidity distributions and the dependence of different observables on p_t and rapidity at LHC energies.
- Most of the physics signals, like hyperon production, high- p_t observables, open charm and beauty, quarkonia, etc even at lower energies, are not exactly reproduced by the existing event generators.
- Simulation of small cross-section observables would demand prohibitively long runs to simulate a number of events that is commensurable with the expected number of detected events in the experiment.
- The existing generators do not simulate correctly some features like momentum correlations, azimuthal flow, etc.

Nevertheless, to allow for efficient simulations we have developed the offline framework such that allows for a number of options:

- The simulation framework provides an interface to several external generators, like for example HIJING [508] as detailed in section 4.2.
- A simple event generator based on parametrized η and p_t distributions can provide a signal-free event with multiplicity as a parameter.
- Rare signals can be generated using the interface to external generators like PYTHIA [509] or simple parametrizations of transverse momentum and rapidity spectra defined in function libraries.
- The framework provides a tool to assemble events from different signal generators (event cocktails).
- The framework provides tools to combine underlying events and signal events on the primary particle level (cocktail) and on the digit level (merging).
- ‘Afterburners’ are used to introduce particle correlations in a controlled way.




Figure 4.4. `AliGenerator` is the base class that has the responsibility of generating the primary particles of an event. Some realizations of this class do not generate the particles themselves but delegate the task to an external generator like PYTHIA through the `TGenerator` interface.

The implementation of these strategies is described below. The predictions of different Monte Carlo generators for heavy-ion collisions at LHC energy are described in section 4.2.

The theoretical uncertainty on the description of heavy-ion collisions at LHC has several consequences for our simulation strategy. A large part of the physics analysis will be the search for rare signals over an essentially uncorrelated background of emitted particles. To avoid being dependent on a specific model, and to gain in efficiency and flexibility, we generate events from a specially developed parametrization of a signal-free final state, see for details section 4.2.1.2. This is based on a parametrization of the HIJING pseudo-rapidity (η) distribution and of the transverse momentum (p_t) distribution of CDF [510] data. To simulate the highest anticipated multiplicities we scale the η -distribution so that 8000 charged particles per event are produced in the range $|\eta| < 0.5$. Events generated from this parametrization are sufficient for a large number of studies, such as optimization of detector and algorithms performance, e.g. studies of tracking efficiency as a function of particle multiplicity and occupancy.

To facilitate the usage of different generators we have developed an abstract generator interface called `AliGenerator` (see figure 4.4). The objective is to provide the user with an easy and coherent way to study a variety of physics signals as well as a full set of tools for testing and background studies. This interface allows the study of full events, signal processes, and a mixture of both, i.e. cocktail events.

Several event generators are available via the abstract ROOT class that implements the generic generator interface, `TGenerator`. Through implementations of this abstract base class we wrap FORTRAN Monte Carlo codes like PYTHIA, HIJING, etc that are thus accessible from the `AliRoot` classes. In particular the interface to PYTHIA includes the use of nuclear structure functions of PDFLIB.

In many cases, the expected transverse momentum and rapidity distributions of particles are known. In other cases the effect of variations in these distributions must be investigated. In both situations it is appropriate to use generators that produce primary particles and their decays sampling from parametrized spectra. To meet the different physics requirements in a modular way, the parametrizations are stored in independent function libraries wrapped into classes that can be plugged into the generator. This is schematically illustrated in figure 4.5 where four different generator libraries can be loaded via the abstract generator interface.

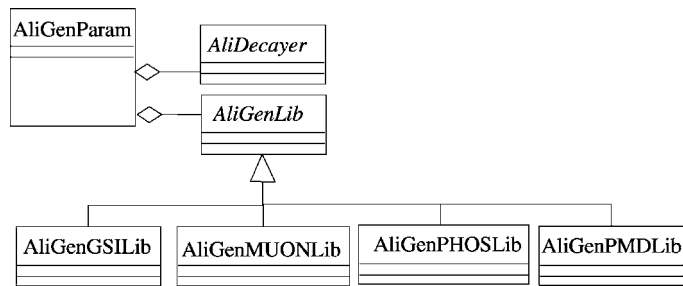


Figure 4.5. `AliGenParam` is a realization of `AliGenerator` that generates particles using parametrized p_t and pseudo-rapidity distributions. Instead of coding a fixed number of parametrizations directly into the class implementations, user defined parametrization libraries (`AliGenLib`) can be connected at run time allowing for maximum flexibility.




Figure 4.6. The `AliCocktail` generator is a realization of `AliGenerator` which does not generate particles itself but delegates this task to a list of objects of type `AliGenerator` that can be connected as entries (`AliGenCocktailEntry`) at run time. In this way different physics channels can be combined in one event.

It is customary in heavy-ion event generation to superimpose different signals on an event to tune the reconstruction algorithms. This is possible in `AliRoot` via the so-called cocktail generator (figure 4.6). This creates events from user-defined particle cocktails by choosing as ingredients a list of particle generators.

4.1.2.2. Afterburner processors and correlation analysis. The modularity of the event generator framework allows easy integration with the simulation steering class `AliRun` and with the objects that are responsible for changing the output of event generators or for assembling new events making use of the input of several events. These processors are generally called Afterburners. They are especially needed to introduce a controlled (parametrized) particle correlation into an otherwise uncorrelated particle sample. In `AliRoot` this task is further simplified by the implementation of a stack class (`AliStack`) that can be connected to both `AliRun` and `AliGenerator`. Currently, afterburners are used for the simulation of the two-particle correlations and azimuthal flow signals.

4.1.2.3. Detector response simulation. Much of the activity described in this report is a large virtual experiment where we generate thousands of events and analyse them in order to produce the presented results. The objectives are to study in detail the ALICE physics capabilities, to evaluate the physics goals of the experiment, and to verify the functionality of our software framework (from processing simulated or raw data to delivering physics results).

To carry out these objectives, it is important to have a high-quality and reliable detector response simulation code. One of the most common programs for full detector simulation is

GEANT3 [500] which, however, is a 20-year old FORTRAN program, not being (officially) further developed since 1993. GEANT4 [505] is being developed by CERN/IT as the OO simulation package for LHC. We also planned to evaluate and use FLUKA [506] as a full detector simulation program. Therefore we decided to build an environment that could take advantage of the maturity and solidity of GEANT3 and, at the same time, protect the investment in the user code when moving to a new Monte Carlo. In order to combine immediate needs and long term requirements into a single framework we wrapped the GEANT3 code in a C++ class (TGeant3) and we introduced a Virtual Monte Carlo abstract interface in AliRoot (see below). We have interfaced GEANT4 with our virtual Monte Carlo interface, and we are now interfacing FLUKA. We will thus be able to change the simulation engine without any modification in the user detector description and signal generation code. This strategy has proved very satisfactory and we are able to assure coherence of the whole simulation process which includes the following steps regardless the detector response simulation package in use:

- *Event generation of final-state particles:* The collision is simulated by a physics generator code (see sections 4.2 and 4.3) or a parametrization and the final-state particles are fed to the transport program.
- *Particle tracking:* The particles emerging from the interaction of the beam particles are transported in the material of the detector, simulating their interaction with it, and the energy deposition that generates the detector response (hits).
- *Signal generation and detector response:* During this phase the detector response is generated from the energy deposition of the particles traversing it. This is the ideal detector response, before the conversion to digital signal and the formatting of the front-end electronics is applied.
- *Digitization:* The detector response is digitized and formatted according to the output of the front-end electronics and the data acquisition system. The results should resemble closely the real data that will be produced by the detector.
- *Fast simulation:* The detector response is simulated via appropriate parametrizations or other techniques that do not require the full particle transport.

Virtual Monte Carlo interface. As explained above, our strategy to isolate the user code from changes of the detector simulation package was to develop a virtual interface to the detector transport code. We call this interface ‘virtual Monte Carlo’. It is implemented [507] via C++ virtual classes and is schematically shown in figure 4.7. The codes that implement the abstract classes are real C++ programs or wrapper classes that interface to FORTRAN programs.

Thanks to the virtual Monte Carlo we have converted all FORTRAN user code developed for GEANT3 into C++, including the geometry definition and the user scoring routines, StepManager. These have been integrated in the detector classes of the AliRoot framework. The output of the simulation is saved directly with ROOT I/O, simplifying the development of the digitization and reconstruction code in C++.

GEANT. GEANT3 is the detector simulation Monte Carlo code used extensively so far by HEP community for simulation of the detector response. However, it is no longer maintained and has several known drawbacks, both in the description of physics processes, particularly hadronic [511] and of the geometry. ALICE has spent considerable effort in evaluating GEANT4 via the virtual Monte Carlo; details on ALICE experience with GEANT4 can be found in [512]. We were able to keep the same geometry definition using the G3toG4 utility to translate from GEANT3 to GEANT4 geometry; in addition we have improved G3toG4 and made it

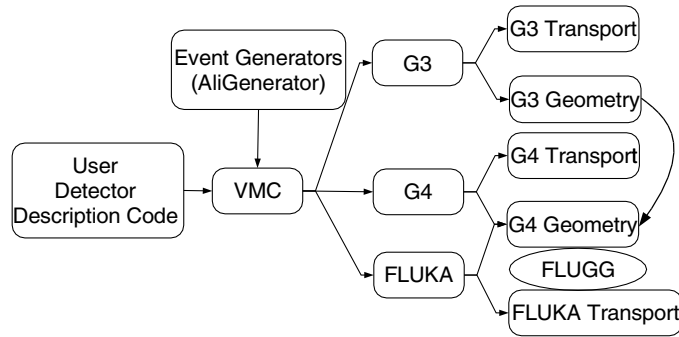


Figure 4.7. Virtual Monte Carlo.

fully operational. The virtual Monte Carlo interface allows us to run full ALICE simulations also with GEANT4 and to compare to the GEANT3 results; the advantage being that both simulation programs use the same geometry and the same scoring routines. An additional advantage is a substantial economy of effort. Using the same geometry description eliminates one of the major sources of uncertainty and errors in the comparison between different Monte Carlo, which comes from the fact that it is rather difficult to make sure that, when comparing two Monte Carlo on a particular experimental configuration, there are no differences in the geometry description and in the scoring. This exercise has exposed the GEANT4 code to a real production environment and we experienced several of its weaknesses. We faced several problems with its functionality that have required substantial user development. In particular the definition of volume or material-specific energy thresholds and mechanism lists are not straightforward as they were in GEANT3.

Before considering the migration to GEANT4, we performed a number of benchmark tests of the hadronic [513] and low-energy neutron transport [514]. The results of these tests point to weaknesses in the GEANT4 physics validation procedures. We have therefore decided to suspend the validation efforts of GEANT4 while maintaining the current interface with AliRoot. GEANT4 will be reconsidered as a possible candidate for full detector simulation when it has solved its problem with hadronic interactions.

An additional step is to replace the geometrical modeler of the different packages with a single one, independent from any specific simulation engine; the aim is to use the same geometrical modeler also for reconstruction and analysis. Thanks to the collaboration between the ALICE Offline project and the ROOT team, we have developed a geometrical modeler that is able to represent the ALICE detector, shown in figure III, and to replace the GEANT3 modeler for navigation in the detector. We are now interfacing it to FLUKA and discussion are under way with the GEANT4 team to interface it to the GEANT4 Monte Carlo.

FLUKA. FLUKA plays a very important role in ALICE for all the tasks where detailed and reliable physics simulation is vital, given its thorough physics validation and its almost unique capability to couple low-energy neutron transport with particle transport in a single program. These include background calculation, neutron fluence, dose rates, and beam-loss scenarios [515]. FLUKA is particularly important for ALICE in the design of the front absorber and beam shield. To ease the input of the FLUKA geometry, ALICE has developed an interactive interface [516], called AliFe, that allows setups described with FLUKA to be combined and modified easily. Figure 4.8 schematically describes the use of AliFe to prepare the input for FLUKA.

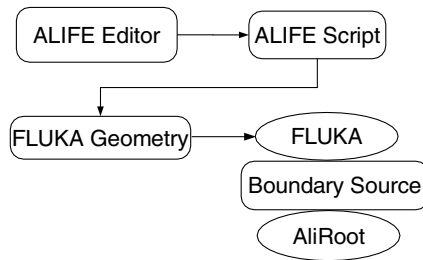


Figure 4.8. Example of the use of AliFe. The AliFe editor allows easy creation of an AliFe script, which is in fact FLUKA input geometry. FLUKA is then used to transport particles, including low-energy neutrons to a virtual boundary surface. The particles are then written to a file that is used as a source for a regular AliRoot simulation to evaluate the detector response to background.

To provide another alternative to GEANT3 for full detector simulation, we are now developing an interface with FLUKA, again via the virtual Monte Carlo, with the help of the FLUGG [516] package that allows running FLUKA with a GEANT4 geometry that, in our case, is generated by G3toG4.

Simulation framework. The AliRoot simulation framework can provide data at different stages of the simulation process [517], as described in figure 4.2. Most of the terminology comes from GEANT3. First, there are the so-called ‘hits’ that represent the precise signal left by the particle in the detector before any kind of instrumental effect, i.e. precise energy deposition and position. These are then transformed into the signal produced by the detector, ‘summable digits’ that correspond to the raw data before digitization and threshold subtraction. The introduction of summable digits is necessary because of the embedding simulation strategy elaborated for the studies of the Physics Performance Report. These summable digits are then transformed into digits that contain the same information as raw data, but in ROOT structures. The output of raw data in DATE (the prototype for the ALICE data acquisition system [518]) format has already been realized during the data challenges but is not yet implemented for all subdetectors.

The ALICE detector is described in great detail (see figure 4.9) including services and support structures, beam pipe, flanges, and pumps. The AliRoot geometry follows the evolution of the baseline design of the detector in order to continuously provide the most reliable simulation of the detector response. AliRoot is also an active part of this process since it has been used to optimize the design, providing different geometry options for each detector. The studies that provided the results presented in the PPR were performed with the baseline geometry.

Geometry of structural elements. The description of the front- and small-angle absorber regions is very detailed due to their importance to the muon spectrometer. The simulation has been instrumental in optimizing their design and in saving costs without a negative impact on the physics performance.

The material distribution and magnetic fields of the L3 solenoidal magnet and of the dipole magnets are also described in detail. The magnetic field description also includes the interference between the two fields. The field distributions are described by three independent maps for 0.2, 0.4 and 0.5 T solenoid L3 magnetic field strengths. Alternatively, it is possible to use simple parametrizations of the fields, i.e. constant solenoidal field in the barrel and a dipole field varying along z -direction for the muon spectrometer.




Figure 4.9. AliRoot simulation of the ALICE detector.

The space frame, supporting the barrel detectors, is described according to its final design taking into account modifications to the initial design such that it allows the eventual addition of a proposed Electromagnetic Calorimeter [519].

The design of the ALICE beam pipe has also been finalized. All elements that could limit the detector performance (pumps, bellows, flanges) are represented in the simulation.

Geometry of detectors. Most of the detectors are described by two versions of their geometry; a detailed one, which is used to accurately simulate the detector response and study their performance, and a coarse version that provides the correct material budget with minimal details, and is used to study the effect of this material budget on other detectors. For some detectors, different versions of the geometry description corresponding to different geometry options are selectable via the input C++ script at run time. In the following we give some examples. We remind the reader that some of this information is still subject to rapid evolution.

Both a detailed and a coarse geometry are available for the ITS. The detailed geometry of the ITS is very complicated and crucially affects the evaluation of impact parameter and electron bremsstrahlung. On the other hand, simulation of the coarse geometry is much faster when ITS hits are not needed.

A lot of attention has been devoted to the correct simulation of detector response and very sophisticated effects can already be studied and optimized via AliRoot simulation.

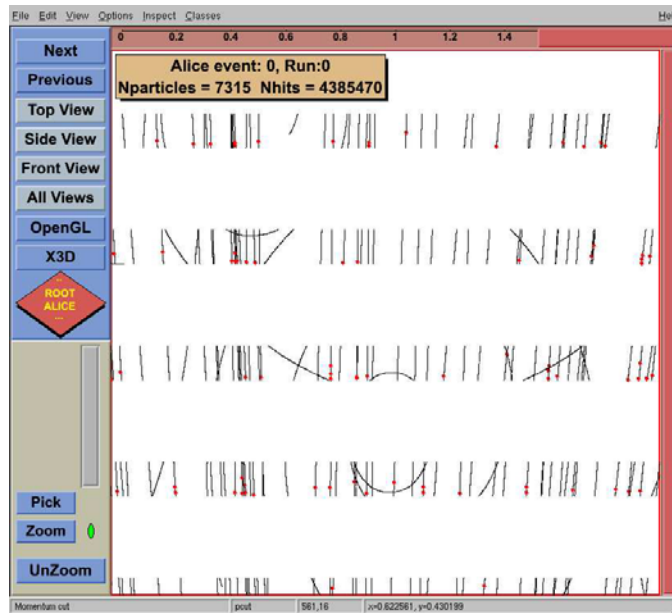


Figure 4.10. AliRoot event display of the TRD response. The photon- and electron-energy deposition in the different layers of the TRD detector as simulated by GEANT3 are shown.

Three configurations are available for the TPC. Version 0 is the coarse geometry, without any sensitive element specified. It is used for the material budget studies and is the version of interest for the outer detectors. Version 1 is the geometry version for the Fast Simulator. The sensitive volumes are thin gaseous strips placed in the Small (S) and Large (L) sectors at the pad-row centres. The hits are produced whenever a track crosses the sensitive volume (pad-row). The energy loss is not taken into account. Version 2 is the geometry version for the slow simulator. The sensitive volumes are S and L sectors. One can specify as sensitive volumes either all sectors or only a few of them, up to 6 S- and 12 L-sectors. The hits are produced in every ionizing collision. The tracking step is calculated for every collision from an exponential distribution. The energy loss is calculated from an $1/E^2$ distribution and the response is parametrized by a Mathieson distribution.

The TRD geometry is now complete, including the correct material budget for electronics and cooling pipes. The full response and digitization have been implemented allowing studies of open questions such as the number of time-bins, the 9- or 10-bit ADC, the gas and electronics gain, the drift velocity, and maximum Lorentz angle. The transition-radiation photon yield is approximated by an analytical solution for a foil stack, with adjustment of the yield for a real radiator, including foam and fibre layers from test beam data. This is quite a challenging detector to simulate, as both normal energy loss in the gas and absorption of transition-radiation photons have to be taken into account (figure 4.10).

During the signal generation several effects are taken into account: diffusion, one-dimensional pad response, gas gain and gain fluctuations, electronics gain and noise, as well as conversion to ADC values. Absorption and $\mathbf{E} \times \mathbf{B}$ effects will be introduced.

A detailed study of the background coming from slow neutron capture in Xe gas was performed [520]. The spectra of photons emitted after neutron capture are not included in standard neutron-reaction databases. An extensive literature search was necessary in order to simulate them. The resulting code is now part of the FLUKA Monte Carlo [521].

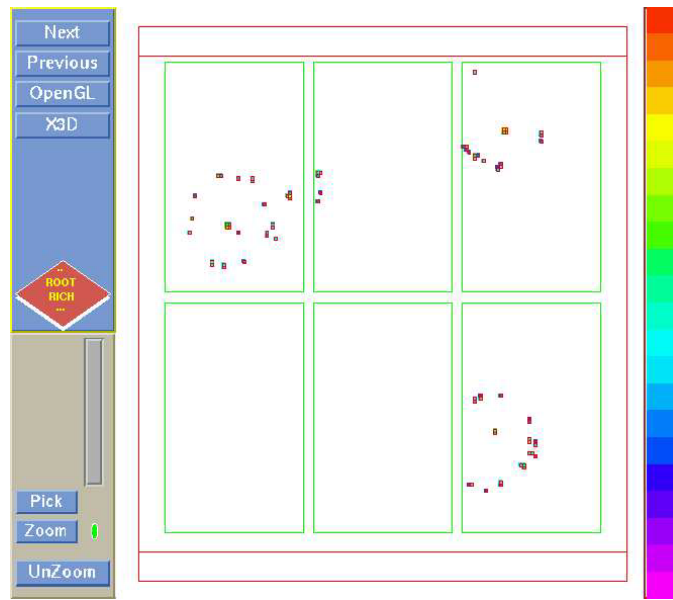


Figure 4.11. AliRoot event display of HMPID response. The rectangles correspond to pads with signals over threshold (digits).

The TOF detector covers a cylindrical surface of polar acceptance $|\theta-90^\circ| < 45^\circ$. Its total weight is 25 tons and it covers an area of about 160 m^2 with about 160 000 total readout channels (pads) and an intrinsic resolution of 60 ps. It has a modular structure corresponding to 18 sectors in φ and to 5 segments in z . All modules have the same width of 128 cm and increasing lengths, adding up to an overall TOF barrel length of 750 cm.

Inside each module the strips are tilted, thus minimizing the number of multiple partial-cell hits due to the obliqueness of the incidence angle. The double stack-strip arrangement (see section 3), the cooling tubes, and the material for electronics have been described in detail. During the development of the TOF design several different geometry options have been studied, all highly detailed.

The HMPID detector also poses a challenge in the simulation of the Cherenkov effect and the secondary emission of feedback photons (see figure 4.11). A detailed simulation has been introduced for all these effects and has been validated both by test-beam data and with the ALICE RICH prototype that has been operating in the STAR experiment.

The PHOS has also been simulated in detail. The geometry includes the Charged Particle Veto (CPV), crystals (EMC), readout (PIN or APD) and support structures. Hits record the energy deposition in one CPV and one EMC cell per entering particle. In the digits the contribution from all particles per event are summed up and noise is added.

The simulation of the ZDC in AliRoot requires tracking of spectator nucleons with Fermi spread, beam divergence, and crossing angle for over 100 m (see figure 4.12). The HIJING generator is used for these studies taking into account the correlations with transverse energy and multiplicity.

The muon spectrometer is composed of five tracking stations and two trigger stations. For stations 1–2 a conservative material distribution is adopted, while for stations 3–5 and for the trigger stations a detailed geometry is implemented. Supporting frames and support structures




Figure 4.12. Tracking of spectator nucleons with Fermi spread, beam divergence, and crossing angle. The horizontal scale is compressed with respect to the vertical one to show the development of the shower along the beam line till the ZDC.




Figure 4.13. Event display of muon-chamber simulation. The picture shows the pad response of the lower quadrant of station 3 for a full event.

are still coarse or missing but they are not very important in the simulation of the signal. The muon chambers have a complicated segmentation that has been implemented during the signal generation via a set of virtual classes. This allows changing the segmentation without modifying the geometry. An event display of a simulation of a muon-chamber quadrant is shown in figure 4.13.

Summable digits (pad hits) are generated taking into account the Mathieson formalism for charge distribution, while work is ongoing on the angular dependence, Lorentz angle and charge correlation.

The complex T0–FMD–V0–PMD forward detector system is still under development and optimization. Several options are provided to study their performance.

The description of ALICE geometry and the generation of simulated data are in place. Hence the offline framework allows the full event reconstruction including the main tracking devices. Figure IV shows an ALICE event display of such a simulated event. The framework also allows comparison with test-beam data. The early availability of a complete simulation has been an important point for the development of reconstruction and analysis code and user interfaces, which is now the focus of the development.

4.1.2.4. Fast simulation. Owing to the expected high particle multiplicity for heavy-ion collisions at LHC typical detector performance studies can be performed with a few thousand events. However, many types of physics analysis, in particular of low cross section observables, such as D meson reconstruction from hadronic decay channels, have to make use of millions of events. Computing resources are in general not available for such high-statistics simulations.

To reach the required sample size, fast simulation methods based on meaningful parametrizations of the results from detailed and consequently slow simulations are applied. The systematic error introduced by the parametrizations is in general small compared to the reduction of the statistical error. This is particularly true for the studies of the invariant-mass continuum below a resonance (cocktail plots).

It is hard to find a common denominator for fast simulation methods since they are very specific to the analysis task. As a minimum abstraction, we have designed base classes that allow for a representation of the detector or detector systems as a set of parametrizations of acceptance, efficiency, and resolution. The Muon Spectrometer fast simulation has been implemented using these classes.

Another interesting development concerns the fast simulation of the resolution and efficiency of tracking in the central barrel. In this approach, resolution and efficiency in TPC are obtained from the track parameters at the inner radius of the TPC, using a parametrization. After this, full tracking is performed for the inner tracking system, which is needed for detailed secondary vertex reconstruction studies. For details see [522].

4.1.2.5. Strategy for physics data challenges. To prepare our software to handle the real data, we decided to generate central events up to $dN_{\text{ch}}/d\eta|_{\eta=0} \approx 8000$, which is higher than all predictions. One such central event needs about 24 h on a typical PC to be fully simulated and digitized. For the study of one rare physics channel $\mathcal{O}(10^7)$ central heavy-ion events are needed requiring 25 000 typical PCs running continuously for 1 year. But beyond the CPU problem, we have also to consider that not all the rare events we are interested in are found in the standard heavy-ion generators. To cope with these two problems, we have developed special simulation procedures.

The first is the underlying-event merging procedure. As briefly explained earlier in this section, this method is based on the fact that the majority of the particles emitted in the final state are largely uncorrelated. It is thus possible to generate a signal-free underlying event of given multiplicity, η - and p_t -distributions, and to superimpose the signal on this event. We need to generate only a very limited number, $\mathcal{O}(10^3)$, of underlying events because we can reuse them several times to superimpose different signals. The typical underlying event is composed of pions and kaons and is digitized into analogue or summable digits. In a second pass,

a signal is generated with PYTHIA or with an appropriate parametrization and the signal is added to a randomly selected underlying event. Then the event is reconstructed and the signal is analysed. This of course requires the generation of several classes of underlying events for different multiplicity bins corresponding to different values of centrality or, equivalently, to different hypotheses of multiplicity.

Another strategy for reducing computing time is to introduce parametrizations. This has been done several times for specific studies, and a general framework will be soon introduced into AliRoot. An example of such a technique is the parametrization of the TPC response, which is based on the interpolation of previously computed track covariance matrices and allows to avoid the costly transport, digitization and reconstruction process.

4.1.3. Reconstruction

4.1.3.1. Tracking. Most of the ALICE detectors are tracking detectors. Each charged particle going through the detectors leaves a number of discrete signals that measure the position of the points in space where it has passed. The task of the reconstruction algorithms is to assign these space points to tracks and to reconstruct their kinematics. This operation is called track finding. In ALICE we require good track-finding efficiency for tracks down to $p_t = 100 \text{ MeV } c^{-1}$ even at the highest track density, with occupancy of the electronics channels exceeding 40% in the TPC inner rows at the maximum expected track multiplicity. Given this situation, most of the development is done for Pb–Pb central events, since lower multiplicities are considered an easier problem once the high-multiplicity ones can be handled. However, the opposite may be true for some quantities, such as the main vertex position, where the high track multiplicity will help to reduce the statistical error.

Several prototypes have been developed for the track reconstruction algorithms. In 1994–1995 a simple road tracker was prototyped in the TPC, starting from the ALEPH TPC reconstruction code and using as simulation input just hits smeared according to Gaussian distributions. The merging with the ITS was done by forcing a vertex constraint on particles leaving the TPC onto the two outer silicon strip layers. The experience gained with this prototype led us to develop a second one in 1996–1997. For the track following we used a Kalman filter as introduced by P Billoir for track fitting in 1983 [523]. This was the first ALICE program entirely written in C++. This program was still using just smeared hits as input and cylindrical geometry (i.e. we approximated both the TPC and the ITS sensitive elements with cylindrical surfaces). In 1998 we started development of the current prototype which was used at that time to optimize the parameters of the TPC. This version is now fully integrated within the AliRoot framework and makes use of a very detailed description of detector response and digitization. It is now extended to the ITS and TRD and will provide the basis for the final ALICE track reconstruction program. The detailed description of the algorithm and the performance will be given in section 5 of the PPR Volume II.

At first, we studied the tracking in the TPC using a classical approach where cluster finding precedes tracking. Clusters in pad-row and time directions provide the space points that are used for tracking with the Kalman filter. Recently we started another development where we defer the cluster finding at each pad row until all track candidates are propagated into its position. This way we know in advance which of the clusters are susceptible to be overlapped and we may attempt cluster deconvolution in a specific place.

The overall tracking starts with the track seeding in the outermost pad rows of the TPC. Different combinations of the pad rows were used with and without a primary vertex constraint. Typically more than one pass is done, starting with a rough vertex constraint, i.e. imposing

a primary vertex with a resolution of a few centimetres and then relaxing it. Then we propagate the track candidates in the TPC using Kalman filtering. When continuing the tracks from the TPC to the ITS during the first pass we impose a rather strict vertex constraint with a resolution of the order of $100\ \mu\text{m}$ or better. The other pass is done without a vertex constraint in order to reconstruct the tracks coming from the secondary vertices well separated from the interaction point.

We thus obtain the estimates of the track parameters and their covariance matrices in the vicinity of the interaction point. We then proceed with the Kalman filter in the outward direction. During this second propagation we remove from the track fit the space points with large χ^2 contributions—so-called outliers. In this way we obtain the track parameters and their covariance matrices at the outer TPC radius. We continue the Kalman filter into the TRD and then propagate the tracks towards the outer detectors, TOF, HMPID, and PHOS.

