Institutt for fysikk og teknologi

The Readout Control Unit of the ALICE TPC

by
Jørgen A. Lien

October 2004

Universitetet i Bergen
Bergen, Norway

Institutt for fysikk og teknologi

The Readout Control Unit of the ALICE TPC

by
Jørgen A. Lien

a thesis submitted to Institutt for fysikk og teknologi,
Universitetet i Bergen,
in partial fulfilment of the requirements for
the degree of Doctor Scientiarum

October 2004
Universitetet i Bergen
Bergen, Norway

ISBN 82–497–0246–8

Abstract

The ALICE Time Projection Chamber (TPC) is the main tracking detector of the central barrel of the ALICE (A Large Ion Collider) Experiment at the Large Hadron Collider (LHC), being constructed at CERN, Geneva. It is a 88 m^3 cylinder filled with gas and divided into two drift regions by the central electrode located at its axial center. The readout chambers of the TPC are multi-wire proportional chambers with cathode pad readout. About 570 000 pads are read-out by an electronics chain of amplification, digitalization and pre-processing. One of the challenges in designing the TPC for ALICE is the design of Front End Electronics (FEE) to cope with the data rates and the channel occupancy. The Readout Control Unit (RCU), which is presented in this work, is designed to control and monitor the Front End Electronics, and to collect and ship data to the High Level Trigger and the Data Acquisition System, via the Detector Data Link (DDL - optical fibre). The RCU must be capable of reading out up to 200 Mbytes/s from the Front End Electronics, operating at a maximum of 40 MHz. The handling and distribution of the central trigger and clock signals are done with an incorporated TTCrx chip. Interfacing the front-end electronics with the Detector Control System (DCS) is done via a dedicated Slow Control bus. The DCS itself communicates with the control layer via high level protocols running on an Ethernet fieldbus. For the prototype of the Readout Control Unit (RCU) the Altera EP20K400 Field Programmable Gate Array (FPGA) has been used for application specific system integration.

Chapter 1 gives an overview of the physics observables and thus defines the requirements for the performance of the ALICE TPC and its FEE. In chapter 2 TPC technologies and their development and use in selected heavy ion experiments are discussed. At the end of the chapter the TPC of the ALICE experiment is presented. Chapter 3 gives an overview of the other detectors and the Trigger and Data Acquisition of the same experiment. Details of the readout chain of the ALICE TPC, including the Front End Electronics and the RCU, is the subject of chapter 4. Chapter 5 is a presentation of the first conceptual design of the logic of the RCU. This formed the basis for the development of the RCU, and is a major part of the work of this thesis. It is presented here in detail, even though a number of changes were made during later stages of the development cycle. A number of prototypes of the RCU have been developed, and are presented in chapter 6, along with their use in different test setups of other equipment of the ALICE detector. The work of this thesis includes the development of the first two prototypes, but some results from the ongoing RCU development are also presented. A summary and outlook is given in chapter 7. Two conference contributions originating from this work are included as appendices.

Acknowledgements

The work presented in this thesis has been carried out within a collaboration of the Microelectronics Research Group and the Section for Subatomic Physics, University of Bergen, and the Faculty of Engineering, Bergen University College during the period of January 2000 to December 2003.

I first of all would like to thank my supervisors Håvard Helstrup, Kjetil Ullaland and Terje Natås for their guidance and support during this period of work. Also I would like to thank Dieter Röhrich for all his help in making this thesis come to life.

The time spent at CERN in The Doctoral Student Programme in Engineering and Applied Science under the guidance of Luciano Musa will always be remembered.

A special thank goes to Werner Olsen for his constant belief in me and his patient explanations of all things electronic.

Ketil Røed, Knut Aurbakken, Torsten Alt, Anne Marie Sandvik, Kristin Fanebust Hetland, Anders Strand Vestbø, Jens Ivar Jørdre and Ulrich Frankenfeld are among the many people who has made my everyday at work worthwhile and filled with smiles. A szia and egészségedre goes to Csaba Soós for times well spent.

Petter Seip became much more than a colleague during my 4 years in Bergen. Thank you for giving me a home away from home, Petter.

Last, but not least, I would like to thank my parents Anny and Svein Roar for their support during this work. My aunt Kari also deserves thanks for all the practical work of designing a home out of my 2 room flat in Barliveien, Fyllingsdalen, and removing it just as efficiently at the end of term. However, the far most important result of this working period in Bergen is my new family. My girlfriend Marianne, who is now also my wife, and my daughter Marte, truly deserves the biggest hugs of you all.

Åsgårdstrand, October 2004

Jørgen A. Lien

Contents

List of Figures

List of Tables

Chapter 1

Studying nuclear matter with heavy ion collisions

1.1 Introduction

Inside the nucleus there are protons and neutrons. There are even smaller particles inside the protons and neutrons - quarks. Different types of quarks exist - named up, down, strange, charm, top and bottom. Protons always have two "up" quarks and one "down" quark, while neutrons have two "down" quarks and one "up" quark. The quarks are strongly bound together inside the nucleons by the strong force mediated by the gluons. In fact they are so strongly bound that no single quark has ever been observed.

1.2 The Quark-Gluon Plasma

Even if we know what the basic building blocks of the universe are, we still have a long way to go until we can describe the complex properties of matter and its various manifestations. Nuclear matter, as described above, is only one possible manifestation. At very high densities and temperatures, the nucleons are expected to dissolve into their constituents and to form a plasma consisting of quarks and gluons - the so-called Quark-Gluon Plasma (QGP). All the matter in the universe is expected to have been in the form of a Quark-Gluon Plasma, when the universe was less than one millisecond old, just after the Big-Bang. Other phases might exist in the interior of neutron stars.

1.3 The challenge of looking for Quark-Gluon Plasma in the laboratory

By colliding two heavy nuclei at ultra-relativistic energies a bulk system with huge density, pressure and temperature is created. This should free the quarks and the gluons of the system into a small bubble of Quark-Gluon Plasma. In this state quarks can interact freely with a large number of other quarks, not confined in "bags" as they are in normal nuclear matter. The Quark-Gluon Plasma, if produced in this way, is however only expected to live for a very short time, about 10^{-23}s. Studying a system so small at these energies

demands the use of large accelerator and detector facilities, as available in SPS (CERN), RHIC (Brookhaven) and the future LHC (CERN).

Heavy ion physics is however not limited to the search for the Quark-Gluon Plasma only. More generally, the focus of this research is to study and understand how collective phenomena and macroscopic properties emerge from the microscopic laws of elementary particle physics.

1.3.1 The phase diagram of nuclear matter

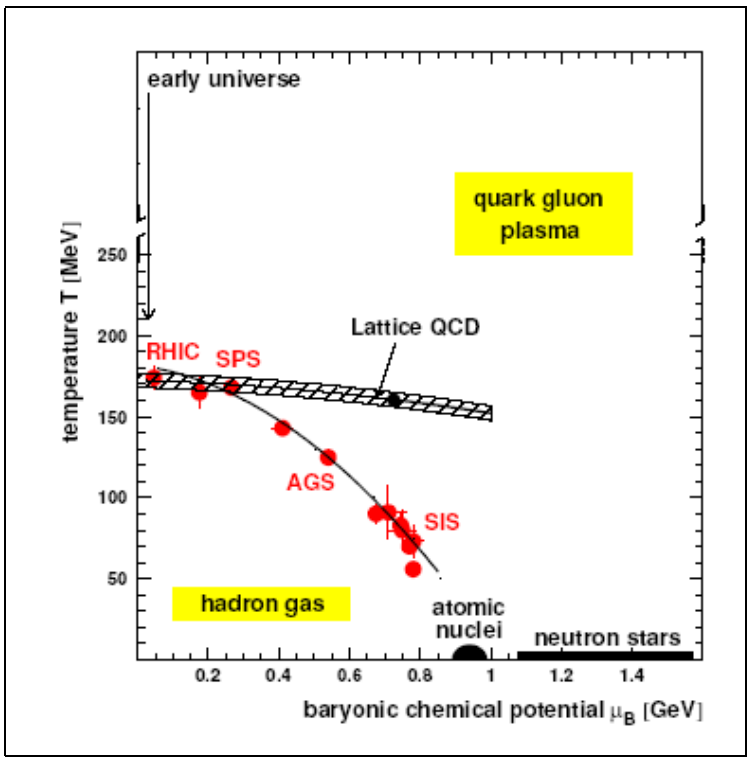

Figure 1.1: Phase diagram summarizing the present understanding about the structure of nuclear matter at different densities and temperatures. The points illustrates the results achieved by the different ultra-relativistic collider experiments, and the dashed line represents the Lattice QCD calculations.

The present experimental and theoretical knowledge about the different phases of strongly interacting matter can be summarized in a generic phase diagram, as shown in figure 1.1. Nuclear matter exists in different phases as a function of temperature and density. At higher temperatures and small net baryon density a cross-over transition from hadronic matter to quark-gluon matter takes place. The transition temperature is about 170 MeV [1]. In highly compressed cold nuclear matter - as it may exist in the interior of neutron stars - the baryons lose their identity and dissolve into quarks and gluons. The data points represent the end points of the evolution of the hot and dense matter created in the collision. Their locations indicate that the boundary is at least reached.

1.3.2 Space-time evolution of the collision of two heavy ions in the laboratory

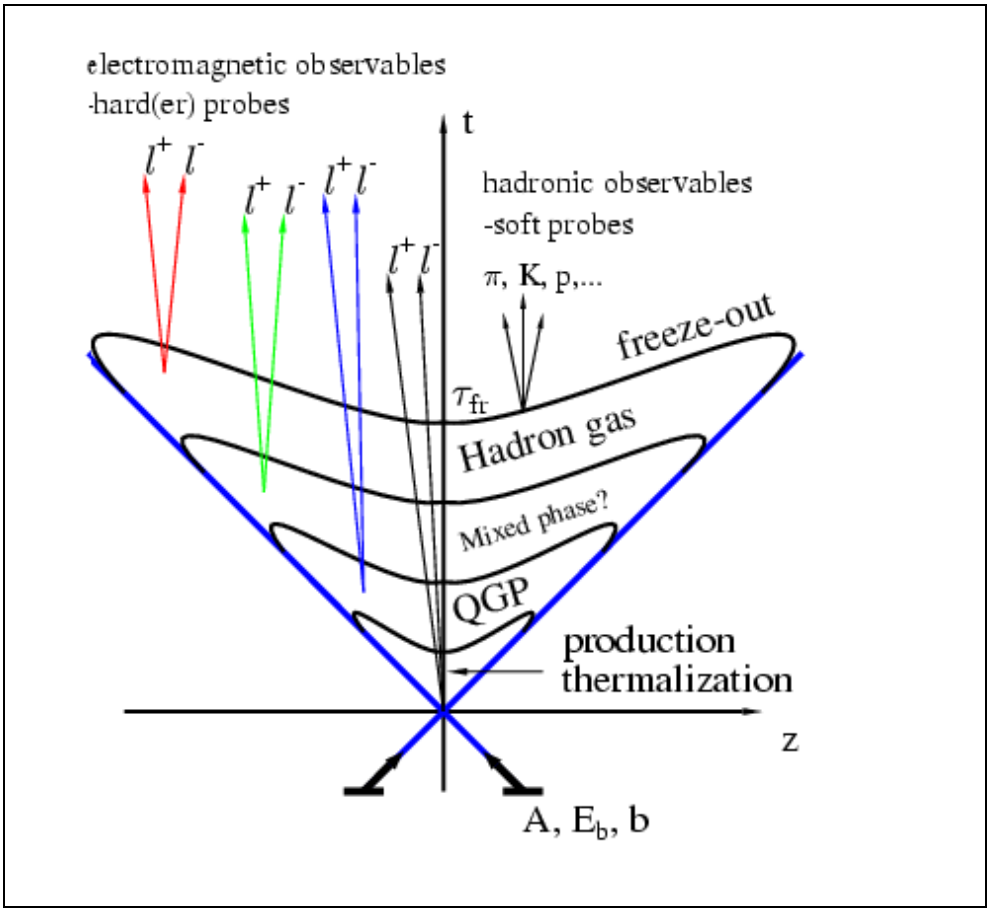

Figure 1.2: Space-time evolution of a strongly interacting system produced in a very high energy A+A collision. The projection is in the plane of the longitudinal coordinate z and time t. Different types of observables and the different stages of evolution they can probe are indicated [2].

The different stages in the space-time evolution of a system produced in ultra-relativistic heavy-ion collisions can be pictured as in figure 1.2. The colliding nuclei are Lorentz-contracted thin disks. The original impact of the nuclear disks takes place almost instantaneously in a region around z, $t \sim 0$. After a rather short equilibration time $\tau_0 \simeq 1$ fm, the presence of a thermalized medium is assumed, and for sufficiently high initial energy densities, this medium would be in the quark-gluon plasma phase. The produced QGP system expands and cools. Eventually the system ends up in a Hadron Gas phase. If the QCD-phase transition is of first order, a stage of a mixed phase with both QGP and HG coexisting at a (local) critical temperature and chemical potential will appear in between. In the end of the HG phase, the average mean free paths of hadrons exceed the homogeneity size of the system, and the Hadron Gas freezes out.

A large number of experiments have been devoted to the study of the nuclear state of matter as a function of temperature and density, and in particular to the search for

a QGP phase, during the last decades. At the CERN SPS accelerator a series of fixed target experiments has collected a wealth of information about nuclear matter at nucleon-nucleon center-of-mass energies of $\sqrt{s_{NN}}$ = 5-20 GeV. A review at the conclusion of the experiments summarized the achievements as that through "a multitude of different observations compelling evidence" was given that a new state of matter has been created [3]. Several reported results are shown in this thesis. However, the energy density obtained only slightly exceeds the critical temperature, T_c.

Currently, four experiments (STAR, BRAHMS, PHENIX and PHOBOS) at RHIC (Relativistic Heavy Ion Collider) [6] at Brookhaven (USA) are dedicated to the study of Au-Au collisions at center-of-mass energies up to $\sqrt{s_{NN}}$ = 200 GeV.

The LHC (Large Hadron Collider) currently being constructed at CERN, will make Pb-Pb nuclei collide at $\sqrt{s_{NN}}$ = 5.5 TeV. These collisions will be studied by ALICE (A Large Ion Collider Experiment). This experiment is presented in further detail later in this thesis. At the energies reached at LHC, nuclear matter is predicted to be transparent enough to form baryon-free matter heated up well beyond the expected phase transition temperature.

1.4 Experimental observables

Colliding two heavy ions at relativistic energies creates a complicated system which evolves through many stages before particles are registered in the detectors. These particles have often undergone many interactions with other particles, or are results of decays. A number of observables are used to try to investigate these highly dynamic, complex systems. A few of them are presented below. It is the sum of results from all observables taken from different energy regimes and system sizes (comparing proton-proton, proton-heavy ion, heavy ion-heavy ion collisions etc.) that will give new insights to the properties of exited nuclear matter. The observables in heavy ion collisions can be divided into four main categories:

- Global event characteristics
- Hadronic observables
- Electromagnetic observables
- Hard probes

1.4.1 Global event characteristics

Numerical calculations of QCD on a lattice predict a phase transition from hadronic matter to deconfined matter at energy densities $\sim$ 1.0 GeV/fm^3 [1]. This is approximately 5 times the density of normal nuclear matter. The energy going into the system created in a heavy ion collision is depending on the centrality of the collision. The collisions taking place in an accelerator is "by accident". Two bunches of heavy ions are circling with opposite directions in the accelerator ring, and by moving them as close together as possible at the center of the detector some ions will eventually collide. But there is no way to ensure that the collision is head-on. There will thus be a number of collisions

taking place, and the distributions of central to non-central collisions will be as illustrated in figure 1.3.

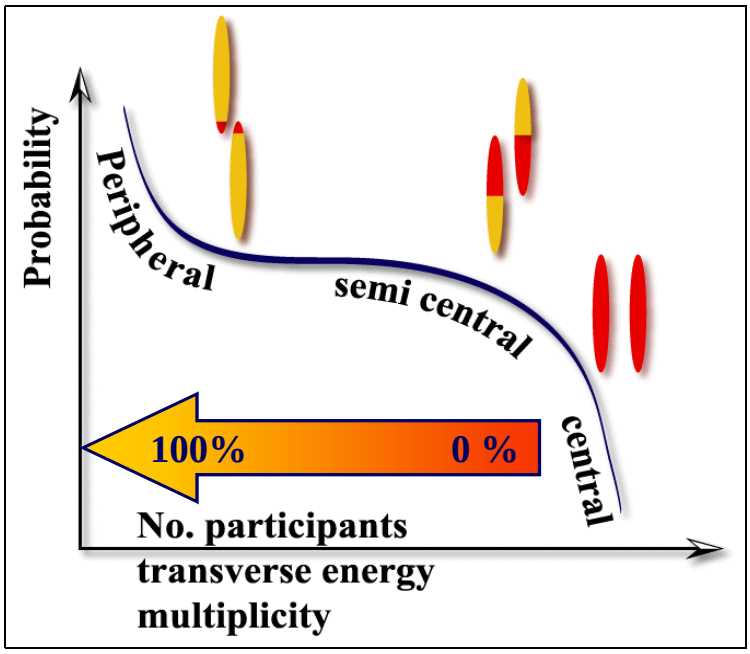

Figure 1.3: Probability distribution of collisions as a function of centrality [4].

Particle multiplicities

The particle multiplicity is a function of the centrality of the collision. Centrality is therefore usually determined in terms of a certain cut in the multiplicity distribution; e. g. a 0 - 10 % centrality cut means that out of a given sample, 10% of the events which have the highest multiplicity have been selected.

The particle multiplicity is also a crucial parameter for the design and performance of the detectors. One of the major challenges for the ALICE experiment in general, and the design of its TPC and Read Out electronics in particular, is to cope with high particle multiplicities and the corresponding large data flow in the read out system. The particle multiplicity in heavy ion collisions are challenging to predict, since it cannot be calculated from first principles. By extrapolating measured quantities from other experiments (at other energies) and using these as inputs to different theoretical models, one gets predictions for the multiplicities to be expected in future experiments. In figure 1.4 data from RHIC and predictions for center of mass energies up to $\sqrt{s_{NN}} = 5.5$ TeV are shown [36]. At RHIC the charged particle multiplicity per unit pseudorapidity, $dN_{ch}/d\eta$, is measured to be 7-800 at $\eta = 0$. The predictions for LHC show that one should expect a value of $\approx$ 2200. Other calculations in [5], where the multiplicity is computed in a two-component soft+semihard string model, gives a slightly higher density of 2600-3200, depending on initial assumptions.

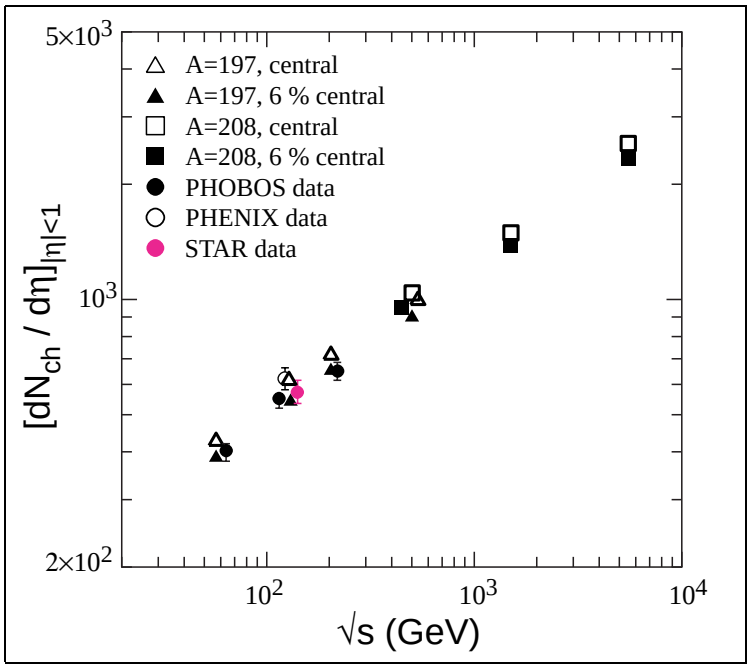

Figure 1.4: Data and predictions for charged particle multiplicity per unit pseudorapidity [36].

Energy density

By using information on the size of the system created in heavy ion collisions (from HBT and flow measuements, see later), combined with measurements of transverse energy, one can use a formula proposed by Bjorken [12] to approximate the energy density in the collision system. At SPS NA49 has estimated the energy density to 3.2 GeV/fm^3 [13]. For RHIC corresponding calculations from PHENIX gives a density of 4.6 GeV/fm^3 [14]. These estimates are well above the predictions for the quoted critical energy density from QCD on a lattice ($\sim$ 1.0 GeV/fm^3) needed for the creation of a QGP phase.

1.4.2 Hadronic observables

The hadrons emerge only in the final stage of the collision after they freeze out from the hadron gas. They carry information about the system at the time of freeze out. Particle production, particle abundancies, spectra and correlations are examples of such observables.

Particle spectra

Most of the particles emitted in heavy ion collisions are hadrons which decouple from the collision region during the hadronic freeze out stage. By measuring the different particle spectra, one obtains information about the chemical and kinematical freeze out conditions. Quantities like the freeze out temperature and chemical potential, radial flow velocity, directed and elliptic flow, the size of the system etc. can then be derived from these observations.

Strangeness production

It has long been suggested [7] that strangeness production (see figure 1.5) is a sensitive probe of a deconfined state. For example, strangeness production may be enhanced by the fast and energetically favorable process of gluon-gluon fusion into strange quark-antiquark pairs or $q\overline{q}$ annihilation. The required energy is $2m_s \approx 300$ MeV.

Hadronic mechanisms, on the other hand, may also enhance strangeness production. In a hadron gas strangeness is produced (and destroyed in the reverse reactions) in processes like $\pi + \pi \rightarrow K + \overline{K}$ and $\pi + N \rightarrow \Lambda + K$. The reaction with the lowest threshold energy ($\sim$ 670 MeV) is the reaction $N + N \rightarrow \Lambda + K + N$.

A systematic investigation of strangeness production is therefore needed to understand different production mechanisms and to discriminate between the quark-gluon plasma and hadron gas scenario in heavy ion physics. Furthermore, enhancement from a QGP source

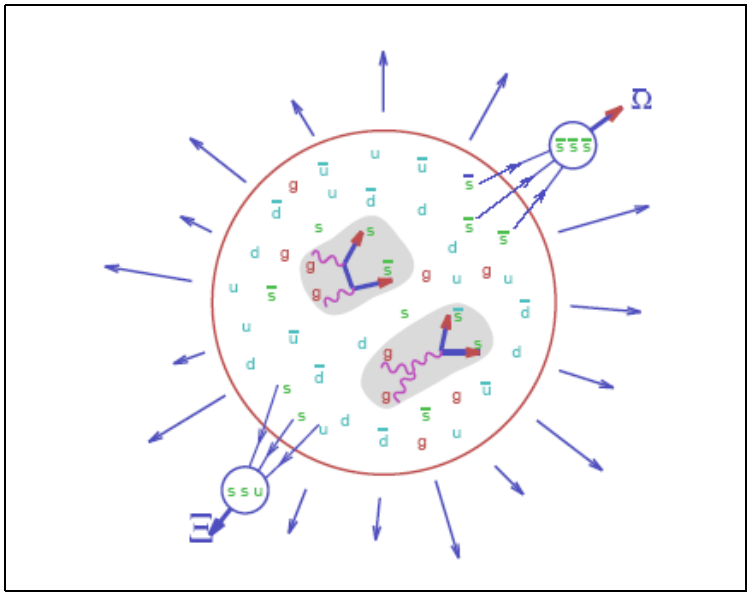

Figure 1.5: Illustration of the two step mechanism of strange hadron formation from QGP [8]. First the gluon fusion into strange- anti-strange quark pairs. Later the recombination during the hadronization.

is expected to increase with the strangeness content of the particle, making the Ω (made of 3 strange quarks) the most sensitive strangeness probe. Such enhancement has been measured by the WA97/NA57 experiments [9], as shown in figure 1.6.

In order to determine whether strangeness is enhanced in heavy ion collisions and to ascertain effects of baryon density on strangeness production, the strange and non-strange particle production as a function of energy and system size is studied. About 70% of the strangeness produced goes into kaons, and the pions are the most frequently detected non-strange particle in heavy ion collisions. In figure 1.7 the K/π ratio is shown. Its energy dependence shows a maximum at center-of-mass energy around 7-9 GeV. In general, particle ratios provide input to thermal models, which describe the chemical freeze-out of the system in terms of a temperature and a baryochemical potential. Typical values obtained at RHIC are T$\approx$170 MeV and $\mu \approx$20-30 MeV.

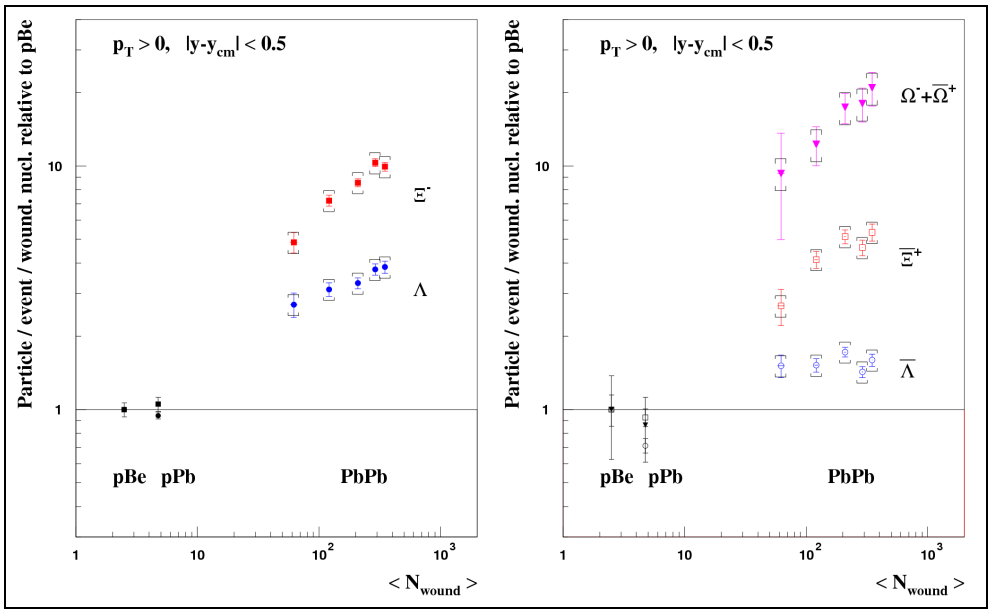

Figure 1.6: Hyperon enhancements as a function of the number of wounded nucleons at 158 A GeV/c [9].

Particle correlations and Hanbury-Twiss-Brown interferometry

Correlations of identical particles with close velocities (small relative momenta) arise as a consequence of a quantum mechanical interference effect. A two-particle *correlation function* $R(p_1, p_2)$ can be defined as a ratio of the measured distribution of the momentum difference $\mathbf{q} = \mathbf{p_1} - \mathbf{p_2}$ using particles from the same event to the reference one which would be the corresponding distribution calculated from particles from different events. For non-interacting identical particles the effect of quantum interference leads to the presence of interference maximum or minimum at small relative momenta.

The correlation function can be parameterized into R_{out},R_{side} and R_{long} components of the relative momentum vector $\overrightarrow{q}$=q_x, q_y, q_z [17]. These HBT radii can be interpreted as (see figure 1.8):

- R_{long} - a measure of the longitudinal width of the emission region. Direction parallel to the beam direction (in the longitudinal co-moving system frame).
- R_{side} - a measure of the width of the emission region in the side direction. Direction perpendicular to beam and total pair momentum $\mathbf{k}$.
- R_{out} - a measure of a combination of the extension in the out direction and the emission duration. Direction perpendicular to q_l and q_s.

Another important parameter is the incoherence parameter λ which measures the strength of the correlation.

The use of HBT to measure sizes of stars works under the assumption of static sources. The particles emitted after heavy ion collisions, however, come from systems which feature strong dynamical evolution both along the beam direction and transverse to it. As a result, two-particle correlation measurements in heavy-ion physics exhibit a much richer

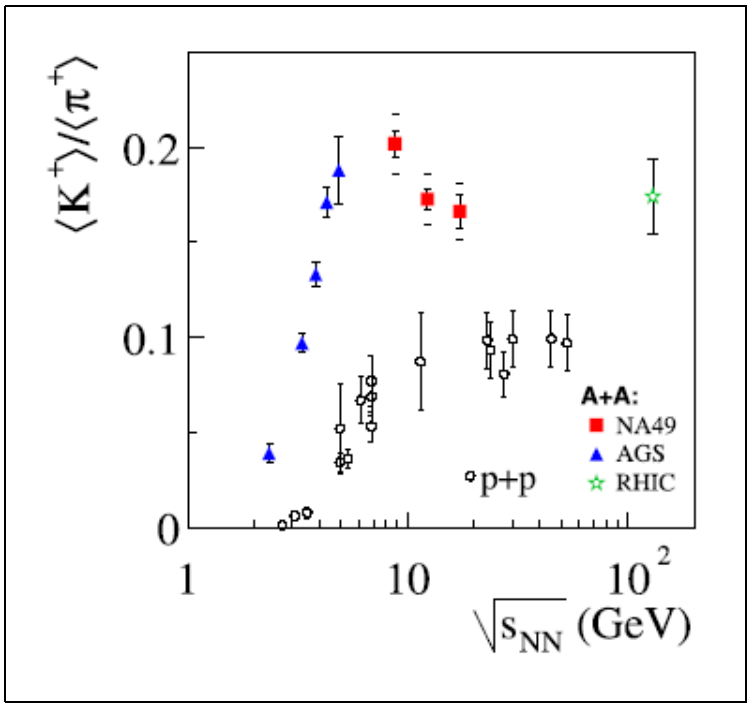

Figure 1.7: Energy dependence of K/π ratio in p+p and central A+A collisions from AGS, NA49 and RHIC [10].

structure, but their interpretation is also considerably more involved. The measurements of the R_{out},R_{side} and R_{long} contains information not only about the geometry of the system fireball (size at freeze-out), but also of the flow and dynamics of the system. The hot and dense early stages are reflected in the measured freeze out pattern. Also included are information on the lifetime of the system (sensitive to nature of phase transition) [18].

Figure 1.9 shows a data set of correlation parameters for mid rapidity, low-p_t ($\sim$ 170 MeV/c) π^- from central Au + Au or Pb + Pb collisions. The parameter λ falls smoothly and rapidly from unity at $\sqrt{s_{NN}} \sim 2$ GeV, to about 0.5 at RHIC energies. This decrease is attributed partially to an increased fraction of π^- arising from long-lived decays as $\sqrt{s_{NN}}$ increases [22]. λ could also be effected by several experiment-specific effects, e.g. e^- rejection efficiency. There is no major increase in the other correlation parameters

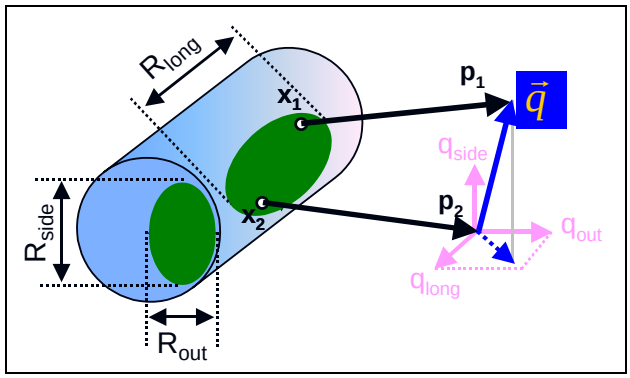

Figure 1.8: Definition of the R-components in two-particle interferometry.

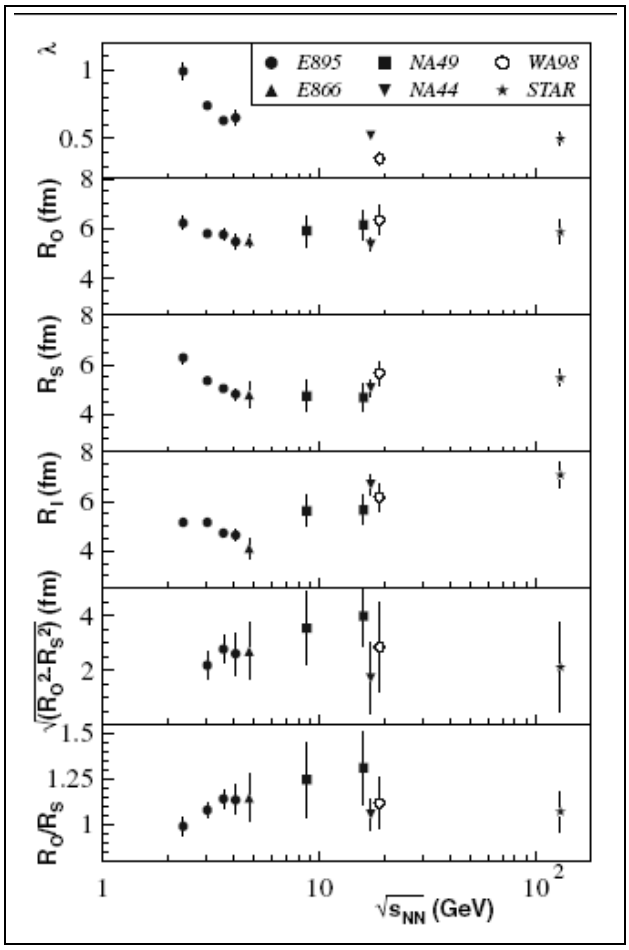

Figure 1.9: Compilation of HBT parameters from experiments as indicated. The energy dependence of π^- correlation parameters for central Au + Au (Pb + Pb) collisions at mid rapidity and $p_t \approx 170$ MeV/c. The SPS data are offset slightly in $\sqrt{s_{NN}}$ for clarity. Error bars on NA44, NA49 and STAR results include systematic uncertainties. Error bars on other results are statistical [19].

going from SPS to RHIC energies, which would reflect an increased time-scale of pion freeze-out. Such an increase has been predicted as a consequence of QGP formation [23].

Most of the emitted particles from a heavy ion collision are pions, the lightest hadrons. However, the high multiplicity and abundance of other particles at SPS and RHIC energies has made it possible to study also:

- Two-particle correlation for other identical particles, like kaons [20], to compare radius parameters.
- Two-particle correlation for non-identical particles, like $K-\pi$ correlations [21]. This gives a possibility to discern which particle species is emitted first and also details of the spatial flow field.
- Three-pion correlations, to seek evidence of coherence in the source [18].

Radial flow

The particles of the system created in a heavy ion collision undergoes numerous rescatterings which leads to a build-up of pressure. *Radial flow* is the term used to denote the collective expansion of the system driven by this pressure. If the matter elements of the system move with a common velocity, then heavier particles will be shifted to higher transverse momenta. Then the flow velocity can be deduced by examining transverse momentum spectra for particles of different masses, by the use of models of the system dynamics. With the blast-wave model [25], a hydrodynamically motivated model, the transverse expansion velocity β_T has been extracted. At SPS such fits gives $\beta_T \sim 0.4-0.5$, and at RHIC $\beta_T \sim 0.5 - 0.6$.

Anisotropic flow

Another important observable characterizing the collective expansion is the *Anisotropic flow* [26]. It is defined as azimuthal asymmetry in particle distribution with respect to the reaction plane (the plane spanned by the beam direction and the impact parameter). It is called flow since it is a collective phenomenon. The anisotropic flow is characterized using Fourier decomposition of the azimuthal distributions. The first harmonic Fourier coefficient, v_1, describes *directed flow*, and the second harmonic coefficient, v_2, corresponds to *elliptic flow*. Non-zero higher harmonics can also be present in the distribution.

While the radial flow is largest in the most central collisions, the elliptic flow depends on the initial asymmetry of the system, and thus is largest for non-central collisions. In a non-central collision, the region of interaction is almond-shaped. With enough reinteraction, the system builds up pressure transforming the initial asymmetry present in coordinate space into a momentum anisotropy. In order for such a momentum anisotropy to be observed, the pressure must build up before the system has expanded and the initial asymmetry has disappeared. Therefore, the elliptic flow reflects the amount of reinteraction in the early times of the collision dynamics.

In figure 1.10 the elliptic flow measured at RHIC [27] is shown. The elliptic flow (v_2) was measured to be ~ 0.03 in mid-central collisions at SPS, and ~ 0.06 at RHIC. This suggests early equilibration at RHIC.

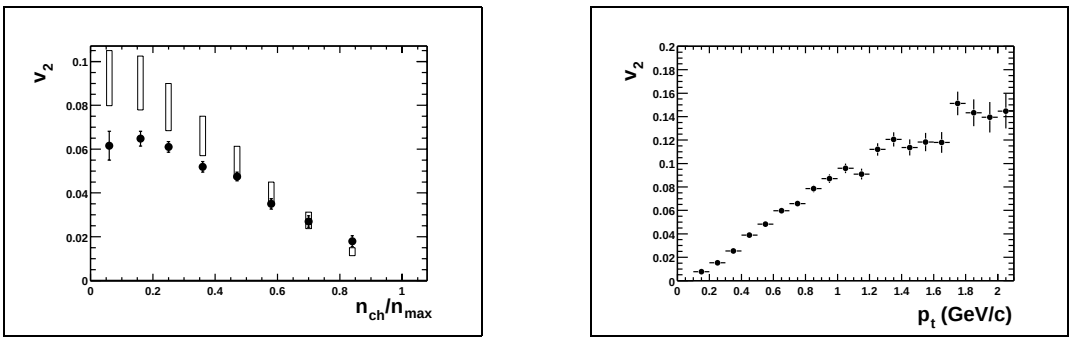

Figure 1.10: Measured elliptic flow at RHIC. Left: Elliptic flow as a function of centrality. The open rectangles show a range of values expected in the hydrodynamical limit. Right: Elliptic flow as a function of transverse momentum for minimum bias events. [27]

Fluctuations

Large multiplicities in single central events allow the study of fluctuations of different observables on event-by-event basis [36]. The primary goal of event-by-event physics is to observe differences between the events with and without QGP formation. The QCD phase transition may cause large fluctuations which will show up in measurable quantities like multiplicities, $< p_t >$ or particle ratios in single events.

1.4.3 Electromagnetic observables

Direct photons

The measurement of photons is motivated by the possibility of:

- Understanding the thermal properties of the early phase of the reaction from *thermal photons*.
- Study the production of *prompt photons* due to hard processes in the dense medium.

Direct photons have very little interaction with the surrounding medium since they are not affected by the strong force. They are therefore not altered (by rescattering) during the final stages of the evolution of the collision. Thus they are a probe of the earliest stages of the evolution.

Measurement of such photons are complicated, especially at low transverse momentum p_t, because of the large background from radiative decay of neutral pions ($\pi^0 \rightarrow \gamma\gamma$). Results from the WA98 experiment at SPS indicates that the task of extracting the direct photons is feasible [39]. The PHENIX collaboration at RHIC reported on direct photon measurements at Quark Matter 2004. The results are shown in figure 1.11.

Dielectrons

Dielectrons are produced during the whole space-time evolution of the interaction by different processes. Due to their electromagnetic character, they can leave the interaction region without further strong interactions and, therefore, the dielectron spectrum measured is the folding of the production rates from each stage of the collision. A careful analysis of the resulting dielectron invariant mass spectrum should, in principle, allow to unfold the whole spacetime history of the nuclear collision.

CERES/NA45 result on dileptons

CERES/NA45 [34] has reported on measurements of correlated electron-positron pairs, also called dileptons. NA45 finds in Pb-Au collisions an excess of dileptons in the mass region between 250 and 700 MeV, by a factor 2.6 $\pm$ 0.5 (statistical) $\pm$ 0.6 (systematic) above expectations from hadron decays scaled from p+p to Pb+Au collisions, as shown in figure 1.12. The measurements could be explained e.g. as a broadening of the ρ peak, resulting from scattering among pions and nucleons in a very dense hadronic fireball, just below the critical energy density for quark-gluon plasma formation.

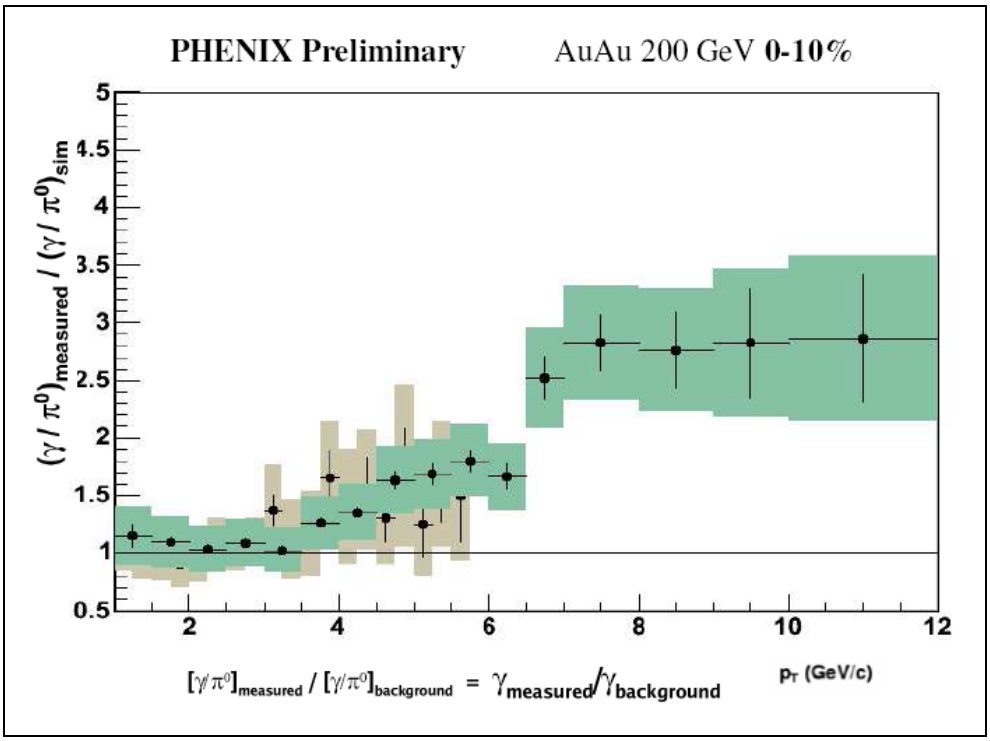

Figure 1.11: Direct photons in RHIC Au+Au collision with PHENIX, reported at QM 2004.

1.4.4 Hard probes

Occasionally hard scatterings between partons of incoming protons and neutrons occur. These are interactions where massive particles and/or particles at high p_t are produced. Hard scattering occurs early in the collision, well before quark matter is expected to form. The scattered partons will sense the full space-time evolution of the collision volume and thus probe the later formed hot and dense phase. In contrast to the hadronic observables, the hard probes involve only a limited number of colliding partons and are theoretically treated by perturbative QCD.

Jets

In the initial stage of the collision, before the hot, dense fireball of the system is formed, quarks and gluons from the incoming nuclei can scatter with high momentum transfer. Scattering with large momentum transfer will result in either the production of high mass particles, like charm quarks (see later), or in high momentum quarks or gluons which will later fragment into jets of hadrons (see figure 1.13).

If the scattered partons penetrate quark matter (like QGP), they will radiate energy. The final fragmentation of the partons into jets of hadrons is then modified relative to the situation in free space, exhibiting reduced jet energies and jets with fewer high p_t particles (This is called *jet quenching*). Therefore, measurements of the properties of jets produced by hard scattered partons in a heavy ion collision can serve as a probe of the produced medium. Since the hard scattered partons experience the entire space-time of the system, this is a promising new probe in the field of heavy ion collisions.

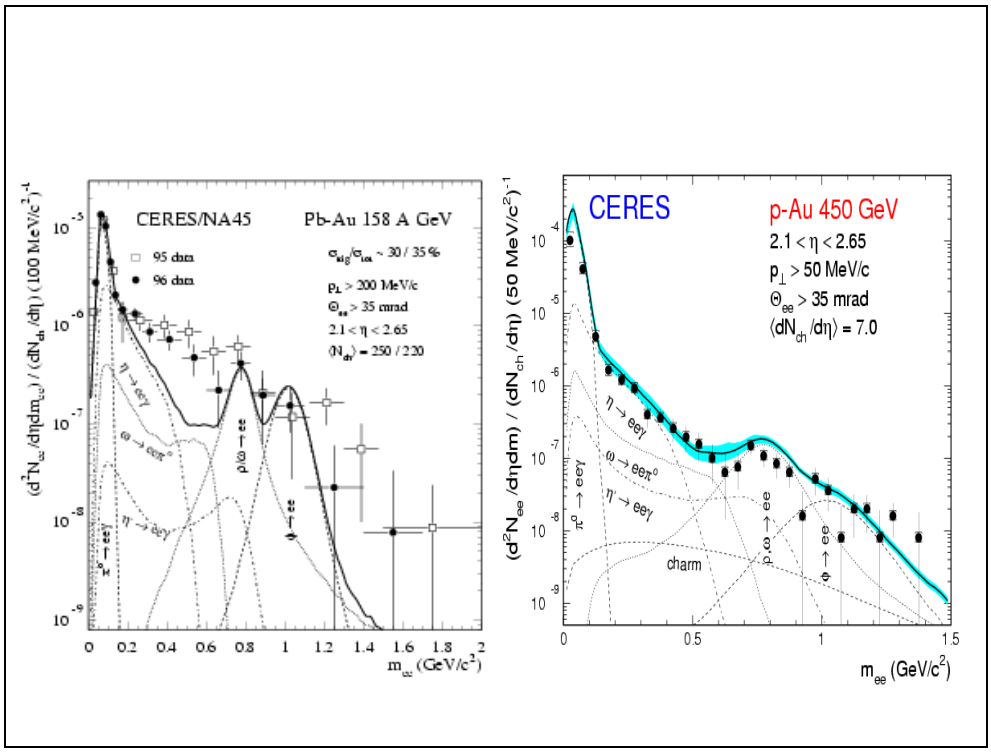

Figure 1.12: The dilepton spectrum in Pb-Au collisions at 158 A GeV from CERES (left). The solid line displays the expected yield from known hadronic sources. The figure on the right shows the same measurements for p-Au collisions [35].