The tracking and propagation algorithms use common base classes in different detectors. The design of the reconstruction program is modular in order to easily exchange and test various parts of the algorithm. For example, we need to exchange different implementations of the clustering algorithm and the track seeding procedure. This modularity also allows us to use smeared positions of the GEANT hits instead of the detailed simulated detector response, very useful for testing purposes. In addition, each hit structure contains the information about the track that originated it. Although, this implies the storage of extra information, it has proved very useful for debugging the track reconstruction code. Thanks to our modular design and the use of common interfaces we can easily combine different versions of algorithms in the TPC and ITS.

As explained above, a vertex constraint is used during various steps of the tracking procedure. This is an important function of the ITS—the exact determination of both the primary- and secondary-vertex positions. The transverse primary-vertex position is defined by the transverse bunch size ($15\ \mu\text{m}$), while its z position is obtained by the correlation of hits in the silicon pixel detectors.

Determination of the primary-vertex position is the object of intense development in order to introduce provisions for alignment. We now also focus on algorithms to find secondary vertices and kinks to reconstruct neutral and charged particle decays and to improve the track reconstruction efficiency. Results have already been obtained for K_s^0 and Λ reconstruction and work is in progress.

Another very important tracking detector in ALICE is the muon arm which is designed to precisely measure the muon momenta behind the absorber. Because of the presence of a thick hadron absorber, the angular information is practically lost and we must impose the vertex constraint in order to obtain the muon directions at primary vertex.

The first study of track-finding algorithms for the muon arm was done in 1995 based on a combinatorial method using tracklet vectors in the various stations. Since then, a better cluster and tracklet finder has been introduced. Recently we have also started to use the Kalman filter for this task, with a great improvement in CPU performance. We are converging on a common software design with the barrel tracking detectors.

In conclusion, we have a well-defined strategy for tracking. The results from the TPC prototype tracker are well within specifications. The tracking efficiency of the ITS is affected by occupancy and more optimization is needed. Track matching between different detectors is basically finished and vertexing and V^0 reconstruction are already at the stage of promising prototypes.

4.1.3.2. Particle correlation analysis. Different analysis packages are currently under development within the AliRoot framework and are going to be touched upon in the different

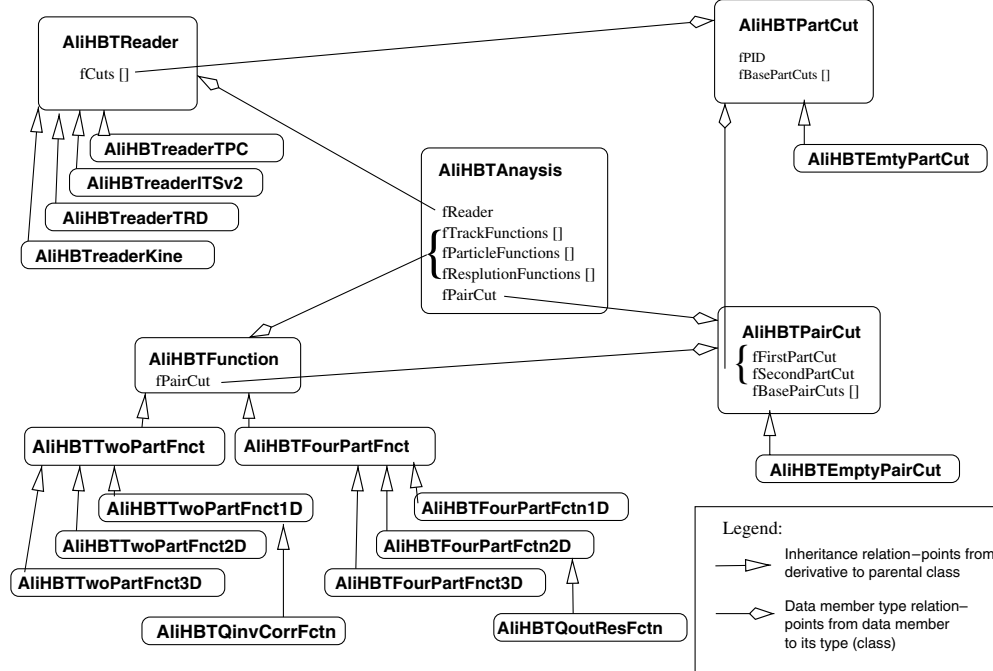


Figure 4.14. Scheme of the HBT Analyser structure.

sections of the PPR Volume II. Here we describe the particle correlation analysis software, called HBT Analyser [524], that was used for the particle correlation studies described in section 6 of the PPR Volume II.

The package was designed to be very modular. A schematic structure of the package is presented in figure 4.14. The main object is `AliHBTAnalysis` that performs the event mixing.

Users have the possibility to transform any given format of input data into the internal format of HBT Analyser using the reader object. The user must specify the reader object that inherits from the `AliHBTReader` class. This pure virtual class defines an abstract interface for readers used by `AliHBTAnalysis` for accessing data to be analysed. The reader has a data member array of cuts. Each cut defines the ranges of properties of particles. Particles that fulfill all criteria defined by any of the cuts are read.

The user sets the appropriate cuts in a macro. This feature allows the selection of a class of particles to be analysed at the moment of data reading which may consequently speed up the calculations. For example, if only positive pions are considered in the analysis, it is enough to set only one cut in the reader on particle type and its properties, like transverse momentum or rapidity. There is no need to apply any additional particle cuts for mixing.

The user defines function objects (`AliHBTFunction`) to be evaluated in the analysis. Two main classes of functions exist: one pair and two pair.

The first one, which is standard type, needs a single pair of particles at once to fill histograms.

The second one needs two pairs: one pair of reconstructed and one of simulated particles corresponding to the reconstructed ones. This function is used for resolution studies and to calculate the correlation functions distorted by the reconstruction procedure using the method of weights. The analysis of three-particle correlations is foreseen as well.

Both kinds of functions described above can be one-, two- or three-dimensional. The skeleton of these virtual classes allows easy creation of new functions. Several function classes are already included in the package.

The classes responsible for particle and pair cuts are `AliHBTPartCut` and `AliHBTPairCut` respectively. These are of very similar design. They have a data member list of basic cuts. Each basic cut examines only one given property of the particle or the pair (e.g. p_t or Q_{inv}), respectively. The user defines and configures cuts in the macro. They can be applied on many levels of the analysis chain. Proper cut configuration can substantially speed up the analysis.

The user can very easily add a new base cut different from those in the standard set. Empty particle and pair cuts that accept all particles are also implemented.

4.1.4. Distributed computing and the Grid. The investment for LHC computing is massive. For example, for ALICE only we foresee:

- data acquisition rate of 1.25 GB s^{-1} in heavy-ion mode;
- approximately 5 PB of data stored on tape each year;
- at least 1.5 PB of data permanently online on disks;
- an amount of CPU power estimated at about 18 MSI2k equivalent to about 25 000 PCs of the year 2003;
- of the order of 10–20 MCHF of hardware (depending on staging), excluding media, personnel, infrastructure, and networking.

It is very likely that funding agencies will prefer to provide computing resources locally in their own country. Also competence is naturally distributed and it would be difficult to move it to CERN. This has already been recognized some time ago by the HEP community and has been formalized in the so-called Monarc model [525] shown in figure 4.15. This is a distributed model where computing resources are concentrated in a hierarchy of centres called Tiers, where Tier-0 is CERN, Tier-1s are the major computing centres, Tier-2s the smaller regional computing centres, Tier-3 the university departmental computing centres and Tier-4 the user workstations.

In such a model the raw data will be stored at CERN where a Tier-1 centre for each experiment will be hosted. Tier-1 centres not at CERN will collectively store a large portion of the raw data, possibly all, providing a natural backup. Reconstruction will be shared by the Tier-1 centres, with the CERN Tier-1 playing a major role because of its privileged situation of being very close to the source of data. Subsequent data reduction, analysis and Monte Carlo production will be a collective operation where all Tiers will participate, with Tier-2s being particularly active for Monte Carlo and analysis.

The basic principle of such a model is that every physicist should have in principle equal access to the data and resources. The resulting system will be extremely complex. We expect several hundreds of components at each site with several tens of sites. A large number of tasks will have to be performed in parallel, some of them following an orderly schedule (reconstruction, large Monte Carlo production, and data filtering) and some being completely unpredictable (single-user Monte Carlo production or analysis). To be usable, the system will have to be completely transparent to the end user, essentially looking like a single system.

Unfortunately, the building blocks to realise such a system are missing: distributed resource management; distributed name-space for files and objects; distributed authentication; local resource management of large clusters; transparent data replication and caching; WAN-LAN monitoring; distributed logging and bookkeeping and so on. However the HEP




Figure 4.15. The Monarc model.

community is not alone in this endeavour. All the above issues are central to the new developments in the US and in Europe under the collective name of Grid [526].

The Grid was born to facilitate the development of new applications based on high-speed coupling of people, computers, databases, instruments, and other computing resource. The Grid should allow ‘dependable, consistent, pervasive access to high-end resources’ for:

- online instruments;
- collaborative engineering;
- parameter studies;
- browsing of remote datasets;
- use of remote software;
- data-intensive computing;
- very large-scale simulation.

This is the latest development of an idea born in the 1980s under the name of meta-computing and complemented by the Gbit test-beds in the early 1990s. The issues that are now the primary aim of the Grid are crucial to the successful deployment of LHC computing in the sense indicated by the Monarc model. Therefore HEP computing has become interested in Grid technologies, resulting in the launching of several Grid projects where HEP plays an important role. ALICE is actively involved in many of them [527].

Grid technology holds the promise of greatly facilitating the exploitation of LHC data for physics research. Therefore ALICE is very active on the different Grid test-beds where the Grid prototype middleware is deployed. The objective is to verify the functionality of the middleware, providing feedback to its authors, and to prototype the ALICE distributed computing environment.

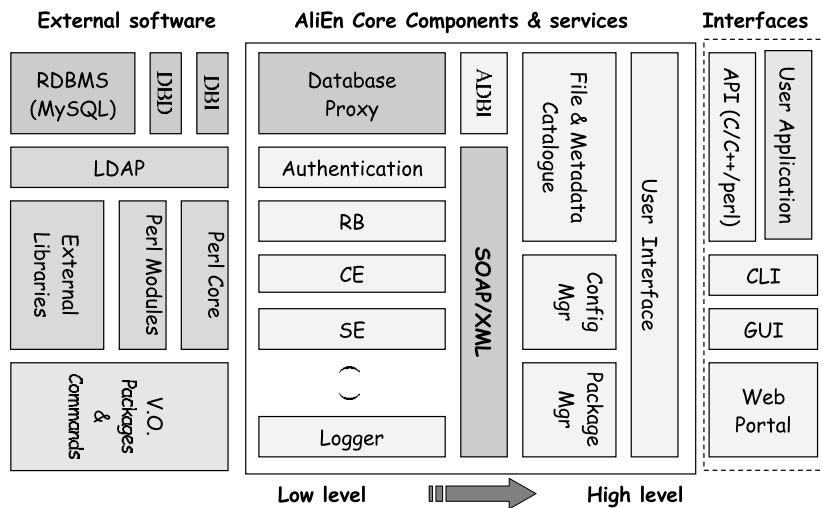


Figure 4.16. Architecture of AliEn—building blocks.

This activity, both very useful and interesting in itself, is hindered by the relative lack of maturity of the middleware. This middleware is largely the result of leading-edge computer science research and is therefore still rather far from production quality software. Moreover, standards are missing and although all the present Grid projects use the GLOBUS toolkit [528], their middleware is very different even if the functionality is similar. This makes it difficult for ALICE to exploit this software to run the productions that are more and more needed for the physics performance studies in a distributed environment, and to harness the resources of different computer centres.

To alleviate these problems, while providing a stable and evolutionary platform, ALICE has developed the AliEn [503] (ALICE ENvironment) framework with the aim of offering to the ALICE user community transparent access to computing resources distributed worldwide. The intention is to provide a functional computing environment that fulfils the needs of the experiment in the preparation phase and, at the same time, defines stable interface to the end users that will remain in place for a long time, shielding the ALICE core software from inevitable changes in the technology that makes distributed computing possible [530].

The system is built around Open Source components (see figure 4.16), uses Web Services model and standard network protocols to implement the computing platform that is currently being used to carry out the production of Monte Carlo data at over 30 sites on four continents. Only 1% (around 30k physical lines of code in PERL) is native AliEn code while 99% of the code has been imported in the form of Open Source packages and PERL modules.

AliEn has been primarily conceived as the ALICE user interface into the Grid world. As new middleware becomes available, we shall interface it with AliEn, evaluating its performance and functionality. Our final objective is to reduce the size of the AliEn code, integrating more and more high-level components from the Grid middleware, while preserving its user environment and possibly enhancing its functionality. If this is found to be satisfactory, then we can progressively remove AliEn code in favour of standard middleware. In particular, it is our intention to make AliEn services compatible with the Open Grid Services Architecture (OGSA) that has been proposed as a common foundation for future Grids. We would be satisfied if, in the end, AliEn would remain as the ALICE interface into the Grid middleware. This would preserve the user investment in AliEn and, at the same time, allow ALICE to




Figure 4.17. Interaction diagram of key AliEn components for typical analysis use case.

benefit from the latest advances in Grid technology. This approach also has the advantage that the middleware will be tested in an existing production environment. Moreover, interfacing different middleware to an existing high-level interface will allow a direct verification of their interoperability and a comparison of their functionality. At the moment, AliEn is interfaced to GLOBUS middleware. Our medium-term plan is to interface it to both EDG and iVDGL, gradually replacing the functionality provided at the moment by AliEn with the middleware tools. The AliEn functionality and UI will be the same across different middleware and new versions will be integrated and tested in the system. AliEn supports various authentication schemes, implemented using the extensible Simple Authentication and Security Layer (SASL) protocol, in particular GLOBUS GSI/GSSAPI that makes it compatible with the EDG security model.

AliEn Web Services play the central role in enabling AliEn as a distributed computing environment. The user interacts with them by exchanging SOAP messages and they constantly exchange messages between themselves behaving like a true Web of collaborating services. AliEn consists of the following key components and services: the authentication, authorization and auditing services; the workload and data management systems; the file and metadata catalogues; the information service; Grid and job monitoring services; storage and computing elements (see figure 4.17).

As opposed to the traditional push architecture, AliEn workload management system is based on a pull approach. A central service manages all the tasks, while computing elements are defined as 'remote queues' and can, in principle, provide an outlet to a single machine dedicated to running a specific task, a cluster of computers, or even an entire foreign Grid. When jobs are submitted, they are sent to the central queue.

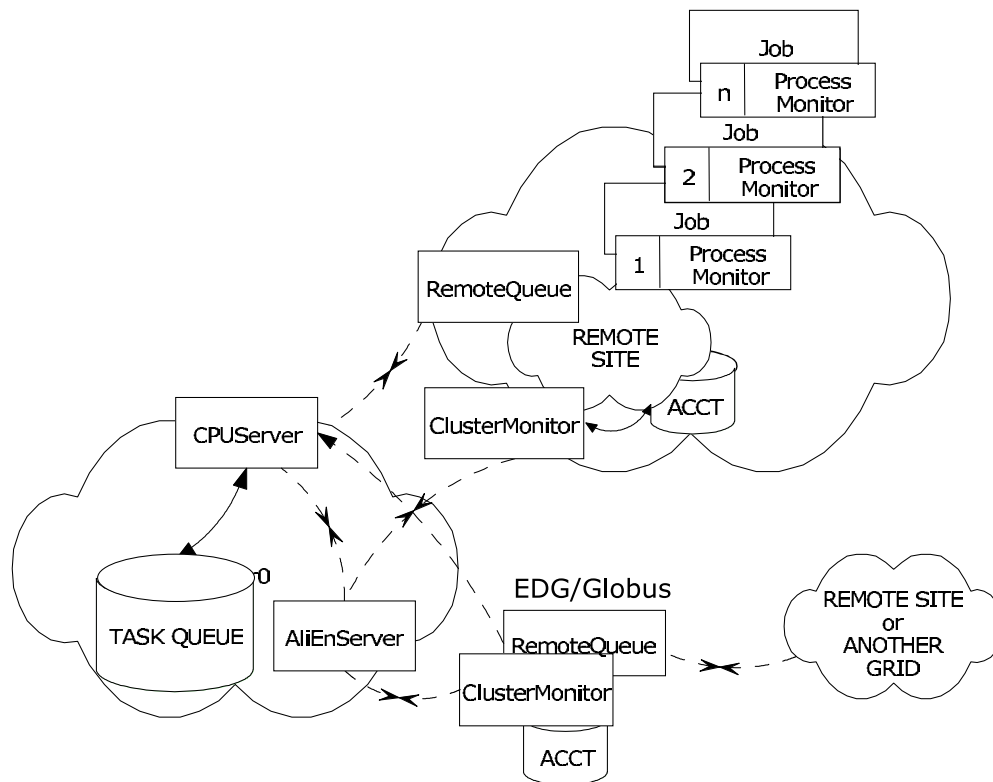


Figure 4.18. AliEn workload management.

The workload manager optimises the queue taking into account job requirements based on input files, CPU time, architecture, disk space, etc (see figure 4.18). It then makes jobs eligible to run on one or more computing elements. The active nodes then get jobs from the queue and start their execution. The queue system monitors the job progression and has access to the standard output and standard error.

To maintain compatibility with EDG–LCG Grid projects, AliEn user interface uses the Condor ClassAds [529] as a Job Description Language (JDL). The JDL defines the executable, its arguments, and the software packages or data that are required by the job. The Workload Management service can modify the job’s JDL entry by adding or elaborating requirements based on the detailed information it gets from the system like the exact location of the dataset and replicas, client and service capabilities.

Input and output associated with any job can be registered in the AliEn file catalogue, a virtual file system in which a logical name is assigned to a file (see figure 4.19). Unlike real file systems, the file catalogue does not own the files; it only keeps an association between the Logical File Name (LFN) and (possibly more than one) Physical File Names (PFN) on a real file or mass storage system. PFNs describe the physical location of the files and include the name of the AliEn storage element and the path to the local file.

The system supports file replication and caching and will use file location information when it comes to scheduling jobs for execution. The directories and files in the file catalogue have privileges for owner, group, and the world. This means that every user can have exclusive read and write privileges for his portion of the logical file namespace (home directory). In

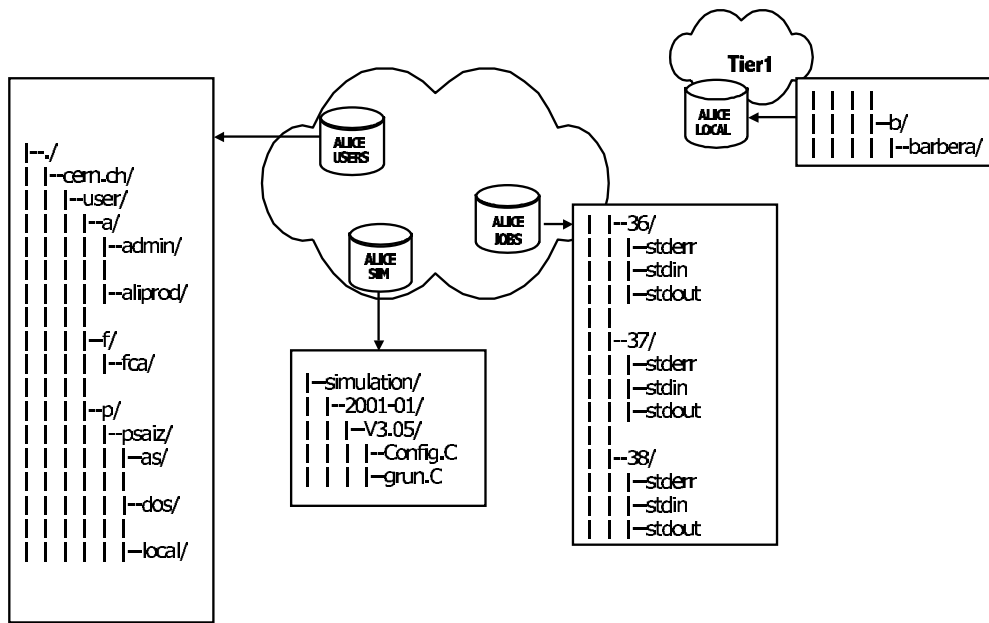


Figure 4.19. AliEn file catalogue.

order to address the problem of scalability, the AliEn file catalogue is designed to allow each directory node in the hierarchy to be supported by different database engines, possibly running on a different host and, in a future version, even having different internal table structures, optimized for a particular branch.

The file catalogue is not meant to support only regular files—we have extended the file system paradigm and included information about running processes in the system (in analogy with the `/proc` directory on Linux systems). Each job sent to AliEn for execution gets a unique id and a corresponding `/proc/id` directory where it can register temporary files, standard input and output, as well as all job products. In a typical production scenario, only after a separate process has verified the output, will the job products be renamed and registered in their final destination in the file catalogue. The entries (LFNs) in the AliEn file catalogue have an immutable unique file id attribute that is required to support long references (for instance in ROOT) and symbolic links.

The hierarchy of files and directories in the AliEn file catalogue reflects the structure of the underlying database tables. In the simplest and default case, a new table is associated with each directory. In analogy to a file system, the directory table can contain entries that represent the files or again subdirectories. With this internal structure, it is possible to attach to a given directory table an arbitrary number of additional tables, each one having a different structure and possibly different access rights while containing metadata information that further describes the content of files in a given directory. This scheme is highly granular and allows fine access control. Moreover, if similar files are always catalogued together in the same directory this can substantially reduce the amount of metadata that needs to be stored in the database. While having to search over a potentially large number of tables may seem ineffective, the overall search scope has been greatly reduced by using the file system hierarchy paradigm and, if data are sensibly clustered and directories are spread over multiple database servers, we could even execute searches in parallel and effectively gain performance while assuring scalability.




Figure 4.20. Federation of collaborating Grids.

Presently, the virtual file system can be seen via the command line, web interface and C++ API (in order to allow access to registered files from the AliRoot framework) as well as via the file system interface (`alienfs`) which plugs into the Linux kernel as a user space module and allows read–write access to the AliEn file catalogue as if it were a real file system. Because of the weak coupling between the resources and the resource brokers in the AliEn Grid model, it is possible to imagine a hierarchical Grid structure that spans multiple AliEn and ‘foreign’ Grids but also includes all resources under the direct control of top level virtual organization.

The connectivity lines in figure 4.20 represent the collaboration and trust relationships. In this picture the entire foreign Grid can be represented as a single computing and storage element (albeit a potentially powerful one). In this sense, we have constructed the AliEn–EDG interface and tested the interoperability. Along the same lines, the AliEn–AliEn interface allows the creation of the federation of collaborating Grids. The resources in this picture can be still shared between various top level virtual organizations according to the local site policy so that the Grid federations can overlap at resource level.

Currently, ALICE is using the system for distributed production of Monte Carlo data at over 30 sites on four continents (see figure 4.21 and figure V). The round of productions run during the last twelve months is aimed at providing data for this report. During this period more than 23 000 jobs have been successfully run under AliEn control worldwide producing 25 TB of data. Computing and storage resources are available in both Europe and the US. The amount of processing needed for a typical production is in excess of $30 \text{ MSI}2\text{k} \times \text{s}$ to simulate and digitize a central event. Some 10^3 events are generated for each major production. This is an average over a very large range since peripheral events may require one order of magnitude less CPU, and pp events two orders of magnitude less. Central events are then reprocessed several times superimposing known signals, in order to be reconstructed and analysed. Again there is a wide spread in the time this takes, depending on the event, but for a central event




Figure 4.21. Job distribution around participating sites during typical production round.

this needs a few $\text{MSI}2\text{k} \times \text{s}$. Each Pb–Pb central event occupies about 2 GB of disk space, while pp events are two orders of magnitude smaller.

Using AliEn we can solve today the ALICE simulation and reconstruction use cases as well as tackle the problem of distributed analysis on the Grid where we follow two approaches: the asynchronous (interactive batch approach) and the synchronous (true interactive) analysis model.

The asynchronous model can be realized by using AliEn as a Grid framework and by extending the ROOT functionality to make it Grid-aware. The `TAlien` class, based on the abstract `TGrid` class, implements the basic methods to connect and disconnect from the Grid environment and to browse the virtual file catalogue. `TAlien` uses the AliEn API for accessing and browsing the file catalogue. The files are handled in ROOT by the `TFile` class, which provides a plug-in mechanism supporting various file access protocols. The `TAlienFile` class inherits from the `TFile` class and provides additional file access protocols using the generic file access interface of the AliEn C++ API. Finally, the `TAlienAnalysis` class provides the analysis framework and overall steering as described in the next paragraph.

As the first step, the analysis framework has to extract a subset of the datasets from the virtual file catalogue using metadata conditions provided by the user. The next and the most difficult part is the splitting of the tasks according to the location of datasets. A trade-off has to be found between best use of available resources and minimal data movements. Ideally, jobs should be executed where the data are stored. Since one cannot expect a uniform storage location distribution for every subset of data, the analysis framework has to negotiate with dedicated Grid services the balancing between local data access and data replication. Once the distribution is decided, the analysis framework spawns sub-jobs. These are submitted to the AliEn workload management with precise job descriptions. The user can control the results during and after data are processed. The framework collects and merges available results from all terminated sub-jobs on request. An analysis object associated with the analysis task remains persistent in the Grid environment so the user can go offline and reload an analysis task at a later date, check the status, merge current results, or resubmit the same task with a modified analysis code.




Figure 4.22. Conventional setup of a SuperPROOF farm.

For implementation of the synchronous analysis model, we need to have much tighter integration between the ROOT and AliEn frameworks. This can be achieved by extending the functionality of PROOF [504]—the parallel ROOT facility—that covers the needs of interactive analysis in the local cluster environment. PROOF is part of the ROOT framework and provides a framework to use in parallel different computing resources while balancing dynamically their workload, with the goal of both optimizing the CPU exploitation and minimizing the data transfers over the network. Rather than transferring all the input files to a single execution node (farm), it is the program to be executed that is transferred to the nodes where the input is locally accessible and then run in parallel. This approach is simplified by the availability of the CINT interpreter, distributed with the ROOT package, which allows the execution of C++ scripts, compiling and linking them on the run. In this way, the amount of data to be moved is of the order of kB while the total size of input files for a typical analysis job can be as large as several tens of TB. The interface to Grid-like services is presently being developed, focusing on GLOBUS authentication and the use of file catalogues, in order to make both accessible from the ROOT shell.

In the conventional setup (see figure 4.22), PROOF slave servers are managed by a PROOF master server, which distributes tasks and collects results. In a multi-site setup each site running a PROOF environment will be seen as a SuperPROOF slave server for a SuperPROOF master server running on the user machine. The PROOF master server has therefore to implement the functionality of a PROOF master server and a SuperPROOF worker server at the same time. AliEn classes used for asynchronous analysis as described earlier can be used for task splitting in order to provide the input data sets for each site that runs PROOF locally.

The SuperPROOF master server assigns these datasets to the PROOF master server on individual sites. Since these datasets are locally readable by all PROOF servers, the PROOF master on each site can distribute the datasets to the PROOF slaves in the same way as for a conventional setup (see figure 4.23).

In a static scenario, each site maintains a static PROOF environment. To become a PROOF master or slave server a farm has to run a `proofd` process. These daemons are always running on dedicated machines on each site. To start a SuperPROOF session the SuperPROOF master

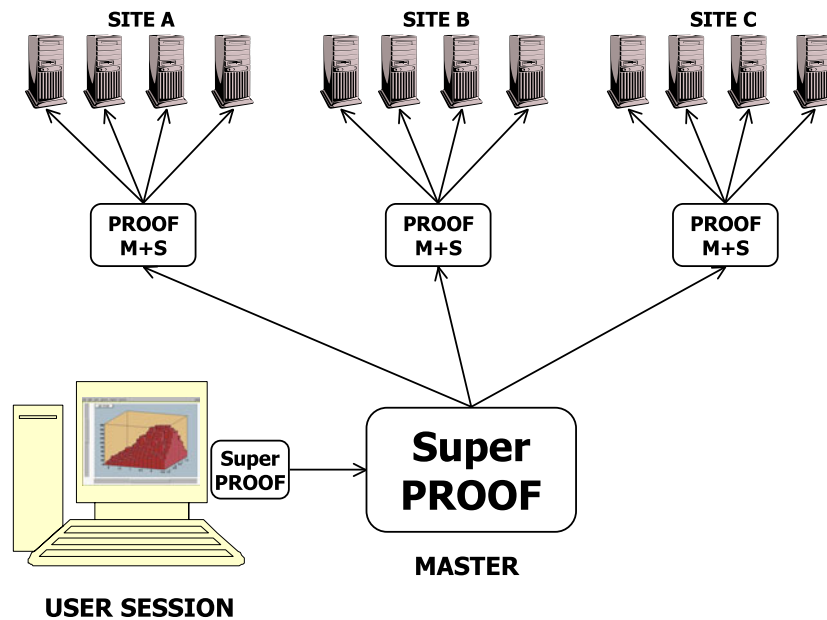


Figure 4.23. Setup of a multi-level distributed PROOF farm.

contacts each PROOF master on individual sites, which start PROOF workers on the configured nodes. The control of the `proofd` processes can be the responsibility of the local site or it can be done by a dedicated AliEn service. In a dynamic environment, once a SuperPROOF session is started, an AliEn Grid service starts `proofd` processes dynamically using dedicated queues in the site batch queue system. This ensures that a minimum set of `proofd` processes is always running. The system can react to an increasing number of requests for SuperPROOF sessions by starting a higher number of `proofd` processes.