Evidence of final-state suppression of high p_t hadrons in Au+Au collisions has been presented by the different RHIC experiments. In order to discriminate between initial state (gluon saturation) and final state (partonic energy loss) effects, data was taken of d+Au collisions. In this type of collisions there will not be produced any dense matter (QGP) for the partons to traverse, in effect turning off any final state effects.

Direct measurement of jets of hadrons resulting from parton fragmentation is difficult in heavy ion collisions due to the high multiplicity environment. A method to detect the presence of jets in such environments is via two-particle angular correlation measurements.

The STAR experiment has observed in the azimuthal distribution that the back-to-back ($\Delta\phi \sim \pi$) correlation peak is strongly suppressed in central Au+Au collisions

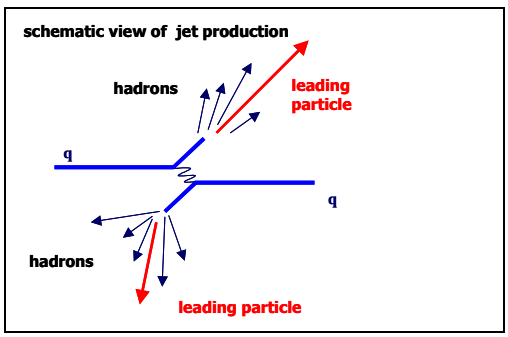

Figure 1.13: Schematic view of jet production.

relative to p+p and d+Au. The near-side ($\Delta\phi \sim 0$) peak is similar in all three collision types, as seen in figure 1.14. This could be interpreted as a jet pair originating from near the surface of the dense medium produced in the Au+Au collision, where the far-side (suppressed) jet has passed through the matter while the near-side has not.

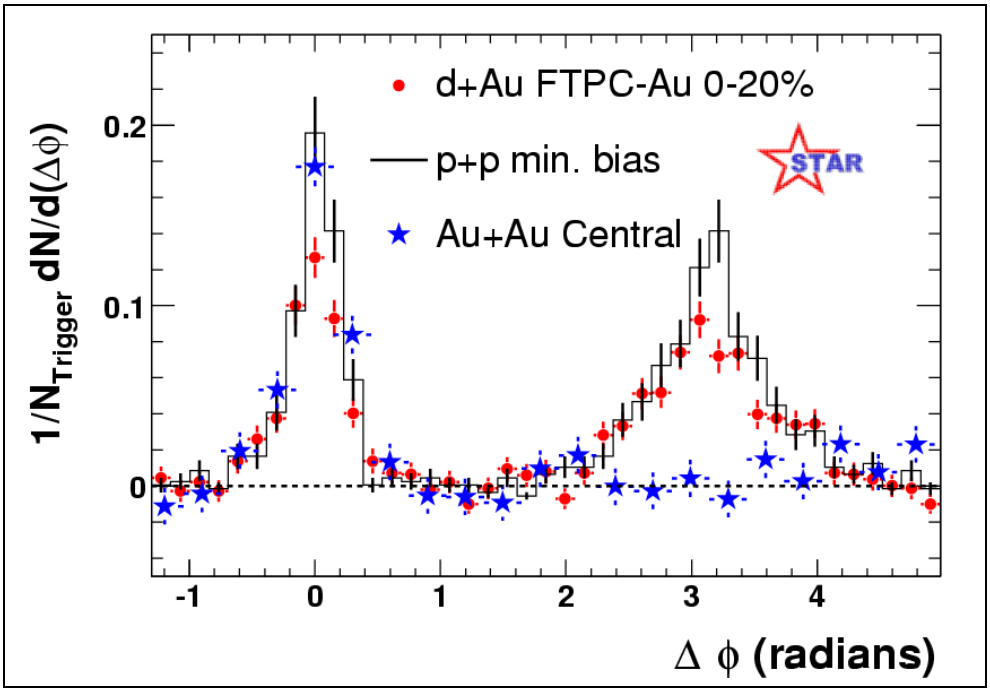

Figure 1.14: Comparison of two-particle azimuthal correlation distributions for central Au+Au collisions to those seen in p+p and central d+Au collisions [28].

If one assumes that the hadrons measured at high p_t emanate from jets, there should be a reduced yield of high momentum hadrons in central A+A collisions. High p_t hadron production are reported by all four experiments at RHIC, and are shown in figure 1.15. The nuclear modification factor (R_{AB}) used in figure 1.15 is defined as the ratio of hadron yield in nucleus-nucleus collisions and the yield of nucleon-nucleon collisions scaled by the number of binary collisions. In absence of nuclear medium effects the nuclear modification factor $R_{AB} \sim 1$. A suppression of high p_t hadrons is observed in central Au+Au collisions. The enhancement of high p_t hadron production in the d+Au collisions indicates that this is not an initial state effect.

Heavy quark production

Heavy quarks like charm and bottom provide a probe which is sensitive to the collision dynamics at both short and long time scales. At SPS only the charmonium states are experimentally accessible. The much higher LHC energy offers the possibility to measure the bottonium yields. The heavy quark production take place shortly after the collision (on the time scale of the order of the inverse of the quark mass), and is a perturbative phenomenon. The relatively long lifetime of the charm and bottom quarks allows them to live through the thermalization phase of the QGP, and to be affected by its presence. Transverse momentum distributions should display the energy loss suffered by a quark

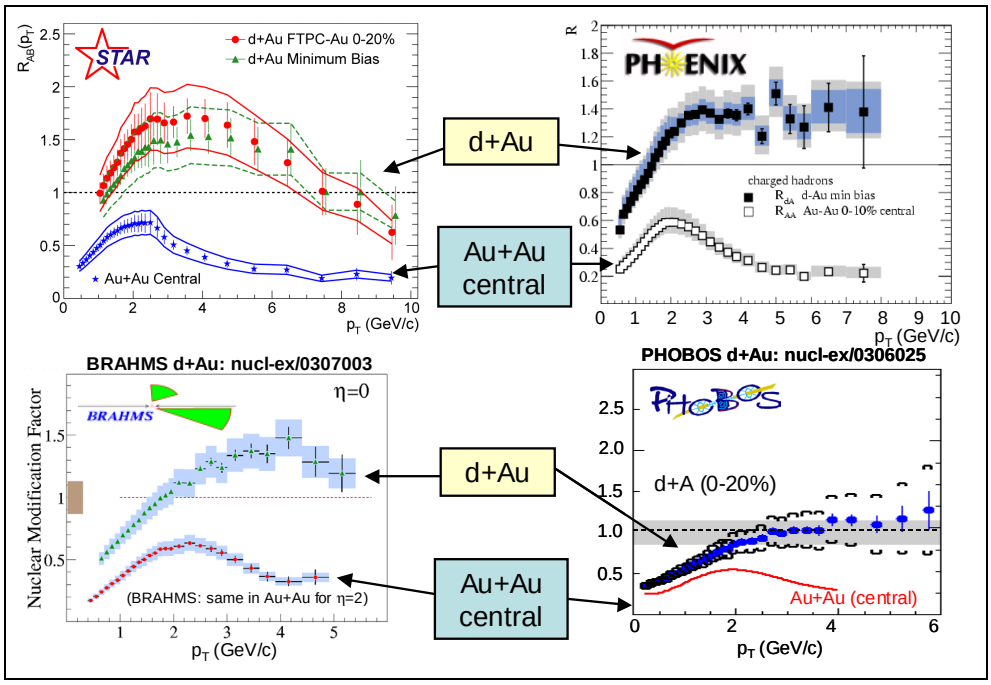

Figure 1.15: Comparison of charged hadron production in d+Au and Au+Au from all 4 RHIC experiments. The nuclear modification factor R_{AB} is plotted as a function of p_t [29].

passing through the plasma. Other important observables are the total production rates and kinematical correlations between the heavy quark and anti-quark. These observables must be compared to measurements from p+p and p+A collisions.

J/Ψ and other quarkonium production

In a deconfined medium the J/Ψ is expected to melt through shielding of the $c\overline{c}$ binding potential by color screening [30], and therefore QGP production should lead to a suppression of J/Ψ-production in nuclear collisions, compared to the rates extrapolated from p+p data. The NA50-experiment [31] at CERN has reported on the J/Ψ-production. They measured muon pairs. These pairs could originate from J/Ψ decaying into two muons or from quark-antiquark annihilation into dileptons, known as Drell-Yan production. The results show that the reference measurement, Drell-Yan events, are produced according to their expected rate from proton-proton up to Pb-Pb reactions. J/Ψ production was found to be suppressed in O-U, S-U and Pb-Pb collisions. The suppression always increases with centrality. By studying p-A collisions it was observed that also normal nuclear matter leads to reduced charmonium production. Extrapolating this to A+B interactions is enough to account for the observed yields up to central S-U collisions. In central Pb-Pb collisions, however, an additional 'anomalous' suppression was observed. Peripheral Pb-Pb collisions follow the normal pattern. Then, with increasing centrality, there is a pronounced onset of further suppression. Centrality is measured either through the associated transverse energy (using an electromagnetic calorimeter) or through the number of participating nucleons (using a zero-degree calorimeter). NA50 results shown

in figure 1.16, shows no visible saturation of the J/Ψ/DY ratio at high E_T.

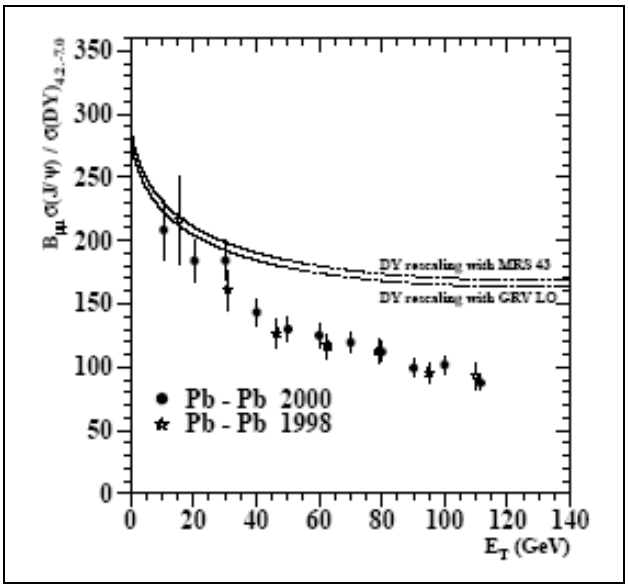

Figure 1.16: Results from the NA50 experiment at CERN SPS [32]. The ratio of J/Ψ to Drell-Yan production in Pb-Pb collisions vs. E_T. The solid lines indicates the extrapolation of the normal nuclear absorption inferred from p-A collisions.

The various models used to explain the J/Ψ-production at SPS energies can be extrapolated to higher energies. The results indicate severe suppression, making even the observation of J/Ψ's problematic [42]. However, the J/Ψ-production in RHIC could be enhanced rather than suppressed. At the high energy of RHIC (and LHC), several $c\overline{c}$ pairs may be present in the same system. These have a finite probability to recombine leading to a J/Ψ yield growing as the square of the number of $c\overline{c}$ pairs produced. Various calculations suggest that this may overcome the expected severe suppression [33].

Because the Υ (1S) dissolves only significantly above the critical temperature, at a value which might only be reachable at an energy above the current one provided by RHIC, the spectroscopy of the Υ family at LHC energies should reveal an unique set of information on the characteristics of the QGP.

1.4.5 Proton-proton physics and proton-nucleus physics

Most of the heavy ion observables require pp measurements of the same observables for comparison. In this way genuine collective effects in A+A collisions can be separated from phenomena already present in pp collisions.

Uncertainty in the initial conditions of the A+A collisions makes it a challenge to separate initial from final state effects. The observables measured in p+A collisions are mostly sensitive to the 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. The measurement of the production and suppression mechanisms in p+A is needed for the interpretation of open charm and charmonium data from A+A collisions. In summary, p+p collision and p+A (or d+A) collisions will provide a compulsory benchmark for the interpretation of the AA data.

Chapter 2

Time Projection Chambers (TPC) in Heavy Ion experiments

In this chapter the main uses (tracking and particle identification via dE/dx) of the TPC in heavy ion experiments is presented. The working principle for the multi wire proportional chamber (MWPC), which are used by all TPCs in heavy ion experiments, is illustrated. A few of the TPCs are presented, along with a short description of the experiment they are used in. Finally the ALICE TPC is presented in further detail, reserving the explanation of the front end readout electronics to the next chapter.

2.1 TPC working principle

A TPC is a three-dimensional tracking detector capable of providing information on many points of a particle track. It also provides information on the specific energy loss, dE/dx, of the particles traversing the drift volume, thus taking the role as an *electronic bubble chamber*. A TPC can be divided into three main parts:

- The drift volume
- The readout chambers
- The readout electronics

The charged particles produced in the collision pass through a gas mixture in the drift volume and releases electrons from the gas atoms. These electrons drift through an external electric field to the readout chambers. The drifting electrons avalanche in the high fields at the anode wires providing an amplification of the electric signal. The electrons disappear immediately. The positive ions created in the avalanche induce a temporary image charge on the readout pads which disappears as the ions move away from the anode wire. The image charge is measured by the readout electronics.

2.1.1 Drift volume

The drift volume is filled with a gas mixture and embedded in an electric field. Some TPCs are operated inside a magnetic field. There is a number of parameters that can to

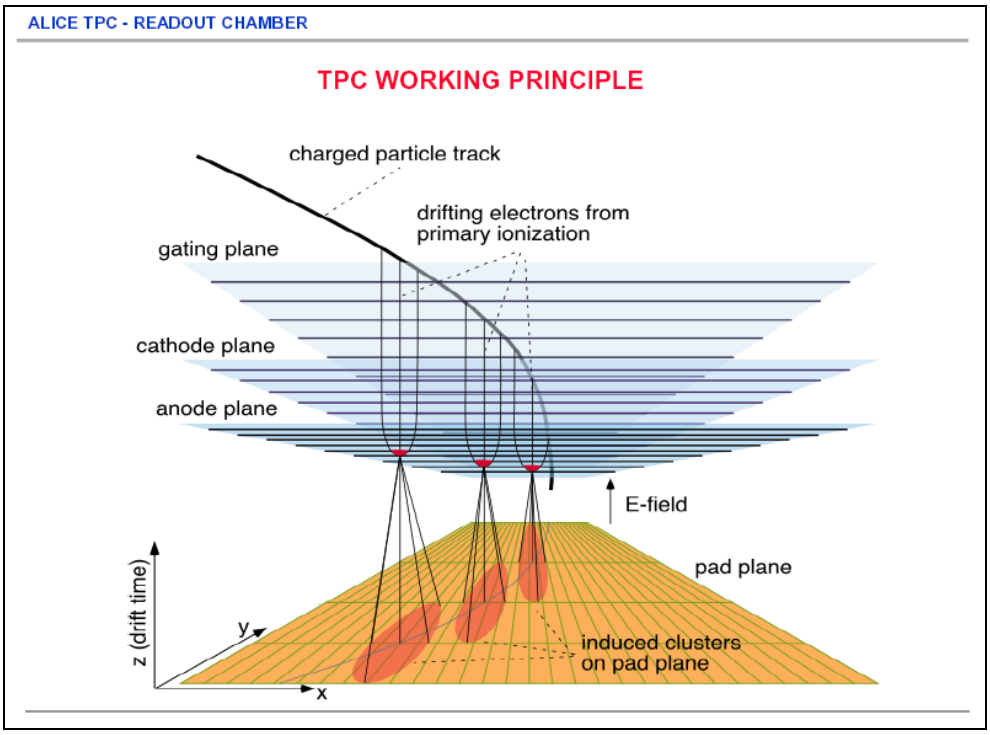

Figure 2.1: The working principle of the ALICE TPC. The gating grid, cathode and anode wires are all in the readout chambers (MWPC chambers with pad readout) [73].

be optimized regarding the choice of gas mixture and the field strengths. These features are discussed below.

Electric drift field

In the presence of an electric field, the electrons and ions freed by ionizing radiation are accelerated along the field lines towards the anode and cathode respectively. This acceleration is interrupted by collisions with the gas molecules which limit the velocity that can be attained by the charge along the field direction. The resulting average velocity attained is known as the *drift velocity* of the charge and is superimposed upon its random movement. Compared to their thermal velocities, the drift speed of the ions is slow. For electrons it can be much higher since they are much lighter.

Magnetic field

Charged particles are bent along a track segment in the presence of a magnetic field. The curvature of the trajectory in the known magnetic field can be used to measure the momentum of the particle. This information can be used to identify the particle. The presence of a magnetic field also influences the drift and diffusion of the electrons.

Gas mixture

The choice of gas mixture [74] in a TPC is based on a number of features:

- **Drift velocity** - Modern, large TPCs have a dead-time dominated by the time needed to collect the ionization electrons from the TPC volume, not by the time the electronics takes to process the data. Therefore a high electron drift velocity is desired. Pressure and temperature variations of the gas modify the density and thus the drift velocity. Special concern has been taken to protect the drift volume from the influence of any heat sources that could cause local and time-related temperature changes of the drift gas (by more than 0.1 degree Celsius) for the ALICE TPC [44]).
- **Electron diffusion** - During the drift the electrons are subject to diffusion.
- **High gain and onset of avalanche at low field** - The onset of the avalanche of electrons in the readout chambers is preferred at a low field. A high gain of the avalanche is desired as this improves the resolution.
- **Low electron attachment** - Losses of electrons between the point where the ionization occurs and the readout chambers affects negatively both the position information and the dE/dx particle identification ability.
- **Mobility of ions** - The total duration of the signals in the readout chambers is determined by the ion mobility. A high ion mobility thus helps to speed up the readout of the chambers.
- **Gas density** - Low gas density and high radiation length reduces multiple scattering of the charged particles traversing the gas.
- **Ageing** - In particular organic components like methane (CH_4) and ethane (C_2H_6) are prone to chemical ageing.
- **Safety** - Inflammable, toxic and radioactive components in the gas mixture impose complicated and expensive safety measures.

A dedicated gas system circulates the gas and maintains purity. The main impurities which accumulate in the system are oxygen and water which capture drifting electrons.

2.1.2 Readout chambers

The chambers consist of a pad plane and three or four wire planes, as illustrated in figure 2.1. Drifting electrons originating from the primary ionization by themselves do not induce a sufficiently large signal in the readout pads. The necessary signal amplification is provided by avalanche creation in the vicinity of the anode wires. The induced charge from an avalanche is shared over several adjacent pads (2-3). The primary purpose of the ground grid is to terminate the field in the avalanche region and provide additional rf shielding for the pads. The outermost wire plane is the gating grid. This grid is a shutter to control entry of electrons from the TPC drift volume into the readout chambers. It also blocks positive ions produced in the readout chamber, keeping them from entering the drift volume where they would distort the drift field. The positive ions are too slow to escape during the open period and get captured during the closed period. The switching of the potential of the gating grid from closed to open mode is controlled by the experiment trigger(s).

Optimalization of the pad and wire layout

The three wire planes have every wire stretched under proper tension, sufficient to ensure that they will not deflect from their nominal positions oriented parallel to the pad rows. The diameter of the wires, the relative positioning and the size, shape and positioning of the pads are, along with the choice of drift gas, the main design parameters that are subject to optimization. In radial TPC (like the forward TPCs of the STAR experiment), the wires are tilted with a small angle so that two or more anode wires cross each pad for the selected pad-wire geometry. This is to avoid the observed periodic shifts in the position measurement that arises when the wire planes are orthogonal to the axial direction of the pads [68]. In linear drift TPCs (i.e. central STAR and ALICE) the wire planes are mounted orthogonal to the pads. The pad size is sometimes different in the outer and inner sectors. The smaller wire spacing and pad size of the inner sectors provides better spatial resolution necessary for tracking resolution near the center of the detector where track density is highest.

Pad response function

The pad response function (PRF) relates the signal induced on a pad to the position of a track parallel to the length of a pad. The PRF depends on both the wire geometry and the size and shape of the pads. The width of the PRF is a key measure of the position resolution of a TPC [57]. Calculations of PRFs have guided in optimalization of readout chambers during the design of TPCs [44].

2.1.3 Track reconstruction

The track reconstruction is done by first finding clusters of charge on neighboring pads. Using an appropriate center-of-gravity algorithm, a precise measurement of the location of the avalanche, in the x and y coordinate, can be obtained. Then one can reconstruct the track since the drift time is known, and the time hits are sampled on the readout pads (gives the z coordinate).

2.1.4 Resolution of TPC

The most desired attributes are dE/dx resolution, momentum resolution and two-track resolution. To guide in the construction phase simulations of track reconstruction and efficiency are carried out, using for instance an event generator like HIJING [58] combined with a tool to simulate particles traversing detector media, like GEANT [59].

Momentum resolution

The momentum of a track can be determined by fitting a helix through the points along the track. The radius of the curvature and the emission angles of the track at the interaction point is used to calculate the momentum [56].

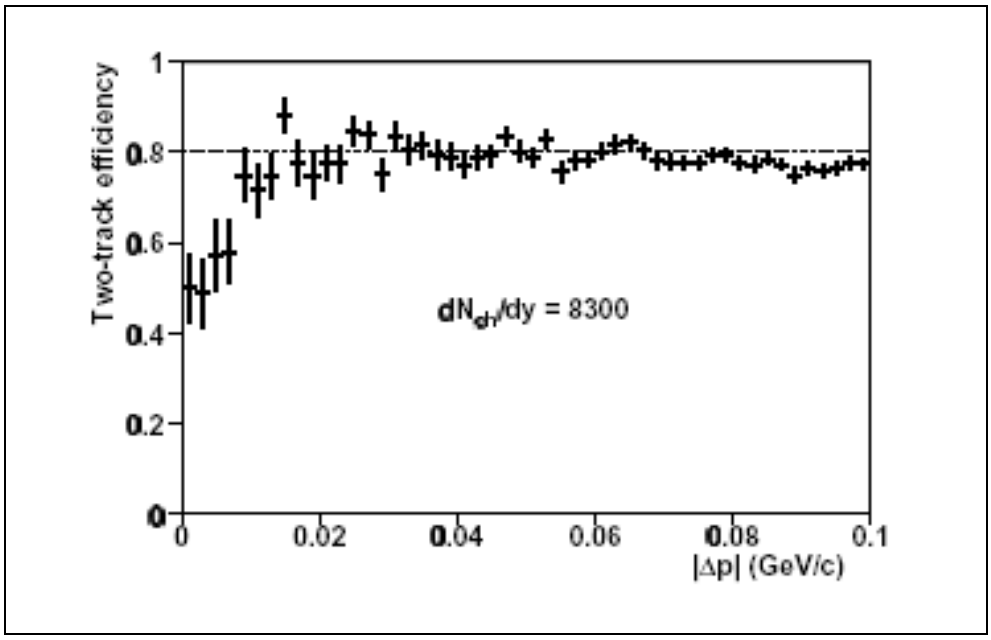

Figure 2.2: Two-track efficiency for particle density dN_{ch}/dy=8300 as a function of absolute value of generated momentum difference of two particles for the ALICE TPC [44].

dE/dx resolution

The charge deposited in the clusters are used to obtain dE/dx information, by calculating the truncated mean over the length of the track.

Two-track resolution

The separation of pairs of close, neighboring tracks is a major concern in the high track density regions of Pb+Pb interactions. In these environments there is a high probability to record hits on single pads during the drift time originating from multiple tracks. This can be tested in simulation by generating particles with small differences in momentum. Fig. 2.2 shows the results of such a simulation for the ALICE TPC.

2.2 The TPCs in NA49

The NA49 detector [60] is a wide acceptance spectrometer for the study of hadron production in p+p, p+A, and A+A collisions at the CERN SPS (fixed target). The main components are 4 large volume TPCs for tracking and particle identification via dE/dx. TOF (Time Of Flight) scintillator arrays complement particle identification. Calorimeters for transverse energy determination and triggering, a detector for centrality selection in p+A collisions, and beam definition detectors complete the setup.

2.2.1 Geometry and drift volume

The TPC system deploys two "Vertex" chambers (VTPC) inside the magnets and two "Main" chambers (MTPC) on both sides of the beam behind the magnets, as seen in

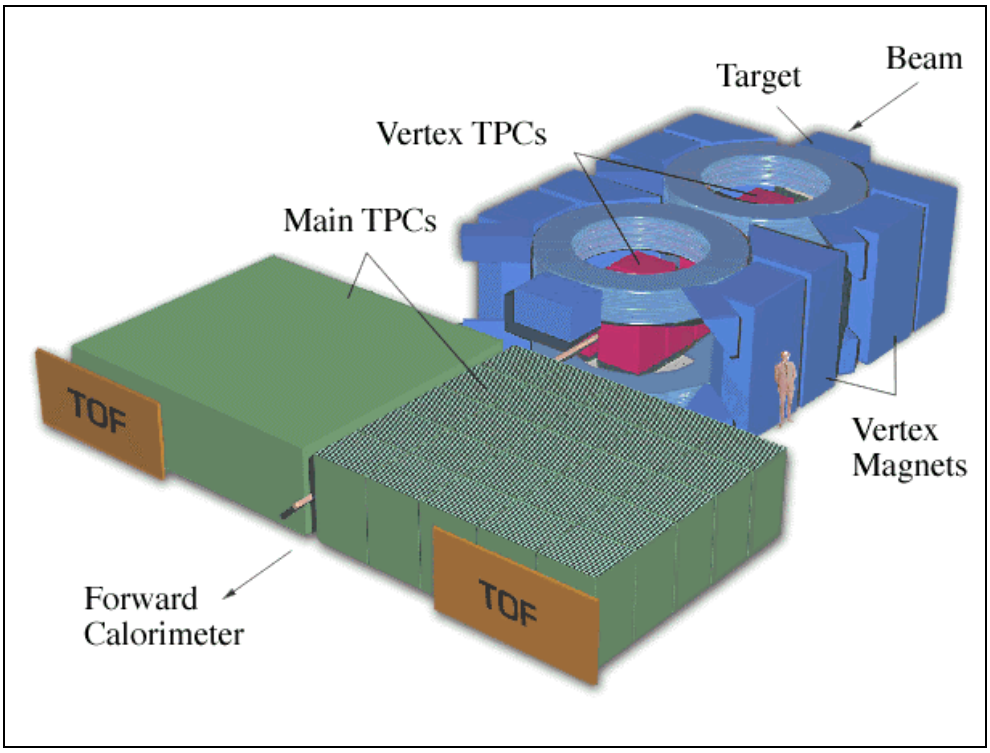

Figure 2.3: Perspective view of the NA49 TPCs.

figure 2.3. The VTPCs measures 2.0 x 2.5 x 0.98 m^3. These are used to derive the particle momentum especially for tracks emerging from secondary vertices. The two "Main" chamber TPCs have a length of 3.9 m. They are placed behind the magnets to measure the ionization energy loss in order to meet the requirements of the dE/dx resolution. Their width is 3.9 m and the height is 1.8 m.

Magnetic field

The two VTPCs are placed inside two dipole magnets with a field measuring 1.5 T and 1.1 T.

Electric drift field

The electric drift field of the Vertex chambers is defined by the HV electrode at a potential of -13 kV and the cathode wires of the readout chambers at ground potential. Drift fields of 200 V/cm for the VTPCs and 175 V/cm for the MTPCs corresponds to drift velocities of 1.4 cm/μs and 2.5 cm/μs. The drift time is thus about 50 μs. The voltage dividers, necessary for a well defined electric field, are constructed with 25 μm aluminized Mylar strips. The divider strips are 12.7 mm wide.

Gas properties

A mixture of 90% Ne and 10% CO_2 is used for the VTPCs, which is the gas mixture also chosen for ALICE (see section 2.6.1). Maximum drift length is 666 mm. A mixture of

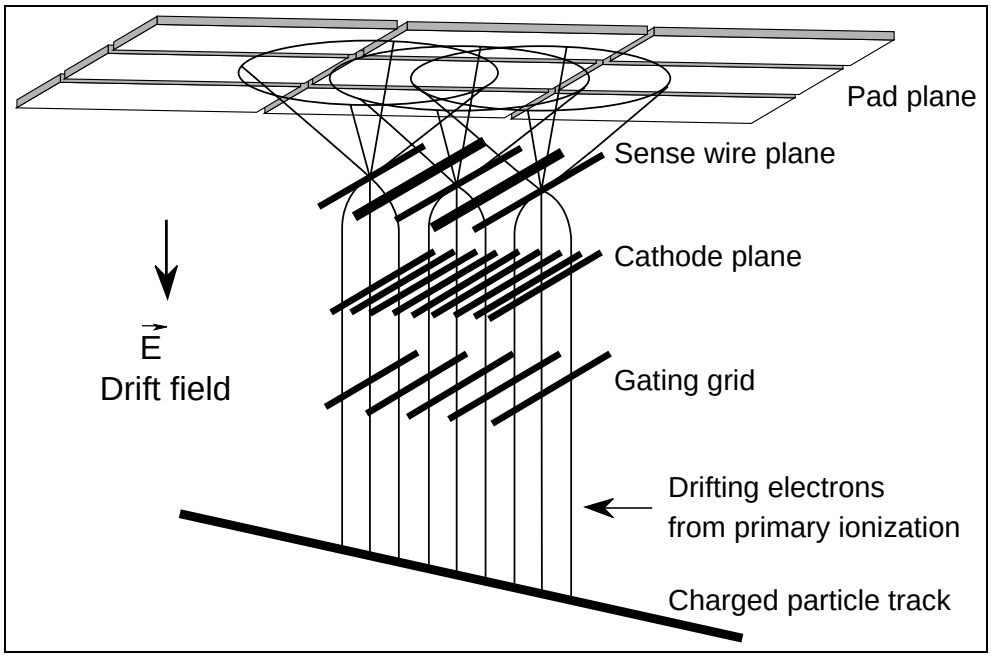

Figure 2.4: Schematic layout of the NA49 TPC readout chambers. The drift direction is vertically upwards. The sense wires are interspaced with field (zero potential) wires [60].

90% Ar, 5% CH_4 and 5% CO_2 is used for the MTPCs. Maximum drift length is 1117 mm. The detector gases have been selected for their low charge diffusion coefficients, providing narrow cluster charge distributions at the read-out plane. The drift velocities are such that the maximum drift times in the VTPC and MTPC are approximately equal.

2.2.2 Readout chambers

The total NA49 TPC system has 62 readout proportional chambers. Their working principle is illustrated in figure 2.4. The pad sizes are given in table 2.1, together with wire diameters and positions.

2.2.3 Readout system including front end electronics

The TPC readout system comprises a total of 182000 pads. 512 time samples are recorded for the analog signal of each channel over the maximum drift time of 50 μs. The analog signals are stored in Switched Capacitor Arrays (SCA), which can be regarded as a series of 512 capacitors with a sample and hold circuit for each capacitor [62]. After the readout cycle of 50 μs, the stored charges are digitized in ADCs with 9 bit resolution.

2.2.4 Selected results

A scatter plot of truncated energy loss against momentum is shown in the left part of figure 2.5 for positive secondaries in p+p and p+Pb collisions. The pions and electrons are measured through the relativistic rise of the Bethe-Bloch function and up to the Fermi plateau.

Dimensions (mm)	VTPC-1	VTPC-2	MTPCs	Wire material
Pad length	16, 28	28	40	
Pad width	3.5	3.5	3.6, 5.5	
Pad angles	12-55°	3-20°	0°, 15°	
Pad/sense wire dist.	3	2	2, 3	
Sense wire diam.	0.02	0.02	0.02	W-Re (gold plated)
Sense wire spacing	4	4	4	
Field wire diam.	0.125	0.125	0.125	Cu-Be (gold plated)
Field wire spacing	4	4	4	
Cathode wire diam.	0.075	0.075	0.075	Cu-Be
Cathode wire spacing	1	1	1	
Gating wire diam.	0.075	0.075	0.075	Cu-Be
Gating wire spacing	1	1	2	

Table 2.1: Wire and pad parameters of the NA49 TPC system [60].

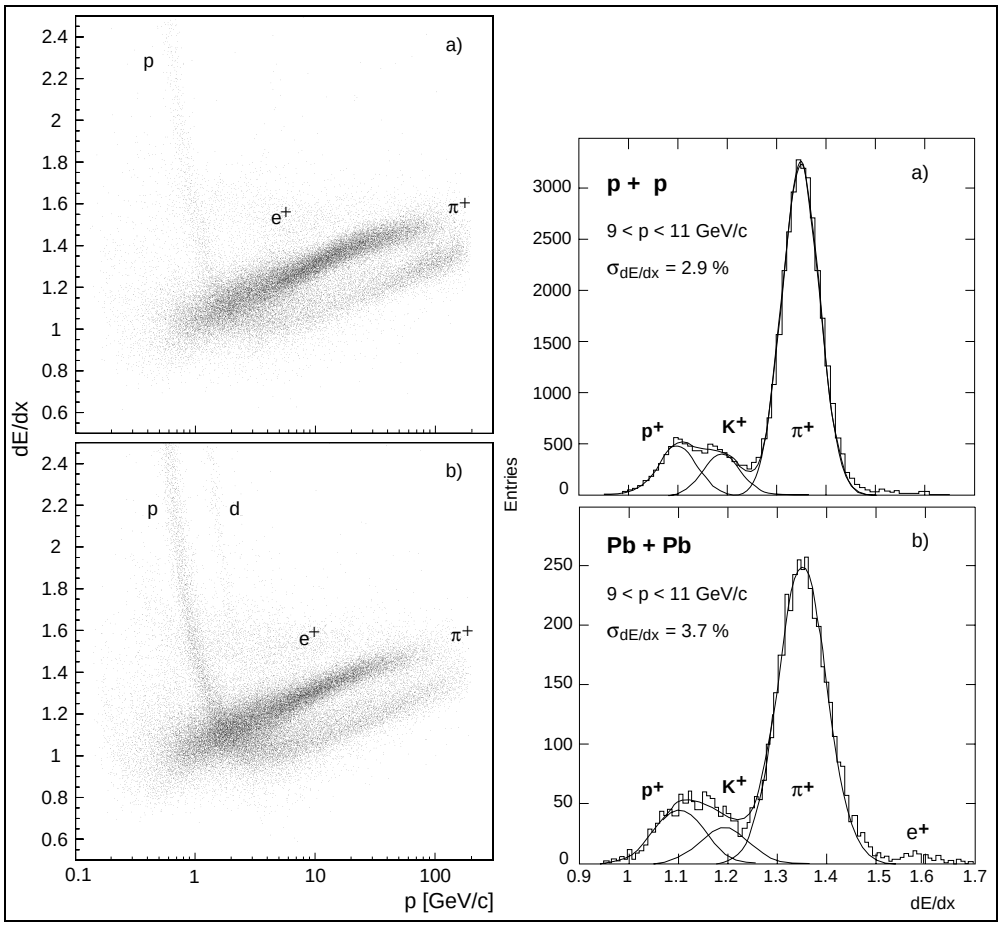

Figure 2.5: dE/dx as a function of momentum for positive particles produced in a) p+p and b) p+Pb collisions in NA49 (left). dE/dx distribution at a fixed momentum of 10 GeV/c for a) p+p interactions and b) central Pb+Pb collisions in NA49 (right) [60].

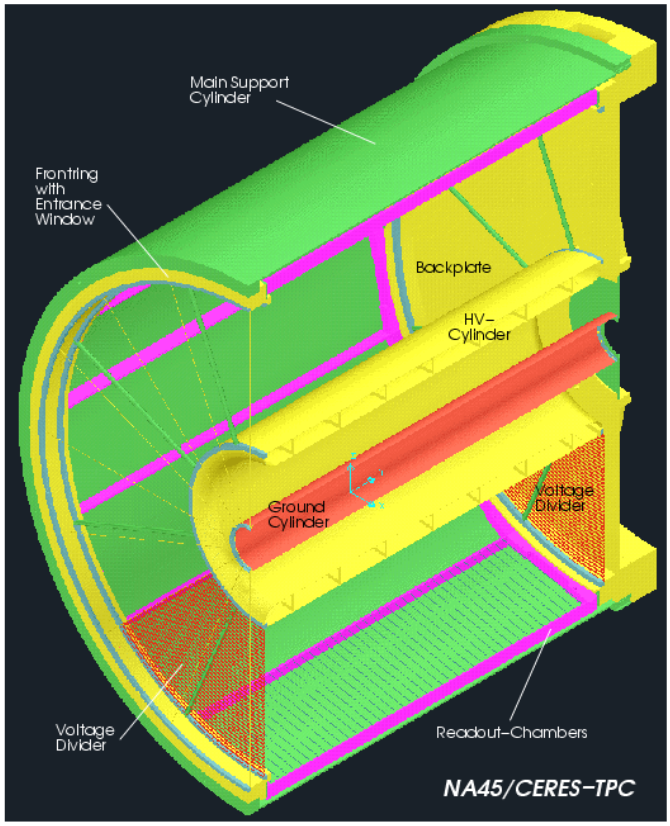

Figure 2.6: Perspective view of the CERES/NA45 TPC.

The right part of figure 2.5 shows a projection around 10 GeV/c momentum. Good separation of electrons, pions and protons is demonstrated.

2.3 The TPC in CERES/NA45

CERES (Cherenkov-Ring Electron Spectrometer) [34] is a dilepton experiment at the CERN SPS (fixed target). The experiment was upgraded with a radial drift TPC to improve the mass resolution (at m$\approx$ 1 GeV/c^2).

2.3.1 Geometry and drift volume

The CERES TPC (figure 2.6) is a cylindrical drift chamber with the drift field in radial direction and segmented pad readout. The sensitive volume is about 9 m^3 and the length 2 m.

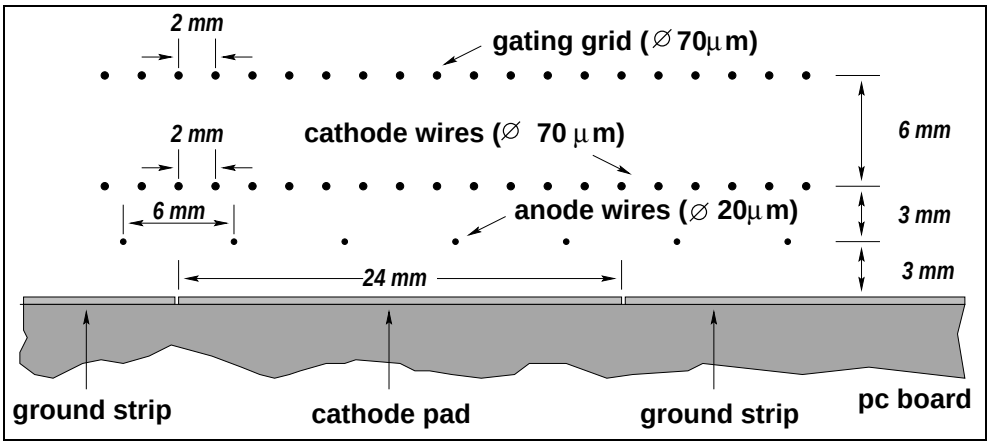

Figure 2.7: Cross section of the CERES TPC readout chamber [61].

Magnetic field

The magnetic field is generated by two warm coils with current floating in opposite directions. The radial component of this field is maximal between the two coils and the deflection of charged particles is mainly in azimuthal direction. The field integral is 0.18 Tm at a polar angle of θ=8° and 0.38 Tm at a polar angle of θ=15°.

Electric drift field

The electric drift field is defined by the inner electrode at a potential of -30 kV and the cathode wires of the readout chambers at ground potential. The resulting field shows approximately a $1/r$ dependence. The associated drift velocities are between 2.4 and 0.7 cm/μs. The maximum drift time is about 71 μs. The voltage dividers, necessary for a well defined electric field, are constructed with 100-200 nm thick copper strips on both sides of the foils enclosing the end caps of the drift volume (see figure 2.6). The divider strips are 15 mm wide and the distance between them is 5 mm.

Gas properties

A mixture of 80% Ne and 20% CO_2 is used.

2.3.2 Readout chambers

Figure 2.7 shows the cross section of the CERES/NA45 TPC readout chamber including wire diameters and positions. The pads are shaped like a chevron, see figure 2.8. The usage of this type of pads enables a precise determination of the charge centroid due to charge sharing between neighboring pads, even for relatively large pads [63].

2.3.3 Readout electronics

The CERES TPC has a granularity of 15360 channels, recording 256 time samples each. Approximately 1000 events during a spill period of 5 s are read out. Online data reduction

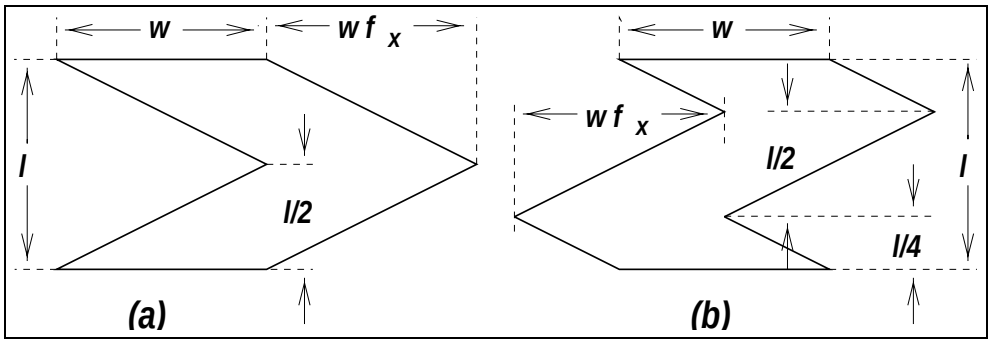

Figure 2.8: The chevron pads of the CERES TPC (type b). w=10.3 mm, l=6 mm and f_x=1.05 [61].

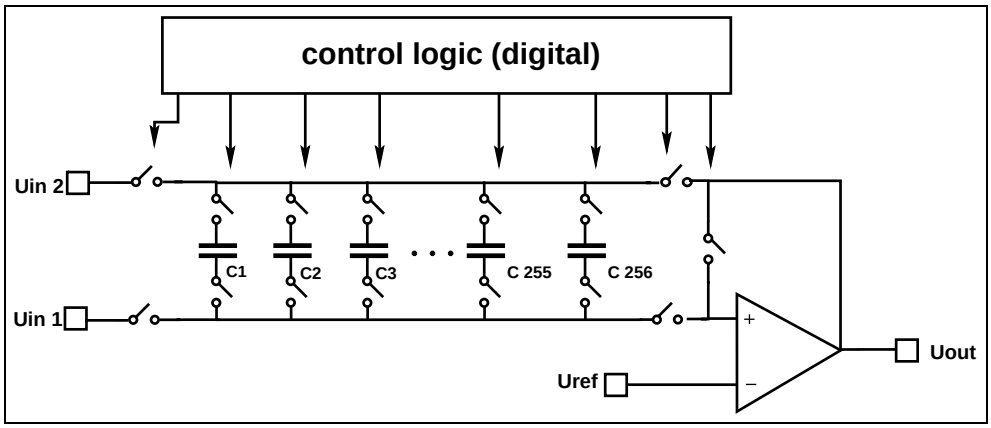

Figure 2.9: Schematic view of the SCA [61].

is done by means of zero suppression and Huffman compression. Each of the 16 readout chambers has its own electronics which works more or less independently. The front-end electronics consists of a charge sensitive amplifier, an analog memory to record the analog output signal of the amplifier (a Switched Capacitor Array (SCA)), a 10-bit ADC and a digital ALTRO chip. The ALICE TPC ReadOut chip is described in section 4.5. A modified version of this chip was used for the CERES readout. The SCA chip contains 16 channels, each with 256 individual samples. Each memory cell consists of a 1.4 pF double poly capacitor. A 16-to-1 analog output multiplexer allows the use of a single external ADC per chip. During sampling, each capacitor is connected to the chip input in turn, by a switch. The 256 switches are closed in turn by a shift register.

2.4 The STAR central TPC

The STAR detector [64] uses the TPC [65] as its primary tracking device. The TPC records the tracks of particles, measures their momenta, and identifies the particles by measuring their ionization energy loss (dE/dx). This is the first big cylindrical TPC in a heavy ion collider experiment.

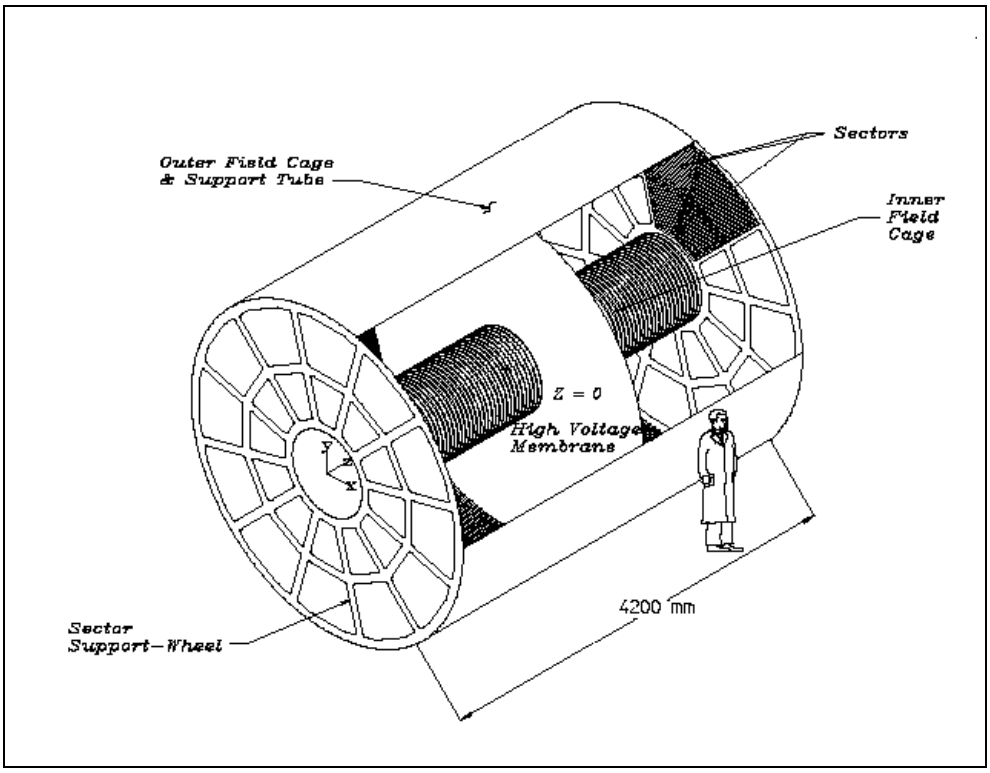

Figure 2.10: Perspective view of the central STAR TPC [65].