The static environment makes sense for sites with large computing capabilities, where farm nodes can be dedicated exclusively to PROOF. If a site requires an efficient resource sharing, a dynamic environment appears to be the best choice, since it can use locally configured job queues to run `proofd` processes. These processes can be killed, if the requests decrease, or after a period of inactivity.

The SuperPROOF concept can have an important impact during the development of the analysis algorithms. The interactive response will simplify the work of the physicist when processing large data sets and it will result in a higher efficiency while prototyping the analysis code. On the other hand, the stateless layout of parallel batch analysis with ROOT will be the method of choice for processing complete datasets with final algorithms. Scheduled jobs allow for a more efficient use of resources since the system load is predictable from the JDLs of queued jobs. In general, the fault tolerance of a highly synchronized system like SuperPROOF with unpredictable user behaviour is expected to be lower than that of an asynchronous system. While the prototype of the analysis system based on AliEn already exists, the first prototype of SuperPROOF is foreseen for the end of year 2003.

4.1.4.1. Relations with the LHC computing Grid project. The LHC Computing Grid (LCG) project [531] was launched in March 2002 [532] following the recommendations of the LHC Computing Review [533]. Its objective is to provide the LHC experiments with the necessary

computing infrastructure to analyse their data. This should be achieved via the realization of common projects in the field of application software and Grid computing, deployed on a common distributed infrastructure. The project is divided into two phases, one phase of preparation and prototyping (2002–2005), and one of commissioning and exploitation of the computing infrastructure (2006–2008).

An overview Software and Computing Committee (SC2) organizes Requirement Technical Assessment Groups (RTAGs) that define requirements. These requirements are received by the Project Execution Board that organizes the project activity in several workpackages.

ALICE has been very active in the setting up of this structure and ALICE members have served as chairs for several RTAGs. ALICE has elaborated a complete solution for handling the data, and intends to continue developing it while collaborating fully with the other experiments in the framework of the LCG project. Of course ALICE will be happy to share its software with the other experiments and it will consider with interest the results of the LCG project.

In the application area the LCG project has decided to make ROOT one of the central elements of its development and ALICE fully supports this decision. However, the LCG project has decided not to base the application area architecture directly on ROOT but to establish with the ROOT team a user–provider relationship. ALICE has expressed some concerns with this approach that it sees as likely to generate ambiguities, duplication of work, and unnecessary complications for the end user. ALICE sees its role in the LCG project as one of close collaboration with the ROOT team in developing base technologies (e.g. geometrical modeler, virtual Monte Carlo) to be included directly in ROOT and to be made available to the other LHC experiments in a user–provider relation.

In the Grid technology area ALICE will continue using its AliEn Grid framework. As the middleware of the other Grid projects (PPDG, iVDGL, GryPhin, DataGrid, CrossGrid, etc) becomes available and stable, AliEn will be interfaced with it. If this interface provides a better and more reliable service compared with the native AliEn implementation of the same services, we shall migrate to it while preserving the AliEn user interface.

In the Grid deployment area we shall collaborate in the establishment of the so-called LCG-1 testbed and in its validation and exploitation with the Computing Data Challenges.

4.1.5. Software development environment. The ALICE software community is composed of small groups of 2–3 people who are geographically dispersed and who do not depend hierarchically on the Computing Coordinator. This affects not only the organization of the computing project but also the software development process itself as well as the tools that have to be employed to aid this development.

This situation is not specific to ALICE. It is similar to the other LHC collaborations and most modern HEP experiments. However, this is a novel situation. In the previous generation of experiments, during the LEP era, although physicists from different institutions developed the software, they did most of the work at CERN. The size of modern collaborations and software teams makes this model impracticable. The experiments' offline programs have to be developed by truly distributed teams. Traditional software engineering methods cannot cope with this reality. If we wanted to apply these, we would need more people, advanced video-conferencing, and frequent travel. However, we would still be in a non-standard situation since all participants in this effort belong to different institutions and none of them primarily develop software.

To cope with this unprecedented situation, we are elaborating a software process inspired by the most recent Software Engineering trends, in particular by the extreme programming. In a nutshell, traditional software engineering aims at reducing the occurrence of change via a complete specification of the requirements and a detailed design before the start of

development. After an attentive analysis of the conditions of software development in ALICE, we have concluded that this strategy has no chance of succeeding.

Change in our environment is unavoidable. Not only do our requirements change but also the necessity to use the most inexpensive hardware causes the implementation technology to evolve continuously as new products reach the market. Moreover, not all our requirements are clear at the start of the development since new requirements arise as our understanding of the problem matures. Change does not have to be averted in an effective development strategy but instead must be integrated in the development.

This calls for the adoption of a new strategy that starts from the same principles but arrives at entirely different conclusions. This inversion of the accepted truths comes from a radicalization of the principles on which traditional Software Engineering is built and it is the reason why the new trends go under the collective name of extreme programming. The main principles guiding ALICE software development are highlighted below.

Requirements are necessary to develop the code. ALICE users express their requirements continuously with the feedback they provide. The offline core team redirects its priorities according to the user feedback. In this way, we make sure that we are working on the highest priority item at every moment, maximizing the efficiency of our work by responding to user needs and expectations.

Design is good for evolving the code. AliRoot code is designed continuously since the feedback received from the users is folded into the continuous design activity of the core offline team. At any given time there is a design of the system considered as a short-term guide to the next development phase and not a long-term static objective. This change is not painfully accommodated but is the essence of the development activity. Static design is useless because it is immediately out of date. However, the design activity is essential since it rationalizes the development and indicates how best to include the features demanded by the users.

Testing is important for quality and robustness. The AliRoot code has a regression test that is run every night on the latest version of the code. While component testing is the responsibility of the groups providing modules to AliRoot, full integration testing is done nightly to make sure that the code is always functional. The results of the tests are reported on the Web and are publicly accessible [534]. Work is in progress to extend the percentage of AliRoot code covered by these tests. We shall also test the simulation and reconstruction algorithms by introducing a set of checkable items to confirm that the code still gives reasonable results.

Integration is needed to ensure coherence. The AliRoot code is collectively owned. We have a single code base handled via the CVS [535] concurrent development tool where all developers store their code, both the production version and the latest development code. For every module, one or two developers can update the new version. The different modules are largely independent and therefore, even if a faulty modification is stored, the other modules still work and the global activity is not stopped. Moreover, every night AliRoot is compiled and tested so that integration of all packages is continuous. Having all the code in a single repository allows a global vision of the code [536]. A hyperized version of the code is also extracted and created every night thanks to the ROOT documentation tools [537].

Discussions are valuable for users and developers. Discussion lists are commonplace in all computing projects. In ALICE, the discussion list is a valid place to express requirements and discuss design. We thus extend the discussions during our offline meetings to the discussion list.

It is not infrequent that important design decisions are taken on the basis of an email discussion thread. All interested ALICE users can then participate in the development of AliRoot and influence its design. Redefinition of planning and sharing of work also is very frequent, adding flexibility to the project and allowing it to address both short-term and long-term needs.

The above development strategy is very close to the one used by successful Open-Source (OS) projects such as Linux, GNU, ROOT, KDE, GSL and so on, making it easy for us to interact with them.

The ALICE offline code is composed of one single framework, OO in design and implemented in C++ where all the users, according to their ability and interest, can participate in the design and implementation. This requires a high degree of uniformity in the code as its structure has to be readable and easily understandable. We realized this very early on and decided to impose a limited set of coding and programming rules [538]. However, it was soon clear that only an automatic tool could verify the compliance of the code with these rules. Therefore we have set up a joint project with an Italian computer science research institute to develop such a tool [539]. This is now used nightly to check the code for compliance and a table of the violations in all modules is published on the Web [540].

The model chosen by ALICE implies that any design becomes quickly obsolete as development is driven by the most urgent needs of the user community. However, to avoid well-known design errors [541] it is important to have a clear view of the current structure of the code to start with. This can be achieved by reverse-engineering the code to produce Unified Modelling Language-like (UML [542]) diagrams. The code-checking tool has therefore been extended to provide reverse-engineering capability and produce active UML diagrams every night on the new code [543].

This development requires a special release policy which has been elaborated during the first two years of existence of AliRoot. It is based on the principle of fixed release dates with flexible scope. Given that the priority list is dynamically rearranged, it is difficult to define the scope of a given release in advance. Instead, we decided to schedule the release cycle in advance. The current branch of the CVS repository is 'tagged' every week with one major release every 6 months. The production branch of the CVS repository is tagged only for new check-ins due to bug fixes.

When a release is approaching, the date is never discussed, only the scope of what is to be included is tailored. Large developments are moved to the current branch and only developments that can be completed within the time remaining are included. The flexible priority scheduling ensures that if a feature is urgent, enough resources are devoted to it for it to be ready in the shortest possible time. If it does not make it for the current release, then it will be postponed till the next one. As soon as it is ready after the release, it will be made available via a CVS tag on the active branch.

We have tried to make installation as fast and reliable as possible. AliRoot is one single package that links only to ROOT. Thanks to this, we have not needed configuration management tools. We try to be as independent of the specific version of the operating system and compiler as possible. To ensure easy portability to any future platform, one of our coding rules states that the code must compile without warning on all supported platforms. At this time, these are HP-UX, Linux, DEC-Unix, and Solaris.

4.2. Monte Carlo generators for heavy-ion collisions

There exists a handful of programs simulating heavy-ion collisions at LHC energy. Most of them comply with the Open Standards for Cascades At RHIC (OSCAR) described in [544]. For the ALICE Technical Proposal, published in 1996, predictions from different generators

for Pb–Pb collisions at LHC energy, $\sqrt{s_{\text{NN}}} = 5.5$ TeV, were compared [545]. In particular, the predicted charged particle density at mid-rapidity varies strongly. The VENUS, HIJING 1.31, DPMJET-II and SFM generators gave $dN_{\text{ch}}/d\eta$ at $\eta = 0$ of 7000, 5200, 3700, 3400 and 1400, respectively. Now these generators have been updated and new ones produced. Here, we summarise the predictions for Pb–Pb central ($b < 3$ fm) collisions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV from HIJING, DPMJET and SFM.

4.2.1. HIJING and HIJING parametrization

4.2.1.1. HIJING. HIJING (Heavy-Ion Jet INteraction Generator) combines a QCD-inspired model of jet production [508] with the Lund model [546] for jet fragmentation. Hard or semi-hard parton scatterings with transverse momenta of a few GeV are expected to dominate high-energy heavy-ion collisions. The HIJING model has been developed with special emphasis on the role of mini jets in pp, pA and A–A reactions at collider energies.

Detailed systematic comparisons of HIJING results with a wide range of data demonstrates a qualitative understanding of the interplay between soft string dynamics and hard QCD interactions. In particular, HIJING reproduces many inclusive spectra, two-particle correlations, and the observed flavour and multiplicity dependence of the average transverse momentum.

The Lund FRITIOF [547] model and the Dual Parton Model [548] (DPM) have guided the formulation of HIJING for soft nucleus–nucleus reactions at intermediate energies, $\sqrt{s_{\text{NN}}} \approx 20$ GeV. The hadronic-collision model has been inspired by the successful implementation of perturbative QCD processes in PYTHIA [509]. Binary scattering with Glauber geometry for multiple interactions are used to extrapolate to pA and A–A collisions.

Two important features of HIJING are jet quenching and nuclear shadowing. Jet quenching is the energy loss by partons in nuclear matter. It is responsible for an increase of the particle multiplicity at central rapidities. Jet quenching is modelled by an assumed energy loss by partons traversing dense matter. A simple colour configuration is assumed for the multi-jet system and the Lund fragmentation model is used for the hadronization. HIJING does not simulate secondary interactions.

Shadowing describes the modification of the free nucleon parton density in the nucleus. At the low-momentum fractions, x , observed by collisions at the LHC, shadowing results in a decrease of the multiplicity. Parton shadowing is taken into account using a parametrization of the modification.

HIJING 1.36 [544] was used to produce central Pb–Pb events ($b < 3$ fm) with and without jet quenching at $\sqrt{s_{\text{NN}}} = 5.5$ TeV. In figure 4.24, pseudo-rapidity density ($dN/d\eta$) and transverse-momentum distributions of charged particles as well as the net baryon pseudo-rapidity-density distribution are shown. Taking into account that jet quenching leads to a factor of two increase in multiplicity at mid-rapidity and a softer transverse-momentum spectrum. Moreover, the pseudo-rapidity-density distribution shows a bump in the range $|\eta| < 3$. This bump is specific to HIJING.

In figure 4.25, the charged particle pseudo-rapidity distribution obtained with default HIJING parameters is shown (solid line). The prediction is compared to a distribution obtained by a modification of the quenching parameters inspired by recent RHIC results [549], to a log E -dependent energy loss and to HIJING without quenching. The LHC distributions are also compared to the corresponding predictions for RHIC. The new parameters decrease the charged multiplicity at $\eta = 0$ by about 25%, i.e. to 4500 compared to the predictions of HIJING 1.36.

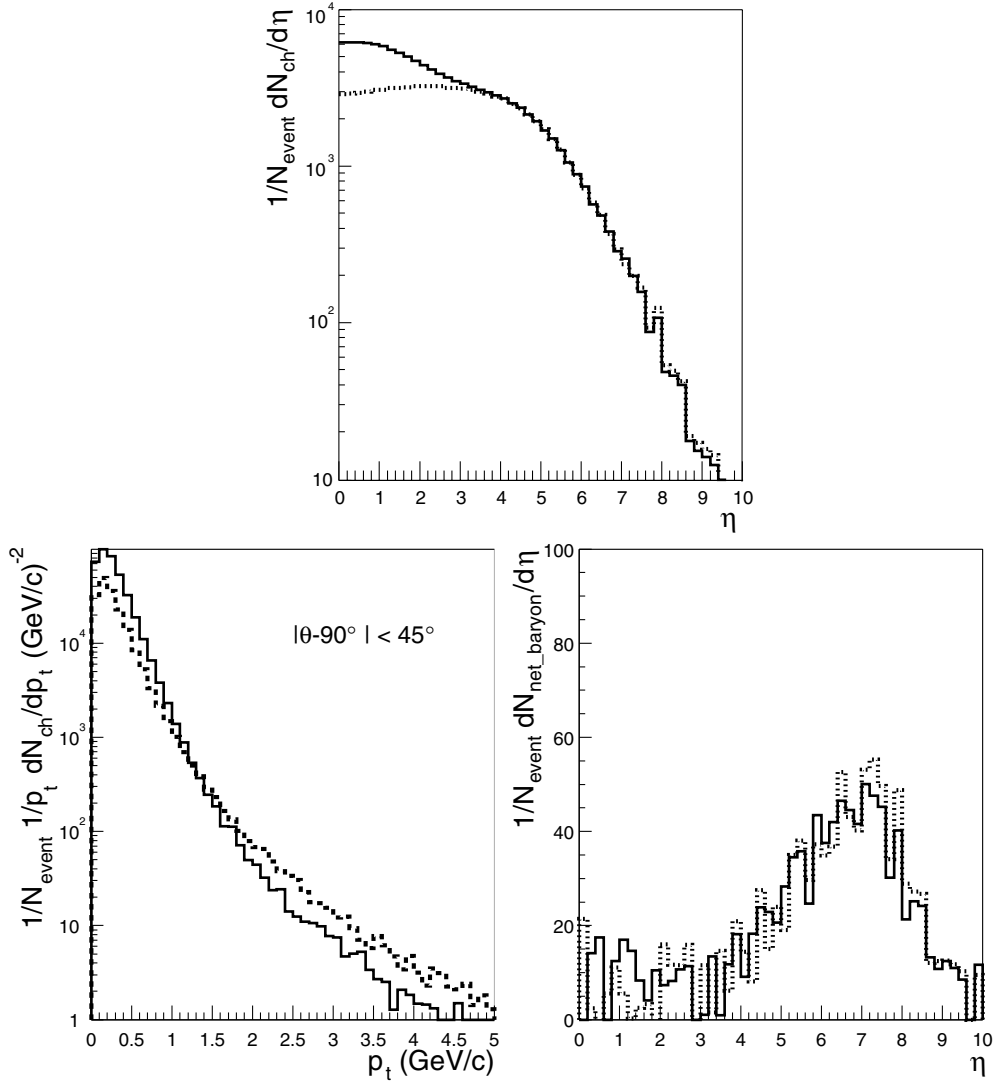


Figure 4.24. HIJING 1.36 predictions for Pb–Pb central, $b < 3$ fm, events at $\sqrt{s_{NN}} = 5.5$ TeV. The solid and dashed lines are the results for simulations with and without jet quenching, respectively.

4.2.1.2. HIJING parametrization in AliRoot. AliGenHIJINGparam [550] is an internal AliRoot [502] generator based on parametrized pseudo-rapidity density and transverse-momentum distributions of charged and neutral pions and kaons. The pseudo-rapidity distribution was obtained from a HIJING simulation of central Pb–Pb collisions and scaled to a charged-particle multiplicity of 8000 in the pseudo-rapidity interval $|\eta| < 0.5$. Note that this is about 10% higher than the corresponding value for a rapidity density with an average dN/dy of 8000 in the interval $|y| < 0.5$.

The transverse-momentum distribution is parametrized from the measured CDF pion p_t -distribution at $\sqrt{s} = 1.8$ TeV. The corresponding kaon p_t -distribution was obtained from the pion distribution by m_t -scaling. See [550] for the details of these parametrizations.




Figure 4.25. HIJING prediction for the η -distributions of charged particles for different quenching scenarios.

4.2.2. Dual-Parton Model (DPM). DPMJET is an implementation of the two-component Dual Parton Model (DPM) for the description of interactions involving nuclei based on the Glauber–Gribov approach. DPMJET treats soft and hard scattering processes in a unified way. Soft processes are parametrized according to Regge phenomenology while lowest-order perturbative QCD is used to simulate the hard component. Multiple-parton interactions in each individual hadron (nucleon, photon)–nucleon interaction are described by the PHOJET event generator. The fragmentation of parton configurations is treated by the Lund model PYTHIA.

Particle production in the fragmentation regions of the participating nuclei is described by a formation zone suppressed intra-nuclear cascade followed by Monte Carlo of evaporation processes of light nucleons and nuclei, high-energy fission, spectator fragmentation (limited to light-spectator nuclei), and de-excitation of residual nuclei by photon emission.

The most important features of DPMJET-II.5 [552] are new diagrams contributing to baryon stopping and a better calculation of Glauber cross sections. A new striking feature of hadron production in nuclear collisions is the large stopping of the participating nucleons in hadron–nucleus and nucleus–nucleus collisions compared to hadron–hadron collisions [553, 554]. The popcorn mechanism implemented in models with independent string fragmentation, like the DPM, is insufficient to explain this enhanced baryon stopping. New DPM diagrams were proposed by Kharzeev [555], Capella and Kopeliovich [556]. DPMJET-II.5 implements these diagrams and obtains an improved agreement with the net-baryon distributions in nuclear collisions.

For some light nuclei, the Glauber model calculations have been modified and the Woods–Saxon nuclear densities are replaced by parametrizations that agree better with the data. Measured nuclear radii are used [557] instead of the nuclear radii given by the previous parametrization. The new option XSECNUC is added to calculate a table of total, elastic, quasi-elastic, and perturbative QCD cross sections using a modified version of the routine XSGLAU adopted from DTUNUC-II [558, 559]. These changes lead to improved agreement with measured nuclear cross sections, such as the p–Air cross sections from cosmic-ray data.

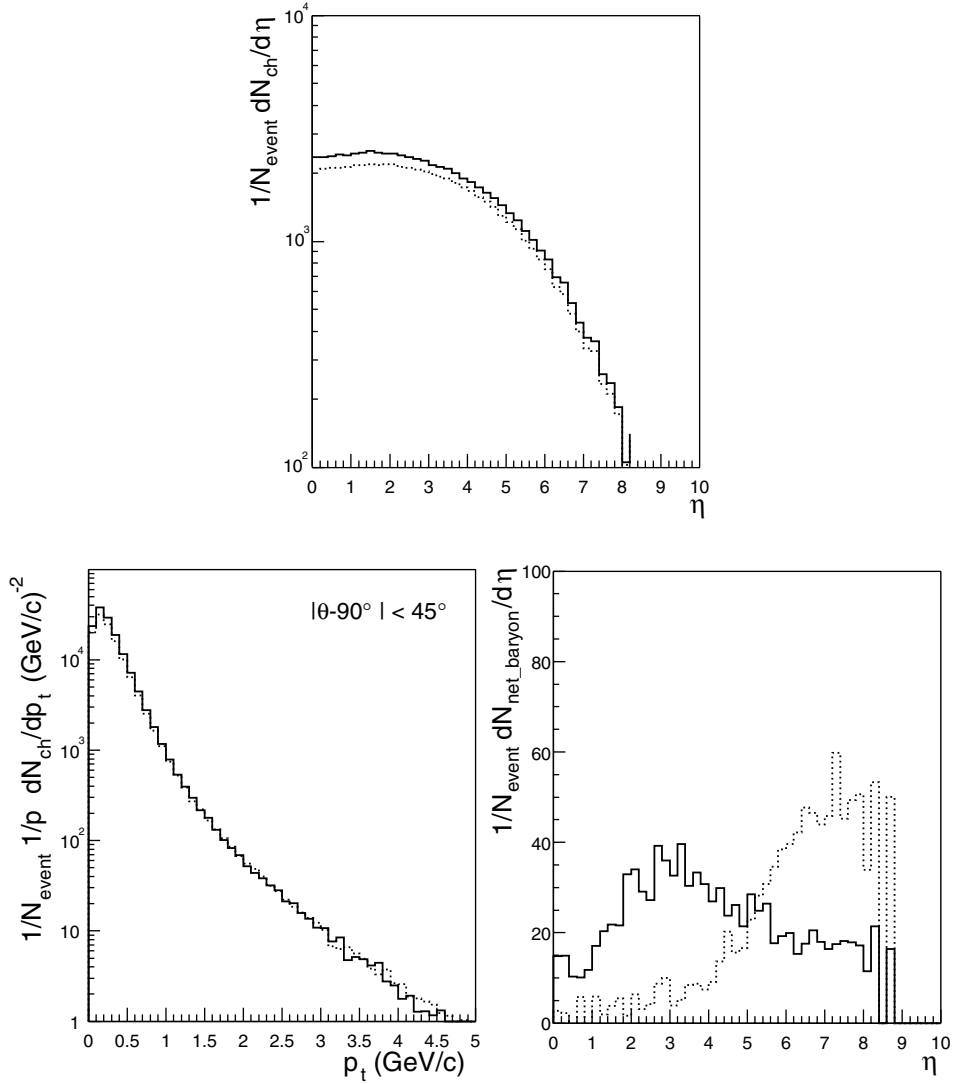


Figure 4.26. DPMJET-II.5 predictions for Pb–Pb central, $b < 3$ fm, events at $\sqrt{s_{\text{NN}}} = 5.5$ TeV. The full and dashed lines are for simulations with and without baryon stopping, respectively.

DPMJET-II.3/4 was extended to higher energies by calculating mini-jet production using as a default the GRVLO94 parton distributions [560]. They are replaced in DPMJET-II.5 by the more recent GRVLO98 [561], which describe the most recent HERA data.

The present version of the model DPMJET-II.5 allows the inclusion or exclusion of single-diffractive events in nucleus–nucleus collisions or the sampling of only single-diffractive events. The diffractive cross section in hadron–nucleus and nucleus–nucleus collisions is calculated by Monte Carlo.

The DPMJET [551] generator version II.5 [552] was used to simulate central, $b < 3$ fm, Pb–Pb events with and without baryon stopping at $\sqrt{s_{\text{NN}}} = 5.5$ TeV. Predictions with and without stopping are shown in figure 4.26. The baryon stopping mechanism increases the multiplicity by $\simeq 15\%$, and changes the net-baryon distribution.

Table 4.1. Charged-particle multiplicity predictions of different event generators.

Generator	Comments	$dN_{\text{ch}}/d\eta$ at $\eta = 0$	N_{ch} in acceptance
HIJING 1.36	With quenching	$\simeq 6200$	$\simeq 10800$
	Without quenching	$\simeq 2900$	$\simeq 5200$
DPMJET-II.5	With baryon stopping	$\simeq 2300$	$\simeq 4000$
	Without baryon stopping	$\simeq 2000$	$\simeq 3500$
SFM	With fusion	$\simeq 2700$	$\simeq 4700$
	Without fusion	$\simeq 3100$	$\simeq 5500$

4.2.3. *String-Fusion Model (SFM).* The features of the String-Fusion Model (SFM) code are [562, 563]:

- The soft interactions are described by the Gribov–Regge theory of multi-pomeron exchange. The hard part of the interactions, included as a new component of the eikonal model, begins to be significant above 50 GeV. The hard part of the interaction is simulated by PYTHIA and the strings formed by gluon splitting are fragmented with JETSET.
- Fusion of soft strings is included. Fragmentation is through the Artru–Mennessier string-decay algorithm.
- Rescattering of produced particles is included. Four different reactions are considered, two annihilation and two exchange reactions:
 - $q\bar{q} \rightarrow q'\bar{q}'$
 - $q\bar{q} \rightarrow s\bar{s}$
 - baryon exchange
 - strangeness exchange.

The latest version of the SFM code, PSM-1.0 [562], was used to generate central, $b < 3$ fm, Pb–Pb events with and without string fusion. Although rescattering between secondaries and spectators is included in the model as an option, these simulations did not take into account rescattering. Predictions at $\sqrt{s_{\text{NN}}} = 5.5$ TeV with and without string fusion are shown in figure 4.27. Including string fusion reduces the multiplicity by $\simeq 10\%$. The p_t -distribution is not strongly affected by the fusion.

4.2.4. *Comparison of results.* Results for charged-particle multiplicities for all three generators are presented in table 4.1. Large differences for $dN_{\text{ch}}/d\eta$ still exist. None of the current event generators reproduce the low multiplicities obtained from extrapolations of RHIC data discussed in section 1.

Detailed comparisons of m_t -spectra for different species predicted by HIJING, DPMJET and SFM codes with default parameters (with quenching for HIJING, with baryon stopping for DPMJET and with string fusion for SFM) are presented in figure 4.28.

In table 4.2 we give the inverse slope T obtained by fitting the transverse mass distribution in the range $0.1 \leq m_t - m_0 \leq 1$ GeV with the function

$$\frac{1}{m_t} \frac{dN}{dm_t} = A \exp\left(-\frac{m_t}{T}\right).$$

The slopes vary a lot, up to about 80%. As observed in data [564], the codes reproduce the increase of the inverse slope with particle mass except for DPMJET-II.5 where $T_{\text{nucleons}} < T_{\text{kaons}}$.




Figure 4.27. SFM predictions for Pb–Pb central, $b < 3$ fm, events at $\sqrt{s_{NN}} = 5.5$ TeV without rescattering. The full and dashed lines are for simulations with and without fusion, respectively.

4.2.5. First results from RHIC. Figure 4.29 summarizes the first results obtained by the PHENIX collaboration at RHIC [565]. Protons and antiprotons become as abundant as pions at p_t above 2 GeV. This effect is not reproduced by any of the generators. The corresponding predictions for the LHC are presented in figure 4.30 for HIJING, DPMJET, and SFM, respectively.

Different theoretical explanations have been proposed to describe such behaviour. It has been suggested that hydrodynamic expansion alone or combined with hadronic rescattering is responsible for baryon dominance at high p_t [566–569]. Protons and antiprotons produced via a baryon-junction mechanism combined with jet quenching for pions exhibit the same

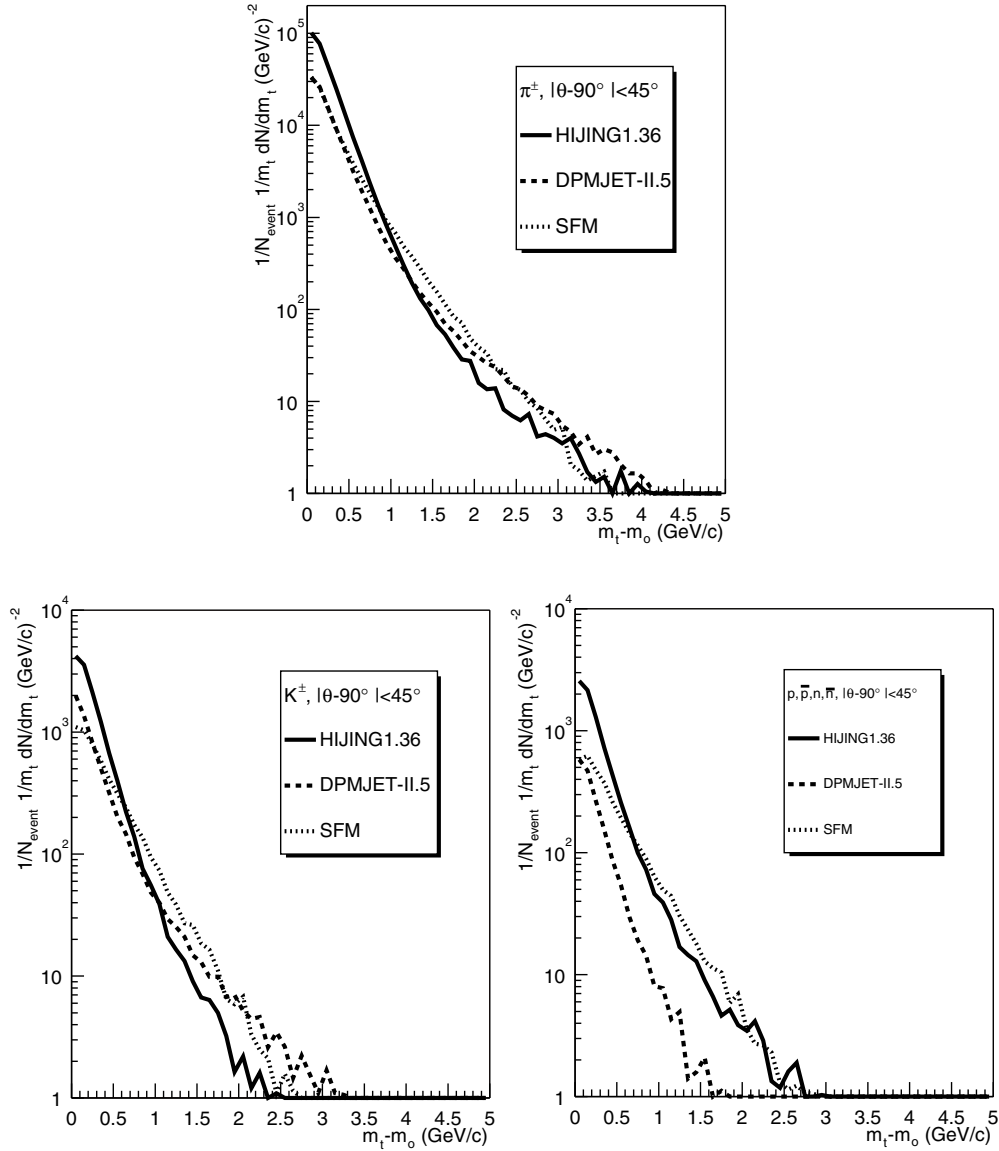



Figure 4.28. Charged-pion, charged-kaon and nucleon transverse-mass distributions in the barrel, $|\theta - 90^\circ| < 45^\circ$, for central Pb–Pb events at $\sqrt{s_{NN}} = 5.5$ TeV.

effect [570]. Intrinsic p_t -broadening in the partonic phase caused by the gluon saturation expected in high-density QCD [571] is an alternative explanation.

4.2.6. Conclusions. Most popular event-generators for ultra-relativistic heavy-ion collisions have been tested at the LHC energy scale. HIJING, DPMJET, and SFM central Pb–Pb events have been simulated at the energy $\sqrt{s_{NN}} = 5.5$ TeV. The charged particle density, $dN_{ch}/d\eta$, predicted by these generators varies widely, from about 2000 (SFM) to about 6000 (HIJING). Moreover, even in the framework of a single generator, the density is strongly model dependent, about 10% for SFM with and without string fusion to about 100% for HIJING with and

Table 4.2. Inverse slope in MeV of particles produced in the barrel, $|\theta-90^\circ| < 45^\circ$, for central Pb–Pb events at $\sqrt{s_{NN}} = 5.5$ TeV. The values in parantheses were previously reported in [545].

Particle	HIJING1.36	DPMJET-II.5	SFM
π^\pm	173 ± 1 (173)	196 ± 1 (209)	224 ± 1 (217)
K^\pm	184 ± 2 (184)	226 ± 5 (270)	336 ± 8 (198)
p, \bar{p}, n, \bar{n}	196 ± 3 (193)	187 ± 7 (241)	350 ± 11 (168)

**Figure 4.29.** Transverse-momentum spectra at mid-rapidity measured by PHENIX for π^+ , K^+ , p (left) and π^- , K^- , \bar{p} (right) for three different centrality selections, indicated in each panel.

without quenching. Predictions of transverse-mass spectra are also very different. In the range $m_t - m_0 < 1$ GeV the inverse slope varies strongly, up to about 80%.

The first RHIC data have also revealed a new effect which is not reproduced by the generators. In central collisions p_t at about 2 GeV, the p and \bar{p} yields are comparable to the pion ones.

4.2.7. Generators for heavy-flavour production. The Monte Carlo generators of nucleus–nucleus interactions described so far have not yet been tested extensively for heavy-flavour production (charm and beauty). On the other hand, several Monte Carlo models of hadron–hadron collisions have been used to describe heavy-flavour production at the LHC and have been compared to lower energy data, see [572].

The heavy-flavour simulations presented in this report, unless stated differently, have been performed using PYTHIA 6.150 [509] to simulate the elementary nucleon–nucleon collisions (section 4.3.1.2) and the results have been scaled to Pb–Pb interactions by the number of binary collisions. The number of collision has been calculated in the Glauber model. The effect of the modification of the parton distribution functions in the nucleus and of the nuclear broadening

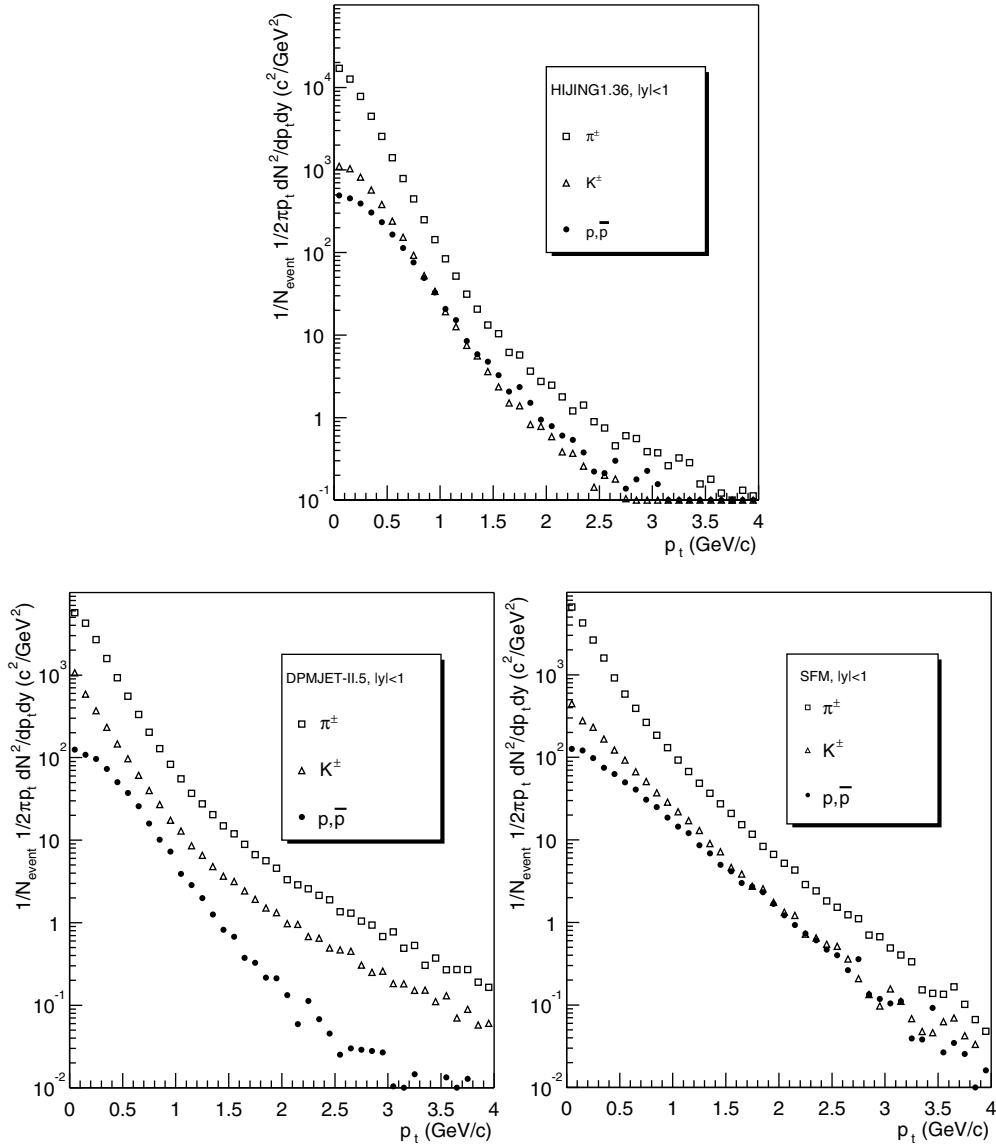


Figure 4.30. Transverse-momentum distributions predicted by HIJING, DPMJET and SFM for central Pb–Pb events at $\sqrt{s_{NN}} = 5.5$ TeV.

of the parton intrinsic p_t have been taken into account. More details on the extrapolation to Pb–Pb collisions will be given in section 6 of the PPR Volume II.

4.2.8. Other generators.

4.2.8.1. MEVSIM. MEVSIM [573] was developed for the STAR experiment to quickly produce a large number of A–A collisions for some specific needs—initially for HBT studies and for testing of reconstruction and analysis software. However, since the user is able to

generate specific signals, it was extended to flow and event-by-event fluctuation analysis. A detailed description of MEVSIM can be found in [573].

MEVSIM generates particle spectra according to a momentum model chosen by the user. The main input parameters are types and numbers of generated particles, momentum-distribution model, reaction-plane and azimuthal-anisotropy coefficients, multiplicity fluctuation, number of generated events, etc. The momentum models include factorized p_t and rapidity distributions, non-expanding and expanding thermal sources, arbitrary distributions in y and p_t and others. The reaction plane and azimuthal anisotropy is defined by the Fourier coefficients (maximum of six) including directed and elliptical flow. Resonance production can also be introduced.

The user can define the acceptance intervals, avoiding the generation of particles that do not enter the detector. The generation time is negligible compared to the analysis time.

MEVSIM was originally written in FORTRAN. It was integrated into AliRoot via the TMeVSim package [576] in order to maintain it as a common code with the STAR offline project. A complete description of the AliRoot implementation of MEVSIM can be found elsewhere [574].

4.2.8.2. GeVSim. GeVSim [575] is a fast and easy-to-use Monte Carlo event generator implemented in AliRoot. It can provide events of similar type configurable by the user according to the specific needs of a simulation project, in particular, that of flow and event-by-event fluctuation studies. It was developed to facilitate detector performance studies and for the testing of algorithms. GeVSim can also be used to generate signal-free events to be processed by afterburners, for example the HBT processor.

GeVSim is based on the MEVSIM [573] event generator developed for the STAR experiment and written in FORTRAN. Since the interfacing of FORTRAN code makes the design complicated and the maintenance more difficult, a new package, GeVSim, written purely in C++, was developed.

GeVSim generates a list of particles by randomly sampling a distribution function. The parameters of single-particle spectra and their event-by-event fluctuations are explicitly defined by the user. Single-particle transverse-momentum and rapidity spectra can either be selected from a menu of four predefined distributions, the same as in MEVSIM, or provided by the user.

Flow can be easily introduced into simulated events. The parameters of the flow are defined separately for each particle type and can either be set to a constant value or parametrized as a function of transverse momentum and rapidity. Two parametrizations of elliptic flow based on results obtained by RHIC experiments are provided.

GeVSim also has extended possibilities for the simulation of event-by-event fluctuations. The model allows fluctuations following an arbitrary analytically defined distribution in addition to the Gaussian distribution provided by MEVSIM. It is also possible to systematically alter a given parameter to scan the parameter space in one run. This feature is useful when analysing performance with respect to, for example, multiplicity or event-plane angle.

The current status and further development of the GeVSim code and documentation can be found in [574].

4.2.8.3. HBT processor. Correlation functions constructed with the data produced by MEVSIM or any other event generator are normally flat in the region of small relative momenta. The HBT-processor afterburner introduces two particle correlations to the set of generated particles. It shifts the momentum of each particle so that the correlation function

of a selected model is reproduced. The imposed correlation effects due to Quantum Statistics (QS) and Coulomb Final State Interactions (FSI) do not affect the single-particle distributions and multiplicities. The event structures before and after passing through the HBT processor are identical. Thus, the event reconstruction procedure with and without correlations is also identical. However, the tracking efficiency, momentum resolution and particle identification need not to be, since correlated particles have a special topology at small relative velocities. We can thus verify the influence of various experimental factors on the correlation functions.

The data structure of reconstructed events with correlations is the same as the structure of real events. Therefore, the specific correlation-analysis software developed at the simulation stage can be directly applied to real data.

The method, proposed by Ray and Hoffmann [577] is based on random shifts of the particle three-momentum within a confined range. After each shift, a comparison is made with correlation functions resulting from the assumed model of the space-time distribution and with the single-particle spectra which should remain unchanged. The shift is kept if the χ^2 -test shows better agreement. The process is iterated until satisfactory agreement is achieved. In order to construct the correlation function, a reference sample is made by mixing particles from some consecutive events. Such a method has an important impact on the simulations when at least two events must be processed simultaneously.

Some specific features of this approach are important for practical use:

- The HBT processor can simultaneously generate correlations of up to two particle types (e.g. positive and negative pions). Correlations of other particles can be added subsequently.
- The form of the correlation function has to be parametrized analytically. One and three dimensional parametrizations are possible.
- A static source is usually assumed. Dynamical effects, related to expansion or flow, can be simulated in a stepwise form by repeating simulations for different values of the space-time parameters associated with different kinematic intervals.
- Coulomb effects may be introduced by one of the three approaches: Gamow factor, experimentally modified Gamow correction and integrated Coulomb wave functions for discrete values of the source radii.
- Strong interactions are not implemented.

The HBT processor is written in FORTRAN and integrated in AliRoot via a C++ interface. The detailed description can be found in [578].

In figure 4.31 are presented (a) the single-particle p_x distributions and (b) the Q_{inv} correlation function before and after applying the HBT processor. Pions were generated by MEVSIM and the HBT processor was then applied.

The shape of the single-particle spectra is unaffected by the shifts of particle momenta. The correlation function, obtained by mixing of particles from the same and different events, is flat when particles are taken before passing through the HBT processor. The correlation effect is clearly visible for particles taken after HBT processor.

4.2.8.4. Flow afterburner. Azimuthal anisotropies, especially elliptic flow, carry unique information about collective phenomena and consequently are important for the study of heavy-ion collisions. Additional information can be obtained by studying different heavy-ion observables, especially jets, relative to the event plane. Therefore it is necessary to evaluate the capability of ALICE to reconstruct the event plane and study elliptic flow.

Since there is not a well-understood microscopic description of the flow effect it cannot be correctly simulated by microscopic event generators. Therefore, to generate events with

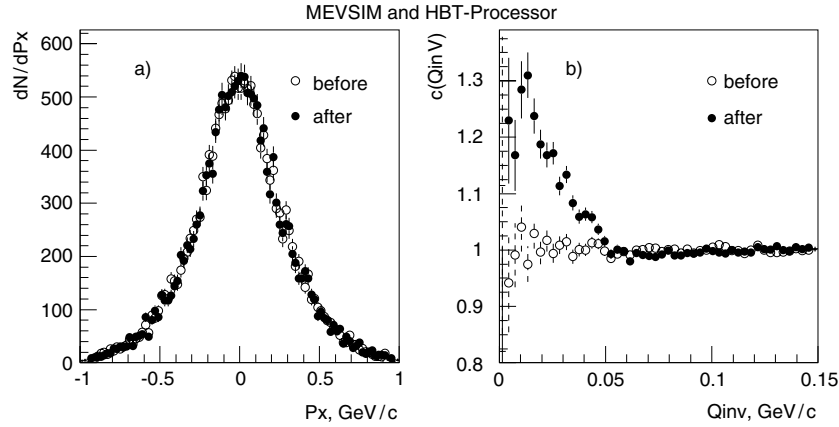


Figure 4.31. Example of HBT processor output: (a) pion single-particle distributions, and (b) correlation functions before and after passing through the HBT processor.

flow one has to use event generators based on macroscopic models, like GeVSim [575] or an afterburner which can generate flow on top of events generated by event generators based on the microscopic description of the interaction. In the AliRoot framework such a flow afterburner is implemented.

The algorithm to apply azimuthal correlation consists in shifting the azimuthal coordinates of the particles. The transformation is given by [579]

$$\varphi \rightarrow \varphi' = \varphi + \Delta\varphi, \quad \Delta\varphi = \sum_n \frac{-2}{n} v_n(p_t, y) \sin n \times (\varphi - \psi),$$

where $v_n(p_t, y)$ is the flow coefficient to be obtained, n the harmonic number and ψ the event-plane angle. Note that the algorithm is deterministic and does not contain any random number generation.

The value of the flow coefficient can be either constant or parametrized as a function of transverse momentum and rapidity. Two parametrizations of elliptic flow are provided as in GeVSim.

4.2.8.5. Generator for e^+e^- pairs in Pb–Pb collisions. In addition to strong interactions of heavy ions in central and peripheral collisions, ultra-peripheral collisions of ions give rise to coherent, mainly electromagnetic, interactions among which the dominant process is the (multiple) e^+e^- -pair production [580]

$$AA \rightarrow AA + n(e^+e^-), \quad (4.1)$$

where n is the pair multiplicity. Most electron–positron pairs are produced in the very forward direction escaping the experiment. However, for Pb–Pb collisions at the LHC, the cross section of this process, about 230 kb, is enormous. A sizeable fraction of pairs produced with large-momentum transfer can contribute to the hit rate in the forward detectors thus increasing the occupancy or trigger rate. In order to study this effect an event generator for e^+e^- -pair production has been implemented in the AliRoot framework [581]. The class `TEpEmGen` is a realization of the `TGenerator` interface for external generators and wraps the FORTRAN code used to calculate the differential cross section. `AliGenEpEmv1` derives from `AliGenerator` and uses the external generator to put the pairs on the AliRoot particle stack.

4.3. Technical aspects of pp simulation

There are many programs to simulate high-energy pp collisions. In order to study the global-event properties and the heavy-flavour production in pp collisions at ALICE, some simulations have been performed with version 6.150 of the PYTHIA event generator [509]. Version 6.1 of HERWIG [582] has been used at the generator level for a first comparison to PYTHIA. Other Monte Carlo generators (ISAJET [583] and PHOJET [584]) have already been interfaced to AliRoot but they were not used for the studies reported here.

4.3.1. PYTHIA. PYTHIA [509] can generate different types of high-energy physics events. It is based on perturbative QCD but also includes models of soft interactions, parton showers and multiple interactions, fragmentation and decays.

Owing to the composite nature of hadrons, several interactions between parton pairs are expected to occur in a typical hadron–hadron collision [585]. Evidence for these interactions has mounted including their direct observation by CDF [586]. However, the understanding of multiple interactions is still primitive. Therefore, PYTHIA contains four different models as options available to users. The main parameter in these models is p_t^{min} , a cut-off introduced to regularize the dominant $2 \rightarrow 2$ QCD cross sections, which diverge as $p_t \rightarrow 0$ and drop rapidly at large p_t . Apart from the default Model 1, all other models assume a continuous turn-off of the cross section at p_t^{min} . Models 3 and 4, originally developed to fit the UA5 data, assume a varying impact parameter between the colliding hadrons. The hadronic matter distributions are assumed to have Gaussian (Model 3) or a double-Gaussian (Model 4) shapes. Several studies have concluded that such models provide a better description of the UA5 multiplicity distributions than Model 1 [587, 588], The Gaussian-type models also better describe the underlying production in beauty events at the Tevatron [589].

To compare PYTHIA to data, p_t^{min} must be adjusted in each case to reproduce the mean charged-particle multiplicity. The result depends strongly on \sqrt{s} , the multiple interaction model and the parton distribution functions. These effects on p_t^{min} have been investigated in [588]. In each case, p_t^{min} appears to increase monotonically with \sqrt{s} , which must be taken into account when extrapolating to LHC energies. A power-law, $p_t^{\text{min}}(\sqrt{s}) = p_t^0(\sqrt{s}/E_0)^{2\varepsilon}$ [588], has been used to fit the tuned p_t^{min} values as a function of \sqrt{s} (where p_t^0 is the p_t^{min} value at a given reference energy E_0). The same functional form of the energy dependence is implemented in PYTHIA, allowing users to select the parameter values. However, in the most recent version of the model (PYTHIA 6.210) the default values of the parameters have been already tuned in order to reproduce all the available collider data.

Figure 4.32 shows the PYTHIA predictions of the charged particle density $dN_{\text{ch}}/d\eta$ at $\eta = 0$ as a function of \sqrt{s} per inelastic pp collisions, compared to UA5 and CDF data. Two different versions of PYTHIA have been used, in association with different sets of parton distributions and of cut-off parameters. PYTHIA version 6.150 (that is implemented as default in the AliRoot framework) has been used with the CTEQ4L parton distributions and Model 3 for multiple interactions with $p_t^0 = 3.47$ GeV at the LHC energy $E_0 = 14$ TeV and $\varepsilon = 0.087$ [588]. PYTHIA version 6.210 has been used instead with CTEQ5L as parton distributions, the Model 4 (with its default settings) for multiple interactions, and with the default values for the parameters that regulate the cut-off energy dependence. From this comparison one can conclude that the PYTHIA 6.210—Model 4 gives the best agreement with the experimental data, up to Tevatron energies, and it is probably the best candidate for a massive production of a sample of simulated events at LHC energies. However, a first round of simulations including the full-detector response modeled in AliRoot, has been performed with the old settings. Results of these simulations will be widely discussed in section 6 of the PPR Volume II.




Figure 4.32. Charged-particle density $dN_{\text{ch}}/d\eta$ at $\eta = 0$ as a function of \sqrt{s} in non-single diffractive inelastic pp collisions. PYTHIA predictions, with two different settings of the input parameters, are compared to the UA5 and CDF data (the fit of those data, performed by CDF Collaboration, is also superimposed).




Figure 4.33. PYTHIA charged-particle pseudo-rapidity (left) and transverse-momentum (right) distributions from inelastic pp collisions. Solid line: PYTHIA 6.150—Model 3; dashed line: PYTHIA 6.210—Model 4; dotted line: PYTHIA 6.210—Model 1.

4.3.1.1. Minimum-bias events. The charged-particle pseudo-rapidity and transverse-momentum distributions produced in inelastic pp collisions at $\sqrt{s} = 14$ TeV, including diffractive events, simulated with three different versions of PYTHIA code and input parameters, are shown in figure 4.33. In addition to the PYTHIA settings described above, Model 1 has also been added for comparison.

The charged-particle density at $\eta = 0$ estimated with PYTHIA 6.210—Model 4 is about 6.7; the values estimated with PYTHIA 6.210—Model 1 and PYTHIA 6.150—Model 3 are




Figure 4.34. Multiplicity distributions for minimum-bias pp collisions at $\sqrt{s} = 14$ TeV (see captions of figure 4.35 for explanation of different line styles).

respectively 7.6 and 5.8. From the fit of the experimental data, the expected value would be 6.1. As regards the transverse-momentum spectra we do not observe large discrepancies between the different PYTHIA versions. However, large differences arise when looking at the multiplicity distributions, as shown in figure 4.34.

4.3.1.2. Heavy-flavour production. PYTHIA has been used to simulate charm and beauty production from fixed-target to collider energies [590]. Heavy-flavour production is based on the concept of factorization between two distinct phases: the perturbative description of hard production and the hadronization phase based on the Lund string model [546].

The hard-production model is exact at leading-order. Heavy-flavour total and differential cross sections can be calculated exactly to next-to-leading order in perturbative QCD (see section 6 in the PPR Volume II for details). Since these calculations are not well suited for event generators, higher order contributions are instead included in PYTHIA in the parton-shower approach [509]. This model is not exact even at next-to-leading order but it catches the leading-log aspects of the multiple-parton-emission phenomenon [590].

In PYTHIA, the next-to-leading order processes such as flavour excitation, $qQ \rightarrow qQ$ and $gQ \rightarrow gQ$, and gluon splitting, $g \rightarrow Q\bar{Q}$, are calculated using massless matrix elements. Consequently, these cross sections diverge as the cut-off, p_t^{hard} , approaches zero. These divergences are regularized by a lower bound on p_t^{hard} . The value of the minimum p_t^{hard} strongly influences the cross section at low p_t , a focus for ALICE physics, and covered by the ALICE acceptance. Our approach has been to tune PYTHIA to reproduce as well as possible the next-to-leading order predictions. Section 6 of the PPR Volume II will describe the details of this tuning and demonstrates that reasonable agreement can be found despite the fundamental differences between the two calculations.

4.3.2. HERWIG. HERWIG is a general-purpose particle-physics event generator [582], which includes the simulation of hard lepton–lepton, lepton–hadron and hadron–hadron scattering and soft hadron–hadron collisions in one package. The following describes the features of this program emphasizing the differences between HERWIG and PYTHIA. The discussion is limited to the production of minimum-bias events.

The main limitation of commonly used general-purpose generators is that normally only leading-order processes are included, with higher-order QCD and QED corrections included by showers but no weak corrections at all. HERWIG uses the parton shower approach for initial- and final-state QCD radiation, including colour coherence effects and azimuthal correlations both within and between jets. The use of parton showers in HERWIG is in the same spirit as in PYTHIA but the details and the implementation differ. HERWIG also uses a cluster model for jet hadronization based on non-perturbative gluon splitting, and a similar model for soft and underlying hadronic events. A complete space-time picture of event development is included.

The non-perturbative QCD sector is not solved, so HERWIG and the other general-purpose generators introduce the hadronization aspects based on models rather than on theory. Here is the biggest difference between generators. Three basic models are available at present: independent fragmentation (ISAJET), Lund string fragmentation (PYTHIA/JETSET), and cluster hadronization (HERWIG). The independent fragmentation model does not take into account the colour connections and it is also less successful than HERWIG and PYTHIA. The concept of colour connection is essential for the string and cluster hadronization models. PYTHIA/JETSET gives better agreement with the experimental data than HERWIG but also contains a large number of parameters which can be tuned. Nevertheless, HERWIG, with its few-parametric cluster hadronization model represents the main alternative to the string models. In the cluster hadronization model colour-less clusters are formed from colour-connected quarks [591, 592]. They consist of quark–antiquark (meson-like clusters), quark–diquark (baryon-like), or antiquark–antidiquark (antibaryon-like) pairs. Only the meson-like clusters can occur in e^+e^- interactions. The basic idea of the model is that the clusters decay according to the phase space available to the decay products.

4.3.2.1. Minimum-bias events in HERWIG. It is commonly accepted that soft interactions are described by Regge theory [593] but the details are still controversial. Most fits of the existing data predict a cross section for LHC energies not far from 110 mb, and this agrees with a cosmic-ray measurement [594]. At the highest centre-of-mass energy reached up to now, there are two measurements [595] which are not consistent with each other. The parameters of the Regge fits depend on which data point (if any of both) is used for the parametrizations.

HERWIG uses the Regge fit of Donachie and Landshoff [596]. Their model is the sum of just two powers, $\sigma_{\text{tot}} = Xs^e + Ys^n$, where X represents the pomeron exchange and it is assumed that the pomeron coupling to a particle, a , and its antiparticle, \bar{a} , are the same because pomeron has the quantum numbers of the vacuum. Then $X_{pp} = X_{\bar{p}\bar{p}}$. For Tevatron (1800 GeV) and LHC (14 000 GeV) energies the above formula predicts σ_{tot} of 73 mb and 102 mb, respectively.

The minimum-bias cross section used in HERWIG is an inelastic non-diffractive cross section. HERWIG calculates this value as 70% of the total cross section: $\sigma_{\text{MB}} = 0.7\sigma_{\text{tot}}$. Note that the definition of minimum-bias events as used in the PYTHIA implementation within AliRoot includes the single- and double-diffractive processes. Furthermore, the total cross section for all these processes in PYTHIA is around 79 mb at the LHC energy, i.e. lower than the expected total cross section predicted by Regge fits. Note also that in PYTHIA the cross section depends on the chosen PDF but this is not so in HERWIG, because the Regge fits do not include PDFs.

The number of charged particles N_{ch} in HERWIG is drawn from a negative binomial distribution depending on two parameters, $\langle n \rangle$, the average number of charged particles in the event, and k , a parameter determining the shape of the distribution. The energy dependence of $\langle n \rangle$ and k are taken from fits to UA5 data [599]. Specifically, HERWIG uses $\langle n \rangle = N_{\text{ch,pp}}(\sqrt{s}) = a \times s^b + c$ with default values $a = 9.11$, $b = 0.115$, and $c = -9.50$; and $1/k = d \times \ln s + e$ with $d = 0.029$ and $e = -0.104$.