2.4.1 Geometry and drift volume

The central TPC is 4.2 m long and 4 m in diameter. The central cathode, the inner and outer field cage cylinders and the readout end caps, defines the electronic drift field.

Magnetic field

The magnetic field is provided by a large solenoidal magnet that operates at 0.5 T. [67]

Electric drift field

The central membrane is operated at 28 kV. The end caps are at ground potential. The field cage cylinders serve the dual purpose of both gas containment and electric field definition. The mechanical design has been optimized to reduce mass, in order to minimize secondary particle production. The metal layers of the field cage are made of flexible printed circuit material, kapton with metal on both sides. The metal is etched to form electrically separated 10 mm strips separated by 1.5 mm. This defines a uniform electric field of $\approx$ 135 V/cm.

Gas properties

P10 gas (10% Methane and 90% Argon) is used, regulated at $\sim$2 mbar above atmospheric pressure [66]. The regulated pressure is chosen in order to modify the density of the gas mixture. Temperature is regulated close to 24°C for the TPC via water cooling.

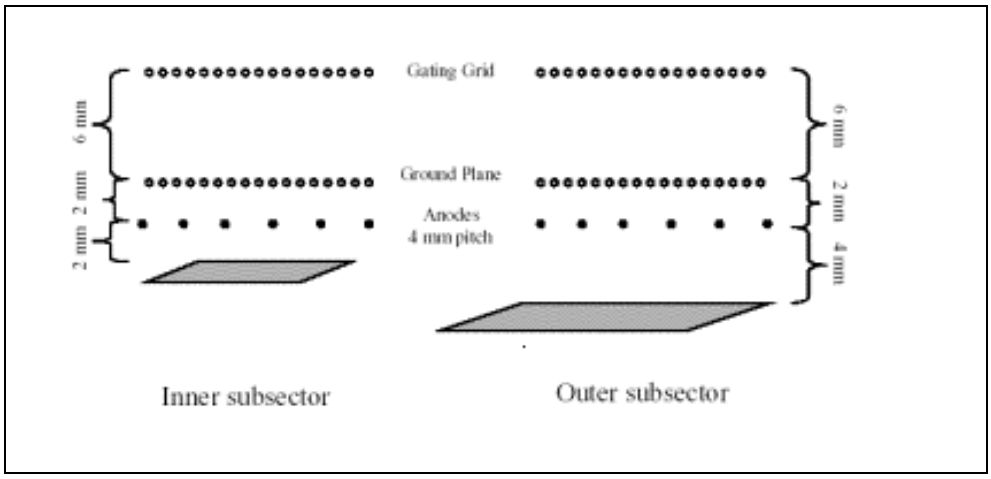

Figure 2.11: Wire geometry of the central STAR TPC outer and inner readout subsectors [65].

Temperature and pressure affects the drift velocity of the gas mixture. The P10 gas is chosen for its relatively high drift velocity at a low electric field. Such low electric fields simplify the design and operation of the field cage. TPCs using these mixtures do however need a longitudinal magnetic field to reduce the, otherwise huge, transverse diffusion. Typically, the electric field in the drift volume is close to 148 V/cm. This is somewhat above the saturation value (135 V/cm), so that the drift velocity can be regulated by varying the field. Organic components like Methane may cause ageing. This is less desirable for ALICE.

2.4.2 Readout chambers

The end caps of the TPC are divided into 12 read out sectors, which are further divided into inner and outer subsectors. The pads of the inner subsectors are 2.85 mm x 11.5 mm. The pads of the outer subsectors measure 6.20 x 19.5 mm. The chosen wire direction is set to best determine the momentum of the highest transverse momentum (p_t) particles. Their tracks are nearly straight radial lines emanating from the interaction point. The momentum of low p_t particles is well determined without special consideration. The sagitta of the high p_t tracks is accurately determined by setting the anode wires roughly perpendicular to the straight radial tracks because position resolution is best along the direction of the anode wire. In the other direction, the resolution is limited by the quantized spacing of the wires (4 mm between anode wires). The dimensions of the rectangular pads are likewise optimized to give the best position resolution perpendicular to the stiff tracks. The width of the pad along the wire direction is chosen such that the induced charge from an avalanche point on the wire shares most of its signal with only 3 pads. Concentrating the avalanche signal on 3 pads gives the best centroid reconstruction using either a 3 point gaussian fit or a weighted mean. The distance from the anode wire to the pad plane (2 mm for the inner subsectors and 4 mm for the outer ones) differs in order to compensate for the reduced pad size, such that the signal fraction seen on each pad is kept at an acceptable level.

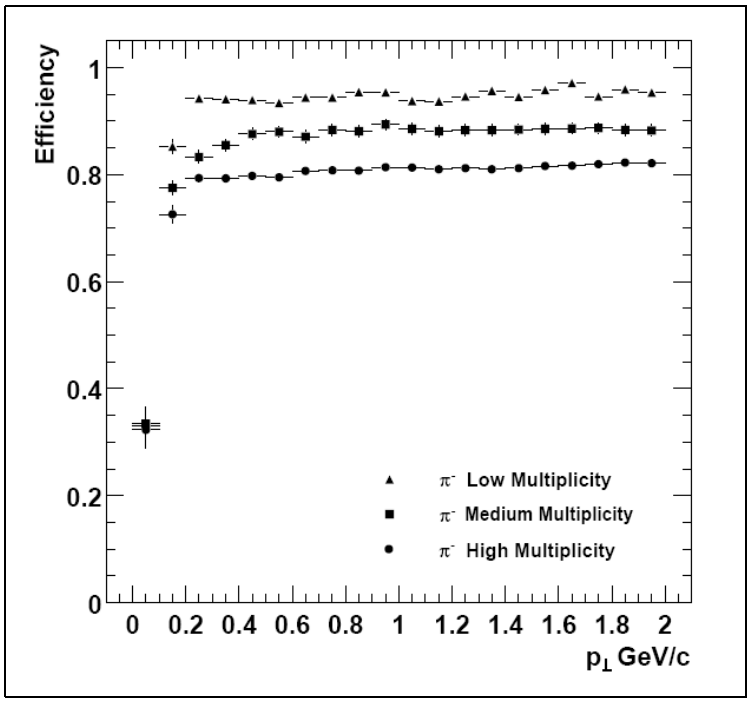

Figure 2.12: The pion tracking efficiency in STAR for central Au+Au events at RHIC for different multiplicities. Tracks with $|y| < 0.5$ were used to generate the figure and the magnetic field was set to 0.25 T [65].

2.4.3 Readout electronics

The total number of pads are 136608. NA49 and STAR share the SCA/ADC chip design. However, the STAR system has a larger dynamic range, lower noise, and a much faster readout rate. A slightly modified version of this electronics is also used for STAR forward TPCs.

2.4.4 Selected results

The tracking efficiency depends on the acceptance of the detector, the electronics detection efficiency, as well as the two-hit separation capability of the system. In order to estimate the tracking efficiency, simulated tracks are embedded inside real events. The number of simulated tracks that are found by the track reconstruction software are counted. The tracking efficiency for pions in STAR is shown in figure 2.12.

The primary vertex is found by considering all the tracks reconstructed in the TPC and then extrapolating them back to the origin. The global average is the vertex position. The primary vertex resolution of the STAR Central TPC is shown in figure 2.13. It is calculated by comparing the position of the vertices that are reconstructed using each side of the TPC, separately. As expected, the resolution decreases as the square root of the number of tracks used in the calculation. A resolution of 350 μm is achieved when there are more than 1,000 tracks.

The momentum resolution of the detector is found by embedding simulated tracks in real events and reconstructing them with the standard reconstruction procedure. Fig-

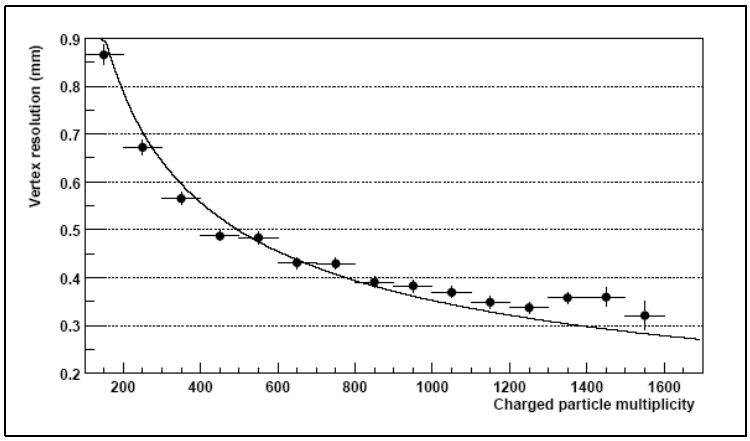

Figure 2.13: Primary vertex resolution in the transverse plane [65].

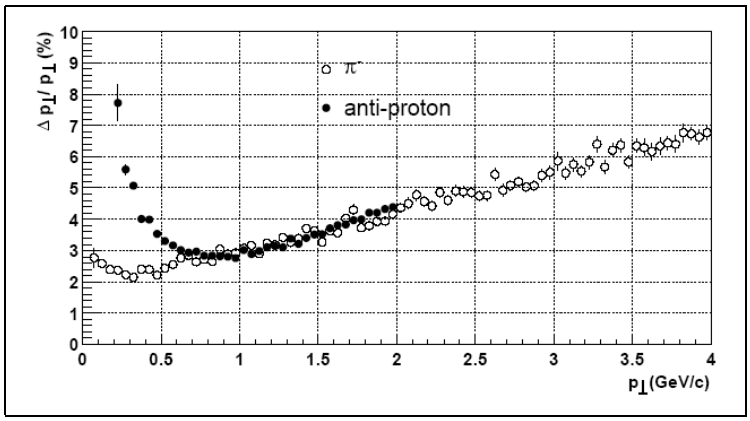

Figure 2.14: Transverse momentum resolution of the STAR TPC for π^- and anti-protons in the 0.25 T magnetic field. Tracks are required to be formed by more than 15 hits. Tracks are embedded in minimum bias events. The momentum resolution is calculated as the Gaussian sigma [65].

ure 2.14 shows the transverse momentum resolution for π^- and anti-protons.

Energy lost in the TPC gas is a valuable tool for identifying particle species. It works especially well for low momentum particles. As the particle energy rises, the energy loss becomes less mass-dependent and it is hard to separate particles with velocities v $> 0.7c$. STAR was designed to be able to separate pions and protons up to 1.2 GeV/c. Figure 2.15 shows the energy loss for particles in the TPC as a function of the particle momentum. The data have been corrected for signal and gain variations and the data are plotted using a 70% truncated mean. The magnetic field setting is 0.25 T. The resolution is 8% for a track that crosses 40 pad-rows. At 0.5 T, the dE/dx resolution improves because the transverse diffusion is smaller and this improves the signal to noise ratio for each cluster. Figure 2.15 includes both primary and secondary particles. The prominent proton, deuteron, and muon bands come from secondary interactions in the beam pipe and Inner Field Cage, and from pion and kaon decays. Pions and protons can be separated from each other up to 1 GeV/c.

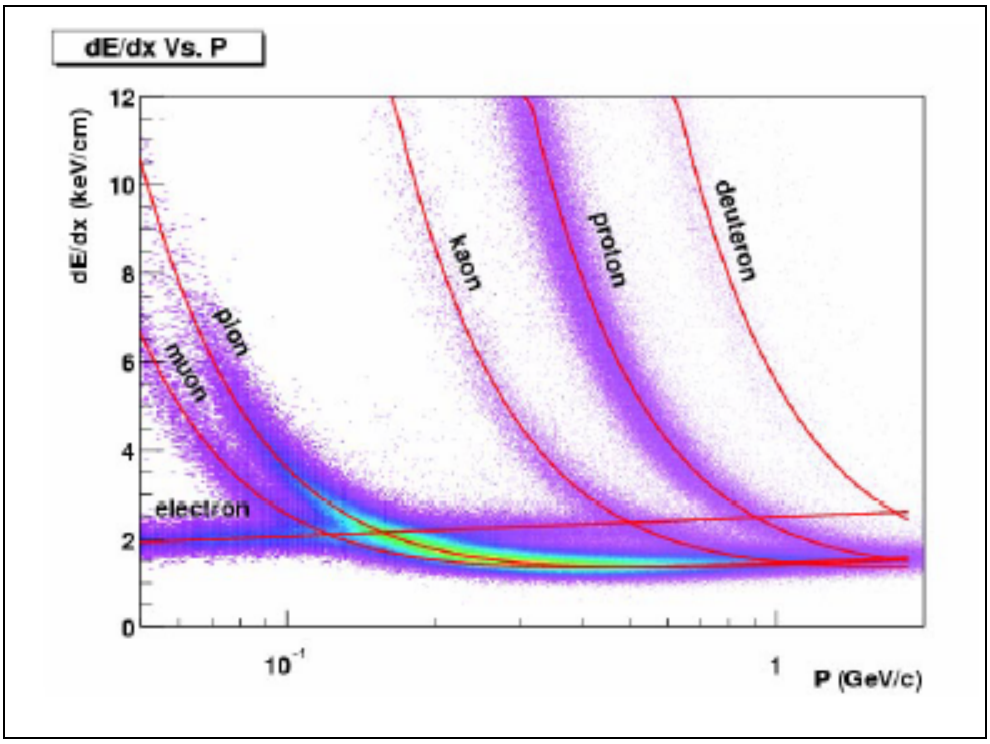

Figure 2.15: The energy loss distribution for primary and secondary particles in the STAR TPC as a function of the p_T of the primary particle. The magnetic field was 0.25 T [65].

2.5 The STAR forward TPC

The Forward Time Projection Chambers (FTPC) [68] were constructed to extend the acceptance of the STAR experiment. The two FTPCs cover the pseudorapidity range of $2.5 < |\eta| < 4.0$ on both sides of STAR and measure momenta and production rates of positively and negatively charged particles as well as neutral strange particles. The increased acceptance improves the general event characterization in STAR and allows the study of asymmetric systems like p+A collisions.

The STAR forward TPCs are constructed with a radial drift field. The main argument to chose a TPC with a radial drift field is the improved two-track resolution in the transverse direction, since the transverse separation of two charge clouds increases linearly with increasing r, but the width of the charge distribution of the cloud only increases by $\sqrt{r}$. The multiplicity in the STAR forward TPC is approximately 1000 charged particles in a central Au + Au collision. Most of these particles follow paths with small angles with respect to the beam direction.

Design features of the FTPCs:

- The electrons drift in a radial electrical field perpendicular to the solenoidal magnetic field.
- Curved readout chambers are used to keep the radial field as ideal as possible.
- A two-track separation of 1-2 mm is expected.

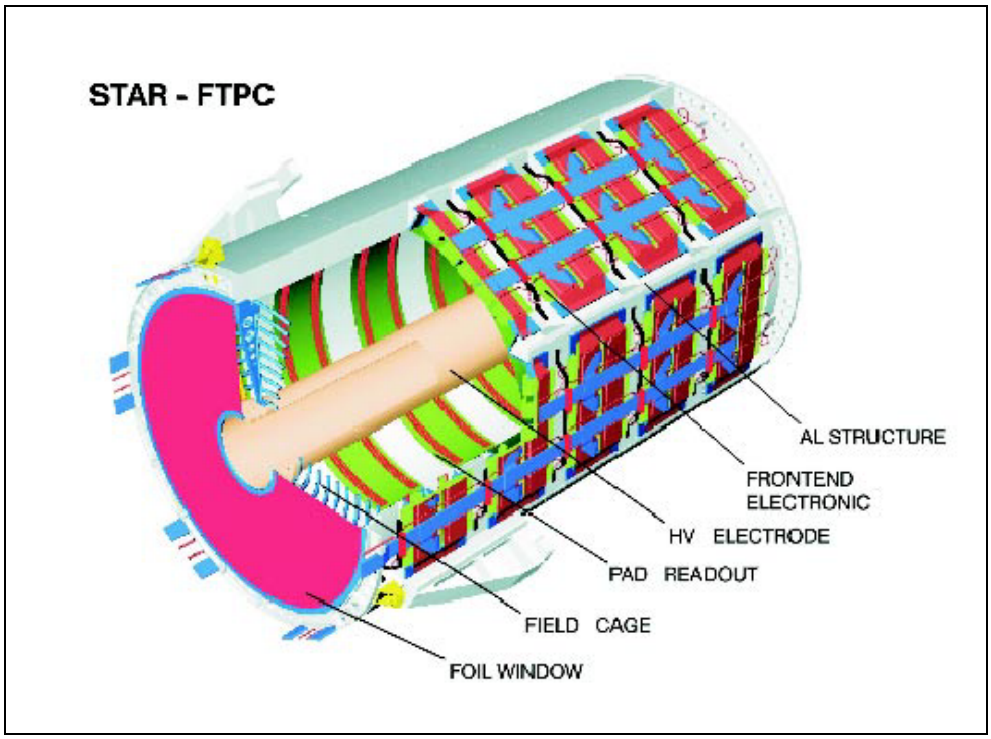

Figure 2.16: Perspective view of the STAR FTPC [68].

2.5.1 Geometry and drift volume

The FTPC has a cylindrical structure, 75 cm in diameter and 120 cm long. The field cage is formed by the inner HV-electrode, a thin metalized plastic tube, and the outer cylinder wall at ground potential. Drift length at maximum is only 23 cm.

Magnetic field

The same magnet as the central STAR TPC [67] defines the magnetic field.

Electric drift field

A radial drift field has been chosen to improve the two-track separation in the region close to the beam pipe where the particle density is highest.

Gas properties

50% Argon and 50% CO_2 is used. The drift time is about 50 μs for the 23 cm maximum drift length.

2.5.2 Readout chambers

The readout chambers (see figure 2.17) are located in 5 rings on the outer cylinder surface. Each ring has two padrows and is subdivided azimuthally into 6 readout chambers. In total there are 9600 pads, each 1.6 x 20 mm^2. The anode wires crosses the pads with an

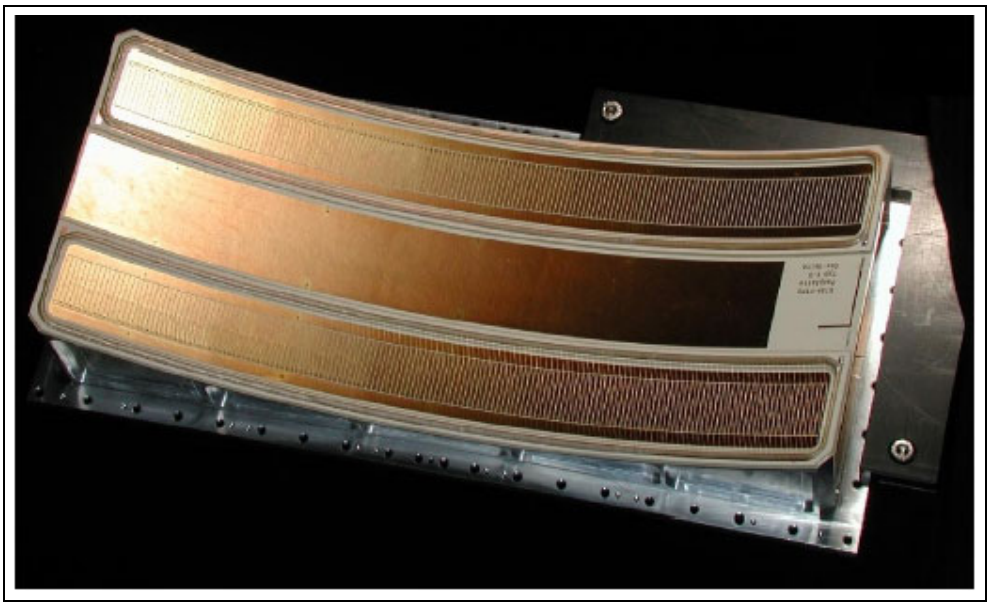

Figure 2.17: Photograph of a STAR FTPC readout chamber. The bending radius is 305 mm and each of the two padrows has 160 pads [68].

angle of 17.4°. This means that three anode wires are crossing each pad. This is to compensate focusing effects which would lead to periodic shifts in the position measurement, if the wires had been orthogonal to the axial direction of the pads [69].

2.5.3 Readout system including front end electronics

The STAR FTPC has a granularity of 9600 channels, recording 256 time samples each. The design of the front end electronics closely follows that of the STAR central TPC, see section 2.4.3.

2.6 The ALICE Time Projection Chamber

The TPC is the main tracking detector of the central barrel of the ALICE experiment at LHC. It provides charged particle momentum measurement with sufficient momentum resolution, particle identification by dE/dx and by decay topology analysis and vertex determination in the region $p_t < 10$ GeV/c and pseudorapidities $\mid \eta \mid < 0.9$. In addition, the TPC will provide - standalone or in combination with the other central barrel detectors (e.g. TRD and ITS) - the input for the High Level Trigger in order to select low cross section signals and rare processes.

2.6.1 Geometry and drift volume

The basic design philosophy of the TPC field cage is to provide a highly uniform electrostatic field in a cylindrical high-purity gas volume to transport primary charges over long distances (2.5 m) towards the readout end-plates. For reasons of symmetry in colliding beam arrangements, two such field configurations are chosen, back-to-back in a common

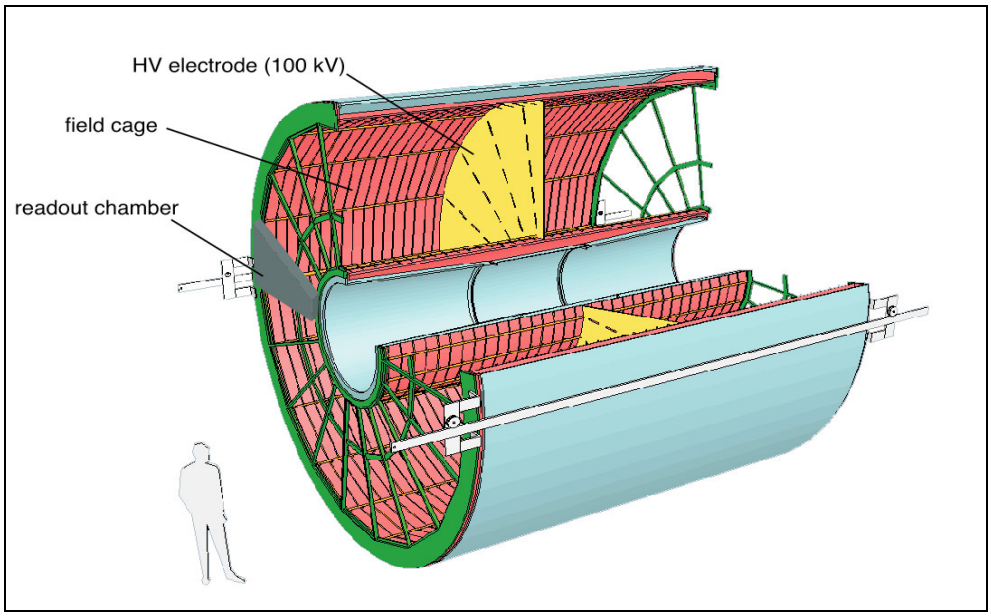

Figure 2.18: Perspective view of the ALICE TPC [73].

gas volume, with a common high-voltage (HV) electrode located at the axial centre of the cylinder.

The ALICE TPC has an inner radius of 90 cm, given by the maximum acceptable hit density (0.1 cm^{-2}). The outer radius, 250 cm, is given by the length required for a dE/dx resolution of < 10 %, necessary for particle identification.

The actual field cage volume is surrounded by an insulating gas envelope (containment). Containment of the drift volume is essential for personnel and operational safety and also for minimizing the amount of material traversed by particles. The construction of the field cage and consequently the choice of material are driven by the following constraints:

- high structural integrity against gravitational and thermal loads;
- very low permeability to atmospheric gas components considered harmful to the drift gas (O_2, N_2 and H_2O);
- negligible vapor pressure of contaminants emanating from material exposed to the drift volume;
- adequate surface smoothness to protect against HV discharges;
- low-density and minimal material near perpendicular to the beam direction to reduce multiple scattering and secondary particle production.

In order to fulfill the requirements on structural accuracy and gas tightness combined with small multiple scattering, composite materials are used throughout.

Magnetic field

The ALICE TPC is placed inside a solenoid magnet with a nominal field of 0.5 T.

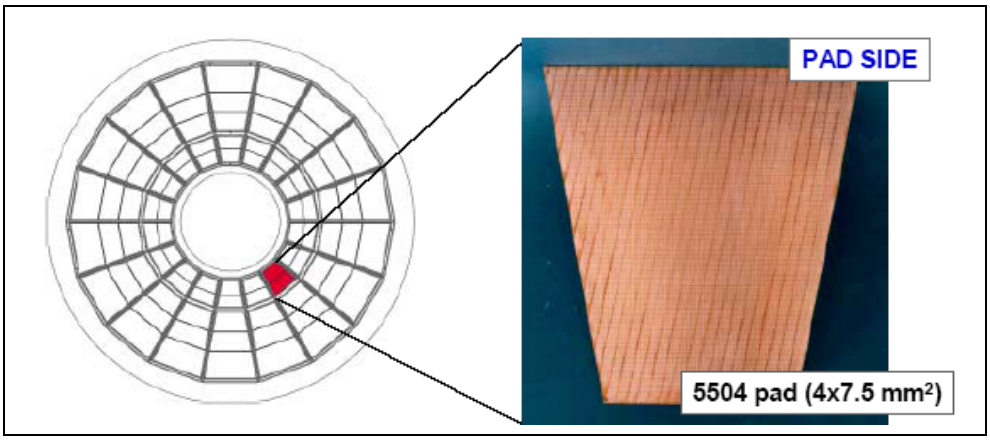

Figure 2.19: Segmentation of the readout plane [73].

Electric drift field

The strips for configuring a uniform field along the axis of the cylinder are made of aluminized Mylar. These are constructed 25 μm thick and 12.7 mm wide. 166 strips are needed over the 2.5 m length from the central electrode to the end planes on either side. The drift gas requires a high drift field (400 V/cm) to secure an acceptable drift time of 88 μs, thus the central electrode must be at 100 kV.

Gas properties

The drift gas [74] (90%Ne-10%CO_2) is optimized for drift speed, low diffusion, low radiation length and multiple scattering, and ageing properties. The drawback of Ne/CO_2 is that it has a steep dependence of drift velocity on temperature. The ALICE TPC is aiming for a thermal stability with $\Delta T \leq 0.1$ K.

2.6.2 Readout chambers

The overall design of the readout plane is chosen to provide full azimuthal coverage, and to optimize the momentum and dE/dx resolution. For practical reasons there is a subdivision of the 2 readout planes into individual modular readout chambers. The readout chambers will be mounted into the sector cutouts of the end planes, see figure 2.19.

Wire planes

As illustrated in figure 2.20 the ALICE TPC employs the commonly used TPC scheme of a grid of anode wires above the pad plane, a cathode wire grid and a gating grid. All wires run in the azimuthal direction. Since the design constraints are different for the inner and outer chambers, their wire geometry is different. The readout chambers of the ALICE TPC are conventional multiwire proportional chambers. The working principle is illustrated in figure 2.1. Incoming energetic ionizing particles passes through the drift region of the TPC. The liberated charge is subject to the electric field, which has direction

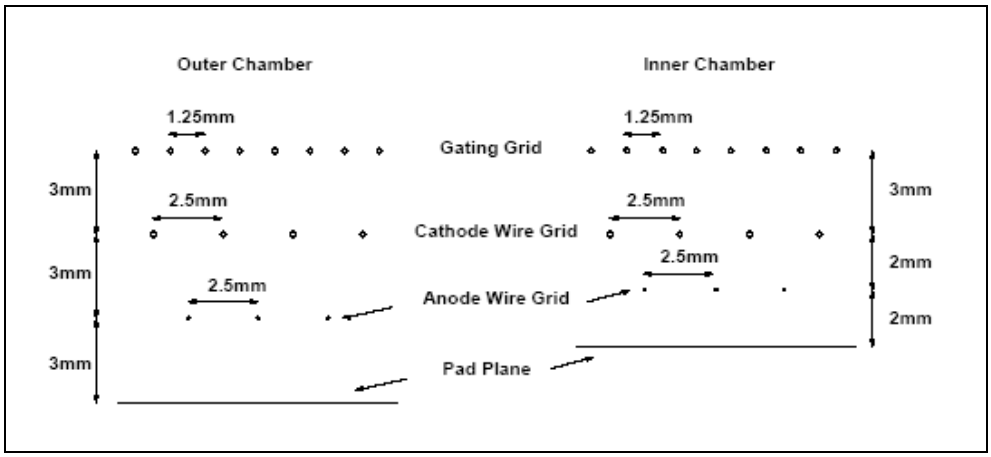

Figure 2.20: Wire geometry of the ALICE TPC outer (left) and inner (right) readout chambers [44].

parallel to the beam. The electrons will drift opposite the direction of $\overrightarrow{E}$, in the direction of the two end-caps.

Since the anode wires are not read out, there are no field wires in between them to prevent crosstalk between adjacent anode wires. Such field wires would reduce the signal coupling to the pads, as they pick up a significant fraction of the signal [44]. The absence of field wires also reduces the mechanical forces on the wire frames. A chamber without field wires requires a higher voltage to achieve the required gas gain [44]. The choice of a 2.5 mm pitch for the anode wires is done to minimize the accumulated charge per unit length of the anode wire and hence the risk of rate-induced gas gain variations in the environment of high particle multiplicity and relatively large gas gains.

The cathode wire pitch is 2.5 mm. No noticeable reduction is observed in electron transmission, while a large amount of the ions produced in the amplification avalanche are collected at the cathode wires.

The gating grid has alternate wires connected together electronically. In the open gate mode all the gating grid wires are held at the same potential, admitting electrons from the drift volume to enter the amplification region. In the absence of a valid trigger, the gating grid is biased with a dipolar field which prevents electrons from the drift volume from entering the amplification region. In addition, the closed gate stops ions created in the avalanche processes of previous events from drifting back into the drift volume. Such ions would accumulate in the drift volume and cause severe distortions of the drift field [44].

To minimize the contribution on the integral charge accumulation from particles traversing the gap between the gating grid and the pad plane, the gap between the gating and cathode wire grid is only 3 mm, sufficient to trap the ions within a typical gate opening time of 100 μs. In order to keep the alternating bias voltages low, the pitch between the gating grid wires is 1.25 mm.

	Anode wires	Cathode wires	Gating grid wires
Material	Au plated W	Cu/Be	Cu/Be
Diameter	20 μm	75 μm	75 μm
Stretching force	0.45 N	1.2 N	1.2 N
Length (inner)	27 - 44 cm	27 - 44 cm	27 - 44 cm
Length (outer)	45 - 84 cm	45 - 84 cm	45 - 84 cm
Total number (inner)	200	201	400
Total number (outer)	456	457	912

Table 2.2: Wire parameters of the ALICE TPC.

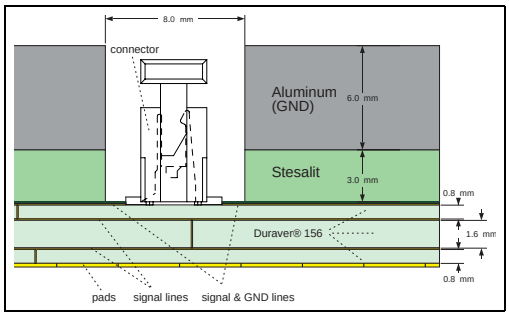

Figure 2.21: Cross section through the pad plane showing the pad and signal layers as well as a cross section of a Harwin connector. The Stesalit plate insulates the connectors layer from the Aluminium body [75].

Readout pads

The pad plane consists of multilayer printed circuit boards (PCBs), with the readout pad pattern exposed to the drift volume. The pads are made of copper. There are three different pad sizes: 4x7.5 mm^2 in the inner chambers, and 6x10 mm^2 and 6x15 mm^2 in the outer chambers. The signals induced on the pads are read out on the rear side of the pad plane PCB. Each pad uses vias (holes in the PCB) to connect it to the back side, where a connector for connecting the cables to the front end electronics is mounted. Figure 2.21 shows a cross section of the pad plane. The advantage of a multi layer board, beside the routing, is that the vias are only to the next layer. The PCB layouts will be arranged such that other layers will cover all holes/vias, thus preventing gas leakage [75].

2.6.3 Readout system including front end electronics

The readout system of the ALICE TPC is presented in detail in chapter 4. The Front End Electronics chain consists of a custom PreAmplifier and Shaper (PASA), an ASIC doing digitizing and data processing (the ALICE TPC Read Out chip - ALTRO) and a data collector and Front End Electronics control unit (Readout Control Unit - RCU). Data processing like baseline correction, tail cancellation, zero suppression and lossless data compression is done before the data is sent to the counting house.

Chapter 3

The ALICE experiment

In this chapter an overview of the detectors being part of the ALICE experiment is given. The present ALICE design is a multi-purpose one, which foresees the measurement and identification of produced hadrons, leptons and photons. Special emphasis is given to the Time Projection Chamber and the experiment parts that has connection to it in general and the Readout Control Unit in particular.

3.1 Physics motivation for the ALICE experiment

The LHC will run with heavy ions about 10 % of its running time, which translates into 10^6 seconds of running time per year. The event rate of Pb-Pb collisions, given the maximum luminosity of 1.0 x 10^{27} cm^{-2} s^{-1} and an inelastic cross-section of 8 b, will be 8000 minimum-bias collisions per second. Only some 5 % of these events are expected to correspond to the most central collisions. This low interaction rate allows the use of slow but high-granularity detectors, like the time projection chamber (TPC) and the silicon drift detectors. The ALICE rapidity acceptance has been chosen to be large enough to allow the study of particle ratios, p_t spectra and HBT (Hanbury-Brown-Twiss) radii on an event-by-event basis. Detecting the decay products of particles at $p_t < m$ requires about 2 units of rapidity (for masses above 1-2 GeV/c^2) and a corresponding coverage in azimuth. ALICE has been specifically designed to maximize momentum coverage, from $\approx$ 100 MeV/c, the lowest values relevant for thermodynamical studies, to $\approx$ 100 GeV/c, the transverse momentum of the leading particles of jets with transverse energy well over 100 GeV. The measurement of numerous precision points over a long measured track length in a moderate magnetic field and with minimal material allows to satisfy both requirements.

3.2 The ALICE detectors - detector overview

An overview of the ALICE detector setup at point 2 of the LHC is seen in figure 3.1. The experiment consists of a central detector system, covering mid-rapidity ($\mid \eta \mid \leq 0.9$) over the full azimuth, and several forward systems for extension of the experiment to large rapidity. The central system is installed inside a large solenoidal magnet which generates a magnetic field of $\leq$ 0.5 T.

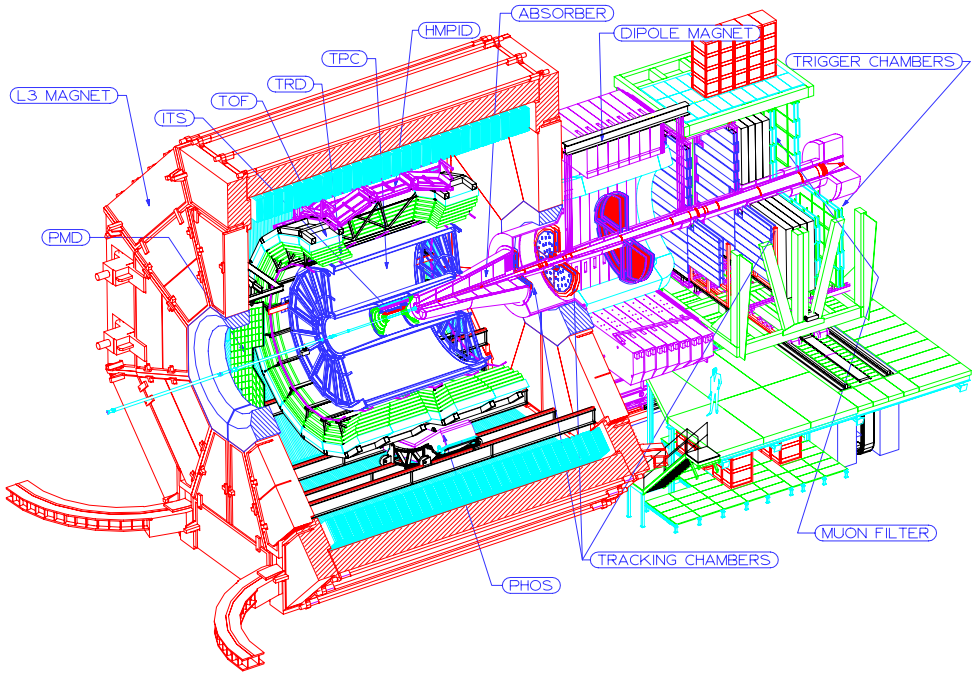

Figure 3.1: Longitudinal view of the ALICE detector.

Tracking is performed combining the information from a large TPC, which is the main tracking detector, with the information coming from the Inner Tracking System (ITS). Additional tracking information is available from the Transition Radiation Detector (TRD). The ITS is a system of six barrel layers of silicon detectors providing high-resolution tracking and precise vertex information. Particle identification is provided by the central tracking detectors through the measurement of the specific energy loss which provides separation of low momenta and some identification capability in the relativistic rise region. Electrons are identified using the TRD information. Short-lived particles are identified by exploiting the tracking and vertex measurement of the experiment. Muons are measured by the Forward Muon spectrometer. Photons are measured in a very-high-granularity crystal calorimeter (PHOton Spectrometer - PHOS). A dedicated high-resolution Time Of Flight (TOF) barrel provides the identification for the bulk of the intermediate-momentum particles. High Momentum Particle Identification Detector (HMPID), a Ring Imaging Cherenkov detector, extends the particle identification to higher momenta. Finally, the event characterization is completed by the measurement of charged particle multiplicity, forward photon multiplicity and the very-forward energy flow (measuring spectators in the collision) by the forward detectors. A High Level Trigger, with an online processing computer farm, is included to enhance the sensibility to low cross section processes.

3.3 The ALICE tracking system

This consists of the Inner Tracking System (ITS) and the Time Projection Chamber (TPC). They are presented below.

3.3.1 Inner Tracking System

The ITS [43] consists of six cylindrical layers of silicon detectors, as illustrated in figure 3.2. Because of the expected high particle density, up to 80 particles per cm^2, the four innermost layers is silicon pixel and silicon drift detectors. These are two-dimensional devices. The outer two layers, where the track densities are below 1 particle per cm^2, will be equipped with double-sided silicon microstrip detectors.

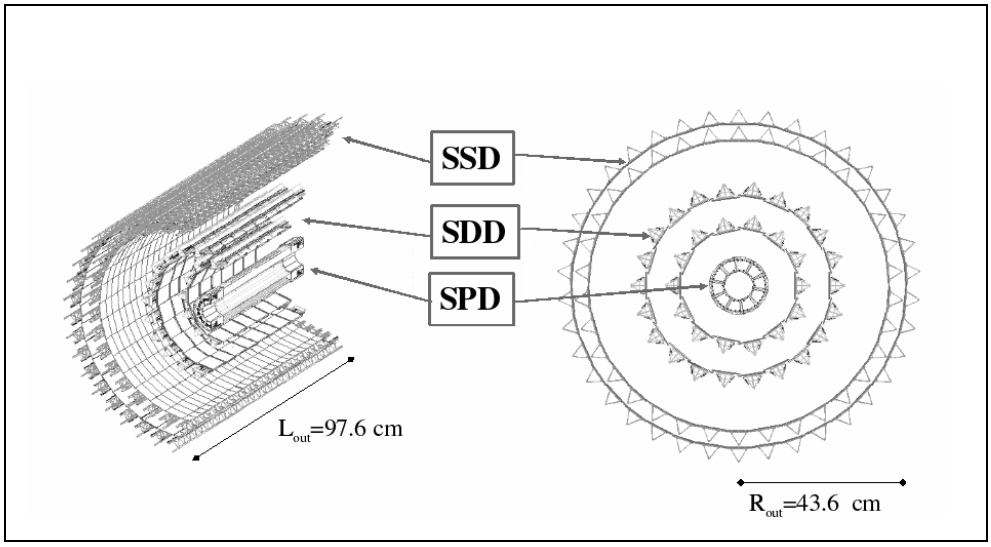

Figure 3.2: Layout of the ITS [43].

The tasks of the ITS are:

- to localize the primary vertex with a resolution better than 100 μm
- 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 (with limited 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.

Silicon Pixel Detector

The Silicon Pixel Detector consists of a two-dimensional matrix (sensor ladder) of reverse-biased silicon detector diodes bump-bonded to readout chips. Each pixel cell has the size of 50 x 425 μm^2 and is connected through a conductive bump (made of solder or indium) to a contact on the readout chip. 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.

Silicon Drift Detector

In the linear SDD a series of parallel drift cathodes is deposited on both sides of the detector. The function of these is to fully deplete the volume and to provide a constant electrostatic field parallel to the wafer surface, forming a drift region. Electrons created by an ionizing particle then drifts through this potential channel towards small capacitance anodes, as illustrated in figure 3.3. 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. The space resolution along the drift direction is better than 38 μm. In the z-direction (anode axis) the precision is better than 30 μm over 94% of the detector surface (reaches 60 μm close to the anodes). The two-track resolution is better than 300 μm.

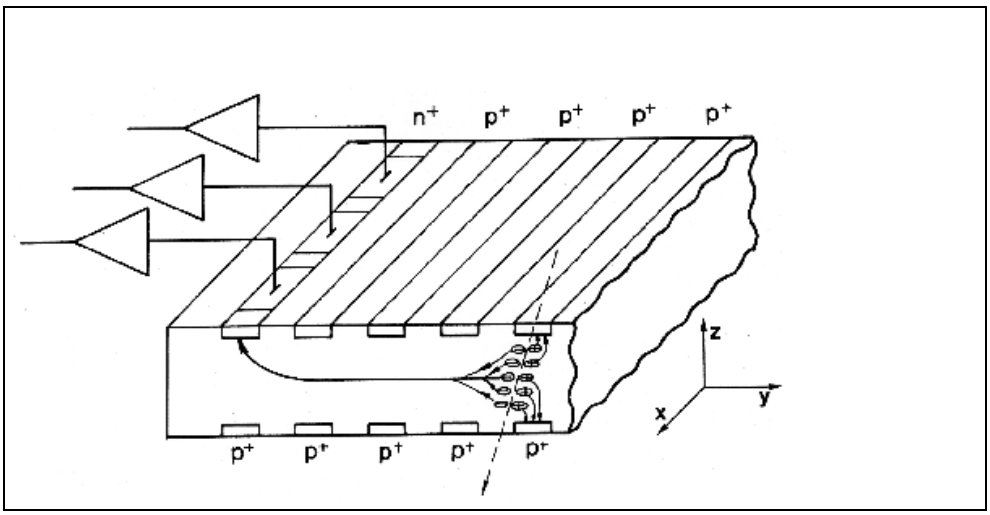

Figure 3.3: Principle of the linear Silicon Drift Detector [43].

Silicon Strip Detector

Both outer layers of the ITS consists of double sided silicon strip detectors, 782 in layer 5 and 988 in layer 6. Each of these sensors has a rectangular shape of 75 x 42 mm^2. All detectors are made of n-type silicon and are 300 μm thick. On each side 768 strips are implanted. 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, as illustrated in figure 3.4. In this way a two dimensional reconstruction of track position is achieved, with a spatial resolution of 20 μm in the rφ direction and 820 μm in the z-direction.

3.3.2 Time Projection Chamber

The TPC is the main tracking detector of the ALICE central barrel. It will provide charged-particle momentum measurements, particle identification and vertex determination together with the ITS, TRD and TOF. The TPC is described in [40] and presented in section 2.6.

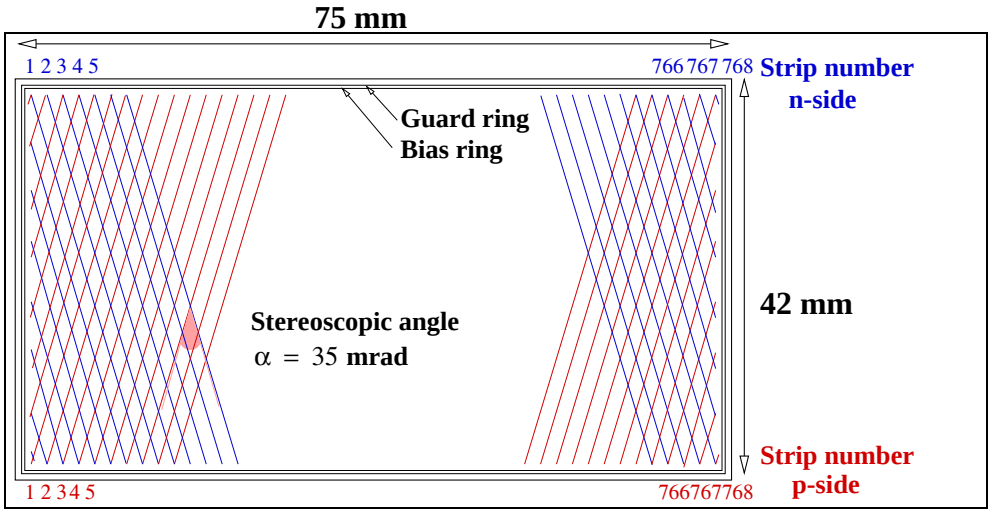

Figure 3.4: Principle of a Silicon Strip Detector (part of) [43].

Physics objectives

The TPC is the main detector for studying hadronic observables in both heavy-ion and pp collisions in ALICE. Combined with data from the TRD and ITS, the TPC will provide sufficient electron identification to measure the production of light and heavy vector-meson resonances and the dilepton continuum in Pb-Pb collisions. Hard probes, such as heavy quark production and high-p_t jets, requires very good momentum resolution at high momenta, which has to be achieved with the help of other tracking detectors, like the ITS and TRD. The electron identification provided by the TPC and TRD for $p_t > 1$ GeV/c can be used to measure open charm and open beauty produced in the collisions. This information must be combined with impact-parameter determination of electron tracks from the ITS. A similar approach 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 suppression due to quark-gluon plasma formation.