As mentioned before, most fits of the existing data predict a cross section for the LHC energy not far from 110 mb but there are different measurements from E710 and from CDF [595]. Depending on which one is chosen, the value of cross section can rise up to 120 mb. So the overall normalization of the total cross section as predicted by HERWIG may vary considerably.

Furthermore, there is the issue of the energy dependence of the percentage of the cross section carried by non-diffractive processes, which in HERWIG was 70%. Some experimental results [595] suggest that elastic cross section increases with energy faster than the total cross section and this will be another source of uncertainty in the overall predictions. HERWIG predictions for LHC energies are between 100–120 mb for a total cross section and 60–74 mb for the minimum-bias cross section.

To compare HERWIG and PYTHIA predictions at LHC energies, pp minimum-bias events have been generated with HERWIG 6.100, using the default values of its parameters. The overall multiplicity, the pseudo-rapidity and the transverse-momentum distributions of charged particles as generated with HERWIG in this configuration are shown in figure 4.35. The figure includes also a comparison with PYTHIA events, generated using the default parameters of the AliRoot implementation of the PYTHIA generator (vers. 6.150).

It must be pointed out that minimum-bias events are defined in HERWIG as non-diffractive inelastic collisions, while in our PYTHIA implementation the diffractive events have also been included. That explains the main difference observed in the overall multiplicity distributions of charged particles. There are two peaks in PYTHIA; the one at low multiplicities is related to diffractive events, which are absent in the HERWIG generation. The presence of these diffractive events is also visible in the small humps shown at high absolute values in the pseudo-rapidity distribution.

Finally, a comparison of the number of charged particles within one unit of the central rapidity, and the spectra of their transverse momenta is shown in the last two plots in figure 4.35. While the number of charged particles predicted at mid-rapidity is similar in both generators, the spectra of their transverse momenta are quite different, since no tracks with p_t larger than 3–4 GeV are produced by HERWIG. Similar differences have been observed as well in CDF [600] studies of minimum-bias data in $p\bar{p}$ interactions at $\sqrt{s} = 630$ GeV and $\sqrt{s} = 1800$ GeV. This is an indication of the lack of a description of hard and semi-hard physics in the HERWIG model of minimum-bias events. On the other hand, PYTHIA describes considerably better than HERWIG the minimum-bias p_t -distributions measured at $\sqrt{s} = 1800$ GeV but still fails to reproduce the high- p_t tail at $\sqrt{s} = 630$ GeV. As regards multiplicity distributions, they are described well by both PYTHIA and HERWIG at $\sqrt{s} = 630$ GeV, while at $\sqrt{s} = 1800$ GeV none of them reproduces the high-multiplicity tail.

After these preliminary comparisons of the Tevatron data with the predictions of HERWIG and PYTHIA, more recent analyses have investigated the best adequate tuning of their parameters to reproduce not only minimum-bias data but also the properties of the underlying event in charged-particle jets [601, 602]. HERWIG fails to describe the underlying event because of a too steep p_t -dependence of the beam–beam remnant component. On the other hand, with an appropriate parameter tuning, PYTHIA fits both the underlying event properties and the minimum-bias data. When extrapolating from the Tevatron to the LHC energy, the tuned PYTHIA predicts an increase in the activity of the underlying event which amounts to a factor two, and which is much larger than the HERWIG prediction.

In summary, none of these Monte Carlo programs is so far able to describe all the properties of the currently available data. Work is still in progress to refine the models or to improve the tuning of their parameters, in order to reduce the uncertainties of their predictions at LHC energies.

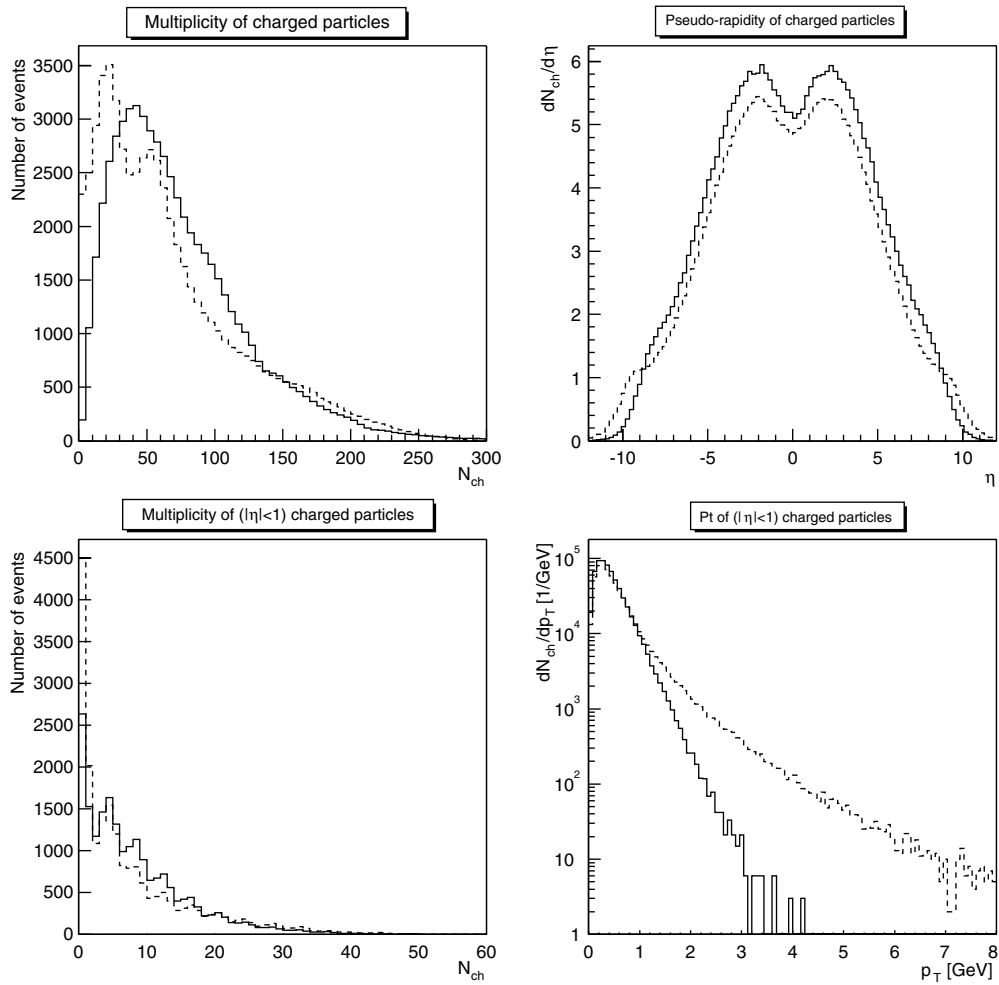


Figure 4.35. Total multiplicity (top left), pseudo-rapidity (top right), central multiplicity (bottom left), and transverse-momentum (bottom right) distributions of charged particles from inelastic pp collisions, as generated with HERWIG (solid line) and PYTHIA (dashed line).

ALICE Collaboration

Alessandria, Italy, Facoltà di Scienze dell'Università: P. Cortese, G. Dellacasa, L. Ramello and M. Sitta.

Aligarh, India, Physics Department, Aligarh Muslim University: N. Ahmad, S. Ahmad, T. Ahmad, W. Bari, M. Irfan and M. Zafar.

Athens, Greece, University of Athens, Physics Department: A. Belogianni, P. Christakoglou, P. Ganoti, A. Petridis, F. Roukoutakis, M. Spyropoulou-Stassinaki and M. Vassiliou.

Bari, Italy, Dipartimento di Fisica dell'Università and Sezione INFN: M. Caselle, D. Di Bari, D. Elia, R. A. Fini, B. Ghidini, V. Lenti, V. Manzari, E. Nappi, F. Navach, C. Pastore, F. Posa, R. Santoro and I. Sgura.

Bari, Italy, Politecnico and Sezione INFN: F. Corsi, D. De Venuto, U. Fratino and C. Marzocca.

Beijing, China, China Institute of Atomic Energy: X. Li, Z. Liu, S. Lu, Z. Lu, Q. Meng, B. Sa, J. Yuan, J. Zhou and S. Zhou.

Bergen, Norway, Department of Physics, University of Bergen: A. Klovning, J. Nystrand, B. Pommeresche, D. Röhrich, K. Ullaland, A. S. Vestbø and Z. Yin.

Bergen, Norway, Bergen University College, Faculty of Engineering: K. Fanebust, H. Helstrup and J. A. Lien.

Bhubaneswar, India, Institute of Physics: R. K. Choudhury, A. K. Dubey, D. P. Mahapatra, D. Mishra, S. C. Phatak and R. Sahoo.

Birmingham, United Kingdom, School of Physics and Space Research, University of Birmingham: D. Evans, G. T. Jones, P. Jovanović, A. Jusko, J. B. Kinson, R. Lietava and O. Villalobos Baillie.

Bologna, Italy, Dipartimento di Fisica dell'Università and Sezione INFN: A. Alici, F. Anselmo, P. Antonioli, G. Bari, M. Basile, Y. W. Baek, L. Bellagamba, D. Boscherini, A. Bruni, G. Bruni, G. Cara Romeo, E. Cerron-Zeballos, L. Cifarelli, F. Cindolo, M. Corradi, D. Falchieri, A. Gabrielli, E. Gandolfi, P. Giusti, D. Hatzifotiadou, G. Laurenti, M. L. Luvisetto, A. Margotti, M. Masetti, S. Morozov, R. Nania, P. Otiougova, F. Palmonari, A. Pesci, F. Pierella, A. Polini, G. Sartorelli, E. Scapparone, G. Scioli, G. P. Vacca, G. Valenti, G. Venturi, M. C. S. Williams and A. Zichichi.

Bratislava, Slovakia, Comenius University, Faculty of Mathematics, Physics and Informatics: V. Černý, R. Janik, S. Kapusta, L. Lúčan, M. Pikna, J. Pišút, N. Pišútová, B. Sitár, P. Strmeň, I. Szarka and M. Zagiba.

Bucharest, Romania, National Institute for Physics and Nuclear Engineering: C. Aiftimiei, V. Catanescu, M. Duma, C. I. Legrand, D. Moisa, M. Petrovici and G. Stoicea.

Budapest, Hungary, KFKI Research Institute for Particle and Nuclear Physics, Hungarian Academy of Sciences: E. Dénes, B. Eged, Z. Fodor, T. Kiss, G. Pála, J. Sulyán and J. Zimányi.

Cagliari, Italy, Dipartimento di Fisica dell'Università and Sezione INFN: S. Basciu, C. Cicalo, A. De Falco, M. Floris, M. P. Macciotta-Serpi, G. Puddu, S. Serci, E. Siddi, L. Tocco and G. Usai.

Cape Town, South Africa, University of Cape Town: J. Cleymans, R. Fearick and Z. Vilakazi.

Catania, Italy, Dipartimento di Fisica dell'Università and Sezione INFN: A. Badalà, R. Barbera, G. Lo Re, A. Palmeri, G. S. Pappalardo, A. Pulvirenti and F. Riggi.

CERN, Switzerland, European Organization for Nuclear Research: Y. Andres, G. Anelli, I. Augustin, A. Augustinus, J. Baechler, P. Barberis, J. A. Belikov¹, L. Betev, A. Boccardi, A. Braem, R. Bramm², R. Brun, M. Burns, P. Buncic², I. Cali³, R. Campagnolo, M. Campbell, F. Carena, W. Carena, F. Carminati, N. Carrer, C. Cheshkov, P. Chochula, J. Chudoba, J. Cruz de Sousa Barbosa⁴, M. Davenport, G. de Cataldo⁵, J. de Groot, A. Di Mauro, R. Dinapoli, R. Divià, C. Engster, S. Evrard, C. Fabjan, A. Fassò, D. Favretto, L. Feng⁶, F. Formenti, E. Futó⁷, A. Gallas-Torreira, A. Gheata, I. Gonzalez-Caballero⁸, C. Gonzalez-Gutierrez, C. Gregory, M. Hoch, H. Hoedelmoser, P. Hristov, M. Ivanov, L. Jirde, A. Junique, W. Klempt, A. Kluge, M. Kowalski⁹, T. Kuhr, L. Leistam, C. Lourenço, J.-C. Marin, P. Martinengo, A. Masoni¹⁰, M. Mast, T. Meyer, A. Mohanty, M. Morel, A. Morsch, B. Mota, H. Muller, L. Musa, P. Nilsson, F. Osmic, D. Perini, A. Peters, D. Picard, F. Piuze, S. Popescu, F. Rademakers, J.-P. Revol, P. Riedler, E. Rosso, K. Šafařík, P. Saiz, J.-C. Santiard, K. Schossmaier, J. Schukraft, Y. Schutz¹¹, E. Schyns, P. Skowronski¹², C. Soos⁷, G. Stefanini, R. Stock², D. Swoboda, P. Szymanski, H. Taureg, M. Tavlet, P. Tissot-Daguette, P. Vande Vyvre, C. Van-Der-Vlugt, J.-P. Vanuxem and A. Vascotto.

Chandigarh, India, Physics Department, Punjab University: M. M. Aggarwal, A. K. Bhati, A. Kumar, M. Sharma and G. Sood.

Clermont-Ferrand, France, Laboratoire de Physique Corpusculaire (LPC), IN2P3-CNRS and Université Blaise Pascal: A. Baldit, V. Barret, N. Bastid, G. Blanchard, J. Castor, P. Crochet, F. Daudon, A. Devaux, P. Dupieux, P. Force, B. Forestier, A. Genoux-Lubain, C. Insa, F. Jouve, L. Lamoine, J. Lecoq, F. Manso, P. Rosnet, L. Royer, P. Saturnini, G. Savinel and F. Yermia.

Columbus, USA, Department of Physics, Ohio State University: T. J. Humanic, I. V. Kotov, M. Lisa, B. S. Nilsen and E. Sugarbaker.

Columbus, USA, Ohio Supercomputer Centre: D. Johnson.

Copenhagen, Denmark, Niels Bohr Institute: I. Bearden, H. Bøggild, P. Christiansen, J. J. Gaardhøje, O. Hansen, A. Holm, B. S. Nielsen and D. Ouerdane.

Cracow, Poland, Henryk Niewodniczanski Institute of Nuclear Physics, High Energy Physics Department: J. Bartke, E. Gładysz-Dziaduś, E. Kornas, A. Rybicki and A. Wroblewski¹³.

¹ On leave from JINR, Dubna, Russia.

² On leave from Frankfurt, Institut für Kernphysik, Johann Wolfgang Goethe-Universität, Germany.

³ On leave from Università del Piemonte Orientale, Alessandria, Italy.

⁴ On leave from Instituto Superior Técnico, Lisbon, Portugal.

⁵ On leave from Dipartimento di Fisica dell'Università and Sezione INFN, Bari, Italy.

⁶ On leave from China Institute of Atomic Energy, Beijing, China.

⁷ On leave from Budapest University, Budapest, Hungary.

⁸ On leave from Universidad de Cantabria, Santander, Spain.

⁹ On leave from Henryk Niewodniczanski Institute of Nuclear Physics, High Energy Physics Department, Cracow, Poland.

¹⁰ On leave from Dipartimento di Fisica dell'Università and Sezione INFN, Cagliari, Italy.

¹¹ On leave from Laboratoire de Physique Subatomique et des Technologies Associées (SUBATECH), Ecole des Mines de Nantes, IN2P3-CNRS and Université de Nantes, Nantes, France.

¹² On leave from University of Technology, Institute of Physics, Warsaw, Poland.

¹³ Cracow Technical University, Cracow, Poland.

Darmstadt, Germany, Gesellschaft für Schwerionenforschung (GSI): A. Andronic¹⁴, D. Antonczyk, H. Appelshäuser, E. Badura, E. Berdermann, P. Braun-Munzinger, O. Busch, M. Ciobanu¹⁴, H. W. Daues, P. Foka¹⁵, U. Frankenfeld, C. Garabatos, H. Gutbrod, C. Lippman, P. Malzacher, A. Marin, D. Miśkowiec, S. Radoski, H. Sako, A. Sandoval, H. R. Schmidt, K. Schwarz, S. Sedykh, R. S. Simon, H. Stelzer, G. Tziledakis and D. Vranic.

Darmstadt, Germany, Institut für Kernphysik, Technische Universität: A. Förster, H. Oeschler and F. Uhlig.

Frankfurt, Germany, Institut für Kernphysik, Johann Wolfgang Goethe-Universität: J. Berger, A. Billmeier, C. Blume, T. Dietel, D. Flierl, M. Gaździcki, Th. Kollegger, S. Lange, C. Loizides, R. Renfordt and H. Ströbele.

Gatchina, Russia, St. Petersburg Nuclear Physics Institute: Ya. Berdnikov, A. Khanzadeev, N. Miftakhov, V. Nikouline, V. Poliakov, E. Rostchine, V. Samsonov, O. Tarasenkova, V. Tarakanov and M. Zhalov.

Havana, Cuba, Centro de Aplicaciones Tecnológicas y Desarrollo Nuclear (CEADEN): E. Lope Torres, A. Abrahantes Quintana and R. Diaz Valdes.

Heidelberg, Germany, Kirchhoff Institute for Physics: V. Angelov, M. Gutfleisch, V. Lindenstruth, R. Panse, C. Reichling, R. Schneider, T. Steinbeck, H. Tilsner and A. Wiebalck.

Heidelberg, Germany, Physikalisches Institut, Ruprecht-Karls Universität: C. Adler, J. Bielčiková, D. Emschermann, P. Glässel, N. Herrmann, Th. Lehmann, W. Ludolphs, T. Mahmoud, J. Milosevic, K. Oyama, V. Petráček, M. Petrovici, I. Rusanov, R. Schicker, H. C. Soltveit, J. Stachel, M. Stockmeier, B. Vulpescu, B. Windelband and S. Yurevich.

Jaipur, India, Physics Department, University of Rajasthan: S. Bhardwaj, R. Raniwala and S. Raniwala.

Jammu, India, Physics Department, Jammu University: S. K. Badyal, A. Bhasin, A. Gupta, V. K. Gupta, S. Mahajan, L. K. Mangotra, B. V. K. S. Potukuchi and S. S. Sambyal.

JINR, Russia, Joint Institute for Nuclear Research: P. G. Akichine, V. A. Arefiev, Ts. Baatar¹⁶, B. V. Batiounia, G. S. Chabratova, V. F. Chepurnov, S. A. Chernenko, V. K. Dodokhov, L. G. Efimov, O. V. Fateev, T. Grigalashvili¹⁷, M. Haiduc¹⁸, D. Hasegan¹⁸, V. G. Kadychevsky, E. K. Koshurnikov, V. Kuznetsov¹⁹, V. L. Lioubochits, V. I. Lobanov, A. I. Malakhov, L. V. Malinina, M. Nioradze²⁰, P. V. Nomokonov, Y. A. Panebrattsev, V. N. Penev, V. G. Pismennaya, I. Roufanov, V. Shestakov¹⁹, A. I. Shklovskaya, P. Smykov, M. K. Suleimanov, Y. Tevzadze²⁰, R. Togoo¹⁶, A. S. Vodopianov, V. I. Yurevich, Y. V. Zanevsky, S. A. Zaporjets and A. I. Zinchenko.

Jyväskylä, Finland, Department of Physics, University of Jyväskylä and Helsinki Institute of Physics: J. Äystö, M. Bondila, V. Lyapin, M. Oinonen, V. Ruuskanen, H. Seppänen and W. Trzaska.

¹⁴ On leave from National Institute for Physics and Nuclear Engineering, Bucharest, Romania.

¹⁵ Also at CERN, Geneva, Switzerland.

¹⁶ Institute of Physics and Technology, Mongolian Academy of Sciences, Ulaanbaatar, Mongolia.

¹⁷ Institute of Physics, Georgian Academy of Sciences, Tbilisi, Georgia.

¹⁸ Institute of Space Sciences, Bucharest, Romania.

¹⁹ Research Centre for Applied Nuclear Physics (RCANP), Dubna, Russia.

²⁰ High Energy Physics Institute, Tbilisi State University, Tbilisi, Georgia.

Karlsruhe, Germany, Institut für Prozessdatenverarbeitung und Elektronik (IPE) (Associate member): T. Blank and H. Gemmeke.

Kharkov, Ukraine, National Scientific Centre, Kharkov Institute of Physics and Technology: G. L. Bochek, A. N. Dovbnya, V. I. Kulibaba, N. I. Maslov, S. V. Naumov, V. D. Ovchinnik, S. M. Potin and A. F. Starodubtsev.

Kharkov, Ukraine, Scientific and Technological Research Institute of Instrument Engineering: V. N. Borshchov, O. Chykalov, L. Kaurova, S. K. Kiprich, L. Klymova, O. M. Listratenko, N. Mykhaylova, M. Protsenko, O. Reznik and V. E. Starkov.

Kiev, Ukraine, Department of High Energy Density Physics, Bogolyubov Institute for Theoretical Physics, National Academy of Sciences of Ukraine: O. Borysov, I. Kadenko, Y. Martynov, S. Molodtsov, O. Sokolov, Y. Sinyukov and G. Zinovjev.

Kolkata, India, Saha Institute of Nuclear Physics: P. Bhattacharya, S. Bose, S. Chattopadhyay, N. Majumdar, S. Mukhopadhyay, A. Sanyal, S. Sarkar, P. Sen, S. K. Sen, B. C. Sinha and T. Sinha.

Kolkata, India, Variable Energy Cyclotron Centre: Z. Ahammed, P. Bhaskar, S. Chattopadhyay, D. Das, S. Das, M. R. Dutta Majumdar, M. S. Ganti, P. Ghosh, B. Mohanty, B. K. Nandi, T. K. Nayak, P. K. Netrakanti, S. Pal, R. N. Singaraju, B. Sinha, M. D. Trivedi and Y. P. Vijoyi.

Košice, Slovakia, Institute of Experimental Physics, Slovak Academy of Sciences and Faculty of Science, P. J. Šafárik University: J. Bán, M. Bombara, S. Fedor, M. Hnatič, I. Králik, A. Kravčáková, F. Kriváň, M. Krivda, G. Martinská, B. Pastirčák, L. Šándor, J. Urbán, S. Vokál and J. Vrláková.

Legnaro, Italy, Laboratori Nazionali di Legnaro: M. Cinausero, E. Fioretto, G. Prete, R. A. Ricci and L. Vannucci.

Lisbon, Portugal, Departamento de Física, Instituto Superior Técnico: P. Branco, R. Carvalho, J. Seixas and R. Vilela Mendes.

Lund, Sweden, Division of Cosmic and Subatomic Physics, University of Lund: H.-A. Gustafsson, A. Oskarsson, L. Osterman, I. Otterlund and E. A. Stenlund.

Lyon, France, Institut de Physique Nucléaire de Lyon (IPNL), IN2P3-CNRS and Université Claude Bernard Lyon-I: B. Cheynis, L. Ducroux, J. Y. Grossiord, A. Guichard, P. Pillot, B. Rapp and R. Tieulent.

Mexico City and Merida, Mexico, Centro de Investigación y de Estudios Avanzados del IPN, Universidad Nacional Autónoma de México, Instituto de Ciencias Nucleares, Instituto de Física: J. R. Alfaro Molina, A. Ayala, A. Becerril, E. Belmont Moreno, J. G. Contreras, E. Cuautle, J. C. D'Olivo, G. Herrera Corral, I. Leon Monzon, J. Martínez Castro, A. Martínez Davalos, A. Menchaca-Rocha, L. M. Montano Zetina, L. Nellen, G. Paic²¹, J. Solano, S. Vergara and A. Zepeda.

Moscow, Russia, Institute for Nuclear Research, Academy of Science: M. B. Goloubeva, F. F. Gouber, O. V. Karavichev, T. L. Karavicheva, E. V. Karpechev, A. B. Kourepin, A. I. Maevskaia, V. V. Marin, I. A. Pshenichnov, V. I. Razine, A. I. Rechetin, K. A. Shileev and N. S. Topilskaja.

²¹ Department of Physics, Ohio State University, Columbus, USA.

Moscow, Russia, Institute for Theoretical and Experimental Physics: A. N. Akindinov, V. Golovine, A. B. Kaidalov, M. M. Kats, I. T. Kiselev, S. M. Kiselev, E. Lioubelev, A. N. Martemiyanov, P. A. Polozov, V. S. Serov, A. V. Smirnitski, M. M. Tchoumakov, I. A. Vetlitski, K. G. Volochine, L. S. Vorobiev and B. V. Zagreev.

Moscow, Russia, Russian Research Center Kurchatov Institute: D. Aleksandrov, V. Antonenko, S. Beliaev, S. Fokine, M. Ippolitov, K. Karadjev, V. Lebedev, V. I. Manko, T. Moukhanova, A. Nianine, S. Nikolaev, S. Nikouline, O. Patarakine, D. Peressounko, I. Sibiriak, A. Tsvetkov, A. Vasiliev, A. Vinogradov, M. Volkov and I. Yushmanov.

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

Münster, Germany, Institut für Kernphysik, Westfälische Wilhelms Universität: C. Baumann, D. Bucher, R. Glasow, H. Gottschlag, N. Heine, K. Reygers, R. Santo, W. Verhoeven, J. Wessels and O. Zaudtke.

Nantes, France, Laboratoire de Physique Subatomique et des Technologies Associées (SUBATECH), Ecole des Mines de Nantes, IN2P3-CNRS and Université de Nantes: L. Aphecette, A. Boucham, K. Boudjemline, G. Conesa-Balbestre, J. P. Cussonneau, H. Delagrangé, M. Dialinas, C. Finck, B. Erazmus, M. Germain, P. Lautridou, F. Lefèvre, L. Luquin, L. Martin, G. Martinez, O. Ravel, C. Roy and A. Tournaire.

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

Oak Ridge, USA, Oak Ridge National Laboratory: T. Awes.

Omaha, USA, Creighton University: M. Cherney and M. Swanger.

Orsay, France, Institut de Physique Nucléaire (IPNO), IN2P3-CNRS and Université de Paris-Sud: M. P. Comets, P. Courtat, B. Espagnon, I. Hřivnáčová, R. Kunne, Y. Le Bornec, M. Mac Cormick, J. Peyré, J. Pouthas, S. Rousseau and N. Willis.

Oslo, Norway, Department of Physics, University of Oslo: L. Bravina, G. Løvhøiden, B. Skaali, T. S. Tvetter, T. Vik, J. C. Wikne and D. Wormald.

Padua, Italy, Dipartimento di Fisica dell'Università and Sezione INFN: F. Antinori, A. Dainese, D. Fabris, M. Lunardon, M. Morando, S. Moretto, A. Pepato, E. Quercigh, F. Scarlassara, G. Segato, R. Turrisi and G. Viesti.

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

Protvino, Russia, Institute for High Energy Physics: M. Yu. Bogolyubsky, G. I. Britvich, G. V. Khaoustov, I. V. Kharlov, S. A. Konstantinov, N. G. Minaev, V. S. Petrov, B. V. Polichtchouk, S. A. Sadovsky, P. A. Semenov, A. S. Soloviev and V. A. Victorov.

Puebla, Mexico, Benemerita Universidad Autonoma de Puebla: A. Fernandez Tellez, E. Gamez Flores, R. Lopez, S. Roman and M. A. Vargas.

Řež u Prahy, Czech Republic, Academy of Sciences of Czech Republic, Nuclear Physics Institute: D. Adamová, S. Kouchpil, V. Kouchpil, A. Kugler, 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 G. M. Urciuoli.

Saclay, France, Centre d'Études Nucléaires, DAPNIA: M. Anfreville, A. Baldisseri, H. Borel, D. Caut, E. Dumonteil, R. Durand, P. De Girolamo, J. Gosset, P. Hardy, V. Hennion, S. Herlant, F. Orsini, Y. Pénichot, H. Pereira, S. Salasca, F. M. Staley and M. Usseglio.

Salerno, Italy, Dipartimento di Fisica 'E. R. Caianiello' dell'Università and Sezione INFN: A. De Caro, S. De Pasquale, A. Di Bartolomeo, M. Fusco Girard, G. Grella, M. Guida, G. Romano, S. Sellitto and T. Virgili.

Sarov, Russia, Russian Federal Nuclear Center (VNIIEF): V. Basmanov, D. Budnikov, V. Demanov, V. Ianowski, R. Ilkaev, L. Ilkaeva, A. Ivanov, A. Khlebnikov, A. Kouryakin, S. Nazarenko, V. Pavlov, S. Philchagin, V. Punin, S. Poutevskoi, I. Selin, I. Vinogradov, S. Zhelezov and A. Zhitnik.

Split, Croatia, Technical University of Split, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture (FESB): S. Gotovac, E. Mudnic and L. Vidak.

St. Petersburg, Russia, V. Fock Institute for Physics of St. Petersburg State University: M. A. Braun, G. A. Feofilov, S. N. Igoikin, A. A. Kolojvari, V. P. Kondratiev, O. I. Stolyarov, T. A. Toulina, F. A. Tsimbal, F. F. Valiev, V. V. Vechernin and L. I. Vinogradov.

Strasbourg, France, Institut de Recherches Subatomiques (IReS), IN2P3-CNRS and Université Louis Pasteur: J. Baudot, D. Bonnet, J. P. Coffin, C. Kuhn, J. Lutz and R. Vernet.

Trieste, Italy, Dipartimento di Fisica dell'Università and Sezione INFN: V. Bonvicini, L. Bosisio, M. Bregant, P. Camerini, E. Fragiaco, N. Grion, R. Grosso, G. Margagliotti, A. Penzo, S. Piano, I. Rachevskaya, A. Rachevski, R. Rui, F. Soramel²² and A. Vacchi.

Turin, Italy, Dipartimenti di Fisica dell'Università and Sezione INFN: B. Alessandro, R. Arnaldi, S. Beolè, P. Cerello, E. Chiavassa, E. Crescio, N. De Marco, A. Ferretti, M. Gallio, P. Giubellino, R. Guernane, M. Idzik, P. Innocenti, A. Marzari-Chiesa, M. Masera, M. Monteno, A. Musso, D. Nouais, C. Oppedisano, A. Piccotti, F. Prino, L. Riccati, E. Scapparini, F. Tosello, E. Vercellin, A. Werbrouck and R. Wheadon.

Utrecht, The Netherlands, Subatomic Physics Department, Utrecht University and National Institute for Nuclear and High Energy Physics (NIKHEF): M. Botje, J. J. F. Buskop, A. P. De Haas, R. Kamermans, P. G. Kuijer, G. Nooren, C. J. Oskamp, Th. Peitzmann, E. Simili, R. Snellings, A. N. Sokolov, A. Van Den Brink and N. Van Eijndhoven.

Wako-shi, Japan, Institute of Research, RIKEN (Associate member): J. Heuser, H. Kano, K. Tanida, M. Tanaka, T. Kawasaki and K. Fujiwara.

Warsaw, Poland, Soltan Institute for Nuclear Studies: A. Deloff, T. Dobrowolski, K. Karpio, M. Kozłowski, H. Malinowski, K. Redlich²³, T. Siemiarczuk, G. Stefanek²⁴, L. Tykarski and G. Wilk.

Warsaw, Poland, University of Technology, Institute of Physics: Z. Chajecki, M. Janik, A. Kisiel, T. J. Pawlak, W. S. Peryt, J. Pluta, M. Slodkowski, P. Szarwas and T. Traczyk.

Worms, Germany, University of Applied Sciences Worms, ZTT (Associate member): E. S. Conner and R. Keidel.

²² Università di Udine, Italy.

²³ Physics Faculty, University of Bielefeld, Bielefeld, Germany and University of Wrocław, Wrocław, Poland.

²⁴ Institute of Physics, Pedagogical University, Kielce, Poland.

Wuhan, China, Institute of Particle Physics, Huazhong Normal University: X. Cai, F. Liu, F. M. Liu, H. Liu, Y. Liu, W. Y. Qian, X. R. Wang, T. Wu, C. B. Yang, H. Y. Yang, Z. B. Yin, D. C. Zhou and D. M. Zhou.

Yerevan, Armenia, Yerevan Physics Institute: M. Atayan, A. Grigorian, S. Grigoryan, H. Gulkanyan, A. Hayrapetyan, A. Harutyunyan, V. Kakoyan, Yu. Margaryan, M. Poghosyan, R. Shahoyan and H. Vardanyan.

Zagreb, Croatia, Ruder Bošković Institute: T. Anticic, K. Kadija and T. Susa.

External Contributors

P. Aurenche²⁵, R. Baier²⁶, F. Becattini²⁷, T. Csörgő²⁸, K. Eggert²⁹, A. Giovannini³⁰, U. Heinz³¹, K. Hencken³², E. Iancu³³, K. Kajantie³⁴, F. Karsch²⁶, V. Koch³⁵, B. Z. Kopeliovich³⁶, M. Laine²⁶, R. Lednický³⁷, M. Mangano²⁹, S. Mrowczynski³⁸, E. Pilon²⁵, R. Rapp³⁹, C. A. Salgado²⁹, B. Tomášik²⁹, D. Treleani⁴⁰, R. Ugoccioni³⁰, R. Venugopalan⁴¹, R. Vogt³⁵, U. A. Wiedemann²⁹.

Acknowledgments

It is a pleasure to acknowledge that the PPR has benefited from a number of contributions from physicists outside the ALICE Collaboration, listed above as External Contributors, in particular in section 1 (Theoretical Overview).

The Collaboration wishes to thank all the administrative and technical staff involved during the preparation of the Physics Performance Report Volume I and in particular: the ALICE secretariat, M Connor, U Genoud, and C Hervet; C Decosse; the CERN Desktop Publishing Service, in particular S Leech O’Neale and C Vanoli.

In addition the Collaboration would like to acknowledge the contributions by S Maridor, L Ray and R Veenhof.

²⁵ Laboratoire d’Annecy-le-Vieux de Physique Théorique LAPTH, Annecy-le-Vieux, France.

²⁶ Fakultät für Physik, Universität Bielefeld, Bielefeld, Germany.

²⁷ Università di Firenze, Firenze, Italy.

²⁸ KFKI Research Institute for Particle and Nuclear Physics, Budapest, Hungary.

²⁹ CERN, Geneva, Switzerland.

³⁰ Dipartimento di Fisica Teorica dell’Università and Sezione INFN, Turin, Italy.

³¹ Ohio State University, Columbus, USA.

³² Universität Basel, Basel, Switzerland.

³³ Service de Physique Théorique, DAPNIA, Saclay, France.

³⁴ Department of Physics, University of Helsinki, Helsinki, Finland.

³⁵ Nuclear Science Division, Lawrence Berkeley National Laboratory, Berkeley, USA.

³⁶ MPI für Kernphysik, Heidelberg, Germany.

³⁷ Laboratory for High Energy (JINR), Dubna, Russia.

³⁸ Institute of Physics, Swietokrzyska Academy, Kielce and Soltan Institute for Nuclear Studies, Warsaw, Poland.

³⁹ NORDITA, Copenhagen, Denmark.

⁴⁰ Dipartimento di Fisica Teorica, Università di Trieste and Sezione INFN, Trieste, Italy.

⁴¹ Brookhaven National Laboratory, Upton, USA.

References

- [1] Hagedorn R 1965 *Nuovo Cim. Suppl.* **3** 147
- [2] Cabibbo N and Parisi G 1975 *Phys. Lett. B* **59** 67
- [3] Pisarski R D and Wilczek F 1984 *Phys. Rev. D* **29** 338
- [4] Karsch F 2002 *Nucl. Phys. A* **698** 199
- [5] Bailin D and Love A 1984 *Phys. Rep.* **107** 325
- [6] Alford M, Rajagopal K and Wilczek F 1998 *Phys. Lett. B* **422** 247
- [7] Rapp R, Schäfer T, Shuryak E V and Velkovsky M 1998 *Phys. Rev. Lett.* **81** 53
- [8] Rajagopal K 2000 *Acta Phys. Polon. B* **31** 3021
- [9] Fodor Z and Katz S D 2002 *J. High Energy Phys.* JHEP03(2002)014
Fodor Z and Katz S D 2004 *J. High Energy Phys.* JHEP04(2004)050
- [10] Allton C R *et al* 2002 *Phys. Rev. D* **66** 074507
- [11] de Forcrand P and Philipsen O 2002 *Nucl. Phys. B* **642** 290
- [12] Cleymans J and Redlich K 1998 *Phys. Rev. Lett.* **81** 5284
- [13] Cleymans J and Redlich K 1999 *Phys. Rev. C* **60** 054908
- [14] Braun-Munzinger P and Stachel J 1996 *Nucl. Phys. A* **606** 320
- [15] Braun-Munzinger P and Stachel J 1998 *Nucl. Phys. A* **638** 3
- [16] Braun-Munzinger P and Stachel J 2002 *J. Phys. G: Nucl. Part. Phys.* **28** 1971
- [17] Wilson K G 1974 *Phys. Rev. D* **10** 2445
- [18] Karsch F, Laermann E and Peikert A 2001 *Nucl. Phys. B* **605** 579
- [19] Ali Khan A *et al* 2001 *Phys. Rev. D* **63** 034502
- [20] Ali Khan A *et al* 2001 *Phys. Rev. D* **64** 074510
- [21] Karsch F, Laermann E and Peikert A 2000 *Phys. Lett. B* **478** 447
- [22] Gottlieb S *et al* 1987 *Phys. Rev. Lett.* **59** 2247
- [23] Gavai R V and Gupta S 2001 *Phys. Rev. D* **64** 074506
- [24] Nakahara Y, Asakawa M and Hatsuda T 1999 *Phys. Rev. D* **60** 091503
- [25] DeTar C and Kogut J B 1987 *Phys. Rev. Lett.* **59** 399
DeTar C and Kogut J B 1987 *Phys. Rev. D* **36** 2828
- [26] Chandrasekharan S *et al* 1999 *Phys. Rev. Lett.* **82** 2463
- [27] Laermann E 1998 *Nucl. Phys. Proc. Suppl. B* **63A–C** 114
- [28] Wetzorke I *et al* 2002 *Nucl. Phys. Proc. Suppl.* **106** 510
- [29] Karsch F *et al* 2002 *Phys. Lett. B* **530** 147
- [30] Blaizot J-P and Iancu E 2002 *Phys. Rep.* **359** 355
Litim D F and Manuel C 2002 *Phys. Rep.* **364** 451
- [31] Arnold P and Zhai C 1994 *Phys. Rev. D* **50** 7603
Arnold P and Zhai C 1994 *Phys. Rev. D* **51** 1906
Zhai C and Kastening B 1995 *Phys. Rev. D* **52** 7232
Braaten E and Nieto A 1996 *Phys. Rev. D* **53** 3421
Braaten E and Nieto A 1996 *Phys. Rev. Lett.* **76** 1417
- [32] Boyd G *et al* 1996 *Nucl. Phys. B* **469** 419
- [33] Pisarski R D 1989 *Phys. Rev. Lett.* **63** 1129
Frenkel J and Taylor J C 1990 *Nucl. Phys. B* **334** 199
- [34] Braaten E and Pisarski R D 1990 *Nucl. Phys. B* **337** 569
- [35] Blaizot J-P and Iancu E 1993 *Nucl. Phys. B* **390** 589
Blaizot J-P and Iancu E 1994 *Nucl. Phys. B* **417** 608
Kelly P F, Liu Q, Lucchesi C and Manuel C 1994 *Phys. Rev. D* **50** 4209
- [36] Andersen J O, Braaten E and Strickland M 2000 *Phys. Rev. D* **61** 074016
Peshier A 2001 *Phys. Rev. D* **63** 105004
- [37] Blaizot J-P, Iancu E and Rebhan A 2001 *Phys. Rev. D* **63** 065003
Blaizot J-P, Iancu E and Rebhan A 2001 *Phys. Lett. B* **523** 143
- [38] Kajantie K *et al* 1997 *Phys. Rev. Lett.* **79** 3130
Kajantie K *et al* 1997 *Nucl. Phys. B* **503** 357
Laine M and Philipsen O 1998 *Nucl. Phys. B* **523** 267
Laine M and Philipsen O 1999 *Phys. Lett. B* **459** 259
- [39] Kajantie K, Laine M, Rummukainen K and Schröder Y 2001 *Phys. Rev. Lett.* **86** 10
- [40] Hart A and Philipsen O 2000 *Nucl. Phys. B* **572** 243
Hart A, Laine M and Philipsen O 2000 *Nucl. Phys. B* **586** 443

- [41] Arnold P, Moore G D and Yaffe L G 2000 *J. High Energy Phys.* JHEP11(2000)001
- [42] Bödeker D 1998 *Phys. Lett. B* **426** 351
Bödeker D 2001 *Phys. Lett. B* **516** 175
- [43] Gribov L V, Levin E M and Ryskin M G 1983 *Phys. Rep.* **100** 1
Mueller A H and Qiu J-W 1986 *Nucl. Phys. B* **268** 427
Blaizot J-P and Mueller A H 1987 *Nucl. Phys. B* **289** 847
- [44] McLerran L and Venugopalan R 1994 *Phys. Rev. D* **49** 2233
McLerran L and Venugopalan R 1994 *Phys. Rev. D* **49** 3352
McLerran L and Venugopalan R 1994 *Phys. Rev. D* **50** 2225
- [45] Golec-Biernat K and Wüsthoff M 1999 *Phys. Rev. D* **60** 114023
- [46] Mueller A H 1999 *Nucl. Phys. B* **558** 285
- [47] Gyulassy M and McLerran L 1997 *Phys. Rev. C* **56** 2219
- [48] Baier R, Mueller A H, Schiff D and Son D T 2001 *Phys. Lett. B* **502** 51
- [49] Jalilian-Marian J, Kovner A, McLerran L and Weigert H 1997 *Phys. Rev. D* **55** 5414
Kovchegov Y V 1996 *Phys. Rev. D* **54** 5463
- [50] Jalilian-Marian J, Kovner A, Leonidov A and Weigert H 1997 *Nucl. Phys. B* **504** 415
Jalilian-Marian J, Kovner A and Weigert H 1999 *Phys. Rev. D* **59** 014015
McLerran L and Venugopalan R 1999 *Phys. Rev. D* **59** 094002
Kovner A and Milhano J G 2000 *Phys. Rev. D* **61** 014012
Iancu E, Leonidov A and McLerran L 2001 *Nucl. Phys. A* **692** 583
- [51] Iancu E, Leonidov A and McLerran L 2001 *Phys. Lett. B* **510** 133
- [52] Gavai R V and Venugopalan R 1996 *Phys. Rev. D* **54** 5795
- [53] Balitsky I 1996 *Nucl. Phys. B* **463** 99
Kovchegov Y 1999 *Phys. Rev. D* **60** 034008
Kovchegov Y 2000 *Phys. Rev. D* **61** 074018
- [54] Kovner A, McLerran L and Weigert H 1995 *Phys. Rev. D* **52** 3809
Kovner A, McLerran L and Weigert H 1995 *Phys. Rev. D* **52** 6231
- [55] Kovchegov Y V 2001 *Nucl. Phys. A* **692** 557
- [56] Krasnitz A and Venugopalan R 1997 *Preprint hep-ph/9706329*
Krasnitz A and Venugopalan R 1999 *Nucl. Phys. B* **557** 237
- [57] Krasnitz A and Venugopalan R 2000 *Phys. Rev. Lett.* **84** 4309
Krasnitz A and Venugopalan R 2001 *Phys. Rev. Lett.* **86** 1717
Krasnitz A, Nara Y and Venugopalan R 2001 *Phys. Rev. Lett.* **87** 192302
- [58] Mueller A H 2000 *Nucl. Phys. B* **572** 227
Mueller A H 2000 *Phys. Lett. B* **475** 220
Bjorker J and Venugopalan R 2001 *Phys. Rev. C* **63** 024609
- [59] Kharzeev D and Nardi M 2001 *Phys. Lett. B* **507** 121
- [60] Kharzeev D and Levin E 2001 *Phys. Lett. B* **523** 79 (*Preprint nucl-th/0108006*)
- [61] Schaffner-Bielich J, Kharzeev D, McLerran L and Venugopalan R 2002 *Nucl. Phys. A* **705** 494
(*Preprint nucl-th/0108048*)
- [62] Eskola K J 2001 *Preprint hep-ph/0104058*
- [63] Hebecker A 2001 *Preprint hep-ph/0111092*
- [64] Eskola K J, Kajantie K, Ruuskanen P V and Tuominen K 2000 *Nucl. Phys. B* **570** 379
- [65] Eskola K J, Ruuskanen P V, Rasanen S S and Tuominen K 2001 *Nucl. Phys. A* **696** 715 (*Preprint hep-ph/0104010*)
- [66] Amelin N S, Armesto N, Pajares C and Sousa D 2001 *Eur. Phys. J. C* **22** 149 (*Preprint hep-ph/0103060*)
- [67] Harris J W and Müller B 1996 *Annu. Rev. Nucl. Part. Sci.* **46** 71
Bass S A, Gyulassy M, Stöcker H and Greiner W 1999 *J. Phys. G: Nucl. Part. Phys.* **25** R1
- [68] Satz H 2000 *Rep. Prog. Phys.* **63** 1511
- [69] Heinz U 1998 *Nucl. Phys. A* **638** 357c
Heinz U 1999 *J. Phys. G: Nucl. Part. Phys.* **25** 263
Heinz U 1999 *Nucl. Phys. A* **661** 140c
- [70] Stock R 1999 *Phys. Lett. B* **456** 277
Stock R 1999 *Prog. Part. Nucl. Phys.* **42** 295
- [71] Stachel J 1999 *Nucl. Phys. A* **654** 119c
- [72] Heinz U 2001 *Nucl. Phys. A* **685** 414c
- [73] Redlich K 2002 *Nucl. Phys. A* **698** 94c

- [74] Schnedermann E, Sollfrank J and Heinz U 1993 *Particle Production in Highly Excited Matter (NATO ASI Series vol B303)* ed H H Gutbrod and J Rafelski (New York: Plenum) p 175
- [75] Heinz U 1995 *Hot Hadronic Matter: Theory and Experiment (NATO ASI Series vol B346)* ed J Letessier *et al* (New York: Plenum) p 413
- [76] Heinz U 1999 *Measuring the Size of Things in the Universe: HBT Interferometry and Heavy Ion Physics* ed S Costa *et al* (Singapore: World Scientific) p 66
- [77] Asakawa M, Heinz U and Müller B 2000 *Phys. Rev. Lett.* **85** 2072
- [78] Jeon S and Koch V 2000 *Phys. Rev. Lett.* **85** 2076
- [79] Ollitrault J-Y 1992 *Phys. Rev. D* **46** 229
- [80] Voloshin S A and Zhang Y 1996 *Z. Phys. C* **70** 665
- [81] Sorge H 1997 *Phys. Rev. Lett.* **78** 2309
Sorge H 1999 *Phys. Rev. Lett.* **82** 2048
- [82] Kolb P F, Sollfrank J and Heinz U 1999 *Phys. Lett. B* **459** 667
- [83] Kolb P F, Sollfrank J and Heinz U 2000 *Phys. Rev. C* **62** 054909
- [84] Kolb P F, Huovinen P, Heinz U and Heiselberg H 2001 *Phys. Lett. B* **500** 232
- [85] Huovinen P *et al* 2001 *Phys. Lett. B* **503** 58
- [86] Teaney D, Lauret J and Shuryak E V 2001 *Phys. Rev. Lett.* **86** 4783
- [87] Teaney D, Lauret J and Shuryak E V 2002 *Nucl. Phys. A* **698** 479c
- [88] Kolb P F *et al* 2001 *Nucl. Phys. A* **696** 197
- [89] Heinz U and Kolb P F 2002 *Nucl. Phys. A* **702** 269c
- [90] Teaney D, Lauret J and Shuryak E V 2001 *Preprint* nucl-th/0110037
- [91] Heinz U and Jacob M 2000 *Preprint* nucl-th/0002042 and <http://cern.web.cern.ch/CERN/Announcements/2000/NewStateMatter/>
- [92] Braun-Munzinger P, Magestro D, Redlich K and Stachel J 2001 *Phys. Lett. B* **518** 41
- [93] Bravina L V *et al* 1999 *Phys. Rev. C* **60** 024904
- [94] Fachini P *et al* [STAR Collaboration] 2002 *J. Phys. G: Nucl. Part. Phys.* **28** 1599 (*Preprint* nucl-ex/0203019)
Xu Z *et al* [STAR Collaboration] 2002 *Nucl. Phys. A* **698** 607c
- [95] Lenkeit B *et al* [CERES Collaboration] 1999 *Nucl. Phys. A* **661** 23c
- [96] Rapp R and Wambach J 2000 *Adv. Nucl. Phys.* **25** 1
- [97] Lee K S, Heinz U and Schnedermann E 1990 *Z. Phys. C* **48** 525
- [98] Schnedermann E, Sollfrank J and Heinz U 1993 *Phys. Rev. C* **48** 2462
- [99] Csörgő T and Lorstad B 1996 *Phys. Rev. C* **54** 1390
- [100] Letessier J *et al* 1995 *Phys. Rev. D* **51** 3408
- [101] Cleymans J and Satz H 1993 *Z. Phys. C* **57** 135
- [102] Braun-Munzinger P, Stachel J, Wessels J P and Xu N 1996 *Phys. Lett. B* **365** 1
Braun-Munzinger P, Heppel I and Stachel J 1999 *Phys. Lett. B* **465** 15
- [103] Becattini F and Heinz U 1997 *Z. Phys. C* **76** 269
- [104] Ackermann K H *et al* [STAR Collaboration] 2001 *Phys. Rev. Lett.* **86** 402
- [105] Adler C *et al* [STAR Collaboration] 2001 *Phys. Rev. Lett.* **87** 262301
- [106] Adcox K [PHENIX Collaboration] 2002 *Phys. Rev. Lett.* **89** 212301 (*Preprint* nucl-ex/0204005)
- [107] Adler C *et al* [STAR Collaboration] 2002 *Phys. Rev. Lett.* **89** 132301
- [108] Wiedemann U A and Heinz U 1999 *Phys. Rep.* **319** 145
- [109] Weiner R M 2000 *Phys. Rep.* **327** 249
- [110] Csörgő T 2002 *Heavy Ion Phys.* **15** 1
- [111] Goldhaber G, Goldhaber S, Lee W Y and Pais A 1960 *Phys. Rev.* **120** 300
- [112] Kopylov G I and Podgoretsky M I 1972 *Sov. J. Nucl. Phys.* **15** 219
Kopylov G I and Podgoretsky M I 1972 *Yad. Fiz.* **15** 392
- [113] Kopylov G I and Podgoretsky M I 1975 *Zh. Eksp. Teor. Fiz.* **69** 414
- [114] Lednický R and Lyuboshits V L 1982 *Sov. J. Nucl. Phys.* **35** 770
Lednický R and Lyuboshits V L 1981 *Yad. Fiz.* **35** 1316
- [115] Koonin S E 1977 *Phys. Lett. B* **70** 43
- [116] Gyulassy M, Kauffmann S K and Wilson L W 1979 *Phys. Rev. C* **20** 2267
- [117] Baym G and Braun-Munzinger P 1996 *Nucl. Phys. A* **610** 286C
- [118] Anchishkin D, Heinz U and Renk P 1998 *Phys. Rev. C* **57** 1428
- [119] Bass S A, Danielewicz P and Pratt S 2000 *Phys. Rev. Lett.* **85** 2689
- [120] Asakawa M, Csörgő T and Gyulassy M 1999 *Phys. Rev. Lett.* **83** 4013
- [121] Sinyukov Y *et al* 1998 *Phys. Lett. B* **432** 248
- [122] Shuryak E V 1973 *Phys. Lett. B* **44** 387

- [123] Heinz U and Zhang Q H 1997 *Phys. Rev. C* **56** 426
- [124] Podgoretskii M I 1983 *Sov. J. Nucl. Phys.* **37** 272
Podgoretskii M I 1983 *Yad. Fiz.* **37** 455
- [125] Herrmann M and Bertsch G F 1995 *Phys. Rev. C* **51** 328
- [126] Chapman S, Scotto P and Heinz U 1995 *Phys. Rev. Lett.* **74** 4400
- [127] Wiedemann U A 1998 *Phys. Rev. C* **57** 266
- [128] Brown D A and Danielewicz P 2001 *Phys. Rev. C* **64** 014902
- [129] Lednický R and Podgoretsky M I 1979 *Yad. Fiz.* **30** 837
- [130] Lednický R and Progulova T B 1992 *Z. Phys. C* **55** 295
- [131] Bolz J *et al* 1993 *Phys. Rev. D* **47** 3860
- [132] Csörgő T, Lörstad B and Zimányi J 1996 *Z. Phys. C* **71** 491
- [133] Wiedemann U A and Heinz U 1997 *Phys. Rev. C* **56** 3265
- [134] Akkelin S V and Sinyukov Y M 1999 *Nucl. Phys. A* **661** 613
- [135] Csörgő T, Lörstad B, Schmid-Sorensen J and Ster A 1999 *Eur. Phys. J. C* **9** 275
- [136] Akkelin S V, Lednický R and Sinyukov Y M 2002 *Phys. Rev. C* **65** 064904
- [137] Cramer J G and Kadija K 1996 *Phys. Rev. C* **53** 908
- [138] Makhlin A N and Sinyukov Y M 1988 *Z. Phys. C* **39** 69
- [139] Bowler M G 1985 *Z. Phys. C* **29** 617
- [140] Andersson B and Hofmann W 1986 *Phys. Lett. B* **169** 364
- [141] Appelschauser H *et al* [NA49 Collaboration] 1998 *Eur. Phys. J. C* **2** 661 (Preprint hep-ex/9711024)
- [142] Tomášik B, Wiedemann U A and Heinz U 2003 *Heavy Ion Phys.* **17** 105 (Preprint nucl-th/9907096)
- [143] Sinyukov Y M, Akkelin S V and Hama Y 2002 *Phys. Rev. Lett.* **89** 052301
- [144] Tomášik B and Wiedemann U A 2003 *Phys. Rev. C* **68** 034905 (Preprint nucl-th/0207074)
- [145] Pratt S 1986 *Phys. Rev. D* **33** 1314
- [146] Rischke D H and Gyulassy M 1996 *Nucl. Phys. A* **608** 479
- [147] Adamová D *et al* [CERES Collaboration] 2003 *Nucl. Phys. A* **714** 124 (Preprint nucl-ex/0207005)
- [148] Adler C *et al* [STAR Collaboration] 2001 *Phys. Rev. Lett.* **87** 082301
- [149] Bjorken J D 1983 *Phys. Rev. D* **27** 140
- [150] Lednický R, Lyuboshits V L, Erasmus B and Nouais D 1996 *Phys. Lett. B* **373** 30
- [151] Lednický R 2001 Preprint nucl-th/0112011
- [152] Gmitro M, Kvasil J, Lednický R and Lyuboshits V L 1986 *Czech. J. Phys. B* **36** 1281
- [153] Appelschauser H 2002 Preprint hep-ph/0204159
- [154] Adamová D *et al* [CERES Collaboration] 2003 *Phys. Rev. Lett.* **90** 022301 (Preprint nucl-ex/0207008)
- [155] Barrette J *et al* [E877 Collaboration] 1997 *Phys. Rev. Lett.* **78** 2916
- [156] Lisa M A *et al* [E895 Collaboration] 2002 *Nucl. Phys. A* **698** 185
- [157] Ferenc D *et al* 1999 *Phys. Lett. B* **457** 347
- [158] Murray M and Holzer B 2001 *Phys. Rev. C* **63** 054901
- [159] Lednický R *et al* 2000 *Phys. Rev. C* **61** 034901
- [160] Heinz U, Scotto P and Zhang Q H 2001 *Ann. Phys.* **288** 325
- [161] Alexander G and Lipkin H J 1995 *Phys. Lett. B* **352** 162
- [162] Lyuboshitz V L and Podgoretsky M I 1996 *Phys. Atomic Nuclei* **59** 476
Lyuboshitz V L and Podgoretsky M I 1997 *Phys. Atomic Nuclei* **60** 39
- [163] Bialas A and Koch V 1999 *Phys. Lett. B* **456** 1
- [164] COBE Collaboration, http://space.gsfc.nasa.gov/astro/cobe/cobe_home.html
- [165] Stodolsky L 1995 *Phys. Rev. Lett.* **75** 1044
Mrówczyński S 1998 *Phys. Lett. B* **430** 9
Shuryak E V 1998 *Phys. Lett. B* **423** 9
- [166] Landau L D and Lifschitz I M 1958 *Course of Theoretical Physics: Statistical Physics* (Oxford: Pergamon Press)
- [167] Stephanov M, Rajagopal K and Shuryak E 1998 *Phys. Rev. Lett.* **81** 4816
Stephanov M, Rajagopal K and Shuryak E 1999 *Phys. Rev. D* **60** 114028
Mrówczyński S 1998 *Phys. Rev. C* **57** 1518
- [168] Pruneau C, Gavin S and Voloshin S 2002 *Phys. Rev. C* **66** 044904 (Preprint nucl-ex/0204011)
- [169] Shuryak E and Stephanov M A 2001 *Phys. Rev. C* **63** 064903
- [170] Heiselberg H and Jackson A 2001 *Phys. Rev. C* **63** 064904
- [171] Randrup J 2002 *Phys. Rev. C* **65** 054906
- [172] Hwa R C and Yang C B 2002 *Phys. Lett. B* **534** 69
- [173] Heiselberg H 2001 *Phys. Rep.* **351** 161