The TPC during proton-proton runs

Some modifications are needed in order to collect data in proton-proton runs. The memory time of the TPC due to the 90 μs drift time is the limiting factor. This is the time it takes to drift the electrons from the high voltage central membrane to the readout chambers. The foreseen interaction rate for pp collisions is about 200 kHz. The TPC records tracks from about 38 events during one readout cycle at this rate. This is illustrated in figure 3.5. Practically the full event is recorded for events close in time to the triggered event (left side). The fraction of the event recorded decreases with increasing distance in time. There is a final cutoff at $\pm$ 90 μs. Total occupancy is lower by about one order of magnitude compared to Pb-Pb since the average pp multiplicity is a factor 200 lower than the average Pb-Pb multiplicity [40].

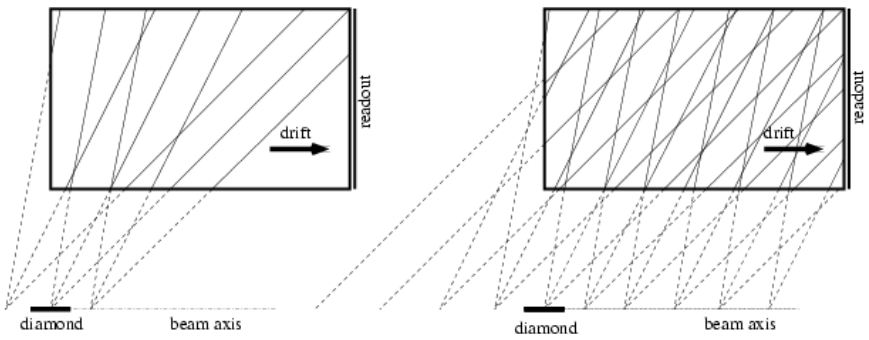

Figure 3.5: The distribution of primary vertices along the vertices along the beam axis as reconstructed from the TPC. The side view is limited to one quadrant only. The events are spaced evenly in time by about 11 μs. Left side includes one interaction before and one after the triggered event. The right side displays all interactions that contribute to the triggered event after the readout is completed [40].

3.4 Particle Identification System

Particle identification is provided in ALICE over a large part of the phase space by the combination of dE/dx measurement in the ITS and TPC with a barrel Time Of Flight (TOF) at r=3.7 m. A Transition Radiation Detector (TRD) allows electron identification above 1 GeV/c. The High Momentum Particle Identification Detector (HMPID) allows the identification of hadrons up to higher momenta (π/K to 3 GeV/c and K/p to 5 GeV/c).

3.4.1 Time Of Flight Detector

The TOF is a 140 m^2, 160000 channel system. The choice of technology for such a large area detector is the Multigap Resistive Plate Chamber (MRPC) [41]. The key aspect is that the electric field is high and uniform over the whole sensitive gaseous volume of the detector. Any ionization produced by a through-going charged particle will immediately start a gas avalanche process which will eventually generate the observed signals on the pick-up electrodes. This system assists in the identification of π, K and protons up to 2.5 GeV/c, on an event-by-event basis.

3.4.2 High Momentum Particle Identification Detector

This is a smaller-area Ring Imaging Cherenkov detector placed at a distance of $\approx$ 4.5 m from the beam axis. The HMPID consists of seven modules, each 1.5x1.5 m^2, for a total of over 160000 readout channels. It will use a liquid C_6F_{14} radiator and a MWPC with pad readout as UV detector.

3.4.3 Transition Radiation Detector

The Transition Radiation Detector (TRD) [45] measures in the same rapidity window as the TPC ($|\eta| \leq 0.9$).

Physics objectives:

- Sufficient electron identification to measure, in the dielectron channel, the production of light and heavy vector meson resonances for Pb+Pb collisions.
- Studies of the dilepton continuum (at the Drell-Yan level).
- Measurements of open charm and open beauty cross sections through their electron decay channels. This measurement is important for understanding the role of color screening in the expected suppression of quarkonium productions in Pb+Pb collisions at LHC energy.
- Identification of J/Ψ mesons resulting from B decay, a potential serious background for measurement of "direct" J/Ψ's.
- The tracking information of the TRD is available a few microseconds after the start of an event. It is foreseen to use this information to select events with high transverse momentum or high mass electron pairs. By opening the gating grid of the TPC only for such events, the yields of recorded collisions with high mass pairs could be enhanced by more than an order of magnitude.

The TRD is a barrel detector surrounding the TPC with 6 layers of a combination of a foil stack and a Xe-filled wire chamber. The basic operational principle of the TRD is sketched in figure 3.6. Transition radiation is emitted when a charged particle crosses the boundary between two media with different indices of refraction. This radiation can be observed in the form of soft X-rays. Since the probability for soft X-ray emission is only about 1% per boundary crossing, practical TRDs use radiators with a large number of interfaces. In the ALICE TRD a radiation stack of 100 polypropylene foils are used. The produced X-rays are detected in the Time Expansion Chambers (TECs). The electron clusters produced in the gas volume of the detector drift in a homogeneous electric field towards the amplification region of the TEC. There, gas multiplication takes place at the anode wires and the corresponding signals induced on a cathode pad plane are read out. The medium is selected such as only electrons create the transition radiation (not pions). Both the particle and the X-ray is detected.

3.5 Photon Spectrometer

The Photon Spectrometer (PHOS) [49] is a high resolution electromagnetic calorimeter consisting of 17920 detection channels based on lead-tungstate crystals ($PbWO_4$, with dimension 22x22x180 mm^3). Energy from photons and electrons are deposited within the crystal volume. The atoms of the crystal are then deexcited by emitting UV radiation. UV radiation is detected at one end of the crystal by an Avalanching PhotoDiode (APD).

PHOS will be positioned on the bottom of the ALICE set-up at a distance of 460 cm from the interaction point. It is optimized for measuring photons, π^0 and η mesons in broad energy ranges. The pseudorapidity coverage is -0.12 $\leq \eta \leq$ 0.12.

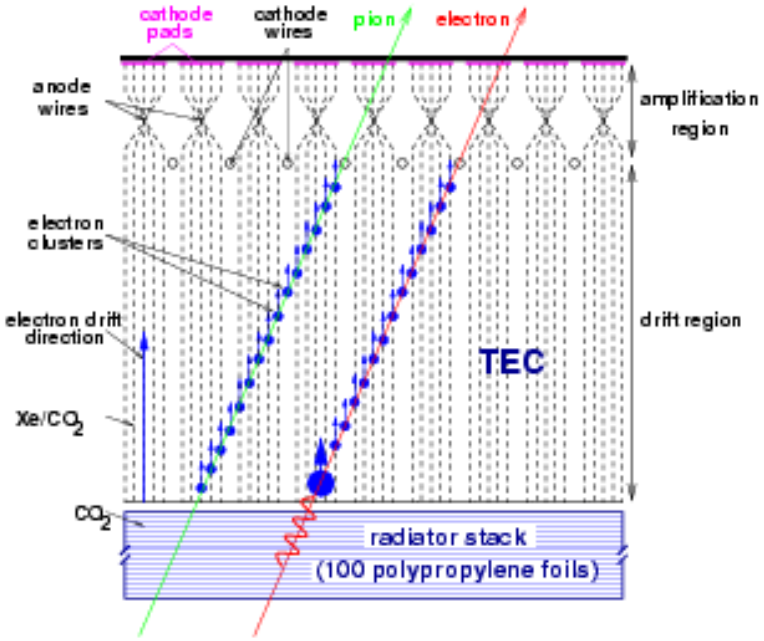

Figure 3.6: Basic operating principle of the ALICE TRD [45].

3.6 Forward Detectors

A number of smaller detector systems (ZDC, PMD, FMD, T0, V0) define and/or trigger on global event characteristics:

- Impact parameter is measured by four small and dense calorimeters (Zero Degree Calorimeter [47]).
- The event reaction plane and elliptic flow is measured by the Photon Multiplicity Detector [48]. The PMD also measures the ratio of photons to charged particles and the transverse energy of neutral particles.
- Multiplicity of charged particles is measured by the Forward Multiplicity Detector.
- Precise time of the collision is measured by the T0 counters which consists of 24 Cherenkov radiators with a precision of less than 50 ps.
- The event vertex is located by the V0 counters, which consists of scintillator paddles. These are used as the main interaction trigger, and also discriminates against beam-gas interactions by requiring the coincidence of the scintillators on both sides of the interaction region.

These detectors are all placed at small angles from the beam.

3.7 The Forward Muon Spectrometer

The complete spectrum of heavy quark mesons (i.e. J/Ψ, Ψ', Υ, Υ', Υ''), as well as the Φ meson, will be measured in the $\mu^+\mu^-$ decay channel by the ALICE Dimuon Spectrometer [46]. The Dimuon Spectrometer is designed to detect muons in the pseudorapidity

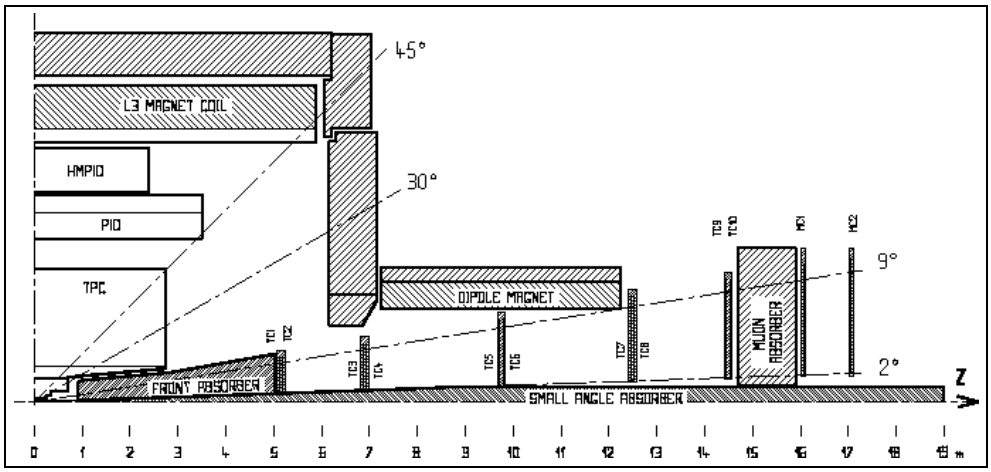

Figure 3.7: Principal layout of the forward muon spectrometer, including absorbers, dipole magnet, chambers, and the muon identifier [46].

range of 2.5 $\leq\eta\leq$ 4.0. The spectrometer consists of several parts, as illustrated in figure 3.7. A Front Absorber is placed closest to the interaction point to absorb hadrons and photons from the interaction vertex. At the same time, the Front Absorber is designed to protect the other ALICE detectors from the secondaries produced within the absorbing material itself. The spectrometer is shielded throughout its length by a dense absorber tube surrounding the beam pipe (called Small Angle Absorber in figure 3.7). The tube (made of tungsten, lead and stainless steel) reduces the background coming from direct particles produced during the collision and from particle showers in the tube itself. The Front Absorber and the Small Angle Absorber sufficiently protects the tracking chambers. A Muon Filter (which is an iron wall about 1 m thick) is placed just in front of the first trigger chamber. This is needed to prevent other particles than muons from reaching the trigger chambers. The tracking system for the muons consists of 10 cathode pad chambers. They are organized in 5 stations (marked as TC1 to TC10 in figure 3.7. The size of the pad planes varies from a few square meters for the TC1 to more than 30 m^2 for TC10. The total number of channels is about 1 million. This allows the detector to achieve the desired space resolution of about 100 μm, which is needed for a mass resolution of 100 MeV/c^2 for the Υ. The trigger system consists of 4 Resistive Plate Chambers arranged in two stations, one meter apart, placed behind the Muon Filter. The total active surface is about 150 m^2. A dimuon trigger signal is issued when at least two tracks above a pre-defined p_t threshold are detected in an event. The p_t cut is done to reduce the huge rate of low-p_t muons from π and K decays.

3.8 Trigger and Data Acquisition

The increasing importance of low-cross section, high-p_t observables (confirmed by resent RHIC results) has guided the design of the ALICE trigger system. It includes a pretrigger, three trigger levels (L0, L1 and L2) and the High Level Trigger (HLT). The pretrigger provides a fast wake-up signal to the Transition Radiation Detector frontend, thus allowing its digital electronics to be in a low-power mode most of the time. This consists of inputs from the T0 and V0 detectors. A copy of the inputs are sent in parallel to the TRD while being sent to the ALICE trigger. The L0 and L1 triggers gate the fast detectors, while only after an L2 decision level has been reached, the slow Time Projection Chamber is read out. The level 0 signal reaches the detectors at 1.2 μs. This is too fast to include all the trigger inputs in the decision, but needed by some detectors (e.g. HMPID using the GASSIPLEX front-end chip). All the remaining fast trigger inputs are picked up and sent as a level 1 signal at 6.5 μs. The L2 includes a past-future protection scheme. The high multiplicities expected in the ALICE environment make events containing more than one central collision unreconstructable. The level 2 waits to the end of the past-future protection interval (88 μs - equalling the TPC drift time) to verify that the event can be taken. A list of trigger inputs for the L0 and L1 trigger is shown in table 3.1. The trigger information is distributed from the ALICE central trigger to dedicated TTCrx-chips on each detector readout boards and synchronized to the LHC clock cycle. For each accepted event, information is sent specifying the event identifier (orbit number and bunch crossing number), trigger type, detector set to be read out, and the list of active trigger classes. These data are sent in the L2A (Level 2 Accept) message distributed to all TTCrx chips. The ALICE DAQ reads out data from the different detectors and archives it to permanent storage for later analysis. A bandwidth of 1.25 GBytes/s to mass storage is consistent with constraints imposed by technology, cost and storage capacity, including the cost of media needed to store the data and the computing needed to reconstruct and analyze these data. The data is read from the front-end electronics of each detector to the DAQ system in parallel over hundreds of optical links (Detector Data Link - DDL). Trigger selectivity (including the High Level Trigger), data compression and partial readout reduces the amount of data to satisfy the physics requirements with the total allowed bandwidth to mass storage.

3.9 High Level Trigger (HLT)

The High Level Trigger (HLT) provides online, real time event reconstruction in ALICE. The Time Projection Chamber (TPC), the Inner Tracking System (ITS) and the muon tracking chambers are tracking drift detectors and need a longer time span after the collision to deliver their data. Their slowness is compensated for by the detailed information they provide. The HLT system is intended to take advantage of this information (e.g. up to 76 MB/event at rates of up to 200 Hz for the TPC) in order to reduce the data rate as far as possible to have reasonable taping cost. After data reduction in the HLT, the data is returned to the DAQ and recorded onto an archival-quality medium for subsequent off line analysis. HLT can accomplish data reduction in many ways:

- It may be used as an advanced data compression device. Data compression can be

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	T0 Right	TRD jet low p_t
4	T0 Multiplicity	T0 Left	TRD jet high p_t
5	T0 beam gas	T0 Multiplicity	TRD electron
6	PHOS MB	T0 beam gas	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			

Table 3.1: List of trigger inputs for Pb-Pb and pp interactions.

implemented most efficiently if tailored towards the particular detector's response function and is therefore implemented in parallel for each sub-detector individually. Examples range from simple lossless or slightly lossy (e.g. vector quantization) compression techniques, pile-up removal in pp collisions up to advanced online track reconstruction and use of a track model to collect and code cluster information efficiently.

- It may generate selected candidate events according to the detected track topologies that correspond to certain specific trigger tasks. Such trigger task could be:

  HLT dimuon trigger More than 80% of dimuon fake events can be rejected within the HLT by tracking the event.

  HLT jet trigger The existence of a number of high p_t particles within a small phase space segment can be used as trigger for jets in the TPC acceptance. This search can possibly be augmented by TRD seeds. The stiff nature of the track candidates and their relatively close proximity allow for the implementation of a specific and fast local tracking in the TPC.

  HLT momentum filter About 50% of all tracks in the TPC are soft, i.e. $p_t < 0.4$ GeV. Tracking all tracks in the TPC (with the emphasis on high efficiency at high pt), and keeping only raw data regions along the high p_t trajectories, can reduce the data volume by a large fraction.

  HLT impact parameter trigger By propagating the TPC tracks (e.g. high p_t tracks) to the ITS, and with knowledge of the interaction point from either ITS or TPC tracking, the track-by-track impact parameter can be determined and used as a criterion for event rejection.

  HLT TRD trigger The TRD trigger will be dominated by background tracks. Therefore, the intention here is to reject fake tracks both by precise tracking (TPC/ITS) and by improved PID (TPC dE/dx) of the high p_t candidates in the TPC. The relevant tracks are identified by the TRD, and a processing command is sent to the appropriate sector processors of the TPC. After having received the event in memory, the processor starts tracking only the defined track candidates.

- It might reduce the full TPC readout to region-of-interest readout.

- It may compress the TPC data by tracklet and cluster modelling, which would reduce the data volume by a factor of 10.

In general first the initial detector specific local pattern recognition is performed. This is for example the cluster finder and the tracklet search functionality in the TPC. Both can be executed highly parallel due to the locality and high degree of independence of the data. The next step of data analysis implements global pattern recognition, such as global tracking in the TPC. The event reconstruction receives processed input of the sub-detectors and reconstructs the event by combining the derived information. The resulting trigger decision is thus based on a fully analyzed and reconstructed event. The classical approach of pattern recognition in a TPC detector is divided into two sequential

steps: cluster finding and track finding/fitting [51]. In general, one seldom finds reason in a traditional off line analysis (such as NA49 or STAR) to return to the raw data once the analysis chain is set up correctly. However, the cluster level is often returned to. Thus, cluster data are almost equivalent to raw data. In a first step in the classical approach a cluster finder reconstructs the cluster centroids, which are interpreted as the three dimensional space points produced by the traversing particle. The input to this cluster finder could be a list of above threshold time-bin sequences for each pad in the format described in section 4.10.2 (i.e. "raw" data). The cluster finder has no knowledge about the tracks, and solves the two dimensional problem of reconstructing the cluster centroids in padrow-time plane. The clusters found are then used by a "track finder", a "track merger" and a "track fitter". The algorithms used for the cluster finder and the track finder are based on the reconstruction scheme which has been implemented and used in the STAR L3 trigger. A first HLT-RORC (64-Bit, 66 MHz PCI Card) has been prototyped in summer 2002. The HLT-RORC receives detector raw data via the DDL-DIU and serves as a co-processor, where the pattern recognition is performed by an FPGA. The TPC cluster finder code has been ported to VHDL [52] and has been simulated and synthesized. Due to the high occupancy in the ALICE TPC, a large fraction of clusters overlap and, therefore, such a simple cluster-finder approach fails for central Pb-Pb collisions. However it is expected to be useable for pp collisions (see section 3.3.2) and heavy ion collisions with lower occupancy ($\mathrm{d}N_{ch}/\mathrm{d}\eta$=2000-3000) in the TPC. During pp running about 20 events will pile up within the drift-time of the TPC. Also there may be a substantial halo, which adds additional background hits to the data stream. These hits can be traced back and removed if clearly identified and if they do not belong to the primary interaction.

An iterative approach which aims to provide the track parameters prior to the cluster finding has been studied and implemented within the HLT framework. It contains an implementation of the Hough Transform which is applied to the raw ADC-values, followed by a subsequent cluster fitting and unfolding procedure which utilizes the obtained track parameters. This HLT tracklet-cluster-modelling approach preserves all information needed for further off line reconstruction and promises a compression factor of 10 [53].

The computer architecture of the ALICE HLT is driven by two constraints: the data has an inherent granularity and parallelism, which can be exploited by local pattern recognition operations (e.g. on-the-fly cluster finder and Hough transformation in the HLT-RORC). On the other hand, one single trigger decision has to be derived from a (complete) event reconstruction. Both requirements demand a hierarchical tree-like topology with a high degree of connectivity. Prototypes implementing such a topology have been built in Bergen and Heidelberg, are operational, and have been used for HLT data challenges. The system is designed to be flexible, open and scalable, such that additional functionalities can be easily added at any point in time, even during its operation.

Chapter 4

The readout chain of the ALICE TPC

4.1 Introduction

In this chapter the readout chain of the ALICE TPC is presented. The different devices are presented in the order of which they are used in the normal read out data flow of the TPC, thus starting with the Front End Electronics and ending with the ALICE DAQ/HLT which sends the data to permanent storage for later off line analysis.

An overview of the architecture and main components of the Front End Electronics is given in figure 4.1. The Readout Control Unit (RCU) is controlling and monitoring the connected Front End Cards, collecting data and sending it to the Data Acquisition System, via the Detector Data Link (DDL - optical fibre). The preamplifier and shaper (PASA) and the ALTRO chip (doing front end data processing) are mounted on the Front End Cards. The data coming from the TPC are read out from the Front End Cards using a Front End Bus, while the monitoring and controlling tasks are done via a dedicated Control Network. The interface of the Front End Electronics to the central ALICE trigger are on the RCU. The ALICE Detector Control System (DCS) also interfaces to the Front End Electronics via the RCU.

4.2 The grouping of readout pads into (sub)sectors

The two cathode planes, located at the TPC end caps (see chapter 2), are divided into 570132 pads. The signals from the pads are passed to 4356 front-end cards located some 10 cm away from the pad plane via flexible capton cables. The readout region of the end caps of the TPC is split into 36 sectors. These sectors are divided into 6 subsectors, with up to 25 Front End Cards connected in each subsector, as seen in figure 4.2. Each Front End Card connects to 128 pads. The data recorded in each subsector are read out by one RCU.

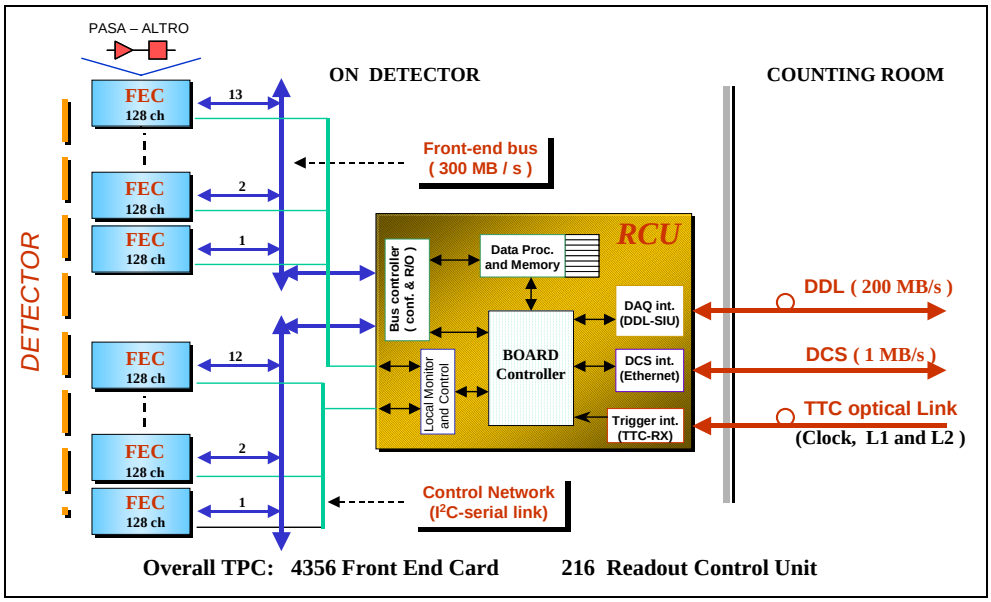

Figure 4.1: Architecture and main components of the Front End Electronics of the ALICE TPC.

4.3 Requirements for the Front End Electronics

The main requirements for the readout electronics are listed in table 4.1. One of the tightest requirements is defined by the extremely high pulse rate. The innermost detector chamber will be exposed to a signal occupancy up to 50%. The large data volume ($\sim$ 700 MByte/event) requires that zero suppression is performed in the FEE. In order to preserve the full resolution of the signal features (amplitude and time), a very accurate cancellation of the signal tail and correction of the baseline has to be performed before the zero suppression. The FEE (Front End Electronics) system has to satisfy many other constraints while meeting the required performance specifications. Mainly, the read-out electronics needs to fit into the overall detector structure; in particular into the available space, which has important consequences for the requirements on reliability, power, and cooling. The front end electronics should be radiation tolerant, as described in section 7.2.

4.4 Preamplifier and shaper (PASA)

On each front-end card a custom made charge sensitive amplifier transforms the charge induced on the pads into a differential semi-gaussian signal that is feed to the input of the ALICE TPC Read Out chip (ALTRO). The current signal delivered on the readout pads is characterized by a fast rise time (1 ns), and a long tail due to the motion of the positive ions. The amplitude has a typical value of 7μA. The signal is delivered on the detector impedance that can be approximated as a pure capacitance of the order of a few pF. The Preamplifier and shaper (PASA) is based on a charge sensitive amplifier (CSA) followed by a semi-Gaussian pulse shaper of the 4th order. These analogue functions are realized by a custom integrated circuit, implemented in CMOS technology 0.35 μm, which contains 16 channels with a power consumption of 11 mW/channel. The circuit

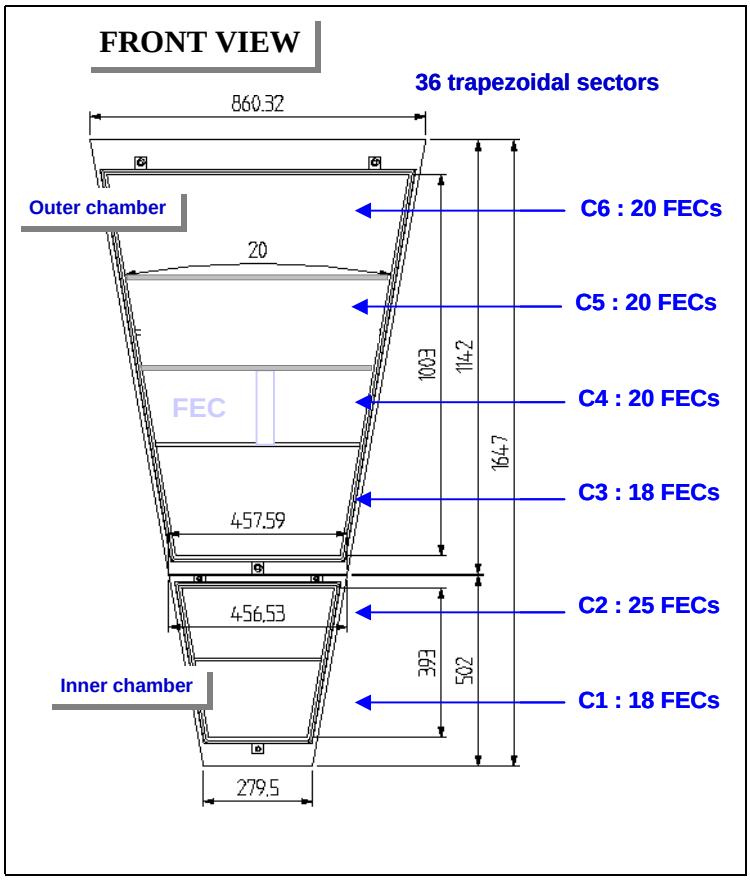

Figure 4.2: One of the 36 sectors (18 on each endcap of the TPC). Each sector is further divided into 6 subsectors (C1-C6), each read out by one RCU. The data is transferred to the counting house by one Detector Data Link (DDL) per RCU. The number of Front End Cards in each subsector is indicated.

has a conversion gain of 12 mV/fC and an output dynamic range of 2V with a differential non-linearity of 0.2%. It produces a pulse with a rise time of 120 ns and with a shaping time (FWHM) of 190 ns. The single channel has a noise value below 560 e^- (r.m.s.) and a channel-to-channel cross-talk below -60 dB [70].

4.5 ALTRO chip

The ALTRO chip [72] contains 16 of the TSA1001 ADC from ST Microelectronics in the form of two 8-ADC IP blocks along with system specific logic developed at CERN. The ADC has differential inputs and can work up to 25 MS/sec with a resolution of 10 bits, while consuming as little as 12 mW. It samples the signal feed from the PASA at a rate of 5-12 MHz. The rest of the data processing logic is divided into the stages indicated in figure 4.3, and explained below.

PARAMETER	VALUE	MEASURED
Nr. of channels	557 568	
Signal-to-noise ratio (MIP)	30:1	
Dynamic range	900:1	
Noise (ENC)	< 1000 e$^-$	560e$^-$
Conversion gain	12 mV/fC	12 mV/fC
Cross-talk	< 0.3 %	< 0.1 %
Shaping time (FWHM)	190 ns	188 ns
Sampling rate	5-12 MHz	
Tail suppression after 1 μs	0.1 %	
Power consumption	< 100 mW/channel	

Table 4.1: Front-end electronics requirements. The measured values are from the engineering run of the PASA [78].

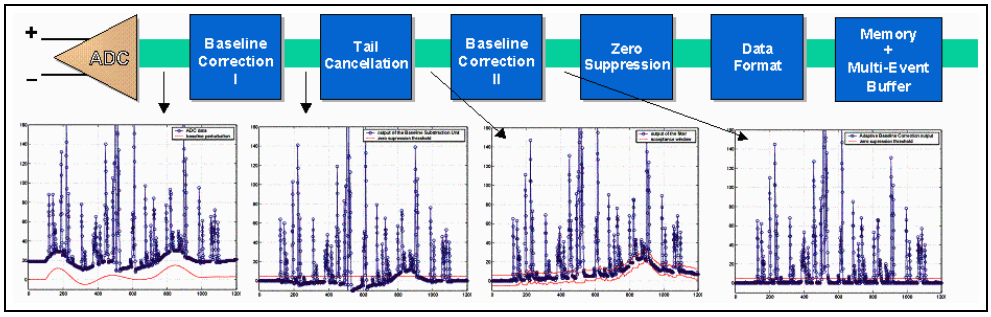

Figure 4.3: The data processing stages of the ALTRO chip [72].

4.5.1 First baseline correction

The first stage in the digital processing chain is the Baseline Correction and Subtraction I unit. The main task of this unit is to prepare the signal for the tail cancellation that takes place in the subsequent stage. To this purpose the signal is corrected in order to remove perturbations of different nature. These perturbations can be:

- Low-frequency spurious signals (in the range of less than one kilohertz). They perturb the detector signal by shifting its baseline by an amount that, inside the processing time window (PTW), is almost constant (less than one ADC count). This type of signal perturbation could be, for instance, the one produced by a temperature variation of the electronics components.
- Signal perturbations created by systematic effects, like those related to the triggering of the detector, which affect the signal in terms of a superimposed noise pattern.

This stage has 3 different modes of operation: subtraction mode, conversion mode and test mode.

Subtraction mode

This mode is further divided into 3 modes of operation as listed below:

- fixed-subtraction mode, the value to be subtracted from the input signal is constant and stored in a user-defined register;
- time-dependent subtraction mode, the time dependent pedestal values to be subtracted are stored in a memory (pedestal memory) that is addressed by a time counter started by the trigger signal;
- self-calibrated subtraction mode, the value to be subtracted is computed as cumulative average of the input signal.

While the fixed-mode and time-dependent-mode are exclusive, any of them can be combined with the self-calibrated mode.

Conversion mode

This mode performs a conversion of the input signal of the type $y_n = F(x_n)$, where x_n is the input sample. The output value y_n is stored in the baseline memory.

Test mode

The look up table is used to generate a pattern to be injected into the processing chain for test purposes.

4.5.2 Tail cancellation

The positive ions, created in the avalanche of the MWPC, induce a positive current signal on the pad plane. This current signal is characterized by a fast rise time (less than 1 ns) and a long tail with a rather complex shape, which depends on the details of the wires and pad geometry. The signal tail increases the superposition of subsequent pulses (pile-up) rendering the zero suppression quite inefficient. In order to minimize such effects, the ALTRO chip incorporates a filter for the cancellation of the signal tail. The tail cancellation filter is a 3-stage IIR digital filter. The filter is composed of 3 first order filters in cascade. The filter operation on the data is selected by changing any of the 6 programmable coefficients of each filter. The use of the Tail Cancellation Filter in the processing chain is optional.

4.5.3 Second baseline correction

A second correction of the baseline is performed based on the moving average of the samples that fall within a certain acceptance window. This mechanism can remove non-systematic perturbations of the baseline that are superimposed on the clusters.

4.5.4 Zero suppression

In order to compress the data stream "zero" data is discarded. By "zero" data is meant sampled values below a certain reference level (pedestal). These data are considered as noise. The Zero Suppression procedure removes all data below the pedestal threshold, except for a specified number of pre- and post-samples around each peak. The number of

pre- and post-samples can vary independently in the range between 0 and 4. A glitch filter checks for a consecutive number of samples above threshold, confirming the existence of a real pulse. The minimum sequence of samples above the threshold which defines a pulse can vary from 1 to 3.

4.5.5 Data format from ALTRO

The stream of zero-suppressed data is formatted by adding, to each set of samples, two extra words, and encoding the 10-bit words and hardware address into a 40-bit set of words. Due to the removal of a variable number of samples between accepted clusters, the timing information would be lost during the zero-suppression process without the addition of time information to each accepted set of samples. Since the maximum length of the data stream that can be processed by the ALTRO chip is 1000, the time information can be encoded in a 10-bit word. The principle is to label each sample with a time-stamp that defines the time distance from the trigger signal. The samples of the processed data stream are numbered starting from 0 to 1000. The time information added to each cluster during the formatting phase corresponds to the time-stamp of the last sample in the cluster. The ALTRO data format does not make use of extra flag bits to distinguish the samples data from the time data, but introduces an extra word for each accepted cluster, which represents the number of words in the cluster without counting the time data. These two 10-bit words are added at the end of the cluster data words. The procedure is summarized in figure 4.4. The 10-bit words are packed in 40-bit words. If

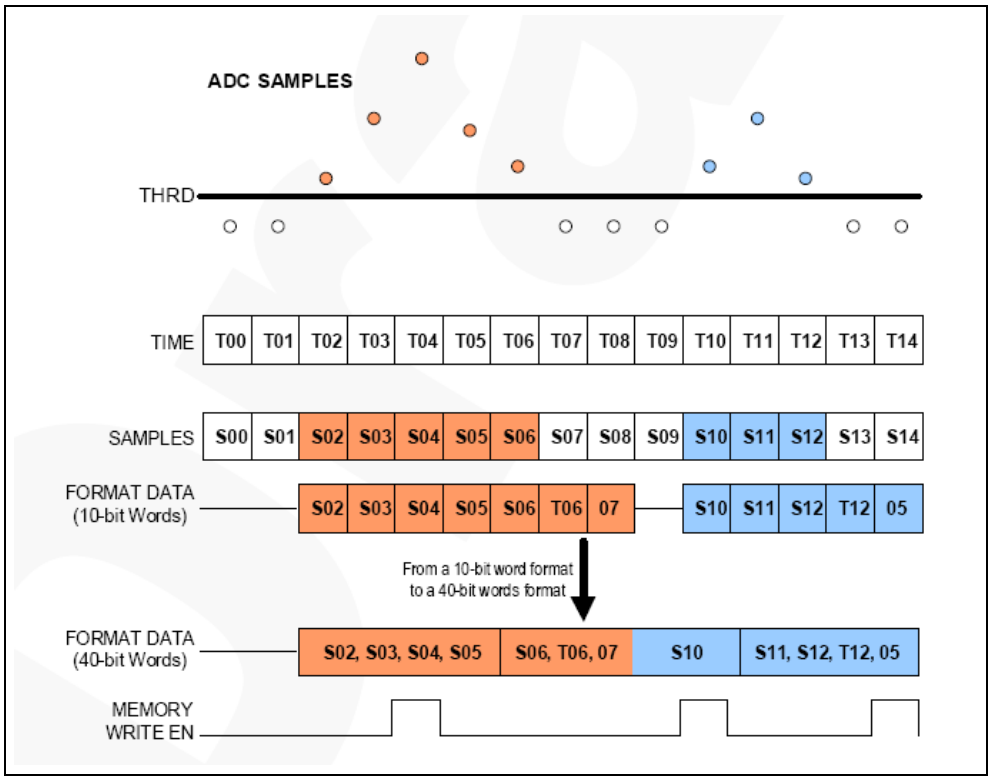

Figure 4.4: Data formatting procedure [71].

some data is missing to complete a 40-bit word an "A" hexadecimal pattern is used. A

trailer completes the data packet, which is the last 40-bit word of the data structure. The 40-bit trailer word is composed of the total number of 10-bit words in the packet (10 bits), and the hardware and channel address (8 and 4 bits respectively). The combined hardware and channel address represents a sort of geographical address and is used in the data packet to identify unambiguously to which channel the data packet is associated. The rest of the trailer word is filled with a pattern ("A" hexadecimal) in order to have the information available in bytes. The format of an ALTRO data block is illustrated in figure 4.5.

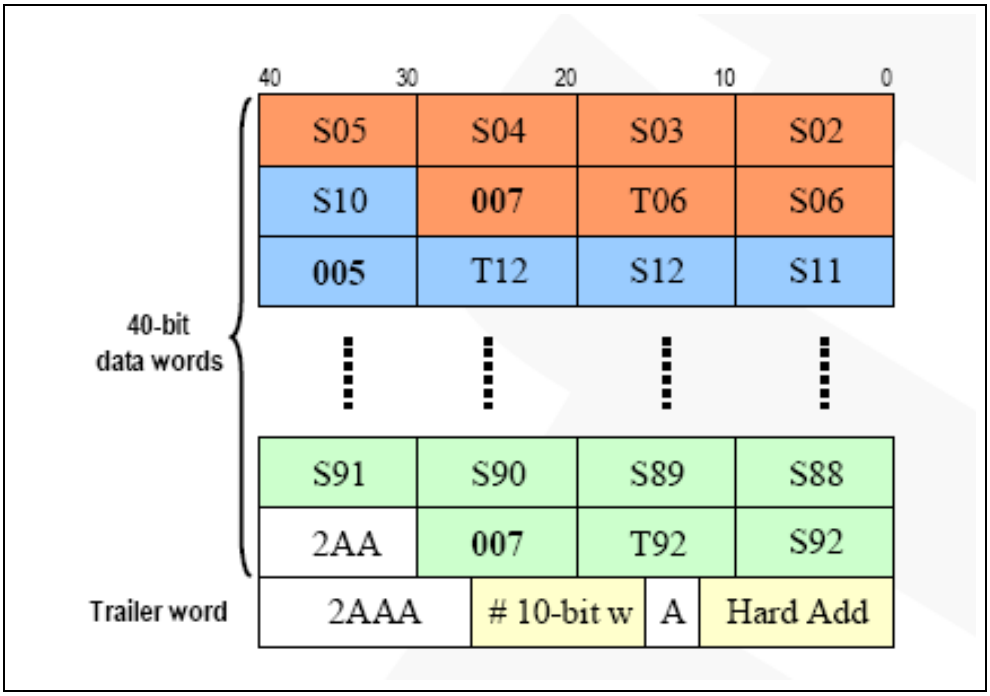

Figure 4.5: ALTRO data block [71].

4.6 Front End Card (FEC)

The Front End Card contains 8 ALTRO custom ASICs, and 8 custom pre-amplifier/shapers (PASAs). Thus the FEC contains the complete chain to amplify, shape, digitize, process, and store the signals from 128 channels. Connectors to the Front End Bus (see section 4.7) and Front End Slow Control (I^2C link) are also mounted on the FEC. The Front End Slow Control provides the RCU with an independent access to the FEC via an FPGA-implemented Board Controller on the FEC. This secondary access is normally used to control the state of the voltage regulators and monitor the board activity, power supplies and temperature. In order to match the position of the connectors on the back of the chamber pad plane, the FEC has a width of 190 mm. The height and thickness are 170 mm and 14 mm respectively, in order to fit into the available space. The FEC has a maximum power consumption of about 6W. In order to minimize the heat transfer to the detector sensitive volume, the FECs are embedded in two copper plates cooled by water. The assembled FEC is shown in figure 4.6

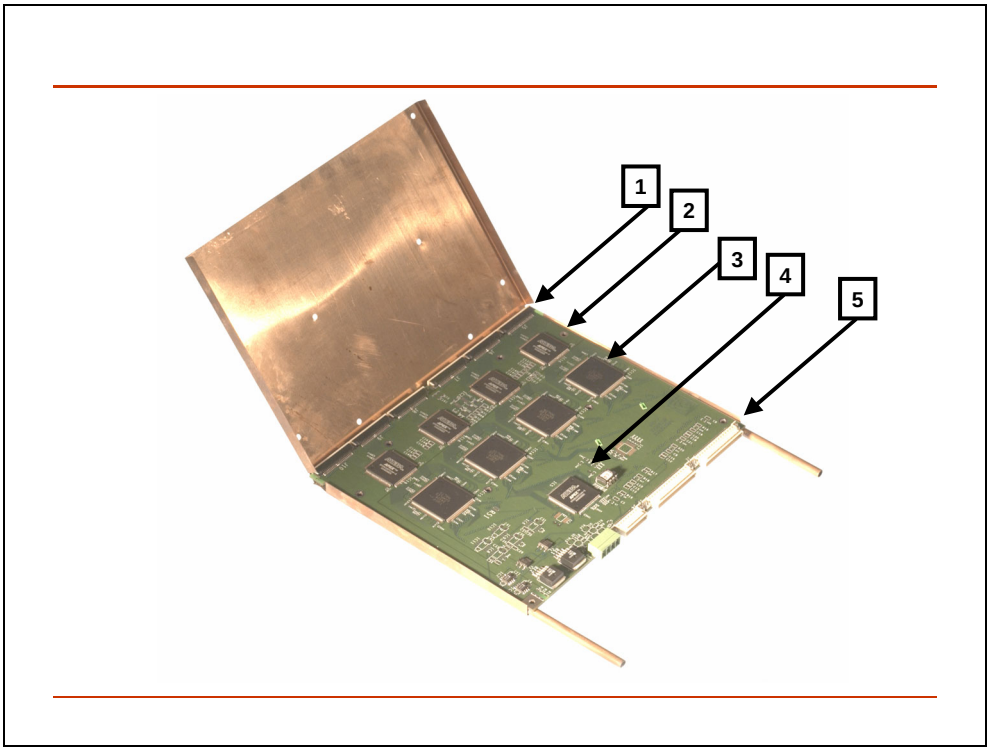

Figure 4.6: The assembled Front End Card enveloped in two copper plates. The plates are opened for illustration purposes. Indicated is 1. The connectors for the capton cables to the detector pads, 2. The 4 PASAs (4 more on the backside of the PCB), 3. 4 ALTROS (again 4 more on the backside), 4. The Board Controller (BC) and 5. the connectors for the Front End Bus.

4.7 Front End Bus

The communication between the RCU and the FECs is implemented via a custom bus, the ALICE TPC Read Out bus (ALTRO bus), based on a ribbon cable and a custom protocol. The ALTRO bus is essentially an extension of the FEC's internal bus and allows the RCU to access the FEC's internal components: the ALTRO chips and the Board Controller (BC).

The ALTRO bus to the front-end provides, with a clock frequency of 40 MHz, a bandwidth of 200 Mbytes/s on each of the two branches, leading to an aggregate bandwidth of 400 Mbytes/s. Depending on the Huffman data compression done on the RCU, the clock frequency of the FEBus read out clock will be adjusted to match the bandwidth of the DDL. However a FEBus with a higher bandwidth allows to support possible future upgrades of the DDL performance. The electrical implementation of the bus transceivers is done using GTL/GTL+ technology [81]. The transceivers are mounted on the RCU and FECs. The ALTRO instructions and the Front End Bus protocol is further discussed in section 5.5.2.

4.8 Readout Control Unit

The Readout Control Unit interfaces the Front End Electronics of the ALICE TPC with the ALICE DAQ, the ALICE CTP (via TTCrx chip) and the ALICE Detector Control System (DCS).

4.8.1 RCU Requirements

The RCU must be capable of reading out up to 400 Mbyte/s from the Front End Electronics, through the Front End Bus, operating at a maximum of 40 MHz, and transferring the data onto the DDL optical link. The physical size of the PCB should be minimized, but a width of 225 mm and a height of 150 mm is sufficient. It should also host the interfaces to the DCS, Front End Slow Control and Trigger according to their chosen technologies and protocols. The RCU has the following functional requirements:

- It should merge digitized data from FECs into the DDL in ordered form (time bin, pad number, pad row),
- decode trigger data and add trigger information to data stream,
- distribute trigger and clock signals to FECs,
- interface Detector Control System to FECs,
- supervise the functionality of itself and the connected front end electronics devices,
- optionally perform lossless data compression on the TPC data before they are sent to the DDL

It should be radiation tolerant. Further details on the inclusion and fulfilment of these requirements are presented below and in the next chapters.

4.8.2 Readout of TPC data

The basic functionality of the RCU is to be a data collector from the Front End Cards (FEC - up to 25 connected to one RCU). The RCU controls the readout procedure by sending ALTRO instructions on the Front End Bus to the Front End Cards. The collected data is transferred to the Detector Data Link (DDL), after being marked by a data header including event identification information.

4.8.3 Trigger interface

The trigger information received on the TTCrx is decoded by the RCU and trigger information added to the data header of the event. Through the Front End Bus the RCU broadcasts the Level-1 and Level-2 trigger information to the connected FECs and ALTROs.

4.8.4 Configuration of ALTRO/FEC

Besides these basic functions, the RCU will also download the FEC's configuration parameters (pedestals, thresholds, tail cancellation coefficients, etc) and run test procedures via the ALTRO bus.

4.8.5 Bandwidth optimizing

The accepted ALICE data transmission standard off the detector is the DDL, supporting 200 Mbyte/s. Assuming black events (no zero-suppression on the FEC) and the maximum 200 Hz TPC readout each Front End card will produce a data stream of about 11 Mbyte/s. However, depending on the radial position, not the total depth of the chamber will be read out and also position dependent compression factors are expected. Therefore the number of FECs merged into one DDL varies from 18 - 25. This functionality is performed by the RCU.