- [174] Jeon S and Koch V *Quark Gluon Plasma 3* ed R C Hwa and X-N Wang (Singapore: World Scientific) p 430
(Preprint hep-ph/0304012)
- [175] Seymour M H 2000 Preprint hep-ph/0007051
- [176] Affolder T *et al* [CDF Collaboration] 2002 *Phys. Rev. D* **65** 092002
- [177] Bjorken J D 1982 Preprint FERMILAB-PUB-82-59-THY
- [178] Thoma M H and Gyulassy M 1991 *Nucl. Phys. B* **351** 491
- [179] Braaten E and Thoma M H 1991 *Phys. Rev. D* **44** 1298
- [180] Wang X-N and Gyulassy M 1991 *Phys. Rev. D* **44** 3501
Gyulassy M and Wang X-N 1994 *Nucl. Phys. B* **420** 583
Wang X-N, Gyulassy M and Plümer M 1995 *Phys. Rev. D* **51** 3436
- [181] Baier R, Dokshitzer Y L, Peigne S and Schiff D 1995 *Phys. Lett. B* **345** 277
- [182] Zakharov B G 1996 *JETP Lett.* **63** 952
Zakharov B G 1997 *JETP Lett.* **65** 615
Zakharov B G 1998 *Phys. At. Nuclei* **61** 838
- [183] Baier R *et al* 1997 *Nucl. Phys. B* **483** 291
Baier R *et al* 1997 *Nucl. Phys. B* **484** 265
- [184] Wiedemann U A 2000 *Nucl. Phys. B* **588** 303
- [185] Gyulassy M, Lévai P and Vitev I 2000 *Nucl. Phys.* **571** 197
- [186] Baier R, Schiff D and Zakharov B G 2000 *Annu. Rev. Nucl. Part. Sci.* **50** 37
- [187] Adler C *et al* 2002 *Phys. Rev. Lett.* **89** 202301 (Preprint nucl-ex/0206011)
- [188] Adcox K *et al* [PHENIX Collaboration] 2002 *Phys. Rev. Lett.* **88** 022301 (Preprint nucl-ex/0109003)
- [189] Adler C *et al* [STAR Collaboration] 2003 *Phys. Rev. Lett.* **90** 082302 (Preprint nucl-ex/0210033)
- [190] Guo X F and Wang X-N 2000 *Phys. Rev. Lett.* **85** 3591
Wang X-N and Guo X F 2001 *Nucl. Phys. A* **696** 788
- [191] Luo M, Qiu J-W and Sterman G 1994 *Phys. Rev. D* **49** 4493
Luo M, Qiu J-W and Sterman G 1994 *Phys. Rev. D* **50** 1951
- [192] Wang X-N, Huang Z and Sarcevic I 1996 *Phys. Rev. Lett.* **77** 231
Wang X-N and Huang Z 1997 *Phys. Rev. C* **55** 3047
- [193] Baier R, Dokshitzer Y L, Mueller A H and Schiff D 2001 *J. High Energy Phys.* JHEP09(2001)033
- [194] Salgado C A and Wiedemann U A 2002 *Phys. Rev. Lett.* **89** 092303 (Preprint hep-ph/0204221)
- [195] Wiedemann U A 2001 *Nucl. Phys. A* **690** 731
- [196] Baier R, Dokshitzer Y L, Mueller A H and Schiff D 1999 *Phys. Rev. C* **60** 064902
Baier R, Dokshitzer Y L, Mueller A H and Schiff D 2001 *Phys. Rev. C* **64** 057902
- [197] Gyulassy M, Lévai P and Vitev I 2002 *Phys. Lett. B* **538** 282 (Preprint nucl-th/0112071)
- [198] Wang X-N 2001 *Phys. Rev. C* **63** 054902
Lokhtin I P and Snigirev A M 2000 *Eur. Phys. J. C* **16** 527
Lokhtin I P and Snigirev A M 1998 *Phys. Lett. B* **440** 163
Gyulassy M, Vitev I and Wang X-N 2001 *Phys. Rev. Lett.* **86** 2537
- [199] Salgado C A and Wiedemann U A 2004 *Phys. Rev. Lett.* **93** 042301 (Preprint hep-ph/0310079)
- [200] Armesto N, Salgado C A and Wiedemann U A 2004 Preprint hep-ph/0405301
- [201] Dokshitzer Y L and Kharzeev D E 2001 *Phys. Lett. B* **519** 199
- [202] Lokhtin I P and Snigirev A M 2001 *Eur. Phys. J. C* **21** 155
- [203] Lin Z W and Vogt R 1999 *Nucl. Phys. B* **544** 339
- [204] Armesto N, Salgado C A and Wiedemann U A 2004 Preprint hep-ph/0405184
- [205] Bourhis L, Fontannaz M and Guillet J-Ph 1998 *Eur. Phys. J. C* **2** 529
- [206] Aurenche P *et al* 1993 *Nucl. Phys. B* **399** 34
Gordon L E and Vogelsang W 1993 *Phys. Rev. D* **48** 3136
- [207] Lai H-L and Li H-N 1998 *Phys. Rev. D* **58** 114020
Laenen E, Sterman G and Vogelsang W 2000 *Phys. Rev. Lett.* **84** 4296
Laenen E, Sterman G and Vogelsang W 2001 *Phys. Rev. D* **63** 114018
- [208] Apanasevich L *et al* 2001 *Phys. Rev. D* **63** 014009
- [209] Aurenche P, Fontannaz M, Guillet J Ph, Kniehl B, Pilon E and Werlen M 1999 *Eur. Phys. J. C* **9** 107
- [210] Kniehl B A, Kramer G and Potter B 2000 *Nucl. Phys. B* **582** 514
- [211] Feinberg E L 1976 *Nuovo Cimento A* **34** 391
- [212] Jalilian-Marian J, Orginos K and Sarcevic I 2002 *Nucl. Phys. A* **700** 523
- [213] Braaten E and Pisarski R D 1990 *Nucl. Phys. B* **339** 310
- [214] Kapusta J, Lichard P and Seibert D 1991 *Phys. Rev. D* **44** 2774
Kapusta J, Lichard P and Seibert D 1991 *Phys. Rev. D* **47** 4171 (erratum)
Baier R, Nakkagawa H, Niegawa A and Redlich K 1992 *Z. Phys. C* **53** 433

- [215] Aurenche P, Gelis F, Kobes R and Petitgirard E 1997 *Z. Phys. C* **75** 315
Aurenche P, Gelis F, Kobes R and Zaraket H 1998 *Phys. Rev. D* **58** 085003
- [216] Arnold P, Moore G D and Yaffe L G 2001 *J. High Energy Phys.* JHEP11(2001)057
Arnold P, Moore G D and Yaffe L G 2001 *J. High Energy Phys.* JHEP12(2001)009
Arnold P, Moore G D and Yaffe L G 2002 *J. High Energy Phys.* JHEP06(2002)030
- [217] Mustafa M G and Thoma M H 2000 *Phys. Rev. C* **62** 014902
Mustafa M G and Thoma M H 2001 *Phys. Rev. C* **63** 069902
Dutta D, Sastry S V, Mohanty A K, Kumar K and Choudhury R K 2002 *Nucl. Phys. A* **710** 415
Sastry S V S 2002 *Preprint* hep-ph/0208103
- [218] Aurenche P, Gelis F and Zaraket H 2002 *J. High Energy Phys.* JHEP05(2002)043
Aurenche P, Gelis F and Zaraket H 2002 *J. High Energy Phys.* JHEP07(2002)063
- [219] Gavin S, McGaughey P L, Ruuskanen P V and Vogt R 1996 *Phys. Rev. C* **54** 2606
- [220] Baur G *et al* 2000 CMS Note-2000/060, available at <http://cmsdoc.cern.ch/documents/00/>
- [221] Eskola K J, Kolhinen V J and Vogt R 2001 *Nucl. Phys. A* **696** 729
- [222] Adcox K *et al* [PHENIX Collaboration] 2002 *Phys. Rev. Lett.* **88** 192303
- [223] Vogt R 2001 *Phys. Rev. C* **64** 044901
- [224] McLerran L D and Toimela T 1985 *Phys. Rev. D* **31** 545
- [225] Shuryak E V 1980 *Phys. Rep.* **61** 71
- [226] Rapp R and Shuryak E V 2000 *Phys. Lett. B* **473** 13
- [227] Rapp R and Wambach J 2000 *Adv. Nucl. Phys.* **25** 1
- [228] Rapp R 2001 *Phys. Rev. C* **63** 054907
- [229] Brown G E and Rho M 2002 *Phys. Rep.* **363** 85
- [230] Rapp R and Wambach J 1999 *Eur. Phys. J. A* **6** 415
- [231] Pisarski R D 1995 *Phys. Rev. D* **52** 3773
- [232] Steele J V, Yamagishi H and Zahed I 1997 *Phys. Rev. D* **56** 5605
- [233] Nason P, Dawson S and Ellis R K 1988 *Nucl. Phys. B* **303** 607
Nason P, Dawson S and Ellis R K 1989 *Nucl. Phys. B* **327** 49
- [234] Beenakker W, Kuijff H, van Neerven W L and Smith J 1989 *Phys. Rev. D* **40** 54
Beenakker W *et al* 1991 *Nucl. Phys. B* **351** 507
- [235] Mangano M L, Nason P and Ridolfi G 1993 *Nucl. Phys. B* **405** 507
- [236] Laenen E, Smith J and van Neerven W L 1992 *Nucl. Phys. B* **369** 543
- [237] Berger E L and Contopanagos H 1996 *Phys. Rev. D* **54** 3085
- [238] Catani S, Mangano M L, Nason P and Trentadue L 1996 *Nucl. Phys. B* **478** 273
- [239] Kidonakis N, Smith J and Vogt R 1997 *Phys. Rev. D* **56** 1553
- [240] Kidonakis N and Vogt R 1999 *Phys. Rev. D* **59** 074014
- [241] Bonciani R, Catani S, Mangano M L and Nason P 1998 *Nucl. Phys. B* **529** 424
- [242] Kidonakis N 2001 *Phys. Rev. D* **64** 014009
- [243] Kidonakis N, Laenen E, Moch S and Vogt R 2001 *Phys. Rev. D* **64** 114001
- [244] McGaughey P L *et al* 1995 *Int. J. Mod. Phys. A* **10** 2999
- [245] Vogt R [Hard Probe Collaboration] 2003 *Int. J. Mod. Phys. E* **12** 211 (*Preprints* LBNL-45350 and hep-ph/0111271)
- [246] Martin A D, Roberts R G, Stirling W J and Thorne R S 1998 *Eur. Phys. J. C* **4** 463
Martin A D, Roberts R G, Stirling W J and Thorne R S 1998 *Phys. Lett. B* **443** 301
- [247] Frixione S, Mangano M L, Nason P and Ridolfi G 1994 *Nucl. Phys. B* **431** 453
- [248] Mangano M L 2004 *Proc. CERN Workshop on Hard Probes in Heavy Ion Collisions at the LHC (October 2003) CERN Yellow Report* CERN-2004-009
- [249] Norrbin E and Sjöstrand T 2000 *Eur. Phys. J. C* **17** 137
- [250] Peterson C, Schlatter D, Schmitt I and Zerwas P 1983 *Phys. Rev. D* **27** 105
- [251] Eskola K J, Kolhinen V J and Salgado C A 1999 *Eur. Phys. J. C* **9** 61
- [252] Eskola K J, Kolhinen V J and Vogt R 2001 *Nucl. Phys. A* **696** 729 (*Preprint* hep-ph/0104124)
- [253] Svetitsky B 1988 *Phys. Rev. D* **37** 2484
Svetitsky B and Uziel A 1997 *Phys. Rev. D* **55** 2616
Mrówczyński S 1991 *Phys. Lett. B* **269** 383
Koike Y and Matsui T 1992 *Phys. Rev. D* **45** 3237
- [254] Mustafa M G, Pal D, Srivastava D K and Thoma M H 1998 *Phys. Lett. B* **428** 234
- [255] Gallmeister K, Kämpfer B and Pavlenko O P 1999 *Eur. Phys. J. C* **8** 473
- [256] Shor A 1988 *Phys. Lett. B* **215** 375
Shor A 1989 *Phys. Lett. B* **233** 231

- [257] Müller B and Wang X-N 1992 *Phys. Rev. Lett.* **68** 2437
- [258] Lévai P and Vogt R 1997 *Phys. Rev. C* **56** 2707
- [259] Braun-Munzinger P and Stachel J 2000 *Phys. Lett. B* **490** 196
- [260] Braun-Munzinger P and Stachel J 2001 *Nucl. Phys. A* **690** 119
- [261] Grandchamp L and Rapp R 2001 *Phys. Lett. B* **523** 60
- [262] Thews R L, Schroedter M and Rafelski J 2001 *Phys. Rev. C* **63** 054905
- [263] Gorenstein M I, Kostyuk A P, Stöcker H and Greiner W 2001 *Phys. Lett. B* **509** 277
- [264] Gorenstein M I, Kostyuk A P, Stöcker H and Greiner W 2001 *J. Phys. G: Nucl. Part. Phys.* **27** L47
- [265] Gorenstein M I, Kostyuk A P, Stöcker H and Greiner W 2000 *Preprint hep-ph/0012292*
- [266] Thews R L 2002 *Nucl. Phys. A* **702** 341 (*Preprint hep-ph/0111015*)
Thews R L 2001 *Proc. on Statistical QCD (Bielefeld, August)* ed K Karsch *et al*
- [267] Tavernier S P K 1987 *Rep. Prog. Phys.* **50** 1439
- [268] Matsui T and Satz H 1986 *Phys. Lett. B* **178** 416
- [269] Karsch F, Mehr M T and Satz H 1988 *Z. Phys. C* **37** 617
- [270] Gunion J F and Vogt R 1997 *Nucl. Phys. B* **492** 301
- [271] Barger V, Keung W Y and Philips R N 1980 *Z. Phys. C* **6** 169
Barger V, Keung W Y and Philips R N 1980 *Phys. Lett. B* **91** 253
- [272] Bodwin G T, Braaten E and Lepage G P 1995 *Phys. Rev. D* **51** 1125
- [273] Schuler G A and Vogt R 1996 *Phys. Lett. B* **387** 181
- [274] Braaten E, Doncheski M A, Fleming S and Mangano M L 1994 *Phys. Lett. B* **333** 548
- [275] Digal S, Petreczky P and Satz H 2001 *Phys. Rev. D* **64** 094015
- [276] Lai H L *et al* 1995 *Phys. Rev. D* **51** 4763
- [277] Beneke M and Rothstein I Z 1996 *Phys. Rev. D* **54** 2005
- [278] Cho P and Leibovich A K 1996 *Phys. Rev. D* **53** 6203
Beneke M and Krämer M 1997 *Phys. Rev. D* **55** 5269
Kniehl B A and Kramer G 1999 *Eur. Phys. J. C* **6** 493
Braaten E, Kniehl B A and Lee J 2000 *Phys. Rev. D* **62** 094005
- [279] Abe F *et al* [CDF Collaboration] 1997 *Phys. Rev. Lett.* **79** 572 and 578
- [280] Affolder T *et al* [CDF Collaboration] 2000 *Phys. Rev. Lett.* **85** 2886
- [281] Vogt R 1999 *Heavy Ion Phys.* **9** 339
- [282] Digal S, Petreczky P and Satz H 2001 *Phys. Lett. B* **514** 57
- [283] Vogt R 1999 *Phys. Rep.* **310** 197
- [284] Kharzeev D and Satz H 1996 *Phys. Lett. B* **366** 316
- [285] Capella A and Sousa D 2001 *Preprint nucl-th/0110072*
Capella A, Kaidalov A B and Sousa D 2002 *Phys. Rev. C* **65** 054908 (*Preprint nucl-th/0105021*)
- [286] Blaizot J-P, Dinh P M and Ollitrault J-Y 2000 *Phys. Rev. Lett.* **85** 4020
- [287] Kharzeev D and Satz H 1994 *Phys. Lett. B* **334** 155
- [288] Martins K, Blaschke D and Quack E 1995 *Phys. Rev. C* **51** 2723
Matinyan S G and Müller B 1998 *Phys. Rev. C* **58** 2994
Haglin K 2000 *Phys. Rev. C* **61** 031902R
Lin Z and Ko C M 2000 *Phys. Rev. C* **62** 034903
- [289] Lin Z and Ko C M 2001 *Phys. Lett. B* **503** 104
- [290] Abreu M C *et al* [NA50 Collaboration] 2000 *Phys. Lett. B* **477** 28
- [291] Abreu M C *et al* [NA50 Collaboration] 1999 *Phys. Lett. B* **450** 456
- [292] Redlich K and Braun-Munzinger P 2000 *Eur. Phys. J. C* **16** 519
- [293] Ko C M, Zhang B, Wang X-N and Zhang X-F 1998 *Phys. Lett. B* **444** 237
- [294] Geiger K and Ellis J R 1995 *Phys. Rev. D* **52** 1500
Geiger K and Ellis J R 1996 *Phys. Rev. D* **54** 1967
- [295] Giubellino P *et al* 2000 *ALICE Internal Note* 2000-28
Revol J-P 2002 *Proc. 3rd Int. Symp. on LHC Physics and Detectors, EPJ Direct C* **4** (S1) 14,
http://194.94.42.12/licensed_materials/10105/bibs/2004001/2004cs01/2004s125.htm
- [296] Afanasev S V *et al* [NA49 collaboration] 2002 *Phys. Rev. C* **66** 054902 (*Preprint nucl-ex/0205002*)
- [297] Adcox K *et al* [PHENIX Collaboration] 2002 *Phys. Rev. Lett.* **88** 242301
- [298] Adler C *et al* [STAR Collaboration] 2004 *Phys. Lett. B* **595** 143 (*Preprint nucl-ex/0206008*)
- [299] Giovannini A and Ugoccioni R 1999 *Phys. Rev. D* **60** 074027
- [300] Giovannini A and Ugoccioni R 1999 *Phys. Rev. D* **59** 094020
- [301] Acosta D *et al* [CDF Collaboration] 2002 *Phys. Rev. D* **65** 072005
- [302] Van Hove L 1982 *Phys. Lett. B* **118** 138

- [303] Alexopoulos T *et al* [E735 Collaboration] 1990 *Phys. Rev. Lett.* **64** 991
- [304] Alexopoulos T *et al* [E735 Collaboration] 1993 *Phys. Rev. D* **48** 984
- [305] Becattini F *et al* 2001 *Phys. Rev. C* **64** 024901
- [306] Kopeliovich B Z and Zakharov B G 1989 *Z. Phys. C* **43** 241
- [307] Kopeliovich B Z and Povh B 1999 *Z. Phys. C* **75** 693
Kopeliovich B Z and Povh B 1999 *Phys. Lett. B* **446** 321
- [308] Rossi G C and Veneziano G 1997 *Nucl. Phys. B* **123** 507
Rossi G C and Veneziano G 1980 *Phys. Rep.* **63** 149
- [309] Calucci G and Treleani D 1998 *Phys. Rev. D* **57** 503
- [310] Calucci G and Treleani D 1999 *Phys. Rev. D* **60** 054023
- [311] Kaidalov A B 1979 *Phys. Rep.* **50** 157
- [312] Alberi G and Goggi G 1981 *Phys. Rep.* **74** 1
- [313] Goulianos K 1983 *Phys. Rep.* **101** 169
- [314] Martin A D 2001 *Preprint hep-ph/0103296*
- [315] Kaidalov A B, Khoze V A, Martin A D and Ryskin M G 2001 *Eur. Phys. J. C* **21** 521
- [316] Abe F *et al* [CDF Collaboration] 1997 *Phys. Rev. Lett.* **79** 584
- [317] Abe F *et al* [CDF Collaboration] 1997 *Phys. Rev. D* **56** 3811
- [318] Martin A D, Roberts R G, Stirling W J and Thorne R S 2002 *Eur. Phys. J. C* **23** 73
- [319] Pumplin J *et al* 2002 *J. High Energy Phys.* JHEP07(2002)012 (*Preprint hep-ph/0201195*)
- [320] Gluck M, Reya E and Vogt A 1998 *Eur. Phys. J. C* **5** 461
- [321] Eskola K J *et al* 2001 *Preprint hep-ph/0110348*
- [322] Eskola K J, Kolhinen V J and Ruuskanen P V 1998 *Nucl. Phys. B* **535** 351 (*Preprint hep-ph/9802350*)
- [323] Hirai M, Kumano S and Miyama M 2001 *Phys. Rev. D* **64** 034003
- [324] Alde D M *et al* 1990 *Phys. Rev. Lett.* **64** 2479
- [325] Arneodo M *et al* [New Muon Collaboration] 1996 *Nucl. Phys. B* **481** 23
- [326] Armesto N 2002 *Eur. Phys. J. C* **26** 35 (*Preprint hep-ph/0206017*)
- [327] Huang Z, Lu H J and Sarcevic I 1998 *Nucl. Phys. A* **637** 79
- [328] Frankfurt L, Guzey V, McDermott M and Strikman M 2002 *J. High Energy Phys.* JHEP02(2002)027
- [329] Li S-Y and Wang X-N 2002 *Phys. Lett. B* **527** 85
- [330] Eskola K J, Kolhinen V J, Salgado C A and Thews R L 2001 *Eur. Phys. J. C* **21** 613
- [331] Eskola K J, Honkanen H, Kolhinen V J and Salgado C A 2002 *Phys. Lett. B* **532** 222
- [332] Paver N and Treleani D 1982 *Nuovo Cimento A* **70** 215
- [333] Strikman M and Treleani D 2002 *Phys. Rev. Lett.* **88** 031801
- [334] Krauss F, Greiner M and Soff G 1997 *Prog. Part. Nucl. Phys.* **39** 503
- [335] Baur G, Hencken K and Trautmann D 1998 *J. Phys. G: Nucl. Part. Phys.* **24** 1657
- [336] Baur G *et al* 2002 *Phys. Rep.* **364** 359
- [337] Baur G *et al* 2002 *Preprint hep-ex/0201034*
Baur G *et al* 2003 *Electromagnetic Probes of Fundamental Physics (Erice, Italy, October 2001)* ed W J Marciano and S White (Singapore: World Scientific) p 235 (*Preprint hep-ex/0201034*)
- [338] Baltz A J, Rhoades-Brown M J and Weneser J 1996 *Phys. Rev. E* **54** 4233
- [339] Chiu M *et al* 2002 *Phys. Rev. Lett.* **89** 012302
- [340] Pshenichnov I A *et al* 2001 *Phys. Rev. C* **64** 024903
- [341] Pshenichnov I A *et al* 2000 *ALICE Internal Note* 2002-07
- [342] Klein S [STAR Collaboration] 2002 *Heavy Ion Phys.* **15** 369 (*Preprint nucl-ex/0104016*)
Klein S [STAR Collaboration] 2001 *Preprint nucl-ex/0108018*
Adler C *et al* [STAR Collaboration] 2002 *Phys. Rev. Lett.* **89** 272302 (*Preprint nucl-ex/0206004*)
- [343] Hencken K and Yepes P (ed) 2002 *Workshop on Ultra-peripheral Collisions in Heavy-Ion Collisions (CERN, March 2002)* <http://quasar.unibas.ch/upc/>
- [344] Ranft J 1999 *Nucl. Phys. Proc. Suppl. B* **71** 228
- [345] Allkofer O C *et al* 1981 *Proc. 17th Int. Conf. on Cosmic Ray* vol 10 (Paris) p 401
Wachsmuth H 1993 *CosmoLep-note* 94000
Ball A *et al* 1994 *Note CERN/LEPC 94-10, LEPC/M 109*
Le Coultre P [L3C Collaboration] 1997 *Proc. 25th Int. Conf. on Cosmic Ray* vol 7 (Durban, South Africa) p 305
L3-Cosmics Experiment 1998 *Experiments at CERN in 1998 (Grey Book, CERN)* p 369
- [346] Taylor C *et al* 1999 *Note CERN/LEPC 99-5, LEPC/P9*
- [347] Avati V *et al* 2003 *Astropart. Phys.* **19** 513
- [348] Ambrosio M *et al* [MACRO Collaboration] 1995 *Phys. Rev. D* **52** 3793

- [349] Klages H O *et al* [KASCADE Collaboration] 1997 *Nucl. Phys. Proc. Suppl. B* **52** 92
- [350] Avati V *et al* 2001 *Note CERN/2001-003, SPSC/P321*; Addendum *CERN/SPSC 2001-012, SPSC/P321 Add1*
- [351] Brandt D 2000 *LHC Project Report* 450
- [352] Meier H *et al* 2001 *Phys. Rev. A* **63** 032713
- [353] Morsch A 2002 *ALICE Internal Note* 2002-32
- [354] Jowett J M *et al* 2003 *LHC Project Report* 642
- [355] Giubellino P *et al* 2000 *ALICE Internal Note* 2000-28
- [356] Revol J-P 2002 *EPJ Direct C* **4** (S1) 14, http://194.94.42.12/licensed_materials/10105/tocs/t2004c-s01a.htm
- [357] Morsch A 1997 *ALICE Internal Note* 97-13
- [358] Baudreghien P 2002 Minutes of the *1st Meeting of the LHC Experiment Machine Data Exchange Working Group*
Muratori B 2002 Minutes of the *4th Meeting of the LHC Experiment Machine Data Exchange Working Group*
- [359] Muratori B 2002 private communication
- [360] Rossi A 2002 private communication
- [361] Rossi A *et al* 2002 *Proc. EPAC 2002 (Paris, France)*
- [362] Collins L R and Malyshev O B 2001 *LHC Project Note* 274
- [363] Rossi A and Hilleret N 2003 *LHC Project Report* 674
- [364] Rossi A 2004 *LHC Project Note* 341
- [365] Baglin V 2003 *Proc. LHC Performance Workshop Chamonix XII (Chamonix, France)*
- [366] Azhgirey I *et al* 2001 *LHC Project Note* 273
- [367] Arduini G *et al* 2003 *LHC Project Report* 645
- [368] Pastirčák B 2002 *ALICE Internal Note* 2002-028
- [369] Fassò A, Ferrari A, Ranft J and Sala P 1997 *Proc. 3rd Workshop on Simulating Accelerator Radiation Environments (SARE-3) (KEK, Tsukuba, Japan)* p 32
- [370] Ranft J 1995 *Phys. Rev. D* **51** 64
Ranft J 1995 Presentaion at *SARE2 Workshop (CERN, 9–11 October)*
Ranft J 1997 *Proc. CERN/TIS-RP/97-05*
- [371] Wang X-N and Gyulassy M 1991 *Phys. Rev. D* **44** 3501
Gyulassy M and Wang X-N 1994 *Comput. Phys. Commun.* **83** 307–31
The code can be found at <http://www-nsdth.lbl.gov/~xnwang/hijing/>
- [372] Tsiledakis G *et al* 2003 *ALICE Internal Note* 2003-010
- [373] Fauss-Golfe A *et al* 2002 *Proc. EPAC 2002 (Paris, France)* p 320
- [374] Kienzle W *et al* [TOTEM Collaboration] 1999 *TOTEM, Total Cross Section, Elastic Scattering and Diffraction Dissociation at the LHC, Technical Proposal CERN-LHCC-99-007, LHCC-P-5*
- [375] Baltz A J, Chasman C and White S N 1998 *Nucl. Instrum. Methods A* **417** 1
- [376] ALICE Collaboration 1999 *Technical Design Report of the Zero Degree Calorimeter CERN/LHCC/1999-05*
- [377] Pshenichnov I A *et al* 2001 *Phys. Rev. C* **64** 024903
- [378] Pshenichnov I A, Bondorf J P, Kurepin A B, Mishustin I N and Ventura A 2002 *ALICE Internal Note* 2002-07
- [379] Dekhissi H *et al* 2000 *Nucl. Phys. A* **662** 207
- [380] Hill J C, Petridis A, Fadern B and Wohn F K 1999 *Nucl. Phys. A* **661** 313c
- [381] Young P G, Arthur E D and Chadwick M B 1998 *Nuclear Reaction Data and Nuclear Reactors vol I* ed A Gandini and G Reffo (Singapore: World Scientific) p 227
- [382] ALICE Collaboration 1995 *Technical Proposal CERN/LHCC/95-71*
- [383] ALICE Collaboration 1996 *Technical Proposal Addendum 1, CERN/LHCC/96-32*
- [384] ALICE Collaboration 1999 *Technical Proposal Addendum 2, CERN/LHCC/99-13*
- [385] ALICE Collaboration 1998 *Technical Design Report of the High-Momentum Particle Identification Detector CERN/LHCC/1998-19*
- [386] ALICE Collaboration 1999 *Technical Design Report of the Photon Spectrometer CERN/LHCC/1999-04*
- [387] ALICE Collaboration 1999 *Technical Design Report of the Zero-Degree Calorimeter CERN/LHCC/1999-05*
- [388] ALICE Collaboration 1999 *Technical Design Report of the Inner Tracking System CERN/LHCC/1999-12*
- [389] ALICE Collaboration 1999 *Technical Design Report of the Forward Muon Spectrometer CERN/LHCC/1999-22*
- [390] ALICE Collaboration 2000 *Technical Design Report of the Forward Muon Spectrometer Addendum-1, CERN/LHCC/2000-46*
- [391] ALICE Collaboration 1999 *Technical Design Report of the Photon Multiplicity Detector CERN/LHCC/1999-32*
- [392] ALICE Collaboration 2003 *Technical Design Report for the Photon Multiplicity Detector Addendum-1, CERN/LHCC 2003-038*

- [393] ALICE Collaboration 2000 *Technical Design Report of the Time-Projection Chamber* CERN/LHCC/2000-01
- [394] ALICE Collaboration 2000 *Technical Design Report of the Time-Of-Flight Detector* CERN/LHCC/2000-12; Addendum CERN/LHCC/2002-16
- [395] ALICE Collaboration 2001 *Technical Design Report of the Transition-Radiation Detector* CERN/LHCC/2001-21
- [396] ALICE Collaboration 2004 *Technical Design Report of the Forward Detectors* CERN/LHCC/2004-025
- [397] ALICE Collaboration 2003 *Technical Design Report of the Trigger, Data Acquisition, High Level Trigger and Control System* CERN/LHCC/2003-062
- [398] <http://www.cern.ch/ALICE/Projects/offline/aliroot/Welcome.html>
- [399] Wang X-N and Gyulassy M 1991 *Phys. Rev. D* **44** 3501
Gyulassy M and Wang X-N 1994 *Comput. Phys. Commun.* **83** 307-31
The code can be found at <http://www-nsdth.lbl.gov/~xnwang/hijing/>
- [400] Brun R, Bruyant F, Maire M, McPherson A C and Zanarini P 1985 *GEANT3 User Guide* CERN Data Handling Division DD/EE/84-1
<http://wwwinfo.cern.ch/asdoc/geantold/GEANTMAIN.html>
Brun R *et al* 1994 *CERN Program Library Long Write-up* W5013, GEANT Detector Description and Simulation Tool
- [401] Antinori F *et al* 1995 *Nucl. Instrum. Methods A* **360** 91
Antinori F *et al* 1995 *Nucl. Phys. A* **590** 139c
Manzari V *et al* 1995 *Nucl. Phys. A* **661** 7161c
- [402] Faccio F *et al* 1998 *Proc. 4th Workshop on Electronics for LHC Experiments (Rome, 21-25 September)* (Rome: INFN) pp 105-113
- [403] Kluge A *et al* 2001 *Proc. 7th Workshop on Electronics for LHC Experiments (Stockholm, 10-14 September)* pp 95-100
Kluge A *et al* 2002 *Proc. PIXEL 2002 Workshop (Carmel, September 2002), CERN Yellow Report* CERN-2001-005, p 95 (published in eConf C 020909)
- [404] Snoeys W *et al* 2001 *Proc. PIXEL 2000 Workshop (Genova, June 2000)* ed L Rossi *Nucl. Instrum. Methods A* **465** 176
- [405] Evans D and Jovanović P 2003 *ALICE Internal Note* 2003-055
- [406] Rashevsky A *et al* 2001 *Nucl. Instrum. Methods A* **461** 133-38
- [407] Nouais D *et al* 2000 *Nucl. Instrum. Methods A* **450** 338-42
- [408] Nouais D *et al* 2001 *CERN-ALI-2001-059*
- [409] Cerello P *et al* 2001 *ALICE Internal Note* 2001-34
- [410] Lopez Torres E and Cerello P 2001 *ALICE Internal Note* 2001-35
- [411] Mazza G *et al* 2001 *CERN-ALI-2001-025*
- [412] Alberici G *et al* 1999 *ALICE Internal Note* 1999-28
Alberici G *et al* 2001 *Nucl. Instrum. Methods A* **471** 281-84
- [413] Antinori S *et al* 2003 *LECC'03, Proc. 9th Workshop on Electronics for LHC Experiments (Amsterdam Holland, 29 September-3 October 2003)* *CERN Yellow Report* CERN-2003-006
- [414] Carena W *et al* 2001 *ALICE Internal Note* 2001-47
- [415] Borshchov V *et al* 2002 *Proc. 8th Workshop on Electronics for LHC Experiments (Colmar, France, 9-13 September 2002)* *CERN Yellow Report* CERN-2002-003
- [416] van den Brink A *et al* 2001 *Proc. 7th Workshop on Electronics for the LHC Experiments (Stockholm, Sweden, 10-14 September 2001)* *CERN Yellow Report* CERN-2001-005
- [417] Hu-Guo C *et al* 2002 *Proc. 8th Workshop on Electronics for LHC Experiments (Colmar, France, 9-13 September 2002)* *CERN Yellow Report* CERN-2002-003
- [418] Kluit R *et al* 2001 *Proc. 7th Workshop on Electronics for the LHC Experiments (Stockholm, Sweden, 10-14 September 2001)* *CERN Yellow Report* CERN-2001-005
- [419] ALICE Collaboration *ALICE Physics Performance Report, Volume II* section 5, to be published
- [420] Meyer T *et al* 2000 *ALICE Internal Note* 2000-011
- [421] Meyer T *et al* 2001 *ALICE Internal Note* 2001-046
- [422] Stelzer H *et al* 2003 *ALICE Internal Note* 2003-017
- [423] Frankenfeld U *et al* 2002 *ALICE Internal Note* 2002-030
- [424] Musa L [ALICE Collaboration] 2003 *Nucl. Phys. A* **715** 843c-48c
- [425] Mota B *et al* 2002 *Proc. ESSCIRC (Florence, September 2002)* ed A Baschirritto and P Malcovati, p 259
- [426] Bosch R E, de Parga A J, Mota B and Musa L 2003 *IEEE Trans. Nucl. Sci.* **50** 2460
- [427] Musa L *et al* 2000 *Proc. 6th Workshop on Electronics for LHC Experiments (Krakow, September 2000)* *CERN Yellow Report* CERN-2000-010

- [428] Lien J *et al* 2002 *Proc. 8th Workshop on Electronics for LHC Experiments (Colmar, 9–13 September 2002)* CERN Yellow Report CERN-2002-003
- [429] Camapgnolo R *et al* 2003 *Proc. 9th Workshop on Electronics for the LHC Experiments (Amsterdam, Netherlands, 29 September–3 October 2003)* CERN Yellow Report CERN-2003-006
- [430] Andronic A *et al* 2001 *IEEE Trans. Nucl. Sci.* **48** 1259
- [431] Andronic A *et al* 2003 *Nucl. Instrum. Methods A* **498** 143
- [432] Cerron-Zeballos E, Crotty I, Hatzifotiadou D, Lamas-Valverde J, Neupane D, Williams M C S and Zichichi A 1996 *Nucl. Instrum. Methods A* **374** 132
- [433] Akindinov A *et al* 2000 *Nucl. Instrum. Methods A* **456** 16
- [434] Akindinov A *et al* 2002 *Nucl. Instrum. Methods A* **490** 58
- [435] Basile M 2003 *ALICE LHCC Comprehensive Review* <http://agenda.cern.ch/fullAgenda.php/?ida=a021989>
- [436] Christiansen J 2001 CERN/EP-MIC, HPTDC High Performance Time to Digital Converter, Version 2.0 for HPTDC version 1.1, September 2001, available online at http://micdigital.web.cern.ch/micdigital/hptdc/hptdc_manual_ver2.0.pdf
- [437] Andres Y *et al* 2002 *EPJ Direct C* **4** (S1) 25, http://194.94.42.12/licensed_materials/10105/tocs/t2004c-s01a.htm
- [438] Piuz F 1996 *Nucl. Instrum. Methods A* **371** 96
Nappi E 2001 *Nucl. Instrum. Methods A* **471** 18
- [439] Di Mauro A 2000 *Preprint CERN-EP-2000-58*
- [440] Santiard J C and Marent K 2001 *ALICE Internal Note* 2001-49
- [441] Witters H, Santiard J C and Martinengo P 2000 *ALICE Internal Note* 2000-010
- [442] Bogolyubsky M Yu, Kharlov Yu V and Sadovsky S A 2003 *Nucl. Instrum. Methods A* **502** 719
- [443] Ippolitov M *et al* 2002 *Nucl. Instrum. Methods A* **486** 121
- [444] Blik A M *et al* 2001 *Instrum. Exp. Tech.* **44** 339
- [445] Bogolyubsky M Yu *et al* 2002 *Instrum. Exp. Tech.* **45** 327
- [446] Thews R L, Schroeder M and Rafelski J 2001 *Phys. Rev. C* **63** 054905
- [447] Braun-Munzinger P and Stachel J 2000 *Phys. Lett. B* **490** 196
- [448] Sansoni A *et al* [CDF Collaboration] 1996 *Nucl. Phys. A* **610** 373c
- [449] Fassò A *et al* 1994 *Proc. IV Int. Conf. on Calorimeters and their Applications* (Singapore: World Scientific) p 493
- [450] Donskoy E N *et al* 1999 *VNIIEF Status Report*; private communications
- [451] Duinker P *et al* 1988 *Nucl. Instrum. Methods A* **273** 814
- [452] Arnaldi R *et al* 2002 *Nucl. Instrum. Methods A* **490** 51 and references therein
- [453] Ferretti A *et al* 2003 *Proc. VI Workshop on Resistive Plate Chambers and Related Detectors (Coimbra, 26–27 November 2001)* *Nucl. Instrum. Methods A* **508** 106
- [454] Arnaldi R *et al* 2001 *Nucl. Instrum. Methods A* **457** 117
- [455] Gorodetzky P *et al* 1994 *Proc. 4th Int. Conf. on Calorimetry in High Energy Physics* (Singapore: World Scientific) p 433
Anzivino G *et al* 1995 *Nucl. Phys. Proc. Suppl. B* **44** 168
Ganel O and Wigmans R 1995 *Nucl. Instrum. Methods A* **365** 104
- [456] Gorodetzky P *et al* 1992 *Radiation Physics and Chemistry* (Oxford: Pergamon Press) p 253
Gorodetzky P *et al* 1995 *Nucl. Instrum. Methods A* **361** 161
- [457] Arnaldi R *et al* 1998 *Nucl. Instrum. Methods A* **411** 1
- [458] Gorodetzky P *et al* 1995 *Nucl. Instrum. Methods A* **361** 161
- [459] Arnaldi R *et al* 1999 *Proc. 8th Int. Conf. on Calorimetry in High Energy Physics (Lisbon, Portugal)* p 362
- [460] Arnaldi R *et al* 2001 *Nucl. Instrum. Methods A* **456** 248
- [461] Arnaldi R *et al* 2001 *Proc. Int. Europhys. Conf. on High Energy Physics (HEP 2001) (Budapest, Hungary)* (*Preprint hep2001/257*)
- [462] Oppedisano C *et al* 2002 *ALICE Internal Note* 2002-08
Oppedisano C *et al* 2001 *PhD Thesis* Università di Torino
- [463] Cheynis B *et al* 2003 *ALICE Internal Note* 2003-040
- [464] Cheynis B *et al* 2000 *ALICE Internal Note* 2000-29
- [465] Grigoriev V *et al* 2000 *Instrum. Exp. Tech.* **43** 750
- [466] Grigoriev V *et al* 2000 *ALICE Internal Note* 2000-17
- [467] Grigoriev V *et al* 2001 *ALICE Internal Note* 2001-38
- [468] Alessandro B *et al* 2003 *Proc. 28th Int. Conf. on Cosmic Ray (Tokyo, Japan)* ed Y Muraki (Tokyo: Universal Academic Press) pp 1203–06
- [469] Adriani O *et al* [L3 + C Collaboration] 2002 *Nucl. Instrum. Methods A* **488** 209
- [470] Dzhelyadin R I *et al* 1986 *DELPHI Internal Note* 86-108, TRACK 42, CERN

- [471] Paic G *et al* 2000 *ALICE Internal Note* 2000-30
- [472] Nicolauic A, Mattavelli M and Carrato S 2002 *Nucl. Instrum. Methods Phys. Res. A* **487** 542
- [473] Lindenstruth V *et al* 2001 *ALICE Internal Note* 2001-19
- [474] Giubellino P *et al* 2000 *ALICE Internal Note* 2000-28
- [475] Rubin G 1998 *ALICE Internal Note* 98-21
- [476] Denes E 2002 *ALICE Internal Note* 2002-009
- [477] Rubin G and Sulyan J 1997 *ALICE Internal Note* 1997-14
- [478] Carena W *et al* 2002 *PCI-based Read-Out Receiver Card (RORC) in the ALICE DAQ system, Proc. 18th JINR Int. Symp. on Nuclear Electronics and Computing (Varna, Bulgaria)*
- [479] CERN ALICE DAQ Group 2002 *ALICE Internal Note* 2002-36
- [480] Baud J P *et al* 2001 *ALICE Internal Note* 2001-36
- [481] Di Marzo G *et al* 2002 *ALICE Internal Note* 2002-01
- [482] Anticic T *et al* 2003 *ALICE Internal Note* 2003-001
- [483] Anticic T *et al* 2001 *Proc. on Computing in High Energy Physics (CHEP 2001) (IHEF Beijing, China, September)*
- [484] Ptolemy Project, <http://ptolemy.eecs.berkeley.edu/ptolemyclassic/pt0.7.1/index.htm>
- [485] ALICE HLT Collaboration 2002 *ALICE High Level Trigger Conceptual Design Report*; http://www.kip.uni-heidelberg.de/ti/L3/CDRFinal_CERN-tilsn.pdf
- [486] Manso F *et al* 2002 *ALICE Internal Note* 2002-04
- [487] Loizides C *et al* 2003 *IEEE Trans. Nucl. Sci.* **51** 3 (*Preprint physics/0310052*)
- [488] Vestbø A *et al* 2002 *Proc. 12th IEEE-NPSS REAL TIME Conf. (Valencia, June 2001)* *IEEE Trans. Nucl. Sci.* **49** 389
- [489] Brinkmann D *et al* 1995 *Nucl. Instrum. Methods A* **354** 419
- [490] Pühlhofer F, Röhrich D and Keidel R 1988 *Nucl. Instrum. Methods A* **263** 360
- [491] Berger J *et al* 2001 *Nucl. Instrum. Methods A* **462** 463
- [492] Grastveit G *et al* 2003 *Proc. 9th Workshop on Electronics for LHC Experiments (Amsterdam, 29 September–3 October 2003)* *CERN Yellow Report* CERN-2003-006
- [493] Grastveit G *et al* 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA, 24–28 March 2003)* (published in eConf C 0303241) (*Preprint physics/0306017*)
- [494] Tilsner H *et al* 2003 *Proc. Int. Europhys. Conf. on High Energy Physics EPS (Aachen, Germany)* to be published
- [495] Reinefeld A and Lindenstruth V 2001 *Proc. 2nd Int. Workshop on Metacomputing Systems and Applications, MSA'2001 (Valencia, 3 September 2001)* p 221
- [496] Steinbeck T *et al* 2002 *IEEE Trans. Nucl. Sci.* **49** 455
- [497] See e.g. Booch G 1993 *Object-oriented Analysis and Design with Applications* 2nd edn (Redwood City: Benjamin Cummings)
- [498] <http://wwwinfo.cern.ch/asd/cernlib>
- [499] <http://wwwinfo.cern.ch/asd/paw/>
- [500] Brun R, Bruyant F, Maire M, McPherson A C and Zanarini P 1985 *GEANT3 User Guide* CERN Data Handling Division DD/EE/84-1
<http://wwwinfo.cern.ch/asdoc/geantold/GEANTMAIN.html>
- [501] <http://root.cern.ch>
- [502] <http://www.cern.ch/ALICE/Projects/offline/aliroot/Welcome.html>
- [503] Saiz P *et al* 2003 *Nucl. Instrum. Methods A* **502** 437
<http://alien.cern.ch/>
- [504] Ballintijn M, Brun R, Rademakers F and Roland G 2004 *Distributed Parallel Analysis Framework with PROOF, Proc. TUCT004*
- [505] Agostinelli S *et al* 2003 *Geant4—A Simulation Toolkit* CERN-IT-20020003, KEK *Preprint* 2002-85, SLAC-PUB-9350; *Nucl. Instrum. Methods A* **506** 250
<http://wwwinfo.cern.ch/asd/geant4/geant4.html>
- [506] Fassò A *et al* 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA)*
<http://www.slac.stanford.edu/econf/C0303241/proc/papers/MOMT004.PDF>
- [507] Hřivnáčová I *et al* 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA)*
<http://www.slac.stanford.edu/econf/C0303241/proc/papers/THJT006.PDF>
- [508] Wang X-N and Gyulassy M 1991 *Phys. Rev. D* **44** 3501
Gyulassy M and Wang X-N 1994 *Comput. Phys. Commun.* **83** 307–31
The code can be found at <http://www-nsdth.lbl.gov/~xnwang/hijing/>
- [509] Bengtsson H-U and Sjostrand T 1987 *Comput. Phys. Commun.* **46** 43
The code can be found at <http://www.thep.lu.se/~torbjorn/Pythia.html>

- Sjostrand T 1994 *Comput. Phys. Commun.* **82** 74
The code can be found at <http://www.thep.lu.se/~torbjorn/Pythia.html>
- [510] Abe F *et al* [CDF Collaboration] 1988 *Phys. Rev. Lett.* **61** 1819
- [511] See e.g. *Note ATL-PHYS-96-086*
<http://preprints.cern.ch/cgi-bin/setlink?base=atlnot&categ=Note&id=phys-96-086>
- [512] González Caballero I *et al* 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA)*
<http://www.slac.stanford.edu/econf/C0303241/proc/papers/MOMT011.PDF>
- [513] Carminati F and González Caballero I 2001 *ALICE Internal Note* 2001-041
- [514] Note in course of publication, see <http://wwwinfo.cern.ch/asd/geant4/reviews/delta-2002/schedule.html>
- [515] Fassò A *et al* 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA)*
<http://www.slac.stanford.edu/econf/C0303241/proc/papers/MOMT005.PDF>
- [516] Campanella M, Ferrari A, Sala P R and Vanini S 1998 Reusing code from FLUKA and GEANT4 geometry
ATLAS Internal Note ATL-SOFT-98-039
Campanella M, Ferrari A, Sala P R and Vanini S 1999 First calorimeter simulation with the FLUGG prototype
ATLAS Internal Note ATL-SOFT-99-004
- [517] Carminati F and Morsch A 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA)*
<http://www.slac.stanford.edu/econf/C0303241/proc/papers/TUMT004.PDF>
- [518] <http://aldwww.cern.ch>
- [519] <http://nuclear.ucdavis.edu/jklay/ALICE/>
- [520] Tsiledakis G *et al* 2003 *ALICE Internal Note* 2003-010
- [521] Fassò A, Ferrari A, Sala P R and Tsiledakis G 2001 *ALICE Internal Note* 2001-28
- [522] Carrer N and Dainese A 2003 *ALICE Internal Note* 2003-019
Dainese A and Turrise R 2003 *ALICE Internal Note* 2003-028
- [523] Billoir P 1984 *Nucl. Instrum. Methods A* **225** 352
Billoir P *et al* 1984 *Nucl. Instrum. Methods A* **241** 115
Fruhwith R 1987 *Nucl. Instrum. Methods A* **262** 444
Billoir P 1989 *Comput. Phys. Commun.* **57** 390
- [524] <http://aliweb.cern.ch/people/skowron>
- [525] <http://www.cern.ch/MONARC/>
- [526] Foster I and Kesselmann C (ed) 1999 *The Grid: Blueprint for a New Computing Infrastructure* (San Francisco, CA: Kaufmann)
Foster I, Kesselmann C and Tuecke S 2001 The Anatomy of the Grid Enabling Scalable Virtual Organizations, see <http://www.globus.org/research/papers/anatomy.pdf>
- [527] <http://www.eu-datagrid.org/>
<http://www.gridpp.ac.uk/>
<http://www.ppdg.net/>
<http://www.griphyn.org/>
<http://www.ivdgl.org/>
<http://www.infn.it/grid/>
- [528] <http://www.globus.org/>
- [529] Condor Classified Advertisements, <http://www.cs.wisc.edu/condor/classad>
- [530] Buncic P 2003 *Proc. Computing in High Energy and Nuclear Physics (La Jolla, CA)*
<http://www.slac.stanford.edu/econf/C0303241/proc/papers/MOAT004.PDF>
- [531] <http://lcg.web.cern.ch/LCG/SC2/>
- [532] <http://lcg.web.cern.ch/LCG/lw2002/>
- [533] <http://lhc-computing-review-public.web.cern.ch/>
- [534] <http://aliweb.cern.ch/offline/aliroot-pro/nightbuilds.html>
- [535] <http://www.cvshome.org>
- [536] <http://aliweb.cern.ch/cgi-bin/cvswweb>
- [537] http://aliweb.cern.ch/offline/aliroot-new/roothtml/USER_Index.html
- [538] <http://aliweb.cern.ch/offline/codingconv.html>
- [539] Potrich A and Tonella P 2000 *Proc. Int. Conf. on Computing in High Energy and Nuclear Physics (Padova, 7–11 February 2000)* pp 758–61
Tonella P and Potrich A 2002 *Sci. Comput. Programming* **42** 229
Tonella P and Potrich A 2001 *Proc. Int. Conf. on Software Maintenance (Florence, 7–9 November 2001)* p 376
Potrich A and Tonella P 2001 Automatic verification of coding standards SCAM'2001, *Int. Workshop on Source Code Analysis and Manipulation*
- [540] <http://aliweb.cern.ch/offline/aliroot-new/log/violatedRules.html>

- [541] Brown W J *et al* 1998 *AntiPatterns: Refactoring Software, Architectures, and Projects in Crisis*
- [542] <http://www.rational.com/uml>
- [543] <http://aliweb.cern.ch/offline/reveng>
- [544] OSCAR, Open Standard Code And Routines <http://nt3.phys.columbia.edu/OSCAR>
- [545] van Eijndhoven N *et al* 1995 *ALICE Internal Note ALICE/GEN 95-32*
- [546] Andersson B *et al* 1983 *Phys. Rep.* **97** 31
- [547] Andersson B *et al* 1987 *Nucl. Phys. B* **281** 289
Nilsson-Almqvist B and Stenlund E 1987 *Comput. Phys. Commun.* **43** 387
- [548] Capella A *et al* 1994 *Phys. Rep.* **236** 225
- [549] Wang X-N 2002 private communications
- [550] Morsch A <http://home.cern.ch/~morsch/AliGenerator/AliGenerator.html> and <http://home.cern.ch/~morsch/generator.html>
- [551] Ranft J 1995 *Phys. Rev. D* **51** 64
- [552] Ranft J 1999 *Preprint hep-ph/9911213*
Ranft J 1999 *Preprint hep-ph/9911232*
Ranft J 2000 *Preprint hep-ph/0002137*
- [553] Alber T *et al* [NA35 Collaboration] 1994 *Z. Phys. C* **64** 195
- [554] Alber T *et al* [NA35 Collaboration] 1998 *Eur. J. Phys. C* **2** 643
- [555] Kharzeev D 1996 *Phys. Lett. B* **378** 238
- [556] Capella A and Kopeliovich B 1996 *Phys. Lett. B* **381** 325
- [557] Barrett R V and Jackson D F 1977 *Nuclear Sizes and Structure* (Oxford: Clarendon Press)
- [558] Roesler S, Engel R and Ranft J 1998 *Phys. Rev. D* **57** 2889
- [559] Roesler S 1999 private communication
- [560] Glück M, Reya E and Vogt A 1995 *Z. Phys. C* **67** 433
- [561] Glück M, Reya E and Vogt A 1998 *Eur. Phys. J. C* **5** 461
- [562] Amelin N S, Armesto N, Pajares C and Sousa D 2001 *Eur. Phys. J. C* **22** 149
Armesto N and Pajares C 2000 *Int. J. Mod. Phys. A* **15** 2019
Ferreiro E G, Pajares C and Sousa D 1998 *Phys. Lett. B* **422** 314
Armesto N, Braun M A, Ferreiro E G and Pajares C 1995 *Phys. Lett. B* **344** 301
Amelin N S *et al* 1995 *Phys. Rev. C* **52** 362
- [563] Amelin N S, Braun M A and Pajares C 1993 *Phys. Lett. B* **306** 312
Amelin N S, Braun M A and Pajares C 1994 *Z. Phys. C* **63** 507
- [564] Broniowski W and Florkowski W 2002 *Acta Phys. Polon. B* **33** 1935 (*Preprint nucl-th/0204025*)
- [565] Velkovska J [PHENIX Collaboration] 2002 *Nucl. Phys. A* **698** 507 (*Preprint nucl-ex/0105012*)
- [566] Teaney D *et al* 2001 *Phys. Rev. Lett.* **86** 4783 (*Preprint nucl-th/0110037*)
- [567] Kolb P *et al* 2001 *Nucl. Phys. A* **696** 197
Huovinen P *et al* 2001 *Phys. Lett. B* **503** 58
- [568] Huovinen P *et al* 2001 *Phys. Lett. B* **503** 58 (*Preprint hep-ph/0101136*)
- [569] Broniowski W and Florkowski W 2001 *Phys. Rev. Lett.* **87** 272302 (*Preprint nucl-th/0106050*)
- [570] Vitev I and Gyulassy M 2002 *Phys. Rev. C* **65** 041902 (*Preprint nucl-th/0104066*)
- [571] Schaffner-Bielich J *et al* 2002 *Nucl. Phys. A* **705** 494 (*Preprint nucl-th/0108048*) (and references therein)
- [572] Mangano M L and Altarelli G (ed) 2000 *CERN Workshop on Standard Model Physics (and more) at the LHC, CERN Yellow Report CERN-2002-004* (*Preprint hep-ph/0003142*)
- [573] Ray L and Longacre R S 2000 *STAR Note* 419
- [574] <http://home.cern.ch/~radomski>
- [575] Radomski S and Foka Y 2002 *ALICE Internal Note* 2002-31
- [576] <http://home.cern.ch/~radomski/AliMevSim.html>
- [577] Ray L and Hoffmann G W 1996 *Phys. Rev. C* **54** 2582
Ray L and Hoffmann G W 1999 *Phys. Rev. C* **60** 014906
- [578] Skowroński P K 2003 *ALICE HBT* <http://aliweb.cern.ch/people/skowron>
- [579] Poskanzer A M and Voloshin S A 1998 *Phys. Rev. C* **58** 1671
- [580] Alscher A, Hencken K, Trautmann D and Baur G 1997 *Phys. Rev. A* **55** 396
- [581] Hencken K, Kharlov Y and Sadovskiy S 2002 *ALICE Internal Note* 2002-27
- [582] Marchesini G *et al* 1992 *Comput. Phys. Commun.* **67** 465
- [583] Paige F and Protopopescu S 1986 *BNL Report* BNL38034, unpublished
- [584] Engel R 1995 *Z. Phys. C* **66** 203
Engel R and Ranft J 1996 *Phys. Rev. D* **54** 4244
- [585] Sjostrand T and van Zijl Z M 1987 *Phys. Rev. D* **36** 2019

- [586] Abe F *et al* [CDF Collaboration] 1997 *Phys. Rev. Lett.* **79** 584
- [587] Alner G J *et al* [UA5 Collaboration] 1983 *Phys. Lett. B* **138** 304
- [588] Mangano M L and Altarelli G (ed) 2000 *CERN Workshop on Standard Model Physics (and more) at the LHC, CERN Yellow Report CERN-2002-004 (Preprint hep-ph/0003142)*
- [589] Abe F *et al* [CDF Collaboration] 1999 *Phys. Rev. D* **59** 32001
- [590] Norrbin E and Sjöstrand T 2000 *Eur. Phys. J.* **17** 137
- [591] Marchesini G and Webber B R 1988 *Nucl. Phys. B* **310** 461
- [592] Webber B R 1984 *Nucl. Phys. B* **238** 492
- [593] Collins P D B 1977 *Introduction to Regge Theory* (Cambridge: Cambridge University Press)
- [594] Particle Data Group 1996 *Phys. Rev. D* **54** 193
- [595] Amos N *et al* [E710 Collaboration] 1990 *Phys. Lett. B* **243** 158
Abe F *et al* [CDF Collaboration] 1994 *Phys. Rev. D* **50** 5550
- [596] Donachie A and Landshoff P V 1992 *Phys. Lett. B* **296** 227
Donachie A and Landshoff P V 1997 *Preprint hep-ph/9703366*
- [597] Cudell J R *et al* 2000 *Phys. Rev. D* **61** 034019
- [598] Esakia S M, Garsevanishvili V R, Jalagania T R, Kuratashvili G O and Tevzade Yu V 1997 *Preprint hep-ph/9704251*
- [599] Alner G J *et al* [UA5 Collaboration] 1986 *Phys. Lett. B* **167** 476
- [600] Tano V 2001 *Proc. 31st Int. Symp. on Multiparticles Dynamics (Datong, China) (Preprint hep-ph/0111412)*
Tano V 2002 *37th Rencontres de Moriond on QCD and Hadronic Interactions (Les Arcs, France, 16–23 March 2002) (Preprint hep-ex/0205023)*
- [601] Field R 2002 Minimum bias and the underlying event at the tevatron and the LHC *talk presented at the Institute for Fundamental Theory (University of Florida, 22 October)*
- [602] Moraes A *et al* 2003 Minimum bias interactions and the underlying event *Talk presented at Monte Carlo at Hadron Colliders Workshop (University of Durham, 14–17 January 2003)*