4.9 Detector Control System

The TPC Detector Control System is a part of the overall ALICE Detector Control System. Figure 4.7 shows the ALICE Online system hierarchy, including the DCS. The

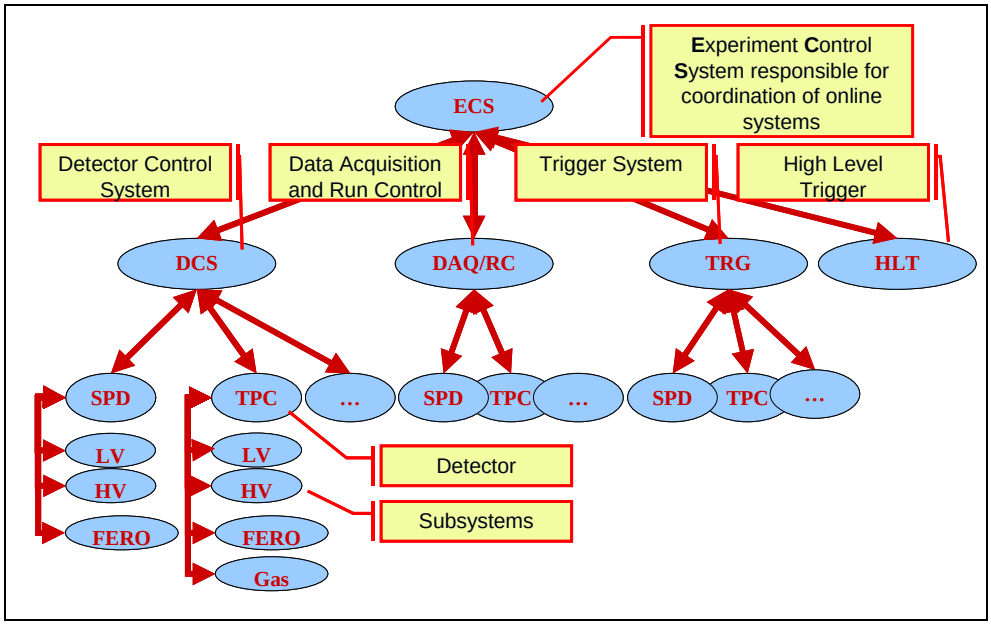

Figure 4.7: ALICE Online systems hierarchy.

control and operation of the experiment takes place on the ECS layer. Through this layer DCS is linked up with the Data Acquisition (DAQ), Trigger (TRG), High Level Trigger (HLT) and with the LHC Machine. The "standard" DCS model consists of a number of hardware and software components on three hierarchical layers: the supervision layer, the control layer and the field layer.

4.9.1 Supervision layer

The supervision and operation of the detectors take place on this layer and some of the tools and facilities needed are:

- SCADA - Supervisory Control And Data Acquisition - framework tools
  - Man-Machine-Interface and graphics
  - Logging and archiving facilities
  - Alarms
  - Configuration databases
- Application programs and data
  - Object library for supervision
  - Supervision and operation programs

4.9.2 Control layer

On the control layer the collection and processing of data takes place and the following components can typically be found here:

- Process Control Units
- Local Operator Workplace
- Local Engineering Workplace
- Data server
- Application programs and data
  - Object library for process control
  - Local operation programs
  - Sub-detector configuration data

PVSS II is a SCADA system. PVSS [82] will be used by the ALICE DCS to connect to hardware (or software) devices, acquire the data they produce and use it for their supervision, i.e. to monitor their behaviour and to initialize, configure and operate them.

4.9.3 Field layer

For the access to the hardware of the front end electronics of the TPC, the DIM - Distributed Information Management system [83] - has been chosen. This uses TCP/IP as the communication protocol, and interfaces to the PVSS. The DIM implementation for the TPC is illustrated in figure 4.8. DIM servers are running on the DCS FPGA on every RCU, and DIM clients are subscribing to services from the DIM servers. The request for service is answered by a DIM DNS (Domain Name Service) server. The hardware layer between the DIM server and client is Ethernet. The architecture proposed is presented in figure 4.9. As illustrated the bandwidth of the Ethernet line to each RCU is 100 Mb/s.

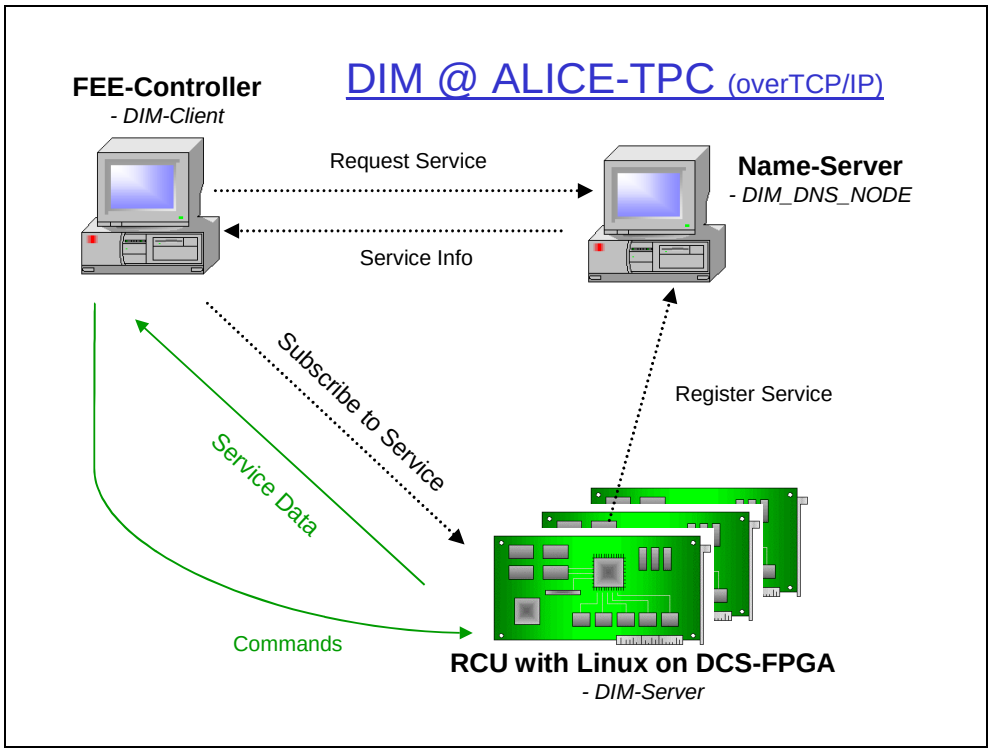

Figure 4.8: The DIM implementation for the TPC-FEE [84].

4.9.4 The RCU DIM server

The DIM server is named FEEserver. It communicates with the Command Engine, as illustrated in figure 4.10. The FEEserver is responsible for receiving and sending blocks of commands using DIM. The Command Engine does the unwrapping of the blocks and execution of the commands in hardware.

The interface of the Command Engine to the FEEserver consists of three functions (see figure 4.10):

- void initializeCE() - Called by FEEserver on startup. Any code that must be executed in the command engine at startup should be placed here.
- void cleanupCE() - Called by FEEserver on shutdown. Any code that must be executed in the command engine at shutdown should be placed here.
- int issue(char* input, char** result, int* resultSize) - Called by FEEserver when a package of commands is to be processed. input is a pointer to this package. result is allocated by the Command Engine and will be the package containing the result blocks. resultSize is the size of the result package in bytes. The function returns 0 if everything is OK, or an error code between -20 and -39 if an error occurred when processing the package of commands. The issue function unwraps the blocks in the input package one by one, and executes the correct command (single read, single write, multiple write, etc. - see section 5.5.1) by accessing hardware.

The Command Engine communicates with the DCS Interface Module. This module is described in section 5.5.1.

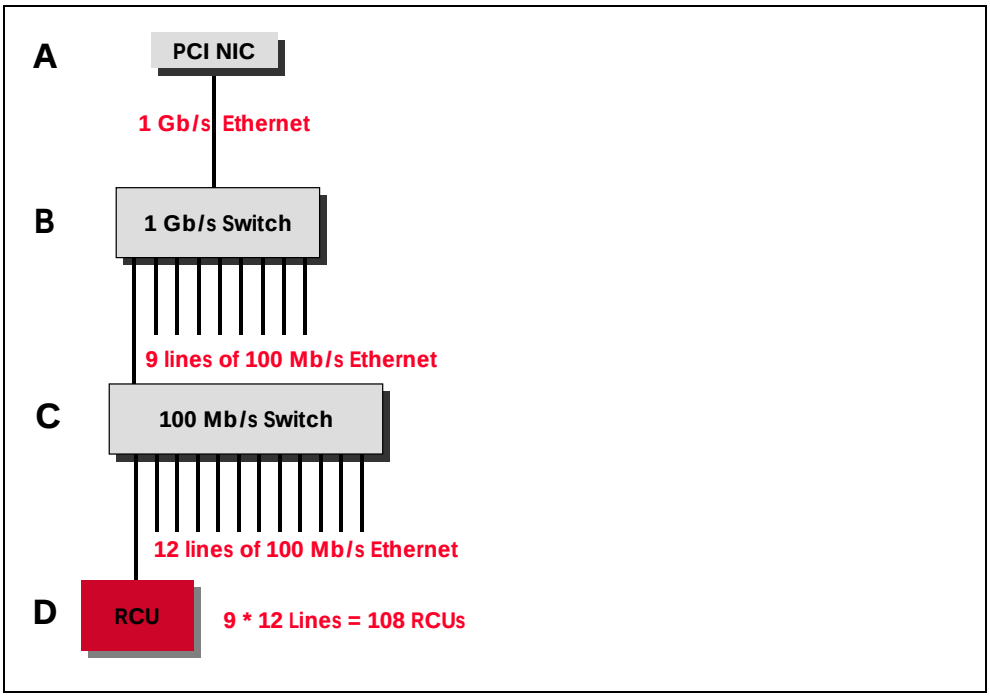

Figure 4.9: Architecture of the DCS Ethernet connection.

4.10 ALICE Data Acquisition (DAQ)

The overall ALICE Trigger and Data Acquisition system is shown in figure 4.11, which will be further explained in the coming sections describing DAQ, Trigger and High Level Trigger (HLT).

4.10.1 Detector Data Link

The interface from the RCU to the ALICE DAQ is the Detector Data Link (DDL). This is based on two multi-mode optical fibers transmitting at 200 MB/s in both directions if used in half duplex mode. The detected bit error rate is less than 10^{-15}. On the detector side, the optical fibre is equipped with a transceiver named the Source Interface Unit. This is a PCB with a form factor based on the Common Mezzanine Card (CMC) format (half size CMC: 149 x 40 x 10 (mm)). The SIU card is driven by a 3.3V LVTTL signal, and connects to the RCU via 2 CMC format 64-pin connectors. The logic of the SIU design is put into an onboard FPGA. On the other end of the fibers the DDL destination interface unit (DIU) are connected to the Read Out Receiver Card (D-RORC), located in the counting room about 200 meters from the detector. The D-RORC interfaces the DDL to the PCI bus of the host LDC computer. The D-RORC includes two DDL interfaces. For the detectors whose data are analyzed using the HLT, one of the DDL interfaces is used to transfer a copy of the raw data to the HLT computers. An event data transmission transaction can be opened by using a Ready to Receive command. The RORC may only send this command to the RCU through the DDL, when it is ready to receive event data. When the RCU receives this command, it can start to transfer event data blocks to the SIU. Any open block transfer transactions must be closed by the RCU, when it receives

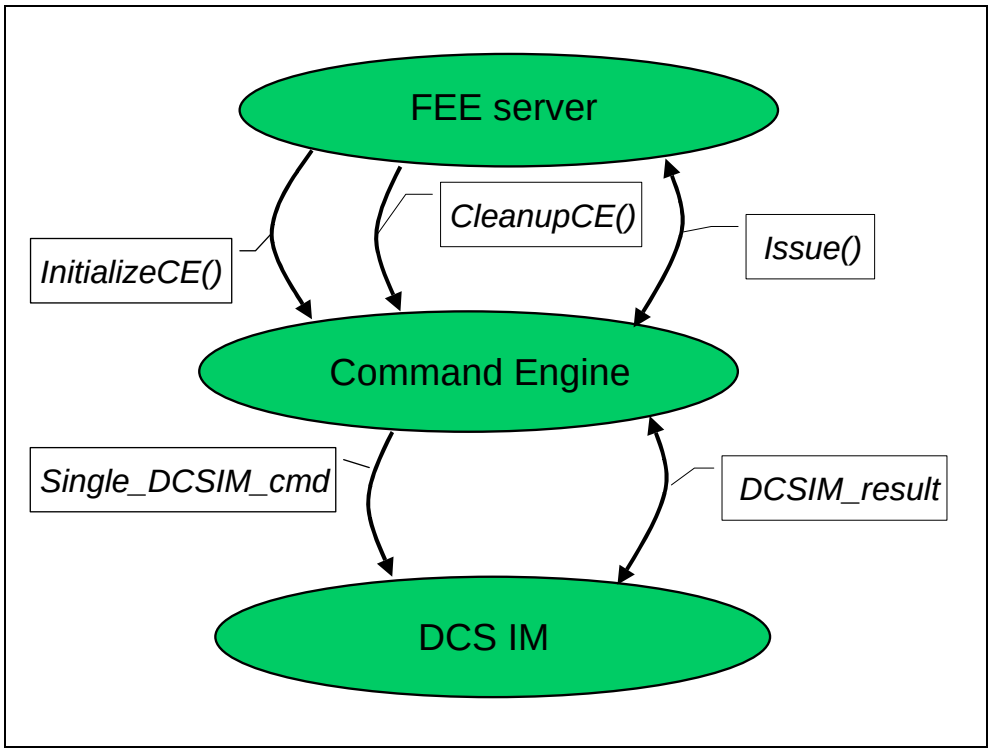

Figure 4.10: Implementation of DIM in the DCS FPGA.

an End of Block Transfer command. Thus the readout of (TPC-)data is started and ended by the ALICE DAQ via the RORC.

4.10.2 Data format over the DDL

The data from the detectors is sent in blocks. The format of these data blocks is described in [80]. A summary of the data block header is given in figure 4.12. As can be seen the RCU has to collect information from the trigger interface and count the number of data words coming from the ALTROs, before adding this information to the data being sent on the DDL. The TPC data itself, which is following the data header, is sent in the data format coming from the ALTROs. This format is described in section 4.5.5.

4.10.3 DAQ-system

The complete ALICE DAQ system (figure 4.11) will contain about 400 DDLs. The acquisition of data is done in a number of steps. First, the data is collected by the RCU (or the dedicated front end electronics for other detectors) after receipt of a L2 trigger. The data transfer from the readout electronics is performed in parallel for all the sub-detectors over the optical links (DDLs). The DAQ system then performs sub-event building inside the Local Data Concentrators (LDC). High level trigger (HLT) can be performed in the HLT farm. Following the decision of the High Level Trigger (if activated), the event building, partial or total is performed by the Global Data Collectors (GDC). After event building, the data are archived. 1.25 GB/s is the bandwidth foreseen

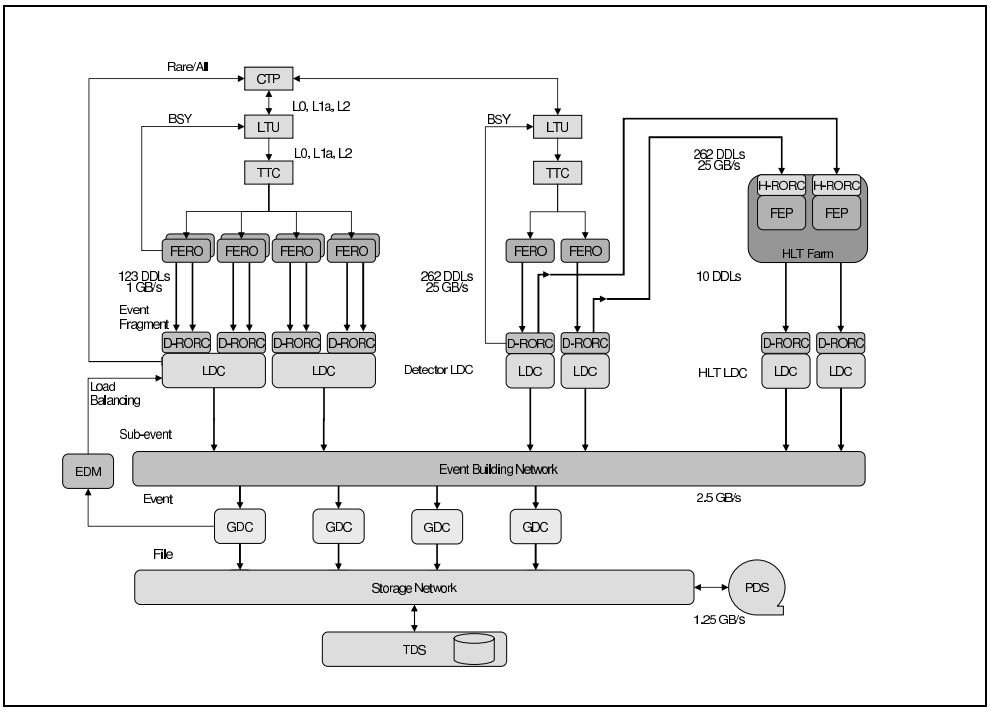

Figure 4.11: ALICE Trigger-DAQ-HLT overall architecture [54].

to mass storage. The data will be stored temporarily at the experimental area in the so called Transient Data Storage (TDS). This allows an independent functioning of the ALICE DAQ for a few hours even in case of network or permanent data storage failures. The Permanent Data Storage (PDS) will be provided by CERN as a facility available through the Storage Network.

4.11 ALICE Trigger

The ALICE Trigger system is designed to handle trigger priority and past-future protection. Some trigger types (for example high p_t dimuon and electron triggers) are much more rare than other triggers in ALICE. To enhance these rare triggers, the trigger system must determine the trigger type of a given interaction and determine if the appropriate detector set is not busy. This is called trigger priority and means that, in general, different detector sets can be read out in consecutive events. The past-future protection is there to protect events in detectors with long sensitive times from being spoilt by pile up. For example, the drift time of the ALICE TPC is 88 μs. Events in which the detector would also register tracks from earlier and later interactions are deemed unreconstructable owing to the very high track densities in Pb-Pb interactions.

4.11.1 Central Trigger Processor (CTP)

The ALICE CTP is designed to combine and synchronize information from all the triggering detectors in ALICE, and to send the correct sequences of trigger signals to all detectors in order to make them read out correctly. It also coordinates calibration re-

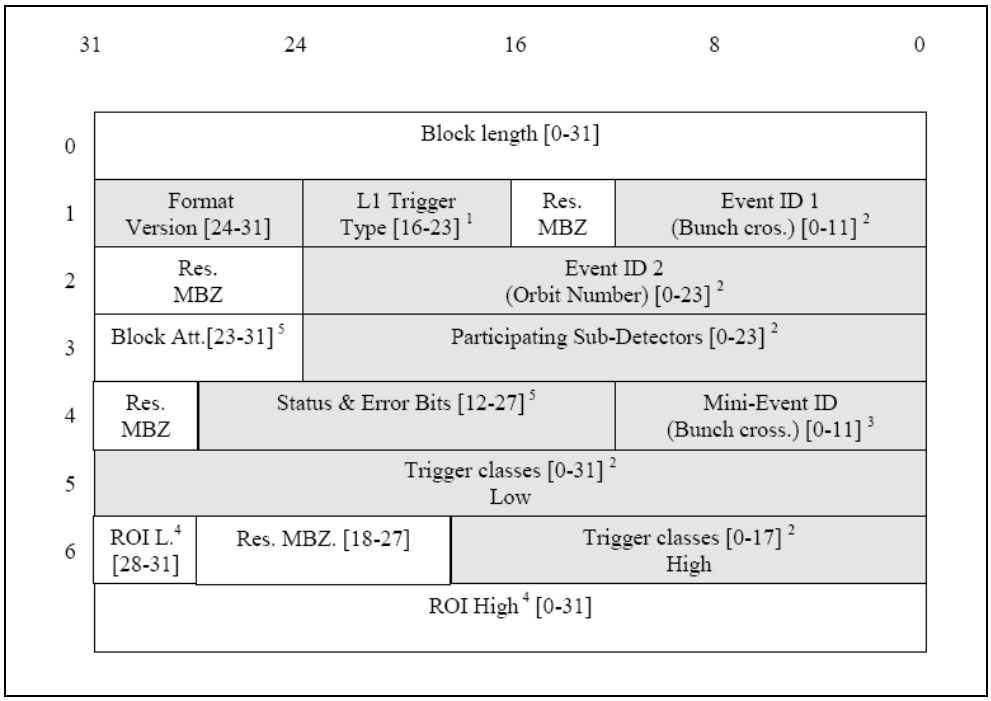

Figure 4.12: The ALICE DAQ common data header layout [80].

quests from sub-detectors and generates data summarizing why a particular trigger has been taken.

4.11.2 Local Trigger Unit (LTU)

The LTU performs several operations with the signals received. Among the most important are:

- re-synchronization of the received data and strobe signals
- de-serialization of the L1 data and L2 data messages, and conversion to the appropriate format for TTC(rx) transmission

4.11.3 TTCrx

The interface to the ALICE trigger is the TTCrx custom integrated circuit, developed by the CERN EP Microelectronics group [79]. This is packaged in a 144 pin BGA, and fabricated in radiation-hard DMILL technology. The trigger data is sent from the Central Trigger Processor to the TTCrx via optical fibre. The most important data coming via the TTCrx is the L1 trigger and the event data being distributed as a L2a message. In this message the event identifier information (orbit number and bunch crossing number), trigger type (physics trigger, software trigger, calibration trigger), detector set to be read out and the list of active trigger classes is included. From this information the L2 trigger signal is generated, and the identifier data to go into the data block header is extracted.

4.11.4 Trigger types

In addition to the physics triggers there are a number of other trigger types. Among them are periodical, random and software triggers. Also included in the ALICE trigger system are the calibration triggers. Most of these trigger types are for test and calibration purposes, but can also be used in combination with physics triggers.

Chapter 5

The RCU conceptual design

5.1 Introduction

The Readout Control Unit interfaces the Front End Electronics of the ALICE TPC with the ALICE DAQ, the ALICE CTP (via TTCrx chip) and the ALICE Detector Control System (DCS). The logic of the RCU is designed to a large extent in VHDL and put in dedicated Field Programmable Gate Arrays (FPGA) on a custom developed RCU Printed Circuit Board (PCB). In this chapter the first conceptual design of the RCU is presented. During the development of the different prototypes presented in the next chapter, changes were made to this design. The changes will be discussed in chapter 7. However the first concept of the RCU is presented here since it forms the basis upon which the later development was done.

5.2 RCU functional overview

The logic of the RCU is split into 8 memory mapped modules as listed in table 5.1. 3 of these modules (the RCU Master Module, DCS Interface Module and Configuration Module) can be masters of the RCU local bus and read and write to the other modules. The 5 interface modules are communicating with the interfaces to the RCU according to their protocols. Each module has a status register. This is used to set the operation mode of the module and for monitoring purposes. All the modules are described in further detail in separate sections below.

5.3 RCU local bus

The 3 Masters can request the RCU local bus. Once granted the bus they can start to do transactions on the bus following the RCU local bus protocol (section 5.3.2). The RCU local bus contains 16 address lines, 32 data lines and 3 control lines (Read not Write (RnW), Control Strobe (CSTB) and Acknowledge (ACKN)). The 3 upper bits of the address bus (A15-A13) are the module addresses, as defined in table 5.1. Thus, each module has an addressable memory size of 8192 words. Each module has a tri-state interface to the bus, and must attach or detach itself to the address and data-lines.

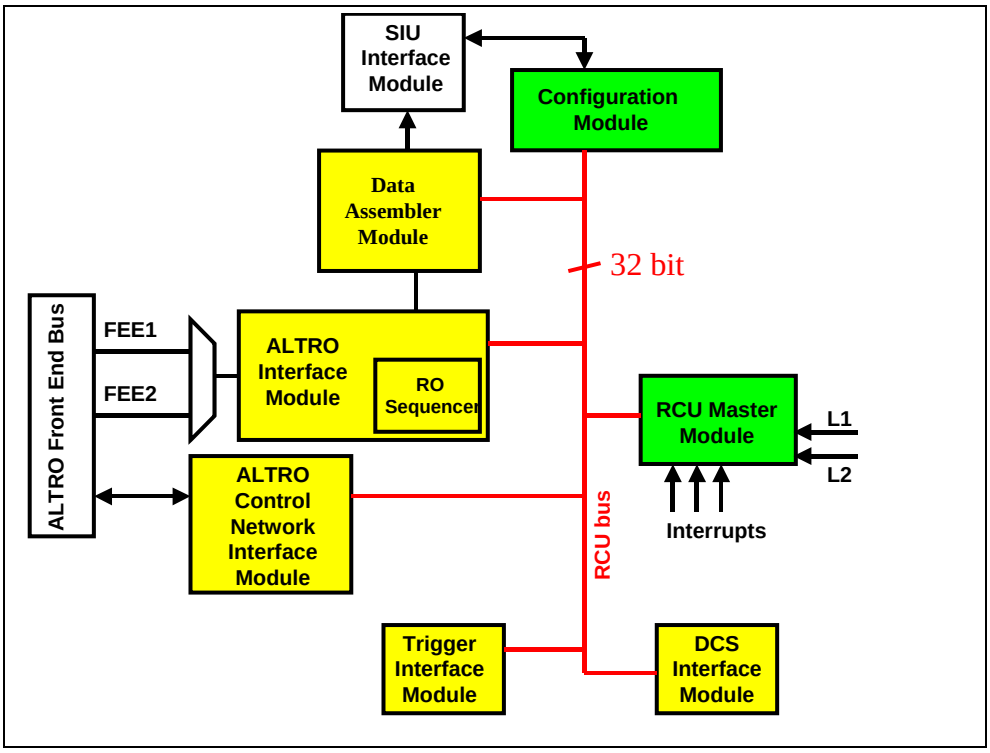

Figure 5.1: Overview of the different modules the control logic of the RCU is organized in.

5.3.1 Bus Arbiter

The Bus Arbiter controls the access to the bus. According to the priority list the three masters are granted access to the bus. As can be seen from the priority list there are multiple levels of requests possible from each master (Error and Warning typically). If a grant has been given, and a new request comes from someone higher up on the priority list, the current grant is removed. However the new grant is not issued by the Bus Arbiter for 10 RCU clock cycles, in order to allow the current grantee to finish the ongoing transaction. Priority list:

- 1. DCS Safety Message
- 2. FE Control Network Error (RCU Master)
- 3. RCU Monitoring (RCU Master)
- 4. L2 (RCU Master)
- 5. DCS Monitor Message
- 6. FE Control Network Warning (RCU Master)
- 7. Configuration Module

Module Address (Hex)	Module Name
0	Configuration Module
1	SIU Interface Module
2	Data Assembler Module
3	ALTRO Interface Module
4	ALTRO Control Network Interface Module
5	Trigger Interface Module
6	DCS Interface Module
7	RCU Master Module

Table 5.1: The 8 memory mapped modules of the RCU design.

5.3.2 Bus protocol

Any transaction is controlled by the 3 control lines: RnW, CSTB and ACKN. Any Address (and Data if write) must be valid as long as the CSTB is asserted. The target will respond with an assertion of ACKN. Only after this assertion is received by the master, CSTB can be deasserted.

Read Instructions

The chronogram of a read instruction is shown in the left part of figure 5.2. The address

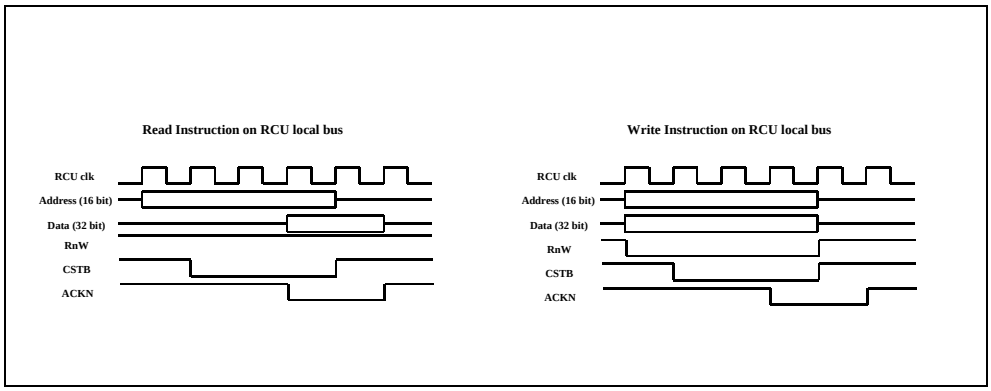

Figure 5.2: Chronograms of a read instruction (left) and a write instruction (right) on the RCU local bus.

is put on the address line. One clock cycle later the CSTB is asserted. The RnW stays high. When the target has received the address and put the data requested on the data lines, it asserts the ACKN. The data is then read by the master, and CSTB deasserted.

Write Instructions

The chronogram of a write instruction is shown the right part of figure 5.2. Address and Data is put on the address and data lines. The RnW is put low. The CSTB is then asserted. When the target has received the address and put the data requested on the

data lines, it asserts the ACKN. The data is then read by the master, and the CSTB deasserted.

5.4 The RCU Master Module

5.4.1 Basic functionality

This module resets all other modules of the RCU logic when necessary. It receives trigger signals (L1 and L2) and copies the trigger messages from the Trigger Module into the Data Assembler Module. The read out of data via the ALTRO Interface Module and the Data Assembler Module is controlled by this module. The RCU Master Module also takes care of the internal monitoring of the RCU main board.

5.4.2 RCU Monitoring

Activities:

- Monitor power consumption on RCU board
- Main watchdog of RCU
- Reset control for all modules
- Single Event Upset monitoring

The reset of the other modules is activated via a dedicated register. Setting the dedicated bit of this register high executes a reset of the module in question.

5.4.3 The Trigger Handler

Level 1 and Level 2 are received from the Trigger Interface Module via dedicated lines. Two separate subprocesses handle Level 1 and Level 2. The Level 1 is counted and passed on to the ALTRO Front End bus via a dedicated line. The incoming Level 2 is generated by the Trigger Interface Module, and signals that a L2accept has been received from the TTCrx. This is also counted by the Trigger Handler process in the RCU Master Module, and checked with local past-future protection. If this is OK the L2 is passed on to the ALTRO Interface Module, else it is marked as an error in the status register. Immediately after an accepted Level 2 the Trigger Handler process also requests access to the RCU data bus. When bus access is granted, the 7 Trigger information words (see table below) are read from the Trigger IM. This information is put in the Event Identification registers in the Data Assembler Module, along with the Status and error bits signalling trigger overlap or trigger missing. When the RCU is in configuration phase, incoming triggers will be ignored and marked as error. They will not be passed to the ALTRO Interface Module during this phase. The Trigger Handler process will not request bus access from the Bus Handler while the RCU is in configuration phase.

5.4.4 The Front End Error Handler

This process is the interface between the RCU Master Module and the Front End Control Network IM. There are two levels of messages coming from the Front End Control Network IM. These are FE Error and FE Warning. Bus access is requested upon receiving any of these on the dedicated signal lines. The bus access needed to read the messages are prioritized as shown in the section describing the Bus Handler process above. When the bus access is granted, the error is read, and action is taken as needed. If the error is serious, the RCU will send information to the DCS, and generate commands to be executed by the Front-End Control Network Interface Module (to the Front End Cards). One typical error message coming from the FE Control Network IM is that one or more channels are drawing too much current. The typical response would then be to switch off a part of, or the whole Front End Card in question. The Front End Error Handler, who requests the bus to write to the Status register bit(s) in the FE Control Network IM, does this. The Front End Error Handler will then update the Active Channel register in the ALTRO Interface Module. Finally it will notify the RCU Monitoring process, which in turn notifies the DCS IM.

5.5 The 5 Interface Modules

5.5.1 DCS Interface Module - Interfacing the DCS

The DCS Interface Module is responsible for communication between the RCU and DCS (Detector Control System). The DCS Interface Module can get access to the internal RCU bus for data transfer to/from other RCU modules. In external communication to the DCS bus, the DCS Interface Module appears as having two FIFO queues (for incoming and outgoing messages). Internal data transfer within the RCU is done using the bus protocol described in section 5.3.2, and at this side the DCS Interface Module memory appears as directly addressable RAM.

DCS requests are sent over the DCS bus as packets containing control information, commands and address/data words. DCS requests can be transmitted as "safety messages", which will interrupt the RCU Master Module, and "monitoring messages" which will be handled when the requested RCU resource is available. When a DCS request is received, the DCS Interface Module interprets the first command word. This contains information on whether this is a "safety" or "monitor" message, and the corresponding signal is sent to the Bus Arbiter. The RCU Master Module/Bus Arbiter will always respond to a Safety Message or Monitoring Message signal within a configurable time. If the DCS Interface Module request has not been granted before 100 RCU clock cycles after the Message signal is set, the DCS IM must consider the message as rejected. The DCS Interface Module then continues interpreting any new block received from the DCS.

If RCU local bus access is granted for a DCS message, the DCS IM continues to interpret the DCS command, and the requested data are stored/loaded to other RCU memory mapped modules. If the message is graded as a safety message, the DCS Interface Module continues to control the RCU bus, until the bus control is released by itself. If the current message is a monitor message, the bus access can be removed with a latency of 10 RCU clock cycles. Should this happen during data transfer of the message, then this

message must be considered as rejected. This will be marked as an error in the response message returned to the DCS.

DCS commands

The DCS Interface Module recognizes and interprets the following commands (with description of the words in the command):

single read address word

single write address word, data word

multiple read address word, number of data words to read

multiple write address word, number of data words, data word 1,.....,data word N

random read address word 1, address word 2,.....,address word N

random write address word 1, data word 1, address word 2, data word 2,.....,address word N, data word N

The single read/write commands will transfer one word of data from/to RCU memory/registers. The multiple read/write commands will transfer a block of data from/to subsequent memory locations in the RCU, but only from one memory module per command. The random read/write commands will transfer one data word from/to a given location, followed by another word to another location etc. Random write requires exactly one data word for each specified address word. Random read/write should normally not take place parallel to other RCU activity (not in Readout Mode or Configuration Mode). A marker word follows each command. Hence in the case of a large number of registers, the random read/write commands will occupy approximately half the number of bytes compared to a set of single read/write commands.

Data formats for DCS communication

Configuration data format

The configuration data is split into programming blocks. The size of these blocks must match the memory in the RCU they are to be loaded into. The format of the configuration data is shown in figure 5.3. The first word in each block is an information word, which contains:

Bit 0-15 Number of words. Number of 32 bit data words in command block, including information word and marker word.

Bit 16-23 Block number (decremented)

Bit 24-25 Not used/Don't care. Available for future use.

Bit 26 Packed format bit (1 if the message contains packed 3×10 bit words. Used for Pedestal Data)

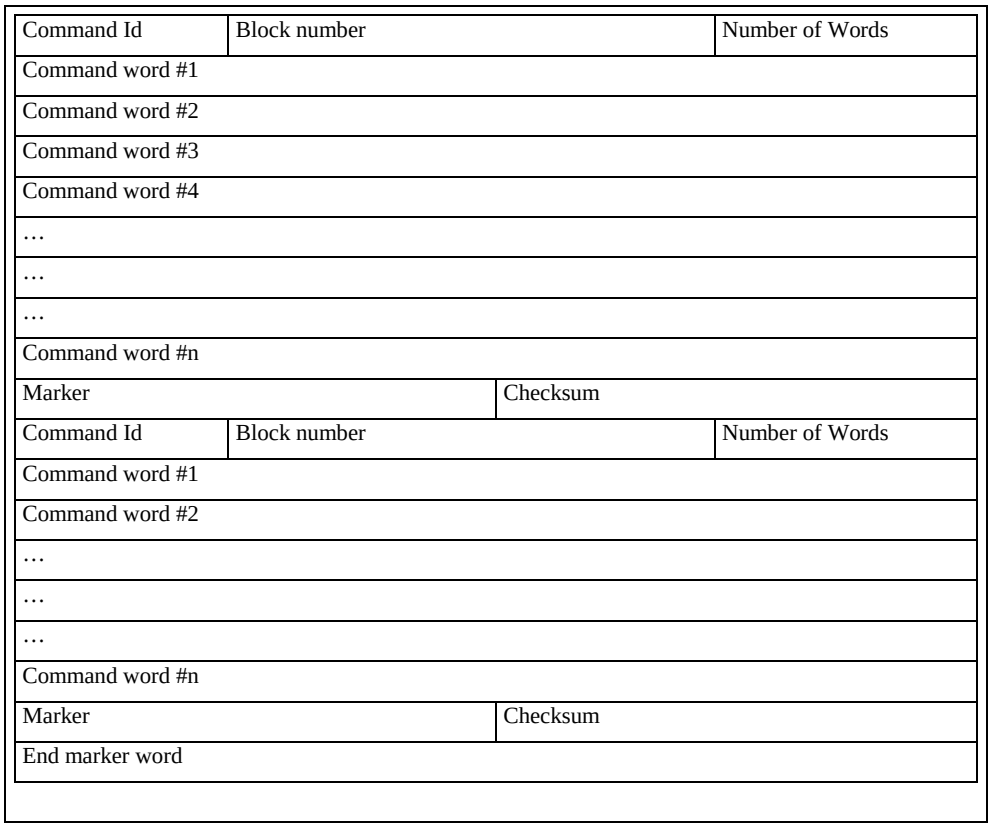

Figure 5.3: Format of the configuration data.

Bit 27 Safety message identifier (1 if the message is a safety message, 0 if not)

Bit 28-31 Command identifier (See figure 5.4 for contents)

The command words then follows. They are coded as described in section 5.5.1. The command words are followed by a marker word. This contains:

Bit 0-15 Checksum, which will be compared with the checksum computed by the Configuration Module/DCS IM.

Bit 16-31 Marker word identifier, coded as AA55 (hex).

After the last programming block (recognized by Block number equal zero) an extra word - the End Marker - is added. This contains:

Bit 0-15 Not used/don't care.

Bit 16-31 End marker identifier, coded as DD33 (hex).

The Command Identifier is coded as shown in figure 5.4

The left column shows the binary coding of the 4 Command identifier bits. The DCS Interface Module or the Configuration Module interprets them. The modules execute all commands, depending on which RCU operational mode is active. The Change RCU Mode command has the same format as the Single write command, but uses an unique identifier.

Identifier	Command
0001	Single read
0010	Single write
0011	Multiple read
0100	Multiple write
0101	Random read
0110	Random write
1101	Reserved (end marker)
1111	Change RCU Mode

Figure 5.4: Coding of the Command Identifier.

RCU response message format

The RCU will answer all messages sent from the DCS. The format of the response is shown in figure 5.5 The response will be sent word by word, starting with the Requested

Command Id	Block Number	Number of words
Status word		
Requested data #1		
Requested data #2		
.....		
Requested data #n		

Figure 5.5: Format of the RCU response message.

data #1 word put on the DCS Data out signal lines with DCS Address out signals equal to the value 2(hex). Then all words containing the requested data is sent incrementally. Following the data, the Status word is sent from the DCS IM (see below for format of status word). Finally a copy of the message first word (command id, block number and number of words) is sent. However, the number of words is now updated with the number of words sent as a response from DCS IM (including this word).

The status word in each response block contains:

Bit 0 OK (1 if message has been executed ok).

Bit 1 Illegal mode (1 if the RCU is in a mode that do not allow this kind of message to be executed).

Bit 2 Timeout (1 if RCU bus access was not received 100 clk-cycles after request was sent).

Bit 3 Block abort (1 if bus access was lost during execution of block. Some requested data may be already received).

Bit 4-31 Not used/Available for future use.

5.5.2 ALTRO Interface Module - Interfacing the ALTRO/FEC

The ALTRO Interface Module is divided into the following parts;

- The interface to the Front End Cards via the Front End Bus.
- The interface to the Data Assembler
- The interface to the RCU data bus
- The Readout Sequence Builder

This module receives ALTRO commands to be executed and runs them on the Front End Bus. There are two ways to deliver such commands via the RCU data bus. Either they are stored in the Instruction Memory, for later execution of the full sequence of commands or they are received in the Instruction Register for immediate execution. The results of read register commands are put in the Result Memory. The module also handles the programming of the pedestals on the ALTRO. The pedestal values are stored in a dedicated pedestal memory.

Trigger signals and TPC Data Readout

The ALTRO interface module receives the L1 and L2 trigger signals from the RCU Master Module. The number of L1 and L2 trigger signals received are counted. These counters can be used to compare with the similar counters in the ALTRO chips. After the receipt of a L2 trigger all necessary ALTRO commands to conduct a full read out of TPC data, from all connected and functional ALTRO channels on the Front End Bus, are generated by the Readout Sequence Builder. The result of a channel read out (the TPC data) is made available for the Data Assembler Module. The clocking of the data to the Data Assembler Module is illustrated in figure 5.7. The ALTRO Interface Module keeps track of the number of events stored in the ALTROs, and checks each ALTRO to ensure that the correct event is readout. Any error information is added to the payload after the transfer for a complete channel. One or several error words is added after the transfer of all channels.

Front End Bus protocol and ALTRO instructions

The Front End Bus is composed of 40 bi-directional lines and 8 control lines. The 40-bit bus contains 20 address bits that define the ALTRO address space and 20 data bits.

The bus protocol is asynchronous for read and write instructions, with a 2-line handshake. When an instruction is issued, the CSTB line must be held low until the ALTRO asserts the line ACK. ACK keeps low until CSTB is de-asserted. The readout of data, however, is a synchronous block transfer. The CSTB and WRITE lines must be held low until ACK is asserted. The upper data lines must be valid during the assertion of CSTB. Three clock cycles after the de-assertion of ACK the chip will start driving the 40 data lines. On the following clock cycle, TRSF will be asserted and output data will

Signal name	Function	# bits	Dir.	Polarity
AD	Address/Data	40	Bi-directional	H
WRITE	Write/Read	1	Input	L
CSTB	Command Strobe	1	Input	L
ACKN	Acknowledge	1	Output	L
ERROR	Error	1	Output	L
TRSF	Transfer	1	Output	L
DSTB	Data Strobe	1	Output	L
LVL1	Level-1 Trigger	1	Input	L
LVL2	Level-2 Trigger	1	Input	L
GRST	Global Reset	1	Input	L
SCLK	Sampling Clock	1	Input	-
RCLK	Readout Clock	1	Input	-

Table 5.2: Front End Bus signal summary.

be valid on each falling edge of DSTB. One clock cycle after the de-assertion of TRSF the data bus will be in high impedance. The ALTRO has 137 registers. Out of these,

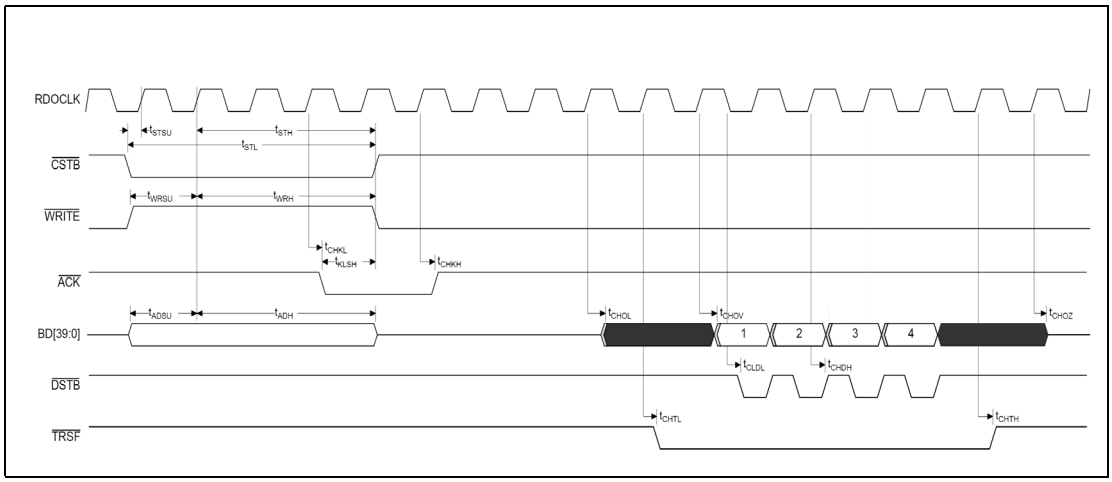

Figure 5.6: Front End Bus readout instruction chronogram.

128 are channel specific, that is, a different version exists for each channel. There are 8 channel-specific registers for each of the 16 channels ($8 \times 16 = 128$). These registers can be written and read to for configuration and status controlling purposes. In addition the listed commands can be used to control the ALTRO chip (see figure 5.3). ALTRO instructions can be broadcasted to all connected ALTROs. A typical readout sequence will consist of a Write Pointer Increase (WPINC) command broadcasted to all connected ALTROs after receipt of a L2 trigger signal. This is followed by a channel by channel readout (CHRDO) of all active channels addressed in sequence. Finally the Read Pointer Increase (RPINC) command is broadcasted in order to release the ALTRO buffer after readout of the TPC data.

Register name	Register Address	Access Type	Allow Broadcasting	Meaning
WPINC	18	W	Y	Write Pointer Increase
RPINC	19	W	Y	Read Pointer Increase
CHRDO	1A	W	N	Channel Readout
SWTRG	1B	W	Y	Software Trigger
TRCLR	1C	W	Y	Clear Trigger Counter
ERCLR	1D	W	Y	Clear Error Flags

Table 5.3: ALTRO instruction set.

5.5.3 SIU Interface Module - Interfacing the DDL

The main use of the DDL link is to transfer data from the TPC front-end electronics to the ALICE DAQ. Before the start of run configuration data can be transferred from the DAQ side to the Front End Electronics.

TPC data readout

The TPC data will be transferred in blocks on the DDL link. One block corresponds to all data connected to one particular L2 accepted event. These data include event header information as described in [80] and the TPC data collected by all the ALTROs connected to the RCU in question.

Configuration of the ALTROs

The different registers and memories of the ALTROs must be configured before starting data collection. The most bandwidth demanding task in the configuration sequence is the uploading of data to the pedestal memories. For a maximum of 25 connected Front End Cards the amount of raw data for the pedestal memories is 32 MB. The SIU Interface Module can be used for configuration tasks like this. The part of this module that can be used for such tasks are very similar to the DCS Interface Module in functionality. The data formats used are the same and are described in detail in section 5.5.1 and subsections.

5.5.4 Trigger Interface Module - Interfacing the TTC

The trigger information from the TTCrx is read into a dual port memory in the FPGA by the Trigger Interface Module. The most urgent data is received via some of the TTCrx output pins, while other data is read out through the I^2C interface in the TTCrx chip. The most urgent data is L1, the event and bunch numbers and the trigger messages L2A or L2R (Level 2 Reject). All registers in the TTCrx are mirrored in the FPGA memory and the RCU master and DCS can read and write to these register mirrors. Any write operation to a mirrored TTCrx register in the FPGA is also transferred to the TTCrx via the I^2C interface. L1 and L2A/L2R is sent to the RCU Master Module on dedicated lines. L1 is just passed on from the TTCrx while L2A/L2R is generated by the trigger

interface upon reception of a complete L2A message or an L2R word. Their outputs are synchronized with the RCU clock. The L1 and L2A trigger signal is directly forwarded to the ALTROs. The receipt of these signals from the Trigger IM initiates readout of trigger information by the RCU Master Module. This module requests bus access and starts to read the trigger information from the Trigger Module. The information is written to dedicated registers in the Data Assembler Module. Later this information is added to the event header information of the event in question.

5.5.5 Front End Control Network Interface Module - Interfacing the control network of the Front End Electronics

This is the local Slow Control for monitoring status, interrupt and error handling of the FECs. It has a dedicated bus connection between the RCU and the FEC Board Controller, which uses a I^2C derived protocol [85]. The responsibilities of the local Slow Control is to:

- configure the power state of all FECs,
- monitor voltages, currents and temperature,
- handle interrupts (up to switching off the cards who asserted the line) in case of:
  - temperature or currents over thresholds
  - voltages under threshold
  - power supply errors
  - missed sclk
  - protocol error,
- read status parameters:
  - L1, L2 counters
  - sclk counter
  - dstb counter

5.6 Data Assembler Module

The Data Assembler Module is responsible for receiving the TPC data from the ALTRO Interface Module, Huffman compressing (optional) and formatting the data to 32-bit format. The TPC data is then stored in the internal memory of the RCU, which can be accessed by either the SIU/DDL and/or DCS (via the RCU local bus). The module also builds the Data Header and puts it in the output memory before the TPC data arrives from the ALTRO interface.

31 - 24	23 - 16	15 - 8	7 - 0
TPC WORD1[31-24]	TPC WORD1[23-16]	TPC WORD1[15-08]	TPC WORD1[07-00]
TPC WORD2[23-16]	TPC WORD2[15-08]	TPC WORD2[07-00]	TPC WORD1[39-32]
TPC WORD3[15-08]	TPC WORD3[07-00]	TPC WORD2[39-32]	TPC WORD2[31-24]
TPC WORD4[07-00]	TPC WORD3[39-32]	TPC WORD3[31-24]	TPC WORD3[23-16]
TPC WORD4[39-32]	TPC WORD4[31-24]	TPC WORD4[23-16]	TPC WORD4[15-08]

Table 5.4: The data format of the TPC data out of the RCU when Data Formatter is activated.

5.6.1 Data formatting

The Data Formatter formats the data from the 40-bit ALTRO output to the 32 bit SIU input. The data is delivered from the ALTRO Interface Module internal memory in a sequence of 4 data words and one clock cycle pause, as illustrated in figure 5.7. This

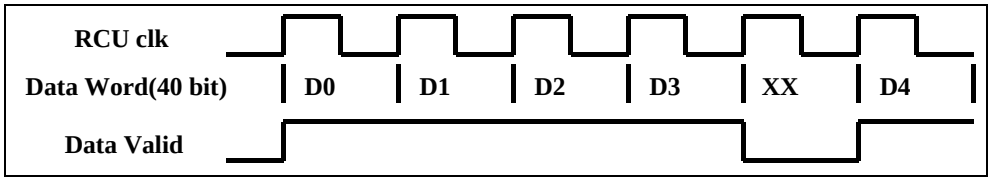

Figure 5.7: TPC data clocking from the ALTRO Interface Module to the Data Assembler Module.

input data is reformatted into 32 bit words as illustrated in table 5.4. The data formatter is only used if Huffman encoding is not active.

5.6.2 Huffman Encoding

Optionally Huffman encoding can be used to do lossless compression of the TPC data before they are sent to the DDL. The data header is not encoded. In this case the data formatting described in the previous section is not necessary. The method of Huffman encoding was first described in [86].

5.7 Configuration Module

The Configuration Module is responsible for receiving data from the DDL, interpreting them and using them to configure the other modules and the ALTROs. This module is heavily based on the DCS IM (refer section 5.5.1) and operates with the same data format.

5.8 PLD technology

A Field Programmable Gate Array (FPGA) was chosen as the main component of the RCU. This programmable logic device (PLD) has the advantage over system specific

circuits (ASICs), that they are reprogrammable both during the development phase and later during use "in field". This means that the logic can be developed in modules, and that the different parts of the logic can be tested in hardware during the development phase. Should new functionality be desired after the installation of the RCU in the ALICE pit, further development can be done and this can be programmed via a dedicated programming interface onto the RCU even during operation. In fact it is expected that the RCU will be routinely reprogrammed during operation as a means of ensuring full functionality in the ALICE radiation environment (see also section 7.2).

The programmable logic devices (PLD) used in the RCU incorporates system-on-a-programmable-chip (SOPC) integration. The architecture of these devices integrates look-up table (LUT) logic, product-term logic, and embedded memory collected in a basic logic element. This makes it possible to construct efficiently any known fundamental digital component, as explained below.

The logic elements are collected in logic blocks. These logic blocks are interconnected through a network of interconnection lines, that are organized in rows and columns as illustrated in figure 5.8. The Input/Output blocks are also connected to this interconnection network. The PLD on the RCU is SRAM (Static Random Access Memory) based, so the interconnection between a row and column line is provided by a transistor. Controlled by the state of an on-the-chip SRAM cell, the connection can be ON (SRAM cell stores a 1) or OFF (SRAM cell stores a 0).

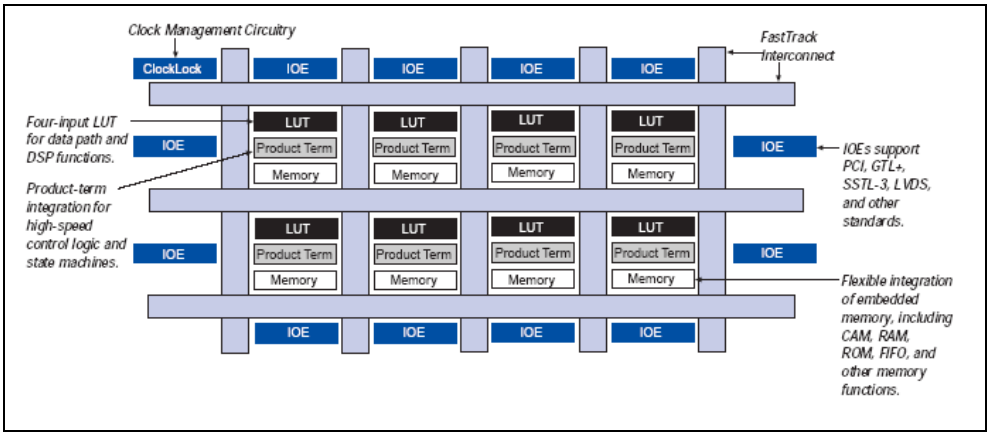

Figure 5.8: APEX 20K Device Block Diagram [87].

5.8.1 Configuration of device

The PLD used in the RCU uses SRAM configuration elements that require configuration data to be loaded each time the circuit powers up. The process of physically loading the SRAM data into the device is called configuration. During initialization, which occurs immediately after configuration, the device resets registers, enables I/O pins, and begins to operate as a logic device. The I/O pins are tri-stated during power-up, and before and during configuration.

SRAM configuration elements allow APEX devices to be reconfigured in-circuit by loading new configuration data into the device. Real-time reconfiguration is performed

by forcing the device into command mode with a device pin, loading different configuration data, reinitializing the device, and resuming user mode operation. In-field upgrades can be performed by distributing new configuration files.

5.8.2 Logic Element

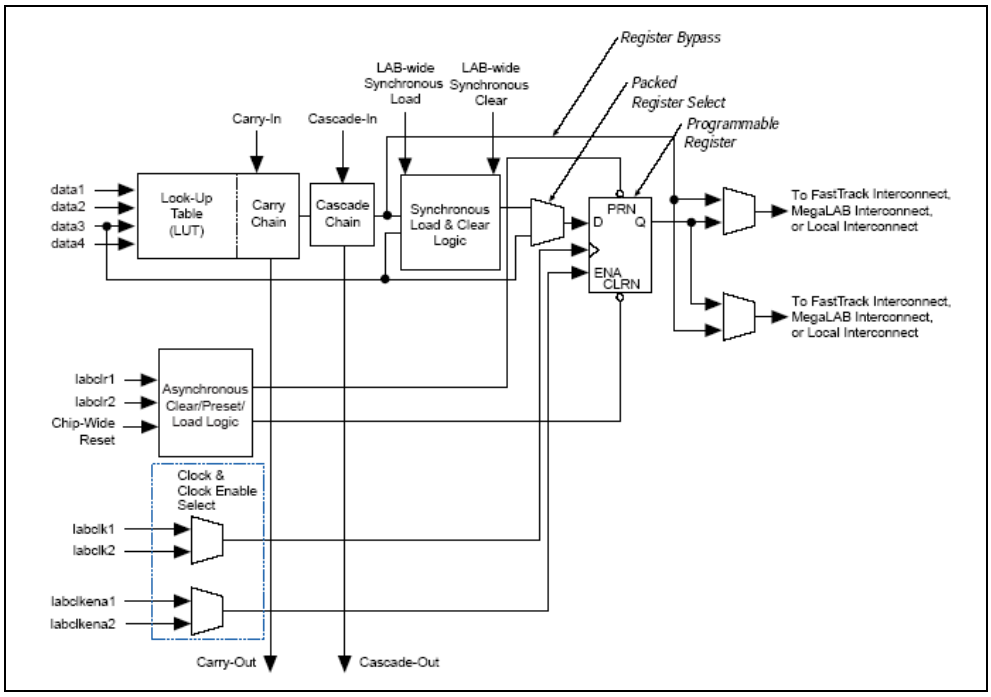

Figure 5.9: APEX 20K Logic Element [87].

The Logic Element (LE) is the smallest unit of logic in the APEX 20K architecture, see figure 5.9. It contains a four-input LUT (Look Up Table), which is a function generator that can quickly implement any function of four variables. Each LEs programmable register can be configured for D, T, JK, or SR flip-flop operation. For combinatorial functions, the register is bypassed and the output of the LUT drives the outputs of the LE. Each LE has two outputs that drive the local interconnect routing structure. Each output can be driven independently by the LUTs or registers output. For example, the LUT can drive one output while the register drives the other output. This feature, called register packing, improves device utilization because the register and the LUT can be used for unrelated functions. The LE can also drive out registered and unregistered versions of the LUT output. The PLD architecture provides two types of dedicated high-speed data paths that connect adjacent LEs without using local interconnect paths: carry chains and cascade chains. A carry chain supports high-speed arithmetic functions such as counters and adders, while a cascade chain implements wide-input functions such as equality comparators with minimum delay.

5.8.3 Embedded System Block

The ESB can implement various types of memory blocks, including dual-port RAM, ROM, FIFO, and CAM blocks. The ESB includes input and output registers; the input registers synchronize writes, and the output registers can pipeline designs to improve system performance. The ESB offers a dual-port mode, which supports simultaneous reads and writes at two different clock frequencies. The product-term portion of the MultiCore architecture is implemented with the ESB.

5.9 Design methodology

Effort is made to keep the RCU logic architecture independent throughout the design and simulation phase. This is done through the use of standard hardware description languages, like VHDL (see section 5.9.1 below). In order to optimize the performance of for example the FIFOs and SRAM memories, LPM (see section 5.9.2) was embedded in the VHDL code.

The VHDL code is verified with the use of testbenches in a simulator (ModelSim) before the code is synthesized to a netlist. This netlist can then be used in the tools of hardware vendors to place and route the logic into the specific architecture of the hardware device in question (for example the ALTERA APEX FPGA by the use of ALTERAs QUARTUS developing tool).

Only then the logic is programmed in the FPGA and tested in hardware. Lately ALTERA has included a tool making it possible to tap the values of the signals taking the role as an embedded logic analyzer, making on-chip debugging of the developed logic possible.

5.9.1 VHDL

There are two major programming languages available today for specification, verification and construction of electronic designs. These are Verilog (VERIfying LOGic) [89] and VHDL (Very high speed integrated circuit Hardware Description Language) [90]. In the RCU project both languages are in use, but the modules developed at University of Bergen and Bergen University College has been developed in VHDL.

VHDL makes it possible to describe an electronics design in a high level programming language. The VHDL specification includes for example the use of IF and CASE statements similar to other high level languages. A hardware description language like VHDL is, however, inherently parallel. Commands, which correspond to logic gates, are executed in parallel, as soon as a new input arrives. A VHDL program mimics the behavior of a physical, usually digital, system. It also allows incorporation of timing specifications (gate delays) as well as to describe a system as an interconnection of different components.

5.9.2 Hierarchical design

VHDL supports hierarchical designs. In the RCU project the design is done top-down, taking the overall requirements and splitting the design into modules, as shown in figure 5.1 and described above. The designer can then concentrate on one of these modules,

and also split the development of this module into further sub-modules that are designed and tested separately.

Library of Parameterized Modules

The Library of Parameterized Modules (LPM) [88] is a standard set of functions supported by most major EDA (Electronic Design Automation) and IC (Integrated Circuit) vendors. This library assists in keeping the design architecture independent throughout the design and simulation phase. LPM functions describes for example the logical operation of memories (SRAM, FIFO etc.). LPM functions used in a design can be directly passed to the IC vendors design implementation software through an EDIF [88] (Electronic Design Interchange Format) netlist file. LPM's have been used in the development of the RCU logics. For example the internal memory of the Data Assembler Module is built around an SRAM LPM.

5.9.3 Simulating the design with testbenches

Testbenches are also written in VHDL, and contains stimuli and/or processes the output of the design under test. This is the standard way to test the behavior of the module itself and the interfaces to the other modules before it is being synthesized and configured in hardware. Testbenches are written at all levels in the hierarchical design to verify the design functionality. The interfacing of logic to other firmware, like for example the SIU firmware in the RCU development, can also be simulated in this way. During the development of the RCU interface to the SIU, the VHDL code of the SIU interface and the testbenches used to develop this was made available by the DDL development team. This made it possible to fully simulate the communication over the DDL link before putting the RCU logic into hardware, both sending data from the RCU and receiving configuration data from the DDL. This drastically reduced the development time of the RCU SIU interface module, and also contributed to verification of the DDL firmware and documentation.

If no such firmware is available from the interfaced logic, testbenches must be written from the information extracted from the documentation of the interfaced logic. This is a procedure that is less recommended and often generates errors in the design that are due to interpretation of the documentation itself.

5.9.4 Synthesis tool

After the simulation of the design the next step is synthesis of the VHDL code into hardware. This is done with the help of synthesis tools from the manufacturer of the FPGA in question:

- Timing requirements are specified by the designer
- Input/Output pin assignment are specified by the designer
- Register Transfer Level (RTL) synthesis are done by the synthesis tool
- Place and Route are done by the synthesis tool, but can be modified by the designer

5.9.5 Testing in hardware

The last step of the implementation cycle is the testing of the developed logic in hardware. The synthesis tools have the option of embedded logic analyzers. This makes the debugging of logic in hardware easier, since the different signals and even states of the state machines can be displayed running at speed under real-world system conditions.

Hardware tests of the different RCU prototypes are described in the next chapter, where the bottom-up approach is emphasized. This means starting with the simplest possible hardware configuration (for example sending of data from the RCU to the DDL and comparing the result with the sent data), and building up to the full ALICE TPC Front-End Electronics chain.

Chapter 6

RCU prototypes

The Readout Control Unit and its control logic has developed through three prototype versions. The development of these three RCU prototypes is described in this chapter along with results achieved. The changes and upgrades made from one prototype to the next are discussed.

The development of the RCU has been done in parallel with the development of the other Front End Electronics of the ALICE TPC and also electronics for other ALICE detectors, like the ALICE TRD. The first prototype was done using a commercial PCI card as a basis. This was done in order to speed up the development process since the RCU was needed for testing the ALTRO and the Front End Card prototypes, and reading out data from the prototype ALICE TPC. For the second prototype, a card also developed for the prototyping of the HLT RORC was used. Collaboration with the ALICE Transition Radiation Detector Electronics group later made it clear that the development of the RCU and the control unit of the TRD Front End Electronics had many identical interfaces. This lead to the splitting of the RCU into a main board with the DCS board (and DDL Link card) connected. The DCS board, developed in Heidelberg in cooperation with the TRD group, has an ALTERA FPGA with an embedded ARM processor and the TTCrx chip mounted. Consequently, the DCS and Trigger logic modules were moved there from the main RCU.

Since the development of the RCU has gone in parallel with the use of the RCU in different test set-ups, all the prototypes has been in use for "real" data taking, and thus proven their functionality.

6.1 Prototype 1

The first prototype of the RCU is based on a commercial PCI-card, see figure 6.1. The choice was made to use such a general purpose PCI card in order to concentrate on the design of the RCU logic that goes into the FPGA. The commercial PCI-card chosen has a large size FPGA from ALTERA (APEX 20K400) mounted. A number of the Input/Output pins of the FPGA is connected to a local bus on the PCI-card. This local bus is again connected to the external memory chips that can be mounted on the card and the 4 CMC connectors that makes it possible to connect custom made cards and devices, and communicate with them. The available Input/Output pins are however limited, so it

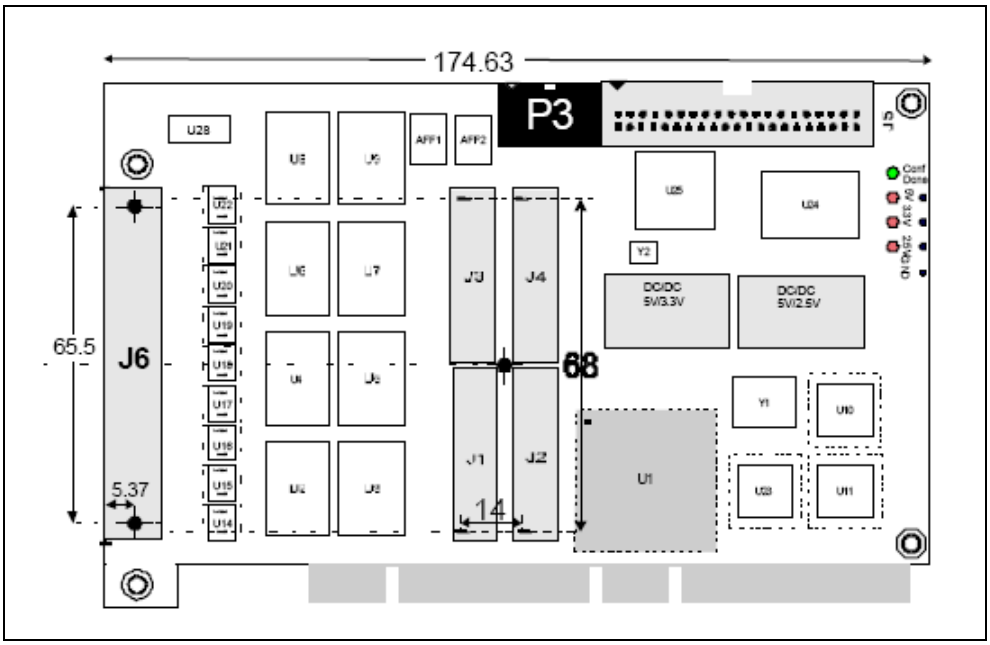

Figure 6.1: PLDA PCI20K PROD b mechanical format and layout (component side). The PCI interface is on the bottom side. The 4 CMC connectors are in the middle (J1-J4). And APEX 20K400 is placed at the U1 position. Left of the CMC connectors the external SRAM are placed. P3 is the interface for the programming of the FPGA. A number of pins of the connector to the right of P3 was also used to connect the Front End Bus via the mezzanine card [91].

was only possible to connect one branch of the Front End Bus. The design of the Front End Electronics of the ALICE TPC has been made such that two branches are needed to connect the RCU to all (maximum) 25 Front End Cards by the use of the GTL standard. This has a limitation of maximum 1 meter length of the Front End Bus cables in order to comply with signal and noise levels and data speed.

6.1.1 Custom mezzanine board

A customized mezzanine board was designed containing GTL transceivers for one branch of the Front End Bus. In addition there was L1 and L2 trigger inputs mounted on this mezzanine. The SIU connector was also mounted on the mezzanine board. The mezzanine card was connected to the main PCI board via 4 CMC connectors as shown in figure 6.2.

6.1.2 RCU logic and software for prototype 1

The RCU logic and software development for prototype 1 was split into logic blocks as shown in figure 6.3. The communication via the PCI bus was done using the commercial PCI core delivered with the PCI card and a PCI driver. A text based User Interface was developed in C where the ALTRO instructions could be transferred from formatted text files to the ALTRO Interface Instruction Memory. This UI also read out the RCU status registers and the Readout Memory of the ALTRO Interface getting the TPC data. These

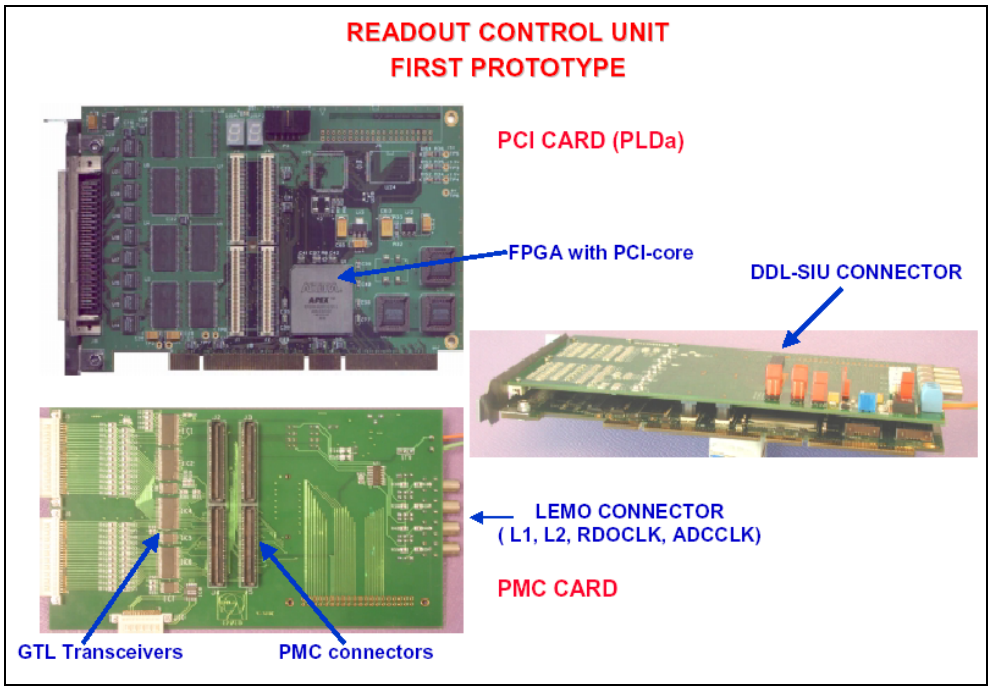

Figure 6.2: The RCU prototype 1 (top left) and the mezzanine card (bottom left) developed at CERN. Right part of the figure shows the two cards connected.

data were stored as text files.

A general interface logic to the SIU was developed in VHDL, using the protocol defined by the ALICE DAQ team. This logic block was tested separately as described in section 6.1.4.

The other major logic block in the RCU prototype 1 FPGA was the ALTRO Interface. This interface was designed to receive ALTRO commands via the PCI interface and execute them on the Front End Bus. With the help of these commands data taking can be done and the ALTROs can be monitored. This was tested as described in section 6.1.3. The ALTRO Interface contains the memories and registers as specified in table 6.1. These were accessible from the PC side via the Base Address Register mapping to the PCI core.

6.1.3 Test of the interface to the Front End Bus

The first test was to be able to communicate with the FECs and ALTROs via the FEBus. This requires the RCU to be able to send commands to the ALTROs and also to read out and store the data being delivered from the ALTROs (when asked for by the RCU). A design was chosen so that separate internal memory locations in the FPGA of the PCI-card would store these commands and data. By using the built-in PCI interface and the (firmware) PCI core these memories could be accessed from the PCI bus of the PC to which the PCI-card was connected. A mailbox register in the PCI core was used to start the execution of the stored commands and put them on the FEBus. By polling this register from the PC via a PCI driver one could also find out when the data memories where filled and ready to be read out over the PCI bus.

This setup was used to test the communication on the FEBus and to the ALTROs

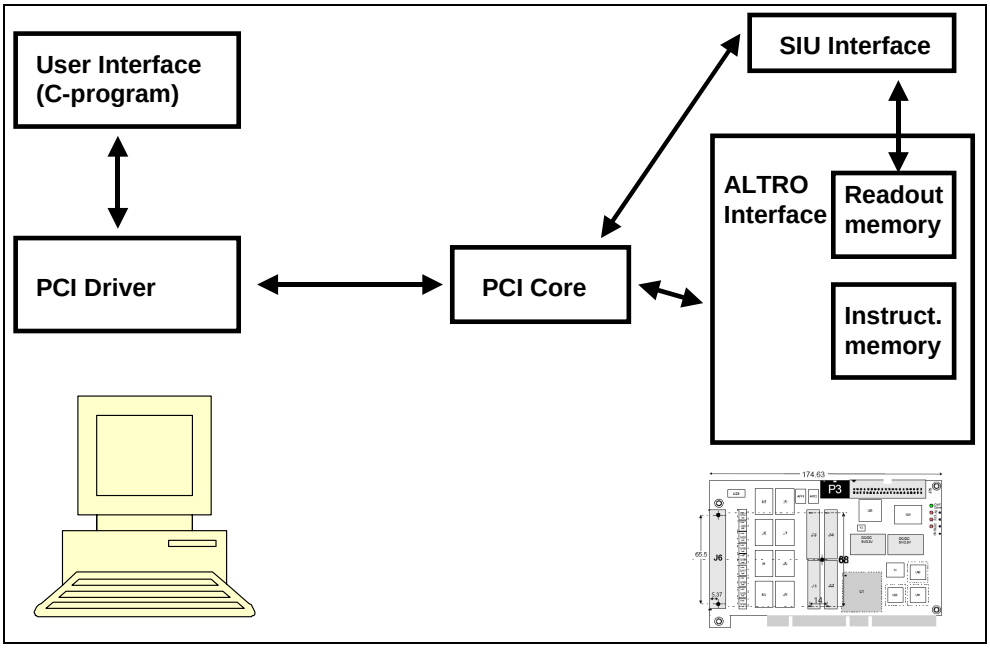

Figure 6.3: Logic block diagram for RCU Prototype 1.

Register name	BAR (mapping)	Offset	Access Access	Size (32 bit words)	Word width (bits used)
Status and config.	BAR 0	0	R/W	12	32
Test program mem.	BAR 1	0	W	512	24
Pedestal ref. mem.	BAR 2	0	W	342	30
Pedestal out mem.	BAR 2	512	R	342	30
Results mem.	BAR 3	0	R	128	24
Readout mem.	BAR 4	0	R	256	32
Readout mem.	BAR 4	256	R	64	32

Table 6.1: RCU mapped memories and registers.

and FECs connected. It was shown that up to 13 Front End Cards could be connected on the Front End Bus, thus fulfilling the requirements. This setup was also used with the prototype TPC to read out the test data, as described in subsection 6.1.6. These data have been collected in a database, available at [92].

6.1.4 Test of the interface to the DDL

The part of the RCU logic including the RCU internal data memory and the state machine responsible for handling and responding to the DDL control lines was isolated, and tested separately. A small application was written to be able to fill the internal RCU memory with a known data pattern. This was read out via the DDL. The result after the readout was compared to the input. The required error rate of the DDL transmission (less than 10^{-15}) was achieved. The data rate (100 MB/s) was also reached, but only by continuously reading out the identical data from the RCU Data Memory, since the writing of this memory via the PCI bus was slowing this test down to about 1/3 of the speed. This was due to the handshake and the fact that single write cycle was used over the PCI bus (not Direct Memory Access).

6.1.5 First FEC-DDL integration test

The next test was to combine all the components of the Front End chain and the RCU logic of the two tests described above for a full chain readout test. This was done using the ALTRO in test mode (where a data pattern can be uploaded into the pedestal memories of the ALTRO and afterwards be read back out). These data were read out using the sequence of ALTRO instructions and trigger signals that are to be used during normal data taking with the TPC. The data were stored in the RCU memory, and then sent to the DDL and collected by the pRORC on the receiver end. From the pRORC the data were read into the PC memory via the PCI bus and were then displayed and compared to the pattern stored in the ALTRO pedestal memories. The system was also connected to the DATE DAQ software. The data rate was again considerably slower than the maximum 100 MB/s of the DDL (about 1/5), due to the fact that the internal data memory of the RCU had to be filled by the ALTROs first and only later being read out. This low data rate was expected, and the important result of this test was that the full chain was cooperating for the first time.

6.1.6 The RCU prototype 1 used for reading out the ALICE TPC prototype

An ALICE TPC prototype has been developed and is installed at a CERN facility, see figure 6.4. In total, 4 front-end cards, in its final version, containing 8 ALTRO chips each, were plugged directly on the detector. Different sets of measurements were taken, in the first place, to solve electrical issues such as grounding, decoupling and spurious noise in the Front End Electronics.

In figure 6.5 can be seen a reconstructed cluster from the TPC data collected by the use of the RCU prototype 1 with the prototype ALICE TPC.

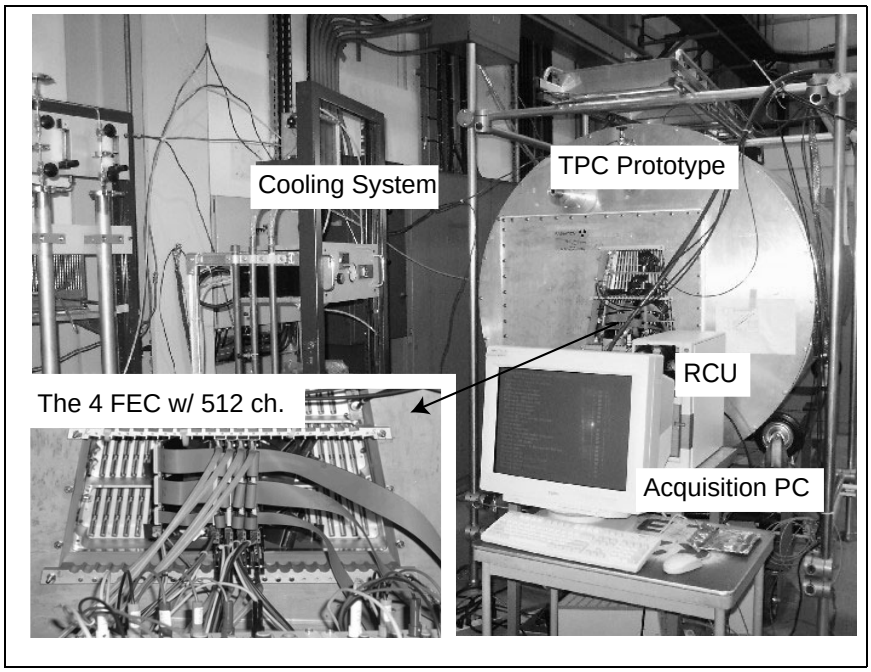

Figure 6.4: The experimental setup for testing the Front End Electronics with the ALICE TPC Prototype. 4 FECs are plugged in the TPC pad plane in adjacent slots. The metal cover around each FEC is in place and the cooling system is functioning. An RCU board, plugged on the mother board of a PC receives data from the 4 FEC.

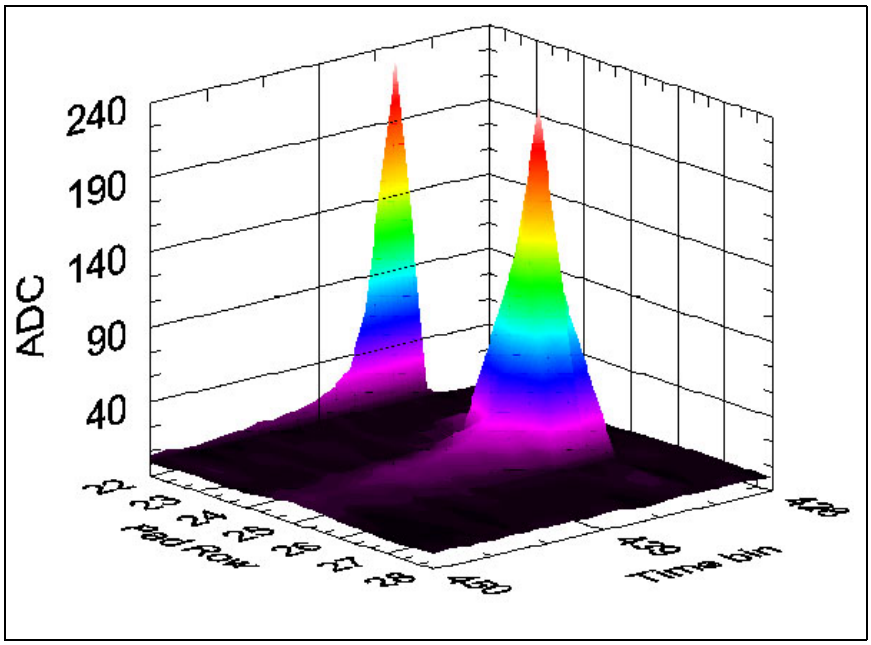

Figure 6.5: Example of single ionization cluster produced by the decay of ^{83}Kr as reconstructed from data recorded by the prototype TPC at CERN [77].

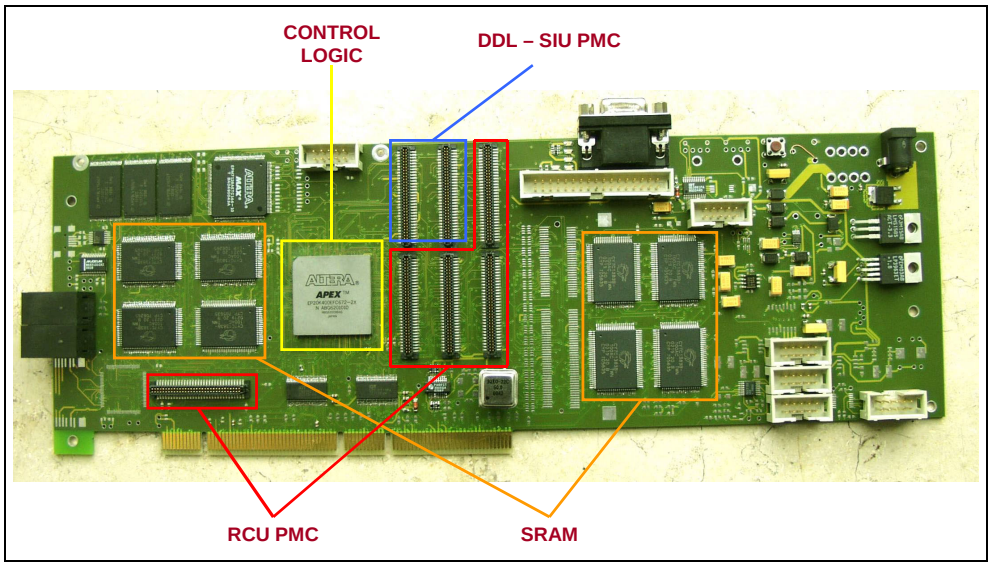

Figure 6.6: The RCU2 prototype.

6.2 Prototype 2

To test various hardware options and to make possible the connection of 2 branches of the Front End Bus, a custom prototype was built in collaboration with the Kirchhoff Institute of Physics, University of Heidelberg [93]. This was a general purpose board also used as a prototype for the HLT-RORC. A picture of the card is shown in figure 6.6. The PCI connector is seen in the bottom side of the card. The logic of the RCU is programmed in an ALTERA 20K400 (same device as used on the first prototype) FPGA shown in the middle of the card. To this card is connected two mezzanine cards, namely the SIU card (later the upgraded and renamed DDL Link Card) and a mezzanine card for connecting to the other interfaces of the RCU (Front End Bus, Trigger and DCS). A picture of the RCU 2 with the two mezzanine cards connected is shown in figure 6.7.

The RCU logic developed for Prototype 1 was ported and expanded to include handling of the two branches of the Front End Bus. The RCU prototype 2 has been used in many settings:

- first tests of the DCS communication chain (via PCI and Ethernet chip)
- in testing equipments set-ups for testing front end electronics (PASA and Front End Cards)
- for testing the feasibility of using the ALTRO/RCU (modified) readout chain for reading out the PHOS detector
- radiation tests of FPGA technology
- further test of the full TPC readout chain (including DDL integration)

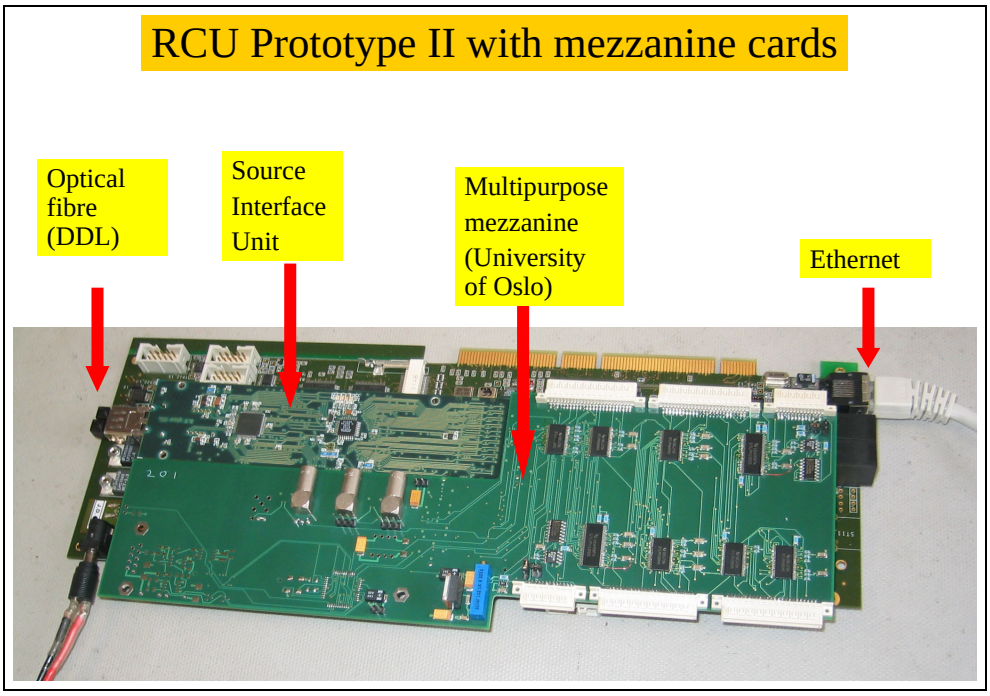

Figure 6.7: The RCU2 with the mezzanine cards mounted (Source Interface Unit for connecting the Detector Data Link and the multipurpose mezzanine for connecting the Front End Bus).

6.2.1 The multipurpose mezzanine card

A mezzanine board was developed in collaboration with the University of Oslo for use with this board. Built on the same philosophy as the mezzanine card for the first RCU prototype, this mezzanine board was carrying the connectors to the two branches of the Front End Bus, the drivers and the termination of the Front End Bus. Connectors for Trigger signals was also included. The read-out and sample clock has to be fed to the ALTROs via the Front End Bus. For the prototype 2 the read out clock could be generated on the RCU 2 prototype card either from the PCI clock or a quartz (using PLL for the derivation). The sample clock was generated from an external source, so it was fed into the RCU as well as the Front End Bus. The connectors on the mezzanine card could be set up as inputs or outputs using jumpers mounted on the card. The mezzanine could also be equipped with quartz for generation of clock signals. Another option on this mezzanine was the possibility to connect Profibus. This was for tests of possible DCS interface technology. Component placement of the mezzanine card is given in appendix A.

6.2.2 The ALTRO-GUI test software

For the testing of the developed hardware a more user-friendly interface was soon needed. Hence, a Graphical User Interface (GUI) was developed [94] using C++ and ROOT [95].

The ALTRO-GUI is an expansion of the text-based interface that connected to the RCU via the PCI-driver as described in section 6.1. It utilizes a three layer architecture, as shown in figure 6.9, which has the advantage that parts of the software can be modified

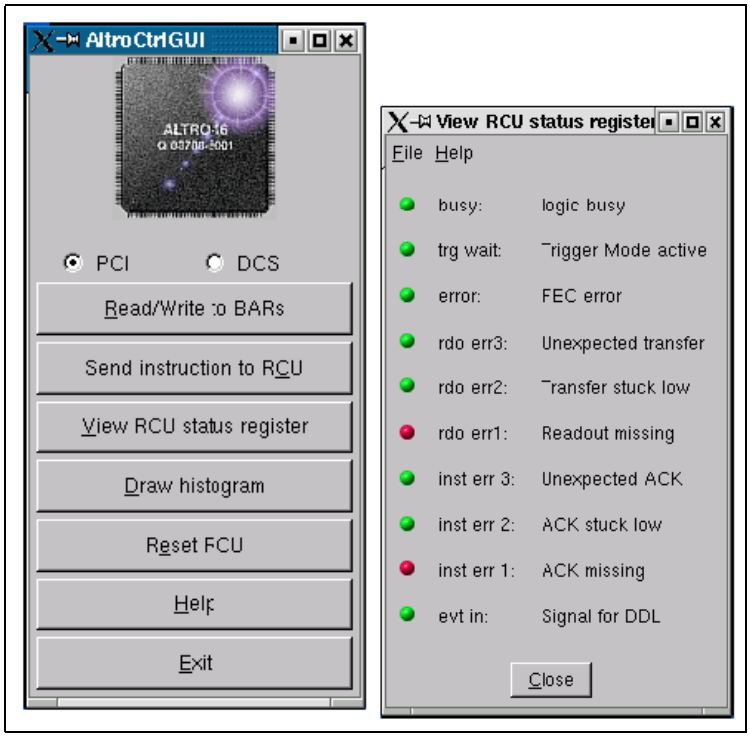

Figure 6.8: Main menu of the ALTRO-GUI test software (left). Display of RCU status register (right).

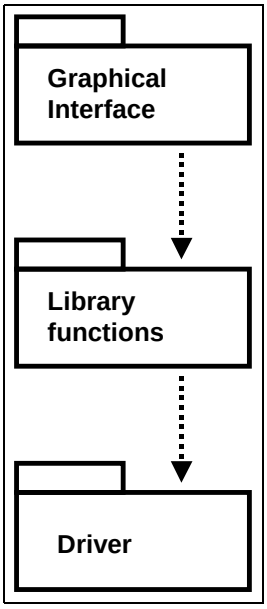

Figure 6.9: The three-layer architecture of the ALTRO-GUI test software [94].

separately. This feature was used when the ALTRO-GUI test software was integrated with the DCS DIM server (see section 4.9.4) using communication over Ethernet instead of the PCI communication.

As can be seen from the main menu screenshot in figure 6.8 the following options are available:

- **Read/Write to BARs** - The Base Addressed Registers in the RCU are listed in table 6.1. This option makes it possible to write or read the different memories and registers in the RCU.
- **Send Instruction to RCU** - RCU instructions along with the necessary parameters are composed and sent to the RCU Test Program Memory either as a single command or a list of commands. This makes it possible to configure the RCU to wait for a trigger signal and readout the TPC data from the ALTROs afterwards in one continuous sequence. The RCU and ALTRO instructions used are further discussed in section 5.5.2.
- **View RCU status register** - This option displays a graphical presentation of the status register in the RCU. The register is read out multiple times per second and each bit is displayed as a lamp next to the interpretation of that bit. If there is an error the lamp turns red, if not the lamp is green. An example of the status of the RCU is shown in figure 6.8.
- **Draw Histogram** - This option displays a histogram of data read out from the ALTROs readout memories. An example of such a histogram is shown in figure 6.10.

6.2.3 RCUs in the readout chain of the PHOS detector

The embedded electronics within each PHOS module consist of FEE (Front End Electronics) and TRU (Trigger Unit) trigger cards. Each crystal is equipped with an APD (Avalanche Photo Diode) and a low-noise, charge sensitive preamplifier. The 32 channel FEE cards are equipped with semi-gaussian shapers and 12 bit digitizers based on the ALTRO-16 ADCs developed for the TPC. Both the digitized, multi-event buffers of the FEE cards and the trigger patterns of the TRU cards are read out using the Level-1 and Level-2 timings received from the external RCU. Each RCU reads 28 FEE cards and their 2 TRU's via a common GTL+ bus which interconnects all embedded electronics cards. The above described front end chain was tested with the PHOS prototype late 2003. The test chain is shown in figure 6.11 and figure 6.12.

6.3 Prototype 3

As the third prototype of the RCU a fully customized card was developed containing only the necessary RCU components and following the physical size and placement constraints of the ALICE TPC FEE. The component placement on the card is shown in appendix B. A picture of the card with the two mezzanine boards connected and mounted on the Front End Bus, is shown in figure 6.13.

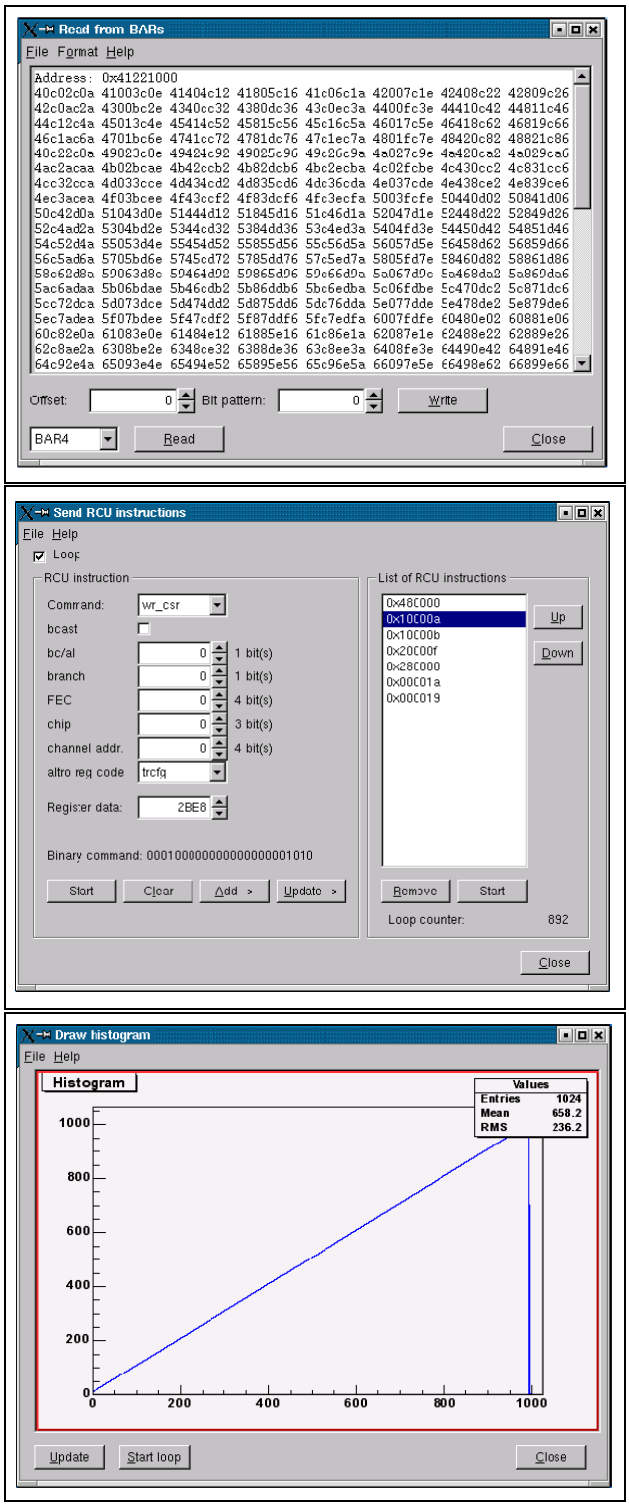

Figure 6.10: Screen shots of window for reading and writing to memory locations on the RCU (top) and window for sending instructions to the RCU (middle). These instructions are containing ALTRO instructions and are interpreted by the RCU and sent as ALTRO command on the FE-Bus. The bottom part of this figure shows a screen shot of the histogram window in the ALTRO-GUI test software, displaying data originating from the ALTRO Readout Memory. In this case the ALTRO was operated in test mode where the data which makes up the ramp was preloaded into the ALTRO from the RCU.

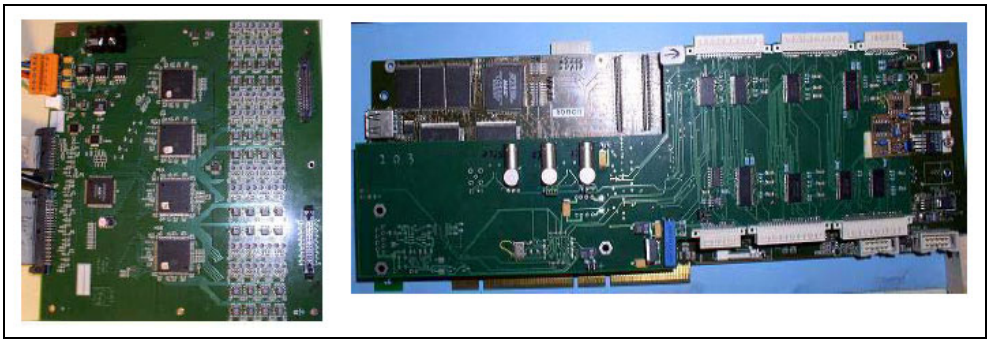

Figure 6.11: Closeup of Front End Electronics board with 4 ALTROs mounted (left), and the RCU2 prototype with mezzanine mounted (right).

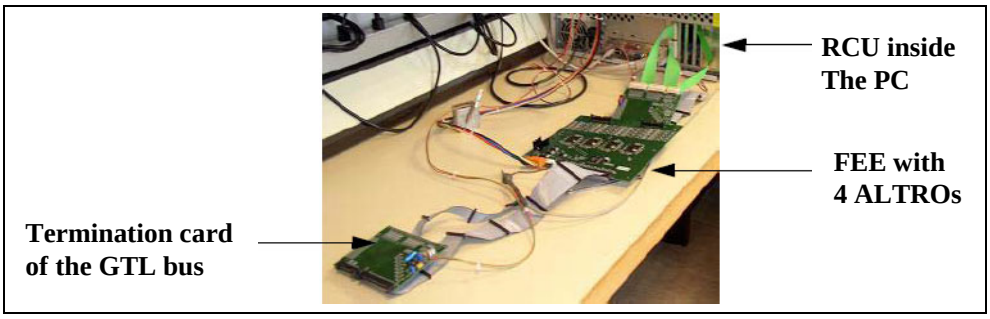

Figure 6.12: Front End Electronics test chain for the PHOS readout tests.

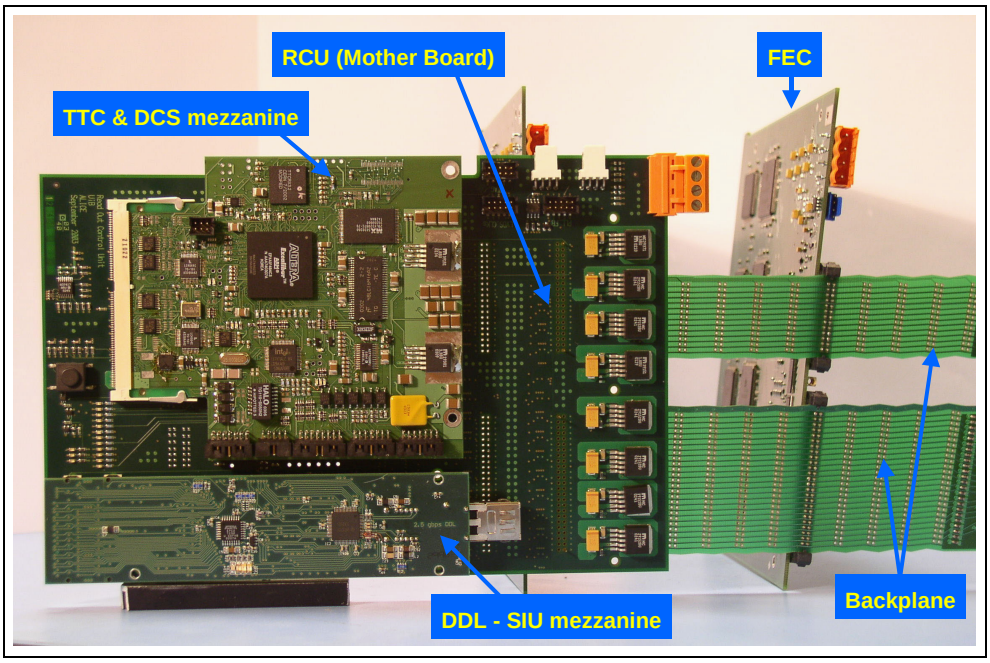

Figure 6.13: The RCU3 with the DCS mezzanine board and DDL Link Card connected and mounted on the Front End Bus (only two FECs mounted).

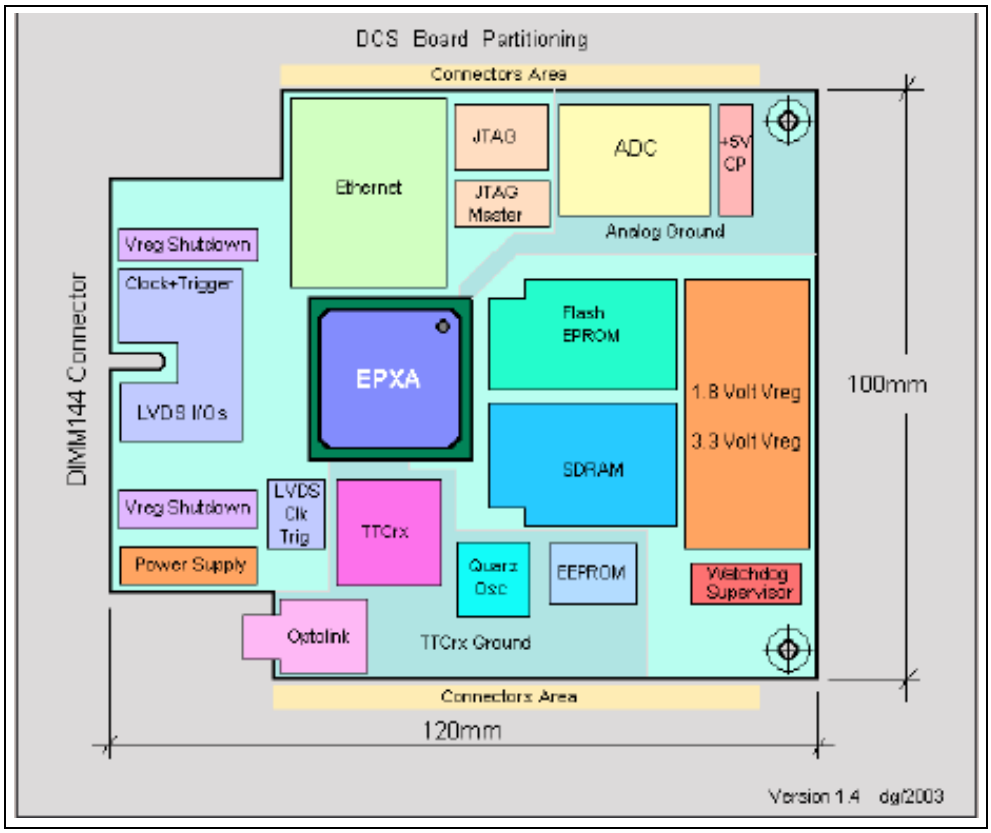

Figure 6.14: The DCS mezzanine board [96].

6.3.1 The two mezzanine cards

2 mezzanine boards are connected to the RCU3 prototype. One is the DDL Link Card, connected to the RCU board via 2 CMC connectors, as it was on the previous two prototypes. The other is a DCS-card (see figure 6.14). This is a new card developed in Heidelberg in cooperation with the TRD group. This card contains an FPGA (Excalibur EPXA1) with an embedded ARM922T hardwired processor, Ethernet (for DCS) and the TTCrx chip. It is connected to the RCU main board via a DIMM144 connector. Some of the components on this board are not used by the RCU, but used when the board is part of the Front End Electronics of the TRD.

6.3.2 RCU logic

The many updates to the logic of the RCU3 from the previous version includes:

- Integration of the RCU local bus (section 5.3) interfacing all memory mapped modules to it,
- The Data Assembler module (section 5.6) which connects the ALTRO Interface and SIU Interface, speeding up data transfer significantly compared to previous versions,
- inclusion of a Command Interpreter and Bus Switch making possible configuration and monitoring of the RCU and FECs from the DCS (via the DCS mezzanine board)

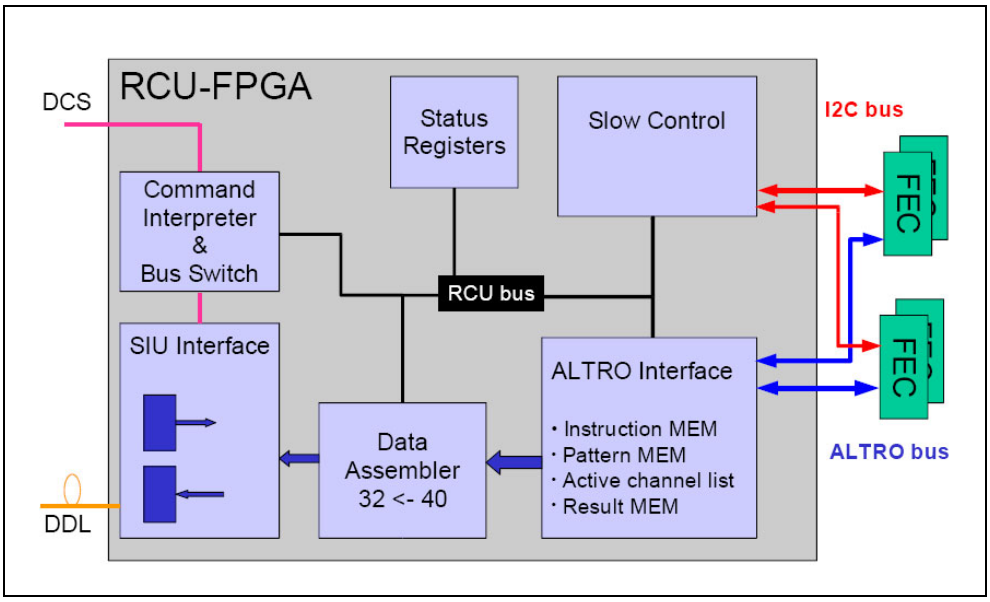

Figure 6.15: Overview of the RCU3 architecture [97].

and giving access to the RCU local bus for the master modules (like from the DDL for configuration purposes),

- the Trigger Module (section 5.5.4) has been moved to the DCS board FPGA and the Trigger signals (L1 and L2a) are sent on dedicated lines on the DIMM144 connector,
- the Slow Control module (section 5.5.5) has been included.

An overview of the RCU3 architecture is shown in figure 6.15.

6.3.3 DCS card architecture

An overview of the firmware of the DCS card is shown in figure 6.16. The functionality of the different modules are [97]:

- TTCrx - Trigger Module:
  - interface to the TTCrx chip
  - maps all the TTCrx registers in a register file in the PLD
  - decodes the L2a/r message. In case of a L2a message the RCU local bus is requested and after access the event header is transferred to the Data Assembler module.
- FPGA-Config - configures the RCU card FPGA
- RCU-Access - message system for configuration and monitoring of the RCU. By the use of 6 different kind of messages as described in section 5.5.1 the registers and memories of the RCU can be read out and/or written to. Each message sent is acknowledged and the answer contains a status report. Since the DCS card and the

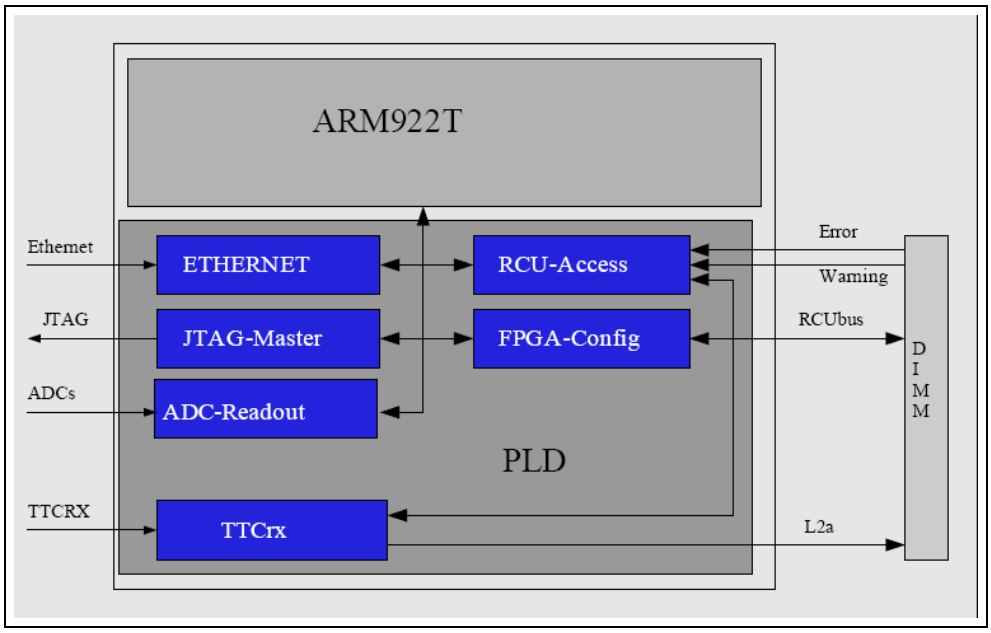

Figure 6.16: Overview of the architecture of the DCS board connected to the RCU [97].

RCU are running on different clock domains a synchronizer has been implemented in the RCU local bus protocol to guarantee data integrity.

- ADC readout - monitors the voltages on the DCS card
- JTAG master - configures the chained DCS cards
- Ethernet - provides access to the Ethernet

6.3.4 The TPC beam test May/June 2004

The RCU third prototype was used in the ALICE TPC beam test in May/June 2004 at CERN. All the major parts of the front end electronics of the TPC, including the Trigger, DDL, FECs and RCU were tested. The setup of the readout chain is shown in figure 6.17 and a picture of the two RCU cards connected to the TPC prototype is shown in figure 6.18. Two RCUs with DCS boards and DDL Link cards was used for reading out 5504 channels from 43 FECs. DATE system was used for data acquisition. D-RORC was used to ship the data to/from the HLT. 533 GB was stored on CASTOR [98] coming from $\sim$ 300k events. The configuration of the RCU and FECs was done over DCS and DDL. Data was readout over the DDL and sent to DAQ/HLT. Slow Control was readout over the DCS. The system was found to be fully operational [99].

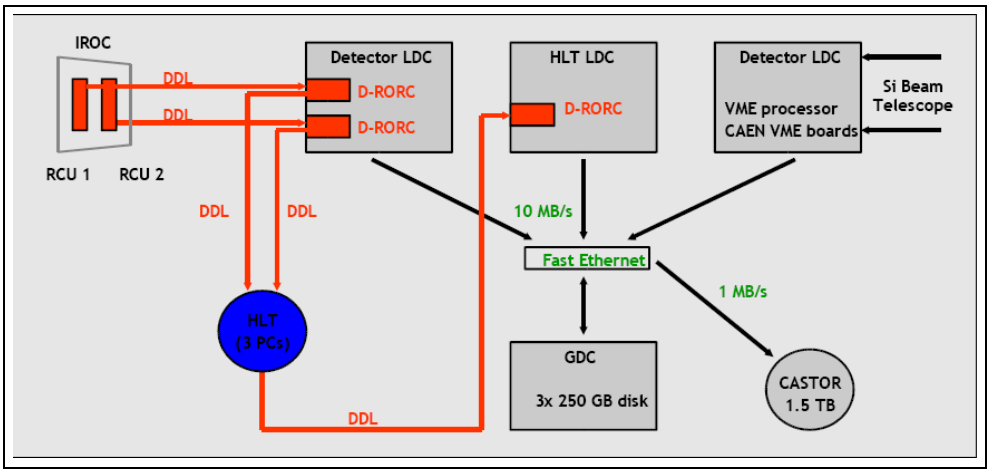

Figure 6.17: The setup of the DAQ during the beam test.

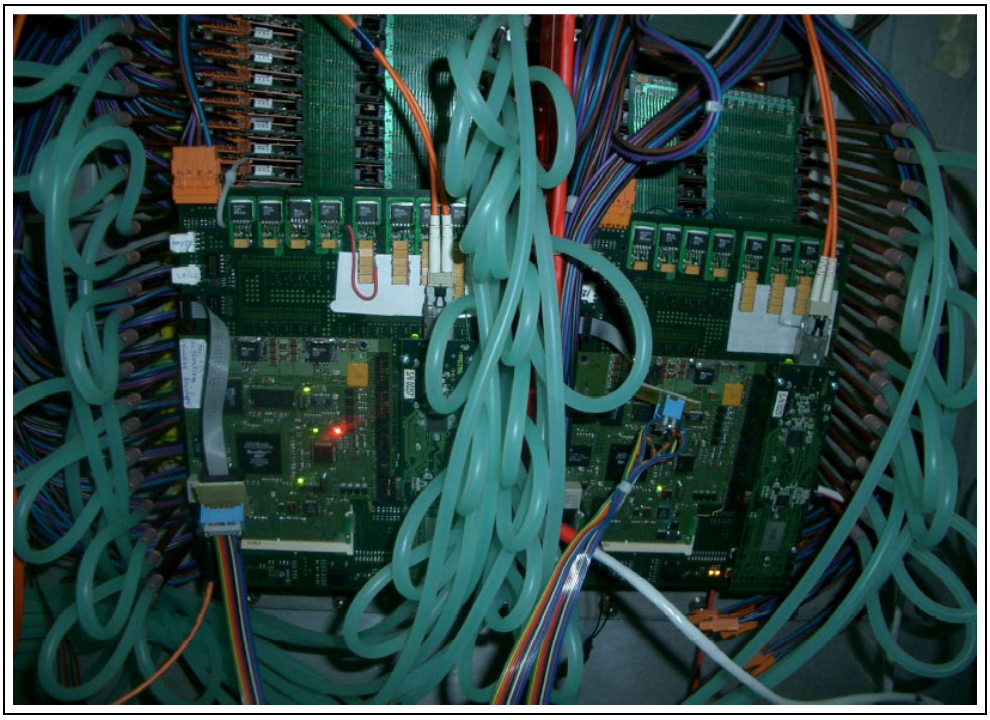

Figure 6.18: The two RCUs with DCS boards and DDL Link cards connected on the prototype TPC during the TPC beam test.

Chapter 7

Summary and Outlook

7.1 Summary

Colliding heavy ions is of interest for studying the properties of nuclear matter. ALICE is a dedicated experiment in the future LHC-accelerator at CERN, designed to study such reactions. A Time Projection Chamber is the main tracking device of the central barrel. This detector is able to reconstruct the large number of tracks being produced in heavy ion collisions. In order to cope with the expected high multiplicity and interaction rate, a new read-out chain of front-end electronics had to be developed.

The Readout Control Unit provides the local on-detector intelligence for controlling and reading out the detector. The RCU interfaces to the Front End Cards and to the DCS, trigger and DAQ systems. The conceptual design of the RCU was developed and partially implemented into three generations of prototypes. These three prototypes, based on a commercial PCI-FPGA-board, a custom PCI board and a custom stand-alone board, have been developed in collaboration with groups at CERN and the Universities in Oslo and Heidelberg. The implementation of the RCU logic and the testing and debugging of the first two systems was part of this work. The systems were later used in the characterization of the front-end boards and in test beams.

One major aspect of the development has been the design of the control logic in VHDL (and Verilog), thus being independent of any hardware providers and being able to migrate to alternative technologies (FPGAs) if necessary or desired.

A second constraint was the splitting of the hardware into a RCU motherboard and a DCS daughter board. This was done to allow a common development of the DCS subsystem for both the TPC and the TRD (Transition Radiation Detector). Since the DCS board has an FPGA mounted with an embedded ARM922T hardwired processor, the logic of the DCS and trigger modules have been moved there.

The logic of the RCU has gone through a number of changes from the first conceptual design. The complex RCU master has been simplified to a number of RCU status registers that can be read out by the DCS. The originally envisioned decision taking by the RCU master has been moved out to the different modules. The interrupt handling and power state configuration of the Front End Cards are now being taken care of by the Slow Control module. The bus arbiter has been removed. A Command Interpreter and Bus Switch instead takes care of the interpretation of the commands coming from the DCS,

and gives access to the DDL when needed, for downloading of configuration data to the ALTRO interface Pattern Memory and other registers. The Slow Control module updates the Active channel list of the ALTRO interface directly so that channels not working on the FECs will not be read out. Trigger information that goes into the data header in the Data Assembler module are written directly from the Trigger Module on the DCS board (not going via the RCU Master module as originally planned).

Further work is needed regarding the protection of the RCU against radiation in the ALICE experiment area, see section 7.2. The work is still ongoing, and the final RCU will be produced in 2005.

7.2 Outlook

The stability of operation taken into account that the RCU is not physically available for maintenance when mounted in the ALICE experiment, and that it will be exposed to a radiation environment, is of major concern in the project.

7.2.1 Radiation tolerant design

Different ways of ensuring the proper functioning of the RCU logic exists. The main parts of this logic are:

- configuration registers for the different modules
- a number of state machines in each module
- memories for intermediate storage of data

The configuration registers could be multiplied, and a voting scheme between for example 3 registers could be applied. Should one of them contain different information than the other 2 this could be discarded. Hamming encoding could also be used, allowing error detection and even correction. To protect critical registers from unintendedly being written to, one could make the demand that such registers must be written to twice with the identical value before the new value is accepted.

Hamming coding of the state machines are done in the ALTRO, and could also be done in the RCU. By assigning each state an identification value, and changing this value by one bit only in adjacent states, Hamming encoding would detect an erroneous jump between states. Parallel identical state machines has also been discussed, where errors would be detected by comparing the current state of the state machines.

The intermediate data stored are either TPC data being read out or configuration data going to the ALTROs (like pedestal values). These could be protected by checksum algorithms, like the CRC16 (Cyclic Redundancy Check) used in the DCS.

The design of the RCU is heavily based on commercial products. Due to the radiation level in the ALICE pit, each of these components must be tested and qualified in simulated working environment, to ensure operating reliability. The major device on the RCU board is the SRAM-based FPGA. There are concerns that such devices might not work in the LHC radiation environment, due to upsets in Configuration RAM and in registers and memory cells.

7.2.2 Radiation environment

The ALICE radiation environment has been simulated, and radiation maps have been created (by the use of FLUKA [100] - a particle physics Monte Carlo simulation package) [102]. As can be seen from figure 7.1 a number of energetic particles will be present in the area where the Front End Electronics of the ALICE TPC will be installed in the ALICE pit. The radiation effects caused by these energetic particles can be both cumulative effects and effects related to single individual interactions in the silicon.

7.2.3 Single Event Upsets

As explained in section 5.8, the FPGAs currently used in the RCU is SRAM based. This makes it vital to test how the SRAM is behaving in the ALICE radiation environment. In particular how resistant the device is for Single Event Upsets (SEU). In figure 7.2 is shown a high energetic proton interacting with a Silicon nucleus and creating secondary particles and a short range heavy recoil ion. This is capable of ionizing the material. The created electrons will easily drift if an electric field is present and be collected at some nearby sensitive node (like the drain of a CMOS transistor). This collected charge can cause the transistor to change its logic state, if it exceeds the critical charge of the transistor. This is called a single event upset (SEU). A SEU is a "soft" error, non-destructive, and a reprogramming of the device will return the FPGA to normal behavior.

In an FPGA SRAM cells are used for both configuration memory and application memory. If the SEU happens in application memory used as temporary storage of TPC data it is not a major concern. This can be dealt with through the use of error detecting and correcting techniques like the ones described in section 7.2.1. However, if the bit flip occurs in the configuration part of the memory, the upset has to be detected immediately and the device must be reprogrammed to correct the error.

7.2.4 Tests at the Oslo Cyclotron Laboratory and the Svedberg laboratory in Uppsala

Tests of the radiation tolerance of the FPGAs and other devices on the RCU is ongoing at the time of writing of this thesis. These tests are done with proton beams at the Oslo Cyclotron Laboratory (29 MeV) and the Svedberg laboratory (38 and 180 MeV) [101].

The tests are done on the same model of FPGA as is planned used on the final RCU board. The challenge to the testers is to separate between errors in the configuration SRAM and in the programmed logic itself (including application memory). The device is typically filled with shift registers that are read out sequentially in these tests. The logic necessary on the FPGA to execute this read out could also experience upsets and then cause errors in the data read out. This logic is therefore minimized in complexity and gate size.

Upsets and even loss of functionality in the FPGA are observed to the degree that the configuration has to be reloaded. The measured configuration SEFI[1] cross-section for the APEX 20K is $6.0 * 10^{-9} \pm 1.1 * 10^{-9}$. Using the flux estimations for different radial

[1]Single Event Functional Interrupt

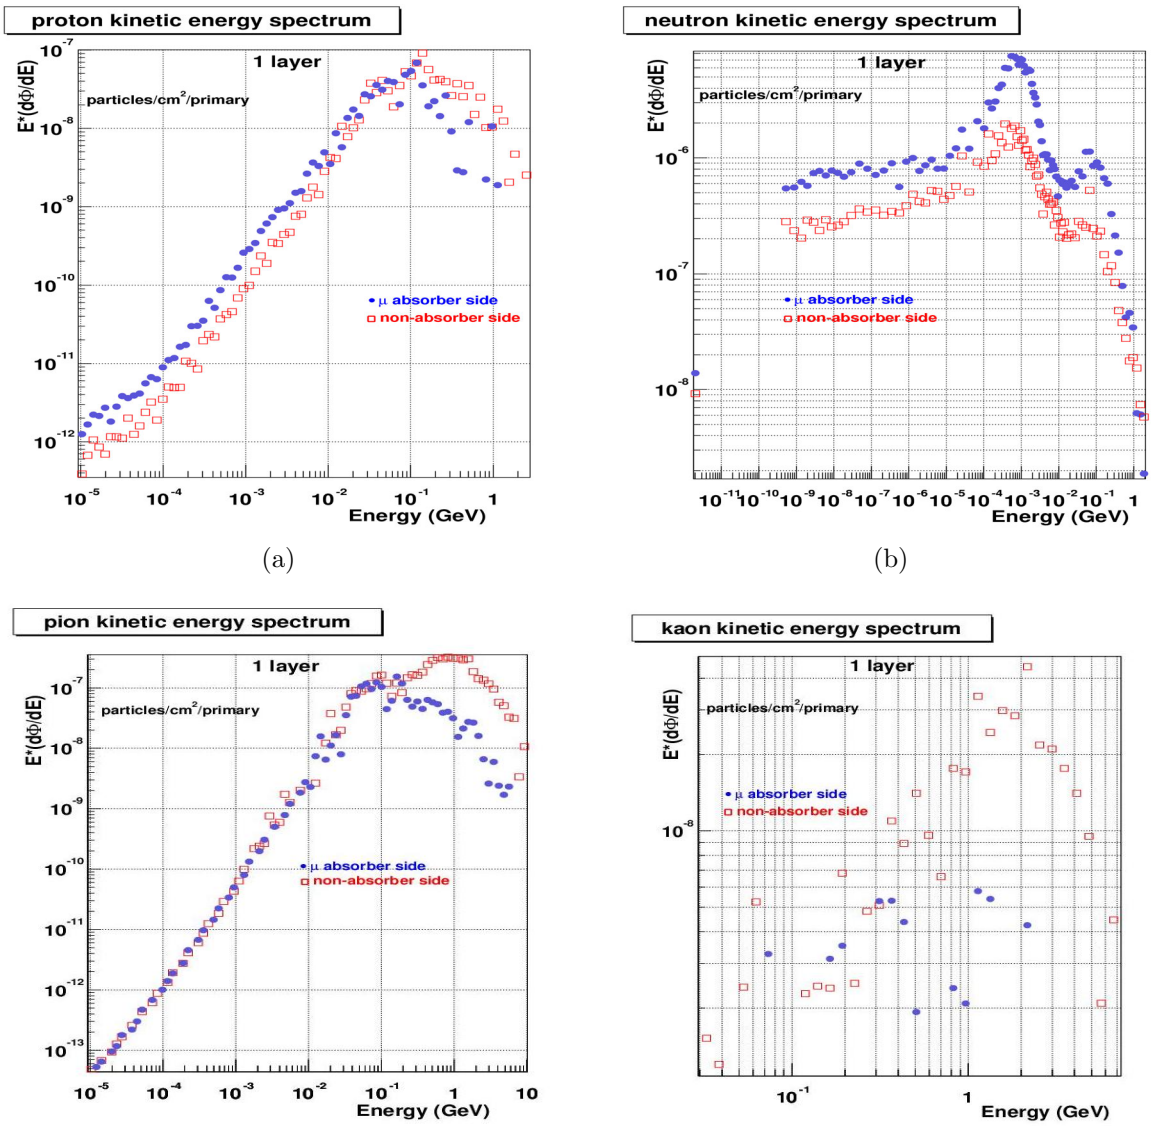

Figure 7.1: Expected kinetic energy spectrum of neutrons, protons, pions and kaons at the position of the ALICE TPC front end electronics [102].

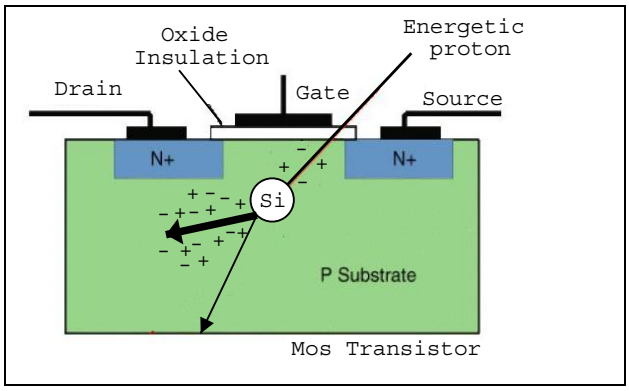

Figure 7.2: Example of ionization of a transistor [101].

positions in the TPC (from the FLUKA simulations [102]) and the knowledge of where each (of the 216) RCU sits in the TPC, the total expected number of SEUs in RCUs per 4 hour run can be estimated. The total expected number of SEUs in RCUs per 4 hour run was estimated to 3.7 ± 0.2 [101].

The result obtained from the irradiation tests so far shows that the expected SEU rate in the SRAM-based FPGAs are at the limit of what can be tolerated. This makes radiation tolerant designs necessary, which is possible, but complicates the firmware development. There are some hardware alternatives available. The FLASH-based FPGAs from Actel have been tested for radiation tolerance. Using these devices would mean a re-design of parts of the front-end electronics. Another alternative is the Xilinx Virtex FPGA which offer the feature of direct reading back the configuration memory. It also supports online partial reconfiguration. This means that if an upset is detected in a part of the configuration memory, this part can be reconfigured separately without interrupting the normal operation. This concept is currently being tested and will be further pursued.

Appendix A

The RCU prototype 2 multipurpose mezzanine card

Component placement of the RCU 2 multipurpose mezzanine card. The top and bottom side of the mezzanine card is shown in figures A.1 and A.2. The connectors to the RCU main card are named PN3, PN6 and PN2A. The connectors to the Front End Bus branches are named JA1, JA2, JA3 and JB1, JB2, JB3. The transceivers are named GTLPxxxx(x). I/O for trigger signals are marked L1 and L2. Below the trigger connectors on the lower left side of the top side figure are placed the S Clk (sample clock) and R clock (Read Out Clock).

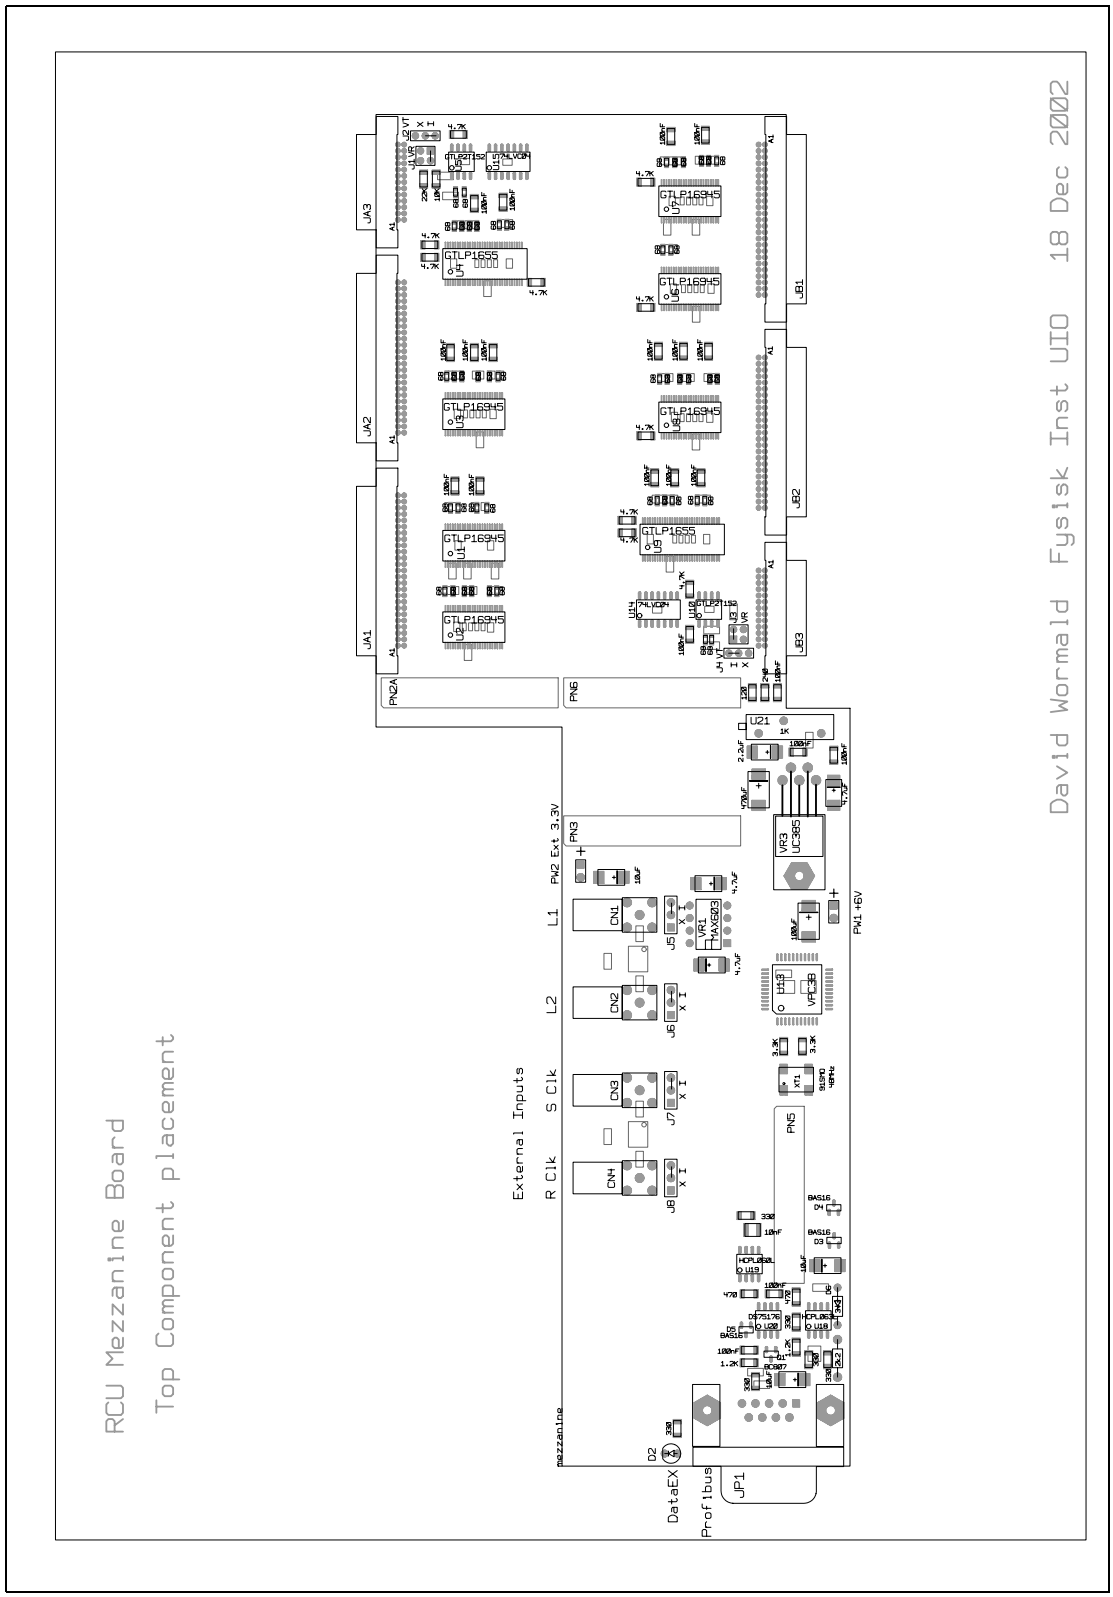

Figure A.1: The RCU2 multipurpose mezzanine card, top side.

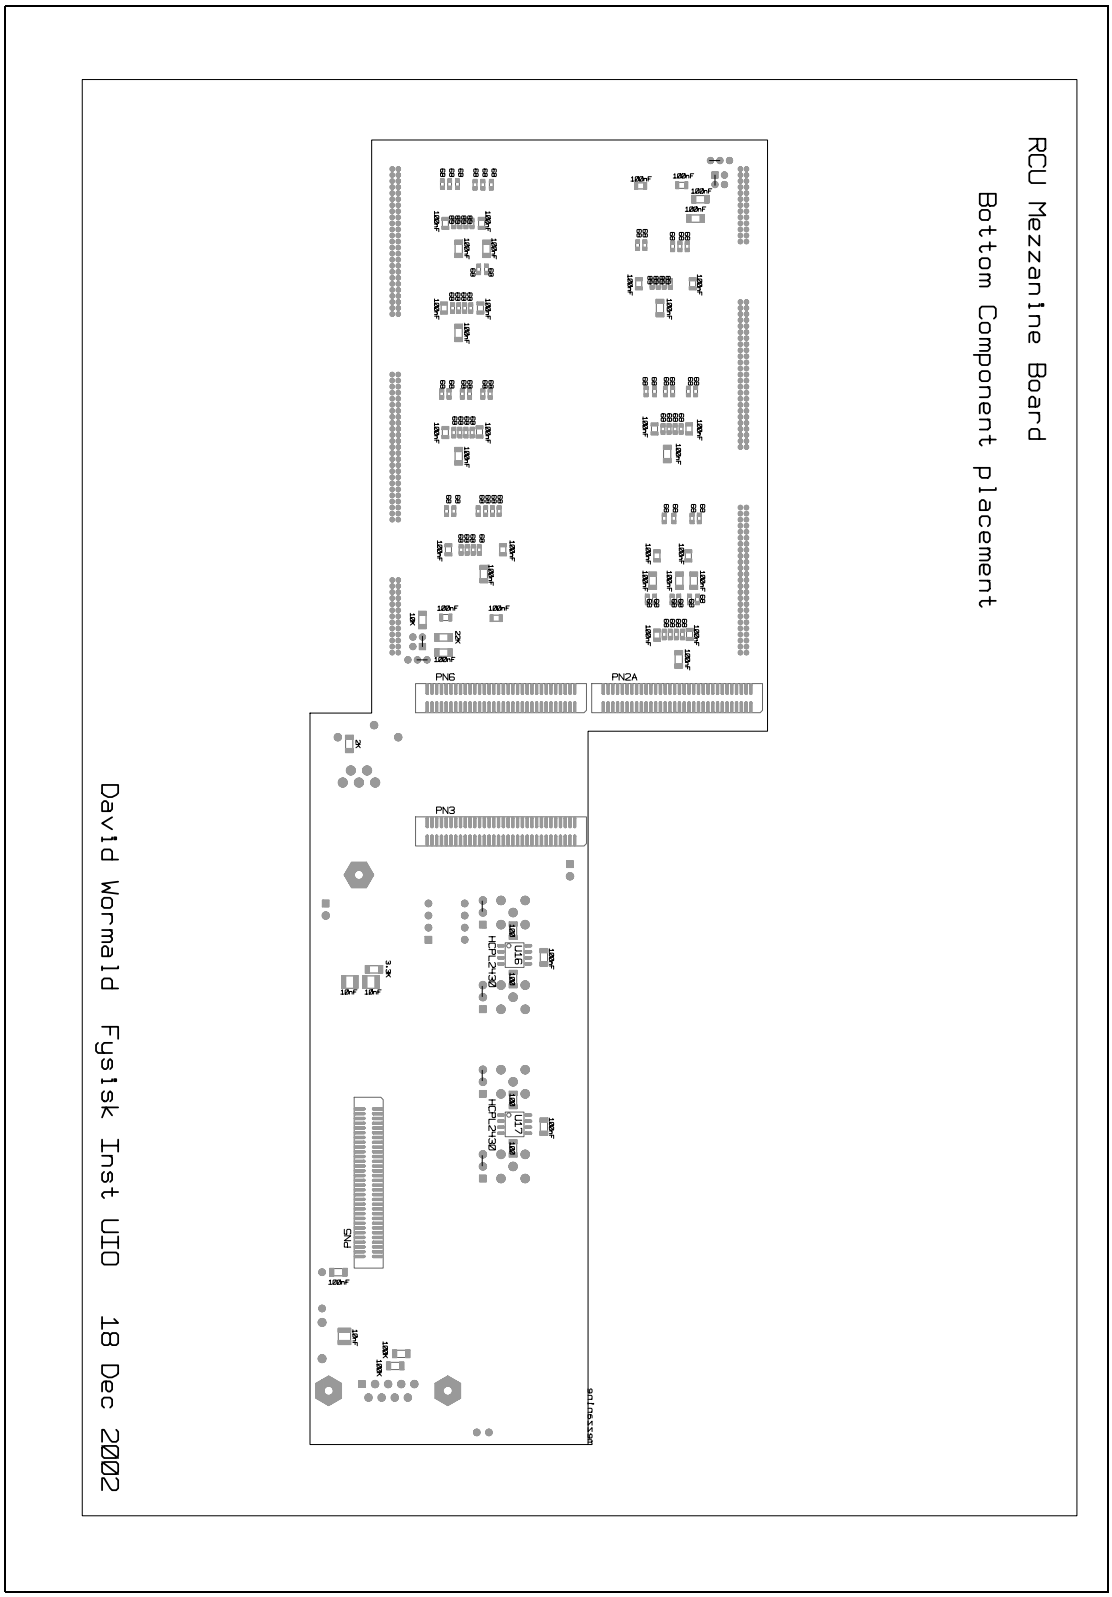

Figure A.2: The RCU2 multipurpose mezzanine card, bottom side.

Appendix B

The RCU prototype 3

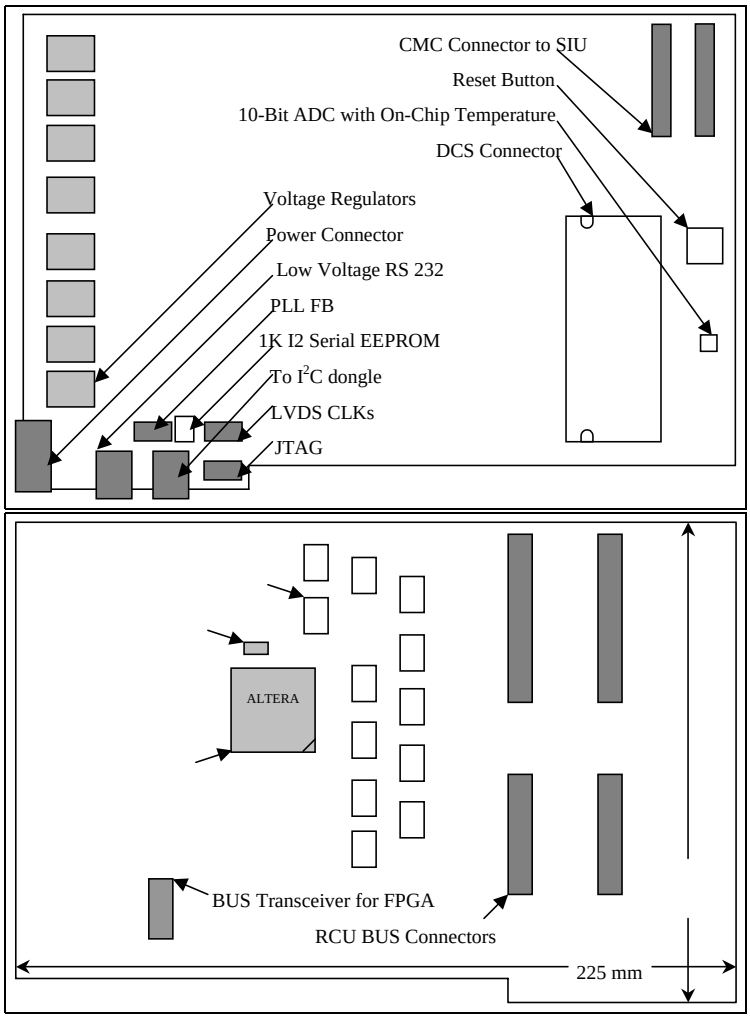

Figure B.1: Component placement on the RCU3 PCB top side (top) and bottom side (bottom).

Appendix C

Conference contributions

Readout Control Unit of the Front End Electronics for the ALICE Time Projection Chamber

R. Esteve Bosch[1], R. Campagnolo[1], H. Helstrup[3], A. Jiménez de Parga [1], J. A. Lien[1,2,3], B. Mota[1], L. Musa[1], D. Röhrich[2], B. Skaali[4], K. Ullaland[2], A. S. Vestbø[2], D. Wormald[4] for the ALICE Collaboration

[1] CERN, Geneva, Switzerland, [2] University of Bergen, Bergen, Norway, [3] Høgskolen i Bergen, Bergen, Norway, [4] University of Oslo, Oslo, Norway

Abstract

The readout electronics for the ALICE TPC detector consists of 4356 front-end cards (FECs) that contain the complete chain to readout the signals coming from 570132 pads. The front-end cards are grouped in 216 readout partitions, each controlled by a Readout Control Unit (RCU) that interfaces the FECs to the DAQ, the Trigger, and the Detector Control System. The RCU broadcast the trigger information to the FECs, collects the trigger related data from the FECs, assembles a sub-event, compresses the data and sends the compressed packed sub-event to the DAQ via the ALICE Detector Data Link (DDL). The RCU is designed to cope with a maximum data throughput of 200 Mbyte/s. In this paper we present an overview of the requirements, the architecture and some of the main design features of the RCU, as well as a description of two RCU prototypes.

INTRODUCTION

The ALICE Time Projection Chamber [1] (TPC) is a large gas cylinder (88 m^3) divided in two drift regions by a central electrode located at its axial centre. A field cage creates a uniform electric field along each half of the chamber. Charged particles traversing the TPC volume ionise the gas along their path, liberating electrons that drift towards the detector end plates where multi-wire proportional chambers, with cathode pad readout, provide the necessary signal amplification. Each of the two readout planes is azimuthally segmented in 18 trapezoidal sectors.

The front-end electronics [2] for the ALICE TPC consists of 570132 channels. A single readout channel is comprised of three basic functional units: 1) a charge sensitive shaping amplifier, which transforms the charge induced in the pad into a differential semi-gaussian signal; 2) a 10-bit 10-MSPS A/D converter; 3) a digital circuit that contains a shortening filter for the cancellation of the signal tail, the baseline subtraction and zero suppression circuits, and a multi-event memory. The analogue functions are implemented by a custom integrated circuit (PASA) that incorporates in a single chip 16 channels. The ADC and the digital circuits are contained in a single chip named ALTRO (**AL**ICE **T**PC **R**ead **O**ut) [3]. The maximum number of samples that can be continuously processed for each trigger (event data stream) is 1000. Upon arrival of a first level trigger, the data stream is stored in the multi-event memory. This memory has the capacity to store either 4 events, if the event data stream is longer than 512 words, or 8 events otherwise. When the second level trigger (accept or reject) is received, the latest event data stream is either frozen in the data memory until its complete readout takes place, or it is discarded. The complete chain is contained in the Front-End Card (FEC) [4], each containing 128 channels, which are located at some 10cm away from the pad plane and are connected to it by capton cables. In total the TPC is equipped with 4356 FECs grouped in 216 readout partitions, with 6 partitions per TPC sector, each controlled by a Readout Control Unit (RCU).

As shown in figure 1, the RCU interfaces the FECs to the Data Acquisition System (DAQ), the Timing and Trigger System (TTC) and the Detector Control System (DCS). The RCU broadcasts the trigger information to the FECs, collects the trigger-related data from the FECs, assembles a sub-event, compresses the data and sends the compressed packed sub-event to the DAQ via optical fibre, the ALICE Detector Data Link (DDL) [5]. Moreover, the RCU has to initialise the FECs and monitor their behaviour reporting to the DCS any detected fault.

The radiation load on the TPC is low, with a total dose received over 10 years of less than 1Krad and a neutron fluency of less than 10^{11} n/cm^2. Thus standard radiation-soft technologies are suitable for the implementation of this electronics. However, some special care should be taken to protect the system against potential damages caused by Single Event Effects (SEEs).

II. RCU ARCHITECTURE

The RCU has the following functional requirements:

I. Distribution of the trigger and clock signals to the FECs.

II. Initialisation of the FECs, including pedestal values, digital filter coefficients, zero suppression threshold, etc.

III. Readout of trigger related data from the FECs, assembly of a sub-event, compression of the data and transfer to the DAQ via the ALICE Detector Data Link (DDL).

IV. Supervision and monitoring the data flow and the status of the FECs: power state, temperature, exception and error conditions, etc.

In this section we discuss each of the above functions and the corresponding interfaces.

A custom bus, the ALTRO bus [6], implements the main communication between the RCU and the FECs. The ALTRO bus is essentially an extension of the FEC's internal bus that allows the RCU to access the FEC's internal components. It is a multi-drop single-master bus where the RCU is the master unit and the FECs are the slaves. In order to minimize the length of the bus cables the RCU supports two branches running in opposite directions. From the electrical point of view the ALTRO bus is based on the GTL/GTL+ technology. All the RCU functions (I-IV) can be performed by means of the ALTRO bus. The ALTRO bus features a bandwidth of up to 200 Mbyte/sec. As shown in figure 1, an independent Control Network [7] establishes a dedicated bus connection between the RCU and the FEC, in view of monitoring and controlling (function IV) the FECs without interfering with the readout procedure. This Control Network is based on the I2C protocol, enhanced with an extra line for handling interrupts from the FECs, and supports a bandwidth of 3.4Mbit/sec.

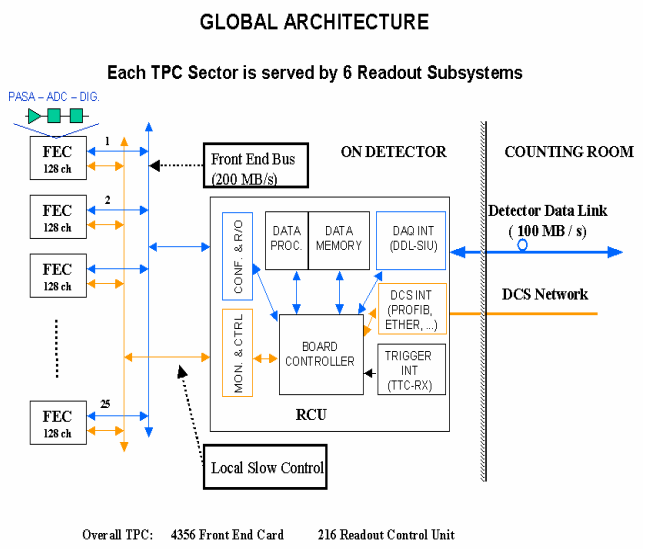

Figure 1: Block diagram of the TPC readout partition. The RCU interfaces a group of up to 25 FECs to the DAQ, the DCS and the TTC. The overall TPC readout consists of 216 readout partitions.

The interfacing of the RCU to the Trigger and the DAQ follows the standard data-acquisition architecture of the ALICE experiment [8]. The ALICE standard to transmit data off detector is the DDL technology (optical fibre). It implements a full duplex interface between the RCU and the Read Out Receiver Card (RORC), which is the first stage of the DAQ and is located in the experiment counting room. The DDL provides the RCU with a bandwidth of 100Mbyte/sec in the RCU→RORC direction and 10Mbyte/sec in the RORC→RCU direction. The DDL is composed of three hardware items: the Source Interface Unit (SIU), which is incorporated in the RCU, the fibre, and the Destination Interface Unit (DIU), which is part of the RORC. The SIU is a CMC-format mezzanine card plugged on the RCU as a daughter board.

The RD12 TTC [9] system is the ALICE standard used to distribute trigger and clock information. Every RCU implements an appropriate receiver (TTCrx), which produces the phase-corrected LHC clock, Level-1 trigger and Level-2 accept/reject. These are used on the RCU, and also distributed to all the FECs via the ALTRO bus. In addition, the TTCrx delivers event identification information to be added to the sub-event header before sending the data to the DDL.

For the communication between the RCU and the DCS two technologies are under investigation: Profibus and Ethernet.

III. RCU CONTROL LOGIC

The control logic of the RCU is divided into multiple state machines (see figure 2). The Resource and Priority Manager is the master state machine, controlling all subsystems.

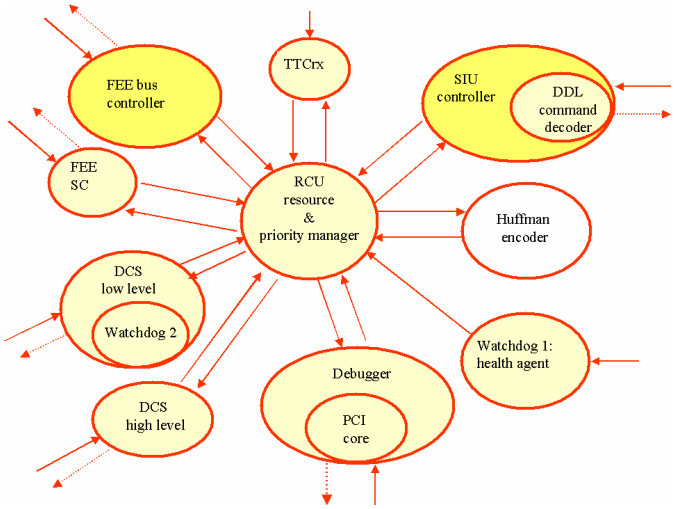

Figure 2: Flowchart of the RCU control logic. The Resource and Priority Manager is the master state machine, controlling all subsystems.

A. *ALTRO Controller*

Uses a custom protocol to service two ALTRO bus branches and all ALTRO chips. It builds and verifies ALTRO commands. The readout sequence of a TPC event is built up from information, stored in look-up tables, about the enabled and disabled channels. The amount of data received is counted and cross-checked with the block length transferred by each individual ALTRO chip.

B. SIU Controller.

It transfers blocks of data form the RCU memory into the DDL SIU by means of the DDL protocol. Each block of data contains one TPC sub-event. It also receives data from the DDL to be stored in the RCU memories or control registers.

C. *Trigger Manager (TTCrx).*

Manages the reception and distribution of triggers from the TTCrx chip. During the dead time of the TPC read out, incoming triggers are blocked. The TTCrx also supplies information that will be included in the Data Header of each sub-event. The Data Format of this header is an ALICE standard described in [10]. The Trigger Manager includes also

the logic to compare the number of triggers received by the RCU and by the ALTRO chips.

D. RCU Resource & Priority Manager.

The *Resource & Priority Manager* defines the priority of the incoming commands to the RCU. The standard operating mode of the RCU is the readout of TPC data from the ALTRO chips. The readout of one event is initiated by the reception of a Level-2 trigger from the TTCrx. The readout is controlled by the RCU. However, the data transfer from the ALTRO to RCU cannot be split in pieces. Thus the memory manager of the RCU must ensure that sufficient memory is available before starting a new data transfer. Whenever needed, the operator can ask to reconfigure either the RCU or the ALTRO chips. A request is sent to the *Resource & Priority Manager* that executes the configuration only after finishing the ongoing readout. Incoming requests are executed, under the control of this state machine, during dead time in readout.

E. *FEE Slow Control.*

Operates a dedicated bus connection between the RCU and the Board Controller in each Front End Card (FEC). This is used for voltage and temperature probing, protocol and multi-event buffer errors reporting, and statistical data read out.

F. Detector Slow Control.

It accesses in write/read mode a set of RCU registers dedicated to store the values of the control/status parameters of the overall system (power state, temperature, error messages, etc.). It connects the FEC's Control to the DCS system.

G. Huffman decoder (optional).

Compression of data can be done on the RCU to make effective use of the available bandwidth and/or memory. An algorithm for Huffman coding has been developed and tested [11].

H. *Watchdogs.*

Health agents, internal and external, that verify and monitor the functionality of the RCU FPGA. SRAM based FPGAs are vulnerable to radiation SEU. If the watchdog reports malfunction, reconfiguration of the FPGA can be done by either onboard or external configuration options.

I. *Debugging interface.*

PCI core for development and debugging. Optional in the final version.

The development of the RCU Control Logics is done using Mentor Tools (FPGA Advantage, ModelSim, Leonardo) and ALTERA Quartus II. Behavioural models (VHDL or Verilog) are used as test benches for the development of the interface state machines.

IV. RCU PROTOTYPE

In this section we describe two RCU prototypes that have been developed and tested with the rest of the TPC electronics components.

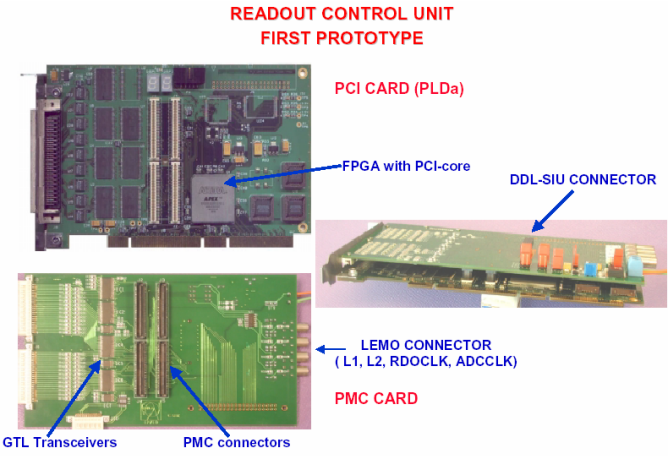

Figure 3: RCU Prototype I. This prototype is based on a comercially available PCI cards (PLDApplications) and a custom-made mezzanine card, which implements the interface to the FECs and the DDL SIU.

A first prototype of the RCU (pRCU-I) has been developed using the ALTERA EP20K400 FPGA on a commercial PCI-board (PLDA) [12]. The ALTRO protocol, a memory controller (FIFO structure) for accessing internal and external banks, and the SIU interface have been implemented in the FPGA and successfully tested. The set-up for the test of the complete readout chain is shown in figure 4.

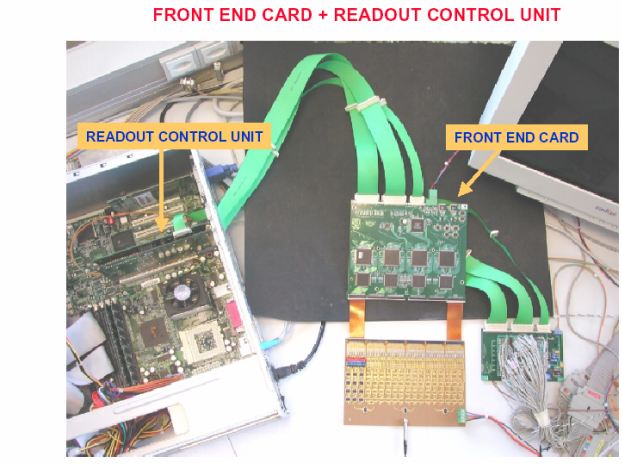

Figure 4: Set-up for the test of the ALTRO electronics. The RCU-I prototype is plugged in a PCI-slot of the PC on the left. One FEC is connected to the RCU via the ALTRO bus prototype.

A second prototype of the RCU has been developed and is being tested. The prototype is based on a PCI motherboard carrying two mezzanine cards, an SIU board and a board interfacing the two FEC-busses, the TTCrx and the DCS system. In this board all components required in the final design are connected to the FPGA. The FPGA I/O-pins are connected to 7 PMC connectors on which are plugged the two mezzanine cards.

The external memory configuration is optimised to meet the RCU requirements. Indeed, the memory is organised in 2 banks with separate data and address lines. This architecture allows the sub-event building of one event while transferring to the DAQ the previous event. The processing of the event includes adding a data header (including event length, identification and status) using the ALICE standard data format. Meanwhile, the sub-event stored in the other memory bank can be Huffmann-encoded (compressed) and pushed into the DDL. Another use of this prototype is to test different hardware options for the Detector Control System (DCS).

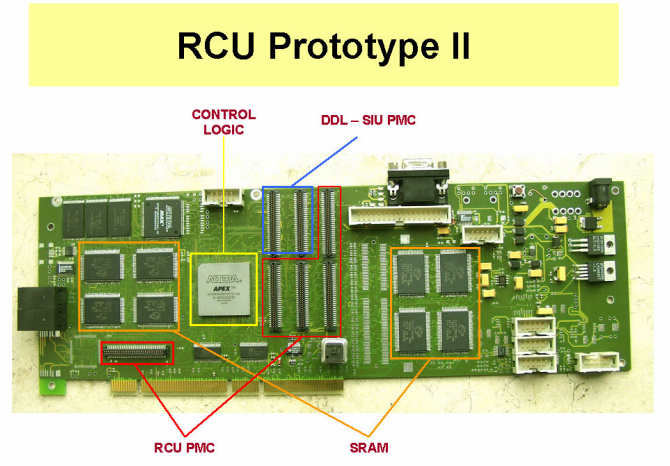

Figure 5: RCU Prototype II layout. SRAM is organised in 2 banks with separate data and address lines. The control logic is contained in the onboard FPGA. The 7 PMC connectors are used to connect to the FPGA two mezzanine boards: the DDL SIU card and the FEC interface mezzanine card.

V. CONCLUSIONS

The RCU plays an important role in the readout and control of the ALICE TPC front-end electronics. It interfaces the front-end cards, which contain the complete readout chain to amplify, digitise process and buffer the detector signals, to the DAQ, the Trigger and DCS. Besides the initialisation and readout functionality the RCU provides features to diagnostic errors and hardware faults.

A first prototype of the RCU has been developed by using a commercial PCI-board (PLDA), based on an EP20K400 FPGA, and a custom PMC board. This first prototype (pRCU-I), which incorporates only the interface to the FECs and to the DAQ, has been successfully used to validate the FEC and ALTRO bus design. A second prototype (pRCU-II), based on a PCI motherboard carrying two mezzanine cards, incorporates all the RCU functions and will be tested with a significant amount of the TPC electronics in the near future.

VI. REFERENCES

[1] *A Large Ion Collider Experiment, ALICE TPC - Technical Design Report*, December 1999, ISBN 92-9083-155-3, Geneva, Switzerland.

[2] L. Musa et al., *Front_End Electronics for the ALICE-TPC Detector*, Proc. of the 4th Workshop on Electronics for LHC Experiments, Rome, Sept. 21-25, 1998.

[3] L. Musa et al., *ALTRO User Manual*, http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/altro_chip.htm

[4] R. Campagnolo et al., *Front End Card*, http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/fec.htm

[5] G. Rubin, P. Vande Vyvre, *ALICE Detector Data Link (DDL) – Interface Control Document*, ALICE Internal Note, INT-96-43, December 1996.

[6] L. Musa et al., *The ALICE TPC Readout Bus*, http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/altro_bus.htm

[7] B. Mota et al., *The ALICE TPC Control Network*, http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/control_network.htm

[8] CERN/LHCC 95-71 LHCC/P3, *A Large Ion Collider Experiment, ALICE-Technical proposal*, Dec. 1995, ISBN 92-9083-077-8, Geneva, Switzerland.

[9] TTC homepage: http://ttc.web.cern.ch/TTC/intro.html

[10] P. Vande Vyvre et al., *ALICE DAQ Data Format*, http://edms.cern.ch/document/340186/

[11] T. Jahnke, S. Schoessel and K. Sulimma, EDA group, Department of Computer Science, University of Frankfurt (2000).

[12] PLDApplications, http://www.plda.com/hardw-pci20k.htm

Eur Phys J C (2004)
Digital Object Identifier (DOI) 10.1140/epjcd/s2004-03-1799-6

EPJ C direct
electronic only

Time projection chambers (TPC) in heavy ion experiments

J.A. Lien[1,2], D. Roehrich[2], H. Helstrup[1], and K. Ullaland[2]; presented by J.A. Lien

[1] Bergen University College, Postboks 7030, 5020 Bergen
[2] University of Bergen, Allegaten 55, 5007 Bergen

Received: 15 Oct 2003 / Accepted: 10 Mar 2004 /
Published Online: 30 Mar 2004 – © Springer-Verlag / Società Italiana di Fisica 2004

Abstract. In this paper the performance of TPCs in heavy ion experiments in terms of tracking and particle identification via dE/dx is reviewed. The principle of operation - drift, readout chambers and front end electronics, is illustrated. A few of the TPCs used are presented, along with a short description of the experiments.

1 TPC working principle

A TPC is a three-dimensional almost continuous tracking detector. It also provides information on the specific energy loss, dE/dx, of the particles traversing the drift volume. A TPC can be divided into three main parts:

- The drift volume
- The readout chambers
- The readout electronics

The charged particles produced in the collision pass through a gas mixture in the drift volume and release electrons from the gas atoms. These electrons drift in an external electric field to the readout chambers. The drifting electrons initiate avalanches in the high fields at the anode wires providing an amplification. The electron cloud disappear immediately. The positive ions created in the avalanche induce a temporary image charge on the readout pads which disappears as the ions move away from the anode wire. The image charge is measured by the readout electronics.

2 Drift volume

2.1 Gas mixture

The gas mixture used differs between experiments, see Table 1. The choice of gas [1] is driven by the drift velocity which should be fast and peak at a low electric field. Another property of the gas mixture is its influence on the diffusion of the charge cloud. A dedicated gas system circulates the gas and maintains purity. The main impurities which accumulate in the system are oxygen and water which capture drifting electrons.

2.2 Electric drift field

In the presence of an electric field, the electrons and ions freed by ionization are accelerated along the field lines towards the anode and cathode respectively. This acceleration is interrupted by collisions with the gas molecules which limit the maximum average velocity which can be attained by the charge along the field direction. The average velocity attained is known as the *drift velocity* of the charge and is superimposed upon its random movement.

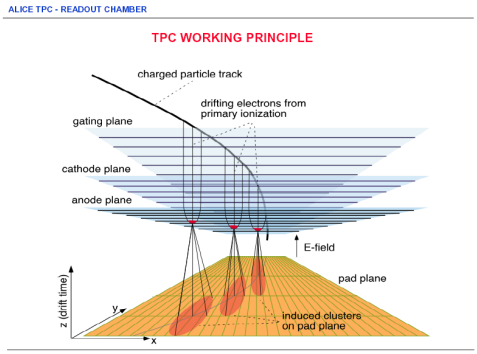

Fig. 1. The working principle of the ALICE TPC. The gating grid, cathode and anode wires are all in the readout chambers (MWPC chambers with pad readout)

Table 1. Gas mixtures of the presented TPCs

Experiment	TPC	Gas mixture
NA49	VTPC	90% Ne and 10% CO_2
	MTPC	90% Ar, 5% CH_4 and 5% CO_2
CERES/NA45	radial	80% Ne and 20% CO_2
STAR central	cylinder	90% Ar and 10% Methane
STAR forward	radial	50% Argon and 50% CO_2
ALICE	cylinder	90% Ne and 10% CO_2

2.3 Magnetic field

The curvature of the trajectory of charged particles in the known magnetic field can be used to measure the mo-

mentum of the particle. The presence of a magnetic field also influences the drift and diffusion of the ionized electrons [2].

3 Readout chambers

The readout chambers used in heavy ion experiments are based on MWPCs. The chambers consist of a pad plane and three or four wire planes. Drifting electrons originating from the primary ionization by themselves do not induce a sufficiently large signal in the readout pads. The necessary signal amplification is provided by avalanche creation in the vicinity of the anode wires. The induced charge from an avalanche is shared over several adjacent pads (2-3). The outermost wire plane is the gating grid. This grid is a shutter to control entry of electrons from the TPC drift volume into the readout chambers. It also blocks positive ions produced in the readout chamber, keeping them from entering the drift volume where they would distort the drift field.

3.1 Optimalization of the pad and wire layout

The diameter of the wires, their relative position, and the size, shape and positioning of the pads are, along with the choice of drift gas, the main design parameters that are subject to optimization. In a radial TPC (like the forward TPCs of the STAR experiment), the wires are tilted with a small angle so that two or more anode wires cross each pad for the selected pad-wire geometry. This is to avoid the observed periodic shifts in the position measurement that arises when the wire planes are orthogonal to the axial direction of the pads [7]. In linear drift TPCs (i.e. central STAR and ALICE) the wire planes are mounted orthogonal to the pads. The pad size is sometimes different in the outer and inner sectors. The smaller wire spacing and pad size of the inner sectors provides better spatial resolution necessary for tracking resolution near the center of the detector where track density is highest.

4 Readout electronics

The readout electronics generally consists of a chain of amplification, digitalization and transfer to DAQ/permanent storage. In ALICE (see Sect. 9.1) and CERES/NA45 the chain also includes fast digitized pre-processing in the front end electronics, done in a dedicated ASIC. For the other TPCs the amplified signals are stored in analog form in Switched Capacitor Arrays (SCA), which can be regarded as a series of capacitors (512 for NA49 and STAR central) with a sample and hold circuit for each capacitor [3]. After the readout cycle, the stored charges are digitized in ADCs and sent to DAQ. The number of channels to be read out is summarized on Table 2.

Table 2. Number of channels of the presented TPCs

Experiment	Pads	Pad size (mm^2)
NA49	182000	16x3.5, 28x3.5, 40x5.5
CERES/NA45	15360	chevron (w=10.3 mm, l=6 mm)
STAR central	136608	2.85x11.5, 6.20x19.5
STAR forward	9600	1.6x20
ALICE	570000	4x7.5, 6x10, 6x15

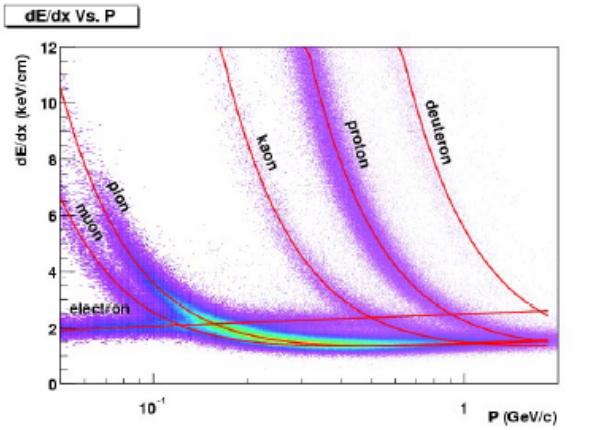

Fig. 2. The energy loss distribution for primary and secondary particles in the STAR TPC as a function of the momentum p of the primary particle. The magnetic field was 0.25 T

5 The NA49 TPCs

The NA49 detector [4] is a wide acceptance spectrometer for the study of hadron production in p+p, p+A, and A+A collisions at the CERN SPS (fixed target). The main components are 4 large volume TPCs for tracking and particle identification via dE/dx. The TPC system deploys two "Vertex" chambers (VTPC) inside the magnets and two "Main" chambers (MTPC) on both sides of the beam behind the magnets.The VTPCs measures 2.0 x 2.5 x .98 m^3. These are used to derive the particle momentum. The MTPCs have a length of 3.9 m and a height of 1.8 m. They are placed behind the magnets to measure the ionization energy loss in the relativistic rise in order to meet the requirements of the dE/dx resolution.

6 The CERES/NA45 TPC

CERES (Cherenkov-Ring Electron Spectrometer) [5] is a dilepton experiment at the CERN SPS (fixed target). The CERES TPC is a cylindrical drift chamber with the drift field in radial direction and segmented pad readout. The sensitive volume is about 9 m^3 and the length 2 m. It is used to determine the momentum of the particles produced in the collision and was added to increase the mass resolution of the experiment.

7 The STAR central TPC

The STAR detector at RHIC uses the TPC [6] inside a solenoidal magnetic field as its primary tracking device. The TPC records the tracks of particles, measures their momenta, and identifies the particles by measuring their ionization energy loss (dE/dx) with a resolution of 8%. This is the first big cylindrical TPC in a heavy ion collider experiment. The central TPC is 4.2 m long and 4 m in diameter. The central cathode, the inner and outer field cage cylinders and the readout end caps, defines the electronic drift field. More than 3000 tracks per event are routinely reconstructed.

8 The STAR forward TPC

The main argument to chose a TPC with a radial drift field is the improved two-track resolution in the transverse direction. The transverse separation of two charge clouds increases linearly with increasing r, but the width of the charge distribution of the cloud only increases by $\sqrt{r}$. The Forward Time Projection Chambers (FTPC) [7] were constructed to extend the acceptance of the STAR experiment. The FTPC has a cylindrical structure, 75 cm in diameter and 120 cm long. The field cage is formed by the inner HV-electrode, a thin metalized plastic tube, and the outer cylinder wall at ground potential. Drift length at maximum is only 23 cm. A two-track resolution of 1-2 mm and a position resolution of 100 μm is expected.

9 The ALICE TPC

The ALICE TPC [8] is the main tracking detector of the central barrel of the ALICE experiment at LHC. It provides charged particle momentum measurement with sufficient momentum resolution, particle identification by dE/dx and by decay topology analysis and vertex determination in the region $p_t < 10$ GeV/c and pseudorapidities $|\eta| < 0.9$. In addition, the TPC will provide - standalone or in combination with the other central barrel detectors (e.g. TRD and ITS) - the input for the High Level Trigger [9] in order to select low cross section signals and rare processes.

The TPC is a 88 m^3 cylinder (5.1 m long and 5.56 m in diameter) divided in two drift regions by the central electrode located at its axial center. It is being constructed based upon the aggregated experience from previous TPCs, and is very similar to the STAR central TPC.

9.1 Front end electronics

One of the challenges in designing the TPC for ALICE is the design of front end electronics to cope with the data rates and the channel occupancy. For this purpose a custom chip has been designed. The ALTRO (ALice Tpc ReadOut) chip [10] contains 16 times the TSA1001 ADC from ST Microelectronics in the form of two 8-ADC IP blocks. This ADC has differential inputs and can work up to 20 MS/sec with a resolution of 10 bits, while consuming as little as 12 mW. After digitization, a Baseline Correction Unit is able to remove systematic perturbations of the baseline by subtracting a pattern stored in a memory. The tail cancellation filter is a 3-stage IIR digital filter that removes the long complex tail of the detector signal. This narrows the clusters to improve their identification. The filter can also remove undershoot that typically distorts the amplitude of the clusters when pile-up occurs. A second correction of the baseline is performed based on the moving average of the samples that fall within a certain acceptance window. Eventually, the Zero Suppression procedure removes all data that is below a certain threshold, except for a specified number of pre- and post-samples around each peak. An example of data processing with the ALTRO chip is shown in Fig. 3.

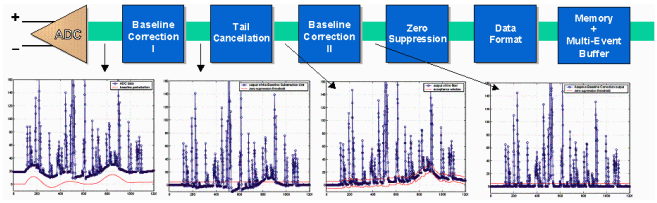

Fig. 3. The data processing stages of the ALTRO chip

References

1. R. Veenhof: "Choosing a gas mixture for the ALICE TPC", ALICE-INT-2003-29, 2003
2. W. Blum and L. Rolandi: *Particle Detection with Drift Chambers* (Springer-Verlag, 1994)
3. S.A. Kleinfelder: IEEE Trans. Nucl. Sci. **37**, 1230 (1990)
4. S.V. Afanasiev et al.: "The NA49 Large Acceptance Hadron Detector", Nucl. Instr. Meth. A **430**, 210-244 (1999)
5. A. Marin for the CERES Collaboration: "First results from the CERES radial TPC", Nucl. Phys. A **661**, 673c (1999)
6. M. Anderson et al.: "The STAR time projection chamber:a unique tool for studying high multiplicity events at RHIC", Nucl. Instr. Meth. A **499**, 659-678 (2003)
7. K.H. Ackermann et al.: "The Forward Time Projection Chamber (FTPC) in STAR", Nucl. Instr. Meth. A **499**, 713-719 (2003)
8. ALICE Collaboration: "Technical Design Report of the Time Projection Chamber", ISBN 92-9083-155-3, (2000)
9. T. Alt et al.: "ALICE High-Level Trigger Conceptional Design", http://www.kip.uni-heidelberg.de/ti/L3, (2002)
10. L. Musa for the ALICE Collaboration: "The Time Projection Chamber for the ALICE experiment", Nucl. Phys. A **715**, 843c-848c (2003)

References

[1] F. Karsch et. al., Phys. Lett. **B478**, 2000 447. Z. Fodor and S. Katz, J. High Energy Phys. JHEP03, 2002 14.

[2] K. J. Eskola, EPS-HEP99, arXiv:hep-ph/9911350

[3] U. Heinz and M. Jacob, arXiv:nucl-th/0002042, 2000

[4] S. Mioduszewski, plenary talk HEP 2003

[5] N. S. Amelin et. al., Eur.Phys.J. **C22**, 2001 149-163

[6] M. Harrison et al. "RHIC project overview", Nucl. Instr. Meth. **A499**, 2003 235-244

[7] P. Koch, B. Muller and J. Rafelski, Phys. Rept. **142**, 1986 167-262

[8] J. Rafelski and J. Letessier, Acta Phys.Polon. **B34**, 2003 5791-5824

[9] G. E. Bruno et al., J.Phys.G **30**, 2004 717-724

[10] M. van Leeuwen for the NA49 Collaboration, "Recent results on spectra and yields from NA49", Nucl. Phys. **A715**, 2003 161c

[11] R. J. Glauber, *Lectures on theoretical Physics* (Inter-Science, New York 1954)

[12] J. D. Bjorken, Phys. Rev. D **27**, 1983 140

[13] T. Alber et al., Phys. Rev. Lett **75**, 1995 3814

[14] K. Adcox et al., Phys. Rev. Lett. **87**, 2001 5

[15] R. Hanbury Brown and R.Q. Twiss "A test of a new type of stellar interferometer on Sirius", Nature **178**, 1956 1046

[16] G. Goldhaber et al., Phys. Rev. **120**, 1960 300

[17] M.I. Podgoretsky, Sov. J. Nucl. Phys. **37**, 1983 272

[18] U. Heinz and B. Jacak "Two-particle correlations in relativistic heavy-ion collisions", Annu. Rev. Nucl. Part. Sci. **49**, 1999 529

[19] STAR Collaboration, "Pion Interferometry of $\sqrt{s_{NN}} = 130$ GeV Au + Au Collisions at RHIC", Phys. Rev. Lett. **87-8**, 2001

[20] I.G. Bearden et al. [NA44 Collaboration] "Two-kaon correlations in central Pb + Pb collisions at 158-A-GeV/c.", Phys. Rev. Lett. **87**, 2001 112301

[21] STAR Collaboration, nucl-ex/0111013

[22] J. Sullivan et al., Nucl. Phys. **A566**, 1994 531c

[23] D.H. Rischke and M. Gyulassy, Nucl. Phys. **A608**, 1996 479

[24] M.A. Lisa, U. Heinz and U.A. Wiedemann, Phys. Lett. **B489**, 2000 287

[25] E. Schnedermann, J. Sollfrank and U. Heinz, Phys. Rev. C **48**, 2001 2462

[26] S. Voloshin and Y. Zhang, Z. Phys. **C70**, 1996 665

[27] R. Snellings et. al., arXiv:nucl-ex/0305001

[28] STAR Collaboration, Phys. Rev. Lett. **91**, 2003 072304

[29] D. Hofman, "Recent Results from RHIC", talk HEP 2003

[30] T. Matsui and H. Satz, Phys. Lett. **B416**, 1986

[31] NA50 "Homepage of NA50 experiment:", http://www.cern.ch/NA50/

[32] G. Borges, Proceedings of the QM04 conference, J. Phys. G, 2004 vol. 8

[33] R. L. Thews, arXiv:hep-ph/0111015

[34] CERES/NA45 "Homepage of CERES/NA45 experiment:", http://www.physi.uni-heidelberg.de/physi/ceres/

[35] B. Lenkeit for the CERES Collaboration, Nucl. Phys. **A661**, 1999 23c

[36] ALICE Collaboration, "ALICE Physics Performance Report, Volume I", CERN/LHCC 2003-049, 2003

[37] K. J. Eskola, "High Energy Nuclear Collisions", Proc. Int. Europhysics Conf. on High Energy Physics **ISBN-0-7503-0661-0**, 1999 231

[38] C. Y. Wong, *Introduction to High-Energy Heavy-Ion Collisions* (World Scientific Publishing, Singapore 1994) 251

[39] M. M. Aggarwal et. al., Phys. Rev. Lett. **85**, 2000 2895-2899

[40] ALICE Collaboration, "ALICE Technical proposal", 1995

[41] E. Cerron-Zeballos et al., Nucl. Instr. Meth. **A374**, 1996 132

[42] J.-P. Blaizot, "Theoretical Conference Summary", Nucl. Phys. **A698**, 2002 360c-371c

[43] ALICE Collaboration, "Technical Design Report of the Inner Tracking System", CERN/LHCC 99-12, ALICE TDR 4, 1999

[44] ALICE Collaboration, "Technical Design Report of the Time Projection Chamber", ISBN 92-9083-155-3, 2000

[45] ALICE Collaboration, "A Transition Radiation Detector", Technical Proposal - Addendum 2 **CERN/LHCC/99-13**, 1999

[46] ALICE Collaboration, "The forward muon spectrometer - Addendum to the ALICE Technical Proposal", CERN/LHCC/96-32, LHCC/P3-Addendum 1, 1996

[47] ALICE Collaboration, "Technical design report of the Zero Degree Calorimeter (ZDC)", CERN/LHCC 99-5, ALICE TDR 3, 1999

[48] ALICE Collaboration, "Technical design report of the Photon Multiplicity Detector (PMD)", CERN/LHCC 99-32, ALICE TDR 6, 1999

[49] ALICE Collaboration, "Technical design report of the Photon Spectrometer (PHOS)", CERN/LHCC 99-4, ALICE TDR 2, 1999

[50] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments* (Springer-Verlag, Berlin 1994) 378

[51] C. Loizides et al., IEEE Trans. Nucl. Sci. **51**, 2004 3

[52] G. Grastveit, "VHDL-implementation of the Cluster Finder algorithm for use in Alice", Cand. Scient. thesis, University of Bergen, 2003.

[53] A. S. Vestbø, "Pattern Recognition and Data Compression for the ALICE High Level Trigger", Dr. Scient. thesis, University of Bergen, May 2004

[54] ALICE Collaboration, "Technical design report of the ALICE TRIGGER, DAQ, HLT, and Control System", ALICE-DOC-2004-001, 2004

[55] F. Sauli "Principles of operation of multiwire proportional and drift chambers", Lectures given in the academic training programme of CERN 1975-76, 1977

[56] R. Gluckstern, Nucl. Instr. and Meth. **24**, 1963 381

[57] W. Blum and L. Rolandi, *Particle Detection with Drift Chambers* (Springer-Verlag, 1994)

[58] M. Gyulassy and X.-N. Wang, "Homepage of the HIJING event generator", http://www-nsdth.lbl.gov/ xnwang/hijing/index.html

[59] "GEANT - Detector Description and Simulation Tool", http://wwwasd.web.cern.ch/wwwasd/geant/

[60] S. V. Afanasiev et al., "The NA49 Large Acceptance Hadron Detector", Nucl. Instr. Meth. **A430**, 1999 210-244

[61] H. Tilsner, "Two-Particle Correlations at 40, 80, and 158 AGeV Pb-Au Collisions", Dr. thesis, Ruperto-Carola University of Heidelberg, June 2002

[62] S.A. Kleinfelder, IEEE Trans. Nucl. Sci. **37**, 1990 1230

[63] B. Yu et al., "Investigation of chevron cathode pads for positioning encoding in very high rate, gas proportional chambers", IEEE Trans. Nucl. Sc. **38**, 1991 2454

[64] STAR Collaboration, "Homepage of the STAR experiment:", http://www.star.bnl.gov

[65] M. Anderson et al., "The STAR time projection chamber:a unique tool for studying high multiplicity events at RHIC", Nucl. Instr. Meth. **A499**, 2003 659-678

[66] L. Kotchenda et al., "The STAR TPC Gas System", Nucl. Instr. Meth. **A499**, 2003 659-678

[67] F. Bergsma et al., "The STAR Magnet System", Nucl. Instr. Meth. **A499**, 2003 659-678

[68] K. H. Ackermann et al., "The Forward Time Projection Chamber (FTPC) in STAR", Nucl. Instr. Meth. **A499**, 2003 713-719

[69] F. Bieser et al., "The Forward Time Projection Chamber for the STAR Detector - Proposal", MPI PhE/98-3, 1998

[70] H. K. Soltveit et. al., "Final version of the Analog Front-end Electronics for the ALICE TPC-Detector and ALICE TRD-Detector", GSI Annual Report 2003, 244

[71] L. Musa et al., "ALICE TPC Readout Chip User Manual", http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/

[72] L. Musa et al., "Homepage of the ALTRO chip", http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/

[73] L. Musa for the ALICE Collaboration, "The Time Projection Chamber for the ALICE experiment", Nucl. Phys. **A715**, 2003 843c-848c

[74] R. Veenhof, "Choosing a gas mixture for the ALICE TPC", ALICE-INT-2003-29, 2003

[75] H. Stelzer et al., "The ALICE TPC Readout Chamber: From Prototypes to Series Production", ALICE-INT-2003-017, 2003

[76] T. Alt et al., "ALICE High-Level Trigger Conceptional Design", http://www.kip.uni-heidelberg.de/ti/L3, 2002

[77] R. Campagnolo et al., "Performance of the ALICE TPC Front End Card", Proceedings of the 9th workshop on Electronics for LHC Experiments, Amsterdam, 29 Sep - 3 Oct 2003

[78] H.-A. Gustafsson, presentation at LHC Comprehensive Review March 22, 2004

[79] B. G. Taylor, "Timing, Trigger and Control (TTC) Systems for LHC Detectors", http://www.cern.ch/TTC/intro.html

[80] P. Vande Vyvre, "ALICE DAQ Data Format", http://edms.cern.ch/document/340186/

[81] "High Performance Backplane Design with GTL+", SCEA011 Texas Instruments, June 1999

[82] "General information about PVSS and Experiments Control Systems", http://cern.ch/itcobe/Services/Pvss/ScadaLab/CSC

[83] "DIM Homepage", http://www.cern.ch/dim

[84] "ZTT Worms ALICE-FEEControl Homepage", http://www.ztt.fh-worms.de/en/projects/Alice-FEEControl/index.shtml

[85] "Control Network documentation page", http://ep-ed-alice-tpc.web.cern.ch/ep-ed-alice-tpc/controlnetwork.htm

[86] D.A. Huffman, "A method for the construction of minimum redundancy codes", Proc. IRE **40**, 1951 1098-1101.

[87] "ALTERA Homepage", http://www.altera.com

[88] "EDIF Homepage", http://www.edif.org/introduction.html

[89] "Verilog Homepage", http://www.verilog.com/

[90] "VHDL Homepage", http://www.vhdl.org/

[91] "PLD Applications Homepage", http://www.plda.com/

[92] "TPC Test Data Homepage", http://mactpc01.cern.ch/ tpc/

[93] "Homepage of the Lehrstuhl fur Technische Informatik", http://www.kip.uni-heidelberg.de/ti/

[94] L. E. Danielsen and M. Solheim Myhre, "Datainnsamling - utlesning og visualisering", Hovedprosjekt ved Høgskolen i Bergen, June 2003

[95] R. Brun and F. Rademakers, "ROOT: An Object-Oriented Data Analysis Framework", http://root.cern.ch/

[96] "TRD DCS webpage", http://www.physi.uni-heidelberg.de/∼petracek/webpage/trd/dcs.html

[97] T. Alt, Presentation at the ALICE week in Sardinia, May 2004

[98] "CASTOR Homepage", http://castor.web.cern.ch/castor/

[99] "TPC Beam Test Homepage", http://test-alicetpc.web.cern.ch/test-alicetpc/

[100] "FLUKA Homepage", http://www.fluka.org/

[101] K. Røed, "Irradiation tests of ALTERA SRAM-based FPGAs", Cand. Scient. thesis, University of Bergen, May 2004

[102] A. Fasso, P. Foka, A. Morsch, A. Sandroval, G. Tsiledakis, "Radiation in the ALICE TPC detector", ALICE oct. 2003 Internal Note-TRD